WHITE PAPER

Communications Service Providers
Network Security

SECUI Tests BLUEMAX NGF 5000 IPsec VPN on Intel® Platform

SECUI BLUEMAX NGF 5000 utilizes unique architecture for multi-processor encryption processing based on Intel® QuickAssist Technology (Intel® QAT) and Intel® Ethernet Flow Director. Results of tests show significant performance increases when Intel QAT is used compared to CPU-based encryption processing.[1]

SECUI

Table of Contents

Overview

Firewalls are a proven security foundation for protecting information assets. With evolving network infrastructure, more network layers, and increased data security threats, firewalls have evolved to become next-generation firewalls (NGFs), which include a broader array of security functions to protect against more elaborate attacks.

NGFs include traditional packet filtering-based firewall technologies, while integrating a number of new and related security functions such as application firewalls that use network and port address translation (NAT), virtual private network (VPN), deep packet inspection (DPI), intrusion prevention system (IPS) technologies, and cryptography support for encrypted traffic.

One way to satisfy increasing performance requirements for this added security functionality, and still use cost-effective general purpose processors, is to use multiple CPU cores and hardware accelerators. The challenge is to maintain the load balancing of encrypted packets when IP addresses cannot be used for this job because they are encrypted. SECUI has developed its BLUEMAX NGF 5000 with a unique design using Intel® Xeon® Scalable processors, Intel® QuickAssist Technology (Intel® QAT), and Intel® Converged Ethernet controllers to overcome these challenges and maintain high throughput. The encryption performance improvement using Intel QAT is up to triple the performance of a CPU-only solution, depending on packet size, and is demonstrated through testing conducted by SECUI and described in this paper.[1]

Introduction to the BLUEMAX NGF Series

The BLUEMAX NGF is a family of enterprise-wide security appliances (BLUEMAX 100, 200, 300, 500, 1000, 1500, 2000, 5000, and 20000) designed to detect and block a wide range of security threats. A virtual version of the software is available for cloud infrastructure applications. Key functions built into the BLUEMAX NGF products include both legacy firewall and NGF filtering, IPsec virtual private network (VPN), security automation, malware protection and device compliance, and threat intelligence services.

These technologies enable the following security benefits:

- App Control: Defines and analyzes applications in order to help prevent the introduction of vulnerabilities through foreign applications and malicious code distribution. A similar feature offers control of software as a service (SaaS) applications.
- User ID Recognition: Ensures the mobility of users by applying the same security policy regardless of when and where they connect to the network. Produces data on when and how users are connected for analysis.

- Device Control: Controls access to internal networks and essential business systems from users whose security settings and updates are not up to date to prevent malware infection. System can help to install required software, and update security status, and backup/encryption settings.
- Open API: System uses API to integrate with third-party vendors for security orchestration and automation.
- Domain Object: For cloud applications, system substitutes domain name for IP address as a firewall object. Systems are able to collect up to 2,048 objects per domain for full picture of cloud environment.
- SSL/TLS Inspection: Automatically detects secure sockets layer (SSL)/transport layer security (TLS) sessions, decrypts packets, and applies various next-generation network security functions to the decrypted packets.

Hardware Platform Utilizes Intel® Technology

The BLUEMAX NGF appliance platform is based on technology from Intel starting with the Intel Xeon Gold processors. Intel Xeon Gold processors are part of the Intel Xeon Scalable processor family that is designed for cloud-optimized networks. The CPUs feature an open architecture that scales and adapts to handle the demands of agile networks that can operate with cloud economics and be highly automated and responsive. In particular, Intel Xeon Gold CPUs are designed to offer workload-optimized performance for general-purpose compute delivering significant improvements for demanding storage and networking workloads. For the BLUEMAX NGF 5000 used in the tests for this paper, SECUI used servers with two 12-core processors per server to utilize 24 CPU cores.

To maximize encryption processing, certain models of the BLUEMAX NGF family also utilize Intel QAT. Intel QAT provides security and compression acceleration capabilities focused on compute-intensive operations such as symmetric cryptography functions, public key functions, and compression and decompression functions, including DEFLATE. The performance of the Intel QAT helps maximize cryptography throughput of the BLUEMAX NGF appliances in both IPsec VPN and SSL/TLS applications.

Ethernet connectivity for the NGF appliance is provided by the Intel® Ethernet Converged Network Adapter X710. This adapter has either two or four 10 GbE ports that offer unique features for improving server and network performance for virtualized networks. Particularly important for the BLUEMAX NGF 5000 is Intel® Ethernet Flow Director (Intel® Ethernet FD) for hardware-based application traffic steering. Intel Ethernet FD can classify received packets and then provide load balancing by directing those data flows to special queues to be matched with specific CPU cores. Similarly, the BLUEMAX NGF 5000 leverages receive-side scaling (RSS), which calculates hash values using information from L3 and L4 packet headers and distributes packets to the CPU core through these values. RSS prevents received packets from being sent to a single CPU core resulting in a network performance bottleneck.

Encryption Performance Challenges

SECUI's use of multiple general-purpose processors provides low costs and system agility, but matching the encryption performance of fixed function appliances proved to be challenging. This encryption performance is most critical with the IPsec VPN. VPNs create a virtual tunnel via encryption to more securely connect two or more networks via a public network such as the internet. VPNs have become more important as companies have shifted from dedicated networks (MPLS, for example) to the use of public networks like the internet.

Issues of Packet Distribution in VPN

IPsec VPNs create a virtual tunnel between two gateway devices (see Figure 1) and use the virtual tunnel to transport encrypted packets.

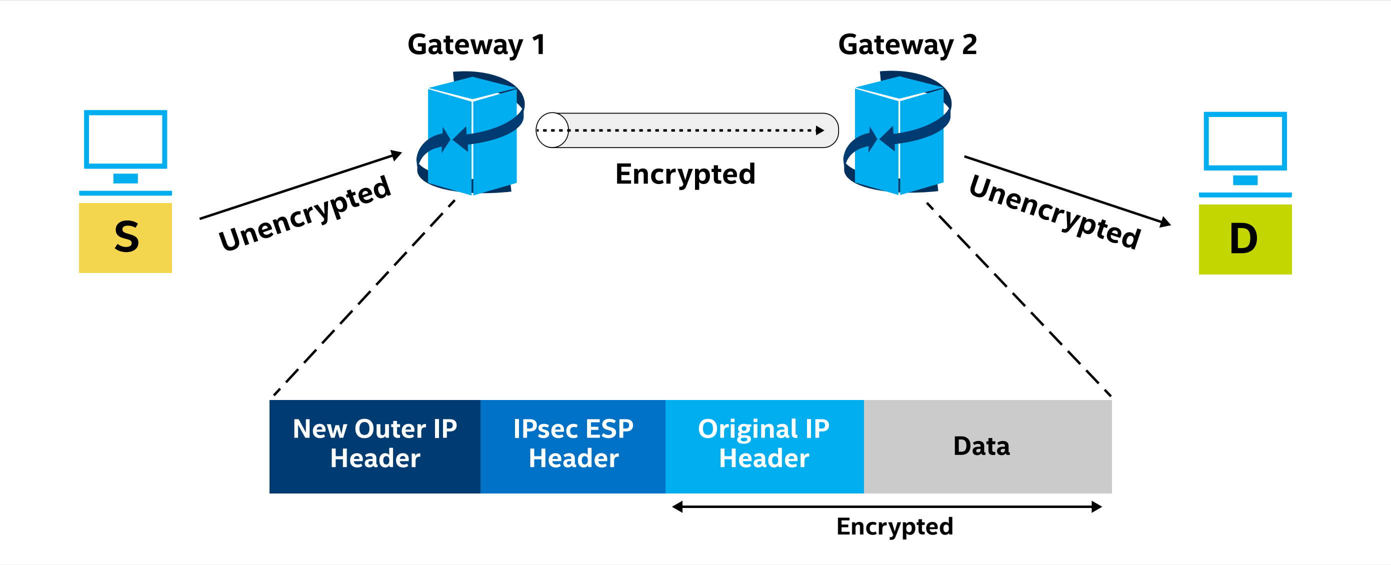

Figure 1. The components of IPsec VPN tunnel

In Figure 2, the IPsec encryption process is shown starting with a packet being sent from host S to host D. Gateway 1 encrypts the IP header and data of the original IP packet and then adds the encapsulating security payload (ESP) header and a new outer IP header before it sends the packet to gateway 2. After encryption, the address of gateway 1 and gateway 2 devices are given instead of the original source and destination address in the new outer IP header. Because of this, issues with ESP packet distribution can occur.

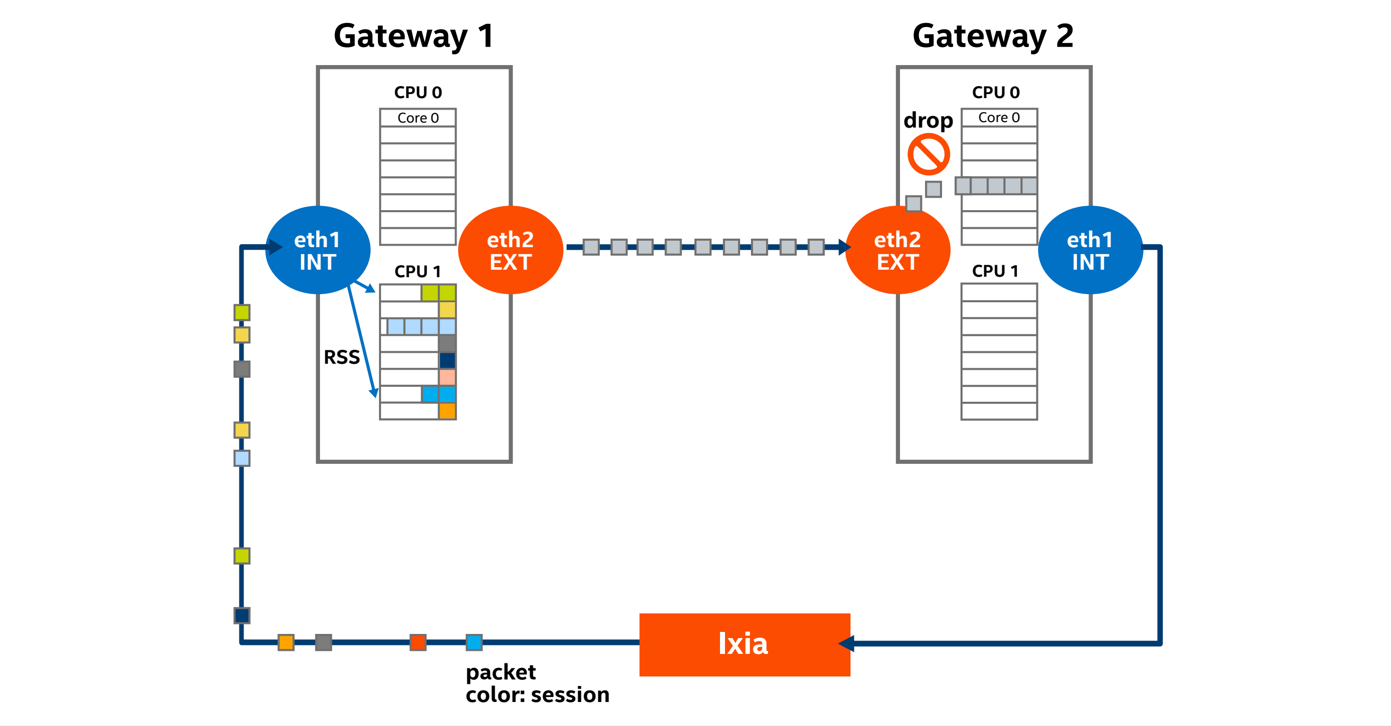

Figure 2. The flow and distribution of packets in IPsec VPN

In Figure 2, traffic flowing into the eth1 interface of gateway 1 is distributed to various CPU cores in session units using RSS. Each CPU core encrypts data in the packets that are assigned to it and adds a new ESP header and new outer IP header and sends it to gateway 2.

In gateway 2, packets flowing into the eth2 interface are distributed using RSS, and because the source and destination address of the encrypted ESP packets are the same, they are not distributed to multiple CPU cores, but are concentrated into a single CPU core. The result is that the decryption of ESP packets is delayed, which causes some packets to be dropped.

SECUI Improves ESP Packet Distribution Method

To solve the CPU distribution issues of ESP packets, SECUI integrated the receive packet steering (RPS) capability built into Linux with Intel Ethernet Flow Director. RPS can distribute packets that are flowing into a single NIC RX queue in Linux to multiple CPU cores. Similarly, Intel Ethernet Flow Director resides in the Intel Converged Ethernet Controller X710 and accelerates the directing of incoming packets to the right processor core.

Instead of RSS that distributes packets using 5 tuple information (source/destination address, source/destination port, protocol) of incoming packets, RPS and Intel Ethernet Flow Director were programmed to look at the specific offset of a packet and assign CPU cores that would process the packets according to the location information. In the specific offset location, the session information of the original packet is included for packet distribution.

Encrypted ESP packets do not include data that can be used to presume the original packet's session information, and therefore, a method to add data to differentiate original packet sessions in the ESP packet is needed. The IV field of the ESP header was used for this.

Figure 3 briefly shows the process of the ESP packet being distributed. It was assumed that packets flow to file server D in host S1, S2, and S3, and are being encrypted and decrypted by the VPN gateway G1, G2 located in the middle of these packets.

In Figure 3, the packet that left host S1 is encrypted into an ESP packet in VPN gateway G1. The source/destination address of ESP packets at this time are changed (G1, G2), and the session information (hash (S1, D)) of the original packets is recorded in the ESP header's IV field. VPN gateway G2 that received the ESP packets checks the session information recorded in the ESP header's IV field, and based on this information, it distributes ESP packets to various CPU cores.

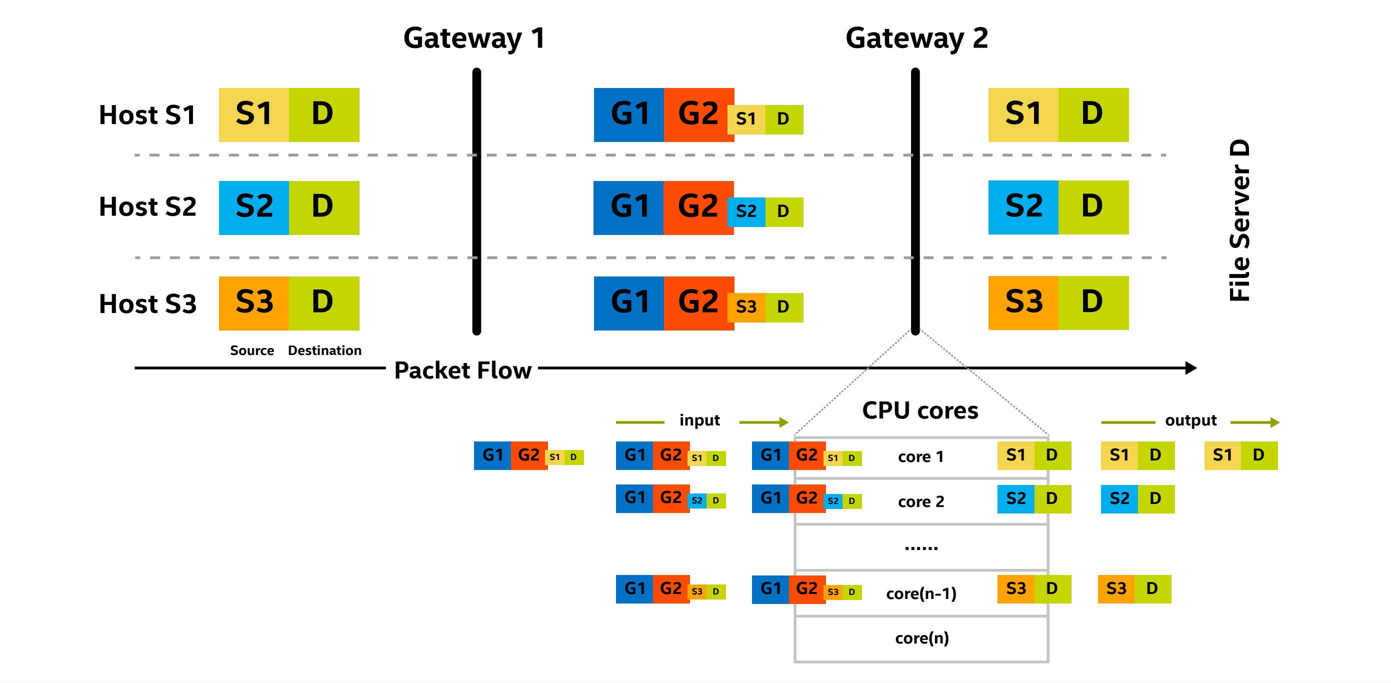

Figure 3. The distribution process of the ESP packet

Saving Session Information Using ESP IV Field

As mentioned earlier, the decryption process must be able to extract session information from the original packet that is encapsulated within the incoming ESP packet to distribute the ESP packet. For this, the hash value of the 5 tuple of the original packet is calculated when encrypting the packet, and this value is recorded in the IV field of the ESP header.

The ESP packet structure is as shown in Figure 4, and the IV field is a type of seed value used to encrypt packets. Values randomly generated for each packet are used and even when there are other network devices existing between the VPN gateways, these do not affect services, so it is suitable to save the session information of the original packet.

The IV field has variable lengths, usually between 8 bytes and 16 bytes according to the encryption algorithm used. This is shown in Figure 5, along with a marker that is used to distinguish whether two bytes among the IV field include session information in the ESP packet. Following the marker, the CPU calculates the hash value for the 5 tuple in two bytes and saves it.

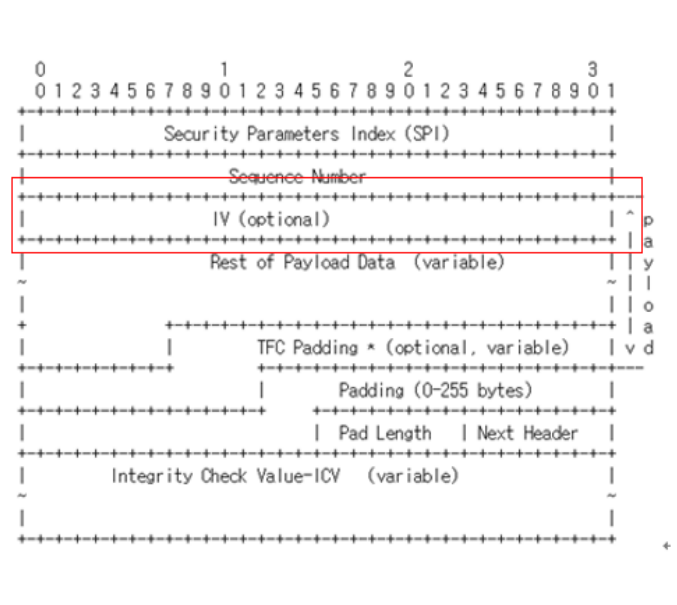

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|               Security Parameters Index (SPI)                 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      Sequence Number                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+---
|                    IV (optional)                              | ^ p
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+ | a
|                    Rest of Payload Data  (variable)           | | y
~                                                               ~ | l
|                                                               | | o
+               +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+ | a
|               |         TFC Padding * (optional, variable)    | v d
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+---
|                         |        Padding (0-255 bytes)        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+     +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                               |  Pad Length   | Next Header   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|         Integrity Check Value-ICV   (variable)                |
~                                                               ~
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 4. The packet architecture[2]

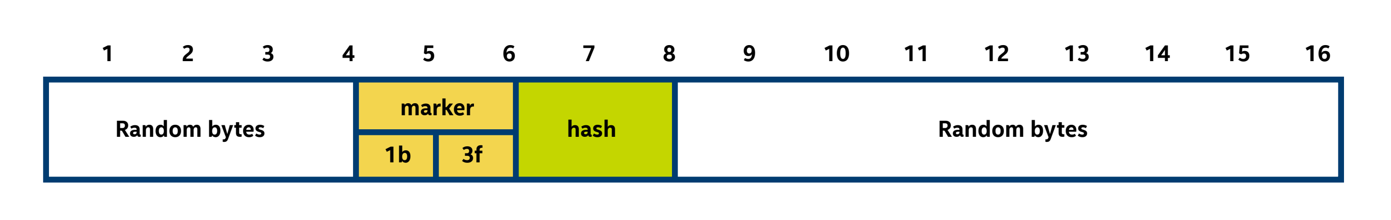

Figure 5. Configuring the IV fields for storing source session information

ESP Packet Distribution Using RPS

In order to distribute ESP packets to several CPU cores, the hash value saved in the ESP header's IV field was looked at and the RPS software filter was expanded to assign CPU cores that can process the ESP packet. It is similar to the Intel Ethernet Flow Director taking into account that RPS is used to distribute incoming packets to multiple CPU cores and that filters can be added to support new distribution methods. In Figure 6, encrypted ESP packets are first sent to a specific CPU core through RSS and then packets are redistributed to RPS to distribute them.

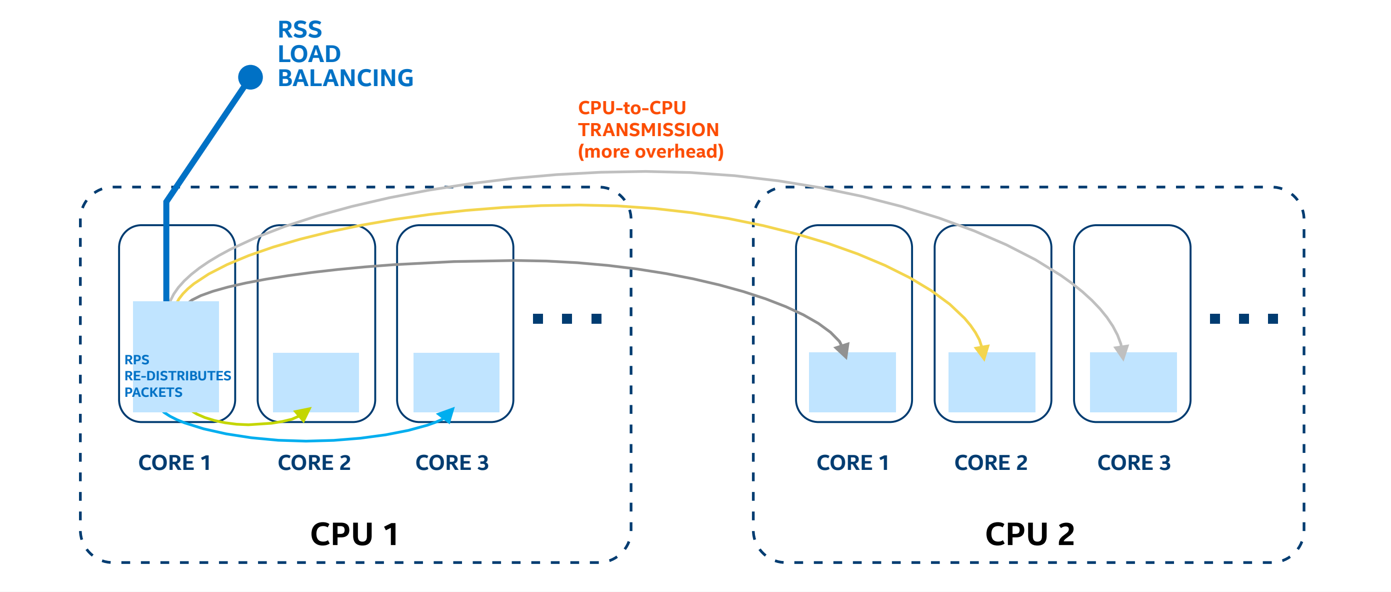

Figure 6. The packet distribution with RPS

ESP Packet Distribution Using Intel® Ethernet Flow Director

Intel Ethernet Flow Director supports packets flowing into the system so that the corresponding packet can be immediately distributed to the CPU core that would consume it (see Figure 7). When receiving ESP packets, they are added to the perfect-match filter rules for distribution in the corresponding CPU core according to the hash value saved in the ESP packet's IV field, and the ESP packet is thereby distributed.

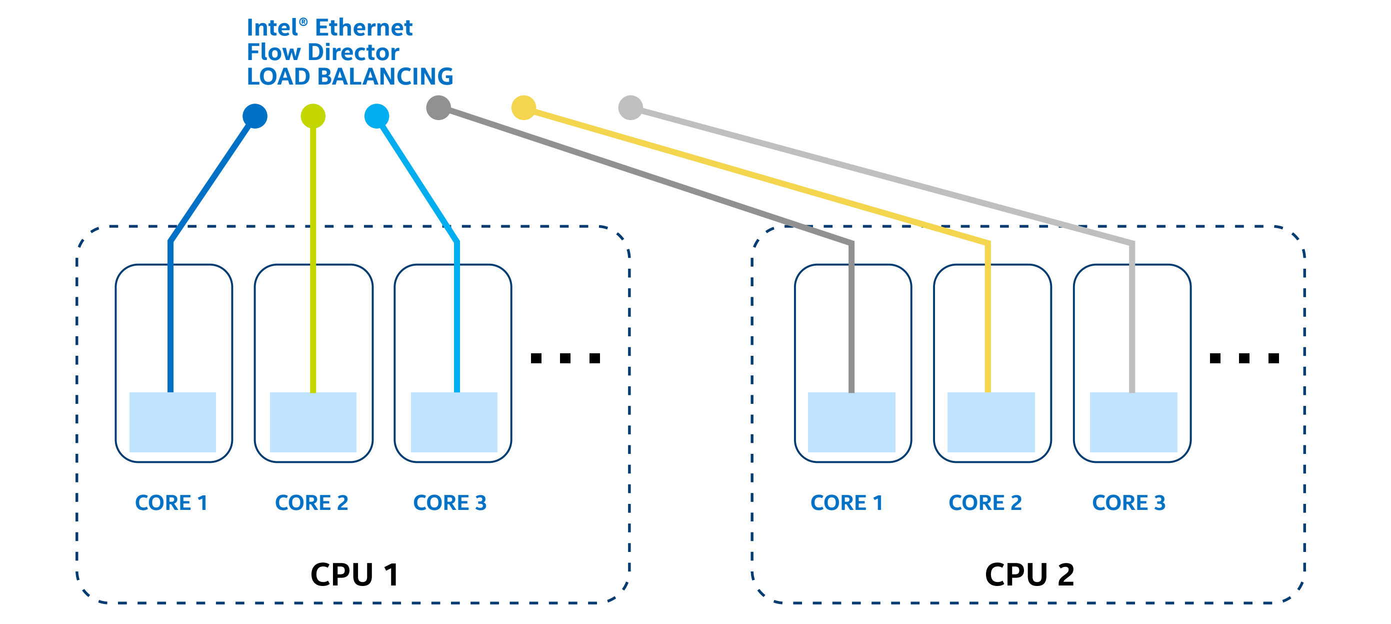

Figure 7. The ESP packet distribution with Intel Ethernet Flow Director

Because Intel Ethernet Flow Director checks the IV field of ESP header before distributing packets with RSS and directly distributes packets into the CPU core, it can safely process packets without unnecessarily consuming CPU resources to redistribute packets unlike the ESP packet distribution method that uses RPS.

Traffic Encryption/Decryption

The process for handling packets in IPsec VPN is shown in Figure 8:

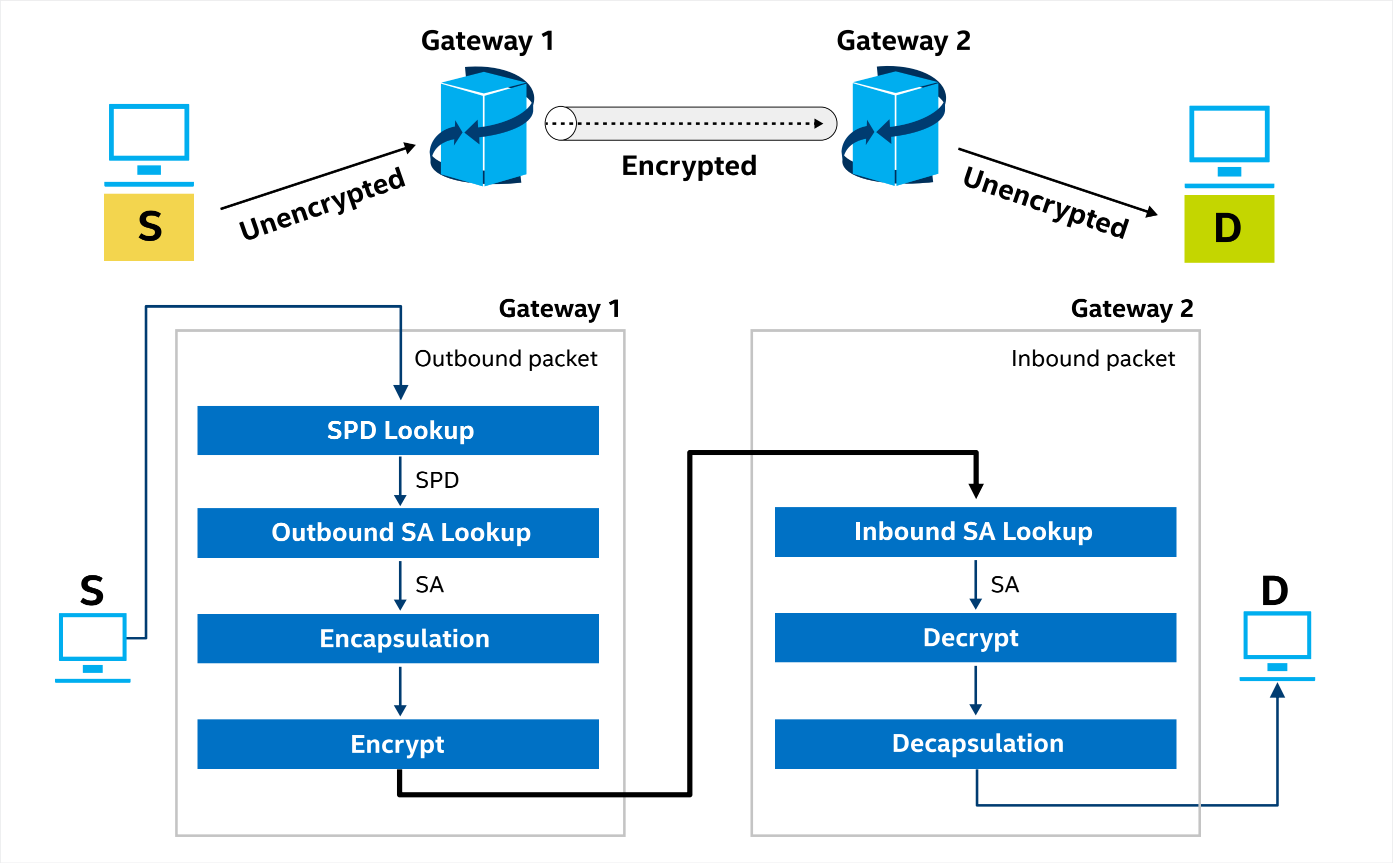

Figure 8. The encryption/decryption process in IPsec VPN

As shown in Figure 8, packets that depart from host S pass through the following process at gateway 1 to go to gateway 2:

1. **Security Policy Database (SPD) lookup:** Select SPD policy to treat packets according to the packet's source/destination address.
2. **Outbound SA lookup:** Search outbound SA connected to the corresponding SPD policy.
3. **Encapsulation:** ESP header/trailers are attached to the packet.
4. **Encrypt:** Auth field is generated to verify integrity of encryption packet for the original packet's data.

Gateway 2 decrypts ESP packets flowing into the system through the following process and sends the decrypted packet to the final destination host D.

1. **Inbound SA Lookup:** Search inbound SA for processing packets from the ESP packet's destination address, protocol and SPI.
2. **Decrypt:** Decrypts the corresponding packet using the selected Inbound SA information and verifies integrity of the received packet.
3. **Decapsulation:** Removes the header/trailers and sends the decrypted packet to the final destination.

Encryption/Decryption Using Intel® QAT

From the IPsec VPN packet handling process explained above, the greatest amount of CPU resources are used during the encryption/decryption process, and in order to reduce CPU resources used here and to enhance IPsec handling performance, Intel QAT is used in the BLUEMAX NGF 5000 to encrypt/decrypt data at the hardware level instead of using software cryptography using the CPU.

Figure 9 below shows Intel QAT services while processing IPsec VPN packets. It also shows how the Data Plane API is used to enable the Intel QAT services.[3] Intel QAT service requests are summoned asynchronously and once the encryption/decryption process is completed, the callback function is summoned to send the encrypted/decrypted data.

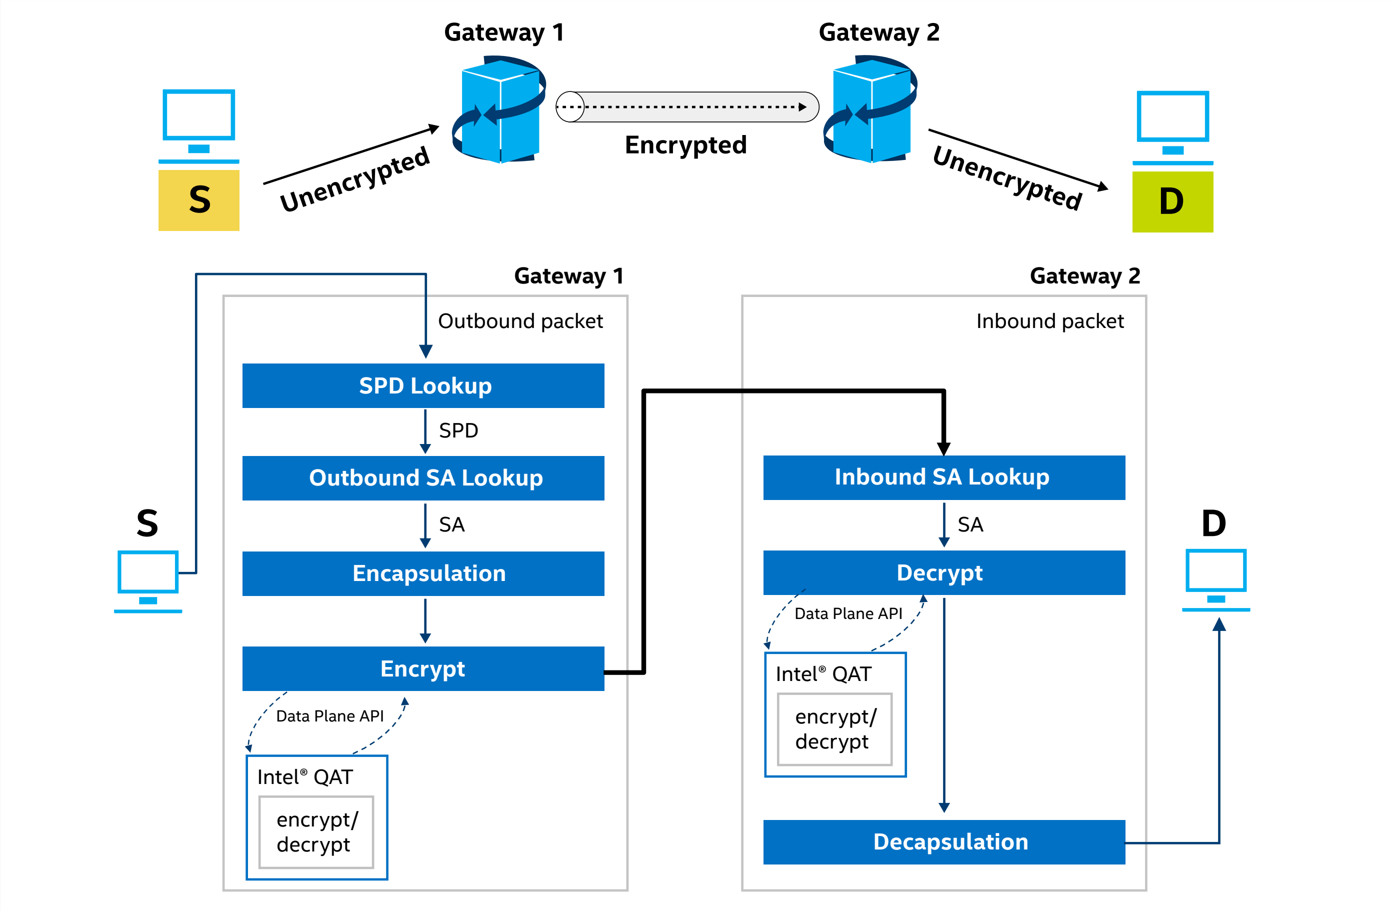

Figure 9. The encryption/decryption process using Intel® QAT

IPsec VPN Benchmark Testing

Intel and SECUI devised a test configuration to demonstrate the performance of the IPsec VPN described above. The configured test environment is shown in Figure 10. VPN 1 and VPN 2 used BLUEMAX NGF 5000 to process IPsec VPN and each BLUEMAX NGF used two 10 G Ethernet ports to send and receive encrypted packets and unencrypted packets.

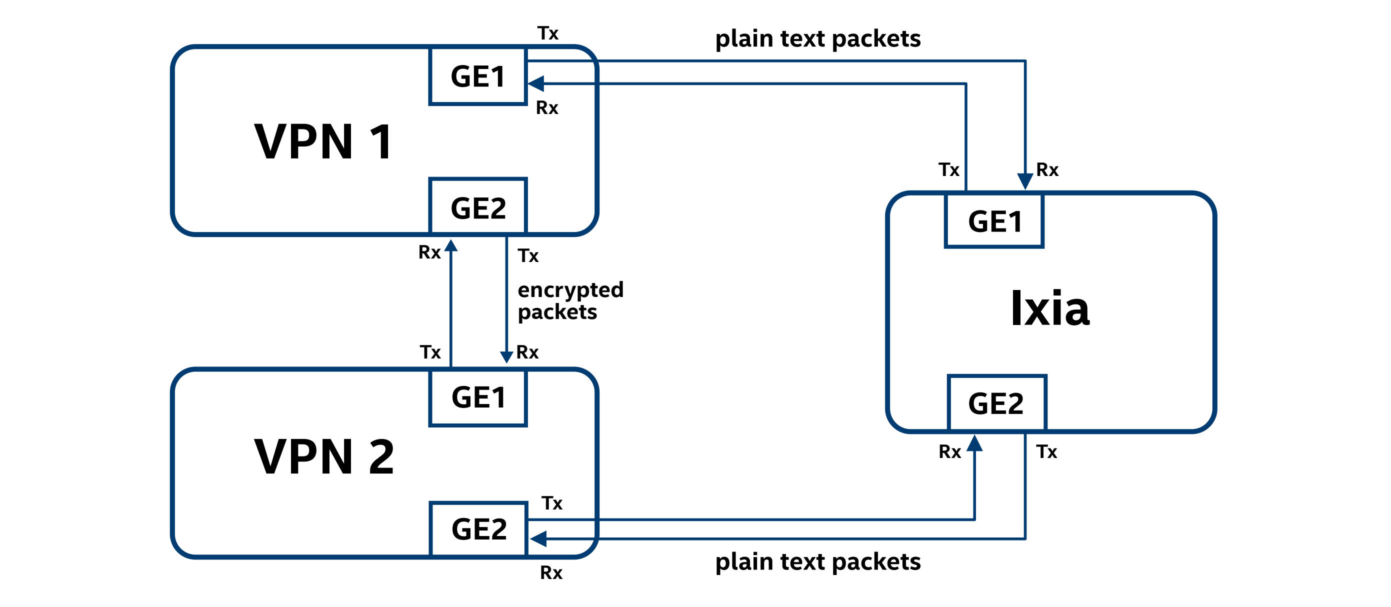

Figure 10. Configuration for Performance Evaluation

The hardware specifications of BLUEMAX NGF 5000 device under test (DUT) are listed in Table 1.

PROCESSOR	Two Intel® Xeon® Gold 6126 CPU (12C/24T) 2.6 GHz Microcode : 0x200004d
MEMORY	12 DDR4 ECC-RDIMM 2666 MHz 8GB (6 Channel + 6 Channel)
NIC	Intel® Ethernet Converged Network Adapter X710 (two 10 G ports)
INTEL® QAT	Intel® C627 Chipset

Table 1. Device under test hardware specifications

While the BLUEMAX NGF 5000 is equipped with two Intel® Xeon® Gold 6126 processors, only one processor was used to receive traffic and process security functions. RCF-2544 bidirectional traffic processing performance was checked using an Ixia measuring device

The performance measurement standards are listed in Table 2.

MEASURING METHOD	RFC 2544
PACKET SIZE	64, 128, 256, 512, 1024, 1280
ALGORITHM	AES256 / SHA256 / CBC
NUMBER OF IPSEC TUNNELS	1
NUMBER OF SESSIONS	250
LOSS TOLERANT	0%

Table 2. Performance measurement standards

Performance According to Packet Distribution Methods

Three packet distribution methods were measured:

1. Using RSS to distribute traffic to a single CPU core
2. Using a combination of RSS and RPS software to distribute traffic
3. Using Intel Ethernet Flow Director to distribute traffic using hardware

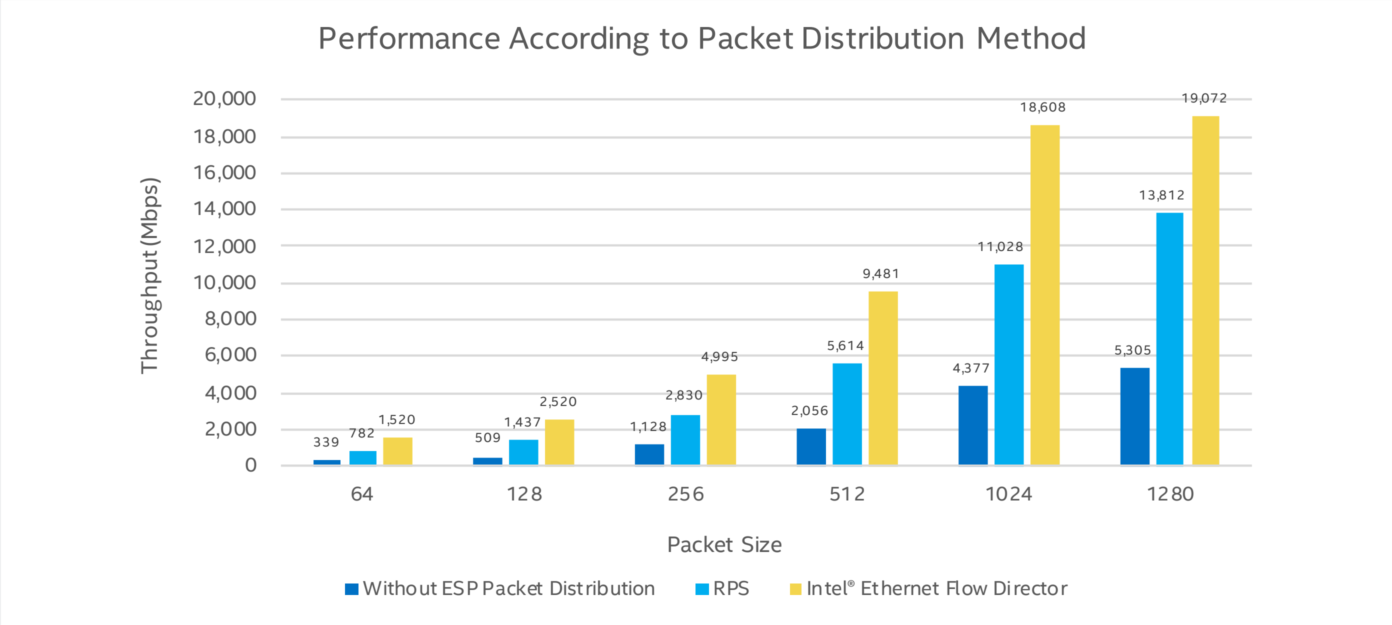

Figure 11. Performance according to packet distribution method[1]

Test results are in Figure 11 and show that when using Intel Ethernet Flow Director, performance was better compared to distribution using Linux RPS (method two) with the improvement difference based on packet sizes and obtained using Intel QAT. Performance doubled at 64 byte packet sizes with significant performance gains at other packet sizes.[1]

Performance Improvements When Using Intel® QAT

Figure 12 shows results of performance tests using Intel QAT and Intel Ethernet Flow Director:

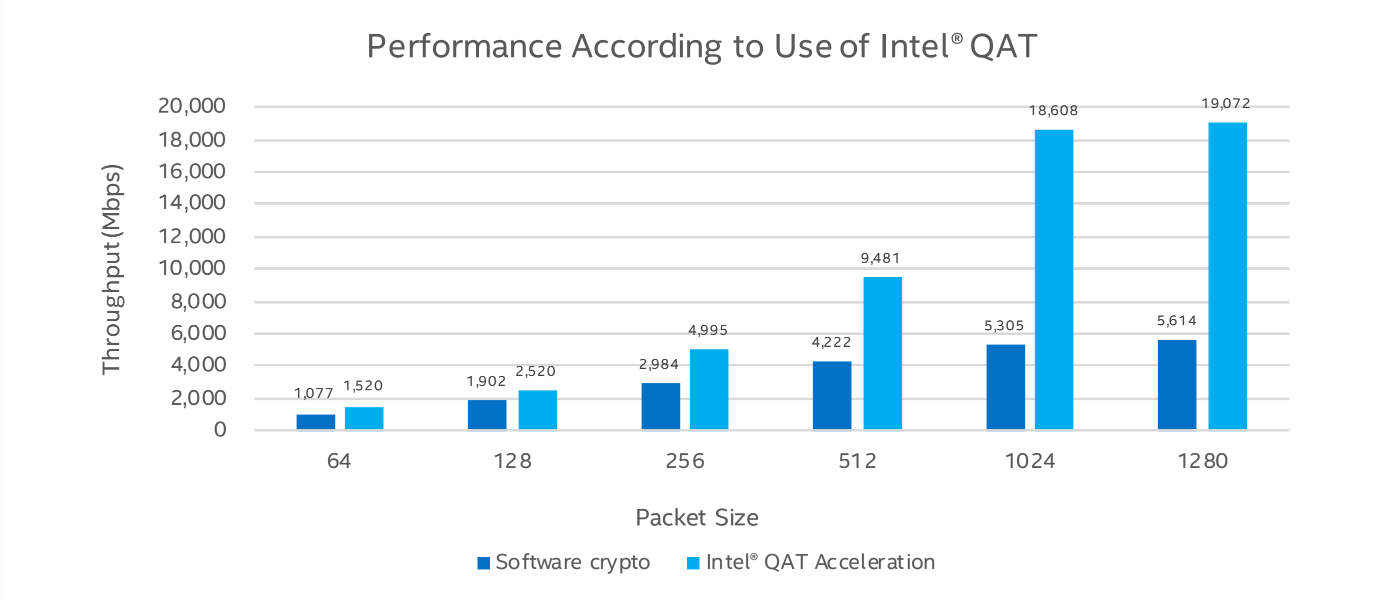

Figure 12. Performance according to use of Intel QAT[1]

When tests used Intel QAT and Intel Ethernet Flow Director, performance improved at all packet sizes when compared to CPU-based crypto processing. Performance more than doubled at 512 bytes packets and more than tripled at both 1024 byte and 1280 byte packet sizes.[1]

When the packet size is small, the gain from the encryption/decryption process is small because latency is smaller than with big packets. This is due to the overhead (some CPU cycles are spent processing the Intel QAT request). The overhead, however, is not big and the overall throughput is improved.

Summary

IPsec performance is an important metric for NGF systems because it requires significant encryption/decryption processing. This can be handled by the CPU, but SECUI set out to optimize it for multi-CPU systems and found that its innovative IPsec VPN design combined with Intel Ethernet Flow Director and Intel QAT resulted in improved traffic processing performance. With results that are up to double and triple that of systems that don't use these technologies, the SECUI BLUEMAX NGF 5000 offers a new level of encryption/decryption performance.[1]

About SECUI

SECUI is a company that specializes in information security and it was established in 2000 with the goal of constructing a perfect information security system. Based on the best technological capacities, it launched next generation firewalls equipped with application recognition and control functions, and it is supplying high quality solutions through strict quality control. It is the number 1 company in Korea's security sector and from 2012 to 2018 for seven consecutive years, it has maintained the largest share in the network security market and firewall/VPN market. Furthermore, it provides various solutions to a number of nations around the world. Technicians with the best technological skills and diverse experiences provide the best security services starting from information protection solution development to security monitoring, security SI services, and all fields related to security.

About Intel® Network Builders

Intel® Network Builders is an ecosystem of infrastructure, software, and technology vendors coming together with communications service providers and end users to accelerate the adoption of solutions based on network functions virtualization (NFV) and software defined networking (SDN) in telecommunications and data center networks. The program offers technical support, matchmaking, and co-marketing opportunities to help facilitate joint collaboration through to the trial and deployment of NFV and SDN solutions. Learn more at http://networkbuilders.intel.com.

[1] Testing conducted by SECUI on May 27, 2019, using the following configuration: Server with two, 2.6 GHz Intel® Xeon® Gold 6126 processors, µcode 0x200004d, (12 cores/ 24 threads) with 12 8GB DDR4 ECC-RDIMM 2666MHz RAM. Intel QAT services were provided by the Intel C627 chipset. Connectivity was provided using 10 GbE ports on two Intel® Ethernet Converged Network Adapters X710.

[2] Figure provided courtesy of Secui.

[3] DP API is explained in chapter 3.2 of the "Intel QuickAssist Technology Performance Optimization Guide" (document number: 330687-004). The guide is available online at https://01.org/intel-quick-assist-technology/intel-quickassist-technology-drivers-and-patches/downloads/intel%C2%AE-quickassist-technology-performance-optimization-guide.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.

Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit www.intel.com/benchmarks.

Performance results are based on testing as of May 27, 2019, and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure.

Intel does not control or audit third-party data. You should review this content, consult other sources, and confirm whether referenced data are accurate.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. Check with your system manufacturer or retailer or learn more at intel.com.

Cost reduction scenarios described are intended as examples of how a given Intel- based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

Intel, the Intel logo, and Xeon are trademarks of Intel Corporation or its subsidiaries.

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