

OPTIMIZED FOR

AI/ML

A cost-effective option for accelerated computing and applications, Sesame, optimized for AI/ML, rack-level solutions combine the flexibility of open hardware with the highest performance for the most demanding workloads. Recertified, upgraded and right-sized, Sesame AI/ML racks are optimized and supercharged for cloud or enterprise applications using proven hyperscale technology.

Power and flexibility are the key to any learning structure. Sesame developed the algorithms and coded training for AI/ML so that your teams could have hardware to leverage the best of accelerated architectures in a manageable and configurable resource available on-premises.

Bringing together high-performance InfiniBand connectivity to large and small accelerated Open Compute Project (OCP) nodes, Sesame for AI/ML can put the power of large-scale infrastructure into the hands of any sized team.

Key Benefits

- **Scalable building blocks** enable configuration and tuning for specific needs
- **High- to ultra-high performance** via application acceleration
- **Flexible configuration** across node types
- **Massive data storage** enables workload flexibility and the ability to adjust data at scale, from small to large data sets
- **Iterative processing** and the ability to adjust as the ML algorithms process data shortens time-to-solution cycles




Available three- or six-year warranty



From 3 to 20 accelerated compute nodes



High- to ultra-high performance via acceleration



Example shown: Scale

Configuration

	BASE 15 Nodes	GROWTH 21 Nodes	MEDIUM 21 Nodes	SCALE 27 Nodes
Compute Nodes				
Storage Nodes	 360TB raw HDD@ 4TB	 360TB raw HDD@ 4TB	 1.22PB raw HDD@ 14TB	 1.8PB raw HDD@ 20TB
Infrastructure / Management Nodes				
Ingress / Egress Nodes				
External Network	Redundant 25G switches with 100G uplinks			
Internal Network	Redundant InfiniBand (IB) switches			

Performance Specs

- From three to 20 accelerated compute nodes
- Single- or dual-socket compute nodes with acceleration and 25 GbE connectivity
- Three dense storage nodes with up to a petabyte of capacity to keep your applications fed with data
- Configurable, cost-effective ingress and infrastructure nodes in groups of three
- Redundant external top-of-rack switches compatible with your existing network environments
- High-bandwidth internal switches for optimized performance and scalable cross-rack connectivity
- Variety of acceleration options and configuration to meet workload needs

Node Specs

	NODES	CONFIGURATION	CATEGORY
Compute 2S	> 20 cores (40 vcpu) 256GB memory	Optimized for two-socket high compute	High CPU
Storage	> 10 cores > 5TB > 1M IOPS	Optimized for high-capacity, high-IOPs flash storage	Fast storage
Compute 1S	> 10 cores	Optimized for one-socket medium compute	Medium CPU
Management, Ingress / Egress, Infrastructure	> 8 cores > 1TB high IOPS	Optimized for infrastructure services	Modest CPU + modest storage

Check out these other Sesame solutions built and tuned for what you need your infrastructure to do



Optimized for Kubernetes

Optimized for open source, cloud-native workloads requiring containerization and Kubernetes orchestration.



Optimized for Hyper-Converged

Optimized for combining storage, compute, and networking to replace complex legacy infrastructure and reduce CapEx.



Optimized for Edge Computing

Optimized for deployment in distributed environments that require smaller form factors running forward-deployed workloads at scale.