Chelsio Terminator 6 ASIC 100GE Crypto Offload

Enabling Secure Network Interconnects

Introduction

Chelsio introduces ground breaking TLS/SSL performance with inline cryptographic functions leveraging Chelsio's proprietary TCP/IP offload engine. Chelsio's full offload TLS/SSL is uniquely capable of 100Gb line rate performance. In addition, accelerator can be used with inline mode for DTLS and in a traditional co-processor Lookaside mode to accelerate TLS/SSL, IPsec, SMB 3.X crypto, data at rest encryption/decryption, and data-deduplication fingerprint computation.

Security Offloads

The T6 supports all the most popular AES/SHA cipher suites in TLS/SSL in-line mode with 100Gbps bandwidth and less than 2µs end-to-end latency. The typical T6 SKU supports 32K simultaneous TLS sessions and the T6 has support for up to 1M simultaneous sessions. The in-line mode achieves TCP/IP processing and TLS/SSL AES/SHA processing in cut-through fashion to achieve optimal bandwidth and latency. A co-processor mode of operation is supported for TLS/SSL, SMB 3.X, IPsec, data at rest encryption/decryption, authentication, and data de-dupe fingerprint generation. The TLS/SSL session key negotiation is performed by software on a host computer. The T6 is therefore ideal for TLS/SSL encryption authenticated media streaming applications and it also has efficient support for SMB 3.X, IPsec, data at rest encryption/decryption, authentication, and data de-dupe fingerprint generation.

The performance of the AES and SHA protocol suites is summarized in the following table:

Cipher	BW	Latency
AES-CBC	Encryption=30Gbps/Decryption=100Gbps	< 10µs
SHA1	40Gbps	< 10µs
SHA224/256/384/512	25-40Gbps	< 10µs
AES-GCM/CTR/XTS	100Gbps	< 1µs

The supported options with the AES and SHA protocols are summarized in the following tables:

Cipher only modes (encryption/decryption only):

Cipher	Key Sizes supported	Protocol Requirement
AES-CBC	128, 192, 256	TLS, IPSEC
AES-CTR	128, 192, 256	IPSEC
AES-XTS	128, 192, 256	Generic Protocol

Combined cipher modes (authentication and encryption/decryption):

Cipher	Key Sizes supported	Protocol Requirement
AES-GCM	128, 192, 256	TLS, IPSEC, SMB 3.1
AES-CCM	128, 192, 256	SMB 3.X (co-processor only)

Authentication and generic hash modes:

Hash Function	Key Sizes supported	ICV Size	Protocol Requirement
SHA1 **SHA224/256/384/512**	Equal to the output of hashing algorithm, it is expected longer keys will be hashed to L bits, refer to RFC2104	Variable	TLS, IPSEC, Generic
SHA1-HMAC **SHA2-224-HMAC** **SHA2-256-HMAC** **SHA2-384-HMAC** **SHA2-512-HMAC**	Equal to the output of hashing algorithm, it is expected longer keys will be hashed to L bits, refer to RFC2104	Variable	TLS, IPSEC

T6 inline TLS/SSL

In the T6 in-line TLS/SSL mode of operation, cleartext is sent from the server and cleartext is received by the server. The T6 uses software on the host e.g. OpenSSL to manage the session key negotiation but once key negotiation is completed, the TLS/SSL operation is fully offloaded to the T6. In the send direction, see below, cleartext (1) is sent to an offloaded TCP/IP connection (2) that is in TLS/SSL mode which causes the T6 egress pipeline to fetch session keys to encrypt the data and to compute authentication codes that are added to the cleartext and a TCP/IP packet with the TLS/SSL format (3) is sent to the wire. The TLS/SSL processing happens in cut-through, as the rest of the TCP/IP processing and therefore adds minimal latency to the processing pipeline.

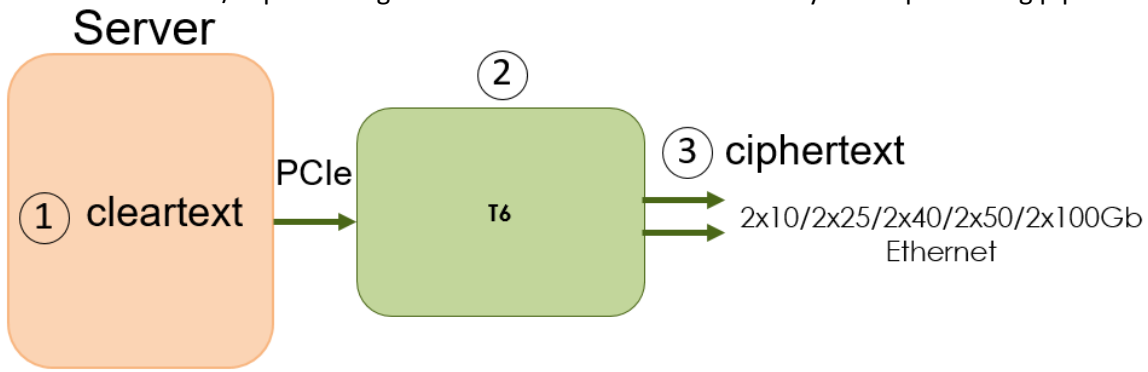

Figure 1 – T6 inline TLS/SSL Send

On receive, see below, a ciphertext TCP/IP segment (1) is received by T6 and if it belongs to a TLS/SSL offloaded connection (2), the session keys are fetched and the ciphertext is decrypted and authenticated and cleartext (3) is delivered to the host. The TLS/SSL processing again happens in cut-through fashion and minimal latency is therefore added, the overall end-to-end latency with TLS/SSL is less than 2µs.

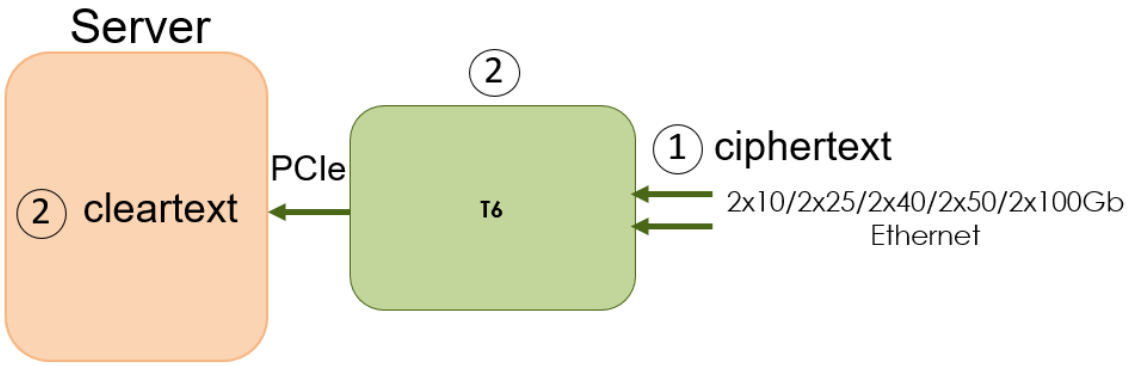

Figure 2 – T6 inline TLS/SSL Receive

There is no assumption that the TLS/SSL segments align with TCP/IP frames and there is no assumption concerning lossless Ethernet behavior or any other limiting assumptions, i.e. T6 in-line TLS/SSL offload works where TCP/IP goes, i.e. within a rack, LAN, MAN or WAN, or wireless. The T6 in-line TLS/SSL performance profile is full-line rate operation up to 100Gbps with less than 1 μs latency added due to TLS/SSL encryption, decryption, and authentication operation.

The following diagram shows an example of the T6 value proposition for the CDN or media streaming use case.

The CDN server (1) delivers 20K 5Mpbs streams of content e.g. video, movie, IPTV using a single T6 that offloads 20K TLS/SSL connections (3) and each of these connections is traffic managed by the integrated T6 traffic manager (2) to proceed at 5Mpbs rate with low jitter. The T6 traffic manager is capable of supporting multiple traffic classes concurrently up to 16 groups e.g. there can be a 25Mbps group in addition to the 5Mpbs group, etc.

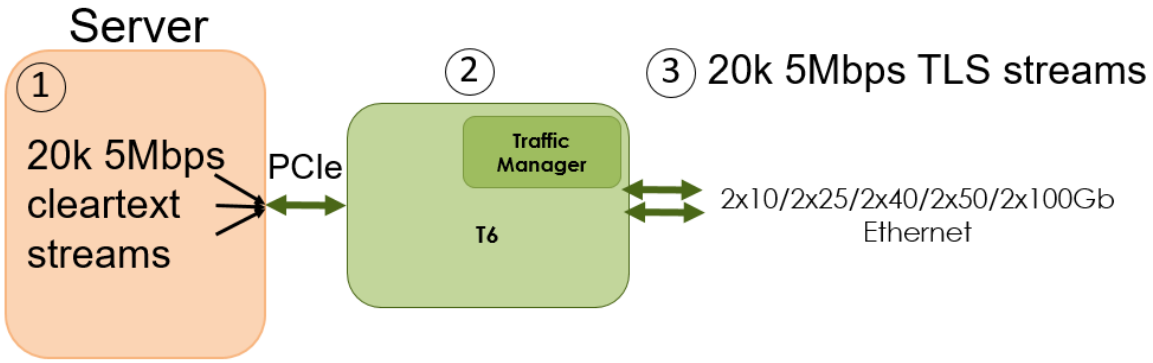

Figure 3 – T6 for CDN or Media Streaming

The T6 in-line TLS/SSL mode of operation delivers an order of magnitude improvement in terms of CAPEX and OPEX for the CDN use case, i.e. it lowers the requirements on the CDN server both in terms of crypto capabilities and in terms of memory sub-system performance. In addition, experience with the Chelsio T5 that has the same traffic manager, shows the offloading of traffic management capability decreases the number of required CDN servers.

T6 In-line DTLS

In the T6 in-line DTLS mode of operation, cleartext is sent from the server and cleartext is received by the server. The T6 uses software on the host for e.g. WebRTC or VXLAN/GENEVE to manage the session key negotiation but once key negotiation is completed, the DTLS operation is fully offloaded to the T6.

Co-processor TLS/SSL, IPsec, and SMB 3.X crypto

In the T6 co-processor mode of operation, either cleartext is sent to the T6 over the PCIe bus for encryption, and authentication code, or encrypted authenticated ciphertext is sent to the T6 for decryption and authentication.

The following diagram shows the encryption process: cleartext (1) is sent to the T6 using Linux crypto API to request any of the supported T6 AES and/or SHA crypto computation be performed on the cleartext. The ciphertext (2) is returned by T6, and finally the cipher text is sent (3) through the T6 to Ethernet (4). In sending the ciphertext, the offload capabilities of T6 can optionally be used to achieve offloaded co-processor mode TLS/SSL or to implement SMB 3.X encryption/authentication. To achieve IPsec encrypted and/or authenticated iSCSI and NVMf, the offload processing is performed first using a loopback connection that loops the offloaded Ethernet/TCP/IP packet back to T6 encryption hardware for IPsec encryption.

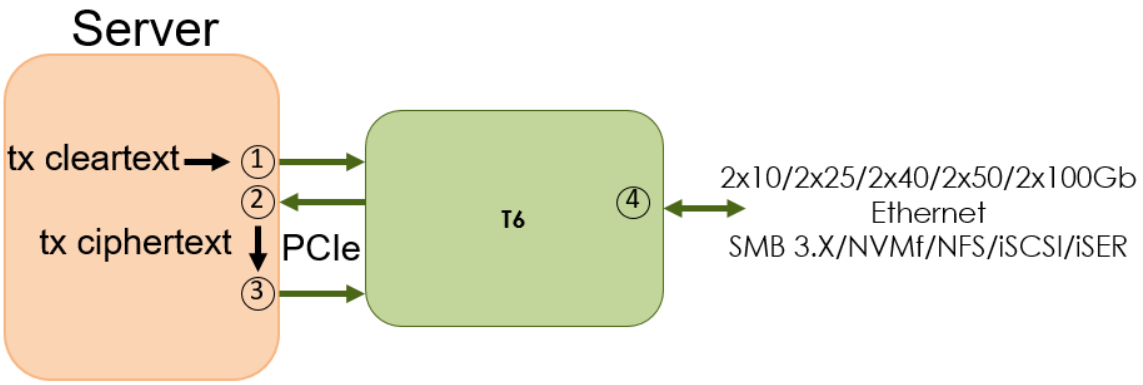

Figure 4 – T6 co-processor Encryption

The following diagram shows the co-processor decryption process: the ciphertext (1) is received from Ethernet and then sent to the host (2) and injected into the T6 crypto hardware (3) for decryption. The T6 crypto returns the cleartext (4). At this point offload processing on the received data can optionally be performed by sending cleartext data to the T6 network processing pipeline. In the case of IPsec of NVMf or iSCSI, the offload processing is performed by an offloaded loopback connection.

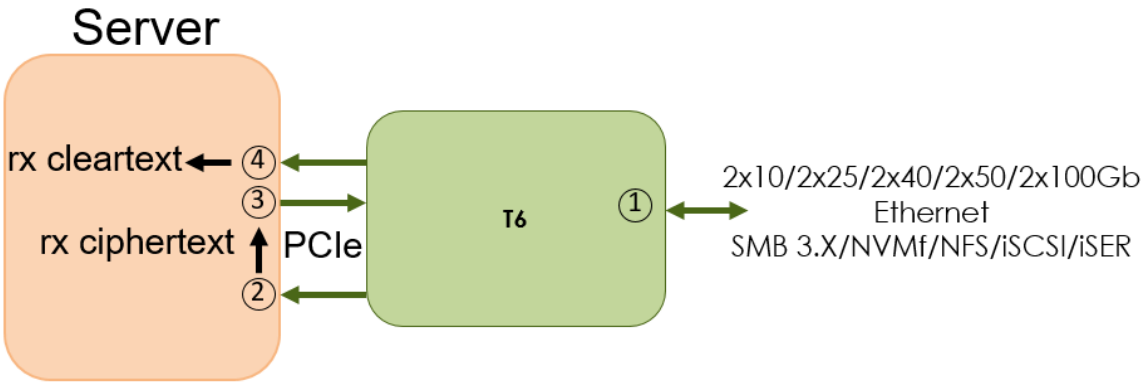

Figure 5 – T6 co-processor decryption

The T6 crypto co-processor mode of operation can be combined with the T6 various offload capabilities to support secure operation always everywhere.

Data at-rest encryption/decryption

The T6 Data at-rest encryption/decryption uses the T6 crypto co-processor mode of operation as shown in the following diagram. The cleartext to be encrypted e.g. with the AES-XTS algorithm is sent (1) to the T6 crypto engine and the ciphertext is returned (2). The decryption proceeds by sending the ciphertext (1) and the cleartext is returned (2).

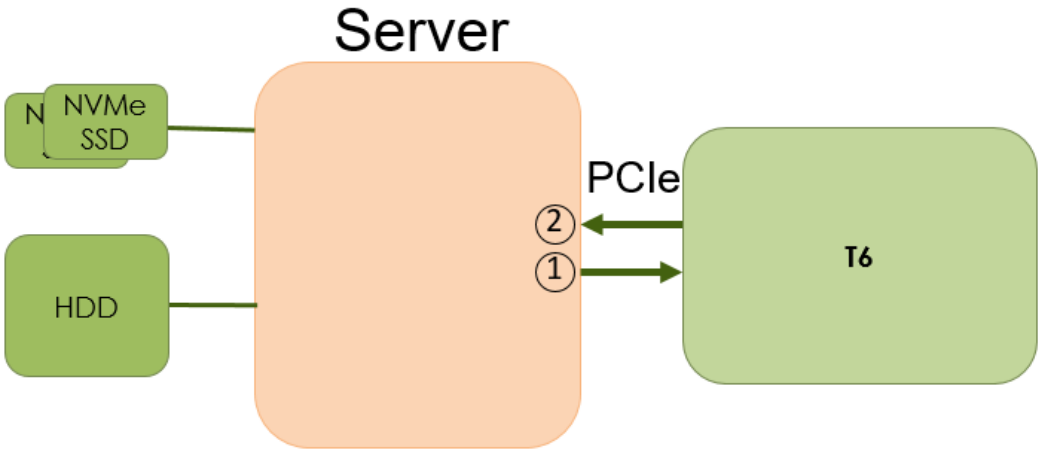

Figure 6 – T6 Data at-rest encryption/decryption

The T6 crypto co-processor mode of operation can be combined with the use of TPM to support highly efficient and low CAPEX and OPEX secure data always everywhere.

Data de-duplication fingerprint

The T6 crypto co-processor can also be used for de-duplication fingerprint generation as shown in the following diagram. In this case e.g. offloaded iSCSI or NVMf has been received (1) and a fingerprint of the received data is then computed by injecting the received data into the T6 crypto co-processor (2) to compute e.g. a SHA hash over storage blocks or objects such as spreadsheets

or PDF documents contained in the received data. The computed fingerprint (3) can then be used to identify opportunities for de-duplication in the storing of the data.

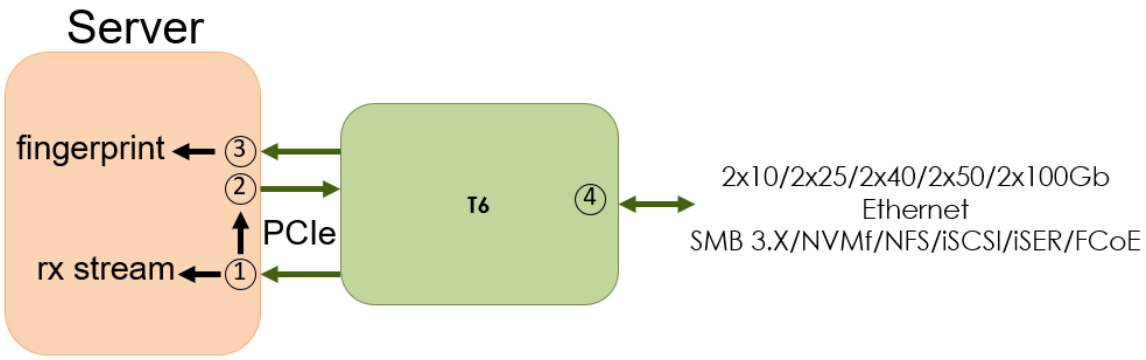

Figure 7 – T6 fingerprint co-processing

The T6 crypto co-processor can lower CAPEX and OPEX by offloading the fingerprint crypto to the T6 NIC rather than investing in a more powerful processor with crypto capabilities.

Conclusions

The concept of secure network convergence around 10, 25, 40, 50, 100GbE has been discussed in the industry for some time. The security aspect is being addressed by the ubiquitous TLS/SSL, the SMB 3.X crypto features, by IPsec, and by DTLS. But changes of this magnitude do not happen overnight.

Against this backdrop of diverse and dynamic requirements, creating a universal secure IP protocol over 10, 25, 40, 50, 100GbE controller offers a superior ROI for the customer. Offloading protocols such as iSCSI and iWARP and TLS/SSL requires a reliable high-performance underlying TCP engine. For secure storage and cluster traffic alike, low-latency/high-IOPS is increasingly important. Virtualization requires significant new hardware functions to support both VM isolation and VM-to-VM communication. Finally, a universal design delivers a high level of integration to meet the total space and cost and power budget requirements of LOM and mezzanine designs.

With its sixth-generation T6 ASIC, Chelsio has taken the unified wire to the next level. T6 delivers an unmatched feature set combined with a single-chip design. No other vendor offers a single SKU for NVMf, NIC, TOE, iSCSI, FCoE, and iWARP RDMA that concurrently support in-line TLS/SSL and that supports SMB 3.X crypto, IPsec, and DTLS. Why settle for partial solutions to server connectivity when Chelsio makes a universal solution today?