Dell Puts 9.83PB in 2U With 245TB KIOXIA SSDs, Adds S3 over RDMA and KV Cache Offload to ObjectScale

by Harold Fritts on August 15, 2026

Enterprise ◇ Enterprise Storage

Dell has qualified KIOXIA's 245.76TB NVMe SSDs for software-defined ObjectScale, putting 9.83PB of raw flash in a single 2U server. The density milestone lands on top of the data-path work Dell detailed earlier this year across ObjectScale 4.2 and 4.3, which targets high-throughput AI training, small-object pipeline efficiency, and inference acceleration through native S3 over RDMA, Key-Value (KV) Cache offload for LLMs, an overhauled chunk-store engine, and direct Apache Iceberg support via S3 Tables.

Density Milestone: 9.83 PB in a 2U Footprint via 245 TB KIOXIA SSDs

As enterprise AI data lakes expand across model training, fine-tuning, and compliance-retention tiers, rack density, thermal envelope management, and floor-space constraints have become primary design variables. To address raw capacity per rack unit, Dell has qualified KIOXIA's 245.76 TB NVMe SSDs for software-defined ObjectScale running on Dell PowerEdge R7725 servers.

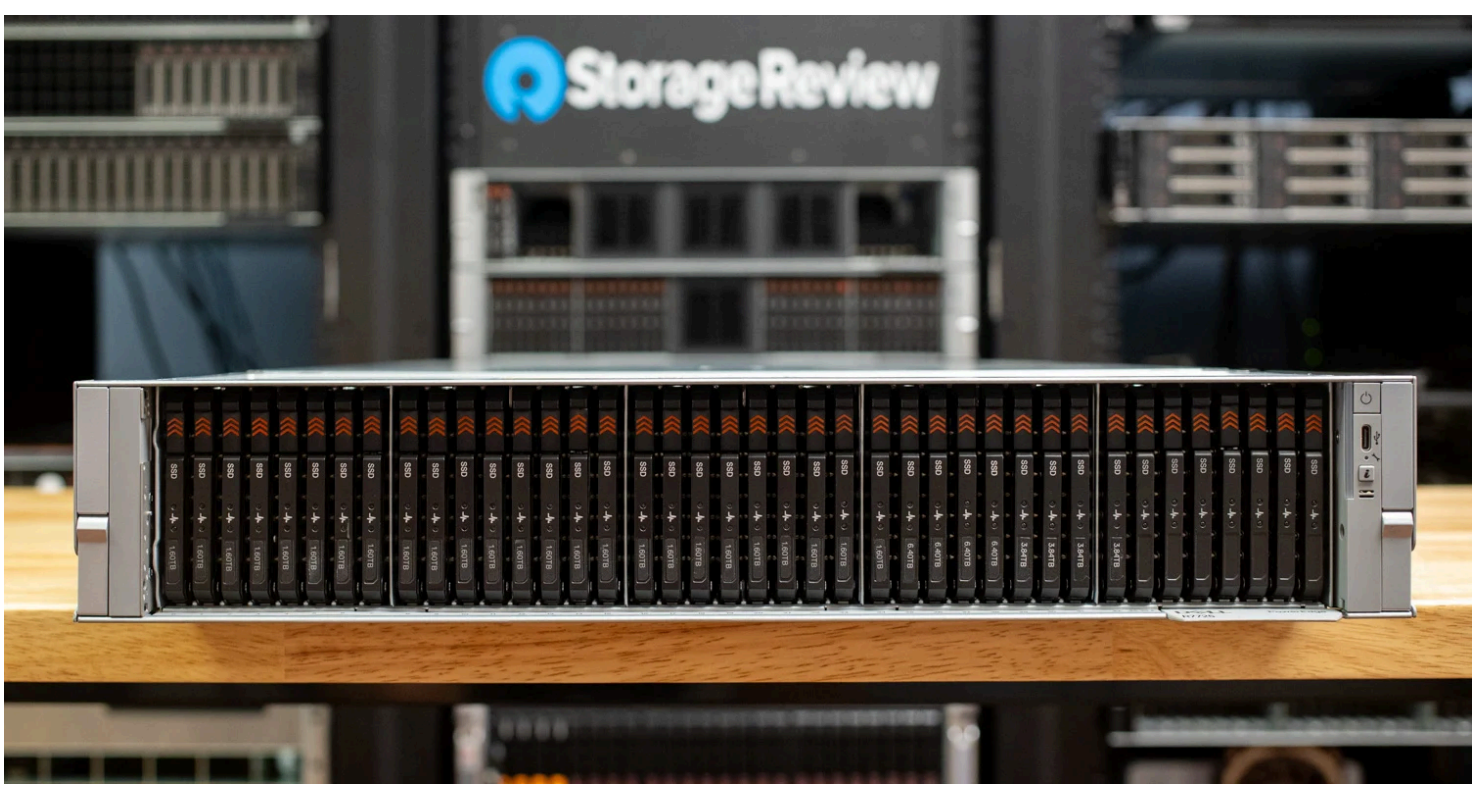

By populating 40 of the 245.76 TB drives in a single 2U chassis, ObjectScale reaches 9.83 PB of raw all-flash capacity; we reviewed the platform last summer. This density is aimed directly at enterprise AI factories, vector databases, large analytics repositories, and cyber-resilience platforms that require multi-petabyte unstructured storage footprints while minimizing power consumption, cooling load, and cabling complexity.

Throughput Gains and Small-Object Architecture

In software-defined deployments on qualified Dell PowerEdge servers, Dell's internal testing indicates per-node read throughput of up to 40 GB/s, an 8x increase over previous-generation all-flash object platforms. Dell also points to a UK high-frequency trading customer sustaining roughly 280 GB/s of aggregate read throughput in production. Those figures line up with what we have measured on the hardware itself: in our R7725xd testing last November, the platform sustained more than 300 GB/s across the internal NVMe pool and 160 GB/s over NVMe-oF RDMA. That work characterized the server directly rather than running ObjectScale on top of it, but it establishes the ceiling the software is working against.

To eliminate metadata bottlenecks common in pipelines dealing with millions of small feature sets, vector chunks, and telemetry logs, ObjectScale uses an optimized chunk-store engine. The architecture aggregates small payloads into unified 128MB chunks before applying erasure coding across nodes. For typical 10KB files, more than 10,000 objects fit into a single chunk, substantially lowering metadata overhead.

This chunking approach also restructures drive failure rebuilds by reducing the total number of shards that must be recreated from billions to millions, cutting recovery times for dense NVMe drives from weeks to hours. Inline checksums verified at the stripe level reduce host CPU cycles spent on background integrity scans. Furthermore, the re-engineered Key-Value Store in ObjectScale 4.2 provides roughly 4x better memory efficiency and cuts metadata disk footprints by 30% to 60%.

S3 over RDMA and Inference Acceleration

ObjectScale 4.3 expands S3 over RDMA across Dell's all-flash portfolio, including the XF960, EXF900, and software-defined PowerEdge R7725xd deployments. Utilizing RoCEv2 networking and Dell's S3-over-RDMA SDK, the platform bypasses kernel TCP overhead and CPU interrupts to establish direct memory paths between host GPUs and backend NVMe media. Internal testing shows this delivers up to 230% higher throughput, an 80% reduction in latency, and up to a 98% reduction in client CPU utilization compared to standard S3 over TCP.

For inference architectures, ObjectScale functions as an external tier for KV Cache offload. Because long-context inference quickly exhausts on-GPU HBM and VRAM, integrating with vLLM, LMCache, and NVIDIA's NIXL library allows inference frameworks to spill active attention states directly to ObjectScale over RDMA-accelerated S3. Dell reports up to a 19x improvement in Time to First Token against standard vLLM without cache offload, along with up to 5.3x higher token throughput and nearly 3x higher multi-turn throughput. Separately, Dell cites a TTFT of roughly 0.86 seconds in its own head-to-head testing against a competing object platform.

Native S3 Tables and Scale-Out Efficiency

ObjectScale 4.3 also brings native S3 Tables, which Dell lists as a tech preview, embedding Apache Iceberg open table formats directly within the storage bucket layer. Query engines including Trino, Starburst, Spark, and Flink can run analytics natively against S3 buckets without requiring dedicated staging or ETL data movement into separate warehouses. Dell reports that this architecture delivers up to 2x faster ingestion and up to 4.5x faster queries compared to traditional warehouse configurations.

Underlying storage efficiency includes 24+2 and 24+4 erasure coding profiles, which Dell says cut write amplification by up to 75%. On high-capacity HDD platforms like the EX500, the company reports up to 25% faster large-object ingest and up to 2x higher mid-size object write performance. Workload-specific bucket compression supports Snappy, LZ4, ZSTD, and Deflate algorithms with telemetry for FinOps cost tracking. Provisioning and orchestration integrate natively with cloud-native tooling, including Kubernetes Container Object Storage Interface (COSI) and Terraform providers.

Engage with StorageReview

Newsletter | YouTube | Podcast iTunes/Spotify | Instagram | Twitter | TikTok | RSS Feed

Harold Fritts

I have been in the tech industry since IBM created Selectric. My background, though, is writing. So I decided to get out of the pre-sales biz and return to my roots, doing a bit of writing but still being involved in technology.

Previous post: **NVIDIA Spectrum-X Ethernet Photonics Enters Full Production With 4x Fewer Lasers and a Five-Vendor CPO Supply Chain**

Ubiquiti

Newegg

Amazon

NEWSLETTER

Subscribe to the StorageReview newsletter to stay up to date on the latest news and reviews. We promise no spam!

Email...

SUBSCRIBE

ADVERTISEMENT

CONTENT CATEGORIES

- Consumer
  - Client Accessories
  - Client HDD
  - Client Software
  - Client SSD
  - Desktop Storage (External)
  - HDD
  - Memory Card
  - NAS
  - PCIe Storage
  - Portable Storage
  - Power
  - Workstation
- Enterprise
  - Accessories
  - AI
  - Attached Storage
  - Cloud
  - Data Protection
  - DPU
  - Enterprise Storage
  - HDD
  - HomeLab
  - Hyperconverged
  - In the Lab
  - Medium NAS
  - Networking
  - Power Management
  - Server
  - Server Rack
  - Small NAS
  - Software
  - SSD
  - Storage Adapters
  - Thin Client

Facebook X Instagram LinkedIn YouTube Email Spotify Reddit Discord RSS TikTok