

iWARP Support in Scalable Xeon Platform

Frequently Asked Questions

What are the highlights vis-à-vis iWARP of the Intel Scalable Xeon platform announcement?

The Intel Scalable Xeon server platform announced on July 11, 2017 provides integrated support of 4 10Gbps iWARP RDMA ports supporting software-defined storage, NVM Express over Fabrics, virtual machine migration, and other high-throughput, low-latency applications.

What is the significance of integrated support of iWARP in Scalable Xeon platform?

- The announcement represents the first instance of Ethernet RDMA networking as an integrated component (as opposed to an add-on adapter) of a server chipset. iWARP is now “free” and “default” with the server purchase.
- It is very significant that Intel chose to integrate the iWARP protocol versus Ethernet RDMA alternatives into the Scalable Xeon platform. This represents strong validation for the ease-of-use, scalability and efficiency advantages of iWARP. This will allow decoupling of switch and server sales from each other since there are no longer special switch requirements.
- Integrated iWARP support in Scalable Xeon heralds pervasive support of the industry-standard RDMA protocol across enterprise and cloud datacenters, dramatically expanding the market for iWARP-enabled hardware and software solutions.

Is the new Intel iWARP feature interoperable with Chelsio T6 iWARP solution?

iWARP is an industry-standard protocol. As such as well as based on extensive testing, Chelsio’s Terminator 5/6 (T5/6) based adapters (which offer an industry-leading, robust fourth-generation implementation of iWARP over 1/10/25/40/50/100Gb Ethernet) are fully interoperable with Intel 10Gb Ethernet in the Scalable Xeon platform.

What are the implications of Intel’s iWARP announcement on the market for iWARP-based offerings?

We expect that the integrated support for iWARP in Scalable Xeon will dramatically expand the market and customer deployments for iWARP.

What is the specific impact of Intel’s new iWARP capability on iWARP deployment for storage applications?

Intel’s announcement positions iWARP to be the de facto standard for high-performance, low-latency storage network applications, including Windows Server 2016 Storage Spaces Direct (S2D) and NVMe over Fabrics (NVMe-oF).

What are the implications of Intel’s iWARP announcement on Chelsio T5/T6 iWARP adapter solution?

Intel’s iWARP announcement complements Chelsio’s iWARP product direction and solution portfolio in several very strong aspects.

- It validates Chelsio's pioneering product directions in the areas of TCP/IP Offload Engine (TOE) and iWARP spanning 15+ years.
- It dramatically expands the market for iWARP directly benefiting deployments of industry-leading Chelsio 1/10/25/40/50/100GbE iWARP product portfolio.
- Chelsio T5/T6 1/10/25/40/50/100GbE iWARP solution offers upgrade and interoperability options for end-users, integrators and OEMs deploying Intel 10GbE iWARP.

Why would a customer want to consider the Chelsio iWARP offering over that of Intel?

There are several scenarios where Chelsio iWARP solution offers value compared to the Intel Scalable Xeon iWARP offering:

- Intel iWARP in Scalable Xeon is limited to 10Gbps Ethernet support, whereas Chelsio iWARP product portfolio supports 1/10/25/40/50/100Gbps Ethernet. Thus, Chelsio iWARP offering provides an upgrade path as well as support for applications requiring higher performance than enabled by the Intel iWARP feature.
- Intel 10G iWARP offering is limited to supporting the new Scalable Xeon scalable platform, while Chelsio iWARP solution can be deployed in servers using earlier generation Xeon processors (e.g. "Sandy Bridge" Xeon processors) as well as the latest-generation of Scalable Xeon processors, and other non-x86 platforms.

References

[Intel Scalable Xeon Platform Product Brief](#)

[Burgeoning Intel Xeon SP Lewisburg PCH Options Overview](#)