From the Curb to the Computer Room

A Practical Walk Through the Spaces and Systems of a Modern Data Center

Publication Date
February 2026

Estimated Reading Time
15 minutes

Prepared By
Jeramy Utara
Volterra Advisors

Discipline
INFRASTRUCTURE

Geography
GLOBAL

Perspective
OPERATIONAL

Sections

Executive Summary

A data center is often described in terms of megawatts, rack density, or square footage. In practice, reliability is built much earlier — in how people, equipment, power, and cooling physically move through the building.

This report is a guided walk through a modern data center, following the same path taken by delivery drivers, engineers, and operators — from the perimeter entrance to the compute floor. Each major space is examined not as an isolated room, but as part of an interlocking system designed to minimize risk, simplify operations, and support long-term scale.

Everything that follows — from security to cooling — exists to manage how energy and people move through the building, and to ensure that those movements remain controlled even under stress.

> *Reliability is not an abstract property. It is built room by room.*

Why This Matters Now

Data center design is becoming more visible — and more contested — than it has ever been.

In many markets, the limiting factors for new data center development are no longer internal engineering challenges. They are external: power availability, permitting timelines, community sentiment, and infrastructure readiness. As facilities grow larger and denser, the physical realities of their operation — truck traffic, generators, cooling equipment, security perimeters — are increasingly part of public discussion rather than industry shorthand.

At the same time, workloads are changing. Higher-density compute, driven by AI and accelerated workloads, is pushing power and cooling systems harder and compressing tolerances that were once forgiving. Layouts that functioned well for traditional enterprise loads can become brittle under new operating regimes, exposing design shortcuts that were invisible on day one.

This shift has implications beyond uptime. Investors, utilities, and municipalities are now asking physical questions: How does this building actually work? Where does equipment move? What happens during maintenance, emergencies, or outages? Understanding the physical systems inside a data center is no longer just an operational concern — it has become a prerequisite for permitting, financing, and long-term acceptance.

Fundamentals: How a Data Center Actually Works

At its core, a data center is a physical system designed to do one thing reliably: convert electricity into computation, and remove the resulting heat continuously, safely, and predictably.

Electricity enters the site, is transformed and distributed through layers of equipment, and ultimately powers IT hardware. Nearly all that electrical energy becomes heat. Cooling systems exist to move that heat away from sensitive equipment and reject it to an external sink — air, water, or both — without interruption. Around those core functions sits a parallel set of systems designed to control movement: who enters the building, how equipment arrives and is staged, how racks travel through corridors, and how engineers operate and maintain the facility under normal and abnormal conditions.

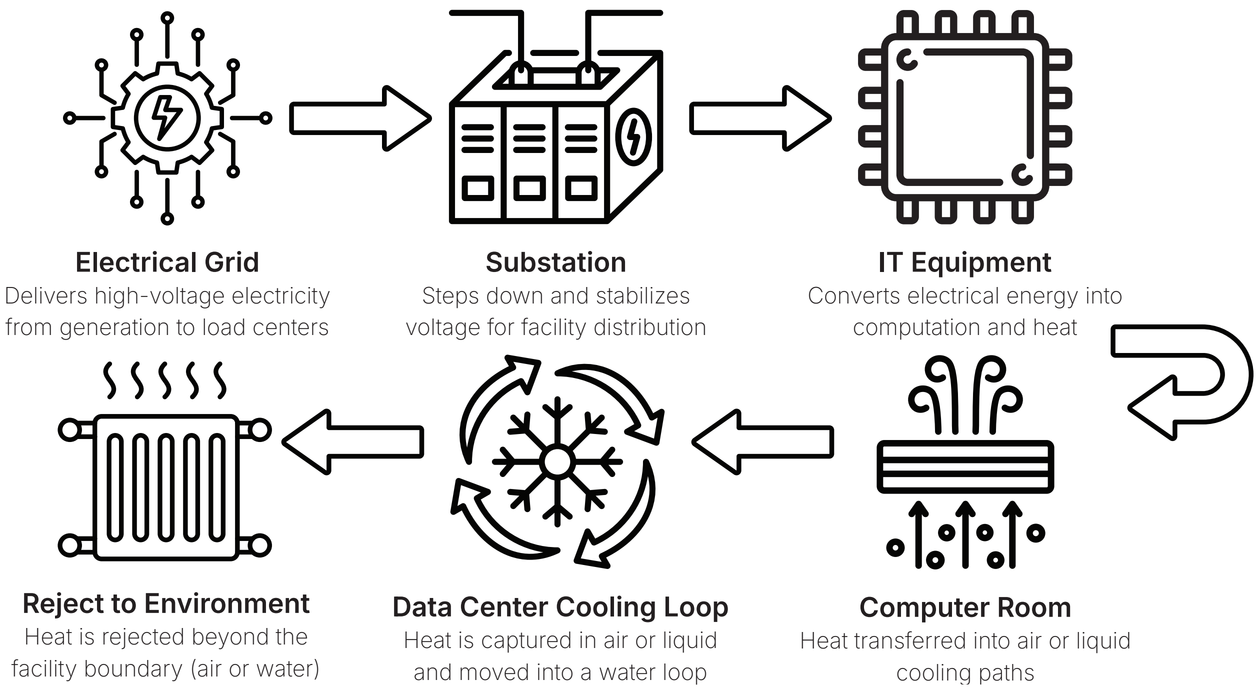

Exhibit A: The Energy Loop Inside a Data Center
Electricity entering a data center does not disappear — it is converted into computation and heat. Cooling systems exist to transport that heat safely beyond the facility boundary. The diagram is conceptual; specific architectures vary by design.

Reliability emerges from how these flows intersect. A perfectly redundant power system can still fail if fuel cannot be delivered or pumps are flooded. A highly efficient cooling plant can still underperform if air paths are constrained or maintenance access is poor. Conversely, well-designed circulation, visibility, and control spaces can absorb operational stress and prevent small issues from becoming outages.

The sections that follow trace these flows in order — from the perimeter inward — to show how reliability is built not through any single component, but through the alignment of many physical decisions made early and lived with for decades.

Arrival and Access

The entry sequence into a data center sets the tone for everything that follows — not just for security, but for daily operations.

Arrival typically begins at a secured perimeter, with most facilities gated and monitored. Visitors request entry through an intercom system, allowing security staff to verify identity and purpose before granting access to the site. This initial layer establishes control well before anyone reaches the building itself.

Inside, the lobby functions as the first interior control space. Access to the lobby itself is typically secured; visitors without a badge must request entry via intercom before doors are unlocked. Once inside, the lobby is where visitors enroll in security, check in, and wait for escorts.

Lobbies are typically furnished with seating and badge enrollment stations and are designed to maintain direct visibility from the Security Operations Center (SOC) rather than relying on a traditional front desk. This layout reinforces centralized control and ensures that visitor movement is continuously observed as individuals transition from the public exterior to the controlled interior of the data center.

The Security Operations Center (SOC) remains the nerve center of the facility and is typically among the most hardened spaces in the building. SOCs are commonly constructed with bullet-resistant glass and blast-resistant walls, reflecting their role as a last line of control during incidents. Importantly, SOCs should be designed with more than one exit, ensuring that guards cannot be trapped inside during an emergency.

From the SOC, security personnel maintain continuous visibility across the facility through extensive CCTV coverage. Most areas of the data center are monitored by cameras that record to on-site DVR systems, with footage retained for a minimum of 60 days[1]. Direct visual adjacency between the SOC and high-traffic areas — especially the loading dock — further simplifies oversight in ways cameras alone cannot, improving response times and reducing operational friction.

Entrance into the secure portion of the data center is typically controlled through a mantrap. These systems prevent piggy-backing and enforce one-person-at-a-time entry. Access commonly requires a combination of card-key authentication and biometrics, and in some facilities includes weight verification on entry and exit to ensure that individuals do not remove or leave behind unauthorized equipment.

1 Continuous video surveillance with on-site retention periods of 60 days or more is a common operational practice in enterprise and hyperscale data centers, particularly those aligned with SOC 2 and ISO/IEC 27001 control frameworks. While retention durations are not universally mandated, extended on-site storage supports incident investigation, audit response, and forensic review, and is widely adopted as a baseline expectation in secure facilities.

> *"Defense-in-depth security includes 24/7 staffed access, mantraps, biometric authentication, & CCTV monitoring to prevent unauthorized hardware access."*
>
> — Charlie Lane (Equinix), quoted in Interconnections on Equinix Security Philosophy

While access control technologies have evolved incrementally, the security model itself has changed little over the past decade. Many facilities still rely on 24/7, 365-day guard staffing. Increasingly, however, operators are exploring opportunities to reduce operating expense through supplemental tools such as autonomous patrol robots and drones, particularly for perimeter rounds and monitoring.

One frequently overlooked requirement is lobby scale. In larger facilities, the lobby must be able to stage 20 or more visitors simultaneously, particularly during customer tours, audits, and multi-team walkthroughs. The same space must also accommodate the daily flow of operations staff during shift changes, when technicians exchange personal IDs for facility badges. Undersized lobbies quickly become choke points at exactly the moments when order matters most.

It is also highly beneficial to include a small conference room directly off the lobby. This allows meetings, interviews, and vendor discussions to occur without requiring guests to enroll in security or enter the secure interior of the data center. In practice, this space proves invaluable — especially for interviews, audits, or short meetings that do not warrant full access processing.

Logistics and Control: The Loading Dock

The loading dock is one of the most operationally active spaces in a data center and the primary entry point for critical equipment.

A well-designed dock includes multiple bays configured for different uses. At least one bay should be reserved to accommodate large 30-yard dumpsters, which are routinely required during equipment refreshes, decommissions, and construction activity. Other bays are typically equipped with truck lifts or dock levelers to handle both full-size trailers and smaller delivery vehicles.

Adjacency between the loading dock and the SOC is equally important. Direct visibility is simpler and more reliable than camera systems that require panning, zooming, or troubleshooting. In high-traffic areas, seconds matter, and the ability for security staff to physically reach the dock within moments is a meaningful operational advantage.

With a strong scheduling system and rules-based dock reservations, deliveries tend to be

straightforward. In practice, most operational issues stem not from delivery volume, but from poor layout or insufficient separation between secure and semi-secure zones.
If there is one universal operational lesson, it is this: always include a restroom at the loading dock.

Delivery personnel require access to basic facilities without entering the fully secure envelope of the data center. Providing a restroom within the dock allows drivers to complete deliveries without triggering escort requirements or security processing — a small design choice that eliminates repeated access requests and reduces guard workload.

Once equipment clears the dock, the quality of internal circulation determines how safely and efficiently it reaches the data hall.

Circulation: Hallways & Vertical Movement

Circulation spaces — corridors, ramps, and elevators — quietly determine how safely and efficiently a data center operates over its full life.

Hallways are typically designed to be wide, tall, and unobstructed, allowing racks and equipment to move freely throughout the building. Clear overhead space is critical; as rack designs evolve and power densities increase, ceiling congestion can quickly become a limiting factor. What feels generous on day one often becomes barely sufficient a decade later.

Elevation changes deserve particular scrutiny. Even minor grade differences between spaces — such as between a loading dock and the computer room floor — can become operational liabilities as rack weights increase over time. What may initially be manageable by hand often becomes unsafe, eventually requiring mechanical lifts or temporary workarounds. These retrofits are expensive, disruptive, and almost always avoidable with early design discipline.

Vertical circulation follows the same logic. Freight elevators are wider and better suited for equipment movement, but passenger elevators are often rated for higher loads and can serve as effective backups if designed with that contingency in mind. Planning for a passenger elevator to act as a secondary freight path can eliminate the need for redundant freight elevators while preserving operational resilience.

Finally, visual clarity matters more than most designers expect. Large data centers can quickly become visually repetitive, increasing the risk of navigation errors during maintenance events, emergencies, or vendor visits. Clear signage, consistent labeling, and visual differentiation between zones reduce confusion and improve safety — especially when time and attention are limited.

Where the Outside Connects In: MMRs & POEs

Meet-Me Rooms (MMRs) and Points of Entry (POEs) are the physical interfaces between the data center and the outside world.

These spaces are typically compact, highly-secure rooms located near the building perimeter, designed to terminate incoming fiber from carriers, long-haul networks, and metro providers. They contain fiber vault penetrations, conduit pathways, patch panels, splice enclosures, and cross-connect infrastructure. While not visually dramatic, these rooms are foundational to network resilience and customer flexibility.

Over the past decade, MMR and POE design has evolved quietly. Rather than a few large interconnection rooms, modern facilities increasingly deploy multiple, smaller entry points. This reflects customer demand for additional redundancy, diverse routing, and reduced shared risk across carriers and fiber paths. In practice, these rooms are becoming smaller but more numerous.

In some architectures, customers bypass traditional POEs altogether, routing fiber directly from exterior vaults to dedicated pathways into the computer room. This approach can reduce latency and simplify certain network topologies, but it places greater emphasis on disciplined exterior routing, physical protection, and clear demarcation between customer-owned and facility-owned infrastructure.

From a leasing and diligence perspective, the quality of MMR and POE design is less about size and more about clarity — clear demarcation points, logical cable management, space for future growth, and the ability to add diverse paths without disruption. Well-designed interconnection spaces quietly enable everything that happens downstream.

The Operational Brain: Facility Operations Center (FOC)

The Facility Operations Center (FOC) is the primary coordination space for building operations and engineering teams.

Physically, the FOC is a centralized room designed to support both routine monitoring and high-consequence decision-making. It typically houses workstations and shared displays that provide real-time visibility into the building's critical systems, allowing teams to understand current conditions at a glance.

In normal operation, the FOC supports daily coordination, shift handoffs, and planning activities. During incidents or major maintenance events, it becomes the command space where issues are assessed, actions are sequenced, and communication is centralized. In those moments, clarity and accessibility matter far more than aesthetics.

The FOC is also commonly the home of several essential control and monitoring systems, including:

- Building Management Systems (BMS) and Building Automation Systems (BAS)
- Fire alarm and life safety monitoring
- Fuel pump controls and fuel polishing system monitoring

Persistent visibility into power, cooling, alarm, and fire systems allows teams to build shared situational awareness quickly. While documentation is increasingly digital, hard copies of Methods of Procedure (MOPs) remain critical. During network or power disruptions, physical binders are often the fastest and most reliable reference and, in practice, easier to navigate under pressure.

Power Infrastructure: From Utility to Rack

Power systems occupy some of the most substantial and deliberately engineered rooms in the data center.

Medium-voltage (MV) rooms and electrical galleries are typically large, high-clearance spaces designed around safety, access, and maintainability. Inside these rooms, the dominant elements are switchgear lineups — long runs of metal-clad cabinets housing breakers, protective relays, and bus bars. Their primary role is not simply to distribute power, but to isolate faults. When something goes wrong, the system is designed to fail locally and predictably, rather than cascading across the facility.

Breakers are sized not just for steady-state load, but for fault interruption — the ability to safely interrupt extremely high currents during abnormal events. This is why electrical equipment appears oversized relative to its apparent function. Physical size reflects arc-flash energy, thermal limits, and safety boundaries as much as amperage.

The generous spacing between equipment is intentional. Clearances are dictated by electrical code, arc-flash protection zones, and the practical reality that maintenance must often occur on energized systems. Tight electrical rooms become operational liabilities over time, increasing both safety risk and downtime during routine service. In well-designed facilities, electrical rooms are built to be worked in, not just walked through.

Uninterruptible Power Supply (UPS) rooms are similarly purposeful. Over time, static UPS systems have largely overtaken rotary designs, driven by scalability, predictable maintenance profiles, and maintainability. In most modern architectures, static UPS systems can provide meaningful ride-through — often up to ~20 minutes depending on configuration and battery capacity — bridging the gap until generators assume load and the facility stabilizes. Rotary

UPS systems, by contrast, are more commonly associated with shorter ride-through windows (often on the order of seconds), relying on inertia and fast generator start rather than extended battery duration.[2]

UPS rooms are defined as much by energy storage as by power electronics. Battery systems occupy significant floor area and impose substantial structural loads, which is why these rooms are often slab-on-grade and isolated from adjacent spaces. Ventilation, fire separation, and access paths are designed around worst-case scenarios rather than normal operation, reflecting the reality that reliability planning assumes things will eventually go wrong.

The physical footprint of UPS infrastructure reinforces a broader theme: redundancy is not abstract. It has weight, volume, and clearance requirements that must be accommodated early if a facility is to remain flexible as loads and technologies evolve.

> *"We watched other people's stuff get submerged and fail. Those [incidents] make you re-think where you're going to put your generators, your fuel pumps, etc."*
>
> — Hunter Newby (Netrality), quoted in Data Center Dynamics on Hurricane Sandy

Generator systems are often housed either in dedicated interior generator rooms or externally, with careful attention paid to exhaust routing, acoustics, and vibration isolation. Fuel systems — including storage tanks, transfer pumps, and polishing equipment — are equally critical and often monitored directly from the FOC.

Fuel can sit in tanks for months or even years. Without regular fuel polishing, degradation can render generators unusable when they are needed most. Maintaining fuel quality ensures generators are immediately deployable during extended outages.

A defining lesson emerged during Hurricane Sandy in the early 2010s. Many data centers in the New York and New Jersey region lost load not because generators failed, but because fuel pumps and generators were located below grade and flooded. Facilities that placed pumps above ground and generators on rooftops or penthouses were able to operate for days — even weeks — without dropping customer load. Elevation decisions, often treated as secondary during design, became decisive in practice.

2 Diesel rotary UPS (DRUPS) systems combine a motor-generator with energy storage and can bridge mains failure only for a short interval (typically tens of seconds) until a backup generator takes over. Static UPS systems paired with battery storage dominate data center deployments and, depending on battery capacity and load, are generally configured to provide extended ride-through durations measured in minutes rather than mere seconds.

Cooling and Heat Rejection

At the plant level, cooling infrastructure typically includes chillers, cooling towers or dry coolers, pumps, and heat exchangers, arranged either within dedicated mechanical rooms or outdoors in secured yards or rooftop enclosures. Some facilities also incorporate air-side or water-side economization, using rooftop equipment or exterior heat exchangers to take advantage of favorable ambient conditions and reduce mechanical cooling demand.

Mechanical rooms are less about individual machines and more about flow. Pipes, headers, and valves dominate these spaces because moving heat is ultimately about moving mass. Pumps consume energy not to create cooling, but to transport heat away from where it is generated and toward a place where it can be rejected or reused.

Equipment is intentionally oversized and redundant. Parallel pumps, bypass loops, isolation valves, and strainers exist so that components can be serviced without interrupting operation. This is why mechanical rooms often appear spacious and repetitive: reliability depends on maintaining flow under maintenance, partial failure, and changing load conditions.

As with electrical rooms, physical layout matters. Poor access to valves, tight pipe corridors, or constrained service clearances can turn routine maintenance into risk events. Well-designed mechanical rooms assume intervention — not just steady operation — as a normal part of the system's life.

Chilled water produced at the plant is distributed throughout the facility to serve the computer rooms. Depending on design, that chilled water may be used to cool air — via CRAHs, fan walls, or air handlers — or delivered more directly to the IT load through liquid-to-chip (L2C) or rack-level cooling systems. In air-cooled designs, chilled water removes heat from the room environment; in liquid-cooled architectures, it removes heat closer to the source, enabling higher rack densities and tighter thermal control.

There is no single "correct" cooling architecture. Different designs rely on different combinations of equipment, and those choices directly influence the layout and footprint of the building. Facilities designed around large chiller plants and cooling towers often dedicate significant yard or roof space to heat rejection, while air-cooled or hybrid systems may shift that footprint vertically or reduce on-site water infrastructure altogether. Increasingly, these early cooling decisions are shaping not just mechanical rooms, but the overall massing and siting of data centers.

Cooling strategy today is shaped by more than pure engineering. Closed-loop, non-evaporative heat rejection systems are becoming more common due to water scarcity, permitting complexity, and community sentiment. These designs reduce dependence on municipal water supplies and simplify long-term compliance, particularly in water-stressed regions.

At the same time, many operators are intentionally designing computer rooms with future cooling flexibility. By selecting distribution systems and room layouts that can support both air cooling and liquid-cooled architectures, facilities can transition from fan-wall or CRAH-based designs to liquid-to-chip systems through equipment changes rather than full mechanical retrofits. This adaptability allows data centers to support traditional enterprise workloads alongside high-density AI deployments without reworking the plant.

For a deeper discussion of how cooling design choices affect water use, energy efficiency, and community considerations, see Volterra's report *Water, Watts, & Workloads*, which examines the trade-offs between evaporative, dry, and liquid-cooled architectures in detail.

The Computer Room

The computer room, or data hall, is the physical endpoint of every upstream design decision.

These rooms are typically expansive, open spaces designed to support long-term flexibility. Raised floors have largely given way to slab-on-grade construction, driven by increasing rack weights and higher electrical distribution densities. Power is commonly delivered via overhead busways or distribution boards, allowing racks to be added or reconfigured without extensive electrical work.

Air delivery is often overhead as well, with hot- and cold-aisle containment used to manage airflow efficiently. As rack densities increase, environmental monitoring becomes more granular. Modern data halls rely on a dense network of temperature, humidity, and pressure sensors to detect hotspots early and maintain tight operating margins.

Fire protection systems — including early smoke detection and pre-action suppression — are integrated throughout the room, designed to respond quickly while minimizing the risk of accidental discharge.

While the computer room may appear visually similar across generations, the infrastructure supporting it has changed dramatically. Fewer racks now consume far more power, placing greater demands on power distribution, cooling responsiveness, and monitoring fidelity.

Beyond the Data Hall: Flex Space & Building Geometry

Modern customers increasingly request flex space adjacent to their deployments. This may range from a few hundred square feet to hundreds of square feet per megawatt, used for offices, secure staging, or storage. These spaces are often bundled into leases but can also represent incremental revenue opportunities.

Building geometry is evolving as well. Where data centers were once long, rectangular structures, square and U-shaped footprints are becoming more common. High-performance computing workloads require extremely low latency between clusters; excessive distance between data halls can prevent systems from operating as a single logical unit. Shortening fiber paths is now influencing not only interior layouts, but exterior building design.

As data center workloads continue to densify, physical design will matter more, not less. Higher power densities compress tolerances around circulation, cooling, and maintenance access, while liquid-cooled architectures shift where heat is captured and how systems interface. At the same time, facilities are being asked to accommodate flexibility — for tenants, future equipment generations, and evolving regulatory expectations — without sacrificing reliability.

These forces don't require a new design philosophy so much as a return to fundamentals: layouts that respect flow, systems that assume change, and buildings designed to operate under stress rather than ideal conditions.

Design Is Experience

The most important design decisions in a data center are rarely the most visible.

They are made early, often quietly, and typically long before a facility is live. Yet those decisions echo for decades — shaping how safely equipment moves, how quickly teams respond to issues, how easily customers expand, and how resilient a building remains under stress.

> *"Data center outages are becoming less frequent and less severe relative to the rapid growth of digital infrastructure."*
>
> — Uptime Institute Survey Summary

From lobby size to loading dock layout, from fiber entry paths to fuel pump elevation, these choices are not abstract best practices. They are the accumulated result of experience — of understanding how buildings are actually used, how they age, and where theory diverges from reality.

Well-designed data centers do not just meet specifications on day one. They continue to function smoothly through equipment refreshes, changing workloads, evolving cooling strategies, and moments of real operational pressure. The difference is rarely one system or one technology, but how thoughtfully the pieces are arranged and how easily they work together.

Key Takeaways

- Reliability is physical first: The layout of rooms and pathways often determines uptime as much as redundancy diagrams do.
- Security begins at the perimeter: Gated entry, controlled lobbies, hardened SOCs, and mantrap discipline set the operational baseline long before anyone reaches the data hall.
- Design for real traffic, not drawings: Oversized lobbies, direct SOC visibility, and dock workflows reduce friction during tours, shift changes, and delivery events.
- The loading dock is a reliability system: Bays, dumpsters, dock levelers, and driver facilities eliminate avoidable access churn and keep high-velocity logistics from bleeding into secure space.
- Circulation must survive the full “rack life”: Wide corridors, clear overhead space, minimal elevation changes, and smart elevator contingencies prevent future retrofits as racks get heavier.
- Interconnection is about clarity and growth: MMR/POE layouts that preserve demarcation, routing discipline, and expansion space protect customer flexibility and future leasing outcomes.
- The FOC earns its keep under pressure: Centralized monitoring, accessible procedures, and cross-system visibility matter most during incidents, maintenance windows, and shift handoffs.
- Fuel systems fail more quietly than generators: Polishing, pump placement, and flood resilience often decide whether backup power works when it’s needed.
- Cooling architecture shapes the building itself: Chillers, towers or dry coolers, economization, and distribution choices drive footprint, roof/yard planning, and long-term flexibility.
- Flex space is part of the product: Offices, secure staging, and storage increasingly influence customer satisfaction and can materially affect leasing terms and velocity.
- HPC is changing geometry, not just equipment: Shorter internal distances and more compact layouts are becoming design drivers as clusters demand ultra-low latency between rooms.

Glossary of Terms

BAS (Building Automation System)
A subset of building controls focused on automation, sequencing, and environmental regulation.

BMS (Building Management System)
Software and controls used to monitor and manage mechanical and electrical systems such as cooling, power, and alarms.

Closed-Loop Cooling
A cooling system that recirculates fluid within a sealed circuit, reducing water consumption and simplifying environmental compliance.

CRAH / Fan Wall
Cooling equipment that circulates conditioned air through the data hall, often used in high-density environments.

Cooling Tower
A heat rejection device that removes heat by evaporating a small portion of recirculating water; efficient electrically, but it consumes water and requires chemical treatment and drift control.

Cross-Connect
A physical fiber or copper connection linking two networks or endpoints within a data center.

Economization (Air-Side / Water-Side)
Using outdoor air or water conditions to reduce reliance on mechanical cooling systems.

Fuel Polishing
Ongoing filtration and treatment of stored diesel fuel to prevent degradation and ensure reliable generator operation during extended outages.

Liquid-to-Chip (L2C) Cooling
A liquid cooling method that delivers coolant directly to processors, enabling higher rack densities and improved thermal control.

MMR (Meet-Me Room)
A secure room where multiple network providers and customers interconnect fiber within a data center, enabling redundancy, diversity, and flexible network expansion.

MOPs (Methods of Procedure)
Step-by-step operational documents used to safely execute maintenance, testing, and changes to critical systems.

Non-Evaporative Cooling
Cooling approaches that reject heat without evaporating water, such as dry coolers or air-cooled heat exchangers.

POE (Point of Entry)
The location where external fiber infrastructure enters the building, typically routed from underground vaults or conduits into protected interior pathways.

Pre-Action Fire Suppression
A fire protection system that requires multiple triggers before releasing water, minimizing accidental discharge risk

Rack Density
The amount of electrical load per rack, typically measured in kilowatts, which drives power and cooling design.

Rotary UPS

A UPS system that uses rotating mass and mechanical inertia to bridge short-duration power events, typically paired with rapid generator start.

Static UPS

A battery-based uninterruptible power supply that conditions incoming power and provides sustained ride-through during utility interruptions until generators stabilize.

Scope & Methodology

Volterra Reports are based on a combination of publicly available data, industry standards, operator and municipal disclosures, and first-hand market experience. Technical examples and metrics are illustrative, reflecting modern hyperscale and colocation facility designs under typical climate, load, and operational assumptions. Where ranges are shown, they are intended to clarify relative trade offs rather than prescribe specific designs.

Analysis focuses on facility-level systems and local, on-site impacts unless otherwise noted, and does not attempt to model upstream grid effects, full life-cycle emissions, or site-specific engineering constraints.

Sources & References

External sources are cited selectively to anchor specific examples and industry observations; the analysis and conclusions reflect Volterra's first-hand operational experience.

Data Center Dynamics. "Hurricane Sandy: Data Center Stories from Manhattan."
Accessed November 2025.
https://www.datacenterdynamics.com/en/news/hurricane-sandy-data-center-stories-from-manhattan/

Interconnections (Equinix). "The Equinix Security Philosophy."
Accessed November 2025.
https://blog.equinix.com/blog/2025/12/02/3-ways-high-performance-data-centers-mitigate-risk/

Uptime Institute. "Global Data Center Survey."
Accessed November 2025.
https://uptimeinstitute.com/research-and-reports

Wikipedia. "Diesel Rotary Uninterruptible Power Supply (DRUPS)."
Accessed November 2025.
https://en.wikipedia.org/wiki/Diesel_rotary_uninterruptible_power_supply

Schneider Electric Blog / FacilitiesNet. "Lessons Learned from Hurricane Sandy."
Accessed November 2025.
https://blog.se.com/datacenter/architecture/2012/11/29/lessons-learned-from-hurricane-sandy/

© 2026 Volterra Advisors LLC.
Portions of this report may be quoted with attribution to Volterra Advisors.
Reproduction of full sections or graphics for commercial use requires permission.

About Volterra

Volterra Advisors is a boutique advisory platform dedicated to accelerating growth across the digital infrastructure and data center sectors. Founded by Jeramy Utara, a recognized industry veteran with nearly two decades of experience, Volterra helps developers, investors, landowners, and operators navigate the decisions that matter most — from power and land strategy to leasing, platform growth, and market expansion.

Jeramy began his career in 2007 at DuPont Fabros Technology (later acquired by Digital Realty), supporting the development of some of the earliest large-scale hyperscale campuses in North America. He later joined CloudHQ as its first employee, where he helped grow the company from concept to one of the world's most successful privately held data center developers. Over more than a decade at CloudHQ, Jeramy played a leading role in leasing more than 1 gigawatt of capacity and driving over $15 billion in revenue, shaping relationships with the world's largest cloud, AI, and enterprise clients.

Through Volterra, Jeramy now partners with digital infrastructure platforms to define and execute strategies around site selection, powered land, leasing, marketing, and organizational growth. His advisory work emphasizes clarity, momentum, and execution — translating complex market forces into tangible outcomes for emerging and established platforms alike.

How We Can Help

Market Intelligence

We analyze power, permitting, and competitive dynamics to identify where the next generation of data center growth will land.

Site Identification & Readiness

We help landowners and developers qualify, position, and advance properties into premier data center ready sites.

Utility & Power Engagement

We maintain direct relationships with utilities nationwide to confirm capacity, align on timing, and support interconnection strategy.

Marketing & Positioning

We produce investor-grade materials — from decks to data sheets — that clearly communicate technical strengths and value.

Sales & Buyer Engagement

We connect landowners and developers directly with hyperscalers, operators, and investors ready to transact.

Partner & Consultant Network

We bring trusted engineering, permitting, and development partners to every opportunity — ensuring readiness, speed, and credibility.

Capital & Transaction Support

We support capital formation, deal structuring, and JV alignment — keeping incentives tied to successful outcomes.

Strategic Advisory

We serve as an extension of your leadership team — providing guidance, relationships, and insight from concept to close.

About Volterra Reports

Volterra Reports examine the systems, markets, and decisions shaping digital infrastructure — connecting power, land, cooling, policy, and technology into a coherent view of how the sector operates and evolves.

Each publication combines technical understanding with strategic perspective to explain how legacy decisions and emerging trends influence today's investments and tomorrow's platforms.

Average Read Time	Release Frequency	Audience	Distribution Format
15 minutes	Monthly	Digital Infrastructure	PDF + web

Example Reports

Water, Watts, & Workloads DEC 2025

Behind every megawatt of compute sits a cooling system that shapes both the energy profile and water footprint of a data center. *Watts, Water, & Workloads* follows the evolution of cooling from traditional air systems to modern evaporative and liquid-cooled architectures — revealing how AI-driven densities are rewriting the rules of heat rejection. With communities and utilities paying closer attention to water use, and operators pushing for efficiency at scale, cooling has become a defining design decision. The result is a landscape where water is emerging as a second 'fuel' for compute, and where liquid cooling is enabling dry, high-temperature pathways once considered impractical.

Waste Heat, Warm Cities JAN 2026

Data centers are increasingly measured by what they consume — megawatts, water, land, and grid capacity — but that focus can miss the most consistent physical output of compute: heat. *Waste Heat, Warm Cities* follows where that heat comes from and where it normally goes, then asks a simple question with complex answers: when does heat stop being a byproduct and become a community benefit? The report traces the practical mechanics of reuse — temperature limits, heat pumps, transfer stations, and district heating networks — and shows why Europe has become the proving ground for repeatable models. Through real deployments and honest constraints, it explains why heat reuse hasn't scaled everywhere, what conditions make it viable, and how new pressures (policy, disclosure, and higher-density cooling architectures) are changing the equation. The result is a clear framework for treating waste heat as infrastructure: engineered, metered, contracted, and planned as early as the substation.

Volterra Reports are written for clarity, momentum, and practical application — bridging technical depth with real-world decisions.