

Upgrading Data Center Network Architecture to 10 Gigabit Ethernet

Upgrading our network architecture will optimize our data center infrastructure to respond faster to business needs while enhancing the services and value IT brings to the business.

Matt Ammann
Senior Network Engineer, Intel IT

Carlos Briceno
Senior Network Engineer, Intel IT

Kevin Connell
Senior Network Engineer, Intel IT

Sanjay Rungta
Principal Engineer, Intel IT

Executive Overview

To accommodate the increasing demands that data center growth places on Intel's network, Intel IT is converting our data center network architecture to 10 gigabit Ethernet (GbE) connections. Existing 100 megabit per second (Mb/s) and 1 GbE connections no longer support Intel's growing business requirements. Our new 10 GbE data center fabric design will enable us to accommodate current growth and meet increasing network demand in the future.

Intel IT is engaged in a verticalization strategy that optimizes data center resources to meet specific business requirements in different computing areas. Data center trends in three of these areas drove our decision to upgrade, including server virtualization and consolidation in Office and Enterprise computing environments, and rapid growth in Design computing applications and their performance requirements. Additionally, we experience 40 percent per year growth in our Internet connection requirements.

While designing the new data center fabric, we tested several 10 GbE connection products and chose those that offered the highest quality performance and reliability. We also balanced ideal design against cost considerations.

The new data center fabric design provides many benefits:

- **Reduced data center complexity.** As virtualization increases, a 10 GbE network allows us to use fewer physical servers and switches.

- **Reduced total cost of ownership in a virtualized environment.** A 10 GbE fabric has the potential to reduce network cost in our virtualized environment by 18 to 25 percent, mostly due to simplifying the LAN and cable infrastructures. The new system also requires fewer data center space, power, and cooling resources.
- **Increased throughput.** Faster connections and reduced network latency provide design engineers with faster workload completion times and improved productivity.
- **Increased agility.** The network can easily adapt to meet changing business needs and will enable us to meet future requirements, such as additional storage capacity.

Upgrading our network architecture will optimize our data center infrastructure to respond faster to business needs while enhancing the services and value IT brings to the business.

Contents

Executive Overview.....	1
Background.....	2
Solution.....	2
Simplifying Virtualization for Office and Enterprise Applications ..	3
Increasing Throughput and Reducing Latency for Design Applications.....	4
Choosing the Right Network Components.....	4
Future Plans for Office and Enterprise I/O and Storage Consolidation	4
Conclusion.....	5
Acronyms.....	5

IT@INTEL

The IT@Intel program connects IT professionals around the world with their peers inside our organization – sharing lessons learned, methods and strategies. Our goal is simple: Share Intel IT best practices that create business value and make IT a competitive advantage. Visit us today at www.intel.com/IT or contact your local Intel representative if you'd like to learn more.

BACKGROUND

Intel's 97 data centers are at the center of a massive worldwide computing environment, occupying almost 460,000 square feet and housing approximately 100,000 servers. These servers support five main computing application areas, often referred to as "DOMES": Design engineering, Office, Manufacturing, Enterprise, and Services.

Intel's rapidly growing business requirements place increasing demands on data center resources. Intel IT is engaged in a virtualization strategy that examines each application area and provides technology solutions that meet specific business needs. We are also developing an Office and Enterprise private cloud, and we see opportunities to expand cloud computing to support manufacturing.

These strategies, combined with the following significant trends in several computing application areas, compelled us to evaluate whether our existing 1 gigabit Ethernet (GbE) network infrastructure was sufficient to meet network infrastructure requirements.

- Large-scale virtualization in the Office and Enterprise computing domains.
- Increasing compute density in the Design computing domain.

In addition, high-performance Intel® processors and clustering technologies are rapidly increasing file server performance. This means that the network, not the file servers, is the

limiting factor in supporting faster throughput. Our Internet connection requirements are growing 40 percent each year as well, which requires faster connectivity than is possible with a 1 GbE data center fabric.

SOLUTION

To meet these demands, we determined it was necessary to convert our data center network architecture from 100 megabit per second (Mb/s) and 1 GbE connections to 10 GbE connections. Our new 10 GbE data center fabric design will meet current needs while positioning us to incorporate new technology to meet future network requirements.

For example, the 10 GbE network will simplify the virtualized host architecture used for Office and Enterprise computing, and will also reduce network component cost. For Design computing, the 10 GbE network will reduce latency and allow faster application response times without the expense associated with alternative low-latency systems such as InfiniBand*. And, although storage area networks (SANs) in the Office and Enterprise application areas currently use a separate Fibre Channel (FC) fabric, we anticipate that as 10 GbE technology matures, we will be able to consolidate the SANs onto the 10 GbE network—thereby reducing network complexity even further.

Simplifying Virtualization for Office and Enterprise Applications

Intel's data center strategy for Office and Enterprise relies on virtualization and consolidation to reduce data center cost and power consumption, while reducing time to provision servers. Our current consolidation level is 12:1, and we are targeting a 20:1 consolidation level or greater with newer dual-socket servers based on Intel® Xeon® processors.

As virtualization increases, so does the number of server connections. Using a 1 GbE network fabric, a single physical server on the LAN requires eight GbE LAN ports, as shown in Figure 1 (left).

By converting to a 10 GbE data center design, we can simplify server connectivity by reducing the number of LAN ports from eight to two 10 GbE and one 1 GbE. This significantly reduces cable and infrastructure complexity, also shown in Figure 1 (right).

Presently, the SAN will continue to use FC connections and host bus adapters (HBAs).

In addition to simplifying physical infrastructure, a 10 GbE network fabric also has the potential to reduce the overall total cost of ownership (TCO) for LAN components per server by about 18.5 percent compared to the 1 GbE fabric, as shown in Table 1.

Table 1. 10 Gigabit Ethernet (GbE) Cost Savings

LAN Component	10 GbE Fabric Effect on Cost
Cable Infrastructure	48% Cost Reduction
LAN Infrastructure (per Port)	50% Cost Reduction
Server	12% Cost Increase
Storage Infrastructure	0% (Same for Both Fabrics)
Total Savings (per Server)	18.5% Overall Cost Reduction




Figure 1. Implementing a 10 gigabit Ethernet data center fabric reduces the number of LAN ports for our virtualized host architecture from eight to three.

Increasing Throughput and Reducing Latency for Design Applications

For some specific silicon design workloads, we needed to build a small, very low-latency cluster between servers. Several parallel applications running on these servers typically carry very large packets. As shown in Table 2, we compared application response times, using several network fabric options.

The 10 GbE network provided acceptable performance at an acceptable price point. For messages 16 megabytes (MB) and larger, the 10 GbE performance was about one-quarter of the 1 GbE response time, and was closer to the performance of InfiniBand, a more expensive low-latency switched fabric subsystem. In the table, “multi-hop” is defined as having to use more than one switch to get through the network.

Choosing the Right Network Components

We surveyed the market to find products that met our technical requirements at the right price point. We discovered that not all equipment and network cards are identical—performance can vary. With extensive testing, we found that we could reduce the cost for

a 10 GbE port by 65 percent by selecting the right architecture and supplier. For example, we had to decide where to place the network switch—on top of the rack, which provides more flexibility but is more expensive, or in the center of the row. We chose to use center-of-the-row switches to reduce cost.

Higher transmission speed requirements have led to new cable technologies, which we are deploying to optimize our implementation of a 10 GbE infrastructure:

Small-form factor pluggable (SFP+) direct-attach cables. These twinaxial cables support 10 GbE connections over short distances of up to 7 meters. Some suppliers are producing a cable with a transmission capability of up to 15 meters.

Connectorized cabling. We are using this technology to simplify cabling and reduce installation cost, because it is supported over SFP+ ports. One trunk cable that we use can support 10 GbE up to 90 meters and provides six individual connections. This reduces the amount of space required to support comparable densities by 66 percent. The trunks terminate on a variety of options, providing for a very flexible system. We also use a Multi-fiber Push-On (MPO) cable, which is a connectorized fiber technology

comprised of multi-strand trunk bundles and cassettes. This technology can support 1 GbE and 10 GbE connections and can be upgraded easily to support 40 and 100 GbE parallel-optic connections by simply swapping a cassette. The current range for 10 GbE is 300 meters on Optical Multimode 3 (OM3) multi-mode fiber (MMF) and 10 kilometers on single-mode fiber (SMF).

To maximize the supportable distances for 10 GbE, and 40 GbE/100 GbE when it arrives, we changed Intel's fiber standard to reflect a minimum of OM3 MMF and OM4 where possible, and we try to use more energy-efficient SFP+ ports.

FUTURE PLANS FOR OFFICE AND ENTERPRISE I/O AND STORAGE CONSOLIDATION

Historically, Ethernet's bandwidth limitations have kept it from being the fabric of choice for some application areas, such as I/O, storage, and interprocess communication (IPC). Consequently, we have used other fabrics to meet high-bandwidth,

Table 2. Application Response Times for Various Packet Sizes

Packet Size in Bytes	Application Response Time in Microseconds			
	Multi-Hop 1 Gigabit Ethernet (GbE) Current Standard	Multi-Hop 10 GbE	One-Hop 1 GbE	Multi-Hop InfiniBand*
8	69.78	62.50	41.78	15.57
128	75.41	62.55	44.77	17.85
1,024	116.99	62.56	64.52	32.25
4,096	165.24	65.28	103.15	60.15
16,384	257.41	62.47	195.87	168.57
32,768	414.52	129.48	348.55	271.95
65,536	699.25	162.30	627.12	477.93
131,072	1,252.90	302.15	1,182.41	883.83

Note: Intel internal measurements, June 2010.

low-latency, no-drop needs, such as FC. The advent of 10 GbE is enabling us to converge all our network needs onto a single, flexible infrastructure.

Several factors are responsible for increasing the I/O demand on our data centers. First, when more servers are added to the data center, it increases input/output operations per second (IOPS), which creates a proportional demand on the network. In addition, as each generation of processors becomes more complex, the amount of data associated with silicon design also increases significantly—again, increasing network demand. Finally, systems with up to 512 gigabytes (GB) of memory are becoming more common, and these systems also need a high-speed network to read, write, and back up large amounts of data.

Our upgrade to 10 GbE products will enable us to consolidate storage for Office and Enterprise applications while reducing our 10 GbE per-port cost. When I/O convergence on Ethernet becomes a reality, multiple traffic types, such as LAN, storage, and IPC, can be consolidated onto a single, easy-to-use network fabric. We have conducted multiple phases of testing, and in the near future, these 10 GbE ports will be carrying multiple traffic types.

CONCLUSION

A high-performance 10 GbE data center infrastructure can simplify the virtualization of Office and Enterprise applications and reduce per-server TCO. In addition, 10 GbE's lower network latency and increased throughput performance can support our design teams' high-density computing needs, improving design engineer productivity.

Our analysis shows that for a virtualized environment, a 10 GbE infrastructure can reduce our network TCO by as much as 18 to 25 percent. And, for design applications, where low latency is required, 10 GbE can play a crucial role without requiring expensive low-latency technology. The new fabric will also reduce data center complexity and increase our network's agility to meet Intel's growing data center needs.

ACRONYMS

GbE	gigabit Ethernet
GB	gigabyte
FC	Fibre Channel
HBAs	host bus adapter
IOPS	input/output operations per second
IPC	interprocess communication
Mb/s	megabits per second
MB	megabytes
MMF	multi-mode fiber
MPO	Multi-fiber Push-On
OM3	Optical Multimode 3
SAN	storage area network
SFP+	small form-factor pluggable
SMF	single-mode fiber
TCO	total cost of ownership

For more information on Intel IT best practices, visit www.intel.com/it.

This paper is for informational purposes only. THIS DOCUMENT IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NON-INFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE. Intel disclaims all liability, including liability for infringement of any proprietary rights, relating to use of information in this specification. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted herein.

Intel, the Intel logo, and Xeon are trademarks of Intel Corporation in the U.S. and other countries.

* Other names and brands may be claimed as the property of others.

Copyright © 2011 Intel Corporation. All rights reserved.

Printed in USA
0111/AC/KC/PDF

 Please Recycle
324159-001US

