maxObject · For teams running Ceph RGW

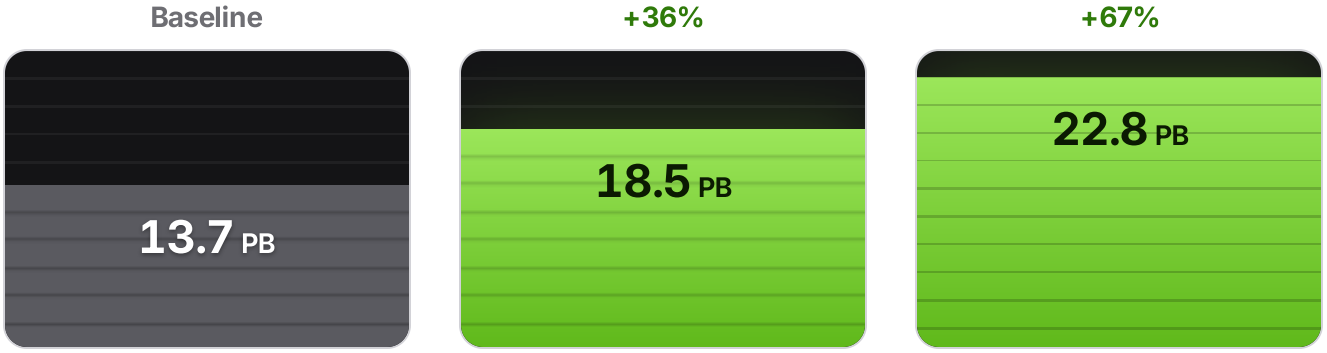
Same 1,000 hard drives. 18.5 PB usable. Ceph RGW: 13.7 PB.

And 5× faster than Ceph RGW.[6]

maxObject is S3 object storage for the same kind of servers and hard drives Ceph RGW runs on. Ceph splits every object across twelve data drives, so one small read keeps twelve drives busy. maxObject packs many objects into each protected stripe, so a small read usually touches one drive. The drives serve more customers, run fuller and rebuild faster.[1]

Same 1,000 drive slots. Count the petabytes.

Usable capacity in the same 1,000 drive slots

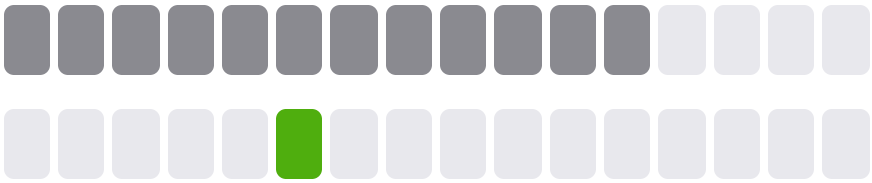


Ceph RGW	maxObject	maxObject
26 TB CMR drives	26 TB CMR drives	32 TB host-managed SMR drives

On the same conventional drives: 4.85 PB more to sell, about $350,000 more revenue a year.[4]

50 servers × 20 drives. Every case survives any four drive losses. Ceph RGW does not support host-managed SMR drives.[1]

One small read. Count the drives.



Across 1,000 drives: about 120,000 small reads a second, against 10,000.[2]

One S3 GET of a typical small object. Ceph RGW splits each object 12 + 4 across sixteen drives and reads twelve; maxObject packs objects into wide stripes and usually reads one.

max**Object** next to Ceph RGW

MONEY

$1.1M

more over five years, after the maxObject licence[4]

Revenue after costs: $4.8M vs $3.7M

Same servers and drives, sold out at $6 per TB a month. The licence at list price is counted; Ceph RGW is counted free.

REBUILD

< 30 min

to rebuild a failed 26 TB drive[3]

Ceph RGW: 2.7 hours

Every healthy drive takes a share of the rebuild, and customers keep reading and writing.

FULLER DRIVES

About 85%

how full the drives run[5]

Ceph RGW: about 70%

Data is placed evenly, and recovery needs little spare room.

INDEX FLASH

1.3 TB

of index (metadata) flash per server[5]

Ceph RGW: about 17 TB

For a server of 20 × 26 TB drives. Leaner servers, lower cost.

Fits what you run

- Conventional (CMR) and host-managed SMR drives, mixed
- S3: AWS SDKs and CLI work as-is
- Tenants, policies, quotas and billing rules
- Hard-drive and flash tiers in one cluster
- One program per server, self-healing

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-hard-drives

1. Usable capacity in the same 1,000 drive slots (50 servers × 20 drives), with protection against any four drive losses in every case. Ceph RGW 12+4 on 26 TB CMR drives; maxObject with 36+4L+3G LRC sets on 26 TB CMR drives, and on 32 TB host-managed SMR drives, which Ceph RGW does not support.
2. Design arithmetic from published drive characteristics and IBM Cloud Object Storage's published traces, not benchmarks: 1,000 drive slots (50 servers × 20 drives), Ceph RGW at 12+4. Small reads: about 120 per drive per second for maxObject and 10 for Ceph RGW. A typical S3 GET reads one drive on maxObject and twelve on Ceph RGW, as described in the maxObject technical brief.
3. Estimate, not a benchmark: a full 26 TB drive, its rebuild shared by every other drive, on an otherwise idle cluster of 1,000 drives. maxObject streams at about 175 MB/s per drive. Ceph RGW at 12+4 reads each 4 MiB RADOS object as pieces of about 340 KiB, each costing an 8 ms seek: about 35 MB/s of useful repair per drive. Ceph's default recovery limits often make it slower in practice; under heavy customer load either rebuild takes longer.
4. Estimate. Ceph RGW at 12+4 and 70% full; maxObject at 36+4+3 and 85% full, on the same 1,000 × 26 TB drives. Five-year cost: servers at $12,000 without drives, hard drives at $15 per TB, metadata flash at $100 per TB, power at $0.12 per kWh including cooling; maxObject licence at list price ($15 per raw TB plus 15% a year, five years counted), Ceph RGW with no licence or support cost; excludes rack space and staff. Revenue after costs: five years of revenue at $6 per TB a month with all usable capacity sold, minus five-year cost.
5. Estimate. Ceph RGW at 12+4 and 70% full; maxObject 85% full. Metadata flash for servers of 20 × 26 TB drives; 17 TB is the middle of the 13–21 TB that Ceph's published guidance implies.
6. 5× faster: design arithmetic, not a benchmark. Weighted by a week of real public-cloud traffic (IBM Cloud Object Storage's published traces), maxObject needs about one sixth of the drive time of Ceph RGW at 12+4 for the same requests, on the same 1,000 drives.
7. Features and figures may change.