

Intel® EP80579 Software for Security Applications on Intel® QuickAssist Technology

Package Version 1.0.2

Release Notes

June 2010



INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

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

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

Any software source code reprinted in this document is furnished under a software license and may only be used or copied in accordance with the terms of that license.


Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See http://www.intel.com/products/processor_number for details.

Code Names are only for use by Intel to identify products, platforms, programs, services, etc. ("products") in development by Intel that have not been made commercially available to the public, i.e., announced, launched or shipped. They are never to be used as "commercial" names for products. Also, they are not intended to function as trademarks.

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

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

Copyright © 2010, Intel Corporation. All rights reserved.



Contents

1	Description of Release	5
1.1	New Features	5
1.2	Supported Operating Systems	5
1.3	Supported Component Versions	5
1.3.1	Version Numbering Scheme	5
1.3.2	Package Versions	6
1.3.3	BIOS/Firmware Version	6
2	Known Issues - Silicon	6
3	Known Issues - Embedded	7
4	Resolved Issues - Embedded	12
5	Known Issues - Security	16
6	Resolved Issues - Security	22
7	Related Documentation	31
7.1	Where to Find Current Software and Documentation	31
7.2	Embedded Documents	32
7.3	Security Documents	32

Tables

1	Intel® EP80579 Integrated Processor Product Line Errata	6
2	Summary of Embedded Software Open Issues	7
3	Summary of Embedded Software Resolved Issues	12
4	Summary of Security Software Open Issues	17
5	Summary of Security Software Resolved Issues	22
6	Embedded Documents	32
7	Security Documents	32



Revision History

Date	Revision	Description
04 June 2010	002	Added "IXA00161211 - Security Vulnerability in Gigabit Ethernet Driver due to Ethernet frames that exceed MTU" on page 9
12 November 2009	001	This document is derived from document number 320181-011 prepared for the 1.0.2 release. Differences between the 1.0.2 version and this version are marked with changebars.

§ §



1 Description of Release

This document describes extensions to and deviations from the release functionality described in Intel® EP80579 Software for Security Applications on Intel® QuickAssist Technology Programmer's Guide.

For instructions on loading and running the release, see the Getting Started Guide for your operating system. See [Section 7, "Related Documentation" on page 31](#) for details.

Note: Prior to installing the EP80579 software package, uninstall the previous installation of the software. Refer to the Getting Started Guide for instructions on uninstalling the software.

These release notes may also include known issues with third-party or reference platform components that affect the operation of the software.

Note: The "Intel® EP80579 Integrated Processor with Intel® QuickAssist Technology Development Board" is referred to as "development board" throughout this document.

1.1 New Features

- EP80579 embedded software drivers version 1.0.3
- Linux* package includes CompactFlash driver
- BIOS version update
- Bug fixes related to open source software (see [Section 5.6](#) and [Section 6.16](#))

See the Getting Started Guide for your operating system for detailed information about the features of this release.

1.2 Supported Operating Systems

This software release has been validated with the following operating systems:

- CentOS 5.2 Linux*, Kernel 2.6.18
- FreeBSD* v7.1

1.3 Supported Component Versions

1.3.1 Version Numbering Scheme

The version numbering scheme used in this software release follows this naming convention:

package.os.major.minor.maintenance-build

where:

- package can be one of the following:
 - Embedded
 - Security
 - Telephony
- os can be one of the following:
 - L = Linux*
 - B = FreeBSD*



1.3.2 Package Versions

Operating System	Package Version
Linux*	Security.L.1.0.2-213
FreeBSD*	Security.B.1.0.2-174

1.3.3 BIOS/Firmware Version

The term BIOS is used to refer to pre-boot firmware which could include legacy BIOS or Extensible Firmware Interface (EFI) compliant firmware.

BIOS Version: TRXTG064.ROM

2 Known Issues - Silicon

The Intel® EP80579 Integrated Processor Product Line Specification Update describes known silicon defects. Defects that are specific to Intel® EP80579 silicon are listed in the section called “Intel® EP80579 Integrated Processor Product Line Errata”.

Table 1 lists defects identified with the Intel® EP80579 silicon that have software workarounds. If the workaround is implemented in the released software, it is indicated by **X** in the appropriate Operating System column in Table 1. For workaround details, see the Intel® EP80579 Integrated Processor Product Line Specification Update.

Table 1. Intel® EP80579 Integrated Processor Product Line Errata

No. ¹	Errata	Operating System		
		Linux*	FreeBSD*	Windows*
3	Gigabit Ethernet MAC Receive Timer interrupt problems	X	X	X
4	Gigabit Ethernet MAC Large Segment Offload (LSO) premature descriptor write back	X	N/A ²	X
5	Gigabit Ethernet MAC XOFF from link partner can pause flow-control (XON/XOFF) transmission ³	N/A	N/A	N/A
6	Gigabit Ethernet MAC transmit descriptor use of Report Status (RS) bit for non-data (Context & Null) descriptors	X	X	X
8	Gigabit Ethernet MAC legacy transmit descriptor write-back may occur before the packet data associated with the descriptor is fetched ⁴		N/A ⁵	
9	Gigabit Ethernet MAC may have EEPROM deadlock when using manual software EEPROM access	X	X	X
Notes: 1. This is the same defect number assigned in the Intel® EP80579 Integrated Processor Product Line Specification Update. 2. LSO is not supported by the FreeBSD driver. 3. This issue depends on the link partner flow control settings. The Flow Control thresholds have been tested extensively with no observed issues. 4. The software workaround defined in the Specification Update is not currently implemented in the Linux or Windows Ethernet driver. 5. TCP Segmentation Offload is not supported by the FreeBSD driver.				



3 Known Issues - Embedded

For supplementary information relating to the Known Issues, please refer to the following document:

- Intel® EP80579 Software Drivers for Embedded Applications Programmer's Guide and API Reference Manual, Number: 320154

Note: If the Affected OS field in an errata table lists Red Hat Enterprise Linux 5.0, readers can assume the errata is also present under CentOS v5.2 Linux which is supported in this release.

Table 2. Summary of Embedded Software Open Issues

IXA00058263 - SATA port 1 not showing populated when CD/DVD ROM attached.....	7
IXA00160881 - Using the smbmsg utility to probe the SMBus may hang the Intel® EP80579 Development Board on FreeBSD	8
IXA00161154 - Booting with EFI from Compact Flash on Intel® EP80579 Development Board fails with 1 GB DIMM installed	8
IXA00161211 - Security Vulnerability in Gigabit Ethernet Driver due to Ethernet frames that exceed MTU .	9
IXA00179772 - IDE mode within the BIOS setup menu should be set to "AHCI" for optimal performance on Red Hat Enterprise Linux 5.0.....	11
IXA00216017 - CompactFlash cards are not supported with embedded software drivers.....	11
IXA00241849 - Installation issues observed on Red Hat Enterprise Linux 5 using DVD SATA drives with IDE mode set to AHCI	11
IXA00309267 - UART port may not come out of S3 hibernation	12

3.1 IXA00058263 - SATA port 1 not showing populated when CD/DVD ROM attached

Title	SATA port 1 not showing populated when CD/DVD ROM attached
Reference #	IXA00058263
Description	In the BIOS Setup Menu under the IDE selection screen the 2nd SATA port (SATA 1) shows "Not Present" when a CD/DVD ROM is plugged into SATA port 1 and the IDE mode is AHCI. Hard Drives will appear correctly. However, the CD/DVD ROM appears to be fully functional and can be selected in the boot order.
Implication	In the AHCI mode for IDE, CD/DVD ROMs may not show up when plugged into the SATA ports. This is viewed in the Advanced->IDE Configuration Screen in the BIOS setup menu. Functionality of the device is not impacted and will still be selectable in the boot order. The CD/DVD ROM devices tested were Plextor PX-72Sa and LG GSA-H62L.
Resolution	No work around available.
Affected OS	FreeBSD 6.2 Red Hat Enterprise Linux 5.0
Driver/Module	Driver-General



3.2 IXA00160881 - Using the smbmsg utility to probe the SMBus may hang the Intel® EP80579 Development Board on FreeBSD

Title	Using the smbmsg utility to probe the SMBus may hang the Intel® EP80579 Development Board on FreeBSD
Reference #	IXA00160881
Description	<p>If the FreeBSD smbmsg utility is used to probe the SMBus, the system may hang. Note that probing the SMBus is risky. Individual devices can perform unwanted actions upon receiving the probe request message. For example, if a particular SMBus device considers any write operation issued to it as a request to power off the system, the probing would trigger this action.</p> <p>For additional information please refer to the smbmsg man page (http://www.ipnom.com/FreeBSD-Man-Pages/smbmsg.8.html)</p>
Implication	Using the smbmsg utility to probe the SMBus may hang the development board on FreeBSD.
Resolution	It is not advisable to use the smbmsg utility to probe SMBus.
Affected OS	FreeBSD 6.2
Driver/Module	SMBus Controller Driver

3.3 IXA00161154 - Booting with EFI from Compact Flash on Intel® EP80579 Development Board fails with 1 GB DIMM installed

Title	Booting with EFI from Compact Flash on Intel® EP80579 Development Board fails with 1 GB DIMM installed
Reference #	IXA00161154
Description	Booting with EFI from CompactFlash on Intel® EP80579 Development Board fails with 1 GB DIMM installed. This problem is not seen when 512MB DIMM is installed or when legacy boot is used with 1 GB DIMM is installed.
Implication	It is not possible to perform EFI boot from Compact Flash on Intel® EP80579 Development Board with 1 GB DIMM installed.
Resolution	Booting from Compact Flash can be done in legacy mode with 1 GB DIMM installed or EFI Boot with 512 MB DIMM installed.
Affected OS	CentOS 5.2 Linux
Driver/Module	Intel® EP80579 Development Board BIOS




3.4

IXA00161211 - Security Vulnerability in Gigabit Ethernet Driver due to Ethernet frames that exceed MTU

Title	Security Vulnerability in Gigabit Ethernet Driver due to Ethernet frames that exceed MTU
Reference #	IXA00161211
Description	A security vulnerability exists in the Gigabit Ethernet Driver. The driver allows remote DOS attack through careful selection of frame size in relation to interface MTU, which causes a denial of service (panic), via a crafted frame size.
Implication	<p>Since the Intel® EP80579 Linux Gigabit Ethernet driver does not support receiving packets that span multiple Rx buffers, it checks the End of Packet bit of every frame, and discards it if it is not set. This creates a situation where the first part of a spanning packet is discarded, but the second part is not (since it is the end of packet and it passes the EOP bit test).</p> <p>If the second part of the frame is small (4 bytes or less), the driver subtracts 4 from it to remove its CRC, underflow the length, and winds up in <code>skb_over_panic</code>, when the driver tries to <code>skb_put</code> a huge number of bytes into the <code>skb</code>. This allows a remote DOS attack through careful selection of frame size in relation to interface MTU, which causes a denial of service (panic), via a crafted frame size.</p> <p>Additional information on this defect is available at: http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2009-4536</p>



Title	Security Vulnerability in Gigabit Ethernet Driver due to Ethernet frames that exceed MTU (Continued)
Resolution	<p>The following updates should be applied to the driver to address the security vulnerability.</p> <p>Corrective Action #1:</p> <p>Embedded/src/GbE/iegbe_main.c:</p> <p>In the function iegbe_clean_rx_irq() find the conditional code fragment:</p> <pre> if(unlikely(!(rx_desc->status & E1000_RXD_STAT_EOP))) { /* All receives must fit into a single buffer */ E1000_DBG("%s: Receive packet consumed multiple" " buffers\n", netdev->name); dev_kfree_skb_irq(skb); goto next_desc; } </pre> <p>Replace it with:</p> <pre> /* !EOP means multiple descriptors were used to store a * single packet, if that is the case we need to toss it. * In fact, we need to toss every packet with the EOP bit * clear and the next frame that _does_ have the EOP bit set, * as it is by definition only a frame fragment */ if (unlikely(!(rx_desc->status & E1000_RXD_STAT_EOP))) adapter->discarding = TRUE; if (adapter->discarding) { /* All receives must fit into a single buffer */ E1000_DBG("%s: Receive packet consumed multiple" " buffers\n", netdev->name); if (rx_desc->status & E1000_RXD_STAT_EOP) adapter->discarding = FALSE; dev_kfree_skb_irq(skb); goto next_desc; } </pre> <p>Corrective Action #2:</p> <p>Embedded/src/GbE/iegbe.h:</p> <p>In the iegbe_adapter structure, add the discarding element to the end of the structure.</p> <p>Ensure that the element is not added within any of the conditional preprocessor statements:</p> <pre> struct iegbe_adapter { struct timer_list tx_fifo_stall_timer; struct timer_list watchdog_timer; . . . uint32_t pci_state[16]; int msg_enable; #ifdef CONFIG_PCI_MSI boolean_t have_msi; #endif boolean_t discarding; } </pre>
Affected OS	Red Hat Enterprise Linux 5.0 CentOS 5.2 Linux
Driver/Module	Gigabit Ethernet Controller Driver



3.5 IXA00179772 - IDE mode within the BIOS setup menu should be set to "AHCI" for optimal performance on Red Hat Enterprise Linux 5.0

Title	IDE mode within the BIOS setup menu should be set to "AHCI" for optimal performance on Red Hat Enterprise Linux 5.0
Reference #	IXA00179772
Description	The performance of the Embedded Gigabit driver is significantly lower when the IDE mode is set to "Legacy". AHCI inherently provides higher performance and setting the system to AHCI mode will result in superior performance results, provided the hard drive used in the system supports AHCI.
Implication	System performance is significantly lower when IDE mode is set to "Legacy" mode.
Resolution	AHCI inherently provides higher performance. Setting the system IDE mode to "AHCI" will result in superior performance results. Refer to the chapter titled Pre-boot (BIOS) Firmware within the Getting Started Guide for instructions to toggle the IDE mode to "AHCI" in the BIOS setup menu.
Affected OS	Red Hat Enterprise Linux 5.0
Driver/Module	Gigabit Ethernet Controller Driver

3.6 IXA00216017 - CompactFlash cards are not supported with embedded software drivers

Title	CompactFlash cards are not supported with embedded software drivers
Reference #	IXA00216017
Description	Currently CompactFlash on the Lower Expansion Bus of the Intel EP80579 Development Board are not supported with drivers in any of the embedded software releases.
Implication	No CompactFlash driver support is available for use of the CompactFlash.
Resolution	As of Embedded Release 1.0.3, CompactFlash is supported under Linux only. No work around available for FreeBSD or Windows XP Embedded.
Affected OS	FreeBSD 6.2 Red Hat Enterprise Linux 5.0
Driver/Module	Not Applicable

3.7 IXA00241849 - Installation issues observed on Red Hat Enterprise Linux 5 using DVD SATA drives with IDE mode set to AHCI

Title	Installation issues observed on Red Hat Enterprise Linux 5 using DVD SATA drives with IDE mode set to AHCI
Reference #	IXA00241849
Description	When IDE mode is set to AHCI, installation of Red Hat Enterprise Linux using some DVD SATA drives will fail. During the installation process the message "Loading AHCI driver" is displayed, and the installation hangs. This has been observed on some DVD SATA drives.
Implication	When IDE mode is set to AHCI, installation of Red Hat Enterprise Linux using some DVD SATA drives will fail.
Resolution	If DVD SATA drives exhibit this behavior, perform installation with IDE mode set to Legacy. If AHCI is desired, perform installation using USB DVD, USB CD drive, or other DVD SATA drive.
Affected OS	Red Hat Enterprise Linux 5.0
Driver/Module	SATA



3.8 IXA00309267 - UART port may not come out of S3 hibernation

Title	UART port may not come out of S3 hibernation
Reference #	IXA00309267
Description	When the development board is suspended via the command 'echo -n mem > /sys/power/state' and brought back into service via pressing the power button, the UART may not come back into service. Crash trace is reported in /var/log/messages relating to IRQ 9 which is associated with ACPI as reported /proc/interrupts
Implication	Affects Embedded release packages only because power management feature is not available on accelerated software.
Resolution	To recover, reboot the system.
Affected OS	Red Hat Enterprise Linux 5.0
Driver/Module	Intel® EP80579 Development Board BIOS

4 Resolved Issues - Embedded

Issues that have been resolved in previous versions of the software are included here.

Table 3. Summary of Embedded Software Resolved Issues

IXA00058114 - Embedded Gigabit Ethernet driver Load/Unload memory leak on Red Hat Enterprise Linux v5.0.....	12
IXA00058233 - Intel® EP80579 Integrated Processor with Intel® QuickAssist Technology Development Board will hang on reset command in FreeBSD	13
IXA00058236 - Gigabit Ethernet devices do not appear after reboot on FreeBSD.....	13
IXA00160925 - Manual setting of full-duplex mode of Embedded Gigabit Ethernet driver defaults to half-duplex on Microsoft Windows XP Embedded & Red Hat Enterprise Linux 5.0.....	13
IXA00160970 - Make install targets do not work when executed within component directory on FreeBSD ..	14
IXA00161027 - Autonegotiate fails to switch modes in Gigabit Ethernet driver	14
IXA00206755 - Driver displays install warnings with Red Hat Enterprise Linux 5.0 distribution - 2.6.18 kernel	15
IXA00343773 - Software lockup may occur in Embedded Gigabit Ethernet driver.....	16

4.1 IXA00058114 - Embedded Gigabit Ethernet driver Load/Unload memory leak on Red Hat Enterprise Linux v5.0

Title	Embedded Gigabit Ethernet driver Load/Unload memory leak on Red Hat Enterprise Linux v5.0
Reference #	IXA00058114
Description	A memory leak has been found while loading/unloading the Gigabit Ethernet driver over 15000 times in a 20 hour period without a reboot of the development board.
Implication	Although any memory leak is a concern, this is an extreme situation that should not impact any known usage model with the embedded software Linux release.
Resolution	This issue was determined to be an Operating System feature. See IXA00320219 for details
Affected OS	Red Hat Enterprise Linux 5.0
Driver/Module	Gigabit Ethernet Controller Driver



4.2 IXA00058233 - Intel® EP80579 Integrated Processor with Intel® QuickAssist Technology Development Board will hang on reset command in FreeBSD

Title	Intel® EP80579 Integrated Processor with Intel® QuickAssist Technology Development Board will hang on reset command in FreeBSD
Reference #	IXA00058233
Description	Commands to reboot the system under FreeBSD result in a system hang on the Development Board. This has been observed with 'reboot' and 'shutdown -r now' commands. The operating system completes the shutdown process, but hangs in the reset mechanism. There have been several reports of this in the public forums and is apparently caused by a defect in the operating system.
Implication	The board must be manually reset.
Resolution	This issue was resolved in FreeBSD 6.3
Affected OS	FreeBSD 6.2
Driver/Module	Not Applicable

4.3 IXA00058236 - Gigabit Ethernet devices do not appear after reboot on FreeBSD

Title	Gigabit Ethernet devices do not appear after reboot on FreeBSD
Reference #	IXA00058236
Description	When there is a dependency between drivers, FreeBSD detects the dependency when loading the drivers manually using kldload command and loads the drivers accordingly. However, when the task is automated to load dependent drivers after each boot, the OS seems to load the drivers in ascending order of PCI bus:device:function:number, as it enumerates devices. The GbE adapters and GCU are on the same PCI bus, but GCU has a function numerically larger than the adapters, hence its driver gets loaded only after the GbE driver loads (which fails to initialize the hardware because it is dependent on GCU).
Implication	After rebooting FreeBSD, Gigabit Ethernet devices do not appear.
Resolution	This issue was resolved in Embedded.B.1.0.1
Affected OS	FreeBSD 6.2
Driver/Module	Gigabit Ethernet Controller Driver

4.4 IXA00160925 - Manual setting of full-duplex mode of Embedded Gigabit Ethernet driver defaults to half-duplex on Microsoft Windows XP Embedded & Red Hat Enterprise Linux 5.0

Title	Manual setting of full-duplex mode of Embedded Gigabit Ethernet driver defaults to half-duplex on Microsoft Windows XP Embedded & Red Hat Enterprise Linux 5.0
Reference #	IXA00160925
Description	When Embedded Gigabit Ethernet driver is manually configured to run at 10Mbps or 100Mbps full duplex, the operation mode defaults to half-duplex. This behavior has been observed under Microsoft Windows XP Embedded and Red Hat Enterprise Linux 5.0. Autonegotiation allows driver to run at 10Mbps and 100Mbps full-duplex.
Implication	It is not possible to manually set the duplex mode to full-duplex with the Embedded Gigabit Ethernet driver. Attempts to do this will result in duplex mode being set to half-duplex. The system will still function, but will do so in half-duplex mode.



Title	Manual setting of full-duplex mode of Embedded Gigabit Ethernet driver defaults to half-duplex on Microsoft Windows XP Embedded & Red Hat Enterprise Linux 5.0 (Continued)
Resolution	<p>The errata text description and implication incorrectly states that it is not possible to manually set duplex mode and speed with the Embedded Gigabit Ethernet driver. Earlier testing included manually setting duplex mode and speed on the Embedded Gigabit Ethernet driver and then connecting to a link partner with autonegotiation enabled. When a link partner is connected to a device that is not using autonegotiation, the autonegotiation process fails. The autonegotiation end of the connection is still able to correctly detect the speed of the other end, but cannot detect the duplex mode. Industry standard requires use of half-duplex mode in these conditions. Because the link partner was placed in half-duplex mode, the Embedded Gigabit Ethernet driver was incorrectly identified as the source of the issue and errata was created to document the behavior.</p> <p>There is no issue with manual setting of duplex mode and speed with the Embedded Gigabit Ethernet driver.</p>
Affected OS	Red Hat Enterprise Linux 5.0
Driver/Module	Gigabit Ethernet Controller Driver

4.5 IXA00160970 - Make install targets do not work when executed within component directory on FreeBSD

Title	Make install targets do not work when executed within component directory on FreeBSD
Reference #	IXA00160970
Description	<p>The Makefiles for 1588, CAN, EDMA, GCU, GPIO, and WDT each contain a "install" target. There is an issue with the "kldload \$(KMOD).ko" line. Executing "make install" when in component directory (such as 1588) the command fails with an error message similar to following:</p> <p>"can't load timesync.ko: No such file or directory"</p> <p>This error message is returned, because in FreeBSD the current directory is not part of the path for security reasons, so the kldload command does not know where the .ko image comes from.</p>
Implication	Individual embedded drivers will not be installed.
Resolution	This issue was resolved in Embedded.B.1.0.1
Affected OS	FreeBSD 6.2
Driver/Module	All

4.6 IXA00161027 - Autonegotiate fails to switch modes in Gigabit Ethernet driver

Title	Autonegotiate fails to switch modes in Gigabit Ethernet driver
Reference #	IXA00161027
Description	When an ethernet interface on the EP80579 is set to autonegotiate and the link partner changes speed, the interface on the EP80579 does not autonegotiate.
Implication	Ifconfig reports the previous speed and there is no network connectivity on that interface.
Resolution	This issue was resolved in Embedded.B.1.0.2
Affected OS	FreeBSD 6.2 FreeBSD 6.3
Driver/Module	Gigabit Ethernet Controller Driver



4.7

IXA00206755 - Driver displays install warnings with Red Hat Enterprise Linux 5.0 distribution - 2.6.18 kernel

Title	Driver displays install warnings with Red Hat Enterprise Linux 5.0 distribution - 2.6.18 kernel
Reference #	IXA00206755
Description	When building the Intel® EP80579 Software Drivers for Embedded Applications using Red Hat Enterprise Linux 5.0, warnings are displayed as follows: Warning: vmlinux - Section mismatch: reference to .exit.text: from .smp_alternatives between '__smp_alt_begin' (at offset...) and '__smp_locks_end'
Implication	These warning messages are from the Red Hat Enterprise Linux 5.0 distribution. They are not from the embedded software drivers. These warnings will not interfere with the building and function of the embedded software drivers
Resolution	Operating System warning messages are not displayed because Red Hat Enterprise Linux is not supported.
Affected OS	Red Hat Enterprise Linux 5.0
Driver/Module	All




4.8 IXA00343773 - Software lockup may occur in Embedded Gigabit Ethernet driver

Title	Software lockup may occur in Embedded Gigabit Ethernet driver
Reference #	IXA00343773
Description	<p>The Embedded Gigabit Ethernet driver uses spin_lock_bh() to protect the MDIO registers. This can cause a soft-lockup issue.</p> <p>The lockup would look similar to:</p> <p>BUG: soft lockup - CPU#0 stuck for 10s! [netstat:4281]</p> <p>Pid: 4281, comm: netstat</p> <p>EIP: 0060: [<c0609590>] CPU: 0</p> <p>EIP is at _spin_lock+0x7/0xf</p> <p>EFLAGS: 00000286 Tainted: PF (2.6.18-GA102_OCF #1)</p> <p>EAX: f48c2f00 EBX: 00000002 ECX: c072bf8e EDX: f152be40</p> <p>ESI: 00000011 EDI: c072bf8e EBP: e604a7c0 DS: 007b ES: 007b</p> <p>CRO: 8005003b CR2: b7fa6000 CR3: 2582a000 CR4: 000006d0</p> <p>[<f48c10b6>] gcu_get_adapter+0x16/0x20 [gcu]</p> <p>[<f48c14e3>] gcu_read_eth_phy+0x23/0x130 [gcu]</p> <p>[<f494c6c5>] iegbe_oem_phy_is_link_up+0x55/0x90 [iegbe]</p> <p>[<f49432a9>] iegbe_watchdog+0x419/0x5c0 [iegbe]</p> <p>[<f4942e90>] iegbe_watchdog+0x0/0x5c0 [iegbe]</p> <p>[<c042dc29>] run_timer_softirq+0xfb/0x151</p> <p>[<c042a77a>] __do_softirq+0x5a/0xbb</p> <p>[<c0407461>] do_softirq+0x52/0x9d</p> <p>[<c0407406>] do_IRQ+0xa5/0xae</p> <p>[<c040592e>] common_interrupt+0x1a/0x20</p> <p>[<c04e5edd>] delay_tsc+0x9/0x13</p> <p>[<c04e5f10>] __delay+0x6/0x7</p> <p>[<f48c1519>] gcu_read_eth_phy+0x59/0x130 [gcu]</p> <p>[<f4946f63>] iegbe_read_phy_reg+0x163/0x180 [iegbe]</p> <p>[<f493f325>] iegbe_update_stats+0x855/0x880 [iegbe]</p> <p>[<f4941a56>] iegbe_get_stats+0x16/0x20 [iegbe]</p> <p>[<c05adeed>] dev_seq_show+0x22/0x8e</p> <p>[<c048bdae>] seq_read+0x191/0x273</p> <p>[<c048bc1d>] seq_read+0x0/0x273</p> <p>[<c0471174>] vfs_read+0x9f/0x141</p> <p>[<c04715c2>] sys_read+0x3c/0x63</p> <p>[<c0404eff>] syscall_call+0x7/0xb</p>
Implication	The Gigabit Ethernet driver's use of spin_lock_bh() to protect the MDIO registers can cause a soft-lockup issue.
Resolution	This issue was resolved with Embedded.L.1.0.3
Affected OS	CentOS 5.2 Linux Red Hat Enterprise Linux 5.0
Driver/Module	Gigabit Ethernet Controller Driver

5 Known Issues - Security

For supplementary information relating to the Known Issues, please refer to the following documents:

- Intel® EP80579 Software for Security Applications on Intel® QuickAssist Technology Programmer's Guide, Order Number: 320183
- Intel® EP80579 Software for Security Applications on Intel® QuickAssist Technology for Linux* Getting Started Guide, Order Number: 322798



- Intel® EP80579 Software for Security Applications on Intel® QuickAssist Technology for FreeBSD* Getting Started Guide, Order Number: 322799

Note: Some open issues listed below require applying patches that are described in the Getting Started Guide. No additional action is required.

Note: If the Affected OS field in an errata table lists Red Hat Enterprise Linux 5.0, readers can assume the errata is also present under CentOS v5.2 Linux which is supported in this release.

Table 4. Summary of Security Software Open Issues

IXA00239764 - OCF cannot use Look Aside Crypto Random number generator without a patch	17
IXA00239767 - Linux kernel will only invoke OCF once to generate a random number	18
IXA00281625 - Keygen: TLS PRF as documented in RFC 2246 for TLS v 1.0	18
IXA00304731 - Synchronous symmetric operation timeout under extreme loads	19
IXA00320219 - Decreasing value of the metric 'MemFree' in /proc/meminfo may not imply a memory loss in Linux or the application	19
IXA00342544 - IPSec L3 forwarding packet order issue	20
IXA00343567 - Computing hash value for NULL buffers fails with an error message	20
IXA00343604 - Firmware limitation - 64K buffer size	21
IXA00345188 - Soft lockup occurs when high rate of traffic is pumped across Openswan tunnel shortly following tunnel being brought up	21
IXA00365361 - Memory leak identified related to mbuf_cluster on FreeBSD 7.1	22

5.1

IXA00239764 - OCF cannot use Look Aside Crypto Random number generator without a patch

Title	OCF cannot use Look Aside Crypto Random number generator without a patch
Reference #	IXA00239764
Description	OCF 20071215 uses a statically declared array of memory to store random data. This memory comes from the heap which is not physically contiguous. The Look-Aside Crypto component requires the data pointer to be physically contiguous.
Implication	The Look-Aside Crypto random number generator cannot be used with OCF unless a patch is applied to OCF to change the random number memory from a static array to kmalloc'd memory.
Resolution	AFTER the OCF patch (ocf-linux-26-20071215.patch) has been applied to the Linux kernel, the following change needs to be applied to \$KERNEL_SOURCE_ROOT/drivers/char/random.c: At line 720: - input_pool.entropy_count < random_write_wakeup_thresh); Should be changed to: + input_pool.entropy_count <= random_write_wakeup_thresh); The kernel must then be re-built. Also, see Errata item IXA00239767. This update does not need to be applied more than once.
Affected OS	Red Hat Enterprise Linux 5.0
Driver/Module	Opensource Software



5.2 IXA00239767 - Linux kernel will only invoke OCF once to generate a random number

Title	Linux kernel will only invoke OCF once to generate a random number
Reference #	IXA00239767
Description	The OCF 20071215 release patches the Linux kernel random.c file with changes that enable OCF to be called to create a random number and save this into the Linux entropy pool. The changes will result in OCF only getting invoked once to fill the Linux entropy pool with random data.
Implication	The Linux kernel patched with OCF will only invoke OCF once to fill the Linux entropy pool with data. When the random data left in the Linux entropy pool goes below a threshold, only the other devices in the system configured to fill the pool with random data will be invoked.
Resolution	AFTER the OCF patch (ocf-linux-26-20071215.patch) has been applied to the Linux kernel, the following change needs to be applied to \$KERNEL_SOURCE_ROOT/drivers/char/random.c: At line 720: - input_pool.entropy_count < random_write_wakeup_thresh); Should be changed to: + input_pool.entropy_count <= random_write_wakeup_thresh); The kernel must then be re-built. Also, see Errata item IXA00239764. This update does not need to be applied more than once.
Affected OS	Red Hat Enterprise Linux 5.0
Driver/Module	Opensource Software

5.3 IXA00281625 - Keygen: TLS PRF as documented in RFC 2246 for TLS v 1.0

Title	Keygen: TLS PRF as documented in RFC 2246 for TLS v 1.0
Reference #	IXA00281625
Description	<p>The key generation API accelerates the TLS PRF, which is defined as part of RFC2246. One of the inputs to this function is a label. The API defines an enumerated type with values that correspond to some of the required labels. However, for some of the operations/labels required by RFC2246, no values are specified. The following are the operations/labels specified by RFC2246 and the corresponding enum values, where defined:</p> <pre>+-----+-----+-----+-----+ :RFC2246: : : : :Section: Operation : Label : API Enum : +-----+-----+-----+-----+ : 6.3 : Derive the key material : "key expansion" : CPA_CY_KEY_TLS_OP_KEY_MATERIAL_DERIVE : : : Derive final client write key : "client write key" : : : : Derive final server write key : "server write key" : : : : Derive IV block : "IV block" : : : 7.4.9 : Client finished : "client finished" : CPA_CY_KEY_TLS_OP_CLIENT_FINISHED_DERIVE : : : Server finished : "server finished" : CPA_CY_KEY_TLS_OP_SERVER_FINISHED_DERIVE : : 8.1 : Computing the master secret : "master secret" : CPA_CY_KEY_TLS_OP_MASTER_SECRET_DERIVE : +-----+-----+-----+-----+</pre>
Implication	For some of the operations and labels required by RFC2246, no supported enum type is provided. A user-defined value must be provided.
Resolution	For those operations/labels above for which no enum value is provided, the client should use the enum value CPA_CY_KEY_TLS_OP_USER_DEFINED, and pass the label using the userLabel field of the CpaCyKeyGenTlsOpData data structure.
Affected OS	Red Hat Enterprise Linux 5.0
Driver/Module	Crypto Module



5.4 IXA00304731 - Synchronous symmetric operation timeout under extreme loads

Title	Synchronous symmetric operation timeout under extreme loads
Reference #	IXA00304731
Description	<p>In some cases, under heavy loads / high data rates, for example running several hundred large random number generations, a synchronous symmetric operation may timeout and report a failure when run simultaneously. This is a rare occurrence and is caused by a synchronous operation's wait-queue timing out. The current wait-queue time limit is set to 1 second for synchronous symmetric operations. The symmetric operation wait-queue timeout period is defined in LAC_SYM_SYNC_CALLBACK_TIMEOUT in the file lac_sync.h.</p> <p>Under 'normal' random number usage this should not be an issue. This may only be evident when flooding the system with large random number requests and synchronous symmetric operations.</p>
Implication	A synchronous symmetric operation may timeout and fail under heavy loads.
Resolution	If a synchronous timeout occurs while performing operations, then increase the timeout of LAC_SYM_SYNC_CALLBACK_TIMEOUT in the file lac_sync.h. Also reduce the number of simultaneous large random number requests and synchronous operations.
Affected OS	Red Hat Enterprise Linux 5.0
Driver/Module	Crypto Module

5.5 IXA00320219 - Decreasing value of the metric 'MemFree' in /proc/meminfo may not imply a memory loss in Linux or the application

Title	Decreasing value of the metric 'MemFree' in /proc/meminfo may not imply a memory loss in Linux or the application
Reference #	IXA00320219
Description	<p>When the performance sample code was run for several days, the metric 'MemFree' in /proc/meminfo indicated a non-trivial loss in free memory. When memory loss was ruled out in Intel drivers and sample code, the same metric was monitored on an idle Linux system over a similar period of time. Once again, the metric pointed to a memory loss.</p> <p>Further investigation on the web, indicated that the apparent loss in memory reported in /proc/meminfo is caused by the kernel not freeing memory released by applications. The kernel does not release memory to realize other optimizations. However, this unused memory can be released non-destructively by setting the tunable parameter as described in the Resolution section.</p>
Implication	When memory loss is suspected or reported by memory tools while running applications, performance suites, or stress tests, the customer should use prescriptions in the Resolution section to rule out memory grabbed by kernel as described in this erratum.
Resolution	<p>Writing to /proc/sys/vm/drop will cause the kernel to drop clean caches, dentries, and inodes from memory, causing that memory to become free.</p> <p>To free pagecache: echo 1 > /proc/sys/vm/drop_caches</p> <p>To free dentries and inodes: echo 2 > /proc/sys/vm/drop_caches</p> <p>To free pagecache, dentries and inodes: echo 3 > /proc/sys/vm/drop_caches</p> <p>As this is a non-destructive operation, and dirty objects are not freeable, the user should run "sync" first in order to make sure all cached objects are freed.</p> <p>This tunable field was added in kernel version 2.6.16.</p>
Affected OS	Red Hat Enterprise Linux 5.0
Driver/Module	Opensource Software



5.6 IXA00342544 - IPSec L3 forwarding packet order issue

Title	IPSec L3 forwarding packet order issue
Reference #	IXA00342544
Description	When Openswan (KLIPS) is patched with OCF patch for asymmetric cryptographic processing of IPSec packets, packets may be sent in a different order than received. This only occurs when a cryptographic driver is available for use with OCF, such as OCF cryptosoft or the OCF shim (enabling Intel® QuickAssist Technology). Packets are not re-ordered due to the drivers themselves, but rather additional contention for a bottom half spinlock due to asymmetric processing.
Implication	Reordering of traffic within an IP flow (IP flow being the pair of source and destination address), is not in violation of the Internet Protocol, but reordering within a router unnecessarily causes unneeded load on the endpoint of the flow and higher layer protocols.
Resolution	<p>The root cause of this issue is in the KLIPS IPSec stack and is caused by contention for a bottom half spin lock in callbacks from OCF (which tend to be grouped together). When a group of callbacks are executed in quick succession, contention for the lock causes a FILO scenario.</p> <p>The ordering issue is worked around using the OCF callback wait queue, which ensures callback order by causing each callback to be executed serially, each waiting for the previous callback to be completed before it is run.</p> <p>This workaround resulted in a soft lockup during the re-keying process of an IPSec tunnel. The system recovers from the soft lockup once traffic is turned off. The exact root cause of this issue was not found.</p> <p>The soft lockup issue's workaround is to use work queues for all packet processing calls in KLIPS before the bottom half lock comes into use (that is, the ipsec_xsm and ipsec_rsm functions).</p>
Affected OS	CentOS 5.2 Linux
Driver/Module	Opensource Software

5.7 IXA00343567 - Computing hash value for NULL buffers fails with an error message

Title	Computing hash value for NULL buffers fails with an error message
Reference #	IXA00343567
Description	It is meaningful to compute a MD5, SHA1/2 hash value for a NULL string; however, the routine cpaCySymPerformOp API fails with an Invalid Parameter message if the input parameter pSrcBuffer stores a zero length CpaFlatBuffer.
Implication	FIPS certification of SHA algorithms requires computation of hash value of NULL strings, and the corresponding test vector would fail to produce the expected results.
Resolution	Pass a single non-NULL CpaFlatBuffer which points to a valid memory location and is contained within the CpaBufferList provided in the parameter pSrcBuffer on the API cpaCySymPerformOp. This will circumvent the existing check within the source code to prevent NULL source data allowing a zero length request. The field messageLenToHashInBytes in the parameter CpaCySymOpData also needs to be set to zero.
Affected OS	FreeBSD 7.1 CentOS 5.2 Linux
Driver/Module	Crypto Module



5.8 IXA00343604 - Firmware limitation - 64K buffer size

Title	Firmware limitation - 64K buffer size
Reference #	IXA00343604
Description	<p>As allowed by the API specification, the implementation of the API on the Intel(R) EP80579 Integrated Processor with Intel(R) QuickAssist Technology does not support buffers of length greater than or equal to 65536 bytes (64K, i.e. where the representation of the length required more than 16 bits). This is true even though the length on the API is a 32-bit number.</p> <p>If a length greater than 64K is specified, an error is printed to /var/log/messages, an 'Invalid API Param' error is returned to the calling function, and the buffer is not processed.</p> <p>Note that OCF supports lengths of greater than 64K. The OCF shim does not check for this condition.</p>
Implication	Clients should not use the API to perform cryptographic operations on buffers whose length is 64K or greater.
Resolution	This is a limitation of the underlying firmware; no resolution is planned.
Affected OS	FreeBSD 7.1 CentOS 5.2 Linux
Driver/Module	QuickAssist (R) Driver

5.9 IXA00345188 - Soft lockup occurs when high rate of traffic is pumped across Openswan tunnel shortly following tunnel being brought up

Title	Soft lockup occurs when high rate of traffic is pumped across Openswan tunnel shortly following tunnel being brought up
Reference #	IXA00345188
Description	A soft lockup occurs when an excessive rate of traffic is pumped across the Openswan tunnel for greater than 30 seconds after initially bringing up an IPSec tunnel. The development board recovers once traffic is stopped and the soft lockup does not occur again if the same amount of traffic is resumed. To reproduce the soft lockup, the IPSec tunnel must be restarted. Traffic sent is 100Mbit/s of traffic @64byte frames, which is 3x time the max throughput at this packet size.
Implication	Excessive rates of traffic (inducing packet loss) should be avoided directly after starting an IPSec tunnel. Initially a low traffic rate should be used, for example, 35 Mbit/s for 64 bit packets.
Resolution	Do not exceed 35 Mbit/s @ 64bit packets for a duration greater than 30 seconds directly after restarting an IPSec tunnel. Note: This issue does not occur during IPSec SA rekeying.
Affected OS	CentOS 5.2 Linux
Driver/Module	Opensource Software



5.10 IXA00365361 - Memory leak identified related to mbuf_cluster on FreeBSD 7.1

Title	Memory leak identified related to mbuf_cluster on FreeBSD 7.1
Reference #	IXA00365361
Description	<p>At high data rates, mbuf_cluster memory is seen to be allocated but not freed. The amount of allocated memory for mbuf clusters may be seen using the 'vmstat -z' command. When high data rates are used, the allocated amount always increases, but never decreases. Once there are no mbuf_clusters available, it is not possible to connect to the machine over Ethernet. COM port connection is still possible.</p> <p>This problem was reproduced as follows:</p> <ul style="list-style-type: none"> a) IPSec setup as described in FreeBSD Getting Started Guide, single tunnel, subnet to subnet connection b) Packet generator used for bi-directional traffic at 100% line rate (1 Gb of traffic) using 64 bit packets c) Traffic generation initiated before IPSec tunnel is up <p>In this scenario, the following things are seen:</p> <ul style="list-style-type: none"> a) mbuf_cluster memory is fully consumed. If traffic is turned off, these do not get freed back up (memory leak). b) QuickAssist (-3) errors are seen due to the perform cookie pool being fully consumed. If traffic is turned off, cookies are freed back to the system (no memory leak). c) If debug is compiled into the OCF shim, some conflicting messages are seen in /var/log/messages, meaning that the OCF shim is not handling the mbuf_clusters correctly.
Implication	IPSec tunnel eventually fails at high data rates. For example, for 64 bit packets the IPSec tunnel eventually fails at data rates above 10 Mb/s.
Resolution	<p>If only a small amount of traffic (a few packets) is used to kick off the tunnel setup and the full line rate is only sent once the tunnel is fully set up, then the system remains stable.</p> <p>Also:</p> <ul style="list-style-type: none"> a) mbuf_cluster memory is only unused if traffic is set above 10 Mb/s for 64 bit packets. Traffic should be kept below this level. b) QuickAssist running out of cookies can be worked around by changing the NUM_CONCURRENT_LAC_SYMMETRIC_REQUESTS parameter in the /etc/icp_asd.conf from 768 to 9000 if 100% line rate is going to be used.
Affected OS	FreeBSD 7.1
Driver/Module	Opensource Software

6 Resolved Issues - Security

Issues that have been resolved in previous versions of the software are included here.

Table 5. Summary of Security Software Resolved Issues

IXA00239768 - Openswan patched with OCF uses the incorrect bottom half spinlock variant.....	23
IXA00241987 - High data rates expose a memory leak in the openswan-2.4.9-ocf-linux-20070727.patch file.....	23
IXA00277532 - IPSec performance improvement modifications to Gigabit Ethernet driver on Linux.....	24
IXA00283570 - Kernel panic when removing icp_ocf driver while passing traffic (>400Mbps)	24
IXA00298242 - Changes required to enable Security Acceleration on A0 Silicon	25
IXA00302585 - CpaCySymCbFunc: Symmetric callback function does not document CPA_STATUS_RETRY status	25
IXA00309248 - Only ACPI mechanism should be used to retrieve [N]CDRAM info from BIOS.....	25
IXA00325998 - FreeBSD system hangs when QuickAssist APIs are called from MOD_LOAD handler	26
IXA00328199 - Multiple reloading of both Security and Embedded modules may result in a crash.....	26
IXA00329534 - PKE functionality issues with OpenSSL patched by OCF on Linux	27



IXA00331333 - Unloading icp_asd module results in interrupts consuming 75% of CPU cycles in FreeBSD .	28
IXA00331336 - Repeated start/stop of QuickAssist modules results in QuickAssist initialization failure.....	28
IXA00331338 - Intel® EP80579 Development Platform does not recover from soft lockup caused by a SSL stress test on Linux	30
IXA00335408 - Repeated start/stop of QuickAssist modules results in a system crash.....	30
IXA00343583 - hmacWithSHA1 algorithm is incorrectly reported to be supported in OpenSSL after being patched for OCF	31
IXA00343599 - FreeBSD 7.1 OpenSSL cryptodev engine is not used for OpenSSH, some PKE acceleration.	31

6.1

IXA00239768 - Openswan patched with OCF uses the incorrect bottom half spinlock variant

Title	Openswan patched with OCF uses the incorrect bottom half spinlock variant
Reference #	IXA00239768
Description	Openswan 2.4.9 is patched with OCF 20070727 with the Openswan 2.4.9 is patched with OCF 20070727 with the ocf-openswan-2.4.9-20070727.patch file. This patch applies the incorrect type of spinlock to the Openswan code.
Implication	The system may crash when running high data rates through a VPN tunnel using Openswan 2.4.9 and OCF 20070727.
Resolution	AFTER the OCF patch (ocf-openswan-2.4.9-20070727.patch) has been applied to the Openswan code, the ocf-openswan-2.4.9-20070727-session-migration-backport.patch file needs to be applied to the Openswan installation folder with the command: patch -p1 < \$ICP_ROOT/OpenSourcePatches/ocf-openswan-2.4.9-20070727-session-migration-backport.patch Also, see Errata item IXA00241987. The patch does not need to be applied more than once. Documentation has been updated to include application of this patch, which is supplied as standard.
Affected OS	Red Hat Enterprise Linux 5.0
Driver/Module	Opensource Software

6.2

IXA00241987 - High data rates expose a memory leak in the openswan-2.4.9-ocf-linux-20070727.patch file

Title	High data rates expose a memory leak in the openswan-2.4.9-ocf-linux-20070727.patch file
Reference #	IXA00241987
Description	When passing traffic at very high data rates over an OCF accelerated Openswan IPsec VPN tunnel, a memory loss can be observed in the system. This is most acute with bi-directional traffic tests with very obvious packet loss.
Implication	All of the available memory in the system can be consumed in a relatively short period of time when running bi-directional tests at very high data rates when packet loss is obvious.
Resolution	AFTER the OCF patch (ocf-openswan-2.4.9-20070727.patch) has been applied to the Openswan code, the ocf-openswan-2.4.9-20070727-session-migrationbackport.patch file needs to be applied to the Openswan installation folder with the command: patch -p1 < \$ICP_ROOT/OpenSourcePatches/ocf-openswan-2.4.9-20070727-session-migration-backport.patch. Also, see Errata item IXA00239768. The patch does not need to be applied more than once. Documentation has been updated to include application of this patch, which is supplied as standard.
Affected OS	Red Hat Enterprise Linux 5.0
Driver/Module	Opensource Software




6.3 IXA00277532 - IPSec performance improvement modifications to Gigabit Ethernet driver on Linux

Title	IPSec performance improvement modifications to Gigabit Ethernet driver on Linux
Reference #	IXA00277532
Description	For IPSec, the Openswan application requires additional memory for the IPSec headers and footers beyond what is originally allocated for the incoming packet. To do this, an additional alloc_skb() is called to create the new buffer and skb_copy_expand() is called to copy the original sk_buff to the new larger buffer.
Implication	These additional memory creation and copy buffer calls impact performance for IPSec.
Resolution	This issue was resolved in Security.L.1.0.2
Affected OS	Red Hat Enterprise Linux 5.0
Driver/Module	Gigabit Ethernet Controller Driver

6.4 IXA00283570 - Kernel panic when removing icp_ocf driver while passing traffic (>400Mbs)

Title	Kernel panic when removing icp_ocf driver while passing traffic (>400Mbs)
Reference #	IXA00283570
Description	<p>Removing icp_ocf driver while passing traffic of more than 400 Mbs generates a kernel panic. The log file shows the following:</p> <pre>icp_ocfDrvSymProcess(): icp_LacSymOpPerform failed, lacStatus = 2 icp_ocfDrvSymProcess(): icp_LacSymOpPerform failed, lacStatus = 2 icp_ocfDrvSymProcess(): icp_LacSymOpPerform failed, lacStatus = 2 icp_ocfDrvSymProcess(): icp_LacSymOpPerform failed, lacStatus = 2 icp_ocfDrvSymProcess(): icp_LacSymOpPerform failed, lacStatus = 2 [error] LacSessionTableDescInvalidate() - : Session cannot be deregistered as there are 128 or more callbacks pending [error] LacSessionTableDescInvalidate() - : Session cannot be deregistered as there are 134 or more callbacks pending icp_ocfDrvExit(): 2 LAC sessions were not deregistered correctly. This is not a clean exit! slab error in kmem_cache_destroy(): cache `ICP Op Data': Can't free all objects [<c1062f7d>] kmem_cache_destroy+0x88/0x133 [<ecca82e6>] icp_ocfDrvFreeCaches+0x36/0xe0 [icp_ocf] [<ecca8558>] icp_ocfDrvExit+0x128/0x140 [icp_ocf] [<c103b677>] sys_delete_module+0x192/0x1bb [<c10464b8>] audit_syscall_entry+0x11c/0x144 [<c1003d0b>] syscall_call+0x7/0xb</pre>
Implication	There will be a kernel panic when the current OCF driver is removed while traffic is passing.
Resolution	<p>AFTER the OCF patch (ocf-linux-26-20071215.patch) has been applied to the Linux kernel, the ocf-linux-20071215-driver-removal.patch patch file needs to be applied to \$KERNEL_SOURCE_ROOT with the command:</p> <pre>patch -p0 < \$ICP_ROOT/OpenSourcePatches/ocf-linux-20071215-driver-removal.patch</pre> <p>Documentation has been updated to include application of this patch, which is supplied as standard.</p>
Affected OS	Red Hat Enterprise Linux 5.0
Driver/Module	Opensource Software



6.5 IXA00298242 - Changes required to enable Security Acceleration on A0 Silicon

Title	Changes required to enable Security Acceleration on A0 Silicon
Reference #	IXA00298242
Description	The current software release will not run correctly on A0 silicon due to a silicon change which corrects the interpretation of the setting of the 'interrupt disable' (INTD) bit for the Gigabit Ethernet driver. The GigE driver uses MSI interrupts; the software currently sets the interrupt disable bit to a 1, to disable INTx interrupts. However, on A0 silicon, setting this bit to a 1 incorrectly disables MSI interrupts too, so the bit needs to be set to 0.
Implication	The code changes detailed below are required to enable correct operation on earlier (A0) silicon revisions.
Resolution	This issue was resolved with B0 silicon.
Affected OS	Red Hat Enterprise Linux 5.0
Driver/Module	QuickAssist (R) Driver

6.6 IXA00302585 - CpaCySymCbFunc: Symmetric callback function does not document CPA_STATUS_RETRY status

Title	CpaCySymCbFunc: Symmetric callback function does not document CPA_STATUS_RETRY status
Reference #	IXA00302585
Description	In some cases, the callback function from cpaCySymPerformOp() may return with a status of CPA_STATUS_RETRY, even though only CPA_STATUS_SUCCESS and CPA_STATUS_FAIL are specified by the documentation. This behavior has been observed for symmetric operations that requires a pre-compute (HMAC, GCM, XCBC).
Implication	This may be visible when stressing the system, for example when performing several hundred random number generations in succession, whilst initializing normal priority sessions and performing operations before the pre-compute has completed. The reason for this is that the random operations fill the request rings resulting in a retry.
Resolution	The client can invoke the operation again for that session and the system will recover. To reduce the likelihood of seeing this behavior, minimize the number of requests to generate random numbers while simultaneously initializing sessions requiring pre-computes.
Affected OS	Red Hat Enterprise Linux 5.0
Driver/Module	Crypto Module

6.7 IXA00309248 - Only ACPI mechanism should be used to retrieve [N]CDRAM info from BIOS

Title	Only ACPI mechanism should be used to retrieve [N]CDRAM info from BIOS
Reference #	IXA00309248
Description	In the current release, if the retrieval of the [N]CDRAM information using ACPI mechanism fails, then the software calls the routine "asd_read_bios_registers" to compute [N]CDRAM information. The returned values are invalid because the routine makes incorrect assumptions about the BIOS settings. NOTE: Only ACPI mechanism should be used in the release package to retrieve [N]CDRAM info from BIOS.
Implication	When the ACPI retrieval of [N]CDRAM information does not succeed, the EP80579 Security Software will fail to load with the error message "Failed to initialize HAL" in Syslog.



Title	Only ACPI mechanism should be used to retrieve [N]CDRAM info from BIOS
Resolution	In the file .../Acceleration/drivers/icp_asd/src/kernel/linux/asd_dram.c, line 175: If the call to <code>asd_get_acpi_vars</code> fails, delete the call to <code>asd_read_bios_registers</code> , display an error instead of a warning, and return FAIL.
Affected OS	Red Hat Enterprise Linux 5.0
Driver/Module	Crypto Module

6.8 IXA00325998 - FreeBSD system hangs when QuickAssist APIs are called from MOD_LOAD handler

Title	FreeBSD system hangs when QuickAssist APIs are called from MOD_LOAD handler
Reference #	IXA00325998
Description	An EP80579-based platform was installed with FreeBSD OS and EP80579 Security application package. The system hung when QuickAssist performance test case was run from MOD_LOAD handler (or other thread owning the Giant mutex). Initial investigation indicated that the OS had assigned the same IRQ to some of the non-MPSafe Embedded drivers, MPSafe Security driver, and OS-released drivers. The sharing of IRQs between non-MPSafe and MPSafe drivers seems to be the cause of the problem. This issue may be a specific to FreeBSD 6.2 and may not be present in later releases of the OS.
Implication	Inappropriate IRQ line sharing amongst MPSafe and non-MPSafe may cause the system hang under FreeBSD.
Resolution	One recommendation is to avoid invoking any APIs, such as QuickAssist APIs, that could result in an interrupt call from MOD_LOAD handler of kernel module or in other words when the Giant mutex is being held. Such APIs could be called from a kernel thread or a task. This problem has not been seen on FreeBSD 7.1.
Affected OS	FreeBSD 6.2
Driver/Module	General

6.9 IXA00328199 - Multiple reloading of both Security and Embedded modules may result in a crash

Title	Multiple reloading of both Security and Embedded modules may result in a crash
Reference #	IXA00328199
Description	The Security driver module crashes randomly if both Security and Embedded modules are loaded and unloaded many times; however, multiple reloadings of just the Security module does not result in a crash. The crash is caused by a bug in FreeBSD 6.2 kernel source <code>/usr/src/sys/vm/vm_contig.c</code> that makes a process sleep in a critical section. This is a known problem and is being addressed by the FreeBSD developer community.
Implication	The user may not be able to load and unload both Security and Embedded modules multiple times.
Resolution	After unloading Security and Embedded modules, the user should reset the OS before reloading Security and Embedded modules. This has not been seen on FreeBSD 7.1
Affected OS	FreeBSD 6.2
Driver/Module	General



6.10

IXA00329534 - PKE functionality issues with OpenSSL patched by OCF on Linux

Title	PKE functionality issues with OpenSSL patched by OCF on Linux
Reference #	IXA00329534
Description	<p>OpenSSL fails to generate RSA Digital Certificates when using OCF based kernel cryptographic operations if the --with-cryptodev-digests switch is used in the configuration command. The failure has been observed both when using the cryptosoft software library and EP80579 OCF driver.</p> <p>Also, the following RSA functions only work if an RSA keysize of 1024, 2048, 3072 or 4096 (the Mod Exp CRT EP80579 OCF driver supported sizes) are used:</p> <pre>openssl dgst -engine cryptodev -sha256 -sign priKey.pem -out mysign.sha1 aFileToSign.txt openssl dgst -engine cryptodev -sha256 -verify pubKey.pem -signature mysign.sha1 aFileToSign.txt</pre> <p>It has also been seen that some certificate generation related PKE functions are not accelerated (although they execute correctly with the OCF EP80579 driver present). The following functions were seen not to be accelerated:</p> <p>Linux (OpenSSL 0.9.8g patched with OCF 2008 patch):</p> <pre>openssl dsaparam -out dsaparam.pem 2048 openssl gendsa -engine cryptodev -out dsaPrivKey.pem dsaparam.pem openssl dsa -engine cryptodev -in dsaPrivKey.pem -noout -text openssl genrsa -engine cryptodev -out priKey.pem openssl rsa -engine cryptodev -in priKey.pem -pubout -out pubKey.pem openssl gendh -engine cryptodev -out key.pem openssl rand -engine cryptodev -out random.text 1024 openssl gendh -engine cryptodev -out key.pem openssl dh -engine cryptodev -in key.pem -noout -text</pre>
Implication	<p>It is not possible to generate an OpenSSL certificate when the full set of OCF-OpenSSL patch functionality is used.</p> <p>DSA, Diffie-Hellman, and RSA certificate generation functions are not accelerated in some cases.</p>
Resolution	<p>Compiling OpenSSL without the '--cryptodev-digests' option allows the RSA certificates to be created correctly using the following command:</p> <pre>openssl req -new -x509 -days 365 -newkey rsa:1024 -nodes -keyout server.key -out server.</pre> <p>Documentation has been updated to reflect that the '--cryptodev-digests' option is not supported.</p>
Affected OS	<p>CentOS 5.2 Linux</p> <p>Red Hat Enterprise Linux 5.0</p>
Driver/Module	Opensource Software



6.11 IXA00331333 - Unloading icp_asd module results in interrupts consuming 75% of CPU cycles in FreeBSD

Title	Unloading icp_asd module results in interrupts consuming 75% of CPU cycles in FreeBSD
Reference #	IXA00331333
Description	<p>When the Security kernel modules and the Gigabit Ethernet driver iegbe.ko are loaded and the following unload command is run: <code>>kldunload icp_asd</code> This results in interrupts consuming 75% of CPU cycles in FreeBSD. The interrupt usage drops to nearly 0% under one of the following conditions:</p> <ol style="list-style-type: none">1. icp_asd module is loaded again and the <code>asd_ctl</code> command is run subsequently; however, the loading process is noticeably slow.2. iegbe.ko is unloaded. <p>The problem lies in the area of IRQ line sharing, interrupt teardown followed by interrupt service routine de-registration by the OS.</p> <p>The problem also occurs in the presence of the driver <code>if_em.ko</code> for the external PCIe NIC Intel PRO/1000.</p> <p>Furthermore, random resource allocation errors (originating in the kernel's <code>vm_contig.c</code> and known to the developer community), leading to QuickAssist initialization failure, occur if other Foundation kernel modules such as <code>dma.ko</code>, <code>timesync.ko</code> or <code>can.ko</code> are loaded.</p>
Implication	The order of loading and loading various Security and Foundation kernel modules in the presence of the external and/or integrated Intel PRO/1000 NIC could result in system's instability.
Resolution	<p>Workaround 1. Unload the iegbe.ko driver first.</p> <p>Workaround 2. In the presence of external NIC, shutdown the system, shift the Intel NIC to another slot to avoid IRQ sharing between the em driver for the Intel NIC and icp_asd and reboot the operating system.</p>
Affected OS	FreeBSD 6.2
Driver/Module	General

6.12 IXA00331336 - Repeated start/stop of QuickAssist modules results in QuickAssist initialization failure

Title	Repeated start/stop of QuickAssist modules results in QuickAssist initialization failure
Reference #	IXA00331336
Description	<p>The FreeBSD/Linux system was built, installed as per the instructions in the Security Getting Started Guide on a platform with 1GB RAM and an external PCIe NIC - Intel PRO/1000 NIC.</p> <p>FreeBSD: The QuickAssist modules were repeatedly started and stopped using the following commands for over 24 hours:</p> <pre>>/etc/rc.d/qat_service_freeBSD start >/etc/rc.d/qat_service_freeBSD stop</pre> <p>The QuickAssist modules eventually fail to install and logs an error in the syslog.</p> <p>Linux: The QuickAssist modules were repeatedly started and stopped about 100 times using the following commands:</p> <pre>>/etc/init.d/qat_service start >/etc/init.d/qat_service stop</pre> <p>The system eventually does a kernel panic due to page fault.</p>
Implication	Starting and stopping the QuickAssist modules repeatedly might result in QuickAssist initialization failure on FreeBSD or a kernel panic on Linux.



Title	Repeated start/stop of QuickAssist modules results in QuickAssist initialization failure (Continued)
Resolution	<p>To eliminate the memory leak, please apply the following diff:</p> <pre> <installation directory>/Acceleration/library/icp_services/RuntimeTargetLibrary/ Target_CoreLibs/uclo/uclo.c 3240,3246d3239 < else < { < if(objHandle->objHdr) < { < ixOsalMemFree(objHandle->objHdr); < } < } 3248d3240 < <installation directory>/Acceleration/library/icp_crypto/QATAL/src/common/ qat_comms/qat_comms.c: 1107a1108 > qatComms_ringInfo[qatReqType].RequestRingID[priority] = 0xFF ; 1217a1219,1220 > > qatComms_ringInfo[qatReqType].ResponseRingID = 0xFF; FreeBSD variant: <installation directory>/Acceleration/drivers/icp_asd/src/kernel/freebsd/ asd_uclo_ldr.c: 85c85,86 < extern int halAe_Init(unsigned int aeMask); --- > extern int halAe_Init(unsigned int aeMask); > extern void halAe_DeLib(void); 223a225,226 > halAe_DeLib(); > Linux variant: <installation directory>/Acceleration/drivers/icp_asd/src/kernel/linux/ asd_uclo_ldr.c: 76,77c76 > extern int halAe_Init(unsigned int aeMask); > extern void halAe_DeLib(void); --- < extern int halAe_Init(unsigned int aeMask); 204,205d202 > halAe_DeLib(); > <installation directory>/Acceleration/library/icp_crypto/QATAL/src/linux/ qatal_symbols.c 113d112 < EXPORT_SYMBOL(halAe_DeLib); </pre>
Affected OS	FreeBSD 6.2 Red Hat Enterprise Linux 5.0
Driver/Module	General



6.13 IXA00331338 - Intel® EP80579 Development Platform does not recover from soft lockup caused by a SSL stress test on Linux

Title	Intel® EP80579 Development Platform does not recover from soft lockup caused by a SSL stress test on Linux
Reference #	IXA00331338
Description	Accelerated Secure web server (Apache & OpenSSL) soft locks up during a SSLSwamp stress test on Linux, and does not recover once traffic is stopped. The test consists of SSLSwamp sending heavy traffic towards a platform running Apache / OpenSSL. The two soft lock ups observed occur after 2 & 5 days of the test running. SSLSwamp command: sslswamp -connect IP:[ipaddress]:443 -update 10 -cipher EDH-RSA-DESCBC3-SHA-time 86400
Implication	High load of HTTPS traffic might trigger this problem.
Resolution	It is not recommended to use the CRB for applications requiring more than 20 secure (https) connections per second. This has not been seen on CentOS 5.2
Affected OS	Red Hat Enterprise Linux 5.0
Driver/Module	Multiple Drivers

6.14 IXA00335408 - Repeated start/stop of QuickAssist modules results in a system crash

Title	Repeated start/stop of QuickAssist modules results in a system crash
Reference #	IXA00335408
Description	Linux: The QuickAssist modules were repeatedly started and stopped for 10-20 minutes using the following commands: >/etc/init.d/qat_service start >/etc/init.d/qat_service stop The system eventually does a kernel panic due to page fault
Implication	A random kernel panic may occur due to /etc/init.d/qat_service stop command.
Resolution	To prevent qat_service stop from crashing please apply the following diff: <installation directory>/Acceleration/drivers/icp_asd/src/kernel/linux/asd_init.c 483,489d482 < if (BIT_IS_SET(asdSubsystemStatus,ISR_RESOURCES_ALLOCATED)) < { < asd_isr_resource_free(accel_dev); < CLEAR_STATUS_BIT(asdSubsystemStatus, ISR_RESOURCES_ALLOCATED); < status = CPA_STATUS_SUCCESS; < } < 503a497,503 > if (BIT_IS_SET(asdSubsystemStatus,ISR_RESOURCES_ALLOCATED)) > { > asd_isr_resource_free(accel_dev); > CLEAR_STATUS_BIT(asdSubsystemStatus, ISR_RESOURCES_ALLOCATED); > status = CPA_STATUS_SUCCESS; > } >
Affected OS	Red Hat Enterprise Linux 5.0
Driver/Module	Crypto Module



6.15 IXA00343583 - hmacWithSHA1 algorithm is incorrectly reported to be supported in OpenSSL after being patched for OCF

Title	hmacWithSHA1 algorithm is incorrectly reported to be supported in OpenSSL after being patched for OCF
Reference #	IXA00343583
Description	hmacWithSHA1 is an algorithm supported by the OCF Shim; however it cannot be called via OpenSSL.
Implication	Due to an issue in the OCF patch to OpenSSL, the hmacWithSHA1 algorithm is shown to be available through the cryptodev engine (after typing the "openssl engine -c" command). However, it cannot be executed through OpenSSL. Using OpenSSL with the hmacWithSHA1 algorithm will result in an "Error: bad input or value" message. The algorithm is not supported by the OpenSSL software.
Resolution	If the --with-cryptodev-digests option is not specified when executing the openssl autoconfigure script, this algorithm is not shown to be available. GSG instructions have been updated to reflect that this option is no longer supported.
Affected OS	CentOS 5.2 Linux Red Hat Enterprise Linux 5.0
Driver/Module	Opensource Software

6.16 IXA00343599 - FreeBSD 7.1 OpenSSL cryptodev engine is not used for OpenSSH, some PKE acceleration

Title	FreeBSD 7.1 OpenSSL cryptodev engine is not used for OpenSSH, some PKE acceleration
Reference #	IXA00343599
Description	There are some situations where acceleration is not used: - Although OpenSSL shows that acceleration is available, it is not used for OpenSSH functions. - Some PKE certificate generation functions are also not accelerated. The following commands were tested, but it is possible similar commands are also not accelerated. openssl dsaparam -out dsaparam.pem 2048 openssl gendsa -engine cryptodev -out dsaPrivKey.pem dsaparam.pem openssl dsa -engine cryptodev -in dsaPrivKey.pem -noout -text
Implication	SSH (secure terminal connection) and SCP (secure file transfer) is not accelerated through Intel software for the FreeBSD 7.1 OS. DSA, Diffie-Hellman and RSA certificate generation functions are not accelerated in some cases.
Resolution	Although in some situations acceleration is not used, the given functionality always works through OpenSSL software.
Affected OS	FreeBSD 7.1
Driver/Module	Opensource Software

7 Related Documentation

7.1 Where to Find Current Software and Documentation

The software release and associated collateral can be found on the Hardware Design resource center.

1. In a web browser, go to <http://www.intel.com/go/soc>
2. For Software and pre-boot firmware: Click on "Tools & Software" tab.



3. For Documentation: Click on “Technical Documents” tab.

The EP80579 security software release package contains encryption software and is subject to export requirements defined by the U.S Department of Commerce. To satisfy these requirements, the End User Certification Form must be filled out and submitted for review/approval. Instructions on this process are included during the download process. Please note that this process may take up to two business days to complete.

7.2 Embedded Documents

The documents in [Table 6](#) provide more information about the Embedded software provided in this release.

Table 6. Embedded Documents

Document Name	Number
Intel® EP80579 Software Drivers for Embedded Applications Programmer's Guide and API Reference Manual	320154
Intel® EP80579 Software Drivers for Embedded Applications on Linux* Getting Started Guide	320151
Intel® EP80579 Software Drivers for Embedded Applications on FreeBSD* Getting Started Guide	320152
Software for Intel® EP80579 Integrated Processor Product Line PHY Porting Guide	320203
Ethernet PHY Selection Criteria for the Intel® EP80579 Integrated Processor Product Line Application Note	320254
Intel® EP80579 Integrated Processor Product Line Datasheet	320066
Intel® EP80579 Integrated Processor Product Line Specification Update	320176

7.3 Security Documents

The documents in [Table 7](#) provide more information about the Security software provided in this release.

Table 7. Security Documents

Document Name	Number
Intel® EP80579 Software for Security Applications on Intel® QuickAssist Technology for Linux* Getting Started Guide	322798
Intel® EP80579 Software for Security Applications on Intel® QuickAssist Technology Programmer's Guide	320183
Intel® EP80579 Software for Security Applications on Intel® QuickAssist Technology Cryptographic API Reference Manual	320184
Intel® EP80579 Software on Intel® QuickAssist Technology Debug Services API Reference Manual	320185
Intel® EP80579 Software for Security Applications on Intel® QuickAssist Technology for FreeBSD* Getting Started Guide	322799
Installing and Using OpenVPN* on Linux* Application Note	321165

