White Paper
**Intel® QuickAssist Technology**
Big Data and Cloud

# Improving Performance and Security of Big Data and Cloud Solutions

## Using Intel® QuickAssist Technology to Accelerate Public Key Cryptography and Compression

As the complexity of networking and security applications continues to grow, systems need more and more computational resources for demanding tasks. Compute-intensive tasks, including cryptography and data compression, use algorithms that can be offloaded from the server processor to accelerators, significantly increasing the throughput of these operations. Intel® QuickAssist Technology is designed to offload security and compression algorithms to allow developers to easily incorporate acceleration into their software.

Intel QuickAssist Technology enables developers to create software solutions that leverage encryption/decryption and compression/decompression acceleration, accessing the technology through APIs in the Intel® QuickAssist Software. Enterprises and cloud service providers can add these solutions to their standard server infrastructure, using add-in accelerator cards, by developing with Intel QuickAssist Technology built into the standard server chipsets. This paper provides developers with information on Intel QuickAssist Technology and presents some key use cases to provide background for them to understand how they can take advantage of the security benefits and performance improvements available with Intel QuickAssist Technology in their solutions.

## Table of Contents

## Meeting the Growing Demand for Encryption and Compression with Intel QuickAssist Technology

Enterprise and cloud data centers are under pressure to continuously expand revenue-generating and value-added services, such as compute intensive and I/O-demanding Big Data solutions, which moves large amounts of data into and out of storage, and sends it across the networked clusters. Compressing these data streams can vastly improve performance by reducing the size of data in motion. At the same time, data protection and secure transmissions remain a key challenge, forcing many of these same environments to adopt a "Secure Sockets Layer (SSL) everywhere" strategy and implement ubiquitous security with strong ciphers, authentication and a Public Key Infrastructure for all transactions.

Compression/decompression and encryption/decryption with public key generation and handshaking are compute-intensive tasks, adding to server processor loads, which SSL can substantially degrade server performance. These operations are ideal candidates for offloading to hardware accelerators in order to reduce the demands on the server processors by requiring significant core resources. Intel architects have shown that offloading such tasks to hardware-based accelerators using Intel QuickAssist technology can significantly reduce Apache Hadoop* MapReduce solution times by as much as 40 percent[1,2] and speed up SSL public key encryption and decryption operations over 27X.[1,2]

This paper provides developers with the following information that will assist them in navigating the performance options available to them as they develop their enterprise, cloud, and Big Data solutions:

- An introduction to the Intel QuickAssist Technology.
- Two key use cases in enterprise and cloud data centers, where Intel QuickAssist Technology offers clear benefits.
- How to implement these use cases in system architectures.
- Additional potential uses for Intel QuickAssist Technology acceleration.

### Intel QuickAssist Technology: Integrated Acceleration

Intel QuickAssist Technology integrates hardware-acceleration in Intel Chipsets for widely used, compute-intensive algorithms, including key exchanges, data compression, SSL Bulk Encryption Algorithms, and SSL ciphers for Intel® chipsets. Hardware vendors offer Intel QuickAssist Technology integrated in chipsets for server motherboard or into accelerator cards, such as the Intel® QuickAssist Adapter 8950/55, which can be installed into a standard server as a PCIe* Plug In Card. Developer software solutions access the technology through APIs and shims provided in the Intel QuickAssist Software.

### The Intel QuickAssist Software Framework

Figure 1 illustrates the Intel QuickAssist Software framework with device drivers, libraries, and APIs to take advantage of Intel QuickAssist Technology. The Intel QuickAssist Technology shims and patches in the Intel QuickAssist Software (such as a zlib patch for offloading compression) offer the best performance by offloading these compute-intensive workloads from the CPU core to the accelerator. Intel contributions to open source software, such as asynchronous OpenSSL available from the OpenSSL repository, provide the broadest applicability across heterogeneous environments and systems. Together, the software helps accelerate development of innovative applications for offloading operations across a wide variety of data center applications.

## Intel QuickAssist Technology Use Cases

Intel QuickAssist Technology has been shown to speed up SSL public key encryption and decryption operations over 27X[1,2] and to reduce Apache Hadoop* MapReduce solution times by as much as 40 percent.[1,2] The following use cases illustrate two examples where Intel QuickAssist Technology has proven a significant benefit in server throughput.

### Use Case 1: Dynamic Web SSL Encryption Offload Boosts Throughput

As more and more information is encrypted across networks using OpenSSL*[3], the workloads in cloud data centers and enterprises demand greater compute power. That's because the SSL and TLS protocols use several security/cryptographic constructs. First, during session initialization, the session keys are exchanged as part of a "handshake" process using public key cryptography to provide confidentiality and authentication. This process can be an expensive operation, which can consume millions of processor core cycles. Subsequently, symmetric cryptography using these session keys is used to provide the cipher algorithms encrypting the data being exchanged between the client and the server. Intel® Data Protection Technology with Advanced Encryption Standard New Instructions available in today's Intel® processors significantly speeds up symmetric cryptography, however, processing public key cryptography can considerably impact server performance. Hardware-accelerating SSL with Intel QuickAssist Technology can improve server throughput by freeing significant processor cycles spent on exponential calculations for key decryption for other useful computations.

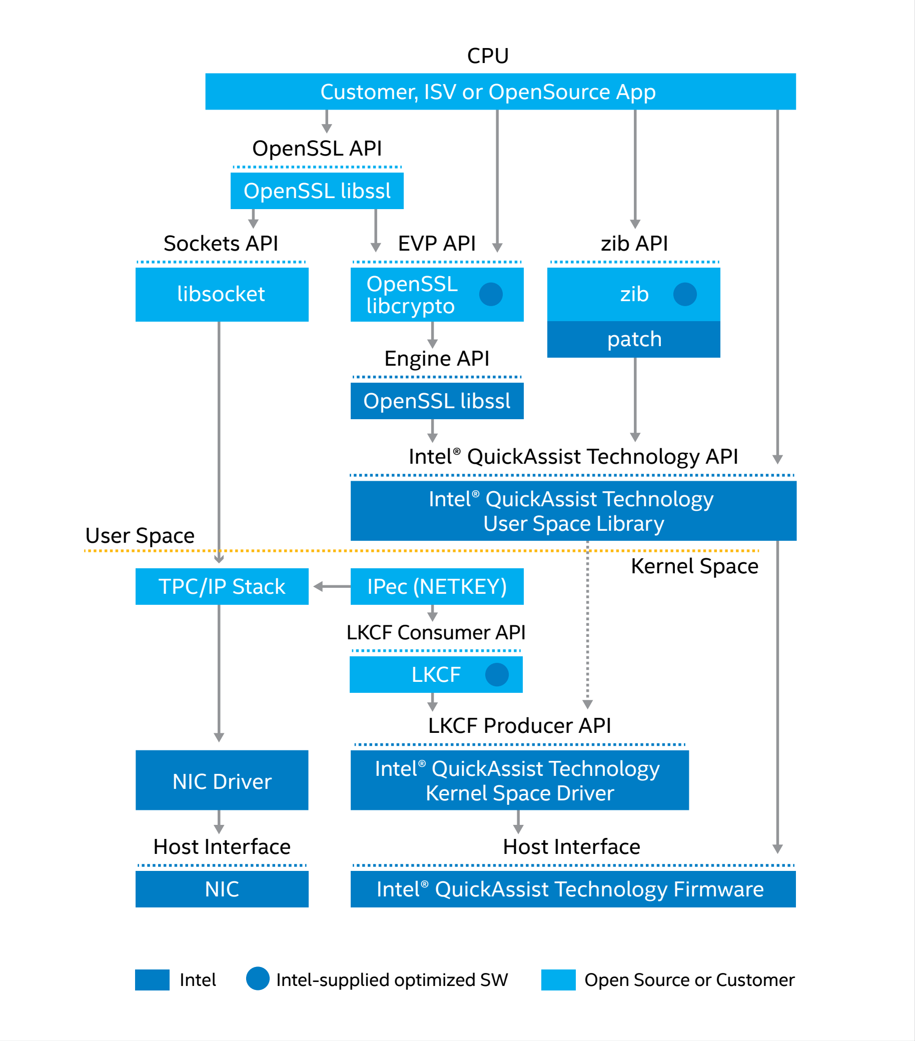

**Figure 1.** Intel® QuickAssist Software stack.

### Synchronous and Asynchronous OpenSSL*

Servers generally handle default OpenSSL operations synchronously. Synchronous architectures imply that the application thread will be blocked until the cryptography operation finishes, requiring the OS to schedule a different application thread on the CPU to continue to get useful work done. The resulting extra context switches waste processing cycles and introduce negative caching effects, seriously impacting the potential performance of the server. With synchronous OpenSSL, launching additional application processes, or multi-threading the code, helps increase encryption and decryption performance, but this also incurs the additional overhead and impact of more resource utilization these threads require.

Asynchronous OpenSSL is a non-blocking approach to encryption/decryption and writing to sockets that allows a single software thread to make multiple simultaneously outstanding cryptography requests. This greatly increases the ability of a single application thread to attain high encrypt/decrypt performance while minimizing the additional resources to do so.

Asynchronous OpenSSL enables improved SSL throughput, freeing additional processor cycles to run other tasks. This operation also enables single-threaded applications to efficiently handle multiple SSL connections, because individual flows are not blocked when using hardware acceleration.

### Accelerating OpenSSL Transactions with Intel QuickAssist Technology

Asynchronous OpenSSL running on Intel QuickAssist Technology significantly improves server throughput over the standard (synchronous) release. Figure 2 compares OpenSSL performance for the two versions, with asynchronous OpenSSL offloaded to Intel QuickAssist Technology. These measurements are indicative of the throughput benefits during the bulk data transfer phase.

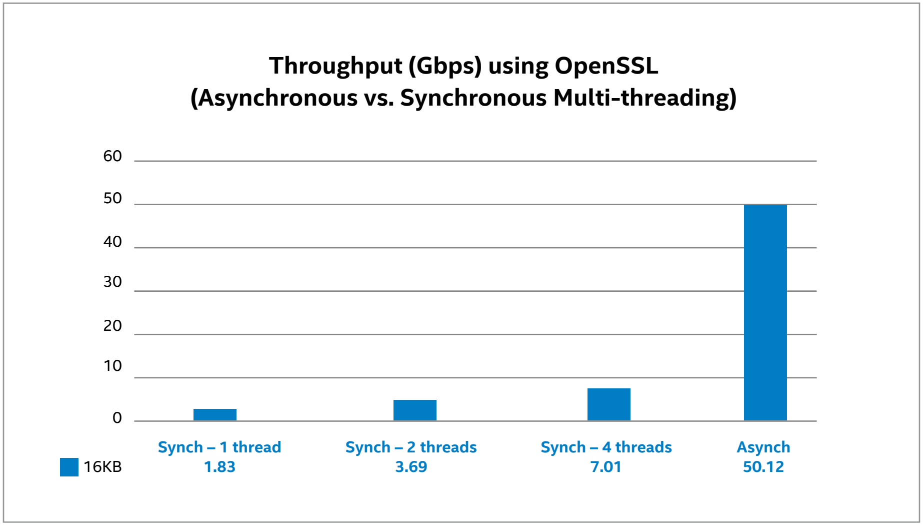

**Figure 2.** OpenSSL* throughput: asynchronous vs. synchronous.[1,2]

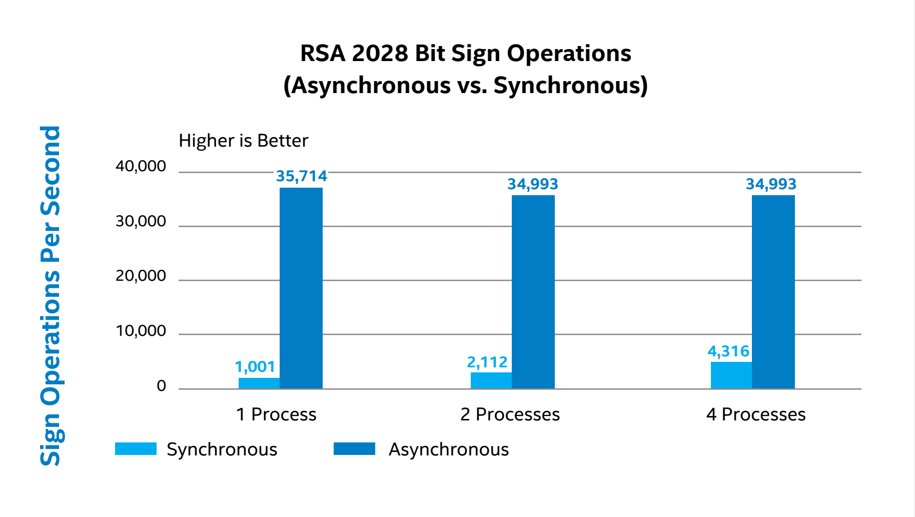

**Figure 3.** RSA 2028 Bit Sign Operations: Asynchronous Versus Synchronous.[1,2]

Asynchronous OpenSSL achieves higher throughput from Intel Quick Assist Technology using a single thread of execution—as much as 27X—compared to synchronous OpenSSL with multiple threads running. To achieve the same throughput from synchronous OpenSSL requires multi-threading, which is more expensive in terms of CPU utilization due to the need for context switching.

Figure 3 shows the number of RSA 2048-bit sign operations per second for asynchronous and synchronous OpenSSL. This operation is used during the initial session initiation and handshake phase, which is very compute-intensive. As discovered in bulk encryption, asymmetric cryptography produces full throughput with only one process, whereas the synchronous version delivers significantly less throughput with four processes. In fact, asynchronous OpenSSL is extracting the maximum performance from the accelerator used in the testing. Consequently, synchronous OpenSSL generally will require significantly more threads than asynchronous OpenSSL to achieve the same performance. This negatively affects CPU utilization because of context switching and other impacts from additional threads using caching resources.

### Adding Intel QuickAssist Technology for SSL Transactions

Asynchronous OpenSSL was developed by the OpenSSL community with contributions from Intel architects. Asynchronous OpenSSL for Intel QuickAssist Technology is available in the OpenSSL stack development branch, apart from the Generally Available (GA) release, at www.openssl.org. Compile and install instructions are provided with the source in the repository.

To gain the benefits of software supporting asynchronous OpenSSL, developers need to add asynchronous SSL to their solutions. IT departments can then add Intel QuickAssist Technology-based accelerator cards to existing Intel® Xeon® processor-based servers, or upgrade to server platforms with Intel chipsets that integrate the technology.

### Intel QuickAssist Technology for OpenSSL Summary

For enterprises and cloud data centers supporting massive SSL transactions, using asynchronous OpenSSL to offload symmetric and asymmetric encryption onto Intel QuickAssist Technology-enabled platforms and plug-in PCie-based accelerator cards can significantly improve performance to customers while reducing TCO.

## Use Case 2: Hadoop* Compression Offload Significantly Improves Big Data Application Performance

Hadoop jobs move massive amounts of data in and out of the Hadoop Distributed File System (HDFS) and across the network. Hadoop I/O results in considerable disk operations and network utilization, such that Hadoop MapReduce jobs are often I/O bound.

### Software-based Compression/ Decompression is Not Enough

Compressing data before writing to the HDFS and during data transmissions can make more efficient use of storage capacity and accelerate data transfers. Figure 4 illustrates an example where Hadoop MapReduce jobs can benefit from data compression. However, processor-based, software-only compression and decompression may take significant CPU cycles as illustrated in Figure 5. In this TextSort job, zlib compression/decompression together consume about 30 percent CPU time on average.[1,2]

### Intel® QuickAssist Technology Offload Improves Hadoop Performance

Intel QuickAssist Technology enables considerable speedup for Hadoop. Intel testing shows that Intel QuickAssist Technology results in MapReduce finishing about 40 percent sooner than software-only zlib compression, while reducing average CPU utilization by about 20 percent.[1,2] Intel QuickAssist Technology makes extra CPU cycles available for other Hadoop tasks, such as Sort and Merge.

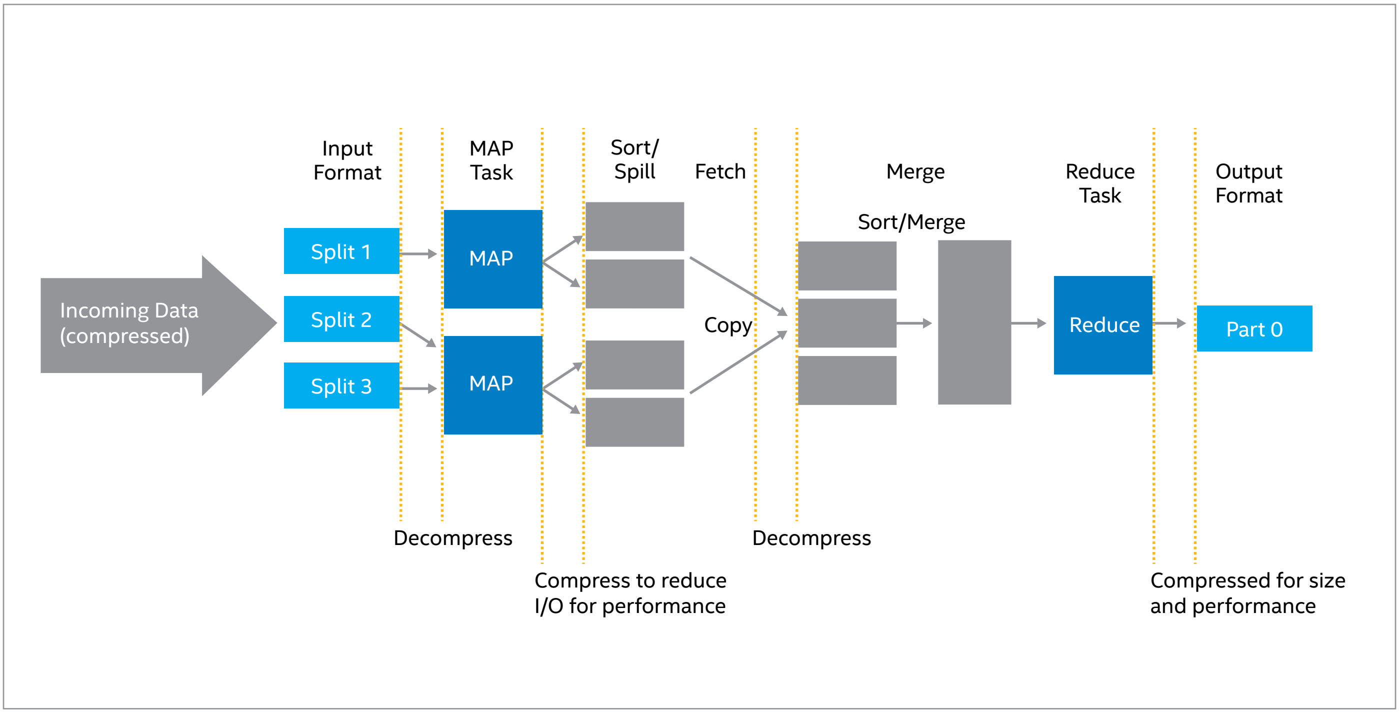

**Figure 4.** Ideal locations in the Hadoop job cycle for compress/decompress.

Furthermore, Intel QuickAssist Technology results in twice the throughput for disk reads and data transmission[1,2] (Figure 6). Disk reads and network throughput are limited by how fast data can be decompressed. Software decompression utilizes significant processor resources, delaying further activities. By offloading decompression to Intel QuickAssist Technology, decompression is no longer the bottleneck. Thus, increasing decompression throughput indirectly increases disk reads and network throughput.[1,2]

### Adding Intel QuickAssist Technology for Big Data Applications

Intel has made the integration of Intel QuickAssist Technology into Hadoop simple. When Hadoop calls the zlib API, an Intel-supplied zlib patch allows the zlib library to interface to the Intel QuickAssist Technology API. If the Intel technology is present in the server, the modified zlib library forwards the compression requests to the Intel QuickAssist Technology compression accelerators, seamlessly integrating hardware-based acceleration into the Hadoop job.

### Intel QuickAssist Technology for Big Data Summary

Compression on Hadoop workloads reduces the size of data on disk and accelerates transfers through the network, but this compression may take significant CPU resources. Offloading compression and decompression to Intel QuickAssist Technology significantly improves Hadoop performance across the CPU, disk reads, and network throughput. Intel makes it simple for developers to support hardware-based Intel QuickAssist Technology acceleration in their applications, reducing the cost and complexity for data centers to add enhanced Hadoop performance to big data services.

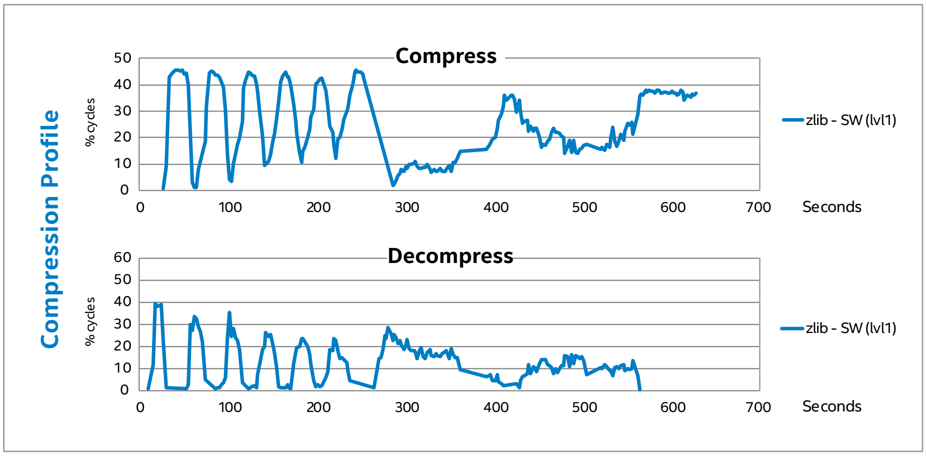

**Figure 5.** CPU Utilization for software-based compression/decompression during Textsort.[1,2]

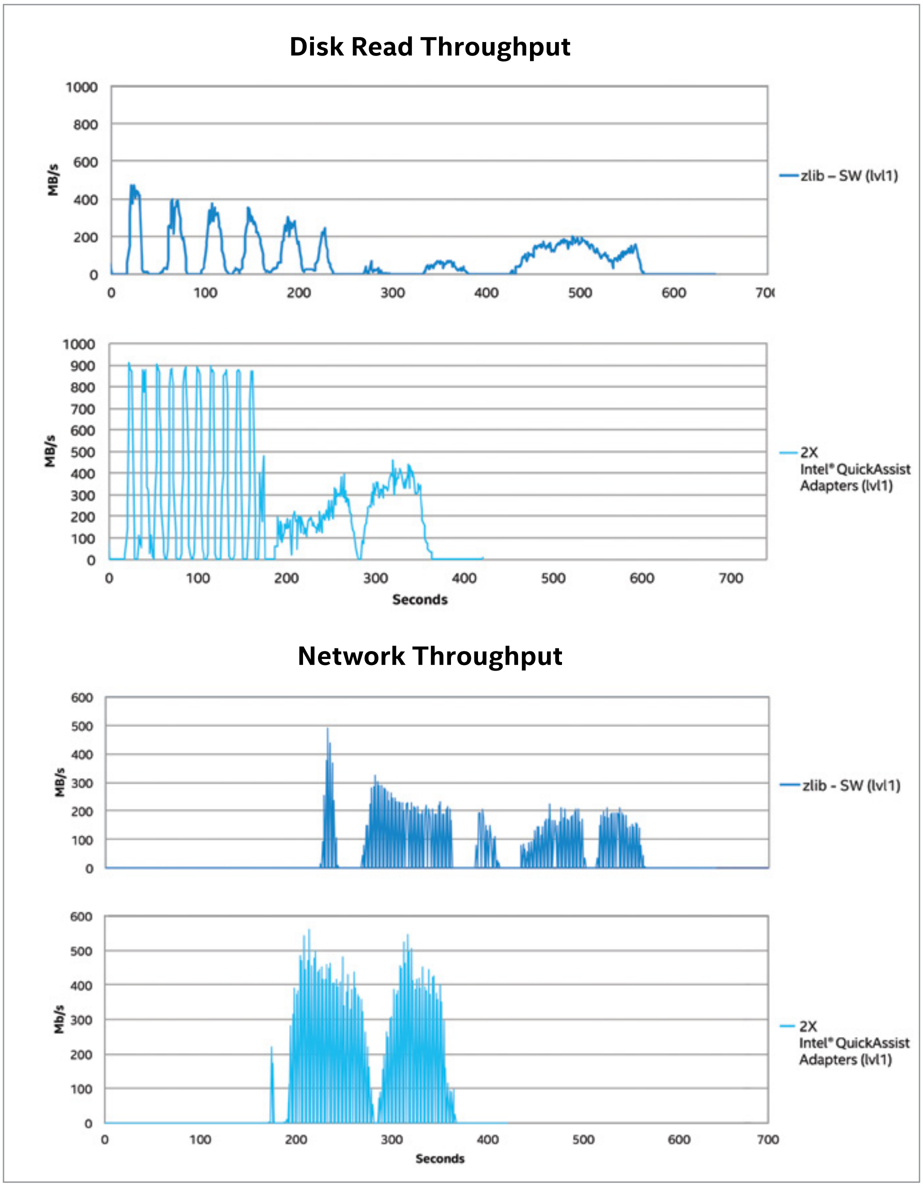

**Figure 6.** Intel® QuickAssist Technology doubles throughput for Hadoop disk reads and network.

## Additional Offload Opportunities for the Cloud Data Center

The improvements provided by offloading encryption and compression can potentially be extended to other applications, which can enhance cloud data center operations and benefit cloud customers. Developers might consider assessing these use cases.

**Hadoop Encryption/Decryption –** With Hadoop being a distributed system, data is expected to traverse servers in the data center, and potentially disparate data centers around the globe. Protecting data in transit requires encryption on top of compression before sending across a possibly unsecure network connection. As seen with asynchronous OpenSSL, offloading encryption and decryption to Intel QuickAssist Technology can deliver improvements in SSL throughput. On large data transfers, this can offer a significant benefit in run times.

**Virtual Machine (VM) Migration –** In the virtualized cloud data center, VM migration can happen at any time, moving a VM image from one server to another within—or even outside—the physical premise. To optimize the transition across platforms, compression can reduce the image size and thus accelerate the transfer over the network. With compression offloaded to Intel QuickAssist Technology, transfer times are shortened, accelerating the migration and launch of the environment on another platform.

## Summary

Intel QuickAssist Technology-enabled accelerators give developers new opportunities to create value-added applications for enterprise and cloud data centers that can improve data movement and cryptography performance. Current Intel developments with both asynchronous OpenSSL and Big Data shows marked benefits when Intel QuickAssist Technology is added to servers. Many further opportunities exist as Cryptography and security solutions become a mandatory baseline for data center, enterprise, and cloud applications.

Intel simplifies the process of adding Intel QuickAssist Technology offload to applications by providing device drivers, libraries, and APIs in the Intel QuickAssist Technology Software, and through work with open source software, like Asynchronous OpenSSL. See the list of Additional Resources in the Appendix for more information.

## Appendix

### Additional Resources

For further information about Intel QuickAssist Technology, visit

http://www.intel.com/content/www/us/en/io/quickassist-technology/quickassist-technology-developer.html

Intel® QuickAssist Technology Accelerates Hadoop* Applications

Intel® QuickAssist Technology Accelerates OpenSSL* Applications

Intel® QuickAssist Technology | 01.org

IPSec Performance Demonstration

Scaling Acceleration Capacity with Intel® QuickAssist Technology

Intel® QuickAssist Adapter Family for Servers

Intel® QuickAssist Adapter 8950 Product Brief

Intel® QuickAssist Technology Performance Optimization Guide

Intel® QuickAssist Technology API: Programmer's Guide

Intel® QuickAssist Technology Cryptographic API: Reference Guide

Intel® QuickAssist Technology Compression API: Reference Guide

### Test Configurations

**Hadoop Compression Testing**

Hadoop testing was completed on a highly tuned and optimized Hadoop cluster with a single name node and two data nodes.

Data Nodes - Two

| Processors | 2x Intel® Xeon® processor E5-2680 @ 2.70 GHz |
|---|---|
| Cores/Threads | 8 cores/socket/Intel® Hyper-Threading technology[4] enabled for 16 hardware threads |
| Chipset | Intel® Communications Chipset 8950 with Intel® QuickAssist technology |
| Memory | 128 GB of DDR3 @ 1600 MHz |
| Network | Intel® 82598EB 10-Gigabit |
| Storage | 1 x 32 GB Intel X25-E Extreme - Operating System<br>8 x 300 GB Intel 710 Series 3Gb/s SATA<br>24 x 32 GB Intel X25-E Extreme – external SAS with a 6Gb/s link |
| Operating System | CentOS* 6.3_64<br>kernel 2.6.32-279.19.1.el6.x86_64 |
| Java Hotspot* | 1.7.0_13 64-Bit Server VM (build 23.7-b01) |
| Apache Hadoop* | 1.0.4 (with several performance fixes) |
| Compression | zlib 1.2.7 |

## Appendix Continued

### *Methodology*

Focus was on I/O-intensive symmetric workloads (data in=data out):

- Terasort: synthetic data; highly compressible (1 TB->160 GB)
- Sort: English dictionary; highly compressible (1 TB->225 GB)

Compression algorithms used:

- Native GZIP - compression level 1 (least compression)
- QuickAssist GZIP - compression level 1

Reducer process start was delayed until all maps completed.

### *Hadoop Settings*

The tests were completed on a highly tuned Hadoop cluster, optimized for performance. Some notable settings include the following:

- Oversubscribe map/reduce slots (40/24)
- Adjust the ratio of reduce slots to shuffle copy threads (concurrent copies=reducers x shuffle threads)
- Two waves of reducers to merge maps output in memory
- HDFS and temp directories not co-located on same disks
- Java* Virtual Machine (JVM) reuse turned on
- Tuned the JVM to avoid stop-the-world garbage collection
- HDFS duplication turned off

### SSL Platform Testing

| Processor | Dual socket Intel® Xeon® processor E5-2680 v2 |
|---|---|
| Chipset | Intel® Communications Chipset 8950 (Intel® DH8920 PCH) with Intel® QuickAssist technology |
| Memory | DDR3 1333 MHz |
| Operating System | Linux* version 3.1.0-7.fc16.x86_64 |
| Compiler | gcc version 4.6.2 20111027 (Red Hat 4.6.2-1) (GCC) |

Results have been estimated based on internal Intel analysis and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance.

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

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

A "Mission Critical Application" is any application in which failure of the Intel Product could result, directly or indirectly, in personal injury or death. SHOULD YOU PURCHASE OR USE INTEL'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD INTEL AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS COSTS, DAMAGES, AND EXPENSES AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT INTEL OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE INTEL PRODUCT OR ANY OF ITS PARTS.

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

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

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

[1] Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel® products as measured by those tests. Any difference in system hardware or software design or configuration, as well as system use patterns including wireless connectivity, may affect actual test results and ratings.

[2] Configurations for all tests are presented in the Appendix A.

[3] OpenSSL: http://www.openssl.org

[4] Intel Hyper Threading Technology available on select Intel® Core™ processors. Requires an Intel® HT Technology-enabled system. Consult your PC manufacturer. Performance will vary depending on the specific hardware and software used. For more information including details on which processors support HT Technology, visit http://www.intel.com/info/hyperthreading.

1114/RS/HBD/PDF 331416-001US