INTELLIGENCE BRIEF

THE DATA CENTER MARKET: PROGNOSIS, RISKS, AND OBSOLESCENCE

What Communities, Investors, and Policymakers Need to Know — February 2026

Prepared by Phi Growth Strategies | Sources verified as of February 16, 2026

EXECUTIVE SUMMARY: THE 60-SECOND VERSION

The data center industry is in a **$3 trillion investment supercycle** that will double global capacity by 2030. But the technology inside these buildings is changing so fast that **facilities built even two years ago are already at risk of functional obsolescence**. The industry is splitting in two: hyperscalers (Amazon, Google, Microsoft, Meta) with fortress balance sheets are building for the long term. Everyone else—speculative developers, neoclouds, and debt-financed operators—is increasingly exposed to stranded-asset risk, community backlash, and a technology shift that could make their buildings useless within five years.

Meanwhile, the architecture of AI computing itself is migrating away from massive centralized data centers toward **distributed, federated, edge-based models** that process data closer to where it's generated. This doesn't eliminate data centers—but it does mean that the biggest, most power-hungry spec builds targeting a single use case (AI training) may be the first to become white elephants.

Sources for this summary: JLL 2026 Global Data Center Outlook (Jan 2026); Moody's "Data Centers: Managing Risk Amid a Market Boom" (Jan 2026); Goldman Sachs Research (Feb 2025); Bloomberg/JPMorgan (Feb 2026); DC Byte 2026 Outlook (Jan 2026); Dell'Oro Group 2026 Predictions (Dec 2025).

GLOSSARY: WHAT THE WORDS MEAN

For a lay audience—county commissioners, community members, local press.

TERM	WHAT IT MEANS	WHO USES IT
Hyperscaler	The giants—Amazon (AWS), Microsoft (Azure), Google Cloud, Meta. They build data centers for their own use and rent capacity to others. They have cash to burn.	*Industry-wide*
Neocloud	Newer, smaller companies (CoreWeave, Lambda, Nebius) that rent GPU capacity to AI companies. They do NOT own their own clouds—they borrow billions to buy chips and lease space. High risk.	*S&P, SDxCentral*
Spec / Speculative Build	A data center built WITHOUT a signed tenant. The developer is betting someone will want it. Like building a shopping mall before any stores sign leases.	*JLL, Moody's*
Colocation (Colo)	A data center where multiple tenants share space. Think of it as an industrial office park for servers. The landlord provides power and cooling; tenants bring their own equipment.	*Industry-wide*
Edge Computing	Processing data close to where it's created (a factory floor, a hospital, a retail store) instead of sending it to a massive centralized facility hundreds of miles away.	*IDC, Cisco, Dell*
Federated Learning	Training AI models across many devices/locations without moving the raw data. Each device learns locally and only shares the lessons, not the data. Preserves privacy, reduces bandwidth.	*MIT, Google, PMC*
Inference	When a trained AI model answers questions or makes predictions. This is the "using" phase (vs. training). Inference is shifting to smaller, distributed facilities closer to users.	*Dell'Oro, IDC, JLL*
Training	Teaching an AI model by feeding it massive data. Requires enormous computing power, concentrated in huge data centers. This is what drives the current mega-build.	*Industry-wide*
GPU	Graphics Processing Unit—originally for video games, now the workhorse chip for AI. Nvidia dominates. These chips cost $30K–40K each and consume enormous power.	*Nvidia, industry*
Rack Density	How much power each server cabinet uses. Traditional: 8–15 kW. AI racks: 50–100+ kW. This 5–10x increase is what makes older data centers obsolete—they literally can't cool the new chips.	*Global DC Hub, Deloitte*
Liquid Cooling	Running liquid (water or coolant) directly over chips to remove heat. Required for AI-density workloads.	*Deloitte, Dell'Oro*

	Most existing data centers use air cooling and CANNOT be easily retrofitted.	
Stranded Asset	An asset that loses value prematurely because demand shifts or technology changes. A data center built for 2023 workloads that can't handle 2026 AI chips is a stranded asset.	*Moody's, Cbus, Invesco*
PJM Interconnection	The largest U.S. grid operator, serving 65M+ people across 13 states. Data centers are overwhelming PJM's capacity—it will be 6 GW short of reliability targets by 2027.	*CNBC, PJM, Monitoring Analytics*
ABS / CMBS	Asset-Backed Securities and Commercial Mortgage-Backed Securities. Wall Street packages data center debt into bonds and sells slices to investors. JPMorgan projects \$30–40B/year in 2026–2027.	*JPMorgan, Bloomberg*
DDTL	Delayed Draw Term Loan. A credit facility where borrowers draw funds over time. CoreWeave's DDTLs carry 9–15% interest rates, collateralized by GPUs that depreciate 50–70% in value.	*CoreWeave SEC filings, Fortune*
Megawatt (MW)	Standard unit for data center capacity. 1 MW powers roughly 750–1,000 server racks. A large data center campus might be 100–500 MW. Austin's entire peak load is about 5 GW.	*JLL, industry*

THE MARKET: CURRENT STATE AND PROGNOSIS

The Numbers Are Staggering—and Real

The data center industry is undergoing what JLL calls an "investment supercycle." The raw numbers, corroborated across multiple tier-one sources: nearly **100 GW of new capacity** will come online between 2026 and 2030, doubling what exists today. Global investment will reach approximately **$3 trillion** by 2030 (JLL, Moody's). Average construction cost has risen from $7.7M to $11.3M per MW (JLL, 2026). AI workloads could represent half of all data center capacity by 2030, up from roughly a quarter in 2025 (JLL). In 2025 alone, AI-related companies tapped debt markets for at least **$200 billion** (Bloomberg/JPMorgan, Feb 2026). 113 data center M&A transactions completed in 2025, totaling over $69 billion (S&P, Feb 2026).

Sources: JLL 2026 Global Data Center Outlook (Jan 7, 2026); Moody's Investors Service (Jan 16, 2026); Bloomberg/JPMorgan "The $3 Trillion AI Data Center Build-Out" (Feb 3, 2026); S&P Global Market Intelligence (Feb 13, 2026); Deloitte 2026 Hardware & Consumer Tech Outlook (Feb 10, 2026).

But the Gap Between Plans and Reality Is Widening

DC Byte's global tracking data reveals a structural problem: **committed supply has increased more than 6x since 2019, while live (operational) capacity has grown far more slowly**. In several key hub markets, committed capacity is more than twice the volume actually under construction. This means a growing share of announced projects will never be built—or will be built too late to matter. Capital is being deployed 24–36 months ahead of expected delivery, often at the land acquisition or permitting stage, increasing exposure to execution risk.

Goldman Sachs Research projects occupancy will peak above 95% in late 2026, **then moderate starting in 2027 as more data centers come online and AI-driven demand growth slows**. This is the window in which oversupply risk becomes material.

Sources: DC Byte 2026 Data Centre Outlook (Jan 8, 2026); Goldman Sachs "AI to drive 165% increase in data center power demand" (Feb 2025).

THE OBSOLESCENCE THREAT: ARE THEY BUILDING YESTERDAY'S TECHNOLOGY?

The Short Answer: Yes, Many of Them Are

The technology shift is brutal and accelerating. Moody's (Jan 2026) explicitly warns that **data center landlords are "grappling with technological advancements so swift that they threaten obsolescence of existing projects and properties."** The core problem is rack density. Traditional data centers were designed for 8–15 kW per rack. AI workloads require 50–100+ kW per rack. Designs optimized for 20 kW racks in 2023 may be incompatible with the 80 kW workloads demanded by 2026 (Global Data Center Hub, June 2025). This is not a minor upgrade—it requires fundamentally different cooling systems, power distribution, floor slab strength, and physical architecture.

Liquid cooling is now the dividing line. Air-cooled facilities—which constitute the vast majority of the existing installed base—cannot handle AI-density workloads. Direct-to-chip liquid cooling, rear-door heat exchangers, and immersion cooling are no longer "advanced options." They are baseline requirements for next-generation tenants. Most facilities built even five years ago lack the physical provisions (chilled water loops, slab strength, piping access) to retrofit without disruptive and costly overhauls.

Nvidia's GPU chip cycle is 18 months. Cbus Super (Australia's major pension fund) told the 2025 Fiduciary Investors Symposium: **"We're talking about and trying to price stranded asset risk... there's the potential that in five years' time this wonderful data centre you bought has been completely superseded by something else and nobody wants to buy it."**

Asset managers are now bifurcating valuations: AI-optimized facilities may trade at 25x+ EBITDA multiples, while **everything else faces compressed valuations, divestiture, or exclusion from strategic portfolios** (Global Data Center Hub). What the industry calls "obsolescence debt"—the deferred capital expenditure required to bring legacy infrastructure to AI-ready standards—builds quietly as operators avoid upgrades to protect margins. When the retrofit becomes urgent, operators face reactive spending under time pressure, lower ROI, and impaired asset value at exit.

Sources: Moody's Investors Service (Jan 16, 2026); Global Data Center Hub "How Technology Risk Is Reshaping Global Data Center Investing" (June 2025); Investment Magazine / Cbus Super (Nov 2025); Invesco (Nov 2025); Deloitte 2026 Hardware Outlook (Feb 2026); Dell'Oro Group 2026 Predictions (Dec 2025).

THE TECHNOLOGY SHIFT: FEDERATED, DISTRIBUTED, AND EDGE

AI Is Moving Away from Mega Data Centers

The most consequential shift for anyone evaluating data center investments: **AI inference—the "using" phase that represents actual business value—is migrating to the edge**. JLL projects that by 2027, inference workloads will overtake training as the dominant AI requirement. IDC's Dave McCarthy states: "As the focus of AI shifts from training to inference, edge computing will be required to address the need for reduced latency and enhanced privacy." Cisco's Jeremy Foster says AI "breaks the old cloud model" by forcing distributed architecture.

What this means in plain language: the giant centralized data centers being proposed in rural communities are optimized for **training** AI models—the initial teaching phase. But the business value comes from **inference**—actually using the trained model—and that requires speed, proximity, and distribution, not concentration. A self-driving car can't wait 200 milliseconds for an answer from a data center 500 miles away. A hospital's diagnostic AI needs to run on-site. A factory's quality control system needs to process in real time.

Federated learning enables AI models to train across many devices without moving raw data—each device learns locally and shares only the model improvements, not the data itself. This reduces the need for massive centralized data gathering and processing. Combined with smaller, more efficient AI models (small language models, or SLMs), the trend is toward **distributed intelligence**: smaller compute nodes placed closer to users, connected by high-speed networks, rather than everything funneled to one mega-facility.

IDC projects edge computing spend will reach **nearly $378 billion by 2028**. The edge AI market alone is projected to grow at 21.7% CAGR through 2030 (VerticalData). Architecture is moving toward what CrispIdea calls "federated multi-cloud"—data, AI services, and applications interacting seamlessly across providers and locations, not locked inside a single massive facility.

Sources: JLL 2026 Global Data Center Outlook (Jan 2026); IDC / Dave McCarthy, Research VP Cloud & Edge Services (Dec 2025); Cisco / Jeremy Foster, SVP Compute (Jan 2026); SDxCentral "AI Inferencing Will Define 2026" (Jan 2, 2026); Dell Technologies / Pierluca Chiodelli, VP Edge Portfolio (Feb 2025); VerticalData "Edge AI Mega-Clusters" (Jan 2026); CrispIdea "Cloud Computing in 2026" (Dec 2025); Akamai / Tom Bradshaw (Jan 2026).

What This Means for Spec Data Centers in Rural Communities

A speculative data center proposed today—one without a signed tenant, designed around current-generation hardware, using air cooling, targeting AI training workloads—faces a **triple obsolescence risk**: (1) The cooling and power architecture may be inadequate for next-generation chips arriving in 18–24 months. (2) The primary workload it's designed for (training) is the segment most likely to consolidate into fewer, larger hyperscaler-owned facilities. (3) The growth workload (inference) is moving in the opposite direction—toward smaller, distributed, edge-based facilities closer to users, not in rural locations far from population centers.

THE DEBT PROBLEM: WHO'S HOLDING THE BAG?

The Numbers Are Eye-Watering

Bloomberg and JPMorgan report that AI-related companies tapped debt markets for **at least $200 billion in 2025**—and that's likely an undercount given private deals. Morgan Stanley projects $250–300 billion in 2026 issuance from hyperscalers alone. JPMorgan projects data center securitization (ABS + CMBS) will reach **$30–40 billion annually in 2026–2027**, representing 7–10% of total combined issuance. The total debt financing needed to support the buildout: roughly $870 billion by 2030 (JLL).

The CoreWeave Case: A Canary in the Coal Mine

CoreWeave—the poster child for neocloud AI infrastructure—illustrates the fragility. The company IPO'd in 2025 at $40, peaked at $180, then **collapsed 62%** to ~$68 by December. Its financial profile: over $14 billion in total debt with interest rates of 9–15%. $3.6 billion in debt payable by June 30, 2026. Interest expenses of $311 million in a single quarter—triple the prior year. Operating margins of 1.6%. Net losses persisting despite revenue growth to $1.36 billion/quarter. Its GPUs—used as loan collateral—have seen market rental rates fall 50–70%, shrinking collateral value just as repayment schedules begin.

62% of CoreWeave's revenue comes from **a single customer (Microsoft)**. Its entire model depends on perfect execution: borrow billions, buy GPUs, get data centers online fast, generate revenue immediately to service debt. Texas rainstorms delayed a single 260 MW facility by 60 days—and wiped out $33 billion in market value. Fortune's AI editor noted that even at projected 2026 revenues of $10.9–14.9 billion, a 1.6% operating margin generates only ~$192 million in income—"not enough to cover interest expenses or make principal repayments."

Sources: Bloomberg/JPMorgan (Feb 3, 2026); Fortune (Nov 8, 2025); CoreWeave SEC filings; Baptista Research (Dec 2025); The Investor Channel / Substack (Nov 2025); ainvest.com (Dec 2025); Level-Headed Investing (Oct 2025); Construction Owners (2025).

The 1990s Telecom Parallel Is Not Hypothetical

Multiple credible sources now explicitly draw the comparison. In the late 1990s, telecom companies invested over $500 billion in fiber optic infrastructure, financed largely with debt, based on projections that internet traffic would grow 1,000% annually. The traffic eventually materialized—but not before the debt markets closed, dozens of companies went bankrupt, and communities were left with dark fiber they didn't need. As Technical.ly noted (Nov 2025): "The 1990s telecom infrastructure left investors holding fiber that no one yet needed. But it was a **much less intrusive stranded asset**. The question now is whether today's rural data centers will follow that pattern." Fiber is invisible underground. A 500,000 sq ft data center campus is not.

Sources: Technical.ly / Christopher Wink (Nov 2025); Outlook India (Feb 16, 2026); Hemispheres Investment Management (Oct 2025); Development Corporate (Nov 2025); Construction Dive (Jan 2026).

WHAT COMMUNITIES ARE EXPERIENCING RIGHT NOW

The Opposition Is Massive and Accelerating

Data Center Watch reports that **$98 billion in data center projects were blocked or delayed in Q2 2025 alone**—more than all previous quarters since 2023 combined. 188 activist groups are now organized across the United States. 66% of protested projects were blocked or delayed in Q2 2025. Opposition rose 125% quarter over quarter. In 2025, St. Charles, Missouri passed the nation's first year-long ban on data center construction. Several Georgia cities, including parts of Atlanta, approved restrictions or moratoriums. Virginia's Prince William County has $24.7 billion in data center projects entangled in lawsuits.

As of February 2026, WALB News reports **more than 140 activist groups across 24 states** are pushing back against data center projects. A coalition of 230 environmental groups—including Food & Water Watch, Greenpeace, and Oil Change International—wrote to Congress calling for a halt to new construction. The backlash is bipartisan: Republican officials focus on tax incentives and grid strain; Democrats focus on environmental impacts.

The Specific Harms

Water: Large data centers consume up to 5 million gallons of water per day—equivalent to the daily needs of 10,000–50,000 people (Washington Post). AI could consume 312–765 billion liters of water in 2025 (The Verge / Dec 2025 study). In Saline Township, Michigan, Oracle/OpenAI's planned facility will use more water than any other single user in the Great Lakes region. In Valdosta, Georgia, residents on private wells fear aquifer depletion will destroy farming, wildlife, and property values.

Power and Electricity Costs: Data centers currently account for 4% of U.S. electricity demand; that figure is expected to triple within three years (Project Censored / MediaJustice). PJM Interconnection will be 6 GW short of reliability targets by 2027. Data center demand drove PJM's capacity auction cost from $2.2 billion to $14.7 billion (a 500%+ increase), then to $16.4 billion—with data centers accounting for 63% of the total increase. In South Carolina, data centers will account for 65–70% of all new energy usage statewide. Austin's city manager found that proposed data centers seek more power than the city's entire peak load. Microsoft received $333 million in sales tax exemptions for a single Washington state data center.

Jobs: The promise of jobs is minimal. Data centers are among the most automated commercial real estate categories. A facility costing hundreds of millions creates as few as 30–50 permanent jobs after construction. Residents in Valdosta, Georgia describe facilities as "brightly lit, disruptive to wildlife, and providing only a handful of long-term jobs."

Farmland and Rural Character: In Appalachia, developers target former coal country, promising economic diversification. But as the Daily Yonder reported (Jan 2026), the pattern is familiar: extract value, externalize costs, leave communities holding the consequences. Data center campuses are industrial facilities—concrete, steel, generator farms, security fencing, 24/7 lighting and noise—incompatible with agricultural and residential surroundings.

Sources: Data Center Watch quarterly reports (2024–2025); WALB News (Feb 10, 2026); Nevada Current (Feb 9, 2026); Washington Post; The Guardian / SourceMaterial (Apr 2025); CNBC (Jan 2026); Common Dreams (Jan 3, 2026); Project Censored / Media Justice (Jan 2026); Daily Yonder (Jan 7, 2026); Austin American-Statesman (Jan 2026); Monitoring Analytics / PJM reports; Mountain State Spotlight (Oct 2025).

THE BOTTOM LINE: ARE SPEC DATA CENTERS ABOUT TO BE STRANDED?

The honest answer, based on the weight of current evidence from the most credible sources in this space:

Hyperscaler-owned, tenant-committed, AI-optimized facilities with liquid cooling will be fine. These are being built by companies with cash reserves measured in hundreds of billions, for workloads they control. They represent the bulk of projected investment value.

Speculative builds by non-hyperscaler developers are in serious trouble. They face a converging set of risks that no amount of optimistic demand projections can paper over:

- **Technology obsolescence** — air-cooled, low-density facilities are already functionally outdated for AI workloads (Moody's, Global DC Hub, Deloitte)
- **Workload migration** — inference is moving to distributed edge, reducing demand for massive centralized capacity (IDC, Cisco, JLL, Dell'Oro)
- **Debt exposure** — GPU collateral is depreciating 50–70% while interest rates run 9–15%; refinancing walls are imminent (CoreWeave filings, Fortune, Bloomberg)
- **Oversupply timing** — Goldman Sachs projects demand moderation starting 2027, just as the current wave of construction delivers (Goldman Sachs Research)
- **Community opposition** — $98 billion in projects blocked or delayed in a single quarter; 188 activist groups nationwide (Data Center Watch)
- **The fiber parallel** — the 1990s telecom bust left invisible underground cable; data centers leave 500,000+ sq ft industrial hulks in communities that didn't want them

Dell'Oro Group (Dec 2025) captures the tension precisely: even if the data center buildout doesn't slow materially in 2026, **"speculation about an AI-driven investment bubble is expected to intensify, as trillions of dollars in critical infrastructure are deployed amid lingering uncertainty about long-term monetization models."**

For communities evaluating proposed data centers: the question to ask is not whether data centers have value. They do. The question is whether **this specific developer, with this specific tenant commitment (or lack thereof), using this specific technology, at this specific location, will still be relevant in five years**—or whether the community will be left with a power-hungry industrial monument to someone else's speculation.

PREPARED BY

Phi Growth Strategies | Lauren Romero, MBA, Principal

Conservation-Commerce Integration | Central Texas Corridor