

The True Cost of Non-Offloaded NICs

Chelsio T5 & T6 Lines of Offloaded NICs

Executive Summary

This paper examines the true cost of using a non-offloaded NIC on a server. Offloaded NICs relieve the server of CPU-cycle consuming network operations and hence allow the server to be able to do more real work while simultaneously supporting a high performance network. This paper will explore the real savings that the user can expect when using an off-loaded NIC in their environment.

Off-loaded NICs, what are they and what do they provide?

Many NICs on the market today tout that they perform offloading - the ability to run certain functions required for network operations on the NIC itself rather than relying on the CPUs. However, the functions that these NICs offload are typically low-level support functions like checksum calculations and large send offload (LSO), which are not exceptionally CPU intensive and a standard feature on all NICs. Since these functions do not require retention of the state of the connection by the NIC, they are commonly referred to as *Stateless Offloads*.

Chelsio not only provides stateless offloads on its NICs, but also includes other more meaningful offloads, such as full TCP Offload-Engine (TOE). A TOE allows for the full TCP/IP stack to run on the NIC itself, including connection set-up and tear-down, and all exception handling, thus saving considerable host resources. Using the TCP/IP offload engine as the foundation, Chelsio's NICs can provide support for various protocols such as iWARP (RDMA over TCP/IP), iSCSI, NVMe-oF (NVMe running on iWARP), SSL/TLS and others that do not use the TOE (such as IPsec), allowing for tremendous CPU savings.

Various studies have been done to show the effectiveness of using offloads. One white paper authored by Chelsio shows that using TOE on FreeBSD allows a Chelsio T62100-CR 100Gb NIC to deliver *100 Gb/sec throughput using netperf on a single connection, utilizing only 1% of host CPU*. In another paper recently released by Demartek, the performance and scalability advantages of using offload with 25 and 100 Gb/s iSCSI are demonstrated.

OK, I get the value of offloading, now show me the money

Let's define a performance metric $CTS = \# \text{ of Cores} * \text{Threads} * \text{Speed}$. Figures 1 & 2 plot the price of the CPU based on Intel's latest pricing tables vs. CTS for two of the most popular CPUs used today in servers. Extrapolating, in figure 1 we see that for every 10% in CPU savings, we save \$613. In figure 2, the savings are even more significant, saving around \$736 for every 10% of CPU.

E5-2600 Series




Figure 1 – CTS vs Price Curve for Intel E5-2600 series

E7-8800 Series




Figure 2 – CTS vs Price Curve for Intel E7-8800 series

What does it mean in a real-world environment?

By using an offloaded NIC, the user can realize higher performance of the overall system, since the CPUs have more time to do real work, as opposed to having to manage the high-performance network. In other words, the user can realize the same performance that they desire while using lower performance, lower power and lower cost CPUs.

In case of a storage target head or other fixed-application, fixed-workload appliance, for example, that runs either iSCSI, TOE or NVMe-oF, the same dual socketed head can provide the same IOPs and bandwidth performance, with at least one or two slower speed grades of an x86 CPU, resulting in thousands of dollars of savings per head. The savings go beyond acquisition costs and into operational costs with larger volumes. For example, on a server running ~9000 hrs/year at \$0.10/kwh, and with ~\$3000 initial server cost and 5-year refresh cycle, if we save ~30% of CPU using offloaded NICs, the user saves as much in power and heat removal costs to buy another server!

Here, it is essential to look at costs through the prism of the cost of the whole solution, as opposed to the cost of the NIC only, added to the cost of the CPU. This is why some vendors for example bundle a stateless NIC essentially for free with a high-end CPU – the result is deceptively more expensive (both the acquisition cost and operational cost).

Conclusions

Using offloaded NICs allows for less total cost of ownership, while achieving higher performance metrics.

Related Links

[Intel Recommended Customer Price List](#)

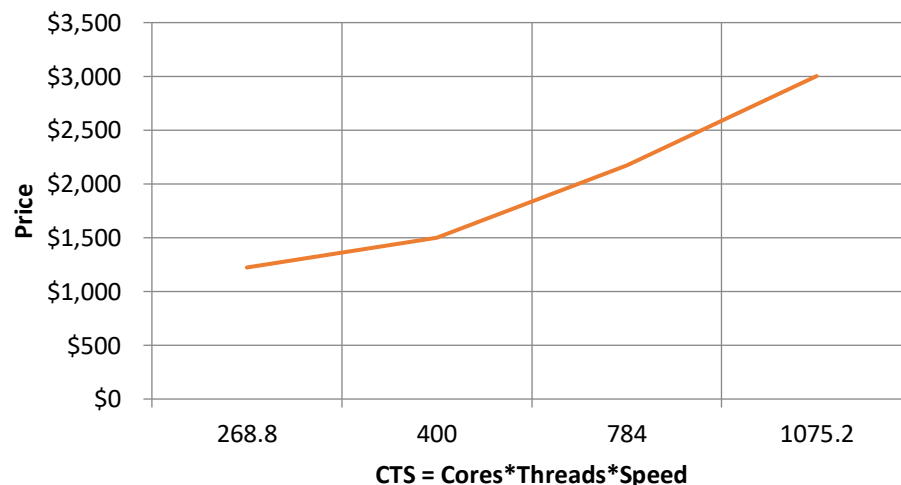
[Demartek Evaluation: Chelsio Terminator 6 \(T6\) Unified Wire iSCSI Offload](#)

[Industry's First 100G Offload with FreeBSD](#)

[Other Chelsio White Papers](#)

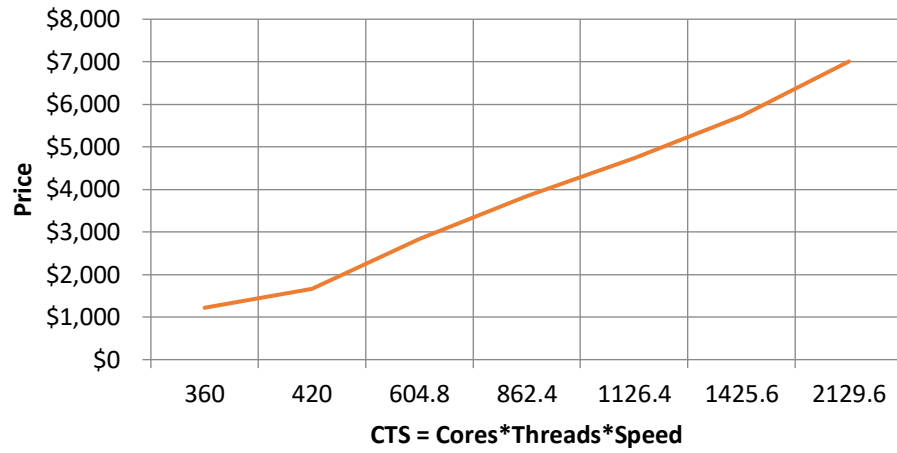
Additional CPU models

E7-4800 Series




Every 10% CPU savings is ~\$667.

E5-4600 Series



Every 10% CPU savings is ~\$527.

E5-1600 Series



Every 10% CPU savings is ~\$263.