

## Introduction

The standard for 10 Gigabit Ethernet (IEEE802.3ae) was ratified in 2002. While 10 GigE deployments have grown every year since then, the technology has primarily been used to interconnect switches and routers. Almost all the server connections in datacenters have remained at 1 Gbps, limiting the amount of network throughput available to each server.

The primary reason for this has been cost-performance. Historically it has been more cost-effective to have multiple GigE connections rather than a single 10 GigE port. In addition, today's servers typically cannot utilize the full bandwidth of a 10 GigE connection. However, both of these factors are changing which will lead to widespread adoption of 10 GigE for server connections over the next few years.

This white paper is an overview of the factors that are driving the growth for 10 GigE in the datacenter.

## Increased Demand for Network Performance

The demand for network performance is driven by increasing application demands, specifically:

- Next generation multi-core CPUs with multi-threaded networking stacks will be able to fully utilize a 10 Gbps connection. Server CPU throughput has historically doubled every 18 month, and there is every reason to expect this trend to continue.
- The increase in CPU throughput drives increased demands in storage I/O. The fastest growing storage architecture are so-called unified storage servers that offer both network file storage as well as network block storage, a natural fit for 10 Gigabit Ethernet.
- Server virtualization is increasing server utilization and therefore the demands for network bandwidth per physical server.
- Cluster computing requires high bandwidth and low-latency between all the nodes in the cluster. 10 GigE offers a 10X improvement in this area compared to 1 GigE.
- Real-time applications such as distribution of financial market data demand the lowest latency. 10 GigE is the natural solution for this requirement.


## Scalability of the Network Fabric

A 10 Gigabit Datacenter network offers many advantages in terms of scalability, including the following:

- Enables terabits of aggregate traffic, without adding more layers to the network
- Simplifies the network design by eliminating congestion points
- Reduces the need for complex QOS schemes
- Supports server virtualization that increases network demands due to server consolidation
- Improves application efficiency, server utilization and power efficiency
- Avoids the need to use multiple 1 GigE NICs and connections to get higher bandwidth

## Lower Latency

10 Gigabit Ethernet offers a compelling improvement in end-to-end latency. For HPC applications and financial trading environments, low latency is critical. Not only does 10 GigE reduce the packet transmit time by a factor of 10X, but 10 GigE switches also offer typically an order of magnitude lower latency than 1 GigE switches. For latency sensitive application, 10 GigE is a major improvement.



## Reliability

Many datacenters require 24x7 network availability and have no provisions for scheduled downtime. To support this on the network side requires a higher level of hardware and software reliability design than the conventional enterprise design practice. Fortunately, next generation 10 Gigabit Ethernet switches offer fundamental reliability improvements for both fixed (1U) and modular switches, including:

- Hot-swap redundant power and cooling
- Highly modular software design to offer fault containment
- Stateful Fault Containment (SFC) and Stateful Fault Repair (SFR)
- In-service software upgrades (ISSU) for the entire network

## Density, Power and Cooling of 10GigE Switches

Density, Power and Cooling are key drivers for deployment of datacenter network infrastructure.

- The latest 10 GigE racktop switches now support the same 48-port density as 1 GigE switches, which means that the customer does not lose valuable rack space when upgrading to 10 GigE.
- Next generation 10GigE racktop switches use front-to-rear cooling to match the server airflow. This is essential for efficient cooling of switches located in server racks.
- Power consumption per 10G port is still a major concern however the latest 10 GigE switches have much improved power consumption compared to earlier models.

## Challenges in Moving to 10GigE

While 10 GigE offers many advantages, there are certain challenges as well.

- Not all switches offer mixed-speed ports. As a result, you are forced to upgrade all your servers and the network at the same time. In many cases, this is hard to do.
- 10 GigE NICs are still somewhat expensive. This will change once 10 GigE NICs are integrated on future server mother boards, which is expected during calendar 2009.
- 10 GigE optical transceivers are still expensive, and only some vendors offer twin-ax copper cables as a more cost effective solution.

The cost per 10G port, including the server NIC and the physical layer, remains the biggest impediment to large-scale adoption of 10G for servers. During 2009, the cost for a 10G server connection is expected to decline below the cost of multiple GigE server connections, which is the present approach to scale server bandwidth beyond 1 Gigabit Ethernet.

## Future Trends

It is clear that the migration to 10 Gigabit Ethernet has started. New datacenter build outs everywhere are considering the advantages of 10 GigE. The percentage of servers directly attached to 10 GigE is expected to increase from low single-digits in 2008 to double digits in 2009. By 2012 it is expected that the majority of all servers will be connected directly to 10 GigE.

A new development is the emergence of 10GBASE-T products, capable of driving 10 Gigabit Ethernet over 100m of Cat-6 twisted pair cabling. While 10GBASE-T draws more power than 10 GigE SFP+, it is expected to be popular with customers that prefer the familiarity and ease of use of twisted pair cabling.

The IEEE802.3 high-speed task force is now defining the standard for 40 and 100 Gigabit Ethernet, which is a key development for future backbone connectivity. It is expected that the first switches that support 40 and 100 Gigabit Ethernet will become available in 2010.

## Summary

For next-generation datacenters, 10 Gigabit and higher speed Ethernet is the interconnect of choice, providing key improvements in terms of bandwidth, latency, scalability, reliability and application performance. If you are planning to upgrade your datacenter fabric, you should consider the latest 10GigE solutions that are cost-effective compared to multiple GigE ports but deliver substantially better throughput and latency.