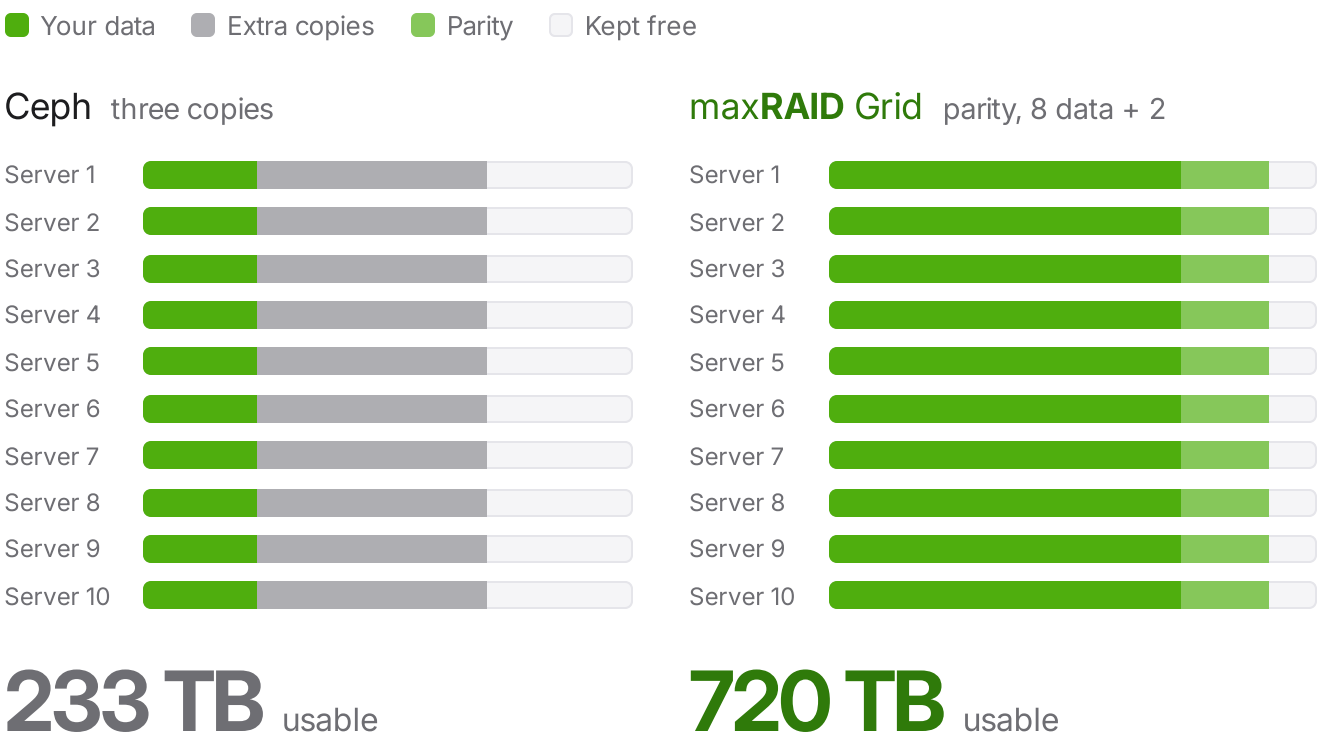
maxRAID Grid · For teams running Ceph

Same 1 PB of flash.
720 TB usable. Ceph: 233 TB.

And 5× faster than Ceph, per server.[2]

maxRAID Grid is block storage for the same kind of servers Ceph runs on. Ceph block storage usually keeps three full copies for speed: its parity mode (erasure coding) is slower for small writes. Grid keeps parity at full speed: a little recovery data spread across the servers, enough to rebuild anything lost. So most of every drive holds your data.[3]

Same flash. Count the usable terabytes.



233 TB usable

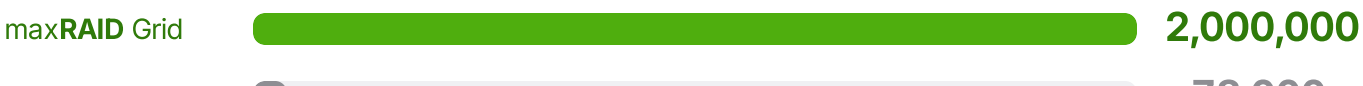
720 TB usable

That's before compression. maxRAID Grid compresses data inline: data that compresses by 30% fits about 940 TB in the same flash.[4]

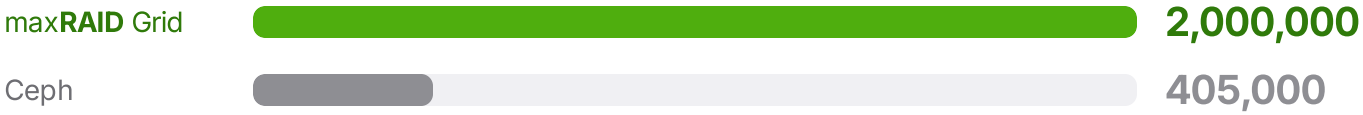
From the same 1 PB of raw flash: 10 servers, 10 SSDs each. Ceph with three copies, as most block clusters run for speed; both survive any two failed drives or servers.[1]

Same servers. Count the writes.

Small writes a second, per server



Small reads a second, per server



4 KB random writes and reads, per server. maxRAID Grid: design targets. Ceph: its published 63-server result, divided by 63. Large blocks: up to 20 GB/s per server on Grid.[2]

maxRAID Grid next to Ceph

COST

$986

per usable TB, five years, all in[1]

Ceph: $2,143

SSDs, licence and maintenance included, before compression.

FAILURES

Any 2 or more

drives or whole servers can fail

and it keeps serving

Shown with 2 parity drives, as resilient as Ceph's three copies. Your choice: 1, 2, 3 or more. A spare server takes a failed server's place, automatically.

Fits what you run

- Standard x86 servers, any SSD, Ethernet
- Inline compression and thin provisioning
- Kubernetes CSI driver
- NVMe-oF or iSCSI

See it run. On your own servers.

We size it with you, set it up on your servers, and you test it with your own workload.

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

Read it online
wildfire-storage.com/resources/more-from-your-ceph-hardware

1. SSDs at $500 per raw TB. maxRAID Grid list prices: licence $120 per raw TB plus maintenance $18 per raw TB a year, five years counted; Grid at 8 data + 2 parity drives (as resilient as three copies), 90% full. Ceph with three copies, 70% full, with no licence or support cost counted. Usable TB = raw × share that holds data × how full it runs: 233 TB vs 720 TB from 1 PB.
2. Ceph: "Ceph: A Journey to 1 TiB/s", ceph.io, January 2024 — 63 servers, each with 10 NVMe SSDs and 2 × 100 GbE, three copies, tested on the storage servers themselves; per server = cluster result ÷ 63; 4 KB random reads and writes. maxRAID Grid: per-server design targets, with the same 10 NVMe SSDs and 200 GbE of networking per server. 5× faster: small reads, 2,000,000 against 405,000 a second per server; small writes, 2,000,000 against 78,000.
3. Ceph supports erasure-coded pools for block storage (RBD, with overwrites enabled), at a cost in small-write performance; most block deployments use three-way replication.
4. Compression is modeled, not measured: data that compresses by 30%, as in the maxRAID Grid planning model: about 1.31 times as much data in the same flash, 720 TB → about 940 TB. The comparison above counts no compression on either side.
5. Features and figures may change.