10Gb Intel® Ethernet Converged Network Adapter Evaluation

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 Converged Network Adapters (CNA) that support storage over Ethernet protocols including NFS, iSCSI, and FCoE. Intel® commissioned Demartek to compare the performance of CNAs from two leading Fibre Channel adapter vendors and the Intel® Ethernet Converged Network Adapter X520 in virtual server environments with various numbers of guest virtual machines.

Evaluation Environment

A variety of application performance workloads were run on each of these three Converged Network Adapters connecting servers to storage targets using FCoE. These workloads were run with various numbers of guests running in a VMware vSphere 5.0 and 5.1 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 data centers, rather than testing "micro-benchmark" 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 CNAs fell into a fairly narrow range. In some tests, one CNA had the highest performance; in other tests, a different CNA had the highest performance.

Because the CNA 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 CNAs should be considered. Lastly, is the CNA supported with robust software, or does it require extensive firmware and driver updates with major updates to system software? For these tests, the Intel® Ethernet CNA provides the best price for performance in these real-world application tests. The additional cost of a CNA that offers an offload for only FCoE or iSCSI traffic type could be applied to the purchase of better CPUs that would benefit not only FCoE and iSCSI but all the applications and traffic types.

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 in recent years following the success of iSCSI as a storage-over-Ethernet SAN protocol. However, the fusion of the Ethernet and Fibre Channel worlds triggers many questions among data center 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 port that handles 10Gb/sec of both Ethernet and Fibre Channel storage traffic. These CNAs take advantage of a new collection of standards-based end-to-end networking technologies that make Ethernet the unified fabric for multiple types of traffic in the data center known collectively as Data Center Bridging (DCB). The new standards provided the ability for Ethernet to carry multiple types of traffic concurrently, in a lossless fashion, giving each traffic type its own priority-based flow control. DCB was architected so that Ethernet could meet Fibre Channel's requirements for a lossless transport and timing requirements.

In the past, vendors of Fibre Channel solutions were isolated from the world of Ethernet. But the trends toward unified networking started long ago with network file systems in the 80's and iSCSI in the late 90's for storage connectivity using Ethernet. iWARP showed how RDMA could be effectively delivered over Ethernet as well, rather than a dedicated HPC fabric. These successes have driven interest in supporting the FC protocols over Ethernet as well and led to the T11 Fibre Channel over Ethernet (FCoE) specifications. That trend caused the traditional FC HBA vendors to enter the Ethernet world with their own solutions for Ethernet-based solutions for LAN, iSCSI, and FCoE. Data center professionals now have to consider different vendors than they previously considered for Ethernet NICs and Fibre Channel HBAs, respectively.

Two Approaches

There are two approaches to providing converged network adapters. One approach is to use proprietary hardware adapters with full protocol offload of FC and FCoE (and in some cases offloaded TCP/IP). 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 calls "intelligent offloads" and has chosen with the Intel® Ethernet Converged Network Adapter X520, 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®

Xeon® 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. In fact, with the Intel® Xeon® processor E5-2600 family, a single dual-processor server can support up to 16 ports of 10GbE.



2 – Real World CNA Testing

For this evaluation, we performed a very large amount of application performance tests using realworld configurations of servers and storage that customers actually use in production. We used 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 of the results for this report were run in a virtual machine environment, using VMware vSphere 5.0 and vSphere 5.1. These results cannot be directly compared to results from non-virtualized environments, as virtualized and non-virtualized environments have different operating overhead factors.

Converged Network Adapters Tested

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

- Emulex OCE11102-FM UCNA
- Intel® Ethernet Converged Network Adapter X520-DA2 and X520-SR2
- QLogic QLE8242 CNA

QLogic Incompatibility with Intel® Server System and VMware vSphere Client

During the initial phase of this testing in the first half of 2012, we found an incompatibility with the Intel® Server System P4308CP4MHGC and the QLogic QLE8242 CNA in the VMware vSphere environment. Thus, FCoE results could not be obtained for the QLogic CNA. There is a support note regarding this issue on the QLogic website at https://support.glogic.com/app/answers/detail/a_id/1816/kw/1816.



Pricing

As enterprises consider converged network adapters, pricing plays a role. Budgets must be considered, especially where large numbers of adapters needed in an enterprise setting are needed. The chart below shows the retail pricing of these adapters from a well-known online retailer of enterprise IT equipment.



Effective September 30, 2012, Intel lowered the manufacturer's suggested retail price (MSRP) for many of the models of the Intel® Ethernet Converged Network Adapter X520. These lower prices are expected to make their way into the distribution channels within 30-60 days.

Optics vs. No optics

The adapters used in this series of tests support the SFP+ connector, and are available with and without optical transceivers, also known as "optics." Host adapters with "no optics" are less expensive than the equivalent adapter "with optics."

For 10GbE connections, there are two types of cables that can be used with the SFP+ connector in a host adapter or a switch. There are Direct Attached Copper (DAC) cables and fiber optic cables.

The 10GbE DAC cables have an SFP+ transceiver built onto the cable that slides into the SFP+ cage on the adapter or switch. The copper cables are used for shorter lengths, such as within a rack or to a nearby rack.

Fiber optic cables generally do not have the SFP+ transceiver attached to the cable. When using fiber optic cables, the SFP+ optical transceiver is mounted into the SFP+ cage on the adapter or switch port. Fiber optic cables are generally used for longer distance connections.

More information on cables, connectors and storage networking interfaces is available on the <u>Demartek Storage Interface Comparison</u> reference page.

3 – Performance Results

For this set of tests, we ran a comprehensive suite of tests including Microsoft Exchange Jetstress 2010, Microsoft SQLIO, Oracle Orion and IOmeter. We ran many of these tests in a VMware environment with one, four, five, ten or twenty guests.

Exchange Jetstress 2010

Exchange Jetstress 2010 is an Exchange Server storage stress tool provided by Microsoft. It is used to measure the performance capabilities of a storage subsystem for Microsoft Exchange database and log files. The purpose of the Jetstress test process is to find the maximum workload that Exchange Server can support while still passing the test.

For Exchange Server 2010, Microsoft has eliminated the usage of the user profile terms "light", "medium" and "heavy" to describe the estimated IOPS per mailbox and simply uses a numeric value for the average number of messages sent and received in a day. Demartek has a reference page explaining some of the Jetstress parameters and comparing Exchange Server 2003, 2007 and 2010 with respect to I/O profiles at

www.demartek.com/Demartek Exchange 2003 2007 2010 I-O Comparison Summary.html.

Key metrics for Jetstress are:

- Achieved IOPS
- Average CPU utilization
- Average database reads latency (must be less than 20 milliseconds)
- Average log writes latency (must be less than 10 milliseconds)



FCoE











<u>iscsi</u>









ORION

Oracle Orion is a test tool that simulates Oracle database workloads.

The following data are the results of the OLTP workload with 4 guests, 8 disks per guest, connected to a native Fibre Channel storage system.





SQLIO

SQLIO is a tool provided by Microsoft that can be used to determine the I/O capacity of a storage configuration. SQLIO is generally used to simulate database workloads in Microsoft environments.

The following data are the results of 8KB block size reads with 4 guests, 8 disks per guest, connected to an iSCSI storage system. For each block size, multiple queue depths were run.





IOmeter

IOmeter is an open source I/O workload generator. IOmeter can be configured to run various read/write mixtures, queue depths, block sizes and other I/O parameters. IOmeter provides I/O operations per second (IOPS), bandwidth, or MegaBytes per second (MBPS) and CPU utilization data.

Hundreds of tests were run with various numbers of guest virtual machines. There were far too many test case results to include in this report, so a sampling of some of the results is shown here. Note that there are multiple runs shown for each adapter configuration. There is variation in the results with the same adapter, reflecting a busy server with busy storage systems.

We selected the results from twenty-guest results tests. These results are as viewed from the guest operating system perspective as provided by IOmeter itself. During the twenty guest runs, all the guests were performing similar IOmeter tests, placing plenty of load on the back-end storage systems. The IOmeter parameters for these results:

- 20 guests, 4 vCPU for each guest
- Queue depth = 4
- Blocks sizes: 512 byte 1 MB
- Guests connected to FCoE storage targets
- Separate runs:
 - o 100% sequential read
 - o 50% sequential read and 50% sequential write
 - o 100% sequential write

The results from each guest were taken and averaged together. Of special interest may be the CPU utilization, which differed only slightly among the adapters tested.

These IOmeter tests were run with VMware vSphere 5.1 with Windows Server 2008 R2 guests.



20 Guests - 4 vCPU









20 Guests - 4 vCPU









20 Guests - 4 vCPU







Apache Webserver

The Apache webserver test was run in a VMware environment with different numbers of guests on a variety of storage systems. These storage systems shown include the vendor, interface and number of guests configured. For example, "EMC-FC-10" refers to EMC Fibre Channel storage with 10 guests running.

The Apache webserver test shows the number of web requests per second that were performed by the web server. We observed mixed results with the CNAs, across the different storage configurations and interfaces.



Summary and Conclusion

In running a broad set of application tests that varied parameters such as queue depth and number of threads, we reproduced many of the types of environments the customers have in their production environments. The performance and overall CPU utilization of the three CNAs in these tests were fairly close and as a result, IT professionals need to consider other aspects such as the flexibility, simplicity and total price of the CNAs when making a purchase decision.

We found that for these real-world application tests, the Intel® Ethernet Converged Network Adapter X520 provides good overall performance at a great price point, for both direct-attached copper and optical implementations of iSCSI and FCoE connectivity. Because the storage protocols are "free" with the Intel CNA, IT professionals can deploy this product broadly for its LAN capabilities and selectively enable its iSCSI and FCoE capabilities where and when desired at no incremental expense.

Additional information regarding the Intel® Ethernet Converged Network Adapter X520 is available at www.intel.com/go/ethernet.

Appendix – Evaluation Environment

The servers used for this project included:

- Intel® Server System P4308CP4MHGC with Intel® Xeon® processor E5-2600
- Dell PowerEdge R720

The Switches used for this project included:

• Cisco Nexus 5010

The Storage targets used for this project included:

- EMC CLARiiON CX4 480 (FC protocol)
- Dell Compellent Series 40 Storage Array (FCoE protocol)
- Dell EqualLogic 6010 (iSCSI protocol)

Operating systems used for this project included:

- VMware vSphere 5.0 and vSphere 5.1
- Microsoft Windows Server 2008 R2
- Red Hat Enterprise Linux 6.1

The original version of this report is available at www.demartek.com/Demartek_Intel_X520_Performance_Evaluation_2012-11.html on the Demartek web site.

Intel and Xeon are registered trademarks of Intel Corporation in the United States and other countries.

Demartek is a registered trademark of Demartek, LLC.

All other trademarks are the property of their respective owners.