SIEMON®

NETWORK CABLING SOLUTIONS

E-BOOK: DATA CENTER CABLING INFRASTRUCTURE

CISCO PARTNER
Technology Developer

CONNECTING THE WORLD TO A HIGHER STANDARD

WWW.SIEMON.COM

SIEMON™

INTRODUCTION:

It would be difficult for any IT professional to point to a single element as the most critical piece of a data center and impossible to do so without an argument in a group of data center experts. Like any complex network environment, the data center is a system of interrelated and interdependent elements, each relying on the performance of the rest — an ecosystem of sorts.

In such an ecosystem, the best approach to understanding the whole is having core knowledge of its parts. This e-book, by noted data center expert Carrie Higbie focuses on one of those critical parts: the cabling infrastructure. From physical infrastructure basics and best practices to the key considerations in selecting cabling types to environmental considerations, Higbie arms the network professional with the knowledge they need to make educated data center cabling decisions.

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 President of the BladeSystems Alliance. She participates in IEEE, the Ethernet Alliance and IDC Enterprise Expert Panels. She has one telecommunications patent and one pending.

10Gb/s Data Center Cabling Solutions

The term data center conjures up different meanings to different people. Some would argue that the data center is the room where the servers are stored. Others visualize quite a different picture. It is true that at one time, the data center was little more than a secure server room. However with technological advances and data centric businesses of today the term could better be expressed as a "mission critical data center". Business models have gone through a complete cycle from centralized data centers to decentralized and now back to centralized. Businesses realize that data is their strongest asset and as such are making strides to assure its availability, security and redundancy.

The data center concept has also grown into its own business model. Companies that provide redundant and offsite storage for other companies are building state of the art facilities on a global scale. At the heart of these facilities is the IT infrastructure. This paper will address infrastructures and components of a data center. Whether a company implements all or part of these components, one core element will always remain, that is the cabling system infrastructure. This planning guide is designed to provide you with a basic roadmap for your data center.

DATA CENTER TRENDS

According to Infonetics Research's latest North American data center market research study, the projected combined data center services and products are expected to grow 47% from $10.6 billion to $15.6 billion between 2003 and 2007. Data centers can represent 50% of an IT budget for an organization. These data centers house the data for ERP (Enterprise Resource Planning) applications, e-commerce applications, SCM (Supply Chain Management), CAD/CAM, rich media, video/voice/data convergence, B2B (Business to Business) applications along with the back-office applications on which companies run. The communications mechanisms for the applications very, but the critical elements of data uptime do not change. According to IT Week, a survey of 80 large US companies conducted by analyst firm Infonetics last year indicated they had an average of 501 hours of network downtime per year, and this cost them almost four percent of their revenue, totalling millions of dollars. In separate research, analyst Gartner estimated a typical business experiencing an average of 87 hours of downtime a year, can result in total losses exceeding $3.6m.

It is not very difficult to see that downtime directly translates into dollars, and lots of them. Companies that provide data center components and equipment are sensitive to this and have made great strides in providing companies with viable, hearty solutions for their growing data stores and requirements.

COMPONENTS OF A DATA CENTER

Data centers are comprised of a high speed, high demand networking communication systems capable of handing the traffic for SAN (Storage Area Networks), NAS (Network Attached Storage), file/application/web server farms, and other components located in the controlled environment. The control of the environment relates to humidity, flood, electrical, temperature, fire controls, and of course, physical access. Communication in and out of the data center is provided by WAN, CAN/MAN and LAN links in a variety of configurations depending upon the needs of the particular center. A properly designed data center will provide availability, accessibility, scalability, and reliability 24 hours a day, 7 days a week, 365 days per year minus any scheduled downtime for maintenance. Telephone companies work for 99.999% uptime and the data center is no different. There are two basic types of data centers: corporate and institutional data centers (CDCs) and Internet Data Centers (IDCs). CDCs are maintained and operated from within the corporation, while IDCs are operated by Internet Service Providers (ISPs). The ISPs provide third party web sites, collocation facilities and other data services for companies such as outsourced email. Critical data centers are monitored by a NOC (Network Operations Center) which may be in-house or outsourced to a third party. The NOC is the first place outages are realized and the starting point for corrective action. NOCs are generally staffed during the data center's hours of operations. In 24 x 7 data centers, the NOC is an around the clock department. Equipment monitoring devices will advise the NOC of problems such as overheating, equipment outages, and component failure via a set of triggers that can be configured on the equipment or via a third party monitoring software which can run over all of the equipment.

DATA CENTER PLANNING AND DESIGN GUIDELINE

Data center planning has become somewhat of a specialty in the architectural world. Most architectural firms either have an RCDD (Registered Communications Distribution Designer) on staff, or acting as a consultant to assist with the specialized equipment not addressed by their Electrical Engineers and Mechanical Engineers. The equipment housed within the center is complex each with specific requirements for heating, cooling, power budgets and spatial considerations. A typical data center contains the following components:

- Computing and network infrastructure (cabling, fiber, and electronics)
- NOC or NOC communications and monitoring
- Power distribution, generation and conditioning systems
    - — Uninterruptible Power Supplies, generators
- Environmental control and HVAC systems
- Fire Detection and Suppression systems (typically halon or other non-water suppression)
- Physical security and access control prevention, allowance, and logging
- Circuit breaker protection (lightning protection in some cases)
- Proper lighting
- Minimum of 8 ft. 5 in. ceiling height

- Grounding
- Racks and cabinets for equipment
- Pathway: Raised access flooring and/or overhead cable tray
- Carrier circuits and equipment
- Telecommunications equipment
- Proper clearances around all equipment, termination panels and racks

Data centers must be carefully planned PRIOR to building to assure compliance with all applicable codes and standards. Design considerations include site and location selection, space, power and cooling capacity planning, floor loading, access and security, environmental cleanliness, hazard avoidance and growth. In order to calculate the above needs, the architect and RCDD must know the components that will be housed in the data center including all electronics, cabling, computers, racks, etc. To provide this list it is important to predict the number of users, application types and platforms, rack units required for rack mount equipment and most importantly, expected or predicted growth.

Anticipating growth and technological changes can be somewhat of a "crystal ball" prediction. With the possible combination of storage islands, application islands, server platforms and electronic components literally being factorial, planning is as important to a data center as the cabling is to a network. The data center will take on a life of its own and should be able to respond to growth and changes in equipment, standards and demands all while remaining manageable and of course, reliable. Larger data centers are designed in tiers or zones (sometimes on different floors) with each tier performing different functions and generally with different security levels. Redundancy may be between different levels or different geographic locations depending on the needs of the users of the facility.

EQUIPMENT IMPROVEMENTS

In an effort to conserve space and lower costs within data centers, KVM switches have been on the market for quite sometime. KVM (Keyboard, Video and Mouse) switches allow a single keyboard, monitor and mouse to control multiple servers in a rack or the new blade servers that are entering the market. Newer versions of these switches allow this control to happen remotely as well as locally through the switch.

SAN (Storage Area Networks) and NAS (Network Attached Storage) devices have made sharing disk drives between servers or over the network a faster and easier alternative to the older server mirroring technologies. These devices can be attached via Fibre Channel, SCSI, or network cabling. IP based products are becoming prevalent that allow for the communications between the storage devices and network components to be either IP based or tunneled through IP. This makes these solutions far more scaleable and reliable than their predecessors. For more information on Storage Area Networks, you may visit www.siemon.com and view the whitepaper on these technologies.

Another plus in the data center world is that electronics are becoming smaller and more compact thereby conserving space on the data center floor. This can be seen in telecommunications switching equipment, servers, UPS solutions and various other components within the data center. Single chassis switches equipped with blades for various tasks replace the older versions where an entire switch unit was needed for each function. Servers and rack mounted appliance servers are also smaller than their counterparts of old.

DATA CENTER CABLING SYSTEM CONSIDERATIONS

The TIA TR-42.1.1 group was tasked with the development of the “Telecommunications Infrastructure Standard for Internet Data Centers.” “The scope of the working group included topologies and performance for copper and fiber cabling, and other aspects of the IT infrastructure that will enable these facilities to rapidly deploy new technologies. Although the standard was published prior to the requirements for 10GBASE-T, the design practices are solid for new technologies. The TIA/EIA has recently adopted TIA/EIA-942 ‘The Telecommunications Infrastructure Standard for Data Centers’. The requirements will consider the need for flexibility, scalability, reliability and space management.” (Source www.tiaonline.org). The National Electric Code (NEC) in Article 645 “Information Technology Equipment” and the National Fire Protection Association (NFPA) in NFPA-75 “The Standard for the Protection of Information Technology” have addressed these important factors. While these standards will provide guidelines, there are specific design elements that will vary with each data center and its housed equipment. General considerations that will apply to all data centers include:

- Standards based open systems
- High performance and high bandwidth with growth factors incorporated
- Support for storage devices (i.e. Fiber channel, SCSI or NAS)
- Support for convergence with growth factors incorporated
- High quality, reliability and scalability
- Redundancy
- High capacity and density
- Flexibility and expandability with easy access for moves, adds and changes
- BAS, voice, video, CCTV and other low voltage systems
- Incorporation of Data Center security and monitoring systems

Cabling may be copper (UTP, F/UTP, S/FTP) or fiber (SM/MM) which will depend on the interface of the equipment to which it is to connect. In many cases a combination of several media types will be used. It is in an end user's best interest to run cabling accommodating growth during the first cabling implementation. Pricing can be negotiated on a project basis saving money. Also moves, adds and changes can be costly and increase the risk of bringing down critical components that are in use. Typical practices allow for dark fiber (unused strands) to be run along with the active fiber. Equipment may be active or passive.

CABLING PATHWAYS

Data centers contain highly consolidated networks and equipment. This high consolidation requires high density cabling systems. Cabling pathways in the data center generally consist of a combination of access under a raised flooring system and overhead cable tray. Raised floors provide the benefit of aesthetic pleasure along with heat management and easy access to the hidden cables. Cables under a raised floor should be run in raceways (cabling channels) to protect them from power cables, security devices and fire suppression systems which may be run in the same environment. Power cables can be run either in conduit or in power raceways and should respect the minimum distances outlined in industry standard specifications. Pathways can help assure that air pressure is maintained throughout the remainder of the data center, facilitate future moves, adds and changes, and assure that cables are properly supported removing the likelihood of damage or degredation of performance.

The fiber cabling pathway and management in the data center should be provided by a dedicated duct system. This provides a safe and protective method for routing and storing optical fiber patchcords, pigtails and riser cables among fiber distribution frames, panels, splice cabinets and termination equipment. Fiber carries different stress and bend radius requirements than copper due to the fact that it carries light rather than electrical signals. Planning is required to assure that proper space allowances are provided.

ENCLOSURES AND RACKS

Equipment enclosures and rack space should be a very early consideration in the entire design process. Identification of equipment and the number of rack units used will determine the number of racks needed for installation. Rack mounted equipment is expressed in xRU, with x representing the number of rack units (1-3/4 in. rack space). Some equipment also carries buffer or air requirements for separation from other equipment. Racks are standardized on a 19 in. equipment mounting width. Larger versions and larger cabinets are available.

All racks should be properly labeled as should all equipment contained therein. All racks/cabinets should be properly labeled as should all equipment contained therein, being careful not to label the spaces with any information that could pose a security risk. In most compliance related industries, it is now a requirement that networks are fully documented and that the documentation is maintained. TIA-942 suggest the use if a grid system so that each cabinet can be identified by its position on the flooring grid. Equipment enclosures and racks should contain the required cabling and should utilize wire management. Equipment enclosures and rack should be placed in locations allowing 4 ft. from the center of the rack to the wall behind with a minimum clearance of 3 ft. in front. Should equipment be contained in the rack, a 6 ft. clearance should be allowed. ANSI TIA/EIA and the NEC codes should all be consulted for proper placement of all components within the data center. In raised floor environments equipment enclosures and rack placement should also consider floor tile layout in order to prevent "a land-locked" situation. Cabinet enclosures will have varied positions and clearances due to the size of the cabinets and any airflow requirements, but typically 4 ft. in front of the cabinet is maintained (two full tiles) with one full tile plus the remaining tile space at the rear of the cabinet comprised the clearance at the rear.

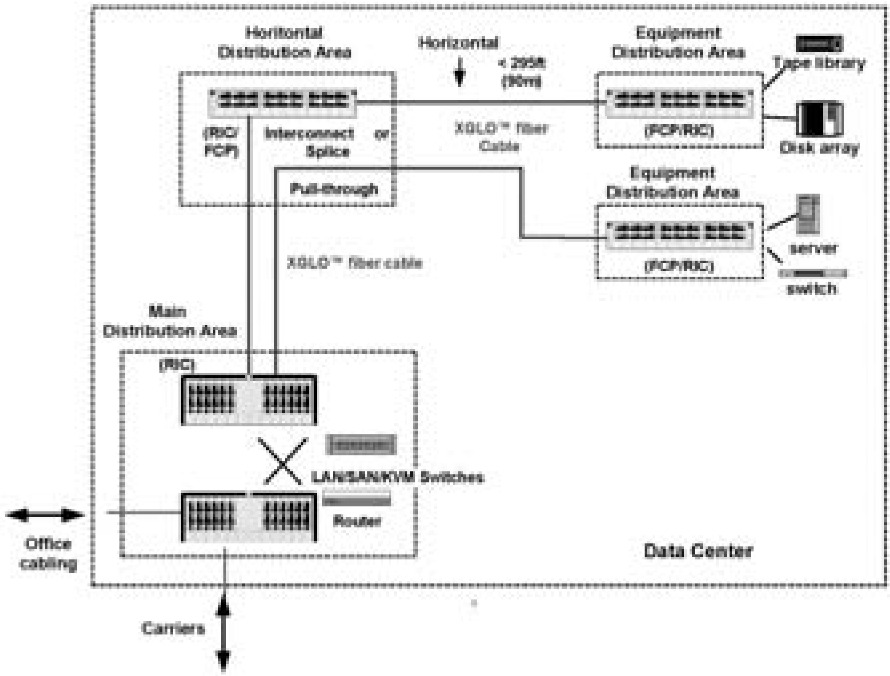

Figure-1: Data Center Centralized Siemon 10G ip™ XGLO Optical Fiber Cabling System Structure

For information on these and other systems from The Siemon Company, please visit www.siemon.com or consult one of our global team of certified installers, resellers or distributors.

Best Practices for Data Center Infrastructure Design

Cabling Distances and Space Planning

When designing and laying out a data center, understanding best practices as well the pros and cons for each type of data center is critical. The TIA 942 data center guidelines are very specific that horizontal and vertical cabling should be run accommodating growth, so that these areas do not have to be revisited. It is also specific about equipment not being directly connected unless it is specifically required by the manufacturer. This is in line with other standards documents such as ANSI/TIA/EIA 568-B that design for open systems architecture. So the question is raised: what is the best way to do this for a 10Gb/s environment?

There are considerations outside of the cable plant and number of connectors alone: usability, scalability, costs and the ability to perform Moves, Adds and Changes (MAC's). Additionally, some limitations exist based on the category of the cabling system. Copper and fiber distances may vary with the type of cabling system selected. We will discuss some of those parameters and their potential impact on data center designs.

All copper channels are based on a worst case, 100 meter, 4 connector model. ISO/IEC 24764 (draft), TIA-942, ISO/IEC 11801 Ed2.0 and recommendations from electronics manufacturers suggest that the fixed horizontal portion of the channel be a minimum of 15m (50 ft.). While some shorter lengths may be supported in other portions of the channels, there is a requirement in zone distribution and consolidation points for this minimum distance. When moving to 10Gb/s electronics, the 15m minimum will likely exist for all horizontal cables due to recommendations from electronics manufacturers and that all models within IEEE are based on a minimum 15m distance.

The 15m length is also dictated by signal strength issues as your signal is strongest in those first 15m. This can create issues with two connectors in close proximity. By providing at least 15m to the first connection point in the channel, you are allowing the attenuation to reduce the signal strength at the receiver or between components. In order to achieve the 15m distance, two options exist: either provide space in the pathway to take up the distance or create service loops under the floor. Service loops should not be a loop, but rather a loosely configured figure 8 for UTP systems. However, this configuration is not a requirement for F/FUTP or S/FTP systems due to their inherent resistance to interference. Bear in mind that the additional cable will consume more pathway space.

Copper distances for category 6A twisted pair cabling are limited to 100m for all channels. 10GBASE-T running on category 6/class E cabling will be limited to less than 37m depending upon the scope of potential mitigation practices to control alien crosstalk. It should be noted that the purpose of TSB 155 is to provide parameters for the qualification of existing category 6/Class E applications for use of 10GBASET, TSB 155 should not be used for designing new installations.

Fiber channel lengths vary based on the grade and type of fiber and type of interface. Understanding these limitations will assist in the design and layout of the data center space. If you are utilizing 10GBASE-CX4 or Infiniband, you are distance limited to a maximum of 15m. The following chart summarizes the distances for all 10Gb/s applications and their associated cabling systems.

Application	Media	Classification	Max. Distance	Wavelength
10GBASE-T	Twisted Pair Copper	Category 6/Class E UTP	up to 55m*	
10GBASE-T	Twisted Pair Copper	Category 6A/Class E_A UTP	100m	
10GBASE-T	Twisted Pair Copper	Category 6A/Class E_A F/UTP	100m	
10GBASE-T	Twisted Pair Copper	Class F/Class F_A	100m	
10GBASE-CX4	Manufactured	N/A	10-15m	
10GBASE-SX	62.5 MMF	160/500	28m	850nm
10GBASE-SX	62.5 MMF	200/500	28m	850nm
10GBASE-SX	50 MMF	500/500	86m	850nm
10GBASE-SX	50 MMF	2000/500	300m	850nm
10GBASE-LX	SMF		10km	1310nm
10GBASE-EX	SMF		40km	1550nm
10GBASE-LRM	All MMF		220m	1300nm
10GBASE-LX4	All MMF		300m	1310nm
10GBASE-LX4	SMF		10km	1310nm

* As defined in 802.3an

THE LAYOUT...WHERE AND HOW TO CONNECT

When designing a cabling infrastructure, too often cost is the deciding characteristic of the channel selected. However, once all elements are considered, a design with higher initial cost may have a lower overall cost of ownership to a company that has a lot of MAC activity. The most important concern is that designers are familiar with all aspects of the different configurations available to make the best selection possible. A listing of cost, flexibility and performance is listed below.

Model	Cost	Flexibility	Performance
2-Connector	Lowest	Lowest	Highest
3-Connector with CP	Medium	Medium	Medium
3-Connector with CC	Medium	Medium	Medium
4-Connector	Highest	Highest	Lowest

SPACE PLANNING OPTIONS

The MDA (Main Distribution Area) is considered the core of the data center. Connectivity will be needed to support the HDA (Horizontal Distribution Area). Following TIA-942 recommendations and utilizing EDA (Equipment Distribution Areas) and ZDA (Zone Distribution Areas), we would like to present four design options for consideration.

OPTION ONE

Option One is to run all fibers and copper from the core horizontal distribution areas and equipment distribution areas to a central patching area equipped with patch panels. This provides one central area for patching all channels.

There are several benefits to this design. First, all cabinets can remain locked. As patching is done in a central area — there is no need to enter a cabinet at any time unless there is an actual hardware change. For industries that are governed by compliance and security related issues, this may provide a greater benefit by reducing physical access to connections. Intelligent patching can be added to the patching field to increase security by automatically monitoring and tracking moves, adds and changes in that environment.

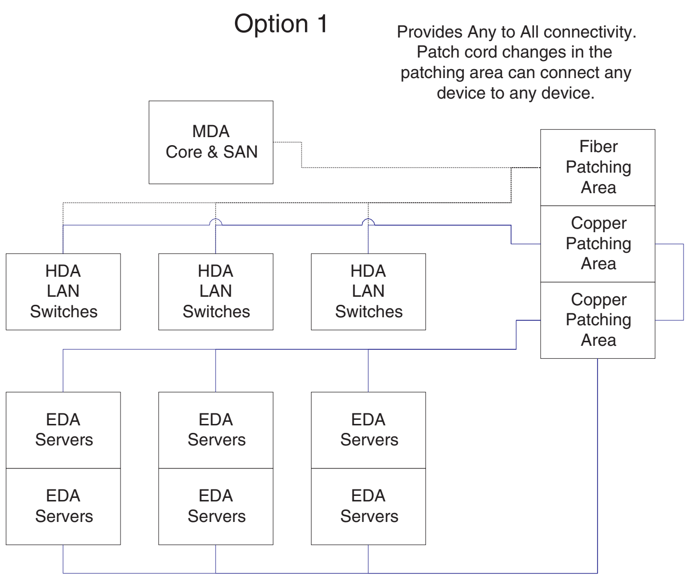

Note: Black lines are Fiber, Blue lines are Copper

Another advantage is that all ports purchased for active gear can be utilized. With the ability to use VLANs, networks can be segmented as needed.

In other scenarios, entire switch blades are likely dedicated to a cabinet of servers. However, if there are insuffcient server NICs to utilize all ports, then the idle ports become costly and inefficient. For instance, if a 48 port blade was dedicated to a cabinet at location XY12 but there was only 6 servers with two connections each, 36 ports were paid for but remain idle. By utilizing a central patching field, the additional 36 ports can be used as needed elsewhere in the network, thereby lowering equipment and maintenance costs which are far more expensive than the cable channels.

OPTION TWO

Option Two is to place patch panels in server cabinets that correspond directly to their counterparts in the switch cabinets. In this scenario, switch blades/ports will be dedicated to server cabinets. This may be easier from a networking perspective, but may not provide the best usage of all ports in the active electronics. Extra ports can be used as spares or simply for future growth. However, if an enterprise is planning to implement blade technology where server density may decrease per cabinet, this may not be a cost effective option.

For the switch cabinets, the type of copper cabling chosen will be a significant factor due to the increased UTP cable diameters required to support 10GBASE-T. In reality, cabinets and cabling (both copper and fiber) are changed far less frequently than the active electronics. But with the new category 6A UTP cable's maximum diameter of .354 inches (9.1mm), pathways within the cabinets may not provide enough room to route cable and still provide the structural stability necessary. It is always recommended that fill percentage calculations be addressed with the cabinet manufacturer. Moving the patch panels to adjacent locations or implementing a lower switch density may be required. While moving switches into open racks with adjacent patch panels provides a solution, this is only recommended if proper access security processes exist and some form of intelligent patching or other monitoring system is used so that network administrators can be notified immediately of any attempt to access switch ports.

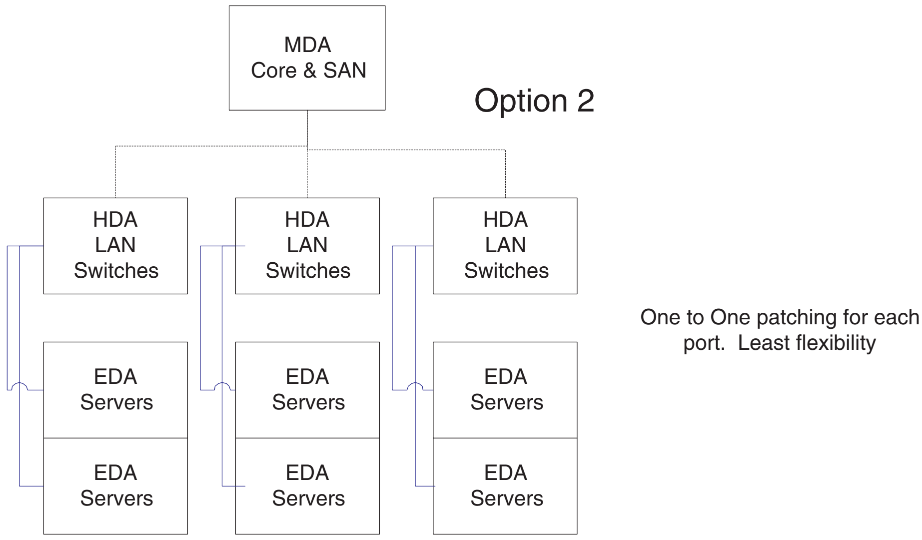

Note: Black lines are Fiber, Blue lines are Copper

OPTION THREE

Option Three consists of providing consolidation points for connections. These can be either connecting blocks or patch panels. This allows for a zoned cabling approach, but may lead to higher moves, adds and changes costs. It is also difficult to design within the parameters of a 4 connector channel when using Zone distribution.

The other disadvantage to the consolidation point model is that the changes take more time than swapping a patch cord if the pair count changes. Depending on the location of the consolidation point, there may be additional risks from loss of static pressure under the floor when removing floor tiles ending up with more than 4 connectors in a channel, or harming existing channels during changes.

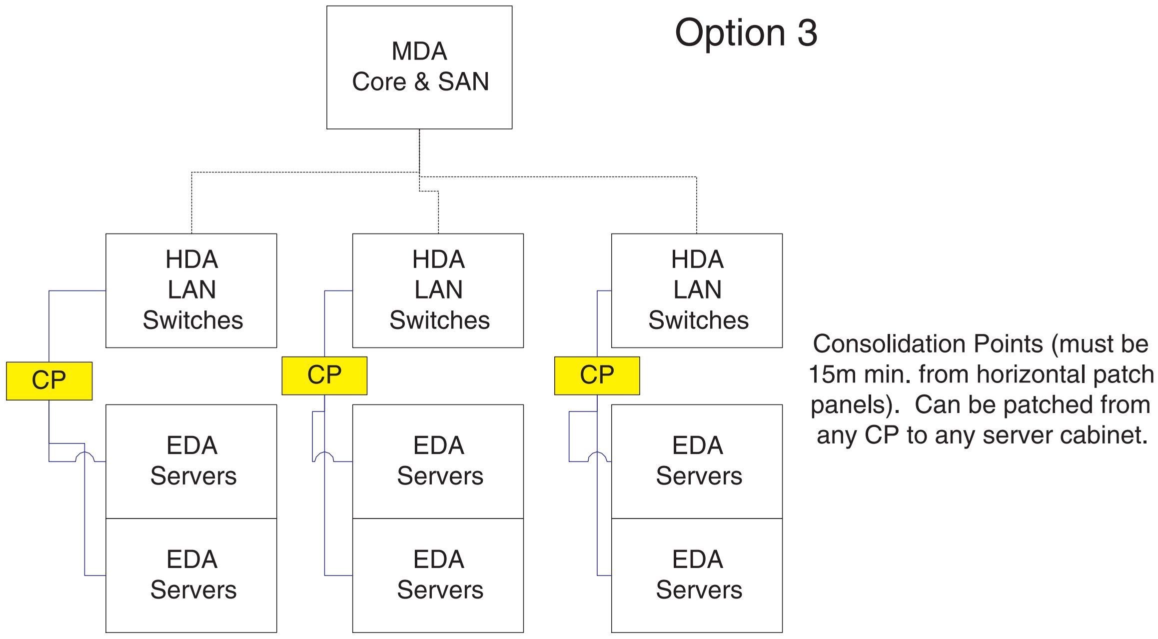

Note: Black lines are Fiber, Blue lines are Copper

OPTION FOUR

A final option is to have all server cabinets and switch cabinets in a row, terminating to a single patching field for the row, rather than to a central location. Core connections from the MDA are brought into this patching field. This option can work well in ISP or other environments where cross department/customer functionality is not desirable or tolerated. This option provides a bit of best of both worlds in that there will be some spare ports, but also the floor tiles will not have to be lifted to perform MAC work. While this is very similar to the first option, the segmentation can make it easier for network administrators and physical plant technicians to coordinate efforts. Additionally this style of design provides for flexibility in the ever changing environment of shrinking and expanding storage/networking requirements over time.

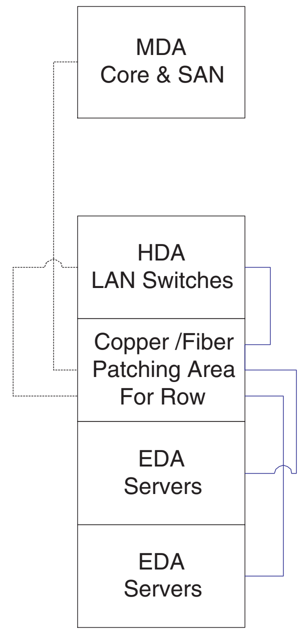

Option 4

All Patching done within respective rows.

Note: Black lines are Fiber, Blue lines are Copper

CONCLUSION

Whichever cabling choice or space option is made, the key step is planning. Siemon has resources to assist in the layout and planning or just as a second pair of eyes for any project. For more information, and additional resources go to www.siemon.com.

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 IDC[1], "281 exabytes of information existed in 2007, or about 45Gb 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, 10Gb/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 2010 [2]. 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."

10Gb/S INFRASTRUCTURE OPTIONS

There are several cabling alternatives available over which 10Gb/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 2.5Gb/s, and 4 lanes result in 10Gb/s operations in SDR (Single Data Rate) mode and 20Gb/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 A5[4]. 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, 10Gb/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 1Gb/s speeds, balanced twisted-pair compatible electronics offer better latency performance than fiber; however, considering latency at 10Gb/s, currently fiber components perform better than balanced twisted-pair compatible 10GBASE-T electronics, but not as well as 10Gb/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 10Gb/s or 40Gb/s copper connectivity and equipment."

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 10GbE 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 100Gb/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.

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

In many cases for both optical fiber and balanced twisted-pair cabling, an upgrade from 1Gb/s to 10Gb/s will require a change of the Ethernet switch, as older switch fabrics will not support multiple 10Gb/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 10Gb/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 5.5 W per port 10GBASE-T chip from Aquantia[6].

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.

It is further noted that IEEE is working on Energy Efficient Ethernet (802.3az) technology that will allow links to autonegotiate down to lower speeds during periods of inactivity – a capability which could reduce power by an estimated 85% when negotiating from 10Gb/s to 1Gb/s, and even further for lower speeds. Average power per 24-hour period will be far less when Energy Efficient Ethernet is built into future generation 10GBASE-T chips. This potential power savings is not available for optical fiber as there is no ability to autonegotiate over optical fiber.

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.

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	$ 379.09	$ 1,185.00	$ 533.25	$ 2,097.34
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 by $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 Infinband, 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 1Gb/s to 10Gb/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 1Gb/s and 10Gb/s and assuming the standards-based 10x the performance at 3x the cost, a single end-to-end channel supporting both speeds is $2,097.34 which translates into a savings of $3,004.44.

In a data center with five hundred (500) 10Gb/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 100Gb/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

The Hidden Costs of 10Gb/s UTP Systems

With both the 10GBASE-T application standard for 10Gb/s Ethernet over copper and the category 6A cabling standard now ratified, there has been much flurry over UTP versus F/UTP and S/FTP cabling systems.

While a consideration, channel costs are only a part of the equation in today's 10Gb/s world. In order to properly evaluate the benefits of each 10Gb/s copper cabling option, it is important to understand all the facts. This paper will provide current facts regarding 10Gb/s channels, pathways, labor and performance.

The systems supported in the 10GBASE-T 802.3an standard are:

- Limited distance category 6 (designed for existing installations)
- Category 6A (screened and unshielded)
- Category 7 (fully shielded only)

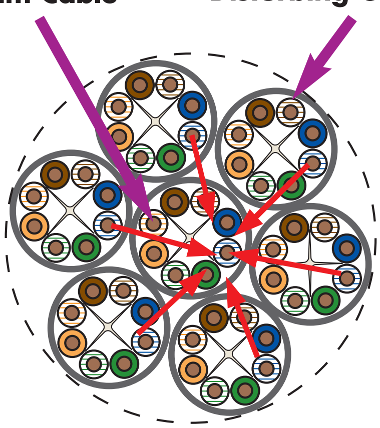

Understanding the feasibility of different types of 10Gb/s capable copper systems requires a brief explanation of alien crosstalk. Alien crosstalk (AXT) is cable-to-cable noise that is seen at higher frequencies and is more problematic in unshielded (UTP) systems. The noise coupling occurs on like pairs (for instance blue-white to blue-white) as they will have roughly the same twist lay. This is not a concern with screened or shielded systems as the shield prevents alien crosstalk from affecting adjacent cables.

Existing category 6 systems are the most costly for 10Gb/s due to the amount of mitigation practices necessary to support 10Gb/s, which translates into labor dollars. Mitigation techniques for category 6 include unbundling all cables, changing to shielded patchcords, reterminating to 6A hardware, and energizing every other switch or patch panel port only, to name a few. Testing these systems requires a significant amount of time because after applying one mitigation technique, you may need to try another and retest, then another and retest, etc. For this reason, it is recommended that new installations utilize a category 6A system or better for 10Gb/s transmissions.

In order to mitigate the effects of alien crosstalk at 10Gb/s, a new cabling category/class was introduced. Designated as augmented category 6 or category 6A by the TIA and,as class E_A within ISO/IEC, the standard applies to screened F/UTP systems as well as UTP. Category 7/Class F was accepted without modification and in fact for years was the only standardized 100 meter 10 Gigabit capable copper solution.

Among the most immediately noticeable differences between category 6 and 10Gb/s capable category 6A is that the maximum allowable diameter of the category 6A cable has been increased to 0.354 in. (9.1mm). This increase allows for the separation between the disturbing cable and the cable that is disturbed in UTP configurations. Because F/UTP cable eliminates AXT through the use of an overall foil screen, this additional separation is not required and has typically allowed category 6A F/UTP cable to be smaller than its UTP counterparts. While there are some newer smaller diameter 10Gb/s capable UTP cables on the market, it is still important that they pass the AXT testing parameters in the 6-around-1 configuration. End users should be wary of waiving AXT field testing for UTP systems. This is being done in some cases because AXT field testing can add up to 3.5 - 4 hours for each 24-port patch panel, but in these cases the manufacturer and the end-user are taking a risk that they wouldn't take if installing a screened or shielded system.

In order for UTP cables to remain viable as a 10Gb/s option, other factors must be considered. One is pathway fill. Should the time come where the pathway fill is greater than 40%, the cables on the bottom may become crushed and the resulting deformation of the cable will undo the separation benefits. This will certainly affect AXT and other performance parameters. Noise will be introduced back into the system and this is one of the hardest problems to troubleshoot as it can be sporadic and intermittent.

Noise, in the form of AXT, occurs not only between cables but also in closely placed connectors. In order to mitigate AXT at the patch panel, various methods are used, such as increased spacing and varying punch down practices, the latter of which may increase installation time and certainly increases the potential for errors. In some installations, one end of the cable is wired differently than the other, further complicating installation. The same AXT protection offered by shielded cable is extended to shielded connectors, eliminating the need for varied termination practices.

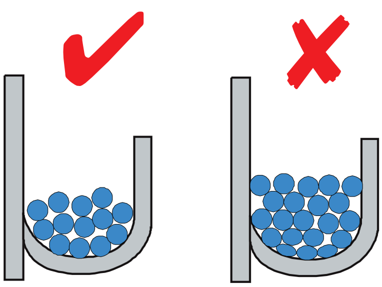

Illustration A

Another potentially costly caveat to a 10Gb/s UTP system is that you cannot mix different categories of unshielded cable in the same pathway. Bundling of category 6A UTP cabling with other category UTP cables can result in alien crosstalk levels that exceed industry requirements and potentially impact operation of the 10GBASE-T application. This means that if you choose category 6A UTP, you must use it for everything within the shared pathway! In a typical installation, there is a combination of cabling categories within a single pathway. In the horizontal, there may be building automation systems, detection systems, access control, analog voice or VoIP. Wherever such systems are installed, the only options are to run them all on the more expensive category 6A UTP cable plant or run additional pathways and spaces to separate the lower categories of UTP cabling.

In a data center that will underrgo a partial upgrade for 10Gb/s as new equipment is phased in, there will likely be access control, fire monitoring, cooling monitoring and other systems that do not require 10Gb/s copper cable systems. Again, in the UTP scenario, you are forced to do a full upgrade for all channels to category 6A UTP or provide separate pathways and spaces. With real estate at a premium, there may not be enough room for additional pathways and spaces, especially in equipment cabinets. Screened or shielded cabling may be bundled and/or share pathways with all other categories of cabling without concern.

If we examine just the cabling installation, without respect to additional pathways, the cost for 500 drops fully installed would look similar to the table at right.

Cat 5e/Class D UTP	$ 85,492.92
Cat 6/Class E UTP	$115,848.75
10G 6A UTP	$169,372.08
10G 6A F/UTP	$188,046.67
TERA - Cat 7_A/Class F_A	$232,386.67

The figures in the table use plenum cable, retail pricing, $65 per hour average installation labor rate, a 50 meter channel, 2 connectors and 2 patch cords. Of course, these figures will vary by installation. A more thorough explanation of the total cost of ownership (TCO) model available at: http://www.siemon.com/us/white_papers/06-05-18-tco.asp. In addition, an interactive whitepaper is available for end users through their local Siemon sales representative.

Assuming these averages for 500 drops, we can see a significant savings over category 6A UTP with a screened system when two categories are needed in the same pathway. It is likely that category 5e and 6 will both support voice applications over the next 10 years as it is not expected that voice applications will go above gigabit in that time.

In new building construction or during an upgrade, in order to use category 6A UTP cabling you would need to run all channels as category 6A UTP or run separate pathways. Assuming two drops per work area, the total cost for 500 category 6A UTP would be $169,372.08. The same scenario utilizing one category 5e UTP channel for voice, modem, or lower speed connections and one category 6A F/UTP channel for higher speed data would decrease costs to $127,432.50. Similarly, one category 6 UTP and one category 6A F/UTP in the same pathway would cost just $151,947.71. This translates into a savings of $41,393.59 and $36,098.96 respectively.

In an existing building, it may be impossible to add new pathways for high speed UTP cable, or to accommodate larger diameter category 6A UTP cable. The same may hold true for a data center that is at capacity. If a full 10Gb/s upgrade is performed, it is likely that KVM, modem connections, monitoring and control systems will continue using category 5e or category 6 UTP for several years. In this case, the same cost savings benefits would apply tif installing a screened system to support 10Gb/s channels within the same pathways that also carry category 5e or 6 UTP channels.

Running secondary trays and pathways to allow category 5e and 6 systems to coexist in a network with category 6A UTP would significantly impact project cost for both materials and labor; costs that are not required with screened or shielded systems used for high speed data.

Those who argue against shielded systems state the additional termination time and cost of grounding/bonding as an additional cost. The fact is that you must have a grounding and bonding (earthing) system for every network, even if it is 100% fiber. The only incremental, albeit small cost increase, is the cost to install one piece of #6 AWG wire per 24 port patch panel that attaches to the equipment rack (which should already be attached to the existing grounding system). While a screened or shielded connector may take slightly longer to terminate than a UTP connector, this is not significent enough to impact the savings and performance benefits realized.

At Siemon, we offer all systems for our end users, however the vast majority of our end users are moving to 10Gb/s and taking advantage of the benefits of screened and shielded systems. In fact, in a recent Data Center study performed by BSRIA, 75% of respondents who plan to choose copper cabling for their 10G links plan for shielded cabling." For more information on screened and shielded systems including the top reasons to specify a screened/shielded cabling system, please visit www.siemon.com or contact your local sales representative.

Light It Up

Fiber Transmissions and Applications

Several grades, speeds and applications exist for fiber optic network cabling plants. Two factors will govern the speeds for your fiber: the grade of fiber and the light source used to launch the data onto the fiber. This can also increase depending on which type of multiplexing is used to carry multiple signals across the same fiber strands.

Fiber applications include fiber backbone operations, fiber to the carrier options and fiber to the desktop operations. There is also a growing number of fiber to the premise (FTTP) applications including fiber-distributed video. Fiber based storage and network interfaces are readily available. The fiber channel industry association has been instrumental in setting operating parameters for fiber-attached storage in data centers. Regardless of your intended use for fiber, the characteristics, losses and bandwidth are critical to the success of your fiber network.

TRANSMITTING ON FIBER

The combination of fiber light source, grade of fiber and core diameter of the fiber will determine not only the functional distance of the channel, but also the speed of the communications. In the recent past, multimode fiber has seen a transition from 62.5/125 (core/cladding) micron fiber in various grades to 50/125 micron fiber, also available in various grades, including the highest performing laser optimized fiber. Multimode fiber is typically the choice for shorter distances. The light is carried is several paths or modes. A mode is the transmission of light in a steady state. Singlemode fiber allows for a single path/mode and is suitable for longer-range applications. The active electronics supporting multimode fiber are less costly than for singlemode.

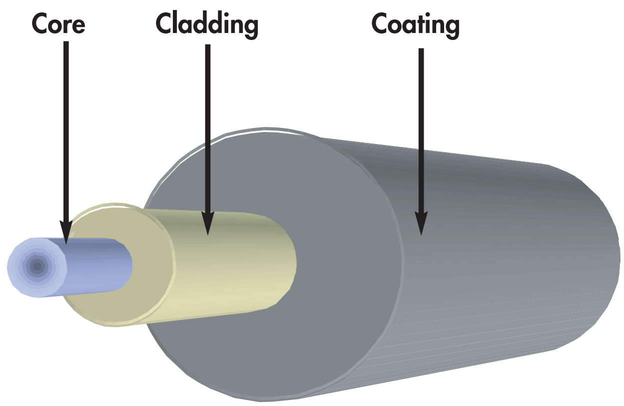

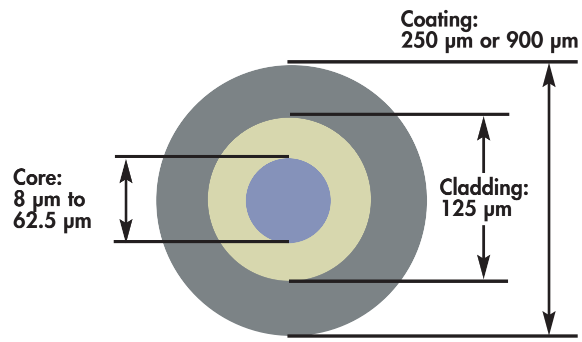

The speed of light in a fiber is measured by its index of refraction. This is an equation that compares the speed of light in one medium to its fastest speed in a vacuum. This light travels in the non-visible spectrum, typically 850nm, 1300nm or 1550nm. The visible spectrum ends at roughly 750nm. This spectrum can carry both analog and digital signals. The devices that launch the light source convert those signals to light pulses within the spectrum required for transmission. At the other end, another device detects the light pulses within each mode and converts them back to analog or digital signals to be understood by its attached device.

The speed of fiber is measured in hertz or cycles per second. Each cycle equates to one pulse or lightwave. One hertz equals one cycle. As fiber transmits at a very high rates of speed, we measure the pulses/cycles that a device launches onto the fiber in megahertz (millions of hertz per second). Because light does not degrade over distance as rapidly as electrical signals, the attenuation or loss of the signal is lower than it's copper counterparts. This allows the light waves to travel greater distances.

FIBER LOSS AND SIGNAL DEGRADATION

Like copper, fiber signal loss or attenuation is measured in decibels (dB). Fiber attenuation will increase with each connector or splice. Typically the splice loss is around 0.2 dB per splice. However, poor terminations can increase this number. A fiber loss budget compares actual loss to the budget (or acceptable loss) based on the number of splices in the channel. In order to measure this loss, a power meter and light source must be used. Testing a fiber strictly with an Optical Time Domain Reflectometer (OTDR) provides characterization of the fiber segment, but does not provide definitive performance.

There are two types of attenuation that are encountered in optical fiber cables: intrinsic and extrinsic. Intrinsic attenuation is inherent to the fiber and is introduced during the manufacturing process. An example of this would be impurities or inhomogineities within the glass. These causes the light signals to either be absorbed or scattered and is the reason that some fibers support greater application distances than others. Advancements in manufacturing have introduced a new grade of multimode fiber known as laser-optimized

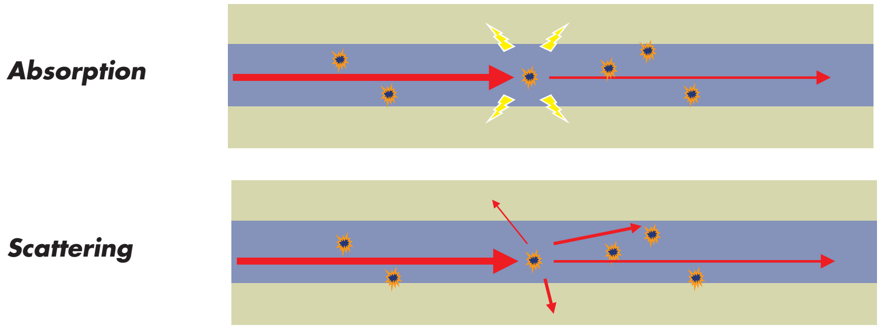

fiber. This fiber incorporates two essential manufacturing improvements. First is the elimination of the anomalies listed above by reducing impurities in the fiber core. Second is increased control of the index of refraction which reduces modal dispersion ensuring that all modes arrive at the receiver at essentially the same time. The combination of these improvements greatly increases the bandwidth capacity of the fiber which results in the support of higher speed applications including 10Gb/s transmissions as well as increased transmission distances. TIA refers to this as laser-optimized fiber while ISO/IEC refers to this as OM3 grade fiber.

Extrinsic attenuation is introduced during cable handling. Examples would be small mechanical stresses (microbends) or bend radius violations (macrobends) and result is light being refracted out of the core. In all fiber installations, the manufacturer-specified bend radii limits must be respected.

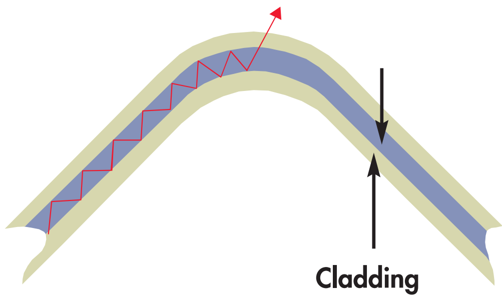

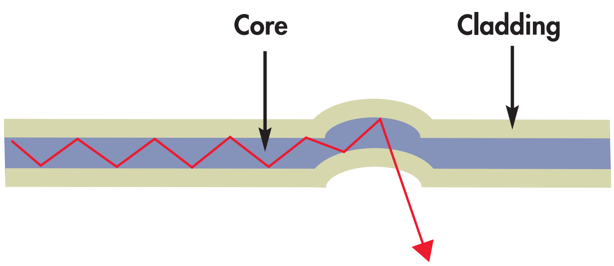

In fiber, the signals and pulses must be understood by the receiver at each end. Running too long a fiber channel can cause errors, as can excessive splices, poor fiber quality and poor installation. Any anomaly that keeps a receiver from registering the pulses will equate to a bit error.

LAUNCHING THE LIGHT

In order to send a signal over a fiber optic cable you need a light source. This can be an LED (Light Emitting Diode), a Laser (Light Amplification by Stimulated Emission of Radiation) or a VCSEL (Vertical Cavity Surface Emitting Laser). Lasers and VCSELs provide a stronger, more focused light source and therefore can transmit greater distances than their LED counterparts. Equipment that generates the signal via the latter two technologies is more expensive than an LED source.

Regardless of the type of fiber, the act of placing the light pulses onto the fiber is called the "launch". The launch method can vary from an overfilled launch to restricted mode launch. As discussed previously, the light path is called the mode. In an overfilled launch, the light introduces a signal that is greater in size than the actual core of the fiber. This allows all modes to become excited. In a restricted mode launch, a smaller core of light is introduced which excites only certain modes on the fiber. In singlemode, only a single path or mode is excited.

Within the fiber, light pulses can spread over distance, called dispersion. When the pulses overlap, they can limit the receiver's ability to register distinct pulses, thereby limiting the bandwidth on a fiber. Light travels at different speeds in different colors as well. In order to counteract some dispersion loss, the light source can provide what is called a restricted launch, typically used for higher speed applications. Rather than filling all modes within a fiber with light, only certain modes are excited thereby restricting the range of pulses and the affects of dispersion.

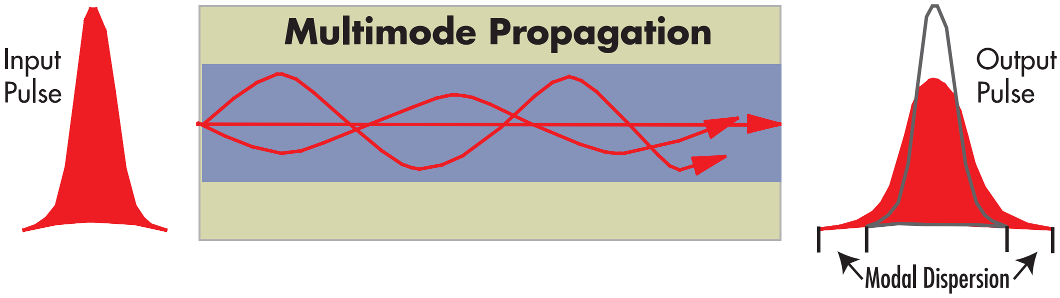

At longer wavelengths for gigabit speeds running on older grades of 62.5/125 micron multimode fiber, the restricted launch causes differential mode delay (the signals do not arrive at the receiver at the same time). For these applications, mode conditioning patch cords must be used. These cords provide an offset so that the light does not enter directly in the center of the fiber core. By offsetting the beam to an area outside the center of the core, dispersion is minimized. One set of restricted launch cords must be used at each end of the system.

The bandwidth of the fiber is the information transmission capacity of the fiber. It is inversely proportional to the amount of dispersion. Thus, the extent to which dispersion can be controlled essentially determines the usable bandwidth of the fiber.

ISO/IEC 11801 Ed2.0 defines three optical types of multimode fiber. OM1 primarily comprises historic 62.5/125 micron fiber. OM2 has an effective bandwidth of 500 MHz·km at both wavelengths and represents standard grade 50/125 micron fiber. OM3 has 1500/500 MHz·km bandwidth for overfilled launches and 2000 MHz·km with a restricted mode launch and is referred to by TIA/EIA as "laser optimized" 50/125 micron fiber.

ISO Multimode Bandwidth Specifications

			Minimum Modal Bandwidth MHz • km	
Optical Fiber Type	**Core Diameter (µm)**	**Wavelength (nm)**	**Overfilled Launch (OFL)**	**Restricted Mode Launch (RML) 850 nm**
OM1	50 or 62.5	850	200	Not Specified
		1300	500	Not Specified
OM2	50 or 62.5	850	500	Not Specified
		1300	500	Not Specified
OM3	50	850	1500	2000
		1300	500	Not Specified

Note: Effective laser launch bandwidth is assured using differential mode delay (DMD) as specified in IEC/PAS 6073-1-49.

TIA Multimode Bandwidth Specifications

		Minimum Modal Bandwidth MHz • km	
Optical Fiber Type	**Wavelength (nm)**	**Overfilled Launch (OFL)**	**Restricted Mode Launch (RML)**
62.5/125µm	850	160	Not Required
Mulitmode	1300	500	Not Required
50/125µm	850	500	Not Required
Multimode	1300	500	Not Required
Laser-Optimized	850	1500	2000
50/125µm Multimode	1300	500	Not Required

UNDERSTANDING THE EQUIPMENT SIDE

Each piece of active electronics will have a variety of light sources used to transmit over the various types of fiber. The distance and bandwidth will vary with light source and quality of fiber. In most networks, fiber is used for uplink/backbone operations and connecting various buildings together on a campus. The speed and distance are a function of the core, modal bandwidth, grade of fiber and the light source, all discussed previously. For gigabit transmissions, the IEEE approved distances are shown in the following table.

GBIC	Wavelength (nm)	Fiber Type	Core Size (micron)	Modal Bandwidth (MHZ/km)	Cable Distance
1000BASE-SX	850	MMF	62.5	160	722 ft. (220m)
			62.5	200	920 ft. (275 m)
			50	400	1640 ft. (500m)
1000BASE-LX	1300	MMF[1]	62.5	500	1804 ft. (550m)
			50	400	1804 ft. (550m)
			9/10	N/A	6.2 miles (10km)
1000BASE-EX	1550	SMF	9/10	N/A	43.4 to 62 miles (70 to 100km)

[1] *A mode conditioning patch cord is required.*

As you can see, depending on the type of fiber and the type of light source, the distances supported vary from 722 feet to multiple kilometers. Some are singlemode only as noted. The cost of each option increases with distance due to the type of light source. Each of these distances are maximums based on the type, quality and installation of the fibers. Additional loss introduced can vary the distances greatly. The GBIC (Gigabit Interface Converters) listed in column one are listed as SX (short haul), LX (long haul) and EX (extended haul). It is also important to note that using singlemode fiber for short distances can cause the receiver to be overwhelmed and an inline attenuator may be needed to introduce attenuation into the channel.

To notice the difference between speeds and quality, the following chart compares distances beginning with 100Mb/s applications through 10Gb/s applications. As you can see, anomalies in fiber, light sources and overall bandwidth speeds all affect the distances that you can carry a signal over the fiber. The distance increases with laser optimized fiber products and therefore would provide the greatest return on investment. The additional distances supported would eliminate the need for more costly electronics and repeaters within the fiber network.

Application	Core Wavelength	62.5µm 160/500	62.5µm 200/500	50µm 500/500	50µm 2000/500	SMF
100BASE-SX	850nm	300m	300m	300m	300m	
1000BASE-SX	850nm	220m	275m	550m	550m	
1000BASE-LX	1300nm	550m	550m	550m	550m	5km
10GBASE-SR*	850nm	28m	28m	86m	300m	
10GBASE-LR*	1310nm					10km
10GBASE-ER*	1550nm					40km
10GBASE-LRM	1300nm	220m	220m	220m	220m	
10GBASE-LX4	1310nm	300m	300m	300m	300m	10km

**These interfaces are also seen noted with an X in place of the R but can be used interchangeably.*

The interface for 10Gb/s fiber is called a Xenpack, as opposed to GBIC for gigabit. Both 10GBASE-SR and 10GBASE-LR have a counterpart for wide area communications to allow them to connect to SONET networks at 9.584640Gb/s (OC-192) through its wide area interface. These interfaces are known as 10GBASE-SW and 10GBASE-LW respectively.

With gigabit to the desktop becoming commonplace, 10Gb/s backbones are also becoming more common. The SR interfaces are also becoming common in data center applications and even some desktop applications. As you can see, the higher quality fiber (or laser optimized fiber) provides for greater flexibility for a fiber plant installation. Although some variations (10GBASE-LRM and 10GBASE-LX4) support older grades of fiber to distances 220m or greater, the equipment is more costly. In many cases, it is less expensive to upgrade fiber than to purchase the more costly components that also carry increased maintenance costs over time.

Cabling Infrastructure and Green Building Initiatives

In a recent study, Gartner analysts revealed that IT activity accounts for 2% of global CO_2 emissions, equivalent to the amount produced by the aviation industry. We typically think of emissions as coming from forms of transportation, heavy industry and power generation, but with recent work to address global issues we see that in fact, IT and IT related products can have an impact in several areas.

Siemon, a network infrastructure specialist, addresses a number of cabling-related areas in which a greener approach to IT is possible.

GLOBALLY GREEN - AN OVERVIEW

The drive to reduce emissions and other environmental harm from more “hidden” activities such as IT has spawned a number of international efforts. The most wide reaching initiatives in the IT marketplace are focused on “Green Buildings” — efforts aimed at reducing the environmental impact of commercial and residential spaces.

The WGBC (World Green Building Council) currently consists of members from the USA, Canada, Mexico, UK, United Arab Emirates, India, Taiwan, Japan, Australia and New Zealand. As of the date of this writing, the following countries have announced plans to participate in green initiatives: Argentina, Brazil, Chile, Egypt, Germany, Greece, Guatemala, Hong Kong, Israel, Korea, Nigeria, Panama, Philippines, Switzerland, Turkey and Vietnam, with more to likely follow.

While the initial focus is on renewable energy sources, power and energy savings and environmental protection of sites for new and existing buildings; further examination indicates that network cabling and infrastructure will impact the overall effort.

In the United States, the USGBC (United States Green Building Council) has issued LEED (Leadership in Energy and Environmental Design) guidelines that provide a road map for measuring and documenting success for every building type and phase of a building life cycle. Although they vary slightly region to region, most global WGBG participants have guidelines that are very similar to LEED. An overview of the USGBC LEED program is provided in the appendix of this document

While a portion of these documents’ scope falls outside of the cabling realm, Siemon wishes to illustrate some ways in which we can assist our end users in their Green efforts globally. In doing so, we primarily address LEED guidelines as put forth by the USGBC, as they are largely representative of related global guidelines such as: Indian, Canadian and Mexican proposed revisions; the BREEAM (BRE Environmental Assessment Methodology) certification; Green Star Program from Australia and New Zealand’s adaptation of the same; Japan’s Comprehensive Assessment System for Building Environmental Efficiency (CASBEE); and EEWH as set forth in Taiwan.

SOME CABLING-SPECIFIC STRATEGIES

Based on the general guidelines put forth by LEED, there are a number of cabling strategies that may be explored as green options and potential contributors to overall facility certification.

DATA CENTERS AND ENERGY CONSUMPTION

Energy conservation efforts are being introduced into the data center space at an ever–increasing rate with good reason. Current studies show that power alone represents from 30-50% of overall data center budgets. While a portion of the energy is consumed by the actual servers, switches, routers and other active gear, an additional power load is needed to cool this equipment. There is a cycle of cooling to power and power requiring cooling.

In order to have the most efficient cooling, cabling must be properly designed, remediated and routed to allow the air to flow in an unobstructed manner. TIA-942 and other complimentary data center standards around the globe suggest that horizontal and vertical cabling be run accommodating growth so that these areas do not need to be revisited. There are several reasons for this recommendation, including: eliminating the adverse affects of removing floor tiles and decreasing static pressure under raised floors during MAC work; assuring that pathways are run in a manner that will allow the flow of cold air in cold aisles to be unobstructed by cabling; and a potential benefit to cooling as the cabling can be installed to provide a baffle of sorts, channeling cool air into cold aisles.

A significant number of older data centers and even telecommunications areas have suffered from ill-managed MACs (moves, adds and changes) over the years, leaving abandoned cabling channels behind. These unused channels often create air dams which obstruct air flow, which could result in higher energy consumption as your cooling equipment will work less efficiently. While that problem alone should be enough to commission the removal of abandoned cabling, there may also be issues with the older cabling jackets not meeting current RoHS (Reduction of Hazardous Substances) requirements. In many cases, these older cables carry significant fuel load which can pose additional fire threats, and can release toxins such as halogens if ignited. Beyond the life and safety issues at risk, the proper removal and disposal/recycle of abandoned cable can remove a significant environmental risk.

Although removing abandoned cable will have a positive green impact, reducing the volume of potentially abandoned channels through proper management is an even better option. Intelligent infrastructure management systems (such as MapIT®), can provide a lights out advantage by allowing detailed monitoring of any MACs made. By providing a consistent and up to date diagram of the physical layer connections, channels can be managed and fully utilized before they become a management headache or a source of unchecked MAC work.

While the ability to keep the cabling channels in check will almost certainly reduce power consumption on the cooling side, intelligent infrastructure management can also reduce power needs of the active network equipment. When designed with a central patching field, an intelligent infrastructure management system can help ensure that all switch ports are utilized — decreasing the power needs for electronics by keeping unused ports to a minimum. The ability to patch into unused ports rather than adding additional switches can provide an energy savings which in turn translates into further cooling savings.

BENEFITS OF INSTALLING SYSTEMS WITH GREATER BANDWIDTH THAN CURRENTLY REQUIRED

When installing data cabling, it is in the end-user's best interest to install systems that will provide the maximum longevity. Currently, category 7/class F cabling is the highest performing cabling system on the market, with a category 7A/class FA standard due to publish soon. The latter is characterized to 1000MHz, or 1GHz per channel, which provides a significant amount of bandwidth above and beyond the latest 10Gb/s network speeds for copper. These higher bandwidth cabling systems are completely backwards compatible with older technology.

A recent whitepaper released by Siemon explores the ROI/TCO (return on investment and total cost of ownership) for cabling plants. It concludes that lower-performing cabling will cost significantly more over the entire life cycle of the cabling plant. When examining the green building initiatives, the reduction of materials that will need to be replaced over time is an even greater incentive to install higher-performing cabling.

For instance, installation of a category 5e system would mean replacement in a few years as 10GBASE-T is implemented to the desktop. Category 6 systems will require remediation (another visit from the installer) and certain replacement of longer channels. Each of these scenarios would have a negative impact on "green" ratings due to the waste of materials and additional site visits by contractors. The significant reduction in cables being removed and reinstalled, likewise results in the conservation of copper, aluminum and other natural resources.

PHY designers are always on the look-out for improvements that can support performance enhancements for their next-generation products. Moving to higher performing class F/FA, fully-shielded cabling systems such as TERA® will significantly reduce noise on the cabling channel which can result in a significant power savings in the active electronics by eliminating Digital Signal Processing (DSP) complexity used to suppress noise levels. A study presented jointly by Siemon and KeyEye Communications indicated that the use of fully-shielded cabling could offer a reduction of approximately 20% in the overall power budget related to 10GBASE-T chip architectures. The bulk of these savings would result from a reduction in the levels of DSP complexity associated with NEXT and FEXT cancellers.

Furthermore, low alien crosstalk levels exhibited in these channels will lead to greater signal to noise ratios which help a system achieve higher levels of robustness and reliability. Class F/FA cabling helps dramatically improve issues with noise budgets that factor into transceiver DSP complexity and power for worst case cabling lengths up to 100m. Processing and level requirements can be reduced with no loses in performance and the additional bandwidth provided by class F/FA cabling offers an end user an upgrade path to even higher signaling rates when needed in the future. While network equipment that is specifically designed to take advantage of the internal noise and SNR benefits provided by class F/FA cabling is not commercially available at this time, research clearly demonstrates the advantage in power utilization and latency that these cabling systems potentially offer to next-generation product designers.

POTENTIAL SOURCES FOR GREEN LEED CREDITS

Potential Contributions to LEED Credits

Product/Svc.	LEED Credit	Explanation
MapIT	MR 2.1 - Construction Waste Mgmt. - 50% MR 2.2 - Construction Waste Mgmt. - 75%	• Reduction of unnecessary channels due to undocumented/poorly managed MAC work
	MR 3.1 - Resource Reuse - 5% MR 3.2 - Resource Reuse - 10%	• Identification and utilization of unused cabling channels to limit installation of new channels
	EA 1 - Optimize Energy Performance	• Maximization of active port usage to limit the installation of unnecessary active equipment • Identification and utilization/elimination of abandoned channels to maximize pathway space/increase airflow for energy-efficient cooling
TERA	MR 2.1 - Construction Waste Mgmt. - 50% MR 2.2 - Construction Waste Mgmt. - 75%	• Cable sharing as a means to reduce number of installed cabling channels • Future-proof performance extends the lifecycle of the cabling, decreasing the frequency of cable removal/disposal and installation of additional cabling
	EA 1 - Optimize Energy Performance	• Shielded construction may limit noise sufficiently to reduce active equipment power consumption through elimination of DSP.
10G 6A F/UTP	MR 2.1 - Construction Waste Mgmt. - 50% MR 2.2 - Construction Waste Mgmt. - 75%	• Future-proof performance extends the lifecycle of the cabling, decreasing the frequency of cable removal/disposal and installation of additional cabling • Reduced cabling diameter reduces pathway infrastructure (cable tray, conduit, j-hooks) • Reduced cabling diameter reduces the use of cable jacket materials
	EA 1 - Optimize Energy Performance	• Smaller cable diameter maximizes pathway space/increase airflow for energy efficient cooling
10G 6A UTP	MR 2.1 - Construction Waste Mgmt. - 50% MR 2.2 - Construction Waste Mgmt. - 75%	• Future-proof performance extends the lifecycle of the cabling, decreasing the frequency of cable removal/disposal and installation of additional cabling
XGLO Fiber	MR 2.1 - Construction Waste Mgmt. - 50% MR 2.2 - Construction Waste Mgmt. - 75%	• Reduced cabling diameter reduces pathway infrastructure (cable tray, conduit, j-hooks) • Reduced cabling diameter reduces the use of cable jacket materials • Future-proof performance extends the lifecycle of the cabling, decreasing the frequency of cable removal/disposal and installation of additional cabling
	EA 1 - Optimize Energy Performance	• Smaller cable diameter maximizes pathway space/increase airflow for energy-efficient cooling
Trunking Cable	MR 2.1 - Construction Waste Mgmt. - 50%	• Factory termination eliminates onsite waste created by field terminations
	MR 2.2 - Construction Waste Mgmt. - 75%	• Faster and more efficient installation of trunk cables requires fewer contractor visits and smaller crews
	MR 3.1 - Resource Reuse - 5% MR 3.2 - Resource Reuse - 10%	• Modular design of trunks allows for on-site re-use
	EA 1 - Optimize Energy Performance	• Well-organized channels eliminate airdams in pathways caused by poorly managed individual channels to maximize airflow for energy-efficient cooling
GPS/Project Asst.	MR 5.1 - Regional Materials - 10% MR 5.2 - Regional Materials - 20%	• GPS/Project Assistance providing logistical efficiencies through use of local stock and labor

MANY PATHS TO GREEN BUILDINGS

While this paper covers a number of ways in which cabling infrastructure decisions may affect a Green Building effort, it is hardly comprehensive. As Siemon continues to align its drive to innovation with its longstanding commitment to the environment, more opportunities for global improvements in sustainable IT practices will arise.

APPENDIX:

INTRODUCTION TO LEED

A "green" building is a building that is constructed in a responsible manner that minimizes or eliminates the negative environmental impact of the building on the environment, its community and on the health of its occupants, and reduces natural resource consumption. Historically, how to define and standardize the green building is a long term challenge until the advent of green building rating systems. (LEED) is the most widely accepted national green building rating system. Through its use as a design guideline and third-party certification tool, LEED aims to improve occupant well-being, environmental performance and economic returns of buildings using established and innovative practices, standards and technologies. In fact, LEED has been the green building standard of choice for Federal agencies and state and local governments nationwide.

LEED promotes integrated, entire building design and construction practices and encourages awareness various green building benefits. LEED-based green design not only makes a positive impact on public health and the environment, it also reduces operating costs, enhances building and organizational marketability, potentially increases occupant productivity, and helps create a sustainable community. LEED typically recognizes performance in six key areas of human and environmental health: sustainable site development, water savings, energy efficiency, materials selection, indoor environmental quality, and innovation & design process.

Since 1998, Members of the U.S. Green Building Council representing all segments of the building industry developed LEED and continue to contribute to its evolution. LEED provides a roadmap for measuring and documenting success for every building types and phase of a building lifecycles.

While understanding LEED is helpful for any enterprise wishing to reduce their environmental impact, it is absolutely critical to gaining green building certification. Fortunately, LEED is a very user-friendly system. In fact, a major contributor to the success of LEED is the simplicity of its credit/point-based the rating system.

For each credit, the LEED standard identifies the detailed intent, requirements, and technologies or strategies to achieve the credit. One or more points are available within each credit, and points are achieved by meeting specified requirements.

The amount of points achieved will determine which level of LEED certification the project is awarded. There are (69) possible points and (4) levels of LEED certification available:

- Certified (26 to 32 points)
- Silver (33 to 38 points)
- Gold (39 to 51 points)
- Platinum (52 to 69 points)

It is important to note that individual products and services do not earn projects points.

DETAILED SUMMARY OF THE LEED AREAS OFFERING POTENTIAL CABLING-RELATED CREDIT OPPORTUNITIES

ENERGY & ATMOSPHERE (EA)	
EA Prerequisites	
1. Fundamental of the Building Energy System 2. Minimum Energy Performance 3. Fundamental Refrigerant Management	
EA Credits	**Eligible Points**
1. Optimize Energy Performance	1-10
2. On-Site Renewable Energy	1-3
3. Enhanced Commissioning	1
4. Enhanced Refrigerant Management	1
5. Measurement & Verification	1
6. Green Power	1
Total Possible Points	17

MATERIALS & RESOURCES (MR)		
MR Prerequisites		
1. Storage & Collection of Recyclables		
MR Credits		**Eligible Points**
1. Building Reuse	1.1: Maintain 75% of Existing Wall, Floors and Roof	1
	1.2: Maintain 95% of Existing Walls, Floors and Roof	1
	1.3: Maintain 50% of Interior Non-Structural Elements	1
2. Construction Waste Management	2.1: Divert 50% From Disposal	1
	2.2: Divert 75% From Disposal	1
3. Materials Reuse	3.1: 5%	1
	3.2: 10%	1
4. Recycle Content	4.1: 10% (post-consumer + 1/2 pre-consumer	1
	4.2: 20% (post-consumer + 1/2 pre-consumer	1
5. Regional Materials	5.1: 10% Extracted, Processed and Manufactured Regionally	1
	5.2: 20% Extracted, Processed and Manufactured Regionally	1
6. Rapidly Renewable Materials		1
7. Certified Wood		1
	Total Possible Points	13

CISCO TECHNOLOGY DEVELOPERS 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: **http://www.cisco.com/pcgibin/ctdp/displayProfile.pl?PARTNER_ID=54708**

ABOUT SIEMON

Established in 1903, Siemon is an industry leader specializing in the manufacture and innovation of high quality, high-performance network cabling solutions. Headquartered in Connecticut, USA, with global offices, manufacturing and service partners throughout the world, Siemon offers the most comprehensive suite of copper (unshielded and shielded twisted-pair) category 5e, category 6 (Class E), category 6A (Class EA) and category 7/7A (Class F/FA), and multimode and singlemode optical fiber cabling systems available. With over 400 active patents specific to structured cabling, from patch cords to patch panels, Siemon Labs invests heavily in R&D and development of industry standards, underlining the company's long-term commitment to its customers and the industry.

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 President of the BladeSystems Alliance. She participates in IEEE, the Ethernet Alliance and IDC Enterprise Expert Panels. She has one telecommunications patent and one pending.


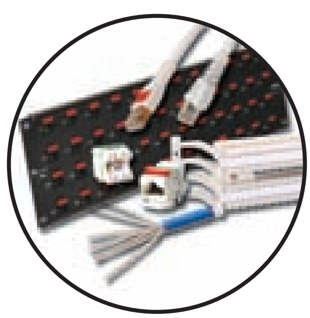
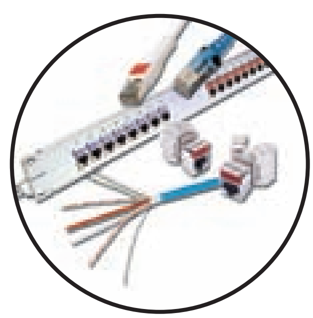
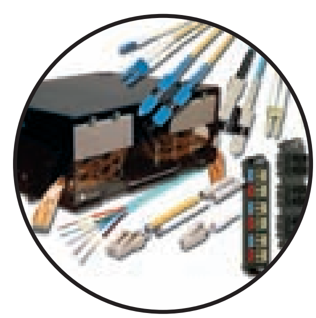

The Americas
Watertown, CT USA
Phone (1) 860 945 4200 US
Phone (1) 888 425 6165 Canada

Europe/Middle East/Africa
Chertsey, England
Phone (44) 0 1932 571771

Asia/Pacific
Shanghai, P.R. China
Phone (86) 21 6390 6778

Central & South America
Bogota, Columbia
Phone (571) 317 2121