



Comparing Copper and Fiber Options in the Data Center

In most data center designs there is a mixture of both copper and fiber infrastructure. This paper is not suggesting that one should replace the other, rather that each should be considered carefully with respect to the applications expected to be supported over the life of the data center. With varied capabilities of networking equipment and cabling options, a thorough analysis should be performed to plan the most cost effective data center infrastructure to maximize your return on investment.

POWER AND COOLING EFFICIENCIES

There are several factors driving data center specifiers and decision makers to revise, remediate, relocate or consolidate current data centers. Power and cooling are two of the more significant factors. In many legacy data centers, older-model air-handling units operate at roughly 80% efficiency at best, measured in terms of electrical use per ton of cooling (kW/ton). Newer units operate at between 95-98% efficiency depending on the manufacturer and model. In some instances, it is more cost effective for companies to write off unrealized depreciation in order to receive the efficiency benefits of the newer units.

But with any cooling equipment, conditions apart from the cooling unit itself can have a significant impact on efficiency. Simple steps like removing abandoned cable from pathways to reduce air dams and maximize air flow, installing brush guards or air pillows to maintain static pressure under the floor, and redressing cabling within cabinets to lessen impedance of front to back airflow, are all beneficial and are forcing companies to look at these and other relatively simple upgrades for improving power and cooling efficiencies. With green/ecological and power reduction initiatives swaying today's decisions, the circular relationship between power consumption and cooling is bringing facilities back into the discussions for selecting network equipment (e.g., servers, switches, SANs).

INCREASING STORAGE AND BANDWIDTH TRENDS

In addition to requirements for faster processing and lower power consumption, recent changes in legislation and mandates for data retention (Sarbanes Oxley for example) are driving storage costs up. While these vary by industry, governance and company policy, there is no question that storage and data retrieval requirements are on the rise. According to IDC¹, "281 exabytes of information existed in 2007, or about 45 Gb for every person on earth." As with any other equipment in the data center, the more data you have and transfer, the more bandwidth you will need. To support faster communications, there are a growing number of high-speed data transmission protocols and cabling infrastructures available, each with varying requirements for power and physical interfaces.

To meet these increasing demands for bandwidth in the data center, 10 Gb/s applications over balanced twisted-pair cabling, twinax cabling and optical fiber cabling are growing. The Dell'Oro Group, a market research firm, predicts that copper-based 10 GbE will expand to represent 42% of the projected 8.8M 10GbE units by 2010². A study by the Linley Group indicated that: "...by 2009, we expect 10GbE shipments to be well in excess of one million ports. The fast-growing blade-server market will drive the demand for 10GbE switches. At the physical layer, the 10GbE market will go through several transitions. . . including a shift to 10GBASE-T for copper wiring."³

10 Gb/S INFRASTRUCTURE OPTIONS

There are several cabling alternatives available over which 10 Gb/s can be accomplished. Infiniband is one option. The single biggest advantage of Infiniband is that it has far lower latency (around one microsecond) than TCP/IP and Ethernet based applications, as there is much less overhead in this transmission protocol. Infiniband is gaining popularity in cluster and grid computing environments not only for storage, but as a low latency, high performance LAN interconnect with power consumption at approximately 5 Watts per port on average.

A single Infiniband lane is 2.5 Gb/s, and 4 lanes result in 10 Gb/s operations in SDR (Single Data Rate) mode and 20 Gb/s in DDR (Dual Data Rate) mode. Interfaces for Infiniband include twinax (CX4) type connectors and optical fiber connectors: even balanced twisted-pair cabling is now supported through Annex A5⁴. The most dominant Infiniband connector today, however, utilizes twinax in either a 4x (4 lane) or 12x (12 lane) serial communication. These applications are limited to 3-15 m depending on manufacturer, which may be a limiting factor in some data centers. Optical Fiber Infiniband consumes approximately 1 Watt per port, but at a port cost of nearly 2x that of balanced twisted-pair. Active cable assemblies are also available that convert copper CX4 cable to optical fiber cable and increase the distance from 3-15 m to 300 m, although this is an expensive option and creates an additional point of failure and introduces latency at each end of the cable. One drawback to the CX4 Infiniband cable is diameter which is 0.549 cm (0.216 in) for 30 AWG and 0.909 cm (0.358 in) for 24 AWG cables.

With the release of the IEEE 802.3an standard, 10 Gb/s over balanced twisted-pair cabling (10GBASE-T) is the fastest growing and is expected to be the most widely adopted 10GbE option. Because category 6_A/class E_A and category 7/class F or category 7_A/class F_A cabling offer much better attenuation and crosstalk performance than existing category 6 cabling, the standard specified Short Reach Mode for these types of cabling systems. Higher performing cabling simplifies power reduction in the PHY devices for Short Reach Mode (under 30 m). Power back off (low power mode) is an option to reduce power consumption compared to category 6 or longer lengths of class E_A, class F or class F_A channels. Data center links less than or equal to 30 meters can take advantage of this power savings expected to roughly 50% depending on manufacturer.

The IEEE 802.3 10GBASE-T criteria states a goal that "the 10GBASE-T PHY device is projected to meet the 3x cost versus 10x performance guidelines applied to previous advanced Ethernet standards". This means that balanced twisted-pair compatible electronics, when they become commercially affordable, and not simply commercially available, will provide multiple speeds at a very attractive price point, relative to the cost of optical fiber compatible electronics. As maintenance is based on original equipment purchase price, not only will day-one costs be lower, but day-two costs will also be lower. Latency on first generation balanced twisted-pair compatible electronics chips is already faster than that written in the standard with latency near 2.5 microseconds.

At 1 Gb/s speeds, balanced twisted-pair compatible electronics offer better latency performance than fiber; however, considering latency at 10 Gb/s, currently fiber components perform better than balanced twisted-pair compatible 10GBASE-T electronics, but not as well as 10 Gb/s Infiniband/CX4. However, this will likely change with future generation 10GBASE-T chips for copper switches. It is important to remember that in optical transmissions, equipment needs to perform an electrical to optical conversion, which contributes to latency.

Balanced twisted-pair remains the dominant media for the majority of data center cabling links. According to a recent BSRIA press release: “. . .survey results highlight a rush to higher speeds in data centers; a broad choice of copper cabling categories for 10G, especially shielded; and a copper / fiber split of 58:42 by volume. 75% of respondents who plan to choose copper cabling for their 10G links plan for shielded cabling, relatively evenly split between categories 6, 6a and 7. OM3 has a relatively low uptake at the moment in U.S. data centers. The choice for fiber is still heavily cost related, but appears to be gaining some traction with those who want to future-proof for 100G and those not willing to wait for 10 Gb/s or 40 Gb/s copper connectivity and equipment.”⁵

Optical fiber-based 10Gb/s applications are the most mature 10GbE option, although designed originally for backbone applications and as an aggregation for gigabit links. Fiber’s longer reach makes the additional cost of fiber electronics worthwhile when serving backbone links longer than 90 meters. But using optical fiber for shorter data center cabling links can be cost prohibitive.

Mixing both balanced twisted-pair cabling and optical fiber cabling in the data center is common practice. The most common 10 GbE optical fiber transmission in use in the data center is 10GBASE-SR. This will support varied distances based on the type of optical fiber cabling installed. For the OM1 optical fiber (e.g., FDDI grade 62.5/125µm multimode fiber), distance is limited to 28 meters. For laser optimized OM3 grade 50/125µm (500/2000) multimode fiber, the distance jumps to 300 m with future proof support for 40 and 100 Gb/s currently under development within IEEE.

In order to increase the distances on OM1 grade optical fiber, two other optical fiber standards have published. 10GBASE-LX4 and 10GBASE-LRM increase allowable distances to 300 m, and 220 m respectively. However it is important to note that LX4 and LRM electronics are more expensive than their SR counterparts, and in most cases, it is less expensive to upgrade your optical fiber cabling to laser optimized (OM3) grade optical fiber as a cabling upgrade would not result in elevated maintenance costs due to the higher cost of the electronics.

10 Gb/S INFRASTRUCTURE OPTIONS PROGRESSION FROM 1Gb/S TO 10 Gb/S

Benefits of Virtualization

With virtualization, efficiency is increased further regardless of the interface used. It is possible to virtualize 8-10 servers on one blade today. That is, 8-10 servers that would have used 8-10 gigabit ports for primary network, 8-10 for secondary, 8-10 keyboard-video-mouse (KVM) connections, 8-10 management or remote monitoring connections, 8-10 sets of power supplies, 8-10 sets of storage connections, etc. Bringing this total down to one server, one set of power supplies, one set of network connections and one set of storage connections is attractive both for real estate and capital expenditures. In order to do this, however, sufficient throughput needs to be available on the network. The overall cost reduction in power is a large driving factor for virtualization.

In many cases for both optical fiber and balanced twisted-pair cabling, an upgrade from 1 Gb/s to 10 Gb/s will require a change of the Ethernet switch, as older switch fabrics will not support multiple 10 Gb/s ports. Prior to selecting balanced twisted-pair or optical fiber for an upgrade to 10 GbE, a study should be completed to ensure that power, cooling, and available space for cabling is adequate. This analysis should also include day one and day two operating and maintenance costs.

Power consumption for 10 Gb/s switches is currently a major factor in the cost analysis of balanced twisted-pair vs. optical fiber cabling in the data center. With first generation 10GBASE-T chips operating at 10-17 Watts per port, lower power consumption is a goal and a challenge for 10GBASE-T PHY manufacturers. This is certainly something to watch as next generation chips are expected to have much lower power demands on par with Infiniband ports or roughly one half of the first iterations. The same was seen in gigabit Ethernet, which from first generation chips to current technologies, saw a 94% decrease in power from 6 Watts per port to the 0.4 Watts per port figure we see today. Supporting this is the recent release of a 3.5 Watts per port 10GBASE-T chip from Aquantia⁶.

It is further noted that IEEE published Energy Efficient Ethernet (802.3az) technology will allow links go into a sleep state during periods of inactivity - a capability which could reduce power by an estimated 1.5 Watts per port. Average power per 24-hour period will be far less when Energy Efficient Ethernet is built into future generation 10GBASE-T chips.

Since optical fiber electronics cannot autonegotiate, a move from 1000BASE-xx to 10GBASE-xx requires a hardware change. In contrast, both 1GbE and 10GbE can be supported by 10GBASE-T balanced twisted-pair compatible equipment. Hardware changes cause downtime and a shortened lifecycle of the network hardware investment. There are several options for optical fiber communications at 10GbE. Each is characterized by range, wavelength and type of optical fiber media. The following table shows an estimated end-to-end cost comparison between various balanced twisted-pair and optical fiber data center applications including estimated 3 year maintenance contract costs.

BENEFITS OF VIRTUALIZATION

With virtualization, efficiency is increased further regardless of the interface used. It is possible to virtualize 8-10 servers on one blade today. That is, 8-10 servers that would have used 8-10 gigabit ports for primary network, 8-10 for secondary, 8-10 keyboard-video-mouse (KVM) connections, 8-10 management or remote monitoring connections, 8-10 sets of power supplies, 8-10 sets of storage connections, etc. Bringing this total down to one server, one set of power supplies, one set of network connections and one set of storage connections is attractive both for real estate and capital expenditures. In order to do this, however, sufficient throughput needs to be available on the network. The overall cost reduction in power is a large driving factor for virtualization.

APPLICATION	RANGE	CHANNEL COST*	MSRP FOR MODULE	MAINTENANCE@ 15% OF COST FOR 3 YEARS	TOTAL END-TO-END COST
OPTICAL FIBER					
1000BASE-SX	220m-550m	\$ 381.64	\$ 500.00	\$ 225.00	\$ 1,106.64
1000BASE-LR	550m	\$ 381.64	\$ 995.00	\$ 447.75	\$ 1,824.39
10GBASE-SR	28m-300m	\$ 381.64	\$ 3,000.00	\$ 1,350.00	\$ 4,731.64
10GBASE-LRM	220m-550m	\$ 381.64	\$ 1,495.00	\$ 672.75	\$ 2,549.39
10GBASE-LX4	300m	\$ 381.64	\$ 2,995.00	\$ 1,347.75	\$ 4,724.39
BALANCED TWISTED-PAIR					
1000BASE-T / 10GBASE-T*	100m	\$ 376.09	\$ 345.00	\$ 155.00	\$ 876.34
10GBASE-CX4	3m-15m	\$ 495.00	\$ 600.00	\$ 270.00	\$ 1,365.00
Infiniband	3m-15m	\$ 495.00	\$ 1,399.00	\$ 629.55	\$ 2,523.55

NOTES:

10GBASE-LRM requires Mode conditioning patch cords for OM1, OM2 increasing the channel cost approximately \$700.00.

10GBASE-T is estimated based on 10x performance at 3x the cost from IEEE 802.3AN.

Prices do not include chassis, power supplies or management modules which will vary with application.

* In this model, laser optimized (OM3) multimode fiber and category 6A F/UTP balanced twisted-pair cabling were used for calculating channel costs including installation with the exception of Infiniband, which uses pre-assembled 10GBASE-CX4 cable assemblies. For details on cost calculations see Total Cost of Ownership White Paper at http://www.siemon.com/us/white_papers/06-05-18-tco.asp MSRP for Modules is based on Cisco® Systems.

The above figures do not include chassis costs, power supplies, management modules, etc. The costs listed are for a single interface only based on pricing available at the time of publication.. The backplane and type of switch will vary with individual configurations. Twinax based Infiniband and 10GBASE-CX4 applications do not run on structured cabling systems. These cable assemblies are typically purchased from the equipment manufacturer and have a limited distance range of 15 meters. The cost of the 10GBASE-CX4 and Infiniband includes the average cost of the CX4 cable assemblies. For 10GBASE-LRM, Mode Conditioning patch cords are needed at each end of the channel if using less than OM3 fiber, increasing this overall cost to approximately \$3,359.30 for each port.

As previously noted, on the optical fiber side there is a network hardware change required to move from 1 Gb/s to 10 Gb/s. Assuming that SR modules were used for both applications, a 1000BASE-SR implementation today upgraded to a 10GBASE-SR implementation tomorrow would have to include the costs for both systems for a total of \$1,824.39 + \$4,731.64 – \$381.64 = \$5,456.64, assuming that a capable optical fiber channel (\$381.64) is installed and will be reused. For 10GBASE-T, since it is able to support both 1 Gb/s and 10 Gb/s and assuming the standards-based 10x the performance at 3x the cost, a single end-to-end channel supporting both speeds is \$876.09 which translates into a savings of \$4580.55.

In a data center with five hundred (500) 10 Gb/s capable ports using 1000BASE-SR today with a planned upgrade to 10GBASE-SR, the total costs including equipment upgrades (not including chassis, downtime or labor) is roughly \$2.7 million. The equivalent using the autonegotiation power of 10GBASE-T copper based gear is roughly \$1.0M. This translates to a 61% savings of roughly \$1.7 million (excluding chassis, power supplies and management modules) when using 10GBASE-T over balanced twisted pair cabling.

It is no wonder that many experts agree that balanced twisted-pair cabling will remain a dominant solution for a long time to come. Most data centers, in reality, will be a mixture of balanced twisted-pair and optical fiber for Ethernet communications. Optical fiber will continue to enjoy its place in the data center for storage applications and for distances beyond 100m or for those users with a higher budget who may wish to future proof for 100 Gb/s.

For design assistance and other tools to help in the decision making process please contact your Siemon sales representative and visit www.siemon.com. Siemon has extensive experience in data centers design assistance and implementation along with a global team to support you in your data center decisions.

FOOTNOTES:

1. "The Diverse and Exploding Digital Universe: An Updated Forecast of Worldwide Information Growth Through 2011" - International Data Corporation, 3/2008.
2. " Short-Reach 10GBaseT Cuts Power Consumption In The Data Center" - Electronic Design, 9/2007
3. "A Guide to Ethernet Switch and PHY Chips, Fourth Edition" - Linley Group, 8/2007
4. Supplement to InfiniBand™ Architecture Specification Volume 2 - Annex A5
5. "U.S. Data Center Structured Cabling & Network Choices" - BSRIA (March 2008)
6. Press Release: "Aquantia Demonstrates Robust Performance of Industry's First Low-Power 10GBASE-T PHY at Interop Las Vegas" – Aquantia, 4/2008 link: www.aquantia.com/pdf/Aquantia_Interop_08_release_Final.pdf



CISCO SOLUTION DEVELOPER'S PROGRAM

– Siemon is a participant in the Cisco Technology Developer Partner Program, with a full range of cabling products to support their technologies. For a listing of these products please visit www.siemon.com and click on the Cisco logo.

ABOUT THE AUTHOR

Carrie Higbie has been involved in the computing and networking for 25+ years in executive and consultant roles. She is Siemon's Global Network Applications Manager supporting end-users and active electronics manufacturers. She publishes columns and speaks at industry events globally. Carrie is an expert on TechTarget's SearchNetworking, SearchVoIP, and SearchDataCenters and authors columns for these and SearchCIO and SearchMobile forums and is on the board of advisors. She is on the BOD and former Pres of the BladeSystems Alliance. She participates in IEEE, the Ethernet Alliance and IDC Enterprise Expert Panels. She has one telecommunications patent and one pending.

THE AMERICAS

USA.....	(1) 866 474 1197
Canada.....	(1) 888 425 6165
Colombia - Central and South America Main.....	(571) 317 2121
Argentina.....	(54) 11 5031 4061
Brasil.....	(55) 11 3831 5552
Mexico.....	(52) 55 2881 0438
Peru.....	(511) 275 1292
Venezuela.....	(58) 212 661 6538

EUROPE, MIDDLE EAST AND AFRICA

United Kingdom.....	(44) (0) 1932 571771
Germany.....	(49) 2241 97398311
France.....	(33) 1 46 46 11 85
Italy.....	(39) 02 64 672 209
Middle East.....	(44) (0) 1932 571771 / (971) 506 441 470

ASIA PACIFIC

Australia (Sydney).....	(61) 2 8977 7500
Australia (Brisbane).....	(61) 7 3854 1200
Australia (Melbourne).....	(61) 3 9866 5277
Australia (Perth).....	(61) 8 9866 5277
Southeast Asia.....	(65) 6345 9119
China (Shanghai).....	(86) 21 5385 0303
China (Beijing).....	(86) 10 6559 8860
China (Guangzhou).....	(86) 20 3882 0055
China (Chengdu).....	(86) 28 6680 1100
India.....	(91) 11 46142162 / (91) 11 46142163
Japan.....	(81) (3) 5798 5790

Visit our website at www.siemon.com for detailed global office contact information