



Intel® QuickAssist Technology (Intel® QAT) Software for Linux*

Programmer's Guide – Customer Enabling Release

Revision 029

December 2023



Legal Notices & Disclaimers

Performance varies by use, configuration and other factors. Learn more on Intel's [Performance Index site](#).

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

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

Code names are used by Intel to identify products, technologies, or services that are in development and not publicly available. These are not "commercial" names and not intended to function as trademarks.

See [Intel's Legal Notices and Disclaimers](#).

© Intel Corporation. Intel, the Intel logo, Atom, Xeon, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

Contents

1	Introduction.....	14
	1.1 Terminology.....	14
	1.2 Typographical Conventions.....	17
2	Software Overview.....	18
	2.1 Intel® Communications Chipset 8925 to 8955 Series Compatibility.....	18
	2.2 Logical Instances.....	18
	2.2.1 Response Processing.....	18
	2.2.1.1 Interrupt Mode.....	18
	2.2.1.2 Epoll Mode.....	20
3	Acceleration Drivers Overview.....	22
	3.1 Hardware/Software Overview.....	22
	3.2 Acceleration Driver Configuration File.....	24
	3.3 Utility for Loading Configuration Files and Sending Events to the Driver - <code>adf_ctl</code>	24
	3.3.1 Usage.....	24
	3.3.2 Examples.....	24
	3.4 Application Payload Memory Allocation.....	25
	3.4.1 Thread Specific USDM.....	25
	3.5 User Space Additional Functions.....	26
	3.6 Managing Intel® QAT Endpoints Using <code>qat_service</code>	27
	3.7 Overview of QAT <i>debugfs</i> entries.....	28
	3.7.1 Entries in <code>/sys/kernel/debug/qat_*</code>	28
	3.7.2 Memory driver queries (<code>qae_mem_slabs</code>).....	28
	3.8 Compression Status Codes.....	29
	3.8.1 Intel® QAT Compression API Errors.....	29
	3.9 Stateful Compression Unsupported.....	33
	3.10 Stateless Compression Level Details.....	33
	3.10.1 Compression Level Mapping.....	33
	3.10.2 Limitation on History Buffer Size (aka Deflate Window Size).....	34
	3.11 Acceleration Driver Return Codes.....	35
	3.12 Batch and Pack Compression Unsupported.....	36
	3.13 Compress and Verify Feature.....	36
	3.14 Running Applications as Non-Root User.....	37
	3.15 Random Number Generation.....	38
	3.16 Huge Pages with the Included Memory Driver.....	38
	3.17 Heartbeat.....	39
	3.17.1 Heartbeat Operation.....	39
	3.17.1.1 Initialization.....	39
	3.17.1.2 Heartbeat Monitoring.....	39
	3.17.1.3 Resetting a Failed Device.....	39
	3.17.2 Incorporating Heartbeat into Intel® QAT Applications.....	40
	3.17.3 Testing Heartbeat.....	42
	3.17.3.1 Simulated Heartbeat Failure Configuration.....	43
	3.17.3.2 Simulating Heartbeat Failure.....	43

	3.17.3.3	System Virtual Files	43
	3.17.3.4	Heartbeat Polling Frequencies.....	43
3.18		Handling Device Failures in a Virtualized Environment	44
	3.18.1	Understanding System Messages and Warnings	45
3.19		Incorporating Dummy Responses into an Intel® QAT Application.....	45
	3.19.1	Reliability, Availability, Serviceability	46
	3.19.2	End to End Data Integrity Support in QAT 1.8:	47
3.20		Rate Limiting	47
	3.20.1	Service Level Agreement (SLA).....	48
	3.20.2	SLA Units.....	48
	3.20.3	SLA Manager Application.....	48
	3.20.3.1	Rate Limiting Commands.....	48
3.21		DU Manager Application	49
	3.21.1	Commands to Fetch Device Utilization.....	50
	3.21.2	Durations.....	50
	3.21.3	Reference Algorithm.....	50
3.22		Cipher-CRC.....	50
3.23		Access to Legacy Algorithms	51
4		Acceleration Driver Configuration File	55
4.1		Configuration File Overview	55
4.2		General Section.....	56
	4.2.1	General Parameters.....	56
4.3		Logical Instances Section	58
	4.3.1	[KERNEL] Section.....	59
	4.3.1.1	Enabling Linux* Kernel Crypto Framework (LKCF)	60
	4.3.2	[KERNEL_QAT] Section	60
	4.3.3	User Process [xxxx] Sections.....	61
	4.3.3.1	Maximum Number of Process Calculations.....	62
	4.3.3.2	Increasing the Maximum Number of Processes/Instances.....	62
	4.3.3.3	Configuring Instances for Virtual Functions.....	63
	4.3.4	Cryptographic Logical Instance Parameters	65
	4.3.4.1	LKCF-supported algorithms:	65
	4.3.5	Data Compression Logical Instance Parameters.....	66
	4.3.6	Setting the Core Affinity Parameter for a Logical Instance	66
4.4		Configuring Multiple Intel® QAT Endpoints in a System	67
4.5		Configuring Multiple Processes on a System with Multiple Intel® QAT Endpoints.....	68
4.6		Sample Configuration File	71
5		Secure Architecture Considerations.....	72
5.1		Terminology.....	72
	5.1.1	Threat Categories	72
	5.1.2	Attack Mechanism	73
	5.1.3	Attacker Privilege	73
	5.1.4	Deployment Models	74
5.2		Threat/Attack Vectors.....	74
	5.2.1	General Mitigation.....	74
	5.2.2	General Threats.....	75
	5.2.2.1	DMA.....	75
	5.2.2.2	Intentional Modification of IA Driver	75

	5.2.2.3	Modification of the QAT Configuration File.....	76
	5.2.2.4	Malicious Application Code.....	76
	5.2.2.5	Denial of Service.....	76
	5.2.3	Threats Specific to Cryptographic Service.....	77
	5.2.3.1	Reading Cryptographic Keys.....	77
6		Supported APIs.....	78
6.1		Intel® QAT APIs.....	78
6.1.1		Intel® QAT API Limitations.....	78
	6.1.1.1	Resubmitting After Getting an Overflow Error.....	80
	6.1.1.2	Dynamic Compression for Data Compression Service.....	82
	6.1.1.3	Maximal Expansion with Auto Select Best Feature for Compression.....	83
	6.1.1.4	Maximal Expansion and Destination Buffer Size in Compression Direction.....	84
6.1.2		Data Plane APIs Overview.....	85
	6.1.2.1	IA Cycle Count Reduction When Using Data Plane APIs.....	85
	6.1.2.2	Usage Constraints on the Data Plane APIs.....	87
	6.1.2.3	Cryptographic and Data Compression API Descriptions.....	87
6.1.3		Recovering from a Compress and Verify Error.....	87
6.1.4		Counting Recovered Compression Errors.....	89
6.1.5		Compress and Verify Error log in Sysfs:.....	89
6.1.6		Supported Algorithms in LKCF.....	89
6.2		Additional APIs.....	90
6.2.1		IOMMU Remapping Functions.....	90
	6.2.1.1	icp_sal_iommu_get_remap_size.....	90
	6.2.1.2	icp_sal_iommu_map.....	91
	6.2.1.3	icp_sal_iommu_unmap.....	91
	6.2.1.4	IOMMU Remapping Function Usage.....	92
6.2.2		Polling Functions.....	92
	6.2.2.1	icp_sal_pollBank.....	93
	6.2.2.2	icp_sal_pollAllBanks.....	93
	6.2.2.3	icp_sal_CyPollInstance.....	94
	6.2.2.4	icp_sal_DcPollInstance.....	95
	6.2.2.5	icp_sal_CyPollDplInstance.....	95
	6.2.2.6	icp_sal_DcPollDplInstance.....	96
6.2.3		User Space Access Configuration Functions.....	97
	6.2.3.1	icp_sal_userStart.....	97
	6.2.3.2	icp_sal_userStop.....	98
6.2.4		Version Information Function.....	99
	6.2.4.1	icp_sal_getDevVersionInfo.....	99
6.2.5		Reset Device Function.....	100
	6.2.5.1	icp_sal_reset_device.....	100
6.2.6		Thread-Less APIs.....	100
	6.2.6.1	icp_sal_poll_device_events.....	100
	6.2.6.2	icp_sal_find_new_devices.....	101
6.2.7		Compress and Verify (CnV) Related APIs.....	101
	6.2.7.1	icp_sal_dc_get_dc_error().....	101
	6.2.7.2	icp_sal_dc_simulate_error().....	102
6.2.8		Heartbeat APIs.....	102
	6.2.8.1	icp_sal_check_device().....	103

	6.2.8.2	icp_sal_check_all_devices()	103
	6.2.8.3	icp_sal_heartbeat_simulate_failure()	104
6.2.9		Device Polling APIs	104
	6.2.9.1	icp_sal_poll_device_events()	104
	6.2.9.2	cpaCylInstanceSetNotificationCb	105
	6.2.9.3	cpaDclInstanceSetNotificationCb	105
6.2.10		Congestion Management APIs	106
	6.2.10.1	icp_sal_SymGetInflightRequests	107
	6.2.10.2	icp_sal_AsymGetInflightRequests	107
	6.2.10.3	icp_sal_dp_SymGetInflightRequests	108
6.2.11		Service Specific Polling APIs	109
	6.2.11.1	icp_sal_CyPollSymRing	109
	6.2.11.2	icp_sal_CyPollAsymRing	109
6.2.12		Check Device Availability APIs	110
	6.2.12.1	icp_sal_userIsQatAvailable	110
7		Application Usage Guidelines	112
7.1		Mapping Service Instances to Engines on the Intel® QAT Endpoint	112
	7.1.1	Processor and Intel® QAT Endpoint Communication	112
	7.1.2	Service Instances and Interaction with the Hardware	112
	7.1.3	Service Instance Configuration	113
	7.1.4	Cryptographic Load Balancing Using Multiple Intel® QAT Instances	113
7.2		Cryptography Applications	114
	7.2.1	IPsec and SSL VPNs	114
	7.2.2	Encrypted Storage	114
	7.2.3	Web Proxy Appliances	115
7.3		Data Compression Applications	116
	7.3.1	Compression for Storage	116
	7.3.2	Data Deduplication and WAN Acceleration	116
8		Black Box Debug Tool	118
8.1		Introduction	118
	8.1.1	Overview	118
		8.1.1.1 Security Considerations	118
		8.1.1.2 Performance Considerations	118
8.2		Detailed Description	119
	8.2.1	Collection Data	119
		8.2.1.1 Data Synchronization	119
		8.2.1.2 Handling QAT Error Events	121
	8.2.2	Post-Processing	121
		8.2.2.1 Physical Addresses Audit	122
		8.2.2.2 Cipher Lengths Audit	122
		8.2.2.3 Return Codes Audit	123
		8.2.2.4 Listing Collected Data in 'Human Readable Form'	123
8.3		Installation	125
	8.3.1	Hardware and Software Compatibility	125
	8.3.2	Installing the Driver	126
	8.3.3	Compiling and Executing Performance Sample Code	126
	8.3.4	Uninstalling the Driver	126
8.4		Configuration	127

8.4.1	Configuration via QAT Device Configuration Files	127
8.4.2	Configuration via sysfs	128
8.4.3	Checking Current Configuration Used by Driver	129
8.5	Usage Examples.....	131
8.5.1	Collecting Data – Sanity Check	131
8.5.1.1	Continuous Sync Enabled	131
8.5.1.2	Continuous Sync Disabled	131
8.5.2	Audit Physical Addresses – Sanity Check.....	132
8.5.2.1	Emulate Uncorrectable Error	132
8.5.2.2	Continuous Sync Enabled	133
8.5.2.3	Continuous Sync Disabled (Crash Dump Based).....	135
8.5.3	Audit Cipher Buffers Alignment – Sanity Check.....	138
8.5.3.1	Emulate Slice Hang Caused by Incorrect Buffers Alignments.....	138
8.5.3.2	Slice Hang Handling with Continuous Sync Enabled.....	139
8.5.3.3	Slice Hang Handling with Continuous Sync Disabled.....	141
8.5.4	Audit Return Codes	144
8.5.4.1	Audit Return Codes – Continuous Sync Option.....	144
8.6	SR-IOV	146
8.6.1	Build instructions	146
8.6.2	Usage.....	148
8.7	Programming Guide	148
8.7.1	Physical to Virtual Translation Callback.....	148

Figures

Figure 1.	Kernel Space Response Ring Processing.....	19
Figure 2.	Intel® C62x Chipset (PCH) Acceleration Endpoint Configuration 1.....	22
Figure 3.	Intel® C62x Chipset (PCH) Acceleration Endpoint Configuration 2	23
Figure 4.	Incorporating Dummy Responses in an Intel® QAT Operation	46
Figure 5.	Dynamic Compression Data Path.....	83
Figure 6.	Amortizing the Cost of an MMIO Across Multiple Requests	86
Figure 7.	Service Instance Configuration	113
Figure 8.	Data Collection Architecture	120
Figure 9.	Typical Crash Dump Scenario.....	121

Tables

Table 1.	Terminology.....	14
Table 2.	Reference Documents and Resources	17
Table 3.	Services	25
Table 4.	Intel® QuickAssist Technology /sys/kernel/debug Entries	28
Table 5.	<i>qae_mem_slabs</i> Commands Supported.....	29
Table 6.	Intel® QAT Compression API Errors	30
Table 7.	Compression Levels for QAT 1.7 Hardware.....	33
Table 8.	Compression Levels for QAT 1.8 Hardware.....	34
Table 9.	Acceleration Driver Return Codes.....	35
Table 10.	Acceleration Driver Return Codes for Linux* Device Driver Operations.....	35
Table 11.	AutoResetOnError Values.....	40
Table 12.	Heartbeat System Virtual Files	43
Table 13.	Supported Legacy Algorithms.....	52



Table 14.	General Default Configuration Parameters	56
Table 15.	General Parameters.....	57
Table 16.	[KERNEL] Section Parameters	60
Table 17.	[KERNEL_QAT] Section Parameters.....	60
Table 18.	[KERNEL_QAT] Section Parameters.....	61
Table 19.	Configuring Physical Functions and Virtual Functions	64
Table 20.	Cryptographic Logical Instance Parameters.....	65
Table 21.	Data Compression Logical Instance Parameters	66
Table 22.	System Threat Categories.....	72
Table 23.	Attack Mechanisms and Examples.....	73
Table 24.	Attacker Privilege	73
Table 25.	Deployment Models	74
Table 26.	Compression/Decompression Overflow Behavior.....	81
Table 27.	API Support for Compress and Verify and Recover.....	88

Revision History

Document Number	Revision Number	Description	Revision Date
336210	029	<p>Updates for Intel® QAT Programmers Guide – Customer Enabling Release:</p> <ul style="list-style-type: none"> Updated Thread Specific USDM Added Enabling Linux* Kernel Crypto Framework (LKCF) Added 3.18.1: Understanding System Messages and Warnings Updated Section 3.10.1: Compression Level Mapping Added note about enabling Black Box Debug Tool (BBDT) to Virtual Functions 	December 2023
336210	028	<p>Updates for Intel® QAT Programmers Guide – Customer Enabling Release:</p> <ul style="list-style-type: none"> Updated Section 6.1.1: Intel® QAT API Limitations with updated Guidance on. Auto-Select-Best feature (ASB, i.e. <code>CPA_DC_ASB_ENABLED</code>) Updated Table II: Access to Legacy Algorithms with Opt-In PKE Algorithms Updated Section 3.21: DU Manager Application for clarity 	May 2023
336210	027	<p>Updates for Intel® QAT Programmers Guide – Customer Enabling Release, with Release v4.21:</p> <ul style="list-style-type: none"> Added 3.4.1: Thread Specific USDM 	March 2023
336210	026	<p>Updates for Intel® QAT Programmers Guide – Customer Enabling Release:</p> <ul style="list-style-type: none"> Updated Section 3.7: Overview of QAT debugfs entries to include new Section 3.7.2: Memory driver queries (<code>qae_mem_slabs</code>) Updated Legal Notices & Disclaimers 	February 2023
336210	025	<p>Updates for Intel® QAT Programmers Guide – Customer Enabling Release:</p> <ul style="list-style-type: none"> Name change, now supports 1.8 HW Gen lookaside features (non-inline) 	December 2022
336210	024	<p>Updates for Intel® QAT Programmers Guide Hardware Version 1.7:</p> <ul style="list-style-type: none"> Added Section 3.22 Access to Legacy Algorithms 	September 2022

Document Number	Revision Number	Description	Revision Date
		<ul style="list-style-type: none"> Removed Section 6.2.1 Dynamic Instance Allocation Functions, as it is unsupported 	
336210	023	Updates for Intel® QAT Programmers Guide Hardware Version 1.7: <ul style="list-style-type: none"> Updated Chapter 8: Black Box Debug Tool, changed numeration, added SR-IOV section, update report tool behavior Updated Sections 7.2.2: Encrypted Storage, & 7.3: Data Compression Applications with minor edits to terminology and grammatical updates for clarity 	April 2022
336210	022	Updates for Intel® QAT Programmers Guide Hardware Version 1.7: <ul style="list-style-type: none"> Updated Section 2.2.1.3: changed first note on epoll mode to improve clarity. Updated Section 3.16 Huge Pages with the Included Memory Driver Updated with IntelOne font 	March 2022
336210	021	Updates for Intel® QAT Programmers Guide Hardware Version 1.7: <ul style="list-style-type: none"> Updated Table 1: Terminology table Updated Section 3.3.1 Usage Added Section 3.3.2 Examples Updated Section 6.2 Additional APIs Added Section 6.2.13 Check Device Availability APIs Updated formatting on Section 6.2 Additional APIs Added Chapter 8: Black Box Debug Tool 	December 2021
336210	020	Updates for Intel® QAT Programmers Guide Hardware Version 1.7: <ul style="list-style-type: none"> Updated Section 3.14 Running Applications as Non-Root User 	November 2021
336210	019	Updates for Intel® QAT Programmers Guide Hardware Version 1.7: <ul style="list-style-type: none"> Updated Section 3.20 Rate Limiting Updated Section 3.21 DU Manager Application Updated Section 3.21.1 Commands to Fetch Device Utilization Updated formatting in Section 3.21.2 Durations Updated Section 4.2.1 General Parameters Updated Table 10. General Default Configuration Parameters 	October 2021

Document Number	Revision Number	Description	Revision Date
		<ul style="list-style-type: none"> Updated formatting for Table 11. General Parameters, Table 12. [KERNEL] Section Parameters, Table 13. [KERNEL_QAT] Section Parameters, Table 14. [KERNEL_QAT] Section Parameters, Table 15. Configuring Physical Functions and Virtual Functions, Table 16. Cryptographic Logical Instance Parameters 	
336210	018	Updates for Intel® QAT Programmers Guide Hardware Version 1.7: <ul style="list-style-type: none"> Section 3.20 Rate Limiting Section 3.20.1 Service Level Agreement (SLA) Section 3.20.2 SLA Units Section 3.20.3.1 Commands to Fetch Device Utilization Section 3.21 DU Manager Application Section 3.21.1 Commands to Fetch Device Utilization Section 6.1.1.3.5 CPA_DC_ASB_ENABLED Section 6.1.1.4 Maximal Expansion and Destination Buffer Size in compression direction 	September 2021
336210	017	Updates for Intel® QAT Programmers Guide Hardware Version 1.7 <ul style="list-style-type: none"> Clarified "exception" vs "error" in the Section 6.1.1.1 title. Add new section 6.1.1.5: "Avoiding a Compression Overflow exception" Updated Table 10 under Section 4.2.1 ("General Parameters") to reflect a new capability of hashing being available with ServicesProfile = COMPRESSION 	May 2021
336210	016	Updated guidance to enable rate limiting	March 2021
336210	015	Updates for Intel® QAT Programmers Guide Hardware Version 1.7 Added new sections: <ul style="list-style-type: none"> Section 5 Secure Architecture Considerations 	December 2020
336210	014	Updates for Intel® QAT Programmers Guide Hardware Version 1.7 Added new sections: <ul style="list-style-type: none"> Section 6.2.11 Congestion Management APIs Section 6.2.12 Service Specific Polling APIs 	September 2020
336210	013	Updates for Intel® QAT software v4.10.0 release: <ul style="list-style-type: none"> Revised Note in Section 2.2.1.3 	June 2020

Document Number	Revision Number	Description	Revision Date
		<ul style="list-style-type: none"> Section 3.8.1 Added Note before Table 5 Revised Table 10 Table 11, removed StorageEnabled and PkeServiceDisabled parameter 	
336210	012	Updates for Intel® QAT software v4.8.0 release: <ul style="list-style-type: none"> Revised Section Chapter, 3.4 Application Payload Memory Allocation 	February 2020
336210	011	Updated: <ul style="list-style-type: none"> Permissions for using huge pages with included memory driver Rate limiting and device utilization measurement impacts performance when active 	November 2019
336210	010	Updates for 4.7.0 release: <ul style="list-style-type: none"> Added virtual functions to list of configurable instances Rate limiting and device utilization measurement features 	October 2019
336210	009	Updated configuration options for concurrent requests (Tables 10, 14, 15)	July 2019
336210	008	Updates for 4.6.0 release: <ul style="list-style-type: none"> Dummy responses added to Heartbeat feature Handling device failures in a ritualized environment 	June 2019
336210	007	Updates for 4.5.0 release: <ul style="list-style-type: none"> Updated list of general parameters Updated list of Intel® QuickAssist entries in <code>/sys/kernel/debug</code> 	March 2019
336210	006	Updates for 4.4.0 release: <ul style="list-style-type: none"> Updated list of Compression API Errors 	December 2018
336210	005	Updates for 4.3.0 release: <ul style="list-style-type: none"> Intel® QuickAssist API in kernel space Added epoll content Updates for the Compress and Verify and Recover feature Other minor changes 	September 2018

Document Number	Revision Number	Description	Revision Date
336210	004	Added description of Compress and Verify and Recover (CnVnR) capability.	June 2018
336210	003	Added Heartbeat description. Clarified explanations of stateless and stateful compression and decompression.	April 2018
336210	002	Stateful compression is no longer supported by default.	April 2018
336210	001	Initial public release.	August 2017

§

1 Introduction

This programmer’s guide provides information on the architecture of the software and usage guidelines., information on the use of Intel® QuickAssist Technology (Intel® QAT) APIs, which provide the interface to the acceleration services (cryptographic and data compression), is documented in the related Intel® QAT software library documentation (refer to [Table 2](#)).

1.1 Terminology

In this document, for convenience:

- The software package is used as a generic term for the Intel® QAT software package for Hardware Version 1.7.
- Acceleration driver is used as a generic term for the software that allows the Intel® QAT Software Library APIs to access the Intel® QAT Endpoint(s).

Table 1. Terminology

Term	Description
ADF	Acceleration Driver Framework
AE	Acceleration Engine
AES	Advanced Encryption Standard
ASIC	Application Specific Integrated Circuit
AU	Acceleration Unit
BDF	Bus Device Function
BMSM	Broad Market Switch Mode
BnP	Batch and Pack
BTS	Base Transceiver Station
CBC	Cipher Block Chaining mode
CCM	Counter with CBC-MAC mode
CnV	Compress and Verify
CnVnR	Compress and Verify and Recover
CPK	Columbia Park
CY	Cryptography
DC	Data Compression
DID	Device ID
DMA	Direct Memory Access

Term	Description
DPDK	Data Plane Development Kit
DRAM	Dynamic Random-Access Memory
DSA	Digital Signature Algorithm
DTLS	Datagram Transport Layer Security
ECC	Elliptic Curve Cryptography
EVP	Envelope (OpenSSL* high-level cryptographic functions)
FW	Firmware
GCM	Galois/Counter Mode
GPL	General Public License
HLP	Highland Park
HMAC	Hash-based Message Authentication Mode
IA	Intel® Architecture
I/O	Input/Output
IDC	Inter Driver Communication
IDS/IPS	Intrusion Detection System/Intrusion Prevention System
IEEE	Institute of Electrical and Electronics Engineers
IKE	Internet Key Exchange
Intel® QAT	Intel® QuickAssist Technology
IOCTL	Input Output Control function
IOMMU	Input-Output Memory Management Unit
IOSF-SB	Intel® On-chip System Fabric Side Band
IPSec	Internet Protocol Security
LKCF	Linux* Kernel Cryptographic Framework
LTTng	Linux* Trace Toolkit Next Generation
MGF	Mask Generation Function
MSI	Message Signaled Interrupts
NAC	Network Acceleration Complex
NUMA	Non-uniform Memory Access
OP Data	Operational Data
PCH	Platform Controller Hub. In this manual, a Platform Controller Hub device includes standard interfaces and Intel® QAT Endpoint and I/O interfaces.

Term	Description
PCI	Peripheral Connect Interface
PCIe*	PCI Express*
PF	Physical Function
PKE	Public Key Encryption
RSA	Rivest-Shamir-Adleman
RTE	Run-Time Environment
SA	Security Association
SADB	Security Association Database
SAL	Service Access Layer
SATA	Serial Advanced Technology Attachment
SGL	Scatter-Gather List
SHA	Secure Hash Algorithm
SKU	Stock Keeping Unit
SoC	System-on-a-Chip
SPI	Serial Peripheral Interconnect
SR-IOV	Single Root I/O Virtualization
SSC	Storage Subsystem Class
SSL	Secure Sockets Layer
SYM	Symmetric Crypto
TCG	Trusted Computing Group
TLS	Transport Layer Security
TPM	Trusted Platform Module
USDMA	User Space DMA-able Memory
VF	Virtual Function
VPN	Virtual Private Network
WAN	Wide Area Network
WQM	Work Queue Manager

Table 2. Reference Documents and Resources

Document	Document Number/ Location
Intel® QuickAssist Technology Software for Linux* CE Release Notes	336211
Intel® QuickAssist Technology Software for Linux* Release Notes (Hardware Version 1.8 for In-line)	613775
Intel® QuickAssist Technology Software for Linux* CE Getting Started	336212
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
Using Intel® Virtualization Technology (Intel® VT) with Intel® QuickAssist Technology Application Note	330689

1.2 Typographical Conventions

The following font conventions are used in this manual:

- **Courier font** - file names, path names, executables, code examples, command line entries, API names, parameter names, and other programming constructs
- *Italic text* – key terms and publication titles
- **Bold text** - graphical user interface entries, buttons, keyboard keys and Intel® software names

§

2 Software Overview

In addition to the hardware mentioned in [Section 3.1, Hardware/Software Overview](#), the respective platforms have critical software components that are part of the offering. The software includes drivers and acceleration code that runs on the Intel® Architecture (IA) CPUs and Intel® QAT Endpoints.

2.1 Intel® Communications Chipset 8925 to 8955 Series Compatibility

While the focus of this document is on Intel® QAT software for Hardware Version 1.7, the Intel® Communications Chipset 8925 to 8955 Series is also supported.

2.2 Logical Instances

A logical instance may be thought of as a channel to the hardware. A logical instance allows an address domain (that is, kernel space and individual user space processes) to configure the rings to be used by that address domain and to define the behavior of that ring.

2.2.1 Response Processing

In the kernel space, each logical instance can be configured to operate in one of the two modes:

- Interrupt mode
- Polled mode

In the user space, each logical instance can be configured to operate in one of the two modes:

- Polled mode
- Epolled mode

2.2.1.1 Interrupt Mode

The interrupt is only supported in Kernel space. In User space it is no longer supported; therefore, the user space instance can no longer be configured with interrupt enabled mode.


When configured in interrupt mode, the Accelerator Driver Framework (ADF) registers an interrupt handler for response ring processing.

As the latency in servicing an interrupt may be costly, the hardware-assisted ring provides a mechanism to amortize the cost of interrupts into a single interrupt that may service multiple responses. The interrupt coalescing section of the configuration file allows the user to select the mechanism to amortize response interrupts using either a time-based interrupt scheme or a number-of-responses-based scheme.

The ADF registers an interrupt handler to service the ring bank interrupt. When an interrupt fires, the ADF services the interrupt and creates an interrupt handler bottom half to consume the responses from the response ring. When MSI-X is supported, the bottom half of the interrupt handler is created and affinitized to the configured core. Callbacks to the application code occur in the context of this taskset. This sequence is shown in the following figure (the full sequence has been reduced for clarity).

NOTE: *Linux* (and other operating systems) split an interrupt handler into two halves. The so-called "top half" is the routine that responds to the interrupt, that is, the one you register with `request_irq`. The "bottom half" is a routine that is scheduled by the top half to be executed later, at a safer time.*

Figure 1. Kernel Space Response Ring Processing



If the cost of servicing an interrupt and scheduling the interrupt handler bottom half is not desired, a user can choose to disable interrupts and poll for responses. This mechanism can be configured on a per logical instance basis by setting the `Dc/ CyXIsPolled` attribute of a logical instance in the configuration file to 1. When configured to 1, the ADF does not service interrupts for that logical instance.

The ADF provides a set of APIs to allow the client to poll a single bank or all banks on a given accelerator:

- `icp_sal_pollBank` - Poll the rings on the given bank number for a given accelerator.
- `icp_sal_pollAllBanks` - Poll the rings on all banks for a given accelerator.

The Service Access Layer (SAL) provides an API to poll on an individual logical instance:

- `icp_sal_CyPollInstance` - Poll a specific Cryptographic (CY) logical instance.

- `icp_sal_DcPollInstance` - Poll a specific Data Compression (DC) logical instance.

Refer to [Section 6.2.2, "Polling Functions"](#) for details on all the polling functions.

2.2.1.2 Epoll Mode

The event-based poll mode is called "epoll mode". The Intel® QAT driver's new mode supports the Linux* epoll interface. The Linux* epoll is a scalable I/O event notification mechanism intended to replace the older select/poll system calls.

NOTE: *For performance reasons, in epoll mode, only one instance (and one process) per bank should be used.*

To use the Linux* epoll, the user space application uses the following APIs:

- `epoll_create()` / `epoll_create1()` - creates an epoll instance and returns a file descriptor referring to that instance.
- `epoll_ctl()` - registers the file descriptors which will be polled.
- `epoll_wait()` - waits for I/O events for the file descriptors registered via `epoll_ctl`, blocking the calling thread if no events are currently available.

For more information, consult the Linux* epoll manuals, here: <http://man7.org/linux/man-pages/man7/epoll.7.html>

NOTE: *The Intel® QAT driver's epoll mode is only used by the user space instances, it is not valid for the kernel space.*

The Intel® QAT driver's epoll mode consists of two parts: the kernel space part and the user space part.

The coalescing fields expose the same behavior for the epoll mode. If the interrupt is delayed by changing the Coalescing fields, the event delivery to user space will be delayed too.

To enable the epoll mode, ensure the following steps are followed:

1. In the configuration file, please use the "`IsPolled = 2`" for the user space instance, for example:

```
Cy0Name = "SSL0"
Cy0IsPolled = 2
```
2. Whether the application uses the driver in a synchronously or asynchronously, it should create a thread to call the Intel® QAT drivers epoll API and the Linux* standard epoll interface.

The Intel® QAT drivers epoll API:

```
Crypto: icp_sal_CyGetFileDescriptor() / icp_sal_CyPutFileDescriptor()
Compression: icp_sal_DcGetFileDescriptor() /
icp_sal_DcPutFileDescriptor()
```

The Linux* standard epoll interface:

```
epoll_create() / epoll_ctl() / epoll_wait()
```

NOTE: For performance reasons, in *epoll* mode, only one instance (and one process) per bank should be used. The instance can be a crypto or compression instance.

For QAT 1.7 Generation Hardware:

When a bank is used for the *epoll* mode, it means there is only one instance (crypto or compression) for this bank. When the instance is used by a process, it means the process is the only user for this bank. Other processes could not use this bank temporarily. But if the process releases this instance, other processes can use this bank. Since there is only one instance for this bank, no more than 16 user space instances are available for 1.7 HW to configure all the banks for the *epoll* mode vs 128 user space instances for 1.8 HW. (For the Intel® Communications Chipset 8925-8955 series, up to 32 user space instances are available.)

For QAT 1.8 Generation Hardware:

If a process needs to provide compression and crypto services at the same time, it will need two instances, which means the process needs two banks. In such a scenario, no more than eight processes can be used for 1.7 HW vs 64 processes for 1.8 HW. (For the Intel® Communications Chipset 8925-8955 series, up to 16 processes can be used.)

For comparison purposes, when the CPU is in the idle state, for the user space instance, the standard poll mode ("`IsPolled = 1`") will poll the empty rings periodically and the polling will consume some CPU cycles (for instance, 2% usage may appear available when the CPU is in the idle state). But if *epoll* mode is used, the usage will stay at 0% when the CPU is in the idle state.

NOTE: The standard poll mode performs better when the CPU is in the high load state.

For user space instances, interrupt mode is no longer supported. Interrupt mode for the user space did not consume CPU cycles when there was no data in the response rings, unlike the polling mode, which continues to check at specified intervals. With the *epoll* support, standard Linux* *epoll* APIs, such as `epoll_create()`/`epoll_ctl()`/`epoll_wait()`, can be used.

Most web servers and socket-based applications, such as Nginx*, Apache*, etc., use one of *epoll*/*select*/*poll* to be notified when a socket is available for reading or writing, and then take appropriate action. With the *epoll* mode, the Intel® QuickAssist Technology driver will have more seamless integration into existing applications, such as Nginx*, as it will be using a standard notification mechanism.

§

3 Acceleration Drivers Overview

Selected Intel® products support Intel® QAT. Depending on the product chosen, Intel® QAT accelerates both or either of two services: cryptography (both symmetric and public key) and data compression.

The Intel® QAT Endpoints are exposed as Peripheral Connect Interface (PCI) devices. Applications running in the user space typically access these services via the Intel® QAT APIs. Support for the applications that run only in the kernel space is planned for a future software release, but driver support for the Linux* Kernel Cryptographic Framework (LKCF) API is present in this software release (default disabled).

3.1 Hardware/Software Overview

Because the hardware is accessed using the Intel® QAT APIs, it is not necessary to know all the hardware and software architecture details, but some knowledge of the underlying hardware and software is helpful for performance optimization and debugging purposes. For example, to support customers with different acceleration performance requirements, the Intel® C62x Chipset is available in different SKUs and supports two different "fabric configurations". [Figure 2](#) and [Figure 3](#) show two possible configurations for the acceleration endpoints in one Intel® C62x Chipset die.

Figure 2. Intel® C62x Chipset (PCH) Acceleration Endpoint Configuration 1

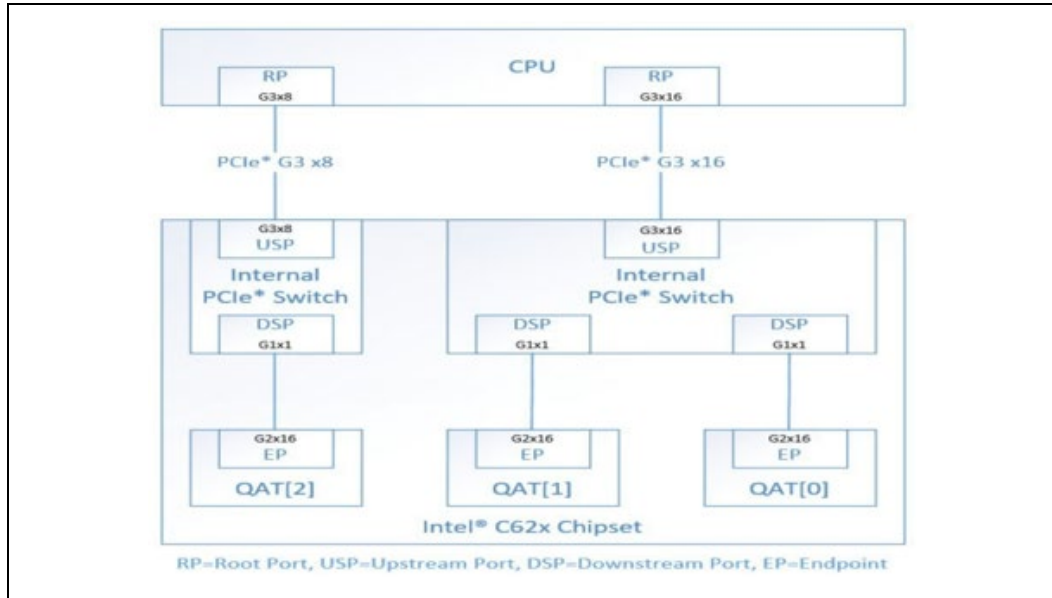
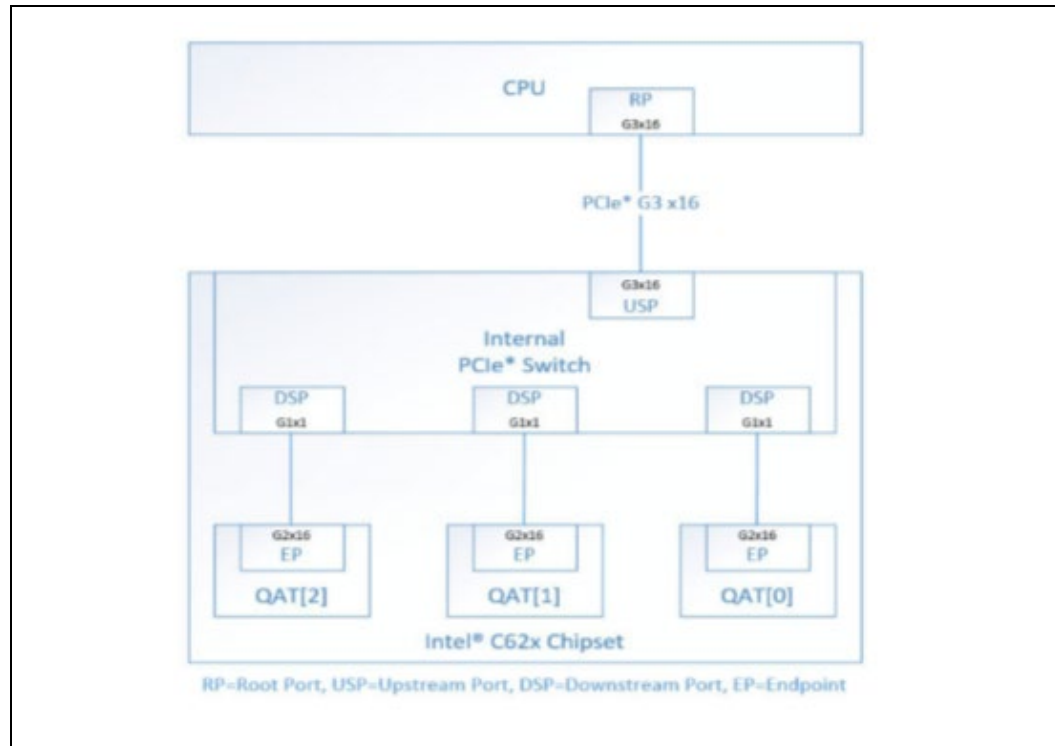


Figure 3. Intel® C62x Chipset (PCH) Acceleration Endpoint Configuration 2



For a given platform, the specific internal connections and number of Intel® QAT Endpoints per die (for instance, up to three for Intel® C62x Chipset) is product dependent, SKU-dependent, routing-dependent (i.e., how many lanes are routed), and configuration-dependent (e.g., with different fabric configuration soft-straps). For each Intel® QAT Endpoint (e.g., `QAT[0]`), hardware-assisted rings are used as the communication mechanism to transfer requests between the CPU and the Intel® QAT Endpoint(s) and vice-versa. The 1.7 HW supports 256 rings vs 1024 rings for 1.8 HW (per Intel® QAT Endpoint), each with head and tail Configuration Status Register (CSR) pointers that are mapped to PCIe* memory on the CPU. Rings are assigned by the provided software based on the Cryptography (CY) and Data Compression (DC) instances declared in the configuration files. Refer to [Section 3.2, Acceleration Driver Configuration File](#) for more information.

Each Intel® QAT Endpoint has multiple computation engines. For a given Intel® QAT Endpoint, all rings associated with that endpoint are shared, and the hardware load balances requests from these rings.

A user can write directly to the Intel® QAT APIs, or the use of Intel® QAT can be done through frameworks that have been enabled by others including Intel® (for example, `zlib*`, `OpenSSL*`, `libcrypto*`, and the `Linux* Kernel Crypto Framework`).

The driver architecture supports simultaneous operation of multiple applications.

3.2 Acceleration Driver Configuration File

An acceleration driver has a configuration file that is used to configure the driver for runtime operation. There is a single configuration file for each Intel® QAT Endpoint in the system. If Single-Root Input/Output Virtualization (SR-IOV) is enabled, a separate configuration file is used for each virtual function, if applicable. The configuration file format is described in [Section 4.1, Configuration File Overview](#).

3.3 Utility for Loading Configuration Files and Sending Events to the Driver - `adf_ctl`

The `adf_ctl` user space utility is separate from the driver and provides a mechanism for:

- Loading configuration file data to the kernel driver. The kernel space driver uses the data and provides it to the user space driver.
- Sending events to the driver to bring devices up and down.

The `adf_ctl` provided with the Intel® QAT 1.7 & 1.8 drivers can be used to interface with Intel® QAT v1.6, 1.7 and 1.8 devices.

3.3.1 Usage

- To bring up, down, restart or reset device(s):

```
./adf_ctl [-c|--config] [config_file_path] [qat_dev<N>]
[up|down|restart|reset]
```
- To print device(s) status:

```
./adf_ctl [qat_dev<N>] status
```
- To use the specified configuration file:

```
-c (--config) [config_file_path]
```

NOTE: If no device (physical or virtual) is selected, this file is used against all existing devices.

3.3.2 Examples

- To bring device 0 down:

```
./adf_ctl qat_dev0 down
```
- To load device configuration from default path `/etc/c4xxx_dev1.conf`, then bring device 1 up:

```
./adf_ctl qat_dev1 up
```
- To load device configuration from specified path `~/user_c4xxx_dev1.conf`, then bring device 1 up:

```
./adf_ctl -c ~/user_c4xxx_dev1.conf qat_dev1 up
```
- To restart all devices with default configuration file `~/user_c4xxx_dev1.conf`:

```
./adf_ctl restart
```


- To restart all devices with specified configuration file `~/user_c4xxx_dev1.conf`:

```
./adf_ctl -c ~/user_c4xxx_dev1.conf restart
```
- To restart device 0 with specified configuration file `~/user_c4xxx_dev1.conf`:

```
./adf_ctl -c ~/user_c4xxx_dev1.conf qat_dev0 restart
```
- To restart device 0:

```
./adf_ctl qat_dev0 reset
```

3.4 Application Payload Memory Allocation

When performing offload operations through the Intel® QAT API, it is required that the payload data be placed in a buffer that is resident, physically contiguous, and Direct Memory Access (DMA) accessible from the acceleration hardware. It is the applications responsibility to provide buffers with these constraints.

Buffers are passed to the API with virtual addresses. The API translates these addresses to the address information required by the hardware (see the following table).

Table 3. Services

Service	API	Reference
Cryptographic service	<code>cpaCySetAddressTranslation</code>	See the Intel® QuickAssist Technology Cryptographic API Reference Manual (refer to Table 2) for details.
Data Compression service	<code>cpaDcSetAddressTranslation</code>	See the Intel® QuickAssist Technology Data Compression API Reference Manual (refer to Table 2) for details.

When the software requires the physical address, it calls the registered function.

NOTE: *This address translation function is called at least once per request. Consequently, for optimal performance, the implementation of this function should be optimized.*

If using the Intel® QAT Data Plane API, buffers are passed to the Intel® QAT API as physical addresses. The library passes this directly to the hardware, without the need for translation.

All these tasks can be performed utilizing the User Space DMA-able Memory (USDM) driver supplied with the Intel® QAT driver package. The driver consists of the kernel-mode and user-mode parts allowing allocation of 1k-aligned memory blocks, setting up address translation, and automatic block deallocation in case of a user application crash."

3.4.1 Thread Specific USDM

By default, memory allocation uses the USDM slab allocator, which gives 2MB contiguous memory. The allocation has locks in the library to prevent a race condition in getting the memory from the slab. This lock has an impact on some multi-threaded applications and use cases, like HAProxy, causing a drop in performance. To mitigate this issue, thread specific

USDM is implemented with the v4.21 release, which allocates and handles memory specific to threads. (For multi-thread apps, allocated memory information will be maintained separately for each thread). This feature can be enabled by configuring with the configure flag `--enable-icp-thread-specific-usdm`.

In some use cases with thread specific USDM, using a 128K slab allocator instead of the default 2MB allocator could improve performance and reduce memory consumption for a large number of threads. This can be enabled by configuring with the configure flag `--enable-128k-slab`.

NOTE: *There is a limitation with thread specific USDM: memory allocated in one thread should be freed only by the thread which allocates it. Incorrect cleanup can lead to a segmentation fault (segfault). Also, memory allocated in a thread is freed automatically when the thread exits/terminates, even if the user does not explicitly free the memory.*

See the *Getting Started Guide* for more information on `./configure` flags.

We have observed poor multithreaded performance with QAT_Engine using OpenSSL* at higher thread counts. Unfortunately, these issues appear to stem from the way OpenSSL* implements its `engine_table_select` and locks. For relevant issues on the OpenSSL* github pages, see the two issues below:

- **OpenSSL* 1.1.1.x: Performance bottleneck with locks in engine_table_select() function #18509**, <https://github.com/openssl/openssl/issues/18509>
- **OpenSSL* 3.0: 3.0 performance degraded due to locking #20286**, <https://github.com/openssl/openssl/issues/20286>

3.5 User Space Additional Functions

To allow a user space access to the Intel® QAT rings, the service access layer must be configured to expose logical instances to the user space process. Logical instances are configured using the per device configuration files.

To allow each process to have separate logical instances, the configuration file groups a set of logical instances by name. The process then must call the `icp_sal_userStart` function (refer to [Section 6.2.4.1](#)) at initialization time with the name associated with the group of logical instances. Similarly, on process exit, to free the resources and make them available to other processes with the same name, the process must call the function `icp_sal_userStop` (refer to [Section 6.2.2.1](#)).

For example, the user can configure the driver to have two crypto logical instances available for the process called "SSL". The user space process may then access these logical instances by calling the `cpaCyGetInstances` function. The application may then initiate a session with these logical instances and perform a cryptographic operation. See the *Intel® QuickAssist Technology Cryptographic API Reference Manual*. Refer to [Table 2](#) of the manual for more information on the API functions available for use.

For this example, the logical instances section of the configuration file is as follows:

```
[SSL]
NumberCyInstances = 2
```

```

NumberDcInstances = 0
NumProcesses = 1
LimitDevAccess = 0

# Crypto - User instance #0
Cy0Name = "SSL0"
Cy0IsPolled = 1
# List of core affinities
Cy0CoreAffinity = 1

# Crypto - User instance #1
Cy1Name = "SSL1"
Cy1IsPolled = 1
# List of core affinities
Cy1CoreAffinity = 2

```

In this example, the user process Secure Sockets Layer (SSL) configures two logical instances (called "SSL0" and "SSL1").

3.6 Managing Intel® QAT Endpoints Using `qat_service`

The `qat_service` script is installed with the software package in the `/etc/init.d/` directory. The script allows a user to start, stop, shutdown or query the status (up or down) of a single Intel® QAT Endpoint or all Intel® QAT Endpoints in the system.

Usage:

```
# ./qat_service start||stop||status||restart||shutdown
```

To view all Intel® QAT Endpoints in the system, use:

```
# ./qat_service status
```

If there are two Intel® QAT Endpoints in the system, the output will be as follows:

```
qat_dev0 - type: c6xx, inst_id: 0, bsf: 06:00:0, #accel: 5
#engines: 10 state: up
qat_dev1 - type: c6xx, inst_id: 1, bsf: 83:00:0, #accel: 5 #engines: 10
state: up
```

For a system with multiple Intel® QAT Endpoints, you can start, stop or restart each device by passing the Intel® QAT Endpoint to be restarted or stopped as a parameter (`qat_dev<N>`). For example:

```
# ./qat_service stop qat_dev0
```

where the device number `<N>` is equal to 0 in this case.

The `shutdown` qualifier enables the user to bring down all Intel® QAT Endpoints and unload driver modules from the kernel. Compared with the `stop` qualifier, which brings down one or more Intel® QAT Endpoints, but does not unload kernel modules, so other Intel® QAT Endpoints can still run.

NOTE: In systems with more than three devices it might be necessary to change the `qat_service_timeout` in `/etc/systemd/system/qat_service.service.d/startup-timeout.conf`.

3.7 Overview of QAT *debugfs* entries

Some useful debugging information for the driver and configuration is available via the Linux* `debugfs` file system, with the entries `/sys/kernel/debug/qat_*` and `/sys/kernel/debug/qae_mem_dbg/qae_mem_slabs`.

For more information, see [Chapter 8: Black Box Debug Tool](#) or the [Intel® QuickAssist Technology Debugging Guide](#).

3.7.1 Entries in `/sys/kernel/debug/qat_*`

This includes:

Table 4. Intel® QuickAssist Technology `/sys/kernel/debug` Entries

Entry	Description
<code>cnv_errors</code>	Indicates number of compress and Verify errors. Refer to Section 6.1.5. Compress and Verify Error log in Sysfs :
<code>dev_cfg</code>	Displays internal device configuration information
<code>frequency</code>	Displays frequency of Acceleration Engines
<code>fw_counters</code>	Displays Acceleration Engine firmware requests/responses
<code>heartbeat heartbeat_failed heartbeat_sent</code>	Refer to Section 3.17.3.3.1 System Virtual Files
<code>transport</code>	Contains firmware request/response data. Available only for kernel space instances.
<code>version</code>	Includes package version information

3.7.2 Memory driver queries (`qae_mem_slabs`)

Debug features are also available by reading and writing the file `/sys/kernel/debug/qae_mem_dbg/qae_mem_slabs`. When *reading* the virtual/physical address, size and slab id together with the pid of the allocating process are shown. Writing a string to the file will start executing debug commands.

For example:

```
# cat /sys/kernel/debug/qae_mem_dbg/qae_mem_slabs
```

```
Pid 78854, Slab Id 10550771712
Virtual address 000000000b39412d, Physical Address 274e00000, Size
2097152
```

```
Pid 78854, Slab Id 10309599232
Virtual address 000000003670dd45, Physical Address 266800000, Size
2097152
. . .
```

There are three commands supported, and the table below shows their output. Writing these strings will give the output when the file is read.

Table 5. *qae_mem_slabs* Commands Supported

Command	Output
"d <pid> <virtual or physical address>"	The 256 byte in hex and ascii from the start address
"c <pid> <slab id>" (pid should be the process id that can be obtained by a previous read)	The allocation bit map for the given slab identifier
"t"	Total size of NUMA memory allocated in kernel space

For example, by combining a *write* to the file and a subsequent *read*, you can see the total allocated NUMA memory, e.g.:

```
# echo "t" > /sys/kernel/debug/qae_mem_dbg/qae_mem_slabs ; cat
/sys/kernel/debug/qae_mem_dbg/qae_mem_slabs
```

```
Total allocated NUMA memory: 0 bytes
```

As above, the "d" and "c" commands will output their respective information.

3.8 Compression Status Codes

The `CpaDcRqResults` structure should be checked for compression status codes in the `CpaDcReqStatus` data field. The mapping of the error codes to the enums is included in the `quickassist/include/dc/cpa_dc.h` file.

3.8.1 Intel® QAT Compression API Errors

The Intel® QAT Compression APIs that send requests to the compression hardware can return the error codes shown in the following table. These APIs are:

- `cpaDcCompressData()`
- `cpaDcDecompressData()`
- `cpaDcDpEnqueueOp()`
- `cpaDcDpEnqueueOpBatch()`

NOTE: *Decompression issues in the table below may also apply to the compression use case due to potential issues encountered during a Compress-and-Verify operation. In this case, the file(s) `/sys/kernel/debug/qat_*/cnv_errors` may show these nested errors. In some cases, the suggested corrective action may need to be to store the block uncompressed or to compress the block with software.*

Table 6. Intel® QAT Compression API Errors

Error Code	Error Type	Description	Suggested Corrective Action(s)
0	CPA_DC_OK	No error detected by compression hardware.	None.
-1	CPA_DC_INVALID_BLOCK_TYPE	Invalid block type (type = 3); invalid input stream detected for decompression	Decompression error. Discard output. For a stateless session, resubmit affected request. For a stateful session, abort the session calling <code>CpaDcRemoveSession()</code> .
-2	CPA_DC_BAD_STORED_BLOCK_LEN	Stored block length did not match one's complement; invalid input stream detected	Decompression error. Discard output. For a stateless session, resubmit affected request. For a stateful session, abort the session calling <code>CpaDcRemoveSession()</code> .
-3	CPA_DC_TOO_MANY_CODES	Too many length or distance codes; invalid input stream detected	Decompression error. Discard output. For a stateless session, resubmit affected request. For a stateful session, abort the session calling <code>CpaDcRemoveSession()</code> .
-4	CPA_DC_INCOMPLETE_CODE_LENS	Code length codes incomplete: invalid input stream detected	Decompression error. Discard output. For a stateless session, resubmit affected request. For a stateful session, abort the session calling <code>CpaDcRemoveSession()</code> .
-5	CPA_DC_REPEATED_LEN	Repeated lengths with no first length; invalid input stream detected	Decompression error. Discard output. For a stateless session, resubmit affected request. For a stateful session, abort the session calling <code>CpaDcRemoveSession()</code> .
-6	CPA_DC_MORE_REPEAT	Repeat more than specified lengths; invalid input stream detected	Decompression error. Discard output. For a stateless session, resubmit affected request. For a stateful session, abort the session calling <code>CpaDcRemoveSession()</code> .

Error Code	Error Type	Description	Suggested Corrective Action(s)
-7	<code>CPA_DC_BAD_LITLEN_CODES</code>	Invalid literal/length code lengths; invalid input stream detected	Decompression error. Discard output. For a stateless session, resubmit affected request. For a stateful session, abort the session calling <code>CpaDcRemoveSession()</code> .
-8	<code>CPA_DC_BAD_DIST_CODES</code>	Invalid distance code lengths; invalid input stream detected	Decompression error. Discard output. For a stateless session, resubmit affected request. For a stateful session, abort the session calling <code>CpaDcRemoveSession()</code> .
-9	<code>CPA_DC_INVALID_CODE</code>	Invalid literal/length or distance code in fixed or dynamic block; invalid input stream detected	Decompression error. Discard output. For a stateless session, resubmit affected request. For a stateful session, abort the session calling <code>CpaDcRemoveSession()</code> .
-10	<code>CPA_DC_INVALID_DIST</code>	Distance is too far back in fixed or dynamic block; invalid input stream detected	Decompression error. Discard output. For a stateless session, resubmit affected request. For a stateful session, abort the session calling <code>CpaDcRemoveSession()</code> .
-11	<code>CPA_DC_OVERFLOW</code>	Overflow detected. This is not an error, but an exception. Overflow is supported and can be handled.	Resubmit with a larger output buffer when appropriate. Table 22 in Section 6.1.1.1 gives details on the various overflow exceptions.
-12	<code>CPA_DC_SOFTERR</code>	Other non-fatal detected.	Discard output. For a stateless session, resubmit affected request. For a stateful session, abort the session calling <code>CpaDcRemoveSession()</code> .
-13	<code>CPA_DC_FATALERR</code>	Fatal error detected.	Discard output and abort the session calling <code>CpaDcRemoveSession()</code> .

Error Code	Error Type	Description	Suggested Corrective Action(s)
-14	<code>CPA_DC_MAX_RESUBMITERR</code>	On an error being detected, the firmware attempted to correct and resubmitted the request, however, the maximum resubmit value was exceeded. Maximal value is internally set in the firmware to 10 attempts. This is a QAT1.6 error only. This error code is considered as a fatal error.	Discard output and abort the session calling <code>CpaDcRemoveSession()</code> .
-15	<code>CPA_DC_INCOMPLETE_FILE_ERR</code>	This decompression error can be reported only by QAT 1.7 devices. However, it is not exposed to the application. The input file is incomplete. This indicates that the request was submitted with a <code>CPA_DC_FLUSH_FINAL</code> . However, a <code>BFINAL</code> bit was not found in the request.	No corrective action is required as it is not exposed to the application.
-16	<code>CPA_DC_WDOG_TIMER_ERR</code>	The request was not completed as a watchdog timer hardware event occurred.	Discard output and resubmit the affected request.
-17	<code>CPA_DC_EP_HARDWARE</code>	This is a recoverable error available only with QAT1.7 devices. Request was not completed as an end point hardware error occurred (for example, a parity error).	Discard output and abort the session calling <code>CpaDcRemoveSession()</code> .
-18	<code>CPA_DC_VERIFY_ERROR</code>	Compress and Verify (CnV). This is a compression direction error only. During the decompression of the compressed payload, an error was detected and the deflate block produced is invalid.	Discard output; resubmit affected request.
-19	<code>CPA_DC_EMPTY_DYM_BLK</code>	Decompression request contained an empty dynamic stored block (not supported).	Discard output.
-20	<code>CPA_DC_CRC_INTEG_ERR</code>	Compression CRC data integrity check error detected.	Discard output; resubmit affected request or abort the session.

NOTE: Except for the errors `CPA_DC_OK`, `CPA_DC_OVERFLOW`, `CPA_DC_FATALERR`, `CPA_DC_MAX_RESUBMITERR`, `CPA_DC_WDOG_TIMER_ERR`, `CPA_DC_VERIFY_ERR`, and `CPA_DC_EP_HARDWARE_ERR`, the rest of the error codes can be considered as invalid input stream errors.

NOTE: When the suggested corrective action is to discard the output, it implies that the application must also ignore the consumed data, the produced data, and the checksum values.

3.9 Stateful Compression Unsupported

Stateful compression is no longer supported.

3.10 Stateless Compression Level Details

The throughput and compression ratio for stateless compression can be adjusted with the compression levels to achieve particular requirements. The most recent software packages now support four compression levels, and the history buffer size is ignored.

3.10.1 Compression Level Mapping

3.10.1.1 QAT 1.7 hardware:

Compression levels 1 to 4 translate to search depth 1, 4, 8, and 16, respectively.

Compression levels 5 to 9 are retained for backward compatibility, but map to level 4.

Table 7. Compression Levels for QAT 1.7 Hardware

Compression Level (at the QAT API)	Search Depth	HB ¹ Size (KB)	Stateful Context Size (KB)			
			HB ¹	HT ²	LL ³	Total
1	1	32	32	16	0	48
2	4	16	16	16	32	64
3	8	16	16	16	32	64
4 through 9	16	16	16	16	32	64

1. History Buffer. For a search depth of 1, this is 32KB and uses Banks A, B, C and D. For other search depths, this is 16KB and uses Banks A and B.
2. Hash Table. Regardless of search depth, this is 16KB and uses Banks F and G.
3. Linked List. For a search depth of 1, this is not used. For other search depths, this is 32KB and uses Banks C, D, H and I.

3.10.1.1.2 QAT 1.8 hardware:

Compression levels 1 to 5 translate to search depth 1, 4, 8, 16, and 128, respectively.

Compression levels 6 to 9 are retained for backward compatibility but map to level 5.

Table 8. Compression Levels for QAT 1.8 Hardware

Compression Level (at the QAT API)	Search Depth	HB ¹ Size (KB)	Stateful Context Size (KB)			
			HB ¹	HT ²	LL ³	Total
1	1	32	32	16	0	48
2	4	16	16	16	32	64
3	8	16	16	16	32	64
4	16	16	16	16	32	64
5 through 9	128	16	16	16	32	64

4. History Buffer. For a search depth of 1, this is 32KB and uses Banks A, B, C and D. For other search depths, this is 16KB and uses Banks A and B.
5. Hash Table. Regardless of search depth, this is 16KB and uses Banks F and G.
6. Linked List. For a search depth of 1, this is not used. For other search depths, this is 32KB and uses Banks C, D, H and I.

3.10.2 Limitation on History Buffer Size (aka Deflate Window Size)

NOTE: These details are specific to QAT 1.x hardware.

NOTE: The history buffer size is also known as the deflate window size.

There are rare use cases where compressible files may have *worse* compression at higher compression levels, and this section explains those rare cases.

The issue is related to the history buffer size used during the compression process.

In level 2 (L2) through level 9 (L9) compression levels, the history buffer size is limited to 16KB. This buffer/window this is used to store previously processed data, and find matches within the data to achieve better compression. For files *smaller* than 16KB, higher compression levels will usually achieve better compression. However, for larger files, with a 16KB history buffer, the compression algorithm will not be able to find matches at lengths greater than 16KB. This limitation affects the compression ratio at L2 and above.

For example, there are specific files with most compressible matches at distances \geq the 16KB buffer size, all of which cannot be accessed at L2 or above. The result is that for these unique files, compression level L1 provides the best compression results. While this is generally uncommon, it could be common to specific datasets.

[Table 7](#) and [Table 8](#) (above) help understand the relationship between compression levels, history buffer size, and compression performance. These shows the history window, search depth, and context size for different compression levels.

3.11 Acceleration Driver Return Codes

The following table shows the return codes used by various components of the acceleration driver, defined in `quickassist/include/cpa.h`.

Table 9. Acceleration Driver Return Codes

Return Type	Return Code	Description
<code>CPA_STATUS_SUCCESS</code>	0	Requested operation was successful.
<code>CPA_STATUS_FAIL</code>	-1	A general or unspecified error occurred. Refer to the console log user space application or to <code>/var/log/messages</code> in kernel space for more details of the failure.
<code>CPA_STATUS_RETRY</code>	-2	Recoverable errors occurred. Refer to relevant sections of the API for specifics on what the suggested course of action.
<code>CPA_STATUS_RESOURCE</code>	-3	Required resource is unavailable. The resource that has been requested is unavailable. Refer to relevant sections of the API for specifics on what the suggested course of action.
<code>CPA_STATUS_INVALID_PARAM</code>	-4	Invalid parameter has been passed in.
<code>CPA_STATUS_FATAL</code>	-5	A fatal error has occurred. A serious error has occurred. The recommended course of action is to shut down and restart the component.
<code>CPA_STATUS_UNSUPPORTED</code>	-6	The function is not supported, at least not with the specific parameters supplied. This may be because the current implementation does not support a particular capability.
<code>CPA_STATUS_RESTARTING</code>	-7	The API implementation is restarting. Restarting may be reported if, for example, a hardware implementation is undergoing a reset.

The following table shows the return codes used by the acceleration driver to handle the Linux* device driver operations.

Table 10. Acceleration Driver Return Codes for Linux* Device Driver Operations

Return Type	Return Code	Description
<code>SUCCESS</code>	0	The operation was successful.

Return Type	Return Code	Description
FAIL	1	A general error occurred. Refer to the console log user space application or to <code>/var/log/</code> messages in kernel space for more details of the failure.
-EPERM	-1	Operation is not permitted. Used during ioctl operations.
-EIO	-5	Input/Output error occurred. Used when copying configuration data to and from user space.
-EBADF	-9	Bad File Number. Used when an invalid file descriptor is detected.
-EAGAIN	-11	Try Again. Used when a recoverable operation occurred.
-ENOMEM	-12	Out of Memory. A memory resource that has been requested is not available.
-EACCES	-13	Permission Denied. Used when the operation failed to connect to a process or open a device.
-EFAULT	-14	Bad Address. Used when an operation detects invalid parameter data.
-ENODEV	-19	No Such Device. Used when an operation detects invalid device id.
-ENOTTY	-25	Invalid Command Type. Used when an ioctl operation detects an invalid command type.

3.12 Batch and Pack Compression Unsupported

Batch and Pack (BnP) compression are no longer supported.

3.13 Compress and Verify Feature

The Compress and Verify (CnV) feature check and ensures data integrity in the compression operation of the Data Compression API. This feature introduces an independent capability to verify the compression transformation.

Refer to Intel® QuickAssist Technology Data Compression API Reference Manual.

NOTE:

1. CnV is always enabled via the `cpaDcCompressData ()` API.
2. CnV supports compression operations only.
3. The `compressAndVerify` flag in the `CpaDcDpOpData` structure should be set to `CPA_TRUE` when using the `cpaDcDpEnqueueOp ()` or `cpaDcDpEnqueueOpBatch ()` API. These APIs are declared in the API file `cpa_dc_dp.h`.

- The `compressAndVerify` flag in the `CpaDcOpData` structure should be set to `CPA_TRUE` when using the `cpaDcCompressData2 ()` API. This API is declared in the API file `cpa_dc.h`.

The CnV functionality is implemented in the Data Compression APIs

`cpaDcCompressData ()`, `cpaDcCompressData2 ()`, `cpaDcDpEnqueueOp ()` and `cpaDcDpEnqueueOpBatch ()` for the compression path only.

These APIs are declared and documented in the API file `cpa_dc.h`.

NOTE: *It is possible to recover from Compress and Verify errors in a seamless manner. Refer to the Compress and Verify and Recover discussion in [Section 6.1.3](#).*

3.14 Running Applications as Non-Root User

The installation of Intel® QAT software package configures the driver to allow applications to run as Non-Root User. The users must be added to the 'qat' group.

When the `make install` is performed at the directory where the Intel® QAT package is installed, the following `udev` file is created that is responsible for setting up non-root access.

```
KERNEL=="qat_adf_ctl" MODE="0660" GROUP="qat" RUN+="/bin/chgrp qat
/usr/local/bin/adf_ctl"
KERNEL=="qat_dev_processes" MODE="0660" GROUP="qat"
KERNEL=="usdm_drv" MODE="0660" GROUP="qat"
ACTION=="add", DEVPATH=="/module/usdm_drv" SUBSYSTEM=="module"
RUN+="/bin/mkdir /dev/hugepages/qat"
ACTION=="add", DEVPATH=="/module/usdm_drv" SUBSYSTEM=="module"
RUN+="/bin/chgrp qat /dev/hugepages/qat"
ACTION=="add", DEVPATH=="/module/usdm_drv" SUBSYSTEM=="module"
RUN+="/bin/chmod 0770 /dev/hugepages/qat"
ACTION=="remove", DEVPATH=="/module/usdm_drv" SUBSYSTEM=="module"
RUN+="/bin/rmdir
/dev/hugepages/qat"
KERNEL=="uio*", ATTRS{vendor}=="0x8086", ATTRS{device}=="0x0435"
MODE="0660" GROUP="qat"
KERNEL=="uio*", ATTRS{vendor}=="0x8086", ATTRS{device}=="0x0443"
MODE="0660" GROUP="qat"
KERNEL=="uio*", ATTRS{vendor}=="0x8086", ATTRS{device}=="0x37c8"
MODE="0660" GROUP="qat"
KERNEL=="uio*", ATTRS{vendor}=="0x8086", ATTRS{device}=="0x37c9"
MODE="0660" GROUP="qat"
KERNEL=="uio*", ATTRS{vendor}=="0x8086", ATTRS{device}=="0x6f54"
MODE="0660" GROUP="qat"
KERNEL=="uio*", ATTRS{vendor}=="0x8086", ATTRS{device}=="0x6f55"
MODE="0660" GROUP="qat"
KERNEL=="uio*", ATTRS{vendor}=="0x8086", ATTRS{device}=="0x19e2"
MODE="0660" GROUP="qat"
KERNEL=="uio*", ATTRS{vendor}=="0x8086", ATTRS{device}=="0x19e3"
MODE="0660" GROUP="qat"
```

The updates to the `udev` rules are performed during the installation of the Intel® QAT driver.

The following steps need to be manually applied:

Change the amount of max locked memory for the username that is included in the group name (which is by default 64). This can be done by specifying the limit in

```
/etc/security/limits.conf.  
@qat - memlock 4096
```

3.15 Random Number Generation

Starting with Intel® QAT Hardware version 1.7, Intel® QAT no longer includes random number generation capability, because this capability is already included in the CPU and is available via the RDRAND and RDSEED instructions.

3.16 Huge Pages with the Included Memory Driver

The included User space DMAable Memory driver (`usdm_drv.ko`) supports 2 MB pages. This allows direct access to main memory by devices other than the CPU and the actual supported maximum memory size in one individual allocation when huge pages is enabled is 2 MB – 5 KB. Where the 5 KB is used for memory management for the memory driver. The use of 2 MB pages provides benefits, but also requires additional configuration. Use of this capability assumes that enough huge pages are allocated in the operating system for the particular use case and configuration.

Here are some examples use cases:

```
# insmod ./usdm_drv.ko
```

Default settings applied.

```
# insmod ./usdm_drv.ko max_mem_numa=32768
```

Maximum amount of Non-uniform Memory Access (NUMA) type memory that the User Space DMA-able Memory (USDMA) driver can allocate is 32 MB in total for all processes. Huge pages are disabled.

```
# insmod ./usdm_drv.ko max_huge_pages=50 max_huge_pages_per_process=5
```

Maximum number of huge pages that the USDMA can allocate is 50 in total and 5 per process (up to 10 processes, 0 for the next processes).

```
# insmod ./usdm_drv.ko max_huge_pages=3  
max_huge_pages_per_process=5
```

An erroneous configuration, maximum number of huge pages that USDMA can allocate is 3 totals: 3 for a first process, 0 for the next processes.

```
# insmod ./usdm_drv.ko max_huge_pages_per_process=5
```

An invalid configuration, huge pages are disabled because `max_huge_pages` is 0 by default.

```
# insmod ./usdm_drv.ko max_huge_pages=5
```

An invalid configuration, huge pages are disabled because `max_huge_pages_per_process` is 0 by default.

NOTE: *The use of huge pages may not be supported for all use cases. For instance, depending on the driver version, some limitations may exist for an Input/Output Memory Management Unit (IOMMU).*

3.17 Heartbeat

Under some circumstances, firmware in the Intel® QAT devices could become unresponsive, requiring a device reset to recover. The Intel® QAT Heartbeat feature provides a mechanism for the customer application to detect and reset unresponsive devices. It also notifies the application processes of the start and end of the reset operation and suspends all Intel® QAT instances between the events.

3.17.1 Heartbeat Operation

A Heartbeat enabled Intel® QAT device firmware periodically writes counters to a specified physical memory location. A pair of counters per thread is incremented at the start and end of the main processing loop within the firmware. Checking for Heartbeat consists of checking the validity of the pair of counter values for each thread. Stagnant counters indicate a firmware hang.

3.17.1.1 Initialization

At startup, the Intel® QAT device driver allocates memory for the counter pairs to be written by the firmware and then sends a message to the firmware to start the heartbeat functionality.

3.17.1.2 Heartbeat Monitoring

Heartbeat check/monitoring refers to the invocation of one of the two API calls that check if the device is responsive. Heartbeat failure refers to the API returning failure.

The Intel® QAT driver does not monitor for Heartbeat. It should be initiated by a Heartbeat management thread calling one of the following APIs periodically:

- `icp_sal_check_device(Cpa32U accelId);`
- `icp_sal_check_all_devices(void);`

A failure return code implies the device has failed or hung.

The Heartbeat management thread should satisfy the following conditions:

- For any given device, only one such process/thread should monitor.
- One process can monitor one or more devices.
- Can be a user application that uses Intel® QAT services, or a separate management/control plane process.
- In virtualized environment, monitoring process(es)/thread(s) must run in the context of the host or hypervisor.

3.17.1.3 Resetting a Failed Device

A device can be configured for automatic reset by the Intel® QAT framework or manually reset by the application by using the `AutoResetOnError` field in the device configuration file `/etc/<device>.conf`, as shown in [the following table](#).

Table 11. AutoResetOnError Values

AutoResetOnError Value	Action on Heartbeat Failure
0 (default)	Do not reset the device
1	Reset the device automatically

If an Intel® QAT device is not configured for automatic reset, the management thread should reset it using the `icp_sal_reset_device(Cpa32U accelId)` API.

The `icp_sal_reset_device()` function starts an asynchronous reset sequence and returns immediately. The reset function should not be called again until the device has completed the reset to avoid a reset storm. The `icp_sal_check_device(<device id>)` function could be called in a loop to check if the device reset is still in progress.

If the application devices are all configured for automatic reset, then the `icp_sal_check_all_devices()` function could be used; otherwise, the function should not be used because it does not return the identity of the failed device, which is a required parameter for the `icp_sal_reset_device()` function.

3.17.1.3.1 Function Signatures

The details of the above functions, parameters, and return values can be found in [Section 6.2, Additional APIs](#).

3.17.2 Incorporating Heartbeat into Intel® QAT Applications

A typical Intel® QAT user application consists of two tasks:

- The first task is typically an application thread that initializes Intel® QAT instances and sessions, and then submits service requests for Intel® QAT crypto or compression.
- If an application employs polling to receive Intel® QAT service responses, then this task is also an application thread. Alternatively, responses are received as an interrupt handler.

Two more tasks are required to support Heartbeat:

- The first is a management task to monitor the devices for failure or hang and then resets them, when required. As discussed earlier, this could be an application thread of an independent management process.
- The second task is an application thread that polls for device reset events:

— `CPA_INSTANCE_EVENT_RESTARTING` (device is restarting)

— `CPA_INSTANCE_EVENT_RESTARTED` (device restart is complete)

If the application employs polling to receive Intel® QAT service responses, then this task could be included in the same polling loop.

The polling for device events is done using the API:

- `icp_sal_poll_device_events()`

The two callback functions for crypto and compression are registered using the following APIs:

- `cpaCyInstanceSetNotificationCb`
- `cpaDcInstanceSetNotificationCb`

The details of the above functions, parameters, and return values can be found in [Section 6.2, Additional APIs](#)

3.17.2.1.1 Restart Sequence

During the restart sequence, the user space library releases the memory used for rings and other data structures as part of the shutdown and reallocates them when the restart is completed. The process is transparent to the user application, so it can continue to use the same logical instance after reset to submit Intel® QAT service requests. Any memory allocated by the user application for the Intel® QAT service is untouched during device reset.

A typical Heartbeat error use-case is as follows:

1. The driver and the firmware are loaded, initialized, and started.
2. The user-space application registers to receive instance notifications by calling `cpaCyInstanceSetNotificationCb` and `cpaDcInstanceSetNotificationCb`.
3. The management thread monitors for the device's Heartbeat. When a device is unresponsive, a device reset is initiated by the management thread or by the Intel® QAT framework depending on the device configuration.
4. The kernel-space process sends the Restarting event to the user-space process.
5. The user-space driver passes the device restarting event to all the registered application instances. It also frees memory and rings associated with the registered instances.
6. The kernel-space driver triggers the device reset.
7. During reset, the Intel® QAT service requests made by the user application returns one of:
 - `CPA_STATUS_FAIL`
 - `CPA_STATUS_RETRY`
 - `CPA_STATUS_RESTARTING`
8. When the device reset is complete, the kernel-space driver sends a device Restarted event to the user space driver.
9. The user space driver allocates the memory and rings and then forwards the device Restarted event to each of the registered instances.

3.17.2.1.2 Status of Packets in Flight (Crypto Applications Only)

When a device has fatal errors, the application ordinarily cannot determine whether or not inflight requests have been processed successfully.

The current Intel® QAT release includes a dummy response feature that creates mock responses to all requests submitted during a fatal error condition, so the application can detect them and, therefore, know which requests need to be resubmitted to the available devices or to the software.

NOTE: The sequence of dummy responses will match the sending request sequence for all requests submitted during a fatal error.

Since the dummy response feature only supports Public Key Encryption (PKE), dummy responses may be generated only when the `icp_sal_CyPollInstance()` function is called, since it is the function for crypto services.

The `icp_sal_poll_device_events()` function should also be called by the application, so that the application get a notification when the device encounters a failure and dummy responses are generated when calling `icp_sal_CyPollInstance()` for the inflight requests.

3.17.2.1.3 Determining Device ID

The `<device id>` that is passed as a parameter to several Heartbeat API is the numeric suffix of the device name displayed by the following command. (Device name: `qat_dev0`)

```
#service qat_service status
There is 1 QAT acceleration device(s) in the system:
qat_dev0 - type: c3xxx, inst_id: 0, node_id: 0, bsf: 01:00.0, #accel:
3 #engines: 6 state: up
```

The Intel® QAT library has no API to discover the device number easily. However, an application can use the IOCTLs `IOCTL_GET_NUM_DEVICES` and `IOCTL_STATUS_ACCEL_DEV` to find the `device_id` of a particular device if they know the Bus Device Function (BDF). Refer to `perform_query_dev()` in `./adf_ctl.cpp`.

3.17.2.1.4 Setting Polling Minimal Period

QAT driver has possibility to set Heartbeat poll period value inside conf file as `HeartbeatTimer` parameter (see [Table 16](#)):

- `HeartbeatTimer` - minimal acceptable value is 100 [ms], due to limitation on firmware
- If value is not set in config file, default heartbeat pool period value is equal 500 [ms]

Reading Heartbeat value (e.g.: ``cat /sys/kernel/debug/qat_c4xxx_0000\:f4\:00.0/heartbeat``) more frequent than once per Heartbeat poll period time, causes return value equal -1 and Kernel log: "HB poll frequency is higher than configured HB timer".

3.17.3 Testing Heartbeat

Two debug capabilities are available to assist the developers incorporating Heartbeat into their applications:

- Simulation of Heartbeat failure
- System virtual files under `/sys/kernel/debug/`

3.17.3.1 Simulated Heartbeat Failure Configuration

The Heartbeat feature is always enabled in the package. However, a debug capability that simulates device failure can be enabled during the configure step as follows:

```
# ./configure --enable-icp-hb-fail-sim
```

3.17.3.2 Simulating Heartbeat Failure

Simulating Heartbeat failure can be accomplished using two methods:

- Using API `icp_sal_heartbeat_simulate_failure(<device id>)`
- Executing the command:

```
# cat /sys/kernel/debug/<device>/heartbeat_sim_fail
```

3.17.3.3 System Virtual Files

NOTE: The Heartbeat `/sys/kernel/debug` files are associated with the QAT Physical Function (PF).

The Heartbeat feature implements the following system virtual files under the `/sys/kernel/debug/qat_cxxx_<your_device_BDF>/` directory.

Table 12. Heartbeat System Virtual Files

File	Content
<code>heartbeat</code>	0: Device is responsive. -1: Device is NOT responsive.
<code>heartbeat_failed</code>	Number of times the device became unresponsive.
<code>heartbeat_sent</code>	Number of times the control process checked if the device is responsive.

A developer could simulate the Heartbeat management process by running the following script in the background:

```
#!/bin/bash while : do
    cat /sys/kernel/debug/<device>/heartbeat > /dev/null    sleep 1
done
```

3.17.3.4 Heartbeat Polling Frequencies

The application developer should decide on the following two Heartbeat polling frequencies:

- Device Heartbeat monitoring
- Checking for device reset events

3.17.3.4.1 Device Heartbeat Monitoring

Consider the following points when determining the frequency of Heartbeat monitoring:

- Increasing Heartbeat monitoring frequency minimize the customer's system downtime
- However, since device unresponsiveness should be an infrequent event, high frequency Heartbeat monitoring wastes CPU cycles.
- Also, if there are large Intel® QAT service requests that take some time to complete, high frequency Heartbeat monitoring could result in false reports of unresponsiveness.

3.17.3.4.2 Checking for Device Reset Events

If the application uses polling for reading Intel® QAT service responses, there is no value in checking for resets more frequently. Since device unresponsiveness is an infrequent occurrence, frequency of checking for reset events could be a fraction of the frequency of polling for Intel® QAT service responses.

3.18 Handling Device Failures in a Virtualized Environment

The Heartbeat feature in the acceleration software can be used in a virtualized environment. Refer to the *Using Intel® Virtualization Technology (Intel® VT) with Intel® QuickAssist Technology Application Note* (refer to [Table 2](#)) for more details on enabling SR-IOV and the creation of Virtual Functions (VFs) from a single Intel® QuickAssist Technology acceleration device to support acceleration for multiple Virtual Machines (VMs).

The following sequence describes a possible use case for using the Heartbeat feature in a virtualized environment:

5. The Intel® QAT Physical Function driver (PF driver) is loaded, initialized and started.
6. The Intel® QAT Virtual Function driver (VF driver) is loaded, initialized and started in the Guest OS in the VM.

NOTE: For Intel® Communications Chipset 8900 to 8920 Series Software (aka Cave Creek) -
- The PF driver detects that the firmware is unresponsive (using either of the following methods: User Proc Entry Read (not Enabled by Default) or User Application Heartbeat APIs (not Enabled by Default)).

7. The PF driver sends the "Restarting" event message to the VF via the internal PF to VF communication messaging mechanism.
8. The VF driver sends the "Restarting" event to the application's registered callback. The callback is registered using either of the Intel® QAT API functions `cpaDcInstanceSetNotificationCb()` or `cpaCyInstanceSetNotificationCb()` in the Guest OS. The application's callback function may perform any application-level cleanup.
9. The PF driver starts the reset sequence (save state, initiate reset, and restore state).
10. The user restarts the Guest OS and loads the VF driver and application in the Guest OS.

NOTE: If the Heartbeat feature in the acceleration software is not enabled, the PF driver will not notify the VF driver that the firmware is unresponsive.

NOTE: *The error detection mechanisms are not available on the VF driver in the VM, but device errors caused by any of the software running on the VM will be detected by the PF driver using the above mechanisms.*

3.18.1 Understanding System Messages and Warnings

During the operation of Intel® QAT hardware, the system may log various messages that help diagnose configuration and performance issues. One such message is:

```
[ 17.730925] QAT: Could not find a device on node 1
```

This message is informational only, and indicates that a kernel application is attempting to use a QAT device on a specific node, but no QAT device is directly attached to that node. As a result, the application may experience reduced performance due to using a QAT device on a remote node.

This message is not indicative of an error, but rather a potential performance consideration. It is most commonly seen during the early stages of driver loading and the crypto self-test. The message is rate-limited and, as of kernel version 6.3, is logged as a debug message to avoid excessive entries in the system logs.

If you do not observe this message on your platform, it may be due to one of the following reasons:

1. The platform does not have remote nodes, i.e., it is a single-socket system.
2. Kernel tests are running on a core that has a local QAT accelerator attached, avoiding the need for remote node access.
3. The kernel configuration suppresses the printing of informational and debug messages.

For optimal performance, it is recommended to run applications on cores that have local access to QAT devices. Please refer to the system topology and QAT device distribution to ensure proper application and QAT device affinity.

3.19 Incorporating Dummy Responses into an Intel® QAT Application

The dummy response feature has been incorporated in a scenario with the Intel® QAT engine and Nginx*. [Figure 4](#) below illustrates how it works. This can be used as a reference to so-called “software fallback.”

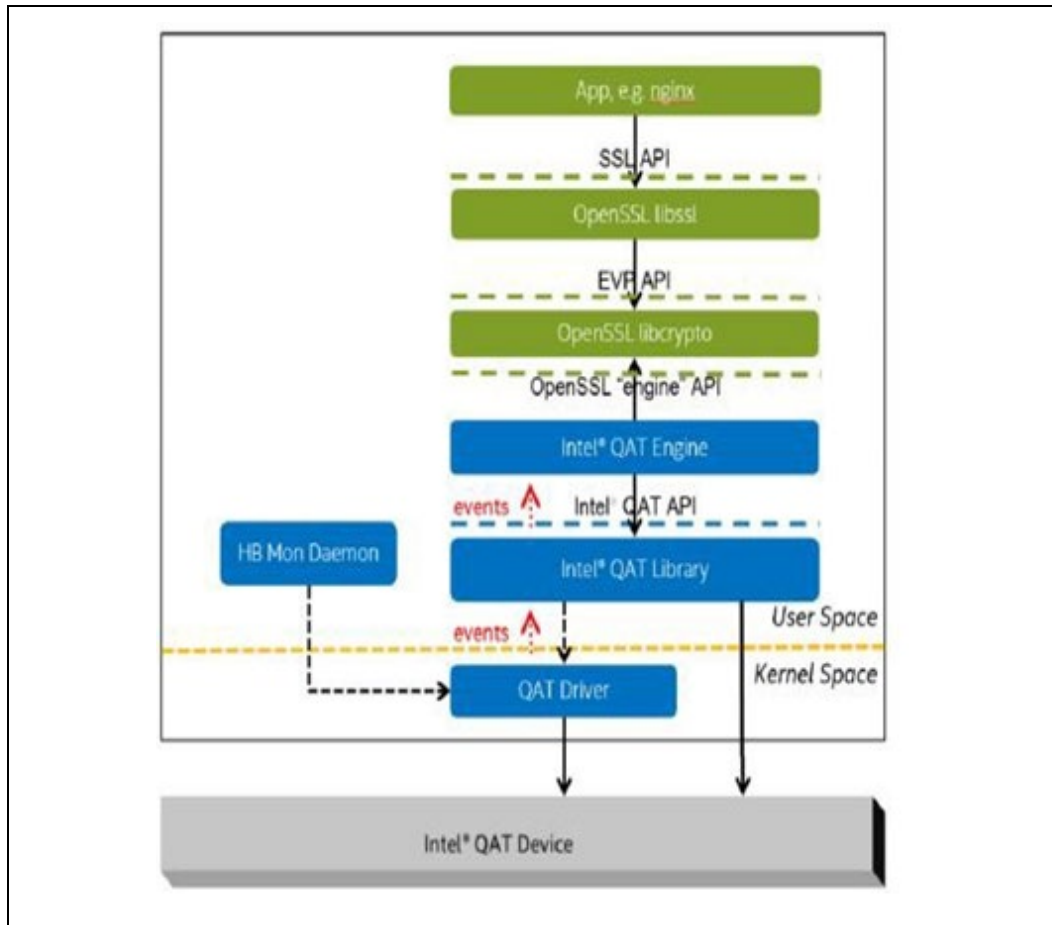
The Intel® QAT engine is a shim layer between OpenSSL* libcrypto* and Intel® QAT Library. The Intel® QAT Library will generate failover responses.

The Heartbeat Monitoring Daemon, a single process, is a daemon which is used to check the device status periodically and trigger the driver to reset the device when Heartbeat failure happens. Its only activity is calling `icp_sal_check_device()` or `icp_sal_check_all_devices()` periodically.

The Intel® QAT Engine polls for and handles “device error” and “device ok” events (via udev). It keeps track of the number of devices which are active.

- If some, but not all, Intel® QAT devices encounter errors, switch to remaining available devices by resubmitting the inflight requests, which are responded to with dummy responses and new requests to the available devices.
- If the number of active Intel® QAT devices goes to zero, switch to software and resubmit the inflight requests which are responded to with dummy responses and new requests to the software.
- If the number of active Intel® QAT devices goes positive again, switch back to hardware.

Figure 4. Incorporating Dummy Responses in an Intel® QAT Operation



3.19.1 Reliability, Availability, Serviceability

The Reliability, Availability, Serviceability (RAS) features are designed to limit the impact of errors within QAT. This section describes the software element required to support the QAT RAS capabilities. As background, the RAS terms are summarized as follows:

- **Reliability:** Refers to how often errors occur in a system and whether the system can recover from an error condition.

- **Availability:** Refers to how flexible the system resources can be allocated or redistributed for the system utilization and system recovery from errors.
- **Serviceability:** Refers to how well the system reports and handles events related to errors.

3.19.2 End to End Data Integrity Support in QAT 1.8:

NOTE: *This End-to-End Data Integrity Support is not available in QAT Hardware Generation 1.7 and earlier devices.*

In QuickAssist Hardware Generation 1.8, additional CRCs have been added to the compression to provide end-to-end data integrity support for performing payload verification throughout the compression pipeline. The CRC for both the input and output data are generated. The Compress and Verify feature supported in previous generations of QAT SW forms part of the overall data integrity feature. The Compress and Verify feature is used to verify that the compressed output from a compression job can be successfully decompressed. The additional CRCs in QAT 1.8 adds protection for the data as it is transferred between Dynamic Random-Access Memory (DRAM) and the QAT and as it flows through the compression processing pipeline.

3.20 Rate Limiting

Rate Limiting is implemented by monitoring the utilization of the device on a per-VF, per-service basis and comparing that to the SLA allocated to that VF and service. Resources are shared across guests and the resource utilization of each guest is measured relative to the capacity of the physical function.

The feature is supported only in rate limiting firmware for cryptographic or compression services.

To enable the Rate Limiting feature:

1. Install the driver package on the host with Single-Root Input/Output Virtualization (SR-IOV) enabled.
2. Update the physical function configure file – depending on your device type set either the `ServicesProfile` parameter to a value that supports rate limiting (e.g., `CRYPTO`, `CUSTOM1`, `COMPRESSION`) or set `RateLimitingEnabled` parameter as 1 to enable the rate limiting.
3. Set `ServicesEnabled` to `cy` or `sym` or `asym` or `dc`.
4. *Perform `qat_service shutdown` and `qat_service start`.*

This procedure also enables Device Utilization measurement (refer to [Section 3.21](#)) Rate limiting requires a virtualized environment, but device utilization can be used without virtualization.

`RateLimitingEnabled` flag is used only for c4xxx driver. For drivers: 200xx, c3xxx, c6xx, d15xx, dh895xcc Rate Limiting feature is enabled by `ServicesProfile` parameter and proper image selection (`CRYPTO`, `COMPRESSION` or `CUSTOM1`),

When a `ServicesProfiles` parameter value is used that supports rate limiting is defined, internal resources are reallocated to administrating Rate Limiting/Device Utilization. This reduces performance for symmetric crypto and data compression by roughly 10%.

3.20.1 Service Level Agreement (SLA)

Service Level Agreement enforcement allocates a specified amount of capacity for a specified service to a specified VF.

Max SLA enforced = (number of VFs) X (number of services) where:

- Number of VFs varies based on device type
- Number of services = 3 (asymmetric or symmetric or compression)

NOTE: *The number of VFs supporting rate limiting is 32 due to firmware limitation.*

3.20.2 SLA Units

SLA units are measured as follows:

- Symmetric Crypto/Compression – 1Mbps of reference operation
- Asymmetric Crypto – 1 operation (ops) of reference operation

NOTE: *Enforced SLAs are rounded up to the next multiple of 1000 units.*

3.20.3 SLA Manager Application

The `sla_mgr` tool is used to create, update, delete, list, and get SLA capabilities.

The SLA Manager executable is available in `$ICP_ROOT/build/sla_mgr` after the package is built and installed using `./configure; make install` commands.

3.20.3.1 Rate Limiting Commands

- Create SLA:
`./sla_mgr create <vf_addr> <rate_in_sla_units> <service>`
- Update SLA:
`./sla_mgr update <pf_addr> <sla_id> <rate_in_sla_units>`
- Delete SLA:
`./sla_mgr delete <pf_addr> <sla_id>`
- Delete all SLAs:
`./sla_mgr delete_all <pf_addr>`
- Query SLA capabilities:
`./sla_mgr caps <pf_addr>`

- Query list of SLAs:
`./sla_mgr list <pf_addr>`

Options:

- `pf_addr` Physical address in bus:device.function(xx:xx.x) format
- `vf_addr` Virtual address in bus:device.function(xx:xx.x) format
- `Service` Asym(=0) or Sym(=1) or Dc(=2)
- `rate_in_sla_units` [0-MAX]. MAX is found by querying the capabilities.

1 rate_in_sla_units is equal to:

- 1 operation per second – for asymmetric service
- 1 Megabits per second – for symmetric service/compression service
- `sla_id` Value returned by `create` command

3.21 DU Manager Application

Device Utilization (DU) is a way to measure utilization of acceleration hardware that corresponds to the throughput of cryptographic or compression services on a given physical or virtual function. This can vary between different device types and generations.

The `du_mgr` tool is used to measure the utilization of cryptographic or compression service for a given physical or virtual function.

The DU execution tool is available in `$ICP_ROOT/build/du_mgr` after the package is built and installed using `./configure; make install` commands.

To enable the Device Utilization feature:

1. Install the driver package on the host with SR-IOV enabled.
2. Update the physical function configure file to set `ServicesProfile` parameter to a value that supports rate limiting (e.g., `CRYPTO`, `CUSTOM1`, or `COMPRESSION`).
3. Set `ServicesEnabled` to `cy` or `sym` or `asym` or `dc`.
4. Perform `qat_service shutdown` and `qat_service start`.

NOTE: When a `ServicesProfiles` parameter value is used that supports rate limiting is defined, internal resources are reallocated to administrating Rate Limiting/Device Utilization. This reduces performance for symmetric crypto or data compression by roughly 10%

NOTE: The maximum SLA that can be set for a device is the maximum DU for that device. For various reasons, the acceptable margin of error for device utilization is 15%; therefore, the tool may report percentages over 100% (allowable range is 85-110%). This margin of error is much greater if durations over 5-seconds are used, as mentioned below.

3.21.1 Commands to Fetch Device Utilization

Start or Stop the device measurement:

```
./du_mgr ( start / stop ) <pf_addr>
```

Query utilization for Physical function:

```
./du_mgr query <pf_addr> <service>
```

Query utilization for Virtual function:

```
./du_mgr query_vf <pf_addr> <vf_addr> <service>
```

Options:

- `pf_addr` Physical address in bus:device.function(xx:xx.x) format
- `vf_addr` Virtual address in bus:device.function(xx:xx.x) format
- `service` Asym(=0) or Sym(=1) or Dc(=2)

3.21.2 Durations

Duration between start and stop commands should be between 5 to 10 seconds.

Duration of more than 10 seconds may give inconsistent query results.

Device utilization `query` and `query_vf` reports utilization between the last start and stop command.

For a given physical or virtual function, the device utilization reported would be in relation to the maximum device capacity.

3.21.3 Reference Algorithm

The Symmetric Crypto Algorithm for Intel® QAT 1.7 devices is AES128-CBC HMACSHA1 with Packet size 1024 bytes.

The Symmetric Crypto Algorithm for Intel® QAT 1.6 devices is AES128-CBC HMACSHA2-256.

The Asymmetric Crypto Algorithm for both systems are RSA with 2048 modulus size.

3.22 Cipher-CRC

Cipher-CRC is a feature that enables offloading of cryptographic processing along with CRC operations to QAT 1.8 device.

This feature is supported **only** by DPDK (Data Plane Development Kit) API and cannot be used with QAT 1.8 package solely. It is supported only in Cipher-CRC firmware for cryptographic (`cy`) service (other services might be used, but cryptographic service is required). Cipher-CRC cannot be used in combination with Rate Limiting feature.

To enable the Cipher-CRC feature:

1. Install DPDK software and QAT 1.8 driver package according to DPDK instructions
2. Update the physical function configuration file by adding `CipherCRCEnabled` parameter in [GENERAL] section and set it to 1 to enable Cipher-CRC.
3. Set `ServicesEnabled` to `cy`. Other services (`inline` or `dc`) might also be set, but `cy` is required for Cipher-CRC feature.
4. Perform `qat_service stop` and `qat_service start`.
5. Follow DPDK instructions to use Cipher-CRC with DPDK API.

NOTE: *In case cryptographic service is not enabled in `ServicesEnabled` the Cipher-CRC is disabled regardless of `CipherCRCEnabled`.*

NOTE: *Cipher-CRC feature cannot be used in combination with Rate Limiting. If both `CipherCRCEnabled` and `RateLimitingEnabled` parameters are set to 1 the device will not start unless Cipher-CRC is disabled due to incorrect platform type or `ServicesEnabled`.*

NOTE: *`CipherCRCEnabled` flag is used **only** for `c4xxx` driver and **only** with DPDK API.*

3.23 Access to Legacy Algorithms

By default, legacy algorithms are now disabled. To enable those algorithms, use the compilation flag `--enable-legacy-algorithms` (Getting Started Guide), which enables all legacy algorithms. Also see associated functions in our Cryptographic API Reference manual: `cpaCyQueryCapabilities()`, `cpaCySymCapabilitiesInfo()`, `cpaCySymQueryCapabilities()`, etc.

The following are the legacy algorithms now disabled by default.

Cipher Algorithms:

- ARC4
- AES-ECB
- AES-F8
- DES-ECB
- DES-CBC
- 3DES-ECB
- 3DES-CBC
- 3DES-CTR
- SM4-ECB

Hash Algorithms:

- MD5
- SHA1
- SHA224
- SHA3_224

PKE Algorithms:

- RSA with key lengths less than 2048 bits

- DSA
- DH
- ECC with curve length less than 256 bits

Table 13. Supported Legacy Algorithms

Cipher Algorithm	QAT1.6	QAT1.7x	QAT1.8	QAT1.9
NULL	Y	Y	Y	Y
ARC4	Opt-in	Opt-in	Opt-in	Opt-in
AES_ECB	Opt-in	Opt-in	Opt-in	Opt-in
AES_CBC	Y	Y	Y	Y
AES_CTR	Y	Y	Y	Y
AES_CCM	Y	Y	Y	Y
AES_GCM	Y	Y	Y	Y
AES_F8	Opt-in	Opt-in	Opt-in	Opt-in
AES_XTS	Y	Y	Y	Y
DES_ECB	Opt-in	Opt-in	Opt-in	Opt-in
DES_CBC	Opt-in	Opt-in	Opt-in	Opt-in
3DES_ECB	Opt-in	Opt-in	Opt-in	Opt-in
3DES_CBC	Opt-in	Opt-in	Opt-in	Opt-in
3DES_CTR	Opt-in	Opt-in	Opt-in	Opt-in
KASUMI_F8	Y	Y	Y	Y
SNOW3G-UEA2	Y	Y	Y	Y
ZUC_EEA3	Y	Y	Y	Y
CHACHA			Y	Y
SM4_ECB			Opt-in	Opt-in
SM4_CBC			Y	Y
SM4_CTR			Y	Y

Hash Algorithm	QAT1.6	QAT1.7x	QAT1.8	QAT1.9
MD5	Opt-in	Opt-in	Opt-in	Opt-in
SHA1	Opt-in	Opt-in	Opt-in	Opt-in
SHA224	Opt-in	Opt-in	Opt-in	Opt-in
SHA256	Y	Y	Y	Y
SHA384	Y	Y	Y	Y
SHA512	Y	Y	Y	Y
SHA3_224			Opt-in	Opt-in

SHA3_256		Y	Y	Y
SHA3_384			Y	Y
SHA3_512			Y	Y
AES_XCBC	Y	Y	Y	Y
AES_CBC_MAC	Y	Y	Y	Y
AES_CCM	Y	Y	Y	Y
AES_GCM	Y	Y	Y	Y
AES_GMAC	Y	Y	Y	Y
AES_CMAC	Y	Y	Y	Y
KASUMI_F9	Y	Y	Y	
SNOW3G_UIA2	Y	Y	Y	Y
ZUC_EIA3	Y	Y	Y	Y
SHAKE_128				
SHAKE_256				
POLY			Y	Y
SM3			Y	Y

PKE	QAT1.6	QAT1.7x	QAT1.8	QAT1.9
RSA-512	Opt-in	Opt-in	Opt-in	Opt-in
RSA-1024	Opt-in	Opt-in	Opt-in	Opt-in
RSA-1536	Opt-in	Opt-in	Opt-in	Opt-in
RSA-2048	Y	Y	Y	Y
RSA-3072	Y	Y	Y	Y
RSA-4096	Y	Y	Y	Y
RSA-8192				
DH	Opt-in	Opt-in	Opt-in	Opt-in
DSA	Opt-in	Opt-in	Opt-in	Opt-in
SM2				Y
ECC key < 256-bit	Opt-in	Opt-in	Opt-in	Opt-in
ECDH Point Multiply		Y	Y	Y
ECDSA Sign		Y	Y	Y
ECDSA Verify		Y	Y	Y
x25519			Y	Y
x448			Y	Y

“Opt-in” means that the algorithm is supported by SW/FW, but is not enabled with the default build configuration. Customers must use the opt-in build flag `--enable-legacy-algorithms` when building the SW library/driver to enable support for these legacy algorithms.

§

4 Acceleration Driver Configuration File

This chapter describes the configuration file(s) that allows the customization of runtime operation. The configuration file(s) must be tuned to meet the performance needs of the target application.

NOTE: *The software package includes a default configuration file, which may not provide optimal performance on all platforms. Consider performance implications as well as the configuration details provided in this chapter if your system requires modifications to the default configuration file.*

4.1 Configuration File Overview

There is a single configuration file for each Intel® QAT Endpoint (and there may be multiple Intel® QAT Endpoints for a given hardware).

NOTE: *Depending on the model number, a device may also contain no Intel® QAT Endpoints.*

The configuration file is split into a number of different sections: a general section and one or more Logical Instance sections.

The **General** section includes parameters that allow the user to specify:

- Which services are enabled?
- Concurrent request default configuration.
- Interrupt coalescing configuration (optional).
- Statistics gathering configuration.

Additional details are included in [Section 4.2, General Section](#).

NOTE: *The concurrent request parameters include both Transmit (**Tx**) and Receive (**Rx**) requests.*

Logical Instances sections (there may be one or more) include parameters that allow the user to set:

- The number of cryptography or data compression instances being managed.
- For each instance, the name of the instance, whether polling is enabled, and the core to which an instance is affinitized.

Additional details are included in [Section 4.3, Logical Instances Section](#).

A sample configuration file is included in the package in the `quickassist/utilities/adf_ctl/conf_files` directory.

Available Sample Configuration per SKUs:

Sample configurations are broadly divided into services and SKUs that Intel® QAT 1.8 platforms can support. Intel® QAT 1.8 offers these services:

- Cryptography ([cy](#))
- Symmetric cryptography ([sym](#))
- Asymmetric cryptography ([asym](#))
- Compression ([dc](#))

4.2 General Section

The General section of the configuration file contains general parameters and statistics parameters.

4.2.1 General Parameters

The [ServicesProfile](#) parameter (see [Table 11](#)) defines the services that are available when the driver loads. For example, if "[ServicesProfile = COMPRESSION](#)" is in the GENERAL section, the compression and decompression are available, along with service chaining, but not cryptography.

NOTE: The [ServicesProfile](#) parameter is used for all drivers excluding *c4xxx*, which uses [RateLimitingEnabled](#).

NOTE: When a [ServicesProfile](#) parameter value is used that supports rate limiting is defined, internal resources are reallocated to administrating Rate Limiting/Device Utilization. This reduces performance by roughly 5%.

Table 14. General Default Configuration Parameters

Service	DEFAULT	CRYPTO	COMPRESSION	CUSTOMI
Asymmetric Crypto	YES	YES		YES
Symmetric Crypto	YES	YES		YES
Hash	YES	YES	YES	YES
Cipher	YES	YES		YES
MGF KeyGen	YES	YES		
SSL/TLS KeyGen	YES	YES		YES
HKDF		YES		YES
Compression	YES		YES	YES
Decompression (stateless)	YES		YES	YES
Decompression (stateful)	YES		YES	

Service	DEFAULT	CRYPTO	COMPRESSION	CUSTOMI
Service Chaining			YES	
Device Utilization		YES	YES	YES
Rate Limiting		YES	YES	YES

NOTE: Set the *ServicesProfile* to determine available features **excluding** *c4xxx*, which uses *RateLimitingEnabled*.

The following table describes the other parameters that can be included in the General section.

Table 15. General Parameters

Parameter	Description	Default	Range
<i>ServicesEnabled</i>	Defines the service(s) available (cryptographic [cy], data compression [dc]), symmetric cryptography only [sym], asymmetric cryptography only [asym] Note: Mutually exclusive with [cy]).	cy;dc	cy, dc, sym and asym Note: Multiple values permitted, use “;” as the delimiter. For exceptions, see Section 4.3.3.2, “Increasing the Maximum Number of Processes/Instances” .
<i>CyNumConcurrentSymRequests</i>	Specifies the number of cryptographic concurrent symmetric requests for cryptographic instances in general.	512	64, 128, 256, 512, 1024, 2048, 4096, 8192, 16384, 32768, or 65536
<i>CyNumConcurrentAsymRequests</i>	Specifies the number of cryptographic concurrent asymmetric requests for cryptographic instances in general.	64	64, 128, 256, 512, 1024, 2048, 4096, 8192, 16384, 32768, or 65536
<i>DcNumConcurrentRequests</i>	Specifies the number of data compression concurrent requests for data compression instances in general.	512	64, 128, 256, 512, 1024, 2048, 4096, 8192, 16384, 32768, or 65536
<i>DcIntermediateBufferSizeIn KB</i>	Specifies the size in KB of each intermediate buffer in on-chip memory for dynamic compression.	64	32 or 64

Parameter	Description	Default	Range
<code>AutoResetOnError</code>	Automatically resets the device in case of fatal error or Heartbeat failure.	0	0 or 1
<code>NumInlineAccelUnits</code>	Define AU number for the inline service	0	Note: Inline feature is only supported on specific packages
<code>NumCyAccelUnits</code>	Define AU number for the crypto service	4,2, or 1 (depends on the SKU)	0 to 6
<code>NumDcAccelUnits</code>	Define AU number for the data compression service	2 or 1 (depends on the SKU)	0 to 6
<code>RateLimitingEnabled</code>	This flag is to enable Rate Limiting	0	0 or 1
<code>HeartbeatTimer</code>	This value set minimal Heartbeat polling period time	500	>= 100

NOTE: Not all parameters listed are available on all device types. `RateLimitingEnabled` parameter is used **only** for `c4xxx` driver and visible **only** inside `c4xxx_.conf.<services>.<SKU>` configuration files.

NOTE: "Default" denotes the value in the configuration file when shipped or the value used if not specified in the configuration file.

For all the services enabled, `NumConcurrentRequests` must be set in the configuration file to one of the following values: 64, 128, 256, 512, 1024, 2048, 4096, 8192, 16384, 32768 and 65536.

The number of concurrent requests registered by the Intel® QAT driver is set to `NumConcurrentRequests - 2`.

This implementation ensures that the request ring will never be full and avoids the need for a Memory Mapped IO (MMIO) read. This implementation maximizes throughput performance.

4.3 Logical Instances Section

This section allows the configuration of logical instances in each address domain (kernel space and individual user space processes).

The address domains are in the following format:

- For the kernel address domain: `[KERNEL]` targeted to Linux* Kernel Crypto Framework (LKCF)
- For the Intel® QAT API in Kernel address domain `[KERNEL_QAT]`
- For user process address domains: `[xxxxxx]`, where `xxxxxx` may be any ASCII value that uniquely identifies the user mode process.

In user space, to allow the driver to configure the logical instances associated with a user process correctly, the process must call the function `icp_sal_userStart` passing the `xxxxxx` string during process initialization. When the user space process is finished, it must call the function `icp_sal_userStop` to free resources. Refer to [Section 6.2.4, User Space Access Configuration Functions](#) for more information.

A single Virtual Function (VF) configured for the SR-IOV use case cannot have both user space instances and kernel space instances. Separate VFs must be created for user space and kernel space.

The `NumProcesses` parameter (in the User Process section) indicates the max number of user space processes within that section name with access to instances on this device. Refer to [Section 6.2.4.2, icp_sal_userStop](#) for more information.

The items that can be configured for a logical instance are:

- The name of the logical instance
- The polling mode
- The core to which the instance is affinitized (optional)

4.3.1 `[KERNEL]` Section

In the `[KERNEL]` section of the configuration file, information about the number and type of kernel instances supporting Linux* Kernel Crypto Framework can be defined.

This section is different from the `[KERNEL_QAT]` section. The `[KERNEL]` section in the configuration file defines instances to register the Intel® QuickAssist Acceleration with Linux* Kernel Crypto Framework (LKCF) while the instances defined in the `[KERNEL_QAT]` section are exclusively targeted to be used with the Intel® QuickAssist API.

LKCF can be used with all devices supported within this software package.

The following table describes the parameters that determine the number of kernel instances for each service.

NOTE: *The maximum number of cryptographic instances supported per Intel® QAT Endpoint is 32; for exceptions, refer to [Section 4.3.3.2, Increasing the Maximum Number of Processes/Instances](#).*

NOTE: *The `NumberDcInstances` is ignored in this section and is set to 0.*

Table 16. [KERNEL] Section Parameters

Parameter	Description	Default	Range
NumberCyInstances	Specifies the number of cryptographic instances. Note: Depends on the number of allocations to other services.	0	0 to 32

4.3.1.1 Enabling Linux* Kernel Crypto Framework (LKCF)

To enable Linux* Kernel Crypto Framework, or LKCF: during the `./configure` step, add the flag `-enable-qat-lkcf`. Also enable at least one `Cy` instance in [KERNEL] section of the configuration file. `Dc` instances are not used in [KERNEL] section.

After installation, to confirm which QAT algorithms were registered with LKCF, run `cat /proc/crypto`, and look for algorithms with their module set to `intel_qat`.

4.3.2 [KERNEL_QAT] Section

The [KERNEL_QAT] section defines instances that can be used by the Intel® QuickAssist API in Kernel space domain.

The table below describes the parameters for the [KERNEL_QAT] section.

NOTE: Intel® QuickAssist API is not supported by Intel® QAT 1.8 devices.

NOTE: The maximum number of cryptographic and data compression instances supported is 32 per Intel® QAT Endpoint; for exceptions, refer to [Section 4.3.3.2, Increasing the Maximum Number of Processes/Instances](#).

Table 17. [KERNEL_QAT] Section Parameters

Parameter	Description	Default	Range
NumberCyInstances	Specifies the number of cryptographic instances. Note: Depends on the number of allocations to other services.	6	0 to 32
NumberDcInstances	Specifies the number of data compression instances. Note: Depends on the number of allocations to other services.	2	0 to 32

1. `NumberCyInstances` depends on the number of allocations to other two services
2. "Default" denotes the value in the configuration file when shipped.

4.3.3 User Process [xxxxx] Sections

There is one [xxxxx] section of the configuration file for each Intel® QAT Endpoint to be configured.

NOTE: Check the SKU information for your specific device to determine how many Intel® QAT Endpoints the device contains. There can be up to three Intel® QAT Endpoints per device.

In each [xxxxx] section of the configuration file, user space access to the Intel® QAT Endpoint can be configured.

The table below shows the parameters in the configuration file that can be set for user process [xxxxx] sections.

Parameters for each user process instance can also be defined. The parameters that can be included for each specific user process instance are like those in [Section 4.3, Logical Instances Section](#).

Table 18. [KERNEL_QAT] Section Parameters

Parameter	Description	Default	Range
NumProcesses	<p>The number of user space processes with section name [xxxxx] that have access to this device.</p> <p>The maximum number of processes that can call <code>icp_sal_userStart</code> and be active at any one time. Refer to Section 6.2.4.1, "icp_sal_userStart" for more information.</p> <p>Caution: Resources are pre-allocated. If this parameter value is set too high, the driver fails to load.</p>	1	<p>For constraints, see Section 4.3.3.1 Maximum Number of Process Calculations.</p> <p>For exceptions, see Section 4.3.3.2, Increasing the Maximum Number of Processes/Instances.</p>
LimitDevAccess	<p>Indicates if the user space processes in this section are limited to only access instances on this Intel® QAT Endpoint.</p>	0	<p>0 (disabled, processes in this section can access multiple Intel® QAT Endpoints) or 1 (enabled, processes in this section can only access this Intel® QAT Endpoint). For additional information, see Section 4.5 Configuring Multiple Processes on a System with Multiple Intel® QAT Endpoints.</p>

Parameter	Description	Default	Range
<code>NumberCyInstances</code>	Specifies the number of cryptographic instances. Note: Depends on the number of allocations to other services.	6	0 to 32. For exceptions, see Section 4.3.3.2, Increasing the Maximum Number of Processes/Instances.
<code>NumberDcInstances</code>	Specifies the number of data compression instances. Note: Depends on the number of allocations to other services.	2	0 to 32

4.3.3.1 Maximum Number of Process Calculations

The `NumProcesses` parameter is the number of user space processes per service within the `[xxxx]` section domain with access to this Intel® QAT Endpoint.

The value to which this parameter can be set is determined by a number of factors, most significantly, the number of cryptography instances and/or data compression instances in the process section. The total number of processes, per service, created by the driver is given by the expression (e.g., for cryptography):

$$(\text{NumProcesses}) \times (\text{NumberCyInstances})$$

In Intel® QAT 1.7 devices, there are 16 ring banks per Intel® QAT Endpoint and a maximum of two cryptography instances and two compression instances per bank. The maximum number of instances per device is 32 for cryptography and 32 for compression. For exceptions, refer to [Section 4.3.3.2, Increasing the Maximum Number of Processes/Instances.](#)

The following code example illustrates the maximum number of possible processes per device in polling mode:

```
NumProcesses = 32
NumberCyInstances = 1
NumberDcInstances = 1
```

4.3.3.2 Increasing the Maximum Number of Processes/Instances

NOTE:

1. One bank is used per Intel® QAT virtual function (VFs).
2. This section only applies when the instances make use of polled mode.

It is possible to increase the number of processes supported by the software. In Intel® QAT 1.7 devices, there are 16 ring banks per Intel® QAT Endpoint where Intel® QAT 1.8 devices have 128 and a maximum of two cryptography instances and two compression instances per bank (or per VF) when the configuration file has `ServicesEnabled` equal to `cy;dc`. However, the maximum number of instances can be increased with the careful selection of the `ServiceEnabled` parameter.

Compression, symmetric cryptography, and asymmetric cryptography each require two rings out of the 16 possible rings for Intel® QAT 1.7 devices versus 8 for Intel® QAT 1.8 devices for a ring bank. By selecting only, the services needed, the number of instances can be increased.

NOTE: *Not all versions of the Intel® QAT software package support the ability to increase the number of processes.*

Here are the variations:

- With `ServicesEnabled` equal to `sym`, only two rings are used for each instance, so for Intel® QAT 1.7 devices, eight instances can be used per bank (or per VF), or 128 instances per Intel® QAT Endpoint. For Intel® QAT 1.8 devices, four instances can be used per bank (or per VF), or 512 instances per Intel® QAT Endpoint. In this case, compression and asymmetric crypto services will not be available.
- With `ServicesEnabled` equal to `asym`, only two rings are used for each instance, so for Intel® QAT 1.7 devices, eight instances can be used per bank (or per VF), or 128 instances per Intel® QAT Endpoint. For Intel® QAT 1.8 devices, four instances can be used per bank (or per VF), or 512 instances per Intel® QAT Endpoint. In this case, compression and symmetric crypto services will not be available.
- With `ServicesEnabled` equal to `cy`, only four rings are used for each instance (two each for asymmetric and symmetric crypto), so for Intel® QAT 1.7 devices, four instances can be used per bank (or per VF), or 64 instances per Intel® QAT Endpoint. For Intel® QAT 1.8 devices, two instances can be used per bank (or per VF), or 256 instances per Intel® QAT Endpoint. In this case, compression services will not be available.
- With `ServicesEnabled` equal to `dc`, only two rings are used for each instance, so for Intel® QAT 1.7 devices, eight instances can be used per bank (or per VF), or 128 instances per Intel® QAT Endpoint. For Intel® QAT 1.8 devices, four instances can be used per bank (or per VF), or 512 instances per Intel® QAT Endpoint. In this case, asymmetric and symmetric crypto services will not be available.
- With `ServicesEnabled` equal to `dc; asym`, only four rings are used for each instance (two each for compression and asymmetric crypto), so for Intel® QAT 1.7 devices, four instances can be used per bank (or per VF), or 64 instances per Intel® QAT Endpoint. For Intel® QAT 1.8 devices, two instances can be used per bank (or per VF), or 256 instances per Intel® QAT Endpoint. In this case, symmetric crypto services will not be available.
- With `ServicesEnabled` equal to `dc; sym`, only four rings are used for each instance (two each for compression and symmetric crypto), so for Intel® QAT 1.7 devices, four instances can be used per bank (or per VF), or 64 instances per Intel® QAT Endpoint. For Intel® QAT 1.8 devices, two instances can be used per bank (or per VF), or 256 instances per Intel® QAT Endpoint. In this case, asymmetric crypto services will not be available.

NOTE: *The `ServicesProfile` parameter value may also need to be changed. See [Section 4.2.1](#).*

4.3.3.3 Configuring Instances for Virtual Functions

To configure the number of instances for a virtual function:

1. Install the driver package on the host with SR-IOV enabled.

2. Update the physical function configuration file to set `ServicesEnabled` (refer to [Section 4.3.3.2, Increasing the Maximum Number of Processes/Instances.](#))
3. Perform `qat_service shutdown` and `qat_service start`.
4. Update the virtual function configuration file to set `ServicesEnabled` (refer to [Section 4.3.3.2, Increasing the Maximum Number of Processes/Instances.](#))
5. Restart `qat_service`.

The value of `ServicesEnabled` in the VF configuration file should be the same as the value of `ServicesEnabled` in the PF configuration file, or a subset of that value as shown in [Table 15](#). For instance, if a PF is configured as `cy`, allowable VF configurations related to that PF can only be `cy`, `asym`, or `sym`. VF device restart will fail if a VF configuration is not allowed for that related PF.

If a VF service is configured to a subset of PF service, the number of VF instances is limited to the number allowed for that PF service as described in [Section 4.3.3.2, Increasing the Maximum Number of Processes/Instances](#). For example, if the PF configuration file has `ServicesEnabled=dc; asym`, only four (not eight) `dc` instances are enabled if the VF is configured for `dc` only.

Table 19. Configuring Physical Functions and Virtual Functions

Configured PF Service	Available VF Services
cy;dc	cy;dc
	cy
	dc
	sym
	asym
	dc;sym
	dc;asym
cy	cy
	sym
	asym
dc;asym	dc;asym
	asym
	dc
dc;sym	dc;sym
	sym
	dc

Configured PF Service	Available VF Services
asym	asym
sym	sym
dc	dc

4.3.4 Cryptographic Logical Instance Parameters

The following table shows the parameters that can be set for cryptographic logical instances.

Table 20. Cryptographic Logical Instance Parameters

Parameter	Description	Default	Range
CyXName	Specifies the name of cryptographic instance number X.	IPSec0 for KERNEL and KERNEL_QAT sections. SSL0 for user section	String (max. 64 characters)
CyXIsPolled	Specifies if cryptographic instance number x works in poll mode, interrupt mode or epoll mode.	0 for kernel space instances 1 for user space instance	0 (interrupt mode) for instances in the KERNEL and KERNEL_QAT sections 1 (poll mode) for instances in the KERNEL_QAT and user space sections 2 (epoll mode eventbased polling mode) for instances in user space section
CyXCoreAffinity	Specifies the core to which the instance should be affinitized.	Varies depending on the value of X.	0 to max. number of cores in the system
CyNumConcurrentSymRequests	Specifies the number of cryptographic concurrent symmetric requests for cryptographic instance X.	512	64, 128, 256, 512, 1024, 2048, or 4096
CyNumConcurrentAsymRequests	Specifies the number of concurrent asymmetric requests for cryptographic instance X.	64	64, 128, 256, 512, 1024, 2048, or 4096

NOTE: “Default” denotes the value in the configuration file when shipped.

4.3.4.1 LKCF-supported algorithms:

See [Supported Algorithms in LKCF](#) for a full list.

4.3.5 Data Compression Logical Instance Parameters

The following table shows the parameters in the configuration file that can be set for data compression logical instances.

NOTE: The maximum number of data compression instances supported is 64.

Table 21. Data Compression Logical Instance Parameters

Parameter	Description	Default	Range
<code>DcXName</code>	Specifies the name of data compression instance number X.	IPComp0	String (max. 64 characters)
<code>DcXIsPolled</code>	Specifies if data compression instance number x works in poll mode, interrupt mode or epoll mode.	0 - kernel space instances 1 - user-space instances	0 (interrupt mode) for instances in the KERNEL and KERNEL_QAT sections 1 (poll mode) for instances in the KERNEL_QAT and user space sections 2 (epoll mode eventbased polling mode) for instances in user space section
<code>DcXCoreAffinity</code>	Specifies the core to which the data compression instance should be affinitized.	Varies depending on the value of X.	0 to max. number of cores in the system
<code>DcXNumConcurrentRequests</code>	The parameter specifies the number of concurrent data requests for compression instance X.	512	64, 128, 256, 512, 1024, 2048, or 4096

NOTE: “Default” denotes the value in the configuration file when shipped.

4.3.6 Setting the Core Affinity Parameter for a Logical Instance

When instances are configured with `IsPolled = 1` (Polling mode), the parameter `CoreAffinity` does not have any impact.

Although not used, it is a valid parameter and applications can query the value using `cpaCyInstanceGetInfo2` (see `coreAffinity` bitmask in `CpaInstanceInfo2`). For example, the sample code affinitizes the thread that uses an instance to the core indicated in `CoreAffinity` the config file for that instance.

For instances configured in Interrupt Mode (`IsPolled = 2` in user space (epoll) and `IsPolled = 1` in kernel space), the value of `CoreAffinity` is used to affinitize the interrupt handler to that core.

4.4 Configuring Multiple Intel® QAT Endpoints in a System

A platform may include more than one Intel® QAT Endpoint. Each device must have its own configuration file. The format and structure of the configuration file is exactly the same for all devices. Consequently, the configuration file for Intel® QAT Endpoint 0, (`c6xx_dev0.conf`, for the Intel® C62x Chipset; `c3xxx_dev0.conf`, for the Intel® Atom® C3000 Processor Family SoC; `d15xx_dev0.conf`, for the Intel® Xeon® Processor D Family), can be cloned for use with other Intel® QAT Endpoints.

All the configuration files are located in the `/etc` folder following the installation of the Intel® QAT package.

Simply make a copy of the file and rename it by changing the `dev0` part of the file name. For example, for a second Intel® C62x Chipset Intel® QAT Endpoint, change the file name to `c6xx_dev1.conf`; for a third Intel® QAT Endpoint, change the Intel® QAT Intel® QAT Endpoint by editing the corresponding configuration file accordingly.

NOTE: *If a configuration file does not exist for an Intel® QAT Endpoint, that endpoint will not start, and an error is displayed indicating that a configuration file was not found.*

To determine the number of Intel® QAT Endpoints in a system, use the `lspci` utility:
`lspci -nn | egrep -e '8086:37c8|8086:19e2|8086:0435|8086:6f54'`

The output from a system with a high-end Intel® C62x Chipset SKU is similar to the following:
`88:00.0 Co-processor [0b40]: Intel Corporation Device [8086:37c8] (rev 03)`
`8a:00.0 Co-processor [0b40]: Intel Corporation Device [8086:37c8] (rev 03)`
`8c:00.0 Co-processor [0b40]: Intel Corporation Device [8086:37c8] (rev 03)`

Then, after the driver is loaded, the user can use the `qat_service` script to determine the name of each Intel® QAT Endpoint and its status. For example:

```
# service qat_service status
qat_dev0 - type: c6xx, inst_id: 0, bsf: 06:00:0, #accel: 5 #engines: 10
state: up qat_dev1 - type: c6xx, inst_id: 1, bsf: 85:00:0, #accel: 5
#engines: 10 state: up qat_dev2 - type: c6xx, inst_id: 2, bsf: 87:00:0,
#accel: 5 #engines: 10 state: up
```

The `qat_service` can start, stop, restart and shutdown each device separately or all Intel® QAT Endpoints together. Refer to [Section 3.6, Managing Intel QuickAssist Technology Endpoints Using qat_service](#) for more information.

Some important configuration file information when using multiple Intel® QAT Endpoints:

- When specifying kernel and user space instances in the configuration file, the `Cy<Number>Name` and `Dc<Number>Name` parameters must be unique in the context of the section name only. For example, it is valid to have a parameter called `Cy0Name` in both a kernel instance section (if supported) and a user instance section in the same configuration file without issue. Also, the parameter names do not need to be unique at a system-wide level. For example, it is valid to have a parameter called `Cy0Name` in both the configuration file for `dev0` and the configuration file for `dev1` without issue.

- For Intel® QAT Endpoints with configuration files that have the same section name (for example, [SSL] and the same data in that section), it is necessary to use the `cpaCyInstanceGetInfo2 ()` function to distinguish between Intel® QAT Endpoints. The `cpaCyInstanceGetInfo2 ()` allows the user of the API to query which Intel® QAT Endpoint a cryptography instance handle belongs to. In addition, for any application domain defined in the configuration files (e.g., [SSL]), a call to `cpaCyGetNumInstances ()` returns the number of cryptography instances defined for that domain across all configuration files. A subsequent call to `cpaCyGetInstances ()` obtains these instance handles.

4.5 Configuring Multiple Processes on a System with Multiple Intel® QAT Endpoints

As an example, consider a system with two Intel® QAT Endpoints where it is necessary to configure two user space sections. One section is identified as `SSL` and the other is identified as Internet Protocol Security (`IPSec`).

- For the `SSL` section, configure eight processes, where each process has access to one acceleration instance.
- For the `IPSec` section, configure one process, with access to eight acceleration instances, four per Intel® QAT Endpoint.

In this scenario, the user space section of the configuration files would look like the following.

For `/etc/c6xx_dev0.conf`:

```
[SSL] #User space section name

NumProcesses=4 # There are 4 user space process with section name SSL with access to this
device

LimitDevAccess=1 # These 4 SSL user space processes only use this device

NumCyInstances=1 # Each process has access to 1 Cy instance on this device

NumDcInstances=0 # Each process has access to 0 Dc instances on this device

# Crypto - User instance #0
Cy0Name = "SSL0"
Cy0IsPolled = 1
Cy0CoreAffinity = 0 # Core affinity not used for polled instance

[IPsec] #User space section name

NumProcesses=1 # There is 1 user space process with section name IPSec with access to this
device

LimitDevAccess=0 # This IPSec user space process may have access to other devices

NumCyInstances=4 # The IPSec process has access to 4 Cy instances on this device
```

```
NumDcInstances=0 # The IPSec process has access to 0 Dc instances on this device
```

```
# Crypto - User instance #0
Cy0Name = "IPSec0"
Cy0IsPolled = 1
Cy0CoreAffinity = 0 # Core affinity not used for polled instance
```

```
# Crypto - User instance #1
Cy1Name = "IPSec1"
Cy1IsPolled = 1
Cy1CoreAffinity = 0 # Core affinity not used for polled instance
```

```
# Crypto - User instance #2
Cy2Name = "IPSec2"
Cy2IsPolled = 1
Cy2CoreAffinity = 0 # Core affinity not used for polled instance
```

```
# Crypto - User instance #3
Cy3Name = "IPSec3"
Cy3IsPolled = 1
Cy3CoreAffinity = 0 # Core affinity not used for polled instance
```

For /etc/c6xx_dev1.conf:

```
[SSL] #User space section name
```

```
NumProcesses=4 # There are 4 user space process with section name SSL with access to this device
```

```
LimitDevAccess=1 # These 4 SSL user space processes only use this device
```

```
NumCyInstances=1 # Each process has access to 1 Cy instance on this device
```

```
NumDcInstances=0 # Each process has access to 0 Dc instances on this device
```

```
# Crypto - User instance #0
Cy0Name = "SSL0"
Cy0IsPolled = 1
Cy0CoreAffinity = 0 # Core affinity not used for polled instance
```

```
[IPsec] #User space section name
```

```
NumProcesses=1 # There is 1 user space process with section name IPSec with access to this device
```

```
LimitDevAccess=0 # This IPSec user space process may have access to other devices
```

```
NumCyInstances=4 # The IPSec process has access to 4 Cy instances on this device
```

```
NumDcInstances=0 # The IPSec process has access to 0 Dc instances on this device
```

```

# Crypto - User instance #0
Cy0Name = "IPSec0"
Cy0IsPolled = 1
Cy0CoreAffinity = 0 # Core affinity not used for polled instance

# Crypto - User instance #1
Cy1Name = "IPSec1"
Cy1IsPolled = 1
Cy1CoreAffinity = 0 # Core affinity not used for polled instance

# Crypto - User instance #2
Cy2Name = "IPSec2"
Cy2IsPolled = 1
Cy2CoreAffinity = 0 # Core affinity not used for polled instance

# Crypto - User instance #3
Cy3Name = "IPSec3"
Cy3IsPolled = 1
Cy3CoreAffinity = 0 # Core affinity not used for polled instance

```

Eight processes (with section name `SSL`) can call the `icp_sal_userStart("SSL")` function to get access to one crypto instance each. One process (with section name `IPSec`) can call the `icp_sal_userStart("IPSec")` function to get access to eight crypto instances.

Internally in the driver, this works as follows:

1. When the driver is configured (that is, the service `qat_service` is called), the driver reads the configuration file for the device and populates an internal configuration table.
2. Reading the configuration file for `dev0`:
 - For the section named `[SSL]`, the driver determines that four processes are required and that these processes limit access to this device only. In this case, the driver creates four internal sections that it labels `SSL_DEVO_INT_0`, `SSL_DEVO_INT_1`, `SSL_DEVO_INT_2` and `SSL_DEVO_INT_3`. Each section is given access to one crypto instance as described.
 - For section name `[IPSec]`, the driver determines that one process is required and that this process does not limit access to this device only (that is, it may access instances on other devices). In this case, the driver creates one internal section that it labels `IPSec_INT_0` and gives this access to four crypto instances on this device.
3. Reading the configuration file for `dev1`:
 - For the section named `[SSL]`, the driver determines that four processes are required and that these processes are limited to access this device only. In this case, the driver creates four internal sections that it labels `SSL_DEV1_INT_0`, `SSL_DEV1_INT_1`, `SSL_DEV1_INT_2` and `SSL_DEV1_INT_3`. Each section is given access to one crypto instance as described.
 - For the section named `[IPSec]`, the driver determines that one process is required and that this process may have access to instances on other

devices. In this case, the driver creates one internal section that it labels IPsec_INT_0 and gives this access to four crypto instances on this device.

NOTE: *This section name now appears in both devices' internal configuration and, therefore, the process that gets assigned this section name will have access to instances on both devices.*

4. In total, there are nine separate sections (`SSL_DEVO_INT_0`, `SL_DEVO_INT_1`, `SSL_DEVO_INT_2`, `SSL_DEVO_INT_3`, `SSL_DEV1_INT_0`, `SL_DEV1_INT_1`, `SSL_DEV1_INT_2`, `SSL_DEV1_INT_3` and `IPsec_INT_0`) with access to crypto instances.

When a process calls the `icp_sal_userStart ("SSL")` function, the driver locates the next available section of the form `SSL_DEV<m>_INT<...>` (of which there are eight in total in this example) and assigns this section to the process. This gives the process access to corresponding crypto instances.

When a process calls the `icp_sal_userStart ("IPsec")` function, the driver locates the next available section of the form `IPsec_INT_<...>` (of which there is only one in total for this example) and assigns this section to the process. This gives the process access to the corresponding crypto instances.

The `icp_sal_userStartMultiProcess()` function has been deprecated. The API still exists, but it simply calls `icp_sal_userStart()`.

4.6 Sample Configuration File

Sample configuration files are available in [quickassist/utilities/adf_ctl/conf_files](#). Depending on the product and configuration, one or more of these will be copied to `/etc` during the package installation.

NOTE: *The previous "v1" configuration file format is not supported.*

§

5 Secure Architecture Considerations

This chapter describes the potential threats identified as part of the secure architecture analysis of the Intel® Quick Assist Technology acceleration complex within the Intel® Communications C62x Chipset family and the actions that can be taken to protect against these threats.

This chapter concentrates on the acceleration complex. First, the terminology covering the main threat categories and mechanisms, attacker privilege and deployment models are presented. Then, some common mitigation actions that can be applied to many of these threat categories and mechanisms are discussed. Finally, more specific threat/attack vectors, including attacks against specific services of the PCH device are described.

5.1 Terminology

Each of the potential threat/attack vectors discussed may be described in terms of the following:

- [Threat Categories](#)
- [Attack Mechanism](#)
- [Attacker Privilege](#)
- [Deployment Models](#)

5.1.1 Threat Categories

System threats can be classified into the categories in the following table.

Table 22. System Threat Categories

Category	Nature of Threat and Examples
Exposure of Data	<ul style="list-style-type: none"> • Attacker reads data to which they should not have read access • Attacker reads cryptographic keys
Modification of Data	<ul style="list-style-type: none"> • Attacker overwrites data to which they should not have write access • Attacker overwrites cryptographic keys
Denial of Service	<ul style="list-style-type: none"> • Attacker causes application or driver software (running on an IA core) to fail or terminate. • Attacker causes Intel® QuickAssist Accelerator firmware to hang, temporarily impeding service

Category	Nature of Threat and Examples
	<ul style="list-style-type: none"> Attacker causes excessive use of resource (IA core, Intel® QuickAssist Accelerator firmware thread, silicon slice, PCIe* bandwidth, and so on), thereby reducing availability of the service to legitimate client

5.1.2 Attack Mechanism

Some of the mechanisms by which an attacker can carry out an attack are listed in the following table.

Table 23. Attack Mechanisms and Examples

Mechanism	Examples
Contrived Packet Stream	Attacker crafts a packet stream that exploits known vulnerabilities in the software, firmware, or hardware. This could include vulnerabilities such as buffer overflow bugs, lack of parameter validation, and so on.
Compromised Application Software	Attacker modifies the application code calling the Intel® QuickAssist Technology API to exploit known vulnerabilities in the driver/hardware.
Application Malware	In an environment where an attacker may be able to run their own application, separate from the main application software, they may invoke the Intel® QuickAssist Technology API to exploit known vulnerabilities in the driver/hardware
Compromised IA driver software	Attacker modifies the IA driver to exploit known vulnerabilities in the driver/hardware.
Defect	It is also possible that the attack is not malicious, but rather an unintentional defect

5.1.3 Attacker Privilege

The following table describes the privileges that an attacker may have. The table describes the case of a non-virtualized system.

Table 24. Attacker Privilege

Privilege	Comments
Physical access	There is no attempt to protect against threats, such as signal probes, where the attacker has physical access to the system. Customers can protect their systems using physical locks, tamper-proof enclosures, Faraday cages, and so on.
Logged in as privileged user	There is no attempt to protect against threats where the attacker is logged in as a privileged user. Customers can protect their systems using strong, frequently changed passwords, and so on.
Logged in as unprivileged user	If the attacker is logged into a platform as an unprivileged user, it is important to ensure that they cannot use the services of the PCH to access (read or write) any data to which they would not otherwise have access.

Privilege	Comments
Ability to send packets	In almost all deployments, attackers have the ability to send arbitrary packets from the network into the system. It is assumed that threats (for example, denial of service attacks) may arrive in this way.

5.1.4 Deployment Models

Some of the possible deployment models are given in the following table.

Table 25. Deployment Models

Deployment Model	Examples
System with no untrusted users	<ul style="list-style-type: none"> • Network security appliance • Server in data center
System with potentially untrusted users	<ul style="list-style-type: none"> • Server in data center

5.2 Threat/Attack Vectors

A thorough analysis has been conducted by considering each of the threat categories, attack mechanisms, attacker privilege levels, and deployment models. As a result, the following threats have been identified. Also described are the steps a user of the PCH chipset can take to mitigate against each threat. Some general practices that mitigate many of the common threats are considered first. Thereafter, threats on specific services and mitigation against those threats are described.

5.2.1 General Mitigation

The following mitigation techniques are generic to different threats and attack vectors:

- Ensure that all software running on the platform that has access to Intel® Quick Assist Technology devices is within the trust boundary of the platform owner. This mitigation includes software running in virtual machines and containers.
- Intel® follows Secure Coding guidelines, including performing code reviews and running static analysis on its driver software and firmware, to ensure its compliance with security guidelines. It is recommended that customers follow similar guidelines when developing application code. This should include the use of tools such as static analysis, fuzzing, and so on.
- Ensure each hardware component, including the PCH chipset, processor, and DRAM, is physically secured from attackers. This can include such examples as physical locks, tamper proofing, and Faraday cages (to prevent side-channel attacks via electromagnetic radiation).
- Ensure that network services not required on the module are not operating and that the corresponding network ports are locked down.

- Use strong passwords to protect against dictionary and other attacks on administrative and other login accounts.

5.2.2 General Threats

General threats include the following:

- [DMA](#)
- [Intentional Modification of IA Driver](#)
- [Modification of the QAT Configuration File](#)
- [Malicious Application Code](#)
- [Denial of Service](#)

5.2.2.1 DMA

Threat: The PCH can perform Direct Memory Access (DMA, the copying of data) between defined memory locations. Once an attacker has sufficient privilege to invoke the Intel® QuickAssist Technology API, or to write to/read from the hardware rings used by the driver to communicate with the device, they can send requests to the Intel® QuickAssist Accelerator to perform such DMA, passing arbitrary physical memory addresses as the source and/or destination addresses, thereby exposing or modifying regions of memory to which they would otherwise not have access.

Mitigation: Ensure that only trusted users are granted permissions to access the Intel® QuickAssist Technology API, or to write to and read from the hardware rings. Specifically, the PCH configuration file describes logical instances of acceleration services and the set of hardware rings to be used for each such instance. User processes can ask the kernel driver to map these rings into their address spaces. To access a given device (identified by the number in the filenames below), the user must be granted read/write access to the following files, which may be in /dev:

- `uio<0..N>` (where "0..N" are the qat uio device numbers)
- `qat*`
- `usdm_drv`

5.2.2.2 Intentional Modification of IA Driver

Threat: An attacker can potentially modify the IA driver to behave maliciously. This may lead to a denial of service of Intel® Quick Assist Technology services.

Mitigation: The driver object/executable file on disk should be protected using the normal file protection mechanisms so that it is writable only by trusted users, for example, a privileged user or an administrator. Specifically, the Intel® QuickAssist Technology kernel objects and libraries should not be writeable by user. If the `qat` user group is being used to provide access to Intel® Quick Assist Technology services, then this group should not have write permission to the binaries.

5.2.2.3 Modification of the QAT Configuration File

Threat: The QAT configuration file is read at initialization time by the driver and specifies what instances of each service (cryptographic, data compression) should be created, and which rings each service instance will use. Modifying this file could lead to denial of service by deleting required instances or could be used to attempt to create additional instances that the attacker could subsequently attempt to access for malicious purposes.

Mitigation: The configuration file should be protected using the normal file protection mechanisms so that it is writable only by trusted users, for example, a privileged user or an administrator.

NOTE: *By default, the configuration file is stored in the /etc directory and may be named something like, c6xxx_dev0.conf. Its default permissions are that it is readable and writeable only by root user and qat group.*

5.2.2.4 Malicious Application Code

Threat: An attacker who can gain access to the Intel® QuickAssist Technology API may be able to exploit the following features of the API:

- Buffers passed to the API have a specified length of up to 32 bits. By specifying excessive lengths, an attacker may be able to cause denial of service by overwriting data beyond the end of a buffer.
- Buffer lists passed to the API consist of a scatter gather list (array of buffers). An attacker may incorrectly specify the number of buffers, causing denial of service due to the reading or writing of incorrect buffers.

Mitigation: Platform management can include the Rate Limiting feature to mitigate against Noisy Neighbors. Only trusted users and applications should be allowed to access the Intel® QuickAssist Technology API, as described in General Mitigations.

5.2.2.5 Denial of Service

Threat: An attacker may construct a service request that does not conform to the specification, resulting in loss of service due to service timeouts, halting of Quick Assist service or undesired platform level conditions.

Mitigation: The current generation of Intel® Quick Assist Technology has been designed for performance, providing direct access to hardware via PCIe* MMIO space. Misuse of hardware registers is to be avoided, and the threat against intentional misuse must be mitigated by ensuring all software on the platform is trusted.

An attacker may attempt to contrive a packet stream that monopolizes the acceleration services, thereby denying service to legitimate users. This may consist of one or more of the following:

- Sending packets that are compressed (for example, using IPComp) or encrypted (for example, using IPsec), thereby reducing the availability of these services to legitimate traffic.
- Sending excessively large packets, causing some latency for legitimate packets.

- Sending small packets at a high packet rate, causing extra bandwidth utilization on the PCI Express* bus connecting the device to the processor.

Mitigation: Proper monitoring of Device Usage (DU) and the construction of Service Level Agreements (SLA) are now available as part of the Rate Limiting feature.

5.2.3 Threats Specific to Cryptographic Service

Threats against the cryptographic service include:

- [Reading of Cryptographic Keys](#)

5.2.3.1 Reading Cryptographic Keys

Threat: Cryptographic keys are stored in DRAM. An attacker who can determine where these are stored could read the DRAM to get access to the keys or could write the DRAM to use keys known by the attacker, thereby compromising the confidentiality of data protected by these keys. Some cryptographic keys have long lives. The impact of an attacker obtaining the key may exist for the lifetime of the key itself.

Mitigation: DRAM is considered inside the cryptographic boundary (as defined by FIPS 140-2). The normal memory protection schemes provided by the Intel® architecture processor and memory controller, and by the operating system, prevent unauthorized access to these memory regions.

§

6 Supported APIs

The supported APIs are described in two categories:

- [Intel® QuickAssist Technology APIs](#)
- [Additional APIs](#)

6.1 Intel® QAT APIs

The platforms described in this manual support the following Intel® QAT API libraries:

- Cryptographic - API definitions are located in: `$ICP_ROOT/quickassist/include/lac`, where `$ICP_ROOT` is the directory where the Acceleration software is unpacked. See the *Intel® QuickAssist Technology Cryptographic API Reference Manual* (refer to [Table 2](#)) for details.
- Data Compression - API definitions are located in: `$ICP_ROOT/quickassist/include/dc`. See the *Intel® QuickAssist Technology Data Compression API Reference Manual* (refer to [Table 2](#)) for details.

Base API definitions that are common to the API libraries are located in: `$ICP_ROOT/quickassist/include`. See also the *Intel® QuickAssist Technology API Programmer's Guide* (refer to [Table 2](#)) for guidelines and examples that demonstrate how to use the APIs.

6.1.1 Intel® QAT API Limitations

The following limitations apply when using the Intel® QAT APIs on the platforms described in this manual:

- For all services, the maximum size of a single perform request is 4 GB.
- For all services, data structures that contain data required by the Intel® QAT Endpoint should be on a 64-byte-aligned address to maximize performance. This alignment helps minimize latency when transferring data from DRAM to an Intel® QAT Endpoint integrated in the PCH device.
- For the key generation cryptographic API, the following limitations apply:

Caution: Secure Sockets Layer (SSL) key generation op-data:

Maximum secret length is 512 bytes
 Maximum `userLabel` length is 136 bytes
 Maximum `generatedKeyLenInBytes` is 248

Caution: Transport Layer Security (TLS) key generation op-data

Secret length must be <128 bytes for TLS v1.0/1.1; <512 bytes for TLS v1.2
`userLabel` length must be <256 bytes
 Maximum seed size is 64 bytes
 Maximum `generatedKeyLenInBytes` is 248 bytes

Caution: Mask Generation Function (MGF) op-data

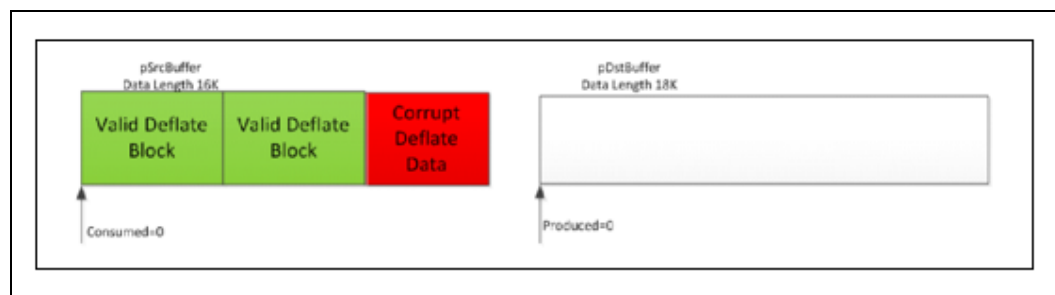
Maximum seed length is 255 bytes
 Maximum `maskLenInBytes` is 65528

- For the cryptographic service, SNOW 3G and KASUMI* operations are not supported when `CpaCySymPacketType` is set to `CPA_CY_SYM_PACKET_TYPE_PARTIAL`. The error returned in this case is `CPA_STATUS_INVALID_PARAM`.
- For the cryptographic service, when using the asymmetric crypto APIs, the buffer size passed to the API should be rounded to the next power of 2, or the next 3- times a power of 2, for optimum performance.
- For the data compression service, the size of all stateful decompression requests have to be a multiple of two with the exception of the last request.
- For the data compression service, the `CpaDcFileType` field in the `CpaDcSessionSetupData` data structure is ignored (previously this was considered for semi-dynamic compression/decompression).
- For static compression, the maximum expansion during compression is ceiling $(9 * \text{Total_Input_Byte} / 8) + 7$ bytes. If `CPA_DC_ASB_ENABLED` is selected, the maximum expansion during compression is the input buffer size plus ceiling $(\text{Total_Input_Byte} / 65535) * 5$ bytes.

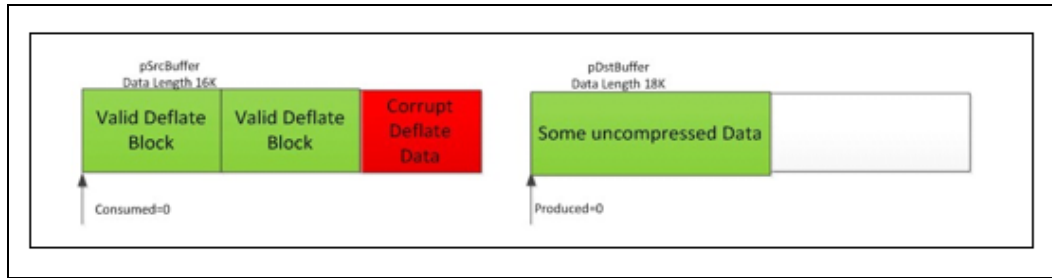
NOTE: Due to the need for a skid pad and the way the checksum is calculated in the stored block case to prevent compression overflow, an output buffer size of ceiling $(9 * \text{Total_Input_Byte} / 8) + 55$ bytes needs to be supplied (even though the stored block output size might be less).

The decompression service can report various error conditions, most of which arise from processing dynamic Huffman code trees that are ill-formed. These soft error conditions are reported at the Intel® QAT API using the `CpaDcReqStatus` enumeration. At the point of soft error, the hardware state will not be accurate to allow recovery. Therefore, in this case, the Intel® QAT software rolls back to the previous known good state and reports that no input has been processed and no output produced. This allows an application to correct the source of the error and resubmit the request.

For example, if the following source and destination buffers were submitted to the Intel® QAT

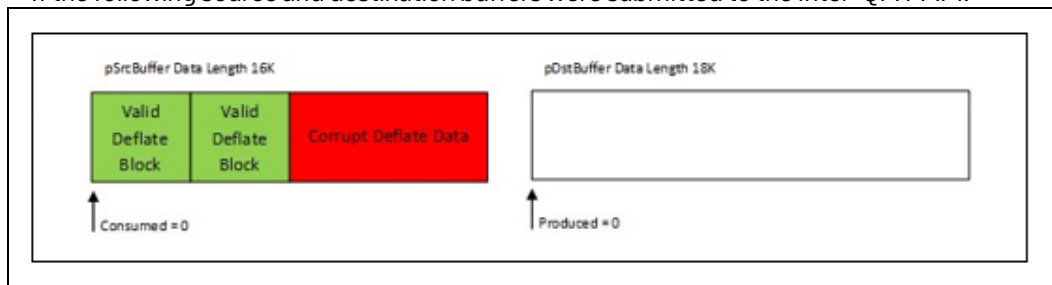


The result would be:

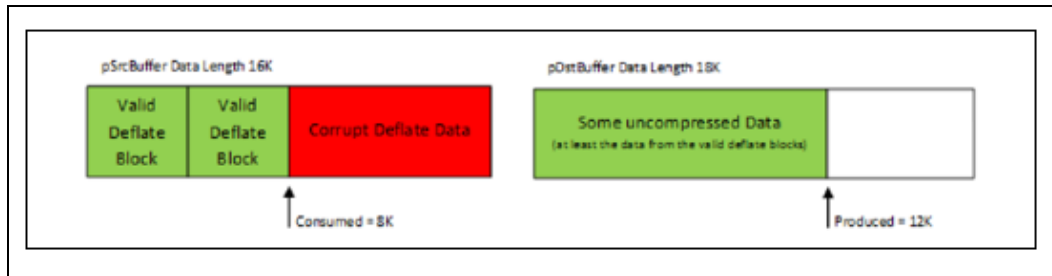


- Behavior when build flag `ICP_DC_RETURN_COUNTERS_ON_ERROR` is defined. In some specialized applications, when a decompression soft error occurs, the application has no way of correcting the source of the error and resubmitting the request. The session will need to be invalidated and terminated. In this case it is more useful to the application to output the uncompressed data up to the point of soft error before terminating the session. There is a compile time build flag (`ICP_DC_RETURN_COUNTERS_ON_ERROR`) to select this mode of operation. This is the behavior of decompression in case of soft error when this build flag is used.

If the following source and destination buffers were submitted to the Intel® QAT API:



The result would be:



It is important to note in this case:

- Caution:** The consumed value returned in the `CpaDcRqResults` structure is not reliable.
- Caution:** No further requests can be submitted on this session.

- For stateful decompression, the maximum output size is 4.29 GB (2³² bytes).

6.1.1.1 Resubmitting After Getting an Overflow Error

The following table describes the behavior of the Intel® QAT compression service when an overflow occurs during a compression or decompression operation.

It describes the expected behavior of an application when an overflow occurs.

Table 26. Compression/Decompression Overflow Behavior

	Operation	Overflow Supported	Input Data Consumed ?	Valid Data Produced?	Status Returned in Results	Note
Traditional API	Stateless compression	YES	Possible - indicated in results consumed field	Possible - indicated in results produced field	-11	Overflow is considered as an exception
	Stateless decompression	NO	NO	NO	-11	Overflow is considered as an error
	Stateful decompression	YES	Possible - indicated in results consumed field	Possible - indicated in results produced field	-11	Overflow is considered as an exception
Data Plane API	Stateless compression	NO	NO	NO	-11	Overflow is considered as an error
	Stateful decompression	NO	NO	NO	-11	Overflow is considered as an error

The Intel® QAT releases enable the Compress and Verify feature by default for compression requests. The Compress and Verify feature imply that sessions can only be **Stateless** in the compression direction.

6.1.1.1 Overflow Exception in the Traditional API

Stateless sessions support overflow as an exception for traditional API in the compression direction only. This means that the application can rely on the `cpaDcRqResults.consumed` to resubmit from where the overflow occurred. An overflow in the decompression direction must be treated as an error.

In this case, the application must resubmit the request with a larger buffer as described in the procedure for handling overflow errors. For stateful sessions, overflow is supported only in the decompression direction.

6.1.1.2 Overflow error in the Data Plane API

The Data Plane API considers overflow status as an error. If an overflow occurs with the data plane API, the driver will output the following error message to the user:

"Unrecoverable error: stateless overflow. You may need to increase the size of your destination buffer"

In this case, `cpaDcRqResults.consumed`, `.produced` and `.checksum` should be ignored. If length and checksum are required, they must be tracked in the application, because they are not maintained in the session.

6.1.1.1.3 Procedure for Handling Overflow Errors

Resubmit the request with the following data:

- Use the same Source buffer.
- Allocate a bigger Destination buffer.
- Put the checksum from the previous successful request into the `cpaDcRqResults` struct.

6.1.1.1.4 Compression Overflow Support in A Virtualized Environment

In a virtual environment, the guest does not download the firmware. Only the host downloads the firmware.

Therefore, if the guest runs a newer Intel® QAT driver than the host, the guest application might experience false CNV errors. The correct course of action would be to update the host with the latest Intel® QAT driver.

6.1.1.1.5 Avoiding a Compression Overflow Exception

Overflow happens for 2 reasons:

1. The application allocated a destination buffer that was too small to receive the compressed data.
2. A recovery occurred after a compress and verified error with an input payload greater than 65,535 bytes.


To minimize the impact of resubmitting data after an overflow exception, the API `cpaDcDeflateCompressBound()` has been added to the Intel® QAT driver. This new API will provide to the application a recommended destination buffer size to avoid the exception. This API must be called by the application before allocating the destination buffer.

The `cpaDcDeflateCompressBound()` API requires the instance handle so that the formula that it uses is tailored to the device generation.

6.1.1.2 Dynamic Compression for Data Compression Service

Dynamic compression involves feeding the data produced by the compression hardware block to the translator hardware block. [Figure 5](#) shows the dynamic compression data path.

Figure 5. Dynamic Compression Data Path



When the application selects the Huffman type to `CPA_DC_HT_FULL_DYNAMIC` in the session and auto-select best feature is set to `CPA_DC_ASB_DISABLED`, the compression service may not always produce a deflate stream with dynamic Huffman trees.

In the case of Stateful decompression requests, if the service returns an exception (e.g., overflow status in the results), it is recommended to examine the bytes consumed and returned in the `CpaDcRqResults` structure to verify if all the data in the source data buffer has been processed. Unprocessed data can be submitted in a subsequent request that uses the offset reported by the `consumed` field in the `CpaDcRqResults` structure.

6.1.1.3 Maximal Expansion with Auto Select Best Feature for Compression

Some input data may lead to a lower-than-expected compression ratio. This is because the input data may not be very compressible.

To achieve a maximum compression ratio, the acceleration unit provides an auto select best (ASB) feature. In this mode, the Intel® QuickAssist Technology hardware will first execute static compression followed by dynamic compression and then select the output that yields the best compression ratio.

However, if the produced data both for dynamic and static operations return a greater value than the uncompressed source data and source block headers, the source data will be used as a stored block.

A 5-byte stored block header is always prepended to the stored block.

To use the ASB feature, configure the `autoSelectBestHuffmanTree` enum during the session creation.

Regardless of the ASB setting selected, dynamic compression will only be attempted if the session is configured for dynamic compression.

There are five possible settings available for the `autoSelectBestHuffmanTree` when creating a session. Based on the ASB settings described below, the produced data returned in the `CpaDcRqResults` structure will vary.

6.1.1.3.1 CPA_DC_ASB_DISABLED

ASB mode is disabled.

6.1.1.3.2 CPA_DC_ASB_STATIC_DYNAMIC

This setting is deprecated. To avoid incompatibility with older applications, it is internally redirected to `CPA_DC_ASB_ENABLED`. Redirecting to `CPA_DC_ASB_ENABLED` effectively means that despite what the old enum name suggests, QAT is now allowed to return an uncompressed block with this option (if it's smaller than compressed block would be).

6.1.1.3.3 CPA_DC_ASB_UNCOMP_STATIC_DYNAMIC_WITH_STORED_HDRS

This setting is deprecated. To avoid incompatibility with older applications, it is internally redirected to `CPA_DC_ASB_ENABLED`.

6.1.1.3.4 CPA_DC_ASB_UNCOMP_STATIC_DYNAMIC_WITH_NO_HDRS

This setting is deprecated. To avoid incompatibility with older applications, it is internally redirected to `CPA_DC_ASB_ENABLED`.

For QAT 1.6/1.7 Hardware, deprecation means it is no longer possible to return an uncompressed stored block without a compliant DEFLATE header.

6.1.1.3.5 CPA_DC_ASB_ENABLED

ASB mode is enabled. When `CPA_DC_ASB_ENABLED` is used, the output will be a format compliant block with a proper header, whether the data is compressed or uncompressed. QAT is allowed to return an uncompressed or compressed (static/dynamic) block, whichever is smaller.

`CPA_DC_ASB_ENABLED` behaves the same as `CPA_DC_ASB_UNCOMP_STATIC_DYNAMIC_WITH_STORED_HDRS` did in previous version of QAT.

6.1.1.4 Maximal Expansion and Destination Buffer Size in Compression Direction

For static compression operations, the worst-case possible expansion can be expressed as:

Max Static Produced data in bytes = $\text{ceil}(9 * \text{Total input bytes} / 8) + 7$

The memory requirement for the destination buffer is expressed by the following formula:

Destination buffer size in bytes = $\text{ceil}((9 * \text{Total input bytes} + (8-1)) / 8) + 55 \text{ bytes} + N \text{ bytes}$

With:

$\text{ceil}(x,y) = (x + (y - 1)) / y$

$N = 8 - (\text{total input byte count})$ when total input byte count < 8

Or

$N = 0$ when total input byte count ≥ 8

The destination buffer size must consider the worst-case possible maximal expansion + 55 bytes + N bytes

Example 1 with an input source size of 111,261 bytes:

```
Memory required for destination buffer = ceil((9 * 111261 + (8 - 1)) / 8)
+ 55 + (111261 < 8 ? (8 - 111261) : 0)
= ceil (125169.5) + 55 + 0
= 125169 + 55 + 0
= 125224 bytes to be allocated
```

Example 2 with a 7-byte input source size:

```
Memory required for destination buffer = ceil((9 * 7 + (8 - 1)) / 8) + 55
+ (7 < 8 ? (8 - 7) : 0)
= ceil (8.75) + 55 + 1
= 8 + 55 + 1
= 64 bytes to be allocated
```

NOTE: Regardless of the ASB settings, the memory must be allocated for the worst case. If an overflow occurs, either using static or dynamic compression, then the returned counters, status, and expected application behavior is as shown per [Table 27](#).

6.1.2 Data Plane APIs Overview

The Intel® QAT Cryptographic API Reference Manual and the Intel® QAT Data Compression API Reference Manual (refer to [Table 2](#)) contain information on the APIs that are specific to data plane applications.

The APIs are recommended for applications that are executing in a data plane environment where the cost of offload (that is, the cycles consumed by the driver sending requests to the hardware) needs to be minimized. To minimize the cost of offload, several constraints have been placed on the APIs. If these constraints are too restrictive for your application, the traditional APIs can be used instead (at a cost of additional IA cycles).

The definition of the Cryptographic Data Plane API's are contained in:
[\\$ICP_ROOT/quickassist/include/lac/cpa_cy_sym_dp.h](#)

The definition of the Data Compression Data Plane APIs is contained in:
[\\$ICP_ROOT/quickassist/include/dc/cpa_dc_dp.h](#)

6.1.2.1 IA Cycle Count Reduction When Using Data Plane APIs

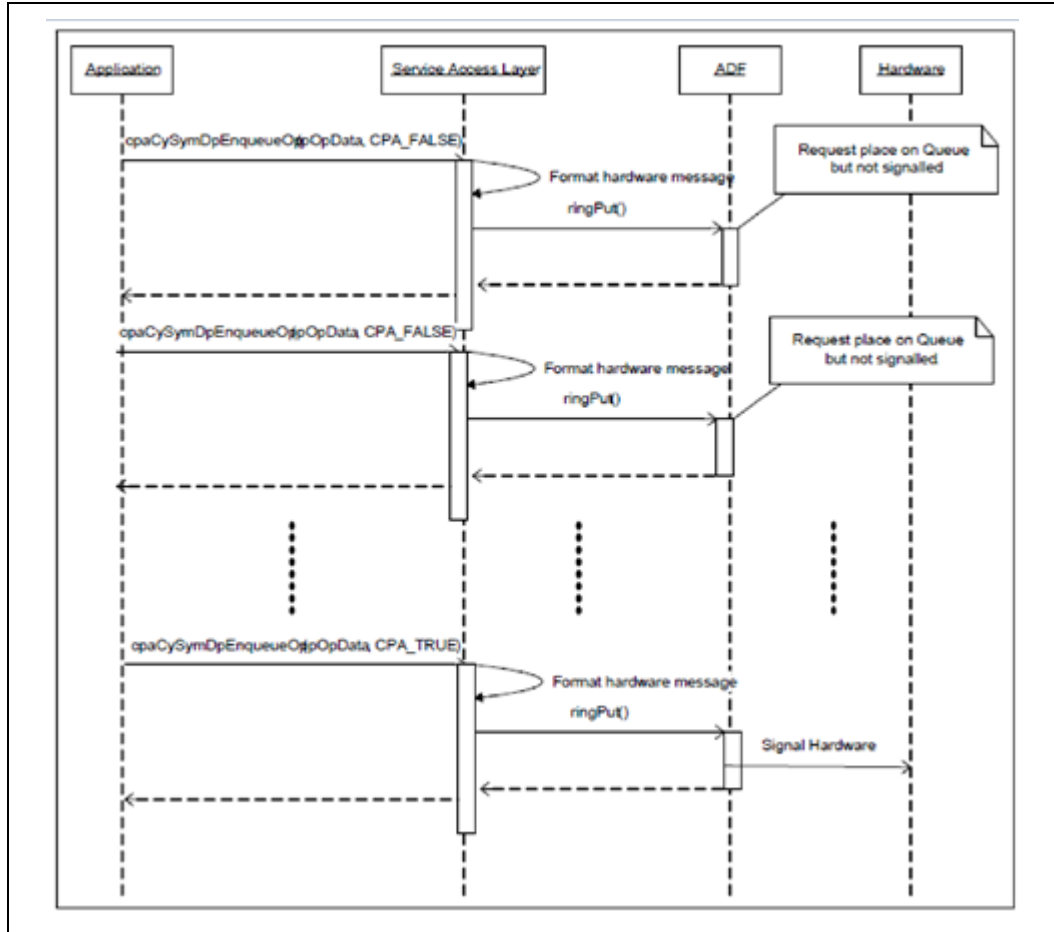
From an IA cycle count perspective, the Data Plane APIs are more performant than the traditional APIs (that is, for example, the symmetric cryptographic APIs defined in [\\$ICP_ROOT/quickassist/include/lac/cpa_cy_sym.h](#)). The majority of the cycle count reduction is realized by the reduction of supported functionality in the Data Plane APIs and the application of constraints on the calling application (refer to [Section 6.1.2.2, Usage Constraints on the Data Plane APIs](#)).

In addition, to further improve performance, the Data Plane APIs attempt to amortize the cost of an MMIO access when sending requests to, and receiving responses from, the hardware.

A typical usage is to call the `cpaCySymDpEnqueueOp()` or the `cpaDcDpEnqueueOp()` function multiple times with requests to process and the `performOpNow` flag set to `CPA_FALSE`. Once multiple requests have been enqueued, the

`cpaCySymDpEnqueueOp()` or `cpaDcDpEnqueueOp()` function may be called with the `performOpNow` flag set to `CPA_TRUE`. This sends the requests to the Intel® QAT Endpoint for processing. This sequence is shown in [Figure 6](#).

Figure 6. Amortizing the Cost of an MMIO Across Multiple Requests



The Intel® QAT API returns a `CPA_STATUS_RETRY` when the ring becomes full.

The number of requests to place on the ring is application dependent and it is recommended that performance testing be conducted with tunable parameter values.

Two functions, `cpaCySymDpPerformOpNow()` and `cpaDcDpPerformOpNow()`, are also provided that allow queued requests to be sent to the hardware without the need for queuing an additional request. This is typically used in the scenario where a request has not been received for some time and the application would like the enqueued requests to be sent to the hardware for processing.

6.1.2.2 Usage Constraints on the Data Plane APIs

The following constraints apply to the use of the Data Plane APIs. If the application can handle these constraints, the Data Plane APIs can be used:

- Thread safety is not supported. Each software thread should have access to its own unique instance ([CpaInstanceHandle](#)) to avoid contention on the hardware rings.
- For performance, polling is supported, as opposed to interrupts (which are comparatively more expensive).
- Polling functions (refer to [Section 6.2.2, Polling Functions](#)) are provided to read responses from the hardware response queue and dispatch callback functions.
- Buffers and buffer lists are passed using physical addresses to avoid virtual-to-physical address translation costs.
- Alignment restrictions are placed on the operation data (that is, the [CpaCySymDpOpData](#) structure) passed to the Data Plane API. The operation data must be at least 8-byte aligned, contiguous, resident, DMA-accessible memory.
- Only asynchronous invocation is supported, that is, synchronous invocation is *not* supported.
- There is no support for cryptographic partial packets. If support for partial packets is required, the traditional Intel® QAT APIs should be used.
- Since thread safety is *not* supported, statistic counters on the Data Plane APIs are not atomic.
- The *default* instance ([CPA_INSTANCE_HANDLE_SINGLE](#)) is not supported by the Data Plane APIs. The specific handle should be obtained using the instance discovery functions ([cpaCyGetNumInstances\(\)](#), [cpaCyGetInstances\(\)](#)).
- The submitted requests are always placed on the high-priority ring.
- The data plane APIs are supported in both user space and polling mode in kernel space, but not supported in interrupt mode in kernel space.

6.1.2.3 Cryptographic and Data Compression API Descriptions

Full descriptions of the Intel® QAT APIs are contained in the Intel® QAT Cryptographic API Reference Manual and the Intel® QAT.

Data Compression API Reference Manual (refer to [Table 2](#)). In addition to the Intel® QAT Data Plane APIs, there are several Data Plane Polling APIs that are described in [Section 6.2.2, Polling Functions](#).

6.1.3 Recovering from a Compress and Verify Error

The Compress and Verify and Recover (CnVnR) feature allow a compression error to be recovered in a seamless manner. It is supported in both the Traditional and in the Data Plane API.

The CnVnR feature is an enhancement of the existing Compress and Verify (CnV) solution. When a compress and verify error is detected, the Intel® QAT software will do a correction without returning a CnV error to the application.

When a recovery occurs, `CpaDcRqResults.status` will return `CPA_DC_OK` or `CPA_DC_OVERFLOW` and the destination buffer will hold valid DEFLATE data.

The application can find out if CnVnR is supported by querying the instance capabilities via the `cpaDcQueryCapabilities` API. On completion, the

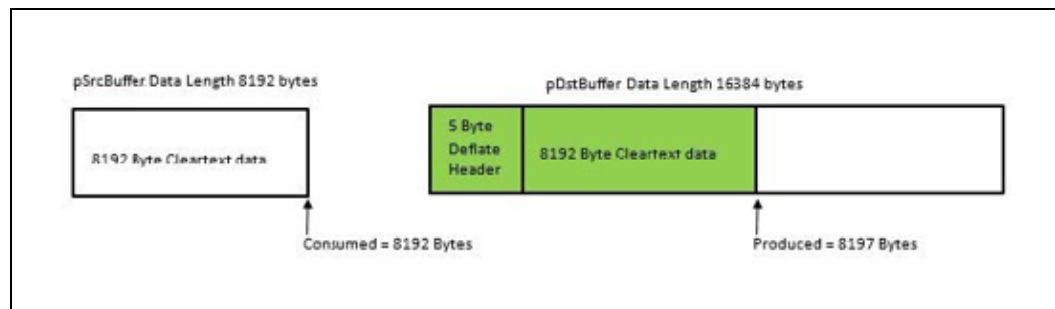
`compressAndVerifyAndRecover` property of the `CpaDcInstanceCapabilities` structure will be set to `CPA_TRUE` if the feature is supported.

The table below provides details on the Intel® QuickAssist APIs supporting the CnVnR feature.

Table 27. API Support for Compress and Verify and Recover

API	CnVnR Behavior
<code>cpaDcCompressData</code>	Enabled by default, no option to disable it.
<code>cpaDcCompressData2</code>	CnVnR is enabled when <code>compressAndVerifyAndRecover</code> property is set to <code>CPA_TRUE</code> in <code>CpaDcOpData</code> structure.
<code>cpaDcDecompressData</code>	Not applicable
<code>cpaDcDecompressData2</code>	Not applicable
<code>cpaDcDpEnqueueOp</code>	CnVnR is enabled when <code>compressAndVerifyAndRecover</code> property is set to <code>CPA_TRUE</code> in <code>CpaDcOpData</code> structure.
<code>cpaDcDpEnqueueOpBatch</code>	CnVnR is enabled when <code>compressAndVerifyAndRecover</code> property is set to <code>CPA_TRUE</code> in <code>CpaDcOpData</code> structure.

When a CnV recovery takes place, the Intel® QAT software creates a stored block out of the input payload that could not be compressed. The maximal size of a stored block allowed by the deflate standard is 65,535 bytes.



When a stored block is created, the DEFLATE header specifies that the data is uncompressed so that the decompressor does not attempt to decode the cleartext data that follows the header. The size of a stored block can be defined as:

Stored block size = Source buffer size + 5 Bytes (used for the deflate header)

If a stored block needs to be created out of a cleartext payload size greater than 65,535 bytes, the Intel® QuickAssist solution creates one stored block of 65,535 bytes and `CpaDcRqResults.status` returns `CPA_DC_OVERFLOW`.

NOTE: If the application uses the Data Plane API, it is responsible for submitting request sizes smaller or equal to 65,530 bytes to avoid meeting the overflow error limit.

6.1.4 Counting Recovered Compression Errors

The Intel® QAT API has been updated to allow the application to track recovered compression errors. The `CpaDcStats` data structure has a new property called `numCompCnvErrorsRecovered` that is incremented every time a compression recovery happens.

The compression recovery process is agnostic to the application.

`CpaDcRqResults.status` returns `CPA_DC_OK` when a compression recovery takes place. The only way to know if a compression recovery took place on the current request is to call the `CpaDcGetStats()` API and to monitor `CpaDcStats.numCompCnvErrorsRecovered`.

6.1.5 Compress and Verify Error log in Sysfs:

The implementation of the Compress and Verify and Recover solution keeps a record of the CnV errors that have occurred since the driver was loaded. The error count is provided on a per Acceleration Engine basis.

The path to the CnV error log is:

```
cat /sys/kernel/debug/qat_dh895xcc_<Bus>\:<device>.<Function>/ cnv_errors
```

Each Acceleration Engine keeps a count of the CnV errors. The CnV error counter is reset when the driver is loaded. The tool also reports the last error type that caused a CnV error.

6.1.6 Supported Algorithms in LKCF

If LKCF is enabled (see section [Enabling Linux* Kernel Crypto Framework \(LKCF\)](#)), the following algorithms and templates are supported:

- `authenc(hmac(sha1),cbc(aes))` - legacy
- `authenc(hmac(sha256),cbc(aes))`
- `authenc(hmac(sha512),cbc(aes))`
- `cbc(aes)`
- `ctr(aes)`
- `xts(aes)`
- `gcm(aes)` - available only with CE release package R4.24.0 and newer, or QAT1.8 driver version 1.11.0, for kernels 4.3.0 and newer
- `rsa`
- `dh` - legacy

For more information on the usage of these algorithms through LKCF, refer to the LKCF API documentation provided by the Linux* kernel, including <https://www.kernel.org/doc/html/v4.15/crypto/architecture.html#ciphers-and-templates>, or <https://www.kernel.org/doc/html/v6.0/crypto/architecture.html#ciphers-and-templates>.

6.2 Additional APIs

There are a number of additional APIs that can serve for optimization and other uses outside of the Intel® QAT services.

NOTE: *Not all additional APIs are supported with all versions of the software package.*

The additional APIs are grouped into the following categories:

- [IOMMU Remapping Functions](#)
 - [Polling Functions](#)
 - [User Space Access Configuration Functions](#)
 - [Version Information Function](#)
 - [Thread-less APIs](#)
 - [Compress and Verify \(CnV\) Related APIs](#)
 - [Heartbeat APIs](#)
 - [Device Polling APIs](#)
1. [Congestion Management APIs](#)
 2. [Service Specific Polling APIs](#)
 3. [Check Device Availability APIs](#)

6.2.1 IOMMU Remapping Functions

These functions are intended for IOMMU remapping operations.

All IOMMU remapping function definitions are in: `$ICP_ROOT/quickassist/lookaside/access_layer/include/icp_sal_iommu.h`.

The IOMMU remapping functions include:

- [Section 6.2.2.1, icp_sal_iommu_get_remap_size](#)
- [Section 6.2.2.2, icp_sal_iommu_map](#)
- [Section 6.2.2.3, icp_sal_iommu_unmap](#)

6.2.1.1 icp_sal_iommu_get_remap_size

Returns the `page_size` rounded for IOMMU remapping.

6.2.1.1.1 Syntax

```
size_t icp_sal_iommu_get_remap_size(size_t size);
```

6.2.1.1.2 Parameters

`size_t` the minimum required page size.

6.2.1.1.3 Return Value

The `icp_sal_iommu_get_remap_size` function returns the `page_size` rounded for IOMMU remapping.

6.2.1.2 icp_sal_iommu_map

Adds an entry to the IOMMU remapping table.

6.2.1.2.1 Syntax

```
CpaStatus icp_sal_iommu_map(Cpa64U phaddr, Cpa64U iova, size_t size);
```

6.2.1.2.2 Parameters

`phaddr` Host physical address.

`iova` Guest physical address.

`size` of the remapped region.

6.2.1.2.3 Return Value

The `icp_sal_iommu_map` function returns one of the following codes:

6.2.1.2.3.1 Code Meaning

`CPA_STATUS_SUCCESS` Successful operation.

`CPA_STATUS_FAIL` Indicates a failure.

6.2.1.3 icp_sal_iommu_unmap

Removes an entry from the IOMMU remapping table.

6.2.1.3.1 Syntax

```
CpaStatus icp_sal_iommu_unmap(Cpa64U iova, size_t size);
```

6.2.1.3.2 Parameters

`iova` Guest physical address to be removed.

`size` Size of the remapped region.

6.2.1.3.3 Return Value

The `icp_sal_iommu_unmap` function returns one of the following codes:

6.2.1.3.3.1 Code Meaning

`CPA_STATUS_SUCCESS` Successful operation.

`CPA_STATUS_FAIL` Indicates a failure.

6.2.1.4 IOMMU Remapping Function Usage

These functions are required when the user wants to access an acceleration service from the Physical Function (PF) when SR-IOV is enabled in the driver. In this case, all I/O transactions from the device go through DMA remapping hardware. This hardware checks 1) if the transaction is legitimate and 2) what physical address the given I/O address needs to be translated to. If the I/O address is not in the transaction table, it fails with a DMA Read error shown as follows:

```
DRHD: handling fault status reg 3
DMAR:[DMA Read] Request device [02:01.2] fault addr <ADDR> DMAR:[fault
reason 06] PTE Read access is not set
```

To make this work, the user must add a 1:1 mapping as follows:

1. Get the size required for a buffer:
`int size = icp_sal_iommu_get_remap_size(size_of_data);`
2. Allocate a buffer:
`char *buff = malloc(size);`
3. Get a physical pointer to the buffer:
`buff_phys_addr = virt_to_phys(buff);`
4. Add a 1:1 mapping to the IOMMU tables:
`icp_sal_iommu_map(buff_phys_addr, buff_phys_addr, size);`
5. Use the buffer to send data to the Intel® QAT Endpoint.
6. Before freeing the buffer, remove the IOMMU table entry:
`icp_sal_iommu_unmap(buff_phys_addr, size);`
7. Free the buffer:
`free(buff);`

The IOMMU remapping functions can be used in all contexts that the Intel® QAT APIs can be used, that is, kernel and user space in a Physical Function (PF) Domain 0, as well as kernel and user space in a Virtual Machine (VM). In the case of VM, the APIs will do nothing. In the PF Domain 0 case, the APIs update the hardware IOMMU tables.

6.2.2 Polling Functions

These functions are intended for retrieving response messages that are on the rings and dispatching the associated callbacks.

All polling function definitions are in:

```
$ICP_ROOT/quickassist/lookaside/access_layer/include/icp_sal_poll.h
```

The polling functions include:

- [Section 6.2.2.1, icp_sal_pollBank](#)
- [Section 6.2.2.2, icp_sal_pollAllBanks](#)
- [Section 6.2.2.3, icp_sal_CyPollInstance](#)
- [Section 6.2.2.4, icp_sal_DcPollInstance](#)
- [Section 6.2.2.5, icp_sal_CyPollDpInstance](#)
- [Section 6.2.2.6, icp_sal_DcPollDpInstance](#)

6.2.2.1 icp_sal_pollBank

Poll all rings on the given Intel® QAT Endpoint on a given bank number to determine if any of the rings contain response messages from the Intel® QAT Endpoint. The `response_quota` input parameter is per ring.

6.2.2.1.1 Syntax

```
CpaStatus icp_sal_pollBank(Cpa32U accelId, Cpa32U bank_number, Cpa32U
response_quota);
```

6.2.2.1.2 Parameters

`accelId` the device number associated with the Intel® QAT Endpoint.

The valid range is 0 to the number of Intel® QAT Endpoint devices in the system.

`bank_number` the number of the memory bank on the Intel® QAT Endpoint that will be polled for response messages. The valid range is 0 to 31.

`response_quota` the maximum number of responses to take from the ring in one call.

6.2.2.1.3 Return Value

The `icp_sal_pollBank` function returns one of the following codes:

6.2.2.1.3.1 Code Meaning

`CPA_STATUS_SUCCESS` Successfully polled a ring with data.

`CPA_STATUS_RETRY` There is no data on any ring on any bank or the banks are already being polled.

`CPA_STATUS_FAIL` Indicates a failure.

6.2.2.2 icp_sal_pollAllBanks

Poll all banks on the given Intel® QAT Endpoint to determine if any of the rings contain response messages from the Intel® QAT Endpoint. The `response_quota` input parameter is per ring.

6.2.2.2.1 Syntax

```
CpaStatus icp_sal_pollAllBanks(Cpa32U accelId, Cpa32U response_quota);
```

6.2.2.2.2 Parameters

`accelId` the device number associated with the Intel® QAT Endpoint. The valid range is 0 to the number of Intel® QAT Endpoints in the system.

`response_quota` the maximum number of responses to take from the ring in one call.

6.2.2.2.3 Return Value

The `icp_sal_pollAllBanks` function returns one of the following codes:

6.2.2.2.3.1 Code Meaning

`CPA_STATUS_SUCCESS` Successfully polled a ring with data.

`CPA_STATUS_RETRY` There is no data on any ring on any bank or the banks are already being polled.

`CPA_STATUS_FAIL` Indicates a failure.

6.2.2.3 icp_sal_CyPollInstance

Poll the Cryptographic (CY) logical instance associated with the `instanceHandle` to retrieve requests that are on response rings associated with that instance and dispatch the associated callbacks. The `response_quota` input parameter is the maximum number of responses to process in one call.

NOTE: The `icp_sal_CyPollInstance()` function is used in conjunction with the `CyXIsPolled` parameter in the acceleration configuration file.

6.2.2.3.1 Syntax

```
CpaStatus icp_sal_CyPollInstance(CpaInstanceHandle instanceHandle, Cpa32U response_quota);
```

6.2.2.3.2 Parameters

`instanceHandle` the logical instance to poll for responses on the response ring.

`response_quota` the maximum number of responses to take from the ring in one call. When set to 0, all responses are retrieved.

6.2.2.3.3 Return Value

The `icp_sal_CyPollInstance` function returns one of the following codes:

6.2.2.3.3.1 Code Meaning

`CPA_STATUS_SUCCESS` The function was successful.

`CPA_STATUS_RETRY` There are no responses on the rings associated with the specified logical instance.

NOTE: *A ring is only polled if it contains data.*

`CPA_STATUS_FAIL` Indicates a failure.

6.2.2.4 icp_sal_DcPollInstance

Poll the Data Compression (DC) logical instance associated with the `instanceHandle` to retrieve requests that are on response rings associated with that instance and dispatch the associated callbacks. The `response_quota` input parameter is the maximum number of responses to process in one call.

NOTE: *The `icp_sal_DcPollInstance()` function is used in conjunction with the `DcXIsPolled` parameter in the acceleration configuration file.*

6.2.2.4.1 Syntax

```
CpaStatus icp_sal_DcPollInstance(CpaInstanceHandle instanceHandle, Cpa32U
response_quota);
```

6.2.2.4.2 Parameters

`instanceHandle` the logical instance to poll for responses on the response ring.

`response_quota` the maximum number of responses to take from the ring in one call. When set to 0, all responses are retrieved.

6.2.2.4.3 Return Value

The `icp_sal_DcPollInstance` function returns one of the following codes:

6.2.2.4.3.1 Code Meaning

`CPA_STATUS_SUCCESS` The function was successful.

`CPA_STATUS_RETRY` There are no responses on the rings associated with the specified logical instance.

NOTE: *A ring is only polled if it contains data.*

`CPA_STATUS_FAIL` Indicates a failure.

6.2.2.5 icp_sal_CyPollDpInstance

Poll a particular Cryptographic (CY) data path logical instance associated with the `instanceHandle` to retrieve requests that are on the high-priority symmetric ring associated with that instance and dispatch the associated callbacks. The `response_quota` input parameter is the maximum number of responses to process in one call.

6.2.2.5.1 Syntax

NOTE: This function is a Data Plane API function and consequently the restrictions in Section 6.1.2.2, “Usage Constraints on the Data Plane APIs” apply.

```
CpaStatus icp_sal_CyPollDpInstance(CpaInstanceHandle instanceHandle,
Cpa32U response_quota);
```

6.2.2.5.2 Parameters

`instanceHandle` the logical instance to poll for responses on the response ring.

`response_quota` the maximum number of responses to take from the ring in one call. When set to 0, all responses are retrieved.

6.2.2.5.3 Return Value

The `icp_sal_CyPollDpInstance()` function returns one of the following codes:

6.2.2.5.3.1 Code Meaning

`CPA_STATUS_SUCCESS` The function was successful.

`CPA_STATUS_RETRY` There are no responses on the rings associated with the specified logical instance.

`CPA_STATUS_FAIL` Indicates a failure.

6.2.2.6 icp_sal_DcPollDpInstance

Poll a particular Data Compression (DC) data path logical instance associated with the `instanceHandle` to retrieve requests that are on the response ring associated with that instance. The `response_quota` input parameter is the maximum number of responses to process in one call.

6.2.2.6.1 Syntax

NOTE: This function is a Data Plane API function and consequently the restrictions in Section 6.1.2.2 apply.

```
CpaStatus icp_sal_DcPollDpInstance(CpaInstanceHandle instanceHandle,
Cpa32U response_quota);
```

6.2.2.6.2 Parameters

`instanceHandle` the logical instance to poll for responses on the response ring.

`response_quota` the maximum number of responses to take from the ring in one call. When set to 0, all responses are retrieved.

6.2.2.6.3 Return Value

The `icp_sal_DcPollDpInstance` function returns one of the following codes:

6.2.2.6.3.1 Code Meaning

`CPA_STATUS_SUCCESS` The function was successful.

`CPA_STATUS_RETRY` There are no responses on the rings associated with the specified logical instance.

`CPA_STATUS_FAIL` Indicates a failure.

6.2.3 User Space Access Configuration Functions

Functions that allow the configuration of user space access to the Intel® QAT services from processes running in user space.

All user space access configuration function definitions are located in `$ICP_ROOT/quickassist/lookaside/access_layer/include/icp_sal_user.h`.

The user space access configuration functions include:

- [Section 6.2.4.1, `icp_sal_userStart`](#)
- [Section 6.2.4.2, `icp_sal_userStop`](#)

6.2.3.1 `icp_sal_userStart`

Initializes user space access to an Intel® QAT Endpoint and starts in the `pProcessName` section in the given section of the configuration file. This function needs to be called before to any call to Intel® QAT API function from the user space process. This function is typically called only once in a user space process.

NOTE: *The `icp_sal_userStartMultiProcess()` function is still supported, but the parameter `limitDevAccess` is ignored because its value is set once in the configuration file and is not allowed to be specified again in the function.*

The configuration format allows the user to create a configuration for many user spaces processes. The driver internally generates unique process names and a valid configuration for each process based on the section name (`pSectionName`) and mode (`limitDevAccess`) provided.

For example, on a system with M number of devices, if all M configuration files contain:

```
[IPSec]
NumProcesses = N LimitDevAccess = 0
```

Then, N internal sections are generated (each with instances on all devices) and N processes can be started at any given time. Each process can call

`icp_sal_userStart("IPSec")` and the driver determines the unique name to use for each process.

Similarly, on an M device system, if all M configuration files contain:

```
[SSL]
NumProcesses = N LimitDevAccess=1
```

Then, M*N internal sections are generated (each with instances on one device only) and M*N processes can be started at any given time. Each process can call `icp_sal_userStart("SSL")` and the driver determines the unique name to use for each process.

Refer to [Section 4.5 Configuring Multiple Processes on a System with Multiple Intel® QAT Endpoints](#) for a detailed example.

6.2.3.1.1 Syntax

```
CpaStatus icp_sal_userStart(const char *pSectionName);
```

6.2.3.1.2 Parameters

`*pSectionName` The section name described in the simplified configuration file format.

`limitDevAccess` Deprecated/ignored.

6.2.3.1.3 Return Value

The `icp_sal_userStart` function returns one of the following codes:

6.2.3.1.3.1 Code Meaning

`CPA_STATUS_SUCCESS` Successfully started user space access to the Intel® QAT Endpoint as defined in the configuration file.

`CPA_STATUS_FAIL` Operation failed.

6.2.3.2 icp_sal_userStop

Closes user space access to the Intel® QAT Endpoint; stops the services that were running and frees the allocated resources. After a successful call to this function, user space access to the Intel® QAT Endpoint from a calling process is not possible. This function should be called once when the process is finished using the Intel® QAT Endpoint and does not intend to use it again.

6.2.3.2.1 Syntax

```
CpaStatus icp_sal_userStop( void);
```

6.2.3.2.2 Parameters

None

6.2.3.2.3 Return Value

The `icp_sal_userStop` function returns one of the following codes:

6.2.3.2.3.1 Code Meaning

`CPA_STATUS_SUCCESS` Successfully stopped user space access to the Intel® QAT Endpoint.

`CPA_STATUS_FAIL` Operation failed.

6.2.4 Version Information Function

A function that allows the retrieval of version information related to the software and hardware being used.

The version information function definition is located in: `$ICP_ROOT/quickassist/lookaside/access_layer/include/icp_sal_versions.h`.

There is only one version information function, that is, `icp_sal_getDevVersionInfo`.

6.2.4.1 `icp_sal_getDevVersionInfo`

Retrieves the hardware revision and information on the version of the software components being run on a given device.

NOTE: The `icp_sal_userStartMultiProcess` (or `icp_sal_userStart`) function must be called before calling this function. If not, calling this function returns `CPA_STATUS_INVALID_PARAM` indicating an error. The `icp_sal_userStartMultiProcess` (or `icp_sal_userStart`) function is responsible for setting up the ADF user space component, which is required for this function to operate successfully.

6.2.4.1.1 Syntax

```
CpaStatus icp_sal_getDevVersionInfo(Cpa32U devId,
icp_sal_dev_version_info_t *pVerInfo);
```

6.2.4.1.2 Parameters

`devId` the ID (number) of the device for which version information is to be retrieved

`*pVerInfo` A pointer to a structure that holds the version information.

6.2.4.1.3 Return Values

The `icp_sal_getDevVersionInfo` function returns one of the following codes:

6.2.4.1.3.1 Code Meaning

`CPA_STATUS_SUCCESS` Operation finished successfully; version information retrieved.

`CPA_STATUS_INVALID_PARAM` Invalid parameter passed to the function.

`CPA_STATUS_RESOURCE` System resource problem.

`CPA_STATUS_FAIL` Operation failed.

6.2.5 Reset Device Function

This API can only be called in user-space.

The device can be reset using this API call. This API call schedules a reset of the device. The device can also be reset using the `adf_ctl` utility, e.g., by calling `adf_ctl qat_dev0 reset`.

6.2.5.1 icp_sal_reset_device

Resets the device.

6.2.5.1.1 Syntax

```
CpaStatus icp_sal_reset_device(Cpa32U accelid);
```

6.2.5.1.2 Parameters

`accelid` the device number.

6.2.5.1.3 Return Value

The `icp_sal_reset_device` function returns one of the following codes:

Code	Meaning
<code>CPA_STATUS_SUCCESS</code>	Successful operation.
<code>CPA_STATUS_FAIL</code>	Indicates a failure.

6.2.6 Thread-Less APIs

These APIs can be used in the user space application.

The thread-less API functions include:

- [Section 6.2.6.1, icp_sal_poll_device_events](#)
- [Section 6.2.6.2, icp_sal_find_new_devices](#)

6.2.6.1 icp_sal_poll_device_events

This reads any pending device events from `icp_dev%d_csr` and forwards to interested subsystems.

6.2.6.1.1 Syntax

```
CpaStatus icp_sal_poll_device_events(void);
```

6.2.6.1.2 Parameters

None

6.2.6.1.3 Return Value

The `icp_sal_poll_device_events` function returns one of the following codes:

6.2.6.1.3.1 Code Meaning

`CPA_STATUS_SUCCESS` Successful operation.

`CPA_STATUS_FAIL` Indicates a failure.

6.2.6.2 `icp_sal_find_new_devices`

This tries to connect to any available devices that the kernel driver has brought up and initialized for use in user space process.

6.2.6.2.1 Syntax

```
CpaStatus icp_sal_find_new_devices(void);
```

6.2.6.2.2 Parameters

None

6.2.6.2.3 Return Value

The `icp_sal_find_new_devices` function returns one of the following codes:

6.2.6.2.3.1 Code	Meaning
------------------	---------

<code>CPA_STATUS_SUCCESS</code>	Successful operation.
---------------------------------	-----------------------

<code>CPA_STATUS_FAIL</code>	Indicates a failure.
------------------------------	----------------------

6.2.7 Compress and Verify (CnV) Related APIs

These APIs can be used in the user space application.

The CnV API functions include:

- [Section 6.2.8.1, `icp_sal_dc_get_dc_error\(\)`](#)
- [Section 6.2.8.2, `icp_sal_dc_simulate_error\(\)`](#)

6.2.7.1 `icp_sal_dc_get_dc_error()`

This API allows the application to return the number of errors that occurred a particular number of times during the lifetime of a process.

6.2.7.1.1 Syntax

```
Cpa64U icp_sal_get_dc_error(Cpa8S dcError);
```

6.2.7.1.2 Parameters

Compression Error code exposed by `CpaDcReqStatus` enum in `cpa_dc.h`

6.2.7.1.3 Return Value

The `icp_sal_get_dc_error()` API returns a 64 bit unsigned integer representing how many times the error type specified by `Cpa8S dcError` occurred in the current process.

6.2.7.2 `icp_sal_dc_simulate_error()`

This API injects a simulated compression error for a defined number of compression or decompression requests. The simulated compression errors can only be applied to the traditional APIs. It must be called prior to the APIs that perform the request.

In the case of a simulated Compress and Verify error for a single request, the application would call `icp_sal_dc_simulate_error()` API as such: `icp_sal_dc_simulate_error(1, CPA_DC_VERIFY_ERROR);`

Followed by a call to:

`CpaDcCompressData()` or `CpaDcCompressData2()`.

To use this API, the driver must be configured and compiled with option `--enabledc-error-simulation`.

6.2.7.2.1 Syntax

```
CpaStatus icp_sal_dc_simulate_error(Cpa8U numErrors, Cpa8S dcError);
```

6.2.7.2.2 Parameters

`Cpa8U numErrors` Number of simulated compression or decompression errors desired.

`Cpa8S dcError` Desired error code to be returned by the compression or decompression API.

6.2.7.2.3 Return Value

The `icp_sal_dc_simulate_error` API returns one of the following codes:

6.2.7.2.3.1 Code Meaning

`CPA_STATUS_SUCCESS` Successful operation.

`CPA_STATUS_FAIL` Indicates that an invalid error type was assigned to `dcError` parameter.

6.2.8 Heartbeat APIs

These APIs check firmware/hardware status for a given device and are used as part of the Heartbeat functionality.

The Heartbeat API functions include:

- [Section 6.2.8.1, icp_sal_check_device\(\)](#)
- [Section 6.2.8.2, icp_sal_check_all_devices\(\)](#)
- [Section 6.2.8.3, icp_sal_heartbeat_simulate_failure\(\)](#)

6.2.8.1 icp_sal_check_device()

This function checks the status of the firmware/hardware for a given device and is used as part of the Heartbeat functionality.

6.2.8.1.1 Syntax

```
CpaStatus icp_sal_check_device(Cpa32U accelID);
```

6.2.8.1.2 Parameters

`accelid` the device ID.

6.2.8.1.3 Return Value

The `icp_sal_check_device` function returns one of the following codes:

6.2.8.1.3.1 Code Meaning

`CPA_STATUS_SUCCESS` Successful operation.

`CPA_STATUS_FAIL` Indicates a failure.

6.2.8.2 icp_sal_check_all_devices()

This function checks the status of the firmware/hardware for all devices and is used as part of the Heartbeat functionality.

6.2.8.2.1 Syntax

```
CpaStatus icp_sal_check_all_devices(void);
```

6.2.8.2.2 Parameters

None

6.2.8.2.3 Return Value

The `icp_sal_check_all_devices` function returns one of the following codes:

6.2.8.2.3.1 Code Meaning

`CPA_STATUS_SUCCESS` Successful operation.

`CPA_STATUS_FAIL` Indicates a failure.

6.2.8.3 icp_sal_heartbeat_simulate_failure()

This function simulates Heartbeat failure for a specific device.

6.2.8.3.1 Syntax

```
CpaStatus icp_sal_heartbeat_simulate_failure(Cpa32U accelID);
```

6.2.8.3.2 Parameters

`Accelid` -- the device ID.

6.2.8.3.3 Return Value

The `icp_sal_heartbeat_simulate_failure` function returns one of the following codes:

6.2.8.3.3.1 Code Meaning

`CPA_STATUS_SUCCESS` Successful operation.

`CPA_STATUS_FAIL` Indicates a failure.

6.2.9 Device Polling APIs

6.2.9.1 icp_sal_poll_device_events()

This function polls for device reset events.

6.2.9.1.1 Syntax

```
CpaStatus icp_sal_poll_device_events(void);
```

6.2.9.1.2 Parameters

None

6.2.9.1.3 Return Value

The `icp_sal_poll_device_events` function returns one of the following codes:

6.2.9.1.3.1 Code Meaning

`CPA_STATUS_SUCCESS` Successful operation.

`CPA_STATUS_FAIL` Indicates a failure.

NOTE: *The events are sent to each instance that has registered a callback function. The callbacks are registered using `cpaCyInstanceSetNotificationCb` and `cpaDcInstanceSetNotificationCb`.*

6.2.9.2 cpaCyInstanceSetNotificationCb

Cryptographic instances use this function to register for device event notifications.

6.2.9.2.1 Syntax

```
CpaStatus cpaCyInstanceSetNotificationCb(
    const CpaInstanceHandle instanceHandle,
    const CpaCyInstanceNotificationCbFunc
    pinstanceNotificationCb,
    void *pCallbackTag);
```

6.2.9.2.2 Parameters

`instanceHandle` Instance handle.

`pinstanceNotificationCb` Instance notification callback function pointer.

`pCallbackTag` Opaque value provided by user.

6.2.9.2.3 Return Values

The `cpaCyInstanceSetNotificationCb()` function returns one of the following codes:

6.2.9.2.3.1 Code Meaning

`CPA_STATUS_SUCCESS` The function was successful.

`CPA_STATUS_FAIL` Indicates a failure.

`CPA_STATUS_INVALID_PARAM` Invalid parameter passed in.

`CPA_STATUS_UNSUPPORTED` Function is not supported.

The signature for the callback function is:

```
typedef void (*CpaCyInstanceNotificationCbFunc)(
    const CpaInstanceHandle instanceHandle,
    void * pCallbackTag,
    const CpaInstanceEvent instanceEvent);
```

6.2.9.2.4 Parameter

```
typedef enum _CpaInstanceEvent
{
    CPA_INSTANCE_EVENT_RESTARTING = 0,
    CPA_INSTANCE_EVENT_RESTARTED,
    CPA_INSTANCE_EVENT_FATAL_ERROR
} CpaInstanceEvent;
```

6.2.9.3 cpaDcInstanceSetNotificationCb

Cryptographic instances use this function to register for device event notifications.

6.2.9.3.1 Syntax

```
CpaStatus cpaDcInstanceSetNotificationCb
    const CpaInstanceHandle instanceHandle,
    const CpaDcInstanceNotificationCbFunc pinstanceNotificationCb,
void *pCallbackTag);
```

6.2.9.3.2 Parameters

`instanceHandle` Instance handle.

`pinstanceNotificationCb` Instance notification callback function pointer.

`pCallbackTag` Opaque value provided by user.

6.2.9.3.3 Return Values

The `cpaDcInstanceSetNotificationCb()` function returns one of the following codes:

6.2.9.3.3.1 Code Meaning

`CPA_STATUS_SUCCESS` The function was successful.

`CPA_STATUS_FAIL` Indicates a failure.

`CPA_STATUS_INVALID_PARAM` Invalid parameter passed in.

`CPA_STATUS_UNSUPPORTED` Function is not supported.

The signature for the callback function is:

```
typedef void (*CpaDcInstanceNotificationCbFunc) (
    const CpaInstanceHandle instanceHandle,
    void * pCallbackTag,
    const CpaInstanceEvent instanceEvent);
```

6.2.9.3.4 Parameter

```
typedef enum _CpaInstanceEvent
{
    CPA_INSTANCE_EVENT_RESTARTING = 0,
    CPA_INSTANCE_EVENT_RESTARTED,
    CPA_INSTANCE_EVENT_FATAL_ERROR
} CpaInstanceEvent;
```

6.2.10 Congestion Management APIs

Congestion Management or Back-pressure mechanism APIs are intended to handle the cases when the device is busy. These APIs ensures there is enough space on the ring before submitting a request.

Applications can query the appropriate ring on each instance and select any instance with enough space without creating any OpData structures.

All these API definitions are located in: `$ICP_ROOT/quickassist/lookaside/access_layer/include/icp_sal_congestion_mgmt.h`.

The Congestion Management APIs include:

- [Section 6.2.10.1, icp_sal_SymGetInflightRequests](#)
- [Section 6.2.10.2, icp_sal_AsymGetInflightRequests](#)
- [Section 6.2.10.3, icp_sal_dp_SymGetInflightRequests](#)

6.2.10.1 icp_sal_SymGetInflightRequests

This function is used to fetch in-flight and max in-flight request counts for the given symmetric instance handle.

6.2.10.1.1 Syntax

```
CpaStatus icp_sal_SymGetInflightRequests (CpaInstanceHandle
instanceHandle,
```

```
Cpa32U *maxInflightRequests,
```

```
Cpa32U *numInflightRequests)
```

6.2.10.1.2 Parameters

`instanceHandle` Symmetric instance handle.

`*maxInflightRequests` A pointer to the max in-flight request count.

`*numInflightRequests` A pointer to the current in-flight request count.

6.2.10.1.3 Return Value

The `icp_sal_SymGetInflightRequests` function returns one of the following codes:

6.2.10.1.3.1 Code Meaning

`CPA_STATUS_SUCCESS` Successfully retrieved the request counts.

`CPA_STATUS_FAIL` Indicates a failure.

`CPA_STATUS_INVALID_PARAM` Invalid parameter.

6.2.10.2 icp_sal_AsymGetInflightRequests

This function is used to fetch in-flight and max in-flight request counts for the given asymmetric instance handle.

6.2.10.2.1 Syntax

```
CpaStatus icp_sal_AsymGetInflightRequests (CpaInstanceHandle
instanceHandle,
```

```
Cpa32U *maxInflightRequests,
```

```
Cpa32U *numInflightRequests)
```

6.2.10.2.2 Parameters

`instanceHandle` Asymmetric instance handle.

`*maxInflightRequests` A pointer to the max in-flight request count.

`*numInflightRequests` A pointer to the current in-flight request count.

6.2.10.2.3 Return Value

The `icp_sal_AsymGetInflightRequests` function returns one of the following codes:

6.2.10.2.3.1 Code Meaning

`CPA_STATUS_SUCCESS` Successfully retrieved the request counts.

`CPA_STATUS_FAIL` Indicates a failure.

`CPA_STATUS_INVALID_PARAM` Invalid parameter.

6.2.10.3 icp_sal_dp_SymGetInflightRequests

This data plane function is used to fetch in-flight and max in-flight request counts for the given symmetric instance handle.

6.2.10.3.1 Syntax

```
CpaStatus icp_sal_dp_SymGetInflightRequests (CpaInstanceHandle
instanceHandle,

Cpa32U *maxInflightRequests,

Cpa32U *numInflightRequests)
```

6.2.10.3.2 Parameters

`instanceHandle` Symmetric instance handle.

`*maxInflightRequests` A pointer to the max in-flight request count.

`*numInflightRequests` A pointer to the current in-flight request count.

6.2.10.3.3 Return Value

The `icp_sal_dp_SymGetInflightRequests` function returns one of the following codes:

6.2.10.3.3.1 Code Meaning

`CPA_STATUS_SUCCESS` Successfully retrieved the request counts.

`CPA_STATUS_FAIL` Indicates a failure.

`CPA_STATUS_INVALID_PARAM` Invalid parameter.

6.2.11 Service Specific Polling APIs

These service specific polling APIs are intended for retrieving response messages that are on the specific ring and dispatching the associated callback.

All these API definitions are located in: `$ICP_ROOT/quickassist/lookaside/access_layer/include/icp_sal_poll.h`.

The Polling APIs include:

- [Section 6.2.11.1, icp_sal_CyPollSymRing](#)
- [Section 6.2.11.2, icp_sal_CyPollAsymRing](#)

6.2.11.1 icp_sal_CyPollSymRing

Poll the symmetric logical instance associated with the `instanceHandle` to retrieve requests that are on the response rings associated with that instance and dispatch the associated callbacks. The `response_quota` input parameter is the maximum number of responses to process in one call.

6.2.11.1.1 Syntax

```
CpaStatus icp_sal_CyPollSymRing(CpaInstanceHandle instanceHandle,
Cpa32U response_quota)
```

6.2.11.1.2 Parameters

`instanceHandle` Instance handle to poll for responses on the response ring.

`response_quota` the maximum number of messages that will be read in one polling. Setting the response quota to zero means that all messages on the ring will be read.

6.2.11.1.3 Return Value

The `icp_sal_CyPollSymRing` function returns one of the following codes:

6.2.11.1.3.1 Code Meaning

`CPA_STATUS_SUCCESS` Successfully polled a ring with data.

`CPA_STATUS_RETRY` There are no responses on the rings associated with the instance.

`CPA_STATUS_FAIL` Indicates a failure.

`CPA_STATUS_INVALID_PARAM` Invalid parameter passed.

`CPA_STATUS_RESTARTING` Device restarting. Resubmit the request.

6.2.11.2 icp_sal_CyPollAsymRing

Poll the asymmetric logical instance associated with the `instanceHandle` to retrieve requests that are on the response rings associated with that instance and dispatch the associated

callbacks. The `response_quota` input parameter is the maximum number of responses to process in one call.

6.2.11.2.1 Syntax

```
CpaStatus icp_sal_CyPollAsymRing(CpaInstanceHandle instanceHandle,  
Cpa32U response_quota)
```

6.2.11.2.2 Parameters

`instanceHandle` Instance handle.

`response_quota` the maximum number of messages that will be read in one poll. Setting the response quota to zero means that all messages on the ring will be read.

6.2.11.2.3 Return Value

The `icp_sal_CyPollAsymRing` function returns one of the following codes:

6.2.11.2.3.1 Code Meaning

`CPA_STATUS_SUCCESS` Successfully polled a ring with data.

`CPA_STATUS_RETRY` There are no responses on the rings associated with this instance.

`CPA_STATUS_FAIL` Indicates a failure.

`CPA_STATUS_INVALID_PARAM` Invalid parameter passed.

`CPA_STATUS_RESTARTING` Device restarting. Resubmit the request.

6.2.12 Check Device Availability APIs

6.2.12.1 icp_sal_userIsQatAvailable

This API allows an application to establish if there is any active QAT device present on system, without calling internal libadf APIs or without a dependency on `icp_sal_userStart()`

6.2.12.1.1 Syntax

```
CpaBoolean icp_sal_userIsQatAvailable(void);
```

6.2.12.1.2 Parameters

None

6.2.12.1.3 Return Value

The `icp_sal_userIsQatAvailable` API returns one of the following codes:

6.2.12.1.3.1 Code Meaning

CPA_TRUE Indicates that there is at least one active device

CPA_FALSE Indicates that there are no active devices

§

7 Application Usage Guidelines

This chapter provides useful guidelines and identifies some of the applications to which the platforms described in this manual are ideally suited.

7.1 Mapping Service Instances to Engines on the Intel® QAT Endpoint

A processor may be connected to one or more Intel® QAT Endpoints. For example, an Intel® Atom® C3000 Processor contains a single integrated Intel® QAT Endpoint, while a single Intel® C620 Series Chipset contains up to three Intel® QAT Endpoints.

Communication between software running on the processor and the Intel® QAT Endpoint is via hardware-assisted rings. Rings are used in pairs; software writes requests onto a request ring and reads responses back from a response ring. The Intel® QAT Endpoint load balances requests from all rings of a given service type across all available hardware "engines" of the corresponding type.

A set of 16 ring banks provides the communication mechanism between a processor and the acceleration complex. Each ring bank contains 16 individual rings for communication.

Intel® provides a software package that abstracts the communication between the host and the rings and presents the high-level Intel® QAT APIs.

7.1.1 Processor and Intel® QAT Endpoint Communication

An acceleration service uses different rings for request and response messages. Communication between the processor and Intel® QAT Endpoint is achieved using the following operations:

- The processor uses a write (PUT) operation to place a request on the request ring.
- The Intel® QAT Endpoint uses a read (GET) operation to retrieve the request from the request ring.
- Once the operation has been performed, the Intel® QAT Endpoint uses a write (PUT) operation to put the response to the response ring.
- The processor uses a read (GET) operation to retrieve the response from the response ring.

7.1.2 Service Instances and Interaction with the Hardware

A ring bank supports two crypto instances and two compression instances. A service instance can be thought of as a channel between an Intel® QAT Endpoint and a core/ thread running on the processor, which uses the rings for communication. The rings are not exposed by an API but are set up using configuration files (one for each Intel® QAT Endpoint).

In general, a service instance uses a pair of rings, one for requests and one for responses. For cryptographic instances, separate request/response pairs are used.

7.1.3 Service Instance Configuration

The configuration of a service instance is done in the configuration file.

The following figure shows an example extract of the relevant section in the configuration file.

Figure 7. Service Instance Configuration

```
##### #
User Space Instances Section

##### #
[proc0]

NumberCyInstances = 1

NumberDcInstances = 0

# Crypto - user space instance #0 Cy0Name="proc0_0"

Cy0IsPolled = 1

Cy0CoreAffinity = 0
```

In the previous figure, the meaning of each numbered item is explained as follows:

- Each named address domain (one domain for the kernel, any number of user space process domains) has its own service instances.
- Specifies a name for the instance.
- Specifies that the instance is using polling.
- Specifies the core affinity for the instance.

7.1.4 Cryptographic Load Balancing Using Multiple Intel® QAT Instances

The application is responsible for load balancing/spreading requests across Intel® QAT Endpoints. Load balancing across the engines computing instances within the Intel® QAT Endpoint is performed by hardware.

In general, the device can be fully utilized from a single instance/ring pair. The main reasons for using multiple instances/ring pairs are:

- Separate software processes each benefit by having their own ring pair to enable the rings to be mapped into the address space of that process.
- Separate threads within a process, possibly on different cores, avoid contention.
- If using interrupts, they can be affinitized from different instances/ring pairs to different cores.

7.2 Cryptography Applications

Cryptography applications supported by the platforms described in this manual include, but are not limited to:

- Virtual Private Networks (VPNs, both IPsec and SSL). Both symmetric and public key cryptography can be offloaded for bulk transfer and key exchange (IKE, SSL handshakes and so on). Refer to [Section 7.2.1, IPsec and SSL VPNs](#) for more information.
- Encrypted Storage. See [Section 7.2.2, Encrypted Storage](#) for more information.
- Web Proxy Appliances. See [Section 7.2.3, Web Proxy Appliances](#).

7.2.1 IPsec and SSL VPNs

Virtual Private Networks (VPNs) allow for private networks to be established over the public Internet by providing confidentiality, integrity, and authentication using cryptography. VPN functionality can be provided by a standalone security gateway box at the boundary between the trusted and untrusted networks. It is also commonly combined with other networking and security functionality in a security appliance, or even in standard routers.

VPNs are typically based on one of two cryptographic protocols, either IPsec or Datagram Transport Layer Security (DTLS). Each has its advantages and disadvantages.

One of the most compute-intensive aspects of a VPN is the cryptographic processing required to encrypt/decrypt traffic for confidentiality, to perform cryptographic hash functionality for authentication, and to perform public key cryptography, based on modular exponentiation of large numbers or elliptic curve cryptography as part of key negotiation and exchange. The PCH provides cryptographic acceleration that can offload this computation from the CPU, thereby freeing up CPU cycles to perform other networking, encryption, or other value-add applications.

The Intel® QAT Endpoint offers its acceleration services through an API, called the Intel® QAT Cryptographic API. This can be invoked from the Linux* kernel or from Linux* user space as well as from other operating systems. Intel® also provides plugins to enable many of the PCH's cryptographic services to be accessed through open-source cryptographic frameworks, such as the Linux* Kernel Crypto Framework/API (also known as the `scatterlist` API) and OpenSSL* libcrypto* (through its EVP API). This facilitates ease of integration with certain open-source implementations of protocol stacks, such as the Linux* kernel's native IPsec stack (called `NETKEY`) or with OpenVPN* (an open source SSL VPN implementation).

7.2.2 Encrypted Storage

In recent years, cases of lost laptops containing sensitive information have made the headlines all too frequently. Full disk encryption has become a standard procedure for many corporate PCs. Safe-guarding critical data, however, is not just a necessity in the client space, it is also a necessity in the data center.

Enterprise-class storage appliances achieve throughput rates in excess of 50 Gbps. Several high-profile cases of data theft have triggered updates to government regulations and industry standards. These regulations/standards now require protection of data-at-rest for applications

involving sensitive data such as medical and financial records, typically using strong encryption. The high computational cost of adding encryption to storage appliances makes offload solutions an attractive value proposition.

Several complimentary standards exist for the encryption of data-at-rest, which, when combined with traditional network security protocols, such as IPsec or SSL/Transport Layer Security (TLS), provide an end-to-end encrypted storage solution, even for data-in-flight.

The IEEE* Security in Storage working group is developing the IEEE 1619 series of standards that deal with cipher algorithms for disk and tape storage devices (AES in CCM and GCM modes). The cryptographic acceleration services of platforms that use the Intel® QAT Endpoints are ideally suited for long-term encrypted storage solutions implementing the IEEE 1619.1 standard, by providing acceleration of the AES-256 cipher in CBC, CCM, and GCM modes and HMAC authentication using SHA-1, SHA-256 and SHA-512 hashes.

The Trusted Computing Group's (TCG) Storage Working Group does not prescribe a particular set of algorithms for the disk encryption. Instead, it defines several Storage Subsystem Classes (SSC) for various usage models, which define services such as enrollment and connection, protected storage (an extension of Trusted Platform Module (TPM)), locking, logging, cryptographic services, authorization, and firmware updates. The cryptographic acceleration services of the platform can help by providing the highest level of encryption for authenticating the host to trusted peripherals implementing the TCG storage standards.

7.2.3 Web Proxy Appliances

Historically, Web Proxy appliances have evolved to present a public or intermediary interface for clients seeking resources from other servers, providing services such as web page caching and load balancing. These appliances are located at the edge of the network, typically at network gateways. Due to their centralized presence in the network, Web Proxy appliances today (referred to with several different names, such as Application Delivery Controllers, Reverse Proxy, and so on) have become a collection of services that include:

- Application Load Balancing (L4-L7)
- SSL Acceleration
- Wide Area Network- (WAN) Acceleration
- Caching
- Traffic Management
- Web Application Firewall

SSL and WAN acceleration have become common place capabilities of the Web Proxy appliance, require computing intensive algorithms for cryptography (SSL) and compression (WAN acceleration). Intel® QAT devices on the platforms described in this manual provide acceleration of asymmetric cryptography (RSA is the most used key negotiation algorithm in SSL), symmetric cryptography (all algorithms defined in the TLS RFCs can be accelerated with the PCH) and compression (DEFLATE algorithm). With the prominence of Web Proxy appliances in typical networks, this use case has applications from cloud computing to small webserver deployments.

7.3 Data Compression Applications

Data compression can be used as part of application delivery networks, data de-duplication, as well as in several crypto applications, for example, VPNs, IDS/IPS and so on.

7.3.1 Compression for Storage

In a time when the amount of online information is increasing dramatically, but budgets for storing that information remain static, compression technology is a powerful tool for improved information management, protection, and access.

Compression appliances can transparently compress data such that clients can keep between two- and five-times more data online and reap the benefit of other efficiencies throughout the data lifecycle. By shrinking the primary data, all subsequent copies of that data, such as backups, archives, snapshots, and replicas are also compressed. Compression is the newest advancement in storage efficiency. Storage compression appliances can shrink primary online data in real time, without performance degradation. Compression can significantly lower storage capital and operating expenses by reducing the amount of data that is stored, and the required hardware that must be powered and cooled.

Compression can help slow the growth of storage, reducing storage costs while simplifying both operations and management. It also enables organizations to keep more data available for use, as opposed to storing data offsite or on harder-to-access media (such as tape).

Compression algorithms are very compute-intensive, which is one of the reasons why the adoption of compression techniques in mainstream applications has been slow. As an example, the DEFLATE Algorithm, which is one of the most used and popular compression techniques today, involves several compute-intensive steps: string search and match, sort logic, binary tree generation, Huffman Code generation. Intel® QAT devices in the platforms described in this manual provide acceleration capabilities in hardware that allow the CPU to offload the compute-intensive DEFLATE algorithm operations, thereby freeing up CPU cycles for other networking, encryption, or other value-add operations.

7.3.2 Data Deduplication and WAN Acceleration

Data Deduplication and WAN Acceleration are coarse-grain data compression techniques centered around the concept of single-instance storage. Identical blocks of data (either to be stored on disk or to be transferred across a WAN link) are only stored/moved once, and any further occurrences are replaced by a reference to the first instance.

While the benefits of deduplication and WAN acceleration obviously depend on the type of data, multi-user collaborative environments are the most suitable due to the amount of naturally occurring replication caused by forwarded emails and multiple (similar) versions of documents in various stages of development.

Deduplication strategies can vary in terms of inline vs post-processing, block size granularity (file-level only, fixed block size or variable block-size chunking), duplicate identification (cryptographic hash only, simple CRC followed by byte-level comparison or hybrids) and duplicate look-up (for example, Bloom filter based index).

Cryptographic hashes are the most suitable techniques for reliably identifying matching blocks with an improbably low risk for false positives, but they also represent the most compute-intensive workload in the application. As such, the cryptographic acceleration services offered by the hardware through the Intel® QAT Cryptographic API can be used to improve the throughput of deduplication/WAN acceleration applications considerably.

Additionally, the compression/decompression acceleration services can be used to further compress blocks for storage on disk, while optionally encrypting the compressed contents.

§

8 Black Box Debug Tool

This chapter provides information on the configuration and use of the Intel® QuickAssist Technology Black Box Debug tool. Information contained includes usage examples, fail signature cases and sample outputs.

8.1 Introduction

8.1.1 Overview

QAT Debug tool was designed to add customer-usable debug solutions that can gather data in order to help with issue diagnosis. It is intended to help the customer to identify root-cause, in a relatively short time and avoid putting significant effort into the whole debugging process.

QAT library does not perform extensive checks or input data validation, which can cause device hangs and other unexpected behavior. Root-causes of these issues are hard to identify without advanced debugging techniques. Using this tool, the customer is given enough information to allow them to find and fix defects caused by probable QAT API misuse. This should be achieved without the QAT-specific technical knowledge required. The QAT Debug tool is released as part of a customer-deployed solution.

8.1.1.1 Security Considerations

QAT Debug tools main aim is to store data captured from traffic generated between the user-space and the QAT device. This useful function could also be considered a potential security risk, given written data could be sensitive in some cases.

Potential sensitive data that is stored by QAT Debug feature:

- Physical addresses of flat buffers (memory at these addresses could contain sensitive data)
- Content descriptor in SYM (symmetric) request that can contain sensitive data
- Members of the QAT group can access stored information contained in logs describing traffic generated by all users inside this group

Potential risk mitigations:

- Sensitive data cleanup function was implemented to be called right after capture
- Storage directories are restricted to QAT group only

8.1.1.2 Performance Considerations

Users should be aware that the QAT Debug feature is able to work in continuous sync mode, which can significantly decrease overall performance when configuration is set to collect all possible data. Types of storage also have a large impact on overall performance. It is recommended to use high-performance storage (high speed NAND memory, RAM-disk, etc.).

8.2 Detailed Description

8.2.1 Collection Data

This feature is intended to collect low-level traffic between the QAT driver and firmware. Collected data is stored in binary file - original form with additional metadata such as: timestamp, process ID of sender and basic slice configuration data (cipher, hash algorithms, data compression type, Huffman tree type).

In order to perform validation of physical addresses and buffer lengths alignments, the content of SGL is additionally captured.

Contents of OP Data can also be captured when the proper log-level is set.

Data collection is supported for:

- QAT driver 'Traditional API':
 - FW (firmware) request descriptors including SGL (Scatter-Gather List), SYM (Symmetric) and DC (Data Compression))
 - FW response descriptors
 - API calls (OP Data content provided by caller)
- QAT driver 'Data Plane API':
 - FW request descriptors including SGL (SYM, DC)
 - FW response descriptors
 - API calls (OP Data content provided by caller)

Data is stored in kernel-managed debug buffers. This approach has been chosen to meet the following requirements:

- The integrity of data must be preserved despite e.g., user-process crash
- Performance degradation should be insignificant
 - Any additional sys-call during request preparation or responses parsing can degrade performance significantly
- Debug buffers should be available either from user-space or kernel (to support e.g., QAT kernel API)

What is not supported:

- Kernel API
- LKCF (Linux* Kernel Cryptography Framework)

8.2.1.1 Data Synchronization

QAT Debug can operate in two synchronization modes:

- Continuous data synchronization

- Event triggered synchronization (crash dump)

8.2.1.1.1 Continuous Synchronization

This optional feature is intended to perform an ongoing data synchronization with persistent storage. The data stored in debug buffers is dumped to continuous synchronization files immediately after receiving the following events from user-space application or kernel module:

- Debug buffer is released
- Debug buffer is full
- User-process crashes (event caught by kernel module)

NOTE: This mode is introducing additional performance degradation closely related to disk performance


8.2.1.1.2 Crash-Dump Mode

This option is intended to use only debug buffers while handling traffic and to dump contents of buffers to persistent storage, only if any of error events occur. When a buffer is full, it is replaced with an empty buffer or a buffer with less recent data.

8.2.1.1.3 Data Collecting Architecture

The high-level design for data collection is presented in the figure below:

Figure 8. Data Collection Architecture



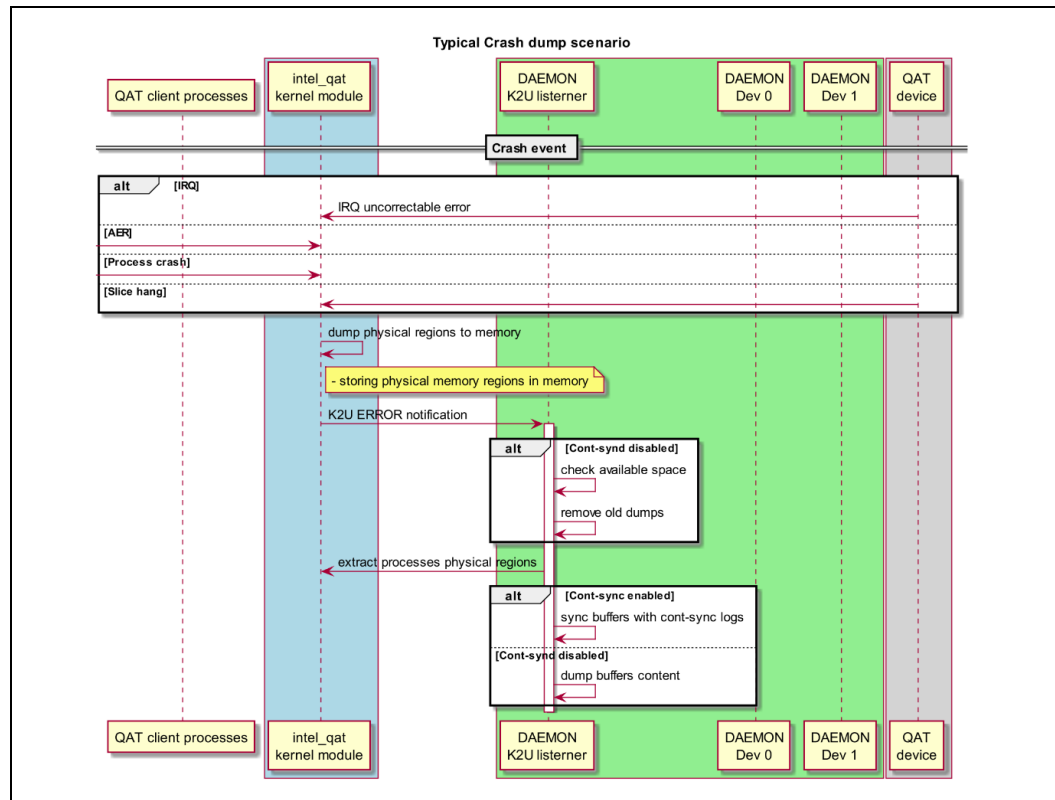
8.2.1.2 Handling QAT Error Events

The following 'error events' are supported:

- IRQ based event caught by QAT kernel driver
- AER (Advance Error Reporting) depending on platform configuration
- Firmware error response (including Slice hang)
- Slice hang caught as an IRQ
- Process crash – event connected to 'orphan ring cleaner'

The handling of error events is presented in the diagram below:

Figure 9. Typical Crash Dump Scenario



8.2.2 Post-Processing

This post-processing tool provides the following utilities:

Audits:

- Physical address used in FW request and SGLs
- Return codes in FW responses
- FLAT buffers and SGL buffers lengths based on cipher algorithm

Listings:

- Lists all collected entries sorted by sent/extraction time

Triggers:

- Manual trigger to dump content of debug buffers to configured location

8.2.2.1 Physical Addresses Audit

NOTE: Usage example is available in Section [8.5.2](#)

The audit is based on:

- Memory map regions - mapped to user space process collected now of 'error event'

NOTE: Huge pages are supported

- Buffers overlapping test
- Basic null checks

User space process memory map example:

```
[root@silpixa00400507 qat_logs]# cat
proc.mmaps.dev00_0000_4d_00_0 | grep 43927
43927:0x0000001cd9c00000:4194304
43927:0x0000001c8f400000:2097152
43927:0x0000001c8f200000:2097152
43927:0x0000001c8f000000:2097152
43927:0x0000001c8ee00000:2097152
43927:0x0000001c8ec00000:2097152
```

Command line interface:

```
#qat_dbg_report command=audit_phy_addresses path=<path>
[dev=<dev>] [bdf=<bdf>]
```

- Path: path to
 - Crash dump directory
 - Continuous sync directory
- Dev: device ID
- Bdf: domain, bus, device, function of a device in a hexadecimal format of 0000:00:00.0

NOTE: Dev or bdf option required only in case of analyzing cont-sync data

8.2.2.2 Cipher Lengths Audit

NOTE: Usage example is available in Section [8.5.3](#)

The audit is based on:

- Cipher algorithm extracted from session and stored in 'content type' field
- Lengths of input buffers (flat buffers and in SGLs)

Command line interface:

```
#qat_dbg_report command=audit_fields_lengths path=<path>
[dev=<dev>] |[bdf=<bdf>]
```

- Path: path to
 - Crash dump directory
 - Continuous sync data directory
- Dev: device ID
- Bdf: domain, bus, device, function of a device in a hexadecimal format of 0000:00:00.0

NOTE: *Dev or bdf option required only in case of analyzing cont-sync data*

8.2.2.3 Return Codes Audit

NOTE: *Usage example is available in Section [8.5.4](#)*

The audit is based on:

- return codes in FW responses (sym/asym/dc)

Command line interface:

```
#qat_dbg_report command= audit_ret_codes path=<path>
[dev=<dev>] |[bdf=<bdf>]
```

Caution: path: path to
Caution: Crash dump directory
Caution: Continuous sync data directory

- dev: device ID
- bdf: domain, bus, device, function of a device in a hexadecimal format of 0000:00:00.0

NOTE: *Dev or bdf option required only in case of analyzing cont-sync data*

NOTE: *In case of any error found – audit tries to match response with corresponding request and prints both to output.*

8.2.2.4 Listing Collected Data in 'Human Readable Form'

This option is intended to display the collected data in human readable form. It can be useful while investigating certain issues. Entries are sorted in descending order according to timestamp.

Command line interface:

```
#qat_dbg_report command=list [path=<path>] [dev=<dev>][bdf=<bdf>]
[<last=<last>]
```

- Path: path to
- Caution:** Crash dump directory
- Caution:** Continuous sync data directory

- Dev: device ID
- bdf: is a domain, bus, device, function of a device in a hexadecimal format of 0000:00:00.0

NOTE: *Dev or bdf option required only in case of analyzing cont-sync data*

- Last: prints last several packets restricted by <last> entities

Example:

```
Entry [REQUEST SYM]: Time-stamp: 2020-11-10 13:58:24.211579144
    Bank: 0 Ring: 2 [2] PID: 43784
        [0.1B] Crypto command ID:
ICP_QAT_FW_LA_CMD_CIPHER_HASH [2]
    [0.2B] Service type: ICP_QAT_FW_COMN_REQ_CPM_FW_LA [4]
    [1.0-1B] LA BULK (SYMMETRIC CRYPTO) COMMAND FLAGS
(0x2c)
        [1.12] ZUC_3G_PROTO: 0
        [1.11] GCM_IV_LEN_FLAG: 0
        [1.10] DIGEST_IN_BUFFER: 0
        [1.7-9] PROTO: 0
        [1.6] CMP_AUTH: 0
        [1.5] RET_AUTH: 1
        [1.4] UPDATE_STATE: 0
        [1.3] CIPH_AUTH_CFG_OFFSET_FLAG: 1
        [1.2] CIPH_IV_FLD_FLAG: 1
        [1.0-1] PARTIAL FLAGS: 0 (FULL)
    [1.2B] Common Request flags: 0x1
        SGL[1] CD_IN [0] BNP [0]
    [1.3B] Extended Symmetric Crypto Command Flags: 0
    [2-3] Content Descriptor (CD) Param Pointer:
0xe864e4540
        [4.2B] Content Descriptor Param Size: 8 [Quad words]
    [6-7] Opaque Data: 0x7fcd7c2e2040
    [8-9] Source phy_addr: 0xe864e4c00
    [10-11] Destination phy_addr: 0xe864e4c00
    [12] Source length: 0
    [13] Destination length: 0
    [14-19] Cipher Request Parameters:
        [14] uint32_t::cipher_offset: 24
        [15] uint32_t::cipher_length: 64
        [16-17] uint64_t::cipher_IV_ptr:
0xadbcefabebafeca
        [18-19] uint64_t::resrvd1: 0x459113d88f8cade
    [27-28.0B] Cipher Request Control Header:
        [27.0B] uint8_t::cipher_state_sz: 2
        [27.1B] uint8_t::cipher_key_sz: 2
        [27.2B] uint8_t::cipher_cfg_offset: 18
        [27.3B] uint8_t::next_curr_id: 0x21 (curr_id:
1, next: 2)
```

```

[28.0B] uint8_t::cipher_padding_sz: 0
[20-26] Authentication Request Parameters:
[20]   uint32_t::auth_off: 0
[21]   uint32_t::auth_len: 88
[22-23] uint64_t::aad_addr/APS: 0xe864e46e0
[24-25] uint64_t::auth_res_addr: 0xe864e5058
[26.0B] uint8_t::aad_sz/inner_prefix_sz: 0
[26.1B] uint8_t::resrvd1: 0
[26.2B] uint8_t::hash_state_sz: 0
[26.3B] uint8_t::auth_res_sz: 0
[27-31] Authentication Request Control Header:
[27]   uint32_t::resrvd1: 0x21120202
[28.0B] uint8_t::resrvd2: 0x0
[28.1B] uint8_t::hash_flags: 0x0
[28.2B] uint8_t::hash_cfg_offset: 46
[28.3B] uint8_t::next_curr_id: 0x42 (curr_id:
2, next: 4)

[29.0B] uint8_t::resrvd3: 0x0
[29.1B] uint8_t::outer_prefix_offset: 0
[29.2B] uint8_t::final_sz: 12
[29.3B] uint8_t::inner_res_sz: 20
[30.0B] uint8_t::resrvd4: 0x0
[30.1B] uint8_t::inner_state1_sz: 24
[30.2B] uint8_t::inner_state2_offset: 3
[30.3B] uint8_t::inner_state2_sz: 24
[31.0B] uint8_t::outer_config_offset: 0
[31.1B] uint8_t::outer_state1_sz: 0
[31.2B] uint8_t::outer_res_sz: 0
[31.3B] uint8_t::outer_prefix_offset: 0
SGL Data:
Source SGL contains 1 flat buffer(s):
[0] Flat buffer: len: 100 phy_addr:
e864e5000
Destination SGL contains 1 flat buffer(s):
[0] Flat buffer: len: 100 phy_addr:
e864e5000

```

8.3 Installation

8.3.1 Hardware and Software Compatibility

Hardware:

- LBG Intel® C62x Chipset
- Intel® Atom® C3000 processor product family
- Intel® QuickAssist Adapter 8960/Intel® QuickAssist Adapter 8970 (formerly known as 'Lewis Hill')
- Intel® Communications Chipset 8925 to 8955 Series
- Intel® C4xxx Series QAT

Supported Operating System:

- As per 4.16 release of the 1.7 QAT Linux* driver

8.3.2 Installing the Driver

NOTE: User must have root privileges to perform the following:

Step 1 - copy package onto the system

Step 2 - extract package:

```
# mkdir /root/QAT
# cd /root/QAT
# tar -xzf <path_to>/QAT<hw_version>.<sw_version>.tar.gz
```

Step 3 - setup the environment to build driver:

```
# ./configure --enable-icp-qat-dbg
```

Step 4 - build and install driver:

```
# make install
```

NOTE: Successful build should end up with message similar to the following:

```
Checking status of all devices.
There is 3 QAT acceleration device(s) in the system:
  qat_dev0 - type: c6xx, inst_id: 0, node_id: 0, bsf:
0000:3d:00.0, #accel: 5 #engines: 10 state: up
  qat_dev1 - type: c6xx, inst_id: 1, node_id: 0, bsf:
0000:3f:00.0, #accel: 5 #engines: 10 state: up
  qat_dev2 - type: c6xx, inst_id: 2, node_id: 1, bsf:
0000:da:00.0, #accel: 5 #engines: 10 state: up
```

Step 5 - change access permissions to kernel debug sysfs directory (for non-root usage):

```
# chmod o+rx /sys/kernel/debug
```

8.3.3 Compiling and Executing Performance Sample Code

Step 1 - build application:

```
# make sample-all
```

Step 2 - install sample applications:

```
# make sample-install
```

Step 3 - run sample code sanity check:

```
# cpa_sample_code signOfLife=1
```

NOTE: Tool execution should end with following message:

```
Sample code completed successfully.
```

8.3.4 Uninstalling the Driver

Step 1 - bring down the driver:

```
# adf_ctl down
```

Step 2 - uninstall driver:

```
# cd /root/QAT/
# make uninstall
```

8.4 Configuration

8.4.1 Configuration via QAT Device Configuration Files

QAT Debug may be configured via dedicated section in QAT device configuration file. The listing below shows example configuration of the debug feature:

```
#####
# QAT Debuggability Section
# Debug levels description:
# 0: no data collection
# 1: API calls data collection
# 2: FW calls data collection
# 3: combined level 1 and 2
#####
[DEBUG]
Enabled = 0
DebugLevel = 2
NumBuffers = 128
BufferSizeMB = 4
LogDir = "/qat_crash"
DumpOnProcessCrash = 0
LogDirMaxSizeMB = 4096
ContSyncEnabled = 1
ContSyncLogDir = "/qat_logs"
ContSyncMaxLogFiles = 10
ContSyncMaxLogSizeMB = 100
```

NOTE: Package is installed with Debug section already added to configuration files – but feature is disabled by default.

NOTE: For Virtual Functions (VFs), the above [DEBUG] section must also exist in the configuration files, i.e. `c4xxxvf_dev0.conf`, `c4xxxvf_dev1.conf`, etc., with `Enabled = 1 | 2 | 3`.

Field descriptions:

- Enabled
 - [0]: Collecting data disabled
 - [1]: Collecting data enabled
- DebugLevel
 - [0]: No data collecting
 - [1]: Collecting API calls only
 - [2]: Collecting FW requests and responses (default)
 - [3]: Collecting all above
- NumBuffers:
 - [50-2000]: Number of buffers for data storage per device

- BufferSizeMB:
[2-4]: Size of each buffers in MB
 - LogDir:
["path"]: Path to directory for crash dumps
 - LogDirMaxSizeMB:
[1024+]: Maximum size of crash dump directory
5. *If there is no space for new crash dump – the oldest crash dump directory present under 'LogDir' path is removed.*
- DumpOnProcessCrash:
[0]: Do not dump buffers in case of user-space process connected to QAT crash
[1]: Dump buffers in case of user-space process connected to QAT crash
 - ContSyncEnabled:
[0]: Do not perform ongoing synchronization of collected data with persistent storage
[1]: Perform ongoing synchronization of collected data with persistent storage

NOTE: *If continuous sync mode is enabled – crash dumps are not performed while handling error events. In such case the post-processing analysis is performed only based on the data collected by continuous sync option. In this case the buffers number and their size must be configured properly to hold generated QAT payload in case of high throughput.*

- ContSyncLogDir:
["path"]: Path to directory for continuous sync data
 - ContSyncMaxLogFiles:
[10-100]: Maximum number of continuous sync files
 - ContSyncMaxLogSizeMB:
[100-1000]: Maximum size in MB of particular continuous sync file
6. *Please reload configuration after each change by using 'adf_ctl restart'.*

8.4.2 Configuration via sysfs

On Linux*, the configuration via `sysfs` is very similar to configuration via QAT configuration files. Main difference is that the parameters are passed to files, which are accessible in `/sys/kernel/debug/<device>/qat_debug/` directory. QAT configuration files stored by default in `/etc/` may stay unmodified. Reloading parameters follows writing to `enabled` file. Thus, all parameters should already be passed before writing to that file.

NOTE: *This functionality is not available on FreeBSD.*

One could define a script named `sysfs_cfg.sh` to perform configuration in a clean manner:

```
#!/bin/bash

[[ $# -ne 1 ]] && echo "Error: Please provide sysfs subfolder" && exit
1;
dev=$1
echo "200" > /sys/kernel/debug/${dev}/qat_debug/buffer_pool_size
echo "4" > /sys/kernel/debug/${dev}/qat_debug/buffer_size_mb
echo "/qat_logs" > /sys/kernel/debug/${dev}/qat_debug/cont_sync_dir
echo "0" > /sys/kernel/debug/${dev}/qat_debug/cont_sync_enabled
echo "10" > /sys/kernel/debug/${dev}/qat_debug/cont_sync_max_files
echo "100" >
/sys/kernel/debug/${dev}/qat_debug/cont_sync_max_file_size_mb
echo "/qat_crash" > /sys/kernel/debug/${dev}/qat_debug/dump_dir
echo "4096" > /sys/kernel/debug/${dev}/qat_debug/dump_dir_size_mb
echo "3" > /sys/kernel/debug/${dev}/qat_debug/level
echo "1" > /sys/kernel/debug/${dev}/qat_debug/dump_on_process_crash

#commit
echo "1" > /sys/kernel/debug/${dev}/qat_debug/enabled
```

Example call to modify QAT 1.7 VF available under `0000:4d:02.0` bdf:
`sh sysfs_cfg.sh qat_c4xxxvf_0000:4d:02.0`

The configuration should end with starting Debuggability daemon with a command:
`qat_dbg_daemon_sync`

8.4.3 Checking Current Configuration Used by Driver

To check details about current configuration used by driver, the following utility can be used:

- `qat_dbg_ctl`

Tool usage looks as follows:

```
-----
USAGE:
-----
# /usr/local/bin/qat_dbg_ctl start||stop||status||restart
-----

To see QAT debuggability configuration in the system use:
/usr/local/bin/qat_dbg_ctl status
To start QAT debuggability synchronization daemon use:
/usr/local/bin/qat_dbg_ctl start
To terminate QAT debuggability synchronization daemon use:
/usr/local/bin/qat_dbg_ctl stop
To restart QAT debuggability synchronization daemon use:
/usr/local/bin/qat_dbg_ctl restart
-----
```

Example (QAT Debug not configured):

```
# qat_dbg_ctl status

QAT debuggability configuration:
```

```
No QAT devices configured with debuggability
QAT debuggability synchronization daemon not running.
```

Example (QAT Debug configured):

```
# qat_dbg_ctl status

QAT debuggability configuration:
  Device: qat_c6xx_0000:3f:00.0
    Debug level: 3
    Buffer pool size: 100
    Buffer size in MB: 4
    Crash dump on client process: 0
    Synchronization mode: dump on crash
    Crash dump directory: /qat_crash
    Crash dump directory max size in MB: 4096
  Device: qat_c6xx_0000:3d:00.0
    Debug level: 3
    Buffer pool size: 128
    Buffer size in MB: 4
    Crash dump on client process: 0
    Synchronization mode: cont-sync
    Cont-sync directory: /qat_logs
    Max number of cont-sync log files 10
    Max size of cont-sync log file: 100
QAT debuggability synchronization daemon running. Pid:
31962
```

7. If feature is enabled – ‘[qat_dbg_sync_daemon](#)’ should be up and running. Daemon is initialized automatically by [adf_ctl](#) during configuration reloading.

QAT debug synchronization daemon ([qat_dbg_sync_daemon](#)) logs to the syslog – you can check daemon activities e.g., by the following command:

```
# grep qat_dbg_sync_daemon /var/log/syslog

Nov 22 12:18:42 ubuntu qat_dbg_sync_daemon[11641]: Starting
daemon...
Nov 22 12:18:42 ubuntu qat_dbg_sync_daemon[11641]: Device 0
configuration:
Nov 22 12:18:42 ubuntu qat_dbg_sync_daemon[11641]: -dump_dir:
/qat_crash
Nov 22 12:18:42 ubuntu qat_dbg_sync_daemon[11641]: -
dump_dir_size_mb: 4096
Nov 22 12:18:42 ubuntu qat_dbg_sync_daemon[11641]: -
buffer_pool_size: 128
Nov 22 12:18:42 ubuntu qat_dbg_sync_daemon[11641]: -buffer_size_mb:
4
Nov 22 12:18:42 ubuntu qat_dbg_sync_daemon[11641]: -level: 2
Nov 22 12:18:42 ubuntu qat_dbg_sync_daemon[11641]: -
dump_on_process_crash:0
Nov 22 12:18:42 ubuntu qat_dbg_sync_daemon[11641]: -sync_mode:
continuous synchronization
Nov 22 12:18:42 ubuntu qat_dbg_sync_daemon[11641]: -
cont_sync_dir:/qat_logs
Nov 22 12:18:42 ubuntu qat_dbg_sync_daemon[11641]: -
cont_sync max file size mb:100
```

```
Nov 22 12:18:42 ubuntu qat_dbg_sync_daemon[11641]: -
cont_sync_max_files:10
Nov 22 12:18:42 ubuntu qat_dbg_sync_daemon[11641]: Creating cont-
sync directory: /qat_logs
Nov 22 12:18:42 ubuntu qat_dbg_sync_daemon[11641]: Initialized
cont-sync mode for 1 devices
Nov 22 12:18:42 ubuntu qat_dbg_sync_daemon[11641]: Daemon started
Nov 22 12:18:42 ubuntu qat_dbg_sync_daemon[11641]: QAT events
listener worker started
```

8.5 Usage Examples

8.5.1 Collecting Data – Sanity Check

Various tools can be used to perform collecting data sanity check. One of them is `cpa_sample_code` (installed directly with QAT package by execution of 'samples-install' target) and this one will be used as an example.

8.5.1.1 Continuous Sync Enabled

NOTE: Ensure that at least one device is configured with Debug feature (Enabled=1) and cont-sync mode enabled (ContSyncEnabled=1) by editing QAT device configuration file and reloading the configuration by running 'adf_ctl restart' command

1. Run test:
`cpa_sample_code signOfLife=1`
2. Check collected data by using 'qat_dbg_report' tool:

```
# qat_dbg_report path=/qat_logs command=list dev=0 last=0
=====
Building index...
DONE
      Overall indexed 3269338 msgs.
      Requests: 150408 (Sym:11280, PKE:138932, DC:196)
      Responses:150408
      API calls:2968522
=====
```

NOTE: You can use different last or audit commands other than 'list' to perform this test as well.

8.5.1.2 Continuous Sync Disabled

NOTE: Ensure that at least one device is configured with Debug feature enabled (Enabled=1) and cont-sync mode disabled (ContSyncEnabled=0)

1. Run test:
`cpa_sample_code signOfLife=1`

2. Trigger crash-dump manually to check collected data (set the *dev* parameter the same as QAT device number configured with cont-sync mode disabled):

```
qat_dbg_report command=dump dev=0
```

3. [Optional] You can check `qat_dbg_sync_daemon` logs if event has been handled:

```
# tail -F /var/log/messages | grep -i qat_dbg_sync
Nov 30 11:56:25 localhost qat_dbg_sync_daemon[38164]: Daemon started
Nov 30 12:00:39 localhost qat_dbg_sync_daemon[38164]: Received QAT
event: manual_dump
Nov 30 12:00:39 localhost qat_dbg_sync_daemon[38164]: Creating crash
dump directory: /qat_crash
Nov 30 12:00:39 localhost qat_dbg_sync_daemon[38164]: Crash dump in
progress ...
Nov 30 12:00:40 localhost qat_dbg_sync_daemon[38164]: Dumping physical
memory regions to file:
/qat_crash/qat_crash_dev_<dev>_<bdf>_<timestamp>/proc.mmaps.<dev>_<bdf>
Nov 30 12:00:40 localhost qat_dbg_sync_daemon[38164]: Crash dump done -
path: /qat_crash/qat_crash_dev_<dev>_<bdf>_<timestamp>
```

4. Check collected data by using `qat_dbg_report` tool:

```
# qat_dbg_report path=/qat_crash/qat_crash_dev_<dev>_<timestamp>
command=list last=0

=====

Building index...
DONE
      Overall indexed 3430983 msgs.
           Requests: 158708 (Sym:11280, PKE:147232, DC:196)
           Responses:158708
           API calls:3113567
=====
```

NOTE: You can use different last or audit commands other than 'list' to perform this test as well.

8.5.2 Audit Physical Addresses – Sanity Check

8.5.2.1 Emulate Uncorrectable Error

You can force QAT FW to crash by using a modified version of the tool provided in the QAT package. To prepare tool to send incorrect data to QAT FW please use the modified following file:

```
$ICP_ROOT/quickassist/lookaside/access_layer/src/sample_code/functional/sym/symdp_sample/cpa_sym_dp_sample.c
```

NOTE: Set *ICP_ROOT* to where you have your package extracted (e.g., `export ICP_ROOT=/root/QAT`)

Following change can be applied to force uncorrectable error:

```
196,199c196,199
```

```

-     pOpData->srcBuffer = sampleVirtToPhys(pSrcBuffer);
-     pOpData->srcBufferLen = bufferSize;
-     pOpData->dstBuffer = sampleVirtToPhys(pSrcBuffer);
-     pOpData->dstBufferLen = bufferSize;
+     pOpData->srcBuffer = pSrcBuffer;
+     pOpData->srcBufferLen = CPA_DP_BUFLIST;
+     pOpData->dstBuffer = pSrcBuffer;
+     pOpData->dstBufferLen = CPA_DP_BUFLIST;

```

Compile modified tool by using following commands:

1. `cd`
`$ICP_ROOT/quickassist/lookaside/access_layer/src/sample_code/function`
`al`
2. `make all`

To significantly improve recovery time after an uncorrectable error event, ensure that the `AutoResetOnError` configuration option (`AutoResetOnError = 1`) is set in the QAT configuration file.

8.5.2.2 Continuous Sync Enabled

Ensure that at least one device is configured with Debug feature (`Enabled = 1`) and cont-sync mode enabled (`ContSyncEnabled = 1`)

1. Execute modified tool:

```

# cd
$ICP_ROOT/quickassist/lookaside/access_layer/src/sample_code/functional
/build
# ./sym_dp_sample

```

NOTE: Please press 'ctrl+c' almost immediately after tool execution. Keeping tool up, can cause timeout on driver side while waiting for client processes to detach from device before restart routine.

2. [Optional] Check if event has been caught and handled properly:

```

# tail -F /var/log/messages | grep -i qat_dbg_sync

Nov  3 15:14:38 localhost qat_dbg_sync_daemon[12940]: Received QAT
event: error
Nov  3 15:14:38 localhost qat_dbg_sync_daemon[12940]: Dumping physical
memory regions to file: /qat_logs/proc.mmmaps.dev00_0000_4d_00_0
Nov  3 15:14:38 localhost qat_dbg_sync_daemon[12940]: Received QAT
event: restarting
Nov  3 15:14:41 localhost qat_dbg_sync_daemon[12940]: Received QAT
event: restarted

```

3. Execute audit:

```

# qat_dbg_report path

=/qat_logs/ command=audit_phy_addresses dev=0
=====
Building index...

```

```

DONE
    Overall indexed 2 msgs.
        Requests: 1 (Sym:1, PKE:0, DC:0)
        Responses:0
        API calls:1
=====
=====

QAT Physical addresses - audit in progress ...

ERROR: Missing SGL source in log entry.
ERROR: Missing SGL destination in log entry.
ERROR: SGL audit failed - check entry below.
ERROR: address overlapping audit failed - check entry below.
ERROR: Physical address (0x7f460961a800) used in request is out of
process pid: 9423 range.
    Check /qat_logs//proc.mmaps.dev00_0000_4d_00_0 to see process
physical addresses ranges.
ERROR: User process memory regions audit failed - check entry below.
    Entry [REQUEST SYM]: Time-stamp: 2021-11-03 15:14:37.849936202
    Bank: 1 Ring: 2 PID: 9423
    [0.1B] Crypto command ID:
ICP_QAT_FW_LA_CMD_CIPHER_HASH [2]
    [0.2B] Service type: ICP_QAT_FW_COMN_REQ_CPM_FW_LA
[4]
    [1.0-1B] LA BULK (SYMMETRIC CRYPTO) COMMAND FLAGS
(0x24)
    [1.12] ZUC_3G_PROTO: 0
    [1.11] GCM_IV_LEN_FLAG: 0
    [1.10] DIGEST_IN_BUFFER: 0
    [1.7-9] PROTO: 0
    [1.6] CMP_AUTH: 0
    [1.5] RET_AUTH: 1
    [1.4] UPDATE_STATE: 0
    [1.3] CIPH_AUTH_CFG_OFFSET_FLAG: 0
    [1.2] CIPH_IV_FLD_FLAG: 1
    [1.0-1] PARTIAL_FLAGS: 0 (FULL)
[1.2B] Common Request flags: 0x1
    SGL[1] CD_IN [0] BNP [0]
[1.3B] Extended Symmetric Crypto Command Flags: 0
[2-3] Content Descriptor (CD) Param Pointer:
0x192752c40
    [4.2B] Content Descriptor Param Size: 15 [Quad words]
[6-7] Opaque Data: 0x7f460961b000
[8-9] Source phy_addr: 0x7f460961a800
[10-11] Destination phy_addr: 0x7f460961a800
[12] Source length: 0
[13] Destination length: 0
[14-19] Cipher Request Parameters:
    [14] uint32_t::cipher_offset: 0
    [15] uint32_t::cipher_length: 96
    [16-17] uint64_t::cipher_IV_ptr:
0xdfc54a821d4c9b7e
    [18-19] uint64_t::resrvd1:
0x27378daa44a14c99
    [27-28.0B] Cipher Request Control Header:

```

```

[27.0B] uint8_t::cipher_state_sz: 2
[27.1B] uint8_t::cipher_key_sz: 4
[27.2B] uint8_t::cipher_cfg_offset: 0
[27.3B] uint8_t::next_curr_id: 0x21
(curr_id: 1, next: 2)
[28.0B] uint8_t::cipher_padding_sz: 0
[20-26] Authentication Request Parameters:
[20] uint32_t::auth_off: 0
[21] uint32_t::auth_len: 96
[22-23] uint64_t::aad_adr/APS: 0
[24-25] uint64_t::auth_res_addr: 0x192753860
[26.0B] uint8_t::aad_sz/inner_prefix_sz: 0
[26.1B] uint8_t::resrvd1: 0
[26.2B] uint8_t::hash_state_sz: 0
[26.3B] uint8_t::auth_res_sz: 0
[27-31] Authentication Request Control Header:
[27] uint32_t::resrvd1: 0x21000402
[28.0B] uint8_t::resrvd2: 0x0
[28.1B] uint8_t::hash_flags: 0x0
[28.2B] uint8_t::hash_cfg_offset: 5
[28.3B] uint8_t::next_curr_id: 0x42
(curr_id: 2, next: 4)
[29.0B] uint8_t::resrvd3: 0x0
[29.1B] uint8_t::outer_prefix_offset: 0
[29.2B] uint8_t::final_sz: 32
[29.3B] uint8_t::inner_res_sz: 32
[30.0B] uint8_t::resrvd4: 0x0
[30.1B] uint8_t::inner_state1_sz: 32
[30.2B] uint8_t::inner_state2_offset: 11
[30.3B] uint8_t::inner_state2_sz: 32
[31.0B] uint8_t::outer_config_offset: 0
[31.1B] uint8_t::outer_state1_sz: 0
[31.2B] uint8_t::outer_res_sz: 0
[31.3B] uint8_t::outer_prefix_offset: 0
SGL Data:
=====
100% [||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||]
Checked 2 records. Found 1 issue(s).

```

8.5.2.3 Continuous Sync Disabled (Crash Dump Based)

NOTE: Ensure that at least one device is configured with Debug feature enabled (`Enabled = 1`) and cont-sync mode disabled (`ContSyncEnabled = 0`)

1. Execute modified tool:

```

# cd
$IACP_ROOT/quickassist/lookaside/access_layer/src/sample_code/functional
/build
# ./sym dp_sample

```

NOTE: Please press 'ctrl+c' almost immediately after tool execution. Keeping tool up, can cause timeout on driver side while waiting for client processes to detach from device before restart routine

2. [Optional] Check if event has been caught and handled properly:

```
# tail -F /var/log/messages|grep -i qat_dbg_sync

Nov  3 15:50:03 localhost qat_dbg_sync_daemon[14822]: Received QAT
event: error
Nov  3 15:50:03 localhost qat_dbg_sync_daemon[14822]: Creating crash
dump directory: /qat_crash
Nov  3 15:50:03 localhost qat_dbg_sync_daemon[14822]: Crash dump in
progress ...
Nov  3 15:50:03 localhost qat_dbg_sync_daemon[14822]: Dumping physical
memory regions to file: /qat_crash/qat_crash_dev_00_2021-11-
03_155003//proc.mmaps.dev00_0000_4d_00_0
Nov  3 15:50:03 localhost qat_dbg_sync_daemon[14822]: Crash dump done
- path: /qat_crash/qat_crash_dev_00_2021-11-03_155003/
Nov  3 15:50:03 localhost qat_dbg_sync_daemon[14822]: Received QAT
event: restarting
Nov  3 15:50:06 localhost qat_dbg_sync_daemon[14822]: Received QAT
event: restarted
```

3. Execute audit:

```
# qat_dbg_report
path=/qat_crash/qat_crash_dev_<dev>_<bdf>_<timestamp>/
command=audit_phy_addresses

=====
====

Building index...
DONE
    Overall indexed 2 msgs.
        Requests: 1 (Sym:1, PKE:0, DC:0)
        Responses:0
        API calls:1

=====
====

=====
====

QAT Physical addresses - audit in progress ...

ERROR: Missing SGL source in log entry.
ERROR: Missing SGL destination in log entry.
ERROR: SGL audit failed - check entry below.
ERROR: address overlapping audit failed - check entry below.
ERROR: Physical address (0x7f950ff52c00) used in request is out of
process pid: <PID> range.
    Check
/qat_crash/qat_crash_dev_<dev>_<bdf>_<timestamp>/proc.mmaps.dev<dev>_<
bdf> to see process physical addresses ranges.
ERROR: User process memory regions audit failed - check entry below.

Entry [REQUEST SYM]: Time-stamp: 2020-11-30 13:04:10.59952422
Bank: 1 Ring: 2 PID: <PID>
```



```

[0.1B] Crypto command ID:
ICP_QAT_FW_LA_CMD_CIPHER_HASH [2]
[0.2B] Service type: ICP_QAT_FW_COMN_REQ_CPM_FW_LA
[4]
[1.0-1B] LA BULK (SYMMETRIC CRYPTO) COMMAND FLAGS
(0x24)
[1.12] ZUC_3G_PROTO: 0
[1.11] GCM_IV_LEN_FLAG: 0
[1.10] DIGEST_IN_BUFFER: 0
[1.7-9] PROTO: 0
[1.6] CMP_AUTH: 0
[1.5] RET_AUTH: 1
[1.4] UPDATE_STATE: 0
[1.3] CIPH_AUTH_CFG_OFFSET_FLAG: 0
[1.2] CIPH_IV_FLD_FLAG: 1
[1.0-1] PARTIAL FLAGS: 0 (FULL)
[1.2B] Common Request flags: 0x1
      SGL[1] CD_IN [0] BNP [0]
[1.3B] Extended Symmetric Crypto Command Flags: 0
[2-3] Content Descriptor (CD) Param Pointer:
0xeb20e4440
[4.2B] Content Descriptor Param Size: 15 [Quad words]
[6-7] Opaque Data: 0x7f950ff53400
[8-9] Source phy_addr: 0x7f950ff52c00
[10-11] Destination phy_addr: 0x7f950ff52c00
[12] Source length: 0
[13] Destination length: 0
[14-19] Cipher Request Parameters:
      [14] uint32_t::cipher_offset: 0
      [15] uint32_t::cipher_length: 96
      [16-17] uint64_t::cipher_IV_ptr:
0xdfc54a821d4c9b7e
      [18-19] uint64_t::resrvd1:
0x27378daa44a14c99
[27-28.0B] Cipher Request Control Header:
      [27.0B] uint8_t::cipher_state_sz: 2
      [27.1B] uint8_t::cipher_key_sz: 4
      [27.2B] uint8_t::cipher_cfg_offset:
0
      [27.3B] uint8_t::next_curr_id: 0x21
(curr_id: 1, next: 2)
      [28.0B] uint8_t::cipher_padding_sz:
0
[20-26] Authentication Request Parameters:
      [20] uint32_t::auth_off: 0
      [21] uint32_t::auth_len: 96
      [22-23] uint64_t::aad_adr/APS: 0
      [24-25] uint64_t::auth_res_addr:
0xeb20e4c60
      [26.0B]
uint8_t::aad_sz/inner_prefix_sz: 0
      [26.1B] uint8_t::resrvd1: 0
      [26.2B] uint8_t::hash_state_sz: 0
      [26.3B] uint8_t::auth_res_sz: 0
[27-31] Authentication Request Control Header:
      [27] uint32_t::resrvd1:
0x21000402
      [28.0B] uint8_t::resrvd2: 0x0
      [28.1B] uint8_t::hash_flags: 0x0

```

```

[28.2B] uint8_t::hash_cfg_offset: 5
[28.3B] uint8_t::next_curr_id: 0x42
(curr_id: 2, next: 4)
[29.0B] uint8_t::resrvd3: 0x0
[29.1B]
uint8_t::outer_prefix_offset: 0
[29.2B] uint8_t::final_sz: 32
[29.3B] uint8_t::inner_res_sz: 32
[30.0B] uint8_t::resrvd4: 0x0
[30.1B] uint8_t::inner_state1_sz: 32
[30.2B]
uint8_t::inner_state2_offset: 11
[30.3B] uint8_t::inner_state2_sz: 32
[31.0B]
uint8_t::outer_config_offset: 0
[31.1B] uint8_t::outer_state1_sz: 0
[31.2B] uint8_t::outer_res_sz: 0
[31.3B]
uint8_t::outer_prefix_offset: 0
      SGL Data:
=====
=====

100% [||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||]
Checked 2 records. Found 1 issue(s).

```

8.5.3 Audit Cipher Buffers Alignment – Sanity Check

8.5.3.1 Emulate Slice Hang Caused by Incorrect Buffers Alignments

NOTE: To execute this test, the package should be compiled with ‘--disable-param-check’ option. To do this, you should uninstall existing package and install it again with extra configuration option mentioned above

You can force QAT slice to hang by using a modified version of the tool provided by the QAT package. To prepare tool to send incorrect data to QAT FW please modify the following file: `$ICP_ROOT/quickassist/lookaside/access_layer/src/sample_code/functional/sym/ipsec_sample/cpa_ipsec_sample.c`

Following change can be applied to buffer lengths alignment error and slice hang:

```

308c308
-           pOpData->messageLenToCipherInBytes =
sizeof(samplePayload);
+           pOpData->messageLenToCipherInBytes = 2;

```

Compile modified tool by using following commands:

1. `cd $ICP_ROOT/quickassist/lookaside/access_layer/src/sample_code/functional`
2. `make all`

8.5.3.2 Slice Hang Handling with Continuous Sync Enabled

Ensure that at least one device is configured with Debug feature (`Enabled = 1`) and cont-sync mode enabled (`ContSyncEnabled = 1`)

1. Execute modified tool:

```
# cd
$ICP_ROOT/quickassist/lookaside/access_layer/src/sample_code/functional
/build
# ./ipsec_sample
main(): Starting IPsec Sample Code App ...
algChainSample(): cpaCyStartInstance
algChainSample(): Encrypt-Generate ICV
algChainPerformOp(): cpaCySymPerformOp
[error] LacSymQat_SymLogSliceHangError() - : slice hang detected on CPM
cipher or auth slice.
[error] LacSymCb_ProcessCallbackInternal() - : Response status value
not as expected
symCallback(): Callback called with status = -1.
symCallback(): Callback verify result error
algChainPerformOp(): Output does not match expected output encrypt
generate
algChainSample(): cpaCyStopInstance
algChainSample(): Sample code failed with status of -1
main():
IPsec Sample Code App failed
```

2. [Optional] Check if event has been caught and handled properly:

```
# tail -F /var/log/messages | grep -i qat_dbg_sync

Nov 30 16:09:34 localhost qat_dbg_sync_daemon[22391]:
Received QAT event: err_resp
Nov 30 16:09:34 localhost qat_dbg_sync_daemon[22391]: Dumping
physical memory regions to file:
/qat_logs/proc.mmmaps.dev00 0000 4d 00 0
```

3. Execute audit:

```
# qat_dbg_report path=/qat_logs command=audit_fields_lengths dev=0

=====
=====

Building index...
DONE

Overall indexed 3 msgs.
Requests: 1 (Sym:1, PKE:0, DC:0)
Responses:1
API calls:1

=====
=====

=====
=====
```

```

QAT request fields length - audit in progress...
ERROR: Cipher data size must be block multiple (Cipher len:2, block
size:16) for alg: CPA_CY_SYM_CIPHER_AES_CBC
    Entry [REQUEST_SYM]: Time-stamp: 2021-11-23 07:37:32.625265272
    Bank: 0 Ring: 1 PID: 23612
        [0.1B] Crypto command ID:
ICP_QAT_FW_LA_CMD_CIPHER_HASH [2]
    [0.2B] Service type: ICP_QAT_FW_COMN_REQ_CPM_FW_LA [4]
    [1.0-1B] LA BULK (SYMMETRIC CRYPTO) COMMAND FLAGS
(0x2c)
        [1.12] ZUC_3G_PROTO: 0
        [1.11] GCM_IV_LEN_FLAG: 0
        [1.10] DIGEST_IN_BUFFER: 0
        [1.7-9] PROTO: 0
        [1.6] CMP_AUTH: 0
        [1.5] RET_AUTH: 1
        [1.4] UPDATE_STATE: 0
        [1.3] CIPH_AUTH_CFG_OFFSET_FLAG: 1
        [1.2] CIPH_IV_FLD_FLAG: 1
        [1.0-1] PARTIAL FLAGS: 0 (FULL)
    [1.2B] Common Request flags: 0x1
        SGL[1] CD_IN [0] BNP [0]
    [1.3B] Extended Symmetric Crypto Command Flags: 0
    [2-3] Content Descriptor (CD) Param Pointer:
0x1af801640
    [4.2B] Content Descriptor Param Size: 8 [Quad words]
    [6-7] Opaque Data: 0x7f9dd0ebbc40
    [8-9] Source phy_addr: 0x1af802000
    [10-11] Destination phy_addr: 0x1af802000
    [12] Source length: 0
    [13] Destination length: 0
    [14-19] Cipher Request Parameters:
        [14] uint32_t::cipher_offset: 24
        [15] uint32_t::cipher_length: 2
        [16-17] uint64_t::cipher_IV_ptr:
0xaddbcefabebafeca
        [18-19] uint64_t::resrvd1: 0x459113d88f8cade
    [27-28.0B] Cipher Request Control Header:
        [27.0B] uint8_t::cipher_state_sz: 2
        [27.1B] uint8_t::cipher_key_sz: 2
        [27.2B] uint8_t::cipher_cfg_offset: 18
        [27.3B] uint8_t::next_curr_id: 0x21 (curr_id:
1, next: 2)
        [28.0B] uint8_t::cipher_padding_sz: 0
    [20-26] Authentication Request Parameters:
        [20] uint32_t::auth_off: 0
        [21] uint32_t::auth_len: 88
        [22-23] uint64_t::aad_adr/APS: 0x1af801820
        [24-25] uint64_t::auth_res_addr: 0x1af802458
        [26.0B] uint8_t::aad_sz/inner_prefix_sz: 0
        [26.1B] uint8_t::resrvd1: 0
        [26.2B] uint8_t::hash_state_sz: 0
        [26.3B] uint8_t::auth_res_sz: 0
    [27-31] Authentication Request Control Header:
        [27] uint32_t::resrvd1: 0x21120202
        [28.0B] uint8_t::resrvd2: 0x0
        [28.1B] uint8_t::hash_flags: 0x0
        [28.2B] uint8_t::hash_cfg_offset: 46
    
```

```

2, next: 4)
[28.3B] uint8_t::next_curr_id: 0x42 (curr_id:
[29.0B] uint8_t::resrvd3: 0x0
[29.1B] uint8_t::outer_prefix_offset: 0
[29.2B] uint8_t::final_sz: 12
[29.3B] uint8_t::inner_res_sz: 20
[30.0B] uint8_t::resrvd4: 0x0
[30.1B] uint8_t::inner_state1_sz: 24
[30.2B] uint8_t::inner_state2_offset: 3
[30.3B] uint8_t::inner_state2_sz: 24
[31.0B] uint8_t::outer_config_offset: 0
[31.1B] uint8_t::outer_state1_sz: 0
[31.2B] uint8_t::outer_res_sz: 0
[31.3B] uint8_t::outer_prefix_offset: 0
SGL Data:
Source SGL contains 1 flat buffer(s):
[0] Flat buffer: len: 100 phy_addr:
0x1af802400
Destination SGL contains 1 flat buffer(s):
[0] Flat buffer: len: 100 phy_addr:
0x1af802400
=====
====

100% [|||||]
Checked 3 records. Found 1 issue(s).

```

8.5.3.3 Slice Hang Handling with Continuous Sync Disabled

Ensure that at least one device is configured with Debug feature enabled (`Enabled = 1`) and cont-sync mode disabled (`ContSyncEnabled = 0`)

- Execute modified tool:

```

# cd
$IACP_ROOT/quickassist/lookaside/access_layer/src/sample_code/functional
/build
# ./ipsec_sample
main(): Starting IPSec Sample Code App ...
algChainSample(): cpaCyStartInstance
algChainSample(): Encrypt-Generate ICV
algChainPerformOp(): cpaCySymPerformOp
[error] LacSymQat_SymLogSliceHangError() - : slice hang detected on CPM
cipher or auth slice.
[error] LacSymCb_ProcessCallbackInternal() - : Response status value
not as expected
symCallback(): Callback called with status = -1.
symCallback(): Callback verify result error
algChainPerformOp(): Output does not match expected output encrypt
generate
algChainSample(): cpaCyStopInstance
algChainSample(): Sample code failed with status of -1
main():
IPSec Sample Code App failed

```

4. [Optional] Check if event has been caught and handled properly:

```
# tail -F /var/log/messages | grep -i qat_dbg_sync

Nov 30 16:17:02 localhost qat_dbg_sync_daemon[23780]: Received QAT
event: err_resp
Nov 30 16:17:02 localhost qat_dbg_sync_daemon[23780]: Creating crash
dump directory: /qat_crash
Nov 30 16:17:02 localhost qat_dbg_sync_daemon[23780]: Crash dump in
progress ...
Nov 30 16:17:02 localhost qat_dbg_sync_daemon[23780]: Dumping physical
memory regions to file: /qat_crash/qat_crash_dev_00_0000_4d_00_0_2020-
11-30_161702//proc.mmaps.dev00_0000_4d_00_0
Nov 30 16:17:02 localhost qat_dbg_sync_daemon[23780]: Crash dump done
- path: /qat_crash/qat_crash_dev_00_0000_4d_00_0_2020-11-30_161702/
```

5. Execute audit:

```
# qat_dbg_report path=/qat_crash/qat_crash_dev_<dev>_<bdf>_<timestamp>
command=audit_fields_lengths

=====
====

Building index...
DONE

Overall indexed 3 msgs.
Requests: 1 (Sym:1, PKE:0, DC:0)
Responses:1
API calls:1

=====
====

=====
====

QAT request fields length - audit in progress...
ERROR: Cipher data size must be block multiple (Cipher len:2, block
size:16) for alg: CPA_CY_SYM_CIPHER_AES_CBC
Entry [REQUEST SYM]: Time-stamp: 2021-11-23 07:39:01.357421282
Bank: 0 Ring: 1 PID: 23661
[0.1B] Crypto command ID:
ICP_QAT_FW_LA_CMD_CIPHER_HASH [2]
[0.2B] Service type: ICP_QAT_FW_COMN_REQ_CPM_FW_LA [4]
[1.0-1B] LA BULK (SYMMETRIC CRYPTO) COMMAND FLAGS
(0x2c)

[1.12] ZUC_3G_PROTO: 0
[1.11] GCM_IV_LEN_FLAG: 0
[1.10] DIGEST_IN_BUFFER: 0
[1.7-9] PROTO: 0
[1.6] CMP_AUTH: 0
[1.5] RET_AUTH: 1
[1.4] UPDATE_STATE: 0
[1.3] CIPH_AUTH_CFG_OFFSET_FLAG: 1
[1.2] CIPH_IV_FLD_FLAG: 1
[1.0-1] PARTIAL FLAGS: 0 (FULL)
[1.2B] Common Request flags: 0x1
SGL[1] CD_IN [0] BNP [0]
[1.3B] Extended Symmetric Crypto Command Flags: 0
```

```

0x18e001640 [2-3] Content Descriptor (CD) Param Pointer:
              [4.2B] Content Descriptor Param Size: 8 [Quad words]
              [6-7]  Opaque Data: 0x7f4681b98c40
              [8-9]  Source phy_addr: 0x18e002000
              [10-11] Destination phy_addr: 0x18e002000
              [12]   Source length: 0
              [13]   Destination length: 0
              [14-19] Cipher Request Parameters:
                  [14]   uint32_t::cipher_offset: 24
                  [15]   uint32_t::cipher_length: 2
                  [16-17] uint64_t::cipher_IV_ptr:
0xaddbcefabebafeca [18-19] uint64_t::resrvd1: 0x459113d88f8cade
                    [27-28.0B] Cipher Request Control Header:
                        [27.0B] uint8_t::cipher_state_sz: 2
                        [27.1B] uint8_t::cipher_key_sz: 2
                        [27.2B] uint8_t::cipher_cfg_offset: 18
                        [27.3B] uint8_t::next_curr_id: 0x21 (curr_id:
1, next: 2)
                    [28.0B] uint8_t::cipher_padding_sz: 0
                    [20-26] Authentication Request Parameters:
                        [20]   uint32_t::auth_off: 0
                        [21]   uint32_t::auth_len: 88
                        [22-23] uint64_t::aad_addr/APS: 0x18e001820
                        [24-25] uint64_t::auth_res_addr: 0x18e002458
                        [26.0B] uint8_t::aad_sz/inner_prefix_sz: 0
                        [26.1B] uint8_t::resrvd1: 0
                        [26.2B] uint8_t::hash_state_sz: 0
                        [26.3B] uint8_t::auth_res_sz: 0
                    [27-31] Authentication Request Control Header:
                        [27]   uint32_t::resrvd1: 0x21120202
                        [28.0B] uint8_t::resrvd2: 0x0
                        [28.1B] uint8_t::hash_flags: 0x0
                        [28.2B] uint8_t::hash_cfg_offset: 46
                        [28.3B] uint8_t::next_curr_id: 0x42 (curr_id:
2, next: 4)
                    [29.0B] uint8_t::resrvd3: 0x0
                    [29.1B] uint8_t::outer_prefix_offset: 0
                    [29.2B] uint8_t::final_sz: 12
                    [29.3B] uint8_t::inner_res_sz: 20
                    [30.0B] uint8_t::resrvd4: 0x0
                    [30.1B] uint8_t::inner_state1_sz: 24
                    [30.2B] uint8_t::inner_state2_offset: 3
                    [30.3B] uint8_t::inner_state2_sz: 24
                    [31.0B] uint8_t::outer_config_offset: 0
                    [31.1B] uint8_t::outer_state1_sz: 0
                    [31.2B] uint8_t::outer_res_sz: 0
                    [31.3B] uint8_t::outer_prefix_offset: 0
                    SGL Data:
                        Source SGL contains 1 flat buffer(s):
                            [0] Flat buffer: len: 100 phy_addr:
0x18e002400
                        Destination SGL contains 1 flat buffer(s):
                            [0] Flat buffer: len: 100 phy_addr:
0x18e002400
=====
=====

```

```
100% [|||||]
Checked 3 records. Found 1 issue(s).
```

8.5.4 Audit Return Codes

To collect data with incorrect return codes – the tool with the same modifications as described in Section [8.5.3.1](#) can be used.

8.5.4.1 Audit Return Codes – Continuous Sync Option

An example audit of return codes for the data collected in Section [8.5.3.2](#) looks as follows:

```
# qat_dbg_report path=/qat_logs dev=0 command=audit_ret_codes

=====
====

Building index...
DONE
    Overall indexed 3 msgs.
        Requests: 1 (Sym:1, PKE:0, DC:0)
        Responses:1
        API calls:1

=====
====

=====
====

QAT Response return codes audit in progress ...

WARNING: Incorrect response RCs. Status: 128 error_code: 0xf0
    Entry [RESPONSE SYM]: Time-stamp: 2021-11-23
07:37:32.646255054
    Bank: 0 Ring: 5 PID: 23612
        [0.1B] Service ID: ICP_QAT_FW_COMN_RESP_SERV_CPM_FW
[1]
        [0.2B] Response type: ICP_QAT_FW_COMN_REQ_CPM_FW_LA
[4]
        [1.3B] Command ID: ICP_QAT_FW_LA_CMD_CIPHER_HASH [2]
        [1.1B] Common error code: 240
        [1.2B] Common status flags: 0x80
            CRYPTO STAT FLAG: 1
            PKE STAT FLAG: 0
            CMP STAT FLAG: 0
            XLAT STAT FLAG: 0
            XLAT APPLIED STAT FLAG: 0
            CMP EOF LAST BLK FLAG: 0
            UNSUPPORTED RQ STAT FLAG: 0
        [2-3] Opaque data: 0x7f9dd0ebbc40
    Entry [REQUEST SYM]: Time-stamp: 2021-11-23 07:37:32.625265272
    Bank: 0 Ring: 1 PID: 23612
```



```

[0.1B] Crypto command ID:
ICP_QAT_FW_LA_CMD_CIPHER_HASH [2]
[0.2B] Service type: ICP_QAT_FW_COMN_REQ_CPM_FW_LA [4]
[1.0-1B] LA BULK (SYMMETRIC CRYPTO) COMMAND FLAGS
(0x2c)
[1.12] ZUC_3G_PROTO: 0
[1.11] GCM_IV_LEN_FLAG: 0
[1.10] DIGEST_IN_BUFFER: 0
[1.7-9] PROTO: 0
[1.6] CMP_AUTH: 0
[1.5] RET_AUTH: 1
[1.4] UPDATE_STATE: 0
[1.3] CIPH_AUTH_CFG_OFFSET_FLAG: 1
[1.2] CIPH_IV_FLD_FLAG: 1
[1.0-1] PARTIAL FLAGS: 0 (FULL)
[1.2B] Common Request flags: 0x1
      SGL[1] CD_IN [0] BNP [0]
[1.3B] Extended Symmetric Crypto Command Flags: 0
[2-3] Content Descriptor (CD) Param Pointer:
0x1af801640
[4.2B] Content Descriptor Param Size: 8 [Quad words]
[6-7] Opaque Data: 0x7f9dd0ebbc40
[8-9] Source phy_addr: 0x1af802000
[10-11] Destination phy_addr: 0x1af802000
[12] Source length: 0
[13] Destination length: 0
[14-19] Cipher Request Parameters:
[14] uint32_t::cipher_offset: 24
[15] uint32_t::cipher_length: 2
[16-17] uint64_t::cipher_IV_ptr:
0xabdbcefabebafeca
[18-19] uint64_t::resrvd1: 0x459113d88f8cade
[27-28.0B] Cipher Request Control Header:
[27.0B] uint8_t::cipher_state_sz: 2
[27.1B] uint8_t::cipher_key_sz: 2
[27.2B] uint8_t::cipher_cfg_offset: 18
[27.3B] uint8_t::next_curr_id: 0x21 (curr_id:
1, next: 2)
[28.0B] uint8_t::cipher_padding_sz: 0
[20-26] Authentication Request Parameters:
[20] uint32_t::auth_off: 0
[21] uint32_t::auth_len: 88
[22-23] uint64_t::aad_adr/APS: 0x1af801820
[24-25] uint64_t::auth_res_addr: 0x1af802458
[26.0B] uint8_t::aad_sz/inner_prefix_sz: 0
[26.1B] uint8_t::resrvd1: 0
[26.2B] uint8_t::hash_state_sz: 0
[26.3B] uint8_t::auth_res_sz: 0
[27-31] Authentication Request Control Header:
[27] uint32_t::resrvd1: 0x21120202
[28.0B] uint8_t::resrvd2: 0x0
[28.1B] uint8_t::hash_flags: 0x0
[28.2B] uint8_t::hash_cfg_offset: 46
[28.3B] uint8_t::next_curr_id: 0x42 (curr_id:
2, next: 4)
[29.0B] uint8_t::resrvd3: 0x0
[29.1B] uint8_t::outer_prefix_offset: 0
[29.2B] uint8_t::final_sz: 12
[29.3B] uint8_t::inner_res_sz: 20

```

```

[30.0B] uint8_t::resrvd4: 0x0
[30.1B] uint8_t::inner_state1_sz: 24
[30.2B] uint8_t::inner_state2_offset: 3
[30.3B] uint8_t::inner_state2_sz: 24
[31.0B] uint8_t::outer_config_offset: 0
[31.1B] uint8_t::outer_state1_sz: 0
[31.2B] uint8_t::outer_res_sz: 0
[31.3B] uint8_t::outer_prefix_offset: 0
SGL Data:
Source SGL contains 1 flat buffer(s):
[0] Flat buffer: len: 100 phy_addr:
0x1af802400
Destination SGL contains 1 flat buffer(s):
[0] Flat buffer: len: 100 phy_addr:
0x1af802400
=====
=====

100% [||||||||||||||||||||||||||||||||||||||||||||||||||||||||||||]
Checked 3 records. Found 1 issue(s).

```

Audit prints warnings in case of any issue found. Every response considered as unsuccessful is tried to be matched to its corresponding request.

NOTE: *In some cases this operation may not be possible because the searched request could already be overwritten by new data.*

8.6 SR-IOV

The Debuggability feature can be used in Single Root – Input Output Virtualization terms.

NOTE: *As the errors occurring on VF utilized on the host is forwarded, unnecessary logs might be generated if the feature is enabled on guest utilizing the same QAT device. Thus, logs might not contain error information.*

8.6.1 Build instructions

The usage flow might be following:

Step No.	Host	Guest
1	Set up BIOS and OS	
2	Install the driver with debuggability flags with configure options	
3	Start the desired KVM	
4	Stop/start QAT PF service	
5	Start QAT VFs service	
6	Attach desired VFs to Guest	

7		Install the driver with debuggability flags with configure options
8		Set up configuration

Example steps to build the environment might be as follows:

1. Set up BIOS and OS:
Options regarding VT-d, VT-x, and SR-IOV should be enabled in BIOS. Host OS needs to enable IOMMU. Details can be found in 2.1 *Updating the BIOS Settings* and 2.2 *Installing and Configuring the Host Operating System* sections of *Using Intel® Virtualization Technology (Intel® VT) with Intel® QuickAssist Technology App Notes*.
2. Install on host:

```
cd /QAT && \
tar -xzomf QAT.L.4.17.0-00002.tar.gz && \
export ICP_ROOT=$PWD && \
./configure --enable-icp-sriov=host --enable-icp-qat-dbg && \
make install -j
```
3. Install and start the desired KVM:

```
virt-install --name Ubuntu_20.04_64bit \
--memory 8096 \
--cpu host \
--vcpus 8 \
--os-type linux \
--os-variant ubuntu20.04 \
--import \
--graphics none \
--disk /PathToImage/Ubuntu_20.04_64_k5.4.0-90_b1_v000.img,format=img,bus=virtio \
--network direct,source=enp4s0,source_mode=bridge,model=virtio \
--network bridge=virbr0,model=virtio && \
virsh start Ubuntu_20.04_64bit
```
4. Change configuration: set *enable* equals 1 in *[DEBUG]* section and reload with *adf_ctl* utility for desired VF. Restart QAT service with:

```
service qat_service stop &&\
service qat_service start
```
5. Start VFs service with:

```
/etc/init.d/qat_service_vfs start
```
6. Attach desired VFs to KVM. It might be done by

```
virsh attach-device Ubuntu_20.04_64bit attachdev.xml
```

Assuming that the attachdev.xml file has following body and it's addresses (domain, bus, slot, and function corresponds to appropriate VF):

```
<?xml version="1.0" encoding="utf-8"?>
<hostdev mode='subsystem' type='pci' managed='yes'>
  <source>
    <address domain='0x0' bus='0x4d'
slot='0x01' function='0x0' />
  </source>
</hostdev>
```

7. On guest: Ensure that the VFs were appropriately attached using:

```
lspci|grep Co-
```

8. Install with *guest* indicating flag:

```
./configure --enable-icp-sriov=guest --enable-icp-qat-dbg && \
make install -j
```
9. Set the configuration (ex. by modifying configuration file in `/etc/` directory and using `adf_ct/utility`).

8.6.2 Usage

After successful installation the tool can be used in multiple ways:

- Usage on SR-IOV guest only is the same as without SR-IOV enabled. The difference is that VFs are utilized, so there is a different module loaded and QAT configuration files names contain 'vf' string.
- For SR-IOV host only, there is a need to insert appropriate QAT VF kernel module. For example, in case of QAT Gen 2 Intel® C62x Chipset, it would be `qat_c62xvf`. In this case, the feature can be used as in a PF only scenario.
- In order to use the detached VFs kernel module, `qat_*vfs` should not be added and `sysfs` should be used to configure the device.
- To utilize the host attached and detached devices, the QAT VF kernel module should be added, required VFs are detached (for example, with `virsh nodedev-detach`) and `sysfs` is used to configure detached VFs with Debuggability.

8.7 Programming Guide

8.7.1 Physical to Virtual Translation Callback

By default, QAT Debug uses common, USDM-based (User Space DMA-able Memory) routine to perform physical to virtual address translation, which can be used with USDM memory allocator only. To provide possibility to use user-selected memory allocator, the user has to provide custom implementation of physical to virtual address translation routine for Debug purposes.

The callback is called in the same context as payload thread, immediately before placing QAT request to DMA-able memory.

Callback setter definition is available in the following header file:

```
$ICP_ROOT/quickassist/lookaside/access_layer/include/icp_sal_user.h
```

and look as follows:

```
/******
 *
 * @ingroup SalUser
 * @description
 * This function sets provided by user callback for translating
 * physical to virtual addresses for QAT Debuggability purposes.
 * The callback will be used for Data Plane API requests with SGL
 * provided.
 *
 * @assumptions
 * None
```

```

* @sideEffects
*     None
* @reentrant
*     Yes
* @threadSafe
*     Yes
*
* @param[in]  instanceHandle      Instance handle
* @param[in]  user_dbg_phys2virt  Function which will be translating
*                                physical addresses to virtual ones
*
* @retval CPA_STATUS_FAIL        Failed to extract transport
handle
* @retval CPA_STATUS_SUCCESS     User callback set successfully
*

```

```

*****
/
CpaStatus icp_sal_userSetDbgPhysToVirtCallback(
    CpaInstanceHandle instanceHandle,
    icp_sal_dbg_phys2virt_callback user_dbg_phys2virt);

```

QAT Debug handle can be extracted after transport initialization by the following function:

```

/*
* icp_adf_transGetLoggerHandle
*
* Description:
* Function returns Debuggability logger handle for the given
trans_handle
*
* Returns:
* CPA_STATUS_INVALID_PARAM  when invalid arguments provided
* CPA_STATUS_SUCCESS        on success
* CPA_STATUS_FAIL           on failure
*/
CpaStatus icp_adf_transGetLoggerHandle(icp_comms_trans_handle
trans_handle,
                                     icp_adf_dbg_handle_t **handle);
$ICP_ROOT/quickassist/lookaside/access_layer/include/icp_adf_transport.h

```

§