Intel® Data Direct I/O Technology Overview

Data center network bandwidth requirements continue to rise, driven by growth in server virtualization, unified data and storage networking, and other bandwidth-intensive applications. Many IT organizations are deploying 10 Gigabit Ethernet (10GbE) to meet these increasing needs. While 10GbE provides greater, more cost-effective bandwidth, it also places a greater burden on server resources. Intel® Data Direct I/O Technology (Intel® DDIO), an innovative, platform-wide approach from Intel, redefines server I/O to deliver substantial improvements in system bandwidth, power usage, and latency.

Inefficiencies of the Classic Approach to Server I/O

The "classic" approach to server I/O requires numerous reads to and writes from memory and dates from the era when I/O was slow and processor caches small. The classic model made perfect sense in that context, but the world has changed; 10GbE and faster I/O data rates are now common, and the Intel® Xeon® processor E5 family supports up to 20MB of last-level cache.

In this model, transferring data from an I/O device to the host for an inbound data flow (Figure 1) involves a direct memory access (DMA) operation directly into the host processors' main memory, bypassing the processor cache (1). When the processor is ready to operate on the data, it copies the data into cache from memory (2), requiring both a memory read and a memory write.

Transferring data from the host to an I/O device for an outbound data flow requires even more memory accesses, as shown in Figure 2. When the host processor creates a buffer in cache for the data, it triggers a read from memory (1). Once the data is created, the host processor informs the I/O device that it can consume this data. At this point, the I/O Hub triggers a speculative second read from main memory because it does not know for certain that the host has already done so. The data will then go from cache to the I/O device (2), but as it is evicted from cache, the data will be written back to main memory (3).

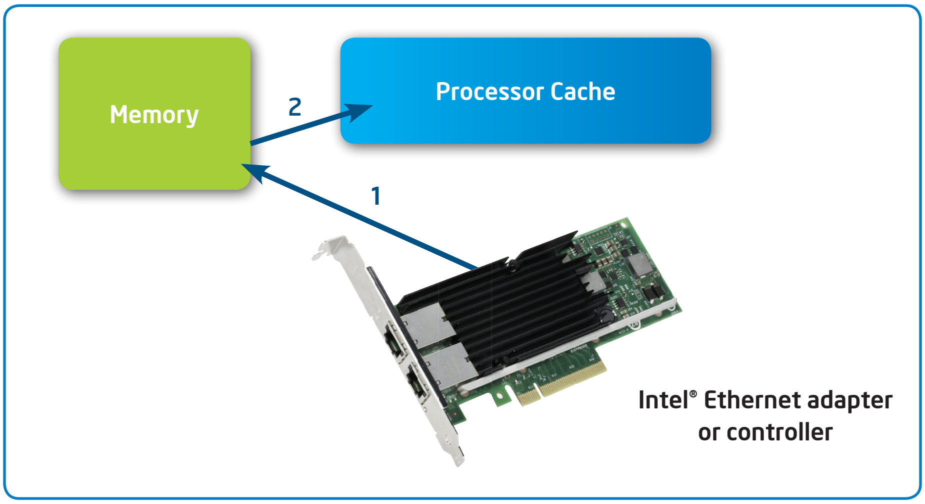

Figure 1. Classic model inbound data flow

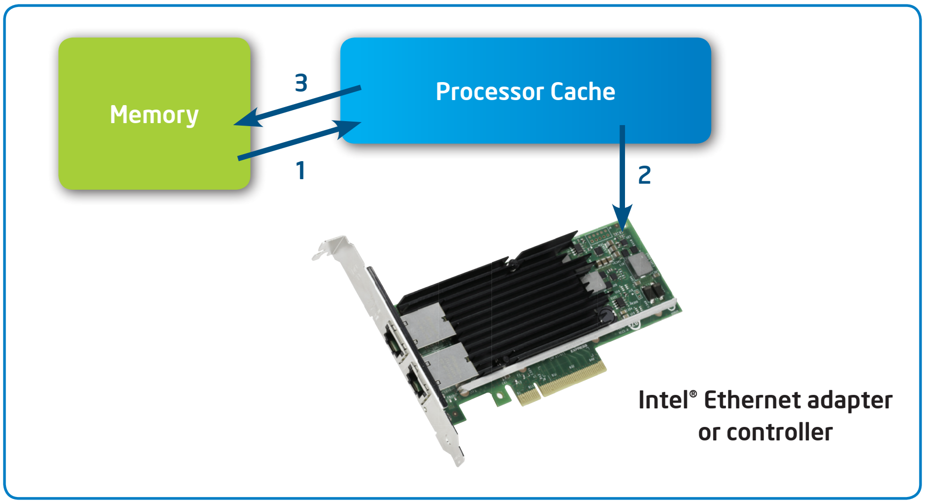

Figure 2. Classic model outbound data flow

These frequent memory accesses degrade system performance and increase system power consumption, highlighting the inefficiencies of the approach in today's context.

Intel DDIO: A Platform-wide Approach to Improving I/O

With Intel DDIO, Intel has rearchitected the Intel Xeon processor E5 family to remove the inefficiencies of the classic model by enabling direct communication between Intel® Ethernet controllers and adapters and host processor cache. With Intel DDIO, an Intel® Ethernet Converged Network Adapter or controller talks directly to the processor's last-level cache, as shown in Figure 3. Eliminating the frequent visits to main memory present in the classic model reduces power consumption, provides greater I/O bandwidth scalability, and lowers latency.

Intel DDIO allows the host processor to operate on inbound data immediately, without invoking any memory accesses. This advancement significantly lowers latency, as the processor no longer has to wait for memory accesses before beginning work, and lowers power consumption by not engaging main memory. A write to main memory will only occur upon eviction from the highest-level cache. Intel DDIO allows inbound I/O to use up to one-tenth of the Intel Xeon processor E5's L3 cache. If this limit is exceeded, new inbound I/O will continue to go directly to cache, but the least-recently used I/O will be written to memory to make room for the new data. For outbound flows, data in the cache is simply transferred to the I/O device without memory reads or writes.

Performance Benefits

Intel Ethernet controllers and adapters, with their high performing, stateless architecture excel with Intel DDIO. Tests conducted by Intel show significant performance benefits in Intel Xeon processor E5-based systems with Intel DDIO enabled and Intel Ethernet 10 Gigabit Converged Network Adapters installed.

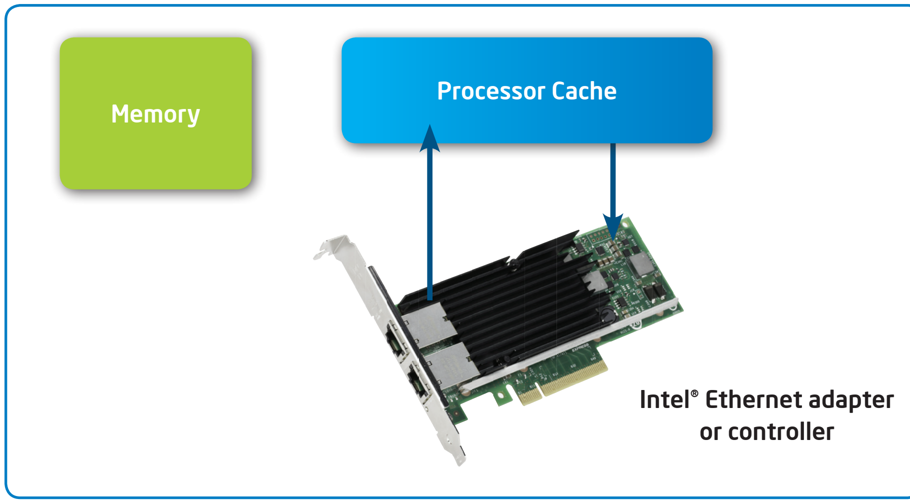

Figure 3. Intel DDIO allows the network device to talk directly to host processor cache

Increased system bandwidth: High-bandwidth applications that tax I/O performance, such as telecom, data plane embedded, and network appliances, can see dramatic performance improvements from Intel DDIO, though specific benefits will vary, depending on workload.

Figure 4 shows L2 forwarding performance on an Intel Xeon processor E5-based server with Intel DDIO enabled and an Intel Ethernet Converged Network Adapter X520 installed achieving more

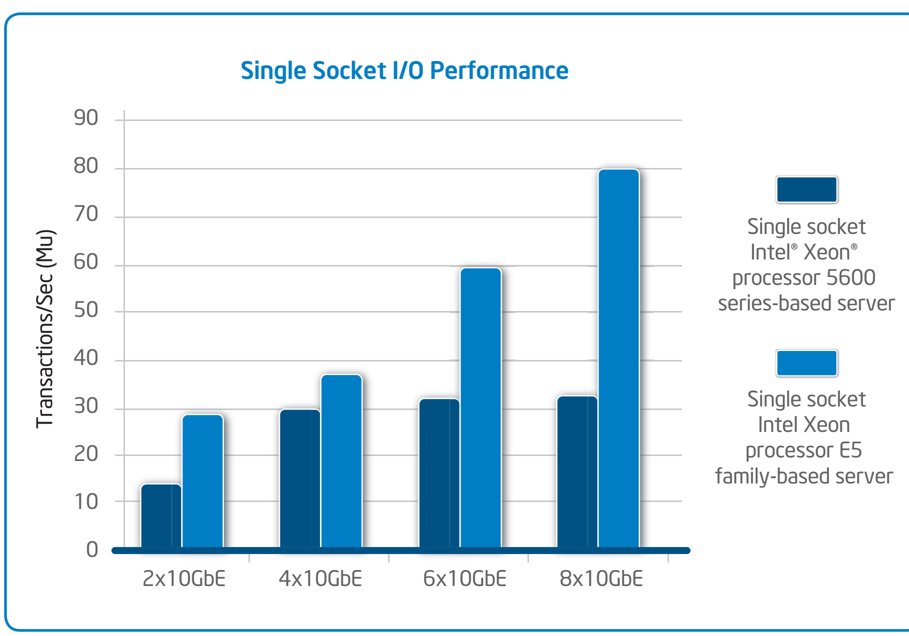

Figure 4. L2 forwarding performance comparison[1]

than twice the bandwidth of a previous-generation server. Note that bandwidth scales as additional 10GbE ports are added to the test machine.

Decreased power consumption: Typical data center workloads, which are not I/O-bound or response-time critical, will benefit primarily from reduced power consumption and modest performance gains. These improvements result from decreased memory accesses, as Intel DDIO essentially eliminates memory utilization for I/O workloads.

Figure 5 summarizes throughput and power utilization results in tests conducted using Intel Xeon processor E5 family-based servers, an Intel Ethernet Converged Network Adapter X520, and VMware vSphere* 5. As shown in the graph on the left, the system with Intel DDIO enabled slightly outperforms the system with Intel DDIO disabled. The graph on the right shows that for a similar level of performance, the Intel DDIO-enabled system uses seven watts less power.

Lower latency: Applications for which response time is critical, including financial trading and high-performance computing, will benefit from the lower latency delivered by Intel DDIO and will experience substantially higher transaction rates. Tests conducted by Intel have shown a reduction in latency of approximately 1.0 μsec on Intel Xeon processor E5 family-based servers, which include Intel DDIO.

A Standard Platform Feature

Intel DDIO is a key component of Intel® Integrated I/O, a suite of architectural enhancements introduced with systems based on the Intel® Xeon® processor E5 family. In addition to Intel DDIO, Intel Integrated I/O incorporates support for the PCI Express* 3.0 interface into the Intel Xeon processor, eliminating the need for an I/O hub and improving system I/O performance.

Intel DDIO is a standard feature on all systems based on the Intel Xeon processor E5 family and is supported by the Intel Ethernet 10 Gigabit Converged Network Adapter family. No changes to drivers, applications, operating systems or firmware are required to realize the benefits of Intel DDIO.

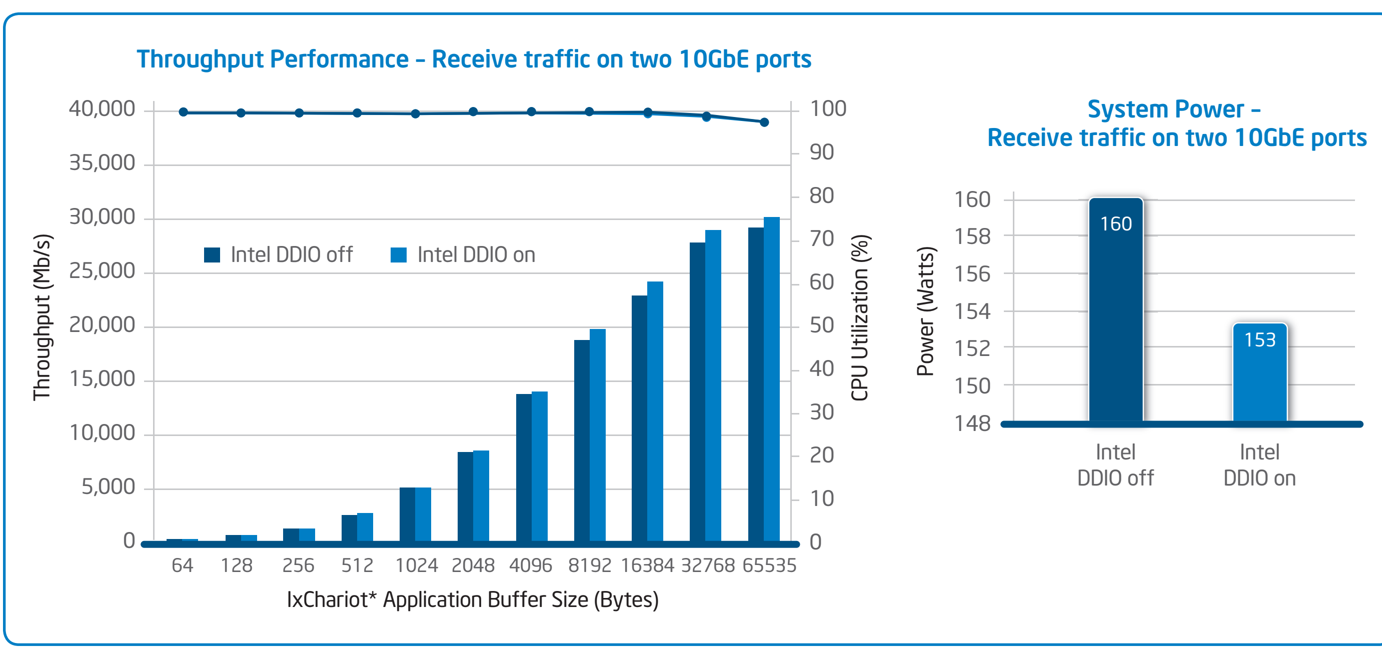

Figure 5. Intel DDIO reduces system power consumption[2]

[1] Source: Intel internal measurements showing performance comparison between Intel® Xeon® processor E5 family and Intel® Xeon® processor 5600 series. Configurations: Pre-release Intel Xeon processor E5 (8 core, 2.8GHz (2.7GHz + turbo), 8x2 GB DDR3-1333 MHz, 1-4x Intel Ethernet Converged Network Adapter X520-DA2, Linux 2.6.32, Ixgbe 2.0.94 (stackless driver with data touch); and Intel® Xeon® processor 5680 (6 core, 3.3 GHz), 6x2 GB DDR3-1333 MHz, 1-4x Intel Ethernet Converged Network Adapter X520-DA2, Linux 2.6.32, Ixgbe 2.0.94 (stackless driver with data touch). Intel Ethernet Labs, April 2011.

[2] Performance comparison using the Chariot 7.1 application. Configuration: Intel Xeon processor E5 family (8 core, 2.6 GHz), 8x2 GB DDR3-1333 MHz, 1 Intel Ethernet Converged Network Adapter X520-DA2, VMware ESX 5.0, ESX 5.0 inbox driver. Intel Ethernet Labs, April 2011.

0612/BY/MESH/PDF 327480-001US