


AXIOM AI MODULE

Factory integrated modular AI infrastructure
of up to 4MW for repeatable AI factories

The Axiom Solutions™ AI module is a prefabricated, factory-integrated AI compute module that combines high-density rack infrastructure, liquid cooling, power distribution, networking, controls, and containment into a repeatable deployment block for AI factories and next-generation data centers. Each module is manufactured, integrated, and tested before shipment, enabling rapid deployment and predictable performance at scale.

Markets

// Hyperscale // Enterprise AI // Colocation
// Neo cloud // Government



Integrated by design.

Power, direct-to-chip liquid cooling, high-density racks, networking, controls, and containment built and tested as one system



AI-ready density.

engineered for current accelerators, with up to 1 MW per rack, for next-generation AI platforms



A repeatable building block.

Standardized modules scale from a single module to a complete AI factory



Validated before shipment.

Manufactured, integrated, and validated in the factory for predictable on-site performance




Images shown are for illustration only. Final product configuration and appearance may vary.

SCALABILITY

Building blocks of the AI factory


Scaling AI infrastructure should not mean redesigning the data center with every deployment. The AI Module is the repeatable unit that makes growth predictable.

Each module is delivered as one engineered system with a known power capacity, a defined cooling envelope, and a consistent interface to the rest of the architecture. Capacity is added by replicating validated blocks rather than re-engineering from scratch, so expansion is faster, simpler, and lower risk. As demand grows, modules combine into rows, rows into data halls, and halls into a full AI factory, all built on the same proven module.




AI Module

One module up to 4 MW




Data Hall

Modules combined tens of MW



AI Factory

100 MW and beyond



Modular infrastructure for the AI factory

AI infrastructure demand is accelerating faster than conventional construction can deliver, and long build timelines, grid constraints, and skilled-labor shortages widen the gap between when capacity is needed and when it comes online.

Axiom closes that gap with a modular approach. The AI factory is assembled from standardized, factory-built modules that arrive integrated and tested, so site preparation and module production advance in parallel and capacity is delivered with speed, consistency, and confidence.

Power Pod/Skid

Modular prefabricated medium and low voltage power distribution sized as per AI module load profile

Cooling/ mechanical Pod

Liquid cooling and air-cooling capacity coordinated to support mixed rack densities


AI Modules

High-density white space with racks, containment, SFNs, networking, CDUs and controls, integrated and validated as a block.



Inside the AI Module

Four integrated subsystems, engineered together and configured to the workload.



Cooling

Direct-to-chip liquid architecture, with centralized Coolant Distribution Units and a looped Secondary Fluid Network integrated with controls to monitor safety and performance of the system. Hybrid air and liquid operation is supported where workloads require it.

Power

Dedicated feeds and overhead busway with rack-level tap-offs, arranged in the four-to-make-three configuration aligned with leading AI accelerator reference architectures.

Structure

A modular bay architecture carries the racks and containment, with upper static and lower hot-aisle containment and self-returning doors for controlled airflow and safe service access. The frame is engineered for shipment and fast, simple assembly.

Controls/IT

Integrated instrumentation across power, cooling, and IT interfaces, with real-time monitoring of thermal, electrical, and flow conditions at module and rack level, for fault detection and lifecycle visibility.

Technical specifications

Parameter	Specification
Capacity and Configuration	
Module type	Prefabricated, factory-integrated AI compute module
Total IT load per module	Up to 4 MW
Configuration	Combined high-density GPU compute and supporting network / storage racks
GPU racks	High density GPU racks, configurable
Support racks	Configurable network, storage, and management
Rack count	Configurable per deployment (typically ~30 racks per module)
Density roadmap	Scalable from 250 kW to 1MW per rack
Compute basis	Aligned to leading AI accelerator platforms
Power	
Rack distribution voltage	480 VAC/ 800 VDC/ ±400 VDC , configurable
Distribution topology	4-to-make-3 redundancy via busway and tap-off units
Feeders	Multiple rated feeders in an N+1 arrangement
Efficiency	Depending on climate and cooling configuration
Monitoring (EPMS)	Branch metering, power quality, alarms, source-to-rack
Cooling	
Architecture	Hybrid: direct-to-chip liquid for high-density racks + air for residual loads
Liquid coolant supply temp	Warm-water operation, typically below ~45°C / 113°F
Heat rejection	Fluid coolers / dry coolers with chiller-assist for peak conditions
Thermal ride-through	Thermal Energy Storage buffers utility loss and source transition
Rack-level cooling	Secondary Fluid Networks (SFN) to shared liquid-to-liquid CDUs
Redundancy	N+1 on pump skids, chillers, CRAHs, CDUs, and fluid coolers
Cooling medium	Water / glycol liquid loops and air
Controls	
Controls integration	Unified BMS, EPMS, DCIM, CDU controls, leak detection, environmental monitoring
Validation	
Validation model	Structured L6–L12 gates: component and board through rack integration to cluster-level FAT
Liquid-loop qualification	Fill, dry, leak test, and nitrogen prep before shipment
Deployment model	Factory-built and tested; site prep and module production run in parallel

Lifecycle support

Axiom provides end-to-end services that extend beyond equipment supply. With full control of the manufacturing process and deep in-house expertise, we help customers plan, deploy, operate, and scale modular power solutions with confidence. Our services are designed to reduce risk, shorten timelines, and help ensure long-term performance across the full asset lifecycle.

■ Cooling

Direct-to-chip liquid architecture, with centralized Coolant Distribution Units and a looped Secondary Fluid Network integrated with controls to monitor safety and performance of the system. Hybrid air and liquid operation is supported where workloads require it.

■ Accelerated project execution

Early engineering engagement allows design risks and project-specific requirements to be resolved up front, reducing rework and shortening lead times.

■ Flexible commercial and warranty options

Commercial terms are tailored to project needs, with optional warranty extensions available for up to five additional years. Preferential manufacturing arrangements support future expansion or replacement with minimal disruption.

■ Comprehensive testing and servicing

Factory and field services with full electrical verification include segregation, earthing, insulation, metering, protection, and coordination, as well as flash, insulation, and primary injection testing.

■ End-to-end delivery assurance

Axiom maintains full control across the manufacturing lifecycle, enabling accurate forecasting, responsive capacity planning, and reliable delivery.

■ Deep systems expertise

Our teams routinely deliver some of the largest and most complex modular power systems globally, providing integration expertise that extends beyond individual components.

■ Global service coverage

24/7 engineering support is backed by continuous manufacturing operations and component stock. Bespoke service contracts range from start-up support to full 24/7 call-out coverage worldwide.



The Axiom AI Module is part of Flex's Cloud and Power Infrastructure (CPI) segment, which Flex plans to separate into Axiom, an independent publicly traded company, in the first quarter of calendar 2027.

axiomsolutions.com