TECH BRIEF

ATTO

Technical Features

- Single– and dual– port configurations
- Optical connectors SFP28(25GbE), QSFP+(40GbE), and QSFP28(50/100GbE) modules included in -000 models
- Up to 100Gb/s throughput per port
- High-performance x16 PCIe 3.0 bus
- Low profile form factor
- Driver support for Windows, Linux, and macOS operating systems.
- RDMA/NVMe over Fabrics support for Windows and Linux
- RDMA over Converged Ethernet (RoCE) enables industry-leading low-latency (1us MPI ping latency) and decreases CPU utilization
- TCP/UDP/IP hardware-based stateless offloads
- SR-IOV technology dedicates adapter resources for VMs within servers
- Guaranteed bandwidth and low-latency services
- Energy Efficient Ethernet enables lowest power draw on the market for a 40GbE NIC

Maximize High-Speed Ethernet Bandwidth with RDMA

With 4K fast becoming the standard for digital projection in movie theaters, and TVs also now supporting 4K, consumer demand for higher-resolution content is on the rise. As a result, more content is being shot and edited in both 4K and 8K resolutions. This development in turn has created a need for high-bandwidth 4K and 8K video workflows with throughput requirements that exceed what is possible with 10 Gigabit Ethernet (GbE).

While 4K and 8K footage can be compressed for storage, media files are typically streamed uncompressed for editing, rendering and special effects work. The demand to support these high-bandwidth video streams is currently driving heavy interest in SmartNICs capable of lossless transport of up to 100GB/s.

Do All SmartNICs Perform the Same?

It is important to note that throughput is at least partially a function of latency, or the lack thereof. In order to achieve the throughput required to support multiple uncompressed streams of 4K or 8K video, latency must be carefully managed since data transfer from computer memory through the CPU and Ethernet TCP/IP stacks introduces significant overhead.

While a latency effect was prevalent with 10GbE technology, most users were still able to achieve 800-900MB/s of throughput via TCP—enough for almost a full 4K-Full DPX stream. With these newer network standards (25/40/50/100GbE), however, the need for latency management increases exponentially. Without proper understanding of this limitation, a user expecting full throughput will be disappointed by the number of video streams the system can support, not to mention the discovery that their high-speed Ethernet connection lacks sufficient bandwidth for uncompressed 8K video.

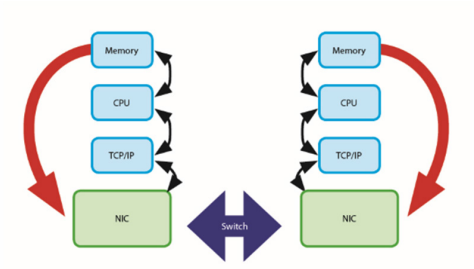

Figure 1 (Red = RoCE transfer; black = TCP transfer)

About ATTO

For over 30 years, ATTO has been a global leader across the IT and media & entertainment markets, specializing in storage and network connectivity and infrastructure solutions for the most data-intensive computing environments. ATTO works with partners to deliver end-to-end solutions to better store, manage and deliver data. Working as an extension of customer's design teams, ATTO manufactures host bus adapters, network adapters, storage controllers, Thunderbolt™ adapters, and software. ATTO solutions provide a high level of connectivity to all storage interfaces, including Fibre Channel, SAS/SATA, iSCSI, Ethernet, NVMe, NVMe over Fabrics and Thunderbolt. ATTO is the Power Behind the Storage.

How Can I Manage Latency with ATTO Smart NICs?

Remote Direct Memory Access (RDMA) is a feature designed to enable memory-to-memory data transfers. It has the benefit of eliminating the intermediate step of utilizing the CPU to transfer data, which decreases latency (and, as a result, increases throughput) while freeing up the CPU to conduct other tasks. There are two RDMA iterations on the market: RDMA over Converged Ethernet (RoCE) and iWARP. In addition to bypassing the CPU during data transfers, the more common version, RoCE, also bypasses the TCP/IP network stacks to further reduce latency and enhance throughput (see Figure 1). The second version, iWARP, operates on similar principles as RoCE, but adds several crucial nanoseconds of latency since it does not bypass the TCP/IP stack when transmitting data.

In a study conducted by NASA's Advanced Supercomputing Division (Boscia & Sidhu, 2013) comparing 40GbE RDMA to 40GbE TCP, a sustained rate of 21Gb/s was the best case scenario for a 40GbE card with a measured 6-7 microseconds of latency over TCP. A 40GbE card configured with RoCE, in contrast, dropped latency below 1 microsecond. This yielded a sustained rate of up to 39Gb/s — a nearly 100% performance improvement. What is the practical upshot of this increased bandwidth utilization? The ability to enhance 4K or 8K video workflow efficiency by supporting additional streams of video.

Video Format	Bandwidth per stream	Streams per 40GbE TCP	Streams per 40GbE RoCE
4K DPX	6.4Gb/s	3	6
4K-Full DPX	9.2Gb/s	2	4
4K-Full EXR	14.7Gb/s	1	2
8K DPX	23.4Gb/s	0	1
8K EXR	28.2Gb/s	0	1

The impact of RoCE for 40GbE vs. TCP in the number of supported video streams

How Do I Implement RDMA/RoCE?

Remote Direct Memory Access (RDMA) is a feature designed to enable memory-to-memory data transfers. It has the benefit of eliminating the intermediate step of utilizing the CPU to transfer data, which decreases latency (and, as a result, increases throughput) while freeing up the CPU to conduct other tasks. There are two RDMA iterations on the market: RDMA over Converged Ethernet (RoCE) and iWARP. In addition to bypassing the CPU during data transfers, the more common version, RoCE, also bypasses the TCP/IP network stacks to further reduce latency and enhance throughput (see Figure 1). The second version, iWARP, operates on similar principles as RoCE, but adds several crucial nanoseconds of latency since it does not bypass the TCP/IP stack when transmitting data.

In a study conducted by NASA's Advanced Supercomputing Division (Boscia & Sidhu, 2013) comparing 40GbE RDMA to 40GbE TCP, a sustained rate of 21Gb/s was the best case scenario for a 40GbE card with a measured 6-7 microseconds of latency over TCP. A 40GbE card configured with RoCE, in contrast, dropped latency below 1 microsecond. This yielded a sustained rate of up to 39Gb/s — a nearly 100% performance improvement. What is the practical upshot of this increased bandwidth utilization? The ability to enhance 4K or 8K video workflow efficiency by supporting additional streams of video.

ATTO FastFrame SmartNICs with RoCE

ATTO Technology Inc. single- and dual-port Smart NICs feature industry-leading latency management with RoCE. With FastFrame™ 40GbE NICs, users can expect significantly reduced CPU utilization/latency and uncompromised 40GbE bandwidth. Given ATTO's long experience in developing solutions optimized for media & entertainment, our FastFrame Smart NICs provide the premier Ethernet connectivity option for demanding 4K and 8K video workflows.