SCALITY

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 the results of a structured performance validation conducted on 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 achieved:

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

Sustained Throughput

encrypted traffic

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

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, they must determine whether object storage built on high-capacity disk infrastructure can sustain large-scale parallel data movement while operating through real infrastructure paths, including encryption, load balancers, and distributed availability zones.

This document presents the results of a structured performance validation campaign conducted on a large-scale on-premises deployment running Scality RING 9.5. The objective of the validation was to measure sustained S3 throughput under high concurrency and distributed multi-availability-zone conditions representative of production environments. The validation focused on large-object workloads of 10 MB under high parallelism, reflecting common patterns in large-scale data movement such as AI dataset distribution and large analytical pipelines. The emphasis was on sustained throughput over extended duration rather than short-lived peak benchmarks.

This document provides a technical summary of the validated environment, test methodology, 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. The objective of the validation was to assess sustained S3 throughput under high concurrency and distributed load conditions representative of hyperscale and enterprise-grade environments.

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 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 testing injectors highlighted network performance limitations below the theoretical capabilities of the underlying network infrastructure. This was quickly addressed by raising the amount 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.

Performance metrics were collected at multiple layers of the system during the validation campaign. Measurements from the Scality RING Supervisor provided S3-layer visibility and were cross-validated with load balancer monitoring to confirm end-to-end throughput behavior across the complete infrastructure path.

Test methodology

The performance validation followed a controlled testing approach designed to measure sustained throughput under realistic production conditions rather than short-duration peak benchmarks.

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. The workload was generated using distributed injector instances and request-level parallelism across multiple buckets and high concurrency settings.

Scality developed an in-house performance testing orchestrator to drive the hundreds of injectors from a single UI, single point of result collection and scenario based execution engine. Data injection/retrieval itself was carried out using a common third-party S3 open-source testing tool. It was a key point for Scality that performance validation was carried out using common, industry-standard, third-party software to guarantee authenticity of the results. This approach ensured that the distributed architecture of the object storage cluster was exercised and that traffic patterns resembled real-world multi-application access behavior.

The primary validation profile focused on large objects of 10 MB, transferred using standard S3 PUT and GET operations. All traffic was directed through the global DNS endpoint in HTTPS mode, ensuring that the complete production request path was included in the measurements, including TLS encryption and load balancing infrastructure.

To validate stability under sustained load, the main benchmark scenario was executed over a continuous two-hour test window. This duration allowed verification that throughput remained stable over time and did not degrade after initial ramp-up. Additional shorter runs were used to collect detailed injector-level statistics and to validate behavior under different concurrency levels.

Performance metrics were collected at multiple points in the system. Internal S3-layer measurements from the Scality RING Supervisor were compared with load balancer monitoring to confirm consistent end-to-end throughput and account for expected differences related to TLS processing and monitoring resolution.

Additional test profiles explored behavior with smaller objects, including 4 KB workloads, to provide further insight into request-rate performance and system stability under different workload characteristics. Error rates were monitored throughout the campaign to ensure that sustained throughput was achieved without request failures.

This methodology ensured that the reported results reflect sustained, distributed S3 performance across a full production infrastructure path rather than isolated component-level measurements.

Measured performance

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 validation scenario used 10 MB objects distributed across hundreds of buckets with high concurrency across all injector instances. All traffic was executed over HTTPS through the global DNS endpoint and load balancer infrastructure, ensuring full end-to-end request processing.

During the sustained validation run of approximately two hours, throughput remained stable throughout the measurement window.

At the load balancer layer, reflecting externally visible encrypted traffic, sustained throughput measured approximately:

- Read throughput around 420 GB per second, peaks approaching 450 GB per second
- Write throughput around 250 GB per second

Shorter measurement windows captured higher instantaneous read throughput at the Scality RING Supervisor layer, exceeding 420 GB per second during peak intervals. Across the full two-hour validation period, internal read throughput remained stable while encrypted traffic measured at the load balancer sustained approximately 420 GB per second.

Sustained Throughput

Injector-level statistics show the system serving approximately 20,000 to 40,000 large-object S3 operations per second under sustained load. No RING errors were reported during the validation window.

Observed performance exceeded the target thresholds, with read throughput exceeding expectations and write throughput significantly above required levels.

Behavior under sustained load

As concurrency increased toward saturation, average request latency increased predictably, which is expected when maximizing aggregate throughput. Despite higher request latency at peak load, overall throughput remained stable. This indicates controlled saturation behavior rather than throughput collapse or instability.

No storage-layer errors were observed during the sustained validation window.

Stable Throughput

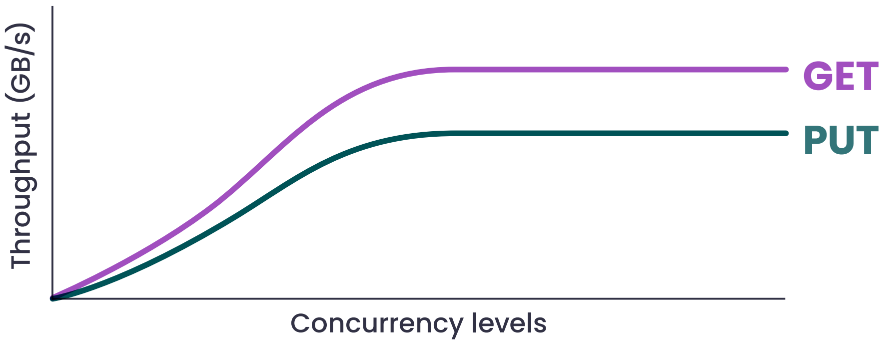

Small object profile: 4 KB objects

Additional testing was performed using 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.

Taken together, these measurements demonstrate sustained, distributed S3 throughput at hyperscale levels across multiple availability zones under encrypted end-to-end traffic conditions.

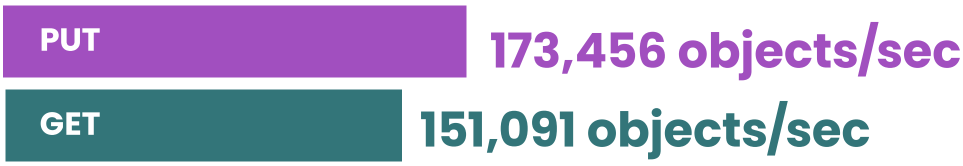

Observations and interpretation

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 validation window, aggregate throughput remained stable once the system reached steady-state load. No throughput degradation was observed over time. This behavior confirms that the platform sustains 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. Minor differences between Scality RING Supervisor measurements and load balancer measurements are expected due to TLS processing and monitoring resolution differences observed during sustained GET traffic. The close alignment between the two layers confirms that the reported throughput reflects externally visible encrypted traffic rather than internal-only system counters.

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. This indicates controlled saturation behavior and confirms that 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, which are typical of metadata-intensive or checkpoint-driven workloads.

Taken together, these observations indicate that the tested deployment delivers hyperscale-level S3 throughput under realistic multi-AZ and encrypted deployment conditions while maintaining predictable performance at high concurrency.

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 in the 380 to 420 GB per second range enables large-scale parallel retrieval while maintaining predictable response times and efficient GPU utilization. This 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 above 230 GB per second enables faster ingestion of parallel 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. Validating sustained throughput under multi-AZ and encrypted conditions confirms that 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 validation window, the tested deployment delivered stable aggregate throughput of approximately 231 GB per second for writes and 384 GB per second for reads at the storage layer, with externally visible encrypted traffic reaching approximately 250 GB per second for writes and 420 GB per second for reads. The results confirm that high aggregate bandwidth can be 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 demonstrate that large-scale HDD-based object storage can sustain the aggregate bandwidth required by modern data-intensive workloads. AI data pipelines, large-scale backup environments, and distributed analytics platforms increasingly depend on predictable high throughput across large datasets. The validated performance confirms that Scality RING provides a scalable data foundation capable of supporting these workloads at hyperscale.