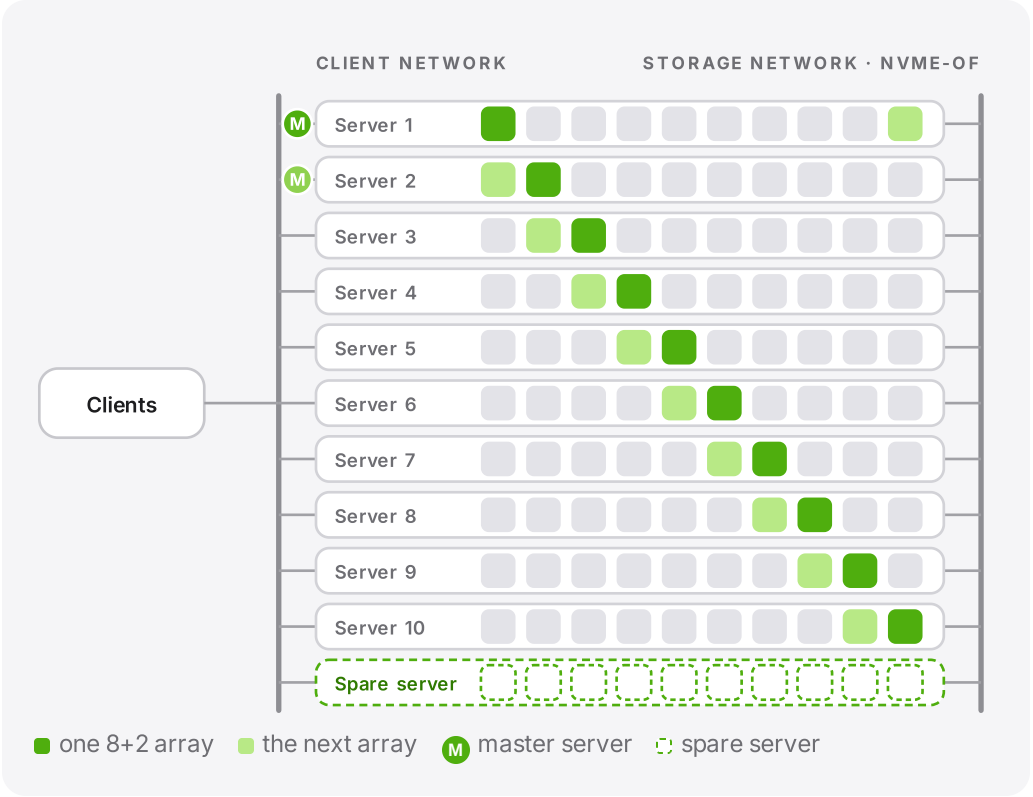
WildFire Storage

AGAINST CEPH WITH THREE COPIES

maxRAID Grid vs Ceph

Resilient block storage across many servers, for cloud and AI.



wildfire-storage.com/grid

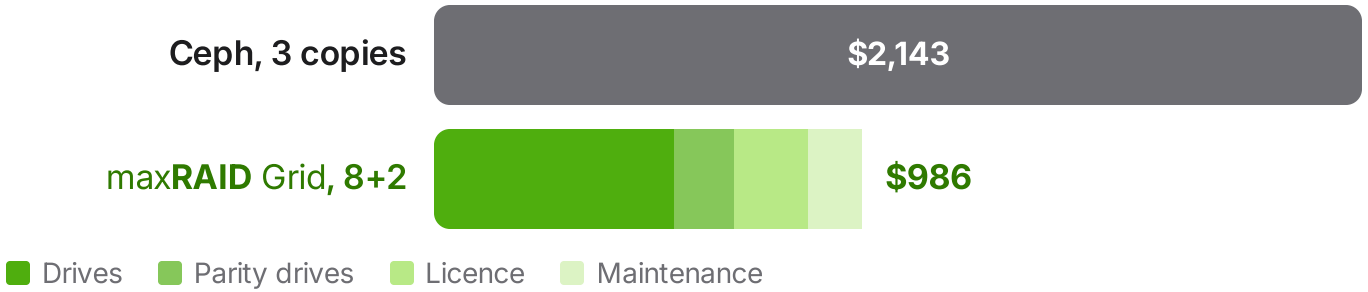
Features and figures may change. maxRAID Grid performance figures are design targets.

WildFire Storage

AT A GLANCE

More flash. More speed. Lower cost.

Cost per usable TB, five years, all in



Shown with 2 parity drives, as resilient as Ceph's 3 copies. **Your choice: 1, 2, 3 or more.**

	Ceph	maxRAID Grid
Protection	3 full copies (33%)	8 data + 2 parity (80%)
How full it runs	70%	90%
Usable share of raw flash	23.3%	**72.0%**
Cost per raw TB	$500	$500 SSD + $120 licence
Maintenance	—	$18 per raw TB a year

$986
per usable TB over five years. Ceph: $2,143.

3.1×
the usable capacity, same flash

26×
the writes per server, and 5× the reads

SSDs at $500 per raw terabyte. Drives and parity drives at that price (in this example, 2 of every 10 drives are parity; you choose how many). maxRAID Grid list prices: licence $120 per raw terabyte, parity drives included, plus maintenance $18 per raw terabyte a year, five years counted. Ceph is counted with no licence or support cost. Usable terabytes = raw × share that holds data × how full it runs. maxRAID Grid speed figures are design targets.

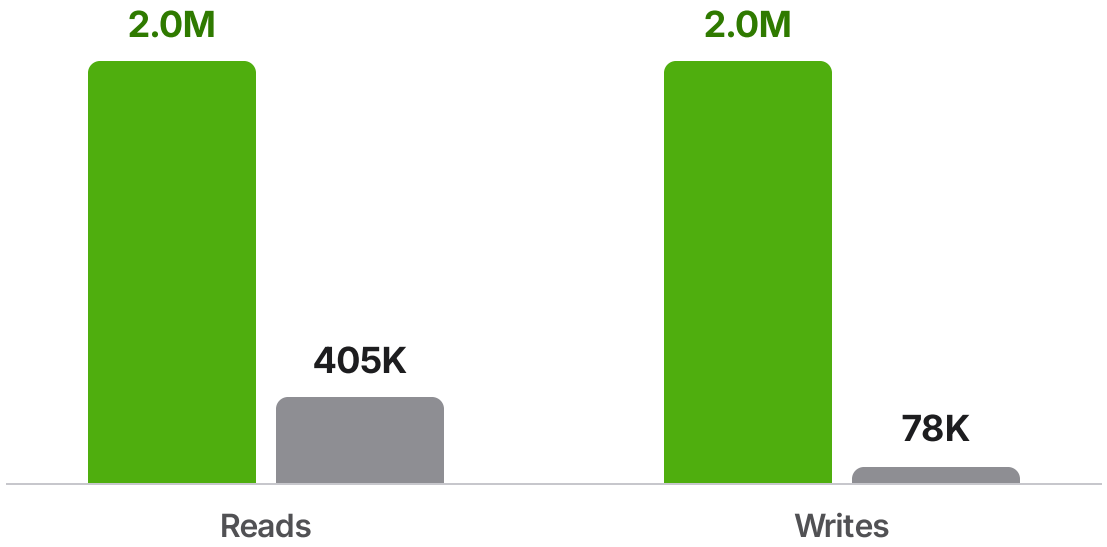
WildFire Storage

PERFORMANCE PER SERVER

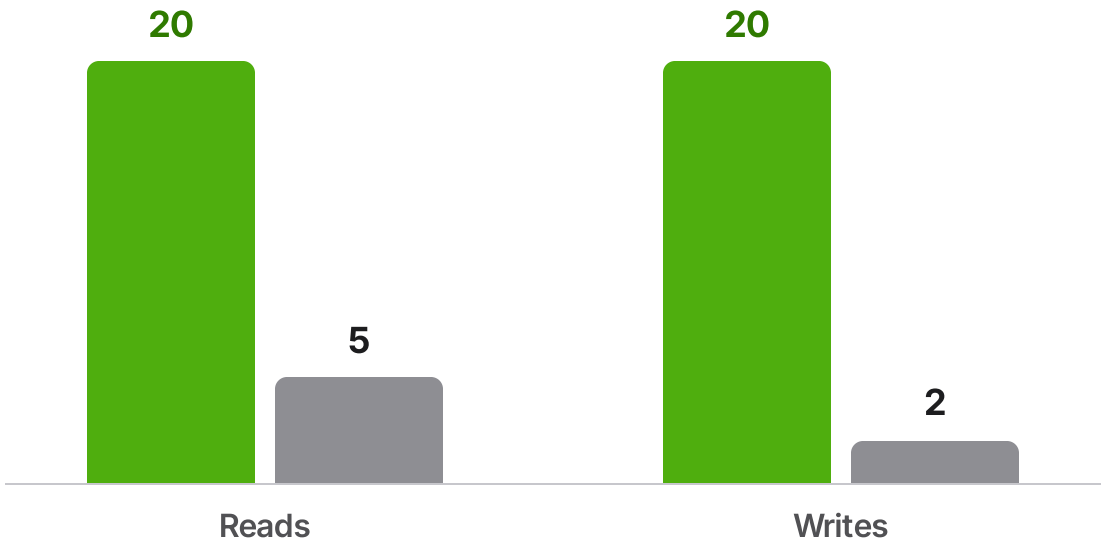
5× the reads. 26× the writes.

Handle the same load with fewer servers.

Small blocks (4 KB): operations per second



Large blocks (64 KB): GB per second



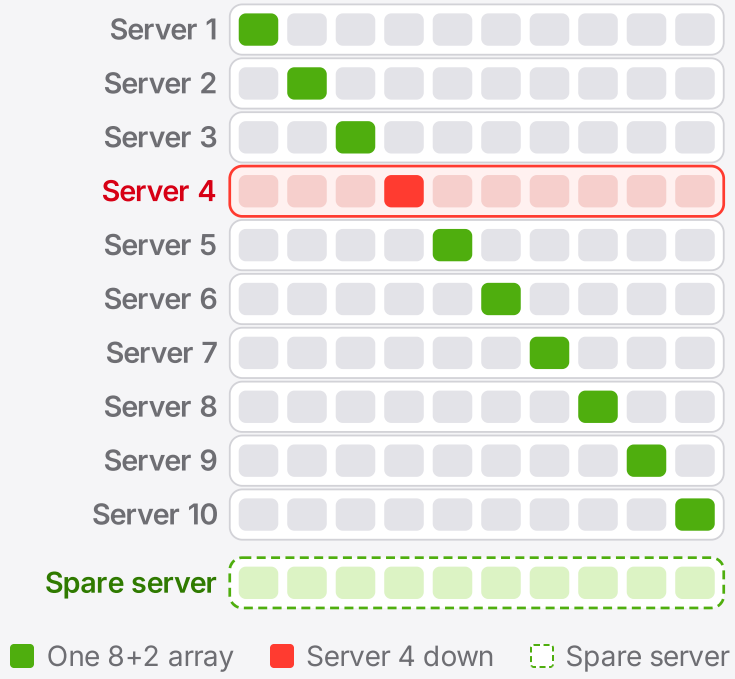
maxRAID Grid Ceph, 3 copies

Ceph (measured): "Ceph: A Journey to 1 TiB/s", ceph.io, January 2024: 63 servers, each with 10 NVMe SSDs and 2 × 100 GbE, 3 copies, tested on the storage servers themselves; per server = cluster result ÷ 63. The Ceph 64 KB bars are estimates: "Ceph Reef Freeze Part 1" (ceph.io, 2023) measured 128 KB random writes at about 2.3 GB/s per server; reads taken as about 3× writes. **maxRAID Grid:** per-server design targets, with the same 10 NVMe SSDs and 200 GbE of networking per server.

WildFire Storage

RESILIENCE

Lose a drive. Lose a server. Keep serving.



One drive in every server.

Losing a server costs each array just one drive.

Protection you choose.

With 2 parity drives, any two drives or servers can fail. Or pick 1, 3 or more.

Back to full strength, automatically.

Drives rebuild from the other servers. A spare server steps in.

10 servers × 10 SSDs, with one spare server. Each logical array (8 data + 2 parity drives) has one drive in every server; a failed drive is rebuilt onto spare capacity from its array's drives in the other servers, while the array keeps serving. The number of parity drives is your choice; 2 is shown because it matches the resilience of Ceph's three copies. Features and figures may change.

HOW IT FITS

Fits what you run. Grows with you.

Your servers. Your SSDs.

Standard x86, any SSD, Ethernet.

Kubernetes, built in.

Volumes on demand, through a CSI driver.

Standard block storage.

NVMe over TCP or RDMA, or iSCSI. Shared volumes for AI.

Add servers, keep growing.

Data rebalances automatically.

The Kubernetes driver ships with general availability. RDMA = direct memory access across the network; iSCSI = block storage over standard IP networks. Features and figures may change.

WildFire Storage

NEXT STEP

See it run. On your servers.

How an evaluation works

1. We size it with you.
2. We set it up on your servers.
3. You test it your way.

Talk to us any time.

wildfire-storage.com/grid

sales@wildfire-storage.com
+1-707-maxRAID (707) 629-7243