

STORAGE FOR HYPERSCALE

In the movie, *The Matrix*, there is a scene where everything happening in the world is presented as just 1s and 0s. In the real world, the proliferation of unstructured data in all its formats has become the essence of everything we do, see, feel, and think. The rug you want to purchase online is a series of bytes, as is your DNA, images of Mars, or songs being heard. All these bytes must live someplace; that someplace must address the key challenges of storing data at scale.

The challenges of storing data at hyperscale can be categorized as:

- 1. Scalability** – eliminate the need for islands of storage; organize and deploy resources reflecting business needs
- 2. Accessibility** – standardize access to data with superior performance
- 3. Resiliency** – assure data availability and integrity

ActiveScale™ storage software architecture addresses the needs of hyperscale environments, delivering an easy-to-scale platform, data resiliency through its Dynamic Data Placement and Repair, and support for industry standards to simplify access and management.

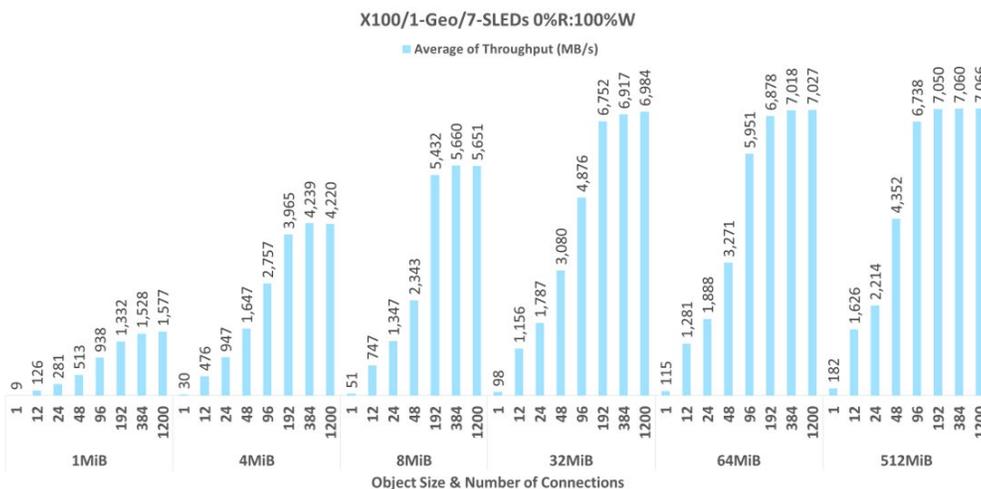
SCALABILITY

- Scalable metadata repository deployed with the storage nodes eliminates bucket file-count or capacity limitations.
- 4,500 concurrent S3 client connections per rack



ACCESSIBILITY

- Performance without resiliency loss is achieved by using Localized Ingest capability to ensure resiliency without loss of accessibility and performance. Writes are dispersed asynchronously to boost local performance, while metadata is dispersed synchronously.
- Access to data in a geo-dispersal configuration sustains performance of 75 GB/s with up to 100 ms of latency.
- No rebalancing is required; Refreshes, maintenance, and fixes are completed in the background without disrupting production workloads.



RESILIENCY

- Dynamic Data Placement (DDP) ensures that encoded objects are dispersed across all available nodes and media for greatest failure tolerance. The system can maintain data accessibility with multiple drive failures in a parity, chassis failure, or site failure in a geo-dispersed configuration. ActiveScale allows the business to dictate the level of acceptable resiliency.
- Dynamic Data Repair (DDR) continually checks the health of the data and corrects objects when found inconsistent with their hashes.



CONCLUSION

Hyperscale environments need to store data efficiently, but cost is only one of the necessary attributes. It is important that the system can scale with the needs of the organization, data stored is accessible over time, and that the data is secure. ActiveScale object storage addresses the needs of hyperscale by delivering increasingly scalable architecture, resiliency through Dynamic Data Placement (DDP) and Dynamic Data Repair (DDR) that disperse erasure encoded chunks across drives, chassis, and geographies, protecting against data loss and data corruption. The distributed, scale-out design supports high-throughput performance even in a geo-dispersed deployment; data is made accessible in multiple geographies with resiliency to lose a site and not lose data. Finally, multi-tenancy, encryption, and physical isolation offer advanced security for data.