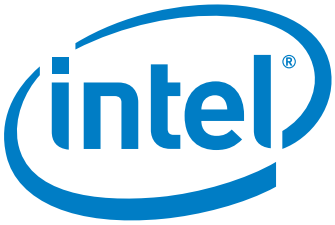

Intel® QuickAssist Technology Software

Package Version: QATmux.L.1.1.0-60

Release Notes

August 2014

Document Number: 330683-002

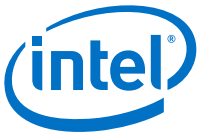

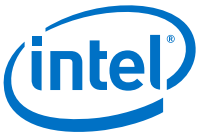

Contents

Tables

Revision History

Date	Revision	Description
July 2014	001	Initial release of document. Based on information in: • Intel® Communications Chipset 8925 to 8955 Series Software Release Notes (523129) • Intel® Communications Chipset 8900 to 8920 Series Software Release Notes (441779) Errata information updated since the last published version of the documents above is shown with change bars.
August 2014	002	• Added the following Erratum for Intel® Communications Chipset 8925 to 8955 Series Software Release Notes: – IXA00387193 - SRIOV: User app needing pfvfcomms APIs and gaeMemDrv doesn't build out-of-box – IXA00387238 - GEN: Heartbeat API will not work in mux env where CCK and CLC devices installed – IXA00387260 - GEN: Device enumeration inconsistency between adf_ctl and the cpa_mux Errata information updated since the last published version of the documents above is shown with change bars.

§ §

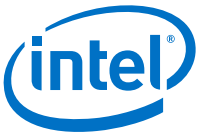

1.0 Description of Release

This document describes extensions and deviations from the release functionality described in the software Programmer's Guides for the various platforms that support Intel® QuickAssist Technology.

For instructions on loading and running the release software, see the Getting Started Guide for your platform (see Section 2.3, "Related Documentation" on page 9).

Note: This software release is intended for platforms that contain:
- Intel® Communications Chipset 8900 to 8920 Series
- Intel® Communications Chipset 8925 to 8955 Series

Note: The Intel® QuickAssist Technology API for kernel space access is currently in the process of being deprecated in favor of making the Intel® QuickAssist Technology services available directly from the Linux kernel, including using the Linux Kernel Crypto framework. User space access is not affected.

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

1.1 Features/Limitations

The main features of the software package are:

Table 1. Features of Platforms Using Intel® QuickAssist Technology

Feature/Limitation	Intel® Communications Chipset 8900 to 8920 Series	Intel® Communications Chipset 8925 to 8955 Series
Cryptographic Services	•	•
Data Compression Services	•	•
Cryptographic Sample Applications	•	•
Data Compression Sample Applications	•	•
Intel® QuickAsssist Technology Data Plane Cryptographic API (cpa_cy_sym_dp.h)	•	•
Intel® QuickAsssist Technology Data Plane Data Compression API (cpa_dc_dp.h)	•	•
Heartbeat Feature	•	•

1.2 Supported Operating Systems

The software in this release has been validated against the operating systems given in the following table on the Customer Reference Boards (CRBs) for the following products:

- Intel® Communications Chipset 8900 to 8920 Series

- Intel® Communications Chipset 8925 to 8955 Series

Note: While the Intel® QuickAssist Accelerator software is validated on Fedora* 16, it should work without change on some other Linux* distributions and kernels.

Table 2. Operating System Support

Operating System	Intel® Communications Chipset 8900 to 8920 Series	Intel® Communications Chipset 8925 to 8955 Series
Fedora* 16 (32-bit and 64-bit)	•	•

1.2.1 Version Numbering Scheme

The software is provided in a top-level package that contains sub-packages for the various supported platforms.

The version numbering scheme for all package levels is the similar:

name.os.major.minor.maintenance-build

Where:

- *name* is the name of the package:
  For the top-level package, the name is “QATmux”
  For sub-packages, name is one of the following:
  - “QAT1.5” for use with Intel® Communications Chipset 8900 to 8920 Series
  - “QAT1.6” for use with Intel® Communications Chipset 8925 to 8955 Series
- *os* is the operating system, in all cases, “Linux*”
- *major* is the major version of the software
- *minor* is the minor version of the software
- *maintenance-build* is the maintenance release and build number

1.2.2 Package Versions

The following table shows the OS-specific package versions for each platform supported in this release.

Chipset or SoC	Package Version
Top-Level Package	QATmux.L.1.1.0-60.tar.gz
Intel® Communications Chipset 8900 to 8920 Series	QAT1.5.L.1.5.0-200.tar.gz
Intel® Communications Chipset 8925 to 8955 Series	QAT1.6.L.2.0.0-102.tar.gz

1.2.3 Licensing for Linux* Acceleration Software

The acceleration software is provided under a dual BSD/GPLv2 license with a few exceptions as listed in Table 3. When using or redistributing duel licensed components, you may do so under either license.

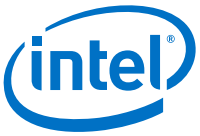

Table 3. Linux* Acceleration Software Licensing Files

Directory Name	License Used	Content Description
./quickassist/utilities/osal/thirdparty/openssl/include/*	OpenSSL	These are OpenSSL header files used for software- based hash pre-computes used with algorithm chaining. These hash pre-computes can be performed in hardware if needed. Refer to the Programmer's Guide for additional information.
./quickassist/lookaside/access_layer/src/sample_code/functional/sym/ssl_sample/*	GPLv2	Functional sample application illustrating the usage of Intel® QuickAssist Technology APIs in an ssl application.
./quickassist/utilities/osal/src/linux/kernel_space/*	GPLv2	Kernel space implementation of OS abstraction layer (OSAL). These functions are used only with the kernel space driver (icp_qa_al.ko). In user space, the OSAL layer source files are dual licensed.

1.2.4 BIOS/Firmware Version

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

It is important to update your platform so that it uses the latest available version of the BIOS/firmware that is available for that platform.

1.2.5 MD5 Checksum Information

The table below gives MD5 checksum information.

	Package	Checksum
Main Package	QATmux.L.1.1.0-60.tar.gz	e56a308c54f9cb98ac321f3d2f11c42e

1.3 Intel® QuickAssist Technology Driver Information

To obtain driver information, use the command below corresponding to your chipset or SoC device:

- For Intel® Communications Chipset 8900 to 8920 Series:
  `cat /proc/icp_dh89xxcc_dev0/version`
- For Intel® Communications Chipset 8925 to 8955 Series:
  `cat /proc/icp_dh895xcc_dev0/version`

The following is example output when using the Intel® Communications Chipset 8925 to 8955 Series. The output is similar for other devices.

```
# cat /proc/icp_dh895xcc_dev0/version

+-------------------------------------------------+
| Hardware and Software versions for device 0      |
+-------------------------------------------------+
 Hardware Version:             A0 SKU2
 Firmware Version:             1.1.0
 MMP Version:                  1.0.0
 Driver Version:               2.0.0
 Lowest Compatible Driver:     2.0
 QuickAssist API CY Version:   1.8
 QuickAssist API DC Version:   1.3
+-------------------------------------------------+
```

1.4 Intel® QuickAssist Technology API Updates

Note: The QAT API version number is different from the software package version number.

For details on any changes to the Intel® QuickAssist Technology APIs, refer to the Revision History pages in the following API reference manuals:

- Intel® QuickAssist Technology Cryptographic API Reference Manual
  API Version 1.8
- Intel® QuickAssist Technology Data Compression API Reference Manual
  API Version 1.3

1.5 Patch Support

This release supports Intel-provided patch software for selected frameworks and applications such as OpenSSL and zlib. Supported patches are available on the 01.org website in the same location as the release software.

§ §

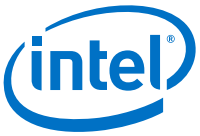

2.0 Where to Find Current Software

The software release and associated collateral can be found on the Intel's Open Source Technology Centre (https://01.org). See the access instructions following.

2.1 Accessing the Software

1. In a web browser, go to https://01.org/packet-processing/intel%C2%AE-quickassist-technology-drivers-and-patches.
2. Scroll down to the table of ATTACHMENTs.
3. Click on the QATmux.L.1.1.0-60.tar.gz link. The browser asks you if you want to download the file.
4. Save the file in the directory of your choice.
5. Unpack and install the software using the instructions in your platform's Getting Started Guide.

Note: The documentation related to the software is also available at the link in step 1.

2.2 List of Files in Release

The Bill of Materials, sometimes referred to as the BOM, is included as a text file in the released software package. This text file is labeled `filelist` and is located at the top directory level for each release.

2.3 Related Documentation

Table 4 lists Intel® QuickAssist Technology generic documentation.

Table 4. Intel® QuickAssist Technology Documentation

Document Name	Reference Number
Intel® QuickAssist Technology API Programmer's Guide	330684
Intel® QuickAssist Technology Cryptographic API Reference Manual	330685
Intel® QuickAssist Technology Data Compression API Reference Manual	330686
Intel® QuickAssist Technology Performance Optimization Guide	330687
Intel® QuickAssist Technology Acceleration Software OS Porting Guide	330688
Using Intel® Virtualization Technology (Intel® VT) with the Intel® QuickAssist Technology	330689

Table 5 lists Intel® Communications Chipset 8900 to 8920 Series specific documentation.

Table 5. Intel® Communications Chipset 8900 to 8920 Series Software Documentation

Document Name	Reference Number
Intel® Communications Chipset 8900 to 8920 Series Software for Linux Getting Started Guide*	330752
Intel® Communications Chipset 8900 to 8920 Series Software Programmer's Guide	330753
Intel® Communications Chipset 89xx Series FIPS Certification Guide	473819

Table 6 lists Intel® Communications Chipset 8925 to 8955 Series specific documentation.

Table 6. Intel® Communications Chipset 8925 to 8955 Series Software Documentation

Document Name	Reference Number
Intel® Communications Chipset 8925 to 8955 Series Software for Linux Getting Started Guide*	330750
Intel® Communications Chipset 8925 to 8955 Series Software Programmer's Guide	330751
Intel® Communications Chipset 89xx Series FIPS Certification Guide Addendum for SKUs 8925 to 8955	523125

§ §

3.0 Intel® Communications Chipset 8925 to 8955 Series Software - Issues

Known and resolved issues relating to the Intel® QuickAssist Technology software are described in this section.

Note: Issue titles follow the pattern:
Identifier - <Component> [Stepping] : Description of issue
where:

<Component> is one of the following:

- CY - Cryptographic
- DC - Compression
- EP - Endpoint
- GEN - General
- SYM DP - Symmetric Cryptography on Data Plane
- SRIOV - Single Root I/O Virtualization
- FIRM - Firmware

[Stepping] is an optional qualifier that identifies if the errata applies to a specific chipset device stepping.

3.1 Known Issues for Intel® Communications Chipset 8925 to 8955 Series

The known issues in the current release are listed below. Change bars indicate additions or updates since the last release of the software for this platform.

Table 7. Summary of Known Issues for Intel® Communications Chipset 8925 to 8955 Series

3.1.1 "IXA00387091 - SRIOV: Issue in pfvfcomms on a guest in a mux installation" on page 20IXA00381319 - CY: digestIsAppended option is not fully supported for Hash-Only operations -

Title	**CY: digestIsAppended option is not fully supported for Hash-Only operations**
Reference #	IXA00381319
Description	Digest Verify is not currently supported for appended digest in Hash-Only operations (i.e. within the CpaCySymSessionSetupData structure, if symOperation=CPA_CY_SYM_OP_HASH then setting both verifyDigest=TRUE AND digestIsAppended=TRUE is not supported).
Implication	Incorrect digest verify result.
Resolution	Do not set digestIsAppended=TRUE for Hash-Only digest verify operations. Instead, ensure that the pDigestResult field of the CpaCySymOpData structure, or digestResult field of the CpaCySymDpOpData structure, is set correctly.
Affected OS	Linux
Driver/Module	CPM IA - Crypto

3.1.2 IXA00382717 - SRIOV: VF will have invalid SKU and capability info for SKU3 -

Title	**SRIOV: VF will have invalid SKU and capability info for SKU3**
Reference #	IXA00382717
Description	When using the Intel® Communications Chipset 8926 (SKU3) that supports data compression (DC) only, the acceleration driver in the VF does not have data about the device SKU and capability set. Consequently, it defaults to assuming that it has full device capabilities, that is, it has support for data compression (DC) and crypto (CY).
Implication	An acceleration driver running on a data compression only SKU (SKU3) on a VF will not behave correctly if configured to run a crypto service. Since the device does not support crypto, the driver should fail to start and report a crypto not supported error. Instead, the device starts, but any crypto content sent to the accelerator using the QA API will be dropped.
Resolution	Ensure that for a SKU3 device (data compression only SKU), the config file has ServicesEnabled = dc.
Affected OS	Linux
Driver/Module	CPM IA - Common

3.1.3 IXA00383477 - DC: Compression sample code does not demonstrate how to handle stateful compression overflow on deflate -

Title	DC: Compression sample code does not demonstrate how to handle stateful compression overflow on deflate
Reference #	IXA00383477
Description	When using stateful compression with data that does not compress, it is possible to get an overflow if the user performs a deflate and has not provided an output buffer large enough to contain all of the data. The stateful sample code does not demonstrate to the user how to handle an overflow in this situation. The stateful sample code does demonstrate how to handle an overflow in a similar situation on an inflate operation, but needs updating to also show how to handle the deflate.
Implication	Without this change, it is unclear to the user that an overflow may occur on a stateful deflate operation and how they should handle the situation.
Resolution	The user currently has to determine what to do from the acceleration driver API and/or from the Zlib Shim.
Affected OS	Linux
Driver/Module	CPM IA - Sample Code

3.1.4 IXA00383480 - GEN: Firmware Mismatch with 32-bit App on 64-bit OS -

Title	GEN: Firmware Mismatch with 32-bit App on 64-bit OS
Reference #	IXA00383480
Description	Note, this bug has not been seen in any DH895xcc release, but is theoretically possible due to a known bug in NUMA memory allocation. This bug will be fixed in a future release. Symptom when seen on a DH89xxcc release: When attempting to run the 32-bit sample code on 64-bit Linux, in some cases the driver may report a firmware mismatch error, for example: ./cpa_sample_code signOfLife=1 [error] LacPke_VerifyMmpLib() - : Error in LAC initialisation. Compiled firmware version (0x972DED54) does not match loaded firmware version (0x00000000)
Implication	If this occurs, the driver will not start.
Resolution	Restarting the driver is a possible workaround for this issue.
Affected OS	Linux
Driver/Module	CPM IA - Common

3.1.5 IXA00383481 - GEN: Installer script fails due to pci.ids file update -

Title	GEN: Installer script fails due to pci.ids file update
Reference #	IXA00383481
Description	The issue can occur if there has been a package update on Fedora 16. The lspci command uses the /usr/share/hwdata/pci.ids file to print additional information about the PCI devices in the system. If that file has been updated to match device ID 0435, the installer script fails to copy the config files to /etc. lspci needs to return the following for the installer script to work correctly: [root@localhost ~]# lspci | grep 0435 02:00.0 Co-processor: Intel Corporation Device 0435 However, with the recent change to the pci.ids, this output is formatted differently causing the installer script to fail.
Implication	The device driver cannot be started with the installer script.
Resolution	Removing the line: 0435 Coleto Creek PCIe Endpoint from the /usr/share/hwdata/pci.ids file allows the installer script to behave correctly.
Affected OS	Linux
Driver/Module	CPM IA - Sample Code

3.1.6 IXA00384101 - GEN: Building the driver with LAC_HW_PRECOMPUTES is not supported in this version of the driver -

Title	GEN: Building the driver with LAC_HW_PRECOMPUTES is not supported in this version of the driver
Reference #	IXA00384101
Description	If the driver is built with the LAC_HW_PRECOMPUTES compiler option, the system may hang and/or crash.
Implication	The LAC_HW_PRECOMPUTES feature should not be used. Software precomputes, which are the default, must be used instead.
Resolution	Do not use the LAC_HW_PRECOMPUTES compiler option.
Affected OS	Linux
Driver/Module	CPM IA - Common

3.1.7 IXA00384236 - GEN: QAT driver build fails on CentOS 6.5 -

Title	GEN: QAT driver build fails on CentOS 6.5
Reference #	IXA00384236
Description	The QAT driver will not build on CentOS 6.5 without copying over the following files from a different kernel (for example, CentOS 6.4): /usr/src/kernels/<kernelver>/include/crypto/md5.h /usr/src/kernels/<kernelver>/include/linux/pci.h /usr/src/kernels/<kernelver>/include/asm-generic/pci.h
Implication	The QAT driver will not build on CentOS 6.5.
Resolution	Copy above files before building.
Affected OS	Linux
Driver/Module	CPM IA - Common

3.1.8 IXA00384677 - Gen: Unnecessary extra interrupts generated in user mode -

Title	Gen: Unnecessary extra interrupts generated in user mode
Reference #	IXA00384677
Description	(Corresponds to IXA00384933 on QAT1.5) In user space, if configured to use interrupts there are more interrupts generated than necessary. It can be seen in /proc/interrupts that the number of interrupts is greater than the number of responses received.
Implication	Unnecessary interrupt processing wastes CPU cycles.
Resolution	None
Affected OS	Linux
Driver/Module	ADF - User Mode

3.1.9 IXA00384678 - GEN: Cannot evenly distribute/affinities interrupts in multi-process environment -

Title	GEN: Cannot evenly distribute/affinities interrupts in multi-process environment
Reference #	IXA00384678

Description	In user space using interrupt mode. Configure to run 1 process per core (32 cores) and give each process access to 1 logical instance on DH895xcc. The device fails to load.
	e.g
	[SSL]
	NumberCyInstances = 1
	NumberDcInstances = 0
	NumProcesses = 32
	LimitDevAccess = 0
	# Crypto - User instance #0
	Cy0Name = "SSL0"
	Cy0IsPolled = 0
	# List of core affinities
	Cy0CoreAffinity = 0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17,18,19,20,21,22,23,24,25,26,27,28,29,30,31
	adf_ctl up fails as follows:
	# ./adf_ctl up
	Processing file: /etc/dh895xcc_qa_dev0.conf
	/etc/dh895xcc_qa_dev0.conf:232: Parse Error - skipping line
	ADF_CONFIG_CTL err: adf_generate_config_params_dh895x: ERROR: Hardware resources unavailable for requested configuration.

Implication	The device fails to load.

Resolution	Use fewer processes, e.g. the following works ok
	[SSL]
	NumberCyInstances = 1
	NumberDcInstances = 0
	NumProcesses = 24
	LimitDevAccess = 0
	# Crypto - User instance #0
	Cy0Name = "SSL0"
	Cy0IsPolled = 0
	# List of core affinities
	Cy0CoreAffinity = 0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17,18,19,20,21,22,23

Affected OS	Linux
Driver/Module	ADF - User Mode

3.1.10 IXA00385318 - DC: Error with stateful compression with CPA_DC_DIR_COMBINED -

Title	**DC: Error with stateful compression with CPA_DC_DIR_COMBINED**
Reference #	IXA00385318
Description	When the session direction is CPA_DC_DIR_COMBINED and the session is STATEFUL, the first decompression operation after at least one compression operation may not succeed.
Implication	Errors may be generated at the start of a session if mixing compression and decompression.
Resolution	Do not use CPA_DC_DIR_COMBINED with a STATEFUL session, or start the CPA_DC_DIR_COMBINED STATEFUL session with a complete compression and a complete decompression.
Affected OS	Linux
Driver/Module	CPM IA - Data Compression

3.1.11 IXA00385777 - GEN: 32-bit apps on 64-bit OS display typo in adf_ctl status -

Title	**GEN: 32-bit apps on 64-bit OS display typo in adf_ctl status**
Reference #	IXA00385777
Description	The ./adf_ctl status prints out inaccurate information when running a 32-bit user application on a 64-bit kernel. For example, on a board with two dh89xxcc and two dh895xcc devices, the following is the output: [root@wgclpixa00381808 build]# ../../QAT1.6/build/adf_ctl s There is 4 acceleration device(s) in the system: icp_dev0 - type=dh89xxcc, inst_id=0, node_id=0, bsf=01:00:0, #accel=2, #engines=8, state=up icp_dev1 - type=xxcc, inst_id=0, node_id=0, bsf=00:00:0, #accel=2, #engines=2, state=up icp_dev2 - type=dh895xcc, inst_id=0, node_id=0, bsf=02:00:0, #accel=6, #engines=12, state=up icp_dev3 - type= , inst_id=0, node_id=0, bsf=00:00:0, #accel=1, #engines=0, state=up
Implication	Invalid data displayed.
Resolution	None
Affected OS	Linux
Driver/Module	CPM IA - Common

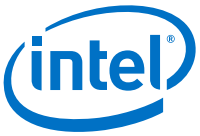

3.1.12 IXA00386283 - GEN: cpaCyGetNumInstances()/ cpaDcGetNumInstances() fails when QATmux loaded with some device types not present -

Title	GEN: cpaCyGetNumInstances()/cpaDcGetNumInstances() fails when QATmux loaded with some device types not present
Reference #	IXA00386283
Description	The QATmux can interface to both QAT1.5 and QAT1.6 to bring up DH89xxcc and DH895xcc devices. If a physical device is not present, the corresponding kernel space driver is usually not loaded, or is loaded but there are no corresponding devices up. In user space, a process will by default attempt to register both QAT1.5 and QAT1.6 drivers. If the QAT1.6 or QAT1.5 dynamic library is present in LD_LIBRARY_PATH this registration will succeed, whether the underlying kernel driver has brought a corresponding device up or not. However, when trying to use CY/DC instances, the cpaCyGetNumInstances()/ cpaDcGetNumInstances() API will fail, preventing the use of the instance in the driver/ devices that are loaded. By removing the dynamic library for the missing devices from the library path, mux registration will throw an error for that particular driver. In this case, the CY/DC instances on the QAT driver which succeeded to register will be available for use. See also IXA00386768
Implication	The cpaCyGetNumInstances()/cpaDcGetNumInstances() API fails, preventing the use of the instance in the driver/devices that are loaded.
Resolution	Remove libqat_1_5_mux_s.so from LD_LIBRARY_PATH if there are no DH89xxcc or C2xxx devices present. Remove libqat_1_6_mux_s.so from LD_LIBRARY_PATH if there are no DH895xcc devices present.
Affected OS	Linux
Driver/Module	CPM IA - Common

3.1.13 IXA00386532 - GEN: Multiprocess test failures observed -

Title	GEN: Multiprocess test failures observed
Reference #	IXA00386532
Description	It has been noticed that the multi-process tests can fail. On failure an unrecoverable error was reported. The test runs for one instance with multiple processes.
Implication	The problem has been root caused to the qaeMemDrv driver that is provided in the sample code. When running user space stress tests using the qaeMemDrv driver, it is possible that an invalid physical address is returned to the upper layer application. This can lead to a firmware hang or an uncorrectable error.
Resolution	This is a sample code issue, not a driver issue. If issue is seen try using a different memory driver.
Affected OS	Linux
Driver/Module	CPM IA - Sample Code

3.1.14 IXA00386643 - GEN: Driver -12 error occurs when configuring for PF/VF concurrency on some platforms -

Title	**GEN: Driver -12 error occurs when configuring for PF/VF concurrency on some platforms**
Reference #	IXA00386643
Description	When starting the driver on some platforms with CentOS6.4, with a build and configuration that supports SR-IOV and with service instances allocated to the PF, -12 error occurs.
Implication	Running sample code fails and generates un-correctable error messages.
Resolution	Under investigation, however, -12 error occurs when calling a kernel function which attaches the virtual function to IOMMU domain fails to find the page directory from the systems page table. This issue was observed on Stargo platform running CentOS 6.4. One alternative is to avoid using the PF for acceleration when using the VF for acceleration.
Affected OS	Linux
Driver/Module	CPM IA - Common

3.1.15 IXA00386644 - GEN: System crash when driver built without SRIOV support is used on a virtualized system -

Title	**GEN: System crash when driver built without SRIOV support is used on a virtualized system**
Reference #	IXA00386644
Description	It has been observed on a virtualization enabled system will crash when trying to bring up the driver which is built without SRIOV support.
Implication	The system will become unresponsive and will require to restart the system.
Resolution	Ensure that driver is built with SRIOV support when running on a virtualized system.
Affected OS	Linux
Driver/Module	CPM IA - Common

3.1.16 IXA00386732 - SRIOV: SIGINT to the sample code application running on host can cause a driver crash -

Title	**SRIOV: SIGINT to the sample code application running on host can cause a driver crash**
Reference #	IXA00386732
Description	It has been observed that when the system has been configured for virtualization (VT-d, SRIOV enabled and intel_iommu=on), sending SIGINT (via Ctrl+C) to the console that is running sample code on the host can cause a driver crash.
Implication	When monitoring "/var/log/messages", the user will observe a DMAR error. The user space application "adf_ctl down" will not be able to bring the driver down. A system reboot will be required.
Resolution	In lieu of application changes, do not use SIGINT to stop the sample code application under these conditions. Instead of doing CTRL+C, we recommend the following approach: - Do Ctrl+Z. This will hang the process. Do not send the process to the background. - Run the "ps" command to list the process ID. - Identify the process ID of the sample code. - Run "kill -9 <process ID of the sample code>".
Affected OS	Linux
Driver/Module	CPM IA - Common

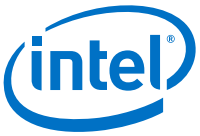

3.1.17 IXA00386733 - GEN: Guest VM reboot issue after icp_reset_dev -

Title	GEN: Guest VM reboot issue after icp_reset_dev
Reference #	IXA00386733
Description	After resetting the Device in the Host without warning the Guest, the Guest VF fails to run Sample Code and the FW hangs. Sequence: - VF running acceleration (userspace sample code signOfLife) - On PF issue reset on dev0 - Reboot VM - Kick off acceleration on the VF Results in: - Host: FW hang detected. - Guest: sampleCodeSemaphoreWait():665 sample_code_semaphoreWait(): errno: 110 waitForResponses():1893 System is not responding.
Implication	Device should not be reset on the Host without first bringing the VF devices down on the VMs.
Resolution	Do adf_ctl down on all VF device on VMs before resetting the device on the Host.
Affected OS	Linux
Driver/Module	CPM IA - Common

3.1.18 IXA00386760 - GEN: Issue while running sample code as a 32-bit application in a 64-bit OS -

Title	GEN: Issue while running sample code as a 32-bit application in a 64-bit OS
Reference #	IXA00386760
Description	Running the user-space sample code with SignOfLife=1 fails for a system configured with 32-bit user space on 64-bit kernel space.
Implication	32-bit user space on 64-bit kernel space package does not work.
Resolution	None
Affected OS	Linux
Driver/Module	CPM IA - Common

3.1.19 IXA00386768 - GEN: MUX package outputs message to stdout -

Title	GEN: MUX package outputs message to stdout
Reference #	IXA00386768
Description	QATmux is designed to work with both QAT1.5 and QAT1.6 drivers and so on user-space process start the icp_sal_userStart API will try to start both of those drivers. However on a platform which is missing one QAT1.x driver, e.g. no QAT1.5 driver loaded and so file removed as no DH89xxcc device present, warnings will be displayed on stdout. e.g. the following warnings are expected: > libqat_1_5_mux_s.so Library not loaded/ libqat_1_6_mux_s.so Library not loaded > qat_1_5 driver not registered with qat_mux / qat_1_6 driver not registered with qat_mux These warnings can be ignored, the user-space process will start with whichever driver is available. The warnings should be removed or moved to stderr. See also IXA00386283
Implication	Applications which depend on parsing stdout may not expect to see these warnings.
Resolution	If only one device type is being used, build and install the driver without QATmux.
Affected OS	Linux
Driver/Module	CPM IA - Common

3.1.20 IXA00387091 - SRIOV: Issue in pfvfcomms on a guest in a mux installation -

Title	SRIOV: Issue in pfvfcomms on a guest in a mux installation
Reference #	IXA00387091
Description	When building the mux package for virtualization, two libqat_1_6_mux so files are created: libqat_1_6_mux_s.so and libqat_1_6_mux_vf_s.so. These are functionally identical; however, if a user-space process on a guest uses the second of these, the system behavior is undefined. This is because the first object is also loaded on the guest during registration with the qat_mux. With both objects loaded, the pfvfComms APIs will not work; there may be also be other problems.
Implication	If using the mux on a guest the libqat_1_6_mux_vf_s.so should not be used.
Resolution	When using the mux on a guest don't use: • libqat_1_6_mux_vf_s.so • libqat_1_5_mux_vf_s.so Instead use: • ibqat_1_6_mux_s.so • libqat_1_5_mux_s.so i.e., copy these objects to /lib64 and link in user-space processes. Delete the ones that shouldn't be used from the /lib64 and the QA build directory.
Affected OS	Linux
Driver/Module	CPM IA - Common

3.1.21 IXA00387193 - SRIOV: User app needing pfvfcomms APIs and gaeMemDrv doesn't build out-of-box -

Title	SRIOV: User app needing pfvfcomms APIs and gaeMemDrv doesn't build out-of-box
Reference #	IXA003870913
Description	A user application compiling in both qae_mem_utils.h and lac_common.h (needed to access pfvfcomms APIs) gets duplicated symbols errors and code doesn't compile.
Implication	qaeMemDrv can't be used out-of-box by app using pfvfcomms.

Resolution	In user application instead of using:
	#include "qae_mem_utils.h"
	add whatever qae_mem_utils fn prototypes are needed directly, e.g.:
	CpaPhysicalAddr qaeVirtToPhysNUMA(void* pVirtAddr);
	void* qaeMemAllocNUMA(Cpa32U size, Cpa32U node, Cpa32U alignment);
	void qaeMemFreeNUMA(void** ptr);

Affected OS	Linux
Driver/Module	CPM IA - Common - Coleto

3.1.22 IXA00387238 - GEN: Heartbeat API will not work in mux env where CCK and CLC devices installed -

Title	GEN: Heartbeat API will not work in mux env where CCK and CLC devices installed
Reference #	IXA00387238
Description	The Heartbeat APIs: * icp_sal_check_device() * icp_sal_check_all_devices() can't be used to access all device types in mux installation where both DH895xcc and DH89xxcc device types are installed in the system. This is because the APIs are not supported in libqat_mux_s.so, just in the individual libqat_1_5_mux_s.so and libqat_1_6_mux_s.so user space objects. If only one of the latter two objects is installed then the user-space app can use the Heartbeat APIs to check the devices supported by that driver. However if both are installed the API will only route to one of the drivers and so can't check devices supported by the other driver.
Implication	Use Heartbeat feature via the /proc interface. This can target any device, e.g. watch -n0.1 cat /proc/icp_dh895xcc_dev0/qat See the applicable Programmer's Guide for full details on how to use the heartbeat feature in this way.

Resolution	In user application instead of using:
	#include "qae_mem_utils.h"
	add whatever qae_mem_utils fn prototypes are needed directly, e.g.:
	CpaPhysicalAddr qaeVirtToPhysNUMA(void* pVirtAddr);
	void* qaeMemAllocNUMA(Cpa32U size, Cpa32U node, Cpa32U alignment);
	void qaeMemFreeNUMA(void** ptr);

Affected OS	OS Linux
Driver/Module	CPM IA - Common - Coleto

3.1.23 IXA00387260 - GEN: Device enumeration inconsistency between adf_ctl and the cpa_mux -

Title	**GEN: Heartbeat API will not work in mux env where CCK and CLC devices installed**
Reference #	IXA00387260
Description	This only affects APIs which take an accelId parameter and are used in a mux installation The adf_ctl utility and the cpa_mux are inconsistent in how they enumerate devices. adf_ctl starts with devices supported by QAT1.5, then QAT1.6. For instance, on a platform with one DH89xxcc and one DH895xcc device: [root@localhost build]# ./adf_ctl s There is 2 acceleration device(s) in the system: icp_dev0 - type=dh89xxcc, inst_id=0, node_id=0, bdf=02:00:0, #accel=2, #engines=8, state=down icp_dev1 - type=dh895xcc, inst_id=0, node_id=1, bdf=82:00:0, #accel=6, #engines=12, state=down However in APIs which take accelId as a parameter, when routed through the qat_mux the enumeration is QAT1.6 devices first, then QAT1.5, so icp_sal_pollBank(0) will be routed to the first device on QAT1.6 and icp_sal_pollBank(1) will be routed to the first device on QAT1.5. These numbers should be aligned, so the X in icp_devX matches the accelID used to call the APIs.
Implication	The adf_ctl status output can't be used directly as the accelID on the API.
Resolution	When calling an API with accelID parameter, device enumeration starts with QAT1.6 devices. adf_ctl status shows all devices available, so the accelID to use can be deduced from that list
Affected OS	OS Linux
Driver/Module	CPM IA - Common - Coleto

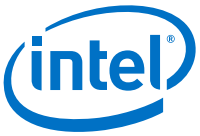

3.2 Resolved Issues for Intel® Communications Chipset 8925 to 8955 Series

This section contains issues resolved since package version 1.0.0. Change bars indicate additions or updates since the last release of the software for this platform.

Table 8. Summary of Resolved Issues for Intel® Communications Chipset 8925 to 8955 Series

3.2.1 IXA00168921 - SRIOV: Physical Function passthrough on VM Guest fails -

Title	SRIOV: Physical Function passthrough on VM Guest fails
Reference #	IXA00168921
Description	The pass-through of physical function to guest OS is not supported in this version of the acceleration driver.
Implication	Physical function passthrough on the VM guest fails. The PF driver (host driver) does not function on a VM.

Title	SRIOV: Physical Function passthrough on VM Guest fails
Resolution	No workaround currently available.
Affected OS	Linux
Driver/Module	CPM IA - Common

3.2.2 IXA00168943 - GEN: Under stress conditions kernel panic can occur -

Title	GEN: Under stress conditions kernel panic can occur
Reference #	IXA00168943
Description	Under stress conditions when running compression and crypto simultaneously kernel panic can occur. Errors similar to the following may be seen in dmesg: BUG: unable to handle kernel paging request at 0000004200000048 Feb 12 18:18:31 localhost kernel: [3454585.660185] IP: [<ffffffffa0375d7d>] dcCompression_ProcessCallback+0x1d/0x550 [icp_qa_al] Or Nov 30 15:04:39 localhost kernel: [193.677963] BUG: unable to handle kernel paging request at 0000004200000040 Nov 30 15:04:39 localhost kernel: [193.678099] IP: [<ffffffffa01ba4f1>] LacSymCb_ProcessCallback+0x21/0x380 [icp_qa_al]
Implication	The QuickAssist driver crashes
Resolution	This issue is resolved in R1.0.1
Affected OS	Linux
Driver/Module	CPM FW - Crypto

3.2.3 IXA00382664 - SRIOV: Instance not in a Running state when QAT is used in host OS and SRIOV is enabled -

Title	SRIOV: Instance not in a Running state when QAT is used in host OS and SRIOV is enabled
Reference #	IXA00382664
Description	The driver prevents the usage of the Physical Function as accelerator if it is built with SRIOV support.
Implication	The driver will report the message "Instance not in a Running state" every time that a function of the QAT API is called from the host OS. Any content sent to the accelerator's Physical Function is dropped.
Resolution	This is resolved with the R1.1 release.
Affected OS	Linux
Driver/Module	CPM IA - Common

3.2.4 IXA00382978 - GEN: Dynamic Instance QAT 32bit application on 64bit OS fail -

Title	GEN: Dynamic Instance QAT 32bit application on 64bit OS fail
Reference #	IXA00382978
Description	Dynamic instance allocation doesn't work if using cross-compilation, i.e. a 32bit user space process running on a 64bit kernel. The following error appears in the log files: allocDcDynamicInstaces():221 Error icp_sal_userDcInstancesAlloc for n dynamic instances For information on using Dynamic instances see Dynamic instance section of Programmer's Guide.
Implication	Instances in [DYN] section of config file should not be used in the cross-compilation case.

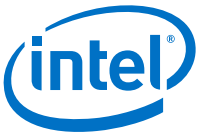

Title	GEN: Dynamic Instance QAT 32bit application on 64bit OS fail
Resolution	There is currently no workaround for this issue. Static instance allocation should be used.
Affected OS	Linux
Driver/Module	CPM IA - Common

3.2.5 IXA00383041 - GEN: Max NumProcesses is limited to 32 -

Title	GEN: Max NumProcesses is limited to 32
Reference #	IXA00383041
Description	The maximum value for NumProcesses is limited to 32. i.e. ## # User Process Instance Section ## [SSL] NumberCyInstances = 0 NumberDcInstances = 1 NumProcesses = 33 LimitDevAccess = 1 gives the following error: [root@sie-lab-214-178 build]# ./adf_ctl icp_dev0 up Processing file: /etc/dh895xcc_qa_dev0.conf ADF_CONFIG_CTL err: adf_read_config_file: ERROR: NumProcesses must be less than or equal to the number of available banks for this device ADF_CONFIG_CTL err: adf_read_config_file: ERROR: dh895xcc has max banks 32. Up to 64 NumProcesses should be allowed as described in the Programmer's Guide, but max allowed is 32 in this release.
Implication	NumProcesses in .conf > 32 will not allow driver to start.
Resolution	This is resolved with the R1.1 release.
Affected OS	Linux
Driver/Module	CPM IA - Common

3.2.6 IXA00383361 - DC: High rates of simultaneous crypto and (de)compression operations may cause (de)compression to report a soft error and write an incorrect amount of data into the output buffer -

Title	**DC: High rates of simultaneous crypto and (de)compression operations may cause (de)compression to report a soft error and write an incorrect amount of data into the output buffer**
Reference #	IXA00383361
Description	A soft error may be reported during some compression/decompression operations. This soft error indicates that an incorrect amount of data was written to the (de)compression output buffer in memory. It may not write enough data to memory or it may write too much data to memory. The issue is most likely to occur when running both symmetric cryptographic operations and compression/decompression operations at the same time with high request/traffic rates. When the issue occurs, the cryptographic operations are unaffected. The issue is highly unlikely to occur when running only compression/decompression operations (i.e. without concurrent cryptographic operations). The issue will not occur when running cryptographic operations alone. In the scenario where not enough data is written to the output buffer, data is missing and hence the compressed/decompressed data is incomplete and hence invalid. In the scenario where too much data is written to the output buffer, the compressed/ decompressed output will be complete and valid, but additional bytes will have been written to the output buffer following the end of the compressed/decompressed data. The extra data written out may be earlier data from this request or previous crypto/ compression requests (potential security vulnerability). Note: The returned output byte counters will be correct in all cases, but will not reflect the amount of data written when the issue occurs. When this issue occurs, the Intel® QuickAssist accelerator driver returns a soft error (CPA_DC_SOFTERR = -12) and also generates an error message in the log that says: "Non-fatal data transfer error detected".
Implication	If too little data is written, the compressed/decompressed data is incomplete and hence invalid. If too much data is written, the compressed/decompressed data is valid, but additional data is present in the output buffer. If the issue occurs, only Intel® QuickAssist accelerator compression/decompression operations are affected.
Resolution	This issue is resolved in R1.0.1
Affected OS	Linux
Driver/Module	CPM IA - Common

3.2.7 IXA00383479 - GEN: Sample code may have errors due to qaeMemDrv issue -

Title	**GEN: Sample code may have errors due to qaeMemDrv issue**
Reference #	IXA00383479
Description	Related to dh89xxcc issue IXA00382836. If running the sample code for 24 hours, there may be errors caused by a known bug in the qaeMemDrv. Note: This has not been seen on this DH895xcc release, only on an earlier DH89xxcc release, but it is theoretically possible. The fix to the qaeMemDrv will be provided in a future release.
Implication	The sample code may fail if run repeatedly.
Resolution	Installing the version of the qae memory driver supplied with the DH89xxCC QAT1.3 package should resolve the issue.
Affected OS	Linux
Driver/Module	CPM IA - Crypto

3.2.8 IXA00383555 - CY: AES XTS mode encrypted data is corrupted if payload not multiple of 16B -

Title	CY: AES XTS mode encrypted data is corrupted if payload not multiple of 16B
Reference #	IXA00383555
Description	AES XTS mode allows for payloads of variable size, not necessarily a multiple of the block size. The encrypted data can be corrupted if the payload is not a multiple of 16B.
Implication	AES XTS mode encrypted data can be corrupted if payload is not a multiple of 16B.
Resolution	Use payloads that are multiples of 16 bytes with AES XTS.
Affected OS	Linux
Driver/Module	CPM FW - Crypto

3.2.9 IXA00383764 - GEN: The acceleration driver does not support the full range of ring size configurations -

Title	GEN: The acceleration driver does not support the full range of ring size configurations
Reference #	IXA00383764
Description	The acceleration driver is currently limiting the max. size of a request/response ring to 256 KB (64K * DWords), which corresponds to 2048 x 128 byte messages. According to the DH895xCC specification, the hardware supports ring buffer sizes up to 1M Dwords (=4 MB) which corresponds to 32K x 128 byte messages.
Implication	This will impact the performance of the acceleration driver. The user will get a RETRY error during packet bursts.
Resolution	The driver now supports ring sizes up to 4M
Affected OS	Linux
Driver/Module	CPM IA - Common

3.2.10 IXA00383882 - GEN: Starting acceleration service on CentOS 6.4 takes a long time -

Title	GEN: Starting acceleration service on CentOS 6.4 takes a long time
Reference #	IXA00383882
Description	Starting the acceleration service usually takes <5s, but can take 30+ seconds if using Linux kernel 2.6.32 on CentOS 6.4. This is due to a secondary bus reset (SBR) of the device being performed as part of the initialization of the driver. After the reset, the driver has to restore PCI configuration space, either manually (via adf_restore_pci_state()) or via kernel API (pci_restore_state()). If using Linux kernel 2.6.32 on CentOS 6.4, the driver has to manually perform the restore of the PCI configuration space, which takes approximately 20+ seconds. The driver doesn't yet support legacy kernel versions. if building CentOS on Linux kernel 2.6.32 the above issue arises.
Implication	adf_ctl up can take >30s.
Resolution	This is resolved with the R1.1 release.
Affected OS	Linux
Driver/Module	CPM IA - Common

3.2.11 IXA00384139 - CY: The wireless firmware can hang -

Title	CY: The wireless firmware can hang
Reference #	IXA00384139
Description	It may be possible under heavy load conditions for the wireless FW to hang.
Implication	The wireless FW may hang.
Resolution	This issue is resolved in R1.0.1
Affected OS	Linux
Driver/Module	CPM IA - Common

3.2.12 IXA00384310 - GEN: Error in osalIOMMUVirtToPhys function -

Title	GEN: Error in osalIOMMUVirtToPhys function
Reference #	IXA00384310
Description	While looking at the IOMMU support in the User Space side of the Osal Memory Driver: OsalUsrKrnProxy.c The function osalIOMMUVirtToPhys is called in the kernel space driver with an unmap call (DEV_MEM_IOC_IOMMUUNMAP) rather than a virt to phys call (DEV_MEM_IOC_IOMMUVTOP).
Implication	Memory deallocation will fail.
Resolution	The call to: - ioctl(fd, DEV_MEM_IOC_IOMMUUNMAP, &info) == 0) ? info.phaddr : 0; can be replaced by: - ioctl(fd, DEV_MEM_IOC_IOMMUVTOP, &info) == 0) ? info.phaddr : 0;
Affected OS	Linux
Driver/Module	OSAL

3.2.13 IXA00384311 - GEN: Dynamic instance allocation issues for 32-bit application on 64-bit OS -

Title	GEN: Dynamic instance allocation issues for 32-bit application on 64-bit OS
Reference #	IXA00384311
Description	The existing solution had an incompatible data structure size for 32 bit user space application running on 64-bit kernel space.
Implication	Dynamic instance would have issues running on 32 bit user space with 64 bit kernel space.
Resolution	This is resolved with the R1.1 release.
Affected OS	Linux
Driver/Module	CPM FW - Crypto

3.2.14 IXA00384312 - DC: Errors not reported if session type is COMBINED and compression is dynamic -

Title	DC: Errors not reported if session type is COMBINED and compression is dynamic
Reference #	IXA00384312
Description	Errors from translation slice are ignored when the session direction is CPA_DC_DIR_COMBINED and huffType is CPA_DC_HT_FULL_DYNAMIC
Implication	If user's application configures the session so that the direction is combined and the compression type is dynamic then the user will not see any potential errors that could occur.

Title	DC: Errors not reported if session type is COMBINED and compression is dynamic
Resolution	This is resolved with the R1.1 release.
Affected OS	Linux
Driver/Module	CPM FW - Crypto

3.2.15 IXA00384313 - CY: Need Param check in crypto APIs for CpaBoolean variables passed as pointers -

Title	CY: Need Param check in crypto APIs for CpaBoolean variables passed as pointers
Reference #	IXA00384313
Description	All the cpaCy APIs that take a boolean parameter as a pointer need to have this parameter check before being used.
Implication	If the user decides to pass a null pointer to this parameter, the driver will report a segmentation fault.
Resolution	This is resolved with the R1.1 release.
Affected OS	Linux
Driver/Module	CPM IA - Crypto

3.2.16 IXA00384314 - GEN: Typo in CONFIG_PCI(E)AER -

Title	GEN: Typo in CONFIG_PCI(E)AER
Reference #	IXA00384314
Description	The linux kernel has the flag CONFIG_PCIEAER set when Advanced Error Reporting is supported. Making a typo on this #define prevents the code within CONFIG_PCIEAER section not to be compiled.
Implication	Advanced Error Reporting feature is not available.
Resolution	This is resolved with the R1.1 release.
Affected OS	Linux
Driver/Module	ADF - Kernel Mode

3.2.17 IXA00384315 - GEN: Error in ring handling when using interrupts in user space -

Title	GEN: Error in ring handling when using interrupts in user space
Reference #	IXA00384315
Description	When using interrupt mode with a user space QA application, the response ring head is not always updated during response processing and so the interrupt condition persists. The root of the issue is that a previous SW optimization coalesces response ring head updates. This optimization should not be applied when the response ring is in interrupt mode.
Implication	The fact that interrupts are re-enabled causes a the same interrupt to fire again multiple times. This leads to significant drop in system performance.
Resolution	This is resolved with the R1.1 release.
Affected OS	Linux
Driver/Module	ADF - User Mode

3.2.18 IXA00384316 - CY: more cleanup code in osal_init is needed -

Title	CY: more cleanup code in osal_init is needed
Reference #	IXA00384316
Description	If function calls with osal_init fail, then the driver will not be unregistered and the crypto interface may not be uninitialized.
Implication	This can lead to memory leaks in the kernel space.
Resolution	This is resolved with the R1.1 release.
Affected OS	Linux
Driver/Module	CPM IA - Crypto

3.2.19 IXA00384317 - CY: Driver lockup with RC4 cipher when hit ring full -

Title	CY: Driver lockup with RC4 cipher when hit ring full
Reference #	IXA00384317
Description	When doing decryption with the RC4 cipher which has lots of threads submitting decrypt requests to a single instance, eventually one thread hits ring full (CPA_STATUS_RETRY). That thread then resubmits the same request but no response will be received.
Implication	For multi threaded operations, if the ring is already full and a re-submitted request receives a CPA_STATUS_RETRY then this request will not receive a response. The bug is in function LacSymQueue_RequestSend() in file lookaside/access_layer/src/common/crypto/sym/lac_sym_queue.c. Before putting the request on the ring, the function assigns: - pSessionDesc->nonBlockingOpsInProgress = CPA_FALSE; This indicates that a blocking operation is in flight. The function then attempts to write the request to the ring. If the write to the ring is not successful, then the function fails to restore the nonBlockingOpsInProgress flag. Subsequent requests are queued in the session's queue, waiting for a pending operation which does not exist.
Resolution	This is resolved with the R1.1 release.
Affected OS	Linux
Driver/Module	CPM IA - Crypto

3.2.20 IXA00384318 - CY: Hit ring full when number of threads << ring size -

Title	CY: Hit ring full when number of threads << ring size
Reference #	IXA00384318
Description	The ring reports full for an invalid reason.
Implication	The driver returns frequent CPA_STATUS_RETRY errors, requests are being resubmitted and performance decreases.
Resolution	This is resolved with the R1.1 release.
Affected OS	Linux
Driver/Module	ADF - User Mode

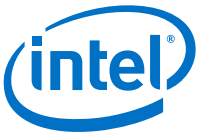

3.2.21 IXA00384319 - DC: Report overflow from XLT when byte count mismatch detected -

Title	DC: Report overflow from XLT when byte count mismatch detected
Reference #	IXA00384319
Description	During a dynamic compression operation, the translator slice will not report an overflow error if the dynamic output is greater than the static output.
Implication	Compressed data returned to the user will be correct. The user will not be aware that XLT has overflowed if the dynamic output is greater than the static output after the re-submit. This has no effect on the user application.
Resolution	This is resolved with the R1.1 release.
Affected OS	Linux
Driver/Module	CPM FW - Data Compression

3.2.22 IXA00384642 - GEN: Heartbeat failure -

Title	GEN: Heartbeat failure
Reference #	IXA00384642
Description	When testing heartbeat, we observed that with a firmware hang, the firmware could not be restarted by heartbeat.
Implication	Under certain conditions, the firmware cannot be restarted by the heartbeat feature.
Resolution	This is resolved with the R1.1 release.
Affected OS	Linux
Driver/Module	CPM IA - 128B

3.2.23 IXA00384681 - CY: Silent drop of messages -

Title	CY: Silent drop of messages
Reference #	IXA00384681
Description	When Crypto partials are in flight, messages can get queued in the driver so they are processed sequentially, that is, the next request is not put on the ring to the firmware until the response from the previous request has been received. If the response ring has been polled 10000 times and a response is not received, the queued requests are silently dropped. Although unlikely to happen, this behavior has been seen for example when submitting requests in the callback.
Implication	An application could end up waiting indefinitely for a response and not be aware that the request has been dropped.
Resolution	This is resolved with the R1.1 release.
Affected OS	Linux
Driver/Module	CPM IA - Crypto

3.2.24 IXA00385003 - GEN: Response ring processing is unnecessarily restricting request submissions -

Title	GEN: Response ring processing is unnecessarily restricting request submissions
Reference #	IXA00385003
Description	Inflight count for ring messages is not updated until after all callbacks handled together have been completed. This can lead to the ring pair reporting full, even though there is space available on the rings.The inflight counter should be decremented as soon as the response message has been picked up in the response ring rather than after a burst of callbacks have been completed.
Implication	A RETRY response can be received when trying to put a message in the ring even though space is available.
Resolution	This is resolved with the R1.1 release.
Affected OS	Linux
Driver/Module	ADF - User Mode

3.2.25 IXA00385148 - CY: Issue with Data Plane API GCM operations when using the acceleration driver -

Title	CY: Issue with Data Plane API GCM operations when using the acceleration driver
Reference #	IXA00385148
Description	According to API documentation, values for hashStartSrcOffsetInBytes and messageLenToHashInBytes are not required for internal purposes for GCM operations. If these values are not set, GCM operations do not work with the Data Plane APIs. The driver is responsible for setting these values and this is not currently being done. This issue does not occur with the traditional GCM operations.
Implication	GCM operations in the Data Plane API fail.
Resolution	This is resolved with the R1.1 release.
Affected OS	Linux
Driver/Module	CPM IA - Common

3.2.26 IXA00385374 - DC: Incorrect compressed data produced in case of level 1 or level 2 compression and CPA_OVERFLOW -

Title	DC: Incorrect compressed data produced in case of level 1 or level 2 compression and CPA_OVERFLOW
Reference #	IXA00385374
Description	Some data compressed with deflate level 1 or level 2 may produce an invalid deflate bitstream in case of overflow. Besides the overflow exception (reported as -11) by the API, there is no other error reported to the user to indicate that the bitstream is invalid. As a consequence, when decompressing the data, a soft error will be reported. Note: This issue has only been seen when an overflow condition occurs.
Implication	Data produced with level 1 or level 2 compression in the case of overflow could be corrupt. In these cases, the output from compression cannot be reliably used to create the original input data.
Resolution	This is resolved with the R1.1 release.
Affected OS	Linux
Driver/Module	CPM IA - Data Compression

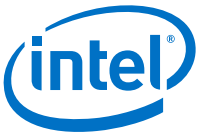

3.2.27 IXA00386072 - DC: Wrong consumed value reported by QAT on compression overflow -

Title	DC: Wrong consumed value reported by QAT on compression overflow
Reference #	IXA00386072
Description	When overflow is encountered during compression a boundary condition can occur which will result in a code representing the last 16 compressed bytes to not be output. The generated block is valid Deflate but that code is missing from the generated stream.
Implication	As the 16 bytes of input has been processed the “bytes consumed” value returned is incremented accordingly. If continuing compression after the overflow event it would appear as if the 16 bytes was skipped. On decompression the 16 bytes would be missing from the produced data.
Resolution	This is resolved with the R1.1 release.
Affected OS	Linux
Driver/Module	CPM FW - Data Compression

3.2.28 IXA00386298 - DC: Mismatch between QAT compressed data and zlib -

Title	DC: Mismatch between QAT compressed data and zlib
Reference #	IXA00386298
Description	Due to an issue on the older version of OSAL, there is a possible mismatch between the data being decompressed by Zlib and the source data compressed by QAT.
Implication	The implication is a reliability issue that can be seen only when stressing the memory management of the driver. In the normal mode of operation, this does not happen. The root cause of the issue is the OSAL module. For some reason, the virtual address being returned is invalid and when doing a virt2phys on it, it is passing an invalid physical address to the firmware. Issues can manifest themselves as a data mismatch or firmware hang. The issue was first observed in the 1.0.5 release.
Resolution	This issue is not seen from 1.1.0 package onwards. The issue is fixed in OSAL module which is now common to DH895xcc and DH89xxcc drivers. This module is different from the OSAL module in the 1.0.x driver releases.
Affected OS	Linux
Driver/Module	CPM FW - Data Compression

4.0 Intel® Communications Chipset 8900 to 8920 Series - Software Issues

Known and resolved issues are described in this section. Unless explicitly stated or noted, each issue relates to the Intel® Communications Chipset 8900 to 8920 Series.

Note: Issue titles follow the pattern:
Identifier - <Component> [Stepping] : Description of issue
where:

<Component> is one of the following:

- CY - Cryptographic
- DC - Compression
- EP - Endpoint
- GEN - General
- SYM DP - Symmetric Cryptography on Data Plane
- SRIOV - Single Root I/O Virtualization
- FIRM - Firmware

[Stepping] is an optional qualifier that identifies if the errata applies to a specific chipset device stepping.

4.1 Known Issues for Intel® Communications Chipset 8900 to 8920 Series

The known issues in the current release are listed below. Changebars indicate additions or updates since the last release of the software for this platform.

Table 9. Summary of Known Issues for Intel® Communications Chipset 8900 to 8920 Series

4.1.1 IXA00372157 - SRIOV: Shutdown of guest with inflight operations results in device hang -

Title	SRIOV: Shutdown of guest with inflight operations results in device hang
Reference #	IXA00372157
Description	A non graceful forced shutdown of a guest with inflight QA operations may result in a device hang. For example, issuing the command 'virsh destroy' for the guest.
Implication	Device hang. Device is unavailable for other guests and the host.
Resolution	Don't force shutdown, use a graceful shutdown, e.g. 'virsh shutdown' or shutdown from within the guest.
Affected OS	Linux
Driver/Module	CPM IA - Common

4.1.2 IXA00382936 - CY: When starting an application that uses a large number (>16) of processes, a timeout error may occur -

Title	CY: When starting an application that uses a large number (>16) of processes, a timeout error may occur
Reference #	IXA00382936
Description	The issue was observed using the openssl speed application patched with libcrypto* (OpenSSL*) Sample Patch for Intel(R) QuickAssist Technology. Libcrypto's openssl speed test uses multiple processes (instead of multiple threads) to measure performance. Each process has access to one crypto instance. When running a test with 16 or 32 processes, the test can fail from time to time reporting: [error] SalCtrl_AdfServicesStartedCheck() - : Sal Ctrl failed to start in given time [error] do_userStart() - : Failed to start services can't use that engine
Implication	Time-out errors may occur when starting a user space application when a large number (>16) of processes is used.
Resolution	No workaround currently available.
Affected OS	Linux
Driver/Module	ADF - User Mode

4.1.3 IXA00384936 - CY: digestIsAppended not fully supported for hash-only -

Title	CY: digestIsAppended not fully supported for hash-only
Reference #	IXA00384936
Description	Digest verify is not currently supported for appended digest in Hash-Only operations (that is, within the CpaCySymSessionSetupData structure, if symOperation=CPA_CY_SYM_OP_HASH then setting both verifyDigest=TRUE AND digestIsAppended=TRUE is not supported). (Corresponds to IXA00381319 on QAT1.6)
Implication	Incorrect digest verify result.
Resolution	Do not set digestIsAppended=TRUE for Hash-Only digest verify operations. Instead, ensure that the pDigestResult field of the CpaCySymOpData structure, or the digestResult field of the CpaCySymDpOpData structure, is set correctly.
Affected OS	Linux
Driver/Module	CPM IA - Crypto

4.1.4 IXA00385637 - DC: Zero-length compression requests are not supported when performing stateless data compression using multiple compress operations. -

Title	DC: Zero-length compression requests are not supported when performing stateless data compression using multiple compress operations.
Reference #	IXA00385637
Description	Zero-length compression requests are not supported when performing stateless data compression using multiple compress operations.
Implication	This feature is not available if the user performs zero-length compression requests.
Resolution	When performing stateless data compression using multiple compress operations, ensure that no requests with a zero length buffer size are submitted to the acceleration driver.
Affected OS	Linux
Driver/Module	CPM IA - Data Compression

4.1.5 IXA00386517 - GEN: Issue in SRIOV Physical passthrough for 2 devices to single VM -

Title	GEN: Issue in SRIOV Physical passthrough for 2 devices to single VM
Reference #	IXA00386517
Description	When two physical function are assigned to a single VM then the sample code fails to run in kernel space.
Implication	The firmware hangs and the following message is displayed. [root@wgclpixa00382221_01 build]# cat /proc/icp_dh89xxcc_dev1/qat0 ERROR: Qat is not responding. Please restart the device
Resolution	The root cause for this issue, is that Qemu does not pass the interrupt request from guest to host. We need to upgrade Qemu to 1.2.0 to get this issue resolved.
Affected OS	Linux
Driver/Module	CPM IA - Common

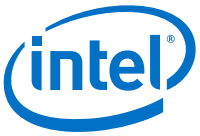

4.1.6 IXA00386756 - GEN: Driver -12 error occurs when configuring for PF/VF concurrency on some platforms -

Title	GEN: Driver -12 error occurs when configuring for PF/VF concurrency on some platforms
Reference #	IXA00386756
Description	When starting the driver on some platforms with CentOS6.4, with a build and configuration that supports SR-IOV and with service instances allocated to the PF, -12 error occurs.
Implication	Running sample code fails and generates un-correctable error messages.
Resolution	Under investigation, however, -12 error occurs when calling a kernel function which attaches the virtual function to IOMMU domain fails to find the page directory from the systems page table. This issue was observed on Stargo platform running CentOS 6.4. One alternative is to avoid using the PF for acceleration when using the VF for acceleration.
Affected OS	Linux
Driver/Module	CPM IA - Common

4.2 Resolved Issues for Intel® Communications Chipset 8900 to 8920 Series

This section contains issues resolved since the following package versions:

- Intel® Communications Chipset 8900 to 8920 Series Software version 1.0.0.

Changebars indicate additions or updates since the last release of the software for this platform.

Table 10. Summary of Resolved Issues for Intel® Communications Chipset 8900 to 8920 Series

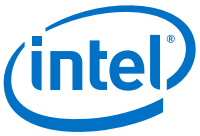

4.2.1 IXA00168573 - CY: Algorithm chaining for AES-GCM can return an incorrect digest when using the Traditional API -

Title	CY: Algorithm chaining for AES-GCM can return an incorrect digest when using the Traditional API
Reference #	IXA00168573
Description	The problem occurs intermittently in a multi-threaded environment and does not affect the cipher text created.
Implication	Incorrect digest can be returned to the user.
Resolution	Addressed in the R1.1.0 Software Package
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.2 IXA00168788 - CY: Application crash when a user space application fails to initialize asymmetric crypto on IRQ mode only -

Title	**CY: Application crash when a user space application fails to initialize asymmetric crypto on IRQ mode only**
Reference #	IXA00168788
Description	When user space app fails to initialize PKE running in IRQ mode, it is possible that the application reports a segmentation fault. This issue is reproduceable when allocating and freeing dynamic instances again and again. The initialization will be done at each allocation. It can also be reproduced calling start/stop process APIs in user space again and again.
Implication	The driver will report a segmentation fault.
Resolution	Addressed in the R1.2 Software Package
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.3 IXA00168886 - DC: Decompression reports -10 soft error after processing stored block of size>=32KB -

Title	**DC: Decompression reports -10 soft error after processing stored block of size>=32KB**
Reference #	IXA00168886
Description	Stateful decompression of >=32K of specific compressed data can result in a -10 error status being returned to the application for a very specific set of compressed data.
Implication	If a user application receives the -10 error, it should split the data into packets of size <32K and resubmit the data.
Resolution	Addressed in the R1.3 Software Package
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.4 IXA00368704 - GEN: icp_sal_userStart hangs when there is no memory in SNB-EN slot -

Title	**GEN: icp_sal_userStart hangs when there is no memory in SNB-EN slot**
Reference #	IXA00368704
Description	When using the Shumway board with the default configuration, where SNB-EP is the host processor, but with only the SNB-EP banks populated with memory modules, the driver is not be able to bring up the acceleration device connected to the SNB-EN processor node. In kernel space, this results in an error logged in the system log. In user space, the icp_sal_userStart call hangs.This is due to an issue in the kernel that is tracked by this bugzilla entry: https://bugzilla.kernel.org/show_bug.cgi?id=42967
Implication	In kernel space, the instances set up on the second device are unavailable.In user space, the user is not able to execute the icp_sal_userStart function; consequently no instances are available.
Resolution	If a user wants to allocate instances on a particular CPU node, then it is required that this node will be populated with memory since the driver will call kmalloc_node (size, node). If no memory plugged into the memory controller for this node, then a crash will occur.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.5 IXA00369518 - DC B0 only: Only marginal increase in compression ratio at certain levels when stateful used and potential fatal error. -

Title	**DC B0 only: Only marginal increase in compression ratio at certain levels when stateful used and potential fatal error.**
Reference #	IXA00369518
Description	Due to current workarounds compression ratio during stateful compression at certain levels is affected. Recommended level is 32K L2 for best ratio gain. A fatal error may happen during stateful compression for level 2, 3, 5, 6, 8 or 9.
Implication	The compression ratio gain from using stateful compression will not be realised at the following levels: 32K History Window: L5,L6,L8,L9 8K History Window: L2,L3,L5,L6,L8,L9 The fatal error will return the error -13 in the compression callback.
Resolution	Resolved with C0 silicon.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.6 IXA00370156 - DC B0 Only - Decompression may hang if dynamic src data is incomplete. -

Title	**DC B0 Only - Decompression may hang if dynamic src data is incomplete.**
Reference #	IXA00370156
Description	The Hardware accelarator may hang during stateless decompression if the src data is incomplete and ends in the middle of a dynamic blocks trees.
Implication	The accelarator will hang and will be unresponsive.
Resolution	Resolved in C0 silicon.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.7 IXA00370174 - DC B0 Only - Decompression can hang if overflow occurs in the middle of a non-empty stored block > 8 Bytes -

Title	**DC B0 Only - Decompression can hang if overflow occurs in the middle of a non-empty stored block > 8 Bytes**
Reference #	IXA00370174
Description	There is a problem in the handling of overflow when it occurs in the middle of a non-empty stored block. Affects stateless and stateful decompression. This issue only affects B0 silicon.
Implication	The Compression slice can hang if an overflow occurs in the middle of a stored block that is greater than 8 bytes.
Resolution	A device reset is required when a hang is encountered. Resolved in C0, this issue only affects B0 silicon.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.8 IXA00371167 - DC B0 Only: Overflow may not be reported when incomplete data sent to the slice for decompression. -

Title	**DC B0 Only: Overflow may not be reported when incomplete data sent to the slice for decompression.**
Reference #	IXA00371167
Description	If overflow occurs during stateless decompression and the last byte is a data byte from a non-empty stored block then overflow will not be reported but the output from the Hardware Accelerator will be incomplete.
Implication	The output from the Hardware Accelerator is incomplete due to a possible overflow.
Resolution	There is currently no workaround for this issue in B0. The issue has been resolved in C0.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.9 IXA00371315 - DC B0 Only - Static output can be corrupted, Dynamic can hang on very rare occasions. -

Title	**DC B0 Only - Static output can be corrupted, Dynamic can hang on very rare occasions.**
Reference #	IXA00371315
Description	On very rare occasions, a hang can occur during stateless and stateful dynamic compression. The output from stateless and stateful static compression can also be corrupted.
Implication	A hang occurs and compression output can be corrupted.
Resolution	There is currently no workaround for this issue. This issue is resolved with chipset C0 silicon.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.10 IXA00371601 - DC B0 Only : Slice hangs during decompression of fixed/dynamic blocks without empty stored blocks padding to a byte -

Title	**DC B0 Only : Slice hangs during decompression of fixed/dynamic blocks without empty stored blocks padding to a byte**
Reference #	IXA00371601
Description	The driver incorrectly handles fixed/dynamic blocks without an empty stored block in between padding each block out to a byte boundary.
Implication	The driver can hang if it encounters this type of input. Affects both stateless and stateful decompression.
Resolution	No known workaround. This issue is resolved with chipset C0 silicon.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.11 IXA00371624 - SRIOV: Device hang on adf_ctl up when guest defines a service not enabled in host services -

Title	SRIOV: Device hang on adf_ctl up when guest defines a service not enabled in host services
Reference #	IXA00371624
Description	If the host only enables service "cy0" and guest attempts to define and use "cy1" then device hangs affecting both host and guest.
Implication	Device hang.
Resolution	All services to be used by guests should be enabled in the host configuration file.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.12 IXA00372583 - DC : Decompression can falsely detect soft error with certain input. -

Title	DC : Decompression can falsely detect soft error with certain input.
Reference #	IXA00372583
Description	During decompression certain dynamic trees can cause the decompression slice to falsely detect soft error. This situation occurs when, during the compression of a data set, the compressor generates a dynamic Huffman tree with the following characteristics: i) Exactly 256 symbols are used. ii) All the symbols have the same, or nearly the same frequency of occurrence. iii) All of the symbols are encoded to 8-bit codes. In this particular situation, during decompression, the decompressor does not process the tree correctly and returns an error that usually indicates a data error.
Implication	Soft error is encountered and decompression of stream cannot be completed.
Resolution	See the Stateful Compression - Dealing with Error Code CPA_DC_BAD_LITLEN_CODES (-7) section in the Intel® Communications Chipset 89xx Series Software Programmer's Guide for more information
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.13 IXA00372698 - DC B0 Only - Decompression of stored blocks may cause CRC error -

Title	DC B0 Only - Decompression of stored blocks may cause CRC error
Reference #	IXA00372698
Description	During decompression of stored blocks an incorrect CRC/Adler may be returned.
Implication	During decompression of stored blocks an incorrect CRC/Adler may be returned.
Resolution	No workaround for this issue although CRC can be calculated in software.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.14 IXA00373407 - DC: Decompression service will continue to increment the consumed value for all data provided in the request after the final deflate block -

Title	**DC: Decompression service will continue to increment the consumed value for all data provided in the request after the final deflate block**
Reference #	IXA00373407
Description	The decompression service processes data up to and including a deflate block with a BFINAL bit set. If there is further data in the source buffer after the end of the BFINAL deflate block it will be not be processed. The behavior of the consumed byte counter (returned to the user via consumed value in CpaDcRqResults structure) cannot be predicted when the compression engine is fed with additional data after the end of deflate stream. Consecutive packets must be submitted as a Start of Packet. The acceleration slice does not understand Zlib or Gzip format therefore it is recommanded for the user not to submit the header nor the footer of the deflate stream to the compression slice.
Implication	If the source buffer provided to the decompression service contains data (e.g. a footer) after the BFINAL deflate block then the consumed value in CpaDcRqResults structure returned to the user is not correct. Subsequent packet sent to the slice will then fail to inflate.
Resolution	Addressed in the R1.5 Software Package
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.15 IXA00373795 - GEN: Segfault on exit of the performance sample user-space application with optimized wireless firmware -

Title	**GEN: Segfault on exit of the performance sample user-space application with optimized wireless firmware**
Reference #	IXA00373795
Description	When running the optimized wireless firmware, there is a segmentation fault on exit when operating in the following environment: - Wireless Firmware - 1024 byte packets or bigger - 4 instances (2 slices on 2 CPMs) - Traditional API
Implication	All tests run successfully, but there is a segmentation fault on exit, resulting in a memory leak.
Resolution	Addressed in the R1.0.1 Software Package
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.16 IXA00378322 - CY: Performance drop for alg chain cipher then hash when append flag set -

Title	**CY: Performance drop for alg chain cipher then hash when append flag set**
Reference #	IXA00378322
Description	There is performance drop for Cipher-Hash operation when the digest is generated and inserted in to the destination buffer. The performance drop is observed comparing to Cipher-Hash operations when digest is generated and added to the separate buffer.
Implication	The algorithm remains functional. However, performance throughput levels for encrypt are reduced by 30-40%. Decrypt is unaffected.

Title	CY: Performance drop for alg chain cipher then hash when append flag set
Resolution	Addressed in the R1.0.1 Software Package
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.17 IXA00378522 - GEN : Cannot have different values of LimitDevAccess across device config files -

Title	GEN : Cannot have different values of LimitDevAccess across device config files
Reference #	IXA00378522
Description	Problems arise when LimitDevAccess is enabled in one config file, but disabled in another. The driver is not currently able to handle this mixed configuration.
Implication	Driver may not find the correct entries in the configuration file if this setting is different for each device in the system.
Resolution	Addressed in the R1.2.0 Software Package
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.18 IXA00378662 - DC: The upper word of the Adler checksum can be calculated incorrectly on very rare occasions -

Title	DC: The upper word of the Adler checksum can be calculated incorrectly on very rare occasions
Reference #	IXA00378662
Description	During decompression, because Adler32 is computed on four bytes at a time, the Compression slice incorrectly calculates the upper word of the Adler checksum. This situation occurs under the following conditions: i) The upper word of the current Adler calculation is in the vicinity of 0xFFF0. ii) The lower word of the current Adler calculation is in the vicinity of 0x7FFC. iii) The Adler32 computation for the next four bytes causes the upper word to exceed 0x2ffe0 before modulo arithmetic is performed. The modulo arithmetic is not done correctly in this case, causing an error in the checksum.
Implication	An erroneous checksum error is encountered after decompression is completed. There is no ability for the decompression solution to examine a checksum and determine that the problem occurred.
Resolution	Addressed in the R1.0.1 Software Package. A software workaround now takes care of this under the API.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.19 IXA00378701 - GEN : When slub_debug is defined in kernel space, seg fault when service is shutdown -

Title	GEN : When slub_debug is defined in kernel space, seg fault when service is shutdown
Reference #	IXA00378701
Description	When slub_debug is enabled at boot using the kernel parameters, the driver's memory allocation function can return non aligned memory thus crashing the driver during ring creation.
Implication	Cannot use slub-debug as a kernel parameter.

Title	**GEN : When slub_debug is defined in kernel space, seg fault when service is shutdown**
Resolution	Do not enable slub_debug at boot. Enable the relevant slub_debug parameters in the /sys/kernel/slub/* files once the system is up-and-running.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.20 IXA00378934 - CY: Crypto RSA segmentation fault while running performance tests in user space -

Title	**CY: Crypto RSA segmentation fault while running performance tests in user space**
Reference #	IXA00378934
Description	when running multiple threads performing primeTests on the same logical instance the application may run into the following error [error] cpaCyPrimeTest() - : Cannot get mem pool entry Segmentation fault (core dumped)
Implication	Performance test fails.
Resolution	Addressed in the R1.0.1 Software Package
Affected OS	Fedora 16
Driver/Module	Feature - Acceleration Driver

4.2.21 IXA00378952 - GEN : Error over-riding a single logical Instance NumConcurrentRequest Value -

Title	**GEN : Error over-riding a single logical Instance NumConcurrentRequest Value**
Reference #	IXA00378952

Description	The V2 configuration file provides a general setting for the NumConcurrentRequests for each service:
	#Default values for number of concurrent requests*/
	CyNumConcurrentSymRequests = 512
	CyNumConcurrentAsymRequests = 64
	DcNumConcurrentRequests = 512
	The user should be able to override a single logical instance NumConcurrentRequest value using: CyXNumConcurrentAsymRequests = 512
	where X= the instance number.
	However, the adf_ctl command or icp_sal_userStartMultiProcess function will report an error that the other instances are missing a corresponding CyXNumConcurrentAsymRequests parameter, but they should just use the default defined in the general section.

Implication	The user cannot overwrite settings for a single instance.

Title	GEN : Error over-riding a single logical Instance NumConcurrentRequest Value

Resolution	If a user wants to overwrite a specific settings for an instance the entry for this needs to be put in a appropriate space in the instance definition. Configuration is processed from up - down and if the parser does not know about the specific configuration for a given instance, it will produce the default value. For example, when the user wants to overwrite the CyXNumConcurrentAsymRequests value for instance #1, it should be configured as above.
	# Crypto - User instance #1
	Cy1Name = "SSL1"
	Cy1NumConcurrentAsymRequests = 128
	Cy1IsPolled = 1
	Cy1AcceleratorNumber = 1
	# List of core affinities
	Cy1CoreAffinity = 1
	In the below configuration, the specific value will not be taken into account and the default value will be generated.
	# Crypto - User instance #1
	Cy1Name = "SSL1"
	Cy1IsPolled = 1
	Cy1AcceleratorNumber = 1
	# List of core affinities
	Cy1CoreAffinity = 1
	Cy1NumConcurrentAsymRequests = 128

Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.22 IXA00379409 - DC: cpaDcDecompressData reqs fail intermittantly on resubmit after CPA_STATUS_RETRY -

Title	DC: cpaDcDecompressData reqs fail intermittantly on resubmit after CPA_STATUS_RETRY
Reference #	IXA00379409
Description	If the user application resubmits both cpaDcCompressData or cpaDcDecompressData calls and if the driver returns status CPA_STATUS_RETRY, then: - The cpaDcCompressData call works fine on resubmits. - The cpaDcDecompressData call fails, not on every resubmit but just on some of them. It doesn't fail on the call itself but on the status returned within the results structure of the callback (the call is asynchronous). The error maybe any of the failure codes (-2, -13 etc).
Implication	The user application will see error codes (-2, -13 etc) and the data will not be decompressed.
Resolution	Addressed in the R1.1 Software Package
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.23 IXA00379630 - CY : NRBG cannot be used on 2nd crypto accelerator engine -

Title	CY : NRBG cannot be used on 2nd crypto accelerator engine
Reference #	IXA00379630
Description	The driver currently does not prevent the user from attempting to use NRBG functions on the second accelerator engine. There should be an error message to that effect, however, the driver behaviour in this situation is unpredictable.
Implication	Application may experience segmentation faults if it attempts to run NRBG on second accelerator engine.

Title	CY : NRBG cannot be used on 2nd crypto accelerator engine
Resolution	Only use Accelerators Engine number 0 on each accelerator. (or 0 and 2 when using V2 config file)
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.24 IXA00379858 - GEN : Heartbeat feature and error detection resets can cause GbE interfaces to go down -

Title	GEN : Heartbeat feature and error detection resets can cause GbE interfaces to go down
Reference #	IXA00379858
Description	When the acceleration driver detects either the firmware being unresponsive or other errors (detected through PCIe Advanced Error Reporting), it saves the PCIe state of the PCH and then performs a secondary bus reset of the Intel® Communications Chipset 89xx Series device. Once the reset is complete, the driver restores the PCIe state of the PCH. At this point, the GbE interfaces have been reset and there were no saving/restoring of its PCIe state and there is no notification of this reset to the igb driver (if it was loaded and running). The GbE interfaces are no longer operational. Any network connections previously opened are now unresponsive.
Implication	The GbE interfaces will hang and will be unresponsive.
Resolution	The "heartbeat" feature and GbE Watchdog are described in the Getting Started Guide and the Programmer's Guide.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.25 IXA00380162 - CY: Kernel Opps Running qat_service shutdown before gige_watchdog_service stop -

Title	CY: Kernel Opps Running qat_service shutdown before gige_watchdog_service stop
Reference #	IXA00380162
Description	If the qat_service service is shutdown before the gige_watchdog_service the kernel will report an opps.
Implication	The kernel will report a kernel oops.
Resolution	Addressed in the R1.1 Software Package
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.26 IXA00380177 - DC: Decompressing files greater than 4 GB can result in an unreported error -

Title	DC: Decompressing files greater than 4 GB can result in an unreported error
Reference #	IXA00380177
Description	Decompressing files greater than 4 GB can result in an unreported error.
Implication	The compression job fails.
Resolution	Addressed in the R1.2 Software Package
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

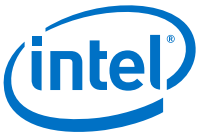

4.2.27 IXA00380411 - CY: Deadlock for Partial Packets in user space -

Title	CY: Deadlock for Partial Packets in user space
Reference #	IXA00380411
Description	When creating multiple sessions all feeding partial packet requests to one instance, a deadlock can occur.
Implication	The driver will lockup.
Resolution	Addressed in the R1.2 Software Package.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.28 IXA00380578 - SRIOV: iommu_domain_alloc() crashes if IOMMU not enabled in the BIOS -

Title	SRIOV: iommu_domain_alloc() crashes if IOMMU not enabled in the BIOS
Reference #	IXA00380578
Description	It has been noticed that if in the BIOS, the fields: "VT-d" and "SRIOV Support" are disabled, the iommu_domain_alloc() function fails. The console reports: Jan 15 08:45:44 localhost kernel: [148.892566] RIP [<ffffffff813b9029>] iommu_domain_alloc+0x3f/0x56 Jan 15 08:45:44 localhost kernel: [148.892680] RSP <ffff8802226d3e78> Jan 15 08:45:44 localhost kernel: [148.892750] CR2: 0000000000000000 Jan 15 08:45:44 localhost kernel: [148.892830] ---[end trace d612ff498c7f024b]--- Jan 15 08:45:45 localhost abrtd: Directory 'oops-2013-01-15-08:45:45-720-0' creation detected Jan 15 08:45:45 localhost abrt-dump-oops: Reported 1 kernel oopses to Abrt
Implication	iommu_domain_alloc() function fails and prevents the icp_qa_al from being insmoded correctly.
Resolution	Ensure that IOMMU is supported and enabled on the system.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.29 IXA00381020 - GEN: IA Driver - Osal - Error in osalIOMMUVirtToPhys function -

Title	GEN: IA Driver - Osal - Error in osalIOMMUVirtToPhys function
Reference #	IXA00381020
Description	While looking at the IOMMU support in the User Space side of the Osal Memory Driver: OsalUsrKrnProxy.c The function osalIOMMUVirtToPhys is called in the kernel space driver with an unmap call (DEV_MEM_IOC_IOMMUUNMAP) rather than a virt to phys call (DEV_MEM_IOC_IOMMUVTOP).
Implication	Memory deallocation will fail.
Resolution	Addressed in the R1.3 Software Package.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.30 IXA00381037 - DC: cpaDcRemoveSession return code check error. -

Title	DC: cpaDcRemoveSession return code check error.
Reference #	IXA00381037
Description	Issue 1: In the cpaDcRemoveSession() API, the spinlock was destroyed if the status to be return to the user was CPA_STATUS_RETRY. The spinlock should only be destroyed when the status is CPA_STATUS_SUCCESS. Issue 2: Before removing the session, it is necessary to verify the status of the rings to make sure these are empty. To do this, the API icp_adf_queueDataToSend() must be used. Originally, the verification was done using icp_adf_isRingEmpty() which was wrong.
Implication	Issue 1: If a retry occured then this could lead to timing problems. Issue 2: Ring may not be empty even so the driver thinks it is.
Resolution	Addressed in the R1.2 Software Package.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.31 IXA00381038 - DC: Potential memory leak in DC (Trad API) -

Title	DC: Potential memory leak in DC (Trad API)
Reference #	IXA00381038
Description	If a call to dcCreateRequest fails in dcCompDecompData, memory associated with the cookie is not freed.
Implication	This can lead to memory leaks when requests fail to be created.
Resolution	Addressed in the R1.2 Software Package.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.32 IXA00381173 - DC: Stateless Cumulative Checksum reports incorrect bytes consumed. -

Title	DC: Stateless Cumulative Checksum reports incorrect bytes consumed.
Reference #	IXA00381173
Description	The issue is present on stateless decompression and shows the consumed byte count being incorrect. The issue was seen with the following session settings: - Level 6 - Dynamic type - Deflate - AutoSelectBest enabled, - Adler32 checksum selected The test case used compressed a 512 block of data split into 2 scatter-gather lists. The compression operation worked fine but when trying to decompress the data, the expected consumed was different of what was produced during the compression operation. Switching the test to non-cumulative mode shows correct bytes consumed.
Implication	The decompressed data is invalid.

Title	DC: Stateless Cumulative Checksum reports incorrect bytes consumed.
Resolution	Addressed in the R1.2 Software Package
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.33 IXA00381337 - SRIOV : Pass-through of Intel® QuickAssist Technology PF and PCH GigE ports to a VM result in the GigE ports being non-responsive -

Title	SRIOV : Pass-through of Intel® QuickAssist Technology PF and PCH GigE ports to a VM result in the GigE ports being non-responsive
Reference #	IXA00381337
Description	Pass-through of Intel® QuickAssist Technology Physical Function (PF) and the PCH device GbE ports to a virtual machine result in GbE ports being nonresponsive after Intel® QuickAssist Technology device initialization. Intel® QuickAssist Technology Virtual Function (VF) pass-through with GbE port passthrough is not affected.
Implication	GigE ports are non-responsive.
Resolution	After Intel® QuickAssist Technology device initialization, removing the igb module and re-inserting it allows both devices to work as normal.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.34 IXA00381487 - GEN: Dynamic instance allocation fails for 32-bit app on 64-bit OS -

Title	GEN: Dynamic instance allocation fails for 32-bit app on 64-bit OS
Reference #	IXA00381487
Description	The exisiting solution had an incompatible data structure size for 32 bit user space application running on 64-bit kernel space.
Implication	Dynamic instance would have issues running on 32 bit user space with 64 bit kernel space.
Resolution	Addressed in the R1.3 Software Package
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.35 IXA00381488 - DC: Errors not reported if session type is COMBINED and compression is dynamic -

Title	DC: Errors not reported if session type is COMBINED and compression is dynamic
Reference #	IXA00381488
Description	Errors from translation slice are ignored when the session direction is CPA_DC_DIR_COMBINED and huffType is CPA_DC_HT_FULL_DYNAMIC
Implication	If user's application configures the session so that the direction is combined and the compression type is dynamic then the user will not see any potential errors that could occur.
Resolution	Addressed in the R1.3 Software Package.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.36 IXA00381962 - CY: Need Param check in crypto APIs for CpaBoolean variables passed as pointers -

Title	**CY: Need Param check in crypto APIs for CpaBoolean variables passed as pointers**
Reference #	IXA00381962
Description	All the cpaCy APIs that take a boolean parameter as a pointer need to have this parameter check before being used.
Implication	If the user decides to pass a null pointer to this parameter, the driver will report a segmentation fault.
Resolution	Parameter check added to verify that the pointer address being passed is not NULL.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.37 IXA00381968 - GEN: Typo in CONFIG_PCI(E)AER -

Title	**GEN: Typo in CONFIG_PCI(E)AER**
Reference #	IXA00381968
Description	The linux kernel has the flag CONFIG_PCIEAER set when Advanced Error Reporting is supported. Making a typo on this #define prevents the code within CONFIG_PCIEAER section not to be compiled.
Implication	Advanced Error Reporting feature is not available.
Resolution	Addressed in the R1.3 Software Package.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.38 IXA00382154 - GEN: Error in ring handling when using interrupts in user space -

Title	**GEN: Error in ring handling when using interrupts in user space**
Reference #	IXA00382154
Description	When using interrupt mode with a user space QA application, the response ring head is not always updated during response processing and so the interrupt condition persists. The root of the issue is that a previous SW optimization coalesces response ring head updates. This optimization should not be applied when the response ring is in interrupt mode.
Implication	The fact that interrupts are re-enabled causes a the same interrupt to fire again multiple times. This leads to significant drop in system performance.
Resolution	In the API adf_ring_ioc_create_handle(): The ring_handle.min_resps_per_head_write is now set 0 in interrupt mode. Thi sprevents the interrupt from firing again. The ring_handle.min_resps_per_head_write = ring_data->minRespsPerHeadWrite in polling mode
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

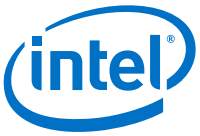

4.2.39 IXA00382531 - CY: more cleanup code in osal_init is needed -

Title	CY: more cleanup code in osal_init is needed
Reference #	IXA00382531
Description	If function calls with osal_init fail, then the driver will not be unregistered and the crypto interface may not be uninitialized.
Implication	This can lead to memory leaks in the kernel space.
Resolution	The driver now calls the API unregister_mem_device_driver() if the osalCryptoInterfaceInit() fails. If the API osalIOMMUInit() fails, the driver now calls both unregister_mem_device_driver() and osalCryptoInterfaceExit() APIs.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.40 IXA00382623 - DRAM Read in counter mode does not consume desc read signal correctly -

Title	DRAM Read in counter mode does not consume desc read signal correctly
Reference #	IXA00382623
Description	During a DRAM read operation of the scatter gather list, the decision to read the next flat buffer is made at the wrong time. This issue is visible when using dynamic compression and under a heavy load of traffic. Timing / IO latency all play a role when it comes to the manifestation of this issue. The problem shows when the compression firmware does not read the entire SGL. Dynamic compression or cases where the compression length in the request parameters is less than the length of the SRC flat buffer (inside the SGL) will be the main culprits. During DRAM read counter mode when not all of the flat buffers are populated with data we want to read we set the num of buffers in local memory to zero and the current buffer size to the dram read counter. The problem occurs because the decision to read the next flat buffer is made *before* the decision to modify these values. For correct operation it needs to happen after.
Implication	Consequences: - In the static compression scenario, we may end up with compressed data but not the ones that we expect - In the case of dynamic comnpression scenario, we may get slice error or history corruption.
Resolution	Firmware has been modified in 1.4 release so the decision to read the next flat buffer is made after the number of flat buffer is set to 0 in local memory.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.41 IXA00382998 - CY: Driver lockup with RC4 cipher when hit ring full -

Title	CY: Driver lockup with RC4 cipher when hit ring full
Reference #	IXA00382998
Description	When doing decryption with the RC4 cipher which has lots of threads submitting decrypt requests to a single instance, eventually one thread hits ring full (CPA_STATUS_RETRY). That thread then resubmits the same request but no response will be received.
Implication	For multi threaded operations, if the ring is already full and a re-submitted request receives a CPA_STATUS_RETRY then this request will not receive a response. The bug is in function LacSymQueue_RequestSend() in file lookaside/access_layer/src/common/crypto/sym/lac_sym_queue.c. Before putting the request on the ring, the function assigns: - pSessionDesc->nonBlockingOpsInProgress = CPA_FALSE; This indicates that a blocking operation is in flight. The function then attempts to write the request to the ring. If the write to the ring is not successful, then the function fails to restore the nonBlockingOpsInProgress flag. Subsequent requests are queued in the session's queue, waiting for a pending operation which does not exist.
Resolution	A check has been added in lac_sym_queue.c so that if the icp_adf_transPutMsg fails then the nonBlockingOpsInProgress is reset to CPA_TRUE to prevent the driver from waiting for a request to be process that has not been sent down to the firmware.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.42 IXA00382999 - CY: Hit ring full when number of threads << ring size -

Title	CY: Hit ring full when number of threads << ring size
Reference #	IXA00382999
Description	The ring reports full for an invalid reason.
Implication	The driver returns frequent CPA_STATUS_RETRY errors, requests are being resubmitted and performance decreases.
Resolution	The number of inflight messages is decremented just before the callback is invoked rather than updating the number of inflight messages after all the callback have been called.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

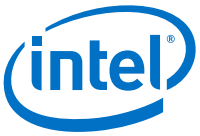

4.2.43 IXA00383390 - DC: High rates of simultaneous crypto and (de)compression operations may cause in correct compression output and write an incorrect amount of data into the output buffer -

Title	DC: High rates of simultaneous crypto and (de)compression operations may cause in correct compression output and write an incorrect amount of data into the output buffer
Reference #	IXA00383390
Description	An incorrect compression output may be generated during some compression / decompression operations. This soft error indicates that an incorrect amount of data was written to the (de)compression output buffer in memory. It may not write enough data to memory or it may write too much data to memory. The issue is most likely to occur when running both symmetric cryptographic operations and compression/decompression operations at the same time with high request/traffic rates. When the issue occurs, the cryptographic operations are unaffected. The issue is highly unlikely to occur when running only compression/decompression operations (i.e. without concurrent cryptographic operations). The issue will not occur when running cryptographic operations alone. In the scenario where not enough data is written to the output buffer, data is missing and hence the compressed/decompressed data is incomplete and hence invalid.
Implication	If too little data is written the compressed/decompressed data is incomplete and hence invalid. If too much data is written the compressed/decompressed data is valid but additional data is present in the output buffer. If the issue occurs, only Intel(r) QuickAssist accelerator compression/decompression operations are affected.
Resolution	Resubmit the failed Intel(r) QuickAssist accelerator compression/decompression request. This issue will be resolved in a future release of the acceleration driver.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.44 IXA00383454 - CY: Performance Sample Code digestAppend setting is not optimal -

Title	CY: Performance Sample Code digestAppend setting is not optimal
Reference #	IXA00383454
Description	The performance sample code for the symmetric and symmetric data plan code is setting the digestAppend flag to true. This does not yeild the best performance results. In general, the difference is small. However, for 2048 bytes buffers and greater the digestAppend=true follows an unoptimized data path and therefore, the impact is more significant. Files affected: cpa_sample_code_sym_perf.c cpa_sample_code_sum_perf_dp.c
Implication	Symmetric AlgChain performance does not maximize the silicon potential.
Resolution	The cpa_sample_code_sym_perf.c and cpa_sample_code_sum_perf_dp.c files contain the following functions: setupAlgChainTest and setupAlgChainDpTest. These two functions call setupSymmetricTest and setupSymmetricDpTest. The last parameter is the digestAppend flag, this needs to be set to CPA_FALSE. Note: Changing this may result in some alg-chain combinations having invalid parameters.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.45 IXA00383572 - DC: Report overflow from XLT when byte count mismatch detected -

Title	DC: Report overflow from XLT when byte count mismatch detected
Reference #	IXA00383572
Description	During a dynamic compression operation, the translator slice will not report an overflow error if the dynamic output is greater than the static output.
Implication	Compressed data returned to the user will be correct. The user will not be aware that XLT has overflowed if the dynamic output is greater than the static output after the re-submit. This has no effect on the user application.
Resolution	This issue is resolved in R1.5.0
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.46 IXA00384087 - GEN: icp_adf_check_device() API fails to detect when firmware hang. -

Title	GEN: icp_adf_check_device() API fails to detect when firmware hang.
Reference #	IXA00384087
Description	Under certain circumstances, an ME hang due to the read() function called from the icp_adf_check_device() API may go undetected. The function checks for the return value however, it is also required to verify the content of the file.
Implication	if the read function does not fail, it may be able to read /proc/icp_dh89xxcc_dev0/qat0 that could be zero length and therefore it would not catch a firmware hang issue.
Resolution	To resolve this issue, we now read the content of the file /proc/icp_dh89xxcc_dev0/qat0 to make sure that the content is valid.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.47 IXA00384581 - GEN: The “NumberConcurrent” options in the configuration file will be overwritten to be half of the requested value. -

Title	GEN: The “NumberConcurrent” options in the configuration file will be overwritten to be half of the requested value.
Reference #	IXA00384581
Description	The configuration file holds the number of concurrent requests for Sym, Asym and DC. The Number of Concurrent request is user configurable, but this user value is divided by 2.
Implication	Mismatch between the number of concurrent requests defined by the user in the configuration file and the number of concurrent requests used by the driver.
Resolution	This issue is resolved in the R1.5.0.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

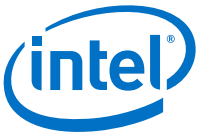

4.2.48 IXA00384651 - CY: When ICP_WITHOUT_THREAD is defined, enabling poll mode will cause a core dump -

Title	CY: When ICP_WITHOUT_THREAD is defined, enabling poll mode will cause a core dump
Reference #	IXA00384651
Description	If defining ICP_WITHOUT_THREAD and setting the instance poll mode to be 0, then a QAT Driver core dump occurs.
Implication	The QAT Driver will core dump.
Resolution	This issue is resolved in R1.5.0
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.49 IXA00384930 - CY: Issue with data plane GCM operations when using the acceleration driver -

Title	CY: Issue with data plane GCM operations when using the acceleration driver
Reference #	IXA00384930
Description	According to API documentation, values for hashStartSrcOffsetInBytes and messageLenToHashInBytes are not required and are intended for internal purposes for GCM operations. If these values are not set, GCM operations do not work with the data plane APIs. The driver is responsible for setting these values and this is not currently being done. This issue does not occur with the traditional GCM operations.
Implication	Crypto operations fail.
Resolution	This issue is resolved in R1.5.0
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.50 IXA00384933 - GEN: Unnecessary extra interrupts generated in user mode -

Title	GEN: Unnecessary extra interrupts generated in user mode
Reference #	IXA00384933
Description	In user space, if configured to use interrupts, more interrupts are generated than are necessary. In /proc/interrupts, the number of interrupts shown is greater than the number of responses received. (Corresponds to IXA00384677 on QAT1.6)
Implication	Unnecessary interrupt processing wastes CPU cycles.
Resolution	Addressed in the R1.5 Software Package
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.51 IXA00384934 - CY: Silent drop of messages -

Title	CY: Silent drop of messages
Reference #	IXA00384934
Description	When Crypto partials are in flight, messages can get queued in the driver so they are processed sequentially, that is, the next request is not put on the ring to the firmware until the response from the previous request has been received. If the response ring has been polled 10,000 times and the response is not received, the queued requests are silently dropped. (Corresponds to IXA00384681 on QAT1.6)
Implication	An application could remain waiting indefinitely for a response and not be aware that the request has been dropped.
Resolution	This issue is resolved in R1.5.0
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.52 IXA00385456 - GEN: Response ring processing is unnecessarily restricting request submissions -

Title	GEN: Response ring processing is unnecessarily restricting request submissions
Reference #	IXA00385456
Description	The inflight count for ring messages is not updated until after all callbacks handled together have been completed. This can lead to the ring pair reporting full, even though there is space available on the rings. The inflight counter should be decremented as soon as the response msg has been picked up in the response ring rather than after a burst of callbacks have been completed.
Implication	A RETRY response can be received when trying to put a message in the ring even though space is available.
Resolution	This issue is resolved in R1.5.0
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.53 IXA00385555 - GEN: The driver does not completely enable all error correction and detection (ECC and Parity) in the accelerator -

Title	GEN: The driver does not completely enable all error correction and detection (ECC and Parity) in the accelerator
Reference #	IXA00385555
Description	All previously released Intel® QuickAssist Technology drivers do not have some of the ECC error correction and some of the parity detection logic enabled in the accelerator.
Implication	In the logic that does not have ECC error correction enabled, a single bit error will be treated as an uncorrectable error. For the limited memory that has parity detection disabled, a parity error will NOT be treated as uncorrectable. Uncorrectable errors may cause the accelerator to halt.
Resolution	This issue is resolved in R1.5.0
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

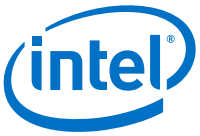

4.2.54 IXA00385634 - DC: Static decompression can falsely detect soft error with certain input. -

Title	DC: Static decompression can falsely detect soft error with certain input.
Reference #	IXA00385634
Description	It was noticed that the fix for IXA00372583 was not enabled if the session type was set to CPA_DC_HT_STATIC. The fix for this IXA was verifying if the session was set to Dynamic. This was wrong as a the fix for IXA00372583 needs to be applied in all types of sessions (dynamic /static).
Implication	As a consequence it could have happened that the IXA00372583 workaround would not be applied. The user would have seen a (-7) error code reporting a soft error.
Resolution	This issue is resolved in R1.5.0
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.55 IXA00385765 - GEN: Heartbeat test fails - platform does not recover and requires a reset -

Title	GEN: Heartbeat test fails - platform does not recover and requires a reset
Reference #	IXA00385765
Description	Under a heavy CPU load, it has been observed that when the driver is compiled with ICP_HEARTBEAT set, the Acceleration Engines may not restart when the firmware hangs. We've observed that the Acceleration Engines are stopped (as they should be) but won't be restarted.
Implication	As a consequence, the DH89xxCC device will halt as the firmware will not be reloaded
Resolution	This issue is resolved in R1.5.0
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.56 IXA00385873 - GEN: free_page usage issues -

Title	GEN: free_page usage issues
Reference #	IXA00385873
Description	OSAL API userMemFreePage() uses free_page() while userMemFreeAllPagePid() and userMemFreeAllPagePid() use __free_page. __free_page() and free_page() are defined as follow: #define __free_page(page) __free_pages((page), 0) #define free_page(addr) free_pages((addr),0) There is also some inconsistency in the parameter being passed to the free_page() calls.
Implication	No known effect to date.
Resolution	This is resolved with the R1.5 release.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.57 IXA00386021 - GEN: Higher than expected CPU cycles used in some low-traffic cases -

Title	GEN: Higher than expected CPU cycles used in some low-traffic cases
Reference #	IXA00386021
Description	This applies to CY / DC acceleration in user-space using polled rings.In response message handling, the CSR write coalescing logic results in the ring head CSR being written more frequently than necessary. This occurs in some low-traffic cases and results in an increased CPU cycle count for those cases. In high traffic cases, many responses are processed by the same poll and the ring head CSR writes are coalesced. In the case where only a small number of responses are processed by a poll, the head CSR is written in each poll, as the CSR write coalescing doesn't kick in.
Implication	More CPU cycles are used.
Resolution	This is resolved with the R1.5 release.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

4.2.58 IXA00386090 - GEN: QAT R1.3.7 driver cause Linux kernel crash -

Title	GEN: QAT R1.3.7 driver cause Linux kernel crash
Reference #	IXA00386090
Description	After powering on/off the system and run performance test, it is possible to see a Linux kernel crash.If the response to a administration message sent from the QAT driver (using QatCtrl_SendAdminMsg()) takes longer that 300 Jiffies to respond then the semaphore associated with the message is deleted. The response arriving after 300 Jiffies will try to modify the semaphore and cause a kernel crash.Here is QAT call flow,1.icp_adf_check_device->PROC file system/proc/icp_dh89xxcc_dev/??/qat0, then QatCtrl_Debug->QatCtrl_FWCountGet->QatCtrl_SendAdminMsg->QatCtrl_AdminMsgSendSync2.QatCtrl_AdminMsgSendSync put msg in RING, then call LacSync_WaitForCallback to wait for response. If after 300 jiffies, No response, then it will remove the pSyncCallbackCookie->sid3.The call back function QatCtrl_AdminSyncCb() will POST pSyncCallbackCookie->sid
Implication	Linux kernel will crash and will require a reboot
Resolution	This is resolved with the R1.5 release.
Affected OS	Linux
Driver/Module	Feature - Acceleration Driver

§ §

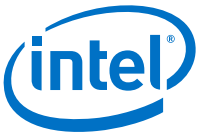

5.0 Frequently Asked Questions

The following codes that prefix each title identify which platform each FAQ applies to:

- A - Intel® Communications Chipset 8900 to 8920 Series
- B - Intel® Communications Chipset 8925 to 8955 Series

5.1 [A, B] I am seeing PCIe Bus Errors when executing the sample code (cpa_sample_code).

These errors could be caused by one of the following:

- incorrect PCIe De-emphasis setting.
  Use -3.5dB for trace lengths < 11” and -6dB for trace lengths > 11”. The Root Complex that the EndPoint is connected to needs to request the -3.5 de-emphasis when training in Gen2.
- Max Payload Size mismatch between the Root Port and each EndPoint.

One way to test this is to transmit traffic and see if this works. If this does, then perform loopback test. If this fails, then one of the above is the likely source.

5.2 [A] I am using Intel® Communications Chipset 8900 to 8920 Series (PCH) SKU2, but the acceleration service does not start correctly. How do I resolve this?

Verify that your configuration file does not declare an index for `AcceleratorNumber` or `ExecutionEngine` that is greater than that which your device supports. If you are running the sample code, copy `dh89xxcc_qa_dev0_single_accel.conf` from `quickassist/config` to `/etc/dh89xxcc_qa_dev0.conf`, restart the acceleration service, and try again.

5.3 [A, B] I have an application called XYZ with the intent to use two cryptography instances from each of two chipset (PCH) devices in the system (a total of four instances). What would the configuration files look like?

In this case, the `NumberCyInstances` parameter should be set to 2 in the configuration file for each PCH device.

5.4 [A, B] Should the Cy<n>Name parameter use unique values for <n> in each configuration file?

The `Cy<n>Name` parameter can be used in different configuration files without issue. In addition, the same `Cy<n>Name` name can be used in different domains within the same configuration file. The same rules apply to the `Dc<n>Name` parameter.

5.5 [A, B] Since the SSL data and the KERNEL sections in the configuration files for two Intel® Communications Chipset 8925 to 8955 Series (PCH) devices are identical, it is unclear how an application is able to use instances from more than one device. How does the application know which device each instance maps to?

The application can use the `cpaCyInstanceGetInfo2()` function to query which physical device an instance handle belongs to. For any application domain defined in the configuration files ([KERNEL], [SSL] and so on), a call to `cpaCyGetNumInstances()` returns the number of instances defined for that domain across all configuration files. A subsequent call to cpaCyGetInstances() can be used to obtain the instance handles. The `cpaCyInstanceGetInfo2()` function can then be used with an instance handle to identify which device a given instance maps to.

5.6 [A, B] Given the configuration below, what do the cpaCyGetNumInstances() and cpaCyGetInstances() functions return for each application and kernel domain?

Given this configuration:
- Application ABC: defines two crypto instances on dev0 only
- Application DEF: defines two crypto instances on dev1 only
- Application XYZ: defines four crypto instances (two on dev0, two on dev1)
- KERNEL: defines eight ctypto instances (four on dev0, four on dev1)

The `cpaCyGetNumInstances()` function returns the following:
- Application ABC: 2
- Application DEF: 2
- Application XYZ: 4
- KERNEL application domain: 8

The `cpaCyGetInstances()` function returns the following:
- Application ABC: two handles on dev0
- Application DEF: two handles on dev1
- Application XYZ: four handles (two on dev0 and two on dev1 in the order given)
- KERNEL application domain: 8 handles (four on dev0, four on dev1 in the order given)

Note: The order in the last two permutations is important.

5.7 [A, B] How can an application use instances from more than one Intel® Communications Chipset 8925 to 8955 Series (PCH) device when the [SSL] and [KERNEL] sections in both configuration files provided in the software package are identical?

An application must call `cpaCyInstanceGetInfo2()` on each handle returned from cpaCyGetHandles() and sort them by device. The `cpaCyInstanceGetInfo2()` function allows the user to query which physical device an instance handle belongs to. To expand somewhat, for any application domain defined in the configuration files ([KERNEL], [SSL] and so on), a call to `cpaCyGetNumInstances()` returns the number of instances defined for that domain across all configuration files, a subsequent call to `cpaCyGetInstances()` would obtain these instance handles.

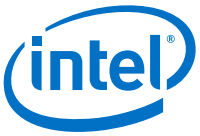

5.8 [A, B] Driver compiles correctly, but acceleration service fails to start. How do I fix this?

A typical reason why the acceleration service does not start is that the `intel_iommu=off` kernel boot option was not specified, as required and documented in the *Getting Started Guide*.

The following errors are seen in this scenario:

```
[error] QatCtrl_InitMsgSendSync() - : Callback timed out
[error] QatCtrl_SendInitMsg() - : Failed to send Init msg 0 to AE 0
[error] SalCtrl_QatStart() - : Sending SET_AE_INFO msg Failed

[error] SalCtrl_QatEventStart() - : Failed to start all qat instances
icp_qa_al err: adf_subsystemStart: Failed to start subservice QAT
icp_qa_al err: adf_do_init: adf_subsystemInit error, stopping and shutting down
```

5.9 [A] The firmware does not load. How can I fix this?

If the firmware does not load, verify that `udev` is available and running, and verify that the kernel was built with `CONFIG_FW_LOADER=y`.

5.10 [A, B] When I try to start the driver, I see errors (including kernel messages) that appear to be related to memory allocation. What can I do to avoid this?

When many instances are declared in the configuration file, it is possible to see these errors. The errors can typically be avoided by using the recommendations in the "Reducing Asymmetric Service Memory Usage" section of the *Intel® QuickAssist Technology Performance Optimization Guide*, by reducing the `NumConcurrentSymRequests` parameters in the configuration file, or by reducing the number of instances declared in the configuration file (see the "Acceleration Driver Configuration File" chapter in the chipset Programmer's Guide).

Another approach is to modify Linux* such that the value in `/proc/sys/vm/max_map_count` is increased (for example, to double the value). That value can be increased by modifying `/etc/sysctl.conf` to include the following line:

```
vm.max_map_count = <large_number_here>
```

Then reboot, and run `cat /proc/sys/vm/max_map_count` to verify that the value has been increased.