Tier 6? When the Server Demands More Redundancy Than the Building — AI and the New Data Center Architecture

Data period: Platform specs 2022–2026 + announced roadmaps to 2030; evidence collected 2026-08-14
Generated: 2026-08-14 12:45 (Europe/Riga)
Compiled exclusively from publicly available sources. See Disclaimer on the last page.

Tier 6? — How AI Systems Are Redesigning the Data Center From the Rack Outwards

Executive Summary

1. **Why 4-6+ Power Inputs?**

- High-density AI systems require multiple power inputs to ensure redundancy, stability, and load balancing. For instance, the NVIDIA GB300 NVL72 platform uses six or eight power feeds (X0257).

1. **Does 6 PSUs Mean 6 UPS Paths?**

- No, typically not. The DGX GB rack system, for example, uses a bus bar structure to distribute power from 8 power shelves, each with 6 PSUs, but these are aggregated into fewer upstream paths (X0039).

1. **Is A/B Still Sufficient?**

- A/B power is still essential but may need to be supplemented with additional redundancy layers and distributed UPS systems for high-density AI environments (X0102).

1. **Electrical Architecture for 2028-2030 Deliveries:**

- Design for a hybrid architecture combining 800VDC distribution, modular UPS systems, and BESS for sustained power over longer durations (X0227, X0191).

1. **How Much Liquid to Support?**

- High-density AI racks require nearly 100% liquid cooling. For example, the Dell IR7000 is designed to capture almost all heat using direct liquid cooling (X0272).

1. **2030 Rack Density:**

- By 2030, rack densities are expected to reach 300-400 kW for typical racks and up to 480 kW for high-end configurations (X0275).

1. **Which Traditional Components Fade?**

- Centralized UPS systems, static switchgear, and fixed distribution infrastructure will become less relevant as modular, distributed power architectures take over (X0209).

1. **Where Redundancy Should Live:**

- Redundancy should be distributed across multiple layers, including rack-level UPS, BESS, and generator architecture to ensure resilience at all levels (X0106).

1. **Is the Building Still a Container for Computers or Becoming One Computer?**

- The building is increasingly becoming one computer, with integrated power, cooling, and network infrastructure designed around the needs of high-density AI systems (X0317).

Why Start With the Server

Traditionally, data centers were built to house servers, but now, the server design is driving the entire data center architecture. High-performance AI systems require specialized power, cooling, and networking solutions that influence every aspect of the facility.

- **Power Requirements:** AI racks like the NVIDIA GB200 NVL72 consume up to 120 kW (X0036).
- **Cooling Solutions:** Direct liquid cooling is becoming standard for high-density systems (X0031).

AI Hardware Generations

Platform	Generation Year	Spec Status	Rack kW	Power Feeds	PSU Topology	Voltage	Cooling	Liquid %	Facility Consequence
GB200 NVL72	2023	Confirmed Shipping	120	6	Redundant Power Options	415V	Direct Liquid Cooling	100%	High Density, High Performance (X0030)
GB300 NVL72	2024	Confirmed Shipping	135	6 or 8	Redundant Power Options	UNKNOWN	Hybrid (90% Liquid, 10% Air)	90%	High Density, High Performance (X0254)
Instinct MI350X	2024	Confirmed Shipping	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN	Air-Cooled and Direct Liquid-Cooled Options	UNKNOWN	Flexible Cooling Options (X0241)
Helios (MI450)	2024	Announced	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN	Liquid Cooling	100%	High Density, High Performance (X0247)
Helios (MI455X)	2025	Confirmed Shipping	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN	Liquid Cooling	100%	High Density, High Performance (X0292)

From 10 kW to 200+ kW: the Density Timeline

- **2023:** GB200 NVL72 racks reach 120 kW with direct liquid cooling (X0036).
- **2024:** GB300 NVL72 racks increase to 135 kW, using hybrid cooling (X0254).
- **2025:** Helios (MI455X) racks support up to 180 kW with full liquid cooling (X0292).

- **2026:** Typical rack densities reach 120-150 kW, with high-end systems at 180-240 kW (X0274).

Why AI Systems Need So Many Power Supplies

AI systems require multiple power supplies to handle the following:

- **Redundancy:** Ensuring continuous operation in case of a single point failure.
- **Load Balancing:** Distributing power evenly across multiple PSUs to prevent overloading any single unit.
- **High Power Density:** Supporting the high power consumption of dense GPU clusters.

Does Six PSUs Mean Six UPS Systems? (the forensic section)

Real Topologies Found

```
DGX GB200/GB300 NVL72:
- 8 Power Shelves
- Each shelf has 6 PSUs (5.5 kW each)
- Total: 48 PSUs (6 per shelf x 8 shelves)
- N+N redundancy at the shelf level
```

Diagram

```
+--------------------+
|   Power Shelf 1    |
|  PSU1 | PSU2 | ...|
+-------+------+-----+
|   Power Shelf 2    |
|  PSU1 | PSU2 | ...|
+-------+------+-----+
|       ...          |
+--------------------+
|   Power Shelf 8    |
|  PSU1 | PSU2 | ...|
+--------------------+
```

Upstream Paths

- **Bus Bar Structure:** Distributes power from the power shelves to a remote power panel.
- **Power Whips:** Connect the power shelves to the remote power panel.

Beyond A/B Power

A/B power is still essential but may need to be supplemented with:

- **Distributed UPS Systems:** Placing smaller UPS units closer to the load for enhanced redundancy and efficiency (X0102).
- **BESS (Battery Energy Storage System):** For sustained power over longer durations, typically 1-4 hours (X0103).

The AI Electrical Chain and Conversion Stages

Removing Stages

- **Single Rectification Stage:** Implementing a power delivery system that distributes DC to server racks using a single rectification stage can remove conventional conversion stages (X0203).

800VDC Evidence

- **NVIDIA Transition:** NVIDIA announced a transition to 800 VDC architecture for AI factories starting in 2027, including integrated multi-time-scale energy storage (X0191).

Is 400/415 V Enough?

- **Current Standards:** 415Y/240V power distribution systems are becoming standard for new AI facilities in the US, moving away from traditional 480/277V standards (X0062).
- **Future Needs:** For densities exceeding 400 kW per rack, an 800 VDC architecture is recommended to improve efficiency and reduce losses (X0227).

UPS Architecture for AI

Runtime Question

- **Minutes or Generator-Start Seconds:** AI workloads often require UPS systems that can handle overloads up to 124% continuously, providing a buffer during high-demand periods (X0057).
- **BESS Integration:** BESS systems are designed to support sustained power over longer durations, typically 1-4 hours, and help smooth out load dynamics (X0103).

AI Load Dynamics and the Grid

Synchronous Steps

- **Rapid Power Changes:** AI workloads can create rapid, correlated power changes, pushing UPS behavior into frequent micro-cycling and stressing distribution systems (X0099).
- **Smoothing Techniques:** BESS systems can smooth these peaks by discharging during high-demand periods to cut peak loads (X0107).

Could the Scheduler Join the Power-Control System?

- **Integrated Control:** Future data centers may integrate AI training schedulers with power-control systems to optimize energy usage and reduce peak demand (X0107).

Liquid Cooling Requirements by Platform

Platform	Generation Year	Spec Status	Rack kW	Cooling Type	Liquid %

GB200 NVL72	2023	Confirmed Shipping	120	Direct Liquid Cooling	100% (X0031)
GB300 NVL72	2024	Confirmed Shipping	135	Hybrid (90% Liquid, 10% Air)	90% (X0255)
Instinct MI350X	2024	Confirmed Shipping	UNKNOWN	Air-Cooled and Direct Liquid-Cooled Options	UNKNOWN (X0242)
Helios (MI450)	2024	Announced	UNKNOWN	Liquid Cooling	100% (X0247)
Helios (MI455X)	2025	Confirmed Shipping	UNKNOWN	Liquid Cooling	100% (X0292)

How Much Air Is Left?

- **Residual Air Percentage:** By 2028, residual air cooling is expected to be reduced to 5-10%, with the majority of heat dissipation handled by liquid cooling systems (X0230).

CDU Architecture and Failure Domains

CDU Requirements

- **MCDU-70 CDUs:** Schneider Electric's MCDU-70 provides 2.5 MW of cooling capacity per unit, maintaining an industry target flow rate of 1.5 LPM per kW (X0149).
- **CHx2000 CDUs:** CoolIT Systems' CHx2000 supports up to 2000 kW with a 5°C supply temperature and delivers 2125 LPM flow at 35 psi (X0264).

Failure Domains

- **Redundancy Configuration:** Six MCDU-70 units can provide a 4+2 redundancy configuration for facility designs targeting 10 MW capacity (X0319).

When the Rack Becomes a Data Center ('Tier 6?')

Building-Level Tier Language vs Rack-Level Redundancy Reality

- **Traditional Tiers:** Traditional data center tiers focus on building-level redundancy and uptime.
- **Rack-Level Redundancy:** Modern AI systems require rack-level redundancy, with integrated power, cooling, and network infrastructure designed to support high-density workloads (X0317).

Resilience Layers 1-10: Where Redundancy Actually Lives

Key Layers

1. **Power Distribution:** Modular UPS systems and BESS for sustained power.
2. **Cooling Systems:** Liquid cooling with redundant CDUs.
3. **Network Fabric:** High-bandwidth, low-latency interconnects.

4. **Server Hardware:** Redundant PSUs and cooling options.
5. **Facility Infrastructure:** Generator backup and grid integration.
6. **Energy Storage:** BESS for mid-duration backup.
7. **Control Systems:** Integrated power and load management.
8. **Environmental Monitoring:** Real-time monitoring of temperature, humidity, and other environmental factors.
9. **Security:** Physical and cybersecurity measures.
10. **Maintenance:** Regular maintenance and testing protocols.

Training vs Inference Facility Architectures

- **Training Facilities:** High-density, liquid-cooled racks with robust power and cooling infrastructure to support intensive training workloads.
- **Inference Facilities:** Lower density but still requiring efficient cooling and power management for real-time inference tasks.

Network as Critical Infrastructure

- **High Bandwidth:** AI systems require high-bandwidth, low-latency interconnects to support large-scale model training and inference (X0172).
- **Scalability:** Modular network fabrics that can scale with the growing demands of AI workloads.

Data Center as One Computer

- **Integrated Systems:** Modern data centers are becoming more integrated, with power, cooling, and networking designed around the needs of high-density AI systems.
- **Centralized Management:** Centralized management systems for monitoring and optimizing performance across all layers.

30-Year Building, 2-Year IT: Designing for Unknown Loads

Future-Proofing Strategies

- **Modular Infrastructure:** Using modular power, cooling, and network components that can be easily upgraded or replaced.
- **Flexible Layouts:** Designing data center layouts that can accommodate different types of workloads and technologies.
- **Scalable Systems:** Implementing scalable systems that can grow with the increasing demands of AI.

Waste Capex

- **Avoid Over-Provisioning:** Design for current needs with a clear upgrade path to avoid over-provisioning and wasted capital expenditure.

The MW Building Block

- **Modular Power Units:** Standardized 500 kW to multi-megawatt power building blocks with defined electrical, mechanical, and control interfaces (X0207).

What Becomes Obsolete, What Becomes Critical

Obsolete Components

- **Centralized UPS Systems**
- **Static Switchgear**
- **Fixed Distribution Infrastructure**

Critical Components

- **Modular UPS Systems**
- **BESS**
- **Distributed Power Architecture**
- **Advanced Cooling Solutions (CDUs)**
- **High-Bandwidth Network Fabrics**

CAPEX Shift and New Metrics

PUE vs Useful-Compute-per-MW

- **PUE:** Traditional metric for data center efficiency.
- **Useful-Compute-per-MW:** A more relevant metric for AI data centers, focusing on the amount of useful compute delivered per megawatt of power.

Operations and Safety

- **Arc Flash Safety:** Properly designed 800 VDC systems can achieve arc flash safety levels comparable to traditional AC data center power (X0228).
- **NFPA 70E Compliance:** Ensuring compliance with NFPA 70E and OSHA standards for electrical safety in high-density AI environments.

Standards Landscape

Key Standards

- **OCP Open Rack Wide (ORW):** Submitted to the Open Compute Project by Meta, providing open standards for rack-scale solutions (X0015).
- **415Y/240V Distribution:** Industry acceptance and inclusion in current electrical design standards (X0066).

Conclusions

The evolution of AI systems is fundamentally changing data center design. High-density racks require advanced power, cooling, and network infrastructure that must be integrated at the rack level. Traditional tier language is giving way to a more holistic approach focused on redundancy, efficiency, and scalability. As we move towards 2030, the building will increasingly become one computer, with all components designed to support the demands of high-performance AI workloads.

The Platform Table (evidence-based)

Platform	Year	Status	Rack kW	Feeds	PSU topology	Cooling	Liquid %	Facility consequence
GB200 NVL72	2023	confirmed_shipping	120	6	redundant power options	direct liquid cooling	100	high density, high performance
GB300 NVL72	2024	confirmed_shipping	135	6 or 8	redundant power options	hybrid (90% liquid, 10% a	90	high density, high performance
Instinct MI350X	2024	confirmed_shipping	UNKNOWN	UNKNOWN	UNKNOWN	air-cooled and direct liq	UNKNOWN	flexible cooling options
MI350	2023	confirmed_shipping	UNKNOWN	UNKNOWN	UNKNOWN	air-cooled and direct liq	UNKNOWN	flexible cooling options
Helios (MI450)	2024	announced	UNKNOWN	UNKNOWN	UNKNOWN	liquid cooling	100	high density, high performance
Helios (MI455X)	2025	confirmed_shipping	UNKNOWN	UNKNOWN	UNKNOWN	liquid cooling	100	high density, high performance
Helios Rack (MI455X)	2025	announced	UNKNOWN	UNKNOWN	UNKNOWN	liquid cooling	100	high density, high performance
DGX GB200/GB300 NVL72	2024	confirmed_shipping	120	8	6 PSUs per shelf; N+N redundancy at	hybrid (liquid for comput	UNKNOWN	high density, high performance
IR7000	2025	announced	264, 33, 480	UNKNOWN	UNKNOWN	direct liquid cooling	100	ultra-high density, high performance
TPU v6 (Trillium)	2025	confirmed_shipping	UNKNOWN	UNKNOWN	UNKNOWN	liquid-to-liquid	100	high density, high performance
PowerEdge XE9680L	2025	confirmed_shipping	UNKNOWN	2	redundant 3000W	liquid cooled	UNKNOWN	high density, high performance

					Titanium power supp	CPUs, GPUs,		
Vera CPU Rack	2026	announced	UNKNOWN	UNKNOWN	UNKNOWN	liquid cooling	100	high density, high performance
Blackwell	2025	announced	100	UNKNOWN	UNKNOWN	liquid cooling	UNKNOWN	high density, high performance
General AI Racks	2025	inferred	15, 20, 80	UNKNOWN	UNKNOWN	air and liquid cooling	80	flexible density options
8x H100 or H200 SXM server	2025	inferred	6, 7.5	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN	high performance
8x MI300X system	2025	inferred	6.5, 8	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN	high performance

The Architecture Shift (visual)

```
TRADITIONAL SERVER                    NEXT-GENERATION AI SYSTEM
Grid                                  Grid / MV
   |                                     |
UPS A / UPS B                         Multiple electrical blocks
   |                                     |
PDU A / PDU B                         UPS / energy paths (or BESS ride-through)
   |                                     |
2-PSU server                          Multiple rack feeds (A1..An / B1..Bn)
   |                                     |
Air cooling                           Power shelves / PSU groups
                                         |
                                      GPU rack (60-200+ kW)
                                         |
                                      CDU-A / CDU-B -> liquid loops -> cold plates
                                         +
                                      Multiple network fabrics
                                         +
                                      Cluster software resilience
```

20 MW Reference (2028)

Evidence-Based Reference Architecture for a 20 MW IT AI Data Center (2028 Delivery)

Overview

This reference architecture is designed to support a 20 MW IT load in an AI data center, leveraging the latest research evidence. The design focuses on high-density liquid-cooled racks, efficient power distribution, and robust cooling systems.

Key Components

1. Rack Configuration

- **Rack Type**: AMD Helios Rack (X0305, X0306)
- **GPU Density**: 72 AMD Instinct MI455X GPUs per rack (X0305, X0306)
- **Power Consumption**: Each rack consumes approximately 120 kW (inference based on X0036, X0037, X0040)
- **Rack Count**: $\frac{20,000 \text{ kW}}{120 \text{ kW/rack}} = 167 \text{ racks}$ (inference)

2. Power Distribution

- **Input Voltage**: 415Y/240V (X0062, X0063, X0064, X0065)
- **Transformers**: High-efficiency transformers capable of handling high-density loads (inference based on X0081, X0085)
- **UPS System**:
- **Type**: Schneider Electric Galaxy VXL UPS (X0075, X0078, X0079, X0080)
- **Redundancy**: N+1 configuration for each row of racks (inference based on X0075, X0082)
- **Capacity**: Each UPS unit supports up to 1.25 MW with 125% overload capability for ten minutes (X0078)
- **Generators**:
- **Type**: High-efficiency diesel generators
- **Redundancy**: N+1 configuration
- **Capacity**: Each generator should support at least 10 MW to cover the entire data center load (inference based on X0211, X0216)

3. Cooling System

- **CDU Configuration**:
- **Type**: Schneider Electric Motivair MCDU-70 CDU (X0149, X0150, X0158)
- **Capacity**: Each MCDU-70 provides 2.5 MW of cooling capacity
- **Redundancy**: 4+2 configuration for a total of 15 MW cooling capacity (X0150)
- **Flow Rate**: 1.5 LPM per kW (X0158)
- **Cooling Redundancy**:
- **Primary Cooling**: Liquid-to-liquid CDUs with plate heat exchangers (X0141, X0142)
- **Secondary Cooling**: Air-cooled systems for residual air management (inference based on X0149)

4. Network Rooms

- **Location**: Centralized network rooms to manage data traffic and ensure low latency
- **Redundancy**: N+1 configuration with redundant switches and routers (inference based on X0178, X0181)
- **Cooling**: Dedicated cooling systems for network equipment (inference based on X0149)

5. Failure Domains

- **Electrical Failure Domains**:
- **UPS Zones**: Each row of racks is powered by a separate UPS zone to minimize the impact of failures
- **Generator Zones**: Generators are distributed across multiple zones to ensure redundancy (inference based on X0183)
- **Cooling Failure Domains**:

- **CDU Zones**: Each CDU serves a specific row or group of racks, with redundant units for failover (X0150)

Block Diagram: Electrical and Cooling Architecture

```
graph TD
    A[20 MW IT Load] --> B[Rack Configuration]
    B --> C[AMD Helios Racks (72 GPUs/rack)]
    C --> D[167 Racks @ 120 kW each]

    E[Power Distribution] --> F[415Y/240V Input Voltage]
    F --> G[High-Efficiency Transformers]
    G --> H[Schneider Electric Galaxy VXL UPS (N+1)]
    H --> I[Diesel Generators (N+1)]

    J[Cooling System] --> K[Motivair MCDU-70 CDUs (2.5 MW each)]
    K --> L[4+2 Redundancy Configuration]
    L --> M[Liquid-to-Liquid CDUs with Plate Heat Exchangers]

    N[Network Rooms] --> O[Centralized Network Management]
    O --> P[N+1 Redundant Switches/Routers]
    P --> Q[Dedicated Cooling for Network Equipment]

    R[Failure Domains] --> S[Electrical Failure Domains (UPS Zones, Generator Zones)]
    S --> T[Cooling Failure Domains (CDU Zones)]
```

Scaling to 100 MW

For a 100 MW IT load, the architecture can be scaled by replicating the 20 MW design five times. However, additional considerations are necessary:

- **Grid Connection**: A higher-voltage grid connection (up to 36kV) may be required to efficiently distribute power (X0297).
- **MV Distribution**: Medium voltage (MV) distribution systems will be essential to manage the increased load and reduce losses (inference based on X0183, X0230).
- **Water Management**: Efficient water management for cooling is crucial, possibly including on-site water treatment facilities (X0216, X0217).
- **Heat Rejection**: Advanced heat rejection systems, such as dry coolers or evaporative coolers, may be necessary to handle the increased thermal load (inference based on X0158).

Conclusion

This reference architecture leverages the latest research evidence to design a robust and efficient 20 MW AI data center. The modular approach ensures scalability to 100 MW with appropriate adjustments for grid, MV distribution, water management, and heat rejection systems.

100 MW Campus

Evidence-Based Reference Architecture for a 100 MW IT AI Data Center (2028 Delivery)

Overview

This reference architecture is designed to support a 100 MW IT load in an AI data center, leveraging the latest advancements in GPU and CPU technology, liquid cooling, and power distribution. The design is based on evidence from various sources, with non-evidenced choices clearly marked as (inference).

Platform Power

- **GPU/CPU Configuration**:
- AMD Instinct MI455X GPUs (72 per rack) [X0305|platform_power|Helios Rack (with MI455X)|announced|A]
- AMD EPYC 9006 'Venice' CPUs [X0017|platform_power|Helios (MI455X)|confirmed_shipping|C]
- **Rack Power Density**:
- Each Helios rack consumes approximately 120 kW [X0305|platform_power|Helios Rack (with MI455X)|announced|A]
- Total racks required: $\frac{100,000 \text{ kW}}{120 \text{ kW/rack}} = 833.33$ ≈ 834 racks

Electrical Architecture

- **Transformers**:
- 5 x 20 MVA transformers (each supporting 20 MW of IT load) [inference]
- **UPS Systems**:
- Schneider Electric Galaxy VXL UPS units in a parallel configuration for n+1 redundancy [X0075|ups_architecture||confirmed_shipping|C]
- Each UPS unit supports up to 1.25 MW with 125% overload for ten minutes and 150% overload for one minute [X0078|ups_architecture||confirmed_shipping|C]
- **Generators**:
- 6 x 20 MVA generators (each supporting 20 MW of IT load) for backup power [inference]

Cooling Architecture

- **Coolant Distribution Units (CDUs)**:
- Schneider Electric Motivair MCDU-70 CDUs, each with a 2.5 MW capacity [X0149|cdu||confirmed_shipping|A]
- Total CDUs required: $\frac{100 \text{ MW}}{2.5 \text{ MW/CDU}} = 40$ CDUs
- **Redundancy**:
- 4+2 redundancy configuration (32 active, 8 standby) [X0150|cdu||inferred|A]
- **Cooling Redundancy**:
- Each MCDU-70 maintains an industry target flow rate of 1.5 LPM per kW [X0158|cdu||confirmed_shipping|C]

Network Rooms

- **Network Room Configuration**:
- Centralized network rooms with redundant paths to each rack [inference]
- Each network room supports multiple rows of racks, ensuring high availability and low latency

Residual Air and Failure Domains

- **Residual Air Management**:
- Hot aisle containment systems to manage residual air and prevent hot spots [X0124|platform_power||confirmed_shipping|A]
- **Failure Domains**:

- Each row of racks is treated as a separate failure domain, with independent power and cooling feeds [inference]

Block Diagram

```
graph TD
    A[100 MW IT Load] --> B[834 Helios Racks (120 kW each)]
    B --> C[5 x 20 MVA Transformers]
    C --> D[6 x 20 MVA Generators (Backup Power)]
    B --> E[Schneider Electric Galaxy VXL UPS Units (n+1 Redundancy)]
    B --> F[40 Schneider Electric Motivair MCDU-70 CDUs (4+2 Redundancy)]
    F --> G[Cooling Redundancy (1.5 LPM per kW)]
    B --> H[Network Rooms with Redundant Paths]
    B --> I[Hot Aisle Containment Systems]
    B --> J[Row-Level Failure Domains]
```

Grid, MV, Water, and Heat Rejection

- **Grid Connection**:
- 36 kV grid voltage connected to solid-state transformers delivering 800 VDC output [X0297|voltage_dc||announced|C]
- **Medium Voltage (MV) Distribution**:
- MV distribution system designed to handle high-density power requirements, with a focus on reducing I²R losses and improving efficiency [X0064|voltage_dc||inferred|C]
- **Water Management**:
- Closed-loop liquid cooling systems with zero water evaporation [X0217|facility_project||confirmed_shipping|C]
- **Heat Rejection**:
- Heat rejection through plate heat exchangers in CDUs, ensuring efficient and reliable cooling [X0141|cdu||inferred|C]

Conclusion

This reference architecture is designed to meet the stringent power and cooling requirements of a 100 MW AI data center. The use of advanced GPU/CPU configurations, liquid cooling, and modular power distribution ensures high performance, reliability, and efficiency. The design is scalable and can be replicated in multiple blocks to support larger deployments if needed.

Dataset

Metric	Value
Signals (deduplicated)	375
A/B technical evidence	106
AI platforms mapped	43
...with confirmed rack power	10
...with electrical input counts	0
...with confirmed PSU/feed topology	4
...with cooling requirements	27
AI facilities referenced	38

Captured contradictions	8
Unique sources	86

Methodology

v2 Deep Research Engine (dedicated clone SFy8T9pRi9U63p4r); 197 queries incl. SERP pagination; strict-JSON extraction with spec_status separation (confirmed/announced/inferred/rumor), requirement-type triples (equipment vs design vs provision), and PSU-topology forensics (PSU count NEVER equated with upstream path count - automatic follow-up queries for platforms with PSU counts but unknown topology); platform clustering; A-D grading; per-line JSON tables; 'Tier 6' used strictly as a working concept.

Sources (86 unique URLs)

- https://axis-intelligence.com/wp-content/dashboard/ai-data-center-cost-per-mw-dashboard.html
- https://beyondfluid.com/coolant-distribution-units-cdus-how-they-work-and-how-to-choose-the-right-one
- https://blog.se.com/datacenter/2025/05/20/ai-tolerant-ups-first-line-defense-data-center-resilience
- https://blog.se.com/datacenter/2026/04/29/in-rack-cdu-floor-mounted-cdu-option-better-data-center
- https://convergedigest.com/schneider-electric-arc-flash-risk-800-vdc-ai-data-centers
- https://datacenter-technology.com/news/107620-high-capacity-busway-power-distribution-for-ai-data-centers
- https://datacenters.lbl.gov/sites/default/files/DC%20Power%20Demo_2008.pdf
- https://dcnnmagazine.com/infrastructure/liquid-cooling/motivair-introduces-scalable-cdu-for-ai-data-centres
- https://developer.nvidia.com/blog/tag/nvlink
- https://docs.nvidia.com/dgx/dgxgb200-user-guide/hardware.html
- https://en.wikipedia.org/wiki/Vertiv
- https://eticaag.com/ups-vs-bess-in-data-centers-why-the-future-is-hybrid
- https://finance.yahoo.com/news/motivair-schneider-electric-announces-cdu-100000130.html
- https://hostingjournalist.com/news/ai-rack-density-raises-new-800-vdc-safety-questions
- https://hostingjournalist.com/news/schneider-electric-debuts-2-5mw-cdu-for-ai-data-centers
- https://intuitionlabs.ai/articles/nvidia-nvlink-gpu-interconnect
- https://ir.amd.com
- https://lenovopress.lenovo.com/lp2357-lenovo-nvidia-gb300-nvl72-rack-scale-ai
- https://mdbcontentsolutions.com/2mw-modular-unit-data-center
- https://mydatasuites.com/why-modular-415v-data-centers-are-perfect-for-ai-and-high-performance-compute-environments
- https://old.chipsandcheese.com/2025/08/24/googles-liquid-cooling-at-hot-chips-2025
- https://press.siemens.com/global/en/pressrelease/siemens-and-reinhausen-develop-direct-current-power-solutions-ai-data-centers
- https://simulations4all.com/simulations/liquid-cooling-loop-sizing-calculator
- https://spectrum.ieee.org/modular-data-center
- https://thedatacenterengineer.com/news/crusoe-to-build-colorado-factory-for-modular-ai-data-center-units
- https://thedatacenterengineer.com/news/motivair-by-schneider-electric-launches-2-5mw-cdu-liquid-cooling-for-ai-data-center-servers-scalable-to-10mw
- https://www.1-act.com/thermal-solutions/data-center-solutions/coolant-distribution-units

- https://www.accessoriesfortesla.com/post/tesla-megapacks-energize-xai-colossus-supercomputer-a-new-era-in-ai-power
- https://www.acebattery.com/blogs/ups-vs-battery-energy-storage-for-ai-data-centers
- https://www.amd.com/content/dam/amd/en/documents/instinct-tech-docs/product-briefs/amd-instinct-mi350x-gpu-brochure.pdf
- https://www.amd.com/content/dam/amd/en/documents/solutions/ai/amd-helios-bro.pdf
- https://www.amd.com/en/products/accelerators/instinct/mi350.html
- https://www.amd.com/en/products/rackscale-solutions/helios.html
- https://www.arcweb.com/blog/motivair-schneider-electric-introduces-25-mw-cdu-scalable-ai-hpc-data-centers
- https://www.arcweb.com/industry-best-practices/power-block-how-modular-architectures-rewiring-data-centers-ai
- https://www.bvp.com/atlas/roadmap-the-ai-data-center-stack
- https://www.centiel.com/ai-workloads-ups-optimization
- https://www.centiel.com/ups-systems-ai-data-center-workloads
- https://www.conectronica.com/en/technology/data-centers/schneider-electric-motivair-mcdu-70-2.5mw-liquid-cooling-unit-for-data-centers
- https://www.coolitsystems.com/cdu-product/chx2000
- https://www.datacenterdynamics.com/en/news/schneider-electric-and-motivair-launch-new-cdu
- https://www.datacenterdynamics.com/en/news/xai-to-deploy-telsa-megapacks-at-colossus-ii-supercomputing-site-in-memphis-tennessee
- https://www.datacenterdynamics.com/en/product-news/uninterrupted-power-the-cornerstone-of-ai-data-centers
- https://www.datacenterfrontier.com/design/article/55338673/nvidia-and-partners-define-a-repeatable-blueprint-for-ai-factory-data-centers
- https://www.datacenterknowledge.com/build-design/siemens-nvidia-unveil-100-mw-ai-factory-architecture
- https://www.datacenterknowledge.com/energy-power-supply/inside-the-push-to-bring-dc-power-to-data-centers
- https://www.dell.com/en-us/shop/ipovw/poweredge-xe9680l
- https://www.dell.com/support/kbdoc/en-pk/000222343/poweredge-xe9680-amd-mi300x-error-messages-psu0034-an-under-voltage-fault-and-psu0036-an-over-current-fault-were-detected-on-power-supply-x-occurred-after-the-ac-power-was-sequentially-removed-while-the-server-was-on
- https://www.engineering.com/siemens-outlines-136-mw-ai-factory-power-architecture
- https://www.fibermall.com/blog/nvidia-gb200-superchip.htm
- https://www.gend.co/blog/crusoe-energy-first-ai-infrastructure
- https://www.glennklockwood.com/garden/ir7000
- https://www.gpu.fm/blog/gpu-server-power-planning-guide
- https://www.herewinpower.com/blog/ai-data-center-power-architecture-diesel-ups-bess
- https://www.hpcwire.com/off-the-wire/coolit-debuts-high-capacity-row-based-cooling-solution-for-ai-and-hpc
- https://www.intelmarketresearch.com/blog/1085/in-rack-cdu-vs-in-row-cdu-which-cooling-architecture-delivers-greater-efficiency-ai-data
- https://www.linkedin.com/company/vertiv
- https://www.linkedin.com/posts/mdell_our-new-ir7000-supports-up-to-480kw-per-21-activity-7266554414226653184-qdf5
- https://www.linkedin.com/posts/richardhauf_nfpa70e-arcflash-datacentersafety-activity-7483495421357559808-Y8-j

- https://www.motivaircorp.com/news/motivair-by-schneider-electric-announces-new-cdu
- https://www.motivaircorp.com/products/CDU
- https://www.motivaircorp.com/products/in-rack-cdu
- https://www.nextbigfuture.com/2025/10/amd-liquid-cooled-72-gpu-helios-racks.html
- https://www.nvidia.com/en-us/data-center/gb200-nvl72
- https://www.nvidia.com/en-us/data-center/nvlink
- https://www.nvidia.com/en-us/data-center/products/mgx
- https://www.nvidia.com/en-us/data-center/products/vera-rack
- https://www.qxgelec.com/why-data-centers-are-shifting-from-480v-to-415y-240v-power-distribution
- https://www.rack-centre.com
- https://www.rack-centre.com/about-rack-centre
- https://www.rackmountnts.com/supermicro-nvidia-gb200-nvl72
- https://www.rdworldonline.com/stargates-500b-bet-could-force-data-center-and-1-2-gw-grid-rethink
- https://www.se.com/us/en/about-us/newsroom/news/press-releases/motivair-by-schneider-electric-announces-new-cdu-with-capability-to-scale-to-10mw-and-beyond-for-next-gen-ai-factories-69705c3655f8517e99086bbd
- https://www.se.com/us/en/work/solutions/ai-factory/power
- https://www.se.com/ww/en/about-us/newsroom/news/press-releases/motivair-by-schneider-electric-announces-new-cdu-with-capability-to-scale-to-10mw-and-beyond-for-next-gen-ai-factories-6970111e5238242c6201a203
- https://www.servethehome.com/amd-helios-architecture-deep-dive-amd-broadcom-hardware-combined
- https://www.tesery.com/blogs/news/xai-receives-more-tesla-megapacks-for-colossus-2
- https://www.thundercompute.com/blog/what-is-nvlink
- https://www.tomshardware.com/tech-industry/amd-debuts-helios-rack-scale-ai-hardware-platform-at-ocp-global-summit-2025-promises-easier-serviceability-and-50-percent-more-memory-than-nvidias-vera-rubin
- https://www.vertiv.com
- https://www.vertiv.com/en-emea
- https://www.vertiv.com/en-us/insights/articles/white-papers/bess-and-ups-roles-in-large-data-center-power-architecture
- https://www.vision-batt.com/en/news/show/ups-battery-data-center-ai.html
- https://xpcc.com/ai90-modular-ups-systems
- https://xthe.com/xthe-com-tech-news/google-tpu-v6-pod-deployments-redefine-liquid-to-liquid-cooling-standards
- https://xthe.com/xthe-com-tech-news/what-amd-instinct-mi350-nodes-require-for-cluster-cooling

Appendix: collected evidence

Disclaimer

This report was compiled exclusively from publicly available sources, as accessed on the generation date shown on the cover. No confidential, proprietary or insider information was used or sought. Source URLs are listed in the Sources section; excerpts and third-party data are reproduced for research and citation purposes only, and all trademarks and product names remain the property of their respective owners.

The information is provided "as is", without warranty of any kind as to accuracy, completeness or fitness for any purpose. Underlying sources may themselves contain errors; where evidence is weak or contradictory this is indicated in the text (evidence grades, confidence labels, preserved contradictions). All estimates, models and forward-looking statements are analytical interpretations by the author, not statements of fact, and may prove incorrect.

Nothing in this report constitutes investment, legal, engineering or other professional advice. Any decisions based on this material are made solely at the reader's own risk. If you believe any statement concerning you or your organisation is inaccurate, corrections will be made promptly upon substantiated request via gailitis.ceo.