



Intel® QuickAssist Technology

API Programmer's Guide

Revision 012

December 2022



You may not use or facilitate the use of this document in connection with any infringement or other legal analysis. You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

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

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

All product plans and roadmaps are subject to change without notice.

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.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at [intel.com](https://www.intel.com).

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

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries.

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

Copyright © 2021, Intel Corporation. All rights reserved.

Contents

1	Introduction	9
1.1	Intended Audience	9
1.2	Related Documents and References.....	9
1.3	Using This Document.....	10
1.4	Terminology	10
2	Base API and API Conventions.....	12
2.1	Intel® QAT Technology Base API	12
2.1.1	Data Buffer Models	12
2.1.1.1	Flat Buffers	12
2.1.1.2	Scatter-Gather Lists	13
2.2	Intel® QuickAssist Technology API Conventions	14
2.2.1	Instance Discovery	14
2.2.2	Modes of Operation	16
2.2.2.1	Asynchronous Operation	16
2.2.2.2	Synchronous Operation	17
2.2.3	Memory Allocation and Ownership	18
2.2.4	Data Plane APIs.....	19
3	Intel® QuickAssist Technology Cryptographic API.....	21
3.1	Overview	21
3.1.1	Sessions.....	22
3.1.2	Priority.....	22
3.2	Using the Symmetric Cryptography API	22
3.2.1	General Concepts	22
3.2.1.1	Session	23
3.2.1.2	In-Place and Out-of-Place Support	23
3.2.1.3	Partial Support	23
3.2.2	Cipher.....	24
3.2.2.1	symCallback.....	24
3.2.2.2	cipherSample	24
3.2.3	Hash	29
3.2.4	Hash a File.....	30
3.2.5	Chained Cipher and Hash.....	32
3.2.6	Chained Cipher and Hash – IPsec Like Use Case.....	33
3.2.7	Chained Cipher and Hash – SSL Like Use Case	36
3.2.8	Chained Cipher and Hash – CCM Use Case.....	41
3.2.9	Chained Cipher and Hash – GCM Use Case.....	44
3.2.10	Chained Cipher and Hash Using the Symmetric Data Plane API.....	46
3.2.11	TLS Key and MGF Mask Generation	51
3.2.11.1	Setting CpaCyKeyGenTlsOpData Structure Fields	51
3.2.12	Session Update for Chained Cipher and Hash Operation.....	53
3.2.12.1	Create and Initialize a Session:	53
3.2.13	HKDF Use Case	55
3.2.13.1	Instance Configuration and Memory Allocation.....	56
3.2.13.2	HKDF Extract Expand Operation.....	56
3.2.13.3	HKDF Extract Expand Label Operation	56
3.2.13.4	HKDF Extract Expand Label and Sublabels operation	57
3.2.14	Perform HKDF operation	58
3.3	Using the Diffie-Hellman API	58



3.3.1	Prime Number Testing.....	60
3.4	Using the SM2 API	61
3.4.1	SM2 Digital Signature Generation and Verification	61
3.4.2	SM2 Public Key Encryption.....	62
3.4.3	SM2 Key Exchange	64
3.4.4	SM2 Elliptic Curve Point.....	67
4	Intel® QuickAssist Technology Data Compression API.....	70
4.1	Overview	70
4.1.1	Session	70
4.2	Sample – Stateful Data Compression.....	71
4.2.1	Session Establishment.....	71
4.2.2	Sample – Stateless Data Compression.....	79
4.3	Sample – Stateless Data Compression Using Multiple Compress Operations .	82
4.4	Sample – Data Compression Data Plane API.....	83
4.5	Sample - Chained Hash and Stateless Compression.....	86

Figures

Figure 1.	Flat Buffer Diagram.....	13
Figure 2.	Scatter-Gather List Diagram.....	14
Figure 3.	Asynchronous Operation.....	17
Figure 4.	Synchronous Operation	18
Figure 5.	IPSec Outbound	33
Figure 6.	IPSec Inbound.....	34
Figure 7.	SSL Outbound	37
Figure 8.	Secure Sockets Layer Inbound	38

Tables

Table 1.	Related Documents and References	9
Table 2.	Terminology.....	10

Listing

Listing 1.	Getting an Instance	14
Listing 2.	Querying and Starting an Instance.....	15
Listing 3.	Callback Function.....	24
Listing 4.	Getting an Instance	25
Listing 5.	Set Address Translation Function	25
Listing 6.	Start up.....	25
Listing 7.	Create and Initialize Cipher Session	26
Listing 8.	Memory Allocation	27
Listing 9.	Set Up Cipher Operational Data	28
Listing 10.	Perform Cipher Operation	28
Listing 11.	Wait for Completion	28
Listing 12.	Wait for Outstanding Requests	29
Listing 13.	Remove Cipher Session	29

Listing 14.	Create and Initialize Hash Session	29
Listing 15.	Set up Hash Operational Data	30
Listing 16.	Hash Session Setup Data.....	30
Listing 17.	Hashing a File	31
Listing 18.	Create and Initialize Session Cipher and Hash	32
Listing 19.	Set up Operational Data Cipher and Hash	32
Listing 20.	Session Setup Data IPsec Outbound	34
Listing 21.	Operational Data IPsec Outbound	35
Listing 22.	Session Setup Data IPsec Inbound	35
Listing 23.	Operational Data IPsec Inbound	36
Listing 24.	Session Data SSL Outbound	38
Listing 25.	Buffer Size SSL Outbound	39
Listing 26.	Buffer Setup SSL Outbound	39
Listing 27.	BufferList Setup SSL Outbound.....	39
Listing 28.	Operational Data SSL Outbound	40
Listing 29.	Session Data SSL Inbound	40
Listing 30.	Calculating padLen SSL Inbound.....	41
Listing 31.	Operational Data SSL Inbound	41
Listing 32.	Session Data CCM Generate-Encrypt.....	42
Listing 33.	Session Data CCM Decrypt-Verify	42
Listing 34.	CCM Allocate IV and AAD Buffers.....	43
Listing 35.	CCM Operational Data	43
Listing 36.	Session Data GCM Auth-Encrypt	44
Listing 37.	Session Data GCM Auth-Decrypt.....	45
Listing 38.	GCM Allocate IV and AAD Buffers.....	45
Listing 39.	GCM Operational Data	46
Listing 40.	Register Callback Function	47
Listing 41.	Create and Initialize Data Plane Session	47
Listing 42.	Data Plane Operational Data	48
Listing 43.	Data Plane Enqueue	50
Listing 44.	Data Plane Perform	50
Listing 45.	Wait for Outstanding Requests	50
Listing 46.	Data Plane Remove Session	50
Listing 47.	Register Callback Function	53
Listing 48.	Create and Initialize Data Plane Session	53
Listing 49.	Data Plane Operational Data	54
Listing 50.	Data Plane Enqueue	55
Listing 51.	Data Plane Perform	55
Listing 52.	Session Update	55
Listing 53.	HKDF Operation Data – Memory Allocation.....	56
Listing 54.	HKDF Extract Expand Operation	56
Listing 55.	HKDF Extract Expand Label Operation	57
Listing 56.	HKDF Extract Expand Label and Sublabels Operation.....	57
Listing 57.	HKDF Operation.....	58
Listing 58.	Allocate Memory and Populate Operational Data	58
Listing 59.	Perform Phase 1 Operation	59
Listing 60.	Perform Phase 2 Operation	60
Listing 61.	Setup Operational Data and Test Prime	60
Listing 62.	Create and Initialize Stateful Session	72
Listing 63.	Stateful Compression Memory Allocation	73
Listing 64.	Create Header.....	74
Listing 65.	Perform Stateful Compression Operation	75
Listing 66.	Create Footer	76
Listing 67.	Perform Stateful Decompression Operation	77
Listing 68.	Remove Stateful Session	79



Listing 69.	Querying and Starting a Compression Instance	79
Listing 70.	Create and Initialize Stateless Session.....	81
Listing 71.	Data Plane Remove Compression Session	82
Listing 72.	Setting the Initial Value of the Checksum	82
Listing 73.	Register Compression Callback Function	83
Listing 74.	Create and Initialize Compression Data Plane Session	83
Listing 75.	Setup Source Buffer	84
Listing 76.	Compression Data Plane Operational Data	85
Listing 77.	Data Plane Enqueue and Submit.....	85
Listing 78.	Data Plane Remove Compression Session	86
Listing 79.	Querying and Starting a Compression Instance	86
Listing 80.	Create and Initialize Session Hash and Compression	87
Listing 81.	Chained Hash and Stateless Compression Memory Allocation	89
Listing 82.	Set Up Operational Data Hash and Compression	90
Listing 83.	Verify the Output of Chained Hash and Stateless Compression	90
Listing 84.	Remove Chained Hash and Stateless Compression Session	92

Revision History

Document Number	Revision Number	Description	Revision Date
330684	012	In this release: <ul style="list-style-type: none"> • Section 2.1.1: Added a note that source and destination buffer types must match. 	December 2022
330684	011	In this release: <ul style="list-style-type: none"> • Updated Section 3.1 – added cpa_cy_ecsm2.h file reference as the API for SM2 	October 2021
330684	010	<ul style="list-style-type: none"> • Updated Section 3.2.11 - TLS Key and MGF Mask Generation • Updated Section 3.4 - Using the SM2 API 	June 2021
330684	009	<ul style="list-style-type: none"> • Updated cpaDcGenerateFooter Content 	April 2021
330684	008	<ul style="list-style-type: none"> • Updated Section 1.3 Using this Document • Updated Section 3.2.2, Cipher AES-256 • Listing 7, Updated to use newer cipher function in place of occurrence 3DES • Updated Section 3.2.3, Hash with SHA-256 • Listing 14, Updated to use a newer hash function in place of occurrence of MD5 • Listing 18, Updated to use newer cipher function in place of occurrence 3DES • Updated Section 3.2.10 Chained Cipher and Hash using the Symmetric Data Plane API • Listing 41 and 42, updated to use newer cipher function in place of occurrence 3DES • Section 3.2.11.2, Setting CpaCyKeyGenTlsOpsData updated code 	May 2020
330684	007	<ul style="list-style-type: none"> • Added Section 3.2.13, HKDF Use Case 	February 2020
330684	006	<ul style="list-style-type: none"> • Added Section 4.5 Chained Hash and Stateless Compression 	March 2019
330684	005	<ul style="list-style-type: none"> • Updated Section 3.1 Overview 	December 2018
330684	004	<ul style="list-style-type: none"> • Updated Listing 13. Remove Cipher Session • Updated Section 4.0 Note • Updated Section 4.2 Sample – Stateful Data Compression • Updated Listing 56. Remove Stateful Session 	June 2018
330684	003	<ul style="list-style-type: none"> • Added requirement that customers using Static or Dynamic Compression must decompress data and verify that it matches original source data. 	February 2018
330684	002	<ul style="list-style-type: none"> • Updates to code samples. 	September 2017
330684	001	<ul style="list-style-type: none"> • First public version of the document. Based on Intel® Confidential document number 442844- 	June 2014

Document Number	Revision Number	Description	Revision Date
330684	012	In this release: <ul style="list-style-type: none"> • Section 2.1.1: Added a note that source and destination buffer types must match. 	December 2022
330684	011	In this release: <ul style="list-style-type: none"> • Updated Section 3.1 – added <code>cpa_cy_ecsm2.h</code> file reference as the API for SM2 	October 2021
		1.3 with the revision history of that document retained for reference purposes.	
330684	1.3	<ul style="list-style-type: none"> • Added new API function, <code>cpaCySymSessionCtxGetDynamicSize()</code>, to Section 3.2.2.2, plus other minor updates. 	May 2014
330684	1.2	<ul style="list-style-type: none"> • Updated Listing 26, 27, and 30 in Section 3.2 	January 2014
330684	1.1	Updates for stateless data compression samples: <ul style="list-style-type: none"> • Updated Section 4.3 • Added Section 4.4 	March 2013
330684	1.0	<ul style="list-style-type: none"> • Initial release of the document 	September 2012

§

1 Introduction

This API programmer's guide describes the sample code that demonstrates how to use the Intel® QuickAssist Technology (Intel® QAT) APIs.

1.1 Intended Audience

This document is intended to be used by software engineers who wish to develop application software that uses the Intel® QAT APIs to accelerate the supported workloads and/or services.

1.2 Related Documents and References

Table 1. Related Documents and References

Document	Document Number /Location
Intel® QuickAssist Technology Cryptographic API Reference Manual	330685
Intel® QuickAssist Technology Data Compression API Reference Manual	330686
Intel® QuickAssist Technology Performance Optimization Guide	330687
Intel® QuickAssist Technology Programmer's Guide for Linux* - Hardware Version 1.7	336210
Intel® Communications Chipset 8925 to 8955 Series Software Programmer's Guide	330751
Intel® Communications Chipset 8900 to 8920 Series Software Programmer's Guide	330753
Intel Atom® Processor C2000 Product Family for Communications Infrastructure Software Programmer's Guide	330755
Intel® QuickAssist Technology Software for Linux* Release Notes	336211
NIST publication SP800-38C Recommendation for Block Cipher Modes of Operation: The CCM Mode for Authentication and Confidentiality	https://nvlpubs.nist.gov/nistpubs/Legacy/sp/nistspecialpublication800-38c.pdf
NIST publication SP800-38D Recommendation for Block Cipher Modes of Operation: Galois/Counter Mode (GCM) and GMAC	https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-38d.pdf

Document	Document Number /Location
NIST SP 800-90, March 2007 Recommendation for Random Number Generation Using Deterministic Random Bit Generators (Revised)	https://csrc.nist.gov/publications/detail/sp/800-90/revised/archive/2007-03-14
GZIP file format specification v4.3	RFC 1952
GZIP file format specification v3.3	RFC 1950
The Transport Layer Security (TLS) Protocol Version 1.0	RFC 2246
The Transport Layer Security (TLS) Protocol Version 1.1	RFC 4346
The Transport Layer Security (TLS) Protocol Version 1.2	RFC 5246
The Secure Sockets Layer (SSL) Protocol Version 3.0	RFC 6106

1.3 Using This Document

This document is structured as follows:

- [Section 2 Base API and API Conventions](#) describes aspects common to all Intel® QuickAssist Technology APIs
- [Section 3 Intel® QuickAssist Technology Cryptographic API](#) describes the Intel® QuickAssist Technology Cryptographic API
- [Section 4 Intel® QuickAssist Technology Data Compression API](#) describes the Intel® QuickAssist Technology Data Compression API

Code for all the examples in this document is contained in the software package and, after installation, can be found in a sub-directory of the following directory:

[quickassist/lookaside/access_layer/src/sample_code/functional](#)

Refer to [Table 2](#) for a list of terms and acronyms used in this manual.

1.4 Terminology

Table 2. Terminology

Term	Description
AAD	Additional Authenticated Data
AES	Advanced Encryption Standard
API	Application Programming Interface
CBC	Cipher Block Chaining
CCM	Counter mode with Cipher-block chaining Message authentication code

Term	Description
CPM	Content Processing Module
CY	Cryptographic
DC	Data Compression
DRBG	Deterministic Random Bit Generator
DSA	Digital Signature Algorithm
ECDH	Elliptic Diffie-Hellman
EC	Elliptic Curve
ESP	Encapsulating Security Payload
GCD	Greatest Common Divisor
GCM	Galois Counter Mode
HKDF	HMAC Key Derivation Function
HMAC	Hashed Message Authenticate Code
ICV	Integrity Check Value
IPSec	Internet Protocol Security
MAC	Message Authentication Code
NRBG	Non-Deterministic Random Bit Generator
PKE	Public Key Encryption
PV	Public Value
Intel® QAT	Intel® QuickAssist Technology
RBG	Random Bit Generation
RSA	A public-key encryption algorithm created by Rivest, Shamir, and Adleman
SSL	Secure Sockets Layer
TLS	Transport Layer Security (SSL successor)

2 Base API and API Conventions

This chapter describes aspects common to all Intel® QAT Technology APIs, starting with the base API and followed by conventions.

2.1 Intel® QAT Technology Base API

The Base API is a top-level API definition for Intel® QAT Technology. It contains structures, data types, and definitions that are common across the interface.

2.1.1 Data Buffer Models

Data buffers are passed across the API interface in one of the following formats:

- Flat Buffers represent a single region of physically contiguous memory and are described in detail in Section [2.1.1.1, Flat Buffers](#).
- Scatter-Gather Lists are essentially an array of flat buffers, for cases where the memory is not all physically contiguous. These are described in detail in Section [2.1.1.2, Scatter-Gather Lists](#).
- **Note: The source and destination buffer types must match.** For example, if the source buffer type is set to flat buffer, the destination buffer type must also be a flat buffer.

2.1.1.1 Flat Buffers

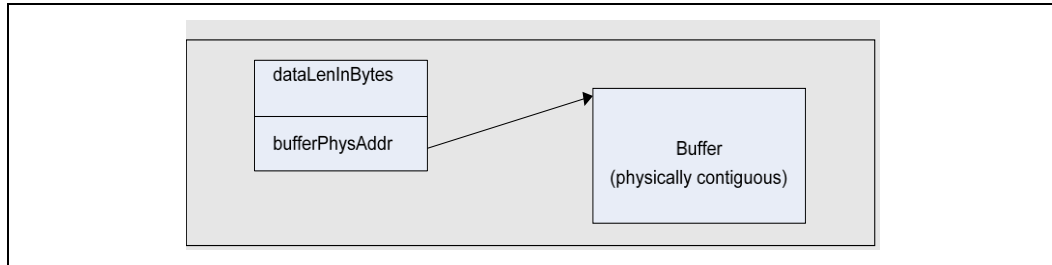
Flat buffers are represented by the type `CpaFlatBuffer`, defined in the file `cpa.h`. It consists of two fields:

- Data pointer `pData`: points to the start address of the data or payload. The data pointer is a virtual address; however, the actual data pointed to is required to be in contiguous and DMAable physical memory. This buffer type is typically used when simple, unchained buffers are needed.
- Length of this buffer: `dataLenInBytes` specified in bytes.

For data plane APIs (`cpa_sym_dp.h` and `cpa_dc_dp.h`), a flat buffer is represented by the type `CpaPhysFlatBuffer`, also defined in `cpa.h`. This is similar to the `CpaFlatBuffer` structure; the difference is that, in this case, the data pointer, `bufferPhysAddr`, is a physical address rather than a virtual address.

[Figure 1](#) shows the layout of a flat buffer.

Figure 1. Flat Buffer Diagram



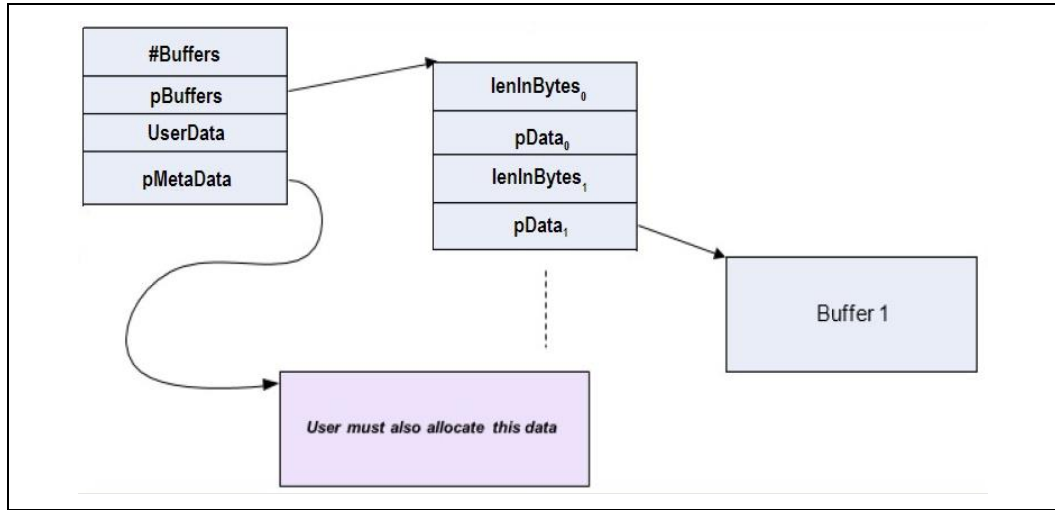
2.1.1.2 Scatter-Gather Lists

A scatter-gather list is defined by the type `CpaBufferList`, also defined in the file `cpa.h`. This buffer structure is typically used where more than one flat buffer can be provided to a particular API. The buffer list contains four fields, as follows:

- The number of buffers in the list.
- Pointer to an unbounded array of flat buffers.
- User Data: an opaque field; is not read or modified internally by the API.
- This field could be used to provide a pointer back into an application data structure, providing the context of the call.
- Pointer to metadata required by the API:
 - The metadata is required for internal use by the API. The memory for this buffer needs to be allocated by the client as contiguous data. The size of this metadata buffer is obtained by calling `cpaCyBufferListGetMetaSize` for crypto, `cpaBufferLists`, and `cpaDcBufferListGetMetaSize` for data compression.
 - The memory required to hold the `CpaBufferList` structure and the array of flat buffers is not required to be physically contiguous. However, the flat buffer data pointers and the metadata pointer are required to reference physically contiguous `DMAable` memory.
 - There is a performance impact when using scatter-gather lists instead of flat buffers. Refer to [Table 1](#), Intel® QAT Performance Optimization Guide for additional information.

[Figure 2](#) shows a graphical representation of a scatter-gather buffer list.

Figure 2. Scatter-Gather List Diagram



For data plane APIs (`cpa_sym_dp.h` and `cpa_dc_dp.h`) a region of memory that is not physically contiguous is described using the `CpaPhysBufferList` structure. This is similar to the `CpaBufferList` structure; the difference, in this case, the individual flat buffers are represented using physical rather than virtual addresses.

2.2 Intel® QuickAssist Technology API Conventions

2.2.1 Instance Discovery

Intel® QAT API supports multiple instances. An instance represents a "channel" to a specific hardware accelerator. Multiple instances can access the same hardware accelerator (that is, the relationship between instances and a hardware accelerator is N:1). The instance is identified using the `CpaInstanceHandle` handle type. This handle type represents a specific instance within the system and is passed as a parameter to all API functions that operate on instances.

Instance discovery is achieved through service-specific API invocations. This section describes the instance discovery for data compression (`dc`); however, the flow of the calls is similar for the cryptographic service.

Listing 1. Getting an Instance

```
void sampleDcGetInstance (CpaInstanceHandle *pDcInstHandle)
{
    CpaInstanceHandle dcInstHandles[MAX_INSTANCES];
    Cpa16U numInstances = 0;
    CpaStatus status = CPA_STATUS_SUCCESS;

    *pDcInstHandle = NULL;

    status = cpaDcGetNumInstances (&numInstances);
}
```

```

    if ((status == CPA_STATUS_SUCCESS) && (numInstances > 0)) {
        status = cpaDcGetInstances(MAX_INSTANCES, dcInstHandles);
        if (status == CPA_STATUS_SUCCESS) {
            *pDcInstHandle = dcInstHandles[0];
        }
    }

    if (0 == numInstances) {
        PRINT_ERR("No instances found for 'SSL'\n");
        PRINT_ERR("Please check your section names in the config
file.\n");
        PRINT_ERR("Also make sure to use config file version 2.\n");
    }
}

```

In this example, the number of `dc` instances available to the Application is queried via the `cpaDcGetNumInstances` call. The Application obtains the instance handle of the first instance.

The next example shows the Application querying the capabilities of the data compression implementation, and verifying the required functionality is present. Each service implementation exposes the capabilities that have been implemented and are available. Capabilities include algorithms, common features, and limits to variables. Each service has a unique capability matrix, and each implementation identifies and describes its particular implementation through its capability's API.

Listing 2. Querying and Starting an Instance

```

status = cpaDcQueryCapabilities(dcInstHandle, &cap);
if (status != CPA_STATUS_SUCCESS) {
    return status;
}
if (!cap.statelessDeflateCompression || !cap.statefulDeflateDecompression
||
    !cap.checksumCRC32 || !cap.dynamicHuffman) {
    PRINT_ERR("Error: Unsupported functionality\n");
    return CPA_STATUS_FAIL;
}
/*
 * Set the address translation function for the instance
 */
status = cpaDcSetAddressTranslation(dcInstHandle, sampleVirtToPhys);
if (CPA_STATUS_SUCCESS == status) {
    /* Start DataCompression component
    * In this example we are performing static compression so
    * an intermediate buffer is not required */
    PRINT_DBG("cpaDcStartInstance\n");
    status = cpaDcStartInstance(dcInstHandle, 0, NULL);
}

```

In the example, the application requires stateless deflate compression with dynamic Huffman encoding and stateful decompression with support for CRC32 checksums. The example also sets the address translation function for the instance. The specified function is used by the API to perform any required translation of a virtual address to a physical address. Finally, the instance is started.

2.2.2 Modes of Operation

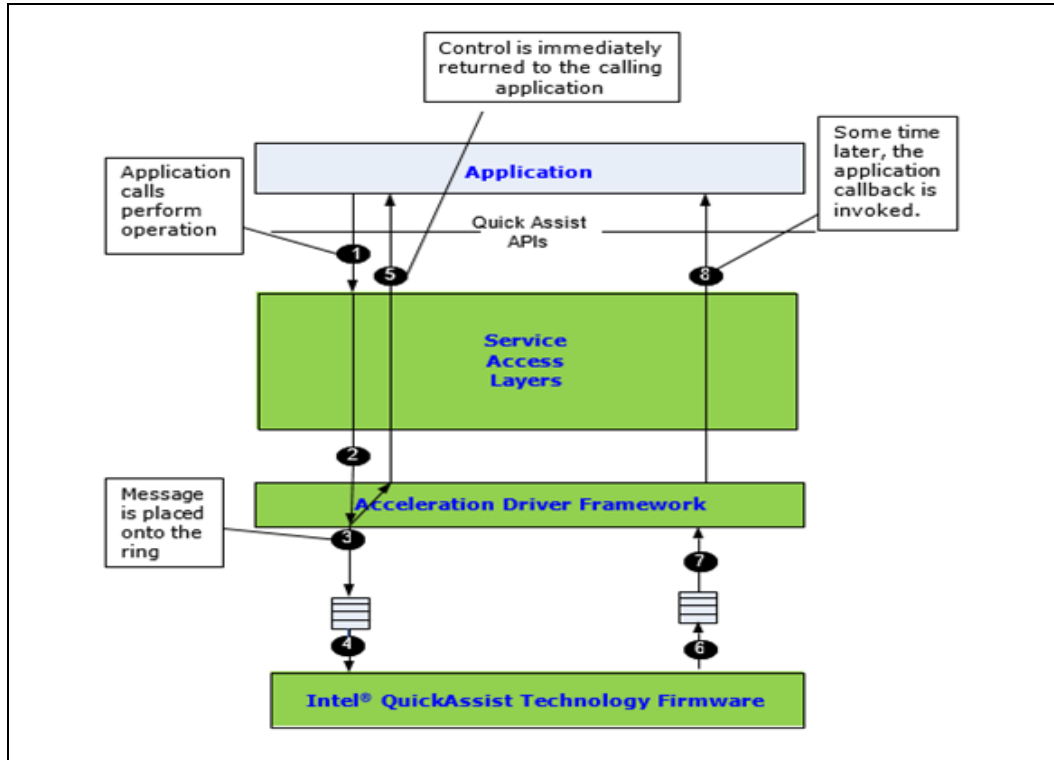
The Intel® QAT API supports both synchronous and asynchronous modes of operation. For optimal performance, the Application should be capable of submitting multiple outstanding requests to the acceleration engines. Submitting multiple outstanding requests minimizes the processing latency on the acceleration engines. This can be done by submitting requests asynchronously or by submitting requests in synchronous mode using multi-threading in the Application.

Developers can select the mode of operation that best aligns with their Application and system architecture.

2.2.2.1 Asynchronous Operation

To invoke the API asynchronously, the user supplies a callback function to the API, as shown in [Figure 3](#). Control returns to the client once the request has been sent to the hardware accelerator, and the callback is invoked when the engine completes the operation. The mechanism used to invoke the callback is implementation-dependent. For some implementations, the callback is invoked as part of an interrupt handler bottom half. For other implementations, the callback is invoked in the context of a polling thread. In this case, the user application is responsible for creating and scheduling this polling thread. Refer to [Table 1](#) for the implementation of specific documentation for more details.

Figure 3. Asynchronous Operation

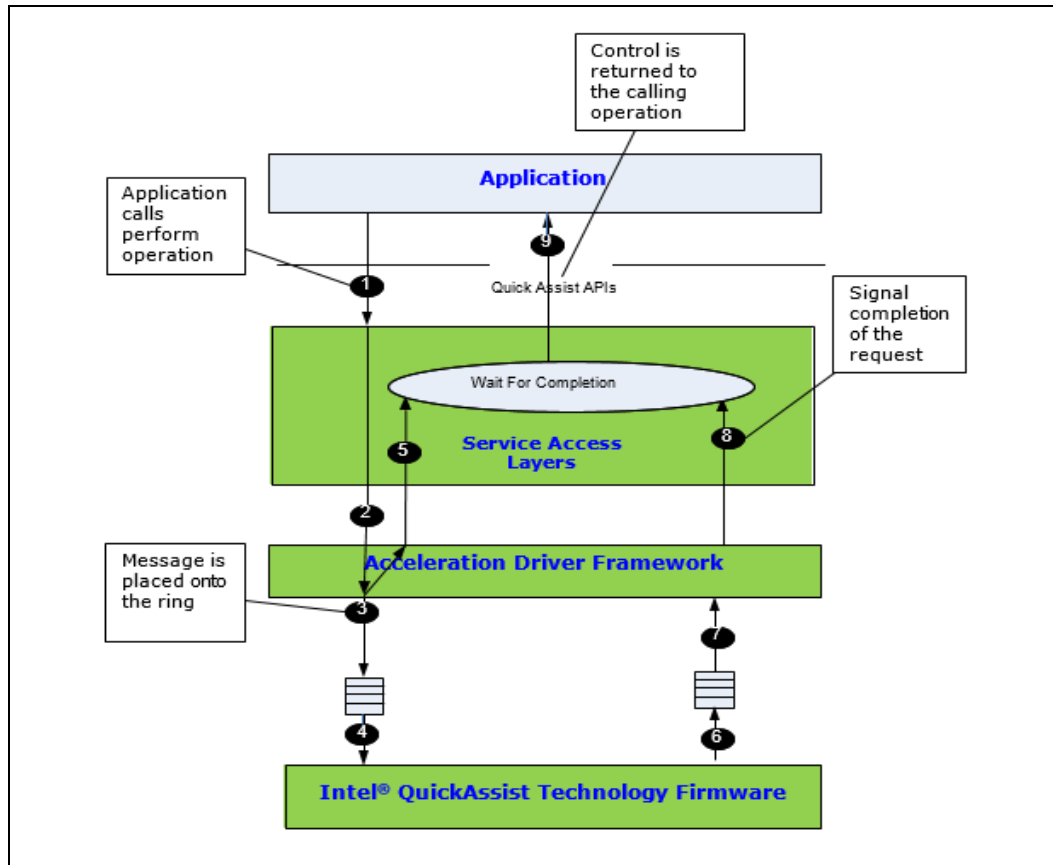


2.2.2.2 Synchronous Operation

Synchronous operation is specified by supplying a NULL function pointer in the callback parameter of the perform API, as shown in [Figure 4](#). In this case, the function does not return until the operation is complete. The calling thread may spend on a semaphore or other synchronization primitive after sending the request to the execution engine.

Upon the completion of the operation, the synchronization primitive unblocks, and execution resumes. Synchronous mode is therefore blocking and should not be used when invoking the function from a context in which sleeping is not allowed (for example, an interrupt context on Linux*).

Figure 4. Synchronous Operation



2.2.3 Memory Allocation and Ownership

The convention is that all memory needed by an API implementation is allocated outside of that implementation. In other words, the APIs are defined such that the memory needed to execute operations is supplied by a client or platform control entity rather than having memory allocated internally.

Memory used for parameters is owned by the side (caller or callee) that allocated the memory. An owner is responsible for de-allocating the memory when it is no longer needed.

Generally, memory ownership does not change. For example, if a program allocates memory and then passes a pointer to the memory as a parameter to a function call, the caller retains ownership and is still responsible for the de-allocation of the memory. Default behavior and any function which deviates from this behavior clearly state so in the function definition.

For optimal performance, data pointers should be 8-byte aligned. In some cases, this is a requirement, while in most other cases, it is a recommendation for performance. Refer to [Table 1](#) for the service-specific API manual for optimal usage of the particular API.

2.2.4 Data Plane APIs

The Intel® QAT APIs for symmetric cryptography and for data compression supports both “traditional” (`cpa_cy_sym.h` and `cpa_dc.h`) and “data plane” APIs (`cpa_cy_sym_dp.h` and `cpa_dc_dp.h`)¹. The data plane APIs are recommended for applications running in a data plane environment where the cost of offload (that is, the cycles consumed by the driver sending requests to the accelerator) needs to be minimized. Several constraints have been placed on these APIs to minimize the cost of offload. If these constraints are too restrictive for a given application, the more general-purpose “traditional” APIs can be used (at an increased cost of offload).

The data plane APIs can be used if the following constraints are acceptable:

- There is no support for partial packets or stateful requests.
- Thread safety is not supported. Each software thread should have access to its unique instance (`CpaInstanceHandle`).
- Only asynchronous invocation is supported.
- Polling is used, rather than interrupts, to dispatch callback functions. Callbacks are invoked in the context of a polling thread.
 - The user application is responsible for creating and scheduling this polling thread. Polling functions are not defined by the Intel® QAT API. Implementations provide their polling functions. Refer to [Table 1](#) for Implementation Specific Documentation containing further information on polling functions.
- Buffers and buffer lists are passed using physical addresses to avoid virtual-to-physical-address translation costs.
- Alignment restrictions may be placed on the operation data (that is, `CpaCySymDpOpData` and `CpaDcDpOpData`) and buffer list (that is, `CpaPhysBufferList`) structures passed to the data plane APIs. For example, the operation data may need to be at least 8-byte aligned, contiguous, resident, DMA-accessible memory. Refer to [Table 1](#) for Implementation Specific Documentation for more details.
- For CCM and GCM modes of the AES, when performing decryption and verification, if the verification fails, then the message buffer is not zeroed.

The data plane APIs distinguish between enqueueing a request and submitting that request to the accelerator to be performed. This allows the cost of submitting a request (which can be expensive, in terms of cycles, for some hardware-based implementations) to be amortized over all enqueued requests on that instance (`CpaInstanceHandle`).

- To enqueue one request and to optionally submit all previously enqueued requests, the function `cpaCySymDpEnqueueOp` (or `cpaDcDpEnqueueOp` for data compression service) can be used.
- To enqueue multiple requests and to optionally submit all previously enqueued requests, the function `cpaCySymDpEnqueueOpBatch` (or `cpaDcDpEnqueueOpBatch` for data compression service) can be used.

¹ There is no “data plane” support for asymmetric cryptography services.

- Use the function `cpaCySymDpPerformOpNow` (or `cpaDcDpPerformOpNow` for data compression service) that can be used to submit all previously enqueued requests.
- Different implementations of this API may have different performance trade-offs. Refer to [Table 1](#) for documentation for implementation details.

§

3 Intel® QuickAssist Technology Cryptographic API

This chapter describes the sample code for the Intel® QuickAssist Technology Cryptographic API, beginning with an API overview, and followed by descriptions of various scenarios to illustrate the usage of the API.

3.1 Overview

The Intel® QuickAssist Technology Cryptographic API can be categorized into the following broad areas:

- **Common:** This is defined by the file `cpa_cy_common.h`. This includes the functionality for the initialization and shutdown of the service.
- **Instance Management:** The file `cpa_cy_im.h` defines the functions for managing instances. A given implementation of the API can present multiple instances of the cryptographic service, each representing a logical or virtual "device". Request order is guaranteed within a given instance of the service.
- **Symmetric:** The following files constitute the symmetric API:
 - The `cpa_cy_sym.h` file contains the symmetric API, used for ciphers, hashing/message digests, "algorithm chaining" (combining cipher and hash into a single call), and authenticated ciphers.
 - The `cpa_sy_sym_dp.h` file also contains the symmetric API, used for ciphers, hashing/message digest, "algorithm chaining" (combining cipher and hash into a single call) and authenticated ciphers. This API is recommended for data plane applications, in which the cost of offload (i.e. the cycles consumed by the API in sending requests to the hardware and processing the responses) needs to be minimized. Several constraints need to be acceptable to the Application. To use this API, these are listed in Section [2.2.4, Data Plane APIs](#).
 - The `cpa_cy_key.h` file contains the API for key generation for Secure Sockets Layer (SSL) and Transport Layer Security (TLS).
- **Asymmetric:** The following files constitute the asymmetric API:
 - The `cpa_cy_rsa.h` file defines the API for RSA.
 - The `cpa_cy_dsa.h` file defines the API for Digital Signature Algorithm (DSA).
 - The `cpa_cy_dh.h` file defines the API for Diffie-Hellman.
 - The `cpa_cy_ec.h` file defines the API for "base" elliptic curve cryptography.
 - The `cpa_cy_ecdsa.h` file defines the API for Elliptic Diffie (EC) DSA.
 - The `cpa_cy_ecdh.h` file defines the API for Elliptic Diffie-Hellman (ECDH).
 - The `cpa_cy_prime.h` file defines the API for prime number testing.
 - The `cpa_cy_ln.h` file defines the API for a large number of math operations, such as modular exponentiation, etc.
 - The `cpa_cy_ecsm2.h` file defines the API for SM2 algorithm which is based on Elliptic Curves Cryptography (ECC)
- **Random Bit Generation (RBG):** The following files constitute the RBG API and have been deprecated because random bit generation can be handled in the CPU:
 - The `cpa_cy_drbg.h` file defines the API for deterministic random bit generation.

- The `cpa_cy_nrbg.h` file defines the API for a non-deterministic random bit generation.

The Cryptographic API uses the base API, which defines base data types used across all services of the Intel® QAT API.

3.1.1 Sessions

The symmetric API is the only API with the concept of sessions. The meaning of a session within the symmetric API is defined below.

3.1.2 Priority

The Cryptographic symmetric API has support for priorities. Priority can be specified on a per-session basis. Two levels of priority are supported: high priority and normal priority. Implementations may use a strict priority order or a weighted round robin-based priority scheme.

3.2 Using the Symmetric Cryptography API

This section contains examples of how to use the symmetric API. It describes general concepts and how to use the symmetric API to perform various types of cipher and hash.

Note: Examples are simplified and demonstrate how to use the APIs and build the structures required for various use cases.

These examples may not demonstrate the optimal way to use the API to get maximum performance for a particular implementation. Refer to [Table 1](#) for Implementation Specific Documentation and performance sample code for a guide on how to use the API for best performance.

Note: All of the symmetric examples follow the same basic steps:

- Define a callback function (if the API is to be invoked asynchronously)
- Discover and start up the cryptographic service instance
- Create and initialize a session
- Invoke multiple symmetric operations (cipher and/or hash) on the session
- Tear down the session
- Stop the Cryptographic service instance

3.2.1 General Concepts

This section describes the following concepts:

- Session
- Place and Out-of-Place Support

- Partial Support

3.2.1.1 Session

In case of the symmetric API, a session is a handle that describes the cryptographic parameters to be applied to several buffers. This might be the buffers within a single file or all the packets associated with a particular Internet Protocol Security (IPSec) tunnel or security association. The data within a session handle includes the following:

- The operation (cipher, hash, or both, and if both, the order in which the algorithms should be applied).
- The cipher setup data, including the cipher algorithm and mode, the key and its length, and the direction (encrypt or decrypt).
- The hash setup data, including the hash algorithm, mode (plain, nested or authenticated), and digest result length (to allow for truncation).
 - The authenticated mode can refer to Hashed Message Authenticate Code (HMAC), which requires that the key and its length are also specified. It is also used for Galois Counter Mode (GCM), and Counter mode with Cipher-block Chaining Message authentication code (CCM) authenticated encryption, in which case the Additional Authenticated Data (AAD) length is also specified.
 - For nested mode, the inner and outer prefix data and length are specified, as well as the outer hash algorithm.

3.2.1.2 In-Place and Out-of-Place Support

An In-Place operation means that the destination buffer is the same as the source buffer. An Out-of-Place operation means that the destination buffer is different from the source buffer.

3.2.1.3 Partial Support

Most of the examples in this chapter operate on full packets, as indicated by the `packetType` of `CPA_CY_SYM_PACKET_TYPE_FULL`. The API also supports operating in partial mode, where, for example, state (e.g., cipher state) needs to be carried forward from one packet/record to the next. In Section [3.2.4, Hash a File](#), there is an example of hashing a file that uses the partial API.

Note:

1. The size of the data to be hashed or ciphered must be a multiple of the block size of the algorithm for all partial packets.
2. For hash/authentication, the digest verify flag only applies to the last partial packet.
3. For algorithm chaining, only the cipher state is maintained between calls. The hash state is not maintained between calls; instead, the hash digest is generated/verified for each call. The size of the data to be ciphered must be a multiple of the block size of the algorithm for all partial packets. The size of the data to be hashed does not have this restriction. If both the cipher state and the hash state need to be maintained between calls, then algorithm chaining cannot be used.

3.2.2 Cipher

This example demonstrates the usage of the symmetric API, specifically using this API to perform a cipher operation. It encrypts some sample text using the AES-256 algorithm in Cipher Block Chaining (CBC) mode.

These samples are located in:

```
quickassist/lookaside/access_layer/src/sample_code/functional/sym/cipher_sample
```

The following subsections describe the main functions in this file.

3.2.2.1 symCallback

A callback function must be supplied to use the API in asynchronous mode, and this function is called back (that is, invoked by the implementation of the API) when the asynchronous operation has completed. The context in which it is invoked depends on the implementation. For example, it could be invoked in the context of a Linux* interrupt handler's bottom half or in the context of a user created polling thread. The context in which this function is invoked places restrictions on what processing can be done in the callback function. On the API, it states that this function should not sleep (since it may be called in a context that does not permit sleeping, for example, a Linux* bottom half).

This function can perform whatever processing is appropriate for the Application. For example, it may free memory, continue the processing of a decrypted packet, etc. In this example, the function only sets the complete variable to indicate it has been called, as illustrated below.

Listing 3. Callback Function

```
static void symCallback(void *pCallbackTag,
                       CpaStatus status,
                       const CpaCySymOp operationType,
                       void *pOpData,
                       CpaBufferList *pDstBuffer,
                       CpaBoolean verifyResult)
{
    PRINT_DBG("Callback called with status = %d.\n", status);

    if (NULL != pCallbackTag) {
        /* indicate that the function has been called */
        COMPLETE((struct COMPLETION_STRUCT *)pCallbackTag);
    }
}
```

3.2.2.2 cipherSample

This is the main entry point for the sample cipher code. It demonstrates the sequence of calls to be made to the API to create a session, perform one or more cipher operations, and then tear down the session. The following is performed:

- Call the instance discovery utility function - `sampleCyGetInstance` - which is a simplified version of instance discovery, in which exactly one instance of a crypto service is discovered. It does this by querying the API for all instances, and returning the first instance, as illustrated in [Listing 4](#).
- This step is described in Section [2.2.1, Instance Discovery](#), but is repeated here for convenience.

Listing 4. Getting an Instance

```
#ifdef DO_CRYPTO
void sampleCyGetInstance(CpaInstanceHandle* pCyInstHandle)
{
    CpaInstanceHandle cyInstHandles[MAX_INSTANCES];
    Cpa16U numInstances = 0;
    CpaStatus status = CPA_STATUS_SUCCESS;

    *pCyInstHandle = NULL;

    status = cpaCyGetNumInstances(&numInstances);
    if ((status == CPA_STATUS_SUCCESS) && (numInstances > 0))
    {
        status = cpaCyGetInstances(MAX_INSTANCES, cyInstHandles);
        if (status == CPA_STATUS_SUCCESS)
        {
            *pCyInstHandle = cyInstHandles[0];
        }
    }
    if (0 == numInstances)
    {
        PRINT_ERR("No instances found for 'SSL'\n");
        PRINT_ERR("Please check your section names in the config
file.\n");
        PRINT_ERR("Also make sure to use config file version 2.\n");
    }
}
#endif
```

Set the address translation function for the instance. This function will be used by the API to convert virtual addresses to physical addresses.

Listing 5. Set Address Translation Function

```
status = cpaCySetAddressTranslation(cyInstHandle,
sampleVirtToPhys);
```

Start the crypto service running as shown below.

Listing 6. Start up

```
status = cpaCyStartInstance(cyInstHandle);
```

The next step is to create and initialize a session. First, populate the fields of the session initialization operational data structure.

Note: The size required to store a session is implementation-dependent, so you must query the API first to determine how much memory to allocate, and then allocate that memory.

One of two available queries can be used:

- `cpaCySymSessionCtxGetSize(const CpaInstanceHandle instanceHandle_in, const CpaCySymSessionSetupData *pSessionSetupData, Cpa32U *pSessionCtxSizeInBytes)`
- This will always return the maximum session context size (i.e., the full size of the session including padding and other session state information) ([see Listing 7 below](#)).
- `cpaCySymSessionCtxGetDynamicSize(const CpaInstanceHandle instanceHandle_in, const CpaCySymSessionSetupData *pSessionSetupData, Cpa32U *pSessionCtxSizeInBytes)`
- This query can be used instead to return a reduced memory size, based on whether the use case meets certain session setup criteria ([see Listing 7 below](#)).
- This query will return one of three values for `pSessionCtxSizeInBytes` as follows:
 - If partial packets are not being used and the Symmetric operation is Auth-Encrypt (i.e., the cipher and hash algorithms are either CCM or GCM), the size returned will be approximately half of the standard size.
 - If partial packets are not being used and the cipher algorithm is not `ARC4`, `Snow3g_UEA2`, `AES_CCM` or `AES_GCM`, and the hash algorithm is not `Snow3G_UIA2`, `AES_CCM` or `AES_GCM`, and Hash Mode is not Auth, the size returned will be between half and one third of the standard size.
 - In all other cases, the standard size is returned.

Note: The following parameter exists in the `CpaCySymSessionSetupData` structure:

```
CpaBoolean partialsNotRequired
```

This flag indicates if partial packet processing is required for the session. If partial packets are not being used and the preference is to use one of the reduced session memory sizes, set this flag to `CPA_TRUE` before calling the `cpaCySymSessionCtxGetDynamicSize()` function.

Note: The equivalent reduced memory context query for Data Plane API (Refer to Section [3.2.10, Chained Cipher and Hash Using the Symmetric Data Plane API](#)) is:

```
cpaCySymDpSessionCtxGetDynamicSize(const CpaInstanceHandle
instanceHandle_in, const CpaCySymSessionSetupData *pSessionSetupData,
Cpa32U *pSessionCtxSizeInBytes)
```

Listing 7. Create and Initialize Cipher Session

```
/* Populate the session setup structure for the operation required */
sessionSetupData.sessionPriority = CPA_CY_PRIORITY_NORMAL;
sessionSetupData.symOperation = CPA_CY_SYM_OP_CIPHER;
sessionSetupData.cipherSetupData.cipherAlgorithm =
CPA_CY_SYM_CIPHER_AES_CBC;
```

```

sessionSetupData.cipherSetupData.pCipherKey = sampleCipherKey;
sessionSetupData.cipherSetupData.cipherKeyLenInBytes =
sizeof(sampleCipherKey);
sessionSetupData.cipherSetupData.cipherDirection =
    CPA_CY_SYM_CIPHER_DIRECTION_ENCRYPT;

/* Determine size of session context to allocate */
PRINT_DBG("cpaCySymSessionCtxGetSize\n");
status =
    cpaCySymSessionCtxGetSize(cyInstHandle, &sessionSetupData,
&sessionCtxSize);
}

if (CPA_STATUS_SUCCESS == status) {
    /* Allocate session context */
    status = PHYS_CONTIG_ALLOC(&sessionCtx, sessionCtxSize);
}

/* Initialize the Cipher session */
if (CPA_STATUS_SUCCESS == status) {
    PRINT_DBG("cpaCySymInitSession\n");
    status = cpaCySymInitSession(cyInstHandle,
/*                                     symCallback,          /* callback function
*/
/*                                     &sessionSetupData, /* session setup data
*/
/*                                     sessionCtx);          /* output of the
function*/
}

```

- Call the function `cipherPerformOp`, which actually performs the cipher operation. This in turn performs the following steps:

Memory Allocation: Different implementations of the API require different amounts of space to store metadata associated with buffer lists. Query the API to find out how much space the current implementation needs, and then allocate space for the buffer metadata, the buffer list, and for the buffer itself. You must also allocate memory for the initialization vector.

Listing 8. Memory Allocation

```

status = cpaCyBufferListGetMetaSize(cyInstHandle, numBuffers,
&bufferMetaSize);
if (CPA_STATUS_SUCCESS == status) {
    status = PHYS_CONTIG_ALLOC(&pBufferMeta, bufferMetaSize);
}
if (CPA_STATUS_SUCCESS == status) {
    status = OS_MALLOC(&pBufferList, bufferListMemSize);
}
if (CPA_STATUS_SUCCESS == status) {

```

```

    status = PHYS_CONTIG_ALLOC(&pSrcBuffer, bufferSize);
}
if (CPA_STATUS_SUCCESS == status) {
    status = PHYS_CONTIG_ALLOC(&pIvBuffer, sizeof(sampleCipherIv));
}

```

- **Memory Allocation:** The memory for the source buffer and initialization vector is populated with the required data.
- **Set Up Operational Data:** Populate the structure containing the operational data that is needed to run the algorithm as shown below.

Listing 9. Set Up Cipher Operational Data

```

pOpData->sessionCtx = sessionCtx;
pOpData->packetType = CPA_CY_SYM_PACKET_TYPE_FULL;
pOpData->pIv = pIvBuffer;
pOpData->ivLenInBytes = sizeof(sampleCipherIv);
pOpData->cryptoStartSrcOffsetInBytes = 0;
pOpData->messageLenToCipherInBytes = sizeof(sampleCipherSrc);

```

- **Perform Operation:** Initialize the completion variable, which is used by the callback function to indicate that the operation is complete, then perform the operation.

Listing 10. Perform Cipher Operation

```

COMPLETION_INIT(&complete);

status = cpaCySymPerformOp(
    cyInstHandle,
    (void *)&complete, /* data sent as is to the callback function*/
    pOpData,           /* operational data struct */
    pBufferList,       /* source buffer list */
    pBufferList,       /* same src & dst for an in-place operation*/
    NULL);

```

- **Wait for Completion:** Because the asynchronous API is used in this example, the callback function must be handled. This example uses a macro that can be defined differently for different operating systems. In a typical real-world application, the calling thread would not block, and the callback would essentially re-inject the (decrypted, decapsulated) packet into the stack.

Listing 11. Wait for Completion

```

if (!COMPLETION_WAIT(&complete, TIMEOUT_MS)) {
    PRINT_ERR("timeout or interruption in cpaCySymPerformOp\n");
    status = CPA_STATUS_FAIL;
}

```

In a normal usage scenario, the session would be reused multiple times to encrypt multiple buffers or packets. In this example, however, the session is torn down.

Listing 12. Wait for Outstanding Requests

```
symSessionWaitForInflightReq(sessionCtx)
```

Since cryptographic API v2.2 before removing the symmetric session context it is recommended to wait for the completion of any outstanding request using `cpaCySymSessionInUse`.

It is executed in the `symSessionWaitForInflightReq` call which polls for the in-flight requests.

Listing 13. Remove Cipher Session

```
sessionStatus = cpaCySymRemoveSession(cyInstHandle, sessionCtx);
```

- Query statistics at this point, which can be useful for debugging.
Some implementations may also make the statistics available through other mechanisms, such as the `/proc` virtual filesystem.
- Finally, clean up by freeing up memory, stopping the instance, etc.
- Since Cryptographic API v2.2 two new functions have been implemented:
`cpaCySymUpdateSession` and `cpaCySymSessionInUse`
- The function `cpaCySymUpdateSession` can be used to update certain parameters of a session like the cipher key, the cipher direction, and the authentication key. `cpaCySymSessionInUse`, indicates whether there are outstanding requests on a given session.

As a result of the implementation of this feature, the behavior of `cpaCySymRemoveSession` has been changed. `cpaCySymRemoveSession` will fail if there are outstanding request for the session that the user is trying to remove.

As a result, it is recommended to wait for the completion of any outstanding request, using `cpaCySymSessionInUse`, before removing a session.

3.2.3 Hash

This example demonstrates the usage of the symmetric API, specifically using this API to perform a hash operation. It performs an SHA-256 hash operation on some sample data.

These samples are located in `/sym/hash_sample`

The example is very similar to the cipher example, so only the differences are highlighted:

When creating and initializing a session, some of the fields of the session initialization operational data structure are different from the cipher case, as shown below.

Listing 14. Create and Initialize Hash Session

```
/* populate symmetric session data structure
 * for a plain hash operation */
sessionSetupData.sessionPriority = CPA_CY_PRIORITY_NORMAL;
sessionSetupData.symOperation = CPA_CY_SYM_OP_HASH;
```

```

sessionSetupData.hashSetupData.hashAlgorithm = CPA_CY_SYM_HASH_SHA256;
sessionSetupData.hashSetupData.hashMode = CPA_CY_SYM_HASH_MODE_PLAIN;
sessionSetupData.hashSetupData.digestResultLenInBytes = DIGEST_LENGTH;
/* Place the digest result in a buffer unrelated to srcBuffer */
sessionSetupData.digestIsAppended = CPA_FALSE;
/* Generate the digest */
sessionSetupData.verifyDigest = CPA_FALSE;

```

When calling the function to perform the hash operation, some of the fields of the operational data structure are again different from the cipher case, as shown below.

Listing 15. Set up Hash Operational Data

```

pOpData->sessionCtx = sessionCtx;
pOpData->packetType = CPA_CY_SYM_PACKET_TYPE_FULL;
pOpData->hashStartSrcOffsetInBytes = 0;
pOpData->messageLenToHashInBytes = sizeof(vectorData);
pOpData->pDigestResult = pDigestBuffer;

```

3.2.4 Hash a File

This example demonstrates the usage of the symmetric API for partial mode, specifically using this API to perform hash operations. It performs a SHA1 hash operation on a file.

These samples are located in `/sym/hash_file_sample`

The example is very similar to the cipher example, so only the differences are highlighted:

When creating and initializing a session, some of the fields of the session initialization operational data structure are different from the cipher case, as shown below.

Listing 16. Hash Session Setup Data

```

/* populate symmetric session data structure
 * for a plain hash operation */
sessionSetupData.sessionPriority = CPA_CY_PRIORITY_NORMAL;
sessionSetupData.symOperation = CPA_CY_SYM_OP_HASH;
sessionSetupData.hashSetupData.hashAlgorithm = CPA_CY_SYM_HASH_SHA1;
sessionSetupData.hashSetupData.hashMode = CPA_CY_SYM_HASH_MODE_PLAIN;
sessionSetupData.hashSetupData.digestResultLenInBytes = DIGEST_LENGTH;
/* Place the digest result in a buffer unrelated to srcBuffer */
sessionSetupData.digestIsAppended = CPA_FALSE;
/* Generate the digest */
sessionSetupData.verifyDigest = CPA_FALSE;

```

- Memory is allocated for the source buffer in a similar way to the cipher case.
- To perform the operation data is read from the file to the source buffer and the symmetric API is called repeatedly with `packetType` set to `CPA_CY_SYM_PACKET_TYPE_PARTIAL`. When the end of the file is reached the API is

called with `packetType` set to `CPA_CY_SYM_PACKET_TYPE_PARTIAL_LAST`. The digest is produced only on the last call to the API.

Listing 17. Hashing a File

```
while (!feof(srcFile)) {
    /* read from file into src buffer */
    pBufferList->pBuffers->dataLenInBytes =
        fread(pSrcBuffer, 1, SAMPLE_BUFF_SIZE, srcFile);

    /* If we have reached the end of file set the last partial flag */
    if (feof(srcFile)) {
        pOpData->packetType = CPA_CY_SYM_PACKET_TYPE_LAST_PARTIAL;
    } else {
        pOpData->packetType = CPA_CY_SYM_PACKET_TYPE_PARTIAL;
    }
    pOpData->sessionCtx = sessionCtx;
    pOpData->hashStartSrcOffsetInBytes = 0;
    pOpData->messageLenToHashInBytes = pBufferList->pBuffers-
>dataLenInBytes;
    pOpData->pDigestResult = pDigestBuffer;

    PRINT_DBG("cpaCySymPerformOp\n");
    /* Perform symmetric operation */
    status = cpaCySymPerformOp(
        cyInstHandle,
        (void *)&complete, /* data sent as is to the callback function*/
        pOpData, /* operational data struct */
        pBufferList, /* source buffer list */
        pBufferList, /* same src & dst for an in-place operation*/
        NULL);

    if (CPA_STATUS_SUCCESS != status) {
        PRINT_ERR("cpaCySymPerformOp failed. (status = %d)\n", status);
        break;
    }

    if (CPA_STATUS_SUCCESS == status) {
        /** wait until the completion of the operation*/
        if (!COMPLETION_WAIT(&complete), TIMEOUT_MS)) {
            PRINT_ERR("timeout or interruption in cpaCySymPerformOp\n");
            status = CPA_STATUS_FAIL;
            break;
        }
    }
}
}
```

3.2.5 Chained Cipher and Hash

This example demonstrates the usage of the symmetric API, specifically using this API to perform a "chained" cipher and hash operation. It encrypts some sample text using the AES-256 algorithm in CBC mode, and then performs an SHA-256 Hashed Message Authenticate Code (HMAC) operation on the ciphertext, writing the Message Authentication Code (MAC) to the buffer immediately after the ciphertext.

These samples are located in `/sym/alg_chaining_sample`

The example is very similar to the cipher and hash examples, above, so only the differences are highlighted:

When creating and initializing a session, some of the fields of the session initialization operational data structure are different, as shown below.

Listing 18. Create and Initialize Session Cipher and Hash

```
sessionSetupData.symOperation = CPA_CY_SYM_OP_ALGORITHM_CHAINING;
sessionSetupData.algChainOrder =
CPA_CY_SYM_ALG_CHAIN_ORDER_CIPHER_THEN_HASH;

sessionSetupData.cipherSetupData.cipherAlgorithm =
CPA_CY_SYM_CIPHER_AES_CBC;
sessionSetupData.cipherSetupData.pCipherKey = sampleCipherKey;
sessionSetupData.cipherSetupData.cipherKeyLenInBytes =
sizeof(sampleCipherKey);
sessionSetupData.cipherSetupData.cipherDirection =
CPA_CY_SYM_CIPHER_DIRECTION_ENCRYPT;

sessionSetupData.hashSetupData.hashAlgorithm = CPA_CY_SYM_HASH_SHA256;
sessionSetupData.hashSetupData.hashMode = CPA_CY_SYM_HASH_MODE_AUTH;
sessionSetupData.hashSetupData.digestResultLenInBytes = DIGEST_LENGTH;
sessionSetupData.hashSetupData.authModeSetupData.authKey =
sampleCipherKey;
sessionSetupData.hashSetupData.authModeSetupData.authKeyLenInBytes =
sizeof(sampleCipherKey);

/* The resulting MAC is to be placed immediately after the ciphertext */
sessionSetupData.digestIsAppended = CPA_TRUE;
sessionSetupData.verifyDigest = CPA_FALSE;
```

When calling the function to perform the chained cipher and hash operation, some of the fields of the operational data structure are again different from the cipher case, as shown below.

Listing 19. Set up Operational Data Cipher and Hash

```
/** Populate the structure containing the operational data that is
 * needed to run the algorithm
 */
```



```

pOpData->sessionCtx = sessionCtx;
pOpData->packetType = CPA_CY_SYM_PACKET_TYPE_FULL;
pOpData->pIv = pIvBuffer;
pOpData->ivLenInBytes = sizeof(sampleCipherIv);
pOpData->cryptoStartSrcOffsetInBytes = 0;
pOpData->hashStartSrcOffsetInBytes = 0;
pOpData->messageLenToCipherInBytes = sizeof(sampleAlgChainingSrc);
pOpData->messageLenToHashInBytes = sizeof(sampleAlgChainingSrc);
    
```

Notice the `digestIsAppended` is set in the session; therefore, the MAC is placed immediately after the region to hash, and the `pDigestResult` parameter of the operational data is ignored.

3.2.6 Chained Cipher and Hash – IPsec Like Use Case

This example demonstrates the usage of the symmetric API for IPsec-like use cases, as described in [Figure 5](#) and [Figure 6](#). For the outbound direction, this example uses the symmetric API to perform a "chained" cipher and hash operation. It encrypts some plaintext using the Advanced Encryption Standard (AES) algorithm in CBC mode, and then performs a SHA1 HMAC operation on the ciphertext, initialization vector, and header, writing the Integrity Check Value (ICV) to the buffer immediately after the ciphertext. For the inbound direction, this example again uses the symmetric API to perform a "chained" hash and cipher operation. It performs a SHA1 HMAC operation on the ciphertext, initialization vector, and Header and compares the Result with the input ICV. Then it decrypts the ciphertext using the AES algorithm in CBC mode.

Figure 5. IPsec Outbound

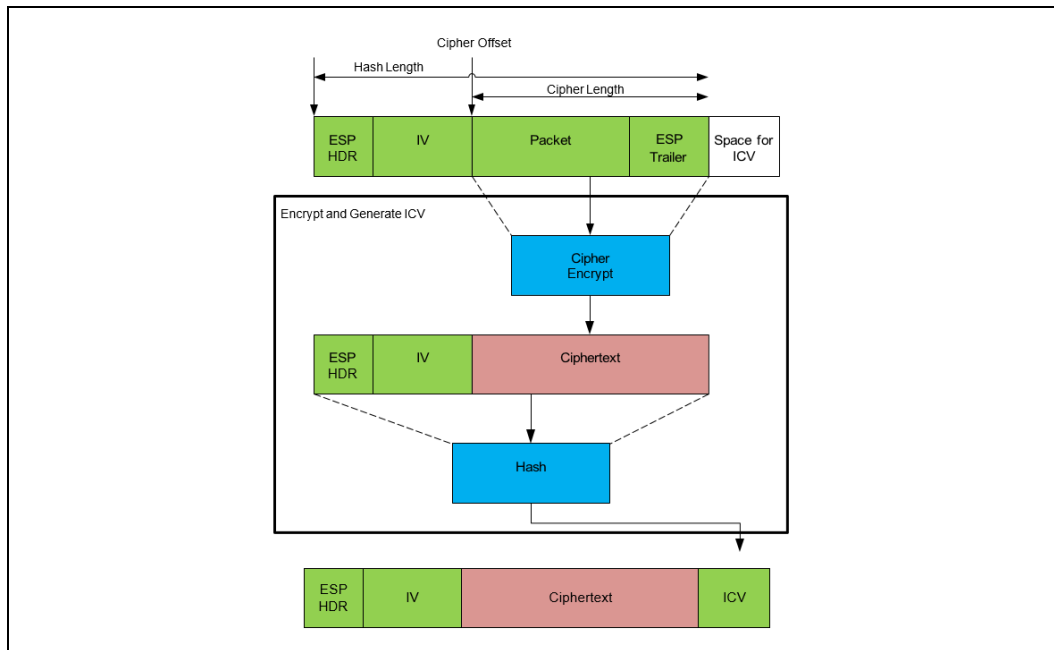
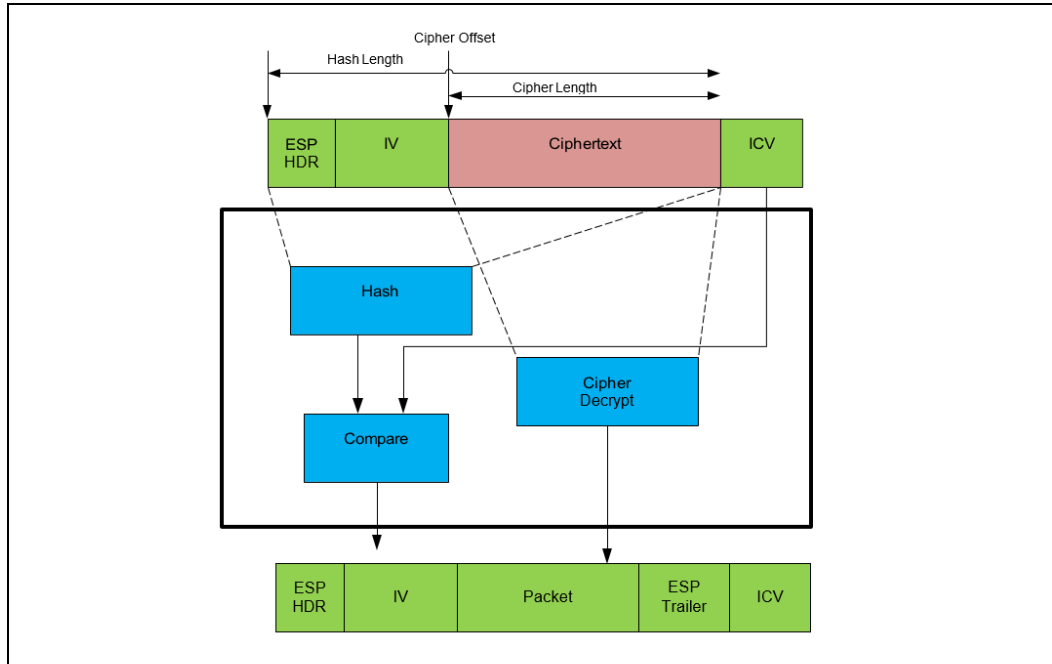


Figure 6. IPSec Inbound



These samples are located in `/sym/ipsec_sample`.

Again, only the differences compared to previous examples are highlighted:

When creating and initializing a session in the outbound direction, the session initialization operational data structure is shown below.

Listing 20. Session Setup Data IPSec Outbound

```

sessionSetupData.symOperation = CPA_CY_SYM_OP_ALGORITHM_CHAINING;
sessionSetupData.algChainOrder =
CPA_CY_SYM_ALG_CHAIN_ORDER_CIPHER_THEN_HASH;

sessionSetupData.cipherSetupData.cipherAlgorithm =
CPA_CY_SYM_CIPHER_AES_CBC;
sessionSetupData.cipherSetupData.pCipherKey = sampleCipherKey;
sessionSetupData.cipherSetupData.cipherKeyLenInBytes =
sizeof(sampleCipherKey);
sessionSetupData.cipherSetupData.cipherDirection =
CPA_CY_SYM_CIPHER_DIRECTION_ENCRYPT;

sessionSetupData.hashSetupData.hashAlgorithm = CPA_CY_SYM_HASH_SHA1;
sessionSetupData.hashSetupData.hashMode = CPA_CY_SYM_HASH_MODE_AUTH;
sessionSetupData.hashSetupData.digestResultLenInBytes = ICV_LENGTH;
sessionSetupData.hashSetupData.authModeSetupData.authKey = sampleAuthKey;
sessionSetupData.hashSetupData.authModeSetupData.authKeyLenInBytes =
sizeof(sampleAuthKey);

```

```

/* Even though ICV follows immediately after the region to hash
 * digestIsAppended is set to false in this case to workaround
 * errata number IXA00378322 */
sessionSetupData.digestIsAppended = CPA_FALSE;
/* Generate the ICV in outbound direction */
sessionSetupData.verifyDigest = CPA_FALSE;

```

When calling the function to perform the chained cipher and hash operation, the fields of the operational data structure are shown below.

Listing 21. Operational Data IPsec Outbound

```

/* Populate the structure containing the operational data that is
 * needed to run the algorithm in outbound direction */
pOpData->sessionCtx = sessionCtx;
pOpData->packetType = CPA_CY_SYM_PACKET_TYPE_FULL;
pOpData->pIv = pIvBuffer;
pOpData->ivLenInBytes = sizeof(sampleCipherIv);
pOpData->cryptoStartSrcOffsetInBytes =
sizeof(sampleEspHdrData) + sizeof(sampleCipherIv);
pOpData->messageLenToCipherInBytes = sizeof(samplePayload);
pOpData->hashStartSrcOffsetInBytes = 0;
pOpData->messageLenToHashInBytes =
sizeof(sampleEspHdrData) + sizeof(sampleCipherIv) +
sizeof(samplePayload);
/* Even though ICV follows immediately after the region to hash
 * digestIsAppended is set to false in this case to workaround
 * errata number IXA00378322 */
pOpData->pDigestResult =
pSrcBuffer +
(sizeof(sampleEspHdrData) + sizeof(sampleCipherIv) +
sizeof(samplePayload));

```

In this example `samplePayload` is the packet data plus the Encapsulating Security Payload (ESP) trailer.

When creating and initializing a session in the inbound direction, the session initialization operational data structure is shown below.

Listing 22. Session Setup Data IPsec Inbound

```

sessionSetupData.symOperation = CPA_CY_SYM_OP_ALGORITHM_CHAINING;
sessionSetupData.algChainOrder =
CPA_CY_SYM_ALG_CHAIN_ORDER_HASH_THEN_CIPHER;

sessionSetupData.cipherSetupData.cipherAlgorithm =
CPA_CY_SYM_CIPHER_AES_CBC;
sessionSetupData.cipherSetupData.pCipherKey = sampleCipherKey;

```

```

sessionSetupData.cipherSetupData.cipherKeyLenInBytes =
sizeof(sampleCipherKey);
sessionSetupData.cipherSetupData.cipherDirection =
CPA_CY_SYM_CIPHER_DIRECTION_DECRYPT;

sessionSetupData.hashSetupData.hashAlgorithm = CPA_CY_SYM_HASH_SHA1;
sessionSetupData.hashSetupData.hashMode = CPA_CY_SYM_HASH_MODE_AUTH;
sessionSetupData.hashSetupData.digestResultLenInBytes = ICV_LENGTH;
sessionSetupData.hashSetupData.authModeSetupData.authKey = sampleAuthKey;
sessionSetupData.hashSetupData.authModeSetupData.authKeyLenInBytes =
sizeof(sampleAuthKey);

/* ICV follows immediately after the region to hash */
sessionSetupData.digestIsAppended = CPA_TRUE;
/* Verify the ICV in the inbound direction */
sessionSetupData.verifyDigest = CPA_TRUE;

```

When calling the function to perform the chained hash and cipher operation, the fields of the operational data structure are listed below.

Listing 23. Operational Data IPsec Inbound

```

/** Populate the structure containing the operational data that is
 * needed to run the algorithm in inbound direction */
pOpData->sessionCtx = sessionCtx;
pOpData->packetType = CPA_CY_SYM_PACKET_TYPE_FULL;
pOpData->pIv = pIvBuffer;
pOpData->ivLenInBytes = sizeof(sampleCipherIv);
pOpData->cryptoStartSrcOffsetInBytes =
    sizeof(sampleEspHdrData) + sizeof(sampleCipherIv);
pOpData->messageLenToCipherInBytes =
    bufferSize -
    (sizeof(sampleEspHdrData) + sizeof(sampleCipherIv) + ICV_LENGTH);
pOpData->hashStartSrcOffsetInBytes = 0;
pOpData->messageLenToHashInBytes = bufferSize - ICV_LENGTH;

```

In the example above, `bufferSize` is the size of the data input (header, iv, ciphertext, and ICV).

3.2.7 Chained Cipher and Hash – SSL Like Use Case

This example demonstrates the usage of the symmetric API for SSL-like use cases, as described in [Figure 7](#) and [Figure 8](#). For the outbound direction, this example employs the symmetric API to perform a "chained" hash and cipher operation. It performs a SHA1 HMAC2 operation on a sequence number, part of the header, and the plaintext.

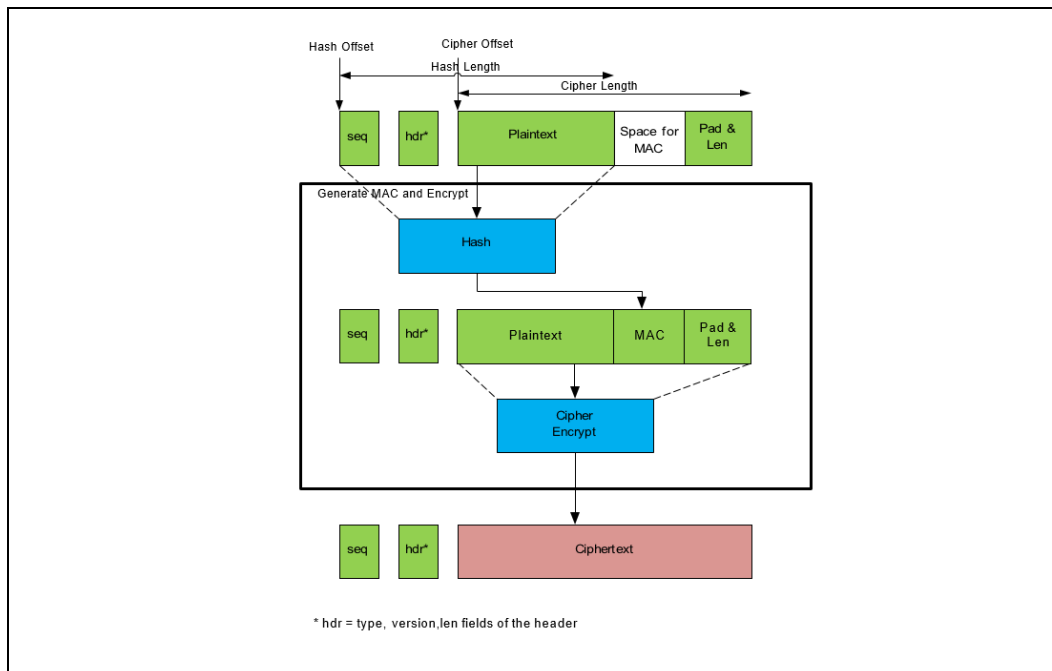
The resultant MAC is placed immediately after the plaintext. Then it encrypts the plaintext, MAC, and padding using the AES algorithm in CBC mode.²

For the inbound direction, this example again employs the use of the symmetric API to perform a "chained" cipher and hash operation. It decrypts the ciphertext using the AES algorithm in CBC mode. Then it performs a SHA1 HMAC operation on the resultant plaintext, sequence number, and part of the Header and compares the Result with the input MAC.

Note: For the inbound direction to use the "chained" API, the length of the plaintext needs to be known before the ciphertext is decrypted to set `messageLenToHashInBytes` and the length field in the Header correctly.

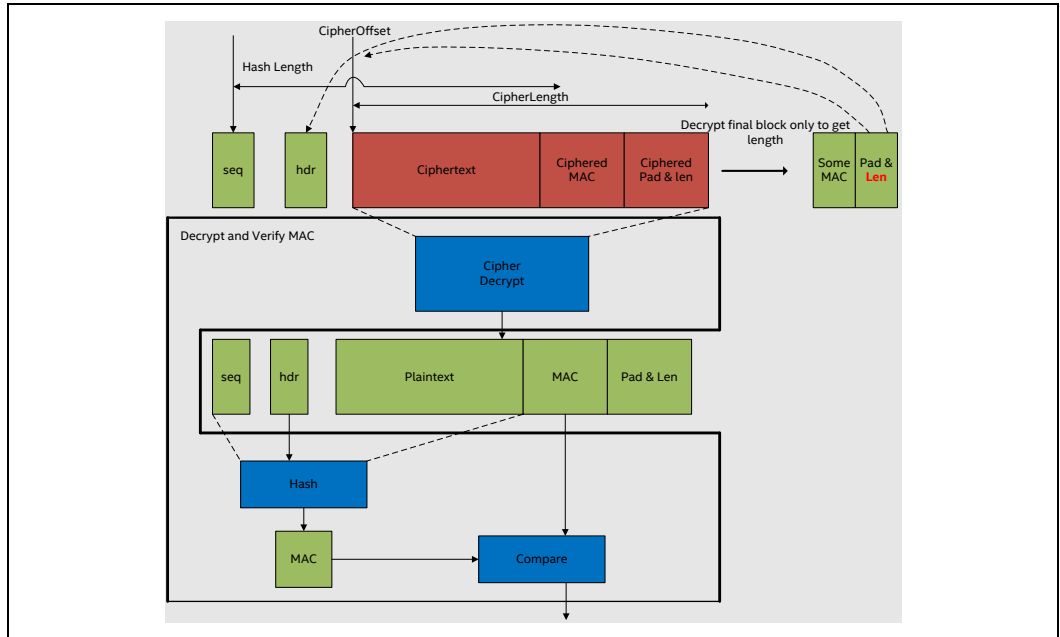
For stream ciphers (e.g., ARC4), there is no padding added in the outbound direction, so the length of the plaintext is simply the length of the ciphertext minus the length of the MAC. However, for block ciphers in CBC mode (as used in this example), the padlen is required to calculate the plaintext length. The final block of the ciphertext needs to be decrypted to discover the padlen. In this example, before calling the "chained" API, the final block of the ciphertext is decrypted to discover the padlen.

Figure 7. SSL Outbound



² Not all SSL use cases use HMAC. For example, The Secure Sockets Layer (SSL) Protocol Version 3.0 (SSL)v3 (RFC 6106) does not use HMAC (in this case the nested hash functionality on the API can be used). However, The Transport Layer Security (TLS) Protocol (TLS) V1.2 (Table 1), for example, does use the HMAC algorithm.

Figure 8. Secure Sockets Layer Inbound



If using a block cipher in CBC mode, then the last ciphertext block is used as the IV for subsequent packets (or records) in Secure Sockets Layer (SSL) and TLSv1.0, whereas in TLSv1.1 and 1.2 an explicit IV is used. However, if using a stream cipher that does not use a synchronization vector (such as ARC4), the stream cipher state from the end of one packet is used to process the subsequent packets. If using the QA API in this case, then partial mode should be used to ensure the stream cipher state is maintained across multiple calls to the API.

Again, these examples are very similar to previous examples, so only the differences are highlighted:

When creating and initializing a session in the outbound direction, the session setup data structure is shown below.

Listing 24. Session Data SSL Outbound

```
sessionSetupData.symOperation = CPA_CY_SYM_OP_ALGORITHM_CHAINING;
sessionSetupData.algChainOrder =
CPA_CY_SYM_ALG_CHAIN_ORDER_HASH_THEN_CIPHER;

sessionSetupData.cipherSetupData.cipherAlgorithm =
CPA_CY_SYM_CIPHER_AES_CBC;
sessionSetupData.cipherSetupData.pCipherKey = sampleCipherKey;
sessionSetupData.cipherSetupData.cipherKeyLenInBytes =
sizeof(sampleCipherKey);
sessionSetupData.cipherSetupData.cipherDirection =
CPA_CY_SYM_CIPHER_DIRECTION_ENCRYPT;

sessionSetupData.hashSetupData.hashAlgorithm = CPA_CY_SYM_HASH_SHA1;
```

```

sessionSetupData.hashSetupData.hashMode = CPA_CY_SYM_HASH_MODE_AUTH;
sessionSetupData.hashSetupData.digestResultLenInBytes = MAC_LENGTH;
sessionSetupData.hashSetupData.authModeSetupData.authKey = sampleAuthKey;
sessionSetupData.hashSetupData.authModeSetupData.authKeyLenInBytes =
sizeof(sampleAuthKey);

/* MAC follows immediately after the region to hash */
sessionSetupData.digestIsAppended = CPA_TRUE;
/* Generate the MAC in outbound direction */
sessionSetupData.verifyDigest = CPA_FALSE;

```

A buffer large enough to hold the plaintext, MAC and padding is required. The size of this buffer will be.

Listing 25. Buffer Size SSL Outbound

```

bufferSize = sizeof(samplePayload) + MAC_LENGTH;

/* bufferSize needs to be rounded up to a multiple of the AES block size
*/
padLen = 16 - bufferSize % 16;
bufferSize += padLen;
/* padLen excludes pad_length field */
padLen--;

```

This buffer is filled with plaintext and padding leaving room for the “chained” API operation to add the MAC.

Listing 26. Buffer Setup SSL Outbound

```

memcpy(pSrcBuffer, samplePayload, sizeof(samplePayload));
/* Leave space for MAC but insert padding data */
for (i = 0; i <= padLen; i++) {
    pSrcBuffer[(sizeof(samplePayload) + MAC_LENGTH + i)] = padLen;
}

```

The session sequence number, the header and the buffer with the plaintext are described using a [CpaBufferList](#):

Listing 27. BufferList Setup SSL Outbound

```

pBufferList->pBuffers = pFlatBuffer;
pBufferList->numBuffers = numBuffers;
pBufferList->pPrivateMetaData = pBufferMeta;

/* Seq number */
pFlatBuffer->dataLenInBytes = SSL_CombinedHeadSize;
pFlatBuffer->pData = pCombinedHeadBuffer;
pFlatBuffer++;
memcpy((char *)pCombinedHeadBuffer + SESSION_SEQ_START,
&sessSeqNum,

```

```

sizeof(sessSeqNum));
memcpy((char *)pCombinedHeadBuffer + HDR_START,
sampleHdrData,
sizeof(sampleHdrData));

/* Data */
pFlatBuffer->dataLenInBytes = bufferSize;
pFlatBuffer->pData = pSrcBuffer;

```

When calling the function to perform the chained hash and cipher operation, the fields of the operational data structure are shown below.

Listing 28. Operational Data SSL Outbound

```

pOpData->sessionCtx = sessionCtx;
pOpData->packetType = CPA_CY_SYM_PACKET_TYPE_FULL;
pOpData->pIv = pIvBuffer;
pOpData->ivLenInBytes = sizeof(sampleCipherIv);
pOpData->cryptoStartSrcOffsetInBytes = SSL_CombinedHeadSize;
pOpData->messageLenToCipherInBytes = bufferSize;
pOpData->hashStartSrcOffsetInBytes = SESSION_SEQ_START;
pOpData->messageLenToHashInBytes = sizeof(sessSeqNum) +
sizeof(sampleHdrData) +
bufferSize - MAC_LENGTH - padLen;

```

When creating and initializing a session in the inbound direction, the session setup data structure is shown below.

Listing 29. Session Data SSL Inbound

```

sessionSetupData.symOperation = CPA_CY_SYM_OP_ALGORITHM_CHAINING;
sessionSetupData.algChainOrder =
CPA_CY_SYM_ALG_CHAIN_ORDER_HASH_THEN_CIPHER;

sessionSetupData.cipherSetupData.cipherAlgorithm =
CPA_CY_SYM_CIPHER_AES_CBC;
sessionSetupData.cipherSetupData.pCipherKey = sampleCipherKey;
sessionSetupData.cipherSetupData.cipherKeyLenInBytes =
sizeof(sampleCipherKey);
sessionSetupData.cipherSetupData.cipherDirection =
CPA_CY_SYM_CIPHER_DIRECTION_ENCRYPT;

sessionSetupData.hashSetupData.hashAlgorithm = CPA_CY_SYM_HASH_SHA1;
sessionSetupData.hashSetupData.hashMode = CPA_CY_SYM_HASH_MODE_AUTH;
sessionSetupData.hashSetupData.digestResultLenInBytes = MAC_LENGTH;
sessionSetupData.hashSetupData.authModeSetupData.authKey = sampleAuthKey;
sessionSetupData.hashSetupData.authModeSetupData.authKeyLenInBytes =
sizeof(sampleAuthKey);

/* MAC follows immediately after the region to hash */

```



```

sessionSetupData.digestIsAppended = CPA_TRUE;
/* Generate the MAC in outbound direction */
sessionSetupData.verifyDigest = CPA_FALSE;

```

In this case the length of the ciphertext is `bufferSize` to calculate the `padLen` the final block is decrypted.

Listing 30. Calculating padLen SSL Inbound

```

Cpa8U resBuff[16];

/* For decrypt direction need to decrypt the final block
 * to determine the messageLenToHashInBytes */
status = sampleCodeAesCbcDecrypt(sampleCipherKey,
sizeof(sampleCipherKey),
(pSrcBuffer + (bufferSize - 32)), /* IV */
(pSrcBuffer + (bufferSize - 16)), /* src */
resBuff);                          /* dest */
/* padLen is the last byte decrypted incremented by one to
 * included the padLen block itself
 */
padLen = resBuff[15] + 1;

```

When calling the function to perform the chained cipher and hash operation, the fields of the operational data structure are.

Listing 31. Operational Data SSL Inbound

```

pOpData->sessionCtx = sessionCtx;
pOpData->packetType = CPA_CY_SYM_PACKET_TYPE_FULL;
pOpData->pIv = pIvBuffer;
pOpData->ivLenInBytes = sizeof(sampleCipherIv);
pOpData->cryptoStartSrcOffsetInBytes = SSL_CombinedHeadSize;
pOpData->messageLenToCipherInBytes = bufferSize;
pOpData->hashStartSrcOffsetInBytes = SESSION_SEQ_START;
pOpData->messageLenToHashInBytes = sizeof(sessSeqNum) +
sizeof(sampleHdrData) +
bufferSize - MAC_LENGTH - padLen;

```

3.2.8 Chained Cipher and Hash – CCM Use Case

This example demonstrates the usage of the symmetric API to perform a CCM operation as described in NIST publication SP800-38C (Recommendation for Block Cipher Modes of Operation: the CCM Mode for Authentication and Confidentiality, refer to [Table 1](#)).

This sample is located in `/sym/ccm_sample`

The example is very similar to the cipher example, so only the differences are highlighted:

For the generation-encryption process the session setup data is shown below:

Listing 32. Session Data CCM Generate-Encrypt

```
sessionSetupData.symOperation = CPA_CY_SYM_OP_ALGORITHM_CHAINING;
sessionSetupData.algChainOrder =
CPA_CY_SYM_ALG_CHAIN_ORDER_HASH_THEN_CIPHER;

sessionSetupData.cipherSetupData.cipherAlgorithm =
CPA_CY_SYM_CIPHER_AES_CCM;
sessionSetupData.cipherSetupData.pCipherKey = sampleKey;
sessionSetupData.cipherSetupData.cipherKeyLenInBytes = sizeof(sampleKey);
sessionSetupData.cipherSetupData.cipherDirection =
    CPA_CY_SYM_CIPHER_DIRECTION_ENCRYPT;

sessionSetupData.hashSetupData.hashAlgorithm = CPA_CY_SYM_HASH_AES_CCM;
sessionSetupData.hashSetupData.hashMode = CPA_CY_SYM_HASH_MODE_AUTH;
sessionSetupData.hashSetupData.digestResultLenInBytes = DIGEST_LENGTH;

/* Notice for CCM authKey and authKeyLen are not required this
 * information is provided by the cipherKey in cipherSetupData
 */
sessionSetupData.hashSetupData.authModeSetupData.aadLenInBytes =
    sizeof(sampleAssocData);
/* For CCM digestAppended and digestVerify are not required. In
 * the encrypt direction digestAppended is CPA_TRUE and
 * digestVerify is CPA_FALSE
 */
```

For the decryption-verification process the session setup data is:

Listing 33. Session Data CCM Decrypt-Verify

```
sessionSetupData.symOperation = CPA_CY_SYM_OP_ALGORITHM_CHAINING;
sessionSetupData.algChainOrder =
CPA_CY_SYM_ALG_CHAIN_ORDER_CIPHER_THEN_HASH;

sessionSetupData.cipherSetupData.cipherAlgorithm =
CPA_CY_SYM_CIPHER_AES_CCM;
sessionSetupData.cipherSetupData.pCipherKey = sampleKey;
sessionSetupData.cipherSetupData.cipherKeyLenInBytes = sizeof(sampleKey);
sessionSetupData.cipherSetupData.cipherDirection =
    CPA_CY_SYM_CIPHER_DIRECTION_DECRYPT;

sessionSetupData.hashSetupData.hashAlgorithm = CPA_CY_SYM_HASH_AES_CCM;
sessionSetupData.hashSetupData.hashMode = CPA_CY_SYM_HASH_MODE_AUTH;
sessionSetupData.hashSetupData.digestResultLenInBytes = DIGEST_LENGTH;
sessionSetupData.hashSetupData.authModeSetupData.aadLenInBytes =
    sizeof(sampleAssocData);
```

The IV and AAD buffers are allocated as shown below:

Listing 34. CCM Allocate IV and AAD Buffers

```

/* Allocate memory to store IV. For CCM this is the counter block
 * ctr0 (size equal to AES block size). The implementation will
 * construct the ctr0 block given the nonce. Space for ctr0 must be
 * allocated here */
        status = PHYS_CONTIG_ALLOC(&pIvBuffer, AES_BLOCK_SIZE);
    }
if (CPA_STATUS_SUCCESS == status) {
    /* Allocate memory for AAD. For CCM this memory will hold the 16 byte
     * B0 block, the 2 bytes encoded length of associated data, the
     * associated data itself and any padding to ensure total size is
     * a multiple of the AES block size
     */
    aadBuffSize = B0_BLOCK_SIZE + ALLEN_ENCODING_SIZE +
sizeof(sampleAssocData);
    if (aadBuffSize % AES_BLOCK_SIZE) {
        aadBuffSize += AES_BLOCK_SIZE - (aadBuffSize % AES_BLOCK_SIZE);
    }
    status = PHYS_CONTIG_ALLOC(&pAadBuffer, aadBuffSize);
}
}

```

The operational data needed to perform the generate-encrypt or decrypt-verify operation is shown below:

Listing 35. CCM Operational Data

```

pOpData->sessionCtx = sessionCtx;
pOpData->packetType = CPA_CY_SYM_PACKET_TYPE_FULL;
pOpData->pIv = pIvBuffer;
/* Even though the iv buffer is 16 bytes the ivLenInBytes
 * is set to the length of the nonce. For CCM valid lengths
 * are in the range 7-13
 */
pOpData->ivLenInBytes = sizeof(sampleNonce);
pOpData->cryptoStartSrcOffsetInBytes = 0;
pOpData->messageLenToCipherInBytes = sizeof(samplePayload);
/* Notice for CCM hash offset and length are not required */
pOpData->pAdditionalAuthData = pAadBuffer;

/* Populate pIv and pAdditionalAuthData buffers with nonce and assoc data
 */
CPA_CY_SYM_CCM_SET_NONCE(pOpData, sampleNonce, sizeof(sampleNonce));
CPA_CY_SYM_CCM_SET_AAD(pOpData, sampleAssocData,
sizeof(sampleAssocData));

```

3.2.9 Chained Cipher and Hash – GCM Use Case

This example demonstrates the usage of the symmetric API to perform a GCM operation as described in *NIST publication SP800-38D (Recommendation for Block Cipher Modes of Operation: Galois/Counter Mode (GCM) and GMAC*, refer to [Table 1](#)).

These samples are located in `/sym/gcm_sample`

An example of the session setup data and operational data for GCM authenticated encryption and decryption is shown below.

```
sessionSetupData.symOperation = CPA_CY_SYM_OP_ALGORITHM_CHAINING;
sessionSetupData.algChainOrder =
CPA_CY_SYM_ALG_CHAIN_ORDER_CIPHER_THEN_HASH;

sessionSetupData.cipherSetupData.cipherAlgorithm =
CPA_CY_SYM_CIPHER_AES_GCM;
sessionSetupData.cipherSetupData.pCipherKey = sampleKey;
sessionSetupData.cipherSetupData.cipherKeyLenInBytes = sizeof(sampleKey);
sessionSetupData.cipherSetupData.cipherDirection =
    CPA_CY_SYM_CIPHER_DIRECTION_ENCRYPT;

sessionSetupData.hashSetupData.hashAlgorithm = CPA_CY_SYM_HASH_AES_GCM;
sessionSetupData.hashSetupData.hashMode = CPA_CY_SYM_HASH_MODE_AUTH;
sessionSetupData.hashSetupData.digestResultLenInBytes = TAG_LENGTH;
/* For GCM authKey and authKeyLen are not required this information
 * is provided by the cipherKey in cipherSetupData */
sessionSetupData.hashSetupData.authModeSetupData.aadLenInBytes =
    sizeof(sampleAddAuthData);
/* Tag follows immediately after the region to hash */
sessionSetupData.digestIsAppended = CPA_TRUE;
/* digestVerify is not required to be set. For GCM authenticated
 * encryption this value is understood to be CPA_FALSE */
```

For authenticated encryption the session setup data is:

Listing 36. Session Data GCM Auth-Encrypt

```
sessionSetupData.symOperation = CPA_CY_SYM_OP_ALGORITHM_CHAINING;
sessionSetupData.algChainOrder =
CPA_CY_SYM_ALG_CHAIN_ORDER_HASH_THEN_CIPHER;

sessionSetupData.cipherSetupData.cipherAlgorithm =
CPA_CY_SYM_CIPHER_AES_GCM;
sessionSetupData.cipherSetupData.pCipherKey = sampleKey;
sessionSetupData.cipherSetupData.cipherKeyLenInBytes = sizeof(sampleKey);
sessionSetupData.cipherSetupData.cipherDirection =
    CPA_CY_SYM_CIPHER_DIRECTION_DECRYPT;

sessionSetupData.hashSetupData.hashAlgorithm = CPA_CY_SYM_HASH_AES_GCM;
```

```

sessionSetupData.hashSetupData.hashMode = CPA_CY_SYM_HASH_MODE_AUTH;
sessionSetupData.hashSetupData.digestResultLenInBytes = TAG_LENGTH;

/* For GCM authKey and authKeyLen are not required this information
 * is provided by the cipherKey in cipherSetupData */
sessionSetupData.hashSetupData.authModeSetupData.aadLenInBytes =
    sizeof(sampleAddAuthData);
/* Tag follows immediately after the region to hash */
sessionSetupData.digestIsAppended = CPA_TRUE;
/* digestVerify is not required to be set. For GCM authenticated
 * decryption this value is understood to be CPA_TRUE */

```

For authenticated decryption the session setup data is:

Listing 37. Session Data GCM Auth-Decrypt

```

sessionSetupData.symOperation = CPA_CY_SYM_OP_ALGORITHM_CHAINING;
sessionSetupData.algChainOrder =
CPA_CY_SYM_ALG_CHAIN_ORDER_HASH_THEN_CIPHER;

sessionSetupData.cipherSetupData.cipherAlgorithm =
CPA_CY_SYM_CIPHER_AES_GCM;
sessionSetupData.cipherSetupData.pCipherKey = sampleKey;
sessionSetupData.cipherSetupData.cipherKeyLenInBytes = sizeof(sampleKey);
sessionSetupData.cipherSetupData.cipherDirection =
    CPA_CY_SYM_CIPHER_DIRECTION_DECRYPT;

sessionSetupData.hashSetupData.hashAlgorithm = CPA_CY_SYM_HASH_AES_GCM;
sessionSetupData.hashSetupData.hashMode = CPA_CY_SYM_HASH_MODE_AUTH;
sessionSetupData.hashSetupData.digestResultLenInBytes = TAG_LENGTH;

/* For GCM authKey and authKeyLen are not required this information
 * is provided by the cipherKey in cipherSetupData */
sessionSetupData.hashSetupData.authModeSetupData.aadLenInBytes =
    sizeof(sampleAddAuthData);
/* Tag follows immediately after the region to hash */
sessionSetupData.digestIsAppended = CPA_TRUE;
/* digestVerify is not required to be set. For GCM authenticated
 * decryption this value is understood to be CPA_TRUE */

```

The IV and AAD buffers are allocated as shown below:

Listing 38. GCM Allocate IV and AAD Buffers

```

/* Allocate memory to store IV. For GCM this is the block J0
 * (size equal to AES block size). If iv is 12 bytes the
 * implementation will construct the J0 block given the iv.
 * If iv is not 12 bytes then the user must construct the J0
 * block and give this as the iv. In both cases space for J0

```

```

    * must be allocated. */
    status = PHYS_CONTIG_ALLOC(&pIvBuffer, AES_BLOCK_SIZE);
}
if (CPA_STATUS_SUCCESS == status) {
    /* Allocate memory for AAD. For GCM this memory will hold the
    * additional authentication data and any padding to ensure total
    * size is a multiple of the AES block size */
    aadBuffSize = sizeof(sampleAddAuthData);
    if (aadBuffSize % AES_BLOCK_SIZE) {
        aadBuffSize += AES_BLOCK_SIZE - (aadBuffSize % AES_BLOCK_SIZE);
    }
    status = PHYS_CONTIG_ALLOC(&pAadBuffer, aadBuffSize);
}
}

```

The operational data needed to perform the encrypt or decrypt operation is:

Listing 39. GCM Operational Data

```

pOpData->sessionCtx = sessionCtx;
pOpData->packetType = CPA_CY_SYM_PACKET_TYPE_FULL;
pOpData->pIv = pIvBuffer;
/* In this example iv is 12 bytes. The implementation
 * will use the iv to generation the J0 block
 */
memcpy(pIvBuffer, sampleIv, sizeof(sampleIv));
pOpData->ivLenInBytes = sizeof(sampleIv);
pOpData->cryptoStartSrcOffsetInBytes = 0;
pOpData->messageLenToCipherInBytes = sizeof(samplePayload);
/* For GCM hash offset and length are not required */
pOpData->pAdditionalAuthData = pAadBuffer;

```

GMAC is supported using the same API and similar data structures as the general GCM case shown above. However, for GMAC, the `messageLenToCipherInBytes` will be set to **0**.

3.2.10 Chained Cipher and Hash Using the Symmetric Data Plane API

This example demonstrates the usage of the data plane symmetric API to perform a “chained” cipher and hash operation. It encrypts some sample text using the AES-256 algorithm in CBC mode, and then performs an SHA-256 HMAC operation on the ciphertext, writing the MAC to the buffer immediately after the ciphertext.

This example has been simplified to demonstrate the basics of how to use the API and build the structures required. This example does not demonstrate the optimal way to use the API to get the maximum performance for a particular implementation. Refer to Implementation Specific Documentation in [Table 1](#) (for example, the Intel® Communications Chipset 8900 to 8920 Series Software Programmer’s Guide) and performance sample code for a guide on how to use the API for best performance.

These samples are located in `/sym/symdp_sample`

Note: Use of the data plane symmetric API follows some of the same basic steps as the traditional symmetric API:

- Discover and start up the cryptographic service instance.
- Register a callback function for the instance.
- Create and initialize a session.
- Enqueue the symmetric operation on the instance.
- Submit the symmetric operation for processing.
- Poll the instance for a response.
- Tear down the session.
- Stop the Cryptographic service instance.

The following are the steps in more detail:

- Cryptographic service instances are discovered and started in the same way and using the same API as the traditional symmetric use cases described in [Listing 4](#), [Listing 5](#), and, [Listing 6](#).
- The next step is to register a callback function for the cryptographic instance:

The function is called back in the context of the polling function when an asynchronous operation has completed. This function can perform whatever processing is appropriate to the Application.

Callback differs from the traditional symmetric API, where the callback function is registered for the session.

Listing 40. Register Callback Function

```
status = cpaCySymDpRegCbFunc(cyInstHandle, symDpCallback);
```

Create and initialize a session:

Listing 41. Create and Initialize Data Plane Session

```
sessionSetupData.sessionPriority = CPA_CY_PRIORITY_HIGH;
sessionSetupData.symOperation = CPA_CY_SYM_OP_ALGORITHM_CHAINING;
sessionSetupData.algChainOrder =
CPA_CY_SYM_ALG_CHAIN_ORDER_CIPHER_THEN_HASH;

sessionSetupData.cipherSetupData.cipherAlgorithm =
CPA_CY_SYM_CIPHER_AES_CBC;
sessionSetupData.cipherSetupData.pCipherKey = sampleCipherKey;
sessionSetupData.cipherSetupData.cipherKeyLenInBytes =
sizeof(sampleCipherKey);
sessionSetupData.cipherSetupData.cipherDirection =
    CPA_CY_SYM_CIPHER_DIRECTION_ENCRYPT;

sessionSetupData.hashSetupData.hashAlgorithm = CPA_CY_SYM_HASH_SHA256;
```

```

sessionSetupData.hashSetupData.hashMode = CPA_CY_SYM_HASH_MODE_AUTH;
sessionSetupData.hashSetupData.digestResultLenInBytes = DIGEST_LENGTH;
sessionSetupData.hashSetupData.authModeSetupData.authKey =
sampleCipherKey;
sessionSetupData.hashSetupData.authModeSetupData.authKeyLenInBytes =
    sizeof(sampleCipherKey);

/* Even though MAC follows immediately after the region to hash
 * digestIsAppended is set to false in this case to workaround
 * errata number IXA00378322 */
sessionSetupData.digestIsAppended = CPA_FALSE;
sessionSetupData.verifyDigest = CPA_FALSE;

/* Determine size of session context to allocate */
PRINT_DBG("cpaCySymDpSessionCtxGetSize\n");
status = cpaCySymDpSessionCtxGetSize(cyInstHandle,
                                     &sessionSetupData,
                                     &sessionCtxSize);
}

if (CPA_STATUS_SUCCESS == status) {
    /* Allocate session context */
    status = PHYS_CONTIG_ALLOC(&sessionCtx, sessionCtxSize);
}

if (CPA_STATUS_SUCCESS == status) {
    /* Initialize the session */
    PRINT_DBG("cpaCySymDpInitSession\n");
    status = cpaCySymDpInitSession(cyInstHandle, &sessionSetupData,
    sessionCtx);
}

#ifdef LAC_HW_PRECOMPUTES
if (CPA_STATUS_SUCCESS == status) {
    /* Poll for hw pre-compute responses. */
    do {
        status = icp_sal_CyPollDpInstance(cyInstHandle, 0);
    } while (CPA_STATUS_SUCCESS != status);
}
#endif

```

In this example, data is stored in flat buffers (as opposed to scatter gather lists). The operational data in this case is shown below.

Listing 42. Data Plane Operational Data

```

sessionSetupData.sessionPriority = CPA_CY_PRIORITY_HIGH;
sessionSetupData.symOperation = CPA_CY_SYM_OP_ALGORITHM_CHAINING;

```



```

sessionSetupData.algChainOrder =
CPA_CY_SYM_ALG_CHAIN_ORDER_CIPHER_THEN_HASH;

sessionSetupData.cipherSetupData.cipherAlgorithm =
CPA_CY_SYM_CIPHER_AES_CBC;
sessionSetupData.cipherSetupData.pCipherKey = sampleCipherKey;
sessionSetupData.cipherSetupData.cipherKeyLenInBytes =
sizeof(sampleCipherKey);
sessionSetupData.cipherSetupData.cipherDirection =
    CPA_CY_SYM_CIPHER_DIRECTION_ENCRYPT;

sessionSetupData.hashSetupData.hashAlgorithm = CPA_CY_SYM_HASH_SHA256;
sessionSetupData.hashSetupData.hashMode = CPA_CY_SYM_HASH_MODE_AUTH;
sessionSetupData.hashSetupData.digestResultLenInBytes = DIGEST_LENGTH;
sessionSetupData.hashSetupData.authModeSetupData.authKey =
sampleCipherKey;
sessionSetupData.hashSetupData.authModeSetupData.authKeyLenInBytes =
    sizeof(sampleCipherKey);

/* Even though MAC follows immediately after the region to hash
 * digestIsAppended is set to false in this case to workaround
 * errata number IXA00378322 */
sessionSetupData.digestIsAppended = CPA_FALSE;
sessionSetupData.verifyDigest = CPA_FALSE;

/* Determine size of session context to allocate */
PRINT_DBG("cpaCySymDpSessionCtxGetSize\n");
status = cpaCySymDpSessionCtxGetSize(cyInstHandle,
                                     &sessionSetupData,
                                     &sessionCtxSize);
}

if (CPA_STATUS_SUCCESS == status) {
    /* Allocate session context */
    status = PHYS_CONTIG_ALLOC(&sessionCtx, sessionCtxSize);
}

if (CPA_STATUS_SUCCESS == status) {
    /* Initialize the session */
    PRINT_DBG("cpaCySymDpInitSession\n");
    status = cpaCySymDpInitSession(cyInstHandle, &sessionSetupData,
    sessionCtx);
}

#ifdef LAC_HW_PRECOMPUTES
if (CPA_STATUS_SUCCESS == status) {
    /* Poll for hw pre-compute responses. */

```

```
do {
    status = icp_sal_CyPollDpInstance(cyInstHandle, 0);
} while (CPA_STATUS_SUCCESS != status);
}
#endif
```

This request is then enqueued on the instance.

Listing 43. Data Plane Enqueue

```
status = cpaCySymDpEnqueueOp(pOpData, CPA_FALSE);
```

Other requests can now be enqueued before submitting all the requests to be processed. Enqueued requests allow the cost of submitting a request (which can be expensive, in terms of cycles, for some hardware-based implementations) to be amortized over all enqueued requests on the instance. Once sufficient requests have been enqueued they are all submitted for processing.

Listing 44. Data Plane Perform

```
status = cpaCySymDpPerformOpNow(cyInstHandle);
```

- An alternative to calling the `cpaCySymDpPerformOpNow` function is to set `performOpNow` to `CPA_TRUE` when calling the enqueue functions (`cpaCySymDpEnqueueOp` or `cpaCySymDpEnqueueOpBatch`). This is illustrated in Sections [4.4, Sample – Data Compression Data Plane API](#) example.
- After submitting several requests and possibly doing other work (e.g., enqueueing and submitting more requests), the Application can poll for responses that invoke the callback function registered with the instance. Refer to [Table 1](#) Implementation Specific Documentation for information on the implementations polling functions.
- Once all requests associated with a session have been completed, the session can be removed.

Listing 45. Wait for Outstanding Requests

```
symSessionWaitForInflightReq(sessionCtx)
```

Since cryptographic API v2.2 before removing the symmetric session context, it is recommended to wait for the completion of any outstanding request using `cpaCySymSessionInUse`.

It is executed in the `symSessionWaitForInflightReq` call, which polls for the in-flight requests.

Listing 46. Data Plane Remove Session

```
sessionStatus = cpaCySymDpRemoveSession(cyInstHandle, sessionCtx);
```

Since Cryptographic API v2.2, two new functions have been implemented:

```
cpaCySymUpdateSession and cpaCySymSessionInUse
```

- The function `cpaCySymUpdateSession` can be used to update certain parameters of a session like a cipher key, the cipher direction, and the authentication key. `cpaCySymSessionInUse` indicates whether there are outstanding requests on a given session.

As a result of the implementation of this feature, the behavior of `cpaCySymRemoveSession` has been changed. The `cpaCySymRemoveSession` fails if there is an outstanding request for the session that the user is trying to remove.

- As a result, it is recommended to wait for the completion of any outstanding request, using `cpaCySymSessionInUse`, before removing a session.

3.2.11 TLS Key and MGF Mask Generation

Refer [Table 1](#) to the API manuals for full details of Key and Mask Generation operations.

1. Define a Flat Buffer callback function as per the API prototype, Refer [Table 1](#) to the API manuals.
If synchronous operation is preferred, instead simply pass NULL to the API for the callback parameter.
2. Allocate memory for the operation.
3. Populate data for the appropriate operation data structure, Refer [Table 1](#) to the API manuals.
 - a. Fill in the Flat Buffers, a pointer to data, and length.
 - b. Fill in the options for the operation required.
4. Call the appropriate key or Mask Generation API.
5. Complete the operation.

The API for TLS key operations is based on the Transport Layer Security (TLS) Protocol Version 1.1 standard, RFC 4346, Backward compatibility is supported with the legacy Transport Layer Security (TLS) Protocol Version 1.0 standard, RFC 2246, refer to [Table 1](#). The user-defined label should be used for backward compatibility with the client write key, server write key, and iv block. Refer to [Table 1](#) Intel® QAT Cryptographic API Reference Manual for details of populating `CpaCyKeyGenTlsOpData`, the operation data structure.

The following sections describe examples of the parameter mapping to the Cryptographic API.

3.2.11.1 Setting CpaCyKeyGenTlsOpData Structure Fields

The Transport Layer Security (TLS) Protocol Version 1.1 standard, RFC 4346, refer to [Table 1](#), Section 6.3 `key_block` is described as:

```
key_block = PRF(SecurityParameters.master_secret,
               "key expansion",
               SecurityParameters.server_random +
               SecurityParameters.client_random);
```

This maps to the Cryptographic API's `CpaCyKeyGenTlsOpData` as follows:

TLS Key-Material Derivation:

```
tlsOp = CPA_CY_KEY_TLS_OP_KEY_MATERIAL_DERIVE
```

```
secret = master secret key

seed = server_random + client_random

userLabel = NULL
```

Setting `CpaCyKeyGenTlsOpData` Structure Fields for Backward Compatibility.

1. In the *Transport Layer Security (TLS) Protocol Version 1.0 standard*, RFC 2246, refer to [Table 1](#), Section 6.3 `final_client_write_key` is described as:

```
final_client_write_key = PRF(client_write_key,
                             "client write key",
                             client_random + server_random) [0..15]
```

This maps to the Cryptographic API's `CpaCyKeyGenTlsOpData` as follows:

```
TLS User Defined Derivation:
tlsOp = CPA_CY_KEY_TLS_OP_USER_DEFINED
secret = client_write_key
seed = client_random + server_random
userLabel = "client write key"
```

2. In the *Transport Layer Security (TLS) Protocol v1.0 standard*, RFC 2246, refer to [Table 1](#), Section 6.3 `final_server_write_key` is described as:

```
final_server_write_key = PRF(server_write_key,
                             "server write key",
                             client_random +
                             server_random) [0..15]
```

This maps to the Cryptographic API's `CpaCyKeyGenTlsOpData` as follows:

```
TLS User Defined Derivation:
tlsOp = CPA_CY_KEY_TLS_OP_USER_DEFINED
secret = server_write_key
seed = client_random + server_random
userLabel = "server write key"
```

3. In the *Transport Layer Security (TLS) Protocol Version 1.0 standard*, RFC 2246, refer to [Table 1.](#), Section 6.3 `iv_block` is described as:

```
iv_block = PRF("", "IV block", client_random +
               server_random) [0..15]
```

This maps to the Cryptographic API's `CpaCyKeyGenTlsOpData` as follows:

```
TLS User Defined Derivation:
tlsOp = CPA_CY_KEY_TLS_OP_USER_DEFINED
secret = NULL
seed = client_random + server_random
userLabel = "IV block"
```

Memory for the user label must be physically contiguous memory allocated by the user. This memory must be available to the API for the duration of the operation.

3.2.12 Session Update for Chained Cipher and Hash Operation

This example demonstrates the usage of the session update together with data plane symmetric API to perform a “chained” cipher and hash operation. It performs a [KASUMI F9](#) hash operation on the sample text and then encrypts the sample text using the [KASUMI F8](#) algorithm. After the operation is complete, the cipher and authentication keys are updated inside the session and the operation is performed again with different keys.

Note: This example is simplified to demonstrate the basics of how to use the API and how to build the structures required. This example does not demonstrate the optimal way to use the API to get maximum performance for a particular implementation. Refer to [Table 1](#) for implementation specific documentation and performance sample code for a guide on how to use the API for best performance.

These samples are located in `/sym/update_sample`

The following are the details of the steps performed in the sample:

- Cryptographic service instances are discovered and started in the same way and using the same API as the traditional symmetric use cases described in [Listing 4](#), [Listing 5](#) and [Listing 6](#).
- Next register a callback function for the cryptographic instance.

The function is called back in the context of the polling function when an asynchronous operation has completed. This function can perform whatever processing is appropriate to the application. Note this differs from the traditional symmetric API where the callback function is registered for the session.

Listing 47. Register Callback Function

```
status = cpaCySymDpRegCbFunc(cyInstHandle, symDpCallback);
```

3.2.12.1 Create and Initialize a Session:

Listing 48. Create and Initialize Data Plane Session

```
sessionSetupData.sessionPriority = CPA_CY_PRIORITY_HIGH;
sessionSetupData.symOperation = CPA_CY_SYM_OP_ALGORITHM_CHAINING;
sessionSetupData.algChainOrder =
    CPA_CY_SYM_ALG_CHAIN_ORDER_HASH_THEN_CIPHER;
sessionSetupData.cipherSetupData.cipherAlgorithm =
    CPA_CY_SYM_CIPHER_KASUMI_F8;
sessionSetupData.cipherSetupData.pCipherKey = pCipherKey;
sessionSetupData.cipherSetupData.cipherKeyLenInBytes = cipherKeyLen;
sessionSetupData.cipherSetupData.cipherDirection =
    CPA_CY_SYM_CIPHER_DIRECTION_ENCRYPT;
sessionSetupData.hashSetupData.hashAlgorithm =
CPA_CY_SYM_HASH_KASUMI_F9;
    sessionSetupData.hashSetupData.hashMode = CPA_CY_SYM_HASH_MODE_AUTH;
    sessionSetupData.hashSetupData.digestResultLenInBytes =
DIGEST_LENGTH;
```

```

sessionSetupData.hashSetupData.authModeSetupData.authKey = authKey;
sessionSetupData.hashSetupData.authModeSetupData.authKeyLenInBytes =
    authKeyLen;
sessionSetupData.hashSetupData.authModeSetupData.aadLenInBytes =
    sizeof(additionalAuthData);
sessionSetupData.digestIsAppended = CPA_TRUE;
sessionSetupData.verifyDigest = CPA_FALSE;

/* Determine size of session context to allocate */
PRINT_DBG("cpaCySymDpSessionCtxGetSize\n");
status = cpaCySymDpSessionCtxGetSize(
    cyInstHandle, &sessionSetupData, &sessionCtxSize);

if (CPA_STATUS_SUCCESS == status)
{
    /* Allocate session context */
    status = PHYS_CONTIG_ALLOC(sessionCtx, sessionCtxSize);
}

if (CPA_STATUS_SUCCESS == status)
{
    /* Initialize the session */
    PRINT_DBG("cpaCySymDpInitSession\n");
    status = cpaCySymDpInitSession(cyInstHandle, &sessionSetupData,
        *sessionCtx);
}

```

In this example, data is stored in flat buffers (as opposed to scatter gather lists). The operational data in this case is:

Listing 49. Data Plane Operational Data

```

pOpData->thisPhys = sampleVirtToPhys(pOpData);
pOpData->instanceHandle = cyInstHandle;
pOpData->sessionCtx = sessionCtx;
pOpData->pCallbackTag = (void *)0;
pOpData->cryptoStartSrcOffsetInBytes = 0;
pOpData->messageLenToCipherInBytes = srcLen;
pOpData->hashStartSrcOffsetInBytes = 0;
pOpData->messageLenToHashInBytes = srcLen;
pOpData->digestResult = sampleVirtToPhys(pSrcBuffer) + srcLen;
pOpData->iv = sampleVirtToPhys(pIvBuffer);
pOpData->pIv = pIvBuffer;
pOpData->ivLenInBytes = ivLen;
pOpData->additionalAuthData =
sampleVirtToPhys(pAdditionalAuthData);
pOpData->pAdditionalAuthData = pAdditionalAuthData;

```

```
pOpData->srcBuffer = sampleVirtToPhys(pSrcBuffer);
pOpData->srcBufferLen = bufferSize;
pOpData->dstBuffer = sampleVirtToPhys(pDstBuffer);
pOpData->dstBufferLen = bufferSize;
```

This request is then enqueued on the instance.

Listing 50. Data Plane Enqueue

```
status = cpaCySymDpEnqueueOp(pOpData, CPA_FALSE);
```

Other requests can now be enqueued before submitting all the requests to be processed. This allows the cost of submitting a request (which can be expensive, in terms of cycles, for some hardware-based implementations) to be amortized over all enqueued requests on the instance. Once sufficient requests have been enqueued they are all submitted for processing.

Listing 51. Data Plane Perform

```
status = cpaCySymDpPerformOpNow(cyInstHandle);
```

- An alternative to calling the `cpaCySymDpPerformOpNow` function is to set `performOpNow` to `CPA_TRUE` when calling the enqueue functions (`cpaCySymDpEnqueueOp` or `cpaCySymDpEnqueueOpBatch`). This is illustrated in the data compression data plane example.
- After submitting a number of requests and possibly doing other work (e.g. enqueueing and submitting more requests) the application can poll for responses which will invoke the callback function registered with the instance. Refer to [Table 1](#) for implementation specific documentation for information on the implementations polling functions.
- After the operation is complete cipher key and authentication key are updated in the existing session via session update API:

Listing 52. Session Update

```
sessionUpdateData.flags = CPA_CY_SYM_SESUPD_CIPHER_KEY;
sessionUpdateData.flags |= CPA_CY_SYM_SESUPD_AUTH_KEY;
sessionUpdateData.pCipherKey = pCipherKey;
sessionUpdateData.authKey = authKey;

status = cpaCySymUpdateSession(sessionCtx, &sessionUpdateData);
```

- With the keys changed, the chained cipher and hash operation is performed again, just as described above.
- Once all requests associated with a session have been completed, the session can be removed.

3.2.13 HKDF Use Case

This section contains sample code that demonstrates the usage of the symmetric API, specifically using this API to perform hash-based message authentication code key derivation function (HKDF) operations. It performs HKDF Extract and Expand, and Extract and Expand Label operation without and with sublabels (KEY and IV).

The simplified code example below is simplified and demonstrates how to use the API and build the structures required. This example does not demonstrate the optimal way to use the API to get maximum performance for implementation.

This sample is located in the directory:

```
/quickassist/lookaside/access_layer/src/sample_code/functional/sym/hkdf_sample
```

3.2.13.1 Instance Configuration and Memory Allocation

Cryptographic service instances are discovered and started in the same way and using the same API as the traditional symmetric use cases.

1. If the instance is polled, start the polling thread.
Polling is done in an implementation-dependent manner.
2. Allocate memory for HKDF operation data:

Listing 53. HKDF Operation Data – Memory Allocation

```
pOpData = qaeMemAllocNUMA(sizeof(CpaCyKeyGenHKDFOpData),
                           instanceInfo2.nodeAffinity,
                           BYTE_ALIGNMENT_64);
```

This structure must be allocated with USDM to be pinned in physical memory.

3. Allocate memory for HKDF output data. Output data is `CpaFlatBuffer` type:

```
PHYS_CONTIG_ALLOC(&pHkdfData, hkdfDataSize);
```

3.2.13.2 HKDF Extract Expand Operation

To perform an Extract Expand operation, go to the `CpaCyKeyGenHKDFOpData` structure, and set `hkdfKeyOp` to `CPA_CY_HKDF_KEY_EXTRACT_EXPAND`.

Note: Provide the lengths `seedLen`, `secretLen`, and `infoLen` and copy all data into the seed, secret, and info tables.

Listing 54. HKDF Extract Expand Operation

```
pOpData->hkdfKeyOp = CPA_CY_HKDF_KEY_EXTRACT_EXPAND;

pOpData->seedLen = sizeof(ikm);
memcpy(pOpData->seed, ikm, pOpData->seedLen);

pOpData->secretLen = sizeof(slt);
memcpy(pOpData->secret, slt, pOpData->secretLen);

pOpData->infoLen = sizeof(inf);
memcpy(pOpData->info, inf, pOpData->infoLen);
```

3.2.13.3 HKDF Extract Expand Label Operation

To perform an Extract, Expand Label operation:

1. Go to the `CpaCyKeyGenHKDFOpData` structure and set `hkdfKeyOp` to `CPA_CY_HKDF_KEY_EXTRACT_EXPAND_LABEL`.

2. Provide the lengths `seedLen`, `secretLen`, and `infoLen` and copy all data into the `seed`, `secret`, and `info` tables.
3. Set the number of labels in the `numLabels` field.
4. Set `label[0].labelLen`.
5. Copy label data into the `label[0].label` table.
6. Finally, set the `label[0].sublabelFlag` field to `0x00` to disable generating sublabels.

Listing 55. HKDF Extract Expand Label Operation

```
pOpData->hkdfKeyOp = CPA_CY_HKDF_KEY_EXTRACT_EXPAND_LABEL;
pOpData->seedLen = sizeof(seed_label);
memcpy(pOpData->seed, seed_label, sizeof(seed_label));

pOpData->secretLen = sizeof(secret_label);
memcpy(pOpData->secret, secret_label, sizeof(secret_label));

pOpData->numLabels = 1;
memcpy(pOpData->label[0].label, label, sizeof(label));

pOpData->label[0].labelLen = sizeof(label);
pOpData->label[0].sublabelFlag = 0x00;
```

3.2.13.4 HKDF Extract Expand Label and Sublabels operation

To perform an Extract, Expand Label and Sublabels operation:

1. Go to the `CpaCyKeyGenHKDFOpData` structure and set `hkdfKeyOp` to `CPA_CY_HKDF_KEY_EXTRACT_EXPAND_LABEL`.
2. Provide the lengths `seedLen`, `secretLen`, and `infoLen` and copy all data into the `seed`, `secret`, and `info` tables.
3. Set the number of labels in the `numLabels` field.
4. Set `label[0].labelLen`.
5. Copy label data into the `label[0].label` table.
6. Set the `label[0].sublabelFlag` and `label[0].sublabelFlag` field as shown below to generate Key and IV sublabels.

Listing 56. HKDF Extract Expand Label and Sublabels Operation

```
pOpData->hkdfKeyOp = CPA_CY_HKDF_KEY_EXTRACT_EXPAND_LABEL;
pOpData->seedLen = sizeof(seed_label);
memcpy(pOpData->seed, seed_label, sizeof(seed_label));

pOpData->secretLen = sizeof(secret_label);
memcpy(pOpData->secret, secret_label, sizeof(secret_label));

pOpData->numLabels = 1;
memcpy(pOpData->label[0].label, label, sizeof(label));
```

```
pOpData->label[0].labelLen = sizeof(label);
pOpData->label[0].sublabelFlag = CPA_CY_HKDF_SUBLABEL_KEY;
pOpData->label[0].sublabelFlag |= CPA_CY_HKDF_SUBLABEL_IV;
```

3.2.14 Perform HKDF operation

The crypto instance must be specified in the `instanceHandle` to execute the HKDF operation. When the operation is performed asynchronously, the callback function and callback tag should be set in the `pKeyGenCb` and `pCallbackTag` arguments. Operational data is provided in `pKeyGenTlsOpData`, and `CpaCyKeyHKDFCIPHERSuite` must be chosen. The output is passed to the `CpaFlatBuffer`. All generated values are arranged one after the other in a single buffer. Depending on what operations are performed, the buffer length should be adjusted.

Listing 57. HKDF Operation

```
cpaCyKeyGenTls3(cyInstHandle, /* Instance handle */
               hkdfSampleCallback, /* Callback function */
               (void *)&complete, /* Callback tag */
               pOpData, /* HKDF operational data */
               CPA_CY_HKDF_TLS_AES_128_GCM_SHA256, /* HKDF cipher suite
*/
               &hkdfOut); /* Output buffer */
```

3.3 Using the Diffie-Hellman API

This example demonstrates the usage of the Diffie-Hellman API.

These samples are located in `/asym/diffie_hellman_sample`

The following steps are carried out:

- The example uses the API asynchronously; therefore, you must define a Diffie-Hellman callback function per the API prototype.
- Instance, discovery, and start-up are made in a way similar to that defined for the symmetric examples above.
- The function `sampleDhPerformOp` is called, which does the following:
 - Allocate memory for the operation and populate data for the appropriate DH phase 1 operation data structure to generate the public value. The fields to be allocated and populated are the prime P, the base G, and the private value X. Space must also be allocated for the output, which is the public value (PV).

Listing 58. Allocate Memory and Populate Operational Data

```
status = OS_MALLOC(&pCpaDhOpDataP1, sizeof(CpaCyDhPhase1KeyGenOpData));

/*
 * Allocate input buffers for phase 1 and copy data. Input to DH
 * phase 1 includes the prime (primeP), the base g (baseG) and
 * a random private value (privateValueX).
 */
if (CPA_STATUS_SUCCESS == status) {
```

```

memset(pCpaDhOpDataP1, 0, sizeof(CpaCyDhPhase1KeyGenOpData));

pCpaDhOpDataP1->primeP.dataLenInBytes = sizeof(primeP_768);
status =
    PHYS_CONTIG_ALLOC(&pCpaDhOpDataP1->primeP.pData,
sizeof(primeP_768));

if (NULL != pCpaDhOpDataP1->primeP.pData) {
    memcpy(pCpaDhOpDataP1->primeP.pData, primeP_768,
sizeof(primeP_768));
}
}

if (CPA_STATUS_SUCCESS == status) {
    pCpaDhOpDataP1->baseG.dataLenInBytes = sizeof(baseG1);
    status = PHYS_CONTIG_ALLOC(&pCpaDhOpDataP1->baseG.pData,
sizeof(baseG1));

    if (NULL != pCpaDhOpDataP1->baseG.pData) {
        memcpy(pCpaDhOpDataP1->baseG.pData, baseG1, sizeof(baseG1));
    }
}

if (CPA_STATUS_SUCCESS == status) {
    pCpaDhOpDataP1->privateValueX.dataLenInBytes = sizeof(privateValueX);
    status = PHYS_CONTIG_ALLOC(&pCpaDhOpDataP1->privateValueX.pData,
sizeof(privateValueX));

    if (NULL != pCpaDhOpDataP1->privateValueX.pData) {
        memcpy(pCpaDhOpDataP1->privateValueX.pData,
privateValueX,
sizeof(privateValueX));
    }
}
}

```

Invoke the phase 1 operation, which performs the modular exponentiation such that $PV = (\text{baseG} \wedge \text{privateValueX}) \bmod \text{primeP}$.

Note: In the case of phase 1, the operation is invoked synchronously, hence the NULL pointer for the callback function.

Listing 59. Perform Phase 1 Operation

```

status = cpaCyDhKeyGenPhase2Secret(
    cyInstHandle,
    (const CpaCyGenFlatBufCbFunc)asymCallback, /* CB function*/
    pCallbackTagPh2, /* pointer to the complete variable*/

```

```
pCpaDhOpDataP2,          /* structure containing p, the public value &
x*/
pOctetStringSecretKey); /* private key (output of the function)*/
```

In a real-world implementation of a key exchange protocol, the public value generated above would now be shared with another party, B. This example uses this public value to go on and invoke the second phase operation. First allocate memory for the secret value, setup the operational data for the phase 2 operation, and then perform that operation. This operation is invoked asynchronously, taking the callback function defined earlier as a parameter:

Listing 60. Perform Phase 2 Operation

```
status = cpaCyDhKeyGenPhase2Secret(
    cyInstHandle,
    (const CpaCyGenFlatBufCbFunc)asymCallback, /* CB function*/
    pCallbackTagPh2,          /* pointer to the complete variable*/
    pCpaDhOpDataP2,          /* structure containing p, the public value &
x*/
    pOctetStringSecretKey); /* private key (output of the function)*/
```

Finally, clean up by freeing up memory, stopping the instance, etc.

3.3.1 Prime Number Testing

This example demonstrates the usage of the prime number testing API.

These samples are located in `/asym/prime_sample`

The following steps are carried out:

- The API is used asynchronously: therefore, a callback function is defined as per the API prototype.
- Instance, discovery, and start-up is made in a way similar to that defined for the symmetric examples above.
- The function `primePerformOp` is called, which does the following:
 - Allocate memory for the operation
 - Populate data for the appropriate input fields and perform the operation. The fields populated include the following:
 - Prime Candidate
 - Whether to perform the greatest common divisor (GCD) test
 - Whether to perform the Fermat test
 - Number of Miller-Rabin rounds
 - Whether to perform Lucas test

Listing 61. Setup Operational Data and Test Prime

```
pPrimeTestOpData->primeCandidate.pData = pPrime;
pPrimeTestOpData->primeCandidate.dataLenInBytes =
sizeof(samplePrimeP_768);
pPrimeTestOpData->performGcdTest = CPA_TRUE;
pPrimeTestOpData->performFermatTest = CPA_TRUE;
```

```

pPrimeTestOpData->numMillerRabinRounds = NB_MR_ROUNDS;
pPrimeTestOpData->millerRabinRandomInput.pData = pMR;
pPrimeTestOpData->millerRabinRandomInput.dataLenInBytes = sizeof(MR);
pPrimeTestOpData->performLucasTest = CPA_TRUE;

status = cpaCyPrimeTest(
    cyInstHandle,
    (const CpaCyPrimeTestCbFunc)primeCallback, /* CB
function */
    (void *)&complete, /* callback tag */
    pPrimeTestOpData, /* operation data */
    &testPassed); /* return value: true if the number is probably
a prime, false if it is not a prime */

```

Finally, statistics are queried and the service stopped.

3.4 Using the SM2 API

This example demonstrates the usage of the SM2 API.

The following steps are carried out:

- The example contains synchronous and asynchronous API. For latter one, you must define a proper callback function per the API prototype for different SM2 operations.
- Instance, discovery, and start-up are made in a way similar to that defined for the symmetric examples above.
- This sample involves sign/verify, encryption/decryption, key exchange, point multiplication/verify.

3.4.1 SM2 Digital Signature Generation and Verification

This operation is to sign a given message and output its digital signature (r,s) then verify (r,s).

- The function `sampleEcdsm2SignPerformOp` provisions parts of example implementation, which does the following:
 - Allocate memory and populate data for input buffer which includes scalar multiplier k, digest of the message e and private key d.
 - Allocate memory for output buffer which includes signature r and s.
 - Call function `cpaCyEcdsm2Sign` for sign operation.

```

status = OS_MALLOC(&pCpaEcdsm2SignOpData, sizeof(CpaCyEcdsm2SignOpData));
status = PHYS_CONTIG_ALLOC(&pCpaEcdsm2SignOpData->k.pData, sizeof(k));
status = PHYS_CONTIG_ALLOC(&pCpaEcdsm2SignOpData->e.pData, sizeof(e));
status = PHYS_CONTIG_ALLOC(&pCpaEcdsm2SignOpData->d.pData, sizeof(d));

status = PHYS_CONTIG_ALLOC(&pR->pData, GFP_SM2_SIZE_IN_BYTE);
status = PHYS_CONTIG_ALLOC(&pS->pData, GFP_SM2_SIZE_IN_BYTE);

```

```

status = cpaCyEcsm2Sign(
    cyInstHandle,
    (const CpaCyEcsm2SignCbFunc)asymSignCallback, /* CB
function*/
    pCallbackTag, /* Opaque user
data */
    pCpaEcsm2SignOpData, /* Structure containing k, d and e */
    &signStatus, /* signStatus indicates if the result is valid
*/
    pR, /* Signature r (function output) */
    pS); /* Signature s (function output) */

```

- The function `sampleEcsm2VerifyPerformOp` provisions parts of example implementation, which does the following:
 - Allocate memory and populate data for input buffer which includes digest of the message e, signature r and s, x coordinate of public key and y coordinate of public key.
 - Call function `cpaCyEcsm2Verify` for signature verification operation.

```

status = OS_MALLOC(&pCpaEcsm2VerifyOpData,
sizeof(CpaCyEcsm2VerifyOpData));

status = PHYS_CONTIG_ALLOC(&pCpaEcsm2VerifyOpData->e.pData, sizeof(e));
status = PHYS_CONTIG_ALLOC(&pCpaEcsm2VerifyOpData->r.pData, sizeof(r));
status = PHYS_CONTIG_ALLOC(&pCpaEcsm2VerifyOpData->s.pData, sizeof(s));
status = PHYS_CONTIG_ALLOC(&pCpaEcsm2VerifyOpData->xP.pData, sizeof(xP));
status = PHYS_CONTIG_ALLOC(&pCpaEcsm2VerifyOpData->yP.pData,
sizeof(yPA));

status = cpaCyEcsm2Verify(
    cyInstHandle,
    (const CpaCyEcsm2VerifyCbFunc)asymVerifyCallback, /* CB
function*/
    pCallbackTag, /* Opaque user data */
    pCpaEcsm2VerifyOpData, /* Verify request data*/
    &verifyStatus); /* verify status */

```

3.4.2 SM2 Public Key Encryption

This operation is to encrypt a given message then decrypt the cipher and compare to the given message.

- The function `sampleEcsm2EncPerformOp` provisions parts of example implementation, which does the following:
 - Allocate memory and populate data for input buffer which includes scalar multiplier k, x coordinate of public key xP and y coordinate of public key yP.
 - Allocate memory for output buffer which includes x coordinate of [k]G x1, y coordinate of [k]G y1, x coordinate of [k]Pb x2 and y coordinate of [k]Pb y2.

- Call function `cpaCyEcsm2Encrypt` for encryption operation.

```

status = OS_MALLOC(&pCpaEcsm2EncOpData, sizeof(CpaCyEcsm2EncryptOpData));
status = PHYS_CONTIG_ALLOC(&pCpaEcsm2EncOpData->k.pData, sizeof(k));
status = PHYS_CONTIG_ALLOC(&pCpaEcsm2EncOpData->xP.pData, sizeof(xP));
status = PHYS_CONTIG_ALLOC(&pCpaEcsm2EncOpData->yP.pData, sizeof(yP));

status = OS_MALLOC(&pCpaEcsm2EncOutputData,
                  sizeof(CpaCyEcsm2EncryptOutputData));
status = PHYS_CONTIG_ALLOC(&pCpaEcsm2EncOutputData->x1.pData,
                          GFP_SM2_SIZE_IN_BYTE);
status = PHYS_CONTIG_ALLOC(&pCpaEcsm2EncOutputData->y1.pData,
                          GFP_SM2_SIZE_IN_BYTE);
status = PHYS_CONTIG_ALLOC(&pCpaEcsm2EncOutputData->x2.pData,
                          GFP_SM2_SIZE_IN_BYTE);
status = PHYS_CONTIG_ALLOC(&pCpaEcsm2EncOutputData->y2.pData,
                          GFP_SM2_SIZE_IN_BYTE);

status = cpaCyEcsm2Encrypt(
    cyInstHandle,
    (const CpaCyGenFlatBufCbFunc)asymEncCallback, /* CB
function*/
    pCallbackTag, /* Opaque user
data */
    pCpaEcsm2EncOpData, /* Encryption request data */
    pCpaEcsm2EncOutputData /* Encryption response data */);

```

- The function `sampleEcsm2DecPerformOp` provisions parts of example implementation, which does the following:
 - Allocate memory and populate data for input buffer which includes private key d , x coordinate of $[k]G$ $x1$ and y coordinate of $[k]G$ $y1$.
 - Allocate memory for output buffer which includes x coordinate of $[k]Pb$ $x2$ and y coordinate of $[k]Pb$ $y2$.
 - Call function `cpaCyEcsm2Decrypt` for decryption operation.
 - Call function `sm3` and `hashCheck` to check correctness of decryption.

```

status = OS_MALLOC(&pCpaEcsm2DecOpData, sizeof(CpaCyEcsm2DecryptOpData));

status = PHYS_CONTIG_ALLOC(&pCpaEcsm2DecOpData->d.pData, sizeof(d));
status = PHYS_CONTIG_ALLOC(&pCpaEcsm2DecOpData->x1.pData,
                          GFP_SM2_SIZE_IN_BYTE);
status = PHYS_CONTIG_ALLOC(&pCpaEcsm2DecOpData->y1.pData,
                          GFP_SM2_SIZE_IN_BYTE);

status = OS_MALLOC(&pCpaEcsm2DecOutputData,
                  sizeof(CpaCyEcsm2DecryptOutputData));
status = PHYS_CONTIG_ALLOC(&pCpaEcsm2DecOutputData->x2.pData,
                          GFP_SM2_SIZE_IN_BYTE);
status = PHYS_CONTIG_ALLOC(&pCpaEcsm2DecOutputData->y2.pData,
                          GFP_SM2_SIZE_IN_BYTE);

```

```

status = cpaCyEcdh2Decrypt(
    cyInstHandle,
    (const CpaCyGenFlatBufCbFunc)asymDecCallback, /* CB
function*/
    pCallbackTag, /* Opaque user
data */
    pCpaEcdh2DecOpData, /* Decryption request data */
    pCpaEcdh2DecOutputData /* Decryption response data */);

sm3(pDecOutputData, MESSAGE_LEN + 2 * GFP_SM2_SIZE_IN_BYTE,
    pHashBuffer);
hashCheck(pC3Buffer, pHashBuffer, &status);

```

3.4.3 SM2 Key Exchange

This operation is to exchange key between A side and B side, and check if the shared keys are the same.

- The function `sampleEcdh2KeyExPerformOp` provisions parts of example implementation, which does the following:
 - Allocate phase 1 input buffer which includes scalar multiplier r for A side and B side separately.
 - Allocate phase 1 output buffer which includes x coordinate of a point on the curve x and y coordinate of a point on the curve y for for A side and B side separately.
 - Call function `cpaCyEcdh2KeyExPhase1` for A side and B side separately.
 - Allocate phase 2 input buffer which includes scalar multiplier r , private key d , x coordinate of a point on the curve from other side x_1 , x coordinate of a point on the curve from phase 1 x_2 , y coordinate of a point on the curve from phase 1 y_2 , x coordinate of public key from other side x_P and y coordinate of public key from other side y_P for A side and B side separately.
 - Allocate phase 2 output buffer which includes x coordinate of a point on the curve x and y coordinate of a point on the curve y for A side and B side separately.
 - Call function `cpaCyEcdh2KeyExPhase2` for A side and B side separately.

```

status = OS_MALLOCS(&pCpaEcdh2KeyExPhase1AOpData,
    sizeof(CpaCyEcdh2KeyExPhase1OpData));
status = PHYS_CONTIG_ALLOCS(&pCpaEcdh2KeyExPhase1AOpData->r.pData,
    sizeof(rA));

status = OS_MALLOCS(&pCpaEcdh2KeyExPhase1BOpData,
    sizeof(CpaCyEcdh2KeyExPhase1OpData));
status = PHYS_CONTIG_ALLOCS(&pCpaEcdh2KeyExPhase1BOpData->r.pData,
    sizeof(rB));

status = OS_MALLOCS(&pCpaEcdh2KeyExPhase1AOutputData,
    sizeof(CpaCyEcdh2KeyExOutputData));

```



```

status = PHYS_CONTIG_ALLOC (&pCpaEcsm2KeyExPhase1AOutputData->x.pData,
                             GFP_SM2_SIZE_IN_BYTE);
status = PHYS_CONTIG_ALLOC (&pCpaEcsm2KeyExPhase1AOutputData->y.pData,
                             GFP_SM2_SIZE_IN_BYTE);

status = OS_MALLOC (&pCpaEcsm2KeyExPhase1BOutputData,
                    sizeof (CpaCyEcsm2KeyExOutputData));
status = PHYS_CONTIG_ALLOC (&pCpaEcsm2KeyExPhase1BOutputData->x.pData,
                             GFP_SM2_SIZE_IN_BYTE);
status = PHYS_CONTIG_ALLOC (&pCpaEcsm2KeyExPhase1BOutputData->y.pData,
                             GFP_SM2_SIZE_IN_BYTE);

status = cpaCyEcsm2KeyExPhase1 (
    cyInstHandle,
    (const CpaCyGenFlatBufCbFunc)
        asymKeyExPhase1Callback, /* CB function*/
    pCallbackTag,                /* Opaque user data */
    pCpaEcsm2KeyExPhase1AOpData, /* Key exchange p1 request data
*/
    pCpaEcsm2KeyExPhase1AOutputData /* Key exchange p1 response
data */);

status = cpaCyEcsm2KeyExPhase1 (
    cyInstHandle,
    (const CpaCyGenFlatBufCbFunc)
        asymKeyExPhase1Callback, /* CB function*/
    pCallbackTag,                /* Opaque user data */
    pCpaEcsm2KeyExPhase1BOpData, /* Key exchange p1 request data
*/
    pCpaEcsm2KeyExPhase1BOutputData /* Key exchange p1 request
data*/);

status = OS_MALLOC (&pCpaEcsm2KeyExPhase2AOpData,
                    sizeof (CpaCyEcsm2KeyExPhase2OpData));
status = PHYS_CONTIG_ALLOC (&pCpaEcsm2KeyExPhase2AOpData->r.pData,
                             sizeof (rA));
status = PHYS_CONTIG_ALLOC (&pCpaEcsm2KeyExPhase2AOpData->d.pData,
                             sizeof (dA));
status = PHYS_CONTIG_ALLOC (&pCpaEcsm2KeyExPhase2AOpData->x1.pData,
                             pCpaEcsm2KeyExPhase2AOpData->x1.dataLenInBytes);
status = PHYS_CONTIG_ALLOC (
    &pCpaEcsm2KeyExPhase2AOpData->x2.pData,
    pCpaEcsm2KeyExPhase1BOutputData->x.dataLenInBytes);
status = PHYS_CONTIG_ALLOC (
    &pCpaEcsm2KeyExPhase2AOpData->y2.pData,
    pCpaEcsm2KeyExPhase1BOutputData->y.dataLenInBytes);

```

```

status = PHYS_CONTIG_ALLOC (&pCpaEcsm2KeyExPhase2AOpData->xP.pData,
                             sizeof (xPB));
status = PHYS_CONTIG_ALLOC (&pCpaEcsm2KeyExPhase2AOpData->yP.pData,
                             sizeof (yPB));

status = OS_MALLOC (&pCpaEcsm2KeyExPhase2BOpData,
                    sizeof (CpaCyEcsm2KeyExPhase2OpData));
status = PHYS_CONTIG_ALLOC (&pCpaEcsm2KeyExPhase2BOpData->r.pData,
                             sizeof (rB));
status = PHYS_CONTIG_ALLOC (&pCpaEcsm2KeyExPhase2BOpData->d.pData,
                             sizeof (dB));
status = PHYS_CONTIG_ALLOC (&pCpaEcsm2KeyExPhase2BOpData->x1.pData,
                             pCpaEcsm2KeyExPhase2BOpData -
>x1.dataLenInBytes);
status = PHYS_CONTIG_ALLOC (
    &pCpaEcsm2KeyExPhase2BOpData->x2.pData,
    pCpaEcsm2KeyExPhase1AOutputData->x.dataLenInBytes);
status = PHYS_CONTIG_ALLOC (
    &pCpaEcsm2KeyExPhase2BOpData->y2.pData,
    pCpaEcsm2KeyExPhase1AOutputData->y.dataLenInBytes);
status = PHYS_CONTIG_ALLOC (&pCpaEcsm2KeyExPhase2BOpData->xP.pData,
                             sizeof (xPA));
status = PHYS_CONTIG_ALLOC (&pCpaEcsm2KeyExPhase2BOpData->yP.pData,
                             sizeof (yPA));

status = OS_MALLOC (&pCpaEcsm2KeyExPhase2AOutputData,
                    sizeof (CpaCyEcsm2KeyExOutputData));
status = PHYS_CONTIG_ALLOC (&pCpaEcsm2KeyExPhase2AOutputData->x.pData,
                             GFP_SM2_SIZE_IN_BYTE);
status = PHYS_CONTIG_ALLOC (&pCpaEcsm2KeyExPhase2AOutputData->y.pData,
                             GFP_SM2_SIZE_IN_BYTE);

status = OS_MALLOC (&pCpaEcsm2KeyExPhase2BOutputData,
                    sizeof (CpaCyEcsm2KeyExOutputData));
status = PHYS_CONTIG_ALLOC (&pCpaEcsm2KeyExPhase2BOutputData->x.pData,
                             GFP_SM2_SIZE_IN_BYTE);
status = PHYS_CONTIG_ALLOC (&pCpaEcsm2KeyExPhase2BOutputData->y.pData,
                             GFP_SM2_SIZE_IN_BYTE);

status = cpaCyEcsm2KeyExPhase2 (
    cyInstHandle,
    (const CpaCyGenFlatBufCbFunc) asymKeyExPhase2Callback,
    pCallbackTag, /* Opaque user data; */
    pCpaEcsm2KeyExPhase2AOpData, /* Key exchange p2 request data,
    containing
r, d, x1, x2, y2, xp, yp*/

```

```

        pCpaEcsm2KeyExPhase2AOutputData /* Key exchange p2 response
*/);

status = cpaCyEcsm2KeyExPhase2 (
    cyInstHandle,
    (const CpaCyGenFlatBufCbFunc) asymKeyExPhase2Callback,
    pCallbackTag, /* Opaque user data */
    pCpaEcsm2KeyExPhase2BOpData, /* Key exchange p2 request data
                                containing
r, d, x1, y1, x2, y2, xp, yp*/
    pCpaEcsm2KeyExPhase2BOutputData /* Key exchange p2 response
data*/);

```

3.4.4 SM2 Elliptic Curve Point

This operation is to calculate a point on the curve according to a given random number and verify if the point (x,y) on the curve or not.

- The function `sampleEcsm2PointMultiply` provisions parts of example implementation, which does the following:
 - Allocate memory and populate data for input buffer which includes scalar multiplier k, x coordinate of a point on the curve x and y coordinate of a point on the curve y.
 - Allocate memory for output buffer which includes x coordinate of the resulting point multiplication pXk and y coordinate of the resulting point multiplication pYk.
 - Call function `cpaCyEcsm2PointMultiply` for point multiply operation.

```

status = OS_MALLOC(&opData, sizeof(CpaCyEcsm2PointMultiplyOpData));
status = PHYS_CONTIG_ALLOC(&opData->k.pData, sizeof(k));
status = PHYS_CONTIG_ALLOC(&opData->x.pData, sizeof(xP));
status = PHYS_CONTIG_ALLOC(&opData->y.pData, sizeof(yP));

status = OS_MALLOC(&pXk, sizeof(CpaFlatBuffer));
status = OS_MALLOC(&pYk, sizeof(CpaFlatBuffer));
status = PHYS_CONTIG_ALLOC(&pXk->pData, GFP_SM2_SIZE_IN_BYTE);
status = PHYS_CONTIG_ALLOC(&pYk->pData, GFP_SM2_SIZE_IN_BYTE);

status = cpaCyEcsm2PointMultiply(
    cyInstHandle,
    (const CpaCyEcPointMultiplyCbFunc)
        asymPointMultCallback, /* CB function*/
    pCallbackTag, /* Opaque user data; */
    opData, /* Point multiplication request
data */
    &multiplyStatus,
    pXk, /* Point multiplication response data */
    pYk /* Point multiplication response data */

```

```
);
```

- The function `sampleEcdsm2GeneratorMultiply` provisions parts of example implementation, which does the following:
 - Allocate memory and populate data for input buffer which includes scalar multiplier `k`.
 - Allocate memory for output buffer which includes `x` coordinate of the resulting point multiplication `pXk` and `y` coordinate of the resulting point multiplication `pYk`.
 - Call function `cpaCyEcdsm2GeneratorMultiply` for generator multiply operation.

```
status = OS_MALLOC(&opData, sizeof(CpaCyEcdsm2GeneratorMultiplyOpData));
status = PHYS_CONTIG_ALLOC(&opData->k.pData, sizeof(k));

status = OS_MALLOC(&pXk, sizeof(CpaFlatBuffer));
status = OS_MALLOC(&pYk, sizeof(CpaFlatBuffer));
status = PHYS_CONTIG_ALLOC(&pXk->pData, GFP_SM2_SIZE_IN_BYTE);
status = PHYS_CONTIG_ALLOC(&pYk->pData, GFP_SM2_SIZE_IN_BYTE);

status = cpaCyEcdsm2GeneratorMultiply(
    cyInstHandle,
    (const CpaCyEcPointMultiplyCbFunc)
        asymGeneratorMultCallback, /* CB function*/
    pCallbackTag, /* Opaque user data; */
    opData, /* Generator multiplication request data */
    &multiplyStatus,
    pXk, /* Generator multiplication response data */
    pYk); /* Generator multiplication response data */
pCallbackTag, /* Opaque user data; */
opData, /* Point multiplication request
data */
    &multiplyStatus,
    pXk, /* Point multiplication response data */
    pYk /* Point multiplication response data */
);
```

- The function `sampleEcdsm2PointVerify` provisions parts of example implementation, which does the following:
 - Allocate memory and populate data for input buffer which includes `x` coordinate of a point on the curve `x` and `y` coordinate of a point on the curve `y`.
 - Call function `cpaCyEcdsm2PointVerify` for EC point verification.

```
status = OS_MALLOC(&opData, sizeof(CpaCyEcdsm2PointVerifyOpData));
status = PHYS_CONTIG_ALLOC(&opData->x.pData, sizeof(x));
status = PHYS_CONTIG_ALLOC(&opData->y.pData, sizeof(y));

status = cpaCyEcdsm2PointVerify(cyInstHandle,
    (const CpaCyEcPointVerifyCbFunc)
        asymPointVerifyCallback, /* CB
function*/
```

```
                                pCallbackTag, /* Opaque user data */
                                opData,        /* Point verify request
data */
                                &verifyStatus);

status = cpaCyEcdsm2GeneratorMultiply(
        cyInstHandle,
        (const CpaCyEcPointMultiplyCbFunc)
        asymGeneratorMultCallback, /* CB function*/
```

Finally, statistics are queried and the service stopped.

§

4 Intel® QuickAssist Technology Data Compression API

This chapter describes the sample code for the Intel® QuickAssist Technology Data Compression API, beginning with an API overview, and followed by descriptions of various scenarios to illustrate the usage of the API.

Note: This document does not cover data integrity concepts. Refer to [Table 1](#) in the Programmer's Guide, Compress and Verify (CnV) Related APIs for your product for important information on data integrity concepts, including the Compress-and-Verify feature.

4.1 Overview

The Intel® QuickAssist Technology Data Compression API can be categorized into three broad areas as follows:

- **Common:** This includes functionality for the initialization and shutdown of the service.
- **Instance Management:** A given implementation of the API can present multiple instances of the compression service, each representing a logical or virtual "device". Request order is guaranteed within a given instance of the service.
- **Transformation:**
 - Compression functionality
 - Decompression functionality

These areas of functionality are defined in [cpa_dc.h](#) and [cpa_dc_dp.h](#).

The Intel® QAT Data Compression API uses the "base" API ([cpa](#)), which defines base data types used across all services of the Intel® QAT Technology API.

4.1.1 Session

Similar to the symmetric cryptography API, the data compression API has the concept of a session. In the case of the compression API, a session is an object that describes the compression parameters to be applied across several requests. These requests might submit buffers within a single file, or buffers associated with a particular data stream or flow. A session object is described by the following:

- **The compression level:** Lower levels provide faster compression and the cost of compression ratio, whereas higher levels provide a better compression ratio as the cost of performance.
- **The compression algorithm:** to use (e.g. deflate) and what type of Huffman trees to use (static or dynamic).
- **The session direction:** If all requests on this session are compression requests, then the direction can be set to compress (and similarly, for decompress). A combined direction is also available if both compression and decompression requests are called using this session.

- **The session state:** a session can be described as stateful or stateless. Stateful sessions maintain history and state between calls to the API, and stateless sessions do not.
 - Stateless compression does not require history data from a previous compression/decompression request to be restored before submitting the request. Stateless sessions are used when the output data is known to be constrained in size. An overflow condition (when the output data is about to exceed the output buffer) is treated as an error condition in the decompression direction. In the compression direction, the Application can keep submitting data from where the overflow was registered in the input stream. The Data Plane API treats overflow as an error. In this case, the overflow is treated as an error rather than an exception. The client application is required to resubmit the job in its entirety with a larger output buffer. Requests are treated independently; state and history are not saved and restored between calls.

Note: When using a stateless session, it is possible to feed a seed checksum to the `cpaDcCompressData()` or the `cpaDcDecompressData()` API when the `CPA_DC_FLUSH_FULL` flush flag is used. The user application is responsible for maintaining the checksum across requests. This feature is also known as Stateful Lite.

- Stateful sessions are required when the data to be decompressed is larger than the buffers being used. This is a standard mode of operation for applications such as GZIP, where the size of the uncompressed data is not known before execution, and therefore the destination buffer may not be large enough to hold the resultant output. Requests to stateful sessions are not treated independently, and state and history can be saved and restored between calls. The amount of history and state carried between calls depends on the compression level. For stateful decompression, only one outstanding request may be in-flight at any one time for that session.

4.2 Sample – Stateful Data Compression

This example demonstrates the usage of the synchronous API, specifically using this API to perform a compression operation. It compresses a file via a stateful session using the deflate compress algorithm with static Huffman trees and using GZIP style headers and footers.

These samples are located in `/dc/stateful_sample`

Note: Stateful data compression is not available in Intel® QAT v1.8 and later releases. However, stateful decompression is available in Intel® QAT v1.8 and later releases.

4.2.1 Session Establishment

This is the main entry point for the sample compression code. It demonstrates the sequence of calls to be made to the API to create a session, perform one or more compress operations, and then tear down the session. At this point, the instance has been discovered and started, and the capabilities of the instance have been queried and found to be suitable.

A session is established by describing a session, determining how much session memory is required, and then invoking the session initialization function `cpaDcInitSession`.

Listing 62. Create and Initialize Stateful Session

```
sd.compLevel = CPA_DC_L4;
sd.compType = CPA_DC_DEFLATE;
sd.huffType = CPA_DC_HT_STATIC;
sd.sessDirection = CPA_DC_DIR_COMBINED;
sd.sessState = CPA_DC_STATEFUL;
#if (CPA_DC_API_VERSION_NUM_MAJOR == 1 && CPA_DC_API_VERSION_NUM_MINOR <
6)
sd.deflateWindowSize = 7;
#endif
sd.checksum = CPA_DC_CRC32;

/* Determine size of session context to allocate */
PRINT_DBG("cpaDcGetSessionSize\n");
status = cpaDcGetSessionSize(dcInstHandle, &sd, &sess_size, &ctx_size);
}
if (CPA_STATUS_SUCCESS == status) {
    /* Allocate session memory */
    status = PHYS_CONTIG_ALLOC(&sessionHdl, sess_size);
}
if ((CPA_STATUS_SUCCESS == status) && (ctx_size != 0)) {
    /* Allocate context bufferlist */
    status = cpaDcBufferListGetMetaSize(dcInstHandle, 1, &buffMetaSize);

    if (CPA_STATUS_SUCCESS == status) {
        status = PHYS_CONTIG_ALLOC(&pBufferMeta, buffMetaSize);
    }

    if (CPA_STATUS_SUCCESS == status) {
        status = OS_MALLOC(&pBufferCtx, bufferListMemSize);
    }

    if (CPA_STATUS_SUCCESS == status) {
        status = PHYS_CONTIG_ALLOC(&pCtxBuf, ctx_size);
    }

    if (CPA_STATUS_SUCCESS == status) {
        pFlatBuffer = (CpaFlatBuffer *) (pBufferCtx + 1);

        pBufferCtx->pBuffers = pFlatBuffer;
        pBufferCtx->numBuffers = 1;
        pBufferCtx->pPrivateMetaData = pBufferMeta;
    }
}
```



```

        pFlatBuffer->dataLenInBytes = ctx_size;
        pFlatBuffer->pData = pCtxBuf;
    }
}
/* Initialize the Stateful session */
if (CPA_STATUS_SUCCESS == status) {
    PRINT_DBG("cpaDcInitSession\n");
    status = cpaDcInitSession(dcInstHandle,
                             sessionHdl, /* session memory */
                             &sd,      /* session setup data */
                             pBufferCtx, /* context buffer */
                             NULL); /* callback function NULL for sync
mode */
}

```

Note: Source and destination buffers must be established.

Listing 63. Stateful Compression Memory Allocation

```

numBuffers = 1; /* only using 1 buffer in this case */
/* allocate memory for bufferlist and array of flat buffers in a
contiguous
* area and carve it up to reduce number of memory allocations required.
*/
bufferListMemSize =
    sizeof(CpaBufferList) + (numBuffers * sizeof(CpaFlatBuffer));

status = cpaDcBufferListGetMetaSize(dcInstHandle, numBuffers,
&bufferMetaSize);

/* Allocate source buffer */
if (CPA_STATUS_SUCCESS == status) {
    status = PHYS_CONTIG_ALLOC(&pBufferMetaSrc, bufferMetaSize);
}
if (CPA_STATUS_SUCCESS == status) {
    status = OS_MALLOC(&pBufferListSrc, bufferListMemSize);
}
if (CPA_STATUS_SUCCESS == status) {
    status = PHYS_CONTIG_ALLOC(&pSrcBuffer, SAMPLE_BUFF_SIZE);
}

/* Allocate destination buffer the same size as source buffer */
if (CPA_STATUS_SUCCESS == status) {
    status = PHYS_CONTIG_ALLOC(&pBufferMetaDst, bufferMetaSize);
}
if (CPA_STATUS_SUCCESS == status) {
    status = OS_MALLOC(&pBufferListDst, bufferListMemSize);
}

```

```

}
if (CPA_STATUS_SUCCESS == status) {
    status = PHYS_CONTIG_ALLOC(&pDstBuffer, SAMPLE_BUFF_SIZE);
}

if (CPA_STATUS_SUCCESS == status) {
    /* Build source bufferList */
    pFlatBuffer = (CpaFlatBuffer *) (pBufferListSrc + 1);

    pBufferListSrc->pBuffers = pFlatBuffer;
    pBufferListSrc->numBuffers = 1;
    pBufferListSrc->pPrivateMetaData = pBufferMetaSrc;

    pFlatBuffer->dataLenInBytes = SAMPLE_BUFF_SIZE;
    pFlatBuffer->pData = pSrcBuffer;

    /* Build destination bufferList */
    pFlatBuffer = (CpaFlatBuffer *) (pBufferListDst + 1);

    pBufferListDst->pBuffers = pFlatBuffer;
    pBufferListDst->numBuffers = 1;
    pBufferListDst->pPrivateMetaData = pBufferMetaDst;

    pFlatBuffer->dataLenInBytes = SAMPLE_BUFF_SIZE;
    pFlatBuffer->pData = pDstBuffer;
}

```

Listing 64. Create Header

At this point, the application has opened an instance, established a session, and allocated buffers. It is time to start some compress operations. To produce GZIP style compressed files, the first thing that needs to be performed is header generation. Create a header using the following code:

```

/* Write RFC1952 gzip header to destination buffer */
status = cpaDcGenerateHeader(sessionHdl, pFlatBuffer, &hdr_sz);
}
if (CPA_STATUS_SUCCESS == status) {
    /* write out header */
    fwrite(pFlatBuffer->pData, 1, hdr_sz, dstFile);
}

```

`cpaDcGenerateHeader` produces a GZIP style header (compliant with *GZIP file format specification* v4.3, RFC 1952, refer to [Table 1](#)) when the session set up data is set such that `compType` is `CPA_DC_DEFLATE` and checksum is `CPA_DC_CRC32`.

Note: Alternatively, a zlib style header (compliant with *ZLIB Compressed Data Format Specification*, v3.3, RFC 1950, refer to [Table 1](#)) can be produced if the session setup data is set such that `compType` is `CPA_DC_DEFLATE` and checksum is `CPA_DC_ADLER32`. This operation

demonstrates looping through a file, reading the data, invoking the data compress operation, and writing the results to the output file.

Listing 65. Perform Stateful Compression Operation

```

pBufferListSrc->pBuffers->dataLenInBytes = 0;
while ((!feof(srcFile)) && (CPA_STATUS_SUCCESS == status)) {
    /* read from file into src buffer */
    pBufferListSrc->pBuffers->pData = pSrcBuffer;
    pBufferListSrc->pBuffers->dataLenInBytes +=
        fread(pSrcBuffer + pBufferListSrc->pBuffers->dataLenInBytes,
            1,
            SAMPLE_BUFF_SIZE - pBufferListSrc->pBuffers->
>dataLenInBytes,
            srcFile);
    if (pBufferListSrc->pBuffers->dataLenInBytes < SAMPLE_BUFF_SIZE) {
        flush = CPA_DC_FLUSH_FINAL;
    } else {
        flush = CPA_DC_FLUSH_SYNC;
    }
    do {
        PRINT_DBG("cpaDcCompressData\n");
        status = cpaDcCompressData(dcInstHandle,
            sessionHdl,
            pBufferListSrc, /* source buffer list
*/
            pBufferListDst, /* destination buffer
list */
            &dcResults, /* results structure
*/
            flush, /* Stateful session */
            NULL);

        if (CPA_STATUS_SUCCESS != status) {
            PRINT_ERR("cpaDcCompressData failed. (status = %d)\n",
status);
            break;
        }

        /*
         * We now check the results
         */
        if ((dcResults.status != CPA_DC_OK) &&
            (dcResults.status != CPA_DC_OVERFLOW)) {
            PRINT_ERR("Results status not as expected (status = %d)\n",
                dcResults.status);
            status = CPA_STATUS_FAIL;
            break;

```

```

    }

    fwrite(pDstBuffer, 1, dcResults.produced, dstFile);

    if (dcResults.consumed <= pBufferListSrc->pBuffers->dataLenInBytes) {
        pBufferListSrc->pBuffers->dataLenInBytes -=
dcResults.consumed;
        pBufferListSrc->pBuffers->pData += dcResults.consumed;
    } else {
        pBufferListSrc->pBuffers->dataLenInBytes = 0;
    }

    if (dcResults.consumed == 0 &&
        pBufferListSrc->pBuffers->dataLenInBytes > 0) {
        memcpy(pSrcBuffer,
            pBufferListSrc->pBuffers->pData,
            pBufferListSrc->pBuffers->dataLenInBytes);
        break;
    }
} while (pBufferListSrc->pBuffers->dataLenInBytes != 0 ||
        dcResults.status == CPA_DC_OVERFLOW);
}

```

Finally, a GZIP footer is generated. Similar to the call to `cpaDcGenerateHeader`, a GZIP footer (compliant with GZIP style header (compliant with *GZIP file format specification* v4.3, RFC 1952, refer to [Table 1](#)) is produced because the session setup data is set such that `compType` is `CPA_DC_DEFLATE` and checksum is `CPA_DC_CRC32`. The call to `cpaDcGenerateFooter` increments the `produced` field of the `CpaDcRqResults` structure by the size of the footer added. In this example, the data produced so far has already been written out to the file. As such, the `produced` field of the `CpaDcRqResults` structure is cleared before calling the `cpaDcGenerateFooter` function.

In the event the destination buffer would be too small to accept the footer, the `cpaDcGenerateFooter()` API will return an invalid parameter error. The `cpaDcGenerateFooter()` API cannot return an overflow exception. It is application's responsibility to ensure that there is enough allocated buffer memory to append the algorithm specific footer.

Listing 66. Create Footer

```

        dcResults.produced = 0;
        /* Write RFC1952 gzip footer to destination buffer */
        status = cpaDcGenerateFooter(sessionHdl, pFlatBuffer,
&dcResults);
    }
    if (CPA_STATUS_SUCCESS == status) {
        /* write out footer */
        fwrite(pFlatBuffer->pData, 1, dcResults.produced, dstFile);
    }
}

```

}

Because this session was created with `CPA_DC_DIR_COMBINED` it can also be used to decompress data.

The Stateful Decompression Operation demonstrates looping through a file, reading the compressed data, invoking the data decompress operation, and writing the results to the output file. In this case, the overflow condition has to be considered.

Listing 67. Perform Stateful Decompression Operation

```
pBufferListSrc->pBuffers->dataLenInBytes = 0;
while ((!feof(srcFile)) && (CPA_STATUS_SUCCESS == status))
{
    /* read from file into src buffer */
    pBufferListSrc->pBuffers->pData = pSrcBuffer;
    pBufferListSrc->pBuffers->dataLenInBytes += fread(
        pSrcBuffer + pBufferListSrc->pBuffers->dataLenInBytes,
        1,
        SAMPLE_BUFF_SIZE - pBufferListSrc->pBuffers->dataLenInBytes,
        srcFile);

    if (pBufferListSrc->pBuffers->dataLenInBytes < SAMPLE_BUFF_SIZE)
    {
        /* FLUSH FINAL flag must be set for last request */
        opData.flushFlag = CPA_DC_FLUSH_FINAL;
    }
    else
    {
        /* FLUSH SYNC flag must be set for intermediate requests */
        opData.flushFlag = CPA_DC_FLUSH_SYNC;
    }
}

do
{
    status = cpaDcDecompressData2(
        dcInstHandle,
        sessionHdl,
        pBufferListSrc, /* source buffer list */
        pBufferListDst, /* destination buffer list */
        &opData,
        &dcResults, /* results structure */
        NULL);
    if (CPA_STATUS_SUCCESS != status)
    {
        PRINT_ERR(
            "cpaDcDecompressData2 failed. (status = %d)\n",
            status);
    }
}
```

```
        break;
    }

    /*
     * We now check the results - in decompress direction the
     * output buffer may overflow
     */
    if ((dcResults.status != CPA_DC_OK) &&
        (dcResults.status != CPA_DC_OVERFLOW))
    {
        PRINT_ERR(
            "Results status not as expected (status = %d)\n",
            dcResults.status);
        status = CPA_STATUS_FAIL;
        break;
    }

    /* The gzip file generated by Deflate algorithm has an
     * 8-byte
     * footer, containing a CRC-32 checksum and the length of
     * the
     * original uncompressed data. The 'endOfLastBlock' flag
     * tells
     * if we have processed the last data block. Break the loop
     * here, otherwise it will keep on reading gzip file.
     */
    if (CPA_TRUE == dcResults.endOfLastBlock)
    {
        break;
    }

    if (dcResults.consumed <=
        pBufferListSrc->pBuffers->dataLenInBytes)
    {
        pBufferListSrc->pBuffers->dataLenInBytes -=
            dcResults.consumed;
        pBufferListSrc->pBuffers->pData += dcResults.consumed;
    }
    else
    {
        pBufferListSrc->pBuffers->dataLenInBytes = 0;
    }

    if (dcResults.consumed == 0 &&
        pBufferListSrc->pBuffers->dataLenInBytes > 0)
    {
```

```

        memcpy(pSrcBuffer,
               pBufferListSrc->pBuffers->pData,
               pBufferListSrc->pBuffers->dataLenInBytes);
        break;
    }
} while (pBufferListSrc->pBuffers->dataLenInBytes != 0 ||
        dcResults.status == CPA_DC_OVERFLOW);
}

```

Once all operations on this session have been completed, the session is torn down using the Remove Stateful Session in [Listing 68](#).

Listing 68. Remove Stateful Session

```
sessionStatus = cpaDcRemoveSession(dcInstHandle, sessionHdl);
```

Query statistics at this point, which can be useful for debugging.

Finally, clean up by freeing up memory, stopping the instance, etc.

4.2.2 Sample – Stateless Data Compression

This example demonstrates the usage of the asynchronous API, specifically using this API to perform a compression operation. It compresses a data buffer through a stateless session using the deflate compress algorithm with dynamic Huffman trees.

The example below compresses a block of data into a compressed block.

These samples are located in `/dc/stateless_sample`

In this example, dynamic Huffman trees are used. The instance can be queried to ensure dynamic Huffman trees are supported, and if an instance-specific buffer is required to perform a dynamic Huffman tree deflate request.

Listing 69. Querying and Starting a Compression Instance

```

status = cpaDcQueryCapabilities(dcInstHandle, &cap);
if (status != CPA_STATUS_SUCCESS) {
    return status;
}

if (!cap.statelessDeflateCompression ||
    !cap.statelessDeflateDecompression ||
    !cap.checksumAdler32 ||
    !cap.dynamicHuffman) {
    PRINT_DBG("Error: Unsupported functionality\n");
    return CPA_STATUS_FAIL;
}

if (cap.dynamicHuffmanBufferReq) {
    status = cpaDcBufferListGetMetaSize(dcInstHandle, 1, &buffMetaSize);
}

```

```

    if (CPA_STATUS_SUCCESS == status) {
        status =
            cpaDcGetNumIntermediateBuffers(dcInstHandle,
&numInterBuffLists);
    }
    if (CPA_STATUS_SUCCESS == status && 0 != numInterBuffLists) {
        status = PHYS_CONTIG_ALLOC(&bufferInterArray,
            numInterBuffLists *
sizeof(CpaBufferList *));
    }
    for (bufferNum = 0; bufferNum < numInterBuffLists; bufferNum++) {
        if (CPA_STATUS_SUCCESS == status) {
            status = PHYS_CONTIG_ALLOC(&bufferInterArray[bufferNum],
                sizeof(CpaBufferList));
        }
        if (CPA_STATUS_SUCCESS == status) {
            status = PHYS_CONTIG_ALLOC(
                &bufferInterArray[bufferNum]->pPrivateMetaData,
buffMetaSize);
        }
        if (CPA_STATUS_SUCCESS == status) {
            status = PHYS_CONTIG_ALLOC(&bufferInterArray[bufferNum]-
>pBuffers,
                sizeof(CpaFlatBuffer));
        }
        if (CPA_STATUS_SUCCESS == status) {
            /* Implementation requires an intermediate buffer
approximately
                twice the size of the output buffer */
            status =
                PHYS_CONTIG_ALLOC(&bufferInterArray[bufferNum]->pBuffers-
>pData,
                2 * SAMPLE_MAX_BUFF);
            bufferInterArray[bufferNum]->numBuffers = 1;
            bufferInterArray[bufferNum]->pBuffers->dataLenInBytes =
                2 * SAMPLE_MAX_BUFF;
        }
    } /* End numInterBuffLists */
}

if (CPA_STATUS_SUCCESS == status) {
    /*
    * Set the address translation function for the instance
    */
    status = cpaDcSetAddressTranslation(dcInstHandle, sampleVirtToPhys);
}

```



```

}

if (CPA_STATUS_SUCCESS == status) {
    /* Start DataCompression component */
    PRINT_DBG("cpaDcStartInstance\n");
    status =
        cpaDcStartInstance(dcInstHandle, numInterBuffLists,
bufferInterArray);
}

```

The create and initialize stateless session demonstrates the sequence of calls to be made to the API to create a session. To establish a session: describing the session, determining how much session memory is required, and then invoke the session initialization function `cpaDcInitSession`.

Listing 70. Create and Initialize Stateless Session

```

sd.compLevel = CPA_DC_L4;
sd.compType = CPA_DC_DEFLATE;
sd.huffType = CPA_DC_HT_FULL_DYNAMIC;
/* If the implementation supports it, the session will be configured
 * to select static Huffman encoding over dynamic Huffman as
 * the static encoding will provide better compressibility.
 */
if (cap.autoSelectBestHuffmanTree) {
    sd.autoSelectBestHuffmanTree = CPA_TRUE;
} else {
    sd.autoSelectBestHuffmanTree = CPA_FALSE;
}
sd.sessDirection = CPA_DC_DIR_COMBINED;
sd.sessState = CPA_DC_STATELESS;
#if (CPA_DC_API_VERSION_NUM_MAJOR == 1 && CPA_DC_API_VERSION_NUM_MINOR <
6)
sd.deflateWindowSize = 7;
#endif
sd.checksum = CPA_DC_ADLER32;

/* Determine size of session context to allocate */
PRINT_DBG("cpaDcGetSessionSize\n");
status = cpaDcGetSessionSize(dcInstHandle, &sd, &sess_size, &ctx_size);
}

if (CPA_STATUS_SUCCESS == status) {
    /* Allocate session memory */
    status = PHYS_CONTIG_ALLOC(&sessionHdl, sess_size);
}

/* Initialize the Stateless session */

```

```

if (CPA_STATUS_SUCCESS == status) {
    PRINT_DBG("cpaDcInitSession\n");
    status = cpaDcInitSession(
        dcInstHandle,
        sessionHdl, /* session memory */
        &sd,         /* session setup data */
        NULL,      /* pContextBuffer not required for stateless
operations */
        dcCallback); /* callback function */
}

```

Source and destination buffers are allocated in a similar way to the stateful example above.

Perform Operation: This listing demonstrates invoking the data compress operation, in the stateless case.

Listing 71. Data Plane Remove Compression Session

```
sessionStatus = cpaDcRemoveSession(dcInstHandle, sessionHdl);
```

4.3 Sample – Stateless Data Compression Using Multiple Compress Operations

This example demonstrates the use of the asynchronous API: specifically, using this API to perform a compression operation. It compresses a data buffer using multiple stateless compression API requests and maintains length and checksum information across the multiple requests without the overhead of maintaining full history information as used in a stateful operation.

The samples are located in: [/dc/stateless_multi_op_checksum_sample](#)

In this sample, session creation is the same as for regular stateless operation. Refer to the previous sample described in Section [4.3, Sample – Stateless Data Compression Using Multiple Compress Operations](#) for details.

Perform Operation: This listing demonstrates the invoking of the data compress operation in the stateless case while maintaining checksum information across multiple compress operations. The key points to note are:

The initial value of `dcResults.checksum` is set to 0 for CRC32 or set to 1 for Adler32 when invoking the first compress or decompress operation for a data set.

Listing 72. Setting the Initial Value of the Checksum

```

if (sd.checksum == CPA_DC_ADLER32) {
    /* Initialize checksum to 1 for Adler32 */
    dcResults.checksum = 1;
} else {
    /* Initialize checksum to 0 for CRC32 */
    dcResults.checksum = 0;
}

```

```
}

```

The value of `dcResults.checksum` when invoking a subsequent compress operation for a data set is set to the `dcResults`. Checksum value returned from the previous compress operation on that data set.

4.4 Sample – Data Compression Data Plane API

This example demonstrates the usage of the data plane data compression API to perform a compression operation. It compresses a data buffer via a stateless session using the deflate compress algorithm with dynamic Huffman trees. This example is simplified to demonstrate the basics of how to use the API and how to build the structures required. This example does not demonstrate the optimal way to use the API to get maximum performance for a particular implementation. Refer to [Table 1](#) Implementation Specific Documentation and performance sample code for a guide on how to use the API for best performance.

These samples are located in `/dc/dc_dp_sample`

The data plane data compression API is used in a similar way to the data plane-symmetric cryptographic API:

Data compression service instances are queried and started in the same way and using the same functions as before (see [Listing 1](#) and [Listing 68](#)).

Listing 73. Register Compression Callback Function

This listing registers a callback function for the data compression instance.

```
status = cpaDcDpRegCbFunc(dcInstHandle, dcDpCallback);
```

Next, create and initialize a session.

Listing 74. Create and Initialize Compression Data Plane Session

```
if (CPA_STATUS_SUCCESS == status) {
    sd.compLevel = CPA_DC_L4;
    sd.compType = CPA_DC_DEFLATE;
    sd.huffType = CPA_DC_HT_FULL_DYNAMIC;
    /* If the implementation supports it, the session will be configured
     * to select static Huffman encoding over dynamic Huffman as
     * the static encoding will provide better compressibility.
     */
    if (cap.autoSelectBestHuffmanTree) {
        sd.autoSelectBestHuffmanTree = CPA_TRUE;
    } else {
        sd.autoSelectBestHuffmanTree = CPA_FALSE;
    }
    sd.sessDirection = CPA_DC_DIR_COMBINED;
    sd.sessState = CPA_DC_STATELESS;
#if (CPA_DC_API_VERSION_NUM_MAJOR == 1 && CPA_DC_API_VERSION_NUM_MINOR <
6)

```

```

    sd.deflateWindowSize = 7;
#endif
    sd.checksum = CPA_DC_CRC32;

    /* Determine size of session context to allocate */
    PRINT_DBG("cpaDcGetSessionSize\n");
    status = cpaDcGetSessionSize(dcInstHandle, &sd, &sess_size,
&ctx_size);
}

if (CPA_STATUS_SUCCESS == status) {
    /* Allocate session memory */
    status = PHYS_CONTIG_ALLOC(&sessionHdl, sess_size);
}

/* Initialize the Stateless session */
if (CPA_STATUS_SUCCESS == status) {
    PRINT_DBG("cpaDcDpInitSession\n");
    status = cpaDcDpInitSession(dcInstHandle,
                                sessionHdl, /* session memory */
                                &sd);      /* session setup data */
}

```

Listing 75. Setup Source Buffer

In this example, input and output data is stored in a scatter gather list. The source and destination buffers are described using the `CpaPhysBufferList` structure. In this example the allocation (which needs to be 8-byte aligned) and setup of the source buffer is shown. The destination buffers can be allocated and set up in a similar way.

```

numBuffers = 2;
/* Size of CpaPhysBufferList and array of CpaPhysFlatBuffers */
bufferListMemSize =
    sizeof(CpaPhysBufferList) + (numBuffers * sizeof(CpaPhysFlatBuffer));

/* Allocate 8-byte aligned source buffer List */
status = PHYS_CONTIG_ALLOC_ALIGNED(&pBufferListSrc, bufferListMemSize,
8);
if (CPA_STATUS_SUCCESS == status) {
    /* Allocate first data buffer to hold half the data */
    status = PHYS_CONTIG_ALLOC(&pSrcBuffer, (sizeof(sampleData)) / 2);
}
if (CPA_STATUS_SUCCESS == status) {
    /* Allocate second data buffer to hold half the data */
    status = PHYS_CONTIG_ALLOC(&pSrcBuffer2, (sizeof(sampleData)) / 2);
}
if (CPA_STATUS_SUCCESS == status) {
    /* copy source into buffer */

```

```

memcpy(pSrcBuffer, sampleData, sizeof(sampleData) / 2);
memcpy(pSrcBuffer2,
       &(sampleData[sizeof(sampleData) / 2]),
       sizeof(sampleData) / 2);

/* Build source bufferList */
pBufferListSrc->numBuffers = 2;
pBufferListSrc->flatBuffers[0].dataLenInBytes = sizeof(sampleData) /
2;
pBufferListSrc->flatBuffers[0].bufferPhysAddr =
    sampleVirtToPhys(pSrcBuffer);
pBufferListSrc->flatBuffers[1].dataLenInBytes = sizeof(sampleData) /
2;
pBufferListSrc->flatBuffers[1].bufferPhysAddr =
    sampleVirtToPhys(pSrcBuffer2);

```

The operational data in this case is:

Listing 76. Compression Data Plane Operational Data

```

/* Allocate memory for operational data. Note this needs to be
 * 8-byte aligned, contiguous, resident in DMA-accessible
 * memory.
 */
status = PHYS_CONTIG_ALLOC_ALIGNED(&pOpData, sizeof(CpaDcDpOpData), 8);
}

if (CPA_STATUS_SUCCESS == status) {
    pOpData->bufferLenToCompress = sizeof(sampleData);
    pOpData->bufferLenForData = sizeof(sampleData);
    pOpData->dcInstance = dcInstHandle;
    pOpData->pSessionHandle = sessionHdl;
    pOpData->srcBuffer = sampleVirtToPhys(pBufferListSrc);
    pOpData->srcBufferLen = CPA_DP_BUFLIST;
    pOpData->destBuffer = sampleVirtToPhys(pBufferListDst);
    pOpData->destBufferLen = CPA_DP_BUFLIST;
    pOpData->sessDirection = CPA_DC_DIR_COMPRESS;
    pOpData->thisPhys = sampleVirtToPhys(pOpData);
    pOpData->pCallbackTag = (void *)0;
}

```

This request is then enqueued and submitted on the instance.

Listing 77. Data Plane Enqueue and Submit

```
status = cpaDcDpEnqueueOp(pOpData, CPA_TRUE);
```

- After possibly doing other work (e.g., enqueueing and submitting more requests), the Application can poll for responses that invoke the callback function registered

with the instance. Refer to [Table 1](#) Implementation Specific Documentation on the implementations polling functions.

- Once all requests associated with a session have been completed, the session can be removed.

Listing 78. Data Plane Remove Compression Session

Clean up by freeing up memory, stopping the instance, etc. using this command:

```
sessionStatus = cpaDcDpRemoveSession(dcInstHandle, sessionHdl);
```

4.5 Sample - Chained Hash and Stateless Compression

This example demonstrates the use of the asynchronous API, specifically, using the data compression chain API to perform chained hash and stateless compression operations. It performs a sha256 hash on the sample text and then compresses the sample text through a stateless session using the deflate compress algorithm with static Huffman trees.

These samples are located in `/dc/chaining_sample`

Listing 79. Querying and Starting a Compression Instance

```
/*
 * In this simplified version of instance discovery, we discover
 * exactly one instance of a data compression service.
 */
sampleDcGetInstance(&dcInstHandle);
if (dcInstHandle == NULL)
{
    PRINT_ERR("Get instance failed\n");
    return CPA_STATUS_FAIL;
}

/* Query Capabilities */
PRINT_DBG("cpaDcQueryCapabilities\n");
status = cpaDcQueryCapabilities(dcInstHandle, &cap);
if (status != CPA_STATUS_SUCCESS)
{
    PRINT_ERR("Query capabilities failed\n");
    return status;
}

if (CPA_FALSE == CPA_BITMAP_BIT_TEST(cap.dcChainCapInfo,
                                     CPA_DC_CHAIN_HASH_THEN_COMPRESS))
{
    PRINT_ERR(
        "Hash + compress chained operation is not supported on logical "
```

```

        "instance.\n");
    PRINT_ERR("Please ensure StorageEnabled=1 in the device configuration
"
        "file.\n");
    return CPA_STATUS_FAIL;
}

if (!cap.statelessDeflateCompression || !cap.checksumCRC32 ||
    !cap.checksumAdler32)
{
    PRINT_ERR("Error: Unsupported functionality\n");
    return CPA_STATUS_FAIL;
}

if (CPA_STATUS_SUCCESS == status)
{
    /* Set the address translation function for the instance */
    status = cpaDcSetAddressTranslation(dcInstHandle, sampleVirtToPhys);
}

if (CPA_STATUS_SUCCESS == status)
{
    /* Start static data compression component */
    PRINT_DBG("cpaDcStartInstance\n");
    status = cpaDcStartInstance(dcInstHandle, 0, NULL);
}

```

Listing 80. Create and Initialize Session Hash and Compression

```

if (CPA_STATUS_SUCCESS == status)
{
    /*
     * If the instance is polled start the polling thread. Note that
     * how the polling is done is implementation-dependent.
     */
    sampleDcStartPolling(dcInstHandle);
    /*
     * We now populate the fields of the session operational data and
create
     * the session. Note that the size required to store a session is
     * implementation-dependent, so we query the API first to determine
how
     * much memory to allocate, and then allocate that memory.
     */

    //<snippet name="initSession">

```

```

/* Initialize compression session data */
dcSessionData.compLevel = CPA_DC_L1;
dcSessionData.compType = CPA_DC_DEFLATE;
dcSessionData.huffType = CPA_DC_HT_STATIC;
dcSessionData.autoSelectBestHuffmanTree = CPA_FALSE;
dcSessionData.sessDirection = CPA_DC_DIR_COMPRESS;
dcSessionData.sessState = CPA_DC_STATELESS;
dcSessionData.checksum = CPA_DC_CRC32;

/* Initialize crypto session data */
cySessionData.sessionPriority = CPA_CY_PRIORITY_NORMAL;
/* Hash operation on the source data */
cySessionData.symOperation = CPA_CY_SYM_OP_HASH;
cySessionData.hashSetupData.hashAlgorithm = CPA_CY_SYM_HASH_SHA256;
cySessionData.hashSetupData.hashMode = CPA_CY_SYM_HASH_MODE_PLAIN;
cySessionData.hashSetupData.digestResultLenInBytes =

GET_HASH_DIGEST_LENGTH(cySessionData.hashSetupData.hashAlgorithm);
/* Place the digest result in a buffer unrelated to srcBuffer */
cySessionData.digestIsAppended = CPA_FALSE;
/* Generate the digest */
cySessionData.verifyDigest = CPA_FALSE;

/* Initialize chaining session data - hash + compression
 * chain operation */
chainSessionData[0].sessType = CPA_DC_CHAIN_SYMMETRIC_CRYPT0;
chainSessionData[0].pCySetupData = &cySessionData;
chainSessionData[1].sessType = CPA_DC_CHAIN_COMPRESS_DECOMPRESS;
chainSessionData[1].pDcSetupData = &dcSessionData;

/* Determine size of session context to allocate */
PRINT_DBG("cpaDcChainGetSessionSize\n");
status = cpaDcChainGetSessionSize(dcInstHandle,
                                  CPA_DC_CHAIN_HASH_THEN_COMPRESS,
                                  NUM_SESSIONS_TWO,
                                  chainSessionData,
                                  &sess_size);
}

if (CPA_STATUS_SUCCESS == status)
{
    /* Allocate session memory */
    status = PHYS_CONTIG_ALLOC(&sessionHdl, sess_size);
}

/* Initialize the chaining session */

```



```

if (CPA_STATUS_SUCCESS == status)
{
    PRINT_DBG("cpaDcChainInitSession\n");
    status = cpaDcChainInitSession(dcInstHandle,
                                   sessionHdl,
                                   CPA_DC_CHAIN_HASH_THEN_COMPRESS,
                                   NUM_SESSIONS_TWO,
                                   chainSessionData,
                                   dcCallback);
}

```

Note: `cySessionData.digestIsAppended` should be always set to `CPA_FALSE` as the digest must not appended in the end of output.

Listing 81. Chained Hash and Stateless Compression Memory Allocation

```

status = cpaDcBufferListGetMetaSize(dcInstHandle, numBuffers,
&bufferMetaSize);
if (CPA_STATUS_SUCCESS != status)
{
    PRINT_ERR("Error get meta size\n");
    return CPA_STATUS_FAIL;
}
bufferSize = sampleDataSize;
if (CPA_STATUS_SUCCESS == status)
{
    status = dcChainBuildBufferList(&pBufferListSrc, numBuffers,
bufferSize, bufferMetaSize);
}

/* copy source data into buffer */
if (CPA_STATUS_SUCCESS == status)
{
    pFlatBuffer = (CpaFlatBuffer *) (pBufferListSrc + 1);
    memcpy(pFlatBuffer->pData, sampleData, bufferSize);
}

/* Allocate destination buffer the four times as source buffer */
if (CPA_STATUS_SUCCESS == status)
{
    status = dcChainBuildBufferList(&pBufferListDst, numBuffers, 4 *
bufferSize, bufferMetaSize);
}

/* Allocate digest result buffer to store hash value */
if (CPA_STATUS_SUCCESS == status)
{

```

```

    status = PHYS_CONTIG_ALLOC(&pDigestBuffer,
GET_HASH_DIGEST_LENGTH(hashAlg));
}

```

Listing 82. Set Up Operational Data Hash and Compression

```

dcOpData.flushFlag = CPA_DC_FLUSH_FINAL;
dcOpData.compressAndVerify = CPA_TRUE;
dcOpData.compressAndVerifyAndRecover = CPA_TRUE;

cySymOpData.packetType = CPA_CY_SYM_PACKET_TYPE_FULL;
cySymOpData.hashStartSrcOffsetInBytes = 0;
cySymOpData.messageLenToHashInBytes = bufferSize;
cySymOpData.pDigestResult = pDigestBuffer;

/* Set chaining operation data */
chainOpData[0].opType = CPA_DC_CHAIN_SYMMETRIC_CRYPTO;
chainOpData[0].pCySymOp = &cySymOpData;
chainOpData[1].opType = CPA_DC_CHAIN_COMPRESS_DECOMPRESS;
chainOpData[1].pDcOp = &dcOpData;

```

Hash and stateless dynamic compression are also supported. Refer to [Listing 64](#) and [Listing 65](#) to add dynamic compression released buffers and session data.

Note: Hash algorithms are not limited to sha1 and sha256. Refer to Intel® QuickAssist Technology Software for Linux* Release Notes ([Table 1](#)) for any limitations on using other hash algorithms in the current release.

Listing 83. Verify the Output of Chained Hash and Stateless Compression

```

/* Use software to calculate digest and verify digest */
if (CPA_STATUS_SUCCESS == status)
{
    PHYS_CONTIG_ALLOC(&pSWDigestBuffer, GET_HASH_DIGEST_LENGTH(hashAlg));
    status = calSWDigest(sampleData,
        bufferSize,
        pSWDigestBuffer,
        GET_HASH_DIGEST_LENGTH(hashAlg),
        hashAlg);

    if (CPA_STATUS_SUCCESS == status)
    {
        if (memcmp(pDigestBuffer,
            pSWDigestBuffer,
            GET_HASH_DIGEST_LENGTH(hashAlg)))
        {
            status = CPA_STATUS_FAIL;
            PRINT_ERR("Digest buffer does not match expected output\n");
        }
    }
}

```

```
    }
    else
    {
        PRINT_DBG("Digest buffer matches expected output\n");
    }
}

PHYS_CONTIG_FREE(pSWDigestBuffer);
}

/* Use zlib to decompress and verify integrity */
//<snippet name="software decompress">
if (CPA_STATUS_SUCCESS == status)
{
    struct z_stream_s stream = {0};
    Cpa8U *pDecompBuffer = NULL;
    Cpa8U *pHWCompBuffer = NULL;
    Cpa8U *pSWCompBuffer = NULL;
    Cpa32U bufferLength = 0;

    status = inflate_init(&stream);
    if (CPA_STATUS_SUCCESS != status)
    {
        PRINT("zlib stream initialize failed");
    }

    bufferLength = pBufferListSrc->numBuffers *
        pBufferListSrc->pBuffers->dataLenInBytes;

    if (CPA_STATUS_SUCCESS == status)
    {
        status = PHYS_CONTIG_ALLOC(&pDecompBuffer, bufferLength);
    }

    if (CPA_STATUS_SUCCESS == status)
    {
        status = PHYS_CONTIG_ALLOC(&pHWCompBuffer, bufferLength);
    }

    if (CPA_STATUS_SUCCESS == status)
    {
        status = PHYS_CONTIG_ALLOC(&pSWCompBuffer, bufferLength);
    }

    if (CPA_STATUS_SUCCESS == status)
    {

```

```
        copyMultiFlatBufferToBuffer(pBufferListDst, pHWCompBuffer);
    }

    if (CPA_STATUS_SUCCESS == status)
    {
        status = inflate_decompress(&stream,
                                    pHWCompBuffer,
                                    bufferLength,
                                    pDecompBuffer,
                                    bufferLength);

        if (CPA_STATUS_SUCCESS != status)
        {
            PRINT_ERR("Decompress data on zlib stream failed\n");
        }
    }

    if (CPA_STATUS_SUCCESS == status)
    {
        /* Compare with original Src buffer */
        if (memcmp(pDecompBuffer, sampleData, bufferSize))
        {
            status = CPA_STATUS_FAIL;
            PRINT_ERR("Decompression does not match source buffer\n");
        }
        else
        {
            PRINT_DBG("Decompression matches source buffer\n");
        }
    }

    inflate_destroy(&stream);

    PHYS_CONTIG_FREE(pSWCompBuffer);
    PHYS_CONTIG_FREE(pHWCompBuffer);
    PHYS_CONTIG_FREE(pDecompBuffer);
}
```

Listing 84. Remove Chained Hash and Stateless Compression Session

```
status = cpaDcChainRemoveSession(dcInstHandle, sessionHdl);
```

§