



Intel® Xeon® Processor D-1500 Product Family

Datasheet - Volume 4 of 4: Intel® Xeon® Processor D-1500 Product
Family LAN Controller

November 2018



No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

The products and services described may contain defects or errors which may cause deviations from published specifications.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting www.intel.com/design/literature.htm.

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

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

© 2018 Intel Corporation.



Revision History

Revision Number	Description	Date
001	Initial Release (Intel Public)	March 2015
002	<p>Updated Table 1.1, Supported SKUs</p> <p>Updated Section 1.2, Supported Modes of Operation</p> <p>Updated Table 1.3, Supported System Configurations</p> <p>Updated Section 1.5, Related Documents</p> <p>Updated Table 1.2, Network Features</p> <p>Updated Table 1.4, Miscellaneous Features</p> <p>Updated Section 2.2, High-Speed Network Data Interface</p> <p>Updated Section 3.4.1.2.1, Flash Identification</p> <p>Updated Table 3.8, Clients and Access Types to the NVM</p> <p>Updated Section 3.4.5.1, Flash Deadlock Avoidance</p> <p>Updated Section 3.4.8.12 Software Flash Read Flow via the Flash-mode Interface</p> <p>Deleted Section 3.4.9.3 Hardware Auto Load from NVM</p> <p>Updated Figure 3.5, SFP+: SoC w/o SFI PHY</p> <p>Updated Figure 3.6. Dual 10/1 GbE BASE-T PHY</p> <p>Updated Section 3.7.1. I2C or MDIO Selection</p> <p>Updated Section 3.8 Network Interface</p> <p>Added Section 4.1.3 Firmware Reset</p> <p>Deleted Section 4.5.10 Security Initialization</p> <p>Updated Section 7.1.3.5.11 Flow Director Filters Initializing Flow</p> <p>Updated Section 8.2.2.3.20 PCIe LCB Data Port - PCI_LCBDATA (0x00011734)</p> <p>Updated Section 8.2.2.17.7 Energy Efficient Ethernet (EEE) Setup Register - EEE_SU (0x00004380)</p> <p>Updated Section 9.2.2.2 Device ID Register (0x2; RO)</p> <p>Updated Chapter 12.0 Glossary and Acronyms</p> <p>Deleted Section B.1.1 XAUI Operating Mode</p>	November 2015
003 (Revision 2.1)	Removed all references to the KX4 interface (not supported).	March 2016
004 (Revision 2.2)	<p>Updated Section 3.4.9.1 (Flow for Updating Secured Modules).</p> <p>Updated Section 6.0 (NVM Header).</p> <p>Updated Section 8.2 (Device Registers - PF).</p> <p>Added NVM recovery mode.</p>	November 2018



NOTE: *This page intentionally left blank.*

§



Contents

1.0	Introduction	17
1.1	Intel® Xeon® Processor D-1500 Product Family LAN Controller Overview	17
1.2	Supported Modes of Operation	18
1.3	Conventions	19
1.4	References	19
1.5	Related Documents	20
1.6	Features	21
2.0	Pin Interface	27
2.1	Signal Type Definition	27
2.2	High-Speed Network Data Interface	28
2.3	Low-Speed Network Data Interface	28
3.0	Interconnects	31
3.1	Host Primary Interface	31
3.1.1	Host Interface Architecture, Transaction and Link Layer Properties	31
3.1.2	PCIe Transaction Layer	32
3.1.3	Error Events and Error Reporting	37
3.2	Management Interfaces	44
3.2.1	SMBus	44
3.3	Sideband Interface (NC-SI)	44
3.3.1	Electrical Characteristics	45
3.3.2	NC-SI Transactions	45
3.4	Non-Volatile Memory (NVM)	45
3.4.1	General Overview	45
3.4.2	Shadow RAM	47
3.4.3	NVM Clients and Interfaces	48
3.4.4	Memory Mapped Host Interface	50
3.4.5	Flash Access Contention	51
3.4.6	Signature Field	52
3.4.7	VPD Support	52
3.4.8	NVM Read, Write, and Erase Sequences	54
3.4.9	Extended NVM Flows	60
3.4.10	NVM Authentication Procedure	61
3.5	Configurable I/O Pins — Software-Definable Pins (SDPs)	62
3.6	LEDs	64
3.7	Network Management Interface (MDIO or I ² C)	65
3.7.1	I ² C or MDIO Selection	65
3.7.2	Management Data I/O Interface (MDIO)	67
3.7.3	I ² C	72
3.8	Network Interface	74
3.8.1	Integrated PHY Support	74
3.8.2	Loopback Support	74
3.8.3	Ethernet Flow Control (FC)	74
3.8.4	Inter Packet Gap (IPG) Control and Pacing	82
3.9	Physical Functions (PF), Virtual Functions (VF) and ARI (Alternative RID Interpretation)	84
4.0	Initialization	87
4.1	Reset Operation	87
4.1.1	Reset Sources/Reset Order	88



- 4.1.2 Reset in a PCI-IOV Environment..... 91
- 4.1.3 Firmware Reset..... 92
- 4.1.4 Reset Effects 92
- 4.2 Queue Disable 95
- 4.3 Function Disable 95
 - 4.3.1 General..... 95
 - 4.3.2 Overview 95
 - 4.3.3 Control Options..... 96
 - 4.3.4 Event Flow for Enable/Disable Functions 97
- 4.4 Device Disable 98
 - 4.4.1 Overview 98
 - 4.4.2 BIOS Disable of the Device at Boot Time by Using the Strapping Option..... 98
- 4.5 Software Initialization and Diagnostics 98
 - 4.5.1 Introduction 98
 - 4.5.2 Power-Up State..... 99
 - 4.5.3 Initialization Sequence 99
 - 4.5.4 Link Initialization..... 100
 - 4.5.5 Initialization of Statistics 100
 - 4.5.6 Interrupt Initialization 100
 - 4.5.7 Receive Initialization 101
 - 4.5.8 Transmit Initialization 103
 - 4.5.9 Virtualization Initialization Flow 104
 - 4.5.10 Alternate MAC Address Support..... 107
- 4.6 Access to Shared Resources 108
- 5.0 Power Management and Delivery 109**
- 5.1 Power Management 109
 - 5.1.1 Introduction to the Intel® Xeon® Processor D-1500 Product Family LAN Controller Power States 109
 - 5.1.2 Auxiliary Power Usage..... 109
 - 5.1.3 Power States 110
- 5.2 Network Interfaces Power Management 114
 - 5.2.1 PHY Power-Down State 115
 - 5.2.2 Low Power Link Up (LPLU) 115
 - 5.2.3 Energy Efficient Ethernet (EEE) 115
- 5.3 Wake Up 118
 - 5.3.1 Advanced Power Management Wake Up 118
 - 5.3.2 ACPI Power Management Wake Up 119
 - 5.3.3 Wake-Up Packets 120
 - 5.3.4 Wake Up and Virtualization 124
- 5.4 Protocol Offload (Proxying) 124
 - 5.4.1 Proxying Filters 124
 - 5.4.2 Proxying and Virtualization 131
 - 5.4.3 Protocol Offload (Proxying) Flow..... 131
- 5.5 DMA Coalescing 133
 - 5.5.1 DMA Coalescing Activation..... 133
 - 5.5.2 DMA Coalescing Operating Mode 134
 - 5.5.3 DMA Coalescing Recommended Settings..... 134
- 5.6 Latency Tolerance Reported (LTR) 135
 - 5.6.1 LTR Algorithm..... 135
 - 5.6.2 LTR Initialization Flow 136



6.0	Non-Volatile Memory Map	137
6.1	NVM Organization	137
6.1.1	Protected Areas	138
6.2	NVM Header	139
6.3	Software Sections	140
6.3.1	Software Compatibility Module — Word Address 0x10-0x14	140
6.3.2	Word Address 0x15-0x16 - Reserved	141
6.3.3	Boot Configuration Block — Word Address 0x17	141
6.3.4	Software Reserved — Words 0x18-0x2E	143
6.3.5	VPD Module Pointer — Word Address 0x2F	144
6.3.6	PXE Configuration Words — Word Address 0x30-0x36	144
6.3.7	Alternate Ethernet MAC Address Pointer — Word Address 0x37	148
6.4	Hardware Sections	149
6.4.1	Hardware Section — Auto-Load Sequence	149
6.4.2	NVM Init Module	149
6.4.3	SoC Configuration Module	152
6.4.4	PCIe General Configuration Module	153
6.4.5	PCIe Configuration Space 0/1 Modules	154
6.4.6	LAN Core 0/1 Modules	155
6.5	PCIe Expansion/Option ROM	156
7.0	Inline Functions	157
7.1	Receive Functionality	157
7.1.1	MAC Layer - Receive	157
7.1.2	Packet Filtering	158
7.1.3	Rx Queues Assignment	162
7.1.4	Receive Data Storage in System Memory	185
7.1.5	Receive Descriptors	185
7.1.6	Receive Offloads	199
7.1.7	Receive Statistics	203
7.2	Transmit Functionality	206
7.2.1	Packet Transmission	206
7.2.2	Transmit Contexts	215
7.2.3	Transmit Descriptors	216
7.2.4	TCP and UDP Segmentation	230
7.2.5	Transmit Checksum Offloading in Non-segmentation Mode	236
7.2.6	Transmit Statistics	239
7.3	Interrupts	241
7.3.1	Interrupt Registers	241
7.3.2	Interrupt Moderation	244
7.3.3	TCP Timer Interrupt	246
7.3.4	Mapping of Interrupt Causes	246
7.4	802.1q VLAN Support	253
7.4.1	802.1q VLAN Packet Format	253
7.4.2	802.1q Tagged Frames	253
7.4.3	Transmitting and Receiving 802.1q Packets	254
7.4.4	802.1q VLAN Packet Filtering	254
7.4.5	Double VLAN and Single VLAN Support	255
7.4.6	E-tag and VLAN	257
7.5	TLP Processing Hints (TPH)	258
7.5.1	Steering Tag and Processing Hint Programming	259



- 7.6 Time SYNC (IEEE1588 and 802.1AS) 260
 - 7.6.1 Overview 260
 - 7.6.2 Flow and Hardware/Software Responsibilities 260
 - 7.6.3 Hardware Time Sync Elements 262
 - 7.6.4 Hardware Time Sync Elements 264
 - 7.6.5 Time SYNC Interrupts 266
 - 7.6.6 PTP Packet Structure 267
- 7.7 Virtualization 269
 - 7.7.1 Overview 269
 - 7.7.2 PCI-SIG SR-IOV Support 274
 - 7.7.3 Packet Switching 283
 - 7.7.4 Security Features 292
 - 7.7.5 Virtualization of Hardware 296
- 7.8 Tunnelling Support 297
- 7.9 Receive Side Coalescing (RSC) 297
 - 7.9.1 Packet Candidacy for RSC..... 299
 - 7.9.2 Flow Identification and RSC Context Matching 301
 - 7.9.3 Processing New RSC 302
 - 7.9.4 Processing Active RSC..... 302
 - 7.9.5 Packet DMA and Descriptor Write Back..... 304
 - 7.9.6 RSC Completion and Aging 306
- 7.10 Reliability 307
 - 7.10.1 Memory Integrity Protection 307
 - 7.10.2 PCIe Error Handling 307
- 8.0 Programming Interface 309**
- 8.1 General 309
 - 8.1.1 Memory-Mapped Access 309
 - 8.1.2 I/O-Mapped Access..... 310
 - 8.1.3 Configuration Access to Internal Registers and Memories 311
 - 8.1.4 Register Terminology 312
 - 8.1.5 VF Registers Allocated per Queue 313
 - 8.1.6 Non-queue VF Registers 313
 - 8.1.7 Access to MAC Registers..... 313
- 8.2 Device Registers - PF 315
 - 8.2.1 BAR0 Registers Summary 315
 - 8.2.2 Detailed Register Description - PF BAR0 328
 - 8.2.3 BAR3 Registers Summary 481
 - 8.2.4 Detailed Register Description - PF BAR3 481
- 8.3 Device Registers - VF 483
 - 8.3.1 VF Registers Mapping PF Space 483
 - 8.3.2 BAR0 Registers Summary 484
 - 8.3.3 Detailed Register Description - VF BAR0 486
- 9.0 PCIe Programming Interface 495**
- 9.1 Overview 495
 - 9.1.1 PCIe Configuration Space in an Integrated I/O Interface Connected System 495
 - 9.1.2 Register Attributes 496
- 9.2 PCIe Register Map 496
 - 9.2.1 PCIe Configuration Space Summary 496
 - 9.2.2 Mandatory PCI Configuration Registers..... 498
 - 9.2.3 PCI Capabilities 504



9.2.4	PCIe Extended Configuration Space	527
9.2.5	Driver Forward Compatibility Register (0x94; RO).....	549
9.2.6	CSR Access Via Configuration Address Space	550
9.3	Virtual Functions Configuration Space	551
9.3.1	Mandatory Configuration Space.....	553
9.3.2	PCI Capabilities.....	554
9.3.3	PCIe Extended Capabilities	555
10.0	System Manageability	557
10.1	Pass-Through (PT) Functionality	557
10.1.1	Supported Topologies	558
10.1.2	Pass Through Packet Routing	558
10.2	Components of the Sideband Interface	559
10.2.1	Physical Layer.....	559
10.2.2	Logical Layer	560
10.3	Packet Filtering	562
10.3.1	Manageability Receive Filtering.....	562
10.3.2	L2 Filters.....	563
10.3.3	L3/L4 Filtering	564
10.3.4	Flexible 128-byte Filter	565
10.3.5	Configuring Manageability Filters	566
10.3.6	Filtering Programming Interfaces.....	570
10.3.7	Possible Configurations	571
10.3.8	Determining Manageability MAC Address	572
10.4	OS-to-BMC Traffic	572
10.4.1	Overview	572
10.4.2	Filtering	573
10.4.3	Blocking of Network-to-BMC Flow	575
10.4.4	OS-to-BMC and Flow Control.....	575
10.4.5	Statistics.....	576
10.4.6	OS-to-BMC Enablement.....	576
10.5	SMBus Pass-Through Interface	576
10.5.1	General.....	577
10.5.2	Pass-Through Capabilities.....	577
10.5.3	Port to SMBus Mapping	577
10.5.4	Automatic Ethernet ARP Operation.....	577
10.5.5	SMBus Transactions.....	578
10.5.6	SMBus Notification Methods	583
10.5.7	Receive Pass Through Flow	585
10.5.8	Transmit Pass Through Flow	586
10.5.9	SMBus Link State Control	587
10.5.10	SMBus ARP Transactions	588
10.5.11	SMBus Pass-Through Transactions.....	591
10.5.12	Example Configuration Steps	614
10.5.13	SMBus Troubleshooting	622
10.6	NC-SI Pass Through Interface	625
10.6.1	Overview	625
10.6.2	NC-SI Standard Support.....	629
10.6.3	NC-SI Mode — Intel Specific Commands.....	632
10.6.4	Asynchronous Event Notifications	676
10.6.5	Querying Active Parameters.....	676



- 10.6.6 Resets 676
- 10.6.7 Advanced Workflows..... 677
- 10.6.8 External Link Control via NC-SI 679
- 10.7 Management Component Transport Protocol (MCTP) 681
 - 10.7.1 MCTP Overview 681
 - 10.7.2 NC-SI to MCTP Mapping 682
 - 10.7.3 MCTP Over SMBus 685
 - 10.7.4 NC-SI over MCTP 686
 - 10.7.5 MCTP Programming 688
- 10.8 Manageability Host Interface 692
 - 10.8.1 HOST CSR Interface (Function 1/0) 692
 - 10.8.2 Host Slave Command Interface to Manageability 692
 - 10.8.3 Host Interface Commands 694
 - 10.8.4 Software and Firmware Synchronization 702
- 10.9 Host Isolate Support 704
- 11.0 Diagnostics 705**
 - 11.1 JTAG Test Mode Description 705
 - 11.2 MAC Loopback Operations 706
 - 11.2.1 MAC Loopback 707
 - 11.2.2 PHY Loopbacks 708
- 12.0 Glossary and Acronyms 709**
- Appendix A Packet Formats..... 721**
 - A.1 Legacy Packet Formats 721
 - A.2 Packet Types for Packet Split Filtering 727
 - A.3 E-tag and S-tag Formats..... 743
- Appendix B Integrated PHY Support..... 745**
 - B.1 10 GbE Interface..... 745
 - B.2 10GBASE-KR Operating Mode 745
 - B.3 1 GbE Interface 747
 - B.4 Auto Negotiation For Backplane Ethernet and Link Setup Features 749
 - B.5 KR Link Configuration Flows..... 749
 - B.6 Low Power Link Up (LPLU)..... 750
 - B.7 **Behavior in Non-DO State** 750

Figures

- Figure 1.1 Intel® Xeon® Processor D-1500 Product Family LAN Controller Overview 17
- Figure 1.2 Intel® Xeon® Processor D-1500 Product Family LAN Controller Logical Block Diagram 18
- Figure 3.1 Error Reporting Mechanism..... 38
- Figure 3.2 NVM Shadow RAM 48
- Figure 3.3 Flash access modes 50
- Figure 3.4 Sign and Verify Procedures for Authenticated NVM Modules..... 62
- Figure 3.5 SFP+: SoC w/o SFI PHY 66
- Figure 3.6 Dual 10/1 GbE BASE-T PHY 67
- Figure 3.7 Basic PHY MAC Connectivity..... 68
- Figure 3.8 MDIO Timing Sourced by the MAC 69
- Figure 3.9 MDIO Timing Sourced by the PHY 69
- Figure 3.10 Behavior of MDIO During TA Field of a Read Transaction..... 71
- Figure 4.1 Rest Tree 87
- Figure 5.1 Power Management State Diagram 109
- Figure 5.2 EEE Operation 116



Figure 5.3	NS IPv6 Packet Flow	128
Figure 6.1	NVM Structure	138
Figure 7.1	Stages in Packet Filtering - Virtualization Mode.....	158
Figure 7.2	Rx Filtering Flow Chart.....	159
Figure 7.3	VLAN Filtering.....	161
Figure 7.4	Manageability / Host Filtering.....	162
Figure 7.5	Rx Queuing Flow (Non-virtualized).....	164
Figure 7.6	Rx Queuing Flow (Virtualization Case).....	166
Figure 7.7	Flow Director Filters Block Diagram.....	172
Figure 7.8	RSS Block Diagram (Multiple_RSS = 0).....	180
Figure 7.9	RSS Block Diagram (Multiple_RSS = 1).....	181
Figure 7.10	Receive Descriptor Ring Structure	196
Figure 7.11	Descriptors and Memory Rings	197
Figure 7.12	Header Splitting Diagram	199
Figure 7.13	Receive Flow Statistics.....	205
Figure 7.14	Tx Packet in Host Memory	206
Figure 7.15	Tx Arbitration Schemes	210
Figure 7.16	Transmit Architecture Intel VT-on — Eight TCs Mode.....	211
Figure 7.17	Transmit Architecture Intel VT-off — Eight TCs Mode.....	212
Figure 7.18	Transmit Architecture Intel VT-on — 32 VFs.....	214
Figure 7.19	Transmit Architecture Intel VT-off.....	215
Figure 7.20	Transmit Descriptor Ring Structure	226
Figure 7.21	Transmit Flow Statistics	240
Figure 7.22	Cause Mapping in Legacy and MSI Modes	247
Figure 7.23	Rx and Tx Queue Mapping to IVAR Registers.....	247
Figure 7.24	Cause Mapping in MSI-X Mode (non-IOV)	249
Figure 7.25	VF Interrupt Cause Mapping (MSI-X, IOV)	250
Figure 7.26	VF Mapping of Rx and Tx Queue to VFIVAR in 64 VFs Mode.....	250
Figure 7.27	VF Mapping of Rx and Tx Queue to VFIVAR in 32 VFs Mode.....	251
Figure 7.28	VF Mapping of Rx and Tx Queue to VFIVAR in 16 VFs Mode.....	251
Figure 7.29	PF / VF MSI-X Vectors Mapping to <i>EITR</i>	252
Figure 7.30	Sync Flow and Offset Calculation	261
Figure 7.31	Time Stamp Point.....	263
Figure 7.32	System Configuration for VMDq2 Mode.....	272
Figure 7.33	System Configuration for IOV Mode	273
Figure 7.34	BARs in an IOV-enabled System.....	275
Figure 7.35	RSC Functionality (No Header Split).....	298
Figure 7.36	RSC Functionality (No Header Split).....	298
Figure 7.37	RSC Event Flow.....	299
Figure 7.38	RSC — Header and Data Buffers.....	304
Figure 10.1	Sideband Interface	557
Figure 10.2	Manageability Decision Filters	568
Figure 10.3	OS-to-BMC Block Diagram	573
Figure 10.4	OS-to-BMC and VM-to-VM Filtering	574
Figure 10.5	SMBus ARP Flow	581
Figure 10.6	Single NC Package, Two NC Channels	627
Figure 10.7	Two NC Packages (Left, with Two NC Channels and Right, with One NC Channel).....	627
Figure 10.8	NC-SI Frames Receive Flow for the NC.....	629
Figure 10.9	Multi-NC Environment.....	677
Figure 10.10	MCTP Connections of the Intel® Xeon® Processor D-1500 Product Family LAN Controller	682



Figure 10.11 MCTP Endpoint Topology 683
Figure 10.12 MCTP Bus Transition State Machine 684
Figure 11-1 Loopback Modes 707
Figure 15.1 Architectural Positioning of 10GBASE-KR 746
Figure 15.2 KR Backplane Ethernet Channel 747
Figure 15.3 Architectural Positioning of 1000BASE-KX 748
Figure

Tables

Table 1.1 Supported SKUs 18
Table 1-3 Supported System Configurations 19
Table 1.2 Network Features..... 21
Table 1.3 Host Interface Features 22
Table 1.4 Miscellaneous Features 22
Table 1.5 LAN Functions Features 23
Table 1.6 LAN Performance Features 23
Table 1.7 Virtualization Features 24
Table 1.8 Manageability Features 24
Table 3.1 Transaction Types Accepted by the Transaction Layer 32
Table 3.2 Transaction Types Initiated by the Transaction Layer 33
Table 3.3 Supported Message in the Intel® Xeon® Processor D-1500 Product Family LAN Controller (as a Receiver) 35
Table 3.4 Response and Reporting of PCIe Error Events 39
Table 3.5 PCIe Errors Reported to Device Software..... 42
Table 3.6 Reset Warn Message Format 43
Table 3.7 Reset Ack Message Format..... 44
Table 3.8 Clients and Access Types to the NVM 49
Table 3.9 Small Resource Structure..... 53
Table 3.10 Large Resource Structure..... 53
Table 3.11 VPD Structure 53
Table 3.12 GPI to SDP Bit Mappings 63
Table 3.13 SDP Settings 64
Table 3.14 NW Management Selection Settings - Per Port 66
Table 3.15 MDC Frequency as Function of Link Speed and MDC Speed Bit 68
Table 3.16 Clause 22 Basic MDIO Frame Format 70
Table 3.17 Clause 45 Indirect Addressing MDIO Frame Format 70
Table 3.18 I²C Read Transaction - Dummy Write 73
Table 3.19 I²C Read Transaction - Word Read 73
Table 3.20 I²C Read Transaction - Dummy Write 73
Table 3.22 I²C Write Transaction - Word Write 74
Table 3.23 I²C Write Transaction - Byte Write 74
Table 3.21 I²C Read Transaction - Byte Read 74
Table 3.24 MAC Control Frame Format 75
Table 3.25 802.3X Packet Format 76
Table 3.26 Transfer of PAUSE Packet to Host (DPF Bit) 77
Table 3.27 Delay Values (DV) Used For FCRTM 80
Table 3.28 Intel® Xeon® Processor D-1500 Product Family LAN Controller FCRTL..... 81
Table 3.29 Some Recommended Rx Packet Buffer Settings 82
Table 3.30 Pacing Speeds at 10 GbE Link Speed 83



Table 4.1	Reset Effects — Common Resets	92
Table 4.2	Reset Effects — Per Function Resets	93
Table 4.3	Reset Effects -Virtual Function Resets	93
Table 4.4	PCI Functions Mapping	96
Table 5.1	EEE Parameters	116
Table 6.1	NVM Header Description	139
Table 6.2	NVM Section Auto-Read	149
Table 7.1	Rx Queuing Schemes Supported (No Virtualization)	163
Table 7.2	Queue Indexing Illustration in Non-virtualization Mode.....	163
Table 7.3	Rx Queuing Schemes Supported with Virtualization	165
Table 7.4	Queue Indexing Illustration in Virtualization Mode	165
Table 7-5	Candidate for Flow Director Per Mode	168
Table 7-6	Lookup fields for Flow Director Per Mode	169
Table 7-7	RSS Type Values According to Flow Director Match	171
Table 7.8	Rx Packet Buffer Allocation.....	172
Table 7.9	Signature Match Filter Parameters.....	176
Table 7.10	Perfect Match Filter Parameters	176
Table 7.11	Input Bit Stream for Hash Calculation	179
Table 7.12	IPv4.....	184
Table 7.13	IPv6.....	185
Table 7.14	Legacy Receive Descriptor (RDESC) Layout	185
Table 7.15	Receive Status (RDESC.STATUS) Layout	186
Table 7.16	Receive Errors (RDESC.ERRORS) Layout	187
Table 7.17	VLAN Tag Field Layout (for 802.1q Packet).....	188
Table 7.18	Descriptor Read Format	188
Table 7.19	Descriptor Write-Back Format.....	189
Table 7.20	Checksum Enable/Disable	191
Table 7.21	Receive Status (RDESC.STATUS) Layout of Last Descriptor.....	192
Table 7.22	Receive Status (RDESC.STATUS) Layout of Non-Last Descriptor	192
Table 7.23	Receive Errors (RDESC.ERRORS) Layout	194
Table 7.24	Behavior in Header Split Modes.....	200
Table 7.25	Supported Receive Checksum Capabilities	202
Table 7.26	SCTP Header.....	203
Table 7.27	Tx Queuing Schemes.....	207
Table 7.28	Queues, Pools and TCs Programming.....	208
Table 7.29	Tx Queues Indexing When Intel VT-on.....	208
Table 7.30	Tx Queues Indexing When Intel VT-off.....	208
Table 7.31	VLAN Tag.....	216
Table 7.32	Transmit Descriptor (TDESC) Layout — Legacy Mode	217
Table 7.33	Transmit Descriptor Write-Back Format — Legacy Mode	217
Table 7.34	Transmit Command (TDESC.CMD) Layout	218
Table 7.35	VLAN Tag Insertion Decision Table for VLAN Mode Enabled	218
Table 7.36	VLAN Field (TDESC.VLAN) Layout.....	219
Table 7.37	Transmit Context Descriptor (TDESC) Layout — (Type = 0010).....	219
Table 7.38	Valid Fields by Offload Option	223
Table 7.39	Advanced Transmit Data Descriptor Read Format	223
Table 7.40	Advanced Transmit Data Descriptor Write-back Format	223
Table 7.41	TCP/IP and UDP/IP Packet Format Sent by Host.....	232
Table 7.42	Packets Format Sent by Device.....	232
Table 7.43	TCP Partial Pseudo-header Checksum for IPv4	233



Table 7.44	TCP Partial Pseudo-header Checksum for IPv6	233
Table 7.45	Checksums Supported by Packet Type	238
Table 7.46	PF Interrupt Registers	241
Table 7.47	PF / VF MSI-X Vectors Mapping Table to EITR Registers	252
Table 7.48	Double VLAN Packet Format	255
Table 7.49	UP Extraction Rules	258
Table 7-50	Steering tag and Processing hint programming	259
Table 7-51	Enabling Receive Timestamp	269
Table 7.52	BARs in the Intel® Xeon® Processor D-1500 Product Family LAN Controller (64-bit BARs)	274
Table 7.53	RID Per VF — ARI Mode	278
Table 7.54	RID Per VF — Non-ARI Mode	278
Table 7.55	Queue to VF Allocation.....	279
Table 7.56	Mailbox Memory	281
Table 7.57	PF-to-VF Messaging Flow	281
Table 7.58	VF-to-PF Messaging Flow	281
Table 7.59	Malicious Driver - Tx Descriptor Checks.....	295
Table 7.60	Reactive Malicious Checks	296
Table 7.61	Packet Format Supported by RSC.....	300
Table 7.62	IP TOS Field — Bit Map	300
Table 7.63	TCP Time-Stamp Option Header (RFC 1323)	300
Table 7.64	RSC Context	301
Table 8.1	Intel® Xeon® Processor D-1500 Product Family Address Regions	309
Table 8-2	IOADDR and IODATA in Configuration Address Space	311
Table 3	BAR0 Registers Summary	315
Table 8.3	BAR3 Registers Summary	481
Table 8.4	BAR0 Registers Summary	484
Table 9.1	PCI Configuration Registers Map - LAN Functions	497
Table 9.2	PCIe Configuration Registers Map - Dummy Function	498
Table 9.3	Configuration Sharing of PCI Configuration Space.....	498
Table 9.4	Header Type Settings	502
Table 9.5	Intel® Xeon® Processor D-1500 Product Family LAN Controller Base Address Registers Description — LAN 0 / LAN 1	503
Table 9.6	Base Address Registers' Fields	503
Table 9.7	Usable Flash Size and CSR Mapping Window Size	504
Table 9.8	PCI Capabilities List for LAN Functions	504
Table 9.10	Sharing the Power Management Capability Registers	505
Table 9.9	PCI Capabilities for Dummy Function	505
Table 9.11	Configuration sharing of the MSI Capability	508
Table 9.12	MSI-X Capability Structure	511
Table 9.13	Configuration sharing of the MSI-X Capability	511
Table 9.14	Configuration sharing of the VPD Capability	513
Table 9.15	Configuration Sharing of the PCIe Capability	515
Table 9.16	Extended Capabilities list	528
Table 9.17	Configuration Sharing of the AER Capability	529
Table 9.18	Configuration Sharing of the Serial Number Capability	534
Table 9.19	Configuration sharing of the ARI Capability	535
Table 9.20	Configuration Sharing of the SR-IOV Capability	537
Table 9.21	Configuration Sharing of the TPH Requester Capability	543
Table 9.22	Lane [(Maximum Link Width – 1):0] Equalization Control Register	549
Table 9.23	VF PCIe Configuration Space	551



Table 10.1	SMBus PEC Modes.....	559
Table 10.2	Filters Accessible to BMC.....	562
Table 10.3	Exclusive traffic behavior	569
Table 10.4	<i>MNGONLY</i> Register Description and Usage.....	569
Table 10.5	Filtering Programming Interfaces	570
Table 10.6	Typical SMBus Transaction	578
Table 10.7	Shorthand Field Names.....	578
Table 10.8	UDID.....	582
Table 10.9	Flag Options During Transmit Packet Transactions	587
Table 10.10	Management Control Command Parameters/Content	595
Table 10.11	Management Receive Filter Parameters.....	596
Table 10.12	Filter Enable Parameters	597
Table 10.13	Set Common Filters Receive Control Bytes	599
Table 10.14	TX Control	602
Table 10.15	SMBus Read Transactions	602
Table 10.16	TCO LAN Packet Status Data	604
Table 10.17	Packet Status Info.....	604
Table 10.18	MNG Status.....	604
Table 10.19	Status Data Byte 1.....	605
Table 10.20	Status Data Byte 2.....	606
Table 10.21	Status Data Byte 2 (Bits 2 and 1).....	606
Table 10.22	Get Controller Information Data	610
Table 10.23	MACsec Owner Status.....	613
Table 10.24	MACsec Host Control Status	613
Table 10.25	TX Control Status:	613
Table 10.26	Example 1 MDEF Results.....	615
Table 10.27	Example 2 MDEF Results.....	617
Table 10.28	Example 3 MDEF Results.....	619
Table 10.29	Example 4 MDEF Results.....	621
Table 10.30	NC-SI Terminology.....	625
Table 10.31	Supported NC-SI Commands	629
Table 10.32	Optional NC-SI Features Support	630
Table 10.33	OEM Specific Command Response Reason Codes.....	633
Table 10.34	OEM Commands Summary	633
Table 10.35	Filter Data Group	639
Table 10.36	Filter Values.....	643
Table 10.37	Extended Filter 1 Values	644
Table 10.38	Filter Data Group	647
Table 10.39	Packet Reduction Field Description	652
Table 10.40	Extended Packet Reduction Field Description	652
Table 10.41	MACsec Host Control	662
Table 10.42	MACsec Owner Status.....	670
Table 10.43	MACsec Host Control Status	670
Table 10.44	MACsec Owner Status:	671
Table 10.45	MACsec Host Control Status	671
Table 10.46	TX Control Status:	671
Table 10.47	Status byte description	675
Table 10.48	Controller information Items.....	676
Table 10.49	MCTP Commands Support	689
Table 10.50	Host Driver Command Structure	693



Table 10.51	Status Structure Returned to Host Driver	693
Table 10.52	Driver Info Host Command	694
Table 10.53	Driver Info Host Status	694
Table 10.54	Disable RXEN Command	695
Table 10.55	OS2BMC Control Status	695
Table 10.56	Flash Read Command	695
Table 10.57	Flash Read Response	696
Table 10.58	Shadow RAM Read Command	696
Table 10.59	Shadow RAM Read Response	696
Table 10.60	Flash Write Command	697
Table 10.61	Flash Write Response	697
Table 10.62	Shadow RAM Write Response	698
Table 10.63	Flash Module Update Command	698
Table 10.64	Flash Module Update Response	699
Table 10.65	Apply Update Command	699
Table 10.66	Flash Block Erase Command	700
Table 10.67	Flash Block Erase Response	700
Table 10.68	Shadow RAM Dump Command	700
Table 10.69	Shadow RAM Dump Response	700
Table 10.70	Flash Info Command	701
Table 10.71	Flash Info Response	701
Table 10.72	Internal PHY Request	701
Table 10.73	Internal PHY Request	702
Table 11-1	TAP Controller Pins	705
Table 11-2	Main TAP Instructions Supported	705



1.0 Introduction

This document describes the external architecture (including device operation, pin descriptions, register definitions, etc.) for the integrated the Intel® Xeon® Processor D-1500 Product Family LAN controller that includes a dual-port 10 GbE Network Interface Controller.

It is intended as a reference for logical design group, architecture validation, firmware development, software device driver developers, board designers, test engineers, or anyone else who might need specific technical or programming information about the Intel® Xeon® Processor D-1500 Product Family LAN controller.

1.1 Intel® Xeon® Processor D-1500 Product Family LAN Controller Overview

The Intel® Xeon® Processor D-1500 Product Family LAN controller contains two independent 10 GbE Media Access Control (MACs) that support an XGMII-like interface link to the following integrated Physical Layer (PHY) device interfaces (see Figure 1.2):

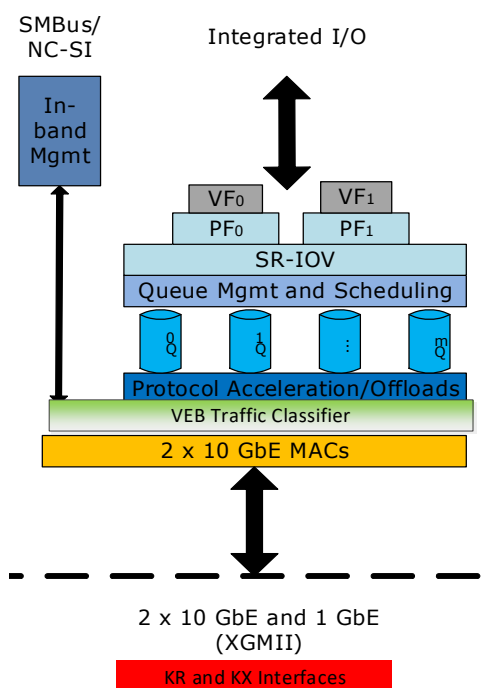


Figure 1.1. Intel® Xeon® Processor D-1500 Product Family LAN Controller Overview

1.2 Supported Modes of Operation

Table 1.1 lists which interfaces that are supported by the various SKUs.

Table 1.1. Supported SKUs

SKU	Ethernet Interface
D-1521, D-1526, D-1528, D-1519, D-1531, D-1541, D-1571, D-1581, D-1508, D-1517, D-1527, D-1537, D-1548, D-1559, D-1567, D-1577, D-1587	10GBASE-KR — Yes 10GBASE-T ¹ — Yes SFP+ ² — Yes 1000BASE-KX — Yes 1000BASE-T ¹ — Yes
D1507, D-1518, D-1557	10GBASE-KR — No 10GBASE-T ¹ — Yes SFP+ ² — Yes 1000BASE-KX — Yes 1000BASE-T ¹ — Yes

1. External PHY is required. Refer to the Platform Design Guide (PDG) for supported PHYs and additional routing requirements.
2. External PHY (Inphi* CS4227) is required. Refer to the PDG for additional layout requirements. SR, LR, and DAC modules are supported.

The Intel® Xeon® Processor D-1500 Product Family LAN controller connects (over XGMII) to two MDI interfaces with supported speeds of 10 GbE and 1 GbE Ethernet (see Figure 1.2).

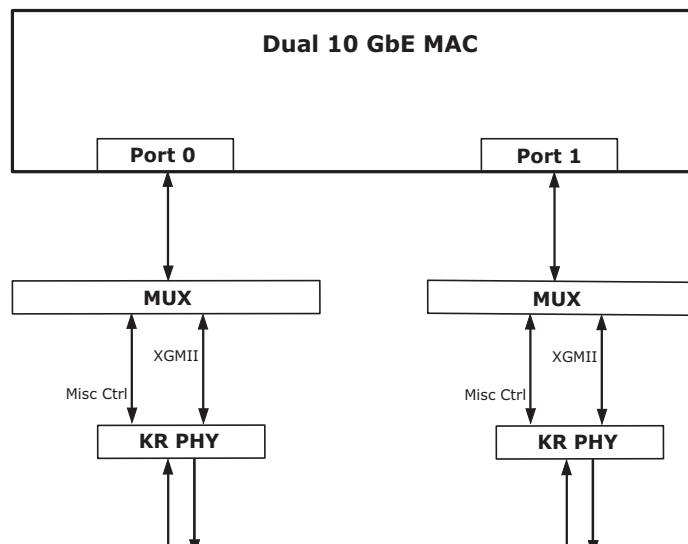


Figure 1.2. Intel® Xeon® Processor D-1500 Product Family LAN Controller Logical Block Diagram

The KR PHY interface provides the following physical interfaces and electrical modes:

- 10GBASE-KR for GbE backplane applications (IEEE802.3 clause 72)
- 10GBASE-KR FEC (IEEE 802.3 Clause 74)
- 1000BASE-KX for GbE backplane applications (IEEE802.3 clause 70)
- Auto-negotiation for backplane Ethernet (IEEE 802.3 Clause 73)
- 10GBASE-KR interface to:
 - 10 Gb SFI PHY (Inphi* CS4227)



— 10GBASE-T PHY (Intel® X557-AT2 (2-port))

Notes: Contact your Intel representative on support time lines for the different configuration.

Ethernet speeds of 100 Mb/s and 10 Mb/s are NOT supported.

Table 1-3 lists all the supported system configurations, operating modes, and link partners for the Intel® Xeon® Processor D-1500 Product Family LAN controller.

Note: Contact your Intel representative for support time lines.

Table 1-3. Supported System Configurations

Connection	Speed	Electrical Interface	3 rd Party PHYs/Switches
Backplane	1 GbE	1000BASE-KX	Intel® FM4000/FM5000/FM6000 Ethernet switches Broadcom* BCM5684x-series switches
Backplane	10 GbE	10GBASE-KR	Intel FM5000/FM6000 Ethernet switches Broadcom* BCM5684x-series switches
1000BASE-T	1 GbE	1000BASE-X	Marvell 88E1512 Marvell 88E1514
10GBASE-T	10 GbE	10GBASE-KR	Intel® X557-AT2 (2-port)
SFP+ SR/LR/ DAC	10 GbE	10GBASE-KR	Inphi CS4227

1.3 Conventions

This section defines the organization of registers and memory transfers as it relates to information carried over the network:

- Any register defined in Big Endian notation can be transferred as is to/from Tx and Rx buffers in the host memory. Big Endian notation is also referred to as being in network order or ordering.
- Any register defined in Little Endian notation must be swapped before it is transferred to/from Tx and Rx buffers in the host memory. Registers in Little Endian order are referred to being in host order or ordering.

Tx and Rx buffers are defined as being in network ordering and are transferred as-is over the network.

Note: Registers not transferred on the wire are defined in Little Endian notation. Registers transferred on the wire are defined in Big Endian notation, unless specified differently.

1.4 References

The Intel® Xeon® Processor D-1500 Product Family LAN controller implements features from the following specifications:



IEEE Specifications

- IEEE 1149.6 standard for Boundary Scan (MDI pins excluded)
- IEEE 802.3-2012, Clauses 70 through 74.
- IEEE standard 1149.1, 2001 Edition (JTAG). Institute of Electrical and Electronics Engineers (IEEE).
- IEEE standard 802.1Q for VLAN.
- IEEE 1588 International Standard, Precision clock synchronization protocol for networked measurement and control systems, 2004-09.
- IEEE P802.1AE/D5.1, Media Access Control (MAC) Security, January 19, 2006.
- IEEE 802.3-2012, Clause 78

PCI-SIG Specifications

- PCI Express 3.0 Base specification
- PCI Bus Power Management Interface Specification, Rev. 1.2, March 2004.
- PICMG3.1 Rev 2.0.
- Single Root I/O Virtualization and Sharing, Revision 1.1, September 8, 2009.

IETF Specifications

- IPv4 specification (RFC 791)
- IPv6 specification (RFC 2460)
- TCP specification (RFC 793)
- UDP specification (RFC 768)
- ARP specification (RFC 826)

Manageability Documents

- DMTF Network Controller Sideband Interface (NC-SI) Specification rev 1.0.0, May 2009
- System Management Bus (SMBus) Specification, SBS Implementers Forum, Ver. 2.0, August 2000

1.5 Related Documents

Note: Contact your Intel representative for document access.

- Intel® FM5000/FM6000 Ethernet Switch Datasheet
- Intel® Ethernet Switch FM6000 Series - Layout Checklist
- Intel® Ethernet Switch FM5000/FM6000 Series - User Guide
- Intel® Ethernet Switch FM5000/FM6000 Series - Specification Update
- Intel® X557-AT/AT2/AT4 10 GbE PHY Datasheet



1.6 Features

Table 1.2 to Table 1.8 list the Intel® Xeon® Processor D-1500 Product Family LAN controller features.

Table 1.2. Network Features

Feature	Intel® Xeon® Processor D-1500 Product Family LAN Controller	82599
Compliant with the 10 Gb/s and 1 Gb/s Ethernet/802.3 (KX) Specifications	Y	Y
Compliant with the 10 Gb/s 802.3 (KR) Specification	Y	Y
Support of 10GBASE-KR FEC	Y	Y
Full-Duplex Operation at all Supported Speeds	Y	N
802.3az Energy Efficient Ethernet (EEE) Support	Y ¹	N
Support Jumbo Frames of up to 15.5 KB	Y ²	Y ²
MDIO Interface Clause 45	Y	N
Flow Control Support: Send/Receive Pause Frames and Receive FIFO Thresholds	Y	Y
Statistics for Management and RMON	Y	Y
802.1q VLAN Support	Y	Y
Double VLAN	Y	Y

1. EEE is ONLY supported for 10GBASE-T interfaces via the X557-AT2 PHY. The Intel® Xeon® Processor D-1500 Product Family LAN controller does not support EEE feature.
2. The Intel® Xeon® Processor D-1500 Product Family LAN controller and the 82599 support full-size 15.5 KB jumbo packets while in a basic mode of operation. When security engines enabled or virtualization is enabled, or OS2BMC is enabled, then the Intel® Xeon® Processor D-1500 Product Family LAN controller and the 82599 support 9.5 KB jumbo packets. Packets to/from the MC longer than 2 KB are filtered out.



Table 1.3. Host Interface Features

Feature	Intel® Xeon® Processor D-1500 Product Family LAN Controller	82599
Integrated I/O	Y	N
64-bit Address Support for Systems Using More Than 4 GB of Physical Memory	Y	Y
Outstanding Requests for Tx Data Buffers	16	16
Outstanding Requests for Tx Descriptors	8	8
Outstanding Requests for Rx Descriptors	8	8
Credits for P-H/P-D/NP-H/NP-D (shared for the two ports)	4/255/4/255 ¹	16/16/4/4
Credits for C-H/C-D (shared for the two ports)	4/255	N
Max Payload Size Supported	64 thru 512 bytes	512 bytes
Max Request Size Supported	64 thru 2 KB	2 KB
Vital Product Data (VPD)	Y	Y
End to End CRC (ECRC)	N/A	Y
TLP Processing Hints (TPH)	Y	N
Latency Tolerance Reporting (LTR)	N	N
ID-Based Ordering (IDO)	Y	N
Access Control Services (ACS)	Y	N
ASPM Optional Compliance Capability	Y	N
PCIe Functions Off Via Pins, While LAN Ports Are On	Y	N

1. Each data credit accounts for 16 bytes.

Table 1.4. Miscellaneous Features

Feature	Intel® Xeon® Processor D-1500 Product Family LAN Controller	82599
Serial Flash Interface	Y	Y
Configurable LED Operation for Software or OEM Customization of LED Displays	Y	Y
Protected NVM Space for Private Configuration	Y	Y
Device Disable Capability	Y	Y
Watchdog Timer	Y	Y
Time Sync (IEEE 1588)	Y	Y



Table 1.5. LAN Functions Features

Feature	Intel® Xeon® Processor D-1500 Product Family LAN Controller	82599
Programmable Host Memory Receive Buffers	Y	Y
Descriptor Ring Management Hardware for Transmit and Receive	Y	Y
ACPI Register Set and Power Down Functionality Supporting D0 and D3 States	Y	Y
Software-Controlled Global Reset Bit (Resets Everything Except the Configuration Registers)	Y	Y
Software-Definable Pins (SDP) (per port)	2	8
SDP Pins Can Be Configured as General Purpose Interrupts	Y	Y
Wake on LAN (WoL)	Y	Y
IPv6 Wake-up Filters	Y	Y
Configurable (through NVM) Wake-up Flexible Filters	Y	Y
Default Configuration by NVM for all LEDs for Pre-Driver Functionality	Y	Y
LAN Function Disable Capability	Y	Y
Programmable Memory Transmit Buffers	160 KB / port	160 KB / port
Programmable Memory Receive Buffers	384 KB / port	512 KB / port

Table 1.6. LAN Performance Features (Sheet 1 of 2)

Feature	Intel® Xeon® Processor D-1500 Product Family LAN Controller	82599
TCP/UDP Segmentation Offload	256 KB in all modes	256 KB in all modes
TSO Interleaving for Reduced Latency	Y	Y
TCP Receive Side Coalescing (RSC)	32 flows / port	32 flows / port
Data Center Bridging (DCB), IEEE Compliance to: Priority Groups (up to 8) and Bandwidth Allocation (ETS) IEEE802.1Qaz Priority-based Flow Control (PFC) IEEE802.1Qbb	N N	Y Y
Rate Limit VM Tx Traffic per TC (per TxQ)	Y	Y
IPv6 Support for IP/TCP and IP/UDP Receive Checksum Offload	Y	Y
Fragmented UDP Checksum Offload for Packet Reassembly	Y	Y
FCoE Tx / Rx CRC Offload	N	Y
FCoE Transmit Segmentation	N	256 KB
FCoE Coalescing and Direct Data Placement	N	512 outstanding read - write requests / port
Message Signaled Interrupts (MSI)	Y	Y
Message Signaled Interrupts (MSI-X)	Y	Y
Interrupt Throttling Control to Limit Maximum Interrupt Rate and Improve CPU Use	Y	Y
Rx Packet Split Header	Y	Y
Multiple Rx Queues (RSS)	Y (multiple modes)	Y (multiple modes)




Table 1.6. LAN Performance Features (Sheet 2 of 2)

Feature	Intel® Xeon® Processor D-1500 Product Family LAN Controller	82599
Flow Director Filters: up to 32 KB Flows by Hash Filters or up to 8 KB Perfect Match Filters	Y	Y
Number of Rx Queues (per port)	128	128
Number of Tx Queues (per port)	128	128
TCP Timer Interrupts	Y	Y
Relax Ordering	Y	Y
DMA Coalescing	Y	N

Table 1.7. Virtualization Features

Feature	Intel® Xeon® Processor D-1500 Product Family LAN Controller	82599
Support for Virtual Machine Device Queues (VMDq1 and Next Generation VMDq)	64	64
Virtual Functions (VFs) in SR-IOV Mode	64	64
L2 Ethernet MAC Address Filters (Unicast and Multicast)	128	128
L2 VLAN Filters	64	64
PCI-SIG SR IOV	Y	Y
Multicast and Broadcast Packet Replication	Y	Y
Packet Mirroring	Y	Y
Packet Loopback	Y	Y
Traffic Shaping	Y	Y
Anti Spoof	MAC, VLAN, Ethertype	N
Forwarding Modes	MAC, VLAN E-tag	N
VEB Support	Y	N


Table 1.8. Manageability Features (Sheet 1 of 2)

Feature	Intel® Xeon® Processor D-1500 Product Family LAN Controller	82599
Advanced Pass Through-compatible Management Packet Transmit/Receive Support	Y	Y
SMBus Interface to an External MC	Y	Y
New Management Protocol Standards Support (NC-SI) Interface to an External MC	Y	Y
L2 Address Filters	4	4
VLAN L2 Filters	8	8
Flex L3 Port Filters	16	16
Flexible TCO Filters	1	4



Table 1.8. Manageability Features (Sheet 2 of 2)

Feature	Intel® Xeon® Processor D-1500 Product Family LAN Controller	82599
L3 Address Filters (IPv4)	4	4
L3 Address Filters (IPv6)	4	4
Host-Based Application-to-BMC Network Communication Patch (OS2BMC)	Y	N
Flexible MAC Address	Y	N
MC Inventory of LOM Device Information	Y	N
iSCSI Boot Configuration Parameters via MC	Y	N
MC Monitoring	Y	N
NC-SI to MC	Y	N
NC-SI Arbitration	Y	N
NC-SI Package ID Via SDP Pins	Y	N
NC-SI Flow Control	Y	N



NOTE: *This page intentionally left blank.*



2.0 Pin Interface

This section describes all the LAN related external pins for the Intel® Xeon® Processor D-1500 Product Family LAN controller.

2.1 Signal Type Definition

Signal	Definition
In	
Out (O)	
T/s	Tri-state is a bi-directional, tri-state input/output pin.
O/d	Open drain enables multiple devices to share as a wire-OR.
A-in	Analog input signals.
A-out	Analog output signals.
A-Inout	Bi-directional analog signals.
B	Input BIAS.
In-Only	
Out-Only	
LVDS-O	Low voltage differential signal - output
Pup	Pull up
Pdn	Pull down
I/O - LVCMOS	



2.2 High-Speed Network Data Interface

Group Name	Signal Name	Type	Name and Function
(n = 0 thru 1 for port#) (ln = 0 thru 3 for lane#)			
High Speed TX I/O and RX I/O	LAN_KR_TX<n>_P LAN_KR_TX<n>_N	OUT	KR PHY Tx Differential Output.
	LAN_KR_RX<n>_P LAN_KR_RX<n>_N	IN	KR PHY Rx Differential Input.
High Speed TX I/O and RX I/O	LAN_KX4_TX<n>_L<ln>_P LAN_KX4_TX<n>_L<ln>_N	OUT	KX4 ¹ PHY Tx Differential Output. Leave these pins as a No Connect.
	LAN_KX4_RX<n>_L<ln>_P LAN_KX4_RX<n>_L<ln>_N	IN	KX4 ¹ PHY Rx Differential Output. Leave these pins as a No Connect.


1. The KX4 interface is not supported.

2.3 Low-Speed Network Data Interface

Group Name	Signal Name	Type	Name and Function
(n = 0 thru 1 for port#) (ln = 0 thru 3 for lane#)			
Management Interface for Physical-Layer (PHY) Components and LED Indicators. Note: MDIO and LED functions share the same interface pins, so only one can be operational at a time. These interfaces can be configured using NVM	LAN_I2C_SCL[n]	I/O - OD	I2C Clock, of the 2-wire management interface used to access the management registers of an external optical module of port n. One clock pulse is generated for each data bit transferred.
	LAN_I2C_SDA[n]	I/O - OD	I2C Data, of the 2-wire management interface used to access the management registers of an external optical module of port n. Stable during the high period of the clock (unless it is a start or stop condition).
	LAN_MDIO<n>_LED<n>_[1:0]	I/O - LVCMOS	Management Data. Bi-directional signal for serial data transfers between the MAC and the external PHY. Tristate buffer requires an external pull-up device. Programmable LEDs that can be used for Link Up indication.
	LAN_MDC<n>_LED<n>_[1:0]	O - LVCMOS	Management Clock. Clock output for accessing the external PHY management registers. MDC clock frequency is proportional to link speed. At 10 Gb/s Link speed, MDC frequency can be set up to 2.4 MHz. Programmable LEDs that can be used for Link Up indication.
	LAN_MDIO_DIR_CTL_<n>	O - LVCMOS	MDIO Direction Control Pin. Used to control the direction of the level shifters.



Group Name	Signal Name	Type	Name and Function
(n = 0 thru 1 for port#) (ln = 0 thru 3 for lane#)			
Software Defined Pins (SDPs) For example, an external PHY topology (10 GbE SFP+ and 10GBASE-T) can detect an interrupt only on the rising edge of the signal from the PHY. Some discrete PHYs (like the X577-AT2 and Inphi CS4227) are designed for active-low signaling so adding a trigger inverter between the PHY and the level shifter on each pin is required to ensure proper functionality.	LAN_SDP<n>_[1:0]	I/O - LVCMOS	2 general purpose SDP Pins per port. Can be used to support IEEE1588 auxiliary devices or input for external interrupts, etc.
Dedicated System Management Interface - SMBus	LAN_SMBALRT_N	I/O - OD	SMBus Alert. Acts as an interrupt pin of a slave device on the SMBus.
	LAN_SMBCLK	I/O - OD	SMBus Clock. One clock pulse is generated for each data bit transferred.
	LAN_SMBDATA	I/O - OD	SMBus Data. Stable during the high period of the clock (unless it is a start or stop condition).
Dedicated Flash Interface	LAN_FLSH_CLK	O - LVCMOS	Flash serial clock. Operates at the maximum frequency of 25 MHz.
	LAN_FLSH_CS_N	O - LVCMOS	Flash chip select output.
	LAN_FLSH_DI	O - LVCMOS	Serial data output to the Flash.
	LAN_FLSH_DO	I - LVCMOS	Serial data input from the Flash.



NOTE: *This page intentionally left blank.*



3.0 Interconnects

3.1 Host Primary Interface

The Intel® Xeon® Processor D-1500 Product Family LAN controller communicates with the host CPUs and host memories over an integrated I/O interface. This interface is an Intel standard bus defined for SoCs and replaces the PCIe interface and functionality of the discrete LAN controller.

3.1.1 Host Interface Architecture, Transaction and Link Layer Properties

- All PCI functions are native PCIe functions
- Credit-based flow control
- Packet sizes/formats:
 - Maximum upstream (write) payload size of 512B
 - Configurable sizes are: 128 thru 512
 - Maximum downstream (read) payload size of 2KB
 - Configurable sizes are: 128 thru 2 KB
- Reset/initialization:
 - Credit negotiation performed by hardware
- Transaction layer mechanisms
 - 64-bit memory address spaces
 - Removal of I/O BAR (optional)
 - Relaxed ordering
 - Flow control update time out mechanism
 - ID-based ordering (IDO)
 - Function-Level Reset (FLR)
 - TLP Processing Hints (TPH)
 - Reliability
 - Advanced Error Reporting (AER)
- Power management
 - Wake capability
 - Latency Tolerance Reporting (LTR)
- DFT and DFM support for high-volume manufacturing
- The Intel® Xeon® Processor D-1500 Product Family LAN controller supports the following extended capabilities:
 - Device Serial Number
 - Alternative RID Interpretation (ARI)



- Single Root I/O Virtualization (SR-IOV)
- TPH Requester
- Access Control Services (ACS)
- Software configuration mechanism:
 - Uses PCI configuration and bus enumeration model
 - PCIe-specific configuration registers mapped via PCI extended capability mechanism

3.1.2 PCIe Transaction Layer

3.1.2.1 Transaction Types Accepted by the Intel® Xeon® Processor D-1500 Product Family LAN Controller

Table 3.1 lists the transactions accepted by the device and their attributes.

Table 3.1. Transaction Types Accepted by the Transaction Layer

Transaction Type	Tx Layer Reaction	Hardware Should Keep Data From Original Packet
Configuration Read Request	CPLH + CPLD	Requester ID, TAG, attribute
Configuration Write Request	CPLH	Requester ID, TAG, attribute
Memory Read Request	CPLH + CPLD	Requester ID, TAG, attribute
Memory Write Request	-	-
I/O Read Request	CPLH + CPLD	Requester ID, TAG, attribute
I/O Write Request	CPLH	Requester ID, TAG, attribute
Read Completions	-	-
Message	-	-

3.1.2.1.1 Size of Target Accesses

3.1.2.1.1.1 Memory accesses

Rules for accesses to the CSR space (both memory BAR and MSI-X BAR):

- Write accesses
 - Zero-length writes have no internal impact (nothing written, no effect such as clear-by-write). The transaction is treated as a successful operation (no error event).
- CSR writes are 32-bit or 64-bit only. Larger or partial CSR writes are handled as completer abort - data is dropped and an error is generated per PCIe rules read accesses
 - Partial reads with at least one byte disabled are handled as a full read. Any side effect of the full read (such as clear by read) is also applicable to partial reads. The completion on PCIe follows the specification rules regarding the number of bytes reported in the completion.
 - Zero-length reads generate a completion, but the register is not accessed and undefined data is returned.
 - CSR reads are 32-bit or 64-bit only. Larger CSR read requests are handled as completer abort - the completion includes a CA status and an error is generated per PCIe rules.
 - Some 64-bit reads are handled atomically (such as not interleaved with any other requests). This applies mainly to reading counters, where all 64 bits need to be read simultaneously. Such registers are explicitly marked in their description.

Rules for accessing the Flash space in the memory BAR or the expansion ROM BAR:



- Write accesses
 - Writes to Flash are 8-bit wide only
 - Any larger write accesses are handled as completer abort - data is dropped and an error is generated per PCIe rules
- Read accesses
 - Reads to Flash are 32-bit wide
 - Partial reads with at least one byte disabled are handled internally as a full read. That is, any side effect of the full read (such as clear by read) is also applicable to partial reads. The completion on PCIe follows the specification rules regarding the number of bytes reported in the completion
 - Larger CSR read requests are handled as completer abort - the completion includes a CA status and an error is generated per PCIe rules

I/O accesses

Rules for accesses to the I/O BAR:

- Write accesses
 - Write accesses are 32-bit wide
 - Zero-length writes have no internal impact (nothing written, no effect such as clear-by-write). The transaction is treated as a successful operation (no error event).
 - Other accesses (partial writes, larger writes) are handled as completer abort - data is dropped and an error is generated per PCIe rules
- Read accesses
 - Reads to the I/O BAR are 32-bit wide
 - Partial reads with at least one byte disabled are handled internally as a full read. That is, any side effect of the full read (such as clear by read) is also applicable to partial reads. The completion on PCIe follows the specification rules regarding the number of bytes reported in the completion.
 - Larger CSR read requests are handled as completer abort - the completion includes a CA status and an error is generated per PCIe rules

3.1.2.1.2 Support for Dynamic Changes

The Intel® Xeon® Processor D-1500 Product Family LAN controller captures the bus number and device number per each configuration write request. However, a dynamic change of the bus number or device number is not supported. Rather, the PCIe link should be quiescent prior to such a change, including reception of all completion for previous requests.

3.1.2.2 Transaction Types Initiated by the Intel® Xeon® Processor D-1500 Product Family LAN Controller

Table 3.2. Transaction Types Initiated by the Transaction Layer

Transaction Type	Payload Size	Tx Layer Reaction
Configuration Read Request Completion	Dword	CPLH + CPLD
Configuration Write Request Completion	-	CPLH
I/O Read Request Completion	Dword	CPLH + CPLD
I/O Write Request Completion	-	CPLH



Table 3.2. Transaction Types Initiated by the Transaction Layer

Transaction Type	Payload Size	Tx Layer Reaction
Read Request Completion	Dword/Qword	CPLH + CPLD
Memory Read Request	-	NPH
Memory Write Request	<= MAX_PAYLOAD_SIZE	PH + PD
Message	64 bytes	PH

Note:

Configuration values:

- Max Payload Size - The value of the Max_Payload_Size Supported field in the Device Capabilities register is loaded from the NVM.
 - Hardware default is 512B.
 - System software then programs the actual value into the Max_Payload_Size field of the Device Control register.
 - Non-ARI mode: If not all functions are programmed with the same value, the max payload size used for all functions is the minimum value programmed among all functions.
 - ARI mode: Max_Payload_Size is determined solely by the setting in Function 0
- Max_Read_Request_Size - The Intel® Xeon® Processor D-1500 Product Family LAN controller supports read requests of up to 2KB.
 - The actual maximum size of a read request is defined as the minimum (2KB, Max_Read_Request_Size field in the Device Control register).

The number of outstanding memory read requests is bounded by the following:

- The total number of outstanding requests is not more than 32 requests. These are shared by all sources for memory reads.

3.1.2.2.1 Data Alignment

Requests must never specify an address/length combination that causes a memory space access to cross a 4 KB boundary. The Intel® Xeon® Processor D-1500 Product Family LAN controller therefore breaks requests into 4 KB-aligned requests (if needed). This does not place any requirement on software. However, if software allocates a buffer across a 4 KB boundary, hardware issues multiple requests for the buffer. Software should consider aligning buffers to a 4 KB boundary in cases where it improves performance. The maximum size of a read request is defined as the minimum (2 KB bytes, Max_Read_Request_Size).

The general rules for packet alignment are as follows. Note that these apply to all Broadwell-DE LAN controller requests (read/write):

- The length of a single request does not exceed the PCIe limit of MAX_PAYLOAD_SIZE(512B) for write and MAX_READ_REQ (2KB) for read.
- The length of a single request does not exceed the Intel® Xeon® Processor D-1500 Product Family LAN controller internal limitations.
- A single request does not span across different memory pages as noted by the 4KB boundary alignment previously mentioned.

If a request can be sent as a single packet and still meet the general rules for packet alignment, then it is not broken at the cache line boundary but rather sent as a single packet. However, if any of the three general rules require that the request is broken into two or more packets, then the request is broken at the cache line boundary.



For requests with data payload, if the payload size is larger than (MAX_PAYLOAD_SIZE - CACHELINE_SIZE), then the request is broken into multiple TLPs starting at the first cache line boundary following the (MAX_PAYLOAD_SIZE - CACHELINE_SIZE) bytes. For example, if MAX_PAYLOAD_SIZE = 256 bytes and CACHELINE_SIZE = 64 bytes, a 1 KB request starting at address 0x...10 is broken into TLPs such that the first TLP contains 240bytes of payload (since 240bytes+ 0x10 = 256bytes is on cache line boundary).

The system cache line size is controlled by the *PCI_CNFG2.CACHELINE_SIZE* bit, loaded from the NVM. Note that the Cache Line Size register in the PCI configuration space is not related to the *PCI_CNFG2.CACHELINE_SIZE* and is solely for software use

3.1.2.3 Messages

3.1.2.3.1 Received Messages

Message packets are special packets that carry a message code. The upstream device transmits special messages to the Intel® Xeon® Processor D-1500 Product Family LAN controller by using this mechanism. The transaction layer decodes the message code and responds to the message accordingly.

Table 3.3. Supported Message in the Intel® Xeon® Processor D-1500 Product Family LAN Controller (as a Receiver)

Message Code [7:0]	Routing r2r1r0	Message	Intel® Xeon® Processor D-1500 Product Family LAN Controller Later Response
0x40, 0x41, 0x43, 0x44, 0x45, 0x47, 0x48	100b	Ignored messages (used to be hot-plug messages)	Silently drop.
0x50	100b	Slot power limit support (has one Dword data)	Silently drop.
0x7E	000b, 010b, 011b, 100b	Vendor defined type 0	Drop and handle as an Unsupported request
0x7F	100b	Vendor defined type 1	Silently drop.
0x7F	010b, 011b, 000b	Vendor defined type 1	Silently drop.
0x00	011b	Unlock	Silently drop.

3.1.2.3.2 Transmitted Messages

The transaction layer is also responsible for transmitting specific messages to report internal/external events (such as interrupts and PMEs).

3.1.2.4 Transaction Attributes

3.1.2.4.1 Traffic Class (TC) and Virtual Channels (VC)

The Intel® Xeon® Processor D-1500 Product Family LAN controller only supports TC = 0b and VC = 0b (default).

3.1.2.4.2 TLP Processing Hints (TPH)

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports the TPH capability defined in the PCIe specification. It does not support extended TPH requests.



Existence of a TLP Processing Hint (TPH) is indicated on the PCIe link by setting the TH bit in the TLP header. Using the PCIe TLP Steering Tag (ST) and Processing Hints (PH) fields, The Intel® Xeon® Processor D-1500 Product Family LAN controller can provide hints to the root complex about the destination (socket ID) and about data access patterns (locality in the cache) when executing DMA memory writes or read operations.

The Intel® Xeon® Processor D-1500 Product Family LAN controller exposes a PCIe TPH capability structure (See [Section 9.2.4.5](#)) with no steering table.

See [Section 7.5](#) for details on the usage of TPH.

3.1.2.5 Ordering Rules

The Intel® Xeon® Processor D-1500 Product Family LAN controller meets the PCIe ordering rules by following the PCI simple device model:

1. Deadlock Avoidance – The Intel® Xeon® Processor D-1500 Product Family LAN controller meets the PCIe ordering rules that prevent deadlocks:
 - a. Posted writes overtake stalled read requests. This applies to both target and master directions. For example, if master read requests are stalled due to lack of credits, master posted writes are allowed to proceed. On the target side, it is acceptable to timeout on stalled read requests in order to allow later posted writes to proceed.
 - b. Target posted writes overtake stalled target configuration writes.
 - c. Completions overtake stalled read requests. This applies to both target and master directions. For example, if master read requests are stalled due to lack of credits, completions generated by the Intel® Xeon® Processor D-1500 Product Family LAN controller are allowed to proceed.
2. Descriptor/Data Ordering — The Intel® Xeon® Processor D-1500 Product Family LAN controller insures that a Rx descriptor is written back on PCIe only after the data that the descriptor relates to is written to the PCIe link.
3. MSI and MSI-X Ordering Rules – System software might change the MSI or MSI-X tables during run-time. Software expects that interrupt messages issued after the table has been updated are using the updated contents of the tables.
 - a. Since software doesn't know when the tables are actually updated in the Intel® Xeon® Processor D-1500 Product Family LAN controller, a common scheme is to issue a read request to the MSI or MSI-X table (a PCI configuration read for MSI and a memory read for MSI-X). Software expects that any message issued following the completion of the read request, is using the updated contents of the tables.
 - b. Once an MSI or MSI-X message is issued using the updated contents of the interrupt tables, any consecutive MSI or MSI-X message does not use the contents of the tables prior to the change.
4. The Intel® Xeon® Processor D-1500 Product Family LAN controller meets the rules relating to independence between target and master accesses:
 - a. The acceptance of a target posted request does not depend upon the transmission of any TLP.
 - b. The acceptance of a target non-posted request does not depend upon the transmission of a non-posted request.
 - c. Accepting a completion does not depend upon the transmission of any TLP.

3.1.2.5.1 Relaxed Ordering

The Intel® Xeon® Processor D-1500 Product Family LAN controller takes advantage of the relaxed ordering rules in PCIe. By setting the relaxed ordering bit in the packet header, the Intel® Xeon® Processor D-1500 Product Family LAN controller enables the system to optimize performance in the following cases:



1. Relaxed ordering for descriptor and data reads — When the Intel® Xeon® Processor D-1500 Product Family LAN controller masters a read transaction, its split completion has no ordering relationship with the writes from the CPUs (same direction). It should be allowed to bypass the writes from the CPUs.
2. Relaxed ordering for receiving data writes — When the Intel® Xeon® Processor D-1500 Product Family LAN controller masters receive data writes, it also enables them to bypass each other in the path to system memory because software does not process this data until their associated descriptor writes are done.
3. The Intel® Xeon® Processor D-1500 Product Family LAN controller cannot relax ordering for receive descriptor writes or an MSI write.

Relaxed ordering is enabled in the Intel® Xeon® Processor D-1500 Product Family LAN controller by clearing the *CTRL_EXT.RO_DIS* bit. Relaxed ordering is further controlled through the *Enable Relaxed Ordering* bit in the PCIe Device Control register.

3.1.2.5.2 ID-based Ordering (IDO)

IDO was introduced in the PCIe rev. 2.1 specification. When enabled, The Intel® Xeon® Processor D-1500 Product Family LAN controller sets IDO in all applicable TLPs defined in the PCIe specification.

This capability enables a supporting root complex to relax ordering rules for TLPs sent by different requesters.

IDO is enabled when all of the following conditions are met:

- The NVM *PCI_CAPSUP.IDO Enable* bit is set ([Section 6.4.4](#) and [Section 8.2.2.6.6](#))
- The PCIe *IDO Request Enable* bit (for requests) or the *IDO Completion Enable* bit (for completions) in Device Control 2 register is set

3.1.3 Error Events and Error Reporting

3.1.3.1 General Description

PCIe defines three error reporting paradigms: the baseline capability, the Advanced Error Reporting (AER) capability, and a proprietary mechanism. The baseline error reporting capabilities are required of all PCIe devices and define the minimum error reporting requirements. The AER capability is defined for more robust error reporting and is implemented with a specific PCIe capability structure. Both mechanisms are supported by the Intel® Xeon® Processor D-1500 Product Family LAN controller. The proprietary error reporting mechanism used for error better handled by the software device driver using internal CSRs is described in [Section 3.1.3.7](#).

The *SERR# Enable* and the *Parity Error* bits from the Legacy Command register also take part in the error reporting and logging mechanism.

In a multi-function device, PCIe errors that are not related to any specific function within the device are logged in the corresponding status and logging registers of all functions in that device. [Figure 3.1](#) shows, in detail, the flow of error reporting in the Intel® Xeon® Processor D-1500 Product Family LAN controller.

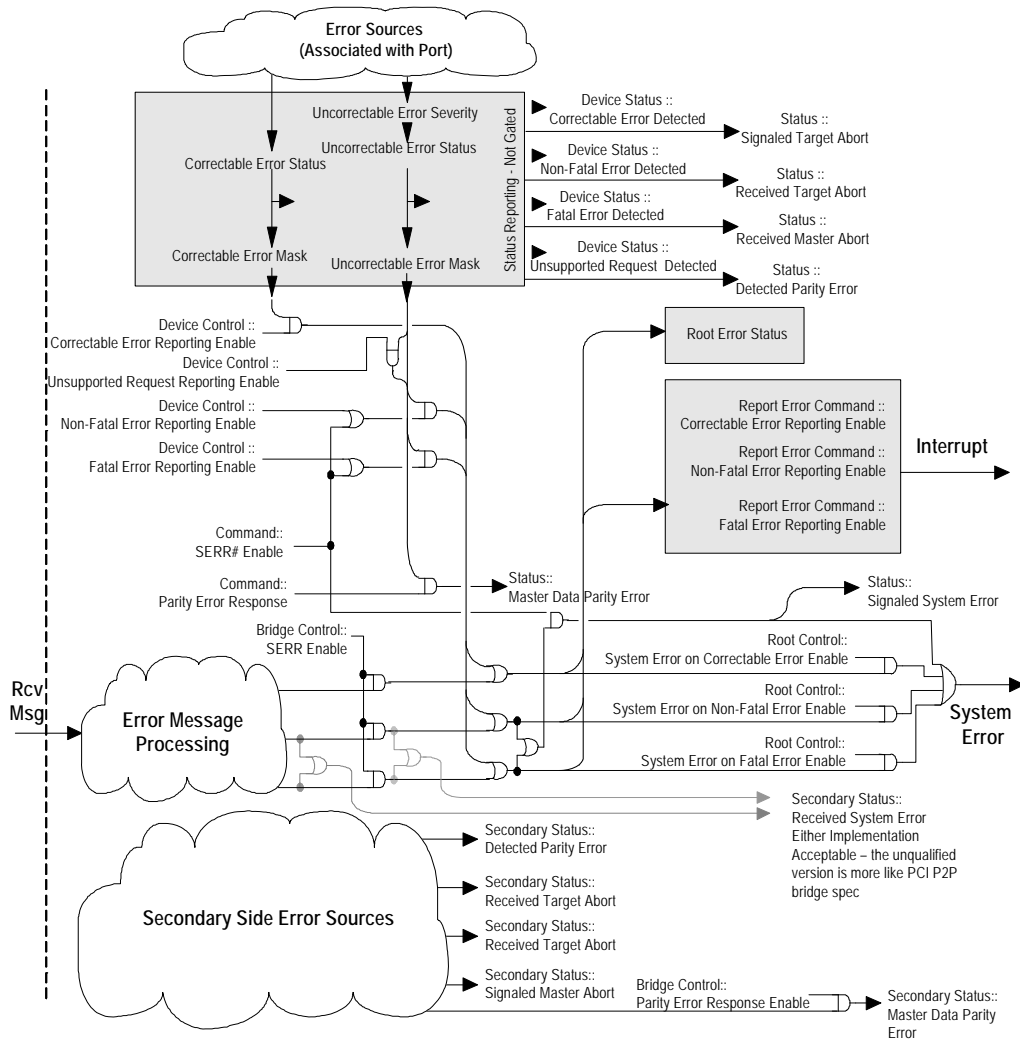


Figure 3.1. Error Reporting Mechanism




3.1.3.2 Error Events

Table 3.4 lists the error events identified by the Intel® Xeon® Processor D-1500 Product Family LAN controller and the response in terms of logging, reporting, and actions taken. Refer to the PCIe specification for the effect on the PCI Status register.

Table 3.4. Response and Reporting of PCIe Error Events

Error Name	Error Events	Default Severity	Action
TLP Errors			
Poisoned TLP Received	<ul style="list-style-type: none"> TLP With Error Forwarding (EP = 1b) Data Parity 	Uncorrectable ERR_NONFATAL Log Header	If completion TLP: Error is non-fatal (default case) <ul style="list-style-type: none"> Send error message if advisory Retry the request once and send advisory error message on each failure If fails, send uncorrectable error message Error is defined as fatal <ul style="list-style-type: none"> Send uncorrectable error message
Unsupported Request (UR)	<ul style="list-style-type: none"> Receipt of TLP with unsupported Request Type Receipt of an Unsupported Vendor Defined Type 0 Message Invalid Message Code Wrong Function Number Received TLP Outside BAR Address Range Receipt of a Request TLP during D3hot, other than Configuration and Message requests 	Uncorrectable ERR_NONFATAL Log header	Send Completion With UR
Completion Timeout	<ul style="list-style-type: none"> Completion Timeout Timer Expired 	Uncorrectable ERR_NONFATAL	Error is non-fatal (default case) <ul style="list-style-type: none"> Send error message if advisory Error is defined as fatal <ul style="list-style-type: none"> Send uncorrectable error message
Completer Abort	<ul style="list-style-type: none"> Received Target Access With illegal data size per Section 3.1.2.1.1 	Uncorrectable. ERR_NONFATAL Log header	Send completion with CA
Unexpected Completion	<ul style="list-style-type: none"> Received Completion Without a Request For It (Tag, ID, etc.) 	Uncorrectable ERR_NONFATAL Log Header	Discard TLP
Malformed TLP (MP)	<ul style="list-style-type: none"> Data Payload Exceed Max_Payload_Size Received TD field value does not correspond with the observed size PM Messages That Don't Use TCO Usage of Unsupported VC Target request crosses a 4KB boundary 	Uncorrectable ERR_FATAL Log Header	Note: Drop the Packet, Free FC Credits
Completion with Unsuccessful Completion Status		No Action (already done by originator of completion)	Free FC Credits
Command Parity	Parity error is detected on the command	Uncorrectable ERR_FATAL	



3.1.3.3 Completion Timeout Mechanism

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports completion time out as defined in the PCIe specification.

The Intel® Xeon® Processor D-1500 Product Family LAN controller controls the following aspects of completion time out:

- Disabling or enabling completion timeout
 - The PCIe *Completion Timeout Disable Supported* bit in the Device Capabilities 2 register is hard wired to 1b to indicate that disabling completion timeout is supported
 - The PCIe *Completion Timeout Disable* bit in Device Control 2 register controls whether completion timeout is enabled
- A programmable range of timeout values
 - The Intel® Xeon® Processor D-1500 Product Family LAN controller supports all four ranges as programmed in the *Completion Time out Ranges Supported* field of the Device Capabilities 2 register. The actual completion time out value is written in the *Completion Time out Value* field of Device Control 2 register.

The following sequence takes place when completion timeout is detected:

- The appropriate message is sent on PCIe as listed in [Table 3.4](#).
- The affected queue or client takes action based on the nature of the original request.
- An interrupt is issued to the respective PF.

3.1.3.4 Error Forwarding (TLP Poisoning)

If a TLP is received with an error-forwarding trailer, the packet is dropped and is not delivered to its destination, the Intel® Xeon® Processor D-1500 Product Family LAN controller then reacts as listed in [Table 3.4](#).

The following sequence takes place when a poisoned TLP is received:

- The appropriate message is sent on PCIe as listed in [Table 3.4](#).
- An interrupt is issued.
- If the TLP is a completion, a completion time out follows at some later time. Processing continues as described in [Section 3.1.3.3](#).

System logic is expected to trigger a system-level interrupt to signal the operating system of the problem. Operating systems can then stop the process associated with the transaction, re-allocate memory to a different area instead of the faulty area, etc.

3.1.3.5 Completion With Unsuccessful Completion Status

A completion arriving with an unsuccessful completion status (either UR or CA) is dropped and not delivered to its destination. A completion time out follows at some later time. Processing continues as described in [Section 3.1.3.3](#).

3.1.3.6 Blocking on Upper Address

The *PCI_UPADD* register blocks master accesses from being sent out on PCIe if the TLP address exceeds some upper limit. Bits [31:1] correspond to bits [63:33] in the PCIe address space, respectively.



When a bit is set in *GLPCI_UPADD[31:1]*, any transaction, in which the corresponding bit in its address is set, is blocked and not sent over PCIe. If all register bits are cleared, there is no effect (such as no TLPs are blocked by this mechanism).

The *PCI_UPADD* register is loaded from the NVM with a value allowing all addresses to pass. The software device driver should override this value with a system dependent value.

Processing a blocked transaction:

- Write transaction:
 - The transaction is dropped.
 - Set the Exceeded upper address limit (write requests) event in the PCIe errors register (see [Section 3.1.3.7](#)).
 - An interrupt is issued as described in [Section 3.1.3.7](#).
- Read transaction:
 - The transaction is dropped.
 - Set the Exceeded upper address limit (read requests) event in the PCIe errors register (see [Section 3.1.3.7](#)).
 - The originating internal client is notified.
 - The affected queue or client takes action based on the nature of the original request. An interrupt is issued to the respective PF.

3.1.3.7 Proprietary Error Reporting

The PCIe specification defines how to report errors to system software. There are, however, error events that the software device driver should be aware of or that the software device driver is in better position to handle and recover from. This section describes the mechanism to report PCIe related errors to software device drivers.

Several CSRs are dedicated to this functionality, with a separate bit allocated per error type (see [Table 3.5](#)):

- The PCIe Errors Reported register (*PCI_PCIERR - RO*) indicates which errors are reported using this mechanism. It is shared by all PFs and loaded from the NVM and all non-reserved errors are enabled.
- The PCIe Interrupt Cause register (*PCI_ICAUSE - RW1C*) indicates pending errors for errors set in the PCIe Errors Reported register. It is dedicated per PF.
- The PCIe Interrupt Enable register (*PCI_IENA - RW*) determines if an interrupt should be issued to the respective PCI function on an error event. It is dedicated per PF.

Reporting an error to the PF driver involves the following steps:

- The Intel® Xeon® Processor D-1500 Product Family LAN controller checks if the respective bit is set in the PCIe Errors Reported register. If cleared, done. Else, continue.
- The respective bit is set in the PCIe Interrupt Cause register.
- If the respective bit is set in the PCIe Interrupt Enable register, an interrupt is issued to the PCI function. The *PCI_EXCEPTION* cause is used (see the *EICR* register - [Section 8.2.2.8.1](#)).

**Table 3.5. PCIe Errors Reported to Device Software**

Error Event	Index	Description and Comments	Function Association
Exceeded upper address limit (read requests)	00	See Section 3.1.3.6	Sent to PF
Exceeded upper address limit (write requests)	01	See Section 3.1.3.6	Sent to PF
Reserved	02	Reserved entries	N/A
Poisoned TLP received	03	See Section 3.1.3.4	Sent to PF
Reserved	04-05	Reserved entries	N/A
Unsupported Request - Request Type	07	Request causes an Unsupported Request due to receipt of TLP with unsupported Request Type	Sent to PF
Unsupported Request - Vendor Message	08	Request causes an Unsupported Request due to receipt of an Unsupported Vendor Defined Type 0 Message	Sent to PF unless r[2:0] = <i>Broadcast from Root Complex</i> , in which case sent to all PFs
Reserved	09	Reserved entries	N/A
Unsupported Request - Function Number	10	Request causes an Unsupported Request due to receipt of a not-supported Function Number	Sent to all PFs
Unsupported Request - D3hot	12 - 11	Reserved entries	N/A
Completer abort - target size	13	Received Target Access with illegal data size per Section 3.1.2.1.1 (CA)	Sent to PF
Reserved	14 - 31	Reserved entries	N/A

3.1.3.8 Reset Warn Message

- Source ID: should match the hard strap value of strap_gbe_rstwarn_isspid
- Destination ID: should match the hard strap value of strap_gbe_sb_pid
- Opcode: should match the hard strap value of strap_gbe_rstwarn_eopcode (0x70)
- For Power down (byte 4): is set according to power down state S5,S4 or S3
- For cold/warm reset (byte 4): is set accordingly
- BYTE 5 -> 0xFF
- BYTE 6 -> 0xFF
- BYTE 7 -> 0x0



Table 3.6. Reset Warn Message Format

MSGDAT_p_RSWARN (Without SAI extended header) message details

DW#	Byte#	Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0
0	0x00	Destination ID							
	0x01	Source ID							
	0x02	Opcode 0x70 = MSGDAT_p_RSWARN							
	0x03	EH = 0	Reserved					Tag	
1	0x04	Target S-State / reset flow information – 1-hot encoding. IP may use this field to modify behavior – eg. skip shut-down interaction with external devices during S5 and cold reset flows							
		Cold reset	Warm Reset	S5	S4	S3	S2 (Rsvd)	S1	S0
	Permitted values 0x01, 0x02, 0x08, 0x10, 0x20, 0x40, 0x80. 00h = Notification that PCI Config space is hidden by PSF (func. disable) 0xF8=Undeclared flow (may be any of S3/4/5/Warm/cold reset). All other values reserved								
	0x05	Suspend I/O Well	Core I/O Well	Suspend logic Well	Core logic Well				
		Capabilities during S-state / flow 00b – Rail will be off 01b – Reserved 10b – Rail will be on with reduced capabilities (voltage/power) 11b – Rail will be on as normal							
	0x06	Suspend well clocks	Core clocks	Platform aux rail	Platform main rail				
Capabilities during S-state / flow 00b – Clocks will be off 01b – Reserved 10b – Clock will be on with at reduced frequency 11b – Clock will be on at normal frequency				Capabilities during S-state / flow 00b – Rail will be off 01b – Reserved 10b – Rail will be on with reduced capabilities (voltage/power) 11b – Rail will be on as normal					
0x07	Reserved					System wake latency indication. Indicates how long it will take to service (OS S0) a wake event caused by a downstream device 0h – Reserved 1h – 1-10uS 2h – 10 - 100uS 3h – 100uS – 1mS 4h – 1 - 10mS 5h – 10 – 100mS 6h – 100mS – 1S 7h - >1s			

3.1.3.9 Reset ACK Message

- Source ID: will be assigned based on the hard strap value of strap_gbe_sb_pid
- Destination ID: will be assigned based on the hard strap value of strap_gbe_rstdone_idpid
- Opcode: will be assigned based on the hard strap value of strap_gbe_rstdone_eopcode (0x71) BYTES 4 thru 7 are ZERO



Table 3.7. Reset Ack Message Format

MSGDAT_p_RSACK (Without SAI extended header) message details

DW #	Byte #	Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0
0	0x00	Destination ID							
	0x01	Source ID							
	0x02	Opcode 0x71 = MSGDAT_p_RSACK							
	0x03	EH = 0	Reserved					Tag	
1	0x04	PME_EN (functions 7:0) state storage							
	0x05	PME_STATUS (functions 7:0) state storage							
	0x06	Reserved							
	0x07	General purpose storage							

3.2 Management Interfaces

The Intel® Xeon® Processor D-1500 Product Family LAN controller contains three possible interfaces to an external BMC.

- SMBus
- NC-SI (over RMII)
- MCTP (over SMBus)

3.2.1 SMBus

SMBus is an optional interface for pass-through and/or configuration traffic between an external BMC and the Intel® Xeon® Processor D-1500 Product Family LAN controller. The SMBus channel behavior and the commands used to configure or read status from the Intel® Xeon® Processor D-1500 Product Family LAN controller are described in [Section 10.5](#).

The Intel® Xeon® Processor D-1500 Product Family LAN controller also enables reporting and controlling the device using the MCTP protocol over SMBus. The MCTP interface is used by the BMC to control the NIC and for pass-through traffic. All network ports are mapped to a single MCTP endpoint on SMBus. For additional information, refer to [Section 10.5](#).

3.2.1.1 Channel Behavior

The SMBus specification defines a maximum frequency of 100 KHz. However, when acting as a slave, the Intel® Xeon® Processor D-1500 Product Family LAN controller can receive transaction with a clock running at up to 1 MHz. When acting as a master, it can toggle the clock at 100 KHz, 400 KHz or 1 MHz. The speed used is set by the *SMBus Connection Speed* field in the SMBus Notification Time out and Flags NVM word.

3.3 Sideband Interface (NC-SI)

The NC-SI interface in the Intel® Xeon® Processor D-1500 Product Family LAN controller is a connection to an external MC. It operates as a single interface with an external BMC, where all traffic between the Intel® Xeon® Processor D-1500 Product Family LAN controller and the BMC flows through the interface.

The Intel® Xeon® Processor D-1500 Product Family LAN controller NC-SI interface meets the NC-SI version 1.0.0 specification as a PHY-side device.



3.3.1 Electrical Characteristics

The Intel® Xeon® Processor D-1500 Product Family LAN controller complies with the electrical characteristics defined in the NC-SI specification.

3.3.2 NC-SI Transactions

The NC-SI link supports both pass-through traffic between the BMC and the Intel® Xeon® Processor D-1500 Product Family LAN controller LAN functions, as well as configuration traffic between the BMC and the Intel® Xeon® Processor D-1500 Product Family LAN controller internal units as defined in the NC-SI protocol. Refer to [Section 10.6.1](#) for information.

3.3.3 MCTP (over SMBus)

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports MCTP protocol for management. MCTP runs over SMBus. The Intel® Xeon® Processor D-1500 Product Family LAN controller implements NC-SI over MCTP protocol for command and pass-through traffic. See [Section 10.7](#) for details.

3.4 Non-Volatile Memory (NVM)


3.4.1 General Overview

The Intel® Xeon® Processor D-1500 Product Family LAN controller uses a Flash device to store product configuration information. The Flash is divided into a few general regions:

- **Hardware Accessed** — Loaded by the the Intel® Xeon® Processor D-1500 Product Family LAN controller hardware after power-up, PCI reset de-assertion, D3 to D0 transition, or software reset. Different hardware sections in the Flash are loaded at different events. For more details on power-up and reset sequences, see [Section 4.1](#) and [Section 4.2](#).
- **Firmware Area** — Includes firmware code and structures used by the firmware for management configuration in its different modes.
- **Software Accessed** — This region is used by software entities such as LAN drivers, option ROM software and tools, PCIe bus drivers, VPD software, etc.

The basic requirements from the NVM access in The Intel® Xeon® Processor D-1500 Product Family LAN controller include the following:

1. Guarantee that only Intel provided firmware code is executed by the internal micro-controllers.
2. Protect NVM update flow from power failure before completion. This implies the image-update procedure of modules uses a double-bank policy. The requirement is to fall back to a valid image after a power failure.
3. Meet reasonable boot time. The following specifications must be met:
 - a. Access to all CSRs via the integrated I/O interface must be made possible within two seconds of PCIe reset in devices with a PCIe upstream port. This means that auto-load from the NVM of all the CSRs must complete within this time frame.
 - b. Device must be available for NC-SI connectivity within two seconds.
 - c. Device must be available for SMBus connectivity within 0.5 seconds.
4. Guarantee that the NVM settings, which are critical for device accessibility from the host are provided only by Intel, but can still be replaced in the field with other Intel provided settings if necessary.



3.4.1.1 NVM Protection

To meet requirements 1. and 4. previously described, the contents of several NVM modules must be protected via authentication.

The NVM protection method implemented in the Intel® Xeon® Processor D-1500 Product Family LAN controller relies on an authenticate on update concept. It means that protected modules are not authenticated after initialization, but prior to committing a module update operation only. NVM protection is guaranteed by an inductive authentication chain, that assumes an initial secured NVM image, and requires that any NVM update must be secure as well. This method mandates the following limitations and restricting working assumptions:

1. An initial good image is loaded into the Flash at the manufacturing site, which is assumed to be safe.
 - a. It assumes customers (OEM and end-user) know the source of the installed components, the supply chain producing these components is not compromised during manufacturing, and that the NIC/LOM is physically protected from modification after deployment.
 - b. The possibility exists that unauthorized firmware might be loaded into the NVM via physical modification post manufacturing, as well as through supply chain vulnerabilities. However, firmware updates via programmatic (software) methods are enhanced to require authentication prior to updating NVM settings. Furthermore, host software can independently detect whether the firmware image has an invalid digital signature.
2. In a normal operating mode, NVM write accesses are controlled by the device (firmware) and cannot be performed by the host. Memory mapped NVM access remains available for NVM read accesses only. For simplicity and flexibility reasons, NVM write accesses from the host can be initiated via Software Host Interface commands (Section 10.8.3.3), VPD write interface, or via a BMC command, which are all handled by firmware. All other direct access write modes are blocked by hardware when the NVM is protected.
3. All the supported Flash parts share the same set of opcodes as described in Section 3.4.1.2. A blank Flash programming mode is provided (besides the normal programming mode previously mentioned in item 2), where the Flash can be programmed directly without firmware involvement via the *FLA* bit banging (only in direct accessible Flashes) or Flash BAR interfaces. This mode is indicated by *FLA.LOCKED* = 0b.

3.4.1.2 Flash Device Requirements

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports Flash devices with a sector erase size of 4 KB and an address width of 24 bit (up to 4 MB). Note that many Flash vendors are using the term sector differently. This document uses the term Flash sector for a logic section of 4 KB.

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports Flash devices that are either write-protected by default after power-up or not. The Intel® Xeon® Processor D-1500 Product Family LAN controller is responsible to remove the protection by sending the write-protection removal OpCode to the Flash after power up.

The following OpCodes are supported by the Intel® Xeon® Processor D-1500 Product Family LAN controller as they are common to all supported Flash devices:

1. Write Enable (0x06)
2. Read Status Register (0x05) - used by hardware internally.
3. Write Status Register (0x01). The written data is 0x00 to cancel the Flash default protection. Used by hardware internally.
4. Read Data (0x03). Burst read is supported.
5. Byte Program (0x02). To program a data byte.
6. 4 KB Sector-Erase (0x20)



7. Chip-Erase (0xC7)
8. Read JEDEC ID (0x9F)

3.4.1.2.1 Flash Identification

The Flash connected to the Intel® Xeon® Processor D-1500 Product Family LAN controller can be identified by its JEDEC ID that can be read using the Flash Info Host Interface command (Section 10.8.3.3.9). This identification is available only if a valid Flash is installed, as the extraction of the ID from the flash is performed by firmware using the RDID (Opcode 0x9F) command. If the Flash is empty or with an invalid signature, then software can read the JEDEC ID by applying an RDID command (opcode 0x9F) via the bit-bang interface.

3.4.2 Shadow RAM

The first eight 4 KB sectors of the Intel® Xeon® Processor D-1500 Product Family LAN controller's Flash are allocated to create two 16 KB sections (section 0 and section 1), for the configuration content. At least one of these two sections must be valid at any given time, otherwise the Intel® Xeon® Processor D-1500 Product Family LAN controller is set based on its hardware default. Following a Power On Reset (POR), the Intel® Xeon® Processor D-1500 Product Family LAN controller copies the valid section of the Flash device into an internal shadow RAM. Any further read accesses by software or firmware to the lower 16 KB addresses of the NVM (not through Flash BAR) are directed to the internal shadow RAM. Modifications made to the shadow RAM contents are copied by the Intel® Xeon® Processor D-1500 Product Family LAN controller to the other 16 KB section of the NVM, circularly flipping the valid section between sections 0 and 1 in the NVM. This mechanism provides the following advantages:

1. A seamless backward-compatible interface for software/firmware to access the first 16 KB of the NVM as if an external EEPROM device was connected. This interface is referred as EEPROM-mode access to the Flash.

Note: Write access to the physical Flash is only available in blank Flash/non-secure modes.

2. A way to protect the image-update procedure from power down events by establishing a double-image policy. See Section 3.4.9.1 for a description of the double-image policy. It relies on having pointers to NVM modules stored in the NVM section mirrored in the internal shadow RAM.
3. A way to ensure hardware auto-load during a PCIe reset event can be completed within PCIe timing constraints (100 ms), even if the Flash is busy performing an erase operation initiated prior to that reset event.

Figure 3.2 shows the shadow RAM mapping and interface.

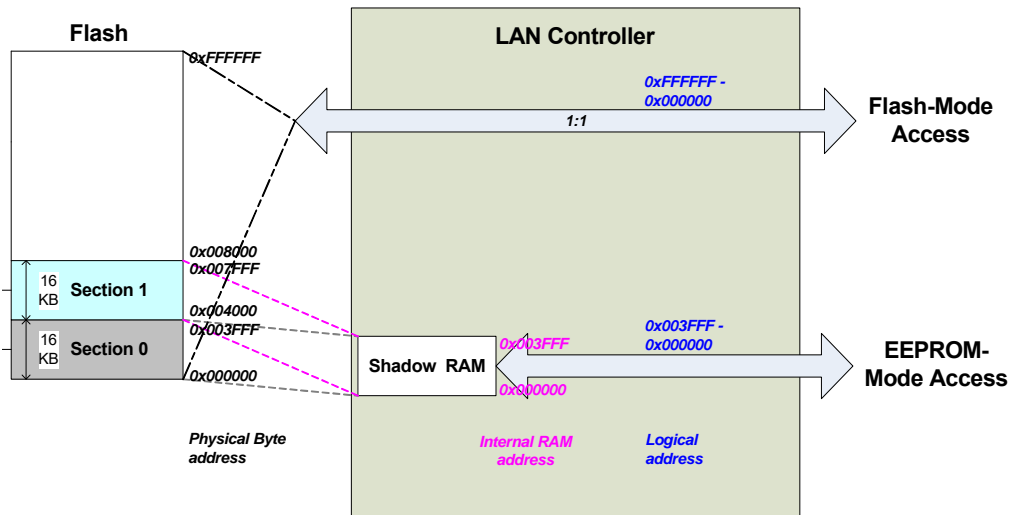


Figure 3.2. NVM Shadow RAM

3.4.2.1 Shadow RAM Update Flow

1. Following a write access by the software device driver to update the shadow RAM, the data should be updated in the Flash as well. The Intel® Xeon® Processor D-1500 Product Family LAN controller updates the Flash from the shadow RAM when software explicitly requests an update using the Shadow RAM Dump Host Interface command (Section 10.8.3.3.8). To reduce Flash update operations, software is expected to request a dump only once its last shadow RAM write access completes. The Intel® Xeon® Processor D-1500 Product Family LAN controller then copies the contents of the shadow RAM to the non-valid configuration section and makes it the valid one.
2. The software device driver should wait for the command completion to make sure the flow succeeded. Software should be aware that programming the Flash might require a long time since the Flash update sequence is handled by firmware. The sector erase command by itself can last hundreds of milliseconds. The regular timeout of the host interface commands should be sufficient to cover this command also.

3.4.3 NVM Clients and Interfaces

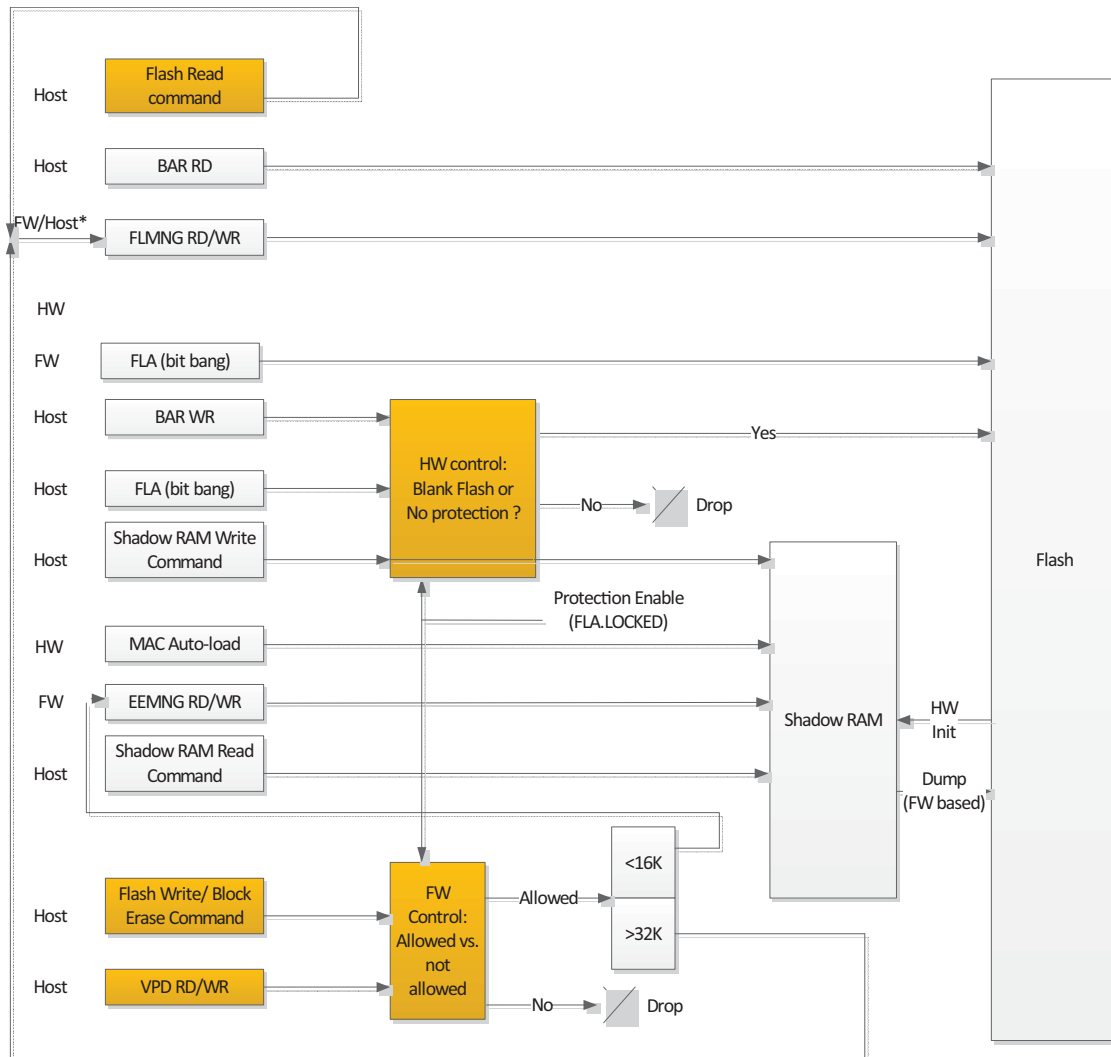
There are different software clients that can access the NVM: driver, tools, BIOS, VPD, etc. Table 3.8 lists the different NVM access methods.



Table 3.8. Clients and Access Types to the NVM

Client	NVM Access Method (Data Width)	NVM Access Target	Logical Byte Address Range	NVM Access Interface	Protection and Enforcement
Host Software	Memory BAR (parallel 32-bit)	Flash	0x000000 - 0xFFFFFFFF	Memory mapped 32-bit read/write via BARs.	<ul style="list-style-type: none"> Write access is limited to a single byte, and is allowed only if hardware protection is disabled. (FLA.LOCKED = 0b).
BIOS	Expansion ROM BAR (parallel 32-bit)	Flash	Expansion ROM Base Address Register	Memory mapped 32-bit read/write via BARs.	<ul style="list-style-type: none"> Write access is limited to a single byte, and is allowed only if hardware protection is disabled. (FLA.LOCKED = 0b).
Host Software	FLA bit banging (serial 1-bit)	Flash	0x000000 - 0xFFFFFFFF	Software accesses to Flash by toggling the SPI pins.	FLA interface access is available only if hardware protection is disabled (FLA.LOCKED = 0b).
Host Software	Host interface Shadow RAM Read/Write command	Shadow RAM	0x000000 - 0x003FFF	Access to shadow RAM through the Shadow RAM Read/Write command	Requires a valid firmware image. Write protection is enforced by firmware.
Host Software	Host interface Flash read command	Flash	0x000000 - 0xFFFFFFFF	Software read from Flash via Flash Read command.	Requires a valid firmware image.
Host Software	Host interface Flash write command/Flash sector erase command	Flash	0x008000 - 0xFFFFFFFF	Software writes and sector erase.	Requires a valid firmware image. Writes and erases to protected areas are dropped when firmware protection is enabled.
Host Software	VPD access (parallel 32-bit)	Shadow RAM	0x000000 - 0x0003FF	VPD Address and Data registers.	Write accesses are enabled to the R/W area of the VPD If the VPD structure is not valid, the entire 1024 bytes area becomes RO to VPD accesses.
Firmware	EEMNG Parallel (16-bits)	Shadow RAM	0x000000 - 0x003FFF	EEMNGCTL, EEMNGDATA accesses to the shadow RAM	
Firmware	FLMNG Parallel (Read: 32-bits Write: 8-bits)	Flash	0x000000 - 0xFFFFFFFF	FLMNGCTL and FLMNGDATA registers accesses to the FLASH	
MAC Hardware	Serial pins	Flash	0x000000 - 0x007FFF	Initial read access to sections 0 and 1 is initiated by hardware for the auto-load process. Consecutive reads are performed to the shadow RAM logic.	No write access (by design).

Figure 3.3. shows the different access modes in and the intervention of firmware in the process.



* Only if protection is disabled

Figure 3.3. Flash access modes

3.4.4 Memory Mapped Host Interface

The Flash is accessed by the Intel® Xeon® Processor D-1500 Product Family LAN controller each time the host CPU performs a read operation to a memory location mapped to the Flash address space, or upon boot via accesses to the space indicated by the Expansion ROM Base Address register. Accesses to the Flash are based on a direct decode of CPU accesses to a memory window defined in either:

- Memory CSR + Flash Base Address register (PCIe Control register at offset 0x10).
- The Expansion ROM Base Address register (PCIe Control register at offset 0x30).
- The Intel® Xeon® Processor D-1500 Product Family LAN controller is responsible to map accesses via the expansion ROM BAR to the physical NVM. The offset in the NVM of the expansion ROM module is defined by the PCIe expansion/option ROM pointer (shadow RAM word address 0x05).



This pointer is loaded by the Intel® Xeon® Processor D-1500 Product Family LAN controller from the NVM before enabling any access to the expansion ROM memory space.

The Intel® Xeon® Processor D-1500 Product Family LAN controller controls accesses to the Flash when it decodes a valid access. Attempts to out of range read access, the PCIe expansion/option ROM module (according to NVM size field in NVM Control Word 1) would return a value of 0xDEADBEEF. Attempts to memory-mapped write accesses to the Flash when protection is enabled or via expansion ROM BAR are ignored.

3.4.5 Flash Access Contention

Flash accesses initiated through different LAN functions might occur concurrently. The Intel® Xeon® Processor D-1500 Product Family LAN controller does not synchronize between entities accessing the Flash, so a contention caused from one entity reading and another modifying the same location is possible.

To avoid such contention between software LANs or between software and firmware accesses, these entities are required to make use of the semaphore registers. Refer to [Section 10.8.4](#). Any read or write access to the NVM made by software/firmware must be preceded by acquiring ownership over the NVM. This is also useful to avoid the time out of a PCIe transaction made to a memory mapped Flash address while the Flash is busy performing a sector erase operation.

However, two software entities cannot use this semaphore mechanism: BIOS access through expansion ROM and VPD software.

- Since VPD software accesses only the VPD module, which is located in the configuration section of the NVM, VPD accesses are always performed against the shadow RAM. Firmware must take NVM ownership before dumping the VPD changes to the Flash. The Shadow RAM dump sequence is described in [Section 3.4.2.1](#).
- No contention can occur between the BIOS access through expansion ROM and other software entities (including VPD) as it accesses the NVM while the operating system is down.
- Contentions between BIOS and firmware can however happen if a system reboot occurs while the MC is accessing the NVM.
 - If a system reboot is caused by a user pressing the standby button, it is required to route the wake-up signal from the standby button to the MC and not to the chipset. The MC issues a system reboot signal to the chipset only after the NVM write access completes. Firmware is responsible to respond with a busy error code to MC NC-SI commands while other NVM writes are in progress.
 - If a system reboot is issued by a local user on the host, there is no technical way to prevent NVM access contentions between the BIOS and the MC.

Caution: It is the user's responsibility when remotely accessing the NVM via the MC, to make sure another user is not currently initiating a local host reboot.

Note: The MAC auto-load from the Flash device itself occurs after power-up only, before host or firmware can attempt to access the Flash. The host must wait until PCIe reset is de-asserted (after ~1 second, which is enough time for the MAC auto-load to complete), and firmware starts its auto-load after the *EEC.MNG_READY* bit is asserted by hardware.

Other MAC auto-load events are performed from the internal shadow RAM and do not compete with memory mapped accesses to the Flash device. During such MAC auto-loads, accesses from other clients via EEPROM-mode registers are delayed until the auto-load process completes.

Software and firmware should avoid holding Flash ownership (via the dedicated semaphore bit) for more than 500 ms.



3.4.5.1 Flash Deadlock Avoidance

The Flash is a shared resource between the following clients:

1. Hardware MAC auto-load of shadow RAM (at power up).
2. LAN port 0 and LAN port 1 software accesses.
3. Manageability/firmware accesses.
4. Software tools.

All clients can access the Flash using parallel access. Hardware/firmware implements the actual access to the Flash. Hardware arbitrates between the different clients and schedules these accesses, avoiding starvation of any client.

However, in the Intel® Xeon® Processor D-1500 Product Family LAN controller, the software and firmware clients can access the Flash using bit banging. In this case, there is a request/grant mechanism that locks the Flash to the exclusive use of one client. If one client is stuck without releasing the lock, the other clients can no longer access the Flash. To avoid this deadlock, the Intel® Xeon® Processor D-1500 Product Family LAN controller implements a time-out mechanism, which revokes the grant from a client that doesn't toggle the Flash bit-bang interface (*FLA.FL_SCK* bit) for more than two seconds. If any client fails to release the Flash interface, hardware clears its grant, enabling the other clients to use the interface.

The deadlock timeout mechanism is enabled by the *Deadlock Timeout Enable* bit in the Control Word 2 in the Flash.

3.4.6 Signature Field

The only way the Intel® Xeon® Processor D-1500 Product Family LAN controller can tell if a Flash is present is by trying to read the Flash. The Intel® Xeon® Processor D-1500 Product Family LAN controller first reads the Control word at word address 0x0 and at word address 0x2000. It then checks the signature value at bits 7 and 6 in both addresses.

If bit 7 is 0b and bit 6 is 1b in (at least) one of the two addresses, it considers the Flash to be present and valid. It then reads the additional Flash words from that section and programs its internal registers based on the values read. Otherwise, it ignores the values read from that location and does not read additional words.

If the signature bits are valid at both addresses the Intel® Xeon® Processor D-1500 Product Family LAN controller assumes that the first section is the valid one.

3.4.7 VPD Support

The Flash image can contain an area for VPD. This area is managed by the OEM vendor and does not influence the behavior of hardware. Word 0x2F of the Flash image contains a pointer to the VPD area in the Flash. A value of 0xFFFF means VPD is not supported and the *PCI_CAPCTRL.VPD_EN* bit should be cleared in the PCI NVM section, to prevent the VPD capability from appearing in the configuration space.

The maximal area size is 1024 bytes but can be smaller. The VPD block is built from a list of resources. A resource can be either large or small. The structure of these resources are listed in the following tables.

**Table 3.9. Small Resource Structure**

Offset	0	1 – n
Content	Tag = 0xxx,xyyyb (Type = Small(0), Item Name = xxxx, length = yy bytes)	Data

Table 3.10. Large Resource Structure

Offset	0	1 – 2	3 – n
Content	Tag = 1xxx,xxxxb (Type = Large(1), Item Name = xxxxxxxx)	Length	Data

The Intel® Xeon® Processor D-1500 Product Family LAN controller parses the VPD structure during the auto-load process following PCIe reset in order to detect the read only and read/write area boundaries. The Intel® Xeon® Processor D-1500 Product Family LAN controller assumes the following VPD fields with the limitations listed:

Table 3.11. VPD Structure

Tag	Length (bytes)	Data	Resource Description
0x82	Length of identifier string	Identifier	Identifier string.
0x90	Length of RO area	RO data	VPD-R list containing one or more VPD keywords.
0x91	Length of RW area	RW data	VPD-W list containing one or more VPD keywords. This part is optional.
0x78	N/A	N/A	End tag.

VPD structure limitations:

- The structure must start with a Tag = 0x82. If the Intel® Xeon® Processor D-1500 Product Family LAN controller does not detect a value of 0x82 in the first byte of the VPD area or the structure does not follow the description of [Table 3.11](#), it assumes the area is not programmed and the entire 1024 bytes area is read only.
- The RO area and RW area are both optional and can appear in any order. A single area is supported per tag type. Refer to Appendix I in the PCI 3.0 specification for details of the different tags.
- If a VPD-W tag is found, the area defined by its size is writable via the VPD structure.
- The structure must end with a Tag = 0x78. The tag must be word aligned.
- The VPD area can be accessed through the PCIe configuration space VPD capability structure listed in [Table 3.11](#). Write accesses to a read only area or any access to an offset outside of the VPD area via this structure are ignored.
- VPD area must be mapped in the first 16 KB section of the Flash mapped to the shadow RAM.
- VPD software does not check the semaphores before attempting to access the Flash via dedicated VPD registers. Even if the Flash is owned by another entity, VPD software read access to the VPD area in the Flash might complete immediately since it is first performed against the shadow RAM. However, VPD software write access might not complete immediately since the VPD changes are committed to the Flash device at the Intel® Xeon® Processor D-1500 Product Family LAN controller's initiative, once the other entity releases Flash ownership, which may take up to several seconds.



3.4.7.1 VPD Access Flows

3.4.7.1.1 First VPD Area Programming

The VPD capability is exposed in the PCIe configuration space only if the *PCI_CAPCTRL.VPD_EN* bit is set, regardless to any other sanity check that is performed on the VPD area contents.

The VPD content and pointer can be written on a blank Flash without any limitation, similar any other NVM module when in the blank Flash programming mode. After protection is enabled, if *VPD Write Enable* bit in NVM control word 1 is cleared, only the RW area of the VPD is writable and only via the VPD interface.

3.4.7.1.2 VPD Area Update Flow

1. The host initiates a VPD write by programming the offset and data fields of the VPD capability register set and then setting the capabilities *Flag* bit (bit 15 in VPD Address register - 0xE2).
2. Firmware checks if the VPD write is allowed making sure that the write offset falls within the VPD-RW area.
 - a. If writing is not allowed, firmware clears the VPD flag in the configuration space to notify the host VPD software that the transaction completed. It then exits the flow.
3. Firmware indicates VPD access completion by clearing the VPD flag in the configuration space.

Note: In case the Flash is busy by a previous sector erase operation, or if NVM ownership is held by software, the flow should not be restarted, for successive VPD write accesses, before the last step is completed.

3.4.8 NVM Read, Write, and Erase Sequences

Refer to [Section 3.4.9.1](#) for the flow required to update secure NVM modules.

Any software or firmware flow described in this section must be preceded by taking NVM ownership via semaphores as described in [Section 10.8.4](#).

3.4.8.1 Flash Erase Flow by Firmware

If the software device driver triggered the erase flow, firmware must not take the semaphore.

3.4.8.1.1 Device Erase Flow

Note: This flow is not used by firmware.

1. Poll the *FLMNGCTL.FLBUSY* and *FLMNGCTL.DATA_VALID* bits until they are cleared, and the *FLMNGCTL.GLDONE* bit until is set.
2. Set *FLMNGCTL.CMD* field to 11b and *FLMNGCTL.CMDV* bit to 1b.
3. Wait until *FLMNGCTL.FLBUSY* bit is cleared before releasing the NVM semaphore.

3.4.8.1.2 Sector Erase Flow

1. Poll the *FLMNGCTL.FLBUSY* and *FLMNGCTL.DATA_VALID* bits until they are cleared, and the *FLMNGCTL.GLDONE* bit until it is set.
2. Set *FLMNGCTL.CMD* field to 10b, *FLMNGCTL.CMDV* bit to 1b, and set the *FLMNGCTL.ADDR* field to any address which belongs to the Flash 4KB sector to be erased.
3. Wait until *FLMNGCTL.FLBUSY* bit is cleared before releasing the NVM semaphore.



3.4.8.2 Flash Erase Flow by Software

Note: The actual erase command is issued by the *FLMNG* register set. Software uses the Flash Block Erase host Interface command (Section 10.8.3.3.7) as a way to communicate the transaction to firmware, which uses the *FLMNG* registers for the actual operation (initiated by software or firmware).

3.4.8.2.1 Device Erase Flow

1. This flow is not implemented in Firmware. Software should use the *FLA.FL_DER* to erase the flash (Section 3.4.8.4). Poll the *FLSWCTL.DONE* bit until it is set
2. Set *FLSWCTL.CMD* fields to 0011b
3. Hardware interrupts Firmware to execute the erase flow as defined in Section 3.4.8.1
4. Wait until *FLSWCTL.DONE* bit is read as 1b and the *FLSWCTL.CMDV* bit is cleared before releasing the NVM semaphore

Note: Firmware sets the *FLSWCTL.DONE* bit and the *FLSWCTL.COMMAND_FAIL* bit without executing the operation if software attempts this command while protection is enabled.

3.4.8.2.2 Sector Erase Flow

1. Send a Flash Block Erase host Interface command (Section 10.8.3.3.7) with the aligned address of the block to erase
2. Hardware interrupts Firmware to execute the erase flow as defined in Section 3.4.8.1 if the sector does not contain Read Only data.
3. Wait for the command to complete before releasing the NVM semaphore

3.4.8.3 Software or Firmware Flow to the Bit Banging Interface

To directly access the Flash when Flash is blank or not protected, software/firmware should follow these steps:

1. Write a 1b to the *Flash Request* bit (*FLA.FL_REQ*).
2. Poll the *Flash Grant* bit (*FLA.FL_GNT*) until it becomes 1b. It remains 0b as long as there are other accesses to the Flash.
3. Write or read the Flash using the direct access to the 4-wire interface as defined in the *FLA* register. The exact protocol used depends on the Flash placed on the board.
4. Write a 0b to the *Flash Request* bit (*FLA.FL_REQ*).
5. Following a write or erase instruction, software/firmware should clear the *Request* bit only after it has checked that the cycles were completed by the NVM. This can be checked by reading the *BUSY* bit in the Flash device *STATUS* register. Refer to Flash data sheet for the OpCode to be used for reading the Status register.

Note: The bit banging interface is blocked during normal operation (protection enabled). Software should use the host interface commands. Firmware can use this interface all the time.

3.4.8.4 Erase Flow Using the *FLA* Register

To directly erase a sector in the Flash when Flash is blank or not protected, software/firmware should follow these steps:

1. Take ownership of NVM semaphore.
2. Set the *Flash Address* (*FLA.FL_ADDR*) to the index of the 4 KB sector to erase and *Sector Erase* bit (*FLA.FL_SER*) bit.



3. Read the FLA register until *Flash Busy* bit (*FLA.FL_BUSY*) is cleared.
4. Release ownership of NVM semaphore.

To directly erase the entire Flash when Flash is blank or not protected, software /firmware should follow these steps:

1. Take ownership of NVM semaphore.
2. Set *Device Erase* bit (*FLA.FL_DER*) bit.
3. Read the FLA register until *Flash Busy* bit (*FLA.FL_BUSY*) is cleared.
4. Release ownership of NVM semaphore.

3.4.8.5 Firmware Word Access Flow to the EEPROM-Mode Interface

3.4.8.5.1 Read Interface

Firmware initiates a read cycle from the shadow RAM as follow:

1. Poll the *EEMNGCTL.DONE* bit until it is set. This step is only needed if the flow is executed following a reset event.
2. Writes the address to be read in the *EEMNGCTL* register and sets the *START* bit in the same register. Clears the *CFG_DONE* and *WRITE* bits.
3. Polls the *EEMNGCTL.DONE* bit until it is asserted.
4. Reads the *EEMNGDATA* register field

Hardware executes the following steps following a write to the *EEMNGCTL.START* bit:

1. Reads the data from the shadow RAM.
2. Puts the data in *RDDATA* field of the *EEMNGDATA* register.
3. Sets the *DONE* bit in the *EEMNGCTL* register.

3.4.8.5.2 Write Interface

Firmware initiates a write cycle to the shadow RAM as follows:

1. Poll the *DONE* bit in the *EEMNGCTL* register until it is set.
2. Writes the data in the *WRDATA* of the *EEMNGDATA* register
3. Writes the address in the *EEMNGCTL* register. Clears the *CFG_DONE* bit and sets the *WRITE* and *START* bits.

As a response, hardware executes the following steps:

1. Clears the *START* bit. If the command can not be executed, go to step 3.
2. Writes the data to the shadow RAM.
3. Sets the *DONE* bit in the *EEMNGCTL* register.

Note: In both cases (read/write), whenever the *DONE* bit is set by hardware, an interrupt is issued to firmware.

3.4.8.6 Software Word Access Flow Through the EEPROM-Mode Interface

Note: These flows are obsolete and are replaced by the flow described in [Section 3.4.8.7](#).

Software must take semaphore ownership before executing these flows.



Writes using this interface are only allowed on blank Flash programming mode or if protection is disabled.

3.4.8.6.1 Read Interface

Software initiates a read cycle to the NVM via the EEPROM-mode as follows:

1. Software writes the address to be read in the *EERD.ADDR* field and sets the *EERD.START* bit.
2. Software polls the *EERD.DONE* bit until it is asserted.
3. Software reads the *EERD.DATA* register field.

Hardware executes the following steps:

1. Clears the *EERD.START* bit.
2. Reads the data from the shadow RAM.
3. Puts the data in *EERD.DATA* field.
4. Sets the *EERD.DONE* bit.

Note: Any word read this way is not loaded into the Intel® Xeon® Processor D-1500 Product Family LAN controller's internal registers. This happens only at a hardware auto-load event. Write Interface

When Flash is blank or protection is disabled Software may initiate a write cycles to the shadow RAM via the *EEWR* register as follow:

Software initiates a write cycle to the NVM via the EEPROM-mode as follows:

1. Poll the *EEWR.DONE* bit until it is set
2. Write the data word and its address in the *EEWR* register and set the *EEWR.START* bit.

As a response, hardware executes the following steps:

1. Clears the *EEWR.START* bit.
2. The Intel® Xeon® Processor D-1500 Product Family LAN controller writes the data to the shadow RAM
3. The Intel® Xeon® Processor D-1500 Product Family LAN controller sets the *DONE* bit in the *EEWR* register

Note: EEPROM-mode writes are performed into the internal shadow RAM. Software can instruct copying of the internal shadow RAM content into the Flash device by setting the *EEC.FLUPD* bit.

Note: When there is no firmware, the shadow RAM dump is not available. In these scenarios, while programming a non-valid Flash, software should access the Flash directly and not use the shadow RAM mechanism ([Section 3.4.2.1](#)).

3.4.8.7 Software Access Flow to Shadow RAM

3.4.8.7.1 Read Interface

Software can read from the shadow RAM using the following flow:

1. Send a Shadow RAM Read Host Interface command ([Section 10.8.3.3.2](#)) with the address and length to read.
2. Wait for the command to complete



3. Read the data from the response buffer.

3.4.8.7.2 Write Interface

Software can write to the shadow RAM using the following flow:

1. Send a Shadow RAM Write host Interface command (Section 10.8.3.3.4) with the address, length and data to write.
2. Wait for the command to complete

3.4.8.8 Flash Access Flow via the Memory Mapped Interface

3.4.8.8.1 Read Access

Software can always use the Flash BAR for read accesses. Example flow:

1. Acquire the NVM ownership semaphore.
2. Perform a memory read from the address range mapped to the NVM.
3. Hardware initiates a read transaction, returning a single Dword.
4. Steps 1-2 can be repeated multiple times, as needed.
5. Release the NVM ownership semaphore.

Note: Software should take semaphore ownership before executing the flow.

3.4.8.8.2 Write Access

When the Flash is blank or protection is disabled, software might initiate a write cycles via the Flash BAR as follows:

1. Take semaphore ownership before executing the flow.
2. Write the data byte to the Flash through the Flash BAR.
3. Poll the *FL_BUSY* flag in the *FLA* register until cleared.
4. Repeat steps 2. and 3. to write additional bytes.
5. Release NVM semaphore ownership.

As a response, hardware executes the following steps for each write access:

1. Set the *FL_BUSY* bit in the *FLA* register.
2. Initiate autonomous write enable instruction.
3. Initiate the program instruction right after the enable instruction.
4. Poll the Flash status until programming completes.
5. Clear the *FL_BUSY* bit in the *FLA* register.

Note: Software must erase the sector prior to programming it.

3.4.8.9 Firmware Flash Program Flow via the Flash-Mode Interface

Internally the following flow is performed for programming flash sectors:

1. Poll the *FLMNGCTL.FL_BUSY* and the *FLMNGCTL.DATA_VALID* bits until they are cleared, and the *FLMNGCTL.GLDONE* bit until it is set.
2. Write the byte of data to bits 7:0 of the *FLMNGDATA* register.



3. Set the *ADDR* field with the byte resolution address in the *FLMNGCTL* register, set the *CMD* field to 01b and set the *CMDV* bit to 1b.
4. Hardware starts accessing the Flash and begins writing the 8 low order data bits from the *FLMNGDATA* register. Once hardware completes writing the 8-bits of data to the Flash, the *FLMNGCTL.GLDONE* register is set.
5. Steps 2. to 4. can be repeated several times.
6. Wait until *FLMNGCTL.FLBUSY* bit is cleared before releasing the NVM semaphore.

3.4.8.10 Firmware Flash Auto-load Flow via the Flash-Mode Interface

The Intel® Xeon® Processor D-1500 Product Family LAN controller provides the Firmware with an engine for reading the Flash in a burst mode, which can be useful to auto-load the entire Firmware image in a more efficient manner:

1. Poll the *FLMNGCTL.FLBUSY* and the *FLMNGCTL.DATA_VALID* bits until they are cleared, and the *FLMNGCTL.GLDONE* bit until it is set.
2. Set the *FLMNGCNT.RDCNT* field with the number of DWords to be read from Flash in a burst mode.
3. Set the *FLMNGCTL.ADDR* field with the byte address of the first DWord to be read, set the *CMD* field to 00b and set the *CMDV* bit to 1b. The *FLMNGCTL.GLDONE* bit is cleared by hardware to indicate a burst read has started.
4. Hardware starts accessing the Flash and clears the *FLMNGCTL.DATA_VALID* bit until it writes the read DWord into the *FLMNGDATA* register. Firmware polls the *FLMNGCTL.DATA_VALID* bit until it is set.
5. Firmware reads the DWord from *FLMNGDATA* register, which is used by hardware to trigger a clear of the *FLMNGCTL.DATA_VALID* bit again.
6. Hardware increments *FLMNGCTL.ADDR* field by four (DWord granularity).
7. Steps 4. to 6. are repeated until the number of DWords programmed in *FLMNGCNT.RDCNT* has been read.
8. Hardware set the *FLMNGCTL.GLDONE* bit and clears the *FLMNGCTL.CMDV* bit to indicate that all the Flash transactions related to the command issued at step 3. were completed.

Firmware has the possibility to abort a burst read as follow:

1. Check *FLMNGCTL.CMDV* bit. If it is cleared, then do nothing, else proceed to the next step.
2. Set the *FLMNGCNT.Abort* bit to 1b.
3. Poll the *FLMNGCNT.Abort* bit until it is cleared.

3.4.8.11 Software Flash/Shadow RAM Program Flow

Software can write to non-write protected areas of the Flash using the following flow:

1. Send a Flash Write Host Interface command (Section 10.8.3.3.3) with the address, length and data to write.
2. Wait for the command to complete.

3.4.8.12 Software Flash Read Flow via the Flash-mode Interface

Software can read from the Flash using the following flow:

1. Send a Flash Read host Interface command (Section 10.8.3.3.1) with the address and length of the data to read.
2. Wait for the command to complete.
3. .Read the data returned by the host IF command.



3.4.9 Extended NVM Flows

3.4.9.1 Flow for Updating Secured Modules

This section describes the flow to use to update the firmware image, option ROM (Section 6.2.1.6) and the internal PHY indirect list (Section 6.2.1.4).

In order to protect the Flash update procedure from power-down events, a double image policy is used for each of the updated modules with the exception of the internal PHY indirect list. The software flow to update a module is as follows:

1. Take ownership over the NVM via the semaphore bits. Refer to Section 10.8.4.
 - a. If *SW_FW_SYNC.NVM_Update_Started* bit is read as clear, then set this bit together with setting NVM semaphore bit. It is used to notify other entities that a long NVM update process that might take up to several minutes has started. During this time, other entities cannot perform a write access to the firmware module, but reading this module in between update write bursts is allowed using the Flash memory mapping. Legacy EEPROM modules are not concerned by this limitation.
 - b. Otherwise, release NVM semaphore ownership and restart the update process later on.
2. Read the pointer to the free provisioning area (NVM word 0x40). Check that the free provisioning area size read from NVM word 0x41 is greater or equal to the size of the new firmware/PXE image to be loaded in the NVM.
 - a. If not, release NVM semaphore ownership, clear the *SW_FW_SYNC.NVM_Update_Started* bit and exit the flow.
3. Initiate sector erase instructions to the entire free space provisioning segment.
 - a. In order to guaranty NVM semaphore ownership time does not exceed the one second timeout, it is recommended to perform at this step no more than four 4 KB sector erase operations at once in a burst, releasing semaphore ownership for 10 ms in between. This way, other entities can insert NVM read accesses in between burst without waiting for the entire update process completion, which might take minutes.
4. Write the new firmware/option ROM to the free provisioning area via Flash-mode access.
 - a. Same as (3.a), it is recommended to write at this step no more than four 4 KB sectors at once in a burst, releasing semaphore ownership for 10 ms in between.
5. Send a Flash module update host interface command (Section 10.8.3.3.5) with the module ID of the section to update. The encoding of the modules is:

Module	ID	Reference
Firmware code	0x1	
Reserved	0x2	
Reserved	0x3	
Internal PHY Indirect List ¹	0x4	Section 6.2.1.4
Reserved	0x4	
Reserved	0x5	
Reserved	0x6:0xFD	
Option ROM	0xFE	Section 6.2.1.6
Reserved (full shadow RAM)	0xFF	

1. The internal PHY indirect list must be the last module to be updated.



6. Release the NVM semaphore and clear the *SW_FW_SYNC.NVM_Update_Started* bit.
 - a. Software must avoid taking the NVM semaphore again until the firmware command has completed. Any attempt to write the NVM until then is not performed by the device.
7. Check that the size field in the CSS header corresponds to a 488 KB long module for the firmware code and the option ROM or 8 KB long for the internal PHY indirect list.
8. If the firmware code or option ROM needs updating, firmware swaps between the free provisioning area pointer (word 0x40) and either the firmware code module pointer or option ROM module pointer located at shadow RAM word addresses 0x3A/0x05, respectively. It uses EEPROM-mode access registers to perform the swap.

If the internal PHY indirect list needs updating, firmware copies the new internal PHY indirect list content directly into the respective module.

Firmware dumps the shadow RAM into the Flash via the flow described in [Section 3.4.2.1](#).

9. Software waits for the command to complete.
 - a. If the update process failed due to a security check failure or a Flash write fault, an Authentication Error (0x80) or Data Error (0x6), respectively is returned. Software must then exit the flow, prior to attempt another update.
10. If the updated module is the firmware image, then firmware should be instructed to reload the image using the *Apply Update* command.

3.4.9.2 Flow for Updating One of the RW Legacy EEPROM Modules

When updating one or several fields from a legacy EEPROM module there is a risk that a hardware auto-load event occurs in the middle of the operation (for example, due to a sudden PCIe reset), leading to the auto-load of an invalid or inconsistent content from the internal shadow RAM into the device registers or memory. Therefore unless the field(s) can be updated by a single EEPROM-mode access, the updating software must repeatedly use the following procedure for each legacy EEPROM module to be updated:

1. Take ownership over the NVM via semaphore bits. Refer to [Section 10.8.4](#).
2. Invalidate the pointer to the module to be modified by setting it to 0xFFFF using Shadow RAM Write command. This way, if a hardware auto-load of the module is attempted, the associated register defaults are loaded instead. Do not invalidate pointers to firmware modules, only to hardware autoloading modules.
3. Modify the contents of the module via Shadow RAM Write command.
4. Restore the pointers modified in step 2. via Shadow RAM Write command.
5. Compute and update the software checksum (word 0x3F) if the contents covered by the software checksum was modified.
6. Release the NVM semaphore.
7. Send a Shadow RAM Dump command ([Section 10.8.3.3.8](#)) to ask the device firmware to load the internal shadow RAM into the Flash.

Note: Depending on the modified RO items, a system reset is generally required for loading the modifications into the device.

3.4.10 NVM Authentication Procedure

The NVM update integrity feature ensures that only Intel-approved firmware code (or other protected NVM module) is able to be updated on the Intel® Xeon® Processor D-1500 Product Family LAN controller devices after manufacturing. This procedure is performed whenever attempting to update one of the protected modules.

Integrity validation of NVM updates is provided by a digital signature. The digital signature is a SHA256 Hash computed over the protected content (long by 256-bits), which is then encrypted by a 2048-bits RSA encryption using an Intel private key. This digital signature is stored in the manifest in the NVM module image. Also stored in the manifest is the corresponding RSA modulus (public key) and RSA exponent to be used to decrypt the digital signature.

To verify the authenticity of the digital signature, firmware must first verify that the RSA Modulus and RSA Exponent fields in the new firmware image loaded are identical to those in the old firmware image. If the RSA Modulus and Exponent fields are the same, firmware decrypts the digital signature using the 2048-bit RSA Modulus and Exponent fields stored in the manifest of the old firmware image to extract the expected SHA256 Hash of content (stored hash). Firmware then performs an independent SHA256 Hash over the protected content (computed hash). If the stored hash matches the computed hash, the digital signature is accepted, and the NVM update is applied.

NVM updates are validated prior to invalidating the old NVM configuration, such that the old NVM configuration is still usable if the update fails to validate. After the new NVM is successfully verified, the updated image is committed to device Flash.

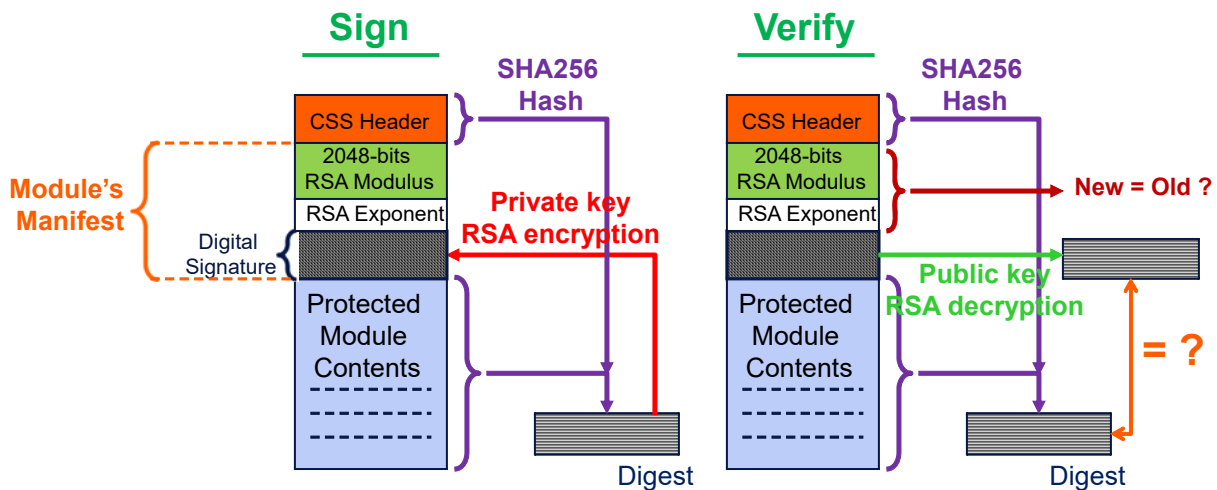


Figure 3.4. Sign and Verify Procedures for Authenticated NVM Modules

3.4.10.1 Digital Signature Algorithm Details

As described the digital signature generation is a hash computation followed by an RSA encryption. This is performed within Intel as part of the NVM update image generation process and not performed by Intel software in the field, nor by the Intel® Xeon® Processor D-1500 Product Family LAN controller.

The different algorithms used are described in the following locations:

- PKCS #1 v2.1: RSA Cryptography Standard, RSA Laboratories, June 14, 2002 - www.rsa.com
- SHA family definition - http://csrc.nist.gov/publications/fips/fips180-3/fips180-3_final.pdf
- SHA usage with digital signatures - <http://csrc.nist.gov/publications/nistpubs/800-107/NIST-SP-800-107.pdf>
- SHA validation vectors - <http://csrc.nist.gov/groups/STM/cavp/documents/shs/SHAVS.pdf>



3.5 Configurable I/O Pins — Software-Definable Pins (SDPs)

The Intel® Xeon® Processor D-1500 Product Family LAN controller has two software-defined pins (SDP pins) per port that can be used for miscellaneous hardware or software-controllable purposes. Unless specified otherwise, these pins and their function are bound to a specific LAN device. The use, direction, and values of SDP pins are controlled and accessed by the Extended SDP Control (ESDP) register. To avoid signal contention, following power up, all four pins are defined as input pins.

Some SDP pins have specific functionality:

- The default direction of the SDP pins is loaded from the SDP Control word in the NVM.
- The SDP pins can also be configured for use as External Interrupt Sources (GPI). To act as GPI pins, the desired pins must be configured as inputs and enabled by the GPIE register. When enabled, an interrupt is asserted following a rising-edge detection of the input pin (rising-edge detection occurs by comparing values sampled at the internal clock rate, as opposed to an edge-detection circuit). When detected, a corresponding GPI interrupt is indicated in the EICR register.

Note: In an external PHY topology (10 GbE SFP+ and 10GBASE-T), the SPD's can detect an interrupt only on the rising edge of the signal from the PHY. Some discrete PHYs (like the X577-AT2 and Inphi CS4227) are designed for active-low signaling so adding a trigger inverter between the PHY and the level shifter on each pin is required to ensure proper functionality. For implementation details, contact your local Intel representative.

Note: An SDP configured as output can also generate interrupts, but this is not a recommended configuration.

The bit mappings are listed in the following table for clarity.

Table 3.12. GPI to SDP Bit Mappings

SDP Pin To Be Used As GPI	ESDP Field Settings		Resulting EICR Bit (GPI)
	Directionality	Enable as GPI interrupt	
1	SDP1_IODIR	SDP1_GPIEN	26
0	SDIPIODIR	SDP0_GPIEN	25

- The lowest SDP pins (SDP0_0) of port 0 can be used to encode the NC-SI package ID of the Intel® Xeon® Processor D-1500 Product Family LAN controller. This ability is enabled by setting bit 15 in NC-SI Configuration 2 word (offset 0x07) of the NVM. The 3-bit package ID is encoded as follows: Package ID = {0, SDP0_0, 0}.
- When the SDP pins are used as IEEE1588 auxiliary signals they can generate an interrupt on any transition (rising or falling edge), refer to [Section 7.6.4](#).

All SDP pins can be allocated to hardware functions. See more details on IEEE1588 auxiliary functionality in [Section 7.6.4](#) while I/O pins functionality are programmed by the TimeSync Auxiliary Control (TSAUXC) register.

If mapping of these SDP pins to a specific hardware function is not required then the pins can be used as general purpose software defined I/Os. For any of the function-specific usages, the SDP I/O pins should be set to native mode by software setting of the SDPxxx_NATIVE bits in the ESDP register. Native mode in those SDP I/O pins, defines the pin functionality at inactive state (reset or power down) while behavior at active state is controlled by the software. The hardware functionality of these SDP I/O pins differs mainly by the active behavior controlled by software.

The following table lists the setup required to achieve each of the possible SDP configurations:



Table 3.13. SDP Settings

SDP	Usage	NVM Setting	ESDP		
			SDPx_NATIVE	SDPx_IODIR	SDP1_Function
0	SDP	N/A	0b	Input/Output	N/A
	NC-SI package ID (Port #0 only)	bit 15 in NC-SI Configuration 2 NVM word	0b	Input	
	1588 functionality as defined by the TSSDP register	N/A	1b	Input/Output	
1	SDP	N/A	0b	Input/Output	N/A
	PCI disable	bit 2 in PCIe Control 3 Word	N/A	N/A	
	1588 functionality as defined by the TSSDP register	N/A	1b	Input/Output	0b

3.6 LEDs

The Intel® Xeon® Processor D-1500 Product Family LAN controller implements two output drivers intended for driving external LED circuits per port. Each of the LED outputs can be individually configured to select the particular event, state, or activity, which is indicated on that output. In addition, each LED can be individually configured for output polarity as well as for blinking versus non-blinking (steady-state) indication.

The configuration for LED outputs is specified via the LEDCTL register. Furthermore, the hardware-default configuration for all LED outputs can be specified via NVM fields thereby supporting LED displays configurable to a particular OEM preference.

Each of the LED's can be configured to use one of a variety of sources for output indication. For more information on the MODE bits refer to the LEDCTL register description in [Section 8.2.2.1.12](#).

The *IVRT* bits enable the LED source to be inverted before being output or observed by the blink-control logic. LED outputs are assumed to normally be connected to the negative side (cathode) of an external LED.

The *BLINK* bits control whether the LED should be blinked (on for 200 ms, then off for 200 ms) while the LED source is asserted. The blink control can be especially useful for ensuring that certain events, such as *ACTIVITY* indication, cause LED transitions, which are sufficiently visible by a human eye.

Note: The *LINK/ACTIVITY* mode ignores the *BLINK* value. The LED's behavior in this mode is:

- Off if there is no *LINK*
- On if there is *LINK* and no *ACTIVITY*
- Blinks if there is *LINK* and *ACTIVITY*



The following mapping is used to specify the LED control source (MODE) for each LED output:

MODE ¹	Selected Mode	Source Indication
0000b	LINK_UP	Asserted or blinking according to the LEDx_BLINK setting when any speed link is established and maintained.
0001b	LINK_10G	Asserted or blinking according to the LEDx_BLINK setting when a 10 Gb/s link is established and maintained.
0010b	MAC_ACTIVITY	Active when link is established and packets are being transmitted or received. In this mode, the LEDx_BLINK must be set.
0011b	FILTER_ACTIVITY	Active when link is established and packets are being transmitted or received that passed MAC filtering. In this mode, the LEDx_BLINK must be set.
0100b	LINK/ACTIVITY	Asserted steady when link is established and there is no transmit or receive activity. Blinking when there is link and receive or transmit activity.
0101b	LINK_1G	Asserted or blinking according to the LEDx_BLINK setting when a 1 Gb/s link is established and maintained.
0110	LINK_100	Reserved.
0111b	LINK_2_5G	Asserted or blinking according to the LEDx_BLINK setting when a 2.5 Gb/s link is established and maintained.
1110b	LED_ON	Always asserted or blinking according to the LEDx_BLINK setting.
1111b	LED_OFF	Always de-asserted.

1. Undefined values are reserved.

3.7 Network Management Interface (MDIO or I²C)

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports MDIO (Section 3.7.2) and/or I²C (Section 3.7.3) interfaces, per port, for control plane connection between the MAC (master side) and external PHY devices. The management interface enables both MAC and software access to the PHY for monitoring and controlling of the PHY's functionality and configurations. The Intel® Xeon® Processor D-1500 Product Family LAN controller is compliant with the IEEE Std 802.3 Clause 45 in 10 GbE and 1 GbE operation. The Intel® Xeon® Processor D-1500 Product Family LAN controller also supports IEEE Std 802.3 Clause 22 frame formats and register address space for accessing legacy PHY registers.

There are various connectivity options to external PHYs some use MDIO, such as Base-T PHYs, others might use I²C, like SFP+ modules while in some special scenarios both might be needed such as when an SoC does not natively support SFI and requires the system designer to use both an external PHY for SFI translation, and an external SFP+ module. In such scenarios the Intel® Xeon® Processor D-1500 Product Family LAN controller might need to use the MDIO to manage the external PHY and the I²C to manage the external SFP+ module.

3.7.1 I²C or MDIO Selection

Since the Intel® Xeon® Processor D-1500 Product Family LAN controller IP supports both MDC/MDIO and I²C management interfaces the final board and device setup must be communicated to the software device driver via an NVM loaded configuration.

The following settings are provided by the IP, in the NW Management Interface Select register, for the software device driver to read. Based on these settings the software device driver is able to understand if there is a companion device connected through either of the port's management interfaces. After power on, the software device driver reads the identity of the device connected through the interface and learns the physical topology.

Table 3.14. NW Management Selection Settings - Per Port

Field Name	Description	SFP+: SoC Without SFI PHY	Dual 10/1 GbE BASE-T PHY
I2C ACT	A value of 1b indicates that the port's I ² C interface is active and connected to an external PHY or module.	1b	0b
MDIO ACT	A value of 1b indicates that the port's MDIO interface is active and connected to an external PHY or module.	0b	1b
MDIO IF MODE	Used to indicate what MDC/MDIO mode is used on Port0 when NW Management Mode is MDC/MDIO. 1b = Use IEEE Clause 22 Frame Structure. 0b = Use IEEE Clause 45 Frame Structure.	0b	0b

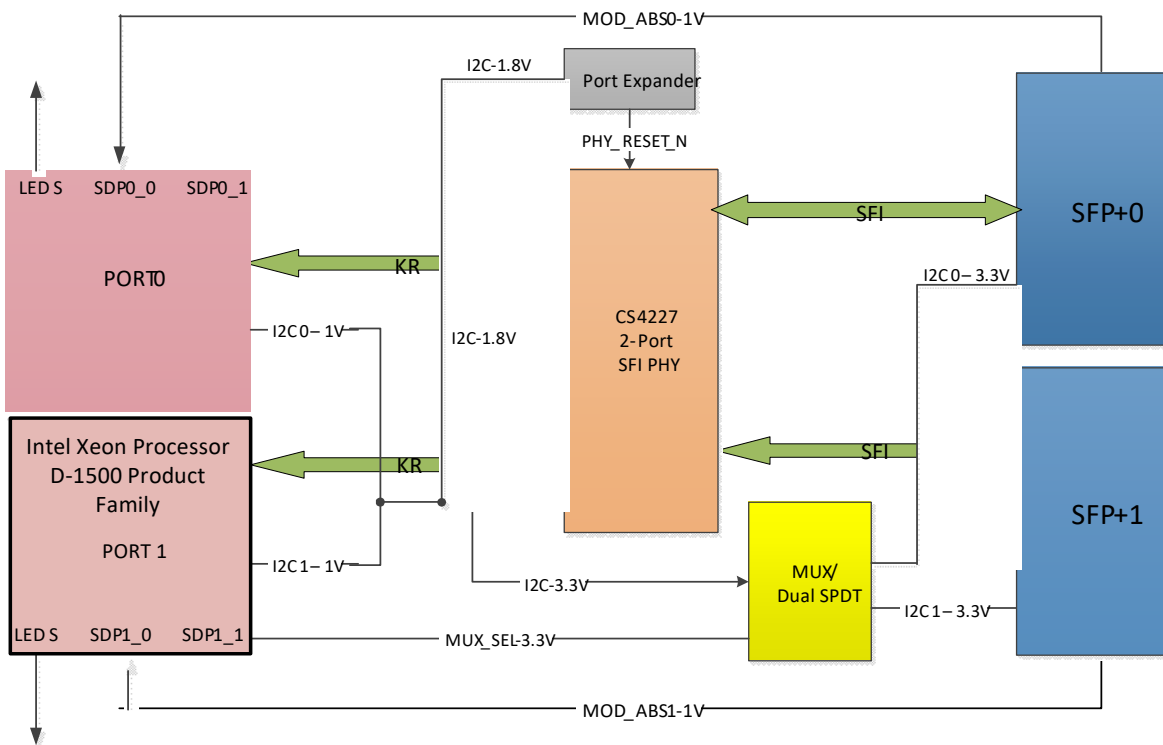


Figure 3.5. SFP+: SoC w/o SFI PHY

Figure 3.5 illustrates a mode where the SoC's Ethernet interface does not support SFI, or a scenario where system design limitations, require adding an external PHY that translates from the SoC's supported interface to SFI.

The management I/F used to configure the External PHY and the SFP+ modules is I²C, therefore an external I²C MUX block is placed on board to connect the SoCs (Broadwell-DE LAN controller's) per port I²C interfaces to both the per port SFP+ module and the shared I²C interface of the External PHY. The MUX is controlled using the SDPs driven from the Intel® Xeon® Processor D-1500 Product Family LAN controller.

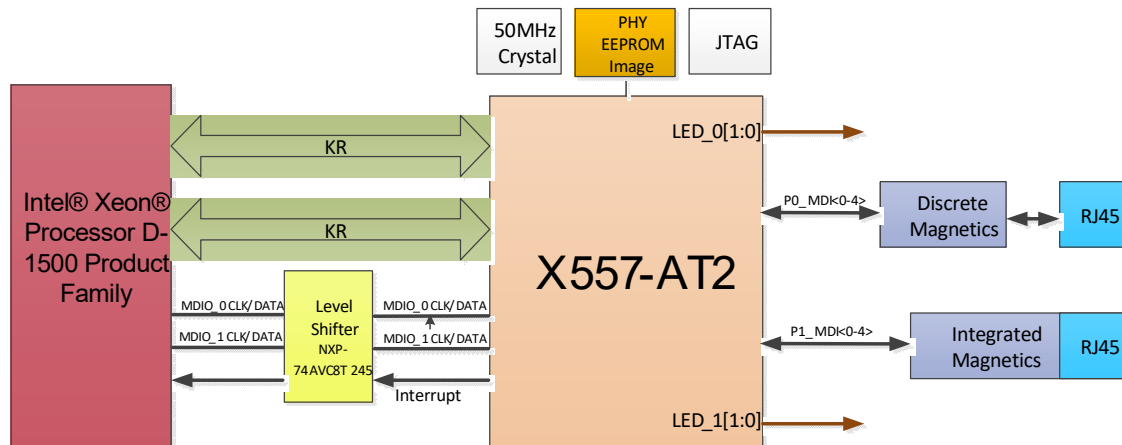



Figure 3.6. Dual 10/1 GbE BASE-T PHY

Figure 3.6 illustrates a BASE-T solution. The management I/F used to configure the external BASE-T PHY is MDC/MDIO. Since the BASE-T PHY supports only a single, shared, MDC/MDIO interface an on-board MDC/MDIO MUX must be added. The MUX is controlled using the Intel® Xeon® Processor D-1500 Product Family LAN controller (SoC) driven SDPs.

Management Interface Sharing

Since the registers that control the management interface can be used by both drivers and firmware the ownership must be centrally managed by firmware. This is done using a host command described in Section 10.8.4.

3.7.2 Management Data I/O Interface (MDIO)

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports the MDIO interface for a control plane connection between the MAC (master side) and PHY devices. The MDIO interface enables both MAC and software access to the PHY for monitoring and controlling of the PHY’s functionality and configurations. The Intel® Xeon® Processor D-1500 Product Family LAN controller is compliant with the IEEE802.3 clause 45 in both 10 GbE and 1 GbE operation. The Intel® Xeon® Processor D-1500 Product Family LAN controller also supports IEEE 802.3 clause 22 frame formats and register address space for accessing legacy PHY registers.

Note: The MDIO interface uses LVTTTL signaling as defined in Clause 22 of the IEEE802.3 standard. To access PHYs that support clause 45 1.2V electrical interface, level translators might be needed on board.

Figure 3.7 shows the basic connectivity between the PHY and MAC.

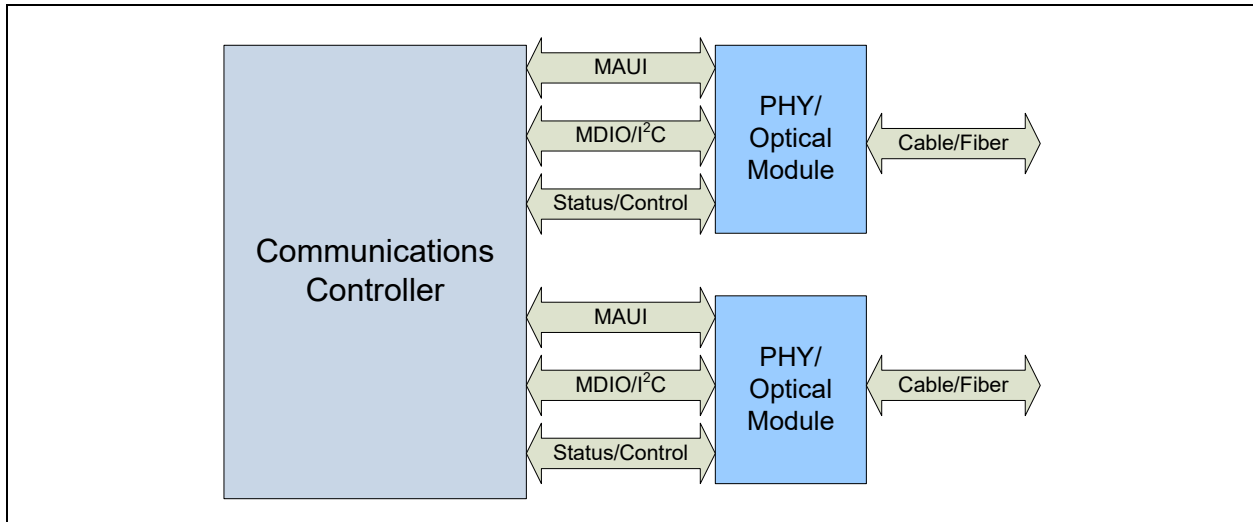


Figure 3.7 Basic PHY MAC Connectivity

The MDIO interface is a simple 2-wire serial interface between the MAC and PHY and is used to access Control and Status registers inside the PHY. The interface is implemented using two LVTTTL I/Os:

1. MDC — MDIO-interface clock signal driven by an external MAC (STA) device.
2. MDIO — Read/write data between an external MAC and PHY.

3.7.2.1 MDIO Timing Relationship to MDC

The MDC clock toggles during a read/write operation at a frequency of 24 MHz, 2.4 MHz or 240 KHz depending on the link speed and register bit HLREG0.MDCSPD as listed in [Table 3.15](#).

Table 3.15 MDC Frequency as Function of Link Speed and MDC Speed Bit

Link Speed	MDCSPD=1b	MDCSPD=0b
10 Gb/s	24 MHz	2.4 MHz
1 Gb/s	2.4 MHz	240 KHz

MDIO is a bidirectional signal that can be sourced by the Station Management Entity (STA) or the PHY. When the STA sources the MDIO signal, the STA must provide a minimum of 10 ns of setup time and a minimum of 10 ns of hold time referenced to the rising edge of MDC, as shown in [Figure 3.8](#) (measured at the MII connector).




Figure 3.8 MDIO Timing Sourced by the MAC

When the MDIO signal is sourced by the PHY, it is sampled by the MAC (STA) synchronously with respect to the rising edge of MDC. The clock to output delay from the PHY, as measured at the MII connector, must be a minimum of 0 ns, and a maximum of 300 ns, as shown in Figure 3.9.




Figure 3.9 MDIO Timing Sourced by the PHY

3.7.2.2 IEEE802.3 Clause 22 and Clause 45 Differences

IEEE802.3 clause 45 provides the ability to access additional device registers while still retaining logical compatibility with interface defined in Clause 22. Clause 22 specifies the MDIO frame format and uses an ST code of 01 to access registers. In clause 45, additional registers are added to the address space by defining MDIO frames that use a ST code of 00.

Clause 45 (MDIO interface) major concepts:

- Preserve management frame structure defined in IEEE 802.3 Clause 22.
- Define mechanism to address more registers than specified in IEEE802.3 Clause 22.
- Define ST and OP codes to identify and control the extended access functions.



3.7.2.3 MDIO Management Frame Structure

The MDIO interface frame structure defined in IEEE802.3 clause 22 and Clause 45 are compatible so that the two systems supporting different formats can co-exist on the same MDIO bus. The Intel® Xeon® Processor D-1500 Product Family LAN controller supports both frame structures to enable interfacing PHYs that support either protocol.

The basic frame format as defined in IEEE802.3 clause 22 can optionally be used for accessing legacy PHY registers is listed in Table 3.16.

Table 3.16 Clause 22 Basic MDIO Frame Format

	Management Frame Fields							
Frame	Pre	ST	OP	PRTAD	REGAD	TA	Data	Idle
Read	1...1	01	10	PPPPP	RRRRR	Z0	DDDDDDDDDDDDDDDD	Z
Write	1...1	01	01	PPPPP	RRRRR	10	DDDDDDDDDDDDDDDD	Z

The MDIO interface defined in clause 45 uses indirect addressing to create an extended address space enabling access to a large number of registers within each MDIO Managed Device (MMD). The MDIO management frame format is listed in Table 3.17.

Table 3.17 Clause 45 Indirect Addressing MDIO Frame Format

	Management Frame Fields							
Frame	Pre	ST	OP	PRTAD	DEVAD	TA	Address / Data	Idle
Address	1...1	00	00	PPPPP	EEEEEE	10	AAAAAAAAAAAAAAAA	Z
Write	1...1	00	01	PPPPP	EEEEEE	10	DDDDDDDDDDDDDDDD	Z
Read	1...1	00	11	PPPPP	EEEEEE	Z0	DDDDDDDDDDDDDDDD	Z
Post-Read Increment Address	1...1	00	10	PPPPP	EEEEEE	Z0	DDDDDDDDDDDDDDDD	Z

- A <10> pattern indicates that the frame is an indirect post-read-increment-address operation.

For Clause 22 compatible frames:

3. When an I²C or MDIO operation completes, release the semaphore by sending a PHY Token Request command with <Command Type> = Token Release.

Note: Token must be released after no longer than one second.

- A <10> pattern indicates a direct read transaction from a register.
- A <01> pattern indicates a direct write transaction to a register.

These patterns ensure a transition from the default value of one on the MDIO signal, and identifies the start of frame.

Operation Code (OP) — The OP field indicates the type of transaction being performed by the frame.

For Clause 45 compatible frames:

- A <00> pattern indicates that the frame payload contains the address of the register to access.
- A <01> pattern indicates that the frame payload contains data to be written to the register whose address was provided in the previous address frame.



- A <11> pattern indicates that the frame is an indirect read operation.
- A <10> pattern indicates that the frame is an indirect post-read-increment-address operation.

For Clause 22 compatible frames:

- A <10> pattern indicates a direct read transaction from a register.
- A <01> pattern indicates a direct write transaction to a register.

Port Address (PRTAD) — The PRTAD is five bits, allowing 32 unique PHY port addresses. The first *PRTAD* bit to be transmitted and received is the MSB of the address. A station management entity must have prior knowledge of the appropriate port address for each port to which it is attached, whether connected to a single port or to multiple ports.

Device Address (DEVAD) — The DEVAD is five bits, allowing 32 unique MMDs per port. The first *DEVAD* bit transmitted and received is the MSB of the address. This field is relevant only in clause 45 compatible frames (ST=<00>).

Register Address (REGAD) — The REGAD is five bits, allowing 32 individual registers to be addressed within each PHY. The first *REGAD* bit transmitted and received is the MSB of the address. This field is relevant only in clause 22 compatible frames (ST=<01>).

Turnaround (TA) — The TA time is a 2-bit time spacing between the *DEVAD* field and the *Data* field of a management frame. This is to avoid contention during a read transaction. For a read or post-read-increment-address transaction, both the STA and the PHY must remain in a high-impedance state for the first bit time of the TA. The PHY must drive a zero bit during the second bit time of the TA of a read or post-read-increment-address transaction. During a write or address transaction, the STA must drive a one bit for the first bit time of the TA and a zero bit for the second bit time of the TA. [Figure 3.10](#) shows the behavior of the MDIO signal during the *TA* field of a read transaction.





Figure 3.10 Behavior of MDIO During TA Field of a Read Transaction

- Clause 45 compatible frames have 16-bit address/data fields. For an auto-negotiation address cycle, it contains the address of the register to be accessed on the next cycle. For the data cycle of a write frame, the field contains the data to be written to the register. For a read or post-read-increment-address frame, the field contains the contents of the register. The first bit transmitted and received must be bit 15.
- Clause 22 compatible frames have 16-bit data fields. The first data bit transmitted and received must be bit 15 of the register being addressed.

3.7.2.4 MDIO Direct Access

The MDI is accessed through registers MSCA and MSRWD. A single management frame is sent by setting bit MSCA.MDICMD to 1b after programming the appropriate fields in the MSCA and MSRWD registers. The MSCA.MDICMD bit is auto cleared after the read or write transaction completes. To execute clause 22 format write operations, the following steps should be done:

1. Data to be written is programmed in field MSRWD.MDIWRDATA.



2. Register MSCA is initialized with the appropriate control information (start, code, etc.) with bit MSCA.MDICMD set to 1b.
3. Wait for bit MSCA.MDICMD to reset to 0b when indicating that the transaction on the MDIO interface is complete.

The steps for clause 22 format read operations are identical to the write operation except that the data in field MSRWD.MDIWRDATA is ignored and the data read from the external device is stored in register field MSRWD.MDIRDDATA bits. Clause 45 format read/write operations must be performed in two steps. The address portion of the pair of frames is sent by setting register field MSCA.MDIADD to the desired address, field MSCA.STCODE to 00b (start code that identifies clause 45 format), and register field MSCA.OPCODE to 00b (clause 45 address register write operation). A second data frame must be sent after the address frame completes. This second frame executes the write or read operation to the address specified in the PHY address register.

3.7.3 I²C

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports 2-wire management interface (I²C) for connectivity to external SFP+ modules. Optical and direct attached copper PHYs use 2-wire management interface (I²C) as described in SFF8431 (SFP+) or SFF8636 (Direct attach Cu).

Note: Some external PHYs also support I²C management interface instead of MDIO. It is up to the board designer to choose which interface to use. See [Section 3.7.1](#).

The I²C interface operates via the I2CCMD and I2CPARAMS register set. Since this register set can be used by either software or firmware in alternation, its ownership must be acquired/released via the semaphore ownership taking/release flows.

The I²C interface can be used in two methods, a hardware based access, where the device initiates a transaction following a software device driver request via the I2CCMD register or a software controlled bit banging using the I2CPARAMS register.

3.7.3.1 Hardware Based I²C Access

The following flows should be used to access an I²C register.

As part of device initialization, or anytime before the actual access, the following parameters should be set:

- *I2CPARAMS.PHYADD* - the address of the device to access.
- *I2CPARAMS.ACCESS_WIDTH* - the width of the data to read or write (byte or word).

To execute a write access, the following steps should be done:

1. Check that register is ready: Poll the *I2CCMD.R* bit until it is read as 1b.
2. Command - The I2CCMD register is initialized with the appropriate PHY register address in the *REGADD* field, the data to write in the *DATA* field and the operation (write) to the *OP* field (0b).
 - a. If an interrupt is required, set the *I2CCMD.I* field
3. Check that command is done: Poll *I2CCMD.R* bit until it is read as 1b.
 - a. Check that no error is indicated in the *I2CCMD.E* field.

To execute a read access, the following steps should be done:

1. Check that the register is ready: Poll *I2CCMD.R* bit until it is read as 1b.
2. Command - The I2CCMD register is initialized with the appropriate PHY register address in *REGADD* field, and the operation (read) to the *OP* field (1b).



- a. If an interrupt is required, set the *I2CCMD.I* field
3. Check that command is done: Poll *I2CCMD.R* bit until it is read as 1b.
 - a. Check that no error is indicated in the *I2CCMD.E* field.
4. Read the data returned from the *I2CCMD.DATA* field. If a byte access is done (*I2CPARAMS.ACCESS_WIDTH* = 0b), only *DATA[7:0]* is valid.

Note: See Section 3.7.3.3 for the I²C commands supported when using the built-in read and write commands. When using the bit bang method any command can be given to the I²C device.

3.7.3.2 Bit Bang Based I²C Access

In this mode, the software device driver controls the I²C interface directly using the *I2CPARAMS* register according to the following table:

Pad	Field Controlling The Output Value	Field Reflecting The Input Value	Field Controlling The Output Enable Value ¹
I2C clock	CLK_OUT	CLK_IN	CLK_OE_N
I2C data	DATA_OUT	DATA_IN	DATA_OE_N

1. 0b = Pad is output. 1b = Pad is input.

3.7.3.3 Supported Commands

Note: The gray columns that follow denotes cycles driven by the I²C device. White columns denotes cycles driven by the Intel® Xeon® Processor D-1500 Product Family LAN controller.

When a word Read command (*I2CPARAMS.ACCESS_WIDTH* = 1b, *I2CCMD.OP* = 1b) is given, the following sequence is done by the Intel® Xeon® Processor D-1500 Product Family LAN controller:

Table 3.18. I²C Read Transaction - Dummy Write

1	7	1	1	8	1	
S	Device Address	Wr	A	Register Address	A	
	From <i>I2CCMD.PHYADD</i>	0	0	From <i>I2CCMD.REGADD</i>	0	

Table 3.19. I²C Read Transaction - Word Read

1	7	1	1	8	1	8	1	1
S	Device Address	Rd	A	Data	A	Data	A	P
	From <i>I2CPARAMS.PHYADD</i>	1	0	Stored in <i>I2CCMD.DATA[7:0]</i>	0	Stored in <i>I2CCMD.DATA[15:8]</i>	0	

When a byte read command (*I2CPARAMS.ACCESS_WIDTH* = 0b, *I2CCMD.OP* = 1b) is given the following sequence is done by the Intel® Xeon® Processor D-1500 Product Family LAN controller:

Table 3.20. I²C Read Transaction - Dummy Write

1	7	1	1	8	1
S	Device Address	Wr	A	Register Address	A
	From <i>I2CPARAMS.PHYADD</i>	0	0	From <i>I2CCMD.REGADD</i>	0



Table 3.21. I²C Read Transaction - Byte Read

1	7	1	1	8	1
S	Device Address	Rd	A	Data	A
	From <i>I2CPARAMS.PHYADD</i>	1	0	Stored in <i>I2CCMD.DATA[7:0]</i>	0

When a word Write command (*I2CPARAMS.ACCESS_WIDTH* = 1b, *I2CCMD.OP* = 0b) is given the following sequence is done by the Intel® Xeon® Processor D-1500 Product Family LAN controller:

Table 3.22. I²C Write Transaction - Word Write

1	7	1	8	1	1	8	1	8	1	1
S	Device Address	Wr	Register Address	A	P	Data	A	Data	A	P
	From <i>I2CPARAMS.PHYADD</i>	0	From <i>I2CCMD.REGADD</i>	0		From in <i>I2CCMD.DATA[7:0]</i>	0	From in <i>I2CCMD.DATA[15:8]</i>	0	

When a byte write command (*I2CPARAMS.ACCESS_WIDTH* = 0b, *I2CCMD.OP* = 0b) is given the following sequence is done by the Intel® Xeon® Processor D-1500 Product Family LAN controller:

Table 3.23. I²C Write Transaction - Byte Write

1	7	1	8	1	1	8	1
S	Device Address	Wr	Register Address	A	P	Data	A
	From <i>I2CPARAMS.PHYADD</i>	0	From <i>I2CCMD.REGADD</i>	0		From in <i>I2CCMD.DATA[7:0]</i>	0

3.8 Network Interface

The Intel® Xeon® Processor D-1500 Product Family LAN controller connects, over XGMII, to two MDI interfaces with supported speeds of 10 GbE and 1 GbE Ethernet. Each interface provides the following physical interfaces and electrical modes:

- Four 3.125 GHz lanes providing support for:
 - 1000BASE-KX for GbE backplane applications (IEEE802.3 clause 70)
- A single lane running at up to 10.3125 GHz providing support for:
 - 10GBASE-KR for GbE backplane applications (IEEE802.3 clause 72)
 - 1000BASE-KX for GbE backplane applications (IEEE802.3 clause 70)

3.8.1 Integrated PHY Support

See Appendix B.

3.8.2 Loopback Support

See Section 11.2.

3.8.3 Ethernet Flow Control (FC)

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports flow control as defined in 802.3x, as well as the specific operation of asymmetrical flow control defined by 802.3z.



Note: The Intel® Xeon® Processor D-1500 Product Family LAN controller can be configured to receive regular FC packets.

FC is implemented to reduce receive buffer overflows, which result in the dropping of received packets. FC also allows for local controlling of network congestion levels. This can be accomplished by sending an indication to a transmitting station of a nearly full receive buffer condition at a receiving station.

The implementation of asymmetric FC allows for one link partner to send FC packets while being allowed to ignore their reception. For example, not required to respond to PAUSE frames.

The following registers are defined for implementing FC.

- MAC Flow Control Register (MFLCN) — Enables FC and passing of control packets to the host.
- Flow Control Configuration (FCCFG) — Determines mode for Tx FC (No FC vs. link-based vs. priority-based). Note that if Tx FC is enabled then Tx CRC by hardware should be enabled as well (*HLREG0.TXCRCE* = 1b).
- Flow Control Source Address Low, High (RAL[0], RAH[0]).
- Flow Control Destination Address Low, High (*FCAMACL*, *FCAMACH*) — 6-byte FC multicast address.
- Flow Control Receive Threshold High (FCRTH[7:0]) — A set of 13-bit high watermarks indicating receive buffer fullness. A single watermark is used in link FC mode.
- Flow Control Receive Threshold Low (FCRTL[7:0]) — A set of 13-bit low watermarks indicating receive buffer emptiness. A single watermark is used in link FC mode.
- Flow Control Transmit Timer Value (FCTTV[3:0]) — A set of 16-bit timer values to include in transmitted PAUSE frame. A single timer is used in link FC mode.
- Flow Control Refresh Threshold Value (FCRTV) — 16-bit PAUSE refresh threshold value (in legacy FC FCRTV[0] must be smaller than FCTTV[0]).

3.8.3.1 Enabling Flow Control

To enable flow control, program the FCTTV, FCRTL, FCRTH, FCRTV and FCCFG registers. If flow control is not enabled, these registers should be written with 0x0 except for FCRTH[n].RTH field which must be always configured.

When Tx flow control is enabled, then Tx CRC by hardware should be enabled as well (*HLREG0.TXCRCE* = 1b).

Refer to [Section 3.8.3.4.2](#) and [Section 3.8.3.4.4](#) for more details.

3.8.3.2 MAC Control Frames and Reception of Flow Control Packets

3.8.3.2.1 MAC Control Frame — Other than FC

The IEEE specification reserved the Ethertype value of 0x8808 for MAC control frames, which are listed in [Table 3.24](#).

Table 3.24. MAC Control Frame Format

DA	The <i>Destination Address</i> field can be an individual or multicast (including broadcast) address. Permitted values for the <i>Destination Address</i> field can be specified separately for a specific control OpCode such as FC packets.
SA	Port Ethernet MAC Address (six bytes).
Type	0x8808 (two bytes).
Opcode	The MAC control OpCode indicates the MAC control function.



Table 3.24. MAC Control Frame Format

DA	The <i>Destination Address</i> field can be an individual or multicast (including broadcast) address. Permitted values for the <i>Destination Address</i> field can be specified separately for a specific control OpCode such as FC packets.
Parameters	The MAC control <i>Parameters</i> field must contain MAC control OpCode-specific parameters. This field can contain none, one, or more parameters up to a maximum of $\text{minFrameSize} = 20$ bytes.
Reserved field = 0x00	The <i>Reserved</i> field is used when the MAC control parameters do not fill the fixed length MAC control frame.
CRC	Four bytes.

3.8.3.2.2 Structure of 802.3X FC Packets

802.3X FC packets are defined by the following three fields (see Table 3.25):

1. A match on the six-byte multicast address for MAC control frames or a match to the station address of the device (Receive Address register 0). The 802.3x standard defines the MAC control frame multicast address as 01-80-C2-00-00-01.
2. A match on the *Type* field. The *Type* field in the FC packet is compared against an IEEE reserved value of 0x8808.
3. A match of the MAC control *Opcode* field has a value of 0x0001.

Frame-based FC differentiates XOFF from XON based on the value of the PAUSE *Timer* field. Non-zero values constitute XOFF frames while a value of zero constitutes an XON frame. Values in the *Timer* field are in units of pause quanta (such as slot time). A pause quanta lasts 64 byte times, which is converted into an absolute time duration according to the line speed.

Note: XON frame signals the cancellation of the pause from that was initiated by an XOFF frame. Pause for zero pause quanta.

Table 3.25. 802.3X Packet Format

DA	01_80_C2_00_00_01 (6 bytes).
SA	Port Ethernet MAC Address (6 bytes).
Type	0x8808 (two bytes).
Opcode	0x0001 (two bytes).
Time	XXXX (two bytes).
Pad	42 bytes.
CRC	Four bytes.

3.8.3.2.3 Operation and Rules

The Intel® Xeon® Processor D-1500 Product Family LAN controller in FC mode.

- Link FC is enabled by the *RFCE* bit in the MFLCN register.

Note: Link FC capability must be negotiated between link partners via the auto-negotiation process. It is the software device driver's responsibility to reconfigure the link FC settings (including *RFCE*) after the auto-negotiation process was resolved.

Once the receiver has validated the reception of an XOFF, or PAUSE frame, the device performs the following:

- Increments the appropriate statistics register(s).



- Initialize the pause timer based on the packet's PAUSE *Timer* field (overwriting any current timer's value).
- Disable packet transmission or schedule the disabling of transmission after the current packet completes.
 - Tx manageability traffic is bound to a specific TC as defined in the MNGTXMAP register, and is thus paused when its TC is paused.

Resumption of transmission can occur under the following conditions:

- Expiration of the PAUSE timer.
- Receiving an XON frame (a frame with its PAUSE timer set to 0b).

Both conditions clear the relevant TXOFF status bits in the Transmit Flow Control Status (TFCS) register and transmission can resume. Hardware records the number of received XON frames.

3.8.3.2.4 Timing Considerations

When operating at 10 GbE line speed, the Intel® Xeon® Processor D-1500 Product Family LAN controller must not begin to transmit a (new) frame more than 74 pause quanta after receiving a valid Link XOFF frame, as measured at the wires (a pause quantum is 512 bit times).

When operating at 1 GbE line speed, the Intel® Xeon® Processor D-1500 Product Family LAN controller must not begin to transmit a (new) frame more than two pause quanta after receiving a valid Link XOFF frame, as measured at the wires.

The 802.1Qbb draft 2.3, proposes that the tolerated response time for priority XOFF frames are the same as Link XOFF frames with extra budget of 19360 bit times if MACSec is used, or of two pause quanta otherwise. This extra budget is aimed to compensate the fact that decision to stop new transmissions from a specific TC must be taken earlier in the transmit data path than for the link flow control case.

3.8.3.3 PAUSE and MAC Control Frames Forwarding


Two bits in the Receive Control register control transfer of PAUSE and MAC control frames to the host. These bits are *Discard PAUSE Frames (DPF)* and *Pass MAC Control Frames (PMCF)*. Note also that any packet must pass the L2 filters as well.

- The *DPF* bit controls the transfer of PAUSE packets to the host. The same policy applies to link FC packets as listed in [Table 3.26](#). Note that any packet must pass the L2 filters as well.
- The *Pass MAC Control Frames (PMCF)* bit controls the transfer of non-PAUSE packets to the host. Note that when link FC frames are not enabled (RFCE = 0b) then link FC frames are considered as MAC control frames for this case.

Note: When virtualization is enabled, forwarded control packets are queued according to the regular switching procedure.

Table 3.26. Transfer of PAUSE Packet to Host (DPF Bit)

RFCE	DPF	Link FC Handling
0b	X	Treat as MAC control (according to <i>PMCF</i> setting).
1b	0b	Accept.
1b	1b	Reject.

**Table 3.26. Transfer of PAUSE Packet to Host (DPF Bit)**

0b	0b	Treat as MAC control (according to <i>PMCF</i> setting).
0b	1b	Treat as MAC control (according to <i>PMCF</i> setting).
1b	X	Unsupported setting.

3.8.3.4 Transmitting PAUSE Frames

The Intel® Xeon® Processor D-1500 Product Family LAN controller generates PAUSE packets to ensure there is enough space in its receive packet buffers to avoid packet drop. The Intel® Xeon® Processor D-1500 Product Family LAN controller monitors the fullness of its receive FIFOs and compares it with the contents of a programmable threshold. When the threshold is reached, the Intel® Xeon® Processor D-1500 Product Family LAN controller sends a PAUSE frame.

Note: Similar to receiving flow control packets previously mentioned, software can enable FC transmission by setting the *FCCFG.TFCE* field only after it is negotiated between the link partners (possibly by auto-negotiation).

3.8.3.4.1 Operation and Rules

The *TFCE* field in the Flow Control Configuration (FCCFG) register enables transmission of PAUSE packets as well as selects FC mode.

The content of the Flow Control Receive Threshold High (FCRTH) register determines at what point the Intel® Xeon® Processor D-1500 Product Family LAN controller transmits the first PAUSE frame. The Intel® Xeon® Processor D-1500 Product Family LAN controller monitors the fullness of the receive FIFO and compares it with the contents of FCRTH. When the threshold is reached, the Intel® Xeon® Processor D-1500 Product Family LAN controller sends a PAUSE frame with its pause time field equal to FCTTV.

At this time, the Intel® Xeon® Processor D-1500 Product Family LAN controller starts counting an internal shadow counter (reflecting the pause time-out counter at the partner end). When the counter reaches the value indicated in FCRTV register, then, if the PAUSE condition is still valid (meaning that the buffer fullness is still above the low watermark), an XOFF message is sent again.

Once the receive buffer fullness reaches the low water mark, the Intel® Xeon® Processor D-1500 Product Family LAN controller sends an XON message (a PAUSE frame with a timer value of zero). Software enables this capability with the XONE field of the FCRTL.

The Intel® Xeon® Processor D-1500 Product Family LAN controller sends a PAUSE frame if it has previously sent one and the FIFO overflows. This is intended to minimize the amount of packets dropped if the first PAUSE frame did not reach its target.

3.8.3.4.2 Flow Control High Threshold — FCRTH

The Intel® Xeon® Processor D-1500 Product Family LAN controller sends a PAUSE frame when the Rx packet buffer is full above the high threshold. The threshold should be large enough to overcome the worst case latency from the time that crossing the threshold is sensed until packets are not received from the link partner.

Referring to Annex O of IEEE802.1Qbb rev 2.3, worst case latency depends on three parameters:

1. Maximum frame size over the TC for which FCRTH is computed. It is referred as MaxFrame(TC).
2. Maximum frame size over the link (all TCs altogether). It is referred as MaxFrame(link).
3. Whether or not MACsec is enabled or disabled over the link.



Three values are envisaged for MaxFrame:

- 1.5 KB (Ethernet - jumbo disabled)
- 9.5 KB (jumbo enabled)

Worst case latency, which is referred as Standard Delay Value (Std DV), is given by:

Std DV = MaxFrame(TC) + MaxFrame(link) + 2 x Cable Delay + 2 x Interface Delay + Higher Layer Delay + (MACSec enabled = yes) x Sec Y Transmit Delay

Std DV (bit time units) = MaxFrame(TC) + MaxFrame(link) + 672 + 2 x 5,556 + 2 x (25,600 + 8,192 + 2 x 2,048) + 6,144 + (MACSec enabled = yes) x (MaxFrame(link) + 3,200)

MaxFrame(TC) term and MaxFrame(link) term included in Sec Y Transmit Delay correspond to worst case scenarios issued by the link partner. All other terms in Std DV formula must take in account worst case incoming traffic pattern which would lead to worst case buffer utilization as per the internal architecture of Rx packet buffer in the Intel® Xeon® Processor D-1500 Product Family LAN controller.

Internal architecture of the Rx packet buffer has the following restrictions:

1. Any packet starts at 32 byte aligned address.
2. Any packet has an internal status of 32 bytes. As a result, the Rx packet buffer is used at worst conditions when the Rx packet includes 65 bytes that are posted to the host memory. Assuming that the CRC bytes are not posted to host memory then in the worst case the Rx packet buffer can be filled at 1.44 higher rate than the wire speed (69-byte packet including CRC + 8-byte preamble + 12-byte back-to-back IFS consumes 4 x 32 bytes = 128 bytes on the Rx packet buffer).
3. An additional packet from the concerned traffic class may be inserted into the Rx packet buffer due to the internal loopback switch just before it is decided to issue XOFF to the link partner.

It leads to the following revised formula for the Intel® Xeon® Processor D-1500 Product Family LAN controller:


The Intel® Xeon® Processor D-1500 Product Family LAN controller DV (bit time units) = 1.44 x [(MaxFrame(link) + 672 + 2 x 5,556 + 2 x (25,600 + 8,192 + 2 x 2,048) + 6,144 + (MACSec enabled = yes) x (3,200)] + MaxFrame(TC) + MaxFrame(TC) + (MACSec enabled = yes) x MaxFrame(link).

FCRTH must be set to the size of the Rx packet buffer allocated to the TC minus the Intel® Xeon® Processor D-1500 Product Family LAN controller DV.

Table 3.27. Delay Values (DV) Used For FCRTH

ReservedMACsec Enabled	9.5 KB Jumbo Enabled	Reserved	Broadwell-DE LAN controller DV
No	No		24 KB
No	No		25 KB
No	Yes		50 KB
No	Yes		35 KB
Yes	No		27 KB
Yes	No		27 KB
Yes	Yes		60 KB
Yes	Yes		45 KB

Note: 9.5 KB jumbo enabled/disabled is a global setting per port which concerns all TCs..



3.8.3.4.3 FC Low Threshold — FCRTL

The low threshold value is aimed to protect against wasted available host bandwidth. There is some latency from the time that the low threshold is crossed until the XON frame is sent and packets are received from the link partner. The low threshold must be set high enough so that the Rx packet buffer does not get empty before any new entire packets are received from the link partner. When considering data movement from the Rx packet buffer to host memory, then large packets represent the worst. Assuming the host bandwidth is about the bandwidth on the wire (when dual ports are active at a given time), and assuming a PCIe round trip is required to get the receive descriptors, we get the following formula for FCRTL:

$$\text{FCRTL} = 2 \times \text{MaxFrame(TC)} + \text{PCIe round trip delay}$$

PCIe round trip delay is assumed to be $\sim 1 \mu\text{s}$ and it must cover for worst case incoming traffic pattern (buffer utilization by 1.44 than wire rate):

$$\text{FCRTL (bit time units)} = 2 \times \text{MaxFrame(TC)} + 1.44 \times 10,000$$

Setting the FCRTL to lower values than expressed by the previous equation is permitted. It might simply result with potential sub-optimal use of the PCIe bus once bandwidth is available.

Table 3.28. Intel® Xeon® Processor D-1500 Product Family LAN Controller FCRTL

9.5 KB Jumbo Enabled	Intel® Xeon® Processor D-1500 Product Family LAN Controller DV
No	5 KB
No	7 KB
Yes	21 KB
Yes	7 KB

3.8.3.4.4 Packet Buffer Size

When FC is enabled, the total size of a TC packet buffer must be large enough for the Low and high thresholds. In order to avoid constant transmission of XOFF and XON frames it is recommended to add some space for hysteresis type of behavior. The difference between the two thresholds is recommended to be at least one frame size (when 9.5 KB jumbo frames are expected over the TC) and larger than a few frames in other cases (4.5 KB for instance). If the available Rx buffer is large enough, it is recommended to increase as much as possible the hysteresis budget. If the available Rx buffer is not large enough it might be required to cut both the low threshold as well as the hysteresis budget.

- For a best effort TC:
 - Rx packet buffer size = FCRTL, as the same considerations than described in [Section 3.8.3.4.3](#) play here to avoid bubbles over PCIe

The total Rx packet buffer size available to a port for all its supported TCs is either 384 KB, 320 KB, or 256 KB, depending on the size allocated to the Flow Director table, 0 KB, 64 KB, or 128 KB, respectively.

The following table lists the recommended settings for the supported combinations. Note that reducing the number of TCs of a port to what is really needed, helps increasing the port's throughput.



Table 3.29. Some Recommended Rx Packet Buffer Settings

Flow Director Table Size	9.5 KB Jumbo Enabled	Packet Buffer Size of Any of the 4 Best Effort TCs
No	No	33 KB
No	Yes	27 KB
No	No	32 KB
No	Yes	22 KB
64 KB	No	25 KB
64 KB	Yes	19 KB
64 KB	No	24 KB
64 KB	Yes	14 KB
128 KB	No	17 KB
128 KB	No	16 KB

Notes: In some of the previous cases, it has been necessary to get compromised on the rules for hysteresis and FCRTL in order to fit the size available for Rx packet buffer.

In some other cases, after having applied all the rules there was an exceeding available Rx packet buffer left that has been used to extend the hysteresis budgets.

In all cases, FCRTL rule has been applied as is, since compromising on it is not allowed and extending it provides no performance benefits.

3.8.4 Inter Packet Gap (IPG) Control and Pacing

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports transmission pacing by extending the IPG (the gap between consecutive packets). The pacing mode enables the average data rate to be slowed in systems that cannot support the full link rate (10 GbE, 1 GbE or 100 Mb/s). As listed in Table 3.30, the pacing modes work by stretching the IPG in proportion to the data sent. In this case the data sent is measured from the end of preamble to the last byte of the packet. No allowance is made for the preamble or default IPG when using pacing mode.

Example 1:

Consider a 64-byte frame. To achieve a 1 GbE data rate when link rate is 10 GbE and packet length is 64 bytes (16 Dwords), add an additional IPG of 144 Dwords (nine times the packet size to reach 1 GbE). When added to the default IPG gives an IPG of 147 Dwords.

Example 2:

Consider a 65-byte frame. To achieve a 1 GbE data rate when link rate is 10 GbE and packet length is 65 bytes (17 Dwords when rounded up) add an additional IPG of 153 Dwords (nine times the packet duration in Dwords). When added to the default IPG gives an IPG of 156 Dwords. Note that in this case, where the packet length counted in Dwords is not an integer, count any fraction of a Dword as an entire Dword for computing the additional IPG.



Table 3.30 lists the pacing configurations supported by the Intel® Xeon® Processor D-1500 Product Family LAN controller at link rates of 10 GbE. When operating at lower link speeds the pacing speed is proportional to the link speed.

Table 3.30. Pacing Speeds at 10 GbE Link Speed

Pacing Speeds (Gb/s)	Delay Inserted into IPG	Register Value
10 (LAN)	None	0000b
9.294196 (WAN)	1 byte for 13 transmitted	1111b
9.0	1 Dword for 9 transmitted	1001b
8.0	1 Dword for 4 transmitted	1000b
7.0	3 Dwords for 7 transmitted	0111b
6.0	2 Dwords for 3 transmitted	0110b
5.0	1 Dwords for 1 transmitted	0101b
4.0	3 Dwords for 2 transmitted	0100b
3.0	7 Dwords for 3 transmitted	0011b
2.0	4 Dwords for 1 transmitted	0010b
1.0	9 Dwords for 1 transmitted	0001b
10	None	Default

Note: Pacing is configured in the *PACE* field of the Pause and Pace (PAP) register.

3.9 Physical Functions (PF), Virtual Functions (VF) and ARI (Alternative RID Interpretation)


The BIOS provided by the project natively supports Address Translation Services (ATS) and ARI implementations, but any BIOS that requires PF and VF support needs ATS and ARI implementations. Both ACS and ARI are required in order for PCIe configuration space to be able to address higher than 7 VFs PCIe configurations traditionally configures PF and VFs using 16 bits. ATS and ARI allow more (higher) than 7 functions to be addressed.

The number of VFs, in addition to the PFs, are indicated in the system NVM section for the LAN controllers LAN0 and LAN1. A function in the NVM image known as I/O Virtualization (IOV) needs to be enabled. These newly created VF shows up in PCIe configuration space as base/device/function.

In short, to enable hardware based Single Root I/O Virtualization (SR-IOV) and benefit from use of VFs and the built-in Virtual Ethernet Bridge (VEB), you must have:

1. Ensure that ACS/ARI are supported in your version of BIOS.
2. Enable SR-IOV in your BIOS.
3. Use an OS that supports virtualization such as Xen OpenSource 4.0+, Xen/Citrix* 5.0+, VMWare* 5.1 or higher, RHEL* 6.x and above, SLES* 11 and above and Windows* 2008 R2 or higher.
4. Enable IOV in the NVM section and indicated the # of VMs desired to create.

For further information regarding SR-IOV, refer to the published SR-IOV Primer titled Intel® Networking and Communications Products - PCI-SIG SR-IOV Primer - Rev. 2.0.



4.0 Initialization

4.1 Reset Operation

The Intel® Xeon® Processor D-1500 Product Family LAN controller reset sources are described in the sections that follow.

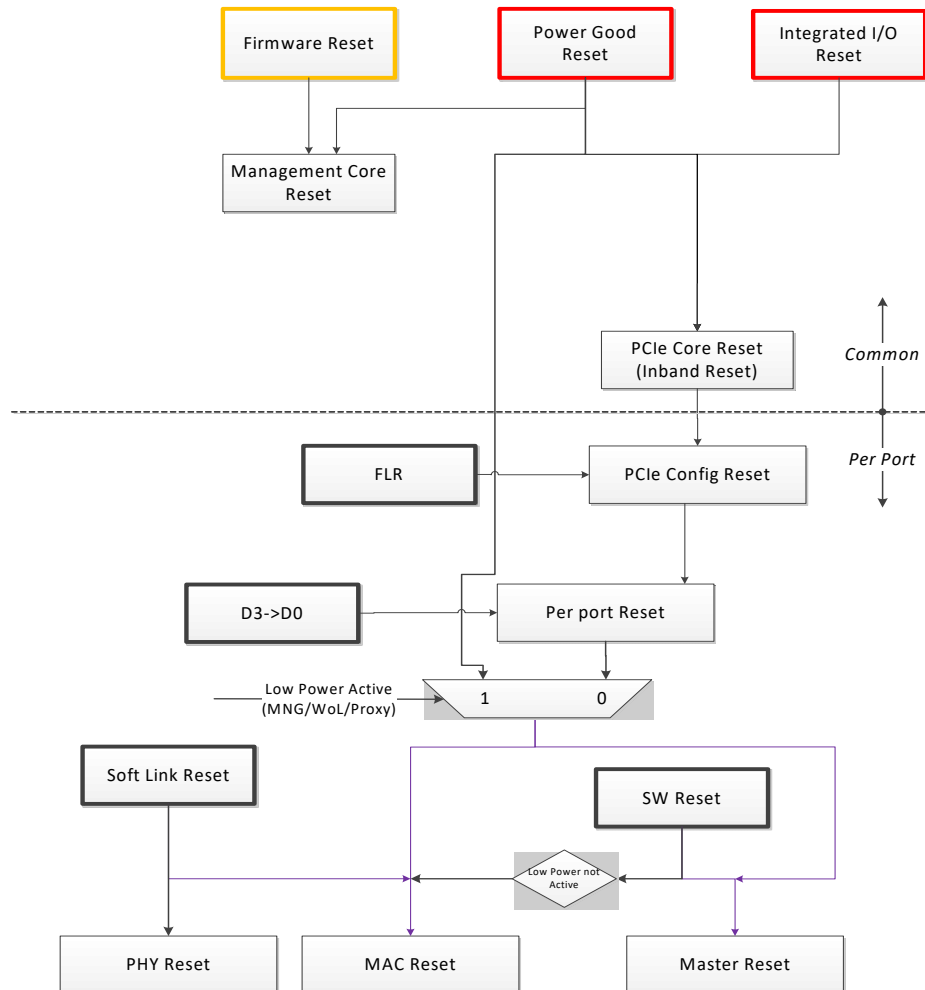


Figure 4.1. Rest Tree



4.1.1 Reset Sources/Reset Order

The different reset sources are organized in the following order (high-to-low):

1. Power Good Reset
2. Integrated I/O Reset
3. Physical Function Reset (PFLR)
4. Virtual Function Reset (VFLR)

A higher-level reset is either followed by a lower-level reset or directly resets the information related with a lower-level reset.

4.1.1.1 Power Good Reset

This reset signal is the master reset of the Intel® Xeon® Processor D-1500 Product Family LAN controller. It is level sensitive, and while it is 0b, all of the registers are held in reset.

Power Good Reset changes state during system power up.

4.1.1.2 Integrated I/O Reset

This is the reset event that replaces the PCIe Reset (PE_RST_N) signal in a discrete LAN controller. Most of the internal blocks of the Intel® Xeon® Processor D-1500 Product Family LAN controller are reset on the rising edge of this reset.

This reset might also reset the integrated PHY depending on WoL and manageability requirements.

4.1.1.2.1 Integrated I/O Cold/Warm Reset

When the SoC is about to assert an Integrated I/O reset, it first issues an internal [Reset Warn](#) message to enable a graceful reset flow. Upon receiving this message, the Intel® Xeon® Processor D-1500 Product Family LAN controller:

- Stops sending new requests out to the Integrated I/O interfaces.
- Waits for all pending completion to arrive.
- Responds to all pending non-posted requests.

Further, during the period between receiving this message and the actual reset event, the Intel® Xeon® Processor D-1500 Product Family LAN controller is expected to respond to new incoming requests either normally or with a Unsupported Request (UR). Once these previous steps complete, the Intel® Xeon® Processor D-1500 Product Family LAN controller sends an acknowledge to the SoC to indicate that it might perform:

- A warm reset is equal to asserting the Integrated I/O reset

Or

- A cold reset. For example, a full power on

4.1.1.3 D3hot to D0 Transition

This is also known as an ACPI reset. The Intel® Xeon® Processor D-1500 Product Family LAN controller generates an internal reset on the transition from D3hot power state to D0 (caused after configuration writes from a D3-to-D0 power state). Note that this reset is per function and resets only the function that transitions from D3hot-to-D0 and does not reset the configuration space of the function. Its effect is equivalent to a software reset ([Section 4.1.1.5.1](#)).



4.1.1.4 Function Level Reset (FLR) Capability

The *FLR* bit is required for the Physical Function (PF) and per Virtual Function (VF). Setting this bit for a VF only resets the part of the logic dedicated to the specific VF and does not influence the shared part of the port. Setting the PF *FLR* bit resets the entire function.

4.1.1.4.1 FLR in Non-IOV Mode

A FLR reset to a function is equivalent to a D0 → D3 → D0 transition with the exception that this reset doesn't require any software device driver intervention in order to stop the master transactions of this function. FLR affects the device 1 parallel clock cycle from FLR assertion by default setting, or any other value defined by the *FLR Delay Disable* and *FLR Delay* fields in the PCIe Init Configuration 2 — Offset 0x02 word in the NVM.

4.1.1.4.2 Physical Function FLR (PFLR)

A FLR reset to the PF function in an IOV mode is equivalent to a FLR reset in non-IOV mode. All VFs in the PCIe function of the PF are affected.

The affected VFs are not notified of the reset in advance. The *RSTD* bit in *VFMailbox[n]* is set following the reset (per VF) to indicate to the VFs that a PF FLR took place. Each VF is responsible to probe this bit (such as after a timeout).

4.1.1.4.3 Virtual Function FLR (VFLR)

A VF operating in an IOV mode can issue a FLR. VFLR resets the resources allocated to the VF (like disabling the queues and masking interrupts). It also clears the PCIe configuration for the VF. There is no impact on other VFs or on the PF.

Tx and Rx flows for the queues allocated to this VF are disabled. All pending read requests are dropped and PCIe read completions to this function can be completed as unsupported requests.

Note: Clearing the *IOV Enable* bit in the IOV structure is equivalent to a VFLR to all VFs in the same port.

A VFLR does not release queues that were blocked due to malicious events. The PF device driver needs to release them using the matching bit in *WQBR_RX* and *WQBR_TX* registers. The matching bits in these registers should be cleared after each VFLR, even if not caused due to a malicious event.

4.1.1.5 Soft Resets

4.1.1.5.1 Software Reset

A software reset is done by writing to the *Device Reset* bit of the Device Control (CTRL.RST) register. The Intel® Xeon® Processor D-1500 Product Family LAN controller re-reads the per-function NVM fields after software reset. Bits that are not normally read from the NVM are reset to their default hardware values.

Note: This reset is per function and resets only the function that received the software reset.

PCI configuration space (configuration and mapping) of the device is unaffected. The MAC might or might not be reset (see [Section 4.1.4](#)).

Prior to issuing a software reset, the software device driver needs to execute the master disable algorithm.



4.1.1.5.2 Physical Function (PF) Software Reset

A software reset by the PF in IOV mode has the same consequences as a software reset in a non-IOV mode.

The procedure for PF software reset is as follows:

1. The PF driver disables master accesses by the device through the master disable mechanism. Master disable affects all VFs traffic.
2. Execute the procedure described in [Section 4.1.2](#) to synchronize between the PF and VFs.

VFs are expected to timeout and check on the *RSTD* bit in order to identify a PF software reset event. The *RSTD* bits are cleared on read.

4.1.1.5.3 VF Software Reset

A software reset applied by a VF is equivalent to a FLR reset to this VF with the exception that the PCIe configuration bits allocated to this function are not reset. It is activated by setting the *VTCTRL.RST* bit.

4.1.1.5.4 Force TCO

This reset is generated when manageability logic is enabled. It is only generated if enabled by the *Force TCO Reset* bit in the Common Firmware Parameters word in the NVM. If enabled by the NVM, the firmware triggers a port reset by setting the *CTRL.RST* bit. In pass-through mode it is generated when receiving a Force TCO SMB or NC-SI command with bit 0 set.

4.1.1.6 Link Reset

Also referred to as MAC reset.

Initiated by writing the *Link Reset* bit of the Device Control register (*CTRL.LRST*).

A link reset is equivalent to a software reset + reset of the MAC + reset of the PHY. The Intel® Xeon® Processor D-1500 Product Family LAN controller re-reads the per-function NVM fields after link reset. Bits that are not normally read from the NVM are reset to their default hardware values. Note that this reset is per function and resets only the function that received the link reset.

The PF in IOV mode can generate a link reset.

Prior to issuing link reset the software device driver needs to execute the master disable algorithm.

4.1.1.7 PHY Resets

A PHY reset event causes a link down and restarts the auto-negotiation process. This can take a few seconds to complete, which might result in the drop of the TCP sessions with the host and/or with a Manageability Controller (MC). Because the PHY can be accessed by the MC (via internal firmware) and by the software device driver concurrently, the software device driver should coordinate any PHY reset with the firmware using the following procedure:

1. Ensure that the *MMNGC.MNG_VETO* bit is cleared. If it is set, the MC requires a stable link and thus the PHY should not be reset at this stage. The software device driver can skip the PHY reset (if it is not mandatory) or wait for this bit to be cleared by the MC. See [Section 4.2](#) for more details on *MNG_VETO* bit.
2. Take ownership of the relevant PHY using the flow described in [Section 10.8.4](#).
3. Reset the PHY using register configuration.



For full details on how to reset the PHY.

4.1.2 Reset in a PCI-IOV Environment

Several mechanisms are provided to synchronize reset procedures between the PF and the VFs.

4.1.2.1 RSTI/RSTD

This mechanism is provided specifically for a PF software reset but can be used in other reset cases as follows.

- One of the following reset cases takes place:
 - Power Good Reset
 - Integrated I/O Reset
 - D3hot-to-D0 transition
 - FLR
 - Software reset by the PF
- The Intel® Xeon® Processor D-1500 Product Family LAN controller sets the *RSTI* bits in all the VF Mailbox registers. Once the reset completes, each VF can read its VF Mailbox register to identify a reset in progress.
 - The VF can poll the *RSTI* bit to detect if the PF is in the process of configuring the device.
- Once the PF completes configuring the device, it sets the CTRL_EXT.PFRSTD bit. As a result, the Intel® Xeon® Processor D-1500 Product Family LAN controller clears the *RSTI* bits in all the VF Mailbox registers and sets the *Reset Done (RSTD)* bits in all the VFMailbox registers.
 - The VF might read the *RSTD* bit to detect that a reset has occurred. The *RSTD* bit is cleared on read.

4.1.2.2 VF Receive Enable — PFVFRE / VF Transmit Enable — PFVFTE

This mechanism insures that a VF cannot transmit or receive before the Tx and Rx path have been initialized by the PF.

- The *PFVFRE* register contains a bit per VF. When the bit is set to 0b, assignment of an Rx packet for the VF's pool is disabled. When set to 1b, the assignment of an Rx packet for the VF's pool is enabled.
- The *PFVFTE* register contains a bit per VF. When the bit is set to 0b, fetching data for the VF's pool is disabled. When set to 1b, fetching data for the VF's pool is enabled. Fetching descriptors for the VF pool is maintained, up to the limit of the internal descriptor queues — regardless of *PFVFTE* settings.

PFVFTE and *PFVFRE* are initialized to zero (VF Tx and Rx traffic gated) following a PF reset. The relevant bits per VF are also initialized by a VF software reset or VFLR.

4.1.3 Firmware Reset

A firmware reset is caused either due to an internal fault in the firmware code or due to a Force TCO SMBus or NC-SI command with bit 2 set.

After a reset, firmware initiates an interrupt in the EICR.MNG field. The software device driver can detect a firmware reset by an increase in the value of the FWRESETCNT register.



If software detects a firmware reset in a 400 MHz Intel® Xeon® Processor D-1500 Product Family (as indicated by the FUSES0_GROUP[0] bit 5 being cleared), it should initiate a software reset as described in Section 4.1.1.5.1.

4.1.4 Reset Effects

The resets listed in Section 4.1.1 affect the following registers and logic:

Table 4.1. Reset Effects — Common Resets

Reset Activation	Power Good Reset	Integrated IO Reset	FW Reset	Force TCO	Notes
NVM read					
PCI configuration registers (RO)	X	X			8.
PCI configuration registers (RW)	X	X			8.
PCIe local registers	X				
Data path	X	X		X	2., 7.
MAC, TimeSync	X	X ^{6.}		X	
Integrated PHY (state only reset)	X	X ^{6.}		X	
Wake up (PM) Context	X	X ^{1.}			3.
Wake up/manageability control/status registers	X				4., 5.
Manageability unit	X		X		
LAN disable strapping pins	X	X			
All other strapping pins	X				
Shadow RAMs in MAC or PHY	X				

Table 4.2. Reset Effects — Per Function Resets


Reset Activation	D3 or Dr	FLR or PFLR	SW Reset	Exit from LAN Disable	Notes
NVM read					
PCI configuration registers (RO)					8.
PCI configuration registers (RW)		X			8., 9.
Data path and memory space, TimeSync	X	X	X	X	2., 7.
MAC	X ^{6.}	X ^{15.}	X ^{15.}	X	
Integrated PHY (state only reset)	X ^{6.}	X ^{15.}		X	
Virtual function resources	X	X	X		10.
Wake up (PM) context					3.
Wake up/manageability control/status registers					4., 5.
Manageability unit					
Shadow RAM					

**Table 4.3. Reset Effects -Virtual Function Resets**

Reset Activation	VFLR	VF SW Reset	Notes
Interrupt registers	X	X	11.
Queue disable	X	X	12.
VF specific PCIe configuration space	X		13.
Data path			
Statistics registers			14.

Table 4.1 Through Table 4.3 Notes:

- If $AUX_PWR = 0b$ the wake up context is reset (PME_Status and PME_En bits should be $0b$ at reset if the Intel® Xeon® Processor D-1500 Product Family LAN controller does not support PME from D3cold).
- The following register fields do not follow the general rules previously described:
 - ESDP registers — Reset on Power Good Reset only.
 - LED configuration registers — Reset on Power Good Reset and on software reset events.
 - The *Aux Power Detected* bit in the PCIe Device Status register is reset on Power Good Reset and Integrated I/O Reset only.
 - FLA — Reset on Power Good Reset only.
 - RAH/RAL[n, where $n > 0$], MTA[n], VFTA[n], WUPM[n], FFMT[n], FFVT[n], TDBAH/TDBAL, and RDBAH/RDVAL registers have no default value. If the functions associated with these registers are enabled they must be programmed by software. Once programmed, their value is preserved through all resets as long as power is applied.
 - Statistic registers (physical function)
- The wake up context is defined in the PCI Bus Power Management Interface specification (sticky bits). It includes:
 - PME_En bit of the Power Management Control/Status Register (PMCSR).
 - PME_Status bit of the PMCSR.
 - Aux_En bit in the PCIe registers.
 - The device requester ID (since it is required for PM_PME TLP).
The shadow copies of these bits in the Wake Up Control (WUC) register are treated identically.
- Refers to bits in the WUC register that are not part of the wake up context (the PME_En and PME_Status bits). The WUFC register is not part of the wake up context and is reset as part of the data path.
- The Wake Up Status (WUS) registers include the following:
 - WUS register.
 - Wake Up Packet Length (WUPL) register.
 - Wake Up Packet Memory (WUPM) register.
- The MAC cluster and the PHY are reset by the appropriate event only if the manageability unit is disabled and the host is in a low-power state with WoL disabled. WoL disabled means either AUX_PWR pin is cleared, or the *APM Enable* bit in NVM Control Word 3 is disabled, or ACPI is disabled (all wake up filters are disabled or PME_EN bit is disabled in PMCSR register).
- The contents of the following memories are cleared to support the requirements of PCIe FLR:
 - The Tx packet buffers.
 - The Rx packet buffers.



8. Sticky bits and hardware initialization bits (indicated as HwInit) in the PCI Configuration registers are cleared only by Power Good Reset reset.
9. The following register fields are not affected by FLR or PFLR:
 - Max_Payload_Size in the Device Control register
10. These registers include:
 - a. VFEICS.
 - b. VFEIMS.
 - c. VFEIAC.
 - d. VFEIAM.
 - e. VFEITR 0-2.
 - f. VFIVAR0.
 - g. VFIVAR_MISC.
 - h. VFPBACL.
 - i. VFMailbox.
11. These registers include:
 - a. VFEICS.
 - b. VFEIMS.
 - c. VFEIMC.
 - d. VFEIAC.
 - e. VFEIAM.
 - f. VFEICR.
 - g. EITR 0-2.
 - h. VFIVAR0.
 - i. VFIVAR_MISC.
 - j. VFPBACL.
 - k. VFMailbox.
 - l. VFMBMEM.
 - m. RSCINT
12. These registers include specific VF bits in the FVRE and FVTE registers are cleared as well.
13. These registers include:
 - a. MSI/MSI-X enable bits.
 - b. BME.
 - c. Error indications.
14. Rx and Tx counters might miss proper counting due to VFLR indicating more packets than those ones actually transferred. It could happen if VFLR happened after counting occurred but before Tx or Rx completed.
15. Will not be reset in case the *VETO* bit is asserted by manageability.

Note: Unless specified otherwise the Intel® Xeon® Processor D-1500 Product Family LAN controller's on-die memories are reset together with the functional block(s) they belong to.



4.2 Queue Disable

See [Section 4.5.7.1.1](#) for details on disabling and enabling an Rx queue.

See [Section 4.5.8.1.1](#) for details on disabling and enabling a Tx queue.

4.3 Function Disable

4.3.1 General

The Intel® Xeon® Processor D-1500 Product Family LAN controller provides support for BIOS to selectively enable or disable one or both LAN device(s) in the system.

4.3.2 Overview

Device presence (or non-presence) must be established early during BIOS execution, in order to ensure that BIOS resource-allocation (of interrupts, of memory or IO regions) is done according to devices that are present only. This is frequently accomplished using a BIOS Configuration Values Driven on Reset (CVDR) mechanism. The Intel® Xeon® Processor D-1500 Product Family LAN controller’s LAN-disable mechanism is implemented in order to be compatible with such a solution. These straps might be controlled by BIOS using a dedicated register in the SoC. Refer to the BIOS Writer’s Guide for more details.

LAN ports and/or PCIe functions can be disabled by:

- The LAN_DIS_n[x] straps (one strap per LAN port) are read by the Intel® Xeon® Processor D-1500 Product Family LAN controller on reset to determine the LAN enablement.
- One of the LAN ports can be disabled using a NVM configuration.

Disabling a LAN port affects the PCI function it resides on. When function 0 is disabled (either LAN0 or LAN1) then it does not disappear from the PCIe configuration space. Rather, the function presents itself as a dummy function. The device ID and class code of this function changes to other values (dummy function device ID, class code 0xFF0000). In addition, the function does not require any I/O space, and does not require an interrupt line. It requires a minimal memory space (4 KB) that is not mapped to internal registers.

Mapping between function and LAN ports is listed in the following table.

Table 4.4. PCI Functions Mapping

PCI Function #	LAN Function Select	Function 0	Function 1
Both LAN functions are enabled.	0	LAN 0	LAN 1
	1	LAN 1	LAN 0
LAN 0 is disabled.	0	Dummy	LAN1
	1	LAN 1	Disable
LAN 1 is disabled.	0	LAN 0	Disable
	1	Dummy	LAN 0
Both LAN functions are disabled.	Both PCI functions are disabled. Device is in low-power mode.		

The following rules apply to function disable:

- When function 0 is disabled, it is converted into a dummy PCI function. Function 1 is not affected.
- When function 1 is disabled, it disappears from the PCI configuration space.



- The disabled LAN port is still available for manageability purposes if disabled through the NVM mechanism.
- The disabled LAN port is still available for manageability purposes if disabled through the NVM. In this case, and if the *LPLU* bit is set, the PHY attempts to create a link at the lowest supported speed. The disabled LAN port is not available for manageability purposes if disabled through the *LAN_DIS_n[x]* strap mechanism.
- Dummy function mode should not be used in PCI IOV mode (since PF0 is required to support certain functionality).

The following NVM Control Word2 bits control function disable:

- One PCI function can be enabled or disabled according to the *LAN_PCI_DISABLE* field in the NVM (reflected in *DEV_FUNC_EN.LAN_PCI_DISABLE* bit).
- The *LAN Disable Select* field in the NVM (reflected in *DEV_FUNC_EN.LAN_DISABLE_SELECT* bit) indicates which function is disabled.
- The *LAN_FUNCTION_SEL* field (reflected in *FACTPS.LAN_FUNCTION_SEL* bit) defines the correspondence between LAN port and PCI function.

When a particular LAN is fully disabled, all internal clocks to that LAN are disabled, the device is held in reset, and the internal PHY for that LAN is powered down. In both modes, the device does not respond to PCI configuration cycles. Effectively, the LAN device becomes invisible to the system from both a configuration and power consumption standpoint.

4.3.3 Control Options

Note: Mapping LAN0 and LAN1 to PCI function 0 and PCI function 1 is controlled by the *FACTPS.LAN_FUNCTION_SEL* field.

LAN0 and LAN 1 can be disabled by BIOS by driving the *LAN_DIS_n[x]* pins respectively to low. These pins are strapping options, sampled at Power Good Reset or at integrated I/O reset.

Note: Mapping LAN0 and LAN1 to PCI function 0 and PCI function 1 is controlled by the *FACTPS.LAN_FUNCTION_SEL* field.

4.3.4 Event Flow for Enable/Disable Functions

This section describes the driving levels and event sequence for device functionality.

Following a Power Good Reset or integrated I/O reset events the *LAN_DIS_n[x]* signals should be driven high for normal operation. If any of the LAN functions are not required statically, its associated disable strapping pin can be configured statically to low.

4.3.4.1 BIOS Disable the LAN Function at Boot Time by the Using Strapping Option

Assume that following a power-up sequence *LAN_DIS_n[x]* signals are driven high.

1. Integrated I/O is established following an Integrated I/O reset.
2. BIOS recognizes that a LAN function in the Intel® Xeon® Processor D-1500 Product Family LAN controller should be disabled.
3. The BIOS drives the *LAN_DIS_n[x]* signal to a low level.
4. BIOS issues an Integrated I/O reset.
5. As a result, the Intel® Xeon® Processor D-1500 Product Family LAN controller samples the *LAN_DIS_n[x]* signals and disables the LAN function and issues an internal reset to this function.



6. The BIOS might start with the device enumeration procedure (the disabled LAN function is invisible; changed to dummy function).
7. Proceed with normal operation.
8. Re-enable could be done by driving the LAN_DIS_n[x] signal high and then requesting the end user to issue a warm boot to initialize new bus enumeration.

4.3.4.2 Multi-Function Advertisement

If one of the LAN devices is disabled and function 0 is the only active function, the Intel® Xeon® Processor D-1500 Product Family LAN controller is no longer a multi-function device. The Intel® Xeon® Processor D-1500 Product Family LAN controller normally reports 0x80 in the *PCI Configuration Header* field (*Header Type*), indicating multi-function capability. However, if a LAN is disabled and only function 0 is active, the Intel® Xeon® Processor D-1500 Product Family LAN controller reports 0x0 in this field to signify single-function capability.

4.3.4.3 Interrupt Use

When both LAN devices are enabled, the Intel® Xeon® Processor D-1500 Product Family LAN controller uses the PCI legacy interrupts of both ports for interrupt reporting. The NVM configuration controls the *Interrupt Pin* field of the PCI configuration header to be advertised for each LAN device to comply with PCI specification requirements.

However, if either LAN device is disabled, then the legacy PCI interrupt of port A must be used for the remaining LAN device, therefore the NVM configuration must be set accordingly. Under these circumstances, the *Interrupt Pin* field of the PCI header always reports a value of 0x1, indicating INTA# pin usage, which means legacy PCI interrupt of port A is used.

4.3.4.4 Power Reporting


When both LAN devices are enabled, the PCI Power Management register block has the capability of reporting a common power value. The common power value is reflected in the *Data* field of the PCI Power Management registers. The value reported as common power is specified via an NVM field, and is reflected in the *Data* field each time the *Data_Select* field has a value of 0x8 (0x8 = common power value select).

When only one LAN port is enabled and the Intel® Xeon® Processor D-1500 Product Family LAN controller appears as a single-function device, the common power value, if selected, reports 0x0 (undefined value), as common power is undefined for a single-function device.

4.4 Device Disable

4.4.1 Overview

In order to disable the device, both LAN_DIS_n[x] signals should be configured statically to low. When sampled at a Power Good Reset or integrated I/O reset events, the Intel® Xeon® Processor D-1500 Product Family LAN controller is disabled. It is held in reset and power-down mode (some clocks are still running), and digital I/O pins are at High-Z if *DEV_FUNC_EN.DEV_OFF_EN* bit was set in NVM control word 1. As an example, digital I/O pins are in an electrical off state where pull-up/pull-down resistors are at their defined values. The manageability interface is also disabled in this state.



4.4.2 BIOS Disable of the Device at Boot Time by Using the Strapping Option

Assume that following a power-up sequence LAN_DIS_n[x] signals are driven high.

1. Integrated I/O is established following an Integrated I/O reset.
2. BIOS recognizes that the Intel® Xeon® Processor D-1500 Product Family LAN controller should be disabled.
3. The BIOS drives the LAN_DIS_n[x] signals to the low level.
4. BIOS issues an Integrated I/O reset.
5. As a result, the Intel® Xeon® Processor D-1500 Product Family LAN controller samples the LAN_DIS_n[x] signals and disables the LAN ports and the Integrated I/O connection.
6. Re-enable can be done by driving high at least one of the LAN_DIS_n[x] signals and then issuing an Integrated I/O reset to restart the device.

4.5 Software Initialization and Diagnostics

4.5.1 Introduction

This section discusses general software notes for the Intel® Xeon® Processor D-1500 Product Family LAN controller, especially initialization steps. These include:

- General hardware power-up state
- Basic device configuration
- Interrupts initialization
- Initialization of transmit
- Receive initialization
- Link configuration
- Software reset capability
- Statistics
- Diagnostic hints
- Virtualization support initialization
- Security (IPSec) support initialization

4.5.2 Power-Up State

When the Intel® Xeon® Processor D-1500 Product Family LAN controller powers up, it automatically reads the NVM. The NVM contains sufficient information to bring the link up and configure the Intel® Xeon® Processor D-1500 Product Family LAN controller for manageability and/or APM wake up. However, software initialization is required for normal operation.

4.5.3 Initialization Sequence

The following sequence of commands is typically issued to the device by the software device driver in order to initialize the Intel® Xeon® Processor D-1500 Product Family LAN controller to normal operation. The major initialization steps are:

1. Disable interrupts.
2. Issue a Power Good Reset and perform general configuration (see [Section 4.5.3.2](#)).



3. Wait for the NVM auto-read completion.
4. Wait for manageability configuration done indication (EEMNGCTL.CFG_DONE0/1).
5. Wait until the DMA initialization completes (RDRXCTL.DMAIDONE).
6. Setup the PHY and the link — see [Appendix B](#).
7. Initialize all statistical counters — see [Section 4.5.5](#).
8. Initialize receive — see [Section 4.5.7](#).
9. Initialize transmit — see [Section 4.5.8](#).
10. Initialize Virtualization support — see [Section 4.5.9](#)
11. Configure Security (IPSec and MACSec) — see [Section 4.5.10](#)
12. Enable interrupts — see [Section 4.5.3.1](#).

4.5.3.1 Interrupts During Initialization

Most software device drivers disable interrupts during initialization to prevent re-entrance. Interrupts are disabled by writing to the EIMC register. Note that the interrupts need to also be disabled after issuing a Power Good Reset reset, so a typical driver initialization flow is:

1. Disable interrupts.
2. Issue a Power Good Reset reset.
3. Disable interrupts (again).

After initialization completes, a typical software device driver enables the desired interrupts by writing to the IMS register.

4.5.3.2 Global Reset and General Configuration

Global reset = software reset + link reset.


Device initialization typically starts with a software reset that puts the device into a known state and enables the software device driver to continue the initialization sequence. Following a global reset the software device driver should poll the CTRL.RST until it is cleared and then wait at least 10 ms to enable a smooth initialization flow.

To enable flow control, program the FCTTV, FCRTL, FCRTH, FCRTV and FCCFG registers. If flow control is not enabled, these registers should be written with 0x0. If Tx flow control is enabled, then Tx CRC by hardware should be enabled as well (HLREG0.TXCRCEN = 1b).

The link inter-connect configuration according to the electrical specification of the relevant electrical interface should be set prior to the link setup. This configuration is done through the PHY image section of the NVM by applying the appropriate settings to the link interconnect block.

4.5.4 Link Initialization

Refer to [Section B.4.1](#) for the initialization and link setup steps.



4.5.5 Initialization of Statistics

Statistics registers are hardware-initialized to values as detailed in each particular register's description. The initialization of these registers begins upon transition to a D0 active power state (when internal registers become accessible, as enabled by setting the *Memory Access Enable* field of the PCIe Command register), and is guaranteed to complete within 1 ms of this transition. Note that access to statistics registers prior to this interval might return indeterminate values.

All statistical counters are cleared on read and a typical software device driver reads them (thus making them zero) as a part of the initialization sequence.

Queue counters are mapped using the RQSMR registers for Rx queues, and TQSM registers for Tx queues. Refer to the RQSMR register section for RQSMR setup, and to the TQSM register section for TQSM setup. Note that if software requires the queue counters, the RQSMR and TQSM registers must be reprogrammed following a device reset.

4.5.6 Interrupt Initialization

4.5.6.1 Working with Legacy or MSI Interrupts

- Software device drivers associate between Tx and Rx interrupt causes and the EICR register by setting the IVAR[n] registers
- Program the SRRCTL[n].RDMTS per receive queue if software uses the Receive Descriptor Minimum Threshold Interrupt (RDMTI).
- All interrupts should be set to zero — no auto clear in the EIAC register. Following an interrupt software can read the EICR register to check for the interrupt causes.
- Set the auto mask in the EIAM register according to the preferred mode of operation.
- Set the interrupt throttling in the EITR[n] and GPIE according to the preferred mode of operation.
- Software enables the required interrupt causes by setting the EIMS register.

4.5.6.2 Operating with MSI-X

- The operating system/BIOS sets hardware to MSI-X mode and programs the MSI-X table as part of the device enumeration procedure.
- The software device driver associates between interrupt causes and MSI-X vectors and the throttling timers EITR[n] by programming the IVAR[n] and IVAR_MISC registers.
- Program the SRRCTL[n].RDMTS (per receive queue) if software uses the receive descriptor minimum threshold interrupt.
- The EIAC[n] registers should be set to auto clear for transmit and receive interrupt causes (for best performance). The EIAC bits that control the other and TCP timer interrupt causes should be set to 0b — no auto clear.
- Set auto mask in the EIAM and EIAM[n] registers according to the preferred mode of operation.
- Set the interrupt throttling in the EITR[n] and GPIE registers according to the preferred mode of operation.
- Software enables the required interrupt causes by setting the EIMS[n] registers.

4.5.7 Receive Initialization

Initialize the following register tables before receive and transmit is enabled:

- Set *CTRL_EXT.Extended VLAN* bit if needed



- Receive Address (*RAL[n]* and *RAH[n]*) for used addresses and Receive Address High — *RAH[n].VAL=0b* for unused addresses
- Unicast Table Array — *PFUTA*
- VLAN Filter Table Array — *VFTA[n]*
- VLAN Pool Filter — *PFVLVF[n]*
- MAC Pool Select Array — *MPSAR[n]*
- VLAN Pool Filter Bitmap — *PFVLVFB[n]*.

Program the Receive Address register(s) (*RAL[n]*, *RAH[n]*) per the station address. This can come from the NVM or from any other means (for example, it could be stored anywhere in the NVM or even in the platform PROM for a LOM design).

Set up the Multicast Table Array — *MTA* registers. Assuming the entire table was zeroed by the last reset, only the desired multicast addresses should be permitted (by writing 0x1 to the corresponding bit location). Set the *MCSTCTRL.MFE* bit if multicast filtering is required.

Set up the VLAN Filter Table Array — *VFTA* if VLAN support is required. Assuming the entire table was zeroed by the last reset, only the desired VLAN addresses should be permitted (by writing 0x1 to the corresponding bit location). Set the *VLNCTRL.VFE* bit if VLAN filtering is required.

Initialize the flexible filters 0...7 — Flexible Host Filter Table (*FHFT_FILTER*) registers.

After all memories in the filter units previously indicated are initialized, enable ECC reporting by setting the *RXFECERR0.ECCFLT_EN* bit.

Program the different Rx filters and Rx off loads via registers *FCTRL*, *VLNCTRL*, *MCSTCTRL*, *RXCSUM*, *RQTC*, *RFCTL*, *MPSAR*, *RSSRK*, *RETA*, *SAQF*, *DAQF*, *SDPQF*, *FTQF*, *SYNQF*, *ETQF*, *ETQS*, *RDRXCTL*, and *RSCDBU*.

Note: Because NFS detection is not supported, the *RFCTL.NFSW_DIS* and *RFCTL.NFSR_DIS* bits should be set to 1b.

Program *RXPBSIZE*, *MRQC*, *PFQDE*, *RTRUP2TC*, and *MFLCN.RFCE* registers according to virtualization mode.


Enable receive jumbo frames by setting *HLREGO.JUMBOEN* in the following case:

1. Jumbo packets are expected. Set *MAXFRS.MFS* to the expected maximum packet size.
2. Enable receive coalescing if required as described in [Section 4.5.7.2](#).

4.5.7.1 Receive Queues Enable

The following should be done for each receive queue:

1. Allocate a region of memory for the receive descriptor list.
2. Receive buffers of appropriate size should be allocated and pointers to these buffers should be stored in the descriptor ring.
3. Program the descriptor base address with the address of the region (registers *RDBAH* and *RDBAL*).
4. Set the length register to the size of the descriptor ring (register *RDLEN*).
5. Program *SRRCTL* associated with this queue according to the size of the buffers and the required header control.
6. Set the *RXDCTL[n].RLPML* field enabled by the *RXDCTL[n].RLPML_EN* limiting the maximum Rx packet size. This setting is optional enabling the software to use smaller buffers than the size



defined by the *SRRCTL[n].BSIZEPACKET*. Software might not use smaller buffers than defined by the *SRRCTL[n]* on Rx queues that enables RSC.

7. If header split is required for this queue, program the appropriate PSRTYPE for the appropriate headers.
8. Program RSC mode for the queue via the RSCCTL register.
9. Program RXDCTL with appropriate values including the queue *Enable* bit. Note that packets directed to a disabled queue are dropped.
10. Poll the RXDCTL register until the *Enable* bit is set. The tail should not be bumped before this bit was read as 1b.
11. Bump the tail pointer (RDT) to enable descriptors fetching by setting it to the ring length minus one.

Enable the receive path by setting *RXCTRL.RXEN*. This should be done only after all other settings are done. If software uses the receive descriptor minimum threshold interrupt, that value should be set.

4.5.7.1.1 Dynamic Enabling and Disabling of Receive Queues

Receive queues can be enabled or disabled dynamically using the following procedure.

4.5.7.1.1.1 Enabling

- Follow the per-queue initialization described in the previous section.

4.5.7.1.1.2 Disabling


- Disable the routing of packets to this queue by re-configuring the Rx filters.
- If RSC is enabled on the specific queue and VLAN strip is enabled as well then wait two ITR expiration times (ensure all open RSCs completed).
- Disable the queue by clearing the *RXDCTL.ENABLE* bit. The Intel® Xeon® Processor D-1500 Product Family LAN controller stops fetching and writing back descriptors from this queue. Any further packet that is directed to this queue is dropped. If a packet is being processed, the Intel® Xeon® Processor D-1500 Product Family LAN controller completes the current buffer write. If the packet spreads over more than one data buffer, all subsequent buffers are not written.
- The Intel® Xeon® Processor D-1500 Product Family LAN controller clears the *RXDCTL.ENABLE* bit only after all pending memory accesses to the descriptor ring are done. The software device driver should poll this bit before releasing the memory allocated to this queue.
- Once the *RXDCTL.ENABLE* bit is cleared the software device driver should wait an additional amount of time (~100 µs) before releasing the memory allocated to this queue.

The Rx path can be disabled only after all the receive queues are disabled.

Note: As there could be additional packets in the receive packet buffer targeted to the disabled queue and the arbitration could be such that it would take a long time to drain these packets, if software re-enables a queue before all packets to that queue were drained, the enabled queue could potentially get packets directed to the old configuration of the queue. For example, VM goes down and a different VM gets the queue. The software device driver should delay the re-enablement of the queue until it is guaranteed there are no more packets directed to this queue in the packet buffer.

4.5.7.2 RSC Enablement

RSC enablement as well as RSC parameter settings are assumed as static. It should be enabled prior to receiving and can be disabled only after the relevant Rx queue(s) are disabled.



4.5.7.2.1 RSC Global Setting

- Enable global CRC stripping via the HLREG0 register (hardware default setting).
- Software device driver should set the *RDRXCTL.RSCACKC* bit that forces RSC completion on any change of the *ACK* bit in the Rx packet relative to the RSC context.
- The *SRRCTL[n].BSIZEHEADER* (header buffer size) bit must be larger than the packet header (even if header split is not enabled). A minimum size of 128 bytes for the header buffer addresses this requirement.

4.5.7.2.2 RSC Per Queue Setting

- Enable RSC and configure the maximum allowed descriptors per RSC by setting the *MAXDESC* and *RSCEN* fields in the *RSCCTL[n]* register.
- Use a non-legacy descriptor type by setting the *SRRCTL[n].DESCTYPE* bit to non-zero values.
- TCP header recognition: the *PSR_type4* bit in the *PSRTYPE[n]* registers should be set.
- The *SRRCTL[n].BSIZEPACKET* (packet buffer size) must be 2 KB at minimum.
- Interrupt setting:
 - Interrupt moderation must be enabled by setting the *EITR[n].ITR_INTERVAL* bit to a value greater than zero. The *ITR Interval* bit must be larger than the *RSC Delay* field described later. Note that if the *CNT_WDIS* bit is cleared (write enable), the *ITR Counter* bit should be set to 0b.
 - The *RSC_DELAY* field in the *GPIE* register should be set to the expected system latency descriptor write-back cycles. 4 to 8 μ s should be sufficient in most cases. If software sees cases where RSC did not complete as expected (following EITR interrupt assertion), then the *RSC Delay* field might need to be increased.
 - Map the relevant Rx queues to an interrupt by setting the relevant *IVAR* registers.

4.5.7.3 Flow Director Initialization

Flow Director initialization is described in [Section 7.1.3.5.11](#).

4.5.8 Transmit Initialization

- Program the HLREG0 register according to the MAC behavior needed.
- Program the TCP segmentation parameters via registers *DMATXCTL* (while keeping the *TE* bit cleared), *DTXTCPFLGL*, *DTXTCPFLGH*, and *TPH* parameters via *TPH_TXCTRL*.
- Program the *DTXMXSZRQ*, *TXPBSIZE*, *TXPBTHRESH*, *MTQC*, and *MNGTXMAP* registers according to virtualization mode.
- Legacy software device drivers that uses queue zero and assumes it is enabled by default, should at this stage, set queue zero parameters as described in the section that follows.
- Enable the transmit path by setting the *DMATXCTL.TE* bit.

4.5.8.1 Transmit Queues Enable

The following steps should be done once for each transmit queue:

1. Allocate a region of memory for the transmit descriptor list.
2. Program the descriptor base address with the address of the region (*TDBAL* and *TDBAH*).
3. Set the length register to the size of the descriptor ring (*TDLEN*).
4. If needed, set *TDWBAL/TWDBAH* to enable head write back.
5. Program the *TXDCTL* register with the desired TX descriptor write-back policy (see the recommended values in the register description).



6. Enable the queue using the `TXDCTL.ENABLE` bit. Poll the `TXDCTL` register until the `Enable` bit is set.

Note: Queue 0 is enabled by default when the `DMATXCTL.TE` bit is set. If transmit is already enabled (`DMATXCTL.TE` bit is set), this queue should be disabled (clear `TXDCTL.ENABLE` bit) before changing the basic queue parameters (`TDBAL`, `TDBAH`, `TDLEN`, `TDWBAL/TWDBAH`).

Note: The tail register of the queue (`TDT`) should not be bumped until the queue is enabled.

4.5.8.1.1 Dynamic Enabling and Disabling of Transmit Queues

Transmit queues can be enabled or disabled dynamically given the following procedure is followed.

Enabling:

Follow the per-queue initialization described in the previous section.

Disabling:

1. Stop storing packets for transmission in this queue.
2. The completion of the last transmit descriptor must be visible to software in order to guarantee that packets are not lost in step 5 (Section 4.5.8). Therefore, its `RS` bit must be set or `WTHRESH` must be greater than zero. If none of the previous conditions are met, software should add a null Tx data descriptor with an active `RS` bit.
3. Wait until the software head of the queue (`TDH`) equals the software tail (`TDT`) indicating the queue is empty.
4. Wait until all descriptors are written back (polling the `DD` bit in ring or polling the `Head_WB` content). It might be required to flush the transmit queue by setting the `TXDCTL[n].SWFLSH` bit if the `RS` bit in the last fetched descriptor is not set or if `WTHRESH` is greater than zero.
5. Disable the queue by clearing `TXDCTL.ENABLE`.
6. Any packets waiting for transmission in the packet buffer would still be sent at a later time.

The transmit path can be disabled only after all transmit queues are disabled.

4.5.9 Virtualization Initialization Flow

4.5.9.1 VMDq Mode

4.5.9.1.1 Global Filtering and Offload Capabilities

- Select one of the VMDQ pooling and queuing methods
 - Pool selection can be based on MAC/VLAN, and optionally E-tag. The filtering mode is defined using the `PFVTCTL.POOLING_MODE` field.
- A queue in a poll can be based on RSS. Potential values for `MRQC.Multiple Receive Queues Enable` to set VMDq modes are 1000b to 1111b.
- Configure the `PFVTCTL` register to define the default pool.
- Enable replication via the `PFVTCTL.RpL_En` bit.
- If needed, enable padding of small packets via the `HLREGO.TXPADEN` bit.
- The `MPSAR` registers are used to associate Ethernet MAC addresses to pools. Using the `MPSAR` registers, software must reprogram `RAL[0]` and `RAH[0]` by their values. Note that software could read these registers and then write them back with the same content.



4.5.9.1.2 Mirroring Rules

For each mirroring rule to be activated:

- Set the type of traffic to be mirrored in the `PFMRCTL[n]` register. Only one type of traffic can be selected in each rule.
- Set the mirror pool by setting the `PFMRCTL[n].MP` bit.
- For pool mirroring, set the `PFMRVM[n]` register with the pools to be mirrored.
- For VLAN mirroring, set the `PFMRVLAN[n]` register with the indexes from the `PFVLVF` registers of the VLANs to be mirrored.

4.5.9.1.3 Security Features


For each pool, the software device driver might activate the MAC, VLAN and Ethertype anti-spoof features via the relevant bit in the `PFVFSPOOF.MACAS`, `PFVFSPOOF.VLANAS` and `PFVFSPOOF.ETHERTYPEAS/ETHERTYPELB` fields, respectively.

In addition, stripping and hiding of VLAN and external tag (E-tag) might also be requested via the `PFQDE.HIDE_VLAN` and `PFQDE.STRIP_TAG` flags, respectively.

4.5.9.1.4 Per-pool Settings

As soon as a pool of queues is associated to a VM, software should set the following parameters:

- Associate the unicast Ethernet MAC address of the VM by enabling the pool in the MPSAR registers.
- If all the Ethernet MAC addresses are used, the Unicast Hash Table (PFUTA) can be used. Pools servicing VMs whose address is in the hash table should be declared as so by setting the `PFVML2FLT.ROPE` bit. Packets received according to this method didn't pass perfect filtering and are indicated as such.
- Enable the pool in all RAH/RAL registers representing the multicast Ethernet MAC addresses this VM belongs to.
- If all the Ethernet MAC addresses are used, the Multicast Hash Table (MTA) can be used. Pools servicing VMs using multicast addresses in the hash table should be declared as so by setting the `PFVML2FLT.ROMPE` bit. Packets received according to this method didn't pass perfect filtering and are indicated as such.
- Define whether this VM should get all multicast/broadcast packets in the same VLAN via `PFVML2FLT.MPE` and `PFVML2FLT.BAM`, and whether it should accept un-tagged packets via `PFVML2FLT.AUPE`.
- Enable the pool in each of the `PFVLVF` and `PFVLVFB` registers this VM belongs to.
- A VM might be set to receive its own traffic in case the source and the destination are in the same pool via the `PFVMTXSW.LLE` bit.
- Whether VLAN header and CRC should be stripped from the packet. Note that even if the CRC is kept, it might not match the actual content of the forwarded packet, because of other offloading applications such as VLAN strip.
 - A striped VLAN might also be hidden from the VM.
- Set which header split is required via the `PSRTYPE` register.
- In RSS mode, define if the pool uses RSS via the proper `MRQC.MRQE` mode.
 - Enable the pool in the `PFVFRE` register to enable Rx filtering.
 - To enable multiple Tx queues, Set MTQC as described in [Section 7.2.1.2.1](#).
 - Enable the pool in the `PFVFTE` register to enable Tx filtering.
- Enable Rx and Tx queues as described in [Section 4.5.7](#) and [Section 4.5.8](#).



- For each Rx queue a drop/no drop flag can be set in *SRRCTL.DROP_EN* and via the PFQDE register, controlling the behavior in cases no receive buffers are available in the queue to receive packets. The usual behavior is to enable a drop in order to avoid head of line blocking. Setting the PFQDE (per queue) is done by using the *Queue Index* field in the PFQDE register.

4.5.9.2 IOV Initialization

4.5.9.2.1 PF Driver Initialization

The PF driver is responsible for the link setup and handling of all the filtering and offload capabilities for all the VFs as described in [Section 4.5.9.1.1](#) and the security features as described in [Section 4.5.9.1.3](#). It should also set the bandwidth allocation per transmit queue for each VF as described in [Section 4.5.9](#).

Note: Link setup might include the authentication process (802.1X or other).

After all the common parameters are set, the PF driver should set all the VFMailbox[n].RSTD bits by setting the CTRL_EXT.PFRSTD bit.

PF enables VF traffic via the PFVFTE and PFVFRE registers after all VF parameters are set as defined in [Section 4.5.9.1.4](#).

Note: If the operating system changes the NumVF setting in the PCIe SR-IOV Num VFs register after the device was active, it is required to initiate a PF software reset following this change.

4.5.9.2.1.1 VF Specific Reset Coordination

After the PF driver receives an indication of a VF FLR via the PFVFLREC register, it should enable the receive and transmit for the VF only once the device is programmed with the right parameters as defined in [Section 4.5.9.1.4](#). The receive filtering is enabled using the PFVFRE register and the transmit filtering is enabled via the PFVFTE register.

Note: The filtering and offloads setup might be based on a central IT settings or on requests from the VF drivers.

4.5.9.2.2 VF Driver Initialization

At initialization, after the PF indicated that the global initialization was done via the VFMailbox.RSTD bit, the VF driver should communicate with the PF, either via the mailbox or via other software mechanisms to assure that the right parameters of the VF are programmed as described in [Section 4.5.9.1.4](#). The PF driver might then send an acknowledge message with the actual setup done according to the VF request and the IT policy.

The VF driver should then setup the interrupts and the queues as described in [Section 4.5.6](#), [Section 4.5.7](#), and [Section 4.5.8](#).

4.5.9.2.3 Full Reset Coordination

A mechanism is provided to synchronize reset procedures between the PF and the VFs. It is provided specifically for PF software reset but can be used in other reset cases as described later in this section.

The procedure is as follows:

One of the following reset cases takes place:

- Power Good Reset



- Integrated I/O reset
- D3hot-to-D0 transition
- FLR
- Software reset by the PF

The Intel® Xeon® Processor D-1500 Product Family LAN controller sets the *RSTI* bits in all VFMailbox registers. Once the reset completes, each VF might read its VFMailbox register to identify a reset in progress.

Once the PF completes configuring the device, it clears the CTRL_EXT.PFRSTD bit. As a result, the Intel® Xeon® Processor D-1500 Product Family LAN controller clears the *RSTI* bits in all VFMailbox registers and sets the *Reset Done* (RSTD) bits are set in all VFMailbox registers.

Until a RSTD condition is detected, the VFs should only access the VFMailbox register and should not attempt to activate the interrupt mechanism or the transmit and receive process.

4.5.10 Alternate MAC Address Support

In some systems, the MAC address used by a port needs to be replaced with a temporary MAC address in a way that is transparent to the software layer. One possible usage is in blade systems, to allow a standby blade to use the MAC address of another blade that failed, so that the network image of the entire blade system does not change.

In order to allow this mode, a management console might change the MAC address in the NVM image. It is important in this case to be able to keep the original MAC address of the device as programmed at the factory.

In order to support this mode, the Intel® Xeon® Processor D-1500 Product Family LAN controller provides the Alternate Ethernet MAC Address structure in the NVM to store the original MAC addresses. This structure is described in [Section 6.2.2](#).

In some systems, it might be advantageous to restore the original MAC address at power on reset, to avoid conflicts where two network controllers would have the same MAC address.


The Intel® Xeon® Processor D-1500 Product Family LAN controller supports replacement of the MAC address with an alternate MAC address via the NC-SI interface using the Update System MAC address Intel OEM Command or via a BIOS CLP interface as described in the BIOS CLP document

4.5.10.1 LAN MAC Address Restore

The Intel® Xeon® Processor D-1500 Product Family LAN controller restores the LAN MAC addresses stored in the Alternate Ethernet MAC Address structure to the regular LAN MAC address location (see [Section 6.4.6.2](#)) if the following conditions are met:

1. The *restore MAC address* bit in the *Common Firmware Parameters* NVM word is set.
2. The value in word 0x37 is not 0xFFFF.
3. The MAC address set in the regular MAC address location is different than the address stored in the Alternate Ethernet MAC Address structure.
4. The addresses stored in the Alternate Ethernet MAC Address structure are valid (not all zeros or all ones).

If the value at word 0x37 is valid, but the MAC addresses in the Alternate MAC structure are not valid (0xFFFFFFFF or all zeros) and *restore MAC address* NVM bit is set, the regular MAC address is backed up in the Alternate MAC structure.



4.5.10.2 Restore Reporting

If after the power up cycle, the LAN address is equal to the LAN Factory MAC addresses (either they were restored by the internal firmware or they were originally equal), the *FWSM.Factory MAC address restored* bit is set. This bit is common to all ports.

4.6 Access to Shared Resources

Part of the resources in the Intel® Xeon® Processor D-1500 Product Family LAN controller are shared between several software entities — namely the software device drivers of the two ports and internal firmware. In order to avoid contentions, a software device driver that needs to access one of these resources should use the *Gaining Control of Shared Resource by Software* flow in order to acquire ownership of this resource, and use the *Releasing a Shared Resource by Software* flow described there in order to relinquish ownership of this resource.

The shared resources are:

- NVM
- PHY 0 and PHY 1 registers
- MAC CSRs
- I²C interface of each port

§



5.0 Power Management and Delivery

This section defines how power management is implemented in the Intel® Xeon® Processor D-1500 Product Family LAN controller.

5.1 Power Management

5.1.1 Introduction to the Intel® Xeon® Processor D-1500 Product Family LAN Controller Power States

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports the D0 and D3 power states defined in the PCI Power Management and PCIe specifications. D0 is divided into two sub-states: D0u (D0 un-initialized), and D0a (D0 active). In addition, the Intel® Xeon® Processor D-1500 Product Family LAN controller supports a Dr state that is entered when its Integrated I/O Reset signals are asserted (including the D3cold state).

Figure 5.1 shows the power states and transitions between them.

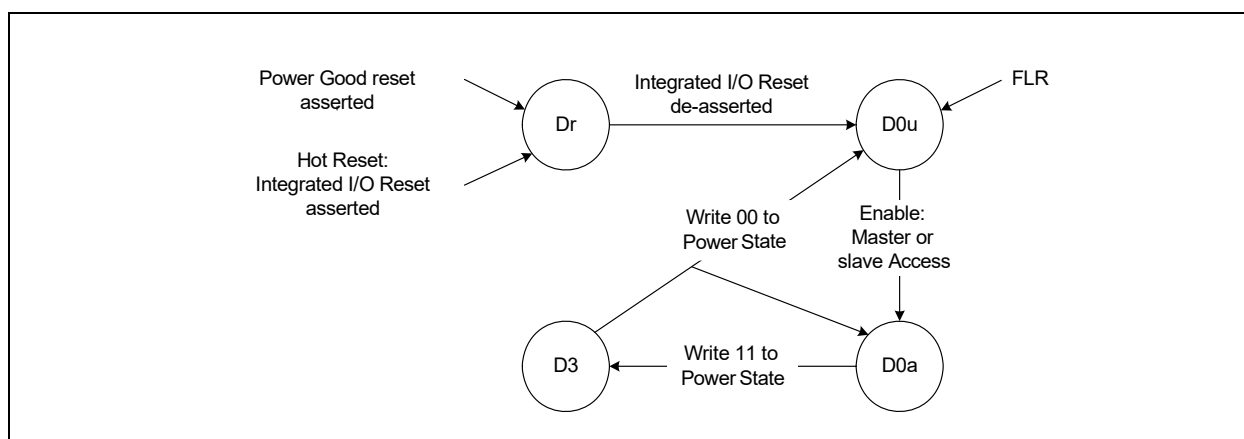



Figure 5.1. Power Management State Diagram

5.1.2 Auxiliary Power Usage

The Intel® Xeon® Processor D-1500 Product Family LAN controller uses the *AUX_PWR* indication that auxiliary power is available to the controller, and therefore advertises D3cold wake up support in the *PMC.PME_Support* field and sets the *Aux_Power_Detected* bit in the PCIe capability structure *Device Status Register*. The *AUX_PWR* pin is a soft strap that is loaded at Power On Reset (POR). The amount of power required for the function, which includes the entire Network Interface Card (NIC), is advertised in the Power Management Data register, which is loaded from the NVM.

Note: In the Intel® Xeon® Processor D-1500 Product Family LAN controller *AUX_PWR* is tied to a fixed value of 1.



The only effect of setting *AUX_PWR* to 1b is advertising D3cold wake up support and changing the reset function of *PME_En* and *PME_Status*. *AUX_PWR* is a strapping option in the Intel® Xeon® Processor D-1500 Product Family LAN controller. If D3cold is supported, the *PME_En* and *PME_Status* bits of the Power Management Control/Status Register (PMCSR), as well as their shadow bits in the Wake Up Control Register (WUC) are reset only by the Power Good Reset. The Intel® Xeon® Processor D-1500 Product Family LAN controller tracks the *PME_En* bit of the PMCSR and the Auxiliary (AUX) power *PM Enable* bit of the PCIe Device Control register to determine the power it can consume (and therefore its power state) in the D3cold state (internal Dr state).

According to the following settings, the Intel® Xeon® Processor D-1500 Product Family LAN controller might consume higher auxiliary power than is allowed by PCIe specifications:

- If the *AUX Power PM Enable* bit of the PCIe Device Control register is set, the Intel® Xeon® Processor D-1500 Product Family LAN controller might consume higher power for any purpose (even if *PME_En* is not set).
- If the *AUX Power PM Enable* bit of the PCIe Device Control register is cleared, higher power consumption is determined by the PCI-PM legacy *PME_En* bit of the PMCSR.

5.1.3 Power States

5.1.3.1 D0uninitialized State

The D0u state is a low-power state used after the integrated I/O interface is de-asserted following power-up (cold or warm) or on D3 exit.

When entering D0u, the Intel® Xeon® Processor D-1500 Product Family LAN controller disables wake ups.

5.1.3.1.1 Entry to a D0u State

D0u is reached from either the Dr state (on de-assertion of integrated I/O interface) or the D3hot state (by configuration software writing a value of 00b to the *Power State* field of the PCI PM registers).

De-asserting the integrated I/O interface means that the entire state of the device is cleared, other than sticky bits. State is re-loaded from the NVM. Once this is done, configuration software can access the device.

On a transition from D3-to-D0u state, the Intel® Xeon® Processor D-1500 Product Family LAN controller requires that software perform a full re-initialization of the function excluding its PCI configuration space which is kept while in D3.

5.1.3.2 D0active State

Once memory space is enabled, the Intel® Xeon® Processor D-1500 Product Family LAN controller enters an active state. It can transmit and receive packets if properly configured by the software device driver. The PHY is enabled or re-enabled by the software device driver to operate/auto-negotiate to full line speed/power if not already operating at full capability. Any APM wake up previously active remains active. The software device driver can deactivate ACPI wake up by writing to the Wake Up Control (WUC) register and APM wake up by writing to the General Receive Control (GRC) register, or activate other wake up filters by writing to the Wake Up Filter Control (WUFC) register.



5.1.3.2.1 Entry to D0a State

D0a is entered from the D0u state by writing a 1b to the *Memory Access Enable* or the *I/O Access Enable* bit of the PCI Command register. The DMA, MAC, and PHY of the appropriate LAN function are enabled.

5.1.3.3 D3 State (PCI-PM D3hot)

The Intel® Xeon® Processor D-1500 Product Family LAN controller transitions to D3 when the system writes a 11b to the *Power State* field of the PMCSR. Any wake-up filter settings that were enabled before entering this reset state are maintained. Upon transition to D3 state, the Intel® Xeon® Processor D-1500 Product Family LAN controller clears the *Memory Access Enable* and *I/O Access Enable* bits of the PCI Command register, which disables memory access decode. In D3, the Intel® Xeon® Processor D-1500 Product Family LAN controller only responds to PCI configuration accesses and does not generate master cycles.

Configuration and message requests are the only PCIe TLPs accepted by a function in the D3hot state. All other received requests must be handled as unsupported requests, and all received completions can optionally be handled as unexpected completions. If an error caused by a received TLP (like an unsupported request) is detected while in D3hot, and reporting is enabled an error message must be sent (See Section 5.3.1.4.1 in the PCIe Base Specification).

A D3 state is followed by either a D0u state (in preparation for a D0a state) or by a transition to Dr state (PCI-PM D3cold state). To transition back to D0u, the system writes a 00b to the *Power State* field of the PMCSR. Transition to Dr state is through the integrated I/O interface assertion.

5.1.3.3.1 Entry to D3 State

Transition to D3 state is through a configuration write to the *Power State* field of the PCI-PM registers.

Prior to transitioning from D0 to the D3 state, the software device driver disables scheduling of further tasks to the Intel® Xeon® Processor D-1500 Product Family LAN controller; it masks all interrupts, it does not write to the Transmit Descriptor Tail (TDT) register or to the Receive Descriptor Tail (RDT) register and operates the master disable algorithm as defined in Section 5.1.3.3.2. If wake-up capability is needed, the software device driver should set up the appropriate wake-up registers and the system should write a 1b to the *PME_En* bit of the PMCSR prior to the transition to D3.

If all PCI functions are programmed into the D3 state the Intel® Xeon® Processor D-1500 Product Family LAN controller suspends scheduling of new TLPs and waits for the completion of all previous TLPs it has sent. The Intel® Xeon® Processor D-1500 Product Family LAN controller clears the *Memory Access Enable* and *I/O Access Enable* bits of the PCI Command register, which disables memory access decode. Any receive packets that have not been transferred into system memory is kept in the device (and discarded later on D3 exit). Any transmit packets that have not be sent can still be transmitted (assuming the Ethernet link is up).

In preparation to a possible transition to D3cold state, and in order to reduce power consumption, whenever entering D3 state, the software device driver might disable one of the LAN ports (LAN disable) and/or the PHY(s) auto-negotiates network link(s) to a lower speed (if supported by the network interface).

5.1.3.3.2 Master Disable

System software might disable master accesses on the integrated I/O interface link by either clearing the *PCI Bus Master* bit or by bringing the function into a D3 state. From that time on, the Intel® Xeon® Processor D-1500 Product Family LAN controller must not issue master accesses for this function. Due



to the full-duplex nature of the integrated I/O interface and the pipelined design in the Intel® Xeon® Processor D-1500 Product Family LAN controller, it might happen that multiple requests from several functions are pending when the master disable request arrives. The protocol described in this section insures that a function does not issue master requests to the integrated I/O interface link after its *Master Enable* bit is cleared (or after entry to D3 state).

Two configuration bits are provided for the handshake between the device function and its software device driver:

- *PCIe Master Disable* bit in the Device Control Register (CTRL) register — When the *PCIe Master Disable* bit is set, the Intel® Xeon® Processor D-1500 Product Family LAN controller blocks new master requests by this function. The Intel® Xeon® Processor D-1500 Product Family LAN controller then proceeds to issue any pending requests by this function. This bit is cleared on master reset all the way to software reset to allow master accesses.
- *PCIe Master Enable Status* bits in the Device Status register — Cleared by the Intel® Xeon® Processor D-1500 Product Family LAN controller when the *PCIe Master Disable* bit is set and no master requests are pending by the relevant function. Set otherwise. Indicates that no master requests are issued by this function as long as the *PCIe Master Disable* bit is set. The following activities must end before the Intel® Xeon® Processor D-1500 Product Family LAN controller clears the *PCIe Master Enable Status* bit:
 - Master requests by the transmit and receive engines
 - All pending completions to the Intel® Xeon® Processor D-1500 Product Family LAN controller are received.

Note: The software device driver disables any reception to the Rx queues as described in [Section 4.5.7.1.1](#). Then the software device driver sets the *PCIe Master Disable* bit when notified of a pending master disable (or D3 entry). The Intel® Xeon® Processor D-1500 Product Family LAN controller then blocks new requests and proceeds to issue any pending requests by this function. The software device driver then reads the change made in the *PCIe Master Disable* bit and then polls the *PCIe Master Enable Status* bit. Once the bit is cleared, it is guaranteed that no requests are pending from this function.

The software device driver might time out if the *PCIe Master Enable Status* bit is not cleared within a given time. Examples for cases that the Intel® Xeon® Processor D-1500 Product Family LAN controller might not clear the *PCIe Master Enable Status* bit for a long time are cases of flow control, link down, or DMA completions not making it back to the DMA block. In these cases, the software device driver should check that the *Transaction Pending* bit (bit 5) in the Device Status register in the PCI configuration space is cleared before proceeding. Intel's device driver software times-out at ~800 μs. The *PCIe Master Disable* bit must be cleared to enable master request to the integrated I/O interface link. This bit should be cleared through reset.

5.1.3.4 Dr State (D3 Cold)

Transition to Dr state is initiated on several occasions:

- On system power up - Dr state begins with the assertion of Power Good Reset and ends with de-assertion of integrated I/O reset.
- On transition from a D0a state- During operation the system might assert the integrated I/O reset at any time. In an ACPI system, a system transition to the G2/S5 state causes a transition from D0a to Dr state. The transition can be orderly (such as user selected a shut down operating system option), in which case the software device driver might have a chance to intervene. Or, it might be an emergency transition (like power button override), in which case, the software device driver is not notified.
- On transition from a D3 state - The system transitions the device into the Dr state by asserting the Integrated I/O Reset.

Any wake-up filter settings that were enabled before entering this reset state are maintained.



The system might maintain the Integrated I/O Reset asserted for an arbitrary time. The de-assertion (rising edge) of the Integrated I/O Reset causes a transition to D0u state.

While in Dr state, the Intel® Xeon® Processor D-1500 Product Family LAN controller might maintain functionality (for WoL or manageability) or might enter a Dr Disable state (if no WoL and no manageability) for minimal device power. The Dr Disable mode is described in the sections that follow.

5.1.3.4.1 Entry to Dr State at Power On

Dr-entry on platform power-up is as follows:

- Asserting Power Good Reset. Device power is kept to a minimum by keeping the PHYs in low power.
- The NVM is then read and determines device configuration.
- If the *APM Enable* bit in the NVM's Control Word 3 is set then APM wake up is enabled (for each port independently).
- If the *MNG Enable* bit in the NVM word is set, pass-through manageability is not enabled.
- Each of the LAN ports can be enabled if required for WoL or manageability. See [Appendix B](#) for the exact condition to enable a port. In such a case, the PHY might auto-negotiate to a lower speed on Dr entry (see [Section 5.2.2](#)).
- The integrated I/O interface link is not enabled in Dr state following system power up (since the Integrated I/O Reset is asserted).

5.1.3.4.2 Transition to Dr (D3Cold) State

When the the Intel® Xeon® Processor D-1500 Product Family LAN controller is in D3 or D0a state the transition to Dr state begins with a dedicated message.

Upon the reception of this message the Intel® Xeon® Processor D-1500 Product Family LAN controller performs the following operations:

- Suspends the integrated I/O interface scheduling of new TLPs and waits for the completion of all previous TLPs it has sent.
- Responds to all pending integrated I/O interface non-posted requests.
- Completes the scheduling of any of the following pending events:
 - PCI Error

Once the previous steps complete, the Intel® Xeon® Processor D-1500 Product Family LAN controller sends out a dedicated message to indicate to the SoC that it is ready for reset.


After the response is sent, the Intel® Xeon® Processor D-1500 Product Family LAN controller is ready to transition to Dr state and therefore takes additional steps to limit its operation:

- Stop sending assert/de-assert INTx messages

To reduce power consumption, if any of manageability, APM wake or PCI-PM PME¹ is enabled, the PHY might auto-negotiate to a lower link speed on D0a to Dr transition (see [Section 5.2.2](#)).

At this point the IP is ready to have its integrated I/O reset asserted.

1. ACPI 2.0 specifies that OSPM does not disable wake events before setting the *SLP_EN* bit when entering the S5 sleeping state. This provides support for remote management initiatives by enabling Remote Power On (RPO) capability. This is a change from ACPI 1.0 behavior.



5.1.3.4.3 Dr Disable Mode

The Intel® Xeon® Processor D-1500 Product Family LAN controller enters a Dr disable mode on transition to D3cold state when it does not need to maintain any functionality. The conditions to enter Dr Disable state are:

- The IP is in Dr state (all LAN functions are in Dr state)
- APM WOL is inactive for all LAN functions
- Pass-through manageability is disabled for all LAN functions as described in the following sections:
 - SMBus: [Section 10.5.9](#).
 - NC-SI: [Section 10.6.2.2](#).
- ACPI PME is disabled for all PCI functions

Entry into Dr disable is usually done on assertion of the Integrated I/O Reset.

It might also be possible to enter Dr disable mode by reading the NVM while already in Dr state. The usage model for this later case is on system power up, assuming that manageability and wake up are not required. Once the Intel® Xeon® Processor D-1500 Product Family LAN controller enters Dr state on power-up, the NVM is read. If the NVM contents determine that the conditions to enter Dr Disable are met, the Intel® Xeon® Processor D-1500 Product Family LAN controller then enters this mode (assuming that the Integrated I/O Reset is still asserted).

Exit from Dr disable is through de-assertion of the Integrated I/O Reset.

5.2 Network Interfaces Power Management

The Intel® Xeon® Processor D-1500 Product Family LAN controller transitions any of the network interfaces into a low-power state in the following cases:

- The respective LAN function is in LAN disable mode.
- The Intel® Xeon® Processor D-1500 Product Family LAN controller is in D3 or Dr state, APM WoL is disabled for the port, ACPI wake is disabled for the port and pass-through manageability is disabled for the port.

Use of the LAN ports for pass-through manageability follows this behavior:

- If manageability is disabled (*MNG Enable* bit in the NVM is cleared), then LAN ports are not allocated for manageability.
- If manageability is enabled:
 - Power-up — Following an NVM read, a single port is enabled for manageability, running at the lowest speed supported by the interface. If APM WoL is enabled on a single port, the same port is used for manageability. Otherwise, manageability protocols (such as teaming) determine which port is used.
 - D0 state — Both LAN ports are enabled for manageability.
 - D3 and Dr states — A single port is enabled for manageability, running at the lowest speed supported by the interface. If WoL is enabled on a single port, the same port is used for manageability. Otherwise, manageability protocols such as teaming) determine which port is used.

Enabling a port as a result of the previous behaviors cause an internal reset of the port including its associated PHY.

When a network interface is in low-power state, the Intel® Xeon® Processor D-1500 Product Family LAN controller MAC asserts an internal signals to notify the integrated PHY that it might power down as well. Alternatively, the Intel® Xeon® Processor D-1500 Product Family LAN controller might assert an



internal signal to trigger LPLU to re-negotiation of link to a lower link speed. When working with an external PHY, the software device driver should power down the PHY using the dedicated SDPs or the side-band management interface.

5.2.1 PHY Power-Down State

Each Broadwell-DE LAN controller port enters a power-down state when none of its clients is enabled and therefore has no need to maintain a link. This can happen in one of the following cases:

1. **D3/Dr state:** Each PHY enters a low-power state if all the following conditions are met:
 - a. The LAN function associated with this PHY is in a non-D0 state.
 - b. APM WoL is inactive.
 - c. Manageability doesn't use this port.
 - d. ACPI PME is disabled for this port.
2. **LAN disable strap:** Each PHY is disabled if the associated LAN disable strap indicates that the port should be disabled.
3. **LAN PCI Disable bit in NVM:** A single LAN port can also be disabled through NVM settings. If the *LAN PCI Disable* bit is set in NVM PCIe Control 2 word, and if the port is not used for WoL or for MC traffic, then the *LAN Disable Select* bit selects the MAC and PHY port that enters power down even in D0 state. Note that if the port is used for WoL or by the MC, setting the *LAN PCI Disable* bit in NVM PCIe Control 2 word does not bring the MAC and PHY into power down, but only the DMA block.

When powered down by one of these means, a significant portion of the MAC and the PHY might be powered down.

When the Intel® Xeon® Processor D-1500 Product Family LAN controller is completely powered down (Dr state), the PHYs might reach a deeper power saving mode, however; they should retain all NVM loaded settings. When the PHY exits power down, it re-initializes all analog functions.

5.2.2 Low Power Link Up (LPLU)

Note: LPLU supports 10GBASE-T interfaces, does not support SFP+ interfaces and is disabled by default for the KR backplane interface.


Normal PHY speed negotiation drives to establish a link at the highest possible speed. The PHY might support an additional mode of operation, where the PHY drives to establish a link at a low speed. The LPLU process enables a link to come up at the lowest possible speed in cases where power is more important than performance. Different behavior is defined for the D0 state and the other non-D0 states.

For more details on LPLU see [Section B.6](#).

5.2.3 Energy Efficient Ethernet (EEE)

Note: EEE is ONLY supported for 10GBASE-T interfaces via the X557-AT2 PHY.

The Intel® Xeon® Processor D-1500 Product Family LAN controller enters EEE Low Power Idle (LPI) mode on transmit each time the Intel® Xeon® Processor D-1500 Product Family LAN controller detects no data is scheduled for transmission.



EEE LPI mode, defined in IEEE802.3 Clause 78, enables power saving by switching off part of the PHY functionality when no data needs to be transmitted or/and received. A decision on whether the Intel® Xeon® Processor D-1500 Product Family LAN controller transmit path should enter LPI mode or exit LPI mode is done according to a need to transmit. Information on whether a link partner has entered LPI mode is detected by PHY and used for power saving in the receive circuitry.

When no data needs to be transmitted, a request to enter transmit LPI is issued on the internal XGMII Tx interface causing the PHY to transmit sleep symbols for a predefined period of time followed by a quiet period. During LPI, the PHY periodically transmits refresh symbols that are used by the link partner to update adaptive filters and timing circuits in order to maintain link integrity. This quiet-refresh cycle continues until transmitting normal inter-frame encoding on the internal XGMII Tx interface. The PHY communicates to the link partner the move to link active state by sending wake symbols for a predefined period of time. The PHY then enters a normal operating state where data or idle symbols are transmitted.

In the receive direction, entering LPI mode is triggered by receiving sleep symbols from the link partner. This signals that the link partner is about to enter LPI mode. After sending the sleep symbols, the link partner ceases transmission. When the link partner initiates the move to LPI, the PHY indicates assert low power idle on the internal XGMII Rx interface and the PHY receiver disables functionality to reduce power consumption.

Figure 5.2 and Table 5.1 show the general principles of EEE LPI operation on the Ethernet link.

Table 5.1. EEE Parameters

Parameter	Description
Sleep Time (T_s)	Duration PHY sends sleep symbols before going quiet.
Quiet Duration (T_q)	Duration PHY remains quiet before it must wake for refresh period.
Refresh Duration (T_r)	Duration PHY sends refresh symbols for timing recovery and coefficient synchronization.
PHY Wake Time (T_w_{PHY})	Minimum duration PHY takes to resume to active state after decision to wake.
Receive System Wake Time ($T_w_{System_rx}$)	Wait period where no data is expected to be received to give the local receiving system time to wake up.
Transmit System Wake Time ($T_w_{System_tx}$)	Wait period where no data is transmitted to give the remote receiving system time to wake up.

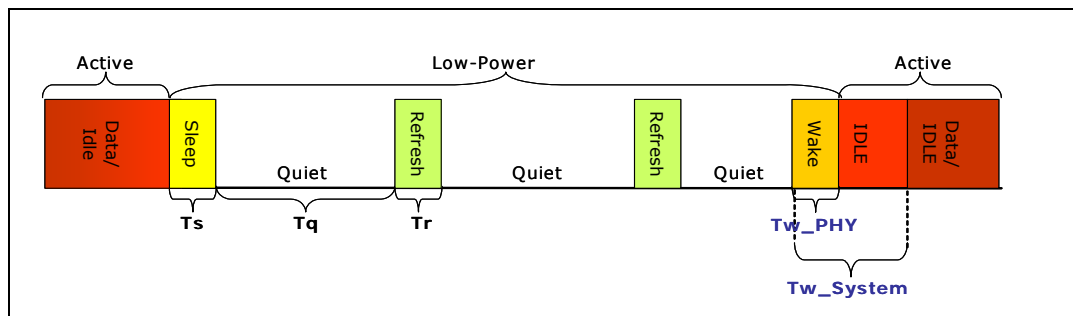


Figure 5.2. EEE Operation

5.2.3.1 Conditions to Enter EEE Tx LPI

In the transmit direction entry into EEE LPI mode of operation is triggered when one of the following conditions exist:



1. No transmission is pending, management does not need to transmit, the internal transmit buffer is empty, and the *EEER.TX_LPI_EN* bit is set to 1b.
2. If the *EEER.TX_LPI_EN* bit is set to 1b and a XOFF flow control packet is received from the link partner.
3. When *EEER.Force_TLPI* is set (even if *EEER.TX_LPI_EN* is cleared).
 - If *EEER.Force_TLPI* is set in mid-packet, the Intel® Xeon® Processor D-1500 Product Family LAN controller completes the packet transmission and then moves Tx to LPI.

When one of the previous conditions to enter Tx LPI state are detected assert LPI is transmitted on the internal xxMII interface and the the Intel® Xeon® Processor D-1500 Product Family LAN controller PHY transmits sleep symbols on the network interface to communicate to the link partner entry into Tx LPI link state. After sleep symbols transmission, the PHY immediately enters the low power quiet mode. In this state the PHY periodically transitions between quiet link state, where link is idle, to sending refresh symbols until a request to transition link back to normal (active) mode is transmitted on the internal xxMII Tx interface (See [Figure 5.2](#)).

Note: Initial *EEER.TX_LPI_EN* configuration is loaded from NVM.

Note: IEEE802.3 Ethernet LPI enables each link direction to enter sleep, refresh or wake states asymmetric from the other direction.

Note: EEE LPI status of a LAN controller port can be found in the *EEE_STAT* register.

5.2.3.2 Transition from Tx LPI to Active Link State

The Intel® Xeon® Processor D-1500 Product Family LAN controller exits Tx LPI link state and transitions the link into active link state when none of the conditions defined in [Section 5.2.3.1](#) exist. To transition into active link state, the Intel® Xeon® Processor D-1500 Product Family LAN controller transmits:


1. Normal inter-frame encoding on the internal xxMII Tx interface for a pre-defined link rate dependent period time of *Tw_sys_tx-min* (defined by the *EEE_SU* register and IEEE802.3az clause 78.5). As a result, the PHY transmits wake symbols for a *Tw_phy* duration followed by idle symbols.
2. If the *Tw_System_tx* duration defined in the *EEER.Tw_system* field is longer than *Tw_sys_tx-min*, the Intel® Xeon® Processor D-1500 Product Family LAN controller continues transmitting the inter-frame encoding on the internal xxMII interface until the time defined in the *EEER.Tw_system* field has expired, before transmitting the actual data. During this period the PHY continues transmitting idle symbols.

Note: When moving out of Tx LPI to transmit a 802.3x flow control frame, the Intel® Xeon® Processor D-1500 Product Family LAN controller waits for the *Tw_sys_tx-min* duration before transmitting the flow control frame. It should be noted that even in this scenario actual data is transmitted only after the *Tw_System_tx* time defined in the *EEER.Tw_system* field has expired.

5.2.3.3 EEE Auto-Negotiation

Auto-negotiation provides the capability to negotiate EEE capabilities with the link partner using the next page mechanism defined in IEEE802.3 Annex 73A. IEEE802.3 auto-negotiation is performed at power up, on command from software, upon detection of a PHY error or following link re-connection.

During the link establishment process, both link partners indicate their EEE capabilities using IEEE802.3 auto-negotiation. If EEE is supported by both link partners for the negotiated PHY type then the EEE function can be used independently in either direction.



EEE capabilities advertised during auto-negotiation can be modified via the Auto-negotiation EEE Advertisement Register in the backplane PHY registers.

5.2.3.4 EEE Link Level (LLDP) Capabilities Discovery

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports LLDP negotiation via software using the EEE IEEE802.1AB Link Layer Discovery Protocol (LLDP) Type, Length, Value (TLV) fields defined in IEEE802.3az clause 78 and clause 79. LLDP negotiation enables negotiation of increased system wake time (Transmit T_w and Receive T_w) to enable improving system energy efficiency.

5.2.3.4.1 LLDP Negotiation Actions

Following negotiation of a new system wake time via EEE LLDP negotiation, the following fields and registers should be updated:

The *EEER.Tw_system* field with the negotiated Transmit T_w time value, to increase the duration where idle symbols are transmitted following a move out of EEE Tx LPI state before actual data can be transmitted.

- A value placed in *EEER.Tw_system* field does not effect transmission of flow control packets. Depending on the technology flow control packet, transmission is delayed following a move out of EEE Tx LPI state only by the minimum *Tw_sys_tx* time as defined in IEEE802.3az clause 78.5. In the Intel® Xeon® Processor D-1500 Product Family LAN controller, the minimum *Tw_sys_tx* time value is defined in the *EEE_SU* register together with time defined in IEEE802.3az clause 78.5. Value varies as a function of link rate and technology.

5.2.3.5 EEE Statistics

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports reporting a number of EEE LPI Tx and Rx events via the RLPIC and TLPIC registers.

5.3 Wake Up

5.3.1 Advanced Power Management Wake Up

Advanced power management wake up, or APM wake up, was previously known as Wake on LAN (WoL). It is a feature that has existed in the 10/100 Mb/s NICs for several generations. The basic premise is to receive a packet with an explicit data pattern, and then to assert a signal to wake up the system. The NIC asserts the signal for approximately 50 ms to signal a wake up.

At power up, the Intel® Xeon® Processor D-1500 Product Family LAN controller reads the *APM Enable* bit from the NVM into the *APM Enable (APME)* bits of the GRC register. This bit control the enabling of *APM Wakeup*.

When *APM Wakeup* is enabled, the Intel® Xeon® Processor D-1500 Product Family LAN controller checks all incoming packets for Magic Packets. See [Section A.1.3](#) for a definition of Magic Packets.

Once the Intel® Xeon® Processor D-1500 Product Family LAN controller receives a matching magic packet, it:

- Asserts the WAKE signal so that clock and power will be resumed to the CPU when it is in Sx or possibly reset when it is in S0¹.

1. Reset on LAN due to Magic Packet is a platform feature. The Intel® Xeon® Processor D-1500 Product Family LAN controller performs the same flow as Wake on LAN.



- Sets the *PME_Status* bit in the *PMCSR*.
- Stores the first 128 bytes of the packet in the Wake Up Packet Memory (WUPM).
- Sets the Magic Packet *Received* bit in the Wake up Status (WUS) register.
- Sets the packet length in the Wake up Packet Length Register (WUPL).
- Once the CPU is powered on and the integrated I/O (P-IOSF) reset is de-asserted, sends a PM_PME P_IOSF message.

The Intel® Xeon® Processor D-1500 Product Family LAN controller maintains the first magic packet received in the Wake Up Packet Memory (WUPM) until the software device driver writes a 0b to the Magic Packet Received *MAG* bit in the Wake Up Status Register (WUS), hence the software device driver has to read the packet from WUPM prior to clearing the WUS indication.

APM Wakeup event might be issued in any¹ power state. It is disabled if a subsequent NVM read results in the *APM Wake Up* bit being cleared or if the software explicitly writes a 0b to the *APM Wake Up (APM)* bit of the GRC register.

5.3.2 ACPI Power Management Wake Up

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports ACPI power management based wake up. It can generate system wake-up events from three sources, regardless of the power state:

- Receiving a Magic Packet.
- Receiving a network wake up packet.
- Detecting a link change of state.
- Firmware reset event.

Activating ACPI power management wake up requires the following steps:

- The operating system (at configuration time) writes a 1b to the *PME_En* bit of the *PMCSR* (*PMCSR.8*).
- The software device driver clears all pending wake-up status in the WUS register by writing 1b to all the status bits.
- The software device driver programs the WUFC register to indicate the packets it needs to wake up and supplies the necessary data to the IPv4/v6 Address Table (IP4AT, IP6AT), Flexible Host Filter Table (FHFT). It can also set the *Link Status Change Wake Up Enable (LNKC)* bit in the WUFC register to cause wake up when the link changes state.
- Once the Intel® Xeon® Processor D-1500 Product Family LAN controller wakes the system, the software device driver needs to clear the WUS and WUFC registers until the next time the system goes to a low-power state with wake up.

Normally, after enabling ACPI wake up, the operating system writes (11b) to the lower two bits of the *PMCSR* to put the Intel® Xeon® Processor D-1500 Product Family LAN controller into low-power mode, and once entering D0, OS disables ACPI wake up.

Once wake up is enabled, the Intel® Xeon® Processor D-1500 Product Family LAN controller monitors incoming packets, first filtering them according to its standard address filtering method, then filtering them with all of the enabled wake-up filters. If a packet passes both the standard address filtering and at least one of the enabled wake-up filters, the Intel® Xeon® Processor D-1500 Product Family LAN controller:

- Sets the *PME_Status* bit in the *PMCSR*.

1. Wake on LAN in D0 (S0) might be used for Reset on LAN at the platform level.



- If the *PME_En* bit in the *PMCSR* is set and device is in Dr state or *PCI_GLBL_CNF.WAKE_PIN_EN* is set in all states, asserts GBE wake pin.
- In non Dr state or when exiting Dr, a *PM_PME* message is also issued.
- Stores the first 128 bytes of the packet in the Wake Up Packet Memory (WUPM) register. Sets one or more of the *Received* bits in the WUS register. Note that the Intel® Xeon® Processor D-1500 Product Family LAN controller sets more than one bit if a packet matches more than one filter.
- Sets the packet length in the Wake Up Packet Length Register (WUPL).

If enabled, a link state change wake up causes similar results, setting *PME_Status*, asserting GBE wake pin and setting the *LNKC* bit in the WUS register when the link goes up or down.

The GBE wake pin remains asserted until the operating system either writes a 1b to the *PME_Status* bit of the *PMCSR* register or writes a 0b to the *PME_En* bit.

If enabled, a firmware reset causes similar results, setting *PME_Status*, asserting the GBE wake pin and setting the *FW_RST_WK* bit in the *WUS* register when the firmware is reset, thus enabling the software device driver to reconfigure proxying parameters that were lost due to the firmware reset event.

The GBE wake pin remains asserted until the operating system either writes a 1b to the *PME_Status* bit of the *PMCSR* register or writes a 0b to the *PME_En* bit. the WUS register.

5.3.3 Wake-Up Packets

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports various wake-up packets using two types of filters:

- Pre-defined filters
- Flexible filters

Each of these filters are enabled if the corresponding bit in the WUFC register is set to 1b.

5.3.3.1 Pre-Defined Filters

The following packets are supported by the Intel® Xeon® Processor D-1500 Product Family LAN controller's pre-defined filters:

- Directed packet (including exact, multicast indexed, and broadcast)
- Magic Packet
- ARP/IPv4 request packet
- Directed IPv4 packet
- Directed IPv6 packet

Each of these filters are enabled if the corresponding bit in the WUFC register is set to 1b.

The explanation of each filter includes a table listing which bytes at which offsets are compared to determine if the packet passes the filter. Both VLAN frames and LLC/SNAP can increase the given offsets if they are present.

5.3.3.1.1 Directed Packets

Unicast — The Intel® Xeon® Processor D-1500 Product Family LAN controller generates a wake-up event after receiving any packet whose destination address matches one of the 128 valid programmed receive addresses if the *Directed Exact Wake Up Enable* bit is set in the WUFC register (WUFC.EX).



For multicast packets, the upper bits of the incoming packet's destination address index a bit vector, the Multicast Table Array (MTA) that indicates whether to accept the packet. If the *Directed Multicast Wake Up Enable* bit set in the WUFC register (WUFC.MC) and the indexed bit in the vector is one then the Intel® Xeon® Processor D-1500 Product Family LAN controller generates a wake-up event. The exact bits used in the comparison are programmed by software in the *Multicast Offset* field of the Multicast Control register (MCSTCTRL.MO).

If the *Broadcast Wake Up Enable* bit in the WUFC register (WUFC.BC) is set, the Intel® Xeon® Processor D-1500 Product Family LAN controller generates a wake-up event when it receives a broadcast packet.

5.3.3.1.2 Magic Packet

A Magic Packet's destination address must match the address filtering enabled in the configuration registers with the exception that broadcast packets are considered to match even if the *Broadcast Accept* bit of the Receive Control register (FCTRL.BAM) is 0b. If *APM Wakeup* is enabled in the NVM, the Intel® Xeon® Processor D-1500 Product Family LAN controller starts up with the Receive Address Register 0 (RAH0, RAL0) loaded from the NVM. This enables the Intel® Xeon® Processor D-1500 Product Family LAN controller to accept packets with the matching IEEE address before the software device driver comes up.

Offset	# of Bytes	Field	Value	Action	Comment
0	6	Destination Address		Compare	MAC header – processed by main address filter.
6	6	Source Address		Skip	
12	S= (0/4)	Possible VLAN Tag		Skip	
12/16	D=(0/8)	Possible Len/LLC/SNAP Header		Skip	
12+S+D	2	Type		Skip	
Any	6	Synchronizing Stream	FF*6+	Compare	
any+6	96	16 Copies of Node Address	A*16	Compare	Compared to Receive Address Register 0 (RAH0, RAL0).

5.3.3.1.3 ARP/IPv4 Request Packet

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports receiving ARP request packets for wake up if the *ARP* bit is set in the WUFC register. Four IPv4 addresses are supported, which are programmed in the IPv4 Address Table (IP4AT). A successfully matched packet must pass L2 address filtering, a Protocol Type of 0x0806, an ARP OpCode of 0x01, and one of the four programmed IPv4 addresses. The Intel® Xeon® Processor D-1500 Product Family LAN controller also handles ARP request packets that have VLAN tagging on both Ethernet II and Ethernet SNAP types.

Offset	# of Bytes	Field	Value	Action	Comment
0	6	Destination Address		Compare	MAC header – processed by main address filter.
6	6	Source Address		Skip	
12	S=(0/4)	Possible VLAN Tag		Compare	
12/16	D=(0/8)	Possible Len/LLC/SNAP Header		Skip	
12+S+D	2	Type	0x0806	Compare	ARP



Offset	# of Bytes	Field	Value	Action	Comment
14+S+D	2	Hardware Type	0x0001	Compare	
16+S+D	2	Protocol Type	0x0800	Compare	
18+S+D	1	Hardware Size	0x06	Compare	
19+S+D	1	Protocol Address Length	0x04	Compare	
20+S+D	2	Operation	0x0001	Compare	
22+S+D	6	Sender Hardware Address	-	Ignore	
28+S+D	4	Sender IP Address	-	Ignore	
32+S+D	6	Target Hardware Address	-	Ignore	
38+S+D	4	Target IP Address	IP4AT	Compare	Can match any of 4 values in IP4AT.

5.3.3.1.4 Directed IPv4 Packet

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports receiving directed IPv4 packets for wake up if the *IPV4* bit is set in the WUFC register. Four IPv4 addresses are supported that are programmed in the IPv4 Address Table (IP4AT). A successfully matched packet must pass L2 address filtering, a Protocol Type of 0x0800, and one of the four programmed IPv4 addresses. The Intel® Xeon® Processor D-1500 Product Family LAN controller also handles directed IPv4 packets that have VLAN tagging on both Ethernet II and Ethernet SNAP types.

5.3.3.1.5 Directed IPv6 Packet

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports receiving directed IPv6 packets for wake up if the *IPV6* bit is set in the WUFC register. One IPv6 address is supported and it is programmed in the IPv6 Address Table (IP6AT). A successfully matched packet must pass L2 address filtering, a Protocol Type of 0x0800, and the programmed IPv6 address. The Intel® Xeon® Processor D-1500 Product Family LAN controller also handles directed IPv6 packets that have VLAN tagging on both Ethernet II and Ethernet SNAP types.

5.3.3.2 Flexible Filter

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports a total of eight host flexible filters. Each filter can be configured to recognize any arbitrary pattern within the first 128 bytes of the packet. To configure the flexible filter, software programs the required values into the Flexible Host Filter Table (FHFT_FILTER). These contain separate values for each filter. Software must also enable the filter in the WUFC register, and enable the overall wake-up functionality must be enabled by setting *PME_En* in the PMCS register or the WUC register.

Structure of the Flexible Host Filter Table:

31	0	31	8	7	0	31	0	31	0
Reserved		Reserved		Mask [7:0]		DW 1		DW 0	
Reserved		Reserved		Mask [15:8]		DW 3		DW 2	
Reserved		Reserved		Mask [23:16]		DW 5		DW 4	
Reserved		Reserved		Mask [31:24]		DW 7		DW 6	



. . .

31	7	6	0	31	8	7	0	31	0	31	0
Reserved		Reserved		Reserved		Mask [127:120]		DW 29		DW 28	
Reserved		Length		Reserved		Mask [127:120]		DW 31		DW 30	

Each of the filters is allocated addresses as follows:

- Filter 0 — 0x09000 — 0x090FF
- Filter 1 — 0x09100 — 0x091FF
- Filter 2 — 0x09200 — 0x092FF
- Filter 3 — 0x09300 — 0x093FF
- Filter 4 — 0x09600 — 0x096FF
- Filter 5 — 0x09700 — 0x097FF
- Filter 6 — 0x09800 — 0x098FF
- Filter 7 — 0x09900 — 0x099FF

The following table lists the addresses used for filter 0.


Field	Dword	Address	Bit(s)	Initial Value
Filter 0 DW0	0	0x09000	31:0	X
Filter 0 DW1	1	0x09004	31:0	X
Filter 0 Mask[7:0]	2	0x09008	7:0	X
Reserved	3	0x0900C		X
Filter 0 DW2	4	0x09010	31:0	X
...				
Filter 0 DW30	60	0x090F0	31:0	X
Filter 0 DW31	61	0x090F4	31:0	X
Filter 0 Mask[127:120]	62	0x090F8	7:0	X
Length	63	0x090FC	6:0	X

Accessing the FHFT_FILTER registers during filter operation might result in a packet being mis-classified if the write operation collides with packet reception. It is therefore advised that the flex filters are disabled prior to changing their setup.

Once enabled, the flexible filters scan incoming packets for a match. If the filter encounters any byte in the packet where the mask bit is one and the byte doesn't match the byte programmed in the Flexible Host Filter Table (FHFT_FILTER) then the filter fails that packet. If the filter reaches the required length without failing the packet, it passes the packet and generates a wake-up event. It ignores any mask bits set to one beyond the required length.

Packets that passed a wake-up flexible filter should cause a wake-up event only if it is directed to the Intel® Xeon® Processor D-1500 Product Family LAN controller (passed L2 and VLAN filtering).

Note: The flexible filters are temporarily disabled when read from or written to by the host. Any packet received during a read or write operation is dropped. Filter operation resumes once the read or write access is done.



5.3.3.3 Wake-Up Packet Storage

The Intel® Xeon® Processor D-1500 Product Family LAN controller saves the first 128 bytes of the wake-up packet in its internal buffer, which can be read through the Wake Up Packet Memory (WUPM) after the system wakes up.

5.3.4 Wake Up and Virtualization

When operating in a virtualized environment, all wake-up capabilities are managed by a single entity (such as the VMM or an IOVM). In an IOV architecture, the physical software device driver controls wake up and none of the Virtual Machines (VMs) has direct access to the wake-up registers. The wake-up registers are not replicated.

5.4 Protocol Offload (Proxying)

In prior operating system releases, ARP and IPv6 neighbor discovery messages were one of the possible wake-up types for the platform. ARP and IPv6 neighbor discovery packets are required to enable other network devices to discover the link layer address used by the system. Supporting these protocols while the host is in low-power state is fundamental to maintain remote network accessibility to the sleeping host. If the host does not respond, other devices in the network eventually are not able to send routable network traffic (such as IPv4 and IPv6) to the sleeping host.

1. Implementing ARP offload as defined in the Power Management specification on the NDIS Program Connect site. Specifically, the offload must respond to an ARP Request (operation = 1) by responding with an ARP Reply (operation = 2) as defined in RFC 826.
2. Implementing IPv6 NS offload as defined in Power Management specification on the NDIS Program, Connect site. Specifically, the offload must respond to an Neighbor Solicitation (operation = 135) by responding with an NS Advertisement (operation = 136) as defined in RFC 4861. Devices must support at least 2 NS offloads, each with up to 2 target IPv6 addresses:
 - a. When Neighbor Solicitation (NS) Protocol off load is enabled, the Intel® Xeon® Processor D-1500 Product Family LAN controller also supports protocol off load of up to two IPv6 Multicast-Address-Specific Multicast Listener Discovery (MLD) Queries (either MLDv1 or MLDv2) per function. In addition, the Intel® Xeon® Processor D-1500 Product Family LAN controller also responds to General MLD queries, used to learn which IPv6 multicast addresses have listeners on an attached link.
 - MLD protocol offload is supported when NS protocol offload is enabled so that IPv6 routers discover the presence of multicast listeners (that is, nodes wanting to receive multicast packets), for packets with the IPv6 NS solicited-node multicast address and continue forwarding these NS requests on the link.
 - MLD protocol off load is supported for either MLD Multicast Listener Query packets or MLD Multicast Address and Source Specific Query packets that check for IPv6 Multicast Listeners with the solicited-node multicast address placed in the IPv6 destination address field of the IPv6 NS packets that are off-loaded by the Intel® Xeon® Processor D-1500 Product Family LAN controller.
 - Responds to the IPv6 MLD queries, with the solicited-node multicast address placed in the IPv6 destination address field of the IPv6 NS packets that are off-loaded by the Intel® Xeon® Processor D-1500 Product Family LAN controller (as defined in RFC 2710 and RFC 3810).

5.4.1 Proxying Filters

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports protocol offload (proxying) of packets using the embedded management controller. The relevant packet is directed to the MC using the pre-defined filters. Proxy filters are enabled if the corresponding bit in the *Proxying Filter Control (PROXYFC)* register is set to 1b.



The following packet types are detected per PF by the Intel® Xeon® Processor D-1500 Product Family LAN controller's pre-defined filters:

- General and directed ARP/IPv4 request packet
- General and directed NS IPv6 packet
- MLD IPv6 packet

Each of these packet type filters is enabled if the corresponding bit in the *Proxying Filter Control (PROXYFC)* register is set to 1b.

All the proxy filters are applied only to traffic that passes the L2 filtering described in [Section 7.1.3](#).

Note: Packets for which the IP or TCP/UDP checksum is wrong are not considered as candidates for proxy.

5.4.1.1 Directed Packets

The Intel® Xeon® Processor D-1500 Product Family LAN controller forwards to proxying any packet whose destination address matches one of the 128 valid programmed receive addresses if the Directed Exact PROXY Enable bit is set in the PROXYFC register (*PROXYFC.EX*).

5.4.1.2 ARP/IPv4 Request Packet

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports receiving ARP request packets for proxying if the *ARP_Directed* bit or the *ARP* bit is set in the *Proxying Filter Control (PROXYFC)* register.

- If the *Directed ARP* bit is set, a successfully matched packet must contain a broadcast/unicast MAC address, a protocol type of 0x0806, an ARP OpCode of 0x01 and the target IP address matches one of the four IPv4 addresses programmed in the *IPv4 Address Table (IP4AT)*, in addition, the corresponding *IPAV.V40*, *IPAV.V41*, *IPAV.V42* or *IPAV.V43* bit should be set.
- If the *ARP* bit is set, a successfully matched packet must contain a broadcast MAC address or one of the unicast addresses configured for the host, a protocol type of 0x0806 and an ARP OpCode of 0x01.

The Intel® Xeon® Processor D-1500 Product Family LAN controller also handles ARP request packets that have VLAN tagging or double VLAN tagging on both Ethernet II and Ethernet SNAP types.



5.4.1.2.1 ARP Request Packet

Offset	# of Bytes	Field	Value	Action	Comment
0	6	Destination Address		Compare	MAC Header – processed by main address filter.
6	6	Source Address		Skip	
12	D=(0/4/6/8)	Outer Tag (Outer VLAN or E-tag)	0x8100 **** 0x893F ***** *****	Skip	
12 + D	S = (0/4)	Possible VLAN Tag		Skip	
12 + S + D	2	Type	0x0806	Compare	ARP
14 + S + D	2	Hardware Type	0x0001	Compare	
16 + S + D	2	Protocol Type	0x0800	Compare	
18 + S + D	1	Hardware Size	0x06	Compare	
19 + S + D	1	Protocol Address Length	0x04	Compare	
20 + S + D	2	Operation	0x0001	Compare	
22 + S + D	6	Sender Hardware Address	-	Ignore	
28 + S + D	4	Sender IP Address	-	Ignore	
32 + S + D	6	Target Hardware Address	-	Ignore	
38 + S + D	4	Target IP Address	IP4AT	Compare	Match IP4AT values or zero.
42 + S + D	18 - S - D	Padding	0x00	Ignore	Padding to 64 bytes.
60	4	CRC		Check	

5.4.1.2.2 ARP Response Packet

Offset	# of Bytes	Field	Value	Action
0	6	Destination Address		Copy from ARP request source address.
6	6	Source Address		Station address.
12	D=(0/4/6/8)	Outer Tag (Outer VLAN or E-tag)	0x8100 **** 0x893F ***** *****	Skip.
12	S = (0/4)	Possible VLAN Tag		Copy from ARP request.
12 + S + D	2	Type	0x0806	Constant (Copy from ARP request).
14 + S + D	2	Hardware Type	0x0001	Constant (Copy from ARP request).
16 + S + D	2	Protocol Type	0x0800	Constant (Copy from ARP request).
18 + S + D	1	Hardware Size	0x06	Constant (Copy from ARP request).
19 + S + D	1	Protocol Address Length	0x04	Constant (Copy from ARP request).
20 + S + D	2	Operation	0x0002	Constant.
22 + S + D	6	Sender Hardware Address		Station address.



Offset	# of Bytes	Field	Value	Action
28 + S + D	4	Sender IP Address		Target IP address from ARP request or valid IP address if target IP was zero.
32 + S + D	6	Target Hardware Address		Sender MAC address from ARP request.
38 + S + D	4	Target IP Address		Sender IP address from ARP request.
42 + S + D	18 - S - D	Padding	0x00	Padding to 64 bytes.
60	4	CRC		Calculate.

5.4.1.3 NS IPv6 Packet

In IPv6 networks, ICMPv6 neighbor solicitation and neighbor advertisement provides the address mapping of the IP address to a corresponding MAC address.

Machines that operate in IPv6 networks are sent an ICMPv6 neighbor solicitation and must respond with their link-layer (MAC) address in their ICMPv6 neighbor advertisement response. The solicitation can be for either the link-local, global, or a temporary IPv6 addresses.

Neighbor discovery messages have both an IPv6 header and the ICMPv6 header. The IPv6 header is a standard one, including the source and destination IP addresses. The network proxy off load does not support IPv6 neighbor discovery messages that also have IPv6 header extensions, these packets are silently discarded with no reply. The Intel® Xeon® Processor D-1500 Product Family LAN controller supports receiving IPv6 NS packets sent by a node to determine the link-layer address of a neighbor, or to verify that a neighbor is still reachable for wake up or proxying.

If the *NS* or *NS_DIRECTED* bits are set in the Proxying Filter Control (PROXYFC) register packet is sent to the MC for protocol offload.

- If the *NS directed* bit is set, a successfully matched packet must contain the station's MAC address (unicast or multicast), an Ethernet type of 0x86DD, a IPv6 header type of ICMPv6 (0x3A), a ICMPv6 type of 0x87 (NS), correct ICMPv6 checksum and the programmed IPv6 address in the IPv6 Address Table (IP6AT) must match the target IPv6 address. In addition, the corresponding *IPAV.V60*, *IPAV.V61*, *IPAV.V62* or *IPAV.V63* bit should be set.
- If the *NS* bit is set, a successfully matched packet must contain the station's MAC address (unicast or multicast), an Ethernet type of 0x86DD, a IPv6 header htype of ICMPv6 (0x3A), a ICMPv6 type of 0x87 (NS) and a correct ICMPv6 checksum.
- If the packet has an ICMPv6 checksum error, then it is considered as a proxy packet only if the *RXCSUM.ICMPV6XSUM* bit is cleared.

The Intel® Xeon® Processor D-1500 Product Family LAN controller also handles NS IPv6 packets that have VLAN tagging or double VLAN tagging on both Ethernet II and Ethernet SNAP types.

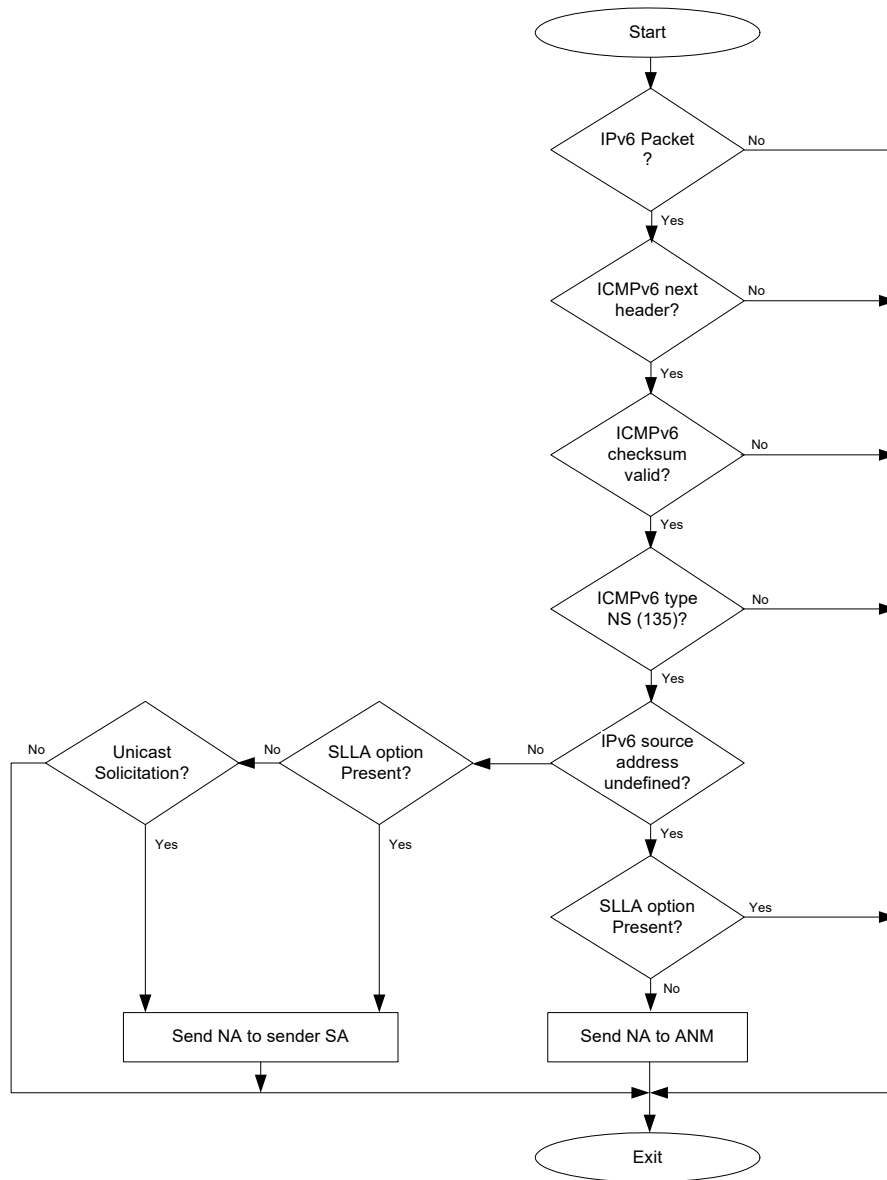


Figure 5.3. NS IPv6 Packet Flow



5.4.1.3.1 IPv6 Neighbor Solicitation Packet

Offset	# of bytes	Field	Value (hex)	Action	Comment
0	6	Destination Address		Compare	MAC Header – processed by main address filter.
6	6	Source Address		Skip	
12	D=(0/4/6/8)	Outer Tag (Outer VLAN or E-tag)	0x8100 **** 0x893F ***** *****	Skip	
12+D	S=(0/4)	Possible VLAN Tag		Skip	
IPv6 header					
12+D+S	2	Type	0x86DD	Compare	IPv6
14+D+S	1	Version/ Traffic Class	0x6	Compare	Check IPv6
15+D+S	3	Traffic Class/Flow Label		Ignore	
18+D+S	2	Payload Length		Ignore	
20+D+S	1	Next Header	0x3A	Check	ICMPv6
21+D+S	1	Hop Limit	0xFF	Compare	
22+D+S	16	Source Address		Ignore	Check if source address is undefined.
38+D+S	16	Destination Address		Ignore	
ICMPv6 header					
54+D+S	1	Type	0x87	Compare	Neighbors solicitation.
55+D+S	1	Code	0x0	Compare	
56+D+S	2	Checksum		Check	
58+D+S	4	Reserved	0x0000	Ignore	
62+D+S	16	Target IP Address	IP6AT	Compare	
78+D+S	1	Type	0x1	Compare	Possible source link layer address option (Should not appear if source address is undefined).
79+D+S	1	Length	0x1	Compare	
80+D+S	6	Link Layer Address		Skip	
86+D+S	4	CRC		Check	



5.4.1.3.2 IPv6 Neighbor Advertisement Packet

Offset	# of bytes	Field	Value (hex)	Action
0	6	Destination Address		Copy from ND packet.
6	6	Source Address		Station address.
12	D=(0/4/6/8)	Outer Tag (Outer VLAN or E-tag)	0x8100 **** 0x893F ***** *****	Skip
12+D	S=(0/4)	Possible VLAN Tag		Copy from ND packet.
IPv6 header				
12+D+S	2	Type	0x86DD	Constant (Copy from ND packet).
14+D+S	1	Version/ Traffic Class	0x6	Constant (Copy from ND packet).
15+D+S	3	Traffic Class/Flow Label		Constant (Copy from ND packet).
18+D+S	2	Payload Length		
20+D+S	1	Next Header	0x3A	Constant.
21+D+S	1	Hop Limit	0xFF	Constant.
22+D+S	16	Source Address		relevant IPv6AT entry (ND target address).
38+D+S	16	Destination Address		Copy from ND packet source address If source address was undefined - send to All Nodes Multicast (FF02::1).
ICMPv6 header				
54+D+S	1	Type	0x88	Constant.
55+D+S	1	Code	0x0	
56+D+S	2	Checksum		Calculate.
58+D+S	4	Flags	0x60000 000	Constant (solicited and override) if the source address was defined.
			0x20000 000	Constant (override) if the source address was undefined.
62+D+S	16	Target IP Address	IPv6AT	Same as source address.
78+D+S	1	Type	0x2	Target link layer address option.
79+D+S	1	Length	0x1	
80+D+S	6	Link Layer Address	From ND	
86+D+S	4	CRC		Calculate.

5.4.1.4 MLD IPv6 Packet

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports receiving IPv6 MLD packets sent by an IPv6 router to discover the presence of multicast listeners (that is, nodes needs to receive multicast packets) on its directly attached links, and to discover specifically which multicast addresses are of interest to those neighboring nodes.

After receiving the following MLD packets (either MLDv1 as defined in RFC 2710 or MLDv2 as defined in RFC 3810):



1. Multicast Listener Query (ICMPv6 Type = decimal 130). Defined in MLDv1 and MLDv2.
2. Multicast Listener Report (ICMPv6 Type = decimal 131). Defined in MLDv1 and MLDv2.
3. Version 2 Multicast Listener Report Message (ICMPv6 Type = decimal 143). Defined in MLDv2 only.

If the *MLD* bit is set in the Proxying Filter Control (PFPM_PROXYFC) register packet is sent to firmware for protocol offload.

If the *MLD* bit is set, a successfully matched packet must contain the station's MAC address (unicast or multicast), a Ethernet type of 0x86DD, a IPv6 header type of ICMPv6 (0x3A), a ICMPv6 type of either 0x82 (130 decimal), 0x83 (131 decimal) or 0x8F (143 decimal) and a correct ICMPv6 checksum.

The Intel® Xeon® Processor D-1500 Product Family LAN controller also handles MLD IPv6 packets that have VLAN tagging on both Ethernet II and Ethernet SNAP types.

If the packet has an ICMPv6 checksum error, then it is considered as an MLD proxy packet only if the *RXCSUM.ICMPV6XSUM* bit is cleared.

5.4.2 Proxying and Virtualization

When operating in a virtualized environment, all proxying capabilities are managed by a single entity (such as the VMM or an IOVM). In an IOV architecture, the physical (PF) driver controls proxying and none of the VMs has direct access to the Proxying registers. The Proxying registers are not replicated per VF.

5.4.3 Protocol Offload (Proxying) Flow

The Proxying register set consists of the following Proxying registers:

1. *PROXYFC* - Proxying Filter Control Register.
2. *PROXYS* - Proxying Status Register.

Proxying address filters:

1. MAC - (RAL/H) L2 address Table.
2. *IP4* - (IP4AT) IPv4 Address Table.
3. *IP6* - (IP6AT) IPv6 Address Table.

5.4.3.1 Protocol Offload Activation

To enable protocol offload, the software device driver should implement the following steps:

1. Clear all pending proxy status bits in the Proxying Status (PROXYS) register.
2. Program the Proxying Filter Control (PROXYFC) register to indicate type of packets that should be forwarded to manageability for proxying and program the necessary data to the IPv4/v6 Address Table (IP4AT, IP6AT).
3. Set the *WUFC.FW_RST_WK* bit to 1b to initiate a wake if firmware reset was issued and proxying information was lost.
4. Take ownership of the management host interface semaphore (*SW_FW_SYNC.SW_MNG_SM* register bit) using the flow defined in [Section 4.6](#) to send protocol offload information to firmware.
5. Read and clear the *FWSTS.FWRI* firmware reset indication bit.
 - If a firmware reset was issued as reported in the *FWSTS.FWRI* bit, the software device driver should clear the bit and then re-initialize the protocol offload list.
6. Verify that the *HICR.En* bit (See [Section 8.2.2.16.6](#)) is set to 1b, which indicates that the shared RAM interface is available.



7. Write proxying information in the shared RAM interface located in addresses 0x15800-0x15EFF using the format defined in [Section 10.8.4](#). All addresses should be placed in networking order.
8. Once information is written into the shared RAM software should set the *HICR.C* bit to 1b.
9. Poll the *HICR.C* bit until bit is cleared by firmware indicating that command was processed and verify that command completed successfully by checking that *HICR.SV* bit was set.
10. Read firmware response from the shared RAM to verify that data was received correctly.
11. Return to [7](#). if additional commands need to be sent to firmware.
12. Release management host interface semaphore (*SW_FW_SYNC.SW_MNG_SM* register bit) using the flow defined in [Section 4.6](#).
13. Verify that a firmware reset was not initiated during the proxying configuration process by reading the *FWSTS.FWRI* firmware reset indication bit. If a firmware reset was initiated return to [1](#).
14. Set *PROXYFC.PPROXYE* bit to 1b and enable proxy.
15. The software device driver might clear the *PROXYFC.PPROXYE* bit, *PROXYS* and *PROXYFC* registers until the next time the system moves to a low power state with proxying enabled.

Once proxying is enabled by setting the *PROXYFC.PPROXYE* bit to 1b, the Intel® Xeon® Processor D-1500 Product Family LAN controller monitors incoming packets, first filtering them according to its standard address filtering method, then filtering them with all of the proxying filters enabled in the *PROXYFC* register. If a packet passes both the standard address filtering and at least one of the enabled proxying filters and does not pass any of the enabled wake-up filters, the Intel® Xeon® Processor D-1500 Product Family LAN controller:

1. Executes the relevant protocol offload for the packet and not forward the packet to the host.
2. Sets one or more bits in the Proxying Status (*PROXYS*) register according to the proxying filters matched.
 - Note that the Intel® Xeon® Processor D-1500 Product Family LAN controller sets more than one bit in the *PROXYS* register if a packet matches more than one filter.
3. Wakes the system and forwards a packet that matches the proxying filters but can't be supported to the host for further processing if configured to do so by the software device driver via the No Match command in the Set Firmware Proxying Configuration command using the shared RAM interface.

Notes:

1. When the Intel® Xeon® Processor D-1500 Product Family LAN controller is in D3 a packet that matches both one of the enabled proxying filters as defined in the *PROXYFC* register and one of the enabled wake-up filters as defined in the *WUFC* register only wakes up the system and protocol offload (proxying) does not occur.
2. Protocol offload is not executed for illegal packets with CRC errors or checksum errors and the packets are silently discarded.

5.4.3.2 Disabling Protocol Offload

3. When the Intel® Xeon® Processor D-1500 Product Family LAN controller transitions from D3-to-D0, the protocol offload is disabled. It can also be disabled by the software device driver using the following steps. It programs the proxying control and status registers:
 - a. Clears *PROXYS*.
 - b. Clears the relevant bits in *PROXYFC* to disable the offload.



5.5 DMA Coalescing

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports DMA coalescing that enables synchronizing port activity and optimizes power consumption of memory and CPU usage. DMA coalescing is a device-based feature and when the conditions to enter DMA coalescing operating mode exists (as defined in [Section 5.5.2.1](#)), the Intel® Xeon® Processor D-1500 Product Family LAN controller:


- Stops initiation of any activity on the integrated I/O interface link. This includes all DMA transactions (if *DMACR.EN_MNG_IND* is set).
- Buffers data received from the Ethernet link in internal Rx buffer until the conditions defined in [Section 5.5.2.2](#) to exit DMA coalescing exist.

5.5.1 DMA Coalescing Activation

Activating DMA coalescing functionality is done by the software device driver. The software device driver can enable or disable the device DMA coalescing functionality by clearing the *DMACR.DMAC_EN* bit.

To activate DMA coalescing functionality the following fields need to be programmed:

1. *DMCTH.DMACRXT* fields to set the per Rx packet buffer receive threshold that causes the Intel® Xeon® Processor D-1500 Product Family LAN controller to move out of DMA coalescing operating mode. When the amount of data in the receive packet buffer exceeds the receive threshold, the Intel® Xeon® Processor D-1500 Product Family LAN controller moves out of DMA coalescing operating mode. A programmed receive watermark should take into account the actual link speed and Latency Tolerance Reported (LTR) to avoid receive buffer overflow when DMA coalescing is enabled. See the LTR description in [Section 5.6](#) for the *DMACR.DMACWT* field to define a maximum time-out value for:
 - a. A receive packet to be stored in the internal receive buffer before the Intel® Xeon® Processor D-1500 Product Family LAN controller moves a packet to host memory.
 - Each time the Intel® Xeon® Processor D-1500 Product Family LAN controller enters DMA coalescing, an internal DMA coalescing watchdog timer is re-armed with the value placed in *DMACR.DMACWT*. When in DMA coalescing, the internal watchdog timer starts to count when one of the following conditions occurs:
 - An incoming Rx packet is received.
 - An interrupt is pending.
 - b. Once an interval defined in the *DMACR.DMACWT* field has passed the Intel® Xeon® Processor D-1500 Product Family LAN controller exits DMA coalescing internal buffers are flushed, RSC flows are closed and interrupts are flushed.
2. *DMCTLX.TTLX* timer field to define the time between detection of DMA idle condition and entry into DMA coalescing state. To limit entry into DMA coalescing state when packet rate is high.
3. *DMACR.DMAC_EN* bit should be set to 1b to enable activating DMA coalescing operating mode.
4. *DMACR.Lx Coalescing Indication* bit defines whether to move in/out of DMA coalescing when the integrated I/O interface moves in/out of IDLE state., when set to 1b, DMA coalescing conditions are met only when the integrated I/O interface is in of IDLE state, when set to 0b DMA coalescing can also start when the integrated I/O interface is not in IDLE state. In addition, when the bit is set to 0b, DMA coalescing stops when any TLP transactions are executed on the integrated I/O interface.
5. *DMACR.EN_MNG_IND* bit should be set to 1b to enable a management indication impact on DMA coalescing mode.
6. *DMCMNGTH.DMCMNGTHR* field to set the threshold for the management data buffer that causes a move out of DMA coalescing operating mode. The manageability indications are ignored if *DMACR.EN_MNG_IND* is cleared.



Note: Any change to a DMA coalescing configuration should be done while the feature is disabled by setting *DMAR.DMAC_EN* to 0b.

5.5.2 DMA Coalescing Operating Mode

Enabling DMA coalescing operation by setting the *DMACR.DMAC_EN* bit to 1b, enables aligning bus master traffic and interrupts from all ports. Power saving is achieved since synchronizing the integrated I/O interface accesses between ports increases the occurrence of idle intervals on the integrated I/O interface bus and also increases the duration of these idle intervals. The Power Management Unit (PMU) on a platform can use these idle intervals to reduce system power.

5.5.2.1 Conditions to Enter DMA Coalescing

The Intel® Xeon® Processor D-1500 Product Family LAN controller enters DMA coalescing when all of the following conditions exist:

1. DMA coalescing is enabled (*DMACR.DMAC_EN* = 1b).
2. Internal receive buffers are empty.
3. There are no pending DMA operations.
4. None of the conditions defined in [Section 5.5.2.2](#) to move out of DMA coalescing exist.

5.5.2.2 Conditions to Exit DMA Coalescing

When the Intel® Xeon® Processor D-1500 Product Family LAN controller is in DMA coalescing operating mode, DMA coalescing mode is exited when one of the following events occurs:

1. Amount of data in the internal receive buffers passed the *DMCTH.DMACRXT* threshold.
2. An interrupt associated to a high priority vector defined by *EITR.HIGH_PRIORITY* was detected.
3. A received packet associated to a high priority traffic class defined by *DMACR.HIGH_PRIORITY* was detected.
4. DMA coalescing watchdog timer expires as a result of the following occurrences not being serviced for the duration defined in the *DMACWT timer* field:
 - a. An incoming Rx packet is received.
 - b. An interrupt is pending.
 - c. An RSC flow was closed due to an ITR expiration.
5. DMA coalescing is disabled (*DMACR.DMAC_EN* = 0b).
6. Integrated I/O interface moves out of idle state.

Notes: Management indications are enabled through *DMACR.EN_MNG_IND* and the amount of data buffered in the management buffer exceeds *DMCMNGTH.DMCMNGTHR*.

5.5.3 DMA Coalescing Recommended Settings

These are the recommended DMA coalescing settings for the Intel® Xeon® Processor D-1500 Product Family LAN controller:

1. DMA coalescing receive threshold per traffic class (*DMCTH.DMACRXT[TC]*) = MAX{(RPB_Size[TC] - 70 μ s of Rx data), Maximum Frame Size (MFS)}.

Note: The receive threshold must be set so that its distance from the effective Rx packet buffer size enables buffering of all Rx traffic that could be received at full blown during 70 μ s, which corresponds to the worst case observed exit time from L1/L0s to L0. It means that the



threshold setting depends on the link speed. A minimum of MFS must be set for the DMA coalescing threshold.

2. DMA coalescing watchdog timer (*DMACR.DMACWT*) must be set to get a worst case packet delay of ~ 20ms.
 - a. If $TTLX \geq \text{max ITR delay over all Qs}$: $DMACWT = 0x262 \sim 20 \text{ ms}$. Flush the entire RPB to the host before re-entering a DMA coalescing state. The longest time a packet can wait in a packet buffer is *DMACWT*.
 - b. If $TTLX < \text{max ITR delay over all Qs}$: $DMACWT = (20 \text{ ms} - \text{Max ITR delay}) / 2$. The device re-enters the DMA coalescing state as soon as the first packet which is waiting in packet buffer is served. Then, when the ITR timer for the second packet expires, the device starts the *DMACWT* timer (again). It comes out that in this case, the longest time a packet can wait in packet buffer is $\text{ITR delay} + 2 \times \text{DMACWT}$.
3. Time to LX request (*DMCTLX.TTLX*) = $0x20 \sim 40 \mu\text{s}$, so that it is smaller than worst case L1/L0s max exit time ($70 \mu\text{s}$), taking advantage of the fact coalescing occurs in L0 too.
4. *DMACR.Lx_Coalescing_Indication* = 0b. For example, coalescing occurs when in L0.

DMA coalescing management threshold (*DMCMNGTH.DMCMNGTHR*) = $0x100$ (equivalent to 4 KB).

5.6 Latency Tolerance Reported (LTR)

The Intel® Xeon® Processor D-1500 Product Family LAN controller generates LTR messages to report service latency requirements for memory reads and writes to the root complex for system power management.

The Intel® Xeon® Processor D-1500 Product Family LAN controller reports either latency tolerance or no latency tolerance requirements as a function of link, LAN port and function status. The reported latency tolerance value is set to optimize platform power consumption without incurring packet loss due to receive buffer overflow.

5.6.1 LTR Algorithm

The Intel® Xeon® Processor D-1500 Product Family LAN controller sends LTR messages according to the following algorithm when the capability is enabled in the LTR capability structure of Function 0 located in the PCIe configuration space:

1. When links on all ports are disconnected or all LAN ports are disabled, the software device driver instructs the Intel® Xeon® Processor D-1500 Product Family LAN controller to send an LTR message with LTR requirement bits cleared, to indicate that no latency tolerance requirements exist.
2. If the Intel® Xeon® Processor D-1500 Product Family LAN controller reported latency tolerance requirements with any requirement bit set in the PCIe LTR message and all enabled functions were placed in D3 low power state via the *PMCSR* register, the Intel® Xeon® Processor D-1500 Product Family LAN controller sends a new LTR message with all the requirement bits clear.
3. If the Intel® Xeon® Processor D-1500 Product Family LAN controller reported latency tolerance requirements with any requirement bit set and the *LTR Mechanism Enable* bit in the PCIe configuration space is cleared, the software device driver instructs the Intel® Xeon® Processor D-1500 Product Family LAN controller to send an LTR PCIe message with all the requirement bits clear.
4. The Intel® Xeon® Processor D-1500 Product Family LAN controller holds the LTR value to send defined in LTR Control (*LTRC*) controlled and configured by the software device driver. The Intel® Xeon® Processor D-1500 Product Family LAN controller sends a new LTR message with the minimum latency tolerance requirement value upon an explicit request from the software device driver (issued by setting *LTRC.Send* bit). The following event can cause the Intel® Xeon® Processor D-1500 Product Family LAN controller's software device driver to change its latency tolerance requirement:



- a. Link speed changed.
- b. Link was disconnected.


The Intel® Xeon® Processor D-1500 Product Family LAN controller conglomerates latency requirements from the different functions and sends a single LTR message in the following manner:

- The acceptable latency values for the message sent upstream by the Intel® Xeon® Processor D-1500 Product Family LAN controller must reflect the lowest latency tolerance values associated with any function.
 - It is permitted that the snooped and non-snooped values reported in the conglomerated message are associated with different functions.
 - If none of the functions have a latency requirement for a certain type of traffic (snoop/non-snoop), the message sent by the Intel® Xeon® Processor D-1500 Product Family LAN controller does not have the *Requirement* bit corresponding to that type of traffic set.
 - The scale parameter should be set to zero by the software device driver if minimal latency is required for a specific type of traffic (latency = 0 and requirement is set).

5.6.2 LTR Initialization Flow

- The BIOS/operating system configures the maximum snoop/non-snoop platform latency in the PCIe (function 0) configuration LTR capability registers.
- As part of software device driver initialization, the software device driver:
 - Reads the LTR capability to get the worst case latency tolerance requirement for the platform that the Intel® Xeon® Processor D-1500 Product Family LAN controller is connected to.
 - Auto negotiates extended LPI exit latencies.
 - The software device driver configures receive buffer watermarks for DMAC as aggressively as possible for the active target latency tolerance value (pre defined value).
 - The software device driver configures the content of the LTR snoop/non-snoop requirements to LTRC.
- The software device driver initiates sending a LTR message from the Intel® Xeon® Processor D-1500 Product Family LAN controller by setting *LTRC.Send*.
- When a software device driver unloads, it should clear its latency tolerance requirements in the LTRC register and set the *LTRC.Send* bit, this could cause a new LTR message as defined in [Section 5.6.1](#).

§



6.0 Non-Volatile Memory Map

6.1 NVM Organization

The Intel® Xeon® Processor D-1500 Product Family LAN controller NVM contains the following high-level modules:

- **Legacy EEPROM Modules.** These modules are mapped over one of the first two 16 KB sections of the Flash device, and cannot be extended beyond them. It is composed of all the NVM modules used by MAC hardware or by manageability firmware, not including the manageability and PHY firmware code images. The sections included in this block are:
 - **NVM Pointers and Generic Words:** This section, starts at the beginning of the valid section. It contains basic device information, pointers to other NVM modules and several software configurations.
 - **PHY Module:** This module, described in [Section 6.4.3](#), is pointed to by NVM word 0x3. It contains the configuration of the PHY link layer (PCS + PMD) block and is loaded at each global (LAN_PWR_GOOD) reset.
 - **PCIe Modules:** These modules, described in [Section 6.4.4](#), contain the parameters required to configure the PCIe configuration space. The pointers to these modules are located at NVM words 0x6 (generic), 0x7 (port 0), and 0x8 (port 1).
 - **LAN Core Modules:** These two modules, described in [Section 6.4.6](#), contain parameters required to configure the LAN port. These include the MAC address, LED, and SDP configuration. The pointers to these modules are located at NVM words 0x9 (port 0) and 0xA (port 1). In addition, each LAN port has a module that enables loading specific CSR values after reset. The pointers to these modules are located at NVM words 0xD (port 0) and 0xE (port 1). These modules are loaded following a software or hardware reset.
 - **Firmware Parameters Module:** Pointed by the firmware module pointer located at NVM word 0x0F. The firmware extension module starts with a list of firmware sub-modules pointers.
 - **Boot Configuration:** This module, pointed by NVM word 0x17, contains the configuration parameters used by the PXE and iSCSI boot code.
 - **VPD:** This module, pointed by NVM word 0x2F, contains the VPD module exposed via the VPD PCIe capability as described in [Section 6.3.5](#).
- **Firmware Image.** This module contains the main code of the firmware loaded to the internal RAM. This image must fit within 512 KB and start at a 4 KB boundary. This module is authenticated upon update.
- **PCIe Expansion/Option ROM.** This module includes the PXE driver (61 KB), iSCSI boot image (116 KB), UEFI network driver (37 KB for x64, 67 KB for IA64), and can also include a CLP module (60 KB). It must fit within 512 KB and start at a 4 KB boundary. It is pointed by the PCIe expansion/option ROM pointer located at NVM word 0x05.
- **Free Space Provisioning Segment.** The NVM structure includes a space used to update the PXE code, firmware image, and PHY image modules via a double image policy. This space is referred as the free space provisioning module or segment. It must be large enough to contain the largest of these three high-level modules. It is pointed by the free space provisioning segment pointer located at NVM word 0x40. See [section 3.4.9.1](#) for the usage model of this Flash area.

Note: Flash device size can be read from NVM Control Word 1 (see [Section 6.4.2.1](#)).

Figure 6.1 shows a general NVM structure and not a required order.

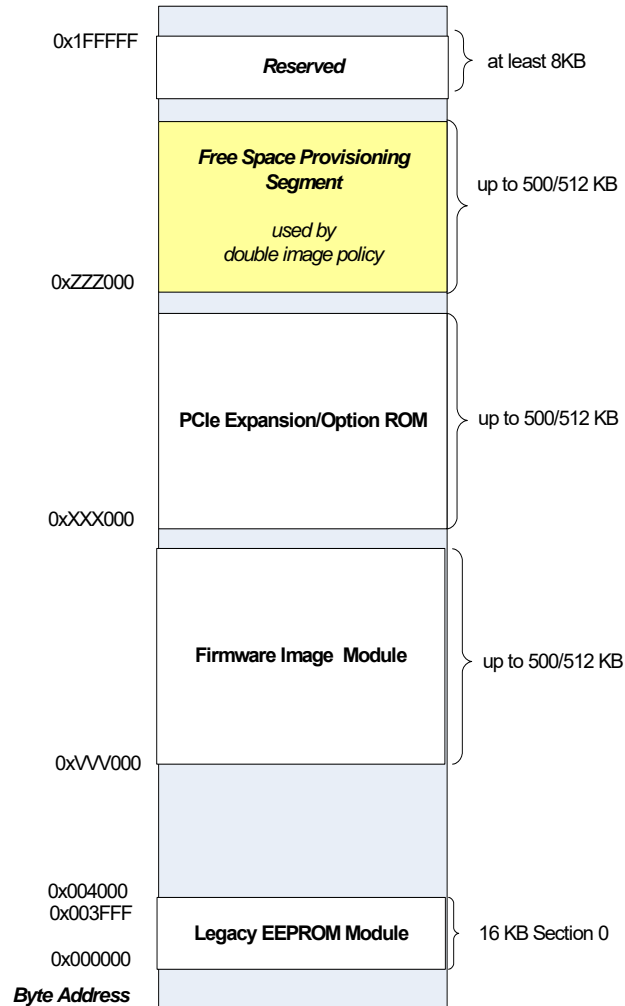


Figure 6.1. NVM Structure

6.1.1 Protected Areas

The following areas are protected from host writes:

- The firmware code area
- The pointers to the different modules.
- The mini loader
- NVM control words 1/2/3



The list of words and modules that are read only to software is listed in [Table 6.1](#).

The firmware code area and option ROM can be updated using the flow described in [Section 3.4.9.1](#).

6.2 NVM Header

Note: Intel configures the reserved NVM fields and they are not intended to be changed beyond the default image provided by Intel.

The following table lists the fixed part of the Legacy EEPROM modules used by the Intel® Xeon® Processor D-1500 Product Family LAN controller. This table lists common modules for the NVM including: hardware pointers, software and firmware. Blocks pointed in this section are detailed in the following sections. All addresses in this table are absolute in word units.

Pointers can be in word units or in 4 KB units. If bit 15 of the pointer is set, then bits [14:0] points to a 4 KB sector and the word address is {pointer[14:0], 000,0000,0000b}. If bit 15 is cleared, then the word address is pointer[14:0].

Table 6.1. NVM Header Description (Sheet 1 of 2)

Word Address	Used By	Word Name	Reference
0x0000	HW	NVM Control Word 1	Section 6.2.1.1
0x0001	HW	NVM Control Word 2	Section 6.2.1.2
0x0002	HW	Reserved/Legacy Module 1 Pointer	Section 6.2.1.3
0x0003	HW	Internal PHY Indirect List Pointer	Section 6.2.1.4
0x0004	HW	Reserved/Legacy Module 2 Pointer	Section 6.2.1.5
0x0005	SW	PCIe Expansion/Option ROM Pointer	Section 6.2.1.6
0x0006	HW	PCIe General Configuration Module Pointer	Section 6.2.1.7
0x0007	HW	PCIe Configuration Space 0 Module Pointer	Section 6.2.1.8
0x0008	HW	PCIe Configuration Space 1 Module Pointer	Section 6.2.1.9
0x0009	HW	LAN Core 0 Module Pointer	Section 6.2.1.10
0x000A	HW	LAN Core 1 Module Pointer	Section 6.2.1.11
0x000B	HW	MAC 0 Module Pointer	Section 6.2.1.12
0x000C	HW	MAC 1 Module Pointer	Section 6.2.1.13
0x000D	HW	CSR 0 Auto Configuration Module Pointer	Section 6.2.1.14
0x000E	HW	CSR 1 Auto Configuration Module Pointer	Section 6.2.1.15
0x000F	FW	Firmware Module Pointer	Section 6.2.1.16
0x0010	SW	Software Compatibility Word1	Section 6.2.1.17
0x0011	SW	Software Compatibility Word2	Section 6.2.1.18
0x0012	SW	Software Compatibility Word3	Section 6.2.1.18
0x0013	SW	Software Compatibility Word4	Section 6.2.1.18
0x0014	SW	Software Compatibility Word5	Section 6.2.1.18
0x0015	SW	PBA Word 1	Section 6.2.1.19
0x0016	SW	PBA Word 2	Section 6.2.1.19
0x0017	SW	Boot Configuration Start Address	Section 6.2.1.20
0x0018	SW	Software Reserved Word 1 - Dev Starter Version	Section 6.2.1.21
0x0019	SW	Software Reserved Word 2 - PHY Image Revision	Section 6.2.1.22



Table 6.1. NVM Header Description (Sheet 2 of 2)

Word Address	Used By	Word Name	Reference
0x001A – 0x0028	SW	Software Reserved Word (3 through 17)	Section 6.2.1.23
0x0029	SW	Software Reserved Word 18 - MAP Version	Section 6.2.1.24
0x002A	SW	Software Reserved Word 19 - IMAGE Version	Section 6.2.1.25
0x002B	SW	Software Reserved Word 20	Section 6.2.1.26
0x002C	SW	Software Reserved Word 21 - Reserved	Section 6.2.1.26
0x002D	SW	Software Reserved Word 22 - EETRACK ID 1	Section 6.2.1.27
0x002E	SW	Software Reserved Word 23 - EETRACK ID2	Section 6.2.1.28
0x002F	SW	VPD Module Pointer	Section 6.2.1.29
0x0030	SW	PXE Setup Options PCI Function 0	Section 6.2.1.30
0x0031	SW	PXE Configuration Customization Options PCI Function 0	Section 6.2.1.31
0x0032	SW	PXE Version	Section 6.2.1.32
0x0033	SW	Flash Capabilities	Section 6.2.1.33
0x0034	SW	PXE Setup Options PCI Function 1	Section 6.2.1.34
0x0035	SW	PXE Configuration Customization Options PCI Function 1	Section 6.2.1.35
0x0036	SW	iSCSI Option ROM Version	Section 6.2.1.36
0x0037	SW	Alternate Ethernet MAC Addresses Pointer	Section 6.2.2
0x0038	SW	NVM Control Word 3	Section 6.2.2.1
0x0039	SW	Reserved	Section 6.2.2.2
0x003A	FW	Firmware Code Pointer	Section 6.2.2.3
0x003B		MNG_IF_SEL LSW	Section 6.2.2.4
0x003C		MNG_IF_SEL MLSW	Section 6.2.2.5
0x003D		MNG_IF_SEL LSW	Section 6.2.2.6
0x003E	HW	MNG_IF_SEL MSW	Section 6.2.2.7
0x003F	SW	Software Checksum, Words 0x00 – 0x3F	Section 6.2.3
0x0040	HW	Free Space Provisioning Segment Pointer	Section 6.2.3.1
0x0041	HW	Free Provisioning Area Size	Section 6.2.3.2
0x0042	HW	Mini Loader Pointer	Section 6.2.3.3
0x0043	HW	Reserved	
0x0044	HW	DL Scripts Pointer	Section 6.2.26
0x0045	HW	NVM Control Word 4	Section 6.2.3.5
0x0046-0x004F	HW	Reserved	Section 6.2.3.4
0x0050	HW	RO Updates Version	Section 6.2.3.6

6.2.1 Init Module Section Summary Table

For the inner structure, see [Section 6.2](#).



6.2.1.1 NVM Control Word 1 — Address 0x0000

Bits	Name	Default	Description
15	Reserved	0	Reserved.
14	NC-SI HW ARB Disable	0x0/0x1	NC-SI Hardware Arbitration Enable. 1b = NCSI_ARB_IN and NCSI_ARB_OUT pads are not used. NCSI_ARB_IN is pulled up internally to provide stable input. 0b = NCSI_ARB_IN and NCSI_ARB_OUT pads are used.
13:12	Reserved	0x0	Reserved.
<u>11</u>	SPI 25 MHz select	1b	This bit is sampled while loading the Flash to the shadow RAM and not as part of a regular autoloading. 1b = SPI runs at 25 MHz. 0b = SPI runs at 12.5 MHz.
10:8	NVM Size	101b	These bits indicate the physical Flash size. Mapped to FLA.FL_SIZE (see field definition in the FLA register section).
7:6	Signature	01b	The <i>Signature</i> field indicates to the Intel® Xeon® Processor D-1500 Product Family LAN controller that there is a valid NVM present. If the <i>Signature</i> field is not 01b, the other bits in this word are ignored, no further NVM read is performed, and the default values are used for the configuration space IDs.
5	VPD Write Enable	1b	When set to 1b, the host can modify the entire VPD area via the Flash Update host I/F command. When set to 0b, any host attempt to modify the VPD area via the Flash Update host I/F command is completed with an error. Note: The RW-VPD area can be modified from the config space register set regardless of the setting of this field.
4:0	Reserved		Reserved.

6.2.1.2 NVM Control Word 2 — Address 0x0001

Bits	Name	Default	Description
15	PCI LAN Function Select - Swap	0x0	When both LAN ports are enabled and LAN Function Sel is set to 0b, LAN 0 is routed to PCI function 0 and LAN 1 is routed to PCI function 1. If the LAN Function Sel is set to 1b, LAN 0 is routed to PCI function 1 and LAN 1 is routed to PCI function 0. This bit is mapped to FACTPS[30].
14:11	Reserved	0x0	Reserved.
10	PCI Device OFF Enable	0b	Device Electrical Off Enable. This bit is relevant only when the device is disabled via strapping during PE_RST_N both LANn_DIS_N pins to 0b at once. 0b = Legacy mode (default). Though the device is disabled, the digital I/O pins are not moved to an electrical off state. 1b = Enable device electrical off. When the device is disabled, the digital I/O pins are put at High-Z For example, electrical off state where pull-ups/downs are at their defined values. Readable as DEV_FUNC_EN[2].
9	LAN PCI Disable Select	0x0	LAN Disable Select. 0b = LAN 0 is disabled. 1b = LAN 1 is disabled. Readable as DEV_FUNC_EN[1].



Bits	Name	Default	Description
8	LAN PCI Disable	0x0	LAN PCI Disable. When set to 1b, one LAN port is disabled. The function that is disabled is determined by the <i>LAN Disable Select</i> bit. If the disabled function is function 0, it acts as a dummy function. If the disabled port is used by MC, only the DMA block and PCIe interface of the port is powered down. Otherwise, the port is powered down up to the PHY included. If the <i>LAN PCI Disable</i> bit is set for a port, the port's respective <i>APM</i> bit in NVM Control Word 3 must be cleared as well. Readable as DEV_FUNC_EN[0].
7:4	Reserved	0x0	Reserved. Formally power management controls for memories.
3	Deadlock Timeout Enable	1b	If set, a device that was granted access to the NVM that does not toggle the interface for 2 seconds will have the grant revoked.
2	Reserved	0b	Reserved
1	Core Clocks Gate Disable	0b	During normal operation this bit should be zero enabling core clock gating in low power state. When set disables the gating of the core clock in low power state.
0	Reserved	0b	Reserved.

6.2.1.3 Reserved/Legacy Module 1 Pointer — Address 0x0002

The field Rsvd Mod1 Pointer points to the Reserved/Legacy Module 1 section. For the Reserved/Legacy Module 1 inner structure, see [Section 6.2.25](#).

Bits	Name	Default	Description
15:0	Rsvd Mod1 Pointer	0xFFFF	

6.2.1.4 Internal PHY Indirect List Pointer — Address 0x0003

The field Internal PHY List Pointer points to the Internal PHY Indirect List section. For the Internal PHY Indirect List inner structure, see [Section 6.2.24](#).

Bits	Name	Default	Description
15:0	Internal PHY List Pointer	0xFFFF	



6.2.1.5 Reserved/Legacy Module 2 Pointer — Address 0x0004

The field Rsvd Mod2 Pointer points to the Reserved/Legacy Module 2 section. For the Reserved/Legacy Module 2 inner structure, see [Section 6.2.33](#).

Bits	Name	Default	Description
15	Pointer Type	0x1	0b = Word units. 1b = 4 KB sector units.
14:0	Reserved Mod2 Pointer	0x7FFF	

6.2.1.6 PCIe Expansion/Option ROM Pointer — Address 0x0005

The field PCIe Expansion/Option ROM Pointer points to the PXE_OROM dummy CSS section. For the PXE_OROM dummy CSS inner structure, see [Section 6.2.30](#).

Bits	Name	Default	Description
15	Pointer Type	0x1	0b = Word units. 1b = 4 KB sector units.
14:0	PCIe Expansion/Option ROM Pointer	0x7FFF	

6.2.1.7 PCIe General Configuration Module Pointer — Address 0x0006

The field PCIe General Configuration Pointer points to the PCIe General Configuration Module section. For the PCIe General Configuration Module inner structure, see [Section 6.2.7](#).

Bits	Name	Default	Description
15:0	PCIe General Configuration Pointer	0xFFFF	

6.2.1.8 PCIe Configuration Space 0 Module Pointer — Address 0x0007

The field PCIe Configuration Space 0 Module Pointer points to the EEPROM PCIe Configuration Space section. For the EEPROM PCIe Configuration Space inner structure, see [Section 6.2.8](#).



Bits	Name	Default	Description
15:0	PCIe Configuration Space 0 Module Pointer	0xFFFF	

6.2.1.9 PCIe Configuration Space 1 Module Pointer — Address 0x0008

The field PCIe Configuration Space 1 Module Pointer points to the EEPROM PCIe Configuration Space section. For the EEPROM PCIe Configuration Space inner structure, see [Section 6.2.8](#).

Bits	Name	Default	Description
15:0	PCIe Configuration Space 1 Module Pointer	0xFFFF	

6.2.1.10 LAN Core 0 Module Pointer — Address 0x0009

The field LAN Core 0 Section Pointer points to the EEPROM LAN Core Module section. For the EEPROM LAN Core Module inner structure, see [Section 6.2.6](#).

Bits	Name	Default	Description
15:0	LAN Core 0 Section Pointer	0xFFFF	

6.2.1.11 LAN Core 1 Module Pointer — Address 0x000A

The field LAN Core 1 Section Pointer points to the EEPROM LAN Core Module section. For the EEPROM LAN Core Module inner structure, see [Section 6.2.6](#).

Bits	Name	Default	Description
15:0	LAN Core 1 Section Pointer	0xFFFF	

6.2.1.12 MAC 0 Module Pointer — Address 0x000B

The field MAC 0 Section Pointer points to the EEPROM MAC section. For the EEPROM MAC inner structure, see [Section 6.2.4](#).



Bits	Name	Default	Description
15:0	MAC 0 Section Pointer	0xFFFF	

6.2.1.13 MAC 1 Module Pointer — Address 0x000C

The field MAC 1 Section Pointer points to the EEPROM MAC section. For the EEPROM MAC inner structure, see [Section 6.2.4](#).

Bits	Name	Default	Description
15:0	MAC 1 Section Pointer	0xFFFF	

6.2.1.14 CSR 0 Auto Configuration Module Pointer — Address 0x000D

The field CSR 0 Auto Configuration Pointer points to the CSR Auto Config section. For the CSR Auto Config inner structure, see [Section 6.2.5](#).

Bits	Name	Default	Description
15:0	CSR 0 Auto Configuration Pointer	0xFFFF	

6.2.1.15 CSR 1 Auto Configuration Module Pointer — Address 0x000E

The field CSR 1 Auto Configuration Pointer points to the CSR Auto Config section. For the CSR Auto Config inner structure, see [Section 6.2.5](#).

Bits	Name	Default	Description
15:0	CSR 1 Auto Configuration Pointer	0xFFFF	

6.2.1.16 Firmware Module Pointer — Address 0x000F

The field Firmware Section Pointer points to the Firmware Module Header section. For the Firmware Module Header inner structure, see [Section 6.2.11](#).



Bits	Name	Default	Description
15:0	Firmware Section Pointer	0x0	

6.2.1.17 Software Compatibility Word 1 — Address 0x0010

Five words (Software Compatibility Word 1 through Word 5) in the NVM image are reserved for compatibility information. New bits within these fields are defined as the need arises for determining software compatibility between various hardware revisions.

Bits	Name	Default	Description
15:12	Reserved	0x0	Reserved
11	LOM	0b	Indicates whether the NVM attached to LAN silicon contains a dedicated module for option ROM. Used by option ROM update applications. 0b = NIC (Attached Flash contains module for option ROM). 1b = LOM (Attached Flash has no module for option ROM).
10	Server	1b	Legacy, not currently used. 1b = Server. 0b = Client.
9	Reserved	0b	Reserved
8	OEM/Retail	0b	Legacy, not currently used. 1b = OEM. 0b = Retail.
7:0	Reserved	0x0	Reserved



6.2.1.18 Software Compatibility Word 2-5 — Address 0x0010-0x0014

Bits	Name	Default	Description
15:0	Reserved	0xFFFF	Reserved

6.2.1.19 PBA Word 1-2 — Address 0x0015-0x0016

The nine-digit Printed Board Assembly (PBA) number used for Intel manufactured Network Interface Cards (NICs) is stored in the NVM.

Note that through the course of hardware ECOs, the suffix field is incremented. The purpose of this information is to enable customer support (or any user) to identify the revision level of a product.

Network driver software should not rely on this field to identify the product or its capabilities.

Current PBA numbers have exceeded the length that can be stored as hex values in these two words. For these PBA numbers the high word is a flag (0xFAFA) indicating that the PBA is stored in a separate PBA block. The low word is a pointer to a PBA block.

Bits	Name	Default	Description
15:0	PBA Word 1/PBA Word 2	0xFAFA	

6.2.1.20 Boot Configuration Start Address — Address 0x0017

Address of the iSCSI Boot configuration module. This is a word pointer. The block length is embedded in the module.

The field iSCSI Boot Configuration Start Address points to the Boot Configuration Block section. For the Boot Configuration Block inner structure, see [Section 6.2.10](#)

Bits	Name	Default	Description
15:0	iSCSI Boot Configuration Start Address	0x0	This module is 1504 bytes long and must be mapped in the first valid 4 KB sector of the Flash.

6.2.1.21 Software Reserved Word 1 — Dev Starter Version Address — 0x0018

Copy of first contents of the MAP version, not auto generated, to trace the source of any image.



Bits	Name	Default	Description
15:12	Major Version	0x0	NVM Minor Version.
11:8	Decimal	0x0	NVM Major Version. 0x0 = Secure. 0xF = Disabled.
7:0	Minor Version	0x05	Human Readable. Hex using only the decimal digits to make the digital version.

6.2.1.22 Software Reserved Word 2 — PHY Image Revision — 0x0019

Bits	Name	Default	Description
15:8	Major	0x68	PHY Major Version.
7:0	Minor	0x34	PHY Minor Version.

6.2.1.23 Software Reserved Word 3-17 — 0x0020-0x0028

Bits	Name	Default	Description
15:0	Reserved	0xFFFF	Reserved



6.2.1.24 Software Reserved Word 18 — MAP Version — 0x0029

Bits	Name	Default	Description
15:0	MAP_Version	0xFFFF	MAP Version. Auto-generated from the same input as the filename.

6.2.1.25 Software Reserved Word 19 — IMAGE Version — 0x002A

Bits	Name	Default	Description
15:0	IMAGE_Version	0xFFFF	IMAGE Version. Auto-generated from the same input as the filename.

6.2.1.26 Software Reserved Word 20-21 — 0x002B-0x002C

Bits	Name	Default	Description
15:0	Reserved	0xFFFF	Reserved

6.2.1.27 Software Reserved Word 22 — EETRACK ID 1 — 0x002D

This word is for the first word of the eTrack_ID number written by EEPROM Manager Tool.

Bits	Name	Default	Description
15:0	eTrack_ID Word 1		The EEPROM Manager Tool writes a unique 32-bit eTrack_ID number in two sequential NVM words. The eTrack_ID is written when the EEPROM Manager Tool creates an image on the Intel network. The eTrack_ID DB tracks NVM images back to a specific SCM build.



6.2.1.28 Software Reserved Word 23 — ETRACK ID 2 — 0x002E

This word is for the second word of the eTrack_ID number written by EEPROM Manager Tool.

Bits	Name	Default	Description
15:0	eTrack_ID Word 2		

6.2.1.29 Vital Product Data (VPD) Module Pointer — 0x002F

Word pointer to the VPD module. Note that the block length is embedded in the module.

The field VPD Pointer points to the VPD Module section. For the VPD Module inner structure, see Section 6.2.21.

Bits	Name	Default	Description
15:0	VPD Pointer	0xFFFF	VPD. 0xFFFF is the default unless the VPD relative section is specified. The VPD section size is usually 64 words and is initialized to 0 or 0xFFFF. During run time, this module is accessible through the VPD capability in the PCI configuration space. This module must be mapped in the first valid 4 KB sector of the Flash.

6.2.1.30 PXE Setup Options PCI Function 0 — Word Address 0x0030

The main setup options for port 0 are stored in this word. These options are those that can be changed by the user using the Control-S setup menu.

Bits	Name	Default	Description
15:13	Reserved	0x0	Reserved. Must be 0x0.
12:10	FSD	0x0	Bits 12-10 control forcing speed and duplex during driver operation. Valid values are: 000b = Auto-negotiate. 001b =Reserved. 010b = 100 Mb/s half duplex. 011b = Not valid (treated as 000b). 100b = Not valid (treated as 000b). 101b =Reserved. 110b = 100 Mb/s full duplex. 111b = 1000 Mb/s full duplex. Note: Only applicable for copper-based adapters. Not applicable to 10 GbE. Default value is 000b.
9	LWS	0b	Legacy OS Wakeup Support. (For the Intel® Xeon® Processor D-1500 Product Family LAN controller-based adapters only) If set to 1b, the agent enables PME in the adapter's PCI configuration space during initialization. This enables remote wake up under legacy operating systems that don't normally support it. Note that enabling this makes the adapter technically non-compliant with the ACPI specification, which is why the default is disabled. Must be set to 0b for 1 GbE and 10 GbE adapters. 0b = Disabled (default). 1b = Enabled.



Bits	Name	Default	Description
8	DSM	1b	Display Setup Message. If the bit is set to 1b, the "Press Control-S" message is displayed after the title message. Default value is 1b.
7:6	PT	0x0	Prompt Time. These bits control how long the Press Control-S setup prompt message is displayed, if enabled by DIM. 00b = 2 seconds (default). 01b = 3 seconds. 10b = 5 seconds. 11b = 0 seconds. Note: The Ctrl-S message is not displayed if zero seconds prompt time is selected.
5	Reserved	0b	Reserved. Must be 0x0.
4:3	DBS	0x0	Default Boot Selection. These bits select which device is the default boot device. These bits are only used if the agent detects that the BIOS does not support boot order selection or if the MODE field of word 0x31 is set to MODE_LEGACY. 00b = Network boot, then local boot (default). 01b = Local boot, then network boot. 10b = Network boot only. 11b = Local boot only.
2:0	PS	0x0	Protocol Select. These bits select the active boot protocol. 00b = PXE (default value). 01b = RPL (only if RPL is in the Flash). 10b = iSCSI boot primary port (only if iSCSI boot is using this device). 11b = iSCSI boot secondary port (only if iSCSI boot is using this device). Only the default value of 00b should be initially programmed into the device. Other values should only be set by configuration utilities.



Bits 2:0 are defined as follows:

Bit(s)	Value	Port Status	CLP (Combo) Executes	iSCSI Boot Option ROM CTRL-D Menu
2:0	0	PXE	PXE	Displays port as PXE. Allows changing to boot disabled, iSCSI primary or secondary.
	1	Boot Disabled	NONE	Displays port as disabled. Allows changing to iSCSI primary/secondary.
	2	iSCSI Primary	iSCSI	Displays port as iSCSI primary. Allows changing to boot disabled, iSCSI secondary.
	3	iSCSI Secondary	iSCSI	Displays port as iSCSI secondary. Allows changing to boot disabled, iSCSI primary.
	4	Reserved	Reserved	Reserved.
	7:5	Reserved	Same as disabled.	Same as disabled.

6.2.1.31 PXE Configuration Customization Options PCI Function 0 - Word Address 0x0031

Word 0x0031 of the NVM contains settings that can be programmed by an OEM or network administrator to customize the operation of the software. These settings cannot be changed from within the Control-S setup menu. The lower byte contains settings that would typically be configured by a network administrator using an external utility; these settings generally control which setup menu options are changeable. The upper byte is generally settings that would be used by an OEM to control the operation of the agent in a LOM environment, although there is nothing in the agent to prevent their use on a NIC implementation. The default value for this word is 0x4000.

Bits	Name	Default	Description
15:14	Signature	0x1	Signature. Must be set to 01b to indicate that this word has been programmed by the agent or other configuration software.
13:12	Reserved	0x0	Reserved. Must be 0x0.
11	Continuous Retry	0b	Selects Continuous Retry operation. If this bit is set, IBA does NOT transfer control back to the BIOS if it fails to boot due to a network error (such as failure to receive DHCP replies). Instead, it restarts the PXE boot process again. If this bit is set, the only way to cancel PXE boot is for the user to press ESC on the keyboard. Retry is not attempted in case of hardware error conditions such as an invalid NVM checksum or failing to establish link. Default value is 0b.



Bits	Name	Default	Description
10:8	Operating Mode	0x0	Selects the agent's boot order setup mode. This field changes the agent's default behavior in order to make it compatible with systems that do not completely support the BBS and PnP Expansion ROM standards. Valid values and their meanings are: 000b = Normal behavior. The agent attempts to detect BBS and PnP Expansion ROM support as it normally does. 001b = Force legacy mode. The agent does not attempt to detect BBS or PnP Expansion ROM support in the BIOS and assumes the BIOS is not compliant. The user can change the BIOS boot order in the Setup Menu. 010b = Force BBS mode. The agent assumes the BIOS is BBS-compliant, even though it might not be detected as such by the agent's detection code. The user CANNOT change the BIOS boot order in the Setup Menu. 011b = Force PnP Int18 mode. The agent assumes the BIOS allows boot order setup for PnP Expansion ROMs and hooks interrupt 0x18 (to inform the BIOS that the agent is a bootable device) in addition to registering as a BBS IPL device. The user CANNOT change the BIOS boot order in the Setup Menu. 100b = Force PnP Int19 mode. The agent assumes the BIOS allows boot order setup for PnP Expansion ROMs and hook interrupt 0x19 (to inform the BIOS that the agent is a bootable device) in addition to registering as a BBS IPL device. The user CANNOT change the BIOS boot order in the Setup Menu. 101b = Reserved for future use. If specified, is treated as a value of 000b. 110b = Reserved for future use. If specified, is treated as a value of 000b. 111b = Reserved for future use. If specified, is treated as a value of 000b.
7:6	Reserved	0x0	Reserved. Must be 0x0.
5	Disable Flash Update	0b	Disable Flash Update. If this bit is set to 1b, the user is not allowed to update the Flash image using PROSet. Default value is 0b.
4	Disable Legacy OS Wakeup Menu	0x0	Disable Legacy Wakeup Support. If this bit is set to 1b, the user is not allowed to change the Legacy OS Wakeup Support menu option. Default value is 0b.
3	Disable Boot Select Menu	0b	Disable Boot Selection. If this bit is set to 1b, the user is not allowed to change the boot order menu option. Default value is 0b.
2	Disable Protocol Selection Menu	0b	Disable Protocol Select. If set to 1b, the user is not allowed to change the boot protocol. Default value is 0b.
1	Disable Title Message Display	0b	Disable Title Message. If this bit is set to 1b, the title message displaying the version of the Boot Agent is suppressed; the Control-S message is also suppressed. This is for OEMs who do not wish the boot agent to display any messages at system boot. Default value is 0b.
0	Disable Setup Menu	0b	Disable Setup Menu. If this bit is set to 1b, the user is not allowed to invoke the setup menu by pressing Control-S. In this case, the NVM can only be changed via an external program. Default value is 0b.

6.2.1.32 PXE Version — Word Address 0x0032

Word 0x0032 of the NVM is used to store the version of the boot agent that is stored in the Flash image. When the boot agent loads, it can check this value to determine if any first-time configuration needs to be performed. The boot agent then updates this word with its version. Some diagnostic tools report the version of boot agent in the Flash also read this word.

Bits	Name	Default	Description
15:12	Major Version	0x0	PXE Boot Agent Major Version. Default value is 0x0.
11:8	Minor Version	0x0	PXE Boot Agent Minor Version. Default value is 0x0.
7:0	Build Number	0x0	PXE Boot Agent Build Number. Default value is 0x0.

6.2.1.33 Flash Capabilities — Word Address 0x0033

Word 0x0033 of the NVM is used to enumerate the boot technologies that have been programmed into the Flash. This is updated by Flash configuration tools and is not updated or read by IBA.



Bits	Name	Default	Description
15:14	Signature	0x1	Signature. Must be set to 01b to indicate that this word has been programmed by the agent or other configuration software.
13:6	Reserved	0x0	Reserved. Must be 0x0.
5	Reserved	0x0	Reserved.
4	iSCSI Boot	0b	iSCSI is present if set to 1b.
3	EFI/UNDI Driver	1b	EFI UNDI driver is present if set to 1b.
2	RPL	0b	RPL module is present if set to 1b. Reserved bit for devices.
1	PXE/UNDI Driver	1b	PXE UNDI driver is present if set to 1b.
0	PXE Base Code	1b	PXE base code is present if set to 1b.

6.2.1.34 PXE Setup Options PCI Function 1 — Word Address 0x0034

This word is the same as word 0x0030, but for PCIe function 1 of the device.

6.2.1.35 PXE Configuration Customization Options PCI Function 1 — Word Address 0x0035

This word is the same as word 0x0031, but for PCIe function 1 of the device.

6.2.1.36 iSCSI Option ROM Version — Word Address 0x0036

Word 0x0036 of the NVM is used to store the version of updated iSCSI option ROM. The value must be above 0x2000. The value below (word 0x1FFF = 16 KB NVM size) is reserved for future expansion for a pointer to a combo option ROM component version structure. iSCSIUtil, FLAUtil, DMiX update iSCSI option ROM version if the value is above 0x2000, 0x0000, or 0xFFFF. The pointer (0x0040 - 0x1FFF) should be kept and not be overwritten.

Bits	Name	Default	Description
15:0	Reserved	0xFFFF	

6.2.2 Alternate Ethernet MAC Address Pointer — Word Address 0x0037

The alternate MAC address does not function if the pointer value is set to 0xFFFF.

Word offset for PCIe function 0 MAC address is: value in word 0x0037 + 0 - 3 words.

The word offset for the port 0 MAC address is: value in word 0x37 + 0 ; - 3 words. The word offset for the port 1 MAC address is: value in word 0x37 + 3 ; - Next 3 words.


The field Pointer points to Alternate Ethernet MAC Address section. For the Alternate Ethernet MAC Address inner structure, see [Section 6.2.2](#).



Bits	Name	Default	Description
15:0	Pointer	0xFFFF	

6.2.2.1 NVM Control Word 3 — Address 0x38

Bits	Name	Default	Description
15	WD_FWRST_EN	1b	Enable FW Reset on Watchdog Timeout. Defines the default value of the FWRSTCTR.WD_FWRST_EN bit. Should be set for normal operation
14	ECC_FWRST_EN	1b	Enable FW Reset on ECC Error. Defines the default value of the FWRSTCTR.ECC_FWRST_EN bit. Should be set for normal operation
13:9	Reserved	0x0	Reserved
8	D10GMP Port 1	1b	Disable 10 GbE in LPLU for LAN Port 1. When set, LAN port 1 never advertises 10 GbE speed capability when in LPLU state (D3/Dr).
7	Enable FW LPLU	0b	Disable 1 GbE in LPLU for LAN Port 1. When set, LAN port 1 never advertises 1 GbE speed capability when in LPLU state (D3/Dr). If set, D10GMP bit must be set as well
6	D10GMP Port 0	1b	Disable 10 GbE in LPLU for LAN Port 0. When set, LAN port 0 never advertises 10 GbE speed capability when in LPLU state (D3/Dr).
5:4	Reserved	0x0	Disable 1 GbE in LPLU for LAN Port 0. When set, LAN port 0 never advertises 1 GbE speed capability when in LPLU state (D3/Dr). If set, D10GMP bit must be set as well.
3	Enable HW LPLU	0b	Enable LPLU. When set, enables a decrease in link speed of the port defined to stay awake in non-D0a states when power policy and power management states dictate it.



Bits	Name	Default	Description
2	Keep_PHY_Link_Up_En	1b	Enables No PHY Link Down when the MC indicates that the PHY should be kept on. When asserted, this bit prevents changes in power management state to be reflected to the PHYs according to the MMNGC.MNG_VETO bit value. When cleared, the MMNGC.MNG_VETO bit is meaningless.
1	APM Enable Port 1	1b	Initial value of advanced power management wake up enable in the General Receive Control register (GRC.APME). Mapped to GRC.APME of port 1.
0	APM Enable Port 0	1b	Initial value of advanced power management wake up enable in the General Receive Control register (GRC.APME). Mapped to GRC.APME of port 0.

6.2.2.2 Reserved Pointer — 0x0039

6.2.2.3 Firmware Code Pointer — 0x003A

The field Firmware Code Pointer points to the firmware code dummy CSS Section. For the firmware code dummy CSS inner structure, see [Section 6.2.28](#).

Bits	Name	Default	Description
15	Pointer Type	0x1	0b = Words in units. 1b = 4 KB sector units.
14:0	Firmware Code Pointer	0x0	



6.2.2.4 MNG_IF_SEL LSW — 0x003B

Bits	Name	Default	Description
15:8	Reserved	0x0	Reserved
7:3	MDIO_PHY_ADDR	0x0	
2:0	I2C MDIO EN	0x0	0x0 = None. 0x1 = I2C EN. 0x2 = MDIO 45. 0x6 = MDIO 22. 0x7 = Legacy.

6.2.2.5 MNG_IF_SEL MSW — 0x003C

Bits	Name	Default	Description
15:10	Link Mode	0x1	0x01 = Reserved. 0x04 = Ext module QSFP+. 0x06 = Ext module QSFP+ plus re-driver. 0x08 = Ext module SFP+. 0x0A = Ext module SFP+ plus re-driver. 0x10 = 10GBASE-T. 0x20 = 1GBASE-T. 0x3F = Legacy.
9	Reserved	0x0	Reserved
8	10GBASE-T Internal PHY Mode	0x0	0x0 = Reserved. 0x1 = KR.
7	Reserved	0x0	100 Mb/s Enable. 0x0 = Disabled. 0x1 = Enabled.
6:0	Reserved	0x0	0x0 = Reserved. 0x1 = Legacy.



6.2.2.6 MNG_IF_SEL LSW — 0x003D

For inner structure, see Section 6.2.2.6

6.2.2.7 MNG_IF_SEL LSW — 0x003E

For inner structure, see Section 6.2.2.7

6.2.3 Software Checksum — Words 0x0000 - 0x003F

The checksum word (0x003F) is used to ensure that the base NVM image is a valid image. It covers contents of all the NVM modules located in the first valid 4 KB sector, except for all firmware modules.

The value of this word should be calculated such that after adding all the concerned words, including the checksum word itself, the sum should be 0xBABA. This word is used strictly by software. Hardware does not calculate nor check its content but rather checks the *Signature* field in the NVM Control Word 1.

Bits	Name	Default	Description
15:0	Checksum		See the psedo code that follows.

The following shows the calculation:

```
#define IXGBE_EEPROM_CHECKSUM 0x3F
#define IXGBE_EEPROM_SUM 0xBABA
#define IXGBE_PCIE_ANALOG_PTR 0x03
#define IXGBE_PHY_PTR 0x04
#define IXGBE_OPTION_ROM_PTR 0x05
#define IXGBE_FW_PTR 0x0F
/**
 * ixgbe_calc_eeprom_checksum_X540 - Calculates and returns the
 * checksum
 * @hw: pointer to hardware structure
 */
u16 ixgbe_calc_eeprom_checksum_X540(struct ixgbe_hw *hw)
{
    u16 i;
    u16 j;
    u16 checksum = 0;
    u16 length = 0;
    u16 pointer = 0;
    u16 word = 0;
    DEBUGFUNC("ixgbe_calc_eeprom_checksum_X540");
    /* Include 0x0-0x3F in the checksum */
    for (i = 0; i < IXGBE_EEPROM_CHECKSUM; i++) {
        if (hw->eeprom.ops.read(hw, i, &word) != IXGBE_SUCCESS) {
            DEBUGOUT("EEPROM read failed\n");
            break;
        }
        checksum += word;
    }
    /*
     * Include all data from pointers 0x3, 0x6-0xE. This excludes the
     * FW, PHY module, and PCIe Expansion/Option ROM pointers.
     */
}
```



```

for (i = IXGBE_PCIE_ANALOG_PTR; i < IXGBE_FW_PTR; i++) {
    if (i == IXGBE_PHY_PTR || i == IXGBE_OPTION_ROM_PTR)
        Continue;
    if (hw->eeprom.ops.read(hw, i, &pointer) != IXGBE_SUCCESS) {
        DEBUGOUT("EEPROM read failed\n");
        break;
    }
    /* Skip pointer section if the pointer is invalid. */
    if (pointer == 0xFFFF || pointer == 0 ||
        pointer >= hw->eeprom.word_size)

        continue;
    if (hw->eeprom.ops.read(hw, pointer, &length) != IXGBE_SUCCESS)
    {
        DEBUGOUT("EEPROM read failed\n");
        break;
    }
    /* Skip pointer section if length is invalid. */
    if (length == 0xFFFF || length == 0) ||
        (pointer + length) >= hw->eeprom.word_size)
        continue;
    for (j = pointer+1; j <= pointer+length; j++) {
        if (hw->eeprom.ops.read(hw, j, &word) != IXGBE_SUCCESS) {
            DEBUGOUT("EEPROM read failed\n");
            break;
        }
        checksum += word;
    }
}
checksum = (u16)IXGBE_EEPROM_SUM - checksum;
return checksum;
}
    
```

6.2.3.1 Free Space Provisioning Segment Pointer — 0x0040

The field Free Space Provisioning Segment Pointer points to the Free Space Provisioning Module_512 section. For the Free Space Provisioning Module_512 inner structure See [Section 6.2.32](#).

Bits	Name	Default	Description
15	Pointer Type	0x1	
14:0	Free Space Provisioning Segment Pointer	0x0	



6.2.3.2 Free Provisioning Area Size — 0x0041

Bits	Name	Default	Description
15:0	Free Provisioning Area Size	0x7A	

6.2.3.3 Mini Loader Pointer — 0x0042

The field Mini Loader Section Pointer points to the Mini Loader Module section. For the Mini Loader Module inner structure, see [Section 6.2.23](#).

Bits	Name	Default	Description
15:0	Mini Loader Section Pointer	0x0	

6.2.3.4 DL Scripts Pointer — 0x0044

Field DL Scripts Section Pointer points to DL Scripts Module Section. For DL Scripts Module inner structure see [Section 6.2.23](#).

Bits	Name	Default	Description
15:0	DL Scripts Section Pointer	0xFFFF	

6.2.3.5 NVM Control Word 4 — 0x0045

Bits	Name	Default	Description
15:1	Reserved	0x7FFFF	
0	BMC_CSR_WR_EN	0x0	



6.2.3.6 RO Updates Version — 0x0050

Bits	Name	Default	Description
15:0	RO Updates Version	0x0	

6.2.4 EEPROM MAC Section Summary Table

Word Offset	Description	Page
0x0000	MAC Configuration	158
0x0001	MAC Configuration Extended	159



6.2.4.1 MAC Configuration — 0x0001

Bits	Name	Default	Description
15:14	Reserved	0x0	Reserved
13:10	FIFOLT	0x3	FIFO High Threshold. Determines the high threshold of the MAC elastic FIFO. Mapped to MACC.FIFOLT.
9:6	FIFOHT	0xD	FIFO High Threshold. Determines the high threshold of the MAC elastic FIFO. Mapped to MACC.FIFOHT.
5	Swap Rx Control	0x0	Swap the MAC internal Rx XGMII controls. 0b = Swap disabled (normal mode). 1b = Swap enabled. Mapped to MACC.Swap Rx Control.
4	SWAP Rx Data	0x0	Swap the MAC internal Rx XGMII lanes. 0b = Swap disabled (normal mode). 1b = Swap enabled. Mapped to MACC.Swap Rx Data.
3	Swizzle Rx Data	0x0	Swizzle the bytes in all four MAC internal Rx XGMII lanes. 0b = Swizzle disabled (normal mode). 1b = Swizzle enabled. Mapped to MACC.Swizzle Rx Data.
2	Swap Tx Control	0x0	Swap the MAC internal Tx XGMII controls. 0b = Swap disabled (normal mode). 1b = Swap enabled. Mapped to MACC.Swap Tx Control.
1	SWAP Tx Data	0x0	Swap the MAC internal Tx XGMII lanes. 0b = Swap disabled (normal mode). 1b = Swap enabled. Mapped to MACC.Swap Tx Data.
0	Swizzle Tx Data	0x0	Swizzle the bytes in all four MAC internal Tx XGMII lanes. 0b = Swizzle disabled (normal mode). 1b = Swizzle enabled. Mapped to MACC.Swizzle Tx Data.



6.2.4.2 MAC Configuration Extended — 0x0002

Bits	Name	Default	Description
15:3	Reserved	0x0	Reserved
2	LDINFITR	0b	Link Down In Fatal Interrupt. 1b = Link indication is gated with the fatal interrupt from the PHY. When the PHY asserts a fatal interrupt, the link is down. 0b = Normal operation. The link indication is the same as comes from the PCS and auto-negotiation machinery.
1:0	BLKXGMII	0x0	Block XGMII interface. 00b = Normal operation. The XGMII interface fed from the PCS while link is down. 01b = Idle state. The XGMII interface fed with idles while link is down. 10b = Local fault state. The XGMII interface fed with local fault sequences while link is down. 11b = Reserved.

6.2.5 CSR Auto Config Section Summary Table

Word Offset	Description	Page
0x0000	Section Length	159
0x0001	0x11178 NW_MNG_IF_SEL addr>>2	160
0x0002	NW_MNG_IF_SEL LSW	160
0x0003	NW_MNG_IF_SEL MSW	160
0x0004	0x11108 MDIO_LED_SEL addr>>2	161
0x0005	NW_MNG_IF_SEL LSW	161
0x0006	NW_MNG_IF_SEL MSW	161
0x0007	CSR RAQ1	161
0x0008	CSR RAW2	161

6.2.5.1 Section Length

The section length word contains the length of the section in words. Note that section length does not include a count for the section length word. Section Length = 3*n (n = number of CSRs to configure).

Bits	Name	Default	Description
15:0	Section Length		Length in: 2 Bytes unit - 1. First Section -> Word: CSR Auto Config -> Section length. Last Section -> Word: CSR Auto Config -> CSR RAW2. Section length in words.



6.2.5.2 0x11178 NW_MNG_IF_SEL addr>>2 — 0x0001

Bits	Name	Default	Description
15:0	0x11178 NW_MNG_IF_SEL addr>>2	0x445E	

6.2.5.3 NW_MNG_IF_SEL LSW — 0x0002

Bits	Name	Default	Description
15:8	Reserved	0x0	Reserved
7:3	MDIO_PHY_ADDR	0x0	
2:0	I2C MDIO EN	0x0	0x0 = None. 0x1 = I2C EN. 0x2 = MDIO 45. 0x6 = MDIO 22.

6.2.5.4 NW_MNG_IF_SEL MSW — 0x0003

Bits	Name	Default	Description
15:10	Link Mode	0x01	0x01 = Reserved. 0x04 = Ext Module QSFP+. 0x06 = Ext Module QSFP+ plus re-driver. 0x08 = Ext Module SFP+ 0x0A = Ext Module SFP+ plus re-driver. 0x10 = 10GBASE-T. 0x20 = 1GBASE-T.
9:0	Reserved	0x0	Reserved



6.2.5.5 0x11108 MDIO_LED_SEL addr>>2 — 0x0004

Bits	Name	Default	Description
15:0	0x11108 NW_MNG_IF_SEL addr>>2	0x4442	

6.2.5.6 MDIO_LED_SEL LSW — 0x0005

Bits	Name	Default	Description
15:2	Reserved	0x0	Reserved
1	MDIO EN Port 1	0b	0x0 = Disabled. 0x1 = Enabled.
0	MDIO EN Port 0	0b	0x0 = Disabled. 0x1 = Enabled.

6.2.5.7 MDIO_LED_SEL MSW — 0x0006

Bits	Name	Default	Description
15:0	Reserved	0x0	Reserved

6.2.5.8 CSR RAW1 — 0x0007

The section length word contains the length of the section in words. Note that section length does not include a count for the section length word. Section Length = 3*n (n = number of CSRs to configure).

6.2.5.9 CSR RAW 2 — 0x0008

The section length word contains the length of the section in words. Note that section length does not include a count for the section length word. Section Length = 3*n (n = number of CSRs to configure).



6.2.6 EEPROM LAN Core Module Section Summary Table

Word Offset	Description	Page
0x0000	Section Length	162
0x0001	Ethernet MAC Address Register1	162
0x0002	Ethernet MAC Address Register2	163
0x0003	Ethernet MAC Address Register3	163
0x0004	LED Control Lower Word	163
0x0005	LED Control Upper Word	165
0x0006	SDP Control	166
0x0007	Filter Control	167

6.2.6.1 Section Length

The section length word contains the length of the section in words. Note that section length does not include a count for the section length word.

Bits	Name	Default	Description
15:0	Section Length		Length in: 2 Bytes unit - 1. First Section -> Word: EEPROM LAN Core Module -> Section length. Last Section -> Word: EEPROM LAN Core Module -> Filter Control. Section length in words.

6.2.6.2 Ethernet MAC Address Register1 — 0x0001

The Ethernet Individual Address (IA) is a 6-byte field that must be unique for each NIC or LOM and must also be unique for each copy of the NVM image. The first three bytes are vendor specific. For example, the IA is equal to [00 AA 00] or [00 A0 C9] for Intel products. The value of this field is loaded into the Receive Address register 0 (RAL0/RAH0).

For the purpose of this document, the numbering convention is as follows:

Vendor	1	2	3	4	5	6
Intel Original	00	AA	00	Variable	Variable	Variable
Intel New	00	A0	C9	Variable	Variable	Variable



Bits	Name	Default	Description
15:8	Eth_Addr_Byte2	0x00	Ethernet MAC address Byte 2.
7:0	Eth_Addr_Byte1	0xC9	Ethernet MAC address Byte 1.

6.2.6.3 Ethernet MAC Address Register2 — 0x0002

Bits	Name	Default	Description
15:8	Eth_Addr_Byte4	0x00	Ethernet MAC address Byte 4.
7:0	Eth_Addr_Byte3	0xC9	Ethernet MAC address Byte 3.

6.2.6.4 Ethernet MAC Address Register3 — 0x0003

Bits	Name	Default	Description
15:8	Eth_Addr_Byte6	0x00	Ethernet MAC address Byte 6.
7:0	Eth_Addr_Byte5	0xC9	Ethernet MAC address Byte 5.

6.2.6.5 LED Control Lower Word — 0x0004

The LEDCTL register defaults are loaded from two words as listed in the tables that follow.



Bits	Name	Default	Description
15	LED1_BLINK	1b	0b = Don't blink. 1b = Blink.
14	LED1_IVRT	0b	0b = Active low. 1b = Active high.
13:12	Reserved	0x0	Reserved
11:8	LED1_MODE	0x1	LED 0 Control. 0x0 = LINK_UP. 0x1 = LINK 10 GbE. 0x2 = MAC ACTIVITY. 0x3 = FILTER ACTIVITY. 0x4 = LINK/ACTIVITY. 0x5 = LINK 1 GbE. 0x6 = LINK 100 Mb/s. 0xE = LED ON. 0xF = LED OFF.
7	LED0_BLINK	0b	0b = Don't blink. 1b = Blink.
6	LED0_IVRT	0b	0b = Active low. 1b = Active high.
5	Global Link Mode	0b	0b = 200 ms 1b = 83 ms
4	Reserved	0b	Reserved
3:0	LED0_MODE	0x0	LED 0 Control. 0x0 = LINK_UP. 0x1 = LINK 10 GbE. 0x2 = MAC ACTIVITY. 0x3 = FILTER ACTIVITY. 0x4 = LINK/ACTIVITY. 0x5 = LINK 1 GbE. 0x6 = LINK 100 Mb/s. 0xE = LED ON. 0xF = LED OFF.



6.2.6.6 LED Control Upper Word — 0x0005

Bits	Name	Default	Description
15	LED3_BLINK	0b	0b = Don't blink. 1b = Blink.
14	LED3_IVRT	0b	0b = Active low. 1b = Active high.
13:12	Reserved	0x0	Reserved
11:8	LED3_MODE	0x5	LED 0 Control. 0x0 = LINK_UP. 0x1 = LINK 10 GbE. 0x2 = MAC ACTIVITY. 0x3 = FILTER ACTIVITY. 0x4 = LINK/ACTIVITY. 0x5 = LINK 1 GbE. 0x6 = LINK 100 Mb/s. 0xE = LED ON. 0xF = LED OFF.
7	LED2_BLINK	1b	0b = Don't blink. 1b = Blink.
6	LED2_IVRT	1b	0b = Active low. 1b = Active high.
5:4	Reserved	0b	Reserved
3:0	LED2_MODE	0x4	LED 0 Control. 0x0 = LINK_UP. 0x1 = LINK 10 GbE. 0x2 = MAC ACTIVITY. 0x3 = FILTER ACTIVITY. 0x4 = LINK/ACTIVITY. 0x5 = LINK 1 GbE. 0x6 = LINK 100 Mb/s. 0xE = LED ON. 0xF = LED OFF.



6.2.6.7 SDP Control — 0x0006

Bits	Name	Default	Description
15	SDP3_NATIVE	0b	Defines SDP1 operating mode that is mapped to ESDP.SDP1_NATIVE and loaded at power up: 0b = Operates as generic software controlled I/O. 1b = Native mode operation (hardware function).
14	SDP2_NATIVE	0b	Defines SDP1 operating mode that is mapped to ESDP.SDP1_NATIVE and loaded at power up: 0b = Operates as generic software controlled I/O. 1b = Native mode operation (hardware function).
13	SDP1_NATIVE	0b	Defines SDP1 operating mode that is mapped to ESDP.SDP1_NATIVE and loaded at power up: 0b = Operates as generic software controlled I/O. 1b = Native mode operation (hardware function).
12	SDP0_NATIVE	0b	Defines SDP1 operating mode that is mapped to ESDP.SDP1_NATIVE and loaded at power up: 0b = Operates as generic software controlled I/O. 1b = Native mode operation (hardware function).
11	SDPDIR[3]	0b	SDP3 Pin. Initial direction is mapped to ESDP.SDP3_IODIR and loaded at power up.
10	SDPDIR[2]	0b	SDP2 Pin. Initial direction is mapped to ESDP.SDP2_IODIR and loaded at power up.
9	SDPDIR[1]	0b	SDP1 Pin. Initial direction is mapped to ESDP.SDP1_IODIR and loaded at power up.
8	SDPDIR[0]	0b	SDP0 Pin. Initial direction is mapped to ESDP.SDP0_IODIR and loaded at power up.
7	Reserved	0b	Reserved
6	REVERSE_DIR	0b	When set, reverses the polarity of the pin direction signal to the pads block such that pads are fed with the output enable signal instead of output disable signal. Should be cleared for normal operation.



Bits	Name	Default	Description
5	SDP23_function	0b	Defines the initial value for ESDP.SDP23_function. 0b = 1588 functionality 1b = I ² C functionality. Relevant only if SDP2_NATIVE or SDP3_NATIVE is set. If this bit is set, then SDP2_NATIVE and SDP3_NATIVE should have the same value.
4	SDP1_function	0b	Defines the initial value for ESDP.SDP1_function. 0b = Time sync functionality. 1b = Reserved. SDPDIR[1] should be configured as an output. Relevant only if SDP1_NATIVE is set.
3	SDPVAL[3]	0b	SDP3 Pin. Initial output value is mapped to ESDP.SDP3_DATA and loaded at power up.
2	SDPVAL[2]	0b	SDP2 Pin. Initial output value is mapped to ESDP.SDP2_DATA and loaded at power up.
1	SDPVAL[1]	0b	SDP1 Pin. Initial output value is mapped to ESDP.SDP1_DATA and loaded at power up.
0	SDPVAL[0]	0b	SDP0 Pin. Initial output value is mapped to ESDP.SDP0_DATA and loaded at power up.

6.2.6.8 Filter Control — 0x0007

Bits	Name	Default	Description
15:0	Reserved	0x0001	Reserved

6.2.7 PCIe General Configuration Module Section Summary Table

Word Offset	Description	Page
0x0000	PCI_CNF2 Low	168
0x0001	PCI_CNF2 High	169
0x0002	PCI_LBARCTRL L	169
0x0003	PCI_LBARCTRL H	170
0x0004	PCI_SERL 1	170
0x0005	PCI_SERL 2	170



Word Offset	Description	Page
0x0006	PCI_SERL 3	171
0x0007	PCI_SERL 4	171
0x0008	PCI_CAPCTRL L	171
0x0009	PCI_CAPCTRL H	171
0x000A	PCI_CAPSUP L	172
0x000B	PCI_CAPSUP H	173
0x000C	PCI_DBGCTL L	173
0x000D	PCI_DBGCTL H	174
0x000E	PCI_UPADD L	174
0x000F	PCI_UPADD H	174
0x0010	PCI_SUBSYSID L	174
0x0011	PCI_SUBSYSID H	175
0x0012	PCI_PWRDATA L	175
0x0013	PCI_PWRDATA H	175
0x0014	PCI_REVID L	175
0x0015	PCI_REVID H	175
0x0016	PCI_VFSUP L	176
0x0017	PCI_VFSUP H	176
0x0018	TPH_CTRL L	176
0x0019	TPH_CTRL H	176
0x001A	PCI_LINKCAP L	177
0x001B	PCI_LINKCAP H	177
0x001C	PCI_PMSUP L	177
0x001D	PCI_PMSUP H	178
0x001E	PCI_GLBL_CNF L	178
0x001F	PCI_GLBL_CNF H	179
0x0020	PCI_VENDORID L	179
0x0021	PCI_VENDORID H	179
0x0022	PCI_PCIERR L	179
0x0023	PCI_PCIERR H	179

6.2.7.1 PCI_CNF2 Low — 0x0000

MSIx and cache line configuration.

Bits	Name	Default	Description
15:0	PCI_CNF2 Low	0x400040FC	



6.2.7.2 PCI_CNF2 High — 0x0001

MSIx and cache line configuration.

Bits	Name	Default	Description
15:0	PCI_CNF2 High		

6.2.7.3 PCI_CNF2 High — 0x0002

PF BAR registers configuration.

Bits	Name	Default	Description
15:14	Reserved	0x0	Reserved
13:11	EXROM_BAR	0x3	0x0 = 64 KB. 0x1 = 128 KB. 0x2 = 256 KB. 0x3 = 512 KB. 0x4 = 1 MB. 0x5 = 2 MB. 0x6 = 4 MB. 0x7 = 8 MB.
10:9	Reserved	0x0	Reserved
8:6	FL_BAR_SIZE	0x5	0x5 = 2 MB. 0x6 = 4 MB. 0x7 = 8 MB.
5:4	Reserved	0x0	Reserved
3	FLASH_EXPOSE	1b	0x0 = Disabled. 0x1 = Enabled.



Bits	Name	Default	Description
2	CRSIZE	0b	0x0 = Disabled. 0x1 = Enabled.
1	BAR32	0b	0x0 = 64BitAddressing. 0x1 = 32BitAddressing.
0	PREFBAR	1b	0x0 = NoPrefech. 0x1 = PrefetchEna.

6.2.7.4 PCI_LBARCTRL H — 0x0003

Bits	Name	Default	Description
15:0	PCI_LBARCTRL H		

6.2.7.5 PCI_SER 1 — 0x0004

Bits	Name	Default	Description
15:0	PCI_SER 1	0x0	

6.2.7.6 PCI_SER 2 — 0x0005

Bits	Name	Default	Description
15:0	PCI_SER 2	0x0	



6.2.7.7 PCI_SER 3 — 0x0006

Bits	Name	Default	Description
15:0	PCI_SER 3	0x0	

6.2.7.8 PCI_SER 4 — 0x0007

Bits	Name	Default	Description
15:0	PCI_SER 4	0x0	

6.2.7.9 PCI_CAPCTRL L — 0x0008

Bits	Name	Default	Description
15:1	PCI_CAPCTRL L	0x0	
0	VPD_EN	1b	

6.2.7.10 PCI_CAPCTRL H — 0x0009

Bits	Name	Default	Description
15:0	PCI_CAPCTRL H	0x0	



6.2.7.11 PCI_CAPSUP L — 0x000A

Bits	Name	Default	Description
15	ECRC_MCTP_GEN	0b	ECRC Generation for MCTP. 0b = Do not add ECRC to MCTP packets even if ECRC is enabled. 1b = Add ECRC to MCTP packets if ECRC is enabled via the <i>ECRC Generation Enable</i> field in PCIe Advanced Error Capabilities and Control register
14:8	Reserved	0x0	Reserved
7	SEC_EN	0b	A value of 1b indicates support for the Secondary PCI Express Extended Capability. 0x0 = Disabled. 0x1 = Enabled.
6	ACS_EN	1b	A value of 1b indicates support for the PCIe ACS Capability. 0b = Disabled. 1b = Enabled.
5	IOV_EN	1b	A value of 1b indicates support for the PCIe SR-IOV Capability. 0b = Disabled. 1b = Enabled.
4	ARI_En	1b	A value of 1b indicates support for the PCIe ARI Capability. 0b = Disabled. 1b = Enabled.
3	TPH_EN	0b	A value of 1b indicates support for the PCIe TPH Requester Capability. 0b = Disabled. 1b = Enabled.
2	LTR_EN	1b	A value of 1b indicates support for the PCIe Latency Tolerance Reporting (LTR) Capability. 0b = Disabled. 1b = Enabled.
1	Reserved	0b	Reserved
0	PCIE_VER	1b	0b = 0x1-Capability version. 1b = 0x2-Capability version.



6.2.7.12 PCI_CAPSUP H — 0x000B

Bits	Name	Default	Description
15	LOAD_DEV_ID	1b	Load Device ID. When set to 1b, indicates that the device loads its PCI device IDs from the NVM. 0b = Disabled. 1b = Enabled.
14	LOAD_SUBSYS_ID	1b	Load Subsystem IDs. When set to 1b, indicates that the device loads its PCIe subsystem ID and subsystem vendor ID from the NVM. 0b = Disabled. 1b = Enabled.
13:5	Reserved	0x0	Reserved
4	CSR_CONF_EN	1b	Enables access to CSRs via the PCI Configuration Space. See the section on Configuration Access to Internal Registers and Memories for more detail. 0b = Disabled. 1b = Enabled.
3	MSI_MASK	1b	MSI Per-vector Masking Setting. This bit is loaded to the masking bit (bit 8) in the Message Control of the MSI Configuration Capability structure. 0b = Disabled. 1b = Enabled.
2	IDO_EN	1b	Enables ID-based Ordering (IDO). 0b = Disabled. 1b = Enabled.
1	ECRC_CHK_EN	1b	Loaded into the ECRC Check Capable bit of the PCIe Configuration registers. 0b = Disabled. 1b = Enabled.
0	ECRC_GEN_EN	1b	Loaded into the ECRC Generation Capable bit of the PCIe Configuration registers. 0b = Disabled. 1b = Enabled.

6.2.7.13 PCI_DBGCTL L — 0x000C

Bits	Name	Default	Description
15:0	PCI_DBGCTL L	0x1	



6.2.7.14 PCI_DBGCTL H — 0x000D

Bits	Name	Default	Description
15:0	PCI_DBGCTL H		

6.2.7.15 PCI_UPADD L — 0x000E

Bits	Name	Default	Description
15:1	ADDRESS	0x0	
0	MODE	0x0	

6.2.7.16 PCI_UPADD H — 0x000F

Bits	Name	Default	Description
15:0	ADDRESS		

6.2.7.17 PCI_SUBSYSID L — 0x0010

Bits	Name	Default	Description
15:0	SUB_VEN_ID	0x8086	



6.2.7.18 PCI_SUBSYSID H — 0x0011

Bits	Name	Default	Description
15:0	SUB_ID	0x0	

6.2.7.19 PCI_PWRDATA L — 0x0012

Bits	Name	Default	Description
15:0	PCI_PWRDATA L	0x0	

6.2.7.20 PCI_PWRDATA H — 0x0013

Bits	Name	Default	Description
15:0	PCI_PWRDATA H	0x0	

6.2.7.21 PCI_REVID L — 0x0014

Bits	Name	Default	Description
15:0	PCI_REVID L	0x0	

6.2.7.22 PCI_PWRDATA H — 0x0015

Bits	Name	Default	Description
15:0	PCI_REVID H	0x0	



6.2.7.23 PCI_VFSUP L — 0x0016

Bits	Name	Default	Description
15:0	PCI_VFSUP L	0x0	

6.2.7.24 PCI_VFSUP H — 0x0017

Bits	Name	Default	Description
15:0	PCI_VFSUP H	0x0	

6.2.7.25 TPH_CTRL L — 0x0018

Bits	Name	Default	Description
15:0	TPH_CTRL L	0x00001800	

6.2.7.26 TPH_CTRL H — 0x0019

Bits	Name	Default	Description
15:0	TPH_CTRL H		



6.2.7.27 PCI_LINKCAP L — 0x001A

Bits	Name	Default	Description
15:13	Reserved	0x0	Reserved
12:9	LINK_MAX_WIDTH	0x7	0x0 = Reserved. 0x1 = x1 limit. 0x2 = reserved. 0x3 = x4 limit. 0x4 = x8 limit. 0x5 = Reserved. 0x6 = Reserved. 0x7 = No limit. 0xF:0x8 = Reserved.
8:6	LINK_MAX_PAYLOAD	0x2	0x0 = 128 bytes. 0x1 = 256 bytes. 0x2 = 512 bytes. 0x3 = 1024 bytes. 0x4 = 2048 bytes. 0x7:0x5 = Reserved
5:0	Reserved	0x0	Reserved

6.2.7.28 PCI_LINKCAP H — 0x001B

Bits	Name	Default	Description
15:0	PCI_LINKCAP H		

6.2.7.29 PCI_PMSUP L— 0x001C

Bits	Name	Default	Description
15:0	PCI_PMSUP L		



6.2.7.30 PCI_PMSUP H — 0x001D

Bits	Name	Default	Description
15:0	PCI_PMSUP H		

6.2.7.31 PCI_GLBL_CNF L — 0x001E

Bits	Name	Default	Description
15:8	PCIE_CLKGATE-TIME	0x10	Clock Gating Idle Timer. This field defines the number of clocks (CSR clock) of idle-detect before gating the PCIe clocks.
7:6	Reserved	0x0	Reserved
5	LINK_MAX_PAYLOAD	0b	When set to 1b, and if PCIE_CLKGATE_DIS is set to 0b, the PCIe PCLK is also dynamically gated.
4	PCIE_CLKGATE_L1_ONLY	0b	When set to 1b, and if PCIE_CLKGATE_DIS is set to 0b, the clock gating of the PCIe clocks is only when PCIe is in the L1 state.
3	PCIE_CLKGAT_DIS	0b	When set to 0b, enables dynamic clock gating of the PCIe clocks (LCB, HIU and CSR).
2	WAKE_PIN_EN	0b	When set to 1b, enables the use of the PE_WAKE_N pin for a PME event in all power states; otherwise, it is enabled only in Dr state.
1:0	Reserved	0x0	Reserved



6.2.7.32 PCI_GLBL_CNF H — 0x001F

Bits	Name	Default	Description
15:0	Reserved	0x0	Reserved

6.2.7.33 PCI_VENDORID L — 0x0020

Bits	Name	Default	Description
15:0	PCI_VENDORID L	0X8086	

6.2.7.34 PCI_VENDORID H — 0x0021

Bits	Name	Default	Description
15:0	PCI_VENDORID H	0x0	

6.2.7.35 PCI_PCIEER L — 0x0022

Bits	Name	Default	Description
15:0	PCI_PCIEER L	0x3FCB	

6.2.7.36 PCI_PCIEER H — 0x0023

Bits	Name	Default	Description
15:0	PCI_PCIEER H	0x000B	



6.2.8 EEPROM PCIe Configuration Space Section Summary Table

Word Offset	Description	Page
0x0000	PCI_PFDEVID Low (LAN)	180
0x0001	Reserved	180
0x0002	PCI_CNF Low	181
0x0003	PCI_CNF High	181
0x0004	PCI_VFDEVID Low (LAN)	181
0x0005	Reserved	182
0x0006	PCI_CLASS Low	182
0x0007	PCI_CLASS High	182

6.2.8.1 PCI_PFDEVID Low — 0x0000

Bits	Name	Default	Description
15:0	PCI_PFDEVID L	0x1563	If the load device ID in offset 0x0007 in the PCIe general configuration section is set, this word is loaded to the device ID of the LAN function.

6.2.8.2 Reserved — 0x0001

Bits	Name	Default	Description
15:0	Reserved	0x1563	Reserved



6.2.8.3 PCI_CNF Low — 0x0002

Bits	Name	Default	Description
15:7	Reserved	0x0	Reserved
6:5	INT_PIN	0x0	Controls the value advertised in the <i>Interrupt Pin</i> field of the PCI configuration header for this function. Values of 00b, 01b, 10b and 11b correspond to INTA#, INTB#, INTC# and INTD#, respectively. The value advertised in the PCI configuration header is the value loaded from the NVM + 1.
4	IO_BAR	0b	I/O BAR Support. 0b = I/O BAR not supported. 1b = I/O BAR supported.
3	EXROM_DIS	0b	Expansion ROM Disable. 0b = The expansion ROM BAR in the PCI configuration space is enabled. 1b = The expansion ROM BAR in the PCI configuration space is disabled.
2:0	Reserved	0x0	Reserved

6.2.8.4 PCI_CNF High — 0x0003

Bits	Name	Default	Description
15:0	Reserved	0x0	Reserved

6.2.8.5 PCI_VFDEVID Low — 0x0004

Bits	Name	Default	Description
15:0	VF_DEV_ID_LAN	0x1565	If the load device ID in offset 0x0007 in the PCIe general configuration section is set, this word is loaded to the device ID of the LAN function.



6.2.8.6 Reserved — 0x0005

Bits	Name	Default	Description
15:0	Reserved	0x1563	Reserved

6.2.8.7 PCI_CLASS Low — 0x0006

Bits	Name	Default	Description
15:1	Reserved	0x0	Reserved
0	STORAGE_CLASS	0b	0b = The class code of this port is set to 0x020000 (LAN). 1b = The class code of this port is set to 0x010000 (SCSI).

6.2.8.8 PCI_CLASS High — 0x0007

Bits	Name	Default	Description
15:1	Reserved	0x0	Reserved
0	STORAGE_CLASS	0b	0b = The class code of this port is set to 0x020000 (LAN). 1b = The class code of this port is set to 0x010000 (SCSI).

6.2.9 Alternate Ethernet MAC Address Section Summary Table

Word Offset	Description	Page
0x0000	MAC Address Word 1 Port 0	183
0x0001	MAC Address Word 2 Port 0	183
0x0002	MAC Address Word 3 Port 0	183
0x0003	MAC Address Word 1 Port 1	183
0x0004	MAC Address Word 2 Port 1	183
0x0005	MAC Address Word 3 Port 1	184



6.2.9.1 MAC Address Word 1 Port 0 — 0x0000

Bits	Name	Default	Description
15:0	MAC Address Word 1 Port 0	0xFFFF	

6.2.9.2 MAC Address Word 2 Port 0 — 0x0001

Bits	Name	Default	Description
15:0	MAC Address Word 2 Port 0	0xFFFF	

6.2.9.3 MAC Address Word 3 Port 0 — 0x0002

Bits	Name	Default	Description
15:0	MAC Address Word 3 Port 0	0xFFFF	

6.2.9.4 MAC Address Word 1 Port 1 — 0x0003

Bits	Name	Default	Description
15:0	MAC Address Word 1 Port 1	0xFFFF	

6.2.9.5 MAC Address Word 2 Port 1 — 0x0004

Bits	Name	Default	Description
15:0	MAC Address Word 2 Port 1	0xFFFF	



6.2.9.6 MAC Address Word 3 Port 1 — 0x0005

Bits	Name	Default	Description
15:0	MAC Address Word 3 Port 1	0xFFFF	


6.2.10 Boot Configuration Block Section Summary Table

Contains the required setup used for boot operations.

Word Offset	Description	Page
0x0000	Boot Signature	186
0x0001	Block Size	186
0x0002	Structure Version	186
0x0003	iSCSI Initiator Name	186
0x0083 + 1*n, n=0...16	Reserved	187
0x0094	iSCSI Flags	187
0x0095 + 1*n, n=0...1	iSCSI Initiator IP	187
0x0097 + 1*n, n=0...1	Subnet Mask	187
0x0099 + 1*n, n=0...1	Gateway IP	188
0x009B	iSCSI Boot LUN	188
0x009C + 1*n, n=0...1	iSCSI Target IP	188
0x009E	iSCSI Target Port	188
0x009F	iSCSI Target Name	188
0x011F	CHAP Password	188
0x0128	CHAP User Name	188
0x0168	VLAN ID	189
0x0169	Mutual CHAP Password	189
0x0173 + 1*n, n=0...2	Reserved	189
0x0176	Target Worldwide Port Name - WWPN	189
0x017A	Boot LUN	189
0x017B	VLAN ID	189
0x017C	Target Boot Order	189
0x017D	Reserved	189
0x017E	Target Worldwide Port Name - WWPN	189
0x0182	Boot LUN	190
0x0183	VLAN ID	190
0x0184	Target Boot Order	190
0x0185	Reserved	190
0x0186	Target Worldwide Port Name - WWPN	190
0x018A	Boot LUN	191



Word Offset	Description	Page
0x018B	VLAN ID	191
0x018C	Target Boot Order	191
0x018D	Reserved	191
0x018E	Target Worldwide Port Name - WWPN	191
0x0192	Boot LUN	192
0x0193	VLAN ID	192
0x0194	Target Boot Order	192
0x0195 + 1*n, n=0...44	Reserved	192
0x01C2	iSCSI Flags	192
0x01C3 + 1*n, n=0...1	iSCSI Initiator IP	192
0x01C5 + 1*n, n=0...1	Subnet Mask	192
0x01C7 + 1*n, n=0...1	Gateway IP	192
0x01C9	iSCSI Boot LUN	192
0x01CA + 1*n, n=0...1	iSCSI Target IP	192
0x01CC	iSCSI Target Port	192
0x01CD	iSCSI Target Name	193
0x024D	CHAP Password	193
0x0256	CHAP User Name	193
0x0296	VLAN ID	193
0x0297	Mutual CHAP Password	193
0x02A1 + 1*n, n=0...2	Reserved	193
0x02A4	Target Worldwide Port Name - WWPN	193
0x02A8	Boot LUN	193
0x02A9	VLAN ID	193
0x02AA	Target Boot Order	193
0x02AB	Reserved	193
0x02AC	Target Worldwide Port Name - WWPN	193
0x02B0	Boot LUN	193
0x02B1	VLAN ID	193
0x02B2	Target Boot Order	194
0x02B3	Reserved	194
0x02B4	Target Worldwide Port Name - WWPN	194
0x02B8	Boot LUN	194
0x02B9	VLAN ID	194
0x02BA	Target Boot Order	194
0x02BB	Reserved	194
0x02BC	Target Worldwide Port Name - WWPN	194
0x02C0	Boot LUN	194
0x02C1	VLAN ID	194
0x02C2	Target Boot Order	194
0x02C3 + 1*n, n=0...44	Reserved	194



6.2.10.1 Boot Signature (0x0000)

Bits	Name	Default	Description
15:0	Boot Signature	0x5369	Boot Signature: 'I', 'S' (0x3569)

6.2.10.2 Block Size (0x0001)

Bits	Name	Default	Description
15:0	Block Size		Length in: 1-byte unit. First Section -> Word: Boot Configuration Block -> Boot Signature Last Section -> Word: Boot Configuration Block -> Reserved

6.2.10.3 Structure Version (0x0002)

Bits	Name	Default	Description
15:8	Reserved	0x00	Reserved.
7:0	Structure Version	0x01	Version of this structure. Should be set to 0x1.

6.2.10.4 iSCSI Initiator Name (0x0003)

Raw data module length: 128 words.

iSCSI initiator name. This field is optional and built by manual input, DHCP host name, or with MAC address.



6.2.10.5 Reserved[n] (0x0083 + 1*n, n=0...16)

Bits	Name	Default	Description
15:0	Reserved	0x0000	Reserved for future use. Should be set to 0x0.

6.2.10.6 iSCSI Flags (0x0094)

Bits	Name	Default	Description
15:10	ARP Timeout	0x0F	Timeout value for each try.
9:8	ARP Retries	01b	Retry value.
7:6	Reserved	00b	Reserved.
5:4	Ctrl-D Entry	00b	Ctrl-D Entry. 00b = Enabled 01b = Reserved 10b = Reserved 11b = Disabled — Skip Ctrl-D entry.
3:2	Authentication Type	00b	Authentication Type. 00b = None 01b = One Way CHAP 10b = Mutual CHAP 11b = Reserved
1	Enable DHCP for iSCSI Target	1b	Enable DHCP for getting iSCSI target information. 0b = Disabled — Use static target configuration. 1b = Enabled — Use DHCP to get target information by the Option 17 Root Path.
0	Enable DHCP	1b	Enable DHC. 0b = Disabled — Use static configurations from this structure. 1b = Enabled — Overrides configurations retrieved from DHCP.

6.2.10.7 iSCSI Initiator IP[n] (0x0095 + 1*n, n=0...1)

Bits	Name	Default	Description
15:0	iSCSI Initiator IP	0x0000	Initiator DHCP flag. Not set = This field should contain the initiator IP address. Set = This field is ignored.

6.2.10.8 Subnet Mask[n] (0x0097 + 1*n, n=0...1)

Bits	Name	Default	Description
15:0	iSCSI Initiator IP	0x0000	Initiator DHCP flag. Not set = This field should contain the subnet mask. Set = This field is ignored.



6.2.10.9 Gateway IP[n] (0x0099 + 1*n, n=0...1)

Bits	Name	Default	Description
15:0	iSCSI Initiator IP	0x0000	Initiator DHCP flag. Not set = This field should contain the gateway IP address. Set = This field is ignored.

6.2.10.10 iSCSI Boot LUN (0x009B)

Bits	Name	Default	Description
15:0	iSCSI Boot LUN	0x0000	Target DHCP flag. Not set = iSCSI target LUN number should be specified. Set = This field is ignored.

6.2.10.11 iSCSI Target IP[n] (0x009C + 1*n, n=0...1)

Bits	Name	Default	Description
15:0	iSCSI Initiator IP	0x0000	Target DHCP flag. Not set = IP address of iSCSI target. Set = This field is ignored.

6.2.10.12 iSCSI Target Port (0x009E)

Bits	Name	Default	Description
15:0	iSCSI Target Port	0x0CBC	Target DHCP flag Not set = TCP port used by iSCSI target. Default is 3260. Set = This field is ignored.

6.2.10.13 iSCSI Target Name (0x009F)

Raw data module length: 128 words.

Target DHCP flag.

Not set = iSCSI target name should be specified.

Set = This field is ignored.

6.2.10.14 CHAP Password (0x011F)

Raw data module length: 9 words.

The minimum CHAP secret must be 12 octets and maximum CHAP secret size is 16. The last 2 bytes are null alignment padding.

6.2.10.15 CHAP User Name (0x0128)

Raw data module length: 64 words.

The user name must be non-null value and maximum size of user name allowed is 127 characters.



6.2.10.16 VLAN ID (0x0168)

Bits	Name	Default	Description
15:0	VLAN ID	0x0001	Reserved area, since the function is disabled due to Microsoft restrictions. VLAN ID to include the tag in iSCSI boot frames. A valid VLAN ID is between 1 and 4094. Zero means no VLAN tag support.

6.2.10.17 Mutual CHAP Password (0x0169)

Raw data module length: 9 words.

The minimum mutual CHAP secret must be 12 octets and maximum mutual CHAP secret size is 16. The last 2 bytes are null alignment padding.

Bits	Name	Default	Description
15:0	Reserved	0x0000	Reserved for future use. Should be set to 0x0.

6.2.10.18 Target Worldwide Port Name - WWPN (0x0176)

Raw data module length: 4 words.

Byte string of target WWPN.

6.2.10.19 Boot LUN (0x017A)

Bits	Name	Default	Description
15:0	Target LUN	0x0000	Target LUN.

6.2.10.20 VLAN ID (0x017B)

Bits	Name	Default	Description
15:0	VLAN ID	0x0000	VLAN ID for the Port. Default is 0x0.

6.2.10.21 Target Boot Order (0x017C)

Bits	Name	Default	Description
15:8	Reserved	0x00	Reserved for future use. Should be set to 0x0.
7:0	Target Boot Order	0x00	Target Boot Order. Valid range is 0-4, with 0 meaning no boot order.


6.2.10.22 Reserved (0x017D)

Bits	Name	Default	Description
15:0	Reserved	0x0000	Reserved.

6.2.10.23 Target Worldwide Port Name - WWPN (0x017E)

Raw data module length: 4 words.

Byte string of target WWPN.



6.2.10.24 Boot LUN (0x0182)

Bits	Name	Default	Description
15:0	Target LUN	0x0000	Target LUN.

6.2.10.25 VLAN ID (0x0183)

Bits	Name	Default	Description
15:0	VLAN ID	0x0000	VLAN ID for the Port. Default is 0x0.

6.2.10.26 Target Boot Order (0x0184)

Bits	Name	Default	Description
15:8	Reserved	0x00	Reserved for future use. Should be set to 0x0.
7:0	Target Boot Order	0x00	Target Boot Order. Valid range is 0-4, with 0 meaning no boot order.

6.2.10.27 Reserved (0x0185)

Bits	Name	Default	Description
15:0	Reserved	0x0000	Reserved.

6.2.10.28 Target Worldwide Port Name - WWPN (0x0186)

Raw data module length: 4 words.

Byte string of target WWPN.



6.2.10.29 Boot LUN (0x018A)

Bits	Name	Default	Description
15:0	Target LUN	0x0000	Target LUN.

6.2.10.30 VLAN ID (0x018B)

Bits	Name	Default	Description
15:0	VLAN ID	0x0000	VLAN ID for the Port. Default is 0x0.

6.2.10.31 Target Boot Order (0x018C)

Bits	Name	Default	Description
15:8	Reserved	0x00	Reserved for future use. Should be set to 0x0.
7:0	Target Boot Order	0x00	Target Boot Order. Valid range is 0-4, with 0 meaning no boot order.

6.2.10.32 Reserved (0x018D)

Bits	Name	Default	Description
15:0	Reserved	0x0000	Reserved.

6.2.10.33 Target Worldwide Port Name - WWPN (0x018E)

Raw data module length: 4 words.

Byte string of target WWPN.



6.2.10.34 Boot LUN (0x0192)

Bits	Name	Default	Description
15:0	Target LUN	0x0000	Target LUN.

6.2.10.35 VLAN ID (0x0193)

Bits	Name	Default	Description
15:0	VLAN ID	0x0000	VLAN ID for the Port. Default is 0x0.

6.2.10.36 Target Boot Order (0x0194)

Bits	Name	Default	Description
15:8	Reserved	0x00	Reserved for future use. Should be set to 0x0.
7:0	Target Boot Order	0x00	Target Boot Order. Valid range is 0-4, with 0 meaning no boot order.

6.2.10.37 Reserved[n] (0x0195 + 1*n, n=0...44)

Bits	Name	Default	Description
15:0	Reserved	0x0000	Reserved.

6.2.10.38 iSCSI Flags (0x01C2)

For detail on the inner structure, see [Section 6.2.10.6](#).

6.2.10.39 iSCSI Initiator IP[n] (0x01C3 + 1*n, n=0...1)

For detail on the inner structure, see [Section 6.2.10.7](#).

6.2.10.40 Subnet Mask[n] (0x01C5 + 1*n, n=0...1)

For detail on the inner structure, see [Section 6.2.10.8](#).

6.2.10.41 Gateway IP[n] (0x01C7 + 1*n, n=0...1)

For detail on the inner structure, see [Section 6.2.10.9](#).

6.2.10.42 iSCSI Boot LUN (0x01C9)

For detail on the inner structure, see [Section 6.2.10.10](#).

6.2.10.43 iSCSI Target IP[n] (0x01CA + 1*n, n=0...1)

For detail on the inner structure, see [Section 6.2.10.11](#).

6.2.10.44 iSCSI Target Port (0x01CC)

For detail on the inner structure, see [Section 6.2.10.12](#).



6.2.10.45 iSCSI Target Name (0x01CD)

For detail on the inner structure, see [Section 6.2.10.13](#).

6.2.10.46 CHAP Password (0x024D)

For detail on the inner structure, see [Section 6.2.10.14](#).

6.2.10.47 CHAP User Name (0x0256)

For detail on the inner structure, see [Section 6.2.10.15](#).

6.2.10.48 VLAN ID (0x0296)

For detail on the inner structure, see [Section 6.2.10.16](#).

6.2.10.49 Mutual CHAP Password (0x0297)

For detail on the inner structure, see [Section 6.2.10.17](#).

6.2.10.50 Reserved[n] (0x02A1 + 1*n, n=0...2)

For detail on the inner structure, see [Section 6.2.10.17](#).

6.2.10.51 Target Worldwide Port Name - WWPN (0x02A4)

For detail on the inner structure, see [Section 6.2.10.18](#).

6.2.10.52 Boot LUN (0x02A8)

For detail on the inner structure, see [Section 6.2.10.19](#).

6.2.10.53 VLAN ID (0x02A9)

For detail on the inner structure, see [Section 6.2.10.20](#).

6.2.10.54 Target Boot Order (0x02AA)

For detail on the inner structure, see [Section 6.2.10.21](#).

6.2.10.55 Reserved (0x02AB)

For detail on the inner structure, see [Section 6.2.10.22](#).

6.2.10.56 Target Worldwide Port Name - WWPN (0x02AC)


For detail on the inner structure, see [Section 6.2.10.23](#).

6.2.10.57 Boot LUN (0x02B0)

For detail on the inner structure, see [Section 6.2.10.24](#).

6.2.10.58 VLAN ID (0x02B1)

For detail on the inner structure, see [Section 6.2.10.25](#).



6.2.10.59 Target Boot Order (0x02B2)

For detail on the inner structure, see [Section 6.2.10.26](#).

6.2.10.60 Reserved (0x02B3)

For detail on the inner structure, see [Section 6.2.10.27](#).

6.2.10.61 Target Worldwide Port Name - WWPN (0x02B4)

For detail on the inner structure, see [Section 6.2.10.28](#).

6.2.10.62 Boot LUN (0x02B8)

For detail on the inner structure, see [Section 6.2.10.29](#).

6.2.10.63 VLAN ID (0x02B9)

For detail on the inner structure, see [Section 6.2.10.30](#).

6.2.10.64 Target Boot Order (0x02BA)

For detail on the inner structure, see [Section 6.2.10.31](#).

6.2.10.65 Reserved (0x02BB)

For detail on the inner structure, see [Section 6.2.10.32](#).

6.2.10.66 Target Worldwide Port Name - WWPN (0x02BC)

For detail on the inner structure, see [Section 6.2.10.33](#).

6.2.10.67 Boot LUN (0x02C0)

For detail on the inner structure, see [Section 6.2.10.34](#).

6.2.10.68 VLAN ID (0x02C1)

For detail on the inner structure, see [Section 6.2.10.35](#).

6.2.10.69 Target Boot Order (0x02C2)

For detail on the inner structure, see [Section 6.2.10.36](#).

6.2.10.70 Reserved[n] (0x02C3 + 1*n, n=0...44)

For detail on the inner structure, see [Section 6.2.10.37](#).

6.2.11 Firmware Module Header Section Summary Table

Word Offset	Description	Page
0x0000	Section Header	195
0x0001	Test Configuration Pointer	195
0x0002	Common Firmware Parameters Pointer	195
0x0003	Pass Through LAN 0 Configuration Pointer	195



Word Offset	Description	Page
0x0004	SideBand Configuration Pointer	196
0x0005	Flexible TCO Filter Configuration Pointer	196
0x0006	Pass Through LAN 1 Configuration Pointer	196
0x0007	Reserved	197
0x0008	LESM Configuration Pointer	197
0x0009	Section Footer	197

6.2.11.1 Section Header — 0x0000

Bits	Name	Default	Description
15:0	Block Length		Length in words of the section covered by CRC. Block length in words.

6.2.11.2 Test Configuration Pointer — 0x0001

The field Test Configuration Pointer points to the Test Configuration Module section. For the Test Configuration Module inner structure, see [Section 6.2.12](#).

Bits	Name	Default	Description
15:0	Test Configuration Pointer	0x0	

6.2.11.3 Common Firmware Parameters Pointer — 0x0002

The field Common Firmware Parameters Pointer points to the Common Firmware Parameters Module section. For the Common Firmware Parameters Module inner structure, see [Section 6.2.13](#).

Bits	Name	Default	Description
15:0	Common Firmware Parameters Pointer	0x0	

6.2.11.4 Pass Through LAN 0 Configuration Pointer — 0x0003

The field Pass Through LAN 0 Configuration Pointer points to the Pass Through Control Words section. For the Pass Through Control Words inner structure, see [Section 6.2.14](#).



Bits	Name	Default	Description
15:0	Pass Through LAN 0 Configuration Pointer	0x0	

6.2.11.5 SideBand Configuration Pointer — 0x0004

This module is 28 bytes long and must be mapped in the first valid 4 KB sector of the Flash.

The field SideBand Configuration Pointer points to the Sideband Configuration Structure section. For the Sideband Configuration Structure inner structure, see [Section 6.2.15](#).

Bits	Name	Default	Description
15:0	SideBand Configuration Pointer	0x0	

6.2.11.6 Flexible TCO Filter Configuration Pointer — 0x0005

This section loads all flexible filters. The control + mask + filter data are repeatable as the number of filters. Section length in offset 0 is for all filters.

The field Flexible TCO Filter Configuration Pointer points to the Flexible TCO Filter Configuration Structure section. For the Flexible TCO Filter Configuration inner structure, see [Section 6.2.19](#).

Bits	Name	Default	Description
15:0	Flexible TCO Filter Configuration Pointer	0x0	

6.2.11.7 Pass Through LAN 1 Configuration Pointer — 0x0006

The field Pass Through LAN 1 Configuration Pointer points to the Pass Through Control Words section. For the Pass Through Control Words inner structure, see [Section 6.2.14](#).

Bits	Name	Default	Description
15:0	Pass Through LAN 1 Configuration Pointer	0x0	



6.2.11.8 Reserved — 0x0007

Bits	Name	Default	Description
15:0	Reserved	0x0	

6.2.11.9 LESM Configuration Pointer — 0x0008

The field LESM Configuration Pointer points to the LESM Configurations section. For the LESM Configurations inner structure, see [Section 6.2.20](#).

Bits	Name	Default	Description
15:0	LESM Configuration Pointer	0x0	

6.2.11.10 Section Footer — 0x0009

Bits	Name	Default	Description
15:8	Block CRC8		CRC-8-CCITT: Start Section -> Word: Firmware Module Header -> Section Header End Section -> Word: Firmware Module Header -> LESM Configuration Pointer
7:0	Reserved	0x0	Block length in words.

6.2.12 Test Configuration Module Section Summary Table

Word Offset	Description	Page
0x0000	Section Header	198
0x0001	Reserved	198
0x0002	Reserved	198
0x0003	Section Footer	199



6.2.12.1 Section Header — 0x0000

Bits	Name	Default	Description
15:0	Block Length		Length in words of the section covered by CRC. Block length in words.

6.2.12.2 Reserved — 0x0001

Bits	Name	Default	Description
15:9	Reserved	0x1	Reserved
8	Reserved	0b	Reserved
7:0	Reserved	0xC6	Reserved

6.2.12.3 Reserved — 0x0002

Bits	Name	Default	Description
15:1	Reserved	0x0001	Reserved
0	Reserved	0b	Reserved



6.2.12.4 Section Footer — 0x0003

Bits	Name	Default	Description
15:8	Block CRC8		CRC-8-CCITT: Start Section -> Word: Test Configuration Module -> Section Header End Section -> Word: Test Configuration Module -> Miscellaneous
7:0	Reserved	0x0	Block length in words.

6.2.13 Common Firmware Parameters Module Section Summary Table

Word Offset	Description	Page
0x0000	Section Header	199
0x0001	Common Firmware Parameters 1	200
0x0002	Common Firmware Parameters 2	201
0x0003	Section Footer	201

6.2.13.1 Section Header — 0x0000

Bits	Name	Default	Description
15:0	Block Length		Length in words of the section covered by CRC. Block length in words.



6.2.13.2 Common Firmware Parameters 1 — 0x0001

Bits	Name	Default	Description
15	Enable Firmware Reset	1b	0b = Disabled. 1b = Enabled.
14:13	Redirection Sideband Interface	0x0	0x0 = SMBus. 0x1 = NC-SI. 0x2 = MCTP, over SMBus only. 0x3 = MCTP, over PCIe and SMBus.
12	Restore MAC Address	1b	0x0 = Do not restore MAC address at power on. 0x1 = Restore MAC address at power on.
11	Reserved	0b	Reserved
10:8	Manageability Mode	0x2	0x0 = None 0x1 = Reserved. 0x2 = Pass Through (PT) mode. 0x7:0x3 = Reserved.
7:6	Secondary Sideband Interface	0x0	0x0 = None 0x1 = Reserved. 0x2 = Pass Through (PT) mode. 0x7:0x3 = Reserved.
5:3	Reserved	0x1	Reserved
2	OS2BMC Capable	0b	0b = Disabled. 1b = Enabled.
1	LAN1 TCO Isolate Disable	1b	0b = Enabled. 1b = Disabled.
0	LAN0 TCO Isolate Disable	1b	0b = Enabled. 1b = Disabled.



6.2.13.3 Common Firmware Parameters 2 — 0x0002

Bits	Name	Default	Description
15:12	Reserved	0x0	Reserved
11	Multi-drop NC-SI	1b	Multi-Drop NC-SI Topology. 0b = Point-to-point. 1b = Multi-drop (default). When this bit is set, the NCSI_CRS_DV and NCSI_RXD[1:0] pins are High-Z following power-up; otherwise, the pins are driven.
10	Proxying Capable	1b	0b = Disable protocol offload. 1b = Enable protocol offload.
9:6	Reserved	0x0	Reserved
5:1	Semaphore Backoff Interval	0x2	Number of 10 ms ticks that firmware must wait before taking semaphore ownership again since it has released it.
0	Reserved	0b	Reserved

6.2.13.4 Section Footer — 0x0003

Bits	Name	Default	Description
15:8	Block CRC8		CRC-8-CCITT: Start Section -> Word: Common Firmware Parameters Module -> Section Header End Section -> Word: Common Firmware Parameters Module -> Common Firmware Parameters 2
7:0	Reserved	0x0	Block length in words.

6.2.14 Pass Through Control Words Section Summary Table

Word Offset	Description	Page
0x0000	Block Length	204



Word Offset	Description	Page
0x0001	LAN IPv4 Address 0 (LSB) MIPAF0	205
0x0002	LAN IPv4 Address 0 (LSB) MIPAF0	205
0x0003	LAN IPv4 Address 1 (LSB) MIPAF1	205
0x0004	LAN IPv4 Address 1 (LSB) MIPAF1	206
0x0005	LAN IPv4 Address 2 (LSB) MIPAF2	206
0x0006	LAN IPv4 Address 2 (LSB) MIPAF2	206
0x0007	LAN IPv4 Address 3 (LSB) MIPAF3	207
0x0008	LAN IPv4 Address 3 (LSB) MIPAF3	207
0x0009	LAN Ethernet MAC Address 0 (LSB) MMAL0	207
0x000A	LAN Ethernet MAC Address 0 (Mid) MMAL0	207
0x000B	LAN Ethernet MAC Address 0 (MSB) MMAL0	208
0x000C	LAN Ethernet MAC Address 1 (LSB) MMAL1	208
0x000D	LAN Ethernet MAC Address 1 (Mid) MMAL1	208
0x000E	LAN Ethernet MAC Address 1 (MSB) MMAL1	209
0x000F	LAN Ethernet MAC Address 2 (LSB) MMAL2	209
0x0010	LAN Ethernet MAC Address 2 (Mid) MMAL2	209
0x0011	LAN Ethernet MAC Address 2 (MSB) MMAL2	210
0x0012	LAN Ethernet MAC Address 3 (LSB) MMAL3	210
0x0013	LAN Ethernet MAC Address 3 (Mid) MMAL3	210
0x0014	LAN Ethernet MAC Address 3 (MSB) MMAL3	211
0x0015	LAN TCP/UDP Flexible Filter Port0	211
0x0016	LAN TCP/UDP Flexible Filter Port1	211
0x0017	LAN TCP/UDP Flexible Filter Port2	212
0x0018	LAN TCP/UDP Flexible Filter Port3	212
0x0019	LAN TCP/UDP Flexible Filter Port4	212
0x001A	LAN TCP/UDP Flexible Filter Port5	212
0x001B	LAN TCP/UDP Flexible Filter Port6	212
0x001C	LAN TCP/UDP Flexible Filter Port7	213
0x001D	LAN TCP/UDP Flexible Filter Port8	213
0x001E	LAN TCP/UDP Flexible Filter Port9	213
0x001F	LAN TCP/UDP Flexible Filter Port10	213
0x0020	LAN TCP/UDP Flexible Filter Port11	213
0x0021	LAN TCP/UDP Flexible Filter Port12	214
0x0022	LAN TCP/UDP Flexible Filter Port13	214
0x0023	LAN TCP/UDP Flexible Filter Port14	214
0x0024	LAN TCP/UDP Flexible Filter Port15	214
0x0025	LAN VLAN Filter 0	215
0x0026	LAN VLAN Filter 1	215
0x0027	LAN VLAN Filter 2	215
0x0028	LAN VLAN Filter 3	216
0x0029	LAN VLAN Filter 4	216



Word Offset	Description	Page
0x002A	LAN VLAN Filter 5	216
0x002B	LAN VLAN Filter 6	217
0x002C	LAN VLAN Filter 7	217
0x002D	MANC Value LSB (LMANC LSB)	217
0x002E	MANC Value MSB (LMANC MSB)	218
0x002F	Receive Enable 1 (LRXEN1)	219
0x0030	Receive Enable 2 (LRXEN2)	219
0x0031	MNGONLY Value (MNGONLY LSB)	220
0x0032	MNGONLY Value (MNGONLY MLSB)	220
0x0033	Manageability Decision Filters - MDEF0 LSB	220
0x0034	Manageability Decision Filters - MDEF0 MSB	222
0x0035	Manageability Decision Filters - MDEF0_EXT0 LSB	222
0x0036	Manageability Decision Filters - MDEF0_EXT0 MSB	222
0x0037	Manageability Decision Filters - MDEF1 LSB	222
0x0038	Manageability Decision Filters - MDEF1 MSB	224
0x0039	Manageability Decision Filters - MDEF2 LSB	224
0x003A	Manageability Decision Filters - MDEF2 MSB	224
0x003B	Manageability Decision Filters - MDEF2_EXT2 LSB	224
0x003C	Manageability Decision Filters - MDEF2_EXT2 MSB	226
0x003D	Manageability Decision Filters - MDEF3 LSB	226
0x003E	Manageability Decision Filters - MDEF3 MSB	226
0x003F	Manageability Decision Filters - MDEF3 LSB	226
0x0040	Manageability Decision Filters - MDEF3 MSB	228
0x0041	Manageability Decision Filters - MDEF_EXT3 LSB	228
0x0042	Manageability Decision Filters - MDEF_EXT3 MSB	228
0x0043	Manageability Decision Filters - MDEF4 LSB	228
0x0044	Manageability Decision Filters - MDEF4 MSB	230
0x0045	Manageability Decision Filters - MDEF4_EXT4 LSB	230
0x0046	Manageability Decision Filters - MDEF4_EXT4 MSB	230
0x0047	Manageability Decision Filters - MDEF5 LSB	230
0x0048	Manageability Decision Filters - MDEF5 MSB	232
0x0049	Manageability Decision Filters - MDEF5_EXT5 LSB	232
0x004A	Manageability Decision Filters - MDEF5_EXT5 MSB	232
0x004B	Manageability Decision Filters - MDEF6 LSB	232
0x004C	Manageability Decision Filters - MDEF6 MSB	234
0x004D	Manageability Decision Filters - MDEF6_EXT6 LSB	234
0x004E	Manageability Decision Filters - MDEF6_EXT6 MSB	234
0x004F	Manageability Ethertype filter 0.1 - METF0 (1)	234
0x0050	Manageability Ethertype filter 0.2 - METF0 (2)	234
0x0051	Manageability Ethertype filter 1.1 - METF1 (1)	235
0x0052	Manageability Ethertype filter 1.2 - METF1 (2)	235



Word Offset	Description	Page
0x0053	Manageability Ethertype filter 2.1 - METF2 (1)	235
0x0054	Manageability Ethertype filter 2.2 - METF2 (2)	235
0x0055	Manageability Ethertype filter 3.1 - METF3 (1)	235
0x0056	Manageability Ethertype filter 3.2 - METF3 (2)	235
0x0057	ARP Response IPv4 Address 0 (LSB)	236
0x0058	ARP Response IPv4 Address 0 (MSB)	236
0x0059	IPv6 Address 0 - 0	236
0x005A	IPv6 Address 0 - 1	236
0x005B	IPv6 Address 0 - 2	236
0x005C	IPv6 Address 0 - 3	237
0x005D	IPv6 Address 0 - 4	237
0x005E	IPv6 Address 0 - 5	237
0x005F	IPv6 Address 0 - 6	237
0x0060	IPv6 Address 0 - 7	237
0x0061	IPv6 Address 1 - 0	237
0x0062	IPv6 Address 1 - 1	237
0x0063	IPv6 Address 1 - 2	238
0x0064	IPv6 Address 1 - 3	238
0x0065	IPv6 Address 1 - 4	238
0x0066	IPv6 Address 1 - 5	238
0x0067	IPv6 Address 1 - 6	238
0x0068	IPv6 Address 1 - 7	238
0x0069	IPv6 Address 2 - 0	238
0x006A	IPv6 Address 2 - 1	239
0x006B	IPv6 Address 2 - 2	239
0x006C	IPv6 Address 2 - 3	239
0x006D	IPv6 Address 2 - 4	239
0x006E	IPv6 Address 2 - 5	239
0x006F	IPv6 Address 2 - 6	239
0x0070	IPv6 Address 2 - 7	239
0x0071	Block CRC8	240

6.2.14.1 Section Header — 0x0000

Bits	Name	Default	Description
15:0	Block Length		Length in words of the section covered by CRC. Block length in words.



6.2.14.2 LAN IPv4 Address 0 (LSB) MIPAFO — 0x0001

Bits	Name	Default	Description
15:8	IPv4 Address 0, Byte 1	0x00	
7:0	IPv4 Address 0, Byte 0	0x0A	

6.2.14.3 LAN IPv4 Address 0 (MSB) MIPAFO — 0x0002

Bits	Name	Default	Description
15:8	IPv4 Address 0, Byte 3	0x01	
7:0	IPv4 Address 0, Byte 2	0x00	

6.2.14.4 LAN IPv4 Address 1 (LSB) MIPAF1 — 0x0003

Bits	Name	Default	Description
15:8	IPv4 Address 1, Byte 1	0x00	
7:0	IPv4 Address 1, Byte 0	0x0A	



6.2.14.5 LAN IPv4 Address 1 (MSB) MIPAF1 — 0x0004

Bits	Name	Default	Description
15:8	IPv4 Address 1, Byte 3	0x02	
7:0	IPv4 Address 1, Byte 2	0x00	

6.2.14.6 LAN IPv4 Address 2 (LSB) MIPAF2 — 0x0005

Bits	Name	Default	Description
15:8	IPv4 Address 2, Byte 1	0x00	
7:0	IPv4 Address 2, Byte 0	0x0A	

6.2.14.7 LAN IPv4 Address 2 (MSB) MIPAF2 — 0x0006

Bits	Name	Default	Description
15:8	IPv4 Address 2, Byte 3	0x03	
7:0	IPv4 Address 2, Byte 2	0x00	



6.2.14.8 LAN IPv4 Address 3 (LSB) MIPAF3 — 0x0007

Bits	Name	Default	Description
15:8	IPv4 Address 3, Byte 1	0x00	
7:0	IPv4 Address 3, Byte 0	0x0A	

6.2.14.9 LAN IPv4 Address 3 (MSB) MIPAF3 — 0x0008

Bits	Name	Default	Description
15:8	IPv4 Address 3, Byte 3	0x04	
7:0	IPv4 Address 3, Byte 2	0x00	

6.2.14.10 LAN Ethernet MAC Address 0 (LSB) MMALO — 0x0009

Bits	Name	Default	Description
15:8	Ethernet MAC Address 0, Byte 1	0x88	
7:0	Ethernet MAC Address 0, Byte 0	0x88	

6.2.14.11 LAN Ethernet MAC Address 0 (Mid) MMALO — 0x000A

This word is loaded by firmware to the 16 bits of the MMAL[0] register.



Bits	Name	Default	Description
15:8	Ethernet MAC Address 0, Byte 3	0x88	
7:0	Ethernet MAC Address 0, Byte 2	0x88	

6.2.14.12 LAN Ethernet MAC Address 0 (MSB) MMALO — 0x000B

This word is loaded by firmware to the MMAH[0] register.

Bits	Name	Default	Description
15:8	Ethernet MAC Address 0, Byte 5	0x01	
7:0	Ethernet MAC Address 0, Byte 4	0x88	

6.2.14.13 LAN Ethernet MAC Address 1 (LSB) MMAL1 — 0x000C

This word is loaded by firmware to the 16 bits of the MMAL[1] register.

Bits	Name	Default	Description
15:8	Ethernet MAC Address 1, Byte 1	0x88	
7:0	Ethernet MAC Address 1, Byte 0	0x88	

6.2.14.14 LAN Ethernet MAC Address 1 (Mid) MMAL1 — 0x000D

This word is loaded by firmware to the 16 bits of the MMAL[1] register.



Bits	Name	Default	Description
15:8	Ethernet MAC Address 1, Byte 3	0x88	
7:0	Ethernet MAC Address 1, Byte 2	0x88	

6.2.14.15 LAN Ethernet MAC Address 1 (MSB) MMAL1 — 0x000E

This word is loaded by firmware to the MMAH[1] register.

Bits	Name	Default	Description
15:8	Ethernet MAC Address 1, Byte 5	0x02	
7:0	Ethernet MAC Address 1, Byte 4	0x88	


6.2.14.16 LAN Ethernet MAC Address 2 (LSB) MMAL2 — 0x000F

This word is loaded by firmware to the 16 bits of the MMAL[2] register.

Bits	Name	Default	Description
15:8	Ethernet MAC Address 2, Byte 1	0x88	
7:0	Ethernet MAC Address 2, Byte 0	0x88	

6.2.14.17 LAN Ethernet MAC Address 2 (Mid) MMAL2 — 0x0010

This word is loaded by firmware to the 16 bits of the MMAL[2] register.



Bits	Name	Default	Description
15:8	Ethernet MAC Address 2, Byte 3	0x88	
7:0	Ethernet MAC Address 2, Byte 2	0x88	

6.2.14.18 LAN Ethernet MAC Address 2 (MSB) MMAL2 — 0x0011

This word is loaded by firmware to the MMAH[2] register.

Bits	Name	Default	Description
15:8	Ethernet MAC Address 2, Byte 5	0x03	
7:0	Ethernet MAC Address 2, Byte 4	0x88	

6.2.14.19 LAN Ethernet MAC Address 3 (LSB) MMAL3 — 0x0012

This word is loaded by firmware to the 16 bits of the MMAL[3] register.

Bits	Name	Default	Description
15:8	Ethernet MAC Address 3, Byte 1	0x88	
7:0	Ethernet MAC Address 3, Byte 0	0x88	

6.2.14.20 LAN Ethernet MAC Address 3 (Mid) MMAL3 — 0x0013

This word is loaded by firmware to the 16 bits of the MMAL[3] register.



Bits	Name	Default	Description
15:8	Ethernet MAC Address 3, Byte 3	0x88	
7:0	Ethernet MAC Address 3, Byte 2	0x88	

6.2.14.21 LAN Ethernet MAC Address 3 (MSB) MMAL3 — 0x0014

This word is loaded by firmware to the MMAH[3] register.

Bits	Name	Default	Description
15:8	Ethernet MAC Address 3, Byte 5	0x04	
7:0	Ethernet MAC Address 3, Byte 4	0x88	

6.2.14.22 LAN TCP/UDP Flexible Filter Port0 — 0x0015

Bits	Name	Default	Description
15:0	LAN TCP/UDP Flexible Filter Port0	0x0	

6.2.14.23 LAN TCP/UDP Flexible Filter Port1 — 0x0016

Bits	Name	Default	Description
15:0	LAN TCP/UDP Flexible Filter Port1	0x0	



6.2.14.24 LAN TCP/UDP Flexible Filter Port2 — 0x0017

Bits	Name	Default	Description
15:0	LAN TCP/UDP Flexible Filter Port2	0x0	

6.2.14.25 LAN TCP/UDP Flexible Filter Port3 — 0x0018

Bits	Name	Default	Description
15:0	LAN TCP/UDP Flexible Filter Port3	0x0	

6.2.14.26 LAN TCP/UDP Flexible Filter Port4 — 0x0019

Bits	Name	Default	Description
15:0	LAN TCP/UDP Flexible Filter Port4	0x0	

6.2.14.27 LAN TCP/UDP Flexible Filter Port5 — 0x001A

Bits	Name	Default	Description
15:0	LAN TCP/UDP Flexible Filter Port5	0x0	

6.2.14.28 LAN TCP/UDP Flexible Filter Port6 — 0x001B

Bits	Name	Default	Description
15:0	LAN TCP/UDP Flexible Filter Port6	0x0	



6.2.14.29 LAN TCP/UDP Flexible Filter Port7 — 0x001C

Bits	Name	Default	Description
15:0	LAN TCP/UDP Flexible Filter Port7	0x0	

6.2.14.30 LAN TCP/UDP Flexible Filter Port8 — 0x001D

Bits	Name	Default	Description
15:0	LAN TCP/UDP Flexible Filter Port8	0x0	

6.2.14.31 LAN TCP/UDP Flexible Filter Port9 — 0x001E


Bits	Name	Default	Description
15:0	LAN TCP/UDP Flexible Filter Port9	0x0	

6.2.14.32 LAN TCP/UDP Flexible Filter Port10 — 0x001F

Bits	Name	Default	Description
15:0	LAN TCP/UDP Flexible Filter Port10	0x0	

6.2.14.33 LAN TCP/UDP Flexible Filter Port11 — 0x0020

Bits	Name	Default	Description
15:0	LAN TCP/UDP Flexible Filter Port11	0x0	



6.2.14.34 LAN TCP/UDP Flexible Filter Port12 — 0x0021

Bits	Name	Default	Description
15:0	LAN TCP/UDP Flexible Filter Port12	0x0	

6.2.14.35 LAN TCP/UDP Flexible Filter Port13 — 0x0022

Bits	Name	Default	Description
15:0	LAN TCP/UDP Flexible Filter Port13	0x0	

6.2.14.36 LAN TCP/UDP Flexible Filter Port14 — 0x0023

Bits	Name	Default	Description
15:0	LAN TCP/UDP Flexible Filter Port14	0x0	

6.2.14.37 LAN TCP/UDP Flexible Filter Port15 — 0x0024

Bits	Name	Default	Description
15:0	LAN TCP/UDP Flexible Filter Port15	0x0	



6.2.14.38 LAN VLAN Filter 0 — 0x0025

Bits	Name	Default	Description
15:12	Reserved	0x0	
11:0	VLAN Filter 0 Value	0x0	

6.2.14.39 LAN VLAN Filter 1 — 0x0026

Bits	Name	Default	Description
15:12	Reserved	0x0	
11:0	VLAN Filter 1 Value	0x0	

6.2.14.40 LAN VLAN Filter 2 — 0x0027

Bits	Name	Default	Description
15:12	Reserved	0x0	
11:0	VLAN Filter 2 Value	0x0	



6.2.14.41 LAN VLAN Filter 3 — 0x0028

Bits	Name	Default	Description
15:12	Reserved	0x0	
11:0	VLAN Filter 3 Value	0x0	

6.2.14.42 LAN VLAN Filter 4 — 0x0029

Bits	Name	Default	Description
15:12	Reserved	0x0	
11:0	VLAN Filter 4 Value	0x0	

6.2.14.43 LAN VLAN Filter 5 — 0x002A

Bits	Name	Default	Description
15:12	Reserved	0x0	
11:0	VLAN Filter 5 Value	0x0	



6.2.14.44 LAN VLAN Filter 6 — 0x002B


Bits	Name	Default	Description
15:12	Reserved	0x0	
11:0	VLAN Filter 6 Value	0x0	

6.2.14.45 LAN VLAN Filter 7 — 0x002C

Bits	Name	Default	Description
15:12	Reserved	0x0	
11:0	VLAN Filter 7 Value	0x0	

6.2.14.46 MANC Value LSB (LMANC LSB) — 0x002D

Bits	Name	Default	Description
15:0	Reserved	0x0	



6.2.14.47 MANC Value MSB (LMANC MSB) — 0x002E

Bits	Name	Default	Description
15:11	Reserved	0x0	Reserved
10	Net Type	0b	This bit is loaded to the NET_TYPE bit in the MANC register.
9	Fixed Net Type	0b	This bit is loaded to the FIXED_NET_TYPE bit in the MANC register.
8	Enable IPv4 Address Filters	0b	When set, the last 128 bits of the MIPAF register are used to store four IPv4 addresses for IPv4 filtering. When cleared, these bits store a single IPv6 filter.
7	Enable Xsum Filtering to MNG	0b	When this bit is set, only packets that pass the L3 and L4 checksum are sent to the manageability block. Note that this feature does not support tunneled IPv4/IPv6 packet inspection.
6:0	Reserved	0x0	Reserved

6.2.14.48 Receive Enable 1 (LRXEN1) — 0x002F

Bits	Name	Default	Description
15:8	Receive Enable Byte 12	0x0	BMC SMBus slave address.
7	Enable BMC Dedicated	0b	
6	Reserved	1b	Reserved



Bits	Name	Default	Description
5:4	Notification Method	0x0	0x0 = SMBus alert. 0x1 = Asynchronous notify. 0x2 = Direct receive. 0x3 = Reserved.
3	Enable ARP Response	0b	
2	Enable Status Reporting	0b	
1	Enable Receive All	0b	
0	Enable Receive TCO	0b	

6.2.14.49 Receive Enable 2 (LRXEN2) — 0x0030

Bits	Name	Default	Description
15:8	Receive Enable Byte 14	0x0	Alert value.
7:0	Receive Enable Byte 13	0x0	Interface data.



6.2.14.50 MNGONLY Value (MNGONLY LSB) — 0x0031

Bits	Name	Default	Description
15:8	Reserved	0x0	Reserved
7:0	Exclusive to MNG	0x0	When set, indicates that packets forwarded by the manageability filters to manageability are not sent to the host. Bit 0 corresponds to decision rule 0, etc.

6.2.14.51 MNGONLY Value (MNGONLY MSB) — 0x0032

Bits	Name	Default	Description
15:0	Reserved	0x0	Reserved

6.2.14.52 Manageability Decision Filters (MDEF0 LSB) — 0x0033

Bits	Name	Default	Description
15:12	Flex port (OR)	0x0	Flex Port. Controls the inclusion of flex port filtering in the manageability filter decision (OR section). Bit 12 corresponds to flex port 0, etc.
11	Port 0x26F (OR)	0x0	Port 0x26F. Controls the inclusion of port 0x26F filtering in the manageability filter decision (OR section).
10	Port 0x298 (OR)	0x0	Port 0x298. Controls the inclusion of port 0x298 filtering in the manageability filter decision (OR section).
9	Neighbor Discovery (OR)	0x0	Neighbor Discovery. Controls the inclusion of neighbor discovery filtering in the manageability filter decision (OR section). The neighbor types accepted by this filter are types 0x86, 0x87, 0x88 and 0x89.
8	ARP Response (OR)	0x0	ARP Response. Controls the inclusion of ARP response filtering in the manageability filter decision (OR section).
7	ARP Request (OR)	0x0	ARP Request. Controls the inclusion of ARP request filtering in the manageability filter decision (OR section).
6	VLAN (OR)	0x0	VLAN. Controls the inclusion of VLAN addresses filtering in the manageability filter decision (OR section).
5	Broadcast (OR)	0x0	Broadcast. Controls the inclusion of broadcast address filtering in the manageability filter decision (OR section).
4	Unicast (OR)	0x0	Unicast. Controls the inclusion of unicast address filtering in the manageability filter decision (OR section).
3	IP Address (AND)	0x0	IP Address. Controls the inclusion of IP address filtering in the manageability filter decision (AND section).



Bits	Name	Default	Description
2	VLAN (AND)	0x0	VLAN. Controls the inclusion of VLAN address filtering in the manageability filter decision (AND section).
1	Broadcast (AND)	0x0	Broadcast. Controls the inclusion of broadcast address filtering in the manageability filter decision (AND section).
0	Unicast (AND)	0x0	Unicast. Controls the inclusion of unicast address filtering in the manageability filter decision (AND section).



6.2.14.53 Manageability Decision Filters – MDEF0 MSB – 0x0034

Bits	Name	Default	Description
15:12	Flex TCO	0x0	Flex TCO. Controls the inclusion of Flex TCO filtering in the manageability filter decision (OR section). Bit 28 corresponds to Flex TCO filter 0, etc.
11:0	Flex port	0x0	Flex Port. Controls the inclusion of flex port filtering in the manageability filter decision (OR section). Bit 12 corresponds to flex port 0, etc.

6.2.14.54 Manageability Decision Filters – MDEF_EXT0 LSB – 0x0035

Bits	Name	Default	Description
15:12	Reserved	0x00	Reserved for additional L2 EtherType OR filters.
11:8	L2 EtherType OR L2 EtherType	0x0	L2 EtherType. Controls the inclusion of L2 EtherType filtering in the manageability filter decision (OR section).
7:4	Reserved	0x0	Reserved for additional L2 EtherType AND filters.
3:0	L2 EtherType AND L2 EtherType	0x0	L2 EtherType. Controls the inclusion of L2 EtherType filtering in the manageability filter decision (AND section).

6.2.14.55 Manageability Decision Filters – MDEF_EXT0 MSB – 0x0036

Bits	Name	Default	Description
15:0	Reserved	0x0000	Reserved

6.2.14.56 Manageability Decision Filters – MDEF1 LSB – 0x0037

Bits	Name	Default	Description
15:12	Flex port (OR)	0x0	Flex Port. Controls the inclusion of flex port filtering in the manageability filter decision (OR section). Bit 12 corresponds to flex port 0, etc.
11	Port 0x26F (OR)	0x0	Port 0x26F. Controls the inclusion of port 0x26F filtering in the manageability filter decision (OR section)
10	Port 0x298 (OR)	0x0	Port 0x298. Controls the inclusion of port 0x298 filtering in the manageability filter decision (OR section).
9	Neighbor Discovery (OR)	0x0	Neighbor Discovery. Controls the inclusion of neighbor discovery filtering in the manageability filter decision (OR section). The neighbor types accepted by this filter are types 0x86, 0x87, 0x88 and 0x89.
8	ARP Response (OR)	0x0	ARP Response. Controls the inclusion of ARP response filtering in the manageability filter decision (OR section).
7	ARP Request (OR)	0x0	ARP Request. Controls the inclusion of ARP request filtering in the manageability filter decision (OR section).



Bits	Name	Default	Description
6	VLAN (OR)	0x0	VLAN. Controls the inclusion of VLAN addresses filtering in the manageability filter decision (OR section).
5	Broadcast (OR)	0x0	Broadcast. Controls the inclusion of broadcast address filtering in the manageability filter decision (OR section).
4	Unicast (OR)	0x0	Unicast. Controls the inclusion of unicast address filtering in the manageability filter decision (OR section).
3	IP Address (AND)	0x0	IP Address. Controls the inclusion of IP address filtering in the manageability filter decision (AND section).
2	VLAN (AND)	0x0	VLAN. Controls the inclusion of VLAN address filtering in the manageability filter decision (AND section).
1	Broadcast (AND)	0x0	Broadcast. Controls the inclusion of broadcast address filtering in the manageability filter decision (AND section).
0	Unicast (AND)	0x0	Unicast. Controls the inclusion of unicast address filtering in the manageability filter decision (AND section).



6.2.14.57 Manageability Decision Filters – MDEF1 MSB – 0x0038

Bits	Name	Default	Description
15:12	Flex TCO	0x0	Flex TCO. Controls the inclusion of Flex TCO filtering in the manageability filter decision (OR section). Bit 28 corresponds to Flex TCO filter 0, etc.
11:0	Flex port	0x0	Flex Port. Controls the inclusion of flex port filtering in the manageability filter decision (OR section). Bit 12 corresponds to flex port 0, etc.

6.2.14.58 Manageability Decision Filters – MDEF_EXT1 LSB – 0x0039


Bits	Name	Default	Description
15:12	Reserved	0x00	Reserved for additional L2 EtherType OR filters.
11:8	L2 EtherType OR L2 EtherType	0x0	L2 EtherType. Controls the inclusion of L2 EtherType filtering in the manageability filter decision (OR section).
7:4	Reserved	0x0	Reserved for additional L2 EtherType AND filters.
3:0	L2 EtherType AND L2 EtherType	0x0	L2 EtherType. Controls the inclusion of L2 EtherType filtering in the manageability filter decision (AND section).

6.2.14.59 Manageability Decision Filters – MDEF_EXT1 MSB – 0x003A

Bits	Name	Default	Description
15:0	Reserved	0x0000	Reserved

6.2.14.60 Manageability Decision Filters – MDEF2 LSB – 0x003B

Bits	Name	Default	Description
15:12	Flex port (OR)	0x0	Flex Port. Controls the inclusion of flex port filtering in the manageability filter decision (OR section). Bit 12 corresponds to flex port 0, etc.
11	Port 0x26F (OR)	0x0	Port 0x26F. Controls the inclusion of port 0x26F filtering in the manageability filter decision (OR section)
10	Port 0x298 (OR)	0x0	Port 0x298. Controls the inclusion of port 0x298 filtering in the manageability filter decision (OR section).
9	Neighbor Discovery (OR)	0x0	Neighbor Discovery. Controls the inclusion of neighbor discovery filtering in the manageability filter decision (OR section). The neighbor types accepted by this filter are types 0x86, 0x87, 0x88 and 0x89.
8	ARP Response (OR)	0x0	ARP Response. Controls the inclusion of ARP response filtering in the manageability filter decision (OR section).
7	ARP Request (OR)	0x0	ARP Request. Controls the inclusion of ARP request filtering in the manageability filter decision (OR section).



Bits	Name	Default	Description
6	VLAN (OR)	0x0	VLAN. Controls the inclusion of VLAN addresses filtering in the manageability filter decision (OR section).
5	Broadcast (OR)	0x0	Broadcast. Controls the inclusion of broadcast address filtering in the manageability filter decision (OR section).
4	Unicast (OR)	0x0	Unicast. Controls the inclusion of unicast address filtering in the manageability filter decision (OR section).
3	IP Address (AND)	0x0	IP Address. Controls the inclusion of IP address filtering in the manageability filter decision (AND section).
2	VLAN (AND)	0x0	VLAN. Controls the inclusion of VLAN address filtering in the manageability filter decision (AND section).
1	Broadcast (AND)	0x0	Broadcast. Controls the inclusion of broadcast address filtering in the manageability filter decision (AND section).
0	Unicast (AND)	0x0	Unicast. Controls the inclusion of unicast address filtering in the manageability filter decision (AND section).



6.2.14.61 Manageability Decision Filters – MDEF2 MSB – 0x003C

Bits	Name	Default	Description
15:12	Flex TCO	0x0	Flex TCO. Controls the inclusion of Flex TCO filtering in the manageability filter decision (OR section). Bit 28 corresponds to Flex TCO filter 0, etc.
11:0	Flex port	0x0	Flex Port. Controls the inclusion of flex port filtering in the manageability filter decision (OR section). Bit 12 corresponds to flex port 0, etc.

6.2.14.62 Manageability Decision Filters – MDEF_EXT2 LSB – 0x003D


Bits	Name	Default	Description
15:12	Reserved	0x00	Reserved for additional L2 EtherType OR filters.
11:8	L2 EtherType OR L2 EtherType	0x0	L2 EtherType. Controls the inclusion of L2 EtherType filtering in the manageability filter decision (OR section).
7:4	Reserved	0x0	Reserved for additional L2 EtherType AND filters.
3:0	L2 EtherType AND L2 EtherType	0x0	L2 EtherType. Controls the inclusion of L2 EtherType filtering in the manageability filter decision (AND section).

6.2.14.63 Manageability Decision Filters – MDEF2_EXT2 MSB – 0x003E

Bits	Name	Default	Description
15:0	Reserved	0x0000	Reserved

6.2.14.64 Manageability Decision Filters – MDEF3 LSB – 0x003F

Bits	Name	Default	Description
15:12	Flex port (OR)	0x0	Flex Port. Controls the inclusion of flex port filtering in the manageability filter decision (OR section). Bit 12 corresponds to flex port 0, etc.
11	Port 0x26F (OR)	0x0	Port 0x26F. Controls the inclusion of port 0x26F filtering in the manageability filter decision (OR section)
10	Port 0x298 (OR)	0x0	Port 0x298. Controls the inclusion of port 0x298 filtering in the manageability filter decision (OR section).
9	Neighbor Discovery (OR)	0x0	Neighbor Discovery. Controls the inclusion of neighbor discovery filtering in the manageability filter decision (OR section). The neighbor types accepted by this filter are types 0x86, 0x87, 0x88 and 0x89.
8	ARP Response (OR)	0x0	ARP Response. Controls the inclusion of ARP response filtering in the manageability filter decision (OR section).
7	ARP Request (OR)	0x0	ARP Request. Controls the inclusion of ARP request filtering in the manageability filter decision (OR section).



Bits	Name	Default	Description
6	VLAN (OR)	0x0	VLAN. Controls the inclusion of VLAN addresses filtering in the manageability filter decision (OR section).
5	Broadcast (OR)	0x0	Broadcast. Controls the inclusion of broadcast address filtering in the manageability filter decision (OR section).
4	Unicast (OR)	0x0	Unicast. Controls the inclusion of unicast address filtering in the manageability filter decision (OR section).
3	IP Address (AND)	0x0	IP Address. Controls the inclusion of IP address filtering in the manageability filter decision (AND section).
2	VLAN (AND)	0x0	VLAN. Controls the inclusion of VLAN address filtering in the manageability filter decision (AND section).
1	Broadcast (AND)	0x0	Broadcast. Controls the inclusion of broadcast address filtering in the manageability filter decision (AND section).
0	Unicast (AND)	0x0	Unicast. Controls the inclusion of unicast address filtering in the manageability filter decision (AND section).



6.2.14.65 Manageability Decision Filters – MDEF3 MSB – 0x0040

Bits	Name	Default	Description
15:12	Flex TCO	0x0	Flex TCO. Controls the inclusion of Flex TCO filtering in the manageability filter decision (OR section). Bit 28 corresponds to Flex TCO filter 0, etc.
11:0	Flex port	0x0	Flex Port. Controls the inclusion of flex port filtering in the manageability filter decision (OR section). Bit 12 corresponds to flex port 0, etc.

6.2.14.66 Manageability Decision Filters – MDEF_EXT3 LSB – 0x0041

Bits	Name	Default	Description
15:12	Reserved	0x00	Reserved for additional L2 EtherType OR filters.
11:8	L2 EtherType OR L2 EtherType	0x0	L2 EtherType. Controls the inclusion of L2 EtherType filtering in the manageability filter decision (OR section).
7:4	Reserved	0x0	Reserved for additional L2 EtherType AND filters.
3:0	L2 EtherType AND L2 EtherType	0x0	L2 EtherType. Controls the inclusion of L2 EtherType filtering in the manageability filter decision (AND section).

6.2.14.67 Manageability Decision Filters – MDEF_EXT3 MSB – 0x0042

Bits	Name	Default	Description
15:0	Reserved	0x0000	Reserved

6.2.14.68 Manageability Decision Filters – MDEF4 LSB – 0x0043

Bits	Name	Default	Description
15:12	Flex port (OR)	0x0	Flex Port. Controls the inclusion of flex port filtering in the manageability filter decision (OR section). Bit 12 corresponds to flex port 0, etc.
11	Port 0x26F (OR)	0x0	Port 0x26F. Controls the inclusion of port 0x26F filtering in the manageability filter decision (OR section)
10	Port 0x298 (OR)	0x0	Port 0x298. Controls the inclusion of port 0x298 filtering in the manageability filter decision (OR section).
9	Neighbor Discovery (OR)	0x0	Neighbor Discovery. Controls the inclusion of neighbor discovery filtering in the manageability filter decision (OR section). The neighbor types accepted by this filter are types 0x86, 0x87, 0x88 and 0x89.
8	ARP Response (OR)	0x0	ARP Response. Controls the inclusion of ARP response filtering in the manageability filter decision (OR section).
7	ARP Request (OR)	0x0	ARP Request. Controls the inclusion of ARP request filtering in the manageability filter decision (OR section).



Bits	Name	Default	Description
6	VLAN (OR)	0x0	VLAN. Controls the inclusion of VLAN addresses filtering in the manageability filter decision (OR section).
5	Broadcast (OR)	0x0	Broadcast. Controls the inclusion of broadcast address filtering in the manageability filter decision (OR section).
4	Unicast (OR)	0x0	Unicast. Controls the inclusion of unicast address filtering in the manageability filter decision (OR section).
3	IP Address (AND)	0x0	IP Address. Controls the inclusion of IP address filtering in the manageability filter decision (AND section).
2	VLAN (AND)	0x0	VLAN. Controls the inclusion of VLAN address filtering in the manageability filter decision (AND section).
1	Broadcast (AND)	0x0	Broadcast. Controls the inclusion of broadcast address filtering in the manageability filter decision (AND section).
0	Unicast (AND)	0x0	Unicast. Controls the inclusion of unicast address filtering in the manageability filter decision (AND section).



6.2.14.69 Manageability Decision Filters – MDEF4 MSB – 0x0044

Bits	Name	Default	Description
15:12	Flex TCO	0x0	Flex TCO. Controls the inclusion of Flex TCO filtering in the manageability filter decision (OR section). Bit 28 corresponds to Flex TCO filter 0, etc.
11:0	Flex port	0x0	Flex Port. Controls the inclusion of flex port filtering in the manageability filter decision (OR section). Bit 12 corresponds to flex port 0, etc.

6.2.14.70 Manageability Decision Filters – MDEF_EXT4 LSB – 0x0045

Bits	Name	Default	Description
15:12	Reserved	0x00	Reserved for additional L2 EtherType OR filters.
11:8	L2 EtherType OR L2 EtherType	0x0	L2 EtherType. Controls the inclusion of L2 EtherType filtering in the manageability filter decision (OR section).
7:4	Reserved	0x0	Reserved for additional L2 EtherType AND filters.
3:0	L2 EtherType AND L2 EtherType	0x0	L2 EtherType. Controls the inclusion of L2 EtherType filtering in the manageability filter decision (AND section).

6.2.14.71 Manageability Decision Filters – MDEF_EXT4 MSB – 0x0046

Bits	Name	Default	Description
15:0	Reserved	0x0000	Reserved

6.2.14.72 Manageability Decision Filters – MDEF5 LSB – 0x0047

Bits	Name	Default	Description
15:12	Flex port (OR)	0x0	Flex Port. Controls the inclusion of flex port filtering in the manageability filter decision (OR section). Bit 12 corresponds to flex port 0, etc.
11	Port 0x26F (OR)	0x0	Port 0x26F. Controls the inclusion of port 0x26F filtering in the manageability filter decision (OR section)
10	Port 0x298 (OR)	0x0	Port 0x298. Controls the inclusion of port 0x298 filtering in the manageability filter decision (OR section).
9	Neighbor Discovery (OR)	0x0	Neighbor Discovery. Controls the inclusion of neighbor discovery filtering in the manageability filter decision (OR section). The neighbor types accepted by this filter are types 0x86, 0x87, 0x88 and 0x89.
8	ARP Response (OR)	0x0	ARP Response. Controls the inclusion of ARP response filtering in the manageability filter decision (OR section).
7	ARP Request (OR)	0x0	ARP Request. Controls the inclusion of ARP request filtering in the manageability filter decision (OR section).



Bits	Name	Default	Description
6	VLAN (OR)	0x0	VLAN. Controls the inclusion of VLAN addresses filtering in the manageability filter decision (OR section).
5	Broadcast (OR)	0x0	Broadcast. Controls the inclusion of broadcast address filtering in the manageability filter decision (OR section).
4	Unicast (OR)	0x0	Unicast. Controls the inclusion of unicast address filtering in the manageability filter decision (OR section).
3	IP Address (AND)	0x0	IP Address. Controls the inclusion of IP address filtering in the manageability filter decision (AND section).
2	VLAN (AND)	0x0	VLAN. Controls the inclusion of VLAN address filtering in the manageability filter decision (AND section).
1	Broadcast (AND)	0x0	Broadcast. Controls the inclusion of broadcast address filtering in the manageability filter decision (AND section).
0	Unicast (AND)	0x0	Unicast. Controls the inclusion of unicast address filtering in the manageability filter decision (AND section).



6.2.14.73 Manageability Decision Filters – MDEF5 MSB – 0x0048

Bits	Name	Default	Description
15:12	Flex TCO	0x0	Flex TCO. Controls the inclusion of Flex TCO filtering in the manageability filter decision (OR section). Bit 28 corresponds to Flex TCO filter 0, etc.
11:0	Flex port	0x0	Flex Port. Controls the inclusion of flex port filtering in the manageability filter decision (OR section). Bit 12 corresponds to flex port 0, etc.

6.2.14.74 Manageability Decision Filters – MDEF_EXT5 LSB – 0x0049

Bits	Name	Default	Description
15:12	Reserved	0x00	Reserved for additional L2 EtherType OR filters.
11:8	L2 EtherType OR L2 EtherType	0x0	L2 EtherType. Controls the inclusion of L2 EtherType filtering in the manageability filter decision (OR section).
7:4	Reserved	0x0	Reserved for additional L2 EtherType AND filters.
3:0	L2 EtherType AND L2 EtherType	0x0	L2 EtherType. Controls the inclusion of L2 EtherType filtering in the manageability filter decision (AND section).

6.2.14.75 Manageability Decision Filters – MDEF5_EXT5 MSB – 0x004A

Bits	Name	Default	Description
15:0	Reserved	0x0000	Reserved

6.2.14.76 Manageability Decision Filters – MDEF6 LSB – 0x004B

Bits	Name	Default	Description
15:12	Flex port (OR)	0x0	Flex Port. Controls the inclusion of flex port filtering in the manageability filter decision (OR section). Bit 12 corresponds to flex port 0, etc.
11	Port 0x26F (OR)	0x0	Port 0x26F. Controls the inclusion of port 0x26F filtering in the manageability filter decision (OR section)
10	Port 0x298 (OR)	0x0	Port 0x298. Controls the inclusion of port 0x298 filtering in the manageability filter decision (OR section).
9	Neighbor Discovery (OR)	0x0	Neighbor Discovery. Controls the inclusion of neighbor discovery filtering in the manageability filter decision (OR section). The neighbor types accepted by this filter are types 0x86, 0x87, 0x88 and 0x89.
8	ARP Response (OR)	0x0	ARP Response. Controls the inclusion of ARP response filtering in the manageability filter decision (OR section).
7	ARP Request (OR)	0x0	ARP Request. Controls the inclusion of ARP request filtering in the manageability filter decision (OR section).



Bits	Name	Default	Description
6	VLAN (OR)	0x0	VLAN. Controls the inclusion of VLAN addresses filtering in the manageability filter decision (OR section).
5	Broadcast (OR)	0x0	Broadcast. Controls the inclusion of broadcast address filtering in the manageability filter decision (OR section).
4	Unicast (OR)	0x0	Unicast. Controls the inclusion of unicast address filtering in the manageability filter decision (OR section).
3	IP Address (AND)	0x0	IP Address. Controls the inclusion of IP address filtering in the manageability filter decision (AND section).
2	VLAN (AND)	0x0	VLAN. Controls the inclusion of VLAN address filtering in the manageability filter decision (AND section).
1	Broadcast (AND)	0x0	Broadcast. Controls the inclusion of broadcast address filtering in the manageability filter decision (AND section).
0	Unicast (AND)	0x0	Unicast. Controls the inclusion of unicast address filtering in the manageability filter decision (AND section).



6.2.14.77 Manageability Decision Filters – MDEF6 MSB – 0x004C

Bits	Name	Default	Description
15:12	Flex TCO	0x0	Flex TCO. Controls the inclusion of Flex TCO filtering in the manageability filter decision (OR section). Bit 28 corresponds to Flex TCO filter 0, etc.
11:0	Flex port	0x0	Flex Port. Controls the inclusion of flex port filtering in the manageability filter decision (OR section). Bit 12 corresponds to flex port 0, etc.

6.2.14.78 Manageability Decision Filters – MDEF_EXT6 LSB – 0x004D

Bits	Name	Default	Description
15:12	Reserved	0x00	Reserved for additional L2 EtherType OR filters.
11:8	L2 EtherType OR L2 EtherType	0x0	L2 EtherType. Controls the inclusion of L2 EtherType filtering in the manageability filter decision (OR section).
7:4	Reserved	0x0	Reserved for additional L2 EtherType AND filters.
3:0	L2 EtherType AND L2 EtherType	0x0	L2 EtherType. Controls the inclusion of L2 EtherType filtering in the manageability filter decision (AND section).

6.2.14.79 Manageability Decision Filters – MDEF6_EXT6 MSB – 0x004E

Bits	Name	Default	Description
15:0	Reserved	0x0000	Reserved

6.2.14.80 Manageability Ethertype filter 0.1 – METF0 (1) – 0x004F

Bits	Name	Default	Description
15:0	METF0_L	0x0000	Loaded to 16 LS bits of METF[0] register

6.2.14.81 Manageability Ethertype filter 0.2 – METF0 (2) – 0x0050

Bits	Name	Default	Description
15:0	METF0_M	0x0000	Loaded to 16 MS bits of METF[0] register (reserved bits in the METF registers should be set in the NVM to the register's default values).



6.2.14.82 Manageability Ethertype filter 1.1 – METF1 (1) – 0x0051

Bits	Name	Default	Description
15:0	METF1_L	0x0000	Loaded to 16 LS bits of METF[1] register.

6.2.14.83 Manageability Ethertype filter 1.2 – METF1 (2) – 0x0052

Bits	Name	Default	Description
15:0	METF1_M	0x0000	Loaded to 16 MS bits of METF[1] register (reserved bits in the METF registers should be set in the NVM to the register's default values).

6.2.14.84 Manageability Ethertype filter 2.1 – METF2 (1) – 0x0053

Bits	Name	Default	Description
15:0	METF2_L	0x0000	Loaded to 16 LS bits of METF[2] register.

6.2.14.85 Manageability Ethertype filter 2.2 – METF2 (2) – 0x0054

Bits	Name	Default	Description
15:0	METF2_M	0x0000	Loaded to 16 MS bits of METF[2] register (reserved bits in the METF registers should be set in the NVM to the register's default values).

6.2.14.86 Manageability Ethertype filter 3.1 – METF3 (1) – 0x0055

Bits	Name	Default	Description
15:0	METF3_L	0x0000	Loaded to 16 LS bits of METF[3] register.

6.2.14.87 Manageability Ethertype Filter 3.2 – METF3 (2)



– 0x0056

Bits	Name	Default	Description
15:0	METF3_M	0x0000	Loaded to 16 MS bits of METF[3] register (reserved bits in the METF registers should be set in the NVM to the register's default values).

6.2.14.88 ARP Response IPv4 Address 0 (LSB) – 0x0057

Bits	Name	Default	Description
15:8	Byte 1	0x00	Firmware use.
7:0	Byte 0	0x00	Firmware use.

6.2.14.89 ARP Response IPv4 Address 0 (MSB) – 0x0058

Bits	Name	Default	Description
15:8	Byte 3	0x00	Firmware use.
7:0	Byte 2	0x00	Firmware use.

6.2.14.90 IPv6 Address 0 - 0 – 0x0059

Bits	Name	Default	Description
15:0	IPv6 Address	0x0000	

6.2.14.91 IPv6 Address 0 – 1 – 0x005A

Bits	Name	Default	Description
15:0	IPv6 Address	0x0000	

6.2.14.92 IPv6 Address 0 – 2 – 0x005B

Bits	Name	Default	Description
15:0	IPv6 Address	0x0000	



6.2.14.93 IPv6 Address 0 – 3 – 0x005C

Bits	Name	Default	Description
15:0	IPv6 Address	0x0000	

6.2.14.94 IPv6 Address 0 – 4 – 0x005D

Bits	Name	Default	Description
15:0	IPv6 Address	0x0000	

6.2.14.95 IPv6 Address 0 – 5 – 0x005E

Bits	Name	Default	Description
15:0	IPv6 Address	0x0000	

6.2.14.96 IPv6 Address 0 – 6 – 0x005F

Bits	Name	Default	Description
15:0	IPv6 Address	0x0000	

6.2.14.97 IPv6 Address 0 – 7 – 0x0060

Bits	Name	Default	Description
15:0	IPv6 Address	0x0000	

6.2.14.98 IPv6 Address 1 – 0 – 0x0061

Bits	Name	Default	Description
15:0	IPv6 Address	0x0000	

6.2.14.99 IPv6 Address 1 – 1 – 0x0062

Bits	Name	Default	Description
15:0	IPv6 Address	0x0000	



6.2.14.100 IPv6 Address 1 – 2 – 0x0063

Bits	Name	Default	Description
15:0	IPv6 Address	0x0000	

6.2.14.101 IPv6 Address 1 – 3 – 0x0064

Bits	Name	Default	Description
15:0	IPv6 Address	0x0000	

6.2.14.102 IPv6 Address 1 – 4 – 0x0065

Bits	Name	Default	Description
15:0	IPv6 Address	0x0000	

6.2.14.103 IPv6 Address 1 – 5 – 0x0066

Bits	Name	Default	Description
15:0	IPv6 Address	0x0000	

6.2.14.104 IPv6 Address 1 – 6 -- 0x0067

Bits	Name	Default	Description
15:0	IPv6 Address	0x0000	

6.2.14.105 IPv6 Address 1 – 7 – 0x0068

Bits	Name	Default	Description
15:0	IPv6 Address	0x0000	

6.2.14.106 IPv6 Address 2 – 0 – 0x0069

Bits	Name	Default	Description
15:0	IPv6 Address	0x0000	



6.2.14.107 IPv6 Address 2 – 1 – 0x006A

Bits	Name	Default	Description
15:0	IPv6 Address	0x0000	

6.2.14.108 IPv6 Address 2 – 2 – 0x006B

Bits	Name	Default	Description
15:0	IPv6 Address	0x0000	

6.2.14.109 IPv6 Address 2 – 3 – 0x006C

Bits	Name	Default	Description
15:0	IPv6 Address	0x0000	

6.2.14.110 IPv6 Address 2 – 4 – 0x006D

Bits	Name	Default	Description
15:0	IPv6 Address	0x0000	

6.2.14.111 IPv6 Address 2 – 5 – 0x006E


Bits	Name	Default	Description
15:0	IPv6 Address	0x0000	

6.2.14.112 IPv6 Address 2 – 6 – 0x006F

Bits	Name	Default	Description
15:0	IPv6 Address	0x0000	

6.2.14.113 IPv6 Address 2 – 7 – 0x0070

Bits	Name	Default	Description
15:0	IPv6 Address	0x0000	



6.2.14.114 Block CRC8 – 0x0071

Bits	Name	Default	Description
15:8	Block CRC8		CRC-8-CCITT: Start Section -> Word: Pass Through Control Words -> Block Length End Section -> Word: Pass Through Control Words -> IPv6 Address 2 - 7
7:0	Reserved	0x0	Block length in words

6.2.15 Sideband Configuration Structure Section Summary Table

Word Offset	Description	Page
0x0000	Block Length	240
0x0001	SMBus Maximum Fragment Size	241
0x0002	SMBus Notification Timeout and Flags	241
0x0003	SMBus Slave Addresses	241
0x0004	NC-SI Configuration 1	242
0x0005	NC-SI Configuration 2	242
0x0006	NCSI Flow Control XOFF	243
0x0007	NCSI Flow Control XON	243
0x0008	NC-SI HW Arbitration TOKEN Timeout	243
0x0009-0x000D	Reserved	243
0x000E	OEM IANA	243
0x000F	NC-SI over MCTP Message Types	243
0x0010	NC-SI over MCTP Configuration	244
0x0011	Traffic Types Parameters	244
0x0012	MCTP Rate Limiter Config 1	244
0x0013	MCTP Rate Limiter Config 2	245
0x0014	NC-SI Channel to Port Mapping	245
0x0015	Block CRC8	245

6.2.15.1 Block Length – 0x0000

Bits	Name	Default	Description
15:0	Block Length	0x13	Length in words of the section covered by CRC. Section length in words.



6.2.15.2 SMBus Maximum Fragment Size – 0x0001


Bits	Name	Default	Description
15:0	Fragment size	0x20	SMBus Maximum Fragment Size (bytes). Supported range is between 32 and 240 bytes. Note: In MCTP mode, this value should be set to 0x45 (64 bytes payload + 5 bytes of MCTP header).

6.2.15.3 SMBus Notification Timeout and Flags – 0x0002

Bits	Name	Default	Description
15:8	SMBus Notification Timeout (ms)	0xFF	
7:6	SMBus Connection Speed	0x0	0x0 Standard SMBus Connection. 0x1 = Reserved 1. 0x2 = Reserved 2. 0x3 = Reserved 3.
5	SMBus Block Read command	0b	0x0 = Block read command is 0xC0. 0x1 Block read command is 0xD0.
4	Reserved, must be 0x1	1b	0x0 Single address mode 0x1 Dual address mode
3	Reserved	0b	Reserved
2	Disable SMBus ARP Functionality	0b	
1	SMBus ARP PEC	1b	SMBus Transactions PEC. Should be set. 0x0 = Disable SMBus ARP PEC. 0x1 = Enable SMBus ARP PEC.
0	SMBus Transaction PEC	0b	SMBus Transactions PEC. Should be set in MCTP modes. 0x0 = Disable PEC If this bit is cleared, PEC is not added to master write or slave read transactions, a slave write transaction with PEC is dropped. 0x1 = Enable PEC If this bit is set, PEC is added for master SMBus write transactions. a PEC is added to slave read transactions and can be received in slave write transaction.

6.2.15.4 SMBus Slave Addresses – 0x0003

Bits	Name	Default	Description
15:9	SMBus 1 Slave Address	0x62	Dual address mode only.
8	Reserved	0x0	
7:1	SMBus 0 Slave Address	0x63	
0	Reserved	0x0	



6.2.15.5 NC-SI Configuration 1 – 0x0004

Bits	Name	Default	Description
15:11	Reserved	0x0	
10	Flow Control	0b	0x0 = Disable. 0x1 = Enable.
9	RMII HW Arbitration Enable	0b	0x0 = Not supported. 0x1 = Supported.
8	RMII HW-based Packet Copy Enable	1b	0x0 = Disable. 0x1 = Enable.
7:5	Package ID	0x0	Meaningful only when bit 15 of NC-SI Configuration 2 word (offset 0x07) is cleared.
4:0	LAN 0 Internal Channel ID	0x0	

6.2.15.6 NC-SI Configuration 2 – 0x0005

Bits	Name	Default	Description
15	Read NC-SI Package ID from SDP	0x0	Read NC-SI Package ID From. 0b = NVM, NC-SI Configuration 1 word bits 7:5 (offset 0x06). 1b = SDP, SDPn_0 pins of LAN ports n=1,0 (where SDP1_0, SDP0 of LAN port 1 is the MS bit of NC-SI Package ID).
14:8	Reserved	0x0	
7:4	Reserved	0x0	Reserved
3:0	Max XOFF Renewal	0x3	NC-SI Flow Control MAX XOFF Renewal (# of XOFF renewals allowed). 0x0 = Disabled (unlimited allowed). 0x1 = Up to 2 consecutive. 0x2 = Up to 0x3 consecutive. 0x3 = Up to 0x4 consecutive. 0x4 = Up to 0x5 consecutive. 0x5 = Up to 0x6 consecutive. 0x6 = Up to 0x7 consecutive. 0x7 = Up to 0x8 consecutive. 0x8 = Up to 0x9 consecutive. 0x9 = Up to 0xA consecutive. 0xA = Up to 0xB consecutive. 0xB = Up to 0xC consecutive. 0xC = Up to 0xD consecutive. 0xD = Up to 0xE consecutive. 0xE = Up to 0xF consecutive. 0xF = Up to 0x10 consecutive.



6.2.15.7 NCSI Flow Control XOFF – 0x0006

Bits	Name	Default	Description
15:0	XOFF Threshold	0xAB8	Tx buffer watermark for sending a XOFF NC-SI flow control packet in bytes. The XOFF threshold value refers to the occupied space in the buffer. Notes: 1. Field relevant for NC-SI operation mode only. 2. To support a maximum packet size of 1.5 KB, the value programmed assuming a Tx buffer size of 6 KB value of field should be 0xAB8 (2,744 bytes).

6.2.15.8 NCSI Flow Control XON – 0x0007

Bits	Name	Default	Description
15:0	XON Threshold	0x1348	Tx buffer water mark for sending a XON NC-SI flow control packet in bytes. The XON threshold value refers to the available space in the TX buffer. Notes: 1. Field relevant for NC-SI operation mode only. 2. To support maximum packet size of 1.5 KB, the value programmed should be a positive value that equals: XOFF threshold + 1536 bytes. Assuming a Tx buffer size is 6 KB and the XOFF threshold is 2,744 bytes value of field should be 0x1348 (4,936 bytes).

6.2.15.9 NC-SI HW Arbitration TOKEN Timeout – 0x0008

Bits	Name	Default	Description
15:0	NC-SI HW Arbitration TOKEN Timeout	0xA000	

6.2.15.10 Reserved – 0x0009-0x000D

6.2.15.11 OEM IANA – 0x000E

Bits	Name	Default	Description
15:0	OEM IANA	0x0	If not zero and not 0x157, device accepts NC-SI OEM commands with this IANA number.

6.2.15.12 NC-SI over MCTP Message Types – 0x000F

Bits	Name	Default	Description
15:8	NC-SI Command Packet Type	0x2	Defines the MCTP packet type used to identify NC-SI control packets.
7:0	NC-SI Pass Through packet type	0x3	Defines the MCTP packet type used to identify NC-SI PT packets



6.2.15.13 NC-SI Over MCTP Configuration – 0x0010

Bits	Name	Default	Description
15:8	Reserved	0x0	Reserved
7	Use Payload Type	1b	If set, a payload type byte is expected in NC-SI over MCTP packets after the packet type. Otherwise, the control and pass through are controlled via the NC-SI over MCTP message types defined in the previous word.
6	Simplified MCTP	0b	If set, only SOM and EOM bits are used for the reassembly process. Relevant only in SMBus mode.
5	Disable ACLs	0b	If set, the ACLs on the PCIe VDMs are disabled.
4:1	Reserved	0x0	Reserved
0	Endpoint Discovery Route	0b	Use Route by ID for Endpoint Discovery Responses. 0b = Use route to root complex type in outgoing messages and broadcast from root complex in incoming messages. 1b = Use route by ID for MCTP discovery messages. Note that this mode should be selected only in debug mode.

6.2.15.14 Traffic Types Parameters – 0x0011

Bits	Name	Default	Description
15:6	Reserved	0x0	Reserved
5:4	Port1 Traffic Types	0x1	00b = Reserved. 01b = Network to BMC traffic only allowed through port 1. 10b = OS2BMC traffic only allowed through port 1. 11b = Both network to BMC traffic and OS2BMC traffic allowed through port 1. This field is valid only if the <i>Port1 Manageability Capable</i> field in the Common Firmware Parameters NVM word is set.
3:2	Reserved	0x0	Reserved
1:0	Port 0 Traffic Types	0x1	00b = Reserved. 01b = Network to BMC traffic only allowed through port 0. 10b = OS2BMC traffic only allowed through port 0. 11b = Both network to BMC traffic and OS2BMC traffic allowed through port 0. This field is valid only if the <i>Port0 Manageability Capable</i> field in the Common Firmware Parameters NVM word is set.

6.2.15.15 MCTP Rate Limiter Config 1 – 0x0012

Bits	Name	Default	Description
15:0	MCTP Rate	0x2800	Defines the number of cycles between accesses of the MCTP send client to the memory arbiter. Default value assumes a clock of 80 MHz and a bus width of 128 bits. This value provides a bit rate of 1 Mb/s.



6.2.15.16 MCTP Rate Limiter Config 2 – 0x0013

Bits	Name	Default	Description
15	Decision Point	0b	Defines if, when credits are available, a full MCTP message is sent or a single VDM is sent.
14:0	MCTP Max Credits	0x5	Defines the maximum number of 16-byte credits that can be accumulated. These credits include the VDM header line (one line for each 64 byte VDM).

6.2.15.17 NC-SI Channel to Port Mapping – 0x0014

Bits	Name	Default	Description
15	Table Valid	0b	If set, the values below are used to define the port-to-channel mapping.
14:7	Reserved	0x0	Reserved
6:5	Port 1 Channel ID	0x1	The channel ID to use for port 1.
4	Port 1 Channel Enable	1b	0x0 = Disabled. 0x1 = Enabled.
3	Reserved	0b	Reserved
2:1	Port 0 Channel ID	0x0	The channel ID to use for port 0.
0	Port 0 Channel Enable	1b	0x0 = Disabled. 0x1 = Enabled.

6.2.15.18 Block CRC8 – 0x0015

Bits	Name	Default	Description
15:8	Block CRC8		CRC-8-CCITT: Start Section -> Word: Sideband Configuration Structure -> Block Length End Section -> Word: Sideband Configuration Structure -> NC-SI Channel to Port Mapping
7:0	Reserved	0x0	Block length in words.

6.2.16 Inventory Parameters Structure Section Summary Table

Word Offset	Description	Page
0x0000	Section Header	246
0x0001	Number of Inventory NVM Items + Media Type ID	246
0x0002	Media Type Length + Media Type Data Byte 1	246
0x0003	Media Type Data Byte 2	246
0x0004	Block CRC8	246



6.2.16.1 Section Header – 0x0000

Bits	Name	Default	Description
15:0	Block Length	0x4	Length in words of the section covered by CRC.

6.2.16.2 Number of Inventory NVM Items + Media Type ID – 0x0001

Bits	Name	Default	Description
15:8	Media Type ID	0x04	
7:0	Number of Inventory NVM Items	0x01	

6.2.16.3 Media Type Length + Media Type Data Byte 1 – 0x0002

Bits	Name	Default	Description
15:8	Media Type Data Byte 1	0x0	
7:0	Media Type Length	0x2	

6.2.16.4 Media Type Data Byte 2 – 0x0003

Bits	Name	Default	Description
15:8	Reserved	0x0	
7:0	Media Type Data Byte 2	0x1	

6.2.16.5 Block CRC8 – 0x0004

Bits	Name	Default	Description
15:8	Block CRC8		CRC-8-CCITT: Start Section -> Word: Inventory Parameters Structure -> Section Header End Section -> Word: Inventory Parameters Structure -> Media Type Data Byte 2
7:0	Reserved	0x0	Block length in words.

6.2.17 Extended Capabilities Structure Section



Summary Table

Word Offset	Description	Page
0x0000	Block Length	247
0x0001	Number of Inventory NVM Items + Capabilities Flags	247
0x0002	Device Name Length + Device Name Data Byte 1	247
0x0003	Capabilities Flags Data Byte 2+3	248
0x0004	Capabilities Flags Data Byte 4 + DCB Capabilities	248
0x0005	DCB Capabilities Length + DCB Capabilities Data	248
0x0006	E-Switch Capabilities ID + E-Switch Capabilities	248
0x0007	E-Switch Capabilities Data + NIC Partitioning Capa	248
0x0008	NIC Partitioning Capabilities Length + Data	250
0x0009	Block CRC8	250

6.2.17.1 Block Length – 0x0000

Bits	Name	Default	Description
15:0	Block Length	0x09	Length in words of the section covered by CRC.

6.2.17.2 Number of Inventory NVM Items + Capabilities Flags – 0x0001

Bits	Name	Default	Description
15:8	Capabilities Flags ID	0x0	
7:0	Number of Inventory NVM Items	0x4	

6.2.17.3 Device Name Length + Device Name Data Byte 1 – 0x0002

Bits	Name	Default	Description
15:8	Capabilities Flags Data Byte 1	0x02	
7:0	Capabilities Flags Length	0x04	



6.2.17.4 Capabilities Flags Data Byte 2+3 – 0x0003

Bits	Name	Default	Description
15:8	Device Name Data Byte 3	0x60	
7:0	Device Name Data Byte 2	0x1C	

6.2.17.5 Capabilities Flags Data Byte 4 + DCB Capabilities – 0x0004

Bits	Name	Default	Description
15:8	DCB Capabilities ID	0x1	
7:0	Capabilities Flags Data Byte 4	0x3B	

6.2.17.6 DCB Capabilities Length + DCB Capabilities Data – 0x0005

Bits	Name	Default	Description
15:8	DCB Capabilities Data	0x0F	
7:0	DCB Capabilities Length	0x1	

6.2.17.7 E-Switch Capabilities ID + E-Switch Capabilities – 0x0006

Bits	Name	Default	Description
15:8	E-Switch Capabilities Length	0x1	
7:0	E-Switch Capabilities ID	0x2	

6.2.17.8 E-Switch Capabilities Data + NIC Partitioning Capa – 0x0007

Bits	Name	Default	Description
15:8	NIC Partitioning Capabilities ID	0x3	



Bits	Name	Default	Description
7:0	E-Switch Capabilities Data	0x1	



6.2.17.9 NIC Partitioning Capabilities Length + Data – 0x0008

Bits	Name	Default	Description
15:8	NIC Partitioning Capabilities Data	0x0	
7:0	NIC Partitioning Capabilities Length	0x1	

6.2.17.10 Block CRC8 – 0x0009

Bits	Name	Default	Description
15:8	Block CRC8		CRC-8-CCITT: Start Section -> Word: Extended Capabilities Structure -> Block Length End Section -> Word: Extended Capabilities Structure -> NIC Partitioning Capabilities Length + Data
7:0	Reserved	0x0	Block length in words.

6.2.18 Partition Information Structure Section Summary Table


Word Offset	Description	Page
0x0000	Block Length	250
0x0001	Number of Inventory NVM Items + LAN Personality ID	250
0x0002	LAN Personality Length + LAN Personality Data	251
0x0003	Block CRC8	252

6.2.18.1 Block Length – 0x0000

Bits	Name	Default	Description
15:0	Block Length	0x26	Length in words of the section covered by CRC.

6.2.18.2 Number of Inventory NVM Items + LAN Personality ID – 0x0001

Bits	Name	Default	Description
15:8	LAN Personality ID	0x0	
7:0	Number of Inventory NVM Items	0x1	



6.2.18.3 LAN Personality Length + LAN Personality Data – 0x0002

Raw data module length: variable



6.2.18.4 Block CRC8 – 0x0003

Bits	Name	Default	Description
15:8	Block CRC8		CRC-8-CCITT: Start Section -> Word: Partition Information Structure -> Block Length End Section -> Word: Partition Information Structure -> LAN Personality Length + LAN Personality Data
7:0	Reserved	0x0	Block length in words.

6.2.19 Flexible TCO Filter Configuration Structure Section Summary Table

Word Offset	Description	Page
0x0000	Block Length	252
0x0001	Flexible Filter Length and Control	252
0x0002 + 1*n, n=0...7	Flexible Filter Enable Mask	252
0x000A + 1*n, n=0...63	Flexible Filter Data	253
0x004A	Block CRC8	253

6.2.19.1 Block Length – 0x0000

Bits	Name	Default	Description
15:0	Block Length		Length in words of the section covered by CRC. Section length in words.

6.2.19.2 Flexible Filter Length and Control – 0x0001

Bits	Name	Default	Description
15:8	Flexible Filter Length (bytes)	0x0	
7:5	Reserved	0x0	
4	Last Filter	1b	
3:2	Filter Index (0-3)	0x1	
1	Apply Filter to LAN 1	0b	
0	Apply Filter to LAN 0	0b	

6.2.19.3 Flexible Filter Enable Mask[n] (0x0002 + 1*n,



n=0...7)

Bits	Name	Default	Description
15:0	Flexible Filter Enable Mask	0x0000	

6.2.19.4 Flexible Filter Data[n] (0x000A + 1*n, n=0...63)

Bits	Name	Default	Description
15:0	Flexible Filter Data	0x0000	

6.2.19.5 Block CRC8 – 0x004A

Bits	Name	Default	Description
15:8	Block CRC8		CRC-8-CCITT: Start Section -> Word: Flexible TCO Filter Configuration Structure -> Block Length End Section -> Word: Flexible TCO Filter Configuration Structure -> Flexible Filter Data
7:0	Reserved	0x0	Block length in words.

6.2.20 LESM Configurations Section Summary Table

Word Offset	Description	Page
0x0000	Section Header	253
0x0001	LESM Global Configurations	254
0x0002	LESM Per State Configurations	254
0x0003	LESM Per State Configurations LSB	254
0x0004	LESM Per State Configurations MSB	254
0x0005	Section Footer	254

6.2.20.1 Section Header – 0x0000

Bits	Name	Default	Description
15:8	Reserved	0x0	Reserved
7:0	Block Length	0x5	Length in words of the section covered by CRC. Block length in words (including the CRC word and not the length word).



6.2.20.2 LESM Global Configurations – 0x0001

Bits	Name	Default	Description
15:6	Reserved	0x0	
5:1	10G Only Counter	0x3	Number of iterations to go through before enabling 1 GbE parallel detection.
0	LESM Global Configurations	1b	1b = LESM is enabled. 0b = LESM is disabled.

6.2.20.3 LESM Per State Configurations – 0x0002

Bits	Name	Default	Description
15:1	[New Field]	0x3	Number of iterations to go through before enabling 1 GbE parallel detection.
0	State Enable	1b	1b = LESM is enabled. 0b = LESM is disabled.

6.2.20.4 LESM Per State Configurations LSB – 0x0003

Bits	Name	Default	Description
15:0	Register Pointer LSB	0x0	

6.2.20.5 LESM Per State Configurations MSB – 0x0004

Bits	Name	Default	Description
15:0	Register Pointer MSB	0x0	

6.2.20.6 Section Footer – 0x0005

Bits	Name	Default	Description
15:8	Section Footer		CRC-8-CCITT: Start Section -> Word: LESM Configurations -> Section Header End Section -> Word: LESM Configurations -> LESM Per State Configurations MSB CRC8 of the previous section.
7:0	Reserved	0x0	Reserved

6.2.21 VPD Module Section Summary Table

6.2.22 PBA Number Module Section Summary



Table

Word Offset	Description	Page
0x0000	PBA Section Length	255
0x0001	Word1	255
0x0002	Word2	255
0x0003	Word3	255
0x0004	Word4	255
0x0005	Word5	256

6.2.22.1 PBA Section Length – 0x0000

Bits	Name	Default	Description
15:0	PBA Section Length Field	0x6	Length in words of the PBA block.

6.2.22.2 Word1 – 0x0001

Bits	Name	Default	Description
15:0	Word1 Field		PBA number stored in hexadecimal ASCII values.

6.2.22.3 Word2 – 0x0002


Bits	Name	Default	Description
15:0	Word2 Field		PBA number stored in hexadecimal ASCII values.

6.2.22.4 Word3 – 0x0003

Bits	Name	Default	Description
15:0	Word3 Field		PBA number stored in hexadecimal ASCII values.

6.2.22.5 Word4 – 0x0004

Bits	Name	Default	Description
15:0	Word4 Field		PBA number stored in hexadecimal ASCII values.



6.2.22.6 Word5 – 0x0005

Bits	Name	Default	Description
15:0	Word5 Field		PBA number stored in hexadecimal ASCII values.

6.2.23 Mini Loader Module Section Summary Table

Word Offset	Description	Page
0x0000	Mini Loader Section Header	256
0x0001	Mini Loader Code	256
0x0002	Mini Loader Section Footer	257

6.2.23.1 Mini Loader Section Header – 0x0000

Bits	Name	Default	Description
15:0	Block Length		Length in words of the section covered by CRC.

6.2.23.2 Mini Loader Code – 0x0001

Raw data module length: variable



6.2.23.3 Mini Loader Section Footer – 0x0002

Bits	Name	Default	Description
15:8	Block CRC8		CRC-8-CCITT: Start Section -> Word: Mini Loader Module -> Mini Loader Section Header End Section -> Word: Mini Loader Module -> Mini Loader Code
7:0	Reserved	0x0	

6.2.24 Internal PHY Indirect List Section Summary Table

Word Offset	Description	Page
0x0000	Section Length	257
0x0001	Reg Write Indirect List 1	257
0x0002	Reg Write Indirect List 2	257
0x0003	Reg Write Indirect List 3	257

6.2.24.1 Section Length – 0x0000

Bits	Name	Default	Description
15:0	Section Length		Length in: 2 Bytes unit - 1 First Section -> Word: Internal PHY Indirect List -> Section Length Last Section -> Word: Internal PHY Indirect List -> Reg Write Indirect List 3

6.2.24.2 Reg Write Indirect List 1 – 0x0001

Raw data module length: variable

6.2.24.3 Reg Write Indirect List 2 – 0x0002

Raw data module length: variable

6.2.24.4 Reg Write Indirect List 3 – 0x0003

Raw data module length: variable



6.2.25 Reserved/Legacy Module 1 Section Summary Table

Word Offset	Description	Page
0x0000	Section Length	258
0x0001	Microcode	258

6.2.25.1 Section Length – 0x0000

Bits	Name	Default	Description
15:0	Section Length		Length in: 2 Bytes unit - 1 First Section -> Word: Reserved/Legacy Module 1 -> Section Length Last Section -> Word: Reserved/Legacy Module 1 -> Microcode

6.2.25.2 Microcode – 0x0001

Raw data module length: variable

6.2.26 DL Scripts Module Section Summary Table

Word Offset	Description	Page
0x0000	DNL Scripts	258

6.2.26.1 DNL Scripts - 0x0000

Raw data module length: variable



6.2.27 2nd Init Module Section Summary Table

Word Offset	Description	Page
0x2000	data	259

6.2.27.1 Data – 0x2000

Bits	Name	Default	Description
15:0	Data	0xFFFF	


6.2.28 Firmware Code Dummy CSS Section Summary Table

64 x 4 KB for 4 MB Flash and 61 x 4 KB for 2 MB Flash.

Word Offset	Description	Page
0x0000	Firmware With Dummy Header	259

6.2.28.1 Firmware With Dummy Header – 0x0000

Raw data module length: variable



6.2.29 FW_Image Section Summary Table

Word Offset	Description	Page
0x0000	moduleTypeL	262
0x0001	moduleTypeH	262
0x0002	headerLenL	262
0x0003	headerLenH	262
0x0004	headerVersionL	262
0x0005	headerVersionH	262
0x0006	moduleIDL	262
0x0007	moduleIDH	263
0x0008	moduleVendorL	263
0x0009	moduleVendorH	263
0x000A	dateL	263
0x000B	dateH	263
0x000C	sizeL	264
0x000D	sizeH	264
0x000E	keySizeL	264
0x000F	keySizeH	264
0x0010	modulusSizeL	264
0x0011	modulusSizeH	264
0x0012	exponentSizeL	264
0x0013	exponentSizeH	265
0x0014	lad_srevL	265
0x0015	lad_srevH	265
0x0016	lad_nvmsizeL	265
0x0017	lad_nvmsizeH	265
0x0018	lad_fw_entry_offsetL	265
0x0019	lad_fw_entry_offsetH	265
0x001A	reservedL	266
0x001B	reservedH	266
0x001C	lad_image_unique_idL	266
0x001D	lad_image_unique_idH	266
0x001E	lad_module_idL	266
0x001F	lad_module_idH	266
0x0020 + 1*n, n=0...31	Reserved	266
0x0040 + 1*n, n=0...127	RSA Public Key	267
0x00C0	RSA ExponentL	267
0x00C1	RSA ExponentH	267



Word Offset	Description	Page
0x00C2 + 1*n, n=0...127	Encrypted SHA256 Hash	267
0x0142	Device Blank NVM Device ID	267
0x0143	Max Module AreaL	267
0x0144	Max Module AreaH	267
0x0145	Current Module AreaL	268
0x0146	Current Module AreaH	268
0x0147	Format Version + CRC	268
0x0148	Code Revision	268
0x0149	Reserved Spare Word	268
0x014A	FW Raw Data	268



6.2.29.1 moduleTypeL – 0x0000

Bits	Name	Default	Description
15:0	moduleType	0x6	

6.2.29.2 moduleTypeH – 0x0001

Bits	Name	Default	Description
15:0	moduleTypeH		

6.2.29.3 headerLenL – 0x0002

Bits	Name	Default	Description
15:0	headerLenL	0xA1	

6.2.29.4 headerLenH – 0x0003

Bits	Name	Default	Description
15:0	headerLenH		

6.2.29.5 headerVersionL – 0x0004

Bits	Name	Default	Description
15:0	headerVersionL	0x00010000	

6.2.29.6 headerVersionH – 0x0005

Bits	Name	Default	Description
15:0	headerVersionH		

6.2.29.7 moduleIDL – 0x0006

Bits	Name	Default	Description
15:0	moduleIDL	0x0000	



6.2.29.8 moduleIDH – 0x0007

Bits	Name	Default	Description
15	signMode	0x1	
14:0	moduleIDH	0x0	

6.2.29.9 moduleVendorL – 0x0008

Bits	Name	Default	Description
15:0	moduleVendorL	0x00008086	

6.2.29.10 moduleVendorH – 0x0009

Bits	Name	Default	Description
15:0	moduleVendorH		

6.2.29.11 dateL – 0x000A

Bits	Name	Default	Description
15:0	DateL	0x20141217	0xMMDD

6.2.29.12 dateH – 0x000B

Bits	Name	Default	Description
15:0	DateH		



6.2.29.13 sizeL – 0x000C

Bits	Name	Default	Description
15:0	sizeL	0x0001E800	

6.2.29.14 sizeH – 0x000D

Bits	Name	Default	Description
15:0	sizeH		

6.2.29.15 keySizeL – 0x000E

Bits	Name	Default	Description
15:0	keySizeL	0x40	

6.2.29.16 keySizeH – 0x000F

Bits	Name	Default	Description
15:0	keySizeH		

6.2.29.17 modulusSizeL – 0x0010

Bits	Name	Default	Description
15:0	modulusSizeL	0x40	

6.2.29.18 modulusSizeH – 0x0011

Bits	Name	Default	Description
15:0	modulusSizeH		

6.2.29.19 exponentSizeL – 0x0012

Bits	Name	Default	Description
15:0	exponentSizeL	0x1	



6.2.29.20 exponentSizeH – 0x0013

Bits	Name	Default	Description
15:0	exponentSizeH		

6.2.29.21 lad_srevL – 0x0014

Bits	Name	Default	Description
15:0	lad_srevL	0x1	

6.2.29.22 lad_srevH – 0x0015

Bits	Name	Default	Description
15:0	lad_srevH		

6.2.29.23 lad_nvmsizeL – 0x0016

Bits	Name	Default	Description
15:0	lad_nvmsizeL	0x3d000	

6.2.29.24 lad_nvmsizeH -- 0x0017

Bits	Name	Default	Description
15:0	lad_nvmsizeH		

6.2.29.25 lad_fw_entry_offsetL – 0x0018

Bits	Name	Default	Description
15:0	lad_fw_entry_offsetL	0x0000014c	

6.2.29.26 lad_fw_entry_offsetH – 0x0019

Bits	Name	Default	Description
15:0	lad_fw_entry_offsetH		



6.2.29.27 reservedL – 0x001A

Bits	Name	Default	Description
15:0	Reserved	0x0	

6.2.29.28 reservedH – 0x001B

Bits	Name	Default	Description
15:0	Reserved	0x0	

6.2.29.29 lad_image_unique_idL – 0x001C

Bits	Name	Default	Description
15:0	lad_image_unique_idL	0x0	

6.2.29.30 lad_image_unique_idH – 0x001D

Bits	Name	Default	Description
15:0	lad_image_unique_idH		

6.2.29.31 lad_module_idL – 0x001E

Bits	Name	Default	Description
15:0	lad_module_idL	0x1	0x1 = Firmware. 0x5 = PHY firmware. 0xFE = Option ROM.

6.2.29.32 lad_module_idH – 0x001F

Bits	Name	Default	Description
15:0	lad_module_idH		

6.2.29.33 reserved[n] (0x0020 + 1*n, n=0...31)

Bits	Name	Default	Description
15:0	Reserved	0x0	



6.2.29.34 RSA Public Key [n] (0x0040 + 1*n, n=0...127)

Bits	Name	Default	Description
15:0	RSA Public Key	0x0	

6.2.29.35 RSA ExponentL – 0x00C0

Bits	Name	Default	Description
15:0	RSA ExponentL	0x0	

6.2.29.36 RSA ExponentH – 0x00C1

Bits	Name	Default	Description
15:0	RSA ExponentH	0x0	

6.2.29.37 Encrypted SHA256 Hash[n] (0x00C2 + 1*n, n=0...127)

Bits	Name	Default	Description
15:0	RSA Public Key	0x0	

6.2.29.38 Device Blank NVM Device ID – 0x0142

Bits	Name	Default	Description
15:0	Device Blank NVM Device ID	0x15A7	

6.2.29.39 Max Module AreaL – 0x0143

Bits	Name	Default	Description
15:0	Max Module AreaL	0x40000	

6.2.29.40 Max Module AreaH – 0x0144

Bits	Name	Default	Description
15:0	Max Module AreaH		



6.2.29.41 Current Module AreaL – 0x0145

Bits	Name	Default	Description
15:0	Current Module AreaL	0x7A000	

6.2.29.42 Current Module AreaH – 0x0146

Bits	Name	Default	Description
15:0	Current Module AreaH		

6.2.29.43 Format Version + CRC – 0x0147

Bits	Name	Default	Description
15	CRC Field Used	1b	0x0 = CRC not used. 0x1 = CRC used.
14:8	Format Version	0x2	
7:0	CRC8		CRC-8-CCITT: Start Section -> Word: FW_Image -> moduleTypeL End Section -> Word: FW_Image -> FW Raw Data

6.2.29.44 Code Revision – 0x0148

Bits	Name	Default	Description
15:8	Major Revision	0x2	
7:0	Minor Revision	0x14	

6.2.29.45 Reserved Spare Word – 0x0149

Bits	Name	Default	Description
15:0	Reserved Spare Word	0x0	

6.2.29.46 FW Raw Data – 0x014A

Raw data module length: variable



6.2.30 PXE_OROM Dummy CSS Section Summary Table

64 x 4 KB for 4 MB Flash and 61 x 4 KB for 2 MB Flash.

Word Offset	Description	Page
0x0000	Dummy OROM	269

6.2.30.1 Dummy OROM – 0x0000

Raw data module length: variable




6.2.31 OROM_500 Section Summary Table

Word Offset	Description	Page
0x0000	moduleTypeL	272
0x0001	moduleTypeH	272
0x0002	headerLenL	272
0x0003	headerLenH	272
0x0004	headerVersionL	272
0x0005	headerVersionH	272
0x0006	moduleIDL	272
0x0007	moduleIDH	273
0x0008	moduleVendorL	273
0x0009	moduleVendorH	273
0x000A	dateL	273
0x000B	dateH	273
0x000C	sizeL	274
0x000D	sizeH	274
0x000E	keySizeL	274
0x000F	keySizeH	274
0x0010	modulusSizeL	274
0x0011	modulusSizeH	274
0x0012	exponentSizeL	274
0x0013	exponentSizeH	275
0x0014	lad_srevL	275
0x0015	lad_srevH	275
0x0016	reserved1L	275
0x0017	reserved1H	275
0x0018	lad_fw_entry_offsetL	275
0x0019	lad_fw_entry_offsetH	275
0x001A	reserved2L	276
0x001B	reserved2H	276
0x001C	lad_image_unique_idL	276
0x001D	lad_image_unique_idH	276
0x001E	lad_module_idL	276
0x001F	lad_module_idH	276
0x0020 + 1*n, n=0...31	Reserved	277
0x0040 + 1*n, n=0...127	RSA Public Key	277
0x00C0	RSA ExponentL	277
0x00C1	RSA ExponentH	277



Word Offset	Description	Page
0x00C2 + 1*n, n=0...127	Encrypted SHA256 Hash	277
0x0142	Device Blank NVM Device ID	277
0x0143	Max Module AreaL	277
0x0144	Max Module AreaH	278
0x0145	Current Module AreaL	278
0x0146	Current Module AreaH	278
0x0147	Format Version + CRC	278
0x0148	Code Revision	278
0x0149	Reserved Spare Word	278
0x014A	OROM Raw Data	278



6.2.31.1 moduleTypeL – 0x0000

Bits	Name	Default	Description
15:0	moduleType	0x6	

6.2.31.2 moduleTypeH – 0x0001

Bits	Name	Default	Description
15:0	moduleTypeH		

6.2.31.3 headerLenL – 0x0002

Bits	Name	Default	Description
15:0	headerLenL	0xa1	

6.2.31.4 headerLenH – 0x0003

Bits	Name	Default	Description
15:0	headerLenH		

6.2.31.5 headerVersionL – 0x0004

Bits	Name	Default	Description
15:0	headerVersionL	0x00010000	

6.2.31.6 headerVersionH – 0x0005

Bits	Name	Default	Description
15:0	headerVersionH		

6.2.31.7 moduleIDL – 0x0006

Bits	Name	Default	Description
15:0	moduleIDL	0x0	



6.2.31.8 moduleIDH – 0x0007

Bits	Name	Default	Description
15	signMode	0x0	
14:0	moduleIDH		

6.2.31.9 moduleVendorL – 0x0008

Bits	Name	Default	Description
15:0	moduleVendorL	0x00008086	

6.2.31.10 moduleVendorH – 0x0009

Bits	Name	Default	Description
15:0	moduleVendorH		

6.2.31.11 dateL – 0x000A

Bits	Name	Default	Description
15:0	DateL	0x20130530	0xMMDD

6.2.31.12 dateH – 0x000B

Bits	Name	Default	Description
15:0	DateH		



6.2.31.13 sizeL – 0x000C

Bits	Name	Default	Description
15:0	sizeL	0x0001D400	

6.2.31.14 sizeH – 0x000D

Bits	Name	Default	Description
15:0	sizeH		

6.2.31.15 keySizeL – 0x000E

Bits	Name	Default	Description
15:0	keySizeL	0x40	

6.2.31.16 keySizeH – 0x000F

Bits	Name	Default	Description
15:0	keySizeH		

6.2.31.17 modulusSizeL – 0x0010

Bits	Name	Default	Description
15:0	modulusSizeL	0x40	

6.2.31.18 modulusSizeH – 0x0011

Bits	Name	Default	Description
15:0	modulusSizeH		

6.2.31.19 exponentSizeL – 0x0012

Bits	Name	Default	Description
15:0	exponentSizeL	0x1	



6.2.31.20 exponentSizeH – 0x0013

Bits	Name	Default	Description
15:0	exponentSizeH		

6.2.31.21 lad_srevL – 0x0014

Bits	Name	Default	Description
15:0	lad_srevL	0x0	

6.2.31.22 lad_srevH – 0x0015

Bits	Name	Default	Description
15:0	lad_srevH		

6.2.31.23 reserved1L – 0x0016

Bits	Name	Default	Description
15:0	reserved	0x0	

6.2.31.24 reserved1H – 0x0017

Bits	Name	Default	Description
15:0	reserved	0x0	

6.2.31.25 lad_fw_entry_offsetL – 0x0018

Bits	Name	Default	Description
15:0	lad_fw_entry_offsetL	0x14C	

6.2.31.26 lad_fw_entry_offsetH – 0x0019

Bits	Name	Default	Description
15:0	lad_fw_entry_offsetH		



6.2.31.27 reserved2L – 0x001A

Bits	Name	Default	Description
15:0	Reserved	0x0	

6.2.31.28 reserved2H – 0x001B

Bits	Name	Default	Description
15:0	Reserved	0x0	

6.2.31.29 lad_image_unique_idL – 0x001C

Bits	Name	Default	Description
15:0	lad_image_unique_idL	0x0	

6.2.31.30 lad_image_unique_idH – 0x001D

Bits	Name	Default	Description
15:0	lad_image_unique_idH		

6.2.31.31 lad_module_idL – 0x001E

Bits	Name	Default	Description
15:0	lad_module_idL	0xFE	0x1 Firmware image. 0x5 PHY firmware. 0xFE Option ROM.

6.2.31.32 lad_module_idH – 0x001F

Bits	Name	Default	Description
15:0	lad_module_idH		



6.2.31.33 reserved[n] (0x0020 + 1*n, n=0...31)

Bits	Name	Default	Description
15:0	reserved	0x0	

6.2.31.34 RSA Public Key [n] (0x0040 + 1*n, n=0...127)

Bits	Name	Default	Description
15:0	RSA Public Key	0x0	

6.2.31.35 RSA ExponentL – 0x00C0

Bits	Name	Default	Description
15:0	RSA ExponentL	0x0	

6.2.31.36 RSA ExponentH – 0x00C1

Bits	Name	Default	Description
15:0	RSA ExponentH	0x0	

6.2.31.37 Encrypted SHA256 Hash[n] (0x00C2 + 1*n, n=0...127)

Bits	Name	Default	Description
15:0	RSA Public Key	0x0	

6.2.31.38 Device Blank NVM Device ID – 0x0142

Bits	Name	Default	Description
15:0	Device Blank NVM Device ID	0x154B	

6.2.31.39 Max Module AreaL – 0x0143

Bits	Name	Default	Description
15:0	Max Module AreaL	0x1000	



6.2.31.40 Max Module AreaH – 0x0144

Bits	Name	Default	Description
15:0	Max Module AreaH	0x0009	

6.2.31.41 Current Module AreaL – 0x0145

Bits	Name	Default	Description
15:0	Current Module AreaL	0x1000	

6.2.31.42 Current Module AreaH – 0x0146

Bits	Name	Default	Description
15:0	Current Module AreaH	0x0000	

6.2.31.43 Format Version + CRC – 0x0147

Bits	Name	Default	Description
15	CRC Field Used	0b	0b = CRC not used. 1b = CRC used.
14:8	Format Version	0x2	
7:0	CRC8	0xFF	

6.2.31.44 Code Revision – 0x0148

Bits	Name	Default	Description
15:8	Major Revision	0x0	
7:0	Minor Revision	0x0	

6.2.31.45 Reserved Spare Word – 0x0149

Bits	Name	Default	Description
15:0	Reserved Spare Word	0x0	

6.2.31.46 OROM Raw Data – 0x014A



Raw data module length: variable

6.2.32 Free Space Provisioning Module_512 Section Summary Table

64 x 4 KB for 4 MB Flash and 61 x 4 KB for 2 MB Flash.


Word Offset	Description	Page
0x0000	Reserved	279

6.2.32.1 Reserved – 0x0000

Bits	Field Name	Default NVM Value	Description
15:0	Reserved	0xFFFF	

6.2.33 Reserved/Legacy Module 2 Section Summary Table

64 x 4 KB for 4 MB Flash and 61 x 4 KB for 2 MB Flash.



NOTE: *This page intentionally left blank.*



7.0 Inline Functions

7.1 Receive Functionality

Packet reception consists of:

- Recognizing the presence of a packet on the wire ([Section 7.1.1](#))
- Parsing of the packet header ([Section 7.1.2](#))
- Performing address filtering ([Section 7.1.2](#))
- Checksum off-loads ([Section 7.1.6](#))
- DMA queue assignment ([Section 7.1.3](#))
- Storing the packet in the receive data FIFO ([Section 7.6.3.1](#))
- Transferring the data to assigned receive queues in host memory ([Section 7.1.4](#))
- Optional Receive Side Coalescing ([Section 7.9](#))
- Updating the state of a receive descriptor ([Section 7.1.5](#)).

7.1.1 MAC Layer - Receive

7.1.1.1 Packet Acceptance Criteria

In addition to the filtering rules described in the next sections, a packet must also meet the following criteria in order to be accepted by the device:

1. Normally, only good packets are received (packets with none of the following errors: Under Size Error, Over Size Error, Packet Error, Length Error and CRC Error). However, if the store-bad-packet bit is set (*FCTRL.SBP*), then bad packets that don't pass the filter function are stored in host memory. Packet errors are indicated by error bits in the receive descriptor (*RDESC.ERRORS*). It is possible to receive all packets, regardless of whether they are bad, by setting the promiscuous enables bit and the store-bad-packet bit. In this case, bad packets are queued according to the same rules as regular packets.
2. Min. Packet Size (Runt packets) — Rx packets, smaller than 48 bytes, cannot be posted to host memory regardless of save bad frame setting.
3. Max Packet Size — Any Rx packet posted from the MAC unit to the DMA unit cannot exceed 15.5 KB. Further restrictions per queue are described in [Section 7.1.4](#).
4. CRC errors before the Start Frame Delimiter (SFD) are ignored. All packets must have a valid SFD in order to be recognized by the device (even bad packets).

7.1.1.2 CRC Strip

The Intel® Xeon® Processor D-1500 Product Family LAN controller potentially strips the L2 CRC on incoming packets.

CRC strip is enabled by the *HLREG0.RXCRCSTRP* bit. When set, CRC is stripped from all received packets.

The global CRC strip bit (*HLREGO.RXCRCSTRP*) must be set in the following cases where the packet changes before being handled to the software device driver:

- RSC is enabled in any queue.
- VLAN is hidden (*PFQDE.HIDE_VLAN = 1b*) in any queue.
- E-tag is removed in any queue (*PFQDE.STRIP_TAG* is set).
- L2 tags are stripped from packets (*PFQDE.STRIP_TAG = 1b*) in any queue. MACSec is enabled.
- Time stamp is added to the packets in any TC (any bit of *TSYNCRXCTL.TSIP_UT_EN* or *TSYNCRXCTL.TSIP_UP_EN* is set).
- Short received packets are padded (*RDRXCTL.PSP* is set).

7.1.2 Packet Filtering

A received packet goes through up to three stages of filtering as depicted in [Figure 7.1](#). It describes a switch-like structure that is used in virtualization mode to route packets between the network port (top of drawing) and one of many virtual ports (bottom of drawing), where each virtual port might be associated with a Virtual Machine (VM), a Virtual Machine Monitor (VMM), or the like. The three stages are:

1. **First Stage — Admission Control:** Ensure that the packet should be received by the port. This is done by a set of L2 filters and is described in detail in this section.
2. **Second Stage — Pooling:** This stage is specific to virtualization environments and defines the virtual ports (called pools in this document) that are the targets for the Rx packet. A packet can be associated with any number of ports/pools and the selection process is described in [Section 7.1.3.2](#). In non-virtualization mode, this stage is skipped and all the queues used in the next stage are considered as part of the same default pool.

Note: A pool is equivalent to a VSI as defined in IEEE 802.1Qbg specification.

3. **Third Stage — Queueing:** A receive packet that successfully passed the Rx filters is associated with one of many receive descriptor queues as described in [Section 7.1.3](#).




Figure 7.1. Stages in Packet Filtering - Virtualization Mode

The receive packet filtering role is to determine which of the incoming packets are allowed to pass to the local machine and which of the incoming packets should be dropped since they are not targeted to the local machine. Received packets that are targeted for the local machine can be destined to the host, to a manageability controller, or to both. This section describes how host filtering is done, and the interaction with management filtering.



As shown in [Figure 7.2](#), host filtering is done in two stages:

1. Packets are filtered by L2 filters (Ethernet MAC address, unicast/multicast/broadcast). See [Section 7.1.2.1](#).
2. Packets are filtered by VLAN if a VLAN tag is present. See [Section 7.1.2.2](#).

A packet is not forwarded to the host if any of the following occurs:

- The packet does not pass L2 filters, as described in [Section 7.1.2.1](#).
- The packet does not pass VLAN filtering, as described in [Section 7.1.2.2](#).
- The packet passes manageability filtering and the manageability filters determine that the packet should not pass to the host (see MNGONLY register and [Section 10.3.5.2](#)).

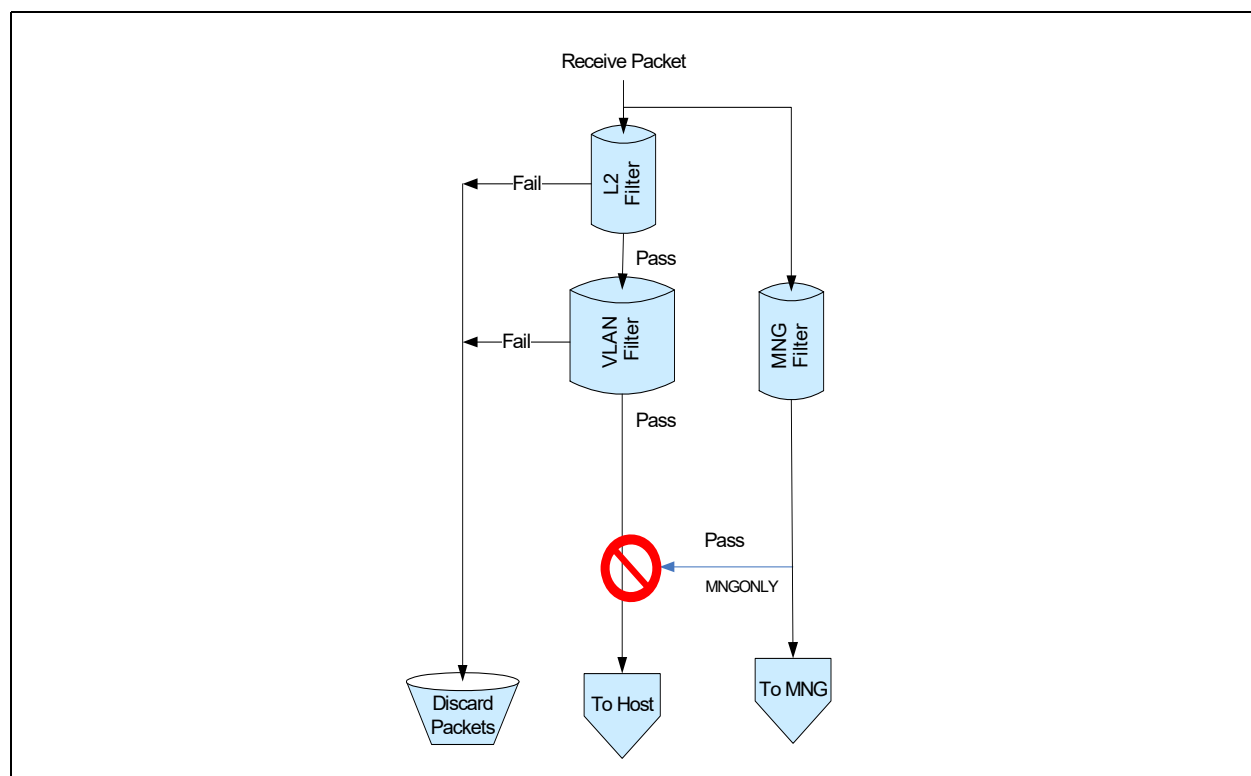


Figure 7.2. Rx Filtering Flow Chart

7.1.2.1 L2 Filtering

A packet passes successfully through L2 Ethernet MAC address filtering if any of the following conditions are met:

- Unicast packet filtering — Promiscuous unicast filtering is enabled ($FCTRL.UPE=1b$) or the packet passes unicast MAC filters (host).
- Multicast packet filtering — Promiscuous multicast filtering is enabled ($FCTRL.MPE=1b$) or the packet matches one of the multicast filters.
- Broadcast packet filtering to host — Promiscuous multicast filtering is enabled ($FCTRL.MPE=1b$) or Broadcast Accept Mode is enabled ($FCTRL.BAM = 1b$).



7.1.2.1.1 Unicast Filter

The Ethernet MAC address is checked against the 128 host unicast addresses and 4 KB hash-based unicast address filters (if enabled). The host unicast addresses are controlled by the host interface. The destination address of an incoming packet must exactly match one of the pre-configured host address filters. These addresses can be unicast or multicast. Those filters are configured through Receive Address Low (RAL), Receive Address High (RAH), In addition, there are 4 KB unicast hash filters used for host defined by the PFUTA registers. The unicast hash filters are useful mainly for virtualization settings in those cases that more than 128 filters might be required.

Promiscuous Unicast — Receive all unicasts. Promiscuous unicast mode is usually used when the LAN device is used as a sniffer.

7.1.2.1.2 Multicast Filter (Partial)

The 12-bit portion of the incoming packet multicast address must be set in the multicast filter address table (MTA) in order to pass the partial multicast filter. The bits (out of 48 bits of the destination address) used to index the MTA table can be selected by the *MO* field in the MCSTCTRL register.

Promiscuous Multicast — Receive all multicasts. Promiscuous multicast mode is usually used when the LAN device is used as a sniffer.

7.1.2.2 VLAN Filtering

The Intel® Xeon® Processor D-1500 Product Family LAN controller provides exact VLAN filtering as follows:

- Host VLAN filters are programmed by the VFTA[n] registers.
- A VLAN match might relate to the *DEI* bit in the VLAN header. It is enabled for host filtering only by the *VLNCTRL.DEIEN* while the expected value is defined by the *VLNCTRL.DEI*.

If double VLAN is enabled (see [Section 7.4.5](#)), filtering is done on the second (internal) VLAN tag. All the filtering functions of the Intel® Xeon® Processor D-1500 Product Family LAN controller ignore the first (external) VLAN in this mode.

A receive packet that passes MAC Address filtering successfully is subjected to VLAN header filtering as illustrated in [Figure 7.3](#):

1. If the packet does not have a VLAN header, it passes to the next filtering stage.
2. Else, if VLAN filters are not enabled (*VLNCTRL.VFE* = 0b), the packet is forwarded to the next filtering stage.
3. Else, if the packet matches an enabled VLAN filter and DEI checking (if enabled), the packet is forwarded to the next filtering stage.
4. Otherwise, the packet is dropped.

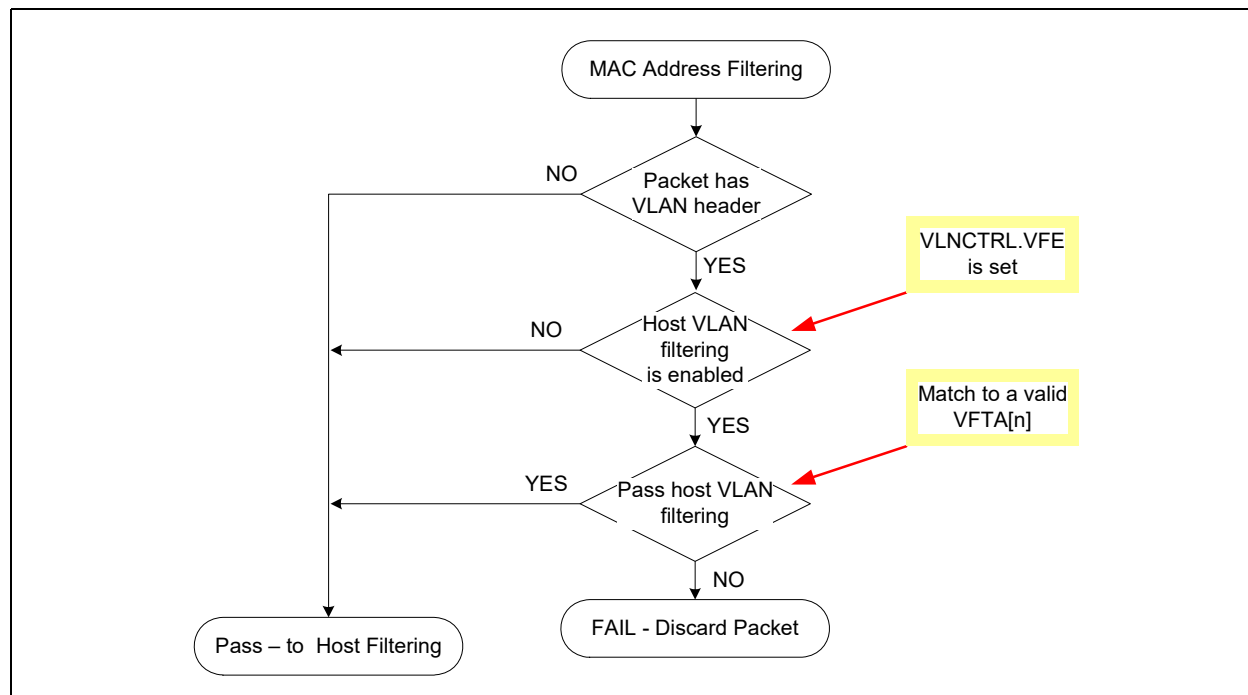


Figure 7.3. VLAN Filtering

7.1.2.3 E-tag Filtering

If the *PFVTCTL.POOLING_MODE* is E-tag (01b), then special L2 filtering rules are applied to tagged packets. If the *VT_mode* is E-tag (01b), then special L2 filtering rules are applied to packets with E-tag.

- If the tag matches one of the RAH registers used for tag filtering (*RAH.ADTYPE* = 1), then it is considered as a packet that passed L2.
- If the *FCTRL.TPE* bit is set (Tag Promiscuous Enable), all the tagged packets passes L2 filtering.
- Otherwise, packets with E-tag (if *PFVTCTL.POOLING_MODE* is E-tag (01b)) are dropped.

Note: E-tag packets are not expected when double VLAN are use, thus *PFVTCTL.POOLING_MODE* should be cleared if *CTRL_EXT.EXTENDED_VLAN* is set.

7.1.2.4 Manageability / Host Filtering

The host and manageability filtering process are mostly independent. Each entity defines which packet it should receive. The only exception is that the manageability filters can define a packet as exclusive and thus prevent it from reaching the host. See [Section 10.3.5.2](#) for details. [Figure 7.4](#) describes this flow.

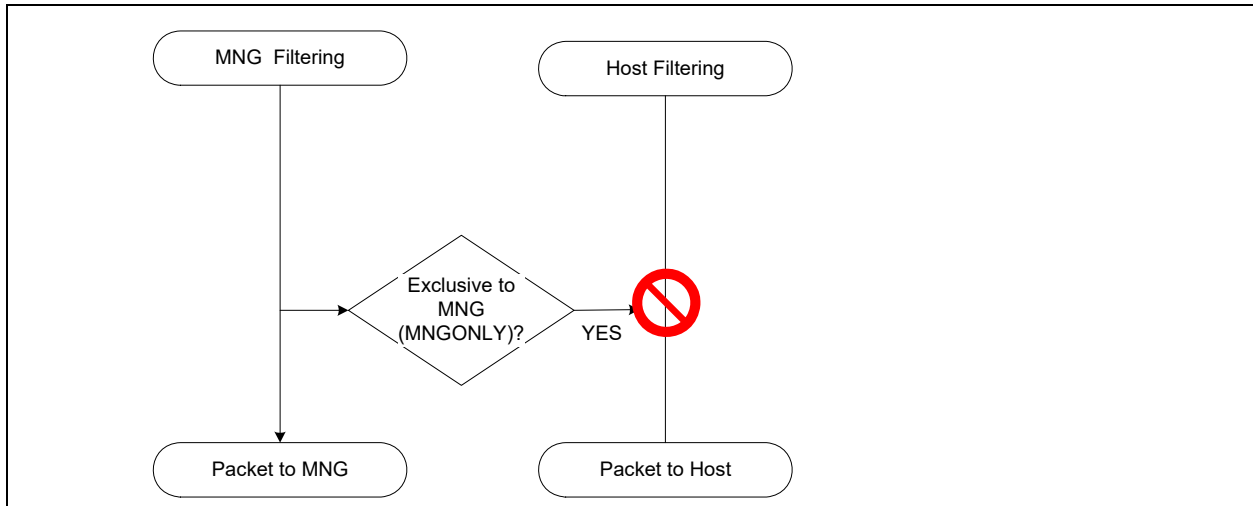


Figure 7.4. Manageability / Host Filtering

7.1.3 Rx Queues Assignment

The following filters/mechanisms determine the destination of a received packet. These filters are described briefly while more detailed descriptions are provided in the following sections:

- **Virtualization** — In a virtual environment, DMA resources are shared between more than one software entity (operating system and/or device driver). This is done by allocating receive descriptor queues to virtual partitions (VMM, IOVM, VMs, or VFs). Allocating queues to virtual partitions is done in sets, each with the same number of queues, called queue pools, or pools. Virtualization assigns to each received packet one or more pool indices. Packets are routed to a pool based on their pool index and other considerations such as RSS. See [Section 7.1.3.2](#) for more on routing for virtualization.
- **Receive Side Scaling (RSS)** — RSS distributes packet processing between several processor cores by assigning packets into different descriptor queues. RSS assigns to each received packet an RSS index. Packets are routed to one queue from a set of Rx queues based on their RSS index and other considerations such as virtualization. See [Section 7.1.3.6](#) for details.
- **L2 Ethertype Filters** — These filters identify packets by their L2 Ethertype and assigns them to receive queues. Examples of possible uses are LLDP packets, and 802.1X packets. See [Section 7.1.3.3](#) for details. The Intel® Xeon® Processor D-1500 Product Family LAN controller incorporates eight Ethertype filters.
- **Flow Director Filters** — These filters provide up to additional 32 K filters. See [Section 7.1.3.5](#) for details.
- **TCP SYN Filters** — The Intel® Xeon® Processor D-1500 Product Family LAN controller might route TCP packets with their SYN flag set into a separate queue. SYN packets are often used in SYN attacks to load the system with numerous requests for new connections. By filtering such packets to a separate queue, security software can monitor and act on SYN attacks. See [Section 7.1.3.4](#) for details.

A received packet is allocated to a queue based on the above criteria and the following order:

- Queue by L2 Ethertype filters (if match)
- Queue by SYN filter (if match)
- Queue by flow director filters
- Queue (in case of virtualization, within a pool) by RSS as described in [Section 7.1.3.1](#) and [Section 7.1.3.2](#).



- Send to queue zero.

7.1.3.1 Queuing in a Non-virtualized Environment

Table 7.1 lists the queuing schemes for packets that do not match any special filters (L2 Ethertype and SYN and flow director filters). Table 7.2 illustrates the queue indexing. Selecting a scheme is done via the *Multiple Receive Queues Enable (MRQE)* field in MRQC register.

Table 7.1. Rx Queuing Schemes Supported (No Virtualization)

RSS	RSS Queues	Special Filters (1)
No	1 queue Rx queue 0	Supported
Yes	64 RSS queues	Supported
No	8 TCs x 1 queue 4 TCs x 1 queue Assign Rx queue 0 of each TC	Supported
Yes	8 TCs x 16 RSS 4 TCs x 32 RSS	Supported

Table 7.2. Queue Indexing Illustration in Non-virtualization Mode

MRQC.MRQE	Queue Index bits	6	5	4	3	2	1	0
0000b	Default only	0						
0001b	RSS	RSS						
0101b	Reserved	TC		RSS ¹				
0011b	Reserved	TC		0				
0100b	Reserved	TC			RSS ¹			
0010b	Reserved	TC			0			

1. The number of bits used for each TC is set according to the *RQTC.RQTCx* fields

A received packet is assigned to a queue according to the ordering shown in Figure 7.5:

- RSS filters — Packets that do not meet any of the previous filtering conditions described in Section 7.1.3 are assigned to one of 128 queues as listed in Table 7.1. The following modes are supported:
 - No RSS — Queue 0 is used for all packets.
 - RSS only — A set of 64 queues is allocated for RSS. The queue is identified through the RSS index. Note that it is possible to use a subset of these queues.
- Example — Assume a 4 TCs x 32 RSS configuration and that the number of RSS queues for TC=3 is set to 4. The queue numbers for TC=3 are 64, 65, 66, and 67 (decimal).

Figure 7.5 shows the flow of allocation of Rx queues by the various queue filters previously described:

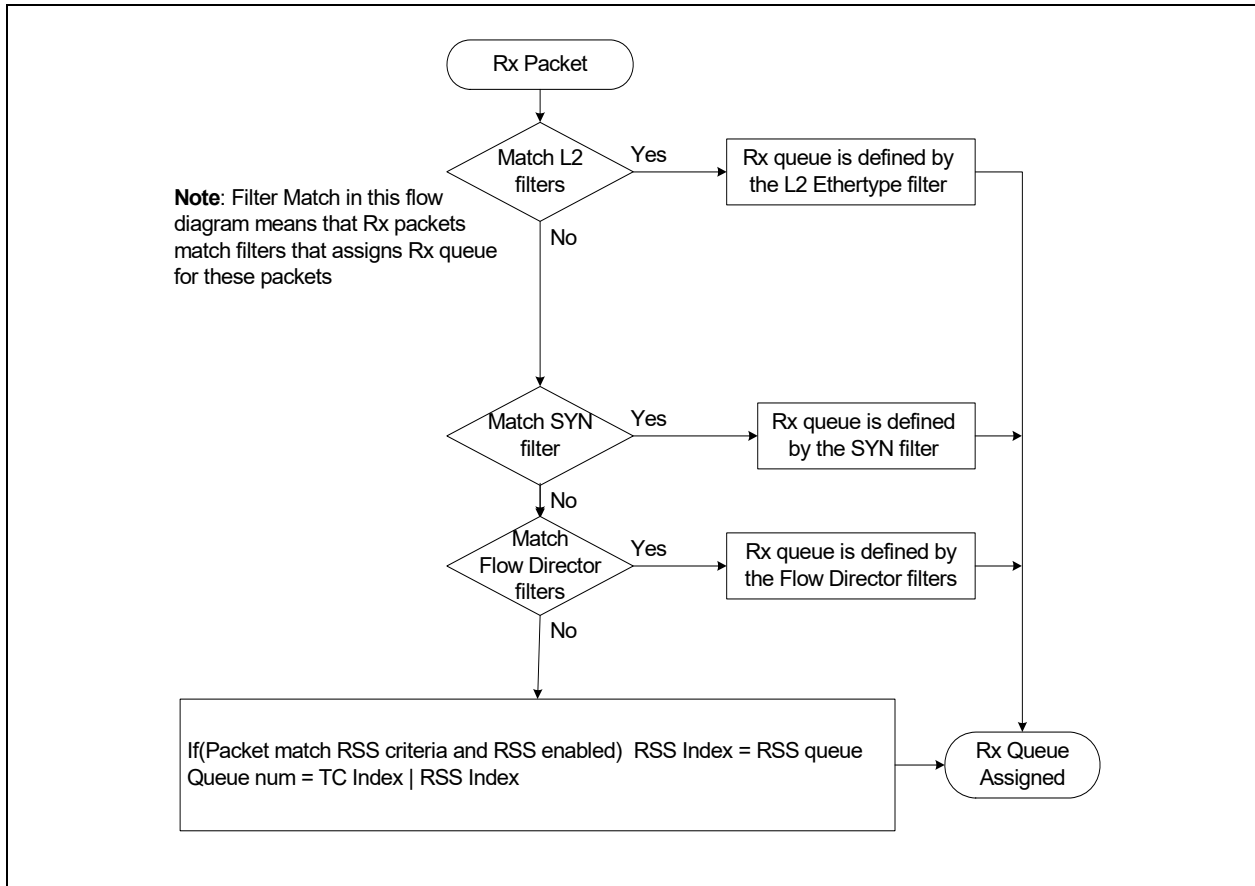


Figure 7.5. Rx Queuing Flow (Non-virtualized)

7.1.3.2 Queuing in a Virtualized Environment

The 128 Rx queues are allocated to a pre-configured number of queue sets, called pools. In non-IOV mode, system software allocates the pools to the VMM or to VMs. In IOV mode, each pool may be associated with a VF.

Incoming packets are associated with pools based on their L2 characteristics as described in Section 7.7.3. This section describes the following stage, where an Rx queue is assigned to each replication of the Rx packet as determined by its pools association.




Table 7.3 lists the queuing schemes supported with virtualization. Table 7.4 lists the queue indexing.

Table 7.3. Rx Queuing Schemes Supported with Virtualization

RSS	RSS Queues	Special Filters (1)
No	16 pools x 1 queue 32 pools x 1 queue 64 pools x 1 queue - Rx queue 0 of each pool	Supported
Yes	32 pools x 4 RSS 64 pools x 2 RSS	Supported
No	16 pools x 8 TCs 32 pools x 4 TCs	Supported
Yes	N/A	N/A

Table 7.4. Queue Indexing Illustration in Virtualization Mode

MRQC.MRQE	Queue Index bits	6	5	4	3	2	1	0	
1000b		Pool Index							0
1011b	Intel VT(64) + RSS	Pool Index							RSS
1001b/1111b	Intel VT(64) + RSS + Mega Pool	Pool Index					VF Index/RSS ¹	RSS	
1010b	Intel VT(32) + RSS	Pool Index					RSS		
N/A	Intel VT(16) + RSS	Not Supported							
1101b	Reserved	Pool Index					TC		
1100b	Reserved	Pool Index				TC			

1. When MRQC.MRQE = 0xF, VF index for pools 0-59, part of RSS for pool 60, 62 and pools 61 and 63 are disabled. When MRQC.MRQE = 0x0, VF index for pools 0-61, part of RSS for pool 62 and pool 63 is disabled.

Selecting a scheme is done in the following manner:

- Non-IOV mode
 - Select one of the above schemes via the *Multiple Receive Queues Enable* field in the MRQC register.
- IOV mode
 - Determine the number of pools: the number must support the value configured by the operating system in the PCIe NumVFs field. Therefore, the number of pools is min. of {16, 32, 64} that is still >= NumVFs.

A received packet is assigned to an absolute queue index according to the ordering shown in Figure 7.6). It is software responsibility to define a queue that belongs to the matched pool:

- RSS filters — The supported modes are listed in Table 7.3 and detailed as follows. The associated queue indexes are listed in Table 7.4.
 - No RSS — A single queue is used as default queue of the pool with either 16, 32, or 64 pools enabled. In a 64 pools setting, queues '2xN'...'2xN+1' are allocated to pool 'N'; In a 32 pools setting, queues '4xN'...'4xN+3' are allocated to pool 'N'. The queues not used as default may be directed to by other filters.
 - RSS only — All 128 queues are allocated to pools. Two configurations supported:
 - 32 pools with 4 RSS queues each

- 64 pools with 2 RSS queues each.
- 63 pools. The first 62 with 2 queues each and the last one with 4 RSS queues. In this mode the last pool is 62 and pool 63 is disabled.
- 62 pools. The first 60 with 2 queues each and the last two with 4 RSS queues each. In this mode the last two pools are 60 and 62 and pools 61/63 are disabled.

Notes: It is possible to use a subset of the RSS queues in each pool. The LS bits of the queue indexes are defined by the RSS index, and the pool index is used as the MS bits. See description below.

In the 62 and 63 pools modes, the pools with 4 queues should be assigned to the PF and can not be assigned to a VF.

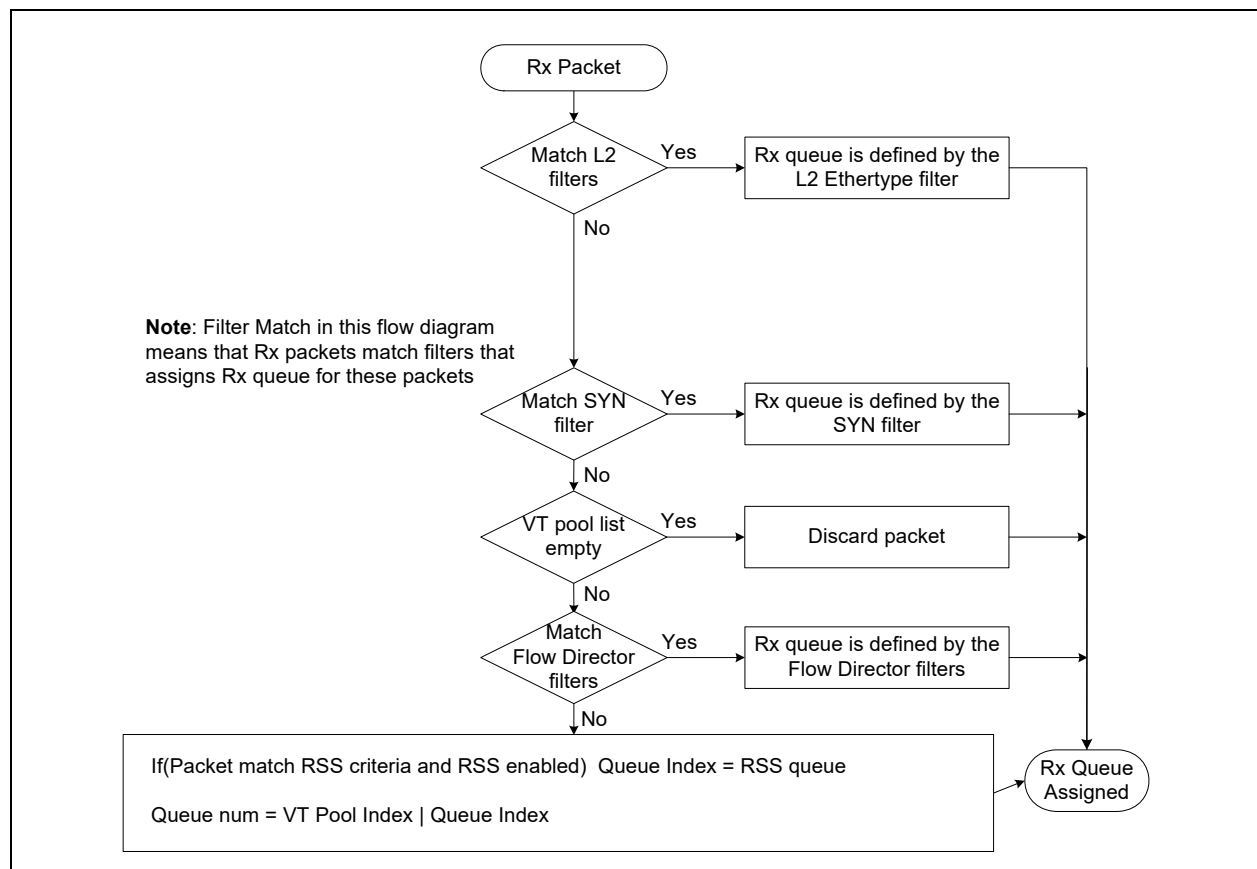


Figure 7.6. Rx Queuing Flow (Virtualization Case)

7.1.3.3 L2 Ethertype Filters

These filters identify packets by their L2 Ethertype, 802.1Q user priority and optionally assign them to a receive queue. The following possible usages have been identified at this time:

- IEEE 802.1X packets — Extensible Authentication Protocol (EAPOL) over LAN
- Time sync packets (such as IEEE 1588) — Identifies Sync or Delay_Req packets
- The L2 type filters should not be set to IP packet type as this might cause unexpected results



The Intel® Xeon® Processor D-1500 Product Family LAN controller incorporates eight Ethertype filters defined by a set of two registers per filter: *ETQF[n]* and *ETQS[n]*.

The L2 packet type is defined by comparing the *Ether-Type* field in the Rx packet with the *ETQF[n].EType* (regardless of the pool and UP matching). The *Packet Type* field in the Rx descriptor captures the filter number that matched with the L2 Ethertype. See [Section 7.1.5.2.2](#) for a description of the *Packet Type* field.

The following flow is used by the Ethertype filters:

1. If the *ETQF.Filter Enable* bit is cleared, the filter is disabled and the following steps are ignored.
2. Receive packet matches any *ETQF* filters if the *EtherType* field in the packet matches the *EType* field of the filter. Note that the following steps are ignored if the packet does not match the *ETQF* filters.

Note: Ethertype Filters should not be configured in a way that may cause a packet to match multiple filters.

3. Packets that match any *ETQF* filters are candidate for the host. If the packet also matches the manageability filters, it is directed to the host as well regardless of the *MNGONLY* register setting.
4. If the *ETQF.1588 Time Stamp* field is set, the packet is identified as an IEEE 1588 packet.
5. If the *ETQS.Queue Enable* bit is cleared, the filter completed its action on the packet. Else, the filter is also used for queuing purposes as described in the sections that follow.
6. If the *ETQF.Pool Enable* field is set, the *Pool* field of the filter determines the target pool for the packet. The packet can still be mirrored to other pools as described in [Section 7.7.3](#). See the sections that follow for more details on the use of the *Pool* field.
7. The *ETQS.RX Queue* field determines the destination queue for the packet. In case of a mirrored packet, only the copy of the packet that is targeted to the pool defined by the *Pool* field in the *ETQF* register is routed according to the *Rx Queue* field.

Setting the *ETQF[n]* registers is described as follows:

- The *Filter Enable* bit enables identification of Rx packets by Ethertype according to this filter. If this bit is cleared, the filter is ignored.
- The *EType* field contains the 16-bit Ethertype compared against all L2 type fields in the Rx packet.
- The *1588 Time Stamp* bit indicates that the Ethertype defined in the *EType* field is identified as IEEE 1588 EType. Packets that match this filter are time stamped according to the IEEE 1588 specification.
- The *Pool* field defines the target pool for a packet that matches the filter.
 - It applies only in virtualization modes. The pool index is meaningful only if the *Pool Enable* bit is set.
 - If the *Pool Enable* bit is set then the *Queue Enable* bit in the *ETQS* register must be set as well. In this case, the *Rx Queue* field in the *ETQS* must be part of the pool number defined in the *ETQF*.

Setting the *ETQS[n]* registers is described as follows:

- The *Queue Enable* bit enables routing of the Rx packet that match the filter to Rx queue as defined by the *Rx Queue* field.
- The *Rx Queue* field contains the destination queue (one of 128 queues) for the packet.

Special considerations for virtualization modes:

- Packets that match an Ethertype filter are diverted from their original pool (as defined by the VLAN and Ethernet MAC address filters) to the pool defined in the *Pool* field in the *ETQF* registers.
- The same applies for multicast packets. A single copy is posted to the pool defined by the filter.



- Mirroring rules
 - A packet sent to a pool by an *ETQF* filter, is still candidate to mirroring using the standard mirroring rules.
 - The Ethertype filter does not take part in the decision on the destination of the mirrored packet.

7.1.3.4 SYN Packet Filters

The Intel® Xeon® Processor D-1500 Product Family LAN controller might route TCP packets whose SYN flag is set into a separate queue. SYN packets are used in SYN attacks to load the system with numerous requests for new connections. By filtering such packets to a separate queue, security software can monitor and act on SYN attacks.

The following rules apply:

- A single SYN filter is provided.

The SYN filter is configured via the SYNQF register as follows:

- The *Queue Enable* bit enables SYN filtering capability.
- The *Rx Queue* field contains the destination queue for the packet (one of 128 queues). In case of mirroring (in virtualization mode), only the original copy of the packet is routed according to this filter.

7.1.3.5 Flow Director Filters

The flow director filters identify specific flows or sets of flows and routes them to specific queues. The flow director filters are programmed by FDIRCTRL and all other FDIR registers. The Intel® Xeon® Processor D-1500 Product Family LAN controller shares the Rx packet buffer for the storage of these filters.

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports three different modes of flow director (set in the *FDIRCTRL.FILTERMODE* field - [Section 8.2.2.11.1](#)):

- IP filtering (000b)
- MAC, VLAN filtering (001b)
- Cloud: NVGRE or VXLAN filtering (010b)

The following table describes the packets that are candidate to compared to the flow director filters according to the different modes:

Table 7-5. Candidate for Flow Director Per Mode

Mode	IP packets TCP/UDP/SCTP (not tunneled)	Non IP Packets	NVGRE Packets	VXLAN Packets	Other IP Packets or Fragmented IP
00 - IP (L4 fields unmasked)	Candidate	Not Candidate	Candidate (inner IP) ¹	Candidate (inner IP) ¹	Not Candidate
00 -IP (L4 fields masked)	Candidate	Not Candidate	Candidate (inner IP) ²	Candidate (inner IP) ²	Candidate ^{2,3}
01 - MAC VLAN	Candidate	Candidate	Candidate	Candidate	Candidate
10 - Cloud	Not Candidate	Not Candidate	Candidate	Candidate	Not Candidate

1. Only if a known L4 header is present (TCP/UDP/SCTP).
 2. VXLAN and NVGRE packets without inner IP header (like tunneled ARP) are not candidate in IP mode.
 3. IP in IP packets are not candidate in IP mode.



Basic rules for the flow director filters are:

In Intel® Virtualization Technology (Intel® VT) mode, the *Pool* field in FDIRCMD must be valid. If the packet is replicated, only the copy that goes to the pool that matches the *Pool* field is impacted by the filter. The flow director filters cover the following fields:

Table 7-6. Lookup fields for Flow Director Per Mode

Mode	Lookup Fields
00 - IP	<ul style="list-style-type: none"> • VLAN header • Source IP and destination IP addresses • Source port and destination port numbers • IPv4 / IPv6¹ and UDP / TCP or SCTP protocol match • Flexible 2-byte tuple anywhere in the first 64 bytes of the packet • Target pool number (relevant only for Intel VT mode)
01 - MAC VLAN	<ul style="list-style-type: none"> • MAC • VLAN • Flexible 2-byte tuple anywhere in the first 64 bytes of the packet • Target pool number (relevant only for Intel VT mode)
10 - Cloud	<ul style="list-style-type: none"> • Tunnel type - NVGRE or VXLAN • TNI (NVGRE) / VNI (VXLAN) • Inner MAC • Inner VLAN • Flexible 2-byte tuple anywhere in the first 64 bytes of the packet • Target pool number (relevant only for Intel VT mode)

1. IPv6 extended headers are parsed by the Intel® Xeon® Processor D-1500 Product Family LAN controller, enabling TCP layer header recognition. Still the IPv6 extended header fields are not taken into account for the queue classification by Flow Director filter. This rule do not apply for security headers and fragmentation header. Packets with fragmentation header miss this filter. Packets with security extended headers are parsed only up to these headers and therefore can match only filters that do not require fields from the L4 protocol.

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports two types of filtering modes (static setting by the FDIRCTRL.Perfect-Match bit):


- Perfect match filters — The hardware checks a match between the masked fields of the received packets and the programmed filters. Masked fields should be programmed as zeros in the filter context. The Intel® Xeon® Processor D-1500 Product Family LAN controller supports up to 8 K - 1 perfect match filters.
- Signature filters — The hardware checks a match between a hash-based signature of the masked fields of the received packet. The Intel® Xeon® Processor D-1500 Product Family LAN controller supports up to 32 K - 1 signature filters. This mode can be used only when the filtering mode is IPMODE (*FDIRCTRL.FILTERMODE* = 000b).
- Denote — The *Perfect Match* fields and *Signature* field are denoted as *Flow ID* fields.

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports masking / range for the previously described fields. These masks are defined globally for all filters in the FDIR...M registers.

- The following fields can be masked per bit enabling power of two ranges up to complete enable / disable of the fields: IPv4 addresses and L4 port numbers.
- The following fields can be masked per byte enabling lower granularity ranges up to complete enable / disable of the fields: IPv6 addresses; Inner MAC; Outer MAC; TNI/VNI field.

Note: In perfect match filters the destination IPv6 address can only be compared as a whole (with no range support) to the *IP6AT* filters. A match to any of the *IP6AT* filter is considered as an IPv6 destination match.

- The following fields can be either enabled or disabled completely for the match functionality: VLAN ID tag; Outer VLAN Priority + DEI bit (the user priority is taken from the outermost tag with UP



bits); Inner or outer VLAN tag; Flexible 2-byte tuple and target pool. Target pool can be enabled by software only when Intel VT is enabled as well.

Flow director filters have the following functionality in virtualization mode:

- Flow director filters are programmed by the registers in the PF described in [Section 7.1.3.5.13](#) and [Section 7.1.3.5.14](#).
- Flow director filters can be utilized in virtualization mode to filter on MAC-VLAN by setting `FDIRCTRL.FILTERMODE = MACVLANMODE` (001b). `MACVLANMODE` is only used in flow director perfect match mode.

7.1.3.5.1 Flow Director Filters Actions

Flow director filters might have one of the following actions programmed per filter in the `FDIRCMD` register:

- Drop packet or pass to host as defined by the `Drop` bit.
 - Matched packets to a flow director filter is directed to the assigned Rx queue only if the packet does not match the L2 filters for queue assignment nor the SYN filter for queue assignment.
 - Packets that match a filter are directed to the Rx queue defined in the filter context as programmed by the `FDIRCMD.Rx-Queue`. The Rx-Queue field is an absolute receive queue index. In a non-Intel VT setting, it can be programmed to any value. In Intel VT mode, the software should set the Rx-Queue to an index that belongs to the matched pool.
 - Packets that match drop filters are directed to the Rx queue defined per all filters in the `FDIRCTRL.DROP-Queue`. The Intel® Xeon® Processor D-1500 Product Family LAN controller drops these packets if software does not enable the specific Rx queue.

7.1.3.5.2 Flow Director Default Action

A default drop action may be applied by the flow director for packets that are candidate to flow director as defined in [Table 7-5](#). If the `FDIRCTRL.DROP_NO_MATCH` bit is set, any packet candidate for flow director that does not match any of the filters will be sent to the queue defined in the `FDIRCTRL.DROP_QUEUE` field.

7.1.3.5.3 Flow Director Filters Status Reporting

Shared status indications for all packets:

- The Intel® Xeon® Processor D-1500 Product Family LAN controller increments the `FDIRMATCH` counter for packets that match a flow director filter. It also increments the `FDIRMISS` counter for packets that do not match any flow director filter.
- The *Flow Director Filter Match (FLM)* bit in the *Extended Status* field of the Rx descriptor is set for packets that match a flow director filter.
- The flow ID parameters are reported in the *Flow Director Filter ID* field in the Rx descriptor if enabled by the `FDIRCTRL.Report-Status`. When the `Report-Status` bit is set, the `RXCSUM.PCSD` bit should be set as well. This field is indicated for all packets that match or do not match the flow director filters.
 - For packets that do not match a flow director filter, if the `FDIRCTRL.REPORT_STATUS_ALWAYS` is set, the *Flow Director Filter ID* field can be used by software for future programming of a



matched filter, otherwise, the RSS hash value is reported. Table 7-7 describes the value of the RSS_TYPE field in the receive descriptor in the different cases.

Table 7-7. RSS Type Values According to Flow Director Match

REPORT_STATUS (FDIRCTRL[5])	REPORT_STATUS_ALWAYS (FDIRCTRL[7])	Packet matches a Flow Director Entry	RSS Type field value
0	0	X	RSS Type
1	0	0	RSS Type
1	0	1	0xF
0	1	X	Illegal configuration
1	1	X	0xF

For packets that match a flow director filter, the *Flow Director Filter ID* field can be used by software to identify the flow of the Rx packet. Too long linked list exception (linked list and too long terms are illustrated in Figure 7.7):

- The maximum recommended linked list length is programmed in the *FDIRCTRL.Max-Length* field
- The length exception is reported in the *FDIRErr* field in the Rx descriptor
- Packets that do not match any flow director filter, reports this exception if the length of the existing linked list is above the maximum recommended length. Software can use it to avoid further programming of additional filters to this linked list before other filters are removed.
- Packets that match a pass filter report this exception if the distance of the matched filter from the beginning of the linked list is higher than the above recommended length.
- Packets that match a drop filter are posted to the Rx queue programmed in the filter context instead of the global *FDIRCTRL.DROP_QUEUE*. The drop exception is reported in addition to the length exception (in the same field in the Rx descriptor).

Collision exception:

- Packets that matches a collided filter report this exception in the *FDIRErr* field in the Rx descriptor.
- Collision events for signature-based filters should be rare. Still it might happen because multiple flows can have the same hash and signature values. Software might leave the setting as is while the collided flows are handled according to the actions of the first programmed flow. Only one flow (out of the collided ones) might remain in the flow director filters. In order to clear the collision indication in the programmed filter, software should remove the filter and then re-program it once again.
- Collision events for a perfect match filter should never happen. A collision error might indicate a programming fault that software might decide to fix.

7.1.3.5.4 Flow Director Filters Block Diagram

The following figure shows a block diagram of the flow director filters. Received flows are identified to buckets by a hash function on the relevant tuples as defined by the FDIR...M registers. Each bucket is organized in a linked list indicated by the hash lookup table. Buckets can have a variable length while the last filter in each bucket is indicated as a last. There is no upper limit for a linked list length during programming; however, a received packet that matches a filter that exceeds the *FDIRCTRL.Max-Length* are reported to software (see Section 7.1.3.5.6).

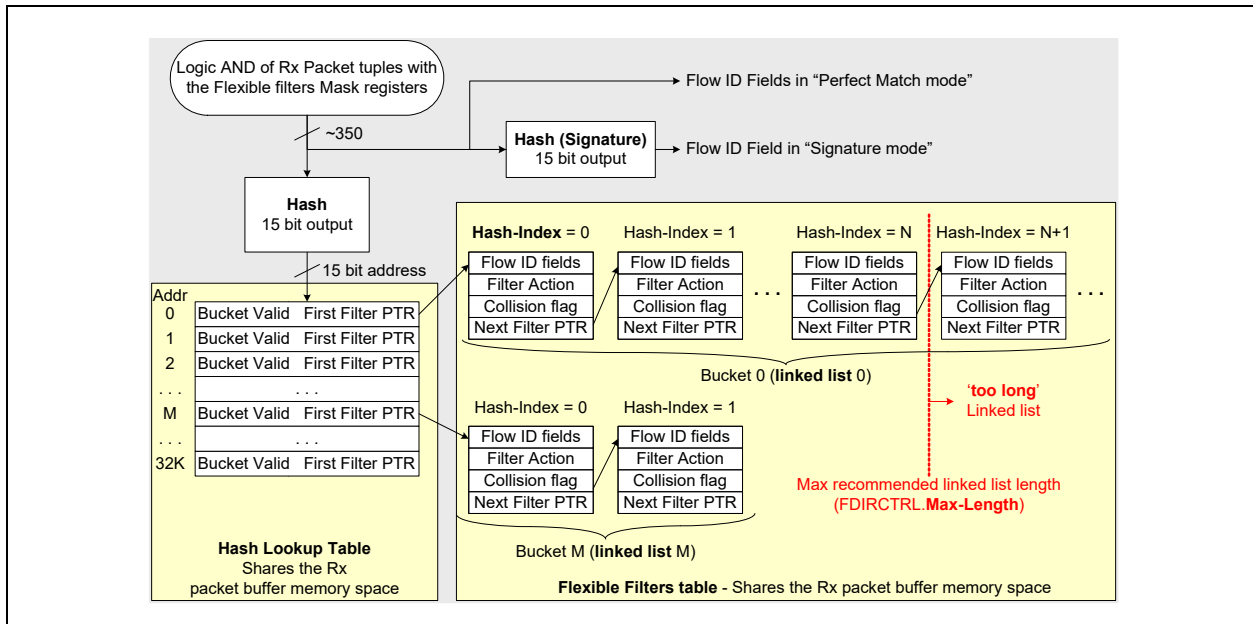


Figure 7.7. Flow Director Filters Block Diagram

7.1.3.5.5 Rx Packet Buffer Allocation

Flow director filters can consume zero space (when disabled) up to ~256 KB of memory. As shown in Figure 7.7, flow director filters share the same memory with the Rx packet buffer. By setting the *PBALLOC* field in the *FDIRCTRL* register, the software can enable and allocate memory for the flow director filters. The memory allocated to received traffic is the remaining part of the Rx packet buffer.

Table 7.8. Rx Packet Buffer Allocation

PBALLOC (2)	Effective Rx Packet Buffer Size (see following note)	Flow Director Filters Memory	Supported Flow Director Filters			
			Signature		Perfect Match	
			Filters	Bucket Hash	Filters	Bucket Hash
00 Flow Director is disabled	384 KB	0	0	N/A	0	N/A
01	320 KB	64 KB	8 K -1 filters	13 bit	2 K -1 filters	11 bit
10	256 KB	128 KB	16 K -1 filters	14 bit	4 K -1 filters	12 bit
11	128 KB	256 KB	32 K -1 filters	15 bit	8 K -1 filters	13 bit

Note: It is the user’s responsibility to ensure that sufficient buffer space is left for reception of traffic. The required buffer space for receive traffic depends on the number of traffic classes, and flow control upper and lower threshold values. If flow director is enabled (*PBALLOC* > 0), software should set the *RXPBSIZE[n]* registers according to the total remaining part of the Rx packet buffer for reception of traffic.

7.1.3.5.6 Flow Director Filtering Reception Flow



- Rx packet is digested by the filter unit which parse the packet extracting the relevant tuples for the filtering functionality.
- The Intel® Xeon® Processor D-1500 Product Family LAN controller calculates a 15-bit hash value out of the masked tuples (logic mask of the tuples and the relevant mask registers) using the hash function described in [Section 7.1.3.5.16](#).
- The address in the hash lookup table points to the selected linked list of flow director filters.
- The Intel® Xeon® Processor D-1500 Product Family LAN controller checks the *Bucket Valid* flag. If it is inactive, then the packet does not match any filter. Otherwise, *Bucket Valid* flag is active, proceed for the next steps.
- The Intel® Xeon® Processor D-1500 Product Family LAN controller checks the linked list until it reaches the last filter in the linked list or until a matched filter is found.
- Case 1: matched filter is found:
 - Increment the *FDIRMATCH* statistic counter.
 - Process the filter's actions (queue assignment) according to queue assignment priority. Meaning, the actions defined in this filter takes place only if the packet did not match any L2 filter or SYN filter that assigns an Rx queue to the packet.
 - Rx queue assignment according to the filter context takes place if *Queue-EN* is set.
 - If the DROP bit is set in the filter context, the packet is sent to the queue pointed by the *FDIRCTRL.DROP_QUEUE* field.
 - Post the packet to host including the flow director filter match indications as described in [Section 7.1.3.5.3](#).
- Case 2: matched filter is not found:
 - Increment the *FDIRMISS* statistic counter.
 - Post the packet to host including the flow director filter miss indications as described in [Section 7.1.3.5.3](#).


7.1.3.5.7 Add Filter Flow

The software programs the filters parameters in the registers described in [Section 7.1.3.5.13](#) and [Section 7.1.3.5.14](#) while keeping the *FDIRCMD.Filter-Update* bit inactive. As a result, the Intel® Xeon® Processor D-1500 Product Family LAN controller checks the bucket valid indication in the hash lookup table (that matches the *FDIRHASH.Hash*) for the presence of an existing linked list. Following are the two programming flows for the case a link list exists or for the case a new list is required.

- Case 1: Add a filter to existing linked list:

The Intel® Xeon® Processor D-1500 Product Family LAN controller checks the linked list until it reaches the last filter in the list or until a matched filter is found. The following cases may occur:

 - Matched filter is found (equal flow ID) with the same action parameters — The programming is discarded silently. This is a successful case since the programmed flow is treated as requested.
 - Matched filter is found (equal flow ID) with different action parameters — The Intel® Xeon® Processor D-1500 Product Family LAN controller keeps the old setting of the filter while setting the *Collision* flag in the filter context (see [Section 7.1.3.5.3](#) for software handling of collision during packet reception).
 - Matched filter is found (equal flow ID) with different action parameters and the *Collision* flag is already set — The programming is discarded silently. Software gets the same indications as the previous case.
 - Matched filter is not found (no collision) — The Intel® Xeon® Processor D-1500 Product Family LAN controller checks for a free space in the flow director filters table.



- No space case — Discard programming; increment the *FADD* counter in the *FDIRFSTAT* register and assert the flow director interrupt. Following this interrupt software should read the *FDIRFSTAT* register and *FDIRFREE.FREE* field, for checking the interrupt cause.
- Free space is found — Good programming case: Add the new filter at the end of the linked list while indicating it as the last one. Program the *Next Filter PTR* field and then clear the *Last* flag in the filter that was previously the last one.
- Case 2 — Create a new linked list:
The Intel® Xeon® Processor D-1500 Product Family LAN controller looks for an empty space in the flow director filters table:
 - Handle no empty space the same as in Case 1.
 - Good programming case: Add the new filter while indicating it as the last one in the linked list. Then, program the hash lookup table entry by setting the *Valid* flag and the *First Filter PTR* pointing to the new programmed filter.

Additional successful add flow indications:

- Increment the *ADD* statistic counter in the *FDIRUSTAT* register.
- Reduce the *FREE* counter in the *FDIRFREE* register and then indicate the number of free filters. If the *FREE* counter crosses the full-thresh value in the *FDIRCTRL* register, then assert the flow director filter interrupt. Following this interrupt software should read the *FDIRFSTAT* register and *FDIRFREE.FREE* field, for checking the interrupt cause.
- Compare the length of the new linked list with *MAXLEN* in the *FDIRLEN* register. If the new linked list is longer than *MAXLEN*, update the *FDIRLEN* by the new flow.

7.1.3.5.8 Update Filter Flow

In some applications, it is useful to update the filter parameters, such as the destination Rx queue. Programming filter parameters is described in [Section 7.1.3.5.7](#).

Setting the *Filter-Update* bit in the *FDIRCMD* register has the following result:


- Case 1: Matched filter does not exist in the filter table — Setting the *Filter-Update* bit has no impact and the command is treated as add filter.
- Case 2: Matched filter already exists in the filter table — Setting the *Filter-Update* bit enables filter parameter's update while keeping the collision indication as is.

When updating an existing filter the software device driver should program the same filter (i.e. the same *FDIRHASH.HASH* and the same *FDIRHASH.SIGNATURE_SW_INDEX*) while keeping the *FDIRCMD.FILTER_UPDATE* = 1b.

7.1.3.5.9 Remove Filter Flow

Software programs the filter Hash and Signature / Software-Index in the *FDIRHASH* register. It then should set the *FDIRCMD.CMD* field to *Remove Flow*. Software might use a single 64-bit access to the two registers for atomic operation. As a result, the Intel® Xeon® Processor D-1500 Product Family LAN controller follows these steps:

- Check if such a filter exists in the flow director filters table.
- If there is no flow, then increment the *FREMOVE* counter in the *FDIRFSTAT* register and skip the next steps.
- If the requested filter is the only filter in the linked list, then invalidate its entry in the hash lookup table by clearing the *Valid* bit.
- Else, if the requested filter is the last filter in the linked list, then invalidate the entry by setting the *Last* flag in the previous filter in the linked list.



- Else, invalidate its entry by programming the Next Filter PTR in the previous filter in the linked list, pointing it to the filter that was linked to the removed filter.

Additional indications for successful filter removal:

- Increment the remove statistic counter in the FDIRUSTAT register.
- Increment the *FREE* counter in the FDIRFREE register.

7.1.3.5.10 Remove all Flow Director Filters

In some cases there is a need to clear the entire flow director table. It might be useful in some applications that might cause the flow director table becoming too occupied. Then, software might clear the entire table enabling its re-programming with new active flows.

Following are steps required to clear the flow director table:

- Poll the *FDIRCMD.CMD* until it is zero indicating any previous pending commands to the flow director table is completed (at worst case the *FDIRCMD.CMD* should be found cleared on the second read cycle). Note that software must not initiate any additional commands (add / remove / query) before this step starts and until this flow completes.
- Clear the FDIRFREE register (set *FREE* field to zero).
- Set *FDIRCMD.CLEARHT* to 1b and then clear it back to 0b
- Clear the FDIRHASH register to zero
- Re-write FDIRCTRL by its previous value while clearing the *INIT-Done* flag.
- Poll the *INIT-Done* flag until it is set to 1b by hardware.
- Clear the following statistic registers: FDIRUSTAT; FDIRFSTAT; FDIRMATCH; FDIRMISS; FDIRLEN (note that some of these registers are read clear and some are read write).

7.1.3.5.11 Flow Director Filters Initializing Flow

Following a device reset, the flow director is enabled by programming the FDIRCTRL register, as follows:

- Set *PBALLOC* to non-zero value according to the required buffer allocation to reception and flow director filter (see [Section 7.1.3.5.5](#)). All other fields in the register should be valid as well (according to required setting) as the FDIRCTRL register is expected to be programmed by a single cycle.
- Poll the *INIT-Done* flag until it is set to 1b by hardware (expected initialization flow should take about 55 μ s at 10 Gb/s and 550 μ s at 1 Gb/s).

7.1.3.5.12 Query Filter Flow

Software might query specific filter settings and bucket length using the Query command.

- Program the filter Hash and Signature/Software-Index in the FDIRHASH register and set the *CMD* field in the FDIRCMD register to 11b (Query Command). A single 64-bit access can be used for this step.
- As a result, the Intel® Xeon® Processor D-1500 Product Family LAN controller provides the query result in the FDIRHASH, FDIRCMD and FDIRLEN registers (described in the sections that follow).
- Hardware indicates query completion by clearing the *FDIRCMD.CMD* field. The following table lists the query result.



Query Outcome	FDIRHASH -> Bucket Valid	FDIRCMD -> Filter Valid	FDIRLEN -> Bucket Length	FDIRCMD -> Filter ID Fields	FDIRCMD -> Filter Action
Empty Bucket	0	0	N/A	N/A	N/A
Valid Bucket, Matched Filter Not Found	1	0	Bucket linked list length	N/A	N/A
Found Signature Filter	1	1	Filter index within the linked list	0	Filter's parameters ¹
Found Perfect Match Filter	1	1	Filter index within the linked list	Filter's parameters	Filter's parameters

1. The pool parameter is not returned as part of the filter action in signature mode.

7.1.3.5.13 Signature Filter Registers

The signature flow director filter is programmed by setting the FDIRHASH and FDIRCMD registers. These registers are located in consecutive 8-byte aligned addresses. Software should use a 64-bit register to set these two registers in a single atomic operation. Table 7.9 lists the recommended setting.

Table 7.9. Signature Match Filter Parameters

Filter Bucket Parameters — FDIRHASH	
Hash	15-bit hash function used to define a bucket of filters. This parameter is part of the flow director filter ID that can be reported in the Rx descriptor. It is shared for signature and perfect match filters.
Valid	Should be set to 1b. It is shared for signature and perfect match filters.
Flow ID — FDIRHASH	
Signature	15-bit hash function used as the flow matching field. This parameter is also part of the flow director filter ID that can be reported in the Rx descriptor.
FDIRCMD — Programming Command and Filter action	

7.1.3.5.14 Perfect Match Filter Registers

Perfect match filters are programmed by the following registers: FDIRSIPv6[n]; FDIRVLAN; FDIRPORT; FDIRIPDA; FDIRIPSA; FDIRHASH; FDIRCMD. Setting the FDIRCMD register, generates the actual programming of the filter. Therefore, write access to this register must be the last cycle after all other registers contain a valid content. Table 7.10 lists the recommended setting.

Note: Software filter programming must be an atomic operation. In a multi-core environment, software must ensure that all registers are programmed in a sequence with no possible interference by other cores.

Table 7.10. Perfect Match Filter Parameters

Filter Bucket Parameters and Software Index — FDIRHASH	
Hash	See Section 7.1.3.5.13.
Valid	See Section 7.1.3.5.13.
Software-Index	16-bit index provided by software at filter programming used by software to identify the matched flow. This parameter is also part of the flow director filter ID that can be reported in the Rx descriptor.
FDIRCMD — Programming Command and Filter Action Set	



Table 7.10. Perfect Match Filter Parameters

	IP Mode FILTERMODE = 000b	MAC, VLAN Mode FILTERMODE = 001b	Cloud Modes FILTERMODE = 010b
IPV6DMATCH	Should be set for IPv6 filters only and only if the destination address should be compared to one of the IPAT filters.	Should be cleared	
L4TYPE	Defines if the filter is for TCP, UDP or SCTP flows.	Should be cleared FDIRM.L4P should be set	
IPV6	Defines if the filter is for IPv4 or IPv6 addresses.	Should be cleared FDIRM.L3P should be set FDIRM.DIPV6 should be set	
TUNNEL_FILTER	Defines if the filter is a VXLAN and NVGRE tunneled packet filter. If set, the parameters of the L3 and L4 header relates to the tunneled (inner) header.	Should be cleared.	
POOL	Pool to which the packet is directed (relevant only in virtualization modes).		
Flow ID — Perfect Match Flow ID Parameters are Listed in the Following Registers and Fields			
	IP Mode FILTERMODE = 000b	MAC, VLAN Mode FILTERMODE = 001b	Cloud Mode FILTERMODE = 010b
FDIRSIPv6[0...2].IP6SA	Three LS DWord of the source IPv6. Meaningful for IPv6 flows depending on the FDIRIP6M.SIPM setting.	{FDIRSIPv6_0[31:0], FDIRSIPv6_1[15:0]} holds the destination MAC address value FDIRSIPv6_1[31:16], and FDIRSIPv6_2[31:0] are reserved and should be set to 0x0. The FDIRIP6M.SIPM field should be set to 0x0xFC0F	FDIRSIPv6_2[31:0] holds the 32 bits of the TNI/VNI field. {FDIRSIPv6_0[31:0], FDIRSIPv6_1[15:0]} holds the destination MAC address value FDIRSIPv6_1[31] holds the cloud mode (0 = NVGRE, 1 = VXLAN). FDIRSIPv6_1[30:16] are reserved and should be set to 0x0. In order to mask one of the fields the matching bits in FDIRIP6M.SIPM should be set: <ul style="list-style-type: none"> • Bits 15:12 control the masking of the TNI/VNI field. All the 32 bits of the key can be used for filtering. • Bit 11 controls the masking of the cloud mode and should be cleared for normal operation • Bits 3:0 and bit 11 should be set. • Bits 9:4 control the masking of the inner MAC address.
FDIRVLAN.VLAN	VLAN fields are meaningful depending on the FDIRM.VLANID and FDIRM.VLANP setting. The VLAN field should be stored in Little endian format.		Inner VLAN fields are meaningful depending on the FDIRM.VLANID and FDIRM.VLANP setting. The VLAN field should be stored in Little endian format.
FDIRVLAN.FLEX	Flexible 2-byte field at offset FDIRCTRL.Flex-Offset. Meaningful depending on FDIRM.FLEX setting.		
FDIRPORT.Source	L4 source port. Meaningful for TCP, UDP and SCTP packets depending on the FDIRTCPM.SportM, FDIRUDPM.SportM and FDIRSCTPM.SportM setting.	Reserved - Should be zero. FDIRM.L4P should be set.	



Table 7.10. Perfect Match Filter Parameters

FDIRPORT.Destination	L4 destination port. Meaningful for TCP, UDP and SCTP packets depending on the FDIRTCPM.DportM, FDIRUDPM.DportM and FDIRSCTPM.DPORTM setting.	Reserved - Should be zero. FDIRM.L4P should be set.
FDIRIPDA.IP4DA	IPv4 destination address. Meaningful depending on the FDIRDIP4M.IPM setting.	Reserved - Should be zero. All bits in FDIRDIP4M.IPM should be set.
FDIRIPSA.IP4SA	IPv4 source address or LS DWord of the source IPv6 address. Meaningful for IPv4 flows depending on the FDIRSIP4M.IPM setting and for IPv6 flows depending on the FDIRIP6M.SIPM setting.	Reserved - Should be zero. All bits in FDIRDSIP4M.IPM should be set. FDIRIP6M.SIPM[3:0] should be set.

7.1.3.5.15 Multiple CPU Cores Considerations

Perfect match filters programming and any query cycles require access to multiple registers. In order to avoid races between multiple cores, software might need to use one of the following programming methods:

- Use a software-based semaphore between the multiple cores for gaining control over the relevant CSR registers for complete programming or query cycles.
- Manage all programming and queries of the flow director filters by a single core.

Programming signature filters requires only the FDIRHASH and FDIRCMD registers. These two registers are located in 8-byte aligned adjacent addresses. Software could use an 8-byte register for the programming of these registers in a single atomic operation, which avoids the need for any semaphore between multiple cores.

7.1.3.5.16 Flow Director Hash Function

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports programmable 16-bit hash functions based on two 32-bit keys, one for the lookup table identifying a bucket of filters and another one for the signature (*FDIRHKEY* and *FDIRSKEY*). The hash function is described in the sections that follows. In some cases, a smaller hash value than 16 bits is required. In such cases, the LS bits of the hash value are used.

```
For (i=0 to 350) {if (Ext_K[i]) then Hash[15: 0] = Hash[15: 0] XOR Ext_S[15+i: i]}
```

While using the following notations:

'XOR' - Bitwise XOR of two equal length strings

If (xxx) - Equals 'true' if xxx = '1' and equals 'false' if xxx = '0'

S[335:0] - The input bit string of the flow director tuples: 42 bytes listed in [Table 7.11](#) AND-logic with the filters masks.

```
Ext_S[n] - S[14:0] | S[335:0] | S[335:321] // concatenated
```

K[31:0] - The hash key as defined by the FDIRHKEY or FDIRSKEY registers.

Tmp_K[11*32-1:0] - (Temp Key) equals K[31:0] | K[31:0]... // concatenated Key 11 times

Ext_K[350:0] - (Extended Key) equals T_K[351:1]



The input bit stream for the hash calculation is listed in [Table 7.11](#) while byte 0 is the MSB (first on the wire) of the VLAN, byte 2 is the MSB of the source IP (IPv6 case) and so on.

Table 7.11. Input Bit Stream for Hash Calculation

Bytes/Mode	Field		
	IP	MAC, VLAN	Cloud Mode
Bytes 0...1	VLAN tag (always the outer VLAN)		Inner VLAN Tag
Bytes 2...5	Zero		
Bytes 6...11	Source IP (16 bytes for IPv6; source IP for IPv4 12 bytes of zero's)	MAC	Inner MAC
Bytes 12...13		Zero	8 bits of zeros bit of tunnel_type 7 bits of zeros
Bytes 14...17		Zero	VNI/TNI
Bytes 18...33	Signature Mode: Destination IP (16 bytes for IPv6; destination IP for IPv4 12 bytes of zero's) Exact Mode: Destination IP (7 zero bits IP6AT match indication 120 zero bits; destination IP for IPv4 12 bytes of zero's) ¹		
34...37	L4 source port number L4 destination port number		
38...39	Flexible bytes		
40	00b pool number (as defined by FDIRCMD.Pool)		
41	[7:5] 000b [4] FDIRCMD.TUNNEL_FILTER (should be zero in MAC, VLAN mode). [3] 0b [2] IPv6/IPv4 type (FDIRCMD.IPV6) [1:0] L4 type (FDIRCMD.L4TYPE)		

1. In VXLAN and NVGRE packets, the IP addresses used are of the tunneled header in other packets the outer IP header is used.

7.1.3.6 RSS

RSS is a mechanism to distribute received packets into several descriptor queues. Software can then assign each queue to a different processor, therefore sharing the load of packet processing among several processors.

As described in [Section 7.1](#), the Intel® Xeon® Processor D-1500 Product Family LAN controller uses RSS as one ingredient in its packet assignment policy (the others are the various filters and virtualization). The RSS output is an RSS index. The Intel® Xeon® Processor D-1500 Product Family LAN controller global assignment uses these bits (or only some of the LSBs) as part of the queue number.

[Figure 7.8](#) and [Figure 7.9](#) show the process of computing an RSS output:

1. The receive packet is parsed into the header fields used by the hash operation (such as IP addresses, TCP port, etc.)
2. A hash calculation is performed. The Intel® Xeon® Processor D-1500 Product Family LAN controller supports a single hash function, as defined by Microsoft* (MSFT) RSS. The Intel® Xeon® Processor D-1500 Product Family LAN controller therefore does not indicate to the software device driver which hash function is used. The 32-bit result is fed into the *RSS Hash* field in the packet receive descriptor.
3. The Intel® Xeon® Processor D-1500 Product Family LAN controller supports two modes of RSS defined by the *MRQC.Multiple_RSS* bit. When set to 0b a single RSS key and redirection table are supported. In this mode (usually used when virtualization is not enabled) the nine LSBs of the hash result are used as an index into a 512-entry redirection table. Each entry provides up to 6-bit RSS output index ([Figure 7.8](#)).
For SR-IOV enabled mode (set by *MRQC.Multiple_RSS* set to 1b), The Intel® Xeon® Processor D-

1500 Product Family LAN controller supports up to 64 (one per pool) RSS keys and redirection tables (both controlled and programmed at the VF space). In this mode, the 6 LSBs of the hash result are used as an index into a 64-entry redirection table. Each entry provides up to 2-bit RSS output index (Figure 7.9).

When RSS is enabled, the Intel® Xeon® Processor D-1500 Product Family LAN controller provides software with the following information as required by Microsoft RSS and provided for device driver assist:

- A Dword result of the MSFT RSS hash function application to the packet header, to be used by the stack for flow classification, is written into the receive packet descriptor. A 4-bit *RSS Type* field conveys the hash function used for the specific packet.
- Packets to which the RSS function can not be applied (for example non IP packets) will return an *RSS Hash* result of zero, an *RSS Type* of zero and an RSS output index of zero.

7.1.3.6.1 Enabling RSS

- RSS is enabled in the MRQC register.
- RSS enabling cannot be done dynamically and must be preceded by a software reset.
- RSS status field in the descriptor write-back is enabled when the *RXCSUM.PCSD* bit is set (fragment checksum is disabled). RSS is therefore mutually exclusive with UDP fragmentation checksum offload.
- Support for RSS is not provided when legacy receive descriptor format is used.

7.1.3.6.2 Disabling RSS

- Disabling RSS on the fly is not allowed, and the Intel® Xeon® Processor D-1500 Product Family LAN controller must be reset after RSS is disabled.
- When RSS is disabled, packets are assigned an RSS output index = zero.

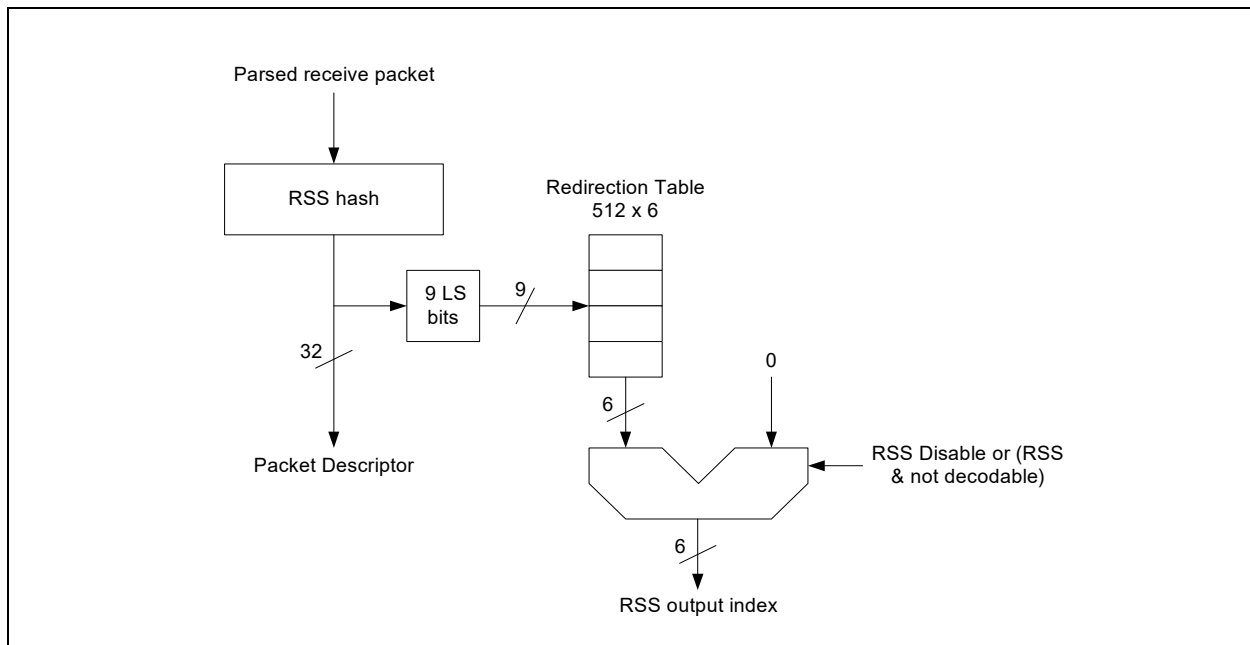


Figure 7.8. RSS Block Diagram (Multiple_RSS = 0)

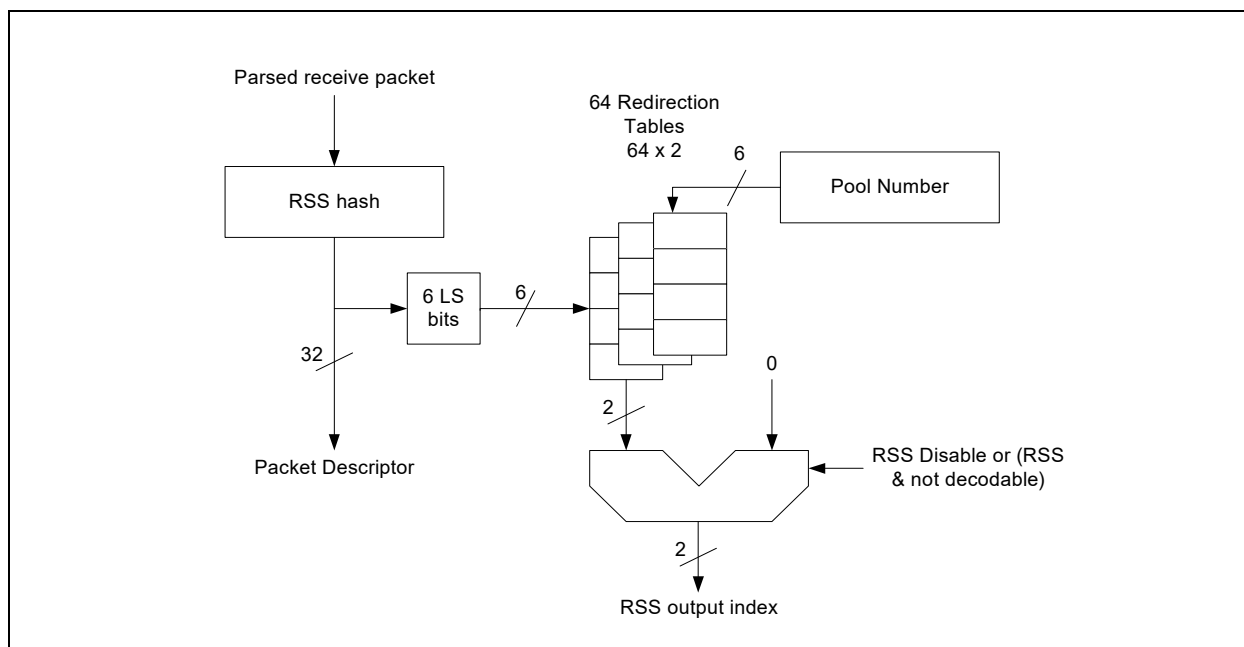


Figure 7.9. RSS Block Diagram (Multiple_RSS = 1)

7.1.3.6.3 RSS Hash Function

This section provides a verification suite used to validate that the hash function is computed according to MSFT nomenclature.

The Intel® Xeon® Processor D-1500 Product Family LAN controller’s hash function follows the MSFT definition. A single hash function is defined with several variations for the following cases:


- **TcpIPv4** — The Intel® Xeon® Processor D-1500 Product Family LAN controller parses the packet to identify an IPv4 packet containing a TCP segment per the criteria described in the section that follows. If the packet is not an IPv4 packet containing a TCP segment, RSS is not done for the packet.
- **IPv4** — The Intel® Xeon® Processor D-1500 Product Family LAN controller parses the packet to identify an IPv4 packet. If the packet is not an IPv4 packet, RSS is not done for the packet.
- **TcpIPv6** — The Intel® Xeon® Processor D-1500 Product Family LAN controller parses the packet to identify an IPv6 packet containing a TCP segment per the criteria described in the section that follows. If the packet is not an IPv6 packet containing a TCP segment, RSS is not done for the packet.
- **IPv6** — The Intel® Xeon® Processor D-1500 Product Family LAN controller parses the packet to identify an IPv6 packet. If the packet is not an IPv6 packet, RSS is not done for the packet.

Tunneled IP to IP packets are considered for the RSS functionality as IP packets. The RSS logic ignores the L4 header while using the outer (first) IP header for the RSS hash.

For NVGRE and VXLAN packets the inner header (IP and TCP/UDP) is used for RSS computations.

The following additional cases are not part of the MSFT RSS specification but are supported RSS modes:

- **UdpIPv4** — The Intel® Xeon® Processor D-1500 Product Family LAN controller parses the packet to identify a packet with UDP over IPv4.



- UdpIPv6 — The Intel® Xeon® Processor D-1500 Product Family LAN controller parses the packet to identify a packet with UDP over IPv6.

A packet is identified as containing a TCP segment if all of the following conditions are met:

- The transport layer protocol is TCP (not UDP, ICMP, IGMP, etc.).
- The TCP segment can be parsed (such as IPv4 options or IPv6 extensions can be parsed, packet not encrypted, etc.).
- The packet is not fragmented (even if the fragment contains a complete L4 header).

Note: IPv6 extended headers are parsed by the Intel® Xeon® Processor D-1500 Product Family LAN controller, enabling TCP layer header recognition. Still the IPv6 extended header fields are not taken into account for the queue classification by RSS filter. This rule do not apply for security headers and fragmentation header. Packets with fragmentation header miss this filter. Packets with security extended headers are parsed only up to these headers and therefore can match only filters that do not require fields from the L4 protocol.

Bits[31:16] of the Multiple Receive Queues Command (MRQC), for single RSS and VFMRQC for multiple RSS, registers enable each of the above hash function variations (several might be set at a given time). If several functions are enabled at the same time, priority is defined as follows (skip functions that are not enabled):

- IPv4 packet
 - Try using the TcpIPv4 function
 - Try using UdpIPv4 function
 - Try using the IPv4 function
- IPv6 packet
 - Try using the TcpIPv6 function.
 - Try using UdpIPv6 function.
 - Try using the IPv6 function

The following combinations are currently supported:

- Any combination of IPv4, TcpIPv4, and UdpIPv4.

And/or:


- Any combination of either IPv6, TcpIPv6, and UdpIPv6.

When a packet cannot be parsed by the previous rules, it is assigned an RSS output index = zero. The 32-bit tag (normally a result of the hash function) equals zero.

The 32-bit result of the hash computation is written into the packet descriptor and also provides an index into the redirection table.

The following notation is used to describe the following hash functions:

- Ordering is little endian in both bytes and bits. For example, the IP address 161.142.100.80 translates into 0xa18e6450 in the signature.
- A “^” denotes bit-wise XOR operation of same-width vectors.
- @x-y denotes bytes x through y (including both of them) of the incoming packet, where byte 0 is the first byte of the IP header. In other words, we consider all byte-offsets as offsets into a packet where the framing layer header has been stripped out. Therefore, the source IPv4 address is referred to as @12-15, while the destination v4 address is referred to as @16-19.
- @x-y, @v-w denotes concatenation of bytes x-y, followed by bytes v-w, preserving the order in which they occurred in the packet.



All hash function variations (IPv4 and IPv6) follow the same general structure. Specific details for each variation are described in the following section. The hash uses a random secret key of length 320 bits (40 bytes); the key is stored in the RSS Random Key Register RSSRK for single RSS and VFRSSRK[63:0] for multiple RSS.

The algorithm works by examining each bit of the hash input from left to right. Our nomenclature defines left and right for a byte-array as follows: Given an array K with k bytes, our nomenclature assumes that the array is laid out as follows:

- K[0] K[1] K[2] ... K[k-1]

K[0] is the left-most byte, and the MSB of K[0] is the left-most bit. K[k-1] is the right-most byte, and the LSB of K[k-1] is the right-most bit.

ComputeHash(input[], N)

For hash-input input[] of length N bytes (8N bits) and a random secret key K of 320 bits

```
Result = 0;
For each bit b in input[] {
if (b == 1) then Result ^= (left-most 32 bits of K);
shift K left 1 bit position;
}
return Result;
```

7.1.3.6.3.1 Pseudo-code Examples

The following four pseudo-code examples are intended to help clarify exactly how the hash is to be performed in four cases, IPv4 with and without ability to parse the TCP header, and IPv6 with and without a TCP header.

Hash for IPv4 with TCP

Concatenate SourceAddress, DestinationAddress, SourcePort, DestinationPort into one single byte-array, preserving the order in which they occurred in the packet: Input[12] = @12-15, @16-19, @20-21, @22-23.

```
Result = ComputeHash(Input, 12);
```

Hash for IPv4 with UDP

Concatenate SourceAddress, DestinationAddress, SourcePort, DestinationPort into one single byte-array, preserving the order in which they occurred in the packet: Input[12] = @12-15, @16-19, @20-21, @22-23.

```
Result = ComputeHash(Input, 12);
```

Hash for IPv4 without TCP

Concatenate SourceAddress and DestinationAddress into one single byte-array.

```
Input[8] = @12-15, @16-19
```

```
Result = ComputeHash(Input, 8)
```

Hash for IPv6 with TCP



Similar to previous:

```
Input[36] = @8-23, @24-39, @40-41, @42-43
Result = ComputeHash(Input, 36)
```

Hash for IPv6 with UDP

Similar to previous:

```
Input[36] = @8-23, @24-39, @40-41, @42-43
Result = ComputeHash(Input, 36)
```

Hash for IPv6 without TCP

```
Input[32] = @8-23, @24-39
Result = ComputeHash(Input, 32)
```

7.1.3.6.4 Redirection Tables

The redirection table used for single RSS mode (RETA - Section 8.2.2.6.18 and ERETA - Section 8.2.2.6.19) is a 512-entry structure, indexed by the nine LSBs of the hash function output.

The redirection tables used for multiple RSS mode (VFRETA) are 64-entry structures indexed by the six LSBs of the hash function output. The table to use is defined by the pool to which the packet is sent.

System software might update the redirection tables during run time. Such updates of the table are not synchronized with the arrival time of received packets. Therefore, it is not guaranteed that a table update takes effect on a specific packet boundary.

7.1.3.6.5 RSS Verification Suite

Assume that the random key byte-stream is:

```
0x6d, 0x5a, 0x56, 0xda, 0x25, 0x5b, 0x0e, 0xc2,
0x41, 0x67, 0x25, 0x3d, 0x43, 0xa3, 0x8f, 0xb0,
0xd0, 0xca, 0x2b, 0xcb, 0xae, 0x7b, 0x30, 0xb4,
0x77, 0xcb, 0x2d, 0xa3, 0x80, 0x30, 0xf2, 0x0c,
0x6a, 0x42, 0xb7, 0x3b, 0xbe, 0xac, 0x01, 0xfa
```

Table 7.12. IPv4

Destination Address/Port	Source Address/Port	IPv4 Only	IPv4 With TCP
161.142.100.80 :1766	66.9.149.187 :2794	0x323e8fc2	0x51ccc178
65.69.140.83 :4739	199.92.111.2 :14230	0xd718262a	0xc626b0ea
12.22.207.184 :38024	24.19.198.95 :12898	0xd2d0a5de	0x5c2b394a
209.142.163.6 :2217	38.27.205.30 :48228	0x82989176	0xafc7327f
202.188.127.2 :1303	153.39.163.191 :44251	0x5d1809c5	0x10e828a2

Note: The IPv6 address tuples are only for verification purposes, and may not make sense as a tuple.




Table 7.13. IPv6

Destination Address/Port	Source Address/Port	IPv6 Only	IPv6 with TCP
3ffe:2501:200:3::1 (1766)	3ffe:2501:200:1fff::7 (2794)	0x2cc18cd5	0x40207d3d
ff02::1 (4739)	3ffe:501:8::260:97ff:fe40:efab (14230)	0x0f0c461c	0xdd51bbf
fe80::200:f8ff:fe21:67cf (38024)	3ffe:1900:4545:3:200:f8ff:fe21:67cf (44251)	0x4b61e985	0x02d1feef

7.1.4 Receive Data Storage in System Memory

The Intel® Xeon® Processor D-1500 Product Family LAN controller posts receive packets into data buffers in system memory.

The following controls are provided for the data buffers:

- The *SRRCTL[n].BSIZEPACKET* field defines the size of the data buffer pointed by each descriptor. The maximum packet size that can be posted to a queue can be limited using the *RXDCTL[n].RLPML* field. This filter enables software to use smaller buffers than the size defined by the *SRRCTL[n].BSIZEPACKET*.

Note: The packet size compared to *RXDCTL[n].RLPML* does not include any parts stripped by the device like CRC VLAN or other tags but include additional TimeStamp appended to the packet.

- The *SRRCTL.BSIZEHEADER* field defines the size of the header buffer pointed by each descriptor (advanced descriptors only).
- Each queue is provided with a separate *SRRCTL* register.

Receive memory buffer addresses are word (2 x byte) aligned (both data and headers).

The internal receive buffers are described in [Section 7.6.3.1](#).

7.1.5 Receive Descriptors

7.1.5.1 Legacy Receive Descriptor Format


A receive descriptor is a data structure that contains the receive data buffer address and fields for hardware to store packet information. Upon receipt of a packet for this device, hardware stores the packet data into the indicated buffer and writes the length, status and errors to the receive descriptor. If *SRRCTL[n].DESCTYPE* = zero, the Intel® Xeon® Processor D-1500 Product Family LAN controller uses the Legacy Rx descriptor as listed in [Table 7.14](#). The shaded areas indicate fields that are modified by hardware upon packet reception (so-called descriptor write-back).

Legacy descriptors should not be used when advanced features are enabled: SCTP, Virtualization, MACsec, time stamp in packet, tunnel packets checksum or RSC. Packets that match these cases might be dropped from queues that use legacy receive descriptors.

Refer to [Table 7.14](#) and the field descriptions that follow.

Table 7.14. Legacy Receive Descriptor (RDESC) Layout

	63	48 47	40 39	2 31	16 15	0
0	Buffer Address [63:0]					
8	VLAN Tag	Errors	Status	Fragment Checksum	Length	



Buffer Address (64-bit offset 0, 1st line)

Physical address in host memory of the received packet buffer.

Length Field (16-bit offset 0, 2nd line)

The length indicated in this field covers the data written to a receive buffer including CRC bytes (if any). Software must read multiple descriptors to determine the complete length for packets that span multiple receive buffers.

Fragment Checksum (16-bit offset 16, 2nd line)

This field is used to provide the fragment checksum value. This field is equal to the unadjusted 16-bit ones complement of the packet. Checksum calculation starts at the L4 layer (after the IP header) until the end of the packet excluding the CRC bytes. In order to use the fragment checksum assist to offload L4 checksum verification, software might need to back out some of the bytes in the packet. For more details see Section 7.1.6.5.

The fragment checksum is always reported in the descriptor with the EOP bit set.

Status Field (8-bit offset 32, 2nd line)

Status information indicates whether the descriptor has been used and whether the referenced buffer is the last one for the packet. Error status information is listed in Table 7.16.

Table 7.15. Receive Status (RDESC.STATUS) Layout

7	6	5	4	3	2	1	0
PIF	IPCS	L4CS	UDPCS	VP	Reserved	EOP	DD

End of Packet (EOP) and Descriptor Done (DD)

Refer to the following table:

DD	EOP	Description
0	0	Software setting of the descriptor when it hands it to the hardware.
0	1	Reserved (invalid option).
1	0	A completion status indication for non-last descriptor of a packet that spans across multiple descriptors. It means that the hardware is done with the descriptor and its buffers while only the length is valid on this descriptor.
1	1	A completion status indication of the entire packet. Software might take ownership of its descriptors while all fields in the descriptor are valid.

VP (VLAN Packet)

When set, the VP field indicates that the incoming packet's type is a VLAN (802.1q, matching the VLNCTRL.VET). If the RXDCTL.VME bit is set as well, then an active VP field also means that the VLAN has been stripped from the packet to the receive descriptor. For a further description of 802.1q VLANs please see Section 7.4.

IPCS (IPv4 Checksum), L4CS (L4 Checksum), UDPCS (UDP Checksum)

These bits are described in the following table.

Note: Switched packets from a local VM that do not use the Tx IP checksum offload by hardware have the IPCS equal to zero; switched packets from a local VM that do not use the Tx L4 checksum offload by hardware have the L4CS and UDPCS equal to zero.



L4CS	UDPCS	IPCS	Functionality
0	0	0	Hardware does not provide checksum offload.
0	0	1	Hardware provides IPv4 checksum offload. Pass/fail indication is provided in the <i>Error</i> field – IPE.
1	0	1 / 0	Hardware provides IPv4 checksum offload if IPCS is active along with TCP checksum offload. Pass/fail indication is provided in the <i>Error</i> field – IPE and TCPE.
1	1	1 / 0	Hardware provides IPv4 checksum offload if IPCS is active along with UDP checksum offload. Pass/fail indication is provided in the <i>Error</i> field – IPE and TCPE.

IPv6 packets do not have the *IPCS* bit set, but might have the *L4CS* bit and *UDPCS* bit set if the Intel® Xeon® Processor D-1500 Product Family LAN controller recognizes the transport header.

PIF (Non Unicast Address)

The *PIF* bit is set on packets with a non-unicast destination Ethernet MAC address — multicast and broadcast.

Error Field (8-bit offset 40, 2nd line)

Table 7.16 and the following text describes the possible errors reported by the hardware.

Table 7.16. Receive Errors (RDESC.ERRORS) Layout

7	6	5	4	3	2	1	0
IPE	TCPE	Reserved	Reserved	Reserved	Reserved	Reserved	RXE

IPE (IPv4 Checksum Error)

The IP checksum error is valid only when the *IPCS* bit in the *Status* field is set (indicating that the hardware validated the IP checksum). This bit is meaningful only on the last descriptor of a packet while the *EOP* bit is set as well. Packets with IP error are posted to host memory regardless of the store bad packet setting (FCTRL.SBP).

TCPE (TCP/UDP Checksum Error)

The TCP/UDP checksum error is valid only when the *L4CS* bit in the *Status* field is set (indicating that the hardware validated the L4 checksum). This bit is meaningful only on the last descriptor of a packet while the *EOP* bit is set as well. Packets with a TCP/UDP error are posted to host memory regardless of the store bad packet setting (FCTRL.SBP).


IPv4/UDP packets that carry a null UDP checksum field are reported with *L4CS*=0 and *TCPE*=0 (valid packet with no checksum).

IPv6/UDP packets that carry a null UDP checksum field are reported with *L4CS*=1 and *TCPE*=1 (UDP checksum is mandatory over IPv6).

RXE

The *RXE* error bit is an indication for any MAC error. It is a logic OR function of the following errors:

- CRC or symbol error might be a result of receiving a /V/ symbol on the TBI interface, /FE/ symbol on the GMII/XGMII interface, RX_ER assertion on GMII interface, bad EOP or loss of sync during packet reception.
- Undersize frames shorter than 64 bytes.
- Oversize frames larger than the MFS definition in the MAXFRS register.
- Length error in 802.3 packet format. Length field is not checked in presence of a VLAN or E-tag.



Packets with an RXE error are posted to host memory only when store bad packet bit (FCTRL.SBP) is set.

VLAN Tag Field (16-bit offset 48, 2nd line)

If the RXDCTL.VME is set and the received packet type is 802.1q (as defined by VLNCTRL.VET) then the VLAN header is stripped from the packet data storage. In this case the 16 bits of the VLAN tag, priority tag and DEI from the received packet are posted to the VLAN Tag field in the receive descriptor. Otherwise, the VLAN Tag field contains 0x0000.

Table 7.17. VLAN Tag Field Layout (for 802.1q Packet)

15	13	12	11	0
PRI		DEI		VLAN

Priority and DEI are part of 802.1Q specifications. The VLAN field is provided in network order.

7.1.5.2 Advanced Receive Descriptors

7.1.5.2.1 Advanced Receive Descriptors — Read Format

Table 7.18 lists the advanced receive descriptor programming by the software. The SRRCTL[n].DESCTYPE should be set to a value other than 000b when using the advanced descriptor format.

Table 7.18. Descriptor Read Format

	63	1	0
0	Packet Buffer Address [63:1]		
8	Header Buffer Address [63:1]		DD

Packet Buffer Address (64)

The physical address in host memory of the packet buffer.

Header Buffer Address (64)

The physical address in host memory of the header buffer with the lowest bit being Descriptor Done (DD). When a packet spans in multiple descriptors, only the header buffer of the first descriptor is used. In subsequent descriptors, only the data buffer is used.

During the programming phase, software must set the DD bit to zero (see the description of the DD bit in this section). This means that header buffer addresses are always word aligned.

Note: The Intel® Xeon® Processor D-1500 Product Family LAN controller does not support null descriptors meaning packet or header addresses are zero.

7.1.5.2.2 Advanced Receive Descriptors — Write-Back Format

When the Intel® Xeon® Processor D-1500 Product Family LAN controller writes back the descriptors, it uses the format listed in Table 7.19. The advanced descriptor write-back format is used when SRRCTL[n].DESCTYPE is set to a value other than 000b.



Table 7.19. Descriptor Write-Back Format

	63	48	47	32	31	30	21	20	17	16	4	3	0
0	RSS Hash / Fragment Checksum / Flow Director Filters ID				SPH	HDR_LEN		RSCCNT		Packet Type			RSS Type
8	VLAN Tag		PKT_LEN		Extended Error				Extended Status / NEXTP				
	63	48	47	32	31	20			19			0	

RSS Type (4-bit offset 0, 1st line)

The Intel® Xeon® Processor D-1500 Product Family LAN controller must identify the packet type and then choose the appropriate RSS hash function to be used on the packet. The RSS type reports the packet type that was used for the RSS hash function.

RSS Type	Description
0x0	No hash computation done for this packet
0x1	HASH_TCP_IPv4
0x2	HASH_IPv4
0x3	HASH_TCP_IPv6
0x4	Reserved
0x5	HASH_IPv6
0x6	Reserved
0x7	HASH_UDP_IPv4
0x8	HASH_UDP_IPv6
0x9 - 0xE	Reserved
0xF	Packet reports flow director filters status (set if a flow director match, even if a packet is forwarded by another filter).

Packet Type (13-bit at offset 4, 1st line)

The *Packet Type* field reports the packet type identified by the hardware as follows. Note that some of the fields in the receive descriptor are valid for specific packet types.

Bit Index	Bit 11 = 0	Bit 11 = 1 (L2 packet matching one of the ETQF filters)
0	IPV4 — IPv4 header present	Ethertype — ETQF register index that matches the packet. Special types are defined for 802.1X, 1588.
1	IPV4O — IPv4 with options	
2	IPV6 — IPv6 header present	Reserved.
3	IPV6E- IPv6 with extensions	
4	TCP — TCP header present	



Bit Index	Bit 11 = 0	Bit 11 = 1 (L2 packet matching one of the ETQF filters)
5	UDP — UDP header present	ReservedMACsec – MACsec encapsulation
6	SCTP — SCTP header	
7	If bit 12 = 0 - Reserved If bit 12 = 1: 0 = NVGRE 1 = VXLAN	
8		
9		
10	MACsec – MACsec encapsulation	
11	0b = non L2 packet	
12	NVGRE or VXLAN Tunnel Packet. If this bit is set, then bit 7 indicates the tunnel type.	Reserved

Notes: UDP, TCP and IPv6 indications are not set in any IPv4 fragmented packet.

In virtualization mode, packets might be received from other local VMs. The Intel® Xeon® Processor D-1500 Product Family LAN controller does not check the L5 header for these packets and does not report NFS header. If such packets carry IP tunneling (IPv4 — IPv6), they are reported as IPV4E. The packets received from local VM are indicated by the *LB* bit in the status field. In order to be identified, the *CC* bit in the transmit descriptor must be set and if it is a tunnel packet, the *TUNNEL.OUTERIPCS* must also be set.

If bit 12 (Tunnel Packet) is set, then all the L3/L4 indications (IPV4, IPV4O, IPV6, IPV6E, TCP, UDP, SCTP) reflects the inner header status.

RSC Packet Count- RSCCNT (4-bit offset 17, 1st line)

The *RSCCNT* field is valid only for RSC descriptors while in non-RSC it equals zero. *RSCCNT* minus one indicates the number of coalesced packets that start in this descriptor. *RSCCNT* might count up to 14 packets. Once 14 packets are coalesced in a single buffer, RSC is closed enabling accurate coalesced packet count. If the *RSCCNTBPP* bit in *RDRXCTL* is set, coalescing might proceed beyond the 14 packets per buffer while *RSCCNT* stops incrementing beyond 0xF.

Note: Software can identify RSC descriptors by checking the *RSCCNT* field for non-zero value.

HDR_LEN (10-bit offset 21, 1st line)

The *HDR_LEN* reflects the size of the packet header in byte units (if the header is decoded by the hardware). This field is meaningful only in the first descriptor of a packet and should be ignored in any subsequent descriptors. Header split is explained in [Section 7.1.6](#) while the packet types for this functionality are enabled by the *PSRTYPE[n]* registers.

Split Header — SPH (1-bit offset 31, 1st line)

When set to 1b, indicates that the hardware has found the length of the header. If set to 0b, the header buffer might be used only in split always mode. If the received header size is greater or equal to 1024 bytes, the *SPH* bit is not set and header split functionality is not supported. The *SPH* bit is meaningful only on the last descriptor (EOP) of a packet. See additional details on *SPH*, *PKT_LEN* and *HDR_LEN* as a function of split modes in [Table 7.24](#).



Packet Type	Header Length (includes all fields up to the field specified)	Header Split
Un-recognized Ethertype only with / without SNAP and with / without VLAN or packets that match the L2 filters (MTQF).	VLAN header(s) if present Else, EtherType field	No
Pv4 only or fragmented IPv4 with any payload including IPv4-IPv6 tunneling	IPv4 header	Enabled
Non-fragmented IPv4, TCP / UDP / SCTP	L4 header	Enabled
IPv4-IPv6, only or fragmented IPv4-IPv6 at IPv6 header with any payload	IPv6 header (up to the fragment extension header if exist)	Enabled
IPv4-IPv6, TCP / UDP / SCTP	L4 header	Enabled
VXLAN or NVGRE	Cloud Header/L2 Header	Enabled ¹

1. If cloud split modes used (PSRTYPE15 or PSRTYPE16 are set). If not, then according to internal header.

RSS Hash or Flow Director Filters ID (32-bit offset 32, 1st line) / Fragment Checksum (16-bit offset 48, 1st line)

This field has multiplexed functionality according to the received packet type (reported on the *Packet Type* field in this descriptor) and device setting.

Fragment Checksum

For fragmented UDP/IP packets, this field holds the UDP fragment checksum (described in [Section 7.1.6.5](#)) if both the *RXCSUM.PCSD* bit is cleared and *RXCSUM.IPPCSE* bit is set. This field is meaningful only for UDP packets where the *UDPV* bit in the Extended Status word is set. When Fragment Checksum is reported, bits 47:32 are invalid.

The checksum does not include any padding or timestamp added by the device.

RSS Hash / Flow Director Filters ID

- If the *FDIRCTRL.Report Status* bit is set, the flow director filters ID is reported;
- If the RSS Type is non zero, the RSS hash is reported.
- Otherwise, the value is not valid.

Table 7.20. Checksum Enable/Disable

RXCSUM.PCSD	0 (Checksum Enable)	1 (Checksum Disable)
	Fragment Checksum and IP Identification are reported in the Rx Descriptor.	RSS Hash value is reported in the Rx Descriptor.

RSS Hash

The RSS hash value is required for RSS functionality as described in [Section 7.1.3.6](#). Note that the RSS hash is meaningful only for 'RSS Type' in the range 0x1 to 0x8.



Flow Director Filters ID

The flow director filters ID is reported only when the received packet matches a flow directory filter (see [Section 7.1.3.5](#)). The flow director filter ID field has a different structure for signature-based filters and perfect match filters as follows:

Filter Type	31	30 29	28 16	15 13	12 0
Hash-based Flow Director Filter ID	Rsv	Bucket Hash		Signature	
Perfect Match Flow Director Filter ID	Rsv		Hash	Rsv	SW-Index

Bucket Hash

A hash value that identifies a flow director bucket. When the Flow Director table is smaller than 32K filters the bucket hash is smaller than 15 bits. In this case the upper bit(s) are set to zero.

Signature

A hash value used to identify flow within a bucket.

SW-Index

The SW-Index that is taken from the filter context, programmed by software. It is meaningful only when the *FLM* bit in the Extended Status is set as well.

Rsv

Reserved.

Extended Status / NEXTP (20-bit offset 0, 2nd line)

Status information indicates whether the descriptor has been used and whether the referenced buffer is the last one for the packet. [Table 7.21](#) lists the extended status word in the last descriptor of a packet (*EOP* bit is set). [Table 7.22](#) lists the extended status word in any descriptor but the last one of a packet (*EOP* bit is cleared).

Table 7.21. Receive Status (RDESC.STATUS) Layout of Last Descriptor

19	18	17	16	15	14	13	12	11	10
BMC	LB	SECP	TS	TSIP	Rsv				UDPV
VEXT	OUTERIPCS	PIF	IPCS	L4I	UDPCS	VP	FLM	EOP	DD
			FCEOFs	FCSTAT					
9	8	7	6	5	4	3	2	1	0

Table 7.22. Receive Status (RDESC.STATUS) Layout of Non-Last Descriptor

19	4	3 : 2	1	0
Next Descriptor Pointer — NEXTP		Rsv	EOP = 0b	DD

Rsv (14:12) — Reserved at zero.

FLM(2) — Flow director filter match indication is set for packets that match these filters.

VP(3), PIF (7) — These bits are described in the legacy descriptor format in [Section 7.1.5](#).

The VP bit is not set even if a VLAN tag is present if the *PFQDE.HIDE_VLAN* bit is set for the queue.

EOP (1) and DD (0) — End of Packet and Done bits are listed in the following table:



DD	EOP	Description
0	X	Software setting of the descriptor when it hands it to hardware.
1	0	A completion status indication for a non last descriptor of a packet (or multiple packets in the case of RSC) that spans across multiple descriptors. In a single packet case the <i>DD</i> bit indicates that the hardware is done with the descriptor and its buffers. In the case of RSC, the <i>DD</i> bit indicates that the hardware is done with the descriptor but might still use its buffers (for the coalesced header) until the last descriptor of the RSC completes. Only the <i>Length</i> fields are valid on this descriptor. In the RSC case, the next descriptor pointer and <i>RSCCNT</i> are valid as well.
1	1	A completion status indication of the entire packet (or the multiple packets in the case of RSC) and software might take ownership of its descriptors. All fields in the descriptor are valid (reported by the hardware).

UDPCS(4), L4I (5) / FCSTAT (5:4) — This field has multiplexed functionality for non-FCoE packets. For non-FCoE packets this field is UDPCS and L4I. The UDPCS (UDP checksum) is set (together with L4CS bit) when hardware provides UDP checksum offload. The L4I (L4 integrity) is set when hardware provides any L4 offload as: UDP checksum, TCP checksum or SCTP CRC offload. **IPCS(6), FCEOFs (6)** — This bit has multiplexed functionality for non-FCoE packets.

For non-FCoE packets it is IPCS as described in Legacy Rx descriptor (in [Section 7.1.5](#)).

OUTERIPCS(8) — Indicates that a checksum was done on the outer IP header of an NVGRE or VXLAN packet.

VEXT (9) — Outer-VLAN is found on a double VLAN packet. This bit is valid only when CTRL_EXT.EXTENDED_VLAN is set. See more details in [Section 7.4.5](#).

UDPV (10) — The *UDP Checksum Valid* bit indicates that the incoming packet contains a valid (non-zero value) checksum field in an incoming fragmented (non-tunneled) UDP IPv4 packet. It means that the *Fragment Checksum* field in the receive descriptor contains the UDP checksum as described in [Section 7.1.6.5](#). When this field is cleared in the first fragment that contains the UDP header, it means that the packet does not contain a valid UDP checksum and the checksum field in the Rx descriptor should be ignored. This field is always cleared in incoming fragments that do not contain the UDP header.

TSIP (15) - Timestamp in packet. The *Timestamp In Packet* bit is set to indicate that the received packet arrival time was captured by the hardware and the timestamp was placed in the receive buffer. For more details see [Section 7.6](#) and [Section 7.1.6.2](#).

TS (16) — The *Time Stamp* bit is set when the device recognized a time sync packet. In such a case the hardware captures its arrival time and stores it in the Time Stamp register. For more details see [Section 7.6](#).

SECP (17) — Security processing bit indicates that the hardware identified the security encapsulation and processed it as configured.

MACsec processing — This bit is set each time MACsec processing is enabled regardless if a matched SA was found.

IPsec processing — This bit is set only if a matched SA was found. Note that hardware does not process packets with an IPv4 option or IPv6 extension header and the SECP bit is not set. This bit is not set for IPv4 packets shorter than 70 bytes, IPv6 ESP packets shorter than 90 bytes, or IPv6 AH packets shorter than 94 bytes (all excluding CRC). Note that these packet sizes are never expected and set the length error indication in the *SECERR* field.

LB (18) — This bit provides a loopback status indication which means that this packet is sent by a local VM (VM to VM switch indication).

BMC (19) - Packet received from BMC. The BMC bit is set to indicate the packet was sent by the local BMC. Bit is cleared if packet arrives from the network. For more details see [Section 10.4.2.2](#).



NEXTP (19:4) — Large receive might be composed of multiple packets and packets might span in multiple buffers (descriptors). These buffers are not guaranteed to be consecutive while the *NEXTP* field is a pointer to the next descriptor that belongs to the same RSC. The *NEXTP* field is defined in descriptor unit (the same as the head and tail registers). The *NEXTP* field is valid for any descriptor of a large receive (the *EOP* bit is not set) except the last one. It is valid even in consecutive descriptors of the same packet. In the last descriptor (on which the *EOP* bit is set), *NEXTP* is not indicated but rather all other status fields previously described in this section.

Extended Error (12-bit offset 21, 2nd line)

Table 7.23 and the following text describe the possible errors reported by hardware.

Table 7.23. Receive Errors (RDESC.ERRORS) Layout

11	10	9	8:7	6	5:4	3	2:0
IPE	L4E	RXE	SECERR	OUTERIPER	Rsv	HBO	FCERR / FDIRERR
FCEOF _e							

FDIRERR (2:0) — This field is relevant for non-FCoE packets when the flow director filters are enabled.

FDIRErr(0) - Length — If the flow director filter matches the *Length* bit, this indicates that the distance of the matched filter from the hash table exceeds the *FDIRCTRL.Max-Length*. If there is no matched filter, the *Length* bit is set if the flow director linked list of the matched hash value exceeds the *FDIRCTRL.Max-Length*.

FDIRErr(1) - Drop — The *Drop* bit indicates that a received packet matched a flow director filter with a drop action. In the case of perfect mode filtering, it is expected to find the drop indication only when the linked list in the flow director bucket exceeds the permitted *Max-Length*. In this case, the packet is not dropped. Instead, it is posted to the Rx queue (indicated in the filter context) for software handling of the *Max-Length* exception. In the case of hash mode filtering, it is expected that the drop queue is always a valid queue so all packets that match the drop filter are visible to software.

FDIRErr(2) - Coll — A matched flow director filter with a collision indication was found. The collision indicates that software attempted to step over this filter with a different action that was already programmed.

HBO (3) — The *Header Buffer Overflow* bit is set if the packet header (calculated by hardware) is bigger than the header buffer (defined by *PSRCTL.BSIZEHEADER*). *HBO* reporting might be used by software to allocate bigger buffers for the headers. It is meaningful only if the *SPH* bit in the receive descriptor is set as well. The *HDR_LEN* field is valid even when the *HBO* bit is set. Packets with *HBO* error are posted to host memory regardless of the store bad packet setting (*FCTRL.SBP*). Packet DMA to its buffers when the *HBO* bit is set, depends on the device settings as follows:

SRCTL.DESCTYPE	DMA Functionality
Header Split (010b)	The header is posted with the rest of the packet data to the packet buffer.
Always Split Mode (101b)	The header buffer is used as part of the data buffers and contains the first <i>SRCTL.BSIZEHEADER</i> bytes of the packet.

Rsv (5:4) — Reserved at zero.

OUTERIPER (6) — Indicates an error was found in the checksum of an outer IP header of a VXLAN or NVGRE packet or that the UDP checksum of the outer UDP header in a VXLAN



packet was not zero. **SECERR (8:7)** — Security error indication for MACsec or IPsec. This field is meaningful only if the SECP bit in the extended status is set.

SECERR	MACsec Error Reporting	IPsec Error Reporting
00	No errors found or no security processing	No errors found while an active SA found or no security processing.
01	No SA match	Invalid IPsec Protocol: IPsec protocol field (<i>ESP</i> or <i>AH</i>) in the received IP header does not match expected one stored in the SA table.
10	Replay error	Length error: ESP packet is not 4-bytes aligned or AH/ESP header is truncated (for example, a 28-byte IPv4 packet with IPv4 header + ESP header that contains only SPI and SN) or AH Length field in the AH header is different than 0x07 for IPv4 or 0x08 for IPv6 or the entire packet size excluding CRC is shorter than 70 bytes for IPv4 or 90 bytes for IPv6 ESP or 94 bytes for IPv6 AH.
11	Authentication failed: Bad signature	Authentication failed: Bad signature

RXE (9) — RXE is described in the legacy descriptor format in [Section 7.1.5](#).

L4E (10) — L4 integrity error is valid only when the *L4I* bit in the *Status* field is set. It is active if L4 processing fails (TCP checksum or UDP checksum or SCTP CRC). Packets with L4 integrity error are posted to host memory regardless of the store bad packet setting (FCTRL.SBP). In case of VXLAN or NVGRE packets, this relates to the internal L4 header.

FCEOF_e(11) / IPE(11) — This bit has multiplexed functionality.

Non-FCoE Packet	Reserved
IPE (IPv4 checksum error) is described in Section 7.1.5 . In case of VXLAN or NVGRE packets, this relates to the internal IP header.	Reserved.

PKT_LEN (16-bit offset 32, 2nd line)

PKT_LEN holds the number of bytes posted to the packet buffer. The length covers the data written to a receive buffer including posted CRC bytes (if any). Software must read multiple descriptors to determine the complete length for packets that span multiple receive buffers. If SRRCTL.DESCTYPE = 2 (advanced descriptor header splitting) and the buffer is not split because the header is bigger than the allocated header buffer, this field reflects the size of the data written to the data buffer (header + data).

When short packets are padded, the PKT_LEN does not include the padding, so that the software device driver can detect the location of a potential time stamp in the packet.

When a time stamp is added (TSIP bit is set), the PKT_LEN includes the timestamp size (+8).

VLAN Tag (16-bit offset 48, 2nd line)

This field is described in the legacy descriptor format in [Section 7.1.5](#). The VLAN tag is set to 0x0000 even if a VLAN tag is present if the *PFQDE.HIDE_VLAN* bit is set for the queue.

7.1.5.3 Receive Descriptor Fetching

The Intel® Xeon® Processor D-1500 Product Family LAN controller implements a fetch-by-demand mechanism for descriptor fetch. Descriptors are not fetched in advance, but rather fetched after a packet is received. Such a strategy eliminates the need to store descriptors on-die for each and every descriptor queue in anticipation for packet arrival.

7.1.5.4 Receive Descriptor Write-Back

The Intel® Xeon® Processor D-1500 Product Family LAN controller writes back the receive descriptor immediately following the packet write into system memory. It is therefore possible for a single descriptor to be written at a time into memory. However, if aggregation occurs during descriptor fetch (see Section 7.1.5.3), then the descriptors fetched in the aggregated operation are written back in a single write-back operation. In Receive Coalescing (RSC), all the descriptors except the last one are written back when they are completed. This does not have to be on packet boundaries but rather when the next descriptor of the same RSC is fetched. See Section 7.9.5.1 for more on RSC.

Note: Software can determine if a packet has been received only by checking the receive descriptor *DD* bit in memory. Checking through *DD* bits (and not by checking the receive head pointer in RDH/RDL registers) eliminates a potential race condition: all descriptor data is posted internally prior to incrementing the head register and a read of the head register could potentially pass the descriptor waiting inside the Intel® Xeon® Processor D-1500 Product Family LAN controller.

7.1.5.5 Receive Descriptor Queue Structure

Figure 7.10 shows the structure of each of the receive descriptor rings. Note that each ring uses a contiguous memory space.

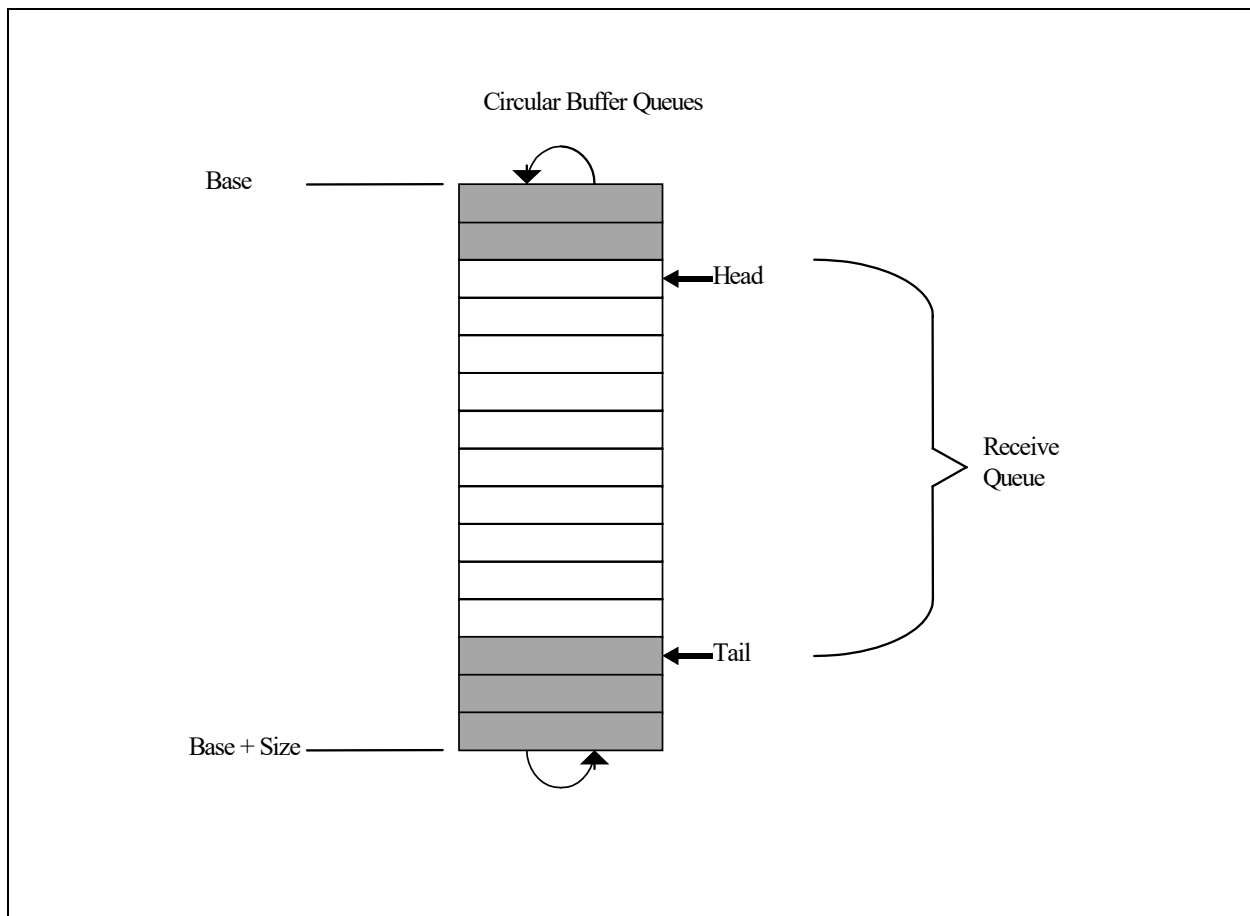



Figure 7.10. Receive Descriptor Ring Structure



Software inserts receive descriptors by advancing the tail pointer(s) to refer to the address of the entry just beyond the last valid descriptor. This is accomplished by writing the descriptor tail register(s) with the offset of the entry beyond the last valid descriptor. The Intel® Xeon® Processor D-1500 Product Family LAN controller adjusts its internal tail pointer(s) accordingly. As packets arrive, they are stored in memory and the internal head pointer(s) is increased by the Intel® Xeon® Processor D-1500 Product Family LAN controller.

When RSC is not enabled, the visible (external) head pointer(s) reflect the internal ones. On any receive queue that enables RSC, updating the external head pointer might be delayed until interrupt assertion. When the head pointer(s) is equal to the tail pointer(s), the queue(s) is empty. The Intel® Xeon® Processor D-1500 Product Family LAN controller stops storing packets in system memory until software advances the tail pointer(s), making more receive buffers available.





Figure 7.11. Descriptors and Memory Rings

The Intel® Xeon® Processor D-1500 Product Family LAN controller writes back used descriptors just prior to advancing the head pointer(s). Head and tail pointers wrap back to base when the number of descriptors corresponding to the size of the descriptor ring have been processed.



The receive descriptor head and tail pointers reference to 16-byte blocks of memory. Shaded boxes in Figure 7.11 represent descriptors that have stored incoming packets but have not yet been recognized by software. Software can determine if a receive buffer is valid by reading descriptors in memory rather than by I/O reads. Any descriptor with a *DD* bit set has been used by the hardware, and is ready to be processed by software.

Note: The head pointer points to the next descriptor that is to be written back. At the completion of the descriptor write-back operation, this pointer is increased by the number of descriptors written back. Hardware owns all descriptors between [head... tail]. Any descriptor not in this range is owned by software.

The receive descriptor rings are described by the following registers:

- Receive Descriptor Base Address registers (RDBA) — This register indicates the start of the descriptor ring buffer; this 64-bit address is aligned on a 16-byte boundary and is stored in two consecutive 32-bit registers. Hardware ignores the lower 4 bits.
- Receive Descriptor Length registers (RDLEN) — This register determines the number of bytes allocated to the circular buffer. This value must be a multiple of 128 (the maximum cache line size). Since each descriptor is 16 bytes in length, the total number of receive descriptors is always a multiple of 8.
- Receive Descriptor Head registers (RDH) — This register holds a value that is an offset from the base, and indicates the in-progress descriptor. There can be up to 64K-8 descriptors in the circular buffer. Hardware maintains a shadow copy that includes those descriptors completed but not yet stored in memory.

Software can determine if a packet has been received by either of two methods: reading the *DD* bit in the receive descriptor field or by performing a Programmed I/O read of the Receive Descriptor Head register. Checking the descriptor *DD* bit in memory eliminates a potential race condition. All descriptor data is written to the I/O bus prior to incrementing the head register, but a read of the head register could pass the data write in systems performing I/O write buffering. Updates to receive descriptors use the same I/O write path and follow all data writes. Consequently, they are not subject to the race.

- Receive Descriptor Tail registers (RDT) — This register holds a value that is an offset from the base, and identifies the location beyond the last descriptor hardware can process. This is the location where software writes the first new descriptor.

If software statically allocates buffers, and uses a memory read to check for completed descriptors, it simply has to zero the status byte in the descriptor to make it ready for re-use by hardware. This is not a hardware requirement, but is necessary for performing an in-memory scan. This is relevant only to legacy descriptors.

All the registers controlling the descriptor rings behavior should be set before receive is enabled, apart from the tail registers which are used during the regular flow of data.

7.1.5.5.1 Low Receive Descriptors Threshold

As previously described, the size of the receive queues is measured by the number of receive descriptors. During run time, software processes descriptors and upon completion of descriptors, increments the Receive Descriptor Tail registers. At the same time, the hardware may post new received packets incrementing the Receive Descriptor Head registers for each used descriptor.

The number of usable (free) descriptors for the hardware is the distance between the Tail and Head registers. When the tail reaches the head, there are no free descriptors and further packets might be either dropped or block the receive FIFO. In order to avoid this situation, the Intel® Xeon® Processor D-1500 Product Family LAN controller might generate a low latency interrupt (associated to the relevant Rx queue) once there are less equal free descriptors than specified by a low level threshold. The threshold is defined in 64 descriptors granularity per queue in the *SRRCTL[n].RDMTS* field.



7.1.6 Receive Offloads

7.1.6.1 Header Splitting

7.1.6.1.1 Purpose

This feature consists of splitting a packet header to a different memory space. This helps the host to fetch headers only for processing: headers are posted through a regular snoop transaction in order to be processed by the host CPU. It is recommended to perform this transaction with TPH enabled (see Section 7.5).

The Intel® Xeon® Processor D-1500 Product Family LAN controller’s support for header split is controlled by the *DESCTYPE* field of the Split Receive Control registers (*SRRCTL*). The following modes exist in both split and non-split modes:

- 000b: Legacy mode - Legacy descriptors are used, headers and payloads are not split.
- 001b: Advanced mode, no split - Advanced descriptors are in use, header and payload are not split.
- 010b: Advanced mode, header split - Advanced descriptors are in use, header and payload are split to different buffers.
- 101b: Advanced mode, split always - Advanced descriptors are in use, header and payload are split to different buffers. If no split is done, the first part of the packet is stored in the header buffer.

The Intel® Xeon® Processor D-1500 Product Family LAN controller uses packet splitting when the *SRRCTL[n].DESCTYPE* is greater than one.

7.1.6.1.2 Description

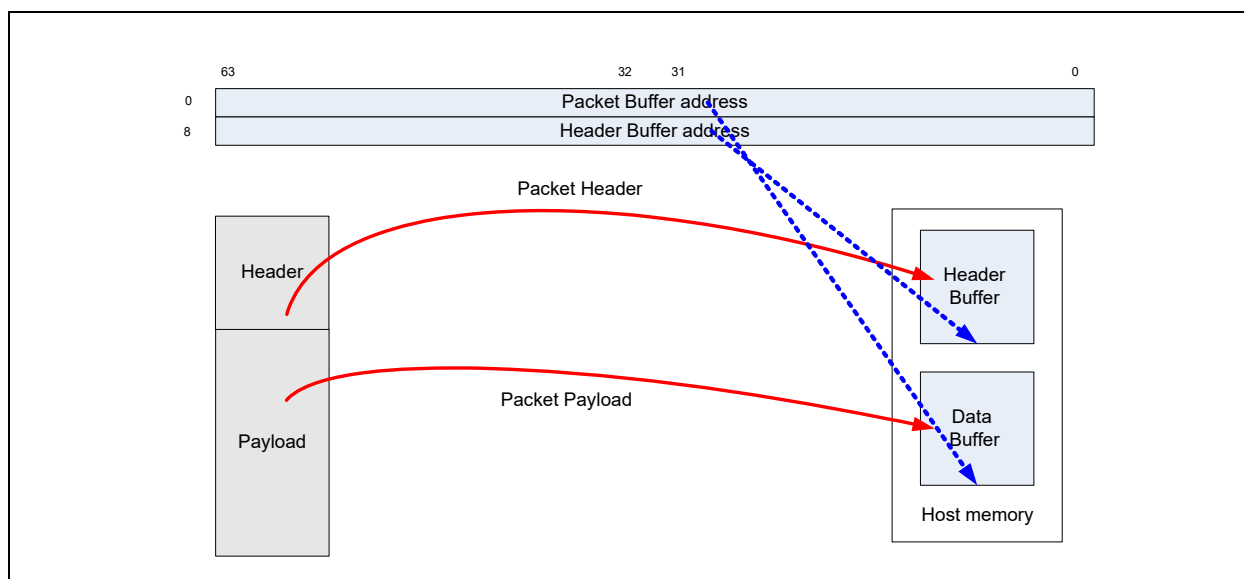


Figure 7.12. Header Splitting Diagram

The physical address of each buffer is written in the *Buffer Addresses* fields:

- The packet buffer address includes the address of the buffer assigned to the packet data.



- The header buffer address includes the address of the buffer that contains the header information. The receive DMA module stores the header portion of the received packets into this buffer.

The sizes of these buffers are statically defined in the SRRCTL[n] registers:

- The *BSIZEPACKET* field defines the size of the buffer for the received packet.
- The *BSIZEHEADER* field defines the size of the buffer for the received header. If header split is enabled, this field must be configured to a non-zero value. The Intel® Xeon® Processor D-1500 Product Family LAN controller only writes the header portion into the header buffer. The header size is determined by the options enabled in the PSRTYPE registers.

When header split is selected, the packet is split only on selected types of packets. A bit exists for each option in PSRTYPE[n] registers, so several options can be used in conjunction. If one or more bits are set, the splitting is performed for the corresponding packet type. See Section 8.2.2.6.4 for details on the possible header types supported. In virtualization mode, a separate PSRTYPE register is provided per pool up to the number of pools enabled. In non-virtualization mode, only PSRTYPE[0] is used.

Rules regarding header split:

- Packets that have headers bigger than 1023 bytes are not split.
- The header of a fragmented IPv6 packet is defined until the *fragment* extension header.
- An IP in IP packet (such as any combination of IPv4 and IPv6 tunneling) is not split. Not relevant for NVGRE and VXLAN packets.
- Packet header cannot span across buffers, therefore, the size of the header buffer must be larger than any expected header size. In case of header split mode (SRRCTL.DESCTYPE = 010b), a packet with a header larger than the header buffer is not split.

Table 7.24 lists the behavior of the Intel® Xeon® Processor D-1500 Product Family LAN controller in the different modes.

Table 7.24. Behavior in Header Split Modes

DESCTYPE	Condition	SPH	HBO ¹	PKT_LEN ²	HDR_LEN ²	Header and Payload DMA ³
Split	1. Header can't be decoded	0	0	Min (packet length, buffer size)	0x0	Header + Payload (+ TimeStamp) (+ Padding) -> Packet Buffer
	2. Header <= BSIZEHEADER and Payload > 0	1	0	Min (payload length, buffer size)	Header size	Header -> Header Buffer Payload (+ TimeStamp) (+ Padding) -> Packet Buffer
	3. Header <= BSIZEHEADER and Payload = 0 (Header only packet)	1	0	0	Header size	Header (+ TimeStamp) (+ Padding) -> Header Buffer ⁴
	4. Header > BSIZEHEADER	1	1	Min (packet length, buffer size)	Header size ⁵	Header + Payload (+ TimeStamp) -> Packet Buffer
Split – always use header buffer	1. Header can't be decoded and packet length <= BSIZEHEADER	0	0	0x0	Packet length	Header + Payload (+ TimeStamp) (+ Padding) -> Header Buffer
	2. Header can't be decoded and packet length > BSIZEHEADER	0	0	Min (packet length – BSIZEHEADER, data buffer size)	BSIZEHEADER	Header + Payload (+ TimeStamp) -> Header + Packet Buffers ⁵
	3. Header <= BSIZEHEADER	1	0	Min (payload length, data buffer size)	Header Size	Header -> Header Buffer Payload (+ TimeStamp) (+ Padding) -> Packet Buffer
	4. Header > BSIZEHEADER	1	1	Min (packet length – BSIZEHEADER, data buffer size)	Header Size ⁶	Header + Payload (+ TimeStamp) -> Header + Packet Buffer

1. HBO is set to 1b if the header size is bigger than BSIZEHEADER and zero otherwise.



2. PKT_LEN and HDR_LEN includes optionally added time stamp if TSIP bit is set and does not include padding added by the device.
3. Partial means up to BSIZEHEADER.
4. Even if there is no payload at all and there is a timestamp in the packet - the timestamp is considered as an 8 bytes payload and is written to the packet buffer.
5. If the packet spans more than one descriptor, only the header buffer of the first descriptor is used.
6. HDR_LEN doesn't reflect the actual data size stored in the header buffer. It reflects the header size determined by the parser.

7.1.6.2 Receive Packet Timestamp in Buffer

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports adding an optional tailored header appended to the end of the packet in the receive buffer. The tailored header includes a 64 bit timestamp composed of the packet reception time measured in the *SYSTIMEL* (Low DW) and *SYSTIMEH* (High DW) registers (See [Section 7.6.3.2](#) for further information on *SYSTIMEL/H* operation). The timestamp starts right at the end of the packet (after the last byte). It is sent as {*SYSTIMEH*, *SYSTEML*} - most significant byte first (closest to packet).

PKT_LEN and *HDR_LEN* includes optionally added time stamp.

When the *TSAUXC.Disable_systemtime* bit is cleared and the *TSYNCRXCTL.TSIP_UP_EN* is set for the UP in the packet (or *TSYNCRXCTL.TSIP_UT_EN* for un-tagged packets), packets received that meet the *TSYNCRXCTL.Type* will be time stamped. A packet that was time stamped is reported as follows:

- Place a 64 bit timestamp, indicating the time a packet was received by the MAC, appended at the end of the received packet within the receive buffer.
- Set the TSIP bit in the RDESC.STATUS field of the last receive descriptor.

Note: When packets are coalesced, the time stamp reflects the reception time of the last coalesced fragment, unless the packet is a pure ACK packet, in which case, the time stamp is the first packet.

7.1.6.3 Receive Checksum Offloading

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports the offloading of three receive checksum calculations: the fragment checksum, the IPv4 header checksum, and the TCP/UDP checksum.

For supported packet/frame types, the entire checksum calculation can be offloaded to the Intel® Xeon® Processor D-1500 Product Family LAN controller. The Intel® Xeon® Processor D-1500 Product Family LAN controller calculates the IPv4 checksum and indicates a pass/fail indication to software via the *IPv4 Checksum Error* bit (RDESC.IPE) in the *ERROR* field of the receive descriptor. Similarly, the Intel® Xeon® Processor D-1500 Product Family LAN controller calculates the TCP or UDP checksum and indicates a pass/fail condition to software via the *TCP/UDP Checksum Error* bit (RDESC.TCPE). For NVGRE or VXLAN packets, the *IPE* and *TCPE* bits relates to the inner IP/TCP header. The outer IPv4 checksum is also checked and a pass/fail indication is indicated to software via the Outer IPv4 Checksum Error bit (*OUTERIPER*). These error bits are valid when the respective status bits indicate the checksum was calculated for the packet (*RDESC.IPCS*, *RDESC.L4CS*, and *RDESC.OUTERIPCS* respectively).

Similarly, if *RFCTL.Ipv6_DIS* and *RFCTL.IP6Xsum_DIS* are cleared to zero, the Intel® Xeon® Processor D-1500 Product Family LAN controller calculates the TCP or UDP checksum for IPv6 packets. It then indicates a pass/fail condition in the *TCP/UDP Checksum Error* bit (RDESC.TCPE).

Note: Only Ethernet II frames are supported (no checksum support for SNAP packets).




Table 7.25. Supported Receive Checksum Capabilities

Packet Type	Hardware IP Checksum Calculation	Hardware TCP/UDP Checksum Calculation
IP header's protocol field contains a protocol # other than TCP or UDP.	Yes	No
IPv4 + TCP/UDP packets.	Yes	Yes
IPv6 + TCP/UDP packets.	No (N/A)	Yes
IPv4 packet has IP options (IP header is longer than 20 bytes).	Yes	Yes
IPv6 packet with next header options: <ul style="list-style-type: none"> • Hop-by-hop options. • Destinations options (without home address). • Destinations options (with home address). • Routing (with segment left 0). • Routing (with segment left > 0). • Fragment. 	No (N/A) No (N/A) No (N/A) No (N/A) No (N/A) No (N/A)	Yes Yes No Yes No No
Packet has TCP or UDP options.	Yes	Yes
IPv4 tunnels: <ul style="list-style-type: none"> • IPv4 packet in an IPv4 tunnel. • IPv6 packet in an IPv4 tunnel. 	Yes (outer IPv4 only) Yes (IPv4)	No No
IPv6 tunnels: <ul style="list-style-type: none"> • IPv4 packet in an IPv6 tunnel. • IPv6 packet in an IPv6 tunnel. 	No No (N/A)	No No
Packet is an IPv4 fragment.	Yes	UDP checksum assist
Packet is greater than 1522 bytes.	Yes	Yes
Packet has 802.3ac tag.	Yes	Yes
NVGRE: <ul style="list-style-type: none"> • Inner Ipv4 packet • Inner Ipv6 packet 	Yes (Inner and outer) Yes (outer)	Yes Yes
VXLAN: <ul style="list-style-type: none"> • Inner Ipv4 packet • Inner Ipv6 packet 	Yes (Inner and outer) Yes (outer)	Yes (inner only) ¹ Yes (inner only)

1. The outer UDP header of VXLAN packets do not use a checksum.

7.1.6.4 SCTP Receive Offload

If a receive packet is identified as SCTP, the Intel® Xeon® Processor D-1500 Product Family LAN controller checks the CRC32 checksum of this packet and identifies this packet as SCTP. Software is notified of the CRC check via the L4I and L4E bits in the *Extended Status* field and *Extended Error* field in the Rx descriptor. The detection of an SCTP packet is indicated via the *SCTP* bit in the *Packet Type* field of the Rx descriptor. SCTP CRC uses the CRC32c polynomial as follows (0x11EDC6F41):

$$X_{32}+X_{28}+X_{27}+X_{26}+X_{25}+X_{23}+X_{22}+X_{20}+X_{19}+X_{18}+X_{14}+X_{13}+X_{11}+X_{10}+X_9+X_8+X_6+X_0$$

The checker assumes the following SCTP packet format.



Table 7.26. SCTP Header

0 1 2 3 4 5 6 7	8 9 0 1 ¹ 2 3 4 5	6 7 8 9 0 1 ² 2 3	4 5 6 7 8 9 0 1 ³
Source Port		Destination Port	
Verification Tag			
CRC Checksum (CRC32c)			
Chunks 1..n			

7.1.6.5 Receive UDP Fragmentation Checksum

The Intel® Xeon® Processor D-1500 Product Family LAN controller might provide a receive fragmented UDP checksum offload for IPv4 non-tunneled packets. The *RXCSUM.PCSD* bit should be cleared and the *RXCSUM.IPPCSE* bit should be set to enable this mode.

The following table lists the outcome descriptor fields for the following incoming packets types.


Incoming Packet Type	Fragment Checksum	UDPV	UDPCS / L4CS
Non-IP packet	0	0	0
IPv6 packet	0	0	Depends on transport UDP: 1 / 1 TCP: 0 / 1
Non fragmented IPv4 packet			
Fragmented IPv4 with protocol = UDP, first fragment (UDP protocol present)	The unadjusted 1's complement checksum of the IP payload if the checksum in a packet header is different than zero and zero, otherwise.	1 if the UDP header checksum is valid (not 0)	0 / 0
Fragmented IPv4, when not first fragment	The unadjusted 1's complement checksum of the IP payload	1	0 / 0
Fragmented IPv4 with protocol = UDP, first fragment (UDP protocol present) within a VXLAN or NVGRE packet.	The unadjusted 1's complement checksum of the inner IP payload	1 if the UDP header checksum is valid (not 0)	1 / 0
Fragmented IPv4 when not first fragment within a VXLAN or NVGRE packet.	The unadjusted 1's complement checksum of the inner IP payload	0	1 / 0

Note: When the driver computes the 16-bit ones complement sum on the incoming packets of the UDP fragments, it should expect a value of 0xFFFF.

7.1.7 Receive Statistics

7.1.7.1 General rules

- All Statistics registers are cleared on read. In addition, they stick at 0xFF..F when the maximum value is reached.
- For the receive statistics it should be noted that a packet is indicated as received if it passes the device filters and is placed into the packet buffer memory. A packet does not have to be DMA'ed to host memory in order to be counted as received.
- Due to divergent paths between interrupt-generation and logging of relevant statistics counts, it might be possible to generate an interrupt to the system for a noteworthy event prior to the associated statistics count actually being increased. This is extremely unlikely due to expected delays associated with the system interrupt-collection and ISR delay, but might be an explanation



for interrupt statistics values that do not quite make sense. Hardware guarantees that any event noteworthy of inclusion in a statistics count is reflected in the appropriate count within 1 μ s; a small time-delay prior to reading the statistics might be required to avoid a potential mismatch between and interrupt and its cause.

- If RSC is enabled, statistics are collected before RSC is applied to the packets.
- All byte (octet) counters composed of 2 registers can be fetched by two consecutive 32-bit accesses while reading the Low 32-bit register first or a single 64-bit access.
- All receive statistic counters count the packets and bytes before coalescing by the RSC logic.
- All receive statistic counters in the Filter unit (listed below) might count packets that might be dropped by the packet buffer or receive DMA. Same comment is valid for the byte counters associated with these packet counters: PRC64; PRC127; PRC255; PRC511; PRC1023; PRC1522; BPRC; MPRC; GPRC; RXNFGPC; RUC; ROC

7.1.7.2 Receive Statistics Hierarchy

The following diagram shows the relationship between the packet flow and the different statistic counters.

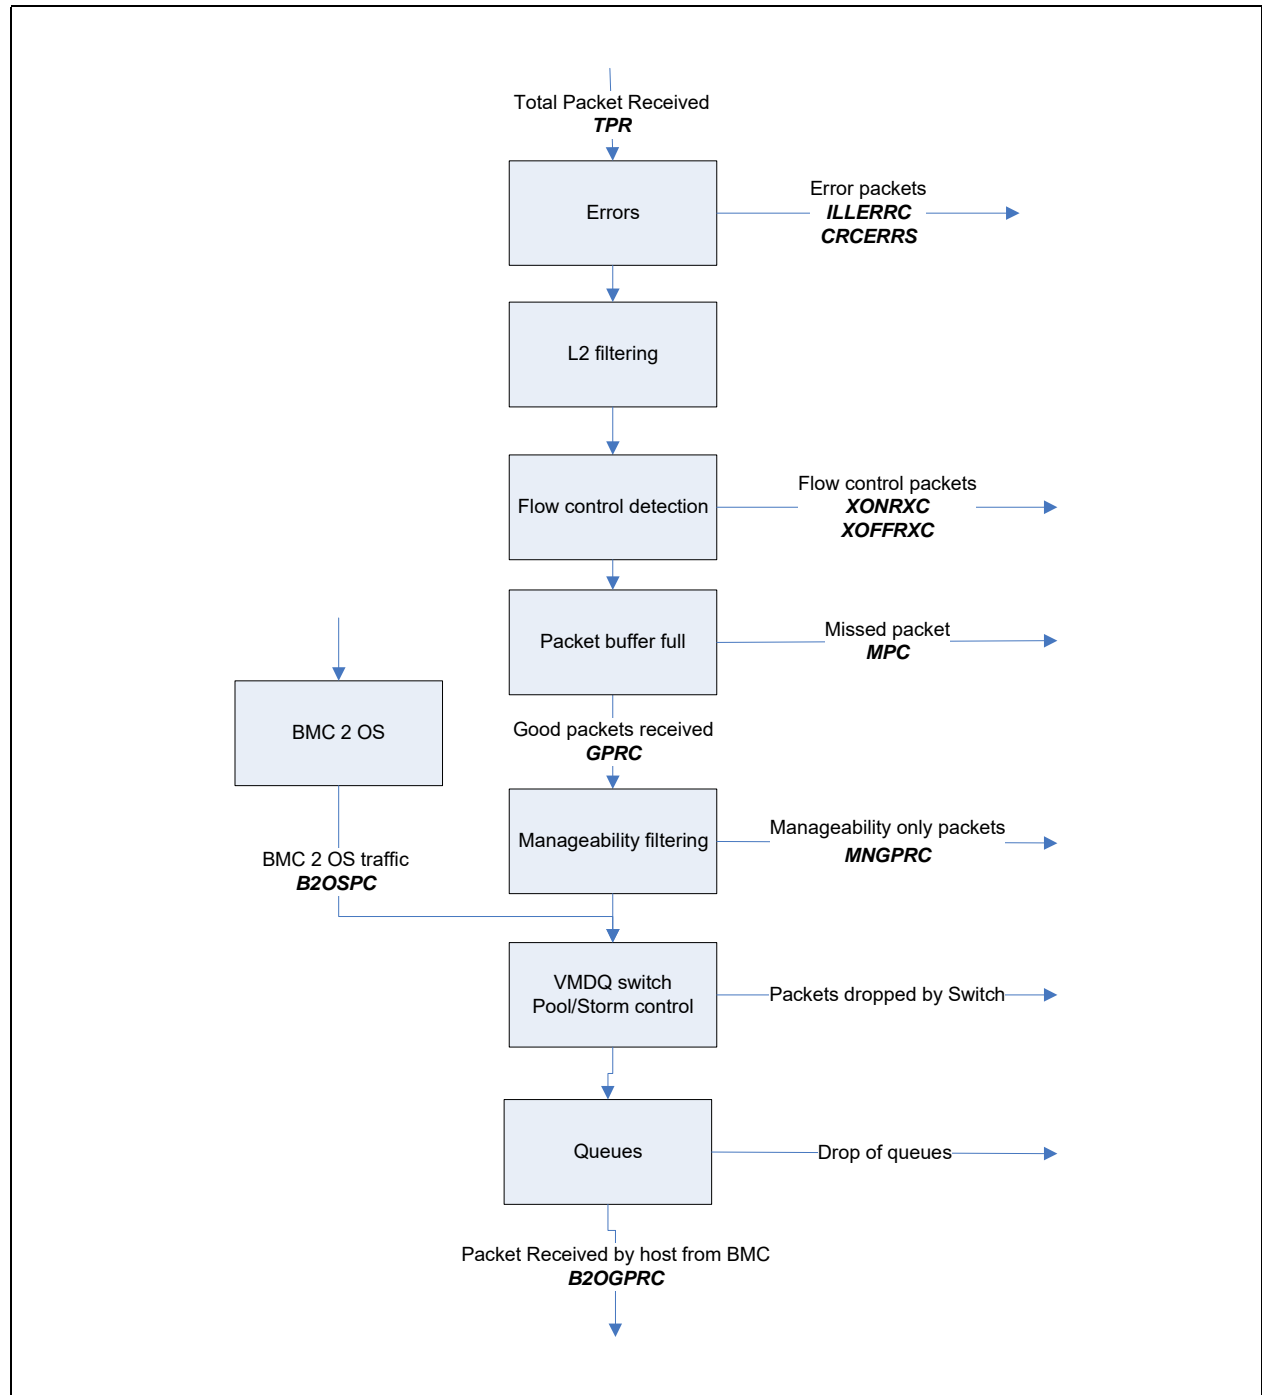



Figure 7.13. Receive Flow Statistics

7.2 Transmit Functionality

7.2.1 Packet Transmission

Transmit packets are made up of data buffers in host memory that are indicated to hardware by pointer and length pairs. These pointer and length pairs are named as transmit descriptors that are stored in host memory as well.

Software prepares memory structures for transmission by assembling a list of descriptors. It then indicates this list to hardware for updating the on-chip transmit tail pointer. Hardware transmits the packet only after it has completely fetched all packet data from host memory and deposited it into the on-chip transmit FIFO. This store and forward scheme enables hardware-based offloads such as TCP or UDP checksum computation, and many other ones detailed in this document while avoiding any potential PCIe under-runs.

7.2.1.1 Transmit Storage in System Memory

A packet (or multiple packets in transmit segmentation) can be composed of one or multiple buffers. Each buffer is indicated by a descriptor. Descriptors of a single packet are consecutive, while the first one points to the first buffer and the last one points to the last buffer (see [Figure 7.14](#)). The following rules must be kept:

- Address alignment of the data buffers can be on any byte boundary.
- Data buffers of any transmitted packet must include at least the 12 bytes of the source and destination Ethernet MAC addresses as well as the 2 bytes of the *Type/Len* field.
- A packet (or multiple packets in transmit segmentation) can span any number of buffers (and their descriptors) up to a limit of 40 minus *WTHRESH* minus 2 (see [Section 7.2.3.3](#) for Tx Ring details and [Section 7.2.3.5.1](#) for *WTHRESH* details). For best performance it is recommended to minimize the number of buffers as possible.

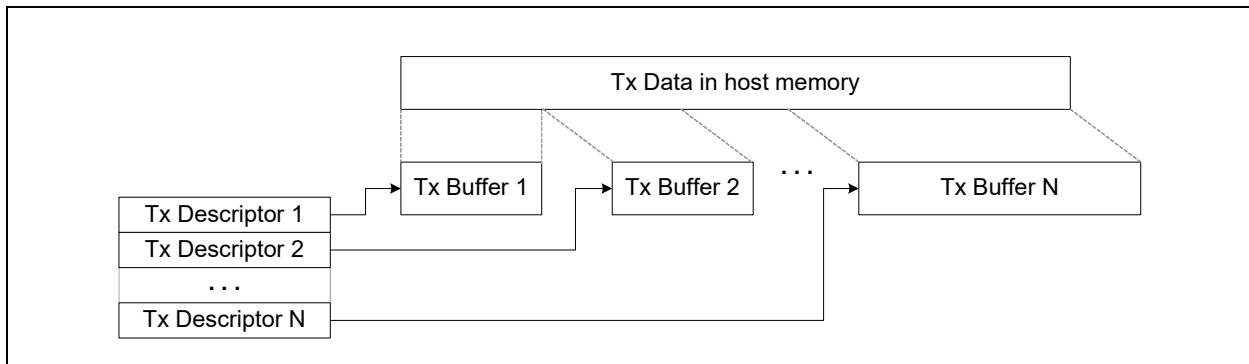



Figure 7.14. Tx Packet in Host Memory

7.2.1.2 Transmit Path in the Intel® Xeon® Processor D-1500 Product Family LAN Controller

The transmit path in the Intel® Xeon® Processor D-1500 Product Family LAN controller consists of the following stages:

- Descriptor plane
 - The Intel® Xeon® Processor D-1500 Product Family LAN controller maintains a set of 128 on-die descriptor queues. Each queue is associated with a single descriptor ring in system memory.



See [Section 7.2.3.3](#) for more details on the Tx descriptor rings. Each on-die descriptor queue contains up to 40 descriptors in order to achieve the desired performance.

- A fetch mechanism loads Tx descriptors from the descriptor rings in system memory to the respective descriptor queues in the Intel® Xeon® Processor D-1500 Product Family LAN controller. A descriptor fetch arbiter determines the order in which descriptors are fetched into the various on-die descriptor queues. See [Section 7.2.3.4](#) for more details on the fetch mechanism.
- An arbitration scheme determines the order in which descriptors are processed and requests are generated for data reads. These requests load packet data from system memory into a set of Tx packet buffers. The arbitration mechanism varies with configuration and is described in [Section 7.7](#).
- Once a packet has been fetched into a packet buffer, status is (optionally) written back into system memory. See [Section 7.2.3.5](#) for more details.
- Packet plane (data plane)
 - Packet data is stored in up to eight packet buffers. The number and size of packet buffers vary with the mode of operation and is described in [Section 7.2.1.2.2](#).
 - If more than a single packet buffer is enabled, an arbitration scheme determines the order in which packets are taken out of the packet buffers and sent to the MAC for transmission. The arbitration mechanism is described in [Section 7.7](#).

7.2.1.2.1 Tx Queues Assignment

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports a total of 128 queues per LAN port. Each Tx queue is associated with a packet buffer and the association varies with the operational mode. The following mechanisms impact the association of the Tx queues. These are described briefly in this section, and in full details in separate sections:

- Virtualization (Intel VT) - In a virtualized environment, DMA resources are shared between more than one software entity (operating system and/or device driver). This is done through allocation of transmit descriptor queues to virtual partitions (VMM, IOVM, VMs, or VFs). Allocation of queues to virtual partitions is done in sets of queues of the same size, called queue pools, or pools. A pool is associated with a single virtual partition. Different queues in a pool can be associated with different packet buffers. The PFVFTE register contains a bit per VF. When the bit is set to 0b, packet transmission from the VF is disabled. When set to 1b, packet transmission from the VF is enabled.
- Transmit fanout — A single descriptor queue might be enough for a given functionality. For example, in an Intel VT system, a single Tx queue can be allocated per VM. However, it is often the case that the data rate achieved through a single buffer is limited. This is especially true with 10 GbE, and traffic needs to be divided into several Tx queues in order to reach the desired data rate. Therefore, multiple queues might be provided for the same functionality.

[Table 7.27](#) lists the queuing schemes. Scheme selection is done via the MTQC register.

Table 7.27. Tx Queuing Schemes

Intel VT	Queues Allocation	Packet Buffers Allocation
No	A single set of 64 queues is assigned to a single packet buffer. Queues 64...127 should not be used.	A single packet buffer for all traffic
No	Eight TCs mode — allocation of 32-32-16-16-8-8-8-8 queues for TC0-TC1-...- TC7, respectively. Four TCs mode — allocation of 64-32-16-16 queues for TC0-TC1-...- TC3, respectively.	A separate packet buffer is allocated to each TC (total of four or eight).
Yes	32 pools x 4 queues, or 64 pools x 2 queues	A single packet buffer for all traffic.
Yes	16 pools x 8 TCs, or 32 pools x 4 TCs	A separate packet buffer is allocated to each TC (total of four or eight).



Note: Software can use any number of queues per each TC or per each pool within the allocated ranges previously described by disabling any unused queue.

Table 7.28. Queues, Pools and TCs Programming

Device Setting (MTQC)		Device Functionality	
VT_Ena	NUM_TC_OR_Q	Total Number of Tx Queues	TC & Pools
0	00	64	-
<> 0	00	Non valid configuration	
0	<> 00	Non valid configuration	
0	01	N/A	
0	10	128	TC0 — TC3
0	11	128	TC0 — TC7
1	01	128	64 pools
1	10	128	32 pools
1	11	N/A	
1	01	N/A	
1	10	128	TC0 — TC3 & 32 pools
1	11	128	TC0 — TC7 & 16 pools

Allocating descriptor queues to VFs uses a consistent indexing over the different Tx queuing schemes. The most significant bits of the queue index represent the VF index, and the least significant bits are either the TC index or are used by software to dispatch traffic according to a Transmit Side Scaling (TSS) algorithm — similar to RSS in the Rx path.

The Tx queue numbers associated with the TCs are listed in the following tables: [Table 7.29](#) and [Table 7.30](#).

Table 7.29. Tx Queues Indexing When Intel VT-on

Intel VT Mode	Allocation of Queue Index Bits According to						
	6	5	4	3	2	1	0
64 VFs + TSS	VF or pool (63..0)						TSS
32 VFs + TSS or 4 TCs	VF or pool (31..0)					TSS / TC	
16 VFs + 8 TCs	VF (15..0)				TC		

Table 7.30. Tx Queues Indexing When Intel VT-off

TC Mode	TCn	# of Qs	Queues Attached to TCn
4 TCs	TC0	64	0xxxxxx
	TC1	32	10xxxxx
	TC2	16	110xxxx
	TC3	16	111xxxx



Table 7.30. Tx Queues Indexing When Intel VT-off

TC Mode	TCn	# of Qs	Queues Attached to TCn
8 TCs	TC0	32	00xxxxx
	TC1	32	01xxxxx
	TC2	16	100xxxx
	TC3	16	101xxxx
	TC4	8	1100xxx
	TC5	8	1101xxx
	TC6	8	1110xxx
	TC7	8	1111xxx

Note: “x” refers to both 0 or 1, and is used by software to dispatch Tx flows via TSS algorithm.

7.2.1.2.2 Tx Packet Buffers

As previously described, the following modes exist for the Intel® Xeon® Processor D-1500 Product Family LAN controller packet buffers:

- A single 160 KB packet buffer that serves all Tx descriptor queues, leading to one single (or no) TC enabled, TC0
- Four 40 KB packet buffers, one per enabled TC, leading to four TCs, TC0 to TC3
- Eight 20 KB packet buffers, one per enabled TC, leading to eight TCs, TC0 to TC7

The size of the Tx packet buffer(s) is programmed via the TXPBSIZE registers, one register per TC. Null-sized packet buffer corresponds to a disabled TC.

Note: Setting the packet buffers’ size leads to a different partition of a shared internal memory and must be done during boot, prior to communicating, and followed by a software reset.

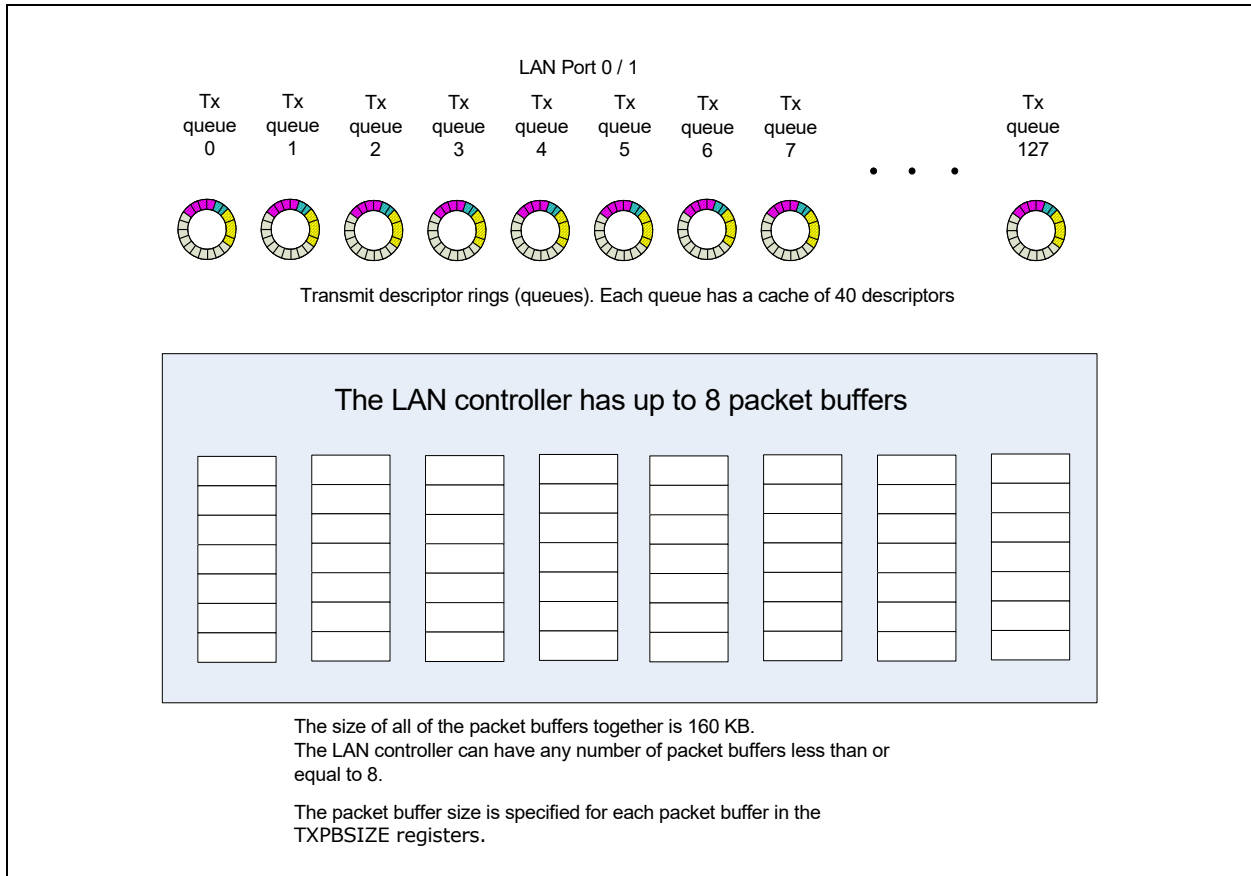


Figure 7.15. Tx Arbitration Schemes

7.2.1.2.3 Tx Arbitration Schemes

There are basically four Tx arbitration schemes, one per each combination of the Virtualization (Intel VT) enabled/disabled modes. They are configured via the *MTQC.MTOE* register field.

7.2.1.2.3.1 Intel VT-on

Virtualization is enabled, queues are allocated to the packet buffers in a fixed manner, the same number of queues per each TC.

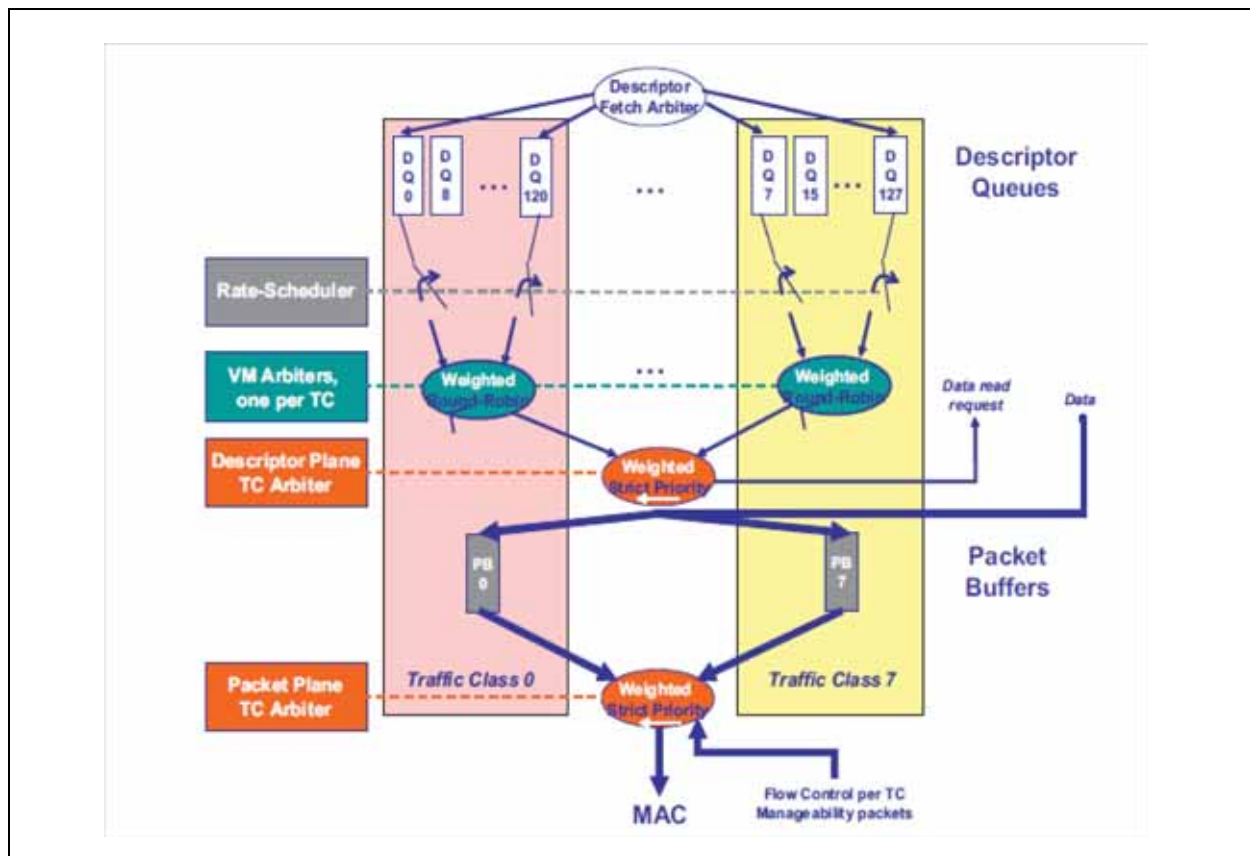


Figure 7.16. Transmit Architecture Intel VT-on — Eight TCs Mode

Descriptor Plane Arbiters and Schedulers:

- **Rate-Scheduler** — Once a frame has been fetched out from a rate-limited queue, the next time another frame could be fetched from that queue is regulated by the rate-scheduler. In the meantime, the queue is considered as if it was empty (such as switched-off) for the subsequent arbitration layers.
- **VM Weighted Round Robin Arbiter** — Descriptors are fetched out from queues attached to the same TC in a frame-by-frame weighted round-robin manner, while taking into account any limitation as previously described. Weights or credits allocated to each queue are configured via the RTTDT1C register. Bandwidth unused by one queue is reallocated to the other queues within the TC, proportionally to their relative bandwidth shares. TC bandwidth limitation is distributed across all the queues attached to the TC, proportionally to their relative bandwidth shares. It is assumed traffic is dispatched across the queues attached to a same TC in a straightforward manner, according to the VF to which it belongs.
- **TC Weighted Strict Priority Arbiter** — Descriptors are fetched out from queues attached to different TCs in a frame-by-frame weighted strict-priority manner. Weights or credits allocated to each TC are configured via RTTDT2C registers. Bandwidth unused by one TC is reallocated to the others, proportionally to their relative bandwidth shares. Link bandwidth limitation is distributed across all the TCs, proportionally to their relative bandwidth shares. It is assumed (each) driver dispatches traffic across the TCs according to the 802.1p *User Priority* field inserted by the operating system and according to a user priority-to-TC Tx mapping table.

Packet Plane Arbiters:

- **TC Weighted Strict Priority Arbiter** — Packets are fetched out from the different packet buffers in a frame-by-frame weighted strict-priority manner. Weights or credits allocated to each TC (such as to each packet buffer) are configured via RTTPT2C registers, with the same allocation done at the descriptor plane. Bandwidth unused by one TC and link bandwidth limitation is distributed over the TCs as in the descriptor plane.
- **Manageability** packets are inserted with strict priority over data packets from the same TC, with respect to the bandwidth allocated to the concerned TC. TCs that belong to manageability packets are controlled by MNGTXMAP.MAP.

Note: Replicating TC arbiters before and after the packet buffers is required to provide arbitration whether PCI bandwidth is smaller or greater than the link bandwidth, respectively.

7.2.1.2.3.2 Intel VT-off

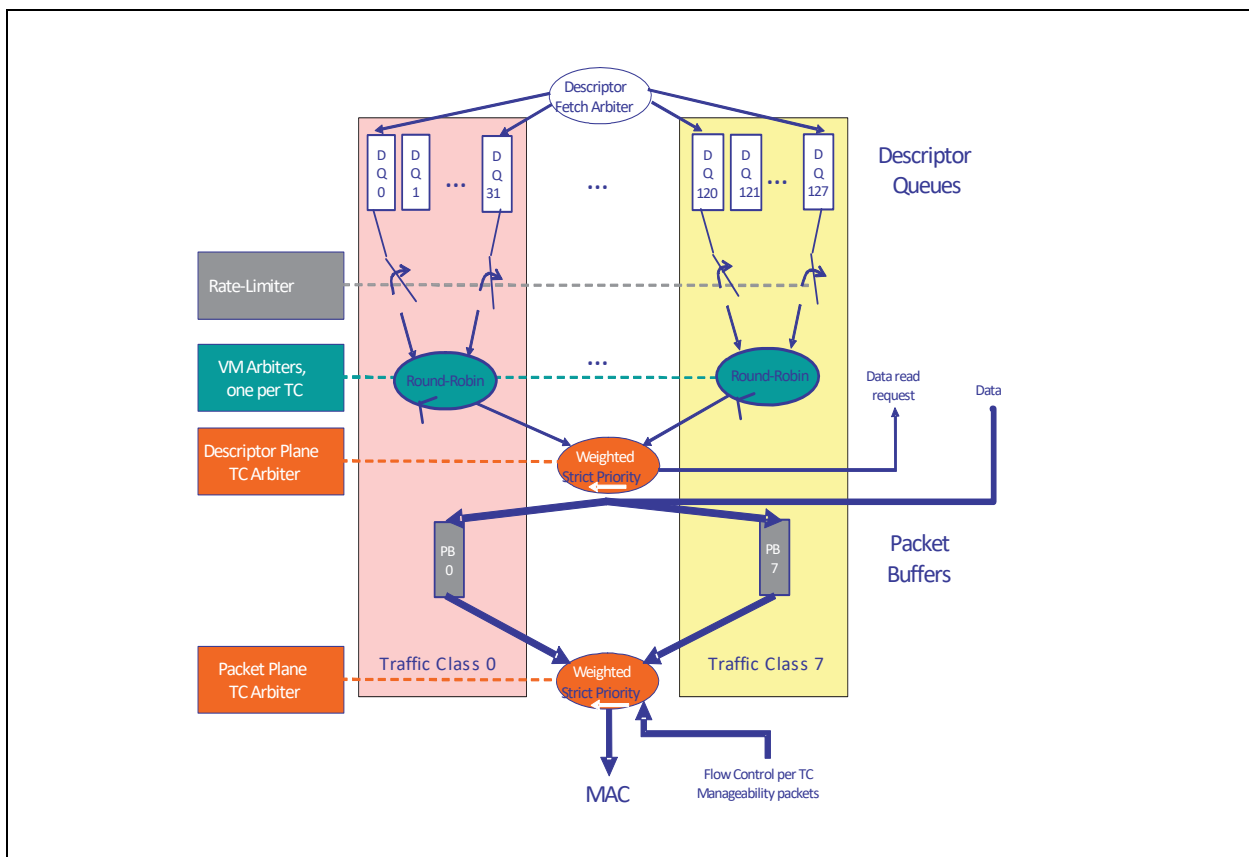


Figure 7.17. Transmit Architecture Intel VT-off — Eight TCs Mode

When virtualization is disabled, queues are allocated to the packet buffers in a fixed manner according to the number of TCs.

- The unique difference with the Intel VT-on arbitration scheme previously described is that the VM weighted round-robin arbiters are degenerated into simple frame-by-frame round-robin arbiters across the queues attached to a same TC. It is assumed driver dispatches traffic across the queues attached to a same TC according to hashing performed on MAC destination addresses. This is aimed to minimize crosstalk between rate-limited and non-rate-limited flows.



7.2.1.2.3.3 Intel VT-on

When virtualization is enabled, all the 128 queues are allocated to a single packet buffer PB(0). Queues are grouped into 32 or 64 pools of 4 or 2 queues, respectively. The number of queue pools corresponds to the number of VFs exposed. Queues are attached to pools according to consecutive indexes.

- For the 32 pools case, queues 0, 1, 2, 3 are attached to VF0, queues 4, 5, 6, 7 are attached to VF1, and so forth up to VF31.
- For the 64 pools case, queues 0 and 1 are attached to VF0, queues 2 and 3 are attached to VF1, and so forth up to VF63.

Descriptor Plane Arbiters:

- **Descriptor Queues Round Robin Arbiter** — Descriptors are fetched out from the internal descriptor queues attached to the same pool in a frame-by-frame round-robin manner. It is assumed driver dispatches traffic across the queues of a same pool according to some Transmit Side Scaling (TSS) algorithm similarly to what is done by hardware in the Rx path with RSS.
- **VM Weighted Round Robin Arbiter** — Descriptors are fetched out from queues attached to different pools in a frame-by-frame weighted round-robin manner. Weights or credits allocated to a pool are those allocated for the lowest queue of the pool via the RTTDT1C register. Bandwidth unused by one pool is reallocated to the others proportionally to their relative bandwidth shares. Link bandwidth limitation is distributed across all the pools, proportionally to their relative bandwidth shares.

Packet Plane Arbiter:

- **Manageability** packets are inserted with strict priority over data packets.

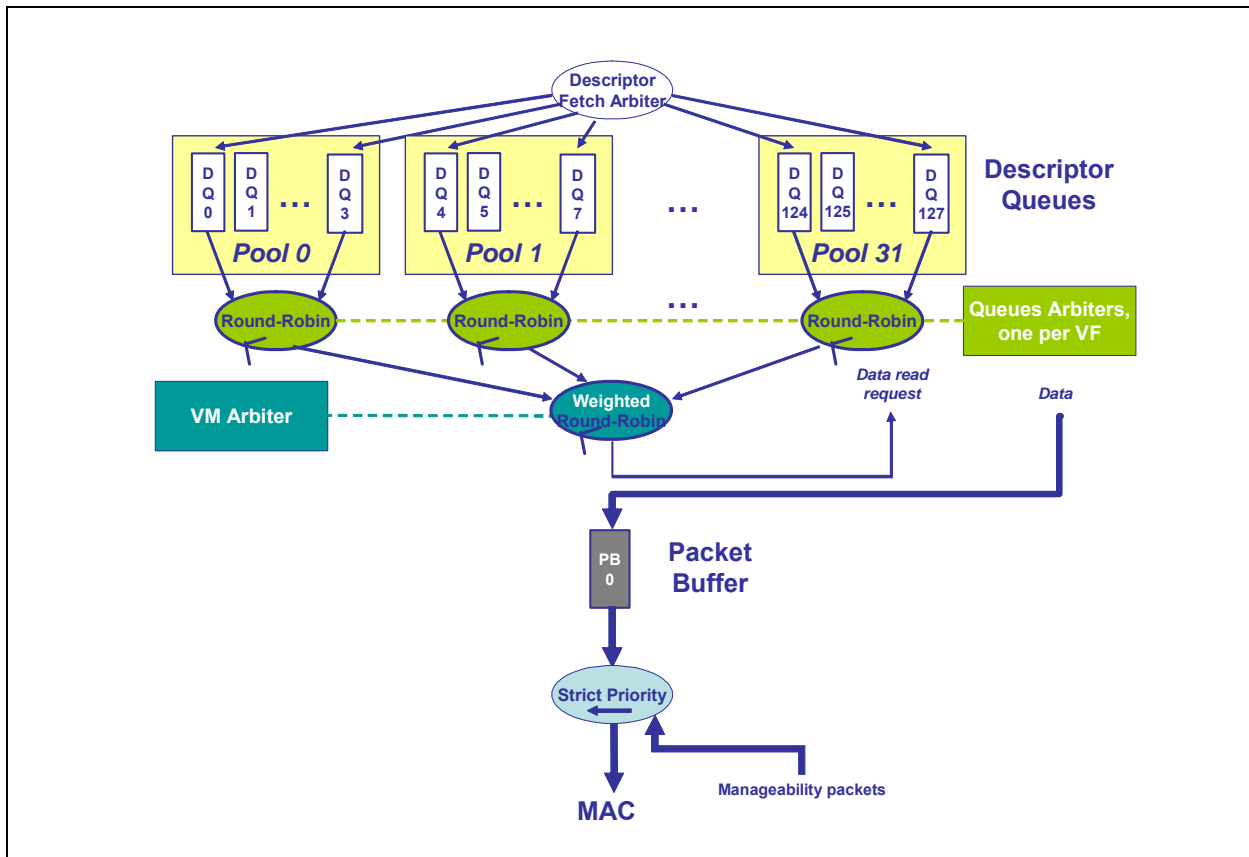



Figure 7.18. Transmit Architecture Intel VT-on — 32 VFs

7.2.1.2.3.4 Intel VT-off

When virtualization is disabled, a single set of up to 64 queues is allocated to a single packet buffer PB(0).



Descriptor Plane Arbiter:

- **Descriptor Queues Round Robin Arbiter** — Descriptors are fetched out from the internal descriptor queues in a frame-by-frame round-robin manner. It is assumed driver dispatches traffic across the queues according to some TSS algorithm similarly to what is done by hardware in the Rx path with RSS.

Packet Plane Arbiter:

- **Manageability** packets are inserted with strict priority over data packets.

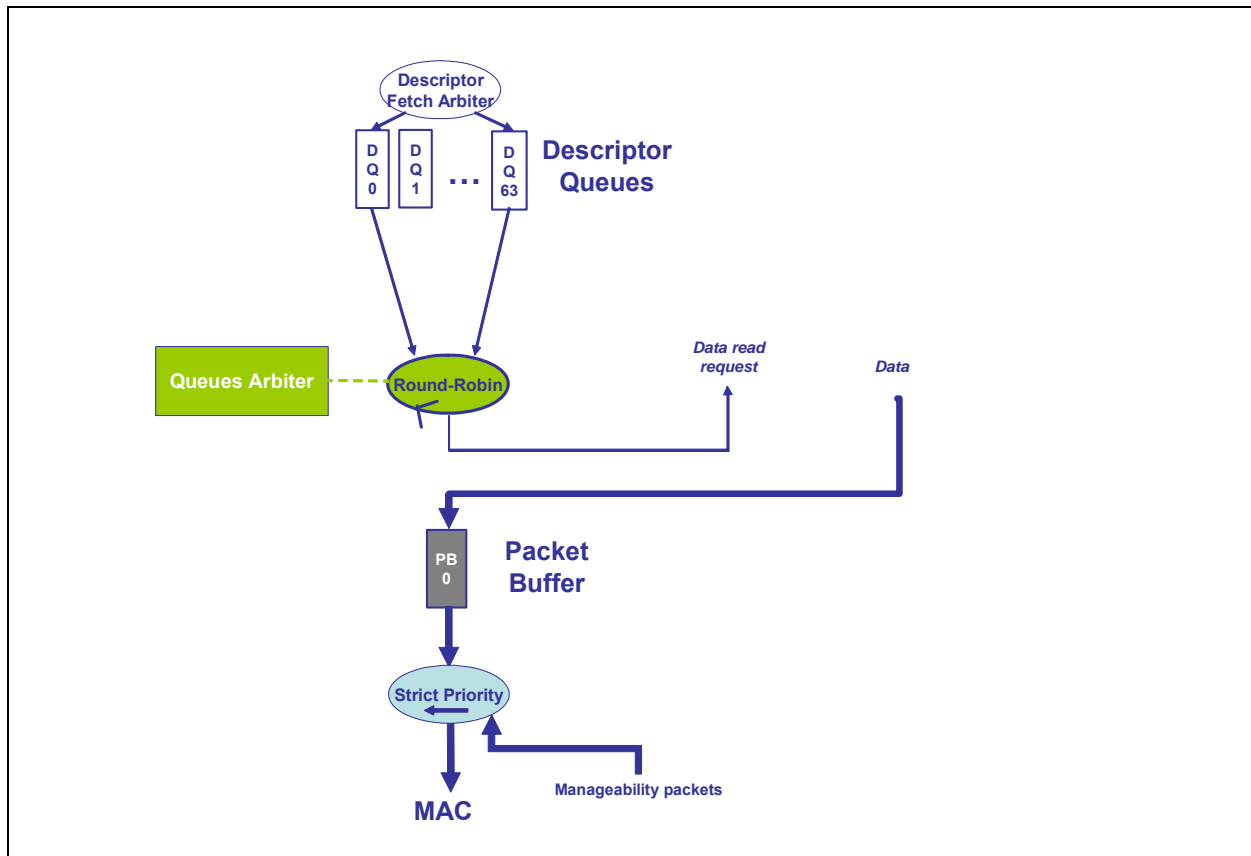


Figure 7.19. Transmit Architecture Intel VT-off

7.2.2 Transmit Contexts

The Intel® Xeon® Processor D-1500 Product Family LAN controller provides hardware checksum offload and TCP segmentation facilities. These features enable TCP and UDP packet types to be handled more efficiently by performing additional work in hardware, thus reducing the software overhead associated with preparing these packets for transmission. Part of the parameters used to control these features are handled through contexts.

A context refers to a set of parameters providing a particular offload functionality. These parameters are loaded by unique descriptors named transmit context descriptors. A transmit context descriptor is identified by the *DTYP* field (described later in this section) equals to 0x2.



The Intel® Xeon® Processor D-1500 Product Family LAN controller supports two contexts for each of its 128 transmit queues. The *IDX* bit contains an index to one of these two contexts. Each advanced data descriptor that uses any of the advanced offloading features must refer to a context by the *IDX* field.

Contexts can be initialized with a transmit context descriptor and then used for a series of related transmit data descriptors. Software can use these contexts as long lived ones, while one of the two contexts is used for checksum offload and the other one for transmit segmentation detailed in the following sections. The contexts should be modified when new offload parameters are required.

7.2.3 Transmit Descriptors

7.2.3.1 Introduction

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports legacy descriptors and advanced descriptors.

Legacy descriptors are intended to support legacy drivers, in order to enable fast platform power up and to facilitate debug. The legacy descriptors are recognized as such based on *DEXT* bit (see the sections that follow). Legacy descriptors are not supported together with virtualization, MACsec. This mode is recognized by a dedicated enable bit each.

In addition, the Intel® Xeon® Processor D-1500 Product Family LAN controller supports two types of advanced transmit descriptors:

1. Advanced transmit context descriptor, DTYP = 0010b
2. Advanced transmit data descriptor, DTYP = 0011b

Note: DTYP = 0000b and 0001b are reserved values.

The transmit data descriptor (both legacy and advanced) points to a block of packet data to be transmitted. The advanced transmit context descriptor does not point to packet data. It contains control/context information that is loaded into on-chip registers that affect the processing of packets for transmission. The following sections describe the descriptor formats.

7.2.3.2 Transmit Descriptors Formats

7.2.3.2.1 Notations

This section defines the structure of descriptors that contain fields carried over the network. At the moment, the only relevant field is the *VLAN Tag* field.

The rule for VLAN tag is to use network ordering (also called big endian). It appears in the following manner in the descriptor:

Table 7.31. VLAN Tag

Byte address N + 1 -> first byte on the wire Bit 7 —		Byte address N -> second byte on the wire Bit 7 -> last on the wire — Bit 0	
PRI (3 bits)	DEI	VID (4 bits)	VID (8 bits)

7.2.3.2.2 Legacy Transmit Descriptor Format



To select legacy mode operation, bit 29 (*TDESC.DEXT*) should be set to 0b. In this case, the descriptor format is defined as listed in Table 7.32. Address and length must be supplied by software on all descriptors. Bits in the command byte are optional, as are the *CSO*, and *CSS* fields.

Table 7.32. Transmit Descriptor (TDESC) Layout — Legacy Mode

	63	48	47	40	39	36	35	32	31	24	23	16	15	0
0	Buffer Address [63:0]													
8	VLAN		CSS		Rsvd		STA		CMD		CSO		Length	

Table 7.33. Transmit Descriptor Write-Back Format — Legacy Mode

	63	48	47	40	39	36	35	32	31	24	23	16	15	0
0	Reserved							Reserved						
8	VLAN		CSS		Rsvd		STA		CMD		CSO		Length	

Buffer Address (64) and Length (16)

The buffer address is a byte-aligned address. Length (*TDESC.LENGTH*) specifies the length in bytes to be fetched from the buffer address provided. The maximum length associated with a single descriptor is 15.5 KB while the total frame size must meet the maximum supported frame size. There is no limitation for the minimum buffer size.

Note: Descriptors with zero length (null descriptors) transfer no data. Null descriptors might appear only between packets and must have their *EOP* bits set.

Checksum Offset and Start — CSO (8) and CSS (8)

A *Checksum Offset* (*TDESC.CSO*) field indicates where, relative to the start of the packet, to insert a TCP checksum if this mode is enabled. A *Checksum Start* (*TDESC.CSS*) field indicates where to begin computing the checksum. Note that *CSO* and *CSS* are meaningful only in the first descriptor of a packet.

Both *CSO* and *CSS* are in units of bytes. These must both be in the range of data provided to the device in the descriptor. This means for short packets that are padded by software, *CSO* and *CSS* must be in the range of the un-padded data length, not the eventual padded length (64 bytes).

CSO must be set to the location of TCP or UDP checksum in the packet. *CSS* must be set to the beginning of the IP header or the L4 (TCP/UDP) header. Checksum calculation is not done if *CSO* or *CSS* are out of range. This occurs if (*CSS* > length) OR (*CSO* > length - 1).

For the 802.1Q header, the offset values depend on the VLAN insertion enable bit — the *VLE* bit. If they are not set (VLAN tagging included in the packet buffers), the offset values should include the VLAN tagging. If these bits are set (VLAN tagging is taken from the packet descriptor), the offset values should exclude the VLAN tagging.

Hardware does not add the 802.1q Ethertype or the VLAN field following the 802.1Q Ethertype to the checksum. So for VLAN packets, software can compute the values to back out only on the encapsulated packet rather than on the added fields.

Note: UDP checksum calculation is not supported by the legacy descriptor because the legacy descriptor does not support the translation of a checksum result of 0x0000 to 0xFFFF needed to differentiate between an UDP packet with a checksum of zero and an UDP packet without checksum.



Because the *CSO* field is eight bits wide, it puts a limit on the location of the checksum to 255 bytes from the beginning of the packet.

Note: *CSO* must be larger than *CSS*.

Software must compute an offsetting entry to back out the bytes of the header that should not be included in the TCP checksum and store it in the position where the hardware computed checksum is to be inserted.

Hardware adds the checksum at the byte offset indicated by the *CSO* field. Checksum calculations are for the entire packet starting at the byte indicated by the *CSS* field. The byte offset is counted from the first byte of the packet fetched from host memory.

Command Byte — CMD (8)

The CMD byte stores the applicable command and has the fields listed in [Table 7.34](#).

Table 7.34. Transmit Command (TDESC.CMD) Layout

7	6	5	4	3	2	1	0
RSV	VLE	DEXT	RSV	RS	IC	IFCS	EOP

- *RSV (bit 7)* — Reserved
- *VLE (bit 6)* — VLAN Packet Enable
When set to 1b, VLE indicates that the packet is a VLAN packet and hardware adds the VLAN header to the Tx packet. The VLAN Ethertype is taken from DMATXCTL.VT and the 802.1q VLAN tag is taken from the *VLAN* field in the Tx descriptor. See [Section 7.4.5](#) for details about double VLAN.

Table 7.35. VLAN Tag Insertion Decision Table for VLAN Mode Enabled

VLE	Action
0	Send generic Ethernet packet.
1	Send 802.1Q packet; the <i>Ethernet Type</i> field comes from the <i>VET</i> field of the VLNCTRL register and the VLAN data comes from the <i>VLAN</i> field of the TX descriptor.

Note: This table is relevant only if PFVMMIR.VLANA = 00b (use descriptor command) for the queue.

- *DEXT (bit 5)* — Descriptor extension (zero for legacy mode)
- *RSV (bit 4)* — Reserved
- *RS (bit 3)* — Report Status - RS signals hardware to report the DMA completion status indication as well as triggering ITR. Hardware indicates a DMA completion by setting the *DD* bit in the transmit descriptor when *TDWBAL[n].Head_WB_En* = 0b or by Head Write-back if *Head_WB_En* = 1b (see [Section 7.2.3.5.2](#)). The *RS* bit is permitted only on descriptors that has the *EOP* bit set (last descriptor of a packet).

Note: Software should not set the *RS* bit when *TXDCTL.WTHRESH* is greater than zero. Instead, the hardware reports the DMA completion according to the *WTHRESH* rules (explained in [Section 7.2.3.5.1](#)). This note is relevant only for descriptor write back while in head write-back mode. *WTHRESH* must also be set to zero.

When *TXDCTL.WTHRESH* = zero, software must set the *RS* bit on the last descriptor of every packet.

There are some exceptions for descriptor completion indication in head write-back mode. For more details see [Section 7.2.3.5.2](#).



- **IC (bit 2)** — Insert Checksum - Hardware inserts a checksum at the offset indicated by the *CSO* field if the *Insert Checksum* bit (IC) is set.
- **IFCS (bit 1)** — Insert FCS - When set, hardware appends the MAC FCS at the end of the packet. When cleared, software should calculate the FCS for proper CRC check. There are several cases in which software must set IFCS as follows:
 - Transmitting a short packet while padding is enabled by the *HLREGO.TXPADEN* bit.
 - Checksum offload is enabled by the *IC* bit in the *TDESC.CMD*.
 - VLAN header insertion enabled by the *VLE* bit in the *TDESC.CMD* or by the *PFVMVIR.VLANA* fields.
 - VN-tag or E-tag insertion enabled by the *PFVMVIR.TAGA* fields.
 - TSO or TCP/IP checksum offload using a context descriptor.
 - MACsec offload is requested.

Note that TSO, MACsec offload are relevant only to advanced Tx descriptors.

- **EOP (bit 0)** — End of Packet - A packet can be composed of multiple buffers (each of them indicated by its own descriptor). When EOP is set, it indicates the last descriptor making up the packet.

Note: *VLE*, *IFCS*, and *IC* fields should be set in the first descriptor of a packet. The *RS* bit can be set only on the last descriptor of a packet. The *DEXT* bit must be set to zero for all descriptors. The *EOF* bit is meaningful in all descriptors.

Transmitted.Status — STA (4)

DD (bit 0) — Descriptor Done Status

This bit provides a status indication that the DMA of the buffer has completed. Software might re-use descriptors with the *DD* bit set and any other descriptors processed by the hardware before this one. The other bits in the *STA* field are reserved.

Rsvd — Reserved (4)

VLAN (16)

The *VLAN* field is used to provide the 802.1q/802.1ac tagging information. The *VLAN* field is qualified on the first descriptor of each packet when the *VLE* bit is set to 1b. The *VLAN* field is provided in network order and is meaningful in the first descriptor of a packet. See [Section 7.2.3.2.1](#) for more details.

Table 7.36. VLAN Field (TDESC.VLAN) Layout

15	13	12	11	0
PRI		DEI	VLAN	

7.2.3.2.3 Advanced Transmit Context Descriptor

Table 7.37. Transmit Context Descriptor (TDESC) Layout — (Type = 0010)

63	56	55	48	47	42	41	32	31	16	15	9	8	0
----	----	----	----	----	----	----	----	----	----	----	---	---	---




Table 7.37. Transmit Context Descriptor (TDESC) Layout — (Type = 0010)

0	TUNNELLEN	OUTERIPLEN	RSV					VLAN				MACLEN	IPLN/HEADLEN				
8	MSS		L4LEN	RSV	IDX	Reserved must be 0x3F		DEXT	RSV	DTYP	TUCMD						
	63	48	47	40	39 37	36	35	30	29	28	24	23	20	19	9	8	0



IPLLEN/HEADLEN (9)

- *IPLLEN* — for IP packets:
This field holds the value of the IP header length for the IP checksum offload feature. If an offload is requested, *IPLLEN* must be greater than or equal to 20, and less than or equal to 511. For IP tunnel packets (IPv4-IPv6) *IPLLEN* must be defined as the length of the two IP headers. The hardware is able to offload the L4 checksum calculation while software should provide the IPv4 checksum.

HEADLEN does not include the MACsec header if it exists. [MACLEN \(7\)](#)

- For non FCoE packets:
This field indicates the length of the MAC header. When an offload is requested, one of the TSE bits (in the advanced transmit data descriptor) or *IXSM* bit or *TXSM* bit are set, *MACLEN* must be larger than or equal to 14, and less than or equal to 127. This field should include only the part of the L2 header supplied by the driver and not the parts added by hardware. The following table lists the value of *MACLEN* in the different cases.

	Regular VLAN	Extended Tag	MACLEN
	By hardware or none	None	14
	By hardware or none	VLAN	18
	By hardware or none	E-tag	22
	By hardware or none	VN-tag	20
	By software	None	18
	By software	VLAN	22
	By software	E-tag	26
	By software	VN-tag	24

VLAN (16)

This field contains the 802.1Q VLAN tag to be inserted in the packet during transmission. This VLAN tag is inserted when a packet using this context has its DCMD.VLE bit is set. This field should include the entire 16-bit VLAN field including DEI and priority fields as listed in [Table 7.36](#).

Note that the *VLAN* field is provided in network order. See [Section 7.2.3.2.1](#).

Ipsec SA IDX (10) – IPsec SA Index. If an IPsec offload is requested for the packet (*IPSEC* bit is set in the advanced Tx data descriptor), indicates the index in the SA table where the IPsec key and SALT are stored for that flow.


- *SOF (bit 2)* — Start of frame delimiter index.
- *ORIS (bit 5)* — Orientation relative to the first frame in an FC sequence.
- *PARINC (bit 3)* — When this bit is set, hardware relates to the *PARAM* field in the FC header as relative offset. In this case, hardware increments the *PARAM* field in TSO by an MSS value on each transmitted packet of the TSO. Software should set the *PARINC* bit when it sets the *Relative Offset Present* bit in the *F_CTL*.

OUTERIPLLEN (8) - This field holds the value of the outer IP header length.

TUNNELLEN (8) - This field holds the length value of the tunnel headers including the inner L2 header.

TUCMD (11)

- *RSV (bit 10-7)* — Reserved



- *Reserved (bit 5:4)*
- *L4T (bit 3:2)* — L4 Packet TYPE (00: UDP; 01: TCP; 10: SCTP; 11: RSV)
- *IPV4 (bit 1)* — IP Packet Type: When 1b, IPv4; when 0b, IPv6

(bit 0) DTYP (4)

This field is always 0010b for this type of descriptor.

RSV(1)

Reserved

DEXT (1) — Descriptor extension (one for advanced mode)

IDX (1)

The context descriptor is posted to a context table in hardware. There are two context tables per queue. The IDX is the index of the context tables.

Note: Because the Intel® Xeon® Processor D-1500 Product Family LAN controller supports only two context descriptors per queue, the two MS bits are reserved and should be set to 0b.

RSV(1)

L4LEN(8)

This field holds the layer 4 header length. If TSE is set, this field must be greater than or equal to 8 and less than or equal to 64. Otherwise, this field is ignored. Note that for UDP segmentation the L4 header size equals 8 and for TCP segmentation (with no TCP options) it equals 20.

MSS (16)

This field controls the maximum segment size. This specifies the maximum protocol payload segment sent per frame, not including any header. MSS is ignored when DCMD.TSE is not set.

TCP / UDP Segmentation

The total length of each frame (or segment) excluding Ethernet CRC as follows. Note that the last packet of a TCP segmentation might be shorter.

$$\text{MACLEN} + 4(\text{if VLE set}) + [4, 8, 14, \text{ or } 16] + \text{OUTERIPLEN} + \text{TUNNELLEN} + \text{IPLLEN} + \text{L4LEN} + \text{MSS}$$

Note: The headers lengths must meet the following: $\text{MACLEN} + \text{IPLLEN} + \text{L4LEN} \leq 512$

The total length of each frame (or segment) excluding Ethernet CRC equals to:

$$\text{MACLEN} + 4(\text{if VLE set}) + [4, 8, 14, \text{ or } 16] + 8 (\text{FC CRC} + \text{EOF})$$

The context descriptor requires valid data only in the fields used by the specific offload options. The following table lists the required valid fields according to the different offload options.



Table 7.38. Valid Fields by Offload Option

	Context Fields ->	Reserved	Reserved	VLAN	MACLEN	OUTERIPLEN/TUNNELLEN	IPLEN/HEADLEN	L4LEN		IPV4	L4T	Encryption	IPSECTYPE	SAIDX	ESP_LEN	MSS	QCNTLEN	CC (data descriptor)
Required Offload	VLAN Insertion			yes													yes	
	IPv4 XSUM				yes		yes			1							yes	0
	L4 XSUM				yes		yes				yes						yes	0
	TCP/UDP Seg				yes		yes	yes		yes	yes					yes	yes	0
	Tunnel XSUM				yes	yes	yes				yes							1
	Tunnel Seg				yes	yes	yes	yes		yes	yes					yes	yes	1
	Tx Switch				yes		yes	yes		yes	yes	n/a	n/a	n/a	n/a		yes	1

Note: All fields that are not used in the context descriptor must be set to zero.

7.2.3.2.4 Advanced Transmit Data Descriptor

Table 7.39. Advanced Transmit Data Descriptor Read Format

0	Address[63:0]											
8	PAYLEN	POPTS	CC	IDX	STA	DCMD	DTYP	MAC	TUNNEL	DTALEN		
63	46	45	40	39	38 36	35 32	31 24	23 20	19 18	17 16	15 0	

Table 7.40. Advanced Transmit Data Descriptor Write-back Format

0	RSV											
8	RSV	STA										RSV
	63 36	35								32		31 0

Address (64)

This field holds the physical address of a data buffer in host memory, which contains a portion of a transmit packet. This field is meaningful in all descriptors.



DTALEN (16)

This field holds the length in bytes of data buffer at the address pointed to by this specific descriptor. This field is meaningful in all descriptors. The maximum length is 16 KB with no limitations on the minimum size. Refer to the comment on descriptors with zero length described in the sections that follow.

TUNNEL (2)

- *Tunnel Type (bit 0)*:
0 = VXLAN
1 = NVGRE
- *OUTERIPCS (bit 1)* — Outer IP checksum enable. Indicates this is a tunnel packet and that the *OUTERPLEN/TUNNELLEN* parameters in the context descriptor are valid.

MAC (2) — see the following. This field is meaningful on the first descriptor of the packet(s).

- *ILSec (bit 0)* — Apply MACsec on packet. When set, hardware includes the MACsec header (SecTAG) and MACsec header digest (signature). The MACsec processing is defined by the Enable Tx MACsec field in the *LSECTXCTRL* register. The *ILSec* bit in the packet descriptor should not be set if MACsec processing is not enabled by the Enable Tx MACsec field. If the *ILSec* bit is set erroneously while the *Enable Tx MACsec* field is set to 00b, then the packet is dropped.
- *1588 (bit 1)* — IEEE1588 time stamp packet.

DTYP (4)

0011b for advanced data descriptor. DTYP should be valid in all descriptors of the packet(s).

DCMD (8) — see the following:

- *TSE (bit 7)* — Transmit Segmentation Enable - This bit indicates TCP segmentation request. When *TSE* is set in the first descriptor of a TCP packet, hardware must use the corresponding context descriptor in order to perform segmentation.

Note: It is recommended that *HLREG0.TXPADEN* be enabled when *TSE* is used since the last frame can be shorter than 60 bytes — resulting in a bad frame.

- *VLE (bit 6)* — VLAN Packet Enable - This bit indicates that the packet is a VLAN packet (hardware must add the VLAN Ethertype and an 802.1q VLAN tag to the packet).
- *DEXT (bit 5)* — Descriptor Extension - This bit must be one to indicate advanced descriptor format (as opposed to legacy).
- *Rsv (bit 4)* — Reserved
- *RS (bit 3)* — Report Status: See the description in the legacy transmit descriptor in [Section 7.2.3.2.2](#).
- *Rsv (bit 2)* — Reserved
- *IFCS (bit 1)* — Insert FCS - When this bit is set, the hardware appends the MAC FCS at the end of the packet. When cleared, software should calculate the FCS for proper CRC check. There are several cases in which software must set *IFCS* as follows:
 - Transmitting a short packet while padding is enabled by the *HLREG0.TXPADEN* bit.
 - Checksum offload is enabled by the either *IC*, *TXSM* or *IXSM* bits in the *TDESC.DCMD*.
 - VLAN header insertion enabled by the *VLE* bit in the *TDESC.DCMD*.
 - TCP segmentation offload enabled by the *TSE* bit in the *TDESC.DCMD*.
- *EOP (bit 0)* — End of Packet - A packet might be composed of multiple buffers (each of them is indicated by its own descriptor). When *EOP* is set, it indicates the last descriptor making up the



packet. In transmit segmentation (explained later on in this section) the *EOP* flag indicates the last descriptor of the last packet of the segmented transmission.

Note: *TSE*, *VLE* and *IFCS* fields should be set in the first descriptor of the packet(s). The *RS* bit can be set only on the last descriptor of the packet. The *EOP* bit is valid in all descriptors. The *DEXT* bit must be set to 1b for all descriptors.

Descriptors with zero length, transfer no data. If the *RS* bit in the command byte is set, then the *DD* field in the status word is not written when hardware processes them.

STA (4)

- *Rsv (bit 3:1)* — Reserved
- *DD (bit 0)* — Descriptor Done: The *DD* bit provides a status indication that the DMA of the buffer has completed. Software might re-use descriptors with the *DD* bit set, and any other descriptors processed by hardware before this one. In TSO, the buffers that include the TSO header are used multiple times during transmission and special considerations should be made as described in Section 7.2.4.2.2.

IDX (3)

This field holds the index into the hardware context table to indicate which of the two per-queue contexts should be used for this request. If no offload is required and the *CC* bit is cleared, this field is not relevant and no context needs to be initiated before the packet is sent. See Table 7.38 for details of which packets requires a context reference. This field is relevant only on the first descriptor of the packet(s).

CC (1)

Check Context bit — When set, a Tx context descriptor indicated by *IDX* index should be used for this packet(s). The *CC* bit should be set in the following cases:

1. Non-zero *QCNTLEN* field is required (defined in the context descriptor).
2. Tx switching is enabled (including anti spoof checks).

POPTS (6)

- *Rsv (bit 5:2)* — Reserved
- *TXSM (bit 1)* — Insert TCP/UDP Checksum: When set to 1b, the L4 checksum must be inserted. In this case, *TUCMD.LP4* indicates whether the checksum is TCP or UDP or SCTP. When *DCMD.TSE* is set, *TXSM* must be set to 1b. If this bit is set, the packet should at least contain a TCP header.
- *IXSM (bit 0)* — Insert IP Checksum: This field indicates that IP checksum must be inserted. In IPv6 mode, it must be reset to 0b. If *DCMD.TSE* and *TUCMD.IPV4* are set, *IXSM* must be set to 1b. If this bit is set, the packet should at least contain an IP header.

This field is relevant only on the first descriptor of the packet(s).

PAYLEN (18)

PAYLEN indicates the size (in byte units) of the data buffer(s) in host memory for transmission.

In a single-send packet, *PAYLEN* defines the entire packet size fetched from host memory. It does not include the fields that hardware adds such as: optional VLAN tagging and Ethernet CRC or Ethernet padding. When MACsec offload is enabled, the *PAYLEN* field does not include the MACsec encapsulation.

In TSO (regardless if it is transmitted on a single or multiple packets), the *PAYLEN* defines the protocol payload size fetched from host memory.

- In TCP or UDP segmentation offload, *PAYLEN* defines the TCP/UDP payload size.

This field is relevant only on the first descriptor of the packet(s). The minimum transmitted packet size excluding VLAN padding and CRC bytes is 17 and the PAYLEN size should meet this limitation. On a single-packet send, the maximum size of the PAYLEN is dictated by the maximum allowed packet size which is 15.5 KB. On TSO, the maximum PAYLEN can be up to $2^{18}-1$.

Note: When a packet spreads over multiple descriptors, all of the descriptor fields are valid only on the first descriptor of the packet, except for *RS* and *EOP* bits, which are set on the last descriptor of the packet.

7.2.3.3 Transmit Descriptor Ring

The transmit descriptor ring structure (shown in Figure 7.20) uses a contiguous memory space. A set of four registers (described later in this section) maintain the transmit descriptor ring in the host memory. Hardware maintains internal circular queues of 64 descriptors per queue to hold the descriptors that were fetched from the software ring.

Descriptors handled to hardware should not be manipulated by software until hardware completes its processing. It is indicated by advancing the head pointer beyond these descriptors.

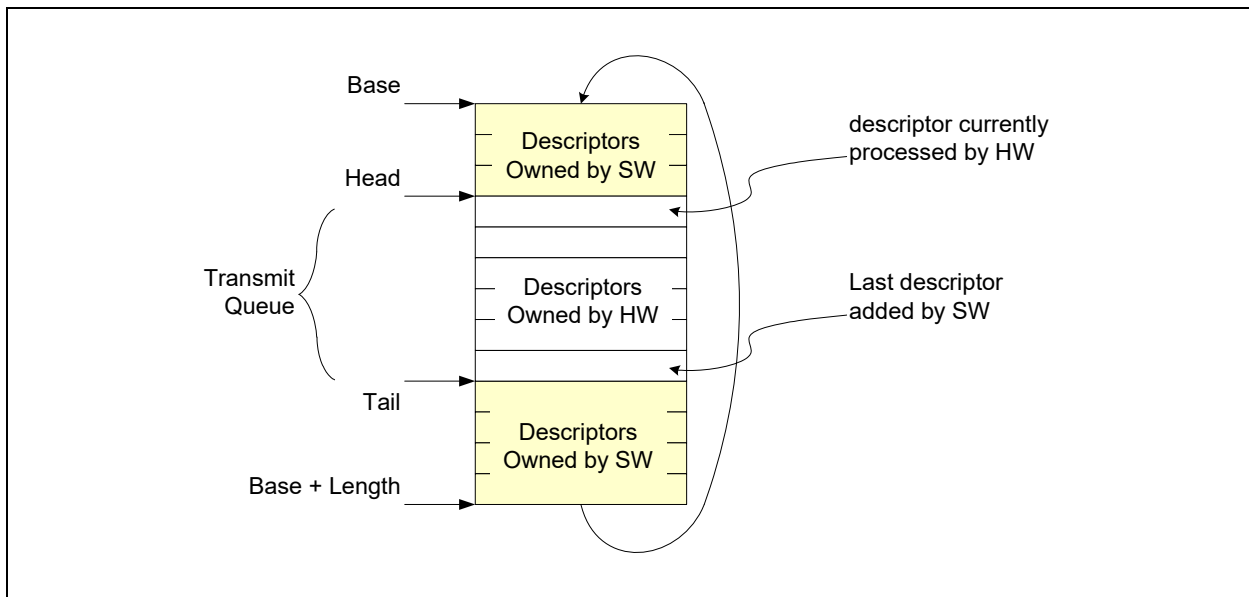


Figure 7.20. Transmit Descriptor Ring Structure

The transmit descriptor ring is defined by the following registers:

- Transmit Descriptor Base Address register (TDBA 0-127) — This register indicates the start address of the descriptor ring buffer in the host memory; this 64-bit address is aligned on a 16-byte boundary and is stored in two consecutive 32-bit registers. Hardware ignores the lower four bits.
- Transmit Descriptor Length register (TDLEN 0-127) — This register determines the number of bytes allocated to the circular buffer. This value must be 0 modulo 128.
- Transmit Descriptor Head register (TDH 0-127) — This register holds a value that is an offset from the base and indicates the in-progress descriptor. There can be up to 64 K minus 8 descriptors in the circular buffer. The transmit queue consists of the descriptors between the head and tail pointers. Transmission starts with the descriptor pointer by the head registers. When the DMA engine processes a descriptor, it might optionally write back the completed descriptor and then



advance the head pointer. It then processes the next descriptor up to the point that the head pointer reaches the tail. Head equals tail means that the transmit queue in host memory is empty. Reading this register indicates the hardware progress to the software. All descriptors behind the head pointer and in front of tail register are owned by the software. The other descriptors are owned by the hardware and should not be modified by the software.

- Transmit Descriptor Tail register (TDT 0-127) — This register holds a value, which is an offset from the base, and indicates the location beyond the last descriptor hardware can process. Software adds new descriptors to the ring by writing descriptors in the circular buffer pointed by the tail pointer. The new descriptor(s) are indicated to hardware by updating the tail pointer one descriptor above the last added descriptor. Note that a single packet or TSO might be composed of multiple descriptors. The transmit tail pointer should never point to the middle of a packet or TSO, which might cause undesired software/hardware races.

Software might detect which packets have already been processed by hardware using the following:

- Read the TDH head register to determine which packets (those logically before the head) have been transferred to the on-chip FIFO or transmitted. This method is not recommended as races between the internal update of the head register and the actual write back of descriptors can occur.
- When head write back is enabled (TDWBAL[n].Head_WB_En = 1b) software might read the image of the head pointer in host memory at the address defined by TDWBAH[n]/TDWBAL[n] pair. Hardware updates the head image in host memory by completed descriptors as described in [Section 7.2.3.5.2](#).
- When head write back is not enabled (TDWBAL[n].Head_WB_En = 0b), software might track the *DD* bits in the descriptor ring. Descriptor write back is controlled by the *RS* bit and the *WTHRESH* setting as well as interrupt assertion.
- Issue an interrupt. An interrupt condition is generated each time a packet was transmitted or received and a descriptor was write back or transmit queue goes empty (EICR.RTxQ[0-19]). This interrupt can either be enabled or masked.

All of the registers controlling the descriptor rings behavior should be set before transmit is enabled.

7.2.3.4 Transmit Descriptor Fetching

The Intel® Xeon® Processor D-1500 Product Family LAN controller fetches new descriptors as required for packet transmission depending on its on-die descriptor buffer state:

Fetch — The on-chip descriptor buffer is empty or contains less descriptors than a complete packet.

- A fetch starts as soon as any descriptors are made available (host writes to the tail pointer).
- A request is issued for any available descriptors up to the size of the on-die buffer.
- Once the sum of on-die descriptors and requested descriptors is more than required for a single packet, the buffer transitions to the pre-fetch state.
- If several on-chip descriptor queues are empty simultaneously, queues are served in round robin arbitration except the queues indicated as strict priority, which are served first.

Pre-Fetch — The on-chip descriptor buffer becomes almost empty while there are enough descriptors in the host memory.

- The on-chip descriptor buffer is defined as almost empty if it contains less descriptors than the threshold defined by TXDCTL[n].PTHRESH
- The transmit descriptor contains enough descriptors if it includes more ready descriptors than the threshold defined by TXDCTL[n].HTHRESH
- In pre-fetch mode, descriptors are fetched only after there are no other DMA activity of greater priority as: transmit descriptor fetch; status write-backs or packet data transfers)
- A request is issued for any available descriptors up to the capacity of the on-die buffer.



- If several on-chip descriptor queues are in this situation simultaneously, queues are served in round robin arbitration except the queues indicated as strict priority which are served first.

Idle — Requests are not issued. This is the state reached when none of the previous states apply.

Note: Software must update the Tail register on packet boundaries. That is, the last valid descriptor might not be a context descriptor and must have the *EOP* bit set.

7.2.3.4.1 Transmit Descriptor Fetch and Write-back Settings

This section describes the settings of transmit descriptor thresholds. It relates to fetch thresholds described above as well as the write-back threshold (*WTHRESH*) when operating in descriptor write-back mode, which is described in [Section 7.2.3.5.1](#).

- Transmit descriptor fetch setting is programmed in the *TXDCTL[n]* register per queue. The default settings of *PTHRESH*, *HTHRESH* and *WTHRESH* are zero's.
- In order to reduce transmission latency, it is recommended to set the *PTHRESH* value as high as possible while the *HTHRESH* and *WTHRESH* as low as possible (down to zero).
- In order to minimize PCIe overhead the *PTHRESH* should be set as low as possible while *HTHRESH* and *WTHRESH* should be set as high as possible.
- The sum of *PTHRESH* plus *WTHRESH* must not be greater than the on-chip descriptor buffer size
- Some practical rules
 - CPU cache line optimization: Assume 'N' equals the CPU cache line divided by 16 (descriptor size). Then, in order to align descriptors pre-fetch to CPU cache line (in most cases), it is advised to set *PTHRESH* to the on-chip descriptor buffer size minus 'N' and *HTHRESH* to 'N'. In order to align descriptor write back to the CPU cache line it is advised to set *WTHRESH* to either 'N' or even 2 times 'N'. Note that partial cache line writes might significantly degrade performance. Therefore, it is highly recommended to follow this advice.
 - Minimizing PCIe overhead: As an example, setting *PTHRESH* to the on-chip descriptor buffer size minus 16 and *HTHRESH* to 16 minimizes the PCIe request and header overhead to ~20% of the bandwidth required for the descriptor fetch.
 - Minimizing transmission latency from tail update: Setting *PTHRESH* to the on-chip descriptor buffer size minus 'N' ('N' previously defined) while *HTHRESH* and *WTHRESH* to zero.

Note: As previously described, device setting is a trade-off between overhead (translated to performance) and latencies. It is expected that some level of optimization is done at software driver development phase. Customers who want better performance might need to adjust the threshold values according to the previous guidelines while optimizing to specific platform and targets.


7.2.3.5 Transmit Write Back

The Intel® Xeon® Processor D-1500 Product Family LAN controller periodically updates software on its progress in processing transmit buffers. Two methods are described for doing so:

- Updating by writing back into the Tx descriptor
- Update by writing to the head pointer in system memory

7.2.3.5.1 Tx Descriptor Write Back

When the *TXDCTL[n].WTHRESH* equals zero, descriptors are written back for those descriptors with the *RS* bit set. When the *TXDCTL[n].WTHRESH* value is greater than zero, descriptors are accumulated until the number of accumulated descriptors equals the *TXDCTL[n].WTHRESH* value, then these descriptors are written back. Accumulated descriptor write back enables better use of the PCIe bus and memory bandwidth.



Any descriptor write back includes the full 16 bytes of the descriptor.

Descriptors are written back in one of three cases:

- $TXDCTL[n].WTHRESH = 0b$ and a descriptor that has RS set is ready to be written back.
- $TXDCTL[n].WTHRESH > 0b$ and $TXDCTL[n].WTHRESH$ descriptors have accumulated.
- $TXDCTL[n].WTHRESH > 0b$ and the corresponding $EITR$ counter has reached zero. The timer expiration flushes any accumulated descriptors and sets an interrupt event (TXDW).

An additional mode in which transmit descriptors are not written back at all and the head pointer of the descriptor ring is written instead is described in the following section.

7.2.3.5.2 Tx Head Pointer Write Back

In legacy hardware, transmit requests are completed by writing the DD bit to the transmit descriptor ring. This causes cache trash since both the driver and hardware are writing to the descriptor ring in host memory. Instead of writing the DD bits to signal that a transmit request is complete, hardware can write the contents of the descriptor queue head to host memory. The driver reads that memory location to determine which transmit requests are complete. In order to improve the performance of this feature, the driver needs to program TPH registers to configure which CPU will be processing each TX queue.

The head pointer is reflected in a memory location that is allocated by software for each queue.

Rules for head pointer write back:

- Head write back occurs if $TDWBAL[n].Head_WB_En$ is set for this queue, and the RS bit is set in the Tx descriptor, following its corresponding data upload into packet buffer.
 - If the head write-back feature is enabled, software must set $WTHRESH$ to $0x0$ while only descriptors with the RS bit set, generate header write back.
 - Note that the head pointer write back does not hold transmission. Instead, if packets with the RS bit are transmitted fast enough, it might happen that the header pointer write back is not updated for each and every packet. In addition, it might happen that the head pointer write back might be updated up to descriptors that do not have the RS bit set. In such cases, hardware might report a completion of a descriptor that might not be the last descriptor in a TSO or even the last descriptor in a single packet.

The driver has control of this feature per queue through the TDWBAL and TDWBAH registers.

The low register's LSB hold the control bits.

- The $Head_WB_EN$ bit enables activation of tail write back. In this case, no descriptor write back is executed.
- The 30 upper bits of this register hold the lowest 32 bits of the head write-back address, assuming that the two last bits are zero.

The high register holds the high part of the 64-bit address.

Note: Hardware writes a full Dword when writing this value, so software should reserve enough space for each head value and make sure the TDWBAL value is Dword-aligned.



7.2.4 TCP and UDP Segmentation

Hardware TCP segmentation is one of the offloading options supported by the Windows* and Linux* TCP/IP stack. This is often referred to as Large Send offloading or TSO. This feature enables the TCP/IP stack to pass to the network device driver a message to be transmitted that is bigger than the Maximum Transmission Unit (MTU) of the medium. It is then the responsibility of the device driver and hardware to divide the TCP message into MTU size frames that have appropriate layer 2 (Ethernet), 3 (IP), and 4 (TCP) headers. These headers must include sequence number, checksum fields, options and flag values as required. Note that some of these values (such as the checksum values) are unique for each packet of the TCP message, and other fields such as the source IP address is constant for all packets associated with the TCP message.

Similar to TCP segmentation, the Intel® Xeon® Processor D-1500 Product Family LAN controller also provides a capability to offload UDP segmentation. Note that current UDP segmentation offload is not supported by any standard operating system.

Note: CRC appending (HLREG0.TXCRCEN) must be enabled in TCP / UDP segmentation mode because CRC is inserted by hardware.
Padding (HLREG0.TXPADEN) must be enabled in TCP / UDP segmentation mode, since the last frame might be shorter than 60 bytes — resulting in a bad frame if TXPADEN is disabled.

The offloading of these mechanisms to the device driver and the Intel® Xeon® Processor D-1500 Product Family LAN controller saves significant CPU cycles. The device driver shares the additional tasks to support these options with the Intel® Xeon® Processor D-1500 Product Family LAN controller.

7.2.4.1 Assumptions and Restrictions

The following assumptions apply to the TCP / UDP segmentation implementation in the Intel® Xeon® Processor D-1500 Product Family LAN controller:

- To limit the internal cache dimensions, software is required to spread the header onto a maximum four descriptors, while still allowed to mix header and data in the last header buffer. This limitation stands for up to Layer 4 header included, and for IPv4 or IPv6 independently.
- The maximum size of a single TSO can be as large as defined by the *PAYLEN* field in the Tx data descriptor (such as up to 256 KB).
- The *RS* bit operation is not changed. Interrupts are set after data in the buffers pointed to by individual descriptors is transferred (DMA'ed) to hardware.
- SNAP packets are not supported for segmentation.
- IP in IP tunneled packets are not supported for offloading under TSO operation. VXLAN and NVGRE tunneled packets can be segmented.
- Software must enable the Ethernet CRC offload in the HLREG0.TXCRCEN register since CRC must be inserted by hardware after the checksum has been calculated.
- Software must initialize the appropriate checksum fields in the packet's header.

7.2.4.2 Transmission Process

The transmission process involves the following:

- The protocol stack receives from an application a block of data that is to be transmitted.
- The protocol stack calculates the number of packets required to transmit this block based on the MTU size of the media and required packet headers.
- The stack interfaces with the device driver and passes the block down with the appropriate header information: Ethernet, IP and TCP / UDP headers.



- The stack interfaces with the device driver and commands the driver to send the individual packet. The device driver sets up the interface to the hardware (via descriptors) for the TCP / UDP segmentation.
- The hardware transfers (DMA's) the packet data and performs the Ethernet packet segmentation and transmission based on offset and payload length parameters in the TCP/IP or UDP/IP context descriptor including:
 - Packet encapsulation
 - Header generation and field updates including IPv4/IPv6 and TCP/UDP checksum generation.
- The driver returns ownership of the block of data to the NOS when the hardware has completed the DMA transfer of the entire data block.

7.2.4.2.1 TCP and UDP Segmentation Data Fetch Control

To perform TCP / UDP segmentation in the Intel® Xeon® Processor D-1500 Product Family LAN controller, the DMA must be able to fit at least one packet of the segmented payload into available space in the on-chip packet buffer. The DMA does various comparisons between the remaining payload and the packet buffer available space, fetching additional payload and sending additional packets as space permits.

The Intel® Xeon® Processor D-1500 Product Family LAN controller enables interleaving between different TSO requests at an Ethernet packet level. In other words, the Intel® Xeon® Processor D-1500 Product Family LAN controller might fetch part of a TSO from a queue, equivalent to one or more Ethernet packets, then transition to another queue and fetch the equivalent of one or more packets (TSO or not), then move to another queue (or the first queue), etc. The Intel® Xeon® Processor D-1500 Product Family LAN controller decides on the order of data fetched based on its QoS requirements (such as bandwidth allocation and priority).

In order to enable interleaving between descriptor queues at the Ethernet frame resolution inside TSO requests, the frame header pointed by the so called header descriptors are re-read from system memory for every TSO segment (once per packet), storing in an internal cache only the header's descriptors instead of the header's content.

7.2.4.2.2 TCP and UDP Segmentation Write-back Modes

TCP / UDP segmentation mode uses the buffers that contain the header of the packet multiple times (once for each transmitted segment). Software should guarantee that the header buffers are available throughout the entire TSO transmission. Therefore, software should not re-use any descriptors of the TSO header during the TSO transmission.

7.2.4.3 TCP and UDP Segmentation Performance

Performance improvements for a hardware implementation of TCP / UDP segmentation offload include:

- The stack does not need to partition the block to fit the MTU size, saving CPU cycles.
- The stack only computes one Ethernet, IP, and TCP / UDP header per segment, saving CPU cycles.
- The stack interfaces with the device driver only once per block transfer, instead of once per frame.
- Larger PCI bursts are used, which improves bus efficiency (such as lowering transaction overhead).
- Interrupts are easily reduced to one per TCP / UDP message instead of one per packet.
- Fewer I/O accesses are required to command the hardware.



7.2.4.4 Packet Format

A TCP / UDP message can be as large as 256 KB and is generally fragmented across multiple pages in host memory. The Intel® Xeon® Processor D-1500 Product Family LAN controller partitions the data packet into standard Ethernet frames prior to transmission. The Intel® Xeon® Processor D-1500 Product Family LAN controller supports calculating the Ethernet, IP, TCP, and UDP headers, including checksum, on a frame-by-frame basis. For tunneled packets (NVGRE, VXLAN) also supports updating the outer IPv4 header.

Table 7.41. TCP/IP and UDP/IP Packet Format Sent by Host

Pseudo Header				Data
Ethernet	Optional tunnel header	IPv4/IPv6	TCP/UDP	DATA (full TCP/UDP message)

Table 7.42. Packets Format Sent by Device

Pseudo Header (updated)	Data (first MSS)	FCS	...	Pseudo Header (updated)	Data (Next MSS)	FCS	...
-------------------------	------------------	-----	-----	-------------------------	-----------------	-----	-----

Frame formats supported by the Intel® Xeon® Processor D-1500 Product Family LAN controller include:

- Ethernet 802.3
- IEEE 802.1Q VLAN (Ethernet 802.3ac)
- Ethernet Type 2
- NVGRE (see frame format in Section A.2.5.1)
- VXLAN (see frame format in Section A.2.5.2)
- IPv4 headers with options
- IPv6 headers with extensions
- TCP with options
- UDP with options

VLAN tag insertion can be handled by hardware.

Note: UDP (unlike TCP) is not a reliable protocol and fragmentation is not supported at the UDP level. UDP messages that are larger than the MTU size of the given network medium are normally fragmented at the IP layer. This is different from TCP, where large TCP messages can be fragmented at either the IP or TCP layers depending on the software implementation.

The Intel® Xeon® Processor D-1500 Product Family LAN controller has the ability to segment UDP traffic (in addition to TCP traffic); however, because UDP packets are generally fragmented at the IP layer, the Intel® Xeon® Processor D-1500 Product Family LAN controller's segmentation capability might not be used in practice for UDP.

7.2.4.5 TCP and UDP Segmentation Indication

Software indicates a TCP / UDP segmentation transmission context to the hardware by setting up a TCP/IP or UDP/IP context transmit descriptor (see Section 7.2.3). The purpose of this descriptor is to provide information to the hardware to be used during the TCP / UDP segmentation offload process.



Setting the *TSE* bit in the DCMD field to one (in the data descriptor) indicates that this descriptor refers to the segmentation context (as opposed to the normal checksum offloading context). This causes the checksum offloading, packet length, header length, and maximum segment size parameters to be loaded from the descriptor into the device.

The TCP / UDP segmentation prototype header is taken from the packet data itself. Software must identify the type of packet that is being sent (IPv4/IPv6, TCP/UDP, other), calculate appropriate checksum off loading values for the desired checksums, and then calculate the length of the header that is pre-appended. The header can be up to 240 bytes in length.

Once the TCP / UDP segmentation context has been set, the next descriptor provides the initial data to transfer. This first descriptor(s) must point to a packet of the type indicated. Furthermore, the data it points to might need to be modified by software as it serves as the prototype header for all packets within the TCP / UDP segmentation context. The following sections describe the supported packet types and the various updates that are performed by hardware. This should be used as a guide to determine what must be modified in the original packet header to make it a suitable prototype header.

The following summarizes the fields considered by the driver for modification in constructing the prototype header.

IP Header

For IPv4 headers:

- Identification field should be set as appropriate for first packet of send (if not already).
- Header checksums of inner and outer IP headers should be zeroed out unless some adjustment is needed by the driver.

TCP Header

- Sequence number should be set as appropriate for first packet of send (if not already).
- PSH, and FIN flags should be set as appropriate for LAST packet of send.
- TCP checksum should be set to the partial pseudo-header checksum as follows (there is a more detailed discussion of this in [Section 7.2.4.6](#)):

Table 7.43. TCP Partial Pseudo-header Checksum for IPv4

IP Source Address		
IP Destination Address		
Zero	Layer 4 Protocol ID	Zero


Table 7.44. TCP Partial Pseudo-header Checksum for IPv6

IPv6 Source Address	
IPv6 Final Destination Address	
Zero	
Zero	Next Header

UDP Header

- Checksum should be set as in TCP header, as previously explained.

The following sections describe the updating process performed by the hardware for each frame sent using the TCP segmentation capability.



7.2.4.6 Transmit Checksum Offloading with TCP and UDP Segmentation

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports checksum offloading as a component of the TCP / UDP segmentation off-load feature and as stand-alone capability. [Section 7.2.5](#) describes the interface for controlling the checksum off-loading feature. This section describes the feature as it relates to TCP / UDP segmentation.

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports IP and TCP header options in the checksum computation for packets that are derived from the TCP segmentation feature.

Two specific types of checksum are supported by the hardware in the context of the TCP / UDP segmentation off-load feature:

- IPv4 checksum
- TCP / UDP checksum

Each packet that is sent via the TCP / UDP segmentation off-load feature optionally includes the IPv4 checksum and/or the TCP / UDP checksum.

All checksum calculations use a 16-bit wide one's complement checksum. The checksum word is calculated on the outgoing data.

Refer to [Table 7.45](#) for the list of supported transmit checksums per packet type.

7.2.4.7 IP/TCP / UDP Header Updating

IP/TCP and IP/UDP header is updated for each outgoing frame based on the header prototype that hardware DMA's from the first descriptor(s). The checksum fields and other header information are later updated on a frame-by-frame basis. The updating process is performed concurrently with the packet data fetch.

The following sections define what fields are modified by hardware during the TCP / UDP segmentation process by the Intel® Xeon® Processor D-1500 Product Family LAN controller.

7.2.4.7.1 TCP/IP/UDP Header for the First Frame

The hardware makes the following changes to the headers of the first packet that is derived from each TCP segmentation context.

Tunnel Header

When tunnel offload is enabled (*Tunnel* bit set in the descriptor) the outer IPv4 header is updated as follow

- $IP\ Total\ Length = OUTERIPLen + TUNNELLEN + IPLen + L4LEN + MSS$
- Calculates the Outer IP Checksum

IPv4 Header

- $IP\ Total\ Length = MSS + L4LEN + IPLen$
- Calculates the IP Checksum

IPv6 Header

- $Payload\ Length = MSS + L4LEN + IPV6_HDR_extension^1$

1. $IPV6_HDR_extension$ is calculated as $IPLen - 40$ bytes.



TCP Header

- Sequence Number: The value is the sequence number of the first TCP byte in this frame.
- The flag values of the first frame are set by logic AND function between the flag word in the pseudo header and the DTXTCPFLGL.TCP_flg_first_seg. The default values of the DTXTCPFLGL.TCP_flg_first_seg are set. The flags in a TSO that ends up as a single segment are taken from the in the pseudo header in the Tx data buffers as is.
- Calculates the TCP checksum.

UDP Header

- Calculates the UDP checksum.

7.2.4.7.2 TCP/IP Header for the Subsequent Frames

The hardware makes the following changes to the headers for subsequent packets that are derived as part of a TCP segmentation context:

Number of bytes left for transmission = PAYLEN — (N * MSS). Where N is the number of frames that have been transmitted.

Tunnel Header

When tunnel offload is enabled (*Tunnel* bit set in the descriptor) the outer IPv4 header is updated as follow

- IP Identification: increased from last value (wrap around based on 16-bit width)
- IP Total Length = OUTERIPLLEN + TUNNELLEN + IPLLEN + L4LEN + MSS
- Calculates the Outer IP Checksum

IPv4 Header

- IP Identification: increased from last value (wrap around)
- IP Total Length = MSS + L4LEN + IPLLEN
- Calculate the IP Checksum

IPv6 Header

- Payload Length = MSS + L4LEN + IPV6_HDR_extension¹

TCP Header

- Sequence Number update: Add previous TCP payload size to the previous sequence number value. This is equivalent to adding the MSS to the previous sequence number.
- The flag values of the subsequent frames are set by logic AND function between the flag word in the pseudo header with the DTXTCPFLGL.TCP_Flg_mid_seg. The default values of the DTXTCPFLGL.TCP_Flg_mid_seg are set.
- Calculate the TCP checksum

UDP Header

- Calculates the UDP checksum.

7.2.4.7.3 TCP/IP Header for the Last Frame

Hardware makes the following changes to the headers for the last frame of a TCP segmentation context:



Last frame payload bytes = PAYLEN — (N * MSS).

Tunnel Header

When tunnel offload is enabled (*Tunnel* bit set in the descriptor) the outer IPv4 header is updated as follow

- IP Identification: increased from last value (wrap around based on 16-bit width)
- IP Total Length = OUTERIPLLEN + TUNNELLEN + IPLLEN + L4LEN + last frame payload bytes.
- Calculates the Outer IP Checksum

IPv4 Header

- IP Total length = last frame payload bytes + L4LEN + IPLLEN
- IP identification: increased from last value (wrap around based on 16-bit width)
- Calculate the IP checksum

IPv6 Header

- Payload length = last frame payload bytes + L4LEN + IPV6_HDR_extension¹

TCP Header

- Sequence number update: Add previous TCP payload size to the previous sequence number value. This is equivalent to adding the MSS to the previous sequence number.
- The flag values of the last frames are set by logic AND function between the flag word in the pseudo header and the DTXTCPFLGH.TCP_Flg_Ist_seg. The default values of the DTXTCPFLGH.TCP_Flg_Ist_seg are set. The flags in a TSO that ends up as a single segment are taken from the in the pseudo header in the Tx data buffers as is.
- Calculate the TCP checksum

UDP Header

- Calculates the UDP checksum.

7.2.5 Transmit Checksum Offloading in Non-segmentation Mode

The previous section on TCP / UDP segmentation offload describes the IP/TCP/UDP checksum offloading mechanism used in conjunction with segmentation. The same underlying mechanism can also be applied as a stand-alone checksum offloading. The main difference in a single packet send is that only the checksum fields in the IP/TCP/UDP headers are calculated and updated by hardware.

Before taking advantage of the Intel® Xeon® Processor D-1500 Product Family LAN controller's enhanced checksum offload capability, a checksum context must be initialized. For a single packet send, DCMD.TSE should be set to zero (in the data descriptor). For additional details on contexts, refer to [Section 7.2.3.3](#).

Enabling checksum offload, software must also enable Ethernet CRC offload by the HLREG0.TXCRCEN since CRC must be inserted by hardware after the checksum has been calculated.

Each checksum operates independently. Insertion of the IP and TCP / UDP checksum for each packet are enabled through the transmit data descriptor POPTS.TXSM and POPTS.IXSM fields, respectively.

7.2.5.1 IP Checksum

Three fields in the transmit context descriptor set the context of the IP checksum offloading feature:

- TUCMD.IPV4



- IPLEN
- MACLEN

TUCMD.IPV4=1 specifies that the packet type for this context is IPv4, and that the IP header checksum should be inserted. TUCMD.IPV4=0 indicates that the packet type is IPv6 (or some other protocol) and that the IP header checksum should not be inserted.

MACLEN specifies the byte offset from the start of the DMA'ed data to the first byte to be included in the checksum, the start of the IP header. The minimal allowed value for this field is 14. Note that the maximum value for this field is 127. This is adequate for typical applications.

Note: The MACLEN+IPLEN value must be less than the total DMA length for a packet. If this is not the case, the results are unpredictable.

IPLEN specifies the IP header length. Maximum allowed value for this field is 511 bytes.

MACLEN+IPLEN specify where the IP checksum should stop. The sum of MACLEN+IPLEN must be smaller equals to the first 638 (127+511) bytes of the packet and obviously must be smaller or equal to the total length of a given packet. If this is not the case, the result is unpredictable.

For IP tunnel packets (IPv4-IPv6), IPLEN must be defined as the length of the two IP headers. Hardware is able to offload the L4 checksum calculation while software should provide the IPv4 checksum.

For NVGRE and VXLAN tunneled packet as defined by the *Tunnel* bit in the descriptor, Outer IPv4 checksum is controlled by OUTERIPCS bit and is determined by MACLEN and OUTERIPLEN similar to the above description. Inner IPv4 checksum is controlled by TUCMD.IPv4 and offset determination should be adjusted by HW such that MACLEN should be replaced by (MACLEN+OUTERIPLEN+TUNNELLEN)

The 16-bit IPv4 header checksum is placed at the two bytes starting at MACLEN+10.

7.2.5.2 TCP and UDP Checksum

Three to five fields in the transmit context descriptor set the context of the TCP / UDP checksum offloading feature:

- MACLEN
- OUTERIPLEN (optional)
- TUNNELLEN (optional)
- IPLEN
- TUCMD.L4T

TUCMD.L4T=01b specifies that the packet type is TCP, and that the 16-bit TCP header checksum should be inserted at byte offset MACLEN+IPLEN+16. TUCMD.L4T=00b indicates that the packet is UDP and that the 16-bit checksum should be inserted starting at byte offset MACLEN+IPLEN+6. If *Tunnel* bit is set, the checksum is inserted at MACLEN+OUTERIPLEN+TUNNELLEN+IPLEN+6

MACLEN+(OUTERIPLEN+TUNNELLEN)+IPLEN specifies the byte offset from the start of the DMA'ed data to the first byte to be included in the checksum, the start of the UDP/TCP header. See MACLEN table in [Section 7.2.3.2.3](#) for its relevant values.

Note: The MACLEN+(OUTERIPLEN+TUNNELLEN)+IPLEN+L4LEN value must be less than the total DMA length for a packet. If this is not the case, the results are unpredictable.

The TCP/UDP checksum always continues to the last byte of the DMA data.



Note: For non-TSO, software still needs to calculate a full checksum for the TCP/UDP pseudo-header. This checksum of the pseudo-header should be placed in the packet data buffer at the appropriate offset for the checksum calculation.

7.2.5.3 SCTP CRC Offloading

For SCTP packets, a CRC32 checksum offload is provided.

Three fields in the transmit context descriptor set the context of the STCP checksum offloading feature:

- MACLEN
- IPLEN
- TUCMD.L4T

TUCMD.L4T=10b specifies that the packet type is SCTP, and that the 32-bit STCP CRC should be inserted at byte offset MACLEN+IPLen+8.

IPLen+MACLEN specifies the byte offset from the start of the DMA'ed data to the first byte to be included in the checksum, the start of the STCP header. The minimal allowed value for this sum is 26. See MACLEN table in [Section 7.2.3.2.3](#) for its relevant values.

The SCTP CRC calculation always continues to the last byte of the DMA data.

The SCTP total L3 payload size (PAYLEN - IPLen - MACLEN) should be a multiple of four bytes (SCTP padding not supported).

Note: TSO is not available for SCTP packets.

Note: The CRC field of the SCTP header must be set by driver to zero prior to requesting a CRC calculation offload.

7.2.5.4 Checksum Supported per Packet Types

The following table lists which checksums are supported per packet type.

Note: TSO is not supported for packet types for which IP checksum and TCP / UDP checksum cannot be calculated.

Table 7.45. Checksums Supported by Packet Type

Packet Type	HW IP Checksum Calculation	HW TCP/UDP/SCTP Checksum Calculation
IPv4 packets	Yes	Yes
IPv6 packets	No (n/a)	Yes
IPv6 packet with next header options:		
• Hop-by-hop options	No (n/a)	Yes
• Destinations options	No (n/a)	Yes
• Routing (with len 0)	No (n/a)	Yes
• Routing (with len >0)	No (n/a)	No
• Fragment	No (n/a)	No
• Home option	No (n/a)	No
• Security option (AH/ESP)	No (n/a)	Yes
IPv4 tunnels:		
• Ipv4 packet in an IPv4 tunnel	No	No
• Ipv6 packet in an IPv4 tunnel	No	Yes



Table 7.45. Checksums Supported by Packet Type

Packet Type	HW IP Checksum Calculation	HW TCP/UDP/SCTP Checksum Calculation
IPv6 tunnels: <ul style="list-style-type: none"> IPv4 packet in an IPv6 tunnel IPv6 packet in an IPv6 tunnel 	No No	No No
Packet is an IPv4 fragment	Yes	No
Packet has 802.3ac tag	Yes	Yes
IPv4 packet has IPSec header without IP options	Yes	Yes
Packet has TCP or UDP options	Yes	Yes
IP header's protocol field contains protocol # other than TCP, UDP, or SCTP	Yes	No
NVGRE: <ul style="list-style-type: none"> Inner Ipv4 packet Inner Ipv6 packet 	Yes (Inner and outer) Yes (outer)	Yes Yes
VXLAN: <ul style="list-style-type: none"> Inner Ipv4 packet Inner Ipv6 packet 	Yes (Inner and outer) Yes (outer)	Yes (inner only) ¹ Yes (inner only)

1. The outer UDP header of VXLAN packets do not use a checksum.

7.2.6 Transmit Statistics

7.2.6.1 General Notes

- All Statistics registers are cleared on read. In addition, they stick at 0xFF..F when the maximum value is reached.
- Due to divergent paths between interrupt-generation and logging of relevant statistics counts, it might be possible to generate an interrupt to the system for a noteworthy event prior to the associated statistics count actually being increased. This is extremely unlikely due to expected delays associated with the system interrupt-collection and ISR delay, but might be an explanation for interrupt statistics values that do not quite make sense. Hardware guarantees that any event noteworthy of inclusion in a statistics count is reflected in the appropriate count within 1 μs; a small time-delay prior to reading the statistics might be required to avoid a potential mismatch between and interrupt and its cause.
- If TSO is enabled, statistics are collected after segmentation.
- All byte (octet) counters composed of 2 registers can be fetched by two consecutive 32-bit accesses while reading the Low 32-bit register first or a single 64-bit access.

7.2.6.2 Transmit Statistics Hierarchy

The following diagram shows the relations between the packet flow and the different statistic counters.

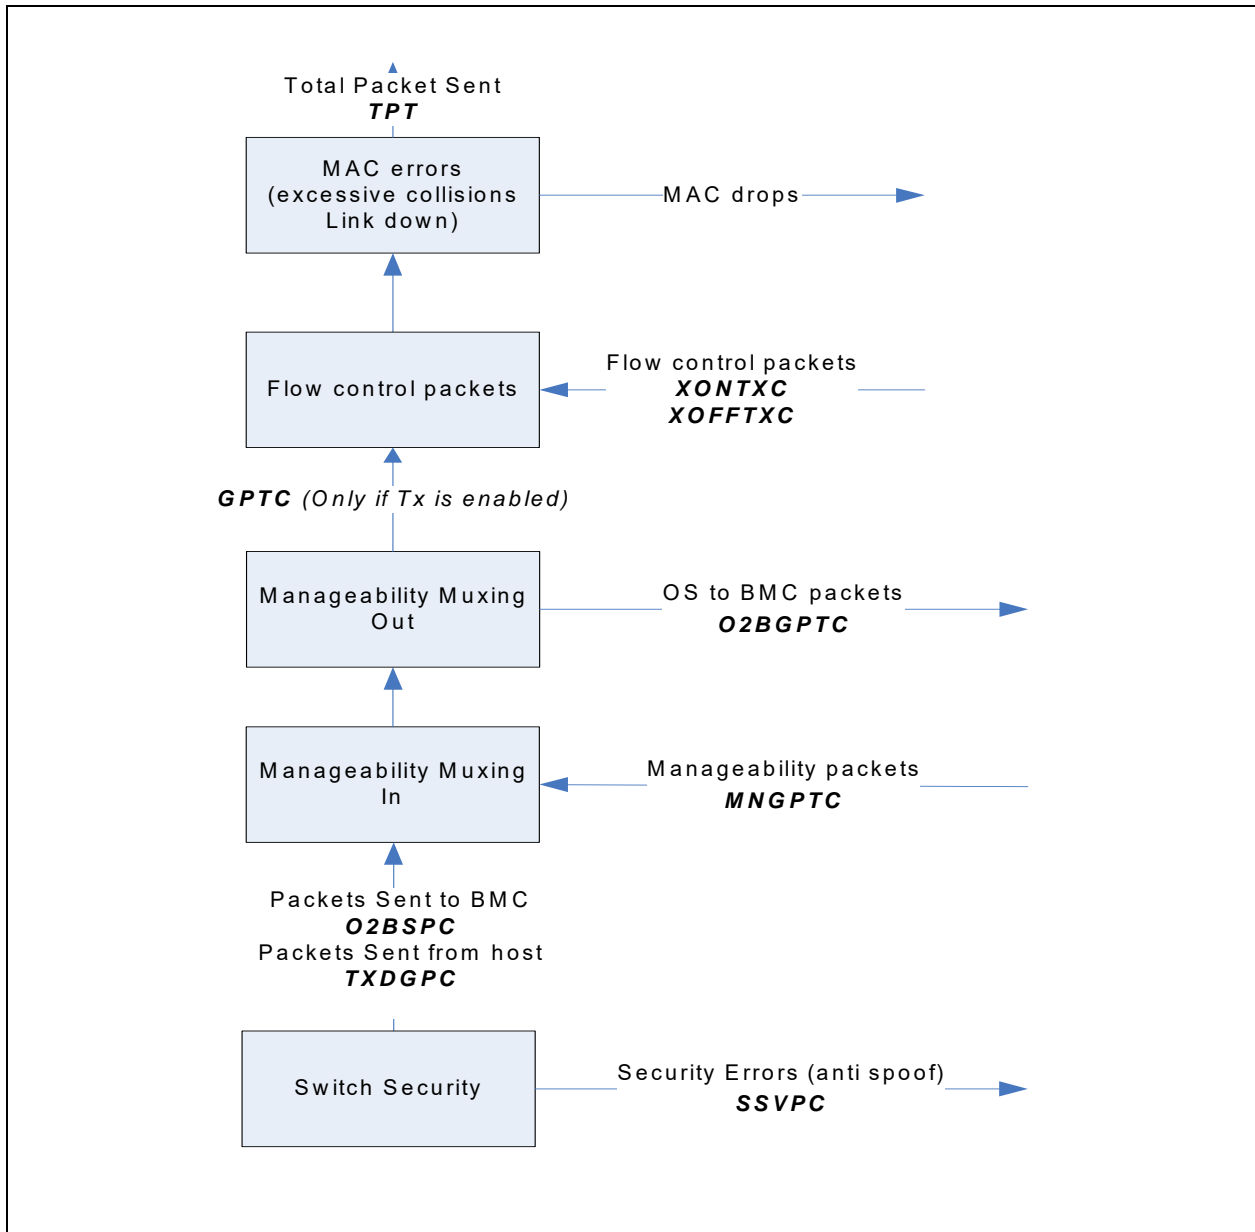


Figure 7.21. Transmit Flow Statistics



7.3 Interrupts

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports the following interrupt modes. Mapping of interrupts causes is different in each of these modes as described in this section.

- PCI legacy interrupts or MSI or MSI-X and only a single vector is allocated — selected when GPIE.Multiple_MSIX is set to 0b.
- MSI-X with multiple MSI-X vectors in non-IOV mode — selected when GPIE.Multiple_MSIX is set to 1b and GPIE.VT_Mode is set to 00b.
- MSI-X in IOV mode — selected when GPIE.Multiple_MSIX is set (as previously stated) and GPIE.VT_Mode DOES NOT equal 00b.

The following sections describe the interrupt registers and device functionality at all operation modes.

7.3.1 Interrupt Registers

Physical Function (PF) Registers


The PF interrupt logic consists of the registers listed in the [Table 7.46](#) followed by their description:

Table 7.46. PF Interrupt Registers

Acronym	Complete Name
EICR	Extended Interrupt Cause register
EICS	Extended Interrupt Cause Set register (enables software to initiate interrupts)
EIMS	Extended Interrupt Mask Set/Read register
EIMC	Extended Interrupt Mask Clear register
EIAC	Extended Interrupt Auto Clear register (following interrupt assertion)
EIAM	Extended Interrupt Auto Mask register (auto set/clear of the EIMS)
EITR	Extended Interrupt Throttling register [throttling]
IVAR	Interrupt Vector Allocation Registers (described in Section 7.3.4)
IVAR_MISC	Miscellaneous Interrupt Vector Allocation Register (described in Section 7.3.4)

These registers are extended to 64 bits by an additional set of two registers. EICR has an additional two registers EICR(1)... EICR(2) and so on for the EICS, EIMS, EIMC, EIAM and EITR registers. The EIAC register is not extended to 64 bits as this extended interrupt causes are always auto cleared. Any reference to EICR... EIAM registers as well as any global interrupt settings in the GPIE register relates to their extended size of 64 bits.

The legacy EICR[15:0] mirror the content of EICR(1)[15:0]. In the same manner the lower 16 bits of EICS, EIMS, EIMC, EIAM mirror the lower 16 bits of EICS(1), EIMS(1), EIMC(1), EIAM(1). For more details on the use of these registers in the various interrupt modes (legacy, MSI, MSI-X) see [Section 7.3.4](#).



Virtual Function (VF) Registers

The VF interrupt logic has the same set of interrupt registers while each of them has three entries for three interrupt causes. The names and functionality of these registers are the same as those of the PF with a prefix of Intel VT as follows: VFEICR, VFEICS, VFEIMS, VFEIMC, VFEIAM, VFEITR. The interrupt causes are always auto cleared. Although each VF can generate up to three interrupts, only the first two registers are capable of interrupt throttling and are associated to VFEITR registers (see [Section 7.3.4.3.2](#) for its proper usage). Each VF also has the mapping registers VFIVAR and VFIVAR_MISC. Note that any global interrupt setting by the GPIE register affect both interrupt settings of the PF as well as the VFs.

7.3.1.1 Extended Interrupt Cause (EICR) Registers

This register records the interrupt causes to provide software information on the interrupt source. Each time an interrupt cause happens, the corresponding interrupt bit is set in the EICR registers. An interrupt is generated each time one of the bits in these registers is set, and the corresponding interrupt is enabled via the EIMS registers. The possible interrupt causes are as follows:

- Each *RTxQ* bit represents the following events: Tx or Rx descriptor write back; Rx queue full and Rx descriptor queue minimum threshold.
- Interrupts can be throttled by ITR as configured in the EITR register
- Mapping the Tx and Rx queues to EICR is done by the IVAR registers as described in [Section 7.3.4](#). Each bit might represent an event on a single Tx or Rx queue or could represent multiple queues according to the IVAR setting. In the later case, software might not be able to distinguish between the interrupt causes other than checking all associated Tx and Rx queues.
- The Multiple_MSIX = 1b setting is useful when multiple MSI-X vectors are assigned to the device. When the GPIE.Multiple_MSIX bit is set, the *RTxQ* bits are associated with dedicated MSI-X vectors. Bit 0 is Tx / Rx interrupt associated with MSI-X vector 0 and bit 15 is Tx / Rx interrupt associated with MSI-X vector 15.

Writing a 1b to any bit in the register clears it. Writing a 0b to any bit has no effect. The EICR is also cleared on read if GPIE.OCD bit is cleared. When the GPIE.OCD bit is set, then only bits 16...29 are cleared on read. The later setting is useful for MSI-X mode in which the Tx and Rx and possibly the timer interrupts do not share the same interrupt with the other causes. Bits in the register can be auto cleared depending on the EIAC register setting (see section [Section 7.3.1.4](#)).

7.3.1.2 Extended Interrupt Cause Set (EICS) Register

This register enables software to initiate a hardware interrupt. Setting any bit on the EICS sets its corresponding bit in the EICR register while bits written to 0b have no impact. It then causes an interrupt assertion if enabled by the EIMS register. Setting any bit generates throttled interrupt depending on the GPIE.EIMEN setting: When the *EIMEN* bit is set, then setting the EICS register causes an LLI interrupt; When the *EIMEN* bit is cleared, then setting the EICS register causes an interrupt after the corresponding interrupt throttling timer expires.

Note: The *EIMEN* bit can be set high only when working in auto-mask mode (*EIAM* bit of the associated interrupt is set).



7.3.1.2.1 EICS Affect on RSC Functionality

Setting *EICS* bits causes interrupt assertion (if enabled). *EICS* settings have the same impact on RSC functionality as nominal operation:

- In ITR mode (GPIE.EIMEN = 0b), setting the *EICS* bits impact the RSC completion and interrupt assertion the same as any Rx packet. The functionality depends on the *EICS* setting schedule relative to the ITR intervals as described in [Section 7.3.2.1.1](#).
- In LLI mode (GPIE.EIMEN = 1b), setting the *EICS* bits impact the RSC completion and interrupt assertion as follow:
 - Interrupt is asserted.
 - Concurrently, hardware triggers RSC completion in all Rx queues associated with the same interrupt.
 - Most likely these RSC(s) are completed to host memory after the interrupt is already asserted. In his case, it is guaranteed that an additional interrupt is asserted when the ITR expires.

7.3.1.3 Extended Interrupt Mask Set and Read (EIMS) Register, and Extended Interrupt Mask Clear (EIMC) Register

The Extended Interrupt Mask Set and Read (EIMS) register enables the interrupts in the EICR. When set to 1b, each bit in the EIMS register, enables its corresponding bit in the EICR. Software might enable each interrupt by setting bits in the EIMS register to 1b. Reading EIMS returns its value. Software might clear any bit in the EIMS register by setting its corresponding bit in the Extended Interrupt Mask Clear (EIMC) register. Reading the EIMC register does not return any meaningful data.

This independent mechanism of setting and clearing bits in the EIMS register saves the need for read modify write and also enables simple programming in multi-thread, multi-CPU core systems.

Note: The EICR register stores the interrupt events regardless of the state of the EIMS register.

7.3.1.4 Extended Interrupt Auto Clear Enable (EIAC) Register

Each bit in this register enables auto clearing of its corresponding bit in EICR following interrupt assertion. It is useful for Tx and Rx interrupt causes that have dedicated MSI-X vectors. When the Tx and Rx interrupt causes share an interrupt with the other or a timer interrupt, the relevant EIAC bits should not be set. Bits in the EICR register that are not enabled by auto clear, must be cleared by either writing a 1b to clear or a read to clear.


Note that there are no EIAC(1)...EIAC(2) registers. The hardware setting for interrupts 16...63 is always auto clear.

Note: Bits 29:16 should never be set to auto clear since they share the same MSI-X vector. Writing to the EIAC register changes the setting of the entire register. In IOV mode, some of the bits in this register might affect VF functionality (VF-56...VF-63). It is recommended that software set the register in PF before VFs are enabled. Otherwise, a software semaphore might be required between the VF and the PF to avoid setting corruption.

7.3.1.5 Extended Interrupt Auto Mask Enable (EIAM) Register

Each bit in this register enables auto clearing and auto setting of its corresponding bit in the EIMS register as follows:

- Following a write of 1b to any bit in the EICS register (interrupt cause set), its corresponding bit in the EIMS register is auto set as well enabling its interrupt.



- A write to clear the EICR register clears its corresponding bits in the EIMS register masking further interrupts.
- A read to clear the EICR register, clears the *EIMS* bits (enabled by the EIAM) masking further interrupts. Note that if the GPIE.OCD bit is set, Tx and Rx interrupt causes are not cleared on read (bits 0:15 in the EICR). In this case, bits 0:15 in the EIMS are not cleared as well.
- In MSI-X mode the auto clear functionality can be driven by MSI-X vector assertion if GPIE.EIAME is set.

Note: Bits 29:16 should never be set to auto clear since they share the same MSI-X vector.

Writing to the EIAM register changes the setting of the entire register. In IOV mode, some of the bits in this register might affect VF functionality. It is recommended that software set the register in PF before VFs are enabled. Otherwise, a software semaphore might be required between the VF and the PF to avoid setting corruption.

If any of the *Auto Mask* enable bits is set in the EIAM registers, the GPIE.EIAME bit must be set as well.


7.3.2 Interrupt Moderation

Interrupt rates can be tuned by the EITR register for reduced CPU utilization while minimizing CPU latency. In MSI or legacy interrupt modes, only EITR register 0 can be used. In MSI-X, non-IOV mode, the Intel® Xeon® Processor D-1500 Product Family LAN controller includes 64 EITR registers 0..63 that are mapped to MSI-X vectors 0..63, respectively. In IOV mode, there are an additional 65 EITR registers that are mapped to the MSI-X vectors of the virtual functions. The mapping of MSI-X vectors to EITR registers are described in [Section 7.3.1.1](#).

7.3.2.1 Time-based Interrupt Throttling — ITR

Time-based interrupt throttling is useful to limit the maximum interrupt rate regardless of network traffic conditions. The ITR logic is targeted for Rx/Tx interrupts only. It is assumed that the software device driver will not moderate the timer, other and mail box (IOV mode) interrupts. In non-IOV mode, all 64 interrupts can be associated with ITR logic. In IOV mode, the ITR logic is shared between the PF and VFs as shown in [Figure 7.22](#). The ITR mechanism is based on the following parameters:

- **ITR Interval** field in the EITR registers — The minimum inter-interrupt interval is specified in 2.048 μ s units (at 1 Gb/s or 10 Gb/s link). When the ITR Interval equals zero, interrupt throttling is disabled and any event causes an immediate interrupt. The field is composed of nine bits enabling a range of 2.048 μ s up to 1046.528 μ s. These ITR interval times correspond to interrupt rates in the range of 488 K INT/sec to 955 INT/sec. When operating at 100 Mb/s link, the ITR interval is specified in 20.48 μ s units.
 - Due to internal synchronization issues, the ITR interval can be shortened by up to 1 μ s at 10 Gb/s or 1 Gb/s link and up to 10 μ s at 100 Mb/s link when it is triggered by packet write back or interrupt enablement.
- **ITR Counter** partially exposed in the EITR registers — Down counter that is loaded by the ITR interval each time the associated interrupt is asserted.
 - The counter is decremented by one each 1.024 μ s (at 1 Gb/s or 10 Gb/s link) and stops decrementing at zero. At 100 Mb/s link, the speed of the counter is decremented by one each 10.24 μ s.
 - If an event happens before the counter is zero, it sets the EICR. The interrupt can be asserted only when the ITR time expires (counter is zero).
 - Else (no events during the entire ITR interval), the EICR register is not set and the interrupt is not asserted on ITR expiration. The next event sets the EICR bit and generates an immediate interrupt. See [Section 7.3.2.1.1](#) for interrupt assertion when RSC is enabled.



- Once the interrupt is asserted, the ITR counter is loaded by the ITR interval and the entire cycle re-starts. The next interrupt can be generated only after the ITR counter expires once again.

7.3.2.1.1 ITR Affect on RSC Functionality

Interrupt assertion is one of the causes for RSC completion (see [Section 7.9.6](#)). When RSC is enabled on specific Rx queues, the associated ITR interval with these queues must be enabled and must be larger (in time units) than RSC delay. The ITR is divided to the two time intervals that are defined by the ITR interval and RSC delay. RSC completion is triggered after the first interval completes and the interrupt is asserted when the second interval completes.

The *RSC Delay* field is defined in the GPIE registers. *RSC Delay* can have one of the following eight values: 4 μ s, 8 μ s, 12 μ s... 32 μ s.


- The first ITR interval equals ITR interval minus RSC delay. The internal ITR counter starts at ITR interval value and counts down until it reaches the RSC delay value. Therefore, the ITR interval must be set to a larger value than the RSC delay.
- The second ITR interval equals RSC delay. The internal ITR counter continues to count down until it reaches zero.
- RSC completion can take some time (usually in the range of a few micro seconds). This time is composed by completing triggering latency and completing process latency. These delays should be considered when tuning the RSC delay. The clock frequency of the RSC completion logic depends on the link speed. As a result, the completion delay can as high as $\sim 0.8 \mu$ s at 10 Gb/s link and $\sim 8 \mu$ s at 1 Gb/s link. The RSC completion logic might take additional ~ 50 ns at 10 Gb/s link and $\sim 0.5 \mu$ s at 1 Gb/s link per RSC. In addition, there is the PCIe bus arbitration latency as well as system propagation latencies from the device up to host memory.
- Recommended RSC delay numbers are: 8 μ s at 10 Gb/s link and 28 μ s at 1 Gb/s link.
- RSC is not recommended when operating at 100 Mb/s link.

Following are cases of packet reception with respect to the ITR intervals:

- Packets are received and posted (including their status) to the Rx queue in the first ITR interval. In this case, RSC completion is triggered at the end of the first ITR interval and the interrupt is asserted at the second interval expiration.
- A packet (and its status) is received and posted to the Rx queue only after the first ITR interval has expired (either on the second interval or after the entire ITR interval has expired). In this case, RSC completion is triggered almost instantly (other than internal logic latencies). The interrupt is asserted at RSC delay time after the non-coalesced Rx status is queued to be posted to the host.
- Due to internal synchronization issues, the RSC delay can be shorten by up to 1 μ s when it is triggered by packet write back.

7.3.2.2 Immediate Interrupt

The Intel® Xeon® Processor D-1500 Product Family LAN controller might initiate an immediate interrupt when the receive descriptor ring is almost empty (Rx descriptors below a specific threshold). The threshold is defined by *SRRCTL[n].RDMTS* per Rx queue. This mechanism can protect against memory resources being used up during reception of a long burst of short packets.



7.3.3 TCP Timer Interrupt

7.3.3.1 Introduction

In order to implement TCP timers, software needs to take action periodically (every 10 ms). Today, the software device driver must rely on software-based timers, whose granularity can change from platform to platform. This software timer generates a software NIC interrupt, which then enables the driver to perform timer functions, avoiding cache thrash and enabling parallelism. The timer interval is system-specific.

It would be more accurate and more efficient for this periodic timer to be implemented in hardware. The software device driver would program a timeout value (usual value of 10 ms), and each time the timer expires, hardware sets a specific bit in the EICR register. When an interrupt occurs (due to normal interrupt moderation schemes), software reads the EICR register and discovers that it needs to process timer events.

The timeout should be programmable by the software device driver, and the driver should be able to disable the timer interrupt if it is not needed.

7.3.3.2 Description

A stand-alone, down-counter is implemented. An interrupt is issued each time the value of the counter is zero.

Software is responsible for setting an initial value for the timer in the *Duration* field. Kick-starting is done by writing a 1b to the *KickStart* bit.

Following kick starting, an internal counter is set to the value defined by the *Duration* field. Then the counter is decreased by one each ms. When the counter reaches zero, an interrupt is issued. The counter re-starts counting from its initial value if the *Loop* field is set.

7.3.4 Mapping of Interrupt Causes

The following sections describe legacy, MSI and MSI-X interrupt modes.

7.3.4.1 Legacy and MSI Interrupt Modes

In legacy and MSI modes, an interrupt cause is reflected by setting one of the bits in the EICR register, where each bit reflects one or more causes. All interrupt causes are mapped to a single interrupt signal: either legacy INTA/B or MSI. This section describes the mapping of interrupt causes (that is a specific Rx or Tx queue event or any other event) to bits in the EICR.

The TCP timer and all other interrupt causes are mapped directly to EICR[30:16]. Note that the IVAR_MISC register is not used in legacy and MSI modes.

Mapping the Tx and Rx queues to interrupt bits in the EICR register is programmed in the IVAR registers as shown in [Figure 7.22](#). Each entry in the IVAR registers is composed of two fields that identify the associated bit in the EICR[15:0] register. Software might map multiple Tx and Rx queues to the same EICR bit.

INT_Alloc — Defines one of the bits (0...15) in the EICR register that reflects the interrupt status indication.

INT_Alloc_val — Valid bit for the this interrupt cause.

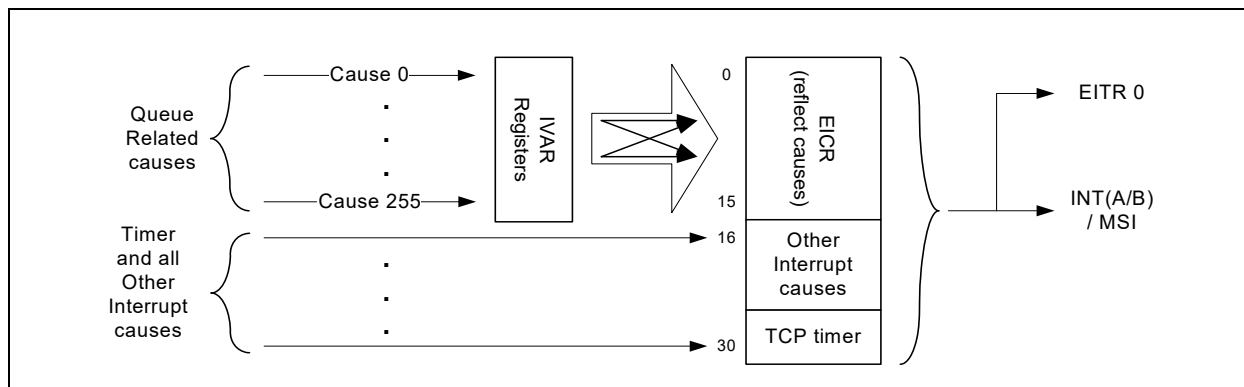


Figure 7.22. Cause Mapping in Legacy and MSI Modes

Mapping between the Tx and Rx queue to the IVAR registers is hard-wired as shown in the Figure 7.23 below:




Figure 7.23. Rx and Tx Queue Mapping to IVAR Registers

7.3.4.2 MSI-X Mode in Non-IOV Mode

- MSI-X defines a separate optional extension to basic MSI functionality. Hardware indicates the number of requested MSI-X vectors in the table size in the MSI-X capability structure in the configuration space. The number of requested MSI-X vectors is loaded from NVM in the *PCI_CNF2.MSI_X_PF_N* field up to maximum of 64 MSI-X vectors. The operating system might allocate any number of MSI-X vectors to the device from a minimum of one up to the requested number of MSI-X vectors.
- Enables interrupts causes allocation to the assigned MSI-X vectors. Interrupt allocation is programmed by the IVAR registers and are described in this section.
- Each vector can use an independent address and data value as programmed directly by the operating system in the MSI-X vector table.
- Each MSI-X vector is associated to an EITR register with the same index (MSI-X 0 to EITR[0], MSI-X 1 to EITR[1],...).

For more information on MSI-X, refer to the PCI Local Bus Specification, Revision 3.0.

MSI-X vectors can be used for several purposes:

1. Dedicated MSI-X vectors per interrupt cause (avoids the need to read the interrupt cause register).
2. Load balancing by MSI-X vectors assignment to different CPUs.
3. Optimized interrupt moderation schemes per MSI-X vector using the EITR registers.



The MSI-X vectors are used for Tx and Rx interrupt causes as well as the other and timer interrupt causes. The remainder of this section describes the mapping of interrupt causes (such as a specific Rx or Tx queue event or any other event) to the interrupts registers and the MSI-X vectors.


The TCP timer and other events are reflected in EICR[30:16] the same as the legacy and MSI mode. It is then mapped to the MSI-X vectors by the IVAR_MISC register as shown in [Figure 7.24](#). The IVAR_MISC register includes two entries for the timer interrupt and an additional entry for all the other causes. The structure of each entry is as follows:

INT_Alloc — Defines the MSI-X vector (0...63) assigned to this interrupt cause.

INT_Alloc_val — Valid bit for the this interrupt cause.

The Tx and Rx queues are associated to the IVAR0...IVAR63 the same as legacy and MSI mode shown in [Figure 7.23](#). The Tx and Rx queues are mapped by the IVAR registers to EICR(1),...EICR(2) registers and MSI-X vectors 0...63 illustrated in [Figure 7.24](#). The IVAR entries have the same structure as the IVAR_MISC register previously shown. Each bit in EICR(1...2) registers is associated to MSI-X vector 0...63 as follows:

- EICR(i).bit_num is associated to MSI-X vector (i x 32 + bit_num).
- The legacy EICR[15:0] mirror the content of EICR(1)[15:0]. In the same manner the lower 16 bits of EICS, EIMS, EIMC, EIAM mirror the lower 16 bits of EICS(1), EIMS(1), EIMC(1), EIAM(1). The use of these registers depends on the number of assigned MSI-X interrupts as follows:
 - **16 Tx and Rx Interrupts** — When using up to 16 Tx and Rx interrupts, software might access the Tx and Rx interrupt bits in the legacy EICR, EICS,... registers.
 - **More than 16 Tx and Rx Interrupts** — When using more than 16 Tx and Rx interrupts, software must use EICS(1)...EICS(2), EIMS(1)...EIMS(2),... In the later case, software should avoid modifying the lower 16 bits in the EICS, EIMS... registers when it accesses the higher bits of these registers as follows:
 - EICR, EICS, EIMS and EIMC — When software programs the higher 16 bits of these registers, it should set their lower 16 bits to zero's keeping the EICR(1), EICS(1), EIMS(1) and EIMC(1) unaffected.
 - EIAM — When software programs the higher 16 bits, it should keep the lower 16 bits at their previous setting so the EIAM(1) is unaffected.
 - EIAC — When software programs the higher 16 bits, it should set the lower 16 bits to one's.



Single MSI-X vector — If the operating system allocates only a single MSI-X vector, the driver might use the non-MSI-X mapping method (setting the GPIE.Multiple_MSIX to 0b). In this case, the *INT_Alloc* field in the IVAR registers might define one of the lower 16 bits in the EICR register while using MSI-X vector 0. The IVAR_MISC should be programmed to MSI-X vector 0.

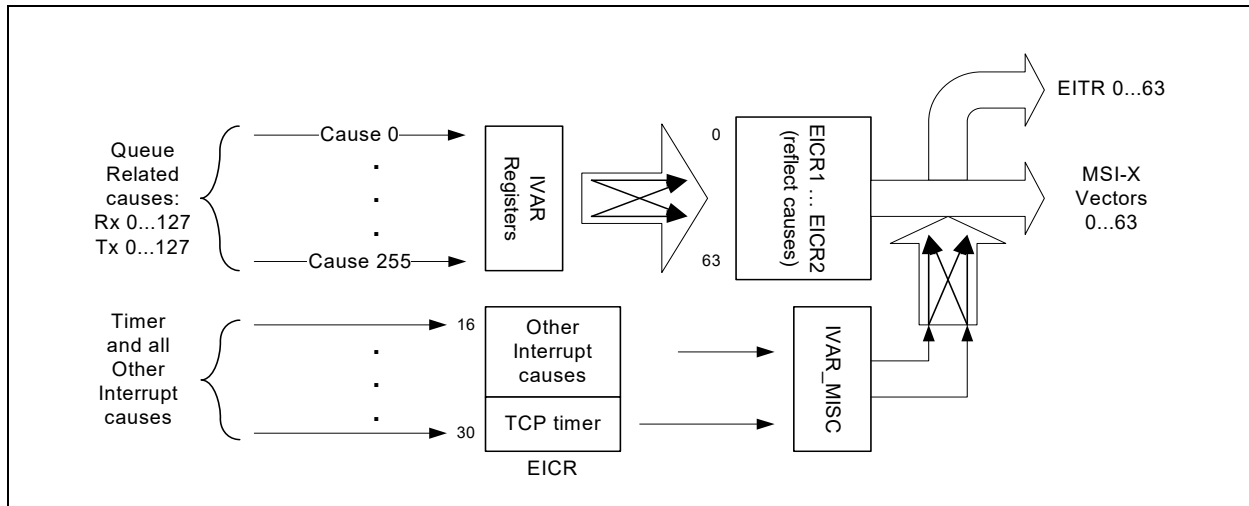


Figure 7.24. Cause Mapping in MSI-X Mode (non-IOV)

7.3.4.3 MSI-X Interrupts In IOV Mode

In IOV mode, interrupts must be implemented by MSI-X vectors. The Intel® Xeon® Processor D-1500 Product Family LAN controller supports up to 64 virtual functions VF(0...63). Each VF can generate up to three MSI-X vectors. The number of requested MSI-X vectors per VF is loaded from NVM in the *PCI_CNF2.MSI_X_VF_N* field. It is reflected in the *Table Size* field in the PCIe MSI-X capability structure of the VFs. In addition, the PF requires its own interrupts. The number of requested MSI-X vectors is loaded from NVM in the *PCI_CNF2.MSI_X_PF_N* field up to maximum of 64 MSI-X vectors. It is reflected in the *Table Size* field in the PCIe MSI-X capability structure.

7.3.4.3.1 MSI-X Vectors Used by Physical Function (PF)

PF is responsible for the timer and other interrupt causes that include the VM to PF mailbox cause (explained in the virtualization sections). These events are reflected in EICR[30:16] and MSI-X vectors are the same as the non-IOV mode (illustrated in Figure 7.22). When there are less than the maximum possible active VFs, some of the Tx and Rx queues can be associated with the PF. These queues can be used for the sake of additional VMs serviced by the hypervisor (the same as VMDq mode) or some Kernel applications handled by the hypervisor. Tx and Rx mapping to the IVAR registers is shown in Figure 7.23 and mapping to the EICR, EICR(1),...EICR(2) registers as well as the MSI-X vectors is shown in Figure 7.24. See Section 7.3.4.3.3 for MSI-X vectors mapping of PF and VFs to the EITR registers.

Note: Software should not assign MSI-X vectors in the PF to Tx and Rx queues that are assigned to other VFs. In the case that VFs become active after the PF used the relevant Tx and Rx queues, it is the responsibility of the PF driver to clear all pending interrupts of the associated MSI-X vectors.

7.3.4.3.2 MSI-X Vectors Used by Virtual Functions (VFs)

Each of the VFs in IOV mode is allocated separate IVAR(s) called VFIVAR registers, and a separate IVAR_MISC called VFIVAR_MISC register. The VFIVAR_MISC maps the mailbox interrupt of the VF to its VFEICR and the MSI-X vector. The VFIVAR registers map the Tx and Rx interrupts of the VF to its VFEICR and the MSI-X vector. The mapping is similar to the mapping in the PF as shown in [Figure 7.25](#) with the following comments:

- Each VF cannot have more than three MSI-X vectors. It has only three active bits in the VFEICR register while VFEICR.bit_num is associated with MSI-X vector (bit_num).
- The Tx and Rx interrupt can be mapped only to MSI-X 0 and MSI-X 1 (associated with VFEICR.0 and VFEICR.1).
- The mailbox interrupt can be mapped to any of the three MSI-X vectors. However, when all three of them are allocated by the operating system, software should map the mailbox to MSI-X 2 (associated with VFEICR.2). This rule should be kept since only VFEICR.0 and VFEICR.1 have ITR registers (VFEITR-0 and VFEITR-1).
- Association between the Tx and Rx queues and the VFIVAR registers is shown in the [Figure 7.25](#), [Figure 7.26](#) and [Figure 7.27](#) for IOV-64 (64 VFs), IOV-32 and IOV-16. The colored boxes in the figures show the mapping between VF Rx and Tx queues to VFIVAR registers while the dashed boxes show the physical IVAR registers and the associated physical Rx and Tx queues.

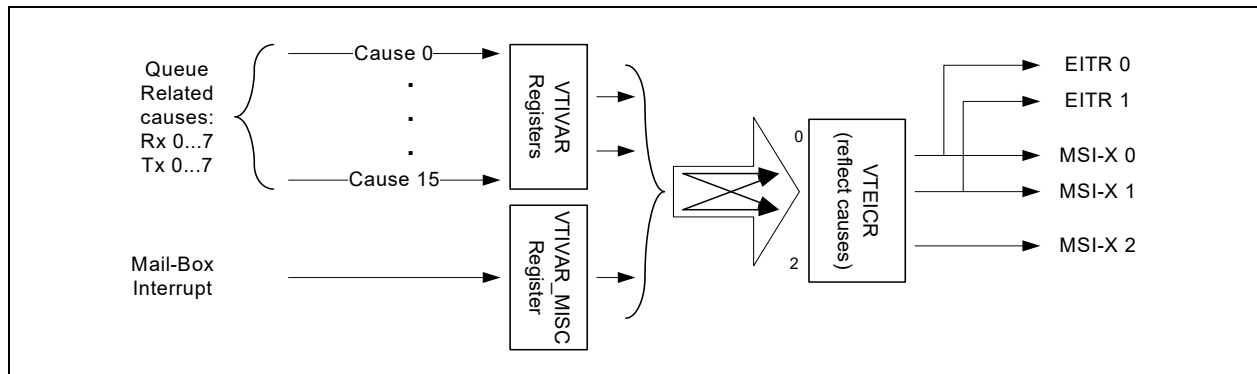


Figure 7.25. VF Interrupt Cause Mapping (MSI-X, IOV)

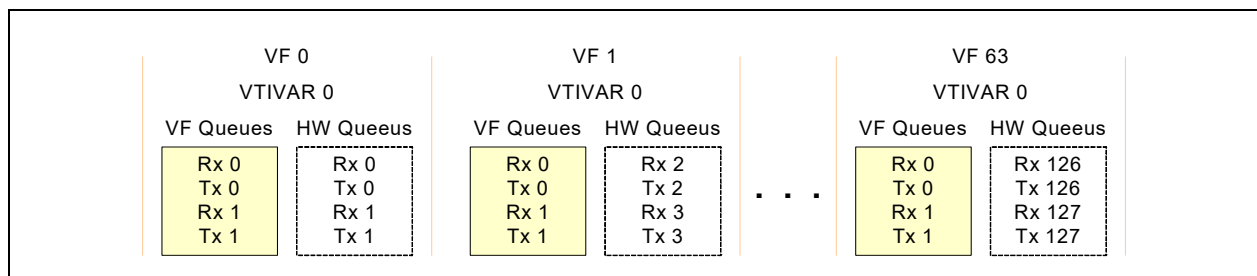


Figure 7.26. VF Mapping of Rx and Tx Queue to VFIVAR in 64 VFs Mode

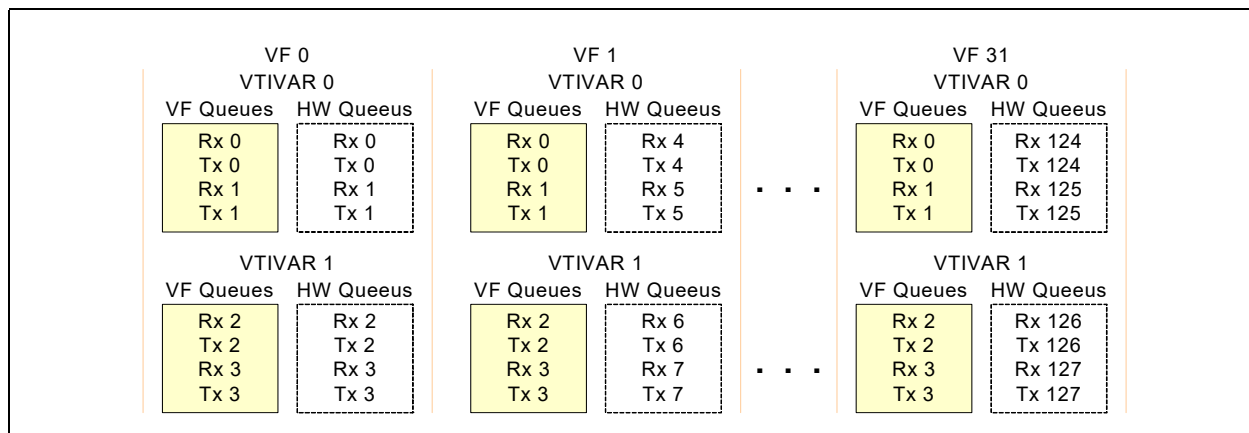


Figure 7.27. VF Mapping of Rx and Tx Queue to VFIVAR in 32 VFs Mode

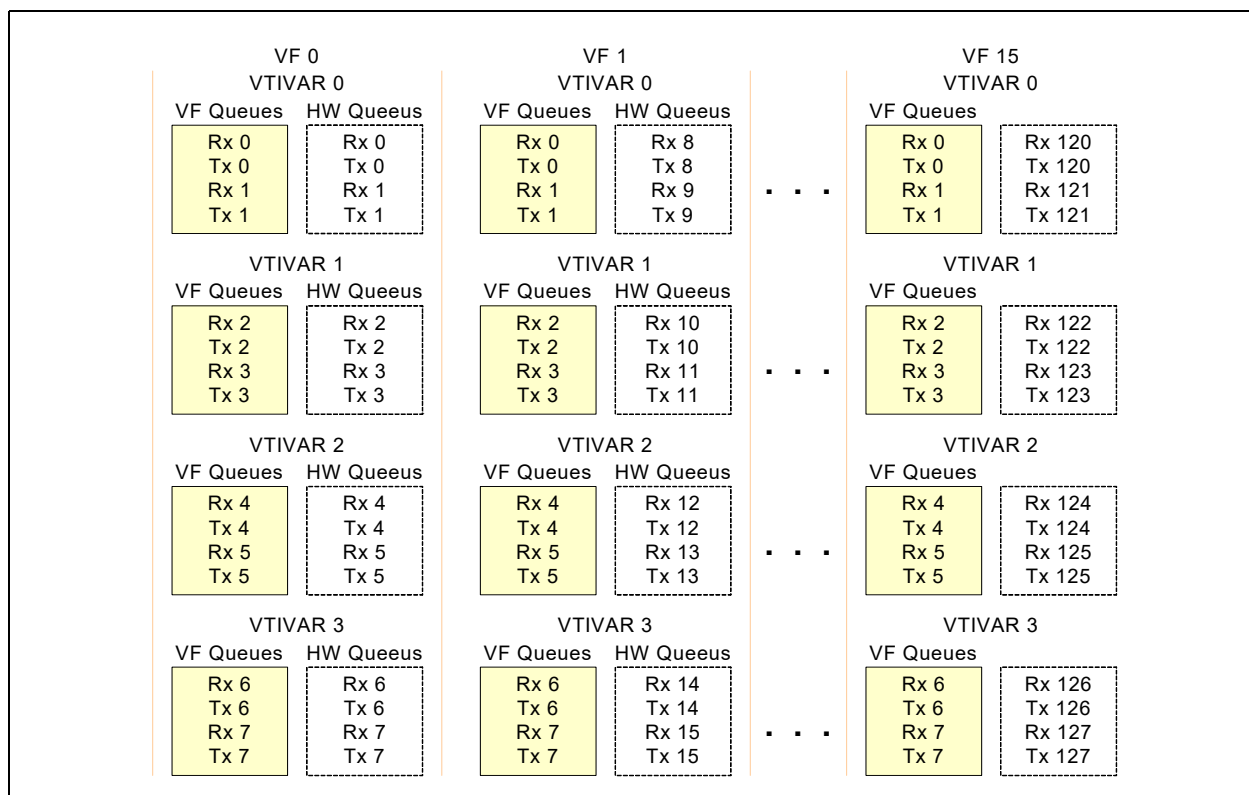


Figure 7.28. VF Mapping of Rx and Tx Queue to VFIVAR in 16 VFs Mode

7.3.4.3.3 MSI-X Vectors Mapping to EITR

EITR registers are aimed for Tx and Rx interrupt throttling. In IOV mode, the Tx and Rx queues might belong to either the PF or to the VFs. EITR(1...63) are multiplexed between the PF and the VFs as configured by the EITRSEL register. Figure 7.29 and Table 7.47 show the multiplexing logic and required software settings. For any active VF (starting from VF32 and above), software should program the matching bit in the EITRSEL to 1b. For any EITR that belongs to a VF, software should not map any interrupt causes in the PF to an MSI-X vector that is associated with the same EITR register.

Any RSCINT[n] register is associated with an MSI-X vector 'n'. As indicated above, the EITRSEL setting affects the MSI-X mapping. It also maps their associated RSCINT registers to either the PF or the VFs.

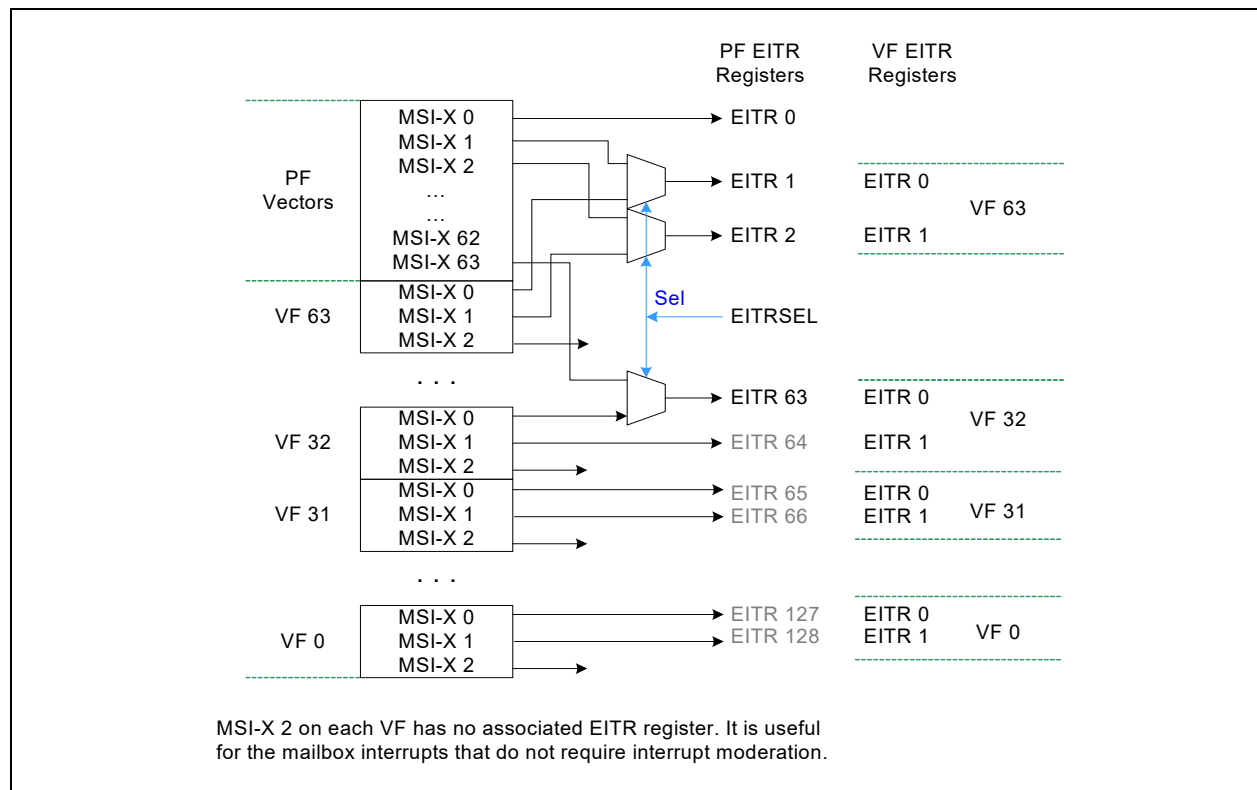


Figure 7.29. PF / VF MSI-X Vectors Mapping to EITR

Table 7.47. PF / VF MSI-X Vectors Mapping Table to EITR Registers

VM Active	EITRSEL.N Setting	MSI-X Routing to EITR
Non-IOV or VF(32...63) inactive	EITRSEL must be set to 0x0000	MSI-X(1...63) -> EITR(1...63)
VF(32) active	EITRSEL[0] must be set to 1b	VF(32) MSI-X(0) -> EITR(63)
VF(33) active	EITRSEL[1] must be set to 1b	VF(33) MSI-X(1) -> EITR(62) VF(33) MSI-X(0) -> EITR(61)
VF(34) active	EITRSEL[2] must be set to 1b	VF(34) MSI-X(1) -> EITR(60) VF(34) MSI-X(0) -> EITR(59)



Table 7.47. PF / VF MSI-X Vectors Mapping Table to EITR Registers

	...	
VF(62) active	EITRSEL[30] must be set to 1b	VF(62) MSI-X(1) -> EITR(4) VF(62) MSI-X(0) -> EITR(3)
VF(63) active	EITRSEL[31] must be set to 1b	VF(63) MSI-X(1) -> EITR(2) VF(63) MSI-X(0) -> EITR(1)

7.4 802.1q VLAN Support

The Intel® Xeon® Processor D-1500 Product Family LAN controller provides several specific mechanisms to support 802.1q VLANs:

- Optional adding (for transmits) and stripping (for receives) of IEEE 802.1q VLAN tags.
- Optional ability to filter packets belonging to certain 802.1q VLANs.

7.4.1 802.1q VLAN Packet Format

The following table compares an untagged 802.3 Ethernet packet with an 802.1q VLAN tagged packet:

802.3 Packet	#Octets	802.1q VLAN Packet	#Octets
DA	6	DA	6
SA	6	SA	6
Type/Length	2	802.1q Tag	4
Data	46-1500	Type/Length	2
CRC	4	Data	46-1500
		CRC*	4

Note: The CRC for the 802.1q tagged frame is re-computed, so that it covers the entire tagged frame including the 802.1q tag header. Also, maximum frame size for an 802.1q VLAN packet is 1522 octets as opposed to 1518 octets for a normal 802.3z Ethernet packet.


7.4.2 802.1q Tagged Frames

For 802.1q, the *Tag Header* field consists of four octets comprised of the Tag Protocol Identifier (TPID) and Tag Control Information (TCI); each taking two octets. The first 16 bits of the tag header makes up the TPID. It contains the protocol type that identifies the packet as a valid 802.1q tagged packet.

The two octets making up the TCI contain three fields as follows:

- User Priority (UP)
- Drop Eligible Indicator (DEI). Should be set to 0b for transmits. For receives, the device has the capability to filter out packets that have this bit set. See the *DEIEN* and *DEI* bits in the *VLNCTRL*
- VLAN Identifier (VID)

Octet 1		Octet 2
UP	DEI	VID



7.4.3 Transmitting and Receiving 802.1q Packets

Since the 802.1q tag is only four bytes, adding and stripping of tags can be done completely in software. (In other words, for transmits, software inserts the tag into packet data before it builds the transmit descriptor list, and for receives, software strips the 4-byte tag from the packet data before delivering the packet to upper layer software). However, because adding and stripping of tags in software adds overhead for the host, the Intel® Xeon® Processor D-1500 Product Family LAN controller has additional capabilities to add and strip tags in hardware. See [Section 7.4.3.1](#) and [Section 7.4.3.2](#).

7.4.3.1 Adding 802.1q Tags on Transmits

The inner VLAN header can be added by software in one of the following methods:

- The header is included in the transmit data buffers.
- Software might instruct the Intel® Xeon® Processor D-1500 Product Family LAN controller to insert an 802.1q VLAN tag on a per-packet basis. If the *VLE* bit in the transmit descriptor is set to 1b, then the Intel® Xeon® Processor D-1500 Product Family LAN controller inserts a VLAN tag into the packet that it transmits over the wire. The Tag Protocol Identifier — TPID (VLAN Ether Type) field of the 802.1q tag comes from the *DMATXCTL.VT*, and the Tag Control Information (TCI) of the 802.1q tag comes from the *VLAN* field of the legacy transmit descriptor or the *VLAN Tag* field of the advanced data transmit descriptor.
- In IOV mode, the priority tag, DEI and VLAN ID can be taken from the *PFVMVIR* (see details in [Section 7.7.4.2](#)).

7.4.3.2 Stripping 802.1q Tags on Receives

Software might instruct the Intel® Xeon® Processor D-1500 Product Family LAN controller to strip 802.1q VLAN tags from received packets. The policy whether to strip the VLAN tag is configurable per queue.


If the *RXDCTL.VME* bit for a given queue is set to 1b, and the incoming packet is an 802.1q VLAN packet (that is, its *Ethernet Type* field matched the *VLNCTRL.VET*), then the Intel® Xeon® Processor D-1500 Product Family LAN controller strips the 4-byte VLAN tag from the packet, and stores the TCI in the *VLAN Tag* field of the receive descriptor.

The Intel® Xeon® Processor D-1500 Product Family LAN controller also sets the *VP* bit in the receive descriptor to indicate that the packet had a VLAN tag that was stripped. If the *RXDCTL.VME* bit is not set, the 802.1q packets can still be received if they pass the receive filter, but the VLAN tag is not stripped and the *VP* bit is not set. If *PFODE.HIDE_VLAN* is set, the VLAN tag is stripped as above, but the *VLAN Tag* field and the *Status.VP* bit in the Rx descriptor are cleared.

7.4.4 802.1q VLAN Packet Filtering

VLAN filtering is enabled by setting the *VLNCTRL.VFE* bit to 1b. If enabled, hardware compares the *Type* field of the incoming packet to a 16-bit field in the VLAN Ether Type (VET) register. If the *VLAN Type* field in the incoming packet matches the VET register, the packet is then compared against the VLAN Filter Table Array for acceptance.

The VLAN filter register VTFA, is a vector array composed of 4096 bits. The VLAN ID (VID) is a 12-bit field in the VLAN tag that is used as an index pointer to this vector. If the VID in a received packet points to an active bit (set to 1b), the packet matches the VLAN filter. The 4096-bit vector is comprised of 128 x 32 bit registers. The upper 7 bits of the VID selects one of the 128 registers while the lower 5 bits map the bit within the selected register.



Two other bits in the VLNCTRL register, *DEIEN* and *DEI*, are also used in conjunction with 802.1q VLAN filtering operations. *DEIEN* enables the comparison of the value of the *DEI* bit in the 802.1q packet to the Receive Control register *DEI* bit as acceptance criteria for the packet.

Note: The *VFE* bit does not effect whether the VLAN tag is stripped. It only effects whether the VLAN packet passes the receive filter.

7.4.5 Double VLAN and Single VLAN Support

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports a mode where all received and sent packets have at least one VLAN tag in addition to the regular tagging that might optionally be added. In this document, when a packet carries two VLAN headers, the first header is referred to as an outer VLAN and the second header as an inner VLAN header (as listed in the table that follows). This mode is used for systems where the near end switch adds the outer VLAN header containing switching information. This mode is enabled by the following configuration:

- This mode is activated by setting the DMATXCTL.GDV and the *Extended VLAN* bit in the CTRL_EXT register.
- The Ethertype of the VLAN tag used for the additional VLAN is defined in the *VET EXT* field in the EXVET and EXVET_T registers.

7.4.5.1 Cross Functionality with Manageability

The Intel® Xeon® Processor D-1500 Product Family LAN controller does not provide any stripping or adding VLAN header(s) to manageability packets. Therefore, packets that are directed to/from the manageability controller should include the VLAN headers as part of the Rx/Tx data. The manageability controller should know if the Intel® Xeon® Processor D-1500 Product Family LAN controller is set to double VLAN mode as well as the VLAN Ethertype(s).

Table 7.48. Double VLAN Packet Format

MAC Address	Outer VLAN	Inner VLAN	L2 Payload	Ethernet CRC
-------------	------------	------------	------------	--------------

7.4.5.2 Transmit Functionality

7.4.5.2.1 Transmit Functionality on the Outer VLAN Header

- A packet with a single VLAN header is assumed to have only the outer VLAN.
- The outer VLAN header must be added by software as part of the Tx data buffers.
- Hardware does not relate to the outer VLAN header other than the capability of skipping it for parsing inner fields.
- Hardware expects that any transmitted packet (see the disclaimer that follows) has at least the outer VLAN added by software. For any offload that hardware might provide in the transmit data path, hardware assumes that the outer VLAN is present. For those packets that an outer VLAN is not present, any offload that relates to inner fields to the Ethertype might not be provided.

7.4.5.2.2 Transmit Functionality on the Inner VLAN Header

- Inner VLAN insertion is handled as described in [Section 7.4.3.1](#).
- Hardware identifies and skips the VLAN header for parsing inner fields.
- Pool Filtering — Destination pool(s) and anti-spoofing functionality is based on the Ethernet MAC address and inner VLAN (if present) as described in [Section 7.7.3.4](#) and [Section 7.7.4.2](#).



7.4.5.3 Receive Handling of Packets with VLAN Header(s)

A received frame is analyzed for the existence of the outer and inner VLAN headers. The procedure is as follows:

```
Check the Ethertype against the outer VLAN ID. If match then
{
    Check the next Ethertype against the inner VLAN ID. If match
    {
        This is the double VLAN case. Process both outer and inner VLANs as described
        below
    }
    Else
    {
        Only an outer VLAN exists. Process the outer as described below. Assume no inner
        VLAN
    }
}
Else
{
    Check the Ethertype against the inner VLAN ID. If match
    {
        This is the case of an inner VLAN only. Handle the frame as an unknown frame -
        device does
        not provide any offloads
    }
    Else
    {
        This is the case of no VLAN. Process the frame (ignore outer and inner VLAN
        processing)
    }
}
```

7.4.5.3.1 Receive Functionality on the Outer VLAN Header

- Hardware checks the Ethertype of the outer VLAN header against the programmed value in the EXVET register. VLAN header presence is indicated in the Status.VEXT bit in the Rx descriptor.
- The outer VLAN header is posted as is to the receive data buffers.

7.4.5.3.2 Receive Functionality on the Inner VLAN Header

- Hardware checks the Ethertype of the inner VLAN header against the programmed value in the VLNCTRL.VET. VLAN header presence is indicated in the Status.VP bit in the Rx descriptor.
- L2 packet filtering is based on the VLAN ID in the inner VLAN header.
- Pool Filtering — Destination pool(s) are defined by the Ethernet MAC address and inner VLAN (if presence) as described in [Section 7.7.3.3](#).
- Inner VLAN tag striping is handled as described in [Section 7.4.3.2](#).



7.4.5.4 Packets With no VLAN Headers in Double VLAN Mode

There are some cases when packets might not carry any VLAN headers, even when extended VLAN is enabled. A few examples for packets that might not carry any VLAN header are: flow control, LACP, LLDP, GMRP, and optional 802.1x packets. When it is expected to transmit untagged packets by software in Double VLAN Mode the software must not enable VLAN anti-spoofing and VLAN validation nor transmit to receive switching.

7.4.5.4.1 Transmit Functionality

Transmit offload functionality — Software should not enable any offload functions.

7.4.5.4.2 Receive Functionality

Receive offload functionality — pool and queue are selected by the Ethernet MAC address or *ETQF/ETQS* registers. Filtering to host and manageability remains functional.

7.4.5.5 Packets With two VLAN Headers not in Double VLAN Mode

When the *Extended VLAN* bit in the *CTRL_EXT* register and *DMATXCTL.GDV* are not set, hardware expects that Rx and Tx packets might not carry a VLAN header or a single VLAN header. Hardware does not relate to the programming of the *VET_EXT* field in the *EXVET* register. Tx and Rx handling of packets with double VLAN headers is unexpected.

7.4.6 E-tag and VLAN

In some systems an additional external tag (E-tag) can be present before the VLAN. This section describes the support for VLANs in presence of external tags. This mode is used for systems where the device adds a tag to identify a subsystem (usually a VM) and the near end switch adds a tag indicating the destination subsystem. External tags may be present on part of the packets and missing in others.


7.4.6.1 Transmit Functionality

7.4.6.1.1 Transmit Functionality on the External Tag

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports insertion of an external tag from a per pool register (*PFVMTIR* - [Section 8.2.2.16.17](#)) and a VLAN tag from a per pool register (*PFVMVIR* - [Section 8.2.2.16.16](#)) or from the descriptor as indicated by the *VLE* bit. The following options are supported:

VLAN Source/External Tag Source	Hardware Register - PFVMTIR/PFVMVIR	Embedded in Packet
Hardware register - PFVMVIR	Supported	Not supported
Descriptor	Supported	Not supported
Embedded in Packet	Supported	Supported

After the tags are inserted in the packet, the Intel® Xeon® Processor D-1500 Product Family LAN controller uses the VLAN tag and the Ethertype as part of the forwarding decision as described in [Section 7.7.3.4](#). In a packet with at most one external tag and one VLAN, the VLAN and Ethertype will be identified correctly.



7.4.6.2 Receive Handling of Packets With External Tags

The parsing of outer tags is described in Section 7.1.2.1.2. External tags can be extracted from the packet according to the *PFODE.STRIP_TAG* field. Inner tag extraction is handled as described above (Section 7.4.3.2).

7.4.6.2.1 Packet Priority in presence of external tags

The user priority used to define the traffic class of a packet is always taken from the inner VLAN.

7.4.6.3 Cross Functionality with Manageability

The Intel® Xeon® Processor D-1500 Product Family LAN controller does not provide any stripping or adding VLAN header(s) to manageability packets. Therefore, packets that are directed to/from the manageability controller should include the L2 headers as part of the Rx/Tx data. The manageability controller should know if it is expected to add/receive an external tag as part of the packet.

7.4.6.4 Packet User Priority (802.1P) Bits Handling

The user priority bits can be used by the device for multiple purposes:

1. Defining the TC to which the traffic is associated (Section 7.6.3.1).
2. Defining which packets get a timestamp in buffer (Section 7.1.6.2).
3. Defining if a packet matches and Ethertype filter (Section 7.1.3.3).

For all these purposes, the UP bits are taken from the outermost tag with UP bits (E-tag or VLAN), unless the *VLNCTRL.UP_FIRST_TAG_EN* bit is cleared, in which case, it is always taken from the inner VLAN.

Table 7.49 lists the tag from which the user priority is extracted from the packets.

Table 7.49. UP Extraction Rules

Packet type	VLNCTRL.UP_FIRST_TAG_EN = 1	VLNCTRL.UP_FIRST_TAG_EN = 0
Un-tagged Packet	User priority = 0	User priority = 0
Packet with a single VLAN	The user priority field from the VLAN header.	Single VLAN mode: The user priority field from the VLAN header. Double VLAN mode: User priority = 0
Packet with a VLAN and an outer tag (outer VLAN or E-tag)	The user priority field from the outer VLAN header	The user priority field from the inner VLAN header.
Packet with E-tag only	The user priority field from the E-tag	User priority = 0

7.5 TLP Processing Hints (TPH)

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports the TPH capability defined in the PCI Express specification. It does not support Extended TPH requests.

On the PCIe link existence of a TLP Process Hint (TPH) is indicated by setting the TH bit in the TLP header. Using the PCIe TLP Steering Tag (ST) and Processing Hints (PH) fields, the Intel® Xeon® Processor D-1500 Product Family LAN controller can provide hints to the root complex about the



destination (socket ID) and about data access patterns (locality in Cache), when executing DMA memory writes or read operations. Supply of TLP processing hints facilitates optimized processing of transactions that target memory space.

To enable TPH usage:

1. For a given function, the *TPH Requester Enable* bit in the PCIe configuration TPH Requester Control register should be set.
2. Appropriate TPH Enable bits in TPH_RXCTRL or TPH_TXCTRL registers should be set.
3. Processing hints should be programmed in the *TPH_CTRL.Desc_PH* and *TPH_CTRL.Data_PH* Processing hints (PH) fields.
4. Steering information should be programmed in the CPUID fields in the TPH_RXCTRL and TPH_TXCTRL registers.

The Processing hints (PH) and Steering Tags (ST) are set according to the characteristics of the traffic as described in [Table 7-50](#).

Note: In order to enable TPH usage, all the memory reads are done without setting any of the byte enable bits.

Note: Per queue, the TPH features are exclusive. Software can enable the TPH feature for a given queue.

7.5.1 Steering Tag and Processing Hint Programming

The following table lists how the steering tag (socket ID) and processing hints are generated and how TPH operation is enabled for different types of DMA traffic.

Table 7-50. Steering tag and Processing hint programming

Traffic type	ST Programming	PH value	Enable
Transmit descriptor write back or head write back	TPH_TXCTRL.CPUID ¹	TPH_CTRL.Desc_PH ²	Tx Descriptor Writeback TPH EN field in TPH_TXCTRL.
Receive data buffers write	TPH_RXCTRL.CPUID ¹	TPH_CTRL.Data_PH ³	Rx Header TPH EN or Rx Payload TPH EN fields in TPH_RXCTRL.
Receive descriptor writeback	TPH_RXCTRL.CPUID ¹	TPH_CTRL.Desc_PH ²	Rx Descriptor Writeback TPH EN field in TPH_RXCTRL.
Transmit descriptor fetch	TPH_TXCTRL.CPUID ⁴	TPH_CTRL.Desc_PH ²	Tx Descriptor Fetch TPH EN field in TPH_TXCTRL.
Receive descriptor fetch	TPH_RXCTRL.CPUID ²	TPH_CTRL.Desc_PH ²	Rx Descriptor fetch TPH EN field in TPH_RXCTRL.
Transmit packet read	TPH_TXCTRL.CPUID ²	TPH_CTRL.Data_PH ³	Tx Packet TPH EN field in TPH_TXCTRL.

1. the software device driver should always set bits [7:3] to zero and place socket ID in bits [2:0].
2. Default is 00b (Bidirectional data structure).
3. Default is 10b (Target).
4. the hints are always zero.

5. During stress when auto learn is enabled some hints learnings might be discarded.



7.6 Time SYNC (IEEE1588 and 802.1AS)

7.6.1 Overview

IEEE 1588 addresses the clock synchronization requirements of measurement and control systems. The protocol supports system-wide synchronization accuracy in the sub-microsecond range with minimal network and local clock computing resources. The protocol is spatially localized and allows simple systems to be installed and operate.

The IEEE802.1AS standard specifies the protocol used to ensure that synchronization requirements are met for time sensitive applications, such as audio and video, across Bridged and Virtual Bridged Local Area Networks consisting of LAN media where the transmission delays are almost fixed and symmetrical; for example, IEEE 802.3 full duplex links. This includes the maintenance of synchronized time during normal operation and following addition, removal, or failure of network components and network re-configuration. It specifies the use of IEEE 1588 specifications where applicable.

Note: The time sync mechanism activation is possible in full-duplex mode only and is not supported at 100 Mb/s link speed.

7.6.2 Flow and Hardware/Software Responsibilities

The operation of a Precision Time Protocol (PTP) enabled network is divided into two stages: initialization and time synchronization.

At the initialization stage, every master-enabled node starts by sending sync packets that include the clock parameters of its clock. Upon receipt of a sync packet, a node compares the received clock parameters to its own and if the received parameters are better, then this node moves to a slave state and stops sending sync packets. While in slave state, the node continuously compares the incoming packet to its currently chosen master and if the new clock parameters are better, than the master selection is transferred to this master clock. Eventually the best master clock is chosen. Every node has a defined time-out interval that if no sync packet was received from its chosen master clock it moves back to a master state and starts sending sync packets until a new best master clock (PTP) is chosen.

The time synchronization stage is different to master and slave nodes. If a node is in a master state it should periodically send a sync packet that is time stamped by hardware on the TX path (as close as possible to the PHY). After the sync packet, a Follow_Up packet is sent that includes the value of the time stamp kept from the sync packet. In addition, the master should time stamp Delay_Req packets on its Rx path and return to the slave that sent the time stamp value using a Delay_Response packet. A node in a slave state should time stamp every incoming sync packet and if it came from its selected master, software uses this value for time offset calculation. In addition, it should periodically send Delay_Req packets in order to calculate the path delay from its master. Every sent Delay_Req packet sent by the slave is time stamped and kept. With the value received from the master with Delay_Response packet, the slave can now calculate the path delay from the master to the slave. The synchronization protocol flow and the offset calculation are shown in [Figure 7.30](#).

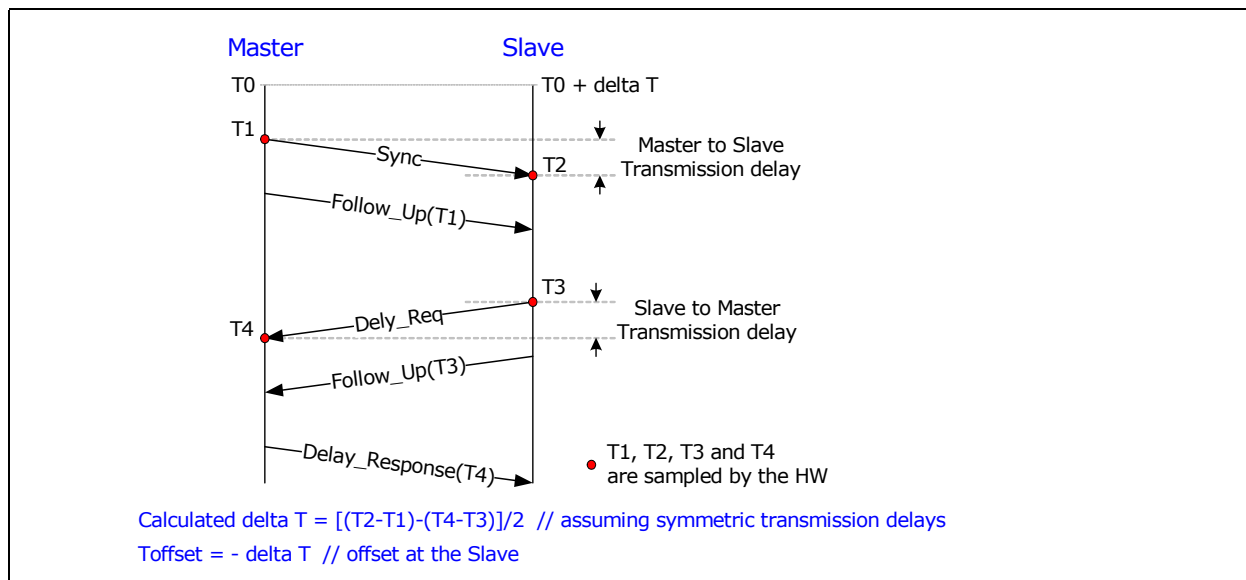


Figure 7.30. Sync Flow and Offset Calculation


Hardware responsibilities are:

1. Identify the packets that require time stamping.
2. Time stamp the packets on both Rx and Tx paths.
3. Store the time stamp value for software.
4. Keep the system time in hardware and give a time adjustment service to software.
5. Maintain auxiliary features related to the system time.

Software responsibilities are:

1. Manageability controller protocol execution, which means defining the node state (master or slave) and selection of the master clock if in slave state.
2. Generate PTP packets, consume PTP packets.
3. Calculate the time offset and adjust the system time using a hardware mechanism for that.
4. Enable configuration and usage of the auxiliary features.

Action	Responsibility	Node Role
Generate a sync packet with time stamp notification in the descriptor.	Software	Master
Time stamp the packet and store the value in registers (T1).	Hardware	Master
Time stamp incoming sync packet, store the value in register	Hardware	Slave
Read the time stamp from register put in a Follow_Up packet and send.	Software	Master
Once received, the Follow_Up store T2 from registers and T1 from Follow-up packet.	Software	Slave
Generate a Delay_Req packet with time stamp notification in the descriptor.	Software	Slave
Time stamp the packet and store the value in registers (T3).	Hardware	Slave
Time stamp incoming Delay_Req packet, store the value in register	Hardware	Master
Read the time stamp from register and send back to slave using a Delay_Response packet.	Software	Master



Action	Responsibility	Node Role
Once received, the Delay_Response packet calculate offset using T1, T2, T3 and T4 values.	Software	Slave

7.6.2.1 Time Sync Indications in Rx and Tx Packet Descriptors

Some indications need to be transferred between software and hardware regarding PTP packets. On the Tx path, software should set the 1588 bit in the Tx packet descriptor (bit 9). On the Rx path, hardware has two indications to transfer to software, one is to indicate that this packet is a PTP packet (whether time stamp is taken or not). This is also for other types of PTP packets needed for management of the protocol and this bit is set only for the L2 type of packets (the PTP packet is identified according to its Ethertype). PTP packets have the *L2 Packet* bit in the packet type field set (bit 11 in the receive descriptor) and the Ethertype matches the filter number set by software in the *ETOF* registers to filter PTP packets. The UDP type of PTP packets don't need such indication since the port number (319 for event and 320 all the rest PTP packets) directs the packets toward the time sync application. The second indication is *RDESC.STATUS.TS* (bit 16) to indicate to software that time stamp was taken for this packet. Software needs to access the time stamp registers to get the time stamp values.


7.6.3 Hardware Time Sync Elements

All time sync hardware elements are reset to their initial values upon Master reset. Upon change in link speed some of the time sync parameters should be changed accordingly.

7.6.3.1 System Time Structure and Mode of Operation

The *SYSTIME* is a 96 bit register is composed of: *SYSTIMR*, *SYSTIMEL* and *SYSTIMEH* registers: The *SYSTIMR* register holds the sub nsec fraction, the *SYSTIMEL* register holds the nsec fraction and the *SYSTIMEH* register holds the sec fraction of the time (note that the upper two bits of the *SYSTIMEL* register are always zero while the max value of this register is 999,999,999 dec).

- **Initial Setting** - Setting the initial time is done by direct write access to the *SYSTIME* register. Software should set first the *SYSTIMEL* and then set the *SYSTIMEH*. Setting the *SYSTIMR* is meaningless while it represents sub nsec units. It is recommended to disable the timer at programming time as follows:
 - **Run Time** - During run time the *SYSTIME* timer value in the *SYSTIMEH*, *SYSTIMEL* and *SYSTIMR* registers, is updated periodically each 12.5 nS clock cycle according to the following formula:
 - Define: $INC_TIME == 12.5 \text{ nsec} +/- TIMINCA.Incvalue * 2^{-32} \text{ nsec}$. Add or subtract the *TIMINCA.Incvalue* is defined by *TIMINCA.ISGN* (while 0b means Add and 1b means Subtract)
 - Then: $SYSTIME = SYSTIME + INC_TIME$
- **Reading the *SYSTIME* register** by software is done by the following sequence:
 - Read the *SYSTIMEL* register
 - Read the *SYSTIMEH* register.
- **Dynamic update of *SYSTIME* registers** is done by using the *TIMADJ* registers by the following flow:
 - Write the *Tadjust* value and its *Sign* to the *TIMADJ* register (the *Sign* bit indicates if the *Tadjust* value should be added or subtracted)
 - Following the write access to the *TIMADJ* register, the hardware repeats the following two steps at each 12.5 nsec clock as long as the *Tadjust* > 0.
 - $SYSTIME = SYSTIME + INC_TIME +/- 1 \text{ nsec}$. Add or subtract 1 nsec is defined by *TIMADJ.Sign* (while '0' means Add and '1' means Subtract)
 - $Tadjust = Tadjust - 1 \text{ nsec}$



- As previously shown, the time adjustment might take multiple clocks. Software might write a new value to the TIMADJ register before the hardware completed the previous adjustment. In such a case, the new value written by software, overrides the above equation. If such a race is not desired, the software could check that the previous adjustment is completed by one of the following methods:
 - Wait enough time before accessing the TIMADJ register which guarantees that the previous update procedure is completed.
 - Poll the matched TSICR.TADJ flag which is set by the hardware each time the update procedure is completed.

7.6.3.2 Time Stamping Mechanism

The time stamping logic is located as close as possible to the PHY. Figure 7.31 shows the exact point in time where the time value is captured by the hardware relative to the packet content. This is to reduce delay uncertainties originated from implementation differences. As the time stamp is sampled at a very late phase in the data path, the Intel® Xeon® Processor D-1500 Product Family LAN controller does not insert it to the transferred packet. Instead, the Intel® Xeon® Processor D-1500 Product Family LAN controller supports the two-step operation as follows for Tx and Rx.

Note: As the time stamping logic is located at the MAC/PHY interface, there is no timestamping of packets when MAC loopback is activated.

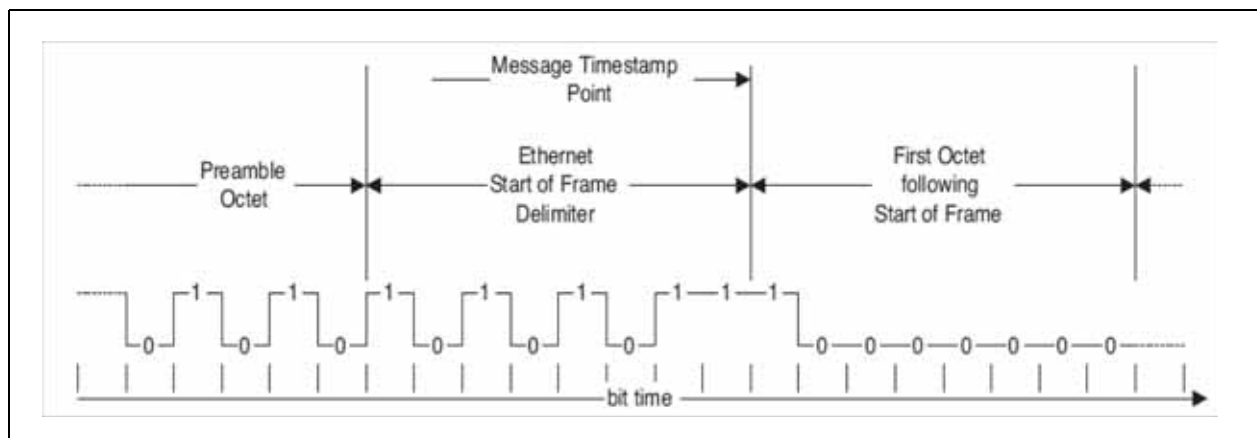


Figure 7.31. Time Stamp Point

7.6.3.2.1 Single Tx Time Stamping

The time stamp logic is activated if enabled by the TSYNCTXCTL.EN bit and the time stamp bit in the packet descriptor is set. In this case, hardware captures the packet's transmission time in the TXSTMPL and TXSTMPH registers. Software is responsible to read the transmission time and append it in the Follow_Up packet as shown in Figure 7.30.

7.6.3.2.2 Single Rx Time Stamping

On the Rx, this logic parses the traversing frame. If it is matching the message type defined in RXMTRL register and the TSYNCRXCTL.TYPE field, and timestamping is enabled by the TSYNCRXCTL.EN bit the reception time stamp is stored in the RXSTMPL and RXSTMPH registers. In addition, the TS status bits is reported in the Rx descriptor to identify that a time stamp was taken for this packet (stored in the RXSTMPL and RXSTMPH registers).



Note: The time stamp values are locked in the RXSTMPL and RXSTMPH registers until software accesses them. As long as software does not read these registers, hardware does not capture the time stamp of further Rx packets. In order to avoid potential deadlocks, it is recommended that software read the Rx time stamp registers at some time after sync or Delay_Req packets are expected. It would overcome erroneous cases on which the hardware latches a packet reception time while the packet's content was not posted properly to the software. Master software must not initiate consecutive sync requests before the previous response is received.

Note: If TSYNCRXCTL.TYPE == 4 (sample time stamp of all packets), then the timestamps are not locked in the RXSTMPL and RXSTMPH registers and each new packet timestamp is stored in these registers.

7.6.3.2.3 Multiple Rx Time Stamping

Packets that are identified to be time stamped by hardware are also indicated by the *TS* or *TSIP* flags in the receive descriptors. If the *TS* flag is set, the packet reception time is sampled by the hardware in the RXSTMPL/H registers (see Single Rx Time Stamping). These registers are locked until the software reads its value. If the *TSIP* flag is set, the packet reception time is posted to the packet buffer in host memory. For more information about posting the time stamp in the receive buffer refer to Section 7.1.6.2.

7.6.4 Hardware Time Sync Elements

All time sync hardware is initialized as defined in the registers section upon MAC reset. The time sync logic is enabled if the *TSAUXC.Disable systime* flag is cleared.

The 1588 logic includes multiple registers larger than 32 bits which are indicated as xxxL (Low portion - LS) and xxxH (High portion - MS). When software accesses these registers (either read or write) it should access first the xxxL register (LS) and only then the xxxH register (MS). Accessing the xxxH might impact the hardware functionality which should be triggered only after both portions of the register are valid.

7.6.4.1 Target Time

The two target time registers TRGTTIML/H0 and TRGTTIML/H1 enable generating a time triggered event to external hardware using one of the SDP pins according to the setup defined in the TSSDP and TSAUXC registers. Each target time register is structured the same as the SYSTIMEL/H registers. If the value of SYSTIMEL/H is equal or larger than the value of the TRGTTIML/H registers, a change in level or a pulse is generated on the matched SDP outputs.

7.6.4.1.1 SYSTIM Synchronized Level Change Generation on SDP Pins

To generate a level change on one of the SDP pins when System Time (SYSTIM) reaches a pre-defined value, the software device driver should do the following:

1. Select a specific SDP pin by setting the *TSSDP.TS_SDPx_EN* flag to 1b (while 'x' is 0, 1, 2 or 3).
2. Assign a target time register to the selected SDP by setting the *TSSDP.TS_SDPx_SEL* field to 00b or 01b if level change should occur based on TRGTTIML/H0 or TRGTTIML/H1, respectively.
3. Define the selected SDPx pin as output, by setting the appropriate *SDPx_IODIR* bit (while 'x' is 0, 1, 2, or 3) in the ESDP register.
4. Define that this SDP is used for TimeSync function by setting the appropriate *SDPx_NATIVE* bit (while 'x' is 0, 1, 2, or 3) in the ESDP register. If used with SDP1, clear *ESDP.SDP1_Function* field



5. Program the target time TRGTTIML/Hx (while 'x' is 0b or 1b) to the required event time.
6. Program the TRGTTIML/Hx to "Level Change" mode by setting the *TSAUXC.PLSG* bit to 0b and *TSAUXC.EN_TTx* bit to 1b (while 'x' is 0b or 1b).
7. To make this a one time operation, the *TSAUXC.DIS_TS_CLEAR* field should be cleared.

When the SYSTIMEL/H registers becomes equal or larger than the selected *TRGTTIML/H* registers, the selected SDP changes its output level.

7.6.4.1.2 SYSTIM Synchronized Pulse Generation on SDP Pins


An output pulse can be generated by using one of the target time registers to define the beginning of the pulse and the other target time registers to define the pulse completion time. To generate a pulse on one of the SDP pins when System Time (SYSTIM) reaches a pre-defined value, the driver should do the following:

1. Select a specific SDP pin by setting the *TSSDP.TS_SDPx_EN* flag to 1b (while 'x' is 0, 1, 2 or 3).
2. Set *TSSDP.TS_SDPx_SEL* field to 00b to define that the TRGTTIML/H0 register defines the start of pulse time and *TRGTTIML/H1* register defines the end of pulse time.
3. Define the selected SDPx pin as output, by setting the appropriate *SDPx_IODIR* bit (while 'x' is 0, 1, 2, or 3) in the ESDP register.
4. Define that this SDP is used for TimeSync function by setting the appropriate *SDPx_NATIVE* bit (while 'x' is 0, 1, 2, or 3) in the ESDP register. If used with SDP1, clear *ESDP.SDP1_Function* field.
5. Program the target time TRGTTIML/Hx (while 'x' is 0b or 1b) to the required event time.
6. *TRGTTIML/H0* should be set to a lower value than *TRGTTIML/H1*.
7. Program TRGTTIML/H0 defined by the *TSSDP.TS_SDPx_SEL* to start of pulse mode by setting the *TSAUXC.PLSGO* bit to 1b and *TSAUXC.EN_TT0* bit to 1b (while 'x' defines the SDP used). The *TRGTTIML/H1* register *should be set* to indicate the end of the pulse and *TSAUXC.EN_TT1* bit should be set to 1b.
8. To make this a one time operation, the *TSAUXC.DIS_TS_CLEAR* field should be cleared.
9. When the SYSTIMEL/H registers becomes equal or larger than the TRGTTIML/H0 registers the selected SDP changes its level. Then, when the SYSTIMEL/H registers becomes equal or larger than TRGTTIML/H1 registers (that define the trailing edge of the pulse), the selected SDP changes its level back.

7.6.4.1.3 Synchronized Output Clock on SDP Pins

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports driving a programmable Clock on the SDP pins (up to two output clocks). The output clocks generated are synchronized to the global System time registers (SYSTIM). The Target Time registers (TRGTTIML/H0 or TRGTTIML/H1) can be used for the clock output generation. To start an clock output on one of the SDP pins when System Time (SYSTIM) reaches a pre-defined value, the driver should do the following:

1. Select a specific SDP pin by setting the *TSSDP.TS_SDPx_EN* flag to 1b (while 'x' is 0, 1, 2 or 3).
2. Select the target time register for a selected SDP, by setting the *TSSDP.TS_SDPx_SEL* field to 10b or 11b if output clock should occur based on TRGTTIML/H0 or TRGTTIML/H1, respectively.
3. Program the matched *FREQOUT0/1* register to define clock half cycle time.
4. Define the selected SDPx pin as output, by setting the appropriate *SDPx_IODIR* bit (while 'x' is 0, 1, 2, or 3) in the ESDP register.
5. Define that this SDP is used for TimeSync function by setting the appropriate *SDPx_NATIVE* bit (while 'x' is 0, 1, 2, or 3) in the ESDP register. If used with SDP1, clear *ESDP.SDP1_Function* field.



6. If the output clock should start at a specific time, the *TSAUXC.STO/1* flag should be set to 1b and the matched *TRGTTIML/Hx* should be set to the required start time.
7. Enable the clock operation by setting the relevant *TSAUXC.EN_CLK0/1* bit to 1b.

An interrupt can be generated from the clock output generated by the device by setting the relevant *TSAUXC.EN_TTO/1* bit to 1b and by setting the *TSAUXC.DIS_TS_CLEAR* to allow it to work continuously. The clock out drives initially a logical zero level on the selected SDP. If the *TSAUXC.STO/1* flag is cleared, it happens instantly when setting the *TSAUXC.EN_CLK0/1* bit. Otherwise it happens when *SYSTIM* is equal or larger than the *TRGTTIM*. Since then, the hardware repeats endlessly the following two steps:

1. Increment the used *TRGTTIML/Hx* by *FREQOUT*.
2. When *SYSTIM* is equal or larger than the *TRGTTIM*, the SDP reverts its output level.

Note: When clearing *TSAUXC.EN_CLK0/1* while *TSAUXC.EN_TTO/1* was set in order to generate interrupts, clear the matching *TSAUXC.EN_TTO/1* too to avoid one unexpected toggle.

7.6.4.2 Time Stamp Events

Upon a change in the input level of one of the SDP pins that was configured to detect Time stamp events using the *TSSDP* register or upon a setting of one of the *TSAUXC.SAMP_AUTx* fields, a time stamp of the system time is captured into one of the two auxiliary time stamp registers (*AUXSTMPL/H0* or *AUXSTMPL/H1*). Software enables the timestamp of input event as follow:

1. Define the sampled SDP on AUX time 'x' ('x' = 0b or 1b) by setting the *TSSDP.AUXx_SDP_SEL* field while setting the matched *TSSDP.AUXx_TS_SDP_EN* bit to 1b.
2. Set also the *TSAUXC.EN_TSx* bit ('x' = 0b or 1b) to 1b to enable timestamping.

Following a transition on the selected SDP, the hardware does the following:

1. The *SYSTIM* registers (low and high) are latched to the selected *AUXSTMP* registers (low and high)
2. The selected *AUTTO* or *AUTT1* flags are set in the *TSICR* and *TSAUXC* registers. If the *AUTT* interrupt is enabled by the *TSIM* register and the 1588 interrupts are enabled by the *Time_Sync* flag in the *ICR* register then an interrupt is asserted as well. After the hardware reports that an event time was latched, the software should read the latched time in the selected *AUXSTMP* registers. Software should read first the Low register and only then the High register. Reading the high register clears the relevant *TSAUC.AUTT_x* field and releases the registers to sample a new event.

The software device driver might initiate a sampling of the current time by setting one of the *TSAUXC.SAMP_AUTx* fields. When one of these bits is written the hardware The *SYSTIM* registers (low and high) are latched to the selected *AUXSTMP* registers (low and high) and the selected *AUTTO* or *AUTT1* flags are set in the *TSAUXC* register. Once the software device driver reads the selected *AUXSTMP* registers as described above, the hardware clears the relevant *TSAUXC.AUTT_x* field and releases the registers to sample a new event.

7.6.5 Time SYNC Interrupts


Time Sync related interrupts can be generated by programming the *TSICR* and *TSIM* registers. The *TSICR* register logs the interrupt cause and the *TSIM* register enables masking specific *TSICR* bits. Occurrence of a Time Sync interrupt sets the *ICR.Time_Sync* interrupt bit.



7.6.6 PTP Packet Structure

The time sync implementation supports both the 1588 V1 and V2 PTP frame formats. The V1 structure can come only as UDP payload over IPv4 while the V2 can come over L2 with its Ethertype or as a UDP payload over IPv4 or IPv6. The 802.1AS uses only the layer 2 V2 format. Note that PTP frame structure over UDP is not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller for IP tunneling packets.

Offset in Bytes	V1 Fields	V2 Fields	
Bits	7 6 5 4 3 2 1 0	7 6 5 4 3 2 1 0	
0	<i>versionPTP</i>	transportSpecific ¹	<i>messageType</i>
1		Reserved	<i>versionPTP</i>
2	versionNetwork	messageLength	
3			
4	Subdomain	SubdomainNumber	
5		Reserved	
6		flags	
7			
8			
9			
10			
11		Correction Field	
12			
13			
14			
15			
16			
17		reserved	
18			
19			
20	messageType	<i>Source Port ID</i>	
21	Source communication technology		
22	<i>Sourceuuid</i>		
23			
24			
25			
26	sourceportid		
27			
28			
29			
30	<i>sequenceId</i>	<i>sequenceId</i>	
31			
32	<i>control</i>	control	
33	reserved	logMessagePeriod	



Offset in Bytes	V1 Fields	V2 Fields
Bits	7 6 5 4 3 2 1 0	7 6 5 4 3 2 1 0
34	fалgs	n/a
35		

1. Should all be zero.

Note: Only the fields with the bold italic format colored red are of interest to hardware.

Ethernet (L2)	VLAN (Optional)	PTP Ethertype	PTP message
Ethernet (L2)	IP (L3)	UDP	PTP message

When a PTP packet is recognized (by Ethertype or UDP port address) on the Rx side the version should be checked if it is V1 then the control field at offset 32 should be compared to message field in the RXMTRL register described in [Section 8.2.2.16.21](#), otherwise the byte at offset 0 should be used for comparison to the rest of the needed field are at the same location and size for both V1 and V2.

Enumeration	Value
PTP_SYNC_MESSAGE	0
PTP_DELAY_REQ_MESSAGE	1
PTP_FOLLOWUP_MESSAGE	2
PTP_DELAY_RESP_MESSAGE	3
PTP_MANAGEMENT_MESSAGE	4
reserved	5-255

Message ID	Message Type	Value (hex)
PTP_SYNC_MESSAGE	Event	0
PTP_DELAY_REQ_MESSAGE	Event	1
PTP_PATH_DELAY_REQ_MESSAGE	Event	2
PTP_PATH_DELAY_RESP_MESSAGE	Event	3
Unused		4-7
PTP_FOLLOWUP_MESSAGE	General	8
PTP_DELAY_RESP_MESSAGE	General	9
PTP_PATH_DELAY_FOLLOWUP_MESSAGE	General	A
PTP_ANNOUNCE_MESSAGE	General	B
PTP_SIGNALLING_MESSAGE	General	C
PTP_MANAGEMENT_MESSAGE	General	D
Unused		E-F

If V2 mode is configured in TSAUXC register ([Section 8.2.2.16.23](#)), then time stamp should be taken on PTP_PATH_DELAY_REQ_MESSAGE and PTP_PATH_DELAY_RESP_MESSAGE for any value in the message field in the RXMTRL register described at [Section 8.2.2.16.21](#).

Note: A PTP over UDP packet encapsulated in a VXLAN or NVGRE tunnel will still be identified as a PTP packet.



7.6.6.1 Time Sync Packets identification configuration

The following table summarize the setting needed to identify Time Sync packets.

Table 7-51. Enabling Receive Timestamp

Functionality	Register	Field	Setting Options
Enable receive timestamp in register	TSYNCRXCTL	En	En = 1b (must be set in all the following options).
Sampled V1 Control value	RXMTRL	CTRLT	The CTRLT defines the recognized V1 Control field. This field must be defined if V1 packets recognition is required. 0x0 = PTP_SYNC_MESSAGE
Sampled V2 MessageType value	RXMTRL	MSGT	The MSGT defines the recognized V2 MessageType field. This field must be defined if V2 packets recognition is required. 0x0 = Sync
Enable all packets for timestamp	TSYNCRXCTL	Type	Type equals to 100b enables sampling all packets. Useful only when posting the timestamp to the packet buffer in host memory, enabled per queue by the SRRCTL[n].Timestamp.
Enable L2 1588 packets for timestamp sampling	TSYNCRXCTL	Type	Type equals to 000b or 010b enable V2 packets with MessageType equals to MSGT Type equals to 101b enable all V2 packets with Message Type bit 3 zero (means any event packets)
	ETQF[n]	EType Filter enable IEEE_1588_TIME_STAMP	The EType on one of the enabled ETQF registers (Filter enable is '1') should be set to the 1588 EtherType (equals to 0x88F7) and IEEE_1588_TIME_STAMP field should be set.
Enable 1588 packets over UDP for timestamp sampling	TSYNCRXCTL	Type	Type equals to 001b enables V1 packets with Control field equals to CTRLT parameter Type equals to 010b enables V2 packets with MessageType fields equals to MSGT parameter Type equals to 101b enables all V2 packets with Message Type bit 3 zero (which means any event packets)
	RXMTRL	UDPT	Defines the UDP port to use for V1 detection
Define specific receive queue for the L2 1588 packets	ETQF[n]	Rx Queue Queue Enable	Setting the "Queue Enable" on the same ETQF register as above, the receive queue is defined by the "Rx Queue" field.

7.7 Virtualization

7.7.1 Overview

I/O virtualization is a mechanism that can be used to share I/O resources among several consumers. For example, in a virtual system, multiple operating systems are loaded and each operates as though the entire system's resources were at its disposal. However, for the limited number of I/O devices, this presents a problem because each operating system might be in a separate memory domain and all the data movement and device management has to be done by a Virtual Machine Monitor (VMM). VMM access adds latency and delay to I/O accesses and degrades I/O performance. Virtualized devices are designed to reduce the burden of the VMM by making certain functions of an I/O device shared among multiple guest operating systems or a Virtual Machine (VM), thereby allowing each VM direct access to the I/O device.

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports two modes of operations of virtualized environments:

1. Direct assignment of part of the port resources to different guest operating systems using the PCI SIG SR IOV standard. Also known as native mode or pass through mode. This mode is referenced as IOV mode throughout this section.



2. Central management of the networking resources by an IOVM or by the VMM. Also known as software switch acceleration mode. This mode is referred to as VMDq2 mode in this section.

The virtualization offloads capabilities provided by the Intel® Xeon® Processor D-1500 Product Family LAN controller apart from the replication of functions defined in the PCI SIG IOV specification are part of VMDq2.

A hybrid model, where part of the VMs are assigned a dedicated share of the port and the rest are serviced by an IOVM is also supported. However, in this case the offloads provided to the software switch might be more limited. This model can be used when parts of the VMs run operating systems for which VF drivers are available and thus can benefit from an IOV and others that run older operating systems for which VF drivers are not available and are serviced by an IOVM. In this case, the IOVM is assigned one VF and receives all the packets with Ethernet MAC addresses of the VMs behind it.

The following section describes the support the Intel® Xeon® Processor D-1500 Product Family LAN controller provides for these modes.

This section assumes a single-root implementation of IOV and no support for multi-root.

7.7.1.1 Direct Assignment Model

The direct assignment support in the Intel® Xeon® Processor D-1500 Product Family LAN controller is built according to the following model of the software environment.

It is assumed that one of the software drivers sharing the port hardware behaves as a master driver (Physical Function or PF driver). This driver is responsible for the initialization and the handling of the common resources of the port. All the other drivers (Virtual Function drivers or VF drivers) might read part of the status of the common parts but cannot change them. The PF driver might run either in the VMM or in some service operating system. It might be part of an IOVM or part of a dedicated service operating system.

In addition, part of the non time-critical tasks are also handled by the PF driver. For example, access to CSR through the I/O space or access to the configuration space are available only through the master interface. Time-critical CSR space like control of the Tx and Rx queue or interrupt handling is replicated per VF, and directly accessible by the VF driver.


Note: In some systems with a thick hypervisor, the service operating system might be an integral part of the VMM. For these systems, each reference to the service operating system in the sections that follow refer to the VMM.

7.7.1.1.1 Rationale

The direct assignment model enables each of the VMs to receive and transmit packets with minimum of overhead. Non time-critical operations such as initialization and error handling can be done via the PF driver. In addition, it is important that the VMs can operate independently with minimal disturbance. It is also preferable that the VM interface to hardware should be as close as possible to the native interface in non-virtualized systems in order to minimize the software development effort.

The main time critical operations that require direct handling by the VM are:

- Maintenance of the data buffers and descriptor rings in host memory. In order to support this, the DMA accesses of the queues associated to a VM should be identified as such on the PCIe using a different requester ID.
- Handling of the hardware ring (tail bump and head updates)
- Interrupts handling



The capabilities needed to provide independence between VMs are:

- Per VM reset and enable capabilities
- Tx rate control
- Allocating separate CSR space per VM. This CSR space is organized as close as possible to the regular CSR space to enable sharing of the base driver code.

Note: The rate control and VF enable capabilities are controlled by the PF.

7.7.1.2 System Overview

The following diagrams show the various elements involved in the I/O process in a virtualized system. [Figure 7.32](#) shows the flow in software VMDq2 mode and [Figure 7.33](#) shows the flow in IOV mode.

This section assumes that in IOV mode, the driver on the guest operating system is aware that it operates in a virtual system (para-virtualized) and there is a channel between each of the VM drivers and the PF driver allowing message passing such as configuration request or interrupt messages. This channel can use the mailbox system implemented in the Intel® Xeon® Processor D-1500 Product Family LAN controller or any other means provided by the VMM vendor.

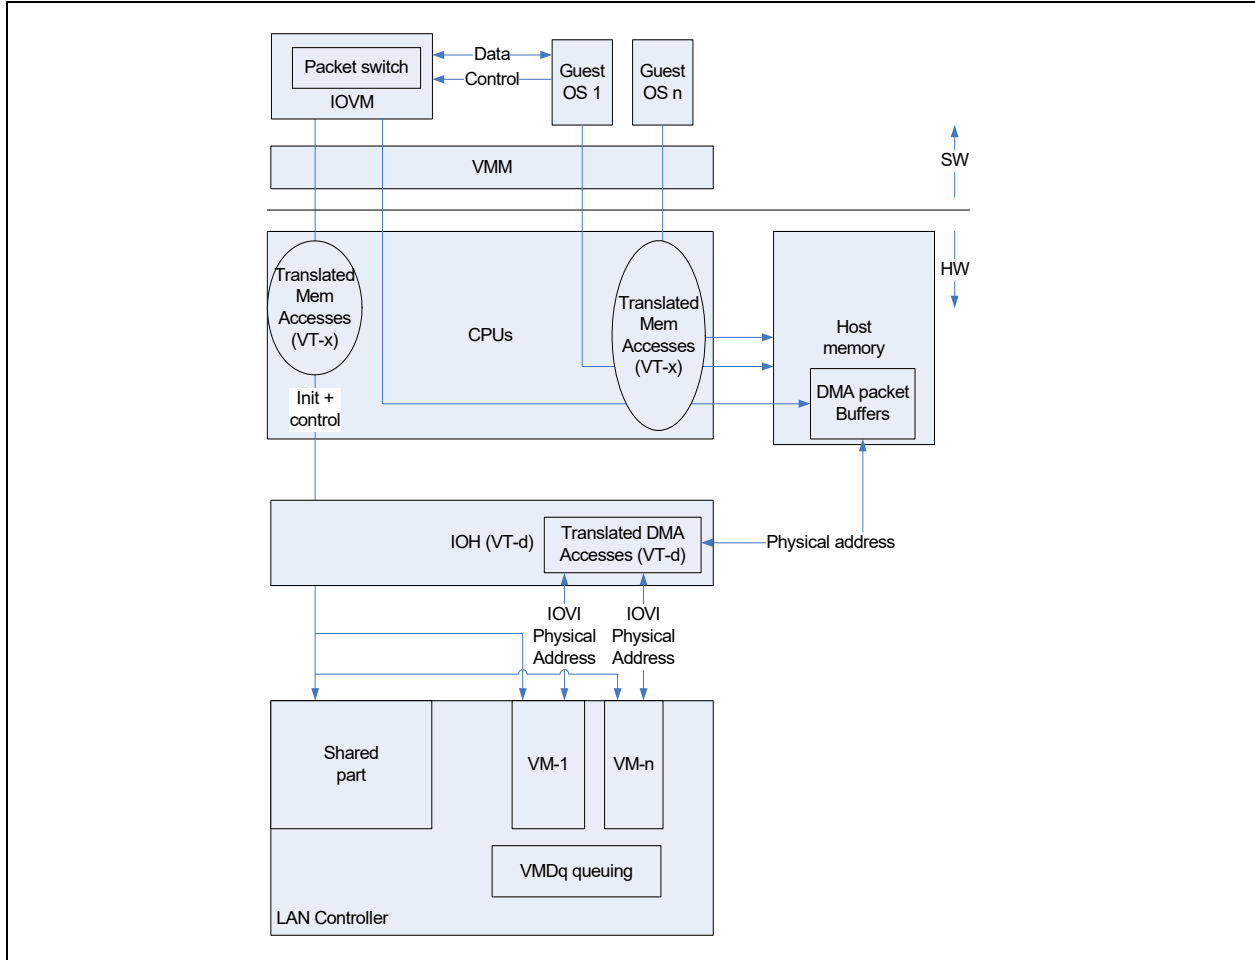


Figure 7.32. System Configuration for VMDq2 Mode

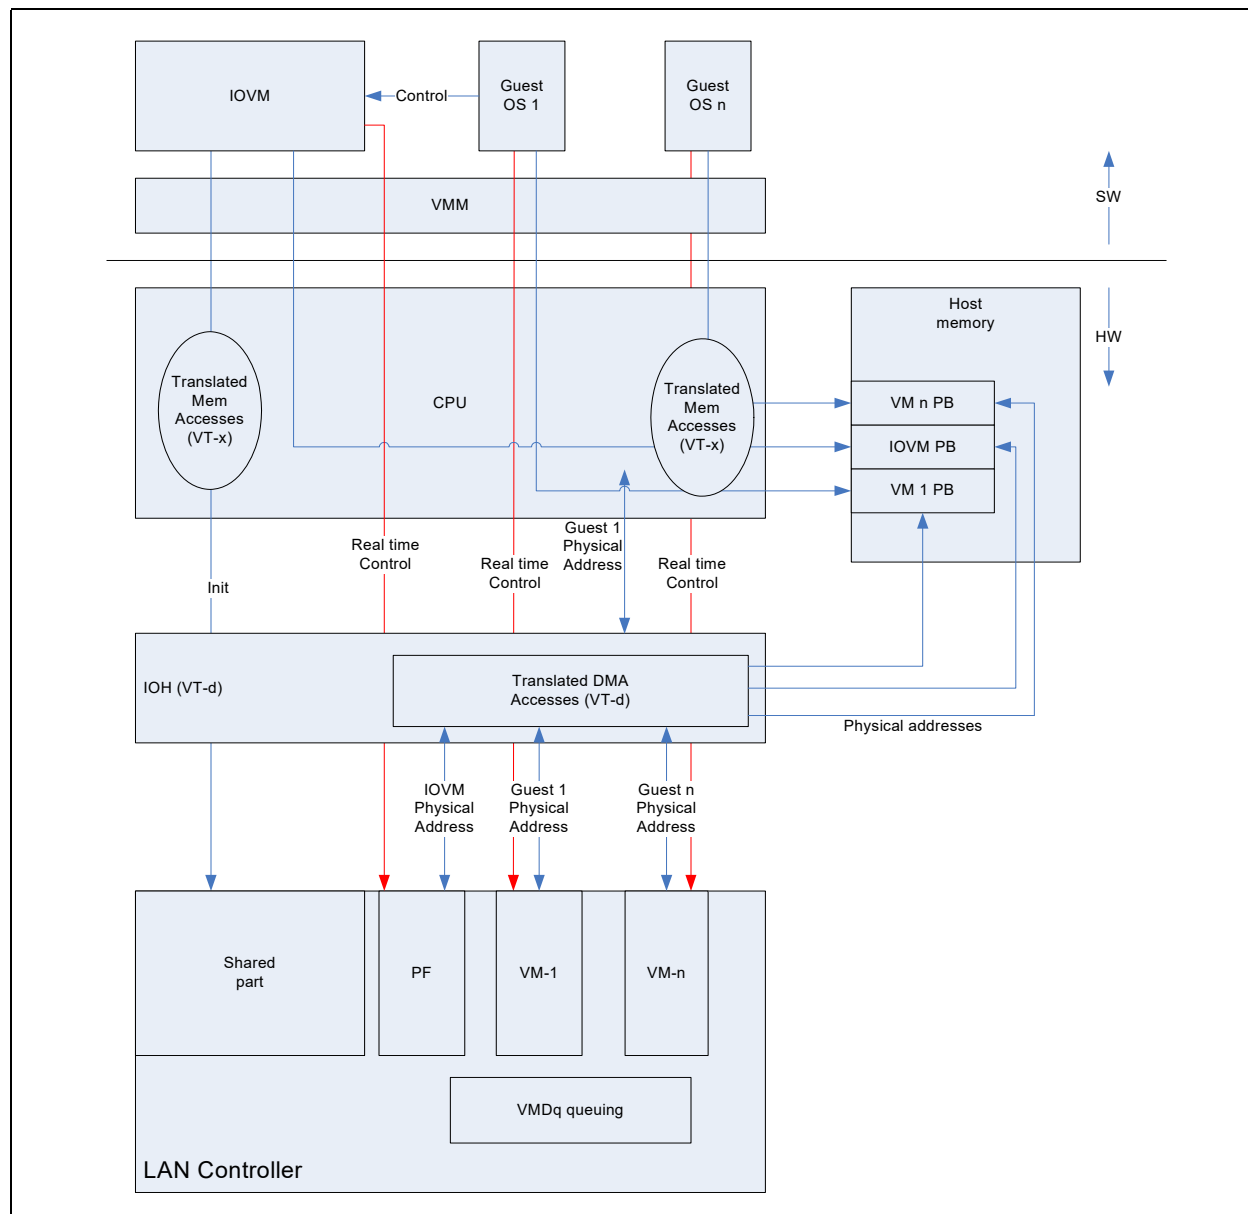


Figure 7.33. System Configuration for IOV Mode



7.7.2 PCI-SIG SR-IOV Support

7.7.2.1 SR-IOV Concepts

SR-IOV defines the following entities in relation to I/O virtualization:

- Virtual Image (VI): Part of the I/O resources are assigned to a VM.
- I/O Virtual Intermediary (IOVI) or I/O Virtual Machine (IOVM): A special VM that owns the physical device and is responsible for the configuration of the physical device.
- Physical function (PF): A function representing a physical instance — One port for the Intel® Xeon® Processor D-1500 Product Family LAN controller. The PF driver is responsible for the configuration and management of the shared resources in the function.
- Virtual Function (VF): A part of a PF assigned to a VI.

7.7.2.2 Configuration Space Replication

The SR-IOV specification defines a reduced configuration space for the virtual functions. Most of the PCIe configuration of the VFs comes from the PF.

This section describes the expected handling of the different parts of the configuration space for virtual functions. It deals only with the parts relevant to the Intel® Xeon® Processor D-1500 Product Family LAN controller.

Details of the configuration space for virtual functions can be found in [Section 9.3](#).

7.7.2.2.1 Legacy PCI Configuration Space

The legacy configuration space is allocated to the PF only and emulated for the VFs. A separate set of BARs and one bus master enable bit is allocated in the SR-IOV capability structure in the PF and is used to define the address space used by the entire set of VFs.

All the legacy error reporting bits are emulated for the VF. See [Section 7.7.2.4](#) for details.

7.7.2.2.2 Memory BARs Assignment

The SR-IOV specification defines a fixed stride for all the VF BARs, so that each VF can be allocated part of the memory BARs at a fixed stride from a basic set of BARs. In this method, only two decoders per replicated BAR per PF are required and the BARs reflected to the VF are emulated by the VMM.

The only BARs that are useful for the VFs are BAR0 and BAR3, so only those are replicated. The following table lists the existing BARs and the stride used for the VFs:

Table 7.52. BARs in the Intel® Xeon® Processor D-1500 Product Family LAN Controller (64-bit BARs)

BAR	Type	Usage	Requested Size per VF (=Stride)
0, 1	Mem	CSR space	Maximum (16 KB, page size). For page size see Section 9.2.4.4.8 for more details.
2	n/a	Not used	n/a
3, 4	Mem	MSI-X	Maximum (16 KB, page size).
5	n/a	Not used	n/a

BAR0 of the VFs are a sparse version of the original PF BAR and include only the register relevant to the VF. For more details see [Section 7.7.2.7](#).



The following figure shows the different BARs in an IOV-enabled system:

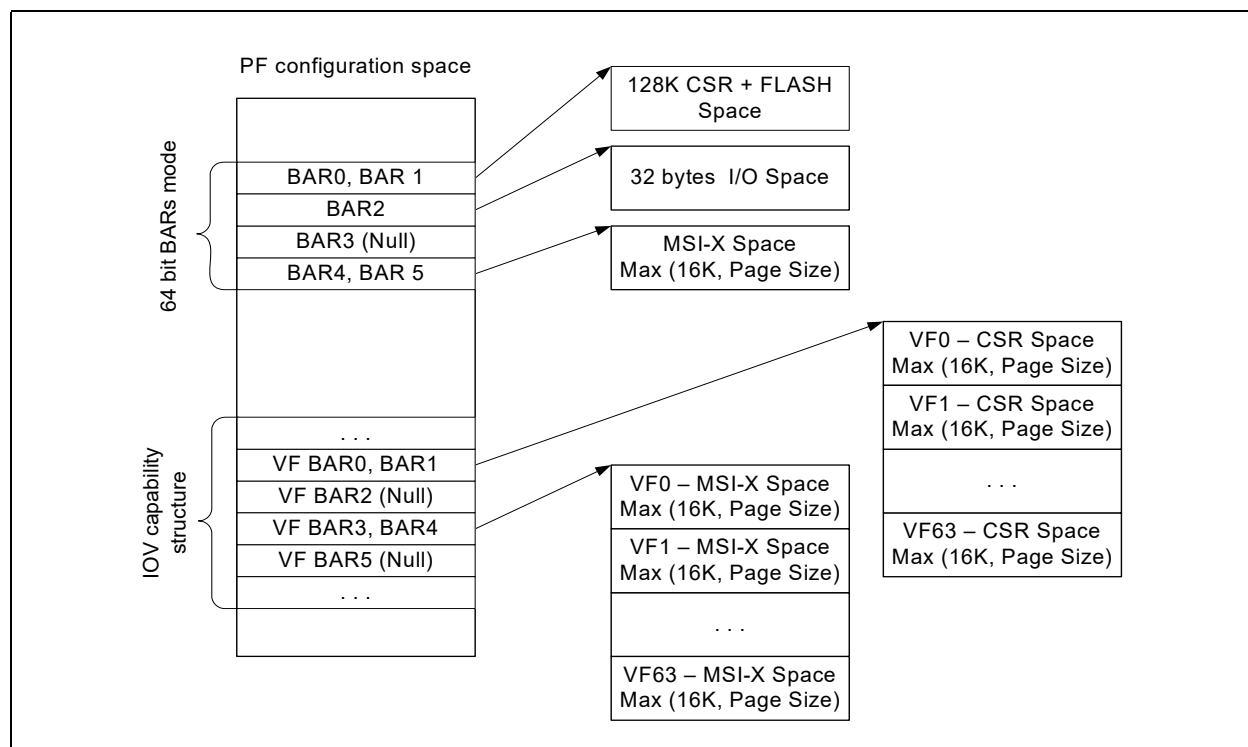


Figure 7.34. BARs in an IOV-enabled System

7.7.2.2.3 PCIe Capability Structure

The PCIe capability structure is shared between the PF and the VFs. The only relevant bits that are replicated are:

1. Transaction pending
2. Function Level Reset (FLR). See Section 7.7.2.3 for details.

7.7.2.2.4 MSI and MSI-X Capabilities


Both MSI and MSI-X are implemented in the Intel® Xeon® Processor D-1500 Product Family LAN controller. MSI-X vectors can be assigned per VF. MSI is not supported for the VFs.

7.7.2.2.5 VPD Capability

VPD is implemented only once and is accessible only from the PF.

7.7.2.2.6 Power Management Capability

The Intel® Xeon® Processor D-1500 Product Family LAN controller does not support power management per VF. The power management registers exist for each VF, but only the D0 power state is supported.



7.7.2.2.7 Serial ID

Serial ID capability is not exposed in VFs.

7.7.2.2.8 Error Reporting Capabilities (Advanced and Legacy)

All the bits in this capability structure are implemented only for the PF. Note that the VMs see an emulated version of this capability structure. See [Section 7.7.2.4](#) for details.

7.7.2.3 FLR Capability

The *FLR* bit is required per VF. Setting of this bit resets only a part of the logic dedicated to the specific VF and does not influence the shared part of the port. This reset should disable the queues, disable interrupts and the stop receive and transmit process per VF.

Setting the PF *FLR* bit resets the entire function.

7.7.2.4 Error Reporting

Error reporting includes legacy error reporting and Advanced Error Reporting (AER) or role-based capability.

The legacy error management includes the following functions:

1. Error capabilities enablement. These are set by the PF for all the VFs. Narrower error reporting for a given VM can be achieved by filtering of the errors by the VMM. This includes:
 - a. SERR# Enable
 - b. Parity Error Response
 - c. Correctable Reporting Enable
 - d. Non-Fatal Reporting Enable
 - e. Fatal Reporting Enable
 - f. UR Reporting Enable
2. Error status in the configuration space. These should be set separately for each VF. This includes:
 - a. Master Data Parity Error
 - b. Signaled Target Abort
 - c. Received Target Abort
 - d. Master Abort
 - e. SERR# Asserted
 - f. Detected Parity Error
 - g. Correctable Error Detected
 - h. Non-Fatal Error Detected
 - i. Unsupported Request Detected

AER capability includes the following functions:

1. Error capabilities enablement. The *Error Mask*, and *Severity* bits are set by the PF for all the VFs. Narrower error reporting for a given VM can be achieved by filtering of the errors by the VMM. These includes:
 - a. Uncorrectable Error Mask Register
 - b. Uncorrectable Error Severity Register



- c. Correctable Error Mask Register
 - d. ECRC Generation Enable
 - e. ECRC Check Enable
2. Non-Function Specific Errors Status in the configuration space.
 - a. Non-Function Specific Errors are logged in the PF
 - b. Error logged in one register only
 - c. VI avoids touching all VFs to clear device level errors
 - d. The following errors are not function specific
 - All Physical Layer errors
 - All Link Layer errors
 - ECRC Fail
 - UR, when caused by no function claiming a TLP
 - Receiver Overflow
 - Flow Control Protocol Error
 - Malformed TLP
 - Unexpected Completion
 3. Function Specific Errors Status in the configuration space.
 - a. Allows Per VF error detection and logging
 - b. Help with fault isolation
 - c. The following errors are function specific
 - Poisoned TLP received
 - Completion Timeout
 - Completer Abort
 - UR, when caused by a function that claims a TLP
 - ACS Violation
 4. Error logging. Each VF has it's own header log.
 5. Error messages. In order to ease the detection of the source of the error, the error messages should be emitted using the requester ID of the VF in which the error occurred.

7.7.2.5 Alternative Routing ID (ARI) and IOV Capability Structures

In order to allow more than eight functions per end point without requesting an internal switch, as usually needed in virtualization scenarios, the PCI-SIG defines the ARI capability structure. This is a new capability that enables an interpretation of the *Device* and *Function* fields as a single identification of a function within the bus. In addition, a new structure used to support the IOV capabilities reporting and control is defined. Both structures are described in sections [Section 9.2.4.3](#) and [Section 9.2.4.4](#). Refer to the following section for details on the Requester ID (RID) allocation to VFs.

7.7.2.6 RID Allocation

RID allocation of the VF is done using the *Offset* field in the IOV structure. This field should be replicated per VF and is used to do the enumeration of the VFs.



Each PF includes an offset to the first associated VF. This pointer is a relative offset to the Bus/Device/Function (BDF) of the first VF. The *Offset* field is added to PF's requester ID to determine the requester ID of the next VF. An additional field in the IOV capability structure describes the distance between two consecutive VF's requester IDs.

7.7.2.6.1 Bus Device Function Layout

7.7.2.6.1.1 ARI Mode

ARI allows interpretation of the device ID part of the RID as part of the function ID inside a device. Thus, a single device can span up to 256 functions. In order to ease the decoding, the least significant bit of the function number points to the physical port number. The *Next* bits indicate the VF number. The following table lists the VF RIDs.

The layout of RIDs used by the Intel® Xeon® Processor D-1500 Product Family LAN controller is reported to the operating system via the PCIe IOV capability structure. See [Section 9.2.4.4.6](#).

Table 7.53. RID Per VF — ARI Mode

Port	VF#	B,D,F	Binary	Notes
0	PF	B,0,0	B,00000,000	PF
1	PF	B,0,1	B,00000,001	PF
0	0	B,16,0	B,10000,000	Offset to first VF from PF is 128.
1	0	B,16,1	B,10000,001	
0	1	B,16,2	B,10000,010	
1	1	B,16,3	B,10000,011	
0	2	B,16,4	B,10000,100	
1	2	B,16,5	B,10000,101	
...				
0	63	B,31,6	B,11111,110	
1	63	B,31,7	B,11111,111	Last

7.7.2.6.1.2 Non-ARI Mode

When ARI is disabled, non-zero devices in the first bus cannot be used, thus a second bus is needed to provide enough RIDs. In this mode, the RID layout is as follows:

Table 7.54. RID Per VF — Non-ARI Mode

Port	VF#	B,D,F	Binary	Notes
0	PF	B,0,0	B,00000,000	PF
1	PF	B,0,1	B,00000,001	PF
0	0	B+1,16,0	B+1,10000,000	Offset to first VF from PF is 384.
1	0	B+1,16,1	B+1,10000,001	
0	1	B+1,16,2	B+1,10000,010	
1	1	B+1,16,3	B+1,10000,011	
0	2	B+1,16,4	B+1,10000,100	
1	2	B+1,16,5	B+1,10000,101	



Table 7.54. RID Per VF — Non-ARI Mode

Port	VF#	B,D,F	Binary	Notes
...				
0	63	B+1,31,6	B+1,11111,110	
1	63	B+1,31,7	B+1,11111,111	Last

Note: When the device ID of a physical function changes (because of LAN disable or LAN function select settings), the VF device IDs changes accordingly.

7.7.2.7 Hardware Resources Assignment

The main resources to allocate per VM are queues and interrupts. The assignment is static. If a VM requires more resources, it might be allocated to more than one VF. In this case, each VF gets a specific Ethernet MAC address/VLAN tag in order to enable forwarding of incoming traffic. The two VFs are then teamed in software.

7.7.2.7.1 PF Resources

A possible use of the PF is for a configuration setting without transmit and receive capabilities. In this case, it is not allocated to any queues but is allocated to one MSI-X vector.

The PF has access to all the resources of all VMs, but it is not expected to make use of resources allocated to active VFs.

7.7.2.7.2 Assignment of Queues to VF

See Section 7.2.1.2.1 for allocating Tx queues.

See Section 7.1.3.2 for allocating Rx queues.

The following table lists the Tx and Rx queues to VF allocation.

Table 7.55. Queue to VF Allocation

VF	Queues in 16 VMs Mode	Queues in 32 VMs Mode	Queues in 64 VMs Mode
0	0-7	0-3	0-1
1	8-15	4-7	2-3
...
15	120-127
...	
31		124-127	...
...			...
63			126-127

7.7.2.7.3 Assignment of MSI-X Vector to VF

See Section 7.3.4.3 for allocating MSI-X vectors in IOV mode.



7.7.2.8 CSR Organization

CSRs can be divided into three types:

- Global Configuration registers that should be accessible only to the PF. For example, link control and LED control. These types of registers also include all of the debug features such as the mapping of the packet buffers and is responsible for most of the CSR area requested by the Intel® Xeon® Processor D-1500 Product Family LAN controller. This includes per VF configuration parameters that can be set by the PF without performance impact.
- Per-VF parameters — For example, per VF reset, interrupt enable, etc. Multiple instances of these parameters are used only in an IOV system and only one instance is needed for non IOV systems.
- Per-queue parameters that should be replicated per queue — For example, head, tail, Rx buffer size, TPH tag, etc. These parameters are used by both a VF in an IOV system and by the PF in a non-IOV mode.

In order to support IOV without distributing the current drivers operation in legacy mode, the following method is used:

- The PF instance of BAR0 continues to contain the legacy and control registers. It is accessible only to the PF. The BAR enables access to all the resources including the VF queues and other VF parameters. However, it is expected that the PF driver does not access these queues in IOV mode.
- The VF instances of BAR0 provide control on the VF specific registers. These BARs have the same mapping as the original BAR0 with the following exceptions:
 - a. Fields related to the shared resources are reserved.
 - b. The queues assigned to a VF are mapped at the same location as the first same number of queues of the PF.
- Assuming some backward compatibility is needed for IOV drivers, The PF/VF parameters block should contain a partial register set as described in VF Device Registers section.

7.7.2.9 SR IOV Control

In order to control the IOV operation, the physical driver is provided with a set of registers. These include:

- The mailbox mechanism described in the next section.
- The switch and filtering control registers described in [Section 7.7.3.9](#).
- *PFVFLREC* register indicating that a VFLR reset occurred in one of the VFs (bitmap).

7.7.2.9.1 VF-to-PF Mailbox

The VF drivers and the PF driver require some means of communication between them. This channel can be used for the PF driver to send status updates to the VFs (such as link change, memory parity error, etc.) or for the VF to send requests to the PF (add to VLAN).

Such a channel can be implemented in software, but requires enablement by the VMM vendors. In order to avoid the need for such an enablement, the Intel® Xeon® Processor D-1500 Product Family LAN controller provides such a channel that enables direct communication between the two drivers.

The channel consists of a mailbox. Each driver can then receive an indication (either poll or interrupt) when the other side wrote a message.

Assuming a maximum message size of 64 bytes (one cache line), a memory of 64 bytes x 64 VMs = 4 KB is provided per port. The RAM is organized as follows:



Table 7.56. Mailbox Memory

RAM Address	Function	PF BAR 0 Mapping ¹	VF BAR 0 Mapping ²
0 – 63	VF0 <-> PF	0 – 63	VF0 + MBO
64 – 127	VF1 <-> PF	64 – 127	VF1 + MBO
....			
(4 KB-64) – (4 KB-1)	VF63<-> PF	(4 KB-64) – (4 KB-1)	VF63 + MBO

1. Relative to mailbox offset.
2. MBO = mailbox offset in VF CSR space.

In addition for each VF, the VFMailbox and PFMailbox registers are defined in order to coordinate the transmission of the messages. These registers contain a semaphore mechanism to enable coordination of the mailbox usage.

The PF driver can decide which VFs are allowed to interrupt the PF to indicate a mailbox message using the PFMBIMR mask register.

The following flows describe the usage of the mailbox:

Table 7.57. PF-to-VF Messaging Flow

Step	PF Driver	Hardware	VF #n driver
1	Set PFMailbox[n].PFU		
2		Set PFU bit if PFMailbox[n].VFU is cleared	
3	Read PFMailbox[n] and check that PFU bit was set. Otherwise wait and go to step 1.		
4	Write message to relevant location in VFMBMEM.		
5	Set the PFMailbox[n].STS bit and wait for ACK ¹ .		
6		Indicate an interrupt to VF #n.	
7			Read the message from VFMBMEM.
8			Set the VFMailbox.ACK bit.
9		Indicate an interrupt to PF.	
10	Clear PFMailbox[n].PFU		

1. The PF might implement a timeout mechanism to detect non-responsive VFs.

Table 7.58. VF-to-PF Messaging Flow

Step	PF Driver	Hardware	VF #n Driver
1			Set VFMailbox.VFU.
2		Set VFU bit if VFMailbox[n].PFU is cleared.	
3			Read VFMailbox [n] and check that VFU bit was set. Otherwise wait and go to step 1.
4			Write message to relevant location in VFMBMEM.
5			Set the VFMailbox.REQ bit.



Table 7.58. VF-to-PF Messaging Flow

Step	PF Driver	Hardware	VF #n Driver
6		Indicate an interrupt to PF.	
7	Read PFMBICR to detect which VF caused the interrupt.		
8	Read the adequate message from VFMBMEM.		
9	Set the PFMailbox.ACK bit.		
10		Indicate an interrupt to VF #n.	
11			Clear VFMailbox.VFU.

The content of the message is hardware independent and is determined by software.

The messages currently assumed by this specification are:

- Registration to VLAN/multicast packet/broadcast packets — A VF can request to be part of a given VLAN or to get some multicast/broadcast traffic.
- Reception of large packet — Each VF should notify the PF driver what is the largest packet size allowed in receive.
- Get global statistics — A VF can request information from the PF driver on the global statistics.
- Filter allocation request — A VF can request allocation of a filter for queuing/immediate interrupt support.
- Global interrupt indication.
- Indication of errors.

7.7.2.10 DMA

7.7.2.10.1 RID

Each VF is allocated a RID. Each DMA request should use the RID of the VM that requested it. See Section 7.7.2.6 for details.


7.7.2.10.2 Sharing of the DMA Resources

The outstanding requests and completion credits are shared between all the VFs. The tags attached to read requests are assigned the same way as in a non-virtualized setting, although in VF systems tags can be re-used for different RIDs.

7.7.2.10.3 TPH

The TPH enable is common to all the devices (all PFs and VFs). Given a TPH enabled device, each VM might decide for each queue, on which type of traffic (data, headers, Tx descriptors, Rx descriptors) the TPH should be asserted and what is the CPU ID assigned to this queue.

Note: There are no plans to virtualize TPH in the IOH. Thus, the physical CPU ID should be used in the programming of the CPUID field.



7.7.2.11 Timers

7.7.2.11.1 TCP Timer

The TCP timer is available only to the PF. It might indicate an interrupt to the VFs via the mailbox mechanism.

7.7.2.11.2 IEEE 1588

IEEE 1588 is a per-link function and thus is controlled by the PF driver. The VMs have access to the real time clock register.

7.7.2.11.3 Free Running Timer

The free running timer is a PF driver resource the VMs can access. This register is read only to all VFs and is reset only by the PCI reset.

7.7.2.12 Power Management and Wake Up

Power management is a PF resource and is not supported per VF.

7.7.2.13 Link Control

The link is a shared resource and as such is controllable only by the PF. This includes interface settings, speed and duplex settings, flow control settings, etc. The flow control packets are sent with the station Ethernet MAC address stored in the NVM. The watermarks of the flow control process and the time-out value are also controllable by the PF only.

MACsec is a per-link function and is controlled by the PF driver.

Double VLAN is a network setting and as such should be common to all VFs.

7.7.2.13.1 Special Filtering Options

Pass bad packets is a debug feature. As such, pass bad packets is available only to the PF. Bad packets are passed according to the same filtering rules of the regular packets.

Note: Pass bad packets might cause guest operating systems to get unexpected packets. As a result, it should be used only for debug purposes of the entire system.


Receiving long packets is enabled separately per Rx queue in the RXDCTL registers. As this impacts the flow control thresholds, the PF should be made aware of the decision of all the VMs. Because of this, the setup of TSO packets is centralized by the PF and each VF might request this setting.

7.7.3 Packet Switching

7.7.3.1 Assumptions

The following assumptions are made:

- The required bandwidth for the VM-to-VM loopback traffic is low. That is, the PCIe bandwidth is not congested by the combination of the VM-to-VM and the regular incoming traffic. This case is handled but not optimized for. Unless specified otherwise, Tx and Rx packets should not be dropped or lost due to congestion caused by loopback traffic.



- If the buffer allocated for the VM-to-VM loopback traffic is full, it is acceptable to back pressure the transmit traffic of the same TC. This means that the outgoing traffic might be blocked if the loopback traffic is congested.
- The decision on local traffic is done only according to the Ethernet DA address and the VLAN tag. There is no filtering according to other parameters (IP, L4, etc.). The switch has no learning capabilities. In case of double VLAN mode, the inner VLAN is used for the switching functionality.
- The forwarding decisions are based on the receive filtering programming.
- No packet switching between TCs.
- Coexistence with TimeSync: time stamp is not sampled for any VM-to-VM loopback traffic.
- Coexistence with Double VLAN: When double VLAN is enabled by DMATXCTL.GDV and it is expected to transmit untagged packets by software, transmit to receive packet switching should not be enabled.

7.7.3.2 Pool Selection

Pool selection is described in the following sections. A packet might be forwarded to a single pool or replicated to multiple pools. Multicast and broadcast packets are cases of replication, as is mirroring.

The following capabilities determine the destination pools of each packet:

- 128 Ethernet MAC address filters (RAH/RAL registers) for both unicast and multicast filtering. These are shared with L2 filtering. For example, the same Ethernet MAC addresses are used to determine if a packet is received by the switch and to determine the forwarding destination. In E-tag pool select mode RAL registers bits [13:0] are used for E-tag filtering {GRP, E-CID_base} and pool selection.
- 64 shared VLAN filters (PFVLVF and PFVLVFB registers) — each VM can be made a member of each VLAN.
- Hash filtering of unicast and multicast addresses (if the direct filters previously mentioned are not sufficient)
- Forwarding of broadcast packets to multiple pools
- Forwarding by Ethertype
- Mirroring by pool, VLAN, or link

Receive pool/queue allocation refer to [Section 7.1.3.2](#).

7.7.3.3 Rx Packets Switching

Rx packet switching is the second of three stages that determine the destination of a received packet. The three stages are defined in [Section 7.1.3](#).

As far as switching is concerned, it doesn't matter whether the Intel® Xeon® Processor D-1500 Product Family LAN controller's virtual environment operates in IOV mode or in VMDq2 mode.

When operating in replication mode, broadcast and multicast packets can be forwarded to more than one pool, and is replicated to more than one Rx queue. Replication is enabled by the Rpl_En bit in the PFVTCTL register.

Note: For the following algorithms packets with no E-tag in *PFVTCTL.POOLING_MODE* = 01b (E-tag mode) or all packets in *PFVTCTL.POOLING_MODE* = 00b (MAC mode) are defined as packets with no external tags.



7.7.3.3.1 Replication Mode Enabled

When replication mode is enabled, each broadcast/multicast packet can go to more than one pool. Finding the pool list of any packet is provided in the following steps:

1. **Exact unicast or multicast match** — If there is a match in one of the exact filters (RAL/RAH), for unicast or multicast packets, use the *MAC Pool Select Array (MPSAR[n])* bits as a candidate for the pool list. Note that MPSAR[n] must not enable more than one pool for unicast RAL/RAH filters. The compared field differs according to the filtering mode (*PFVTCTL.POOLING_MODE*):
 - a. If *POOLING_MODE* = 00b (MAC mode) or for packets with no external tags, the Destination MAC address is compared to {RAH[15:0], RAL[31:0]} for all the registers for which RAH.ADTYPE = 0b (MAC)
 - b. If *POOLING_MODE* = 01b (E-tag mode) and the packet has an E-tag, the E-tag E-PID (as extracted from offset 47:34, starting from the Ethertype, in the tag identified by the *ETAG_ETYPE* register) is compared to RAL[13:0] for all the registers for which RAH.ADTYPE = 1b (E-tag).
 2. **PFVFRE** — If any bit in the *PFVFRE* register is cleared, clear the respective bit in the pool list.
 3. **Broadcast** — If the packet is a broadcast packet with no external tag, add pools for which their *PFVML2FLT.BAM* bit (*Broadcast Accept Mode*) is set.
 4. **Unicast hash** — If the packet is a unicast packet with no external tag, and the prior steps yielded no pools, check it against the Unicast Hash Table (*PFUTA*). If there is a match, add pools for which their *PFVML2FLT.ROPE* bit (*Accept Unicast Hash*) is set.
 5. **Multicast hash** — If the packet is a multicast packet with no external tag and the prior steps yielded no pools, check it against the Multicast Hash Table (*MTA*). If there is a match, add pools for which their *PFVML2FLT.ROMPE* bit (*Receive Multicast Packet Enable*) is set.
 6. **Multicast promiscuous** — If the packet is a multicast packet with no external tag, take the candidate list from prior steps and add pools for which their *PFVML2FLT.MPE* bit (*Multicast Promiscuous Enable*) is set.
 7. **Unicast Promiscuous** - If the packet is a unicast packet with no external tag, take the candidate list from prior steps and add pools for which their *PFVML2FLT.UPE* bit (*Unicast Promiscuous Enable*) is set.
 8. **VLAN groups** — This step is relevant only when VLAN filtering is enabled by the *VLNCTRL.VFE* bit.
 - a. Tagged packets: enable only pools in the packet's VLAN group as defined by the VLAN filters — *PFVLVF[n]* and their pool list — *PFVLVFB[n]* or pools for which the *PFVML2FLT.VPE* (VLAN Promiscuous Enable) is set.
 - b. Untagged packets: enable only pools with their *PFVML2FLT.AUPE* bit set.
 - c. If there is no match, the pool list should be empty.
- Note:** In a VLAN network, un-tagged packets are not expected. Such packets received by the switch should be dropped, unless their destination is a virtual port set to receive these packets. The setting is done through the *PFVML2FLT.AUPE* bit. It is assumed that VMs for which this bit is set are members of a default VLAN and thus only MAC queuing is done on these packets.
9. **Default Pool** — If the pool list is empty at this stage and the *PFVTCTL.Dis_Def_Pool* bit is cleared, then set the default pool bit in the target pool list (from *PFVTCTL.DEF_PL*).
 10. **Ethertype filters** — If one of the Ethertype filters (ETQF) is matched by the packet and pooling action is requested and the *Pool Enable* bit in the ETQF is set, the pool list is set to the pool pointed to by the filter.
 11. **Filter Local Packets (source address pruning)** - The pruning operation depends on the filtering mode (*PFVTCTL.POOLING_MODE*):
 - a. If *POOLING_MODE* = 00b (MAC mode) or for packets with no external tag, the Source MAC address is compared to {RAH[15:0], RAL[31:0]} for all the registers for which RAH.ADTYPE = 0b (MAC).



- b. If *POOLING_MODE* = 01b (E-tag mode) and the packet has an E-tag, {00,E-tag ingress-E-CID_base} (as extracted from offset 31:20, starting from the Ethertype, in the tag identified by the *ETAG_ETYPE* register) is compared to RAL[13:0] for all the registers for which *RAH.ADTYPE* = 1b (E-tag).

If there is a match to one of the RAL/RAH(n) clear all the pools that are set in both the matching MPSAR[n] and the PFFLP from the pool list. This will prune multicast packets from getting back to the sender in VEPA mode.


12. **PFVFRE** — If any bit in the PFVFRE register is cleared, clear the respective bit in the pool list. The PFVFRE register blocks reception by a VF while the PF configures its registers.
13. **Mirroring** — Each of the four mirroring rules adds its destination pool (PFMRCTL.MP) to the pool list if the following applies:
 - a. **Pool mirroring** — *PFMRCTL.VPME* is set and one of the bits in the pool list matches one of the bits in the PFMRVM register.
 - b. **VLAN port mirroring** — *PFMRCTL.VLME* is set and the index of the VLAN of the packet in the PFVLF table matches one of the bits in the *PFMRVLAN* register.
 - c. **Uplink port mirroring** — *PFMRCTL.UPME* is set, the pool list is not empty.
14. **PFVFRE** — If any bit in the *PFVFRE* register is cleared, clear the respective bit in the pool list. The PFVFRE register blocks reception by a VF while the PF configures its registers. Note that this stage appears twice in order to handle mirroring cases.

7.7.3.3.2 Replication Mode Disabled

When replication mode is disabled, software should take care of multicast and broadcast packets and check which of the VMs should get them. In this mode, the pool list of any packet always contains one pool only according to the following steps:

1. **Exact unicast or multicast match** — The compared field differs according to the filtering mode (*PFVTCTL.POOLING_MODE*):
 - a. If *POOLING_MODE* = 00b (MAC mode) or for packets with no external tags, the Destination MAC address is compared to {RAH[15:0], RAL[31:0]} for all the registers for which *RAH.ADTYPE* = 0b (MAC).
 - b. If *POOLING_MODE* = 01b (E-tag mode) and the packet has an E-tag, the E-tag E-PID (as extracted from offset 47:34, starting from the Ethertype, in the tag identified by the *ETAG_ETYPE* register) is compared to RAL[13:0] for all the registers for which *RAH.ADTYPE* = 1b (E-tag).

If the compared field matches one of the exact filters (RAL/RAH), use the *MAC Pool Select Array* (MPSAR[n]) bits as a candidate for the pool list. Note that MPSAR[n] must not enable more than one pool for unicast RAL/RAH filters.
2. **PFVFRE** — If any bit in the PFVFRE register is cleared, clear the respective bit in the pool list.
3. **Unicast hash** — If the packet is a unicast packet with no external tag, and the prior steps yielded no pools, check it against the Unicast Hash Table (PFUTA). If there is a match, add the pool for which their PFVML2FLT.ROPE (*Accept Unicast Hash*) bit is set. Refer to the software limitations described after step 7.
4. **VLAN groups** — This step is relevant only when VLAN filtering is enabled by the *VLNCTRL.VFE* bit.
 - a. Tagged packets: enable only pools in the packet's VLAN group as defined by the VLAN filters — PFVLF[n] and their pool list — PFVLF[n] or pools for which the PFVML2FLT.VPE (VLAN Promiscuous Enable) is set.
 - b. Untagged packets: enable only pools with their *PFVML2FLT.AUPE* bit set.
 - c. If there is no match, the pool list should be empty.
5. **Default Pool** — If the pool list is empty at this stage and the *PFVTCTL.Dis_Def_Pool* bit is cleared, then set the default pool bit in the target pool list (from *PFVTCTL.DEF_PL*).



6. **Multicast or Broadcast** — If the packet is a multicast or broadcast packet with no external tag and was not forwarded in step 1 and 2, set the default pool bit in the pool list (from PFVTCTL.DEF_PL).
7. **Ethertype filters** — If one of the Ethertype filters (ETQF) is matched by the packet and queuing action is requested and the *Pool Enable* bit in the ETQF is set, the pool list is set to the pool pointed by the filter.
8. **PFVFRE** — If any bit in the PFVFRE register is cleared, clear the respective bit in the pool list. The PFVFRE register blocks reception by a VF while the PF configures its registers.

The following software limitations apply when replication is disabled:

- Software must not set more than one bit in the bitmaps of the exact filters. Note that multiple bits can be set in an RAH register as long as it's guaranteed that the packet is sent to only one queue by other means (such as VLAN).
- Software must not set per-VM promiscuous bits (multicast or broadcast).
- Software must not set the *ROPE* bit in more than one PFVML2FLT register.
- Software should not activate mirroring.
- Software should not filter out local packets by setting PFFLP bits

7.7.3.4 Tx Packets Switching

Tx switching is used only in a virtualized environment to serve VM-to-VM traffic. Packets that are destined to one or more local VMs are directed back (loopback) to the Rx packet buffers. Enabling Tx switching is done by setting the PFDTXGSWC.LBE bit. Tx to Rx switching always avoids packet drop as if flow control is enabled. Therefore, the software must set the FCRTTH[n].RTH fields regardless if flow control is activated on the Intel® Xeon® Processor D-1500 Product Family LAN controller.

Tx switching rules are very similar to Rx switching in a virtualized environment, with the following exceptions:

- If a target pool is not found, the default pool is used only for broadcast and multicast packets.
- A unicast packet that matches an exact filter is not sent to the LAN.
- Broadcast and multicast packets are always sent to the external LAN.
- A packet might not be sent back to the originating pool (even if the destination address is equal to the source address) unless loopback is enabled for that pool by the PFVMTXSW[n] register.

The detailed flow for pool selection as well as the rules that apply to loopback traffic is as follows:

- Loopback is disabled when the network link is disconnected. It is expected (but not required) that system software (including VMs) does not post packets for transmission when the link is disconnected. Loopback is disabled when the RXEN (*Receive Enable*) bit is cleared.
- Loopback packets are identified by the *LB* bit in the receive descriptor.

Notes: When Tx switching is enabled, the host shall avoid sending packets longer than 9.5 KB
Tx switching should be enabled only if *PFVTCTL.POOLING_MODE* = 00b (MAC mode).

7.7.3.4.1 Replication Mode Enabled

When replication mode is enabled, the pool list for any packet is determined according to the following steps:

1. **Exact unicast or multicast match** — If there is a match in one of the exact filters (RAL/RAH), for unicast or multicast packets, take the MPSAR[n] bits as a candidate for the pool list. Note that MPSAR[n] must not enable more than one pool for unicast RAL/RAH filters.



2. **Ether-type filters** - if one of the enabled (*ETQF.TX_ANTISPOOF*) Transmit Ether-type filters (*ETQF.ETYPE*) is matched by the packet and the appropriate Pool loopback/anti spoof bit in the *PFVFSPOOF* is set (*PFVFSPOOF.ETHERTYPELB/PFVFSPOOF.ETHERTYPEAS*), the pool list is set to the pool pointed to by the ETQS matching register (loopback) or the packet is dropped (anti spoof). All the subsequent steps after step 3. are skipped.

Note: Loopback on ETQF match is possible even if the general loopback enable bit (*PFDTXGSWC.LBE*) is clear.

3. **PFVFRE** — If any bit in the PFVFRE register is cleared, clear the respective bit in the pool list.
4. **Broadcast** — If the packet is a broadcast packet, add pools for which their *PFVML2FLT.BAM* bit (*Broadcast Accept Mode*) is set.
5. **Unicast hash** — If the packet is a unicast packet, and the prior steps yielded no pools, check it against the Unicast Hash Table (PFUTA). If there is a match, add pools for which their *PFVML2FLT.ROPE* bit (*Accept Unicast Hash*) is set.
6. **Multicast hash** — If the packet is a multicast packet and the prior steps yielded no pools, check it against the Multicast Hash Table (MTA). If there is a match, add pools for which their *PFVML2FLT.ROMPE* bit (*Receive Multicast Packet Enable*) is set.
7. **Unicast Promiscuous** - If the packet is a unicast packet, take the candidate list from prior steps and add pools for which their *PFVML2FLT.UPE* bit (*Unicast Promiscuous Enable*) is set.
8. **Multicast Promiscuous** — If the packet is a multicast packet, take the candidate list from prior steps and add pools for which their *PFVML2FLT.MPE* bit (*Multicast Promiscuous Enable*) is set.
9. **Filter source pool** — The pool from which the packet was sent is removed from the pool list unless the *PFVMTXSW.LLE* bit is set.
10. **VLAN groups** — This step is relevant only when VLAN filtering is enabled by the *VLNCTRL.VFE* bit.
 - a. Tagged packets: enable only pools in the packet's VLAN group as defined by the VLAN filters — *PFVLVF[n]* and their pool list — *PFVLVFB[n]* or pools for which the *PFVML2FLT.VPE* (VLAN Promiscuous Enable) is set.
 - b. Untagged packets: enable only pools with their *PFVML2FLT.AUPE* bit set.
 - c. If there is no match, the pool list should be empty.
11. **Forwarding to the network** — Packets are forwarded to the network in the following cases:
 - a. All broadcast and multicast packets.
 - b. Unicast packets that do not match any exact filter. A match of an exact filter that also points to a pool disabled via PFVFRE is not considered a match.
12. PFVFRE — If any bit in the PFVFRE register is cleared, clear the respective bit in the pool list (pre mirroring step).
13. **Mirroring** — Each of the following three mirroring rules adds its destination pool (*PFMRCTL.MP*) to the pool list if the following applies:
 - a. **Pool mirroring** — *PFMRCTL.VPME* is set and one of the bits in the pool list matches one of the bits in the PFMRVM register.
 - b. **VLAN port mirroring** — *PFMRCTL.VLME* is set and the index of the VLAN of the packet in the PFVLVF table matches one of the bits in the PFMRVLAN register.
 - c. **Downlink port mirroring** — *PFMRCTL.DPME* is set and the packet is sent to the network.
14. PFVFRE — If any bit in the PFVFRE register is cleared, clear the respective bit in the pool list (post mirroring step).



7.7.3.4.2 Replication Mode Disabled

When replication mode is disabled, software should take care of multicast and broadcast packets and check which of the VMs should get them. In this mode the pool list for any packet always contains one pool only according to the following steps:

1. **Exact unicast or multicast match** — If the packet DA matches one of the exact filters (RAL/RAH), take the *MPSAR[n]* bits as a candidate for the pool list. Note that *MPSAR[n]* must not enable more than one pool for unicast RAL/RAH filters.
2. **Ether-type filters** - if one of the enabled (*ETQF.TX_ANTISPOOF*) Transmit Ether-type filters (*ETQF.ETYPE*) is matched by the packet and the appropriate Poll loopback/anti spoof bit in the *PFVFSPOOF* is set (*PFVFSPOOF.ETHERTYPELB/PFVFSPOOF.ETHERTYPEAS*), the pool list is set to the pool pointed to by the *ETQS* matching register (loopback) or the packet is dropped (anti spoof). All the subsequent steps after step 3. are skipped.

Note: Loopback on Ethertype match is possible even if the general loopback enable bit (*PFDTXGSWC.LBE*) is clear.


3. **PFVFRE** — If any bit in the *PFVFRE* register is cleared, clear the respective bit in the pool list.
4. **Unicast hash** — If the packet is a unicast packet, and the prior steps yielded no pools, check it against the Unicast Hash Table (*PFUTA*). If there is a match, add the pool for which their *PFVML2FLT.ROPE* bit (*Accept Unicast Hash*) is set. Refer to the software limitations that follow.
5. **VLAN groups** — This step is relevant only when VLAN filtering is enabled by the *VLNCTRL.VFE* bit.
 - a. Tagged packets: enable only pools in the packet's VLAN group as defined by the VLAN filters — *PFVLVF[n]* and their pool list — *PFVLVFB[n]* or pools for which the *PFVML2FLT.VPE* (VLAN Promiscuous Enable) is set.
 - b. Untagged packets: enable only pools with their *PFVML2FLT.AUPE* bit set.
 - c. If there is no match, the pool list should be empty.
6. **Multicast or Broadcast** — If the packet is a multicast or broadcast packet and was not forwarded in step 1 and 2, set the default pool bit in the pool list (from *PFVTCTL.DEF_PL*).
7. **Filter source pool** — The pool from which the packet was sent is removed from the pool list unless the *PFVMTXSW.LLE* bit is set.
8. **Forwarding to the Network** — Packets are forwarded to the network in the following cases:
 - a. All broadcast and multicast packets.
 - b. Unicast packets that do not match any exact filter. A match of an exact filter that also points to a pool disabled via *PFVFRE* is not considered a match.
9. **PFVFRE** — If any bit in the *PFVFRE* register is cleared, clear the respective bit in the pool list.

The following software limitations apply when replication is disabled:

1. It is software's responsibility not to set more than one bit in the bitmaps of the exact filters. Note that multiple bits can be set in an *RAH* register as long as it is guaranteed that the packet is sent to only one queue by other means (such as VLAN).
2. Software must not set per-VM promiscuous bits (multicast or broadcast).
3. Software must not set the *ROPE* bit in more than one *PFVML2FLT* register.
4. Software should not activate mirroring.

7.7.3.5 Mirroring Support

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports four separate mirroring rules, each associated with a destination pool (mirroring can be done in up to four pools). Each rule is programmed with one of the four mirroring types:



1. Pool mirroring — reflect all the packets received to a pool from the network.
2. Uplink port mirroring — reflect all the traffic received from the network.
3. Downlink port mirroring — reflect all the traffic transmitted to the network.
4. VLAN mirroring — reflect all the traffic received from the network in a set of given VLANs (either from the network or from local VMs).

Note: Reflecting all the traffic received by any of the pools (either from the network or from local VMs) is supported by enabling mirroring of all pools.

Note: In order to get all traffic irrespective of the filtering rules, the mirroring pool should be set to promiscuous mode and promiscuous VLAN mode.

Mirroring modes are controlled by a set of rule control registers:

- *PFMRCTL* — controls the rules to be applied and the destination port.
- *PFMRVLAN* — controls the VLAN ports as listed in the *PFVLVF* table taking part in the VLAN mirror rule.
- *PFMRVM* — controls the pools taking part in the pool mirror rule.

7.7.3.6 Offloads

The general rule is that offloads are executed as configured for the pool and queue associated with the receive packet. Some special cases:

- If a packet is directed to a single pool, then offloads are determined by the pool and queue for that packet.
- If a packet is replicated to more than one pool, then each copy of the packet is offloaded according to the configuration of its pool and queue.
- If replication is disabled, offloads are determined by the unique destination of the packet.

The following subsections describe exceptions to the previously described special cases.

7.7.3.6.1 Local Traffic Offload

The following capabilities are not supported on the loopback path:

- The RSS and VLAN strip offload capabilities are only supported if the *CC* bit in the transmit descriptor of the packet is set and if it is a tunnel packet, the *TUNNEL.OUTERIPCS* is also set. The reason is that when these bits are not set, software does not provide the necessary offload offsets with the Tx packet.
- Receive Side Coalescing (RSC) is not supported.


7.7.3.6.2 Rx Traffic Offload

- CRC offload is a global policy. CRC strip is enabled or disabled for all received packets.

7.7.3.7 Rate Control Features

7.7.3.7.1 Congestion Control

Tx packets going through the local switch are stored in the Rx packet buffer, similar to packets received from the network. Tx to Rx switching always avoids packet drop as if flow control is enabled. Therefore, the software must set the *FCRTH[n].RTH* fields regardless if flow control is activated on the Intel® Xeon® Processor D-1500 Product Family LAN controller.



The Intel® Xeon® Processor D-1500 Product Family LAN controller guarantees that one TC flow is not affected by congestion in another TC.

Receive and local traffic are provided with the same priority and performance expectations. Packets from the two sources are merged in the Rx packet buffers, which can in general support both streams at full bandwidth. Any congestion further in the pipeline (such as lack of PCIe bandwidth) evenly affects Rx and local traffic.

7.7.3.7.2 Tx Queue Arbitration and Rate Control

In order to guarantee each pool with adequate bandwidth, a per-pool bandwidth control mechanism is added to the Intel® Xeon® Processor D-1500 Product Family LAN controller. Each Tx pool gets a percentage of the transmit bandwidth and is guaranteed it can transmit within its allocation. This arbitration is combined with the TC arbitration.

7.7.3.7.3 Receive Priority

As the switch might decide to loopback packets from the transmit path to the receive path, in case the receive path is full, the transmit path might be blocked (including the traffic to the LAN). The Intel® Xeon® Processor D-1500 Product Family LAN controller guarantees that packets are not dropped because of that and that one traffic class flow is not affected by congestion in another traffic class. The PF driver might decide to program the Intel® Xeon® Processor D-1500 Product Family LAN controller to drop packets from receive queues without available descriptors.

In order to keep the congestion effect locality, receive traffic from the LAN have higher priority than loop back traffic. This way large loopback traffic does not impact the network.

7.7.3.8 Small Packets Padding

In Virtualized systems, the software device driver receiving the packet in the VM might not be aware of all the hardware offloads applied to the packet. Thus, in case of stripping actions by the hardware (VLAN strip), it might receive packets which are smaller than a legal packet. The Intel® Xeon® Processor D-1500 Product Family LAN controller provides an option to pad small packets in such cases so that all packets have a legal size. This option can be enabled only if the CRC is stripped. In these cases, all small packets are padded to 60 bytes (legal packet - 4 bytes CRC). The padding is done with zero data. This function is enabled via the *RDRXCTL.PSP* bit.

Packet padding is done after tag extractions and do not take into account if tags are exposed in the Rx descriptor status.

7.7.3.9 Switch Control

The PF driver has some control of the switch logic. The following registers are available to the PF for this purpose:

PFVCTL: - PF Virtual Control register — contains the following fields:

- Replication Enable (*Rpl_En*) — enables replication of multicast and broadcast packets — both in incoming and local traffic. If this bit is cleared, Tx multicast and broadcast packets are sent only to the network and Rx multicast and broadcast packets are sent to the default pool.
- Default Pool (*DEF_PL*) — defines the target pool for packets that passed L2 filtering but didn't pass any of the pool filters. This field is invalid when the *Dis_Def_Pool* bit is set.
- Disable Default Pool (*Dis_Def_Pool*) — disables acceptance of packets that failed all pool filters.
- *PFVFRE* — Enables/disables reception of packets from the link to a specific VF. Used during initialization of the VF. See [Section 4.1.2.2](#) for more details.



- PFDTXGSWC (LBE) — VMDQ loopback enables switching of Tx traffic to the Rx path for VM-to-VM communication.
- PFVFSPOOF — MAC Anti-spoof Enable (MACAS) — enables filtering of Tx packet for anti-spoof.
- Local Loopback Enable (LLE) — defines whether or not to allow loopback of a packet from a certain pool into itself.
- Queue Drop Enable (PFQDE) register — A register defining global policy for drop enable functionality when no descriptors are available. It lets the PF override the per-queue SRRCTL[n] Drop_En setting. PFQDE should be used in SR-IOV mode.
- PFVML2FLT — Receive Overflow Multicast Packets (ROMPE) — accept multicast hash — Defines whether or not a pool accepts packets that match the multicast MTA table.
- Receive MAC Filters Overflow (ROPE) — accept unicast hash — Defines whether or not a pool accepts packets that match the unicast PFUTA table.
- Broadcast Accept (BAM) — Defines whether or not a pool accepts broadcast packets.
- Multicast Promiscuous (MPE) — Defines whether or not a pool accepts all multicast packets.
- Accept Untagged Packets Enable (AUPE) — Defines whether or not a pool accepts untagged VLAN packets.
- Mirror Control — See [Section 7.7.3.5](#).
- PFVFTE — Enables/disables transmission of packets to the link to a specific VF. Used during initialization of the VF. See [Section 4.1.2.2](#) for more details.
- PFVLFV/PFVLVFB — VLAN queuing table — A set of 64 VLAN entries with an associated bitmap, one bit per pool. Bits are set for each pool that participates in this VLAN.
- Unicast Table Array (PFUTA) — a 4 Kb array that covers all combinations of 12 bits from the MAC destination address. A received unicast packet that misses the MAC filters is compared against the PFUTA. If the relevant bit in the PFUTA is set, the packet is routed to all pools for which the *ROPE* bit is set.
- Multicast Table Array (MTA) — a 4 Kb array that covers all combinations of 12 bits from the MAC destination address. A received multicast packet that misses the MAC filters is compared against the MTA. If the relevant bit in the MTA is set, the packet is routed to all pools for which the *ROMPE* bit is set.

7.7.4 Security Features

The Intel® Xeon® Processor D-1500 Product Family LAN controller allows some security checks on the inbound and outbound traffic of the switch.


7.7.4.1 Inbound Security

Each incoming packet (either from the LAN or from a local VM) is filtered according to the VLAN tag so that packets from one VLAN cannot be received by pools that are not members of that VLAN.

7.7.4.2 Outbound Security

7.7.4.2.1 MAC Anti-spoofing

Each pool is associated with one or more Ethernet MAC addresses on the receive path. The association is determined through the MPSAR registers. The MAC anti-spoofing capability insures that a VM always uses a source Ethernet MAC address on the transmit path that is part of the set of Ethernet MAC addresses defined on the Rx path. A packet with a non-matching SA is dropped, preventing spoofing of the Ethernet MAC address. This feature is enabled in the *PFVFSPOOF.MACAS* field, and can be enabled per Tx pool.



Note: Anti-spoofing is not available for VMs that hide behind other VMs whose Ethernet MAC addresses are not part of the RAH/RAL Ethernet MAC Address registers. In this case, anti-spoofing should be done by software switching, handling these VMs.

7.7.4.2.2 VLAN Anti-spoofing

Each pool is associated with one or more VLAN tags on the receive path. The association is determined through the PFVLVF and PFVLVFB registers. The VLAN anti-spoofing capability insures that a VM always uses a VLAN tag on the transmit path that is part of the set of VLAN tags defined on the Rx path. A packet with a non-matching VLAN tag is dropped, preventing spoofing of the VLAN tag. This feature is enabled in the PFVFSPOOF.VLANAS field, and can be enabled per Tx pool.

Notes: If VLAN anti-spoofing is enabled, then MAC anti-spoofing must be enabled as well. When double VLAN is enabled by DMATXCTL.GDV and it is expected to transmit untagged packets by software, VLAN anti-spoofing should not be enabled. Un-tagged packets are not checked by VLAN anti-spoofing.

7.7.4.2.3 VLAN Tag Validation

In PCI-SIG IOV scenarios the driver might be malicious, and thus might fake a VLAN tag. The Intel® Xeon® Processor D-1500 Product Family LAN controller provides the ability to force a specific VLAN value for a VM. The possible behaviors are controlled by the PFVMVIR[n] registers as follows:

- Use descriptor value — to be used in case of a trusted VM that can decide which VLAN to send. This option should also be used in case one VM is member of multiple VLANs.
- Always insert default VLAN — this mode should be used for non-trusted or non-VLAN aware VMs. In this case, any VLAN insertion command from the VM is ignored. If a packet is received with a VLAN, the packet should be dropped.
- Never insert VLAN — This mode should be used in a non-VLAN network. In this case, any VLAN insertion command from the VM is ignored. If a packet is received with a VLAN, the packet should be dropped.

Notes: The VLAN insertion settings should be done before any of the queues of the VM are enabled. When double VLAN is enabled by DMATXCTL.GDV and it is expected to transmit untagged packets by software, VLAN validation should not be enabled.

7.7.4.2.4 E-tag Insertion

In addition to the VLAN insertion previously described, the device can add an E-tag to the packets sent by a VF. The PFVMVIR[n].TAGA bit defines if a tag should be added. If this bit is set, the added tag is as follows:

- TAGA = 00b (MAC mode) or 11b (reserved): Insert no tag.
- TAGA = 01b (E-tag mode): An E-tag is added as follow: {TAG_ETYPE.ETAG_ETHERTYPE[15:0], PFVMTIR.PORT_TAG_ID[31:0], 0x0000}.

Notes: The E-tag insertion settings should be done before any of the queues of the VM are enabled.

7.7.4.2.5 Ether-type Anti-spoofing/Pruning

Each pool can be set to drop packets associated with one or more (up to 8) Ether-types on the transmit path. The Ether-types are defined using the ETQF registers. The Ether-type pruning (anti spoof) capability ensures that a VM doesn't send out LLDP/ECP control frames to external switches. The Ether-type loopback capability enables the loopback of link layer packets to a pre defined pool destination.



Ether-type pruning (anti spoof) feature is enabled in the *PFVFSPOOF.ETHERTYPEAS* field, and can be enabled per transmit pool. Ether-type loopback feature is enabled in the *PFVFSPOOF.ETHERTYPELB* field, and can be enabled per transmit pool.

7.7.4.3 Malicious Driver Detection

Hardware can be programmed to take some action as a result of some misbehavior of a VM. These actions might hint to the fact that some VM is malicious and the VMM should remedy the situation. In order to inform the VMM of this fact, The Mailbox bit (*EICR.MAILBOX* bit) might be set to indicate the occurrence of such behavior. The *LMVM_RX* and *LMVM_TX* register contains information on which queue (*LMVM_TX.Last_Q* or *LMVM_RX.Last_Q*) and port (*LMVM_TX.Ma_PF* and *LMVM_RX.Ma_PF*) the malicious behavior was detected. The *LVMMC_TX* and *LVMMC_RX* registers contain information on the type of error detected and are clear by read.

Malicious driver behavior detection is enabled by setting the *DMATXCTL.MDP_EN* and *RDRXCTL.MDP_EN* bits to 1b. This capability should be enabled before queues are exposed to non-trusted virtual machines.

On detection of a malicious driver event, the Intel® Xeon® Processor D-1500 Product Family LAN controller stops activity of the offending queue, and sets the bit matching the offending queue in the *WQBR_RX* and *WQBR_TX* registers. If *DMATXCTL.MBINTEN* or *RDRXCTL.MBINTEN* are set, the Intel® Xeon® Processor D-1500 Product Family LAN controller generates an interrupt for transmit or receive errors respectively by asserting the *EICR.MAILBOX* bit. Cause of malicious driver activation is reported in the *LVMMC_TX* and *LVMMC_RX* registers. To re-activate the offending queue, the software device driver should release it by setting the matching bit in *WQBR_RX* or *WQBR_TX* register.

The *WQBR_RX* and *WQBR_TX* registers keeps track of all the queues on which a malicious activity was detected and can be used by the software device driver to track multiple events.

After the queue is disabled the initialization flow defined in [Section 4.5.8.1](#) for receive queues and in [Section 4.5.9.1](#) for transmit queues should be applied.

7.7.4.3.1 Queue Context Validation

The Intel® Xeon® Processor D-1500 Product Family LAN controller checks that the queue context submitted when a queue is enabled is valid. It also prevents change to static configuration while the queue is enabled. These checks are done both for receive and transmit queues. The following table lists the checks done on the queue context:

Check type	Description	Bit in LVMMC_TX/RX	Action
Receive queue context validity	Check that: <ul style="list-style-type: none"> RDLEN > 0 SRRCTL.BSIZEPACKET > 0 SRRCTL.BSIZEHEADER > 0 if DESC TYP requires header split SRRCTL.BSIZEHEADER <= 1024. SRRCTL.DESCTYPE is valid 	INVALID_RXQ_CONTEXT	Prevent Enable of queue
Receive queue context change on the fly	Attempt to write RDBAL, RDBAH, RDLEN, or SRRCTL while queue is enabled		Disable queue.
Transmit queue context validity	Check that: <ul style="list-style-type: none"> TDLEN > 0 	INVALID_TXQ_CONTEXT	Prevent Enable of queue
Transmit queue context change on the fly	Attempt to write TDBAL, TDBAH, TDLEN, TDWBAL, or TDWBAH while queue is enabled		Disable queue.



Note: Queue zero is enabled by default. Thus, if malicious driver protection is enabled, in order to change Queue zero configuration, the queue should be disabled.

7.7.4.3.2 Transmit Descriptor Validity Checks

The following table lists the checks are done by the Intel® Xeon® Processor D-1500 Product Family LAN controller to define if a transmit packet descriptor is valid. All the checks are done on the descriptors. The checks on the packet header are described in the previous sections.

Table 7.59. Malicious Driver - Tx Descriptor Checks

Check type	Description	Bit in LVMMC_TX	Action
Mac Header size	Checks that the MAC header size in the context descriptor is at least 14 (or 18 in case of offloaded packet and double VLAN).	MAC_HEADER	Drop packet and stop offending queue
IPV4 Header	If a checksum or TSO offload is required, checks that the IPv4 header size in the context descriptor is at least 20.	IPV4_HEADER	Drop packet and stop offending queue
IPV6 Header	If a TSO offload is required, checks that the IPv6 header size in the context descriptor is at least 40.	IPV6_HEADER	Drop packet and stop offending queue
Wrong MAC_IP	Check that the MAC+ IP (+ OUTERIP + TUNNEL) header size is not bigger than the packet size.	WRONG_MAC_IP	Drop packet and stop offending queue
TCP header size	If a TCP TSO offload is required, checks that the TCP header size in the context descriptor is at least 20.	TCP_LSO	Drop packet and stop offending queue
UDP header size	If a UDP TSO offload is required, checks that the TCP header size in the context descriptor is at least 8.	UDP_LSO	Drop packet and stop offending queue
SCTP data size	If a SCTP checksum offload is required, checks that the SCTP L4 packet size (including header and data) is at least 12.	STCP_CS	Drop packet and stop offending queue
Packet too big	In case of a single send, check that the packet is not larger than 15.5KBytes.	SIZE	Drop packet and stop offending queue
Packet too small	Check that the total length of the packet transmitted, not including FCS is at least 13 bytes.	N/A	Silently drop packet.
Illegal offload request.	Check that TSO is no requested for SCTP or that a checksum offload is not requested for IPv6 packets.	OFF_ILL	Drop packet and stop offending queue
Illegal IPsec offload request by VF	Check that a VF had not requested IPsec offload.	IPSEC_OFFLOAD	Drop packet and stop offending queue
SCTP alignment	If an SCTP CRC offload is requested, check that the data size is 4 byte aligned.	SCTP_ALIGNED	Drop packet and stop offending queue
Zero MSS	Check that the MSS size is larger than zero.	ZERO_MSS	Drop packet and stop offending queue
Context in middle of packet	Check that a context descriptor is not sent in the middle of a packet.	CONTEXT_IN_PACKET	Drop packet and stop offending queue
Number of large send header buffers	Check that the Large send header is contained in at most 4 buffers.	LSO_MORE_THAN_4	Drop packet and stop offending queue
Buffers size and length match - Single Send	For single send, check that the total of all buffers size and the packet length match.	OOS_SSO	Drop packet and stop offending queue
Buffers size and length match - Large Send	For LSO, check that the total of all buffers size and the packet length match.	OOS_LSO	Drop packet and stop offending queue
LSO wrong header size	For LSO, check that the size of the header fetched actually match the expected header size.	N/A	Silently drop packet. Counted in SSVPC
UDP data size	If a UDP checksum offload is required, checks that the UDP L4 packet size in the context descriptor is at least 8.	SSO_UDP	Drop packet and stop offending queue
TCP data size	If a TCP checksum offload is required, checks that the TCP L4 packet size in the context descriptor is at least 20.	SSO_TCP	Drop packet and stop offending queue



Table 7.59. Malicious Driver - Tx Descriptor Checks

Check type	Description	Bit in LVMMC_TX	Action
Descriptor Type	Check that only descriptor types 2 (context) or 3 (advanced data descriptor) are used.	DESC_TYPE	Drop packet and stop offending queue
CC bit not set when needed	Check Context (CC) bit in descriptor is not set in virtualization mode.		Drop packet and stop offending queue
Null packet check	Check that a Null packet has the EOP bit set.	WRONG_NULL	Drop packet and stop offending queue
Packet without EOP	Check that a only entire packets are provided by the software device driver.	NO_EOP	Drop packet and stop offending queue
Burst of contexts	Check that less than 4 contiguous context descriptor are sent by the driver.	CONTEXT_BURST	Drop packet and stop offending queue
Legacy Descriptor in SR-IOV	Legacy descriptor used when SR-IOV is enabled.	LEGACY_DESC_IOV	Drop packet and stop offending queue

7.7.4.3.3 Reactive Malicious Behavior Detection

Table 7.60 lists the checks are done by the Intel® Xeon® Processor D-1500 Product Family LAN controller to detect a malicious behavior, even if the packet seems valid.

Table 7.60. Reactive Malicious Checks

Check type	Description	Bit in LVMMC_TX/RX	Action
Malicious VF memory access	A PCIe DMA access initiated by a VF ended with Unsupported Request (UR) or Completer Abort (CA). This check is done for both Tx and Rx queues.	INV_MACC	Drop packet and stop offending queue
DMA access outside of active memory ¹	A queue attempt to read memory outside of the active memory range of the system. See Section 3.1.3.6 for details.	INV_MACC	Drop PCIe read transaction. Stop queue that requested the transaction
	A queue attempt to write memory outside of the active memory range of the system. See Section 3.1.3.6 for details.	N/A	Silently drop PCIe write transaction
VLAN not expected in packet	A packet where VLAN should be added by device is in Tx descriptor. See Section 7.7.4.2.3	VLAN_IERR	Drop packet
Anti spoof checks	See Section 7.7.4.2.1 and Section 7.7.4.2.2 and Section 7.7.4.2.5 for details	MAC_VLAN_SPOOF	Drop packets

1. A PCIe error interrupt can be asserted when such a transaction is dropped. See Section 3.1.3 for details.

7.7.5 Virtualization of Hardware

This section describes additional features used in both IOV and VMDq2 modes.

7.7.5.1 Per-pool Statistics

Part of the statistics are by definition shared and cannot be allocated to a specific VM. For example, CRC error count cannot be allocated to a specific VM, as the destination of such a packet is not known if the CRC is wrong.

All the non-specific statistics are handled by the PF driver in the same way it is done in non-virtualized systems. A VM might request a statistic from the PF driver but might not access it directly.



The conceptual model used to gather statistics in a virtualization context is that each queue pool is considered as a virtual link and the Ethernet link is considered as the uplink of the switch. Thus, any packet sent by a pool is counted in the Tx statistics, even if it was forwarded to another pool internally or was dropped by the MAC for some reason. In the same way, a replicated packet is counted in each of the pools receiving it.

The following statistics are provided per pool:

- Good packet received count
- Good packet transmitted count
- Good octets received count
- Good octets transmitted count
- Multicast packets received count

Note: All the per VF statistics are read only and wrap around after reaching their maximum value.

7.8 Tunnelling Support

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports the VXLAN and NVGRE tunneling packet formats. As part of this support the following features are supported:

- LSO and transmit checksum offloads for tunneled packets as described in [Section 7.2.4](#) and [Section 7.2.5](#).
- RSS forwarding based on inner L3/L4 header as described in [Section 7.1.3.6](#)
- Flow director forwarding based on Tenant ID, inner MAC and inner VLAN as described in [Section 7.1.3.5](#)
- New indications in the receive descriptor with information on the tunnel headers and on the outer header checksum as described in [Section 7.1.5.2](#)
- New packet split modes based on the tunnel header or outer L2 header as described in [Section A.2.5](#)

7.9 Receive Side Coalescing (RSC)

The Intel® Xeon® Processor D-1500 Product Family LAN controller can merge multiple received frames from the same TCP/IP connection (referred to as flow in this section) into a single structure. The Intel® Xeon® Processor D-1500 Product Family LAN controller does this by coalescing the incoming frames into a single or multiple buffers (descriptors) that share a single accumulated header. This feature is called RSC. Note that the term Large Receive is used to describe a packet construct generated by RSC.

The Intel® Xeon® Processor D-1500 Product Family LAN controller digests received packets and categorizes them by their TCP/IP connections (flows). For each flow, hardware coalesces the packets as shown in [Figure 7.35](#) and [Figure 7.36](#) (the colored parameters are explained in the RSC context table and receive descriptor sections). The Intel® Xeon® Processor D-1500 Product Family LAN controller can handle up to 32 concurrent flows per LAN port at any given time. Each flow handled by RSC offload has an associated context. The Intel® Xeon® Processor D-1500 Product Family LAN controller opens and closes the RSC contexts autonomously with no need for any software intervention. Software needs only to enable RSC in the selected receive queues.

Note: When RSC is enabled, advanced receive descriptors should be used and CRC strip should be enabled.

Figure 7.35 shows a top level flow diagram that is used for RSC functionality. The following sections provide a detailed explanation of this flow as well as the memory structures and device settings that support the RSC functionality.

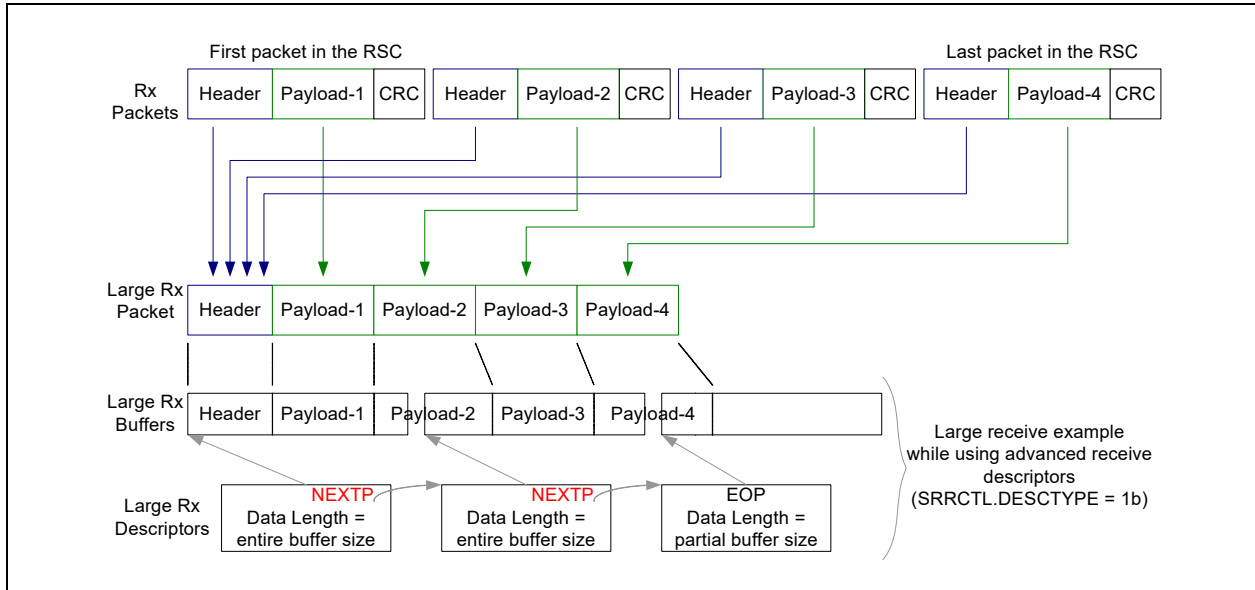


Figure 7.35. RSC Functionality (No Header Split)

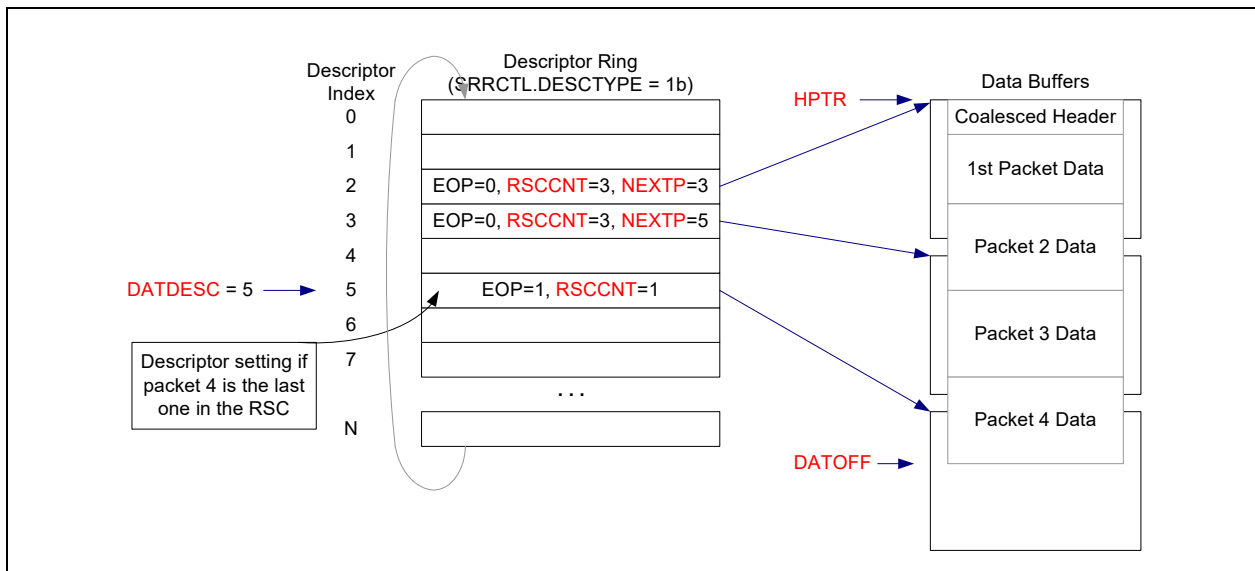


Figure 7.36. RSC Functionality (No Header Split)

Note: Software might abort reception to any queue at any time. For example: VFLR or queue disable. Following these settings, hardware aborts further DMA(s) and descriptor completions. Specifically, active RSC(s) in the specific queue(s) are not completed. In such cases there could be completed packets and RSC(s) hidden from software by prior incomplete RSC(s).

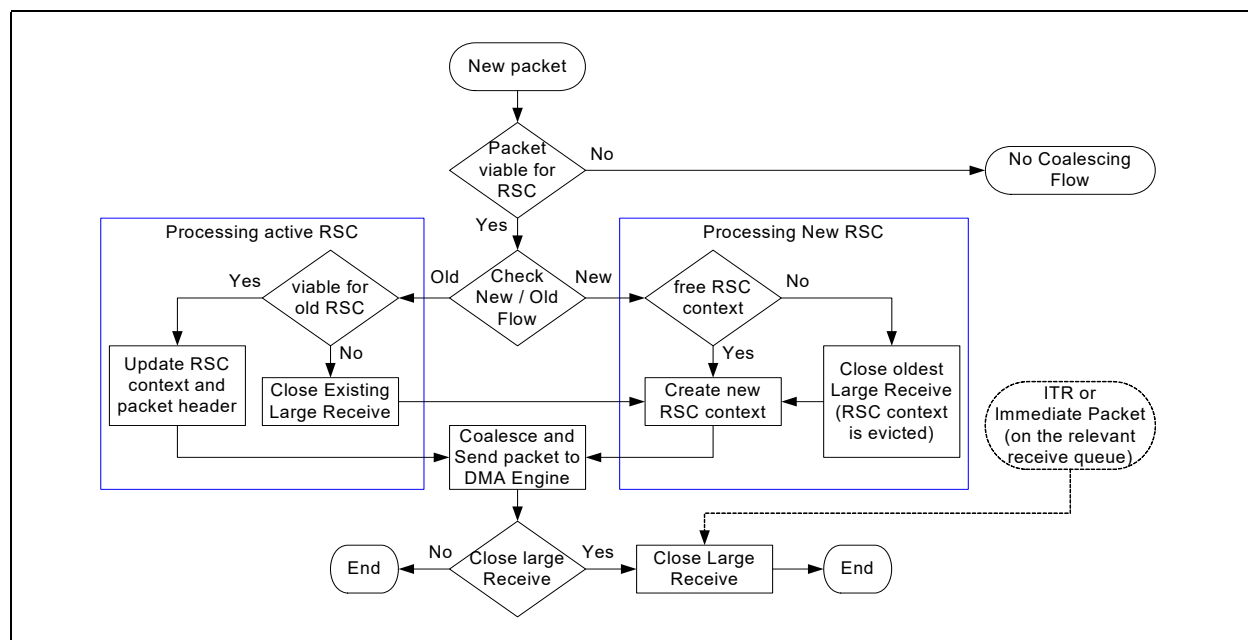


Figure 7.37. RSC Event Flow

7.9.1 Packet Candidacy for RSC

Incoming packets can be good candidates for RSC offload when the following conditions are met. If any of these conditions are not met, and assuming the queue is configured the received packet is processed in the legacy (non-coalescing) scheme.

- RSC is enabled in the destination receive queue by the *RSCCTL.RSCEN*. In this case, software must set the *SRRCTL.DESCTYPE* field in the relevant queues to advanced descriptor modes.
- RSC is further enabled by the *RSCINT.RSCEN* for the receive queues associated to the interrupts defined by the RSCINT registers.
- The *SRRCTL[n].BSIZEHEADER* (header buffer size) must be larger than the packet header (even if header split is not enabled). A minimum size of 128 bytes for the header buffer addresses this requirement.
- The received packet has no MAC errors and no TCP/IP checksum errors. MAC errors are: CRC error or undersize frame received or oversize frame received or error control byte received in mid-packet or illegal code byte received in mid-packet.
- If the *Length* field in the IP header does not cover the entire packet (as the case for padding bytes) then the received packet is not a candidate for RSC.
- The packet type is TCP/IPv4 (non-SNAP) with optional VLAN header(s). RSC is not supported for IPv6 packets.
- The packet is not an NVGRE or VXLAN packet.
- IP header does not carry any option headers.
- The TCP segment is not fragmented.
- The following TCP flags are inactive: FIN, SYN, RST, URG, ECE, CWR, NS and the other three reserved TCP flags (see TCP Flags mapping in [Table 7.61](#)).
- The ECT and CE bits in the TOS field in the IP header are not equal to 11b (see the flags in [Table 7.62](#)).



- Packets with PSH TCP flag are coalesced but also close the large receive.
- The packet does not carry any TCP option headers.
- RSC is not supported for a switched packet transmitted from a local VM.
- When a Rx packet is replicated or mirrored, it can be coalesced only on the Rx queue that belongs to the source VM.
- Note that there are no limitations on the maximum packet length including jumbo packets.
- If there is already an active RSC for the matched flow, then a few additional conditions should be met as listed in [Section 7.9.4](#).

The supported packet format is as follows:

Size	Packet Fields
6 bytes	Destination Ethernet MAC address.
6 bytes	Source Ethernet MAC address.
[4/6/8 bytes]	Optional external tag (VLAN or E-tag).
[4 bytes]	Optional VLAN.
2 bytes	Ethernet type field equals 0x0800 (MS byte first on the wire).
20 bytes	IPv4 header with no options.
20 bytes	Basic TCP header (no options — refer to the rows that follow).
[10 bytes]	Optional TCP time stamp header: 1 byte Kind 0x08 1 byte Length 0x0A 4 bytes TS value variable 4 bytes TS echo reply variable
[1 byte]	Optional TCP no operation header: 1 byte Kind 0x01
[1 byte]	Optional TCP end of option header list: 1 byte Kind 0x00
Variable length	TCP payload (RSC candidate must have payload size greater than zero).

Table 7.61. Packet Format Supported by RSC


11	10	9	8	7	6	5	4	3	2	1	0
Reserved			NS	CWR	ECE	URG	ACK	PSH	RST	SYN	FIN

Table 7.62. IP TOS Field — Bit Map

7	6	5	4	3	2	1	0
TOS (DS)						ECT	CE

Table 7.63. TCP Time-Stamp Option Header (RFC 1323)

1 byte: First on the wire	1 byte	4 byte	4 bytes: Last on the wire
Kind = 0x8	Length = 10	TS Value (TSval)	TS Echo Reply (TSecr)



7.9.2 Flow Identification and RSC Context Matching

TCP/IP packet's flow is identified by its four tuples: Source / Destination IP addresses and Source / Destination TCP port numbers. These tuples are compared against the *Flow Identification* fields stored in the active RSC contexts (listed in [Table 7.64](#)). Comparison is done in two phases:

- Hash compare — Hardware computes a hash value of the four tuples for each flow. The hash value is stored in the RSC context table. It is used for silicon optimization of the compare logic. The hash value of the incoming packet is compared against the hash values of all RSC contexts. No match between the two hash values means that there is no valid context of the same flow.
- Perfect Match — Hardware checks the four tuples of the RSC context that passed the first step with the received frame.
 - A match between the two means that an active RSC context is found.
 - Mismatch between the two indicates a hash collision, which causes a completion of the collided RSC.
- In any case of context mismatch, a new context might be opened as described in [Section 7.9.3](#).
- If the packet's flow matches an active RSC context then the packet might be appended to the existing RSC as described in [Section 7.9.4](#).

Table 7.64. RSC Context

Size	Name	Description
Flow Identification¹		
1 bit	CVALID	Context valid indication. Set to 1b by hardware when a new context is defined. Cleared to 0b when RSC completes.
1 byte	CHASH	Context hash value (logic XOR of all bytes of the four tuples).
16 bytes	IPDADDR	IP destination address (set to zero for inactive context).
16 bytes	IPSADDR	IP source address (set to zero for inactive context).
1 bit	IP4TYPE	Defines IP version type (set to 1b for IPv4).
2 bytes	TCPDPORT	TCP destination port.
2 bytes	TCPSPORT	TCP source port.
37 bytes		Total.
RSC Header²		
2 bytes	RSCIPLN	Total <i>Length</i> field in the IP header defines the size of the IP datagram (IP header and IP payload) in bytes. Dynamic parameter updated by each received packet.
5 bits	IPOFF	The word offset of the IP header within the packet that is transferred to the DMA unit.
1 bit	RSCTS	TCP time stamp header presence indication.
1 bit	RSCACK	ACK bit in the TCP header is a dynamic parameter taken from the last coalesced packet.
1 bit	RSCACKTYPE	ACK packet type indication (ACK bit is set while packet does not has TCP payload).
1 bit	RSCPSH	PSH bit in the TCP header is a dynamic parameter which is a logic OR function of the PSH bits in all packets within the large receive.
2 bits	CE, ECT	ECN bits in the IP.TOS header: <i>CE</i> and <i>ECT</i> .
4 bytes	RSCSEQ	Non-RSCACKTYPE case: Expected sequence number in the TCP header of the next packet. RSCACKTYPE case: The ACK sequence number in the last good packet. Dynamic parameter updated by each received packet.
8 bytes		Total.
DMA Parameters		



Table 7.64. RSC Context (Continued)

7 bits	RXQUEUE	Receive queue index. This parameter is set by the first packet in the RSC and expected to be the same for all packets in the RSC.
4 bits	RSCDESC	Remaining descriptors of this context. The device initialized RSCDESC by the <i>MAXDESC</i> field in the RSCCTL register of the associated receive queue.
4 bits	RSCCNT	Count the number of packets that are started in the current descriptor. The counter starts at 0x1 for each new descriptor. RSCCNT stops incrementing when it reaches 0xF.
8 bytes	HPTR	Header buffer pointer defines the address in host memory of the large receive header (see Section 7.9.5.2).
2 bytes	DATDESC	Data descriptor is the active descriptor index. Initialized by the first packet in the RSC to the first descriptor. It is updated to the active descriptor at a packet DMA completion.
2 bytes	DATOFF	Offset within the data buffer. The data of the first packet in a large receive is the same as the legacy (non-coalescing) definition. Following a DMA completion, it points to the beginning of the data portion of the next packet.
13 bytes		Total.

1. These parameters are extracted from the first packet that opens (activate) the context.
2. All parameters are set by the first packet that opens the context while some are dynamic.

7.9.3 Processing New RSC

Defining the RSC context parameters activates a new large receive. If a received packet does not match any active RSC context, the packet starts (opens) a new one. If there is no free context, the oldest active large receive is closed and its evicted context is used for the new large receive.


7.9.3.1 RSC Context Setting

The Intel® Xeon® Processor D-1500 Product Family LAN controller extracts the flow identification and RSC header parameters from the packet that opens the context (the first packet in a large receive that activates an RSC context). The context parameters can be divided into categories: flow identification; RSC header and DMA parameters.

7.9.4 Processing Active RSC

Received packets that belong to an active RSC can be added to the large receive if all the following conditions are met:

- The L2 header size equals the size of previous packets in the RSC as recorded in the internal IPOFF parameter in the RSC context table.
- The packet header length as reported in the HDR_LEN field is assumed to be the same as the first packet in the RSC (not checked by hardware).
- The ACK flag in the TCP header is equal to the RSCACK bit in the RSC context (The value of the ACK flag should be constant in all the coalesced packets).
- The packet type remains the same as indicated by the RSCACKTYPE bit in the RSC context. Packet type can be either pure ACK packet (with no TCP payload) or other.
- For non-RSCACKTYPE (packet with TCP payload): The sequence number in the TCP header matches the expected value in the RSC context (RSCSEQ).
- For RSCACKTYPE: The acknowledgment number in the TCP header is greater than the RSCSEQ number in the RSC context. Note that the Intel® Xeon® Processor D-1500 Product Family LAN controller does not coalesce duplicated ACK nor ACK packets that only updates the TCP window.
- ECN handling: The value of the CE and ECT bits in the *IP.TOS* field remains the same as the RSC context and different than 11b.
- The target receive queue matches the RXQUEUE in the RSC context.



- The packet does not include a TCP time stamp header unless it was included on the first packet that started the large receive (indicated by the RSCTS). Note that if the packet includes other option headers than time stamp, NOP or end of option header, the packet is not processed by RSC flow at all.
- The packet fits within the RSC buffer(s).
- If the received packet does not meet any of the previous conditions, the matched active large receive is closed. Then hardware opens a new large receive by that packet. Note that since the Intel® Xeon® Processor D-1500 Product Family LAN controller closes the old large receive it is guaranteed that there is at least one free context.

Note: See [Section 7.1.6.2](#) for impact of time stamp in packet on RSC decisions.

If the received packet meets all the previous conditions, the Intel® Xeon® Processor D-1500 Product Family LAN controller appends this packet payload to the active large receive and updates the context and header as follows. The packet is then DMA'ed to the RSC buffers (as described in [Section 7.9.5](#)).

- Update the TCP PSH: The *PSH* bit in the large receive context gets the value of the *PSH* bit in the TCP header of the received packet.
- Update the expected sequence number for non-RSCACKTYPE: The RSCSEQ in the large receive context is increased by the value of the TCP payload size of the received packet.
- Update the expected sequence number for RSCACKTYPE: The RSCSEQ in the large receive context is updated to the value of the ACK sequence number field in the received packet.
- Update the total length: The RSCIPLN in the large receive context is increased by the value of the TCP payload size of the received packet. The value of the *Total Length* field in the IP header in the received packet gets the updated RSCIPLN. Note that in RSCACKTYPE packets the received payload size is zero.
- Acknowledgment number: If the ACK flag is set, the *Acknowledgment Number* field in the header is updated. IP header checksum is modified to reflect the changes in the *Total Length* field as follows (note that there is no special process for RSCACKTYPE packets):

1's $\{(RSCIPLN - \text{Packet total length}) + 1\}$'s (Packet IP header checksum) while...

- Packet total length is the total length value in the received packet.
- Packet IP header checksum stands for the IP header checksum field in the received packet.
- 1's operation defines a ones complement.
- Plus (+) operation is a cyclic plus while the carry out is fed as a carry in.
- TCP header checksum is left as is in the first packet in the RSC and is set to zero on any succeeding packets.
- Update the DMA parameters.
 - The RSCCNT is initialized to 0x1 on each new descriptor. It is then increased by one on each packet that starts on the same descriptor as long as it does not exceed a value of 0xF. When the RSCCNT is set to 0xF (14 packets) the RSC completes.
 - Decrement by one the Remaining Descriptors (RSCDESC) for each new descriptor.
 - Update the receive descriptor index (DATDESC) for each new descriptor.
 - Update the offset within the data buffer (DATOFF) at the end of the DMA to its valid value for the next packet.
- All other fields are kept as defined by the first packet in the large receive.

7.9.5 Packet DMA and Descriptor Write Back

The Figure 7.38 shows a top view of the RSC buffers using advanced receive descriptors and header split descriptors.

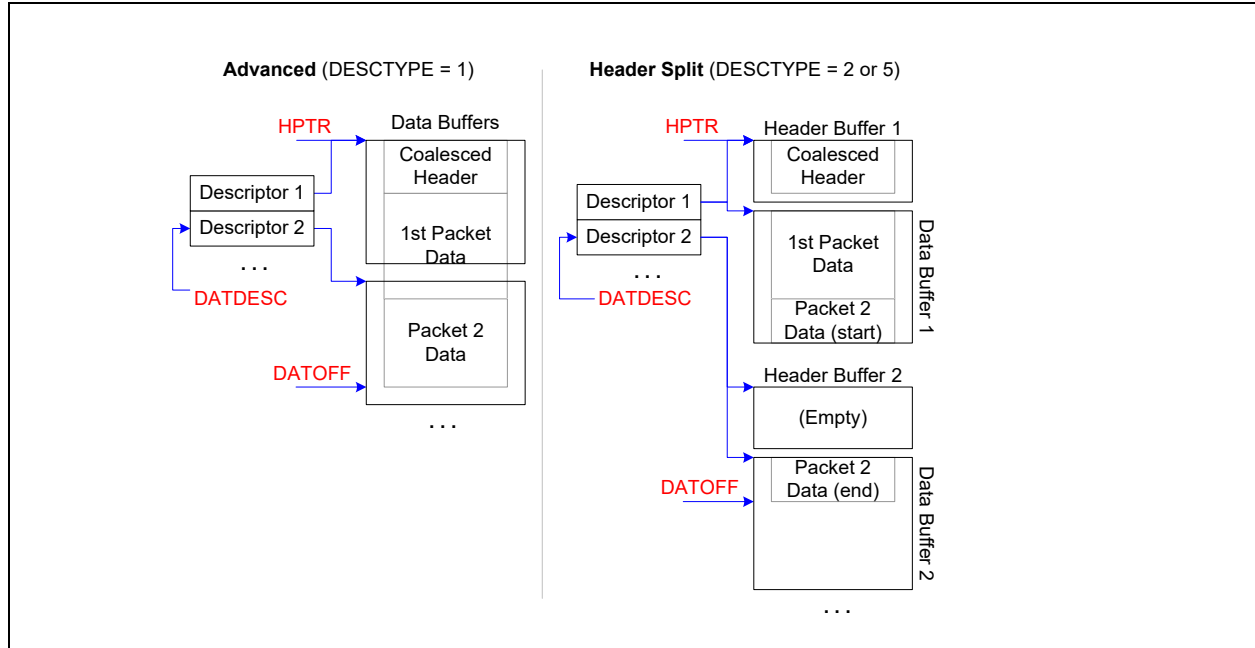


Figure 7.38. RSC — Header and Data Buffers

7.9.5.1 RSC Descriptor Indication (Write Back)

After receiving each packet, the Intel® Xeon® Processor D-1500 Product Family LAN controller posts the packet data to the data buffers and updates the coalesced header in its buffer. Any completed descriptor is indicated (write back) by setting the fields listed in the following table. A descriptor is defined as the last one when an RSC completes. Section 7.9.6 summarizes all the causes for RSC completion. Any other descriptor in the middle of the RSC is indicated (write back) when the hardware requires the next descriptor so it can report the NEXTP field explained in the table that follows.

Fields on the Last Descriptors of Large Receive	Fields on All Descriptors Except for the Last One
EOP: End of packet, and all other fields that are reported together with the EOP.	NEXTP: Points to the next descriptor of the same large receive.
DD: indicates that this descriptor is completed by the hardware and can be processed by the software. The descriptor and the header of the packet may still be updated by the hardware.	
RSCCNT: indicates the number of coalesced packets in this descriptor.	

7.9.5.2 Received Data DMA

On the first packet of a large receive, the entire packet is posted to its buffers in host memory. On any other packet, the packet's header and data are posted to host memory as detailed in Section 7.9.5.3 and Section 7.9.5.4 that follows.



7.9.5.3 RSC Header

The RSC header is stored at the beginning of the first buffer when using advanced receive descriptors, or at the header buffer of the first descriptor when using header split descriptors (it is defined by the internal HPTR parameter in the RSC context). See [Figure 7.38](#) for more details.

The packet's header is posted to host memory after it is updated by the RSC context as follow:

Packets with payload coalescing (RSCACKTYPE=0) - The TCP sequence number is taken from the TCP context (it is taken from the first packet). The *Total Length* field in the IP header is taken from the RSC context (it represent the length of all coalesced packets). The IP checksum is re-calculated. The TCP checksum is set to zero.

ACK no payload coalescing (RSCACKTYPE=1) - The received packet header is posted as is to host memory. Note that if the received packet includes padding bytes, these bytes are posted in the host memory as well.

7.9.5.4 Large Receive Data

The data of a coalesced packet is posted to its buffer by the DMA engine as follows.

Ethernet CRC.

- When RSC is enabled on any queue, the global CRC strip must be set (HLREG0.RXCRCSTRP =1b).

Packet data spans on a single buffer.

- The data of the received packet spans on a single buffer if buffer has the required space.
- The DMA engine posts the packet data to its buffer pointed to by DATDESC descriptor at an offset indicated by the DATOFF.

Packet data spans on multiple buffers.

- The data of the received packet spans across multiple buffers when it is larger than a single buffer or larger than the residual size of the current buffer.
- When a new buffer is required (new descriptor) the DMA engine writes back to the completed descriptor linking it to the new one ([Section 7.9.5.1](#) details the indicated descriptor fields).
- Decrement the RSCDESC parameter by one and update the DATDESC for each new opened descriptor.

DMA completion.


- Following DMA completion, set the DATOFF to the byte offset of the next packet.
- If the PSH TCP flag is active in the coalesced packet then the large receive is completed.

Not enough descriptors in the receive ring buffer.

- If the *SRRCTL[n].Drop_En* bit on the relevant queue is set, The large receive completes and the new packet is discarded.
- Otherwise (the *Drop_En* bit is cleared), the packet waits inside the internal packet buffer until new descriptors are added (indicated by the relevant Tail register).

Not enough descriptors due to RSCDESC exhaust.

- If the received packet requires more descriptors than indicated by the internal RSCDESC parameter, then the Intel® Xeon® Processor D-1500 Product Family LAN controller completes the current large receive while the new packet starts a new large receive.



7.9.6 RSC Completion and Aging

This section summarizes all causes of large receive completion (the first three cases repeat previous sections).

- A packet of a new flow is received while there are no free RSC contexts. the Intel® Xeon® Processor D-1500 Product Family LAN controller completes (closes) the oldest large receive (opened first). The new packet starts a new large receive using the evicted context.
- Packets with PSH TCP flags can be only the last ones in active large receive.
- The received packet cannot be added to the active large receive due to one of the following cases (indicated also in [Section 7.9.4](#)). In these cases the existing RSC completes and the received packet opens a new large receive.
 - The sequence number does not meet expected value.
 - The receive packet includes a time stamp TCP option header while there was no time stamp TCP option header in the first packet in the RSC.
 - There is not enough space in the RSC buffer(s) for the packet data. Meaning, the received packet requires a new buffer while the RSC already exhausted all permitted buffers defined by the RSCCTL[n].MAXDESC.
 - The received packet requires a new buffer while its descriptor wraps around the descriptor ring.
- When a packets is received while there are no more descriptors in the receive queue and the *SRRCTL.Drop_En* bit is set, the large receive completes and the new packet is discarded.
- EITR expiration while interrupt is enabled — RSC completion is synchronized with interrupt assertion to the host. It enables software to process the received frames since the last interrupt. See more details and EITR setting in [Section 7.3.2.1.1](#).
- EITR expiration while interrupt is disabled — The ITR counter continues to count even when its interrupt is disabled. Every time the timer expires it triggers RSC completion on the associated Rx queues.
- Low number of available descriptors — Whenever crossing the number of free Rx descriptors, the receive descriptor minimum threshold size defined in the SRRCTL[n] registers.
- Interrupt assertion by setting the EICS register has the same impact on packet reception as described in [Section 7.3.1.2.1](#).

Note: In some cases packets that do not meet coalescing conditions might have active RSC of the same flow. As an example: received packets with ECE or CWR TCP flags. Such packets bypass completely the RSC logic (posted as single packets), and do not cause a completion of the active RSC. The active RSC would eventually be closed by either reception of a legitimate packet that is processed by the RSC logic but would not have the expected TCP sequence number. Or, an interrupt event closes all RSC’s in its Rx queue. When software processes the packets, it gets them in order even though the RSC completes after the previous packet(s) that bypassed the RSC logic.

Any interrupt closes all RSC’s on the associated receive queues. Therefore, when ITR is not enabled any receive packet causes an immediate interrupt and receive coalescing should not be enabled on the associated Rx queues.

Software Note: Each time an interrupt is generated all active large receives on the relevant queues are completed. The large receives are indicated to host memory before the interrupt is asserted.

Cross Function	Requirements
Ethernet CRC insertion	Software must enable Ethernet CRC insertion by setting the <i>IFCS</i> bit in the transmit data descriptor. The Ethernet CRC covers the entire packet.
VLAN header	In the case of double VLAN mode, the packet must have the two VLAN headers.
SNAP packet	Not supported.



Cross Function	Requirements
Traffic rate control	FC traffic relies on a high quality link that guarantees no packet loss. It is expected that any lost traffic protocols supported by the network are enabled by the Intel® Xeon® Processor D-1500 Product Family LAN controller as well.
FC	
Virtualization	Not supported.
TCP/IP and UDP/IP offload	Any setting of TCP/IP and UDP/IP offload capabilities are not applicable.
Transmit descriptors	Software must use the advanced transmit descriptor to activate either FC CRC offload or TSO functionality.

7.10 Reliability


7.10.1 Memory Integrity Protection

All Intel® Xeon® Processor D-1500 Product Family LAN Controller internal memories are protected against soft errors. Most of them are covered by ECC that correct single error per memory line and detect double errors per memory line. Few of the smaller memories are covered by parity protection that detects a single error per memory line.

Single errors in memories with ECC protection are named also as correctable errors. Such errors are silently corrected. Two errors in memories with ECC protection or single error in memories with parity protection are also named as un-correctable errors. Un-correctable errors are considered as fatal errors. If an un-correctable error is detected in Tx packet data, the packet is transmitted with a CRC error. If un-correctable error is detected in Rx packet data, the packet is reported to the host (or manageability) with a CRC error. If an un-correctable error is detected anywhere else, the Intel® Xeon® Processor D-1500 Product Family LAN controller halts the traffic and sets the ECC error interrupt. Software is then required to initiate a complete initialization cycle to resume normal operation.

7.10.2 PCIe Error Handling

For PCIe error events and error reporting see [Section 3.1.3](#).



NOTE: *This page intentionally left blank.*



8.0 Programming Interface

8.1 General

The Intel® Xeon® Processor D-1500 Product Family LAN controller's address space is mapped into four regions with PCI Base Address registers listed in [Table 8.1](#) and explained more in [Section 9.2.2.16](#) and [Section 9.2.2.17](#).

Table 8.1. Broadwell-DE Address Regions

Addressable Content	Mapping Style	Region Size
Internal registers memories and NVM (Memory BAR)	Direct memory mapped	128 KB + NVM Size ¹
Flash (optional) ²	Direct memory-mapped	2 MB - 4 MB
Expansion ROM (optional)	Direct memory-mapped	64 KB - 512 KB
Internal registers and memories (optional) ³	I/O window mapped	32 bytes
MSI-X (optional)	Direct memory mapped	16 KB ⁴

1. The Flash size is defined by the BARCTRL register.
2. The Flash space in the memory CSR and expansion ROM base address are mapped to different Flash memory regions. Accesses to the memory BAR at offset 256 KB are mapped to the Flash device at offset 0x0, while accesses to the expansion ROM at offset 0x0 are mapped to the Flash region pointed by the PXE driver module pointer (read from NVM word address 0x05). See [Section 6.1](#). The expansion ROM region has a size limited to 512 KB.
3. The internal registers and memories can be accessed through I/O space as explained in the sections that follow.
4. See [Section 9.2.2.16](#) for the MSI-X BAR offset in 32-bit and 64-bit BAR options.

8.1.1 Memory-Mapped Access

8.1.1.1 Memory-Mapped Access to Internal Registers and Memories

The internal registers and memories can be accessed as direct memory-mapped offsets from the memory CSR BAR. See the following sections for detailed description of the Device registers.

In IOV mode, this area is partially duplicated per VF. All replications contain only the subset of the register set that is available for VF programming.

8.1.1.2 Memory-Mapped Accesses to Flash

The external Flash can be accessed using direct memory-mapped offsets from the CSR base address register (BAR0/BAR1). The Flash is only accessible if enabled through the `PCI_LBARCTRL.FLASH_EXPOSE` bit in the PCIe general configuration NVM module. 256 KB should be added to the physical address within the external Flash device to get the offset from the BAR.



8.1.1.3 Memory-Mapped Access to MSI-X Tables

The MSI-X tables can be accessed as direct memory-mapped offsets from the base address register (BAR3). The MSIX registers are described in the MSI-X Table registers section.

In IOV mode, this area is duplicated per VF.

8.1.1.4 Memory-Mapped Access to Expansion ROM

The option ROM module located in the external Flash can also be accessed as a memory-mapped expansion ROM. Accesses to offsets starting from the expansion ROM base address reference the Flash, provided that access is enabled through the NVM Initialization Control Word, and if the Expansion ROM Base Address register contains a valid (non-zero) base memory address.

8.1.2 I/O-Mapped Access

All internal registers and memories can be accessed using I/O operations. I/O accesses are supported only if an I/O base address is allocated and mapped (BAR2), the BAR contains a valid (non-zero value), and I/O address decoding is enabled in the PCIe configuration.

When an I/O BAR is mapped, the I/O address range allocated opens a 32-byte window in the system I/O address map. Within this window, two I/O addressable registers are implemented: IOADDR and IODATA. The IOADDR register is used to specify a reference to an internal register or memory, and then the IODATA register is used to access it at the address specified by IOADDR:

Offset	Abbreviation	Name	RW	Size
0x00	IOADDR	Internal Register, Internal Memory, or Flash Location Address. 0x00000-0x1FFFF – Internal registers/memories 0x20000-0x7FFFF – Undefined	RW	4 bytes
0x04	IODATA	Data field for reads or writes to the internal register, internal memory, or Flash location as identified by the current value in IOADDR. All 32 bits of this register are read/write-able.	RW	4 bytes
0x08-0x1F	Reserved	Reserved.	0	4 bytes

8.1.2.1 IOADDR (I/O Offset 0x00; RW)

The IOADDR register must always be written as a Dword access. Writes that are less than 32 bits are ignored. Reads of any size returns a Dword of data; however, the chipset or CPU might only return a subset of that Dword.

For software programmers, the IN and OUT instructions must be used to cause I/O cycles to be used on the PCIe bus. Because writes must be to a 32-bit quantity, the source register of the OUT instruction must be EAX (the only 32-bit register supported by the OUT command). For reads, the IN instruction can have any size target register, but it is recommended that the 32-bit EAX register be used.

Bits 31 through 20 are ignored by the Intel® Xeon® Processor D-1500 Product Family LAN controller and should be set to zero by software. At hardware reset (Power Good Reset) or PCI reset, this register value resets to 0x00000000. Once written, the value is retained until the next write or reset.



8.1.2.2 IODATA (I/O Offset 0x04; RW)

The IODATA register must always be written as a Dword access when the IOADDR register contains a value for the internal register and memories (such as 0x00000-0x1FFFC). In this case, writes that are less than 32 bits are ignored.

Writes and reads to IODATA when the IOADDR register value is in an undefined range (0x20000-0x7FFFC) should not be performed. Results cannot be determined.

Note: There are no special software timing requirements on accesses to IOADDR or IODATA. All accesses are immediate except when data is not readily available or acceptable. In this case, the Intel® Xeon® Processor D-1500 Product Family LAN controller delays the results through normal bus methods (like split transaction or transaction retry).

Because a register/memory read or write takes two I/O cycles to complete, software must provide a guarantee that the two I/O cycles occur as an atomic operation. Otherwise, results can be non-deterministic from the software viewpoint.

8.1.2.3 Undefined I/O Offsets

I/O offsets 0x08 through 0x1F are considered to be reserved offsets with the I/O window. Dword reads from these addresses return 0xFFFF; writes to these addresses are discarded.

8.1.3 Configuration Access to Internal Registers and Memories

To support legacy pre-boot 16-bit operating environments without requiring I/O address space, the Intel® Xeon® Processor D-1500 Product Family LAN controller enables accessing CSRs via the configuration address space by mapping *IOADDR* and *IODATA* registers into the configuration address space. The register mapping in this case is listed in [Table 8-2](#).

Table 8-2. IOADDR and IODATA in Configuration Address Space

Configuration Address	Abbreviation	Name	RW	Size
0x98	IOADDR	Internal register or internal memory location address. 0x00000-0x1FFFF – Internal registers and memories. 0x20000-0x7FFFF – Undefined.	RW	4 bytes
0x9C	IODATA	Data field for reads or writes to the internal register or internal memory location as identified by the current value in IOADDR. All 32 bits of this register can be read from or written to.	RW	4 bytes

Software writes data to an internal CSR via the configuration space in the following manner:

1. CSR address is written to the IOADDR register where:
 - a. Bit 31 (*IOADDR.Configuration IO Access Enable*) of the IOADDR register should be set to 1b.
 - b. Bits 30:0 of IOADDR should hold the actual address of the internal register or memory being written to.
2. Data to be written is written into the IODATA register.
 - The IODATA register is used as a window to the register or memory address specified by IOADDR register. As a result, the data written to the IODATA register is written into the CSR pointed to by bits 30:0 of the IOADDR register.
3. *IOADDR.Configuration IO Access Enable* is cleared to avoid un-intentional CSR read operations (that might cause clear by read) by other applications scanning the configuration space.



Software reads data from an internal CSR via the configuration space in the following manner:

1. CSR address is written to the IOADDR register where:
 - a. Bit 31 (*IOADDR.Configuration IO Access Enable*) of the IOADDR register should be set to 1b.
 - b. Bits 30:0 of IOADDR should hold the actual address of the internal register or memory being read.
2. CSR value is read from the IODATA register.
 - a. The IODATA register is used as a window to the register or memory address specified by the IOADDR register. As a result, the data read from the IODATA register is the data of the CSR pointed to by bits 30:0 of the IOADDR register.
3. *IOADDR.Configuration IO Access Enable* is cleared to avoid unintentional CSR read operations (that might cause clear by read) by other applications scanning the configuration space.

Notes:

- When functioning in a D3 state, software should not attempt to access CSRs via the IOADDR and IODATA configuration registers.
- To enable CSR access via configuration space, software should set bit 31 to 1b (*IOADDR.Configuration IO Access Enable*) in the IOADDR register. Software should clear bit 31 of the IOADDR register after completing CSR access to avoid an unintentional clear-by-read operation by another application scanning the configuration address space.
- Bit 31 of the IOADDR register (*IOADDR.Configuration IO Access Enable*) has no effect when initiating access via I/O address space.
- Configuration access to 0x9C (IODATA) without setting bit 31 of the IOADDR register (*IOADDR.Configuration Access Enable*) must not result in an unsupported request.
- I/O-mapped access and CSR access via PCIe configuration space are mutually exclusive operating modes, and therefore *PCIE_CNFG.IO_SUP* NVM bit must be cleared when the *PCIE_CAPSUP.CSR_CONF_EN* NVM bit is set, and vice versa.

8.1.4 Register Terminology

Shorthand	Description
RW	Read/Write. A register with this attribute can be read and written. If written since reset, the value read reflects the value written.
RO	Read Only. If a register is read only, writes to this register have no effect.
ROS	Read Only Status. Writes to ROS fields has no effect. The value of the field can change due to Hardware events.
WO	Write Only. Reading this register might not return a meaningful value.
RW1C	Read/Write Clear. A register with this attribute can be read and written. However, a write of a 1b clears (sets to 0b) the corresponding bit and a write of a 0b has no effect.
RC	Read Clear. A register bit with this attribute is cleared after read. Writes have no effect on the bit value.
RW/RC	Read/Write and Read Clear.
RWS	Read Write Set: Register that is set to 1b by software by writing a 1b to the register, and cleared to 0b by hardware.
Reserved/RSV	Reserved. field can return any value on read access and must be set to its initialized value on write access unless specified differently in the field description.



8.1.5 VF Registers Allocated per Queue

Depending on configuration, each pool has 2, 4, or 8 queues allocated to it. Note that in IOV mode, any queues not allocated to a VF are allocated to the PF. The registers assigned to a queue are accessible both in its VF address space and in the PF address space. This section describes the address mapping of registers that belong to queues.

Section 7.7.2.7.2 defines the correspondence of queue indices between the PF and the VFs. For example, when in configuration for 32 VFs, queues 124-127 in the PF correspond to queues [3:0] of VF# 31.

The queues are enumerated in each VF from 0 (such as [1:0], [3:0], or [7:0]). If a queue is allocated to a VF, then its corresponding registers are accessible in the VF CSR space. Each register is allocated an address in the VF (relative to its base) according to its index in the VF space. Therefore, the registers of queue 0 in each VF are allocated the same addresses, which equal the addresses of the same registers for queue 0 in the PF. For example, RDH[0] in the VF space has the same relative address in each VF and in the PF (address 0x01010).

8.1.6 Non-queue VF Registers

Registers that do not correspond to a specific queue are allocated addresses in the VF space according to these rules:

- Registers that are read-only by the VF (such as STATUS) have the same address in the VF space as in the PF space.
- Registers allocated per pool are accessed in the VF in the same location as pool [0] in the PF address space.
- Registers that are RW by the VF (like CTRL) are replicated in the PF, one per VF, in adjacent addresses.


Note: Since the VF address space is limited to 16 KB, any register that resides above that address in the PF space cannot reside in the same address in the VF space and is therefore allocated in another location in the VF.

8.1.7 Access to MAC Registers

When working in SGMII 100 Mb/s or 10 Mb/s speeds, the core and MAC layer have separate clock frequencies therefore there is a delay between the time a posted write is issued by software until the time the actual hardware operation is completed.

In order to prevent coherency problems, software is required to check and insure that the register is ready to be accessed.

This limitation is not applicable to all registers. The relevant registers that are subject to this limitation, each have a BUSY signal in the MAC SGMII Busy register.



NOTE: *This page intentionally left blank.*



8.2 Device Registers - PF

8.2.1 BAR0 Registers Summary

Table 8.3. BAR0 Registers Summary

Offset / Alias Offset	Abbreviation	Name	Section
General Control Registers			
0x00000000	CTRL	Device Control Register	Section 8.2.2.1.1
0x00000008	STATUS	Device Status Register	Section 8.2.2.1.2
0x00000018	CTRL_EXT	Extended Device Control Register	Section 8.2.2.1.3
0x00000020	ESDP	Extended SDP Control	Section 8.2.2.1.4
0x00000028	PHY_GPIO	PHY GPIO Register	Section 8.2.2.1.5
0x00000030	MAC_GPIO	MAC GPIO Register	Section 8.2.2.1.6
0x00000100	PHYINT_STATUS0	PHY Interrupt Status Register 0	Section 8.2.2.1.7
0x00000104	PHYINT_STATUS1	PHY Interrupt Status Register 1	Section 8.2.2.1.8
0x00000108	PHYINT_STATUS2	PHY Interrupt Status Register 2	Section 8.2.2.1.9
0x00000200	LEDCTL	LED Control	Section 8.2.2.1.10
0x00005078	EXVET	Extended VLAN Ether Type - Receive	Section 8.2.2.1.11
0x00008224	EXVET_T	Extended VLAN Ether Type - Transmit	Section 8.2.2.1.12
0x00010200	GRC	General Receive Control	Section 8.2.2.1.13
0x00015F58	I2CCMD	SFP I ² C Command	Section 8.2.2.1.14
0x00015F5C	I2CPARAMS	SFP I ² C Parameters	Section 8.2.2.1.15
0x00010208	DEV_FUNC_EN	Device and Functions Enable Control	Section 8.2.2.1.16
0x00010150	FACTPS	Function Active and Power State to Manageability	Section 8.2.2.1.17
NVM Registers			
0x00010010	EEC	EEPROM Mode Control Register	Section 8.2.2.2.1
0x00010014	EERD	EEPROM Mode Read Register	Section 8.2.2.2.2
0x00010018	EEWR	EEPROM Mode Write Register	Section 8.2.2.2.3
0x0001001C	FLA	Flash Access Register	Section 8.2.2.2.4
0x00015F54	FLUPDATE	Flash Firmware Code Update	Section 8.2.2.2.5
0x00015F48	FLSWCTL	Software Flash Burst Control Register	Section 8.2.2.2.6
0x00015F2C	JEDEC_ID_1	JEDEC ID 1	Section 8.2.2.2.7
0x00010118	FLMNGCTL	Manageability Flash Control Register	Section 8.2.2.2.8
0x00015F4C	FLSWDATA	Software Flash Burst Data Register	Section 8.2.2.2.9
0x00015F50	FLSWCNT	Software Flash Burst Access Counter	Section 8.2.2.2.10
0x00010110	EEMNGCTL	Manageability EEPROM-Mode Control Register	Section 8.2.2.2.11
0x00010114	EEMNGDATA	Manageability EEPROM-Mode Read/Write Data	Section 8.2.2.2.12
0x0001011C	FLMNGDATA	Manageability Flash Read/Write Data	Section 8.2.2.2.13
Flow Control Registers			
0x00004294	MFLCN	MAC Flow Control Register	Section 8.2.2.2.14
0x00003200 + 0x4*n, n=0...3	FCTTVN[n]	Flow Control Transmit Timer Value n	Section 8.2.2.2.15



Table 8.3. BARO Registers Summary (Continued)

Offset / Alias Offset	Abbreviation	Name	Section
0x00003220 + 0x4*n, n=0...7	FCRTL[n]	Flow Control Receive Threshold Low	Section 8.2.2.2.16
0x00003260 + 0x4*n, n=0...7	FCRTH[n]	Flow Control Receive Threshold High	Section 8.2.2.2.17
0x000032A0	FCRTV	Flow Control Refresh Threshold Value	Section 8.2.2.2.18
0x0000CE00	TFCS	Transmit Flow Control Status	Section 8.2.2.2.19
0x00003D00	FCCFG	Flow Control Configuration	Section 8.2.2.2.20
0x0000431C	PFCTOP	Priority Flow Control Type Opcode	Section 8.2.2.2.21
PCIe Registers			
0x00011028	PCI_STATUS1	PCIe Function Status 1	Section 8.2.2.3.1
0x00011050	GCR_EXT	PCIe Control Extended Register	Section 8.2.2.3.2
0x00011098	PCI_REVID	PCIe Revision ID	Section 8.2.2.3.3
0x00011540	PCI_DREVID	PCIe Default Revision ID	Section 8.2.2.3.4
0x00011070	FUNC_RID	Function Requester ID Information Register	Section 8.2.2.3.5
0x00011530	PCI_VMINDEX	PCIe VM Pending Index	Section 8.2.2.3.6
0x00011538	PCI_VMPEND	PCIe VM Pending Status	Section 8.2.2.3.7
0x00011528	PCI_IENA	PCIe Interrupts Enable	Section 8.2.2.3.8
0x00011520	PCI_ICAUSE	PCIe Interrupt Cause	Section 8.2.2.3.9
0x00011140	PCI_PCIERR	PCIe Errors Reported	Section 8.2.2.3.10
0x00011800	PCI_GSCL_1	PCIe Statistic Control Register #1	Section 8.2.2.3.11
0x00011804	PCI_GSCL_2	PCIe Statistic Control Registers #2	Section 8.2.2.3.12
0x00011820 + 0x4*n, n=0...3	PCI_GSCN_0_3[n]	PCIe Statistic Counter Registers #0...#3	Section 8.2.2.3.13
0x00011810 + 0x4*n, n=0...3	PCI_GSCL_5_8[n]	PCIe Statistic Control Register #5...#8	Section 8.2.2.3.14
0x00011548	PCI_BYTCTL	PCIe Byte Counter Low	Section 8.2.2.3.15
0x00011544	PCI_BYTCTH	PCIe Byte Counter High	Section 8.2.2.3.16
0x00011720	PCI_LATCT	PCIe Latency Counter	Section 8.2.2.3.17
0x00011740	PCI_PKTCT	PCIe Packet Counter	Section 8.2.2.3.18
0x00011788	PCI_LCBADD	PCIe LCB Address Port	Section 8.2.2.3.19
0x00011734	PCI_LCBDATA	PCIe LCB Data Port	Section 8.2.2.3.20
0x00011040	PCI_FWCTRL	PCIe Firmware Control	Section 8.2.2.3.21
0x00011044	PCI_CSRTO	PCIe CSR Access Timeout	Section 8.2.2.3.22
0x00011054	PCI_FLASHTO	PCI Flash Access Timeout	Section 8.2.2.3.23
PCIe Configuration Space Setting Registers			
0x000110A8	PCI_VFSUP	PCIe VF Capabilities Support	Section 8.2.2.4.1
0x000110F8	PCI_CNF2	PCIe Global Config 2	Section 8.2.2.4.2
0x00011048	PCI_LBARCTRL	PCI BAR Control	Section 8.2.2.4.3
0x000110B8	PCI_GLBL_CNF	PCIe Global Config	Section 8.2.2.4.4
0x00011088	PCI_CAPSUP	PCIe Capabilities Support	Section 8.2.2.4.5
0x00011090	PCI_LINKCAP	PCIe Link Capabilities	Section 8.2.2.4.6
0x00011080	PCI_CAPCTRL	PCIe Capabilities Control	Section 8.2.2.4.7
0x000110F0	PCI_SERL	PCIe Serial Number MAC Address Low	Section 8.2.2.4.8




Table 8.3. BAR0 Registers Summary (Continued)

Offset / Alias Offset	Abbreviation	Name	Section
0x00011078	PCI_SERH	PCIe Serial Number MAC Address High	Section 8.2.2.4.9
0x00011038	PCI_CLASS	PCIe Storage Class	Section 8.2.2.4.10
0x00011000	PCI_CNF	PCIe PF Configuration	Section 8.2.2.4.11
0x000110A0	PCI_PMSUP	PCIe PM Support	Section 8.2.2.4.12
0x00011060	PCI_PWRDATA	PCIe Power Data Register	Section 8.2.2.4.13
0x0001103C	PCI_VENDORID	PCIe Vendor ID	Section 8.2.2.4.14
0x00011008	PCI_PFDEVID	PCIe PF Device ID	Section 8.2.2.4.15
0x00011010	PCI_VFDEVID	PCIe VF Device ID	Section 8.2.2.4.16
0x00011058	PCI_SUBSYSID	PCIe Subsystem ID	Section 8.2.2.4.17
0x000110E8	PCI_UPADD	PCIe Upper Address	Section 8.2.2.4.18
Interrupt Registers			
0x00000800	EICR	Extended Interrupt Cause Register	Section 8.2.2.4.19
0x00000808	EICS	Extended Interrupt Cause Set Register	Section 8.2.2.4.20
0x00000880	EIMS	Extended Interrupt Mask Set/Read Register	Section 8.2.2.4.21
0x00000888	EIMC	Extended Interrupt Mask Clear Register	Section 8.2.2.4.22
0x00000810	EIAC	Extended Interrupt Auto Clear Register	Section 8.2.2.4.23
0x00000890	EIAM	Extended Interrupt Auto Mask Enable register	Section 8.2.2.4.24
0x00000A90	EICS1	Extended Interrupt Cause Set Registers 1	Section 8.2.2.4.25
0x00000A94	EICS2	Extended Interrupt Cause Set Registers 2	Section 8.2.2.4.26
0x00000AA0	EIMS1	Extended Interrupt Mask Set/Read Registers	Section 8.2.2.4.27
0x00000AA4	EIMS2	Extended Interrupt Mask Set/Read Registers	Section 8.2.2.4.28
0x00000AB0	EIMC1	Extended Interrupt Mask Clear Registers 1	Section 8.2.2.4.29
0x00000AB4	EIMC2	Extended Interrupt Mask Clear Registers 2	Section 8.2.2.4.30
0x00000AD0	EIAM1	Extended Interrupt Auto Mask Enable registers 1	Section 8.2.2.4.31
0x00000AD4	EIAM2	Extended Interrupt Auto Mask Enable registers 2	Section 8.2.2.4.32
0x00000894	EITRSEL	MSIX to EITR Select	Section 8.2.2.4.33
0x00000820 + 0x4*n, n=0...23 and 0x00012300 + 0x4*(n-24), n=24...128	EITR[n]	Extended Interrupt Throttle Registers	Section 8.2.2.4.34
0x00000900 + 0x4*n, n=0...63	IVAR[n]	Interrupt Vector Allocation Registers	Section 8.2.2.4.35
0x00000A00	IVAR_MISC	Miscellaneous Interrupt Vector Allocation	Section 8.2.2.4.36
0x00012000 + 0x4*n, n=0...128	RSCINT[n]	RSC Enable Interrupt	Section 8.2.2.4.37
0x00000898	GPIE	General Purpose Interrupt Enable	Section 8.2.2.4.38
MSI-X Table Registers			
0x000110C8 + 0x4*n, n=0...5	VFPBACL[n]	VF MSI-X PBA Clear	Section 8.2.2.5.1
0x000110C0 + 0x4*n, n=0...1	PBACL[n]	MSI-X PBA Clear	Section 8.2.2.5.2
Receive Registers			
0x00005080	FCTRL	Filter Control Register	Section 8.2.2.6.1
0x00005088	VLNCTRL	VLAN Control Register	Section 8.2.2.6.2



Table 8.3. BARO Registers Summary (Continued)

Offset / Alias Offset	Abbreviation	Name	Section
0x00005090	MCSTCTRL	Multicast Control Register	Section 8.2.2.6.3
0x0000EA00 + 0x4*n, n=0...63	PSRTYPE[n]	Packet Split Receive Type Register	Section 8.2.2.6.4
0x00005000	RXCSUM	Receive Checksum Control	Section 8.2.2.6.5
0x00005008	RFCTL	Receive Filter Control Register	Section 8.2.2.6.6
0x00005200 + 0x4*n, n=0...127	MTA[n]	Multicast Table Array	Section 8.2.2.6.7
0x0000A200 + 0x8*n, n=0...127	RAL[n]	Receive Address Low	Section 8.2.2.6.8
0x0000A204 + 0x8*n, n=0...127	RAH[n]	Receive Address High	Section 8.2.2.6.9
0x00005400 + 0x8*n, n=0...15	RAL_ALIAS[n]	Receive Address Low	Section 8.2.2.6.10
0x00005404 + 0x8*n, n=0...15	RAH_ALIAS[n]	Receive Address High	Section 8.2.2.6.11
0x0000A600 + 0x4*n, n=0...255	MPSAR[n]	MAC Pool Select Array	Section 8.2.2.6.12
0x0000A000 + 0x4*n, n=0...127	VFTA[n]	VLAN Filter Table Array	Section 8.2.2.6.13
0x0000EC80	MRQC	Multiple Receive Queues Command Register	Section 8.2.2.6.14
0x00005818	MRQC_ALIAS	Multiple Receive Queues Command Register	Section 8.2.2.6.15
0x0000EC70	RQTC	RSS Queues per Traffic Class Register	Section 8.2.2.6.16
0x0000EB80 + 0x4*n, n=0...9	RSSRK[n]	RSS Random Key Register	Section 8.2.2.6.17
0x0000EB00 + 0x4*n, n=0...31	RETA[n]	Redirection Table	Section 8.2.2.6.18
0x0000EE80 + 0x4*n, n=0...95	ERETA[n]	Extended Redirection Table	Section 8.2.2.6.19
0x00018000 + 0x4*n + 0x40*m, n=0...15, m=0...63	VFRSSRK[n,m]	Per Pool RSS Random Key Register	Section 8.2.2.6.20
0x00019000 + 0x4*n + 0x40*m, n=0...15, m=0...63	VFRETA[n,m]	Per Pool Redirection Table	Section 8.2.2.6.21
0x0000EC30	SYNQF	SYN Packet Queue Filter	Section 8.2.2.6.22
0x00005128 + 0x4*n, n=0...7	ETQF[n]	EType Queue Filter	Section 8.2.2.6.23
0x0000EC00 + 0x4*n, n=0...7	ETQS[n]	EType Queue Select	Section 8.2.2.6.24
0x00005084	ETAG_ETYPE	E-tag Ethertype Register	Section 8.2.2.6.25
0x0000507C	VXLANCTRL	VXLAN Control	Section 8.2.2.6.26
Receive DMA Registers			
0x00001000 + 0x40*n, n=0...63 and 0x0000D000 + 0x40*(n-64), n=64...127	RDBAL[n]	Receive Descriptor Base Address Low	Section 8.2.2.6.27
0x00001004 + 0x40*n, n=0...63 and 0x0000D004 + 0x40*(n-64), n=64...127	RDBAH[n]	Receive Descriptor Base Address High	Section 8.2.2.6.28
0x00001008 + 0x40*n, n=0...63 and 0x0000D008 + 0x40*(n-64), n=64...127	RDLEN[n]	Receive Descriptor Length	Section 8.2.2.6.29
0x00001010 + 0x40*n, n=0...63 and 0x0000D010 + 0x40*(n-64), n=64...127	RDH[n]	Receive Descriptor Head	Section 8.2.2.6.30



Table 8.3. BARO Registers Summary (Continued)

Offset / Alias Offset	Abbreviation	Name	Section
0x00001018 + 0x40*n, n=0...63 and 0x0000D018 + 0x40*(n-64), n=64...127	RDT[n]	Receive Descriptor Tail	Section 8.2.2.6.31
0x00001028 + 0x40*n, n=0...63 and 0x0000D028 + 0x40*(n-64), n=64...127	RXDCTL[n]	Receive Descriptor Control	Section 8.2.2.6.32
0x00002100 + 0x4*n, n=0...15	SRRCTL_ALIAS[n]	Split Receive Control Registers	Section 8.2.2.6.33
0x00001014 + 0x40*n, n=0...63 and 0x0000D014 + 0x40*(n-64), n=64...127	SRRCTL[n]	Split Receive Control Registers	Section 8.2.2.6.34
0x00002F00	RDRXCTL	Receive DMA Control Register	Section 8.2.2.6.35
0x00003C00 + 0x4*n, n=0...7	RXPBSIZE[n]	Receive Packet Buffer Size	Section 8.2.2.6.36
0x00003000	RXCTRL	Receive Control Register	Section 8.2.2.6.37
0x0000102C + 0x40*n, n=0...63 and 0x0000D02C + 0x40*(n-64), n=64...127	RSCCTL[n]	RSC Control	Section 8.2.2.6.38
Transmit Registers			
0x00008100	DTXMXSZRQ	DMA Tx TCP Max Allow Size Requests	Section 8.2.2.7.1
0x00004A80	DMATXCTL	DMA Tx Control	Section 8.2.2.7.2
0x00004A88	DTXTCPFLGL	DMA Tx TCP Flags Control Low	Section 8.2.2.7.3
0x00004A8C	DTXTCPFLGH	DMA Tx TCP Flags Control High	Section 8.2.2.7.4
0x00006000 + 0x40*n, n=0...127	TDBAL[n]	Transmit Descriptor Base Address Low	Section 8.2.2.7.5
0x00006004 + 0x40*n, n=0...127	TDBAH[n]	Transmit Descriptor Base Address High	Section 8.2.2.7.6
0x00006008 + 0x40*n, n=0...127	TDLEN[n]	Transmit Descriptor Length	Section 8.2.2.7.7
0x00006010 + 0x40*n, n=0...127	TDH[n]	Transmit Descriptor Head	Section 8.2.2.7.8
0x00006018 + 0x40*n, n=0...127	TDT[n]	Transmit Descriptor Tail	Section 8.2.2.7.9
0x00006028 + 0x40*n, n=0...127	TXDCTL[n]	Transmit Descriptor Control	Section 8.2.2.7.10
0x00006038 + 0x40*n, n=0...127	TDWBAL[n]	Tx Descriptor Completion Write Back Address Low	Section 8.2.2.7.11
0x0000603C + 0x40*n, n=0...127	TDWBAH[n]	Tx Descriptor Completion Write Back Address High	Section 8.2.2.7.12
0x0000CC00 + 0x4*n, n=0...7	TXPBSIZE[n]	Transmit Packet Buffer Size	Section 8.2.2.7.13
0x0000CD10	MNGTXMAP	Manageability Transmit TC Mapping	Section 8.2.2.7.14
0x00008120	MTQC	Multiple Transmit Queues Command Register	Section 8.2.2.7.15
0x00004950 + 0x4*n, n=0...7	TXPBTHRESH[n]	Tx Packet Buffer Threshold	Section 8.2.2.7.16
0x00017100	TAG_ETYPE	Tags Ethertypes	Section 8.2.2.7.17
TPH Registers			
0x00002200 + 0x4*n, n=0...15	TPH_RXCTRL_ALIAS[n]	Rx TPH Control Register	Section 8.2.2.8.1



Table 8.3. BARO Registers Summary (Continued)

Offset / Alias Offset	Abbreviation	Name	Section
0x0000100C + 0x40*n, n=0...63 and 0x0000D00C + 0x40*(n-64), n=64...127	TPH_RXCTRL[n]	Rx TPH Control Register	Section 8.2.2.8.2
0x0000600C + 0x40*n, n=0...127	TPH_TXCTRL[n]	Tx TPH Control Registers	Section 8.2.2.8.3
Timers Registers			
0x0000004C	TCPTIMER	TCP Timer	Section 8.2.2.9.1
Flow Director Registers			
0x0000EE00	FDIRCTRL	Flow Director Filters Control Register	Section 8.2.2.10.1
0x0000EE68	FDIRHKEY	Flow Director Filters Lookup Table HASH Key	Section 8.2.2.10.2
0x0000EE6C	FDIRSKEY	Flow Director Filters Signature HASH Key	Section 8.2.2.10.3
0x0000EE3C	FDIRDIP4M	Flow Director Filters DIPv4 Mask	Section 8.2.2.10.4
0x0000EE40	FDIRSIP4M	Flow Director Filters Source IPv4 Mask	Section 8.2.2.10.5
0x0000EE44	FDIRTCPM	Flow Director Filters TCP Mask	Section 8.2.2.10.6
0x0000EE48	FDIRUDPM	Flow Director Filters UDP Mask	Section 8.2.2.10.7
0x0000EE78	FDIRSCTPM	Flow Director Filters SCTP Mask	Section 8.2.2.10.8
0x0000EE74	FDIRIP6M	Flow Director Filters IPv6 Mask	Section 8.2.2.10.9
0x0000EE70	FDIRM	Flow Director Filters Other Mask	Section 8.2.2.10.10
0x0000EE38	FDIRFREE	Flow Director Filters Free	Section 8.2.2.10.11
0x0000EE4C	FDIRLEN	Flow Director Filters Length	Section 8.2.2.10.12
0x0000EE50	FDIRUSTAT	Flow Director Filters Usage Statistics	Section 8.2.2.10.13
0x0000EE54	FDIRFSTAT	Flow Director Filters Failed Usage Statistics	Section 8.2.2.10.14
0x0000EE58	FDIRMATCH	Flow Director Filters Match Statistics	Section 8.2.2.10.15
0x0000EE5C	FDIRMISS	Flow Director Filters Miss Match Statistics	Section 8.2.2.10.16
0x0000EE0C + 0x4*n, n=0...2	FDIRSIPV6[n]	Flow Director Filters Source IPv6	Section 8.2.2.10.17
0x0000EE18	FDIRIPSA	Flow Director Filters IP SA	Section 8.2.2.10.18
0x0000EE1C	FDIRIPDA	Flow Director Filters IP DA	Section 8.2.2.10.19
0x0000EE20	FDIRPORT	Flow Director Filters Port	Section 8.2.2.10.20
0x0000EE24	FDIRVLAN	Flow Director Filters VLAN and FLEX bytes	Section 8.2.2.10.21
0x0000EE28	FDIRHASH	Flow Director Filters Hash Signature	Section 8.2.2.10.22
0x0000EE2C	FDIRCMD	Flow Director Filters Command register	Section 8.2.2.10.23
MAC Registers			
0x00011148	PHY_INDIRECT_DATA	PHY Indirect Data	Section 8.2.2.11.1
0x00011144	PHY_INDIRECT_CTRL	PHY Indirect Control	Section 8.2.2.11.2
0x00004240	HLREG0	Core Control 0 Register	Section 8.2.2.11.3
0x00004244	HLREG1	Core Status 1 Register	Section 8.2.2.11.4
0x00004248	PAP	Pause and Pace Register	Section 8.2.2.11.5
0x0000425C	MSCA	MDI Single Command and Address <i>Note:</i> Values in this register should not be changed.	



Table 8.3. BARO Registers Summary (Continued)

Offset / Alias Offset	Abbreviation	Name	Section
0x00004260	MSRWD	MDI Single Read and Write Data <i>Note:</i> Values in this register should not be changed.	
0x00004268	MAXFRS	Max Frame Size	Section 8.2.2.11.6
0x000042A4	LINKS	Link Status Register	Section 8.2.2.11.7
0x000042D0	MMNGC	MAC Manageability Control Register	Section 8.2.2.11.8
0x00004330	MACC	MAC Control register	Section 8.2.2.11.9
Statistic Registers			
0x00004018	MBSDC	Bad SFD Count	Section 8.2.2.12.1
0x00004000	CRCERRS	CRC Error Count	Section 8.2.2.12.2
0x00004004	ILLERRC	Illegal Byte Error Count	Section 8.2.2.12.3
0x00004008	ERRBC	Error Byte Count	Section 8.2.2.12.4
0x00004034	MLFC	MAC Local Fault Count	Section 8.2.2.12.5
0x00004038	MRFC	MAC Remote Fault Count	Section 8.2.2.12.6
0x00004040	RLEC	Receive Length Error Count	Section 8.2.2.12.7
0x00008780	SSVPC	Switch Security Violation Packet Count	Section 8.2.2.12.8
0x000041A4	LXONRXCNT	Link XON Received Count	Section 8.2.2.12.9
0x000041A8	LXOFFRXCNT	Link XOFF Received Count	Section 8.2.2.12.10
0x00004140 + 0x4*n, n=0...7	PXONRXCNT[n]	Priority XON Received Count	Section 8.2.2.12.11
0x00004160 + 0x4*n, n=0...7	PXOFFRXCNT[n]	Priority XOFF Received Count	Section 8.2.2.12.12
0x0000405C	PRC64	Packets Received [64 Bytes] Count	Section 8.2.2.12.13
0x00004060	PRC127	Packets Received [65-127 Bytes] Count	Section 8.2.2.12.14
0x00004064	PRC255	Packets Received [128-255 Bytes] Count	Section 8.2.2.12.15
0x00004068	PRC511	Packets Received [256-511 Bytes] Count	Section 8.2.2.12.16
0x0000406C	PRC1023	Packets Received [512-1023 Bytes] Count	Section 8.2.2.12.17
0x00004070	PRC1522	Packets Received [1024 to Max Bytes] Count	Section 8.2.2.12.18
0x00004078	BPRC	Broadcast Packets Received Count	Section 8.2.2.12.19
0x0000407C	MPRC	Multicast Packets Received Count	Section 8.2.2.12.20
0x00004074	GPRC	Good Packets Received Count	Section 8.2.2.12.21
0x00004088	GORCL	Good Octets Received Count Low	Section 8.2.2.12.22
0x0000408C	GORCH	Good Octets Received Count High	Section 8.2.2.12.23
0x000041B0	RXNFGPC	Good Rx Non-Filtered Packet Counter	Section 8.2.2.12.24
0x000041B4	RXNFBCL	Good Rx Non-Filter Byte Counter Low	Section 8.2.2.12.25
0x000041B8	RXNFBCH	Good Rx Non-Filter Byte Counter High	Section 8.2.2.12.26
0x00002F50	RXDGPC	DMA Good Rx Packet Counter	Section 8.2.2.12.27
0x00002F54	RXDGBCL	DMA Good Rx Byte Counter Low	Section 8.2.2.12.28
0x00002F58	RXDGBCH	DMA Good Rx Byte Counter High	Section 8.2.2.12.29
0x00002F5C	RXDDPC	DMA Duplicated Good Rx Packet Counter	Section 8.2.2.12.30
0x00002F60	RXDBCL	DMA Duplicated Good Rx Byte Counter Low	Section 8.2.2.12.31
0x00002F64	RXDBCH	DMA Duplicated Good Rx Byte Counter High	Section 8.2.2.12.32
0x00002F68	RXLPBKPC	DMA Good Rx LPBK Packet Counter	Section 8.2.2.12.33



Table 8.3. BARO Registers Summary (Continued)

Offset / Alias Offset	Abbreviation	Name	Section
0x00002F6C	RXLPKBCL	DMA Good Rx LPBK Byte Counter Low	Section 8.2.2.12.34
0x00002F70	RXLPKBCH	DMA Good Rx LPBK Byte Counter High	Section 8.2.2.12.35
0x00002F74	RXDLPBKPC	DMA Duplicated Good Rx LPBK Packet Counter	Section 8.2.2.12.36
0x00002F78	RXDLPKBCL	DMA Duplicated Good Rx LPBK Byte Counter Low	Section 8.2.2.12.37
0x00002F7C	RXDLPKBCH	DMA Duplicated Good Rx LPBK Byte Counter High	Section 8.2.2.12.38
0x00004080	GPTC	Good Packets Transmitted Count	Section 8.2.2.12.39
0x00004090	GOTCL	Good Octets Transmitted Count Low	Section 8.2.2.12.40
0x00004094	GOTCH	Good Octets Transmitted Count High	Section 8.2.2.12.41
0x000087A0	TXDGPC	DMA Good Tx Packet Counter	Section 8.2.2.12.42
0x000087A4	TXDGBCL	DMA Good Tx Byte Counter Low	Section 8.2.2.12.43
0x000087A8	TXDGBCH	DMA Good Tx Byte Counter High	Section 8.2.2.12.44
0x000040A4	RUC	Receive Undersize Count	Section 8.2.2.12.45
0x000040A8	RFC	Receive Fragment Count	Section 8.2.2.12.46
0x000040AC	ROC	Receive Oversize Count	Section 8.2.2.12.47
0x000040B0	RJC	Receive Jabber Count	Section 8.2.2.12.48
0x000040B4	MNGPRC	Management Packets Received Count	Section 8.2.2.12.49
0x000040B8	MNGPDC	Management Packets Dropped Count	Section 8.2.2.12.50
0x000040C0	TORL	Total Octets Received Low	Section 8.2.2.12.51
0x000040C4	TORH	Total Octets Received High	Section 8.2.2.12.52
0x000040D0	TPR	Total Packets Received	Section 8.2.2.12.53
0x000040D4	TPT	Total Packets Transmitted	Section 8.2.2.12.54
0x000040D8	PTC64	Packets Transmitted (64 Bytes) Count	Section 8.2.2.12.55
0x000040DC	PTC127	Packets Transmitted [65-127 Bytes] Count	Section 8.2.2.12.56
0x000040E0	PTC255	Packets Transmitted [128-255 Bytes] Count	Section 8.2.2.12.57
0x000040E4	PTC511	Packets Transmitted [256-511 Bytes] Count	Section 8.2.2.12.58
0x000040E8	PTC1023	Packets Transmitted [512-1023 Bytes] Count	Section 8.2.2.12.59
0x000040EC	PTC1522	Packets Transmitted [Greater than 1024 Bytes] Count	Section 8.2.2.12.60
0x000040F0	MPTC	Multicast Packets Transmitted Count	Section 8.2.2.12.61
0x000040F4	BPTC	Broadcast Packets Transmitted Count	Section 8.2.2.12.62
0x00004010	MSPDC	MAC Short Packet Discard Count	Section 8.2.2.12.63
0x00004120	XEC	XSUM Error Count	Section 8.2.2.12.64
0x00002300 + 0x4*n, n=0...31	RQSMR[n]	Receive Queue Statistic Mapping Registers	Section 8.2.2.12.65
0x00002F40	RXDSTATCTRL	Rx DMA Statistic Counter Control	Section 8.2.2.12.66
0x00008600 + 0x4*n, n=0...31	TQSM[n]	Transmit Queue Statistic Mapping Registers	Section 8.2.2.12.67
0x00001030 + 0x40*n, n=0...15	QPRC[n]	Queue Packets Received Count	Section 8.2.2.12.68
0x00001430 + 0x40*n, n=0...15	QPRDC[n]	Queue Packets Received Drop Count	Section 8.2.2.12.69
0x00001034 + 0x40*n, n=0...15	QBRC_L[n]	Queue Bytes Received Count Low	Section 8.2.2.12.70



Table 8.3. BARO Registers Summary (Continued)

Offset / Alias Offset	Abbreviation	Name	Section
0x00001038 + 0x40*n, n=0...15	QBRC_H[n]	Queue Bytes Received Count High	Section 8.2.2.12.71
0x00006030 + 0x40*n, n=0...15	QPTC_ALIASE[n]	Queue Packets Transmitted Count	Section 8.2.2.12.72
0x00008680 + 0x4*n, n=0...15	QPTC[n]	Queue Packets Transmitted Count	Section 8.2.2.12.73
0x00008700 + 0x8*n, n=0...15	QBTC_L[n]	Queue Bytes Transmitted Count Low	Section 8.2.2.12.74
0x00008704 + 0x8*n, n=0...15	QBTC_H[n]	Queue Bytes Transmitted Count High	Section 8.2.2.12.75
0x00005118	FCCRC	Fiber Channel CRC Error Count	Section 8.2.2.12.76
0x000041C0	B2OSPC	BMC2OS Packets Sent By BMC	Section 8.2.2.12.77
0x00002F90	B2OGPRC	BMC2OS Packets Received By Host	Section 8.2.2.12.78
0x000041C4	O2BGPTC	OS2BMC Packets Received By BMC	Section 8.2.2.12.79
0x000087B0	O2BSPC	OS2BMC Packets Transmitted By Host	Section 8.2.2.12.80
0x00004180	BUPRC	Total Unicast Packets Received (BMC Copy)	Section 8.2.2.12.81
0x00004184	BMPRC	BMC Total Multicast Packets Received	Section 8.2.2.12.82
0x00004188	BBPRC	Total Broadcast Packets Received (BMC Copy)	Section 8.2.2.12.83
0x0000418C	BUPTC	Total Unicast Packets Transmitted (BMC Copy)	Section 8.2.2.12.84
0x00004190	BMPTC	BMC Total Multicast Packets Transmitted	Section 8.2.2.12.85
0x00004194	BBPTC	Total Broadcast Packets Transmitted (BMC Copy)	Section 8.2.2.12.86
0x00004198	BCRCERRS	BMC FCS Receive Errors	Section 8.2.2.12.87
0x0000419C	BXONRXC	BMC Pause XON Frames Received	Section 8.2.2.12.88
0x000041E0	BXOFFRXC	BMC Pause XOFF Frames Received	Section 8.2.2.12.89
0x000041E4	BXONTXC	BMC Pause XON Frames Transmitted	Section 8.2.2.12.90
0x000041E8	BXOFFTXC	BMC Pause XOFF Frames Transmitted	Section 8.2.2.12.91
0x000041F0	B2OSDPC	Sideband Receive Dropped Packet Count	Section 8.2.2.12.92
Wake-Up and Proxy Control Registers			
0x00005800	WUC	Wake Up Control Register	Section 8.2.2.13.1
0x00005808	WUFC	Wake Up Filter Control Register	Section 8.2.2.13.2
0x00005810	WUS	Wake Up Status Register	Section 8.2.2.13.3
0x00005838	IPAV	IP Address Valid	Section 8.2.2.13.4
0x00005840 + 0x8*n, n=0...3	IP4AT[n]	IPv4 Address Table	Section 8.2.2.13.5
0x00005880 + 0x4*n, n=0...3	IP6AT[n]	IPv6 Address Table	Section 8.2.2.13.6
0x00005990 + 0x4*n, n=0...11	IP6AT_EXT[n]	IPv6 Address Table Extended	Section 8.2.2.13.7
0x00005A00 + 0x4*n, n=0...31	WUPM[n]	Wake Up Packet Memory (128 Bytes)	Section 8.2.2.13.8
0x00009000 + 0x10*n + 0x100*m, n=0...15, m=0...3 and 0x00009600 + 0x10*(n- 16) + 0x100*m, n=16...31, m=0...3	FHFT_FILTER_DW_ EVEN[n,m]	Filter DW Even	Section 8.2.2.13.9




Table 8.3. BARO Registers Summary (Continued)

Offset / Alias Offset	Abbreviation	Name	Section
0x00009004 + 0x10*n + 0x100*m, n=0...15, m=0...3 and 0x00009604 + 0x10*(n-16) + 0x100*m, n=16...31, m=0...3	FHFT_FILTER_DW_ODD[n,m]	Filter DW Odd	Section 8.2.2.13.10
0x00009008 + 0x10*n + 0x100*m, n=0...15, m=0...3 and 0x00009608 + 0x10*(n-16) + 0x100*m, n=16...31, m=0...3	FHFT_FILTER_MASK[n,m]	Filter Mask	Section 8.2.2.13.11
0x0000900C + 0x10*n + 0x100*m, n=0...15, m=0...3 and 0x0000960C + 0x10*(n-16) + 0x100*m, n=16...31, m=0...3	FHFT_FILTER_LENGTH[n,m]	Filter Length	Section 8.2.2.13.12
0x00005F60	PROXYS	Proxying Status Register	Section 8.2.2.13.13
0x00005F64	PROXYFC	Proxying Filter Control Register	Section 8.2.2.13.14
Management Filters Registers			
0x00005864	MNGONLY	Manageability Only Traffic	Section 8.2.2.14.1
0x00005010 + 0x4*n, n=0...7	MAVTV[n]	Management VLAN TAG Value	Section 8.2.2.14.2
0x00005030 + 0x4*n, n=0...7	MFUTP[n]	Management Flex UDP/TCP Ports	Section 8.2.2.14.3
0x00005190 + 0x4*n, n=0...3	METF[n]	Management Ethernet Type Filters	Section 8.2.2.14.4
0x00005820	MANC	Management Control Register	Section 8.2.2.14.5
0x00005890 + 0x4*n, n=0...7	MDEF[n]	Manageability Decision Filters	Section 8.2.2.14.6
0x00005160 + 0x4*n, n=0...7	MDEF_EXT[n]	Manageability Decision Filters Ext	Section 8.2.2.14.7
0x00005050 + 0x4*n, n=0...3	BMCIP[n]	BMC IP Address Register	Section 8.2.2.14.8
0x00005060	BMCIPVAL	BMC IP Valid Register	Section 8.2.2.14.9
0x000058B0 + 0x4*n + 0x10*m, n=0...3, m=0...3	MIPAF[n,m]	Manageability IP Address Filter	Section 8.2.2.14.10
0x00005910 + 0x8*n, n=0...3	MMAL[n]	Manageability Ethernet MAC Address Low	Section 8.2.2.14.11
0x00005914 + 0x8*n, n=0...3	MMAH[n]	Manageability Ethernet MAC Address High	Section 8.2.2.14.12
0x00009400 + 0x10*n, n=0...15	FTFT_FILTER_EVEN[n]	FTFT Filter DW Even Words	Section 8.2.2.14.13
0x00009404 + 0x10*n, n=0...15	FTFT_FILTER_ODD[n]	FTFT Filter DW Odd Words	Section 8.2.2.14.14
0x00009408 + 0x10*n, n=0...15	FTFT_FILTER0_MASK[n]	FTFT Filter Mask	Section 8.2.2.14.15
0x0000940C + 0x10*n, n=0...15	FTFT_FILTER0_LENGTH[n]	FTFT Filter Length	Section 8.2.2.14.16
Manageability (ARC subsystem) HOST Interface			
0x00010140	SWSM	Software Semaphore Register	Section 8.2.2.15.1
0x00010148	FWSM	Firmware Semaphore Register	Section 8.2.2.15.2
0x00010160	SW_FW_SYNC	Software-Firmware Synchronization	Section 8.2.2.15.3
0x00010164	SW_FW_SYNC_MIRR	Software-Firmware Synchronization	Section 8.2.2.15.4
0x00015F40	FWRESETCNT	Firmware Resets Count	Section 8.2.2.15.5



Table 8.3. BARO Registers Summary (Continued)

Offset / Alias Offset	Abbreviation	Name	Section
0x00015F00	HICR	HOST Interface Control Register	Section 8.2.2.15.6
0x00015800 + 0x4*n, n=0...447	ARCRAM[n]	Host ARC Data RAM	Section 8.2.2.15.7
Time Sync (IEEE 1588) Registers			
0x00008C0C	SYSTIMEL	System Time Register Low	Section 8.2.2.15.8
0x00008C10	SYSTIMEH	System Time Register High	Section 8.2.2.15.9
0x00008C58	SYSTIMR	System Time Register Residue	Section 8.2.2.15.10
0x00008C14	TIMEINCA	Increment Attributes Register	Section 8.2.2.15.11
0x00008C18	TIMADJ	Time Adjustment Offset Register	Section 8.2.2.15.12
0x00008C68	TSIM	Time Sync Interrupt Mask Register	Section 8.2.2.15.13
0x00008C60	TSICR	Time Sync Interrupt Cause Register	Section 8.2.2.15.14
0x00008C00	TSYNCTXCTL	Tx Time Sync Control Register	Section 8.2.2.15.15
0x00008C04	TXSTMPL	Tx Timestamp Value Low	Section 8.2.2.15.16
0x00008C08	TXSTMPH	Tx Timestamp Value High	Section 8.2.2.15.17
0x00005188	TSYNCRXCTL	Rx Time Sync Control Register	Section 8.2.2.15.18
0x000051E8	RXSTMPL	Rx Timestamp Low	Section 8.2.2.15.19
0x000051A4	RXSTMPH	Rx Timestamp High	Section 8.2.2.15.20
0x00005120	RXMTRL	Rx Message Type Register Low	Section 8.2.2.15.21
0x0000003C	TSSDP	Time Sync SDP Configuration Register	Section 8.2.2.15.22
0x00008C20	TSAUXC	TimeSync Auxiliary Control Register	Section 8.2.2.15.23
0x00008C24	TRGTTIMELO	Target Time Register 0 Low	Section 8.2.2.15.24
0x00008C28	TRGTTIMEHO	Target Time Register 0 High	Section 8.2.2.15.25
0x00008C2C	TRGTTIMEL1	Target Time Register 1 Low	Section 8.2.2.15.26
0x00008C30	TRGTTIMEH1	Target Time Register 1 High	Section 8.2.2.15.27
0x00008C34	FREQOUT0	Frequency Out 0 Control Register	Section 8.2.2.15.28
0x00008C38	FREQOUT1	Frequency Out 1 Control Register	Section 8.2.2.15.29
0x00008C3C	AUXSTMPL0	Auxiliary Time Stamp 0 Register Low	Section 8.2.2.15.30
0x00008C40	AUXSTMPH0	Auxiliary Time Stamp 0 Register High	Section 8.2.2.15.31
0x00008C44	AUXSTMPL1	Auxiliary Time Stamp 1 Register Low	Section 8.2.2.15.32
0x00008C48	AUXSTMPH1	Auxiliary Time Stamp 1 Register High	Section 8.2.2.15.33
Virtualization PF Registers			
0x000050B4	PFFLPH	Filter Local Packets High	Section 8.2.2.16.1
0x000050B0	PFFLPL	Filter Local Packets Low	Section 8.2.2.16.2
0x000051B0	PVVTCTL	PF Virtual Control Register	Section 8.2.2.16.3
0x00004B00 + 0x4*n, n=0...63	PFMAILBOX[n]	PF Mailbox	Section 8.2.2.16.4
0x00000710 + 0x4*n, n=0...3	PFBICR[n]	PF Mailbox Interrupt Causes Register	Section 8.2.2.16.5
0x00000720 + 0x4*n, n=0...1	PFBIMR[n]	PF Mailbox Interrupt Mask Register	Section 8.2.2.16.6
0x00000700 + 0x4*n, n=0...1	PVFLREC[n]	PF VFLR Events Clear	Section 8.2.2.16.7
0x000051E0 + 0x4*n, n=0...1	PVVFRE[n]	PF VF Receive Enable	Section 8.2.2.16.8



Table 8.3. BARO Registers Summary (Continued)

Offset / Alias Offset	Abbreviation	Name	Section
0x00008110 + 0x4*n, n=0...1	PFVFTE[n]	PF VF Transmit Enable	Section 8.2.2.16.9
0x00008790	PFVMECM0	PF VM 0:31 Error Count Mask	Section 8.2.2.16.10
0x00008794	PFVMECM1	PF VM 32:63 Error Count Mask	Section 8.2.2.16.11
0x00002F04	PFQDE	PF Queue Drop Enable Register	Section 8.2.2.16.12
0x00005180 + 0x4*n, n=0...1	PFVMTXSW[n]	PF VM Tx Switch Loopback Enable	Section 8.2.2.16.13
0x00008200 + 0x4*n, n=0...7	PFVFSPOOF[n]	PFVF Anti Spoof Control	Section 8.2.2.16.14
0x00008220	PFDTXGSWC	PF DMA Tx General Switch Control	Section 8.2.2.16.15
0x00008000 + 0x4*n, n=0...63	PFVMVIR[n]	PF VM VLAN Insert Register	Section 8.2.2.16.16
0x00017000 + 0x4*n, n=0...63	PFVMTIR[n]	PF VM Tag Insert Register	Section 8.2.2.16.17
0x0000F000 + 0x4*n, n=0...63	PFVML2FLT[n]	PF VM L2Control Register	Section 8.2.2.16.18
0x0000F100 + 0x4*n, n=0...63	PFVLVF[n]	PF VM VLAN Pool Filter	Section 8.2.2.16.19
0x0000F200 + 0x4*n, n=0...127	PFVLVFB[n]	PF VM VLAN Pool Filter Bitmap	Section 8.2.2.16.20
0x0000F400 + 0x4*n, n=0...127	PFUTA[n]	PF Unicast Table Array	Section 8.2.2.16.21
0x0000F600 + 0x4*n, n=0...3	PFMRCTL[n]	PF Mirror Rule Control	Section 8.2.2.16.22
0x0000F610 + 0x4*n, n=0...7	PFMRVLAN[n]	PF Mirror Rule VLAN	Section 8.2.2.16.23
0x0000F630 + 0x4*n, n=0...7	PFMRVM[n]	PF Mirror Rule Pool	Section 8.2.2.16.24
0x00002FA8	LVMMC_RX	Last VM Misbehavior Cause - Rx	Section 8.2.2.16.25
0x00008108	LVMMC_TX	Last VM Misbehavior Cause - Tx	Section 8.2.2.16.26
0x00002FA4	LMVM_RX	Last Malicious VM - Rx	Section 8.2.2.16.27
0x00008124	LMVM_TX	Last Malicious VM - Tx	Section 8.2.2.16.28
0x00002FB0 + 0x4*n, n=0...3	WQBR_RX[n]	Wrong Queue Behavior Register - Rx	Section 8.2.2.16.29
0x00008130 + 0x4*n, n=0...3	WQBR_TX[n]	Wrong Queue Behavior Register - Tx	Section 8.2.2.16.30
Power Management			
0x00015F20	DMCMNGTH	DMA Coalescing Management Threshold	Section 8.2.2.17.1
0x00002400	DMACR	DMA Coalescing Control Register	Section 8.2.2.17.2
0x00003300 + 0x4*n, n=0...7	DMCTH[n]	DMA Coalescing Threshold	Section 8.2.2.17.3
0x00002404	DMCTLX	DMA Coalescing Time to Lx Request	Section 8.2.2.17.4
0x00011708	LTRC	Latency Tolerance Reporting (LTR) Control	Section 8.2.2.17.5
0x000043A0	EEER	Energy Efficient Ethernet (EEE) Register	Section 8.2.2.17.6
0x00004380	EEE_SU	Energy Efficient Ethernet (EEE) Setup Register	Section 8.2.2.17.7
0x00004398	EEE_STAT	Energy Efficient Ethernet (EEE) STATUS	Section 8.2.2.17.8
0x000041F4	TLPIC	EEE TX LPI Count	Section 8.2.2.17.9
0x000041F8	RLPIC	EEE RX LPI Count	Section 8.2.2.17.10
Security Registers			
0x00008800	SECTXCTRL	Security Tx Control	Section 8.2.2.18.1
0x00008804	SECTXSTAT	Security Tx Status	Section 8.2.2.18.2



Table 8.3. BARO Registers Summary (Continued)

Offset / Alias Offset	Abbreviation	Name	Section
0x00008808	SECTXBUFFAF	Security Tx Buffer Almost Full	Section 8.2.2.18.3
0x00008810	SECTXMINIFG	Security Tx Buffer Minimum IFG	Section 8.2.2.18.4
0x00008D00	SECRXCTRL	Security Rx Control	Section 8.2.2.18.5
0x00008D04	SECRXSTAT	Security Rx Status	Section 8.2.2.18.6
IPSec Registers			
0x00008900	IPSTXIDX	IPsec Tx Index	Section 8.2.2.20.1
0x00008908 + 0x4*n, n=0...3	IPSTXKEY[n]	IPsec Tx Key Registers	Section 8.2.2.20.2
0x00008904	IPSTXSALT	IPsec Tx Salt Register	Section 8.2.2.20.3
0x00008E00	IPSRXIDX	IPsec Rx Index	Section 8.2.2.20.4
0x00008E04 + 0x4*n, n=0...3	IPSRXIPADDR[n]	IPsec Rx IP Address Register	Section 8.2.2.20.5
0x00008E14	IPSRXSPI	IPsec Rx SPI Register	Section 8.2.2.20.6
0x00008E18	IPSRXIPIDX	IPsec Rx SPI Register IP Index	Section 8.2.2.20.7
0x00008E1C + 0x4*n, n=0...3	IPSRXKEY[n]	IPsec Rx Key Register	Section 8.2.2.20.8
0x00008E2C	IPSRXSALT	IPsec Rx Salt Register	Section 8.2.2.20.9
0x00008E30	IPSRXMOD	IPsec Rx Mode Register	Section 8.2.2.20.10
VF Registers Mapping in the PF Space			
0x00000300 + 0x4*n, n=0...63	VFCTRL[n]	VF Control Register	Section 8.2.2.21.1
0x000004C00 + 0x4*n, n=0...63	VFMAILBOX[n]	VF Mailbox	Section 8.2.2.21.2
0x00013000 + 0x4*n + 0x40*m, n=0...15, m=0...63	PFMBMEM[n,m]	PF Mailbox Memory	Section 8.2.2.21.3
0x00000B00 + 0x4*n, n=0...63	VFEICR[n]	VF Extended Interrupt Cause	Section 8.2.2.21.4
0x00000C00 + 0x4*n, n=0...63	VFEICS[n]	VF Extended Interrupt Cause Set	Section 8.2.2.21.5
0x00000D00 + 0x4*n, n=0...63	VFEIMS[n]	VF Extended Interrupt Mask Set/Read	Section 8.2.2.21.6
0x00000E00 + 0x4*n, n=0...63	VFEIMC[n]	VF Extended Interrupt Mask Clear	Section 8.2.2.21.7
0x00004D00 + 0x4*n, n=0...63	VFEIAM[n]	VF Extended Interrupt Auto Mask Enable	Section 8.2.2.21.8
0x00012500 + 0x4*n, n=0...63	VFIVAR[n]	VF Interrupt Vector Allocation Registers	Section 8.2.2.21.9
0x00004E00 + 0x4*n, n=0...63	VFIVAR_MISC[n]	VF Interrupt Vector Allocation Registers Misc	Section 8.2.2.21.10
0x00003400 + 0x4*n, n=0...63	VFMRQC[n]	VF Multiple Receive Queues Command Register	Section 8.2.2.21.11
0x00008300 + 0x4*n, n=0...63	VFGPTC[n]	VF Good Packets Transmitted Count	Section 8.2.2.21.12
0x00008400 + 0x8*n, n=0...63	VFGOTC_LSB[n]	VF Good Octets Transmitted Count LSB	Section 8.2.2.21.13
0x00008404 + 0x8*n, n=0...63	VFGOTC_MSB[n]	VF Good Octets Transmitted Count MSB	Section 8.2.2.21.14
0x00001020 + 0x40*n, n=0...63	VFGORC_LSB[n]	VF Good Octets Received Count Low	Section 8.2.2.21.15



Table 8.3. BAR0 Registers Summary (Continued)

Offset / Alias Offset	Abbreviation	Name	Section
0x0000101C + 0x40*n, n=0...63	VFGPRC[n]	VF Good Packets Received Count	Section 8.2.2.21.16
0x0000D020 + 0x40*n, n=0...63	VFGORC_MSB[n]	VF Good Octets Received Count High	Section 8.2.2.21.17
0x0000D01C + 0x40*n, n=0...63	VFMPRC[n]	VF Multicast Packets Received Count	Section 8.2.2.21.18

8.2.2 Detailed Register Description - PF BAR0

8.2.2.1 General Control Registers

8.2.2.1.1 Device Control Register - CTRL (0x00000000)

CTRL is also mapped to address 0x00004 to maintain compatibility with predecessors.

Field	Bit(s)	Init.	Access Type	Description
RESERVED	1:0	0x0	RSV	Reserved.
PCIE_MASTE R_DISABLE	2	0b	RW	When set, the Intel® Xeon® Processor D-1500 Product Family LAN controller blocks new master requests, including manageability requests, by using this function. Once no master requests are pending by using this function, the PCIe Master Enable Status bit is cleared. Note: After doing any change to this bit the host must read that the bit has been modified as expected before reading STATUS.PCIe Master Enable Status bit.
LRST	3	0b	RW	Link Reset. This bit performs a reset of the MAC, PHY, and the entire Broadwell-DE LAN controller (software reset) resulting in a state nearly approximating the state following a power-up reset or integrated I/O reset, except for the system PCI configuration. Normally 0b, writing 1b initiates the reset. This bit is self-clearing. Also referred to as MAC reset.
RESERVED	25:4	0x0	RSV	Reserved.
RST	26	0b	RW	Device Reset. This bit performs a complete reset of the Intel® Xeon® Processor D-1500 Product Family LAN controller, resulting in a state nearly approximating the state following a power-up reset or integrated I/O reset, except for the system PCI configuration. Normally 0b, writing 1b initiates the reset. This bit is self-clearing. Also referred to as a software reset or global reset.
RESERVED	31:27	0x0	RSV	Reserved.

Notes: LRST and RST can be used to globally reset the entire Broadwell-DE LAN controller. This register is provided primarily as a software mechanism to recover from an indeterminate or suspected hung hardware state. Most registers (receive, transmit, interrupt, statistics, etc.) and state machines are set to their power-on reset values, approximating the state following a power-on or PCI reset. However, PCIe configuration registers are not reset, thereby leaving the device mapped into system memory space and accessible by a software device driver. To ensure that a global device reset has fully completed and that the Intel® Xeon® Processor D-1500 Product Family LAN controller responds to subsequent accesses, programmers must wait approximately 1 ms after setting before attempting to check if the bit has cleared or to access (read or write) any other device register.

8.2.2.1.2 Device Status Register - STATUS (0x00000008)



Field	Bit(s)	Init.	Access Type	Description
RESERVED	1:0	0x0	RSV	Reserved.
LAN_ID	3:2	0x0	RO	LAN ID. Provides software a mechanism to determine the Intel® Xeon® Processor D-1500 Product Family LAN controller's LAN identifier for this MAC. Read as: [0,0] LAN 0; [0,1] LAN 1.
RESERVED	6:4	0x0	RO	Reserved.
LINKUP	7	0b	RO	Linkup Status Indication. This bit is useful for IOV mode. The PF software driver sets it according to the LINKS register and PHY state. It is reflected in the VFSTATUS register indicating link up to the VF drivers.
RESERVED	9:8	0x0	RSV	Reserved.
NUM_VFS	17:10	0x0	RO	Num VFs. This field reflects the value of the Num VFs in the IOV capability structure (note that bit 17 is always 0b).
IOV_ACTIVE	18	0b	RO	IOV Active. This bit reflects the value of the VF Enable (VFE) bit in the IOV Control/Status register.
PCIE_MASTER_ENABLE_STATUS	19	1b	RO	This is a status bit of the appropriate <i>CTRL.PCIE Master Disable</i> bit. 1b = Associated LAN function can issue master requests. 0b = Associated LAN function does not issue any master request and all previously issued requests are complete.
RESERVED	31:21	0x0	RSV	Reserved. Reads as 0b.

8.2.2.1.3 Extended Device Control Register - CTRL_EXT (0x00000018)

Field	Bit(s)	Init.	Access Type	Description
RESERVED	13:0	0x00	RSV	Reserved.
PFRSTD	14	0b	SC	PF Reset Done. When set, the <i>RSTI</i> bit in all the VFMailbox registers are cleared and the <i>RSTD</i> bit in all the VFMailbox registers are set.
RESERVED	16:15	0x0	RSV	Reserved.
RO_DIS	17	0b	RW	Relaxed Ordering Disable. When set to 1b, the Intel® Xeon® Processor D-1500 Product Family LAN controller does not request any relaxed ordering transactions. When this bit is cleared and the <i>Enable Relaxed Ordering</i> bit in the Device Control register is set, the Intel® Xeon® Processor D-1500 Product Family LAN controller requests relaxed ordering transactions per queues as configured in the <i>TPH_RXCTRL[n]</i> and <i>TPH_TXCTRL[n]</i> registers.
RESERVED	25:18	0x0	RSV	Reserved.
EXTENDED_VLAN	26	0b	RW	Extended VLAN. When set, all incoming Rx packets are expected to have at least one VLAN with the Ethertype as defined in the EXVET register. This bit should only be reset by a PCIe reset and should only be changed while Tx and Rx processes are stopped. Should be set to the same value as <i>DMATXCTL.GDV</i>
RESERVED	27	0b	RSV	Reserved.
DRV_LOAD	28	0b	RW	Software device driver loaded and the corresponding network interface is enabled. This bit should be set by the software device driver after it was loaded and cleared when it unloads or at PCIe reset. The Manageability Controller (MC) loads this bit as an indication that the software device driver successfully loaded to it.
RESERVED	31:29	0x0	RSV	Reserved.



8.2.2.1.4 Extended SDP Control - ESDP (0x00000020)

This register is initialized only at Power Good Reset preserving the SDP states across software and PCIe resets. Some specific I/O pins are initialized in other resets in native mode as expected for the specific behavior and described explicitly as follows.

Field	Bit(s)	Init.	Access Type	Description
SDP0_DATA	0	0b	RW	SDP0 Data Value. Used to read (write) a value of the software-controlled I/O pin SDP0. If SDP0 is configured as an output (SDP0_IODIR = 1b), this bit controls the value driven on the pin. If SDP0 is configured as an input, all reads return the current value of the pin. See Notes 1 and 2.
SDP1_DATA	1	0b	RW	SDP1 Data Value. Used to read (write) a value of the software-controlled I/O pin SDP1. If SDP1 is configured as an output (SDP1_IODIR = 1b), this bit controls the value driven on the pin. If SDP1 is configured as an input, all reads return the current value of the pin. See Notes 1 and 3.
SDP2_DATA	2	0b	RW	SDP2 Data Value. Used to read (write) a value of software-controlled I/O pin SDP2. If SDP2 is configured as an output (SDP2_IODIR = 1b), this bit controls the value driven on the pin. If SDP2 is configured as an input, all reads return the current value of the pin. See Note 1.
SDP3_DATA	3	0b	RW	SDP3 Data Value. Used to read (write) a value of the software-controlled I/O pin SDP3. If SDP3 is configured as an output (SDP3_IODIR = 1b), this bit controls the value driven on the pin. If SDP3 is configured as an input, all reads return the current value of the pin. See Note 1.
RESERVED	7:4	0x0	RSV	Reserved.
SDP0_IODIR	8	0b	RW	SDP0 Pin Directionality. Controls whether or not software-controlled pin SDP0 is configured as an input or output. 0b = Input. 1b = Output. See Notes 1 and 2.
SDP1_IODIR	9	0b	RW	SDP1 Pin Directionality. Controls whether or not software-controlled pin SDP1 is configured as an input or output. 0b = Input. 1b = Output. See Notes 1 and 3.
SDP2_IODIR	10	0b	RW	SDP2 Pin Directionality. Controls whether or not software-controlled pin SDP2 is configured as an input or output. 0b = Input. 1b = Output. See Note 1.
SDP3_IODIR	11	0b	RW	SDP3 Pin Directionality. Controls whether or not software-controlled pin SDP3 is configured as an input or output. 0b = Input. 1b = Output. See Note 1.
RESERVED	15:12	0x0	RSV	Reserved.




Field	Bit(s)	Init.	Access Type	Description
SDP0_NATIVE	16	0b	RW	SDP0 Operating Mode. 0b = Generic software controlled I/O by SPD0_DATA and SDP0_IODIR. 1b = Native mode operation (connected to hardware function) is IEEE 1588 functionality. Refer to Section 7.7.4 . See Note 2.
SDP1_NATIVE	17	0b	RW	SDP1 Operating Mode. 0b = Generic software controlled I/O by SPD1_DATA and SDP1_IODIR. 1b = Native mode operation (connected to hardware function) is IEEE 1588 functionality according to the SDP1_Function bit. See Notes 1 and 3.
SDP2_NATIVE	18	0b	RW	SDP2 Operating Mode. 0b = Generic software controlled I/O by SPD2_DATA and SDP2_IODIR. 1b = Native mode operation (connected to hardware function). 1588 functionality or I ² C functionality according to SDP23_FUNCTION.
SDP3_NATIVE	19	0b	RW	SDP3 Operating Mode. 0b = Generic software controlled I/O by SPD3_DATA and SDP3_IODIR. 1b = Native mode operation (connected to hardware function). 1588 functionality or I ² C functionality according to SDP23_FUNCTION.
RESERVED	24:20	0x0	RSV	Reserved.
SDP1_FUNCTION	25	0b	RW	SDP1 Native Mode Functionality (SDP1_NATIVE = 1b). 0b = 1588 functionality. SDP1_IODIR should be configured as an output. 1b = Reserved. See Note 3.
SDP23_FUNCTION	26	0b	RW	Defines the usage of SDP2 & SDP3 when SDP[23]_NATIVE is set to 1: 0b = Use for 1588 functionality. 1b = Use for I ² C functionality. If this bit is set to 1b, then SDP3_NATIVE should be equal to SDP2_NATIVE.
RESERVED	30:27	0x0	RSV	Reserved.

Notes:

- Initial values are read from the NVM.
- It is assumed that bit 15 of NC-SI Configuration 1 word in the NVM is cleared; otherwise, the SDP0 pin of function 0 is used as input pins that encode the NC-SI Package ID of the Intel® Xeon® Processor D-1500 Product Family LAN controller.
- It is assumed that bit SDP_FUNC_OFF_EN of NVM Control word 2 is cleared; otherwise, SDP1 pins are strapped during PE_RST_N to determine that both PCIe functions are disabled.

8.2.2.1.5 PHY GPIO Register - PHY_GPIO (0x00000028)

Field	Bit(s)	Init.	Access Type	Description
PHY_GPIO	3:0	0x0	RO	PHY-to-MAC GPIO. Used to read the four internal PHY to MAC general purpose signals.
RESERVED	31:4	0x0	RSV	Reserved.



8.2.2.1.6 MAC GPIO Register - MAC_GPIO (0x00000030)

Field	Bit(s)	Init.	Access Type	Description
MAC_GPIO	3:0	0x0	RW	MAC-to-PHY GPIO. Used to set the four internal MAC-to-PHY general purpose signals.
RESERVED	31:4	0x0	RSV	Reserved.

8.2.2.1.7 PHY Interrupt Status Register 0 - PHYINT_STATUS0 (0x00000100)

Register contents is valid once the corresponding EEMNGCTL.CFG_DONE0/1 bit is asserted by firmware. Bits are set by firmware to notify the host of the PHY interrupts that were triggered. This register is reset by hardware only at power-up events. The host is responsible to clear the PHY interrupts once it completes the PHY interrupt handling routine.

Field	Bit(s)	Init.	Access Type	Description
PMA_RECEIVE_LINK_STATUS	0	0b	RW	Reflects the inverse state of PHY register bit 1.1.2 before it was set by a firmware read.
PMA_TRANSMIT_FAULT	1	0b	RW	Reflects the state of PHY register bit 1.8.B before it was cleared by a firmware read.
PMA_RECEIVE_FAULT	2	0b	RW	Reflects the state of PHY register bit 1.8.A before it was cleared by a firmware read.
PMA_RESERVED	7:3	0x0	RW	Reserved. Read as written.
PCS_RECEIVE_LINK_STATUS	8	0b	RW	Reflects the inverse state of PHY register bit 3.1.2 before it was set by a firmware read.
PCS_TRANSMIT_FAULT	9	0b	RW	Reflects the state of PHY register bit 3.8.B before it was cleared by a firmware read.
PCS_RECEIVE_FAULT	10	0b	RW	Reflects the state of PHY register bit 3.8.A before it was cleared by a firmware read.
PCS_10GBASE_T_BLOCK_LOCK_LATCHED	11	0b	RW	Reflects the inverse state of PHY register bit 3.21.F before it was set by a firmware read.
PCS_10GBASE_T_HIGH_BER_LATCHED	12	0b	RW	Reflects the state of PHY register bit 3.21.E before it was cleared by a firmware read.
PCS_CRC_ERROR	13	0b	RW	Reflects the state of PHY register bit 3.EC00.F before it was cleared by firmware read.
PCS_LDPC_DECODE_FAILURE	14	0b	RW	Reflects the state of PHY register bit 3.EC00.E before it was cleared by a firmware read.
PCS_INVALID_65B_BLOCK	15	0b	RW	Reflects the state of PHY register bit 3.EC00.8 before it was cleared by a firmware read.
PCS_CHANGE_IN_AUXILIARY_BIT	16	0b	RW	Reflects the state of PHY register bit 3.EC00.0 before it was cleared by a firmware read.
PCS_RESERVED	23:17	0x0	RW	Reserved. Read as written.
PHY_XS_RESERVED	31:24	0x0	RW	Reserved. Read as written.



8.2.2.1.8 PHY Interrupt Status Register 1 - PHYINT_STATUS1 (0x00000104)

Register contents is valid once the corresponding EEMNGCTL.CFG_DONE0/1 bit is asserted by firmware. Bits are set by firmware to notify the host of the PHY interrupts that were triggered. This register is reset by hardware only at power-up events. The host is responsible to clear the PHY interrupts once it completes the PHY interrupt handling routine.

Field	Bit(s)	Init.	Access Type	Description
AUTO_NEGOTIATION_EXTENDED_NEXT_PAGE_RECEIVED	0	0b	RW	Reflects the state of PHY register bit 7.1.6 before it was cleared by a firmware read.
AUTO_NEGOTIATION_REMOTE_FAULT	1	0b	RW	Reflects the state of PHY register bit 7.1.4 before it was cleared by a firmware read.
AUTO_NEGOTIATION_LINK_STATUS	2	0b	RW	Reflects the inverse state of PHY register bit 7.1.2 before it was cleared by a firmware read.
AUTO_NEGOTIATION_MASTER_SLAVE_CONFIGURATION_FAULT	3	0b	RW	Reflects the state of PHY register bit 7.21.F before it was cleared by a firmware read.
AUTO_NEGOTIATION_COMPLETED_FOR_NON_SUPPORTED_RATE	4	0b	RW	Reflects the state of PHY register bit 7.CC00.3 before it was cleared by a firmware read.
AUTO_NEGOTIATION_COMPLETED_FOR_SUPPORTED_RATE	5	0b	RW	Reflects the state of PHY register bit 7.CC00.2 before it was cleared by a firmware read.
AUTO_NEGOTIATION_AUTOMATIC_DOWNSHIFT	6	0b	RW	Reflects the state of PHY register bit 7.CC00.1 before it was cleared by a firmware read.
AUTO_NEGOTIATION_CONNECTION_STATE_CHANGE	7	0b	RW	Reflects the state of PHY register bit 7.CC00.0 before it was cleared by a firmware read.
AUTO_NEGOTIATION_10BASE_TX_DEVICE_DETECT	8	0b	RW	Reflects the state of PHY register bit 7.EC01.F before it was cleared by a firmware read.
AUTO_NEGOTIATION_ENERGY_ON_LINE_DETECT	9	0b	RW	Reflects the state of PHY register bit 7.EC01.E before it was cleared by a firmware read.
AUTO_NEGOTIATION_NEXT_PAGE_3RD_RECEIVED	10	0b	RW	Reflects the state of PHY register bit 7.EC01.3 before it was cleared by a firmware read.
AUTO_NEGOTIATION_NEXT_PAGE_2ND_RECEIVED	11	0b	RW	Reflects the state of PHY register bit 7.EC01.2 before it was cleared by a firmware read.



Field	Bit(s)	Init.	Access Type	Description
AUTO_NEGOTIATION_NEXT_PAGE_1ST_RECEIVED	12	0b	RW	Reflects the state of PHY register bit 7.EC01.1 before it was cleared by a firmware read.
AUTO_NEGOTIATION_BASE_PAGE_RECEIVED	13	0b	RW	Reflects the state of PHY register bit 7.EC01.0 before it was cleared by a firmware read.
AUTO_NEGOTIATION_10BASE_T_DEVICE_DETECT	14	0b	RW	Reflects the inverse state of PHY register bit 7.EC02.2 before it was cleared by a firmware read.
AUTO_NEGOTIATION_PROTOCOL_ERROR	15	0b	RW	Reflects the state of PHY register bit 7.EC01.D before it was cleared by a firmware read.
FLP_IDLE_ERROR	16	0b	RW	Reflects the state of PHY register bit 7.EC01.C before it was cleared by a firmware read.
AUTO_NEGOTIATION_GBE_PHY_RESERVED	27:17	0x0	RW	Reserved. Read as written.
PCIE_SERDES_LOSS_OF_SIGNAL_3_0	31:28	0x0	RW	Reflects the state of PHY register bit 4.CC02.F:C before it was cleared by a firmware read.

8.2.2.1.9 PHY Interrupt Status Register 2 - PHYINT_STATUS2 (0x00000108)

Register contents is valid once the corresponding EEMNGCTL.CFG_DONE0/1 bit is asserted by firmware. Bits are set by firmware to notify the host of the PHY interrupts that were triggered. This register is reset by hardware only at power-up events. The host is responsible to clear the PHY interrupts once it completes the PHY interrupt handling routine.

Field	Bit(s)	Init.	Access Type	Description
GLOBAL_MAC_RESET	0	0b	RW	Reflects the inverse state of PHY register bit 1E.1200.0 before it was cleared by a firmware read.
GLOBAL_MAC_LOW_POWER_LINK_UP_MODE	1	0b	RW	Reflects the state of PHY register bit 1E.1204.4 before it was cleared by a firmware read.
GLOBAL_MAC_PHY_DISABLE_MODE	2	0b	RW	Reflects the inverse state of PHY register bit 1E.1204.1 before it was cleared by a firmware read.
GLOBAL_MAC_LOW_POWER_MODE	3	0b	RW	Reflects the inverse state of PHY register bit 1E.1204.0 before it was cleared by a firmware read.
GLOBAL_MAC_SPI_GRANT	4	0b	RW	Reflects the state of PHY register bit 1E.1206.2 before it was cleared by a firmware read.
GLOBAL_MAC_SPI_CONTROL	5	0b	RW	Reflects the inverse state of PHY register bit 1E.1206.1 before it was cleared by a firmware read.
GLOBAL_HIGH_TEMPERATURE_FAILURE	6	0b	RW	Reflects the state of PHY register bit 1E.CC00.E before it was cleared by a firmware read.
GLOBAL_LOW_TEMPERATURE_FAILURE	7	0b	RW	Reflects the state of PHY register bit 1E.CC00.D before it was cleared by a firmware read.



Field	Bit(s)	Init.	Access Type	Description
GLOBAL_HIGH_TEMPERATURE_WARNING	8	0b	RW	Reflects the state of PHY register bit 1E.CC00.C before it was cleared by a firmware read.
GLOBAL_LOW_TEMPERATURE_WARNING	9	0b	RW	Reflects the state of PHY register bit 1E.CC00.B before it was cleared by a firmware read.
GLOBAL_RESET_COMPLETED	10	0b	RW	Reflects the state of PHY register bit 1E.CC00.6 before it was cleared by a firmware read.
GLOBAL_DEVICE_FAULT	11	0b	RW	Reflects the state of PHY register bit 1E.CC00.4 before it was cleared by a firmware read.
GLOBAL_PAIR_A_CHANGE_OF_STATUS	12	0b	RW	Reflects the state of PHY register bit 1E.CC00.3 before it was cleared by a firmware read.
GLOBAL_PAIR_B_CHANGE_OF_STATUS	13	0b	RW	Reflects the state of PHY register bit 1E.CC00.2 before it was cleared by a firmware read.
GLOBAL_PAIR_C_CHANGE_OF_STATUS	14	0b	RW	Reflects the state of PHY register bit 1E.CC00.1 before it was cleared by a firmware read.
GLOBAL_PAIR_D_CHANGE_OF_STATUS	15	0b	RW	Reflects the state of PHY register bit 1E.CC00.0 before it was cleared by a firmware read.
GLOBAL_MEDIUM_BER	16	0b	RW	Reflects the state of PHY register bit 1E.CC01.E before it was cleared by a firmware read.
GLOBAL_RESERVED	31:17	0x0	RW	Reserved. Read as written.



8.2.2.1.10 LED Control - LEDCTL (0x00000200)

Field	Bit(s)	Init.	Access Type	Description
LED0_MODE	3:0	0x0	RW	LED0 Mode. This field specifies the control source for the LED0 output. An initial value of 0000b selects the LINK_UP indication.
RESERVED	4	0b	RSV	Reserved.
GLOBAL_BLINK_MODE	5	0b	RW	Global Blink Mode. This field specifies the blink mode of all LEDs. 0b = Blink at 200 ms on and 200 ms off. 1b = Blink at 83 ms on and 83 ms off.
LED0_IVRT	6	0b	RW	LED0 Invert. This field specifies the polarity/inversion of the LED0 source prior to output or blink control. By default the output drives the cathode of the LED0 so when the LED output is 0b the LED0 is on. 0b = LED0 output is active low. 1b = LED0 output is active high.
LED0_BLINK	7	0b	RW	LED0 Blink. This field specifies whether or not to apply blink logic to the (inverted) LED0 control source prior to the LED0 output. 0b = Do not blink LED0 output. 1b = Blink LED0 output.
LED1_MODE	11:8	0x1	RW	LED1 Mode. This field specifies the control source for the LED1 output. An initial value of 0001b selects the 10 Gb/s link indication.
RESERVED	13:12	0x0	RSV	Reserved.
LED1_IVRT	14	0b	RW	LED1 Invert. This field specifies the polarity/inversion of the LED1 source prior to output or blink control. By default the output drives the cathode of the LED so when the LED1 output is 0b the LED1 is on. 0b = LED1 output is active low. 1b = LED output is active high.
LED1_BLINK	15	1b	RW	LED1 Blink. This field specifies whether or not to apply blink logic to the (inverted) LED1 control source prior to the LED output. 0b = Do not blink LED output. 1b = Blink LED output.
RESERVED	31:20	0x0	RSV	Reserved.

Notes: All init bits in this register are read from the NVM. See LED Configuration section in NVM section.



8.2.2.1.11 Extended VLAN Ether Type - Receive - EXVET (0x00005078)

Field	Bit(s)	Init.	Access Type	Description
RESERVED	15:0	0x00	RSV	Reserved.
VET_EXT	31:16	0x8100	RW	Outer VLAN Ether Type. The VLAN Tag Protocol Identifier (TPID). Note: This field appears in little endian (MS byte first on the wire).

8.2.2.1.12 Extended VLAN Ether Type - Transmit - EXVET_T (0x00008224)

Field	Bit(s)	Init.	Access Type	Description
RESERVED	15:0	0x00	RSV	Reserved.
VET_EXT	31:16	0x8100	RW	Outer VLAN Ether Type. The VLAN Tag Protocol Identifier (TPID). Note: This field appears in little endian (MS byte first on the wire).

8.2.2.1.13 General Receive Control - GRC (0x00010200)

Field	Bit(s)	Init.	Access Type	Description
RESERVED	0	0b	RSV	
APME	1	0b	RW	Advance Power Management Enable. If set to 1b, manageability wake up is enabled. The Intel® Xeon® Processor D-1500 Product Family LAN controller sets the PME_Status bit in the Power Management Control/Status Register (PMCSR), asserts PE_WAKE_N when manageability wake up is enabled, and when it receives a matching magic packet. It is a single read/write bit in a single register, but has two values depending on the function that accesses the register. The value of this bit is loaded from NVM control word 3 (bit 0 for port 0 and bit 1 for port 1).
RESERVED	31:2	0x0	RSV	Reserved.

8.2.2.1.14 SFP I²C Command - I2CCMD (0x00015F58)

This register is used by software to read or write to the configuration registers in an SFP module when the SDP23_function bit and the SDP[23]_NATIVE bit are set in ESDP register.

Field	Bit(s)	Init.	Access Type	Description
DATA	15:0	X	RW	Data in a write command, software places the data bits and then the MAC shifts them out to the I ² C bus. Data in a read command, the MAC reads these bits serially from the I ² C bus and then software reads them from this location. This field is interpreted as two consecutive bytes unless the I2CPARAMS.I2C_DATA_ORDER bit is set, in which case the data is considered as a single 16 bit word. When I2CPARAMS.ACCESS_WIDTH = 0b, bits 15:8 are not used.
REGADD	23:16	0x0	RW	I ² C Register Address. For example, register 0, 1, 2... 255.
RESERVED	26:24	0x0	RSV	Reserved.



Field	Bit(s)	Init.	Access Type	Description
OP	27	0b	RW	OpCode. 0b = I ² C write. 1b = I ² C read.
RESET	28	0b	RW	Reset Sequence. If set, sends a reset sequence before the actual read or write. This bit is self clearing. A reset sequence is defined as nine consecutive stop conditions.
R	29	1b	RW	Ready Bit. Set to 1b by the Intel® Xeon® Processor D-1500 Product Family LAN controller at the end of the I ² C transaction. Indicates the read or write completed. Reset by a software write of a command.
RESERVED	30	0b	RSV	Reserved.
E	31	0b	RW	Error. This bit set is to 1b by hardware when it fails to complete an I ² C read. Reset by a software write of a command. Note: This bit is valid only when the Ready (R) bit is set.

8.2.2.1.15 SFP I²C Parameters - I2CPARAMS (0x00015F5C)

This register is used to set the parameters for I²C access and to enable bit-banging access to the I²C interface.

Field	Bit(s)	Init.	Access Type	Description
WRITE_TIME	4:0	0x6	RW	Write Time. Defines the delay between a write access and the next access. The value is in milliseconds. Note that a value of zero is not valid.
READ_TIME	7:5	0x2	RW	Read Time. Defines the delay between a read access and the next access. The value is in microseconds. Note that a value of zero is not valid.
I2CBB_EN	8	0b	RW	I ² C Bit-bang Enable. If set, the I2C_CLK and I2C_DATA lines are controlled via the CLK, DATA, and DATA_OE_N fields of this register. Otherwise, they are controlled by hardware once activated via the I2CCMD register.
CLK_OUT	9	0b	RW	I ² C Clock. While in bit-bang mode, controls the value driven on the I2C_CLK pad of this port.
DATA_OUT	10	0b	RW	I ² C_DATA. While in bit-bang mode and when the DATA_OE_N field is zero, controls the value driven on the I2C_DATA pad of this port.
DATA_OE_N	11	0b	RW	I2C_DATA_OE_N. While in bit-bang mode, controls the direction of the I2C_DATA pad of this port. 0b = Pad is output. 1b = Pad is input.
DATA_IN	12	0b	RO	I ² C Data In. Reflects the value of the I2C_DATA pad. While in bit-bang mode when the DATA_OE_N field is zero, this field reflects the value set in the DATA_OUT field.
CLK_OE_N	13	0b	RW	I ² C Clock Output Enable. While in bit-bang mode, controls the direction of the I2C_CLK pad of this port. 0b = Pad is output. 1b = Pad is input.
CLK_IN	14	0b	RO	I ² C Clock In Value. Reflects the value of the I2C_CLK pad. While in bit-bang mode when the CLK_OE_N field is zero, this field reflects the value set in the CLK_OUT field.
CLK_STRETCH_DIS	15	0b	RW	Clock Stretch Disable. 0b = Enable slave clock stretching support in an I ² C access. 1b = Disable clock stretching support in an I ² C access.



Field	Bit(s)	Init.	Access Type	Description
ACCESS_WIDTH	16	0b	RW	I ² C Access width: 0b: Byte access 1b: Word access
RESERVED	23:17	0x0	RSV	Reserved. Write 0x0, ignore on read.
PHYADD	30:24	0x0	RW	Device Address Bits 7:1. The actual address used is b{PHYADD[6:0], 0}.
I2C_DATA_ORDER	31	0b	RW	I ² C Data Order. 0b = I2CCMD.DATA field read in byte order. 1b = I2CCMD.DATA field read in word order.

Notes: This register is reset only on Power Good Reset.

8.2.2.1.16 Device and Functions Enable Control - DEV_FUNC_EN (0x00010208)

Field	Bit(s)	Init.	Access Type	Description
LAN_PCI_DISABLE	0	0b	RO	LAN PCI Disable. When set, one LAN port is disabled. The function that is disabled is determined by the LAN <i>Disable Select</i> bit. If the disabled function is function 0, it acts as a dummy function or the other LAN function depending on the dummy function enable setting. If the disabled port is used for WoL or by MC, only the DMA block of the port is powered down. Otherwise, the port is powered down including the PHY. If the LAN <i>PCI Disable</i> bit is set for a port, the port's respective APM bit in NVM Control Word 3 must be cleared as well.
LAN_DISABLE_SELECT	1	0b	RO	LAN Disable Select 0b = LAN 0 is disabled. 1b = LAN 1 is disabled.
DEV_OFF_EN	2	0b	RO	Device Electrical Off Enable. This bit is relevant only when the device is disabled via strapping during PE_RST_N both LANn_DIS_N pins to 0b at once. 0b = Legacy mode (default). Though the Intel® Xeon® Processor D-1500 Product Family LAN controller is disabled, the digital I/O pins are not moved to an electrical off state. 1b = Enable device electrical off. When the Intel® Xeon® Processor D-1500 Product Family LAN controller is disabled, the digital I/O pins are put at High-Z. For example, electrical off state where pull-ups/downs are at their defined values.
SDP_FUNC_OFF_EN	3	0b	RO	PCIe Function Off via SDP Pins Enable. 0b = Legacy mode (default), SDPn_1 pins does not control PCIe functions off. 1b = Two SDP pins are used in conjunction by strapping (sampled by PE_RST_N) to disable the two PCIe functions altogether. Note: If MNG is present, MC-to-LAN paths are not disabled. PCIe function off pins can also work independently, where each pin controls the corresponding PCIe function. However, independent mode is not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
IGNORE_PHY_FW_VALID	4	0b	RO	Reflects the ignore OHY firmware valid indication in NVM control word 2.
RESERVED	31:5	0x0	RSV	Reserved.

8.2.2.1.17 Function Active and Power State to Manageability - FACTPS (0x00010150)



Register for use by the device firmware for configuration.

Field	Bit(s)	Init.	Access Type	Description
FUNC0_POWER_STATE	1:0	0x0	RO	Power state indication of function 0. 00b = DR. 01b = D0u. 10b = D0a. 11b = D3.
LAN0_VALID	2	0b	RO	LAN 0 Enable. When this bit is set to 0b, it indicates that the LAN 0 function is disabled. When the function is enabled, the bit is set to 1b. This status of this bit reflects whether or not the function is disabled through the external pad.
FUNC0_AUX_EN	3	0b	RO	Function 0 Auxiliary (AUX) Power PM Enable bit shadow from the configuration space.
RESERVED	5:4	0x0	RSV	Reserved.
FUNC1_POWER_STATE	7:6	0x0	RO	Power state indication of function 1: 00b = DR. 01b = D0u. 10b = D0a. 11b = D3.
LAN1_VALID	8	0b	RO	LAN 1 Enable. When this bit is set to 0b, it indicates that the LAN 1 function is disabled. When the function is enabled, this bit is set to 1b. This status of this bit reflects whether or not the function is disabled through the external pad.
FUNC1_AUX_EN	9	0b	RO	Function 1 Auxiliary (AUX) Power PM Enable bit shadow from the configuration space.
RESERVED	29:10	0x0	RSV	Reserved.
LAN_FUNCTION_SEL	30	0b	RO	When both LAN ports are enabled and the LAN Function Sel equals 0b, LAN 0 is routed to PCI function 0 and LAN 1 is routed to PCI function 1. If the LAN Function Sel equals 1b, LAN 0 is routed to PCI function 1 and LAN 1 is routed to PCI function 0. This bit is loaded from NVM.
PM_STATE_CHANGED	31	0b	RO	An indication that one or more of the functional power states had changed. This bit is also a signal to manageability to create an interrupt. This bit is cleared on read and is not set for at least eight cycles after it was cleared.

Notes: In regular mode (port swap disabled - LAN_FUNCTION_SEL = 0b), the power state indication of port 0 is mapped to the FUNC0_POWER_STATE field and the power state indication of port 1 is mapped to the FUNC1_POWER_STATE field. And vice versa when port swap mode is enabled (LAN_FUNCTION_SEL = 1).

8.2.2.2 NVM Registers

8.2.2.2.1 EEPROM Mode Control Register - EEC (0x00010010)

Field	Bit(s)	Init.	Access Type	Description
RESERVED	7:0	0x0	RSV	Reserved. Reads as 0b.
EE_PRESENT	8	0b	RO	NVM Present. When this bit is set, indicates that an NVM is present and has the correct signature field. This bit is read-only.



Field	Bit(s)	Init.	Access Type	Description
AUTO_RD	9	0b	RO	NVM Auto-Read Done. When set to 1b, this bit indicates that the auto-read by hardware from the NVM caused by the latest reset for this port is done. This bit is also set when the NVM is not present or when its signature field is not valid. This bit does not reflect the status of the PHY image auto-load process. Use <i>Reset Completed</i> (bit 6) in PHY register 1E.CC00 to get this indication.
MNG_READY	10	0b	RO	Management Autoload Done: When set, indicates the sections needed for manageability are valid.
EE_SIZE	14:11	0x7	RO	NVM Size Via EEPROM Mode. This field defines the size of the NVM that is accessible via EEPROM mode. This is equal to the size of the internal shadow RAM, fixed to 16 KB (0x7 defines 16 KB - legacy).
PCI_ANA_DONE	15	0b	RO	PCIe Analog Done. When set to 1b, indicates that the PCIe analog section read from NVM is done. This bit is cleared when auto-read starts. This bit is also set when the NVM is not present or when its signature field is not valid.
PCI_CORE_DONE	16	0b	RO	PCIe Core Done. When set to 1b, indicates that the core analog section read from NVM is done. This bit is cleared when auto-read starts. This bit is also set when the NVM is not present or when its signature field is not valid. Note: This bit returns the relevant done indication for the function that reads the register.
PCI_GENERAL_DONE	17	0b	RO	PCIe General Done. When set to 1b, indicates that the PCIe general section read from the NVM is done. This bit is cleared when auto-read starts. This bit is also set when the NVM is not present or when its signature field is not valid.
PCI_FUNC_DONE	18	0b	RO	PCIe Function Done. When set to 1b, indicates that the PCIe function section read from the NVM is done. This bit is cleared when auto-read starts. This bit is also set when the NVM is not present or when its signature field is not valid. Note: This bit returns the relevant done indication for the function that reads the register.
CORE_DONE	19	0b	RO	Core Done. When set to 1b, indicates that the core section read from the NVM is done. This bit is cleared when auto-read starts. This bit is also set when the NVM is not present or when its signature field is not valid. Note: This bit returns the relevant done indication for the function that reads the register.
CORE_CSR_DONE	20	0b	RO	Core CSR Done. When set to 1b, indicates that the core CSR section read from the NVM is done. This bit is cleared when auto-read starts. This bit is also set when the NVM is not present or when its signature field is not valid. Note: This bit returns the relevant done indication for the function that reads the register.
MAC_DONE	21	0b	RO	MAC Done. When set to 1b, indicates that the MAC section read from the NVM is done. This bit is cleared when auto-read starts. This bit is also set when the NVM is not present or when its signature field is not valid. Note: This bit returns the relevant done indication for the function that reads the register.
LCB_DONE	22	0b	RO	LCB Done. When set, indicates that the LCB section read from the NVM is done. This bit is cleared when auto-read starts. This bit is also set when the NVM is not present or when its signature field is not valid.
FLUPD	23	0b	SC	Flash Update. Writing 1b to this bit causes the content of the internal shadow RAM to be written into one of the first two sections of the Flash device (Section 0 or Section 1). The bit is self cleared. Note: Setting this bit issues an interrupt to firmware.
RESERVED	24	0b	RSV	Reserved. Reads as 0b.



Field	Bit(s)	Init.	Access Type	Description
SEC1VAL	25	0b	RO	Sector 1 Valid. When set to 1b, indicates that the content of the second shadow RAM section (from byte address 0x4000 to 0x7FFF) of the Flash device is valid. When set to 0b, indicates that the content of the first section (from byte address 0x0000 to 0x3FFF) is valid. Meaningful only when the EE_PRES bit is read as 1b. This bit can be written by firmware.
FLUDONE	26	1b	RO	Flash Update Done. When set to 1b, indicates that the last Flash update process completed. This bit is set by firmware, and auto clears when EEC.FLUPD is set.
RESERVED	31:28	0x0	RSV	Reserved. Reads as 0b.

8.2.2.2.2 EEPROM Mode Read Register - EERD (0x00010014)

Field	Bit(s)	Init.	Access Type	Description
START	0	0b	RW	Start Read. Writing a 1b to this bit causes the read of a 16-bit word from the shadow RAM at the address stored in the ADDR field and then stores the result in the DATA field. This bit is cleared at the end of the read operation when the DONE bit is set.
DONE	1	0b	RO	Read Done. Set this bit to 1b when the EEPROM mode read completes. Set this bit to 0b when the EEPROM mode read is in progress. Note that writes by software are ignored.
ADDR	14:2	0x0	RW	Read Address. This field is written by software along with Start Read to indicate the address of the word to read.
RESERVED	15	0b	RSV	Reserved. Reads as 0b.
DATA	31:16	0x0	RO	Read Data. Data returned from the EEPROM mode read.

Note: This register is used by software to read individual words from the internal shadow RAM that reflects the first valid 16 KB of the NVM. To read a word, software writes the address to the *Read Address* field and simultaneously writes a 1b to the *Start Read* field. The Intel® Xeon® Processor D-1500 Product Family LAN controller reads the word from the internal shadow RAM and places it in the *Read Data* field, setting the *Read Done* field to 1b. Software can poll this register, looking for a 1b in the *Read Done* field and then using the value in the *Read Data* field. When this register is used to read a word from the NVM via EEPROM mode, that word is not written to any of the Intel® Xeon® Processor D-1500 Product Family LAN controller's internal registers even if it is normally a hardware-accessed word.

8.2.2.2.3 EEPROM Mode Write Register - EEWR (0x00010018)

Field	Bit(s)	Init.	Access Type	Description
START	0	0b	RW	Start Write. Writing a 1b to this bit causes the write of a 16-bit word from the <i>Data</i> field into the shadow RAM at the address stored in the ADDR field. This bit is cleared at the end of the write operation when the DONE bit is set.
DONE	1	1b	RO	Write Done. Cleared when the NVM write starts by setting the START bit. Set when the NVM write completes.
ADDR	14:2	0x0	RW	Write Address. This field is written by software along with <i>Start Write</i> to indicate the address of the word to write.
RESERVED	15	0b	RSV	Reserved. Reads as 0b.



Field	Bit(s)	Init.	Access Type	Description
DATA	31:16	0x0	RW	Write Data. Data to be written into the shadow RAM.

Note: This register is used by software to write individual words in the internal shadow RAM that reflects the first valid 16 KB of the NVM. To write a word, software writes the address to the *Write Address* field, the data to the *Write Data* field, and simultaneously writes a 1b to the *Start Write* field. The Intel® Xeon® Processor D-1500 Product Family LAN controller writes the word into the internal shadow RAM, setting the *Write Done* field to 1b. Software can poll this register, looking for a 1b in the *Write Done* field before the next write. The data is effectively copied into the Flash device using the EEC.FLUPD command. When this register is used to write a word into the NVM, that word is not written to any of the Intel® Xeon® Processor D-1500 Product Family LAN controller's internal registers even if it is normally a hardware-accessed word.

8.2.2.2.4 Flash Access Register - FLA (0x0001001C)

Field	Bit(s)	Init.	Access Type	Description
FL_SCK	0	0b	RW	Clock input to the Flash. When FL_GNT is set to 1b, the FL_SCK output signal is mapped to this bit and provides the serial clock input to the Flash. Software clocks the Flash via toggling this bit with successive writes.
FL_CE	1	0b	RW	Chip select input to the Flash. When FL_GNT is set to 1b, the FL_CE output signal is mapped to the chip select of the Flash device. Software enables the Flash by writing a 0b to this bit.
FL_SI	2	0b	RW	Data input to the Flash. When FL_GNT is set to 1b, the FL_SI output signal is mapped directly to this bit. Software provides data input to the Flash via writes to this bit.
FL_SO	3	0b	ROS	Data output bit from the Flash. The FL_SO input signal is mapped directly to this bit in the register and contains the Flash serial data output. This bit is read-only from a software perspective. Note that writes to this bit have no effect.
FL_REQ	4	0b	RW	Request Flash Access. Software must write a 1b to this bit to get direct Flash access. It has access when FL_GNT is set to 1b. When software completes the access, it must then write a 0b.
FL_GNT	5	0b	RO	Grant Flash Access. When this bit is set to 1b, software can access the Flash using the FL_SCK, FL_CE, FL_SI, and FL_SO bits.
LOCKED	6	0b	RO	When set to 1b, this bit indicates that the NVM is in secure mode. When set to 0b, the NVM is in non-secure mode. This bit is set if the SRAMREL.FLEEP_RO_MODE is set and the NVM_SEC_DIS strap is not set
NVM_SEC_DIS	7	0b	RO	Reflects the value of the NVM_SEC_DIS strap.
RESERVED	11:8	0x0	RSV	Reserved. Reads as 0b.
FL_SIZE	14:12	0x5	RO	Flash Size. Indicates Flash size of 64 KB * 2 ** "FL_SIZE". The Flash size impacts the requested memory space for the Flash and expansion ROM BARs in PCIe configuration space (see CSRSize and FLSize fields in the PCIE_LBARCTRL register). Supported Flash sizes: 101b = 2 MB. 110b = 4 MB. 111b = 8 MB.
RESERVED	15	0b	RSV	Reserved. Reads as 0b.
FL_SADDR	28:16	0x0	RW	Flash Sector Erase Address. Determines which 4 KB sector is erased when the FL_SER command is used. This address is expressed in sector index units, starting from sector index 0.



Field	Bit(s)	Init.	Access Type	Description
FL_SER	29	0b	RW	Flash Sector Erase Command. This bit is auto-cleared and reads as 0b. The 4 KB sector index to be erased is determined by the FL_SADDR field.
FL_BUSY	30	0b	RO	Flash Busy. This bit is set to 1b by hardware while Flash access has been granted to a client. Note: This bit is read-only from a software perspective.
FL_DER	31	0b	RW	Flash Device Erase Command. This bit is auto-cleared and reads as 0b. The entire Flash device is erased.

Note: When NVM is locked, this register is RO to software.

8.2.2.2.5 Flash Firmware Code Update - FLUPDATE (0x00015F54)

Field	Bit(s)	Init.	Access Type	Description
MODULEID	15:0	0x0	RW	The module in the Flash to be updated from the free area when the FLUPDATE.Update bit is set. Should be compared to the LSB of the module ID in the CSS authentication header.
RESERVED	28:16	0x0	RSV	Reserved.
AUT_DONE	29	0b	ROS	Authentication Cycle Done. Set to 1b by firmware when done. This bit is self-cleared once the update request is set to 1b.
AUT_FAIL	30	0b	ROS	Authentication failed. Set to 1b by firmware when authentication failed.
UPDATE	31	0b	SC	Request authentication of the new secure section written. If the authentication succeeds, firmware resets itself to load its new code. This bit is self-cleared, always read as 0b.

8.2.2.2.6 Software Flash Burst Control Register - FLSWCTL (0x00015F48)

Flash access control. When software writes this register, firmware gets an interrupt.

Field	Bit(s)	Init.	Access Type	Description
ADDR	23:0	0x0	RW	Address. This field is written by software along with CMD to indicate the Flash address to do operation (read/write/erase etc.). See command description that follows.
CMD	27:24	0x0	RW	Command. Indicates which command should be executed. Valid only when the CMDV bit is set. 0000b = Read. 0001b = Write. 0010b = Sector erase. 0011b-1111b = Reserved.
CMDV	28	0b	RO	Command Valid. Set by hardware when a successful write of this register is received. When set, indicates that software issued a new command and the command is in progress. Cleared by firmware at the end of the command. If software tries to set this bit while a command is in progress (DONE is not set), then hardware clears the bit.



Field	Bit(s)	Init.	Access Type	Description
COMMAND_FAIL	29	0b	RO	When set, means the command was not executed due to security issues. For example, attempt to write to a read only area or request of an unknown command. This bit is valid only if <i>DONE</i> bit is set.
DONE	30	1b	RO	Partial Command Done. This bit clears after the register is written by software and is set back again when the Flash transaction completes. When writing a burst transaction, the bit is cleared by hardware every time software writes FLSWDATA. When doing a read burst transaction, the bit is cleared by hardware every time software reads FLSWDATA. Firmware writes sets this bit after the data was written/read from the Flash.
GLDONE	31	1b	RO	Global Done. This bit is cleared by hardware after the register is written by software and is set back again by firmware when all Flash read/write transactions completes (FLSWCNT.NVCNT bytes where read/written).

8.2.2.2.7 JEDEC ID 1 - JEDEC_ID_1 (0x00015F2C)

Exposes the JEDEC ID of the connected Flash.

Field	Bit(s)	Init.	Access Type	Description
BANK	7:0	0x0	ROS	Defines the bank number of the manufacturer ID (number of 0x7F read).
MANUFACTURER_ID	15:8	0x0	ROS	Returns the manufacturer ID read by RDID command.
DEVICE_ID_1	23:16	0x0	ROS	First byte of the device ID read by RDID command (<i>Family</i> and <i>Density</i> fields).
DEVICE_ID_2	31:24	0x0	ROS	Second byte of the device ID read by RDID command (sub- version and revision).

Note: This register is valid only if a valid firmware is present.

8.2.2.2.8 Manageability Flash Control Register - FLMNGCTL (0x00010118)

Note: This register can be read/write by manageability firmware and is read-only to host software.

Field	Bit(s)	Init.	Access Type	Description
ADDR	23:0	0x0	RW	Address. This field is written by manageability along with CMD and CMDV to indicate which Flash address to read or write. For the Sector Erase command (CMD = 10b), it must point to any address in the 4 KB sector to be erased.
CMD	25:24	0x0	RW	Command. Indicates which command should be executed. Valid only when the CMDV bit is set. 00b = Read command. 01b = Write command (single byte). 10b = Sector erase. The ADDR field determines the 4 KB sector index to be erased by this command. 11b = Device erase.



Field	Bit(s)	Init.	Access Type	Description
CMDV	26	0b	RW	Command Valid. When set, indicates that the manageability firmware issues a new command and is cleared by hardware at the end of the command.
FLBUSY	27	0b	RW	Flash Busy. This bit indicates that the Flash is busy processing a Flash transaction and should not be accessed.
RESERVED	29:28	0x0	RSV	Reserved.
DATA_VALID	30	0b	RW	Read Done and data in FLMNGDATA is valid. This bit is cleared by firmware when it sets the CMDV bit. It is set by hardware for each DWord read that completes. This bit is read/clear by hardware enabling the multiple DWord read flow.
GLDONE	31	1b	RW	Global Done. This bit clears after the CMDV bit is set by manageability and is set back again when all Flash transactions complete. For example, the Flash device finished reading all the requested read or other accesses (write and erase).

8.2.2.2.9 Software Flash Burst Data Register - FLSWDATA (0x00015F4C)

This register holds the data read/written to the Flash. Writing or reading of this register sends an interrupt to Firmware and clears the FLSWCTL.DONE bit.

Field	Bit(s)	Init.	Access Type	Description
NVDATA	31:0	0x0	RW	NVM Data. Data written or read to/from the NVM.

8.2.2.2.10 Software Flash Burst Access Counter - FLSWCNT (0x00015F50)

Field	Bit(s)	Init.	Access Type	Description
NVCNT	12:0	0x0	RW	NVM Burst Counter. This counter holds the size in bytes of the Flash burst read or write.
RESERVED	31:13	0x0	RSV	Reserved.

8.2.2.2.11 Manageability EEPROM-Mode Control Register - EEMNGCTL (0x00010110)

Note: This register can be read/write by manageability firmware and is read-only to host software. The transactions performed through this register are directed to/from the internal shadow RAM. The write data is effectively copied into the Flash device by use of the EEC.FLUPD command.

Field	Bit(s)	Init.	Access Type	Description
ADDR	12:0	0x0	RW	Address. This field is written by manageability along with the <i>START</i> bit and the <i>WRITE</i> bit to indicate which NVM word address to read or write. Only the first valid 16 KB is accessible through this EEPROM mode interface dedicated to manageability.
RESERVED	14:13	0x0	RSV	Reserved. Reads as 0b.



Field	Bit(s)	Init.	Access Type	Description
START	15	0b	RW	Start. Writing a 1b to this bit causes the Intel® Xeon® Processor D-1500 Product Family LAN controller to start the read or write operation of the shadow RAM according to the write bit. This bit is self cleared by hardware. If one of the CFG_DONE0/1 bits values is set in this write, writing of this bit is ignored.
WRITE	16	0b	RW	Write. This bit signals the Intel® Xeon® Processor D-1500 Product Family LAN controller if the current operation is read or write. 0b = Read. 1b = Write.
EEBUSY	17	0b	RW	EEPROM Mode Busy. This bit indicates that the internal shadow RAM is busy and should not be accessed.
CFG_DONE0	18	0b	RW	Manageability configuration cycle of port 0 completed. This bit indicates that the manageability configuration cycle (configuration of PCIe, core or PHY) completed. This bit is set to 1b by manageability firmware to indicate that the configuration completed and is cleared by software. Writing a 0b by firmware does not affect the state of this bit. Writing a 1b by software does not affect the state of this bit. Note: Software should not try to access the PHY for configuration before this bit is set. Refer to Section 3.0 .
CFG_DONE1	19	0b	RW	Manageability configuration cycle of port 1 completed. This bit indicates that the manageability configuration cycle (configuration of PCIe, core or PHY) completed. This bit is set to 1b by manageability firmware to indicate that the configuration completed and is cleared by software. Writing a 0b by firmware does not affect the state of this bit. Writing a 1b by software does not affect the state of this bit. Note: Software should not try to access the PHY for configuration before this bit is set. Refer to Section 3.0 .
TRANS_ABORTED	20	0b	RO	When read as 1b, it indicates that the current EEMNGCTL access was aborted due to a firmware reset. The bit is cleared by hardware on the next write access to the EEMNGCTL register, regardless of the written value.
RESERVED	30:21	0x0	RSV	Reserved.
DONE	31	1b	RW	Transaction Done. This bit is cleared when the <i>START</i> bit and <i>WRITE</i> bit are set by manageability and is set back again when the shadow RAM write or read transaction completes.

8.2.2.2.12 Manageability EEPROM-Mode Read/Write Data - EEMNGDATA (0x00010114)

Note: This register can be read/write by manageability firmware and is read-only to host software.

Field	Bit(s)	Init.	Access Type	Description
WRDATA	15:0	0x0	RW	Write Data. Data to be written to the shadow RAM.
RDDATA	31:16	X	RW	Read Data. Data returned from the read command. Note: This field is read only.

8.2.2.2.13 Manageability Flash Read/Write Data - FLMNGDATA (0x0001011C)

Note: This register can be read/write by manageability firmware and is read-only to host software.



Field	Bit(s)	Init.	Access Type	Description
DATA	31:0	0x0	RW	Read/Write Data On a Read Transaction. This register contains the data returned from the Flash read. On write transactions, bits 7:0 are written to the Flash.

8.2.2.2.14 MAC Flow Control Register - MFLCN (0x00004294)

Field	Bit(s)	Init.	Access Type	Description
PMCF	0	0b	RW	Pass MAC Control Frames. Filter out unrecognized pause (flow control OpCode doesn't match) and other control frames. 0b = Filter Unrecognized Pause Frames. 1b = Pass/forward unrecognized pause frames.
DPF	1	0b	RW	Discard Pause Frame. When set to 0b, pause frames are sent to the host. Setting this bit to 1b causes pause frames to be discarded only when <i>RFCE</i> is set to 1b. If <i>RFCE</i> is set to 0b, this bit has no effect on incoming pause frames.
Reserved	2	0b	RW	Reserved.
RFCE	3	0b	RW	Receive Link Flow Control Enable. Indicates that the Intel® Xeon® Processor D-1500 Product Family LAN controller responds to the reception of link flow control packets. If auto-negotiation is enabled, this bit should be set by software to the negotiated flow control value. Note: This bit should not be set if bit 2 is set.
RESERVED	31:4	0x0	RSV	Reserved.

8.2.2.2.15 Flow Control Transmit Timer Value n - FCTTVN[n] (0x00003200 + 0x4*n, n=0...3)

Each 32-bit register (n=0... 3) refers to two timer values (register 0 refers to timer 0 and 1; register 1 refers to timer 2 and 3, etc.).

Field	Bit(s)	Init.	Access Type	Description
TTV_2N	15:0	0x0	RW	Transmit Timer Value 2n. Timer value included in XOFF frames as Timer (2n). The same value must be set to User Priorities (UPs) attached to the same TC, as defined in the RTTUP2TC register. For legacy 802.3X flow control packets, TTV0 is the only timer that is used.
TTV_2N_1	31:16	0x0	RW	Transmit Timer Value 2n+1. Timer value included in XOFF frames as Timer 2n+1. The same value must be set to UPs attached to the same TC, as defined in the RTTUP2TC register.

Note: The 16-bit value in the *TTV* field is inserted into a transmitted frame (either XOFF frames or any pause frame value in any software transmitted packets). It counts in units of slot time (usually 64 bytes).

Note: The Intel® Xeon® Processor D-1500 Product Family LAN controller uses a fixed slot time value of 64-byte times.

8.2.2.2.16 Flow Control Receive Threshold Low - FCRTL[n] (0x00003220 + 0x4*n, n=0...7)



Each 32-bit register (n=0... 7) refers to a different receive packet buffer.

Field	Bit(s)	Init.	Access Type	Description
RESERVED	4:0	0x0	RSV	Reserved.
RTL	18:5	0x0	RW	Receive Threshold Low n. Receive packet buffer n FIFO low-water mark for flow control transmission (32 bytes granularity).
RESERVED	30:19	0x0	RSV	Reserved.
XONE	31	0b	RW	XON Enable n. Per the receive packet buffer XON enable. 0b = Disabled. 1b = Enabled.

Note: This register contains the receive threshold used to determine when to send an XON packet and counts in units of bytes. The lower four bits must be programmed to 0x0 (16-byte granularity). Software must set XONE to enable the transmission of XON frames. Each time incoming packets cross the receive high threshold (become more full), and then crosses the receive low threshold, with XONE enabled (1b), hardware transmits an XON frame.

8.2.2.2.17 Flow Control Receive Threshold High - FCRTV[n] (0x00003260 + 0x4*n, n=0...7)

Each 32-bit register (n=0... 7) refers to a different receive packet buffer.

Field	Bit(s)	Init.	Access Type	Description
RESERVED	4:0	0x0	RSV	Reserved.
RTH	18:5	0x0	RW	Receive Threshold High n. Receive packet buffer n FIFO high-water mark for flow control transmission (32 bytes granularity).
RESERVED	30:19	0x0	RSV	Reserved.
FCEN	31	0b	RW	Transmit Flow control enable for packet buffer n. Should be set only for TCs to which UPs are mapped (in RTRUP2TC register).

Note: This register contains the receive threshold used to determine when to send an XOFF packet and counts in units of bytes. This value must be at least eight bytes less than the maximum number of bytes allocated to the receive packet buffer and the lower five bits must be programmed to 0x0 (32-byte granularity). Each time the receive FIFO reaches the fullness indicated by RTH, hardware transmits a pause frame if the transmission of flow control frames is enabled.

8.2.2.2.18 Flow Control Refresh Threshold Value - FCRTV (0x000032A0)

Field	Bit(s)	Init.	Access Type	Description
FC_REFRESH_TH	15:0	0x0	RW	Flow Control Refresh Threshold. This value is used to calculate the actual refresh period for sending the next pause frame if conditions for a pause state are still valid (buffer fullness above low threshold value). The formula for the refresh period for priority group N is: FCTTV[N/2].TTV[Nmod2] - FCRTV.FC_refresh_th Note: The FC_refresh_th must be smaller than TTV of the TC and larger than the maximum packet size in the TC + FC packet size + link latency and Tx latency and Rx latency in 64-byte units.



Field	Bit(s)	Init.	Access Type	Description
RESERVED	31:16	0x0	RSV	Reserved.

8.2.2.2.19 Transmit Flow Control Status - TFCS (0x0000CE00)

Field	Bit(s)	Init.	Access Type	Description
TC_XON	7:0	0xFF	RO	TC is in FC XON state.
RESERVED	31:8	0x0	RSV	Reserved.

8.2.2.2.20 Flow Control Configuration - FCCFG (0x00003D00)

Field	Bit(s)	Init.	Access Type	Description
RESERVED	2:0	0x0	RSV	Reserved.
TFCE	4:3	0x0	RW	Transmit Flow Control Enable. These bits Indicate that the Intel® Xeon® Processor D-1500 Product Family LAN controller transmits flow control packets (XON/XOFF frames) based on receive fullness. If auto-negotiation is enabled, this bit should be set by software to the negotiated flow control value. 00b = Transmit flow control disabled. 01b = Link flow control enabled. 10b = Reserved. 11b = Reserved.
RESERVED	31:5	0x0	RSV	Reserved.

8.2.2.2.21 Priority Flow Control Type Opcode - PFCTOP (0x0000431C)

Field	Bit(s)	Init.	Access Type	Description
RESERVED	2:0	0x0	RSV	Reserved.
FCT	15:0	0x8808	RW	Priority Flow Control EtherType. Note that this field appears in big endian (MS byte first on the wire).
FCOP	31:16	0x0101	RW	Priority Flow Control Opcode. Note that this field appears in big endian (LS byte first on the wire).

8.2.2.3 PCIe Registers

This section contains the registers used to control the PCIe core behavior.

8.2.2.3.1 PCIe Function Status 1 - PCI_STATUS1 (0x00011028)



Field	Bit(s)	Init.	Access Type	Description
FUNC_VALID	0	0b	RO	Func valid: 0b = Function is disabled. 1b = Function is enabled. Note: This bit is valid to firmware even when the function is disabled.
RESERVED	31:1	0x0	RSV	Reserved.

8.2.2.3.2 PCIe Control Extended Register - GCR_EXT (0x00011050)

Field	Bit(s)	Init.	Access Type	Description
VT_MODE	1:0	0x0	RW	VT mode of operation defines the allocation of physical registers to the VFs. Software must set this field the same as GPIE. VT_Mode.00 = noVT - Reserved. For the case that STATUS. IOV Ena is not set. 01b = VT16 - Resources are allocated to 16 VFs. 10b = VT32 - Resources are allocated to 32 VFs 11b = VT64 - Resources are allocated to 64 VFs.
RESERVED	3:2	0x0	RSV	Reserved.
APBACD	4	0b	RW	Auto PBA Clear Disable. When set to 1b, software can clear the PBA only by direct write to clear access to the PBA bit. When set to 0b, any active PBA entry is cleared on the falling edge of the appropriate interrupt request to the PCIe block. The appropriate interrupt request is cleared when software sets the associated interrupt mask bit in the EIMS (re-enabling the interrupt) or by direct write to clear to the PBA.
RESERVED	30:5	0x00	RSV	Reserved

8.2.2.3.3 PCIe Revision ID - PCI_REVID (0x00011098)

Field	Bit(s)	Init.	Access Type	Description
NVM_REVID	7:0	0x00	RW	Value of rev ID loaded from NVM.
RESERVED	31:8	0x00	RSV	Reserved.

8.2.2.3.4 PCIe Default Revision ID - PCI_DREVID (0x00011540)

Field	Bit(s)	Init.	Access Type	Description
DEFAULT_REVID	7:0	0x00	RO	Mirroring of default rev ID prior to NVM load.
RESERVED	31:8	0x00	RSV	Reserved.

8.2.2.3.5 Function Requester ID Information Register - FUNC_RID (0x00011070)



Field	Bit(s)	Init.	Access Type	Description
FUNCTION_NUMBER	2:0	0x0	RO	Function number assigned to the function based on BIOS/OS enumeration.
DEVICE_NUMBER	7:3	0x0	RO	Device Number. Device number assigned to the function based on BIOS/OS enumeration.
BUS_NUMBER	15:8	0x0	RO	Bus Number. Bus number assigned to the function based on BIOS/OS enumeration.
RESERVED	31:16	0x0	RSV	Reserved.

8.2.2.3.6 PCIe VM Pending Index - PCI_VMINDEX (0x00011530)

Field	Bit(s)	Init.	Access Type	Description
VMINDEX	8:0	0x0	RW	VM Index. Software sets the VMINDEX that its transaction pending flag should be reflected in the PFPCI_VMPEND register. The VM index is an absolute index in the range of 0 through 63. It can be set by the software only to VMs that the PF owns and only to VMs which are not assigned to a VF.
RSVD	31:9	0x0	RSV	Reserved.

8.2.2.3.7 PCIe VM Pending Status - PCI_VMPEND (0x00011538)

Field	Bit(s)	Init.	Access Type	Description
PENDING	0	0b	RO	PCIe transaction pending status. The reported VM is controlled by the VMINDEX field in the PFPCI_VMINDEX register. This flag is set to 1b as long as there is at least one PCIe transaction, pending for its completion.
RSVD	31:1	0x0	RSV	Reserved

8.2.2.3.8 PCIe Interrupts Enable - PCI_IENA (0x00011528)

Field	Bit(s)	Init.	Access Type	Description
PCI_ERR_EN	31:0	0x0	RW	Each bit corresponds to a particular error event as described in section Proprietary Error Reporting.

8.2.2.3.9 PCIe Interrupt Cause - PCI_ICAUSE (0x00011520)

Field	Bit(s)	Init.	Access Type	Description
PCI_ERR_CAUSE	31:0	0x0	RW1C	Each bit corresponds to a particular error event as described in section Proprietary Error Reporting.

8.2.2.3.10 PCIe Errors Reported - PCI_PCIERR (0x00011140)

This register indicates which PCIe errors are reported to device software.



Field	Bit(s)	Init.	Access Type	Description
PCI_ERR_REP	31:0	0x0	RO	Each bit corresponds to a particular error event as described in section Proprietary Error Reporting.

8.2.2.3.11 PCIe Statistic Control Register #1 - PCI_GSCL_1 (0x00011800)

This register controls the operation of the PCIe performance counters.

Field	Bit(s)	Init.	Access Type	Description
GIO_COUNT_EN_0	0	0b	RW	Enables PCIe statistic counter number 0.
GIO_COUNT_EN_1	1	0b	RW	Enables PCIe statistic counter number 1.
GIO_COUNT_EN_2	2	0b	RW	Enables PCIe statistic counter number 2.
GIO_COUNT_EN_3	3	0b	RW	Enables PCIe statistic counter number 3.
LBC_ENABLE_0	4	0b	RW	When set, statistics counter 0 operates in leaky bucket mode. When cleared, leaky bucket mode is disabled and the counter is incremented by one for each event.
LBC_ENABLE_1	5	0b	RW	When set, statistics counter 1 operates in leaky bucket mode. When cleared, leaky bucket mode is disabled and the counter is incremented by one for each event.
LBC_ENABLE_2	6	0b	RW	When set, statistics counter 2 operates in leaky bucket mode. When cleared, leaky bucket mode is disabled and the counter is incremented by one for each event.
LBC_ENABLE_3	7	0b	RW	When set, statistics counter 3 operates in leaky bucket mode. When cleared, leaky bucket mode is disabled and the counter is incremented by one for each event.
PCI_COUNT_LAT_EN	8	0b	RW	PCIe Count Latency Enable. Enables the latency counter 0b = Disable the latency counter. 1b = Enable the latency counter.
PCI_COUNT_LAT_EV	13:9	0x00	RW	PCIe Count Latency Event. Selects the event to be measured.
PCI_COUNT_BW_EN	14	0b	RW	PCIe Count Bandwidth Enable. Enables the bandwidth counter. 0b = Disable the bandwidth counter. 1b = Enable the bandwidth counter.
PCI_COUNT_BW_EV	19:15	0x00	RW	PCIe Count Bandwidth Event. Selects the event to be measured.
RESERVED	27:20	0x0	RSV	Reserved.
GIO_64_BIT_EN	28	0b	RW	Enables two 64-bit counters instead of four 32-bit counters.
GIO_COUNT_RESET	29	0b	RWS	Reset indication of PCIe statistic counters. Writing zero has no effect.
GIO_COUNT_STOP	30	0b	RWS	Stop indication of PCIe statistic counters. Writing zero has no effect.
GIO_COUNT_START	31	0b	RWS	Start indication of PCIe statistic counters. Writing zero has no effect.

8.2.2.3.12 PCIe Statistic Control Registers #2 - PCI_GSCL_2 (0x00011804)



This register defines the events counted by the performance counters.

Field	Bit(s)	Init.	Access Type	Description
GIO_EVENT_NUM_0	7:0	0x0	RW	Event number that counter 0 counts (GSCN_0).
GIO_EVENT_NUM_1	15:8	0x0	RW	Event number that counter 1 counts (GSCN_1).
GIO_EVENT_NUM_2	23:16	0x0	RW	Event number that counter 2 counts (GSCN_2).
GIO_EVENT_NUM_3	31:24	0x0	RW	Event number that counter 3 counts (GSCN_3).

8.2.2.3.13 PCIe Statistic Counter Registers #0...#3 - PCI_GSCN_0_3[n] (0x00011820 + 0x4*n, n=0...3)

These registers contain the performance counters 0-3.

GSCL_0 corresponds to n=0.

GSCL_1 corresponds to n=1.

GSCL_2 corresponds to n=2.

GSCL_3 corresponds to n=3.

Field	Bit(s)	Init.	Access Type	Description
EVENT_COUNTER	31:0	0x0	RO	Event counter as defined in GSCL_2.GIO_EVENT_NUM fields. These registers are stuck at their maximum value of 0xFF..F.

8.2.2.3.14 PCIe Statistic Control Register #5...#8 - PCI_GSCL_5_8[n] (0x00011810 + 0x4*n, n=0...3)

These registers control the operation of the leaky bucket counter n.

GSCL_5 corresponds to n=0.

GSCL_6 corresponds to n=1.

GSCL_7 corresponds to n=2.

GSCL_8 corresponds to n=3.

Field	Bit(s)	Init.	Access Type	Description
LBC_THRESH_OLD_N	15:0	0x00	RW	Threshold for the leaky bucket counter n.
LBC_TIMER_N	31:16	0x00	RW	Time period between decrementing the value in leaky bucket counter n. The time period is defined in μ s units.

8.2.2.3.15 PCIe Byte Counter Low - PCI_BYTCTL (0x00011548)

A byte counter used by the PCIe performance counters.



Field	Bit(s)	Init.	Access Type	Description
PCI_COUNT_BW_BCT	31:0	0x00	RO	This register contains the low Dword of a 64-bit counter that counts PCIe payload bytes. This register gets stuck at its maximum value of 0xFF...F.

8.2.2.3.16 PCIe Byte Counter High - PCI_BYTCTH (0x00011544)

A byte counter used by the PCIe performance counters.

Field	Bit(s)	Init.	Access Type	Description
PCI_COUNT_BW_BCT	31:0	0x00	RO	This register contains the high Dword of a 64-bit counter that counts PCIe payload bytes. This register gets stuck at its maximum value of 0xFF...F.

8.2.2.3.17 PCIe Latency Counter - PCI_LATCT (0x00011720)

A latency counter used by the PCIe performance counters.

Field	Bit(s)	Init.	Access Type	Description
PCI_COUNT_LAT_CT	31:0	0x00	RO	A 32-bit counter that counts time in multiples of 100 ns. This register gets stuck at its maximum value of 0xFF...F.

8.2.2.3.18 PCIe Packet Counter - PCI_PKTCT (0x00011740)

A packet counter used by the PCIe performance counters.

Field	Bit(s)	Init.	Access Type	Description
PCI_COUNT_BW_PCT	31:0	0x0	RO	A 32-bit counter that counts PCIe packets. This register gets stuck at its maximum value of 0xFF...F.

8.2.2.3.19 PCIe LCB Address Port - PCI_LCBADD (0x00011788)

Field	Bit(s)	Init.	Access Type	Description
ADDRESS	17:0	0x0	RW	An 18-bit address register, part of a port, used to load the NVM configuration into the PCIe LCB unit. Operates together with the PCI_LCBDATA register.
RESERVED	19:18	0x0	RSV	Reserved.



Field	Bit(s)	Init.	Access Type	Description
BLOCK_ID	27:20	0x0	RW	An ID of a sub-unit in the PCIe unit. Supported values are: 0. NPQ: RLAN. 1. NPQ: TLAN. 2. NPQ: RX_PE. 3. NPQ: TX_PE. 4. NPQ: PMAT. 5. NPQ: MNG. 6. NPQ: TDPU. 7. PQ: RLAN. 8. PQ: TLAN. 9. PQ: RX_PE. 10. PQ: TX_PE. 11. PQ: PMAT. 12. PQ: MNG. 13. PQ: RDPU. 14. VDM. 15. Don't care. 16. HIU: CFG/Slave/MSG/MSI-X. 126: LCB's internal config space registers. 127: LCB's internal memory space registers.
RESERVED	31:28	0x0	RSV	Reserved

8.2.2.3.20 PCIe LCB Data Port - PCI_LCBDATA (0x00011734)

Field	Bit(s)	Init.	Access Type	Description
LCB_DATA	31:0	0x0	RW	A 32-bit data register used for indirect access to the PCIe core module's registers. The address is in PCI_LCBADDR.

8.2.2.3.21 PCIe Firmware Control - PCI_FWCTRL (0x00011040)

Field	Bit(s)	Init.	Access Type	Description
TCO_ISOLATE	0	0b	RO	Set by firmware to block all host write access if instructed to do so by the BMC.
RESERVED	31:1	0x0	RW	Reserved

8.2.2.3.22 PCIe CSR Access Timeout - PCI_CSRT0 (0x00011044)

Field	Bit(s)	Init.	Access Type	Description
CSR_TO	15:0	0x000F	RO	Defines the timeout value in μ s for CSR accesses.
MSIX_TO	31:16	0x0028	RO	Defines the timeout value in μ s for MSI-X table accesses.

8.2.2.3.23 PCI Flash Access Timeout - PCI_FLASHTO (0x00011054)



Field	Bit(s)	Init.	Access Type	Description
PCI_FLASHTO	31:0	0x00000FA0	RO	Defines the timeout value in micro second for flash accesses

8.2.2.4 PCIe Configuration Space Setting Registers

This section contains registers used to set the default setting of the PCIe configuration space. These registers are reset and loaded from the NVM at PCIe reset.

8.2.2.4.1 PCIe VF Capabilities Support - PCI_VFSUP (0x000110A8)

Field	Bit(s)	Init.	Access Type	Description
VF_PREFETCH_H	0	1b	RW	VF Prefetchable. 0b = IOV memory BARs are declared as non-prefetchable. 1b = IOV memory BARs are declared as prefetchable.
RESERVED	31:1	0x0	RSV	Reserved

8.2.2.4.2 PCIe Global Config 2 - PCI_CNF2 (0x000110F8)

This register contains global status fields of the PCIe configuration.

Field	Bit(s)	Init.	Access Type	Description
RESERVED	0	0b	RSV	Reserved.
CACHELINE_SIZE	1	0b	RW	Cache Line Size.Determines the system cache line size. 0b = 64 bytes. 1b = 128 bytes. This field is loaded from the NVM.
MSI_X_PF_N	12:2	0x3F	RW	MSI_X PF Table Size.System software reads this field to determine the MSI-X table size N, which is encoded as N-1. This field is loaded from the NVM <i>MSI_X_N_PF</i> field and reflects the same field in the PCIe MSI-X configuration (Message Control register).
MSI_X_VF_N	23:13	0x02	RW	MSI_X VF Table Size.System software reads this field to determine the MSI-X table size N for VFs, which is encoded as N-1. This field is loaded from the NVM <i>MSI_X_N_VF</i> field and reflects the same field in the PCIe MSI-X configuration (VF MSI-X Control register)
NUM_VFS	31:24	0x40	RW	The number of VFs requested by the function. Valid values are 0 to 64. If the value is zero, then SR-IOV capability is hidden.

8.2.2.4.3 PCI BAR Control - PCI_LBARCTRL (0x00011048)

Field	Bit(s)	Init.	Access Type	Description
PREFBAR	0	1b	RW	Prefetchable bit indication in the memory BARs (should be set when 64-bit BARs are used). 0b = BARs are marked as non prefetchable. 1b = BARs are marked as prefetchable.



Field	Bit(s)	Init.	Access Type	Description
BAR32	1	0b	RW	BAR 32-bit Enable. 0b = 64-bit BAR addressing mode is selected. 1b = 32-bit BARs are enabled.
CSRSIZE	2	0b	RW	The <i>CSRSIZE</i> and <i>FLSIZE</i> fields define the usable Flash size and CSR mapping window size as described in Section 9.0 .
FLASH_EXPOSE	3	1b	RW	When set, the Flash memory is accessible through the memory BAR.
RESERVED	5:4	0x0	RSV	Reserved.
FL_BAR_SIZE	8:6	0x5	RW	This field indicates the size of the external Flash as = 64 KB x (2 ** FL_BAR_SIZE) as used to define exposure in the BAR. The following values are supported: 101b = 2 MB. 110b = 4 MB. 111b = 8 MB. Else = Reserved.
RESERVED	10:9	0x0	RW	Reserved.
EXROM_BAR_SIZE	13:11	0x3	RW	This field indicates the size of the expansion ROM BAR as = 64 KB x (2 ** EXROM_BAR_SIZE). 000b corresponds to a 64 KB size. 111b corresponds to a 8 MB size. Default value is 512 KB.
RESERVED	31:14	0x0	RSV	Reserved.

8.2.2.4.4 PCIe Global Config - PCI_GLBL_CNF (0x000110B8)

Field	Bit(s)	Init.	Access Type	Description
RESERVED	1:0	0x0	RSV	Reserved.
WAKE_PIN_EN	2	0b	RW	When set to 1b, enables the use of PE_WAKE_N pin for PME event in all power states.
PCI_CLKGATE_DIS	3	1b	RW	When set to 0b, enables dynamic clock gating of the PCIe clocks (LCB, HIU and CSR).
PCI_CLKGATE_L1_ONLY	4	1b	RW	When set to 1b, and if PCI_CLKGATE_DIS is 0b, the clock gating of the PCIe clocks are only when the PCIe is in L1 state.
PCI_PCLK_GATE_EN	5	0b	RW	When set to 1b, and if PCI_CLKGATE_DIS is 0b, the PCIw PCLK is also be dynamically gated.
RESERVED	7:6	0x0	RSV	Reserved.
PCI_CLKGATE_TIMER	15:8	0x10	RW	Clock Gating Idle Timer. This field defines the number of clocks (CSR clock) of idle-detect before gating the PCIe clocks.
RESERVED	31:16	0x0	RSV	Reserved.

8.2.2.4.5 PCIe Capabilities Support - PCI_CAPSUP (0x00011088)

Determines PCIe capabilities supported by the Intel® Xeon® Processor D-1500 Product Family LAN controller.



Field	Bit(s)	Init.	Access Type	Description
PCIE_VER	0	1b	RW	Determines the PCIe capability version. 0b = Capability version: 0x1. 1b = Capability version: 0x2.
RESERVED	1	0b	RSV	Reserved.
LTR_EN	2	1b	RW	A value of 1b indicates support for the PCIe Latency Tolerance Reporting (LTR) capability.
TPH_EN	3	1b	RW	A value of 1b indicates support for the PCIe TPH Requester capability.
ARI_EN	4	1b	RW	A value of 1b indicates support for the PCIe ARI capability.
IOV_EN	5	1b	RW	A value of 1b indicates support for the PCIe SR-IOV capability.
ACS_EN	6	1b	RW	A value of 1b indicates support for the PCIe ACS capability.
SEC_EN	7	0b	RW	A value of 1b indicates support for the secondary PCIe extended capability.
RESERVED	14:8	0x0	RSV	Reserved
RESERVED	15	0b	RW	Reserved.
ECRC_GEN_EN	16	1b	RW	Loaded into the ECRC <i>Generation Capable</i> bit of the PCIe Configuration registers.
ECRC_CHK_EN	17	1b	RW	Loaded into the ECRC <i>Check Capable</i> bit of the PCIe Configuration registers.
IDO_EN	18	1b	RW	Enables ID-based Ordering (IDO).
MSI_MASK	19	1b	RW	MSI Per-vector Masking Setting. This bit is loaded to the masking bit (bit 8) in the message control of the MSI configuration capability structure.
CSR_CONF_EN	20	1b	RW	Enables Access to CSRs via the PCI configuration space. See Section Configuration Access to Internal Registers and Memories.
RESERVED	29:21	0x0	RSV	Reserved.
LOAD_SUBSYS_ID	30	0b	RW	Load Subsystem IDs. When set to 1b, indicates that the Intel® Xeon® Processor D-1500 Product Family LAN controller loads its PCIe subsystem ID and sub-system vendor ID from the NVM.
LOAD_DEV_ID	31	0b	RW	Load Device ID. When set to 1b, indicates that the Intel® Xeon® Processor D-1500 Product Family LAN controller loads its PCI device ID's from the NVM.

8.2.2.4.6 PCIe Link Capabilities - PCI_LINKCAP (0x00011090)

Determines PCIe Link capabilities supported by the Intel® Xeon® Processor D-1500 Product Family LAN controller.

Field	Bit(s)	Init.	Access Type	Description
MAX_PAYLOAD	8:6	0x2	RW	Max Payload Size Supported. Loaded to the PCIe Device Capabilities register Supported values are 000b to 100b (128bytes to 2 KB). Otherwise, keep the hardware default.



Field	Bit(s)	Init.	Access Type	Description
MAX_LINK_WIDTH_IDTH	12:9	0x07	RW	Max Link Width. Loaded to the PCIe Link Capabilities register Values: 0001b = Limit max link width to x1. 0011b = Limit max link width to x4. 0100b = Limit max link width to x8. 0111b = Do not limit max link width. negotiate to the max width supported by the link. Else = Reserved.
RESERVED	31:13	0x0	RSV	Reserved.

8.2.2.4.7 PCIe Capabilities Control - PCI_CAPCTRL (0x00011080)

Determines PCIe capabilities supported by the Intel® Xeon® Processor D-1500 Product Family LAN controller and that software is allowed to enable or disable.

Field	Bit(s)	Init.	Access Type	Description
VPD_EN	0	0b	RW	0b = The PCIe VPD Capability is not present and is not exposed. 1b = The PCIe VPD Capability is present and exposed.
RESERVED	31:1	0x0	RSV	Reserved.

8.2.2.4.8 PCIe Serial Number MAC Address Low - PCI_SERL (0x000110F0)

Field	Bit(s)	Init.	Access Type	Description
SER_NUM_L	31:0	0x0	RW	The low Dword of the Ethernet MAC address used to generate the PCIe serial number.

8.2.2.4.9 PCIe Serial Number MAC Address High - PCI_SERH (0x00011078)

Field	Bit(s)	Init.	Access Type	Description
SER_NUM_H	15:0	0x0	RW	The high word of the Ethernet MAC address used to generate the PCIe serial number.
RESERVED	31:16	0x0	RSV	Reserved.

8.2.2.4.10 PCIe Storage Class - PCI_CLASS (0x00011038)

Contains the per-PF configuration loaded from the NVM.

Field	Bit(s)	Init.	Access Type	Description
STORAGE_CLASS	0	0b	RW	0b = The class code of this port is set to 0x020000 (LAN). 1b = The class code of this port is set to 0x010000 (SCSI).
RESERVED	31:1	0x0	RSV	Reserved.

8.2.2.4.11 PCIe PF Configuration - PCI_CNF (0x00011000)



Contains the per-PF configuration loaded from the NVM.

Field	Bit(s)	Init.	Access Type	Description
RESERVED	2:0	0x0	RW	Reserved.
EXROM_DIS	3	1b	RW	Expansion ROM Disable. 0b = The Expansion ROM BAR in the PCI configuration space is enabled. 1b = The Expansion ROM BAR in the PCI configuration space is disabled.
IO_BAR	4	0b	RW	I/O BAR Support. 0b = I/O BAR is not supported. 1b = I/O BAR is supported.
INT_PIN	6:5	0x0	RW	Controls the value advertised in the <i>Interrupt Pin</i> field of the PCI configuration header for this function. Values of 00b, 01b, 10b and 11b correspond to INTA#, INTB#, INTC# and INTD#, respectively. The value advertised in the PCI configuration header is the value loaded from NVM + 1. The default value for port 1 is 0x1.
RESERVED	31:7	0x0	RSV	Reserved.

8.2.2.4.12 PCIe PM Support - PCI_PMSUP (0x000110A0)

This register contains parameters that define PCIe power management support.

Field	Bit(s)	Init.	Access Type	Description
ASPM_SUP	1:0	0x3	RW	Active State Link PM Support. Loaded to the PCIe Link Capabilities register.
L0S_EXIT_LAT	4:2	0x5	RW	L0s Exit Latency. Loaded to the L0s <i>Exit Latency</i> field in the PCIe Link Capabilities register.
L1_EXIT_LAT	7:5	0x4	RW	L1 Exit Latency. Loaded to L1 <i>Exit Latency field</i> in the PCIe Link Capabilities register.
L0S_ACC_LAT	10:8	0x3	RW	Loaded to the Endpoint L0s <i>Acceptable Latency</i> field in the PCIe Device Capabilities register.
L1_ACC_LAT	13:11	0x6	RW	Loaded to the Endpoint L1 <i>Acceptable Latency</i> field in the PCIe Device Capabilities register.
SLOT_CLK	14	1b	RW	Slot Clock Configuration. Loaded to the PCIe Link Status register.
OBFF_SUP	16:15	0x0	RW	Loaded to the OBFF <i>Supported</i> field in the PCIe Device Capabilities 2 register. Must be set to 0b in the NVM (OBFF is not supported).
RESERVED	31:17	0x0	RSV	Reserved.

8.2.2.4.13 PCIe Power Data Register - PCI_PWRDATA (0x00011060)

Field	Bit(s)	Init.	Access Type	Description
D0_POWER	7:0	0x0	RW	The value in this field is reflected in the PCI Power Management Data register of the LAN functions for D0 power consumption and dissipation (Data_Select = 0 or 4).



Field	Bit(s)	Init.	Access Type	Description
COMM_POWER	15:8	0x0	RW	The value in this field is reflected in the PCI Power Management Data register of function 0 when the Data_Select field is set to 8 (common function).
D3_POWER	23:16	0x0	RW	The value in this field is reflected in the PCI Power Management Data register of the LAN functions for D3 power consumption and dissipation (Data_Select = 3 or 7).
RESERVED	31:24	0x0	RSV	Reserved.

8.2.2.4.14 PCIe Vendor ID - PCI_VENDORID (0x0001103C)

The Vendor ID exposed in the configuration space. A value of 0xFFFF is not loaded to the configuration space. Shared for all PFs.

Field	Bit(s)	Init.	Access Type	Description
VENDOR_ID	15:0	0x8086	RW	Contains the vendor ID exposed in offset 0x0 in the configuration space of all functions. A value of 0xFFFF is ignored.
RESERVED	31:16	0x0	RSV	Reserved.

8.2.2.4.15 PCIe PF Device ID - PCI_PFDEVID (0x00011008)

Contains the per-PF Device ID.

Field	Bit(s)	Init.	Access Type	Description
PF_DEV_ID_LAN	15:0	0x0	RW	Contains the device ID for this PF when the function has an Ethernet device class code. The actual default value of this field is defined per device.
PF_DEV_ID_SAN	31:16	0x0	RW	Contains the device ID for this PF when the function has a SCSI device class code. The actual default value of this field is defined per device.

8.2.2.4.16 PCIe VF Device ID - PCI_VFDEVID (0x00011010)

Contains the per-PF device IDs for its VFs.

Field	Bit(s)	Init.	Access Type	Description
VF_DEV_ID_LAN	15:0	0x0	RW	Contains the device ID for this PF's VFs when the function has an Ethernet device class code. The actual default value of this field is defined per device.
VF_DEV_ID_SAN	31:16	0x0	RW	Contains the device ID for this PF's VFs when the function has a SCSI device class code. The actual default value of this field is defined per device.

8.2.2.4.17 PCIe Subsystem ID - PCI_SUBSYSID (0x00011058)



Field	Bit(s)	Init.	Access Type	Description
SUB_VEN_ID	15:0	0x8086	RW	Loaded to the PCI configuration Subsystem Vendor ID register.
SUB_ID	31:16	0x0	RW	Loaded to the PCI configuration Subsystem ID register.

8.2.2.4.18 PCIe Upper Address - PCI_UPADD (0x000110E8)

This register is used to block PCIe master accesses above some address.

Field	Bit(s)	Init.	Access Type	Description
RESERVED	0	0b	RSV	Reserved.
ADDRESS	31:1	0x0	RW	Address.Bits [31:1] correspond to bits [63:33] in the PCIe address space, respectively.

8.2.2.4.19 Extended Interrupt Cause Register - EICR (0x00000800)

The EICR register is RW1C and can be optionally cleared on a read depending on the ODC flag setting in the GPIE register.

Field	Bit(s)	Init.	Access Type	Description
RTXQ	15:0	0x0	RW1C	Receive/Transmit Queue Interrupts. One bit per queue or a bundle of queues, activated on receive/transmit events. The mapping of queue to the RTXQ bits is done by the IVAR registers.
FLOW_DIRECTOR	16	0b	RW1C	Flow director exception is activated by one of the following events: <ol style="list-style-type: none"> 1. Filter removal failed (there was no matched filter to be removed). 2. The number of remaining free filters in the flexible filter table exceeds (goes below) the FDIRCTRL.Full-Thresh. 3. Filter programming failed due to no space in the flow director table. Note that this case should not happen if the software device driver handles the FDIRCTRL.Full-Thresh event.
RX_MISS	17	0b	RW1C	Missed packet interrupt is activated for each received packet that overflows the Rx packet buffer (overrun). Note that the packet is dropped and also increments the associated RXMPC[n] counter.
PCI_EXCEPTION	18	0b	RW1C	The PCI exception is activated by one of the events described in the Proprietary Error Reporting section. The specific PCI event is error reported in the PCI_ICAUSE register:
MAILBOX	19	0b	RW1C	VF-to-PF Mailbox Interrupt. Caused by a VF write access to the PF mailbox or by a malicious event detection.
LSC	20	0b	RW1C	Link Status Change. This bit is set each time the link status changes (either from up to down, or from down to up).
LINKSERVED	21	0b	RW1C	Indicates that the Tx MACsec packet counter reached the threshold requiring key exchange. Reserved.
MNG	22	0b	RW1C	Manageability Event Detected. Indicates that a manageability event happened. When the Intel® Xeon® Processor D-1500 Product Family LAN controller is in power-down mode, the BMC might generate a PME for the same events that would cause an interrupt when the Intel® Xeon® Processor D-1500 Product Family LAN controller is at the D0 state. This interrupt bit is also used to alert the host when the BMC IP address was changed or when EEMNGCTL.CFG_DONE0/1 bit is set by firmware.



Field	Bit(s)	Init.	Access Type	Description
RESERVED	23	0b	RW1C	Reserved.
TIMESYNC	24	0b	RW1C	TimeSync Interrupt. Indicates that a Time Sync event has occurred. Check TSIM register for precise cause.
GPI_SDP0	25	0b	RW1C	General Purpose Interrupt on SDP0. If GPI interrupt detection is enabled on this pin (via GPIE), this interrupt cause is set when a transition to high is sampled on SDP0.
GPI_SDP1	26	0b	RW1C	General Purpose Interrupt on SDP1. If GPI interrupt detection is enabled on this pin (via GPIE), this interrupt cause is set when a transition to high is sampled on SDP1.
GPI_SDP2	27	0b	RW1C	General Purpose Interrupt on SDP2. If GPI interrupt detection is enabled on this pin (via GPIE), this interrupt cause is set when the SDP2 is sampled high.
ECC	28	0b	RW1C	Unrecoverable ECC error. This bit is set when one of the following occurs: <ul style="list-style-type: none"> An unrecoverable error is detected in one of the device memories. A CRC error occurred on the second attempt to load the PHY image from NVM. PHY micro-controller watchdog failure. Software should issue a software reset following this error.
PHY_GLOBAL_INTERRUPT	29	0b	RW1C	PHY Interrupt (non fatal).
TCP_TIMER	30	0b	RW1C	TCP Timer Expired. This bit is set when the timer expires.
OTHER_CAUSE	31	0b	ROS	Other Cause Interrupt. Activated when any bit (29:16) in the Extended Interrupt Cause Register (EICR) is set and its relevant mask bit in the EIMS is enabled.

8.2.2.4.20 Extended Interrupt Cause Set Register - EICS (0x00000808)

Field	Bit(s)	Init.	Access Type	Description
INTERRUPT_CAUSE_SET	30:0	0x0	WO	Setting any bit in this field, sets its corresponding bit in the EICR and generates an interrupt if enabled by EIMS register.
RESERVED	31	0b	RSV	Reserved.

8.2.2.4.21 Extended Interrupt Mask Set/Read Register - EIMS (0x00000880)

Field	Bit(s)	Init.	Access Type	Description
INTERRUPT_ENABLE	30:0	0x0	RWS	Each bit that is set to 1b enables its corresponding interrupt in the EICR. Writing a 1b to any bit sets it. Writing 0b has no impact. Reading this register provides a map of those interrupts that are enabled.
RESERVED	31	0b	RSV	Reserved.

8.2.2.4.22 Extended Interrupt Mask Clear Register - EIMC (0x00000888)



Field	Bit(s)	Init.	Access Type	Description
INTERRUPT_MASK	30:0	0x0	WO	Writing a 1b to any bit clears its corresponding bit in the EIMS disabling the corresponding interrupt in the EICR. Writing 0b has no impact. Reading this register provides no meaningful data.
RESERVED	31	0b	RSV	Reserved.

8.2.2.4.23 Extended Interrupt Auto Clear Register - EIAC (0x00000810)

Field	Bit(s)	Init.	Access Type	Description
RTXQ_AUTO_CLEAR	15:0	0x0	RW	At 1b each bit enables auto clear of the corresponding RTXQ bits in the EICR following interrupt assertion. At 0b the corresponding bits in the EICR are not auto cleared.
RESERVED	29:16	0x0	RSV	Reserved.
TCP_TIMER_AUTO_CLEAR	30	0b	RW	When set to 1b, this bit enables auto clear of the TCP timer interrupt cause in the EICR following interrupt assertion. When set to 0b, auto-clear is not enabled.
RESERVED	31	0b	RSV	Reserved.

Note: Bits 29:16 should never be set to auto clear since they share the same MSI-X vector.

8.2.2.4.24 Extended Interrupt Auto Mask Enable Register - EIAM (0x00000890)

Field	Bit(s)	Init.	Access Type	Description
AUTO_MASK	30:0	0x0	RW	When set to 1b, each bit enables auto set and clear of its corresponding bits in the EIMS. Note that if any of the <i>Auto Mask</i> enable bits is set, the GPIE.EIAME bit must be set as well.
RESERVED	31	0b	RSV	Reserved.


8.2.2.4.25 Extended Interrupt Cause Set Registers 1 - EICS1 (0x00000A90)

Field	Bit(s)	Init.	Access Type	Description
INTERRUPT_CAUSE_SET	31:0	0x0	WO	Setting any bit in these registers, sets its corresponding bit in the EICR[n] and generates an interrupt if enabled by EIMS[n] register. Reading this register provides no meaningful data.

8.2.2.4.26 Extended Interrupt Cause Set Registers 2 - EICS2 (0x00000A94)

Field	Bit(s)	Init.	Access Type	Description
INTERRUPT_CAUSE_SET	31:0	0x0	WO	Setting any bit in these registers, sets its corresponding bit in the EICR[n] and generates an interrupt if enabled by EIMS[n] register. Reading this register provides no meaningful data.

8.2.2.4.27 Extended Interrupt Mask Set/Read Registers - EIMS1 (0x00000AA0)



Field	Bit(s)	Init.	Access Type	Description
INTERRUPT_ENABLE	31:0	0x0	RWS	Each bit when set to 1b enables its corresponding interrupt in the EICR[n]. Writing 1b to any bit sets it. Writing 0b has no impact. Reading this register provides a map of those interrupts that are enabled. Bits 15:0 of EIMS1 are mirrored in EIMS bits 15:0.

8.2.2.4.28 Extended Interrupt Mask Set/Read Registers - EIMS2 (0x00000AA4)

Field	Bit(s)	Init.	Access Type	Description
INTERRUPT_ENABLE	31:0	0x0	RWS	When set to 1b,each bit enables its corresponding interrupt in the EICR[n]. Writing 1b to any bit sets it. Writing 0b has no impact. Reading this register provides a map of those interrupts that are enabled. Bits 15:0 of EIMS1 are mirrored in EIMS bits 15:0.



8.2.2.4.29 Extended Interrupt Mask Clear Registers 1 - EIMC1 (0x00000AB0)

Field	Bit(s)	Init.	Access Type	Description
INTERRUPT_MASK	31:0	0x0	WO	Writing 1b to any bit clears its corresponding bit in the EIMS[n] disabling the corresponding interrupt in the EICR[n]. Writing 0b has no impact. Reading this register provides no meaningful data.

8.2.2.4.30 Extended Interrupt Mask Clear Registers 2 - EIMC2 (0x00000AB4)

Field	Bit(s)	Init.	Access Type	Description
INTERRUPT_MASK	31:0	0x0	WO	Writing 1b to any bit clears its corresponding bit in the EIMS[n] disabling the corresponding interrupt in the EICR[n]. Writing 0b has no impact. Reading this register provides no meaningful data.

8.2.2.4.31 Extended Interrupt Auto Mask Enable registers 1 - EIAM1 (0x00000AD0)

Field	Bit(s)	Init.	Access Type	Description
AUTO_MASK	31:0	0x0	RW	When set to 1b, each bit enables auto set and clear of its corresponding bits in the EIMS[n]. Bits 15:0 of EIAM1 are mirrored in EIAM bits 15:0. Note that if any of the <i>Auto Mask</i> enable bits is set, the GPIE.EIAME bit must be set as well.

8.2.2.4.32 Extended Interrupt Auto Mask Enable Registers 2 - EIAM2 (0x00000AD4)

Field	Bit(s)	Init.	Access Type	Description
AUTO_MASK	31:0	0x0	RW	When set to 1b, each bit enables auto set and clear of its corresponding bits in the EIMS[n]. Bits 15:0 of EIAM1 are mirrored in EIAM bits 15:0. Note that if any of the <i>Auto Mask</i> enable bits is set, the GPIE.EIAME bit must be set as well.

8.2.2.4.33 MSIX to EITR Select - EITRSEL (0x00000894)

Field	Bit(s)	Init.	Access Type	Description
VFSELECT	31:0	0x0	RW	Each bit n in this register selects the VF index (32+n) or PF interrupt source for the EITR registers (VF 0-31 are not multiplexed as described in Section 7.3.4.3.3). At 0x0 it selects the PF and at 0x1 it selects the VF.

8.2.2.4.34 Extended Interrupt Throttle Registers - EITR[n] (0x00000820 + 0x4*n, n=0...23 and 0x00012300 + 0x4*(n-24), n=24...128)

Mapping of the EITR registers to the MSI-X vectors is described in MSI-X Vectors Mapping to EITR. See Section 7.3.4.3.3.



Field	Bit(s)	Init.	Access Type	Description
RESERVED	2:0	0x0	RSV	Reserved.
ITR_INTERVAL	11:3	0x0	RW	Minimum inter-interrupt interval specified in 2.048 μ s units at 1 GbE and 10 GbE link. At 100 Mb/s link, the interval is specified in 20.48 μ s units. At 0x0 interrupt throttling is disabled while any event causes an immediate interrupt.
RESERVED	13:12	0x0	RSV	Reserved.
HIGH_PRIORITY	14	0b	RW	High Priority Vector. Setting this bit causes assertion of this vector to break DMA coalescing.
LLI_MODERATION	15	0b	RW	When set, LLI moderation is enabled. Otherwise, any LLI packet generates an immediate interrupt. LLI moderation can be set only if interrupt throttling is enabled by the ITR <i>Interval</i> field in this register and LLI moderation is enabled by the LL <i>Interval</i> field in the GPIE register.
LLI_CREDIT	20:16	0x0	RW	Reflects the current credits for associated interrupt. When CNT_WDIS is not set on a write cycle, this field must be set to zero.
ITR_COUNTER	27:21	0x0	RW	This field represents the 7 MSb (out of 9 bits) of the ITR counter. It is a down counter that is loaded with an ITR interval value each time the associated interrupt is asserted. When the ITR counter reaches zero it stops counting and triggers an interrupt. On a write cycle, software must set this field to 0b if CNT_WDIS in this register is cleared (write enable to the ITR counter).
RESERVED	30:28	0x0	RSV	Reserved.
CNT_WDIS	31	0b	RW	Write disable to the LLI credit and ITR counter. When set, the LLI credit and ITR counter are not overwritten by the write access. When cleared, software must set the LLI credit and ITR counter to zero, which enables an immediate interrupt on packet reception. This bit is write only and always read as zero.

Notes: Additional address(es): 0x012300 + 4*(n-24), n=24...128.

8.2.2.4.35 Interrupt Vector Allocation Registers - IVAR[n] (0x00000900 + 0x4*n, n=0...63)

These registers map interrupt causes into EICR entries (legacy/MSI modes) or into MSI-X vectors (MSI-X modes). See [Section 7.3.4](#) for mapping and use of these registers.

Field	Bit(s)	Init.	Access Type	Description
INT_ALLOC_0	5:0	X	RW	The interrupt allocation for Rx queue (2xN for IVAR register N).
RESERVED	6	0b	RSV	Reserved.
INT_ALLOC_VALID_0	7	0b	RW	Interrupt allocation valid indication for INT_Alloc[0].
INT_ALLOC_1	13:8	X	RW	The interrupt allocation for Tx queue (2xN for IVAR register N).
RESERVED	14	0b	RSV	Reserved.
INT_ALLOC_VALID_1	15	0b	RW	Interrupt allocation valid indication for INT_Alloc[1].
INT_ALLOC_2	21:16	X	RW	The interrupt allocation for Rx queue (2xN+1 for IVAR register N).
RESERVED	22	0b	RSV	Reserved.
INT_ALLOC_VALID_2	23	0b	RW	Interrupt allocation valid indication for INT_Alloc[2].



Field	Bit(s)	Init.	Access Type	Description
INT_ALLOC_3	29:24	X	RW	The interrupt allocation for Tx queue (2xN+1 for IVAR register N).
RESERVED	30	0b	RSV	Reserved.
INT_ALLOC_VAL_3	31	0b	RW	Interrupt allocation valid indication for INT_Alloc[3].

8.2.2.4.36 Miscellaneous Interrupt Vector Allocation - IVAR_MISC (0x0000A00)

These register maps interrupt causes into MSI-X vectors (MSI-X modes). See [Section 7.3.4](#) for mapping and use of these registers.

Field	Bit(s)	Init.	Access Type	Description
INT_ALLOC_0	6:0	X	RW	Defines the MSI-X vector assigned to the TCP timer interrupt cause. The value must be in the 0-63 range.
INT_ALLOC_VAL_0	7	0b	RW	Valid bit for INT_Alloc[0].
INT_ALLOC_1	14:8	X	RW	Defines the MSI-X vector assigned to the other interrupt cause. The value must be in the 0-63 range.
INT_ALLOC_VAL_1	15	1b	RW	Valid bit for INT_Alloc[1].
RESERVED	31:16	0x0	RSV	Reserved.

Note: The INT_ALLOC_VAL[1] bit default value is one EmDash to enable legacy driver functionality.

8.2.2.4.37 RSC Enable Interrupt - RSCINT[n] (0x00012000 + 0x4*n, n=0...128)

Field	Bit(s)	Init.	Access Type	Description
RSCEN	0	1b	RW	RSC Enable. This bit enables RSC on the receive queues associated with interrupt vector n.
RESERVED	31:2	0x0	RSV	Reserved.

8.2.2.4.38 General Purpose Interrupt Enable - GPIE (0x0000898)

Field	Bit(s)	Init.	Access Type	Description
RESERVED	0	0b	RSV	Reserved.
SDP0_GPIEN	1	0b	RW	General Purpose Interrupt Detection Enable for SDP0. If software-controllable I/O pin SDP0 is configured as an input, this bit (when 1b) enables use for GPI interrupt detection. Resulting in setting EICR[25] bit.
SDP1_GPIEN	2	0b	RW	General Purpose Interrupt Detection Enable for SDP1. If software-controllable I/O pin SDP1 is configured as an input, this bit (when 1b) enables use for GPI interrupt detection. Resulting in setting EICR[26] bit.



Field	Bit(s)	Init.	Access Type	Description
SDP2_GPIEN	3	0b	RW	General Purpose Interrupt Detection Enable for SDP2. If software-controllable I/O pin SDP2 is configured as an input, this bit (when 1b) enables use for GPI interrupt detection. Resulting in setting EICR[27] bit.
MULTIPLE_MSIX	4	0b	RW	MSI-X Mode. Selects between MSI-X interrupts and other interrupt modes. 0b= Legacy and MSI mode (non-MSI-X mode). 1b= MSI-X mode.
OCD	5	0b	RW	Other Clear Disable. When set indicates that only bits 16...29 of the EICR are cleared on read. When cleared, the whole EICR is cleared on read.
EIMEN	6	0b	RW	EICS Immediate Interrupt Enable. When set, setting this bit in the EICS causes a LLI. If not set, the EICS interrupt waits for EITR expiration.
RSC_DELAY	13:11	0x0	RW	Delay from RSC completion triggered by ITR and interrupt assertion. The delay equals = (RSC Delay + 1) x 4 μs = 4, 8, 12... 32 μs in 10 GbE or 1 GbE and (RSC Delay + 1) x 40 μs = 40, 80, 120... 320 μs in 100 Mb/s mode.
VT_MODE	15:14	0x0	RW	Specify the number of active virtual functions. Software must set this field the same as GCR_Ext.VT_Mode. 00= Non IOV mode. 01 = 16 VF mode. 10 = 32 VF mode. 11 = 64 VF mode.
RESERVED	29:17	0x0	RSV	Reserved.
EIAME	30	0b	RW	Extended Interrupt Auto Mask Enable. When set, the EIMS register can be auto-cleared (depending on EIAM setting) upon interrupt assertion. In any case, the EIMS register can be auto-cleared (depending on EIAM setting) following a write-to-clear (or read) to the EICR register. Software might set the EIAME only in MSI-X mode.
PBA_SUPPORT	31	0b	RW	PBA Support. When set, setting one of the extended interrupts masks via EIMS causes the PBA bit of the associated MSI-X vector to be cleared. Otherwise, the Intel® Xeon® Processor D-1500 Product Family LAN controller behaves in a way supporting legacy INT-x interrupts. Note: Should be cleared when working in INT-x or MSI mode and set in MSI-X mode.

8.2.2.5 MSI-X Table Registers

The MSI-X capability is described in [Section 9.2.3.3](#). The MSI-X table is described in [Section 9.3.2.1.2](#) and the Pending Bit Array (PBA) is described in [Section 9.3.2.1.3](#). These registers are located in the MSI-X BAR.

8.2.2.5.1 VF MSI-X PBA Clear - VFPBACL[n] (0x000110C8 + 0x4*n, n=0...5)

Field	Bit(s)	Init.	Access Type	Description
PENBITCLR	31:0	0x0	RW	MSI-X Pending Bits Clear. Writing 1b to any bit clears it's content; writing 0b has no effect. Reading this register returns the MSIPBA.PENBIT value.

Note: These registers reflect the VFPBACL bits of the VFs. The PBA is a vector of 192 bits. The vector starts at bit 31 of register p=5 and ends at bit 0 of register p=0 (reverse ordering). Each VF has 3 bits in this vector while, PENBIT[2...0] of



VF=vi are mapped to bits $vi * 3 \dots vi * 3 + 2$. Explicitly, PENBIT[2] of VF0 is at bit 31 of register p=5 and so on...

8.2.2.5.2 MSI-X PBA Clear - PBAACL[n] (0x000110C0 + 0x4*n, n=0...1)

PBAACL[0] is also mapped to address 0x11068 to maintain compatibility with previous products.

Field	Bit(s)	Init.	Access Type	Description
PENBITCLR	31:0	0x0	RW	MSI-X Pending Bits Clear. Writing 1b to any bit clears it's content; writing 0b has no effect. Reading this register returns the MSIPBA.PENBIT value.

8.2.2.6 Receive Registers

8.2.2.6.1 Filter Control Register - FCTRL (0x00005080)

Field	Bit(s)	Init.	Access Type	Description
RESERVED	0	0b	RSV	Reserved.
SBP	1	0b	RW	Store bad packets. 0b = Do not store. 1b = store. Note that CRC errors before the SFD are ignored. Any packet must have a valid SFD (RX_DV with no RX_ER in the XGMII/GMII i/f) in order to be recognized by the device (even bad packets). Note: Packets with errors are not routed to manageability even if this bit is set. Note: Erroneous packets might be routed to the default queue rather than the originally intended queue. Note: In packets shorter than 64 bytes, the checksum errors might be hidden while MAC errors are reported. Note: Packet with a valid error (caused by byte error or illegal error) might have data contamination in the last 8 bytes when stored in the host memory if the store-bad-packet bit is set.
RESERVED	6:2	0x0	RSV	Reserved.
TPE	7	0b	RW	Tag Promiscuous Enable. When set, any packet with active tag is accepted. The active tag to accept is defined by the PFVTCTL.POOLING_MODE field: PFVTCTL.POOLING_MODE= 01b - accept all packets with E-tag. Other = Ignore this bit.
MPE	8	0b	RW	Multicast Promiscuous Enable. 0b = disabled. 1b = enabled. When set, all received multicast and broadcast packets pass L2 filtering and can be directed to manageability or host by a one of the decision filters.
UPE	9	0b	RW	Unicast Promiscuous Enable. 0b = disabled; 1b = enabled.
BAM	10	0b	RW	Broadcast Accept Mode. 0b = Ignore broadcast packets to host. 1b = Accept broadcast packets to host.
RESERVED	31:12	0x0	RSV	Reserved.



Note: Before receive filters are updated/modified, the *RXCTRL.RXEN* bit should be set to zero. After the proper filters have been set the *RXCTRL.RXEN* bit can be set to 1b to re-enable the receiver.

8.2.2.6.2 VLAN Control Register - VLNCTRL (0x00005088)

Field	Bit(s)	Init.	Access Type	Description
VET	15:0	0x8100	RW	VLAN Ether Type (the VLAN Tag Protocol Identifier - TPID). This register contains the type field hardware that's matched against to recognize an 802.1Q (VLAN) Ethernet packet. For proper operation, software must not change the default setting of this field (802.3ac standard defines it as 0x8100). This field must be set to the same value as the <i>V7</i> field in the <i>DMATXCTL</i> register. Note: This field appears in little endian order (the upper byte is first on the wire (VLNCTRL.VET[15:8])).
RESERVED	26:16	0x0	RSV	Reserved.
UP_FIRST_TAG_EN	27	1b	RW	If set, the UP used for TC determination is taken from the first tag with UP; otherwise, it is taken from the inner VLAN.
DEI	28	0b	RW	Drop Eligible Indicator Bit Value. If <i>DEIEN</i> is set to 1b, then .1q packets with DEI equal to this field are accepted; otherwise, the .1q packet is discarded.
DEIEN	29	0b	RW	Drop Eligible Indicator Enable. 0b = Disabled (DEI bit not compared to decide packet acceptance). 1b = Enabled (DEI from packet must match the DEI field to accept a .1q packet).
VFE	30	0b	RW	VLAN Filter Enable. 0b = Disabled (VFTA filter table does not decide packet acceptance). 1b = Enabled (VFTA filter table decides packet acceptance for .1q packets).
RESERVED	31	0b	RSV	Reserved.

8.2.2.6.3 Multicast Control Register - MCSTCTRL (0x00005090)

Field	Bit(s)	Init.	Access Type	Description
MO	1:0	0x0	RW	Multicast Offset. This determines which bits of the incoming multicast address are used in looking up the bit vector. 00b = [47:36]. 01b = [46:35]. 10b = [45:34]. 11b = [43:32].
MFE	2	0b	RW	Multicast Filter Enable. 0b = the multicast table array filter (MTA[n]) is disabled. 1b = the multicast table array (MTA[n]) is enabled.
RESERVED	31:3	0x0	RSV	Reserved.

8.2.2.6.4 Packet Split Receive Type Register - PSRTYPE[n] (0x0000EA00 + 0x4*n, n=0...63)

Notes:

- This register must be initialized by software.



- Packets are split according to the lowest-indexed entry that applies to the packet and that is enabled. For example, if bits 4 and 8 are set, then an IPv4 packet that is not TCP is split after the IPv4 header.
- Exception to this rule is for tunnel packets. If PSR_TYPE15 (split on tunnel) or PSR_TYPE16 (split on outer L2 header) is set, they take precedence on the other filters. At most one of those should be set.
- This bit mask table enables or disables each type of header to be split. A value of 1b enables an entry.
- In virtualization mode, a separate PSRTYPE register is provided per pool up to the number of pools enabled. In non-virtualization mode, only PSRTYPE[0] is used.

Field	Bit(s)	Init.	Access Type	Description
PSR_TYPE0	0	0b	RW	Reserved.
PSR_TYPE1	1	0b	RW	Split received NFS packets after NFS header.
PSR_TYPE2	2	0b	RW	Reserved.
PSR_TYPE3	3	0b	RW	Reserved.
PSR_TYPE4	4	0b	RW	Split received TCP packets after TCP header.
PSR_TYPE5	5	0b	RW	Split received UDP packets after UDP header.
PSR_TYPE6	6	0b	RW	Reserved.
PSR_TYPE7	7	0b	RW	Reserved.
PSR_TYPE8	8	0b	RW	Split received IPv4 packets after IPv4 header.
PSR_TYPE9	9	0b	RW	Split received IPv6 packets after IPv6 header.
PSR_TYPE10	10	0b	RW	Reserved.
PSR_TYPE11	11	0b	RW	Reserved.
PSR_TYPE12	12	0b	RW	Split received L2 packets after L2 header.
PSR_TYPE13	13	0b	RW	Reserved.
PSR_TYPE14	14	0b	RW	Reserved.
PSR_TYPE15	15	0b	RW	Split on tunnel header.
PSR_TYPE16	16	0b	RW	Split on outer L2 header.
PSR_TYPE17	17	0b	RW	Reserved.
PSR_TYPE18	18	0b	RW	Reserved.
RESERVED	28:19	X	RSV	Reserved.
RQPL	31:29	X	RW	Defines the number of bits to use for RSS redirection of packets dedicated to this pool. Valid values are zero, 0001b and 0010b. The default value should be 0010b, meaning that up to 4 queues can be enabled for this pool. A value of 0001b means that up to 2 queues can be enabled for this pool. A value of zero means that all the traffic of the pool is sent to queue 0 of the pool. This field is used only if MRQC.MRQE equals 1001b, 1010b, 1011b, 1110b or 1111b.

Note: Additional address(es): 0x05480 + 4*n, n=0... 15.

8.2.2.6.5 Receive Checksum Control - RXCSUM (0x00005000)

The Receive Checksum Control register controls the receive checksum off loading features of the Intel® Xeon® Processor D-1500 Product Family LAN controller.



Field	Bit(s)	Init.	Access Type	Description
RESERVED	9:0	0x00	RSV	Reserved.
ICMPV6XSUM	10	1b	RW	ICMPv6 Checksum Enable. 0b = Disable ICMPv6 checksum calculation. 1b = Enable ICMPv6 checksum calculation. Note: ICMPv6 checksum offload is supported only for packets sent to firmware for proxying.
RESERVED	11	0b	RSV	Reserved.
IPPCSE	12	0b	RW	IP Payload Checksum Enable.
PCSD	13	0b	RW	RSS/Fragment Checksum Status Selection. The fragment checksum and IP identification fields are mutually exclusive with the RSS hash. Only one of the two options is reported in the Rx descriptor. When set to 1b, the extended descriptor write back has the RSS field. When set to 0b, it contains the fragment checksum.
RESERVED	31:14	0x0	RSV	Reserved.

Note: This register should only be initialized (written) when the receiver is not enabled (for example, only write this register when RXCTRL.RXEN = 0b).

8.2.2.6.6 Receive Filter Control Register - RFCTL (0x00005008)

Field	Bit(s)	Init.	Access Type	Description
RESERVED	4:0	0x0	RSV	Reserved.
RESERVED	5	0b	RW	Reserved.
NFSW_DIS	6	0b	RW	NFS Write disable. Disable filtering of NFS write request headers.
NFSR_DIS	7	0b	RW	NFS Read disable. Disable filtering of NFS read reply headers.
NFS_VER	9:8	0x0	RW	NFS Version recognized by hardware. 00b= NFS version 2. 01b= NFS version 3. 10b= NFS version 4. 11b= Reserved for future use.
IPV6_DIS	10	0b	RW	IPv6 Disable. Disable IPv6 packet filtering. Should not be set to 1b.
IP6XSUM_DISS	11	0b	RW	IPv6 Xsum Disable Disable XSUM on IPv6 packets. Should not be set to 1b.
RESERVED	13:12	0x0	RSV	Reserved.
IPFRSP_DIS	14	0b	RW	IP Fragment Split Disable. When this bit is set, the header of IP fragmented packets are not set. Should not be set to 1b.
RESERVED	15	0b	RSV	Reserved.



Field	Bit(s)	Init.	Access Type	Description
IPV6_EXDIS	16	0b	RW	IPv6 Extension Header Disable. Chicken bit to disable the IPv6 extension headers parsing for XSUM offload, header split and filtering. 0b = Parse and recognize allowed IPV6 extension headers (hop-by-hop, destination options, and routing). 1b = Do not recognize above extension headers. Should not be set to 1b.
RESERVED	17	0b	RSV	Reserved.
RESERVED	31:18	0x0	RSV	Reserved. Should be written with 0x0 to ensure future compatibility.

8.2.2.6.7 Multicast Table Array - MTA[n] (0x00005200 + 0x4*n, n=0...127)

This table should be initialized by software before transmit and receive are enabled.

Field	Bit(s)	Init.	Access Type	Description
BIT_VECTOR	31:0	X	RW	Word wide bit vector specifying 32 bits in the multicast address filter table. The Intel® Xeon® Processor D-1500 Product Family LAN controller provides multicast filtering for 4096 multicast addresses by providing single bit entry per multicast address. Those 4096 address locations organized in multicast table array 128 registers of 32 bits each one. Only 12 bits out of 48 bit destination addresses are considered as multicast address. Those 12 bits can be selected by the <i>MO</i> field of MCSTCTRL register. The 7 MSbs of the Ethernet MAC address (out of the 12 bits) selects the register index while the 5 LSbits (out of the 12 bits) selects the bit within a register.

8.2.2.6.8 Receive Address Low - RAL[n] (0x0000A200 + 0x8*n, n=0...127)

These registers contain the lower bits of the 48-bit Ethernet MAC address. All 32 bits are valid. If the NVM is present the first register (RAL0) is loaded from the NVM.

Field	Bit(s)	Init.	Access Type	Description
RAL	31:0	X	RW	Receive Address Low. The lower 32 bits of the 48-bit Ethernet MAC address. Note: Field is defined in big endian (LSB of RAL is first on the wire).

Note: The first 16 MAC addresses are also mapped to CSR addresses 0x5400 to 0x5478 for backward compatibility.

8.2.2.6.9 Receive Address High - RAH[n] (0x0000A204 + 0x8*n, n=0...127)

These registers contain the upper bits of the 48-bit Ethernet MAC address. The complete address is (RAH, RAL). *AV* determines whether this address is compared against the incoming packet. The *AV* field is cleared by a master reset.



Field	Bit(s)	Init.	Access Type	Description
RAH	15:0	X	RW	Receive Address High. The upper 16 bits of the 48-bit Ethernet MAC address. Note: Field is defined in big endian (MSB of RAH is last on the wire).
RESERVED	29:16	0x0	RSV	Reserved.
ADTYPE	30	0b	RW	Address Type. 0b= MAC. 1b= E-tag. Fixed to zero in RAH[0].
AV	31	0b	RW	Address Valid. All receive addresses are not initialized by hardware and software should initialize them before receive is enabled. If the NVM is present, the receive address[0] is loaded from the NVM and its <i>Address Valid</i> field is set to 1b after a software or PCI reset or NVM read.

Note: The first receive address register (RAR0) is also used for exact match pause frame checking (DA matches the first register). RAR0 should always be used to store the individual Ethernet MAC address of the Intel® Xeon® Processor D-1500 Product Family LAN controller. After reset, if the NVM is present, the first register (Receive Address register 0) is loaded from the IA field in the NVM, its *Address Select* field is 00b, and its *Address Valid* field is 1b. If no NVM is present, the *Address Valid* field is 0b. The *Address Valid* field for all of the other registers are 0b.
The first 16 MAC addresses are also mapped to CSR addresses 0x5404 to 0x547C for backward compatibility.

8.2.2.6.10 Receive Address Low - RAL_ALIAS[n] (0x00005400 + 0x8*n, n=0...15; RW)

These registers are aliases to the first 16 RAL addresses.

Fields definitions are the same as defined on [Section 8.2.2.6.8](#).

8.2.2.6.11 Receive Address High - RAH_ALIAS[n] (0x00005404 + 0x8*n, n=0...15; RW)

These registers are aliases to the first 16 RAH addresses.

Fields definitions are the same as defined on [Section 8.2.2.6.9](#).

8.2.2.6.12 MAC Pool Select Array - MPSAR[n] (0x0000A600 + 0x4*n, n=0...255)

Software should initialize these registers before transmit and receive are enabled.

Field	Bit(s)	Init.	Access Type	Description
POOL_ENA	31:0	X	RW	Pool Enable Bit Array. Each couple of registers 2*n and 2*n+1 are associated with Ethernet MAC address filter n as defined by RAL[n] and RAH[n]. Each bit when set, enables packet reception with the associated pools as follows: Bit i in register 2*n is associated with pool i. Bit i in register 2*n+1 is associated with pool 32+i.

8.2.2.6.13 VLAN Filter Table Array - VFTA[n] (0x0000A000 + 0x4*n, n=0...127)

This table should be initialized by software before transmit and receive are enabled.



Field	Bit(s)	Init.	Access Type	Description
VLAN_FLT	31:0	X	RW	VLAN Filter. Each bit i in register n affects packets with VLAN VID equal to 32*n+i. 128 VLAN Filter registers compose a table of 4096 bits that cover all possible VIDs. Each bit when set, allows packets with the associated VID to pass. Each bit when cleared, blocks packets with this VID.

8.2.2.6.14 Multiple Receive Queues Command Register - MRQC (0x0000EC80)

Field	Bit(s)	Init.	Access Type	Description
MRQE	3:0	0x0	RW	Multiple Receive Queues Enable. Defines allocation of the Rx queues per RSS, virtualization. 0000b = RSS disabled. 0001b = RSS only - Single set of RSS 64 queues. 0010b = RSS disabled - 8 TCs, each allocated 1 queue. 0011b = RSS disabled - 4 TCs, each allocated 1 queue. 0100b = RSS - 8 TCs, each allocated 16 RSS queues. 0101b = RSS - 4 TCs, each allocated 32 RSS queues. 0110b = Reserved. 0111b = Reserved. 1000b = Virtualization only - 64 pools, no RSS, each pool allocated 2 queues. 1001b = Virtualization and RSS - 62 pools, each allocated 2 RSS queues - 1 pool allocated 4 queues. 1010b = Virtualization and RSS - 32 pools, each allocated 4 RSS queues. 1011b = Virtualization and RSS - 64 pools, each allocated 2 RSS queues. 1100b = Virtualization - 16 pools, each allocated 8 TCs. 1101b = Virtualization - 32 pools, each allocated 4 TCs. 1110b = Virtualization and RSS - 16 pools, each allocated 4 TCs, each allocated 2 RSS queues. 1111b = Virtualization and RSS - 60 pools, each allocated 2 RSS queues - 2 pools, each allocated 4 queues.
RESERVED	12:4	0x0	RSV	Reserved.
MULTIPLE_RS S	13	0b	RW	Multiple RSS Enable. When set to 1b, the Intel® Xeon® Processor D-1500 Product Family LAN controller uses up to 64 RSS keys (one per pool) each RSS has a redirection table with 64 entry of 2 bits each. When set to 0b, the Intel® Xeon® Processor D-1500 Product Family LAN controller uses a single RSS key (legacy). This bit is relevant only if MRQC.MRQE = 1001b, 1010b, 1011b, 1110b, or 1111b.
RSC_DIS_LP	14	0b	RW	RSC Disable LLI and Push Coalescing. If set, RSC does not coalesce packets with LLI or PSH.
RESERVED	15	0b	RW	Reserved.



Field	Bit(s)	Init.	Access Type	Description
RSS_FIELD_ENABLE	31:16	0x0	RW	Each bit, when set, enables a specific field selection to be used by the hash function. Several bits can be set at the same time. Bit[16] = Enable TcpIPv4 hash function. Bit[17] = Enable IPv4 hash function. Bit[18] = Reserved. Bit[19] = Reserved. Bit[20] = Enable IPv6 hash function. Bit[21] = Enable TcpIPv6 hash function. Bit[22] = Enable UdpIPv4. Bit[23] = Enable UdpIPv6. Bit[24] = Reserved. Bits[31:25] = Reserved (zero). Note that on tunnel packets IPv4/IPv6 only the IPv4 header might be used for the RSS filtering.

8.2.2.6.15 Multiple Receive Queues Command Register - MRQC_ALIAS (0x00005818; RW)

This is an alias to the MRQC register.

Fields definitions are the same as defined on [Section 8.2.2.6.14](#).

8.2.2.6.16 RSS Queues Per Traffic Class Register - RQTC (0x0000EC70)

Field	Bit(s)	Init.	Access Type	Description
RQTC0	2:0	0x4	RW	Defines the number of bits to use for RSS redirection of packets dedicated to TC 0. A value of zero means that all the traffic of TC 0 is sent to queue 0 of the TC. This field is used only if MRQC.Multiple Receive Queues Enable = 0100b or 0101b. A value of seven in this field is not valid and is considered as if zero is written.
RESERVED	3	0b	RSV	Reserved.
RQTC1	6:4	0x4	RW	Defines the number of bits to use for RSS redirection of packets dedicated to TC 1. A value of zero means that all the traffic of TC 1 is sent to queue 0 of the TC. This field is used only if MRQC.Multiple Receive Queues Enable = 0100b or 0101b. A value of seven in this field is not valid and is considered as if zero is written.
RESERVED	7	0b	RSV	Reserved.
RQTC2	10:8	0x4	RW	Defines the number of bits to use for RSS redirection of packets dedicated to TC 2. A value of zero means that all the traffic of TC 2 is sent to queue 0 of the TC. This field is used only if MRQC.Multiple Receive Queues Enable = 0100b or 0101b. A value of seven in this field is not valid and is will be considered as if zero is written.
RESERVED	11	0b	RSV	Reserved.
RQTC3	14:12	0x4	RW	Defines the number of bits to use for RSS redirection of packets dedicated to TC 3. A value of zero means that all the traffic of TC 3 is sent to queue 0 of the TC. This field is used only if MRQC.Multiple Receive Queues Enable = 0100b or 0101b. A value of seven in this field is not valid and is considered as if zero is written.
RESERVED	15	0b	RSV	Reserved.



Field	Bit(s)	Init.	Access Type	Description
RQTC4	18:16	0x4	RW	Defines the number of bits to use for RSS redirection of packets dedicated to TC 4. A value of zero means that all the traffic of TC 4 is sent to queue 0 of the TC. This field is used only if MRQC.Multiple Receive Queues Enable = 0100b. A value of seven in this field is not valid and is will be considered as if zero is written.
RESERVED	19	0b	RSV	Reserved.
RQTC5	22:20	0x4	RW	Defines the number of bits to use for RSS redirection of packets dedicated to TC 5. A value of zero means that all the traffic of TC 5 is sent to queue 0 of the TC. This field is used only if MRQC.Multiple Receive Queues Enable = 0100b. A value of seven in this field is not valid and is will be considered as if zero is written.
RESERVED	23	0b	RSV	Reserved.
RQTC6	26:24	0x4	RW	Defines the number of bits to use for RSS redirection of packets dedicated to TC 6. A value of zero means that all the traffic of TC 6 is sent to queue 0 of the TC. This field is used only if MRQC.Multiple Receive Queues Enable = 0100b. A value of seven in this field is not valid and is will be considered as if zero is written.
RESERVED	27	0b	RSV	Reserved.
RQTC7	30:28	0x4	RW	Defines the number of bits to use for RSS redirection of packets dedicated to TC 7. A value of zero means that all the traffic of TC 7 is sent to queue 0 of the TC. This field is used only if MRQC.Multiple Receive Queues Enable = 0100b. A value of seven in this field is not valid and is will be considered as if zero is written.
RESERVED	31	0b	RSV	Reserved.

8.2.2.6.17 RSS Random Key Register - RSSRK[n] (0x0000EB80 + 0x4*n, n=0...9)

The RSS random key is 40 bytes wide (see RSS hash in [Section 7.1.3.6.3](#)).

Field	Bit(s)	Init.	Access Type	Description
K0	7:0	0x0	RW	RSS Key Byte 4*n+0 of the RSS random key, for each register n.
K1	15:8	0x0	RW	RSS key byte 4*n+1 of the RSS random key, for each register n.
K2	23:16	0x0	RW	RSS key byte 4*n+2 of the RSS random key, for each register n.
K3	31:24	0x0	RW	RSS key byte 4*n+3 of the RSS random key, for each register n.

Note: Additional address(es): 0x05C80 + 4*n, n=0...9.

8.2.2.6.18 Redirection Table - RETA[n] (0x0000EB00 + 0x4*n, n=0...31)

he redirection table has 128-entries in 32 registers.

Field	Bit(s)	Init.	Access Type	Description
ENTRY0	5:0	X	RW	Entry0 defines the RSS output index for hash value 4 4*n+0. While n is the register index, equals to 0...31.
RESERVED	7:6	0x0	RSV	Reserved.



Field	Bit(s)	Init.	Access Type	Description
ENTRY1	13:8	X	RW	Entry1 defines the RSS output index for hash value $4*n+1$. While n is the register index, equals to 0...31.
RESERVED	15:14	0x0	RSV	Reserved.
ENTRY2	21:16	X	RW	Entry2 defines the RSS output index for hash value $4*n+2$. While n is the register index, equals to 0...31.
RESERVED	23:22	0x0	RSV	Reserved.
ENTRY3	29:24	X	RW	Entry3 defines the RSS output index for hash value $4*n+3$. While n is the register index, equals to 0...31.
RESERVED	31:30	0x0	RSV	Reserved.

Note: The contents of the redirection table are not defined following reset of the Memory Configuration registers. System software must initialize the table prior to enabling multiple receive queues. It might also update the redirection table during run time. Such updates of the table are not synchronized with the arrival time of received packets. Therefore, it is not guaranteed that a table update takes effect on a specific packet boundary.

8.2.2.6.19 Extended Redirection Table - ERETA[n] (0x0000EE80 + 0x4*n, n=0...95)

The extended redirection table adds 384-entries to extend RETA in 96 registers.

Field	Bit(s)	Init.	Access Type	Description
ENTRY0	5:0	X	RW	Entry0 defines the RSS output index for hash value $4*[n+32]$. While n is the register index, equals to 0...95.
RESERVED	7:6	0x0	RSV	Reserved.
ENTRY1	13:8	X	RW	Entry1 defines the RSS output index for hash value $4*[n+32]+1$. While n is the register index, equals to 0...95.
RESERVED	15:14	0x0	RSV	Reserved.
ENTRY2	21:16	X	RW	Entry2 defines the RSS output index for hash value $4*[n+32]+2$. While n is the register index, equals to 0...95.
RESERVED	23:22	0x0	RSV	Reserved.
ENTRY3	29:24	X	RW	Entry3 defines the RSS output index for hash value $4*[n+32]+3$. While n is the register index, equals to 0...95.
RESERVED	31:30	0x0	RSV	Reserved.

Note: The contents of the extended redirection table are not defined following reset of the Memory Configuration registers. System software must initialize the table prior to enabling multiple receive queues. It might also update the redirection table during run time. Such updates of the table are not synchronized with the arrival time of received packets. Therefore, it is not guaranteed that a table update takes effect on a specific packet boundary.

8.2.2.6.20 Per Pool RSS Random Key Register - VFRSSRK[n,m] (0x00018000 + 0x4*n + 0x40*m, n=0...15, m=0...63)

The RSS Random Key is 40 bytes wide.

Field	Bit(s)	Init.	Access Type	Description
K0	7:0	0x0	RW	RSS Key Byte $4*n+0$ of the RSS random key, for each register n.



Field	Bit(s)	Init.	Access Type	Description
K1	15:8	0x0	RW	RSS Key Byte $4*n+1$ of the RSS random key, for each register n.
K2	23:16	0x0	RW	RSS Key Byte $4*n+2$ of the RSS random key, for each register n.
K3	31:24	0x0	RW	RSS Key Byte $4*n+3$ of the RSS random key, for each register n.

Notes: Only the 10 first registers in each VF array are implemented (0x00 to 0x24).

8.2.2.6.21 Per Pool Redirection Table - VFRETA[n,m] (0x00019000 + 0x4*n + 0x40*m, n=0...15, m=0...63)

The redirection table has 64 entries in 16 registers.

Field	Bit(s)	Init.	Access Type	Description
ENTRY0	1:0	X	RW	Entry0 defines the RSS output index for hash value $4*4*n+0$. While n is the register index, equals to 0...31.
RESERVED	7:2	0x0	RSV	Reserved.
ENTRY1	9:8	X	RW	Entry1 defines the RSS output index for hash value $4*4*n+1$. While n is the register index, equals to 0...31.
RESERVED	15:10	0x0	RSV	Reserved.
ENTRY2	17:16	X	RW	Entry2 defines the RSS output index for hash value $4*4*n+2$. While n is the register index, equals to 0...31.
RESERVED	23:18	0x0	RSV	Reserved.
ENTRY3	25:24	X	RW	Entry3 defines the RSS output index for hash value $4*4*n+3$. While n is the register index, equals to 0...31.
RESERVED	31:26	0x0	RSV	Reserved.

Note: The content of the redirection tables is not defined following reset of the Memory Configuration registers.

Note: System software must initialize the table prior to enabling multiple receive queues. It might also update the redirection table during run time. Such updates of the table are not synchronized with the arrival time of received packets. Therefore, it is not guaranteed that a table update takes effect on a specific packet boundary.

8.2.2.6.22 SYN Packet Queue Filter - SYNQF (0x0000EC30)

Field	Bit(s)	Init.	Access Type	Description
QUEUE_ENABLE	0	0b	RW	When set, enables routing of Rx packets to the queue indicated in this register.
RX_QUEUE	7:1	0x00	RW	Identifies an Rx queue associated with SYN packets.
RESERVED	31:9	0x0	RSV	Reserved.



8.2.2.6.23 EType Queue Filter - ETQF[n] (0x00005128 + 0x4*n, n=0...7)

Field	Bit(s)	Init.	Access Type	Description
ETYPE	15:0	0x0000	RW	Identifies the protocol running on top of IEEE 802. Used to route Rx packets containing this EtherType to a specific Rx queue. Note: Field is defined in little endian (MSB is first on the wire).
POOL	25:20	0x00	RW	In virtualization modes, determines the target pool for the packet.
POOL_ENABLE	26	0b	RW	In virtualization modes, enables the <i>Pool</i> field.
RESERVED	27	0b	RW	Reserved.
CNM	28	0b	RW	When set, a packet with this EType is identified as a CNM packet.
TX_ANTISPOOF	29	0b	RSV	If set, this EtherType is candidate for anti-spoof action on transmitted packets. The actual action applied is set according to PFVFSPOOF.ETHERTYPEAS or PFVFSPOOF.ETHERTYPELB.
IEEE_1588_TIMESTAMP	30	0b	RW	When set, packets with this EType are time stamped according to the IEEE 1588 specification.
FILTER_ENABLE	31	0b	RW	0b = The filter is disabled for any functionality. 1b = The filter is enabled. Exact actions are determined by separate bits.

8.2.2.6.24 EType Queue Select - ETQS[n] (0x0000EC00 + 0x4*n, n=0...7)

Field	Bit(s)	Init.	Access Type	Description
RESERVED	15:0	0x00	RSV	Reserved.
RX_QUEUE	22:16	0x00	RW	Identifies the Rx queue associated with this EType.
RESERVED	24:23	0x00	RSV	Reserved. for future extension of the Rx Queue field.
RESERVED	28:25	0x00	RSV	Reserved.
RESERVED	30	0b	RSV	Reserved.
QUEUE_ENABLE	31	0b	RW	When set, enables queueing of Rx packets by the EType defined in the matching ETQF register to the queue indicated in this register.

8.2.2.6.25 E-tag EtherType Register - ETAG_ETYPE (0x00005084)

Field	Bit(s)	Init.	Access Type	Description
ETAG_ETYPE	15:0	0x893F	RW	Etag EtherType Value.
RESERVED	30:16	0x0	RSV	Reserved.
VALID	31	0b	RW	Valid Bit. 0b = Entry is disabled. 1b = Entry is enabled.



8.2.2.6.26 VXLAN Control - VXLANCTRL (0x0000507C)

Field	Bit(s)	Init.	Access Type	Description
UDPPORT	15:0	0x0	RW	Defines the UDP port used to identify VXLAN traffic.
RESERVED	31:16	0x0	RSV	Reserved.

8.2.2.6.27 Receive Descriptor Base Address Low - RDBAL[n] (0x00001000 + 0x40*n, n=0...63 and 0x0000D000 + 0x40*(n-64), n=64...127)

This register contains the lower bits of the 64-bit descriptor base address. The lower seven bits are always ignored. The receive descriptor base address must point to a 128-byte aligned block of data.

Field	Bit(s)	Init.	Access Type	Description
ZERO	6:0	0x0	RW	Ignored on writes. Returns 0x0 on reads.
RDBAL	31:7	X	RW	Receive Descriptor Base Address Low.

8.2.2.6.28 Receive Descriptor Base Address High - RDBAH[n] (0x00001004 + 0x40*n, n=0...63 and 0x0000D004 + 0x40*(n-64), n=64...127)

This register contains the upper 32 bits of the 64-bit descriptor base address.

Field	Bit(s)	Init.	Access Type	Description
RDBAH	31:0	X	RW	Receive Descriptor Base Address [63:32].

8.2.2.6.29 Receive Descriptor Length - RDLEN[n] (0x00001008 + 0x40*n, n=0...63 and 0x0000D008 + 0x40*(n-64), n=64...127)

Field	Bit(s)	Init.	Access Type	Description
LEN	19:0	0x0	RW	Descriptor Ring Length. This register sets the number of bytes allocated for descriptors in the circular descriptor buffer. It must be 128-byte aligned (7 LSB must be set to zero). Note: Validated lengths up to 128 K (8 K descriptors).
RESERVED	31:20	0x0	RSV	Reads as 0. Should be written to 0 for future compatibility.

Notes: Additional address(es): 0x0D008 + 0x40*(n-64), n=64...127.

8.2.2.6.30 Receive Descriptor Head - RDH[n] (0x00001010 + 0x40*n, n=0...63 and 0x0000D010 + 0x40*(n-64), n=64...127)

Field	Bit(s)	Init.	Access Type	Description
RDH	15:0	0x0	RO	Receive Descriptor Head. This register holds the head pointer for the receive descriptor buffer in descriptor units (16-byte datum). The RDH is controlled by hardware.



Field	Bit(s)	Init.	Access Type	Description
RESERVED	31:16	0x0	RSV	Reserved. Should be written with 0x0.

Note: Additional address(es): 0x0D010 + 0x40*(n-64), n=64...127

8.2.2.6.31 Receive Descriptor Tail - RDT[n] (0x00001018 + 0x40*n, n=0...63 and 0x0000D018 + 0x40*(n-64), n=64...127)

This register contains the tail pointer for the receive descriptor buffer. The register points to a 16-byte datum. Software writes the tail register to add receive descriptors to the hardware free list for the ring. Additional address(es): 0x0D018 + 0x40*(n-64), n=64...127.

Field	Bit(s)	Init.	Access Type	Description
RDT	15:0	0x0	RW	Receive Descriptor Tail.
RESERVED	31:16	0x0	RSV	Reads as 0x0. Should be written to 0x0 for future compatibility.

Note: Software should write an even number to the tail register, if the packet split feature is used. The tail pointer should be set to one descriptor beyond the last empty descriptor in host descriptor ring.

8.2.2.6.32 Receive Descriptor Control - RXDCTL[n] (0x00001028 + 0x40*n, n=0...63 and 0x0000D028 + 0x40*(n-64), n=64...127)

Field	Bit(s)	Init.	Access Type	Description
RLPML	13:0	0x0	RW	Long packet size filter defined in bytes. Packets larger than the RLPML are discarded. The filter is enabled by the <i>RLMPL_EN</i> bit in this register. This field might not be supported per Rx queue in future products. Note: The Intel® Xeon® Processor D-1500 Product Family LAN controller closes active RSC flows when an LLI packet was dropped due to RLPML exceeded. Note: The packet size compared to RLPML includes any optional timestamp added into the packet.
RESERVED	14	0b	RSV	Reserved. Software might read and write maintaining backward compatibility.
RLPML_EN	15	0b	RW	Enable Long Packet Size Filter. 1b = Enable.
RESERVED	22:16	0x00	RSV	Reserved. Software might read and write maintaining backward compatibility.
RESERVED	24:23	0x0	RSV	Reserved.
ENABLE	25	0b	RW	Receive Queue Enable. When set, the <i>Enable</i> bit enables the operation of the specific receive queue, upon read, gets the actual status of the queue (internal indication that the queue is actually enabled/disabled).
RESERVED	29:26	0x00	RSV	Reserved.
VME	30	0b	RW	VLAN Mode Enable. 1b = Strip VLAN tag from received 802.1Q packets destined to this queue. 0b = Do not strip VLAN tag.
RESERVED	31	0b	RSV	Reserved



Note: Additional address(es): $0x0D028 + 0x40*(n-64)$, $n=64...127$.

8.2.2.6.33 Split Receive Control Registers - SRRCTL_ALIAS[n] (0x00002100 + 0x4*n, n=0...15; RW)

Fields definitions are the same as defined on Section 8.2.2.6.34.

8.2.2.6.34 Split Receive Control Registers - SRRCTL[n] (0x00001014 + 0x40*n, n=0...63 and 0x0000D014 + 0x40*(n-64), n=64...127)

Field	Bit(s)	Init.	Access Type	Description
BSIZEPACKET	4:0	0x2	RW	Receive Buffer Size for Packet Buffer. The value is in 1 KB resolution. Value can be from 1 KB to 16 KB. The default buffer size is 2 KB. This field should not be set to 0x0. This field should be greater or equal to two in queues which RSC is enabled.
RESERVED	7:5	0x0	RSV	Reserved. Should be written with 0x0 to ensure future compatibility.
BSIZEHEADER	13:8	0x4	RW	Receive Buffer Size for Header Buffer. The value is in 64 byte resolution. Value can be from 64 bytes to 1024 bytes. Note that the max supported header size is limited to 1023. The default buffer size is 256 bytes. This field must be greater than zero if the value of DESCTYPE is greater or equal to two.
RESERVED	21:14	0x0	RSV	Reserved.
RDMTS	24:22	0x0	RW	Receive Descriptor Minimum Threshold Size. An interrupt associated with this queue is asserted each time the number of free descriptors is decreased to $RDMTS * 64$ (this event is considered as Rx ring buffer almost empty).
DESCTYPE	27:25	0x0	RW	Define the descriptor type in Rx000: legacy. 001 = Advanced descriptor one buffer. 010 = Advanced descriptor header splitting. 011b= Reserved. 100b= Reserved. 101b= Advanced descriptor header splitting always use header buffer. 110b - 111 = Reserved.
DROP_EN	28	0b	RW	Drop Enabled. If set to 1b, packets received to the queue when no descriptors are available to store them are dropped.
RESERVED	31:29	0x0	RSV	Reserved. Should be written with 0x0 to ensure future compatibility.

Note: Additional address(es): $0x0D014 + 0x40*(n-64)$, $n=64...127$ / $0x02100 + 4*n$, $[n=0...15]$.BSIZEHEADER must be bigger than zero if DESCTYPE is equal to 010b, 011b 100b or 101b.

8.2.2.6.35 Receive DMA Control Register - RDRXCTL (0x00002F00) DMA_RX

Field	Bit(s)	Init.	Access Type	Description
RESERVED	0	0b	RW	Reserved.
PSP	2	0b	RW	Pad Small Receive Packets. If this field is set in virtualized operating mode, strip CRC (HLREG0.RXCRCSTRP) should be set.
DMAIDONE	3	0b	RO	DMA Init Done. When read as 1b indicates that the DMA init cycle is done (RO).



Field	Bit(s)	Init.	Access Type	Description
RESERVED	4	0b	RSV	Reserved.
RSC_PUSH_D IS	5	0b	RW	Disable coalescing of packets with PUSH flag asserted.
TPH_AUTOLE ARN	8	0b	RSV	TPH Hints Auto Learn If set, the Intel® Xeon® Processor D-1500 Product Family LAN controller learns the CPUID value in this register from the completer ID in read completion of receive descriptors fetch. Note: If TPH is disabled, this bit should not be set, as the learn capability is independent of the TPH capability.
RESERVED	24:7			Reserved.
MBINTEN	28	0b	RW	Malicious behavior interrupt enable for DMA Rx errors.
MDP_EN	29	0b	RW	Malicious driver protection enable for DMA Rx. 0b = Mechanism is disabled. 1b = Mechanism is enabled.
RESERVED	31:30	0x0	RSV	Reserved.

8.2.2.6.36 Receive Packet Buffer Size - RXPBSIZE[n] (0x00003C00 + 0x4*n, n=0...7)

Field	Bit(s)	Init.	Access Type	Description
RESERVED	9:0	0x0	RSV	Reserved.
SIZE	19:10	0x180	RW	Receive packet buffer size for TC n while n is the register index. The size is defined in KB units and the default size of PB[0] is 384 KB. The default size of PB[1-7] is also 384 KB. Another possible setting of 4 x TCs is 0x60 (96 KB) for all PB[0-3] and 0x0 for PB[4-7] is allowed. See Section 3.8.3.4.4 (packet buffer size) for other optional settings with/without the flow director filters. Note: In any setting, RXPBSIZE[0] must always be enabled (greater than zero). Default value is 0x180 (384 KB).
RESERVED	31:20	0x0	RSV	Reserved.

8.2.2.6.37 Receive Control Register - RXCTRL (0x00003000)

Field	Bit(s)	Init.	Access Type	Description
RXEN	0	0b	RW	Receive Enable. When set to 0b, packets are not received.
RESERVED	31:1	0x0	RSV	Reserved.

8.2.2.6.38 RSC Control - RSCCTL[n] (0x0000102C + 0x40*n, n=0...63 and 0x0000D02C + 0x40*(n-64), n=64...127)

Field	Bit(s)	Init.	Access Type	Description
RSCEN	0	0b	RW	RSC Enable. When the RSCEN bit is set, RSC is enabled on this queue.



Field	Bit(s)	Init.	Access Type	Description
MAXDESC	3:2	0x0	RW	Maximum descriptors per large receive is as follows: 00b = Maximum of 1 descriptor per large receive. 01b = Maximum of 4 descriptors per large receive. 10b = Maximum of 8 descriptors per large receive. 11b = Maximum of 16 descriptors per large receive. Note: MAXDESC * SRRCTL.BSIZEPKT must not exceed 64 KB minus 1 which is the maximum total length in the IP header and must be larger than the expected received MSS.
RESERVED	31:4	0x00	RSV	Reserved.

Note: Additional address(es): 0x0D02C + 0x40*(n-64), n=64...127.

8.2.2.7 Transmit Registers

8.2.2.7.1 DMA Tx TCP Max Allow Size Requests - DTXMXSZRQ (0x00008100)

This register limits the total number of data bytes that might be in outstanding PCIe requests from the host memory. This enables requests to send low latency packets to be serviced in a timely manner, as this request is serviced right after the current outstanding requests are completed.

Field	Bit(s)	Init.	Access Type	Description
MAX_BYTES_NUM_REQ	11:0	0x10	RW	Max Allowed Number of Bytes Requests. The maximum allowed amount of 256 bytes outstanding requests. If the total size request is higher than the amount in the field no arbitration is done and no new packet is requested.
RESERVED	31:12	0x0	RSV	Reserved.

8.2.2.7.2 DMA Tx Control - DMATXCTL (0x00004A80)

Field	Bit(s)	Init.	Access Type	Description
TE	0	0b	RW	Transmit Enable. When set, this bit enables the transmit operation of the DMA-Tx.
RESERVED	1	0b	RSV	Reserved.
RESERVED	2	1b	RW	Reserved.
GDV	3	0b	RW	Global Double VLAN Mode. When set, this bit enables double VLAN mode. Should be set to the same value as CTRL_EXT.EXTENDED_VLAN.
RESERVED	4	1b	RW	Reserved.
MDP_EN	5	0b	RW	Malicious Driver Protection Enable for DMA Tx. 0b = Mechanism is disabled. 1b = Mechanism is enabled.
MBINTEN	6	0b	RW	Malicious behavior interrupt enable for DMA Tx errors.
TPH_AUTOLEARN	7	0b	RW	TPH Hints Auto Learn. If set, the Intel® Xeon® Processor D-1500 Product Family LAN controller learns the CPUID value in this register from the completer ID in read completion of transmit descriptors fetch. Note: If TPH is disabled, this bit should not be set, as the learn capability is independent of the TPH capability.



Field	Bit(s)	Init.	Access Type	Description
RESERVED	15:8	0x0	RSV	Reserved.
VT	31:16	0x8100	RW	VLAN Ether-Type (the VLAN Tag Protocol Identifier - TPID). For proper operation, software must not change the default setting of this field (802.3ac standard defines it as 0x8100). This field must be set to the same value as the VET field in the VLNCTRL register.

8.2.2.7.3 DMA Tx TCP Flags Control Low - DTXTCPFLGL (0x00004A88)

This register holds the *Mask* bits for the TCP flags in Tx segmentation (described in Section 7.2.4.7.1 and Section 7.2.4.7.2).

Field	Bit(s)	Init.	Access Type	Description
TCP_FLG_FIR ST_SEG	11:0	0xFF6	RW	TCP Flags First Segment. Bits that make AND operation with the TCP flags at TCP header in the first segment.
RESERVED	15:12	0x00	RSV	Reserved.
TCP_FLG_MI D_SEG	27:16	0xFF6	RW	TCP Flags Middle Segments. The low bits that make AND operation with the TCP flags at TCP header in the middle segments.
RESERVED	31:28	0x00	RSV	Reserved.

8.2.2.7.4 DMA Tx TCP Flags Control High - DTXTCPFLGH (0x00004A8C)

This register holds the *Mask* bits for the TCP flags in Tx segmentation (described in Section 7.2.4.7.3).

Field	Bit(s)	Init.	Access Type	Description
TCP_FLG_LST _SEG	11:0	0xF7F	RW	TCP Flags Last Segment. Bits that make AND operation with the TCP flags at TCP header in the last segment.
RESERVED	31:12	0x00	RSV	Reserved.

8.2.2.7.5 Transmit Descriptor Base Address Low - TDBAL[n] (0x00006000 + 0x40*n, n=0...127)

Field	Bit(s)	Init.	Access Type	Description
ZERO	6:0	0x0	RW	Ignored on writes. Returns 0x0 on reads.
TDBAL	31:7	X	RW	Transmit Descriptor Base Address Low.

Note: This register contains the lower bits of the 64-bit descriptor base address. The lower seven bits are ignored. The transmit descriptor base address must point to a 128-byte aligned block of data.



8.2.2.7.6 Transmit Descriptor Base Address High - TDBAH[n] (0x00006004 + 0x40*n, n=0...127)

Field	Bit(s)	Init.	Access Type	Description
TDBAH	31:0	X	RW	Transmit Descriptor Base Address [63:32].

Note: This register contains the upper 32 bits of the 64-bit descriptor base address.

8.2.2.7.7 Transmit Descriptor Length - TDLEN[n] (0x00006008 + 0x40*n, n=0...127)

Field	Bit(s)	Init.	Access Type	Description
LEN	19:0	0x0	RW	Descriptor Ring Length. This register sets the number of bytes allocated for descriptors in the circular descriptor buffer. It must be 128-byte aligned (7 LSB must be set to zero). Note: Validated lengths up to 128 K (8 K descriptors).
RESERVED	31:20	0x0	RSV	Reads as 0x0. Should be written to 0x0.

8.2.2.7.8 Transmit Descriptor Head - TDH[n] (0x00006010 + 0x40*n, n=0...127)

Field	Bit(s)	Init.	Access Type	Description
TDH	15:0	0x0	RO	Transmit Descriptor Head.
RESERVED	31:16	0x0	RSV	Reserved. Should be written with 0x0.

Note: This register contains the head pointer for the transmit descriptor ring. It points to a 16-byte datum. Hardware controls this pointer.

Note: The values in these registers might point to descriptors that are still not in the host memory. As a result, the host cannot rely on these values in order to determine which descriptor to release.

Note: The only time that software should write to this register is after a reset (hardware reset or CTRL.RST) and before enabling the transmit function (TXDCTL.ENABLE). If software were to write to this register while the transmit function was enabled, the on-chip descriptor buffers might be invalidated and hardware might become lost.

8.2.2.7.9 Transmit Descriptor Tail - TDT[n] (0x00006018 + 0x40*n, n=0...127)

Field	Bit(s)	Init.	Access Type	Description
TDT	15:0	0x0	RW	Transmit Descriptor Tail.
RESERVED	31:16	0x0	RSV	Reads as 0x0. Should be written to 0x0 for future compatibility.

Note: This register contains the tail pointer for the transmit descriptor ring. It points to a 16-byte datum. Software writes the tail pointer to add more descriptors to the transmit ready queue. Hardware attempts to transmit all packets referenced by descriptors between head and tail.



8.2.2.7.10 Transmit Descriptor Control - TXDCTL[n] (0x00006028 + 0x40*n, n=0...127)

Field	Bit(s)	Init.	Access Type	Description
PTHRESH	6:0	0x00	RW	Pre-fetch Threshold. Controls when a pre-fetch of descriptors is considered. This threshold refers to the number of valid, unprocessed transmit descriptors are in the on-chip buffer. If this number drops below PTHRESH, the algorithm considers pre-fetching descriptors from host memory. However, this fetch does not happen unless there are at least HTHRESH valid descriptors in host memory to fetch. Note: HTHRESH should be given a non-zero value each time PTHRESH is used.
RESERVED	7	0b	RSV	Reserved.
HTHRESH	14:8	0x00	RW	Host Threshold.
RESERVED	15	0b	RSV	Reserved.
WTHRESH	22:16	0x00	RW	Write-back Threshold. Controls the write back of processed transmit descriptors. This threshold refers to the number of transmit descriptors in the on-chip buffer that are ready to be written back to host memory. In the absence of external events (explicit flushes), the write back occurs only after at least WTHRESH descriptors are available for write back. Note: Since the default value for a write-back threshold is 0b, descriptors are normally written back as soon as they are processed. WTHRESH must be written to a non-zero value to take advantage of the write back bursting capabilities. Note: When WTHRESH is set to a non-zero value, the software device driver should not set the RS bit in the Tx descriptors. When WTHRESH is set to zero, the software device driver should set the RS bit in the last Tx descriptors of every packet (if TSO is the last descriptor of the entire large send). Note: When head write back is enabled (TDWBAL[n].Head_WB_En = 1b), the WTHRESH must be set to zero.
RESERVED	24:23	0x00	RSV	Reserved.
ENABLE	25	0b	RW	Transmit Queue Enable. When set, this bit enables the operation of a specific transmit queue. The default value for all queues is 0b. Setting this bit initializes all the internal registers of a specific queue. Until then, the state of the queue is kept and can be used for debug purposes. When disabling a queue, this bit is cleared only after all activity at the queue stopped. Note: When setting the global Tx enable DMATXCTL.TE, the <i>Enable</i> bit of Tx queue zero is enabled as well. Note: Following a write cycle by software, this bit is set by hardware only when the queue is actually enabled. When read, it gets the actual status of the queue (internal indication that the queue is actually enabled/disabled).
SWFLSH	26	0b	RW	Transmit Software Flush. This bit enables software to trigger descriptor write-back flushing, independent of other conditions. This bit is self cleared by hardware.
RESERVED	27	0b	RSV	Reserved (was priority).
RESERVED	28	0b	RW	Reserved.
RESERVED	29	0b	RSV	Reserved.
RESERVED	31:30	0x00	RSV	Reserved.

Note: This register controls the fetching and write back of transmit descriptors. The three threshold values are used to determine when descriptors are read from and written to host memory.



Note: When WTHRESH = 0b, only descriptors with the RS bit set are written back. For PTHRESH and HTHRESH recommended settings, see Section 7.2.3.4.

8.2.2.7.11 Tx Descriptor Completion Write Back Address Low - TDWBAL[n] (0x00006038 + 0x40*n, n=0...127)

Field	Bit(s)	Init.	Access Type	Description
HEAD_WB_EN	0	0b	RW	Head Write-Back Enable. 1b= Head write back is enabled. 0b= Head write back is disabled. When head_WB_en is set, Tx descriptors are not written back.
RESERVED	1	0b	RW	Reserved.
HEADWB_LOW	31:2	0x0	RW	Lowest 32 bits of the head write-back memory location (DWord aligned). The last two bits of this field are ignored and are always read as 0.0, meaning that the actual address is QWord aligned.

8.2.2.7.12 Tx Descriptor Completion Write Back Address High - TDWBAH[n] (0x0000603C + 0x40*n, n=0...127) DMA_TX

Field	Bit(s)	Init.	Access Type	Description
HEADWB_HIGH	31:0	0x0	RW	Highest 32 bits of head write-back memory location (for 64-bit addressing).

8.2.2.7.13 Transmit Packet Buffer Size - TXPBSIZE[n] (0x0000CC00 + 0x4*n, n=0...7) DBU_TX

Field	Bit(s)	Init.	Access Type	Description
RESERVED	9:0	0x0	RSV	Reserved.
SIZE	19:10	0xA0	RW	Transmit Packet Buffer Size of TCn. At default, only packet buffer 0 is enabled and TXPBSIZE values for TC 1-7 are meaningless. Symmetrical 8 TCs partitioning: 0x14 (20KB) for TXPBSIZE[0...7]. Symmetrical 4 TCs partitioning: 0x28 (40KB) for TXPBSIZE[0...3] and 0x0 (0KB) for TXPBSIZE[4...7]. Non-symmetrical partitioning are supported as well. In order to enable wire speed transmission, Intel recommends to set the transmit packet buffers to: 1. At least two times MSS plus PCIe latency (approximate 1 KB) when IPSec AH is not enabled (security block is not enabled or operates in path through mode). 2. At least three times MSS plus PCIe latency when IPSec AH is enabled (security block operates in store and forward mode). The default value is 0xA0 (160 KB).
RESERVED	31:20	0x0	RSV	Reserved.



8.2.2.7.14 Manageability Transmit TC Mapping - MNGTXMAP (0x0000CD10)

Field	Bit(s)	Init.	Access Type	Description
MAP	2:0	0x0	RW	MAP value indicates the TC that the transmit manageability traffic is routed to.
RESERVED	31:3	0x0	RSV	Reserved.

8.2.2.7.15 Multiple Transmit Queues Command Register - MTQC (0x00008120)

Field	Bit(s)	Init.	Access Type	Description
RESERVED	0	0b	RW	Reserved.
VT_ENA	1	0b	RW	Virtualization Enabled Mode. When set, the Intel® Xeon® Processor D-1500 Product Family LAN controller supports either 16, 32, or 64 pools. This bit should be set the same as PFVTCTL.VT_Ena.
NUM_TC_OR_Q	3:2	0x0	RW	Number of TCs or Number of Tx Queues per Pools. See functionality in the Tx Queues Assignment section.
RESERVED	31:4	0x0	RSV	Reserved.

Note: Permitted value and functionality of: VT_Ena; NUM_TC_OR_Q. For Tx queue assignment in VT mode refer to Tx Queuing Schemes Table in Tx Queues Assignment section. This register can be modified only as part of the init phase.

8.2.2.7.16 Tx Packet Buffer Threshold - TXPBTHRESH[n] (0x00004950 + 0x4*n, n=0...7)

Field	Bit(s)	Init.	Access Type	Description
THRESH	9:0	0x96	RW	THRESHold used for checking room place in the Tx packet buffer of TCn. Threshold is in KB units when the packet buffer is filled up with payload over that threshold; no more data read request is sent. Default values: 0x96 (150 KB) for TXPBSIZE0. 0x0 (0 KB) for TXPBSIZE1-7. It should be set to: packet buffer size - MSS. For example, if the packet buffer size is set to 20 KB in the corresponding TXPBSIZE.SIZE and if an MSS of 9.5 KB jumbo frames is supported for TCn, it is set to: 0xA (10 KB).
RESERVED	31:10	0x0	RSV	Reserved.

8.2.2.7.17 Tags Ethertypes - TAG_ETYPE (0x00017100)

Field	Bit(s)	Init.	Access Type	Description
ETAG_ETHER_TYPE	15:0	0x893F	RW	The Ethertype to use to identify E-tags.



8.2.2.8 TPH Registers

8.2.2.8.1 Rx TPH Control Register - TPH_RXCTRL_ALIAS[n] (0x00002200 + 0x4*n, n=0...15)

Field	Bit(s)	Init.	Access Type	Description
RX_DESC_FETCH_TPH_EN	0	0b	RW	Receive Descriptor Fetch TPH Enable. When set, hardware enables TPH for all Rx descriptors fetch from memory. When cleared, hardware does not enable TPH for descriptor fetches. This bit is cleared as a default.
RX_DESC_WRITE_BACK_TPH_EN	1	0b	RW	Receive Descriptor Write Back TPH Enable. When set, hardware enables TPH for all Rx descriptors written back into memory. When cleared, hardware does not enable TPH for descriptor write-backs. This bit is cleared as a default. The hint used is the hint set in the <i>Socket ID</i> field.
RX_HEADER_TPH_EN	2	0b	RW	Receive Header TPH Enable. When set, hardware enables TPH for all received header buffers. When cleared, hardware does not enable TPH for Rx headers. This bit is cleared as a default. The hint used is the hint set in the <i>Socket ID</i> field.
RX_PAYLOAD_TPH_EN	3	0b	RW	Receive Payload TPH Enable. When set, hardware enables TPH for all Ethernet payloads written into memory. When cleared, hardware does not enable TPH for Ethernet payloads. This bit is cleared as a default. The hint used is the hint set in the <i>Socket ID</i> field.
RESERVED	8:4	0x0	RSV	Reserved.
RXDESCREADROEN	9	1b	RW	Rx Descriptor Read Relax Order Enable.
RESERVED	10	0b	RSV	Reserved.
RXDESCRWBROEN	11	0b	RW	Rx Descriptor Write Back Relax Order Enable. This bit must be 0b to allow correct functionality of the descriptors write back.
RESERVED	12	0b	RSV	Reserved.
RXDATAWRITEROEN	13	1b	RW	Rx Data Write Relax Order Enable.
RESERVED	14	0b	RSV	Reserved.
RXREPHEADEROEN	15	1b	RW	Rx Split Header Relax Order Enable. This bit impacts only the relax order setting of the header part writes. The relax order setting of the data part is controlled via the RXDATAWRITEROEN bit.
RESERVED	23:16	0x0	RSV	Reserved.
CPUID	31:24	0x00	RW	Physical ID. TPH capable platforms - the software device driver programs a value, based on the relevant <i>Socket ID</i> , associated with this receive queue. Note that for TPH platforms, bits 31:27 of this field should always be set to 0x0.

Note: This is an alias to the TPH_RXCTRL registers.

8.2.2.8.2 Rx TPH Control Register - TPH_RXCTRL[n] (0x0000100C + 0x40*n, n=0...63 and 0x0000D00C + 0x40*(n-64), n=64...127)



Field	Bit(s)	Init.	Access Type	Description
RX_DESC_FETCH_TPH_EN	0	0b	RW	Receive Descriptor Fetch TPH Enable. When set, hardware enables TPH for all Rx descriptors fetch from memory. When cleared, hardware does not enable TPH for descriptor fetches. This bit is cleared as a default.
RX_DESC_WRITEBACK_TPH_EN	1	0b	RW	Receive Descriptor Write Back TPH Enable. When set, hardware enables TPH for all Rx descriptors written back into memory. When cleared, hardware does not enable TPH for descriptor write backs. This bit is cleared as a default. The hint used is the hint set in the <i>Socket ID</i> field.
RX_HEADER_TPH_EN	2	0b	RW	Receive Header TPH Enable. When set, hardware enables TPH for all received header buffers. When cleared, hardware does not enable TPH for Rx headers. This bit is cleared as a default. The hint used is the hint set in the <i>Socket ID</i> field.
RX_PAYLOAD_TPH_EN	3	0b	RW	Receive Payload TPH Enable. When set, hardware enables TPH for all Ethernet payloads written into memory. When cleared, hardware does not enable TPH for Ethernet payloads. This bit is cleared as a default. The hint used is the hint set in the <i>Socket ID</i> field.
RESERVED	8:4	0x0	RSV	Reserved.
RXDESCREADROEN	9	1b	RW	Rx Descriptor Read Relax Order Enable.
RESERVED	10	0b	RSV	Reserved.
RXDESCRWBROEN	11	0b	RW	Rx Descriptor Write Back Relax Order Enable. This bit must be 0b to allow correct functionality of the descriptors write back.
RESERVED	12	0b	RSV	Reserved.
RXDATAWRITEROEN	13	1b	RW	Rx Data Write Relax Order Enable.
RESERVED	14	0b	RSV	Reserved.
RXREPHEADEROEN	15	1b	RW	Rx Split Header Relax Order Enable. This bit impacts only the relax order setting of the header part writes. The relax order setting of the data part is controlled via the RXDATAWRITEROEN bit.
RESERVED	23:16	0x0	RSV	Reserved.
CPUID	31:24	0x00	RW	Physical ID. TPH capable platforms - the software device driver programs a value, based on the relevant <i>Socket ID</i> , associated with this receive queue. Note that for TPH platforms, bits 31:27 of this field should always be set to 0x0.

Note: Additional address(es): 0x0D00C + 0x40*(n-64), n=64...127TPH_RXCTRL[0...15] are also mapped to address 0x02200.

8.2.2.8.3 Tx TPH Control Registers - TPH_TXCTRL[n] (0x0000600C + 0x40*n, n=0...127)

Field	Bit(s)	Init.	Access Type	Description
TX_DESC_FETCH_TPH_EN	0	0b	RW	Transmit Descriptor Fetch TPH Enable. When set, hardware enables TPH for all Tx descriptors fetch from memory. When cleared, hardware does not enable TPH for descriptor fetches. This bit is cleared as a default.



Field	Bit(s)	Init.	Access Type	Description
TX_DESC_WB_TPH_EN	1	0b	RW	Transmit Descriptor Write Back TPH Enable. When set, hardware enables TPH for all Tx descriptors written back into memory. When cleared, hardware does not enable TPH for descriptor write backs. This bit is cleared as a default. The hint used is the hint set in the <i>Socket ID</i> field.
RESERVED	2	0b	RSV	Reserved.
TX_PACKET_TPH_EN	3	0b	RW	Transmit Packet TPH Enable. When set, hardware enables TPH for all Ethernet payloads read from memory. When cleared, hardware does not enable TPH for Ethernet payloads. This bit is cleared as a default.
RESERVED	8:4	0x0	RSV	Reserved.
TXDESCREADROEN	9	1b	RW	Tx Descriptor Read Relax Order Enable.
RESERVED	10	0b	RSV	Reserved.
TXDESCRWBROEN	11	0b	RW	Tx Descriptor Write Back Relax Order Enable.
RESERVED	12	0b	RSV	Reserved.
TXDATAREADROEN	13	1b	RW	Tx data read Relax Order Enable.
RESERVED	23:14	0x0	RSV	Reserved.
CPUID	31:24	0x0	RW	Physical ID. TPH capable platforms - the software device driver programs a value, based on the relevant <i>Socket ID</i> , associated with this transmit queue. Note that for TPH platforms, bits 31:27 of this field should always be set to 0x0.

8.2.2.9 Timers Registers

8.2.2.9.1 TCP Timer - TCPTIMER (0x0000004C)

Field	Bit(s)	Init.	Access Type	Description
DURATION	7:0	0x0	RW	Duration of the TCP interrupt interval in ms.
KICKSTART	8	0b	RW	Counter Kick-start. Writing 1b to this bit kick-starts the counter down count from the initial value defined in <i>Duration</i> field. Writing 0b has no effect.
TCPCOUNTEN	9	0b	RW	TCP Count Enable. When set to 1b, TCP timer counting is enabled. When set to 0b, it is disabled. Upon enabling, TCP counter must count from its internal state. If the internal state is equal to zero, down count does not restart until kick-start is activated. If the internal state is not 0b, down count continues from the internal state. This enables a pause of the counting for debug purposes.
TCPCOUNTFINISH	10	0b	RW	TCP Count Finish. This bit enables software to trigger a TCP timer interrupt, regardless of the internal state. Writing 1b to this bit triggers an interrupt, and resets the internal counter to its initial value. Down count does not restart until either kick-start is activated or loop is set. Writing 0b to this bit has no effect.



Field	Bit(s)	Init.	Access Type	Description
LOOP	11	0b	RW	TCP Loop. When set to 1b, TCP counter must reload duration each time it reaches zero, and must go on down counting from this point without kick-starting. When set to 0b, TCP counter must stop at a zero value, and must not re-start until kick start is activated.
RESERVED	31:12	0x0	RSV	Reserved.

8.2.2.9.2 FC Receive Control - FCRXCTRL (0x00005100)

Field	Bit(s)	Init.	Access Type	Description
RESERVED	1	0b	RW	Reserved.
RESERVED	2	0b	RW	Reserved.
RESERVED	3	0b	RW	Reserved.
RESERVED	4	0b	RW	Reserved.
RESERVED	5	0b	RW	Reserved.
RESERVED	7	1b	RW	Reserved.
RESERVED	11:8	0x0	RW	Reserved.
SMAC_EN	12	1b	RW	If cleared, <i>SMAC</i> is ignored.
DID_EN	13	1b	RW	If cleared, <i>D_ID</i> is ignored.
RESERVED	31:14	0x0	RSV	Reserved.

8.2.2.9.3 FC Indirect DMA Context: User Descriptor PTR Low - FCPTRL (0x00002410)

Field	Bit(s)	Init.	Access Type	Description
PTR_LOW	31:0	X	RW	User Descriptor PTR Low. 4 LSBs of the physical pointer to the user descriptor list. The pointer must be 16-byte aligned so the 4 LS bits are read only as zeros.

8.2.2.9.4 FC Indirect DMA Context: User Descriptor PTR High - FCPTRH (0x00002414)

Field	Bit(s)	Init.	Access Type	Description
PTR_HI	31:0	X	RW	User Descriptor PTR High. 4 MSBs of the physical pointer to the user descriptor list.

8.2.2.9.5 FC Indirect DMA Context: Buffer Control - FCBUFF (0x00002418)

Field	Bit(s)	Init.	Access Type	Description
VALID	0	0b	RW	DMA Context Valid. When set to 1b, indicates that the context is valid. If software clears the <i>Context Valid</i> bit, the software should poll it until it is actually cleared by hardware before unlocking the user buffers.



Field	Bit(s)	Init.	Access Type	Description
FIRST	1	0b	RW	DMA First. This bit is a status indication. Software should clear it during FC context programming. The DMA unit sets this bit when it receives a frame that matches the context and marked by the filter unit as first.
LAST	2	0b	RW	DMA Last. This bit is a status indication. Software should clear it during FC context programming. Hardware sets this bit when it exhausts the last user buffer.
BUFFSIZE	4:3	0x0	RW	Buffer Size. This field defines the user buffer size used in this context as follows: 00b = 4 KB. 01b = 8 KB. 10b = 16 KB. 11b = 64 KB.
WRCONTX	5	0b	RW	Write DDP Context. This bit should be set to 1b for write exchange context aimed for target (responder) usage. This bit should be set to 0b for read exchange context aimed for initiator (originator) usage.
BUFFCNT	15:6	0x0	RW	Buffer Count. Defines the number of the user buffers while 0x00 equals 1024. It is programmed by the software and updated by hardware during reception.
OFFSET	31:16	0x0	RW	User Buffer Offset. Byte offset within the user buffer to which the FC data of large FC receive should be posted.

8.2.2.9.6 FC Indirect DMA Context: Receive DMA RW - FCDMARW (0x00002420)

Field	Bit(s)	Init.	Access Type	Description
RESERVED	10:0	0x0	RW	Reserved.
RESERVED	12:11	0x0	RSV	Reserved.
WE	14	0b	RW	This bit should never be set together with the <i>RE</i> bit in this register.
RE	15	0b	RW	This bit should never be set together with the <i>WE</i> bit in this register.
LASTSIZE	31:16	0x0	RW	Last User Buffer Size. Defined the size in bytes of the last user buffer.

8.2.2.9.7 FC Indirect Filter Context: Control - FCFLT (0x00005108)

Field	Bit(s)	Init.	Access Type	Description
VALID	0	0b	RW	Filter Context Valid. When set to 1b indicates that the context is valid.
FIRST	1	0b	RW	Filter First. This bit is a status indication. Software should clear it during FC context programming. The filter unit sets this bit when it receives a first frame that matches the context.
RESERVED	7:2	X	RSV	Reserved.
SEQ_ID	15:8	X	RW	Sequence ID. The sequence ID of the last received frame. Set to 0b by the software device driver at context programming.
SEQ_CNT	31:16	X	RW	Sequence Count. The sequence count of the expected received frame. Should be set to 0b for read contexts. Should be set to SEQ_CNT + 1 of the last packet of the same exchange received from the initiator for write contexts.



8.2.2.9.8 FC Indirect Filter Context: Offset Parameter - FCPARAM (0x000051D8)

Field	Bit(s)	Init.	Access Type	Description
PARAM	31:0	0x0	RW	FC Parameter. This field contains the expected FC parameter in the next received frame. Set to 0b by the driver at context programming. Note: Field is defined in big endian (LSB is first on the wire).

8.2.2.9.9 FC Indirect Filter Context: D_ID Register - FCD_ID (0x00005114)

Field	Bit(s)	Init.	Access Type	Description
FCD_ID	23:0	0x0	RW	Reserved.
RESERVED	31:24	0x0	RSV	Reserved.

8.2.2.9.10 FC Indirect Filter Context: Source MAC Address - FCSMAC (0x0000510C)

Field	Bit(s)	Init.	Access Type	Description
FCSMAC	31:0	0x0	RW	Reserved.

Note: This register is stored in network ordering (big endian).

8.2.2.9.11 FC Indirect Filter Context: RW Control - FCFLTRW (0x00005110)

Field	Bit(s)	Init.	Access Type	Description
RESERVED	10:0	0x0	WO	Reserved.
RESERVED	12:11	0x0	RSV	Reserved.
RE_VALIDATE	13	0b	WO	Fast Re-validation of the Filter Context. Setting this bit together with the <i>WE</i> bit in this register validates the selected filter context. Hardware sets the <i>Valid</i> bit and clears the <i>First</i> bit (described in the FCFLT register) while keeping all other filter parameters intact.
WE	14	0b	WO	This bit should never be set together with the <i>RE</i> bit in this register.
RE	15	0b	WO	This bit should never be set together with the <i>WE</i> bit in this register.
SMAC_HIGH	31:16	0x0	RW	This field is stored in network ordering (big endian).

8.2.2.10 Flow Director Registers

These registers are used to control the flow director functionality.

8.2.2.10.1 Flow Director Filters Control Register - FDIRCTRL (0x0000EE00)

Note: This register should be configured ONLY as part of the flow director initialization flow or clearing the flow director table. Programming of this register with non-zero value PBALLOC initializes the flow director table.



Field	Bit(s)	Init.	Access Type	Description
PBALLOC	1:0	0x0	RW	Memory Allocation for the Flow Director Filters. 00b = No memory allocation - flow director filters are disabled. 01b = 64 KB (8 KB minus 1 signature filters or 2 KB minus 1 perfect match filters). 10b= 128 KB (16 KB minus 1 signature filters or 4 KB minus 1 perfect match filters). 11b= 256 KB (32 KB minus 1 signature filters or 8 KB minus 1 perfect match filters).
RESERVED	2	0b	RSV	Reserved.
INIT_DONE	3	0b	RO	Flow Director Initialization Completion Indication (read only status). Indicates that hardware initialized the flow director table according to the PBALLOC setting. Software must not access any other flow director filters registers before the <i>INIT-Done</i> bit is set. When flow director filters are enabled (PBALLOC > 0b, the software must wait for the <i>INIT-Done</i> indication before Rx is enabled.
PERFECT_MATCH	4	0b	RW	Flow Director Filters Mode of Operation. When set to 1b, hardware support perfect match filters according to <i>PBALLOC</i> . When cleared to 0b, hardware support signature filters according to the <i>PBALLOC</i> .
REPORT_STATUS	5	0b	RW	Report flow director filter's status in the RSS field of the Rx descriptor for packets that matches a flow director filter. Enabling the flow director filter's status, the <i>RXCSUM.PCSD</i> bit should be set as well (disabling the fragment checksum). Note that the flow director filter status and error bits in the <i>Extended Status</i> and <i>Error</i> fields in the Rx descriptor are always enabled.
RESERVED	6	0b	RSV	Reserved.
REPORT_STATUS_ALWAYS	7	0b	RW	Report flow director status in the RSS field of the Rx descriptor on any packet that can be candidates for the flow director filters. This bit can be set to 1b only when both the <i>RXCSUM.PCSD</i> bit and the <i>Report-Status</i> bit in this register are set.
DROP_QUEUE	14:8	0x0	RW	Absolute Rx queue index used for the dropped packets. Software might set this queue to an empty one by setting the <i>RDLEN[n]</i> to 0x0.
DROP_NO_MATCH	15	0b	RW	If set, send to the <i>DROP_QUEUE</i> packets that where candidate for flow director and did not match any filters. Candidates of packets are defined in the candidate for flow director per mode table in flow director filters section.
FLEX_OFFSET	20:16	0x0	RW	Offset within the first 64 bytes of the packet of a flexible 2-byte tuple. The offset is defined in word units counted from the first byte of the destination Ethernet MAC address.
FILTERMODE	23:21	0x0	RW	Filter Mode. Defines the fields on which the Flow director acts: 000b= IPMODE - L3/L4 tuple. 001b= MACVLANMODE - acts on the MAC and VLAN. 010b= Cloud: NVGRE - based on TNI, inner MAC, inner VLAN. VXLAN - based on VNI, inner MAC, inner VLAN. 011b - 111b= Reserved.



Field	Bit(s)	Init.	Access Type	Description
MAX_LENGTH	27:24	0x0	RW	Max Linked List Length. This field defines the maximum recommended linked list associated to any hash value. Packets that match filters that exceed MAX_LENGTH are reported with an active <i>Length</i> bit in the <i>Extended Error</i> field. In addition, drop filters that exceed the MAX_LENGTH are posted to the Rx queue defined in the filter context rather than the drop-queue defined in this register. MAX_LENGTH is defined in units of two filters and exceeding it reflects the addition of two more filters. Note: Software should set this field to a value that indicates exceptional long buckets. Supporting 32 KB filters with good hash scheme key, it is expected that a value of 0xA might be a good choice.
FULL_THRES H	31:28	0x0	RW	Full threshold is a recommended minimum number of flows that should remain unused (defined in units of 16 filters). When software exceeds this threshold (too low of a number of unused flows), hardware generates the flow director full interrupt. Software should avoid additional programming following this interrupt.

8.2.2.10.2 Flow Director Filters Lookup Table HASH Key - FDIRHKEY (0x0000EE68)

Field	Bit(s)	Init.	Access Type	Description
KEY	31:0	0x80000001	RW	Programmable Hash lookup Table Key.

8.2.2.10.3 Flow Director Filters Signature HASH Key - FDIRSKEY (0x0000EE6C)

Field	Bit(s)	Init.	Access Type	Description
KEY	31:0	0x80800101	RW	Programmable Signature Key.

8.2.2.10.4 Flow Director Filters DIPv4 Mask - FDIRDIP4M (0x0000EE3C)

Field	Bit(s)	Init.	Access Type	Description
IPM	31:0	0x0	RW	Mask Destination IPv4 Address. Each cleared bit means that the associated bit of the destination IPv4 address is meaningful for the filtering functionality. Each bit set to 1b means that the associated bit of the destination IPv4 address is ignored (masked out). The Lsb of this register matches the first byte on the wire.



8.2.2.10.5 Flow Director Filters Source IPv4 Mask - FDIRSIP4M (0x0000EE40)

Field	Bit(s)	Init.	Access Type	Description
IPM	31:0	0x0	RW	Mask Source IPv4 Address. Each cleared bit means that the associated bit of the source IPv4 address is meaningful for the filtering functionality. Each bit set to 1b means that the associated bit of the source IPv4 address is ignored (masked out). The LSB of this register matches the first byte on the wire.

8.2.2.10.6 Flow Director Filters TCP Mask - FDIRTCPM (0x0000EE44)

Field	Bit(s)	Init.	Access Type	Description
SPORTM	15:0	0x0	RW	Mask TCP Source Port. Each cleared bit means that the associated bit of the TCP Source port is meaningful for the filtering functionality. Each bit set to 1b means that the associated bit of the TCP source port is ignored (masked out). Note that this register is swizzle as follows: bit 0 in the mask affects bit 15 of the source port as defined in the <i>FDIRPORT.Source</i> . bit 1 in the mask affects bit 14 in the <i>FDIRPORT.Source</i> and so on while bit 15 in the mask affects bit 0 in the <i>FDIRPORT.Source</i> .
DPORTM	31:16	0x0	RW	Mask TCP Destination Port. Each cleared bit means that the associated bit of the TCP destination port is meaningful for the filtering functionality. Each bit set to 1b means that the associated bit of the TCP destination port is ignored (masked out). Note that this register is swizzle and is the same as FDIRTCPM.SPortM.

8.2.2.10.7 Flow Director Filters UDP Mask - FDIRUDPM (0x0000EE48)

Field	Bit(s)	Init.	Access Type	Description
SPORTM	15:0	0x0	RW	Mask UDP Source Port. Each cleared bit means that the associated bit of the UDP source port is meaningful for the filtering functionality. Each bit set to 1b means that the associated bit of the UDP source port is ignored (masked out). Note that this register is swizzle and is the same as FDIRTCPM.SPortM.
DPORTM	31:16	0x0	RW	Mask UDP Destination Port. Each cleared bit means that the associated bit of the UDP destination port is meaningful for the filtering functionality. Each bit set to 1b means that the associated bit of the UDP destination port is ignored (masked out). Note that this register is swizzle and is the same as FDIRTCPM.SPortM.



8.2.2.10.8 Flow Director Filters SCTP Mask - FDIRSCTPM (0x0000EE78)

Field	Bit(s)	Init.	Access Type	Description
SPORTM	15:0	0x0	RW	Mask SCTP Source Port. Each cleared bit means that the associated bit of the SCTP source port is meaningful for the filtering functionality. Each bit set to 1b means that the associated bit of the SCTP source port is ignored (masked out). Note that this field is layed out as follows: Bit 0 in the mask affects bit 15 of the source port as defined in <i>FDIRPORT.Source</i> . Bit 1 in the mask affects bit 14 in <i>FDIRPORT.Source</i> Bit 15 in the mask affects bit 0 in <i>FDIRPORT.Source</i> .
DPORTM	31:16	0x0	RW	Mask SCTP Destination Port. Each cleared bit means that the associated bit of the SCTP destination port is meaningful for the filtering functionality. Each bit set to 1b means that the associated bit of the SCTP destination port is ignored (masked out). Note that this field is layed out as follows: Bit 0 in the mask affects bit 15 of the source port as defined in <i>FDIRPORT.DESTINATION</i> . Bit 1 in the mask affects bit 14 in <i>FDIRPORT.DESTINATION</i> Bit 15 in the mask affects bit 0 in <i>FDIRPORT.DESTINATION</i> .

8.2.2.10.9 Flow Director Filters IPv6 Mask - FDIRIP6M (0x0000EE74)

Field	Bit(s)	Init.	Access Type	Description
SIPM	15:0	0x0	RW	Mask Source IPv6 Address. Each cleared bit means that the associated byte of the source IPv6 address is meaningful for the filtering functionality. Each bit set to 1b means that the associated byte of the source IPv6 address is ignored (masked out). The LSB of this register matches the first byte on the wire. When operating in MACVLAN mode, this field defines the mask to be used to compare the MAC address. When operating in VXLAN or NVGRE mode, it should be used to enable comparison of the TNI/VNI and inner MAC address:.
DIPM	31:16	0x0	RW	Mask Destination IPv6 Address. Each cleared bit means that the associated byte of the destination IPv6 address is meaningful for the filtering functionality. Each bit set to 1b means that the associated byte of the destination IPv6 address is ignored (masked out). The entire field is meaningful only for the hash function and the signature based filters. The <i>DIPv6</i> bit in the FDIRM register is meaningful for perfect match filters. The LSB of this register matches the first byte on the wire.

8.2.2.10.10 Flow Director Filters Other Mask - FDIRM (0x0000EE70)

Field	Bit(s)	Init.	Access Type	Description
VLANID	0	0b	RW	Mask VLAN ID Tag. When cleared, the 12 bits of the VLAN ID tag are meaningful for the filtering functionality.
VLANP	1	0b	RW	Mask VLAN Priority Tag. When cleared, the three bits of the VLAN priority are meaningful for the filtering functionality.
POOL	2	0b	RW	Mask Pool. When cleared, the target pool number is meaningful for the filtering functionality.



Field	Bit(s)	Init.	Access Type	Description
L4P	3	0b	RW	Mask L4 Protocol. When cleared, the UDP/TCP/SCTP protocol type is meaningful for the filtering functionality. Note that for the flow director filtering aspects, SCTP is treated as if it is TCP.
FLEX	4	0b	RW	Mask Flexible Tuple. When cleared, the two bytes of the flexible tuple are meaningful for the filtering functionality.
DIPV6	5	0b	RW	Mask Destination IPv6. When cleared, the compare against the IP6AT filter is meaningful for IPv6 packets.
L3P	6	0b	RW	Mask IP type comparison (in order to block all IP comparison, the <i>DIPV6</i> field, <i>FDIRIP6M</i> and <i>FDIRSIP4M</i> and <i>FDIRDIP4M</i> should be set also).
RESERVED	31:7	0x0	RSV	Reserved.

8.2.2.10.11 Flow Director Filters Free - FDIRFREE (0x0000EE38)

Field	Bit(s)	Init.	Access Type	Description
FREE	15:0	0x0000	RW	Number of free (non programmed) filters in the flow director filters logic.
RESERVED	31:16	0x0	RSV	Reserved.

8.2.2.10.12 Flow Director Filters Length - FDIRLEN (0x0000EE4C)

Field	Bit(s)	Init.	Access Type	Description
MAXLEN	5:0	0x0	RC	Longest Linked List of Filters in the Table. This field records the length of the longest linked list that is updated since the last time this register was read by the software. The longest bucket reported by this field includes MAXLEN + 1 filters.
RESERVED	7:6	0x0	RSV	Reserved.
BUCKET_LEN GTH	13:8	0x0	RC	The length of the linked list indicated by a query command. This field is valid following a query command completion.
RESERVED	15:14	0x0	RSV	Reserved.
MAXHASH	30:16	0x0	RC	The lookup HASH value of the added filter that updated the value of the MAXLEN field in this register.
RESERVED	31	0b	RSV	Reserved.

8.2.2.10.13 Flow Director Filters Usage Statistics - FDIRUSTAT (0x0000EE50)

Field	Bit(s)	Init.	Access Type	Description
ADD	15:0	0x0	RC	Number of Added Filters. This field counts the number of added filters to the flow director filters logic. The counter is stuck at 0xFFFF and cleared on read.
REMOVE	31:16	0x0	RC	Number of Removed Filters. This field counts the number of removed filters to the flow director filters logic. The counter is stuck at 0xFFFF and cleared on read.



8.2.2.10.14 Flow Director Filters Failed Usage Statistics - FDIRFSTAT (0x0000EE54)

Field	Bit(s)	Init.	Access Type	Description
FADD	7:0	0x0	RC	Number of failed added filters due to no space in the filter table. The counter is stuck at 0xFF and cleared on read.
FREMOVE	15:8	0x0	RC	Number of failed removed filters. The counter is stuck at 0xFF and cleared on read.
RESERVED	31:16	0x0	RSV	Reserved.

8.2.2.10.15 Flow Director Filters Match Statistics - FDIRMATCH (0x0000EE58)

Field	Bit(s)	Init.	Access Type	Description
PCNT	31:0	0x0	RC	Number of packets that matched any flow director filter. The counter is stacked at 0xFF..F and cleared on read. Note that this counter might include packets that match the L2 filters or 5 tuple filters or syn filters even if they are enabled for queue assignment.

8.2.2.10.16 Flow Director Filters Miss Match Statistics - FDIRMISS (0x0000EE5C)

Field	Bit(s)	Init.	Access Type	Description
PCNT	31:0	0x0	RC	Number of packets that missed matched any flow director filter. The counter is stacked at 0xFF..F and cleared on read.

Note: Flow Programming registers.

8.2.2.10.17 Flow Director Filters Source IPv6 - FDIRSIPV6[n] (0x0000EE0C + 0x4*n, n=0...2)

Field	Bit(s)	Init.	Access Type	Description
IP6SA	31:0	0x0	RW	Three LS Dwords of the Source IPv6 Address. The FDIRIPSA contains the MS DWord of the IP6 address. The LS byte of FDIRIPSA is first on the wire and the MS byte of FDIRSIPV6[2] is last on the wire. In MACVLANMODE FDIRSIPV6_0[31:0] and FDIRSIPV6_1[15:0] hold, the destination MAC address value FDIRSIPV6_1[31:16] and FDIRSIPV6_2[31:0] are reserved and should be set to 0x0.

8.2.2.10.18 Flow Director Filters IP SA - FDIRIPSA (0x0000EE18)

Field	Bit(s)	Init.	Access Type	Description
IP4SA	31:0	0x0	RW	Source IPv4 address or MS Dword of the source IPv6 address. While the field is defined in big endian (LSB is first on the wire). In MACVLANMODE, this field is reserved and should be set to 0x0.



8.2.2.10.19 Flow Director Filters IP DA - FDIRIPDA (0x0000EE1C)

Field	Bit(s)	Init.	Access Type	Description
IP4DA	31:0	0x0	RW	Destination IPv4 Address. This field is defined in big endian (LSB is first on the wire). In MACVLANMODE, this field is reserved and should be set to 0x0.

8.2.2.10.20 Flow Director Filters Port - FDIRPORT (0x0000EE20)


Field	Bit(s)	Init.	Access Type	Description
SOURCE	15:0	0x0	RW	Source port number while the field is defined in little endian (MSB is first on the wire). In MACVLANMODE and cloud modes (NVGRE/VXLAN), this field is reserved and should be set to 0x0.
DESTINATION	31:16	0x0	RW	Destination port number while the field is defined in little endian (MSB is first on the wire). In MACVLANMODE and cloud modes (NVGRE/VXLAN), this field is reserved and should be set to 0x0.

8.2.2.10.21 Flow Director Filters VLAN and FLEX Bytes - FDIRVLAN (0x0000EE24)

Field	Bit(s)	Init.	Access Type	Description
VLAN	15:0	0x0	RW	VLAN Tag. This field is defined in little endian (MSB is first on the wire). The <i>CFI</i> bit must be set to zero and it is not checked by hardware.
FLEX	31:16	0x0	RW	Flexible tuple data as defined by the <i>Flex-Offset</i> field in the FDIRCTRL register while the field is defined in big endian (LSB is first on the wire).

8.2.2.10.22 Flow Director Filters Hash Signature - FDIRHASH (0x0000EE28)

Field	Bit(s)	Init.	Access Type	Description
HASH	14:0	0x0	RW	Bucket hash value that identifies a filter's linked list.
BUCKET_VALID	15	0b	RW	The <i>Valid</i> bit is set by hardware each time there is at least one filter assigned to this hash.
SIGNATURE_SW_INDEX	30:16	0x0	RW	Flow Director Filter Signature for Signature Filters and SW-index for Perfect Match Filters. In perfect match mode, the two MS bits should be zero.
RESERVED	31	0b	RSV	Reserved.



8.2.2.10.23 Flow Director Filters Command Register - FDIRCMD (0x0000EE2C)

Field	Bit(s)	Init.	Access Type	Description
CMD	1:0	0x0	RW	Flow Director Filter Programming Command. 00b = No action. 01b = Add flow. 10b = Remove flow. 11b = Query command. Following a command completion, hardware clears the <i>CMD</i> field. In a query command, all other parameters are valid when the <i>CMD</i> field is zero.
FILTER_VALID	2	0b	RW	A valid filter is found by the query command. This bit is set by the Intel® Xeon® Processor D-1500 Product Family LAN controller following a query command completion.
FILTER_UPDATE	3	0b	RW	Filter Update Command. This bit is relevant only for the add flow command and must be set to zero in any other commands. When cleared, the filter parameters do not override existing ones if they exist while setting only the collision bit. When set to 1b the new filter parameters override existing ones if exist keeping the collision bit as is.
IPV6DMATCH	4	0b	RW	IP Destination Match to IP6AT Filter. This bit is meaningful only for perfect match IPv6 filters. Otherwise, it should be cleared by software at programming time. When set to 1b, the destination IPv6 address should match the IP6AT. When cleared, the destination IPv6 address should not match the IP6AT. This field might never match local VM-to-VM traffic.
L4TYPE	6:5	0x0	RW	L4 Packet Type. Defines the packet as one of the following L4 types: 00b = Reserved. 01b = UDP 10b = TCP 11b = SCTP
IPV6	7	0b	RW	This field determines the IPv6 packet type when set to 1b and IPv4 packet type when set to 0b. Relevant only if FDIRM.L3P is not set.
CLEARHT	8	0b	RW	Clear the Flow Director Head and Tail Registers. This bit is set only as part of flow director initialization. During normal operation it must be kept at 0b.
LAST	11	0b	RW	Last Filter Indication in the Linked List. At flow programming, software should set the last bit to 1b. Note that hardware might modify this bit when adding or removing flows from the same linked list.
COLLISION	12	0b	RW	Collision Indication. This field is set to 1b when software programs the same filter multiple times. In signature-based filtering it is set when software programs a filter with the same hash and signature multiple times. It should be cleared by software when it adds a flow. It might be set by hardware when two flows collide with the same hash and signature. During reception, this bit is reported on the Rx descriptor of packets that match the filter. See bit 7 for a description of the query-type.
RESERVED	14:13	0x0	RSV	Reserved.
QUEUE_EN	15	0b	RW	Enable routing matched packet to the queue defined by the Rx-Queue. Note that packets redirection to the FDIRCTRL.Drop-Queue is not gated by the <i>Queue-EN</i> bit.
RX_QUEUE	22:16	0x0	RW	Rx Queue Index. This field defines the absolute Rx queue index in all modes of operation; regardless of VT enablement.



Field	Bit(s)	Init.	Access Type	Description
TUNNEL_FILTER	23	0b	RSV	If set, the filter is matched by NVGRE or VXLAN packets only. In this case, the L3 and L4 parameters programmed should be of the tunneled (inner) header.
POOL	29:24	0x0	RW	Pool number is meaningful when VT mode is enabled. When both VT is not enabled, this field must be set by software to 0x0.
RESERVED	30	0b	RSV	Reserved.
FD_EXCEPTION	31	0b	RW	If set after a command was given, indicates a flow director exception occurred. This bit is a reflection of the EICR.FLOW_DIRECTOR bit. Should be set to 0b when a new command is given.

8.2.2.11 MAC Registers

8.2.2.11.1 PHY Indirect Data - PHY_INDIRECT_DATA (0x00011148)

This register is used for indirect access to the PHY and PCS registers. A write to this register triggers a write operation that would write the data written to this register based on the information in the PHY Indirect Ctrl register. A read from this register triggers a read operation that would read data based on the information in the PHY Indirect Ctrl register.

Field	Bit(s)	Init.	Access Type	Description
DATA	31:0	0x0	RW	The indirect data to write or the data received following an indirect read.

8.2.2.11.2 PHY Indirect Control - PHY_INDIRECT_CTRL (0x00011144)

This register is used for indirect access to the PHY and PCS registers. It is used to set the address plus other control settings.

Field	Bit(s)	Init.	Access Type	Description
ADDR	15:0	0x0	RW	16-bit address to be used for write or read transaction.
RESERVED	16	0b	RSV	Reserved.
RESERVED	17	0b	RSV	Reserved.
RESP_STAT	19:18	0x0	RO	Response Status: 00b = Successful. 01b = Unsuccessful / not supported. 10b = Reserved. 11b = Powered down. This field reflects the response status for the previously completed transaction. The value of this register is only meaningful if busy = 0x0.



Field	Bit(s)	Init.	Access Type	Description
C MPL_ERR	27:20	0x0	RO	When the transaction is unsuccessful, this field contains the error type. [20]= Wrong source integrated I/O interface ID. [21]= Wrong destination integrated I/O interface ID. [22]= EH not set. [23]= Wrong SAI. [24]= Wrong TAG. [25]= Wrong OP CODE. [26]= Data too long.
PHY_SELECT	30:28	0x0	RWS	Select the destination PHY or PCS block within the attached PHY block(s). Supports up to eight separate endpoints. Current values supported are: 0x0 = KR. 0x1 = Reserved. 0x2 = Reserved.
BUSY	31	0b	RO	The PHY Indirect Control register should not be updated when this bit is set. Hardware sets this bit when it starts to execute a transaction to the PHY. Hardware clears this bit when the transaction completes.

8.2.2.11.3 Core Control 0 Register - HLREG0 (0x00004240)

Field	Bit(s)	Init.	Access Type	Description
RESERVED	1	1b	RW	Reserved.
JUMBOEN	2	0b	RW	Jumbo Frame Enable. Allows frames up to the size specified in Reg MAXFRS (31:16). 1b = Enable jumbo frames. 0b = Disable jumbo frames (default).
RESERVED	9:3	0x7F	RSV	Reserved. must be set to 1111111b.
TXPADEN	10	1b	RW	Tx Pad frame Enable. Pad short Tx frames to 64 bytes if requested by the user. 1b = Pad frames (default). 0b = Transmit short frames with no padding.
RESERVED	11	1b	RSV	Reserved. must be set to 1b.
RESERVED	12	0b	RW	Reserved.
RESERVED	13	1b	RSV	Reserved. must be set to 1b.
LPBK	15	0b	RW	Loopback. Turn on loopback where transmit data is sent back through the receiver. 1b = Loopback enabled. 0b = Loopback disabled (default).
MDCSPD	16	1b	RW	MDC Speed. High Or low speed MDC clock frequency to PCS, XGXS, WIS, etc. MDCSPD Freq at 10 Gb/s Freq at 1Gb/s Freq at 100 Mb/s 0b = (default) 2.4 MHz 240 KHz 240 KHz 1b= 24 MHz 2.4 MHz 240 KHz
CONTMDC	17	0b	RW	Continuous MDC. Turn off MDC between MDIO packets, 1b = Continuous MDC. 0b = MDC off between packets (default).



Field	Bit(s)	Init.	Access Type	Description
PREPEND	23:20	0x0	RW	Prepend Value. Number of 32-bit words starting after the preamble and SFD to exclude from the CRC generator and checker (default 0x-0).
RESERVED	24	0b	RSV	Reserved.
RESERVED	26:25	0x0	RSV	Reserved.
RXLNGTHERR EN	27	1b	RW	Rx Length Error Reporting. 1b = Enable reporting of rx_length_err events if length field < 0x0600. 0b = Disable reporting of all rx_length_err events.
RXPADSTRIP EN	28	0b	RW	Rx Padding Strip Enable 1b = Strip padding from Rx packets with length field < 64 (debug only). 0b = Do not strip padding from Rx packets with length field < 64 (default). Note: This functionality should be used as debug mode only. If Rx pad stripping is enabled, then the Rx CRC stripping needs to be enabled as well.
RESERVED	31:29	0x0	RSV	Reserved.

8.2.2.11.4 Core Status 1 Register - HLREG1 (0x00004244)

Field	Bit(s)	Init.	Access Type	Description
RESERVED	3:0	0x1	RSV	Reserved.
RESERVED	4	0b	RSV	Reserved.
RXERRSYM	5	0b	RO	Rx Error Symbol. Error symbol during Rx packet (latch high, clear on read). 1b = Error symbol received. 0b = No error symbol (default).
RXILLSYM	6	0b	RO	Rx Illegal Symbol. Illegal symbol during Rx packet (latch high, clear on read). 1b = Illegal Symbol Received. 0b = No Illegal Symbol Received (default).
RXIDLERR	7	0b	RO	Rx Idle Error. Non-idle symbol during idle period (latch high, clear on read). 1b = Idle error received. 0b = No idle errors received (default).
RXLCLFLT	8	0b	RO	Rx Local Fault. Fault reported from PMD, PMA, or PCS (latch high, clear on read). 1b = Local fault is or was active. 0b = No local fault (default).
RXRMTFLT	9	0b	RO	Rx Remote Fault. Link partner reported remote fault (latch high, clear on read). 1b = Remote fault is or was active. 0b = No remote fault (default).
RESERVED	31:10	0x0	RSV	Reserved.



8.2.2.11.5 Pause and Pace Register - PAP (0x00004248)

Field	Bit(s)	Init.	Access Type	Description
PACE	19:16	0x0	RW	0000b = 10 Gb/s (LAN). 0001b = 1 Gb/s. 0010b = 2 Gb/s. 0011b = 3 Gb/s. 0100b = 4 Gb/s. 0101b = 5 Gb/s. 0110b = 6 Gb/s. 0111b = Gb/s 1000b = 8 Gb/s 1001b = 9 Gb/s. 1111b = 9.294196 Gb/s (WAN).
RESERVED	31:20	0x0	RSV	Reserved.

8.2.2.11.6 Max Frame Size - MAXFRS (0x00004268)

Field	Bit(s)	Init.	Access Type	Description
RESERVED	15:0	0x0	RSV	Reserved.
MFS	31:16	0x5EE	RW	This field defines the Maximum Frame Size (MFS) in byte units from the Ethernet MAC addresses up to inclusive the CRC. Frames received that are larger than this value are dropped. This field is meaningful when jumbo frames are enabled (HLREG0.JUMBOEN = 1b). When jumbo frames are not enabled, the Intel® Xeon® Processor D-1500 Product Family LAN controller uses a hard-wired value of 1518 for this field. The MFS does not include the four bytes of the VLAN header. Packets with a VLAN header might be as large as MFS + 4. When double VLAN is enabled, the Intel® Xeon® Processor D-1500 Product Family LAN controller adds eight to the MFS for any packet. When E-tags are enabled, the device adds 12 to the MFS for any packet. This value has no effect on transmit frames; it is the responsibility of software to limit the size of transmit frames. Note: Packets to/from the MC are limited to 2 KB even when jumbo frames are enabled.

8.2.2.11.7 Link Status Register - LINKS (0x000042A4)

Field	Bit(s)	Init.	Access Type	Description
FIFO_UNDER RUN	0	0b	RO	Indicates an under-run condition in MAC elastic FIFO.
FIFO_OVERR UN	1	0b	RO	Indicates an over-run condition in MAC elastic FIFO.
RF_STATE	2	0b	RO	MAC is in remote fault state.
LF_STATE	3	0b	RO	MAC is in local fault state.
RESERVED	6:4	0x0	RSV	Reserved.



Field	Bit(s)	Init.	Access Type	Description
LINK_STATUS	7	0b	RO	1b = Link is up and there was no link down from last time read. 0b = Link is currently down or link was down since last time read. Self cleared upon read if the link is low and set if the link is up.
RESERVED	26:8	0x0	RSV	Reserved.
NON_STANDARD_SPEED	27	0b	RO	If set, the <i>LINK_SPEED</i> field reflects non-standard speeds (2.5/5 Gb/s)
LINK_SPEED	29:28	0x0	RO	MAC Link Speed Status. If NON_STANDARD_SPEED = 0b: 00b = Reserved. 01b = 100 Mb/s. 10b = 1 Gb/s. 11b = 10 Gb/s. If NON_STANDARD_SPEED = 1b: 00b = Reserved. 01b = 5 Gb/s. 10b = Reserved. 11b = 2.5 Gb/s.
LINK_UP	30	0b	RO	Link Up. 1b = Link is up. 0b = Link is down.
RESERVED	31	0b	RSV	Reserved.

8.2.2.11.8 MAC Manageability Control Register - MMNGC (0x000042D0)

Field	Bit(s)	Init.	Access Type	Description
MNG_VETO	0	0b	RO	MNG_VETO (default 0) access read/write by the management controller; read only to the host. 0b = No specific constraints on link from the MC. 1b = Hold off any low-power link mode changes. This is done to avoid link loss and interrupting management traffic/activity.
RESERVED	31:1	0x0	RSV	Reserved.

8.2.2.11.9 MAC Control Register - MACC (0x00004330)

Field	Bit(s)	Init.	Access Type	Description
FLU	0	0b	RW	Force Link Up. 0b = Normal mode. 1b = MAC is forced to the link up state regardless to the PHY link status.
MAC_RX2TX_LPBK_EN	1	0b	RW	Enable MAC Rx-to-Tx Loopback. 0b = No loopback (normal mode). 1b = Loopback enabled. Transmit path is driven from the receive path; at the MAC internal XGMII interface.
SWIZZLE_TX_DATA	2	0b	RW	Swizzle the bytes in all 4 MAC internal Tx XGMII lanes. 0b = Swizzle disabled (normal mode). 1 - Swizzle enabled.
SWAP_TX_DATA	3	0b	RW	Swap MAC Internal Tx XGMII Lanes. 0b = Swap disabled (normal mode). 1b = Swap enabled.



Field	Bit(s)	Init.	Access Type	Description
SWAP_TX_CONTROL	4	0b	RW	Swap MAC Internal Tx XGMII Controls. 0b = Swap disabled (normal mode). 1b = Swap enabled.
SWIZZLE_RX_DATA	5	0b	RW	Swizzle Bytes in all Four MAC Internal Rx XGMII Lanes. 0b = Swizzle disabled (normal mode). 1b = Swizzle enabled.
SWAP_RX_DATA	6	0b	RW	Swap MAC Internal Rx XGMII Lanes. 0b = Swap disabled (normal mode). 1b = Swap enabled.
SWAP_RX_CONTROL	7	0b	RW	Swap the MAC Internal Rx XGMII Controls. 0b = Swap disabled, normal mode. 1b = Swap enabled.
FIFOHT	11:8	0xD	RW	FIFO High Threshold. Determines the high threshold of the MAC elastic FIFO.
FIFOLT	15:12	0x3	RW	FIFO Low Threshold. Determines the low threshold of the MAC elastic FIFO.
RESERVED	22:16	0x0	RW	Reserved.
RESERVED	31:23	0x0	RSV	Reserved.

8.2.2.12 Statistic Registers

General notes:

- All statistics registers are cleared on read. In addition, they stick at 0xFF..F when the maximum value is reached.
- For the receive statistics, it should be noted that a packet is indicated as received if it passes the device filters and is placed into the packet buffer memory. A packet does not have to be DMA'ed to host memory in order to be counted as received.
- Due to divergent paths between interrupt generation and logging of relevant statistics counts, it might be possible to generate an interrupt to the system for a noteworthy event prior to the associated statistics count actually being incremented. This is extremely unlikely due to expected delays associated with the system interrupt collection and ISR delay, but might be an explanation for interrupt statistics values that do not quite make sense. Hardware guarantees that any event noteworthy of inclusion in a statistics count is reflected in the appropriate count within 1 μ s; a small time-delay prior to reading the statistics might be required to avoid a potential mismatch between and interrupt and its cause.
- If RSC is enabled, statistics are collected before RSC is applied to the packets.
- If TSO is enabled, statistics are collected after segmentation.
- All byte (octet) counters composed of two registers can be fetched by two consecutive 32-bit accesses while reading the low 32-bit register first or a single 64-bit access.
- All receive statistic counters count the packets and bytes before coalescing by the RSC logic.
- All receive statistic counters in the filter unit (listed later in this section) might count packets that might be dropped by the packet buffer or receive DMA. Same comment is valid for the byte counters associated with these packet counters: PRC64, PRC127, PRC255, PRC511, PRC1023, PRC1522, BPRC, MPRC, GPRC, RXNFGPC, RUC, and ROC.

8.2.2.12.1 Bad SFD Count - MBSDC (0x00004018)



Field	Bit(s)	Init.	Access Type	Description
MBSDC	31:0	0x0	RW	Counts the number of packets received with bad SFD. Does not count short packets not received.

8.2.2.12.2 CRC Error Count - CRCERRS (0x00004000)

Field	Bit(s)	Init.	Access Type	Description
CEC	31:0	0x0	RC	CRC Error Count. Counts the number of receive packets with CRC errors. In order for a packet to be counted in this register, it must be 64 bytes or greater (from <Destination Address> through <CRC>, inclusively) in length. This registers counts all packets received, regardless of L2 filtering and receive enable. This register doesn't counts packets with a bad SFD, error control byte or illegal code byte.



8.2.2.12.3 Illegal Byte Error Count - ILLERRC (0x00004004)

Field	Bit(s)	Init.	Access Type	Description
IBEC	31:0	0x0	RC	Illegal Byte Error Count. Counts the number of receive packets with illegal bytes errors. For example, there is an illegal symbol in the packet. This register counts all packets received, regardless of L2 filtering and receive enablement.

8.2.2.12.4 Error Byte Count - ERRBC (0x00004008)

Field	Bit(s)	Init.	Access Type	Description
EBC	31:0	0x0	RC	Error Byte Count. Counts the number of receive packets with error bytes. For example, there is an error symbol in the packet. This register counts all packets received, regardless of L2 filtering and receive enablement.

8.2.2.12.5 MAC Local Fault Count - MLFC (0x00004034)

Field	Bit(s)	Init.	Access Type	Description
MLFC	31:0	0x0	RC	Number of faults in the local MAC. This register is valid only when the link speed is 10 Gb/s.

8.2.2.12.6 MAC Remote Fault Count - MRFC (0x00004038)

Field	Bit(s)	Init.	Access Type	Description
MRFC	31:0	0x0	RC	Number of faults in the remote MAC. This register is valid only when the link speed is 10 Gb/s.

8.2.2.12.7 Receive Length Error Count - RLEC (0x00004040)



Field	Bit(s)	Init.	Access Type	Description
RLEC	31:0	0x0	RC	Number of packets with receive length errors. A length error occurs if an incoming packet length field in the MAC header doesn't match the packet length. To enable the receive length error count, the HLRG.RXLNGTHERREN bit needs to be set to 1b. This register counts all packets received, regardless of L2 filtering and receive enablement.

8.2.2.12.8 Switch Security Violation Packet Count - SSVPC (0x00008780)

Field	Bit(s)	Init.	Access Type	Description
SSVPC	31:0	0x0	RC	Switch Security Violation Packet Count. This register counts all Tx packets dropped. For example, due to switch security violations such as SA or VLAN anti-spoof filtering or packet that has (inner) VLAN that contradicts with PFVMVIR register definitions. Valid only in VMDq or IOV mode. This counter also includes traffic sent to the MC and blocked due to security violations.or packets dropped due to malicious events detection.

8.2.2.12.9 Link XON Received Count - LXONRXCNT (0x000041A4)

Field	Bit(s)	Init.	Access Type	Description
XONRXC	15:0	0x0	RC	Number of XON Packets Received. Sticks to 0xFFFF. XON packets can use the global address or the station address. This register counts any XON packet whether it is a legacy XON or a priority XON. Each XON packet is counted once even if it designated to a few priorities. If a priority FC packet contains both XOFF and XON, only the LXOFFRXCNT counter is incremented.
RESERVED	31:16	0x0	RSV	Reserved.

8.2.2.12.10 Link XOFF Received Count - LXOFFRXCNT (0x000041A8)

Field	Bit(s)	Init.	Access Type	Description
XOFFRXC	15:0	0x0	RC	Number of XOFF Packets Received. Sticks to 0xFFFF. XOFF packets can use the global address or the station address. This register counts any XOFF packet whether it is a legacy XOFF or a priority XOFF. Each XOFF packet is counted once even if it designated to a few priorities. If a priority FC packet contains both XOFF and XON, only this counter is incremented.
RESERVED	31:16	0x0	RSV	Reserved.

8.2.2.12.11 Priority XON Received Count - PXONRXCNT[n] (0x00004140 + 0x4*n, n=0...7)

Field	Bit(s)	Init.	Access Type	Description
XONRXC	15:0	0x0	RC	Number of XON packets received per UP. Sticks to 0xFFFF.
RESERVED	31:16	0x0	RSV	Reserved.



8.2.2.12.12 Priority XOFF Received Count - PXOFFRXCNT[n] (0x00004160 + 0x4*n, n=0...7)

Field	Bit(s)	Init.	Access Type	Description
XOFFRXC	15:0	0x0	RC	Number of XOFF packets received per UP. Sticks to 0xFFFF.
RESERVED	31:16	0x0	RSV	Reserved.

8.2.2.12.13 Packets Received [64 Bytes] Count - PRC64 (0x0000405C)

Field	Bit(s)	Init.	Access Type	Description
PRC64	31:0	0x0	RW	Number of good packets received that are 64 bytes in length (from <Destination Address> through <CRC>, inclusively). This register counts packets that pass L2 filtering regardless on receive enablement and does not include received flow control packets.

8.2.2.12.14 Packets Received [65-127 Bytes] Count - PRC127 (0x00004060)

Multicast: No

Field	Bit(s)	Init.	Access Type	Description
PRC127	31:0	0x0	RW	Number of packets received that are 65-127 bytes in length (from <Destination Address> through <CRC>, inclusively). This register counts packets that pass L2 filtering regardless on receive enablement and does not include received flow control packets.

8.2.2.12.15 Packets Received [128-255 Bytes] Count - PRC255 (0x00004064)

Field	Bit(s)	Init.	Access Type	Description
PRC255	31:0	0x0	RW	Number of packets received that are 128-255 bytes in length (from <Destination Address> through <CRC>, inclusively). This register counts packets that pass L2 filtering regardless on receive enablement and does not include received flow control packets.

8.2.2.12.16 Packets Received [256-511 Bytes] Count - PRC511 (0x00004068)

Field	Bit(s)	Init.	Access Type	Description
PRC511	31:0	0x0	RW	Number of packets received that are 256-511 bytes in length (from <Destination Address> through <CRC>, inclusively). This register counts packets that pass L2 filtering regardless on receive enablement and does not include received flow control packets.



8.2.2.12.17 Packets Received [512-1023 Bytes] Count - PRC1023 (0x0000406C)

Field	Bit(s)	Init.	Access Type	Description
PRC1023	31:0	0x0	RW	Number of packets received that are 512-1023 bytes in length (from <Destination Address> through <CRC>, inclusively). This register counts packets that pass L2 filtering regardless on receive enablement and does not include received flow control packets.

8.2.2.12.18 Packets Received [1024 to Max Bytes] Count - PRC1522 (0x00004070)

Field	Bit(s)	Init.	Access Type	Description
PRC1522	31:0	0x0	RW	Number of packets received that are 1024-Max bytes in length (from <Destination Address> through <CRC>, inclusively). This register counts packets that pass L2 filtering regardless on receive enablement and does not include received flow control packets. The maximum is dependent on the current receiver configuration and the type of packet being received. If a packet is counted in receive oversized count, it is not counted in this register. Due to changes in the standard for maximum frame size for VLAN tagged frames in 802.3, this device accepts packets that have a maximum length of 1522 bytes. The RMON statistics associated with this range has been extended to count 1522 byte long packets.

8.2.2.12.19 Broadcast Packets Received Count - BPRC (0x00004078)

Field	Bit(s)	Init.	Access Type	Description
BPRC	31:0	0x0	RO	Number of good (non-erred) broadcast packets received. This register does not count received broadcast packets when the broadcast address filter is disabled. The counter counts packets regardless on receive enablement.

8.2.2.12.20 Multicast Packets Received Count - MPRC (0x0000407C)

Field	Bit(s)	Init.	Access Type	Description
MPRC	31:0	0x0	RO	Number of good (non-erred) multicast packets received that pass L2 filtering (excluding broadcast packets). This register does not count received flow control packets. This registers counts packets regardless on receive enablement.

8.2.2.12.21 Good Packets Received Count - GPRC (0x00004074)

Field	Bit(s)	Init.	Access Type	Description
GPRC	31:0	0x0	RO	Number of good (non-erred) Rx packets (from the network) that pass L2 filtering and has legal length as defined by <i>LongPacketEnable</i> . This register counts packets regardless of receive enable. This register doesn't count flow control packets.



8.2.2.12.22 Good Octets Received Count Low - GORCL (0x00004088)

Field	Bit(s)	Init.	Access Type	Description
CNT_L	31:0	0x0	RC	Lower 32 bits of the good octets received counter. The GORCL and GORCH registers make up a logical 36-bit octet counter of the packets counted by the GPRC. This register includes bytes received in a packet from the <Destination Address> field through the <CRC> field, inclusively.

8.2.2.12.23 Good Octets Received Count High - GORCH (0x0000408C)

Field	Bit(s)	Init.	Access Type	Description
CNT_H	3:0	0x0	RC	Higher four bits of the good octets received counter.
RESERVED	31:4	0x0	RSV	Reserved.

8.2.2.12.24 Good Rx Non-filtered Packet Counter - RXNFGPC (0x000041B0)

Field	Bit(s)	Init.	Access Type	Description
GPC	31:0	0x0	RC	Number of good (non-erred with legal length) Rx packets (from the network) regardless of packet filtering and receive enablement.

8.2.2.12.25 Good Rx Non-Filter Byte Counter Low - RXNFBCL (0x000041B4)

Field	Bit(s)	Init.	Access Type	Description
BCL	31:0	0x0	RC	Low 32 bits of the 36-bit byte counter of good (non-erred) Rx packets that match the RXNFGPC. The counter counts all bytes from the <Destination Address> field through the <CRC> field, inclusively.

8.2.2.12.26 Good Rx Non-Filter Byte Counter High - RXNFBCH (0x000041B8)

Field	Bit(s)	Init.	Access Type	Description
BCH	3:0	0x0	RC	High four bits of the 36-bit byte counter associated to RXFBCL.
RESERVED	31:4	0x0	RSV	Reserved.



8.2.2.12.27 DMA Good Rx Packet Counter - RXDGPC (0x00002F50)

Field	Bit(s)	Init.	Access Type	Description
GPC	31:0	0x0	RC	Number of good (non-erred) Rx packets from the network posted to the host memory. In the case of packet replication (or mirrored), the counter counts each packet only once. The counter might count packets directed to all Rx queues or specific Rx queues as defined by the RXDSTATCTRL register.

8.2.2.12.28 DMA Good Rx Byte Counter Low - RXDGBCL (0x00002F54)X

Field	Bit(s)	Init.	Access Type	Description
GBCL	31:0	0x0	RC	Low 32 bits of the 36-bit byte counter of good (non-erred) Rx packets that match RXDGPC. The counter counts all bytes posted to the host before VLAN strip. Furthermore, bytes of RSC are counted before coalescing or DDP.

8.2.2.12.29 DMA Good Rx Byte Counter High - RXDGBCH (0x00002F58)

Field	Bit(s)	Init.	Access Type	Description
GBCH	3:0	0x0	RC	High four bits of the 36-bit byte counter associated to RXDGBCL.
RESERVED	31:4	0x0	RSV	Reserved.

8.2.2.12.30 DMA Duplicated Good Rx Packet Counter - RXDDPC (0x00002F5C)

Field	Bit(s)	Init.	Access Type	Description
GPC	31:0	0x0	RC	Number of replicated or mirrored packets that meet the RXDGPC conditions. The sum of RXDDPC and RXDGPC is the total good (non-erred) Rx packets from the network that are posted to the host. Note: The counter might count packets directed to all Rx queues or specific Rx queues as defined by the RXDSTATCTRL register.

8.2.2.12.31 DMA Duplicated Good Rx Byte Counter Low - RXDDBCL (0x00002F60)

Field	Bit(s)	Init.	Access Type	Description
GBCL	31:0	0x0	RC	Low 32 bits of the 36-bit byte counter of good (non-erred) Rx packets that match RXDDPC. The counter counts all bytes posted to the host before VLAN strip. Furthermore, bytes of RSC are counted before coalescing or DDP.



8.2.2.12.32 DMA Duplicated Good Rx Byte Counter High - RXDDBCH (0x00002F64)

Field	Bit(s)	Init.	Access Type	Description
GBCH	3:0	0x0	RC	High four bits of the 36-bit byte counter associated to RXDDBCL.
RESERVED	31:4	0x0	RSV	Reserved.

8.2.2.12.33 DMA Good Rx LPBK Packet Counter - RXLPBKPC (0x00002F68)

Field	Bit(s)	Init.	Access Type	Description
GPC	31:0	0x0	RC	Number of good (non-erred) Rx packets from a local VM posted to the host memory. In the case of packet replication (or mirrored), the counter counts each packet only once. The counter might count packets directed to all Rx queues or specific Rx queues as defined by the RXDSTATCTRL register. The counter is not affected by RSC since both functions are not supported for loopback traffic.

8.2.2.12.34 DMA Good Rx LPBK Byte Counter Low - RXLPKBCL (0x00002F6C)

Field	Bit(s)	Init.	Access Type	Description
GBCL	31:0	0x0	RC	Low 32 bits of the 36-bit byte counter of good (non-erred) Rx packets that match RXLPBKPC. The counter counts all bytes posted to the host before VLAN strip. Furthermore, bytes of RSC are counted before coalescing or DDP.

8.2.2.12.35 DMA Good Rx LPBK Byte Counter High - RXLPKBCH (0x00002F70)

Field	Bit(s)	Init.	Access Type	Description
GBCH	3:0	0x0	RC	High four bits of the 36-bit byte counter associated to RXLPKBCL.
RESERVED	31:4	0x0	RSV	Reserved.

8.2.2.12.36 DMA Duplicated Good Rx LPBK Packet Counter - RXDLPBKPC (0x00002F74)

Field	Bit(s)	Init.	Access Type	Description
GPC	31:0	0x0	RC	Number of replicated or mirrored packets that meet the RXLPBKPC conditions. The sum of RXDLPBKPC and RXLPBKPC is the total good (non-erred) Rx packets from a local VM posted to the host. Note: The counter might count packets directed to ALL Rx queues or specific Rx queues as defined by the RXDSTATCTRL register.



8.2.2.12.37 DMA Duplicated Good Rx LPBK Byte Counter Low - RXDLPBKBCL (0x00002F78)

Field	Bit(s)	Init.	Access Type	Description
GBCL	31:0	0x0	RC	Low 32 bits of the 36-bit byte counter of good (non-erred) Rx packets that match RXDLPBKPC. The counter counts all bytes posted to the host before VLAN strip. Furthermore, bytes of RSC are counted before coalescing or DDP.

8.2.2.12.38 DMA Duplicated Good Rx LPBK Byte Counter High - RXDLPBKBCH (0x00002F7C)

Field	Bit(s)	Init.	Access Type	Description
GBCH	3:0	0x0	RC	High four bits of the 36-bit byte counter associated to RXDLPBKBCL.
RESERVED	31:4	0x0	RSV	Reserved.

8.2.2.12.39 Good Packets Transmitted Count - GPTC (0x00004080)

Field	Bit(s)	Init.	Access Type	Description
GPTC	31:0	0x0	RC	Number of good packets transmitted. This register counts good (non-erred) transmitted packets. A good transmit packet is considered one that is 64 or more bytes (from the <Destination Address> through <CRC>, inclusively) in length. The register counts transmitted clear packets, secure packets and FC packets.

8.2.2.12.40 Good Octets Transmitted Count Low - GOTCL (0x00004090)

Field	Bit(s)	Init.	Access Type	Description
CNT_L	31:0	0x0	RC	Lower 32 bits of the good octets transmitted counter. See complete description in the next register GOTCH.

8.2.2.12.41 Good Octets Transmitted Count High - GOTCH (0x00004094)

Field	Bit(s)	Init.	Access Type	Description
CNT_H	3:0	0x0	RC	Higher four bits of the good octets transmitted counter. The GOTCL and GOTCH registers make up a logical 36-bit counter of successfully transmitted octets (in packets counted by GPTC). This register includes transmitted bytes in a packet from the <Destination Address> field through the <CRC> field, inclusively.
RESERVED	31:4	0x0	RSV	Reserved.



8.2.2.12.42 DMA Good Tx Packet Counter - TXDGPC (0x000087A0)

Field	Bit(s)	Init.	Access Type	Description
GPTC	31:0	0x0	RC	Number of Tx packets from host memory. This counter includes packets that are transmitted to the external network as well as packets that are transmitted only to local VMs. The latter occurs only in VT mode when the local switch is enabled. Dropped packets counted in SSVPC register are not counted here.

8.2.2.12.43 DMA Good Tx Byte Counter Low - TXDGBCL (0x000087A4)

Field	Bit(s)	Init.	Access Type	Description
BCL	31:0	0x0	RC	The low 32 bits of the 36-bit byte counter of the Tx packets that match the TXDGPC. The counter counts all bytes posted by the host and the VLAN (if bytes are added by hardware). Dropped packets counted in SSVPC register are not counted here.

8.2.2.12.44 DMA Good Tx Byte Counter High - TXDGBCH (0x000087A8)

Field	Bit(s)	Init.	Access Type	Description
BCH	3:0	0x0	RC	High four bits of the 36-bit byte counter associated to TXDGBCL.
RESERVED	31:4	0x0	RSV	Reserved.

8.2.2.12.45 Receive Undersize Count - RUC (0x000040A4)

Field	Bit(s)	Init.	Access Type	Description
RUC	31:0	0x0	RC	Receive Undersize Error. This register counts the number of received frames that are shorter than minimum size (64 bytes from <Destination Address> through <CRC>, inclusively), and had a valid CRC. This register counts packets regardless of L2 filtering and receive enablement.

8.2.2.12.46 Receive Fragment Count - RFC (0x000040A8)

Field	Bit(s)	Init.	Access Type	Description
RFC	31:0	0x0	RC	Number of receive fragment errors (frame shorted than 64 bytes from <Destination Address> through <CRC>, inclusively) that have bad CRC (this is slightly different from the Receive Undersize Count register). This register counts packets regardless of L2 filtering and receive enablement.



8.2.2.12.47 Receive Oversize Count - ROC (0x000040AC)

Field	Bit(s)	Init.	Access Type	Description
ROC	31:0	0x0	RC	Receive Oversize Error. This register counts the number of received frames that are longer than maximum size as defined by MAXFRS.MFS (from <Destination Address> through <CRC>, inclusively) and have valid CRC. This register counts packets regardless of L2 filtering and receive enablement.

8.2.2.12.48 Receive Jabber Count - RJC (0x000040B0)

Field	Bit(s)	Init.	Access Type	Description
RJC	31:0	0x0	RC	Number of Receive Jabber Errors. This register counts the number of received packets regardless of L2 filtering and receive enablement, and are greater than maximum size and have bad CRC (this is slightly different from the Receive Oversize Count register). The packets length is counted from <Destination Address> through <CRC>, inclusively. This register counts packets regardless of L2 filtering and receive enablement.

8.2.2.12.49 Management Packets Received Count - MNGPRC (0x000040B4)

Field	Bit(s)	Init.	Access Type	Description
MNGPRC	31:0	0x0	RO	Number of management packets received. This register counts the total number of packets received that pass the management filters management packets include RMCP and ARP packets. SFD errors and short packets are not counted, except that packets dropped because the management receives a jumbo packet or because the receive FIFO is full are counted.

8.2.2.12.50 Management Packets Dropped Count - MNGPDC (0x000040B8)

Field	Bit(s)	Init.	Access Type	Description
MPDC	31:0	0x0	RO	Number of management packets dropped. This register counts the total number of packets received that pass the management filters and then are dropped because: <ol style="list-style-type: none"> 1. The management receive FIFO is full or disabled. 2. A packet bigger than the maximal allowed size is received. 3. A packet length mismatch. The maximal allowed size is 1518 + the size of any recognized L2 header (VLAN, extended VLAN, E-tag, etc). Management packets include any packet directed to the manageability console (such as RMCP and ARP packets).



8.2.2.12.51 Total Octets Received Low - TORL (0x000040C0)

Field	Bit(s)	Init.	Access Type	Description
CNT_L	31:0	0x0	RC	Lower 32 bits of the total octets received counter. See complete description in the next register TORH.

8.2.2.12.52 Total Octets Received High - TORH (0x000040C4)

Field	Bit(s)	Init.	Access Type	Description
CNT_H	3:0	0x0	RC	Higher four bits of the total octets received counter. The TORL and TORH registers make up a logical 36-bit counter of the total received octets (in the packets counted by the TPR counter). This register includes bytes received in a packet from the <Destination Address> field through the <CRC> field, inclusively.
RESERVED	31:4	0x0	RSV	Reserved.

8.2.2.12.53 Total Packets Received - TPR (0x000040D0)

Field	Bit(s)	Init.	Access Type	Description
TPR	31:0	0x0	RC	Number of all packets received. This register counts the total number of all packets received. All packets received are counted in this register, regardless of their length, whether they are erred, regardless on L2 filtering and receive enablement but excluding flow control packets. TPR might count packets interrupted by link disconnect although they have a CRC error. This register doesn't count packets with SFD errors or packets that discarded by MAC.

8.2.2.12.54 Total Packets Transmitted - TPT (0x000040D4)

Field	Bit(s)	Init.	Access Type	Description
TPT	31:0	0x0	RC	Number of all packets transmitted. This register counts the total number of all packets transmitted. This register counts all packets, including standard packets, secure packets, FC packets, and manageability packets.

8.2.2.12.55 Packets Transmitted (64 Bytes) Count - PTC64 (0x000040D8)

Field	Bit(s)	Init.	Access Type	Description
PTC64	31:0	0x0	RC	Number of packets transmitted that are 64 bytes in length (from <Destination Address> through <CRC>, inclusively). This register counts all packets, including standard packets, secure packets, FC packets, and manageability packets.



8.2.2.12.56 Packets Transmitted [65-127 Bytes] Count - PTC127 (0x000040DC)

Field	Bit(s)	Init.	Access Type	Description
PTC127	31:0	0x0	RC	Number of packets transmitted that are 65-127 bytes in length (from <Destination Address> through <CRC>, inclusively). This register counts all packets, including standard packets, secure packets, and manageability packets.

8.2.2.12.57 Packets Transmitted [128-255 Bytes] Count - PTC255 (0x000040E0)

Field	Bit(s)	Init.	Access Type	Description
PTC255	31:0	0x0	RC	Number of packets transmitted that are 128-255 bytes in length (from <Destination Address> through <CRC>, inclusively). This register counts all packets, including standard packets, secure packets, and manageability packets.

8.2.2.12.58 Packets Transmitted [256-511 Bytes] Count - PTC511 (0x000040E4)

Field	Bit(s)	Init.	Access Type	Description
PTC511	31:0	0x0	RC	Number of packets transmitted that are 256-511 bytes in length (from <Destination Address> through <CRC>, inclusively). This register counts all packets, including standard packets, secure packets, and manageability packets.

8.2.2.12.59 Packets Transmitted [512-1023 Bytes] Count - PTC1023 (0x000040E8)

Field	Bit(s)	Init.	Access Type	Description
PTC1023	31:0	0x0	RC	Number of packets transmitted that are 512-1023 bytes in length (from <Destination Address> through <CRC>, inclusively). This register counts all packets, including standard packets, secure packets, and manageability packets.

8.2.2.12.60 Packets Transmitted [Greater than 1024 Bytes] Count - PTC1522 (0x000040EC)

Field	Bit(s)	Init.	Access Type	Description
PTC1522	31:0	0x0	RC	Number of packets transmitted that are 1024 or more bytes in length (from <Destination Address> through <CRC>, inclusively). This register counts all packets, including standard packets, secure packets, and manageability packets. Due to changes in the standard for maximum frame size for VLAN tagged frames in 802.3, the Intel® Xeon® Processor D-1500 Product Family LAN controller transmits packets that have a maximum length of 1522 bytes. The RMON statistics associated with this range has been extended to count 1522 byte long packets. This register counts all packets, including standard and secure packets.



8.2.2.12.61 Multicast Packets Transmitted Count - MPTC (0x000040F0)

Field	Bit(s)	Init.	Access Type	Description
MPTC	31:0	0x0	RC	Number of multicast packets transmitted. This register counts the number of multicast packets transmitted. This register counts all packets, including standard packets, secure packets, FC packets and manageability packets.

8.2.2.12.62 Broadcast Packets Transmitted Count - BPTC (0x000040F4)

Field	Bit(s)	Init.	Access Type	Description
BPTC	31:0	0x0	RC	Number of broadcast packets transmitted count. This register counts all packets, including standard packets, secure packets, FC packets and manageability packets

8.2.2.12.63 MAC Short Packet Discard Count - MSPDC (0x00004010)

Field	Bit(s)	Init.	Access Type	Description
MSPDC	31:0	0x0	RC	Number of MAC short packet discard packets received.

8.2.2.12.64 XSUM Error Count - XEC (0x00004120)

Field	Bit(s)	Init.	Access Type	Description
XEC	31:0	0x0	RC	Number of receive IPv4, TCP, UDP or SCTP XSUM errors.

Note: XSUM errors are not counted when a packet has any MAC error (CRC, length, under-size, over-size, byte error or symbol error).

8.2.2.12.65 Receive Queue Statistic Mapping Registers - RQSMR[n] (0x00002300 + 0x4*n, n=0...31)

These registers define the mapping of the receive queues to the per queue statistics. Several queues can be mapped to a single statistic register. Each statistic register counts the number of packets and bytes of all the queues that are mapped to that statistics. The registers counting Rx queue statistics are: QPRC, QBRC, and QPRDC. For example, setting RQSMR[0].Q_MAP[0] to 3 maps Rx queue 0 to the counters QPRC[3], QBRC[3], and QPRDC[3]. Setting RQSMR[2].Q_MAP[1] to 5 maps Rx queue 9 to the QPRC[5], QBRC[5], and QPRDC[5].

Field	Bit(s)	Init.	Access Type	Description
Q_MAP_0	3:0	0x0	RW	For each register n, Q_MAP[0] defines the per queue statistic registers that are mapped to Rx queue 4*n+0. (see examples that follow).
RESERVED	7:4	0x0	RSV	Reserved.



Field	Bit(s)	Init.	Access Type	Description
Q_MAP_1	11:8	0x0	RW	For each register n, Q_MAP[1] defines the per queue statistic registers that are mapped to Rx queue 4*n+1. (see examples that follow).
RESERVED	15:12	0x0	RSV	Reserved.
Q_MAP_2	19:16	0x0	RW	For each register n, Q_MAP[2] defines the per queue statistic registers that are mapped to Rx queue 4*n+2. (see examples that follow).
RESERVED	23:20	0x0	RSV	Reserved.
Q_MAP_3	27:24	0x0	RW	For each register n, Q_MAP[3] defines the per queue statistic registers that are mapped to Rx queue 4*n+3. (see examples that follow).
RESERVED	31:28	0x0	RSV	Reserved.

8.2.2.12.66 Rx DMA Statistic Counter Control - RXDSTATCTRL (0x00002F40)

Field	Bit(s)	Init.	Access Type	Description
QSEL	4:0	0x0	RW	The <i>Queue Select</i> field controls which Rx queues are considered for the DMA Good Rx and DMA duplicated counters as follows: 00000- 01111b = The counters relate to the same queues that are directed to the QPRC[QSEL] counter as defined by the RQSMR[n] registers. 10000b = The counters relate to all Rx queues. Else= Reserved.
RESERVED	31:5	0x0	RSV	Reserved.

8.2.2.12.67 Transmit Queue Statistic Mapping Registers - TQSM[n] (0x00008600 + 0x4*n, n=0...31)

These registers define the mapping of the transmit queues to the per queue statistics. Several queues can be mapped to a single statistic register. Each statistic register counts the number of packets and bytes of all the queues that are mapped to that statistics. The registers counting Tx queue statistics are: QPTC and QBTC.

Field	Bit(s)	Init.	Access Type	Description
Q_MAP_0	3:0	0x0	RW	For each register n, Q_MAP[0] defines the per queue statistic registers that are mapped to Tx queue 4*n+0.
RESERVED	7:4	0x0	RSV	Reserved.
Q_MAP_1	11:8	0x0	RW	For each register n, Q_MAP[1] defines the per queue statistic registers that are mapped to Tx queue 4*n+1.
RESERVED	15:12	0x0	RSV	Reserved.
Q_MAP_2	19:16	0x0	RW	For each register n, Q_MAP[2] defines the per queue statistic registers that are mapped to Tx queue 4*n+2.
RESERVED	23:20	0x0	RSV	Reserved.
Q_MAP_3	27:24	0x0	RW	For each register n, Q_MAP[3] defines the per queue statistic registers that are mapped to Tx queue 4*n+3.
RESERVED	31:28	0x0	RSV	Reserved.

**8.2.2.12.68 Queue Packets Received Count - QPRC[n] (0x00001030 + 0x40*n, n=0...15)**

Field	Bit(s)	Init.	Access Type	Description
PRC	31:0	0x0	RC	Number of packets received for the queue.

8.2.2.12.69 Queue Packets Received Drop Count - QPRDC[n] (0x00001430 + 0x40*n, n=0...15)

Field	Bit(s)	Init.	Access Type	Description
PRDC	31:0	0x0	RC	Number of receive packets dropped for the queue. Packets are dropped per queue in one of three cases: <ol style="list-style-type: none">1. Rx Queue is disabled in the RXDCTL[n] register (this includes packets sent to this queue due to the flow control drop no match mechanism or drop action).2. No free descriptors in the Rx queue while hardware is set to Drop En in the SRRCTL[n] register or in the PFQDE register.3. Packet size is larger than RLPML while RLPML_EN is set in the RXDCTL[n] register.

8.2.2.12.70 Queue Bytes Received Count Low - QBRC_L[n] (0x00001034 + 0x40*n, n=0...15)

Field	Bit(s)	Init.	Access Type	Description
BRC_L	31:0	0x0	RC	Lower 32 bits of the statistic counter. The QBRC_L[n] and QBRC_H[n] registers make up a logical 36-bit counter of received bytes that were posted to the programmed Rx queues of the packets counted by the QPRC[n]. The counter counts all bytes posted to the host before VLAN strip. Furthermore, bytes of RSC are counted before coalescing or DDP.

8.2.2.12.71 Queue Bytes Received Count High - QBRC_H[n] (0x00001038 + 0x40*n, n=0...15)

Field	Bit(s)	Init.	Access Type	Description
BRC_H	3:0	0x0	RC	Higher four bits of the statistic counter described in QBRC_L.
RESERVED	31:4	0x0	RSV	Reserved.

8.2.2.12.72 Queue Packets Transmitted Count - QPTC_ALIAS[n] (0x00006030 + 0x40*n, n=0...15; RC)

Fields definitions are the same as defined in Section 8.2.2.12.73.



8.2.2.12.73 Queue Packets Transmitted Count - QPTC[n] (0x00008680 + 0x4*n, n=0...15)

Field	Bit(s)	Init.	Access Type	Description
PTC	31:0	0x0	RC	Number of packets transmitted for the queue. A packet is considered as transmitted if it was forwarded to the MAC unit for transmission to the network and/or is accepted by the internal Tx to Rx switch enablement logic. Packets dropped due to anti-spoofing filtering, or traffic sent to the BMC and blocked due to security violations, or loopback packets that are rejected by the Tx to Rx switch and are not counted.

Note: Additional address(es): 0x06030 + 0x40*n, n=0...15

8.2.2.12.74 Queue Bytes Transmitted Count Low - QBTC_L[n] (0x00008700 + 0x8*n, n=0...15)

Field	Bit(s)	Init.	Access Type	Description
BTC_L	31:0	0x0	RC	Lower 32 bits of the statistic counter. The QBTC_L and QBTC_H registers make up a logical 36-bit counter of transmitted bytes of the packets counted by the matched QPTC counter. These registers count all bytes in the packets from the <Destination Address> field through the <CRC> field, inclusively. These registers must be accessed as two consecutive 32-bit entities while QBTC_L register is read first, or a single 64-bit read cycle. Each register is read cleared. In addition, it sticks at 0xFF..F to avoid overflow.

8.2.2.12.75 Queue Bytes Transmitted Count High - QBTC_H[n] (0x00008704 + 0x8*n, n=0...15)

Field	Bit(s)	Init.	Access Type	Description
BTC_H	3:0	0x0	RC	Higher four bits of the statistic counter described in QBTC_L.
RESERVED	31:4	0x0	RSV	Reserved.

8.2.2.12.76 Fiber Channel CRC Error Count - FCCRC (0x00005118)

Field	Bit(s)	Init.	Access Type	Description
CRC_CNT	15:0	0x0	RC	FC CRC Count. Count the number of packets with good Ethernet CRC and bad FC CRC.
RESERVED	31:16	X	RSV	Reserved.

8.2.2.12.77 BMC2OS Packets Sent by BMC - B2OSPC (0x000041C0)

This register counts the total number of transmitted packets sent from the manageability path that were sent to host. This includes packets received by the host and packet dropped in the device due to congestion conditions.

Counter is cleared when read by the software device driver. The counter is also cleared by PCIe reset and software reset. When reaching the maximum value, the counter does not wrap-around.



Field	Bit(s)	Init.	Access Type	Description
B2OSPC	31:0	0x0	RC	BMC2OS packets sent by BMC.

8.2.2.12.78 BMC2OS Packets Received by Host - B2OGPRC (0x00002F90)

This register counts the total number of packets originating from the BMC that reached the host. When the internal switch is enabled, each replication of a BMC to host packet is counted.

The counter is cleared when read by the software device driver. The counter is also cleared by PCIe reset and software reset. When reaching the maximum value, the counter does not wrap-around.

Field	Bit(s)	Init.	Access Type	Description
B2OGPRC	31:0	0x0	RC	BMC2OS packets received by host.

8.2.2.12.79 OS2BMC Packets Received by BMC - O2BGPTC (0x000041C4)

This register counts the total number of packets originating from the host that reached the NC-SI interface. the counter is cleared when read by the software device driver. The counter is also cleared by PCIe reset and software reset. When reaching maximum value counter does not wrap-around.

Field	Bit(s)	Init.	Access Type	Description
O2BGPTC	31:0	0x0	RC	OS2BMC Good Packets Transmitted Count. This includes packets sent from the host to the MC or consumed by the internal manageability engine.

8.2.2.12.80 OS2BMC Packets Transmitted by Host - O2BSPC (0x000087B0)

This register counts the total number of packets originating from the function that were sent to the manageability path. This includes packets received by the MC and packets dropped in the Intel® Xeon® Processor D-1500 Product Family LAN controller due to congestion conditions or due to anti-spoof check.

The counter is cleared when read by the software device driver. The counter is also cleared by PCIe reset and software reset. When reaching the maximum value, the counter does not wrap-around.

Field	Bit(s)	Init.	Access Type	Description
O2BPC	31:0	0x0	RC	OS2BMC Packets Transmitted Count.

8.2.2.12.81 Total Unicast Packets Received (BMC Copy) - BUPRC (0x00004180)

This register counts the number of good (no errors) unicast packets received from the network. This register does not count unicast packets received that fail to pass address filtering. This register does not count packets counted by the Missed Packet Count (MPC) register. This register also doesn't count flow control packets. Packets sent to the manageability engine are included in this counter.

This register is available to firmware only.



Field	Bit(s)	Init.	Access Type	Description
BUPRC	31:0	0x0	RC	Number of Unicast Packets Received.,

8.2.2.12.82 BMC Total Multicast Packets Received - BMPRC (0x00004184)

This register counts the same events as the MPRC register for BMC usage. This register is available to firmware only.

Field	Bit(s)	Init.	Access Type	Description
BMPRC	31:0	0x0	RC	Number of good (non-erred) multicast packets received that pass L2 filtering (excluding broadcast packets). This register does not count received flow control packets. This register counts packets regardless on receive enablement.

8.2.2.12.83 Total Broadcast Packets Received (BMC Copy) - BBPRC (0x00004188)

This register counts the same events as the BPRC register for BMC usage. This register is available to firmware only.

Field	Bit(s)	Init.	Access Type	Description
BBPRC	31:0	0x0	RC	Number of good (non-erred) broadcast packets received to the MC. This register does not count received broadcast packets when the broadcast address filter is disabled. The counter counts packets regardless on receive enablement.

8.2.2.12.84 Total Unicast Packets Transmitted (BMC copy) - BUPTC (0x0000418C)

This register counts the number of unicast packets transmitted. This register is available to firmware only.

Field	Bit(s)	Init.	Access Type	Description
BUPTC	31:0	0x0	RC	Number of unicast packets transmitted.

8.2.2.12.85 BMC Total Multicast Packets Transmitted - BMPTC (0x00004190)

This register counts the same events as the MPTC register for BMC usage. This register is available to firmware only.

Field	Bit(s)	Init.	Access Type	Description
BMPTC	31:0	0x0	RC	Number of multicast packets transmitted. This register counts the number of multicast packets transmitted. This register counts all packets, including standard packets, secure packets, FC packets and manageability packets.



8.2.2.12.86 Total Broadcast Packets Transmitted (BMC Copy) - BBPTC (0x00004194)

This register counts the same events as the BPTC register for BMC usage. This register is available to firmware only.

Field	Bit(s)	Init.	Access Type	Description
BBPTC	31:0	0x0	RC	Number of broadcast packets transmitted count. This register counts all packets, including standard packets, secure packets, FC packets and manageability packets.

8.2.2.12.87 BMC FCS Receive Errors - BRCERRS (0x00004198)

This register counts the same events as the CRCERRS register for BMC usage. This register is available to the firmware only.

Field	Bit(s)	Init.	Access Type	Description
BCEC	31:0	0x0	RC	CRC Error Count. Counts the number of receive packets with CRC errors. In order for a packet to be counted in this register, it must be 64 bytes or greater (from <Destination Address> through <CRC>, inclusively) in length. This register counts all packets received, regardless of L2 filtering and receive enablement.

8.2.2.12.88 BMC Pause XON Frames Received - BXONRXC (0x0000419C)

This register counts the same events as the XONRXC register for BMC usage. This register is available to firmware only.

Field	Bit(s)	Init.	Access Type	Description
BXONRXC	15:0	0x0	RC	Number of XON Packets Received. Sticks to 0xFFFF. XON packets can use the global address or the station address. This register counts any XON packet whether it is a legacy XON or a priority XON. Each XON packet is counted once even if it designated to a few priorities. If a priority FC packet contains both XOFF and XON, only the LXOFFRXCNT counter is incremented.
RESERVED	31:16	0x0	RC	Reserved.

8.2.2.12.89 BMC Pause XOFF Frames Received - BXOFFRXC (0x000041E0)

This register counts the same events as the XOFFRXC register for BMC usage. This register is available to firmware only.

Field	Bit(s)	Init.	Access Type	Description
BXOFFRXC	15:0	0x0	RC	Number of XOFF Packets Received. Sticks to 0xFFFF. XOFF packets can use the global address or the station address. This register counts any XOFF packet whether it is a legacy XOFF or a priority XOFF. Each XOFF packet is counted once even if it designated to a few priorities. If a priority FC packet contains both XOFF and XON, only this counter is incremented.
RESERVED	31:16	0x0	RC	Reserved.



8.2.2.12.90 BMC Pause XON Frames Transmitted - BXONTXC (0x000041E4)

This register counts the same events as the XONTXC register for BMC usage. This register is available to firmware only.

Field	Bit(s)	Init.	Access Type	Description
BXONTXC	15:0	0x0	RC	Number of XON Packets Transmitted. Sticks to 0xFFFF. XONTXC is incremented by one for each Link XON packet when MFLCN.RFCE is set.
RESERVED	31:16	0x0	RC	Reserved.

8.2.2.12.91 BMC Pause XOFF Frames Transmitted - BXOFFTXC (0x000041E8)

This register counts the same events as the XOFFTXC register for the BMC usage. This register is available to the firmware only.

Field	Bit(s)	Init.	Access Type	Description
BXOFFTXC	15:0	0x0	RC	Number of XOFF Packets Transmitted. Sticks to 0xFFFF. XOFFTXC is incremented by one for each Link XOFF packet when MFLCN.RFCE is set.
RESERVED	31:16	0x0	RC	Reserved.

8.2.2.12.92 Sideband Receive Dropped Packet Count - B2OSDPC (0x000041F0)

Field	Bit(s)	Init.	Access Type	Description
B2OSDPC	31:0	0x0	RW	Counts the number of packets received from the BMC interface receive that were dropped. The following causes can create a drop: 1. Memory buffer full. 2. Pause packet. 3. PT packet without SA match or requires an enable. 4. Reserved. 5. Length error. 6. Reject by specific interface (SMBus abort or RMII reject).

8.2.2.13 Wake-Up and Proxy Control Registers

8.2.2.13.1 Wake Up Control Register - WUC (0x00005800)

Field	Bit(s)	Init.	Access Type	Description
RESERVED	0	0b	RSV	Reserved.
PME_EN	1	0b	RW	PME_En. This bit is used by the software device driver to read the PME_En bit of the Power Management Control/Status Register (PMCSR) without writing to the PCIe configuration space. Writing a 1b to this bit clears it. Note: Software should not modify this bit while PME enablement is active.



Field	Bit(s)	Init.	Access Type	Description
PME_STATUS	2	0b	RW1C	PME_Status. This bit is set when the Intel® Xeon® Processor D-1500 Product Family LAN controller receives a wake-up event. It is the same as the PME_Status bit in the PMCSR. Writing a 1b to this bit clears it. The <i>PME_Status</i> bit in the PMCSR is also cleared.
RESERVED	3	0b	RSV	Reserved.
WKEN	4	1b	RW	WKEN. This bit can be cleared to disable the PE_WAKE_N pin assertion (= 0b) even if APM is enabled in the NVM. In this case, PMCSR and WUS wake-up statuses are invalid. Note: This bit should not be cleared while in ACPI mode.
RESERVED	31:5	0x0	RSV	Reserved.

Note: When AUX_PWR=0b, these bits are also reset by the assertion of PE_RST_N.

8.2.2.13.2 Wake Up Filter Control Register - WUFC (0x00005808)

Field	Bit(s)	Init.	Access Type	Description
LNKC	0	0b	RW	Link Status Change Wake Up Enable.
MAG	1	0b	RW	Magic Packet Wake Up Enable.
EX	2	0b	RW	Directed Exact Wake Up Enable.
MC	3	0b	RW	Directed Multicast Wake Up Enable. Setting this bit does not enable broadcast packets, which are enabled by the <i>BC</i> bit in this register.
BC	4	0b	RW	Broadcast Wake Up Enable.
ARP	5	0b	RW	ARP/IPv4 Request Packet Wake Up Enable.
IPV4	6	0b	RW	Directed IPv4 Packet Wake Up Enable.
IPV6	7	0b	RW	Directed IPv6 Packet Wake Up Enable.
RESERVED	14:8	0x0	RSV	Reserved.
NOTCO	15	0b	RW	Ignore TCO/Managements Packets For Wake up. 0b = Ignore all TCO/management packets for wake up, except packets that meet the criteria defined in the MNGONLY register via the <i>Host Enable</i> field. For example, criteria intended for the host as well as to the MC. 1b = Ignore all TCO/management packets for wake up. While in normal operation this ignore is forwarded to the host as well as to the MC.
FLX0	16	0b	RW	Flexible Filter 0 Enable: Controls the usage of the FHFT[0] register (0x9000).
FLX1	17	0b	RW	Flexible Filter 1 Enable. Controls the usage of the FHFT[1] register (0x9100).
FLX2	18	0b	RW	Flexible Filter 2 Enable. Controls the usage of the FHFT[2] register (0x9200).
FLX3	19	0b	RW	Flexible Filter 3 Enable. Controls the usage of the FHFT[3] register (0x9300).
FLX4	20	0b	RW	Flexible Filter 4 Enable. Controls the usage of the FHFT[0] register (0x9600).
FLX5	21	0b	RW	Flexible Filter 5 Enable. Controls the usage of the FHFT[5] register (0x9700).
FLX6	22	0b	RW	Flexible Filter 6 Enable. Controls the usage of the FHFT[6] register (0x9800).



Field	Bit(s)	Init.	Access Type	Description
FLX7	23	0b	RW	Flexible Filter 7 Enable. Controls the usage of the FHFT[7] register (0x9900).
RESERVED	30:24	0x0	RSV	Reserved.
FW_RST_WK	31	0b	RW	Enable Wake on Firmware Reset Assertion. When set, a firmware reset causes a system wake up that enables the software device driver to resend proxying information to firmware.

Note: This register is used to enable each of the pre-defined and flexible filters for wake-up support. A value of one means the filter is turned on, and a value of zero means the filter is turned off.

8.2.2.13.3 Wake Up Status Register - WUS (0x00005810)

Field	Bit(s)	Init.	Access Type	Description
LNKC	0	0b	RW1C	Link Status Changed. This bit can be set even if another event is already set.
MAG	1	0b	RW1C	Magic Packet Received.
EX	2	0b	RW1C	Directed Exact Packet Received. The packet's address matched one of the 16 pre-programmed exact values in the Receive Address registers.
MC	3	0b	RW1C	Directed Multicast Packet Received. The packet was a multicast packet that's hashed to a value that corresponds to one bit in the Multicast Table Array.
BC	4	0b	RW1C	Broadcast Packet Received.
ARP	5	0b	RW1C	ARP/IPv4 Request Packet Received.
IPV4	6	0b	RW1C	Directed IPv4 Packet Received.
IPV6	7	0b	RW1C	Directed IPv6 Packet Received.
MNG	8	0b	RW1C	Indicates that a manageability event that should cause a PME to happen.
RESERVED	15:9	0x0	RSV	Reserved.
FLX0	16	0b	RW1C	Flexible Filter 0 Match.
FLX1	17	0b	RW1C	Flexible Filter 1 Match.
FLX2	18	0b	RW1C	Flexible Filter 2 Match.
FLX3	19	0b	RW1C	Flexible Filter 3 Match.
FLX4	20	0b	RW1C	Flexible Filter 4 Match.
FLX5	21	0b	RW1C	Flexible Filter 5 Match.
FLX6	22	0b	RW1C	Flexible Filter 6 Match.
FLX7	23	0b	RW1C	Flexible Filter 7 Match.
RESERVED	30:24	0x0	RSV	Reserved.
FW_RST_WK	31	0b	RW1C	Wake Due to Firmware Reset Assertion Event. When set to 1b, indicates that asserting a firmware reset causes the system wake up so the software device driver can re-send proxying information to firmware. This bit can be set even if another event is already set.

Note: This register is used to record statistics about all wake-up packets received. If a packet matches multiple criteria than multiple bits are set by hardware. Software writing a 1b to any bit clears that bit.



Note: This register is not cleared when PE_RST_N is asserted. It is only cleared when Power Good Reset is de-asserted or when cleared by the software device driver.

8.2.2.13.4 IP Address Valid - IPAV (0x00005838)

The IP address valid indicates whether the IP addresses in the IP address table are valid.

Field	Bit(s)	Init.	Access Type	Description
V40	0	0b	RW	IPv4 Address 0 Valid. Loaded from NVM.
V41	1	0b	RW	IPv4 Address 1 Valid.
V42	2	0b	RW	IPv4 Address 2 Valid.
V43	3	0b	RW	IPv4 Address 3 Valid.
RESERVED	15:4	0x0	RSV	Reserved.
V60	16	0b	RW	IPv6 Address 0 Valid.
V61	17	0b	RW	IPv6 Address 1 Valid.
V62	18	0b	RW	IPv6 Address 2 Valid.
V63	19	0b	RW	IPv6 Address 3 Valid.
RESERVED	31:20	0x0	RSV	Reserved.

8.2.2.13.5 IPv4 Address Table - IP4AT[n] (0x00005840 + 0x8*n, n=0...3)

4 x IPv4 addresses for ARP/IPv4 request packet and directed IPv4 packet wake up. IPv4[0] is loaded from MIPAF words in the NVM.

Field	Bit(s)	Init.	Access Type	Description
IPV4ADDR	31:0	X	RW	IPv4 Address n, n = 0...3.

8.2.2.13.6 IPv6 Address Table - IP6AT[n] (0x00005880 + 0x4*n, n=0...3)

First IPv6 addresses for neighbor discovery packet filtering and directed IPv6 packet wake up.

Field	Bit(s)	Init.	Access Type	Description
IPV6ADDR	31:0	X	RW	4 x Register IPv6 Filter. Register n contains bytes 4*n up to 4*n+3 of the IPv6 address. The LSB of register 0 is first on the wire.

8.2.2.13.7 IPv6 Address Table Extended - IP6AT_EXT[n] (0x00005990 + 0x4*n, n=0...11)

3 x IPv6 addresses for neighbor discovery packet filtering.

Field	Bit(s)	Init.	Access Type	Description
IPV6ADDR	31:0	X	RW	4 x Register IPv6 Filter. Register n contains bytes 4*n up to 4*n+3 of the IPv6 address. The LSB of register 0 is first on the wire.



8.2.2.13.8 Wake Up Packet Memory (128 Bytes) - WUPM[n] (0x00005A00 + 0x4*n, n=0...31)

Field	Bit(s)	Init.	Access Type	Description
WUPD	31:0	X	RO	Wake Up Packet Data. This register is used to store the first 128 bytes of the wake up packet for software retrieval after the system wakes up. It should not be cleared by any reset including master reset.

8.2.2.13.9 Filter DW Even - FHFT_FILTER_DW_EVEN[n,m] (0x00009000 + 0x10*n + 0x100*m, n=0...15, m=0...3 and 0x00009600 + 0x10*(n-16) + 0x100*m, n=16...31, m=0...3)

Field	Bit(s)	Init.	Access Type	Description
FHFT_FILTER_0_DW0	31:0	X	RW	Even Dwords of the flex filter.

8.2.2.13.10 Filter DW Odd - FHFT_FILTER_DW_ODD[n,m] (0x00009004 + 0x10*n + 0x100*m, n=0...15, m=0...3 and 0x00009604 + 0x10*(n-16) + 0x100*m, n=16...31, m=0...3)

Field	Bit(s)	Init.	Access Type	Description
FHFT_FILTER_0_DW1	31:0	X	RW	Odd Dwords of the flex filter.

8.2.2.13.11 Filter Mask - FHFT_FILTER_MASK[n,m] (0x00009008 + 0x10*n + 0x100*m, n=0...15, m=0...3 and 0x00009608 + 0x10*(n-16) + 0x100*m, n=16...31, m=0...3)

Field	Bit(s)	Init.	Access Type	Description
FHFT_FILTER_0_MASK	7:0	X	RW	Bit mask for the eight bytes of the filter (bit per byte).
RESERVED	31:8	0x0	RSV	Reserved.

8.2.2.13.12 Filter Length - FHFT_FILTER_LENGTH[n,m] (0x0000900C + 0x10*n + 0x100*m, n=0...15, m=0...3 and 0x0000960C + 0x10*(n-16) + 0x100*m, n=16...31, m=0...3)

Field	Bit(s)	Init.	Access Type	Description
LENGTH	7:0	0x0	RW	Flex Filter Length The length field is only valid on the last Dword of the filter; all other length fields are reserved.
RESERVED	31:8	0x0	RSV	Reserved.



8.2.2.13.13 Proxying Status Register - PROXYS (0x00005F60)

This register is used to record statistics about all proxying packets received. If a packet matches multiple criteria then multiple bits could be set. Writing a 1b to any bit clears that bit. This register is not cleared when RST# is asserted. It is only cleared when Power Good Reset is deasserted or when cleared by the software device driver.

Note: If additional packets are received that matches one of the wake-up filters, after the original wake-up packet is received, the PROXYS register is updated with the matching filters accordingly.

Field	Bit(s)	Init.	Access Type	Description
RESERVED	1:0	0x0	RSV	Reserved.
EX	2	0b	RW1C	Exact packet received.
RESERVED	4:3	0x0	RSV	Reserved.
ARP_DIRECTED	5	0b	RW1C	ARP Request Packet with IP4AT Filter Match Received. When set to 1b, indicates a match on any ARP request packet that passed main filtering and target IP address also matches one of the valid IP4AT filters.
RESERVED	8:6	0x0	RSV	Reserved.
NS	9	0b	RW1C	IPv6 Neighbor Solicitation Received. When set to 1b, indicates a match on a NS packet that passed main filtering.
NS_DIRECTED	10	0b	RW1C	IPv6 Neighbor Solicitation With Directed DA Match Received. When set to 1b, indicates a match on a NS packet and target IP address also matches IPV6AT filter.
ARP	11	0b	RW1C	ARP Request Packet Received. When set to 1b, indicates a match on ARP request packet that passed main filtering.
MLD	12	0b	RW1C	IPv6 Multicast Listener Discovery (MLD) Packet Received. When set to 1b, indicates a match on any of the following MLD packet types that passed main filtering: <ol style="list-style-type: none"> Multicast Listener Query (ICMPv6 Type = decimal 130). Defined in MLDv1 and MLDv2. Multicast Listener Report (ICMPv6 Type = decimal 131). Defined in MLDv1 and MLDv2. Version 2 Multicast Listener Report Message (ICMPv6 Type = decimal 143). Defined in MLDv2 only.
RESERVED	31:13	0x0	RSV	Reserved.

8.2.2.13.14 Proxying Filter Control Register - PROXYFC (0x00005F64)

This register is used to enable each of the pre-defined filters for proxying support. This register is not cleared when RST# is asserted. It is only cleared when Power Good Reset is deasserted or when cleared by the software device driver.



Field	Bit(s)	Init.	Access Type	Description
PPROXYE	0	0b	RW	Port Proxying Enable. When set to 1b, proxying of packets is enabled when the Intel® Xeon® Processor D-1500 Product Family LAN controller is in the D3 low power state. Proxy information and requirements is passed by the software device driver to firmware via the Admin queue.
RESERVED	1	0b	RSV	Reserved.
EX	2	0b	RW	Exact Proxying Enable.
RESERVED	4:3	0x0	RSV	Reserved.
ARP_DIRECTED	5	0b	RW	ARP Request Packet and IP4AT match Proxy Enable. If set to 1b, forward to management for proxying on match of any ARP request packet that passed main filtering and target IP address also matches one of the valid IP4AT filters.
RESERVED	8:6	0x0	RSV	Reserved.
NS	9	0b	RW	IPv6 Neighbor Solicitation Proxy Enable. If set to 1b, forward to management for proxying on match of any NS packet (ICMPv6 type 135) that passed main filtering.
NS_DIRECTED	10	0b	RW	IPv6 Neighbor Solicitation and Directed DA Match Proxy Enable. If set to 1b, forward to management for proxying on match of NS packet and target IP address also matches valid IPV6AT filter.
ARP	11	0b	RW	ARP Request Packet Proxy Enable. If set to 1b, forward to management for proxying on match of any ARP request packet that passed main filtering.
MLD	12	0b	RW	IPv6 Multicast Listener Discovery (MLD) Proxy Enable. If set to 1b, forward to management for proxying on match of any of the following MLD packet types that passed main filtering: 1. Multicast listener query (ICMPv6 Type = decimal 130). Defined in MLDv1 and MLDv2. 2. Multicast listener report (ICMPv6 Type = decimal 131). Defined in MLDv1 and MLDv2. 3. Version 2 multicast listener report message (ICMPv6 Type = decimal 143). Defined in MLDv2 only.
RESERVED	14:13	0x0	RSV	Reserved.
NOTCO	15	0b	RW	Ignore TCO/Management Packets For Proxying. 0b = Ignore only TCO/management packets for proxying that meet the criteria defined in the MNGONLY register. for example, are intended only for the BMC and not the host. 1b = Ignore any TCO/management packets for proxying, even if in normal operation it is forwarded to the host in addition to the BMC.
RESERVED	31:16	0x0	RSV	Reserved.

8.2.2.14 Management Filters Registers

The Management Filters registers are RO for the host. These registers are initialized at Power Good Reset and can be loaded from the NVM by the manageability firmware.



8.2.2.14.1 Manageability Only Traffic - MNGONLY (0x00005864)

Field	Bit(s)	Init.	Access Type	Description
EXCLUSIVE_T O_MNG	7:0	0x0	RW	Exclusive to MNG. When set, indicates that packets forwarded by the manageability filters to manageability are not sent to the host. Bits 0...7 correspond to decision rules defined in registers MDEF[0...7] and MDEF_EXT[0...7].
RESERVED	31:8	0x0	RSV	Reserved.

8.2.2.14.2 Management VLAN TAG Value - MAVTV[n] (0x00005010 + 0x4*n, n=0...7)

Field	Bit(s)	Init.	Access Type	Description
VID	11:0	0x0	RW	Contains the VLAN ID that should be compared with the incoming packet if the corresponding bit in MFVAL.VLAN is set.
RESERVED	31:12	0x0	RSV	Reserved.

8.2.2.14.3 Management Flex UDP/TCP Ports - MFUTP[n] (0x00005030 + 0x4*n, n=0...7)

Field	Bit(s)	Init.	Access Type	Description
MFUTP_2N	15:0	0x0	RW	(2n)-th Management Flex UDP/TCP port.
MFUTP_2N_1	31:16	0x0	RW	(2n+1)-th Management Flex UDP/TCP port.

Note: Each 32-bit register (n=0,...,7) refers to two port filters (register 0 refers to ports 0&1, register 2 refers to ports 2&3, etc.). Note that SCTP packets do not match the MFUTP filters. MFUTP filters are programmed in network order.

8.2.2.14.4 Management Ethernet Type Filters - METF[n] (0x00005190 + 0x4*n, n=0...3)

Field	Bit(s)	Init.	Access Type	Description
ETYPE	15:0	0x0	RW	EtherType value to be compared against the L2 <i>EtherType</i> field in the Rx packet. Note: Appears in little endian order (high byte first on the wire).
RESERVED	29:16	0x0	RSV	Reserved.
POLARITY	30	0b	RW	0b = Positive filter. Filter enters the decision filters if a match occurred. 1b = Negative filter. Filter enters the decision filters if a match did not occur.
RESERVED	31	0b	RSV	Reserved.



8.2.2.14.5 Management Control Register - MANC (0x00005820)

Field	Bit(s)	Init.	Access Type	Description
FC_DISCARD	0	0b	RW	0b = Apply filtering rules to packets with Flow Control EtherType. 1b = Discard packets with Flow Control EtherType. Note: Flow control EtherType is 0x8808
NCSI_DISCARD	1	0b	RW	0b = Apply filtering rules to packets with NC-SI EtherType. 1b = Discard packets with NC-SI EtherType. Note: NC-SI EtherType is 0x88F8.
RESERVED	16:2	0x0	RSV	Reserved.
RCV_TCO_EN	17	0b	RW	Receive TCO Packets Enabled. When this bit is set, it enables the receive flow to the manageability block. This bit should be set only if at least one of MANC.EN_BMC2OS or MANC.EN_BMC2NET bits are set.
RESERVED	18	0b	RSV	Reserved.
RCV_ALL	19	0b	RW	Receive All Enable. When set, all packets are received from the wire and passed to the manageability block.
RESERVED	22:20	0x0	RW	Reserved.
EN_XSUM_FILTER	23	0b	RW	When set, this bit enables Xsum filtering to manageability. Meaning, only packets that pass the L3, L4 checksum are sent to the manageability block. Note that this capability is not provided for tunneled packets.
EN_IPV4_FILTER	24	0b	RW	Enable IPv4 Address Filters. When set, the last 128 bits of the MIPAF register are used to store four IPv4 addresses for IPv4 filtering. When cleared, these bits store a single IPv6 filter.
FIXED_NET_TYPE	25	0b	RW	Fixed Next Type. If set, only packets matching the net type defined by the <i>NET_TYPE</i> field passes to manageability. Otherwise, both tagged and un-tagged packets might be forwarded to the manageability engine.
NET_TYPE	26	0b	RW	Net Type: 0= Pass only un-tagged packets. 1= Pass only VLAN tagged packets. Valid only if FIXED_NET_TYPE is set.
MACSEC_MODE	27	0b	RW	When set, only packets that match one of the following threetwo conditions are forwarded to manageability: 1. The packet is a MACSec packet authenticated and/or decrypted adequately by hardware. 1. The packet Ethertype matches 0x88E5 (kay packets). 2. The packet Ethertype matches 0x888E (EAPOL packets). This bit only applies to network traffic and not OS-to-BMC traffic.
EN_BMC2OS	28	0b	RW	Enable BMC2OS and OS2BMC Traffic. 0b = The BMC can not communicate with the operating system. 1b = The BMC can communicate with the operating system. When cleared, BMC traffic is not forwarded to the operating system, even if the host MAC address filter and VLANs (RAH/L, MTA, VFTA and PFVLVF registers) indicate that it should. When cleared, the operating system traffic is not forwarded to the BMC even if the decision filters indicates it should. This bit does not impact the BMC-to-network traffic. Note: This bit can change while the host is sending or receiving traffic.



Field	Bit(s)	Init.	Access Type	Description
EN_BMC2NET	29	0b	RW	Enable BMC-to-network and network-to-BMC traffic. 0b = The BMC can not communicate with the network. 1b = The BMC can communicate with the network. When cleared, the BMC traffic is not forwarded to the network and the network traffic is not forwarded to the BMC even if the decision filters indicates it should. This bit does not impact the host-to-BMC traffic. Note: This bit can change while the host is sending or receiving traffic.
PROXYEN	30	0b	RW	Set by firmware to indicate that proxy offload is supported.
RESERVED	31	0b	RSV	Reserved.

8.2.2.14.6 Manageability Decision Filters - MDEF[n] (0x00005890 + 0x4*n, n=0..7)

Field	Bit(s)	Init.	Access Type	Description
EXACT_AND	3:0	0x0	RW	Exact. Controls the inclusion of exact MAC address 0-to-3. In the manageability filter decision (AND section), bit 0 corresponds to exact MAC address 0 (MMAL0 and MMAH0), etc.
BROADCAST_AND	4	0b	RW	Broadcast. Controls the inclusion of broadcast address filtering in the manageability filter decision (AND section).
VLAN_AND	12:5	0x0	RW	VLAN. Controls the inclusion of VLAN tag 0-to-7, respectively. In the manageability filter decision (AND section), bit 5 corresponds to VLAN tag0, etc.
IPV4_ADDRESS	16:13	0x0	RW	IPv4 Address. Controls the inclusion of IPV4 address 0-to-3 (MIPAF[3,n]), respectively in the manageability filter decision (AND section). Bit 13 corresponds to IPV4 address 0, etc. Notes: 1. This field is relevant only if the MANC.EN_IPv4_FILTER is set. 2. Supported only for network traffic.
IPV6_ADDRESS	20:17	0x0	RW	IPv6 Address. Controls the inclusion of IPV6 address 0-to-3, respectively (MIPAF[0:3,n]) in the manageability filter decision (AND section). Bit 17 corresponds to IPV6 address 0, etc. Notes: 1. Bit 20 is relevant only if MANC.EN_IPv4_FILTER is cleared. 2. Supported only for network traffic.
EXACT_OR	24:21	0x0	RW	Exact. Controls the inclusion of exact MAC address 0-to-3. In the manageability filter decision (OR section), bit 21 corresponds to exact MAC address 0 (MMAL0 and MMAH0), etc.
BROADCAST_OR	25	0b	RW	Broadcast. Controls the inclusion of broadcast address filtering in the manageability filter decision (OR section).
MULTICAST_AND	26	0b	RW	Multicast. Controls the inclusion of multicast address filtering in the manageability filter decision (AND section). Broadcast packets are not included by this bit.
ARP_REQUEST	27	0b	RW	ARP Request. Controls the inclusion of ARP request filtering in the manageability filter decision (OR section). Note: Supported only for network traffic.
ARP_RESPONSE	28	0b	RW	ARP Response. Controls the inclusion of ARP response filtering in the manageability filter decision (OR section). Note: Supported only for network traffic.



Field	Bit(s)	Init.	Access Type	Description
ND_134	29	0b	RW	Neighbor Discovery. Controls the inclusion of router advertisement neighbor discovery filtering in the manageability filter decision (OR section). Note: 1. Supported only for network traffic. Neighbor discovery types supported by this bit is 0x86 (134d) - router advertisement.
PORT_0X298	30	0b	RW	Port 0x298. Controls the inclusion of port 0x298 filtering in the manageability filter decision (OR section). Note: Supported only for network traffic.
PORT_0X26F	31	0b	RW	Port 0x26F. Controls the inclusion of port 0x26F filtering in the manageability filter decision (OR section). Note: Supported only for network traffic.

8.2.2.14.7 Manageability Decision Filters Ext - MDEF_EXT[n] (0x00005160 + 0x4*n, n=0...7)

Field	Bit(s)	Init.	Access Type	Description
L2_ETHERTYPE_AND	3:0	0x0	RW	L2 EtherType. Controls the inclusion of L2 EtherType filtering in the manageability filter decision (AND section).
L2_ETHERTYPE_OR	7:4	0x0	RW	L2 EtherType. Controls the inclusion of L2 EtherType filtering in the manageability filter decision (OR section).
FLEX_PORT	23:8	0x0	RW	Flex port. Controls the inclusion of Flex port filtering in the manageability filter decision (OR section). Bit 8 corresponds to flex port 0, etc.
FLEX_TCO	24	0b	RW	Flex TCO. Controls the inclusion of Flex TCO filtering in the manageability filter decision (OR section). Bit 24 corresponds to Flex TCO filter 0. Note: Supported only for network traffic.
ND_135	25	0b	RW	Neighbor Discovery. Controls the inclusion of neighbor solicitation neighbor discovery filtering in the manageability filter decision (OR section). Note: 1. Supported only for Network traffic. Neighbor Discovery types supported by this bit is 0x87 (135d) - neighbor solicitation.
ND_136	26	0b	RW	Neighbor Discovery. Controls the inclusion of neighbor advertisement neighbor discovery filtering in the manageability filter decision (OR section). Note: 1. Supported only for network traffic. Neighbor Discovery types supported by this bit is 0x88 (136d) - neighbor advertisement.
ND_137	27	0b	RW	Neighbor Discovery. Controls the inclusion of redirect neighbor discovery filtering in the manageability filter decision (OR section). Notes: 1. Supported only for network traffic. 2. Neighbor discovery types supported by this bit is 0x89 (137d) - redirect.
RESERVED	28	0b	RSV	Reserved.
MLD	29	0b	RW	MLD. Controls the inclusion of MLD packets. These are ICMPv6 packets with the following types: 130, 131, 132, 143.



Field	Bit(s)	Init.	Access Type	Description
APPLY_TO_NETWORK_TRAFFIC	30	0b	RW	0b = Do not apply this decision filter to traffic received from the network. 1b = Apply this decision filter to traffic received from the network.
APPLY_TO_HOST_TRAFFIC	31	0b	RW	0 = This decision filter does not apply to traffic received from the host. 1 = This decision filter applies to traffic received from the host.

8.2.2.14.8 BMC IP address Register - BMCIP[n] (0x00005050 + 0x4*n, n=0...3)

These registers contains the BMC IP address table.

Field	Bit(s)	Init.	Access Type	Description
IPADDR	31:0	0x0	RW	Four bytes of 16-byte destination IP address of the BMC. n=0 contains the MSB for an IPv6 IP address. n=3 contains an IPv4 IP address or the LSB for an IPv6 IP address. For an IPv4 address, BMCIP 0...2 must be written with zeros. Note: This field is defined in big endian (LSB is first on the wire).

8.2.2.14.9 BMC IP Valid Register - BMCIPVAL (0x00005060)

This register indicates the type of IP address stored in the IPVAL register and indicates if a valid address is stored.

Field	Bit(s)	Init.	Access Type	Description
IPADDR_TYPE	0	0b	RW	0b = IPv4 1b = IPv6.
IPADDR_VALID	1	0b	RW	0b = IP address in BMCIP is not valid. 1b = IP address in BMCIP is valid.
RESERVED	31:2	0x0	RSV	Reserved.

8.2.2.14.10 Manageability IP Address Filter - MIPAF[n,m] (0x000058B0 + 0x4*n + 0x10*m, n=0...3, m=0...3)

Field	Bit(s)	Init.	Access Type	Description
IP_ADDR	31:0	X	RW	Manageability IP Address Filters. For each n, m, m=0...3, n=0...3 while MANC.EN_IPv4_FILTER = 0b MIPAF[m,n] register holds Dword n of IPv6 filter m (4 x IPv6 filters). For each n, m, m=0...3, n=0...3 while MANC.EN_IPv4_FILTER = 1b, MIPAF[m,n] registers for m=0,1,2 is the same as the previous case (3 x IPv6 filters). MIPAF[3,n] registers also hold IPv4 filter n (4 x IPv4 filters). Note: These registers appear in big endian order (LSB, LS address is first on the wire).



8.2.2.14.11 Manageability Ethernet MAC Address Low - MMAL[n] (0x00005910 + 0x8*n, n=0...3)

Field	Bit(s)	Init.	Access Type	Description
MMAL	31:0	X	RW	Manageability Ethernet MAC Address Low. The lower 32 bits of the 48-bit Ethernet MAC address. Note: Appear in big endian order (LSB of MMAL is first on the wire).

8.2.2.14.12 Manageability Ethernet MAC Address High - MMAH[n] (0x00005914 + 0x8*n, n=0...3)

Field	Bit(s)	Init.	Access Type	Description
MMAH	15:0	X	RW	Manageability Ethernet MAC Address High. The upper 16 bits of the 48-bit Ethernet MAC address. Note: Appear in big endian order (MSB of MMAH is last on the wire).
RESERVED	31:16	0x0	RSV	Reserved. Reads as 0. Ignored on write.

8.2.2.14.13 FTFT Filter DW Even Words - FTFT_FILTER_EVEN[n] (0x00009400 + 0x10*n, n=0...15)

The Flexible TCO Filter Table registers (FTFT) contains a 128-byte pattern and a corresponding 128-bit mask array. If enabled, the first 128 bytes of the received packet are compared against the non-masked bytes in the FTFT register.

Each 128-byte filter is composed of 32 Dword entries (FTFT_FILTER_ODD and FTFT_FILTER_EVEN), where each 2 Dwords are accompanied by an 8-bit mask (FTFT_FILTER_MASK), one bit per filter byte. 15:8], etc. The *Mask* field is set so that bit 0 in the mask masks byte 0, bit 1 masks byte 1 etc. A value of one in the mask field means that the appropriate byte in the filter should be compared to the appropriate byte in the incoming packet.

The FTFT_FILTER_LENGTH register indicates the number of bytes to compare.

Field	Bit(s)	Init.	Access Type	Description
FTFT_FILTER_0_DW0	31:0	X	RW	Even Dword filter value.

Note: The mask field must be 8-byte aligned even if the length field is not 8-byte aligned as hardware implementation compares eight bytes at a time so it should get extra masks until the end of the next Qword. Any mask bit that is located after the length should be set to zero indicating no comparison should be done.

Note: In case the actual length that is defined by the length field register and the mask bits is not 8-byte aligned there might be a case that a packet which is shorter than the actual required length passes the flexible filter. This might happen due to comparison of up to seven bytes that come after the packet but are not a real part of the packet.

Note: The last Dword of each filter contains a length field defining the number of bytes from the beginning of the packet compared by this filter. If actual packet length is less than length



specified by this field, the filter fails. Otherwise, it depends on the result of actual byte comparison. The value should not be greater than 128.

Note: FTFT registers are configured by firmware.

8.2.2.14.14 FTFT Filter DW Odd Words - FTFT_FILTER_ODD[n] (0x00009404 + 0x10*n, n=0...15)

Field	Bit(s)	Init.	Access Type	Description
FHFT_FILTER_0_DW1	31:0	X	RW	Odd Dword filter value.

8.2.2.14.15 FTFT Filter Mask - FTFT_FILTER0_MASK[n] (0x00009408 + 0x10*n, n=0...15)

Field	Bit(s)	Init.	Access Type	Description
FHFT_FILTER_0_MASK	7:0	X	RW	Mask for the eight bytes of the filter (bit per byte).
RESERVED	31:8	0x0	RSV	Reserved.

8.2.2.14.16 FTFT Filter Length - FTFT_FILTER0_LENGTH[n] (0x0000940C + 0x10*n, n=0...15)

Field	Bit(s)	Init.	Access Type	Description
LENGTH	7:0	0x0	RW	Note: Length field is valid only on the last Dword of the filter all other fields are reserved.
RESERVED	31:8	0x0	RSV	Reserved.

8.2.2.15 Manageability (ARC Subsystem) Host Interface

Host interface to the ARC subsystem is described in the Manageability Host Interface section in the Manageability section.

8.2.2.15.1 Software Semaphore Register - SWSM (0x00010140)

Note: This register is shared by both LAN ports.

Field	Bit(s)	Init.	Access Type	Description
SMBI	0	0b	RW	Semaphore Bit. This bit is set by hardware. When this register is read by the software device driver, one of the two PCIe functions is cleared when the host driver writes 0b to it. The first time this register is read, the value is 0b. In the next read the value is 1b (hardware mechanism). The value remains 1b until the software device driver clears it. This bit can be used as a semaphore between the two device's drivers. This bit is cleared on PCIe reset.
RESERVED	31:1	0x0	RSV	Reserved.



8.2.2.15.2 Firmware Semaphore Register - FWSM (0x00010148)

Note: This register is shared by both LAN ports.

Field	Bit(s)	Init.	Access Type	Description
RESERVED	0	0b	RSV	Reserved.
FLASH_UPDATE_ENABLED	1	0b	RW	If set, the Flash update functionality is active.
PASS_THROUGH_ENABLED	2	0b	RW	If set, the pass through functionality is active.
PROXY_ENABLED	3	0b	RW	If set, the proxy functionality is active.
HOST_INTERFACE_ENABLED	4	0b	RW	If set, the host interface functionality is active.
FWSM	5	0b	RW	NVM Recovery Mode. Defines the operation mode of firmware. 0b = Normal operation mode. 1b = NVM recovery mode.
EEP_RELOAD_IND	6	0b	RW	NVM Reloaded Indication. Set to 1b after firmware reloaded the NVM. Cleared by firmware at firmware reset only.
RESERVED	14:7	0x0	RSV	Reserved.
FW_VAL_BIT	15	0b	RW	Firmware Valid Bit. Hardware clears this bit in reset de-assertion so software can know firmware modes (bits 1-4) are invalid. Firmware should set it to 1b when it is ready (end of boot sequence).
RESERVED	18:16	0x0	RW	Reserved
EXT_ERR_IND	24:19	0x0	RW	External error indication firmware uses this register to store the reason that the firmware has reset / clock gated (such as NVM, patch corruption). Possible values: 0x00 = No error. 0x01 = Unspecified error. 0x02 = No manageability (no NVM). 0x03 = TCO isolate mode active. 0x05 = Shadow RAM dump fault. 0x06 = Bad Flash contents (for FW/RO update). 0x10 = NVM CRC error in the test configuration module. 0x11 = NVM CRC error in the common FW parameters module. 0x12 = NVM CRC error in the PT LAN 0 configuration module. 0x13 = NVM CRC error in the sideband configuration module. 0x14 = NVM CRC error in the flexible TCO filter configuration module. 0x15 = NVM CRC error in PT LAN 1 configuration module. 0x16 = NVM CRC error in OEM support structure module. 0x20 = Management memory parity/ECC error. 0x21 = SR ECC error. 0x22-0x3E = Reserved. 0x3F = Reserved (max error value).
RESERVED	25	0b	RSV	Reserved. Formally PCIE_CONFIG__ERR_IND (PCIe Configuration Error Indication). Set to 1b by firmware when it fails to configure the PCIe interface. Cleared by firmware upon successful configuration of PCIe interface.



Field	Bit(s)	Init.	Access Type	Description
PHY0_CONFIG_ERR_IND	26	0b	RW	PHY0 Configuration Error Indication. Set to 1b by firmware when it fails to configure PHY of LAN0. Cleared by firmware upon successful configuration of PHY of LAN0.
PHY1_CONFIG_ERR_IND	27	0b	RW	PHY1 Configuration Error indication. Set to 1b by firmware when it fails to configure PHY of LAN1. Cleared by firmware upon successful configuration of PHY of LAN1
RESERVED	30:28	0x0	RSV	Reserved.
FACTORY_MAC_ADDRESS_RESTORED	31	0b	RW	When set, it indicates to software that the factory MAC address was successfully restored after the last power-up event. This bit is cleared by the device at power up.

Note: This register should be written only by manageability firmware. The software device driver should only read this register. The firmware ignores the NVM semaphore in operating system hung states. Bits 15:0 are cleared on firmware reset.

8.2.2.15.3 Software-Firmware Synchronization - SW_FW_SYNC (0x00010160)

Each bit represents different software semaphore agreed between software and firmware as follows. Bits 4:0 and bits 10:12 are owned by software while bits 9:5 and bits 13:15 are owned by firmware. Note that hardware does not lock access to these bits.

Field	Bit(s)	Init.	Access Type	Description
SW_NVM_SM	0	0b	RW	If set, NVM access is owned by software.
SW_PHY0_SM	1	0b	RW	If set, PHY 0 access is owned by software.
SW_PHY1_SM	2	0b	RW	If set, PHY 1 access is owned by software.
SW_MAC_CSR_SM	3	0b	RW	If set, MAC CSR access is owned by software.
RESERVED	4	0b	RSV	Reserved.
FW_NVM_SM	5	0b	RW	If set, NVM access is owned by firmware.
FW_PHY0_SM	6	0b	RW	If set, PHY 0 access is owned by firmware.
FW_PHY1_SM	7	0b	RW	If set, PHY 1 access is owned by firmware.
FW_MAC_CSR_SM	8	0b	RW	If set, MAC CSR access is owned by firmware.
NVM_UPDATE_STARTED	9	0b	RW	If set, NVM update started, Software should not write to NVM.
SW_MNG_SM	10	0b	RW	If set, manageability host interface is owned by software.
SW_I2C0_SM	11	0b	RW	If set, I ² C of port 0 is owned by software.
SW_I2C1_SM	12	0b	RW	If set, I ² C of port 1 is owned by software.
FW_I2C0_SM	13	0b	RW	If set, I ² C of port 0 is owned by firmware.
FW_I2C1_SM	14	0b	RW	If set, I ² C of port 1 is owned by firmware.
RESERVED	30:15	0x0	RSV	Reserved. for future use.
REGSMP	31	0b	RW	Register Semaphore. This bit is used to semaphore the access to this register between the firmware and software with no hardware enforcement. When the bit value is 0b and the register is read, the returned value is zero and the bit setting reverts to 1b (for the next read cycle). Writing 0b to this bit clears it. A software driver that reads this register and gets the value of zero for this bit locks the access to this register until it clears this bit.



Note: This register is shared by both LAN ports. Refer to the Software and Firmware Synchronization section for more details on software and firmware synchronization.

8.2.2.15.4 Software-Firmware Synchronization - SW_FW_SYNC_MIRR (0x00010164; RO)

This register is a mirror of the SW_FW_SYNC register enabling to read it without taking the semaphore bit.

Fields definitions are the same as defined on [Section 8.2.2.15.3](#).

8.2.2.15.5 Firmware Resets Count - FWRESETCNT (0x00015F40)

Field	Bit(s)	Init.	Access Type	Description
FWRESETCNT	31:0	0x0	RW	Firmware resets count. Updated by hardware. Stuck at 0xFFFF,FFFF.

8.2.2.15.6 Host Interface Control Register - HICR (0x00015F00)

Field	Bit(s)	Init.	Access Type	Description
EN	0	0b	RW	Enable. When set, it indicates that a RAM area is provided for software device driver accesses. This bit is read only for the software device driver.
C	1	0b	RW	Command. The software device driver sets this bit when it has finished putting a command block in the ARC internal data RAM. This bit should be cleared by firmware when the command's processing is completed. Setting this bit causes an interrupt to the ARC.
SV	2	0b	RW	Status Valid. Indicates that there is a valid status in CSR area that the software device driver can read. 0x1 = status valid. 0x0 = status not valid. The value of the bit is valid only when the C bit is cleared. Only the software device driver reads this bit.
SMB	3	0b	RW	Semaphore Bit. This bit is set when this register is read by the host driver and cleared when the host driver writes zero to it. This bit can be used as semaphore between the two LAN drivers in the Intel® Xeon® Processor D-1500 Product Family LAN controller. This bit is reset at PCIe reset.
SWSMB	4	0b	RW	Software/Firmware Semaphore - Software Bit. This bit is set only by host software (read only to firmware). The bit is not set if bit 8 in the Software Status register (CSR 0x15F10) is set. This bit can be used by the software and firmware to synchronize mutual resources. This bit is cleared by hardware on a core reset or the manageability watch dog expires scenario.
RESERVED	5	0b	RSV	Reserved.
RESERVED	6	0b	RW	Reserved.
RESERVED	7	0b	RW	Reserved.
RESERVED	31:8	0x0	RSV	Reserved.



8.2.2.15.7 Host ARC Data RAM - ARCRAM[n] (0x00015800 + 0x4*n, n=0...447)

Field	Bit(s)	Init.	Access Type	Description
RAM_SPACE	31:0	X	RW	RAM space area that spans on the CSR space: 0x15800 - 0x15EFC.

8.2.2.15.8 System Time Register Low - SYSTIMEL (0x00008C0C)

Field	Bit(s)	Init.	Access Type	Description
STL	29:0	0x0	RW	System Time LSB register, defined in ns. Writes of values above a second (above 0x3B9AC9FF) are ignored.
RESERVED	31:30	0x0	RW	Reserved.

8.2.2.15.9 System Time Register High - SYSTIMEH (0x00008C10)

Field	Bit(s)	Init.	Access Type	Description
STH	31:0	0x0	RW	System time MSB register, defined in seconds.

8.2.2.15.10 System Time Register Residue - SYSTIMR (0x00008C58)

Field	Bit(s)	Init.	Access Type	Description
SYSTIMR	31:0	0x0	RW	System time residue value (defined in $2^{(-32)}$ ns resolution).

8.2.2.15.11 Increment Attributes Register - TIMEINCA (0x00008C14)

Field	Bit(s)	Init.	Access Type	Description
INCVALUE	30:0	0x0	RW	Increment the value to the SYSTIME registers on each 12.5 ns.
ISGN	31	0b	RW	Increment Sign. 0b = Each 12.5 ns cycle add to SYSTIME a value of 12.5 ns + Incvalue * $2^{(-32)}$ ns. 1b = Each 12.5 ns cycle add to SYSTIME a value of 12.5 ns - Incvalue * $2^{(-32)}$ ns.

8.2.2.15.12 Time Adjustment Offset Register - TIMADJ (0x00008C18)

Field	Bit(s)	Init.	Access Type	Description
TADJL	30:0	0x00	RW	Time adjustment value (defined in ns units).
SIGN	31	0b	RW	Sign (0b=+, 1b=-)



8.2.2.15.13 Time Sync Interrupt Mask Register - TSIM (0x00008C68)

Field	Bit(s)	Init.	Access Type	Description
SYS_WARP	0	0b	RW	SYSTIM Warp Around Mask. 0= No Interrupt generated when TSICR.SWARP is set. 1= Interrupt generated when TSICR.SWARP is set.
TXTS	1	0b	RW	Transmit Time Stamp Mask. 0b = No Interrupt generated when TSICR.TXTS is set. 1b = Interrupt generated when TSICR.TXTS is set.
RXTS	2	0b	RW	Receive Time Stamp Mask. 0b = No Interrupt generated when TSICR.RXTS is set. 1b = Interrupt generated when TSICR.RXTS is set.
TT0	3	0b	RW	Target Time 0 Trigger Mask. 0b = No Interrupt generated when TSICR.TT0 is set. 1b = Interrupt generated when TSICR.TT0 is set.
TT1	4	0b	RW	Target Time 1 Trigger Mask. 0b = No Interrupt generated when TSICR.TT1 is set. 1b = Interrupt generated when TSICR.TT1 is set.
AUTT0	5	0b	RW	Auxiliary Timestamp 0 Taken Mask. 0b = No Interrupt generated when TSICR.AUTT0 is set. 1b = Interrupt generated when TSICR.AUTT0 is set.
AUTT1	6	0b	RW	Auxiliary Time Stamp 1 Taken Mask. 0b = No Interrupt generated when TSICR.AUTT1 is set. 1b = Interrupt generated when TSICR.AUTT1 is set.
TADJ	7	0b	RW	Time Adjust 0 Done Mask. 0b = No Interrupt generated when TSICR.TADJ is set. 1b = Interrupt generated when TSICR.TADJ is set.
RESERVED	31:8	0x0	RSV	Reserved. Write 0x0, ignore on read.

8.2.2.15.14 Time Sync Interrupt Cause Register - TSICR (0x00008C60)

Note: The value of this register is always read as 0x0. Once ICR.Time_Sync is set, the internal value of the register should be cleared by a write of 1b to all bits or cleared by read to enable reception of an additional ICR.Time_Sync interrupt.

Field	Bit(s)	Init.	Access Type	Description
SYS_WARP	0	0b	RW1C	SYSTIMEL Warp Around. Set when SYSTIME ns counter warps around (SYSTIMEL). Warp around occurrence can be used by software to update software time. This event should happen every second.
TXTS	1	0b	RW1C	Transmit time stamp set when new a time stamp is loaded into the TXSTMP register.
RXTS	2	0b	RW1C	Receive Time Stamp. Set when a new time stamp is loaded into the RXSTMP register.
TT0	3	0b	RW1C	Target Time 0 Trigger. Set when a target time 0 (TRGTTIML/H0) trigger occurs.
TT1	4	0b	RW1C	Target Time 1 Trigger. Set when a target time 1 (TRGTTIML/H1) trigger occurs.



Field	Bit(s)	Init.	Access Type	Description
AUTT0	5	0b	RW1C	Auxiliary Time Stamp 0 Taken. Set when a new time stamp is loaded into the AUXSTMP 0 (auxiliary timestamp 0) register.
AUTT1	6	0b	RW1C	Auxiliary Time Stamp 1 Taken. Set when a new time stamp is loaded into the AUXSTMP 1 (auxiliary timestamp 1) register.
TADJ	7	0b	RW1C	Time Adjust 0 Done. Set when time adjust to clock out 0 or 1 completes.
RESERVED	31:8	0x0	RSV	Reserved. Write 0, ignore on read.

8.2.2.15.15 Tx Time Sync Control Register - TSYNCTXCTL (0x00008C00)

Field	Bit(s)	Init.	Access Type	Description
TXTT	0	0b	ROS	Tx time Stamp Valid (= 1b when a valid value for Tx time stamp is captured in the Tx time stamp register; clear by read of Tx timestamp register TXSTMPH).
RESERVED	3:1	0x0	RSV	Reserved.
EN	4	0b	RW	Enable Tx Time Stamp. 0x0 = Time stamping disabled. 0x1 = Time stamping enabled.
RESERVED	31:6	0x0	RSV	Reserved.

8.2.2.15.16 Tx Timestamp Value Low - TXSTMPL (0x00008C04)

Field	Bit(s)	Init.	Access Type	Description
TXSTMPL	29:0	0x0	RO	Tx Time Stamp LSB Value (defined in ns units). The value in this field wraps around at 999999999 decimal.
RESERVED	31:30	0x0	RO	Reserved.

8.2.2.15.17 Tx Timestamp Value High - TXSTMPH (0x00008C08)

Field	Bit(s)	Init.	Access Type	Description
TXSTMPH	31:0	0x0	RO	Tx times tamp MSB value defined in seconds.

8.2.2.15.18 Rx Time Sync Control register - TSYNCRXCTL (0x00005188)

Field	Bit(s)	Init.	Access Type	Description
RXTT	0	0b	RO	Rx Time Stamp Valid (= 1b when a valid value for Rx timestamp is captured in the Rx timestamp register; clear by read of Rx timestamp register RXSTMPH).



Field	Bit(s)	Init.	Access Type	Description
TYPE	3:1	0x0	RW	Type Of Packets To Timestamp 000b = Time stamp L2 (V2) packets only (sync or Delay_req depends on RXMTRL.MSGT and packets with message ID 2 and 3). 001b = Time stamp L4 (V1) packets only (sync or Delay_req depends on RXMTRL.CTRLT). 010b = Time stamp V2 (L2 and L4) packets (sync or Delay_req depends on RXMTRL.MSGT and packets with message ID 2 and 3). 100b = Time stamp all packets. 101b - Time stamp V2 packets in which message ID bit 3 is zero, which means time stamp all event packets. 011b, 110b and 111b = Reserved.
EN	4	0b	RW	Enable Rx Timestamp In TXSTMPH/L Registers. 0x0 = Time stamping to TXSTMPH/L disabled. 0x1 = Time stamping to TXSTMPH/L enabled.
RESERVED	20:5	0x0	RSV	Reserved.
TSIP_UT_EN	23	0b	RW	Defines if untagged packets are appended a time stamp.
TSIP_UP_EN	31:24	0x0	RW	Defines which UP time stamp is appended to the received packet (per UP bitmap). For example, to require a time stamp on packets received with UP = 0 or UP = 3, bits 24 and 26 should be set.

8.2.2.15.19 Rx Timestamp Low - RXSTMPL (0x000051E8)

Field	Bit(s)	Init.	Access Type	Description
RXSTMPL	29:0	0x0	RO	Rx Time Stamp LSB Value (defined in ns units). The value in this field wraps around at 999999999 decimal.
RESERVED	31:30	0x0	RO	Reserved.

8.2.2.15.20 Rx Timestamp High - RXSTMPH (0x000051A4)

Field	Bit(s)	Init.	Access Type	Description
RXSTMPH	31:0	0x0	RO	Rx time stamp MSB value. defined in seconds.

8.2.2.15.21 Rx Message Type Register Low - RXMTRL (0x00005120)

Field	Bit(s)	Init.	Access Type	Description
CTRLT	7:0	0x0	RW	V1 control to time stamp.
MSGT	15:8	0x0	RW	V2 message ID to time stamp.
UDPT	31:16	0x13F	RW	UDP port number to time stamp.

8.2.2.15.22 Time Sync SDP Configuration Register - TSSDP (0x0000003C)

This register defines the assignment of SDP pins to the time sync auxiliary capabilities.




Field	Bit(s)	Init.	Access Type	Description
AUX0_SDP_SEL	1:0	0x0	RW	Select one of the SDPs to serve as the trigger for auxiliary time stamp (AUXSTMPL0 and AUXSTMPH0 registers). 00b = SDP0 is assigned. 01b = SDP1 is assigned. 10b = SDP2 is assigned. 11b = SDP3 is assigned.
AUX0_TS_SDP_EN	2	0b	RW	When set, indicates that one of the SDPs can be used as an external trigger to aux timestamp 0. Note that if this bit is set to one of the SDP pins, the corresponding pin should be configured to input mode using SPD_DIR).
AUX1_SDP_SEL	4:3	0x0	RW	Select one of the SDPs to serve as the trigger for auxiliary time stamp 1 (in AUXSTMPL1 and AUXSTMPH1 registers). 00b = SDP0 is assigned. 01b = SDP1 is assigned. 10b = SDP2 is assigned. 11b = SDP3 is assigned.
AUX1_TS_SDP_EN	5	0b	RW	When set indicates that one of the SDPs can be used as an external trigger to Aux timestamp 1 (note that if this bit is set to one of the SDP pins, the corresponding pin should be configured to input mode using SPD_DIR)
TS_SDP0_SEL	7:6	0x0	RW	SDP0 Allocation To Tsync Event. When TS_SDP0_EN is set, these bits select the Tsync event that is routed to SDP0: 00b = Target time 0 is output on SDP0. 01b = Target time 1 is output on SDP0. 10b = Freq clock 0 is output on SDP0. 11b = Freq clock 1 is output on SDP0.
TS_SDP0_EN	8	0b	RW	When set indicates that SDP0 is assigned to Tsync.
TS_SDP1_SEL	10:9	0x0	RW	SDP1 Allocation To Tsync Event. When TS_SDP1_EN is set, these bits select the Tsync event that is routed to SDP1. 00b = Target time 0 is output on SDP1. 01b = Target time 1 is output on SDP1. 10b = Freq clock 0 is output on SDP1. 11b = Freq clock 1 is output on SDP1.
TS_SDP1_EN	11	0b	RW	When set indicates that SDP1 is assigned to Tsync.
TS_SDP2_SEL	13:12	0x0	RW	SDP2 Allocation To Tsync Event. When TS_SDP2_EN is set, these bits select the Tsync event that is routed to SDP2. 00b = Target time 0 is output on SDP2. 01b = Target time 1 is output on SDP2. 10b = Freq clock 0 is output on SDP2. 11b = Freq clock 1 is output on SDP2.
TS_SDP2_EN	14	0b	RW	When set indicates that SDP2 is assigned to Tsync.
TS_SDP3_SEL	16:15	0x0	RW	SDP3 Allocation T Tsync Event. When TS_SDP3_EN is set, these bits select the Tsync event that is routed to SDP3. 00b = Target time 0 is output on SDP3. 01b = Target time 1 is output on SDP3. 10b = Freq clock 0 is output on SDP3. 11b = Freq clock 1 is output on SDP3.
TS_SDP3_EN	17	0b	RW	When set, indicates that SDP3 is assigned to Tsync.
RESERVED	31:18	0x0	RSV	Reserved. Write 0, ignore on read.



8.2.2.15.23 TimeSync Auxiliary Control Register - TSAUXC (0x00008C20)

Field	Bit(s)	Init.	Access Type	Description
EN_TT0	0	0b	RW	Enable Target Time 0. Enable bit is set by software to 1b to enable pulse or level change generation as a function of the <i>TSAUXC.PLSG</i> bit. The bit is auto-cleared by hardware when the target time hits if the <i>DIS_TS_CLEAR</i> flag in this register is cleared. During normal operation, the <i>DIS_TS_CLEAR</i> flag is set and this bit is not auto-cleared by the hardware.
EN_TT1	1	0b	RW	Enable Target Time 1. Enable bit is set by software to 1b to enable pulse or level change generation as a function of the <i>TSAUXC.PLSG</i> bit. The bit is auto-cleared by hardware when the target time hits if the <i>DIS_TS_CLEAR</i> flag in this register is cleared. During normal operation, the <i>DIS_TS_CLEAR</i> flag is set and this bit is not auto-cleared by hardware.
EN_CLK0	2	0b	RW	Enable Configurable Frequency Clock 0. Clock is generated according to frequency defined in the <i>FREQOUT0</i> register on the SDP pin (0 to 3) that has both: 1. <i>TSSDP.TS_SDPx_SEL</i> field with a value of 10b. 2. <i>TSSDP.TS_SDPx_EN</i> value of 1b.
SAMP_AUTO	3	0b	SC	When setting the <i>SAMP_AUTO</i> flag the <i>SYSTIMEL/H</i> registers are latched to the <i>AUXSTMPLO/ AUXSTMPH0</i> registers. Then the <i>SAMP_AUTO</i> flag is auto-cleared by hardware.
ST0	4	0b	RW	Start Clock 0 Toggle on Target Time 0. Enable clock 0 toggle only after target time 0, that's defined in the <i>TRGTTIML0</i> and <i>TRGTTIMH0</i> registers has passed.
EN_CLK1	5	0b	RW	Enable Configurable Frequency Clock 1. Clock is generated according to frequency defined in the <i>FREQOUT1</i> register on the SDP pin (0 to 3) that has both: 1. <i>TSSDP.TS_SDPx_SEL</i> field with a value of 11b. 2. <i>TSSDP.TS_SDPx_EN</i> value of 1b.
SAMP_AUT1	6	0b	SC	When setting the <i>SAMP_AUT1</i> flag the <i>SYSTIMEL/H</i> registers are latched to the <i>AUXSTMP1/ AUXSTMPH1</i> registers. Then the <i>SAMP_AUT1</i> flag is auto-cleared by hardware.
ST1	7	0b	RW	Start Clock 1 Toggle on Target Time 1. Enable clock 1 toggle only after target time 1, that's defined in the <i>TRGTTIML1</i> and <i>TRGTTIMH1</i> registers has passed.
EN_TS0	8	0b	RW	Enable Hardware Time Stamp 0. Enable Time stamping occurrence of change in SDP pin into the <i>AUXSTMPLO</i> and <i>AUXSTMPH0</i> registers. SDP pin (0 to 3) is selected for time stamping, if the SDP pin is selected via the <i>TSSDP.AUX0_SDP_SEL</i> field and the <i>TSSDP.AUX0_TS_SDP_EN</i> bit is set to 1b.
AUTT0	9	0b	RW	Auxiliary Time Stamp Taken. This is a read only bit field. At 0b the auxiliary time stamp 0 is enabled. This bit is set to 1b by hardware following a level change of the input SDP. Reading the <i>AUXSTMPH0</i> register clears this bit.
EN_TS1	10	0b	RW	Enable hardware time stamp 1.
AUTT1	11	0b	RW	Auxiliary Time Stamp Taken. This is a read only bit field. At 0b the auxiliary time stamp 1 is enabled. This bit is set to 1b by hardware following a level change of the input SDP. Reading the <i>AUXSTMPH1</i> register clears this bit.
RESERVED	16:12	0x0	RSV	Reserved.



Field	Bit(s)	Init.	Access Type	Description
PLSG0	17	0b	RW	Use target time 0 to generate start of pulse and target time 1 to generate end of pulse. SDP pin selected to drive pulse or level change is set according to the <i>TSSDP.TS_SDPx_SEL</i> field with a value of 00b and the <i>TSSDP.TS_SDPx_EN</i> bit with a value of 1b. 0b= Target time 0 generates change in SDP level. 1= Target time 0 generates start of pulse on SDP pin. Note: Pulse or level change is generated when TSAUXC.EN_TT0 is set to 1b.
RESERVED	29:18	0x0	RSV	Reserved.
DISABLE_SYSTIME	31	1b	RW	Disable SYSTIME Count Operation. 0b = SYSTIME timer activated. 1b = SYSTIME timer disabled. Value of SYSTIMEH, SYSTIMEL and SYSTIMER remains constant.

8.2.2.15.24 Target Time Register 0 Low - TRGTTIMELO (0x00008C24)

Field	Bit(s)	Init.	Access Type	Description
TTL	29:0	0x0	RW	Target time 0 LSB register (defined in ns units).
RESERVED	31:30	0x0	RW	Reserved.

8.2.2.15.25 Target Time Register 0 High - TRGTTIMEH0 (0x00008C28)

Field	Bit(s)	Init.	Access Type	Description
TTH	31:0	0x0	RW	Target time 0 MSB register (defined in second units).

8.2.2.15.26 Target time register 1 low - TRGTTIMEL1 (0x00008C2C)

Field	Bit(s)	Init.	Access Type	Description
TTL	29:0	0x0	RW	Target time 1 LSB register (defined in ns units).
RESERVED	31:30	0x0	RW	Reserved.

8.2.2.15.27 Target Time Register 1 High - TRGTTIMEH1 (0x00008C30)

Field	Bit(s)	Init.	Access Type	Description
TTH	31:0	0x0	RW	Target time 1 MSB register (defined in second units).



8.2.2.15.28 Frequency Out 0 Control Register - FREQOUT0 (0x00008C34)

Field	Bit(s)	Init.	Access Type	Description
CHCT	29:0	0x0	RW	Clock Out Half Cycle Time. Defines the half cycle time of clock 0 in ns units. When clock output is enabled, permitted values are any value larger than 25 and any value below 1 second.
RESERVED	31:30	0x0	RSV	Reserved

8.2.2.15.29 Frequency Out 1 Control Register - FREQOUT1 (0x00008C38)

Field	Bit(s)	Init.	Access Type	Description
CHCT	29:0	0x0	RW	Clock out half cycle time defines the half cycle time of clock 1 in ns units. When clock output is enabled, permitted values are any value larger than 13 and up to including 999,999,900 decimal (slightly below 1 second).
RESERVED	31:30	0x0	RSV	Reserved.

8.2.2.15.30 Auxiliary Time Stamp 0 Register Low - AUXSTMPLO (0x00008C3C)

Field	Bit(s)	Init.	Access Type	Description
TST_LOW	29:0	0x0	RO	Auxiliary time stamp 0 LSB value (defined in ns units).
RESERVED	31:30	0x0	RO	Reserved

8.2.2.15.31 Auxiliary Time Stamp 0 Register High - AUXSTMPHO (0x00008C40)

Field	Bit(s)	Init.	Access Type	Description
TST_HI	31:0	0x0	RO	Auxiliary time stamp 0 MSB value (defined in second units).

8.2.2.15.32 Auxiliary Time Stamp 1 Register Low - AUXSTMPL1 (0x00008C44)

Field	Bit(s)	Init.	Access Type	Description
TST_LOW	29:0	0x0	RO	Auxiliary time stamp 1 LSB value (defined in ns units).
RESERVED	31:30	0x0	RO	Reserved

8.2.2.15.33 Auxiliary Time Stamp 1 Register High - AUXSTMPH1 (0x00008C48)

Field	Bit(s)	Init.	Access Type	Description
TST_HI	31:0	0x0	RO	Auxiliary time stamp 1 MSB value (defined in second units).



8.2.2.16 Virtualization PF Registers

8.2.2.16.1 Filter Local Packets High - PFFLPH (0x000050B4)

Field	Bit(s)	Init.	Access Type	Description
FLP	31:0	0x0	RW	Filter Local Packets. Filter incoming packets whose MAC source address matches one of the LAN port DA MAC addresses. If the SA of the received packet matches one of the DA in the RAH/RAL registers, then the VM tied to this DA does not receive the packet. Other VMs can still receive it. (bit per pool)

8.2.2.16.2 Filter Local Packets Low - PFFLPL (0x000050B0)

Field	Bit(s)	Init.	Access Type	Description
FLP	31:0	0x0	RW	Filter Local Packets. Filter incoming packets whose MAC source address matches one of the LAN port DA MAC addresses. If the SA of the received packet matches one of the DA in the RAH/RAL registers, then the VM tied to this DA does not receive the packet. Other VMs can still receive it. (bit per pool)

8.2.2.16.3 PF Virtual Control Register - PFVCTL (0x000051B0)

Field	Bit(s)	Init.	Access Type	Description
VT_ENA	0	0b	RW	Virtualization Enabled Mode. When set, the Intel® Xeon® Processor D-1500 Product Family LAN controller supports either 16, 32, or 64 pools. When cleared, Rx traffic is handled internally as if it belongs to VF zero while VF zero is enabled. This bit should be set the same as MTQC.VT_Ena.
RESERVED	6:1	0x0	RSV	Reserved.
DEF_PL	12:7	0x0	RW	Default Pool. Pool assignment for packets that do not pass any pool queuing decision. Enabled by the <i>Dis_Def_Pool</i> bit.
RESERVED	15:13	0x0	RSV	Reserved.
POOLING_MODE	17:16	0x0	RW	Pooling Mode. 00b= Pool select by MAC address. 01b= Pool select by MAC address or E-tag. 10b= Reserved. 11b= Reserved.
RESERVED	28:18	0x0	RSV	Reserved.
DIS_DEF_POOL	29	0b	RW	Disable Default Pool. Determines the behavior of an Rx packet that does not match any Rx filter and is therefore not allocated a destination pool. 0b= Packet is assigned to the Default Pool (see DEF_PL previously described). 1b= Packet is dropped.
RPL_EN	30	0b	RW	Replication enable when set to 1b. When set, MRQC.MRQE should be set to one of the virtualization modes (1000b:1111b).
RESERVED	31	0b	RSV	Reserved.



8.2.2.16.4 PF Mailbox - PFMALBOX[n] (0x00004B00 + 0x4*n, n=0...63)

Field	Bit(s)	Init.	Access Type	Description
STS	0	0b	WO	Status/Command From PF Ready. Setting this bit, causes an interrupt to the relevant VF. This bit is always read as zero. Setting this bit sets the <i>PFSTS</i> bit in VFMailbox.
ACK	1	0b	WO	VF Message Received. Setting this bit, causes an interrupt to the relevant VF... This bit is always read as zero. Setting this bit sets the <i>PFAACK</i> bit in VFMailbox.
VFU	2	0b	RW	Buffer Is Taken By VF. This bit is RO for the PF and is a mirror of the <i>VFU</i> bit of the VFMailbox register.
PFU	3	0b	RW	Buffer Is Taken By PF. This bit can be set only if the <i>VFU</i> bit is cleared and is mirrored in the <i>PFU</i> bit of the VFMailbox register.
RVFU	4	0b	WO	Reset VFU. Setting this bit clears the <i>VFU</i> bit in the corresponding VFMailbox register. This bit should be used only if the VF driver is stuck. Setting this bit also resets the corresponding bits in the <i>PFMBICR</i> , <i>VFREQ</i> , and <i>VFAACK</i> fields.
RESERVED	31:5	0x0	RSV	Reserved.

8.2.2.16.5 PF Mailbox Interrupt Causes Register - PFMBICR[n] (0x00000710 + 0x4*n, n=0...3)

Each register handles 16 VFs as follows.


Field	Bit(s)	Init.	Access Type	Description
VFREQ	15:0	0x0	RW1C	Each bit in the <i>VFREQ</i> field is set when VF number (16*n+j) wrote a message in its mailbox. While n is the register index, n=0...3 and j is the index of the bits in the <i>VFREQ</i> , j=0...15.
VFAACK	31:16	0x0	RW1C	Each bit in the <i>VFAACK</i> field is set when VF number (16*n+j) acknowledged a PF message. While n is the register index, n=0...3 and 16+j is the index of the bits in the <i>VFAACK</i> , j=0...15.

8.2.2.16.6 PF Mailbox Interrupt Mask Register - PFMBIMR[n] (0x00000720 + 0x4*n, n=0...1)

Field	Bit(s)	Init.	Access Type	Description
VFIM	31:0	0xFF	RW	Bit j. Mailbox indication from VF # (32*n+j) might cause an interrupt to the PF...

8.2.2.16.7 PF VFLR Events Clear - PFVFLREC[n] (0x00000700 + 0x4*n, n=0...1)

Field	Bit(s)	Init.	Access Type	Description
CLEAR_VFLE	31:0	0x0	RW1C	When set, bit i in register n reflects an FLR event on VF# 32*n+i. These bits are accessible only to the PF and are cleared by writing 0x1.



8.2.2.16.8 PF VF Receive Enable - PFVFRE[n] (0x000051E0 + 0x4*n, n=0...1)

This register is reset on common reset cases and on per-function reset cases. Respective bits per VF are reset on VFLR, BME clear or on VF software reset. See VF Receive Enable - PFVFRE / VF Transmit Enable - PFVFTE for more details.

Field	Bit(s)	Init.	Access Type	Description
PFVFRE	31:0	0x0	RW	Bit j. Enables receiving packets to VF# (32*n+j). Each bit is cleared by the relevant VFLR.

8.2.2.16.9 PF VF Transmit Enable - PFVFTE[n] (0x00008110 + 0x4*n, n=0...1)

This register is reset on common reset cases and on per-function reset cases. Respective bits per VF are reset on VFLR, BME clear or on VF software reset. See VF Receive Enable - PFVFRE / VF Transmit Enable - PFVFTE for more details.

Field	Bit(s)	Init.	Access Type	Description
PFVFTE	31:0	0x0	RW	Bit j. Enables transmitting packets from VF# (32*n+j). Each bit is cleared by the relevant VFLR.

8.2.2.16.10 PF VM 0:31 Error Count Mask - PFVMECM0 (0x00008790)

Field	Bit(s)	Init.	Access Type	Description
FILTER	31:0	0x0	RW	Defines if a packet dropped from pools 0-to-31, respectively is counted in the SSVPC counter.

8.2.2.16.11 PF VM 32:63 Error Count Mask - PFVMECM1 (0x00008794)

Field	Bit(s)	Init.	Access Type	Description
FILTER	31:0	0x0	RW	Defines if a packet dropped from pools 32-to-63, respectively is counted in the SSVPC counter.

8.2.2.16.12 PF Queue Drop Enable Register - PFQDE (0x00002F04)

Field	Bit(s)	Init.	Access Type	Description
PFQDE	0	0b	RW	Enable drop of packets from Rx queue Queue_Index. This bit overrides the <i>SRRCTL.drop_en</i> bit of each queue. For example, if either of the bits are set, a packet received when no descriptor is available is dropped.
HIDE_VLAN	1	0b	RW	If this bit is set, the <i>RXDCTL.VME</i> setting is ignored, the VLAN is always stripped, a value of zero is written in the <i>RDESC.VLAN</i> tag and in the <i>RDESC.STATUS.VP</i> fields of the received descriptor.



Field	Bit(s)	Init.	Access Type	Description
STRIP_TAG	2	0b	RW	If this bit is set, the E-tag is stripped from the packet. If only one type of tag should be stripped, the other tag Ethertype should be invalidated using VNTAGTBL.VALID or ETAG_ETYPE.VALID field.
RESERVED	3	0b	RSV	Reserved. Refer to the <i>WE</i> and <i>RE</i> bits described later in this section.
RESERVED	7:4	0x0	RSV	Reserved.
QUEUE_INDEX	14:8	0x0	RW	Indicates the queue referenced upon <i>WE/RE</i> commands.
RESERVED	15	0b	RSV	Reserved.
WE	16	0b	RW	Write Enable. When this bit is set, the content of bits 2:0 are written into the relevant queue context. Bit 3 is reserved. This bit should never be set together with the <i>RE</i> bit in this register.
RE	17	0b	RW	Read Enable. When this bit is set, the content of bits 2:0 are read from the relevant queue context. Bit 3 is reserved. This bit should never be set together with the <i>WE</i> bit in this register.
RESERVED	31:18	0x0	RSV	Reserved.

8.2.2.16.13 PF VM Tx Switch Loopback Enable - PFVMTXSW[n] (0x00005180 + 0x4*n, n=0..1)

Field	Bit(s)	Init.	Access Type	Description
LLE	31:0	0x0	RW	Local Loopback Enable. For each register <i>n</i> , and bit <i>i</i> , <i>i</i> =0..31, enables local loopback for pool $32*n+1$. When set, a packet originating from a specific pool and destined to the same pool is allowed to be looped back. If cleared, the packet is dropped.

8.2.2.16.14 PFVF Anti Spoof Control - PFVFSPOOF[n] (0x00008200 + 0x4*n, n=0..7)

Field	Bit(s)	Init.	Access Type	Description
MACAS	7:0	0x00	RW	For each register <i>n</i> , and bit <i>i</i> , <i>i</i> =0..7, enables anti-spoofing filter on Ethernet MAC addresses for $VF(8*n+i)$.
VLANAS	15:8	0x00	RW	For each register <i>n</i> , and bit $8+i$, <i>i</i> =0..7, enables anti-spoofing filter on VLAN tag for $VF(8*n+i)$. Note: If <i>VLANAS</i> is set for a specific pool, then the respective <i>MACAS</i> bit must be set as well.
ETHERTYPEAS	23:16	0x0	RW	For each register <i>n</i> , and bit $16+i$, <i>i</i> =0..7, enables anti-spoofing filter on Ethertype for $VF(8*n+i)$.
ETHERTYPELB	31:24	0x0	RW	For each register <i>n</i> , and bit $24+i$, <i>i</i> =0..7, enables loopback filter on Ethertype for $VF(8*n+i)$.



8.2.2.16.15 PF DMA Tx General Switch Control - PFDTXGSWC (0x00008220)

Field	Bit(s)	Init.	Access Type	Description
LBE	0	0b	RW	Enables VMDQ Loopback.
RESERVED	31:1	0x0	RSV	Reserved.

8.2.2.16.16 PF VM VLAN Insert Register - PFVMVIR[n] (0x00008000 + 0x4*n, n=0...63)

Field	Bit(s)	Init.	Access Type	Description
PORT_VLAN_ID	15:0	0x0	RW	Port VLAN tag to insert if the VLANA field = 01b.
RESERVED	26:16	0x0	RSV	Reserved.
TAGA	28:27	0x0	RW	Tag Action. 00b = Do not insert any outer tag. 01b = Insert an E-tag. 10b = Reserved. 11b = Reserved.
RESERVED	29	0b	RSV	Reserved.
VLANA	31:30	0x0	RW	VLAN Action. 00b = use descriptor command. 01b = always insert default VLAN. 10b = Never insert VLAN. 11b = Reserved.

8.2.2.16.17 PF VM Tag Insert Register - PFVMTIR[n] (0x00017000 + 0x4*n, n=0...63)

Field	Bit(s)	Init.	Access Type	Description
PORT_TAG_ID	31:0	0x0	RW	Port tag to insert if the <i>PFVMIR.TAGA</i> field = 01b or 10b. TAGA = 01b - insert E-tag (such as Etag Etype-PFVMTIR[31:0]-0x0000). TAGA = Reserved. TAGA = Other - ignore this field.

8.2.2.16.18 PF VM L2Control Register - PFVML2FLT[n] (0x0000F000 + 0x4*n, n=0...63)

This register controls per VM Inexact L2 filtering.

Field	Bit(s)	Init.	Access Type	Description
RESERVED	21:0	0x00	RSV	Reserved.
UPE	22	0b	RW	Unicast Promiscuous.
VPE	23	0b	RW	VLAN Promiscuous Enable.
AUPE	24	0b	RW	Accept Untagged Packets Enable. When set, packets without a VLAN tag can be forwarded to this queue, assuming they pass the Ethernet MAC address queuing mechanism.



Field	Bit(s)	Init.	Access Type	Description
ROMPE	25	0b	RW	Receive Overflow Multicast Packets. Accept packets that match the MTA table.
ROPE	26	0b	RW	Receive MAC Filters Overflow. Accept packets that match the PFUTA table.
BAM	27	0b	RW	Broadcast Accept.
MPE	28	0b	RW	Multicast Promiscuous.
RESERVED	31:29	0x00	RSV	Reserved.

8.2.2.16.19 PF VM VLAN Pool Filter - PFVLVF[n] (0x0000F100 + 0x4*n, n=0...63)

Software should initialize these registers before transmit and receive are enabled.

Field	Bit(s)	Init.	Access Type	Description
VLAN_ID	11:0	X	RW	Defines a VLAN tag for pool VLAN filter n. The bitmap defines which pools belong to this VLAN. Note: Appears in little endian order (LSB last on the wire).
RESERVED	30:12	X	RSV	Reserved.
VI_EN	31	0b	RW	VLAN ID Enable. This filter is valid.

8.2.2.16.20 PF VM VLAN Pool Filter Bitmap - PFVLVFB[n] (0x0000F200 + 0x4*n, n=0...127)

Software should initialize these registers before transmit and receive are enabled.

Field	Bit(s)	Init.	Access Type	Description
POOL_ENA	31:0	X	RW	Pool Enable Bit Array. Each couple of registers 2*n and 2*n+1 enables routing of packets that match a PFVLVF[n] filter to a pool list. Each bit when set, enables packet reception with the associated tools as described in the sections that follow: Bit i in register 2*n is associated with pool i+ bit i in register 2*n+1 is associated with pool 32+i.

8.2.2.16.21 PF Unicast Table Array - PFUTA[n] (0x0000F400 + 0x4*n, n=0...127)

There is one register per 32 bits of the unicast address table for a total of 128 registers (the PFUTA[127:0] designation). Software must mask to the desired bit on reads and supply a 32-bit word on writes. The first bit of the address used to access the table is set according to the *MCSTCTRL.MO* field. The 7 MS bits of the Ethernet MAC address (out of the 12 bits) selects the register index while the 5 LS bits (out of the 12 bits) selects the bit within a register.

Note: All accesses to this table must be 32 bit. The lookup algorithm is the same one used for the MTA table. This table should be zeroed by software before start of work.

Field	Bit(s)	Init.	Access Type	Description
BIT_VECTOR	31:0	X	RW	Word wide bit vector specifying 32 bits in the unicast destination address filter table.



8.2.2.16.22 PF Mirror Rule Control - PFMRCTL[n] (0x0000F600 + 0x4*n, n=0...3)

This register defines mirroring rules for each of 4 destination pools.

Field	Bit(s)	Init.	Access Type	Description
VPME	0	0b	RW	Virtual Pool Mirroring Enable. Enables mirroring of certain pools as defined in the PFMRVM registers.
UPME	1	0b	RW	Uplink Port Mirroring Enable. Enables mirroring of all traffic received from the network.
DPME	2	0b	RW	Downlink Port Mirroring Enable. Enables mirroring of all traffic transmitted to the network.
VLME	3	0b	RW	VLAN Mirroring Enable. Enables mirroring of a set of given VLANs as defined in the PFMRVLAN registers.
RESERVED	7:4	0x0	RSV	Reserved.
MP	13:8	0x0	RW	Mirror Pool. Defines the destination pool for this mirror rule.
RESERVED	31:14	0x0	RSV	Reserved.

8.2.2.16.23 PF Mirror Rule VLAN - PFMRVLAN[n] (0x0000F610 + 0x4*n, n=0...7)

This register defines the VLAN values as listed in the PFVLVF table taking part in the VLAN mirror rule. Registers 0, 4 correspond to rule 0, registers 1, 5 correspond to rule 1, etc. Registers 0-3 correspond to the LSB in the PFVLVF table. For example, register 0 corresponds to VLAN filters 31:0, while register 4 corresponds to VLAN filters 63:32.

Field	Bit(s)	Init.	Access Type	Description
VLAN	31:0	0x0	RW	Bitmap listing which VLANs participate in the mirror rule.

8.2.2.16.24 PF Mirror Rule Pool - PFMRVM[n] (0x0000F630 + 0x4*n, n=0...7)

This register defines which pools are being mirrored to the destination pool. Registers 0, 4 correspond to rule 0, registers 1, 5 correspond to rule 1, etc. Registers 0-3 correspond to the LSB in the pool list. For example, register 0 corresponds to pools 31:0, while register 4 corresponds to pools 63:32.

Field	Bit(s)	Init.	Access Type	Description
POOL	31:0	0x0	RW	Bitmap listing which pools participate in the mirror rule.

8.2.2.16.25 Last VM Misbehavior Cause - RX - LVMMC_RX (0x00002FA8)

Bits in LVMMC_RX register define the cause for blocking the malicious queue that was reported in the LMVM_RX. MALICIOUS_QUEUE field when RDRXCTL.MDP_EN is set. Refer to Interrupt on Misbehavior of VM (Malicious Driver Detection) section for details of the different bits.



Field	Bit(s)	Init.	Access Type	Description
INV_MACC	0	0b	RC	Invalid Memory Access A PCIe DMA access initiated by a VF ended with Unsupported Request (UR) or Completer Abort (CA) or a read access was blocked due to out of range address. When a malicious DMA access is detected, the Rx queue (VF) that initiated the access is disabled and corresponding WQBR_RX bit is set.
INVALID_RX_Q_CONTEXT	1	0b	RC	An invalid receive queue context was detected when queue was enabled or an attempt to change the static part of the queue context on the fly was detected.
RESERVED	31:2	0x0	RSV	Reserved.

8.2.2.16.26 Last VM Misbehavior Cause - TX - LVMMC_TX (0x00008108)

Bits in LVMMC_TX register define the cause for blocking the malicious queue that was reported in the LMVM_TX. MALICIOUS_QUEUE field when DMATXCTL.MDP_EN is set. Refer to Interrupt on Misbehavior of VM (Malicious Driver Detection) section for details of the different bits.

Field	Bit(s)	Init.	Access Type	Description
MAC_HEADER	0	0b	RC	Illegal MAC header size.
IPV4_HEADER	1	0b	RC	Illegal IPv4 header size.
IPV6_HEADER	2	0b	RC	Illegal IPv6 header size.
WRONG_MAC_IP	3	0b	RC	Wrong MAC +IP header size.
TCP_LSO	4	0b	RC	Illegal TCP header was detected in a large send operation.
IPSEC_OFFLOAD	5	0b	RC	A VF requested an IPSec offload.
UDP_LSO	6	0b	RC	Illegal UDP header was detected in a large send operation.
STCP_CS	7	0b	RC	Illegal STCP header was detected in an SCTP checksum offload operation,
SIZE	8	0b	RC	Illegal packet size (> 15.5K).
RESERVED	9	0b	RSV	Reserved.
OFF_ILL	10	0b	RC	Illegal offload request.
SCTP_ALIGN	11	0b	RC	SCTP CRC request of non 4-byte aligned data.
ZERO_MSS	12	0b	RC	A request for large send with zero MSS was detected.
CONTEXT_IN_PACKET	13	0b	RC	A context descriptor was detected in the middle of a packet.
LSO_MORE_THAN_4	14	0b	RC	Large send with more than four header buffers misbehavior was detected.
OOS_SSO	15	0b	RC	Out of sync during single send.
OOS_LSO	16	0b	RC	Out of sync during large send.
SSO_UDP	17	0b	RC	Wrong parameter of headers for UDP SSO.
SSO_TCP	18	0b	RC	Wrong parameter of headers for TCP SSO.



Field	Bit(s)	Init.	Access Type	Description
INV_MACC	19	0b	RC	Invalid Access. A PCIe DMA access initiated by a VF ended with Unsupported Request (UR) or Completer Abort (CA) or was prevented from being sent if out of range. When a malicious DMA access is detected, the Tx queue (VF) that initiated the access is disabled and the corresponding WQBR_TX bit is set.
DESC_TYPE	20	0b	RC	Wrong descriptor type (other than 2, 3).
WRONG_NULL	21	0b	RC	Null without EOP.
NO_EOP	22	0b	RC	Packet without EOP (such as bigger than the ring size).
CONTEXT_BURST	23	0b	RC	Contiguous context descriptor burst that exceeds three contexts.
RESERVED	24	0b	RSV	Reserved.
MAC_VLAN_SPOOF	25	0b	RC	A MAC spoof or VLAN spoof attempt was detected.
VLAN_IERR	26	0b	RC	VLAN Insertion Error. The VLE bit set in Tx descriptor when the VMVIR[n].VLANA register field is not 0b causing a drop of Tx packets.
LEGACY_DESC_IOV	27	0b	RC	A legacy descriptor in IOV was detected.
RESERVED	28	0b	RC	Reserved.
INVALID_TXQ_CONTEXT	29	0b	RC	An invalid transmit queue context was detected when the queue was enabled or an attempt to change the static part of the queue context on the fly was detected.
RESERVED	31:30	0x0	RSV	Reserved.

Note: Only the first malicious event is registered for each packet, so if a bit is not set it doesn't mean that this event didn't occur, only that another malicious behavior was detected first.

8.2.2.16.27 Last Malicious VM - RX - LMVM_RX (0x00002FA4)

Field	Bit(s)	Init.	Access Type	Description
MALICIOUS_QUEUE	6:0	0x0	RW	The Rx queue on which the malicious behavior reported in LVMMC was detected. The queue number is the absolute number in the PF space.
RESERVED	16:7	0x0	RSV	Reserved.
MAL_PF	17	0b	RW	Malicious Driver Behavior Detected On Current PF. 0b = Malicious event was on a queue belonging to a VF. 1b = Malicious event was on a queue belonging to the PF.
RESERVED	31:18	0x0	RSV	Reserved.

8.2.2.16.28 Last Malicious VM - TX - LMVM_TX (0x00008124)

Field	Bit(s)	Init.	Access Type	Description
MALICIOUS_QUEUE	6:0	0x0	RW	The Tx queue on which the malicious behavior reported in LVMMC was detected. The queue number is the absolute number in the PF space.
RESERVED	16:7	0x0	RSV	Reserved.



Field	Bit(s)	Init.	Access Type	Description
MAL_PF	17	0b	RW	Malicious Driver Behavior Detected On Current PF. 0b = Malicious event was on a queue belonging to a VF. 1b = Malicious event was on a queue belonging to the PF.
RESERVED	31:18	0x0	RSV	Reserved.

8.2.2.16.29 Wrong Queue Behavior Register - Rx - WQBR_RX[n] (0x00002FB0 + 0x4*n, n=0...3)

Field	Bit(s)	Init.	Access Type	Description
WVBR	31:0	0x0	RW1C	Bitmap indicating against which Rx queue a malicious action was taken. The queue is released only by a reset of the VF or by clearing the corresponding bit in WQBR_RX register.

8.2.2.16.30 Wrong Queue Behavior Register - Tx - WQBR_TX[n] (0x00008130 + 0x4*n, n=0...3)

Field	Bit(s)	Init.	Access Type	Description
WVBR	31:0	0x0	RW1C	Bitmap indicating against which Tx queue an anti-spoof action or malicious action was taken. The queue is released only by a reset of the VF or by clearing the corresponding bit in the WQBR_TX register.

8.2.2.17 Power Management

8.2.2.17.1 DMA Coalescing Management Threshold - DMCMNGTH (0x00015F20)

Field	Bit(s)	Init.	Access Type	Description
RESERVED	3:0	0x0	RSV	Reserved.
DMCMNGTHR	19:4	0x100	RW	DMA Coalescing Management Threshold. BMC TX DMAC Threshold. This value defines the DMA coalescing management threshold in 16-byte units. When the amount of empty space in the internal transmit buffer exceeds the DMCMNGTHR value, DMA coalescing is stopped and the integrated I/O interface ISM moves out of idle state. Note: If the value is 0x0, the condition to move out of DMA coalescing due to passing the DMA coalescing management threshold level is disabled.
RESERVED	31:20	0x0	RSV	Reserved.

8.2.2.17.2 DMA Coalescing Control Register - DMACR (0x00002400)



Field	Bit(s)	Init.	Access Type	Description
DMACWT	15:0	0x20	RW	DMA Coalescing Watchdog Timer. When in DMA coalescing, defines the upper limit in 40.960 μs units between arrival of a coalescing exit event condition and actual exit from DMA coalescing. The programmed value should be a non-zero.
HIGH_PRI_TC	23:16	0x0	RW	High Priority Traffic Class. This field defines which TC is considered high priority and a packet received for this TC will causes an exit from DMA coalescing. When a TC bit is set, it indicates this TC is high priority. Bit 16 TC 0. Bit 17 TC 1. ... Bit 23 TC 7.
RESERVED	24	0b	RSV	Reserved.
RESERVED	28	1b	RW	Reserved.
RESERVED	29	0b	RSV	Reserved.
LX_COALESCING_INDICATION	30	0b	RW	Lx Coalescing Indication. Defines whether to move in/out of DMA coalescing when the integrated I/O interface ISM moves in/out of idle state. 1b = DMA coalescing conditions are met only when the integrated I/O interface ISM is in the idle state. 0b = DMA coalescing can also start when the integrated I/O interface is not in idle. DMA coalescing stops when any TLP transactions are executed on the integrated I/O interface.
DMAC_EN	31	0b	RW	DMA Coalescing Enable. 0b= Disable DMA coalescing. 1b= Enable DMA coalescing.

8.2.2.17.3 DMA Coalescing Threshold - DMCTH[n] (0x00003300 + 0x4*n, n=0...7)

Field	Bit(s)	Init.	Access Type	Description
DMACRXT	8:0	0x20	RW	DMA Coalescing Receive Threshold Rx PB TC[n]. This value defines the DMA coalescing receive threshold in 1KB units. When the amount of data in the internal receive buffer exceeds the DMACRXT[n] value, DMA coalescing is stopped. Notes: Values written to this field should take into account: 1. The Latency Tolerance Requirements (LTR) sent over the PCIe and XOFF receive threshold values to avoid needless generation of flow control packets when in DMA coalescing operating mode and flow control is enabled. 2. The maximum size of the receive buffer.
RESERVED	31:9	0x0	RSV	Reserved.



8.2.2.17.4 DMA Coalescing Time to Lx Request - DMCTLX (0x0002404)

Field	Bit(s)	Init.	Access Type	Description
TTLX	11:0	0x20	RW	Time To Lx Request. Controls the time between detecting a DMA coalescing condition to an actual entry into a DMA coalescing state. Timer counts is in 1.28 μ s intervals. Note: The timer adds a delay to a decision as to when to enter a DMA coalescing state. Note: For a link speed of 10 Gb/s, a minimum value should be 0x1. For a link speed of 1000 Mb/s a minimum value should be 0x2. For link speed of 100 Mb/s, a minimum value should be 0x20. This limitation is enforced by hardware.
RESERVED	31:12	0x0	RSV	Reserved. Write 0x0, ignore on read.

8.2.2.17.5 Latency Tolerance Reporting (LTR) Control - LTRC (0x00011708)

Field	Bit(s)	Init.	Access Type	Description
SLTRV	9:0	0x5	RW	Snoop Latency Value. Along with the <i>Max Snoop Latency Scale</i> field, this register specifies the maximum snoop latency this function needs to request. Software should set this field only if it is set to a lower number than the platform's maximum supported latency.
SSCALE	12:10	0x2	RW	Snoop Latency Scale. This field provides a scale for the value contained within the <i>Snoop Latency Value</i> field. Encoding: 000b = Value times 1 ns. 001b = Value times 32 ns. 010b = Value times 1,024 ns. 011b = Value times 32,768 ns. 100b = Value times 1,048,576 ns. 101b = Value times 33,554,432 ns. 110b - 111b = Not permitted. The software device driver is signalled to use the same scale as configured in LTR capability.
RESERVED	14:13	0x0	RSV	Reserved. Write 0x0, ignore on read.
LTRS_REQUIREMENT	15	0b	RW	LTR Snoop Requirement. 0b= No Latency requirements in snoop memory access. 1b= Latency tolerance in snoop memory access specified in the <i>LTRC.SLTRV</i> field.
NSLTRV	25:16	0x5	RW	No Snoop Latency Value. Along with the <i>Max No Snoop Latency Scale</i> field, this register specifies the maximum no snoop latency this function needs to request. Software should set this field only if it is set to a lower number than the platform's maximum supported latency.



Field	Bit(s)	Init.	Access Type	Description
NSSCALE	28:26	0x2	RW	No Snoop Latency Scale. This field provides a scale for the value contained within the <i>No Snoop Latency Value</i> field. Encoding: 000b = Value times 1 ns. 001b = Value times 32 ns. 010b = Value times 1,024 ns. 011b = Value times 32,768 ns. 100b = Value times 1,048,576 ns. 101b = Value times 33,554,432 ns. 110b - 111b = Not permitted. The software device driver is signalled to use the same scale as configured in LTR capability.
RESERVED	29	0b	RSV	Reserved. Write 0b, ignore on read.
LTR_SEND	30	0b	RW	An indication to send an LTR message. This bit is set by firmware (indicates to hardware to send an LTR message update to the system) and cleared by hardware after the message was sent. This bit is not sent as part of the LTR message.
LTRNS_REQUIREMENT	31	0b	RW	LTR No Snoop Requirement. 0b = No latency requirements in no snoop memory access. 1b = Latency tolerance in no snoop memory access specified in the <i>LTRC.NSLTRV</i> field.

8.2.2.17.6 Energy Efficient Ethernet (EEE) Register - EEER (0x000043A0)

Field	Bit(s)	Init.	Access Type	Description
TW_SYSTEM	15:0	0x0	RW	Time expressed in μ s that no data is transmitted following a move from a EEE Tx LPI link state to a link active state. This field holds the transmit <i>Tw_sys_tx</i> value negotiated during EEE LLDP negotiation. Notes: 1. If the value is lower than the minimum <i>Tw_sys_tx</i> value defined in IEEE802.3az clause 78.5 then the interval where no data is transmitted following a move out of a EEE Tx LPI state defaults to a minimum <i>Tw_sys_tx</i> . 2. Following link disconnect or auto-negotiation value of this field returns to default value, until software re-negotiates a new <i>tw_sys_tx</i> value via EEE LLDP. 3. Fast retrain and local/remote fault indication is not considered a link disconnect and do not cause the field to return to the default value. 4. When transmitting flow control frames, the Intel® Xeon® Processor D-1500 Product Family LAN controller waits the minimum time defined in the IEEE802.3az standard before transmitting the flow control packet. The Intel® Xeon® Processor D-1500 Product Family LAN controller does not wait the <i>Tw_system</i> time following an exit of LPI before transmitting the flow control frame.



Field	Bit(s)	Init.	Access Type	Description
TX_LPI_EN	16	0b	RW	Enable Entry Into EEE LPI on TX Path 0b = Disable entry into EEE LPI on the Tx path. 1b = Enable entry into EEE LPI on the Tx path. Notes: 1. Even when <i>TX_LPI_EN</i> is 1b, the Intel® Xeon® Processor D-1500 Product Family LAN controller does not enable entry into a Tx LPI state for at least one second following the change of link_status to OK as defined in IEEE802.3az clause 78.1.2.1. 2. Even if the <i>TX_LPI_EN</i> bit is set, the Intel® Xeon® Processor D-1500 Product Family LAN controller initiates an entry into a Tx EEE LPI link state only if EEE support at the link speed was negotiated during auto-negotiation.
RX_LPI_EN	17	1b	RW	Enable Entry Into EEE LPI on RX Path 0b = Disable entry into EEE LPI on the Rx path. 1b = Enable entry into EEE LPI on the Rx path. Note: Even if the <i>RX_LPI_EN</i> bit is set, the Intel® Xeon® Processor D-1500 Product Family LAN controller recognizes an entry into a Rx EEE LPI link state only if EEE support at the link speed was negotiated during auto-negotiation.
RESERVED	25:19	0x0	RSV	Reserved. Write 0x0 ignore on read.
EEE_FRC_AN	28	0b	RW	Force EEE Auto-negotiation. When this bit is set to 1b, it enables EEE operation in the internal MAC logic even if a link partner does not support EEE. Note: Should be set to 1b to enable testing of EEE operation via MAC loopback.
RESERVED	31:29	0x0	RSV	Reserved. Write 0x0 ignore on read.

8.2.2.17.7 Energy Efficient Ethernet (EEE) Setup Register - EEE_SU (0x00004380)

Field	Bit(s)	Init.	Access Type	Description
DTW_MIN	7:0	0x0	RW	Time to add to the minimum <i>Tw_sys_tx</i> value defined in IEEE802.3az clause 78.5 This additional time should be configured according to the link speed/type and is in addition to the specification definition. The basic value of <i>Tw_sys_tx</i> is: <ul style="list-style-type: none"> 7.36 μs when link speed is 10 GbE. An additional 10 μs should be added for worst case normal operation by setting this field to 0x64. 16.5 μs when link speed is 1 GbE. This value can be kept (keep this field at zero) for normal operation. The minimum <i>Tw_sys_tx</i> value defines the duration where no data is transmitted following a move out of the EEE LPI Tx state. Time to add defined in this field is expressed in 0.1 μ s resolution. Note: The idle time value defined by this field plus the <i>Tw_sys_tx</i> value defined in IEEE802.3az clause 78.5 is used when moving out of the EEE TX LPI state to transmit flow control frames even if the value specified in <i>EEER.Tw_system</i> field is higher.
RESERVED	15:8	0x0	RSV	Reserved.



Field	Bit(s)	Init.	Access Type	Description
TW_WAKE_MIN	21:16	0x0	RW	Minimum time expressed in one μ s, between sending a request to move into a EEE Tx LPI and sending a request to move back to an active state. Should be set to 0x7 for normal operation. Note: If conditions to exit LPI during the Tw_wake_min interval cease to exist then the Intel® Xeon® Processor D-1500 Product Family LAN controller does not move out of Tx LPI after the timer has expired.
RESERVED	23:22	0x0	RW	Reserved.
TX_LU_LPI_DELAY	25:24	0x3	RW	Delay to enable entry of a Tx EEE LPI state following a link-up indication. 00b = No delay. 01b = 10 ms. 10b = 100 ms. 11b = 1 second. Note: IEEE802.3az clause 78.1.2.1 defines delay of one second following link-up.
TEEE_DLY	31:26	0x2	RW	TX EEE LPI Entry Delay. This field defines delay to EEE entry once conditions to enter EEE LPI are detected. Field resolution is one μ s. Note: If conditions to enter LPI during the TEEE_DLY interval cease to exist the Intel® Xeon® Processor D-1500 Product Family LAN controller does not enter a Tx LPI and continues normal operation.

8.2.2.17.8 Energy Efficient Ethernet (EEE) STATUS - EEE_STAT (0x00004398)

Field	Bit(s)	Init.	Access Type	Description
RESERVED	28:0	0x0	RSV	Reserved. Write 0x0, ignore on read.
EEE_NEG	29	0b	RO	EEE Support Negotiated On Link. 0b = EEE operation not supported on link. 1b = EEE operation supported on link.
RX_LPI_STATUS	30	0b	RO	Rx Link in LPI state. 0b = Rx in Active state. 1b = Rx in LPI state.
TX_LPI_STATUS	31	0b	RO	Tx Link in LPI state. 0b = Tx in Active state. 1b = Tx in LPI state.

8.2.2.17.9 EEE TX LPI Count - TLPIC (0x000041F4)

This register counts EEE Tx LPI entry events. A EEE Tx LPI event occurs when the transmitter enters a EEE (IEEE802.3az) LPI state. This register only increments if transmits are enabled and EEE operation is enabled.

Field	Bit(s)	Init.	Access Type	Description
ETLPIC	31:0	0x0	RC	Number of EEE TX LPI Events. Register cleared on read, register doesn't wrap around at 0xFFFFFFFF.



8.2.2.17.10 EEE RX LPI Count - RLPIC (0x000041F8)

This register counts EEE Rx LPI entry events. A EEE Rx LPI event occurs when the receiver detects a link partner entry into a EEE (IEEE802.3az) LPI state. This register only increments if receives are enabled and EEE operation is enabled.

Field	Bit(s)	Init.	Access Type	Description
ERLPIC	31:0	0x0	RC	Number of EEE RX LPI events. Register cleared on read, register doesn't wrap around at 0xFFFFFFFF.

8.2.2.18 Diagnostic

The Intel® Xeon® Processor D-1500 Product Family LAN controller contains several diagnostic registers. These registers enable software to directly access the contents of the internal Packet Buffer Memory (PBM), also referred to as FIFO space. These registers also give software visibility into other locations in the design that are not for external use.

8.2.2.18.1 Security Tx Control - SECTXCTRL (0x00008800)

Field	Bit(s)	Init.	Access Type	Description
SECTX_DIS	0	1b	RW	Tx Security Offload Disable Bit. When set, the AES crypto engine used in Tx by MACsec and IPsec offloads is disabled. This mode must be used to save the Intel® Xeon® Processor D-1500 Product Family LAN controller's power consumption when no security offload is enabled. When cleared, the AES crypto engine used in Tx by MACsec or IPsec offload is enabled. Normal operating mode when a security offload is enabled. This bit is RW /RO if fused off.
TX_DIS	1	0b	RW	Disable Sec Tx Path. When set, no new packet is fetched out from the Tx packet buffers, so that the Tx security block can be internally emptied prior to changing the security mode. The <i>SECTXSTAT.SECTX_RDY</i> bit is deasserted until the path is emptied by hardware. When cleared, Tx data path is enabled. Normal operating mode.
STORE_FOR_WARD	2	0b	RW	Tx Sec Buffer Mode. When set, a complete frame is stored in the internal security Tx buffer prior to being forwarded to the MAC. Normal operating mode when IPsec offload is enabled (as requested to overwrite <i>ICV</i> field in AH frames). Note that it increases the Tx internal latencies (for all TCs). When cleared, the Tx sec buffer is operated in pass-through mode. Normal operating mode when MACsec is enabled or when no security offload is enabled.
RESERVED	31:4	0x0	RSV	Reserved.

8.2.2.18.2 Security Tx Status - SECTXSTAT (0x00008804)



Field	Bit(s)	Init.	Access Type	Description
SECTX_RDY	0	0b	RO	Tx Security Block Ready For Mode Change. When set, it indicates that the internal data path from the Tx packet buffers to the Tx security block has been emptied, and thus the security mode can be changed by software. When cleared, it indicates that the internal data path from the Tx packet buffers to the Tx security block is not empty, and thus software cannot change the security mode. This bit is polled by software once <i>SECTXCTRL.TX_DIS</i> bit was set.
SECTX_OFF_DIS	1	0b	RO	Tx Security Offload Disabled. When set, it indicates that the Tx security offload feature is disabled by fuse or strapping pin.
ECC_TXERR	2	0b	RO	Unrecoverable ECC Error in the Tx SA Table or SEC Tx FIFO Occurred. When set, it indicates that an unrecoverable ECC error occurred when accessing internally the Tx SA table. The ECC interrupt is set as well, until the Intel® Xeon® Processor D-1500 Product Family LAN controller is reset by software. When cleared, no ECC error occurred on the Tx SA table from the last time the Intel® Xeon® Processor D-1500 Product Family LAN controller has reset.
RESERVED	31:3	0x0	RSV	Reserved.

8.2.2.18.3 Security Tx Buffer Almost Full - SECTXBUFFAF (0x0008808)

Field	Bit(s)	Init.	Access Type	Description
FULLTHRESH	9:0	0x250	RW	Tx Security Buffer Almost Full Threshold (relatively to full capacity). The size of the security buffer is 0x274 lines of 16 bytes. In MACsec offload, the buffer operates in pass-through mode and the recommended threshold is 0x250. It means that the almost full indication is generated very soon while only a fraction of a packet is stored in the buffer. In IPsec mode, the buffer operates in a store and forward mode and the recommended threshold is 0x15. It means that the almost full indication is generated only after the buffer contains at least an entire jumbo packet.
RESERVED	31:10	0x0	RSV	Reserved.

8.2.2.18.4 Security Tx Buffer Minimum IFG - SECTXMINIFG (0x0008810)

Field	Bit(s)	Init.	Access Type	Description
MINSECIFG	3:0	0x1	RW	Minimum IFG between packets. It is the minimum gap between consecutive frames from the DBU-Tx required for the security block. The MINSECIFG is measured in wake DMA clock units (equal to 6.4 ns in 10 GbE).
RESERVED	7:4	0x0	RSV	Reserved.
MRKRINSERT	18:8	0x10	RW	This field is used to configure the security Tx buffer.
RESERVED	31:19	0x0	RSV	Reserved.

8.2.2.18.5 Security Rx Control - SECRXCTRL (0x0008D00)



Field	Bit(s)	Init.	Access Type	Description
SECRX_DIS	0	1b	RW	Rx Security Offload Disable Bit. When set, the AES crypto engine used in Rx by MACsec and IPsec offloads is disabled. This mode must be used to save the Intel® Xeon® Processor D-1500 Product Family LAN controller's power consumption when no security offload is enabled. When cleared, the AES crypto engine used in Rx by MACsec or IPsec offload is enabled. Normal operating mode when a security offload is enabled.
RX_DIS	1	0b	RW	Disable Sec Rx Path. When set, any new packet received from the Rx MAC is filtered out, so that the Rx security block can be internally emptied prior to changing the security mode. SECRXSTAT.SECRX_RDY bit is deasserted until the path is emptied by hardware. When cleared, Rx data path is enabled. Normal operating mode. This bit is RW / RO if fused off.
RESERVED	31:4	0x0	RSV	Reserved.

8.2.2.18.6 Security Rx Status - SECRXSTAT (0x00008D04)

Field	Bit(s)	Init.	Access Type	Description
SECRX_RDY	0	0b	RO	Rx security block ready for mode change. When set, it indicates that the internal data path from the Rx MAC to the Rx security block has been emptied, and thus the security mode can be changed by software. When cleared, it indicates that the internal data path from the Rx MAC to the Rx security block is not empty, and thus software cannot change the security mode. This bit is polled by software once SECRXCTRL.RX_DIS bit was set.
SECRX_OFF_DIS	1	0b	RO	Rx Security Offload Disabled. When set, it indicates that the Rx security offload feature is disabled by the internal fuse or by strapping pin.
ECC_RXERR	2	0b	RO	Unrecoverable ECC Error in an Rx SA Table Occurred. When set, it indicates that an unrecoverable ECC error occurred when accessing internally one Rx SA table. The ECC interrupt is set as well, until the Intel® Xeon® Processor D-1500 Product Family LAN controller is reset by software. When cleared, no ECC error occurred on the Rx SA table from the last time the Intel® Xeon® Processor D-1500 Product Family LAN controller has reset.
RESERVED	31:3	0x0	RSV	Reserved.

8.2.2.19 IPsec Registers

8.2.2.20 MACSec Registers

The MACSec registers are initialized at software reset. When MACSec is disabled, the MACSec statistic registers are meaningless and their values are unpredictable.

8.2.2.20.1 IPsec Tx Index - IPSTXIDX (0x00008900)



Field	Bit(s)	Init.	Access Type	Description
IPS_TX_EN	0	0b	RW	IPsec Tx Offload Enable Bit. When set, IPsec offload ability is enabled for Tx path. When cleared, IPsec offload ability is disabled for Tx path, regardless of the contents of the Tx SA table.
RESERVED	2:1	0x0	RSV	Reserved.
SA_IDX	12:3	0x0	RW	SA index for indirect access into the Tx SA table.
RESERVED	29:13	0x0	RSV	Reserved.
READ	30	0b	RW	Read Command. When set, the contents of the Tx SA table entry pointed by the SA_IDX field is loaded into the IPSTXKEY 0...3 and IPSTXSALT registers. Immediately self cleared by hardware once the entry contents has been loaded into the registers.
WRITE	31	0b	RW	Write command. When set, the contents of the IPSTXKEY 0...3 and IPSTXSALT registers are loaded into the Tx SA table entry pointed to by the SA_IDX field. Immediately self cleared by hardware once the entry contents have been loaded into the memory.

Note: Write and Read bits must not be set at the same time by software. IPS_TX_EN is RW, but it is RO if fused-off and/or if SECTXCTRL.SECTX_DIS is set to 1b.

8.2.2.20.2 IPsec Tx Key Registers - IPSTXKEY[n] (0x00008908 + 0x4*n, n=0...3)

Field	Bit(s)	Init.	Access Type	Description
AES_128_KEY	31:0	0x0	RW	4 bytes of 16 bytes key that has been read/written from/into the Tx SA entry pointed to by SA_IDX. n=0 contains the LSB of the key. n=3 contains the MSB of the key.

8.2.2.20.3 IPsec Tx Salt Register - IPSTXSALT (0x00008904)

Field	Bit(s)	Init.	Access Type	Description
AES_128_SALT	31:0	0x0	RW	4 bytes salt that has been read/written from/into the Tx SA entry pointed to by SA_IDX.

8.2.2.20.4 IPsec Rx Index - IPSRXIDX (0x00008E00)

Field	Bit(s)	Init.	Access Type	Description
IPS_RX_EN	0	0b	RW	IPsec Rx Offload Enable Bit. When set, IPsec offload ability is enabled for Rx path. When cleared, IPsec offload ability is disabled for Rx path, regardless of the contents of Rx SA tables.



Field	Bit(s)	Init.	Access Type	Description
TABLE	2:1	0x0	RW	TABLE select bits. 00b = No Rx SA table is accessed. 01b = IP address table is accessed. 10b = SPI table is accessed. 11b = KEY table is accessed.
TB_IDX	12:3	0x0	RW	Table index bits for indirect access into the Rx SA table selected by Table bits. When accessing the IP address table, only the 7 LS bits of this field are meaningful.
RESERVED	29:13	0x0	RSV	Reserved.
READ	30	0b	RW	READ Command. When set, the contents of the Rx SA table entry as pointed to by the [TABLE and TB_IDX] fields is loaded into the corresponding registers. Immediately self cleared by hardware once the entry contents have been loaded into the corresponding registers. For instance, if this bit is set together with Table=10 and TB_IDX=0x9, then the SPI value stored in entry 9 is loaded into the IPSRXSPI 0...3 registers. Rx SA registers related to another Rx SA table (e.g. IPSRXKEY 0...3 registers) must not be read when Table=01.
WRITE	31	0b	RW	WRITE Command. When set, the contents of the registers concerned by the Rx SA table pointed to by the Table field is loaded into the table entry pointed to by the TB_IDX field. Immediately self cleared by hardware once the entry contents have been loaded into the memory. For instance, if this bit is set together with Table=10 and TB_IDX=0x9, then the value written in IPSRXSPI 0...3 registers is loaded into the SPI table entry 9.

Note: WRITE and READ bits must not be set at the same time by software. IPS_RX_EN is RW, but it is RO if fused-off and/or if SECRXCTRL.SECRX_DIS is set to 1b, software is not allowed to write/read access registers that belong to different Rx SA tables without writing the IPSRXIDX register in between for setting the WRITE/READ bit. Refer to Rx SA tables access rules. Software should not make changes in the Rx SA tables while changing the IPSEC_EN bit.

8.2.2.20.5 IPsec Rx IP address Register - IPSRXIPADDR[n] (0x00008E04 + 0x4*n, n=0...3)

These registers are related to the IP address table.

Field	Bit(s)	Init.	Access Type	Description
IPADDR	31:0	0x0	RW	4 bytes of 16 bytes destination IP address for the associated Rx SA(s). n=0 contains the MSB for an IPv6 IP address. n=3 contains an IPv4 IP address or the LSB for an IPv6 IP address. For an IPv4 address, IPSRXIPADDR 0...2 must be written with zeros. Note: Field is defined in big endian (LS byte is first on the wire).

8.2.2.20.6 IPsec Rx SPI Register - IPSRXSPI (0x00008E14)

This register is related to the Rx SPI table.



Field	Bit(s)	Init.	Access Type	Description
SPI	31:0	0x0	RW	SPI field for the SPI entry. Note: Field is defined in big endian (LSB is first on the wire).

8.2.2.20.7 IPsec Rx SPI Register IP Index - IPSRXIPIDX (0x00008E18)

This register is related to the Rx SPI table.

Field	Bit(s)	Init.	Access Type	Description
IP_IDX	6:0	0x0	RW	IP Index. Index in the IP address table where the destination IP address associated to that SPI entry is found.
RESERVED	31:7	0x0	RSV	Reserved.

8.2.2.20.8 IPsec Rx Key Register - IPSRXKEY[n] (0x00008E1C + 0x4*n, n=0...3)

These registers are related to the Rx KEY table.

Field	Bit(s)	Init.	Access Type	Description
AES_128_KEY	31:0	0x0	RW	4 bytes of 16 bytes key of the KEY entry. n=0 contains the LSB of the key. n=3 contains the MSB of the key.

8.2.2.20.9 IPsec Rx Salt Register - IPSRXSALT (0x00008E2C)

This register is related to the Rx KEY table.

Field	Bit(s)	Init.	Access Type	Description
AES_128_SALT	31:0	0x0	RW	4 bytes salt associated to the KEY entry.

8.2.2.20.10 IPsec Rx Mode Register - IPSRXMOD (0x00008E30)

This register is related to the Rx KEY table.

Field	Bit(s)	Init.	Access Type	Description
VALID	0	0b	RW	Valid bit. When set, the KEY entry is valid. When cleared, the KEY entry is not valid.
RESERVED	1	0b	RSV	Reserved.
PROTO	2	0b	RW	IPsec Protocol Select. When set, the KEY entry off loads ESP packets. When cleared, the KEY entry off loads AH packets.



Field	Bit(s)	Init.	Access Type	Description
DECRYPT	3	0b	RW	Decryption Bit. When set, hardware performs decryption offload for this KEY entry. Meaningful only if PROTO bit is set (such as ESP mode).
IPV6	4	0b	RW	IPv6 Type. When set, only matched IPv6 packet are offloaded for that KEY entry. When cleared, only matched IPv4 packets are offloaded for that KEY entry.
RESERVED	31:5	0x0	RSV	Reserved.

8.2.2.21 VF Registers Mapping In The PF Space

8.2.2.21.1 VF Control Register - VFCTRL[n] (0x00000300 + 0x4*n, n=0...63; WO)

This register array is the mapping of the VFs VFCTRL registers.

Fields definitions are the same as defined in [Section 8.3.3.1.1](#).

8.2.2.21.2 VF Mailbox - VFMAILBOX[n] (0x00004C00 + 0x4*n, n=0...63; RW)

This register array is the mapping of the VFMAILBOX registers of the VFs.

Fields definitions are the same as defined in [Section 8.3.3.1.3](#).

8.2.2.21.3 PF Mailbox Memory - PFMBMEM[n,m] (0x00013000 + 0x4*n + 0x40*m, n=0...15, m=0...63; RW)

Mailbox memory for PF and VF driver communications. The mailbox size for each VM is 64 bytes accessed by 16 x 32-bit registers. Locations can be accessed as 32- or 64-bit words. This is the mapping of this memory in the PF space.

Fields definitions are the same as defined in [Section 8.3.3.1.4](#).

8.2.2.21.4 VF Extended Interrupt Cause - VFEICR[n] (0x00000B00 + 0x4*n, n=0...63; RC/W1C)

Fields definitions are the same as defined in [Section 8.3.3.2.1](#).

8.2.2.21.5 VF Extended Interrupt Cause Set - VFEICS[n] (0x00000C00 + 0x4*n, n=0...63; WO)

Fields definitions are the same as defined in [Section 8.3.3.2.2](#).

8.2.2.21.6 VF Extended Interrupt Mask Set/Read - VFEIMS[n] (0x00000D00 + 0x4*n, n=0...63; RWS)

Fields definitions are the same as defined in [Section 8.3.3.2.3](#).

8.2.2.21.7 VF Extended Interrupt Mask Clear - VFEIMC[n] (0x00000E00 + 0x4*n, n=0...63; WO)

Fields definitions are the same as defined in [Section 8.3.3.2.4](#).



8.2.2.21.8 VF Extended Interrupt Auto Mask Enable - VFEIAM[n] (0x00004D00 + 0x4*n, n=0...63; RW)

Fields definitions are the same as defined in [Section 8.3.3.2.5](#).

8.2.2.21.9 VF Interrupt Vector Allocation Registers - VFIVAR[n] (0x00012500 + 0x4*n, n=0...63; RW)

These registers map VF interrupt causes into MSI-X vectors (PF mirror). See the Mapping of Interrupt Causes section for more details.

Fields definitions are the same as defined in [Section 8.3.3.2.7](#).

8.2.2.21.10 VF Interrupt Vector Allocation Registers Misc - VFIVAR_MISC[n] (0x00004E00 + 0x4*n, n=0...63; RW)

These registers map the mailbox interrupt into MSI-X vector (PF mirror). See the Mapping of Interrupt Causes section for more details.

Fields definitions are the same as defined in [Section 8.3.3.2.8](#).

8.2.2.21.11 VF Multiple Receive Queues Command Register - VFMRQC[n] (0x00003400 + 0x4*n, n=0...63; RW)

PF mirror of VFMRQC registers of the VFs.

Fields definitions are the same as defined in [Section 8.3.3.3.10](#).

8.2.2.21.12 VF Good Packets Transmitted Count - VFGPTC[n] (0x00008300 + 0x4*n, n=0...63; RO)

PF mirror of VFGPTC.

Fields definitions are the same as defined in [Section 8.3.3.6.2](#).

8.2.2.21.13 VF Good Octets Transmitted Count LSB - VFGOTC_LSB[n] (0x00008400 + 0x8*n, n=0...63; RO)

PF mirror of VFGOTC_LSB.

Fields definitions are the same as defined in [Section 8.3.3.6.5](#).

8.2.2.21.14 VF Good Octets Transmitted Count MSB - VFGOTC_MSB[n] (0x00008404 + 0x8*n, n=0...63; RO)

PF mirror of VFGOTC_MSB.

Fields definitions are the same as defined in [Section 8.3.3.6.6](#).

8.2.2.21.15 VF Good Octets Received Count Low - VFGORC_LSB[n] (0x00001020 + 0x40*n, n=0...63; RW)

PF mirror of VFGORC_LSB.

Fields definitions are the same as defined in [Section 8.3.3.6.3](#).



8.2.2.21.16 VF Good Packets Received Count - VFGPRC[n] (0x0000101C + 0x40*n, n=0...63; RW)

PF mirror of VFGPRC.

Fields definitions are the same as defined in [Section 8.3.3.6.1](#).

8.2.2.21.17 VF Good Octets Received Count High - VFGORC_MSB[n] (0x0000D020 + 0x40*n, n=0...63; RW)

PF mirror of VFGORC_MSB.

Fields definitions are the same as defined in [Section 8.3.3.6.4](#).

8.2.2.21.18 VF Multicast Packets Received Count - VFMPRC[n] (0x0000D01C + 0x40*n, n=0...63; RO)

PF mirror of VFMPRC.

Fields definitions are the same as defined in [Section 8.3.3.6.7](#).

8.2.3 BAR3 Registers Summary

Table 8.3. BAR3 Registers Summary


Offset / Alias Offset	Abbreviation	Name	Block	Section
MSI-X Table Registers				
0x00000000 + 0x10*n, n=0...255	MSIXTADD[n]	MSI-X Table Entry Lower Address	MSIX	Section 8.2.4.1.1
0x00000004 + 0x10*n, n=0...255	MSIXTUADD[n]	MSI-X Table Entry Upper Address	MSIX	Section 8.2.4.1.2
0x00000008 + 0x10*n, n=0...255	MSIXTMSG[n]	MSI-X Table Entry Message	MSIX	Section 8.2.4.1.3
0x0000000C + 0x10*n, n=0...255	MSIXVCTRL[n]	MSI-X Table Entry Vector Control	MSIX	Section 8.2.4.1.4

8.2.4 Detailed Register Description - PF BAR3

8.2.4.1 MSI-X Table Registers

The MSI-X capability is described in [Section 9.3.2.1](#). The MSI-X table is described in [Section 9.3.2.1.2](#) and the Pending Bit Array (PBA) is described in [Section 9.3.2.1.3](#). These registers are located in the MSI-X BAR.

8.2.4.1.1 MSI-X Table Entry Lower Address - MSIXTADD[n] (0x00000000 + 0x10*n, n=0...255)



Field	Bit(s)	Init.	Access Type	Description
MSIXTADD10	1:0	0x0	RW	Message Address. For proper Dword alignment, software must always write zeroes to these two bits; otherwise, the result is undefined. The state of these bits after reset must be 0b. These bits are permitted to be read-only or read/write.
MSIXTADD	31:2	0x0	RW	Message Address. System-specified message lower address. For MSI-X messages, the contents of this field from an MSI-X table entry specifies the lower portion of the Dword-aligned address (AD[31:02]) for the memory write transaction. This field is read/write.



8.2.4.1.2 MSI-X Table Entry Upper Address - MSIXTUADD[n] (0x00000004 + 0x10*n, n=0...255)

Field	Bit(s)	Init.	Access Type	Description
MSIXTUADD	31:0	0x0	RW	Message Upper Address. System-specified message upper address bits. If this field is zero, Single Address Cycle (SAC) messages are used. If this field is non-zero, Dual Address Cycle (DAC) messages are used. This field is read/write.

8.2.4.1.3 MSI-X Table Entry Message - MSIXTMSG[n] (0x00000008 + 0x10*n, n=0...255)

Field	Bit(s)	Init.	Access Type	Description
MSIXTMSG	31:0	0x0	RW	Message Data. System-specified message data. For MSI-X messages, the contents of this field from an MSI-X table entry specifies the data driven on AD[31:0] during the memory write transaction's data phase. This field is read/write.

8.2.4.1.4 MSI-X Table Entry Vector Control - MSIXVCTRL[n] (0x0000000C + 0x10*n, n=0...255)

Field	Bit(s)	Init.	Access Type	Description
MSIXVCTRL	0	1b	RW	Mask Bit. When this bit is set, the function is prohibited from sending a message using this MSI-X table entry. However, any other MSI-X table entries programmed with the same vector are still capable of sending an equivalent message unless they are also masked. This bit's state after reset is 1b (entry is masked). This bit is read/write.
RESERVED	31:1	0x0	RSV	Reserved. After reset, the state of these bits must be 0b. However, for potential future use, software must preserve the value of these reserved bits when modifying the value of other Vector Control bits. If software modifies the value of these reserved bits, the result is undefined.

8.3 Device Registers - VF

8.3.1 VF Registers Mapping PF Space

Abbreviation	Virtual Address	Physical Address
VFSRRCTL	0x00001014 + 0x40*n, n=0...7	0x00002100 + 0x4*n, n=0...15
VFTPH_RXCTRL	0x0000100C + 0x40*n, n=0...7	0x00002200 + 0x4*n, n=0...15



8.3.2 BAR0 Registers Summary

Table 8.4. BAR0 Registers Summary

Offset / Alias Offset	Abbreviation	Name	Block	Section
General Control Registers - VF				
0x00000000	VFCTRL	VF Control Register	Target	Section 8.3.3.1.1
0x00000008	VFSTATUS	VF Device Status Register	Target	Section 8.3.3.1.2
0x000002FC	VFMAILBOX	VF Mailbox	Target	Section 8.3.3.1.3
0x00000200 + 0x4*n, n=0...15	VFMBMEM[n]	VF Mailbox Memory	Target	Section 8.3.3.1.4
0x00000010	VFLINKS	Link Status Register	MAC	Section 8.3.3.1.5
Interrupt Registers - VF				
0x00000100	VFEICR	VF Extended Interrupt Cause	Interrupt	Section 8.3.3.2.1
0x00000104	VFEICS	VF Extended Interrupt Cause Set	Interrupt	Section 8.3.3.2.2
0x00000108	VFEIMS	VF Extended Interrupt Mask Set/Read	Interrupt	Section 8.3.3.2.3
0x0000010C	VFEIMC	VF Extended Interrupt Mask Clear	Interrupt	Section 8.3.3.2.4
0x00000114	VFEIAM	VF Extended Interrupt Auto Mask Enable	Interrupt	Section 8.3.3.2.5
0x00000820 + 0x4*n, n=0...1	VFEITR[n]	VF Extended Interrupt Throttle Registers	Interrupt	Section 8.3.3.2.6
0x00000120 + 0x4*n, n=0...3	VFIVAR[n]	VF Interrupt Vector Allocation Registers	Interrupt	Section 8.3.3.2.7
0x00000140	VFIVAR_MISC	VF Interrupt Vector Allocation Registers Misc	Interrupt	Section 8.3.3.2.8
0x00000180 + 0x4*n, n=0...1	VFRSCINT[n]	VF RSC Enable Interrupt	Interrupt	Section 8.3.3.2.9
0x00000148	VFPBACL	VF MSI-X PBA Clear	PCIe	Section 8.3.3.2.10
Receive Registers - VF				
0x00001000 + 0x40*n, n=0...7	VFRDBAL[n]	VF Receive Descriptor Base Address Low	DMA_RX	Section 8.3.3.3.1
0x00001004 + 0x40*n, n=0...7	VFRDBAH[n]	VF Receive Descriptor Base Address High	DMA_RX	Section 8.3.3.3.2
0x00001008 + 0x40*n, n=0...7	VFRDLEN[n]	VF Receive Descriptor Length	DMA_RX	Section 8.3.3.3.3
0x00001010 + 0x40*n, n=0...7	VFRDH[n]	VF Receive Descriptor Head	DMA_RX	Section 8.3.3.3.4
0x00001018 + 0x40*n, n=0...7	VFRDT[n]	VF Receive Descriptor Tail	DMA_RX	Section 8.3.3.3.5
0x00001028 + 0x40*n, n=0...7	VFRXDCTL[n]	VF Receive Descriptor Control	DMA_RX	Section 8.3.3.3.6
0x00001014 + 0x40*n, n=0...7	VFSRRCTL[n]	VF Split and Replication Receive Control Registers	DMA_RX	Section 8.3.3.3.7
0x00000300	VFPSRTYPE	VF Replication Packet Split Receive Type Register	DBU_RX	Section 8.3.3.3.8
0x0000102C + 0x40*n, n=0...7	VFRSCCTL[n]	VF RSC Control	DMA_RX	Section 8.3.3.3.9



Table 8.4. BARO Registers Summary (Continued)

Offset / Alias Offset	Abbreviation	Name	Block	Section
0x00003000	VFMRQC	VF Multiple Receive Queues Command Register	DBU_RX	Section 8.3.3.3.10
0x00003100 + 0x4*n, n=0...9	VFRSSRK[n]	VF RSS Random Key Register	DBU_RX	Section 8.3.3.3.11
0x00003200 + 0x4*n, n=0...15	VFRETA[n]	VF Redirection Table	DBU_RX	Section 8.3.3.3.12
Transmit Registers - VF				
0x00002000 + 0x4*n, n=0...7	VFTDBAL[n]	VF Transmit Descriptor Base Address Low	DMA_TX	Section 8.3.3.4.1
0x00002004 + 0x4*n, n=0...7	VFTDBAH[n]	VF Transmit Descriptor Base Address High	DMA_TX	Section 8.3.3.4.2
0x00002008 + 0x4*n, n=0...7	VFTDLEN[n]	VF Transmit Descriptor Length	DMA_TX	Section 8.3.3.4.3
0x00002010 + 0x4*n, n=0...7	VFTDH[n]	VF Transmit Descriptor Head	DMA_TX	Section 8.3.3.4.4
0x00002018 + 0x4*n, n=0...7	VFTDT[n]	VF Transmit Descriptor Tail	DMA_TX	Section 8.3.3.4.5
0x00002028 + 0x4*n, n=0...7	VFTXDCTL[n]	VF Transmit Descriptor Control	DMA_TX	Section 8.3.3.4.6
0x00002038 + 0x4*n, n=0...7	VFTDWBAL[n]	VF Tx Descriptor Completion Write Back Address Low	DMA_TX	Section 8.3.3.4.7
0x0000203C + 0x4*n, n=0...7	VFTDWBAL[n]	VF Tx Descriptor Completion Write Back Address High	DMA_TX	Section 8.3.3.4.8
TPH Registers - VF				
0x0000100C + 0x4*n, n=0...7	VFTPH_RXCTRL[n]	VF Rx TPH Control Register	DMA_RX	Section 8.3.3.5.1
0x0000200C + 0x4*n, n=0...7	VFTPH_TXCTRL[n]	VF Tx TPH Control Registers	DMA_TX	Section 8.3.3.5.2
Statistic Register Descriptions - VF				
0x0000101C	VFGPRC	VF Good Packets Received Count	DMA_RX	Section 8.3.3.6.1
0x0000201C	VFGPTC	VF Good Packets Transmitted Count	STAT	Section 8.3.3.6.2
0x00001020	VFGORC_LSB	VF Good Octets Received Count Low	DMA_RX	Section 8.3.3.6.3
0x00001024	VFGORC_MSB	VF Good Octets Received Count High	DMA_RX	Section 8.3.3.6.4
0x00002020	VFGOTC_LSB	VF Good Octets Transmitted Count LSB	STAT	Section 8.3.3.6.5
0x00002024	VFGOTC_MSB	VF Good Octets Transmitted Count MSB	STAT	Section 8.3.3.6.6
0x00001034	VFMPRC	VF Multicast Packets Received Count	DMA_RX	Section 8.3.3.6.7



8.3.3 Detailed Register Description - VF BAR0

8.3.3.1 General Control Registers - VF

8.3.3.1.1 VF Control Register - VFCTRL (0x00000000)

Field	Bit(s)	Init.	Access Type	Description
RESERVED	25:0	0x0	RSV	Reserved.
RST	26	0b	WO	VF Reset. This bit performs a reset of the queue enable and the interrupt registers of the VF.
RESERVED	31:27	0x0	RSV	Reserved.

8.3.3.1.2 VF Device Status Register - VFSTATUS (0x00000008; RO)

This register is a mirror of the PF status register.

Fields definitions are the same as defined in [Section 8.2.2.1.2](#).

8.3.3.1.3 VF Mailbox - VFMAILBOX (0x000002FC)

This register is cleared by VLFR (except to the RSTI bit).

Field	Bit(s)	Init.	Access Type	Description
REQ	0	0b	WO	Request for PF Ready. Setting this bit, causes an interrupt to the PF. This bit always reads as 0b. Setting this bit sets the corresponding bit in the VFREQ field in PFMBICR register.
ACK	1	0b	WO	PF message received. Setting this bit, causes an interrupt to the PF. This bit always reads as 0b. Setting this bit sets the corresponding bit in the VFACK field in PFMBICR register.
VFU	2	0b	RW	Buffer is Taken by VF. This bit can be set only if the PFU bit is cleared and is mirrored in the VFU bit of the PF Mailbox register.
PFU	3	0b	RW	Buffer is taken by PF. This bit is RO for the VF and is a mirror of the PFU bit of the PF Mailbox register.
PFSTS	4	0b	RC	PF wrote a message in the mailbox.
PFAK	5	0b	RC	PF acknowledged the VF previous message.
RSTI	6	1b	RO	Indicates that the PF reset the shared resources and the reset sequence is in progress. This bit is not affected by VLFR.
RSTD	7	0b	RC	Indicates that a PF software reset completed.
RESERVED	31:8	0x0	RSV	Reserved.

8.3.3.1.4 VF Mailbox Memory - VFMBMEM[n] (0x00000200 + 0x4*n, n=0...15)

Mailbox memory for PF and VF driver communications. The mailbox size for each VM is 64 bytes accessed by 16 x 32-bit registers. Locations can be accessed as 32- or 64-bit words.



Field	Bit(s)	Init.	Access Type	Description
MAILBOX_DATA	31:0	X	RW	Mailbox data is composed of 16 x 4 byte registers.

8.3.3.1.5 Link Status Register - VFLINKS (0x00000010; RO)

This register is the mapping of the PF's LINKS register.

Fields definitions are the same as defined in [Section 8.2.2.11.7](#).

8.3.3.2 Interrupt Registers - VF

8.3.3.2.1 VF Extended Interrupt Cause - VFEICR (0x00000100)

Field	Bit(s)	Init.	Access Type	Description
MSIX	2:0	0x0	RW1C	Indicates an interrupt cause mapped to MSI-X vectors 2:0.
RESERVED	31:3	0x0	RSV	Reserved.

8.3.3.2.2 VF Extended Interrupt Cause Set - VFEICS (0x00000104)

Field	Bit(s)	Init.	Access Type	Description
MSIX	2:0	0x0	WO	Sets to the corresponding EICR bit of MSI-X vectors 2:0.
RESERVED	31:3	0x0	RSV	Reserved.

8.3.3.2.3 VF Extended Interrupt Mask Set/Read - VFEIMS (0x00000108)

Field	Bit(s)	Init.	Access Type	Description
MSIX	2:0	0x0	RWS	Sets the Mask bit for the corresponding EICR bit of MSI-X vectors 2:0.
RESERVED	31:3	0x0	RSV	Reserved.

8.3.3.2.4 VF Extended Interrupt Mask Clear - VFEIMC (0x0000010C)

Field	Bit(s)	Init.	Access Type	Description
MSIX	2:0	0x0	WO	Clears the Mask bit for the corresponding EICR bit of MSI-X vectors 2:0.
RESERVED	31:3	0x0	RSV	Reserved.

**8.3.3.2.5 VF Extended Interrupt Auto Mask Enable - VFEIAM (0x00000114)**

Field	Bit(s)	Init.	Access Type	Description
MSIX	2:0	0x0	RW	Auto Mask bit for the corresponding EICR bit of MSI-X vectors 2:0.
RESERVED	31:3	0x0	RSV	Reserved.

8.3.3.2.6 VF Extended Interrupt Throttle Registers - VFEITR[n] (0x00000820 + 0x4*n, n=0...1; RW)

Fields definitions are the same as defined in [Section 8.2.2.4.34](#).

8.3.3.2.7 VF Interrupt Vector Allocation Registers - VFIVAR[n] (0x00000120 + 0x4*n, n=0...3)

These registers map VF interrupt causes into MSI-X vectors. See the Mapping of Interrupt Causes Section for more details.

Field	Bit(s)	Init.	Access Type	Description
INT_ALLOC_0	0	0b	RW	Defines the MSI-X vector (0 or 1) assigned to Rx queue 2*N for IVAR 'N' register (N=0...3).
RESERVED	6:1	0x0	RSV	Reserved.
INT_ALLOC_V AL_0	7	0b	RW	Valid bit for INT_Alloc[0].
INT_ALLOC_1	8	0b	RW	Defines the MSI-X vector (0 or 1) assigned to Tx queue 2*N for IVAR 'N' register (N=0...3).
RESERVED	14:9	0x0	RSV	Reserved.
INT_ALLOC_V AL_1	15	0b	RW	Valid bit for INT_Alloc[1].
INT_ALLOC_2	16	0b	RW	Defines the MSI-X vector (0 or 1) assigned to Rx queue 2*N+1 for IVAR 'N' register (N=0...3).
RESERVED	22:17	0x0	RSV	Reserved.
INT_ALLOC_V AL_2	23	0b	RW	Valid bit for INT_Alloc[2].
INT_ALLOC_3	24	0b	RW	Defines the MSI-X vector (0 or 1) assigned to Tx queue 2*N+1 for IVAR 'N' register (N=0...3).
RESERVED	30:25	0x0	RSV	Reserved.
INT_ALLOC_V AL_3	31	0b	RW	Valid bit for INT_Alloc[3].

8.3.3.2.8 VF Interrupt Vector Allocation Registers Misc - VFIVAR_MISC (0x00000140)

This register maps the mailbox interrupt into MSI-X vector. See the Mapping of Interrupt Causes section for more details.

Field	Bit(s)	Init.	Access Type	Description
INT_ALLOC_0	1:0	X	RW	Defines the MSI-X vector assigned to the mailbox interrupt.



Field	Bit(s)	Init.	Access Type	Description
RESERVED	6:2	0x0	RSV	Reserved.
INT_ALLOC_V AL_0	7	0b	RW	Valid bit for INT_Alloc[0].
RESERVED	31:8	0x0	RSV	Reserved.

8.3.3.2.9 VF RSC Enable Interrupt - VFRSCINT[n] (0x00000180 + 0x4*n, n=0...1; RW)

Fields definitions are the same as defined in [Section 8.2.2.4.37](#).

8.3.3.2.10 VF MSI-X PBA Clear - VFPBACL (0x00000148)

Field	Bit(s)	Init.	Access Type	Description
RESERVED	2:0	0x0	RW1C	Reserved.
RESERVED	31:3	0x0	RSV	Reserved.

8.3.3.3 Receive Registers - VF

8.3.3.3.1 VF Receive Descriptor Base Address Low - VFRDBAL[n] (0x00001000 + 0x40*n, n=0...7; RW)

Fields definitions are the same as defined in [Section 8.2.2.6.27](#).

8.3.3.3.2 VF Receive Descriptor Base Address High - VFRDBAH[n] (0x00001004 + 0x40*n, n=0...7; RW)

Fields definitions are the same as defined in [Section 8.2.2.6.28](#).

8.3.3.3.3 VF Receive Descriptor Length - VFRDLEN[n] (0x00001008 + 0x40*n, n=0...7; RW)

Fields definitions are the same as defined in [Section 8.2.2.6.29](#).

8.3.3.3.4 VF Receive Descriptor Head - VFRDH[n] (0x00001010 + 0x40*n, n=0...7; RO)

Fields definitions are the same as defined in [Section 8.2.2.6.30](#).

8.3.3.3.5 VF Receive Descriptor Tail - VFRDT[n] (0x00001018 + 0x40*n, n=0...7; RW) DMA_RX

Fields definitions are the same as defined in [Section 8.2.2.6.31](#).

8.3.3.3.6 VF Receive Descriptor Control - VFRXDCTL[n] (0x00001028 + 0x40*n, n=0...7; RW)

Fields definitions are the same as defined in [Section 8.2.2.6.32](#).



8.3.3.3.7 VF Split and Replication Receive Control Registers - VFSRRCTL[n] (0x00001014 + 0x40*n, n=0...7; RW)

Fields definitions are the same as defined in Section 8.2.2.6.34.

8.3.3.3.8 VF Replication Packet Split Receive Type Register - VFPSRTYPE (0x00000300; RW)

Fields definitions are the same as defined in Section 8.2.2.6.4.

8.3.3.3.9 VF RSC Control - VFRSCCTL[n] (0x0000102C + 0x40*n, n=0...7; RW)

Fields definitions are the same as defined in Section 8.2.2.6.38.

8.3.3.3.10 VF Multiple Receive Queues Command Register - VFMRQC (0x00003000)

Field	Bit(s)	Init.	Access Type	Description
RSSE	0	0b	RW	RSS Enable. This bit enables the VF RSS mechanism. 0b = VF RSS disabled. 1b = VF RSS enable.
RESERVED	15:1	0x0	RSV	Reserved.
RSS_FIELD_ENABLE	31:16	0x0	RW	Each bit, when set, enables a specific field selection to be used by the hash function. Several bits can be set at the same time. Bit[16] = Enable TcpIPv4 hash function. Bit[17] = Enable IPv4 hash function. Bit[18] = Reserved. Bit[19] = Reserved. Bit[20] = Enable IPv6 hash function. Bit[21] = Enable TcpIPv6 hash function. Bit[22] = Enable UdpIPv4. Bit[23] = Enable UdpIPv6. Bit[24] = Reserved. Bits[31:25] = Reserved; 0x0. Note that on tunnel packets IPv4-IPv6, only the IPv4 header might be used for the RSS filtering.

8.3.3.3.11 VF RSS Random Key Register - VFRSSRK[n] (0x00003100 + 0x4*n, n=0...9; RW)


PF mirror of VFRSSK registers of the VFs.

Fields definitions are the same as defined in Section 8.2.2.6.20.

8.3.3.3.12 VF Redirection Table - VFRETA[n] (0x00003200 + 0x4*n, n=0...15; RW)

VF version of VFRETA registers. The redirection table has 64 entries in 16 registers.

Fields definitions are the same as defined in Section 8.2.2.6.21.



8.3.3.4 Transmit Registers - VF

8.3.3.4.1 VF Transmit Descriptor Base Address Low - VFTDBAL[n] (0x00002000 + 0x40*n, n=0...7; RW)

See the TDBAL description.

Fields definitions are the same as defined in [Section 8.2.2.7.5](#).

8.3.3.4.2 VF Transmit Descriptor Base Address High - VFTDBAH[n] (0x00002004 + 0x40*n, n=0...7; RW)

See the TDBAH description.

Fields definitions are the same as defined in [Section 8.2.2.7.6](#).

8.3.3.4.3 VF Transmit Descriptor Length - VFTDLEN[n] (0x00002008 + 0x40*n, n=0...7; RW)

See the TDLEN description.

Fields definitions are the same as defined in [Section 8.2.2.7.7](#).

8.3.3.4.4 VF Transmit Descriptor Head - VFTDH[n] (0x00002010 + 0x40*n, n=0...7; RO)

See the TDH description.

Fields definitions are the same as defined in [Section 8.2.2.7.8](#).

8.3.3.4.5 VF Transmit Descriptor Tail - VFTDT[n] (0x00002018 + 0x40*n, n=0...7; RW)

See the TDT description.

Fields definitions are the same as defined in [Section 8.2.2.7.9](#).

8.3.3.4.6 VF Transmit Descriptor Control - VFTXDCTL[n] (0x00002028 + 0x40*n, n=0...7; RW)

See the TXDCTL description.

Fields definitions are the same as defined in [Section 8.2.2.7.10](#).

8.3.3.4.7 VF Tx Descriptor Completion Write Back Address Low - VFTDWBAL[n] (0x00002038 + 0x40*n, n=0...7; RW)

See the TDWBAL description.

Fields definitions are the same as defined in [Section 8.2.2.7.11](#).

8.3.3.4.8 VF Tx Descriptor Completion Write Back Address High - VFTDWBAH[n] (0x0000203C + 0x40*n, n=0...7; RW)

See the TDWBAH description.

Fields definitions are the same as defined in [Section 8.2.2.7.12](#).



8.3.3.5 TPH Registers - VF

8.3.3.5.1 VF Rx TPH Control Register - VFTPH_RXCTRL[n] (0x0000100C + 0x40*n, n=0...7; RW)

See the TPH_RXCTRL description.

Fields definitions are the same as defined in Section 8.2.2.8.2.

8.3.3.5.2 VF Tx TPH Control Registers - VFTPH_TXCTRL[n] (0x0000200C + 0x40*n, n=0...7; RW)

See the TPH_TXCTRL description.

Fields definitions are the same as defined in Section 8.2.2.8.3.

8.3.3.6 Statistic Register Descriptions - VF

Registers in this section are RO by VF and RW by PF. Statistics are reset by PF clearing the register.

8.3.3.6.1 VF Good Packets Received Count - VFGPRC (0x0000101C)

Field	Bit(s)	Init.	Access Type	Description
VFGPRC	31:0	0x0	RW	Number of good packets received for this VF (of any length). This counter includes loopback packets or replications of multicast packets. The counter is not cleared on read. Furthermore, the register is a cyclic counter incrementing from 0xFFFF to 0x0000.

8.3.3.6.2 VF Good Packets Transmitted Count - VFGPTC (0x0000201C)

Field	Bit(s)	Init.	Access Type	Description
VFGPTC	31:0	0x0	RO	Number of good packets sent by the queues allocated to this VF. A packet is considered as transmitted if it was forwarded to the MAC unit for transmission to the network and/or is accepted by the internal Tx to Rx switch enablement logic. Packets dropped due to anti-spoofing filtering or loopback packets that are rejected by the Tx to Rx switch are not counted. The counter is not cleared on read. Furthermore, the register is a cyclic counter incrementing from 0xFFFF to 0x0000.

8.3.3.6.3 VF Good Octets Received Count Low - VFGORC_LSB (0x00001020)

Field	Bit(s)	Init.	Access Type	Description
VFGORC_LSB	31:0	0x0	RW	Number of good octets received (32 LSB of a 36-bit counter) by the queues allocated to this VF. The counter includes loopback packets or replications of multicast packets. This register includes bytes received in a packet from the <Destination Address> field through the <CRC> field, inclusively. Octets are counted on the VF interface rather than on the network interface. For example, MACsec octets are not counted. Bytes of RSC are counted before coalescing. The counter is not cleared on read. Furthermore, the register is a cyclic counter incrementing from 0xFFFF to 0x0000.



8.3.3.6.4 VF Good Octets Received Count High - VFGORC_MSB (0x00001024)

Field	Bit(s)	Init.	Access Type	Description
VFGORC_MSB	3:0	0x0	RW	Number of good octets received (4 MSB of a 36-bit counter) by the queues allocated to this VF. See the VFGORC_LSB register description for more details. The counter is not cleared on read. Furthermore, the register is a cyclic counter incrementing from 0xF to 0x0.
RESERVED	31:4	0x0	RSV	Reserved.

8.3.3.6.5 VF Good Octets Transmitted Count LSB - VFGOTC_LSB (0x00002020)


Field	Bit(s)	Init.	Access Type	Description
VFGOTC_LSB	31:0	0x0	RO	Number of good octets transmitted (32 LSB of a 36-bit counter) by the queues allocated to this VF. This register includes bytes transmitted in a packet from the <Destination Address> field through the <CRC> field, inclusively. This register counts octets of the packets counted by the VFGPTC register. Octets are counted on the VF interface rather than on the network interface. For example, MACsec octets are not counted. The counter is not cleared on read. Furthermore, the register is a cyclic counter incrementing from 0xFFFF to 0x0000.

8.3.3.6.6 VF Good Octets Transmitted Count MSB - VFGOTC_MSB (0x00002024)

Field	Bit(s)	Init.	Access Type	Description
VFGOTC_MSB	3:0	0x0	RO	Number of good octets transmitted (4 MSB of a 36-bit counter) by the queues allocated to this VF. See the description of the VFGOTC_LSB register for more details. The counter is not cleared on read. Furthermore, the register is a cyclic counter incrementing from 0xF to 0x0.
RESERVED	31:4	0x0	RSV	Reserved.

8.3.3.6.7 VF Multicast Packets Received Count - VFMPRC (0x00001034)

Field	Bit(s)	Init.	Access Type	Description
VFMPRC	31:0	0x0	RO	Number of multicast good packets received by this VF (of any length) that pass the Ethernet MAC address filtering (excluding broadcast packets). The counter does not count received flow control packets. This register increments only if receives are enabled. This register does not count packets counted by the Missed Packet Count (MPC) register. This counter also includes loopback packets or replications of multicast packets. The counter is not cleared on read. Furthermore, the register is a cyclic counter incrementing from 0xFFFF to 0x0000.





9.0 PCIe Programming Interface

9.1 Overview

The Intel® Xeon® Processor D-1500 Product Family LAN controller is a multi-function device with the following functions:

- LAN 0
- LAN 1

Different parameters affect how LAN functions are exposed on PCIe.

Both functions contain the following regions of the PCI configuration space (some of them are enabled by NVM settings as detailed in the following sections):

- Mandatory PCI configuration registers - [Section 9.2.2](#)
- Legacy PCI capabilities - [Section 9.2.3](#)
 - Power management capabilities
 - MSI / MSI-X capabilities
 - Vital Product Data (VPD) capability
- PCIe extended capabilities - [Section 9.2.4](#)
 - Advanced Error Reporting (AER)
 - Serial ID
 - Alternate requester ID.
 - Single root IOV (SR-IOV)
 - Latency Tolerance Reporting (LTR)
 - TPH Requester
 - Access Control Services (ACS)
 - Secondary PCIe

9.1.1 PCIe Configuration Space in an Integrated I/O Interface Connected System

When the integrated Broadwell-DE LAN controller is connected using the integrated I/O interface and not an actual physical PCIe link, the PCIe configurations are subject to some modifications. One such example are the PCIe link registers that are not relevant anymore and should reflect a static setup.

In the following sections these differences are highlighted where relevant.



9.1.2 Register Attributes

The following table lists the register attributes used in this section.

RD/WR	Description
RO	Read-only register: Register bits are read-only and cannot be altered by software.
RW	Read-write register: Register bits are read-write and can be either set or reset.
RW1C	Read-only status, Write-1b-to-clear status register, Writing a 0b to RW1C bits has no effect.
ROS	Read-only register with sticky bits: Register bits are read-only and cannot be altered by software. Bits are not cleared by reset and can only be reset with the PWRGOOD signal. Devices that consume AUX power are not allowed to reset sticky bits when AUX power consumption (either via AUX power or PME Enable) is enabled.
RWS	Read-write register: Register bits are read-write and can be either set or reset by software to the desired state. Bits are not cleared by reset and can only be reset with the PWRGOOD signal. Devices that consume AUX power are not allowed to reset sticky bits when AUX power consumption (either via AUX power or PME Enable) is enabled.
RW1CS	Read-only status, Write-1b-to-clear status register: Register bits indicate status when read, a set bit, indicating a status event, can be cleared by writing a 1b to it. Writing a 0b to RW1C bits has no effect. Bits are not cleared by reset and can only be reset with the PWRGOOD signal. Devices that consume AUX power are not allowed to reset sticky bits when AUX power consumption (either via AUX power or PME Enable) is enabled.
HwInit	Hardware initialized: Register bits are initialized by firmware or hardware mechanisms such as pin strapping or serial NVM. Bits are read-only after initialization and can only be reset (for write-once by firmware) with the PWRGOOD signal.
RsvdP	Reserved and preserved: Reserved for future read/write implementations; software must preserve value read for writes to these bits.
RsvdZ	Reserved and zero: Reserved for future RW1C implementations; software must use 0b for writes to these bits.

9.2 PCIe Register Map

9.2.1 PCIe Configuration Space Summary

Table 9.1 lists the PCIe configuration registers while their detailed description is given in the sections that follow. PCI configuration fields in the summary table are presented by the following marking:

- Fields that have meaningful default values are indicated in parenthesis — (**value**).
- Dotted fields indicates the same value for both LAN functions
- Light-blue fields indicate read-only fields (loaded from the NVM)
- Magenta fields indicate hard-coded values.
- Other fields contain RW attributes.



Table 9.1. PCI Configuration Registers Map - LAN Functions

Section	Byte Offset	Byte 3	Byte 2	Byte 1	Byte 0	
Mandatory PCI Register	0x0	Device ID		Vendor ID		
	0x4	Status Register		Control Register		
	0x8	Class Code (0x020000/0x010000)			Revision ID	
	0xC	Reserved	Header Type (0x0/0x80)	Latency Timer	Cache Line Size (0x0)	
	0x10	Base Address Register 0				
	0x14	Base Address Register 1				
	0x18	Base Address Register 2				
	0x1C	Base Address Register 3				
	0x20	Base Address Register 4				
	0x24	Base Address Register 5				
	0x28	CardBus CIS pointer (0x0000)				
	0x2C	Subsystem ID		Subsystem Vendor ID		
	0x30	Expansion ROM Base Address				
	0x34	Reserved			Cap Ptr (0x40)	
	0x38	Reserved				
	Power Management Capability	0x3C	Max Latency (0x00)	Min Grant (0x00)	Interrupt Pin (0x01...0x04)	Interrupt Line (0x00)
0x40		Power Management Capabilities		Next Pointer (0x50)	Capability ID (0x01)	
PCI / PCIe Capabilities	0x44	Data	Bridge Support Extensions	Power Management Control & Status		
	0x50...0x67	MSI Capability				
	0x70...0x7B	MSI-X Capability				
	0xA0...0xDB	PCIe Capability				
	0xE0...0xE7	VPD Capability				
Extended PCIe Configuration	0x100...0x12B	AER Capability				
	0x140...0x14B	Serial ID Capability				
	0x150...0x157	ARI Capability				
	0x160...0x19F	SR-IOV Capability				
	0x1A0...0x1AB	TPH Capability				
	0x1B0...0x1B7	ACS Capability				
	0x1C0...0x1C7	LTR Capability				
	0x1D0...0x1E3	Secondary PCIe Extended Capability				




Table 9.2. PCIe Configuration Registers Map - Dummy Function

Section	Byte Offset	Byte 3	Byte 2	Byte 1	Byte 0
Mandatory PCI Register	0x0	Device ID		Vendor ID	
	0x4	Status Register		Control Register	
	0x8	Class Code (0xFF0000)			Revision ID
	0xC	Reserved	Header Type (0x0/0x80)	Latency Timer	Cache Line Size (0x0)
	0x10	Base Address Register 0			
	0x14	Base Address Register 1 (0x0)			
	0x18	Base Address Register 2 (0x0)			
	0x1C	Base Address Register 3 (0x0)			
	0x20	Base Address Register 4 (0x0)			
	0x24	Base Address Register 5 (0x0)			
	0x28	CardBus CIS pointer (0x0000)			
	0x2C	Subsystem Device ID		Subsystem Vendor ID	
	0x30	Expansion ROM Base Address (0x0)			
	0x34	Reserved			Cap Ptr (0x40)
	0x38	Reserved			
	0x3C	Max Latency (0x00)	Min Grant (0x00)	Interrupt Pin (0x00)	Interrupt Line (0x00)
PCI / PCIe Capabilities	0x40...0x47	Power Management Capability			
	0xA0...0xDB	PCIe Capability			
	0xE0...0xE7	VPD Capability			
Extended PCIe Configuration	0x100...0x12B	AER Capability			
	0x140...0x14B	Serial ID Capability			
	0x1B0...0x1B7	ACS Capability			

9.2.1.1 Sharing Among PCI Functions

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports multiple PCI functions. As each function exposes a PCIe configuration space, each register and each field is either shared among the functions or is replicated per each PCI function. In each section a table lists configuration sharing of the registers described in this section. Also, the description of each field describes special considerations regarding configuration sharing.

9.2.2 Mandatory PCI Configuration Registers

Table 9.3. Configuration Sharing of PCI Configuration Space


Field	Sub-field	Shared?	Replicated?	Comments
Vendor ID	Vendor ID	x		
Device ID	Device ID		x	
Command Register	I/O Access Enable		x	Issue UR per PF if disabled.

**Table 9.3. Configuration Sharing of PCI Configuration Space (Continued)**

Field	Sub-field	Shared?	Replicated?	Comments
	Memory Access Enable		x	Issue UR per PF if disabled.
	Bus Master Enable		x	
	Parity Error Response		x	Enables certain error reporting per PF.
	SERR# Enable		x	Controls error reporting per PF.
	Interrupt Disable		x	Selection of interrupt method per PF.
Status Register	Interrupt Status		x	
	Capabilities List	x		Hardwired to 1b.
	Data Parity Reported / Master Data Parity Error		x	Reports poisoned packets per PF.
	Signaled Target Abort		x	Reports completer abort per PF.
	Received Target Abort		x	Reports receiving a completer abort per PF.
	Received Master Abort		x	Reports receiving an UR per PF.
	Signaled System Error		x	Reports Fatal / non-fatal message per PF.
	Detected Parity Error		x	Reports receiving a poisoned TLP per PF.
Revision Register		x		
Class Code Register			x	Per function type.
Cache Line Size Register			x	Does not affect device behavior.
Latency Timer		x		Hardwired to 00h in PCIe.
Header Type Register		x		
BIST		x		Not supported (0x00).
BARs	Memory BAR		x	
	I/O BAR		x	
	MSI-X BAR		x	See MSI-X capability.
I/O BAR Mapping	IOADDR, IODATA		x	
Subsystem Vendor ID		x		
Subsystem ID		x		
Expansion ROM			x	
Cap_Ptr Register		x		
Interrupt Line Register			x	Just store the register value.
Interrupt Pin Register			x	Separate interrupt number (A-D) per PF.
Min. Grant		x		
Max Latency		x		

9.2.2.1 Vendor ID Register (0x0; RO)

This is a read-only register that has the same value for all PCI functions. It identifies unique Intel products. The value of this field is loaded from the PCI_VENDORID register loaded from NVM. The default value, if not loaded, is 0x8086.



9.2.2.2 Device ID Register (0x2; RO)

This is a read-only register that identifies individual the Intel® Xeon® Processor D-1500 Product Family LAN controller PCI functions. Both ports have the same default value equals as defined in the following table, and can be auto-loaded from the NVM during initialization with different values for each port as well as the dummy function (See Section 4.3 for dummy function relevance).

The device ID values available for different the Intel® Xeon® Processor D-1500 Product Family LAN controller SKUs are listed in the following table:

Note: At initialization this register is loaded from the NVM if GLPCI_CAPSUP_LOAD_DEV_ID bit is set. In which case the value of each PF is loaded from the PCI_PFDEVID field in the NVM.

Integrated PCS/ PHY	Internal Electrical Interface	External Media Type	External PHYs	Device ID
KR PHY	KR	Backplane	Supported switches: FM6000 and BRCM* Trident+	0x15AB
KR PHY	KR	10GBASE-T PHY	X557-AT2	0x15AD
KR PHY	KR	SFI-to-SFP+	Inphi* CS4227 (DP)	0x15AC
KR PHY	KX	1000BASE-T PHY	Marvell* 88E1512 Marvell 88E1514	0x15AE

9.2.2.3 Command Register (0x4; RW)

Shaded bits are not used by this implementation and are hard wired to 0b. Each function has its own Command register. Unless explicitly specified, functionality is the same in both functions.

Bit(s)	Initial Value	Description
0	0b	I/O Access Enable. If I/O BAR is not requested, this bit is hard-wired to zero.
1	0b	Memory Access Enable.
2	0b	Enable Mastering, also named Bus Master Enable (BME). <ul style="list-style-type: none"> LAN functions RW field. Dummy function RO as zero field.
3	0b	Special Cycle Monitoring – Hard-wired to 0b.
4	0b	MWI Enable – Hard-wired to 0b.
5	0b	Palette Snoop Enable – Hard-wired to 0b.
6	0b	Parity Error Response.
7	0b	Wait Cycle Enable – Hard-wired to 0b.
8	0b	SERR# Enable.
9	0b	Fast Back-to-Back Enable – Hard-wired to 0b.
10	0b	Interrupt Disable. When set, devices are prevented from generating legacy interrupt messages.
15:11	0b	Reserved.

9.2.2.4 Status Register (0x6; RO)

Shaded bits are not used by this implementation and are hardwired to 0b. Each function has its own Status register. Unless explicitly specified, functionality is the same in both functions.



Bits	Initial Value	RW	Description
2:0	0b		Reserved.
3	0b	RO	Interrupt Status. ¹
4	1b	RO	New Capabilities: Indicates that a device implements extended capabilities. The Intel® Xeon® Processor D-1500 Product Family LAN controller sets this bit and implements a capabilities list to indicate that it supports PCI power management, Message Signaled Interrupts (MSI), Enhanced Message Signaled Interrupts (MSI-X), VPD and the PCIe extensions.
5	0b		66 MHz Capable – Hardwire to 0b.
6	0b		Reserved.
7	0b		Fast Back-to-Back Capable – Hardwire to 0b.
8	0b	RW1C	Data Parity Reported.
10:9	00b		DEVSEL Timing – Hardwire to 0b.
11	0b	RW1C	Signaled Target Abort.
12	0b	RW1C	Received Target Abort.
13	0b	RW1C	Received Master Abort.
14	0b	RW1C	Signaled System Error.
15	0b	RW1C	Detected Parity Error.

1. The *Interrupt Status* field is a RO field that indicates that an interrupt message is pending internally to the Intel® Xeon® Processor D-1500 Product Family LAN controller.

9.2.2.5 Revision Register (0x8; RO)

The default revision ID of this device is based on a hard-strapped value.

The value of the rev ID is a logic XOR between the default value and the value in the NVM PCIe Device Revision ID (*PCI_REVID*).

Note that LAN 0 and LAN 1 functions have the same revision ID.

9.2.2.6 Class Code Register (0x9; RO)

The class code is a read-only value that identifies the device functionality according to the value of the *Storage Class* bit in the NVM *PCI_CLASS* NVM register.

- Class Code = 0x020000 (Ethernet Adapter) if NVM->*Storage Class* = 0b
- Class Code = 0x010000 (SCSI Storage device) if NVM->*Storage Class* = 1b

In the dummy function the class code equals to 0xFF0000.

9.2.2.7 Cache Line Size Register (0xC; RW)

This field is implemented by PCIe devices as a read/write field for legacy compatibility purposes but has no impact on any PCIe device functionality. The default value is zero.

9.2.2.8 Latency Timer (0xD; RO)

Not used. Hardwire to 0b.



9.2.2.9 Header Type Register (0xE; RO)

This indicates if a device is single- or multi-function. If a single LAN function is the only active one then this field has a value of 0x00 to indicate a single function device. If other functions are enabled then this field has a value of 0x80 to indicate a multi-function device. Table 9.4 lists the different options to set the header type field.

Table 9.4. Header Type Settings

Lan 0 Enable	Lan 1 Enable	Cross-Mode Enable	Header Type Expected Value
0	0	X	N/A (no function)
1	0	0	0x00
0	1	0	0x80 (dummy exist)
1	1	X	0x80 (dual function)
1	0	1	0x80 (dummy exist)
0	1	1	0x00

9.2.2.10 Subsystem Vendor ID Register (0x2C; RO)

This value can be loaded automatically from the NVM at power up or reset. A value of 0x8086 is the default for this field at power up if the NVM does not respond or is not programmed. All functions are initialized to the same value.

9.2.2.11 Subsystem ID Register (0x2E; RO)

This value can be loaded automatically from the NVM at power up with a default value of 0x0000.

PCI Function	Default Value	NVM Address
LAN Functions	0x0000	0x0B

9.2.2.12 Cap_Ptr Register (0x34; RO)


The *Capabilities Pointer* field (Cap_Ptr) is an 8-bit field that provides an offset in the Intel® Xeon® Processor D-1500 Product Family LAN controller's PCI configuration space for the location of the first item in the capabilities linked list. The Intel® Xeon® Processor D-1500 Product Family LAN controller sets this bit and implements a capabilities list to indicate that it supports PCI power management, MSIs, and PCIe extended capabilities. Its value is 0x40, which is the address of the first entry: PCI power management.

9.2.2.13 Interrupt Line Register (0x3C; RW)

Read/write register programmed by software to indicate which of the system interrupt request lines the Intel® Xeon® Processor D-1500 Product Family LAN controller's interrupt pin is bound to. Refer to the PCI definition for more details. Each PCI function has its own register.

9.2.2.14 Interrupt Pin Register (0x3D; RO)

Read-only register. LAN 0 / LAN 1 — A value of 0x1...0x4 indicates that this function implements a legacy interrupt on INTA#...INTD# respectively. Loaded from the *PCI_CNF* NVM word per function. For dummy function the returned value is 0x0 (Function uses no legacy interrupt Message).



Note: If only a single device/function of the Intel® Xeon® Processor D-1500 Product Family LAN controller component is enabled, this value is ignored and the *Interrupt Pin* field of the enabled device reports INTA# usage.

9.2.2.15 Max_Lat and Min_Gnt (0x3E;RO)

Not used. Hardwired to 0b.

For dummy functions this register is RO - zero.

9.2.2.16 Memory and IO Base Address Registers (0x10...0x27; RW)

Base Address Registers (BARs) are used to map the Intel® Xeon® Processor D-1500 Product Family LAN controller register space of the device functions. The Intel® Xeon® Processor D-1500 Product Family LAN controller has a memory BAR, I/O BAR and MSI-X BAR described in [Table 9.5](#). The BARs location and sizes are listed in the [Table 9.5](#) and [Table 9.6](#). The fields within each BAR are then listed in [Table 9.6](#).

Table 9.5. Intel® Xeon® Processor D-1500 Product Family LAN Controller Base Address Registers Description — LAN 0 / LAN 1

Mapping Windows	Mapping Description
Memory BAR	The internal registers memories and external Flash device are accessed as direct memory mapped offsets from the BAR. Software can access a Dword or 64 bits. The Flash space in this BAR is enabled by the <i>FLSize</i> and <i>CSRSIZE</i> fields in the <i>PCI_LBARCTRL</i> register. Address 0 in the Flash device is mapped to address 256 K in the memory BAR. When the usable Flash size + CSR space are smaller than the memory BAR, then accessing addresses above the top of the Flash wraps back to the beginning of the Flash.
I/O BAR	All internal registers and memories can be accessed using I/O operations. There are two 4-byte registers in the I/O mapping window: Addr Reg and Data Reg accessible as Dword entities. The I/O BAR is supported depending on the <i>IO_Sup</i> bit in the NVM at word PCIe Control 3 – Offset 0x07.
MSI-X BAR	The MSI-X vectors and Pending Bit Array (PBA) structures are accessed as direct memory mapped offsets from the MSI-X BAR. Software can access Dword entities.

Table 9.6. Base Address Registers' Fields

Field	bits	RW	Description	
Memory and I/O Space Indication	0	RO	0b = Indicates memory space. 1b = Indicates I/O.	
Memory Type	2:1	RO	00b = Reserved. 10b = 64-bit BAR	
Prefetch Memory	3	RO	0b = Non-prefetchable space. 1b = Prefetchable space. This bit should be set only in systems that do not generate prefetchable cycles. This bit is loaded from the <i>PREFBAR</i> bit in the NVM because it is required for 64-bit memory BARs.	
Address Space (low register for 64-bit memory BARs)	31:4	RW	The length of the RW bits and RO 0b bits depend on the mapping window size. Initial value of the RW fields is 0x0. The size of the memory BAR is described in Table 9.7 .	
			Mapping Window	RO bits
			Memory CSR + Flash BAR size depends on <i>PCI_LBARCTRL.FLSize</i> and <i>PCI_LBARCTRL.CSRSIZE</i> fields.	17:4 for 256 KB 18:4 for 512 KB and so on...
			MSI-X space is 16 KB.	13:4
			I/O space size is 32 bytes.	4:0



Table 9.7. Usable Flash Size and CSR Mapping Window Size

FLSize	CSRSize	Resulted CSR + Flash BAR Size	Installed Flash Device	Usable Flash Space
000b-100b	X	Reserved		
101b	0	2 MB	2 MB	2 MB minus 256 KB
101b	1	4 MB	2 MB	2 MB
110b	0	4 MB	4 MB	4 MB minus 256 KB
110b	1	8 MB	4 MB	4 MB
111b	0	8 MB	8 MB	8 MB minus 256 KB
111b	1	16 MB	8 MB	8 MB

9.2.2.17 Expansion ROM Base Address Register (0x30; RW)

This register is used to define the address and size information for boot-time access to the expansion ROM module in the Flash memory. It is enabled by the *PCI_LBARCTRL.EXROM_DIS* register field. This register returns a zero value for functions without an expansion ROM window and for dummy functions.

Field	Bit(s)	RW	Initial Value	Description
En	0	RW	0b	1b = Enables expansion ROM access. 0b = Disables expansion ROM access.
Reserved	10:1	R	0b	Always read as 0b. Writes are ignored.
Address	31:11	RW	0b	Read-write bits are hard wired to 0b and dependent on the memory mapping window size. The LAN Expansion ROM spaces can be either 64 KB or up to 512 KB in powers of 2. Mapping window size is set by <i>PCI_LBARCTRL.EXROM_BAR_SIZE</i> field

9.2.3 PCI Capabilities

The first entry of the PCI capabilities link list is pointed to by the Cap_Ptr register. Table 9.8 lists the capabilities supported by the Intel® Xeon® Processor D-1500 Product Family LAN controller.

Table 9.8. PCI Capabilities List for LAN Functions

Address	Item	Next Pointer
0x40:4F	PCI Power Management	0x50 / 0xA0 ¹
0x50:6F	MSI	0x70
0x70:8F	MSI-X	0xA0
0xA0:DF	PCIe Capabilities	0xE0 / 0x00
0xE0:0xEF	VPD Capability	0x00

1. In the dummy function, the power management capability points to the PCIe capabilities.

Table 9.9. PCI Capabilities for Dummy Function

Address	Item	Next Pointer
0x40:47	PCI Power Management	0x50
0xA0:DB	PCIe Capabilities	0x00



9.2.3.1 PCI Power Management Capability

All fields are reset at full power up. All fields except PME_En and PME_Status are reset after exiting from the D3cold state. If AUX power is not supplied, the PME_En and PME_Status fields also reset after exiting from the D3cold state. Refer to the detailed description for registers loaded from the NVM at initialization.

Byte Offset	Byte 3	Byte 2	Byte 1	Byte 0
0x40	Power Management Capabilities		Next Pointer (0x50 / 0xA0)	Capability ID (0x01)
0x44	Data	Bridge Support Extensions	Power Management Control & Status	

Table 9.10 lists sharing of the power management capability registers among the different PCI functions.

Table 9.10. Sharing the Power Management Capability Registers

Field	Sub-field	Shared?	Replicated?	Comments
Capability ID		x		
Next Pointer			x	
Power Management Capabilities	PME_Support	x		
	D2_Support	x		
	D1_Support	x		
	AUX Current	x		
	DSI	x		
	PME Clock	x		Hardwired to 0b.
	Version	x		
Power Management Control / Status	PME_Status		x	
	Data_Scale	x		
	Data_Select		x	
	PME_En		x	
	No_Soft_Reset	x		
	PowerState			x
Data Register			x	

9.2.3.1.1 Capability ID Register (0x40; RO)

This field equals 0x01 indicating the linked list item as being the PCI Power Management registers.

9.2.3.1.2 Next Pointer Register (0x41; RO)

This field provides an offset to the next capability item in the capability list. This field equals for both LAN ports to 0x50 pointing to the MSI capability. In dummy function, it equals to 0xA0 pointing to the PCIe capabilities.



9.2.3.1.3 Power Management Capabilities — PMC Register (0x42; RO)

This field describes the device functionality during the power management states as listed in the following table. Note that each device function has its own register.

Bits	Default	RW	Description
15:11	01001b	RO	PME_Support. This 5-bit field indicates the power states in which the function can assert PME#. Condition Functionality Values: <ul style="list-style-type: none"> No AUX Pwr - PME at D0 and D3hot = 01001b. AUX Pwr - PME at D0, D3hot, and D3cold = 11001b. Note: For dummy function, this field is RO - zero.
10	0b	RO	D2_Support – The Intel® Xeon® Processor D-1500 Product Family LAN controller does not support the D2 state.
9	0b	RO	D1_Support – The Intel® Xeon® Processor D-1500 Product Family LAN controller does not support the D1 state.
8:6	000b	RO	AUX Current – Required current defined in the Data register.
5	1b	RO	DSI – The Intel® Xeon® Processor D-1500 Product Family LAN controller requires its device driver to be executed following a transition to the D0 uninitialized state.
4	0b	RO	Reserved.
3	0b	RO	PME_Clock – Disabled. Hardwire to 0b.
2:0	011b	RO	Version – The Intel® Xeon® Processor D-1500 Product Family LAN controller complies with the PCI PM specification revision 1.2.

9.2.3.1.4 Power Management Control / Status Register — PMCSR (0x44; RW)

This register (listed in the following table) is used to control and monitor power management events in the Intel® Xeon® Processor D-1500 Product Family LAN controller. Note that each device function has its own PMCSR.

Bits	Default	RW	Description
15	0b at power up	RW1CS	PME_Status. This bit is set to 1b when the function detects a wake-up event independent of the state of the <i>PME_En</i> bit. Writing a 1b clears this bit.
14:13	01b	RO	Data_Scale. This field indicates the scaling factor that’s used when interpreting the value of the Data register. This field equals 01b (indicating 0.1 watt/units) if the Data_Select field is set to 0, 3, 4, 7, (or 8 for function 0 in multi-function device). Otherwise, it equals 00b.
12:9	0000b	RW	Data_Select. This 4-bit field is used to select which data is to be reported through the Data register and <i>Data_Scale</i> field. These bits are writable only when power management is enabled via the NVM.
8	0b at power up	RWS	PME_En. If power management is enabled in the NVM, writing a 1b to this register enables wake up.
7:4	0000b	RO	Reserved.
3	1b	RO	No_Soft_Reset. This bit is always set to 1b to indicate that the Intel® Xeon® Processor D-1500 Product Family LAN controller does not perform an internal reset upon transition from D3hot to D0 via software control of the <i>PowerState</i> bits. Configuration context is kept as part of the transition from the D3hot to the D0 state, thus a full re-initialization sequence of the configuration space is not needed to return the Intel® Xeon® Processor D-1500 Product Family LAN controller to the D0 Initialized state.
2	0b	RO	Reserved for PCIe.
1:0	00b	RW	PowerState. This field is used to set and report the power state of a function as follows: 00b = D0. 01b = D1 (cycle ignored if written with this value). 10b = D2 (cycle ignored if written with this value). 11b = D3.



9.2.3.1.5 PMCSR_BSE Bridge Support Extensions Register (0x46; RO)

This register is not implemented in the Intel® Xeon® Processor D-1500 Product Family LAN controller; values set to 0x00.

9.2.3.1.6 Data Register (0x47; RO)

This optional register is used to report power consumption and heat dissipation. The reported register is controlled by the *Data_Select* field in the PMCSR; the power scale is reported in the *Data_Scale* field in the PMCSR. The data for this field is loaded from the NVM via the *PCI_PWRDATA* register if power management is enabled in the NVM or with a default value of 0x00. The values for the Intel® Xeon® Processor D-1500 Product Family LAN controller's functions are as follows (the relevant column is selected based on the value of the *Data_Select* field):

Function	D0 (Consume/ Dissipate)	D3 (Consume/ Dissipate)	Common	Data_Scale
<i>Data_Select</i>	(0x0/0x4)	(0x3/0x7)	(0x8)	
0	PCI_PWRDATA.D0_POWER	PCI_PWRDATA.D3_POWER	Function zero of a multi-function device: PCI_PWRDATA.COMM_POWER Single-function device: 0x00	01b
1	PCI_PWRDATA.D0_POWER	PCI_PWRDATA.D3_POWE	0x00	01b

Note: For other *Data_Select* values the Data register output is reserved (0b).

9.2.3.2 MSI Capability

Note: This capability is not available for dummy functions.

This structure is required for PCIe devices.

Byte Offset	Byte 3	Byte 2	Byte 1	Byte 0
0x50	Message Control (0x0080)		Next Pointer (0x70)	Capability ID (0x05)
0x54	Message Address			
0x58	Message Upper Address			
0x5C	Reserved		Message Data	
0x60	Mask Bits			
0x64	Pending Bits			

Table 9.11 lists configuration sharing of the MSI Capability registers among the different PCI functions.

Table 9.11. Configuration sharing of the MSI Capability

Field	Sub-field	Shared?	Replicated?	Comments
Capability ID		x		
Next Pointer			x	
Message Control	MSI Enable		x	



Table 9.11. Configuration sharing of the MSI Capability (Continued)

Field	Sub-field	Shared?	Replicated?	Comments
	Multiple Messages Capable	x		
	Multiple Message Enable		x	
	64-bit Capable	x		
	MSI per-vector masking	x		
Message Address Low			x	
Message Address High			x	
Message Data			x	
Mask Bits			x	
Pending Bits			x	

9.2.3.2.1 Capability ID Register (0x50; RO)

This field equals 0x05 indicating that the linked list item as being the MSI registers.

9.2.3.2.2 Next Pointer Register (0x51; RO)

This field provides an offset to the next capability item in the capability list. Its value of 0x70 and points to MSI-X capability.

9.2.3.2.3 Message Control Register (0x52; RW)

These register fields are listed in the following table. Note that there is a dedicated register (per PCI function) to separately enable its MSI.

Bits	Default	RW	Description
0	0b	RW	MSI Enable. 1b = Message Signaled Interrupts. The Intel® Xeon® Processor D-1500 Product Family LAN controller generates an MSI for interrupt assertion instead of INTx signaling.
3:1	000b	RO	Multiple Messages Capable. The Intel® Xeon® Processor D-1500 Product Family LAN controller indicates a single requested message per function.
6:4	000b	RW	Multiple Message Enable. Since the Intel® Xeon® Processor D-1500 Product Family LAN controller requests a single vector in the <i>Multiple Message Capable</i> field, Software is expected to write 000b to this field.
7	1b	RO	64-bit Capable. A value of 1b indicates that the Intel® Xeon® Processor D-1500 Product Family LAN controller is capable of generating 64-bit message addresses.
8	1b ¹	RO	MSI per-vector masking. A value of 0b indicates that the Intel® Xeon® Processor D-1500 Product Family LAN controller is not capable of per-vector masking. A value of 1b indicates that the Intel® Xeon® Processor D-1500 Product Family LAN controller is capable of per-vector masking.
15:9	0b	RO	Reserved. Reads as 0b

1. The value is loaded from the MSI *Mask* bit in the NVM.

9.2.3.2.4 Message Address Low Register (0x54; RW)

Written by the system to indicate the lower 32 bits of the address to use for the MSI memory write transaction. The lower two bits always return 0b regardless of the write operation.

9.2.3.2.5 Message Address High Register (0x58; RW)



Written by the system to indicate the upper 32 bits of the address to use for the MSI memory write transaction.

9.2.3.2.6 Message Data Register (0x5C; RW)

Written by the system to indicate the lower 16 bits of the data written in the MSI memory write Dword transaction. The upper 16 bits of the transaction are written as 0b.

9.2.3.2.7 Mask Bits Register (0x60; RW)

The Mask Bits and Pending Bits registers enable software to disable or defer message sending on a per-vector basis. As the Intel® Xeon® Processor D-1500 Product Family LAN controller supports only one message, only bit 0 of these registers are implemented.

Bits	Default	RW	Description
0	0b	RW	MSI Vector 0 Mask. If set, the Intel® Xeon® Processor D-1500 Product Family LAN controller is prohibited from sending MSI messages.
31:1	0x0	RO	Reserved

9.2.3.2.8 Pending Bits Register (0x64; RW)

Bits	Default	RW	Description
0	0b	RO	If set, the Intel® Xeon® Processor D-1500 Product Family LAN controller has a pending MSI message.
31:1	0x0	RO	Reserved

9.2.3.3 MSI-X Capability

Note: This capability is not available for dummy functions.

More than one MSI-X capability structure per function is prohibited while a function is permitted to have both an MSI and an MSI-X capability structure.

In contrast to the MSI capability structure, which directly contains all of the control/status information for the function's vectors, the MSI-X capability structure instead points to an MSI-X table structure and an MSI-X Pending Bit Array (PBA) structure, each residing in memory space.

Each structure is mapped by a BAR belonging to the function that begins at 0x10 in the configuration space. A BAR Indicator Register (BIR) indicates which BAR and a Qword-aligned offset indicates where the structure begins relative to the base address associated with the BAR. The BAR is 64-bit, but must map to the memory space. A function is permitted to map both structures with the same BAR or map each structure with a different BAR.

The MSI-X table structure (Section 9.2.3.4) typically contains multiple entries, each consisting of several fields: *Message Address*, *Message Upper Address*, *Message Data*, and *Vector Control*. Each entry is capable of specifying a unique vector.

The PBA structure [MSI-X Table Offset Register (0x74; RW)] contains the function's pending bits, one per table entry, organized as a packed array of bits within Qwords. The last Qword is not necessarily fully populated.

To request service using a given MSI-X table entry, a function performs a Dword memory write transaction using:

- The contents of the *Message Data* field entry for data



- The contents of the *Message Upper Address* field for the upper 32 bits of the address
- The contents of the *Message Address* field entry for the lower 32 bits of the address

A memory read transaction from the address targeted by the MSI-X message produces undefined results.

The MSI-X table and MSI-X PBA are permitted to co-reside within a naturally aligned 4 KB address range, though they must not overlap with each other.

MSI-X table entries and *Pending* bits are each numbered 0 through N-1, where N-1 is indicated by the *Table Size* field in the MSI-X Message Control register. For a given arbitrary MSI-X table entry K, its starting address can be calculated with the formula:

$$\text{Entry starting address} = \text{Table base} + K * 16$$

For the associated *Pending* bit K, its address for Qword access and bit number within that Qword can be calculated with the formulas:

$$\text{Qword address} = \text{PBA base} + (K \text{ div } 64) * 8$$

$$\text{Qword bit\#} = K \text{ mod } 64$$

Software that chooses to read *Pending* bit K with Dword accesses can use these formulas:

$$\text{Dword address} = \text{PBA base} + (K \text{ div } 32) * 4$$

$$\text{Dword bit\#} = K \text{ mod } 32$$

Table 9.12. MSI-X Capability Structure

Byte Offset	Byte 3	Byte 2	Byte 1	Byte 0
0x70	Message Control (0x00090)		Next Pointer (0xA0)	Capability ID (0x11)
0x74	Table Offset			
0x78	PBA Offset			

Table 9.13 lists configuration sharing of the MSI-X Capability registers among the different PCI functions.

Table 9.13. Configuration sharing of the MSI-X Capability

Field	Sub-field	Shared?	Replicated?	Comments
Capability ID		x		
Next Pointer			x	
Message Control	Table Size	x		
Function Mask			x	
MSI-X Enable			x	
MSI-X Table Offset	Table BIR		x	
	Table Offset		x	
MSI-X Pending Bit Array	PBA BIR		x	

**Table 9.13. Configuration sharing of the MSI-X Capability**

Field	Sub-field	Shared?	Replicated?	Comments
	PBA Offset		x	
MSI-X Table			x	
MSI-X PBA Structure			x	

9.2.3.3.1 Capability ID Register (0x70; RO)

This field equals 0x11 indicating that the linked list item as being the MSI-X registers.

9.2.3.3.2 Next Pointer Register (0x71; RO)

This field provides an offset to the next capability item in the capability list. Its value of 0xA0 points to PCIe capability.

9.2.3.3.3 Message Control Register (0x72; RW)

These register fields are listed in the following table. Note that there is a dedicated register (per PCI function).

Bits	Default	RW	Description
10:0	0x3F	RO	Table Size. System software reads this field to determine the MSI-X Table Size N, which is encoded as N-1. The Intel® Xeon® Processor D-1500 Product Family LAN controller supports up to 64 different interrupt vectors per function. This field is loaded from the PCI_CNF2.MSI_X_PF_N register field.
13:11	0b	RO	Always returns 0b on a read. A write operation has no effect.
14	0b	RW	Function Mask. If 1b, all of the vectors associated with the function are masked, regardless of their per-vector <i>Mask</i> bit states. If 0b, each vector's <i>Mask</i> bit determines whether the vector is masked or not. Setting or clearing the MSI-X <i>Function Mask</i> bit has no effect on the state of the per-vector <i>Mask</i> bits.
15	0b	RW	MSI-X Enable. If 1b and the MSI <i>Enable</i> bit in the MSI Message Control register is 0b, the function is permitted to use MSI-X to request service and is prohibited from using its INTx# pin. System configuration software sets this bit to enable MSI-X. A device driver is prohibited from writing this bit to mask a function's service request. If 0b, the function is prohibited from using MSI-X to request service.

9.2.3.3.4 MSI-X Table Offset Register (0x74; RW)

These register fields are listed in the following table.

Bits	Default	RW	Description
2:0	0x3/0x4	RO	Table BIR. Indicates which one of a function's BARs, beginning at 0x10 in the configuration space, is used to map the function's MSI-X table into the memory space. while BIR values: 0...5 correspond to BARs 0x10...0x 24, respectively.
31:3	0x000	RO	Table Offset. Used as an offset from the address contained in one of the function's BARs to point to the base of the MSI-X table. The lower three <i>Table BIR</i> bits are masked off (set to 0b) by software to form a 32-bit Qword-aligned offset. Note that this field is read only.

9.2.3.3.5 MSI-X Pending Bit Array — PBA Offset (0x78; RW)



This register fields are listed in the following table.

Bits	Default	RW	Description
2:0	0x4	RO	PBA BIR. Indicates which one of a function's BARs, beginning at 0x10 in the configuration space, is used to map the function's MSI-X PBA into the memory space. while BIR values: 0...5 correspond to BARs 0x10...0x 24, respectively.
31:3	0x0400	RO	PBA Offset. Used as an offset from the address contained in one of the functions BARs to point to the base of the MSI-X PBA. The lower three PBA BIR bits are masked off (set to 0b) by software to form a 32-bit Qword-aligned offset. This field is read only.

9.2.3.4 MSI-X Table Structure

Dword3 — MSIXVCTRL	Dword2 — MSIXTMSG	Dword1 — MSIXTUADD	Dword0 — MSIXTADD	Entry Number	BAR 3 — Offset
Vector Control	Msg Data	Msg Upper Addr	Msg Lower Addr	0	Base (0x0000)
Vector Control	Msg Data	Msg Upper Addr	Msg Lower Addr	1	Base + 1*16
Vector Control	Msg Data	Msg Upper Addr	Msg Lower Addr	2	Base + 2*16
...	
Vector Control	Msg Data	Msg Upper Addr	Msg Lower Addr	63	Base + 63*16
Vector Control	Msg Data	Msg Upper Addr	Msg Lower Addr	64	Base + 64*16
...	
Vector Control	Msg Data	Msg Upper Addr	Msg Lower Addr	255	Base + 255*16

Note: All MSI-X vectors > MSI-X 63, are usable only by the Virtual Functions (VFs) in IOV mode. These vectors are not exposed to the operating system by the *Table Size* field in the MSI-X Message Control word.

See [Section 8.2.2.5](#) for details of the MSI-X registers in BAR 3 of the PF.

9.2.3.5 VPD Registers

Note: This capability is not available for dummy functions.

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports access to a VPD structure stored in the NVM using the following set of registers.

Initial values of the configuration registers are marked in parenthesis.

Note: The VPD structure is available through both ports functions. As the interface is common to the two functions, accessing the VPD structure of one function while an access to the NVM is in process on the other function can yield to unexpected results.

Byte Offset	Byte 3	Byte 2	Byte 1	Byte 0
0xE0	VPD Address		Next Pointer (0x00)	Capability ID (0x03)
0xE4	VPD Data			



Table 9.14 lists configuration sharing of the VPD Capability registers among the different PCI functions.

Table 9.14. Configuration sharing of the VPD Capability

Field	Shared?	Replicated?	Comments
Capability ID	x		
Next Pointer	x		
VPD Address/F		x	
VPD Data		x	

9.2.3.5.1 Capability ID Register (0xE0; RO)

This field equals 0x3 indicating the linked list item as being the VPD registers.

9.2.3.5.2 Next Pointer Register (0xE1; RO)

Offset to the next capability item in the capability list. A 0x00 value indicates that it is the last item in the capability-linked list.

9.2.3.5.3 VPD Address Register (0xE2; RW)

Word-aligned byte address of the VPD area in the NVM to be accessed. The register is read/write, and the initial value at power-up is indeterminate.

Bits	Default	Rd/Wr	Description
14:0	X	RW	Address: Dword-aligned byte address of the VPD area in the NVM to be accessed. The register is read/write, and the initial value at power-up is indeterminate. The two LSBs are RO as zero.
15	0b	RW	F: A flag used to indicate when the transfer of data between the VPD Data register and the storage component completes. The Flag register is written when the VPD Address register is written. 0b = Read. Set by hardware when data is valid. 1b = Write. Cleared by hardware when data is written to the NVM. The VPD address and data should not be modified before the action is done.

9.2.3.5.4 VPD Data Register (0xE4; RW)

VPD read/write data.

Bits	Default	Rd/Wr	Description
31:0	X	RW	VPD Data: VPD data can be read or written through this register. The LS byte of this register (at offset 4 in this capability structure) corresponds to the byte of VPD at the address specified by the VPD Address register. The data read from or written to this register uses the normal PCI byte transfer capabilities. Four bytes are always transferred between this register and the VPD storage component. Reading or writing data outside of the VPD space in the storage component is not allowed. In a write access, the data should be set before the address and the flag is set.



9.2.3.6 PCIe Capability

The Intel® Xeon® Processor D-1500 Product Family LAN controller implements the PCIe capability structure linked to the legacy PCI capability list for endpoint devices as follows:

Byte Offset	Byte 3	Byte 2	Byte 1	Byte 0
0xA0	PCI Express Capability Register (0x0002)		Next Pointer (0xE0/0x00)	Capability ID (0x10)
0xA4	Device Capability			
0xA8	Device Status		Device Control	
0xAC	Link Capability			
0xB0	Link Status		Link Control	
0xB4	Reserved			
0xB8	Reserved		Reserved	
0xBC	Reserved			
0xC0	Reserved		Reserved	
0xC4	Device Capability 2			
0xC8	Reserved		Device Control 2	
0xCC	Reserved			
0xD0	Link Status 2		Link Control 2	
0xD4	Reserved			
0xD8	Reserved		Reserved	

Table 9.14 lists configuration sharing of the PCIe Capability registers among the different PCI functions.

Table 9.15. Configuration Sharing of the PCIe Capability

Field	Sub-field	Shared?	Replicated?	Comments
Capability ID		x		
Next Pointer			x	
PCIe Capabilities		x		
Device Capabilities	Max Payload Size Supported	x		
	Phantom Functions Supported	x		Not supported.
	Extended Tag Field Supported	x		
	Endpoint L0s Acceptable Latency	x		
	Endpoint L1 Acceptable Latency	x		
	Function Level Reset Capability	x		
Device Control	Correctable Error Reporting Enable		x	
	Non-Fatal Error Reporting Enable		x	



Table 9.15. Configuration Sharing of the PCIe Capability (Continued)

Field	Sub-field	Shared?	Replicated?	Comments
	Fatal Error Reporting Enable		x	
	Unsupported Request Reporting Enable		x	
	Enable Relaxed Ordering		x	
	Max Payload Size		x	Use minimum of all configured values. In ARI mode, use value in function 0.
	Extended Tag field Enable		x	
	Auxiliary Power PM Enable		x	Same policy for all PFs (Logical OR of the PFs' bits).
	Enable No Snoop		x	
	Max Read Request Size		x	Use minimum of all configured values.
	Initiate Function Level Reset		x	
Device Status	Correctable Detected		x	
	Non-Fatal Error Detected		x	
	Fatal Error Detected		x	
	Unsupported Request Detected	x		
	Aux Power Detected	x		
	Transactions Pending		x	
Link Capabilities	Supported Link Speeds	x		
	Max Link Width	x		
	Active State Link PM Support	x		
	L0s Exit Latency	x		
	L1 Exit Latency	x		
	Clock Power Management	x		
	Port Number	x		
Link Control	Active State Link PM Control		x	Same policy for all PFs (Logical AND of the PFs' bits). In ARI mode, use value in function 0.
	Read Completion Boundary (RCB)		x	
	Common Clock Configuration		x	Same policy for all PFs (Logical AND of the PFs' bits). In ARI mode, use value in function 0.
	Extended Sync		x	Same policy for all PFs (Logical OR of the PFs' bits).
Link Status	Current Link Speed	x		
	Negotiated Link Width	x		
	Slot Clock Configuration	x		
Device Capabilities 2	Completion Timeout Ranges Supported	x		



Table 9.15. Configuration Sharing of the PCIe Capability (Continued)

Field	Sub-field	Shared?	Replicated?	Comments
	Completion Timeout Disable Supported	x		
	LTR Mechanism Supported	x		
	TPH Completer Supported	x		
	Extended Fmt Field Supported	x		
	OBFF Supported	x		
Device Control 2	Completion Timeout Value		x	Completion timeout decision per PF or use the largest configured value among PFs.
	Completion Timeout Disable		x	Completion timeout mechanism enabled per PF.
	IDO Request Enable		x	
	IDO Completion Enable		x	
	LTR Mechanism Enable		x	PF0 only. RsvdP on other functions.
	OBFF Enable		x	PF0 only. RsvdP on other functions.
Link Capabilities 2		x		
Link Control 2		x		PF0 only. RsvdP on other functions.
Link Status 2		x		

9.2.3.6.1 Capability ID Register (0xA0; RO)

This field equals 0x10 indicating that the linked list item as being the PCIe Capabilities registers.

9.2.3.6.2 Next Pointer Register (0xA1; RO)

Offset to the next capability item in the capability list. Its value of 0xE0 points to the VPD structure. If VPD is disabled or for a dummy function, a value of 0x00 value indicates that it is the last item in the capability-linked list.

9.2.3.6.3 PCIe Capabilities Register (0xA2; RO)

The PCIe Capabilities register identifies PCIe device type and associated capabilities. This is a read-only register identical to all functions.

Bits	Default	RW	Description
3:0	0010b	RO	Capability Version. Indicates the PCIe capability structure version. The Intel® Xeon® Processor D-1500 Product Family LAN controller supports PCIe version 2 (also loaded from the <i>PCI_CAPSUP.PCIE_VER</i> bit in the NVM).
7:4	0000b	RO	Device/Port Type. Indicates the type of PCIe functions. This field’s default value is based on a hard-strap. 0 - 0000b: Integrated PCIe endpoint 1 - 1001b: Integrated RC endpoint Note: Should be always 0000b.
8	0b	RO	Slot Implemented. The Intel® Xeon® Processor D-1500 Product Family LAN controller does not implement slot options. Therefore, this field is hard wired to 0b.



Bits	Default	RW	Description
13:9	00000b	RO	Interrupt Message Number. The Intel® Xeon® Processor D-1500 Product Family LAN controller does not implement multiple MSI per function. As a result, this field is hard wired to 0x0.
15:14	00b	RO	Reserved.

9.2.3.6.4 Device Capabilities Register (0xA4; RO)

This register identifies the PCIe device specific capabilities. It is a read-only register with the same value for the two LAN functions and for all other functions.

Bits	Rd/Wr	Default	Description
2:0	RO	010b/111b	Max Payload Size Supported. This field indicates the maximum payload that the Intel® Xeon® Processor D-1500 Product Family LAN controller can support for TLPs. This field's default value is affected by the IMPS hard-strap. If IMPS is set to 64 bytes (111b), then the Max payload size reported is 0 (128 bytes). Else, this field is loaded from the NVM with a value of 512 bytes.
4:3	RO	00b	Phantom Function Supported. Not supported by the Intel® Xeon® Processor D-1500 Product Family LAN controller.
5	RO	0b	Extended Tag Field Supported. Maximum supported size of the <i>Tag</i> field. The Intel® Xeon® Processor D-1500 Product Family LAN controller supports a 5-bit <i>Tag</i> field for all functions.
8:6	RO	011b	Endpoint L0s Acceptable Latency. This field indicates the acceptable latency that the Intel® Xeon® Processor D-1500 Product Family LAN controller can withstand due to the transition from L0s state to the L0 state. All functions share the same value loaded from the NVM PCIe Init Configuration 1 bits [8:6]. A value of 011b equals 512 ns.
11:9	RO	110b	Endpoint L1 Acceptable Latency. This field indicates the acceptable latency that the Intel® Xeon® Processor D-1500 Product Family LAN controller can withstand due to the transition from L1 state to the L0 state. A value of 110b equals 32 μ s-64 μ s. All functions share the same value loaded from the NVM.
12	RO	0b	Attention Button Present. Hard wired in the Intel® Xeon® Processor D-1500 Product Family LAN controller to 0b for all functions.
13	RO	0b	Attention Indicator Present. Hard wired in the Intel® Xeon® Processor D-1500 Product Family LAN controller to 0b for all functions.
14	RO	0b	Power Indicator Present. Hard wired in the Intel® Xeon® Processor D-1500 Product Family LAN controller to 0b for all functions.
15	RO	1b	Role Based Error Reporting. Hard wired in the Intel® Xeon® Processor D-1500 Product Family LAN controller to 1b for all functions.
17:16	RO	000b	Reserved 0b.
25:18	RO	0x00	Slot Power Limit Value. Used in upstream ports only. Hardwired in the Intel® Xeon® Processor D-1500 Product Family LAN controller to 0x00 for all functions.
27:26	RO	00b	Slot Power Limit Scale. Used in upstream ports only. Hardwired in the Intel® Xeon® Processor D-1500 Product Family LAN controller to 0b for all functions.
28	RO	1b	Function Level Reset Capability – A value of 1b indicates the Function supports the optional Function Level Reset (FLR) mechanism.
31:29	RO	0000b	Reserved.

9.2.3.6.5 Device Control Register (0xA8; RW)

This register controls the PCIe specific parameters. Note that there is a dedicated register per each function.



Bits	RW	Default	Description
0	RW	0b	Correctable Error Reporting Enable. Enable error report.
1	RW	0b	Non-Fatal Error Reporting Enable. Enable error report.
2	RW	0b	Fatal Error Reporting Enable. Enable error report.
3	RW	0b	Unsupported Request Reporting Enable. Enable error report.
4	RW	1b	Enable Relaxed Ordering. If this bit is set, the Intel® Xeon® Processor D-1500 Product Family LAN controller is permitted to set the <i>Relaxed Ordering</i> bit in the <i>Attribute</i> field of write transactions that do not need strong ordering. Refer to the CTRL_EXT register bit RO_DIS for more details.
7:5	RW	000b (128 bytes)	Max Payload Size. This field sets the maximum TLP payload size for the Intel® Xeon® Processor D-1500 Product Family LAN controller functions. As a receiver, the Intel® Xeon® Processor D-1500 Product Family LAN controller must handle TLPs as large as the set value. As a transmitter, the Intel® Xeon® Processor D-1500 Product Family LAN controller must not generate TLPs exceeding the set value. The <i>Max Payload Size</i> field supported in the Device Capabilities register indicates permissible values that can be programmed. In ARI mode, <i>Max Payload Size</i> is determined solely by the field in function 0 (even when it is a dummy function) while it is meaningless in the other function(s). This field's value might be overridden according to the IMPS hard-strap. If the IMPS is set to 64 bytes (111b), then this field is ignored and the IMPS value of 64 bytes is used.
8	RW	0b	Extended Tag field Enable. Not implemented in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
9	RW	0b	Phantom Functions Enable. Not implemented in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
10	RWS	0b	Auxiliary Power PM Enable. When set, enables the Intel® Xeon® Processor D-1500 Product Family LAN controller to draw AUX power independent of PME AUX power. The Intel® Xeon® Processor D-1500 Product Family LAN controller is a multi-function device, therefore allowed to draw AUX power if at least one of the functions has this bit set.
11	RW	0b	Enable No Snoop.No-snoop is not used by the Intel® Xeon® Processor D-1500 Product Family LAN controller.
14:12	RW	010b	Max Read Request Size. This field sets maximum read request size for the Intel® Xeon® Processor D-1500 Product Family LAN controller as a requester. 000b = 128 bytes. 001b = 256 bytes. 010b = 512 bytes. 011b = 1024 bytes. 100b = 2048 bytes. 101b = 4096 bytes. 110b = Reserved. 111b = Reserved. This field's value might be overridden according to the IMRS hard-strap. If IMRS is configured to 64 bytes (111b) then this field is ignored and the Intel® Xeon® Processor D-1500 Product Family LAN controller does not issue a read request as a master for which the address plus length crosses a 64-byte naturally aligned boundary.
15	RW	0b	Initiate FLR – A write of 1b initiates FLR to the function. The value read by software from this bit is always 0b.

9.2.3.6.6 Device Status Register (0xAA; RW1C)

This register provides information about PCIe device specific parameters. Note that there is a dedicated register per each function.




Bits	RW	Default	Description
0	RW1C	0b	Correctable Detected. Indicates status of correctable error detection.
1	RW1C	0b	Non-fatal Error Detected. Indicates status of non-fatal error detection.
2	RW1C	0b	Fatal Error Detected. Indicates status of fatal error detection.
3	RW1C	0b	Unsupported Request Detected. Indicates that the Intel® Xeon® Processor D-1500 Product Family LAN controller received an unsupported request. This field is identical in all functions. The Intel® Xeon® Processor D-1500 Product Family LAN controller can't distinguish which function causes the error.
4	RO	0b	Aux Power Detected. If Aux Power is detected, this field is set to 1b. It is a strapping signal from the periphery and is identical for all functions. Resets on LAN Power Good and PE_RST_N only.
5	RO	0b	Transaction Pending. Indicates whether the Intel® Xeon® Processor D-1500 Product Family LAN controller has ANY transactions pending. (transactions include completions for any outstanding non-posted request for all used traffic classes).
15:6	RO	0x00	Reserved.

9.2.3.6.7 Link Capabilities Register (0xAC; RO)

This register identifies PCIe link-specific capabilities. This is a read-only register identical to all functions.

Note: A root complex integrated endpoint MUST not implement this register and return a value of zero when accessed.

Bits	RW	Discrete Default	Integrated EP Default	Description
3:0	RO	0001b	0001b	Supported Max Link Speed. This field indicates the supported Link speed(s) of the associated link port. Defined encodings are: 0001b = 2.5 GT/s link speed supported. 0010b = 5 GT/s and 2.5 GT/s link speeds supported. 0011b = 8 GT/s and 5 GT/s and 2.5 GT/s link speeds supported.
9:4	RO	0x04 - 17x17 0x08 - 25x25	0x1	Max Link Width. Indicates the maximum link width. The Intel® Xeon® Processor D-1500 Product Family LAN controller supports a x1, x2, x4 and x8-link width. This field is loaded from the PCIe Analog Configuration NVM module by interpreting the masked lanes, with a default value of eight lanes for the 25x25 package and four for the 17x17 package. Defined encoding: 000000b = Reserved. 000001b = x1. 000010b = x2. 000100b = x4. 001000b = x8.
11:10	RO	11b	11b	Active State Link PM Support. Indicates the level of the active state of power management supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller. Defined encodings are: 00b = No ASPM Support. 01b = L0s Entry Supported. 10b = L1 Supported. 11b = L0s and L1 Supported. All functions share the same value loaded from the NVM Act_Stat_PM_Sup field in the NVM PCIe Init Configuration 3 offset 0x3.




Bits	RW	Discrete Default	Integrated EP Default	Description
14:12	RO	101b (1 μ s – 2 μ s)	000b	L0s Exit Latency. Indicates the exit latency from L0s to L0 state. 000b = Less than 64 ns. 001b = 64 ns – 128 ns. 010b – 128ns – 256 ns. 011b – 256 ns – 512 ns. 100b = 512 ns – 1 μ s. 101b = 1 μ s – 2 μ s. 110b = 2 μ s – 4 μ s. 111b = Reserved. All functions share the same value loaded from the NVM.
17:15	RO	100b (8 μ s – 16 μ s)	000b	L1 Exit Latency. Indicates the exit latency from L1 to L0 state. 000b = Less than 1 μ s. 001b = 1 μ s – 2 μ s. 010b = 2 μ s – 4 μ s. 011b = 4 μ s – 8 μ s. 100b = 8 μ s – 16 μ s. 101b = 16 μ s – 32 μ s. 110b = 32 μ s – 64 μ s. 111b = L1 transition not supported. All functions share the same value loaded from the NVM.
18	RO	0	0	Clock Power Management.
19	RO	0	0	Surprise Down Error Reporting Capable. Hard wired to 0b.
20	RO	0	0	Data Link Layer Link Active Reporting Capable.
21	RO	0	0	Link Bandwidth Notification Capability. Hard wired to 0b.
22	RO	1b	0	ASPM Optional Compliance. This bit must be set to 1b. Components that were implemented according to an earlier PCIe specification version has this bit set to 0b. Software is permitted to use the value of this bit to help determine whether to enable ASPM or whether to run ASPM compliance tests.
23	RO	0b	0	Reserved.
31:24	HwInit	0x0	Strap	Port Number. The PCIe port number for the given PCIe link. This field is set in the link training phase. For an integrated EP this value is based on a strapping wire.

9.2.3.6.8 Link Control Register (0xB0; RO)

This register controls PCIe link specific parameters. There is a dedicated register per each function.

Note: A root complex integrated endpoint MUST not implement this register and return a value of zero when accessed.



Bits	RW	Discrete Default	Integrated EP Default	Description
1:0	RW	00b	xxb	Active State Link PM Control. This field controls the active state PM supported on the link. Link PM functionality is determined by the lowest common denominator of all functions. For non-ARI mode, only capabilities enabled in all functions are enabled for the component as a whole. When ARI support is exposed, ASPM control is determined solely by the setting in Function 0 (even when it is a dummy function), regardless of Function 0's D-state. The settings in the other functions always return whatever value software programmed for each, but otherwise are ignored by the Intel® Xeon® Processor D-1500 Product Family LAN controller. Defined encodings are: 00b = PM disabled. 01b = L0s Entry supported. 10b = L1 entry enabled. 11b = L0s and L1 supported. In ARI mode, the ASPM is determined solely by the field in function 0 while it is meaningless in the other function(s). For an integrated EP: R/W scratch pad (no effect on agent).
2	RO	0b	0b	Reserved.
3	RO	0b	0b	Read Completion Boundary. For an integrated EP: R/W.
4	RO	0b	0b	Link Disable. Reserved for endpoint devices. Hardwired to 0b.
5	RO	0b	0b	Retrain Clock. Not applicable for endpoint devices. Hardwire to 0b. For an integrated EP: Software writes this bit, but it must always return 0 on reads (no effect on agent).
6	RW	0b	xb	Common Clock Configuration. When set, indicates that the Intel® Xeon® Processor D-1500 Product Family LAN controller and the component at the other end of the link are operating with a common reference clock. A value of 0b indicates that they are operating with an asynchronous clock. In ARI mode, the common clock configuration is determined solely by the field in function 0 (even when it is a dummy function) while it is meaningless in the other function(s). For an integrated EP: R/W scratch pad (no effect on agent).
7	RW	0b	xb	Extended Sync. When set, this bit forces an extended Tx of the FTS ordered set in FTS and an extra TS1 at the exit from L0s prior to entering L0. For an integrated EP: R/W scratch pad (no effect on agent).
8	RO	0b	0b	Enable Clock Power Management. Not supported - hardwired to 0b. For an integrated EP: Same.
9	RW	0b		Reserved. Returns the value that was written.
10	RO	0b	0b	Link Bandwidth Management Interrupt Enable. Not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller. Hardwired to 0. For an integrated EP: Hardwire to 0.
11	RO	0b	0b	Link Autonomous Bandwidth Interrupt Enable. Not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller. Hardwired to 0b. For an integrated EP: Hardwire to 0b.
15:12	RO	0x0	0x0	Reserved

9.2.3.6.9 Link Status Register (0xB2; RO)

This register provides information about PCIe Link specific parameters. This is a read only register identical to all functions.




Bits	RW	Discrete Default	Integrated EP Default	Description
3:0	RO	X	001b	Current Link Speed. This field indicates the negotiated link speed of the given PCIe link. Defined encodings are: 0001b = 2.5 GT/s PCIe link. 0010b = 5 GT/s PCIe link. 0011b = 8 GT/s PCIe link. All other encodings are reserved. The value in this field is undefined when the Link is not up. For an integrated EP: Hardwire to 001b.
9:4	RO	X	001b	Negotiated Link Width. Indicates the negotiated width of the link. Relevant encodings for the Intel® Xeon® Processor D-1500 Product Family LAN controller are: 000001b = x1. 000010b = X2. 000100b = x4. 001000b = x8. The value in this field is undefined when the Link is not up. For an integrated EP: Hardwire to 001b.
10	RO	0b	0b	Undefined.
11	RO	0b	0b	Link Training. Indicates that link training is in progress. This field is not applicable and is reserved for endpoint devices, and is hardwired to 0b. For an integrated EP: Hardwire to 0b.
12	HwInit	1b	1b	Slot Clock Configuration. When set, indicates that the Intel® Xeon® Processor D-1500 Product Family LAN controller uses the physical reference clock that the platform provides at the connector. This bit must be cleared if the Intel® Xeon® Processor D-1500 Product Family LAN controller uses an independent clock. The <i>Slot Clock Configuration</i> bit is loaded from the <i>Slot_Clock_Cfg</i> NVM bit. For an integrated EP: Hardwire to 1b.
13	RO	0b	0b	Data Link Layer Link Active. Not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller. Hardwire to 0b. For an integrated EP: Hardwire to 0b.
14	RO	0b	0b	Link Bandwidth Management Status. Not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller. Hardwire to 0b. For an integrated EP: Hardwire to 0b.
15	RO	0b	0b	Link Autonomous Bandwidth Status. This bit is not applicable and is reserved for endpoints. For an integrated EP: Hardwire to 0b.

The following registers are supported only if the capability version is two and above.


9.2.3.6.10 Device Capability 2 Register (0xC4; RO)

This register identifies the PCIe device-specific capabilities. It is a read-only register with the same value for both LAN functions.



Bits	RW	Default	Description
3:0	RO	1111b	<p>Completion Timeout Ranges Supported. This field indicates the Intel® Xeon® Processor D-1500 Product Family LAN controller's support for the optional completion timeout programmability mechanism.</p> <p>Four time value ranges are defined:</p> <ul style="list-style-type: none"> Range A: 50 μs to 10 ms. Range B: 10 ms to 250 ms. Range C: 250 ms to 4 s. Range D: 4 s to 64 s. <p>Bits are set according to the following values to show the timeout value ranges that the Intel® Xeon® Processor D-1500 Product Family LAN controller supports.</p> <ul style="list-style-type: none"> 0000b = Completion timeout programming not supported. The Intel® Xeon® Processor D-1500 Product Family LAN controller must implement a timeout value in the range of 50 μs to 50 ms. 0001b = Range A. 0010b = Range B. 0011b = Ranges A and B. 0110b = Ranges B and C. 0111b = Ranges A, B and C. 1110b = Ranges B, C and D. 1111b = Ranges A, B, C and D. All other values are reserved.
4	RO	1b	<p>Completion Timeout Disable Supported</p> <p>A value of 1b indicates support for the completion timeout disable mechanism.</p> <p>Note: For dummy functionality, a completion timeout is not relevant as a dummy function because it never sends non-posted requests.</p>
5	RO	0b	<p>ARI Forwarding Supported. Applicable only to switch downstream. Ports and Root Ports; must be 0b for other function types.</p>
10:6	RO	0x00	Not supported - hardwired to 0x00.
11	RO	1b	<p>LTR Mechanism Supported – A value of 1b indicates support for the optional Latency Tolerance Reporting (LTR) mechanism.</p> <p>For a multi-function device associated with an upstream port, each function must report the same value for this bit.</p> <p>Note: Value loaded from LTR_EN bit in the NVM.</p>
13:12	RO	00b	<p>TPH Completer Supported - Value indicates Completer support for TPH or Extended TPH.</p> <p>This capability is not supported.</p>
17:14	RO	0x0	Reserved.
19:18	RO	00b	<p>OBFF Supported.</p> <p>00b = OBFF Not Supported.</p> <p>01b = OBFF supported using message signaling only.</p> <p>10b = OBFF supported using WAKE# signaling only.</p> <p>11b = OBFF supported using WAKE# and Message signaling.</p>
20	RO	0b	<p>Extended Fmt Field Supported - If Set, the function supports the 3-bit definition of the Fmt field. If Clear, the function supports a 2-bit definition of the Fmt field.</p> <p>Not supported by this device.</p>
21	RO	0b	<p>End-End TLP Prefix Supported – Indicates whether End-End TLP Prefix support is offered by a Function.</p> <p>Not supported by this device.</p>
23:22	RsvdP	0b	<p>Max End-End TLP Prefixes – Indicates the maximum number of End-End TLP Prefixes supported by this function.</p> <p>Reserved for this device.</p>
31:24	RO	0x00	Reserved.

9.2.3.6.11 Device Control 2 Register (0xC8; RW)



This register controls the PCIe specific parameters. Note that there is a dedicated register per each function.

Bits	RW	Default	Description
3:0	RW	0x0	<p>Completion Timeout Value. For devices that support completion timeout programmability, this field enables system software to modify the completion timeout value.</p> <p>Defined encodings:</p> <ul style="list-style-type: none"> 0000b = Default range: 16 ms to 32 ms. <p>Note: It is strongly recommended that the completion timeout mechanism not expire in less than 10 ms.</p> <p>Values available if Range A (50 μs to 10 ms) programmability range is supported:</p> <ul style="list-style-type: none"> 0001b = 50 μs to 100 μs. 0010b = 1 ms to 2 ms. <p>Values available if Range B (10 ms to 250 ms) programmability range is supported:</p> <ul style="list-style-type: none"> 0101b = 16 ms to 32 ms. 0110b = 65 ms to 130 ms. <p>Values available if Range C (250 ms to 4 s) programmability range is supported:</p> <ul style="list-style-type: none"> 1001b = 260 ms to 520 ms. 1010b = 1 s to 2 s. <p>Values available if the Range D (4 s to 64 s) programmability range is supported:</p> <ul style="list-style-type: none"> 1101b = 4 s to 8 s. 1110b = 17 s to 34 s. <p>Values not defined are reserved.</p> <p>Software is permitted to change the value of this field at any time. For requests already pending when the completion timeout value is changed, hardware is permitted to use either the new or the old value for the outstanding requests and is permitted to base the start time for each request either on when this value was changed or on when each request was issued.</p> <p>Note: For dummy function, this field is RO - zero.</p>
4	RW	0b	<p>Completion Timeout Disable. When set to 1b, this bit disables the completion timeout mechanism.</p> <p>Software is permitted to set or clear this bit at any time. When set, the completion timeout detection mechanism is disabled. If there are outstanding requests when the bit is cleared, it is permitted but not required for hardware to apply the completion timeout mechanism to the outstanding requests. If this is done, it is permitted to base the start time for each request on either the time this bit was cleared or the time each request was issued.</p> <p>Note: For dummy function, this field is RO - zero.</p>
5	RO	0b	ARI Forwarding Enable. Applicable only to switch devices.
7:6	RO	00b	Not supported - hardwired to 00b.
8	RW	0b	<p>IDO Request Enable – If this bit is set, the function is permitted to set the ID-Based Ordering (<i>IDO</i>) bit (Attribute[2]) of requests it initiates.</p> <p>Default value of this bit is 0b.</p>
9	RW	0b	<p>IDO Completion Enable – If this bit is set, the function is permitted to set the ID-Based Ordering (<i>IDO</i>) bit (Attribute[2]) of completions it returns.</p> <p>Default value of this bit is 0b.</p>
10	RW / RsvdP	0b	<p>LTR Mechanism Enable – When set to 1b, this bit enables upstream ports to send LTR messages.</p> <p>For a multi-function device, the bit in function 0 is RW, and only function 0 controls the component’s link behavior. In all other functions of that device, this bit is RsvdP.</p> <p>If the <i>LTR_EN</i> bit in the NVM is 0b, then this bit is RO with a value of 0b.</p> <p>Default value of this bit is 0b.</p>
12:11	RO	00b	Reserved.



Bits	RW	Default	Description
14:13	RsvdP	00b	OBFF Enable. 00b = Disabled. 01b = Enabled using Message signaling [Variation A]. 10b = Enabled using Message signaling [Variation B]. 11b = Enabled using WAKE# signaling.
15	RsvdP	0b	End-End TLP Prefix Blocking. Not applicable to endpoints.

9.2.3.6.12 Link Capabilities 2 Register (0xCC)

Bits	RW	Default	Description
0	RO	0b	Reserved.
7:1	RO	0x071	Supported Link Speeds Vector – This field indicates the supported Link speed(s) of the associated Port. For each bit, a value of 1b indicates that the corresponding Link speed is supported; otherwise, the Link speed is not supported. Bit definitions are: Bit 1 - 2.5 GT/s. Bit 2 - 5.0 GT/s. Bit 3 - 8.0 GT/s. Bits 7:4 - RsvP. Multi-function devices associated with the same upstream port must report the same value in this field for all Functions. This field is loaded from NVM and is reflected in the <i>PCI_LINKCAP</i> register.
8	RO	0b	Crosslink Supported – When set to 1b, this bit indicates that the associated port supports cross links. It is recommended that this bit be set in any port that supports crosslinks even though doing so is only required for ports that also support operating at 8.0 GT/s or higher Link speeds.
31:9	RO	0x00	Reserved.

9.2.3.6.13 Link Control 2 Register (0xD0; RWS)

All RW fields in this register affect the device behavior only through function 0. In function 1 these fields are reserved read as zeros.

Note: For an integrated EP, this register should be RO and return a value of ZERO always.

Bits	RW	Default	Description
3:0	RWS (func 0) / RsvdP (else)	0001b (func 0) 0000b (else)	Target Link Speed. This field is used to set the target compliance mode speed when software is using the <i>Enter Compliance</i> bit to force a link into compliance mode. The encoding is the binary value of the bit in the Supported Link Speeds Vector (in the Link Capabilities 2 register) that corresponds to the desired target link speed. All other encodings are Reserved. For example, 5.0 GT/s corresponds to bit 2 in the Supported Link Speeds Vector, so the encoding for a 5.0 GT/s target Link speed in this field is 0010b. If a value is written to this field that does not correspond to a speed included in the <i>Supported Link Speeds</i> field, the result is undefined. The default value of this field is the highest link speed supported by the Intel® Xeon® Processor D-1500 Product Family LAN controller (as reported in the <i>Supported Link Speeds</i> field of the Link Capabilities register).
4	RWS (func 0) / RsvdP (else)	0b	Enter Compliance. Software is permitted to force a link to enter compliance mode at the speed indicated in the <i>Target Link Speed</i> field by setting this bit to 1b in both components on a link and then initiating a hot reset on the link. The default value of this field following a fundamental reset is 0b.



Bits	RW	Default	Description
5	RWS (func 0) / RsvdP (else)	0b	Hardware Autonomous Speed Disable. When set to 1b, this bit disables hardware from changing the link speed for reasons other than attempting to correct unreliable link operation by reducing link speed.
6	RO	0b	Selectable De-Emphasis. This bit is not applicable and reserved for endpoints.
9:7	RWS (func 0) / RsvdP (else)	000b	Transmit Margin. This field controls the value of the non de emphasized voltage level at the Transmitter pins. Encodings: 000b = Normal operating range. 001b = 800-1200 mV for full swing and 400-700 mV for half-swing. 010b = (n-1) — Values must be monotonic with a non-zero slope. The value of n must be greater than 3 and less than 7. At least two of these must be below the normal operating range of n: 200-400 mV for full-swing and 100-200 mV for half-swing. 111b= (n) reserved.
10	RWS (func 0) / RsvdP (else)	0b	Enter Modified Compliance. When this bit is set to 1b, the device transmits modified compliance pattern if the LTSSM enters Polling.Compliance state. The default value of this bit is 0b.
11	RWS (func 0) / RsvdP (else)	0b	Compliance SOS-When set to 1b, the LTSSM is required to send SOS periodically in between the (modified) compliance patterns. This bit is applicable when the Link is operating at 2.5 GT/s or 5 GT/s data rates only. The default value of this bit is 0b.
15:12	RWS (func 0) / RsvdP (else)	0x0	Compliance Preset/De-emphasis. For 8.0 GT/s Data Rate: This field sets the Transmitter Preset in Polling.Compliance state if the entry occurred due to the Enter Compliance bit being 1b. For 5.0 GT/s Data Rate: This field sets the de-emphasis level in Polling.Compliance state if the entry occurred due to the Enter Compliance bit being 1b. When the Link is operating at 2.5 GT/s, the setting of this bit field has no effect. Defined Encodings are: 0001b -3.5 dB. 0000b -6 dB. For a multi-function device associated with an upstream port, the bit field in function 0 is of type RWS, and only function 0 controls the component's link behavior. In all other functions of that device, this bit field is of type RsvdP. This bit field is intended for debug, and compliance testing purposes. The default value of this field is 0000b.

9.2.3.7 Link Status 2 Register (0xD2; RW)

Note: For an integrated EP, this register should be RO and return a value of ZERO always.

Bits	RW	Default	Description
0	RO	0b	Current De-emphasis Level. When the link is operating at 5 GT/s speed, this bit reflects the level of de-emphasis. it is undefined when the Link is operating at 2.5 GT/s speed. Encodings: 1b -3.5 dB. 0b -6 dB.
1	ROS (Func 0) RsvdZ (Func 1)	0b	Equalization Complete – When set to 1b, this bit indicates that the transmitter equalization procedure has completed.
2	ROS (Func 0) RsvdZ (Func 1)	0b	Equalization Phase 1 Successful. When set to 1b, this bit indicates that phase 1 of the transmitter equalization procedure has successfully completed.



Bits	RW	Default	Description
3	ROS (Func 0) RsvdZ (Func 1)	0b	Equalization Phase 2 Successful. When set to 1b, this bit indicates that phase 2 of the transmitter equalization procedure has successfully completed.
4	ROS (Func 0) RsvdZ (Func 1)	0b	Equalization Phase 3 Successful. When set to 1b, this bit indicates that phase 3 of the transmitter equalization procedure has successfully completed.
5	RW1C (Func 0) RsvdZ (Func 1)	0b	Link Equalization Request. This bit is set by hardware to request the link equalization process to be performed on the link. This bit is available only in function zero.
15:6	RsvdZ	0x00	Reserved.

9.2.4 PCIe Extended Configuration Space

PCIe configuration space is located in a flat memory-mapped address space. PCIe extends the configuration space beyond the 256 bytes available for PCI to 4096 bytes. The Intel® Xeon® Processor D-1500 Product Family LAN controller decodes an additional four bits (bits 27:24) to provide the additional configuration space as shown. PCIe reserves the remaining four bits (bits 31:28) for future expansion of the configuration space beyond 4096 bytes.

The configuration address for a PCIe device is computed using a PCI-compatible bus, device, and function numbers as follows:

31	28	27	20	19	15	14	12	11	2	1	0
0000b		Bus #		Device #		Fun #		Register Address (offset)		00b	

PCIe extended configuration space is allocated using a linked list of optional or required PCIe extended capabilities following a format resembling PCI capability structures. The first PCIe extended capability is located at offset 0x100 in the device configuration space. The first Dword of the capability structure identifies the capability/version and points to the next capability.

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports the following PCIe extended capabilities:

Table 9.16. Extended Capabilities list

Capability	Offset	Next Header	Section
Advanced Error Reporting Capability	0x100	Any of the below / 0x000	9.2.4.1
Serial Number	0x140	Any of the below / 0x000 ¹	9.2.4.2
Alternative RID Interpretation (ARI)	0x150	Any of the below / 0x000 ¹	9.2.4.3
IOV support ²	0x160	Any of the below / 0x000	9.2.4.4
TPH Requester	0x1A0	Any of the below / 0x000	9.2.4.5
Access Control Services (ACS) ³	0x1B0	Any of the below / 0x000	9.2.4.6
Latency Tolerance Reporting (LTR)	0x1C0	Any of the below / 0x000	9.2.4.7
Secondary PCIe	0x1D0	0x000	9.2.4.8

1. Depends on NVM settings enabling the ARI/IOV structures.
2. In a dummy function, the IOV structure is not exposed.
3. When in Single function mode (function 1 is disabled), the ACS capability is not exposed.



9.2.4.1 Advanced Error Reporting Capability (AER)

The PCIe advanced error reporting capability is an optional extended capability to support advanced error reporting. The tables that follow list the PCIe advanced error reporting extended capability structure for PCIe devices.

Byte Offset	Byte 3	Byte 2	Byte 1	Byte 0
0x100	Next Capability offset		Version (0x1)	AER Capability ID (0x0001)
0x104	Uncorrectable Error Status			
0x108	Uncorrectable Error Mask			
0x10C	Uncorrectable Error Severity			
0x110	Correctable Error Status			
0x114	Correctable Error Mask			
0x118	Advanced Error Capabilities and Control Register			
0x11C... 0x128	Header Log			

Table 9.17 lists configuration sharing of the AER Capability registers among the different PCI functions.

Table 9.17. Configuration Sharing of the AER Capability

Field	Sub-field	Shared?	Replicated?	Comments
Enhanced Capability Header Register	Extended Capability ID	x		
	Capability Version	x		
	Next Capability Offset		x	
Uncorrectable Error Status			x	
Uncorrectable Error Mask			x	
Uncorrectable Error Severity			x	
Correctable Error Status			x	
Correctable Error Mask			x	
Advanced Error Capabilities and Control	First Error Pointer		x	
	ECRC Generation Capable	x		
	ECRC Generation Enable		x	ECRC insertion is per PF.
	ECRC Check Capable	x		
	ECRC Check Enable		x	
Header Log			x	



9.2.4.1.1 Advanced Error Reporting Enhanced Capability Header Register (0x100; RO)

Bit Location	Attribute	Default Value	Description
15:0	RO	0x0001	Extended Capability ID. PCIe extended capability ID indicating advanced error reporting capability.
19:16	RO	0x2	Version Number. PCIe advanced error reporting extended capability version number.
31:20	RO	See description	Next Capability Offset. Next PCIe extended capability offset. See Table 9.16 and Table 9.9 for possible values of the next capability offset.

9.2.4.1.2 Uncorrectable Error Status Register (0x104; RW1CS)

The Uncorrectable Error Status register reports error status of individual uncorrectable error sources on a PCIe device. An individual error status bit that is set to 1b indicates that a particular error occurred; software can clear an error status by writing a 1b to the respective bit. Register is cleared by LAN_PWR_GOOD.

Bit Location	Attribute	Default Value	Description
0	RO	0b	Reserved.
3:1	RsvdP	0b	Reserved.
4	RW1CS	0b	Data Link Protocol Error Status.
5	RO	0b	Reserved.
11:6	RsvdP	0b	Reserved.
12	RW1CS	0b	Poisoned TLP Status.
13	RW1CS	0b	Flow Control Protocol Error Status.
14	RW1CS	0b	Completion Timeout Status.
15	RW1CS	0b	Completer Abort Status.
16	RW1CS	0b	Unexpected Completion Status.
17	RW1CS	0b	Receiver Overflow Status.
18	RW1CS	0b	Malformed TLP Status.
19	RW1CS	0b	ECRC Error Status.
20	RW1CS	0b	Unsupported Request Error Status.
21	RO	0b	ACS Violation Status. Not supported. Hardwired to 0b.
25:22	RO	0x0	Not Supported.
31:26	RsvdP	0b	Reserved.

9.2.4.1.3 Uncorrectable Error Mask Register (0x108; RWS)

The Uncorrectable Error Mask register controls reporting of individual uncorrectable errors by device to the host bridge via a PCIe error message. A masked error (respective bit set in mask register) is not reported to the host bridge by an individual device. Note that there is a mask bit per bit of the Uncorrectable Error Status register.

Bit Location	Attribute	Default Value	Description
0	RO	0b	Reserved.



Bit Location	Attribute	Default Value	Description
3:1	RsvdP	0b	Reserved.
4	RWS	0b	Data Link Protocol Error Mask.
5	RO	0b	Reserved.
11:6	RsvdP	0b	Reserved.
12	RWS	0b	Poisoned TLP Mask.
13	RWS	0b	Flow Control Protocol Error Mask.
14	RWS	0b	Completion Timeout Mask.
15	RWS	0b	Completer Abort Mask.
16	RWS	0b	Unexpected Completion Mask.
17	RWS	0b	Receiver Overflow Mask.
18	RWS	0b	Malformed TLP Mask.
19	RWS	0b	ECRC Error Mask.
20	RWS	0b	Unsupported Request Error Mask.
21	RO	0b	ACS Violation Mask.
25:22	RO	0x0	Not Supported.
31:26	RsvdP	0b	Reserved.

9.2.4.1.4 Uncorrectable Error Severity Register (0x10C; RWS)

The Uncorrectable Error Severity register controls whether an individual uncorrectable error is reported as a fatal error. An uncorrectable error is reported as fatal when the corresponding error bit in the severity register is set. If the bit is cleared, the corresponding error is considered non-fatal.

Bit Location	Attribute	Default Value	Description
0	RO	0b	Reserved.
3:1	RsvdP	0b	Reserved.
4	RWS	1b	Data Link Protocol Error Severity.
5	RO	0b	Reserved.
11:6	RsvdP	000000b	Reserved.
12	RWS	0b	Poisoned TLP Severity.
13	RWS	1b	Flow Control Protocol Error Severity.
14	RWS	0b	Completion Timeout Severity.
15	RWS	0b	Completer Abort Severity.
16	RWS	0b	Unexpected Completion Severity.
17	RWS	1b	Receiver Overflow Severity.
18	RWS	1b	Malformed TLP Severity.
19	RWS	0b	ECRC Error Severity.
20	RWS	0b	Unsupported Request Error Severity.
21	RO	0b	ACS Violation Severity.
25:22	RO	0x0	Not Supported.
31:26	RsvdP	0x0	Reserved.



9.2.4.1.5 Correctable Error Status Register (0x110; RW1CS)

The Correctable Error Status register reports error status of individual correctable error sources on a PCIe device. When an individual error status bit is set to 1b it indicates that a particular error occurred; software can clear an error status by writing a 1b to the respective bit. Register is cleared by LAN_PWR_GOOD.

Bit Location	Attribute	Default Value	Description
0	RW1CS	0b	Receiver Error Status.
5:1	RsvdZ	0b	Reserved.
6	RW1CS	0b	Bad TLP Status.
7	RW1CS	0b	Bad DLLP Status.
8	RW1CS	0b	REPLAY_NUM Rollover Status.
11:9	RsvdZ	0b	Reserved.
12	RW1CS	0b	Replay Timer Timeout Status.
13	RW1CS	0b	Advisory Non-Fatal Error Status.
15:14	RO	0b	Reserved.

9.2.4.1.6 Correctable Error Mask Register (0x114; RWS)

The Correctable Error Mask register controls reporting of individual correctable errors by device to the host bridge via a PCIe error message. A masked error (respective bit set in mask register) is not reported to the host bridge by an individual device. There is a mask bit per bit in the Correctable Error Status register.

Bit Location	Attribute	Default Value	Description
0	RWS	0b	Receiver Error Mask.
5:1	RsvdP	0b	Reserved.
6	RWS	0b	Bad TLP Mask.
7	RWS	0b	Bad DLLP Mask.
8	RWS	0b	REPLAY_NUM Rollover Mask.
11:9	RsvdP	0b	Reserved.
12	RWS	0b	Replay Timer Timeout Mask.
13	RWS	1b	Advisory Non-Fatal Error Mask. This bit is set by default to enable compatibility with software that does not comprehend Role-Based Error Reporting.
15:14	RO	0b	Reserved.

9.2.4.1.7 Advanced Error Capabilities and Control Register (0x118; RO)

Bit Location	Attribute	Default Value	Description
4:0	ROS	0b	Vector pointing to the first recorded error in the Uncorrectable Error Status register. This is a read-only field that identifies the bit position of the first uncorrectable error reported in the Uncorrectable Error Status register.
5	RO	1b	ECRC Generation Capable. If set, this bit indicates that the function is capable of generating ECRC. This bit is loaded from NVM. It is reflected in the PCI_CAPSUP register.



Bit Location	Attribute	Default Value	Description
6	RWS	0b	ECRC Generation Enable. When set, ECRC generation is enabled.
7	RO	1b	ECRC Check Capable. If set, this bit indicates that the function is capable of checking ECRC. This bit is loaded from NVM. It is reflected in the PCI_CAPSUP register.
8	RWS	0b	ECRC Check Enable. When set, ECRC checking is enabled.
9	RO	0b	Multiple Header Recording Capable. Not Supported. hardwired to 0b.
10	RO	0b	Multiple Header Recording Enable. Not Supported. hardwired to 0b.
11	RsvdP	0b	TLP Prefix Log Present. Not supported. hardwired to 0b.
15:12	RsvdP	0x0	Reserved.

9.2.4.1.8 Header Log Register (0x11C:128; RO)

The header log register captures the header for the transaction that generated an error. This register is 16 bytes.

Bit Location	Attribute	Default Value	Description
127:0	ROS	0b	Header of the packet in error (TLP or DLLP).

9.2.4.2 Serial Number

The PCIe device serial number capability is an optional extended capability that can be implemented by any PCIe device. The device serial number is a read-only 64-bit value that is unique for a given PCIe device.

The Intel® Xeon® Processor D-1500 Product Family LAN controller implements this capability on all the functions and returns the same value in both.

Byte Offset	Byte 3	Byte 2	Byte 1	Byte 0
0x140	Next Capability Offset	Version (0x1)	Serial ID Capability ID (0x0003)	
0x144	Serial Number Register (Lower Dword)			
0x148	Serial Number Register (Upper Dword)			



Table 9.18 lists configuration sharing of the Serial ID Capability registers among the different PCI functions.

Table 9.18. Configuration Sharing of the Serial Number Capability

Field	Sub-field	Shared?	Replicated?	Comments
Enhanced Capability Header Register	Extended Capability ID	x		
	Capability Version	x		
	Next Capability Offset		x	
Serial Number Registers		x		

9.2.4.2.1 Device Serial Number Enhanced Capability Header Register (0x140; RO)

Bit(s)	Default value	Attributes	Description
15:0	0x0003	RO	PCIe Extended Capability ID. This field is a PCI-SIG defined ID number that indicates the nature and format of the extended capability. The extended capability ID for the device serial number capability is 0x0003.
19:16	0x1	RO	Capability Version. This field is a PCI-SIG defined version number that indicates the version of the capability structure present. Note: Must be set to 0x1 for this version of the specification.
31:20	See description	RO	Next Capability Offset. This field contains the offset to the next PCIe capability structure or 0x000 if no other items exist in the linked list of capabilities. See Table 9.16 and Table 9.9 for possible values of the next capability offset.

9.2.4.2.2 Serial Number Registers (0x144:0x148; RO)

The Serial Number register is a 64-bit field that contains the IEEE defined 64-bit Extended Unique Identifier (EUI-64*). The register at offset 0x144 holds the lower 32 bits and the register at offset 0x148 holds the higher 32 bits. The following figure details the allocation of register fields in the Serial Number register. The table that follows provides the respective bit definitions.

Bit(s) Location	Attributes	Description
63:0	RO	PCIe Device Serial Number. This field contains the IEEE defined 64-bit EUI-64*. This identifier includes a 24-bit company ID value assigned by IEEE registration authority and a 40-bit extension identifier assigned by the manufacturer.

The serial number uses the Ethernet MAC address according to the following definition:

Field	Company ID			Extension Identifier				
Order	Addr+0	Addr+1	Addr+2	Addr+3	Addr+4	Addr+5	Addr+6	Addr+7
Most Significant Byte			Least Significant Byte					
Most Significant Bit			Least Significant Bit					

The serial number can be constructed from the 48-bit Ethernet MAC address in the following form:



Field	Company ID			MAC Label		Extension identifier		
Order	Addr+0	Addr+1	Addr+2	Addr+3	Addr+4	Addr+5	Addr+6	Addr+7
Most Significant Bytes						Least Significant Byte		
Most Significant Bit						Least Significant Bit		

In this case, the MAC label is 0xFFFF.

For example, assume that the company ID is (Intel) 00-A0-C9 and the extension identifier is 23-45-67. In this case, the 64-bit serial number is:

Field	Company ID			MAC Label		Extension Identifier		
Order	Addr+0	Addr+1	Addr+2	Addr+3	Addr+4	Addr+5	Addr+6	Addr+7
	00	A0	C9	FF	FF	23	45	67
Most Significant Byte						Least Significant Byte		
Most Significant Bit						Least Significant Bit		

The Ethernet MAC address for the serial number capability is loaded from the NVM (not the same field that is loaded from NVM into the *RAL* and *RAH* registers). It is reflected in the *PCI_SERL* and *PCI_SERH* registers. The default value in case of no NVM is 0x0.

Note: The official document that defines EUI-64* is: <http://standards.ieee.org/regauth/oui/tutorials/EUI64.html>

9.2.4.3 Alternate Routing ID Interpretation (ARI) Capability Structure

In order to allow more than eight functions per endpoint without requesting an internal switch, as is usually needed in virtualization scenarios, the PCI-SIG defines a new capability that allows a different interpretation of the *Bus*, *Device*, and *Function* fields. The capability is exposed when the *PCI_CAPSUP.ARI_EN* bit is set from NVM.

The ARI capability structure is as follows:

Byte Offset	Byte 3	Byte 2	Byte 1	Byte 0
0x150	Next Capability Offset		Version (0x1)	ARI Capability ID (0x000E)
0x154	ARI Control Register		ARI Capabilities	

Table 9.19 lists configuration sharing of the ARI Capability registers among the different PCI functions.

Table 9.19. Configuration sharing of the ARI Capability

Field	Sub-field	Shared?	Replicated?	Comments
Enhanced Capability Header Register	Extended Capability ID	x		



Table 9.19. Configuration sharing of the ARI Capability

Field	Sub-field	Shared?	Replicated?	Comments
	Capability Version	x		
	Next Capability Offset		x	
ARI capability Register	Next Function Pointer		x	

9.2.4.3.1 PCIe ARI Header Register (0x150; RO)

Field	Bit(s)	Initial Value	Access	Description
ID	15:0	0x000E	RO	PCIe Extended Capability ID. PCIe extended capability ID for the alternative RID interpretation.
Version	19:16	1b	RO	Capability Version. This field is a PCI-SIG defined version number that indicates the version of the capability structure present. Must be 0x1 for this version of the specification.
Next Capability Offset	31:20	See description	RO	Next Capability Offset. This field contains the offset to the next PCIe extended capability structure. See Table 9.16. and Table 9.9 for possible values of the next capability offset.

9.2.4.3.2 PCIe ARI Capabilities and Control Register (0x154; RO)

Field	Bit(s)	Initial Value	Access	Description
Reserved	0	0b	RO	Not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
Reserved	1	0b	RO	Not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
Reserved	7:2	0b	RO	Reserved.
NFP	15:8	0x1 (func 0) 0x0 (func 1) ¹	RO	Next Function Pointer. This field contains the pointer to the next physical function configuration space or 0x0000 if no other items exist in the linked list of functions. Function 0 is the start of the link list of functions.
Reserved	16	0b	RO	Not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
Reserved	17	0b	RO	Not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
Reserved	19:18	00b	RO	Reserved.
Reserved	22:20	0b	RO	Not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
Reserved	31:23	0b	RO	Reserved.

1. Even if port 0 and port 1 are switched or function zero is a dummy function, this register should keep its attributes according to the function number. If LAN1 is disabled, the value of this field in function zero should be zero.

9.2.4.4 IOV Capability Structure

Note: This capability structure is not exposed in a dummy function.

This is a structure used to support the SR-IOV capabilities reporting and control. The capability is exposed when the *PCI_CAPSUP.IOV_EN* bit is set from NVM and the *PCI_CNF2.NUM_VFS* is non-zero.



The following tables show the implementation of this structure in the Intel® Xeon® Processor D-1500 Product Family LAN controller.

Byte Offset	Byte 3	Byte 2	Byte 1	Byte 0
0x160	Next Capability Offset (0x0)	Version (0x1)	IOV Capability ID (0x0010)	
0x164	SR IOV Capabilities			
0x168	SR IOV Status		SR IOV Control	
0x16C	Total VFs (RO)		Initial VF (RO)	
0x170	Reserved	Function Dependency Link (RO)	Num VF (RW)	
0x174	VF Stride (RO)		First VF Offset (RO)	
0x178	VF Device ID		Reserved	
0x17C	Supported Page Size (0x553)			
0x180	System Page Size (RW)			
0x184	VF BAR0 — Low (RW)			
0x188	VF BAR0 — High (RW)			
0x18C	VF BAR2 (RO)			
0x190	VF BAR3- Low (RW) (64 bit)			
0x194	VF BAR3 High (RW) (64 bit)			
0x198	VF BAR5 (RO)			
0x19C	VF Migration State Array Offset (RO)			

Table 9.20 lists configuration sharing of the SR-IOV Capability registers among the different PCI functions.

Table 9.20. Configuration Sharing of the SR-IOV Capability

Field	Sub-field	Shared?	Replicated?	Comments
Enhanced Capability Header Register	Extended Capability ID	x		
	Capability Version	x		
	Next Capability Offset		x	
SR-IOV Capabilities	VF Migration Capable	x		Not supported.
	ARI Capable Hierarchy Preserved		x	PF0 only. RO zero in all other functions.
	VF Migration Interrupt Message Number			Not supported.
SR-IOV Control	VF Enable		x	
	Memory Space Enable		x	
	ARI Capable Hierarchy	x		PF0 only. RO zero in all other functions.
InitialVFs			x	
TotalVFs			x	
Num VFs			x	



Table 9.20. Configuration Sharing of the SR-IOV Capability (Continued)

Field	Sub-field	Shared?	Replicated?	Comments
Function Dependency Link			x	Each PF indicates its PF number here.
First VF Offset			x	
VF Stride			x	
VF Device ID			x	
Supported Page Size		x		
System Page Size			x	
VF BARs			x	

9.2.4.4.1 PCIe SR-IOV Header Register (0x160; RO)

Field	Bit(s)	Initial Value	Access	Description
ID	15:0	0x0010	RO	PCIe Extended Capability ID. PCIe extended capability ID for the SR-IOV capability.
Version	19:16	0x1	RO	Capability Version. This field is a PCI-SIG defined version number that indicates the version of the capability structure present. Must be 0x1 for this version of the specification.
Next offset	31:20	See description	RO	Next Capability Offset. This field contains the offset to the next PCIe extended capability structure or 0x000 if no other items exist in the linked list of capabilities. See Table 9.16, and Table 9.9 for possible values of the next capability offset.

9.2.4.4.2 PCIe SR-IOV Capabilities Register (0x164; RO)

Field	Bit(s)	Initial Value	Access	Description
Reserved	0	0b	RO	Not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
ARICHP	1	1b/0b ¹	RO	ARI Capable Hierarchy Preserved - If set, the <i>ARI Capable Hierarchy</i> bit is preserved across certain power state transitions.
Reserved	20:2	0x0	RO	Reserved.
Reserved	31:21	0x0	RO	Not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller.

1. Set on first function where SR-IOV is enabled (see *PCI_CAPSUP.IOV_EN* bit) and Read Only Zero in the other PF.



9.2.4.4.3 PCIe SR-IOV Control/Status Register (0x168; RW)

Field	Bit(s)	Initial Value	Access	Description
VFE	0	0b	RW	<p>VF Enable/Disable.</p> <p>VF Enable manages the assignment of VFs to the associated PF. If <i>VF Enable</i> is set to 1b, VFs must be enabled, associated with the PF, and exists in the PCIe fabric. When enabled, VFs must respond to and can issue PCIe transactions following all other rules for PCIe functions.</p> <p>If set to 0b, VFs must be disabled and not visible in the PCIe fabric; VFs cannot respond to or issue PCIe transactions.</p> <p>In addition, if <i>VF Enable</i> is cleared after having been set, all of the VFs must no longer:</p> <ul style="list-style-type: none"> • Issue PCIe transactions • Respond to configuration space or memory space accesses. <p>The behavior must be as if an FLR was issued to each of the VFs. Specifically, VFs must not retain any context after <i>VF Enable</i> has been cleared. Any errors already logged via PF error reporting registers, remain logged. However, no new VF errors must be logged after <i>VF Enable</i> is cleared.</p>
Reserved	1	0b	RO	Not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
Reserved	2	0b	RO	Not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
VF MSE	3	0b	RW	<p>Memory Space Enable for Virtual Functions.</p> <p>VF MSE controls memory space enable for all VFs associated with this PF as with the <i>Memory Space Enable</i> bit in a functions PCI command register. The default value for this bit is 0b.</p> <p>When VF Enable is 1, virtual function memory space access is permitted only when VF MSE is Set. VFs must follow the same error reporting rules as defined in the base specification if an attempt is made to access a virtual functions memory space when VF Enable is 1 and VF MSE is zero.</p> <p>Implementation Note: Virtual functions memory space cannot be accessed when VF Enable is zero. Thus, VF MSE is don't care when VF Enable is zero, however, software might choose to set VF MSE after programming the VF BARn registers, prior to setting VF Enable to 1b.</p>
ARI Capable Hierarchy	4	0b	RW (first function where SR-IOV is enabled) ROS (other function) ¹	<p>ARI Capable Hierarchy.</p> <p>The Intel® Xeon® Processor D-1500 Product Family LAN controller is permitted to locate VFs in function numbers 8 to 255 of the captured bus number.</p> <p>This field is R/W in the lowest numbered PF. Other functions use the PF0 value as sticky.</p> <p>Note: If either <i>ARI Capable Hierarchy Preserved</i> is set (see Section 9.2.4.4.2) or <i>No_Soft_Reset</i> is set, a power state transition of this PF from D3hot to D0 does not affect the value of this bit.</p>
Reserved	15:5	0x0	RO	Reserved.
Reserved	16	0b	RO	Not implemented in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
Reserved	31:17	0b	RO	Reserved.

1. Even if port 0 and port 1 are switched or function zero is a dummy function, this field should keep its attributes according to the function number.



9.2.4.4.4 PCIe SR-IOV Max/Total VFs Register (0x16C; RO)

Field	Bit(s)	Initial Value	Access	Description
InitialVFs	15:0	64	RO	InitialVFs indicates the number of VFs that are initially associated with the PF. If <i>VF Migration Capable</i> is cleared, this field must contain the same value as TotalVFs. In the Intel® Xeon® Processor D-1500 Product Family LAN controller this parameter is equal to the TotalVFs in this register.
TotalVFs	31:16	64	RO	TotalVFs defines the maximum number of VFs that can be associated with the PF. This field is loaded from the <i>Max VFs</i> field in the NVM. Reflected in <i>PCI_CNFG2.TOTAL_VFS</i> .


9.2.4.4.5 PCIe SR-IOV Num VFs Register (0x170; RW)

Field	Bit(s)	Initial Value	Access	Description
NumVFs	15:0	0x0	RW	Num VFs defines the number of VFs software has assigned to the PF. Software sets NumVFs to any value between one and the TotalVFs as part of the process of creating VFs. NumVFs VFs must be visible in the PCIe fabric after both NumVFs is set to a valid value <i>VF Enable</i> is set to 1b.
FDL	23:16	0x0 (func 0) 0x1 (func 1) ¹	RO	Function Dependency Link. Defines dependencies between physical functions allocation. In the Intel® Xeon® Processor D-1500 Product Family LAN controller there are no constraints.
Reserved	31:24	0	RO	Reserved.

1. Even if port 0 and port 1 are switched or function zero is a dummy function, this register should keep it's attributes according to the function number.

9.2.4.4.6 PCIe SR-IOV VF RID Mapping Register (0x174; RO)

Field	Bit(s)	Initial Value	Access	Description
FVO	15:0	0x180	RO	First VF offset defines the requestor ID (RID) offset of the first VF that is associated with the PF that contains this capability structure. The first VFs 16-bit RID is calculated by adding the contents of this field to the RID of the PF containing this field. The content of this field is valid only when <i>VF Enable</i> is set. If <i>VF Enable</i> is 0b, the contents are undefined. If the <i>ARI Enable</i> bit is set, this field changes to 0x80.
VFS	31:16	0x2	RO	VF stride defines the requestor ID (RID) offset from one VF to the next one for all VFs associated with the PF that contains this capability structure. The next VFs 16-bit RID is calculated by adding the contents of this field to the RID of the current VF. The contents of this field is valid only when <i>VF Enable</i> is set and <i>NumVFs</i> is non-zero. If <i>VF Enable</i> is 0b or if <i>NumVFs</i> is zero, the contents are undefined.



9.2.4.4.7 PCIe SR-IOV VF Device ID Register (0x178; RO)



Field	Bit(s)	Initial Value	Access	Description
DEVID	31:16	See Section 9.2.2.2	RO	VF Device ID. This field contains the device ID that should be presented for every VF to the Virtual Machine (VM). The value of this field can be read from the IOV Control Word 2 in the NVM.
Reserved	15:0	0x0	RO	Reserved.

9.2.4.4.8 PCIe SR-IOV Supported Page Size Register (0x17C; RO)

Field	Bit(s)	Initial Value	Access	Description
Supported page Size	31:0	0x553	RO	For PFs that supports the stride-based BAR mechanism, this field defines the supported page sizes. This PF supports a page size of $2^{(n+12)}$ if bit n is set. For example, if bit 0 is Set, the Endpoint (EP) supports 4KB page sizes. Endpoints are required to support 4 KB, 8 KB, 64 KB, 256 KB, 1 MB and 4 MB page sizes. All other page sizes are optional.

9.2.4.4.9 PCIe SR-IOV System Page Size Register (0x180; RW)

Field	Bit(s)	Initial Value	Access	Description
Page size	31:0	0x1	RW	This field defines the page size the system uses to map the PF's and associated VFs' memory addresses. Software must set the value of the <i>System Page Size</i> to one of the page sizes set in the <i>Supported Page Sizes</i> field. As with <i>Supported Page Sizes</i> , if bit n is set in <i>System Page Size</i> , the PF and its associated VFs are required to support a page size of $2^{(n+12)}$. For example, if bit 1 is set, the system is using a 8 KB page size. The results are undefined if more than one bit is set in <i>System Page Size</i> . The results are undefined if a bit is set in <i>System Page Size</i> that is not set in <i>Supported Page Sizes</i> . When <i>System Page Size</i> is set, the PF and associated VFs are required to align all BAR resources on a <i>System Page Size</i> boundary. Each BAR size, including <i>VF BARn Size</i> (described later) must be aligned on a <i>System Page Size</i> boundary. Each BAR size, including <i>VF BARn Size</i> must be sized to consume a multiple of <i>System Page Size</i> bytes. All fields requiring page size alignment within a function must be aligned on a <i>System Page Size</i> boundary. <i>VF Enable</i> must be zero when <i>System Page Size</i> is set. The results are undefined if <i>System Page Size</i> is set when <i>VF Enable</i> is set.

9.2.4.4.10 PCIe SR-IOV BAR 0 — Low Register (0x184; RW)

Field	Bit(s)	Initial Value	Access	Description
Mem	0	0b	RO	0b indicates memory space.
Mem Type	2:1	10b	RO	Indicates the address space size. 10b = 64-bit. This bit is loaded from the NVM. It is reflected in the <i>PCI_VFSUP</i> register.
Prefetch Mem	3	0b*	RO	0b = Non-prefetchable space. 1b = Prefetchable space. This bit is loaded from the NVM. It is reflected in the <i>PCI_VFSUP</i> register.
Memory Address Space	31:4	0x0	RW	Which bits are RW bits and which are RO to 0x0 depend on the memory mapping window size. The size is a maximum between 16 KB and the page size.



9.2.4.4.11 PCIe SR-IOV BAR 0 — High Register (0x188; RW)

Field	Bit(s)	Initial Value	Access	Description
BAR0 — MSB	31:0	0x0	RW	MSB part of BAR0.

9.2.4.4.12 PCIe SR-IOV BAR 2 Register (0x18C; RO)

Field	Bit(s)	Initial Value	Access	Description
BAR2	31:0	0x0	RO	This BAR is not used.

9.2.4.4.13 PCIe SR-IOV BAR 3 — Low Register (0x190; RW)

Field	Bit(s)	Initial Value	Access	Description
Mem	0	0b	RO	0b indicates memory space.
Mem Type	2:1	10b	RO	Indicates the address space size. 10b = 64-bit. This bit is loaded from the NVM. It is reflected in the <i>PCI_VFSUP</i> register.
Prefetch Mem	3	0b*	RO	0b = Non-prefetchable space. 1b = Prefetchable space. This bit is loaded from the NVM. It is reflected in the <i>PCI_VFSUP</i> register.
Memory Address Space	31:4	0x0	RW	Which bits are RW bits and which are RO to 0x0 depend on the memory mapping window size. The size is a maximum between 16 KB and page size.

9.2.4.4.14 PCIe SR-IOV BAR 3 — High Register (0x194; RW)

Field	Bit(s)	Initial Value	Access	Description
BAR4 — MSB	31:0	0x0	RW	MSB part of BAR3.

9.2.4.4.15 PCIe SR-IOV BAR 5 Register (0x198; RO)

Field	Bit(s)	Initial Value	Access	Description
BAR2	31:0	0x0	RO	This BAR is not used.

9.2.4.4.16 PCIe SR-IOV VF Migration State Array Offset Register (0x19C; RO)

Field	Bit(s)	Initial Value	Access	Description
Reserved	2:0	0x0	RO	Not implemented in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
Reserved	31:3	0x0	RO	Not implemented in the Intel® Xeon® Processor D-1500 Product Family LAN controller.



9.2.4.5 TPH Requester Capability

The TPH Requester capability is an optional extended capability to support TLP processing hints. The following table lists the TPH extended capability structure for PCIe devices.

Byte Offset	Byte 3	Byte 2	Byte 1	Byte 0
0x1A0	Next Capability Offset		Version (0x1)	TPH Capability ID (0x000E)
0x1A4	TPH Requester Capability Register			
0x1A8	TPH Requester Control Register			

Table 9.2.4.5.3 lists configuration sharing of the TPH Requester Capability registers among the different PCI functions.

Table 9.21. Configuration Sharing of the TPH Requester Capability


Field	Sub-field	Shared?	Replicated?	Comments
Enhanced Capability Header Register	Extended Capability ID	x		
	Capability Version	x		
	Next Capability Offset		x	
TPH Requester Capability		x		
TPH Requester Control			x	
TPH ST Table			x	The <i>Steering Table Upper</i> fields are not supported.

9.2.4.5.1 TPH Requester Extended Capability Header (0x1A0; RO)

Bit Location	Attribute	Default Value	Description
15:0	RO	0x17	Extended Capability ID. PCIe extended capability ID indicating TPH capability.
19:16	RO	0x1	Capability Version. PCIe TPH extended capability version number.
31:20	RO	See Description	Next Capability Offset. This field contains the offset to the next PCIe capability structure. See Table 9.16, and Table 9.9 for possible values of the next capability offset.

9.2.4.5.2 TPH Requester Capability Register (0x1A4; RO)

Bit Location	Attribute	Default Value	Description
0	RO	1	No ST Mode Supported. If set indicates that the function supports the no ST mode of operation.
1	RO	0	Interrupt Vector Mode Supported. Cleared to indicates that the Intel® Xeon® Processor D-1500 Product Family LAN controller does not support interrupt vector mode of operation.



2	RO	1	Device Specific Mode. Set to indicate that the Intel® Xeon® Processor D-1500 Product Family LAN controller supports device specific mode of operation.
7:3	RO	0	Reserved.
8	RO	0	Extended TPH Requester Supported. Cleared to indicate that the function is not capable of generating requests with extended TPH TLP prefix.
10:9	RO	00b	ST Table Location. Value indicates if and where the ST table is located. Defined encodings are: 00b = ST table is not present. 01b = ST table is located in the TPH requester capability structure. 10b = ST table is located in the MSI-X table structure. 11b = Reserved. ST table is not supported.
15:11	RO	0x0	Reserved.
26:16	RO	0x0	ST Table Size. System software reads this field to determine the ST_Table_Size N, which is encoded as N-1. For example, a returned value of "00000000011" indicates a table size of 4. The value in this field is undefined since the Intel® Xeon® Processor D-1500 Product Family LAN controller does not support an ST Table.
31:27	RO	0x0	Reserved.

9.2.4.5.3 TPH Requester Control Register (0x1A8; R/W)

Bit Location	Attribute	Default Value	Description
2:0	RW	0x0	ST Mode Select. Indicates the ST mode of operation selected. Defined encodings are: 000b = No table mode. 001b = Interrupt vector mode (not supported by the Intel® Xeon® Processor D-1500 Product Family LAN controller). 010b = Device specific mode. Others = Reserved for future use. The default value of 000b indicates no table mode of operation. Functions that support only the no ST mode of operation must hardwire this field to 000b.
7:3	RO	0x0	Reserved.
9:8	RW	0x0	TPH Requester Enable. Controls the ability to issue request TLPs using either TPH or extended TPH. Defined encodings are: 00b = The Intel® Xeon® Processor D-1500 Product Family LAN controller is not permitted to issue transactions with TPH or extended TPH as requester. 01b = The Intel® Xeon® Processor D-1500 Product Family LAN controller is permitted to issue transactions with TPH as requester and is not permitted to issue transactions with extended TPH as requester 10b = Reserved. 11b = The Intel® Xeon® Processor D-1500 Product Family LAN controller is permitted to issue transactions with TPH and extended TPH as requester (the Intel® Xeon® Processor D-1500 Product Family LAN controller does not issue transactions with extended TPH). The default value of this field is 00b.
31:10	RO	0x0	Reserved.



9.2.4.6 Access Control Services (ACS) Capability

The PCIe ACS defines a set of control points within a PCIe topology to determine whether a TLP should be routed normally, blocked, or redirected. ACS is applicable to RCs, switches, and multifunction devices and is not exposed in Single function mode (when function 1 is disabled). The capability is exposed when the PCI_CAPSUP.ACS_EN bit is set from NVM.

The ACS capability structure is shared and exposed to all PFs.

The following table lists the PCIe ACS extended capability structure for PCIe devices.

Byte Offset	Byte 3	Byte 2	Byte 1	Byte 0
0x1B0	Next Capability Offset		Version (0x1)	ACS Capability ID (0x0D)
0x1B4	ACS Control Register (0x0)		ACS Capability Register (0x0)	

9.2.4.6.1 ACS CAP ID (0x1B0; RO)

Bit Location	Attribute	Default Value	Description
15:0	RO	0x0D	ACS Capability ID. PCIe extended capability ID indicating ACS capability.
19:16	RO	0x1	Version Number. PCIe ACS extended capability version number.
31:20	RO	See description	Next Capability Pointer. See Table 9.16. and Table 9.9 for possible values of the next capability offset.

9.2.4.6.2 ACS Control and Capabilities (0x1B4; RO)

Bit Location	Attribute	Default Value	Description
0	RO	0b	ACS Source Validation (V). Hardwired to zero, not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
1	RO	0b	ACS Translation Blocking (B). Hardwired to zero, not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
2	RO	0b	ACS P2P Request Redirect (R). Hardwired to zero, not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
3	RO	0b	ACS P2P Completion Redirect (C). Hardwired to zero, not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
4	RO	0b	ACS Upstream Forwarding (U). Hardwired to zero, not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
5	RO	0b	ACS P2P Egress Control (E). Hardwired to zero, not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller.



Bit Location	Attribute	Default Value	Description
6	RO	0b	ACS Direct Translated P2P (T). Hardwired to zero, not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
7	Rsrv	0b	Reserved.
15:8	RO	0x0	Egress Control Vector Size. Hardwired to zero, not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
16	RO	0b	ACS Source Validation Enable (V). Hardwired to zero, not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
17	RO	0b	ACS Translation Blocking Enable (B). Hardwired to zero, not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
18	RO	0b	ACS P2P Request Redirect Enable (R). Hardwired to zero, not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
19	RO	0b	ACS P2P Completion Redirect Enable (C). Hardwired to zero, not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
20	RO	0b	ACS Upstream Forwarding Enable (U). Hardwired to zero, not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
21	RO	0b	ACS P2P Egress Control Enable (E). Hardwired to zero, not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
22	RO	0b	ACS Direct Translated P2P Enable (T). Hardwired to zero, not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
31:23	Rsrv	0b	Reserved.

9.2.4.7 Latency Tolerance Reporting (LTR) Capability Structure

The LTR capability is an optional extended capability that enables software to provide platform latency information to devices with upstream ports (endpoints and switches). This capability structure is required if the device supports LTR. The capability is exposed when the *PCI_CAPSUP.LTR_EN* bit is set from NVM.

The LTR capability structure is implemented only in function 0 even when function 0 is a dummy function and controls the component's link behavior on behalf of all the functions of the device.

The following table lists the PCIe LTR extended capability structure for PCIe devices.

Byte Offset	Byte 3	Byte 2	Byte 1	Byte 0
0x1C0	PCIe Extended Capability Header			
0x1C4	Max No-Snoop Latency Register		Max Snoop Latency Register	



9.2.4.7.1 LTR Extended Capability Header (0x1C0; RO)

Bit Location	Attribute	Default Value	Description
15:0	RO	0x18	PCIe Extended Capability ID. PCIe extended capability ID indicating LTR capability.
19:16	RO	0x1	Capability Version. PCIe LTR extended capability version number.
31:20	RO	See Description	Next Capability Offset. See Table 9.16 . and Table 9.9 for possible values of the next capability offset.

9.2.4.7.2 Max Snoop Latency Register (0x1C4; RW)

Bit Location	Attribute	Default Value	Description
9:0	RW	0x0	Maximum Snoop Latency Value. Along with the <i>Max Snoop Latency Scale</i> field, this register specifies the maximum snoop latency that a device is permitted to request. Software should set this to the platform's maximum supported latency or less. Field is also an indicator of the platform maximum latency, should an endpoint send up LTR latency values with the requirement bit not set.
12:10	RW	0x0	Max Snoop Latency Scale. This field provides a scale for the value contained within the <i>Maximum Snoop Latency Value</i> field. Encoding: 000b = Value times 1 ns. 001b = Value times 32 ns. 010b = Value times 1,024 ns. 011b = Value times 32,768 ns. 100b = Value times 1,048,576 ns. 101b = Value times 33,554,432 ns. 110b - 111b = Not permitted.
15:13	RW	0x0	Reserved.

9.2.4.7.3 Max No-Snoop Latency Register (0x1C6; RW)

Bit Location	Attribute	Default Value	Description
9:0	RW	0x0	Maximum No-Snoop Latency Value. Along with the <i>Max No-Snoop Latency Scale</i> field, this register specifies the maximum no-snoop latency that a device is permitted to request. Software should set this to the platform's maximum supported latency or less. Field is also an indicator of the platform's maximum latency, should an endpoint send up LTR Latency Values with the Requirement bit not set.
12:10	RW	0x0	Max No-Snoop Latency Scale. This field provides a scale for the value contained within the <i>Maximum No-Snoop Latency Value</i> field. Encoding: 000b = Value times 1 ns. 001b = Value times 32 ns. 010b = Value times 1,024 ns. 011b = Value times 32,768 ns. 100b = Value times 1,048,576 ns. 101b = Value times 33,554,432 ns. 110b-111 = Not permitted.
15:13	RW	0x0	Reserved.



9.2.4.8 Secondary PCIe Extended Capability

The secondary PCIe extended capability structure is required for all ports and RCRBs that support a Link speed of 8.0 GT/s or higher. For multi-function upstream ports, this capability must be implemented in function 0 and must not be implemented in other functions. The capability is exposed when the PCI_CAPSUP.SEC_EN bit is set from NVM.

The following table lists the secondary PCIe extended capability structure for PCIe devices.

Byte Offset	Byte 3	Byte 2	Byte 1	Byte 0
0x1D0	PCIe Extended Capability Header			
0x1D4	Link Control 3 Register			
0x1D8	Lane Error Status Register			
0x1DC	Equalization Control Register (Sized by Maximum Link Width)			
0x1E0	Equalization Control Register (Sized by Maximum Link Width)			

9.2.4.8.1 Secondary PCIe Extended Capability Header (0x1D0)

Bit Location	Attribute	Default Value	Description
15:0	RO	0x19	PCIe Extended Capability ID. This field is a PCI-SIG defined ID number that indicates the nature and format of the extended capability. PCIe extended capability ID for the secondary PCIe extended capability is 0x0019.
19:16	RO	0x1	Capability Version. This field is a PCI-SIG defined version number that indicates the version of the capability structure present. Must be 0x1 for this version of the specification.
31:20	RO	See Description	Next Capability Offset. See Table 9.16 . and Table 9.9 for possible values of the next capability offset.

9.2.4.8.2 Link Control 3 Register (0x1D4)

Bit Location	Attribute	Default Value	Description
0	RW/RsvdP	0b	Perform Equalization. When this bit is 1b and a 1b is written to the Link Retrain register with <i>Target Link Speed</i> set to 8.0 GT/s, the downstream port must perform transmitter equalization. This bit is RW for upstream ports when crosslink supported is 1b. This bit is not applicable and is RsvdP for upstream ports when the crosslink supported bit is 0b. The default value is 0b.
1	RW/RsvdP	0b	Link Equalization Request Interrupt Enable. When set, this bit enables the generation of interrupt to indicate that the link equalization Request bit has been set. This bit is RW for upstream ports when crosslink supported is 1b. This bit is not applicable and is RsvdP for upstream ports when the crosslink supported bit is 0b. The default value for this bit is 0b.
31:2	RsvdZ	0x00	Reserved



9.2.4.8.3 Lane Error Status Register (0x1D8)

The Lane Error Status register consists of a 32-bit vector, where each bit indicates if the corresponding PCI Express Lane detected an error.

Bit Location	Attribute	Default Value	Description
7:0	RW1CS	0x00	Lane Error Status Bits. Each bit indicates if the corresponding lane detected a lane-based error. A value of 1b indicates that a lane based-error was detected on the corresponding lane number. The default value of this field is 0b. This field is intended for debug purposes only.
31:8	RsvdZ	0x00	Reserved.

9.2.4.8.4 Lane Equalization Control Register (0x1DC: 0x1E3)

The Equalization Control register consists of control fields required for per Lane equalization and number of entries in this register are sized by Maximum Link Width.




Table 9.22. Lane [(Maximum Link Width – 1):0] Equalization Control Register

Bit Location	Attribute	Default Value	Description
3:0	RsvdP	0000b	Downstream Port Transmitter Preset. For an upstream port if crosslink supported is 0b, this field is RsvdP.
6:4	RsvdP	000b	Downstream Port Receiver Preset Hint. For an upstream port if crosslink supported is 0b, this field is RsvdP.
11:8	RO	1111b	Upstream Port Transmitter Preset. Field contains the transmit preset value sent or received during link equalization. Since crosslink is not supported, the field is intended for debug and diagnostics. It contains the value captured from the associated lane during link equalization. Field is RO. The default value is 1111b.
14:12	HwInit/RO	111b	Upstream Port Receiver Preset Hint. This field contains the receiver preset hint value sent or received during link equalization. Field usage varies as follows: Since crosslink is not supported, this field is intended for debug and diagnostics. It contains the value captured from the associated lane during link equalization. Field is RO. The default value is 111b.
15	Rsvd	0b	Reserved.

9.2.5 Driver Forward Compatibility Register (0x94; RO)

This register is an advisory register that returns a fixed value indicating the type of software device driver class needed for the Intel® Xeon® Processor D-1500 Product Family LAN controller.



Bit(s)	R/W	Initial Value	Description
15:0	RO	0xFFFF	Pointer to admin. queues. A value of 0xFFFF indicates that the Intel® Xeon® Processor D-1500 Product Family LAN controller does not support admin. queues.
19:16	RO	0x0	Status. Not implemented in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
31:20	RO	0xA0A	Signature indicating a driver class register.

9.2.6 CSR Access Via Configuration Address Space

The registers described in this section are not part of the PCIe standard configuration and can be used to access the CSR space before memory BARs are allocated. When this mechanism is used, there is no need to expose an I/O BAR for pre boot operation.

Note: These registers are not available for dummy functions.

9.2.6.1 IOADDR Register (0x98; R/W)

This is a read/write register. Each function has its own IOADDR register. Functionality is the same in all functions. Register is cleared at power-up (LAN_PWR_GOOD) or PCIe reset.

Note: When functioning in a D3 state, software should not attempt to access CSRs via the IOADDR and IODATA registers.

Bit(s)	R/W	Initial Value	Description
30:0	R/W ¹	0x0	Internal Register or Internal Memory Location Address. 0x00000-0x1FFFF – Internal registers and memories. 0x20000-0x7FFFFFFF – Undefined.
31	R/W	0b	Configuration I/O Access Enable. 0b = CSR configuration read or write disabled. 1b = CSR configuration read or write enabled. When this bit is set, accesses to the IODATA register actually generate transactions to the Intel® Xeon® Processor D-1500 Product Family LAN controller. Otherwise, accesses to the IODATA register are don't-cares (writes are discarded silently, reads return arbitrary results).

1. In the event that the *PCI_CAPSUP.CSR_CONF_EN* bit is cleared, accesses to the IOADDR register via the configuration address space is ignored and has no effect on the register and the CSRs referenced by the IOADDR register.

9.2.6.2 IODATA Register (0x9C; R/W)

This is a read/write register. Each function has its own IODATA register. Functionality is the same in all functions. Register is cleared at power-up (LAN_PWR_GOOD) or PCIe reset.

Bit(s)	R/W	Initial Value	Description
31:0	R/W ¹	0x0	Data field for reads or writes to the Internal register or internal memory location as identified by the current value in IOADDR. All 32 bits of this register can be read from or written to.

1. In the event that the *IO_by_cfg* bit in the PCIe Init Configuration 2 EEPROM word is cleared, access to the IODATA register via the configuration address space is ignored and has no effect on the register and the CSRs referenced by the IOADDR register.



9.3 Virtual Functions Configuration Space

The configuration space reflected to each of the VF is a sparse version of the physical function configuration space. The following table lists the behavior of each register in the VF configuration space.


Table 9.23. VF PCIe Configuration Space

Section	Offset	Name	VF Behavior	Notes
PCI Mandatory Registers	0	Vendor ID	RO — 0xFFFF	
	2	Device ID	RO — 0xFFFF	
	4	Command	Per VF	See Section 9.3.1.1 .
	6	Status	Per VF	See Section 9.3.1.2 .
	8	RevisionID	RO as PF	
	9	Class Code	RO as PF	
	C	Cache Line Size	RO — 0x0	
	D	Latency Timer	RO — 0x0	
	E	Header Type	RO — 0x0	
	F	Reserved	RO — 0x0	
	10 — 27	BARs	RO — 0x0	Emulated by VMM.
	28	CardBus CIS	RO — 0x0	Not used.
	2C	Sub Vendor ID	RO as PF	
	2E	Sub System	RO as PF	
	30	Expansion ROM	RO — 0x0	Emulated by VMM.
	34	Cap Pointer	RO — 0x70	Next = MSI-X capability.
	3C	Int Line	RO — 0x0	
	3D	Int Pin	RO — 0x0	
3E	Max Lat/Min Gnt	RO — 0x0		
MSI-X Capability	70	MSI-X Header	RO — 0xA011	Next = PCIe capability.
	72	MSI-x Message Control	per VF	See Section 9.3.2.1 .
	74	MSI-X table Address	RO — as PF	See Section 9.3.2.1.2 .
	78	MSI-X PBA Address	RO	See Section 9.3.2.1.3 .
PCIe Capability	A0	PCIe Header	RO — 0x0010	Next = Last capability.
	A2	PCIe Capabilities	RO — as PF	
	A4	PCIe Dev Cap	RO — as PF	
	A8	PCIe Dev Ctrl	RW	As PF apart from FLR — See Table 9.3.2.2.1 .
	AA	PCIe Dev Status	per VF	See Table 9.3.2.2.2 .
	AC	PCIe Link Cap	RO — as PF	
	B0	PCIe Link Ctrl	RO — 0x0	
	B2	PCIe Link Status	RO — 0x0	
	C4	PCIe Dev Cap 2	RO — as PF	
	C8	PCIe Dev Ctrl 2	RO — 0x0	
	D0	PCIe Link Ctrl 2	RO — 0x0	
	D2	PCIe Link Status 2	RO — 0x0	



Table 9.23. VF PCIe Configuration Space (Continued)

Section	Offset	Name	VF Behavior	Notes
AER Capability	100	AER — Header	RO — 0x15010001	Next = ARI structure.
	104	AER — Uncorr Status	per VF	See Section 9.3.3.1.1.
	108	AER — Uncorr Mask	RO — 0x0	
	10C	AER — Uncorr Severity	RO — 0x0	
	110	AER — Corr Status	Per VF	See Section 9.3.3.1.2.
	114	AER — Corr Mask	RO — 0x0	
	118	AER — Cap/Ctrl	RO as PF	
	11C — 128	AER — Error Log	Shared two logs for all VFs	Same structure as in PF. In case of overflow, the header log is filled with ones.
ARI Capability	150	ARI — Header	0x1A01000E	Next = TPH.
	154	ARI — Cap/Ctrl	RO — 0X0	
TPH Capability	1A0	TPH - Header	RO - 0x1B010017	Next = ACS.
	1A4	TPH - Capability	RO - 0x00000005	No ST Table.
	1A8	TPH - Control	Per VF	Same structure as PF.
ACS capability	1B0	ACS - Header	RO - 0x0001000D	Last extended capability.
	1B4	ACS - Capability	RO - 0x00000000	



9.3.1 Mandatory Configuration Space

9.3.1.1 VF Command Register (0x4; RW)

Bit(s)	Initial Value	Rd/Wr	Description
0	0b	RO	IOAE: I/O Access Enable. RO as zero field.
1	0b	RO	MAE: Memory Access Enable. RO as zero field.
2	0b	RW	BME: Bus Master Enable. Disabling this bit prevents the associated VF from issuing any memory or I/O requests. Note that as MSI/MSI-X interrupt messages are in-band memory writes, disabling the bus master enable bit disables MSI/MSI-X interrupt messages as well. Requests other than memory or I/O requests are not controlled by this bit. Note: The state of active transactions is not specified when this bit is disabled after being enabled. The device can choose how it behaves when this condition occurs. Software cannot count on the device retaining state and resuming without loss of data when the bit is re-enabled. Transactions for a VF that has its <i>Bus Master Enable</i> set must not be blocked by transactions for VFs that have their <i>Bus Master Enable</i> cleared.
3	0b	RO	SCM: Special Cycle Enable. Hard wired to 0b.
4	0b	RO	MWIE: MWI Enable. Hard wired to 0b.
5	0b	RO	PSE: Palette Snoop Enable. Hard wired to 0b.
6	0b	RO	PER: Parity Error Response. Zero for VFs.
7	0b	RO	WCE: Wait Cycle Enable. Hard wired to 0b.
8	0b	RO	SERRE: SERR# Enable. Zero for VFs.
9	0b	RO	FB2BE: Fast Back-to-Back Enable. Hard wired to 0b.
10	0b	RO	INTD: Interrupt Disable. Hard wired to 0b.
15:11	0b	RO	RSV: Reserved.

9.3.1.2 VF Status Register (0x6; RW)

Bits	Initial Value	Rd/Wr	Description
2:0	0x0	RO	RSV: Reserved.
3	0b	RO	IS: Interrupt Status. Hardwired to 0b.
4	1b	RO	NC: New Capabilities. Indicates that the Intel® Xeon® Processor D-1500 Product Family LAN controller VFs implement extended capabilities. The Intel® Xeon® Processor D-1500 Product Family LAN controller VFs implement a capabilities list, to indicate that it supports MSI-X and PCIe extensions.
5	0b	RO	66E: 66 MHz Capable. Hardwired to 0b.
6	0b	RO	RSV: Reserved.
7	0b	RO	FB2BC: Fast Back-to-Back Capable. Hardwired to 0b.
8	0b	RW1C	MPERR: Data Parity Reported.
10:9	00b	RO	DEVSEL: DEVSEL Timing. Hardwired to 0b.
11	0b	RW1C	STA: Signaled Target Abort.
12	0b	RW1C	RTA: Received Target Abort.
13	0b	RW1C	RMA: Received Master Abort.
14	0b	RW1C	SSERR: Signaled System Error.
15	0b	RW1C	DSERR: Detected Parity Error.



9.3.2 PCI Capabilities

9.3.2.1 MSI-X Capability

The only registers with a different layout than the PF for MSI-X, is the control register.

Note: The message address and data registers in enhanced mode use the first MSI-X entry of each VF in the regular MSI-X table.

9.3.2.1.1 VF MSI-X Control Register (0x72; RW)

Bits	Initial Value	Rd/Wr	Description
10:0	0x002 ¹	RO	TS: Table Size.
13:11	0x0	RO	RSV: Reserved.
14	0b	RW	Mask: Function Mask.
15	0b	RW	En: MSI-X Enable.

1. Default value is read from the NVM. This field is loaded from the *PCI_CNF2.MSI_X_VF_N* register field.

9.3.2.1.2 MSI-X Table Offset Register (0x74; RW)

Bits	Default	RW	Description
31:3	0x000	RO	Table Offset. Used as an offset from the address contained in one of the function's BARs to point to the base of the MSI-X table. The lower three <i>Table BIR</i> bits are masked off (set to 0b) by software to form a 32-bit Qword-aligned offset. Note that this field is read only.
2:0	0x3	RO	Table BIR. Indicates which one of a function's BARs, beginning at 0x10 in the configuration space, is used to map the function's MSI-X table into the memory space. BIR values: 0...5 correspond to BARs 0x10...0x24, respectively. BIR equals 3 to indicate BAR 3 and 4 are used for MSI-X vectors.

9.3.2.1.3 MSI-X PBA Register (0x78; RO)

Bits	Default	Type	Description
31:3	0x400	RO	PBA Offset. Used as an offset from the address contained by one of the function's BARs to point to the base of the MSI-X PBA. The lower three PBA BIR bits are masked off (set to zero) by software to form a 32-bit Qword-aligned offset. This value is changed by hardware to be half of the value programmed to the IOV System Page Size register.
2:0	0x3	RO	PBA BIR. Indicates which one of a function's BARs, located beginning at 0x10 in configuration space, is used to map the function's MSI-X PBA into memory space. BIR values: 0...5 correspond to BARs 0x10...0x24, respectively. BIR equals 3 to indicate BAR 3 and 4 are used for MSI-X PBA



9.3.2.2 PCIe Capability Registers

The device control and device status registers have some fields which are specific per VF.

9.3.2.2.1 VF Device Control Register (0xA8; RW)

Bits	Rd/Wr	Default	Description
0	RO	0b	Correctable Error Reporting Enable. Zero for VFs.
1	RO	0b	Non-Fatal Error Reporting Enable. Zero for VFs.
2	RO	0b	Fatal Error Reporting Enable. Zero for VFs.
3	RO	0b	Unsupported Request Reporting Enable. Zero for VFs.
4	RO	0b	Enable Relaxed Ordering. Zero for VFs.
7:5	RO	0b	Max Payload Size. Zero for VFs.
8	RO	0b	Not implemented in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
9	RO	0b	Not implemented in the Intel® Xeon® Processor D-1500 Product Family LAN controller.
10	RO	0b	Auxiliary Power PM Enable. Zero for VFs.
11	RO	0b	Enable No Snoop. Zero for VFs.
14:12	RO	000b	Max Read Request Size. Zero for VFs.
15	RW	0b	Initiate Function Level Reset. Specific to each VF.

9.3.2.2.2 VF Device Status Register (0xAA; RW1C)

Bits	Rd/Wr	Default	Description
0	RW1C	0b	Correctable Detected. Indicates status of correctable error detection. Zero for VF.
1	RW1C	0b	Non-Fatal Error Detected. Indicates status of non-fatal error detection. Zero for VF.
2	RW1C	0b	Fatal Error Detected. Indicates status of fatal error detection. Zero for VF.
3	RW1C	0b	Unsupported Request Detected. Indicates that the Intel® Xeon® Processor D-1500 Product Family LAN controller received an unsupported request. This field is identical in all functions. The Intel® Xeon® Processor D-1500 Product Family LAN controller can't distinguish which function caused an error. Zero for VF.
4	RO	0b	Aux Power Detected. Zero for VFs.
5	RO	0b	Transaction Pending. Specific per VF. When set, indicates that a particular function (PF or VF) has issued non-posted requests that have not been completed. A function reports this bit cleared only when all completions for any outstanding non-posted requests have been received.
15:6	RO	0x00	Reserved.

9.3.3 PCIe Extended Capabilities

9.3.3.1 AER Registers

The following registers in the AER capability have a different behavior in a VF function.

Unlike the PF AER registers, these registers are not sticky since the VF is reset on FLR and on in-band reset.




9.3.3.1.1 Uncorrectable Error Status Register (0x104; RW1C)

Bit Location	Attribute	Default Value	Description
3:0	RO	0x0	Reserved.
4	RO	0b	Data Link Protocol Error Status.
5	RO	0b	Surprise Down Error Status (Optional).
11:6	RO	0x0	Reserved.
12	RW1C	0b	Poisoned TLP Status.
13	RO	0b	Flow Control Protocol Error Status.
14	RW1C	0b	Completion Timeout Status.
15	RW1C	0b	Completer Abort Status.
16	RW1C	0b	Unexpected Completion Status.
17	RO	0b	Receiver Overflow Status.
18	RO	0b	Malformed TLP Status.
19	RO	0b	ECRC Error Status.
20	RW1C	0b	Unsupported Request Error Status — when caused by a function that claims a TLP.
21	RO	0b	ACS Violation Status.
31:21	RO	0x0	Reserved.

9.3.3.1.2 Correctable Error Status Register (0x110; RW1C)

The Correctable Error Status register reports error status of individual correctable error sources on a PCIe device. When an individual error status bit is set to 1b it indicates that a particular error occurred; software can clear an error status by writing a 1b to the respective bit.

Bit Location	Attribute	Default Value	Description
0	RW1C	0b	Receiver Error Status.
5:1	RO	0x0	Reserved.
6	RW1C	0b	Bad TLP Status.
7	RW1C	0b	Bad DLLP Status.
8	RW1C	0b	REPLAY_NUM Rollover Status.
11:9	RO	0x0	Reserved.
12	RW1C	0b	Replay Timer Timeout Status.
13	RW1C	0b	Advisory Non-Fatal Error Status.
31:14	RO	0b	Reserved.



NOTE: *This page intentionally left blank.*



10.0 System Manageability

Network management is an important requirement in today's networked computer environment. Software-based management applications provide the ability to administer systems while the operating system is functioning in a normal power state (not in a pre-boot state or powered-down state). The Intel® Out of Band Management fill the management void that exists when the operating system is not running or fully functional. This is accomplished by providing mechanisms by which manageability network traffic can be routed to and from a Management Controller (BMC).

This section describes the supported management interfaces and hardware configurations for platform system management. It describes the interfaces to an external BMC, the partitioning of platform manageability among system components, and the functionality provided by the Intel® Xeon® Processor D-1500 Product Family LAN controller in each platform configuration.

10.1 Pass-Through (PT) Functionality

Pass-Through (PT) is the term used when referring to the process of sending and receiving Ethernet traffic over the sideband interface. The Intel® Xeon® Processor D-1500 Product Family LAN controller has the ability to route Ethernet traffic to the host operating system as well as the ability to send Ethernet traffic over the sideband interface to an external BMC. See [Figure 10.1](#).

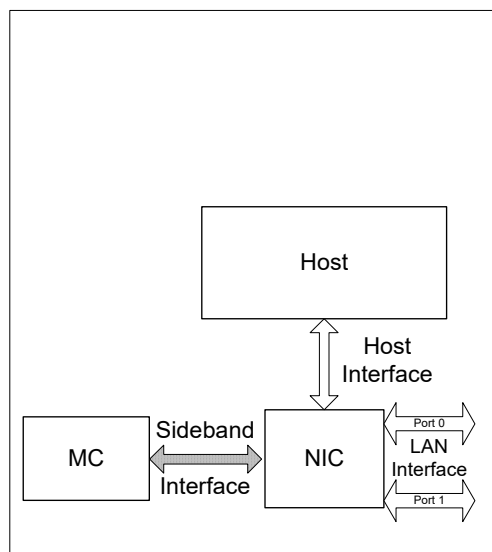



Figure 10.1. Sideband Interface



The sideband interface provides a mechanism by which the Intel® Xeon® Processor D-1500 Product Family LAN controller can be shared between the host and the BMC. By providing this sideband interface, the BMC can communicate with the LAN without requiring a dedicated Ethernet controller. The Intel® Xeon® Processor D-1500 Product Family LAN controller supports three sideband interfaces:

- SMBus
- NC-SI

The usable bandwidth for either direction is up to 1 Mb/s when using SMBus and 100 Mb/s for the NC-SI interface. When working over the integrated I/O interface, the bandwidth is limited only by the link's bandwidth and the the Intel® Xeon® Processor D-1500 Product Family LAN controller processing capabilities and can sustain any network bandwidth. Only one mode of sideband can be active at any given time. The configuration is done using an NVM setting.

The maximal packet size supported for traffic received from the LAN to the BMC is 1518 byte and additional VLAN or E-tag. For traffic from the BMC to the LAN the maximal supported packet size is 1536 bytes including all tags.

Note: In MCTP mode, SMBus interface can receive MCTP commands. For example, the MCTP enumeration process can be done over SMBus. However, only the SMBus can receive NC-SI commands or pass-through traffic.

10.1.1 Supported Topologies

The Intel® Xeon® Processor D-1500 Product Family LAN controller can support some management topologies. The following connections are available:

- Single connection via legacy SMBus (See [Section 10.5](#)).
- Single connection via NC-SI over RMII (See [Section 10.6](#))
- Single connection via NC-SI over MCTP. This connection can be over the SMBus. This connection can be used either for pass through or for control only (see [Section 10.7](#)).

The topology used is defined in the *Redirection Sideband Interface* field in the Common Firmware Parameters 1 NVM word and is common to all the ports in the Intel® Xeon® Processor D-1500 Product Family LAN controller.

10.1.2 Pass Through Packet Routing

When an Ethernet packet reaches the Intel® Xeon® Processor D-1500 Product Family LAN controller, it is examined and compared to a number of configurable filters. These filters are configurable by the BMC and include, but not limited to, filtering on:

- MAC Address
- IP Address
- UDP/IP Ports
- VLAN Tags
- EtherType

If the incoming packet matches any of the configured filters, it is passed to the BMC. Otherwise it is not passed.

The packet filtering process is described in [Section 10.3](#).



10.2 Components of the Sideband Interface

There are two components to a sideband interface:

- Physical Layer
- Logical Layer

10.2.1 Physical Layer

This is the electrical connection between the Intel® Xeon® Processor D-1500 Product Family LAN controller and BMC.

10.2.1.1 SMBus

The SMBus physical layer is defined by the SMBus specification. The interface is made up of two connections: data and clock. There is also an optional third connection: the alert line. This line is used by the Intel® Xeon® Processor D-1500 Product Family LAN controller to notify the BMC that there is data available for reading. Refer to the SMBus specification for details.

The SMBus can run at three speeds: 100 KHz (standard SMBus), or 400 KHz (I²C fast mode) or 1 MHz (I²C fast mode plus). The speed used is selected by the *SMBus Connection Speed* in *SMBus Notification Timeout and Flags* NVM word.

10.2.1.1.1 PEC Support

SMBus transactions can be protected by using Packet Error Code (PEC). Packet Error Checking, whenever applicable, is implemented by appending a PEC byte at the end of each message transfer. The PEC byte is a CRC8 calculation on all the message bytes.

PEC is added in transmit and expected in receive for the following SMBus packets:

- ARP packets
- MCTP over SMBus transactions.

For ARA cycles and legacy SMBus transactions, a PEC is not expected.

The following table lists the behavior of the Intel® Xeon® Processor D-1500 Product Family LAN controller in each PEC configured mode for transactions directly handled by the hardware upon reception of packets with or without PEC.

Table 10.1. SMBus PEC Modes¹

Transaction/Function		Target PEC Mode	
SMBus Transaction (Relative to the Broadwell-DE LAN controller)	Broadwell-DE LAN controller PEC Mode	PEC Enabled	PEC Disabled
Master Write ²	Enabled	(A) Target ACKs the PEC byte.	(A) Target NACKs the PEC byte.
Master Write ²	Disabled	(A) Target receives stop before expected PEC byte.	(A) PEC byte is not expected.
Slave Write ³	Enabled	(A) Target ACKs last data byte; PEC byte is NACKed.	(A) Target NACKs last data byte; No PEC byte is written by a slave.



Table 10.1. SMBus PEC Modes¹

Transaction/Function		Target PEC Mode	
SMBus Transaction (Relative to the Broadwell-DE LAN controller)	Broadwell-DE LAN controller PEC Mode	PEC Enabled	PEC Disabled
Slave Write ³	Disabled	(A) Target ACKs last data byte; PEC byte is 0xFF.	(A) Target NACKs last data byte and generates a stop after that.
Slave Read ⁴	Enabled	(A) Target sends PEC byte; PEC byte is ACKed by the slave.	(A) Target does not send PEC byte and generates a stop after that.
Slave Read ⁴	Disabled	(R) Target sends PEC byte; PEC byte is NACKed by the slave.	(A) Target does not send PEC byte and generates a stop after that.

1. (A) - Accept Transaction (R) - Reject Transaction.
2. Used in Legacy SMBus write commands (direct receive) and in MCTP over SMBus (transmitted transactions).
3. Used in Legacy SMBus Read commands.
4. Used in Legacy SMBus mode (alert/async-notify) and in MCTP over SMBus (received transactions).

Note: In both SMBus ARP and MCTP, the specification indicates that PEC must be used. However, if PEC is not used by the master, the transaction is still accepted and processed by the Intel® Xeon® Processor D-1500 Product Family LAN controller.

The PEC behavior is controlled by the *SMBus transaction PEC* bit in the SMBus Notification Timeout and Flags NVM word: If this bit is set, PEC is added for master SMBus write transactions. a PEC is added to slave read transactions and can be received in slave write transaction. If this bit is cleared, PEC is not added to master write or slave read transactions, a slave write transaction with PEC is dropped. This bit should be set for MCTP mode and should be cleared in legacy SMBus mode.

10.2.1.2 NC-SI


The Intel® Xeon® Processor D-1500 Product Family LAN controller uses the DMTF standard sideband interface. This interface consists of six lines for transmission and reception of Ethernet packets and two optional lines for arbitration among more than one physical network controller.

The physical layer of NC-SI is very similar to the RMI interface, although not an exact duplicate. Refer to the NC-SI specification for details of the differences.

10.2.2 Logical Layer

10.2.2.1 Legacy SMBus

The protocol layer for SMBus consists of commands the BMC issues to configure filtering for the Intel® Xeon® Processor D-1500 Product Family LAN controller management traffic and the reading and writing of Ethernet frames over the SMBus interface. There is no industry standard protocol for sideband traffic over SMBus. The protocol layer for SMBus on the Intel® Xeon® Processor D-1500 Product Family LAN controller is Intel proprietary. The Legacy SMBus protocol is described in [Section 10.5](#).



10.2.2.2 NC-SI

The DMTF also defines the protocol layer for the NC-SI interface. NC-SI compliant devices are required to implement a minimum set of commands. The specification also provides a mechanism for vendors to add additional capabilities through the use of OEM commands. Intel OEM NC-SI commands for the Intel® Xeon® Processor D-1500 Product Family LAN controller are discussed in [Section 10.6.3](#). For information on base NC-SI commands, see the NC-SI specification.

NC-SI traffic can run on top of three different physical layers:

1. NC-SI physical layer as described in [Section 10.2.1.2](#).
2. MCTP over SMBus. As previously described, this layer supports Sx modes.

The Intel® Xeon® Processor D-1500 Product Family LAN controller exposes one NC-SI package with two channels, one per port. The Intel® Xeon® Processor D-1500 Product Family LAN controller implement a type C NC-SI interface (Single package, common bus buffers and shared RX queue) as described in section 5.2 of the NC-SI specification.

10.2.2.2.1 Package ID setting

The package ID can be set either from the NVM *Package ID* field in the NC-SI Configuration 1 NVM word or from an SDP pin. If set from SDP, the Package ID is {0,SDP0_0 value, 0}. The mode used is set by the *NC-SI Package ID from SDP* field in the NC-SI Configuration 2 NVM word. Note that when the package ID is set from the SDP pin, SDP0_0 should be set as input is *ESDP.SDP0_IODIR* field.

The internal channel ID matches the lowest numbered PCIe function number through which this port is exposed to the host.

10.2.2.2.2 Channel ID mapping

The mapping of the channels to physical ports is according to the NC-SI Channel to Port Mapping NVM word if the *NC-SI Channel to Port Mapping Table Valid* bit is set. If this bit is not set, the following algorithm should be used:

```
CHANNEL_ID = 0

If (FACTPS.LAN_FUNCTION_SEL == 0 ) // no swap
    For (x = 0 to 1) {
        PORT = x;
        If (FACTPS.LANx_VALID) NC-SI_Channel[PORT] = CHANNEL_ID++;
    }
Else // swap
    For (x = 0 to 1) {
        PORT = 1 - x;
        If (FACTPS.LANx_VALID) NC-SI_Channel[PORT] = CHANNEL_ID++;
    }
```



10.3 Packet Filtering

Since both the host operating system and BMC use the Intel® Xeon® Processor D-1500 Product Family LAN controller to send and receive Ethernet traffic, there needs to be a mechanism by which incoming Ethernet packets can be identified as those that should be sent to the BMC rather than the host operating system.

There are two different types of filtering available. The first is filtering based upon the MAC address. With this filtering, the BMC has at least one dedicated MAC address and incoming Ethernet traffic with the matching MAC address(es) are passed to the BMC. This is the simplest filtering mechanism to utilize and it allows an BMC to receive all types traffic (including, but not limited to, IPMI, NFS, HTTP etc).

The other mechanism available utilizes a highly configurable mechanism by which packets can be filtered using a wide range of parameters. Using this method, the BMC can share a MAC address (and IP address, if desired) with the host OS and receive only specific Ethernet traffic. This method is useful if the BMC is only interested in specific traffic, such as IPMI packets.

10.3.1 Manageability Receive Filtering

This section describes the manageability receive packet filtering flow. Packet reception by the Intel® Xeon® Processor D-1500 Product Family LAN controller can generate one of the following results:

- Discarded
- Sent to host memory
- Sent to the external BMC
- Sent to both the BMC and host memory

The decisions regarding forwarding of packets to the host and to the BMC are separate and are configured through two sets of registers. However, the BMC might define some types of traffic as exclusive. This traffic is forwarded only to the BMC, even if it passes the filtering process of the host. These types of traffic are defined using the MNGONLY register.

An example of packets that might be necessary to send exclusively to the BMC might be specific TCP/UDP ports of a shared MAC address or a MAC address dedicated to the BMC. If the BMC configures the manageability filters to send these ports to the BMC, it should configure the settings to not send them to the host, otherwise, these ports will be received and handled by the host operating system.

The BMC controls the types of packets that it receives by programming receive manageability filters. The following filters are accessible to the BMC:

Table 10.2. Filters Accessible to BMC (Sheet 1 of 2)

Filters	Functionality	When Reset?
Filters Enable	General configuration of the manageability filters.	
Manageability Only	Enables routing of packets exclusively to the manageability.	
Manageability Decision Filters [7:0]	Configuration of manageability decision filters.	Power Good Reset.
MAC Address [3:0]	Four unicast MAC manageability addresses.	Power Good Reset and Firmware Reset.
VLAN Filters [7:0]	Eight VLAN tag values.	Power Good Reset and Firmware Reset.

**Table 10.2. Filters Accessible to BMC (Sheet 2 of 2)**

Filters	Functionality	When Reset?
UDP/TCP Port Filters [15:0]	16 destination port values.	Power Good Reset and Firmware Reset
Flexible 128 bytes TCO Filter	Length and values for one flex TCO filter.	Power Good Reset and Firmware Reset
IPv4 and IPv6 Address Filters [3:0]	IP address for manageability filtering.	Power Good Reset and Firmware Reset

All filtering capabilities are available on both the NC-SI and legacy SMBus interfaces. However, in NC-SI mode, in order to program part of the capabilities, the Intel OEM commands described in [Section 10.6.3](#) should be used.

All filters are reset only on Internal Power On Reset. Register filters that enable filters or functionality are also reset by firmware reset in NC-SI mode. These registers can be loaded from the NVM following a reset in SMBus mode.

The high-level structure of manageability filtering is done using two steps.

1. The packet is parsed and fields in the header are compared to programmed filters.
2. A set of decision filters are applied to the result of the first step.

Some general rules apply:

- Fragmented packets are passed to manageability but not parsed beyond the IP header.
- Packets with L2 errors (CRC, alignment, etc.) are not forwarded to the BMC.
- Packets longer than 2KB are filtered out.
- The filtering relates only to the outer L3/L4 header. In tunneled packets, the inner L2, L3 and L4 header parameters are not considered for manageability filtering.

The following sections describe the manageability filtering, followed by the final filtering rules.

The filtering rules are created by programming the decision filters as described in [Section 10.3.4](#).

10.3.2 L2 Filters

10.3.2.1 MAC and VLAN Filters

The manageability MAC filters allow comparison of the destination MAC address to one of 4 filters defined in the MMAH and MMAL registers.


The VLAN filters allow comparison of the 12 bit VLAN tag to one of 8 filters defined in the MAVTV registers.

10.3.2.2 EtherType Filters

Manageability L2 EtherType filters enable filtering of received packets based on the Layer 2 EtherType field. The L2 type field of incoming packets is compared against the EtherType filters programmed in the Manageability EtherType Filter (METF; up to 4 filters); the result is incorporated into decision filters.

Each manageability EtherType filter can be configured as pass (positive) or reject (negative) using a polarity bit. In order for the reverse polarity mode to be effective and block certain type of packets, the EtherType filter should be part of all the enabled decision filters.

An examples for usage of L2 EtherType filters is to determine the destination of 802.1X control packets. The 802.1X protocol is executed at different times in either the management controller or by the host. L2 EtherType filters are used to route these packets to the proper agent.



In addition to the flexible EtherType filters, the Intel® Xeon® Processor D-1500 Product Family LAN controller supports 2 fixed EtherType filters used to block NC-SI control traffic (0x88F8) and flow control traffic (0x8808) from reaching the manageability interface. The NC-SI EtherType is used for communication between the management controller on the NC-SI link and the Intel® Xeon® Processor D-1500 Product Family LAN controller. Packets coming from the network are not expected to carry this EtherType and such packets are blocked to prevent attacks on the management controller. Flow control packets should be consumed by the MAC and as such are not expected to be forwarded to the management interface.

10.3.3 L3/L4 Filtering

The manageability filtering stage combines checks done at previous stages with additional L3/L4 checks to make a the decision on whether to route a packet to the BMC. The following sections describe the manageability filtering done at layers L3/L4 and final filtering rules.

Note: The L3 filters will not match tunneled fields (L3, L4 or ARP).

10.3.3.1 ARP Filtering

ARP filtering — The Intel® Xeon® Processor D-1500 Product Family LAN controller supports filtering of ARP request packets (initiated externally) and ARP responses (to requests initiated by the BMC).

In legacy SMBus mode, the ARP filters can be used as part of the ARP off load described in [Section 10.5.4](#). ARP off load is not specifically available when using NC-SI. However, the general filtering mechanism is utilized to filter incoming ARP traffic as requested using the Enable Broadcast Filtering NC-SI command.

In order to limit the reception of ARP packets to the ARP packets dedicated to this station (ARP target IP = BMC IP), the ARP request/response filter can be bind to specific IP address, by setting both the ARP request/response and the IP AND bits in an MDEF filter, as the IP bit is set also if there is a match on the target IP (the *TPA* field in the ARP packet) of an ARP request or an ARP response.

Note: If the OR section of the MDEF is all cleared and one of the IPv4 address are set, then ARP packets matching the IP address passes the filter. If these packets should be dropped, then an OR Ethertype filter with a value of 0x0800 (IPv4) should be added.

10.3.3.2 Neighbor Discovery Filtering and MLD


The Intel® Xeon® Processor D-1500 Product Family LAN controller supports filtering of the following ICMPv6 packets.

Neighbor discovery packets:

1. 0x86 (134d) - Router Advertisement.
2. 0x87 (135d) - Neighbor Solicitation.
3. 0x88 (136d) - Neighbor Advertisement.
4. 0x89 (137d) - Redirect.

MLD packets:

1. 0x82 (130d) - MLD Query
2. 0x83 (131d) - MLDv1 Report
3. 0x84 (132d) - MLD Done
4. 0x8F (143d) - MLDv2 Report



The neighbor discovery packets has dedicated enables for each type in the decision filters. For MLD, a single enable controls the forwarding of all the MLD packets. This means that either all the MLD packet types are selected for reception or none of them.

10.3.3.3 RMCP Filtering

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports filtering by fixed destination port numbers, port 0x26F and port 0x298. These ports are IANA reserved for RMCP.

In SMBus mode, there are filters that can be enabled for these ports. When using NC-SI, they are not specifically available. However, the general filtering mechanism can be utilized to filter incoming ARP traffic.

10.3.3.4 Flexible Port Filtering

The Intel® Xeon® Processor D-1500 Product Family LAN controller implements 16 flex destination port filters. The Intel® Xeon® Processor D-1500 Product Family LAN controller directs packets whose L4 destination port matches to the BMC. The BMC must ensure that only valid entries are enabled in the decision filters.

10.3.3.5 IP Address Filtering

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports filtering by destination IP address using IPv4 and IPv6 address filters. These are dedicated to manageability. The Intel® Xeon® Processor D-1500 Product Family LAN controller provides four IPv6 address filters or three IPv6 addresses and four IPv4 address filters.

The IPv4 match also rises for ARP packets for which the target IP matches the IP address in the MIPAF4 register.

10.3.3.6 Checksum Filtering

If bit `MANC.EN_XSUM_FILTER` is set, the Intel® Xeon® Processor D-1500 Product Family LAN controller directs packets to the BMC only if they match all other filters previously described as well as pass L3/L4 checksum (if it exists).

The Intel® Xeon® Processor D-1500 Product Family LAN controller be instructed to direct packets to the BMC only if they pass the L3/L4 checksum (if they exist) in addition to matching other filters previously described.

Enabling the XSUM filter when using the SMBus interface is accomplished by setting the *Enable XSUM Filtering to Manageability* bit. This is done using the Update Management Receive Filter Parameters command. See [Section 10.5.11.1.6](#).

To enable the XSUM filtering when using NC-SI, use the Enable Checksum Offloading command. See [Section 10.6.3.15](#).

10.3.4 Flexible 128-byte Filter

The Intel® Xeon® Processor D-1500 Product Family LAN controller provides one flex TCO filter. This filter looks for a pattern match within the first 128 bytes of the packet.

Flex filters are temporarily disabled when read from or written to by the host. Any packet received during a read or write operation is dropped. Filter operation resumes once the read or write access completes.



10.3.4.1 Flexible Filter Structure

The filter is composed of the following fields:

1. Flexible Filter length — This field indicates the number of bytes in the packet header that should be inspected. The field also indicates the minimal length of packets inspected by the filter. Packet below that length will not be inspected. Valid values for this field are: $8*n$, where $n=1...16$.
2. Data — This is a set of up to 128 bytes comprised of values that header bytes of packets are tested against.
3. Mask — This is a set of 128 bits corresponding to the 128 data bytes that indicate for each corresponding byte if is tested against its corresponding byte. The general filter is 128 bytes that the BMC configures; all of these bytes may not be needed or used for the filtering, so the mask is used to indicate which of the 128 bytes are used for the filter.

Each filter tests the first 128 bytes (or less) of a packet, where not all bytes must necessarily be tested.

10.3.4.2 TCO Filter Programming

Programming each filter is done using the following commands (NC-SI or SMBus) in a sequential manner:

1. Filter Mask and Length — This command configures the following fields:
 - a. Mask — A set of 16 bytes containing the 128 bits of the mask. Bit 0 of the first byte corresponds to the first byte on the wire.
 - b. Length — A 1-byte field indicating the length.
2. Filter Data — The filter data is divided into groups of bytes. as follows:

Group	Test Bytes
0x0	0-29
0x1	30-59
0x2	60-89
0x3	90-119
0x4	120-127


Each group of bytes need to be configured using a separate command, where the group number is given as a parameter. The command has the following parameters:

- a. Group number — A 1-byte field indicating the current group addressed
- b. Data bytes — Up to 30 bytes of test-bytes for the current group

10.3.5 Configuring Manageability Filters

There are a number of pre-defined filters that are available for the BMC to enable, such as ARPs and IPMI ports 0x298/0x26F. These are generally enabled by setting the appropriate bit within the MANC register using specific commands.

For more advanced filtering needs, the BMC has the ability to configure a number of configurable filters. It is a two-step process to use these filters. They must first be configured and then enabled.



10.3.5.1 Manageability Decision Filters

Manageability Decision Filters (MDEF) are a set of eight filters, each with the same structure. The filtering rule for each decision filter is programmed by the BMC and defines which of the L2, VLAN, EtherType and L2/L3 filters participate in decision making. Any packet that passes at least one rule is directed to manageability and possibly to the host.

The inputs to each decision filter are:

- Packet passed a valid management L2 exact address filter.
- Packet is a broadcast packet.
- Packet has a VLAN header and it passed a valid manageability VLAN filter.
- Packet matched one of the valid IPv4 or IPv6 manageability address filters.
- Packet is a multicast packet.
- Packet passed ARP filtering (request or response).
- Packet passed neighbor solicitation filtering.
- Packet passed MLD filtering.
- Packet passed 0x298/0x26F port filter.
- Packet passed a valid flex port filter.
- Packet passed a valid flex TCO filter.
- Packet passed or failed an L2 EtherType filter.
- Packet passed or failed flow control or NC-SI L2 EtherType discard filter.

The structure of each decision filter is shown in [Figure 10.2](#). A boxed number indicates that the input is conditioned by a mask bit defined in the MDEF register and MDEF_EXT register for this rule. Decision filter rules are as follows:

- At least one bit must be set in a register. If all bits are cleared (MDEF/MDEF_EXT = 0x0000), then the decision filter is disabled and ignored.
- All enabled AND filters must match for the decision filter to match. An AND filter not enabled in the MDEF/MDEF_EXT registers is ignored. If an AND filter is preceded by a OR filter, then at least one of the enabled OR inputs must match for the filter to pass.
- If no OR filter is enabled in the register, the OR filters are ignored in the decision (the filter might still match).
- If one or more OR filters are enabled in the register, then at least one of the enabled OR filters must match for the decision filter to match.

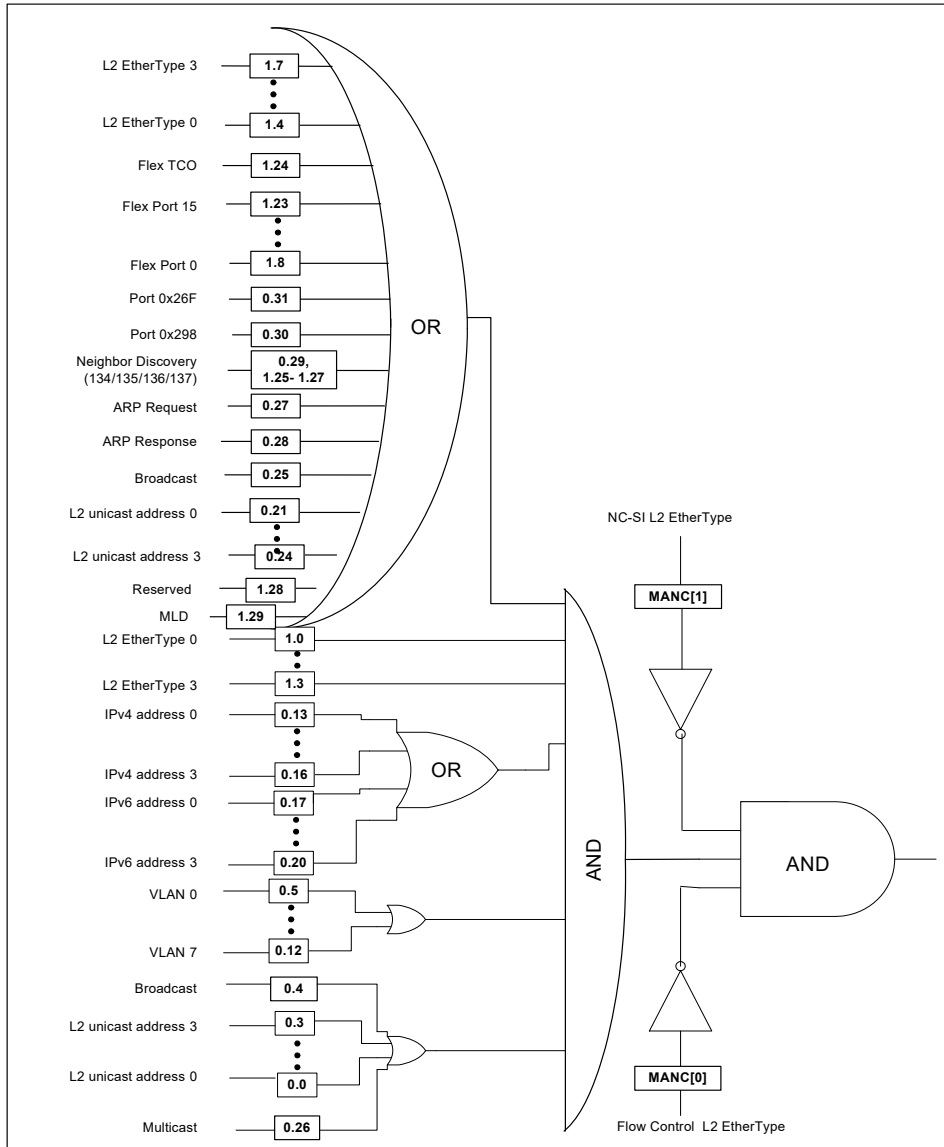


Figure 10.2. Manageability Decision Filters

A decision filter (for any of the eight filters) defines which of the above inputs is enabled as part of a filtering rule. The BMC programs two 32-bit registers per rule (MDEF[7:0] and MDEF_EXT[7:0]) with the settings as described in [Section 8.2.2.15.6](#) and [Section 8.2.2.15.7](#). A set bit enables its corresponding filter to participate in the filtering decision.



In addition to the controls previously described, the *MDEF_EXT.apply_to_host_traffic* and *MDEF_EXT.apply_to_network_traffic* bits defines which traffic is compared to this filter. At least one of these bits must be set for the filter to be valid.

If the *MDEF_EXT.apply_to_host_traffic* bit is set, the traffic from the host is candidate for this filter. If the *MDEF_EXT.apply_to_network_traffic* bit is set, the traffic from the network is candidate for this filter. If both bits are set, this filter is applied to all traffic.

10.3.5.2 Exclusive Traffic

The decisions regarding forwarding of packets to the host for LAN traffic or to the LAN for host traffic are independent from the management decision filters. However, the BMC might define some types of traffic as exclusive. The behavior for such traffic is defined by the using the bits corresponding to the decision filter in the *MNGONLY* register (one bit per each of the eight decision rules) and the *MDEF_EXT.apply_to_host_traffic* and *MDEF_EXT.apply_to_network_traffic* bits. [Table 10.3](#) lists the behavior in each case. If one or more filters match the traffic and at least one of the filters is set as exclusive, the traffic is treated as exclusive.


Table 10.3. Exclusive traffic behavior

	Filter Match		Filter Don't Match
traffic source	<i>MNGONLY</i> = 0	<i>MNGONLY</i> = 1	N/A
From network	Traffic is forwarded to the manageability. Traffic is forwarded to the host according to host filtering	Traffic is forwarded only to manageability.	Traffic is forwarded to the host according to host filtering.
From host	Traffic is forwarded to the manageability and to the LAN	Traffic is forwarded only to manageability.	Traffic is forwarded to the LAN.

Any traffic matching any of the configurable filters (see [Section 10.3.5.1](#)) can be used as filters to pass traffic to the host.

Table 10.4. *MNGONLY* Register Description and Usage

Bits	Description	Default
0	Decision Filter 0	Determines if packets that have passed decision filter 0 are sent exclusively to the manageability path.
1	Decision Filter 1	Determines if packets that have passed decision filter 1 are sent exclusively to the manageability path.
2	Decision Filter 2	Determines if packets that have passed decision filter 2 are sent exclusively to the manageability path.
3	Decision Filter 3	Determines if packets that have passed decision filter 3 are sent exclusively to the manageability path.
4	Decision Filter 4	Determines if packets that have passed decision filter 4 are sent exclusively to the manageability path.
5	Unicast and Mixed	NC-SI mode: Determines if unicast and mixed packets are sent exclusively to the manageability path. SMBus mode: Determines if packets that have passed decision filter 5 are sent exclusively to the manageability path.
6	Global Multicast	NC-SI mode: Determines if multicast packets are sent exclusively to the manageability path. SMBus mode: Determines if packets that have passed decision filter 6 are sent exclusively to the manageability path.
7	Broadcast	NC-SI mode: Determines if broadcast packets are sent exclusively to the manageability path. SMBus mode: Determines if ARP packets are sent exclusively to the manageability path.
31:8	Reserved	Reserved.



When using the SMBus interface, the BMC enables these filters by issuing the Update Management Receive Filter Parameters command (see [Section 10.5.11.1.6](#)) with the parameter of 0x0F.

The MNGONLY is also configurable when using NC-SI using the Set Intel Filters — Manageability Only Command (see [Section 10.6.3.7.2](#)).

All manageability filters are controlled by the BMC only and not by the LAN device driver.

10.3.5.3 Global Controls

On top of the MDEF filters, the MANC register contains some global controls applied to all the packets in order to be candidate for manageability filtering:


- Receive Enable bits:
 - The *RCV_TCO_EN* field controls the reception of manageability traffic. It should be set only if one of the following bits is set also.
 - The *EN_BMC2OS* bit controls the reception of manageability traffic from the host.
 - The *EN_BMC2NET* bit controls the reception of manageability traffic from the network.
- VLAN filtering: In order to support the NC-SI VLAN modes the following controls are provided:
 - The *FIXED_NET_TYPE* field controls if only VLAN tagged or VLAN un-tagged traffic is received. If this bit is cleared both types are received. If it is set, only the type described by the *NET_TYPE* field is accepted.
 - If set, the *NET_TYPE* field indicates that only VLAN tagged traffic is received, if cleared only packets without VLAN is accepted. This field is validated by the *FIXED_NET_TYPE* field.
 Both fields relates to the inner VLAN.

10.3.6 Filtering Programming Interfaces

The Intel® Xeon® Processor D-1500 Product Family LAN controller provides multiple options to program the forwarding filters, depending on the interface used and the level of flexibility needed. The following table lists the different options and points to the description of the relevant commands.

Table 10.5. Filtering Programming Interfaces

Interface	Flexible/ Abstract	Description
NC-SI (over RMIID or over MCTP)	Abstract (dedicated MAC address)	The regular NC-SI commands can be used to allow forwarding based on a dedicated MAC address. The list of supported commands can be found in Section 10.6.2.1 . When using these commands, one of the two other modes can be used to add finer grain filtering.
	Flexible	This interface described in most of the subsections of Section 10.6.3.4 . It uses the packet reduction commands to reduce the forwarding scope of the filters set by the regular NC-SI commands and the packet addition commands to add new packet types to the forwarding rules.
SMBus	Abstract	The Set Common filter command (Section 10.5.11.1.7) can be used to set the most common filters. When using this commands the flexible filtering interface should not be used. When sending this command, all previous filtering requests are cleared.
	Flexible	The Update MNG RCV Filter Parameters commands (Section 10.5.11.1.6) can be used to define the exact filtering rules to be applied.



10.3.7 Possible Configurations

This section describes ways of using management filters. Actual usage might vary.

10.3.7.1 Dedicated MAC Packet Filtering

- Select one of the eight rules for dedicated MAC filtering.
- Load host MAC address to one of the management MAC address filters and set the appropriate bit in field 3:0 of the MDEF register.
- Set other bits to qualify which packets are allowed to pass through. For example:
 - Set bit 5 in MDEF to qualify with the first manageability VLAN.
 - Set relevant bits 13 to 20 in MDEF to qualify with a match to one of the IP addresses.
 - Set any L3/L4 bits (bits 27 to 31 in MDEF and bits 16 to 23 in MDEF_EXT) to qualify with any of a set of L3/L4 filters.

10.3.7.2 Broadcast Packet Filtering

- Select one of the eight rules for broadcast filtering.
- Set bit 25 in MDEF of the decision rule to enforce broadcast filtering.
- Set other bits to qualify which broadcast packets are allowed to pass through. For example:
 - Set bit 5 in MDEF to qualify with the first manageability VLAN.
 - Set relevant bits 13 to 20 in MDEF to qualify with a match to one of the IP addresses.
 - Set any L3/L4 bits (bits 27 to 31 in MDEF and bits 16 to 23 in MDEF_EXT) to qualify with any of a set of L3/L4 filters.


10.3.7.3 VLAN Packet Filtering

- Select one of the eight rules for VLAN filtering.
- Set bit 5 to 12 in MDEF to qualify with the relevant manageability VLANs.
- Set other bits to qualify which VLAN packets are allowed to pass through. For example:
 - Set any L3/L4 bits (bits 27 to 31 in MDEF and bits 16 to 23 in MDEF_EXT) to qualify with any of a set of L3/L4 filters.

10.3.7.4 IPv6 Filtering

IPv6 filtering is done using the following IPv6-specific filters:

- IP Unicast filtering — requires filtering for Link Local address and a Global address. Filtering setup might depend on whether or not the MAC address is shared with the host or dedicated to manageability:
 - Dedicated MAC address (for example, dynamic address allocation with DHCP does not support multiple IP addresses for one MAC address). In this case, filtering can be done at L2 using two dedicated unicast MAC filters.
 - Shared MAC address (for example, static address allocation sharing addresses with host). In this case, filtering needs to be done at L3, requiring two IPv6 address filters, one per address.
- A neighbor Discovery filter — The Intel® Xeon® Processor D-1500 Product Family LAN controller supports IPv6 neighbor Discovery protocol. Since the protocol relies on multicast packets, the Intel® Xeon® Processor D-1500 Product Family LAN controller supports filtering of these packets. IPv6 multicast addresses are translated into corresponding Ethernet multicast addresses in the form of 33-33-xx-xx-xx-xx, where the last 32 bits of address are taken from the last 32 bits of the



IPv6 multicast address. As a result, two direct MAC filters can be used to filter IPv6 solicited-node multicast packets as well as IPv6 all node multicast packets.

10.3.7.5 Receive Filtering with Shared IP

When using the Legacy SMBus interface or the MCTP interface, it is possible to share the host MAC and IP address with the BMC. This functionality is also available when using base NC-SI using Intel OEM commands.

When the BMC shares the MAC and IP address with the host, receive filtering is based on identifying specific flows through port allocation. The following setting might be used when using the legacy SMBus interface:

- Select one of the eight rules.
- Set a manageability dedicated MAC filter to the host MAC address and set the matching bit (0-3) in the MDEF register.
- If VLAN is used for management, load one or more management VLAN filters and set the matching bit (5- 12) in the MDEF register.

ARP filter/neighbor discovery filter is enabled when the BMC is responsible for handling the ARP protocol. Set bit 27 or bit 28 in the MDEF register for this functionality.

In NC-SI over MCTP, dedicated commands are used to allow shared IP filtering.

10.3.8 Determining Manageability MAC Address

If the BMC needs to use a dedicated MAC address or configure the automatic ARP response mechanism (only available in SMBus mode), it might be beneficial for the BMC to be able to determine the MAC address used by the host.

Both the NC-SI and SMBus interfaces provide an Intel OEM command to read the System MAC address.

A possible use for this is that the MAC address programmed at manufacturing time does not increment by one each time, but rather by two. In this way, the BMC can read the system MAC address and add one to it and be guaranteed of a unique MAC address.


Determining the IP address being used by the host is beyond the scope of this document.

10.4 OS-to-BMC Traffic

10.4.1 Overview

Traditionally, the communication between a host and the local BMC is not handled through the network interface and requires a dedicated interface such as an IPMI KCS interface. The Intel® Xeon® Processor D-1500 Product Family LAN controller enables the host and the local BMC communication via the regular pass-through interface, and thus enable management of a local console using the same interface used to manage any BMC in the network.

When this flow is used, the host sends packets to the BMC through the network interface. The Intel® Xeon® Processor D-1500 Product Family LAN controller examines these packets and it then decides if they should be forwarded to the BMC. On the inverse path, when the BMC sends a packet on the pass-through interface, the Intel® Xeon® Processor D-1500 Product Family LAN controller checks if it should be forwarded to the network, the host, or both. [Figure 10.3](#) describes the flow for OS-to-BMC traffic for the NC-SI over RMII case. It is not supported in legacy SMBus mode.



The OS-to-BMC flow can be enabled using the *OS2BMC enable* field for the relevant port in the OS 2 BMC configuration structure of the NVM.

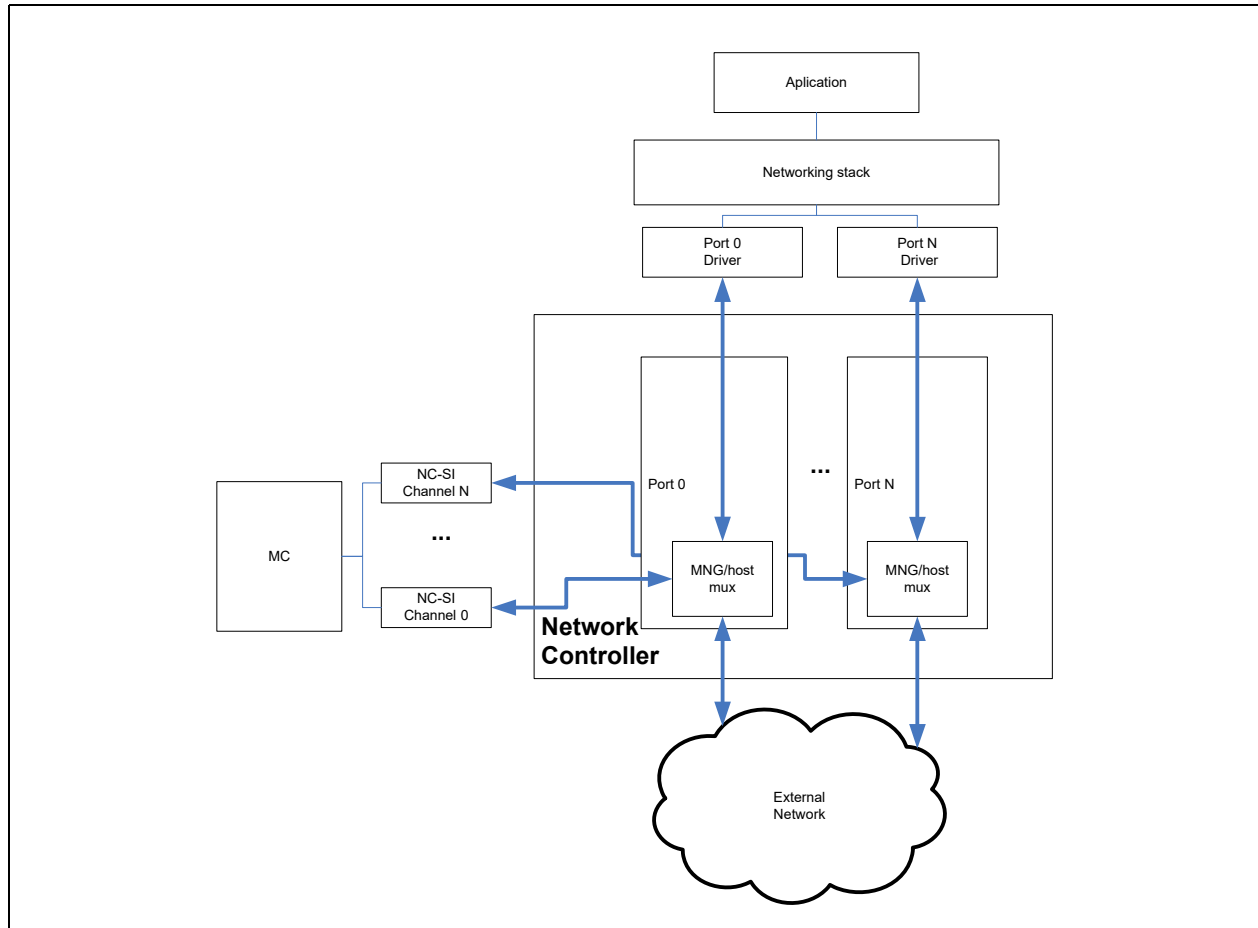


Figure 10.3. OS-to-BMC Block Diagram

The OS-to-BMC flow is enabled only for ports enabled by the NC-SI Enable Channel command or via the *OS to BMC Enable* field for the relevant port in the OS-to-BMC configuration structure of the NVM.

OS-to-BMC traffic must comply with NC-SI specifications and is therefore limited to maximum sized frames of 1536 bytes (in both directions).

10.4.2 Filtering

When OS-to-BMC traffic is enabled, the filters used for network to BMC traffic are also used for OS-to-BMC traffic. Traffic considered as exclusive to the BMC (Relevant bit in MNGONLY is set) is also considered as exclusive to the BMC when sent from the Host and not forwarded to the network.

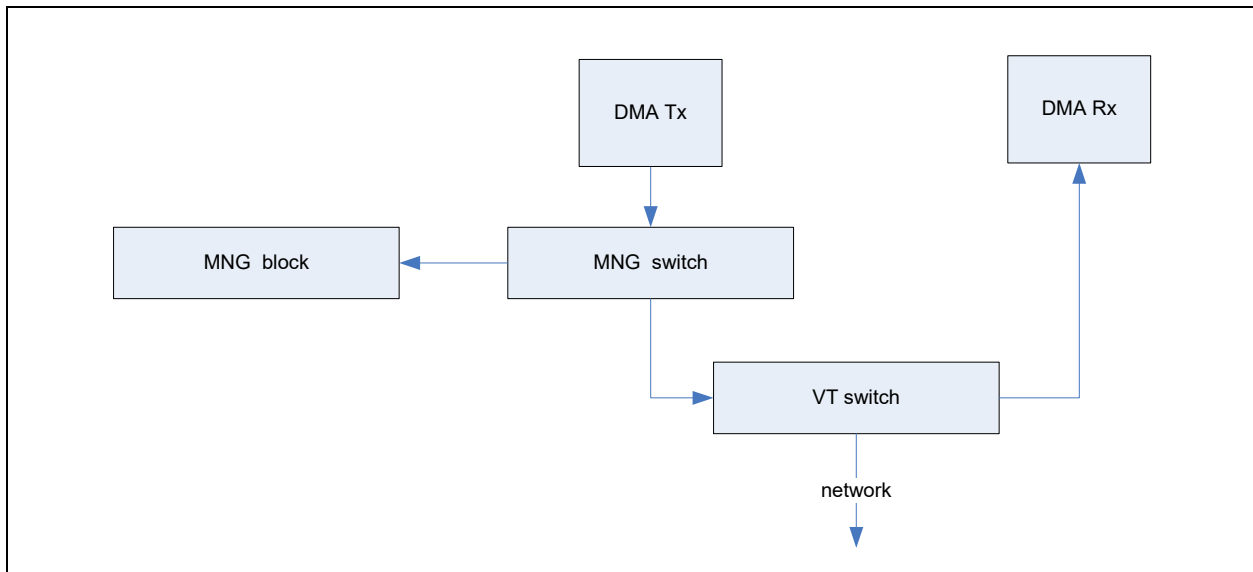


Figure 10.4. OS-to-BMC and VM-to-VM Filtering

10.4.2.1 Handling of OS-to-BMC Packets

All the regular transmit offloads are also available for OS-to-BMC packets.

10.4.2.2 BMC-to-OS Filtering

When OS-to-BMC is enabled, as with regular BMC transmit traffic, the port (operating system or network) to which the packet is sent is fixed according to the source MAC address of the packet. After that, the BMC traffic is filtered according to the L2 host filters of the selected port (as described in [Section 7.1.3](#)). According to the results of the filtering the packet can be forwarded to the operating system, the network or both.


The following rules apply to the forwarding of operating system packets:

- If *BMC to net* is disabled, all the traffic from the BMC is sent to the host.
- If *BMC to host* is disabled, all the traffic from the BMC is sent to the network.
- If both *BMC to net* and *BMC to host* are enabled, the packets are forwarded only according to the destination MAC address and VLAN tag. Unicast packets that match one of the exact filters (RAH/RAL) are sent only to the host. Other packets that pass the L2 host filtering are sent to both the host and the network. Packets that do not pass the L2 host filtering are sent only to the network.

Note: In virtualization mode, if a packet is received by host only due to default pooling (*PFVTCTL.Dis_Def_Pool* bit is cleared), then it is not sent to the host (even if *BMC to net* is disabled).

10.4.2.3 Queuing of Packets Received From the BMC

Packets received from the BMC are queued in the default queue.



10.4.2.4 Offloads of Packets Received from the BMC

Packets received from the BMC and forwarded to the operating system do not pass the same path as regular network packets. Thus parts of the offloads provided for the network packets are not available for the BMC packets. Packet received from the BMC are identified by the *RDESC.STATUS.BMC* bit.

The following list describes which offloads are available for BMC packets:

- CRC is checked and removed on the BMC packets.
- The BMC packets are not detected as time sync packet. The *RDESC.STATUS.TS* always clears for these packets.
- In systems where the double VLAN feature is enabled (*CTRL_EXT.EXTENDED_VLAN* is set), the VEXT bit is valid for BMC packets and reflects the presence of a tag with an Ethertype matching the value in *EXVET* register.

Note: In systems that uses double VLAN, the BMC is expected to send all packets (apart from NC-SI commands) with the outer VLAN included. Failing to do so might cause corruptions to the packet received by the operating system.

- The *RDESC.ERRORS* field is always cleared for these packets.

Note: Traffic sent from the BMC does not cause a PME event, even if it matches one of the wake-up filters set by the port.

10.4.3 Blocking of Network-to-BMC Flow

In some systems the BMC might have its own private connection to the network and might use the Intel® Xeon® Processor D-1500 Product Family LAN controller port only for the OS-to-BMC traffic. In this case, the BMC-to-network flow should be blocked while enabling the OS-to-BMC and OS-to-network flows.

This can be done by clearing the *MANC.EN_BMC2NET* bit for the relevant port. The BMC can control this functionality using the Enable Network to BMC flow and Disable Network to BMC flow NC-SI OEM commands. This can also be controlled using the *Network to BMC disable* field in the NVM OS2BMC Configuration Structure.

Notes: When network to BMC flow is blocked and OS-to-BMC flow is enabled, all the traffic from the BMC is sent to the operating system without any check. The operating system traffic filtering is still done using the regular decision filters.

The NC-SI channel should not be enabled for receive or transmit before at least one of the *EN_BMC2NET* or *EN_BMC2OS* fields is set, unless used for AEN transmissions only. In this case, the channel might be enabled for receive, but all receive filters should be cleared.

10.4.4 OS-to-BMC and Flow Control

The traffic between the host and manageability uses the same buffers as any loopback traffic. Thus it flows through the transmit buffer and then through the receive buffer. If the transmit buffer is flow controlled, then the host to BMC traffic is also stopped. If the receive buffer is full, the traffic is dropped or the transmit is stopped according to the flow control policy of this traffic class.

Packets received to the manageability (either from host or from network) might be dropped if the manageability internal buffers are full.



10.4.5 Statistics

Packets sent from the OS to the BMC should be counted by all statistical counters as packets sent by the operating system. If they are sent to both the network and to the BMC, then they are counted once.

Packets sent from the BMC to the host are counted as packets received by the host. If they are sent to the host and to the network, then they are counted both as received packets and as packet transmitted to the network.

In addition, the Intel® Xeon® Processor D-1500 Product Family LAN controller supports the following statistical counters that measure just the BMC-to-OS and OS-to-BMC traffic:

- O2BGPTC: OS-to-BMC packets received by BMC
- O2BSPC: OS-to-BMC packets transmitted by operating system and received by manageability buffer.
- B2OSPC: BMC-to-OS packets sent by BMC
- B2OGPRC: BMC-to-OS packets received by the operating system.

The software device driver can use these statistics to count packets dropped by the Intel® Xeon® Processor D-1500 Product Family LAN controller during the transfer between the operating system and the BMC or the LAN and the BMC as follows:

- Dropped packets in BMC-to-OS path = $B2OSPC - B2OGPRC$
- Dropped packet in BMC-to-LAN path = $B2OSDPC - (B2OSPC - B2OGPRC)$
- Dropped packets in OS-to-BMC path = $O2BSPC - O2BGPTC$
- Dropped packets in LAN-to-BMC path = $MNGPDC$

See [Section 7.1.7.2](#) and [Section 7.2.6.2](#) for details of the statistics hierarchy.

10.4.6 OS-to-BMC Enablement

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports the unified network software model for OS-to-BMC traffic, where the OS-to-BMC traffic is shared with the regular traffic. In this model, there is no need for a special configuration of the operating system networking stack or the BMC stack, but if the link is down, then the OS-to-BMC communication is stopped.


In order to enable OS-to-BMC either:

- Enable OS2BMC in the port traffic type field in the Traffic type Parameters NVM word for the relevant port.
- Send an Enable Network-to-BMC Command

When OS2BMC is enabled, the operating system must avoid sending packets longer than 1.5 KB to the BMC.

10.5 SMBus Pass-Through Interface

SMBus is the system management bus defined by Intel. It is used in personal computers and servers for low-speed system management communications. This section describes how the SMBus interface operates in legacy pass-through mode.



10.5.1 General

The SMBus sideband interface includes standard SMBus commands used for assigning a slave address and gathering device information as well as Intel proprietary commands used specifically for the pass-through interface.

10.5.2 Pass-Through Capabilities

This section details manageability capabilities the Intel® Xeon® Processor D-1500 Product Family LAN controller provides while in SMBus mode. Pass-through traffic is carried by the sideband interface as described in [Section 10.1](#).

These services are not available in NC-SI mode.

When operating in SMBus mode, in addition to exposing a communication channel to the LAN for the BMC, the Intel® Xeon® Processor D-1500 Product Family LAN controller provides the following manageability services to the BMC:

- ARP handling — The Intel® Xeon® Processor D-1500 Product Family LAN controller can be programmed to auto-ARP replying for ARP request packets to reduce the traffic over the BMC interconnect.
- Default configuration of filters by NVM - When working in SMBus mode, the default values of the manageability receive filters can be set according to the PT LAN and flex TCO NVM structures.

10.5.3 Port to SMBus Mapping

The Intel® Xeon® Processor D-1500 Product Family LAN controller is presented on the SMBus manageability link as two different devices (for example, via two different SMBus addresses on which each device is connected to a different LAN port). There is no logical connection between the devices.

The fail-over between the LAN ports is done by the BMC (by sending/receiving packets through different devices). The status report to the BMC, ARP handling, DHCP, and other pass-through functionality are unique for each port and configured by the BMC.

10.5.4 Automatic Ethernet ARP Operation

The Intel® Xeon® Processor D-1500 Product Family LAN controller can offload the Ethernet Address Resolution Protocol (ARP) for the BMC in order to reduce the bandwidth required on the SMBus link.

Automatic Ethernet ARP parameters are loaded from the NVM when the Intel® Xeon® Processor D-1500 Product Family LAN controller is powered up or configured through the sideband management interface. The following parameters should be configured in order to enable ARP operation:

- ARP auto-reply enabled
- ARP IP address (to filter ARP packets)
- ARP MAC addresses (for ARP responses)

These are all configurable over the sideband interface using the advanced version of the Receive Enable command.

When an ARP request packet is received and ARP auto-reply is enabled, the Intel® Xeon® Processor D-1500 Product Family LAN controller checks the targeted IP address (after the packet has passed L2 checks and ARP checks). If the targeted IP matches the IP configuration for the Intel® Xeon® Processor D-1500 Product Family LAN controller, it replies with an ARP response.



The Intel® Xeon® Processor D-1500 Product Family LAN controller responds to ARP request targeted to the ARP IP address with the configured ARP MAC address. In case that there is no match, the Intel® Xeon® Processor D-1500 Product Family LAN controller silently discards the packets. If the Intel® Xeon® Processor D-1500 Product Family LAN controller is not configured to do auto-ARP response, it can be configured to forward the ARP packets to the BMC (which can respond to ARP requests).

When the external BMC uses the same IP and MAC address of the OS, the ARP operation should be coordinated with the host operating system.

Note: If sharing the MAC and IP with the host operating system is possible, the Intel® Xeon® Processor D-1500 Product Family LAN controller provides the ability to read the stem MAC address, allowing the BMC to share the MAC address. There is no mechanism however provided by the Intel® Xeon® Processor D-1500 Product Family LAN controller to read the IP address. The host operating system (or an agent within) and BMC must coordinate the sharing of IP addresses.

10.5.5 SMBus Transactions

This section gives a brief overview of the SMBus protocol. Following is an example for a format of a typical SMBus transaction.

Table 10.6. Typical SMBus Transaction

1	7	1	1	8	1	8	1	1
S	Slave Address	Wr	A	Command	A	PEC	A	P
	1100 001	0	0	0000 0010	0	[Data Dependent]	0	


The top row of the table identifies the bit length of the field in a decimal bit count. The middle row (bordered) identifies the name of the fields used in the transaction. The last row appears only with some transactions, and lists the value expected for the corresponding field. This value can be either hexadecimal or binary.

The SMBus controller is a master for some transactions and a slave for others. The differences are identified in this document.

Shorthand field names are listed in [Table 10.7](#) and are fully defined in the SMBus specification.

Table 10.7. Shorthand Field Names

Field Name	Definition
S	SMBus START Symbol
P	SMBus STOP Symbol
PEC	Packet Error Code
A	ACK (Acknowledge)
N	NACK (Not Acknowledge)
Rd	Read Operation (Read Value = 1b)
Wr	Write Operation (Write Value = 0b)



10.5.5.1 SMBus Addressing

The SMBus is presented as up to two SMBus devices on the SMBus (two addresses). All pass-through functionality is duplicated on the SMBus address, where each SMBus address is connected to a different LAN port. Note that it is not permitted to configure multiple ports to the same SMBus address. When a LAN function is disabled, the corresponding SMBus address is not presented to the BMC.

SMBus addresses (enabled from the NVM) can be re-assigned using the SMBus ARP protocol.

In addition to the SMBus address values, all parameters of the SMBus (channel selection, address mode, and address enable) can be set only through NVM configuration. Note that the NVM is read at the Intel® Xeon® Processor D-1500 Product Family LAN controller's power up and resets.

10.5.5.2 SMBus ARP Functionality

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports the SMBus ARP protocol as defined in the SMBus 2.0 specification. The Intel® Xeon® Processor D-1500 Product Family LAN controller is a persistent slave address device so its SMBus address is valid after power-up and loaded from the NVM. The Intel® Xeon® Processor D-1500 Product Family LAN controller supports all SMBus ARP commands defined in the SMBus specification both general and directed.

SMBus ARP capability can be disabled through the NVM.

10.5.5.3 SMBus ARP Flow

SMBus ARP flow is based on the status of two flags:

- AV (Address Valid): This flag is set when the Intel® Xeon® Processor D-1500 Product Family LAN controller has a valid SMBus address.
- AR (Address Resolved): This flag is set when the Intel® Xeon® Processor D-1500 Product Family LAN controller SMBus address is resolved (SMBus address was assigned by the SMBus ARP process).


These flags are internal Intel® Xeon® Processor D-1500 Product Family LAN Controller flags and are not exposed to external SMBus devices.

Since the Intel® Xeon® Processor D-1500 Product Family LAN controller is a Persistent SMBus Address (PSA) device, the AV flag is always set, while the AR flag is cleared after power up until the SMBus ARP process completes. Since AV is always set, the Intel® Xeon® Processor D-1500 Product Family LAN controller always has a valid SMBus address.

When the SMBus master needs to start an SMBus ARP process, it resets (in terms of ARP functionality) all devices on SMBus by issuing either Prepare to ARP or Reset Device commands. When the Intel® Xeon® Processor D-1500 Product Family LAN controller accepts one of these commands, it clears its AR flag (if set from previous SMBus ARP process), but not its AV flag (the current SMBus address remains valid until the end of the SMBus ARP process).

Clearing the AR flag means that the Intel® Xeon® Processor D-1500 Product Family LAN controller responds to SMBus ARP transactions that are issued by the master. The SMBus master issues a Get UDID command (general or directed) to identify the devices on the SMBus. The Intel® Xeon® Processor D-1500 Product Family LAN controller always responds to the Directed command and to the General command only if its AR flag is not set.

After the Get UDID, The master assigns the Intel® Xeon® Processor D-1500 Product Family LAN controller SMBus address by issuing an Assign Address command. The Intel® Xeon® Processor D-1500 Product Family LAN controller checks whether the UDID matches its own UDID and if it matches, it



switches its SMBus address to the address assigned by the command (byte 17). After accepting the Assign Address command, the AR flag is set and from this point (as long as the AR flag is set), the Intel® Xeon® Processor D-1500 Product Family LAN controller does not respond to the Get UDID General command. Note that all other commands are processed even if the AR flag is set. If the address changed, from the one previously stored in the NVM, The Intel® Xeon® Processor D-1500 Product Family LAN controller stores the SMBus address that was assigned in the SMBus ARP process in the NVM, so at the next power up, it returns to its assigned SMBus address. This process uses the NVM update flow described in [Section 3.4.2.1](#).

[Figure 10.5](#) shows the Intel® Xeon® Processor D-1500 Product Family LAN controller SMBus ARP flow.

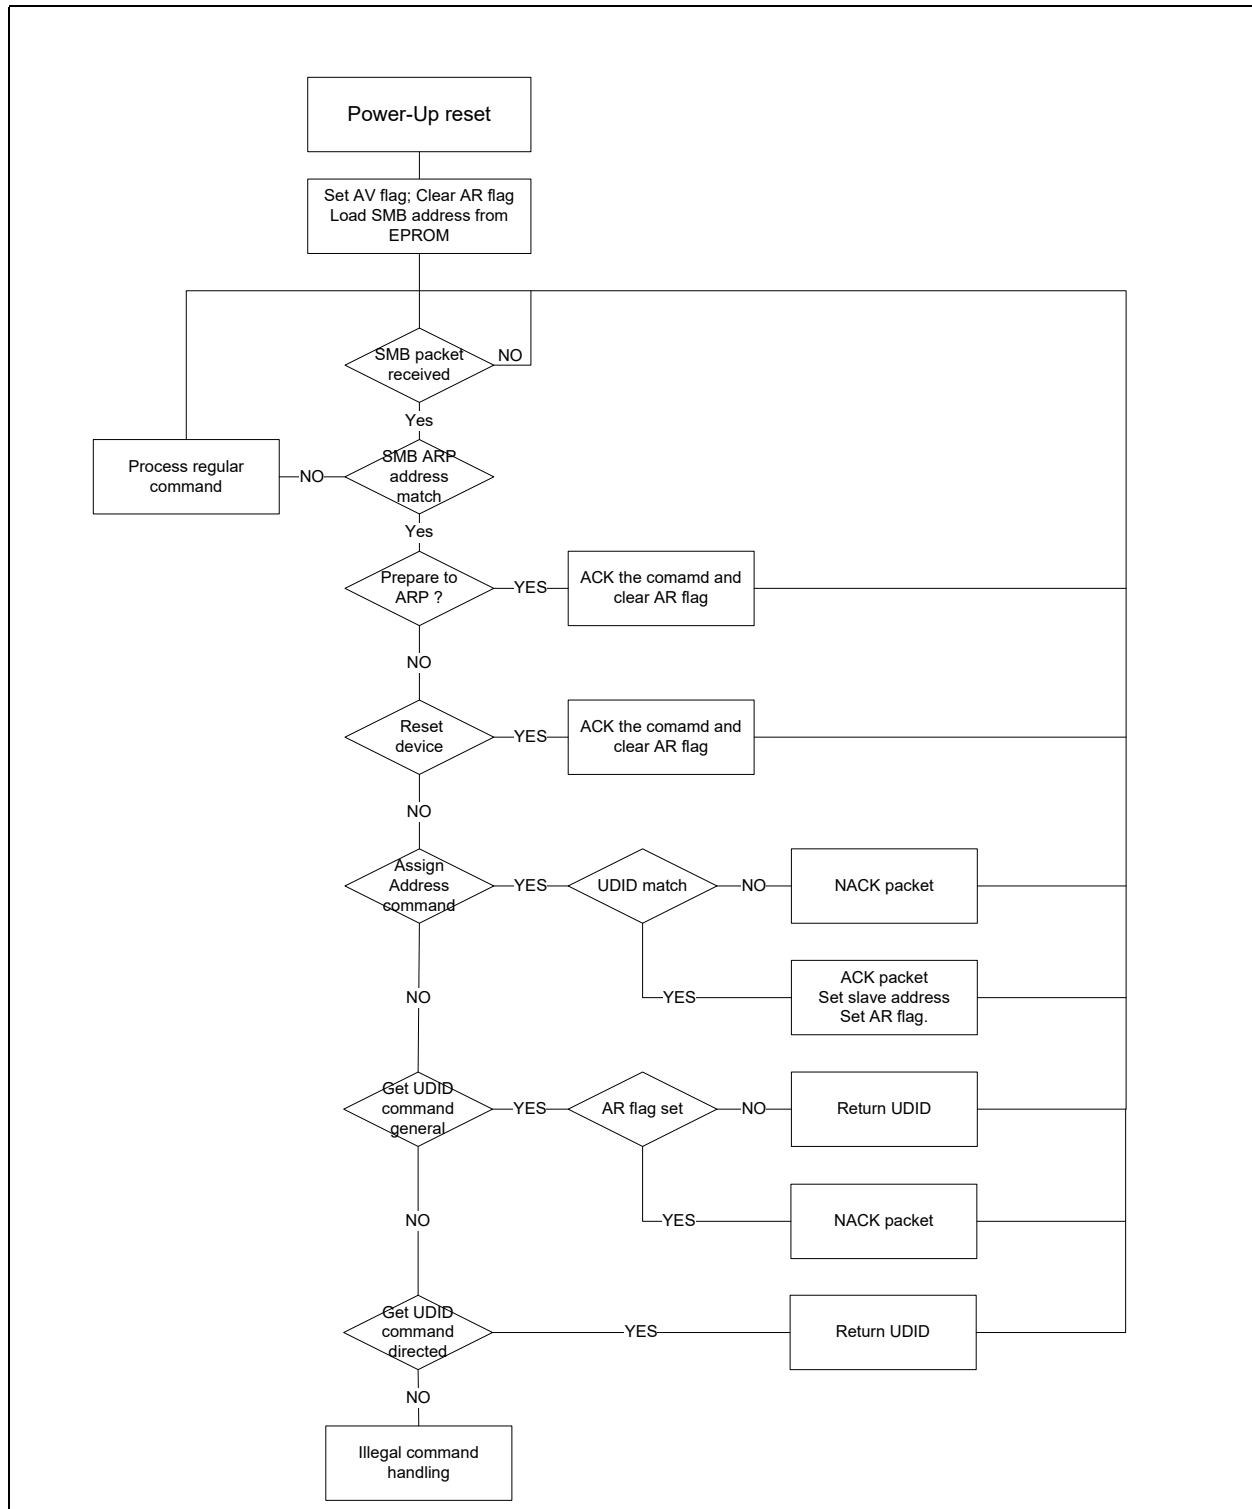



Figure 10.5. SMBus ARP Flow



10.5.5.4 SMBus ARP UDID Content

The UDID provides a mechanism to isolate each device for the purpose of address assignment. Each device has a unique identifier. The 128-bit number is comprised of the following fields:

Table 10.8. UDID

1 Byte	1 Byte	2 Bytes	2 Bytes	2 Bytes	2 Bytes	2 Bytes	4 Bytes
Device Capabilities	Version/Revision	Vendor ID	Device ID	Interface	Subsystem Vendor ID	Subsystem Device ID	Vendor Specific ID
See notes that follow	See notes that follow	0x8086	See notes that follow	0x0004/ 0x0024	0x0000	0x0000	See notes that follow
MSB							LSB

Where:

- Vendor ID: The device manufacturer’s ID as assigned by the SBS Implementers’ Forum or the PCI SIG.
Constant value: 0x8086
- Device ID: The device ID as assigned by the device manufacturer (identified by the Vendor ID field).
Constant value: see the defined silicon default value in [Section 9.2.2.2](#).
- Interface: Identifies the protocol layer interfaces supported over the SMBus connection by the device.
Bits 3:0 = 0x4 indicates SMBus Version 2.0
Bit 5 (ASF bit) = 1 in MCTP mode.
- Subsystem Fields: These fields are not supported and return zeros.

Device Capabilities: Dynamic and Persistent Address, PEC Support bit:

7	6	5	4	3	2	1	0
Address Type		Reserved (0)	Reserved (0)	Reserved (0)	Reserved (0)	Reserved (0)	PEC Supported
0b	1b	0b	0b	0b	0b	0b	0/1b ¹
MSB							LSB

1. The value is set according to the *SMBus Transaction PEC* bit in the NVM.

Version/Revision: UDID Version 1, Silicon Revision:

7	6	5	4	3	2	1	0
Reserved (0)	Reserved (0)	UDID Version			Silicon Revision ID See Section 9.2.2.5 .		
0b	0b	001b					
MSB							LSB



Vendor Specific ID: Four LSB bytes of the device Serial Number combined with the port number. The Serial Number is taken from the NVM and is reflected in the PCI_SERL and PCI_SERH registers.

1 Byte	1 Byte	1 Byte	1 Byte
Port Number, MAC Address, Byte 3[6:0]	MAC Address, Byte 2	MAC Address, Byte 1	MAC Address, Byte 0
MSB			LSB

10.5.5.5 SMBus ARP and Multi-port

The Intel® Xeon® Processor D-1500 Product Family LAN controller responds as two SMBus devices having two sets of AR/AV flags (one for each port). The Intel® Xeon® Processor D-1500 Product Family LAN controller responds two time to the SMBus ARP master, once each for each port. All SMBus addresses are taken from the SMBus ARP address word of the NVM.

Note that the Unique Device Identifier (UDID) is different for the different ports in the version ID field. The Intel® Xeon® Processor D-1500 Product Family LAN controller first respond as port 0, and only when an address is assigned, then start responding as port 1 to the Get UDID command.

10.5.5.6 Concurrent SMBus Transactions

The SMBus interface is single threaded. Thus, concurrent SMBus transactions are not permitted. Once a transaction is started, it must be completed before additional transaction can be initiated.

A transaction is defined as:

- All the SMBus commands used to receive a packet.
- All the SMBus commands used to send a packet.
- The read and write SMBus commands used as part of read parameters described in [Section 10.5.11.2](#).
- The single write SMBus commands described in [Section 10.5.11.1](#).

10.5.6 SMBus Notification Methods

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports three methods of notifying the BMC that it has information that needs to be read by the BMC:

- SMBus alert - Refer to [Section 10.5.6.1](#).
- Asynchronous notify - Refer to [Section 10.5.6.2](#).
- Direct receive - refer to [Section 10.5.6.3](#).

The notification method used by the Intel® Xeon® Processor D-1500 Product Family LAN controller can be configured from the SMBus using the Receive Enable command. The default method is set by the NVM in the Notification Method field in LAN Receive Enable 1.

Note: The SMBus notification method used must be the same for all ports.

The following events cause the Intel® Xeon® Processor D-1500 Product Family LAN controller to send a notification event to the BMC:

- Receiving a LAN packet that is designated to the BMC.
- Firmware was reset and requires re-initialization.
- Receiving a Request Status command from the BMC initiates a status response.



- The Intel® Xeon® Processor D-1500 Product Family LAN controller is configured to notify the BMC upon status changes (by setting the EN_STA bit in the Receive Enable command) and one of the following events happen:
 - TCO Command Aborted
 - Link Status changed
 - Power state change

There can be cases where the BMC is hung and not responding to the SMBus notification. The Intel® Xeon® Processor D-1500 Product Family LAN controller has a time-out value (defined in the NVM) to avoid hanging while waiting for the notification response. If the BMC does not respond until the time out expires, the notification is de-asserted and all pending data is silently discarded.

Note that the SMBus notification time-out value can only be set in the NVM. The BMC cannot modify this value.

10.5.6.1 SMBus Alert and Alert Response Method

The SMBus Alert# (SMBALERT_N) signal is an additional SMBus signal that acts as an asynchronous interrupt signal to an external SMBus master. The Intel® Xeon® Processor D-1500 Product Family LAN controller asserts this signal each time it has a message that it needs the BMC to read and if the chosen notification method is the SMBus alert method. Note that the SMBus alert method is an open-drain signal which means that other devices besides the Intel® Xeon® Processor D-1500 Product Family LAN controller can be connected on the same alert pin. As a result, the BMC needs a mechanism to distinguish between the alert sources.

The BMC can respond to the alert by issuing an ARA Cycle command to detect the alert source device. The Intel® Xeon® Processor D-1500 Product Family LAN controller responds to the ARA cycle with its own SMBus slave address (if it was the SMBus alert source) and de-asserts the alert when the ARA cycle completes. Following the ARA cycle, the BMC issues a read command to retrieve the Intel® Xeon® Processor D-1500 Product Family LAN controller message.

Some BMCs do not implement the ARA cycle transaction. These BMCs respond to an alert by issuing a Read command to the Intel® Xeon® Processor D-1500 Product Family LAN controller (0xC0/0xD0 or 0xDE). The Intel® Xeon® Processor D-1500 Product Family LAN controller always responds to a Read command, even if it is not the source of the notification. The default response is a status transaction. If the Intel® Xeon® Processor D-1500 Product Family LAN controller is the source of the SMBus Alert, it replies the read transaction and then de-asserts the alert after the command byte of the read transaction.

Note: In SMBus Alert mode, the SMBALERT_N pin is used for notification. Each port generate alerts on events that are independent of each other. If two ports have events to notify, the second alert is asserted only after the first event is handled.

The ARA cycle is an SMBus receive byte transaction to SMBus Address 0001-100b. Note that the ARA transaction does not support PEC. The ARA transaction format is as follows:

1	7	1	1	8	1	1	1
S	Alert Response Address	Rd	A	Slave Device Address		A	P
	0001 100	1	0	Manageability Slave SMBus Address	0	1	



10.5.6.2 Asynchronous Notify Method

When configured using the asynchronous notify method, the Intel® Xeon® Processor D-1500 Product Family LAN controller acts as a SMBus master and notifies the BMC of one of the events listed in Section 10.5.6 by issuing a modified form of the write word transaction. The asynchronous notify transaction SMBus address and data payload is configured using the Receive Enable command (Section 10.5.11.1.3) or using the NVM defaults. Note that the asynchronous notify is not protected by a PEC byte.

1	7	1	1	7	1	1	8	1	8	1	1
S	Target Address	Wr	A	Sending Device Address		A	Data Byte Low	A	Data Byte High	A	P
	BMC Slave Address	0	0	MNG Slave SMBus Address	0	0	Interface	0	Alert Value	0	

The target address and data byte low/high are taken from the Receive Enable command or NVM configuration.

If the BMC does not read the status in the delay defined by the *SMBus Notification Timeout* NVM field, the status word is cleared and the packet is dropped.

10.5.6.3 Direct Receive Method

If configured, the Intel® Xeon® Processor D-1500 Product Family LAN controller has the capability to send a message it needs to transfer to the external BMC as a master over the SMBus instead of alerting the BMC and waiting for it to read the message.

The message format follows. Note that the command that is used is the same command that is used by the external BMC in the Block Read command. The opcode that the Intel® Xeon® Processor D-1500 Product Family LAN controller puts in the data is also the same as it put in the Block Read command of the same functionality. The rules for the *F* and *L* flags (bits) are also the same as in the Block Read command.

1	7	1	1	1	1	6	1	
S	Target Address	Wr	A	F	L	Command	A	...
	BMC Slave Address	0	0	First Flag	Last Flag	Receive TCO Command 01 0000b	0	

8	1	8	1		1	8	1	1
Byte Count	A	Data Byte 1	A	...	A	Data Byte N	A	P
N	0		0		0		0	

10.5.7 Receive Pass Through Flow

The Intel® Xeon® Processor D-1500 Product Family LAN controller is used as a channel for receiving packets from the network link and passing them to the external BMC. The BMC configures the Intel® Xeon® Processor D-1500 Product Family LAN controller to pass these specific packets to the BMC. Once



a full packet is received from the link and identified as a manageability packet that should be transferred to the BMC, the Intel® Xeon® Processor D-1500 Product Family LAN controller starts the receive TCO flow to the BMC.

The Intel® Xeon® Processor D-1500 Product Family LAN controller uses the SMBus notification method to notify the BMC that it has data to deliver. Since the packet size might be larger than the maximum SMBus fragment size, the packet is divided into fragments, where the Intel® Xeon® Processor D-1500 Product Family LAN controller uses the maximum fragment size allowed in each fragment (configured via the NVM). The last fragment of the packet transfer is always the status of the packet. As a result, the packet is transferred in at least two fragments. The data of the packet is transferred as part of the receive TCO LAN packet transaction.

When SMBus alert is selected as the BMC notification method, the Intel® Xeon® Processor D-1500 Product Family LAN controller notifies the BMC on each fragment of a multi-fragment packet. When asynchronous notify is selected as the BMC notification method, the Intel® Xeon® Processor D-1500 Product Family LAN controller notifies the BMC only on the first fragment of a received packet. It is the BMC's responsibility to read the full packet including all the fragments.

Any timeout on the SMBus notification results in discarding the entire packet. Any NACK by the BMC causes the fragment to be re-transmitted to the BMC on the next Receive Packet command.

The maximum size of the received packet is limited by the Intel® Xeon® Processor D-1500 Product Family LAN controller to 1536 bytes. Packets larger than 1536 bytes are silently discarded. Any packet smaller than 1536 bytes is processed.

10.5.8 Transmit Pass Through Flow

The Intel® Xeon® Processor D-1500 Product Family LAN controller is used as the channel for transmitting packets from the external BMC to the network link. The network packet is transferred from the BMC over the SMBus and then, when fully received by the Intel® Xeon® Processor D-1500 Product Family LAN controller, is transmitted over the network link.

Each SMBus address is connected to a different LAN port. When a packet is received during a SMBus transaction using SMBus address #0, it is transmitted to the network using LAN port #0; it is transmitted through LAN port #1 if received on SMBus address #1, etc.

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports packets up to an Ethernet packet length of 1536 bytes. Since SMBus transactions can only be up to 240 bytes in length, packets might need to be transferred over the SMBus in more than one fragment. This is achieved using the *F* and *L* bits in the command number of the transmit TCO packet Block Write command. When the *F* bit is set, it is the first fragment of the packet. When the *L* bit is set, it is the last fragment of the packet. When both bits are set, the entire packet is in one fragment. The packet is sent over the network link only after all its fragments are received correctly over the SMBus. The maximum SMBus fragment size is defined within the NVM and cannot be changed by the BMC.

The minimum packet length defined by the 802.3 spec is 64 bytes. The Intel® Xeon® Processor D-1500 Product Family LAN controller pads packets that are less than 64 bytes to meet the specification requirements (there is no need for the external BMC to pad packets less than 64 bytes). If the packet sent by the BMC is larger than 1536 bytes, the Intel® Xeon® Processor D-1500 Product Family LAN controller silently discards the packet.

The Intel® Xeon® Processor D-1500 Product Family LAN controller calculates the L2 CRC on the transmitted packet and adds its four bytes at the end of the packet. Any other packet field (such as XSUM or VLAN) must be calculated and inserted by the BMC (the Intel® Xeon® Processor D-1500 Product Family LAN controller does not change any field in the transmitted packet, other than adding padding and CRC bytes).



If the network link is down when the Intel® Xeon® Processor D-1500 Product Family LAN controller has received the last fragment of the packet from the BMC, it silently discards the packet. Note that any link down event during the transfer of any packet over the SMBus does not stop the operation since the Intel® Xeon® Processor D-1500 Product Family LAN controller waits for the last fragment to end to see whether the network link is up again.

10.5.8.1 Transmit Errors in Sequence Handling

Once a packet is transferred over the SMBus from the BMC to the Intel® Xeon® Processor D-1500 Product Family LAN controller, the *F* and *L* flags should follow specific rules. The *F* flag defines the first fragment of the packet; the *L* flag that the transaction contains the last fragment of the packet. Table 10.9 lists the different flag options in transmit packet transactions.

Table 10.9. Flag Options During Transmit Packet Transactions

Previous	Current	Action/Notes
Last	First	Accept both.
Last	Not First	Error for the current transaction. Current transaction is discarded and an abort status is asserted.
Not Last	First	Error in previous transaction. Previous transaction (until previous First) is discarded. Current packet is processed. No abort status is asserted.
Not Last	Not First	Process the current transaction.

Note: Since every other Block Write command in TCO protocol has both *F* and *L* flags on, they cause flushing any pending transmit fragments that were previously received. When running the TCO transmit flow, no other Block Write transactions are allowed in between the fragments.

10.5.8.2 TCO Command Aborted Flow

The Intel® Xeon® Processor D-1500 Product Family LAN controller indicates to the BMC an error or an abort condition by setting the *TCO Abort* bit in the general status. The Intel® Xeon® Processor D-1500 Product Family LAN controller might also be configured to send a notification to the BMC (see Section 10.5.11.1.3.3).

Following is a list of possible error and abort conditions:

- Any error in the SMBus protocol (NACK, SMBus timeouts, etc.).
- If the BMC does not respond until the notification timeout (programmed in the EEPROM) expires.
- Any error in compatibility between required protocols to specific functionality (for example, RX Enable command with a byte count not equal to 1/14, as defined in the command specification).
- If the Intel® Xeon® Processor D-1500 Product Family LAN controller does not have space to store the transmitted packet from the BMC (in its internal buffer space) before sending it to the link, the packet is discarded and the external BMC is notified via the *Abort* bit.
- Error in the *F/L* bit sequence during multi-fragment transactions.
- An internal reset to the Intel® Xeon® Processor D-1500 Product Family LAN controller's firmware.

10.5.9 SMBus Link State Control

While in SMBus mode, the default setting of the link is defined by the *Enable All PHYs in D3 N* bit in Common Firmware Parameters 1 NVM word.

When a channel is enabled through NVM setting or through the *RCV_EN* option of the Receive Enable Command, the link is established (if not already required for other purposes).



If the channel is disabled by clearing of the *RCV_EN* option, then the link may move back to the default defined by the Enable All PHYs in D3 N if not needed for other purposes.

Note: When the link is taken down due to *RCV_EN* being cleared the transmit traffic from the BMC is also stopped.

Note: Before transitioning to D3 it is the responsibility of the driver to request the PHY to be active for wake-up activities.

10.5.10 SMBus ARP Transactions

All SMBus ARP transactions include the PEC byte.

10.5.10.1 Prepare to ARP

This command clears the *Address Resolved* flag (set to false). It does not affect the status or validity of the dynamic SMBus address and is used to inform all devices that the ARP master is starting the ARP process:

1	7	1	1	8	1	8	1	1
S	Slave Address	Wr	A	Command	A	PEC	A	P
	1100 001	0	0	0000 0001	0	[Data Dependent Value]	0	

10.5.10.2 Reset Device (General)

This command clears the *Address Resolved* flag (set to false). It does not affect the status or validity of the dynamic SMBus address.

1	7	1	1	8	1	8	1	1
S	Slave Address	Wr	A	Command	A	PEC	A	P
	1100 001	0	0	0000 0010	0	[Data Dependent Value]	0	


10.5.10.3 Reset Device (Directed)

The Command field is NACKed if bits 7:1 do not match the current SMBus address. This command clears the *Address Resolved* flag (set to false) and does not affect the status or validity of the dynamic SMBus address.

1	7	1	1	8	1	8	1	1
S	Slave Address	Wr	A	Command	A	PEC	A	P
	1100 001	0	0	Targeted Slave Address 0	0	[Data Dependent Value]	0	

10.5.10.4 Assign Address

This command assigns SMBus address. The address and command bytes are always acknowledged.



The transaction is aborted (NACKed) immediately if any of the UDID bytes is different from the Intel® Xeon® Processor D-1500 Product Family LAN controller UDID bytes. If successful, the manageability system internally updates the SMBus address. This command also sets the *Address Resolved* flag (set to true).

1	7	1	1	8	1	8	1	
S	Slave Address	W	A	Command	A	Byte Count	A	...
	1100 001	0	0	0000 0100	0	0001 0001	0	

8	1	8	1	8	1	8	1	
Data 1	A	Data 2	A	Data 3	A	Data 4	A	...
UDID Byte 15 (MSB)	0	UDID Byte 14	0	UDID Byte 13	0	UDID Byte 12	0	

8	1	8	1	8	1	8	1	
Data 5	A	Data 6	A	Data 7	A	Data 8	A	...
UDID Byte 11	0	UDID Byte 10	0	UDID Byte 9	0	UDID Byte 8	0	


8	1	8	1	8	1	
Data 9	A	Data 10	A	Data 11	A	...
UDID Byte 7	0	UDID Byte 6	0	UDID Byte 5	0	

8	1	8	1	8	1	8	1	
Data 12	A	Data 13	A	Data 14	A	Data 15	A	...
UDID Byte 4	0	UDID Byte 3	0	UDID Byte 2	0	UDID Byte 1	0	

8	1	8	1	8	1	1
Data 16	A	Data 17	A	PEC	A	P
UDID Byte 0 (LSB)	0	Assigned Address	0	[Data Dependent Value]	0	

10.5.10.5 Get UDID (General and Directed)

The general get UDID SMBus transaction supports a constant command value of 0x03 and, if directed, supports a Dynamic command value equal to the dynamic SMBus address.



If the SMBus address has been resolved (*Address Resolved* flag set to true), the manageability system does not acknowledge (NACK) this transaction. If it's a General command, the manageability system always acknowledges (ACKs) as a directed transaction.

This command does not affect the status or validity of the dynamic SMBus address or the *Address Resolved* flag.

S	Slave Address	Wr	A	Command	A	S	...
	1100 001	0	0	See Below	0		

7	1	1	8	1	...
Slave Address	Rd	A	Byte Count	A	...
1100 001	1	0	0001 0001	0	

8	1	8	1	8	1	8	1	...
Data 1	A	Data 2	A	Data 3	A	Data 4	A	...
UDID Byte 15 (MSB)	0	UDID Byte 14	0	UDID Byte 13	0	UDID Byte 12	0	

8	1	8	1	8	1	8	1	...
Data 5	A	Data 6	A	Data 7	A	Data 8	A	...
UDID Byte 11	0	UDID Byte 10	0	UDID Byte 9	0	UDID Byte 8	0	

8	1	8	1	8	1	...
Data 9	A	Data 10	A	Data 11	A	...
UDID Byte 7	0	UDID Byte 6	0	UDID Byte 5	0	

8	1	8	1	8	1	8	1	...
Data 12	A	Data 13	A	Data 14	A	Data 15	A	...
UDID Byte 4	0	UDID Byte 3	0	UDID Byte 2	0	UDID Byte 1	0	

8	1	8	1	8	1	1
Data 16	A	Data 17	A	PEC	~Ä	P
UDID Byte 0 (LSB)	0	Device Slave Address	0	[Data Dependent Value]	1	

The Get UDID command depends on whether or not this is a Directed or General command.



The General Get UDID SMBus transaction supports a constant command value of 0x03.

The Directed Get UDID SMBus transaction supports a Dynamic command value equal to the dynamic SMBus address with the LSB bit set.

Note: Bit 0 (LSB) of Data byte 17 is always 1b.

10.5.11 SMBus Pass-Through Transactions

This section details commands (both read and write) that the Intel® Xeon® Processor D-1500 Product Family LAN controller SMBus interface supports for pass-through.

10.5.11.1 Write SMBus Transactions

This section details the commands that the BMC can send to the Intel® Xeon® Processor D-1500 Product Family LAN controller over the SMBus interface. The SMBus write transactions table lists the different SMBus write transactions supported by the Intel® Xeon® Processor D-1500 Product Family LAN controller.

TCO Command	Transaction	Command	Fragmentation	Section
Transmit Packet	Block Write	First: 0x84 Middle: 0x04 Last: 0x44	Multiple	10.5.11.1.1
Transmit Packet	Block Write	Single: 0xC4	Single	10.5.11.1.1
Request Status	Block Write	Single: 0xDD	Single	10.5.11.1.2
Receive Enable	Block Write	Single: 0xCA	Single	10.5.11.1.3
Force TCO	Block Write	Single: 0xCF	Single	10.5.11.1.4
Management Control	Block Write	Single: 0xC1	Single	10.5.11.1.5
Update MNG RCV Filter Parameters	Block Write	Single: 0xCC	Single	10.5.11.1.6
Set Common Filters	Block Write	Single: 0xC2	Single	10.5.11.1.7
Clear All Filters	Block Write	Single: 0xC3	Single	10.5.11.1.8
Update MACsec parameters	Block Write	Single: 0xC9	Single	10.5.11.1.9

10.5.11.1.1 Transmit Packet Command

The Transmit Packet command behavior is detailed in [Section 10.5.8](#). The Transmit Packet fragments have the following format.

The payload length is limited to the maximum payload length set in the NVM. If the overall packet length is bigger than 1536 bytes, the packet is silently discarded.

Function	Command	Byte Count	Data 1	...	Data N
Transmit first fragment	0x84	N	Packet data MSB	...	Packet data LSB
Transmit middle fragment	0x04				
Transmit last fragment	0x44				
Transmit single fragment	0xC4				



10.5.11.1.2 Request Status Command

An external BMC can initiate a request to read the Intel® Xeon® Processor D-1500 Product Family LAN controller manageability status by sending a Request Status command. When received, the Intel® Xeon® Processor D-1500 Product Family LAN controller initiates a notification to an external BMC when status is ready. After this, the external controller will be able to read the status, by issuing a read status command (see Section 10.5.11.2.2).

The format is as follows:


Function	Command	Byte Count	Data 1
Request Status	0xDD	1	0

10.5.11.1.3 Receive Enable Command

The Receive Enable command is a single fragment command used to configure the Intel® Xeon® Processor D-1500 Product Family LAN controller. This command has two formats: short, 1-byte legacy format (providing backward compatibility with previous components) and long, 14-byte advanced format (allowing greater configuration capabilities). The Receive Enable command format is as follows:

Function	CMD	Byte Count	Data 1	Data 2	...	Data 7	Data 8	...	Data 11	Data 12	Data 13	Data 14
Legacy Receive Enable	0xCA	1	Receive Control Byte	-	...	-	-	...	-	-	-	-
Advanced Receive Enable		14 (0x0E)		MAC Addr MSB		MAC Addr LSB	IP Addr MSB		IP Addr LSB	BMC SMBus Addr	I/F Data Byte	Alert Value Byte

Field	Bit(s)	Description
RCV_EN	0	Receive TCO Enable. 0b= Disable receive TCO packets. 1b= Enable Receive TCO packets. Setting this bit enables all manageability receive filtering operations. Enabling specific filters is done via the NVM or through special configuration commands. Note: When the <i>RCV_EN</i> bit is cleared, all receive TCO functionality is disabled, not just the packets that are directed to the BMC (also auto ARP packets).
RCV_ALL	1	Receive All Enable. 0b = Disable receiving all packets. 1b = Enable receiving all packets. Forwards all packets received over the wire that passed L2 filtering to the external BMC. This flag has no effect if bit 0 (Enable TCO packets) is disabled.
EN_STA	2	Enable Status Reporting. 0b = Disable status reporting. 1b = Enable status reporting.




EN_ARP_RES	3	<p>Enable ARP Response.</p> <p>0b = Disable the Intel® Xeon® Processor D-1500 Product Family LAN controller ARP response. The Intel® Xeon® Processor D-1500 Product Family LAN controller treats ARP packets as any other packet, for example, packet is forwarded to the BMC if it passed other (non-ARP) filtering.</p> <p>1b = Enable the Intel® Xeon® Processor D-1500 Product Family LAN controller ARP response. The Intel® Xeon® Processor D-1500 Product Family LAN controller automatically responds to all received ARP requests that match the IP address programmed by the BMC. The BMC IP address is provided as part of the Receive Enable message (bytes 8:11). If a short version of the command is used, the Intel® Xeon® Processor D-1500 Product Family LAN controller uses IP address configured in the most recent long version of the command in which the EN_ARP_RES bit was set. If no such previous long command exists, then the Intel® Xeon® Processor D-1500 Product Family LAN controller uses the IP address configured in the NVM as ARP Response IPv4 Address in the pass-through LAN configuration structure.</p> <p>If the <i>CBDM</i> bit is set, the Intel® Xeon® Processor D-1500 Product Family LAN controller uses the BMC dedicated MAC address in ARP response packets. If the <i>CBDM</i> bit is not set, the BMC uses the host MAC address.</p> <p>When the ARP off load feature is activated, the Intel® Xeon® Processor D-1500 Product Family LAN controller uses the following registers to filter the ARP traffic addressed to the BMC. BMC should not modify these registers:</p> <ul style="list-style-type: none"> Manageability Decision Filter – MDEF6 (and corresponding bit 6 in Management Only traffic Register – <i>MNGONLY</i>). <p>IPv4 address - MIPAF4[3]</p>
NM	5:4	<p>Notification Method. Define the notification method the Intel® Xeon® Processor D-1500 Product Family LAN controller uses.</p> <p>00b = SMBUS Alert.</p> <p>01b = Asynchronous notify.</p> <p>10b = Direct receive.</p> <p>11b = Not supported.</p> <p>Note: Changing the notification method in any port updates the notification method of all ports.</p>
Reserved	6	Reserved. Must be set to 1b.
CBDM	7	<p>Configure the BMC Dedicated MAC Address.</p> <p>Note: This bit should be 0b when the <i>RCV_EN</i> bit (bit 0) is not set.</p> <p>0b = The Intel® Xeon® Processor D-1500 Product Family LAN controller shares the MAC address for MNG traffic with the host MAC address, which is specified in NVM words 0x0-0x2.</p> <p>1b = The Intel® Xeon® Processor D-1500 Product Family LAN controller uses the BMC dedicated MAC address as a filter for incoming receive packets.</p> <p>The BMC MAC address is set in bytes 2-7 in this command.</p> <p>If a short version of the command is used, the Intel® Xeon® Processor D-1500 Product Family LAN controller uses the MAC address configured in the most recent long version of the command in which the <i>CBDM</i> bit was set.</p> <p>When the dedicated MAC address feature is activated, the Intel® Xeon® Processor D-1500 Product Family LAN controller uses the following registers to filter in all the traffic addressed to the BMC MAC. BMC should not modify these registers:</p> <ul style="list-style-type: none"> Manageability Decision Filter – MDEF7 (and corresponding bit 7 in Management Only traffic Register – <i>MNGONLY</i>) Manageability MAC Address Low – <i>MMAL</i>[3] Manageability MAC Address High – <i>MMAH</i>[3] <p>Note: When the dedicated MAC address feature is cleared, these registers are not programmed and the BMC might use other filters to enforce MAC filtering using the Update Management Receive Filter Parameters command.</p>

10.5.11.1.3.1 Management MAC Address (Data Bytes 7:2)

Ignored if the *CBDM* bit is not set. This MAC address is used to configure the dedicated MAC address. In addition, it is used in the ARP response packet when the EN_ARP_RES bit is set. This MAC address is also used when *CBDM* bit is set in subsequent short versions of this command.

10.5.11.1.3.2 Management IP Address (Data Bytes 11:8)

This IP address is used to filter ARP request packets.



10.5.11.1.3.3 Asynchronous Notification SMBus Address (Data Byte 12)

This address is used for the asynchronous notification SMBus transaction and for direct receive. The SMBus address is stored in bit 7:1 of this byte. Bit 0 is always 0.

10.5.11.1.3.4 Interface Data (Data Byte 13)

Interface data byte used in asynchronous notification.

10.5.11.1.3.5 Alert Value Data (Data Byte 14)

Alert Value data byte used in asynchronous notification.

10.5.11.1.4 Force TCO Command

This command causes the Intel® Xeon® Processor D-1500 Product Family LAN controller to perform a TCO reset, TCO isolate, or Firmware Reset

TCO Reset: If Force TCO reset is enabled in the NVM, the force TCO reset clears the data path (Rx/Tx) of the Intel® Xeon® Processor D-1500 Product Family LAN controller to enable the BMC to transmit/receive packets through the Intel® Xeon® Processor D-1500 Product Family LAN controller by assertion of a global reset. Force TCO reset is asserted only to the port related to the SMBus address the command. This command should only be used when the BMC is unable to transmit receive and suspects that the Intel® Xeon® Processor D-1500 Product Family LAN controller is inoperable. The command also causes the LAN device driver to unload. It is recommended to perform a system restart to resume normal operation.

Firmware Reset: This command causes re-initialization of all the embedded controller functions and a re-load of related NVM words. A firmware reset is achieved by setting the *GSCR.SET_FWRST Aux* bit.

The Intel® Xeon® Processor D-1500 Product Family LAN controller considers the Force TCO reset command as an indication that the operating system is hung. The Force TCO command format is as follows:

Function	Command	Byte Count	Data 1
Force TCO Reset	0xCF	1	TCO Mode

Where TCO mode is:

Field	Bit(s)	Description
DO_TCO_RST	0	Perform TCO Reset. 0b = Do nothing. 1b = Perform TCO reset.
DO_TCO_ISOLATE ¹	1	Do TCO Isolate 0b = Enable PCIe write access to LAN port. 1b = Isolate Host PCIe write operation to the port. Note: Should be used for debug only.
RESET_MGMT	2	Reset manageability; re-load manageability NVM words. 0b = Do nothing. 1b = Issue firmware reset to manageability. Setting this bit generates a one-time firmware reset. Following the reset, management related data from NVM is loaded.
Reserved	7:3	Reserved (set to 0x00).



1. TCO isolate host write operation enabled in NVM.

Note: Only one of the fields should be set in a given command. Setting more than one field might yield unexpected results.

10.5.11.1.5 Management Control

This command is used to set generic manageability parameters. The parameters list is shown in Table 10.10. The command is 0xC1 stating that it is a Management Control command. The first data byte is the parameter number and the data afterwards (length and content) are parameter specific as shown in Management Control Command Parameters/Content.

Note: If the parameter that the BMC sets is not supported by the Intel® Xeon® Processor D-1500 Product Family LAN controller. The Intel® Xeon® Processor D-1500 Product Family LAN controller does not NACK the transaction. After the transaction ends, the Intel® Xeon® Processor D-1500 Product Family LAN controller discards the data and asserts a transaction abort status.

The Management Control command format is as follows:

Function	Command	Byte Count	Data 1	Data 2	...	Data N
Management Control	0xC1	N	Parameter Number	Parameter Dependent		

Table 10.10. Management Control Command Parameters/Content

Parameter	#	Parameter Data
Keep PHY Link Up	0x00	A single byte parameter: Data 2: Bit 0= Set to indicate that the PHY link for this port should be kept up throughout system resets. This is useful when the server is reset and the BMC needs to keep connectivity for a manageability session. Bit [7:1] Reserved. 0b= Disabled. 1b= Enabled.

10.5.11.1.6 Update Management Receive Filter Parameters

This command is used to set the manageability receive filters parameters. The command is 0xCC. The first data byte is the parameter number and the data that follows (length and content) are parameter specific as listed in management RCV filter parameters.

If the parameter that the BMC sets is not supported by the Intel® Xeon® Processor D-1500 Product Family LAN controller, then the Intel® Xeon® Processor D-1500 Product Family LAN controller does not NACK the transaction. After the transaction ends, the Intel® Xeon® Processor D-1500 Product Family LAN controller discards the data and asserts a transaction abort status.

The update management RCV receive filter parameters command format is as follows:



Function	Command	Byte Count	Data 1	Data 2	...	Data N
Update Manageability Filter Parameters	0xCC	N	Parameter Number	Parameter Dependent		

Table 10.11 lists the different parameters and their content.

Table 10.11. Management Receive Filter Parameters

Parameter	Number	Parameter Data
Filters Enables	0x1	Defines the generic filters configuration. The structure of this parameter is four bytes as the Manageability Control (MANC) register Note: The general filter enable is in the Receive Enable command that enables receive filtering.
MNGONLY configuration	0xF	This parameter defines which of the packets types identified as manageability packets in the receive path will never be directed to the host memory. Data 2:5: MNGONLY register bytes - Data 2 is the MSB.
Flex Filter 0 Enable Mask and Length	0x10	Flex Filter 0 Mask. Data 17:2 = Mask. Bit 0 in data 2 is the first bit of the mask. Data 19:18 = Reserved. Should be set to 00b. Data 20 = Flexible filter length.
Flex Filter 0 Data	0x11	Data 2 — Group of flex filter's bytes: 0x0 = bytes 0-29. 0x1 = bytes 30-59. 0x2 = bytes 60-89. 0x3 = bytes 90-119. 0x4 = bytes 120-127. Data 3:32 = Flex filter data bytes. Data 3 is LSB. Group's length is not a mandatory 30 bytes; it might vary according to filter's length and must NOT be padded by zeros.
Decision Filters	0x61	This command is obsolete and should not be used anymore. Please use 0x68 instead.
VLAN Filters	0x62	Three bytes are required to load the VLAN tag filters. Data 2: VLAN filter number. Data 3: MSB of VLAN filter. Data 4: LSB of VLAN filter.
Flex Port Filters	0x63	Three to four bytes are required to load the manageability flex port filters. Data 2: Flex port filter number. Data 3: MSB of flex port filter. Data 4: LSB of flex port filter.
IPv4 Filters	0x64	Five bytes are required to load the IPv4 address filter. Data 2: IPv4 address filter number (3:0). Data 3: LSB of IPv4 address filter. ... Data 6: MSB of IPv4 address filter.
IPv6 Filters	0x65	17 bytes are required to load the IPv6 address filter. Data 2 — IPv6 address filter number (3:0). Data 3 — LSB of IPv6 address filter. ... Data 18 — MSB of IPv6 address filter.

**Table 10.11. Management Receive Filter Parameters (Continued)**

MAC Filters	0x66	Seven bytes are required to load the MAC address filters. Data 2 — MAC address filters pair number (3:0). Data 3 — MSB of MAC address. ... Data 8: LSB of MAC address.
EtherType Filters	0x67	5 bytes to load EtherType Filters (METF). Data 2 — METF filter index (valid values are 0, 1, 2, 3). Data 3 — MSB of METF. ... Data 6 — LSB of METF.
Extended Decision Filter	0x68	9 bytes to load the extended decision filters (MDEF_EXT & MDEF). Data 2 — MDEF filter index (valid values are 0...5). Data 3 — MSB of MDEF_EXT (DecisionFilter1). ... Data 6 — LSB of MDEF_EXT (DecisionFilter1). Data 7 — MSB of MDEF (DecisionFilter0). ... Data 10 — LSB of MDEF (DecisionFilter0). The command must overwrite any previously stored value.

Table 10.12. Filter Enable Parameters

Bit	Name	Description
16:0	Reserved	Reserved.
17	RCV_TCO_EN	Receive TCO Packets Enabled. When this bit is set it enables the receive flow to the manageability block. This bit should be set only if at least one of EN_BMC2OS or EN_BMC2NET bits are set. This bit is usually set using the receive enable command (see Section 10.5.11.1.3).
18	Reserved	Reserved.
22:19	Reserved	Reserved.
23	Enable Xsum Filtering to MNG	When this bit is set, only packets that pass the L3 and L4 checksum are send to the manageability block.
24	EN_IPV4_FILTER	Enable IPv4 address Filters - when set, the last 128 bits of the MIPAF register are used to store 4 IPv4 addresses for IPv4 filtering. When cleared, these bits store a single IPv6 filter
25	FIXED_NET_TYPE	Fixed net type: If set, only packets matching the net type defined by the NET_TYPE field passes to manageability. Otherwise, both tagged and un-tagged packets can be forwarded to the manageability engine.
26	NET_TYPE	NET TYPE: 0b = pass only un-tagged packets. 1b = pass only VLAN tagged packets. Valid only if FIXED_NET_TYPE is set.
31:27	Reserved	Reserved.

10.5.11.1.7 Set Common Filters Command

The Set Common Filters command is a single fragment command capable of configuring the most common filters.



Note: If this command is used, all the other commands that programs forwarding filters should not be used (apart from the Clear All Filters command). When this command is received, an implied Clear All Filters command is done before the application of this command.

The Set Common Filters command has two possible formats:

IPv4 Format:

Function	Command	Byte Count	Data 1	Data 2:4	5:10	Data 11	Data 12	Data 13	Data 14:17
Set Common Filters	0xC2	17	Opcode = 0	Receive Control - see Table 10.13	MAC Address	BMC Alert Address	Interface Data Byte	Alert Value Byte	IPv4 Address

IPv6 Format:

Function	Command	Byte Count	Data 1	Data 2:4	5:10	Data 11	Data 12	Data 13	Data 14:29
Set Common Filters	0xC2	29	Opcode = 0	Receive Control - see Table 10.13	MAC Address	BMC Alert Address	Interface Data Byte	Alert Value Byte	IPv6 Address



Table 10.13. Set Common Filters Receive Control Bytes (Sheet 1 of 2)

Byte	Bit	Field	Description
1	0	RCV_EN	Receive TCO Packets Enabled. When this bit is set it enables the receive flow to the manageability block. This bit should be set only if at least one of EN_BMC2O or EN_BMC2NET bits are set.
	1	EN_STA	Enable Status Reporting. 0b = Disable status reporting. 1b = Enable status reporting.
	2	Auto ARP	Automatically respond to ARP packets. Ignored in IPv6 mode. If this bit is set, broadcast ARP packets will be handled by the Intel® Xeon® Processor D-1500 Product Family LAN controller and ARP requests to the IP address set in the command will be responded. Note: Mutually exclusive to Configure ARP/ Neighborhood Filter bit. If this bit is set, the IP address must be valid and contain the IP address of the MC. This bit is ignored if RCV_EN is cleared. Note: When set, ARP requests to the BMC IP defined below (unicast or Broadcast) will be sent to the internal Firmware for processing. ARP response will be dropped.
	3	Enable Xsum Filtering to MNG	When this bit is set, only packets that pass the L3 and L4 checksum are send to the manageability block. This bit is ignored if RCV_EN is cleared.
	5:4	Reserved	Reserved.
	7:6	Notification Method	Notification Method. Define the notification method the Intel® Xeon® Processor D-1500 Product Family LAN controller uses. 00b = SMBUS Alert. 01b = Asynchronous notify. 10b = Direct receive. 11b = Not supported.



Table 10.13. Set Common Filters Receive Control Bytes (Sheet 2 of 2)

Byte	Bit	Field	Description
2	8	CBDM	<p>Configure the BMC Dedicated MAC Address.</p> <p>0b= The Intel® Xeon® Processor D-1500 Product Family LAN controller shares the MAC address for MNG traffic with the host MAC address, which is specified in RAH/RAL[0] registers</p> <p>1b= The Intel® Xeon® Processor D-1500 Product Family LAN controller uses the BMC dedicated MAC address as a filter for incoming receive packets.</p> <p>The BMC MAC address is set in bytes 5-10 in this command.</p> <p>This bit is ignored if RCV_EN is cleared.</p> <p>If this bit is cleared, at least one of bits 9,10 or 11 must be set. If only bit 9 is set, the IP address should be different than the address used by the host.</p>
	9	Configure IP Address Filter	Automatically configure an IP Address Filter. If this bit is set, only packets matching this IP address and the MAC address defined by the CBDM bit is forwarded. This bit is ignored if RCV_EN is cleared
	10	Configure RMCP 26Fh Filter	Automatically configure standard IPMI port 0x26F filters. If this bit is set, only packets matching this port and the MAC address defined by the CBDM bit is forwarded. If the Configure IP address <i>Filter</i> bits is set, only packets matching the IP address and this port is forwarded. The other port enable bit (11) might add additional forwarding condition. This bit is ignored if RCV_EN is cleared.
	11	Configure RMCP 298h Filter	Automatically configure standard IPMI port 0x298 filter. If this bit is set, only packets matching this port and the MAC address defined by the CBDM bit is forwarded. If the Configure IP address <i>Filter</i> bits is set, only packets matching the IP address and this port is forwarded. The other port enable bit (10) might add additional forwarding condition. This bit is ignored if RCV_EN is cleared.
	12	Configure ARP/ Neighborhood Filter	<p>Automatically configure filters to enable this traffic to the BMC (mutually exclusive to Auto ARP bit). If this bit is set, broadcast or unicast ARP packets are forwarded to the BMC. These packets might also be sent to the host.</p> <p>In IPv4 mode, setting this bit enables forwarding of broadcast or unicast ARP request and response. If IP address is set, only request or response to this address is forwarded.</p> <p>In IPv6 mode, setting this bit enables forwarding of all types of neighbor discovery and MLD ICMPv6 packet types:</p> <ul style="list-style-type: none"> • 0x86 (134d) - Router advertisement. • 0x87 (135d) - Neighbor solicitation. • 0x88 (136d) - Neighbor advertisement. • 0x89 (137d) - Redirect. • 0x82 (130d) - MLD query. • 0x83 (131d) - MLDv1 report. • 0x84 (132d) - MLD done. • 0x8F (143d) - MLDv2 report. <p>This bit is ignored if RCV_EN is cleared.</p>
	13	Configure DHCP port 44h Filter (DHCP server packets)	Automatically configure DHCP port 44 filter to the BMC. If this bit is set, broadcast packets matching this port is forwarded. Otherwise, this port is not added to the broadcast filtering option. This bit is ignored if RCV_EN is cleared or in IPv6 mode.
	15:14	Reserved	Reserved.
3	16	Disable Host ARP	Configure ARP requests and network neighborhood packets not to go to the host. This bit should be cleared in regular operation. Ignored if both bit 12 and bit 2 are cleared or if RCV_EN is cleared.
	17	Disable Host DHCP	Configure DHCP packets (port 0x44) not to go to the host. This bit should be cleared in regular operation. Ignored if bit 13 is cleared, RCV_EN is cleared, or in IPv6 mode.
	24:18	Reserved	Reserved.



10.5.11.1.8 Clear all Filters Command

The Clear all Filters command is a single fragment command capable of clearing all the receive filters currently programmed for manageability traffic.

Function	Command	Byte Count	Data
Clear all Filters	0xC3	1	0x00

10.5.11.1.9 Update MACsec Parameters

This command is used to set the manageability MACsec parameters. The parameters list is shown in the table below. The first data byte is the parameter number and the data after-words (length and content) are parameter specific as shown in the table.

This is the format of the Update MACsec parameters command:

Function	Command	Byte Count	Data 1	Data 2	...	Data N
Update MACsec filter parameters	0xC9	N	Parameter number (PN#)	Parameter dependent		

The table below shows the different parameters and their contents:

Parameter	PN#	Parameter Data
Transfer MACsec ownership to BMC	0x10	Data 2: Host Control: Bit 0 – Reserved Bit 1 – Allow host traffic (0b – blocked, 1b – allowed) Bit 2...31 – Reserved.
Transfer MACsec ownership to Host	0x11	No data needed
Initialize MACsec RX	0x12	Data 2: RX Port Identifier (MSB)... Data 3: (LSB) RX Port Identifier – the port number by which the Intel® Xeon® Processor D-1500 Product Family LAN controller will identify RX packets. It is recommended that the BMC uses 0x0 as the port identifier. Note: The BMC should use the same port identifier when performing the key-exchange. Data 4: RX MAC SecY (MSB) ... Data 9: (LSB)
Initialize MACsec TX	0x13	Data 2: TX Port Identifier (MSB)... Data 3: (LSB) — must be set to zero Data 4: TX SCI (MSB) ... Data 7: TX SCI (LSB) TX SCI – A 6 bytes unique identifier for the MACsec TX CA. It is recommended that the BMC uses its Ethernet MAC Address value for this field. Data 8: Reserved Data 9: Reserved Data 10: Packet Number Threshold (MSB) ... Data 12: (LSB) PN Threshold – When a new key is programmed, the Packet Number is reset to 0x1. With each TX packet, The Packet Number is increased by 1 and inserted to the packet (to avoid replay attacks). The PN Threshold value is 3 MSBytes of the TX Packet number after which a "Key Exchange Required" AEN will be sent to the BMC. Example: a PN Threshold of 0x123456 means that when the PN reaches 0x12345600 a notification will be sent Data 22: TX Control — See Table 10.14
Set MACsec RX Key	0x14	Data 2: Reserved Data 3: RX SA AN (The Association Number to be used with this key) Data 4: RX MACsec Key (MSB) ... Data 19: (LSB) — (16 bytes key to be used)



Set MACsec TX Key	0x15	Data 3: TX SA AN (The Association Number to be used with this key) Data 4: TX MACsec Key (MSB) ... Data 19: (LSB) — (16 bytes key to be used)
Enable MACsec Network TX encryption	0x16	No data needed
Disable MACsec Network TX encryption	0x17	No data needed

Table 10.14. TX Control

Bit	Description
0..4	Reserved
5	Always Include SCI in TX: 0b – Do not include SCI in TX packets 1b – Include SCI in TX packets
6..7	Reserved

10.5.11.2 Read SMBus Transactions

This section details the pass-through read transactions that the BMC can send to the Intel® Xeon® Processor D-1500 Product Family LAN controller over SMBus.

SMBus read transactions lists the different SMBus read transactions supported by the Intel® Xeon® Processor D-1500 Product Family LAN controller. All the read transactions are compatible with SMBus read block protocol format.

Table 10.15. SMBus Read Transactions

TCO Command	Transaction	Write Command ¹	Read Command	Opcode	Fragments	Section
Receive TCO Packet	Block Read	N/A	0xD0 or 0xC0	First: 0x90 Middle: 0x10 Last ² : 0x50	Multiple	10.5.11.2.1
Read Status	Block Read	N/A	0xD0 or 0xC0 or 0xDE	Single: 0xDD	Single	10.5.11.2.2
Get System MAC Address	Block Read	N/A	0xD4	Single: 0xD4	Single	10.5.11.2.3
Read Management Parameters	Block Read	0xC1	0xD1	Single: 0xD1	Single	10.5.11.2.4
Read Management RCV Filter Parameters	Block Read	0xCC	0xCD	Single: 0xCD	Single	10.5.11.2.5
Get Controller Information	Block Read	0xD5	0xD5	Single: 0xD5	Single	10.5.11.2.6
Get Common filters	Block Read	0xD3	0xD3	Single: 0xD3	Single	10.5.11.2.7
Read MACsec parameters	Block Read	0xC9	0xD9	Single: 0xD9	Single	10.5.11.2.8
Read Receive Enable Configuration	Block Read	N/A	0xDA	Single: 0xDA	Single	10.5.11.2.9

1. In some commands, a preliminary Write command is sent to signal firmware to prepare the data for the upcoming Read command. This column describes the OpCode used for the Write command.
2. The last fragment of the receive packet is the packet status.



0xC0 or 0xD0 commands are used for more than one payload. If BMC issues these read commands, and the Intel® Xeon® Processor D-1500 Product Family LAN controller has no pending data to transfer, it always returns as default opcode 0xDD with the Intel® Xeon® Processor D-1500 Product Family LAN controller status and does not NACK the transaction.

If an SMBus Quick Read command is received, it is handled as a the Intel® Xeon® Processor D-1500 Product Family LAN controller Request Status command (See [Section 10.5.11.1.2](#) for details).

10.5.11.2.1 Receive TCO LAN Packet Transaction

The BMC uses this command to read packets received on the LAN and its status. When the Intel® Xeon® Processor D-1500 Product Family LAN controller has a packet to deliver to the BMC, it asserts the SMBus notification for the BMC to read the data (or direct receive). Upon receiving notification of the arrival of a LAN receive packet, the BMC begins issuing a receive TCO packet command using the block read protocol.

A packet can be transmitted to the BMC in at least two fragments (at least one for the packet data and one for the packet status). As a result, BMC should follow the *F* and *L* bit of the OpCode.

The OpCode can have these values:

- 0x90 — First fragment
- 0x10 — Middle fragment
- When the OpCode is 0x50, this indicates the last fragment of the packet, which contains packet status.

If a notification timeout is defined (in the NVM) and the BMC does not finish reading the entire packet within the timeout period, since the packet has arrived, the packet is silently discarded. The time spent in ARA cycle or in reading the packet is not counted by the timeout counter.

Following is the receive TCO packet format and the data format returned from the Intel® Xeon® Processor D-1500 Product Family LAN controller.

Function	Command
Receive TCO Packet	0xC0 or 0xD0

Function	Byte Count	Data 1 (Op-Code)	Data 2	...	Data N
Receive TCO First Fragment	N	0x90	Packet Data Byte	...	Packet Data Byte
Receive TCO Middle Fragment		0x10			
Receive TCO Last Fragment	9 (0x9)	0x50	See Section 10.5.11.2.1.1		

10.5.11.2.1.1 Receive TCO LAN Status Payload Transaction

This transaction is the last transaction that the Intel® Xeon® Processor D-1500 Product Family LAN controller issues when a packet received from the LAN is transferred to the BMC. The transaction contains the status of the received packet.

The format of the status transaction is as follows:



Function	Byte Count	Data 1 (OpCode)	Data 2 – Data 17 (Status Data)
Receive TCO Long Status	9 (0x9)	0x50	See Below

The status is 8 bytes where byte 0 (bits 7:0) is set in Data 2 of the status and byte 7 in Data 9 of the status. Table 10.16 lists the content of the status data.

Table 10.16. TCO LAN Packet Status Data

Name	Bits	Description
Packet Length	13:0	Packet length including CRC, only 14 LSB bits.
Reserved	15:14	Reserved.
Packet status	31:16	See Table 10.17.
VLAN	47:32	The two bytes of the VLAN header tag.
MNG status	63:48	See Table 10.18. This field should be ignored if receive is not enabled,

Table 10.17. Packet Status Info

Field	Bit(s)	Description
Reserved	15:4	Reserved.
LAN#	3:2	Indicates the source port of the packet. 00b = Port 0. 01b = Port 1. 10b - 11b = Reserved.
VP	1	VLAN stripped (indicates if the VLAN is part of the packet, or was removed).
CRC stripped	0	Insertion of CRC is needed.

Table 10.18. MNG Status

Name	Bits	Description
Reserved	15:9	Reserved.
Decision Filter match	8	Set when there is a match to one of the Decision filters.
Decision Filter index	7:4	Indicates which of the decision filters match the packet. (allows for up to 16 filters - although only 8 are currently supported).
MNG VLAN Address Match	3	Set when the MNG packet matches one of the MNG VLAN filters.
Pass MNG VLAN Filter Index	2:0	Indicates which of the VLAN filters match the packet.

10.5.11.2.2 Read Status Command

The BMC should use this command after receiving a notification from the Intel® Xeon® Processor D-1500 Product Family LAN controller (such as SMBus Alert). The Intel® Xeon® Processor D-1500 Product Family LAN controller also sends a notification to the BMC in either of the following two cases:

- The BMC asserts a request for reading the status.
- The Intel® Xeon® Processor D-1500 Product Family LAN controller detects a change in one of the Status Data 1 bits (and was set to send status to the BMC on status change) in the Receive Enable command.



Note: Commands 0xC0/0xD0 are for backward compatibility and can be used for other payloads. The Intel® Xeon® Processor D-1500 Product Family LAN controller defines these commands in the OpCode as well as which payload this transaction is. When the 0XDE command is set, the Intel® Xeon® Processor D-1500 Product Family LAN controller always returns opcode 0XDD with the Intel® Xeon® Processor D-1500 Product Family LAN controller status. The BMC reads the event causing the notification, using the Read Status command as follows.

The Intel® Xeon® Processor D-1500 Product Family LAN controller response to one of the commands (0xC0 or 0xD0) in a given time as defined in the SMBus Notification Timeout and Flags word in the NVM.

Function	Command
Read Status	0XC0 or 0XD0 or 0XDE

Function	Byte Count	Data 1 (OpCode)	Data 2 (Status Data 1)	Data 3 (Status Data 2)
Receive TCO Partial Status	3	0XDD	See as follows	

This command can also be executed using the I²C quick read format as follows:

1	7	1	1	8	1	8	1	8	1	1
Start	Slave Address	Rd	Ack	Byte Count	Ack	Status Data 1	Ack	Status Data 2	Ack	Stop
		1	0	0000 0002	0		0		1	

Table 10.19 lists the status data byte 1 parameters.

Table 10.19. Status Data Byte 1

Bit	Name	Description
7	LAN Port	LAN port: defines the port that sent status.
6	TCO Command Aborted	1b = A TCO command abort event occurred since the last read status cycle. 0b = A TCO command abort event did not occur since the last read status cycle.
5	Link Status Indication	0b = LAN link down. 1b = LAN link up.
4	PHY Link Forced Up	Contains the value of the <i>PHY_Link_Up</i> bit. When set, indicates that the PHY link is configured to keep the link up.
3	Initialization Indication	0b = An NVM reload event has not occurred since the last Read Status cycle. 1b = An NVM reload event has occurred since the last Read Status cycle ¹ .
2	Reserved	Reserved.
1:0	Power State	00b = Dr state. 01b = D0u state. 10b = D0 state. 11b = D3 state.

1. This indication is asserted when the Intel® Xeon® Processor D-1500 Product Family LAN controller manageability block reloads the NVM and its internal database is updated to the NVM default values. This is an indication that the external BMC should reconfigure the Intel® Xeon® Processor D-1500 Product Family LAN controller, if other values other than the NVM default should be configured.

Status data byte 2 is used by the BMC to indicate whether the LAN device driver is up and running.



The LAN device driver valid indication is a bit set by the LAN device driver during initialization; the bit is cleared when the LAN device driver enters a Dx state or is cleared by the hardware on a PCI reset.

Bits 2 and 1 indicate that the LAN device driver is stuck. Bit 2 indicates whether the interrupt line of the LAN function is asserted. Bit 1 indicates whether the LAN device driver dealt with the interrupt line before the last Read Status cycle. Table 10.20 lists status data byte 2.

Table 10.20. Status Data Byte 2

Bit	Name	Description
7:4	Reserved	Reserved.
4	MACsec indication	If set — indicates that a MACsec event has occurred. Use the "Read MACsec parameters" with "MACsec interrupt cause" parameter to read the interrupt cause
3	Driver Valid Indication	0b = LAN driver is not alive. 1b = LAN driver is alive.
2:1	Reserved	Reserved.
0	Reserved	Reserved.

Table 10.21 lists the possible values of bits 2 and 1 and what the BMC can assume from the bits:

Table 10.21. Status Data Byte 2 (Bits 2 and 1)

Previous	Current	Description
Don't Care	00b	Interrupt is not pending (OK).
00b	01b	New interrupt is asserted (OK).
10b	01b	New interrupt is asserted (OK).
11b	01b	Interrupt is waiting for reading (OK).
01b	01b	Interrupt is waiting for reading by the driver for more than one read cycle (not OK). Possible drive hang state.
Don't Care	11b	Previous interrupt was read and current interrupt is pending (OK).
Don't Care	10b	Interrupt is not pending (OK).

BMC reads should consider the time it takes for the LAN device driver to deal with the interrupt (in μ s). Note that excessive reads by the BMC can give false indications.

10.5.11.2.3 Get System MAC Address

The Get System MAC Address returns the system MAC address over to the SMBus. This command is a single-fragment Read Block transaction that returns the following the MAC address configured in RAL0, RAH0 registers.

Get system MAC address format:

Function	Command
Get System MAC Address	0xD4

Data returned from the Intel® Xeon® Processor D-1500 Product Family LAN controller:



Function	Byte Count	Data 1 (Op-Code)	Data 2	...	Data 7
Get System MAC Address	7	0xD4	MAC Address MSB	...	MAC Address LSB

10.5.11.2.4 Read Management Parameters

In order to read the management parameters the BMC should execute two SMBus transactions. The first transaction is a block write that sets the parameter that the BMC wants to read. The second transaction is block read that reads the parameter.

Block write transaction:

Function	Command	Byte Count	Data 1
Management Control Request	0xC1	1	Parameter Number

Following the block write the BMC should issue a block read that reads the parameter that was set in the Block Write command:

Function	Command
Read Management Parameter	0xD1

Data returned:

Function	Byte Count	Data 1 (Op-Code)	Data 2	Data 3	...	Data N
Read Management Parameter	N	0xD1	Parameter Number	Parameter Dependent		

The returned data is in the same format of the BMC command.

The returned data is as follows:

Parameter	#	Parameter Data
Keep PHY Link Up	0x00	A single byte parameter: Data 2 — Bit 0 set to indicate that the PHY link for this port should be kept up. Sets the keep_PHY_link_up bit. When cleared, clears the keep_PHY_link_up bit. Bit [7:1] Reserved.
Wrong Parameter Request	0xFE	Returned by the Intel® Xeon® Processor D-1500 Product Family LAN controller only. This parameter is returned on read transaction, if in the previous read command the BMC sets a parameter that is not supported by the Intel® Xeon® Processor D-1500 Product Family LAN controller.
Intel® Xeon® Processor D-1500 Product Family LAN Controller is Not Ready	0xFF	Returned by the Intel® Xeon® Processor D-1500 Product Family LAN controller only, on read parameters command when the data that should have been read is not ready. This parameter has no data. The BMC should retry the read transaction. This value is also returned if the byte count is illegal or if the read command is not preceded by a write command.



The parameter that is returned might not be the parameter requested by the BMC. The BMC should verify the parameter number (default parameter to be returned is 0x1).

If the parameter number is 0xFF, it means that the data that was requested from the Intel® Xeon® Processor D-1500 Product Family LAN controller is not ready yet, or that the adequate Write command was not given. The BMC should retry the read transaction or send the write transaction.

It is responsibility of the BMC to follow the procedure previously defined. When the BMC sends a Block Read command (as previously described) that is not preceded by a Block Write command with bytcount=1, the Intel® Xeon® Processor D-1500 Product Family LAN controller sets the parameter number in the read block transaction to be 0xFF.

10.5.11.2.5 Read Management Receive Filter Parameters

In order to read the management receive filter parameters, the BMC should execute two SMBus transactions. The first transaction is a block write that sets the parameter that the BMC wants to read. The second transaction is block read that read the parameter.

Block write transaction:

Function	Command	Byte Count	Data 1	Data 2
Update MNG RCV Filter Parameters	0xCC	1 or 2	Parameter Number	Parameter Data

The different parameters supported for this command are the same as the parameters supported for update management receive filter parameters.

Following the block write the BMC should issue a block read that reads the parameter that was set in the Block Write command:

Function	Command
Request MNG RCV Filter Parameters	0xCD

Data returned from the Intel® Xeon® Processor D-1500 Product Family LAN controller:

Function	Byte Count	Data 1 (Op-Code)	Data 2	Data 3	...	Data N
Read MNG RCV Filter Parameters	N	0xCD	Parameter Number	Parameter Dependent		

The parameter that is returned might not be the parameter requested by the BMC. The BMC should verify the parameter number (default parameter to be returned is 0x1).

If the parameter number is 0xFF, it means that the data that was requested from the Intel® Xeon® Processor D-1500 Product Family LAN controller is not ready yet or that the adequate Write command was not given. The BMC should retry the read transaction or send the write transaction.

It is BMC responsibility to follow the procedure previously defined. When the BMC sends a Block Read command (as previously described) that is not preceded by a Block Write command with bytcount=1, the Intel® Xeon® Processor D-1500 Product Family LAN controller sets the parameter number in the read block transaction to be 0xFF.



Parameter	#	Parameter Data
Filters Enable	0x01	None.
MNGONLY Configuration	0x0F	None.
Flex Filter Enable Mask and Length	0x10	None.
Flex Filter Data	0x11	Data 2 — Group of Flex Filter's Bytes: 0x0 = bytes 0-29. 0x1 = bytes 30-59. 0x2 = bytes 60-89. 0x3 = bytes 90-119. 0x4 = bytes 120-127.
Decision Filters	0x61	This command is obsolete. Please use 0x68 instead.
VLAN Filters	0x62	One byte to define the accessed VLAN tag filter (MAVTV). Data 2 — VLAN Filter number.
Flex Ports Filters	0x63	One byte to define the accessed manageability flex port filter (MFUTP). Data 2 — Flex Port Filter number.
IPv4 Filter	0x64	One byte to define the accessed IPv4 address filter (MIPAF4). Data 2 — IPv4 address filter number.
IPv6 Filters	0x65	One byte to define the accessed IPv6 address filter (MIPAF6). Data 2 — Pv6 address filter number.
MAC Filters	0x66	One byte to define the accessed MAC address filters pair (MMAL, MMAH). Data 2 — MAC address filters pair number (0-3).
EtherType Filters	0x67	1 byte to define Ethertype filters (METF). Data 2 — METF filter index (valid values are 0 - 3).
Extended Decision Filter	0x68	1 byte to define the extended decisions filters (MDEF_EXT & MDEF). Data 2 — MDEF filter index (valid values are 0 - 5).
Wrong parameter request	0xFE	Returned by the Intel® Xeon® Processor D-1500 Product Family LAN controller only. This parameter is returned on read transaction, if in the previous read command the BMC sets a parameter that is not supported by the Intel® Xeon® Processor D-1500 Product Family LAN controller.
Intel® Xeon® Processor D-1500 Product Family LAN Controller is not ready	0xFF	Returned by the Intel® Xeon® Processor D-1500 Product Family LAN controller only, on read parameters command when the data that should have been read is not ready. This parameter has no data. This value is also returned if the byte count is illegal or if the read command is not preceded by a write command.


10.5.11.2.6 Get Controller Information Command

The BMC uses this command to get the controller identification. Each parameter is returned using a different OpCode.

In order to read the controller information, the BMC should execute two SMBus transactions. The first transaction is a block write that sets the parameter that the BMC wants to read. The second transaction is block read that read the parameter.

Block write transaction:

Function	Command	Byte Count	Data 1 (OpCode)
Get Controller Information	0xD5	1	Parameter Number



Following the block write the BMC should issue a block read that reads the parameter that was set in the Block Write command:

Function	Command
Get Controller Information	0xD5

Data returned from the Intel® Xeon® Processor D-1500 Product Family LAN controller:

Function	Byte Count	Command	Data 2 (Opcode)	Data 3 -n
Get Controller Information	Per Table 10.22	0xD5	Per Table 10.22	See Table 10.22 for the data for each OpCode.

Table 10.22. Get Controller Information Data

Opcode	Byte Count	Description	Notes
0x00	5	Data 4:3: Generic device ID. Data 5: Silicon Revision (RevID).	This is the hardware default value, not any value programmed via NVM.
0x0B	4	Data 4:3 NVM Image version.	
0x0C	6	Data 6:3: Firmware ROM Internal version.	
0x0D	6	Data 6:3: Firmware Flash Internal version.	
0x0E	4	Data 4:3: PXE Firmware version.	MajorVersion.MinorVersion.Build.SubBuild.
0x0F	4	Data 4:3: iSCSI Firmware version.	
0x10	4	Data 4:3: uEFI Firmware version.	
0x16	4	Reserved.	
0x17	6	Data 6:3: Firmware Mini Loader Internal version.	
0xFE	2	Wrong parameter request.	Returned by the Intel® Xeon® Processor D-1500 Product Family LAN controller only. This parameter is returned on read transaction, if in the previous read command the BMC sets a parameter that is not supported by the Intel® Xeon® Processor D-1500 Product Family LAN controller.
0xFF	2	Broadwell-DE LAN controller is not ready.	Returned by the Intel® Xeon® Processor D-1500 Product Family LAN controller only, on read parameters command when the data that should have been read is not ready. This parameter has no data. The BMC should retry the read transaction. This value is also returned if the byte count is illegal or if the Read command is not preceded by a Write command.

10.5.11.2.7 Get Common Filters Command

The BMC uses this command to get the common filters setting. This data can be configured when using Set Common Filters command. The first transaction is a block write that alerts that the BMC wants to read the filters configuration. The second transaction is block read that read the configuration.



Block write transaction:

Function	Command	Byte count	Data
Get Common Filters	0xD3	1	0x00

Following the block write the BMC should issue a block read that reads the filter settings:

Function	Command
Get Common Filters	0xD3

Data returned from the Intel® Xeon® Processor D-1500 Product Family LAN controller:

Function	Byte Count	Command	Data 1	Data 2:4	5:10	Data 11	Data 12	Data 13	Data 14:17
Get Common Filters	18	0xD3	0	Receive Control - see Table 10.13	MAC Address	BMC Alert Address	Interface Data Byte	Alert Value Byte	IPv4 Address

Function	Byte Count	Command	Data 1	Data 2:4	5:10	Data 11	Data 12	Data 13	Data 14:29
Get Common Filters	30	0xD3	0	Receive Control - see Table 10.13	MAC Address	BMC Alert Address	Interface Data Byte	Alert Value Byte	IPv6 Address

If case of error the following answers might be returned:

Function	Command	Byte Count	Data 1
Get Common Filters	0xD3	1	0xFF

This response is by the Intel® Xeon® Processor D-1500 Product Family LAN controller, on read common filter command when the data that should have been read is not ready. This parameter has no data. The BMC should retry the read transaction.

This value is also returned if the byte count is illegal, if the Read command is not preceded by a Write command, or if the filters were not previously configured with a Set Common Filters command (Section 10.5.11.1.7).



10.5.11.2.8 Read MACSec Parameters

The BMC uses this command to get the MACSec setting. The first transaction is a block write that alerts that the BMC wants to read the MACSec settings. The second transaction is block read that read the settings.

Block write transaction:

Function	Command	Byte Count	Data
Read MACSec Parameters	0xC9	1	Parameter Number (PN#)

Following the block write, the BMC uses this command to read the MACsec status & parameters.

Function	Command	Byte Count	Data 1	Data 2 – n
Read MACsec parameters	0xD9	2,18 or 22	Parameter number (PN#)	Parameter data

The table below shows the different parameters and their contents:

Parameter	PN#	Parameter Data
MACsec Interrupt Cause	0x0	This command shall return 1 byte (Data2). This byte contains the MACsec Interrupt cause, according to the following values: Data2: Bit 0 – TX Key Packet Number (PN) threshold met Bit 1 – Host requested ownership Bit 2 – Host released ownership Bit 3...31 – Reserved
MACsec RX parameters	0x1	Data 2: Reserved Data 3: MACsec Ownership Status. See Table 10.23 Data 4: MACsec Host Control Status. See Table 10.24 Data 5: RX Port Identifier (MSB) Data 6: RX Port Identifier (LSB) Data 7: RX SCI (MSB) ... Data 12: (LSB) Data 13: Reserved Data 14: RX SA AN — The association number currently used for the active SA Data 15: RX SA Packet Number (MSB) ... Data 18: (LSB) RX SA Packet Number is the last packet number, as read from the last valid RX MACsec packet



Parameter	PN#	Parameter Data
MACsec TX parameters	0x2	Data 2: Reserved Data 3: MACsec Ownership Status. See Table 10.23 below Data 4: MACsec Host Control Status. See Table 10.24 below Data 5: TX Port Identifier (MSB) Data 6: TX Port Identifier (LSB) Note: TX Port Identifier is reserved to 0x0 for this implementation. Data 7: TX SCI (MSB) ... Data 12: (LSB) Data 13: Reserved Data 14: TX SA AN — The association number currently used for the active SA Data 15: TX SA Packet Number (MSB) ... Data 18: (LSB) Data 19: Packet Number Threshold (MSB) ... Data 21: (LSB) TX SA Packet Number is the last packet number, as read from the last valid TX MACsec packet. Data 22: TX Control Status. See Table 10.25 below

Table 10.23. MACsec Owner Status

Value	Description
0x0	Host is MACsec owner
0x1	BMC is MACsec owner

Table 10.24. MACsec Host Control Status

Bit	Description
0	Reserved
1	Allow Host Traffic: 0b = Host traffic is blocked 1b = Host traffic is allowed
2..7	Reserved

Table 10.25. TX Control Status:

Bit	Description
0..4	Reserved
5	Include SCI: 0b = Do not include SCI in TX packets 1b = Include SCI in TX packets
6..7	Reserved

10.5.11.2.9 Read Receive Enable Configuration

The BMC uses this command to read the receive configuration data. This data can be configured when using Receive Enable command or through the NVM.

Read Receive Enable Configuration command format (SMBus Read Block) is as follows:

Function	Command
Read Receive Enable	0xDA



Data returned from the Intel® Xeon® Processor D-1500 Product Family LAN controller:

Function	Byte Count	Data 1 (Op-Code)	Data 2	Data 3	...	Data 8	Data 9	...	Data 12	Data 13	Data 14	Data 15
Read Receive Enable	15 (0x0F)	0xDA	Receive Control Byte	MAC Addr MSB	...	MAC Addr LSB	IP Addr MSB	...	IP Addr LSB	BMC SMBus Addr	I/F Data Byte	Alert Value Byte

The detailed description of each field is specified in the receive enable command description in [Section 10.5.11.1.3](#).

10.5.12 Example Configuration Steps

This section provides sample configuration settings for common filtering configurations. Three examples are presented. The examples are in pseudo code format, with the name of the SMBus command followed by the parameters for that command and an explanation.

10.5.12.1 Example 1 - Shared MAC and RMCP Only Ports

This example is the most basic configuration. The MAC address filtering is shared with the host operating system and only traffic directed the RMCP ports (0x26F and 0x298) is filtered. For this example, the BMC must issue gratuitous ARPs because no filter is enabled to pass ARP requests to the BMC.

10.5.12.1.1 Example 1 Pseudo Code

Step 1: Disable existing filtering.

Receive Enable[00]

Utilizing the simple form of the Receive Enable command, this prevents any packets from reaching the BMC by disabling filtering:

Receive Enable Control 00h:

- Bit 0 [0] – Disable Receiving of packets

Step 2: Configure MDEF[0].

Update Manageability Filter Parameters [68, 0, C0000000, 00000000]

Use the Update Manageability Filter Parameters command to update Decision Filters (MDEF) (parameter 68h). This will update MDEF[0], as indicated by the 2nd parameter (0).

MDEF[0] value of C0000000h:

- Bit 30 [1] – port 0x298
- Bit 31 [1] – port 0x26F

MDEF_EXT[0] value of 0x0000000:

Step 3: Configure MNGONLY.

Update Manageability Filter Parameters [F, 0, 00000001]

Use the Update Manageability Filter Parameters command to update Manageability Only (MNGONLY) (parameter 0xF) so that port 0x298 and 0x26F would not be sent to the host.

- Bit [0] - MDEF[0] is exclusive to the BMC.



Step 4: - Enable Filtering.

Receive Enable [05]

Using the simple form of the Receive Enable command:

Receive Enable Control 0x05:

- Bit 0 [1] – Enable Receiving of packets
- Bit 2 [1] – Enable status reporting (such as link lost)
- Bit 5:4 [00] – Notification method = SMB Alert
- Bit 7 [0] – Use shared MAC

The resulting MDEF filters are as follows:

Table 10.26. Example 1 MDEF Results

		Manageability Decision Filter (MDEF)							
Filter		0	1	2	3	4	5	6	7
L2 Unicast Address[3:0]	AND								
Broadcast	AND								
Manageability VLAN[7:0]	AND								
IPv6 Address[3:0]	AND								
IPv4 Address[3:0]	AND								
L2 Unicast Address[3:0]	OR								
Broadcast	OR								
Multicast	AND								
ARP Request	OR								
ARP Response	OR								
Neighbor Discovery	OR								
Port 0x298	OR	X							
Port 0x26F	OR	X							
Flex Port 7:0	OR								
Flex TCO	OR								

10.5.12.2 Example 2 - Dedicated MAC, Auto ARP Response and RMCP Port Filtering

This example shows a common configuration; the BMC has a dedicated MAC and IP address. Automatic ARP responses is enabled as well as RMCP port filtering. By enabling Automatic ARP responses the BMC is not required to send the gratuitous ARPs as it did in Example 1.

For demonstration purposes, the dedicated MAC address is calculated by reading the system MAC address and adding 1 do it, assume the System MAC is AABBCDC. The IP address for this example will be 1.2.3.4. Additionally, the XSUM filtering is enabled.

Note that not all Intel Ethernet Controllers support automatic ARP responses, please refer to product specific documentation.

10.5.12.2.1 Example 2 - Pseudo Code

Step 1: Disable existing filtering.



Receive Enable[00]

Utilizing the simple form of the Receive Enable command, this prevents any packets from reaching the BMC by disabling filtering:

Receive Enable Control 0x00:

- Bit 0 [0] – Disable Receiving of packets

Step 2: Read system MAC address.

Get System MAC Address []

Reads the system MAC address. Assume returned AABBCDC for this example.

Step 3: Configure XSUM filter.

Update Manageability Filter Parameters [01, 00800000]

Use the Update Manageability Filter Parameters command to update filters enable settings (parameter 1). This set the Manageability Control (MANC) register.

MANC Register 0x00800000:

- Bit 23 [1] - XSUM Filter enable

Note that some of the following configuration steps manipulate the MANC register indirectly, this command sets all bits except XSUM to 0b. It is important to either do this step before the others, or to read the value of the MANC and then write it back with only bit 32 changed. Also note that the XSUM enable bit might differ between Ethernet controllers, refer to product specific documentation.

Step 4: Configure MDEF[0].

Update Manageability Filter Parameters [68, 0, C0000000, 00000000]

Use the Update Manageability Filter Parameters command to update Decision Filters (MDEF) (parameter 0x68). This updates MDEF[0], as indicated by the 2nd parameter (0).

MDEF value of 0x00000C00:

- Bit 30 [1] – port 0x298
- Bit 31 [1] – port 0x26F

MDEF_EXT[0] value of 0x00000000:

Step 5: Configure MDEF[1].

Update Manageability Filter Parameters [68, 1, 10000000, 00000000]

Use the Update Manageability Filter Parameters command to update Decision Filters (MDEF) (parameter 0x61). This updates MDEF[1], as indicated by the 2nd parameter (1).

MDEF value of 10000000:

- Bit 28 [1] – ARP Requests

MDEF_EXT[1] value of 0x00000000:

When enabling automatic ARP responses, the ARP requests still go into the manageability filtering system and as such needs to be designated as also needing to be sent to the host. For this reason a separate MDEF is created with only ARP request filtering enabled.

Refer to the next step for more details.

Step 6: Configure Manageability only.

Update Manageability Filter Parameters [F, 0, 00000001]



Use the Update Manageability Filter Parameters command to update Manageability Only (MNGONLY) (parameter 0xF) so that port 0x298 and 0x26F would not be sent to the host.

- Bit [0] - MDEF[0] is exclusive to the BMC.

This enables ARP requests to be passed to both manageability and to the host. Specified separate MDEF filter for ARP requests. If ARP requests had been added to *MDEF[0]* and then *MDEF[0]* specified in management only configuration then not only would RMCP traffic (ports 0x26F and 0x298) be sent only to the BMC, ARP requests would have also been sent to the BMC only.

Step 7: Enable filtering.

Receive Enable [8D, AABBCDD, 01020304, 00, 00, 00]

Using the advanced version Receive Enable command, the first parameter:

Receive Enable Control 0x8D:

- Bit 0 [1] – Enable receiving of packets
- Bit 2 [1] – Enable status reporting (such as link lost)
- Bit 3 [1] – Enable automatic ARP responses
- Bit 5:4 [00] – Notification method = SMB alert
- Bit 7 [1] - Use dedicated MAC

Second parameter is the MAC address (AABBCDD).


Third parameter is the IP address(01020304).

The last three parameters are zero when the notification method is SMB alert.

The resulting MDEF filters are as follows:

Table 10.27. Example 2 MDEF Results

		Manageability Decision Filter (MDEF)							
Filter		0	1	2	3	4	5	6	7
L2 Unicast Address[3:0]	AND								
Broadcast	AND								
Manageability VLAN[7:0]	AND								
IPv6 Address[3:0]	AND								
IPv4 Address[3:0]	AND								
L2 Unicast Address[3:0]	OR								
Broadcast	OR								
Multicast	AND								
ARP Request	OR		X						
ARP Response	OR								
Neighbor Discovery	OR								
Port 0x298	OR	X							
Port 0x26F	OR	X							
Flex Port 7:0	OR								
Flex TCO	OR								



10.5.12.3 Example 3 - Dedicated MAC and IP Address

This example provides the BMC with a dedicated MAC and IP address and enables it to receive ARP requests. The BMC is then responsible for responding to ARP requests.

For demonstration purposes, the dedicated MAC address is calculated by reading the system MAC address and adding one to it (assume the system MAC is AABBCDCD). The IP address for this example is 1.2.3.4. For this example, the Receive Enable command is used to configure the MAC address filter.

In order for the BMC to be able to receive ARP requests, it needs to specify a filter for this, and that filter needs to be included in the manageability-to-host filtering so that the host operating system might also receive ARP requests.

10.5.12.3.1 Example 3 - Pseudo Code

Step 1: Disable existing filtering.

Receive Enable[00]

Utilizing the simple form of the Receive Enable command, this prevents any packets from reaching the BMC by disabling filtering:

Receive enable control 0x00:

- Bit 0 [0] – Disable receiving of packets

Step 2: Read System MAC Address.

Get System MAC Address []

Reads the system MAC address. Assume returned AABBCDCD for this example.

Step 3: Configure IP Address Filter

Update Manageability Filter Parameters [64, 00, 01020304]

Use the update manageability filter parameters to configure an IPv4 filter.

The 1st parameter (0x64) specifies that we are configuring an IPv4 filter.

The 2nd parameter (0x00) indicates which IPv4 filter is being configured, in this case filter 0.

The 3rd parameter is the IP address – 1.2.3.4.

Step 4: Configure MAC address filter.

Update Manageability Filter Parameters [66, 00, AABBCDD]

Use the update manageability filter parameters to configure a MAC address filter.

The 1st parameter (0x66) specifies that we are configuring a MAC address filter.

The 2nd parameter (0x00) indicates which MAC address filter is being configured, in this case filter 0.

The 3rd parameter is the MAC Address - AABBCDD

Step 5: Configure MDEF[0] for IP and MAC filtering.

Update Manageability Filter Parameters [68, 0, 00002001, 00000000]

Use the Update Manageability Filter Parameters command to update Decision Filters (MDEF) (parameter 0x68). This updates MDEF[0], as indicated by the 2nd parameter (0).

MDEF value of 00002001:



- Bit 0 [1] – MAC[0] address filtering
- Bit 13 [1] – IP[0] address filtering

MDEF_EXT[0] value of 0x0000000:

Step 6: Configure MDEF[1].

Update Manageability Filter Parameters [68, 1, 10000000]

Use the Update Manageability Filter Parameters command to update Decision Filters (MDEF) (parameter 0x68). This updates MDEF[1], as indicated by the 2nd parameter (1).

MDEF value of 10000000:

- Bit 28 [1] – ARP requests

MDEF_EXT[1] value of 0x0000000:

Step 7: Configure the management-to-host filter

Update Manageability Filter Parameters [F, 0, 00000001]

Use the Update Manageability Filter Parameters command to update Manageability Only (MNGONLY) (parameter 0xF) so that the dedicated MAC/IP traffic would not be sent to the host. Note that given the host does not program this address in it’s L2 filtering, this step is not a must, unless the host chooses to work in promiscuous mode.

- Bit [0] - MDEF[0] is exclusive to the BMC.

Step 8: Enable filtering.

Receive Enable [05]

Using the simple form of the Receive Enable command,:

Receive Enable Control 0x05:

- Bit 0 [1] – Enable Receiving of packets
- Bit 2 [1] – Enable status reporting (such as link lost)
- Bit 5:4 [00] – Notification method = SMB alert

The resulting MDEF filters are as follows:

Table 10.28. Example 3 MDEF Results (Sheet 1 of 2)

		Manageability Decision Filter (MDEF)							
Filter		0	1	2	3	4	5	6	7
L2 Unicast Address[3:0]	AND	0001							
Broadcast	AND								
Manageability VLAN[7:0]	AND								
IPv6 Address[3:0]	AND								
IPv4 Address[3:0]	AND	0001							
L2 Unicast Address[3:0]	OR								
Broadcast	OR								
Multicast	AND								
ARP Request	OR		X						
ARP Response	OR								
Neighbor Discovery	OR								



Table 10.28. Example 3 MDEF Results (Sheet 2 of 2)

		Manageability Decision Filter (MDEF)							
Filter		0	1	2	3	4	5	6	7
Port 0x298	OR								
Port 0x26F	OR								
Flex Port 7:0	OR								
Flex TCO	OR								

10.5.12.4 Example 4 - Dedicated MAC and VLAN Tag

This example shows an alternate configuration; the BMC has a dedicated MAC and IP address, along with a VLAN tag of 0x32 is required for traffic to be sent to the BMC. This means that all traffic with VLAN a matching tag is sent to the BMC.

For demonstration purposes, the dedicated MAC address is calculated by reading the system MAC address and adding 1 to it (assume the system MAC is AABBCDC). The IP address for this example is 1.2.3.4 and the VLAN tag is 0x0032.

Additionally, the XSUM filtering is enabled.

10.5.12.4.1 Example 4 - Pseudo Code

Step 1: Disable existing filtering.

Receive Enable[00]

Utilizing the simple form of the Receive Enable command, this prevents any packets from reaching the BMC by disabling filtering:

Receive enable control 0x00:

- Bit 0 [0] – Disable receiving of packets

Step 2: - Read system MAC address.

Get System MAC Address []

Reads the system MAC address. Assume returned AABBCDC for this example.

Step 3: Configure XSUM filter.

Update Manageability Filter Parameters [01, 00800000]

Use the Update Manageability Filter Parameters command to update Filters Enable settings (parameter 1). This sets the Manageability Control (MANC) register.

MANC register 0x00800000:

- Bit 23 [1] – XSUM filter enable

Note that some of the following configuration steps manipulate the MANC register indirectly, this command sets all bits except XSUM to 0b. It is important to either do this step before the others, or to read the value of the MANC and then write it back with only bit 32 changed. Also note that the XSUM enable bit might differ between Ethernet controllers, refer to product specific documentation.

Step 4: Configure VLAN 0 filter.

Update Manageability Filter Parameters [62, 0, 0032]



Use the Update Manageability Filter Parameters command to configure VLAN filters. Parameter 0x62 indicates update to VLAN filter, the 2nd parameter indicates which VLAN filter (0 in this case), the last parameter is the VLAN ID (0x0032).

Step 5: Configure MDEF[0].

Update Manageability Filter Parameters [68, 0, 00000020, 00000000]

Use the Update Manageability Filter Parameters command to update Decision Filters (MDEF) (parameter 0x68). This updates MDEF[0], as indicated by the 2nd parameter (0).

MDEF value of 00000020:

- Bit 5 [1] – VLAN[0] AND

MDEF_EXT[0] value of 0x00000000:

Step 6: Enable filtering.

Receive Enable [85, AABBCDD, 01020304, 00, 00, 00]

Using the advanced version Receive Enable command, the first parameter:

Receive Enable Control 0x85:

- Bit 0 [1] – Enable receiving of packets
- Bit 2 [1] – Enable status reporting (such as link lost)
- Bit 5:4 [00] – Notification method = SMB alert
- Bit 7 [1] – Use dedicated MAC

Second parameter is the MAC address: AABBCDD.

Third parameter is the IP address: 01020304.

The last three parameters are zero when the notification method is SMBus alert.

The resulting MDEF filters are as follows:

Table 10.29. Example 4 MDEF Results (Sheet 1 of 2)

		Manageability Decision Filter (MDEF)							
Filter		0	1	2	3	4	5	6	7
L2 Unicast Address[3:0]	AND								0001
Broadcast	AND								
Manageability VLAN[7:0]	AND	X							
IPv6 Address[3:0]	AND								
IPv4 Address[3:0]	AND								
L2 Unicast Address[3:0]	OR								
Broadcast	OR								
Multicast	AND								
ARP Request	OR								
ARP Response	OR								
Neighbor Discovery	OR								
Port 0x298	OR								



Table 10.29. Example 4 MDEF Results (Sheet 2 of 2)

		Manageability Decision Filter (MDEF)							
Filter		0	1	2	3	4	5	6	7
Port 0x26F	OR								
Flex Port 7:0	OR								
Flex TCO	OR								

10.5.13 SMBus Troubleshooting

This section outlines the most common issues found while working with pass-through using the SMBus sideband interface.

10.5.13.1 TCO Alert Line Stays Asserted After a Power Cycle

After the Intel® Xeon® Processor D-1500 Product Family LAN controller resets, all its ports indicates a status change. If the BMC only reads status from one port (slave address), the other ones continue to assert the TCO alert line.

Ideally, the BMC should use the ARA transaction (see [Section 10.5.10](#)) to determine which slave asserted the TCO alert. Many customers only wish to use one port for manageability thus using ARA might not be optimal.

An alternate to using ARA is to configure part of the ports to not report status and to set its SMBus timeout period. In this case, the SMBus timeout period determines how long a port asserts the TCO alert line awaiting a status read from a BMC; by default this value is zero (indicates an infinite timeout).

The SMBus configuration section of the NVM has a SMBus notification timeout (ms) field that can be set to a recommended value of 0xFF (for this issue). Note that this timeout value is for all slave addresses. Along with setting the SMBus notification timeout to 0xFF, it is recommended that the other ports be configured in the NVM to disable status alerting. This is accomplished by having the *Enable Status Reporting* bit set to 0b for the desired ports in the LAN configuration section of the NVM.

The third solution for this issue is to have the BMC hard-code the slave addresses to always read from all ports. As with the previous solution, it is recommend that the other ports have status reporting disabled.

10.5.13.2 When SMBus Commands Are Always NACKed

There are several reasons why all commands sent to the Intel® Xeon® Processor D-1500 Product Family LAN controller from a BMC could be NACKed. The following are most common:

- Invalid NVM Image — The image itself might be invalid or it could be a valid image and is not a pass-through image, as such SMBus connectivity is disabled.
- The BMC is not using the correct SMBus address — Many BMC vendors hard-code the SMBus address(es) into their firmware. If the incorrect values are hard-coded, the Intel® Xeon® Processor D-1500 Product Family LAN controller does not respond.
 - The SMBus address(es) can be dynamically set using the SMBus ARP mechanism.
- The BMC is using the incorrect SMBus interface — The NVM might be configured to use one physical SMBus port; however, the BMC is physically connected to a different one.
- Bus Interference — The bus connecting the BMC and the Intel® Xeon® Processor D-1500 Product Family LAN controller might be unstable.



10.5.13.3 SMBus Clock Speed Is 16.6666 KHz

This can happen when the SMBus connecting the BMC and the Intel® Xeon® Processor D-1500 Product Family LAN controller is also tied into another device (such as an ICH) that has a maximum clock speed of 16.6666 KHz. The solution is to not connect the SMBus between the Intel® Xeon® Processor D-1500 Product Family LAN controller and the BMC to this device.

10.5.13.4 A Network Based Host Application Is Not Receiving Any Network Packets

Reports have been received about an application not receiving any network packets. The application in question was NFS under Linux. The problem was that the application was using the RMPC/RMCP+ IANA reserved port 0x26F (623) and the system was also configured for a shared MAC and IP address with the operating system and BMC.

The management control-to-host configuration, in this situation, was setup not to send RMCP traffic to the operating system (this is typically the correct configuration). This means that no traffic sent to port 623 was being routed.

The solution in this case is to configure the problematic application NOT to use the reserved port 0x26F.

10.5.13.5 Unable to Transmit Packets from the BMC

If the BMC has been transmitting and receiving data without issue for a period of time and then begins to receive NACKs from the Intel® Xeon® Processor D-1500 Product Family LAN controller when it attempts to write a packet, the problem is most likely due to the fact that the buffers internal to the Intel® Xeon® Processor D-1500 Product Family LAN controller are full of data that has been received from the network but has yet to be read by the BMC.

Being an embedded device, the Intel® Xeon® Processor D-1500 Product Family LAN controller has limited buffers that are shared for receiving and transmitting data. If a BMC does not keep the incoming data read, the Intel® Xeon® Processor D-1500 Product Family LAN controller can be filled up. This prevents the BMC from transmitting more data, resulting in NACKs.

If this situation occurs, the recommended solution is to have the BMC issue a Receive Enable command to disable more incoming data, read all the data from the Intel® Xeon® Processor D-1500 Product Family LAN controller, and then use the Receive Enable command to enable incoming data.

10.5.13.6 SMBus Fragment Size

The SMBus specification indicates a maximum SMBus transaction size of 32 bytes. Most of the data passed between the Intel® Xeon® Processor D-1500 Product Family LAN controller and the BMC over the SMBus is RMCP/RMCP+ traffic, which by its very nature (UDP traffic) is significantly larger than 32 bytes in length. Multiple SMBus transactions might therefore be required to move data from the Intel® Xeon® Processor D-1500 Product Family LAN controller to the BMC or to send a data from the BMC to the Intel® Xeon® Processor D-1500 Product Family LAN controller.

Recognizing this bottleneck, the Intel® Xeon® Processor D-1500 Product Family LAN controller handles up to 240 bytes of data in a single transaction. This is a configurable setting in the NVM. The default value in the NVM images is 32, per the SMBus specification. If performance is an issue, increase this size.



During initialization, firmware within the Intel® Xeon® Processor D-1500 Product Family LAN controller allocates buffers based upon the SMBus fragment size setting within the NVM. The Intel® Xeon® Processor D-1500 Product Family LAN controller firmware has a finite amount of RAM for its use: the larger the SMBus fragment size, the fewer buffers it can allocate. Because this is true, BMC implementations must take care to send data over the SMBus efficiently.

For example, the Intel® Xeon® Processor D-1500 Product Family LAN controller firmware has 3 KB of RAM it can use for buffering SMBus fragments. If the SMBus fragment size is 32 bytes then the firmware could allocate 96 buffers of size 32 bytes each. As a result, the BMC could then send a large packet of data (such as KVM) that is 800 bytes in size in 25 fragments of size 32 bytes apiece.

However, this might not be the most efficient way because the BMC must break the 800 bytes of data into 25 fragments and send each one at a time.

If the SMBus fragment size is changed to 240 bytes, the Intel® Xeon® Processor D-1500 Product Family LAN controller firmware can create 12 buffers of 240 bytes each to receive SMBus fragments. The BMC can now send that same 800 bytes of KVM data in only four fragments, which is much more efficient.

The problem of changing the SMBus fragment size in the NVM is if the BMC does not also reflect this change. If a programmer changes the SMBus fragment size in the Intel® Xeon® Processor D-1500 Product Family LAN controller to 240 bytes and then wants to send 800 bytes of KVM data, the BMC can still only send the data in 32 byte fragments. As a result, firmware runs out of memory.

This is because firmware created the 12 buffers of 240 bytes each for fragments; however, the BMC is only sending fragments of size 32 bytes. This results in a memory waste of 208 bytes per fragment. Then when the BMC attempts to send more than 12 fragments in a single transaction, the Intel® Xeon® Processor D-1500 Product Family LAN controller NACKs the SMBus transaction due to not enough memory to store the KVM data.

In summary, if a programmer increases the size of the SMBus fragment size in the NVM (recommended for efficiency purposes) take care to ensure that the BMC implementation reflects this change and uses that fragment size to its fullest when sending SMBus fragments.

10.5.13.7 Losing Link

Normal behavior for the Intel® Xeon® Processor D-1500 Product Family LAN controller when the system powers down or performs a reset is for the link to temporarily go down and then back up again to re-negotiate the link speed. This behavior can have adverse affects on manageability.

For example if there is an active FTP or serial over LAN session to the BMC, this connection might be lost. In order to avoid this possible situation, the BMC can use the Management Control command detailed in [Section 10.5.11.1.5](#) to ensure the link stays active at all times.

This command is available when using the NC-SI sideband interface as well.

Care should be taken with this command, if the software device driver negotiates the maximum link speed, the link speed remains the same when the system powers down or resets. This might have undesirable power consumption consequences. Currently, when using NC-SI, the BMC can re-negotiate the link speed. That functionality is not available when using the SMBus interface.

10.5.13.8 Enable Checksum Filtering

If checksum filtering is enabled, the BMC does not need to perform the task of checking this checksum for incoming packets. Only packets that have a valid checksum is passed to the BMC. All others are silently discarded.



This is a way to offload some work from the BMC.

10.5.13.9 Still Having Problems?

If problems still exist, contact your field representative. Be prepared to provide the following:

- A SMBus trace if possible.
- A dump of the NVM image. This should be taken from the actual Broadwell-DE LAN controller, rather than the NVM image provided by Intel. Parts of the NVM image are changed after writing (such as the physical NVM size).

10.6 NC-SI Pass Through Interface

The Network Controller Sideband Interface (NC-SI) is a DMTF industry standard protocol for the sideband interface. NC-SI uses a modified version of the industry standard RMII interface for the physical layer as well as defining a new logical layer.

The NC-SI specification supported by the Intel® Xeon® Processor D-1500 Product Family LAN controller can be found at:

http://www.dmtf.org/sites/default/files/standards/documents/DSP0222_1.0.1.pdf

10.6.1 Overview

10.6.1.1 Terminology

The terminology in this document is taken from the NC-SI specification.

Table 10.30. NC-SI Terminology (Sheet 1 of 2)

Term	Definition
Frame Versus Packet	Frame is used in reference to Ethernet, whereas packet is used everywhere else.
External Network Interface	The interface of the network controller that provides connectivity to the external network infrastructure (port).
Internal Host Interface	The interface of the network controller that provides connectivity to the host OS running on the platform.
Management Controller (BMC)	An intelligent entity comprising of HW/Firmware/Software, that resides within a platform and is responsible for some or all management functions associated with the platform (BMC, service processor, etc.).
Network Controller (NC)	The component within a system that is responsible for providing connectivity to the external Ethernet network world.
Remote Media	The capability to allow remote media devices to appear as if they were attached locally to the host.
Network Controller Sideband Interface	The interface of the network controller that provides connectivity to a management controller. It can be shorten to sideband interface as appropriate in the context.
Interface	This refers to the entire physical interface, such as both the transmit and receive interface between the management controller and the network controller.
Integrated Controller	The term integrated controller refers to a network controller device that supports two or more channels for NC-SI that share a common NC-SI physical interface. For example, a network controller that has two or more physical network ports and a single NC-SI bus connection.
Multi-Drop	Multi-drop commonly refers to the case where multiple physical communication devices share an electrically common bus and a single device acts as the master of the bus and communicates with multiple slave or target devices. In NC-SI, a management controller serves the role as the master, and the network controllers are the target devices.



Table 10.30. NC-SI Terminology (Sheet 2 of 2)

Term	Definition
Point-to-Point	Point-to-point commonly refers to the case where only two physical communication devices are interconnected via a physical communication medium. The devices might be in a master/slave relationship, or could be peers. In NC-SI, point-to-point operation refers to the situation where only a single management controller and single network controller package are used on the bus in a master/slave relationship where the management controller is the master.
Channel	The control logic and data paths supporting NC-SI pass-through operation on a single network interface (port). A network controller that has multiple network interface ports can support an equivalent number of NC-SI channels.
Package	One or more NC-SI channels in a network controller that share a common set of electrical buffers and common buffer control for the NC-SI bus. Typically, there will be a single, logical NC-SI package for a single physical network controller package (chip or module). However, the specification allows a single physical chip or module to hold multiple NC-SI logical packages.
Control Traffic/Messages/Packets	Command, response and notification packets transmitted between BMC and the Intel® Xeon® Processor D-1500 Product Family LAN controller for the purpose of managing NC-SI.
Pass-Through Traffic/Messages/Packets	Non-control packets passed between the external network and the BMC through the Intel® Xeon® Processor D-1500 Product Family LAN controller.
Channel Arbitration	Refer to operations where more than one of the network controller channels can be enabled to transmit pass-through packets to the BMC at the same time, where arbitration of access to the RXD, CRS_DV, and RX_ER signal lines is accomplished either by software or hardware means.
Logically Enabled/Disabled NC	Refers to the state of the network controller wherein pass-through traffic is able/unable to flow through the sideband interface to and from the management controller, as a result of issuing Enable/Disable Channel command.
NC RX	Defined as the direction of ingress traffic on the external network controller interface
NC TX	Defined as the direction of egress traffic on the external network controller interface
NC-SI RX	Defined as the direction of ingress traffic on the sideband enhanced NC-SI Interface with respect to the network controller.
NC-SI TX	Defined as the direction of egress traffic on the sideband enhanced NC-SI Interface with respect to the network controller.

10.6.1.2 System Topology

In NC-SI each physical endpoint (NC package) can have several logical slaves (NC channels).

NC-SI defines that one management controller and up to four network controller packages can be connected to the same NC-SI link.



Figure 10.6 shows an example topology for a single BMC and a single NC package. In this example, the NC package has two NC channels.

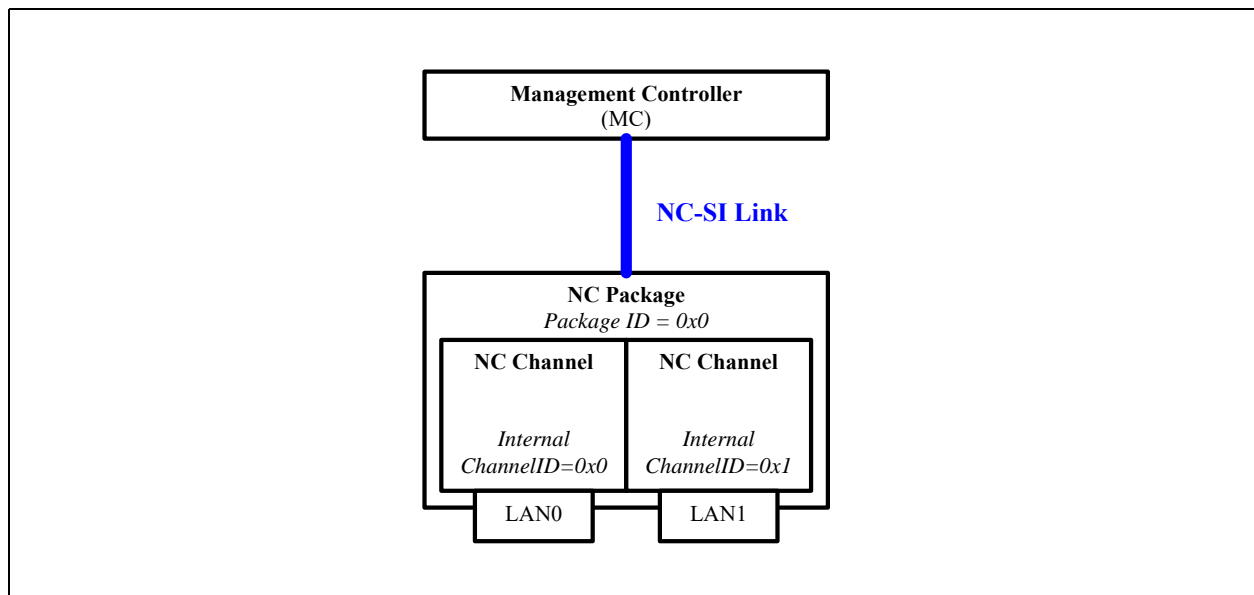


Figure 10.6. Single NC Package, Two NC Channels

Figure 10.7 shows an example topology for a single BMC and two NC packages. In this example, one NC package has two NC channels and the other has only one NC channel. Scenarios in which the NC-SI lines are shared by multiple NCs (Figure 10.7) mandate an arbitration mechanism. The arbitration mechanism is described in Section 10.6.7.1.

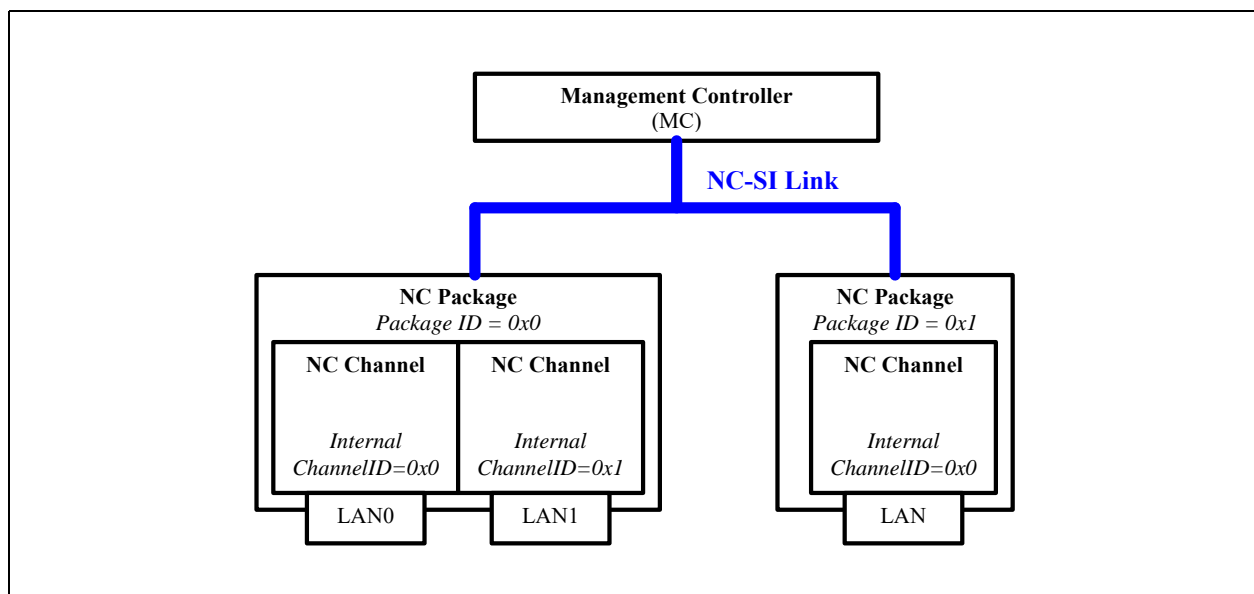


Figure 10.7. Two NC Packages (Left, with Two NC Channels and Right, with One NC Channel)

Note: Channel numbers should match PCI function numbers. If more than one function is defined on a port, the function with the lowest value associated with this port is used (it is assumed



that the first function numbers are assigned to different ports). The association of functions to ports is reflected in the PFGEN_PORTNUM registers.

10.6.1.3 Data Transport

Since NC-SI is based upon the RMIi transport layer, data is transferred in the form of Ethernet frames.

NC-SI defines two types of transmitted frames:

1. Control frames:
 - a. Configures and control the interface.
 - b. Identified by a unique EtherType in their L2 header.
2. Pass-through frames:
 - a. Actual LAN pass-through frames transferred from/to the BMC.
 - b. Identified as not being a control frame.
 - c. Attributed to a specific NC channel by their source MAC address (as configured in the NC by the BMC).

10.6.1.3.1 Control Frames

NC-SI control frames are identified by a unique NC-SI EtherType (0x88F8).

Control frames are used in a single-threaded operation, meaning commands are generated only by the BMC and can only be sent one at a time. Each command from the BMC is followed by a single response from the NC (command-response flow), after which the BMC is allowed to send a new command.

The only exception to the command-response flow is the Asynchronous Event Notification (AEN). These control frames are sent unsolicited from the NC to the BMC.

AEN functionality by the NC must be disabled by default, until activated by the BMC using the Enable AEN commands.

In order to be considered a valid command, a control frame must:

1. Comply with the NC-SI header format.
2. Be targeted to a valid channel in the package via the Package ID and Channel ID fields. For example, to target a NC channel with package ID of 0x2 and internal channel ID of 0x5, the BMC must set the channel ID inside the control frame to 0x45. The channel ID is composed of three bits of package ID and five bits of internal channel ID.
3. Contain a correct payload checksum (if used).
4. Meet any other condition defined by NC-SI.

There are also commands (such as select package) targeted to the package as a whole. These commands must use an internal channel ID of 0x1F.

For details, refer to the NC-SI specification.

10.6.1.3.2 NC-SI Frames Receive Flow



Figure 10.8 shows the flow for frames received on the NC from the BMC.

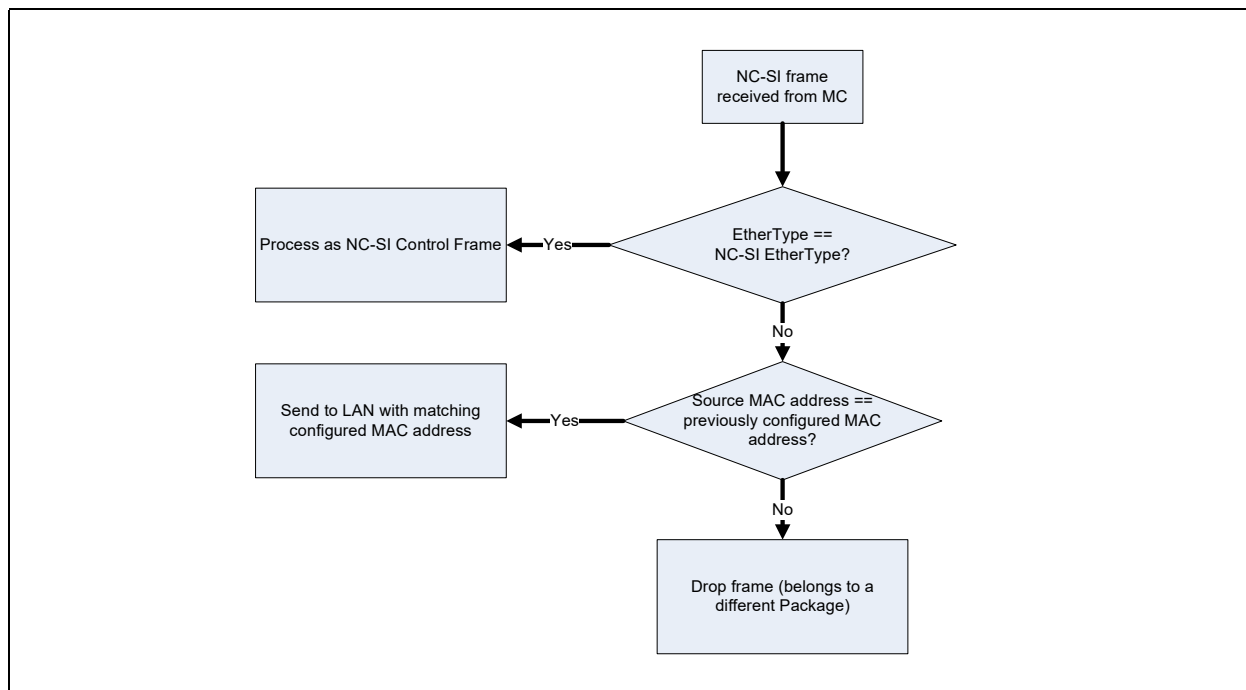


Figure 10.8. NC-SI Frames Receive Flow for the NC

10.6.2 NC-SI Standard Support

10.6.2.1 Supported Features

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports all the mandatory features of the NC-SI specification (rev 1.0.1).

Table 10.31 lists the supported commands.

Table 10.32 lists optional features supported.

Table 10.31. Supported NC-SI Commands (Sheet 1 of 2)

Command	Supported Over RMI I	Supported Over MCTP With Pass Through	Supported Over MCTP Without Pass Through
Clear Initial State	Yes	Yes	Yes
Get Version ID	Yes	Yes	Yes
Get Parameters	Yes	Yes	Yes
Get Controller Packet Statistics	Yes, partially	Yes, partially	Yes, partially
Get Link Status	Yes	Yes	Yes
Enable Channel	Yes	Yes	Yes
Disable Channel	Yes	Yes	Yes
Reset Channel	Yes	Yes	Yes



Table 10.31. Supported NC-SI Commands (Sheet 2 of 2)

Command	Supported Over RMI I	Supported Over MCTP With Pass Through	Supported Over MCTP Without Pass Through
Enable VLAN	Yes ^{1,2}	Yes ¹	No ³
Disable VLAN	Yes	Yes	No ³
Enable Broadcast Filter	Yes	Yes	No ³
Disable Broadcast Filter	Yes	Yes	No ³
Set MAC Address ⁴	Yes	Yes	No ³
Get NC-SI Statistics	Yes, partially	Yes, partially	Yes, partially
Set NC-SI Flow Control	Yes	No	No ³
Set Link Command	Yes	Yes	Yes
Enable Global Multicast Filter	Yes	Yes	No ³
Disable Global Multicast Filter	Yes	Yes	No ³
Get Capabilities	Yes	Yes	Yes
Set VLAN Filters	Yes	Yes	No ³
AEN Enable	Yes	Yes	Yes
Get NC-SI Pass-Through Statistics	Yes, partially	Yes, partially	No ³
Select Package	Yes	Yes	Yes
Deselect Package	Yes	Yes	Yes
Enable Channel Network Tx ⁵	Yes	Yes	No
Disable Channel Network Tx	Yes	Yes	No
OEM Command ⁶	Yes	Yes	Yes

1. In cases that one of the LAN devices is assigned for the sole use of the manageability and its LAN PCIe function is disabled, using the NC-SI Set Link command while advertising multiple speeds and enabling auto-negotiation, results in the lowest possible speed chosen. To enable link of higher a speed, the BMC should not advertise speeds that are below the desired link speed. When doing it, changing the power state of the LAN device has not effect and the link speed is not re-negotiated.
2. The Intel® Xeon® Processor D-1500 Product Family LAN controller does not support filtering of user priority/DEI bits of VLAN.
3. In MCTP without pass through mode, only control commands are supported and not pass through traffic.As a result, many of the regular NC-SI commands are not supported or are supported in a limited manner, only to enable control and status reporting for the Intel® Xeon® Processor D-1500 Product Family LAN controller.
4. Set MAC Address command fails with a 0x002 (Parameter Is Invalid, Unsupported, or Out-of-Range) reason code if received on a Tx enabled port with a unicast MAC address already configured on another Tx enabled port.
5. Enable Channel Network Tx command fails with a 0x002 (Parameter Is Invalid, Unsupported, or Out-of-Range) reason code if received on a port configured with an unicast MAC address equal to the MAC address of another Tx enabled port
6. See [Section 10.6.3](#) for details.

Table 10.32. Optional NC-SI Features Support (Sheet 1 of 2)

Feature	Implement	Details
AENs	Yes	The software device driver state AEN might be emitted up to one minute after actual software device driver change if the software device driver was taken down unexpectedly. A re-configuration required AEN is sent before a firmware reset initiated due to a firmware code update
Get Controller Packet Statistics command	Yes, partially	Supports the following counters ¹ : 2-8,13-16 ² The statistics are cleared between reads.
Get NC-SI statistics	Yes	Supports all counters.
Get NC-SI Pass-Through Statistics	Yes, partially	Support the following counters: 1, 2, 6 ³ , 7.



Table 10.32. Optional NC-SI Features Support (Sheet 2 of 2)

Feature	Implement	Details
VLAN Modes	Yes, partially	Support only modes 1, 3.
Buffering Capabilities	Yes	8 Kb
MAC Address Filters	Yes	Supports two MAC addresses per port.
Channel Count	Yes	Supports four channels.
VLAN Filters	Yes	Supports eight VLAN filters per port. Filtering is ignoring the <i>DEI</i> bit and the 802.1P priority bits.
Broadcast Filters	Yes	Support the following filters: ARP. DHCP. Net BIOS.
Multicast Filters	Yes	Supports the following filters: IPv6 neighbor advertisement. IPv6 router advertisement. DHCPv6 relay and server multicast.
Hardware Arbitration	Yes	Supports NC-SI hardware arbitration.

1. *TCTL.EN* should be set to 1b to activate Tx related counters and *RCTL.RXEN*, *MANC.RCV_EN* or *GRC.APME* should be set to enable RX related counters.
2. As described in the get controller packet statistics counter numbers table in NC-SI specification.
3. The total pass-through Rx packets received on the LAN interface counter also includes OS2BMC traffic.

10.6.2.2 Allow Link Down (ALD) Support

NC-SI PHY power down conditions:

In NC-SI mode, the Intel® Xeon® Processor D-1500 Product Family LAN controller might dynamically change the PHY power mode according to the NC-SI channel state assuming no other functionality requires the PHY to be active (host, proxy or wake up).

The following algorithm is used to define if PHY activity is required:

- At initialization time, if the manageability mode is NC-SI, a PHY is required to be active only if the *Enable All PHYs in D3 N* bit in Common Firmware Parameters NVM word is set.
- Once a channel is enabled via Enable Channel NC-SI command, The PHY is powered up.
- If the channel is disabled via a Disable Channel command with ALD bit set and the PHY is disabled.
- If the channel is disabled via a Reset Channel command, the PHY power state is set back to the initialization value as defined by the *All PHYs in D3 N* bit.

Note: Before transitioning to D3 it is the responsibility of the software device driver to request the PHY to be active for walk up or proxy activities using WUC register (setting *WUC.PROXYE* and *MNG.MPROXYE* (*MANC* bit 30) forces the PHY to be active in D3).

10.6.2.3 AEN Handling

Asynchronous events might occur when the Intel® Xeon® Processor D-1500 Product Family LAN controller is not allowed to send them. The following rules defines the behavior of the Intel® Xeon® Processor D-1500 Product Family LAN controller in these cases:

1. While the Intel® Xeon® Processor D-1500 Product Family LAN controller is disabled, for each type of AEN only the last event is kept
2. Outstanding AENs that occurred while package was deselected is transmitted when package is selected.



- On a transition from channel disabled to channel enabled, all outstanding events are erased to prevent stale events notifications.

10.6.3 NC-SI Mode — Intel Specific Commands

In addition to regular NC-SI commands, the following Intel vendor specific commands are supported. The purpose of these commands is to provide a means for the BMC to access some of the Intel-specific features present in the Intel® Xeon® Processor D-1500 Product Family LAN controller.

10.6.3.1 Overview

The following features are available via the NC-SI OEM specific commands:

- Receive filters:
- Packet addition decision filters 0x0...0x4
- Packet reduction decision filters 0x5...0x7
- MNGONLY register (controls the forwarding of manageability packets to the host)
- Flex 128 filters
- Flex TCP/UDP port filters 0x0...0x2
- IPv4/IPv6 filters
- Get System MAC Address — This command enables the BMC to retrieve the system MAC address used by the NC. This MAC address can be used for a shared MAC address mode.
- Keep PHY Link Up (*Veto* bit) Enable/Disable — This feature enables the BMC to block PHY reset, which might cause session loss.
- TCO Reset — Enables the BMC to reset the Intel® Xeon® Processor D-1500 Product Family LAN controller.
- Checksum offloading — Offloads IP/UDP/TCP checksum checking from the BMC.
- OS2BMC Control commands.
- Firmware Version commands.
- MACsec logic programming

These commands are designed to be compliant with their corresponding SMBus commands (if existing). All of the commands are based on a single DMTF defined NC-SI command, known as OEM command. This command is as follows.

10.6.3.2 OEM Command (0x50)

The OEM command can be used by the BMC to request the sideband interface to provide vendor-specific information. The Vendor Enterprise Number (VEN) is the unique MIB/SNMP private enterprise number assigned by IANA per organization. Vendors are free to define their own internal data structures in the vendor data fields.

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...	Intel Command Number	Optional Data		



	Bits			
Bytes	31:24	23:16	15:08	07:00
...	...			
...	Optional Data		Padding to 32 bits (0x00)	
...	Checksum			

10.6.3.2.1 OEM Response (0xD0)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...27	Intel Command Number	Optional Return Data		
...	...			
...	Optional Return Data		Padding to 32 bits (0x00)	
...	Checksum			

Note: Responses have no command-specific reason code, unless otherwise specified within the command.

Note: The commands/responses described in the sections that follow includes only the part up to the data. The padding and checksum are implied.

10.6.3.3 OEM Commands Summary

Table 10.33. OEM Specific Command Response Reason Codes

Response Code		Reason Code	
Value	Description	Value	Description
0x1	Command Failed	0x5081	Invalid Intel Command Number.
		0x5082	Invalid Intel Command Parameter Number.
		0x5085	Internal Network Controller Error.
		0x5086	Invalid Vendor Enterprise Code.

Table 10.34. OEM Commands Summary (Sheet 1 of 3)

Intel Command	Parameter	Command Name	Supported in MCTP Without Pass Through	Description
0x00	0x00	Set IP Filters Control	No	10.6.3.5
0x01	0x00	Get IP Filters Control	No	10.6.3.6



Table 10.34. OEM Commands Summary (Sheet 2 of 3)

Intel Command	Parameter	Command Name	Supported in MCTP Without Pass Through	Description
0x02	0x0F	Set Manageability Only	No	10.6.3.7.2
	0x10	Set Flexible 128 Filter Mask and Length		10.6.3.7.3
	0x11	Set Flexible 128 Filter Data		10.6.3.7.4
	0x63	Set Flex TCP/UDP Port Filters		10.6.3.7.5
	0x64	Set Flex IPv4 Address Filters		10.6.3.7.6
	0x65	Set Flex IPv6 Address Filters		10.6.3.7.7
	0x67	Set EtherType Filter		10.6.3.7.8
	0x68	Set Packet Addition Extended Filter		10.6.3.7.9
0x03	0x0F	Get Manageability Only	No	10.6.3.8.2
	0x10	Get Flexible 128 Filter Mask and Length		10.6.3.8.3
	0x11	Get Flexible 128 Filter Data		10.6.3.8.4
	0x63	Get Flex TCP/UDP Port Filters		10.6.3.8.5
	0x64	Get Flex IPv4 Address Filters		10.6.3.8.6
	0x65	Get Flex IPv6 Address Filters		10.6.3.8.7
	0x67	Get EtherType Filter		10.6.3.8.8
	0x68	Get Packet Addition Extended Filter		10.6.3.8.9
0x04	0x10	Set Extended Unicast Packet Reduction	No	10.6.3.9.1
	0x11	Set Extended Multicast Packet Reduction		10.6.3.9.2
	0x12	Set Extended Broadcast Packet Reduction		10.6.3.9.4
0x05	0x10	Get Extended Unicast Packet Reduction	No	10.6.3.10.1
	0x11	Get Extended Multicast Packet Reduction		10.6.3.10.2
	0x12	Get Extended Broadcast Packet Reduction		10.6.3.10.3
0x06	N/A	Get System MAC Address	Yes	10.6.3.11
0x20	N/A	Set Intel Management Control	No	10.6.3.12
0x21	N/A	Get Intel Management Control	No	10.6.3.13
0x22	N/A	Perform TCO Reset	Yes	10.6.3.14
0x23	N/A	Enable IP/UDP/TCP Checksum Offloading	No	10.6.3.15.1
0x24	N/A	Disable IP/UDP/TCP Checksum Offloading	No	10.6.3.15.3
0x30	0x10	Transfer MACSec Ownership to BMC	No	10.6.3.16.1
	0x11	Transfer MACSec Ownership to Host		10.6.3.16.2
	0x12	Initialize MACsec RX		10.6.3.16.3
	0x13	Initialize MACsec TX		10.6.3.16.4
	0x14	Set MACsec RX Key		10.6.3.16.5
	0x15	Set MACsec TX Key		10.6.3.16.6
	0x16	Enable Network TX Encryption		10.6.3.16.7
	0x17	Disable Network TX Encryption		10.6.3.16.8
	0x18	Enable Network RX Decryption		10.6.3.16.9
	0x19	Disable Network RX Decryption		10.6.3.16.10



Table 10.34. OEM Commands Summary (Sheet 3 of 3)

Intel Command	Parameter	Command Name	Supported in MCTP Without Pass Through	Description
0x31	0x01	Get MACsec Rx Parameters	No	10.6.3.16.12
	0x02	Get MACsec Tx Parameters		10.6.3.16.13
0x40	0x01	Enable OS2BMC Flow	No	10.6.3.17.1
	0x02	Enable Network-to-BMC Flow		10.6.3.17.2
	0x03	Enable Both Network-to-BMC and Host-to-BMC Flow		10.6.3.17.3
	0x04	Set BMC IP Address	Yes	10.6.3.17.4
0x41	N/A	Get OS2BMC Parameters	No	10.6.3.17.5
0x48	0x1	Get Controller Information	Yes	10.6.3.18.1

Note: All the commands are supported both over RMII NC-SI and over MCTP.


10.6.3.4 Proprietary Commands Format

10.6.3.4.1 Set Intel Filters Control Command (Intel Command 0x00)

Bytes	Bits			
00...15	31...24	23...16	15...08	07...00
16...19	NC-SI Header			
20...23	Manufacturer ID (Intel 0x157)			
	0x00	Filter Control Index		

10.6.3.4.2 Set Intel Filters Control Response Format (Intel Command 0x00)

Bytes	Bits			
00...15	31...24	23...16	15...08	07...00
16...19	NC-SI Header			
20...23	Response Code		Reason Code	
24...27	Manufacturer ID (Intel 0x157)			
	0x00	Filter Control Index		



10.6.3.5 Set Intel Filters Control — IP Filters Control Command (Intel Command 0x00, Filter Control Index 0x00)

This command controls different aspects of the Intel filters.

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...23	0x00	0x00	IP Filters control (3-2)	
24...27	IP Filters Control (1-0)			

Where IP Filters Control has the following format:

Bit #	Name	Description	Default Value
0	IPv4/IPv6 Mode	IPv6 (0b) = There are zero IPv4 filters and four IPv6 filters. IPv4 (1b) = There are four IPv4 filters and four IPv6 filters.	1b
1...31	Reserved	Reserved.	

10.6.3.5.1 Set Intel Filters Control — IP Filters Control Response (Intel Command 0x00, Filter Control Index 0x00)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...27	0x00	0x00		

10.6.3.6 Get Intel Filters Control Commands (Intel Command 0x01)

10.6.3.6.1 Get Intel Filters Control — IP Filters Control Command (Intel Command 0x01, Filter Control Index 0x00)

This command controls different aspects of the Intel filters:

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...21	0x01	0x00		



10.6.3.6.1.1 Get Intel Filters Control — IP Filters Control Response (Intel Command 0x01, Filter Control Index 0x00)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...27	0x01	0x00	IP Filters Control (3-2)	
28...29	IP Filters Control (1-0)			

10.6.3.7 Set Intel Filters Formats

10.6.3.7.1 Set Intel Filters Command (Intel Command 0x02)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...21	0x02	Parameter Number	Filters Data (optional)	

10.6.3.7.1.1 Set Intel Filters Response (Intel Command 0x02)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...	0x02	Filter Control Index	Return Data (Optional)	

10.6.3.7.2 Set Intel Filters — Manageability Only Command (Intel Command 0x02, Filter Parameter 0x0F)

This command sets the MNGONLY register. The MNGONLY register controls whether pass-through packets destined to the BMC are not forwarded to the host operating system. The MNGONLY register is listed in [Table 10.4](#).



	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...23	0x02	0x0F	Manageability Only (3-2)	
24...25	Manageability Only (1-0)			

10.6.3.7.2.1 Set Intel Filters — Manageability Only Response (Intel Command 0x02, Filter Parameter 0x0F)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...25	0x02	0x0F		

10.6.3.7.3 Set Intel Filters — Flex Filter Enable Mask and Length Command (Intel Command 0x02, Filter Parameter 0x10)

The following command sets the Intel flex filters mask and length:

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...23	0x02	0x10	Mask Byte 1	Mask Byte 2
24...27
28...31
32...35
36...37	Mask Byte 15	Mask Byte 16	Reserved	Reserved
38	Length			



10.6.3.7.3.1 Set Intel Filters — Flex Filter Enable Mask and Length Response (Intel Command 0x02, Filter Parameter 0x10)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...25	0x02	0x10		

10.6.3.7.4 Set Intel Filters — Flex Filter Data Command (Intel Command 0x02, Filter Parameter 0x11)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...	0x02	0x11	Filter Data Group	Filter Data 1
	...	Filter Data N		

The Filter Data Group parameter defines which bytes of the flex filter are set by this command:

Table 10.35. Filter Data Group

Code	Bytes Programmed	Filter Data Length
0x0	bytes 0-29	1 - 30
0x1	bytes 30-59	1 - 30
0x2	bytes 60-89	1 - 30
0x3	bytes 90-119	1 - 30
0x4	bytes 120-127	1 - 8

Note: Using this command to configure the filters data must be done after the flex filter mask command is issued and the mask is set.



10.6.3.7.4.1 Set Intel Filters — Flex Filter Data Response (Intel Command 0x02, Filter Parameter 0x11)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...25	0x02	0x11		

10.6.3.7.5 Set Intel Filters — Flex TCP/UDP Port Filter Command (Intel Command 0x02, Filter Parameter 0x63)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...23	0x02	0x63	Port filter index	TCP/UDP Port MSB
24	TCP/UDP Port LSB			

Filter index range: 0x0...0xA.

If the filter index is bigger than 10, a command failed Response Code is returned with Invalid Intel Parameter Number reason (0x5082).

10.6.3.7.5.1 Set Intel Filters — Flex TCP/UDP Port Filter Response (Intel Command 0x02, Filter Parameter 0x63)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...25	0x02	0x63		



10.6.3.7.6 Set Intel Filters — IPv4 Filter Command (Intel Command 0x02, Filter Parameter 0x64)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...23	0x02	0x64	IP filter index	IPv4 Address (3)
24...26	IPv4 Address (2-0)			

Note: The filters index range can vary according to the IPv4/IPv6 mode setting in the Filters Control command.

IPv4 Mode: Filter index range: 0x0...0x3.

IPv6 Mode: This command should not be used in IPv6 mode.

10.6.3.7.6.1 Set Intel Filters — IPv4 Filter Response (Intel Command 0x02, Filter Parameter 0x64)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...25	0x02	0x64		

If the IPv4 address equals all zero, a command failed Response Code is returned, with Invalid Intel Parameter Number reason (0x5082).

10.6.3.7.7 Set Intel Filters — IPv6 Filter Command (Intel Command 0x02, Filter Parameter 0x65)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...23	0x02	0x65	IP filter index	...IPv6 Address (MSB, byte 15)
24...27
28...31
32...35
36...37	...		IPv6 Address (LSB, byte 0)	



Note: The filters index range can vary according to the IPv4/IPv6 mode setting in the Filters Control command.

IPv4 Mode: Filter index range: 0x1...0x3.

IPv6 Mode: Filter index range: 0x0...0x3.

10.6.3.7.7.1 Set Intel Filters — IPv6 Filter Response (Intel Command 0x02, Filter Parameter 0x65)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...25	0x02	0x65		

If the IP filter index does not match the previous ranges, a command failed Response Code is returned, with Invalid Intel Parameter Number reason (0x5082).

If the IPv6 address equals all zero, a command failed Response Code is returned, with Invalid Intel Parameter Number reason (0x5082).

10.6.3.7.8 Set Intel Filters - EtherType Filter Command (Intel Command 0x02, Filter parameter 0x67)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...23	0x02	0x67	EtherType Filter Index	EtherType Filter MSB
24...27	EtherType Filter LSB	

Where the EtherType filter has the format as described in [Section 8.2.2.15.3](#).



10.6.3.7.8.1 Set Intel Filters - EtherType Filter Response (Intel Command 0x02, Filter parameter 0x67)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...25	0x02	0x67		

If the Ethertype filter Index is larger than three, a command failed Response Code is returned with Invalid Intel Parameter Number reason (0x5082).

10.6.3.7.9 Set Intel Filters - Packet Addition Extended Decision Filter Command (Intel Command 0x02, Filter parameter 0x68)

See Figure 10.2 for description of the decision filters structure.

The command must overwrite any previously stored value. The value set is not checked.

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...23	0x02	0x68	Extended Decision filter Index	Extended Decision filter 1 MSB
24...27	Extended Decision filter 1 LSB	Extended Decision filter 0 MSB
28...30	Extended Decision filter 0 LSB	

Extended Decision filter Index Range: 0...4

Filter 0: See Table 10.36.

Filter 1: See Table 10.37.

Table 10.36. Filter Values (Sheet 1 of 2)

Bit #	Name	Description
1:0	Unicast (AND)	If set, packets must match unicast filter 0 to 1, respectively.
3:2	Reserved	Reserved.
4	Broadcast (AND)	If set, packets must match the broadcast filter.
12:5	VLAN (AND)	If set, packets must match VLAN filter 0 to 7, respectively.
16:13	IPv4 Address (AND)	If set, packets must match IPv4 filter 0 to 3, respectively
20:17	IPv6 Address (AND)	If set, packets must match IPv4 filter 0 to 3, respectively



Table 10.36. Filter Values (Sheet 2 of 2)

Bit #	Name	Description
22:21	Unicast (OR)	If set, packets can pass if match unicast filter 0 to 1, respectively or a different OR filter.
24:23	Reserved	Reserved.
25	Broadcast (OR)	If set, packets can pass if match the broadcast filter or a different OR filter.
26	Multicast (AND)	If set, packets must match the multicast filter.
27	ARP Request (OR)	If set, packets can pass if match the ARP request filter or a different OR filter.
28	ARP Response (OR)	If set, packets can pass if match the ARP response filter or a different OR filter.
29	Neighbor Discovery - 134 (OR)	If set, packets can pass if match the neighbor discovery filter (type134 - router advertisement) or a different OR filter.
30	Port 0x298 (OR)	If set, packets can pass if match a fixed TCP/UDP port 0x298 filter or a different OR filter.
31	Port 0x26F (OR)	If set, packets can pass if match a fixed TCP/UDP port 0x26F filter or a different OR filter.

Table 10.37. Extended Filter 1 Values

Bit #	Name	Description
3:0	Ethertype 0 -3 (AND)	If set, packets must match the Ethertype filter 0 to 3, respectively.
7:4	Ethertype 0 -3 (OR)	If set, packets must match the Ethertype filter 0 to 3, respectively or a different OR filter.
18:8	Flex port 10:0 (OR)	If set, packets can pass if match the TCP/UDP port filter 10:0.
19	DHCPv6 (OR)	If set, packets can pass if match the DHCPv6 port (0x0223).
20	DHCP Client (OR)	If set, packets can pass if match the DHCP server port (0x0043).
21	DHCP Server (OR)	If set, packets can pass if match the DHCP client port (0x0044).
22	NetBIOS Name Service (OR)	If set, packets can pass if match the NetBIOS name service port (0x0089).
23	NetBIOS Datagram Service (OR)	If set, packets can pass if match the NetBIOS datagram service port (0x008A).
24	Flex TCO (OR)	If set, packets can pass if match the flex 128 TCO filter.
25	Neighbor Discovery - 135 (OR)	If set, packets must also match the neighbor discovery filter (type135 - neighbor solicitation. or a different OR filter.
26	Neighbor Discovery - 136 (OR)	If set, packets must also match the neighbor discovery filter (type136 - neighbor advertisement) or a different OR filter.
27	Neighbor Discovery - 137 (OR)	If set, packets must also match the neighbor discovery filter (type137 - redirect) or a different OR filter.
28	Reserved	Reserved.
29	MLD (OR)	If set, packets must also match one of the MLD ICMPv6 types or a different OR filter.
31:30	Reserved	Reserved.



10.6.3.7.9.1 Set Intel Filters – Packet Addition Extended Decision Filter Response (Intel Command 0x02, Filter parameter 0x68)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...25	0x02	0x68		

If the Extended Decision filter index is bigger than 5, a command failed Response Code is returned with Invalid Intel Parameter Number reason (0x5082).

10.6.3.8 Get Intel Filters Formats

10.6.3.8.1 Get Intel Filters Command (Intel Command 0x03)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...21	0x03	Parameter Number		

10.6.3.8.1.1 Get Intel Filters Response (Intel Command 0x03)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...25	0x03	Parameter Number	Optional Return Data	



10.6.3.8.2 Get Intel Filters — Manageability Only Command (Intel Command 0x03, Filter Parameter 0x0F)

This command retrieves the MNGONLY register. The MNGONLY register controls whether pass-through packets destined to the BMC are also be forwarded to the host operating system.

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...21	0x03	0x0F		

10.6.3.8.2.1 Get Intel Filters — Manageability Only Response (Intel Command 0x03, Filter Parameter 0x0F)

The MNGONLY register structure is listed in [Table 10.4](#).

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...27	0x03	0x0F	Manageability to Host (3-2)	
28...29	Manageability to Host (1-0)			

10.6.3.8.3 Get Intel Filters — Flex Filter 0 Enable Mask and Length Command (Intel Command 0x03, Filter Parameter 0x10)

The following command retrieves the Intel flex filters mask and length. See [Section 10.3.3.5](#) for details of the values returned by this command.

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...21	0x03	0x10		



10.6.3.8.3.1 Get Intel Filters — Flex Filter 0 Enable Mask and Length Response (Intel Command 0x03, Filter Parameter 0x10)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...27	0x03	0x10	Mask Byte 1	Mask Byte 2
28...31
32...35
36...39
40...43	...	Mask Byte 16	Reserved	Reserved
44	Flexible Filter Length			

10.6.3.8.4 Get Intel Filters — Flex Filter 0 Data Command (Intel Command 0x03, Filter Parameter 0x11)

The following command retrieves the Intel flex filters data.

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...21	0x03	0x11	Filter Data Group 0...4	

The Filter Data Group parameter defines which bytes of the Flex filter are returned by this command:

Table 10.38. Filter Data Group

Code	Bytes Returned
0x0	bytes 0-29
0x1	bytes 30-59
0x2	bytes 60-89
0x3	bytes 90-119
0x4	bytes 120-127



10.6.3.8.4.1 Get Intel Filters — Flex Filter 0 Data Response (Intel Command 0x03, Filter Parameter 0x11)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...	0x03	0x11	Filter Group Number	Filter Data 1
	...	Filter Data N		

10.6.3.8.5 Get Intel Filters — Flex TCP/UDP Port Filter Command (Intel Command 0x03, Filter Parameter 0x63)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...22	0x03	0x63	TCP/UDP Filter Index	

Filter index range: 0x0...0xA.

10.6.3.8.5.1 Get Intel Filters — Flex TCP/UDP Port Filter Response (Intel Command 0x03, Filter Parameter 0x63)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...27	0x03	0x63	TCP/UDP Filter Index	TCP/UDP Port (1)
28	TCP/UDP Port (0)			

Filter index range: 0x0...0xA.

10.6.3.8.6 Get Intel Filters — IPv4 Filter Command (Intel Command 0x03, Filter Parameter 0x64)

	Bits			
Bytes	31...24	23...16	15...08	07...00



00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...22	0x03	0x64	IPv4 Filter Index	

Note: The filters index range can vary according to the IPv4/IPv6 mode setting in the Filters Control command.

IPv4 Mode: Filter index range: 0x0...0x3.

IPv6 Mode: This command should not be used in IPv6 mode.

10.6.3.8.6.1 Get Intel Filters — IPv4 Filter Response (Intel Command 0x03, Filter Parameter 0x64)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...27	0x03	0x64	IPv4 Filter Index	IPv4 Address (3)
28...29	IPv4 Address (2-0)			

10.6.3.8.7 Get Intel Filters — IPv6 Filter Command (Intel Command 0x03, Filter Parameter 0x65)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...22	0x03	0x65	IPv6 Filter Index	

Note: The filters index range can vary according to the IPv4/IPv6 mode setting in the Filters Control command

IPv4 Mode: Filter index range: 0x0...0x2.

IPv6 Mode: Filter index range: 0x0...0x3.



10.6.3.8.7.1 Get Intel Filters — IPv6 Filter Response Intel Command 0x03, Filter parameter 0x65)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...27	0x03	0x65	IPv6 Filter Index	IPv6 Address (MSB, Byte 16)
28...31
32...35
36...39
40...42	IPv6 Address (LSB, Byte 0)	

10.6.3.8.8 Get Intel Filters - EtherType Filter Command (Intel Command 0x03, Filter parameter 0x67)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...22	0x03	0x67	EtherType Filter Index	

Valid indices: 0...3

10.6.3.8.8.1 Get Intel Filters - EtherType Filter Response (Intel Command 0x03, Filter parameter 0x67)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...27	0x03	0x67	EtherType Filter Index	EtherType Filter MSB
28...30	EtherType Filter LSB	

If the Ethertype filter index is larger than three, a command failed Response Code is returned with Invalid Intel Parameter Number reason (0x5082).



10.6.3.8.9 Get Intel Filters – Packet Addition Extended Decision Filter Command (Intel Command 0x03, Filter parameter 0x68)

This command enables the BMC to retrieve the Extended Decision filter.

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...22	0x03	0x68	Extended Decision Filter Index	

10.6.3.8.9.1 Get Intel Filters – Packet Addition Extended Decision Filter Response (Intel Command 0x03, Filter parameter 0x68)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...27	0x03	0x68	Decision Filter Index	Decision Filter 1 MSB
28...31	Decision Filter 1 LSB	Decision Filter 0 MSB
32...34	Decision Filter 0 LSB	

Where Decision Filter 0 and Decision Filter 1 have the structure as detailed in the respective Set commands.

If the Extended Decision Filter Index is bigger than four, a command failed Response Code is returned with Invalid Intel Parameter Number reason (0x5082).

10.6.3.9 Set Intel Packet Reduction Filters Formats

The non-extended commands are obsolete. The extended commands ([Section 10.6.3.9.2](#) to [Section 10.6.3.9.4.1](#)) should be used instead.

10.6.3.9.1 Set Intel Packet Reduction Filters Command (Intel Command 0x04)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...23	0x04	Packet Reduction Index	Packet Reduction Data...	



Note: Intel recommends that the BMC only use the Extended Packet Reduction commands.

The *Packet Reduction Data* field has the following structure:

Table 10.39. Packet Reduction Field Description

Bit #	Name	Description
12:0	Reserved	Reserved.
16:13	IPv4 Address (AND)	If set, packets must match IPv4 filter 0 to 3, respectively.
20:17	IPv6 Address (AND)	If set, packets must match IPv4 filter 0 to 3, respectively.
27:21	Reserved	Reserved.
28	ARP Response (OR)	If set, packets can pass if match the ARP response filter or a different OR filter.
29	Reserved	Reserved.
30	Port 0x298	If set, packets can pass if match a fixed TCP/UDP port 0x298 filter.
31	Port 0x26F	If set, packets can pass if match a fixed TCP/UDP port 0x26F filter.

Table 10.40. Extended Packet Reduction Field Description

Bit #	Name	Description
3:0	Ethertype 0 -3 (AND)	If set, packets must match the Ethertype filter 0 to 3, respectively.
7:4	Ethertype 0-3 (OR)	If set, packets can pass if match the Ethertype filter 0 to 3, respectively.
15:12	Reserved	Reserved.
8:18	Flex port 10:0 (OR)	If set, packets can pass if match the TCP/UDP port filter 10:0.
23:19	Reserved	Reserved.
24	Flex TCO (OR)	If set, packets can pass if match the flex 128 TCO filter.
28:25	Reserved	Reserved.
29	MLD (OR)	If set, packets must also match one of the MLD ICMPv6 types or a different OR filter.
31:30	Reserved	Reserved.

The filtering is divided into two decisions:

- Bit 20:13 in [Table 10.39](#) and bits 3:2 in [Table 10.40](#) works in an AND manner; it must be true in order for a packet to pass (if was set).

Bits 28 in [Table 10.39](#) and bits 24:10 in [Table 10.40](#) work in an OR manner; at least one of them must be true for a packet to pass (if any were set).

10.6.3.9.1.1 Set Intel Packet Reduction Filters Response (Intel Command 0x04)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...	0x04	Packet Reduction Index		



10.6.3.9.2 Set Extended Unicast Packet Reduction Command (Intel Command 0x04, Reduction Filter Index 0x10)

The command must have the following format:

	Bits			
Bytes	31:24	23:16	15:08	07:00
00..15	NC-SI Header			
16..19	Manufacturer ID (Intel 0x157)			
20..23	0x04	0x10	Extended Unicast Reduction Filter MSB	..
24..27	..	Extended Unicast Reduction Filter LSB	Unicast Reduction Filter MSB	..
28..29	..	Unicast Reduction Filter LSB		

The command must overwrite any previously stored value.

Note: See [Table 10.39](#) and [Table 10.40](#) for description of the Unicast Extended Packet Reduction format.

10.6.3.9.2.1 Set Extended Unicast Packet Reduction Response (Intel Command 0x04, Reduction Filter Index 0x10)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..23	Manufacturer ID (Intel 0x157)			
24..25	0x04	0x10		

10.6.3.9.3 Set Extended Multicast Packet Reduction Command (Intel Command 0x04, Reduction Filter Index 0x11)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00..15	NC-SI Header			
16..19	Manufacturer ID (Intel 0x157)			
20..23	0x04	0x11	Extended Multicast Reduction Filter MSB	..
24..27	..	Extended Multicast Reduction Filter LSB	Multicast Reduction Filter MSB	..
28..29	..	Multicast Reduction Filter LSB		



Note: See Table 10.39 and Table 10.40 for description of the Multicast Extended Packet Reduction format.

The command must overwrite any previously stored value.

10.6.3.9.3.1 Set Extended Multicast Packet Reduction Response (Intel Command 0x04, Reduction Filter Index 0x11)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..23	Manufacturer ID (Intel 0x157)			
24..25	0x04	0x11		

10.6.3.9.4 Set Extended Broadcast Packet Reduction Command (Intel Command 0x04, Reduction Filter Index 0x12)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00..15	NC-SI Header			
16..19	Manufacturer ID (Intel 0x157)			
20..23	0x04	0x12	Extended Broadcast Reduction Filter MSB	..
24..27	..	Extended Broadcast Reduction Filter LSB	Broadcast Reduction Filter MSB	..
28..29	..	Broadcast Reduction Filter LSB		

Note: See Table 10.39 and Table 10.40 for description of the Broadcast Extended Packet Reduction format.

The command must overwrite any previously stored value.

10.6.3.9.4.1 Set Extended Broadcast Packet Reduction Response (Intel Command 0x04, Reduction Filter Index 0x12)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..23	Manufacturer ID (Intel 0x157)			
24..25	0x04	0x12		



10.6.3.10 Get Intel Packet Reduction Filters Formats

Note: The non-extended commands are not supported anymore. Use the extended commands (Section 10.6.3.10.1 to Section 10.6.3.10.3.1) instead.

10.6.3.10.1 Get Extended Unicast Packet Reduction Command (Intel Command 0x05, Reduction Filter Index 0x10)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...21	0x05	0x10		

10.6.3.10.1.1 Get Extended Unicast Packet Reduction Response (Intel Command 0x05, Reduction Filter Index 0x10)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...27	0x05	0x00	Extended Unicast Packet Reduction (3-2)	
28...29	Extended Unicast Packet Reduction (1-0)		Unicast Packet Reduction (3-2)	
30...31	Unicast Packet Reduction (1-0)			

10.6.3.10.2 Get Extended Multicast Packet Reduction Command (Intel Command 0x05, Reduction Filter Index 0x11)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...21	0x05	0x11		



10.6.3.10.2.1 Get Extended Multicast Packet Reduction Response (Intel Command 0x05, Reduction Filter Index 0x11)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...27	0x05	0x11	Extended Multicast Packet Reduction (3-2)	
28...29	Extended Multicast Packet Reduction (1-0)		Multicast Packet Reduction (3-2)	
30...31	Multicast Packet Reduction (1-0)			

10.6.3.10.3 Get Extended Broadcast Packet Reduction Command (Intel Command 0x05, Reduction Filter Index 0x12)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...21	0x05	0x12		

10.6.3.10.3.1 Get Extended Broadcast Packet Reduction Response (Intel Command 0x05, Reduction Filter Index 0x12)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...27	0x05	0x12	Extended Broadcast Packet Reduction (3-2)	
28...29	Extended Broadcast Packet Reduction (1-0)		Broadcast Packet Reduction (3-2)	
30...31	Broadcast Packet Reduction (1-0)			

10.6.3.11 System MAC Address

10.6.3.11.1 Get System MAC Address Command (Intel Command 0x06)



In order to support a system configuration that requires the Intel® Xeon® Processor D-1500 Product Family LAN controller to hold the MAC address for the BMC (such as shared MAC address mode), the following command is provided to enable the BMC to query the Intel® Xeon® Processor D-1500 Product Family LAN controller for a valid MAC address.

The NC must return the system MAC addresses. The BMC should use the returned MAC addressing as a shared MAC address by setting it using the Set MAC Address command as defined in NC-SI 1.0.

It is also recommended that the BMC use packet reduction and Manageability-to-host command to set the proper filtering method.

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20	0x06			


10.6.3.11.2 Get System MAC Address Response (Intel Command 0x06)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...27	0x06	MAC Address		
28...30	MAC Address			

10.6.3.12 Set Intel Management Control Formats

10.6.3.12.1 Set Intel Management Control Command (Intel Command 0x20)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...22	0x20	0x00	Intel Management Control 1	



Where Intel Management Control 1 is as follows:

Bit #	Default value	Description
0	0b	Enable Critical Session Mode (Keep PHY Link Up and Veto Bit) 0b = Disabled. 1b = Enabled. When critical session mode is enabled, the following behaviors are disabled: <ul style="list-style-type: none"> The PHY is not reset on PE_RST# and PCIe resets (in-band and link drop). Other reset events are not affected — Internal_Power_On_Reset, device disable, Force TCO, and PHY reset by software. The PHY does not change its power state. As a result, link speed does not change. The Intel® Xeon® Processor D-1500 Product Family LAN controller does not initiate configuration of the PHY to avoid losing link.
7:1	0x0	Reserved.

10.6.3.12.2 Set Intel Management Control Response (Intel Command 0x20)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...25	0x20	0x00		

10.6.3.13 Get Intel Management Control Formats

10.6.3.13.1 Get Intel Management Control Command (Intel Command 0x21)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...21	0x21	0x00		

Where Intel Management Control 1 is as described in [Section 10.6.3.12.2](#).



10.6.3.13.2 Get Intel Management Control Response (Intel Command 0x21)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...26	0x21	0x00	Intel Management Control 1	

10.6.3.14 TCO Reset

Depending on the bit set in the TCO mode field this command will cause the Intel® Xeon® Processor D-1500 Product Family LAN controller to perform either:

1. TCO Reset

- If Force TCO reset is enabled in the NVM. The Force TCO reset clears the data-path (Rx/Tx) of the Intel® Xeon® Processor D-1500 Product Family LAN controller to enable the BMC to transmit/receive packets through the Intel® Xeon® Processor D-1500 Product Family LAN controller.
- If the BMC has detected that the operating system is hung and has blocked the Rx/Tx path The Force TCO reset clears the data-path (Rx/Tx) of the Intel® Xeon® Processor D-1500 Product Family LAN controller to enable the BMC to transmit/receive packets through the Intel® Xeon® Processor D-1500 Product Family LAN controller.
- When this command is issued to a channel in a package, it applies only to the specific channel.
- After successfully performing the command the Intel® Xeon® Processor D-1500 Product Family LAN controller considers Force TCO command as an indication that the operating system is hung and clears the DRV_LOAD flag (disable the driver). If TCO reset is disabled in NVM, the Intel® Xeon® Processor D-1500 Product Family LAN controller clears the *CTRL_EXT.DRV_LOAD* bit but does not reset the data-path and notifies BMC on successful completion.

2. TCO Isolate.

- If TCO isolate is enabled in the NVM. The TCO Isolate command disables PCIe write operations to the LAN port.
- If TCO Isolate is disabled in NVM, the Intel® Xeon® Processor D-1500 Product Family LAN controller does not execute the command but sends a response to the BMC with successful completion.
- Following TCO Isolate, management isolates the function related to the port on which the command was received.

3. Firmware Reset.

- This command causes re-initialization of all the manageability functions and re-load of manageability related NVM words.
- When the BMC has loaded new management related NVM image the Firmware Reset command loads management related NVM information without need to power down the system.
- This command is issued to the package and affects all channels. After the Firmware Reset the Firmware Semaphore register (FWSM) is re-initialized.

Note: TCO Isolate affects only the channel (port) that the command was issued to. Force TCO resets the entire the Intel® Xeon® Processor D-1500 Product Family LAN controller (all channels in the package).



Following firmware reset, BMC needs to re-initialize all ports. Only one of the fields should be set in a given command. Setting more than one field might yield unexpected results.

10.6.3.14.1 Perform Intel TCO Reset Command (Intel Command 0x22)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20	0x22	TCO Mode		

Where TCO mode is:

Field	Bit(s)	Description
DO_TCO_RST	0	Perform TCO Reset. 0b= Do nothing. 1b= Perform TCO reset.
DO_TCO_ISOLATE ₁	1	Do TCO Isolate. 0b = Enable PCIe write access to LAN port. 1b = Isolate host PCIe write operation to the port. Note: Should be used for debug only. Note: The TCO isolate do not impact MCTP traffic. Note: When Isolate is set, the OS2BMC flow is disabled also.
RESET_MGMT	2	Reset manageability; re-load manageability NVM words. 0b = Do nothing. 1b = Issue firmware reset to manageability. Setting this bit generates a one-time firmware reset. Following the reset, management related data from NVM is loaded.
Reserved	7:3	Reserved (set to 0x00).

Note: For compatibility, the TCO Reset command without the TCO mode parameter is accepted (TCO reset is performed).

- TCO Isolate Host Write operation enabled in NVM.

10.6.3.14.2 Perform Intel TCO Reset Response (Intel Command 0x22)

Note: When a Firmware Reset is requested (TCO Mode = RESET_MGMT), there is no response, as the Intel® Xeon® Processor D-1500 Product Family LAN controller goes to Initial State as part of the command execution.

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...26	0x22			



10.6.3.15 Checksum Offloading

This command enables the checksum offloading filters in the Intel® Xeon® Processor D-1500 Product Family LAN controller.

When enabled, these filters block any packets that did not pass IP, UDP or TCP checksum from being forwarded to the BMC.

10.6.3.15.1 Enable Checksum Offloading Command (Intel Command 0x23)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20	0x23			

10.6.3.15.2 Enable Checksum Offloading Response (Intel Command 0x23)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...26	0x23			

10.6.3.15.3 Disable Checksum Offloading Command (Intel Command 0x24)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20	0x24			

10.6.3.15.4 Disable Checksum Offloading Response (Intel Command 0x24)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			



16...19	Response Code	Reason Code
20...23	Manufacturer ID (Intel 0x157)	
24...26	0x24	

10.6.3.16 MACsec Support Commands

The following commands may be used by the MC to control the different aspects of the MACsec engine.

10.6.3.16.1 Transfer MACsec Ownership to BMC Command (Intel Command 0x30, Parameter 0x10)

This command shall cause the Intel® Xeon® Processor D-1500 Product Family LAN controller to clear all MACsec parameters, forcefully release Host ownership and grant the ownership to the BMC. The BMC may allow the Host to use the BMC’s key for traffic by setting the “Host Control – Allow Host Traffic” bit. Activating this command will clear all the MACsec parameters.

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Manufacturer ID (Intel 0x157)			
20..22	0x30	0x10	Host Control	

Table 10.41. MACsec Host Control

Bit	Description
0	Reserved
1	Allow Host Traffic: 0b – Host traffic is blocked 1b – Host traffic is allowed
2..7	Reserved

10.6.3.16.1.1 Transfer MACsec Ownership to BMC Response (Intel Command 0x30, Parameter 0x10)

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..23	Manufacturer ID (Intel 0x157)			
24..25	0x30	0x10		



10.6.3.16.2 Transfer MACsec Ownership to Host Command (Intel Command 0x30, Parameter 0x11)

This command shall cause the Intel® Xeon® Processor D-1500 Product Family LAN controller to clear all MACsec parameters, release BMC ownership and grant ownership to the host.

In this scenario traffic from/to the MC shall be validated by the host’s programmed keys. It is recommended that the MC will try to establish network communication with a remote station to verify that the Host was successful in programming the keys.

Activating this command will clear all the MACsec parameters.

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Manufacturer ID (Intel 0x157)			
20..21	0x30	0x11		

10.6.3.16.2.1 Transfer MACsec Ownership to Host Response (Intel Command 0x30, Parameter 0x11)

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..23	Manufacturer ID (Intel 0x157)			
24..25	0x30	0x11		

10.6.3.16.3 Initialize MACsec RX Command (Intel Command 0x30, Parameter 0x12)

This command may be used by the MC to initialize the MACsec RX engine. This command should be followed by a “Set MACsec RX Key” command to establish a MACsec environment.

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Manufacturer ID (Intel 0x157)			
20..23	0x30	0x12	RX Port Identifier	
24..27	RX SCI [0..3]			
28..29	RX SCI [4..5]			

Where:



- **RX Port Identifier** – the port number by which the NC will identify RX packets. It is recommended that the MC uses 0x0 as the port identifier. Note: The MC should use the same port identifier when performing the key-exchange.
- **RX SCI** – A 6 bytes unique identifier for the MACsec TX CA. It is recommended that the MC uses its Ethernet MAC Address value for this field.

10.6.3.16.3.1 Initialize MACsec RX Response (Intel Command 0x30, Parameter 0x12)

		Bits			
Bytes		31..24	23..16	15..08	07..00
00..15	NC-SI Header				
16..19	Response Code			Reason Code	
20..23	Manufacturer ID (Intel 0x157)				
24..25	0x30	0x12			

10.6.3.16.4 Initialize MACsec TX Command (Intel Command 0x30, Parameter 0x13)

This command may be used by the MC to initialize the MACsec TX engine. This command should be followed by a "Set MACsec TX Key" command to establish a MACsec environment.

		Bits			
Bytes		31..24	23..16	15..08	07..00
00..15	NC-SI Header				
16..19	Manufacturer ID (Intel 0x157)				
20..23	0x30	0x13	TX Port Identifier		
24..27	TX SCI [0..3]				
28..31	TX SCI [4..5]			Reserved	
32..35	Packet Number Threshold				
36	TX Control				

- **TX Port Identifier** – For this implementation this field is a "don't care" and is automatically set to 0x0.
- **TX SCI** – A 6 bytes unique identifier for the MACsec TX CA. It is recommended that the MC uses its Ethernet MAC Address value for this field.
- **PN Threshold** – When a new key is programmed, the Packet Number is reset to 0x1. With each TX packet, The Packet Number increments by 1 and is inserted to the packet (to avoid replay attacks). The PN Threshold value is the 3 MSBytes of the TX Packet number after which a "Key Exchange Required" AEN will be sent to the MC. Example: a PN Threshold of 0x123456 means that when the PN reaches 0x123456FF a notification will be sent. The fourth byte of the PN Threshold can be seen as a reserved bit, because it will always be treated as 0xFF by the NC.
- **TX Control:**



Bit	Description
0..4	Reserved
5	Always Include SCI in TX: 0b – Do not include SCI in TX packets 1b – Include SCI in TX packets
6..7	Reserved

10.6.3.16.4.1 Initialize MACsec TX Response (Intel Command 0x30, Parameter 0x13)

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..23	Manufacturer ID (Intel 0x157)			
24..25	0x30	0x13		

10.6.3.16.5 Set MACsec RX Key Command (Intel Command 0x30, Parameter 0x14)

This command may be used by the MC to set a new MACsec RX key. Upon receiving this command the NC shall switch to the new RX key and send the response.

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Manufacturer ID (Intel 0x157)			
20..23	0x30	0x14	Reserved	RX SA AN
24..27	RX MACsec Key MSB
28..31
32..35
36..39	RX MACsec Key LSB

Where:

- **RX SA AN** – The Association Number to be used with this key.
- **RX MACsec Key** – the 128 bits (16 bytes) key to be used for RX

10.6.3.16.5.1 Set MACsec RX Key Response (Intel Command 0x30, Parameter 0x14)

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	NC-SI Header			



16..19	Response Code		Reason Code
20..23	Manufacturer ID (Intel 0x157)		
24..25	0x30	0x14	

10.6.3.16.6 Set MACsec TX Key Command (Intel Command 0x30, Parameter 0x15)

This command may be used by the MC to set a new MACsec TX key. Upon receiving this command the NC shall switch to the new TX key and send the response.

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Manufacturer ID (Intel 0x157)			
20..23	0x30	0x15	Reserved	TX SA AN
24..27	TX MACsec Key MSB
28..31
32..35
36..39	TX MACsec Key LSB

Where:

- **TX SA AN** – The Association Number to be used with this key.
- **TX MACsec Key** – the 128 bits (16 bytes) key to be used for TX

10.6.3.16.6.1 Set MACsec TX Key Response (Intel Command 0x30, Parameter 0x15)

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..23	Manufacturer ID (Intel 0x157)			
24..25	0x30	0x15		

10.6.3.16.7 Enable Network TX Encryption Command (Intel Command 0x30, Parameter 0x16)

This command may be used by the MC to (re)enable Encryption of outgoing Pass-Through packets.

After this command is issued and until a response is received, the state of any outgoing packets is undetermined.

By default Network TX Encryption is enabled.



	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Manufacturer ID (Intel 0x157)			
20..21	0x30	0x16		

10.6.3.16.7.1 Enable Network TX Encryption Response (Intel Command 0x30, Parameter 0x16)

Following sending this response the NC shall stop encrypting outgoing Pass-Through packets.

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..23	Manufacturer ID (Intel 0x157)			
24..25	0x30	0x16		

10.6.3.16.8 Disable Network TX Encryption Command (Intel Command 0x30, Parameter 0x17)

This command may be used by the MC to disable Encryption of outgoing Pass-Through packets.

After this command is issued and until a response is received, the state of any outgoing packets is undetermined.

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Manufacturer ID (Intel 0x157)			
20..21	0x30	0x17		

10.6.3.16.8.1 Disable Network TX Encryption Response (Intel Command 0x30, Parameter 0x17)

Following sending this response the NC shall start encrypting outgoing Pass-Through packets.

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			



16..19	Response Code		Reason Code
20..23	Manufacturer ID (Intel 0x157)		
24..25	0x30	0x17	

10.6.3.16.9 Enable Network RX Decryption Command (Intel Command 0x30, Parameter 0x18)

This command may be used by the MC to (re)enable decryption of incoming Pass-Through packets. This will cause the NC to execute MACsec offload and to post the frames to the MC (or host) only if the MACsec operation succeeds.

After this command is issued and until a response is received, the state of any incoming packets is undetermined.

By default Network RX Decryption is disabled.

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Manufacturer ID (Intel 0x157)			
20..21	0x30	0x18		

10.6.3.16.9.1 Enable Network RX Decryption Response (Intel Command 0x30, Parameter 0x18)

Following sending this response the NC shall begin decrypting incoming Pass-Through packets.

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..23	Manufacturer ID (Intel 0x157)			
24..25	0x30	0x18		

10.6.3.16.10 Disable Network RX Decryption Command (Intel Command 0x30, Parameter 0x19)

This command may be used by the MC to disable decryption of incoming Pass-Through packets.

After this command is issued and until a response is received, the state of any incoming packets is undetermined.

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Manufacturer ID (Intel 0x157)			
20..21	0x30	0x19		



10.6.3.16.10.1 Disable Network RX Decryption Response (Intel Command 0x30, Parameter 0x19)

Following sending this response the NC shall stop decrypting incoming Pass-Through packets.

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..23	Manufacturer ID (Intel 0x157)			
24..25	0x30	0x19		

10.6.3.16.11 Get MACsec Parameters format (Intel Command 0x31)

The following commands may be used by the MC to retrieve the different MACsec parameters.

These commands responses are valid only if the BMC owns the MACsec.

10.6.3.16.12 Get MACsec RX Parameters Command (Intel Command 0x31, Parameter 0x01)

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Manufacturer ID (Intel 0x157)			
20..22	0x31	0x01		

10.6.3.16.12.1 Get MACsec RX Parameters Response (Intel Command 0x31, Parameter 0x01)

This command allows the MC to retrieve the currently configured set of RX MACsec parameter.

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..23	Manufacturer ID (Intel 0x157)			
24..27	0x31	0x01	Reserved	
28..31	MACsec Owner Status	MACsec Host Control Status	RX Port Identifier	
32..35	SCI [0..3]			
36..39	SCI [4..5]		Reserved	RX SA AN
40..43	RX SA Packet Number			

Where:




Table 10.42. MACsec Owner Status

Value	Description
0x0	Host is MACsec owner
0x1	BMC is MACsec owner

Table 10.43. MACsec Host Control Status

Bit	Description
0	Reserved
1	Allow Host Traffic: 0b- Host traffic is blocked 1b - Host traffic is allowed
2..7	Reserved

- RX Port Identifier – The RX Port identifier
- RX SCI – The RX SCI identifier.
- RX SA AN – The association number associated with the active SA (for which the last valid RX MACsec packet was received).
- RX SA Packet Number – Is the last packet number, as read from the last valid RX MACsec packet.

10.6.3.16.13 Get MACsec TX Parameters Command (Intel Command 0x31, Parameter 0x02)

This command allows the MC to retrieve the currently configured set of TX MACsec parameter.

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Manufacturer ID (Intel 0x157)			
20..22	0x31	0x02		

10.6.3.16.13.1 Get MACsec TX Parameters Response (Intel Command 0x31, Parameter 0x02)

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..23	Manufacturer ID (Intel 0x157)			
24..27	0x31	0x2	Reserved	
28..31	MACsec Owner Status	MACsec Host Control Status	TX Port Identifier	
32..35	SCI [0..3]			
36..39	SCI [4..5]		Reserved	TX SA AN



	Bits			
Bytes	31..24	23..16	15..08	07..00
40..43	TX SA Packet Number			
44..47	Packet Number Threshold			
48	TX Control Status			

Where:

Table 10.44. MACsec Owner Status:

Value	Description
0x0	Host is MACsec owner
0x1	BMC is MACsec owner

Table 10.45. MACsec Host Control Status

Bit	Description
0	Reserved
1	Allow Host Traffic: 0b- Host traffic is blocked 1b - Host traffic is allowed
2..7	Reserved

- TX Port Identifier – Reserved to 0x0 for this implementation.
- TX SCI – The RX SCI identifier.
- TX SA AN – The association number currently used for the active SA.
- TX SA Packet Number – Is the last packet number, as read from the last valid RX MACsec packet.
- Packet Number Threshold.

Table 10.46. TX Control Status:

Bit	Description
0..4	Reserved
5	Include SCI: 0b – Do not include SCI in TX packets 1b – Include SCI in TX packets
6..7	Reserved

10.6.3.17 OS 2 BMC configuration

These commands control enabling of the OS-to-BMC flow.



10.6.3.17.1 EnableOS2BMC Flow Command (Intel Command 0x40, Index 0x1)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...21	0x40	0x01		

10.6.3.17.1.1 EnableOS2BMC Flow Response (Intel Command 0x40, Index 0x1)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...25	0x40	0x01		

10.6.3.17.2 Enable Network to BMC Flow Command (Intel Command 0x40, Index 0x2)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...21	0x40	0x02		

10.6.3.17.2.1 Enable Network to BMC Flow Response (Intel Command 0x40, Index 0x2)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...25	0x40	0x02		



10.6.3.17.3 Enable both Host and Network to BMC flows Command (Intel Command 0x40, Index 0x3)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...21	0x40	0x03		

10.6.3.17.3.1 Enable both Host and Network to BMC flows Response (Intel Command 0x40, Index 0x3)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...25	0x40	0x03		


10.6.3.17.4 Set BMC IP Address Command (Intel Command 0x40, Index 0x4)

This command is used to expose the BMC IP address to the host.

The IP type entry indicate whether the IP address is an IPv4 or an IPv6 address:

- 0 = Ipv4
- 1 = IPv6
- 2 = No IP address, then the command should not include an IP address.

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Manufacturer ID (Intel 0x157)			
20..23	0x40	0x04	IP type	IPv6 Address (MSB, byte 15)/IPv4 Address (MSB, byte 3)
24..27	IPv6 Address (byte 14)/IPv4 Address (byte 2)	IPv6 Address (byte 13)/IPv4 Address (byte 1)	IPv6 Address (byte 12)/IPv4 Address (LSB, byte 0)	IPv6 Address (byte 11)/Reserved
28..31
32..35
36..38	IPv6 Address (LSB, byte 0)/Reserved	



10.6.3.17.4.1 Set BMC IP Address Response (Intel Command 0x40, Index 0x4)

	Bits			
Bytes	31...24	23...16	15...08	07...00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...25	0x40	0x04		

10.6.3.17.5 Get OS2BMC Parameters Command (Intel Command 0x41)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20	0x41			

10.6.3.17.5.1 Get OS2BMC Parameters Response (Intel Command 0x41)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...27	0x41	Status	IPv6 Address (MSB, byte 15)/IPv4 Address (MSB, byte 3)	IPv6 Address (byte 14)/IPv4 Address (byte 2)
28..31	IPv6 Address (byte 13)/IPv4 Address (byte 1)	IPv6 Address (byte 12)/IPv4 Address (LSB, byte 0)	IPv6 Address (byte 11)/Reserved	..
32..35
36..39
40..41	..	IPv6 Address (LSB, byte 0)/Reserved		



Where the status byte partition is as follow:

Table 10.47. Status byte description

Bits	Content
0	0b = IPv4. 1b = IPv6. Relevant only if the IP address valid bit is set.
1	IP address valid.
2	Network to BMC status. 0b = network 2 BMC flow is disabled. 1b = network 2 BMC flow is enabled.
3	OS2BMC status. 0b = OS 2 BMC flow is disabled. 1b = OS 2 BMC flow is enabled.
7:4	Reserved.


10.6.3.18 Get Controller information Command (Intel Command 0x48, Index 0x1)

This command gather the controller identification information and return it back to the BMC.

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Manufacturer ID (Intel 0x157)			
20...23	0x48	0x1		

10.6.3.18.1 Get Controller Information response (Intel Command 0x48, Index 0x1)

	Bits			
Bytes	31:24	23:16	15:08	07:00
00...15	NC-SI Header			
16...19	Response Code		Reason Code	
20...23	Manufacturer ID (Intel 0x157)			
24...27	0x48	0x01	Reserved	Number of Inventory entries
28...31	Controller Info Item 1 ID	Controller Info Item 1 length	Controller Info Item 1 Data	
...			
...	Controller Info Item 2 ID	Controller Info Item 2 length	Controller Info Item 2 Data	
...			
...	Controller Info Item n ID	Controller Info Item n length	Controller Info Item n Data	
...			



Where the possible inventory items are described as follows. Note that not all the inventory items would be present in all the implementations of this command.

Table 10.48. Controller information Items

ID	Length (in bytes)	Data	Notes
0x00	3	Device ID (2 bytes) + RevID	This is the hardware default value, not any value programmed via NVM.
0x0B	2	NVM Image Version	
0x0C	4	EMP ROM Internal Version	
0x0D	4	Flash Firmware Image Internal version	
0x0E	2	PXE Firmware version	MajorVersion.MinorVersion.Build.
0x0F	2	iSCSI Firmware version	
0x10	2	uEFI Firmware version	
0x16	2	Reserved	
0x17	4	Firmware Mini Loader Internal version	

10.6.4 Asynchronous Event Notifications

The asynchronous event notifications are unsolicited messages sent from the NC to the BMC to report status changes (such as link change, operating system state change, etc.).

Recommendations:

- The BMC firmware designer should use AENs. To do so, the designer must take into account the possibility that a NC-SI response frame (such as a frame with the NC-SI EtherType), arrives out-of-context (not immediately after a command, but rather after an out-of-context AEN).
- To enable AENs, the BMC should first query which AENs are supported, using the Get Capabilities command, then enable desired AEN(s) using the Enable AEN command, and only then enable the channel using the Enable Channel command.

10.6.5 Querying Active Parameters

The BMC can use the Get Parameters command to query the current status of the operational parameters.

10.6.6 Resets

In NC-SI there are two types of resets defined:

1. Synchronous entry into the initial state.
2. Asynchronous entry into the initial state.

Recommendations:

- It is very important that the BMC firmware designer keep in mind that following any type of reset, all configurations are considered as lost and thus the BMC must re-configure everything.
- As an asynchronous entry into the initial state might not be reported and/or explicitly noticed, the BMC should periodically poll the NC with NC-SI commands (such as Get Version ID, Get Parameters, etc.) to verify that the channel is not in the initial state. Should the NC channel respond to the command with a Clear Initial State Command Expected reason code, the BMC should consider the channel (and most probably the entire NC package) as if it underwent a (possibly unexpected) reset event. Thus, the BMC should re-configure the NC. See the NC-SI specification section on Detecting Pass-through Traffic Interruption.



- The Intel recommended polling interval is 2-3 seconds.

For exact details on the resets, refer to NC-SI specification.

10.6.7 Advanced Workflows

10.6.7.1 Multi-NC Arbitration

As described in [Section 10.6.1.2](#), in a multi-NC environment, there is a need to arbitrate the NC-SI lines.

[Figure 10.9](#) shows the system topology of such an environment.

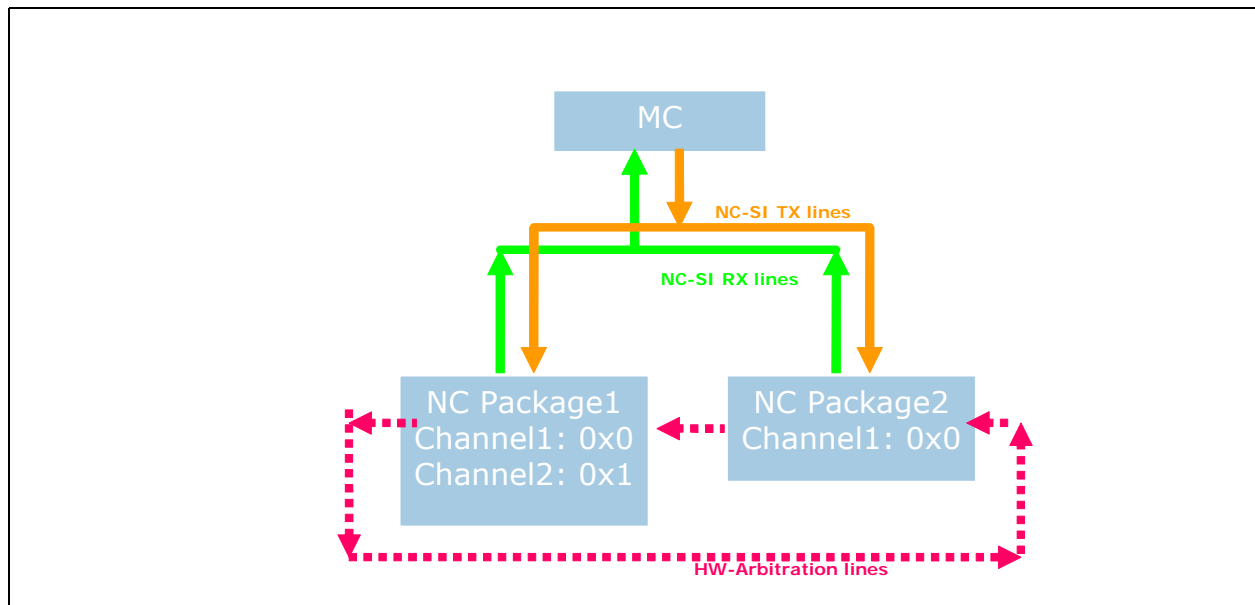


Figure 10.9. Multi-NC Environment

See [Figure 10.9](#). The NC-SI Rx lines are shared between the NCs. To enable sharing of the NC-SI Rx lines, NC-SI has defined an arbitration scheme.

The arbitration scheme mandates that only one NC package can use the NC-SI Rx lines at any given time. The NC package that is allowed to use these lines is defined as selected. All the other NC packages are de-selected.

NC-SI has defined two mechanisms for the arbitration scheme:

1. Package selection by the BMC. In this mechanism, the BMC is responsible for arbitrating between the packages by issuing NC-SI commands (Select/De-Select Package). The BMC is responsible for having only one package selected at any given time.
2. Hardware arbitration. In this mechanism, two additional pins on each NC package are used to synchronize the NC package. Each NC package has an ARB_IN and ARB_OUT line and these lines are used to transfer Tokens. A NC package that has a token is considered selected.

Note: Hardware arbitration is enabled by the NC-SI HW Arbitration Enable configuration bit in the NC-SI Configuration 1 NVM word.

For details, refer to the NC-SI specification.



10.6.7.2 Package Selection Sequence Example

Following is an example work flow for a BMC and occurs after the discovery, initialization, and configuration.

Assuming the BMC needs to share the NC-SI bus between packages, the BMC should:

1. Define a time-slot for each device.
2. Discover, initialize, and configure all the NC packages and channels.
3. Issue a De-Select Package command to all the channels.
4. Set active_package to 0x0 (or the lowest existing package ID).
5. At the beginning of each time slot the BMC should:
 - a. Issue a De-Select Package to the active_package. The BMC must then wait for a response and then an additional timeout for the package to become de-selected (200 μ s). See the NC-SI specification table 10 — parameter NC Deselect to Hi-Z Interval.
 - b. Find the next available package (typically active_package = active_package + 1).
 - c. Issue a Select Package command to active_package.

10.6.7.3 Multiple Channels (Fail-Over)

In order to support a fail-over scenario, it is required from the BMC to operate two or more channels. These channels might or might not be in the same package.

The key element of a fault-tolerance fail-over scenario is having two (or more) channels identifying to the switch with the same MAC address, but only one of them being active at any given time (such as switching the MAC address between channels). To accomplish this, NC-SI provides the following commands:

1. Enable Network Tx command — This command enables shutting off the network transmit path of a specific channel. This enables the BMC to configure all the participating channels with the same MAC address but only enable one of them.
2. Link Status Change AEN or Get Link Status command.

10.6.7.3.1 Fail-Over Algorithm Example

The following is a sample workflow for a fail-over scenario for the Intel® Xeon® Processor D-1500 Product Family LAN controller quad-port 10 GbE controller (one package and four channels):

1. BMC initializes and configures all channels after power-up. However, the BMC uses the same MAC address for all of the channels.
2. The BMC queries the link status of all the participating channels. The BMC should continuously monitor the link status of these channels. This can be accomplished by listening to AENs (if used) and/or periodically polling using the Get Link Status command.
3. The BMC then only enables channel 0 for network transmission.
4. The BMC then issues a gratuitous ARP (or any other packet with its source MAC address) to the network. This packet informs the switch that this specific MAC address is registered to channel 0's specific LAN port.
5. The BMC begins normal workflow.
6. Should the BMC receive an indication (AEN or polling) that the link status for the active channel (channel 0) has changed, the BMC should:
 - a. Disable channel 0 for network transmission.
 - b. Check if a different channel is available (link is up).
 - c. If found:



- Enable network Tx for that specific channel.
- Issue a gratuitous ARP (or any other packet with its source MAC address) to the network. This packet informs the switch that this specific MAC address is registered to channel 0's specific LAN port.
- Resume normal workflow.
- If not found, report the error and continue polling until a valid channel is found.

The previous algorithm can be generalized such that the start-up and normal workflow are the same. In addition, the BMC might need to use a specific channel (such as channel 0). In this case, the BMC should switch the network transmit to that specific channel as soon as that channel becomes valid (link is up).

Recommendations:

- Wait for a link-down-tolerance timeout before a channel is considered invalid. For example, a link re-negotiation might take a few seconds (normally 2 to 3 or might be up to 9). Thus, the link must be re-established after a short time.
- Typically, this timeout is recommended to be three seconds.
- Even when enabling and using AENs, periodically poll the link status, as dropped AENs might not be detected.

10.6.7.4 Statistics

The BMC might use the statistics commands as defined in NC-SI. These counters are meant mostly for debug purposes and are not all supported.

The statistics are divided into three commands:

1. Controller statistics — These are statistics on the network interface (to the host operating system and the pass through traffic). See the NC-SI specification for details.
2. NC-SI statistics — These are statistics on the NC-SI control frames (such as commands, responses, AENs, etc.). See the NC-SI specification for details.

NC-SI pass-through statistics — These are statistics on the NC-SI pass-through frames. See the NC-SI specification for details.

10.6.8 External Link Control via NC-SI

The BMC can use the NC-SI Set Link command to control the external interface link settings. This command enables the BMC to set the auto-negotiation, link speed, duplex, and other parameters.

This command is only available when the host operating system is not present. Indicating the host operating system status can be obtained via the Get Link Status command and/or Host OS Status Change AEN command.

Recommendation:

- Unless explicitly needed, it is not recommended to use this feature. The NC-SI Set Link command does not expose all the possible link settings and/or features. This might cause issues under different scenarios. Even if you decided to use this feature, use it only if the link is down (trust the Intel® Xeon® Processor D-1500 Product Family LAN controller until proven otherwise).
- It is recommended that the BMC first query the link status using the Get Link Status command. The BMC should then use this data as a basis and change only the needed parameters when issuing the Set Link command.

For details, refer to the NC-SI specification.



10.6.8.1 Set Link While LAN PCIe Functionality is Disabled

In cases where the Intel® Xeon® Processor D-1500 Product Family LAN controller is used solely for manageability and its LAN PCIe function is disabled, using the NC-SI Set Link command while advertising multiple speeds and enabling auto-negotiation results in the lowest possible speed chosen.

To enable link of higher a speed, the BMC should not advertise speeds that are below the desired link speed, as the lowest advertised link speed is chosen.

When the Intel® Xeon® Processor D-1500 Product Family LAN controller is only used for manageability and the link speed advertisement is configured by the BMC, changes in the power state of the LAN device is not effected and the link speed is not re-negotiated by the LAN device.

10.6.8.2 Set Link Error codes


The following rules are used to define the error code returned for Set Link command in case an invalid configuration is requested:

1. Host Driver Check: If host device driver is present, return a Command Specific Response (0x9) with a Set Link Host OS/Driver Conflict Reason (0x1).
2. Speed Present Check: If no speed is selected, return a General Reason Code for a failed command (0x1) with Parameter Is Invalid, Unsupported, or Out-of-Range Reason (0x2).
3. Parameter Validity:
 - a. Auto-Negotiation Parameter Validation: If auto-negotiation is requested and none of the selected parameters are valid for the device, return a General Reason Code for a failed command (0x1) with a Parameter Is Invalid, Unsupported, or Out-of-Range Reason (0x2).

Note: This means that, for example, a command requesting 10 GbE on a 1 GbE device succeeds provided that the command requests at least one other supported speed.

The same goes for an unsupported duplex setting (a device with no HD support accepts a command with both FD and HD set), and also for HD being requested with speeds of 1 GbE and higher as long as a speed below 1 GbE is also requested (and is supported in HD). The Intel® Xeon® Processor D-1500 Product Family LAN controller ignore the unsupported parameters.

- b. Force Mode Parameter Validation:
 - If more than one link speed is being forced, then return a General Reason Code for a failed command (0x1) and a Command Specific Reason with a Set Link Speed Conflict Error (0x0905).
 - If more than one duplex setting is being forced, then return a General Reason Code for a failed command (0x1) with Parameter Is Invalid, Unsupported, or Out-of-Range Reason (0x2).
 - If 1 GbE and above is requested with HD, then return a General Reason Code for a failed command (0x1) and a Command Specific Reason with Set Link Parameter Conflict Reason (0x0903).
4. Media Type Compatibility Check: If current media type is not compatible for the requested link parameters, return a General Reason Code for a failed command (0x1) and a Command Specific Reason with Set Link Media Conflict Error (0x0902).
5. Power State Compatibility Check: If current power state does not allow for the requested link parameters, return a General Reason Code for a failed command (0x1) and a Command Specific Reason with Set Link Power Mode Conflict Reason (0x0904).
6. If for some reason the hardware cannot perform the flow required for the command, return a General Reason Code for a failed command (0x1) and a Command Specific Response (0x9) with Link Command Failed-Hardware Access Error (0x6).



10.7 Management Component Transport Protocol (MCTP)

10.7.1 MCTP Overview

MCTP defines a communication model intended to facilitate communication between:

- Management controllers and other management controllers
- Management controllers and management devices

The communication model includes a message format, transport description, message exchange patterns, and configuration and initialization messages.

The basic MCTP specification is described in DMTF's DSP0236 document.

MCTP is designed so that it can potentially be used on many bus types. The protocol is intended to be used for intercommunication between elements of platform management subsystems used in computer systems, and is suitable for use in mobile, desktop, workstation, and server platforms.

Currently, a specification exists for MCTP over SMBus (DMTF's DSP0237). A specification for MCTP over USB is also planned.

Management controllers such as a baseboard management controller (BMC) can use this protocol for communication between one another, as well as for accessing management devices within the platform.

10.7.1.1 NC-SI Over MCTP

MCTP is a transport layer protocol that do not include the functionality required to control the pass through traffic required for BMC connection to the network. This functionality is provided by encapsulating NC-SI traffic as defined in DMTF's DSP0222 document.

The details of NC_SI over MCTP protocol are defined in the NC-SI Over MCTP Specification.

The NC-SI over MCTP specification defines two types of MCTP message types: NC-SI (0x2) and Ethernet (0x3). The Intel® Xeon® Processor D-1500 Product Family LAN controller supports both messages. When used only for control, then only the NC-SI (0x2) message type is supported.

In addition to the above message types supported by the Intel® Xeon® Processor D-1500 Product Family LAN controller, the PCIe based VDM message type is also supported over PCIe to support ACL commands.

Details of the NC-SI over MCTP can be found in [Section 10.7.4](#).

10.7.1.2 MCTP Usage Model

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports NC-SI over MCTP protocol over the SMBus bus. The Intel® Xeon® Processor D-1500 Product Family LAN controller can connect through MCTP to a BMC or the ME engine in the chipset as shown in [Figure 10.10](#).

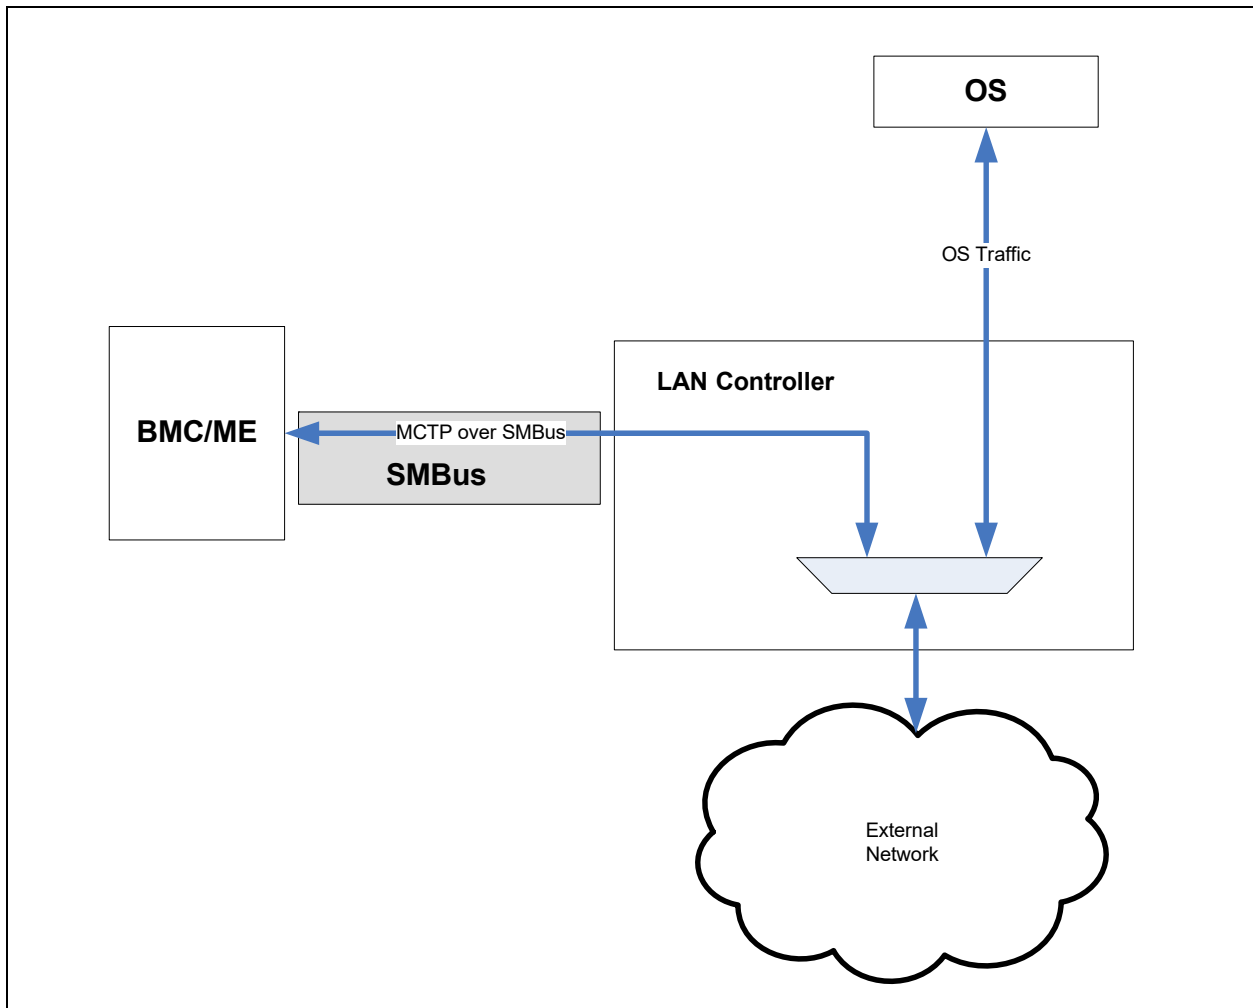


Figure 10.10. MCTP Connections of the Intel® Xeon® Processor D-1500 Product Family LAN Controller

10.7.2 NC-SI to MCTP Mapping

The two network ports of the Intel® Xeon® Processor D-1500 Product Family LAN controller (mapped to two NC-SI channels) are mapped to a single MCTP endpoint.

Section 10.7.2.1 describes the transition between the bus.

For SMBus, the Intel® Xeon® Processor D-1500 Product Family LAN controller should expect MCTP commands from two sources: the bus owner and the BMC. In addition, it should expect pass-through traffic. Thus, it should be able to process up to three interleaved commands/data:



- An MCTP control/OEM command from the SMBus bus owner (single packet message).
- An MCTP control/OEM command from the BMC over SMBus (single packet message).
- An NC-SI command or Ethernet packet from the BMC over the active channel.

A single source should not interleave packets it sends.

The topology used for MCTP connection is shown in [Figure 10.11](#).

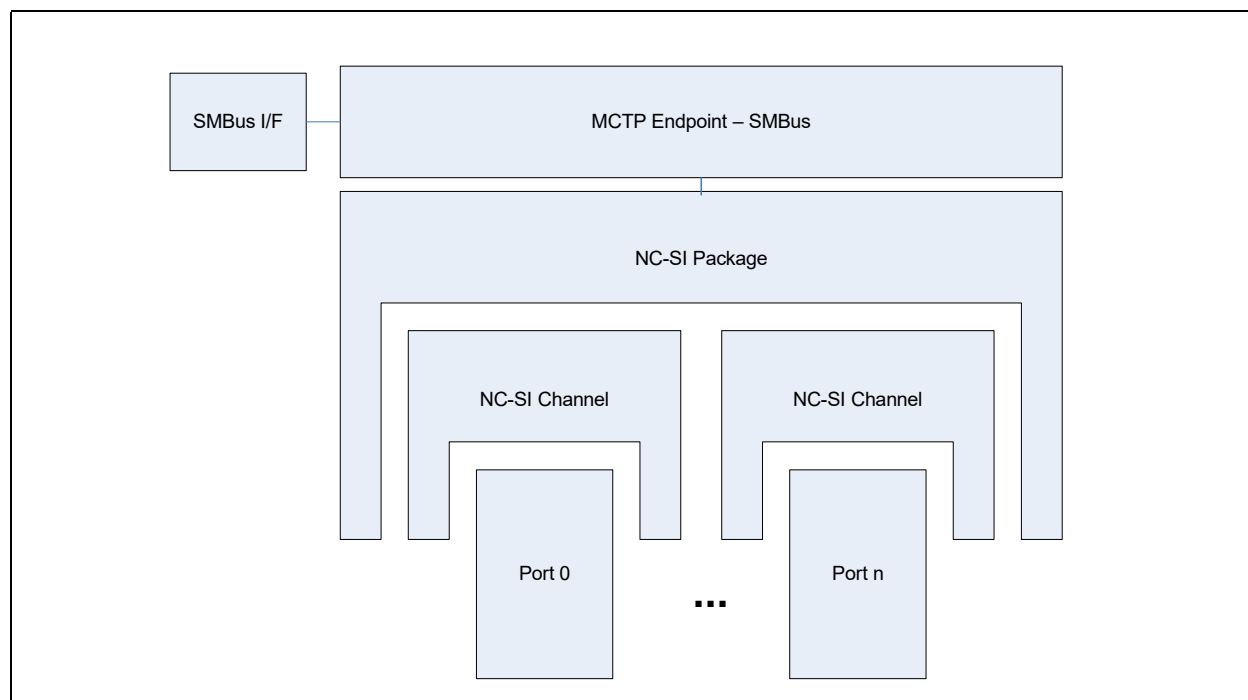


Figure 10.11. MCTP Endpoint Topology

10.7.2.1 Detection of BMC EID and Physical Address

In order to enable transactions between the BMC and the Intel® Xeon® Processor D-1500 Product Family LAN controller, the bus physical address (SMBus) and the EID of the partner needs to be discovered. The Intel® Xeon® Processor D-1500 Product Family LAN controller does not try to discover the BMC and assume the BMC initiates the connection. If the Intel® Xeon® Processor D-1500 Product Family LAN controller is in NC-SI initial state, then the EID and the physical address of the BMC are extracted from the *Clear Initial State* command parameters or any other NC-SI command received later with a channel ID of the Intel® Xeon® Processor D-1500 Product Family LAN controller. Subsequent pass through traffic is received from or sent to this address only.

If the EID or the physical address of the Intel® Xeon® Processor D-1500 Product Family LAN controller changes, it indicates the changes to bus owner so that the routing tables can be updated. There is no attempt to directly send an indication to the BMC about the change.

10.7.2.2 Bus Transition

The following section defines the transition flow between the SMBus bus on which MCTP flows. Figure 10.12 describes the flow to transition between the SMBus. The following parameters are used to define the flow:

- Broadwell-DE LAN controller EID on SMBus
- Bus Owner EID on SMBus
- Bus Owner SMBus address
- BMC EID on SMBus
- BMC SMBus address
- Broadwell-DE LAN controller SMBus address

All these variables are initialized to zero at power on apart from the SMBus address of the endpoint (the Intel® Xeon® Processor D-1500 Product Family LAN controller) which might be initialized from a NVM value.

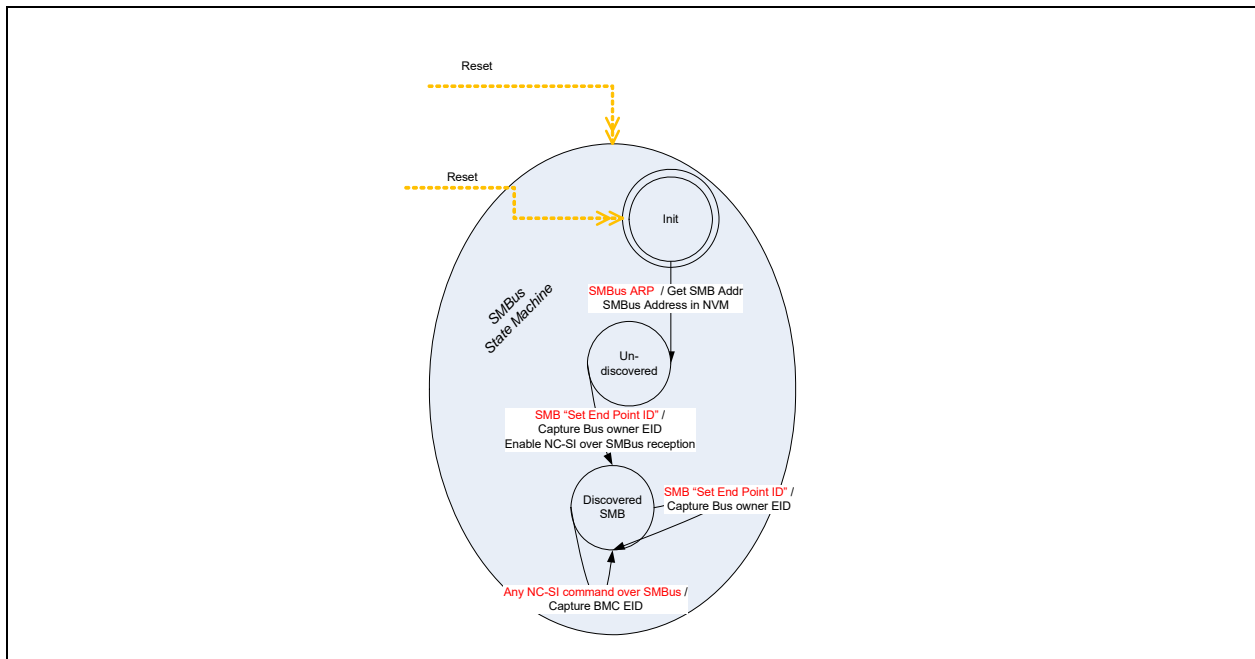


Figure 10.12.MCTP Bus Transition State Machine

10.7.2.2.1 Initial Assignment Flow

- At power on, the Intel® Xeon® Processor D-1500 Product Family LAN controller or BMC MCTP channel is connected to the SMBus, is not assigned an EID and is in undiscovered state.
- The bus owner might preform an SMBus ARP cycle to assign an SMBus address to the Intel® Xeon® Processor D-1500 Product Family LAN controller or to the BMC. Otherwise, a fixed address may be used. It is assumed that the SMBus address does not change after initialization time.
- The bus owner performs an EID assignment using a *Set Endpoint ID* MCTP command. The Intel® Xeon® Processor D-1500 Product Family LAN controller or the BMC captures the SMBus address of the bus owner from the *SMBus Source Slave address* field, the bus owner EID from the source



endpoint ID field and the Intel® Xeon® Processor D-1500 Product Family LAN controller/BMC EID from the destination endpoint ID field in the MCTP header as described in section 10.3 of DSP0236. The Intel® Xeon® Processor D-1500 Product Family LAN controller/BMC is now in discovered state.

- The BMC might detect the Intel® Xeon® Processor D-1500 Product Family LAN controller EID using one of the two following modes:
 - Static configuration of the Intel® Xeon® Processor D-1500 Product Family LAN controller SMBus address in the BMC database and Get Routing Table Entries command to find the EID matching the SMBus address.
 - Get all endpoints through a Get Routing Table Entries command and find endpoints supporting NC-SI using the Get Message Type Support command for each endpoint.
- Once the Intel® Xeon® Processor D-1500 Product Family LAN controller is found, the BMC might send a Clear Initial State command to the Intel® Xeon® Processor D-1500 Product Family LAN controller to start the NC-SI configuration. The the Intel® Xeon® Processor D-1500 Product Family LAN controller captures the BMC SMBus address and BMC EID from any NC-SI command received.
- After the NC-SI channels are enabled, traffic might be sent using the BMC and the Intel® Xeon® Processor D-1500 Product Family LAN controller addresses previously discovered.
- The BMC might send a Get UUID command to get a unique identifier of the the Intel® Xeon® Processor D-1500 Product Family LAN controller that might be used later for re-connection upon topology changes.
- If a Firmware Reset occurs, a Discovery Notify MCTP message should be sent by the Intel® Xeon® Processor D-1500 Product Family LAN controller to restart the flow.


10.7.3 MCTP Over SMBus

The message format used for NC-SI over MCTP over SMBus is as follows:

+0								+1								+2								+3									
7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0		
Destination Slave Address								0	Command Code = MCTP = 0Fh								Byte count								Source Slave Address								1
MCTP Reserved				Header version = 1				Destination endpoint ID								Source endpoint ID								S	E	SEQ#	T	Tag					
I	C	I	Message Type = 0x02					NC-SI Command/Pass Through data																									
.....																																	
NC-SI Command/Pass Through data																																	
PEC																																	

1. IC = 0

SMBus header/PEC
MCTP header
NC-SI header and payload



10.7.3.1 SMBus Discovery Process

The Intel® Xeon® Processor D-1500 Product Family LAN controller follows the discovery process described in section 6.5 of the MCTP SMBus/I2C Transport Binding Specification (DSP0237). It indicates support for ASF in the SMBus getUID command (see [Section 10.5.5.4](#)). It responds to any SMBus command using the MCTP command code - so that the bus owner knows the Intel® Xeon® Processor D-1500 Product Family LAN controller supports MCTP.

Note: MCTP commands over SMBus are received from any master address and are answered to the sender. There is no capturing of the bus owner address from any specific command.

10.7.3.2 MCTP over SMBus Special Features

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports the following optional feature of MCTP when running over SMBus:

1. Simplified MCTP mode.
2. Fairness arbitration.

10.7.3.2.1 Simplified MCTP Mode

For some point-to-point implementations of MCTP the assembly process is simplified. In this mode, the Destination EID, Source EID, Packet sequence number, Tag Owner (TO) bit and Message tag are ignored and the assembly is based only on the SOM and EOM bits. This bit is set according to the *Simplified MCTP* bit in the MCTP configuration word in the NVM.

This mode is relevant only for MCTP over SMBus traffic and when the Redirection Sideband Interface is set to 10b (MCTP over SMBus only - no pass through).

10.7.3.2.2 Fairness Arbitration

When sending MCTP messages over SMBus and when fairness arbitration is enabled, the Intel® Xeon® Processor D-1500 Product Family LAN controller should respect the fairness arbitration as defined in section 5.13 of DSP0237 when sending MCTP messages.

10.7.4 NC-SI over MCTP

The Intel® Xeon® Processor D-1500 Product Family LAN controller support for NC-SI over MCTP is similar to the support for NC-SI over RMII with the following exceptions:

1. A set of new NC-SI OEM commands used to expose the NC-SI over MCTP capabilities.
2. The format of the packets is modified to account for the new transport layer described in the sections that follow.

10.7.4.1 NC-SI Packets Format

NC-SI over MCTP defines two different message type for pass through and for control packets.

Packets with a message type equal to the *Control packets message type* field (default = 0x02) in the NVM are NC-SI control packets (commands, responses and AENs) and packets with a message type equal to the *Pass through packets message type* field (default = 0x03) in the NVM are NC-SI pass through packets



10.7.4.1.1 Control Packets

The format used for control packets (commands, responses and AENs) is as follows:

+0								+1								+2								+3									
7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0		
SMBus or PCIe header																																	
MCTP Reserved				Header version = 1				Destination endpoint ID								Source endpoint ID								S O M		E O M		SEQ#		T O = 1		Tag	
I C = 0		Message Type = Control Packets Message type (0x02)						MC ID = 0x00								Header revision								Reserved									
IID								Command								Channel ID ¹								Reserved				Payload Length[11:8]					
Payload Length[7:0]								Reserved																									
Reserved								Reserved																									
Reserved								Command Data																									
....																																	
Command Data																Checksum																	
Checksum																																	

1. The channel ID is defined as described in [Section 10.2.2.2](#).



Note that the MAC header and MAC FCS present when operating over NC-SI are not part of the packet in MCTP mode.

10.7.4.1.2 Pass-Through Packets

The format used for pass-through packets is as follows. This format is the same for either packets received from the network or packets received from the host.



The CRC is never included in the packet. In receive, the CRC is checked and removed by the Intel® Xeon® Processor D-1500 Product Family LAN controller in transmit, the CRC is added by the Intel® Xeon® Processor D-1500 Product Family LAN controller.

+0								+1								+2								+3													
7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0						
SMBus or PCIe header																																					
MCTP Reserved				Header version = 1				Destination endpoint ID								Source endpoint ID								S O M		E O M		SEQ#		T O = 1		Tag					
I C = 0		Message Type = Pass Through Packets Control Type						DA																													
DA																SA																					
SA																																					
SA								Ether type																Ethernet Packet													
Ethernet packet																																					
....																																					
....																																					

10.7.5 MCTP Programming

The MCTP programming model is based on:

1. A set of MCTP commands used for the discovery process and for the link management. The list of supported commands is described in section [Section 10.7.5.1](#).
2. A subset of the NC-SI commands used in the regular NC-SI interface, including all the OEM commands as described in [Section 10.6.2](#) (NC-SI programming I/F). The specific commands supported are listed in [Table 10.31](#) and [Table 10.34](#).

Note: For all MCTP commands (both native MCTP commands and NCSI over MCTP), the response uses the Msg tag received in the request with *TO* bit cleared.



10.7.5.1 MCTP Commands Support


Table 10.49 lists the MCTP commands supported by the Intel® Xeon® Processor D-1500 Product Family LAN controller.

Table 10.49. MCTP Commands Support

Command Code	Command Name	General Description	Intel® Xeon® Processor D-1500 Product Family LAN Controller Support as Initiator	Intel® Xeon® Processor D-1500 Product Family LAN Controller Support as Responder
0x00	Reserved	Reserved.	–	–
0x01	Set Endpoint ID	Assigns an EID to the endpoint at the given physical address.	N/A	Yes
0x02	Get Endpoint ID	Returns the EID presently assigned to an endpoint. Also returns information about what type the endpoint is and its level of use of static EIDs. See Section 10.7.5.1.1 for details.	No	Yes
0x03	Get Endpoint UUID	Retrieves a per-device unique UUID associated with the endpoint. See Section 10.7.5.1.2 for details.	No	Yes
0x04	Get MCTP Version Support	Lists which versions of the MCTP control protocol are supported on an endpoint. See Section 10.7.5.1.3 for details.	No	Yes
0x05	Get Message Type Support	Lists the message types that an endpoint supports. See Section 10.7.5.1.3.1 for details.	No	Yes
0x06	Get Vendor Defined Message Support	Used to discover an MCTP endpoint's vendor specific MCTP extensions and capabilities. See Section 10.7.5.1.4 for details.	No	Yes ¹
0x07	Resolve Endpoint ID	Used to get the physical address associated with a given EID.	No	N/A
0x08	Allocate Endpoint IDs	Used by the bus owner to allocate a pool of EIDs to an MCTP bridge.	N/A	N/A
0x09	Routing Information Update	Used by the bus owner to extend or update the routing information that is maintained by an MCTP bridge.	N/A	N/A
0x0A	Get Routing Table Entries	Used to request an MCTP bridge to return data corresponding to its present routing table entries.	No	N/A
0x0B	Prepare for Endpoint Discovery	Used to direct endpoints to clear their discovered flags to enable them to respond to the Endpoint Discovery command.	N/A	Yes ¹
0x0C	Endpoint Discovery	Used to discover MCTP-capable devices on a bus, provided that another discovery mechanism is not defined for the particular physical medium.	No	Yes ¹
0x0D	Discovery Notify	Used to notify the bus owner that an MCTP device has become available on the bus.	Yes ¹	N/A
0x0E	Get Network ID	Used to get the MCTP network ID.	No	No
0x0F	Query Hop	Used to discover what bridges, if any, are in the path to a given target endpoint and what transmission unit sizes the bridges will pass for a given message type when routing to the target endpoint.	No	No

1. These commands are supported only for MCTP over PCIe.

10.7.5.1.1 Get Endpoint ID



The Get Endpoint ID response of the Intel® Xeon® Processor D-1500 Product Family LAN controller is described in the following table:

Byte	Description	Value
1	Completion Code	
2	Endpoint ID	0x00 - EID not yet assigned. Otherwise - returns EID assigned using Set Endpoint ID command.
3	Endpoint Type	0x00 (Dynamic EID, Simple Endpoint).
4	Medium Specific	SMBUs: 0x01 - Fairness arbitration protocol supported. PCIe: 0x00.

10.7.5.1.2 Get Endpoint UUID

The UUID returned is calculated according to the following function:

Time Low = Read from NVM words at offset 0x9 and 0xA of Sideband Configuration Structure.

Time mid = Read from NVM word at offset 0xB of Sideband Configuration Structure

Time High and version = Read from NVM word at offset 0xC of Sideband Configuration Structure

Clock Sec and Reserved = Read from NVM word at offset 0xD of Sideband Configuration Structure

Node = MAC address as taken from the *PCI_SERL* and *PCI_SERH* registers.

10.7.5.1.3 Get MCTP Version Support

The following table lists the returned value according to the requested message type.

Byte	Description	Message Type						
		0xFF (Base)	0x00 (Control Protocol Message)	0x02 (NC-SI Over MCTP)	0x03 (Ethernet)	0x7E (PCIe Based VDM Messages)	All Other or Unsupported Messages	
1	Completion Code	0x0						0x80
2	Version Number entry count	3	3	1	1	2	0	
6:3	Version number entry	0xF1F0FF00 (1.0)	0xF1F0FF00 (1.0)	0xF1F0FF00 (1.0)	0xF1F0FF00 (1.0)	0xF1F0FF00 (1.0)	0	
9:7	Version number entry 2	0xF1F1F000 (1.1.0)	0xF1F1F000 (1.1.0)			0xF1F1F000 (1.1.0)		
13:10	Version number entry 3	0xF1F2F000 (1.2.0)	0xF1F2F000 (1.2.0)					



10.7.5.1.3.1 Get Message Type Support Command

The Get Message type support response of the Intel® Xeon® Processor D-1500 Product Family LAN controller is listed in the following table:

Byte	Description	Value
1	Completion Code	0x00.
2	MCTP Message Type Count	0x01/0x02/0x03 - The Intel® Xeon® Processor D-1500 Product Family LAN controller supports up to three additional message types, depending on the mode of operation and the bus used.
3:5	List of Message Type numbers	0x02 (NC-SI over MCTP).
		0x03 (Ethernet). If pass through is supported.
		0x7E (PCIe based VDM messages) - over PCIe only.

10.7.5.1.4 Get Vendor Defined Message Support Command

This command is supported only for MCTP over PCIe.

The Get Vendor Defined Message type support response of the Intel® Xeon® Processor D-1500 Product Family LAN controller is listed in the following table if the Vendor ID Set Selector equal 0x00:


Byte	Description	Value
1	Completion Code	0x00.
2	Vendor ID Set Selector	0xFF = No more capability sets.
2:4	Vendor ID	0x008086 (PCI ID indicator + Intel vendor ID).
5:6	Version	0x0100 (version 1.0).

10.7.5.1.5 Set Endpoint ID Command

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports the Set EID and Force EID operations defined in the Set Endpoint ID command. When operating over PCIe, the Set Discovered Flag operation is also supported. As endpoints in the Intel® Xeon® Processor D-1500 Product Family LAN controller can be set only through their own interface, Set EID and Force EID are equivalent. The Reset EID operation is not supported by the Intel® Xeon® Processor D-1500 Product Family LAN controller.

The Set Endpoint ID response of the Intel® Xeon® Processor D-1500 Product Family LAN controller is listed in the following table:

Byte	Description	Value
1	Completion Code	0x00.
2	Completion Status	[7:6] = 00 - Reserved.
		[5:4] = 00 - EID assignment accepted.
		[3:2] = 00 - Reserved.
		[1:0] = 00 - Device does not use an EID pool.
3	EID Setting	If the EID setting was accepted, this value will match the EID passed in the request. Otherwise, this value returns the present EID setting.
4	EID Pool Size	Always return a zero.



10.8 Manageability Host Interface

This section details host interaction with the manageability portion of the Intel® Xeon® Processor D-1500 Product Family LAN controller. The information within this section is only available to the host driver; the BMC does not have access.

10.8.1 HOST CSR Interface (Function 1/0)

The software device driver of all functions communicates with the manageability block through CSR access. The manageability is mapped to address space 0x15800 to 0x15FFF on the slave bus of each function.

Note: Writing to address 0x15800 from any function is targeted to the same address in the RAM.

10.8.2 Host Slave Command Interface to Manageability

This interface is used by the software device driver for several of the commands and for delivering various types of data in both directions (Manageability-to-Host and Host-to-Manageability).

The address space is separated into two areas:

- Direct access to the internal data RAM: The internal shared (between Firmware and Software) RAM is mapped to address space 0x15800 to 0x15EFF. Writing/reading to this address space goes directly to the RAM.
- Control register located at address 0x15F00.

10.8.2.1 Host Slave Command Interface Low Level Flow

This interface is used for the external host software to access the manageability subsystem. Host software writes a command block or read data structure directly from the data RAM. Host software controls these transactions through a slave access to the HICR control register (see [Section 8.2.2.16.6](#)).

The following flow shows the process of initiating a command to the manageability block:

1. Software clears the *FWSTS.FWRI* flag (clear by write one) to clear any previous firmware reset indications.
2. Software device driver takes ownership of the management host interface using the flow described in [Section 4.6](#).
3. The software device driver reads the HOST Interface Control register (See [Section 8.2.2.16.6](#)) and checks that the *Enable (HICR.En)* bit is set.
4. The software device driver writes the relevant command block into the RAM area that is mapped to addresses 0x15800-0x15EFF.
5. The software device driver sets the *Command (HICR.C)* bit in the HOST Interface Control register (See [Section 8.2.2.16.6](#)). Setting this bit causes an interrupt to the ARC (can be masked).
6. The software checks the *FWSTS.FWRI* flag to make sure a firmware reset didn't occur during the command processing. If this bit is set, the command might have failed.
7. The software device driver polls the HOST Interface Control register for the *Command (HICR.C)* bit to be cleared by firmware. The command should complete within half a second.
8. When firmware finishes with the command, it clears the *Command (HICR.C)* bit (if firmware replies with data, it should clear the bit only after the data is placed in the shared RAM area where the software device driver can read it).



If the software device driver reads the HOST Interface Control register and the *HICR.SV* bit is set to 1b, then there is a valid status of the last command in the shared RAM. If the *HICR.SV* bit is not set, then the command has failed with no status in the RAM.

On completion of access to the shared RAM, the software device driver should release ownership of the shared RAM using the flow described in [Section 4.6](#).

10.8.2.2 Host Interface Structure

10.8.2.2.1 Host Interface Command Structure

[Table 10.50](#) lists the structure used by the software device driver to send a command to firmware using the Host A Slave command interface (shared RAM mapped to addresses 0x15800-0x15EFF).

Table 10.50. Host Driver Command Structure

#Byte	Description	Bit	Value	Description
0	Command	7:0	Command Dependent	Specifies which host command to process.
1	Buffer Length	7:0	Command Length	Command data buffer length: 0 to 252, not including 32 bits of header.
2				
	Reserved	7:0	0x0	Reserved.
3	Checksum	7:0	Defined Below	Checksum signature. If the value is 0xFF, the checksum is not checked by the firmware.
255:4	Data Buffer	7:0	Command Dependent	Command Specific Data. Minimum buffer size: 0. Maximum buffer size: 252.

10.8.2.2.2 Host Interface Status Structure

[Table 10.51](#) lists the structure used by firmware to return a status to the software device driver via the Host Slave command interface. A status is returned after a command has been executed.

Table 10.51. Status Structure Returned to Host Driver (Sheet 1 of 2)

#Byte	Description	Bit	Value	Description
0	Command	7:0	Command Dependent	Command ID.
1	Buffer Length	7:0	Status Dependent	Status buffer length: 252:0.
2	Return Status	7:0	Depends on Command Executing Results.	0x1 = Status OK. 0x2 = Illegal command ID. 0x3 = Unsupported command. 0x4 = Illegal payload length. 0x5 = Checksum failed. 0x6 = Data error. 0x7 = Invalid parameter. 0x8 -0x7F Reserved. 0x80 - 0xFF Command Specific Errors. Might be used by individual commands for command specific errors.
3	Checksum	7:0	Defined in the sections that follow.	Checksum signature.



Table 10.51. Status Structure Returned to Host Driver (Sheet 2 of 2)

#Byte	Description	Bit	Value	Description
255:4	Data Buffer		Status Dependent.	Status configuration parameters. Minimum Buffer Size: 0. Maximal Buffer Size: 252. Note: If return status is not Status OK the data buffer is empty.

10.8.2.2.3 Checksum Calculation Algorithm

The host command/status structure is summed with this field cleared to 0b. The calculation is done using 8-bit unsigned math with no carry. The inverse of this sum is stored in this field (0b minus the result). Result: The current sum of this buffer (8-bit unsigned math) is 0b.

10.8.3 Host Interface Commands

10.8.3.1 Driver Info Host Command

This command is used to provide the driver information in NC-SI mode.

Table 10.52. Driver Info Host Command

Byte	Name	Bit	Value	Description
0	Command	7:0	0xDD	Driver info command.
1	Buffer Length	7:0	0x5	Port Number + 4 bytes of the driver information.
2	Reserved	7:0	0x0	Reserved.
3	Checksum	7:0		Checksum signature of the Host command.
4	Function Number	7:0	Function Number	Indicates the function currently reporting its driver information.
8:5	Driver Version	7:0	Driver Version	Numerical for driver version - should be: Byte 8:Major. Byte 7:Minor. Byte 6:Build. Byte 5:SubBuild.

Following is the status returned on this command:

Table 10.53. Driver Info Host Status

Byte	Name	Bit	Value	Description
0	Command	7:0	0xDD	Driver Info command.
1	Buffer Length	7:0	0x0	No data in return status.
2	Return Status	7:0	0x1	See Table 10.51
3	Checksum	7:0		Checksum signature.

10.8.3.2 Disable RXEN Command

This command is used to enable the software device driver to request a safe disable of receive enable.

**Table 10.54. Disable RXEN Command**

Byte	Name	Bit	Value	Description
0	Command	7:0	0xDE	Driver info command.
1	Buffer Length	7:0	0x1	1 data byte attached to this command (the port number).
2	Reserved	7:0	0x0	Reserved.
3	Checksum	7:0		Checksum signature of the Host command.
4	Port Number	7:0	Port Number	Indicates the port for which the operation is requested.

Following is the status returned on this command:

Table 10.55. OS2BMC Control Status

Byte	Name	Bit	Value	Description
0	Command	7:0	0xDE	Driver Info command.
1	Buffer Length	7:0	0x0	No data in return status.
2	Return Status	7:0	0x1	See Table 10.51 .
3	Checksum	7:0		Checksum signature.

Firmware should execute the following flow when receiving this command:

- Store OS2BMC enable bit.
- Clear OS2BMC enable bit.
- Clear RXEN.
- Restore OS2BMC enable bit to original value.
- Return from command.

10.8.3.3 Flash I/F interface

These commands enable buffers of up to 1 KB of data. In order to do this, the *buffer length* field is expanded to 2 bytes.

Note: This field in the command is in little endian order. For example, byte 1 contains the MSB portion of the buffer length and byte 2 contains the LSB portion of the buffer length.

10.8.3.3.1 Flash Read

This command is used to request a read from the Flash. This command enables access to the region of the Flash owned by the device (in case of non-shared SPI, it is the entire Flash).

This command always returns the value from the Flash, even if read within the shadow RAM boundaries.

Table 10.56. Flash Read Command (Sheet 1 of 2)

Byte	Name	Bit	Value	Description
0	Command	7:0	0x30	Flash read.
2:1	Buffer Length	15:0	0x6	



Table 10.56. Flash Read Command (Sheet 2 of 2)

Byte	Name	Bit	Value	Description
3	Checksum	7:0		Checksum signature.
7:4	Address	31:0		Address to read from the Flash.
9:8	Length to Read	15:0		How many bytes to read. Can be up to 1024.

Table 10.57. Flash Read Response

Byte	Name	Bit	Value	Description
0	Command	7:0	0x30	Flash read.
1	Buffer Length (LS Byte)	7:0		Buffer length[7:0].
2	Buffer Length (MS Byte)	7:5		Buffer length[10:8].
2	Return Status	4:0		See Table 10.51 .
3	Checksum	7:0		Checksum signature.
7:4	Address	31:0		Address to read from the Flash.
9:8	Length Read	15:0		How many bytes where actually read.
11:10	Reserved			Reserved.

10.8.3.3.2 Shadow RAM Read

This command is used to request a read from the shadow RAM. Requesting addresses above shadow RAM boundaries causes an error.

Table 10.58. Shadow RAM Read Command

Byte	Name	Bit	Value	Description
0	Command	7:0	0x31	Shadow RAM read.
2:1	Buffer Length	15:0	0x6	
3	Checksum	7:0		Checksum signature.
7:4	Address	31:0		Address to read from the Flash shadow RAM.
9:8	Length to Read	15:0		How many bytes to read. Can be up to 1024.

Table 10.59. Shadow RAM Read Response (Sheet 1 of 2)

Byte	Name	Bit	Value	Description
0	Command	7:0	0x31	Shadow RAM read.
1	Buffer Length (LS Byte)	7:0		Buffer length[7:0].
2	Buffer Length (MS bits)	7:5		Buffer length[10:8].
2	Return Status	4:0		See Table 10.51 .
3	Checksum	7:0		Checksum signature.
7:4	Address	31:0		Address to read from the Flash shadow RAM.



Table 10.59. Shadow RAM Read Response (Sheet 2 of 2)

Byte	Name	Bit	Value	Description
9:8	Length Read	15:0		How many bytes were actually read.
11:10	Reserved			Reserved.
12: Length Read+12	Read Data			The requested data read from Flash.

10.8.3.3.3 Flash Write

This command is used to update the Flash sections out of the shadow RAM. This command enables access to the writable region of the Flash owned by the device that is not part of the shadow RAM. If not all the area to erase is writable, the command returns an error.

Table 10.60. Flash Write Command

Byte	Name	Bit	Value	Description
0	Command	7:0	0x32	Flash write.
2:1	Buffer Length	15:0	0x68 + Data Write Length	
3	Checksum	7:0		Checksum signature.
7:4	Address	31:0		Address to write to the Flash.
9:8	Length to Write	15:0		How many bytes to write. Can be up to 1024.
11:10	Reserved			Reserved.
12: Length to Write +12	Write Data			The data to write to the Flash.

Table 10.61. Flash Write Response

Byte	Name	Bit	Value	Description
0	Command	7:0	0x32	Flash write.
1	Buffer Length	7:0	0x6	
2	Return Status	7:0		See Table 10.51 .
3	Checksum	7:0		Checksum signature.
7:4	Address	31:0		Address to write to the Flash.
9:8	Data Written	15:0		How many bytes were actually written.



10.8.3.3.4 Shadow RAM Write

This command is used to update the shadow RAM. It enables write to the writable parts of the shadow RAM. If not all the area to write is writable, or not all the area is in the shadow RAM range, the command returns an error.

Table 1-1. Shadow RAM Write Command

Byte	Name	Bit	Value	Description
0	Command	7:0	0x33	Shadow RAM write.
2:1	Buffer Length	15:0	0x68 + Data Write Length	
3	Checksum	7:0		Checksum signature.
7:4	Address	31:0		Address to write to the Flash shadow RAM.
9:8	Length to Write	15:0		How many bytes to write. Can be up to 1024.
11:10	Reserved			Reserved.
12: Length to Write+12	Write Data			The data to write to the shadow RAM.

Table 10.62. Shadow RAM Write Response

Byte	Name	Bit	Value	Description
0	Command	7:0	0x33	Shadow RAM write.
1	Buffer Length	7:0	0x6	
2	Return Status	7:0		See Table 10.51 .
3	Checksum	7:0		Checksum signature.
7:4	Address	31:0		Address to write to the Flash shadow RAM.
9:8	Data Written	15:0		How many bytes were actually written.

10.8.3.3.5 Flash Module Update

This command is used to update a secured module. After a successful response to this command, for a firmware code update, an Apply Update command ([Section 10.8.3.3.6](#)) must be sent to activate the new firmware.

Table 10.63. Flash Module Update Command (Sheet 1 of 2)

Byte	Name	Bit	Value	Description
0	Command	7:0	0x34	Flash module update.



Table 10.63. Flash Module Update Command (Sheet 2 of 2)

2:1	Buffer Length	15:0	0x1	
3	Checksum	7:0		Checksum signature
4	Module ID	7:0		Which module to update: firmware code 0x1. PHY firmware image 0x5. Option ROM 0xFE.

Table 10.64. Flash Module Update Response

Byte	Name	Bit	Value	Description
0	Command	7:0	0x34	Flash module update.
1	Buffer Length	7:0	0x0	
2	Return Status	7:0		See Table 10.51. If the authentication failed, an authentication error (0x80) is returned.
3	Checksum	7:0		Checksum signature.

10.8.3.3.6 Apply Update

This command is used to request the firmware to switch to the new uploaded code. This command involves a firmware reset, so no response should be expected after this command is given. The software device driver might read the *FWRESETCNT* register before and after the command is given to check if the reset took place.

Note: An Apply Update command sent not after a successful Flash Module Update command is ignored.

If after 100 ms the *FWRESETCNT* register has not increased, software should assume the command failed.

Table 10.65. Apply Update Command

Byte	Name	Bit	Value	Description
0	Command	7:0	0x38	Apply update.
2:1	Buffer Length	15:0	0x0	
3	Checksum	7:0		Checksum signature.

10.8.3.3.7 Flash Block Erase

This command is used to erase some of the Flash sections. This command enables access to the writable region of the Flash owned by the device that is not part of the shadow RAM. If not all the area to erase is writable, the command returns an error.

**Table 10.66. Flash Block Erase Command**

Byte	Name	Bit	Value	Description
0	Command	7:0	0x35	Flash block erase.
2:1	Buffer Length	15:0	0x5	
3	Checksum	7:0		Checksum signature.
7:4	Address	31:0		Address of block to erase – must be 4 KB aligned.
8	Number of Sectors to Erase	7:0		How many 4 KB sectors to erase. In order to avoid a too long command, this number should be no larger than three.

Table 10.67. Flash Block Erase Response

Byte	Name	Bit	Value	Description
0	Command	7:0	0x35	Flash block erase.
1	Buffer Length	7:0	0x5	
2	Return Status	7:0		See Table 10.51 .
3	Checksum	7:0		Checksum signature.
7:4	Address	31:0		Address of block to erase.
8	Erased Blocks	7:0		How many blocks were actually erased.

10.8.3.3.8 Shadow RAM Dump

This command is used to trigger a shadow RAM dump. It should be used after updating a module in Shadow RAM.

Table 10.68. Shadow RAM Dump Command

Byte	Name	Bit	Value	Description
0	Command	7:0	0x36	Shadow RAM dump.
2:1	Buffer Length	15:0	0x0	
3	Checksum	7:0		Checksum signature.

Table 10.69. Shadow RAM Dump Response

Byte	Name	Bit	Value	Description
0	Command	7:0	0x36	Shadow RAM dump.
1	Buffer Length	7:0	0x0	
2	Return Status	7:0		See Table 10.51 .
3	Checksum	7:0		Checksum signature.



10.8.3.3.9 Flash Info

This command is used to provide information about the Flash

Table 10.70. Flash Info Command

Byte	Name	Bit	Value	Description
0	Command	7:0	0x37	Flash information.
2:1	Buffer Length	15:0	0x0	
3	Checksum	7:0		Checksum signature.

Table 10.71. Flash Info Response

Byte	Name	Bit	Value	Description
0	Command	7:0	0x37	Flash information.
1	Buffer Length	7:0	0x8	
2	Return Status	7:0		See Table 10.51 .
3	Checksum	7:0		Checksum signature.
7:4	Flash Size	31:0		Size of Flash available to this device. In case on non-shared SPI, it is the size of the Flash.
8	JEDEC ID information – Bank	6:0		Defines the bank number of the manufacturer ID (number of 0x7F read).
	JEDEC ID information – Valid	7		The JEDEC ID information is valid (is not set in shared SPI modes).
9	JEDEC ID information – Manufacturer ID	7:0		Returns the manufacturer ID read by RDID command.
10	JEDEC ID information – Device ID 1 st byte	7:0		First byte of the device ID read by RDID command (<i>Family</i> and <i>Density</i> fields).
11	JEDEC ID information – Device ID 2 nd byte	7:0		Second byte of the device ID read by RDID command (sub-version and revision).

10.8.3.4 Internal PHY Register Host Command

This host command is used for accessing the internal PHY registers.

Table 10.72. Internal PHY Request

Byte	Name	Bit	Value	Description
0	Command	7:0	0x	Internal PHY request.
1	Buffer length	7:0	0x	
2	Reserved	7:0		
3	Checksum	7:0		Checksum signature.
4	Port Number	7:0	Port Number	Indicates the port number whose PHY should be accessed.
5	Command Type	7:0		0b = Read 1b = Write



If <Command Type> = Read				
13:6	Reserved	31:0		Reserved
If <Command Type> = Write				
13:6	Reserved	31:0		Reserved

Table 10.73. Internal PHY Request

Byte	Name	Bit	Value	Description
0	Command	7:0	0x	Internal PHY response.
1	Buffer length	7:0	0x	
2	Status	7:0		
3	Checksum	7:0		Checksum signature.
7:4	Data	7:0		If the request was a read request then this is the data read else this field is 0x0.

10.8.4 Software and Firmware Synchronization


Software and firmware synchronize accesses to shared resources in the Intel® Xeon® Processor D-1500 Product Family LAN controller through a semaphore mechanism and a shared configuration register between the host interface of the two ports and the firmware. This semaphore enables synchronized accesses to the following shared resources:

- NVM
- PHY 0 and PHY 1 registers
- MAC (LAN controller) shared registers

The *SW_FW_SYNC.REGSMP* bit is used as a semaphore mechanism between software and firmware. Once software or firmware takes control over this semaphore flag, it can access the *SW_FW_SYNC* register and claim ownership over specific resources. The *SW_FW_SYNC* includes pairs of bits (one owned by software and the other by firmware), where each pair of bits controls a different resource. A resource is owned by software or firmware when its respective bit is set. It is illegal to have both pair bits set at the same time. Following are the required sequences for gaining and releasing control over shared resources:

10.8.4.1 Gaining Control of Shared Resource by Software

- The software device driver checks that the software device driver of the other LAN function does not use the software/firmware semaphore
 - The software device driver polls the *SWSM.SMBI* bit until it is read as 0b or time expires (recommended expiration is ~10 ms + expiration time used for the *SW_FW_SYNC.REGSMP*).
 - If the *SWSM.SMBI* is found at 0b, the semaphore is taken. Note that following this read cycle the hardware auto sets the bit to 1b.
 - If time expired, it is assumed that the software device driver of the other function malfunctioned. The software proceeds to the next step.
- The software device driver checks that the firmware does not use the software/firmware semaphore and then takes its control.



- Software polls the `SW_FW_SYNC.REGSMP` bit until it is read as 0b or time expires (recommended expiration is ~50 ms). If time has expired, the software assumes that the firmware malfunctioned and proceeds to the next step, while ignoring the firmware bits in the `SW_FW_SYNC` register.
- Software takes control of the requested resource(s).
 - The software device driver reads the firmware and software bit(s) of the requested resource(s) in the `SW_FW_SYNC` register. If the bit(s) is cleared, the resource(s) is accessible [such as no other entity owns the resource(s)]. In this case the software device driver sets the software bit(s) of the requested resource(s) in the `SW_FW_SYNC` register. Software then clears the `SW_FW_SYNC.REGSMP` and `SWSM.SMBI` bits (releasing the software/firmware semaphore register), and can use the specific resource(s).
 - Otherwise (either firmware or software of the other LAN function owns the resource), software clears the `SW_FW_SYNC.REGSMP` and `SWSM.SMBI` bits and then repeats the entire process after some delay (recommended 5-10 ms).
 - If the resources are not released by the software device driver of the other LAN function in a timely manner (recommended expiration time is ~1 second), the software device driver can assume that the other software device driver malfunctioned. In that case, the software device driver should clear all software flags that it does not own (including `SW_FW_SYNC.REGSMP` bit) and then repeat the entire process once again.
 - If the resource is not released by the firmware (recommended expiration time for firmware is ~50 ms) software can assume that the firmware malfunctions. In that case, the software device driver should set the software bit(s) of the requested resource(s) while ignoring the corresponding firmware bits in the `SW_FW_SYNC` register.

Note that the firmware initializes its semaphore flags as part of its initialization flow. The software semaphores are not reset.

10.8.4.2 Releasing a Shared Resource by Software

- The software device driver takes control over the software/firmware semaphore as previously described for gaining shared resources.
- The software device driver clears the bit(s) of the released resource(s) in the `SW_FW_SYNC` register.
- The software device driver releases the software/firmware semaphore by clearing the `SW_FW_SYNC.REGSMP` and `SWSM.SMBI` bits.
- Software should delay (recommended 5-10 ms) before trying to gain the semaphore immediately following its release.

10.8.4.3 Gaining Control of Shared Resource by Firmware

- The firmware takes control over the software/firmware semaphore (`SW_FW_SYNC` register)
 - The firmware polls the `SW_FW_SYNC.REGSMP` bit until it is read as '0b or timeout expires (recommended expiration time is ~10 ms).
 - If timeout has expired, the firmware clears the `SW_FW_SYNC.REGSMP` bit (it is assumed the software device driver is not functional).
- The firmware takes ownership of the requested resources.
 - The firmware reads the software bit(s) corresponding to the requested firmware resource(s) in the `SW_FW_SYNC` register.
 - If the *Software* bit is cleared (such as the software device driver does not own the resource), firmware sets the firmware bit(s) of the requested resource(s). Firmware then clears the `SW_FW_SYNC.REGSMP` bit (releasing the software/firmware semaphore) and can use the specific resource(s).



- Otherwise (software owns the resource), the firmware clears the *SW_FW_SYNC.REGSMP* bit and then repeats the previous process after some delay (recommended delay of 5-10 ms).
 - If the resources owned by software are not released in a timely manner (~1 second), the firmware forces its ownership over the requested resources. Firmware clears the software flags of the requested resources in the *SW_FW_SYNC* register (assuming the software that set those flags is not functional).

10.8.4.4 Releasing a Shared Resource by Firmware

- Firmware takes control over the software/firmware semaphore as previously described for gaining shared resources.
- Firmware clears the bit(s) of the selected resource(s) in the *SW_FW_SYNC* register.
- Firmware releases the software/firmware semaphore by clearing the *SW_FW_SYNC.REGSMP* bit.
- Firmware should delay before trying to gain the selected resource semaphore immediately following its release (recommended 5-10 ms).

10.9 Host Isolate Support

If a BMC decides that a malicious software prevents its usage of the LAN, it might decide to isolate the Intel® Xeon® Processor D-1500 Product Family LAN controller from its software device driver. This is done using the TCO reset command ([Section 10.6.3.14](#)).

If TCO isolate is enabled in the NVM, The TCO Isolate command disables PCIe write operations to the LAN port. As the software device driver needs to access the CSR space in order to provide descriptors to the Intel® Xeon® Processor D-1500 Product Family LAN controller, this operation also stops the network traffic including OS-to-BMC and BMC-to-OS traffic as soon as the existing transmit and receive descriptor queues are exhausted.

§



11.0 Diagnostics

11.1 JTAG Test Mode Description

The Intel® Xeon® Processor D-1500 Product Family LAN controller includes a JTAG (TAP) port compliant with the *IEEE Standard Test Access Port and Boundary Scan Architecture 1149.1 Specification*. The TAP controller is accessed serially through five dedicated pins: TCK, TMS, TDI, TDO, and TRST_N. TMS, TDI, and TDO operate synchronously with TCK. TCK is independent of all other device clocks.

This interface can be used for test and debug purposes. System board interconnects can be DC tested using the boundary scan logic in pads. [Table 11-1](#) shows TAP controller related pin descriptions. [Table 11-2](#) describes the TAP instructions supported.

Table 11-1. TAP Controller Pins

Signal	I/O	Description
TCK	In	Test clock input for the test logic defined by IEEE1149.1. If using JTAG, connect to this signal ground through a 1 K Ω pull-down resistor.
TDI	In	Test Data Input. Serial test instructions and data are received by the test logic at this pin. If using JTAG, connect this signal to VCC33 through a 1 1 K Ω pull-up resistor.
TDO	O/D	Test Data Output. The serial output for the test instructions and data from the test logic defined in IEEE1149.1. If using JTAG, connect this signal to VCC33 through a 1 1 K Ω pull-up resistor.
TMS	In	Test Mode Select input. The signal received at JTMS is decoded by the TAP controller to control test operations.
TRST_N	In	JTAG Reset Input. Active low reset for the JTAG port.

Table 11-2. Main TAP Instructions Supported (Sheet 1 of 2)

Instruction	Description	Comment
BYPASS	The BYPASS command selects the Bypass register, a single-bit register connected between TDI and TDO pins. This enables more rapid movement of test data to and from other components in the system.	IEEE 1149.1 Std. Instruction
EXTEST	The EXTEST Instruction enables circuitry or wiring external to the devices to be tested. Boundary-scan register cells at outputs are used to apply stimulus while boundary-scan cells at input pins are used to capture data.	IEEE 1149.1 Std. Instruction



Table 11-2. Main TAP Instructions Supported (Sheet 2 of 2)

Instruction	Description	Comment
SAMPLE / PRELOAD	<p>The SAMPLE/PRELOAD instruction is used to enable scanning of the boundary scan register without causing interference to the normal operation of the device. Two functions can be performed by use of the sample/preload instruction.</p> <p>SAMPLE – enables a snapshot of the data flowing into and out of a device to be taken without affecting the normal operation of the Intel® Xeon® Processor D-1500 Product Family LAN controller.</p> <p>PRELOAD – enables an initial pattern to be placed into the boundary scan register cells. This enables initial known data to be present prior to the selection of another boundary-scan test operation.</p>	IEEE 1149.1 Std. Instruction
IDCODE	<p>The IDCODE instruction is forced into the parallel output latches of the instruction register during the Test-Logic-Reset TAP state. This enables the Intel® Xeon® Processor D-1500 Product Family LAN controller identification register to be selected by manipulation of the broadcast TMS and TCK signals for testing purposes, as well as by a conventional instruction register scan operation.</p> <p>The ID code value for the Intel® Xeon® Processor D-1500 Product Family LAN controller A0 is 0x11512013 (Intel's Vendor ID = 0x013, Device ID = 0x1562, Rev ID = 0x1).</p> <p>The ID code value for the Intel® Xeon® Processor D-1500 Product Family LAN controller B0 is 0x21512013 (Intel's Vendor ID = 0x013, Device ID = 0x1562, Rev ID = 0x2).</p>	IEEE 1149.1 Std. Instruction
HIGHZ	<p>The HIGHZ instruction is used to force all outputs of the device (except TDO) into a high impedance state. This instruction must select the Bypass register to be connected between TDI and TDO in the Shift-DR controller state.</p>	IEEE 1149.1 Std. Instruction

11.2 MAC Loopback Operations

Loopback operations are supported by the Intel® Xeon® Processor D-1500 Product Family LAN controller to assist with system and the Intel® Xeon® Processor D-1500 Product Family LAN controller debug. Loopback operation can be used to test transmit and receive aspects of software device drivers, as well as to verify electrical integrity of the connections between the Intel® Xeon® Processor D-1500 Product Family LAN controller and the system (such as PCIe bus connections, etc.)

The following types of loopback in the LAN interfaces might be configured:

- MAC Loopbacks (Point 1 in [Figure 11-1](#))
- Internal PHY Loopbacks (Point 2 in [Figure 11-1](#))
- External PHY Loopback (Point 3 in [Figure 11-1](#))

By setting the Intel® Xeon® Processor D-1500 Product Family LAN controller to loopback mode, packets that are transmitted towards the line are looped back to the host. The Intel® Xeon® Processor D-1500 Product Family LAN controller is fully functional in these modes, just not transmitting data over the lines if the loopback is internal. [Figure 11-1](#) shows the points of loopback.



For more details on the usage and loopback test setup, refer to the *Intel® Ethernet Controllers Loopback Modes* application note.

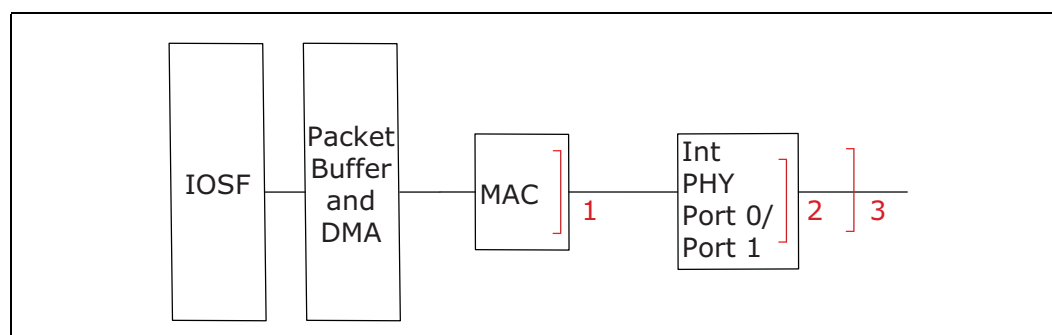


Figure 11-1. Loopback Modes

11.2.1 MAC Loopback

The Intel® Xeon® Processor D-1500 Product Family LAN controller supports three types of MAC loopbacks:

- Tx data is looped back to the MAC before entering the PHY. The PHY is not functioning and is just providing the clocks to the MAC at a rate that corresponds to the forced link speed.
- Tx data is looped back to the MAC before entering the PHY. Tx data is sent over the lines by the PHY, which is functioning.
- Rx data is looped back to Tx path at the MAC level.

11.2.1.1 Tx->Rx MAC Loopback

This loopback is closed on the internal XGMII interface of the MAC core.

To configure the Intel® Xeon® Processor D-1500 Product Family LAN controller for Tx->Rx loopback operation:

- Disable auto-negotiation in PHY register bit 7.0.C.
- Operate only at 10 GbE speed while no link partner is present. For other speeds, establish the link with a link partner at the desired speed while performing Tx -> Rx MAC loopback.
- In the MACC register, set the *FLU* bit to 1b in order to force link up.
- In the HLREG0 register, set the LPBK bit to 1b.

11.2.1.2 Rx->Tx MAC Loopback


This loopback is closed in the internal XGMII interface.

To configure the Intel® Xeon® Processor D-1500 Product Family LAN controller for Rx->Tx loopback operation, the MAC RX2TX LPBK EN bit in the MACC register should be set to 1b.

For loopback to be functional a functional link (with the partner) should be achieved (sync and alignment).

Link configuration should be done as in regular functional mode (see [Section 4.5.3](#)). All link modes can be configured.

Loopback limitation notes:



- Short preamble with minimal IPG is not supported with loopback operation.
- Transmitted data might violate the minimum IPG specification requirements.

11.2.2 PHY Loopbacks

See [Section B.8](#)

11.3 NVM Recovery Mode

NVM recovery mode is intended to recover from a misconfiguration, which can be caused by an interrupted firmware update, power failure, host software access or interrupted BMC configuration access.

Note: Device firmware implements a recovery mode that requires driver and firmware update tool software assistance to complete the recovery process. For a full discussion of recovery mode, Refer to the *Intel® Ethernet Controller X710/XXV710/XL710 NVM Recovery Mode Application Note*.



12.0 Glossary and Acronyms

Term	Definition
1's compliment	A system known as ones' complement can be used to represent negative numbers in a binary system. The ones' complement form of a negative binary number is the bitwise NOT applied to it.
2's compliment	A system of two's-complement arithmetic represents negative integers by counting backwards and wrapping around. Any number whose left-most bit is 1 is considered negative.
1000BASE-BX	1000BASE-BX is the PICMG 3.1 electrical specification for transmission of 1 Gb/s Ethernet or 1 Gb/s fibre channel encoded data over the backplane.
1000BASE-CX	1000BASE-X over specialty shielded 150 W balanced copper jumper cable assemblies as specified in IEEE 802.3 Clause 39.
1000BASE-T	1000BASE-T is the specification for 1 Gb/s Ethernet over category 5e twisted pair cables as defined in IEEE 802.3 clause 40.
AAD	Additional Authentication Data input, which is authenticated data that must be left un-encrypted.
ACK	Acknowledgment
ACPI	Advanced Configuration and Power Interface — ACPI reset is also known as D3hot-D0 transition.
AEN	Address Enable
AER	Advanced Error Reporting
AFE	Analog Front End
AH	IP Authentication Header — An IPsec header providing authentication capabilities defined in RFC 2402 For an example of an AH packet diagram see below: - Next Header: Identifies the protocol of the transferred data. - Payload Length: Size of AH packet. - RESERVED: Reserved for future use (all zero until then). - Security Parameters Index (SPI): Identifies the security parameters, which, in combination with the IP address, then identify the Security Association implemented with this packet. - Sequence Number: Monotonically increasing number, used to prevent replay attacks. Authentication Data: Contains the integrity check value (ICV) necessary to authenticate the packet; it may contain padding.
AN	Auto negotiation
AN	Association Number
APIC	Advanced Programming Interrupt Controller
APM	Advanced Power Management
APT	Advanced Pass Through mode
ARI	Alternative Routing ID capability structure— This is a new capability that allows an interpretation of the Device and Function fields as a single identification of a function within the bus.
ARI	Alternate Requester ID Interpretation
ARP	Address Resolution Protocol
b/w or BW	Bandwidth
backbone	Aa bus shared by many clients for example a management backbone or a host backbone
BAR	Base Address Register



Term	Definition
BDF	Bus/Device/Function
BER	Bit Error Rate
BIOS	Basic Input/Output System.
BIST	Built-In Self Test
BKM	Best Known Method
BMC	Baseboard Management Controller
BME	Bus Master Enable
BT	Byte Time.
BYTE alignment	Implies that the physical addresses can be odd or even. Examples: 0FECBD9A1h, 02345ADC6h.
BWG	Bandwidth Group.
CA	Secure Connectivity Association (CA): A security relationship, established and maintained by key agreement protocols, that comprises a fully connected subset of the service access points in stations attached to a single LAN that are to be supported by MACsec.
CAM	Content Addressable Memory
CCS	Current Cipher Suite
Ciphertext	Encrypted data, whose length is exactly that of the plaintext.
CNM	Congestion Notification Message
concurrency	The concurrent (simultaneous) execution of multiple interacting computational tasks. These tasks may be implemented as separate programs, or as a set of processes or threads created by a single program.
corner case	Is a problem or situation that occurs only outside of normal operating parameters — specifically one that manifests itself when multiple environmental variables or conditions are simultaneously at extreme levels. For example, a computer server may be unreliable, but only with the maximum complement of 64 processors, 512 GB of memory, and over 10,000 signed-on users. From Wiki.
CPID	Congestion Point Identifier —which should include the congestion point Ethernet MAC Address, as well as a local identifier for the local congestion entity, usually a queue in the switch.
CRC	Cyclic Redundancy Check A cyclic redundancy check (CRC) is a type of function that takes as input a data stream of unlimited length and produces as output a value of a certain fixed size. The term CRC is often used to denote either the function or the function's output. A CRC can be used in the same way as a checksum to detect accidental alteration of data during transmission or storage. CRCs are popular because they are simple to implement in binary hardware, are easy to analyze mathematically, and are particularly good at detecting common errors caused by noise in transmission channels. From Wiki
CRS	Carrier Sense Indication.
CSMA/CD	802.3 Carrier Sense Multiple Access / Collision Domain Ethernet LCI-2 Interface to an external LAN Connected Device to provide wired LAN connectivity.
CSR	Control / Status Register
CTS	Cisco Trusted Security
D0a D0 Active	Active fully operational state. Once memory space is enabled all internal clocks are activated and the LAN Controller enters an active state.
D0u D0 uninitialized	The D0u state is a low-power state used after PCI Reset (SPXB Reset) is de-asserted following power-up (cold or warm), or on D3 exit.
D3Hot	In D3 the LAN Controller only responds to PCI configuration accesses and does not generate master cycles.
D3Cold	Power Off If Vcc is removed from the device all of its PCI functions transition immediately to D3 cold. When Power is restored a PCI Reset must be asserted.



Term	Definition
Dr	Internal Power management state when minimal function is provided (WoL, Manageability)
DA	Destination Address
DAC	Digital to Analog Converter
DAC	Dual Address Cycle messages
Data Frame	FC Frames that carry read or write data.
DBU	Data Buffer Unit
DCA	Direct Cache Access
DDP	Direct Data placement
DEI	Drop Eligible Indicator (802.1Q)
DFT	Testability.
DFX	Design for *
DHCP	Dynamic Host Configuration Protocol (protocol for automating the configuration of computers that use TCP/IP)
DLLP	Data Link Layer Packet /PCIe
DMA	Direct Memory Access
DMTF NC-SI	Distributed Management Task Force BMC-NIC interconnect for management
DQ	Descriptor Queue.
DSP	Digital Signal Processor
DUT	Device Under Test
DWORD (Double-Word) alignment	Implies that the physical addresses may only be aligned on 4-byte boundaries; i.e., the last nibble of the address may only end in 0, 4, 8, or Ch. For example, 0FECBD9A8h.
EAPOL	Extensible Authentication Protocol over LAN
EAS	External Architecture Specification.
ECC	Error Correction Coding
ECRC	End to End CRC
EDB	End Data Bit
ECC	Error Correction Coding
EEPROM	Electrically Erasable Programmable Memory. A non-volatile memory located on the LAN controller that is directly accessible from the host.
EHS	External Heat Sink
EOP	End-Of-Packet; when set indicates the last descriptor making up the packet.
EP	End point
ESN	Extended Sequence Number



Term	Definition
ESP	IP Encapsulating Security Payload — An IPsec header providing encryption and authentication capabilities defined in RFC 4303. The Encapsulating Security Payload (ESP) extension header provides origin authenticity, integrity, and confidentiality protection of a packet. ESP also supports encryption-only and authentication-only configurations, but using encryption without authentication is strongly discouraged. Unlike the AH header, the IP packet header is not accounted for. ESP operates directly on top of IP, using IP protocol number 50. ESP fields: <ul style="list-style-type: none"> - Security Parameters Index (SPI): See AH - Sequence Number: See AH - Payload Data: See AH - Padding: Used with some block ciphers to pad the data to the full length of a block. - Pad Length: Size of padding in bytes. - Next Header: Identifies the protocol of the transferred data. - Authentication Data: Contains the data used to authenticate the packet.
EUI	IEEE defined 64-bit Extended Unique Identifier
Extension Header	IPv6 protocol.
Fail-over	Fail-over is the ability to detect that the LAN connection on one port is lost, and enable the other port for traffic.
FC	Fiber Channel
FC	Flow Control.
FC Exchange	Complete Fiber Channel Read or Fiber Channel Write flow. It starts with the read or write requests by the initiator (the host system) till the completion indication from the target (the remote disk).
FC Sequence	A Fiber Channel Exchange is composed of multiple Fiber Channel sequences. Fiber Channel Sequence can be a single or multiple frames that are sent by the initiator or the target. Each FC Sequence has a unique "Sequence ID".
FC Frame	Fiber Channel Frames are the smallest units sent between the initiator and the target. The FC-FS-2 spec define the maximum frame size as 2112 bytes. Each Fiber Channel frame includes an FC header and optional FC payload. It can also may include Extended headers and FC optional headers. Extended headers are not expected and FC optional headers might not be used as well.
FCP_RSP Frame	Fiber Channel control Frames that are sent from the target to the initiator which defines the completion of an FC read or write exchange.
FCS	Frame Check Sequence of Ethernet frames (a.k.a. CRC)
FEC	Forward Error Correction
FEXT	Far End Crosstalk
Firmware (FW)	Embedded code on the LAN controller that is responsible for the implementation of the NC-SI protocol and pass through functionality.
FLR	Function level reset An OS in a VM must have complete control over a device, including its initialization, without interfering with the rest of the functions.
FML	Fast Management Link
Fragment Header	An IPv6 extension Header
Frame	A unit composed of headers, data and footers that are sent or received by a device. Same as a Packet
FSM	Finite State Machine
FTS	Fast Training Sequence
GbE	Gigabit Ethernet (IEEE 802.3z-1998)
GMRP	GARP Multicast Registration Protocol (Cisco)
GPIO	General Purpose I/O
GSP	Group Strict Priority
HBA	Host Bus Adapters
Host Interface	RAM on the LAN controller that is shared between the firmware and the host. RAM is used to pass commands from the host to firmware and responses from the firmware to the host.



Term	Definition
HPC	High — Performance Computing.
HT core option	Hyper Thread Intel's trademark for implementation of the simultaneous multi-threading technology on the Pentium 4 micro architecture. It is a more advanced form of Super-threading that debuted on the Intel Xeon processors and was later added to Pentium 4 processors. The technology improves processor performance under certain workloads by providing useful work for execution units that would otherwise be idle, for example during a cache miss. A Pentium 4 with Hyper-Threading enabled is treated by the operating system as two processors instead of one. From Wiki
IANA	Internet Assigned Number Authority
ICV	128-bits Integrity Check Value (referred also as authentication tag). used for MACsec header and signature
IDS	intrusion detection systems
IFCS	Insert Frame Check Sequence of Ethernet frames
IFS	Inter Frame Spacing
IKE	Internet Key Exchange
Intel AMT	Intel Active Management Technology (Intel® AMT)
IOAT	I/O Acceleration Technology
IOH	I/O Hub
IOV	Input Output Virtualization
IOV mode	Operating through an IOVM or IOVI
IOVI	I/O Virtual Intermediary: A special virtual machine that owns the physical device and is responsible for the configuration of the physical device. Also Known As IOVM
IOVM	I/O Virtual Machine: A special virtual machine that owns the physical device and is responsible for the configuration of the physical device. Also Known As IOVI
IP tunneling	IP tunneling is the process of embedding one IP packet inside of another, for the purpose of simulating a physical connection between two remote networks across an intermediate network. IP tunnels are often used in conjunction with IPSec protocol to create a VPN between two or more remote networks across a "hostile" network such as the Internet.
IPC	Inter Processor Communication.
IP — CPMP	Carrier Performance Measurement Plan
IPG	Inter Packet Gap.
IP Sec	<p>IP security) is a suite of protocols for securing Internet Protocol (IP) communications by authenticating and/or encrypting each IP packet in a data stream. IPsec also includes protocols for cryptographic key establishment.</p> <p>IPsec is implemented by a set of cryptographic protocols for (1) securing packet flows and (2) internet key exchange. There are two families of key exchange protocols.</p> <p>The IP security architecture uses the concept of a security association as the basis for building security functions into IP. A security association is simply the bundle of algorithms and parameters (such as keys) that is being used to encrypt a particular flow. The actual choice of algorithm is left up to the users. A security parameter index (SPI) is provided along with the destination address to allow the security association for a packet to be looked up.</p> <p>For multicast, therefore, a security association is provided for the group, and is duplicated across all authorized receivers of the group. There may be more than one security association for a group, using different SPIs, thereby allowing multiple levels and sets of security within a group. Indeed, each sender can have multiple security associations, allowing authentication, since a receiver can only know that someone knowing the keys sent the data. Note that the standard doesn't describe how the association is chosen and duplicated across the group; it is assumed that a responsible party will make the choice. From Wiki</p>



Term	Definition
iSCSI	Internet SCSI (iSCSI) is a network protocol standard, officially ratified on 2003-02-11 by the Internet Engineering Task Force, that allows the use of the SCSI protocol over TCP/IP networks. iSCSI is a transport layer protocol in the SCSI-3 specifications framework. Other protocols in the transport layer include SCSI Parallel Interface (SPI), Serial Attached SCSI (SAS) and Fibre Channel. From Wiki.
ISR	Interrupt Service Routine
ITR	Interrupt Throttling
IV	Integrity Value
IV	Initialization Vector
IV	Initial Value
KaY	Key agreement entity (KaY – in 802.1AE spec terminology) i.e. control and access the off loading engine (SecY in 802.1AE spec terminology)
KVM	Keyboard — Video — Mouse
LACP	Link Aggregation Control Protocol
LAN auxiliary Power-Up	The event of connecting the LAN controller to a power source (occurs even before system power-up).
landing Zone requirements	General targets for the product.
LF	Local Fault
LDPC	Low density parity check. Coding used in 802.3an (10GBASE-T) to protect transmitted data.
LLC header	802.2 defines a special header that includes a SNAP (subnetwork access protocol) header + EtherType. Some protocols, particularly those designed for the OSI networking stack, operate directly on top of 802.2 LLC, which provides both datagram and connection-oriented network services. This 802.2 header is currently embedded in modern 802.3 frames (Ethernet II frames, aka. DIX frames). The LLC header includes two additional eight-bit address fields, called service access points or SAPs in OSI terminology; when both source and destination SAP are set to the value 0xAA, the SNAP service is requested. The SNAP header + EtherType allows EtherType values to be used with all IEEE 802 protocols, as well as supporting private protocol ID spaces. In IEEE 802.3x-1997, the IEEE Ethernet standard was changed to explicitly allow the use of the 16-bit field after the Ethernet MAC Addresses to be used as a length field or a type field. This definition is from Wiki
LLDP	Link Layer Discovery Protocol
LLINT	Low Latency Interrupt
Local Traffic	In a virtual environment traffic between virtual machines.
LOM	LAN on Motherboard.
LP	Link Partner
LSC	Link Status Change
LS	Least significant / Lowest order (for example: LS bit = Least significant bit)
LSO	Large Send Offload (a.k.a. TSO)
LSP	Link Strict Priority
LTSSM	Link Training and Status State Machine Defined in the PCIe specs.
MAC	Media Access Control.
MCH	Memory Controller Hub
MDC	Management Data Clock over MDC/MDIO lines.
MDI	Media Dependent Interface
MDIO	Management Data Input/Output Interface over MDC/MDIO lines.
MFVC	Multi-Function Virtual Channel Capability structure
MIB	Management Interface Bus



Term	Definition
MIFS/MIPG	Minimum Inter Frame Spacing/Minimum Inter Packet Gap.
MMW	Maximum Memory Window.
Mod / Modulo	In computing, the modulo operation finds the remainder of division of one number by another.
MPA	Marker PDU Aligned Framing for TCP
MPDU	MACSEC Protocol Data Unit including SecTag, User Data and ICV
MRQC	Multiple Receive Queues Command register
MS	Most significant / Highest order (for example: MS byte = Most significant byte)
MSFT	Microsoft ®
MSI	Message Signaled Interrupt
MSS	Maximum Segment Size
MTA	Multicast Table Array
MTU	Maximum Transmission Unit
NACK	Negative Acknowledgment
native mode	Used for GPIO pin that is set to be controlled by the internal logic rather than by software.
NC-SI	Network Controller — Sideband Interface
NEXT	Near End Crosstalk
NIC	Network Interface Controller.
NFTS	Number of Fast Training Signals
NFS	Network File Server
Nonce	96-bits initialization vector used by the AES-128 engine, which is distinct for each invocation of the encryption operation for a fixed key. It is formed by the AES-128 SALT field stored for that IPsec flow in the Tx SA Table, appended with the Initialization Vector (IV) field included in the IPsec packet:
NOS	Network Operating System
NPRD	Non-Posted Request Data
NRZ	non-return-to-zero signaling
NSE	No Snoop Enable
NTL	No Touch Leakage
NTP	Network Time Protocol
NVM	Non Volatile Memory
OEM	Original Equipment manufacturer
Core	Network Interface Registers
Packet	A unit composed of headers, data and footers that are sent or received by a device. Also known as a frame.
Pass Filters	Needs Definition Packets that match this type of filter continue on to their destination?
PB	Packet Buffer
PBA	The nine-digit (Printed Board Assembly) number used for Intel manufactured adapter cards.
PBA	Pending Bit Array (in MSI-X context)
PBA	Printed Board Assembly (in NVM or board context)
PCS	Physical Coding Sub layer.
PDU	Protocol Data Units



Term	Definition
PF	Physical Function (in a virtualization context).
PHY	Physical Layer Device.
Plain text	Data to be both authenticated and encrypted.
PMA	Physical Medium Attachment
PMC	Power Management Capabilities
PMD	Physical Medium Dependent.
PME	Power Management Event
PN	Packet Number (PN) in a MACsec context: Monotonically increasing value used to uniquely identify a MACsec frame in the sequence.
Pool	Virtual ports
Power State D0a	Active fully operational state. Once memory space is enabled all internal clocks are activated and the LAN Controller enters an active state.
Power State D0u	The D0u state is a low-power state used after SPXB Reset is de-asserted following power-up (cold or warm), or on D3 exit.
Power State D3Hot	A Power down state with the PCI continuing to receive a proper power supply.
Power State D3Cold	A Power down state with the PCI also in a power down state.
Power State Dr	Device state when PCIe reset is asserted.
Power State Sx	LAN Connected Device: SMBus Active and PCI Powered down.
PPM	Part Per Million
PRBS	Pseudo-Random Binary Sequence
PT	Pass Through
PTP	Precision Time Protocol
QoS	Quality of Service
CNM	Congestion Notification Message
QWORD (Quad-Word) alignment	Implies that the physical addresses may only be aligned on 8byte boundaries; i.e., the last nibble of the address may only end in 0, or 8. For example, 0FECBD9A8h.
Receive latency	measured from packet reception from the wire and until the descriptor is updated on PCIe.
RDMA	Remote Direct Memory Access
RDMAP	Remote Direct Memory Access Protocol
Relax ordering	When the strict order of packets is not required, the device can send packets in an order that allows for less power consumption and greater CPU efficiency.
RID	Requester ID
RLT	Rate-limited flag bit
RMCP	Remote Management and Control Protocol (Distributed Management Task Force)
RMII	Reduced Media Independent Interface (Reduced MII)
RMON statistics	Remote Network Monitoring or Remote Monitoring
RPC header	Remote Procedure Call
RS	Rate Scheduler
RSC	Receive Side Coalescing coalesces incoming TCP/IP (and potentially UDP/IP) packets into larger receive segments



Term	Definition
RSS	Receive-Side Scaling is a mechanism to distribute received packets into several descriptor queues. Software then assigns each queue to a different processor, therefore sharing the load of packet processing among several processors
RSTD	Reset Sequence Done
RSTI	Reset Sequence in Process
Rx, RX	Receive
SA	Security Association (in IPsec or MACSec context).
SA	Source Address (In Ethernet frame context).
SA (in a MACsec context)	Secure Association (SA): A security relationship that provides security guarantees for frames transmitted from one member of a CA to another. Each SA is supported by a single secret key, or a single set of keys where the cryptographic operations used to protect one frame require more than one key.
SAC	Single Address Cycle (SAC) messages
SAK	Security Associations Key
salt	In cryptography, a salt consists of random bits used as one of the inputs to a key derivation function. Sometimes the initialization vector, a previously generated (preferably random) value, is used as a salt. The other input is usually a password or pass phrase. The output of the key derivation function is often stored as the encrypted version of the password. A salt value can also be used as a key for use in a cipher or other cryptographic algorithm. A salt value is typically used in a hash function. from Wiki
SAP	Service Access Point –an identifying label for network endpoints used in OSI networking.
SC	Secure Channel – Authentication and key exchange
SC	Secure Channel (SC): A security relationship used to provide security guarantees for frames transmitted from one member of a CA to the others. An SC is supported by a sequence of SAs thus allowing the periodic use of fresh keys without terminating the relationship.
SCI	Secure Channel Identifier A globally unique identifier for a secure channel, comprising a globally unique Ethernet MAC Address and a Port Identifier, unique within the system allocated that address.
SCSI	Small Computer System Interface is a set of standards for physically connecting and transferring data between computers and peripheral devices. The SCSI standards define commands, protocols, and electrical and optical interfaces. SCSI is most commonly used for hard disks and tape drives, but it can connect a wide range of other devices, including scanners, and optical drives (CD, DVD, etc. From Wiki.
SCL signal	SM Bus Clock
SCTP	Stream Control Transmission Protocol
SDA signal	SM Bus Data
SDP	Software-Definable Pins
SecY	802.1AE spec terminology Security entity
Segment	subsections of a packet
SerDes	Serializer and De-Serializer Circuit.
SFD	Start Frame Delimiter
SFI	Serial Flash Interface
SGMII	Serialized Gigabit Media Independent Interface.
SKU	subsets of features of a chip that can be disabled for marketing purposes.
SNMP	Standard Network Management Protocol
SMB	Semaphore Bit
SMBus	System Management Bus. A bus that carries various manageability components, including the LAN controller, BIOS, sensors and remote-control devices.



Term	Definition
SN	Sequence Number — contains a counter value that increases by one for each Ethernet frame sent.
SNAP	Subnetwork Access Protocol
SoL	Serial Over LAN Serial Over LAN is a mechanism that enables the input and output of the serial port of a managed system to be redirected via an IPMI (Internet Protocol Multicast Initiative) session over IP.
SPD	Smart Power Down
SPI	The Security Parameter Index is an identification tag added to the header while using IPSec for tunneling the IP traffic. This tag helps the kernel discern between two traffic streams where different encryption rules and algorithms may be in use. The SPI (as per RFC 2401) is an essential part of an IPSec SA (Security Association) because it enables the receiving system to select the SA under which a received packet will be processed. An SPI has only local significance, since is defined by the creator of the SA; an SPI is generally viewed as an opaque bit string. However, the creator of an SA may interpret the bits in an SPI to facilitate local processing. from Wikipedia
SPXB interface	PCI Express Backbone
Spoofing	In computer networking, the term IP address spoofing is the creation of IP packets with a forged (spoofed) source IP address with the purpose to conceal the identity of the sender or impersonating another computing system. IP stands for Internet Protocol. from Wiki
SR-IOV	PCI-SIG single-root I/O Virtualization initiative
SW Switch acceleration mode	Central management of the networking resources by an IOVM or by the VMM. Also known as VMDq2 mode.
SWIZZLE	To convert external names, array indices, or references within a data structure into address pointers when the data structure is brought into main memory from external storage (also called pointer swizzling);
Sx	LAN Connected Device: SMBus Active and PCI Powered down.
SYN Attack	A SYN attack is a form of denial-of-service attack in which an attacker sends a succession of SYN (synchronize) requests to a target's system.
TC	Traffic Class
TCI	For 802.1q, Tag Header field Tag Control Information (TCI); 2 octets.
TCO	Total Cost of Ownership (Management)
TCP/IP	Transmission Control Protocol/Internet Protocol
TDESC	Transmit Descriptor
TDP	Total Device Power
TDR	Time Domain Reflectometry
TFCS	Transmit Flow Control Status
TLP	Transaction layer Packets
ToS	Type of Service
TPID	For 802.1q, Tag Header field Tag Protocol Identifier; 2 octets.
TPPAC	Transmit Packet Plane Arbitration Control
Transmit latency	Measured from Tail update until the packet is transmitted on the wire. It is assumed that a single packet is submitted for this traffic class and its latency is then measured in presence of traffic belonging to other traffic classes.
TS	Time Stamp
TSO	TCP or Transmit Segmentation offload — A mode in which a large TCP/UDP I/O is handled to the device and the device segments it to L2 packets according to the requested MSS.
TSS	Transmit Side Scaling
Tx, TX	Transmit



Term	Definition
UBWG	User Bandwidth Group
ULP	Upper Layer Protocol
UP	User Priority
UR	Error Reporting Unsupported Request Error
VF	Virtual Function– A part of a PF assigned to a VI
VI	Virtual Image – A virtual machine to which a part of the I/O resources is assigned. Also known as a VM.
VM	Virtual Machine
VMDq2	SW switch acceleration mode–Central management of the networking resources by an IOVM or by the VMM. Virtual Machine Devices queue (VMDq) is a mechanism to share I/O resources among several consumers. For example, in a virtual system, multiple OSs are loaded and each executes as though the whole system’s resources were at its disposal. However, for the limited number of I/O devices, this presents a problem because each OS may be in a separate memory domain and all the data movement and device management has to be done by a VMM (Virtual Machine Monitor). VMM access adds latency and delay to I/O accesses and degrades I/O performance. VMDs (Virtual Machine Devices) are designed to reduce the burden of VMM by making certain functions of an I/O device shared and thus can be accessed directly from each guest OS or Virtual Machine (VM). From Nahum
VMM	Virtual Machine Monitor
VPD	Vital Product Data (PCI protocol).
VT	Virtualization
WB	Write Back
WC	Worst Case
WoL	Wake-on-LAN Now called APM Wake up or Advanced power management Wake up.
WORD alignment	Implies that physical addresses must be aligned on even boundaries; i.e., the last nibble of the address may only end in 0, 2, 4, 6, 8, Ah, Ch, or Eh. For example, 0FECBD9A2h.
WRR	Weighted Round-Robin
WSP	Weighted Strict Priority
XAUI	10 Gigabit Attachment Unit Interface
XFP	10 Gigabit Small Form Factor Pluggable modules
XGMII	10 Gigabit Media Independent Interface
XGXS	XGMII Extender Sub layer
XMT	Transmit



NOTE: *This page intentionally left blank.*



Appendix A Packet Formats

A.1 Legacy Packet Formats

A.1.1 ARP Packet Formats

A.1.1.1 ARP Request Packet

Offset	# Of Bytes	Field	Value (In Hex)	Action
0	6	Destination Address		Compare
6	6	Source Address		Stored
12	E=(0/4/6/8)	Outer Tag (Outer VLAN or E-tag)	0x8100 **** 0x893F *****	Ignore
12 + E	S=(0/4)	Possible VLAN Tag		Stored
12 + E + S	2	Type	0x0806	Compare
14 + E + S	2	HW Type	0001	Compare
16 + E + S	2	Protocol Type	0x0800	Compare
18 + E + S	1	Hardware Size	06	Compare
19 + E + S	1	Protocol Address Length	04	Compare
20 + E + S	2	Operation	0001	Compare
22 + E + S	6	Sender HW Address	-	Stored
28 + E + S	4	Sender IP Address	-	Stored
32 + E + S	6	Target HW Address	-	Ignore
38 + E + S	4	Target IP Address	ARP IP address	Compare

A.1.1.2 ARP Response Packet

Offset	# of bytes	Field	Value
0	6	Destination Address	ARP Request Source Address.
6	6	Source Address	Programmed from NVM or BMC.
12	E=(0/4/6/8)	Outer Tag (Outer VLAN or E-tag)	
12 + E	S=(0/4)	Possible VLAN Tag	From ARP Request.



Offset (Continued)	# of bytes	Field	Value
12 + E + S	2	Type	0x0806
14 + E + S	2	HW Type	0x0001
16 + E + S	2	Protocol Type	0x0800
18 + E + S	1	Hardware Size	0x06
19 + E + S	1	Protocol Address Length	0x04
20 + E + S	2	Operation	0x0002
22 + E + S	6	Sender HW Address	Programmed from NVM or BMC.
28 + E + S	4	Sender IP Address	Programmed from NVM or BMC.
32 + E + S	6	Target HW Address	ARP Request Sender HW Address.
38 + E + S	4	Target IP Address	ARP Request Sender IP Address.

A.1.1.3 Gratuitous ARP Packet

Offset	# of bytes	Field	Value
0	6	Destination Address	Broadcast address.
6	6	Source Address	
12	E=(0/4/6/8)	Outer Tag (Outer VLAN or E-tag)	
12 + E	S=(0/4)	Possible VLAN Tag	
12 + E + S	D=(0/8)	Possible Length + LLC/SNAP Header	
12 + E + S + D	2	Type	0x0806
14 + E + S + D	2	Hardware Type	0x0001
16 + E + S + D	2	Protocol Type	0x0800
18 + E + S + D	1	Hardware Size	0x06
19 + E + S + D	1	Protocol Address Length	0x04
20 + E + S + D	2	Operation	0x0001
22 + E + S + D	6	Sender Hardware Address	
28 + E + S + D	4	Sender IP Address	
32 + E + S + D	6	Target Hardware Address	
38 + E + S + D	4	Target IP Address	

A.1.2 IP and TCP/UDP Headers for TSO

This section outlines the format and content for the IP, TCP and UDP headers. The Intel® Xeon® Processor D-1500 Product Family LAN controller requires baseline information from the device driver in order to construct the appropriate header information during the segmentation process.

Note: The IP header is first shown in the traditional (like RFC 791) representation, and because byte and bit ordering is confusing in that representation, the IP header is also shown in Little Endian format. The actual data is fetched from memory in Little Endian format.



0 1 2 3 4 5 6 7								8 9 0 ¹ 1 2 3 4 5								6 7 8 9 0 ² 1 2 3								4 5 6 7 8 9 0 ³ 1							
Version				IP Hdr Length				TYPE of service				Total length (IP header + payload length)																			
Identification																Flags				Fragment Offset											
Time to Live								Layer 4 Protocol ID								Header Checksum															
Source Address																															
Destination Address																															
Options																															

Figure A.1 IPv4 Header (Traditional Representation - most left byte first on the wire)

Byte3								Byte2								Byte1								Byte0																	
7 6 5 4 3 2 1 0								7 6 5 4 3 2 1 0								7 6 5 4 3 2 1 0								7 6 5 4 3 2 1 0																	
Total length LSB ----- MSB																TYPE of service								Version				IP Hdr Length													
Fragment Offset Low												R E S		N F		M F		Fragment Offset High								Identification LSB ----- MSB															
Header Checksum LSB ----- MSB																Layer 4 Protocol ID								Time to Live																	
Source Address LSB																MSB																									
Destination Address LSB																MSB																									
Options																																									

Figure A.2 IPv4 Header (Little Endian Order - Byte 0 first on the wire)

Identification is increased on each packet.

Flags Field Definitions:

The Flags field is defined as follows. Note that hardware does not evaluate or change these bits.

- MF - More Fragments
- NF - No Fragments
- Reserved

The Intel® Xeon® Processor D-1500 Product Family LAN controller does TCP segmentation, not IP Fragmentation. IP Fragmentation might occur in transit through a network's infrastructure.

0 1 2 3 4 5 6 7								8 9 0 ¹ 1 2 3 4 5								6 7 8 9 0 ² 1 2 3								4 5 6 7 8 9 0 ³ 1							
Version				Priority				Flow Label																							
Payload Length (excluding the IP header length)																Next Header Type								Hop Limit							
Source Address MSB																LSB															
Destination Address MSB																LSB															



Extensions (if any)

Figure A.3 IPv6 Header (Traditional Representation - most left byte first on the wire)

Byte3	Byte2	Byte1	Byte0
7 6 5 4 3 2 1 0	7 6 5 4 3 2 1 0	7 6 5 4 3 2 1 0	7 6 5 4 3 2 1 0
Flow Label LSB		MSB	
Version		Priority	
Hop Limit	Next Header Type	Payload Length (excluding the IP header length) LSB ----- MSB	
Source Address LSB		MSB	
Destination Address LSB		MSB	
Extensions			

Figure A.4 IPv6 Header (Little Endian Order - byte 0 first on the wire)

A TCP or UDP frame uses a 16 bit wide one's complement checksum. The checksum word is computed on the outgoing TCP or UDP header and payload, and on the pseudo header. Details on checksum computations are provided in [Section 7.2.4.6](#).

Note: TCP and UDP over IPv6 requires the use of checksum, where it is optional for UDP over IPv4.

Note: The TCP header is first shown in the traditional (such as RFC 793) representation, and because byte and bit ordering is confusing in that representation, the TCP header is also shown in Little Endian format. The actual data is fetched from memory in Little Endian format.

0 1 2 3 4 5 6 7								1 8 9 0 1 2 3 4 5								2 6 7 8 9 0 1 2 3								3 4 5 6 7 8 9 0 1							
Source Port																Destination Port															
Sequence Number																															
Acknowledgement Number																															
TCP Header Length				Reserved				U R G	A C K	P S H	R S T	S Y N	F I N	Window																	
Checksum																Urgent Pointer															
Options																															

Figure A.5 TCP Header (Traditional Representation)

Byte3	Byte2	Byte1	Byte0
7 6 5 4 3 2 1 0	7 6 5 4 3 2 1 0	7 6 5 4 3 2 1 0	7 6 5 4 3 2 1 0
Destination Port		Source Port	



LSB		Sequence Number								MSB			
Acknowledgement Number													
Window					RE S	U R G	A C K	PSH	R S T	S Y N	FIN	TCP Header Length	Reserved
Urgent Pointer					Checksum								
Options													

Figure A.6 TCP Header (Little Endian)

The TCP header is always a multiple of 32-bit words. TCP options might occupy space at the end of the TCP header and are a multiple of 8 bits in length. All options are included in the checksum.

The checksum also covers a 96-bit pseudo header prefixed to the TCP header (see [Figure A.7](#)). For IPv4 packets, this pseudo header contains the IP source address, the IP destination address, the *IP Protocol* field, and TCP length. Software pre-calculates the partial pseudo header sum, that includes IPv4 SA, DA and protocol types, but NOT the TCP length, and stores this value into the TCP checksum field of the packet. For both IPv4 and IPv6, hardware needs to factor in the TCP length to the software supplied pseudo header partial checksum.

Note: When calculating the TCP pseudo header, the byte ordering can be tricky. One common question is whether the *Protocol ID* field is added to the lower or upper byte of the 16-bit sum. The *Protocol ID* field should be added to the Least Significant Byte (LSB) of the 16-bit pseudo header sum, where the Most Significant Byte (MSB) of the 16-bit sum is the byte that corresponds to the first checksum byte out on the wire.

The *TCP Length* field is the TCP header length including option fields plus the data length in bytes, which is calculated by hardware on a frame-by-frame basis. The TCP length does not count the 12 bytes of the pseudo header. The TCP length of the packet is determined by hardware as:

$$\text{TCP Length} = \min(\text{MSS}, \text{PAYLOADLEN}) + \text{L5_LEN}$$

The two flags that might be modified are defined as:

- PSH: receiver should pass this data to the application without delay
- FIN: sender is finished sending data

The handling of these flags is described in [Section 7.2.4.7](#).

Payload is normally MSS except for the last packet where it represents the remainder of the payload.

IPv4 Source Address		
IPv4 Destination Address		
Zero	Layer 4 Protocol ID	TCP/UDP Length

Figure A.7 TCP/UDP Pseudo Header Content for IPv4 (Traditional Representation)

IPv6 Source Address	
IPv6 Final Destination Address	
TCP/UDP Packet Length	
Zero	Next Header

Figure A.8 TCP/UDP Pseudo Header Content for IPv6 (Traditional Representation)

Note: From RFC2460:

- If the IPv6 packet contains a routing header, the destination address used in the pseudo-header is that of the final destination. At the originating node, that address is in the last element of the routing header; at the recipient(s), that address is in the *Destination Address* field of the IPv6 header.
- The next header value in the pseudo-header identifies the upper-layer protocol (such as 6 for TCP, or 17 for UDP). It differs from the next header value in the IPv6 header if there are extension headers between the IPv6 header and the upper-layer header.
- The upper-layer packet length in the pseudo-header is the length of the upper-layer header and data (such as TCP header plus TCP data). Some upper-layer protocols carry their own length information (like the *Length* field in the UDP header); for such protocols, that is the length used in the pseudo- header. Other protocols (such as TCP) do not carry their own length information, in which case the length used in the pseudo-header is the payload length from the IPv6 header, minus the length of any extension headers present between the IPv6 header and the upper-layer header.
- Unlike IPv4, when UDP packets are originated by an IPv6 node, the UDP checksum is not optional. That is, whenever originating a UDP packet, an IPv6 node must compute a UDP checksum over the packet and the pseudo-header, and, if that computation yields a result of zero, it must be changed to hex FFFF for placement in the UDP header. IPv6 receivers must discard UDP packets containing a zero checksum, and should log the error.

A type 0 routing header has the following format:

Next Header	Hdr Ext Len	Routing Type "0"	Segments Left "n"
Reserved			
Address[1]			
Address[2]			
...			
Final Destination Address [n]			

Figure A.9 IPv6 Routing Header (Traditional Representation)

- Next Header - 8-bit selector. Identifies the type of header immediately following the routing header. Uses the same values as the *IPv4 Protocol* field [RFC-1700 et seq.].
- Hdr Ext Len - 8-bit unsigned integer. Length of the routing header in 8-octet units, not including the first 8 octets. For the type 0 routing header, *Hdr Ext Len* is equal to two times the number of addresses in the header.
- Routing Type - 0.
- Segments Left - 8-bit unsigned integer. Number of route segments remaining. For example, the number of explicitly listed intermediate nodes still to be visited before reaching the final destination. Equal to "n" at the source node.

Reserved - 32-bit reserved field. Initialized to zero for transmission; ignored on reception.

- Address[1...n] - Vector of 128-bit addresses, numbered 1 to n.

The UDP header is always 8 bytes in size with no options.



0 1 2 3 4 5 6 7	8 9 0 ¹ 1 2 3 4 5	6 7 8 9 0 ² 1 2 3	4 5 6 7 8 9 0 ³ 1
Source Port		Destination Port	
Length		Checksum	

Figure A.10 UDP Header (Traditional Representation)

Byte3	Byte2	Byte1	Byte0
0 1 2 3 4 5 6 7	0 1 2 3 4 5 6 7	0 1 2 3 4 5 6 7	0 1 2 3 4 5 6 7
Destination Port		Source Port	
Checksum		Length	

Figure A.11 UDP Header (Little Endian Order)

UDP pseudo header has the same format as the TCP pseudo header. The pseudo header prefixed to the UDP header contains the IPv4 source address, the IPv4 destination address, the IPv4 protocol field, and the UDP length (same as the TCP length previously discussed). This checksum procedure is the same as is used in TCP.

Unlike the TCP checksum, the UDP checksum is optional (for IPv4). Software must set the *TXSM* bit in the TCP/IP context transmit descriptor to indicate that a UDP checksum should be inserted. Hardware does not overwrite the UDP checksum unless the *TXSM* bit is set.

A.1.3 Magic Packet

A Magic Packet is a broadcast frame, but could also be a multicast or unicast Ethernet MAC addresses. The Intel® Xeon® Processor D-1500 Product Family LAN controller accepts this packet if it matches any of its pre-programmed Ethernet MAC addresses. Magic packet can be sent over a variety of connection-less protocols (usually UDP or IPX). The Magic Packet pattern is composed of the following sequences:

- Synchronization stream composed of 6 bytes equal to 0xFF 0xFF 0xFF 0xFF 0xFF 0xFF
- Unique pattern composed of 16 times the end node Ethernet MAC address. The Intel® Xeon® Processor D-1500 Product Family LAN controller expects the Ethernet MAC Address stored in the RAL[0] and RAH[0] registers.

The Intel® Xeon® Processor D-1500 Product Family LAN controller looks for the synchronization pattern and the sequence of 16 Ethernet MAC addresses. It does not check the packet content and the length of the header that precedes the magic pattern nor any data that follows it.

A.2 Packet Types for Packet Split Filtering

The following packet types are supported by the packet split feature. This section describes the packets from the split-header point of view. This means that when describing the different fields that are checked and compared, it emphasizes only the fields that are needed to calculate the header length. This document describes the checks that are done after the decision to pass the packet to the host memory was done.

Terminology:

- Compare - The field values are compared and must be exactly equal to the value specified in this document.



- Checked - The field values are checked for calculation (header length ...).
- Ignore - The field values are ignored but the field is counted as part of the header.

A.2.1 Type 1.1: Ethernet (VLAN/SNAP) IP Packets

A.2.1.1 Type 1.1: Ethernet, IP, Data

This type contains only Ethernet header and Ipv4 header while the payload header of the IP is not Ipv6/TCP/UDP.

Offset	# of bytes	Field	Value	Action	Comment
0	6	Destination Address		Ignore	MAC Header – processed by main address filter, or broadcast.
6	6	Source Address		Ignore	
12	D=(0/4/6/8)	Outer Tag (Outer VLAN or E-tag)	0x8100 **** 0x893F *****	Ignore	
12+D	S=(0/4)	Possible VLAN Tag	0x8100 ****	Compare	
12+D+S	D=(0/8)	Possible Length + LLC/ SNAP Header		Compare	
12+D+S	2	Type	0800h	Compare	IP
IPv4 Header					
14+D+S	1	Version/ HDR Length	0x4X	Compare	Check IPv4 and header length.
15+D+S	1	Type of Service	-	Ignore	
16+D+S	2	Packet Length	-	Ignore	
18+D+S	2	Identification	-	Ignore	
20+D+S	2	Fragment Info	>0 or MF bit is set	Check	Check that the packet is fragmented.
22+D+S	1	Time to Live	-	Ignore	
23+D+S	1	Protocol		Ignore	Has no meaning if the packet is fragmented.
24+D+S	2	Header Checksum	-	Ignore	
26+D+S	4	Source IP Address	-	Ignore	
30+D+S	4	Destination IP Address	-	Ignore	
34+D+S	N	Possible IP Options		Ignore	

A.2.1.2 Type 1.2: Ethernet (SNAP/VLAN), IPv4, UDP

This type contains only Ethernet header, Ipv4 header, and UDP header.



Offset	# of bytes	Field	Value	Action	Comment
0	6	Destination Address		Ignore	MAC Header – processed by main address filter, or broadcast filter.
6	6	Source Address		Ignore	
12	D=(0/4/6/8)	Outer Tag (Outer VLAN or E-tag)	0x8100 **** 0x893F *****	Ignore	
12+D	S=(0/4)	Possible VLAN Tag	0x8100 ****	Check	
12+D+S	2	Type	0x0800	Compare	IP
IP Header					
14+D+S	1	Version/ HDR Length	0x4X	Compare	Check IPv4 and header length.
15+D+S	1	Type of Service	-	Ignore	
16+D+S	2	Packet Length	-	Ignore	
18+D+S	2	Identification	-	Ignore	
20+D+S	2	Fragment Info	(xx00) 0x000	Compare	
22+D+S	1	Time to Live	-	Ignore	
23+D+S	1	Protocol	0x11	Compare	UDP header.
24+D+S	2	Header Checksum	-	Ignore	
26+D+S	4	Source IP Address	-	Ignore	
30+D+S	4	Destination IP Address	-	Ignore	
34+D+S	N	Possible IP Options		Ignore	
UDP Header					
34+D+S+N	2	Source Port	Not (0x801)	Check	Not NFS packet.
36+D+S+N	2	Destination Port	Not (0x801)	Check	Not NFS packet.
38+D+S+N	2	Length	-	Ignore	
40+D+S+N	2	Checksum	-	Ignore	

In this case, the packet is split after (42+D+S+N) bytes.

A.2.1.3 Type 1.3: Ethernet (VLAN/SNAP) IPv4 TCP

This type contains only Ethernet header, Ipv4 header, and TCP header.



Offset	# of bytes	Field	Value	Action	Comment
0	6	Destination Address		Ignore	MAC Header – processed by main address filter, or broadcast filter.
6	6	Source Address		Ignore	
12	D=(0/4/6/8)	Outer Tag (Outer VLAN or E-tag)	0x8100 **** 0x893F *****	Ignore	
12+D	S=(0/4)	Possible VLAN Tag	0x8100 ****	Check	
12+D+S	2	Type	0x0800	Compare	IP
IPv4 Header					
14+D+S	1	Version/ HDR Length	0x4X	Compare	Check IPv4 and header length.
15+D+S	1	Type of Service	-	Ignore	
16+D+S	2	Packet Length	-	Ignore	
18+D+S	2	Identification	-	Ignore	
20+D+S	2	Fragment Info	0x00	Compare	
22+D+S	1	Time to Live	-	Ignore	
23+D+S	1	Protocol	0x06	Compare	TCP header.
24+D+S	2	Header Checksum	-	Ignore	
26+D+S	4	Source IP Address	-	Ignore	
30+D+S	4	Destination IP Address	-	Ignore	
34+D+S	N	Possible IP Options		Ignore	
TCP Header					
34+D+S+N	2	Source Port	Not (0x801)	Check	Not NFS packet.
36+D+S+N	2	Destination Port	Not (0x801)	Check	Not NFS packet.
38+D+S+N	4	Sequence Number	-	Ignore	
42+D+S+N	4	Acknowledge Number	-	Ignore	
46+D+S+N	1/2	Header Length		Check	
46.5+D+S+N	1.5	Different Bits	-	Ignore	
48+D+S+N	2	Window Size	-	Ignore	
50+D+S+N	2	TCP Checksum	-	Ignore	
52+D+S+N	2	Urgent Pointer	-	Ignore	
54+D+S+N	F	TCP Options	-	Ignore	

In this case, the packet is split after (54+D+S+N+F) bytes.

- N = (IP HDR length - 5) * 4.
- F = (TCP header length - 5) * 4.




A.2.1.4 Type 1.4: Ethernet IPv4 IPv6

A.2.1.4.1 IPv6 Header Options Processing

This type of processing looks at the next-header field and header length in order to determine the identity of the next-header processes, the IPv6 options, and it's length.

If the next header in the IPv6 header is equal to 0x00/0x2B/0x2C/0x3B/0x3c it means that the next header is an IPv6 option header and this is its structure:



Header Len determines the length of the header while the next header field determines the identity of the next header (should be any IPv6 extension header for another IPv6 header option).

A.2.1.4.2 IPv6 Next Header Values

When parsing an IPv6 header does not parse all possible extension headers and if there is an extension header that is not supported then the packet is treated as an unknown payload after the IPv6 header.

Value	Header type
0x00	Hop-by-Hop
0x2B	Routing
0x2C	Fragment
0x3B	No next header (EOL)
0x3C	Destination option header

- The next header in a fragment header is ignored and this extension header is expected to be the last header.

A.2.1.4.3 Type 1.4.1: Ethernet IPv4 IPv6 Data

This type contains only Ethernet header, IPv4 header, and IPv6 header.

Offset	# of bytes	Field	Value	Action	Comment
0	6	Destination Address		Ignore	MAC Header – processed by main address filter, or broadcast filter.
6	6	Source Address		Ignore	
12	D=(0/4/6/8)	Outer Tag (Outer VLAN or E-tag)	0x8100 **** 0x893F *****	Ignore	
12+D	S=(0/4)	Possible VLAN Tag	0x8100	Check	
12+D+S	2	Type	0x0800	Compare	IP
IP4 Header					



Offset	# of bytes	Field	Value	Action	Comment
14+D+S	1	Version/ HDR Length	0x4X	Compare	Check IPv4 and header length.
15+D+S	1	Type of Service	-	Ignore	
16+D+S	2	Packet Length	-	Ignore	
18+D+S	2	Identification	-	Ignore	
20+D+S	2	Fragment Info	0x00	Compare	
22+D+S	1	Time to Live	-	Ignore	
23+D+S	1	Protocol	0x29	Compare	IPv6
24+D+S	2	Header Checksum	-	Ignore	
26+D+S	4	Source IP Address	-	Ignore	
30+D+S	4	Destination IP Address	-	Ignore	
34+D+S	N	Possible IP Options		Ignore	
IPv6 Header					
34+D+S+N	1	Version/ Traffic Class	0x6X	Compare	Check IPv6
35+D+S+N	3	Traffic Class/Flow Label	-	Ignore	
38+D+S+N	2	Payload Length	-	Ignore	
40+D+S+N	1	Next Header	IPv6 extension headers	Check	
41+D+S+N	1	Hop Limit	-	Ignore	
42+D+S+N	16	Source Address	-	Ignore	
48+D+S+N	16	Destination Address		Ignore	
74+D+S+N	B	Possible IPv6 Next Headers	-	Ignore	

In this case the packet is split after (74+D+S+N+B) bytes.

- $N = (\text{IP HDR length} - 5) * 4$.
- One of the extension headers of the IPv6 packets must be a “fragment header” in order for the packet to be parsed.

A.2.1.4.4 Type 1.4.2: Ethernet (VLAN/SNAP) Ipv4 Ipv6 UDP

This type contains only Ethernet header, Ipv4 header, IPv6 header and UDP header.

Offset	# of bytes	Field	Value	Action	Comment
0	6	Destination Address		Ignore	MAC Header – processed by main address filter, or broadcast filter.
6	6	Source Address		Ignore	
12	D=(0/4/6/8)	Outer Tag (Outer VLAN or E-tag)	0x8100 **** 0x893F *****	Ignore	
12+D	S=(0/4)	Possible VLAN Tag	0x8100	Check	



Offset	# of bytes	Field	Value	Action	Comment
12+D+S	2	Type	0x0800	Compare	IP
IPv4 Header					
14+D+S	1	Version/ HDR length	0x4X	Compare	Check IPv4 and header length.
15+D+S	1	Type of Service	-	Ignore	
16+D+S	2	Packet Length	-	Ignore	
18+D+S	2	Identification	-	Ignore	
20+D+S	2	Fragment Info	0x00	Compare	
22+D+S	1	Time to Live	-	Ignore	
23+D+S	1	Protocol	0x29	Compare	IPv6
24+D+S	2	Header Checksum	-	Ignore	
26+D+S	4	Source IP Address	-	Ignore	
30+D+S	4	Destination IP Address	-	Ignore	
34+D+S	N	Possible IP Options		Ignore	
IPv6 Header					
34+D+S+N	1	Version/ Traffic Class	0x6X	Compare	Check IPv6
35+D+S+N	3	Traffic Class/Flow Label	-	Ignore	
38+D+S+N	2	Payload Length	-	Ignore	
40+D+S+N	1	Next Header	IPv6 extension header or 0x11	Check	IPv6 extension headers.
41+D+S+N	1	Hop Limit	-	Ignore	
42+D+S+N	16	Source Address	-	Ignore	
58+D+S+N	16	Destination Address		Ignore	
74+D+S+N	B	Possible IPv6 Next Headers	-	Ignore	
UDP Header					
74+D+S+N+B	2	Source Port	Not (0x801)	Check	Not NFS packet.
76+D+S+N+B	2	Destination Port	Not (0x801)	Check	Not NFS packet.
78+D+S+N+B	2	Length	-	Ignore	
80+D+S+N+B	2	Checksum	-	Ignore	

In this case the packet is split after (82+D+S+N+B) bytes.

$$N = (\text{IP HDR length} - 5) * 4.$$

A.2.1.4.5 Type 1.4.3: Ethernet (VLAN/SNAP) Ipv4 Ipv6 TCP

This type contain only Ethernet header, IPv4 header, IPv6 header and TCP header.



Offset	# of bytes	Field	Value	Action	Comment
0	6	Destination Address		Ignore	MAC Header – processed by main address filter, or broadcast filter.
6	6	Source Address		Ignore	
12	D=(0/4/6/8)	Outer Tag (Outer VLAN or E-tag)	0x8100 **** 0x893F *****	Ignore	
12+D	S=(0/4)	Possible VLAN Tag	0x8100	Check	
12+D+S	2	Type	0x0800	Compare	IP
IPv4 Header					
14+D+S	1	Version/ HDR Length	0x4X	Compare	Check IPv4 and header length.
15+D+S	1	Type of Service	-	Ignore	
16+D+S	2	Packet Length	-	Ignore	
18+D+S	2	Identification	-	Ignore	
20+D+S	2	Fragment Info	0x00	Compare	
22+D+S	1	Time to live	-	Ignore	
23+D+S	1	Protocol	0x2A	Compare	IPv6
24+D+S	2	Header Checksum	-	Ignore	
26+D+S	4	Source IP Address	-	Ignore	
30+D+S	4	Destination IP Address	-	Ignore	
34+D+S	N	Possible IP Options		Ignore	
IPv6 Header					
34+D+S+N	1	Version/ Traffic Class	0x6X	Compare	Check IPv6.
35+D+S+N	3	Traffic Class/Flow Label	-	Ignore	
38+D+S+N	2	Payload Length	-	Ignore	
40+D+S+N	1	Next Header	IPv6 extension header Or 0x06	Check	IPv6 extension headers.
41+D+S+N	1	Hop Limit	-	Ignore	
42+D+S+N	16	Source Address	-	Ignore	
58+D+S+N	16	Destination Address		Ignore	
74+D+S+N	B	Possible IPv6 Next Headers	-	Ignore	
TCP Header					
74+T	2	Source Port	Not (0x801)	Check	Not NFS packet.
76+T	2	Destination Port	Not (0x801)	Check	Not NFS packet.
78+T	4	Sequence Number	-	Ignore	
82+T	4	Acknowledge Number	-	Ignore	



Offset	# of bytes	Field	Value	Action	Comment
86+T	1/2	Header Length		Check	
86.5+T	1.5	Different Bits	-	Ignore	
88+T	2	Window Size	-	Ignore	
90+T	2	TCP Checksum	-	Ignore	
92+T	2	Urgent Pointer	-	Ignore	
94+T	F	TCP Options	-	Ignore	

In this case the packet is split after (94+D+S+N+B+F) bytes.

- T = D+S+N+B
- N = (IP HDR length - 5) * 4.
- F = (TCP HDR length - 5)*4

A.2.2 Type 2: Ethernet, IPv6

A.2.2.1 Type 2.1: Ethernet, IPv6 data

This type contains only an Ethernet header and an IPv6 header while the packet should be a fragmented packet. If the packet is not fragmented and the next header is not supported then the header is not split. The supported packet types for header split are programmed in PSRTYPE register (per VF).

Offset	# of bytes	Field	Value (hex)	Action	Comment
0	6	Destination Address		Ignore	MAC Header – processed by main address filter, or broadcast filter.
6	6	Source Address		Ignore	
12	D=(0/4/6/8)	Outer Tag (Outer VLAN or E-tag)	0x8100 **** 0x893F ***** ****	Ignore	
12+D	S=(0/4)	Possible VLAN Tag	0x8100	Check	
IPv6 Header					
12+D+S	2	Type	0x86DD	Compare	IP
14+D+S	1	Version/ Traffic Class	0x6X	Compare	Check IPv6.
15+D+S	3	Traffic Class/Flow Label	-	Ignore	
18+D+S	2	Payload Length	-	Ignore	
20+D+S	1	Next Header	IPv6 next header types.	Check	The last header must be fragmented header in order for the header to be split.
21+D+S	1	Hop Limit	-	Ignore	



Offset	# of bytes	Field	Value (hex)	Action	Comment
22+D+S	16	Source Address	-	Ignore	
38+D+S	16	Destination Address		Ignore	
54+D+S	N	Possible IPv6 Next Headers	-	Ignore	

In this case the packet is split after (54+D+S+N) bytes.

- The last next header field of the IP section field should not be 0x11/0x06 (TCP/UDP).

A.2.2.1.1 Type 2.2: Ethernet (VLAN/SNAP) IPv6 UDP

This type contains only Ethernet header, IPv6 header, and UDP header.

Offset	# of bytes	Field	Value (hex)	Action	Comment
0	6	Destination Address		Ignore	MAC Header – processed by main address filter, or broadcast filter.
6	6	Source Address		Ignore	
12	D=(0/4/6/8)	Outer Tag (Outer VLAN or E-tag)	0x8100 **** 0x893F ***** ***	Ignore	
12+D	S=(0/4)	Possible VLAN Tag		Check	
IPv6 Header					
12+D+S	2	Type	0x86DD	Compare	IP
14+D+S	1	Version/ Traffic Class	0x6X	Compare	Check IPv6.
15+D+S	3	Traffic Class/Flow Label	-	Ignore	
18+D+S	2	Payload Length	-	Ignore	
20+D+S	1	Next Header	IPv6 next header types Or 0x11	Check	
21+D+S	1	Hop Limit	-	Ignore	
22+D+S	16	Source Address	-	Ignore	
38+D+S	16	Destination Address		Ignore	
54+D+S	N	Possible IPv6 Next Headers	-	Ignore	
UDP Header					
54+D+S+N	2	Source Port	Not (0x801)	Check	Not NFS packet.



Offset	# of bytes	Field	Value (hex)	Action	Comment
56+D+S+N	2	Destination Port	Not (0x801)	Check	Not NFS packet.
58+D+S+N	2	Length	-	Ignore	
60+D+S+N	2	Checksum	-	Ignore	


In this case the packet is split after (62+D+S+N) bytes.

- The last *next-header* field of the last header of the IP section must be 0x06.

A.2.2.2 Type 2.3: Ethernet (VLAN/SNAP) IPv6 TCP

This type contains only Ethernet header, IPv6 header, and UDP header.

Offset	# of bytes	Field	Value (hex)	Action	Comment
0	6	Destination Address		Ignore	MAC Header – processed by main address filter, or broadcast filter.
6	6	Source Address		Ignore	
12	D=(0/4/6/8)	Outer Tag (Outer VLAN or E-tag)	0x8100 **** 0x893F ***** ***	Ignore	
12+D	S=(0/4)	Possible VLAN Tag		Check	
IPv6 Header					
12+D+S	2	Type	0x86DD	Compare	IP
14+D+S	1	Version/ Traffic Class	0x6X	Compare	Check IPv6.
15+D+S	3	Traffic Class/Flow Label	-	Ignore	
18+D+S	2	Payload Length	-	Ignore	
20+D+S	1	Next Header	IPv6 next header types Or TCP	Check	
21+D+S	1	Hop Limit	-	Ignore	
22+D+S	16	Source Address	-	Ignore	
38+D+S	16	Destination Address		Ignore	
54+D+S	N	Possible IPv6 Next Headers	-	Ignore	
TCP Header					
54+D+S+N	2	Source Port	Not (0x801)	Check	Not NFS packet.
56+D+S+N	2	Destination Port	Not (0x801)	Check	Not NFS packet.
58+D+S+N	4	Sequence number	-	Ignore	



Offset	# of bytes	Field	Value (hex)	Action	Comment
62+D+S+N	4	Acknowledge Number	-	Ignore	
66+D+S+N	1/2	Header Length		Check	
66.5+D+S+N	1.5	Different Bits	-	Ignore	
68+D+S+N	2	Window Size	-	Ignore	
70+D+S+N	2	TCP Checksum	-	Ignore	
72+D+S+N	2	Urgent Pointer	-	Ignore	
74+D+S+N	F	TCP Options	-	Ignore	

In this case the packet is split after (54+D+S+N+F) bytes.

- $F = (\text{TCP header length} - 5) * 4$.
- The last 'next-header' field of the last header of the IP section must be 0x11.

A.2.3 Type 3: Reserved

Type 3 used to be iSCSI packets (header split is not supported for iSCSI packets).

A.2.4 Type 4: Reserved

A.2.5 Type 5: Cloud Packets

A.2.5.1 Type 5.1: Ethernet, IPv4, NVGRE, IPv4/6, TCP/UDP

This type contains an Ethernet header an IPv4 header, a GRE header an IPv4/6 header and a TCP or UDP header. The supported packet types for header split are programmed in PSRTYPE register (per VF).

Offset	# of bytes	Field	Value (hex)	Action	Comment
0	6	Destination Address		Ignore	MAC Header – processed by main address filter, or broadcast filter.
6	6	Source Address		Ignore	
12	D=(0/4/6/8)	Outer Tag (Outer VLAN or E-tag)	0x8100 **** 0x893F ***** ****	Ignore	
12+D	S=(0/4)	Possible VLAN Tag	0x8100	Check	
12+D+S	2	Type	0x0800	Compare	IP Split on outer L2 is done after this field.
IPv4 Header					
14+D+S	1	Version/ HDR length	0x4X	Compare	Check IPv4 and header length.
15+D+S	1	Type of Service	-	Ignore	



Offset	# of bytes	Field	Value (hex)	Action	Comment
16+D+S	2	Packet Length	-	Ignore	
18+D+S	2	Identification	-	Ignore	
20+D+S	2	Fragment Info	0x00	Compare	
22+D+S	1	Time to Live	-	Ignore	
23+D+S	1	Protocol	0x2F	Compare	GRE
24+D+S	2	Header Checksum	-	Ignore	
26+D+S	4	Source IP Address	-	Ignore	
30+D+S	4	Destination IP Address	-	Ignore	
34+D+S	N	Possible IP Options		Ignore	
GRE Header					
34 + D + S + N	2	Flags/Version	0x2000	Compare	
36 + D + S + N	2	Protocol Type	0x6558	Compare	
38+ D + S + N	3	TNI		Store	Used for flow director.
41 + D + S + N	1	Reserved	0x0	Ignore	Split on cloud header splits after this field.
Inner MAC Header					
42 + D + S + N	6	Inner Destination Address		Ignore	Used for flow director.
48 + D + S + N	6	Inner Source Address		Ignore	
54 + D + S + N	I=(0/4)	Possible Inner VLAN Tag	0x8100	Store	Used for flow director.
54 + D + S + N + I	2	Type	0x0800/ 0x86DD	Compare	Split on L2 is done after this field.
IPv4/6 Header - as previously described.					Split on L3 is done after these fields.
UDP/TCP - as previously described.					Split on L4 is done after these fields.

A.2.5.2 Type 5.2: Ethernet, IPv4, VXLAN, IPv4/6, TCP/UDP

This type contains an Ethernet header an IPv4 header, a UDP header with a specific UDP destination port a VXLAN header an IPv4/6 header and a TCP or UDP header. The supported packet types for header split are programmed in PSRTYPE register (per VF).

Offset	# of bytes	Field	Value (hex)	Action	Comment
0	6	Destination Address		Ignore	MAC Header – processed by main address filter, or broadcast filter.
6	6	Source Address		Ignore	
12	D=(0/4/6/8)	Outer Tag (Outer VLAN or E-tag)	0x8100 **** 0x893F ***** ****	Ignore	
12+D	S=(0/4)	Possible VLAN Tag	0x8100	Check	



Offset	# of bytes	Field	Value (hex)	Action	Comment
12+D+S	2	Type	0x0800	Compare	IP Split on outer L2 is done after this field.
IPv4 Header					
14+D+S	1	Version/ HDR Length	0x4X	Compare	Check IPv4 and header length.
15+D+S	1	Type of Service	-	Ignore	
16+D+S	2	Packet Length	-	Ignore	
18+D+S	2	Identification	-	Ignore	
20+D+S	2	Fragment Info	0x00	Compare	
22+D+S	1	Time to Live	-	Ignore	
23+D+S	1	Protocol	0x11	Compare	UDP
24+D+S	2	Header Checksum	-	Ignore	
26+D+S	4	Source IP Address	-	Ignore	
30+D+S	4	Destination IP Address	-	Ignore	
34+D+S	N	Possible IP Options		Ignore	
UDP Header					
34+D+S+N	2	Source Port		Ignore	
36+D+S+N	2	Destination Port		Compare	Compare to VXLANCN-TRL.UDPPORT.
38+D+S+N	2	Length	-	Ignore	
40+D+S+N	2	Checksum	-	Ignore	
VXLAN Header					
42 + D + S + N	1	Flags	0x08	Compare	VNI exists.
43 + D + S + N	3	Reserved	0x000000	Compare	
46 + D + S + N	3	VNI		Store	Used for flow director.
49 + D + S + N	1	Reserved	0x00	Compare	Split on cloud header splits after this field.
Inner MAC Header					
50 + D + S + N	6	Inner Destination Address		Ignore	Used for Flow director
56 + D + S + N	6	Inner Source Address		Ignore	
62 + D + S + N	I=(0/4)	Possible Inner VLAN Tag	0x8100	Store	Used for Flow director
62 + D + S + N + I	2	Type	0x0800/ 0x86DD	Compare	Split on L2 is done after this field
IPv4/6 Header - as previously described.					Split on L3 is done after these fields
UDP/TCP - as previously described.					Split on L4 is done after these fields



A.2.5.3 optional MACsec encapsulation (expected to be offloaded by hardware) and without Ethernet MAC Addresses

L2 destination and source Ethernet MAC addresses (each of them is six bytes long). The Ethernet MAC address of the target is assumed to be assigned by the network. The mechanism that is used for Ethernet MAC address assignment and Ethernet MAC address detection is outside of the scope of this document.

A.2.6 FC Frame Format

Note: This section is provided as a background on FC and is not required for hardware implementation. For a complete description of the FC fields please refer to FC-FS-2 specification.

The FC frame as defined in FC-FS-2 specification is shown in the Figure A.14 while relevant fields are detailed in this section.

SOF	Extended Header	FC Header	Optional Headers	FC Payload (FC Data & optional padding)	FC CRC	EOF
-----	-----------------	-----------	------------------	---	--------	-----

Figure A.14 FC Frame Format

A.2.6.1 FC SOF and EOF

FC SOF delimiter and EOF delimiter.

A.2.6.2 FC CRC

The Cyclic Redundancy Check (CRC) is a four-byte field that follows the *Data* field. It enables end-to-end integrity checking on the entire FC frame. The FC CRC offload is described in more detail in Section 7.10.

A.2.6.3 FC Optional Headers

Note: Most of the following section is provided as a background on FC and is not required for hardware implementation. The reader can skip the detailed explanation of the optional headers provided and concentrate in the tables and figures that follow the text.

The following table and figures describe the FC frame structure with optional headers and lists the optional headers. The optional headers (that are present) are always ordered as shown in Figure A.17 and Figure A.18. Their presence is indicated in the Data Field Control (DF_CTL) field in the FC header as indicated in Table A.5.

Maximum FC frame size: The sum of the length in bytes of the FC Payload, the number of fill bytes, and the lengths in bytes of all optional headers shall not exceed 2112.

Table A.5 FC Optional Headers

DF_CTL	Optional Header	Length
bit 6	ESP Header / ESP Trailer	Variable



DF_CTL	Optional Header	Length
bit 5	Network Header	16 bytes
bit 4	Association Header	32 byte
bits 1:0	Device Header	0, 16, 32, or 64 bytes




Figure A.17 FC Frame format with Optional Headers (without ESP Header)




Figure A.18 FC Frame format with Optional Headers (with ESP Header)

ESP Header

This is the first optional header that covers the entire FC frame other than the FC header, which is transmitted on the clear (as plain text). When an ESP header is present, there is also the ESP trailer. If required, software is responsible for the cryptic calculation and preparing the ESP header and trailer. Its presence is indicated by bit 6 in the DF_CTL field being set to 1b. Hardware does not support large send offload when the ESP optional header is used. When present, the ESP header and trailer are present in all frames of the exchange.

Network Header

The network header, if used, must be present only in the first data frame of a sequence. A bridge or a gateway node that interfaces to an external network might use the network header. The network header, is an optional header 16 bytes long within the FC data field content. Its presence is indicated by bit 5 in the DF_CTL field being set to 1b. The network header might be used for routing between FC networks of different fabric address spaces, or FC and non-FC networks. The network header contains name identifiers for the network destination address and network source address.

Association Header

The association header, if used, must be present only in the first data frame of a sequence. The association header is an optional header 32 bytes long within the data field content. Its presence is indicated by bit 4 in the DF_CTL field being set to 1b. The association header might be used to identify a specific process or group of processes within a node associated with an exchange. When an Nx_Port has indicated during login that an initial process associator is required to communicate with it, the association header should be used by that Nx_Port to identify a specific process or group of processes within a node associated with an exchange. The Intel® Xeon® Processor D-1500 Product Family LAN controller does not use the association for any filtering purposes but rather uses the OX_ID.

Device Header

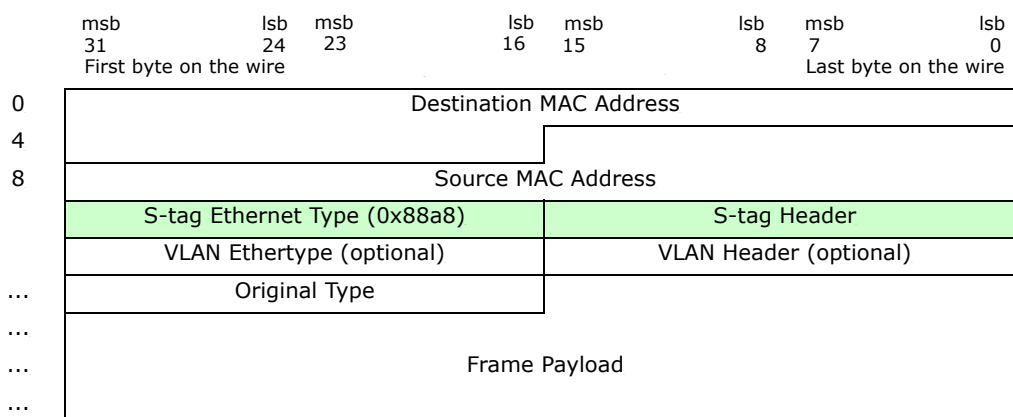


The device header, if present, must be present either in the first data frame or in all data frames of a sequence. If large send off load is used then the device header, if present, is present only in the first frame of the same large send. The device header, if present, must be 16, 32, or 64 bytes in size as defined by bits 1:0 in the DF_CTL field. The contents of the device header are controlled at a level above FC-2. Upper Layer Protocol (ULP) might use a device header, requiring the device header to be supported. The device header might be ignored and skipped, if not needed. If a device header is present for a ULP that does not require it, the related FC-4 might reject the frame with the reason code of TYPE not supported.

A.3 E-tag and S-tag Formats

Note: S-tags are not supported in the Intel® Xeon® Processor D-1500 Product Family LAN controller.

Packets with S-tag has the following format:



The S-tag format is the format defined for S-tags in the IEEE 802.1ad specification described as follows:





Figure A.19 S-tag Format

Packets with E-tag has the following format:





12	E-tag Ethernet Type (0x893F)	E-tag Header
16	E-tag Header (cont.)	
...	VLAN Ethertype (optional)	VLAN Header (optional)
...	Original Type	
...	Frame Payload	

The E-tag format is the format defined in the IEEE 802.1BR specification described as follows:




Figure A.20 E-tag Format

The Ingress_E-CID_ext and E-CID_ext are always zero for endpoints and are effectively reserved.

§



Appendix B Integrated PHY Support

B.1 10 GbE Interface

The SoC provides complete functionality to support the 10 Gb/s ports. The Intel® Xeon® Processor D-1500 Product Family LAN controller performs all functions required for transmission and reception defined in the various standards.

A lower-layer PHY interface is included to attach either to an external PMA or Physical Medium Dependent (PMD) components.

The SoC enables the following connectivity modes:

- IEEE802.3 clause 72 10GBASE-KR
- IEEE Std 802.3 Clause 70 1000BASE-X

B.2 10GBASE-KR Operating Mode

The KR interface supports data rates of 10 Gb/s over copper traces in improved FR4 PCBs. Data is transferred over a single differential path in each direction for a total of two pairs, with each path operating at 10.3125 Gbaud \pm 100 ppm to support overhead of 64B/66B coding. The interface is used to connect the Intel® Xeon® Processor D-1500 Product Family LAN controller to a KR switch port over the backplane.

The high-speed interface is configured as a KR interface while auto-negotiation to a KR link partner is detected. When in 10GBASE-KR operating mode, high-speed interface lane 0 is used for receive and transmit activity while lanes 1 to 3 of the high-speed interface are powered down.

B.2.1 KR Overview

10GBASE-KR definition enables 10 Gb/s operation over a single differential path in each direction for a total of two pairs, or four connections. This system uses the 10GBASE-KR PCS as defined in IEEE802.3 Clause 49 with amendments for auto-negotiation specified in IEEE802.3 Clause 73 and 10 Gigabit PMA as defined in IEEE802.3 clause 51. The 10GBASE-KR PMD is defined in IEEE802.3 Clause 72. The 10GBASE-KR PHY includes 10GBASE-KR Forward Error Correction (FEC), as defined in IEEE802.3 Clause 74. FEC support is optional and is negotiated between Link partners during auto-negotiation. Activating FEC improves link quality (2dB coding gain) by enabling correction of up to 11 bit-burst errors.

KR is a full-duplex interface that uses a single self-clocked serial differential link in each direction to achieve 10 Gb/s data throughput. The serial link transfers scrambled data at 10.3125 Gbaud to accommodate both data and the overhead associated with 64B/66B coding.

Following initialization and auto-negotiation 10GBASE-KR defines a start-up protocol, where link partners exchange continuous fixed length training frames using differential Manchester Encoding (DME) at a signaling rate equal to one quarter of the 10GBASE-KR signaling rate. This protocol

facilitates timing recovery and receive equalization while also providing a mechanism through which the receiver can tune the transmit equalizer to optimize performance over the backplane interconnect. Successful completion of the start-up protocol enables transmission of data between the link partners.

Figure B.1 shows the architectural positioning of 10GBASE-KR.

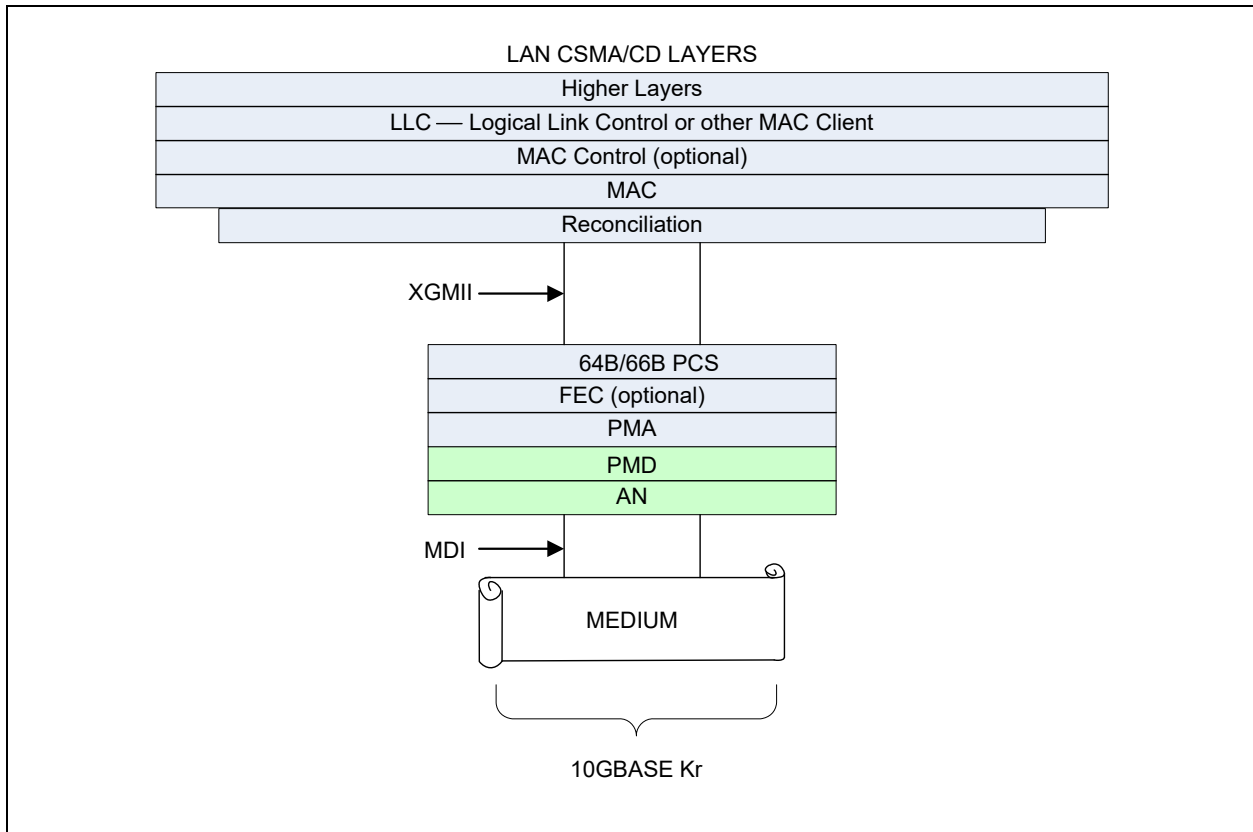


Figure B.1. Architectural Positioning of 10GBASE-KR

B.2.2 KR Electrical Characteristics

The KR lane is a low swing AC coupled differential interface using NRZ signaling. AC coupling allows for inter-operability between components operating from at different supply voltages. Low swing differential signaling provides noise immunity and improved reduced EMI. Differential signal swings defined specifications depend on several factors, such as transmitter pre-equalization and transmission line losses.

The KR signal paths are point-to-point connections. Each path corresponds to a KR lane and is comprised of two complementary signals making a balanced differential pair. There is a single differential path in each direction for a total of two pairs, or four connections.

The 10GBASE-KR link requires a nominal 100 Ω differential source and load terminations with AC coupling on the receive side as shown in Figure B.1.




Figure B.2 KR Backplane Ethernet Channel

B.2.3 KR Reverse Polarity

The KR PHY supports reverse polarity of the KR transmit and receive lanes via the NVM. Contact your Intel representative for more details.

B.3 1 GbE Interface

The Intel® Xeon® Processor D-1500 Product Family LAN controller provides complete support for up to two 1 Gb/s port implementations. The device performs all functions required for transmission and reception defined by the different standards.

A lower-layer PHY interface is included to attach either to external PMA or Physical Medium Dependent (PMD) components.

The SoC enables 1 GbE operation compliant with:

- IEEE802.3 Clause 70 1000BASE-KX

B.3.1 1000BASE-KX Operating Mode

The high-speed interface, when operating as a KX Interface, supports data rates of 1 Gb/s over copper traces on improved FR4 PCBs. Data is transferred over a single differential path in each direction for a total of two pairs, with each path operating at 1.25 Gbaud to support overhead of 8B/10B coding. The interface is used to connect the Intel® Xeon® Processor D-1500 Product Family LAN controller to a KX compliant switch port over the backplane or to KX compliant 1 GbE PHY device. In the event of auto-negotiation defined in IEEE802.3 Clause 73 ending with 1 Gb/s as the HCD, the high-speed interface is configured as a KX interface. KX operating mode can also be forced by software and disabling auto-negotiation (see Section B.4.1).

Note: The KR PHY uses a single KX lane.

B.3.1.1 KX Overview

1000BASE-KX extends the family of 1000BASE-X Physical Layer signaling systems. KX specifies operation at 1 Gb/s over two differential, controlled impedance pairs of traces (one pair for transmit, one pair for receive). This system uses the 1000BASE-X PCS and PMA as defined in IEEE802.3 Clause 36. The 1000BASE-KX PMD is defined in IEEE802.3ap Clause 70.

KX is a full-duplex interface that uses a single serial differential link in each direction to achieve 1 Gb/s data throughput. Each serial link operates at 1.25 GBaud to accommodate both data and the overhead associated with 8B/10B coding. The self-clocked nature eliminates skew concerns between clock and data, and enables a functional reach of up to one meter.

Figure B.3 shows the architecture positioning of 1000BASE-KX.

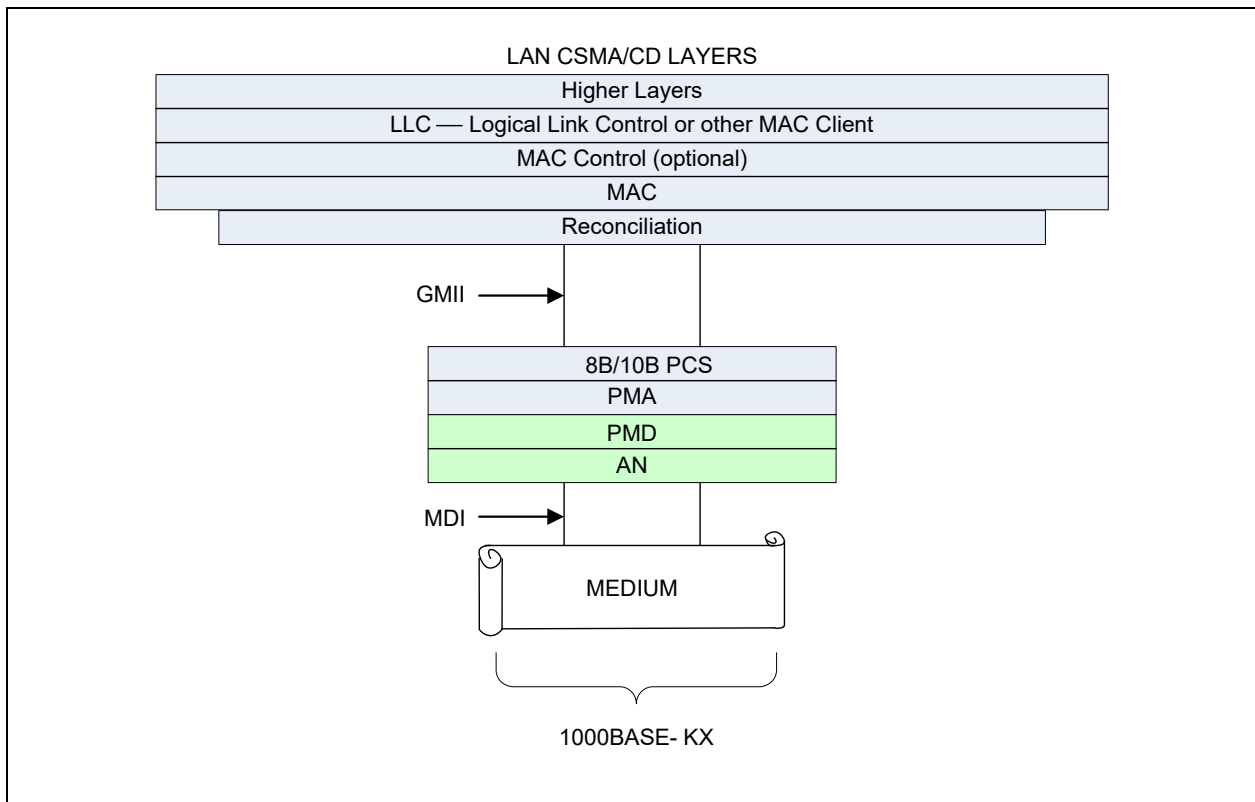


Figure B.3 Architectural Positioning of 1000BASE-KX

B.3.1.2 KX Electrical Characteristics

The KX lane is a low swing AC coupled differential interface using NRZ signaling. AC coupling allows for inter-operability between components operating from at different supply voltages. Low swing differential signaling provides noise immunity and improved reduced electromagnetic interference (EMI). Differential signal swings defined specifications depend on several factors, such as transmitter pre-equalization and transmission line losses.



The KX signal paths are point-to-point connections. Each path corresponds to a KX lane and is comprised of two complementary signals making a balanced differential pair. There is one differential path in each direction for a total of two pairs, or four connections.

B.4 Auto Negotiation For Backplane Ethernet and Link Setup Features

Auto-negotiation provides a linked device with the capability to detect the abilities (modes of operation) supported by the device at the other end of the link, determine common abilities, and configure for joint operation.

Auto-negotiation for backplane Ethernet is defined in IEEE802.3 Clause 73 and is based on IEEE802.3 clause 28 definition of auto-negotiation for twisted-pair link segments.

Auto-negotiation for backplane Ethernet uses an extended base page and next page format and modifies the timers to allow rapid convergence. Furthermore, auto-negotiation does not use Fast Link Pulses (FLPs) for link code word signaling and instead uses Differential Manchester Encoding (DME) signaling, which is more suitable for electrical backplanes. Since DME provides a DC balanced signal. Auto-negotiation for backplane Ethernet also includes support for parallel detection of 1000BASE-KX links in addition to transmission and reception of extended base page and next page auto-negotiation frames.

B.4.1 MAC Link Setup and Auto Negotiation

The MAC block in the Intel® Xeon® Processor D-1500 Product Family LAN controller supports both 10 GbE and 1 GbE link modes and the appropriate functionality specified in the standards for these link modes.

Each of these link modes can use different PMD sub-layer and base band medium types.

In 10 GbE operating mode, the Intel® Xeon® Processor D-1500 Product Family LAN controller supports 10GBASE-KR. While in 1 GbE operating mode, the Intel® Xeon® Processor D-1500 Product Family LAN controller supports 1000Base-KX protocols. The different protocols supported in 10 GbE operating mode and 1 GbE operating mode affect only the configuration of the high-speed AFE and the high-speed interface PHY logic blocks (PCS, FEC, etc.) while the MAC supports rates of either 1 Gb/s or 10 Gb/s, without the need to know the electrical medium actually being interfaced.

Link speed and link characteristics can be determined through static configuration, parallel detect and auto-negotiation or forced operation for diagnostic purposes. The auto-negotiation processes defined in IEEE802.3 Clause 73 enables selection between KR (10 GbE) and KX (1 GbE) compliant link partners and defining link characteristics and link speed.

Link setting is done by configuring the speed configuration, defining the appropriate physical interface and restarting auto-negotiation.


B.5 KR Link Configuration Flows

Note: All register access to the KR link configuration flows are done using the operation described in [Section 3.8.1](#).

B.5.1 Link Setup Flow - KR/KX

This flow describes the operations required for setting the link to a 10 GbE/1GbE backplane.

1. Configure the following settings:



- KRM_KR_PCS_PORT<n>.LINK_CNTL_1[14] FEC requested (F1 bit in the Backplane Auto-negotiation Base Link Code Word)
 - 1b = Request FEC support
 - 0b = Don't request FEC support
 - KRM_KR_PCS_PORT<n>.LINK_CNTL_1[15] FEC ability (F0 bit in the Backplane Auto-negotiation base Link Code Word)
 - 1b = Advertise FEC support
 - 0b = Don't advertise FEC support
 - KRM_KR_PCS_PORT<n>.LINK_CNTL_1[16] Backplane Auto-negotiation Capability bits: KX
 - 1b = Advertise KX support
 - 0b = Don't advertise KX support
 - KRM_KR_PCS_PORT<n>.LINK_CNTL_1[18] Backplane Auto-negotiation Capability bits: KR
 - 1b = Advertise KR support
 - 0b = Don't advertise KR support
 - KRM_KR_PCS_PORT<n>.LINK_CNTL_1[24] Backplane Auto-negotiation EEE bits: KX
 - 1b = Advertise KX EEE support
 - 0b = Don't advertise KX EEE support
 - KRM_KR_PCS_PORT<n>.LINK_CNTL_1[26] Backplane Auto-negotiation EEE bits: KR
 - 1b = Advertise KR EEE support
 - 0b = Don't advertise KR EEE support
 - KRM_KR_PCS_PORT<n>.LINK_CNTL_1[29] Backplane Auto-negotiation Enable = 1b (MUST)
2. Restart Backplane Auto-negotiation:
- KRM_KR_PCS_PORT<n>.LINK_CNTL_1[31] Restart AN = 1b

B.6 Low Power Link Up (LPLU)

Normal PHY speed negotiation drives to establish a link at the highest possible speed. The PHY supports an additional mode of operation, where the PHY drives to establish a link at a low speed. The LPLU process enables a link to come up at the lowest possible speed in cases where power is more important than performance. Different behavior is defined for the D0 state and the other non-D0 states.

B.7 Behavior in Non-D0 State

If the *LPLU* bit is set in the NVM, the PHY negotiates to a low speed while in non-D0 states (Dr or D3). This applies only when the link is required by one of the following: SMBus or NC-SI manageability, APM wake, or PME. Otherwise, the PHY is disabled during the non-D0 state.

Link negotiation begins with the PHY trying to negotiate at the lowest speed it is allowed to advertise. If link establishment fails, the PHY tries to negotiate at additional speeds, such as all speeds allowed up to the lowest speed supported by the partner. For example, the PHY advertises 100 Mb/s only and the partner supports 1000 Mb/s only. After the first try fails, PHY enables 100/1000 Mb/s and tries again. The PHY continues to try and establish a link until it succeeds or until it is instructed otherwise.



B.8 PHY Loopback

The following types of loopback in the LAN interfaces might be configured:

- Internal PHY Loopbacks (Point 2 in [Figure B.4](#))
- External Loopback (Point 3 in [Figure B.4](#))

By setting the device to loopback mode, packets that are transmitted towards the line are looped back to the host. The device is fully functional in these modes, just not transmitting data over the lines if the loopback is internal. [Figure B.4](#) shows the points of loopback.

For more details on the usage and loopback test setup, refer to the *Intel® Ethernet Controllers Loopback Modes* application note.




Figure B.4. Magnolia Park Loopback Modes

B.8.1 MAC Loopbacks

See [Section 11.2.1](#).

B.8.2 Internal PHY Loopback (Kerem)

1. Disable AN and force speed to 10G serial
 - `KRM_KR_PCS_PORT0.LINK_CNTL_1.cfg_teth_an_enable = 1'b0`
 - `KRM_KR_PCS_PORT0.LINK_CNTL_1.cfg_teth_force_speed = 3'b100`
2. Set near-end loopback clocks:
 - `KRM_PORT_CAR_PORT0.PORT_CAR_GEN_CNT.port_dfx_32b_nelb_sel = 1'b1`
 - `KRM_PORT_CAR_PORT0.PORT_CAR_GEN_CNT.port_dfx_krpcs_nelb_sel = 1'b1`
3. Set loopback enable (select one)
 - Force tx->rx lpbk: `KRM_PMD_DFX_PORT0.burninctrl.tx_rx_kr_lpbk_en_reg = 1'b1`
 - Force rx->tx lpbk: `KRM_PMD_DFX_PORT0.burninctrl.dig_rx_tx_lpbk_en_reg = 1'b1`
4. Training Bypass
 - `KRM_KR_PMD_RX_TRN_PORT0.PMD_RX_TRN_linkup_ctrl.training_enable = 1'b0`



B.8.3 Internal Loopback (UniPHY)

The UniPHY provides two modes of internal PHY loopback - Digital Near End Loopback (DNELB) mode and Analog Near End Loopback (ANELB) mode:

- The DNELB Tx → Rx loopback path is in the “shim” fub between the PCS and the UniPhy. This loopback path is in the 10b domain before it has been serialized in the UniPhy, so it effectively bypass the entire analog portion of the IO implementation.
- The ANELB Tx → Rx loopback path is almost at the IO pins – implemented as a mux return path to the CTLE. This loopback path effectively includes the full UniPhy implementation – digital and analog portions of the IO design.

§