Data Center E-Book

Deploying, Managing and Securing an Efficient Physical Infrastructure

www.siemon.com

SIEMON™

Table of Contents

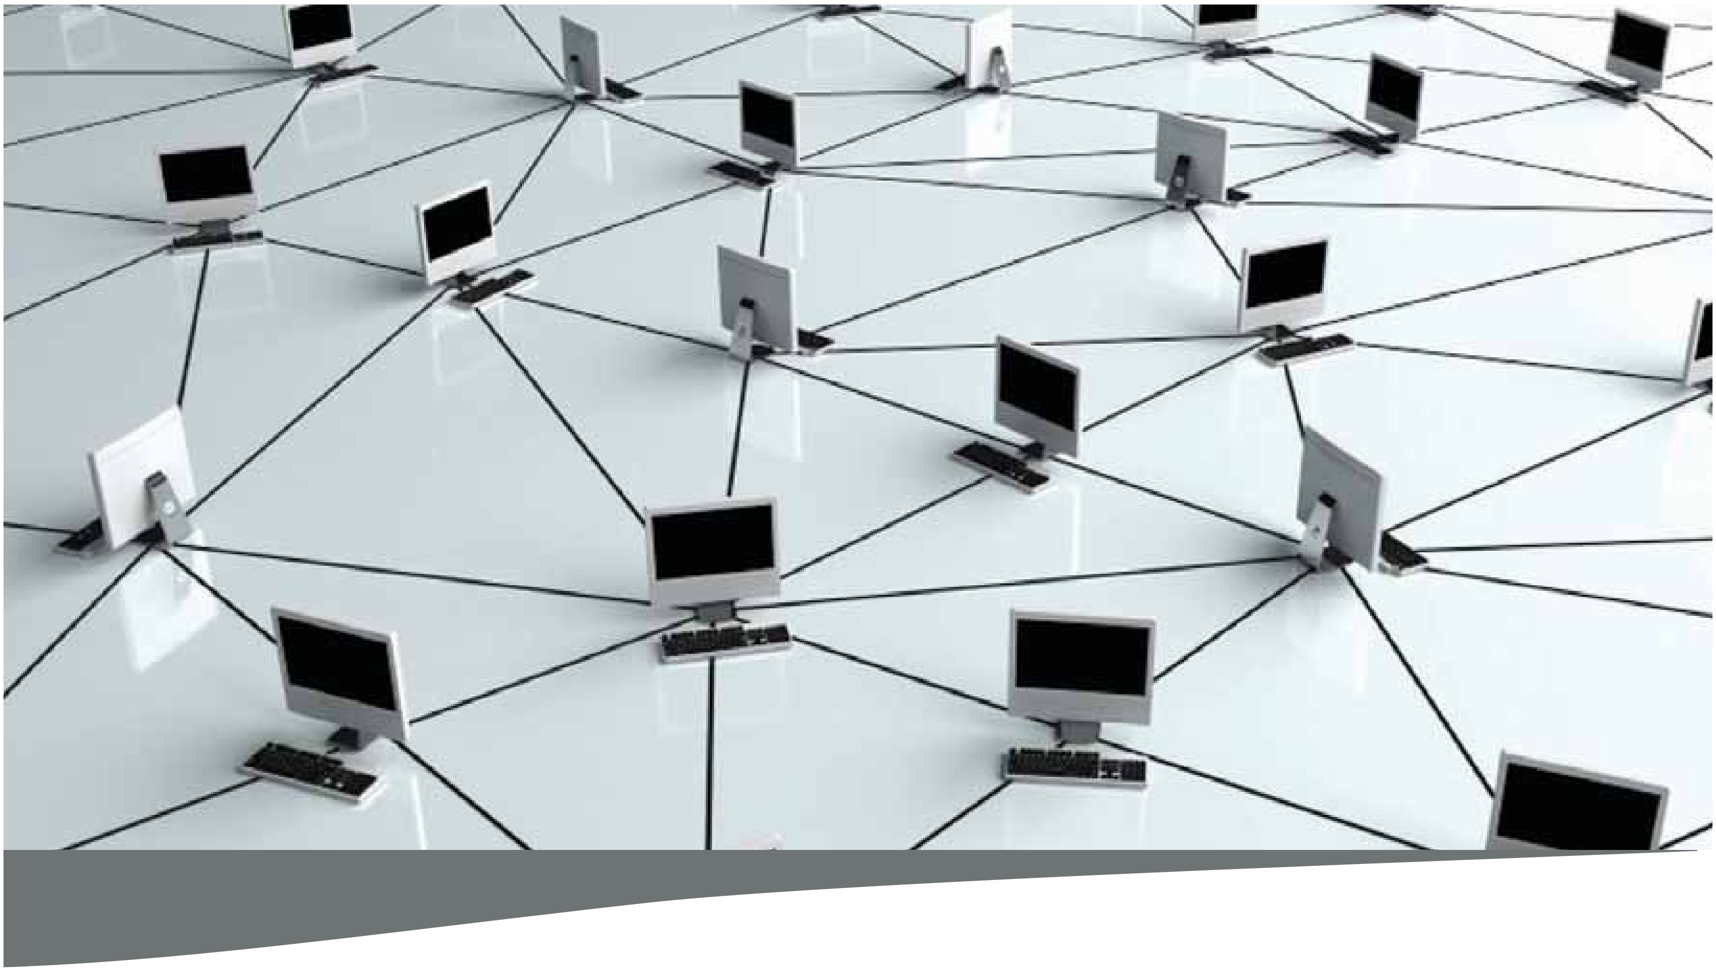

10GBASE-T

10GBase-T for Broad 10Gigabit Adoption in the Data Center

Contributors:
Carl G. Hansen, Intel
Carrie Higbie, Siemon

10 Gigabit Ethernet: Drivers for Adoption

The growing use of virtualization in data centers to address the need to reduce IT costs has caused many administrators to take a serious look at 10Gb Ethernet (10GbE) as a way to reduce the complexities they face when using the existing 1Gb Ethernet (1GbE) infrastructures. The server consolidation associated with virtualization has had significant impact on network I/O because they combine the network needs of several physical machines and the other background services, such as live migration, over the Ethernet network onto a single machine.

Together with trends such as unified networking, the ability to use a single Ethernet network for both data and storage traffic, are increasing I/O demands to the point where a 1GbE network can be a bottleneck and a source of complexity in the data center. The move to implement unified networking requires rethinking of data center networks. While 1GbE connections might be able to handle the bandwidth requirements of a single traffic type, they do not have adequate bandwidth for multiple traffic types during peak periods. This creates a need for multiple 1GbE connections.

Moving to 10 Gigabit Ethernet (10GbE) addresses these network problems by providing more bandwidth and simplifies the network infrastructure by consolidating multiple gigabit ports into a single 10 gigabit connection. Data Center Administrators have a number of 10GbE interfaces to choose from including CX4, SFP+ Fiber, SFP+ Direct Attach Copper (DAC), and 10GBASE-T. Today, most are choosing either 10GbE Optical or SFP+ DAC. However, limitations with each of these interfaces have kept them from being broadly deployed across the data center.

Fiber connections are not cost-effective for broad deployment, and SFP+ DAC is limited by its seven meter reach, and requires a complete infrastructure upgrade. CX4 is an older technology that does not meet high density requirements. For 10GBASE-T, the perception to date has been that it required too much power and was too costly for broad deployments. These concerns are being addressed with the latest manufacturing processes that are significantly reducing both the power and the cost of 10GBASE-T.

Widespread deployment requires a cost-effective solution that is backward compatible and has the flexibility capable of reaching the majority of switches and servers in the data center. This white paper looks at what is driving choices for deploying 10GbE and how 10GBASE-T will lead to broader deployment, including its integration into server motherboards. It also outlines the advantages of 10GBASE-T in the data center, including improved bandwidth, greater flexibility, and infrastructure simplification, ease of migration, and cost reduction.

The Need for 10 Gigabit Ethernet

A variety of technological advancements and trends are driving the increasing need for 10GbE in the data center. For instance, the widespread availability of multi-core processors and multi-socket platforms is boosting server performance. That performance allows customers to host more applications on a single server resulting in multiple applications competing for a finite number of I/O resources on the server. Customers are also using virtualization to consolidate multiple servers onto a single physical server, reducing their equipment and power costs. Servers using the latest Intel® Xeon® processors can support server consolidation ratios of up to fifteen to one .

However, server consolidation and virtualization have a significant impact on a server's network bandwidth requirements, as the I/O needs of several servers now need to be met by a single physical server's network resources. To match the increase in network I/O demand, IT has scaled their network by doubling, tripling, or even quadrupling the number of gigabit Ethernet connections per server. This model has led to increased networking complexity, as it requires additional Ethernet adapters, network cables and switch ports.

The transition to unified networking adds to the increasing demand for high bandwidth networking. IT departments are moving to unified networking to help simplify network infrastructure by converging LAN and SAN traffic, including iSCSI, NAS, and FCoE for a single Ethernet data center protocol. This convergence does simplify the network but significantly increases network I/O demand by enabling multiple traffic types to share a single Ethernet fabric.

Continuing down the GbE path is not sustainable, as the added complexity, power demands, and cost of additional GbE adapters will not allow customers to scale to meet current and future I/O demands. Simply put, scaling GbE to meet these demands significantly increases the cost and complexity of the network. Moving to 10GbE addresses the increased bandwidth needs while greatly simplifying the network and lowering power consumption by replacing multiple gigabit connections with a single or dual port 10GbE connection.

1 Source: Results have been estimated based on internal Intel analysis and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance.

Media Options for 10 Gigabit Ethernet

Despite industry consensus regarding the move to 10GbE, the broad deployment of 10GbE has been limited, due to a number of factors. Understanding this dynamic requires an examination at the pros and cons of current 10GbE media options.

10GbE Media Options

Technology	Pros	Cons	Cable
10GBASE-SR (SFP+ Fiber)	• 3.2w / port • Latency: < 1 µs • 300 meter reach (SR) • 10km reach (LR) • < $699 / port	• Requires new SFP+ switches • High cost of purchase and deployment • Not compatible with installed 1000BASE-T switches • Due to the higher purchase price ongoing maintenance costs are higher	• Multimode fiber
10GBASE-DAC (SFP+ Direct-Attach Copper)	• 2.9w / port • Latency: < 1 µs • < $399 / port	• Limited to 7 meters • Not compatible with existing 1000BASE-T switches • Requires new SFP+ switches • Can lead to a significant oversubscription of ports due to limited reach	• Twinax copper
10GBASE-CX4	• 2.25w / port	• Limited to 15 meters • Large connector form factor limits high density • Requires 10GBASE-CX4 switches, with few new products available • Not compatible with existing 1000BASE-T • Larger diameter cabling is difficult to manage in legacy cabinets	• Twinax copper
10GBASE-T	• <$399 / port • 9.2w / port (<80m) • 10.1w / port (>80m) • CAT6 55m • CAT6a 100m • CAT7 100m • Backward compatibility with 1000BASE-T • Uses existing high performance structured cabling	• Higher power than SFP+ • Latency: 2 - 4 µs	• UTP copper • FTP copper

The challenge IT managers face with 10GbE currently is that each of the current options has a downside, whether in terms of cost, power consumption, or reach.

10GBASE-CX4

10GBASE-CX4 was an early favorite for 10GbE deployments, however its adoption was limited by the bulky and expensive cables, and its reach is limited to 15 meters. The size of the CX4 connector prohibited higher switch densities required for large scale deployment. Larger diameter cables are purchased in fixed lengths resulting in challenges to manage cable slack. Pathways and spaces may not be sufficient to handle the larger cables.

SFP+

SFP+'s support for both fiber optic cables and DAC make it a better (more flexible) solution than CX4. SFP+ is ramping today, but has limitations that will prevent this media from moving to every server.

10GBASE-SR (SFP+ Fiber)

Fiber is great for latency and distance (up to 300 meters), but it is expensive. Fiber offers low power consumption, but the cost of laying fiber networking everywhere in the data center is prohibitive due largely to the cost of the electronics. The fiber electronics can be 4-5 times more expensive than their copper counterparts meaning that ongoing active maintenance, typically based on original equipment purchase price, is also more expensive. Where a copper connection is readily available in a server, moving to fiber creates the need to purchase not only the fiber switch port, but also a fiber NIC for the server.

10GBASE-SFP+ DAC

DAC is a lower cost alternative to fiber, but it can only reach 7 meters and it is not backward compatible with existing GbE switches. DAC requires the purchase of an adapter card and requires a new top of rack (ToR) switch topology. The cables are much more expensive than structured copper channels, and cannot be field terminated. This makes DAC a more expensive alternative to 10GBASE-T. The adoption of DAC for LOM will be low since it does not have the flexibility and reach of BASE-T.

10GBASE-T

10GBASE-T offers the most flexibility, is the lowest cost media type, and is backward compatible with existing 1GbE networks.

REACH

Like all BASE-T implementations, 10GBASE-T works for lengths up to 100 meters, giving IT managers a far-greater level of flexibility in connecting devices in the data center. With flexibility in reach, 10GBASE-T can accommodate either top of the rack, middle of row, or end of the row network topologies. This gives IT managers the most flexibility in server placement since it will work with existing structured cabling systems.

For higher grade cabling plants (category 6A and above) 10GBASE-T operates in low power mode (also known as data center mode) on channels under 30m. This equates to a further power savings per port over the longer 100m mode. Data centers can create any-to-all patching zones to assure less than 30m channels to realize this savings.

Backward Compatibility

Because 10GBASE-T is backward-compatible with 1000BASE-T, it can be deployed in existing 1GbE switch infrastructures in data centers that are cabled with CAT6, CAT6A or above cabling, allowing IT to keep costs down while offering an easy migration path to 10GbE.

Power

The challenge with 10GBASE-T is that the early physical layer interface chips (PHYs) have consumed too much power for widespread adoption. The same was true when gigabit Ethernet products were released. The original gigabit chips were roughly 6.5 Watts/ port. With process improvements, chips improved from one generation to the next. The resulting GbE ports are now under 1W / port. The same has proven true for 10GBASE-T. The good news with 10GBASE-T is that these PHYs benefit greatly from the latest manufacturing processes. PHYs are Moore's Law friendly, and the newer process technologies will continue to reduce both the power and cost of the latest 10GBASE-T PHYs.

When 10GBASE-T adapters were first introduced in 2008, they required 25w of power for a single port. Power has been reduced in successive generations of using newer and smaller process technologies. The latest 10GBASE-T adapters require only 10w per port. Further improvements will reduce power even more. By 2011, power will drop below 5 watts per port making 10GBASE-T suitable for motherboard integration and high density switches.

Latency

Depending on packet size, latency for 1000BASE-T ranges from sub-microsecond to over 12 microseconds. 10GBASE-T ranges from just over 2 microseconds to less than 4 microseconds, a much narrower latency range. For Ethernet packet sizes of 512B or larger, 10GBASE-T's overall throughout offers an advantage over 1000BASE-T. Latency for 10GBASE-T is more than 3 times lower than 1000BASE-T at larger packet sizes. Only the most latent sensitive applications such as HPC or high frequency trading systems would notice any latency.

The incremental 2 microsecond latency of 10GBASE-T is of no consequence to most users. For the large majority of enterprise applications that have been operating for years with 1000BASE-T latency, 10GBASE-T latency only gets better. Many LAN products purposely add small amounts of latency to reduce power consumption or CPU overhead. A common LAN feature is interrupt moderation. Enabled by default, this feature typically adds ~100 microseconds of latency in order to allow interrupts to be coalesced and greatly reduce the CPU burden. For many users this trade-off provides an overall positive benefit.

Cost

As power metrics have dropped dramatically over the last three generations, cost has followed a similar downward curve. First-generation 10GBASE-T adapters cost $1000 per port. Today's third-generation dual-port 10GBASE-T adapters are less than $400 per port. In 2011, 10GBASE-T will be designed as LAN on Motherboard (LOM) and will be included in the price of the server. By utilizing the new resident 10GBASE-T LOM modules, users will see a significant savings over the purchase price of more expensive SFP+ DAC and fiber optic adapters and will be able to free up and I/O slot in the server.

Data Center Network Architecture Options for 10 Gigabit Ethernet

The chart below lists the typical data center network architectures applicable to the various 10GbE technologies. The table clearly shows 10GBASE-T technology provides greater design flexibility than its two copper counter-parts.

Technology	Data Center Network Architectures	Connectivity
10GBASE-SR SFP+ Fiber	• Top of Rack (ToR)	Uplinks from ToR switches to aggregation layer switches
	• Middle of Row (MoR)	Inter-cabinet connectivity from servers to MoR switches
	• End of Row (EoR)	Inter-cabinet connectivity from servers to EoR switches
	• Core network	Backbone
10GBASE-SFP+ DAC	• Top of Rack	Intra-cabinet connectivity from servers to ToR switches
10GBASE-CX4	• Top of Rack	Intra-cabinet connectivity from servers to ToR switches
10GBASE-T	• Top of Rack (ToR)	Intra-cabinet connectivity from servers to ToR switches
	• Middle of Row (MoR)	Inter-cabinet connectivity from servers to MoR switches
	• End of Row (EoR)	Inter-cabinet connectivity from servers to EoR switches

THE FUTURE OF 10GBASE-T

TIntel sees broad deployment of 10GbE in the form of 10GBASE-T. In 2010 fiber represents 44% of the 10GbE physical media in data centers, but this percentage with continue to drop to approximately 12% by 2013. Direct-attach connections will grow over the next few years to 44% by 2013 with large deployments in IP Data Centers and for High Performance Computing. 10GBASE-T will grow from only 4% of physical media in 2010 to 44% in 2013 and eventually becoming the predominate media choice

10GBASE-T as LOM

Sever OEMs will standardize on BASE-T as the media of choice for broadly deploying 10GbE for rack and tower servers. 10GBASE-T provides the most flexibility in performance and reach. OEMs can create a single motherboard design to support GbE, 10GbE, and any distance up to 100 meters. 1GBASE-T is the incumbent in the vast majority of data centers today, and 10GBASE-T is the natural next step.

Conclusion

Broad deployment on 10GBASE-T will simplify data center infrastructures, making it easier to manage server connectivity while delivering the bandwidth needed for heavily virtualized servers and I/O-intensive applications. As volumes rise, prices will continue to fall, and new silicon processes have lowered power and thermal values. These advances make 10GBASE-T suitable for integration on server motherboards. This level of integration, known as LAN on Motherboard (LOM) will lead to mainstream adoption of 10GbE for all server types in the data center.

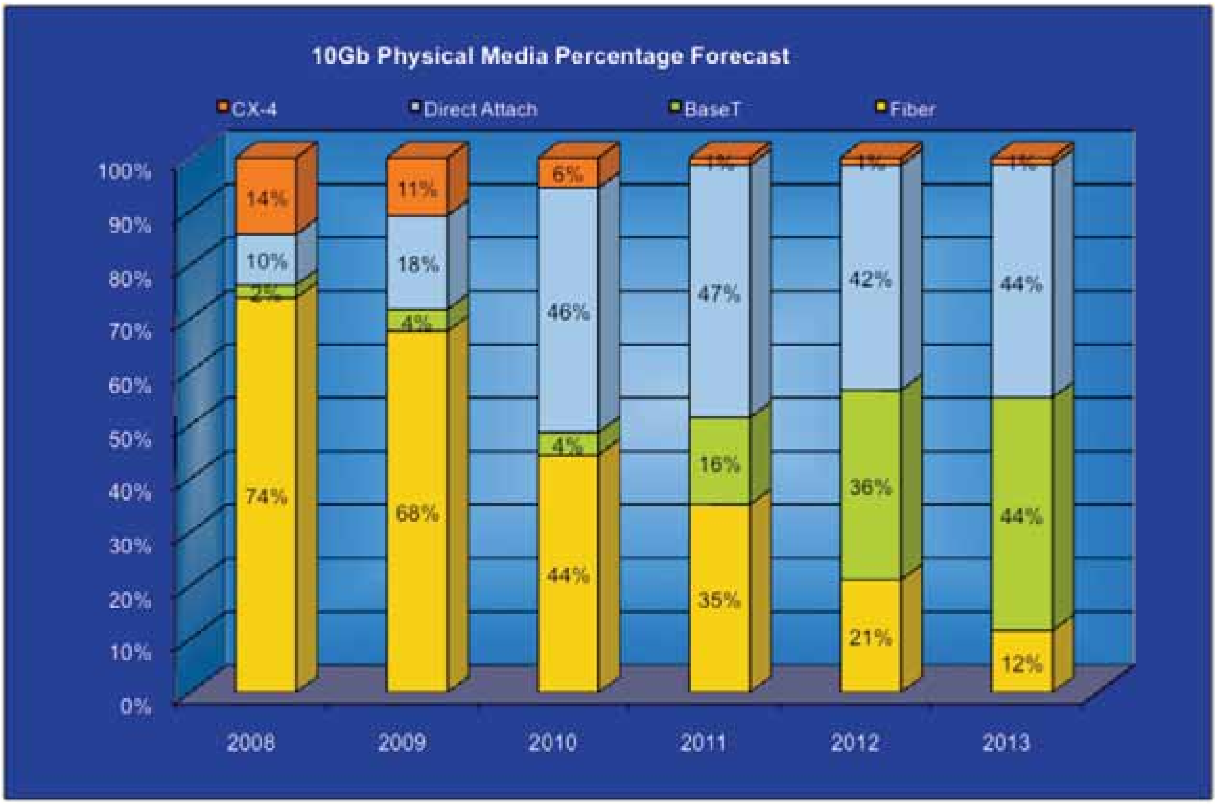

Source: Intel Market Forecast

SIEMON

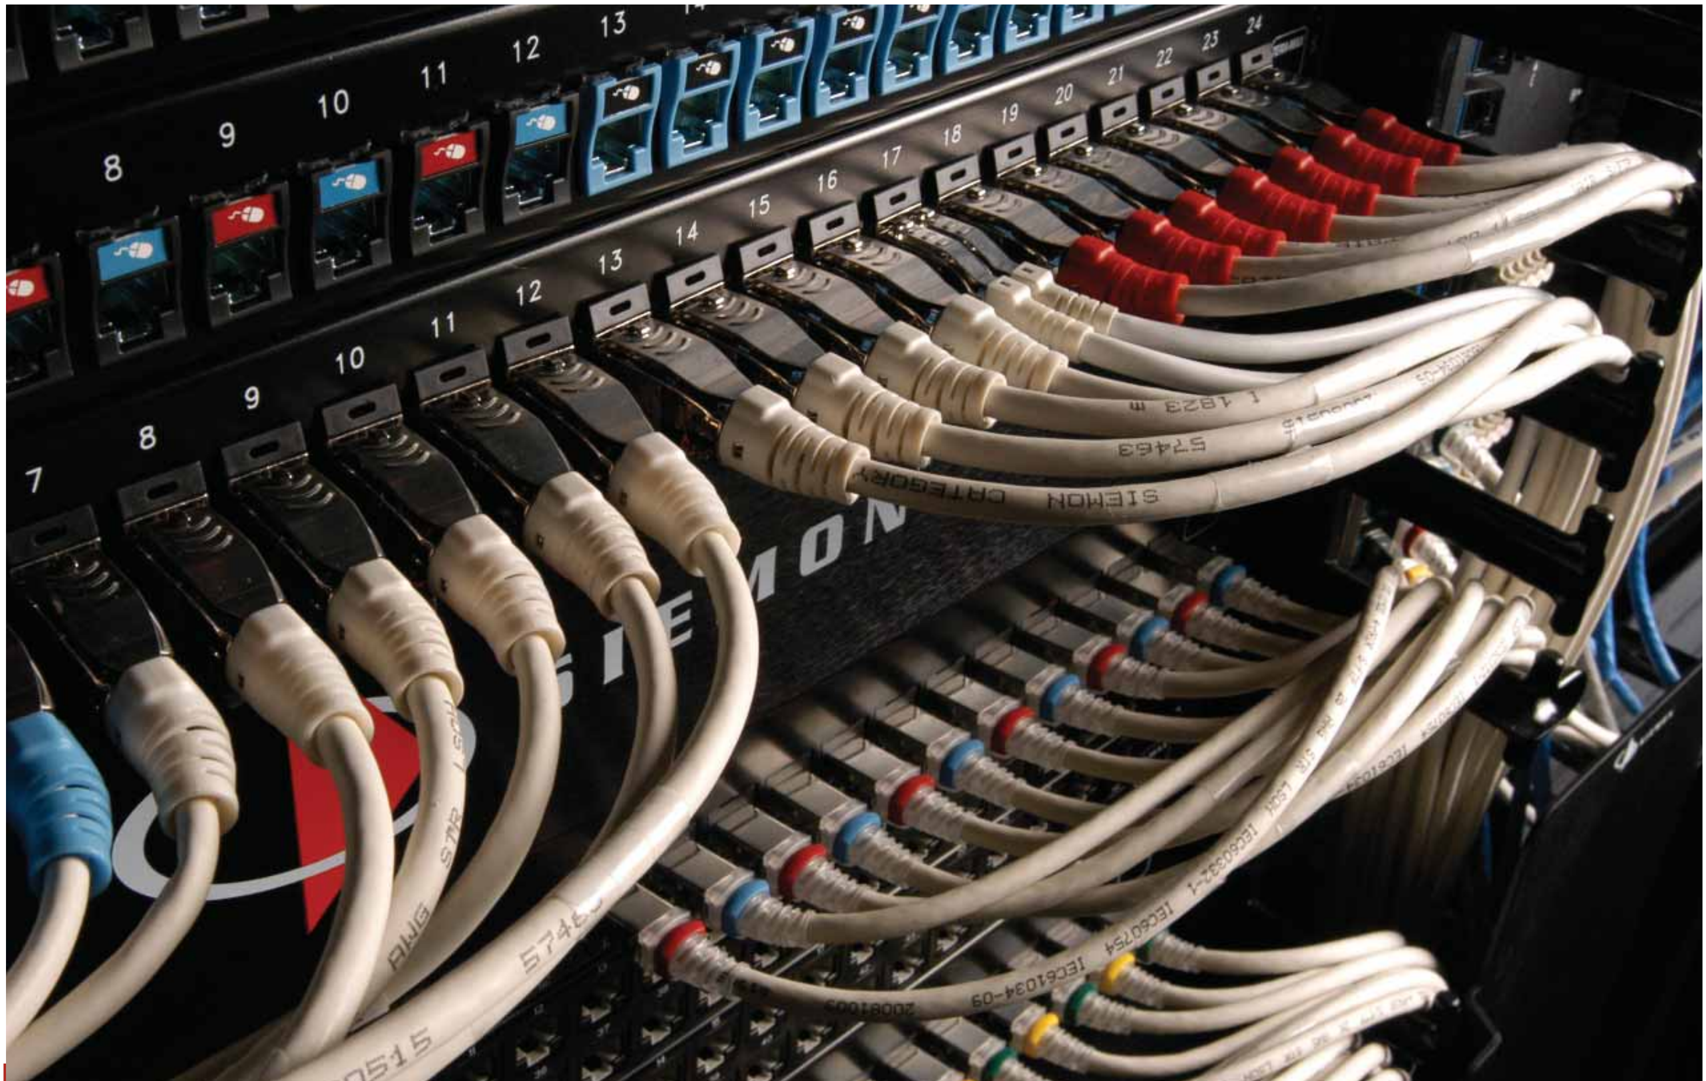

Hosted, Outsourced, and Cloud Data Centers -

Considerations for Overall SLA's for Facility Owners and Hosting Providers

Hosted and Outsourced Facility Definitions

Hosted data centers, both outsourced/managed and co-location varieties, provide a unique benefit for some customers through capital savings, employee savings and in some cases an extension of in-house expertise. Traditionally, these facilities were thought of as more SME (Small to Medium Enterprise) customers. However, many Global 500 companies - have primary, secondary or ancillary data centers in outsourced locations. Likewise, co-location data centers are becoming increasingly popular for application hosting such as web hosting and SaaS (Software as a Service), Infrastructure as a Service (IaaS), Platform as a Service (PaaS) in Cloud computing. These models allow multiple customers to share redundant telecommunications services and facilities while their equipment is colocated in a space provided by their service provider. In-house bandwidth may be freed up at a company's primary site for other corporate applications.

Hosted and outsourced/managed data centers are growing rapidly for both companies' primary and hot site (failover ready) data centers, redundant sites and for small to medium enterprises. Similarly, outsourced data center services are on the rise and allow a company to outsource data center operations and locations, saving large capital requirements for items like generators, UPS/Power conditioning systems and air handling units. As data center services increase, many providers can supply one or all of these models depending on a tenants needs. The various combinations of hosted/co-location and cloud services available from hosting providers are blending terms and services.

Considerations for the Hosted/Cloud Facilities Owner

The challenges for a hosted or cloud facility owner are similar to the user considerations mentioned above, but for different reasons. While most facilities are built with the expectation of full occupancy, the reconfiguration of occupancy due to attrition and customer changes can present the owner with unique challenges. The dynamic nature of a tenant-based data center exacerbates problems such as cable abatement (removal of abandoned cable), increasing power demand and cooling issues.

Data centers that have been in operation for several years have seen power bills increase and cooling needs change - all under fixed contract pricing with their end-user, tenant customers. The dynamic nature of the raised floor area from one tenant to the next compounds issues. Some co-location owners signed fixed long-term contracts and find themselves trying to recoup revenue shortfalls from one cage by adjusting new tenant contracts. Renegotiating contracts carries some risk and may lead to termination of a long-term contract.

Contracts that are based on power per square foot plus a per square foot lease fee are the least effective if the power number is based on average wattage and the contract does not have inflationary clauses to cover rising electricity costs. Power usage metering can be written into contracts, however in some areas this requires special permission from either the power company or governing regulatory committees as it may be deemed as reselling power. As environmental considerations gain momentum, additional focus is being placed on data centers that use alternative energy sources such as wind and solar.

There are however, additional sources of revenue for owners that have traditionally been overlooked. These include packets passed, credits for power saving measures within tenant cages, lease of physical cabinets and cabling (both of which can be reused from one tenant to the next) and monitoring of physical cabling changes for compliance and/or security along with traditional network monitoring.

For new spaces, a co-location owner can greatly mitigate issues over time with proper space planning. By having at least one area of preconfigured cages (cabinets and preinstalled cabling), the dynamic nature in that area and the resulting problems are diminished. This allows a center to better control airflow. Cabling can be leased as part of the area along with the cabinets, switch ports, etc. This allows the cabinets to be move-in ready for quicker occupancy. This rapidly deployed tenancy area will provide increased revenue as the space does not need to be reconfigured for each new tenant. This area can also be used by more transient short term tenants that need space while their new data center or redundant site is built.

If factory terminated and tested trunking cable assemblies aren't used, it is important to use quality cabling so that the cable plant does not impact Service Level Agreements (SLAs). Both TIA 942 and ISO 24764 recommend a minimum of category 6A/Class EA cabling. The minimum grade of fiber is OM3 for multimode. Singlemode is also acceptable for longer distances and may be used for shorter distances, although the singlemode electronics will be higher priced.

Owners must insist on quality installation companies if they allow tenants to manage their own cabling work. An owner may want to maintain a list of approved or certified installers. One bad installer in one cage can compromise other users throughout the facility. Approved installers provide the owner with an additional control over pathways and spaces. Further, owners want to insist on high performing standards-based and fully tested structured cabling systems within the backbone networks and cages. Higher performing systems can provide a technical and marketing advantage over other owners that

While co-location owners historically stop their services at the backbone, distributed switching via a centralized cabling plant and patching area can provide significant power savings through lower switch counts, enhanced pathway control and decreased risk of downtime during reconfigurations. All the while, the additional network distribution services provide increased revenue for the co-location owner. Managed and leased cabling ports can be an additional revenue stream.

Understanding that some tenants will have specific requirements, a combination of preconfigured and non-preconfigured cages may be required. For more dynamic non-preconfigured areas, trunking assemblies, which are factory terminated and tested, allow the owner to offer various cabling performance options, such as category 6 or 6A UTP, 6A shielded or category 7A fully shielded, to best suit the end-user's needs. The owner can lease these high performing cabling channels and, on the greener side, the cabling can be reused from one tenant to the next, eliminating on site waste and promoting recycling.

Whether pre-cabled or cabled upon move in, owner leased or customer installed, category 6A or higher copper and OM3/OM4 fiber or better should be used. Higher performing cabling conforms to the minimum recommended standards, allows for higher speed applications while providing backwards compatibility to lower speed technologies. Category 6A/Class EA, 7/Class F and 7A/Class FA allow short reach (lower power mode) for 10GBASE-T communications under 30m for an additional power savings to the owner. Category 7/7A and class F/FA also provides the most noise immunity and meets strict government TEMPEST/EMSEC emissions tests, meaning they are suitable for use in highly classified networks alongside fiber. Installing the highest performing cabling up front will result in longer cabling lifecycles thus reducing the total cost of ownership and maximizing return on investment.

For non-configured areas, the backbone can be distributed into zones. The zone distribution area can be connected to pods or modular units within a space. This modular approach allows customers to move equipment into their areas one pod at a time. Preterminated copper and fiber trunking cables are easily configured to known lengths allowing the location owner to have stock on hand for rapid deployment of customer areas. These trunks can be reused and leased from tenant to tenant increasing revenue and enabling near instant occupation.

Facility owners are typically under some type of SLA requirements. SLA's can be for performance, uptime, and services. There are some network errors that are caused by poorly performing or underperforming cabling plants. Selecting high performing quality cabling solutions is only partial protection. The quality of the installation company is key for pathways, spaces, performance and error free operation. Cabling has historically been an afterthought or deemed to be the tenant's decision. By taking control of the cabling in hosted spaces, the building owner removes the cabling issues that can cause SLA violations, pathway problems, and ensure proper recycling of obsolete cabling.

While network monitoring can pinpoint ports that cause bit errors and retransmission, determining if the cause is cabling related can be difficult. Noise is harder to troubleshoot as it is intermittent. Testing the cable requires that a circuit is down for the period of testing, but may be necessary when SLAs are in dispute. While intermittent retransmissions are relatively benign in normal data retrieval, poorly performing cabling can make this intermittent issue more constant. This can slow down transmissions, or in the case of voice and video, can become audible and visible. In short, cabling is roughly 3-5% of the overall network spend, but that 3-5% can keep the remaining 95-97% from functioning properly and efficiently.

Modularized Deployment for the Co-location/Hosted Facilities Owner

Hosted and co-location facilities lend themselves well to modular POD-type scalable build outs. It is rare that these centers are built with full occupancy on day one unless there is a sizeable anchor tenant/tenants. Spatial planning for tenant considerations can sometimes be problematic due to varied size, power and rack space required by customers. These facilities are generally an open floor plan to start. Configuring spaces in a cookie cutter manner allows the owner to divide space in parcels while addressing hot/cold aisle requirements, cabling, and most importantly scalability and versatility within the floor plan space. In a typical scenario, the space is allocated based on cage layouts. The rows can be further subdivided for smaller tenants, or cage walls can be removed for larger

Cloud facilities are generally highly occupied day one. A modularized design approach in these environments allows rows of cabinets to be deployed in a cookie cutter fashion. A structured cabling system that is pre-configured within cabinets, or ready for connection to banks of cabinets allows the owner to have a highly agile design that accommodates a wide variety of equipment changes without the need to run additional cabling channels in the future. There are two ways to deploy a modularized cloud or co-location data center. The first entails pre-cabling cabinets and rows to a centralized patching area. The second involves pre-cabling to zones within the data center. Once the zones are cabled, the addition of rows of cabinets within the zone becomes a matter of moving in the new populated cabinets, and connecting them via patch cords to the zone cabling distribution area. One common complaint with high density centers, such as clouds, is that equipment is often moved in with little to no notice. By pre-cabling the data center to a centralized patching area or to zones, the reactionary and often expensive last minute rush is eliminated.

If a centralized patching area is used, equipment changes become a patch cord or fiber jumper change, allowing rapid deployment. In a central patching (any to all) configuration, copper and/or fiber patch panels are provided in the central patching area that corresponds to patch panels in each cabinet. Connections to switching, servers, SAN, etc., are achieved via patch cords rather than having to run new channels as new cabinets are deployed.

The Need for Space Planning

One historical problem in open non-configured spaces has been the varied customer configuration requirements and the need to fit as many customers into the floor space as possible. As with any data center, growth without planning can cause serious issues in a co-location/shared space. One cage's equipment running perpendicular to another cage can cause undesirable hot air to be introduced into cold aisle of adjacent spaces. Haphazard and inconsistent cabling practices can block air flow. Improper use of perforated tiles can cause loss of static pressure at the far sides of the space. In short, in a hosted space that is not properly planned, problems can arise quickly.

For space planning, an owner typically defines zones within the open space. Due to deeper equipment, a minimum of 3 feet (800 mm) should be allowed in all aisles, or slider cage doors should be installed that will provide full access. If that is not possible, deeper equipment should be housed in the cabinets in front of the sliding doors so that cage walls don't block access. A facility owned and operated cage can provide facility wide networking, monitoring and connectivity services to other cages via preconfigured, pre-cabled, cabinets allowing servers to be moved in and plugged in on demand. The cabinets and networking services become part of the tenant lease.

To allow for a variety of customer size requirements, a set of caged areas can be provided with 2-4 preconfigured cabinets for smaller tenants. By preplanning the spaces, cages do not need to move, pathways and spaces are predefined and airflow can be optimized in hot/cold aisles. In reality, there may be tenants that move into one of these areas that do not need to fill the cabinets provided. Some facilities allow for subleasing within cages. This allows underutilized cabinets to be occupied by another tenant as long as access to the area is supervised and cabinets have segmented security access via different combinations and/or key locks. Even in a tenant designed space it is common for a cabinet or partial cabinet to go unused. The benefit over time in pre-configured areas is that the floor will remain unchanged from one tenant to the next.

Another area with 8-10 cabinets is preconfigured for medium size tenants. And another section/area is left blank for those tenants that require their own configuration. The layout of that area should be completed by the building owner to assure that hot aisle/cold aisle planning is consistent throughout the floor area.

Figure 1 – Sample space plan

In the sample space plan above, we see caged areas of various sizes. Cage walls are static, cabling is centralized, and air flow is optimized. By providing varied numbers of cabinets within each cage, the floor plan can accommodate a variety of tenants. Tenants can occupy one or more cages depending on needs. For smaller tenants, individual cabinets or smaller spaces can be leased providing room for growth. The static cage configuration provides a significant cost savings over time. Centralized patching may be provided for the entire floor or in each zone with connections to core services. This keeps cable lengths shorter, less expensive, and easier to manage..

The above plan takes advantage of Siemon's VersaPOD cabinet line. The VersaPOD is available with a variety of integrated Zero U vertical patch panels (VPP) for support of copper and fiber patching. The VPP's supply up to 12U of patching and cable management in the front and/or rear vertical space between two bayed cabinets without consuming critical horizontal mounting space. By utilizing the vertical space adjacent to the vertical mounting rails, the VPP's provides ideal patching proximity to active equipment, minimizing patch cord runs and slack congestion. Zero-U vertical patching areas can also be used to mount PDU's to service the equipment mounted in the adjacent 45 U of horizontal mounting space. This increases versatility and eliminates cabling obstructions and swing arms within equipment areas which can block air flow from the equipment. The Zero-U patching and cable management channels further free up horizontal rack mount space and provides better managed and controlled pathways.

The highly perforated (71%) doors allow greater airflow into equipment whether it be from an underfloor system or if cooling is supplemented by an in row cooling unit. To increase heat egress, optional fans can be installed in the top of the cabinets.

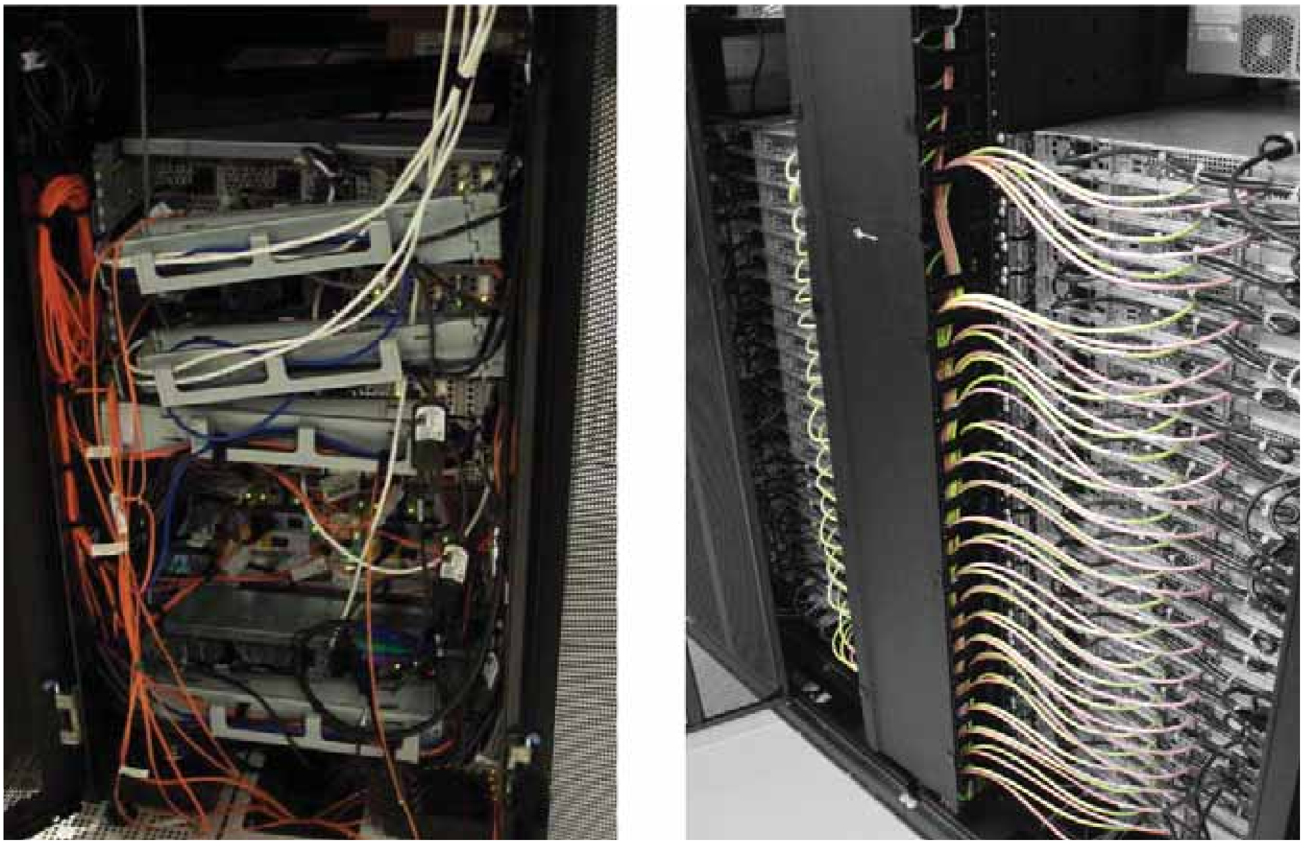

Figure 2 - Swing-arm cable managers issues vs. VersaPOD Zero-U vertical patching channels

Cabinets in all areas should be outfitted with blanking panels that can be removed/moved as equipment is installed. An overall cooling plan must include intra-cage support. Blanking panels can have a significant impact on cooling expenses. Likewise, brush guards where cabling penetrations pass through floor tiles can help to maintain static pressure under the raised floor.

IIM (Intelligent Infrastructure Management)

By using a central patching area or zone patching areas, Intelligent Infrastructure Management can be deployed in a very cost effective manner. It is understood that the equipment that moves in and out of cabinets will vary over time regardless if there is one continuous tenant or several changing tenants.

The connections in the central patching area are monitored dynamically and in real time by analyzers that monitor continuity via a 9th pin on the patch cords and fiber jumpers. Because the software can see the equipment at the end of each channel via SNMP, it really doesn't matter what that the equipment is or if it changes.

Using Cross Connections in a Central patching area eliminates the need for sensor strips that attach to active equipment in each cabinet. Without a cross connect, sensor strips must be replaced as equipment changes either due to failure, upgrade, replacement or new deployment. As new equipment is introduced into the market, there may be a void in time between equipment deployment and the corresponding sensor strip being available.

With IIM, moves, adds and changes are logged for date and time (necessary for most compliance requirements), and can be accompanied by photographs of the person making the change if the central patching area/zone is outfitted with either cameras or video equipment. For companies that have requirements for HIPAA, Sox, CFR-11, and other data

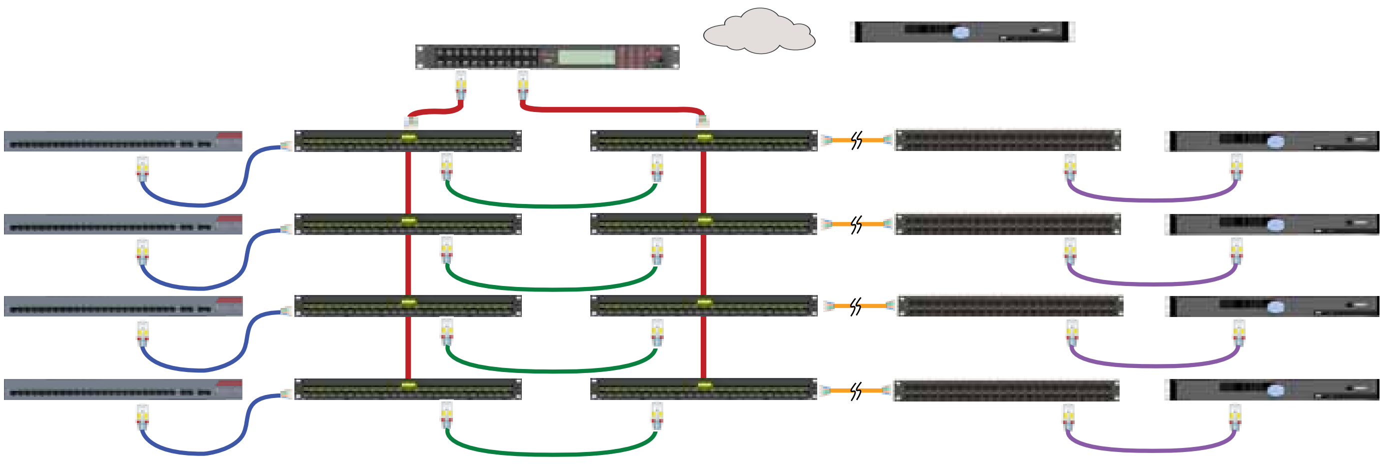

Figure 3 - IIM in cross-connect configuration

protection laws, this audit trail maintains networking documentation.

For the facility owner, this software will also allow visibility into switch ports that are patched but not passing traffic. This enables better asset/port utilization reducing the need to add equipment and the resulting additional power consumption. Because the cabling channel is added to overall troubleshooting, it becomes much easier to identify and locate equipment for repair. The faster reaction times for troubleshooting can increase SLA performance while providing necessary audit trails. A premium may also be charged for Intelligent Infrastructure monitoring.

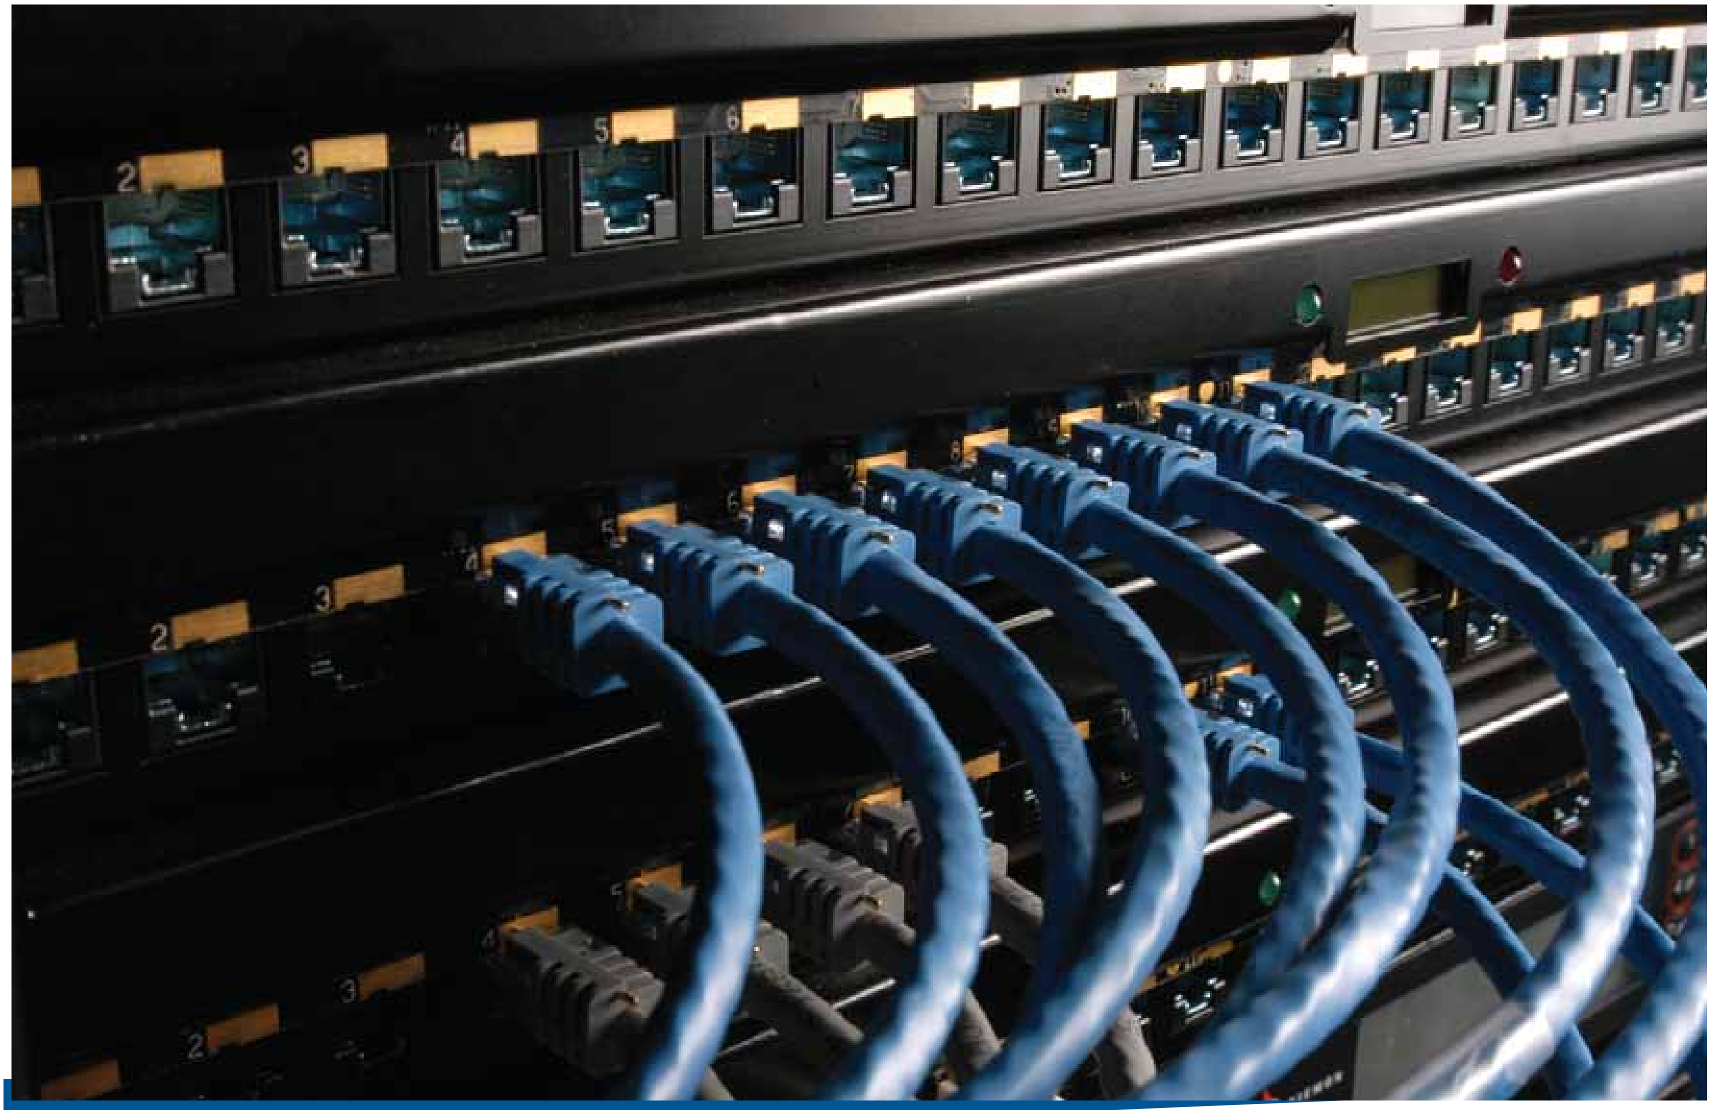

Hosted, Outsourced, and Cloud Data Centers -

Strategies and Considerations for Co-Location Tenants

Hosted and Outsourced Facility Definitions

Hosted data centers, both outsourced/managed and co-location varieties, provide a unique benefit for some customers through capital savings, employee savings and in some cases an extension of in-house expertise. Traditionally, these facilities have been thought of as more SME (Small to Medium Enterprise) customers. However, many Global 500 companies have primary, secondary or ancillary data centers in outsourced locations. Likewise, co-location data centers are becoming increasingly popular for application hosting such as web hosting and SaaS (Software as a Service), Infrastructure as a Service (IaaS), Platform as a Service (PaaS) in Cloud computing. These models allow multiple customers to share redundant telecommunications services and facilities while their equipment is colocated in a space provided by their service provider.
In house bandwidth may be freed up at a company's primary site for other corporate applications.

Hosted and outsourced/managed data centers are growing rapidly for both companies' primary and hot site (failover ready) data centers, redundant sites and for small to medium enterprises. Similarly outsourced data center services are on the rise and allow a company to outsource data center operations, locations, saving large capital requirements for items like generators and UPS/Power conditioning systems and air handling units. As data center services increase, many providers can supply one or all of these models depending on a tenants needs.

Outsourced Data Centers

In an outsourced data center, the tenant basically rents some combination of space, talent and facilities from a larger facility provider for all or part of their corporate applications and data center operations. There are several pricing options including per port, per square foot, and for power consumed, but in general a combination thereof. With power costs and demand on the rise, most newer contracts include a fee that is assessed when a tenant's kilowatt threshold is exceeded, or by power supplied. In the latter case, a tenant typically pays for more power than they need as power is averaged across the square footage of the tenant space.

Outsourced data centers are an attractive option for companies that have a myriad of platforms and applications alleviating the need for constant multivendor training and upgrades, patches, hardware changes, software platform changes, etc. In a typical company environment that has migrated from mainframe type applications to several server platforms just the cost and time for training can be a manpower and financial drain. As outsourced (managed) data centers have the needed expertise on site.
A company utilizing this type of model will see a shift in employee responsibilities from IT/upgrade tasks to more fruitful and beneficial tasks. Outsourced data centers may be for a sole tenant or multi-tenant, and in the case of the latter will share the same concerns as the co-location facilities below.

Co-location Facilities

Co-location facilities are typically divided into cages, cabinet space or in some cases, subdivided cabinets to accommodate smaller computing needs. As a co-location owner, division of space is a prime consideration. While these environments tend to be fluid, critical infrastructures (cabling, cages, power and cooling) that can remain unchanged provide advantages to the owner and tenants alike. There are very few existing outsourced locations that have not felt some pain over time as tenants move in and out leaving cabling messes in pathways that can be detrimental to air flow and cooling. Likewise, changing cabinet locations affects airflow directions, and equipment power loads can create hotspots and have adverse affects from one cage to another. Moving cage walls can render some spaces unusable. Reconfiguration of each space from tenant to tenant can be costly over time.

In a hosted only data center, a tenant leases square feet/meters of space and services including security, facilities (power and cooling), telecommunications and backup systems such as UPS's and generators. In a hosted space, a tenant generally uses their own resources for equipment maintenance, patch management, infrastructure, etc. Co-location scenarios can be an attractive option for redundant hot (instant failover) or cold (manual failover) spare sites, in the interim during a consolidation or new build, when primary data center site space has reached capacity, or when resources such as power, cooling, and space are at capacity. Similarly, if major upgrades are going to occur at a main end-user site (i.e. new chillers, reconfigured or new space) a temporary hosted or outsourced site may provide a solution. The dividing lines between co-location and hosted sites are becoming increasingly blurred as operators are beginning to offer blended services based on customer needs.

While some companies choose to build operate and maintain their own data centers, there is a large segment of companies that either wholly or partially take advantage of hosted/outsourced facilities. Global companies may choose to house a main facility and perhaps it's redundant counterpart in their own buildings. However as operations grow or new countries are added to the company's portfolio, a hosted/managed facility may serve well on an interim basis until it is cost justified to add another data center of their own. Small to medium enterprises which have a harder time attracting and keeping talented IT staff can in some cases, have a much better data center and support by utilizing already trained talent

Cloud Facilities

Cloud computing is a new buzzword that is all encompassing, and can be either IaaS, SaaS, PaaS, or a combination thereof. In most cloud scenarios, an end user is renting space, bandwidth or computing power on an on demand, as needed basis. Each cloud provider has a set of tools that allow them to interface with the hardware installed within their site. Some of their software is proprietary, and there are still some security concerns, but as these facilities and their applications mature, they can offer valuable resources to companies.

Cloud provider offerings may be in co-location facilities, managed facilities, or housed in provider owned facilities. Clouds can also reside in private corporate data centers or as a hybrid combination of public (in a cloud facility) and private (company owned). Clouds can be thought of as clusters of services that are not location dependant to provide processing, storage and/or a combination of these offerings.

An example of cloud computing is Amazon's EC2 (Elastic Compute Cloud) platform. This service allows rapid provisioning of computing and storage needs on demand. For instance, if a customer needs to provision a new server, the server is already there in one of Amazon's facilities. The customer does not need to justify, purchase, configure, power and maintain the server. If a customer only needs the server for a short period of time, it can be commissioned and decommissioned on demand for temporary computing needs. One primary advantage of public cloud computing is that when temporary cloud resources are no longer needed, the bill goes to zero. Public cloud resources are billed on a per use, as needed basis. This allows companies to have burstable resources without having to build networks that support peak loads, but rather build to support baseline or average loads. Public and private clouds allow applications to burst into the cloud when needed and return to normal when peak loads are no longer required.

If a customer is looking at any of the above solutions, Service Level Agreements (SLA's), reliability and confidence in security are the largest factors in the decision making process. It is not as easy to manage what you don't control. Further, end users must trust that the sites are well maintained so that good service doesn't turn into a loss of service over time.

Hosted Space Evaluation for Tenants

When evaluating outsourced space security is a prime consideration. Security should include biometrics, escorted access, after hours access, concrete barriers, and video surveillance at a minimum. Some spaces utilize cages to section off equipment with each tenant having the ability to access only their cage. However, should multiple tenants occupy the same floor; it may be possible to access another tenant's equipment either under the raised floor or over the top of the cage. This may make the space undesirable if personal/confidential information is stored on the servers housed within the cages. Escorted access for service personnel and company employees provides an extra level of assurance that data will remain uncompromised in these spaces.

VersaPOD Zero-U Vertical Patch Panel

Personnel working in adjacent spaces may also provide a risk to equipment and services where pathways cross caged environments. Intelligent Infrastructure Management solutions, such as Siemon's MapIT G2 system, provide real time monitoring of connections to critical equipment, an audit trail of moves, adds and changes, and an extra level of troubleshooting support. While these factors may not apply to all situations, certainly where critical and sensitive information is being stored this additional level can ease troubleshooting and provide assurances for the physical infrastructure. Intelligent infrastructure management can be implemented for either the hosted facility backbone operations, inside cages for customer connections, or both. Due to the real time physical connection monitoring, accidental or unauthorized disconnects can trigger alarms and escalations assuring that services are restored in a timely manner.

Maintenance of the facility and its equipment is also a factor. Determining how often generators are tested, UPS systems are tested and failover mechanisms are tested is critical. The same is true for fire suppression and detection systems. The data center service provider should be able to provide you with reports from cooling and PDU units and explain their processes and procedures for testing and auditing all systems as well as their disaster recovery plans. The power systems should have enough capacity to support all circuits and power demands including planned growth for the entire floor..

It is in a customer's and site's best interests to utilize power supplies that provide power usage monitoring, not just power output monitoring. Without usage monitoring, a tenant may be paying for more power than they use. Power utilization management also helps with provisioning. Power systems that are over capacity may not be able to provide enough power in the event of a failure when redundant equipment powers up. If a user is paying based on port connections and/or power utilization, a risk assessment should performed. This assures that equipment that does not require redundancy for critical business operations does not consume more power and network than necessary. As environmental considerations gain focus, additional importance is being placed on centers that use alternative energy sources such as wind and solar.

Ineffective cooling units may create not only cooling problems, but if not serviced regularly may cause excessive vibration or other harmful effects. It is important to ascertain how often the unit filters are changed, how failover happens, service schedules, etc.

Pathways and spaces within the data center should be properly managed. There should be a standard within the facility for cabling placed in air spaces or overhead. It is worth checking to see what cable management policies are practiced and enforced, not just written. Improperly placed copper and fiber, either overhead or under floor, and overfilled pathways can create air flow and cooling issues either in your area or adjacent cages over which you do not have control.

A tenant should be allowed to use their preferred cabling and installation company provided that the installation company adheres to center's pathway rules. If the space owner requires the use of their own installation company, you will want a listing of credentials and test results upon completion of the work. As some facility owners do not see cabling as critical to core services, installations may be done by the least expensive bidder using the least expensive components which may not provide high quality installation and/or sufficient performance margins which can create issues and finger pointing with SLAs. Copper and Fiber trunking assemblies are an excellent choice in these spaces as links are factory terminated and tested and can be reused should a tenant relocate. Trunking cables also offer an easy cabling upgrade path as they can be quickly removed and replaced with higher category trunking cable assemblies of the same length. For example, Siemon's Z-MAX Trunks are available in category 6 and category 6A shielded and unshielded and any of these assemblies can be used within the Z-MAX 24 or 48-port 1U shielded patch panels, allowing cabling to be upgraded without changing the patch panel.

It is important to ensure that enterprise and campus copper and fiber cabling systems outside of the data center are robust and certified to the specified category. Some Cloud providers are requiring customers to have their enterprise and campus cabling systems tested, certified and even upgraded to a higher performance category to eliminate the possibility that SLA problems are not caused out-

Future growth should also be considered. In some facilities it may be difficult or impossible to provide growth into adjacent spaces resulting in a tenant's equipment being located on multiple floors in multiple cages. This can have an adverse effect on higher speed applications that may have distance limitations which can result in cage reconfiguration, additional and/or more expensive equipment costs.

Growth potential in adjacent spaces may also create airflow and cooling issues in your space. This is particularly problematic if adjacent cages do not conform to hot aisle, cold aisle configurations that remain consistent throughout the floor. If the hot aisle, cold aisle arrangements are not maintained throughout all spaces, a company's equipment may suffer from the heat exhausted into their space from nearby cages. The best centers will have proper space and growth planning in place.

Many data centers today are moving towards shielded cabling systems due to noise immunity, security concerns and the robust performance of these cabling systems. As networking application speeds increase to 10 Gigabit Ethernet and beyond, they are more susceptible to external noise such as alien crosstalk. External noise is eliminated with a category 7A shielded cabling system and because of its noise immunity, can provide twice the data capacity as an unshielded cabling system in support of 10GBASE-T. Likewise, category 6A shielded systems eliminate noise concerns and are more popular than their UTP counterparts. As co-location facilities increase temperatures to save energy, tenants need to evaluate the length derating of their cabling systems. Hotter air provided to equipment means hotter air exhausted from equipment. Increased air intake temperatures are supported by active equipment. In the rear of cabinets where the hotter air is exhausted, is typically where cabling is routed. The derating factor for unshielded twisted pair (UTP) cabling is 2x greater than for shielded systems. Increasing temperatures provides a significant cost savings to the tenant and the facility owner.

Whether planning a shielded system or not, there is a requirement for bonding/earthing connections for your equipment, cabinets, pathways and telecommunications circuits, the center's maintenance plan should include a simple check for voltage transients through the bonding/earthing/grounding system since you will be sharing the single ground reference with other tenants.

Ecological planning and options are becoming increasingly important to end users. Customers are demanding sustainable energy, better performing equipment, ISO 14001 certification and RoHS compliance from their vendors, and in some cases LEED, BREAM, Green Star and other Green building certifications depending on the country location. A service provider should be able to provide documentation for a tenant to determine if the site conforms to environmental sustainability expectations.

Finally, space evaluation should include a check to be sure that all of the telecommunications services are available that you currently use, or that there are suitable alternatives. This includes link speed, redundancy, carrier and protocol requirements, available IP addresses, and critical circuit monitoring.

Some end-users are moving to co-location facilities strictly due to lower power costs in some areas of the country, and some are moving due to increased bandwidth needs or better power and carrier infrastructures being available, while others are moving just to get away from their current mess. With all things considered, an outsourced space may be a good solution either permanently or in the interim. With some facilities providing administrative services, this may be an extra benefit to free up company personnel. Either way, the above guidelines should be considered when evaluating use of outsourcing space and services. If needed, Simeon can provide additional information and assistance with your outsourcing plans.

Additional Cloud Considerations for the End User

Business continuity depends on the reliability of the services you place in the cloud. While an email outage is unfortunate and disruptive, database disruptions can cause serious business harm. As an end user, you will want to ask pointed questions about the service, configurations, SLAs, suppliers, etc. While there is some level of confidentiality that cloud providers want to protect, they will be the custodians of whatever you chose to place in their cloud.

A cloud provider should be able to provide you with a listing of suppliers, typical design configuration in their facilities, and what their maintenance and monitoring procedures are throughout the facilities. If a Cloud provider is using outsourced space, then this same information from their provider should also be provided. It may be advantageous to review a site’s SAS 70 (Statement on Auditing Standard 70). SAS 70 is a "Report on the Processing of Transactions by Service Organizations." It provides prospective clients an assurance that the service organization has been thoroughly checked and deemed to have satisfactory controls and safeguards for hosting specific information or processing information.

In several countries in Europe, due to data privacy laws, customer or any private data must reside in country. The cloud provider should be able to provision within a country and provide an assurance that the data will reside there. In country or not, security and monitoring is an important factor.

It is also important to ascertain whether or not a provider is operating via industry standard-compliant infrastructures (defined as cabling, networking, servers and software). Some providers are proprietary only meaning that once applications are developed in that cloud, they may not be able to be ported to another cloud provider. Bandwidth upgrade plans should also be part of the evaluation. Some cloud providers are already built out for 40/100G Ethernet in the backbone and 10G Ethernet in the horizontal. This means there will be less likelihood of downtime or reliance on other sites during upgrades. In short, if they are going to control all or part of your data center, you want to be sure they are using the latest technologies from the start, and that the facility conforms to the latest industry standards.

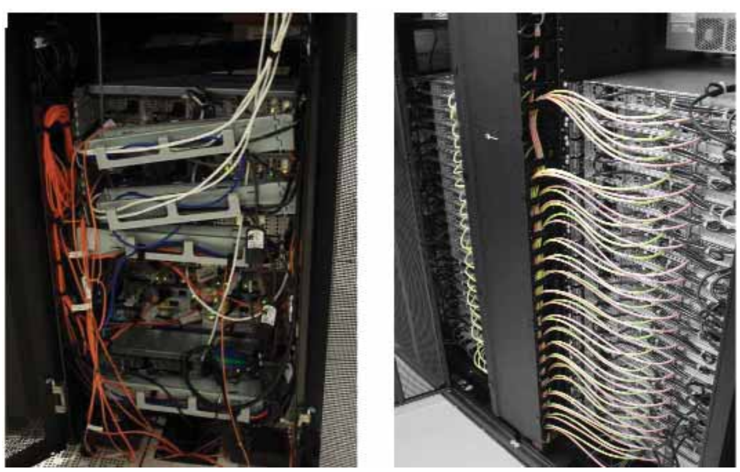

Swing arm cable managers vs.
VersaPOD Zero-U vertical cable management

Data Center Cabling Considerations:

Point-to-Point vs Structured Cabling

The old adage that history repeats itself is very true. If we don't learn from history, we are doomed to repeat it. Many data centers today are victims of historical point-to-point cabling practices.

Direct connections - "Point-to-Point" (i.e. from switches to servers, servers to storage, servers to other servers, etc.) are problematic and costly for a variety of reasons. In the best of data center ecosystems, a standards-based structured cabling system will provide functionality and scalability with the maximum available options for current and future equipment.   While Top of Rack (ToR) and End of Row (EoR) equipment mounting options are now available, these should supplement, not replace, a structured cabling system. ToR and EoR equipment placement both rely heavily on point to point cables, typically fiber jumpers and either twinax copper assemblies or stranded patch cords to connect the network or storage equipment ports to servers.

Data centers are evolving in a rather cyclical manner. When data centers (the original computer rooms) were first built, computing services were provided via a mainframe (virtualized) environment. End users' dumb terminals were connected via point to point with coax or bus cabling using twinax. Enter the PC and Intel based server platforms, and new connections were needed. We have gone through several generations of possible cabling choices: coax (thicknet, thin net), category 3, 4, 5, 5e, 6. Now, the recommended 10 Gigabit capable copper choices for a data center are category 6A, 7 and 7_A channels, OM3 grade fiber for multimode capable electronics and single mode fiber for longer range electronics.

In some data centers, samples of each of these systems can still be found under the raised floor or in overhead pathways, many of which originally were point-to-point. Today however, the "from" point and "to" point are a mystery, making cable abatement (removal of abandoned cable) problematic at best. Compounding this problem was a lack of naming conventions. If the cables were labeled at both ends, the labeling may not make sense anymore. For instance, a cable may be labeled "Unix Row, Cabinet 1." Years later, the Unix row may have been replaced and new personnel may not know where the Unix row was.

There are two standards for structured cabling systems in a data center: TIA 942 and draft ISO 24764, the latter of which is slated to publish in September, 2009.

These standards were created out of need. Both data center standards have language stating that cabling should be installed to accommodate growth over the life of the data center. Moves, adds and changes for a single or a few runs are expensive compared to the same channels run as part of an overall multi-channel installation project. For the larger projects, the end user realizes benefits from project pricing, economies of scale, and lower labor rates per channel. Single channels are typically more expensive, as it is more expensive to send personnel to run one channel. The risk of downtime increases with continual moves, adds and changes. Pathways and spaces can be properly planned and sized up front, but can become unruly and overfilled with additional channels being added on a regular basis.

Data centers that have issues with cable plant pathways typically suffer from poor planning. Growth and new channels were added out of need without regard to pathways. In some cases, pathways do not accommodate growth or maximum capacity over the life of the data center. Overfilled pathways cause problems with airflow, and in some cases cabling becomes deformed due to the weight load, which can adversely affect transmission properties of the channel. This is particularly true in point-to-point systems that have grown into spaghetti-like conditions over time. Likewise, data centers that have not practiced cable abatement or removal of old cabling as newer, higher performing systems are installed experience the same disheveled pathways.

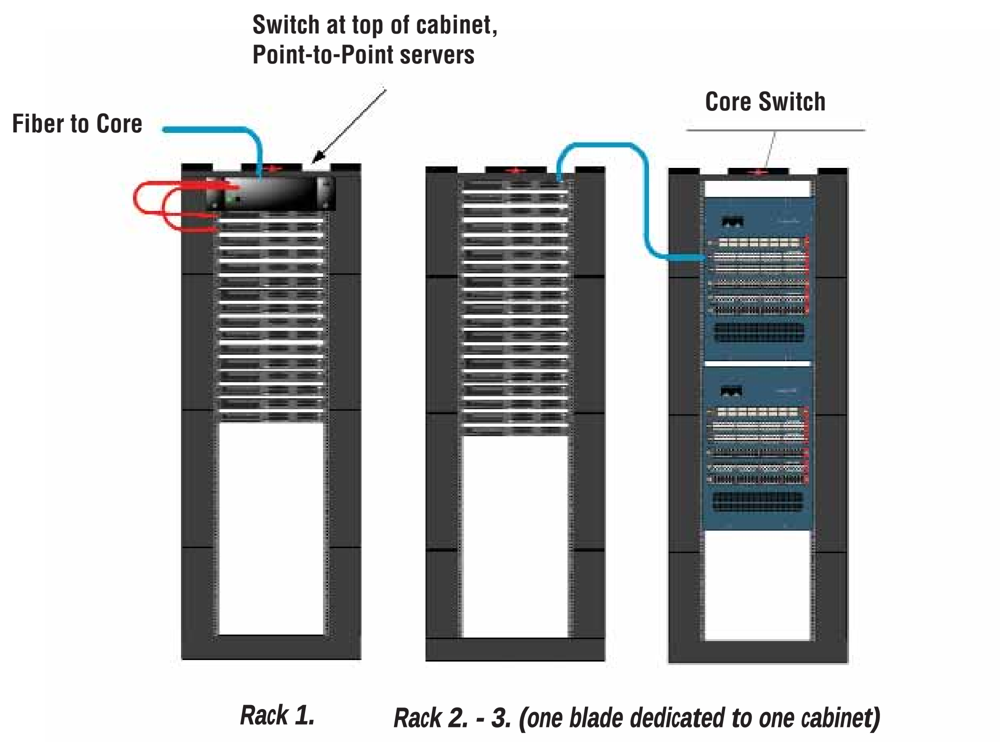

Figure 1: Top of Rack View - Point-to-Point Connections

Copper

Fiber

Figure1. Depicts a ToR patching scenario between switch ports and servers without a structured cabling system. **Rack 2** to **Rack 3** connections are indicative of point-to-point server-to-switch connections, also without a structured system. While proponents of these systems tout a decrease in cabling as a cost offset, further examination may negate such savings.

If a central KVM switch is used, the centralized structured cabling system would need to co-exist anyway, albeit with less channels day one. Newer electronics may have different channel minimum/maximum lengths resulting in the need for new channels. As electronics progress, the structured system may need to be added back to the data center to support future equipment choices, completely negating the savings.

It will cost more to add the structured system later as pathways, spaces, and channels were not planned for and must be installed in a live environment increasing labor costs and the likelihood of downtime. When adding pathways and spaces, fire suppression systems and lighting may need to be moved to accommodate added overhead pathway systems. Floor voids may need to be increased and cabinets may need to be moved to allow new pathways to be routed in a non-obstructive manner for proper airflow.

Further examination highlights other disadvantages of ToR and Point-to - Point methodologies beyond the limitations outlined previously. In either the Rack 1 or Rack 2 -> Rack 3 scenario above, switch ports are dedicated to servers within a particular cabinet. This can lead to an oversubscription of ports. Suppose rack/cabinet 1 had the need for only 26 server connections for the entire rack. If a 48 port switch (ToR switching) or 48 port blade (point-to-point server to switch) is dedicated to the cabinet, this means that 22 additional ports are purchased and maintenance is being paid on those unused ports.

A greater problem occurs when the full 48 ports are used. Adding even one new server will require the purchase of another 48 port switch. In this case, assuming two network connections for the new server, an oversubscription of 46 ports will be added to the cabinet. Even in an idle state, these excess ports consume power. Two power supplies are added to the cabinet. Active maintenance and warranty costs are also associated with the additional switch and ports.

Many of these ToR technologies have limitations for cabling length. Maximum lengths range from 2-15m and are more expensive than a structured cabling channel. Short channel lengths may limit locations of equipment within the shorter cable range. With a structured cabling system, 10GBASE-T can be supported up to 100 meters of category 6A, 7 and 7_A cabling and allows more options for equipment placement within the data center.

Any-to-All Structured Cabling System

The concept behind any-to-all is quite simple. Copper and fiber panels are installed in each cabinet which correspond to copper patch panels installed in a central patching area. All fiber is run to one section of cabinets/racks in that same central patching area. This allows any equipment to be installed and connected to any other piece of equipment via either a copper patch cord or a fiber jumper. The fixed portion of the channel remains unchanged. Pathways and spaces are planned up front to properly accommodate the cabling. While this method may require more cabling up front, it has significant advantages over the life of the data center. These channels are passive and carry no reoccurring maintenance costs as realized with the addition of active electronics. If planned properly, structured cabling systems will last at least 10 years,supporting 2 or 3 generations of active electronics. The additional equipment needed for a point-to-point system will require replacement/upgrade multiple times before the structured cabling system needs to be replaced. The equipment replacement costs, not including ongoing maintenance fees, will negate any up front savings from using less cabling in a point-to-point system.

Figure 2: Racks / Cabinets in Equipment Rows - Central Patching Area

Example of Any-to-All Structured Cabling

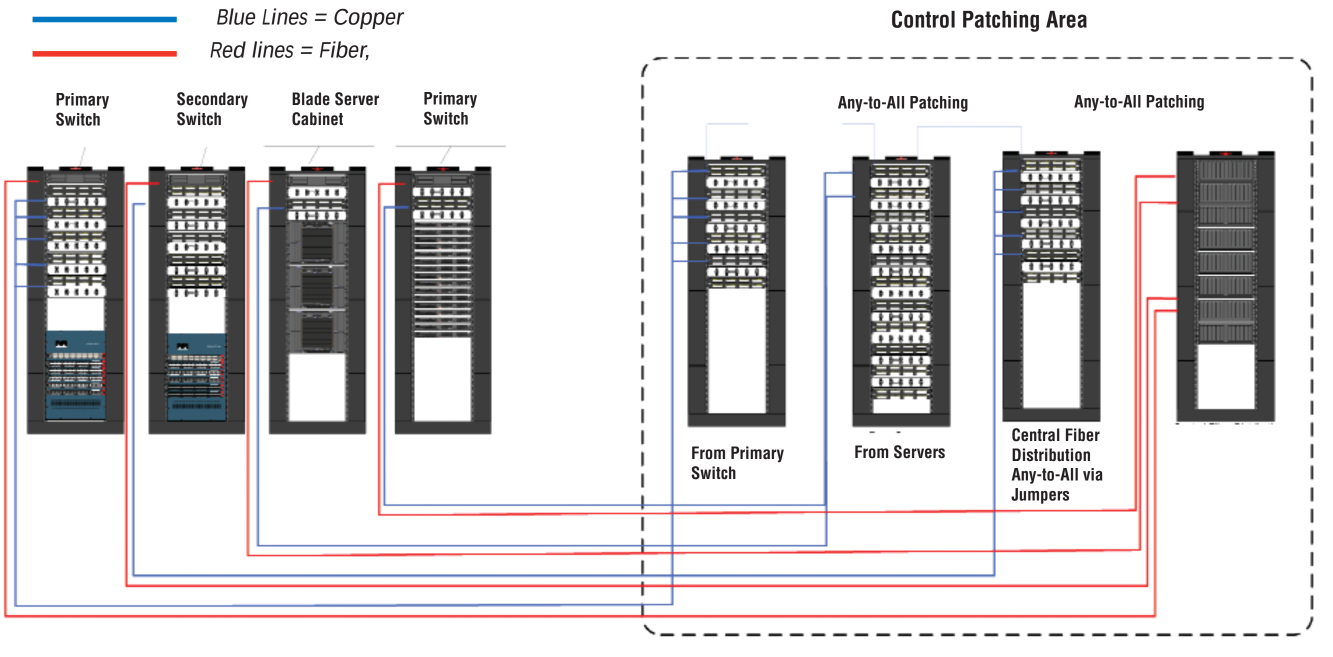

The red lines (fiber connections) all arrive in the central patching area in one location. This allows any piece of equipment requiring a fiber connection to be connected to any other fiber equipment port. For instance, if a cabinet has a switch that requires a fiber connection for a SAN on day one, but needs to be changed to fiber switch connection at a later date, all that is required to connect the two ports is a fiber jumper change in the central patching area. The same is true for copper, although some data centers zone copper connections into smaller zones by function, or based on copper length and pathway requirements. As with the fiber, any copper port can be connected to any other copper port in the central patching area or within the zone.

Cabling standards are written to support 2-3 generations of active electronics. An "any-to-all" configuration assures that the fixed portion of the channels is run once and remains highly unchanged if higher performing fiber and copper cabling plants are used. As a result, there will be less contractor visits to the site for MAC work as the channels already exist. Faster deployment times for equipment will be realized as no new cabling channels have to be run. They are simply connected via a patch cord. Predefined pathways and spaces will not impact cooling airflow or become overfilled as they can be properly sized for the cabling installed. Bearing in mind that the standards recommend installation of cabling accommodating growth, not only will day-one connectivity needs be supported, but also anticipated future connectivity growth needs are already accounted for.

With central patching, switch ports are not dedicated to cabinets that may not require them; therefore, active ports can be fully utilized as any port can be connected to any other port in the central patching area. Administration and documentation are enhanced as the patch panels are labeled (according to the standards) with the location at the opposite end of the channel. Patch cords and jumpers are easy to manage in cabinets rendering a more aesthetically pleasing appearance as cabinets will be tidier. In contrast, with point-to-point cabling, labeling is limited to a label attached to the end of a cable assembly.

With a structured high performing copper and fiber cabling infrastructure, recycling of cabling is minimized as several generations of electronics can utilize the same channels. Being able to utilize all switch ports lowers the number of switches and power supplies. All of these help contribute to green factors for a data center.

To further explain the power supply and switch port impact, contrasting the point-to -point, ToR scenario in section 1, in an "any-to-all" scenario, the 48 ports that would normally be dedicated to a single cabinet (ToR) can now be divided up, on demand, to any of several cabinets via the central patching area. Where autonomous LAN segments are required, VLANs or address segmentation can be used to block visibility to other segments.

Figure 3: Point-to-Point Connections

Top of the Rack view

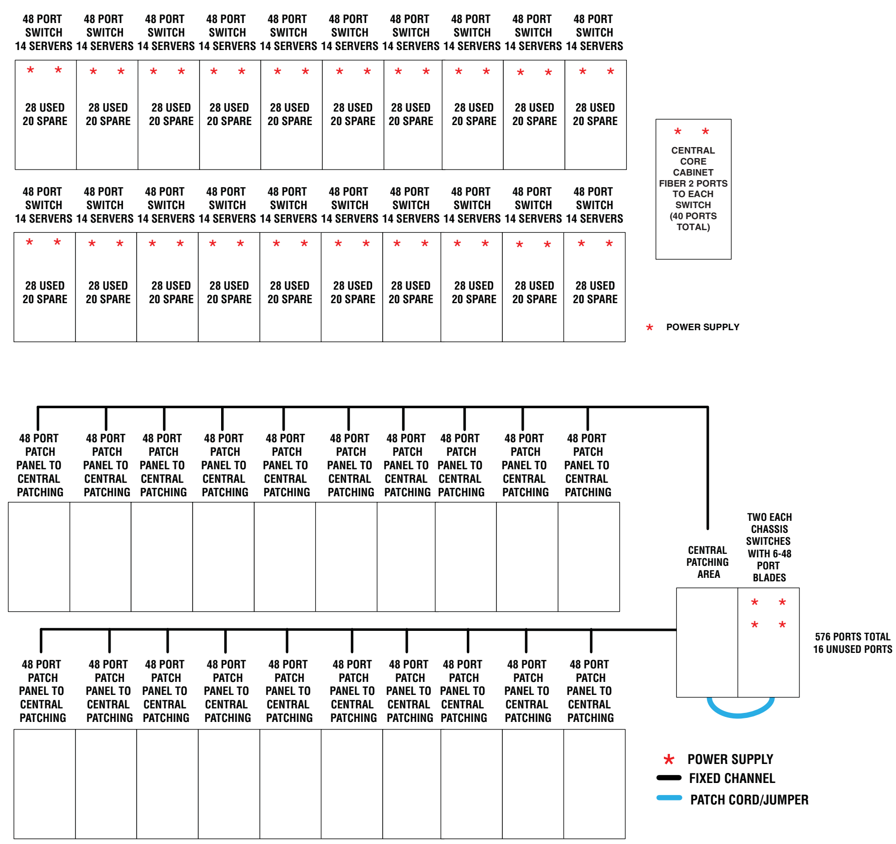

For example: In a data center with 20 server cabinets each housing 14 servers and requiring two network connections each (560 total ports required) the port comparison is shown below.

Note: Table assumes redundant power supplies and VLANs to segment primary and secondary networks. Counts will double if redundant switches are used.

	Number of Switches	Number of Power Supplies (redundant)	Total Ports	Oversubscribed ports
Point-to-Point (ToR)	20 (one 48 port switch per cabinet) 28 con-nections used per cab	40	960	400
Central Any-to-All	2 chassis based with 6 ea 48 port blades	4	576	16

Additional Power Requirements

The real limitation to equipment services within a cabinet is power. Currently in the US, the average power supplied to a cabinet is roughly 6kW[1] with a trend to move towards cabinets that have 18-20kW capacity. As switch ports reach full utilization, the power supplied to the cabinet may not be able to handle the load of a new server and additional switch. This may mean that new power is needed at the cabinet. A complete picture of the power required should be examined before adoption. It may not be possible from a facilities standpoint to provide enough additional power for two devices (4 power supplies in a redundant configuration). According to the Uptime Institute, one of their clients justified a $22 million investment for new blade servers which turned into $76 million after the necessary power and cooling capacity upgrade of $54 million which was required for them to run. [2]

In “Improving Power Supply Efficiency, The Global Perspective” by Bob Mammano, Texas Instruments, “Today there are over 10 billion electronic power supplies in use worldwide, more than 3.1 billion just in the United States.” Increasing the average efficiency of these power supplies by just 10% would reduce lost power by 30 billion kWhrs/year, save approximately $3 billion per year which is equivalent to building 4 to 6 new generating plants.[3] Having a greater number of power supplies (as in ToR) for switches and servers will make it more difficult to upgrade to more efficient power supplies as they are introduced due to the high number of power supplies increasing replacement costs. In a collapsed scenario (central switching, central patching), fewer power supplies are needed and therefore cost less to upgrade.

Virtualization is being implemented in many data centers to decrease the number of server power supplies and to increase the operating efficiency (kW/bytes processed or IT Productivity per Embedded Watt IT-PEW) ratios within equipment. Virtualization also reduces the number of servers and the "floor space" needed to support them. This also reduces the power load to cool the room. Increasing the number of power supplies (ToR) can negate virtualization savings. Further, as servers are retired, the number of needed switch ports decreases. In a ToR configuration, this can increase the number of oversubscribed ports. In an any-to-all scenario dark fiber or non-energized copper cables may exist, but these are passive, require no power, have no reoccurring maintenance/warranty costs, and can be reused for other equipment in the future.

The efficiency of the power supply is only one power factor. To properly examine overall switch to server connections, percentage of processing load, efficiency of the power supply under various loads, required cooling, and voltage required for the overall communications must be factored into overall data center power and efficiency numbers. According to the Uptime Institute the cost to power and cool servers over the next 3 years will equal 1.5 times the price of the server hardware. Future projections extending out to 2012 show this multiplier increasing to almost 3 times even under best case assumptions, 22 times under worst case.[4]

Every port, network, storage, management, etc. contribute to the overall power requirements of a server. According to the US Government Data Center Energy study from Public Law 109-431 signed December 20, 2006, approximately 50% of data center power consumption is power and cooling, 29% is server consumption, and only 5% is attributed to networking equipment. The remainder is divided into storage (a highly variable factor), lighting and other systems. From a networking stand point, looking at port consumption or power draw varies greatly between various architectures (i.e. SFP+, 10GBASE-T and Fiber). Many of these reported power statistics from the manufacturers do not show the entire switch consumption, but rather make a particular architecture sound attractive by only reporting power based on consumption of an individual port, exclusive of the rest of the switch and the higher power server network interface card at the other end of the channel. For instance, a switch might report power consumption of less than 1 watt but the server NIC required can be 15-24 watts.

According to Kevin Tolly of the Tolly Group,[5] “companies that are planning for power studies and including power efficiencies in their RFP documents have difficulties in analyzing the apples to oranges comparisons in response documents. This is because numbers can be reported in a variety of ways. There has been a lack of a standard test methodology leading to our Common RFP project (www.commonrfp.com).” In testing at the Tolly Group, functionality in switching can vary power loads as some switches offload processing from the ASICs chips to CPU which will function at higher power. Edge switches (as those used in ToR configurations) process more instructions in CPU resulting in power spikes that may not be seen without proper testing. The goal of common RFP is to supply end users with some test methodologies to review and compare various architectures and manufacturers.

The switch port power consumption is far less, in most cases, than the server NIC at the opposite end of the channel. There has been a shift in networking led by some vendors for short point to point connections within the racks or near racks as shown in Figure 1. This shift is due in large part due to a need for 10GbE copper connections and a lack of mass manufactured low power 10GBASE-T counterparts using a structured system. The original 10GBASE-T chips had a power requirement of 10-17W per port irrespective of the switch and server power requirements. This is rapidly changing as each new version of silicon manufactured for 10GBASE-T is significantly lower power than the previous iteration. If point-to-point (currently lower power) are used for copper 10GbE communications, coexistence with a structured any-to-all system allows new technologies such as lower power 10GBASE-T to be implemented simply by installing it and connecting it via a patch cord.

End to end power and various power efficiency matrixes are provided by Tolly and The Uptime Institute amongst others. Vendor power studies may not provide a complete picture of what is required to implement the technology. Both of these groups address not only the power consumption of the device, but also the cooling required.

Figure 3
Measured temperatures below the floor and at cabinet heights.

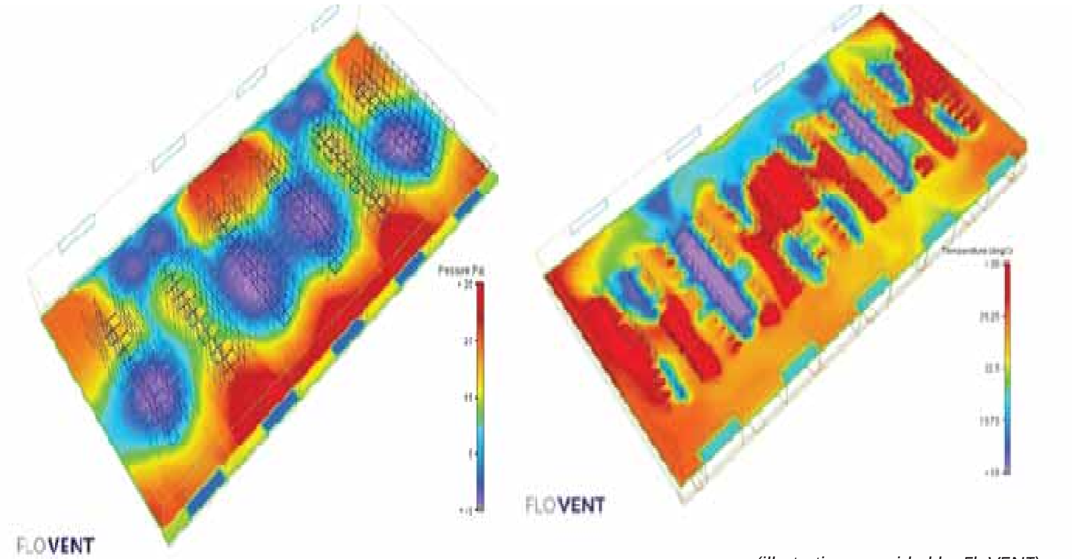

(illustrations provided by FloVENT)

Cooling Considerations

Cooling requirements are critical considerations. Poor data center equipment layout choices can cut usability by 50%.[4] Cooling requirements are often expressed as a function of power, but improper placement of equipment can wreak havoc on the best cooling plans. Point to point systems can land-lock equipment placement.

In Figure 3 above, we can see measured temperatures below the floor and at half cabinet heights, respectively. The ability to place equipment where it makes most sense for power and cooling can save having to purchase additional PDU whips, and in some cases, supplemental or in row cooling for hot spots. In point-to-point configurations, placement choices may be restricted to cabinets where open switch ports exist in order to avoid additional switch purchases rather than as part of the ecosystem decisions within the data center. This can lead to hot spots. Hot spots can have detrimental affects to neighboring equipment within that same cooling zone. Hot spots can be reduced with an any-to-all structured cabling system by allowing equipment to be placed where it makes the most sense for power and cooling instead of being land-locked by ToR restrictions. According to the Uptime Institute, the failure rate for equipment in the top 1/3 of the rack is 3 times greater than that of equipment at the lower 2/3's. In a structured cabling system, the passive components (cabling) are placed in the upper position leaving the cooler spaces below for the equipment. If a data center does not have enough cooling for equipment, placing the switches in a ToR position may cause them to fail prematurely due to heat as cold air supplied from under a raised floor will warm as it rises.

In conclusion, while there are several instances where point-to-point Top of Rack or End of Row connections make sense, an overall study including total equipment cost, port utilization, maintenance and power cost over time should be undertaken including both facilities and networking to make the best overall decision.

.

Simeon has developed several products to assist data center personnel in developing highly scalable, flexible and easy to maintain systems to support various generations of equipment singularly or in conjunction with ToR of Rack systems. Siemon's VersaPOD is an excellent example of one such innovation.

The VersaPOD™ system utilizes a central Zero-U patching zone between bayed cabinets. This space allows for any combination of copper and fiber patching and 19-inch rack-mount PDU's. Should the customer mount the switch in the top of one cabinet, the corner posts are recessed allowing cabinet to cabinet connections and allowing a switch to support multiple server cabinets increasing utilization of the switch ports. This can lower the number of switches required and save energy while providing versatile high density patching options for both copper and fiber.

For information on other Simeon innovations including category 7_A TERA, Z-MAX, category 6A UTP and shielded fiber plug and play and preterminated copper and fiber trunking solutions as well as Siemon's Data Center design assistance services, please visit: **www.simeon.com** or contact your local Simeon representative

Figure 4
VersaPOD™

References:

[1] DataCenter Dynamics, Data Center Trends US, 2008
[2] Data Center Energy Efficiency and Productivity, Kenneth G. Brill, (www.uptimeinstitute.com)

[3] Power Supply Efficiency, The Global Perspective" by Bob Mammano, Texas Instruments
[4] The Economic Meltdown of Moore's Law, The Uptime Institute (www.uptimeinstitute.com)
[5] www.tolly.com and www.commonRFP.com
[6] www.simeon.com/us/versapod and www.simeon.com

Data Center Cooling Best Practices:

-Maximizing power efficiency through smart planning and design

By: Carrie Higbie

Simple mainframe data centers have grown to full fledged Data Centers with a myriad of servers, storage, switching and routing options. As we continue to add equipment to these "rooms" we increase the heat generation while reaching peak capacity. In order to maximize cooling efficiency within Data Centers there are best practices provided by organizations such as ASHRAE (American Society of Heating, Refrigerating, and Air-Conditioning Engineers), which are followed or echoed in many of the industry standards. While some seem to be common sense, others are sometimes neglected.

Addressing Cabling and Pathways

First, and most simply, in order to increase chiller efficiency, it is mandatory to get rid of the old abandoned cabling under raised floors. While cable abatement is a code requirement in some countries due to fuel loads, in all instances and all countries, it makes sense to remove blockages having an impact on air flow to equipment. While working on cable abatement strategies, it is a great time to look at upgrade projects to higher performing cabling which can be either wholly or partially funded through recycling of older copper cable.

While a properly designed under floor cable plant will not cause cooling inefficiencies, when the under floor void is full of cable, a reverse vortex can be created causing the under floor void to pull air from the room rather than push cool air up to the equipment. When pathways and spaces are properly designed, the cable trays can act as a baffle to help maintain the cold air in the cold aisles, or channel the air. Problems occur when there is little or no planning for pathways, They become over filled as many years of abandoned cable fills the pathways and air voids. Overfilling pathways can also cause performance issues. In designing an under floor system, it is critical to look at airflow, void space, cable capacity accommodating growth and other under floor systems such as power, chiller pipes, etc.

In both TIA-942 and the pending ISO 24764 data center standards, it is recommended that structured cabling systems are used and designed accommodating growth so that revisiting the cabling and pathways will not be necessary for the lifecycle of the cable plant. The reasoning behind this is to limit moves, adds and changes, which contribute to the spaghetti we see in many data centers today. In an ideal environment, the permanent link for the channels are run between all necessary cabinets and other central patching locations allowing moves adds and changes to be completed via patch cord changes instead of running new links. Using the highest performing copper cable plant available (currently 7A) assures a longer lifecycle and negates the need for a cable abatement project again in the foreseeable future.

The largest issue with cable abatement is determining which cables can safely be removed. This is compounded in older data centers that have more spaghetti than structure under the floor. One common practice is to upgrade existing copper and fiber cabling utilizing pre-terminated and tested trunking cables. Since cables are combined in a common sheath, once installed and all equipment is cut over to the new system, cables that are not in the common sheath/binder are easily identified for removal. In abatement projects, trunking cables provide the benefit of rapid deployment as the cables are factory terminated to custom lengths eliminating the need for time consuming and labor intensive field terminations.

In some cases, companies move to opposite conveyance systems, i.e. under floor to overhead systems. If moving to an overhead system for abatement, the pathways should be run so that they do not block the natural rise of heat from the rear of cabinets. It is important to consult the proper structural and fire specialties to assure that the ceiling can handle the additional weight, holes for support rods and that the overhead system will not obstruct the reach of fire suppression systems. Just as it is important to plan to accommodate growth under the floor, it is equally important in an overhead system to assure that there is enough room for layers of tray that may be required for overhead pathways.

In order to determine whether an under floor system should be used, the largest factors to consider are the amount of floor void, cooling provided, and layout of the room. For overhead systems, the ceiling height, structural ability to hold mounting brackets, and placement of lighting and fire suppression are the key factors. In both cases, it is important to note that with today's higher density requirements, several layers of trays may be needed in either or both locations.

Running a combination of overhead and under floor systems may be necessary. The past practices of running day one cable tray and/or sizing cable tray based on previous diameters and density requirements can be detrimental to a data center's efficiency during periods of growth. Anticipated growth must be accommodated in day one designs to assure that they will handle future capacity.

Examination of the cabling pathways also includes addressing floor penetrations where the cabling enters cabinets, racks and wire managers. Thinking back to the old bus and tag days in data centers, the standard was to remove half a floor tile for airflow. In many data centers today, that half a tile is still missing and there is nothing blocking the openings to maintain the static pressure under the data center floor. Where the cable penetrations come through the raised floor tiles a product such as brush guards, air pillows or some other mechanism to stop the flow of air into undesirable spaces is paramount.

When you consider that most of the cable penetrations are in the hot aisle and not the cold aisle, the loss of air via these spaces can negatively affect the overall cooling of a data center. In an under floor system, cable tray can act as a baffle to help channel the cold air into the cold aisles if properly configured. While some would prefer to do away with under floor systems if these systems are well designed and not allowed to grow unmanaged, they can provide excellent pathways for cabling.

Cabling pathways inside cabinets are also critical to proper air flow. Older cabinets are notoriously poor at cable management, in large part because that they were not designed to hold the higher concentration of servers that are required today. Older cabinets were typically designed for 3 or 4 servers per cabinet when cabling and pathways were an afterthought. Newer cabinets such as the Simeon VersaPOD™ were designed specifically for data center cabling and equipment providing enhanced Zero-U patching and vertical and horizontal cable management assuring that the cabling has a dedicated without impacting equipment airflow. The same can be said for extended depth wire management for racks such as Siemon's VPC-12.

PODs are changing the face of data centers. According to Carl Claunch of Gartner as quoted in Network World…

"A new computing fabric to replace today's blade servers and a "pod" approach to building data centers are two of the most disruptive technologies that will affect the enterprise data center in the next few years, Gartner said at its annual data center conference Wednesday. Data centers increasingly will be built in separate zones or pods, rather than as one monolithic structure, Gartner analyst Carl Claunch said in a presentation about the Top 10 disruptive technologies affecting the data center. Those zones or pods will be built in a fashion similar to the modular data centers sold in large shipping containers equipped with their own cooling systems. But data center pods don't have to be built within actual containers. The distinguishing features are that zones are built with different densities, reducing initial costs, and each pod or zone is self-contained with its own power feeds and cooling, Claunch says. Cooling costs are minimized because chillers are closer to heat sources; and there is additional flexibility because a pod can be upgraded or repaired without necessitating downtime in other zones, Claunch said."

Lastly, a clean data center is a much better performer. Dust accumulation can hold heat in equipment, clog air filtration gear, and although not heat related, contribute to highly undesirable static. There are companies that specialize in data center cleaning. This simple step should be included yearly and immediately after any cable abatement project.

Inside the cabinets, one essential component that is often overlooked is blanking panels. Blanking panels should be installed in all cabinets where there is no equipment. Air flow is typically designed to move from front to back. If there are open spaces between equipment the air intakes on equipment can actually pull the heated air from the rear of the cabinet forward. The same can be said for spaces between cabinets in a row. Hot air can be pulled to the front either horizontally (around cabinets) or vertically (within a cabinet) supplying warmer than intended air to equipment which can result in failure. In a recent study of a data center with approximately 150 cabinets, an 11 degree temperature drop was realized in the cold aisles simply by installing blanking panels.

Planning for Cooling

Hot aisle, cold aisle arrangements were made popular after the ASHRAE studied cooling issues within data centers. ASHRAE Technical Committee 9.9 characterized and standardized the recommendations.[1] This practice is recommended for either passive or active cooling or a combination of the two. The layout in Figure 1 shows four rows of cabinets with the center tiles between the outer rows representing a cold aisle (cold air depicted by the blue arrows). And the rear faces of the cabinets are directed towards the hot aisles (warmed air depicted by the red arrows). In the past, companies arranged all cabinets facing the same direction to allow an esthetically pleasing showcase of equipment. Looks, however, can be more than deceiving; they can be completely disruptive to airflow and equipment temperatures.

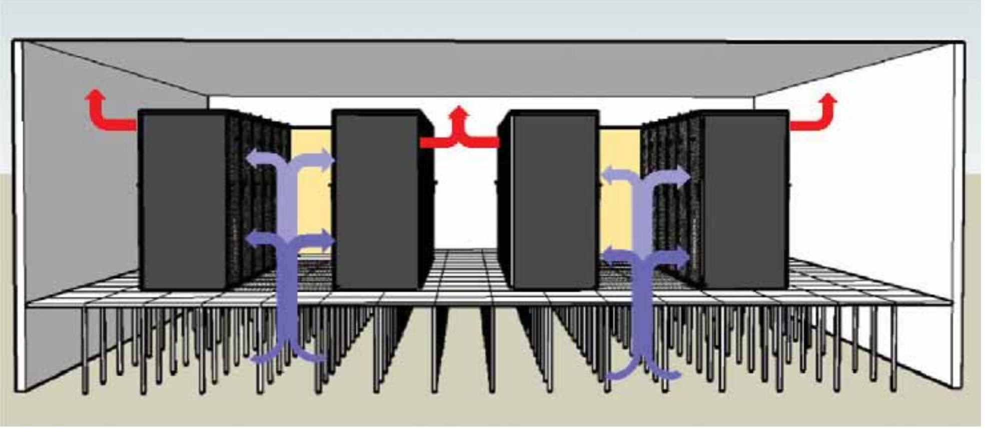

Figure 1: Passive cooling, utilizing airflow in the room and door perforations.

In a passive cooling system, the data center airflow utilizes either perforated doors or intakes in the bottom of cabinets for cold air supply to equipment and perforated rear doors to allow the natural rise of heated / discharged air from the rear of the cabinets into the CRAC (Computer Room Air Conditioner) intake for cooling and reintroduction into the raised floor.

Active cooling systems may be a combination of fans (to force cold air into the faces of cabinets or pull hot air out of the rear roof of cabinets), supplemental cooling systems such as in row cooling, etc. For the purposes of this paper, only passive cooling systems are addressed as the factors for active cooling are as varied as the number of solutions. In order to fully understand the capabilities of each, individual studies and modeling should be performed before any are implemented. ASHRAE recommends pre-implementation CFD (Computational Fluid Dynamics) modeling for the various solutions.

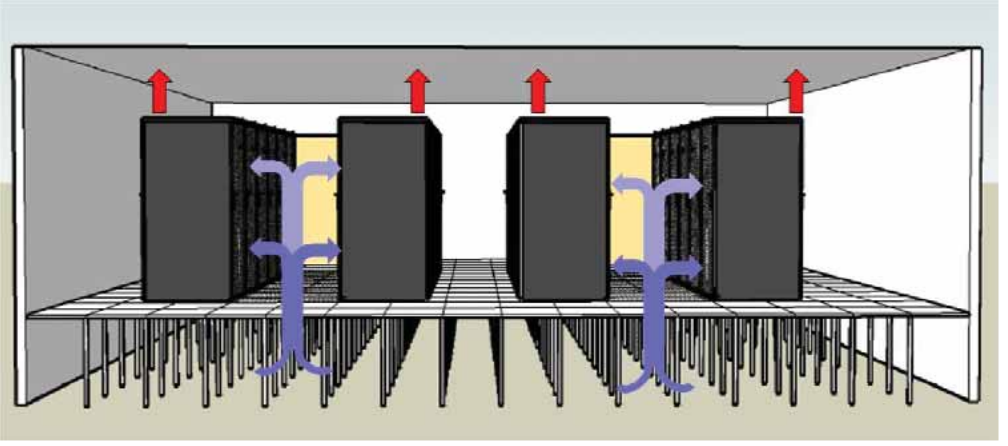

Figure 2: One example of active cooling utilizing fans to pull hot air through the roof

In order to determine the cooling needed, several factors must be known:

- Type of equipment
- Power draw of equipment
- Placement of equipment
- Power density (W/m^2 , W/ft^2)
- Required computer area (m^2, ft^2)

"Computer room floor area totals in the data center would incorporate all of the computing equipment, required access for that equipment, egress paths, air-conditioning equipment, and power distribution units (PDU's). The actual power density is defined as the actual power used by the computing equipment divided by the floor area occupied by the equipment plus any supporting space." [2] This can be defined by the following formula:

Actual power density (W/ft^2) = Computer Power Consumption (W) / required computer area (ft^2)

White space should not be used in the calculations for actual power density. This figure is important when planning a data center. 1U servers have significantly different power density requirements than Blade chassis, storage towers and mainframes. Distribution of this equipment will change the requirements of the various areas of a data center. For instance if a single zone is selected for Blade servers with a greater power density, passive cooling may not provide adequate air temperatures.

In Table 1. IT Equipment Power consumption, it is obvious that one single solution may not address all power needs unless the varied densities are in the initial design. Data Centers using primarily legacy equipment operate at power densities as low as 30W/ft^2 (~320 W/m^2) as compared to more modern higher processing equipment which falls closer to the 60-1000W/ft^2 (~645 to 1,075 W/m^2).

Equipment	W/ft^2 Power Range (~W/m^2)
3U Legacy Rack Server	525 – 735 (~5,645 – 7,900)
4U Legacy Rack Server	430 – 615 (~4,620 – 6,610)
1U Present Rack Server	805 – 2,695 (~8,655 – 28,980)
2U Present Rack Server	750 – 1,050 (8,065 – 11,290)
4U Present Rack Server	1,225 – 1,715 (13,170 – 18,440)
3U Blade Chassis	1,400 – 2,000 (15,050 – 21,500)
7U Blade Chassis	1,200 – 2,300 (12,900 – 24,730)
Mainframe (Large Partitioned Server)	1,100 – 1,700 (11,830 –18,280)

Table 1. IT Equipment Power Consumption[2]

Power consumption can be determined in several ways. Not all will provide an accurate depiction of power needs which in turn would not provide an adequate prediction of cooling demand. Past practices utilized the nameplate rating which as defined by IEC 60950[7] clause 1.7 states "Equipment shall be provided with a power rated marking, the purpose of which is to specify a supply of correct voltage and frequency, and of adequate current-carrying capacity." This rating is a maximum rating as listed by the manufacturer and very rarely will ever be realized. Utilizing this rating will cause oversizing of air conditioning systems and cause a waste in both cooling and money. Most equipment operates at 65-75% of this listing. The correct number to use is measured power consumption. If you will be incorporating new equipment into your data center, equipment manufacturers can provide you with this number.

In addition to the Watts required for equipment, you will also need to determine other sources of heat to be cooled in the data center. This includes lighting, humans, etc., APC has developed a simple spreadsheet to assist with these equations: [3]

Item	Data Required	Heat Output Calcu-lation	Heat Output Subtotal
IT Equipment	Total IT Load Power in Watts	Same as Total IT Load Power in Watts	________Watts
UPS with Battery	Power System Rated Power in Watts	(0.04 x Power System Rating) + (0.05 x Total IT Load Power)	________Watts
Power Distribution	Power System Rated Power in Watts	(0.01 x Power System Rating) + (0.02 x Total IT Load Power)	________Watts
Lighting	Floor Area in Square Feet or Square Meters	2.0 x floor area (sq ft), or 21.53 x floor area (sq m)	________Watts
People	Max # of Personnel in Data Center	100 x Max # of personnel	________Watts
Total	Subtotals from Above	Sum of Heat Output Subtotals	________Watts

Table 2. Data Center Heat Source Calculation Worksheet (Courtesy of APC)

According to APC, cooling capacity is generally about 1.3% of your power load for data centers under 4,000 square feet. For larger data centers, other factors may need to be taken into account such as walls and roof surfaces exposed to outside air, windows, etc. But in general this will give a good indication of overall cooling needs for an average space.

With that said, this is assuming an overall cooling to floor ratio with a similar load at each cabinet. The question gets asked "What cooling can your cabinet support" The variants are significant. Some variants to consider for cabinet cooling include equipment manufacturer recommendations. Many blade manufacturers for instance do not recommend filling cabinets with blades due to cooling and power constraints. According to the Uptime Institute, equipment failures in the top 1/3 of a cabinet is roughly 3x greater than at the lower portion of cabinets. This is due in part to the natural warming of air as heat rises. In order to increase equipment load in high density areas, some form of supplemental cooling may be required. That does not mean that you need to build in-row cooling into every single row, but rather evaluation for high density areas may makes sense. The same may be true for SAN areas and other hotter equipment.

Percentage of door perforation will also be a factor. According to the Industrial Perforators Association, measured air velocity through perforated doors varies with the percentage of perforation. The lower the perforation percentage, the more impact to airflow into the cabinet,as shown in Figure 3.(4) Siemon's VersaPOD™ doors have 71% O.A perforation allowing maximum air flow from cold aisle to hot aisle.

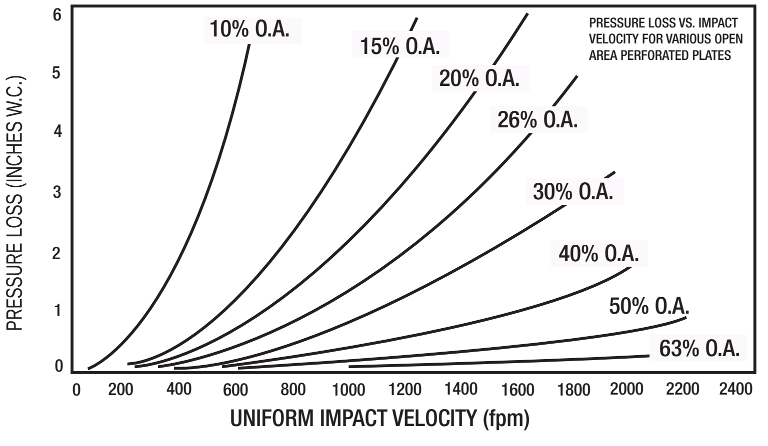

Figure 3: Pressure Loss vs Impact Velocity for Perforated Plates

There are supplemental (active) cooling methods that can be added to cabinets to enhance the airflow either forcing cool air into the cabinets or forcing hot air out. All of these cooling methodologies rely on blanking panels and other steps as outlined earlier in this. There are also workarounds for legacy equipment that utilize side discharge heated airflow, such as legacy Cisco® 6509 and 6513 switches. The newer switch models from Cisco use front to rear airflow.

In side air discharge scenarios, equipment should be isolated cabinet to cabinet so that heated air does not flow into the adjacent cabinet. Some data centers chose to place this equipment in open racks. The Simeon VersaPOD has internal isolation baffles or side panels to assist with this isolation.

Effectiveness of Cooling

Effectiveness of cooling is a necessary test to assure that assumptions made during design are providing the benefits expected. It can also be a good measurement to determine the efficiency of existing data centers and provide a roadmap for remediation on a worst case/first solved basis. The "Greeness" of a data center utilizes two metrics:

1. Data Center Infrastructure Efficiency (DCIE) (a reciprocal of PUE below) is a function of total data center power. This does not just mean servers, but rather includes storage, KVM switches, monitors, control PC's, monitoring stations, etc. Added to the electronics components are all supporting systems such as UPS, PDU's, switch gear, pumps, cooling systems, lighting and the like. The resulting total divided by Total Facility Power will result in DCIE. This is the preferred method used by IBM®. A DCIE of 44% means that for every 100 dollars spent, 44% is actually used by the data center. Improvements in efficiency can bring this number closer to the 100% ideal number.

2. Power Usage Effectiveness (PUE) is another calculation used by some manufacturers. Simply, DCIE = 1/PUE where PUE = Total Facility Power/IT equipment Power. In both cases, the higher the DCIE percentage, the better the data center is on a green scale.

These numbers will not, however, tell you individually how efficient a particular piece of equipment is on the same scale. To determine this, you will need to monitor power at the port for each piece of equipment. New power supplies exist that allow this type of monitoring. When planning for more energy efficient equipment, this can be an invaluable tool.

Another way to measure effectiveness of cooling is to measure cold aisle air temperature throughout the facility. Air is typically measured every other or every third cabinet along the cold aisle. It is normal to see fluctuations in temperature in the hot aisles due to various equipment heat discharge temperatures. But assuring that cool air supply is a consistent temperature will provide you with a clear indication of how well air circulation and conditioning is working. It will also allow you to plan where to put hotter equipment if supplemental cooling will not be introduced.

When active cooling is not an option, a data center will see the best consistency in air temperatures by spacing the hottest equipment around the data center rather than concentrating it all in a single "hot spot" area. Planning is a necessary roadmap for today's hotter equipment. While it may seem logical to have a blade server area, SAN area, etc. In practice, it may be more efficient to have this equipment distributed throughout the data center. It is important to consult your various equipment manufacturers for recommendations.

Regardless of the design methodologies one chooses to follow for their data center, Simeon has resources globally to help. For more information on data center best practices, copper and fiber cabling systems, or the VersaPOD, please visit www.simeon.com or contact your local Simeon representative.

References

(1) Thermal Guidelines for Data Processing Environments. Atlanta: ASHRAE, Inc.

(2) Air-Conditioning Design for Data Centers-Accommodating Current Loads and Planning for the Future; Christopher Kurkjian, PE, Jack Glass, PE, ASHRAE, Inc.

(3) Calculating Total Cooling Requirements for Data Centers, APC, www.apc.com; http://www.apcmedia.com/salestools/NRAN-5TE6HE_R2_EN.pdf

(4) Industrial Perforators Association, http://www.iperf.org/IPRF_DES.pdf

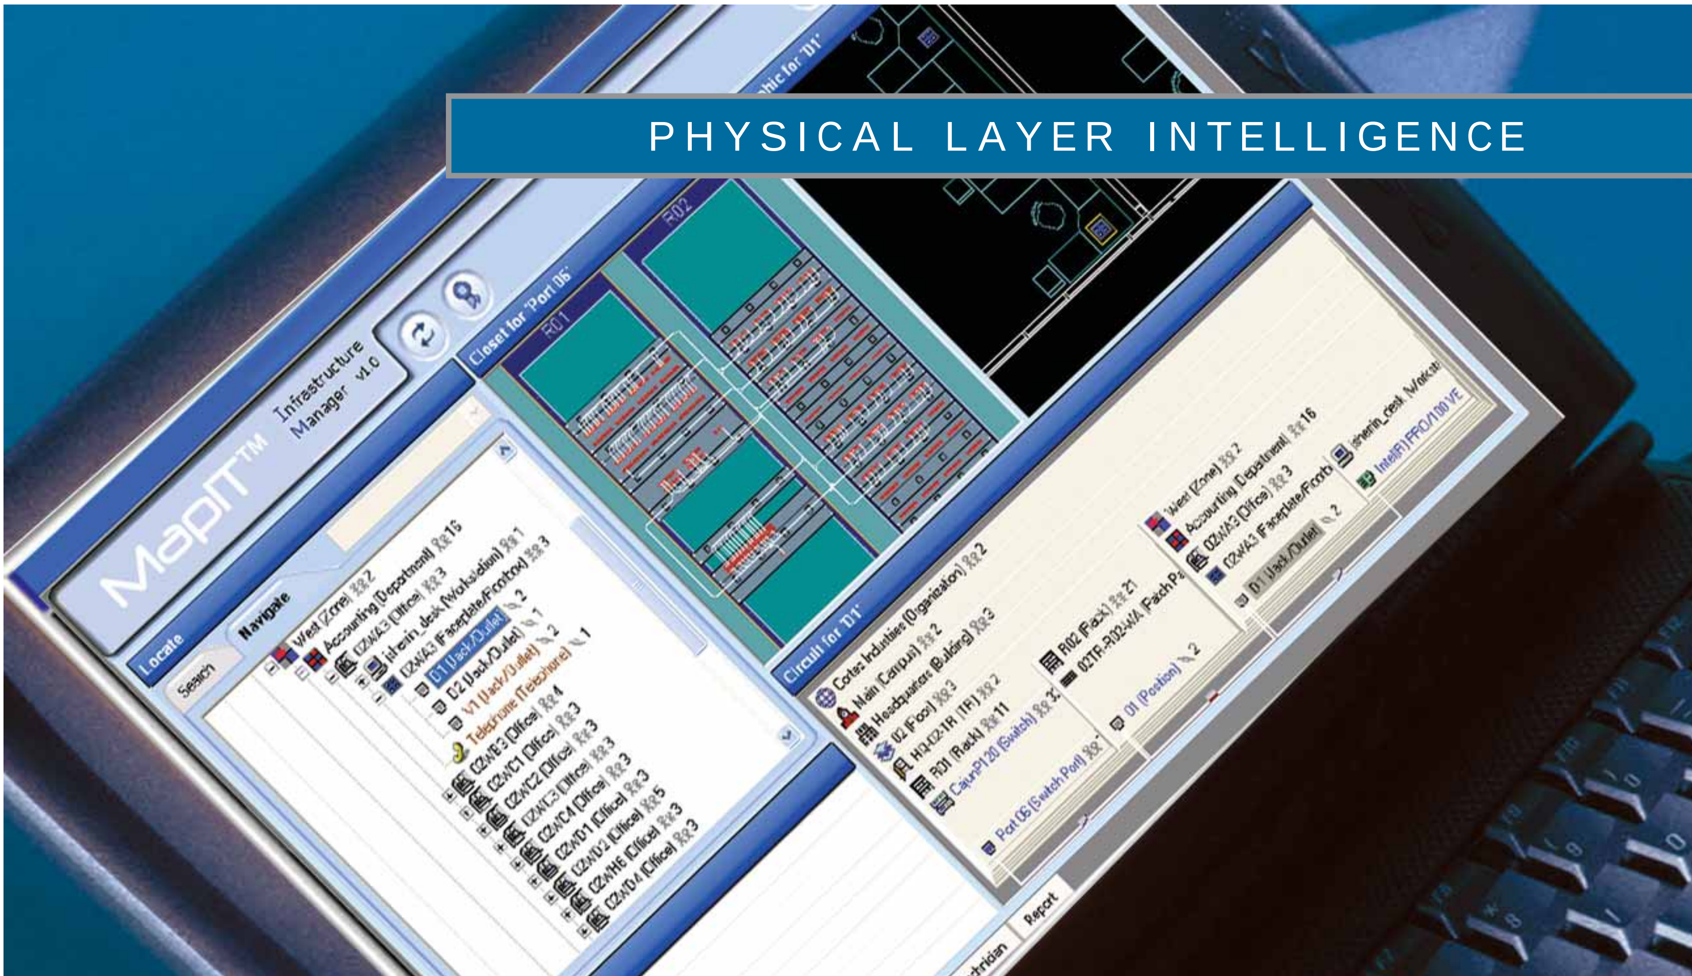

Intelligence at the Physical Layer

Smart Cabling — Better Security

Any network manager will tell you the importance of a fully documented network. This documentation should include all workstations, IP addresses, router configurations, firewall parameters, etc. But this documentation may fall short at the physical layer. In particular, older networks that have gone through many Moves, Adds and Changes (MAC work) are not likely to have current documentation. In real time – during a crisis, this can mean the difference between quickly solving and wasting precious time locating the source of the problem

Perhaps the best illustration is an example taken from a customer that had an issue with a errant device on the network. To provide some background, the company had 5 buildings in the campus. A laptop was creating a denial of service attack from the inside due to a virus. The switch would shut down the port, IT would go to the telecommunications area to determine the location of the misbehaving device. But when IT got to the physical location of the switch, the physical layer (largely undocumented) became an issue – because short of tracing the cable, there was no way to find the location of the laptop. They began tracing the cables only to find that the laptop was no longer there. The laptop user felt that his loss of connectivity was due to a problem with the network. Each time he was disconnected, he moved to another location only to find that after a period of time, he would quickly lose his connection again.

In this scenario, the switches were doing their job by shutting down his port. The user was troubleshooting his own problems. IT was having trouble finding him to correct the problem… and the cycle continued. At one point, the user decided that it must have something to do with the equipment on that particular floor, and moved to another floor. After being disconnected again, he decided that it must be security settings for that building. He then moved to another building. And again, the cycle continued. Roughly 5 hours later, the laptop and user were found and the problems were corrected. For the IT staff, this was 5 hours of pure chaos! For the user, this was 5 hours of pure frustration.

In other scenarios, compliance and overall network security can also be compromised at the physical layer. Most companies have some desks and cubicles that are largely unoccupied and used by more transient staff members. Conference rooms with available ports can also pose a risk. In many vertical markets where compliance is required, these open ports can cause a company to fail their audits unless they are shut down completely or a means exists to allow only certain users can gain access to the network through these connections. The only other option is to firewall these ports from the actual network, which would mean a reconfiguration each time that an authorized network user wanted to utilize the port. All of these risks and their remedies can be burdensome to an IT manager.

In the data center and telecommunications areas, technicians provide an additional risk if they accidentally unplug something that should not be unplugged. Suppose the accidental disconnect was a VoIP switch or a critical server. What if a piece of equipment leaves a facility that contains critical information, as reported many times in the news recently? How does a network manager know who has accessed the network? Where did they access the network? How is access documented? And finally, how are moves, adds and changes managed?

THE INTELLIGENT ANSWER

Intelligent patching has been around for some time, however, the functionality has improved from the original releases. In any of the scenarios above, an intelligent infrastructure management system, such as Siemon's MapIT™ G2 would have allowed the network manager to right click on the offending device, view the entire channel and even locate the device on a graphical map. (See Figure 1).

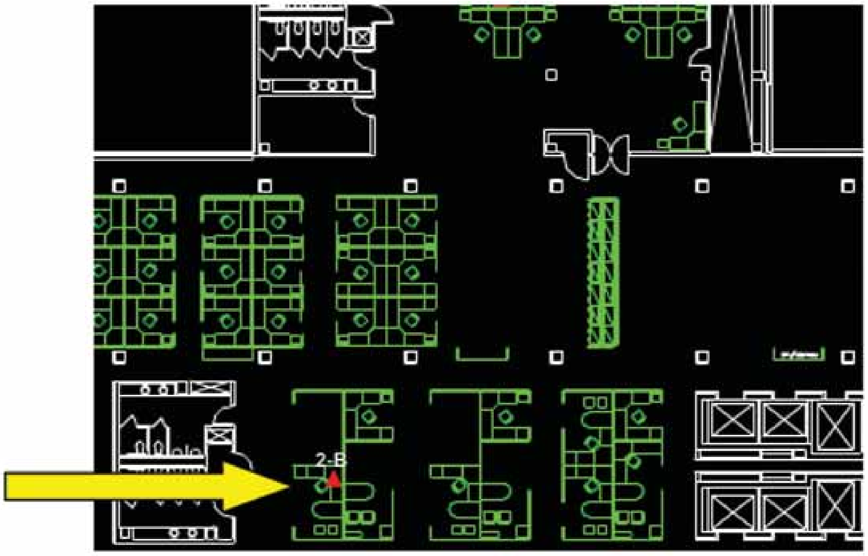

Figure 1: Graphical Layout of Building With Outlet Locations

In the figure above, you will notice that the outlet location is clearly marked on the drawing. By adding the physical layer, network managers are no longer limited to upper layer information only. While knowing MAC address, IP address and logon information is certainly helpful, should physical layer documentation be out of sync with the actual infrastructure, finding problem devices can be a daunting. MapIT™ G2 intelligent patching bridges that gap.

HOW THE SYSTEM WORKS

The system works through a combination of sensor-enabled hardware and software. On the hardware side, MapIT G2 smart patch panels and fiber enclosures are configured with a sensor pad above each port. MapIT G2 patch cords and jumpers have a standard RJ45 interface or a standard fiber connector, and includes a “9th conductor” and contact pin designed to engage the sensor pad. This additional connection allows the system to detect any physical-layer changes in real time. This info is first processed in the smart panels and fiber enclosures and displayed in an on-board graphic LCD for patch cord tracing, diagnostics and technician guidance. A single, twisted-pair cable channel connects the smart panel to a 1U MapIT G2 Master Control Panel, which can monitor up to 2880 ports, relaying the information to the central database running MapIT Im software.

The software is purchased on a per port basis and is written to work either as a standalone application, or can be integrated with an existing network management package. In an integrated configuration, a device and its channel can be traced from within a network management package such as HP OpenView. A simple right click on the device and the MapIT IM software can be launched showing an immediate trace of the physical cable. The trace includes all the information about the channel including patch cords, where the channel terminates, the number of connectors in the channel, and can show the physical location of the device on a CAD drawing.

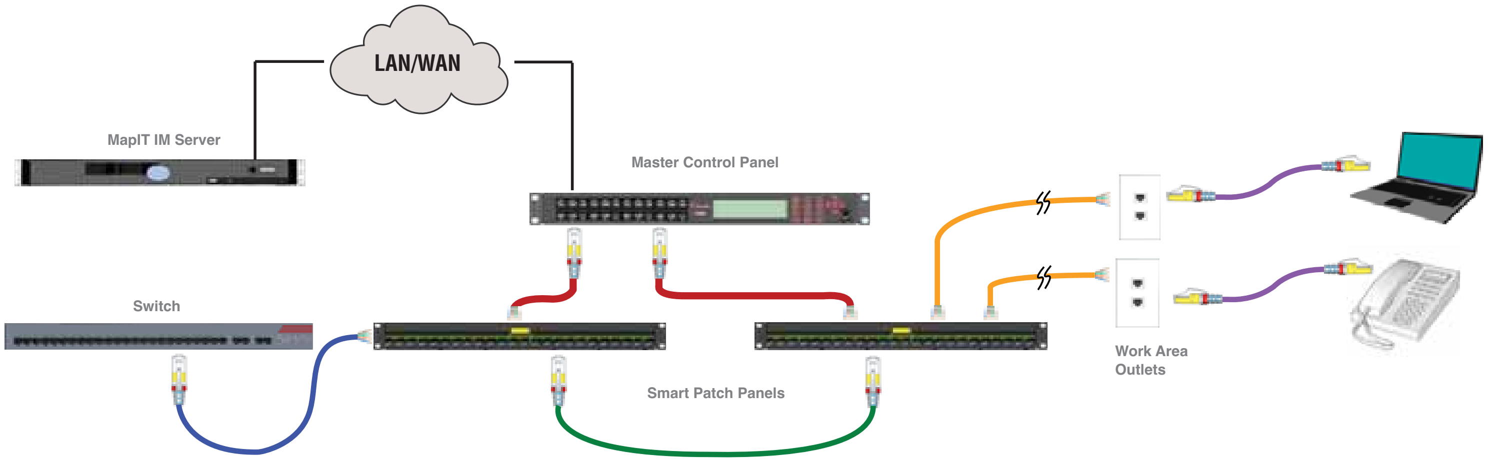

The software reads the object identification information for network devices through SNMP and can also send SNMP (including version 3) traps to shut down ports based on user defined parameters. This provides great benefit when the physical layer is included. For instance, if you wanted to know the location of every PC on your network that was running Windows 2000, you could have it displayed graphically as well as in report format.

The Virtual Wiring Closet (VWC) module provides documentation on the telecommunications rack including connectivity, patch cord length, where each device is connected, etc. It becomes a data dictionary for your racks and/or cabinets. The benefit of MapIT G2 that it will track MAC work without having to update spreadsheets and documentation manually. It also includes a work order module for work order creation. Work orders can be dispatched, displayed onsite on smart anel displays and the changes made are automatically tracked, allowing a manager to know when the work was completed.

This can also be integrated with other security systems such as NetBotz® (owned by APC®) or video cameras. Based on user defined triggers, for instance when someone unplugs a VoIP switch, a camera can snap a picture, write it to the log, and as you would expect from management software, can provide alarms via email, cell, pager, complete with escalation for unanswered alarms. Contacts can be placed on doors to rooms, cabinets, etc. As soon as the contact is broken, the same logging can occur including a photo in the log indicating not only date and time, but additionally photographic/video evidence of the culprit.

While these are only a few of the benefits of MapIT G2, as one can see they are significant. If we go back to the examples at the beginning – in the campus scenario, a simple right click would have saved 5 hours of chasing down a user. Not only would the documentation be up to date, allowing the network manager to know where that switch port terminated in the building, it could also have shown this graphically. They very likely would have gotten to the user before his frustration started and he moved the first time.

Where security and compliance related issues are concerned, the additional documentation and logging abilities not only enhance a company's security position, but also answer many of the compliance related requirements of documentation and access logging. After all, most troubleshooting and investigations start with who, what, where, when, why and how. By adding the physical layer to your overall management the answers to these questions are much easier and more thorough.

For a demonstration of MapIT G2 intelligent patching that provides the full capabilities of the system, please contact your Simeon sales representative. Isn't it time to document and monitor ***all*** of your network?

Appendix

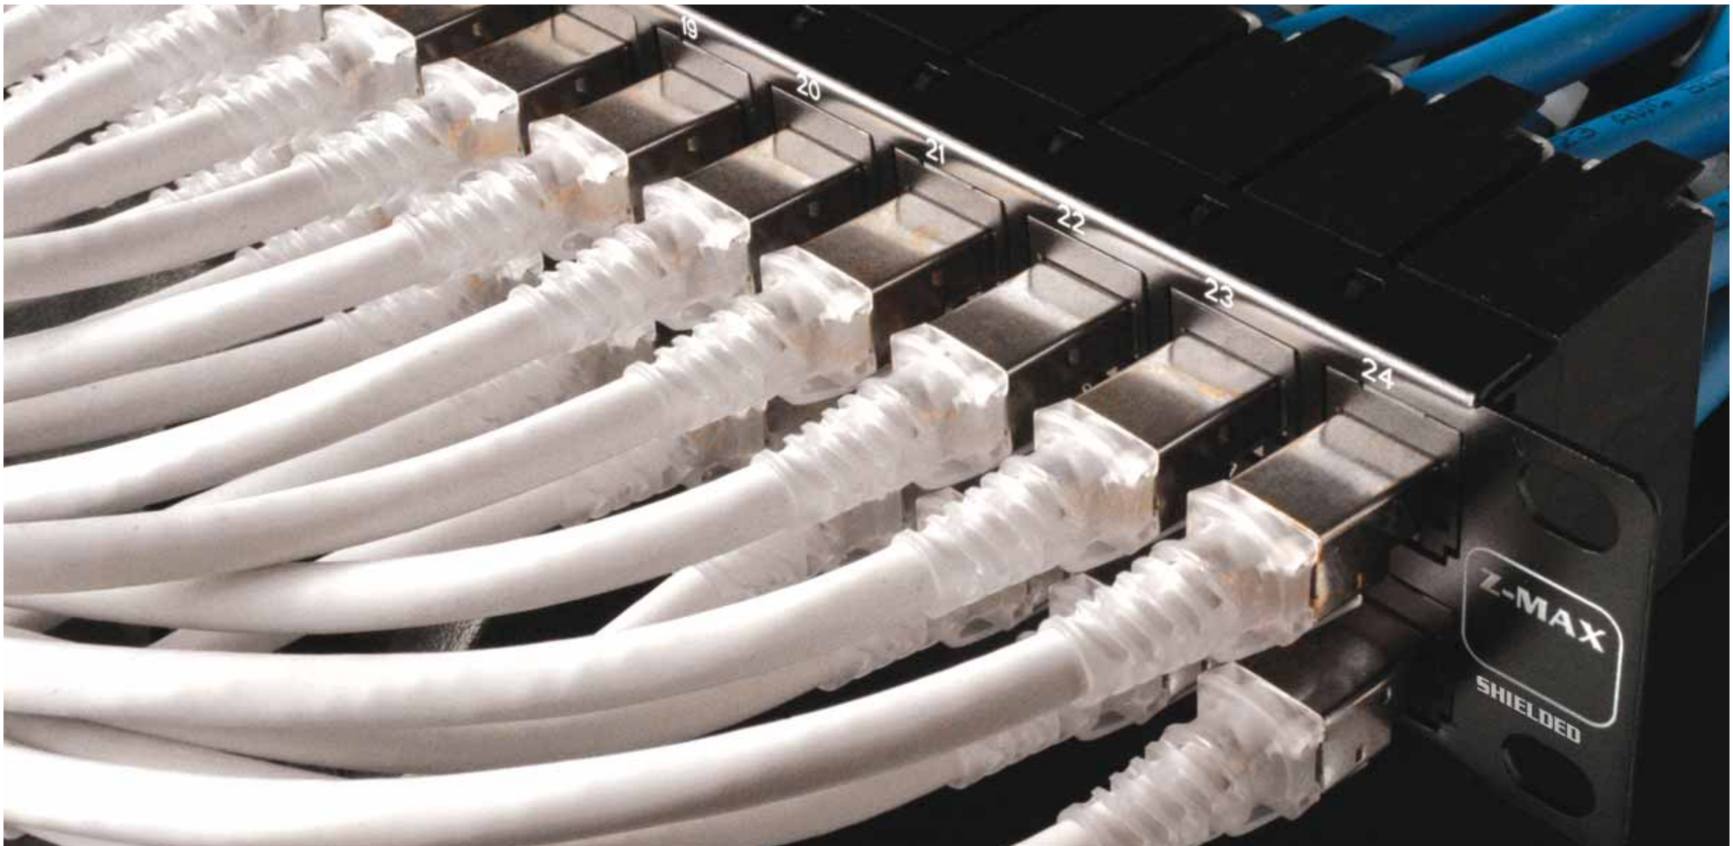

Screened and Shielded Cabling

– Noise Immunity, Grounding, and the Antenna Myth

Screened and shielded twisted-pair copper cabling has been around for quite awhile. A global standard in the 1980s, varieties of screened and shielded have remained a mainstay in some markets, while many others migrated largely to unshielded (UTP) cables.

Recently, however, the ratification of the 10GBASE-T standard for 10Gb/s Ethernet over copper cabling has reestablished the commercial viability of screened and shielded systems and fueled greater adoption of these systems in previously UTP centric markets.

In this competitive landscape, many confusing and often contradictory messages are finding their way to the marketplace, challenging both cabling experts and end-users alike. This whitepaper addresses the most common questions, issues and misconceptions regarding screened and shielded cabling:

CHAPTER 1: Introduction and History of Shielding

In the 1980's, LAN cabling emerged to support the first computer networks beginning to appear in the commercial building space. These first networks were typically supported by IBM Token Ring transmission, which was standardized as IEEE 802.5 in 1985. Cabling for the Token Ring network consisted of "IBM Type 1" cable mated to unique hermaphroditic connectors. IBM Type 1 cable consists of 2 loosely twisted, foil shielded, 150 ohm pairs surrounded by an overall braid as shown in figure 1. This media was an optimum choice for the support of first generation LAN topologies for several reasons. Its design took advantage of the twisted-pair transmission protocol's ability to maximize distance (Token Ring served distances up to 100 meters) and data rates using cost effective transceivers. In addition, the foils and braid improved crosstalk and electromagnetic compatibility (EMC) performance to levels that could not yet be realized by early generation twisted-pair design and manufacturing capability. Not surprisingly, a handful of buildings are still supported by this robust cabling type today.

FIGURE 1: IBM. TYPE 1 CABLE

By 1990, LAN industry experts were beginning to recognize the performance and reliability that switched Ethernet provided over Token Ring. Concurrently, twisted-pair design and manufacturing capabilities had progressed to the point where individual foils were no longer required to provide internal crosstalk isolation and overall shields were not necessary to provide immunity against outside noise sources in the 10BASE-T and 100BASE-T bands of operation. The publication of both the 10BASE-T application in 1990 and the first edition ANSI/EIA/TIA-568 generic cabling standard in 1991, in conjunction with the lower cost associated with unshielded twisted-pair (UTP) cabling, firmly established UTP cabling as the media of choice for new LAN network designs at that time.

15 years later, as Ethernet application technology has evolved to 10Gbps transmit rates, a marked resurgence in the specification of screened and fully-shielded twisted-pair cabling systems has occurred. This guidebook addresses the practical benefits of screens and shields and how they can enhance the performance of traditional UTP cabling designs intended to support high bandwidth transmission. It also dispels common myths and misconceptions regarding the behavior of screens and shields.

CHAPTER 2: Balanced Transmission

The benefit of specifying balanced twisted-pair cabling for data transmission is clearly demonstrated by examining the types of signals that are present in building environments. Electrical signals can propagate in either common mode or differential (i.e. "balanced") mode. Common mode describes a signal scheme between two conductors where the voltage propagates in phase and is referenced to ground. Examples of common mode transmission include dc circuits, building power, cable TV, HVAC circuits, and security devices. Electromagnetic noise induced from disturbers such as motors, transformers, fluorescent lights, and RF sources, also propagates in common mode. Virtually every signal and disturber type in the building environment propagates in common mode, with one notable exception: twisted-pair cabling is optimized for balanced or differential mode transmission. Differential mode transmission refers to two signals that have equal magnitudes, but are 180º out of phase, and that propagate over two conductors of a twisted-pair. In a balanced circuit, two signals are referenced to each other rather than one signal being referenced to ground. There is no ground connection in a balanced circuit and, as a result, these types of circuits are inherently immune to interference from most common mode noise disturbers.

In theory, common mode noise couples onto each conductor of a perfectly balanced twisted-pair equally. Differential mode transceivers detect the difference in peak-to-peak magnitude between the two signals on a twisted-pair by performing a subtraction operation. In a perfectly balanced cabling system, the induced common mode signal would appear as two equal voltages that are simply subtracted out by the transceiver, thereby resulting in perfect noise immunity.

In the real world, however, twisted-pair cables are not perfectly balanced and their limitations must be understood by application developers and system specifiers alike. TIA and ISO/IEC committees take extreme care in specifying balance parameters such as TCL (transverse conversion loss), TCTL (transverse converse transfer loss) and ELTCTL (equal level transverse converse transfer loss) in their standards for higher grade (i.e. category 6 and above) structured cabling. By examining the performance limits for these parameters and noting when they start to approach the noise isolation tolerance required by various Ethernet applications, it becomes clear that the practical operating bandwidth defined by acceptable levels of common mode noise immunity due to balance is approximately 30 MHz. While this provides more than sufficient noise immunity for applications such as 100BASE-T and 1000BASE-T, Shannon capacity modeling demonstrates that this level provides no headroom to the minimum 10GBASE-T noise immunity requirements. Fortunately, the use of shielding significantly improves noise immunity, doubles the available Shannon capacity, and substantially increases practical operating bandwidths for future applications.

An effect of degraded twisted-pair signal balance above 30 MHz is modal conversion, which occurs when differential mode signals convert to common mode signals and vice versa. The conversion can adversely impact noise immunity from the environment, as well as contribute to crosstalk between pairs and balanced cables and must be minimized whenever possible. Shielding can decrease the potential for modal conversion by limiting noise coupled onto the twisted-pair from the environment.

CHAPTER 3: Fundamentals of Noise Interference

All applications require positive signal-to-noise (SNR) margins to transmit within allocated bit error rate (BER) levels. This means that the data signal being transmitted must be of greater magnitude than all of the combined noise disturbers coupled onto the transmission line (i.e. the structured cabling). As shown in figure 2, noise can be coupled onto twisted-pair cabling in any or all of three ways:

1. Differential noise (V_d): Noise induced from an adjacent twisted-pair or balanced cable
2. Environmental noise (V_e): Noise induced by an external electromagnetic field
3. Ground loop noise (V_g): Noise induced by a difference in potential between conductor ends

FIGURE 2: LAN NOISE SOURCES

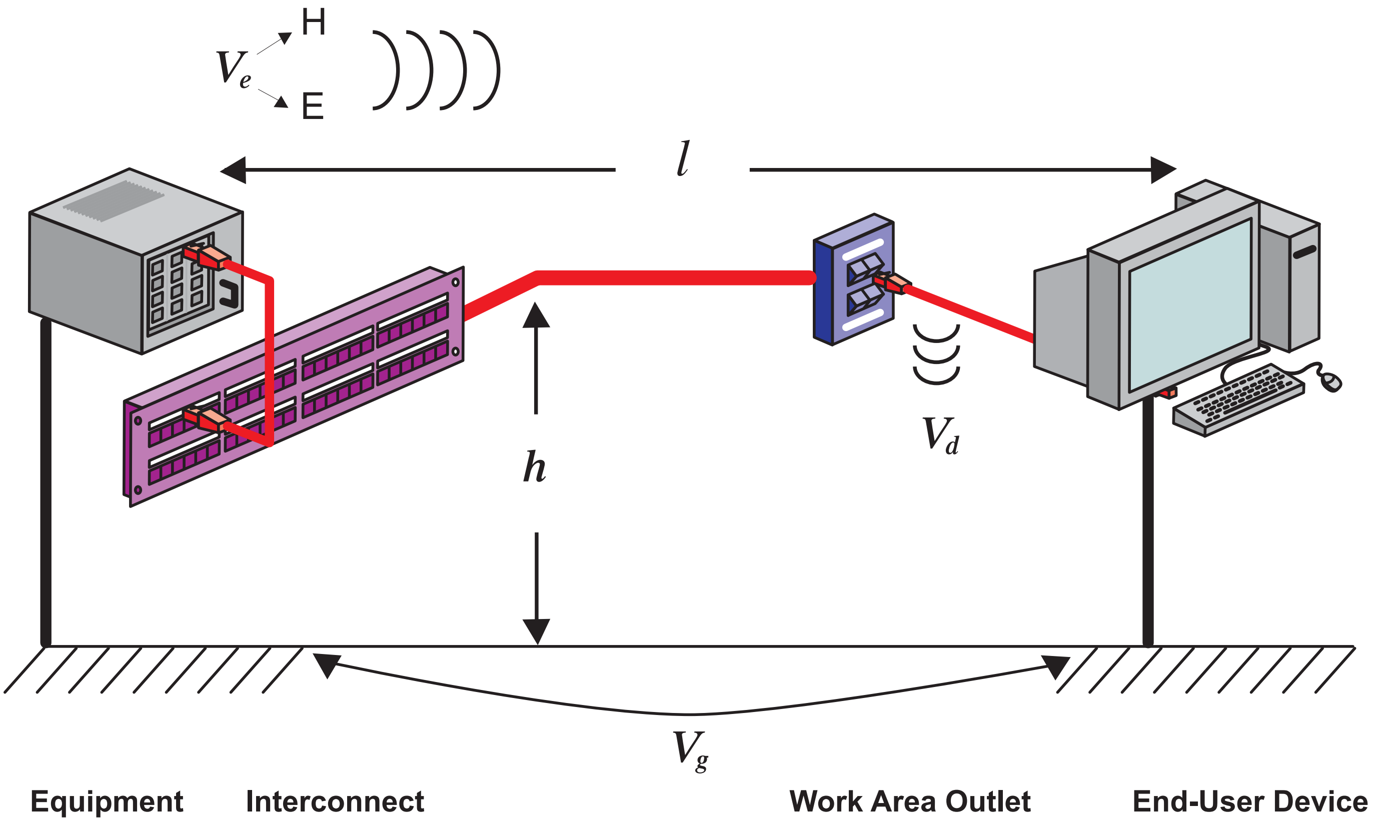

Different applications have varying sensitivity to interference from these noise sources depending upon their capabilities. For example, the 10GBASE-T application is commonly recognized to be extremely sensitive to alien crosstalk (differential mode cable-to-cable coupling) because its digital signal processing (DSP) capability electronically cancels internal pair-to-pair crosstalk within each channel. Unlike pair-to-pair crosstalk, alien crosstalk cannot be cancelled by DSP. Conversely, since the magnitude of alien crosstalk is very small compared to the magnitude of pair-to-pair crosstalk, the presence of alien crosstalk minimally impacts the performance of other applications, such as 100BASE-T and 1000BASE-T that employ partial or no crosstalk cancelling algorithms.

Electromagnetic compatibility (EMC) describes both a system's susceptibility to interference from (immunity) and potential to disturb (emissions) outside sources and is an important indicator of a system's ability to co-exist with other electronic/electrical devices. Noise immunity and emissions performance is reciprocal, meaning that the cabling system's ability to maintain immunity to interference is proportional to the system's potential to radiate. Interestingly, while much unnecessary emphasis is placed on immunity considerations, it is an understood fact that structured cabling systems do not radiate or interfere with other equipment or systems in the telecommunications environment!

Differential noise disturbers: Alien crosstalk and internal pair-to-pair crosstalk are examples of differential mode noise disturbers that must be minimized through proper cabling system design. Susceptibility to interference from differential mode sources is dependent upon system balance and can be improved by isolating or separating conductors that are interfering with each other. Cabling with improved balance (i.e. category 6 and above) exhibits better internal crosstalk and alien crosstalk performance. Since no cable is perfectly balanced, strategies such as using dielectric material to separate conductors or using metal foil to isolate conductors are used to further improve crosstalk performance. For example, category 6A F/UTP cabling is proven to have substantially superior alien crosstalk performance than category 6A UTP cabling because its overall foil construction reduces alien crosstalk coupling to virtually zero. Category 7 S/FTP is proven to have substantially superior pair-to-pair and alien crosstalk performance than any category 6A cabling design because its individual foiled twisted-pair construction reduces pair-to-pair and alien crosstalk coupling to virtually zero. These superior crosstalk levels could not be achieved solely through compliant balance performance.

Environmental noise disturbers: Environmental noise is electromagnetic noise that is comprised of magnetic fields (H) generated by inductive coupling (expressed in A/m) and electric fields (E) generated by capacitive coupling (expressed in V/m). Magnetic field coupling occurs at low frequencies (i.e. 50Hz or 60 Hz) where the balance of the cabling system is more than sufficient to ensure immunity, which means that its impact can be ignored for all types of balanced cabling. Electric fields, however, can produce common mode voltages on balanced cables depending on their frequency. The magnitude of the voltage induced can be modeled assuming that the cabling system is susceptible to interference in the same manner as a loop antenna [1]. For ease of analysis, equation (1) represents a simplified loop antenna model that is appropriate for evaluating the impact on the electric field generated due to various interfering noise source bandwidths as well as the distance relationship of the twisted-pairs to the ground plane. Note that a more detailed model, which specially includes the incidence angle of the electric fields, is required to accurately calculate actual coupled noise voltage.

$$V_e = \frac{2\pi AE}{\lambda} \quad (1)$$

Where: λ is the wavelength of the interfering noise source

A = the area of the loop formed by the disturbed length of the cabling conductor (l) suspended an average height (h) above the ground plane

E = the electric field intensity of the interfering source

The wavelength, λ, of the interfering source can range anywhere from 500,000m for a 60 Hz signal to shorter than 1m for RF signals in the 100 MHz and higher band. The electric field strength density varies depending upon the disturber, is dependent upon proximity to the source, and is normally reduced to null levels at a distance of .3m from the source. The equation demonstrates that a 60 Hz signal results in an electric field disturbance that can only be measured in the thousandths of mV range, while sources operating in the MHz range can generate a fairly large electric field disturbance. For reference, 3V/m is considered to be a reasonable approximation of the average electric field present in a light industrial/commercial environment and 10V/m is considered to be a reasonable approximation of the average electric field present in an industrial environment.

The one variable that impacts the magnitude of the voltage coupled by the electric field is the loop area, A, that is calculated by multiplying the disturbed length of the cabling (*l*) by the average height (*h*) from the ground plane. The cross-sectional view in figure 3 depicts the common mode currents that are generated by an electric field. It is these currents that induce unwanted signals on the outermost conductive element of the cabling (i.e. the conductors themselves in a UTP environment or the overall screen/shield in a screened/fully-shielded environment). What becomes readily apparent is that the common mode impedance, as determined by the distance (*h*) to the ground plane, is not very well controlled in UTP environments. This impedance is dependent upon factors such as distance from metallic raceways, metallic structures surrounding the pairs, the use of non-metallic raceways, and termination location. Conversely, this common mode impedance is well defined and controlled in screened/fully-shielded cabling environments since the screen and/or shield acts as the ground plane. Average approximations for (*h*) can range anywhere from 0.1 to 1 meter for UTP cabling, but are significantly more constrained (i.e. less than 0.001m) for screened and fully-shielded cabling. This means that screened and fully-shielded cabling theoretically offers 100 to 1,000 times the immunity protection from electric field disturbances than UTP cabling does!

FIGURE 3: COMMON MODE CURRENTS

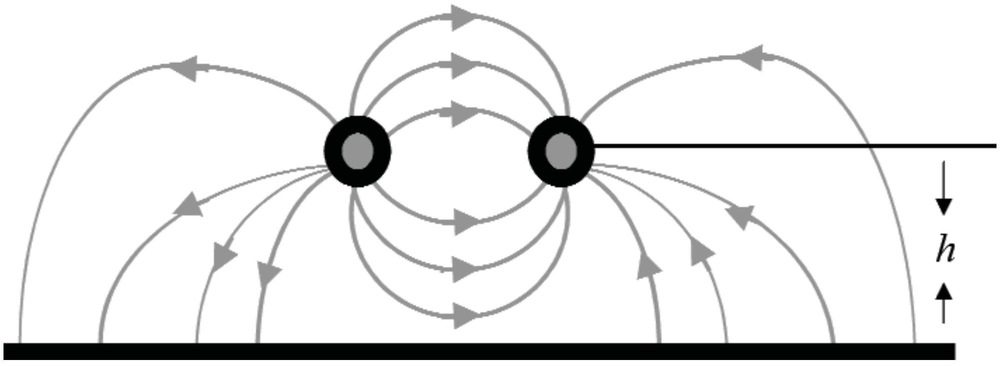

It is important to remember that the overall susceptibility of twisted-pair cables to electric field disturbance is dependent upon both the balance performance of the cabling and the presence of a screen or shield. Well balanced (i.e. category 6 and above) cables should be immune to electromagnetic interference up to 30 MHz. The presence of a shield or screen is necessary to avoid electromagnetic interference at higher frequencies, which is an especially critical consideration for next generation applications. For example, it is reasonable to model that an emerging application using DSP techniques will require a minimum SNR of 20 dB at 100MHz. Since the minimum isolation yielded by balance alone is also 20 dB at 100 MHz, the addition of a screen or shield is necessary to ensure that this application has sufficient noise immunity headroom for operation.

CHAPTER 4: Ground Loops

Ground loops develop when there is more than one ground connection and the difference in common mode voltage potential at these ground connections introduces (generates) noise on the cabling as shown in figure 4. It is a misconception that common mode noise from ground loops can only appear on screens and shields; this noise regularly appears on the twisted-pairs as well. One key point about the voltage generated by ground loops is that its waveform is directly related to the profile of the building AC power. In the US, the primary noise frequency is 60 Hz and its related harmonic, which is often referred to as AC "hum". In other regions of the world, the primary noise frequency is 50 Hz and its related harmonic.

FIGURE 4: INTRODUCTION OF GROUND LOOPS

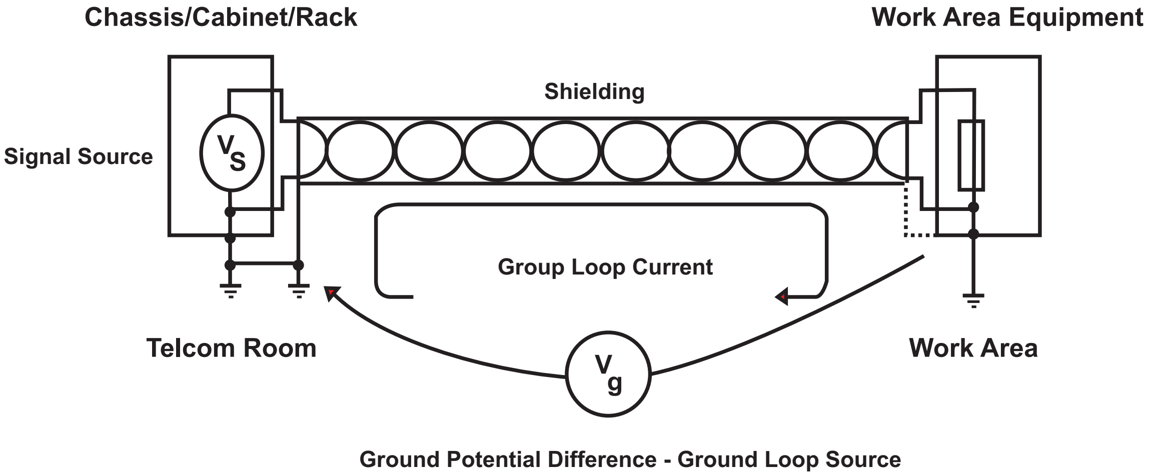

Note: Shield grounded at the TR.
Note: At the WA there is a ground path to shield due to the equipment chassis or cabinet.

Since each twisted-pair is connected to a balun transformer and common mode noise rejection circuitry at both the NIC and network equipment ends, differences in the turns ratios and common mode ground impedances can result in common mode noise. The magnitude of the induced noise on the twisted-pairs can be reduced, but not eliminated, through the use of common mode terminations, chokes, and filters within the equipment.

Ground loops induced on the screen/shield typically occur because of a difference in potential between the ground connection at the telecommunications grounding busbar (TGB) and the building ground connection provided through the network equipment chassis at the work area end of the cabling. Note that it is not mandatory for equipment manufacturers to provide a low impedance building ground path from the shielded RJ45 jack through the equipment chassis. Sometimes the chassis is isolated from the building ground with a protective RC circuit and, in other cases, the shielded RJ45 jack is completely isolated from the chassis ground.

TIA and ISO standards identify the threshold when an excessive ground loop develops as when the difference in potential between the voltage measured at the shield at the work area end of the cabling and the voltage measured at the ground wire of the electrical outlet used to supply power to the workstation exceeds 1.0 Vrms. This difference in potential should be measured and corrected in the field to ensure proper network equipment operation, but values in excess of 1.0 Vrms are very rarely found in countries, such as the US, that have carefully designed and specified building and grounding systems. Furthermore, since the common mode voltage induced by ground loops is low frequency (i.e. 50 Hz or 60 Hz and their harmonic), the balance performance of the cabling plant by itself is sufficient to ensure immunity regardless of the actual voltage magnitude.

CHAPTER 5: Design of Screens and Shields

Shielding offers the benefits of significantly improved pair-to-pair crosstalk performance, alien crosstalk performance, and noise immunity that cannot be matched by any other cabling design strategy. Category 6A and lower rated F/UTP cables are constructed with an overall foil surrounding four twisted-pairs as shown in figure 5. Category 7 and higher rated S/FTP cables are constructed with an overall braid surrounding four individually foil shielded pairs as shown in figure 6. Optional drain wires are sometimes provided.

Shielding materials are selected for their ability to maximize immunity to electric field disturbance by their capability to reflect the incoming wave, their absorption properties, and their ability to provide a low impedance signal path. As a rule, more conductive shielding materials yield greater amounts of incoming signal reflection. Solid aluminum foil is the preferred shielding media for telecommunications cabling because it provides 100% coverage against high frequency (i.e. greater than 100 MHz) leakage, as well as low electrical resistance when properly connected to ground. The thickness of the foil shield is influenced by the skin effect of the interfering noise currents. Skin effect is the phenomenon where the depth of penetration of the noise current decreases as frequency increases. Typical foil thicknesses are 1.5 mils (0.038mm) to 2.0 mils (0.051mm) to match the maximum penetration depth of a 30 MHz signal. This design approach ensures that higher frequency signals will not be able to pass through the foil shield. Lower frequency signals will not interfere with the twisted-pairs as a result of their good balance performance. Braids and drain wires add strength to cable assemblies and further decrease the end-to-end electrical resistance of the shield when the cabling system is properly connected to ground.

FIGURE 5: F/UTP CONSTRUCTION

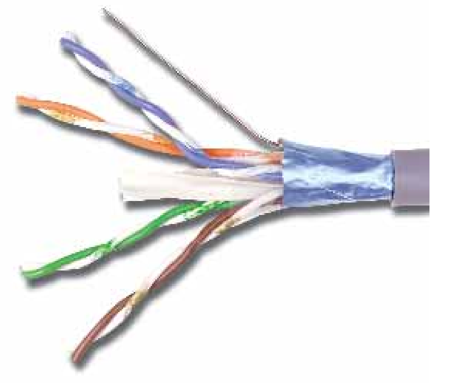

FIGURE 6: S/FTP CONSTRUCTION

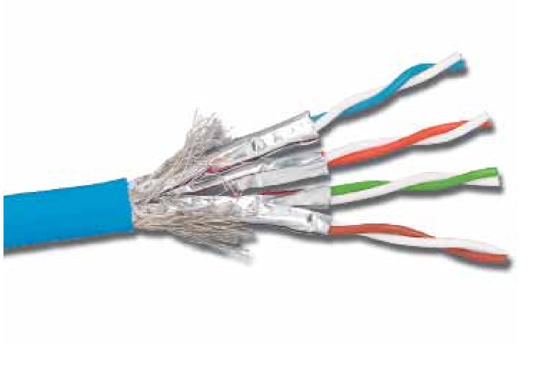

CHAPTER 6: Grounding and Cabling Systems

ANSI-J-STD-607-A-2002 defines the building telecommunications grounding and bonding infrastructure that originates at the service equipment (power) ground and extends throughout the building. It is important to recognize that the infrastructure applies to both UTP and screened/fully-shielded cabling systems. The Standard mandates that:

1. The telecommunications main grounding busbar (TMGB) is bonded to the main building service ground. Actual methods, materials and appropriate specifications for each of the components in the telecommunications grounding and bonding system vary according to system and network size, capacity and local codes.
2. If used, telecommunications grounding busbars (TGB's) are bonded to the TMGB via the telecommunications bonding backbone.
3. All racks and metallic pathways are connected to the TMGB or TGB.
4. The cabling plant and telecommunications equipment are grounded to equipment racks or adjacent metallic pathways.

TIA and ISO standards provide one additional step for the grounding of screened and shielded cabling systems. Specifically, clause 4.6 of ANSI/TIA-568-B.1 and clause 11.3 of ISO/IEC 11801:2002 state that the cable shield shall be bonded to the TGB in the telecommunications room and that grounding at the work area may be accomplished through the equipment power connection. This procedure is intended to support the optimum configuration of one ground connection to minimize the appearance of ground loops, but recognizes that multiple ground connections may be present along the cabling. Since the possibility that grounding at the work area through the equipment may occur was considered when the grounding and bonding recommendations specified in ANSI-J-STD-607-A-2002 were developed, there is no need to specifically avoid grounding the screened/shielded system at the end user's PC or device.

It is important to note the difference between a ground connection and a screen/shield connection. A ground connection bonds the screened/shielded cabling system to the TGB or TMGB, while a screened/shield connection maintains electrical continuation of the cable screen/shield through the screened/shielded telecommunication connectors along the full length of cabling. Part of the function of the screen or shield is to provide a low impedance ground path for noise currents that are induced on the shielding material. Compliance to the TIA and ISO specifications for the parameters of cable and connecting hardware transfer impedance and coupling attenuation ensures that a low impedance path is maintained through all screened/shielded connection points in the cabling system. For optimum alien crosstalk and noise immunity performance, shield continuity should be maintained throughout the end to end cabling system The use of UTP patch cords in screened/shielded cabling systems should be avoided.

It is suggested that building end-users perform a validation to ensure that screened and shielded cabling systems are properly ground to the TGB or TMGB. A recommended inspection plan is to:

1. Visually inspect to verify that all equipment racks/cabinets/metallic pathways are bonded to the TGB or TGMB using a 6 AWG conductor.
2. Visually inspect to verify that all screened/shielded patch panels are bonded to the TGB or TGMB using a 6 AWG conductor.
3. Perform a DC resistance test to ensure that each panel and rack/cabinet grounding connection exhibits a DC resistance measurement of <1 Ω between the bonding point of the panel/rack and the TGB or TMGB. (Note: some local/regional standards specify a maximum DC resistance of <5 Ω at this location.)
4. Document the visual inspection, DC test results, and all other applicable copper/fiber test results.

CHAPTER 7: The Antenna Myth

It is a common myth that screens and shields can behave as antennas because they are long lengths of metal. The fear is that screens and shields can "attract" signals that are in the environment or radiate signals that appear on the twisted-pairs. The fact is that both screens and shields and the copper balanced twisted-pairs in a UTP cable will behave as an antenna to some degree. The difference is that, as demonstrated by the simplified loop antenna model, the noise that couples onto the screen or shield is actually 100 to 1,000 times smaller in magnitude than the noise that is coupled onto an unshielded twisted-pair in the same environment. This is due to the internal pairs' well-defined and controlled common mode impedance to the ground plane that is provided by the screen/shield. Following is an analysis of the two types of signal disturbers that can affect the noise immunity performance of balanced twisted-pair cabling: those below 30 MHz and those above 30 MHz.

At frequencies below 30 MHz, noise currents from the environment can penetrate the screen/shield and affect the twisted-pairs. However, the simplified loop antenna model shows that the magnitude of these signals is substantially smaller (and mostly attenuated due to the absorption loss of the aluminum foil), meaning that unshielded twisted-pairs in the same environment are actually subjected to much a higher electric field strength. The good news is that the balance performance of the cable itself is sufficient up to 30 MHz to ensure minimum susceptibility to disturbance from these noise sources regardless of the presence of an overall screen/shield.

FIGURE 7:
UTP VS. F/UTP SUSCEPTIBILITY

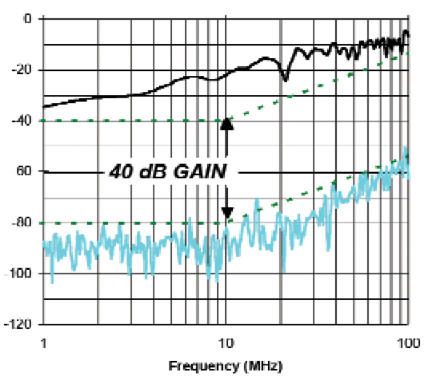

** Data provided courtesy of NEXANS/Berk-Tek*

At frequencies above 30 MHz, noise currents from the environment cannot penetrate the screen/shield due to skin effects and the internal twisted-pairs are fully immune to interference. Unfortunately, balance performance is no longer sufficient to ensure adequate noise immunity for UTP cabling at these higher frequencies. This can have an adverse impact on the cabling system's ability to maintain the SNR levels required by applications employing DSP technology.

The potential for a cable to behave as an antenna can be experimentally verified by arranging two balanced cables in series, injecting a signal into one cable to emulate a transmit antenna across a swept frequency range, and measuring the interference on an adjacent cable to emulate a receiving antenna[2]. As a rule of thumb: the higher the frequency of the noise source, the greater the potential for interference. As shown in figure 7, the coupling between two UTP cables (shown in black) is a minimum of 40 dB worse than the interaction between two properly grounded F/UTP cables (shown in blue). It should be noted that 40 dB of margin corresponds to 100 times less voltage coupling, thus confirming the modeled predictions. Clearly, the UTP cable is radiating and receiving (i.e. behaving like an antenna) substantially more than the F/UTP cable!

A second antenna myth is related to the inaccurate belief that common mode signals appearing on a screen or shield can only be dissipated through a low impedance ground path. The fear is that an ungrounded screen will radiate signals that are “bouncing back and forth” and “building up” over the screen/shield. The fact is that, left ungrounded, a screen/shield will still substantially attenuate higher frequency signals because of the low-pass filter formed by its resistance, distributed shunt capacitance, and series inductance. The effects of leaving both ends of a foil twisted-pair cable ungrounded can also be verified using the previous experimental method. As shown in figure 8, the coupling between two UTP cables (shown in black) is still a minimum of 20 dB worse than the interaction between two ungrounded F/UTP cables (shown in blue). It should be noted that 20 dB of margin corresponds to 10 times less voltage coupling. Even under worst-case, ungrounded conditions, the UTP cable behaves more like an antenna than the F/UTP cable!

FIGURE 8:
UTP VS. UNGROUNDED F/UTP SUSCEPTIBILITY

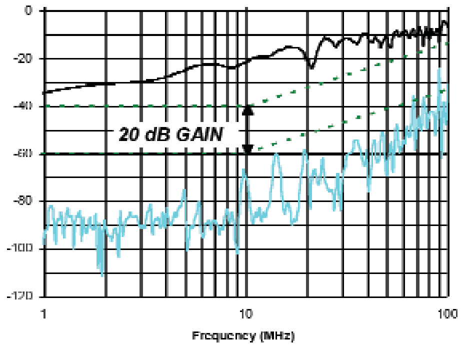

** Data provided courtesy of NEXANS/Berk-Tek*

Modeled and experimental results clearly dispel the antenna myth. It is a fact that screens and shields offer substantially improved noise immunity compared to unshielded constructions above 30 MHz… even when improperly grounded.

CHAPTER 8: The Ground Loop Myth

It is a common myth that ground loops only appear on screened and shielded cabling systems. The fear is that ground loops resulting from a difference in voltage potential between a screen/shielded cabling system’s ground connections cause excessive common mode currents that can adversely affect data transmission. The fact is that both screens and shields and the balanced twisted-pairs in a UTP cable are affected by differences in voltage potential at the ends of the channel.

The difference in the transformer common mode termination impedance at the NIC and the network equipment naturally results in common mode noise current being induced on each twisted-pair. Grounding of the screened/shielded system in multiple locations can also result in common mode noise current induced on the screen/shield. However, these common mode noise currents do not affect data transmission because, regardless of their voltage magnitude, their waveform is always associated with the profile of the building AC power (i.e. 50 Hz or 60 Hz). Due to the excellent balance of the cabling at low frequencies, common mode currents induced onto the twisted-pair either directly from equipment impedance differentials or coupled from a screen/shield are simply subtracted out by the transceiver as part of the differential transmission algorithm.

CHAPTER 9: Why use Screened/Fully-Shielded Cabling

The performance benefits of using screened and fully-shielded systems are numerous and include:

1. Reduced pair-to-pair crosstalk in fully-shielded designs
2. Reduced alien crosstalk in screened and fully-shielded designs
3. Screened category 6A cable diameters are generally smaller than 6A UTP cables allowing greater pathway fill/utilization
4. Substantially improved noise immunity at all frequencies and especially above 30 MHz when cable balance starts to significantly degrade
5. Significantly increased Shannon capacity for future applications

CONCLUSIONS

Achievable SNR margin is dependent upon the combined properties of cabling balance and the common mode and differential mode noise immunity provided by screens and shields. Applications rely on positive SNR margin to ensure proper signal transmission and minimum BER. With the emergence of 10GBASE-T, it's become clear that the noise isolation provided by good balance alone is just barely sufficient to support transmission objectives. The alien crosstalk and noise immunity benefits provided by F/UTP and S/FTP cabling designs have been demonstrated to offer almost double the Shannon capacity and this performance advantage has caught the attention of application developers and system specifiers. It's often said that the telecommunications industry has come full circle in the specification of its preferred media type. In actuality, today's screened and fully-shielded cabling systems represent a fusion of best features of the last two generations of LAN cabling: excellent balance to protect against low frequency interference and shielding to protect against high frequency interference.

BIBLIOGRAPHY

[1] B. Lord, P. Kish, and J. Walling, Nordx/CDT, "Balance Measurements of UTP Connecting Hardware", 1996
[2] M. Pelt, Alcatel Cabling Systems, "Cable to Cable Coupling", 1997
[3] M. Pelt, D. Hess, Alcatel Cabling Systems, "The Relationship between EMC Performance and Applications", 1998
[4] Alcatel Cabling Systems, "The Impact of Cabling Installation Practices on High Speed Performance", 1999
[5] L. Halme and R. Kyoten, "Background and Introduction to EM screening (Shielding) Behaviors and Measurements of Coaxial and Symmetrical Cables, Cable Assemblies, and Connectors", IEE Colloquium on Screening Effectiveness Measurements (Ref. No. 1998/452), pages 8/1-8/7, 1998
[6] S. Hamada, T. Kawashima, J. Ochura, M. Maki, Y. Shimoshio, and M. Tokuda, "Influence of Balance-Unbalance Conversion Factor on Radiated Emission Characteristics of Balanced Cables", IEEE International Symposium on Electromagnetic Compatibility, vol. 1, pages 31-36, 2001
[7] M. Maki, S. Hamada, M. Tokuda, Y. Shimoshio, and H. Koga, "Immunity of Communications Systems using a Balanced Cable", IEEE International Symposium on Electromagnetic Compatibility, vol. 1, pages 37-42, 2001

DEFINITIONS

absorption loss: Signal loss in a metallic media due to impedance losses and heating of the material
alien crosstalk: Undesired differential mode signal coupling between balanced twisted-pair cables
balance: The relationship between the differential signal and common mode signals on a twisted-pair
common mode: Signals that are in phase and are measured referenced to ground
differential mode: Signals that are 180º out of phase and measured referenced to each other
electromagnetic compatibility: The ability of a system to reject interference from noise sources (immunity) and operate without interfering with other devices or equipment (emissions)
equal level transverse conversion transfer loss: The ratio of the measured common mode voltage on a pair relative to a differential mode voltage applied on another pair and normalized to be independent of length
fully-shielded: A construction, applicable to category 7 and 7A cabling, where each twisted-pair is enclosed within an individual foil screen and the screened twisted-pairs are enclosed within an overall braid or foil
ground loop: A difference in voltage potential between two ground termination points that results in an induced common mode noise current
modal conversion: Undesired conversion of differential mode signal to common mode signal and vice versa that results from poor balance
screen: A metallic covering consisting of a longitudinally applied aluminum foil tape
screened: A construction, applicable to category 6A and lower-rated cabling, where an assembly of twisted-pairs is enclosed within an overall metal foil.
Shannon capacity model: A calculation to compute the maximum theoretical amount of error-free digital data that can be transmitted over an analog communications channel within a specified transmitter bandwidth and power spectrum and in the presence of known noise (Gaussian) interference
shield: A metallic covering consisting of an aluminum braid
shielded: See fully-shielded
transfer impedance: A measure of shield effectiveness
transverse conversion loss: The ratio of the measured common mode voltage on a pair relative to a differential mode voltage applied on the same pair
transverse conversion transfer loss: The ratio of the measured common mode voltage on a pair relative to a differential mode voltage applied on another pair

ACRONYMS

BER: Bit error rate
DSP: Digital signal processing
ELTCL: Equal level transverse conversion transfer loss
EMC: Electromagnetic compatibility
F/UTP: Foil unshielded twisted-pair (applicable to category 6A and lower-rated cabling)
IEEE: Institute of Electrical and Electronics Engineers
LAN: Local area network
NIC: Network interface card
S/FTP: Shielded foil twisted-pair (applicable to category 7 and 7A cabling)
SNR: Signal-to-noise margin
TCL: Transverse conversion loss
TGB: Telecommunications grounding busbar
TGMB: Telecommunications main grounding busbar
UTP: Unshielded twisted-pair (applicable to category 6A and lower-rated cabling)
Vrms: Volts, root mean square

IEEE 802.3at PoE Plus Operating Efficiency:

How to Keep a Hot Application Running Cool

The development of the pending PoE Plus standards brings to light a significant new challenge in delivering power over a structured cabling system. The higher power delivered by PoE Plus devices causes a temperature rise within the cabling which can negatively impact system performance. The information in this paper will allow readers to be better equipped to make PoE Plus-ready cabling choices that will support reduced current-induced temperature rise and minimize the risk of degraded physical and electrical performance due to elevated temperature.

HIGHLIGHTS AND CONCLUSIONS:

- Although safe for humans, the 600mA currents associated with the PoE Plus application generate heat in the installed cabling plant.
- Excessive temperature rise in the cabling plant cannot be tested or mitigated in the field
- Excessive temperature rise in the cabling plant can result in an increase in insertion loss and premature aging of jacketing materials.
- Choosing media with improved heat dissipation performance can minimize the risks associated with excessive temperature rise.
- Category 6A F/UTP cabling systems dissipate almost 50% more heat than category 5e cabling.
- Category 7_A S/FTP cabling systems dissipate at least 60% more heat than category 5e cabling.
- It is reasonable to anticipate that category 6A and higher-rated cabling will be the targeted media for the support of tomorrow's high performance telecommunications powering applications.

MARKET OVERVIEW:

The allure of deploying power concurrent with data over telecommunications cabling is undeniable. The benefits of IEEE 802.3af[1] Power over Ethernet (PoE) equipment include simplified infrastructure management, lowered power consumption, reduced operational costs in the case of applications such as voice over internet protocol (VoIP), and even improved safety due to separation from the building's main AC power ring. Market research indicates that the PoE market is on the cusp of significant growth and the numbers are impressive! According to the market research firm Venture Development Corporation[2], approximately 47 million PoE-enabled switch ports were shipped in 2007. Looking forward, the firm expects PoE-enabled switch port shipments to grow at almost double the rate of overall Ethernet port shipments and reach more than 130 million ports by the year 2012.

With it's capability to deliver up to 12.95 watts (W) to the powered device (PD) at a safe nominal 48 volts direct current (VDC) over TIA category 3/ISO class C and higher rated structured cabling, IEEE 802.3af PoE, (soon to be known as "Type 1") systems can easily support devices such as:

- IP-based voice and video transmission equipment,
- IP-based network security cameras,
- Wireless access points (WAPs),
- Radio frequency identification (RFID) tag readers,
- Building automation systems (e.g. thermostats, smoke detectors, alarm systems, security access, industrial clocks/timekeepers, and badge readers),
- Print servers, and bar code scanners

INTRODUCING PoE PLUS:

In 2005, IEEE recognized an opportunity to enhance the capabilities of power sourcing equipment (PSEs) to deliver even more power to potentially support devices such as:

- Laptop computers
- Thin clients (typically running web browsers or remote desktop software applications)
- Security cameras with Pan/Tilt/Zoom capabilities
- Internet Protocol Television (IPTV)
- Biometric sensors
- WiMAX[3] transceivers providing wireless data over long distances (e.g. point-to-point links and mobile cellular access), and high volumes of other devices that require additional power

In support of this need, the IEEE 802.3at[4] task force initiated specification of a PoE Plus or "Type 2" system that can deliver up to 29.5 watts to the powered device (PD) at a safe nominal 53 VDC over legacy TIA category 5/ISO class D:1995 and higher rated structured cabling (note that, for new installations, cabling should meet or exceed TIA category 5e/ISO class D:2002 requirements). Type 2 classification requirements are anticipated to publish as IEEE 802.3 at in mid-2009. Refer to table 1 for a detailed comparison of the capabilities of Type 1 (PoE) and Type 2 (PoE Plus) systems.

TABLE 1:
Overview of PoE and PoE Plus system specifications

	Type 1 - PoE	Type 2 – PoE Plus
Minimum Category of Cabling	Category 3/Class C	Category 5/Class D:1995 with DC loop resistance < 25Ω
Maximum Power Available to the PD	12.95 W	29.5 W
Minimum Power at the PSE Output	15.4 W	30 W
Allowed PSE Output Voltage	44 – 57 VDC	50 – 57 V
Nominal PSE Output Voltage	48 VDC	53 VDC
Maximum DC Cable Current	350 mA per pair	600 mA per pair
Maximum Ambient Operating Temperature	60° C	50° C
Installation Constraints	None	Maximum 5kW delivered power per cable bundle

POE PLUS CHALLENGES:

The development of the pending PoE Plus requirements brought to light a significant new challenge in the specification of power delivery over structured cabling. For the first time, due to the higher power delivered by Type 2 PSE devices, IEEE needed to understand the temperature rise within the cabling caused by applied currents and subsequently specify the PoE Plus application operating environment in such a way as to ensure that proper cabling system transmission performance is maintained. In order to move forward, IEEE enlisted the assistance of the TIA and ISO cabling standards development bodies to characterize the current carrying capacity of various categories of twisted-pair cables.

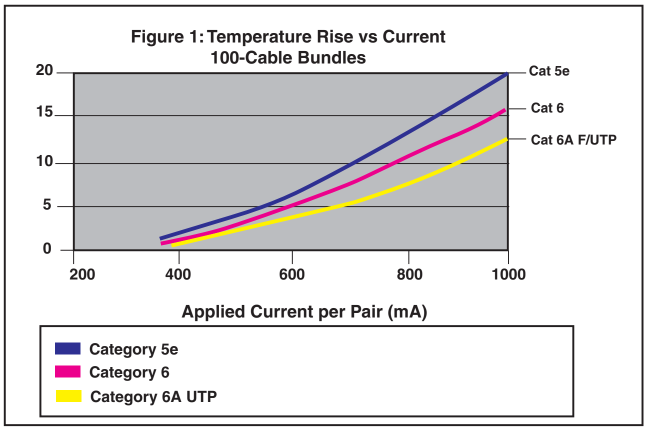

After extensive study and significant data collection, TIA was able to develop profiles of temperature rise versus applied current per pair for category 5e, 6, and 6A cables configured in 100-cable bundles as shown in **Figure 1**. Interestingly, these profiles were created primarily based upon analysis of the performance of unshielded twisted-pair (UTP) cables. They were later corroborated by data submitted to the ISO committee. As expected, since category 5e cables have the smallest conductor diameter, they also have the worst heat dissipation performance and exhibit the greatest temperature rise due to applied current. Note that category 5 cables were excluded from the study since category 5 cabling is no longer recommended by TIA for new installations. IEEE adopted the baseline profile for category 5e cables as representative of the worst-case current carrying capacity for cables supporting the PoE Plus application.

Sidebar No. 1

TEMPERATURE DE-RATING OF UTP VERSUS F/UTP AND S/FTP CABLING SYSTEMS:

It is well known that insertion loss increases (signals attenuate more) as the ambient temperature in the cabling environment increases. To address this issue, both TIA and ISO specify a temperature dependent de-rating factor for use in determining the length that the maximum horizontal cable distance should be reduced by to ensure compliance with specified channel insertion loss limits at temperatures above ambient (20 ºC).

What is not well known is that the de-rating adjustment that is made for UTP cabling allows for a much greater increase in insertion loss (0.4% increase perºC from 20ºC to 40ºC and 0.6% increase per ºC from 40ºC to 60ºC) than the de-rating adjustment that is specified for F/UTP and S/FTP systems (0.2% increase perºC from 20ºC to 60ºC). This means that F/UTP and S/FTP cabling systems have more stable transmission performance at elevated temperatures and are more suited to support applications such as PoE Plus than UTP cabling systems.

Additional TIA guidance recommended that a maximum temperature increase of 10ºC, up to an absolute maximum temperature of 60ºC, would be an acceptable operating environment for cabling supporting PoE Plus applied current levels. In consideration of this input, IEEE chose to reduce the maximum temperature for Type 2 operation to 50ºC, which eliminated the need for complicated power de-rating at elevated temperatures. Next, IEEE had to identify a maximum DC cable current that would not create a temperature rise in excess of 10ºC. An analysis of the worst case category 5e current carrying capacity profile led IEEE PoE Plus system specifiers to target 600 mA as the maximum DC cable current for Type 2 devices, which, according to the TIA profile, results in a 7.2ºC rise in cable temperature. Although this temperature rise is less than the maximum 10ºC value recommended, it provides valuable system headroom that helps to offset additional increases in insertion loss due to elevated temperatures **(See sidebar No. 1)** and minimize the risk of premature aging of the jacketing materials. Operating margin against excessive temperature rise is especially critical because this condition cannot be ascertained in the field.

DISPELLING THE HEAT DISSIPATION MYTH:

Since metal has a higher conductivity than thermoplastic jacketing materials, a thermal model can be used to predict that screened and fully-shielded cables have better heat dissipation than UTP cables. Siemon's data substantiates the model and clearly demonstrates that screened cables exhibit better heat dissipation than UTP cables and fully-screened cables have the best heat dissipation properties of all copper twisted-pair media types. ***Unfortunately, the misconception that screened and fully-shielded systems will "trap" the heat generated by PoE and PoE Plus applications still exists in the industry today. This notion is completely false and easily dispelled by models and laboratory data.***

MEDIA SELECTION:

Interestingly, the PoE Plus application is targeted to be compatible with 10BASE-T, 100BASE-T, and 1000BASE-T, while compatibility with 10GBASE-T is noted as not being precluded by the new Standard. Thus, in an attempt to operate over the largest percentage of the installed cabling base possible, the pending 802.3at Standard specifies ISO '11801 class D:1995[5] and TIA '568-B.2 category 5[6] compliant cabling systems having DC loop resistances less than or equal to 25 ohms as the minimum grade of cabling capable of supporting PoE Plus. Note that these are legacy grades of 100 MHz cabling; TIA recognizes '568-B.2 category 5e[6] cabling and ISO recognizes class D:2002 cabling for new installations. While these objectives represent good news for end-users with an installed base of category 5/category 5e or class D:1995/class D:2002[7] cabling, these cabling systems typically have poor heat dissipation properties and much better choices exist for those specifying new or retrofit cabling plants today.

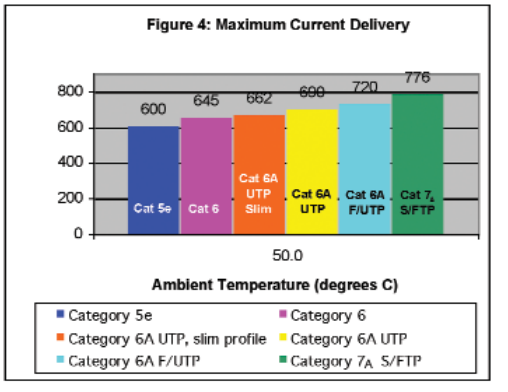

To emphasize, specifying cabling with better heat dissipation characteristics means that:

- Operating temperatures are less likely to exceed 50ºC,
- Certain common installation practices, such as bundling, are less likely to impact overall temperature rise,
- Undesirable increases in insertion loss due to elevated temperatures will be minimized
- The risk of premature aging of cabling jacket materials is reduced.

Good heat dissipation performance exhibited by the cabling plant is especially critical since no methods exist today for monitoring temperature rise in an installation or mitigating a high-temperature environment. Historically, a comfortable level of performance margin is considered to be 50% headroom to Standards-specified limits (this would be equivalent to 6 dB headroom for a transmission performance parameter). Following these guidelines, ***the solutions that offer the most desirable levels of heat dissipation headroom in support of the PoE Plus application are category 6A F/UTP and category 7$_A$ S/FTP cabling systems.*** In fact, category 7$_A$ S/FTP cabling systems dissipate at least 60% more heat than category 5e cables!

BEYOND PoE PLUS:

With the many functional and cost-savings advantages associated with the PoE Plus application, it's easy to predict that the need to supply even more power to the PD is just a few years away. Fortunately, an element of improved heat dissipation is also the ability to support more current delivery within the IEEE maximum 10ºC temperature rise constraint. **Figure 4** shows the maximum current that can be applied over different media types at 50ºC without exceeding maximum temperature rise constraints. Based upon their vastly superior current carrying ability, it's a safe bet that category 6A and higher-rated cabling will be the targeted media for the support of tomorrow's high performance telecommunications powering applications.

DEFINITIONS:

The development of the pending PoE Plus requirements brought to light a significant new challenge in the specification of power delivery over structured cabling. For the first time, due to the higher power delivered by Type 2 PSE devices, IEEE needed to understand the temperature rise within the cabling caused by applied currents and subsequently specify the PoE Plus application operating environment in such a way as to ensure that proper cabling system transmission performance is maintained. In order to move forward, IEEE enlisted the assistance of the TIA and ISO cabling standards development bodies to characterize the current carrying capacity of various categories of twisted-pair cables.

Insertion Loss: The decrease in amplitude and intensity of a signal (often referred to as attenuation).
Type 1: PoE delivery systems and devices
Type 2: PoE Plus delivery systems and devices

ACRONYMS:

º C: . . . . .Degrees Celsius
A: . . . . .Ampere or Amp, unit of current
AC: . . . . .Alternating Current
DC: . . . . .Direct Current
dB: . . . . .Decibel
IP: . . . . .Internet Protocol
IPTV: . . . . .Internet Protocol Television
kW: . . . . .Kilowatt
MHz: . . . . .Megahertz
PD: . . . . .Powered Device
PoE: . . . . .Power over Ethernet, published IEEE 802.3af
PoE Plus: . . . . .Power over Ethernet Plus, pending IEEE 802.3at
PSE: . . . . .Power Sourcing Equipment
F/UTP: . . . . . . .Foil around Unshielded Twisted-Pair (applicable to category 6A and lower-rated cabling)
IEEE: . . . . .Institute of Electrical and Electronics Engineers
ISO: . . . . .International Standards Organization
m: . . . . .Meter
mA: . . . . .Milliampere or Milliamp, unit of current
RFID: . . . . .Radio Frequency Identification
S/FTP: . . . . . . .Shield around Foil Twisted-Pair (applicable to category 7 and 7_A cabling)
TIA: . . . . .Telecommunications Industry Association
UTP: . . . . .Unshielded Twisted-Pair
VDC: . . . . .Volts, Direct Current
VoIP: . . . . .Voice over Internet Protocol
W: . . . . .Watt, unit of power
WAP: . . . . .Wireless Access Point

REFERENCES:

[1]**IEEE 802.3-2005, "IEEE Standard for Information technology:**
Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications", Section Two, Clause 33 (incorporates the content of IEEE Std 802.3af-2003), December 2005
[2]**Venture Deployment Corporation (www.vdc-corp.com), "Power Over Ethernet (PoE):**
Global Market Demand Analysis, Third Edition", March 2008
[3]**Worldwide Interoperability for Microwave Access, Inc.**
[4]**IEEE 802.3at, "IEEE Standard for Information technology:**
Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications Amendment: Data Terminal Equipment (DTE) Power via Media Dependent Interface (MDI) Enhancements", pending publication
[5]**ISO/IEC 11801, 1st edition, "Information technology:**
Generic cabling for customer premises", 1995
[6]**ANSI/TIA/EIA-568-B.2, "Commercial Building Telecommunications Cabling Standard Part 2:**
Balanced Twisted-Pair Cabling Components", May 2001
[7]**ISO/IEC 11801, 2nd edition, "Information technology:**
Generic cabling for customer premises", 2002

WORLD WIDE LOCATIONS

THE AMERICAS

USA..........(1) 866 474 1197
Canada..........(1) 888 425 6165
Columbia - Central and South America Main..........(571) 317 2121
Argentina..........(54) 11 4706 0697
Brasil..........(55) 11 3831 5552
Mexico..........(52) 55 2881 0438
Peru..........(511) 275 1292
Venezuela..........(58) 212 992 5884

EUROPE, MIDDLE EAST AND AFRICA

United Kingdom..........(44) (0) 1932 571771
Germany..........(49) (0) 69 97168 184
France..........(33) 1 46 46 11 85
Italy..........(39) 02 64 672 209

ASIA PACIFIC

Australia (Sydney)..........(61) 2 8977 7500
Australia (Brisbane)..........(61) 7 3854 1200
Australia (Melbourne)..........(61) 3 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 66629661..........(91) 11 66629662
Japan..........(81) (3) 5798 5790