SESAME
BY ITRENEW

OPTIMIZED FOR

Kubernetes

Sesame offers one of the most cost-effective Kubernetes cluster platforms available. Open-source, cloud-native orchestration via Kubernetes has become the go-to approach for automating application development, operations and management. Sesame for Kubernetes rack-level solutions deliver performance and efficiency at scale using proven high performance, hyperscale technology.

The modern DevOps teams you work with know the value of cloud-native applications. The ability to build clusters, automate deployments and deliver consistent infrastructure is what we had in mind when engineering our Sesame for Kubernetes series clusters.

By leveraging the ease of scaling and the ability to cluster using flexible software infrastructure, Sesame racks are designed to enable and optimize cloud-native applications of all kinds.

Key Benefits

- **Built for Kubernetes** clusters, tuned for a variety of container or VM solutions
- **Cloud-native apps** with a focus on cloud-scale distributed infrastructure
- **Flexible configuration** with single or dual-socket compute nodes, infrastructure nodes and networking nodes
- **Easy scalability** based on your needs, from small (six nodes in a single rack) to very large (600 nodes across 20 racks)
- **Production-ready** using field-proven hardware
- **Fully open** configuration and manageability tools, qualified and optimized for Linux

Available three- or six-year warranty

From 6 to 600 compute or storage nodes

Trusted performance using proven technology

sesame@itrenew.com +1-833-573-7263 itrenew.com/sesame

Tech Specs

Example shown: **Scale**

Configuration

	BASE **18 Nodes**	GROWTH **24 Nodes**	MEDIUM **30 Nodes**	SCALE **48 Nodes**
Compute Nodes	6	8	10	15
Storage Nodes	6	6	6	6
Infrastructure / Management Nodes	3	3	3	3
Ingress / Egress Nodes	3	3	3	3
External Network	Redundant 10G/40G switches with 40G uplinks			
Internal Network	Redundant 10G/25G switches with 100G cross-rack links			

Performance Specs

- From six to 600 compute and storage nodes
- Single- or dual-socket compute nodes with 25 GbE nodes connectivity and 100 GbE rack connectivity
- Flash-based storage nodes with millions of IOPS and terabytes of capacity
- Configurable, cost-effective management and infrastructure nodes for system services
- Redundant external top-of-rack switches compatible with your existing network environments
- High-bandwidth internal switches for optimized in-cluster performance and scalable cross-rack connectivity up to 20 racks

Node Specs

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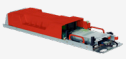	**Compute 2S** > 20 cores (40 vcpu) 256GB memory	Optimized for two-socket high compute	High CPU
	Compute 1S > 10 cores	Optimized for one-socket medium compute	Medium storage
Storage 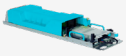	> 10 cores > 5TB > 1M IOPS	Optimized for high-capacity, high-IOPs flash storage	Fast storage
Management, Ingress / Egress, Infrastructure 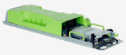	> 8 cores > 1TB high IOPS	Optimized for infrastructure services	Modest CPU + modest storage

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

Optimized for AI/ML

Workloads that require massive data volumes in demanding learning environments using models like TensorFlow.

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.