Feeding AI at Scale

Validated hyperscale object storage performance on HDD infrastructure

Executive summary

Enterprise data infrastructure is evolving to support workloads that move massive volumes of data at high speed. AI data pipelines, large-scale analytics environments, and consolidated backup platforms increasingly rely on object storage as their primary data layer. These workloads require sustained high-bandwidth data movement across distributed environments operating under real production conditions, including encryption, load balancers, and multi-availability-zone architectures.

This paper presents a structured performance validation of a large-scale on-premises deployment running Scality RING 9.5 on HDD-based infrastructure. The objective was to evaluate whether object storage built on high-capacity disk systems can sustain hyperscale S3 throughput under high concurrency and full end-to-end request paths, rather than short-lived peak benchmarks.

Under a large-object workload profile using 10 MB objects and highly parallel access, the validated deployment sustained the following, as measured on the externally visible, encrypted production path:

- Read throughput of approximately 420 GB per second
- Write throughput of approximately 250 GB per second
- Stable system behavior with no storage-layer errors during sustained load

Sustained throughput (externally visible, encrypted traffic) Read ~420 GB/s · Write ~250 GB/s

These results demonstrate that HDD-based object storage can deliver sustained hyperscale S3 throughput under realistic deployment conditions. In a distributed multi-availability-zone environment with encrypted traffic and production infrastructure in the data path, Scality RING maintained stable aggregate bandwidth levels historically associated with parallel file systems. This validation confirms that large-scale AI data pipelines, analytics platforms, and backup environments can rely on object storage as a high-throughput data foundation while operating on high-capacity disk infrastructure.

Introduction

© 2026 Scality

The scale of modern data infrastructure is increasing rapidly. AI workloads, analytics platforms, and large data pipelines increasingly move data at aggregate rates measured in hundreds of gigabytes per second.

At the same time, organizations are deploying multi-tier storage architectures that combine different media types to balance performance, cost, and capacity at scale. Flash provides extreme performance for latency-sensitive workloads, while tape remains essential for deep archive. Hard disk drives (HDD) continue to play a critical role as the high-capacity tier that supports massive data volumes while delivering substantial aggregate throughput at a lower cost than flash.

As a result, object storage systems are no longer evaluated solely on capacity, durability, or API compatibility. They are increasingly evaluated on their ability to deliver sustained throughput under production conditions. Organizations need to understand whether HDD-based S3 object storage can support high-throughput workloads without becoming a bottleneck, in particular whether it can sustain large-scale parallel data movement while operating through real infrastructure paths, including encryption, load balancers, and distributed availability zones.

This document summarizes the validated environment, the test methodology, the measured results, and the implications of those results for high-throughput data workloads on HDD-based object storage infrastructure.

Environment overview

The performance validation was conducted on a large-scale on-premises deployment running Scality RING version 9.5. The system was deployed across three availability zones in a distributed object storage cluster built on x86 servers using HDD-based storage media in a 100+ node configuration.

Each availability zone exposed local S3 endpoints that were integrated behind a load-balancing layer and a global DNS entry point. Testing was executed through the DNS endpoint in HTTPS mode so that traffic traversed the same infrastructure components present in production environments, including TLS encryption, load balancers, and cross-zone routing.

Load generation was distributed across injector instances deployed in all three availability zones. A total of 120 injector virtual machines were used, evenly distributed across the zones. The primary large-object validation scenario generated up to 35,000 concurrent S3 operations distributed across 696 buckets in parallel. This configuration

© 2026 Scality

allowed requests to be issued simultaneously from multiple network segments and failure domains, reflecting realistic multi-application and multi-tenant access patterns.

The injector systems ran Red Hat Enterprise Linux (RHEL) 9.6 virtual machines provisioned with sufficient CPU, memory, and network capacity to avoid client-side bottlenecks. The injectors operated within a virtualized infrastructure where jumbo frames were not available, and traffic used a standard MTU of 1500 bytes. Network-capability validation on the injectors highlighted performance limitations below the theoretical capabilities of the underlying network infrastructure; this was quickly addressed by raising the number of injectors.

For the primary validation workload, 10 MB objects were transferred across multiple buckets in parallel, distributing traffic evenly across the cluster and avoiding artificial serialization. This workload size reflects common data-movement patterns observed in large-scale AI dataset distribution and analytical pipeline transfers.

Test methodology

The validation followed a controlled testing approach designed to measure sustained throughput under realistic production conditions. After validating baseline connectivity and bandwidth between injector systems and S3 endpoints, load generation was applied progressively to drive the system toward sustained high-throughput operation, using distributed injector instances and request-level parallelism across multiple buckets at high concurrency.

Scality developed an in-house performance-testing orchestrator to drive the hundreds of injectors from a single UI, with a single point of result collection and a scenario-based execution engine. Data injection and retrieval itself was carried out using a common third-party, open-source S3 testing tool. It was a key point for Scality that validation be carried out using common, industry-standard software to guarantee the authenticity of the results.

The primary validation profile used 10 MB objects transferred with standard S3 PUT and GET operations, directed through the global DNS endpoint in HTTPS mode so that the complete production request path, including TLS encryption and load-balancing infrastructure, was included in every measurement. To validate stability, the main benchmark scenario ran continuously over a two-hour window; additional shorter runs collected detailed injector-level statistics and validated behavior under different concurrency levels. Error rates were monitored throughout.

© 2026 Scality

Measurement points

Throughput was captured at three points in the request path:

- Performance-testing tool on the client side, showing what actual clients can expect in terms of performance.
- Externally visible encrypted traffic, measured at the frontend of the load-balancer layer, reflecting the full production path including TLS.
- Internal storage layer, measured at the Scality RING Supervisor, reflecting S3-layer throughput inside the cluster.

Both the client-side and external load-balancer measurements report the same figures: **approximately 420 GB/s read and 250 GB/s write**. This is expected, and it represents the throughput a production workload actually experiences. The corresponding internal storage-layer figures were **approximately 400 GB/s read and 240 GB/s write** over the same window.

This difference of approximately 5% is expected, because the measurement points sit at different layers of the protocol stack. The RING Supervisor counts the pure data payload, whereas the external measurement also includes S3 overhead such as authentication, headers and multipart-upload (MPU) packing, as well as the minimal overhead that TLS encapsulation adds.

Measured performance

Results are presented for the primary large-object validation profile and supplemented with observations from smaller-object testing.

Large-object profile: 10 MB objects

The primary scenario used 10 MB objects distributed across hundreds of buckets with high concurrency across all injectors, executed over HTTPS through the global DNS endpoint and load-balancer infrastructure. During the sustained two-hour run, throughput remained stable throughout the measurement window.

At the load-balancer layer (externally visible encrypted traffic), sustained throughput measured approximately 420 GB/s for reads, with peaks approaching 450 GB/s, and approximately 250 GB/s for writes. At the internal storage layer (RING Supervisor), sustained throughput over the same window measured approximately 400 GB/s for reads and 240 GB/s for writes; shorter measurement windows captured higher instantaneous read throughput during peak intervals.

© 2026 Scality
Scality · San Francisco · Paris · Washington, D.C. · Tokyo · London · www.scality.com

Injector-level statistics show the system serving approximately 23,000 to 43,000 large-object S3 operations per second under sustained load. No RING errors were reported during the validation window.

Measurement point	Read (GET)	Write (PUT)
Production path: encrypted, at the load balancer (externally visible)	~420 GB/s (peaks ~450)	~250 GB/s
Internal storage layer: RING Supervisor	~400 GB/s	~240 GB/s
Storage-layer errors (RING)	0	0

Table 1. Sustained throughput over the two-hour validation run, measured at both the production (load-balancer) and internal storage layers. The two layers differ for the reasons described under Measurement points.

Small-object profile: 4 KB objects

Additional testing used 4 KB objects to evaluate behavior under high object-per-second conditions. The system achieved 173,456 PUT operations per second and 151,091 GET operations per second with zero reported storage-layer errors. While overall throughput in gigabytes per second is naturally lower for smaller objects, the results confirm that the system maintains stable performance characteristics across different workload profiles.

Small-object request rate (4 KB) 173,456 PUT ops/s · 151,091 GET ops/s

Observations and interpretation

The measured results demonstrate sustained, distributed S3 throughput under multi-availability-zone, encrypted, end-to-end conditions. Several technical observations emerge from the validation campaign.

Sustained throughput stability

Across the two-hour window, aggregate throughput remained stable once the system reached steady-state load, with no degradation over time. The platform sustained high aggregate bandwidth under continuous load rather than relying on short peak bursts.

End-to-end measurement integrity

Performance was measured at both the storage layer and the load-balancer layer. Capturing throughput at both points shows the reported figures reflect externally visible

© 2026 Scality

encrypted traffic rather than internal-only system counters; the two layers track closely, with the small remaining differences explained under Measurement points.

Controlled saturation behavior

As concurrency increased, request latency rose predictably, which is expected when operating a distributed system near maximum throughput. Importantly, aggregate bandwidth remained stable and did not collapse under load, so the system continues to deliver consistent throughput even as request pressure increases.

Workload relevance

The 10 MB object profile aligns with common large-object workloads such as backup ingestion, data-lake transfers, AI dataset distribution, and large-scale content delivery. Sustained write throughput demonstrates the ability to ingest large data volumes efficiently, while sustained read throughput confirms the capacity to serve highly parallel read workloads. Additional testing with 4 KB objects shows that the system also maintains stable behavior under high object-per-second conditions, typical of metadata-intensive or checkpoint-driven workloads.

Enterprise use case implications

Sustained hyperscale S3 throughput is not simply a benchmark metric. It directly influences system efficiency and workload scalability in large enterprise environments.

AI data pipelines and retrieval workloads

Enterprise AI increasingly relies on large-scale data. Inference pipelines and retrieval-augmented generation architectures continuously access embeddings, documents, checkpoints, and intermediate artifacts stored in object storage. When multiple services operate concurrently, aggregate read bandwidth becomes a limiting factor. Sustained read throughput of approximately 420 GB per second enables large-scale parallel retrieval while maintaining predictable response times and efficient GPU utilization, which is relevant when object storage is the primary data layer for AI.

Large-scale backup and data protection

Enterprise backup strategies increasingly consolidate petabyte-scale datasets into distributed object storage platforms across availability zones. In these environments, write throughput directly influences backup windows and recovery objectives. Sustained write throughput of approximately 250 GB per second enables faster ingestion of parallel

© 2026 Scality

backup streams and improves the scalability of large backup environments. Combined with high read bandwidth, this also accelerates large-scale restore operations.

Data lakes and analytics platforms

Modern analytics platforms generate highly parallel read and write patterns across distributed compute clusters. Sustained throughput under multi-AZ and encrypted conditions means large-scale analytics workloads can operate without storage becoming a bottleneck. This enables consistent performance across large analytics clusters while maintaining a unified object storage architecture for AI, analytics, and data protection workloads.

Architectural implications

The significance of this validation extends beyond the absolute bandwidth achieved. It demonstrates that hyperscale S3 throughput can be sustained under realistic deployment conditions, including encryption, load balancing, and multi-availability-zone access. For infrastructure architects, this confirms that large-scale workloads such as AI data pipelines, analytics platforms, and backup environments can rely on object storage as a high-throughput data foundation without introducing additional storage tiers purely for bandwidth reasons.

Conclusion

This validation demonstrates that Scality RING can sustain hyperscale S3 throughput in a distributed multi-availability-zone environment under realistic deployment conditions. Over a continuous two-hour window, the tested deployment delivered stable aggregate throughput of approximately 420 GB per second for reads and 250 GB per second for writes on the externally visible encrypted path, corresponding to approximately 400 GB per second for reads and 240 GB per second for writes at the internal storage layer. High aggregate bandwidth was maintained under sustained load without throughput collapse or storage-layer errors.

Importantly, the measurements reflect full end-to-end request handling, including TLS encryption, load balancers, and multi-AZ traffic paths. The validation therefore represents production-like behavior rather than isolated component-level benchmarking. Additional testing with smaller objects further confirms stable system behavior under high request rates and diverse workload characteristics.

Taken together, these results show that large-scale HDD-based object storage can sustain the aggregate bandwidth required by modern data-intensive workloads. AI data

© 2026 Scality

pipelines, large-scale backup environments, and distributed analytics platforms increasingly depend on predictable high throughput across large datasets, and the validated performance confirms that Scality RING provides a scalable data foundation capable of supporting these workloads at hyperscale.

© 2026 Scality