Accelerated Inline IPsec Communication with T6 25G

Executive Summary

Chelsio Crypto Accelerator is a co-processor designed specifically to perform computationally intensive cryptographic operations more efficiently than general-purpose CPUs. Servers with system load, comprising of cryptographic operations, see great performance improvement by offloading crypto operations on to the Chelsio Unified Wire adapter. Chelsio's solution uses the standard crypto API framework provided by the operating system and enables the offloading of crypto operations on to the adapter.

This paper highlights Chelsio T6 adapters' unique accelerating capabilities for secure Inline IPsec-based connections by comparing its bandwidth and CPU usage with Intel AES-NI. T6 consumes less CPU cycles and provides consistently higher throughput across the range of I/O sizes compared to Intel AES-NI. Chelsio T6's Inline IPsec solution provides enterprises with secure remote connection to access corporate applications and resources without sacrificing on performance and speed.

Chelsio Inline IPSEC Acceleration

The Terminator 6 (T6) ASIC from Chelsio Communications, Inc. is a sixth generation, high performance 1/10/25/40/50/100 Gbps, unified wire engine which offers crypto offload capability for AES and SHA variants. Internet Protocol Security (IPsec) is an end-to-end security scheme that provides protocols to ensure the authenticity, privacy and integrity of data in transit. Inline IPsec can be used to implement IPsec-aware systems that have a better latency than lookaside-assisted and accelerated hardware, providing that the algorithm supported is suitable. Chelsio solution provides an accelerated Inline IPsec tunnel which is well suited for site-to-site security over WAN.

Chelsio crypto accelerator secures data using AES (Advanced Encryption Standard) - the strongest encryption algorithm available. Encryption and decryption processing for IPsec is offloaded on to the T6 adapter freeing CPU resources for other tasks. Chelsio crypto driver registers with the kernel crypto framework with high priority and ensures that encryption request is offloaded and processed by T6. IPsec protocol integrated in the kernel calls the crypto API framework which transforms the API into Chelsio supported crypto routines. The data is encrypted and decrypted in the loopback mode for both Tx and Rx paths.

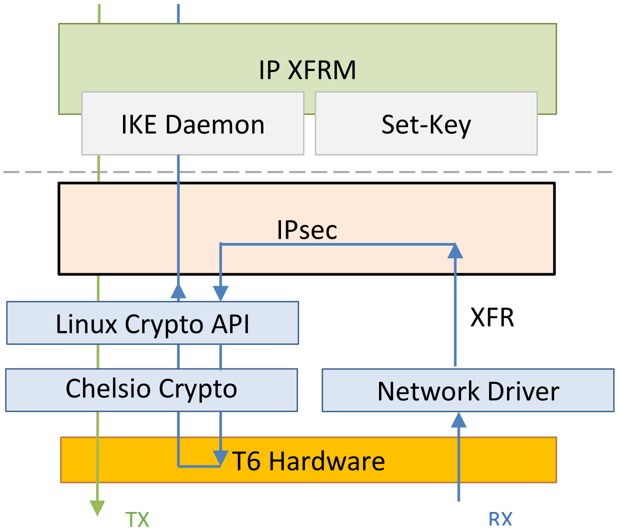

Figure 1 – T6 IPsec Acceleration

Test Results

The following graphs compares the throughput and CPU Usage of Chelsio crypto (offload) and AES-NI modes.

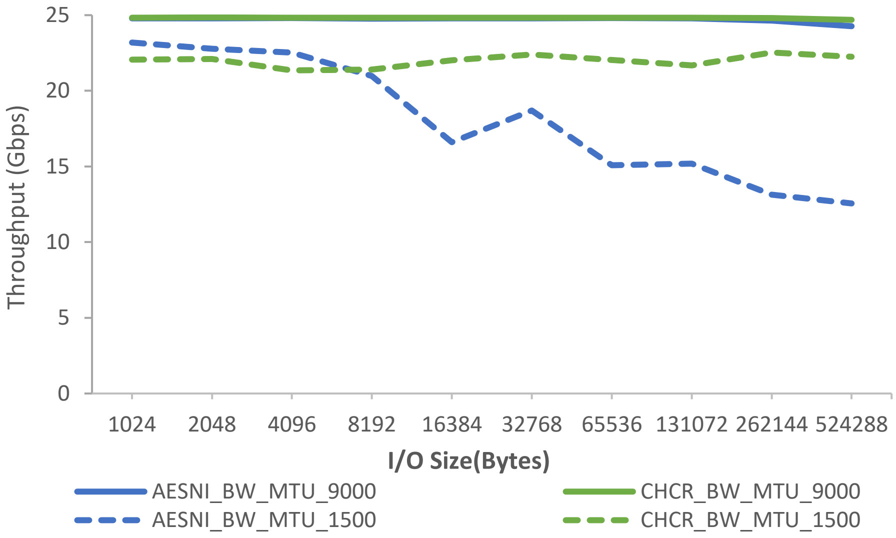

Figure 2 – Throughput: Chelsio crypto vs. Intel AES-NI

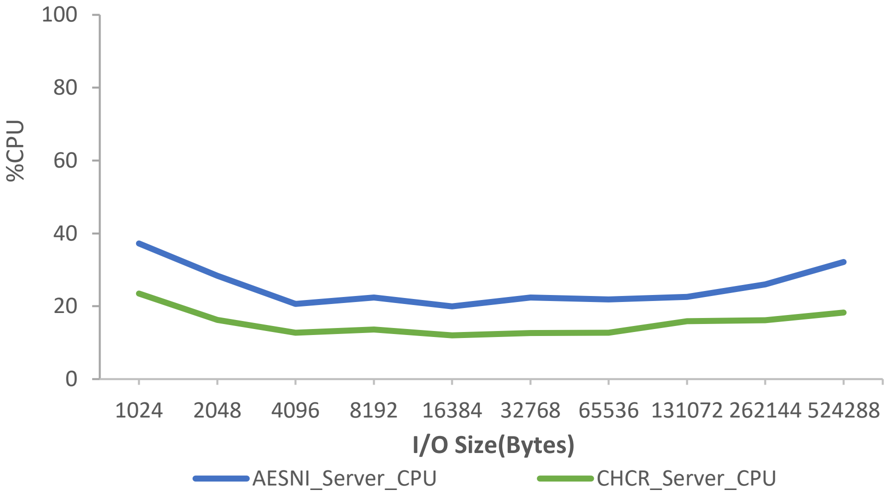

Figure 3 – Server CPU usage: Chelsio crypto vs. Intel AES-NI

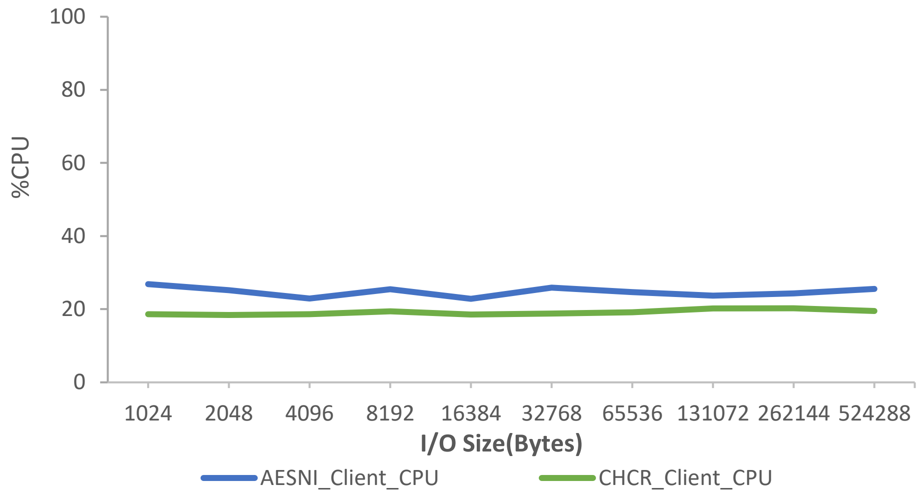

Figure 4 – Client CPU usage: Chelsio crypto vs. Intel AES-NI

Chelsio Inline IPsec solution maintains a flat throughput profile, delivering more than 24Gbps, consuming very less CPU cycles comparing to Intel AES-NI and while preserving the privacy and integrity of the data.

Test Configuration

The setup consists of two machines connected back-to-back using a single 25GbE port- a Server and a Client. Each system was configured with 2 Intel(R) Xeon(R) CPU E5-2687W v2 16-core processors clocked at 3.40GHz (HT enabled), 78GB of RAM and RHEL 7.4 operating system (kernel 5.0-rc8). Chelsio T6225-LL-CR adapter was installed in each system.

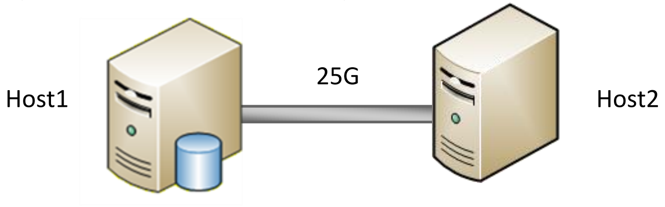

Figure 5 – Back to Back topology

Test Configuration

Following steps were executed on both the Hosts:

i. Install the 5.0-rc8 kernel with the below crypto components enabled.

```
CONFIG_CRYPTO_CCM=y
CONFIG_CRYPTO_GCM=y
CONFIG_CRYPTO_CTS=y
CONFIG_CRYPTO_XTS=y
CONFIG_CRYPTO_GHASH=y
CONFIG_CRYPTO_SHA512=y
CONFIG_CRYPTO_DEFLATE=y
CONFIG_CRYPTO_ANSI_CPRNG=y
CONFIG_CRYPTO_USER_API_RNG=y
CONFIG_CRYPTO_USER_API_AEAD=m
```

```
CONFIG_CRYPTO_DEV_CHELSIO=m
CONFIG_INET_ESP=m
CONFIG_INET_ESP_OFFLOAD=m
CONFIG_CHELSIO_IPSEC_INLINE=y
CONFIG_INET6_ESP_OFFLOAD=m
```

ii. Reboot the machine into newly installed kernel.

iii. Chelsio network driver was loaded.

```
[root@host~]# modprobe cxgb4
```

iv. Respective drivers were loaded for Chelsio crypto modes.

Chelsio crypto:

```
[root@host~]# modprobe -v chcr
[root@host~]# modprobe -v esp4_offload
[root@host~]# modprobe -v esp6_offload
[root@host~]# modprobe -v authenc
```

Intel AES-NI:

```
[root@host ~]# modprobe -v aesni_intel
```

v. IP alias was configured on Chelsio interfaces for 8 tunnels.

```
[root@host~]# for i in `seq 1 8`;do ifconfig ethX:$i $i.0.0.2/24 up;done
```

vi. Turn off the TSO & GSO to ensure the traffic flows through the inline path.

```
[root@host~]# ethtool -K ethX tso off
[root@host~]# ethtool -K ethX gso off
```

vii. CPU affinity was set.

```
[root@host~]# t4_perftune.sh -n -Q nic
[root@host~]# t4_perftune.sh -n -Q crypto
[root@host~]# cpupower frequency-set --governor performance
[root@host~]# tuned-adm profile network-throughput
```

viii. The following *sysctl* parameters were set:

```
sysctl -w net.ipv4.tcp_timestamps=0
sysctl -w net.ipv4.tcp_sack=1
sysctl -w net.core.netdev_max_backlog=250000
sysctl -w net.core.rmem_max=4194304
sysctl -w net.core.wmem_max=4194304
sysctl -w net.core.rmem_default=419430
sysctl -w net.core.wmem_default=4194304
sysctl -w net.core.optmem_max=4194304
sysctl -w net.ipv4.tcp_rmem="4096 87380 4194304"
sysctl -w net.ipv4.tcp_wmem="4096 65536 4194304"
sysctl -w net.ipv4.tcp_low_latency=1
sysctl -w net.ipv4.tcp_adv_win_scale=1
sysctl -w net.ipv4.tcp_timestamps=0
```

ix. ip-xfrm was configured using below scripts on Server and Client.

Server

```
 for i in `seq 1 8`
do
  j=`expr $i - 1`
  local_ip=$i.0.0.1
  remote_ip=$i.0.0.2
  taskset -c $j ip xfrm state add src $local_ip dst $remote_ip proto esp spi
0x53fa1f$i    reqid    16386    mode    transport    aead    "rfc4106(gcm(aes))"
0x010203047aeaca3f87d060a12f4a4487d5a5c335  96  sel  src  0.0.0.0/0  dst  0.0.0.0/0
offload dev enp4s0f4 dir out
  taskset -c $j ip xfrm state add src $remote_ip dst $local_ip proto esp spi
0x54fa1f$i    reqid    16385    mode    transport    aead    "rfc4106(gcm(aes))"
0x010203047aeaca3f87d060a12f4a4487d5a5c335  96  sel  src  0.0.0.0/0  dst  0.0.0.0/0
offload dev enp4s0f4 dir out
  taskset -c $j ip xfrm policy add src $local_ip dst $remote_ip dir out tmpl src
$local_ip dst $remote_ip proto esp reqid 16385 mode transport
  taskset -c $j ip xfrm policy add src $local_ip dst $remote_ip dir in tmpl src
$local_ip dst $remote_ip proto esp reqid 16386 mode transport
done
```

Client

```
for i in `seq 1 8`
do
  j=`expr $i - 1`
  local_ip=$i.0.0.1
  remote_ip=$i.0.0.2
  taskset -c $j ip xfrm state add src $remote_ip dst $local_ip proto esp spi
0x53fa1f$i reqid 16386 mode transport aead "rfc4106(gcm(aes))"
0x010203047aeaca3f87d060a12f4a4487d5a5c335 96 sel src 0.0.0.0/0 dst 0.0.0.0/0
offload dev enp4s0f4 dir out
  taskset -c $j ip xfrm state add src $local_ip dst $remote_ip proto esp spi
0x54fa1f$i reqid 16385 mode transport aead "rfc4106(gcm(aes))"
0x010203047aeaca3f87d060a12f4a4487d5a5c335 96 sel src 0.0.0.0/0 dst 0.0.0.0/0
offload dev enp4s0f4 dir out
  taskset -c $j ip xfrm policy add src $local_ip dst $remote_ip dir out tmpl
src $local_ip dst $remote_ip proto esp reqid 16385 mode transport
  taskset -c $j ip xfrm policy add src $local_ip dst $remote_ip dir in tmpl
src $local_ip dst $remote_ip proto esp reqid 16386 mode transport
done
```

x. The xfrm policies were verified:

```
[root@host~ ]# ip xfrm state list
```

xi. iperf servers with 8 different IP alias were started on Host1.

```
[root@host~]# for i in `seq 1 8`; do taskset -c 0-8 iperf -s -w 512k -p 500$i
& done
```

xii. Connections to the listening servers were established from the Host2.

```
[root@host~]# for i in `seq 1 8`; do taskset -c 0-8 iperf -c $i.0.0.2 -p 500$i
-w 512K -l <I/O size> -t 60 & done
```

Conclusion

This paper presented performance comparison of Chelsio's T6 Inline IPsec solution and Intel's AES-NI using T6225-LL-CR adapter in Linux. With a consistent line-rate 25Gbps performance and CPU savings, Chelsio's solutions proves to be the best choice for clients looking for highest level of data security and integrity, without compromising performance. In addition, Chelsio's CPU usage is significantly low compared to Intel's, indicative of a more efficient processing path. Chelsio's T6 low-cost Inline IPsec solution provides data transfer with high accuracy and speed without affecting integrity and confidentiality.

Related Links

Concurrent Offload & Encryption at 100GbE
Disk Encryption with 100GbE Crypto Accelerator
T6 100GbE Crypto Offload Performance
Chelsio Terminator 6 ASIC 100GE Crypto Offload
The Chelsio Terminator 6 ASIC
Chelsio T6 Crypto Offload Video