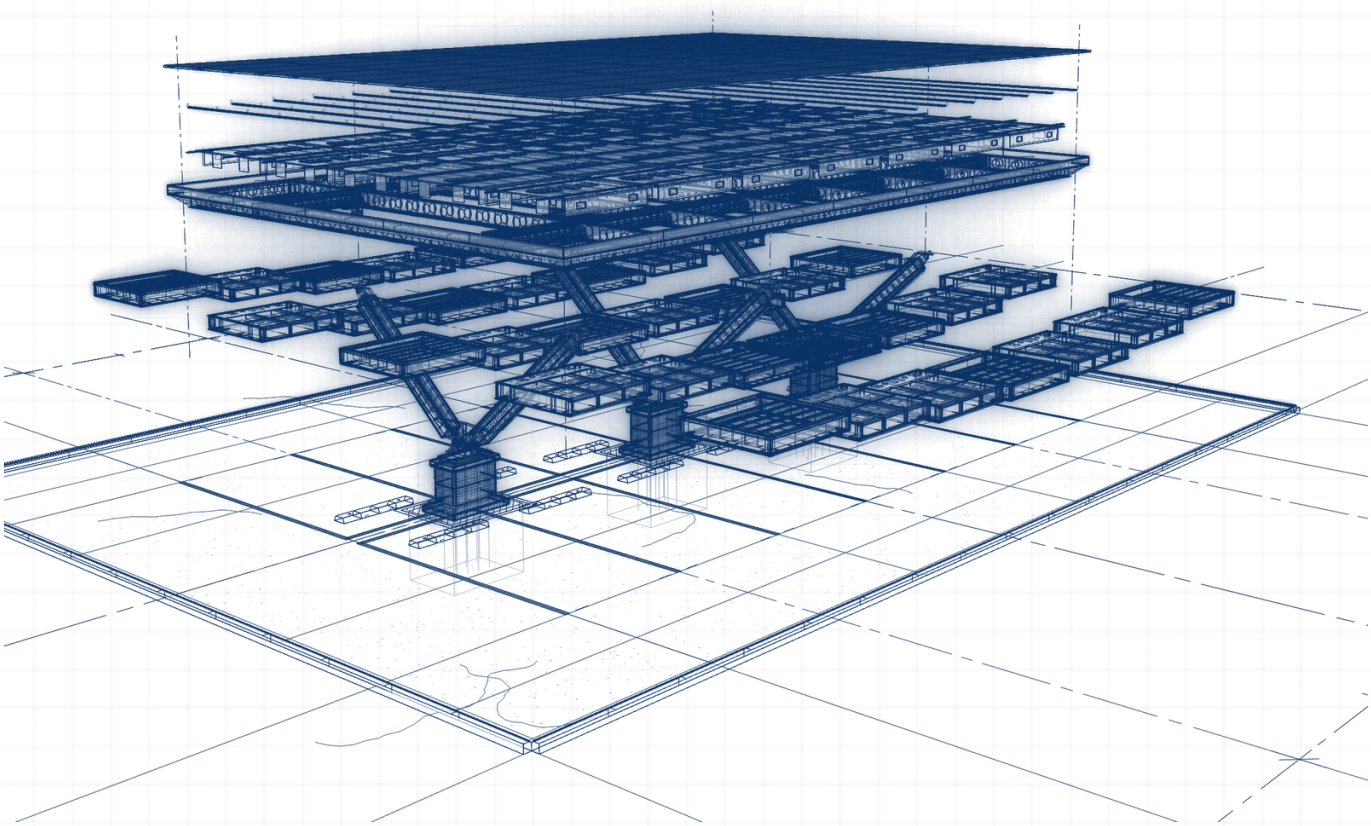
DISTRIBUTED AI DATA CENTERS

Distributed, Power-Sovereign AI Infrastructure for the Inference Economy

Why power, land, and latency now constrain traditional data centers, and how distributed, power-sovereign AI canopies answer all three.



ISSUED BY	DOCUMENT	REVISION	PUBLISHED	REVISED
LT350, LLC	LT350-WP-01	B	2026-03-30	2026-11-12

CONTENTS

01 The AI Infrastructure Crunch

02 LT350's Architecture: Modular AI Canopies in Existing Parking Lots

03 Structural Advantages of LT350's Model

04 Market Alignment: Why LT350 Is Positioned for the Next Decade

05 Conclusion

N Notes and Sources

Executive Summary

AI is driving unprecedented growth in global compute demand, and traditional data-center models are hitting hard constraints in power, land, and latency. Electricity consumption from data centers is projected to rise sharply as AI workloads scale.[1] Meanwhile, power delivery timelines are stretching, sites that meet power and fiber requirements are scarce, and more workloads require real-time inference close to where data is generated.[2]

LT350 introduces a fundamentally different architecture: distributed, power-sovereign, modular AI data centers deployed in existing parking lots, integrating GPU cartridges, memory cartridges, solar generation, battery storage, and local fiber backhaul. This model directly addresses the structural constraints facing hyperscalers, enterprises, hospitals, and autonomous-vehicle fleets.

AT A GLANCE

2,600 GW

of generation and storage waiting in U.S. interconnection queues at the end of 2023[6]

4+ years

median time for a queued project to reach operation[6]

7 to 12%

of U.S. electricity that data centers could use by 2028, up from 4.4% in 2023[5]

33% a year

growth in demand for AI-ready data-center capacity through 2030[4]

01
The AI Infrastructure Crunch

1.1 Explosive Growth in Data-Center Power Demand

Global data-center electricity consumption is expected to double or more over the next decade as AI training and inference workloads scale. [3] Demand for AI-ready data-center capacity is projected to grow even faster, about 33 percent a year through 2030, outpacing the market overall. [4]

In the U.S., data centers used 4.4 percent of total electricity in 2023 and could use roughly 7 to 12 percent by 2028. [5]

1.2 Grid Interconnection Delays and Regional Strain

By the end of 2023, nearly 2,600 GW of generation and storage sat in U.S. interconnection queues, and the median project took more than four years to reach operation. Only about one in five projects that applied ultimately reached commercial operation. [6]

For data-center developers, this means:

- Long, uncertain timelines
- High upfront risk
- Exposure to shifting regulatory frameworks

These delays are already pushing operators toward behind-the-meter power and on-site generation. [7]

1.3 Land Scarcity and Siting Friction

Suitable land for traditional data centers is increasingly scarce in mature markets:

- Large parcels near substations and fiber
- Compatible zoning
- Community acceptance

Data-center land prices have risen sharply, up 23 percent for large U.S. parcels in 2024, as developers compete with buyers from outside the industry. Zoning, environmental review, and community opposition can add years to project timelines. [8]

1.4 The Shift Toward Distributed Data Centers

AI inference is moving closer to where data is generated. As AI shifts from training to inference, enterprises in manufacturing, healthcare, and retail are accelerating edge-computing investment to cut latency and keep data local. [9]

This shift favors distributed, high-performance nodes that deliver:

- Millisecond-level latency
- Local data sovereignty
- Local power solutions
- High availability independent of regional grid stress

LT350 is purpose-built for this emerging landscape.

02

LT350's Architecture: Modular AI Canopies in Existing Parking Lots

LT350 builds modular AI data-center canopies that install directly over existing parking lots without impacting parking. The ceiling of each canopy leverages the parking lot airspace and LT350's proprietary design to become a self-contained micro-datacenter by integrating with:

GPU CARTRIDGES

- Each cartridge houses a small bundle of high-end GPUs
- Hot-swappable and incrementally deployable

MEMORY CARTRIDGES

- Dedicated memory tiers for KV-cache offload
- Optimized for long-context and agentic inference workloads

BATTERY CARTRIDGES

- Behind-the-meter storage
- Peak-shaving and curtailment resilience

SOLAR INTEGRATION

- Canopy-mounted solar generation
- Directly feeds local storage and compute loads

LOCAL FIBER + SECURITY

- Short-run, high-bandwidth connectivity
- Segmented networking for regulated workloads

PHYSICAL ISOLATION

- Dedicated infrastructure per customer
- Supports healthcare, financial, and defense-aligned deployments

This creates a power-sovereign, latency-optimized, modular inference node deployable where traditional data centers cannot go.

03
Structural Advantages of LT350's Model

3.1 Power Sovereignty via Solar + Battery

LT350's behind-the-meter architecture provides:

- Peak-shaving
- Curtailment resilience
- Predictable power cost
- Reduced interconnection burden

As regulators push large loads to "bring their own power," hybrid solar-plus-storage models are gaining traction. [10]

3.2 Near-Zero Land Cost and Rapid Deployment

Parking lots are:

- Already paved
- Already utility-connected
- Already zoned
- Widely distributed

This allows LT350 to:

- Avoid land acquisition
- Avoid greenfield construction
- Minimize permitting friction
- Deploy in weeks or months, not years

This dual-use approach mirrors France, where a 2023 law requires large outdoor parking lots to be at least half covered by solar canopies or, since 2025, shade trees. [11]

3.3 Modular, Demand-Matched GPU Deployment

Traditional data centers deploy multimegawatt increments, creating utilization risk. LT350's 3-GPU cartridge model enables:

- Small, revenue-aligned scaling
- Rapid refresh cycles
- Reduced stranded capex

AI hardware is turning over faster: Amazon shortened the useful life of some servers in 2025, citing the pace of AI development. [12] Swappable cartridges keep LT350's fleet closer to the performance frontier than hyperscale deployments.

3.4 Proximity to High-Value, Regulated Customers

LT350 can deploy canopies within tens to hundreds of feet of:

- Hospitals
- Financial institutions
- Defense facilities
- Biotech campuses
- Autonomous-vehicle depots

This proximity enables:

- Deterministic low latency
- Physical data sovereignty
- Dedicated hardware
- Simplified compliance

Regulated customers face data-residency and sovereignty requirements, and sovereign, dedicated deployments are sold at higher prices than standard public cloud. [13]

3.5 Memory-Augmented Inference Nodes

As context windows expand and agentic workflows grow, inference bottlenecks shift to:

- KV-cache size
- Memory bandwidth
- Cross-GPU communication

LT350's memory cartridges:

- Offload KV-cache
- Increase throughput
- Support long-context workloads
- Improve economics for high-value inference

Inference is moving from general-purpose GPUs to workload-specific silicon, with providers routing work across specialized hardware. [14] This positions LT350 as a specialized inference fabric, not just a GPU host.

3.6 Alignment with Autonomous-Vehicle Infrastructure

Autonomous vehicles require:

- High-bandwidth data offload
- Frequent model updates
- Local inference
- Resilient power

LT350's canopies can be deployed in:

- Mall parking lots
- Airport structures
- Hospital campuses
- Fleet depots
- Convenience stores

Fleets offload large sensor-log volumes at depots, ideally while vehicles charge, and depots need to sit near where vehicles operate. [15] This creates a distributed compute mesh aligned with AV fleet distribution.

04

Market Alignment: Why LT350 Is Positioned for the Next Decade

4.1 Inference Is Moving to the Edge

Real-time AI requires:

- Millisecond latency
- Local data processing
- Local compliance
- Local power availability

Moving data to a centralized compute location is increasingly inefficient for real-time AI, and enterprises are bringing AI to the data instead.[16] Centralized cloud regions cannot meet all these constraints simultaneously.

4.2 Power Scarcity Favors Behind-the-Meter Models

Data centers are becoming a major driver of electricity demand growth. Regulators are signaling that large loads must either:

- Pay more for grid upgrades, or
- Provide their own power flexibility

Data centers are adding on-site generation and battery storage alongside grid supply.[17] LT350's architecture aligns with this shift.

4.3 Land Scarcity Favors Dual-Use Sites

Land near power is scarce, expensive, and politically sensitive.[18] Parking-lot deployment:

- Avoids land competition
- Avoids zoning battles
- Improves the underlying asset

This is a structural advantage over hyperscalers.

05
Conclusion

LT350 is a new category of AI infrastructure:

- Distributed
- Power sovereign
- Memory augmented
- Latency optimized
- Modular
- Parking lot deployable

As AI shifts from centralized training to pervasive, real-time inference, LT350 believes it is positioned to become the default edge-compute fabric for hospitals, financial institutions, defense and biotech campuses, and autonomous vehicle fleets.

Notes and Sources

1. International Energy Agency (IEA), Electricity 2024: Analysis and forecast to 2026, January 2024, executive summary. Source · iea.org
2. CBRE, North America Data Center Trends H1 2024, August 19, 2024. Source · cbre.com; McKinsey & Company, “AI power: Expanding data center capacity to meet growing demand,” October 29, 2024. Source · mckinsey.com
3. International Energy Agency (IEA), Energy and AI, April 2025, executive summary. Source · iea.org
4. McKinsey & Company, “AI power: Expanding data center capacity to meet growing demand,” October 29, 2024. Source · mckinsey.com
5. Shehabi et al., 2024 United States Data Center Energy Usage Report (LBNL-2001637), Lawrence Berkeley National Laboratory, December 2024, executive summary. Source · eta-publications.lbl.gov
6. Rand et al., Queued Up: 2024 Edition, Characteristics of Power Plants Seeking Transmission Interconnection As of the End of 2023, Lawrence Berkeley National Laboratory, April 2024. Source · emp.lbl.gov
7. DataCenterDynamics, “Diversity of power: the biggest data center energy stories of 2024,” January 11, 2025. Source · datacenterdynamics.com; JLL, 2026 Global Data Center Outlook, January 2026. Source · jll.com
8. Building Design + Construction, “Data center demand pivots to larger sites, emerging markets amid power and cost challenges,” on the Cushman & Wakefield Data Center Development Cost Guide, December 13, 2024. Source · bdcnetwork.com; JLL, 2026 Global Data Center Outlook, January 2026. Source · jll.com
9. IDC, “Worldwide Spending on Edge Computing Forecast to Reach $378 Billion in 2028,” September 10, 2024. Source · idc.com
10. JLL, 2026 Global Data Center Outlook, January 2026. Source · jll.com; BloombergNEF, “Battery Storage Costs Hit Record Lows as Costs of Other Clean Power Technologies Increased,” February 18, 2026. Source · about.bnef.com
11. pv magazine, “French parliament eases solar mandate for large parking lots,” May 19, 2025. Source · pv-magazine.com
12. The Register, “Hardware quality problems and server supply chain kinks slow Amazon’s $100 billion AI build,” February 7, 2025. Source · theregister.com
13. Deloitte, “Keeping it local: Cloud sovereignty a major focus of the future,” TMT Predictions 2024, November 29, 2023. Source · deloitte.com
14. Mark Donnigan, “The Era of Specialized Hardware: AI Inference and Video Encoding Shifts,” NETINT, January 20, 2026. Source · netint.com
15. Kevin Chen, “Real estate is one of the hardest open problems in scaled self driving,” September 3, 2024. Source · kevinchen.co
16. Deloitte, “Hardware is eating the world,” Tech Trends 2025, December 11, 2024. Source · deloitte.com
17. Utility Dive, “BNEF ups US energy storage deployment forecast by 30%,” October 27, 2026. Source · utilitydive.com
18. JLL, “Search for power pushes emerging data center markets into the spotlight,” February 28, 2024. Source · jll.com; Bisnow, “Rising Land Prices Spell Trouble For Some Data Center Developers,” July 7, 2024. Source · bisnow.com



DISTRIBUTED AI DATA CENTERS

lt350.com · info@LT350.com · Boulder, Colorado

Distributed, Power-Sovereign AI Infrastructure for the Inference Economy. LT350-WP-01, revision B, revised November 12, 2026. Every figure in this paper is sourced in its notes, and the current edition is published at lt350.com/whitepaper. © 2026 LT350, LLC.