

Intel® 10GbE Adapter Performance Evaluation for FCoE and iSCSI

Evaluation report prepared under contract with Intel® Corporation

Introduction

As the interest in converged networks grows, and as the vendors from the traditional Ethernet adapter and Fibre Channel adapter marketplaces produce converged products, questions arise about the capabilities and performance of each type of adapter. Intel® commissioned Demartek to compare the performance of their Intel® Ethernet Server Adapter X520 that supports Ethernet and Fibre Channel over Ethernet (FCoE) protocols with those of the two leading Fibre Channel adapter vendors.

Evaluation Environment

A series of application performance workloads were run and repeated for each of the three competing converged adapters connecting servers to storage in a Data Center Bridging (DCB)/FCoE environment. The goal of this testing was to evaluate products in environments similar to actual customer environments. As a result, these tests were performed with well-known disk storage arrays with spinning disk drives in FCoE and iSCSI configurations similar to those typically found in customer datacenters, rather than testing theoretical performance with specialized hardware not typically found in customer environments.

Evaluation Summary

In examining the results of these comprehensive performance tests, we found that for the most part, the performance of all three brands of adapters fell into a fairly narrow range. In some tests, one adapter had the highest performance; in other tests, a different adapter had the highest performance.

Because the adapter performance was reasonably close in most of these tests, IT professionals need to consider the cost of these adapters, especially in environments where many adapters are required. In addition, Ethernet features and functions in converged adapters should be considered. For these tests, the Intel® adapter provides the best price for performance in these real-world application tests. In some environments, the additional cost of more expensive adapters could be applied towards faster CPUs or more CPU cores that would benefit all applications running on that host.

1 – Converged Networks and Adapters

“Converged” or “unified” networks have been garnering considerable attention recently, especially as the 10-Gb/sec Ethernet and Fibre Channel over Ethernet (FCoE) standards have been made official within the last calendar year¹. However, the measured fusion of the Ethernet and Fibre Channel worlds triggers many questions among datacenter directors, managers and administrators. They seek answers regarding the technology, planning, implementation, configuration, performance and price, among others.

One of the components of a converged infrastructure is the adapter in the server and the concept of a single adapter that handles both 10Gb/sec Ethernet and 10Gb/sec Fibre Channel storage traffic. These adapters take advantage of the updated version of Ethernet known as Data Center Bridging (DCB) and its ability to carry multiple types of traffic concurrently, in a lossless fashion, giving each traffic type its own priority-based flow control.

For many, running at 10Gb/sec for Ethernet or Fibre Channel storage is a new breakthrough in performance. Because of the fusion of Ethernet and Fibre Channel, it is not surprising that vendors who previously provided only Ethernet adapters and vendors who previously provided only Fibre Channel adapters are now providing converged adapters that support both Ethernet and Fibre Channel. Datacenter professionals now have to consider more vendors than they previously considered for Ethernet NICs and Fibre Channel HBAs, respectively.

This report is a real-world evaluation of the performance of some of these adapters provided by vendors from “both camps” as it were, the Ethernet adapter vendors and the Fibre Channel adapter vendors.

Two Approaches


There are two approaches to providing a converged or unified adapter. One approach is to use proprietary hardware adapters with offloaded FC and FCoE (and in some cases TCP/IP) protocols embedded in the hardware. This is the traditional way that Fibre Channel adapter vendors provide their solutions. They include their own drivers, interfaces and management software.

The second approach, which Intel® has chosen, is to take advantage of native FCoE initiators in operating systems and build the adapter to work in a complementary way with the platform hardware and operating system to enable FCoE traffic, all at lower cost than competitive adapters. Intel® believes that native FCoE operating system support will develop similar to the way that iSCSI support has developed with native initiators in operating systems. With current multi-core Intel® processor-based platforms able to sustain two ports of 10Gb/sec Ethernet in these environments at well under 10% CPU utilization, there is plenty of headroom for the Intel® approach.


¹ Comparisons and roadmaps for Ethernet, FCoE and other storage interfaces are available at:
http://www.demartek.com/Demartek_Interface_Comparison.html.

2 – Intel® Installation and Management Interface

Installation of the Intel® Ethernet Server Adapter X520 is performed in the same manner as other Intel® Ethernet adapters. The software and drivers use the standard Intel® installation process known as Intel® Network Connections, and is available via download or on a distribution CD.



A new option in the installation process allows for the installation of the DCB/FCoE features. This is not checked by default, but was used for this installation.




Intel® uses the same management interface for its Intel® Ethernet Server Adapter X520 as it does for its other Ethernet adapters. The Ethernet settings are in the same places administrators have come to expect, while the FCoE settings have their own page. Each are accessible through the familiar “Device Manager” management console under the “network adapters” and “storage controllers” sections, respectively. The FCoE properties can also be accessed from the Ethernet properties page under the “Data Center” tab.

Each port of the Intel® Ethernet Server Adapter X520 is managed separately for both the Ethernet and FCoE parameters. Each port on an adapter can have different settings.

Ethernet


The Ethernet management properties include several tabs for managing various Ethernet functions such as link speed, NIC teaming, advanced settings and more. The advanced settings include several Ethernet parameters relating to checksum offload, Receive Side Scaling (RSS) and other functions. Information about each of these functions is provided in the lower half of the properties window to help administrators make the best choices. In many cases, the default setting is the best choice. Examples are shown below.

Diagnostics are available that check the health of the connection and the hardware under the “Link Speed” tab.




DCB/FCoE

The DCB/FCoE properties can be accessed from the Ethernet properties by using the “Data Center” tab and clicking on the “FCoE Properties” button, as shown below.



The FCoE properties show parameters such as the World-Wide-Names (WWNs), link speed, and other basic Fibre Channel parameters that are typical in this environment. In addition, some parameters can be changed, such as the device I/O queue depth.

The “Port Properties” button launches the Ethernet properties window that has already been discussed.



3 – Real-world Adapter Testing

For this evaluation, we performed a spectrum of application performance tests using real-world configurations of servers and storage that customers actually use in production. We used two different, well-known enterprise storage targets that use a large number of spinning disks configured as they are typically configured in customer environments. We tested single-port and dual-port adapter configurations because both are used in production environments and we wanted to ensure that the adapters performed well in both configurations. We also tested iSCSI and FCoE performance for these adapters.

All the Data

Further, we are publishing all the results of the tests, not cherry-picking a few results that might look favorable to one particular product or vendor. In the real world, there are many different configurations used in datacenters and we want to show as much data as possible in order to represent as many environments as possible.

Converged Adapters Tested

We chose current adapters from three of the most popular Ethernet and Fibre Channel adapter vendors in today's market.

- Emulex OCE10102-FM UCNA
- Intel® Ethernet Server Adapter X520-SR2
- QLogic QLE8142 CNA


These adapters represent the Ethernet vendor perspective and the Fibre Channel vendor perspective.

These adapters were tested in single-path and multi-path configurations in order to show different ways that customers might choose to deploy these adapters, and to provide some general expectations of performance in these two configurations. The multi-path tests were configured with load balancing and failover functions in mind.

Pricing

In addition to the basic functions of the adapters, IT managers and administrators are also interested in the price of these products. The chart shows the list pricing for these brands of adapters for copper and optical connections. We tested the optical versions of these adapters.

These prices are for quantities of one adapter. Pricing may vary by supplier, quantity and any negotiated discounts.



4 – Evaluation Environment

The tests were conducted at the Demartek lab in Arvada, Colorado and the Intel® lab in Hillsboro, Oregon.

Basic functionality tests were performed in the Demartek lab. The performance tests were performed in the Intel® lab.

Servers:

- Demartek lab: Supermicro X8DTH-6F, Dual Intel® Xeon processor E5540, 2.53 GHz, 5.86 QPI, 48GB RAM, Windows Server 2008 Enterprise.
- Intel® lab: Dell PowerEdge R710, Dual Intel® Xeon processor X5550, 2.67 GHz, 6.40 QPI, 24GB RAM, Windows Server 2008 Enterprise.

Switches:

- Demartek lab: Cisco Nexus 5020, 10GbE, supporting DCB/FCoE with 4Gb FC blade
- Intel® lab: Cisco Nexus 5020, 10GbE, supporting DCB/FCoE with 4Gb FC blade

Storage targets:

- Demartek lab: NetApp FAS3040
- Intel® lab: NetApp FAS3170, EMC CX4
 - **NetApp FAS3170:** Dual-controller, 4xAMD Opteron processors, 16GB RAM, four disk shelves populated with 136GB, 15K RPM, FC disk drives for a total of 56 disk drives, 7TB raw capacity. One QLogic QLE8152 dual-port FCoE CNA. One dual-port 10G Ethernet Controller T320E. Disk storage configures as RAID-DP groups.
 - **EMC CX4:** Dual-controller, 16GB RAM, four disk shelves populated with 45x133.7GB, 15K RPM FC disk drives and 15x268.4GB, 15K RPM FC disk drives for a total of 60 disk drives, 10TB raw capacity. 4x4Gb FC host ports. Disk storage configured as RAID10 groups.

5 – Performance Results

Two sets of tests, Microsoft Exchange JetStress 2007 and Microsoft SQLIO, were run using iSCSI and FCoE configurations. Both performance and CPU utilization are important, so we tracked both. The fundamental performance is critical to determine if the storage can support the application at a particular workload level. The CPU utilization shows how much load that workload is placing on the CPU.

Performance is a vital measurement when testing converged adapters because some adapters can sustain higher performance than others for various workloads. The CPU utilization is important to measure because this shows the load that the adapter driver and interface place on the platform and will have an effect on the number of simultaneous applications or virtualized operating systems that can be run on a given platform.

JetStress

Microsoft Exchange Jetstress 2007 simulates the Exchange Server disk input/output (I/O) load. This tool, provided by Microsoft, verifies the performance and stability of a storage subsystem and its suitability for Microsoft Exchange Server. Jetstress is generally used by customers before deploying Exchange Server to ensure that the storage subsystem can sustain a production workload.

The Jetstress configuration we ran represents a company with 5000 employees. These JetStress configurations used 5000 mailboxes of size 125MB using the “heavy” user profile. There were 8 storage groups spread across 8 LUNs on the storage arrays, all accessed simultaneously. The JetStress tests were run for a minimum of 2 hours each. Several different thread counts were specified for each test run, showing increasing workloads.

SQLIO

SQLIO is a tool provided by Microsoft that can be used to determine the I/O capacity of a given workload. SQLIO can run I/O workloads of various types and is often used by SQL Server administrators to test a storage subsystem before deploying a production SQL Server application.

We tested 8KB random reads and random writes, which represent OLTP workloads, and we tested several block sizes of sequential reads and writes that represent workloads such as batch processing, database backups, database log files, some decision support systems, etc. We varied the queue depth (number of simultaneous I/O requests) and the thread count (number of “workers” issuing I/O requests) in order to represent different customer environments and different workloads, from light to heavy. Though we varied the thread count for these tests, we are reporting the results for one representative set of thread counts in order to save space in this report. All the SQLIO tests accessed 8 LUNs simultaneously across the storage arrays.

Performance Data Comments

The following pages provide the performance results for each of the tests. The data shown are the results of the specific application workloads. Exchange JetStress output provides IOPS and CPU utilization. SQLIO output provides IOPS and latencies.




In most cases, the performance numbers were fairly close for all three adapters. In some cases, one adapter performed better and in other cases a different adapter performed better.




The SQLIO output includes thread count ("T") and queue depth ("Q") for each data point.


FCoE vs. iSCSI

One interesting result of most of these tests was that the performance of the single-path FCoE and iSCSI connections to the NetApp storage yielded very similar results. The multi-path FCoE configuration performed 10%-15% higher than the same iSCSI multi-path configuration.

We used the Microsoft iSCSI software initiator on the host server for all the iSCSI testing for each of the adapters.


JetStress – NetApp FCoE Single Path

JetStress – NetApp iSCSI Single Path

JetStress – NetApp FCoE Multi-Path



JetStress – NetApp iSCSI Multi-Path

JetStress – EMC FCoE Multi-Path

JetStress – EMC FCoE Single Path



SQLIO - NetApp FCoE Single Path


SQLIO - NetApp iSCSI Single Path




SQLIO - NetApp FCoE Multi-Path


SQLIO - NetApp iSCSI Multi-Path


SQLIO - EMC FCoE Multi-Path


SQLIO - EMC FCoE Single Path



Summary and Conclusion

In running a fairly broad set of application tests that varied parameters such as queue depth and number of threads, we attempted to reproduce many of the types of environments the customers have in their production environments. The performance of the three competing adapters in these tests was fairly close and as a result, IT professionals need to consider other aspects including the price of the adapters.

We found that for these real-world application tests, the Intel® Ethernet Server Adapter X520 provides good overall performance at a great price point, for both copper and optical implementations. This adapter fits well into an existing infrastructure that includes other Intel Ethernet adapters.

Intel is a registered trademark of Intel Corporation.

Emulex is a registered trademark of Emulex Corporation.

QLogic is a registered trademark of QLogic Corporation.

Demartek is a registered trademark of Demartek, LLC.

All other trademarks are the property of their respective owners.