



Facility Considerations for the Data Center

Version 2.0

Abstract

The Cisco® Enterprise Data Center Network Architecture—a comprehensive, adaptive network architecture designed by Cisco Systems® for agile IT support of business needs—critically depends upon underlying physical facilities to provide the power, cooling, physical housing, cabling, physical security, and fire protection that allow IT to function reliably and seamlessly. This paper presents strategies for designing those facilities to support the emerging virtualized computing environment.

Introduction

Data center managers are embracing trends that better align IT with business strategies, increase operational effectiveness, and provide a technology platform for continued growth. The Cisco Enterprise Data Center Network Architecture is a comprehensive, adaptive network architecture that supports an IT organization's immediate demands for consolidation, business continuance, and security while enabling emerging service-oriented architectures, infrastructure virtualization, and on-demand computing. This architecture allows IT managers to deploy technologies that best support their business goals today while enabling efficient implementation of future services and applications.

A crucial part of the Cisco Enterprise Data Center Network Architecture are the physical facilities—the power, cooling, physical housing, cabling, physical security, and fire protection—that allow IT to function. Some in the industry, including APC, use the term Network-Critical Physical Infrastructure (NCPI) for this set of components:

- **Power**—Elements of the power infrastructure include the electrical service entrance of the building, main distribution, generator(s), uninterruptible power supply (UPS) systems and batteries, surge protection, transformers, distribution panels, and circuit breakers.
- **Cooling**—Systems that remove heat from the data center include computer room air conditioners (CRACs) and their associated subsystems—chillers, cooling towers, condensers, ductwork, pump packages, piping—and rack- or row-level cooling or air distribution devices.
- **Cabling**—Data cables use different materials and connectors to optimize performance and flexibility, while the management of the systems maintains this optimization for the long haul. Power cables are also addressed in this paper.
- **Racks and physical structure**—The most critical of these elements are the rack structures housing IT equipment, physical room elements such as dropped ceiling and raised floors, and pathways to manage cabling considerations.
- **Management**—To have reliable facilities, it is important to have visibility of all of their physical components. Management includes systems, such as building management systems, network management systems, element managers, and other monitoring hardware and software.



-
- Grounding—This covers the common bonding network and grounding gear that protect your data center investment from electrostatic discharge.
 - Physical security and fire protection—Subsystems included here are physical security devices at the room and rack level and fire detection/suppression systems.

If these elements are implemented in isolation, as traditionally has been the case, the result is typically a complex and unpredictable system of components that haven't been designed to work together. Management of these components becomes unnecessarily complex because different management systems must be cobbled together, and even then may not provide the full visibility and control necessary for reliable business-critical operation.

If, however, these elements are integrated into a seamless end-to-end system, supported by a systemwide management system, they can provide the essential facilities infrastructure needed to support the Cisco Enterprise Data Center Network Architecture.

Power and UPS Considerations for the Data Center

Today's IT systems have new power-related problems not foreseen when the powering principles for the modern data center were developed over 30 years ago. Five essential requirements have been identified to address the problems with today's power systems:

1. A modular power system that can easily grow and adapt to changing power requirements. System requirements are difficult to predict, and the majority of systems are significantly oversized. Research shows that typical data centers today are utilized to less than 50 percent of their infrastructure capacities. In addition, industry projections show power density requirements that are increasing and unpredictable, yet new data centers must meet requirements for 10 years.
2. Pre-engineered, standardized power solutions that eliminate or simplify the planning and custom engineering to accelerate deployment. The planning and unique engineering involved in traditional power systems takes 6 to 12 months, which is too long compared with the planning horizon of most organizations. Engineering is time-consuming, expensive, and a source of downstream quality problems, making it very difficult to expand or modify the installation later.
3. A power system with mistake-proofing features and decreased single points of failure that improves system availability. According to the Uptime Institute, 40 percent of all downtime in data centers is caused by human error. In addition, traditional UPS systems placed far upstream of the IT loads result in more circuit breakers that act as single points of failure.
4. A management system that provides visibility and control of power at the rack and outlet level. Dynamic power variations among servers and constantly changing and reconfigured loads at the rack level cause unexpected overload conditions and overheated racks. As power densities per rack continue to rise, these problems will only get worse.
5. A power system using standardized, hot-swappable, user-serviceable modules that reduces mean time to recover (MTTR). With traditional systems, spare parts are not readily available, and diagnosis and repair are

invasive. These systems can be so complex that service technicians and maintenance staff make errors and drop loads when operating and maintaining the system.

Meeting these challenges requires a number of changes to current design practice, both in the technology and design of power equipment and in how power is determined and specified in the data center. Integration of the components of the power system must move away from the current practice of unique system designs toward pre-engineered and even pre-manufactured solutions.

UPS System Design Configurations

UPS system designs that distribute power from the utility source of a building to the critical loads of a data center fall into one of five configurations. Which one to select for a particular application is determined by the availability needs, risk tolerance, type of loads in the data center, budgets, and existing power infrastructure. Many variables affect a system's availability, including human error, reliability of components, maintenance schedules, and recovery time. The effect that each of these variables has on the overall system's availability is determined, to a large degree, by the configuration chosen. Table 1 presents the five configurations with their associated availability ranking, "tier" classifications, and cost.

Table 1. Availability and Cost for UPS Configurations

UPS Configuration	Description	Availability Ranking	Tier Classification*	Cost per Rack (US\$)
Capacity (N)	Single UPS module or parallel set of modules	1 = Lowest	Tier I	\$13,500 to \$18,000
Isolated redundant	A primary UPS that normally feeds the load and a secondary ("isolation") UPS that feeds the static bypass of the main UPS	2	Tier II	\$18,000 to \$24,000
Parallel redundant (N+1)	Multiple paralleled, same-size UPSs on a common output bus	3		
Distributed redundant	Three or more UPSs with independent input and output feeders	4	Tier III	\$24,000 to \$30,000
System-plus-system (2N, 2N+1)	Two entirely separate power paths, each able to sustain the load independently	5 = Highest	Tier IV	\$36,000 to \$42,000

*Tiers classify availability based on specific goals defined by the Uptime Institute (www.upsite.com).

Calculating Power Requirements for Data Centers

In addition to selecting the UPS configuration, it is necessary to size the electrical service for the data center. This requires data on the power required by the cooling system, the UPS system, and the IT loads. The power requirements of these elements may vary substantially, but they can be accurately estimated once the power requirements of the planned IT load are determined. In addition to estimating the size of the electrical service, this calculation can also be used to estimate the power output capacity needed for a standby generator.

Table 2 is a worksheet that provides a reasonable estimate of data center power requirements. Once the sizing determination is made, planning can go forward with the assistance of a competent facilities systems supplier or, in the case of larger data centers, a consulting engineer.

Table 2. Data Center Power Requirement Worksheet

Item	Data Required	Calculation	Subtotal kW
Power Requirement—Electrical			
Critical load-sizing calculator value from APC Website	Rating of each IT device	(Calc total in VA x 0.67) / 1000	#1 _____ kW
For equipment not listed in the sizing calculator, critical load—nameplate	Subtotal VA (include fire, security, and monitoring systems)	(Subtotal VA x 0.67) / 1000	#2 _____ kW
Future loads	VA of nameplate of each anticipated IT device	[(Add VA rating of future devices) x 0.67] / 1000	#3 _____ kW
Peak power draw due to variation in critical loads	Total steady-state critical load power draw	(#1 + #2 + #3) x 1.05	#4 _____ kW
UPS inefficiency and battery charging	Actual load plus future loads	(#1 + #2 + #3) x 0.32	#5 _____ kW
Lighting	Total floor area associated with the data center	0.002 x floor area (sq ft) or 0.0215 x floor area (sq m)	#6 _____ kW
Total power to support electrical demands	Total from #4, #5 and #6	#4 + #5 + #6	#7 _____ kW
Power Requirement – Cooling			
Total power to support cooling demands	Total from #7	For chiller systems, #7 x 0.7 For DX systems, #7 x 1.0	#8 _____ kW
Total Power Requirement			
Total power to support electrical and cooling demands	Total from #7 and #8	#7 + #8	#9 _____ kW
Size of Electrical Service Estimate			
Requirements to meet NEC and other regulators	Total from #9	#9 x 1.25	#10 _____ kW
Three-phase AC voltage provided at service entrance	AC voltage		#11 _____ VAC
Electrical service (amperage) required from utility company	Total from #10 and AC voltage in #11	(#10 x 1000) / (#11 x 1.73)	_____ A
Size of Standby Generator Estimate (If applicable)			
Critical loads requiring generator backup	Total from #7	#7 x 1.3*	#12 _____ kW
Cooling loads requiring generator backup	Total from #8	#8 x 1.5	#13 _____ kW
Size of generator needed	Total from #12 and #13	#12 + #13	_____ kW

* This 1.3 factor applies to a fully power-factor-corrected UPS. A 3.0 factor must be used when using traditional double conversion UPS with input harmonic filters.

Cooling Considerations for the Data Center

The design of cooling infrastructure for data centers has changed very little since 1965. As a result, cooling-related problems have become more evident, especially with the advent of high-density computing. Today's cooling systems must address the five key requirements listed in Table 3.

Table 3. Five Key Cooling Requirements

Requirement	Description
Scalability and adaptability	Cooling system requirements are difficult to predict and are generally oversized in hopes of meeting future demand, since it is difficult to add cooling capacity to an existing space. Loads are frequently changed without knowing if cooling has been affected.
Standardization	Customer engineering is time-consuming, expensive, and a key source of later quality problems partly because of the numerous vendors involved in a typical installation. The planning and unique engineering involved takes 6 to 12 months, which is too long compared to the planning horizon of most organizations. It is difficult to transfer knowledge gained from uniquely engineered systems, as customized solutions lead to customized problems.
Simplification	Complex cooling systems lead to a greater likelihood of downtime due to human error, particularly when repairs are complex and time-consuming. In addition, it is difficult to plan and verify redundancy when dealing with customized cooling solutions.
Intelligence	The temperature from rack top to bottom can vary up to 18°F (10°C), placing unexpected stress on individual items of IT equipment, which can result in premature failure of equipment.
Management	Traditional cooling management systems rarely provide information that helps in diagnosing faults at the component level, with reported data often bearing little relation to actual symptoms. Cooling performance data is often not summed from separate CRAC units, providing poor insight into overall system performance.

As with power systems, cooling challenges require changes to current design practice. These include changes in the technology and design of cooling equipment and changes in how cooling requirements are determined and specified in the data center. Standardizing and integrating the components of the cooling system—particularly the air distribution and return systems—can greatly improve availability of the data center.

Comfort versus Precision Cooling

Today's technology rooms require precise, stable environments for sensitive electronics to operate optimally. IT equipment produces an unusual, concentrated heat load and at the same time is very sensitive to changes in temperature or humidity. Standard air conditioning is ill-suited for data centers, leading to system shutdowns and component failures.

Design conditions should be 72 to 75°F (22 to 24°C) and 35 to 50 percent relative humidity. As damaging as the wrong ambient conditions can be, rapid temperature change can also have a negative effect on IT equipment. This is one of the reasons hardware is left powered up, even when not processing data. Precision air conditioning is designed to maintain temperature within 1°F (0.56°C) and humidity within 3 to 5 percent at all times. Ordinary "comfort" systems, on the other hand, are designed to maintain 80°F (27°C) and 50 percent humidity, but only during outside summer conditions of 95°F (35°C) and 48 percent humidity. A poorly maintained data center environment can adversely affect data processing and storage operations:

- High or low temperature—High, low, or rapidly varying temperature can corrupt data processing and shut down an entire system. Temperature variations can alter the electrical and physical characteristics of electronic chips and other board components, causing faulty operation or failure. These problems may be transient or may last for days. Even transient problems can be difficult to diagnose and repair.

- High humidity—High humidity can result in tape and surface deterioration, head crashes, condensation, corrosion, paper handling problems, and gold and silver migration leading to component and board failure.
- Low humidity—Low humidity greatly increases the possibility of static electric discharges, which can corrupt data and damage hardware.

Precision air systems are designed for tight control over temperature and humidity. They provide year-round operation with the ease of service, system flexibility, and redundancy necessary to keep the data center up and running 24 hours a day.

Calculating Cooling Requirements for Data Centers

Sizing a precision cooling system requires an understanding of the amount of heat produced by the IT equipment and by other heat sources in the data center. Common measurements include BTUs per hour, tons per day, and watts. The mixed use of these different units causes unnecessary confusion for users and specifiers. Fortunately, there is a worldwide trend among standards organizations to move toward a common cooling standard—watts. The archaic terms of BTUs and tons (which refers to the cooling capacity of ice) will be phased out over time.

Since the power transmitted by IT equipment through data lines is negligible, the power consumed from AC service mains is essentially all converted to heat. (Power over Ethernet or PoE devices may transmit up to 30 percent of their power consumption to remote terminals, but this paper assumes for simplicity that all electrical power is dissipated locally.) This fact allows for the straightforward representation of IT thermal *output* as watts, equal to its power *consumption* in watts. For further simplicity, the total heat output of a system—therefore, the total cooling requirement—is the sum of the heat output of the components, which includes the IT equipment plus other items such as UPS, power distribution, air conditioning units, lighting, and people. Fortunately, the heat output rates of these items can be easily determined using simple and standardized rules.

The heat output of the UPS and power distribution systems consists of a fixed loss plus a loss proportional to operating power. Conveniently, these losses are sufficiently consistent across equipment brands and models to be approximated without significant error. Lighting and people can also be readily estimated using standard values. The only user-specific parameters needed are a few readily available values, such as the floor area and the rated electrical system power.

Although air conditioning units create significant heat from fans and compressors, this heat is exhausted to the outside and does not create a thermal load inside the data center. This unavoidable energy loss does, however, detract from the efficiency of the air conditioning system and is normally accounted for when the air conditioner is sized.

One can do a detailed thermal analysis using thermal output data for each item in the data center, but quick estimates using simple rules are within the typical margin of error of the more detailed analysis. They have the advantage that they can be done by anyone, without specialized knowledge or training. The worksheet in Table 4 makes it possible to determine the total heat output of a data center quickly and reliably. If walls or the roof get a lot of sunlight or there are other sources of environmental heat, an HVAC consultant should assess how that affects the thermal requirement.

Table 4. Data Center Heat Output Worksheet

Item	Data Required	Heat Output Calculation	Heat Output Subtotal
IT equipment	Total IT load (sum of the power inputs of all IT equipment)	Same as total IT load power in watts	_____ W
UPS with battery	Power system rated power (rating of UPS systems, excluding redundant modules)	$(0.04 \times \text{power system rating}) + (0.06 \times \text{total IT load power})$	_____ W
Power distribution	Power system rated power	$(0.02 \times \text{power system rating}) + (0.02 \times \text{total IT load power})$	_____ W
Lighting	Floor area in square feet or square meters, converted to watts	$2.0 \times \text{floor area (sq ft)}$ or $21.53 \times \text{floor area (sq m)}$	_____ W
People	Maximum number of personnel in data center, converted to watts	$100 \times \text{number of personnel}$	_____ W
TOTAL	_____ W

Cooling Large Node Environments

With the adoption of blade server technology, new challenges arise in air distribution as well as cooling capacity.

Because blade servers concentrate power and cooling requirements in a small form factor, server clusters present a new problem within the data center facility. Table 5 suggests five strategies for cooling these environments.

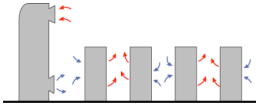
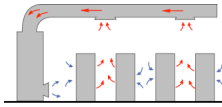
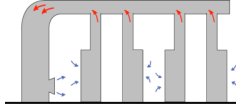
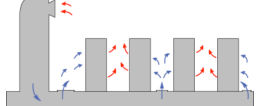
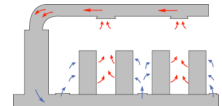
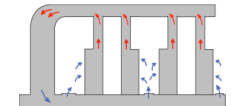
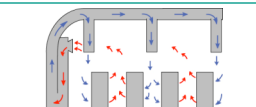
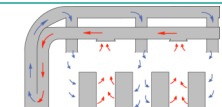
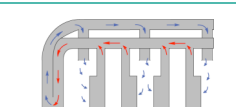
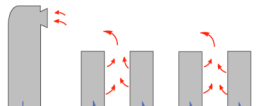
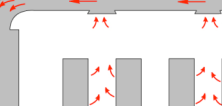

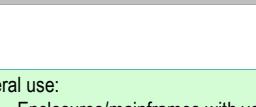
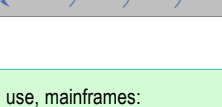
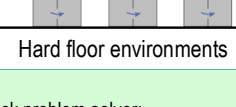
Table 5. Strategies for Cooling Large Node Environments

Strategy	Description
Load spreading	Power and cool to an average value below the peak enclosure value, and spread out the load of any proposed enclosures whose load exceeds the design average value by splitting the equipment among multiple rack enclosures.
Rules-based borrowed cooling	Power and cool to an average value below the peak enclosure value, and use rules to allow high-density racks to borrow adjacent underutilized cooling capacity.
Supplemental cooling	Power and cool to an average value below the peak enclosure value, and use supplemental cooling equipment, as needed, to cool racks with a density greater than the design average value.
Dedicated high-density areas	Power and cool to an average value below the peak enclosure value, provide a special limited area within the room that has high cooling capacity, and limit the location of high density enclosures to that area.
Whole-room cooling	Power and cool any and every rack to the peak expected enclosure density.

Air Distribution Configurations

In designing the cooling system of a data center, the objective is to create a clear path from the source of the cooled air to the intakes of the servers. There are nine basic ways to provide airflow to and from IT equipment that vary in performance, cost, and ease of implementation. For both supply and return, there are three basic methods to convey air between the CRAC and the load: (1) flooded, (2) locally ducted, and (3) fully ducted. Table 6 illustrates the nine possible combinations. APC White Paper 55, "Air Distribution Architecture Options for Mission Critical Facilities," describes each of the nine air distribution methods and the advantages of each.

Table 6. Types of Cooling Systems

	Flooded Return	Locally Ducted Return	Fully Ducted Return
Flooded Supply	 <ul style="list-style-type: none"> Small LAN rooms <40kW: <ul style="list-style-type: none"> Cools racks to 3kW Simple installation Low cost 	 <ul style="list-style-type: none"> General use: <ul style="list-style-type: none"> Cools racks to 3kW No raised floor needed Low cost/ease of install 	 <ul style="list-style-type: none"> Hot rack problem solver: <ul style="list-style-type: none"> Cools racks to 8kW No raised floor needed Increased CRAC efficiencies
Locally Ducted Supply	 <p>Raised floor environment</p>	 <p>Raised floor environment</p>	 <p>Raised floor environment</p>
	 <p>Hard floor environment</p>	 <p>Hard floor environment</p>	 <p>Hard floor environment</p>
Fully Ducted Supply	 <ul style="list-style-type: none"> General use: <ul style="list-style-type: none"> Cools racks to 3kW 	 <ul style="list-style-type: none"> General use: <ul style="list-style-type: none"> Cools racks to 5kW High performance/high efficiency 	 <p>Raised floor environments</p>
	 <ul style="list-style-type: none"> General use: <ul style="list-style-type: none"> Enclosures/mainframes with vertical airflow Raised floor with poor static pressure 	 <ul style="list-style-type: none"> General use, mainframes: <ul style="list-style-type: none"> Enclosures/mainframes with vertical airflow Raised floor with poor static pressure 	 <p>Hard floor environments</p> <ul style="list-style-type: none"> Hot rack problem solver: <ul style="list-style-type: none"> Cools racks up to 15kW Specialized installation

Ten Best Practices to Solving Cooling Problems Caused by High-Density Deployment

Once a cooling system is designed and installed, follow-up checks are important to ensure effective performance of the system. Most importantly, a clear path needs to be maintained from the hot exhaust air of the servers to the return air duct of the CRAC unit. Various factors can reduce the operating efficiency and power density capability of existing cooling systems—10 best practices are presented in Table 7 with the simplest and most cost-effective steps first.

Table 7. Ten Best Practices to Solving Cooling Problems Caused by High-Density Deployment

Practice	Description
Perform a “health check”	<ul style="list-style-type: none"> • Check the overall cooling capacity to ensure that it is not exceeded by the IT equipment in the data center. • Check that all fans are operating properly and that alarms are functioning. Ensure that all filters are clean. • Check condition of the chillers, external condensers, pumping systems, and primary cooling loops. • Check temperature at strategic positions in the aisles of the data center. • Record air intake temperatures at the bottom, middle, and top of each rack and compare with the manufacturer’s recommendations. • Check for gaps within racks (unused rack space without blanking panels, empty blade slots without blanking blades, unsealed cable openings) or excess cabling that may affect cooling performance.
Initiate a cooling system maintenance regime	A regular maintenance regime should be implemented to meet the manufacturer’s recommended guidelines.
Install blanking panels and implement cable management regime	Unused vertical space in rack enclosures allows hot exhaust from equipment to take a shortcut back to the intake and causes the equipment to heat up unnecessarily.
Remove under-floor obstructions and seal floor openings	Under-floor obstructions, such as network and power cables, obstruct airflow and impair the cooling supply to the racks.
Separate high-density racks	When high-density racks are clustered together, most cooling systems become ineffective.
Implement hot-aisle/cold-aisle arrangement	If all rows of racks are arranged with the server intakes facing the same way, progressively hotter intake air is passed from one aisle to the next.
Align CRAC units with hot aisles	Cooling efficiency is greatest when CRAC units are aligned with hot aisles, allowing hot exhaust air to flow directly to the return ducts with little opportunity to mix with cold air streams.
Manage floor vents	Rack airflow and rack layout are key elements in cooling performance. Poorly located supply or return vents can negate nearly all the benefits of a hot-aisle/cold-aisle design.
Install airflow-assisting devices	Where the overall average cooling capacity is adequate but high-density racks create hot spots, fan-assisted devices can improve airflow to high-density racks and can increase cooling capacity to between 3kW and 8kW per rack.
Install self-contained high-density devices	With power densities near or above 8kW per rack, cool air needs to be directly supplied to all levels of the rack—not from the top or the bottom—to ensure an even temperature at all levels.

These practices help keep the data center operating at peak efficiency to maintain the business processes it supports and to prevent future problems. The first eight guidelines help keep a typical data center operating within its original design limits. The last two help the practical design limit for cooling density of a typical data center to be effectively exceeded—without major redesign and construction—by installing self-contained cooling solutions to deal with high-density server applications.

Cabling Considerations for the Data Center

Cabling Topology

Data centers house large numbers of devices with complex networking schemes, so a cabling topology is essential. Network designers are familiar with structured cabling, thanks to the ANSI/TIA/EIA-568 standard, which was first issued in 1991. The Telecommunications Industry Association (TIA) recently completed a standard called TIA-942, “Telecommunications Infrastructure Standard for Data Centers.” This standard uses the topology shown in Figure 1 to describe a data center. The following lists the basic elements that are used in this topology.

Horizontal and Backbone Cabling

As shown in Figure 1, there are two types of cabling: backbone and horizontal.

Backbone cabling connects the main distribution area (MDA), the horizontal distribution area (HDA), including the telecom room, and the entrance room (ER). Backbone cabling consists of backbone cables, main cross-connects, horizontal cross-connects, mechanical terminations, and patch cords, or jumpers, used for backbone-to-backbone cross-connection. Horizontal cabling connects the equipment to the cross-connect located in the MDA or HDA. It consists of horizontal cables, mechanical terminations, and patch cords or jumpers, and may include a zone outlet.

Cross-Connect in the ER or MDA

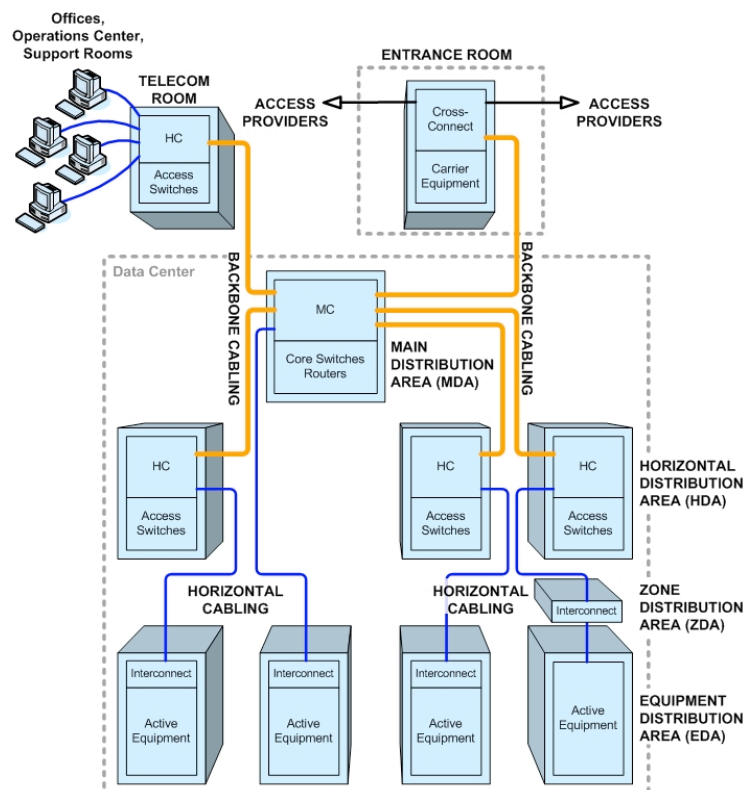
The access provider cables and the data center cables come together in some form of cross-connect. Ideally they are

terminated on a single cross-connect lineup, which effectively becomes “the demarcation.” A centralized location for demarcation to all access providers is often called a “meet-me room.” However, some access providers prefer to terminate their cables in their own racks. The demarcation cross-connect is usually located in the ER, but if a data center does not have one it is located in the MDA.

Main Cross-Connect in the MDA

According to TIA-942, every data center must have an MDA that contains core Ethernet switches and routers, possible core storage area network (SAN) switches, PBX equipment, network management gear, and access provider equipment. The MDA also contains the main cross-connect (MC), used for making connections between the backbone cabling and the equipment cables from the core switches and routers. Because backbone cabling is often fiber, the MC must be designed to handle fiber cables, patch cords, and adapter panels. The cross-connect panels

Figure 1. TIA-942 Basic Data Center Cabling Topology



and network equipment are mounted in cabinets or on racks, depending on the desired security, appearance, and cable management.

Horizontal Cross-Connect in the Telecom Room, HDA, or MDA

The horizontal cross-connect (HC) is a large patch field that distributes horizontal cables out to the equipment. Typically the main job of the HC is to provide a means to patch the horizontal cables to the ports on the access switches. TIA-942 also gives guidance for a centralized optical fiber cabling method, in which the horizontal cables are connected directly to a multifiber backbone cable rather than a switch.

Zone Outlet or Consolidation Point in the Zone Distribution Area

Often it is advantageous to route a group of horizontal cables out to a common termination point. As new equipment is rolled in, it can be connected to the network at this outlet. In comparison to an equipment outlet, which is fixed to the equipment cabinet, the zone outlet offers greater flexibility. The zone distribution area is used to “precable” areas for future use or to facilitate frequent reconfigurations.

Outlet in the Equipment Distribution Area

For many data centers, the majority of the equipment distribution area consists of servers mounted in server cabinets. The horizontal cables are typically routed out to the cabinets and terminated on patch panels, which are then referred to as “outlets” or “interconnects” as shown in Figure 1.

TIA-942 permits certain variations on the topology shown in Figure 1. For example, a small corporate data center may combine the ER, HC, MDA, and telecom room into one area of the computer room. Alternatively, a large Internet data center may require dual ERs for redundancy or to accommodate a diversity of access providers. It may also place the MDA and HDAs in separate rooms or cages for increased security. The topology is determined by factors such as size, scalability, and reliability.

Cabling Media

For today’s data-center LAN and SAN environments, the switch-to-server connections require cabling that has high performance, flexibility, and headroom for future high-bandwidth applications (Table 8). There is a broad base of support for Gigabit Ethernet connectivity: switch ports, network interface cards, and Category 5e and Category 6 UTP cabling. However, Ethernet is constantly evolving, and many organizations are looking to the benefits of 10 Gigabit Ethernet to support new application requirements. Network planners should carefully consider how these new cabling technologies can be employed to infuse bandwidth advantages into the cabling infrastructure.

Table 8. Types of Cabling Media

Cabling Media	Connection Type	Diameter	Minimum Bend Radius	Maximum Data Rate	Maximum Distance*	Common Applications
Category 5E	RJ-45	0.193 in.	1.00 in.	1 Gbps	100 m	Older LANs
Category 6	RJ-45	0.260 in.	1.04 in.	10 Gbps	55 m	IP telephony
Category 6A Category 7	RJ-45	0.310 in.	1.24 in.	10 Gbps	100 m	High-end workstations
Infiniband (CX4 twin-axial)	XENPAK	0.371 in.	4.00 in.	10 Gbps	15 m	Server clusters
Multimode fiber (OM3)	LC, SC, ST, FC, FJ, MPO, XENPAK	0.059 in.	2.00 in.	10 Gbps	220 m	SANs
Single mode fiber (OS1)	LC, SC, ST, FC, FJ, MPO, XENPAK	0.118 in.	2.00 in.	10 Gbps	40 km	WAN

* Maximum distance where maximum data rate can be maintained

Category 5E UTP Cabling

This type of cabling was made popular by the advent of Gigabit Ethernet. Because it was the earliest standard developed by the IEEE, new faster types of cabling have since surpassed it. A Category 5E cabling system is probably sufficient if your business computing is dominated by word processing and simple spreadsheets, and the industry is not expected to drastically change.

Category 6 UTP Cabling

For a flexible infrastructure that can support a wide range of services, Category 6 cabling works appropriately. It provides more than enough bandwidth for gigabit speeds, and it supports 10 Gbps for up to 55 m. For small computer rooms or for modular data centers having server rows tied to nearby HDAs, it may be possible to design horizontal runs that are 55 m or less. However, many medium-to-large data centers need longer cable runs. For example, it may be desirable to consolidate all of the networking equipment into a single MDA that distributes cables to a large number of server cabinets (for hosting or colocation).

Augmented Category 6 UTP Cabling

When cable runs exceed 55 m, the only UTP technology that achieves 10-Gbps transmission for up to 100 m is a new generation of cabling called Augmented Category 6 (C6A). With the development of the 802.3an standard, the IEEE aims to support 10-Gbps data rates using twisted-pair copper cabling over a 4-connector 100-m channel on C6A copper cabling. The standard requires Category 6 electrical channel parameters to be extended from the current 250 MHz to 500 MHz. To support 10 Gbps data rates, a new cable must be designed to improve cable separation in bundles, and new connectors must also be designed so that gains achieved by the cable improvements are not lost in the channel. All components in a 10 Gigabit Ethernet system are important: jack modules, copper cable, patch panels, and patch cords must be precisely tuned to achieve 10 Gbps speeds.

InfiniBand and CX4

Although InfiniBand and CX4 are similar, CX4 has been modified slightly to better address crosstalk and return loss. The two media are discussed together here because they are targeting 15 m using 24 AWG cable and both use the same connectors scheme. The cable is shielded and contains eight twin-axial shielded pairs. The connectors and cabling are based on the InfiniBand standard, developed by the InfiniBand Trade Association.

Although the distance of 15 m hinders longer runs, InfiniBand has found increased usage in server clusters and grid computing. These types of applications require short runs to interconnect multiple servers and require the 10 Gbps bandwidth for a robust server environment. InfiniBand enhances application cluster connections through its utilization of quality-of-service (QoS) mechanisms, which helps prioritize high-priority transactions. It is not suitable for a full-scale structured cabling plant as defined by TIA-942, but for this server cluster purpose, InfiniBand is both recognized and often used by industry leaders like Cisco and IBM.

InfiniBand connections may also be used in SANs to avoid costly host bus adapters (HBAs) used with Fibre Channel. The advantages of removing an HBA from your cabling topology should be measured against the bandwidth and reliability provided with Fibre Channel.

Fiber Cabling

As listed in Table 8, the IEEE 802.3 standards offer a number of choices for 10-Gbps transmission over fiber. The common transmitters for fiber optics can fit into three simplified categories:

- Light-emitting diodes—LEDs are low-cost, but they are limited to lower bit rates that fall well below 1 Gbps.
- Long-wavelength lasers—These lasers operate at the 1310 nm and 1550 nm wavelengths, with much higher speeds and much higher cost.
- Short-wavelength lasers—Vertical Cavity Surface Emitting Laser (VCSEL) technology was developed as a means to manufacture a low-cost laser. VCSELs for 10 Gbps are currently classified as short-wavelength lasers and operate at 850 nm. The next development for 10 Gbps will be long-wavelength (1310 nm) VCSELs.

The 802.3ae standards use the following letters:

S = short wavelength at 850 nm
L = long wavelength at 1310 nm
E = long wavelength at 1550 nm

The letter "W" means the standard is for use with WANs

The three fiber categories, combined with short- and long-wavelength laser technology, give users several choices that have trade-offs between distance and cost.

- Multimode fiber (OM3)—OM3 fibers are classified as laser optimized, because the cable manufacturer must subject cables to 100% inspection as only the best fibers are good enough to be sorted as "laser-optimized." To be classified as OM3, a fiber must have an effective modal bandwidth greater than or equal to 2000 MHz/km. Differential mode delays (DMD) still occur with laser-optimized fibers, but the OM3 classification places a limit on DMD. As each pulse is transmitted through the fiber, its original shape is better preserved, and the receiver is able to decide whether the received pulse represents a zero or a one. The more advanced Bit Error Rate (BER) test actually measures the frequency of misinterpreted pulses. When deploying multimode fiber in a data center, OM3 fibers are highly recommended.
- Single-mode fiber (OS1)—The remaining option listed for 10 Gbps transmission is single-mode fiber. 10GBASE-LR and 10GBASE-ER are based on standard single-mode fiber and long-wavelength lasers, so the maximum distances are quite large in comparison to OM3. These technologies are more expensive because they use edge-emitting long-wavelength lasers. 10GBASE-LR and ER are typically used for long-distance applications across a campus or between sites.

Practical Deployment of Fiber in the Data Center

Business applications generate large amounts of data that must be placed in storage systems that are always accessible, secure, and extremely reliable. SANs satisfy these requirements, and by following a strategy of consolidation, the SAN fabric becomes a flexible, scalable resource that can adapt quickly to changing data requirements. SAN connectivity is supported by fiber optic cabling, so the cabling infrastructure should supply fiber ports to any equipment areas that house or will house applications requiring access to the SAN, whether today or in the future. A recommended strategy is to run a certain number of fiber pairs to every server cabinet—it could be 12, 24 or 48, according to the maximum projected need. This projected need depends on server form factor, power/heat loads, and projected need for SAN access. Likewise, fiber pairs must also be run from the SAN switches to the storage devices themselves.

During the initial build, distribution cables may be routed to each cabinet and terminated onsite. Field termination is intricate, time-consuming, and messy, so many data centers prefer to install pre-terminated trunk cables or plug-and-play fiber solutions. With plug-and-play solutions, the distribution cables are factory terminated with MPO or MTP connectors, typically 12 fibers per connector. The installer simply plugs the MTP cable into an LC or SC breakout cassette at the horizontal cross-connect and at the equipment outlet.

Cable Pathways

The cabling infrastructure should be viewed as an investment that requires functional, protective containment. It is no longer acceptable to lay data cables on the slab beneath the raised floor each time a new server cabinet is added to the network. IT professionals, who are responsible for growth and change, depend on cable pathways that are well-placed, flexible, and accessible. Mission-critical applications require that cables be protected from sudden damage or long-term degradation, and increasing heat loads demand careful consideration, so that pathways do not obstruct cooling. When planning cable pathways, there are no universal solutions. The designer must consider whether to place the pathways under the floor or overhead. The criteria in Table 9 dictate how cable pathways are implemented.

Table 9. Criteria for Cable Pathway Implementation




Criteria	Description
Architectural considerations	<ul style="list-style-type: none">• Use what you have—Raised floors, ceiling height, and structural integrity dictate how you design your cable pathways.• Organized and visible—Layer overhead pathways so there is clear separation between power and data cables.• Flexible—Allow for cabinets to be added and removed.• Maintain pathway-rows ratios—One pathway will feed two rows of equipment for below-floor pathways. For overhead pathways, typically one pathway feeds one row of equipment cabinets.
Capacity	<ul style="list-style-type: none">• Maximum 50 percent fill ratio—A calculated fill ratio of 50 percent will physically fill the entire tray due to spaces between cables and random placement.• Cable trays at 25 percent fill ratio—Maximum pathway fill should be 25 percent to leave room for future growth. This applies to cable tray, wire basket, ladder rack, and rigid cable tray.
Separation of power and data	<ul style="list-style-type: none">• UTP cables and unshielded power cables—100 to 600 mm (4 to 24 in.)• UTP cables and shielded power cables—50 to 300 mm (2 to 12 in.)
Accessibility and security	<ul style="list-style-type: none">• Balancing act—Protect cables from tampering or damage while allowing for changes to the cabling infrastructure.
Fire risk	<ul style="list-style-type: none">• Authority having jurisdiction—This authority has final say on acceptable practices from plenum-rated cabling to fire safety measures.
Fiber protection	<ul style="list-style-type: none">• 40 to 60 percent fill ratio—Copper and fiber have different characteristics, so different cable pathway strategies must be applied.

Rack and Cabinet Considerations for the Data Center

Criteria Used to Decide on a Rack

The data center designer has many choices when it comes to open racks and closed cabinets for mounting equipment and patch panels, each with advantages and disadvantages. As with pathways, there is no universal solution. Regarding width, 19 inches is the norm for servers, switches, and patch panels; the 23-inch width can be used for carrier equipment in the entrance room and any other equipment that is wider. The preferred height is 7 feet, and the maximum allowed height is 8 feet. Other topics such as seismic considerations, placement relative to tiles, adjustable rails, and power strips are also addressed in the standard. The designer should choose open racks or cabinets based on the criteria listed in Table 10.

Table 10. Types of Racks

Rack Type	Description	Maximum Weight Load	Applications
	2-post open rack	800 lb	<ul style="list-style-type: none">• Low-capacity switching/routers• Patching
	4-post open rack	1000-2000 lb	<ul style="list-style-type: none">• High capacity switching/routers• Servers
	Cabinet	2000 lb	<ul style="list-style-type: none">• Server clusters• Storage arrays• Storage switching

Cable Entry

With open racks, cables are easily fed from the overhead or under-floor pathways into the vertical cable managers. With cabinets, it is important to compare the quantity of cables to the size of the cable entry holes. Equipment with a high cable density, such as switches and one-rack-unit servers, can require cable entry holes large enough for hundreds of data cables, plus power cables. The necessary opening A is calculated using a 40 percent fill rate (0.4) as follows:

$$A = \frac{336 \times (\pi \frac{D^2}{4})}{0.4}$$

For Category 6, the diameter D is 0.25 inch and A is calculated to be 41 square inches. For Augmented Category 6, D is 0.35 inch and A is 81 square inches.

Cable Management

Racks can be easily equipped with spacious vertical cable managers that are easy to access. Many cabinets do not provide enough space for vertical cable management. Often the only space available is the narrow, vertical corner channels and the spaces between adjacent cabinets, which are difficult to access. The formula in the previous section applies not only to cable entry holes, but to cable management cross-sectional areas as well. Whenever patch panels are installed in cabinets, there should be a minimum of 100 mm (4 in.) depth between the rail and the door to provide room for cable management.

Security

Cabinets have doors that can be locked. Open racks do not, but they can be secured in a cage or a separate room. Data center managers are concerned about access to equipment by intruders as well as unauthorized personnel. For mission-critical applications, work done on equipment should be carefully controlled with strict access policies. For example, technicians in the networking group are given access to the switch and cross-connect cages, while those in the server group have keys to the server cabinets, and the SAN group has codes for the SAN cabinets. In colocation environments, renters expect the provider to provide secure cabinets and spaces so that other renters cannot tamper with their equipment and the supporting network infrastructure.

Load Capacity and Depth

Cabinets and four-post racks have an advantage, in terms of load capacity. For example, one manufacturer's standard 19-inch aluminum rack with 3-inch channels can hold 800 pounds. The same manufacturer's heavy-duty rack with 6-inch channels holds up to 1500 pounds. Four-post racks and cabinets can be rated for 2000 pounds or more. Cabinets can support deeper equipment. Most network switches are shallow enough to be mounted on two-post racks, depending on what is permitted by the installation instructions. Rack-optimized servers have depths ranging from 16 to 33 inches. Most fall between 26 and 30 inches and are generally not compatible with two-post racks because this would create a structurally unstable cantilever and a possible safety hazard. Cabinets and four-post racks provide rear mounting rails to increase support and surround deep equipment. Mounting rails should be adjustable to accommodate a range of equipment depths.

Pitch

In some cases, it is important to maintain a 24-inch pitch to match the floor panels. Matching the rack or cabinet to the floor panels makes it possible to have standard cable cutouts for all floor panels. Common cabinet widths are 24, 28, 29.5 (750 mm), and 32 inches. For server applications, particularly with long rows of cabinets, the 24-inch width is most common. Wider pitches are accepted for switching or patching applications, where management of large cable bundles is a priority. Racks with vertical cable managers exceed the 24-inch floor panel dimension.

Power Accommodations

Cabinets are designed to hold power strips in the vertical support channels. Long power strips with 20 or more plug outlets are a necessity when cabinets are packed with one- and two-rack-unit servers, using redundant power supplies. There is a perception that open racks do not support long vertical power strips. Depending on the width of the power strip or rack channel, it may be possible to mount a vertical power strip on the rear of an open rack, provided the power strip and its mounting brackets do not interfere with one another. Typically open racks are not designed to route power cables, nor do they have covers or doors to conceal the power strip.

Mounting Rails

Mounting rails should come with a hole pattern that complies with TIA-310-D. For racks, the holes are typically tapped, which facilitates installation of a large number of components such as patch panels. Cabinets, on the other hand, typically come with square holes for caged nuts. Equipment cannot be mounted until the caged nuts are snapped into the square holes. Caged nuts are available in various thread patterns including #12-24, #10-32, and M6. Rails should be marked to show rack unit boundaries, and numbering of the rack units facilitates installation.

Cable Management

Once the horizontal cables have been installed, they should be left alone—protected and undisturbed. From then on, any changes to the network configuration should be done by moving patch cords at the horizontal cross-connect. That patch field becomes the focal point for management of the physical layer, so it should be designed to provide long-term benefits such as manageability, reliability, security, and scalability. The eight best practices listed in Table 11 can elevate the functionality of the horizontal cross-connect beyond the level of a simple patch field to a sophisticated change-center with provisions for cable management, hot-swapability, cooling, and intelligent tracking of patch field changes.

Table 11. Cable Management Best Practices

Practice	Description	
Include sufficient horizontal and vertical cable management	<ul style="list-style-type: none">Vertical cable managers between racks must be at least 83 mm (3.25 in.) wide; 250 mm (10 in.) is recommended for rows having two or more racks.Vertical cable managers at the ends of a row of racks should be at least 150 mm (6 in.) wide.	
Provide bend radius control wherever cables turn corners	<ul style="list-style-type: none">Slack managers and transitions into pathways should also be designed with the proper bend radius.Cables should be guided into the horizontal and vertical cable managers by fingers that are engineered with radiused edges.	
Make the most of the space available	<ul style="list-style-type: none">High-density solutions like angled patch panels and vertical cable management with matched fingers can fit more connections into a smaller footprint.	
Protect critical infrastructure	<ul style="list-style-type: none">Create different levels of security with cabinet locks and cages.	
Route cables to allow hot-swapability	<ul style="list-style-type: none">When routing cables away from the equipment, consider how the fan assemblies, modules, and power supplies are removed and inserted when it comes time for upgrades or replacements.	
Respect airflow patterns	<ul style="list-style-type: none">Place your patch fields in alternating pattern with your switching infrastructure to maximize airflow.	
Document and manage changes to the physical lay	<ul style="list-style-type: none">The patch field contains hundreds or thousands of ports, so it is essential that the patch field be labeled to allow technicians to quickly identify what each port represents.	
Weigh the tradeoffs for interconnection versus cross-connection	Reasons to use interconnection: <ul style="list-style-type: none">Less spaceFewer connections, thus lower insertion lossLower upfront costEasier to trace	Reasons to use cross-connection: <ul style="list-style-type: none">Less possibility of damaging the switchOnly choice for switch cabinetsMore flexibilityCompatible with physical-layer management

Grounding Considerations for the Data Center

The grounding system is not just an insurance policy against a lightning strike. It is an active, functioning system that provides protection for personnel and equipment. Proper grounding is essential for efficient system performance. Surges that are not properly dissipated by the grounding system introduce electrical noise on data cables. They cause faulty data signals and dropped packets, thus decreasing the throughput and overall efficiency of your network.

According to insurance industry data, improper grounding of communication systems leads to \$500 million per year in lightning damage to property and equipment. The Information Technology Industry Council states that grounding is the most important factor in reliable network equipment performance. The component cost of repairing damaged equipment is substantial for complex circuit boards, especially when labor and downtime are considered.

According to the IEEE, the typical AC third-prong ground is almost never sufficient to prevent damage to network equipment. Personal injury from electric shock due to improper grounding can cause immeasurable human suffering and significant expense. Potential fire hazards exist when heat is generated from electrical surges that find a high-resistance path to ground.

Standards for Data Center Grounding

Data center grounding is governed by following documents:

- TIA-942, Telecommunications Infrastructure Standard for Data Centers—TIA-942 defines practical methods to ensure electrical continuity throughout the rack materials and proper grounding of racks and rack-mounted equipment. This is the only specification that addresses problems specific to data centers.
- J-STD-607-A-2002, Commercial Building Grounding (Earthing) and Bonding Requirements for Telecommunications—This standard focuses on grounding for telecommunications. It defines a system that begins at the entrance facility, in the telecommunications main grounding busbar (the TMGB), and ends at the local telecommunications grounding busbars (TGBs) located in the telecommunications rooms.
- IEEE Std 1100 (IEEE Emerald Book), IEEE Recommended Practice for Powering and Grounding Electronic Equipment—IEEE provides further detail on how to design the grounding structure for a computer room environment through a Common Bonding Network (CBN). The CBN is the set of metallic components that are intentionally or incidentally interconnected to provide the principal bonding and grounding inside a telecommunications building. These components include structural steel or reinforcing rods, metallic plumbing, AC power conduit, cable racks, and bonding conductors. The CBN is connected to the exterior grounding electrode system.

Characteristics of the Data Center Grounding System

The purpose of the grounding system is to create a low-impedance path to earth ground for electrical surges and transient voltages. Lightning, fault currents, circuit switching (motors turning on and off), and electrostatic discharge are the common causes of these surges and transient voltages. An effective grounding system minimizes the detrimental effects of these surges. A properly designed grounding system has the following characteristics:

- It should be intentional; that is, each connection must be engineered properly. The grounding system is no more reliable than its weakest link.
- It should be visually verifiable.
- It should be adequately sized.
- It should direct damaging currents away from equipment.
- All metallic components in the data center should be bonded to the grounding system (Figure 2).

Along with these characteristics, all grounding conductors should be copper, components should be listed by an approved test lab such as UL, and local electrical codes must be adhered to.

To ensure long-term integrity of the grounding system, always use compression connectors, not mechanical. A mechanical connector holds the conductor in place with a setscrew, and when exposed to vibration (e.g., nearby fans or humming equipment), the setscrew can back off and loosen. A loose connection is a high-resistance connection that can fail in a surge event. A compression connector is permanently deformed and does not come loose with vibration.

There should be a logical flow as you follow the grounding structure, for example, from the equipment chassis to the rack, from the rack to the data center grounding infrastructure, then over to the local TGB, which feeds into the telecommunications bonding backbone (TBB) that runs back to the main TGB (TMGB), which is tied to earth ground. Table 12 summarizes the important points of this subject.

Figure 2. Typical Data Center Grounding




Table 12. Key Points of Data Center Grounding

Key Point	Description
Equipment chassis	<ul style="list-style-type: none"> • Third prong is never sufficient. • Always follow the grounding/earthing recommendations of the manufacturer when installing equipment.
Rack/cabinet continuity	<ul style="list-style-type: none"> • Hardware typically supplied with bolt-together racks is not designed for grounding/earthing purposes. • Racks should be assembled with paint-piercing grounding washers, under the head of the bolt and between the nut and rack, to provide electrical continuity. • A full-length rack grounding strip should be attached to the rear of the side rail with thread-forming screws to ensure metal-to-metal contact.
Rack/cabinet grounding	<ul style="list-style-type: none"> • Use a #6 AWG or larger bonding conductor to bond each rack or cabinet with the grounding strip to the data center grounding infrastructure. • Do not bond racks or cabinets serially.
Data center grounding infrastructure	<ul style="list-style-type: none"> • A common method for constructing the data center grounding infrastructure is to create a copper conductor grid on 0.6 to 3 m (2 to 10 ft) centers that covers the entire computer room space.
Telecommunications grounding bar	<ul style="list-style-type: none"> • Use a #1 AWG or larger conductor to bond the data center grounding infrastructure to the TGB. • Two-hole copper compression lugs are preferred because they are irreversible and resist loosening when twisted (bumped) or exposed to vibration.
Telecommunications bonding bar	<ul style="list-style-type: none"> • The TBB should be installed as a continuous conductor, avoiding splices where possible. • Avoid routing grounding/earthing conductors in metal conduits. • Although the building steel and metallic water piping must be bonded to the grounding/earthing system for safety reasons, neither may be substituted for the TBB.
Telecommunication main grounding busbar	<ul style="list-style-type: none"> • The TMGB is bonded to the service equipment (power) ground, which connects to earth ground (the grounding electrode system).

Routine Inspection of the Grounding System

Mission-critical facilities should implement a plan to inspect all points along the grounding infrastructure on an annual or semiannual basis. An inspection that follows a line-by-line work order allows early detection of potential problems such as loosened or corroded connections, missing labels, conductors that have been damaged, cut, or removed, and new metallic elements that require connections to the CBN. To facilitate inspection, the grounding system should use connectors, busbars, and conductors from end to end that allow visual verification of the bond.

Other Considerations

Physical Security

Physical security—controlling personnel access to facilities—is critical for high data center availability. As new technologies such as biometric identification and remote management of security data become more widely available, traditional card-and-guard security is being supplanted by security systems that can provide positive identification and tracking of human activity in and around the data center. Before investing in equipment, IT managers must carefully evaluate their specific security needs and determine the most appropriate and cost-effective security measures for their facility.

The first step in a security plan is drawing a map of the physical facility and identifying the areas and entry points that need different rules of access, or levels of security. These areas might have concentric boundaries or side-by-side boundaries. For instance, the computer area would be shown as a square within a larger area, such as the building perimeter. Examples of side-by-side boundaries include visitor areas, offices, and utility rooms. Concentric areas can have different or increasingly stringent access methods, providing added protection called “depth of security.” With depth of security, an inner area is protected by both its own access methods and those of the areas that enclose it. In addition, any breach of an outer area can be met with another access challenge at a perimeter further in.

Once the security plan is mapped, the next step is to create the access criteria. A person’s authority for access to a secure area is based mainly on identity, purpose, and the need to know, plus any criteria specific to an organization. Methods of identifying people fall into three general categories of increasing reliability—and cost:

- *What you have* is something you wear or carry—a key, a card, or a token that can be worn or attached to a key ring. It can be as “dumb” as a brass door key or as “smart” as a card having an onboard processor that exchanges information with a reader (a smart card). This is the least reliable form of identification, because there is no guarantee it is being used by the correct person—it can be shared, stolen, or lost and found.
- *What you know* is a password, code, or procedure for something such as opening a coded lock, verification at a card reader, or keyboard access to a computer. A password/code presents a security dilemma: if it’s easy to remember, it will likely be easy to guess; if it’s hard to remember, it will likely be hard to guess—but also likely to be written down, reducing its security. This is more reliable than what you have, but passwords and codes can still be shared, and if written down they carry the risk of discovery.
- *Who you are* refers to identification by recognition of unique physical characteristics—this is the natural way people identify one another with nearly total certainty. Biometric scanning techniques have been developed for a number of human features including fingerprint, hand, iris, and face. Biometric devices are generally very reliable, if recognition is achieved—that is, if the device recognizes you, then it almost certainly is you. The main source of unreliability for biometrics is not incorrect recognition or spoofing by an impostor, but the possibility that a legitimate user may fail to be recognized (false rejection).

The final step involves selecting the optimal security scheme. A typical security scheme uses methods of increasing reliability—and expense—in progressing from the outermost (least sensitive) areas to the innermost (most sensitive) areas. For example, entry into the building might require a combination of swipe card plus PIN; entry to the computer room might require a keypad code plus a biometric. Combining methods at an entry point increases reliability at that point; using different methods for each level significantly increases security at inner levels, since each is secured by its own methods plus those of outer levels that must be entered first.

Technologies are in place, and getting less expensive, to implement solutions based on these three identification categories. By combining an assessment of risk tolerance with an analysis of access requirements and available technologies, an effective security system can be designed to provide a realistic balance of protection and cost.

Fire Suppression

According to the U.S. National Fire Protection Association (NFPA), there were 125,000 nonresidential fires in 2001 with a total of \$3.231 billion in losses. Industry studies show that 43 percent of businesses that are closed by a fire

never reopen, and 29 percent of those that do open fail within 3 years. Fires in data centers are typically caused by power problems in raceways, raised floors, and other concealed areas—one reason why raised floors are not recommended for data centers and network rooms. Fires have also been caused by arson and corporate sabotage and electrical events, such as lightning and power surges. Like any other critical system in a data center, fire protection must be redundant and fault tolerant.

Fire prevention provides more protection than any type of fire detection or suppression equipment—if an environment is incapable of breeding a fire, there is no threat of fire damage. If a fire does occur the next step is to detect it. Advanced detectors can detect fire in its incipient stages and notify a central control center that in turn notifies personnel and suppression systems. In a data center, the main goal of the fire protection system is to get the fire under control without disrupting the flow of business and without threatening the safety of personnel.

A fire protection solution for a data center has three primary goals: (1) identify the presence of a fire, (2) communicate the existence of that fire to the occupants and proper authorities, and (3) contain the fire and extinguish it if possible.

The following fire detection/suppression elements are required to meet a data center goal of 24-hour uptime:

- Linear heat detection (heat-sensing cable)
- Intelligent spot-type detection
- Air sampling smoke detection
- Portable fire extinguishers
- Total flooding clean-agent fire suppression system
- Pull stations, signaling devices, and control system

Linear heat detection cable should be placed along all wire trays and electrical pathways above and below the raised floor. As with all fire-related sensors, an alarm here should not directly trigger the suppression system—rather, it should prompt the control system to sound an alarm. To further safeguard against accidental discharge of the fire-suppression agent, air-sampling smoke detector systems should be placed beneath as well as above a raised floor—and both detection systems must enter an alarm state before the total flooding suppression system discharges. It is also recommended that intelligent spot-type detectors be positioned at every CRAC unit intake and exhaust. If any other ductwork enters the data center, duct smoke detectors should be installed. The control system should be able to manage all these sensors and ensure that no single alarm will trigger a discharge.

Many Halon-alternative, clean agent systems are available that offer a variety of trade-offs in equipment space, ceiling height constraints, distance of sprinklers from the tank, and environmental friendliness. For example, FM-200 is favored for its small storage footprint. Other commonly used fire-suppression systems, each with its own advantages, are INERGEN and Novec.

In addition to the total flooding fire-extinguishing system, state or local fire codes may require a sprinkler system. If this is the case, it must be a pre-action system to prevent accidental water damage to the data center. FE-36 clean agent fire extinguishers should be placed throughout the data center, according to local fire codes. There should be pull stations and written procedures at every exit, and signaling devices throughout the building capable of notifying all personnel of a fire.

The control system is vital to the effectiveness of the suppression system. It should be fault tolerant, programmable, and capable of monitoring all devices. It should also allow manual override of automatic operation. All detectors should be addressable to allow the control panel to identify the precise location of any alarm. The control system must coordinate the sequence of events after the initial alarm: the sounding of a separate evacuation alarm prior to discharge, closing ventilation dampers to prevent air from escaping, discharging the agent, and notifying the local authorities. The control system must be supported by well-written and effective emergency procedures, reinforced with regular training of all data center employees.

The best way to protect a data center from fire is to implement best practices to prevent it, some of which are listed below:

- Ensure that the data center is built far from any other buildings that may pose a fire threat.
- Free all electrical panels of any obstructions.
- Enforce a strict no-smoking policy in IT and control rooms.
- Keep the data center free of any trash receptacles.
- Ensure that all office furniture in the data center is constructed of metal (except chairs may have seat cushions).
- Keep essential supplies such as paper, disks, and wire ties completely enclosed in metal cabinets.
- Isolate tape storage libraries (burning tapes generate dangerous fumes).
- Do not exceed 15 feet in length for UL-approved extension cords used to connect computer equipment to branch circuits, and avoid running power cords under equipment, mats, or other covering.
- Do not use acoustical materials, such as foam or fabric, in a data center.
- Do not permit air ducts from other parts of the building to pass through the data center—if this is not possible, use fire dampers to prevent fire from spreading to the data center.
- Use dry type transformers (or those filled with noncombustible dielectric) in the data center.
- Protect all cables passing through the raised floor against chafing by installing edge trim around all openings.
- Separate computer areas from other rooms in the building by fire-resistant rated construction extending from the structural floor slab to the structural floor above (or roof).
- Avoid locating computer rooms adjacent to areas where hazardous processes take place.
- Train all data center personnel in the operation of fire protection systems and extinguishers.

Most fires in mission-critical facilities can be prevented if common mistakes are avoided and fire detection is properly designed and monitored. Human error plays a large part in fire hazards and must be eliminated through training and strictly enforced procedures.

Conclusion

Power, cooling, cabling, racks, and other physical elements are the essential foundation of any data center. In the Cisco Enterprise Data Center Network Architecture, these facilities must be strategically designed to provide the critical reliability, agility, and efficiency needed by on-demand IT operations. New innovations in standardization and modularity for data center facilities, combined with a growing body of experience and research in this field, stand ready to provide effective strategies to guide the design of mission-critical facilities for the data center.

The facilities design information in this paper is the result of more than six years of research by APC and Panduit into the challenges that arise in the data center between IT and facilities. The selection, organization, and presentation of this information was developed in collaboration with Cisco's Data Center Network Architecture group based in San Jose, California.

References

For more about the topics in this application note, see the following APC white papers, available at www.apc.com:

White Paper No.	Title
--------------------------------	--------------

Financial considerations in data center facilities

117	Network-Critical Physical Infrastructure: Optimizing Business Value
116	Standardization and Modularity in Network-Critical Physical Infrastructure
37	Avoiding Costs from Oversizing Data Center and Network Room Infrastructure
6	Determining Total Cost of Ownership for Data Center and Network Room Infrastructure

Power and UPS

75	Comparing UPS System Design Configurations
3	Calculating Total Power Requirements for Data Centers
113	Electrical Efficiency Modeling for Data Centers

Cooling

56	How and Why Mission Critical Cooling Systems Differ from Common Air Conditioners
55	Air Distribution Architecture for Mission Critical Facilities
25	Calculating Total Cooling Requirements for Data Centers
49	Avoidable Mistakes that Compromise Cooling Performance in Data Centers and Network Rooms
42	Ten Steps to Solving Cooling Problems Caused by High-Density Server Deployment
44	Improving Rack Cooling Performance Using Blanking Panels

Physical security

82	Physical Security in Mission Critical Facilities
----	--

Fire suppression

83	Mitigating Fire Risks in Mission Critical Facilities
----	--

Management

100	Management Strategy for Network-Critical Physical Infrastructure
-----	--

The following **Panduit** white papers are available at www.panduit.com

Title

Cabling topologies

Bringing Manageability to the Data Center

Cabling media

Consortium White Paper Details Category 6 Standard
Certifying Multimode Fiber Channel Links for 10 Gig
Gigabit Ethernet on Copper Cabling Structures
Evolution of Copper Cabling Systems

Cabling management

Keep Those Cables Under Control

Grounding

Minimize Service Interruptions - Power Connectors

Cisco, Cisco Systems, and the Cisco Systems logo are registered trademarks or trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.