



10GBASE-T Ecosystem Is Ready for Broad Adoption

White Paper

Yinglin (Frank) Yang, CommScope, Inc.

Sudeep Goswami, Cisco Systems

Carl G. Hansen, Intel Corporation

Contents

Executive Summary	3
Business Case	4
Cost	5
Power	6
Latency	7
Availability	8
Applications	9
Conclusion	11

Executive Summary

10GBASE-T is the standard technology that enables 10 Gigabit Ethernet operations over balanced twisted-pair copper, including Category 6A (Category 6A) unshielded and shielded cabling. 10GBASE-T provides great flexibility in network design due to its 100-meter reach capability. 10GBASE-T also provides the requisite backward compatibility that allows most end users to transparently upgrade from their 10/100/1000-Mbps networks. The ecosystem is now in place, and adoption of 10GBASE-T is poised to surge. Many analysts are predicting rapid growth of 10GBASE-T in the coming year. The latest data shows that the price of 10GBASE-T technology is declining dramatically, which is helping increase product demand.

New 10GBASE-T physical layers (PHYs) allow lower-cost and lower-power high-density designs. The latency of the 10GBASE-T PHY has also been improved and allows the building of 10GBASE-T networks to support most of today's applications.

10GBASE-T equipment from multiple vendors is available in the marketplace. The equipment includes high-port-density switches, modular-chassis-based switches, network interface cards, Small computer System Interface over IP (iSCSI) storage arrays, and cabling. 10GBASE-T LAN on motherboard (LOM) and higher-port-density switches are expected to become available in the near future, further aiding the adoption of 10GBASE-T.

An immediate use for 10GBASE-T is to build the data center access-layer network that connects servers to access switches. The data in this article shows a vast number of data center customers need to deploy networks reaching between 18 and 46 meters. 10GBASE-T, along with its corresponding cabling, is the most cost-effective and flexible solution in this length range. With the zone distribution model discussed in this document, 10GBASE-T can be used to build modular and scalable data center networks.

Using the data and analysis in the following sections, this document concludes that the economic and technical barriers to end-user adoption of 10GBASE-T have been overcome, and 10GBASE-T is moving toward massive deployment.

Business Case

Over the past several years, the data center has been playing a pivotal role in enabling IT to become a strategic asset rather than simply a cost center. Strategic initiatives such as server consolidation and server virtualization are gaining widespread momentum, which in turn is propelling a bandwidth transition from traditional 1 Gigabit Ethernet to higher-speed 10 Gigabit Ethernet at the server access layer.

New LOM Form Factors Coming

The traditional LAN on motherboard, or LOM, is soldered-down networking silicon on the motherboard. In the past, all servers came configured with some networking solution. With the upcoming Intel server refresh in Q1 2012, some original equipment manufacturers (OEMs) will offer 10GBASE-T as traditional LOM or will give customers the option of purchasing 10GBASE-T as a flexible LOM (FLEX LOM). FLEX LOM is a new form factor that defines an open networking socket on the motherboard instead of the traditional soldered-down networking silicon.

For the customer, the clear advantage is choice. Customers purchasing higher-end platforms supporting FLEX LOM will now be able to choose the LOM that is plugged into the motherboard. With this new delivery mechanism, customers will be able to choose their preferred network vendors, the number of ports, and the speed of the FLEX LOM. Depending on the OEM vendor, customers may have an option to upgrade their FLEX LOMs in the future. They can increase the speed from 1 Gigabit Ethernet to 10 Gigabit Ethernet or add ports. Some vendors may choose to offer a combination of 1 and 10 Gigabit Ethernet in the same FLEX LOM.

Offering more than simply a tenfold increase in performance, 10 Gigabit Ethernet helps enable exciting new use models including unified networking (iSCSI, Fibre Channel over Ethernet [FCoE], and LAN) and virtualization. These new use models allow customers to consolidate their network infrastructure by running iSCSI, FCoE, and LAN traffic on the same Ethernet fabric, scale their network with virtualization, simplify the network by consolidating the number of cables and switch connections, and partition their 10 Gigabit Ethernet ports into multiple channels. These new options give the customer more choice and flexibility than ever before.

10GBASE-T is poised to become a compelling choice for 10 Gigabit Ethernet server connectivity. The technology is based on the familiar and well-understood RJ-45 cabling scheme and is fully backward compatible with 1 Gigabit Ethernet speeds, thus providing a cost-effective and transparent migration path to 10 Gigabit Ethernet.

"10GBASE-T is an important technology as it will allow greater adoption of 10 Gigabit Ethernet in customers, especially those in small and medium businesses where having a single switch connect to new 10 Gigabit Ethernet servers as well as existing gigabit servers is an important purchasing decision," said Alan Weckel, director of Dell'Oro Consulting Group.

As the 10 Gigabit Ethernet market expands into more general use cases for server applications in which latency is not as critical, 10GBASE-T can be a cost-effective option even with a higher power envelope. As with any new standards-based technology, the speed and trajectory of adoption is dictated by the ecosystem support for that technology.

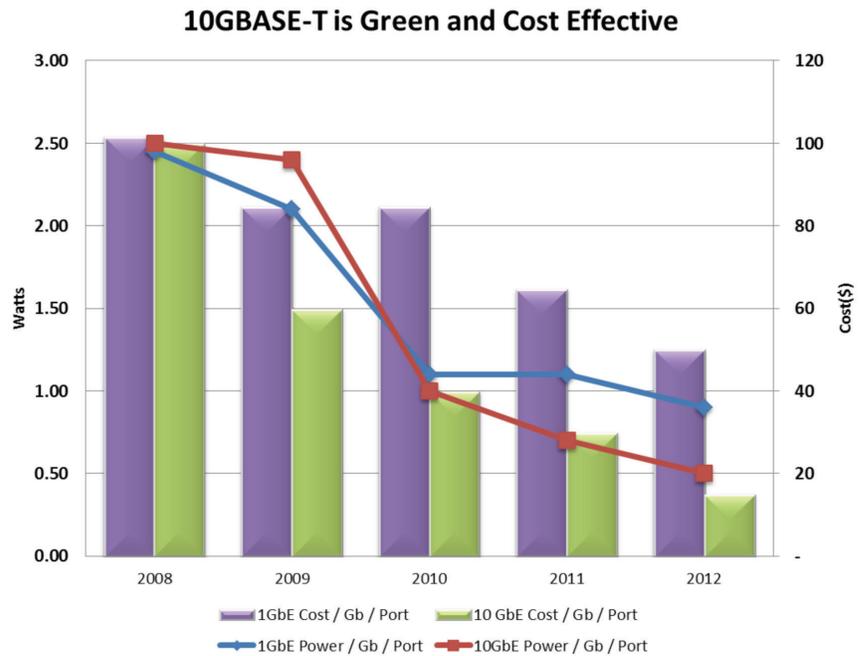
According to Crehan Research, 10GBASE-T is starting to gain traction on server adapter cards and top-of-rack (ToR) switches. *"With the impending introduction of 10GBASE-T LAN-on-motherboard solutions, in conjunction with numerous high port-density 10GBASE-T server access switches, I expect the adoption of this technology to accelerate rapidly, driving 10GBASE-T to eventually become the most popular datacenter connection,"* said Seamus Crehan, president of Crehan Research.

Several products are already available today that take advantage of this new technology, and the coming year promises to deliver additional offerings across the server, switch, adapter, and storage vendor ecosystem. Therefore, customers should start planning now to integrate 10 Gigabit Ethernet into their next-generation data center designs.

Cost

The cost of 10GBASE-T has been dramatically declining. Figure 1 illustrates 10GBASE-T's cost and power consumption per Gbps per port in comparison to 1000BASE-T over several years. The left axis shows the power consumed (watts [W]), and the right axis shows the cost (US\$). The bar graph is a comparison of cost, and the lines compare power use. Figure 1 shows that 10GBASE-T cost has declined approximately 70 percent since 2008, with further decline forecast for the next year. 10GBASE-T currently is much cheaper than 1000BASE-T on a per-Gbps per-port basis, and this cost advantage is expected to continue in the coming years.

Figure 1: 10GBASE-T Cost and Power Comparison with 1000BASE-T



Source: Intel

The continuous decline of cost to end users will be one of the major factors making 10GBASE-T a desirable media for broad adoption. With 10 times the data rate capacity, 10GBASE-T provides end users with the next logical step of upgrade from 10/100/1000-Mbps Ethernet networks.

Some perceptions remain that the 10GBASE-T PHY transceiver consumes too much power and is too costly. This perception may have been correct for the first generation of 10GBASE-T PHY in 2006 and 2007, but PHYs are following Moore's Law, and new processes have significantly decreased both cost and power use, and subsequent process improvements will continue to enhance these decreases.

Power

The power consumption of 10GBASE-T has declined significantly as shown in Figure 1. The figure shows that the power consumption of 10GBASE-T has been reduced by approximately 80 percent. In general, a 10GBASE-T PHY consumes 6.5W or less per port at the full-reach capability. A 10GBASE-T PHY can consume approximately 60 percent less power when running in the short-reach mode. Both the full-reach and short-reach modes are defined in the 10GBASE-T standard. The full-reach mode is specified to support 10GBASE-T up to 100 meters, and the short-reach mode is defined to support 10GBASE-T up to 30 meters. With 40-nanometer (nm) or even future 28-nm semiconductor technology, the next generation of 10GBASE-T PHY is expected to reduce the power consumption and heat dissipation even further.

Note that 10GBASE-T consumes much less power than 1000BASE-T on a per-Gbps per-port basis, making 10GBASE-T more power efficient than 1000BASE-T.

The power consumption and heat dissipation of 10GBASE-T PHY supports the designs of high-port density 10GBASE-T switches and LOMs.

Latency

Today's 10GBASE-T PHY has back-to-back latency (two PHYs) of approximately 2.5 microseconds in the full-reach mode. The latency can be lowered to about 1.5 microseconds in the short-reach (or data center) mode. The 2.5- or 1.5-microsecond latency seems significantly higher than the 0.1-microsecond latency seen on 10 Gigabit Ethernet over multimode fiber (MMF) or 10-Gbps Enhanced Small Form-Factor Pluggable (SFP+) direct-attach cable (DAC). Table 1 shows that the 2.5- microsecond latency is just a tiny fraction of the delay seen in typical end-to-end communications.

TABLE 1: 10GBASE-T PHY LATENCY COMPARED TO THE DELAY IN END-TO-END COMMUNICATIONS

	10GBASE-T PHY	Ping Response	TCP File Transfer	Web Server Response
Description	2 PHY devices back to back without cabling	2 servers back to back with a 10 Gigabit Ethernet fiber connection, round trip	TCP file transfer across 10 Gigabit Ethernet network (3 switching hops), round trip	From a browser to a web server, round trip
Latency	2.5 microseconds	90 microseconds	0.3 to 1 millisecond (ms)	50 ms
Share (%) of latency from 10GBASE-T PHY	–	5%	1.7 to 0.5%	0.01%

Source: CommScope's research

For 1000BASE-T, latency ranges from less than 1 to 12 microseconds depending on buffer size. For 10GBASE-T, the range is limited to just over 2 to 4 microseconds. Therefore, 10GBASE-T has more than three times lower latency than Gigabit Ethernet with buffer sizes of 512 bytes or more.

The initial delay caused by 10GBASE-T latency is higher than that of 1000BASE-T. However, if you take into account the throughput of 10GBASE-T and the Ethernet packet size, you can see that at any Ethernet packet size of 512 bytes or more, 10GBASE-T's overall throughput offers a significant advantage over 1000BASE-T, even with its initial delay.

With jumbo Ethernet frames, the initial 10GBASE-T latency has essentially no impact. Even at packet sizes smaller than 512 bytes, only the most latency-sensitive applications such as high-performance computing (HPC) or high-frequency trading systems would notice any latency. The overwhelming majority of enterprise applications will never notice any effect of latency.

In the real world, however, the communication delay from server to server in the data center environment can be affected by many other factors: processor workload, memory access, switch workload, software interface, etc. Software typically contributes latency on the order of milliseconds (ms). A millisecond is 1000 microseconds. Most real-world applications are not sensitive to the 2.5-microsecond latency.

Some extreme HPC applications, such as high-end scientific, engineering, and financial supercomputers that require very low latency, will use high-speed proprietary or non-Ethernet fabric to meet their interconnect needs. 10GBASE-T will be used for management networks or other non-HPC needs.

Availability

The 10GBASE-T ecosystem continues to grow and offer a robust number of options in the marketplace (Table 2). These options are from different leading vendors. It is certainly desirable to have more high-port-density 10GBASE-T fixed or chassis based modular switches available. 10GBASE-T LOM will be another important factor in promoting massive adoption.

TABLE 2: 10GBASE-T EQUIPMENT IN THE MARKET

Type	Equipment	Availability
Switch	40-port 10GBASE-T fixed switch	Yes
	32-port 10GBASE-T fixed switch	Yes
	24-port 10GBASE-T fixed switch	Yes
	Chassis-based modular switch with medium 10GBASE-T port density (8 or 16 ports per module)	Yes
	Chassis-based modular switch with high 10GBASE-T port density (more than 16 ports per module)	2012
Network interface card (NIC) or LOM	Switch module for blade server	Yes
	Dual-port 10GBASE-T NIC	Yes
	Single-port 10GBASE-T NIC	Yes
Storage device	iSCSI storage array with 10GBASE-T ports	Yes
Cabling	Unshielded and shielded cabling	Yes

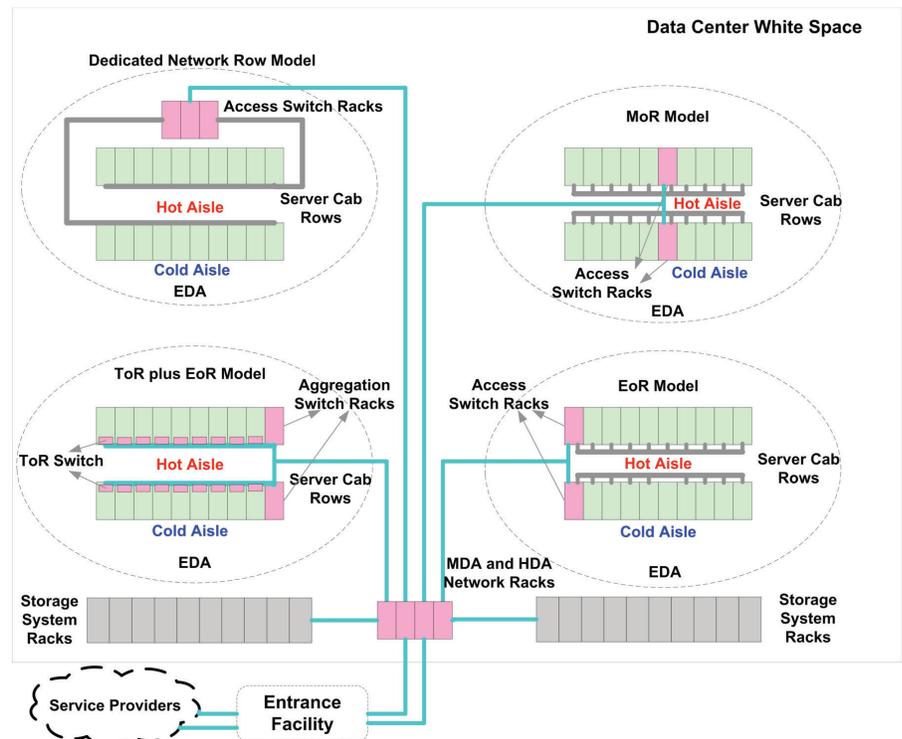
Applications

An immediate use for 10GBASE-T is to provide servers with connectivity to access Ethernet networks in data centers: that is, 10GBASE-T can be used to build access-layer networks in the data center.

One of 10GBASE-T's benefits to end users is the up-to-100-meter reach, which provides great flexibility in various data center network designs. 10GBASE-T fits a zone distribution cabling architecture very well. A zone may be constructed with an end-of-row (EoR), middle-of-row (MoR), dedicated network row, or ToR design combined with an EoR server access connectivity scheme.

Figure 2 shows the zone distribution model. The figure is a top view of the data center white space. A zone, sometimes called a pod, represents a portion of the data center white space. The four ellipses represent four zones, or pods. The zone can be repeated throughout the entire data center white space, so each zone may have the same configuration of server cabinet rows, network racks, and cabling routes. The four popular types of zone are illustrated: EoR, MoR, dedicated network row, and ToR combined with EoR or MoR. In each zone, servers are connected to their access switches through Category 6 and Category 6A cable, represented by the gray lines in the figure. The access switches are connected to aggregation-layer switches through fiber or copper cabling, presented by aqua lines. Aggregation and core switches can be located in the horizontal distribution area (HDA) or main distribution area (MDA). The rows of server cabinets are in the equipment distribution area (EDA).

Figure 2: Zone Distribution Model



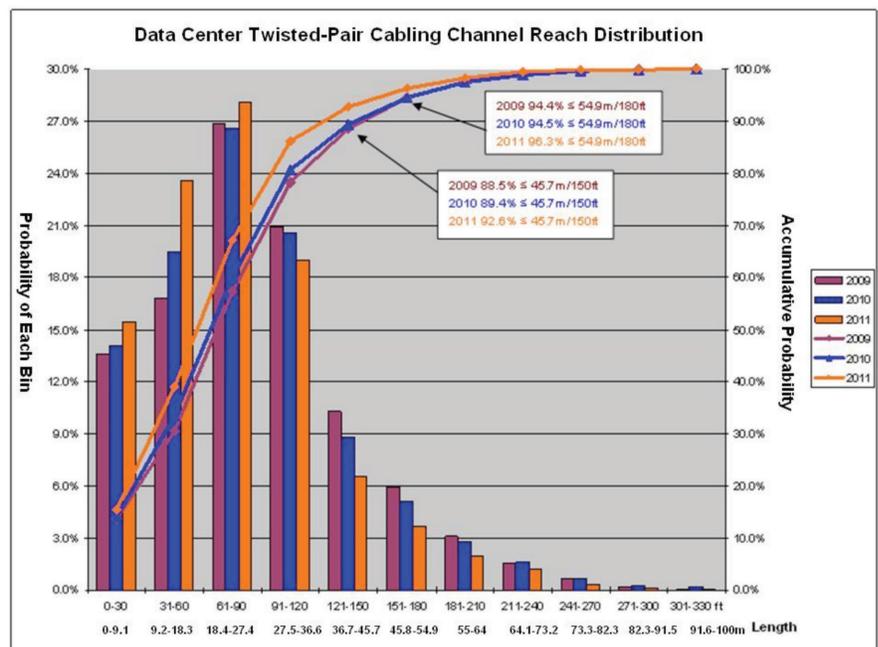
Source: CommScope

The zone distribution model is one of the most widely adopted cabling deployment models in data centers, especially in large and very large data centers. The zone distribution model can provide excellent modularity, scalability, and consistency in the data center.

The size of a zone can be affected by a number of factors, such as the shape and size of the white space, the end user's design preference, and the cable length for optimal cable management and maintenance.

Figure 3 shows the results of a study on the distribution of balanced twisted-pair cabling channel lengths in data centers. A cabling channel is defined as a single end-to-end cabling system that can be composed of horizontal cable, patch cords and connections, or simply a patch cord. A cabling channel can connect a NIC port to a switch port, or a switch port to another switch port. The length data in Figure 3 is from the CommScope sales database of preterminated copper cabling solutions used in data centers. Figure 3 reflects the deployment in the real world. The data includes Category 6A, Category 6, and Category 5E cabling.

Figure 3: Balanced Twisted-Pair Cabling Channel Length Distribution in Data centers



Source: CommScope

Table 3 highlights several important findings from Figure 3. It is clear that many end users need to deploy network connections 18 to 46 meters long. Passive direct-attach Twinax cable, such as SFP+ DAC, does not support the 18- to 46-meter range. OM3/OM4 multimode fiber cabling can support this length range, but is not as cost effective as Category 6A. The active SFP+ DAC cable can support this length range, but it is very hard to pull the active cables for more than a few meters. Therefore, 10GBASE-T is the best cost-effective solution for meeting customers' needs in this length range.

TABLE 3: MAIN FINDINGS FROM FIGURE 3

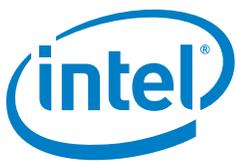
Main Findings from Figure 3	Cabling Deployment Model
56% of the channel lengths are between 18 and 46 meters	Zone distribution model fits in this length range • Most suitable for 10GBASE-T
34% of the channel lengths are less than or equal to 18 meters	ToR model fits well in this length range • Suitable for 10GBASE-T and other 10 Gigabit Ethernet designs
95% of the channel lengths are less than or equal to 55 meters	Any deployment model

The access layer cabling length in a zone is likely between 18 and 46 meters long. 10GBASE-T along with the corresponding cabling can satisfy the end user's needs in the zone distribution model in data centers.

In brief, 10GBASE-T provides end users with backward compatibility with their existing networks and flexibility in network designs. 10GBASE-T can be used for long-reach networks up to 100 meters or short-reach networks as short as a few feet. 10GBASE-T can be used in the zone distribution model for scalability and consistency, and it also can be used in the ToR model for modularity and good cable management.



Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership between Cisco and any other company. (1110R)



Copyright © 2012 Intel Corporation. All rights reserved. Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries.

For more information on Intel® Ethernet products, visit www.intel.com/go/ethernet

Conclusion

10GBASE-T provides the following benefits to end users:

- Backward compatibility with 10/100/1000-Mbps Ethernet networks, which have the largest installed base and highest annual ship rate of any networks today in the world
- Flexibility in design and deployment: 10GBASE-T networks are suitable for zone distribution and ToR models
- Capability to use standard testing procedures in the field
- The simplicity of Gigabit Ethernet with 10 times the speed

Although the 10GBASE-T ecosystem is robust, an increased number of high-port-density fixed and chassis-based switches in the marketplace is certainly desirable, and more will become available over the coming year. The upcoming 10GBASE-T LOM is also seen as one of the major factors influencing adoption.

With the cost decline, technology advancements, product availability, and appropriate applications, the 10GBASE-T ecosystem is ready for broad adoption by end users.



www.commscope.com

Visit our Web site or contact your local CommScope representative for more information.

© 2011 CommScope, Inc. All rights reserved.

All trademarks identified by ® or ™ are registered trademarks or trademarks, respectively, of CommScope, Inc.

This document is for planning purposes only and is not intended to modify or supplement any specifications or warranties relating to CommScope products or services.