


SYSTEM-LEVEL MODULAR INFRASTRUCTURE FOR AI FACTORIES

AXIOM REFERENCE DESIGN FOR NVIDIA VERA RUBIN NVL72



Platform	NVIDIA Vera Rubin NVL72
Reference number	AXIOM-RD-VR-NVL72-1
Deployment unit	13.44 MW compute load across four 3.36 MW AI Modules
Document type	Technical reference design for engineering and proposal use
Classification	External release, planning-level basis of design

Planning basis and validation note

This VR#1 reference design is based on 13.44 MW of IT load across four 3.36 MW AI Modules using nominal 200 kW GPU racks and 35 kW support racks. Values shown are planning-level assumptions and require validation against final OEM rack power, heat capture ratio, National Electrical Code (NEC) continuous-load application, feeder and switchgear ratings, short-circuit and coordination studies, computational fluid dynamics (CFD), hydraulic modelling, transient TES analysis, controls sequencing and site-specific ambient conditions.

Power Pod-level 4-to-make-3 redundancy is not confirmed unless the final capacity model demonstrates that any three Power Pods can support the required operating load with the required design margin.

About This Document

This reference design defines a factory-integrated modular AI data hall organized into repeatable deployment blocks for compute, support, power, hybrid cooling, controls and lifecycle services. It is written for engineering, design and proposal teams who require a common technical baseline for high-density AI data halls based on the NVIDIA Vera Rubin NVL72.

The design is a standard electrical and modular infrastructure baseline that supports customer engagement, proposal development, engineering standardization and roadmap decisions. Values presented are planning-level assumptions intended to guide detailed design. Final ratings and performance require validation against OEM equipment data, applicable codes and standards, and site-specific conditions.

What this reference design establishes

- A documented planning baseline across electrical, mechanical, controls, and services, as the common starting point for engineering, proposal development, and site adaptation
- Capacity aligned to modular blocks, consumed in whole increments rather than stranded
- Power and cooling engineered as one system through shared controls and sequencing
- Assembly, wiring, and test scope completed and documented in the factory before shipment, reducing onsite variability and commissioning duration
- Standardized blocks, interfaces, and acceptance criteria, allowing deployment across multiple locations with predictable outcomes
- A defined path as density rises, without redesigning the facility

Introduction

The Axiom Solutions AI Factory reference design (AXIOM-RD-VR-NVL72-1) describes a repeatable, factory-integrated modular AI data hall. Infrastructure is organized into standard deployment blocks so that compute, support, power, cooling, controls and services scale together as coordinated units rather than as separately engineered systems.

The NVIDIA-based basis of design supports 13.44 MW of AI load across 56 GPU racks at a nominal 200 kW and 64 support racks at 35 kW. At a planning PUE of approximately 1.27, total site load is estimated at approximately 17.1 MW, including 2.8 MW for mechanical loads.

Compute load is arranged in four repeatable AI Modules. Each module combines one compute GPU module and one support module and delivers 3.36 MW of AI load as a standard unit for scale, manufacturing and deployment planning. Rack power distribution uses a 4-to-make-3 configuration at the rack or feeder level through busway and tap-off units. Four Power Pods or skids, each approximately 4.9 MW, provide upstream capacity using a 5 MVA 34.5 kV-to-480 V transformer, LV switchgear, feeder panels, busway interfaces and controls.

The architecture is designed to be scalable for GPU racks with higher power up to 600 kW and beyond while reducing rack count, and is expected to provide a defined migration path toward next-generation 800 VDC power distribution.



Accelerated deployment

Factory-integrated power, cooling and controls reduce onsite construction and commissioning delays, enabling faster time to capacity.



Repeatable, predictable execution

Factory-built and tested solutions improve quality, deployment consistency and schedule certainty.



Scalable path to 800 VDC

Future-ready architecture supports increasing AI power demands while providing a migration path to next-generation power distribution.



Optimized power utilization

Modular capacity blocks reduce stranded power while improving scalability and infrastructure efficiency.



Flexible cooling for evolving AI

Integrated liquid and air cooling with coordinated controls support mixed rack densities and changing compute requirements.



Resilient, AI-ready

Modular power distribution, coordinated controls and NVIDIA DSX™ compatible designs balance reliability, efficiency and operational flexibility.

Data Hall Layout and AI Module Architecture

AI Data Hall layout

Each AI Module aligns to a dedicated infrastructure zone for power, cooling, controls and service access.


- **Power block:** modular prefabricated low-voltage power distribution sized around the AI Module load profile.
- **Cooling block:** liquid cooling and air-cooling capacity coordinated to support mixed rack densities.
- **Compute block:** AI Module is the white-space IT building block, combining a GPU module and a support module into one coordinated unit.
- **Controls zone:** module-level monitoring, controls and alarm management.
- **Service boundary:** clear maintenance access, CDU service zones, busway access and rack replacement aisles.

Discrete capacity blocks reduce stranded capacity and simplify phased deployment, operating ownership and future expansion.

Per AI Module Configuration

Each AI Data Hall is configured as follows:


Parameter	Per AI Module	Data Hall (×4)
GPU racks (200 kW nominal)	14	56
Support racks (35 kW nominal)	16	64
GPU IT load	2.80 MW	11.20 MW
Support IT load	0.56 MW	2.24 MW
Total IT load	3.36 MW	13.44 MW



Data Hall Planning Basis

The table below summarizes planning basis at a glance. Site-level power is derived from the IT load and planning PUE, and is provided for capacity planning.

Parameter	Specification
Platform	NVIDIA Vera Rubin NVL72
Reference number	AXIOM-RD-VR-NVL72-1
Total IT load (deployment unit)	13.44 MW
AI Module	4 modules at 3.36 MW each
GPU racks	56 racks at 200 kW nominal (11.20 MW)
Support racks	64 racks at 35 kW nominal (2.24 MW)
Planning PUE	1.27
Estimated total site load	17.1 MW
Total Mechanical load	2.8 MW
MV intake	34.5 kV medium voltage
Distribution voltage	480 V
Power Pods / Skids	4 blocks at ~4.9 MW (5 MVA transformer each)
Rack power distribution	Busway and tap-off units, 4-to-make-3
Cooling topology	Hybrid direct-to-chip liquid cooling and CRAH air cooling
Mechanical redundancy	N+1 (pump skids, chillers, CRAHs, CDUs, fluid coolers)



Electrical Power Architecture

The electrical architecture uses a 34.5 kV medium-voltage intake feeding modular 34.5 kV to 480 V transformer-based Power Pods or Skids. Each Power Pod or Skid is a repeatable electrical capacity block built around a 5 MVA transformer, supporting approximately 4.9 MW of usable compute capacity at 0.99 power factor, and includes a transformer interface, LV switchgear, busway connections, tap-off units, metering, and controls integration. Four Power Pods provide sufficient installed capacity for the 13.44 MW facility IT load under normal operation, made up of four AI Modules at 3.36 MW each, comprising 2.80 MW GPU load and 0.56 MW support load per pod. Power Pod-level 4-to-make-3 redundancy requires additional Power Pod capacity or a defined redundancy strategy.

LV Distribution

Each modular Power Pod (outdoor deployment) or Power Skid (indoor deployment) distributes 480 V power through LV switchgear to PDU/RPP panels and overhead busway. Busway and tap-off units provide rack-level delivery, allowing rack clusters to be added, reconfigured, or isolated without redesigning upstream MV and LV infrastructure.

Note: Validate the LV architecture through power flow analysis, short-circuit analysis, protective device coordination studies, concurrent maintainability review, and arc flash hazard analysis under high-density AI load profiles.

Rack-Level Redundancy

Rack-level redundancy should use a 4-to-make-3 distribution model with multiple independent rack power paths. The design must state the level at which 4-to-make-3 applies

- Rack level: four rack-side power paths, with any three capable of supporting the required IT load.
- Feeder level: four 1600 A feeders, with any three capable of supporting the connected workload.
- Power Pod level: four Power Pods or Skids, with any three capable of supporting the full facility IT load. and Safety

Protection and Safety

The electrical design should cover short-circuit performance, selective coordination, arc flash, breaker settings, grounding, fault isolation, and rack-level isolation logic, with leak detection and CDU fault signals interlocking with isolation devices where applicable. The system should include selective coordination, short-circuit study, arc flash analysis, ground fault protection, fault isolation, breaker settings validation, leak detection and CDU fault interlocks, and rack-level or row-level isolation where applicable.


Monitoring and Controls

EPMS should be a core layer of the modular electrical architecture. Metering from transformer through rack distribution supports capacity planning, power quality monitoring, predictive maintenance, load balancing, and rapid fault isolation. The reference design integrates EPMS, branch-level metering, power quality monitoring, alarm management, event capture, and predictive maintenance analytics across the full distribution path, covering MV incoming, transformers, main switchgear, feeder breakers, RPP/PDU Panels, busway, tap-off units, rack-level circuits, CDU electrical feeds, and leak detection and environmental monitoring systems.


Design Options

- **UPS and generator:** optional, by criticality and site backup strategy
- **Redundancy:** 4-to-make-3, N+1 or 2N, by tier target
- **Battery chemistry and runtime:** by footprint and TCO
- Standards and environmental rating configured to region future expansion.

Facility Power Block Diagram



Parameter	Specification
Transformer voltage	34.5 kV / 480Y/277 V
Transformer power	5 MVA
Input voltage (IT main bus)	480 V
Main switchboard nominal current	6,000 A
Main switchboard interrupting current capacity	85 kA
Main switchboard configuration	3PH, 4W, UL891
UPS distribution switchboard (UDS) current	6,000 A
Feeder switchgear current	1,600 A
Number of feeder panels	4
Power factor	0.99
Power available per Compute module	4.9 MW
Power redundancy (IT power path)	4-to-make-3
UPS capacity per unit	2,000 kW / 2,000 kVA
UPS output voltage	480 V
Battery cabinet rating	600 A / 480 VDC
Generator rating	5 MW



Mechanical Plant and Hybrid Cooling Architecture

The cooling architecture supports high-density AI and HPC workloads through hybrid cooling, combining direct-to-chip liquid cooling for high-density racks with CRAH-based air cooling for residual whitespace loads. The reference design is based on a 13.44 MW IT deployment unit with nominal 200 kW GPU racks and scalability to higher densities, subject to OEM equipment, heat capture ratio, site conditions, and operating envelope validation. The architecture comprises rack-level Secondary Fluid Networks, liquid-to-liquid CDUs, a high-temperature dry cooler or fluid cooler loop, a chilled-water loop serving CRAHs and supplemental needs, and Thermal Energy Storage for short-duration ride-through, with N+1 redundancy applied to pump skids, chillers, CRAHs, CDUs, and fluid coolers.

Liquid Cooling and Heat Rejection

The primary liquid loop serves direct-to-chip cooling for AI and network racks. Heat transfers from rack-level SFNs to shared CDUs, then through liquid-to-liquid exchangers to the facility dry cooler or fluid cooler loop. Under normal conditions the dry cooler loop is the primary heat rejection path, operating at 104°F supply and 122°F return with an 18°F ΔT . This high-temperature approach is intended to maximize economizer hours and reduce chiller runtime, with final performance requiring climate-based energy modelling and OEM validation. IT coolant supply is assumed to remain below approximately 45°C / 113°F at a ΔT of approximately 10°C, an assumption that directly determines flow rate, pipe sizing, pump head, CDU capacity, and TES volume. The dry cooler loop interfaces with TES to buffer thermal load during utility interruptions, generator startup, cooling source transitions, and short-duration load changes.

Chiller and CRAH Loops

The chiller loop is configured as N+1 and performs two functions, supplying chilled water to CRAHs and supplementing or backing up the high-temperature liquid loop when dry cooler capacity is limited. Reference temperatures are 68°F supply and 82.4°F return with a 14.4°F ΔT . CRAHs manage residual air-cooled heat within the AI Module, including non-liquid-cooled IT equipment, network gear, and room-level thermal loads, receiving chilled water at approximately 68°F / 20°C. The CRAH loop should include air separation, expansion control, isolation valves, balancing valves, and instrumentation for stable operation and maintainability. Final air temperatures, airflow rates, coil approach, humidity strategy, and fan power must be validated through thermal modelling or CFD.

Shared CDU Loop

The shared CDU loop interfaces between facility cooling and rack-level liquid distribution, with redundant pumps providing N+1 availability and maintaining coolant flow to liquid-cooled IT equipment. The system should include duty and standby pumps, air separation, expansion control, filtration as required, fluid quality monitoring, differential pressure control, temperature monitoring, leak detection and isolation coordination, and integration with BMS, EPMS, and cooling controls.

Thermal Energy Storage and Maintainability

The design includes high-temperature side TES with an estimated 31,397 gallons, providing approximately 6 minutes of cooling ride-through during utility loss or cooling source transition. TES maintains cooling continuity while generators start, stabilize, and restore power to pumps, chillers, dry coolers, and controls, with UPS-backed CDU and distribution pumps continuing to circulate coolant through critical paths during the transition. N+1 redundancy across major cooling components is supported by service access, isolation valves, bypasses, and removable components for concurrent maintainability. Detailed design must confirm that piping, valves, controls, heat exchangers, and electrical dependencies do not create single points of failure.


Note: Validate rack density and heat capture ratio, OEM coolant limits, TES volume and ride-through, dry cooler economizer hours and the 82°F switch point, controls sequencing across economizer, chiller-assist and recovery modes, hydraulic stability, CRAH airflow and capacity, concurrent maintainability, fluid quality and materials compatibility, and controls integration across CDU, BMS, EPMS, TES, pumps, chillers, and dry coolers.

Design Options


- Heat rejection selected as dry cooler or fluid cooler, based on climate
- Chiller configured in a supplemental or dominant role, based on ambient profile
- Loop temperatures and delta T set as per OEM limits and site conditions
- TES volume and ride-through sized to the restart profile
- CDU capacity and count selected against rack flow and density

Parameter	Specification
Total IT load (compute + support module)	13.44 MW
Voltage	480 V
Heat rejection medium	PG and Air
Pump skids	N+1
Chiller	N+1
CRAH	N+1
CDU	N+1
Chiller supply / return temperature	68 / 82.4 °F (14.4 °F ΔT)
Dry / fluid cooler supply / return temperature	104 / 122 °F (18 °F ΔT)
IT coolant supply (assumed, OEM validation)	< 45 °C / 113 °F (~10 °C ΔT)
TES storage capacity (high-temp side)	31,397 gal
Ride-through time	6 min
Dry bulb switch point (dry cooler to chiller-assist)	82 °F

Block diagram of hybrid cooling architecture



P&ID of facility cooling schematics




Controls and Monitoring

Controls and monitoring are treated as a core layer of the modular architecture, integrating BMS, EPMS, DCIM, CDU controls, leak detection, power quality monitoring, and environmental alarms into one operating framework from utility source to rack. Electrical telemetry (switchgear, feeders, RPP/PDU panels, busway, tap-offs, rack feeds) and mechanical telemetry (cooling systems, CDUs, leak detection, environmental sensors, fluid loops) are unified so operating teams can monitor system health, manage capacity, isolate faults, and tune performance by modular block.

Telemetry is organized by operating layer and coordinated power-and-cooling control sequences track rack and AI Module demand, CDU performance, temperature and flow trends, and available electrical capacity as GPU loads change. Events are captured, time-stamped, alarmed, and tied to defined response actions, including leak-detection interlocks, CDU fault response, and localized rack or row isolation, for faster fault isolation, safer maintenance, and higher availability. The control logic is designed for the fast-changing load profiles of AI training and inference, and supports predictive maintenance, capacity tracking, load balancing, and trend-based optimization.

Telemetry by Operating Layer

Layer	Monitored parameters
Rack	Power feed status, branch circuit loading, leak detection, temperature, rack-level alarms
Pod	AI Module load, CDU performance, support rack load, busway and tap-off status
Plant	Cooling plant status, pump operation, water temperatures, thermal storage
Power	Transformer loading, switchgear status, feeder load, power quality, event capture, breaker position, alarms




Factory Integration, Testing and Commissioning

Factory integration is applied where practical to Power Pods, skids, switchgear, busway, tap-off units, controls panels, CDU skids, and monitoring interfaces. Modular blocks are assembled, wired, labelled, tested, and documented in a controlled environment, shifting defined integration work from the construction site to the factory to reduce onsite variability, commissioning risk, and deployment duration, and to support repeatable global deployment. Each Power Pod or skid ships as a repeatable electrical block (transformer interface, LV switchgear, feeder distribution, UPS and batteries where applicable, controls, metering, busway). Controls panels are pre-wired and factory-tested for network connectivity, wiring, device addressing, alarm points, metering, CDU interface points, leak detection, and EPMS/BMS readiness. Busway and tap-off units are standardized by pod type and rack density, with defined normal and degraded operating modes.

Rack Integration and Validation

Axiom extends the architecture into the rack, integrating compute, rack-level power, direct-to-chip liquid cooling, controls, imaging, security hardening, and validation before shipment. As rack densities rise toward 250 kW, 500 kW, and 1 MW, thermal integrity becomes a first-order reliability requirement.


Liquid-cooled racks complete a defined coolant qualification before functional verification of cooling paths, sensors, alarms, and CDU interfaces:



Cluster-level FAT can then run against the customer’s actual rack configuration, covering layout verification, power and load verification, leak and alarm validation, simulated AI workload testing, and documented results with exceptions and corrective actions.

L6-L12 Validation Path

The L6–L12 validation model provides a structured path from board-level manufacturing (L6–L8) through rack integration and functional test (L9–L10) to cluster-level acceptance, each stage building on the prior gate. Racks verify the factory deployment-ready, aligned to the site architecture, with full power, cooling, controls, network, and documentation packages (rack configuration, leak and pressure-test records, network and interconnect test records, FAT reports and exception logs, and startup and commissioning support).



Test and Commissioning playbook

Level	Service	What is included
1	Preliminaries, factory testing and delivery	Project management, health and safety planning, factory witness testing, delivery coordination, site logistics
2	Interconnection and static testing	Positioning, interconnection, torque verification, insulation resistance, continuity, dead testing, labelling verification
3	On-site commissioning	Energization, functional testing of breakers/isolators/protection, control wiring verification, I/O point-to-point, cause and effect, UPS and cooling checks
4	Site acceptance testing (SAT)	SCADA/HMI verification, live data validation, alarm and event testing, metering accuracy, fire detection and suppression, record sign-off
5	Integrated systems testing (IST)	Cross-system integration, load shedding and sequencing, network redundancy and failover, scenario simulation, thermal imaging post energization
6	Final acceptance and client handover	Performance and reliability testing, final start-up checks, room integrity, client training, documentation handover, formal sign-off

Project and Lifecycle Services

Axiom delivers four connected service families spanning initial deployment through long-term operation and team enablement. The model combines deep in-house engineering expertise with authorized partners, providing a single, accountable service interface for complex, live operational environments, from first energization through full operational life.



Project Services

- Commissioning
- Installation
- Integration
- Acceptance testing



Lifecycle Services

- Preventative maintenance
- Emergency response
- Extended warranty
- Critical spares



Advanced Services

- Power management systems (PMS)
- SCADA and EMS
- Engineering studies
- Environmental monitoring



Training Services

- Client product training
- Partner certification
- Operations and maintenance documentation
- System handover and familiarisation

Lifecycle Services

Beyond factory build and commissioning, Axiom lifecycle services carry the same infrastructure across five phases, reducing technical risk before deployment, improving quality through factory integration, supporting reliable operation from source to rack, and preserving an upgrade path as densities, cooling, and power architectures evolve.

Plan - validate the architecture before procurement Basis of design, load and rack-density validation, electrical studies, CFD and airflow review, liquid cooling assessment, constructability review, controls architecture definition, commissioning planning.

Build - convert the design into field-ready modular infrastructure Factory integration of Power Pods/Skids, switchgear, busway, CDU skids, RPP/PDU and controls panels, with FAT of electrical, mechanical, controls, metering, labelling, communications, and documentation readiness.

Operate - support reliable daily operation Remote monitoring, preventive and condition-based maintenance, alarm response, event review, spare-parts planning, field service, drawing on integrated EPMS, BMS, DCIM, CDU controls, leak detection, environmental, and power quality data.

Optimize - improve performance using operating data Power utilization, cooling performance, rack load balancing, PUE, alarm thresholds, controls sequences, and maintenance strategy tuned at rack, AI Module, power block, cooling block, and facility levels.

Operate - preserve an upgrade path CDU expansion, busway augmentation, RPP/PDU upgrades, Power Pod capacity expansion, controls and monitoring upgrades, rack-level isolation improvements, and future 800 VDC evaluation.

