

Meta Locks 6.6 GW of Nuclear Power Through Vistra, TerraPower, and Oklo Deals—Largest Corporate Nuclear Commitment for AI Infrastructure

Meta just committed to more nuclear capacity than Switzerland's entire grid produces. The AI energy arms race has moved past renewables—it's now about who can secure baseload gigawatts first.

The Deals: Dissecting 6.6 Gigawatts of Corporate Nuclear Ambition

On January 9, 2026, Meta announced three nuclear energy partnerships that collectively represent the largest corporate nuclear commitment in history. The company secured agreements with Vistra, TerraPower, and Oklo totaling up to 6.6 GW of nuclear capacity by

2035—enough to power approximately 5 million homes.

The numbers break down across three distinct technical approaches:

Vistra: 2.6 GW from existing and uprated plants. Meta signed 20-year power purchase agreements covering the Perry and Davis-Besse nuclear plants in Ohio (2,176 MW existing capacity) plus 433 MW in plant uprates. Deliveries start late 2026, making this the fastest path to electrons. The uprates alone add over 15% new capacity to the PJM grid.

TerraPower: 2.8 GW baseload plus 1.2 GW storage from advanced reactors. Meta is funding two Natrium reactors delivering 690 MW of firm power, with options for six additional units totaling 2.1 GW by 2035. First reactors come online as early as 2032. The Natrium design uses molten salt thermal storage, enabling the system to peak at 4 GW total output during demand spikes—a critical capability for bursty AI training workloads.

Oklo: 1.2 GW from a scalable fast reactor campus. This partnership targets a first-phase deployment by 2030, reaching full capacity by 2034. Oklo's compact fast reactors represent the most experimental portion of the portfolio but offer the densest power delivery per unit of land.

These deals originated from Meta's December 2024 request for proposals seeking 1-4 GW of new nuclear capacity, primarily targeting the PJM grid region. Meta exceeded its own stated requirements by 65%, signaling either unexpected demand projections or a strategic decision to lock up capacity before competitors could act.

Why This Changes the AI Infrastructure Calculus

The shift from renewable energy certificates to firm baseload nuclear represents a fundamental acknowledgment: AI training at scale requires power profiles that intermittent sources cannot reliably deliver.

The Load Profile Problem

Modern AI training clusters operate at sustained utilization rates exceeding 90% for weeks or months at a time. A single H100 cluster consuming 35 MW cannot tolerate supply intermittency without either curtailing training runs or maintaining expensive backup

systems. When your training job costs $50 million and takes six weeks, you cannot afford power variability.

Renewables with storage work for many enterprise workloads. They do not work for continuous 24/7 baseload at gigawatt scale—not yet, not economically. The storage requirements alone would rival the cost of the compute infrastructure being powered.

Nuclear provides what grid operators call "firm capacity"—power that shows up exactly when scheduled, every hour, regardless of weather or time of day. For AI infrastructure operators, this translates directly into compute availability and training job completion guarantees.

The Capacity Arms Race

This deal signals that hyperscalers now view power procurement as a core competitive advantage rather than an operational expense. Microsoft's arrangement to restart Three Mile Island Unit 1 for Copilot infrastructure, Amazon's nuclear investments, and now Meta's commitment indicate a clear pattern: the companies building the largest AI systems are systematically locking up the power needed to run them.

The competitive dynamics here matter enormously. PJM Interconnection—the regional transmission organization covering 13 states from Illinois to New Jersey—has faced severe capacity constraints and interconnection queue backlogs. By securing 2.6 GW from existing Vistra plants with late-2026 delivery, Meta bypassed the multi-year wait for new generation projects. Competitors seeking similar capacity in the same region now face a thinner market.

This is infrastructure as moat. Every gigawatt Meta secures is a gigawatt unavailable to competitors at any price.

Winners and Losers

Winners:

- Existing nuclear plant operators. Vistra's stock implications are obvious—20-year PPAs at what are likely premium rates transform plant economics.

- Advanced reactor developers. TerraPower and Oklo now have the offtake agreements and funding certainty to accelerate construction timelines.
- PJM grid reliability. An additional 15% capacity from Vistra uprates alone helps address the region's growing power crunch.
- Domestic energy security advocates. These deals keep critical AI infrastructure tethered to U.S.-based power generation.

Losers:

- Renewable-only corporate sustainability strategies. The narrative that hyperscaler growth can be powered purely by wind and solar PPAs has been definitively complicated.
- Late movers in AI infrastructure. Companies that delayed power procurement strategy now face higher costs and reduced availability.
- Natural gas peaker plants. As nuclear provides baseload and storage handles peaks, the middle ground that gas occupies becomes increasingly squeezed.

Technical Architecture: Why These Three Approaches Matter

Meta's portfolio diversification across three vendors with distinct reactor technologies is not just risk hedging—it reflects a sophisticated understanding of how different nuclear technologies map to different infrastructure needs.

Vistra: The Pragmatic Play

The Perry and Davis-Besse plants are boiling water and pressurized water reactors respectively—proven designs with decades of operational history. The 433 MW in uprates comes from equipment upgrades that extract more output from existing licensed thermal capacity. This is the lowest-risk, fastest-to-deploy portion of the portfolio.

For infrastructure architects, the key insight is lead time. Late 2026 delivery means Meta is adding gigawatt-scale firm power within 11 months of announcement. No new reactor construction achieves this timeline. When you need power fast, existing plants with upgrade potential are the only answer.

The 20-year PPA structure also reveals Meta's planning horizon. These agreements lock in power costs through 2046, spanning multiple generations of AI hardware. Meta is betting that its power needs will remain concentrated in the PJM region for the next two decades—a significant site commitment.

TerraPower Natrium: The Flexibility Architecture

TerraPower's Natrium design pairs a 345 MW sodium-cooled fast reactor with a molten salt energy storage system. The reactor runs continuously at steady state, while the storage system absorbs excess heat during low-demand periods and releases it during peaks—potentially doubling effective output to 690 MW per unit.

For AI workloads, this architecture solves a real problem. Training clusters run flat-out, but inference demand varies dramatically by time of day and day of week. A Natrium installation can provide steady 345 MW during training jobs while flexing up to 690 MW for inference peak hours, all from the same fuel cycle.

The sodium coolant operates at atmospheric pressure (unlike pressurized water reactors) and at higher temperatures (500°C vs 300°C), improving thermal efficiency and reducing containment requirements. The tradeoff: sodium reacts violently with water and air, requiring specialized handling and maintenance procedures.

Meta's option for six additional Natrium units beyond the initial two suggests confidence in the design's commercial viability. If the first units perform as specified, the subsequent six could be construction-line optimized with significantly reduced per-unit costs.

Oklo: The Density Play

Oklo's Aurora fast reactor design targets a different niche entirely. These are small, compact units optimized for high power density per unit of land area. The campus approach—multiple small reactors on a single site—offers incremental capacity additions without the all-or-nothing investment profile of large reactor construction.

For infrastructure planners, Oklo's value proposition is optionality. A 1.2 GW campus built from smaller units can be scaled more precisely to actual demand growth. If Meta's power needs in a given region plateau at 800 MW, they avoid paying for 400 MW of unused

capacity.

The fast reactor design also closes the nuclear fuel cycle more efficiently, potentially reducing spent fuel volumes and extending uranium resource utilization. This matters less for near-term economics but significantly for long-term sustainability claims.

First-phase delivery by 2030 is aggressive but not implausible given Oklo's head start on NRC licensing. The real question is construction execution—Oklo has not yet built a commercial-scale reactor, and first-of-a-kind projects historically face cost and schedule overruns.

The Contrarian Take: What Most Coverage Gets Wrong

This Is Not Primarily About Decarbonization

Press coverage has emphasized Meta's sustainability commitments and carbon-free energy targets. This framing misses the point. If decarbonization were the primary driver, Meta would be signing more solar and wind PPAs with storage—technologies that are commercially mature and rapidly declining in cost.

The nuclear commitment is about power density, reliability, and competitive positioning. Carbon-free electricity is a welcome side effect, not the primary motivation. Meta needs firm baseload power for AI infrastructure at a scale and reliability level that renewables cannot currently provide at competitive economics. The decision matrix leads to nuclear regardless of carbon considerations.

Hyperscalers will publicly emphasize sustainability messaging because it plays well with regulators, communities, and ESG-focused investors. But the boardroom calculus is simpler: nuclear is the only way to get multiple gigawatts of 24/7 power in the United States without burning fossil fuels and within the required timeline.

The Cost Numbers Are Missing—On Purpose

Notably absent from all announcements: specific pricing terms. We know the Vistra deals are 20-year PPAs, but neither the per-MWh rate nor total contract value has been

disclosed. The TerraPower and Oklo arrangements involve “funding” and “investment” without quantification.

This opacity is strategic. If Meta locked in nuclear power at premium rates, disclosure would invite criticism about overpaying. If they secured below-market rates, it would reveal negotiating leverage that competitors could attempt to exploit with other suppliers.

Reasonable estimates suggest Meta is paying $60-90/MWh for the Vistra existing plant output (competitive with current PJM wholesale rates plus risk premium) and potentially $100-150/MWh for the advanced reactor capacity (reflecting first-of-a-kind risk). Total contract value across all three deals likely exceeds $20 billion over the full term—placing this among the largest single energy procurement decisions in corporate history.

Execution Risk Is Substantial

The 6.6 GW headline number assumes full delivery across all three vendors by 2035. Historical context suggests skepticism.

Vistra’s existing plant output (2.6 GW) is essentially risk-free—these plants are already operating. The 433 MW in uprates carry modest execution risk, as uprate projects have a strong track record.

TerraPower’s Natrium reactors face meaningful first-of-a-kind risk. While the DOE-backed demonstration reactor in Wyoming provides design validation, commercial construction at scale introduces new variables. A 2032 online date is achievable but not guaranteed.

Oklo’s 1.2 GW by 2034 is the highest-risk commitment. Oklo has yet to operate a commercial reactor, and its previous NRC license application was denied, requiring resubmission. The company is building capability rapidly, but projecting 1.2 GW of deployment within eight years from a standing start requires aggressive assumptions about licensing, construction, and commissioning timelines.

A realistic baseline scenario delivers 3.5-4.0 GW by 2035, with the remaining capacity arriving in 2036-2038. This still represents a transformational amount of nuclear capacity but illustrates why portfolio diversification across three vendors was structurally necessary.

What This Means for Technical Leaders

Power Procurement Is Now a Core Infrastructure Competency

For CTOs and infrastructure architects at companies building AI capabilities, Meta's announcement carries a clear message: power strategy cannot be delegated to facilities management or outsourced to cloud providers without significant strategic risk.

If your AI roadmap requires hundreds of megawatts within five years, you need to be making power procurement decisions today. Interconnection queues in major grid regions now stretch 4-6 years. PPAs for existing nuclear or large-scale renewables require 18-24 months of negotiation and regulatory approval. Waiting until you need the power means waiting years past when you needed it.

Practical steps for technical leaders:

- Audit your projected AI compute growth and translate it to power requirements. Most organizations underestimate by 2-3x.
- Map your data center footprint to grid regions and assess local capacity constraints.
- Engage power procurement specialists now, even if deployment is years away. The market is rapidly tightening.
- Consider co-location with hyperscalers as an alternative to independent power procurement—but recognize you inherit their constraints.

Nuclear Proximity Becomes a Site Selection Factor

Data center site selection has historically prioritized connectivity, cooling, land cost, and tax incentives. Nuclear power availability is now a first-tier consideration.

The PJM region—where Meta concentrated its deals—offers the highest density of commercial nuclear plants in the United States. Companies seeking firm baseload power should evaluate sites in PJM, MISO (Midwest), and Southeast regions where existing nuclear capacity is concentrated.

For greenfield developments, proximity to planned advanced reactor deployments may offer long-term advantages. TerraPower's Wyoming site, Oklo's Ohio campus, and other

projects in development represent future anchor power sources. Site selection decisions made today will determine access to these resources in 2030 and beyond.

Technology Architecture Implications

Power profile constraints should influence AI infrastructure architecture decisions. Several patterns emerge:

Training vs. Inference Clustering: Training jobs require sustained power draw for extended periods. Inference demand is bursty and variable. Separating these workloads to facilities with different power profiles—baseload nuclear for training, flexible generation for inference—may optimize costs.

Workload Scheduling for Power Optimization: If your facility includes both firm and intermittent power sources, workload schedulers should incorporate power availability forecasts. Run large training jobs when baseload capacity is fully available; defer to inference and smaller jobs during periods of renewable variability.

Thermal Integration: Advanced reactors like Natrium operate at higher temperatures, producing waste heat that could be captured for district heating or industrial processes. Data centers co-located with these facilities might reduce cooling costs by rejecting heat to a higher-temperature sink.

The Next 12 Months: What to Watch

Competitive Response

Microsoft, Amazon, and Google will respond to Meta's announcement within months. None can afford to cede the power procurement advantage Meta has established. Expect announcements of comparable scale by mid-2026, likely targeting different geographic regions to avoid direct competition for the same capacity.

Watch for nuclear deals in the Southeast (dominated by Southern Company and Duke Energy nuclear fleets) and Texas (where Comanche Peak and South Texas Project offer available capacity). International nuclear procurement—particularly in France, where EDF

operates the world's largest commercial nuclear fleet—may also accelerate.

Regulatory Acceleration

The NRC licensing backlog is now the primary constraint on advanced reactor deployment. Meta, Microsoft, and other hyperscalers have significant lobbying presence and will push for accelerated licensing pathways.

The Oklo license resubmission, expected in 2026, will be a bellwether. A streamlined approval would signal NRC adaptation to commercial pressure. Prolonged review would raise questions about whether the regulatory framework can support the deployment pace hyperscalers require.

Congress is also likely to revisit nuclear permitting reform. The bipartisan ADVANCE Act passed in 2024 provided some acceleration, but hyperscaler demand may justify more aggressive intervention.

Advanced Reactor Execution

TerraPower's Wyoming demonstration reactor, currently under construction, will reach key milestones in 2026-2027. Construction progress, cost tracking, and regulatory interactions will provide leading indicators for the commercial reactors Meta has funded.

Oklo's path to first commercial reactor deployment will clarify by year-end. Watch for NRC license application status, site preparation activities, and manufacturing partnerships. Delays in any of these areas push the 2030 first-phase target at risk.

Grid Integration Challenges

Adding 2.6+ GW to the PJM grid requires transmission upgrades and interconnection agreements. The speed at which Vistra can complete uprates and Meta can take delivery depends heavily on grid operator coordination.

PJM's interconnection queue reform, implemented in 2023, improved throughput but remains a bottleneck. If Meta's Vistra capacity faces interconnection delays, it signals systemic constraints that affect all large-scale power additions.

Pricing Transparency

As more hyperscaler nuclear deals close, pricing norms will emerge. If Meta secured below-market rates due to first-mover advantage, later deals will face higher costs. If nuclear developers are currently subsidizing deployments to build track record, prices may decline as execution risk decreases.

Watch for any disclosure in SEC filings, state regulatory proceedings, or industry conference presentations that reveals pricing benchmarks. These numbers will inform corporate power procurement strategy across the industry.

Beyond 2035: The Structural Shift

Meta's announcement marks a phase transition in how technology companies relate to energy infrastructure. For three decades, hyperscalers treated power as a utility input—something purchased from others at market rates. That era is ending.

The companies building the largest AI systems are now directly financing generation assets, signing multi-decade purchase commitments, and influencing reactor development roadmaps. They are becoming, in effect, the anchor tenants of a new nuclear power industry.

This vertical integration has precedent. Aluminum smelters built their own hydroelectric dams. Steel mills integrated with coal mines. When an industry's power requirements exceed what the grid can readily supply, vertical integration becomes necessary.

For the technology industry, this means power infrastructure joins semiconductors, networking, and custom silicon as domains where hyperscaler scale justifies proprietary solutions. The implications cascade across the industry:

- Cloud customers will increasingly face power-linked capacity constraints as hyperscalers prioritize internal AI workloads.
- Startup AI companies without power procurement capability face a new form of infrastructure dependency.
- Enterprise AI deployments may need to consider power availability as a geographic

constraint.

The AI infrastructure buildout of the next decade will be shaped as much by power availability as by chip production. Meta has made an aggressive bet that nuclear energy—in both existing and advanced forms—represents the winning path through that constraint.

The strategic lesson is unmistakable: in the AI era, power procurement is no longer someone else's problem—it is core infrastructure strategy, and the companies that secure capacity first will build the AI systems that matter.



WRITTEN BY

Artur Markus

Software engineer in Basel building AI infrastructure and agentic systems, with a background in pharmaceutical automation, GxP compliance, and validation. I help teams turn AI from a chatbot into a reliable operational layer.

Book a meeting