GLOBAL NETWORKING

Independent Network Infrastructure Advisory

WHITE PAPER · VOL. 01

Hyperscaler Datacenter Strategies

Metro Clusters vs. Photonic-Remote Deployment

Why moving photons is becoming cheaper, cleaner, and more scalable than moving electrons — and what that means for where the next generation of hyperscale compute actually gets built.

AUTHOR

Lorin

Principal, Global Networking · ld@globalnetworking.ai

PUBLISHED · APRIL 2026

Executive Summary

Hyperscale cloud operators (AWS, Google, Microsoft, Meta, Oracle) are siting tens of billions of dollars of new compute capacity over the next three years. The default answer for the last fifteen years has been to build on the edges of major metros — Ashburn, Santa Clara, Dallas — where existing fiber is dense and users are close. That default is breaking.

The reason is simple physics. Computation is energy intensive. Moving bits is not. As rack densities climb, the bottleneck on a new hyperscale campus is no longer floor space — it is megawatts of stable, affordable, clean power. The metros that host today's datacenter alleys are running out of the only thing that actually matters at scale.

> **THE THESIS**
>
> *Moving photons (data) is cheaper, greener, and more scalable than moving watts (electricity). Put the compute next to the power source and ship the bits to the user — don't ship megawatts of grid power to the user and then burn it in an urban compute site.*

This paper compares the two strategies — metro clusters and photonic-remote deployment — across grid efficiency, energy per bit, land and infrastructure cost, processing density, and sustainability. The evidence, from hyperscale operators' own choices and from the peer-reviewed literature, points decisively in one direction. Metro will not disappear; latency-sensitive services still belong there. But the next wave of bulk compute — AI training, batch workloads, storage replication — is going to follow the power, not the other way around.

Strategy 1 — Metro Clusters

Description. Hyperscalers build on the outskirts of major metros — Northern Virginia, Silicon Valley, Dallas–Fort Worth, Chicago, Phoenix. These campuses take advantage of proximity to end users and to internet backbone hubs, but they must draw their electrical power from the regional grid. Generation is off-site, often hundreds of miles away. Urban land is expensive, so the facilities go high-density: multi-story halls, liquid cooling, very aggressive rack packing to maximize compute per square foot.

Power Supply and Grid Efficiency

Metro datacenters are consumers of grid electricity, not co-located with generation. The utility delivers power from distant plants via the transmission and distribution network. On average, about 5% of US electricity is lost as heat in T&D — so a 100 MW datacenter actually consumes roughly 105 MW at the source. That inefficiency is small, persistent, and entirely avoidable when the load sits next to the generator.

Urban grids can also be congested. Northern Virginia has publicly flagged power availability as the binding constraint for new construction, and utilities there have floated the possibility of moratoria. Metro sites offer robust infrastructure and redundancy, but at the cost of higher demand charges, competition with every other commercial load in the market, and real stress on regional transmission.

Energy Per Bit — Processing vs. Data Movement

A metro strategy minimizes data travel distance for local users, but it does not eliminate data movement energy costs. Internal networking inside a large cluster is itself meaningful, and large deployments still replicate data across regions. Within a single datacenter, moving bits between servers and storage can rival computation in energy consumed.

The important asymmetry is at the long-haul layer. Per-bit processing energy for intensive workloads — AI training, dense inference — runs in the picojoule to nanojoule range. Sending a bit over optical fiber runs in femtojoules. Fiber transport is three to six orders of magnitude more efficient than the compute it feeds. The metro strategy does not benefit from this asymmetry; it still has to bring every joule of compute power in through the grid.

Land, Real Estate, and Infrastructure

Land cost in major metros is the single largest force shaping the physical design of a metro campus. In Loudoun County (Ashburn, VA — Data Center Alley), parcels that sold for around $300,000 per acre a decade ago now trade at $2–3 million per acre. Even secondary markets near cities price prime datacenter land at $200,000–$500,000 per acre. Multi-story construction and aggressive density become necessary simply to justify the land cost.

Beyond land, metro builds face higher labor costs, more protracted permitting, and growing community pushback on noise, water use, and perceived tax shifting. Metros also require substantial investment in grid infrastructure — new high-capacity transmission lines and substations — to serve new load. Northern Virginia's cluster has necessitated significant utility upgrades even as operators enjoy energy rates that sit roughly 28% below the US commercial average. Those subsidies are already under political scrutiny.

Processing Power Density

Because urban land is scarce and expensive, metro facilities tend to push compute density as hard as the cooling plant will allow. Racks at 20–40 kW are standard. Liquid cooling and rear-door heat exchangers are increasingly common. Some urban sites are multi-story. Power draw can exceed 100 MW on a single campus, dissipating enormous heat in a small footprint. Redundancy at that density — UPS, standby generation, on-site fuel — is expensive and operationally demanding. Remote sites can spread the same compute across a much larger footprint at lower per-square-foot stress; metro sites must engineer their way to the same load.

CASE STUDY —NORTHERN VIRGINIA

Loudoun County is the canonical example of the metro strategy. More than 118 facilities sit across roughly 43 million square feet, underpinned by dense fiber and favorable power rates. The power itself, though, is not generated in Loudoun — it rides in on PJM, a grid that has historically leaned on coal and is still fossil-heavy overall. Each megawatt delivered to Ashburn carries a larger carbon footprint than, say, a megawatt of Pacific Northwest hydro or California solar. Ashburn's latency and economic ecosystem are real; its energy profile is not the one you would design from scratch.

Strategy 2 — Remote Photonic Datacenters

Description. Put the compute next to the generation. Rural parcels with cheap land and abundant renewable supply — the Columbia Basin, the Panhandle wind belt, Iowa, the Dakotas, rural Ontario and Quebec — become the anchor. Fiber replaces transmission as the primary linkage back to users. The core idea is *move bits, not watts*: rather than transporting vast electrical power over the grid to an urban compute site, transport the digital workload to a site with ample power and then ship results back at light speed.

Photonic Transport Efficiency

Modern optical fiber carries terabits of data with almost trivial incremental energy cost. Sending a single bit optically can require as little as femtojoules of energy; sending the same bit over copper uses picojoules, several orders of magnitude more. Long-haul links add overhead for lasers, amplifiers, and regeneration every ~80 km, but the per-bit energy cost over continental fiber still comes out well below a nanojoule per bit per kilometer. That is negligible compared with the joules of compute it takes to produce that bit in the first place.

> **THE ASYMMETRY**
>
> *Computing a bit is expensive. Moving a bit is almost free. Any strategy that trades the cheap thing for the expensive thing wins.*

In a peer-reviewed analysis of China's "Eastern Data, Western Computing" initiative, authors quantified this directly: migrating compute to renewable-rich western regions saved energy because the cost of transmitting data was far lower than the loss of transmitting equivalent electrical power. The same logic applies to North America, where the renewable-rich regions sit in the Pacific Northwest, the wind belt, the Midwest, and select hydro-rich provinces.

Grid Efficiency and On-Site Power

A remote photonic datacenter is typically co-located with generation or tied directly into high-voltage transmission at a plant gate. Transmission losses effectively collapse to zero at the site level. A datacenter built adjacent to a Columbia River dam sees electrons move tens to hundreds of meters from turbine to server — not hundreds of miles from plant to city. The wider grid benefits too: regions with surplus generation (Pacific Northwest hydro at spring runoff, West Texas wind at night, California solar at noon) can absorb load that would otherwise be curtailed.

A UCSB study modeled one version of this: if California midday solar surplus were used to run data jobs relocated from coal-heavy grids, it could absorb up to 600,000 MWh of otherwise-curtailed clean energy per year. That only works because the bits can be moved easily and the watts cannot.

Real Estate and Infrastructure

Land is where the remote strategy quietly wins every time. DFW-adjacent datacenter parcels run $50,000–$150,000 per acre — roughly 5% of Loudoun County pricing. Truly remote parcels can be cheaper still, and campus-scale acquisitions (hundreds of acres at a time) are routine for forward expansion. Rural jurisdictions typically offer tax abatements and power-rate incentives to attract investment. Buildings can be single-story warehouse-style, which shortens schedules and lowers per-MW construction cost.

Grant County, Washington sells datacenter power at roughly $0.03–$0.04 per kWh directly from hydroelectric supply — approximately half of even discounted Northern Virginia pricing and well below the U.S. commercial average. Over a decade of 24/7 operation, that delta compounds into hundreds of millions of dollars on a single hyperscale campus.

The offset is fiber. A remote site needs diverse, high-capacity fiber routes to at least two major internet exchange points, ideally three. Hyperscalers build or lease this — increasingly in the form of owned long-haul routes backed by IRUs on diverse physical paths. Fiber infrastructure replaces transmission infrastructure as the primary capital investment, and at modern DWDM capacities it is an excellent trade.

Cooling and Climate

Remote siting allows strategic placement in climates that do most of the cooling work for free. The Pacific Northwest, eastern Washington, Minnesota, Iowa, and the Dakotas all offer extended periods where ambient air or evaporative cooling alone is sufficient. PUEs below 1.2 are routine; Nordic-style builds push below 1.1. Metro sites, in contrast, often sit in warmer latitudes or urban heat islands where free cooling contributes very little. The cooling delta alone can be worth 10–15% of total facility energy.

CASE STUDY — QUINCY, WASHINGTON

Quincy is the canonical hydro-powered hyperscale hub. Microsoft, Sabey, Yahoo, and others built large campuses on cheap farmland specifically to tap abundant reliable hydropower at roughly $0.038/kWh. The climate is semi-arid and cool in winter. The fiber ties back to Seattle and onward to California at 2–5 ms and 10–15 ms respectively — negligible for almost all cloud workloads. Sabey's Quincy campus runs PUE ~1.2, and backup generators fire only about eight hours per year because the underlying hydro base load is essentially always available. The tax revenue has paid for schools and hospitals. This is the model working at maturity.

CASE STUDY — CALIFORNIA EXCESS SOLAR

A modeled case study asked what happens if workloads are shifted dynamically from coal-heavy grids to California datacenters during solar overproduction. The answer: roughly 62% of otherwise-curtailed solar could be absorbed into compute. The operators get near-free electricity; up to 240,000 metric tons of CO_2 per year are avoided. The only real barriers identified were organizational — cross-region workload scheduling and commercial trust between operators — not technical. The physical infrastructure already exists.

CASE STUDY —"EASTERN DATA, WESTERN COMPUTING"

China's national initiative is this thesis at country scale. An assessment found that moving compute from eastern megacities to western renewable-rich regions cut datacenter energy use by 4.8%–12.5% annually via cooling savings and eliminated transmission losses. Where the remote power was cleaner — Shanghai (coal) to Sichuan (hydro) — emissions fell by ~80%. Where the remote power was itself coal-based — Beijing to Inner Mongolia in the wrong corridor — emissions rose by ~25%. The lesson is not that the strategy fails; the lesson is that site selection inside the strategy matters enormously.

Comparative Analysis

The table below distills the difference between the two strategies across the dimensions that actually drive siting decisions for hyperscale operators today.

FACTOR	Metro Clusters	Remote Photonic Datacenters
GRID EFFICIENCY & POWER DELIVERY	Relies on long-distance power transmission from plants to city. ~5% T&D losses on average. Urban grids can be congested; large datacenters strain local substations (Northern Virginia saw power constraints as demand surged). Power drawn from regional mix, which can include high-carbon sources.	Draws power on-site or adjacent to generation. Virtually no long-distance loss. Offloads strain on urban grid. Can be scheduled against surplus renewable output, absorbing power that would otherwise be curtailed. Requires robust fiber rather than additional transmission.
ENERGY PER BIT — COMPUTE VS. MOVEMENT	Majority of energy goes to computing and cooling. Local networking is cheap per bit, but every joule of compute power must still be delivered electrically. Computation is on the order of pico- to nanojoules per operation; moving bits within a metro region is a minor overhead. Overall energy per task is hostage to the grid mix.	Moving data over fiber is exceptionally efficient — femtojoules to picojoules per bit at the photonic layer. Shipping large datasets to a remote site barely dents the energy budget. The remote site then uses locally-generated power more efficiently for compute. A few kWh spent sending data can save tens of kWh in better power use on the compute side.
REAL ESTATE & INFRASTRUCTURE	High land and construction cost. Metro land in Data Center Alley now trades at roughly $2–3M per acre. High-density, multi-story builds increase design complexity. Advanced cooling and redundant electrical plant drive capex. Community and zoning friction is real. Upside: dense existing fiber and short hops to exchanges.	Low land cost and easier expansion. Rural parcels near DFW-style markets trade at ~$50–150k per acre; truly remote parcels are cheaper still. Single-story warehouse designs with modular builds. Fiber buildout replaces transmission buildout as the primary infrastructure investment,

FACTOR	Metro Clusters	Remote Photonic Datacenters
		and many rural jurisdictions offer tax and power-rate incentives.
PROCESSING DENSITY	Pushed to maximum because space is expensive. Very high rack density (tens of kW/rack), sometimes vertical stacking. More compute per acre but also concentrated heat and heavier cooling loads. Once a site is filled, adding capacity means acquiring more expensive land or building up.	More freedom to distribute load. Moderate power density across a larger footprint enables simpler cooling (ambient air, evaporative). Hundreds of MW can sit across a multi-building campus. Scale flexibility is much higher — pre-purchasing adjacent parcels is routine. Slightly longer internal data paths, handled easily by modern fiber.
SUSTAINABILITY & CARBON	Dependent on the metro grid mix. Many of the largest US metro hubs sit on grids that are 60%+ fossil-fueled today. Urban sites lack space for meaningful on-site renewables; sustainability programs rely on PPAs and credits to offset usage rather than truly green on-site power. Waste-heat reuse is possible but rare in North America.	Can be sited directly onto low-carbon generation: hydro, wind, solar, geothermal. Real-world studies show emissions reductions of ~80% when compute moves from coal-heavy grids to hydro-rich regions. Cooler climates further reduce cooling energy. The caveat: a remote site near a coal plant offers none of these benefits — location selection matters.

Discussion and Implications

The comparison supports the thesis: moving photons to the power source is more energy-efficient and more sustainable than moving power to the data source. Photonic transport is so efficient per bit that continental distance becomes almost irrelevant. Generating power far from where it is used still incurs real losses, and it locks the compute into whatever the metro grid mix happens to be.

Grid

If even a meaningful fraction of North American hyperscale growth shifted to power-rich areas, the reduced strain on urban grids would improve overall grid stability and defer new urban transmission. US grid operators already project tens of thousands of miles of new transmission needed by 2035 to meet projected demand. A credible portion of that projection assumes load growth in exactly the metros that are hardest to serve. Siting compute at generation directly addresses that.

Compute vs. Transport

As AI and big data workloads grow, compute energy is scaling faster than anything else on the operator's balance sheet. Network technology is moving the other direction — more bits per watt, every generation. The gap between those two curves is the entire business case for photonic-remote siting. Operators should treat geographic load balancing as an energy optimization tactic, not just a reliability or cost tactic. If a workload can run in a region at half the carbon per kWh, the transport cost and the handful of milliseconds of added latency are cheap compared with what they buy.

Cost

A concrete example: 100 acres in a rural Pacific Northwest or Midwest market at $50,000/acre costs $5M. The same 100 acres in Ashburn costs $250–300M. That $245–295M delta pays for a lot of fiber, a lot of substation work, and most of the capacity build for Year 1. The trade-off is latency, and for the majority of bulk compute workloads — training, batch analytics, storage replication, rendering, encoding — a few dozen extra milliseconds is not a constraint.

Sustainability

Datacenters are ~1–2% of global electricity today and credibly headed to 8–10% over the next decade. Whether that load lands on clean or dirty generation is the single biggest lever the industry has on its own carbon trajectory. Buying renewable PPAs to offset a metro datacenter does not physically change what the local grid burns at 5pm on a hot Tuesday. Siting the load on the renewable asset does.

Honest Challenges

Photonic-remote is not free. Three real challenges: first, network reliability becomes as operationally critical as power reliability, which requires genuine route diversity and not the carrier-sales-slide version of diversity. Second, cross-regional workload scheduling requires organizational maturity that not every operator has today. Third, data residency and compliance requirements restrict certain workloads geographically. The strategy is not about moving 100% of compute to rural campuses; it is about recognizing that the growth marginal workload — especially AI training — has very different siting constraints than the legacy latency-sensitive services that anchored the metro build-out.

Conclusion

Building in expensive metro areas made sense historically for connectivity and latency. The maturity of optical networks and the pressing need for grid efficiency and clean compute have shifted the calculus. By placing datacenters near sources of cheap, clean power and using photonic links to transport workload, hyperscalers can achieve higher grid efficiency, lower per-bit energy cost, materially lower real estate and infrastructure expense, and large sustainability gains — provided the remote power is itself clean.

Metro-adjacent datacenters will not disappear. They remain essential for latency-sensitive services and network-hub functions. The optimal architecture is tiered: keep a footprint in metro for edge and caching, but shift bulk compute — AI training, non-urgent batch, storage replication, rendering — to large-scale remote campuses engineered from the ground up for efficiency. This is already the direction of travel. The next five years of hyperscale capacity additions will decide how far and how fast the industry commits to it.

> **THE BOTTOM LINE**
>
> *The next generation of hyperscale compute will be built on the foundation of energy, not the foundation of geography. Follow the photons.*

Sources

01 Tasoff, H. (2020). Moving Bits, Not Watts. UCSB Current.

02 Zhang, Y. et al. (2025). Moving bits vs. moving watts — Eastern Data / Western Computing. Applied Energy.

03 Datacenters.com (2024). Impact of Geographic Location on Data Center Energy Costs.

04 Datacenters.com (2025). Top 5 Data Center Land Markets.

05 Sabey Data Centers (2023). SDC Quincy — Low-Cost Hydropower.

06 UPSTACK (2025). Why Ashburn is Data Center Alley.

07 UCSB / Joule via Tasoff (2020). Using excess solar for datacenters.

08 U.S. Energy Information Administration (2023). Transmission & Distribution losses.

09 Datacenters.com (2024). Climate and Cooling — Free Cooling vs. Cooling Costs.

10 UCSB Current (2020). PJM coal vs. California renewables example.

ABOUT GLOBAL NETWORKING

Global Networking is an independent network and fiber infrastructure advisory.

Contact · ld@globalnetworking.ai · globalnetworking.ai