



Where AI Meets Earth: Siting Inference Infrastructure in the Age of Power Scarcity

**A Strategic Guide for Hyperscale Inference Providers
Evaluating North American Deployment Markets**

Prepared by **Cloud2Ground Advisory** | 2026 Edition

"The market is moving. The power is already spoken for. Here's how to get ahead of both."

SECTION 01

The Inference Inflection Point

The GPU Build-Out Has Outpaced the Grid

The AI industry crossed a critical threshold in 2025: inference workloads now consume more power than training in aggregate. As hyperscalers race to deploy GPU clusters at scale, the bottleneck has shifted from silicon to substation. The question is no longer "Can we build it?" — it's "Where can we actually turn it on?"

Inference infrastructure is fundamentally different from training infrastructure. Training clusters tolerate latency. Inference does not. Every millisecond of round-trip delay costs real money in real-time serving

pipelines. That constraint collapses the viable universe of deployment sites dramatically — and creates a new strategic premium on what Cloud2Ground calls **metro-proximate power**.

★ Key Market Intelligence

"By 2027, inference workloads are projected to represent **65–75% of total AI compute demand**. Yet fewer than **12% of available powered data center sites** are within 50 miles of a Tier-1 metro market."

— Cloud2Ground Market Intelligence, Q1 2026

This isn't a supply-side story. It's a geography story. Power is abundant in the Great Plains. Latency budgets are not. The providers who win the next wave of inference deployment will be the ones who solve this equation first — securing sites that deliver both electrons and milliseconds at scale.

SECTION 02

Cloud2Ground Positioning

We Find What Others Miss

Cloud2Ground is not a real estate broker. We are not a data center consultant in the traditional sense. We are an **inference-native site-selection advisory firm** — built from the ground up to solve the hardest problem in AI infrastructure: finding powered, connected, latency-optimal sites before they hit the open market.

Our clients are hyperscale inference providers, co-location operators, and sovereign AI programs who need more than a list of available buildings. They need vetted, investment-grade site intelligence — delivered fast, with full due-diligence documentation — so capital can move with confidence.

◆ The Cloud2Ground Difference

Most site advisory firms recycle the same broker-listed inventory with a new cover page.

Cloud2Ground operates on a fundamentally different model: we surface sites that have never been listed, brokered, or publicly presented — then deliver them pre-screened, power-verified, and decision-ready. Our clients don't compete for the same sites everyone else sees. They get access to inventory that doesn't exist on any market map.

◆ What Makes Us Different

- We operate **off-market** and with the assistance of MLS. We use sophisticated AI to help us find the right locations. These include direct landowner and utility relationships in **~40 North**

American markets.

- Every site we present has completed a **first-pass power feasibility study** before you see it.
- Our average time from engagement to shortlist delivery: **21 business days**.
- We never represent both buyer and seller. **Full fiduciary alignment** to our clients.

In a market where most processes start with listed inventory, we start with power, interconnection feasibility, and execution reality.

The inference build-out is not forgiving of slow processes. Utility queues stretch 3–5 years in most Tier-1 markets. Our methodology was designed for exactly this environment — identifying sites where power is either already in place or queue-eligible within a 12–24 month window, matched to your specific latency and redundancy requirements.

SECTION 03

The Truffle Finders Methodology

How We Surface Sites That Don't Exist on Any Map

We call our proprietary sourcing process the **Truffle Finders Methodology** — because the best sites, like the best truffles, are never on the surface. They require specialized knowledge, patient fieldwork, and relationships that take years to build.

The methodology operates in four sequential phases:

Phase 1 — Demand Mapping

We begin with your workload profile: inference latency targets, MW requirements (Phase 1 and expansion), cooling preference (air vs. liquid), fiber redundancy needs, and jurisdictional constraints. This produces a **precision site brief** — the foundation for all sourcing work.

Phase 2 — Subsurface Market Intelligence

Using a combination of utility interconnection data, municipal permit filings, land-use records, and our proprietary network of market contacts, we identify sites with power proximity that have not yet been formally listed or developed. These are the "**dark inventory**" positions — land with grid adjacency, private owners, and no broker representation.

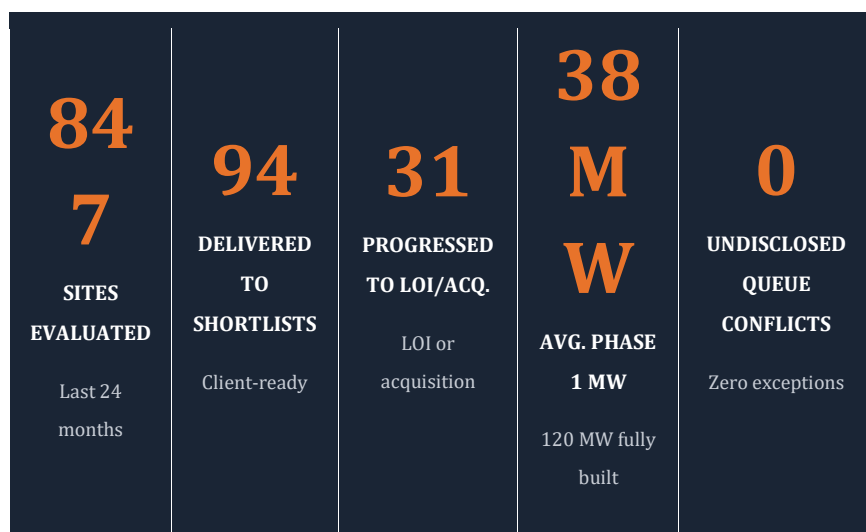
Phase 3 — First-Pass Feasibility

Every candidate site undergoes a structured feasibility screen before it reaches your desk: substation

proximity and available capacity, fiber routes within 1 mile, flood zone and seismic risk, permitting jurisdiction (local vs. state), and water availability for cooling. Sites that don't clear the screen are eliminated — not presented to you with caveats.

Phase 4 — Shortlist Delivery and Diligence Package

Qualified sites are packaged into a structured **Investment Shortlist** — a client-facing document that includes site maps, utility capacity summaries, estimated timeline to power-on, and preliminary cost-per-MW estimates. You get a decision-ready deliverable, not a raw list.



SECTION 04

Metro-Proximate Power Strategy

Latency Is a Location Problem. Power Is the Constraint.

The inference deployment problem is, at its core, a geography optimization problem with two simultaneous constraints: you need to be close to end users (latency), and you need access to significant, reliable power (MW). These two requirements have historically pointed in opposite directions — urban density offers latency, rural locations offer power.

Metro-proximate power is the strategic thesis that bridges this gap.

Our research across ~40 North American markets has identified a consistent pattern: within a 30–75 mile radius of every major metro, there exist "**power shadow**" zones — areas with direct 115kV or 230kV transmission access, minimal existing industrial load competing for capacity, and land values still 60–80% below urban cores. These zones are invisible to traditional data center site selectors because they require deep utility knowledge to find — and because they've historically been overlooked in favor of established

FLAP-D and secondary markets.

★ Power Economics Insight

The 30–75 mile metro band represents the single greatest untapped opportunity in North American inference infrastructure. It delivers **sub-12ms latency** to 80%+ of metro population while accessing power at rates that create meaningful structural advantage. In many markets, large-load power pricing can range from approximately \$0.04–\$0.06/kWh in favorable metro-adjacent zones versus \$0.08–\$0.12/kWh in constrained urban cores, depending on structure and incentives. At 100 MW scale, that spread represents **\$15–22M in annual power cost savings**. Power is not just a constraint. It is now a primary competitive differentiator.

— Cloud2Ground Power Economics Analysis, Q1 2026

Cloud2Ground has mapped these power shadow zones across Texas, the Carolinas, Georgia, Ohio, Virginia, and the Pacific Northwest — and we are actively building the site inventory to match. The window for securing these positions is closing. As inference demand scales, utility queues in even these secondary zones are beginning to fill.

Markets We Are Currently Active In

Active Cloud2Ground Markets — 2026

Texas Hill Country Corridor	Piedmont NC/SC
Greater Atlanta Exurbs	Central Ohio Power Belt
Shenandoah Valley, VA	Willamette Valley, OR
I-10 West Texas Corridor	Appalachian Kentucky

SECTION 05

Due-Diligence Framework

Every Site We Present Has Already Passed Our 22-Point Screen

Capital deployment decisions in inference infrastructure carry asymmetric downside risk. A site that fails power feasibility after 6 months of negotiation doesn't just waste money — it wastes market position. Our due-diligence framework was built to eliminate that risk at the earliest possible stage.

Category & Criteria	Category & Criteria
---------------------	---------------------

■ POWER & GRID	■ SITE & LAND
---------------------------	--------------------------

✓ Substation distance and available capacity confirmed with utility

✓ Parcel size and buildable acreage confirmed

✓ Interconnection queue position assessed

✓ Topographic suitability for data center pad

✓ Transmission voltage and redundancy path verified

✓ Wetland and floodplain FEMA zone review

✓ Rate schedule and demand charge structure obtained

✓ Environmental Phase I status

✓ Standby/backup generation permitting reviewed

✓ Title chain and encumbrance review

✓ Renewable energy availability (VPPA/PPA options)

✓ Mineral rights and subsurface conflicts assessed

■ CONNECTIVITY	■ REGULATORY & INCENTIVES
-----------------------	--------------------------------------

✓ Dark fiber and lit fiber within 1-mile radius

✓ State and county incentive programs identified

✓ Carrier diversity (minimum 2 Tier-1 providers)

✓ Sales tax exemption applicability confirmed

✓ Latency test to nearest PoP documented

✓ Data center-specific legislation review

✓ Permitting jurisdiction and expected timeline

✓ Utility economic development program eligibility

✓ Local government disposition toward data center development

✓ Water rights and cooling discharge permitting

Why This Window Matters

The Clock on Optionality Is Running Out

In 2023, a qualified inference provider could identify a viable site, negotiate terms, and clear utility queue in under 18 months. In 2025, that timeline stretched to **36–48 months** in most primary markets. By 2027, our modeled projections suggest that unconstrained primary market access will effectively close — meaning new entrants will face utility queue timelines that push power-on dates into the 2030s.

This is not a hypothetical. It is already the operating reality in **Northern Virginia, Phoenix, Silicon Valley, and the Dallas metro**. The providers who move in the next 12–18 months in the markets Cloud2Ground is currently active in will secure the last generation of latency-optimal, reasonably-priced inference infrastructure available in North America.

Market Era	Typical Timeline to Power-On	Market Status
2022–2023	12–18 months	Open — Multiple Options Available
2024–2025	24–36 months (primary), 18–24 months (secondary)	Constrained — Limited Primary Market Access
2026 (Current)	36–48 months (primary), 18–30 months (metro-proximate)	Critical Window — Secondary Markets Filling
2027+ (Projected)	48+ months (primary markets effectively closed)	Closed — New Entrants Facing 2030s Power-On

▲ Strategic Reality Check

"We are not in a market where you can wait and react. You are either building your site pipeline now, or you are conceding the market to providers who are."

Working With Cloud2Ground

What an Engagement Looks Like

We work with a select number of inference providers at any given time — deliberately. Our methodology depends on deep market access and relationship-based sourcing, both of which degrade with volume. We are not a high-throughput advisory shop. We are a precision instrument.

Phase	Timeline	Activity & Output
Discovery & Brief	Week 1–2	Structured intake session covering workload profile, MW requirements, latency targets, expansion roadmap, and capital structure. Output: Precision Site Brief
Market Activation	Week 2–6	Active outreach across our subsurface network in your priority markets. Utility capacity verification calls. Off-market landowner contact. Output: Live candidate pipeline
Shortlist Delivery	Week 6–8	Structured Investment Shortlist with all 22-point diligence screen results. Narrative summary of top 3 recommended sites. Optional: client site tour coordination. Output: Investment Shortlist Package
Post-Shortlist Support	Optional	LOI structuring support, utility negotiation advisory, third-party consultant coordination, and closing-phase due diligence management. Output: Transaction-ready position

◆ Our Engagement Is Not Speculative

We work with a contractual arrangement. This structure exists for one reason: it aligns our incentives entirely with your outcome, not with transaction volume. When we recommend a site, it is because it is the **right site** — not because it closes a deal.

Call to Action

Ready to Build Your Inference Site Pipeline?

Cloud2Ground is accepting a limited number of new client engagements for **Q3 2026**. If you are an inference provider, co-location operator, or sovereign AI program evaluating North American deployment markets, we want to hear from you. Our intake process is designed to be fast: a 45-minute qualification call, followed by a written brief within 5 business days.

Option 1	Option 2	Option 3
<p>01</p> <p>Schedule a Discovery Call</p> <p>Book a 45-minute intake session with our senior advisory team. We'll assess your requirements and determine whether Cloud2Ground is the right fit.</p>	<p>02</p> <p>Request a Market Brief</p> <p>We publish quarterly off-market power availability briefs for active markets. Request the brief for your priority geography.</p>	<p>03</p> <p>Download Our Site Criteria Worksheet</p> <p>Use our structured worksheet to sharpen your site requirements before any advisory engagement.</p>

© 2026 Cloud2Ground Advisory LLC. All rights reserved. This document is prepared for qualified recipients only and may not be redistributed without written consent. Information contained herein is based on Cloud2Ground proprietary research and third-party sources believed to be reliable but not guaranteed. | cloud2ground.tech

Market perspectives reflect Cloud2Ground proprietary deal intelligence, partner ecosystem inputs, and third-party research. Site-specific conclusions require confirmatory diligence.

DISCLAIMER & IMPORTANT DISCLOSURES

This white paper is prepared by Cloud2Ground Advisory LLC solely for informational purposes and is intended exclusively for qualified infrastructure investors, hyperscale computing operators, and institutional recipients who have received it directly from Cloud2Ground. It does not constitute and should not be construed as investment advice, legal counsel, engineering guidance, or a solicitation to engage Cloud2Ground Advisory or any affiliated entity.

All market intelligence, site data, power economics estimates, timeline projections, and financial figures contained herein are based on Cloud2Ground's proprietary research, publicly available utility interconnection filings, and third-party sources believed to be reliable as of the date of publication.

Cloud2Ground makes no representation or warranty, express or implied, as to the accuracy, completeness, or fitness for any particular purpose of the information contained in this document. Market conditions, utility queue positions, land availability, and regulatory environments are subject to rapid change.

Past performance of Cloud2Ground engagements, including site identification timelines, shortlist conversion rates, and MW delivery statistics, is not indicative of future results. Individual engagement outcomes will vary based on client requirements, target markets, capital structure, and prevailing market conditions at the time of engagement.

This document is **CONFIDENTIAL** and may not be reproduced, distributed, or transmitted in any form or by any means — electronic, mechanical, photocopying, recording, or otherwise — without the prior written consent of Cloud2Ground Advisory LLC. Unauthorized distribution may violate applicable securities laws and confidentiality obligations.

© 2026 Cloud2Ground Advisory LLC. All rights reserved. Cloud2Ground Advisory LLC is not a registered investment advisor, broker-dealer, or licensed real estate professional in any jurisdiction. Prospective clients should conduct their own independent due diligence and consult qualified legal, financial, and engineering advisors before making any infrastructure investment or site acquisition decision.

[/cdn-cgi/l/email-protection](#)