



SIEMENS

DATA CENTERS

# 100 MW Hyperscale AI Blueprint

Tier III fault-tolerant architecture for Siemens industrial-grade electrical systems, NVIDIA GB200 NVL72 systems, and nVent liquid cooling technology in UL markets.

In partnership with



# Abstract:

This technical paper outlines a reference blueprint for a 100 MW hyperscale AI data center designed to support NVIDIA GB200 NVL72-based systems like NVIDIA DGX GB200 or systems from NVIDIA partners. It presents a Tier III capable architecture that integrates Siemens’ industrial-grade electrical systems with NVIDIA GB200 NVL72 systems and nVent liquid cooling technology in the UL market.

The paper covers pod-level modularity, power distribution, cooling, control and automation, and operational scalability. It also provides definitions, terminology, schematics, and an appendix to support technical planning, design coordination, and stakeholder alignment.

# 1. Introduction

## Accelerating AI outcomes with scalable, efficient infrastructure.

In the race to unlock AI’s full potential, infrastructure must deliver more than raw power. It must deliver speed, efficiency, and resilience.

For hyperscale environments, the key performance indicators are clear: **how fast can I deploy a data center to begin generating tokens, and how many tokens can I produce per watt consumed.** This blueprint presents a reference architecture for a 100 MW AI data center purpose-built to meet those demands.

Designed around **NVIDIA GB200 NVL72 racks**, this Tier III-capable solution integrates **Siemens industrial-grade electrical systems and nVent liquid cooling to** enable rapid deployment, high compute density, and operational continuity. With modular pod-level scalability, intelligent automation, and fault-tolerant design, this architecture empowers operators to minimize time-to-compute and maximize **tokens-per-watt**, delivering predictable performance across globally distributed sites.


# Contents

|                                                       |    |
|-------------------------------------------------------|----|
| 1. Introduction                                       | 3  |
| 2. Overview: Trends, Challenges, and Design Rationale | 4  |
| 3. Power Architecture                                 | 8  |
| 4. LV Block Architecture and Pod layout               | 12 |
| 5. Cooling Architecture                               | 16 |
| 6. Automation and Control                             | 17 |
| 7. Closing Remarks                                    | 20 |
| 8. Appendix: Terminology, Definitions                 | 21 |



## 2. Overview: Trends, Challenges, and Design Rationale


To support the next generation of digital infrastructure, data centers must evolve on many angles - from power and cooling to global deployment and operational resilience. As workloads grow more intensive and distributed, infrastructure strategies must balance performance, efficiency, and scalability through increasingly intelligent, adaptive systems. The following considerations highlight key architectural shifts shaping high-density, high-availability environments.



### Increasing rack-level power densities

As AI drives growing compute demands, rack power densities are rising beyond 100 kW. This requires rethinking power distribution through high-voltage delivery, fast control response on the OT side by having a deep integration into the IT side and dynamic load balancing systems that can respond in real time to shifting workloads.

These capabilities help maintain stability and efficiency while enabling scalable, high-density deployments.



### Dense thermal loads requiring advanced liquid-cooling solutions

As workloads become more compute-intensive, traditional air cooling methods struggle to keep up. Of the various approaches, single-phase liquid cooling has emerged as the preferred solution for high-density loads. Integrated control systems continuously adapt cooling and power supply operations to workload and environmental conditions, helping optimize energy use and maintain thermal stability.



### Maintaining uptime with modular, high-density pods

Modular high-density pods integrate power, cooling, and network systems into scalable, fault-isolated units designed for rapid deployment and operational continuity.

These environments support real-time monitoring, predictive maintenance, and streamlined provisioning, helping provide high availability while maximizing tokens-per-watt and enabling consistent deployment across core, edge, and hybrid environments.



### Global deployment considerations

Geographic climate plays a significant role in energy efficiency, with hot, humid regions increasing cooling demands and colder environments enabling greater use of free cooling.

Infrastructure must adapt to these variations through intelligent, responsive systems that optimize energy use and maintain consistent performance. These capabilities support broader sustainability objectives, including emissions reduction, compliance with regional efficiency standards, and alignment with corporate ESG targets.



Picture credits: NVIDIA Corporation and nVent Management Company

In response to these design challenges, this architecture presents a scalable, high-density data center blueprint purpose-built for hyperscale AI workloads. At its core are NVIDIA GB200 NVL72 clusters, each housing 72 NVIDIA GPUs and 36 NVIDIA Grace CPUs, interconnected via NVIDIA NVLink™. Rated at **127 kW per rack**, these high-performance AI systems require advanced thermal and electrical infrastructure. This reference design supports up to **100 MW of IT load**, distributed across several modular pods engineered for **Tier III fault tolerance** in the UL market.

This **reference architecture** integrates **Siemens industrial-grade electrical and automation systems**, **nVent direct-to-chip liquid cooling**, and **NVIDIA GB200 NVL72 systems** to deliver predictable performance in globally deployable configurations.

Working in concert, these technologies enable real-time power and thermal optimization, low PUE, high operational efficiency, and can advance enterprise sustainability goals.


The following sections provide a detailed description of the physical layout, electrical and cooling topologies, and deployment models that realize the proposed architecture.

### At a glance:

| Parameter                | Specification                                              | Description                                                                                                                      |
|--------------------------|------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| Compute Density          | NVIDIA GB200 NVL72 racks at 127 kW; support racks at 14 kW | High-density AI compute with liquid-cooled GPU racks.                                                                            |
| Pod IT Load              | 2,256 kW per 32-rack (compute and network) pod             | Total IT load for one modular pod.                                                                                               |
| Total Racks              | 640 compute racks (~100 MW)                                | Full-buildout capacity across all pods.                                                                                          |
| Redundancy Level         | Tier III Fault Tolerance                                   | Provides concurrent maintainability and fault isolation.                                                                         |
| Total Pods               | 40 pods                                                    | Modular units enabling scalable capacity expansion.                                                                              |
| Cooling Strategy         | ASHRAE W32–W45 warm-water liquid cooling                   | Maximizes free-cooling hours and efficiency.                                                                                     |
| Electrical Resilience    | 4-to-make-3 redundancy, N+1/Tier III UPS                   | Provides high availability and power continuity with concurrent maintainability and fault isolation.                             |
| Future-proof Scalability | POD-level modular design                                   | Suitable for phased build-out and scale-up to next generation of higher compute racks such as NVIDIA GB300 NVL72-based clusters. |

# Reference architecture schematic

100 MW Hyperscale AI Data Center Blueprint





### 3. Power Architecture

The power system is designed for **Tier III** (Concurrently Maintainable) operation, **using a 4-to-make-3 redundancy model** to support continuous uptime during maintenance or failure scenarios. Each pod is supported by modular UPS systems and busway distribution, enabling scalable deployment and simplified maintenance.

At the rack level, **415 V AC low-voltage distribution** is used throughout, with main switchboards feeding directly into racks that support **415 V input**. Each pod includes four **415 V, 1500 kW UPS modules**, with power distributed via either overhead busways or power panels, depending on site constraints and preferences.


Utility and generator feeds are sized between **2,000–3,000 A**, supporting high-density compute and redundancy requirements. Within each pod, power is allocated as follows:

- **Eight 60 A** circuits per compute rack (to support NVIDIA GB200 NVL72 loads)
- **Two 30 A** circuits per network and support rack
- **Four 30 A** circuits per CDU

This modular, high-capacity power design provides flexibility for future scaling, while maintaining fault isolation and serviceability at the pod level.


## MW Power architecture schematic

100 MW Hyperscale AI Data Center Blueprint



# Sample Single-Line Diagram (SLD) with Main-Breaker-Busbar (MBB)

100 MW Hyperscale AI Data Center Blueprint





## 4. LV Block Architecture and Pod layout

Each pod is designed with a robust low-voltage distribution system to support high-density compute and cooling loads. The electrical distribution architecture includes **four 415 V, 3,000 A main switchboards** per pod. Power is delivered to IT and support racks via **overhead busways or power panels**. Busplugs are spaced on **2-foot centers**, available on one or both sides of the busway, and support **1 to 4 breakers per plug**.

Each network and support rack typically receives **two 30 A circuits**, while **CDUs are fed by four 30A circuits per unit**.

The pod layout incorporates **N+1 power redundancy**, with **8 power shelves per NVIDIA GB200 NVL72 system** to support the 127-kW load. Network and support racks are provisioned with **2 power feeds**, while **CDUs are configured with 4 feeds** to support continuous cooling operation.

This modular LV block design enables consistent power delivery, simplifies installation, and supports scalable deployment across globally distributed sites. This electrical architecture can be scaled to accommodate next-generation racks with higher density, including the NVIDIA GB300 NVL72.

### Key Electrical Specs per Pod:

| Component         | Specification                       | Description                                           |
|-------------------|-------------------------------------|-------------------------------------------------------|
| Main Switchboards | 4 × 415 V, 3,000 A                  | Primary power input and distribution for the pod.     |
| Busway            | 4 overhead busway runs per row      | Overhead power distribution routes.                   |
| Busplugs          | 1–4 breakers per plug, 2 ft spacing | Tap-offs providing modular power connections.         |
| Network Racks     | 2 × 30 AT breakers                  | Dual-feed power supply for network and support racks. |
| NVL72 Racks       | 8 power shelves (N+1)               | Redundant power shelves per NVIDIA rack.              |
| CDUs              | 4 feeds                             | Redundant power feeds for cooling units.              |




# 5. Cooling Architecture

Each rack receives dedicated liquid cooling, with modular CDUs scaled to meet the IT liquid heat load of high-density AI systems. Cooling distribution is implemented at the row or hall level, providing dual-path (A/B) redundancy to ensure continuous operation during maintenance or fault conditions.

At the pod level, the cooling infrastructure is designed to support up to 1.63 MW of liquid IT heat load in an NVIDIA GB200 NVL72 system configuration down to a 4K approach temperature differential (ATD) with three CDU800 CDUs. Each CDU800 uses two fully independent external power feeds, one per pump, eliminating the single input breaker that represents an inherent availability risk in traditional single-feed CDUs. Because each CDU800 delivers its full 544 kW rated capacity for the NVIDIA GB200 NVL72 system with only one active pump, the loss of a pump, drive, or power feed does not derate cooling performance. The remaining independently powered pump remains fully operational,

uniquely decoupling electrical availability from thermal capacity without requiring N+1 CDU oversizing. Operating within **ASHRAE W32–W45** warm-water classes, the system enables chiller-free operation for most of the year in temperate climates, reducing reliance on mechanical cooling, improving overall energy efficiency, and lowering total cost of ownership. This inherently redundant and service-efficient design simplifies deployment and minimizes maintenance touchpoints at scale.


The same liquid cooling architecture seamlessly scales to future platforms such as **NVIDIA GB300 NVL72 based clusters**, utilizing the same **three CDU800s** deployed in the GB200 NVL72 configuration. When additional capacity or redundancy is desired, a **fourth CDU800** can be integrated without modification to the existing liquid cooling infrastructure, preserving design consistency while providing flexible expansion headroom for next-generation systems.



nVent Dual-Feed CDU800 Power Topology

# Cooling architecture schematic

100 MW Hyperscale AI Data Center Blueprint





## 6. Automation and Control

The **automation and control architecture**, implemented within the **Integrated Data Center Management Suite (IDCMS)**, provides a flexible and scalable foundation for subsystem integration and interoperability through open communication protocols. The IDCMS serves as a unified platform for **Building Management** and **Energy and Power Management**, featuring deep integration with **NVIDIA Mission Control** for comprehensive visibility and control across compute, network, and infrastructure domains.

This integration establishes a **high-availability, low-latency**, and **unified monitoring environment** that seamlessly bridges **operational technology (OT)** and **information technology (IT)** layers. It can provide safe, efficient, and optimized operation of the **NVIDIA GB200 NVL72-based data center**.

The automation layer employs a **PLC-based, hot-standby N+1 redundant architecture** with distributed remote I/Os, enabling resilient control of HVAC systems for air-based cooling across the PODs. It integrates with **nVent Coolant Distribution Units (CDUs)** to manage **direct-to-chip liquid cooling** for the compute trays. Intelligent motorized valves and multi-level leak detection sensors provide precise coolant flow control and immediate response to fluid anomalies, capable of providing system integrity and continuous operation.

Each liquid-cooled rack is equipped with a dedicated **Rack Telemetry & Control Gateway** that continuously monitors local sensors and governs the operation of rack-level control and modulation valves. This intelligent gateway serves as the digital bridge between the **CDU control system** and the **IDCMS**, enabling coordinated optimization of liquid flow, temperature, and pressure across the entire cooling network.


In the event of a leak or anomaly at the rack, the gateway instantly communicates with both the IDCMS and CDU, triggering automatic rack-level isolation to contain the issue while preserving cooling continuity for the remaining IT equipment. This integrated control strategy enhances performance, safeguards uptime, and provides the responsive protection necessary for hyperscale and AI data center environments.

On the electrical side, IDCMS interfaces with **advanced power metering and quality monitoring systems** to deliver real-time insights into energy use and power stability. This enables early detection of potential disturbances that could affect compute or network performance, supporting **dynamic load management** and **automated corrective actions** under varying demand conditions. Additionally, integration with protection devices allows for **remote control and event-driven response**, strengthening overall reliability and operational safety.



# Automation architecture schematic

100 MW Hyperscale AI Data Center Blueprint



## 7. Closing Remarks

From blueprint to deployment: enabling AI at scale

This reference architecture is designed to help operators overcome some of the most pressing challenges in AI infrastructure: **power availability and energy efficiency.**

As compute demands surge, the ability to deploy quickly and operate efficiently becomes a competitive advantage. This blueprint addresses both by offering a modular, fault-tolerant design that minimizes downtime and accelerates time-to-compute.

With high-density pods, intelligent automation, and integrated liquid cooling, operators can maximize **tokens-per-watt** and deliver more AI output per unit of energy consumed.

This is especially critical in environments where power is constrained or where operational efficiency directly impacts throughput.

**Next Steps:** Use this blueprint as a foundation to plan, scale, and optimize your AI data center deployments. Whether you're building new capacity or upgrading existing infrastructure, this design offers a clear path to high-performance, energy-aware operations.



## Appendix: Terminology, Definitions

**Busway & Busplugs:** Overhead electrical distribution system delivering power from switchboards/UPS to racks, with modular connection points and breakers.

**CDU** (Cooling Distribution Unit): Liquid-cooling unit distributing chilled liquid to racks and returning heat to facility heat rejection system.

**NVIDIA GB200 NVL72:** NVIDIA rack-scale system with 72 GPUs, designed for ~120 kW per rack, using direct-to-chip liquid cooling.

**Pod:** A self-contained modular unit of ~16 compute racks, support racks, and CDUs with integrated electrical and cooling paths.

**PUE** (Power Usage Effectiveness): Ratio of total facility energy to IT energy; lower values indicate higher energy efficiency.

**IT Racks:** refer to the combined group of compute and network racks within a data center.

**Tier-level system:** The Uptime Institute Tier Standard classifies data center infrastructure based on levels of redundancy, maintainability, and fault tolerance. This reference architecture supports both Tier III and Tier IV configurations, depending on site requirements and workload criticality. The four tiers are:

### Tier I – Basic Capacity

Single path for power and cooling, with no redundancy; any maintenance or failure results in downtime.

### Tier II – Redundant Components

Adds some redundancy (e.g., backup UPS or cooling), but still relies on a single distribution path.

### Tier III – Concurrently Maintainable

Multiple distribution paths and redundant components allow any single system to be taken offline for maintenance without impacting IT operations.

### Tier IV – Fault Tolerant

Fully redundant, dual-active power and cooling paths designed to support uninterrupted operation even during unplanned failures.

### Let's talk!

Fill in this contact form to start a conversation with our data center experts.



### Visit our website

Discover the Siemens data center offering and use cases that align with the strategic business goals.



### Visit use case page

Discover more information and offerings about how to manage infrastructure for AI workloads.



#### **Published by Siemens Switzerland Ltd**

Smart Infrastructure  
Global Headquarters  
Theilerstrasse 1a  
6300 Zug  
Switzerland  
Tel +41 58 724 24 24

#### **For the U.S. published by Siemens Industry Inc.**

3617 Parkway Lane  
Peactree Corners, GA 30092  
United States

Article-No: SI-FlyerRefArch100MWAI-WWEN-2025-11  
Status: 11/2025

Subject to changes and errors. The information given in this document only contains general descriptions and/or performance features which may not always specifically reflect those described, or which may undergo modification in the course of further development of the products. The requested performance features are binding only when they are expressly agreed upon in the concluded contract.