CORNING

Data Center Design Guide

Contents

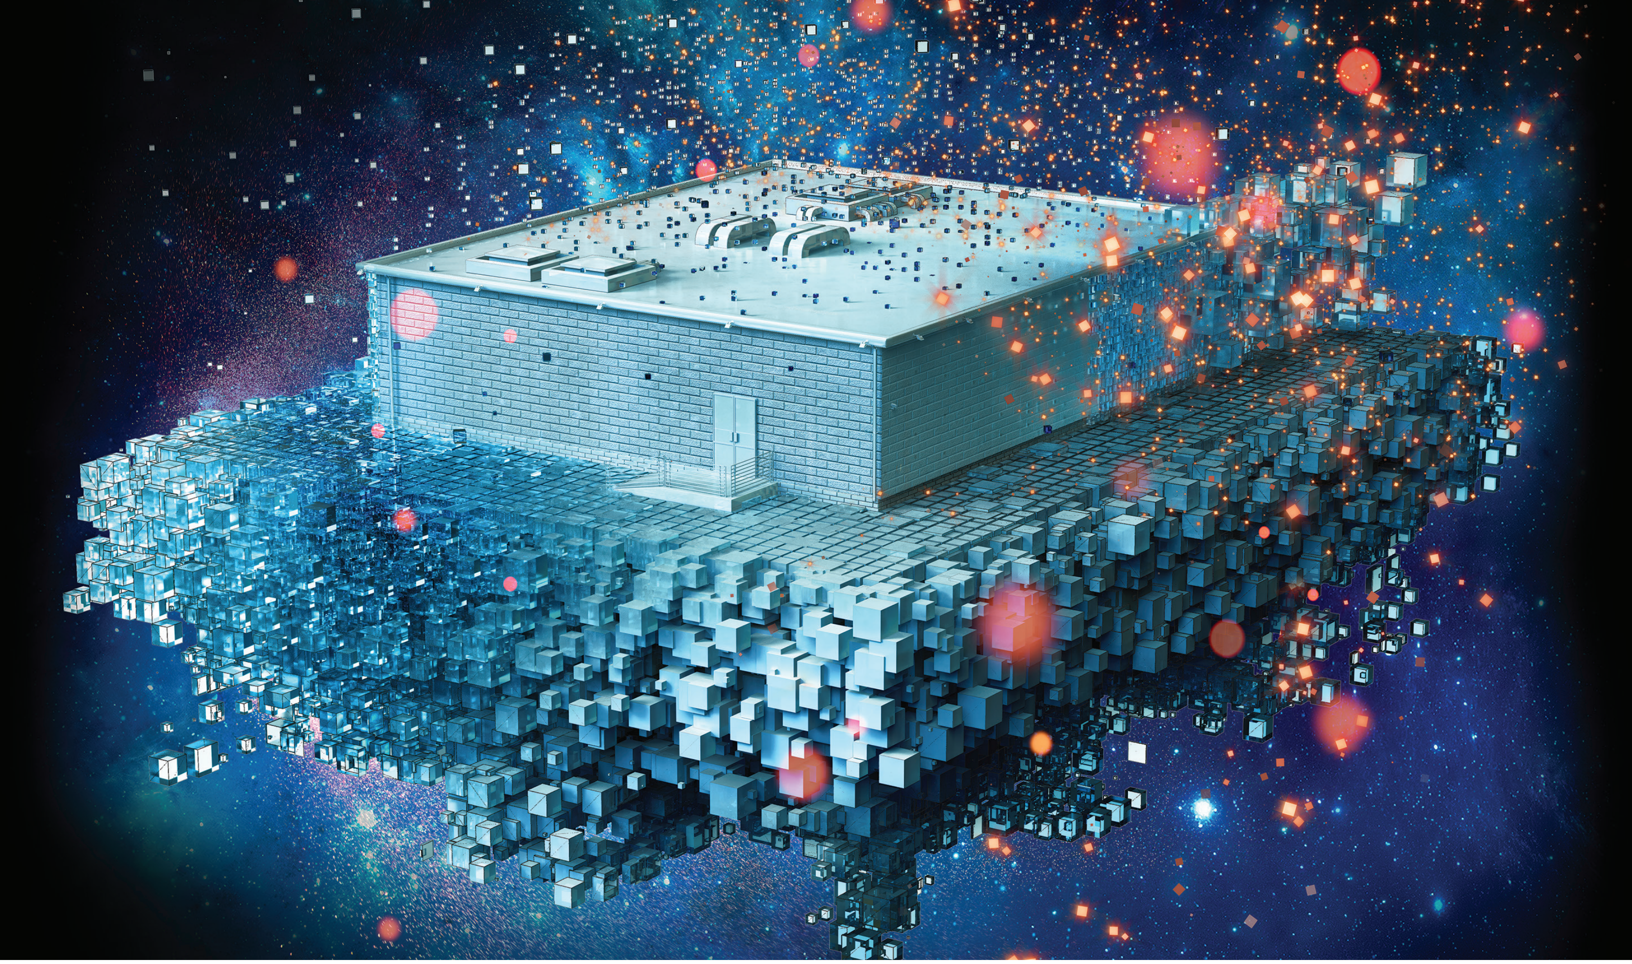

Welcome

The data center industry is evolving at a tremendous pace. This guide is designed for businesses and organizations that are considering building or leasing data center space for the first time or making a change in their data center models.

The information in the guide will help you understand what data centers are, the different types of data centers in the marketplace, and the pros and cons of each. We'll also take a look at the various functional areas of a data center, highlight a few key products, and provide tools and resources that will help ease your discovery and decision-making processes.

As always, our data center experts are here to help you with planning, design, product selection, training, deployment, and trouble-shooting. Please reach out to your Corning representative with any questions.

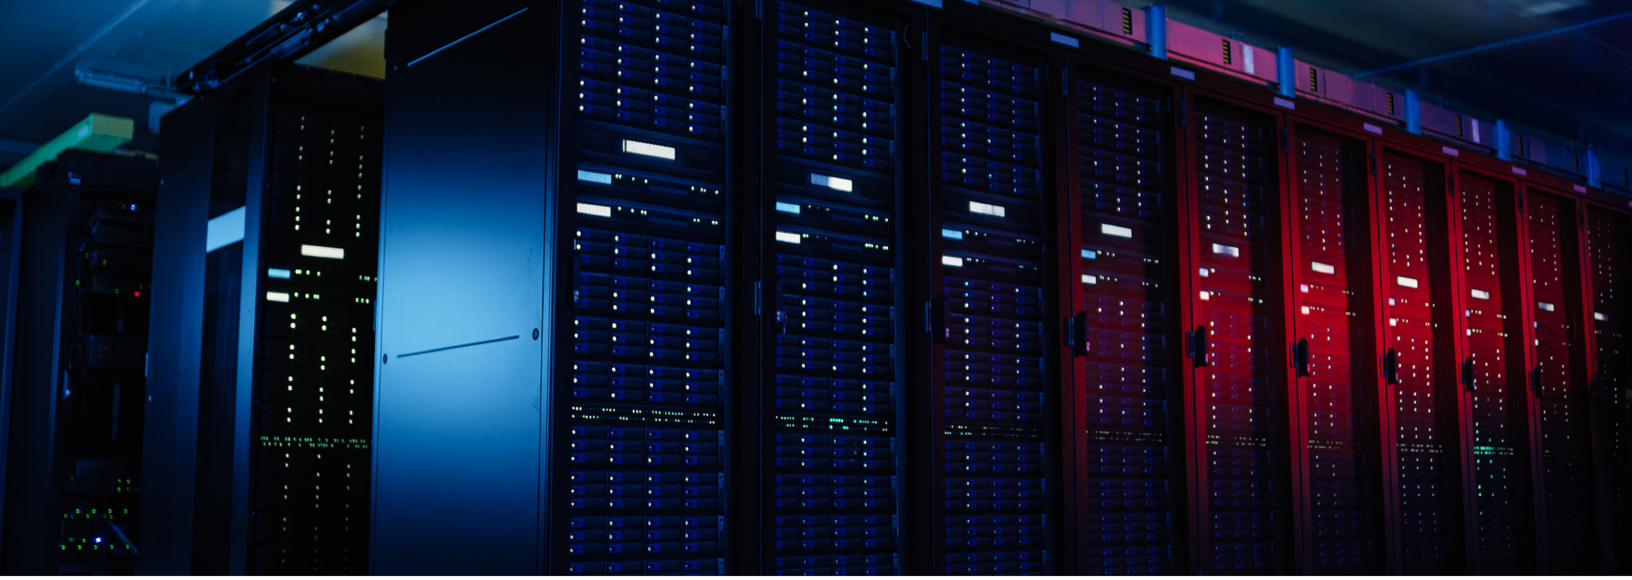

The Evolution of Data Centers

A data center's main functions are to consolidate, manage, and disseminate information technology (IT) or information communication technology (ICT) resources. This includes housing and supporting a network for data processing and operational hosting that provides services such as application and internet access, content distribution, database management, file storage, and data backup.

In addition to these traditional services, today's data centers also enable search technologies, social networking, mobile applications, and e-commerce. They're able to do this by providing e-business services, aka "cloud computing." In the early 2000s, these services focused on Infrastructure-, Platform-, and Software-as-a-Service (IaaS, PaaS, and SaaS). They've since expanded to include the fast-growing AI-as-a-Service (XaaS) sector.

e-Business Services Enabled by Data Centers

IaaS, **PaaS**, **SaaS**, and **XaaS** operations provide a wide gamut of services, including applications, data management, runtime optimization, middleware, O/S, virtualization, server hosting, storage, complete networking, AI, and generative AI. Each e-business type requires a different level of control and management by the customer.

Infrastructure as a Service (IaaS)

Administrators & Providers

Users are able to rent storage and computing capabilities from online service providers who enable the resources from their servers.

Platform as a Service (PaaS)

Developers & Integrators

This type of cloud computing allows businesses and organizations to rent a platform for software development without having to maintain the servers needed to enable it.

Software as a Service (SaaS)

End Users & Clients

This cloud-based application delivery model is also called on-demand software. Users can run SaaS programs from the provider's servers without having to manage and maintain the software.

AI as a Service (XaaS)

Output to the Customer

There is no maintenance associated with this service. The operator simply provides the program, on their terms, to the customer.

More — **Customer Control & Management** — Less

Data Center Design Consideration Checklists

As technology has evolved over the decades, data centers have as well—branching into three main specialized types in the marketplace today: **Enterprise**, **Multitenant (MTDC) or Colocations**, and **Hyperscale**. Which type you'll need to build, own, or lease will be determined by your space and data capacity requirements. Learn more about the **build process**, including the many stakeholders you'll want to consider.

Choosing Your Infrastructure

Manageability

Allows growth and evolution while maintaining order and functionality.

Flexibility

Enables quick MACs, minimizes operational impact, and slows obsolescence.

Scalability

Plans for migration to higher data rates and capacity.

Reliability & Availability

Ensures stability, required uptime, and resiliency for the unexpected.

Learn more **here**.

Top 10 Data Center Requirements

Fiber optic cable infrastructure	Power	Cooling	Security	Sustainability
Physical space & infrastructure	Energy efficiency	Code compliance	Network monitoring & management	Skilled staffing

See **More Tips for Data Center Success** for more information.

Data Center Build Process

Following a comprehensive approach to facility planning, design, and construction enables data center owners to effectively support the computing and storage needs of businesses while ensuring scalability, efficiency, and security.*

Planning & Feasibility Phase

- Key Design Criteria
- Computer Room, Facility, & IT Survey
- CFD Baseline Modeling
- Support & IT Infrastructure Assessment
- Program Requirements
- Alternate Approach Review
- Financial Analysis
- Operation & Maintenance Review
- Conceptual Design & Budget

Engineering & Design Phase

Design development documents.

- Architecture & Structural Engineering
- M/E/P Engineering
- Fire Protection & Alarm Design
- Security & Access Control Design
- IT Systems Design
- Structured Cabling Design
- CFD Modeling

Engineering Administration Phase

Design documents for bid, permit, and construction.

- Construction Award Administration
- Bid Leveling
- Bid Document Review
- Long Lead Item Procurement
- Written Specifications

Construction Administration Phase

- Permits
- 100% Construction Documents
- Shop Drawings Review
- Value Engineering
- Change Orders

Construction Phase

- Construction
- Data Cabling
- Punch List
- Clean Up
- As-Built Drawings

Commissioning Phase

- Commission Report
- Integrated System Testing
- Commission Plan
- Equipment Start-Up
- Cabling Certification

IT Implementation Phase

- IT & Support Infrastructure Monitoring
- Testing
- Installation

Ongoing Monitoring & Maintenance

*See **More Tips for Data Center Success** for more information.

Private Enterprise

Enterprise data centers are privately owned and operated by companies, institutions (e.g., universities), governments, or other businesses and organizations to provide internal data transactions and processing and web-based services through either intranets or extranets. They are supported and managed by internal IT support teams. A private data center can be located on the same premise as the enterprise it serves; however, some enterprises choose a different location based on security or power resources. Private edge networks are also an option for businesses that could benefit from a decentralized solution. (See below for details.)

Pros and Cons of Private Enterprise Data Centers

Enterprises that benefit from private data centers generally have unique or highly specialized network needs and can take advantage of economies of scale. As their workloads increased, many enterprises that initially operated a private solution chose to switch to leased, off-site space at an **MTDC**.

Pros	Cons
Dedicated computing system customized to specific needs.	Scalability limited by physical space and resources.
Full control (data sovereignty) over infrastructure, operations, security protocols, monitoring, and compliance.	Requires on-site hardware, aka "white space," like racks, servers, storage, and network.
Extremely low latency for time-critical operations.	Requires on-site IT staff.
Minimal dependence on external networks or internet connectivity.	Electrical and mechanical operations costs may be higher than at an outsourced data center.

Private Edge Data Centers

Edge data centers are networks of smaller, decentralized facilities that provide essential computing and storage capabilities in proximity to where data is generated. Running processes locally ensures low latency, enabling better performance and more efficient task execution. Acting as hubs for multiple network and service providers, edge data centers can offer localized computing resources that are particularly beneficial for cloud-driven functions, like edge computing and machine learning.

Multitenant Data Centers (MTDC)

A multitenant data center (MTDC), or colocation (colo), is building or campus designed to host data center equipment and services. It is connected to its customers through an internet service provider. How an MTDC space is equipped and operated depends on customer needs that may vary widely. For example, some customers may choose to simply lease empty space and provide their own equipment, some may choose to lease stock space and services all provided by the MTDC, while others may have specific custom requirements for either equipment or services or both.

How much space a customer leases depends on their needs and how those evolve over time. A business might start by leasing a single stock server rack, grow to a complete, purpose-built cage, and later fill an entire data hall. Along with space and equipment needs, costs are an important factor in determining if and how customers use MTDCs. For example, a business using an MTDC may not want the upfront costs of owning a private data center, yet want more control than “the cloud” can offer.

Pros and Cons of Multitenant Data Centers

Enterprises that benefit from MTDCs are those looking to reduce capital expenditure and operational costs while having reliable infrastructure, reduced maintenance, and the flexibility to scale. Some larger businesses that have a private data center also lease MTDC space to cover disaster recovery scenarios and predictable peak usage volumes that could overload their primary system.

Pros	Cons
No owned equipment, no maintenance.	Less control than a private data center.
Infrastructure frequently updated by MTDC owner to maintain competitive position.	Requires adherence to lease terms.
IT staff provided by MTDC.	
High level of flexibility to scale up or down as necessary.	

4 Types of MTDCs

Retail Colocation (< 250kW)

For small- and medium-sized businesses with workloads and applications that require a smaller footprint. Retail colocations accommodate hundreds of customers in cabinets and cages with shared cooling and extensive fiber optic interconnectivity to public and private clouds, internet service providers, data center backhauls, internet exchange services, etc.

Wholesale Colocation (> 250kW to 15MW)

For medium and large businesses with high data space, computing, and power requirements. Depending on their specific needs, these businesses may lease a caged white space, a data hall, or the entire data center.

Hyperscale (15MW – 1GW)

Hyperscale data centers are designed to scale with demand to support businesses with the largest data needs and workloads.

AI Leased (15MW – 1GW)

For businesses that require utilities and components suitable for the massive data and workload needs of AI networks.

Although it's not called an MTDC, a **public cloud** is in fact a leased multitenant environment that serves customers of all sizes. Computing resources "in the cloud" are shared by individuals, businesses, and other entities. The data is secured by state-of-the-art encryption and accessed via the internet alone. Tenants have no physical access to the cloud or control over what infrastructure its owner chooses.

Hyperscale

As their name suggests, hyperscales are very large data centers that have network architectures designed to scale in response to increasing workload demands. There are two types of scaling: vertical and horizontal. Vertical scaling adds more processing power to existing machines, whereas horizontal scaling adds to the number of machines in the computing environment or application spaces. Housed in a single building, on a campus of interconnected buildings, and even supplemented by MTDC leased space, these facilities may have up to several hundred thousand servers to support a wide variety of external customers and clients, or they may be dedicated to a single enterprise, where scalability is just as critical as capacity.

Pros and Cons of Hyperscale

Enterprises that benefit from hyperscale data centers generally need to accommodate exponential data growth.

Pros	Cons
Precise customization and control.	Large upfront costs.
Extreme scalability.	Massive power needs.
	Specialized components.
	Specialized labor.

2 Types of Hyperscale Data Centers

Hyperscale data centers are either leased or owned. Having these two options offers greater flexibility for companies that require what hyperscale data centers are designed to provide: the ability to scale quickly for growth, vast data storage and computing capabilities, and space to operate their Central Processing Units (CPUs), air cooling, and general-purpose servers.

Leased Hyperscale Data Centers

Two types of customers lease hyperscale space: 1) tenants who can't afford to own or don't prefer to operate their own hyperscale data center but still require enough space for the large number of GPUs necessary to enable high-capacity processing for AI, and 2) tenants, generally dominant e-commerce, web service, and cloud service providers, that need leased space in addition to their owned data centers to accommodate rapid growth as they evolve to assimilate AI into their businesses.

Owned Hyperscale Data Centers

Businesses that need data centers designed specifically for AI generally own their own hyperscale data centers. To support AI's complex and intense computations, these hyperscales rely on cutting-edge Graphic Processing Unit (GPU) clusters that require 10x more optical fiber than traditional data center networks. Accommodating these additional connections demands extreme density enabled by innovative **cables**, **structured cabling strategies**, Very Small Form Factor (VSFF) **connectors**, and upcoming **co-packaged optics**. Along with component miniaturization, AI hyperscales require massive amounts of bandwidth and power, plus liquid cooling and highly specialized design, build, and maintenance teams.

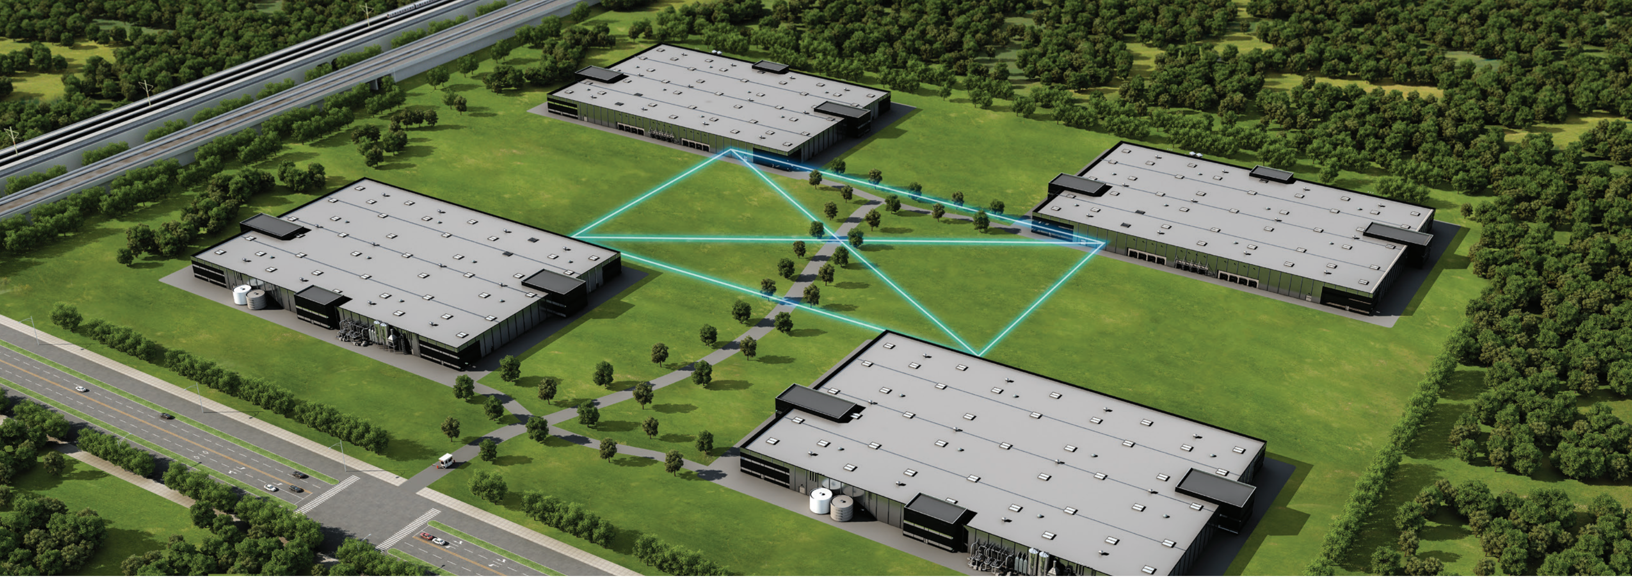

Data Center Functional Areas

Within a data center there are three main functional areas: Meet-Me Room, Main Distribution Area, and Caged Area. Additional areas, the Storage Area and Interconnect, also play key roles where they are utilized. See **TIA-942-C** for detailed definitions of every area and sub-area of a data center.

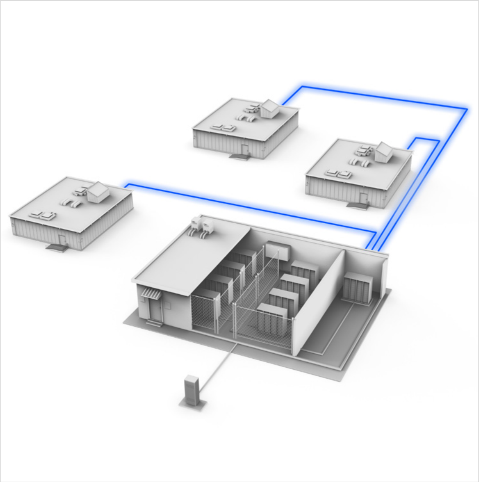

Data Center Interconnect (DCI)

Data Center Interconnect (DCI)

The Data Center Interconnect (DCI) links multiple data center campuses, enabling them to function as a cohesive network to support cloud computing, data redundancy, disaster recovery, and global operations. Although the DCI is not considered a main functional area, due to businesses' increasing reliance on digital services, it has emerged as a pivotal technology in the digital infrastructure landscape. Private, MTDC, and hyperscale data centers have begun adopting DCI solutions to catalyze growth, enhance scalability and redundancy while expanding service offerings — all while facilitating seamless connectivity. Corning has developed a suite of innovative DCI solutions, including small-form-factor fibers and preterminated, high-fiber-count cables that minimize duct size, reduce deployment complexity, and increase speed to market.

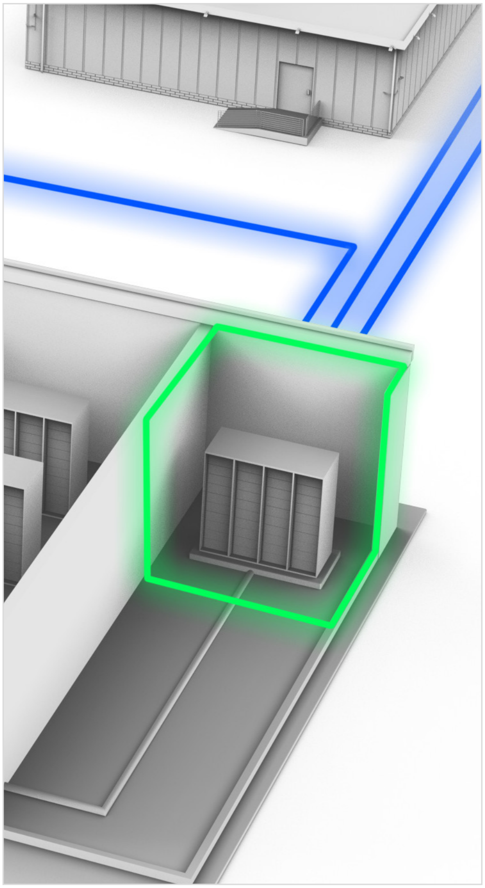

Meet-Me Room

Meet-Me Room

The central hub of a data center can be referred to by several names: Meet-Me Room (MMR), Point of Entry, Point of Presence (PoP), or Entrance Access Zone. Generally, it is a secure space where fibers and cables from data center racks come together. More specifically, it is a physical Entrance Room (ER) where the service provider's interbuilding cabling connects to the customer-owned intrabuilding cabling, enabling easy passage to the provider's telecommunications network.

The MMR may contain the provider's demarcation hardware and equipment, including internet gateways and routers, as well as traditional PBX or voice gateways, and T-3 multiplexers. For this reason, this zone is historically associated with the "core" layer of the network architecture. Data centers may have MMRs for redundancy purposes or to accommodate access for provider-provisioned circuit cable length restrictions. This zone is often outside the main computer room for administration and security purposes, however, in smaller data centers, the MMR may be in the same space as the main computer room or distribution areas.

The MMR is primarily used for multitenant data centers. It enables an efficient point-of-customer connection to ISPs, cable companies, carriers, and other tenants in a data center to exchange data between systems in different networks. This creates a seamless framework for connectivity at a reduced cost. Data is secure in the facility – mitigating the risk of intrusion, due to MMRs being surrounded by secure walls and around-the-clock monitoring.

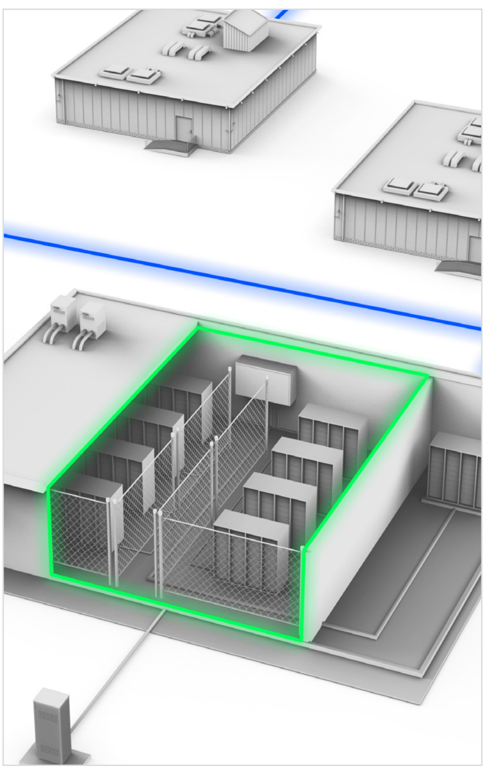

Main Distribution Area

Main Distribution Area

The Meet-Me Room interfaces with the rest of the data center, also known as the "computer room," through the Main Distribution Area (MDA). The MDA may be adjacent to or part of the MMR, depending on the size of the data center. Its infrastructure determines the ability to cross-connect and how quickly customers can be brought online. The MDA's main function is to provide the front-end connection to the primary functional components of the data center, the servers. Server types connected to the MDA may include: application, catalog, database, gaming, file, mail, media, virtual, and web.

Network equipment located in the MDA consists of core routers, LAN and SAN switches (including **"spine" and "interconnection" or "leaf" switches**), high-performance computing (HPC) switches, and software-based virtual machines (VMs) or hypervisors for server virtualization. The MDA also serves as the main cross-connect (MC), which is the central point of distribution for the facility's cabling infrastructure. The MC may also be known as the Central or Main Patching Location. If necessary, the MDA can connect to lower-layer distribution areas as well. This creates a hierarchy designed to efficiently and effectively establish connectivity between all the data center components.

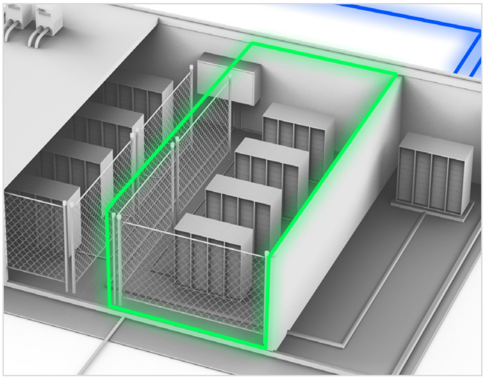

Caged Area

Caged Area

A caged area, also known as “white space”, is a physically secured (literally “caged”), restricted-access zone within a data center used to manage access to sensitive infrastructure. The individual cages can be sized to meet customer needs. The caged area is designed to house critical equipment, including servers, network devices, and storage systems and is used for performance-critical applications that require low latency and high throughput. It functions like a dedicated "storage highway" connecting servers to the Storage Area Network (SAN), a high-speed network that connects servers to storage devices, centralizing block storage in a shared pool.

Storage Area Network

Storage Area Network

Primarily used by private and hyperscale data centers, this area provides a back-end connection, transferring data between computer systems and storage devices and among storage devices. Because SANs offer administrative flexibility, consolidation, management, and scalability advantages, plus access, back-up, movement, automation, virtualization*, and software-defined storage (SDS) capabilities, they may be preferred over conventional direct-attached storage (DAS) and network-attached storage (NAS) file-based devices.

SANs establish switches and directors between servers and storage devices, enabling comprehensive network connectivity, allowing servers to share a common storage utility, which may be remotely located. SANs typically contain devices based on block-level storage, such as disk (JBOD/RAID), tape, and optical storage drives.

Modern SANs also function based on the concept of object storage, in which the data interaction occurs at the application level through an application programming interface (API) instead of direct access via the operating system. Object storage devices act as modular units that can be aggregated into larger storage pools across locations. They use the same types of disk units found in conventional SANs but are accessed via HTTP. They are ideal for unstructured, read-only data, including static web pages, backups, images, photos, music, video, and other multimedia files. Several online collaboration services have emerged to take advantage of these types of databases.

***Storage Virtualization** (per SNIA) is the act of abstracting, hiding, or isolating the internal function of a storage (sub) system or service from applications, compute servers, or general network resources for the purpose of enabling application- and network-independent management of storage or data. Storage virtualization is a conceptual extension of server virtualization.

Network Infrastructure Considerations

When choosing a data center network infrastructure, there are four key considerations: **manageability**, **flexibility**, **scalability**, and **reliability/availability**. These factors are crucial in building a network that meets your current demands, can effortlessly scale for future growth, and maximizes uptime, security, safety, and efficiency throughout its lifetime. Fiber optic cable infrastructure is vital and should be prioritized, even before power or cooling. Due to its flexibility, durability, density, bandwidth, and ease of use, it's the preferred choice, particularly as data centers are increasingly being tasked with supporting AI applications and advanced computational workloads that require robust data transfer capabilities.

Manageability

As networks evolve, what starts as an orderly, highly functional system can quickly become unmanageable, leading to a chaotic buildup of patch cords and unidentifiable cables. Implementing a structured cabling strategy from the outset increases control over data center infrastructure throughout a lifetime of moves, adds, and changes (MACs).

Structured cabling is a building or campus telecommunications infrastructure with standardized, smaller elements called sub-systems. These sub-systems are organized to make individual fibers easy to locate, manage MACs efficiently, and allow ample airflow around cabling. It's particularly beneficial for AI-driven operations that demand high reliability and minimal tolerance for downtime or network failure.

Flexibility

Whether they're part of planned maintenance or a response to unexpected issues, MACs must be executed quickly and efficiently with minimal impact on day-to-day operations and at the lowest risk of human error. A flexible solution provides the best balance between minimizing signal degradation in passive systems and simplifying MACs.

The term "network efficiency" often refers to energy usage, but it also encompasses the speed and ease with which MACs can be performed incrementally at the lowest capital and operational cost. With the growing prevalence of AI applications, data centers will continue to expand in size and number, as networks accommodate data rates of 400Gb, 800Gb, and even 1.6T. Simple, efficient cabling solutions can help maximize space and reduce time and costs for installation, maintenance, and upgrades, allowing data centers to effectively support AI and other advanced technology needs.

Overhead cable trays suspended directly from the ceiling facilitate the addition and removal of cabinets and racks of various heights while freeing up space below raised tiles to meet crucial airflow and cooling requirements. Redundancy in cable routes further enhances flexibility. This can be designed according to different **redundancy levels or tiers** to meet specific needs. Data centers rely on fiber optic cabling to simplify management and support new applications, including AI workloads.

Scalability

Scalability is crucial for efficiently migrating to higher data rates and adding capacity without significant disruption to operations. This is particularly important as AI applications demand more processing power and data throughput. There are two types of scaling: vertical and horizontal. Vertical scaling involves adding more processing power to existing machines. Horizontal scaling means increasing the number of machines within computing environments or application spaces. Both types directly impact cable infrastructure, particularly the number of data rates available on transmission links or channels.

With deployments of 400Gb, 800Gb, and 1.6T, platforms are now scaling up to thousands of ports using 8F to 32F multifiber connectors. This high density of fibers per chassis and MTP®/MPO-connector-based patching will demand unprecedented patch panel density and patch cable management at the interface between network cabling infrastructure and electronics. This can be most efficiently managed through a rack-mountable network cabling infrastructure. Higher data rates are also driving a shift from traditional duplex transmission to parallel optics. While wavelength division multiplexing (WDM) implementations may offer exceptions in some transmission schemes, those high data rates generally necessitate more optical fibers, especially in dual-layer-mesh or "spine-and-leaf" architectures, where network switches are fully cross-linked for increased data throughput and resiliency.

Reliability and Availability

Data center reliability and availability is associated with the likelihood that facility systems or components are operational and available at specified times and can be expressed as hours per year or related percentages. This allows operators to quantify acceptable risks regarding downtime and plan infrastructure accordingly.

The **Uptime Institute** segments uptime as Tier 1 to Tier 4, with Tier 4 data centers requiring uptime of 99.995% or less than half an hour of downtime per year. BICSI defines "reliability" as the probability that a piece of equipment or system will perform its intended function under specified conditions for a certain period without failure. Reliability can be expressed as a percentage, using mean-time-between-failures (MTBF) data associated with electronics or other ancillary equipment or active systems. Data center components' maintainability can be described with mean-time-to-repair (MTTR) data, which quantifies the probability that a component can be repaired within a specified time.

Beyond current system resiliency and redundancy, overall data center resources must be responsive to ever-changing business requirements, including rapid, unanticipated growth and technological innovations, such as those driven by AI. This responsiveness ensures that data centers can adapt and scale according to evolving demands, maintaining high reliability and availability.

Resources

Tools

- Bill of Materials (BOM)
- System Design Calculators

Products

- Contour™ Flow Cable
- Fiber Optic Cables and Ducts/Conduits
- Centrix™ system
- EDGE™ solutions
- EDGE8® solutions
- EDGE Distribution System
- EDGE Rapid Connect
- Optical Splice Enclosure
- CCH Hardware
- Co-packaged optics
- Optical Communications Product Catalog

Standards and Professional Organizations

- ANSI/TIA-942-C Telecommunications Infrastructure Standard for Data Centers
- BICSI Essentials of Data Center Projects
- ANSI/BICSI 002 Standard for Data Center Design
- American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE)
- Uptime Institute
- Leadership in Energy and Environmental Design (LEED)
- Advancing Data Center and IT Infrastructure Professionals (AFCOM)
- Storage Networking Industry Association (SNIA)

General Reference

- Corning GlassWorks AI™ Solutions
- Spine and Leaf Architecture
- Optimizing AI Data Centers with Structured Cabling Strategies
- Optical Communications Data Center Home Web Page
- Multitenant Data Center Solutions Guide

More Tips for Data Center Success

This guide is not intended to be a comprehensive handbook on facility construction and planning or the design and maintenance of the relevant systems. These are comprehensive and coordinated project management and professional engineering undertakings that bear on multiple industry standards, codes, and regulations. Adequate planning and consideration during facility construction or renovation are crucial, including calculating data center capital cost and exploring different cost scenarios. These calculations are usually used to compare how various architectures will operate for specific locations and configurations. Such analyses are typically carried out through the normal bidding process by general construction contractors, subcontractors, and equipment vendors at the design documentation completion levels. For data center project stakeholders or design and construction teams, the cost is naturally the paramount consideration.

With the preceding disclaimer stated, several design parameters are crucial to an overall successful design, implementation, and operation of a data center over its life. These include metrics such as availability, reliability, maintainability, and scalability, as well as other criteria such as capacity, efficiency, density, and of course, capital and operational costs. In today's ever-connected, always-on world of e-commerce, quantifiable attributes can be established to allow some predictive control over acceptable network downtime regarding facility resources.

Still have questions?

Learn more about our data center solutions, find helpful tools, and contact a rep **here**.

CORNING

Corning Optical Communications LLC • 4200 Corning Place • Charlotte, NC 28216 USA
800-743-2675 • FAX: 828-325-5060 • International: +1-828-901-5000 • www.corning.com/opcomm