



Legacy Fiber, Lasting Impact

How 1990s Telecom Overbuild Made Ashburn the Data Center Capital

PUBLICATION DATE

November 2025

ESTIMATED READING TIME

15 minutes

PREPARED BY

Jeremy Utara
Volterra Advisors

DISCIPLINE

TELECOM

GEOGRAPHY

GLOBAL

PERSPECTIVE

HISTORICAL

SECTIONS

Executive Summary	2
Why This Matters Now	2
Background & Context	3
The Setup	3
The Build	4
The Bust	4
The Hand-Off	5
Fundamentals	5
Global Examples	6
Why It Hasn't Scaled	9
How It Could Scale	10
What Changes Next	10
Closing	11
Key Takeaways	12
Glossary	13
Scope & Methodology	14
Sources & References	14
About Volterra	16
About Volterra Reports	17

Executive Summary

Ashburn's rise is easier to understand as an infrastructure inheritance than as a single "data center boom." The region's advantage emerged in four acts: the setup (early interconnection gravity), the build (aggressive fiber deployment), the bust (bankruptcies that reset prices), and the hand-off (cloud and colocation scaling on top of an already-dense network layer).

In the 1990s, Northern Virginia's "setup" wasn't just a cluster of tech firms; it was a practical place for networks to meet. Early interconnection in the region – especially around Metropolitan Area Exchange – East (MAE-East) – made it rational for carriers to route capacity through the corridor because it increased the chances that new fiber would be monetized through peering, transit, and redundancy demand.

THE TELECOM BUST DIDN'T REMOVE THE NETWORK ADVANTAGE – IT MADE IT CHEAP ENOUGH TO BECOME EVERYONE'S STARTING POINT.

The "build" phase produced what looked, in hindsight, like excess: conduit and fiber installed ahead of realized utilization. When the telecom downturn arrived (the boom peaked and then broke around the 2000 era), the damage was financial – firms failed and restructured – but much of the physical network remained in place. That mismatch created "stranded" infrastructure that didn't disappear; it waited to be repurposed.

The "hand-off" is the part that's easy to miss: the bust effectively repriced the network layer, allowing later operators to buy, lease, and light capacity without rebuilding it from scratch. That lowered time-to-market and reduced dependency risk, and it made dense, multi-carrier connectivity a standing feature of the region's economics – not a one-time accident.

Why This Matters Now

Many regions want "the next Ashburn," but most replication attempts start with the visible pieces – land, incentives, buildings – and underweight the invisible one: interconnection as market structure. Ashburn's moat is not simply "fiber exists." It is fiber optionality – dense, flexible, multi-carrier connectivity that makes it easier to connect to everyone else with fewer bespoke projects and fewer single-provider dependencies.

The constraint set is also shifting. In mature hubs, the gating item is increasingly power availability and delivery timelines, not whether you can physically pull fiber to a site.

Finally, the sector is no longer "invisible infrastructure." Communities can't easily see buried conduit; they can see substations, transmission lines, building massing, diesel tanks, and traffic. That public legibility changes permitting dynamics – and it makes the historical lesson from

Northern Virginia practical: clusters compound when governance, timelines, and infrastructure planning stay aligned over decades, not quarters.

Background & Context

In the 1990s the Dulles Corridor (around Washington-D.C., including Ashburn, Loudoun County, and adjacent areas) was emerging as a key hub for technology, fiber-optic networks and internet infrastructure. The area was dubbed the "Netplex" in the 1990s and later the "Silicon Valley of the East".

During that period, many firms – large telecoms and new-age carriers – invested heavily in laying out long-haul fiber-optic cables, building out dark-fiber networks¹, and constructing interconnection points. However, many of these investments came at the tail-end of the dotcom/telecom build-out boom. When demand didn't immediately materialize at the expected scale, some of these fiber carriers went bankrupt (for example AboveNet, formerly Metromedia Fiber Network, filed bankruptcy in 2002).

"WHAT LOOKED LIKE OVER-INVESTMENT IN THE 1990S BECAME THE SINGLE GREATEST HEAD-START ANY DIGITAL REGION EVER HAD."

– LOUDOUN COUNTY ECONOMIC DEVELOPMENT, "DAWN OF DATA"

Thus, while many investments looked over-built and risky at the time, they left behind a physical infrastructure "foundation" (lots of fiber, conduit, interconnection points) that would prove pivotal.

The Setup

In the early commercial internet period, being near a trusted "meeting place" mattered. If multiple carriers interconnect in one place, that location becomes a logical point for redundancy routes, regional aggregation, and enterprise access. In Northern Virginia, MAE-East served as one of those early gravity points – an exchange ecosystem that pulled networks into proximity and made the region legible as "where the internet meets."

In that era, "meeting place" was literal. It meant a small number of facilities where networks could show up, connect quickly, and leave with immediate reach to other networks. Once a

¹ UNUSED OR "UNLIT" OPTICAL FIBER THAT HAS BEEN INSTALLED BUT IS NOT CURRENTLY CARRYING DATA TRAFFIC. TELECOM FIRMS ORIGINALLY LAID EXTENSIVE FIBER NETWORKS DURING THE 1990S TECH BOOM, MUCH OF WHICH WENT UNUSED AFTER THE DOT-COM CRASH. THIS SURPLUS "DARK FIBER" CAN LATER BE LEASED OR "LIT" BY NEW OPERATORS, PROVIDING HIGH-CAPACITY, LOW-LATENCY CONNECTIVITY WITHOUT THE NEED FOR NEW CONSTRUCTION.

location became known as a place where interconnection was routine – not exceptional – it started to attract the next participant simply because it reduced the amount of custom work required to get online.

A critical nuance is that the setup was not primarily branding; it was topology. Northern Virginia became a practical place for carriers to meet, exchange traffic, and build redundancy – and that made later infrastructure bets more likely to find counterparties.

That early interconnection adjacency did not guarantee that every fiber dollar would pay off. It did something more subtle: it increased the probability that any new route landing in the corridor could find counterparties, because counterparties were already there – or were being pulled there. That's why the next act – building ahead of demand – tended to happen in places like this first.

The Build

The late 1990s and early 2000s fiber buildout was shaped by a belief that demand would grow fast enough to justify route coverage and early capacity. Technical advances in optical transmission raised the ceiling on how much traffic a given physical network could eventually carry, reinforcing the logic that “building ahead” might be rational even if near-term utilization lagged.

There's also a physical reality to this phase that matters later: once conduit is in the ground and routes are established, much of the hardest work is already done. Even if strands sit dark for years, the civil construction, routing, and rights-of-way decisions can outlive the business plan that funded them.

This wasn't “waste” so much as a timing mismatch: the infrastructure created future capability, but the financing assumed immediate monetization. And when demand doesn't arrive on the timeline debt requires, the story stops being about engineering and becomes about ownership.

The Bust

When the downturn hit, the damage showed up in debt schedules, equity collapses, and bankruptcies – not in the physical disappearance of conduit or fiber. Importantly, carriers failed or restructured, but the assets remained and the market entered a long digestion period.

That digestion is the under-told part of the story. Networks change hands, contracts get renegotiated, and prices reset in ways that don't make headlines – but those shifts determine whether the infrastructure becomes stranded forever or becomes usable “base layer” for the next cycle.

Dark-fiber networks and related assets were often unloaded at extremely low prices, and then – slowly – traffic began to fill what had been built ahead of demand.

This is where “bankruptcies” become a counterintuitive enabler. They reset ownership and pricing. The region’s network layer didn’t have to be rebuilt; it had to be reheld by entities whose time horizons matched infrastructure. That reset sets up the hand-off: a cheaper, more available network layer becomes the platform the next industry wave can stand on².

The Hand-Off

By the late 2000s and into the 2010s, the demand driver changed. It was less about speculative route builds and more about cloud adoption, streaming, enterprise digitalization, and low-latency interconnection needs. In that world, Northern Virginia’s differentiator was not simply fiber quantity; it was multi-carrier optionality that reduced deployment friction and shortened time-to-revenue.

What changed was how buyers valued the region. They weren’t only buying square meters of building or megawatts of power; they were buying the ability to connect to many networks and partners quickly, with fewer one-off construction projects and fewer single-provider dependencies. Optionality became a product.

The best-known historical marker of the hand-off is the rise of carrier-neutral colocation adjacent to early interconnection. Equinix’s founders chose in 1998 to build the company’s first data center in Ashburn, driven by MAE-East and a demand for a separate carrier-neutral building where carriers could interconnect and reach long-haul backbones³.

Once a region has enough infrastructure and activity, data centers cluster near each other to share information quickly, and an ecosystem of specialized firms and processes forms around repeated development. At that point, the advantage becomes self-reinforcing: being “near the meeting place” turns into being the meeting place.

Fundamentals

Ashburn’s dominance can be explained without turning the story into a fiber engineering spec. The core mechanism is connectivity density compounding: once enough carriers, cloud on-ramps, enterprises, and colocation providers are present in one place, the next participant gets lower friction simply by locating inside the cluster.

2 DUPONT FABROS TECHNOLOGY (ACQUIRED BY DIGITAL REALTY IN 2015) PURCHASED THEIR FIRST DATA CENTER FROM ABOVE NET CREDITORS FOR BELOW 10% OF REPLACEMENT COST, BEGINNING THEIR INVESTMENT IN THE DATA CENTER SECTOR.

3 EQUINIX’S DC2 DATA CENTER IN ASHBURN HOSTS MORE INTERCONNECTIONS THAN ANY OTHER FACILITY GLOBALLY.

What this buys, in practice, is not “more bandwidth” in the abstract. It buys simpler decisions. In a mature interconnection hub, a new entrant can assume that multiple carriers are already present, multiple routes are already available, and “connecting to everyone else” is a matter of standard facility work rather than a custom construction program. That is why Ashburn’s advantage keeps showing up as schedule confidence: the connective tissue is not a one-time asset, it is a standing capability.

That compounding shows up in three fundamentals:

First, optional connectivity is a product, not just an input. A site with one carrier connection has “a link.” A site with many carriers, routes, and counterparties has a marketplace: redundancy options, competitive pricing, and short paths to partners or customers. That marketplace is what makes a region sticky – because leaving it recreates cost and complexity elsewhere.

Second, carrier-neutral interconnection reduces bespoke construction. In a dense hub, networks can connect through standardized cross-connects in shared facilities rather than building one-off links to every counterparty. As counterparties multiply, the number of possible useful connections grows faster than linearly, which is why mature interconnection hubs tend to stay mature.

Third, the overbuild dividend is time. The stranded-fiber era left slack – conduit and unlit strands that later entrants could activate. The value is not only lower cost; it is reduced schedule risk. In the cloud era, schedule often dominates: missing a launch window can cost more than the premium paid for the “right” location.

The four historical acts map onto these fundamentals in a simple way. The setup created a meeting point. The build created physical optionality. The bust broke balance sheets but left the network layer in place. The hand-off is when the market learned to treat that inherited network layer as usable infrastructure rather than speculative surplus.

Global Examples

In the cloud era, schedule often dominates siting decisions. Missing a launch window can cost more than paying a premium for a location that is already connected, already legible to carriers, and already capable of delivering the “first day” network package without a bespoke construction project. That schedule logic is why connectivity hubs stay sticky even when other inputs - land, taxes, or short-term incentives - look more attractive elsewhere.

Northern Virginia (Ashburn): Capacity at Scale, Scarcity in the Present

Ashburn’s story is not “it grew fast once.” It is “it stayed easy long enough that growth kept choosing it.” The four-act sequence matters here because it explains durability: early

interconnection created a credible meeting point; aggressive fiber buildout created physical optionality; the bust reset economics without removing the assets; and the hand-off allowed cloud and colocation operators to scale on top of a network layer that was already dense and already routable.

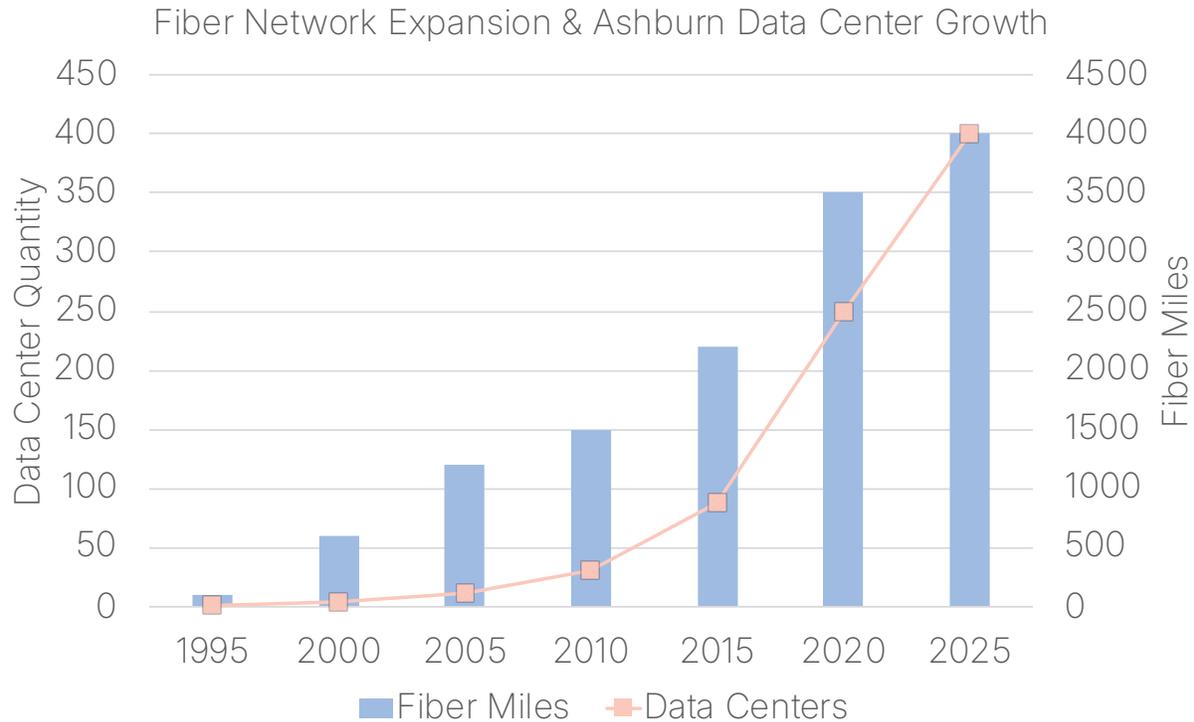


FIGURE A: NORTHERN VIRGINIA’S FIBER NETWORK EXPANDED WELL BEFORE THE REGION’S DATA CENTER INVENTORY. THE GAP BETWEEN THE TWO CURVES IS THE POINT: CONNECTIVITY CAPACITY ARRIVED FIRST, THEN BECAME AN ADVANTAGE LATER WHEN LARGE-SCALE COMPUTE DEMAND ACCELERATED AND OPERATORS COULD ACTIVATE EXISTING ROUTES INSTEAD OF REBUILDING THEM.

That durability shows up in an unintuitive pairing: massive capacity and persistent tightness. Even at today’s scale, market tracking still shows extremely low vacancy – evidence that the cluster is not just large, but still treated as the default place to land demand when time, counterparties, and execution risk matter. In parallel, the region’s earlier network expansion – fiber capacity arriving before data centers – helps explain the later “hockey-stick” build phase: the connective tissue was already in place when large-scale compute demand finally arrived.

The practical lesson is that interconnection hubs don’t just attract buildings; they attract commitments. Customers precommit because they are buying a launch schedule and a menu of connectivity options, not just square meters of white space.

Amsterdam: A Neutral Exchange as a Durable Attractor

Amsterdam illustrates a different path to the same outcome: rather than being defined by a single corridor's overbuild, it is defined by exchange density and the discipline of neutrality. A stable exchange ecosystem gives networks a reason to co-locate even when no single participant controls the venue. Over time, that "meeting place" becomes less about novelty and more about continuity: once many parties depend on the same rendezvous point, leaving creates friction and cost.

The modern indicator is not hype language; it is throughput. When an exchange like AMS-IX can still set peak traffic records in a mature market, it signals that the hub is functioning as infrastructure, not as a one-cycle artifact. The city remains attractive because participation itself compounds: every new entrant increases the value of everyone else being there.

The lesson Amsterdam offers Ashburn is simple: neutrality and stability are not branding. They are the conditions that let a marketplace keep forming inside physical infrastructure year after year.

Singapore: Cable Landings as a Connectivity Substrate

Singapore's advantage begins with geography and submarine infrastructure: dense international ingress and egress points create a base layer of optionality for networks and, by extension, for data centers. Where Ashburn's early story revolves around terrestrial interconnection and corridor routing, Singapore's revolves around being a practical junction for regional and global cables.

What matters is that cable density is necessary but not sufficient. Landings create the opportunity to be the place where traffic arrives and departs, but the conversion of that opportunity into data center scale still depends on local pathways for land, power delivery, and permitting. Put differently: cable landings can make a city "connected," but only the local build system can make it "buildable."

Singapore shows the same flywheel principle as Ashburn, but with a different first step: if you start with international optionality, the next competitive question becomes whether the city can reliably translate that optionality into deployable compute on predictable timelines.

Marseille: Cable Adjacency Turning Into Compute Adjacency

Marseille shows how cable concentration can pull data infrastructure activity inland. When landing points accumulate, the scarce resource shifts: it is no longer "can we get connectivity into the city?" but "where can we turn that connectivity into reliable compute, quickly, with enough power and enough operational control?"

This is the Ashburn pattern in miniature, with a different trigger. Cable adjacency creates the conditions for interconnection gravity, and gravity creates pressure for nearby compute. Over time, the ecosystem starts to look familiar: more carrier presence, more facilities that act as neutral meeting points, and more demand for repeatable expansion rather than one-off projects.

The lesson is not that every market should chase submarine cables. It is that optionality concentrates first, and compute follows where the conversion from connectivity to operational capacity can happen with the least friction.

Taken together, these examples reinforce the report's central claim: hubs form where optionality is concentrated and kept usable over time. The details differ – exchange governance in Amsterdam, cable geography in Singapore, landing concentration in Marseille – but the mechanism is consistent. And that consistency is what makes the next section unavoidable: if the pattern is so powerful, why hasn't it spread everywhere?

Why It Hasn't Scaled

The Ashburn pattern is repeatable in principle, but it is not "copy-paste" in practice – because the hardest parts are the slow, politically exposed ones.

Network effects are path-dependent. Interconnection hubs are not purely engineered; they are also agreed into existence. Carriers must believe a location will stay neutral enough, stable enough, and commercially rational enough to justify co-locating. That trust is built over years and can be broken faster than it can be built. Focusing on buildings and incentives without solving neutrality and friction tends to produce a site, not a flywheel.

There is also a coordination problem that incentives can't fully solve. A carrier does not want to be "first" unless it expects customers and peers to be present; customers do not want to commit unless they see credible route diversity and a menu of providers; and facility operators cannot justify building interconnection-forward space without both. Mature hubs escape this loop because early participation created enough daily utility that the ecosystem stopped needing to be sold and started being used. Most would-be hubs fail not because they lack ambition, but because they never cross the threshold where co-location becomes rational for multiple parties at once.

The "overbuild dividend" is hard to manufacture intentionally. Northern Virginia benefited from a historical anomaly: overbuilt fiber and conduit that survived the bust and got repriced. You can plan for slack and redundancy, but you cannot responsibly plan for a bankruptcy-driven fire sale as an economic development strategy.

That is why "copying Ashburn" is less about replicating a single asset and more about sustaining a sequence: convene participation, keep access neutral, and build enough slack that the next

participant can join without triggering a bespoke rebuild.

How It Could Scale

The transferable lesson from Ashburn is not “build more fiber.” It is design the system so connectivity can compound – then protect the conditions that let compounding continue.

In most places, the early phase looks less like a sudden leap and more like a bridge. The first facilities often “borrow” connectivity from a nearby hub – using regional backbones, cloud on-ramps, and a small number of carrier – while they prove that demand will show up reliably. The turning point comes when local interconnection becomes useful in its own right: when enough parties are close enough that exchanging traffic locally is cheaper, faster, and operationally simpler than routing everything back through a distant hub.

Replication begins with an anchor meeting point. In Northern Virginia’s early history, MAE-East and the push for carrier-neutral interconnection created a reason for networks to co-locate. Equinix’s origin underscores that neutrality was not a branding choice; it was a structural solution to scalability and control concerns around exchange infrastructure.

Protecting neutrality is not just a statement of intent; it is a pattern of behavior that reduces fear of capture. It means clear, predictable rules for access, cross-connects, and pricing; it means making it easy for new carriers to enter without negotiating bespoke terms each time; and it means resisting the temptation to turn the “meeting place” into a bottleneck controlled by any single participant. When neutrality is credible, participation becomes self-reinforcing – and once participation is self-reinforcing, incentives matter less.

From there, scaling requires treating fiber and interconnection as long-lived civic/industrial infrastructure: build with slack, enforce neutrality, and make it operationally easy for the next participant to join without bespoke negotiation every time.

The playbook also has a quiet but decisive ingredient: repeatable execution capacity. Over time, clusters develop specialized contractors, standardized build patterns, and permitting muscle memory. JLARC explicitly points to clustering dynamics and an ecosystem of specialized firms around data centers in Northern Virginia.

What Changes Next

Power becomes the dominant siting argument, even in fiber-rich regions. As power delivery timelines stretch, growth pressure can shift to adjacent or emerging markets that can deliver energy sooner, even if they have to “borrow” interconnection from a nearby hub at first.

Auditable metrics will matter more than slogans. As public scrutiny rises, the industry will benefit

from using measurable anchors (capacity, vacancy, delivered MW, grid upgrades) rather than viral traffic claims.

Permitting and community negotiations become part of the infrastructure stack. The public process is now a gating variable in ways it wasn't when fiber was the main bet. Regions that want Ashburn-like compounding will need clearer benefit framing, more legible mitigation, and better-aligned infrastructure planning – or the system will stall even if the fiber is present.

Taken together, these shifts change what “infrastructure advantage” looks like. In the 1990s corridor era, the scarce resource was routes and interconnection. In the next build cycle, the scarce resource is increasingly the ability to deliver power and approvals on a predictable timeline – and to demonstrate, with plain numbers, what a community receives in exchange. The hubs that stay dominant will be the ones that can translate structural advantages into visible outcomes: faster delivery, fewer single-point dependencies, and a public story that can be audited rather than argued.

Closing

Ashburn's dominance is not a mystery and it is not magic. It is the long shadow of a network-era overbuild that survived long enough to be re-owned, repriced, and repurposed – then amplified by carrier-neutral interconnection and cloud-era clustering.

ASHBURN DIDN'T WIN BY BUILDING THE MOST BUILDINGS FIRST – IT WON BY BECOMING THE EASIEST PLACE FOR NETWORKS TO MEET, AND BY STAYING BUILDABLE LONG AFTER THE FIRST BOOM ENDED.

The balanced lesson is also the useful one: dense connectivity is hard to create quickly, but easy to waste if power, land, and governance don't keep pace. Northern Virginia's lead has persisted because the region repeatedly translated structural advantage into operational outcomes customers could feel – optional connectivity, resilience, and speed to market – while continuing to execute the unglamorous parts of infrastructure delivery.



Key Takeaways

- ▶ Ashburn's dominance is best explained as a four-act handoff: early interconnection setup, fiber buildout, financial bust, and a later hand-off to cloud and colocation scale.
- ▶ The telecom downturn broke balance sheets, but the physical network layer largely remained – and later got repriced into an advantage.
- ▶ Carrier-neutral interconnection converts a cluster into a marketplace, reducing bespoke connectivity work through standardized local cross-connects.
- ▶ "Fiber exists" is not the differentiator; fiber optionality (multi-route, multi-carrier choice) is what reduces deployment friction.
- ▶ Equinix's 1998 Ashburn decision illustrates how neutrality and exchange adjacency can become a long-lived geographic anchor.
- ▶ The most defensible "dominance" metric is capacity, not viral traffic slogans; JLARC reports Northern Virginia at 13% of global operational capacity and 25% of the Americas.
- ▶ CBRE's Q1 2025 figures (3,046.1 MW inventory; 0.76% vacancy) show a market that remains tight even at massive scale.
- ▶ The replication mistake is focusing on land and buildings while underinvesting in neutrality, route diversity, and interconnection governance.
- ▶ In 2026, power availability and delivery timelines increasingly determine where capacity can actually be delivered, even in connectivity-rich regions.
- ▶ Global analogs (Amsterdam, Singapore, Marseille) reinforce the pattern: hubs form where optionality is concentrated and kept usable over time.

Glossary

CAPACITY (MW)

In data center market reporting, “capacity” typically refers to critical IT load expressed as power – megawatts (MW) – rather than floor area. Definitions vary (operational vs under construction vs planned), which is why sources should be explicit about what they count. A useful relationship: 1 MW = 1,000 kW.

CARRIER-NEUTRAL

A facility or campus design/operating model that allows multiple network providers to offer services on equal terms, without a single carrier controlling access. Carrier neutrality reduces dependency risk and increases optionality for customers who want redundancy or competitive pricing.

CROSS-CONNECT

A physical connection (often fiber) inside a data center or interconnection facility that links two parties (for example, two networks, or a customer and a cloud on-ramp). Cross-connects matter because they turn “nearby” into “cheap and fast to connect.”

DARK FIBER

Installed fiber optic strands that are not currently lit (not carrying active optical signals). Dark fiber can be a strategic asset because it represents latent capacity that can be activated without rebuilding conduit or routes.

DWDM (DENSE WAVELENGTH-DIVISION MULTIPLEXING)

A technology that increases fiber capacity by sending multiple wavelengths (channels)

of light through a single fiber strand. For non-specialists: DWDM is one reason fiber networks built in the 1990s could later carry dramatically more traffic without new trenching.

IXP (INTERNET EXCHANGE POINT)

A physical location where networks interconnect to exchange traffic, often via “peering” arrangements. IXPs reduce latency and transit costs by keeping traffic local between networks that would otherwise route through more distant paths.

JLARC (JOINT LEGISLATIVE AUDIT AND REVIEW COMMISSION)

Virginia’s nonpartisan legislative oversight agency that publishes research reports on state policy issues. Its “Data Centers in Virginia” materials are a commonly cited public reference for the scale and drivers of Northern Virginia’s current market position.

MAE-EAST (METROPOLITAN AREA EXCHANGE – EAST)

A historically important set of interconnection locations on the U.S. East Coast that helped carriers exchange traffic in the early commercial internet era. It is frequently cited as part of Northern Virginia’s early interconnection gravity.

TBPS (TERABITS PER SECOND)

A throughput unit used to describe exchange traffic volume at IXPs. 1 Tbps = 1,000 Gbps (gigabits per second).

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

AMS-IX. "AMS-IX: worldwide internet milestones in 2024."

Accessed January 2026.

<https://www.ams-ix.net/ams/documentation/infographics/ams-ix-worldwide-internet-milestones-in-2024>

CBRE. "Global Data Center Trends 2025."

Accessed January 2026.

<https://www.cbre.com/insights/reports/global-data-center-trends-2025>

Carnegie Endowment for International Peace. "Subsea Communication Cables in Southeast Asia: A Comprehensive Approach is Needed."

Accessed January 2026.

<https://carnegieendowment.org/research/2024/12/southeast-asia-undersea-subsea-cables?lang=en>

Data Center Knowledge. "Equinix Doubles Down in One of Internet's Most Important Locations."

Accessed January 2026.

<https://www.datacenterknowledge.com/data-center-site-selection/equinix-doubles-down-in-one-of-internet-s-most-important-locations>

DE-CIX. "2024 highlights from DE-CIX."

Accessed January 2026.

<https://www.de-cix.net/en/about-de-cix/news/2024-highlights-from-de-cix>

Federal Reserve Bank of Richmond. "Boom and Bust in Telecommunications."

Accessed January 2026.

https://www.richmondfed.org/publications/research/economic_quarterly/2003/fall/couperhejkalwolman

JLARC. "Data Centers in Virginia."

Accessed January 2026.

<https://jlarc.virginia.gov/landing-2024-data-centers-in-virginia.asp>

Optics & Photonics News. "Boom, Bubble, Bust: The Fiber Optic Mania."

Accessed January 2026.

https://internethistory.org/wp-content/uploads/2020/01/OSA_Boom.Bubble.Bust_Fiber.Optic_.Mania_.pdf

Orange. "Orange hosts the Medusa submarine cable at its infrastructure in Marseille for its first landing."

Accessed January 2026.

<https://newsroom.orange.com/orange-hosts-the-medusa-submarine-cable-at-its-infrastructure-in-marseille-for-its-first-landing/>

Securities and Exchange Commission. "Metromedia Fiber Network, Inc. Form 8-K Exhibit 99.1 (June 11, 2002)."

Accessed January 2026.

<https://www.sec.gov/Archives/edgar/data/1043533/000095014202000608/ex991frm8k61102.txt>

TeleGeography. "Does 70% of the World's Internet Traffic Flow Through Virginia?"

Accessed January 2026.

<https://blog.telegeography.com/does-70-of-the-worlds-internet-traffic-flow-through-virginia>

The Associated Press. "Virginia county approves data center project after 27-hour public hearing."

Accessed January 2026.

<https://apnews.com/article/ea3d26b1746d251909dc1e9d5f9034d2>

© 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

15 minutes



Release Frequency

Monthly



Audience

Digital Infrastructure



Distribution Format

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.