WildFire Storage

More performance. Less hardware. Lower cost.

Scan for details
wildfire-storage.com/grid

Scale-out block storage for cloud providers and AI workloads

maxRAID Grid

Add nodes. Add performance.

WHY IT'S DIFFERENT

No central metadata server: every pool performs like a local array, and performance grows with every pool you add.

Beta Q4 2026 · general availability Q1 2027

20 GB/s
per storage pool, read and write[1]
Grow when you need to. Performance grows with you.

144 vs 361
SSDs for 1.8 PB usable, vs Ceph with three copies[2]
A fraction of the flash.

2 failures
any two drives per pool, or two whole servers, with 8+2 pools — add parity for more[1]
Servers fail. Service doesn't.

What you get

Pools that span the grid.
N nodes with N NVMe drives form N pools, each with one drive in every node.

No central metadata server.
Each pool runs independently and performs like a local array.

Spares, standing by.
A spare drive in every node, a spare node, or both — rebuild starts at once.

Built for AI.
NVMe-oF over TCP and RDMA, GPUDirect, and shareable multi-mount volumes.

Multi-tenant.
Thousands of volumes per server; QoS by bandwidth or IOPS, burstable.

Your hardware.
White-box x86_64 servers and standard Ethernet.

Specifications

Performance	Up to 20 GB/s and 2M IOPS per pool (NVMe), read and write, network permitting; volumes stripe across pools
Protection	Parity you choose (+2 or more); each pool survives that many failed drives. 8+2: any two drives or two servers
Protocols	NVMe-oF (TCP, RDMA, GPUDirect) and iSCSI; shareable multi-mount volumes
Self-healing	Automatic SSD replacement; server failover to hot spares via virtual target IPs
Geometry	Square: N nodes × N drives = N pools; grow by adding grids
Scale	50 TB to 10+ PB
Data services	Snapshots, backups, inline compression, thin provisioning
Hardware	x86_64 servers, standard Ethernet, NVMe SSDs (TLC, QLC, PLC)

Works with Kubernetes

Grid ships with a CSI driver: thin-provisioned persistent volumes on demand, with snapshots and shared volumes, over NVMe-oF.

Request beta access — or size your grid at wildfire-storage.com/grid.

wildfire-storage.com/grid · +1-707-maxRAID · sales@wildfire-storage.com

1. Pre-release: performance and resilience are design targets, provided the network can carry the traffic.
2. Modeled: 1.8 PB usable on 15 TB SSDs, 12 × 12 grid, no compression (110 SSDs with typical data).
3. Kubernetes is a registered trademark of The Linux Foundation.

© 2026 EasyCo LLC dba WildFire Storage · maxRAID Grid data sheet · Specifications may change without notice