

AXIOM CORE GRID(TM)

EDGEVAULT(TM)

MODULAR AI INFRASTRUCTURE PLATFORM

Enterprise-grade, rapidly deployable infrastructure for high-density AI, private cloud, edge computing, and mission-critical workloads.

DEPLOY IN WEEKS, NOT YEARS

DATA + SPECIFICATION BRIEF

Capacity | Power Profiles | Capabilities | Client Applications | Scalability

Platform Overview

EdgeVault(TM) is a modular AI infrastructure platform engineered for organizations that need secure, scalable, high-density computing without the schedule and capital burden of a traditional ground-up data center. The platform can support artificial intelligence, private cloud, edge computing, high-performance computing, storage, networking, and mission-critical applications in a compact, repeatable deployment model.

Core Value Proposition

Rapid Deployment Factory-coordinated modular architecture designed to compress deployment schedules.	High-Density Compute Supports concentrated GPU, CPU, storage, and network infrastructure.
Enterprise Resiliency Designed for controlled power, cooling, security, monitoring, and serviceability.	Scalable Growth Deploy one pod initially, then expand into a multi-pod AI campus.

Ideal Client Markets

Client Segment	Typical Need
Artificial Intelligence Companies	Model training, inference, and GPU compute clusters
Enterprise Organizations	Private cloud, virtualization, ERP, and mission-critical applications
Cloud & Managed Service Providers	GPU-as-a-Service, colocation, private hosting, and regional capacity
Government & Defense	Secure and rapidly deployable computing environments
Healthcare Systems	Medical imaging, AI diagnostics, electronic records, and research
Hospitality & Commercial Real Estate	Property-wide IT, surveillance, resilience, and business continuity
Manufacturing & Industrial	Automation, digital twins, analytics, and edge processing
Telecommunications Providers	Edge compute, points of presence, and network infrastructure
Universities & Research Institutions	HPC, simulation, modeling, and scientific research
Financial Services	Low-latency processing, private infrastructure, and disaster recovery

GPU Capacity & Power Profiles

EdgeVault(TM) can be configured around the target workload, selected GPU platform, rack density, electrical service, cooling architecture, redundancy target, and site conditions. The following profiles are planning-level examples for a configuration supporting up to 64 GPUs.

Configuration	Approx. GPU Count	Total Power Draw	Infrastructure Rating	Typical Applications
Mid-Range AI / RTX-Class	Up to 64	22.4 kW	0.022 MW	AI inference, rendering, virtualization, engineering workloads
Enterprise AI / NVIDIA H100-Class	Up to 64	44.8 kW	0.045 MW	Large language models, AI training, HPC, enterprise AI
Extreme Density / Next-Generation	Up to 64	57.6+ kW	0.058+ MW	Frontier AI research, advanced simulation, hyperscale inference, scientific computing

Planning Assumptions

Power figures above represent estimated GPU load only and do not include full-facility overhead such as CPUs, storage, networking, fans, pumps, lighting, controls, battery charging, conversion losses, or cooling parasitics. Final utility sizing should be based on the complete IT load, target PUE, redundancy level, ambient design conditions, and manufacturer data.

Technical Configuration Envelope

System Domain	Configurable Options
Compute	GPU, CPU, accelerator, and mixed-density rack configurations
Power	Utility, generator, UPS, PDU, busway, and branch-circuit options based on site requirements
Cooling	Air-cooled, rear-door heat exchanger, in-row, direct-to-chip liquid cooling, or hybrid approaches
Networking	High-speed Ethernet and InfiniBand-ready architectures, fiber uplinks, and redundant switching
Storage	Local NVMe, scale-out storage, object storage, backup, and archival configurations
Security	Physical access control, surveillance integration, segmented networks, and monitored environments
Operations	Remote telemetry, lifecycle monitoring, alerting, and maintenance planning

Platform Capabilities

EdgeVault(TM) combines modular infrastructure, high-density compute readiness, enterprise networking, secure operations, and repeatable expansion into one deployable platform.

Modular, Rapidly Deployable Infrastructure A repeatable platform approach designed to reduce field construction and shorten time to compute.	High-Density GPU Computing Supports concentrated accelerator deployments for AI training, inference, rendering, simulation, and analytics.
AI Training and Inference Configured for model development, fine-tuning, batch inference, real-time inference, and private AI environments.	Private Cloud and Virtualization Supports virtual machines, containers, private cloud stacks, and enterprise application hosting.
High-Performance Computing Enables simulation, scientific modeling, digital twins, research, and other parallel workloads.	Enterprise Storage and Networking Integrates high-speed fabrics, resilient switching, local and scale-out storage, backup, and data protection.
Edge Computing Places compute closer to data sources, facilities, campuses, industrial operations, or regional users.	Disaster Recovery and Business Continuity Can support secondary compute, replicated data, failover systems, and continuity strategies.
Secure and Regulated Workloads Can be engineered around access control, network segmentation, monitoring, and applicable compliance requirements.	Hybrid Cloud Integration Supports connectivity between on-premises, private cloud, public cloud, and colocation environments.
Scalable Multi-Pod Expansion Allows customers to add compute, storage, and network capacity in repeatable phases.	Remote Monitoring and Lifecycle Management Supports telemetry, alarms, utilization visibility, preventive maintenance, and asset lifecycle planning.

Operational Benefits

Benefit	Business Impact
Schedule	Accelerates deployment compared with traditional data center development.
Capital Efficiency	Supports phased expansion so capacity can be added as demand grows.
Site Flexibility	Can be adapted for enterprise campuses, industrial properties, research sites, commercial properties, or regional edge locations.
Serviceability	Organized, monitored infrastructure supports planned maintenance and equipment replacement.
Resiliency	Can incorporate redundant power, network, and cooling elements according to the project requirement.

Built to Scale

EdgeVault(TM) can be deployed as a single modular pod or expanded into a multi-pod AI campus. The architecture is intended to let customers add capacity in stages while maintaining consistent standards for power distribution, cooling, networking, security, monitoring, and operations.

Illustrative Expansion Path

Growth Stage	Deployment Objective
Phase 1 - Initial Pod	Deploy core compute, networking, storage, and support infrastructure for the first production workload.
Phase 2 - Capacity Expansion	Add racks, GPUs, storage, or a second pod as utilization and customer demand increase.
Phase 3 - Resiliency Enhancement	Add redundant utility paths, generation, UPS, network diversity, cooling redundancy, or disaster-recovery capability.
Phase 4 - Multi-Pod Campus	Standardize multiple pods around shared utilities, operations, security, monitoring, and lifecycle management.

Deployment Considerations

Area	Validation Requirement
Site and civil readiness	Pad, access, drainage, setbacks, fire access, structural requirements, and local permitting.
Electrical infrastructure	Available utility capacity, voltage, transformation, switchgear, grounding, generation, and UPS strategy.
Thermal design	Peak IT load, ambient conditions, cooling method, water availability, heat rejection, and redundancy.
Network connectivity	Carrier access, fiber diversity, latency requirements, cloud connectivity, and security architecture.
Compliance and safety	Applicable building, electrical, fire, environmental, cybersecurity, and industry-specific requirements.
Final equipment selection	Actual GPU servers, racks, PDUs, cooling systems, storage, switches, controls, and monitoring platform.

EDGEVAULT(TM)

Secure. Scalable. High-Density. Rapidly Deployable.

Important: This document is a conceptual sales and planning brief, not a stamped engineering specification. Final performance, capacity, power, cooling, redundancy, code compliance, and deployment schedule must be confirmed through project-specific engineering and manufacturer submittals.