

S3 Is the Future, S3 Is the Past

Viktor Leis



Announcing Amazon S3 - Simple Storage Service

Posted on: Mar 13, 2006

[Amazon S3](#) is storage for the Internet. It is designed to make web-scale computing easier for developers. Amazon S3 provides a simple web services interface that can be used to store and retrieve any amount of data, at any time, from anywhere on the web. It gives any developer access to the same highly scalable, reliable, fast, inexpensive data storage infrastructure that Amazon uses to run its own global network of web sites.

Built to host files on the web, became the foundation of the modern cloud:

- Cheap, durable storage
- Illusion of infinite capacity
- Metered on-demand pricing
- Strongly consistent (since 2020), conditional writes (2024)

S3 and Analytics

- Snowflake pioneered S3 for analytics
- Redshift, Microsoft Fabric, etc. followed
- Modern data lakes and lakehouses (Parquet + Iceberg) made S3 the standard for everybody
- Enables serverless compute
- System of record and long-term storage



S3 Is Everywhere

- Snowflake pioneered S3 for analytics
- Redshift, Microsoft Fabric, etc. followed
- Modern data lakes and lakehouses (Parquet + Iceberg) made S3 the standard for everybody
- Enables serverless compute
- System of record and long-term storage



WarpStream

Kafka



Turbopuffer

vectors



InfluxDB

time series



GreptimeDB

metrics



Tesla

IoT data

...

Simple model:

- “1 million magnetic disks”
- Files are hash partitioned across disks
- Redundancy (erasure coding) for durability

Simple model:

- “1 million magnetic disks”
- Files are hash partitioned across disks
- Redundancy (erasure coding) for durability

Explains many of S3's performance characteristics:

~10 ms

latency per request

80 MB/s

throughput per request

High throughput from one node with many large parallel requests

From HDDs to S3 to Cloud Architecture

Good:

HDDs	⇒ low capacity cost	⇒ lowest storage tier
Wide erasure coding	⇒ high durability/availability	⇒ system of record


From HDDs to S3 to Cloud Architecture


Good:

HDDs $\Rightarrow$ low capacity cost $\Rightarrow$ lowest storage tier
Wide erasure coding $\Rightarrow$ high durability/availability $\Rightarrow$ system of record

Not so good:

Wide erasure coding $\Rightarrow$ no in-place writes $\Rightarrow$ log-structuring necessary
HDD seek times $\Rightarrow$ 10+ ms latency $\Rightarrow$ caches necessary
 ≈ 100 IOPS/HDD $\Rightarrow$ operation quotas, high I/O cost $\Rightarrow$ large objects + caches








\$250,000/year

cost for 20K GETs/s






SSDs



$100\times$
bandwidth

$1/100\times$
latency

$10000\times$
IOPS



What Would an SSD-Based S3 Change Architecturally?

- No need for tiering
- No need for SSD caches
- No need to have large sequential reads
- No need to batch writes

But What About Capacity Cost?

- Current HDD/SSD cost gap is $\approx 5\times$
- Long-term trends: NAND flash capacity scales faster than HDDs

But What About Capacity Cost?

- Current HDD/SSD cost gap is $\approx 5\times$
- Long-term trends: NAND flash capacity scales faster than HDDs
- 1TB in S3 for 5 years: \$1,260
- 1TB NVMe data center SSD (in 2026): \$500
- 1TB NVMe data center SSD (in 2025): \$125

- Iceberg on S3 achieves 1 tx/s [Ginter and Leis, VLDB'26]

- Iceberg on S3 achieves 1 tx/s [Ginter and Leis, VLDB'26]
- Pretty much every operational system implements a flash-based storage/caching service
- Aurora, MS Socrates, Neon, SDB, etc. all have it
- Why re-invent the wheel again and again?



btrblocks.com

Conclusion

S3 is the future:

- Truly non-incremental innovation: durable, elastic, metered storage service
- Became the foundation of cloud data systems

Conclusion

S3 is the future:

- Truly non-incremental innovation: durable, elastic, metered storage service
- Became the foundation of cloud data systems

S3 is the past:

- Its performance characteristics have little to do with the cloud per se, but are inherited from legacy disk technology
- Yet these characteristics shape the standard cloud data architecture
- Ironically, this happened just as flash is taking over

Conclusion

S3 is the future:

- Truly non-incremental innovation: durable, elastic, metered storage service
- Became the foundation of cloud data systems

S3 is the past:

- Its performance characteristics have little to do with the cloud per se, but are inherited from legacy disk technology
- Yet these characteristics shape the standard cloud data architecture
- Ironically, this happened just as flash is taking over

An SSD-based primitive could radically simplify cloud data systems.