

FUNGIBLE STORAGE CLUSTER

Scalable Workload-Centric NVMe over TCP Storage Cluster

BENEFITS



Unrivaled economics, performance, density and power efficiency: 10x the performance and 5x the ROI compared to existing solutions¹.



Best-in-class performance: 105M aggregate IOPS and 420 GB/s throughput, provided by seven FS1600 storage target nodes.



Fastest in-line data services: Data reduction, data durability, and data security at full line rate.



Cost optimized: 5x media savings compared to triple-replication with erasure coding, compression, and higher utilization.



Elastic scale-out, disaggregated storage: Shared and pooled capacity for business flexibility, high utilization, and efficiency.



Super-linear performance at cloud scale: Extends to thousands of storage target nodes.



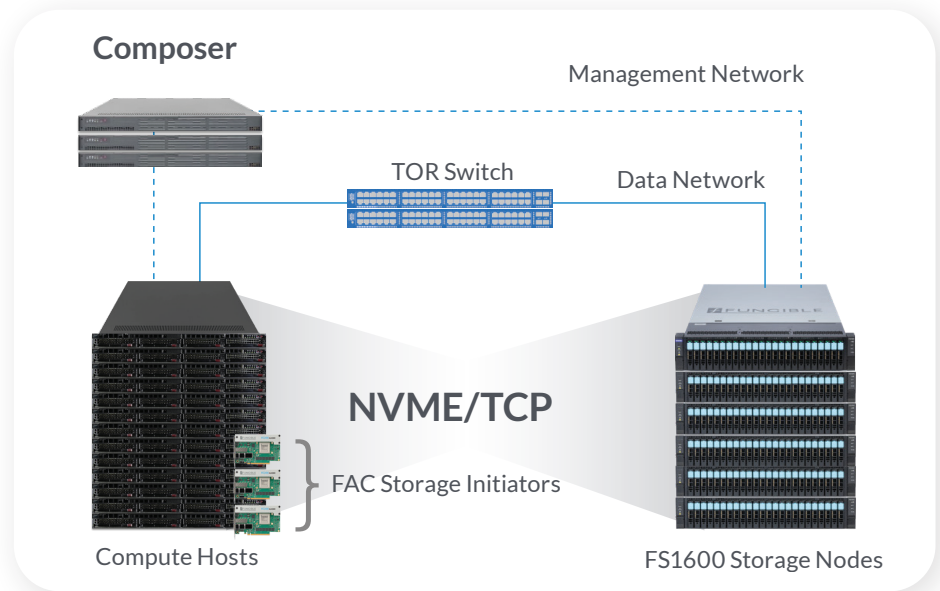
No-compromise security: Immutable root of trust, line rate end-to-end encryption, and fine-grained QoS.



Robust software stack: Optimized for maximum performance and flexibility.



Industry standard compliance: Smooth deployment in existing data centers.



OVERVIEW

Imagine what is possible with a single, massively scalable storage platform that provides NVMe performance to thousands of servers at a time. Experience the Fungible Storage Cluster (FSC)—the world’s fastest All-Flash NVMe over TCP disaggregated storage cluster. Designed to overcome every legacy limitation and address the most demanding applications, FSC delivers best-in-class performance characteristics for IOPS, latency, and throughput, in the smallest footprint and lowest power envelope.

FSC comprises a cluster of two or more FS1600 storage target nodes, three Fungible Composer nodes, and optional Storage Initiator cards. All control is managed and orchestrated by the Fungible Composer software.

Each storage node is powered by The Fungible Data Processing Unit™ (DPU), a new class of microprocessor purpose-built to execute data-centric infrastructure services an order of magnitude more efficiently than general purpose CPUs. In a cluster, the FS nodes implement a highly performant network-protected durable data plane that scales to petabytes of capacity and thousands of nodes.

The Fungible Composer software enables a highly available control plane that supports data plane management, automated provisioning, telemetry and monitoring, alerts and notifications, and more, streamlining administration and reducing operational costs.

SOFTWARE RELEASE 4.1 AVAILABLE NOW

With the launch of 4.1, you can harness the efficiency of erasure coding (EC) with as few as 3 storage nodes and build to any size you need. Fungible now also offers EC 2+1. The FSC 4.1 software release combines the unparalleled performance of raw volumes with the rich data services of durable volumes for RF1 volumes. RF1 volumes are single replicated but offer compression and encryption services to offer workload-centric flexibility.

The Storage Initiator, leveraging either the FC50 and FC200 as the physical platform, delivers performance of ~2M IOPS per card and supports up to 128 volumes per card.

¹ Internal comparison with software-defined storage solutions, which provide durability through triple-replication.

UNRIVALED PERFORMANCE DENSITY

Built in a 2RU form factor, each FS1600 node provides 12M IOPS¹ and 75 GB/s of throughput. A single FS1600 offers a breakthrough 10x performance compared to SDS-based solutions, making it the fastest and densest solution in the market today.

Unlike other storage solutions in which data services such as compression and encryption reduce overall system performance, the FSC provides sustained line rate performance when these data services are simultaneously enabled, scaling linearly to thousands of nodes.

ELASTIC HIGH PERFORMANCE, COST EFFECTIVE STORAGE POOL

Traditional storage solutions have compelled IT administrators to adopt a siloed approach to storage, with silos optimized for IOPS and latency, throughput, or capacity. With the performance density of the FSC, diverse workloads can be consolidated into a high performance, cost-efficient storage pool, eliminating the need for storage silos and the trade-offs that have traditionally existed between performance and cost. For example, a single FSC deployment can serve both IOPS and latency-sensitive applications, such as transactional databases, and throughput-hungry applications, such as modern AI/ML workloads. For capacity optimized storage, the FSC data reduction capabilities significantly reduce capacity requirements and allow data reduction to be applied without affecting storage performance.

SUPER LINEAR CHARACTERISTICS AT CLOUD SCALE

Many existing solutions work well for small clusters but suffer from performance degradation and unpredictability as cluster size increases. The FS node connected to host compute nodes via NVMe/TCP adds less than 10µs of average network latency compared to direct-attached storage (DAS). In a cluster, the FSC achieves super-linear line rate performance and throughput, maintaining superior average and tail latencies even in clusters with thousands of nodes.

LINE RATE, IN-LINE DATA SERVICES

FSC fully implements the storage, networking, and security stack on the Fungible DPU™, displacing CPU-based storage controllers used in today's storage target nodes. The FS node also leverages hardware accelerators in the Fungible DPU to deliver true in-line line rate data services such as:

Compression/Decompression: Each FS node supports DEFLATE (GZIP, ZLIB) and LZMA compression algorithms. Hardware accelerators deliver a combined compression/decompression throughput of 256 GB/s. With such a highly performant compression/decompression engine, data can be compressed in the first pass, removing the need for deep compression during garbage collection. In addition, data can be decompressed for read operations without affecting read performance. Unlike traditional storage solutions in which compression ratios are compromised to support higher throughput, the FS1600 achieves line rate throughput at high compression ratios.

Encryption/Decryption: Each FS node implements 256bit AES-GCM/XTS cryptographic algorithms, the best available, at a sustained throughput of 250 GB/s. Data is encrypted in flight and at rest to protect within-system and cross-system data traffic, with no effect on performance.

Network Erasure Coding (EC): The FSC employs a modern approach to data durability, leveraging a more cost-efficient erasure coding technique. Erasure coding is not used pervasively today because it is a compute-intensive technique that leverages Galois Field Arithmetic. This technique overburdens CPU-based storage systems that already struggle to cope with other data services. The FSC EC is implemented across FS nodes. Cross-node EC protects against not only SSD and processor failures, but also against node, TOR, and rack failures, all without costly, complicated, and failure-prone high availability nodes. With this level of performance, EC can be implemented not just on cold data, but for the first time, on hot and warm data as well. FSC supports 2+1, 4+2, and 8+2 erasure coding schemes.

Replication: For smaller clusters, the FSC offers two-way replication with excellent performance, a significant improvement compared to other CPU based storage arrays.

SECURE MULTI-TENANCY WITH FINEGRAINED QOS

Root of Trust: The FS node incorporates a built-in secure boot processor (SBP) that executes boot code from an embedded ROM and serves as an immutable root of trust. Using a physically unclonable function (PUF)-based authentication method, the secure boot process provides a physically-defined "digital fingerprint" that serves as a unique identifier for every DPU that is fabricated. This feature protects the system against supply chain attacks, ensuring the authenticity of the firmware.

Encryption: The FSC allows encryption to be set on a per-volume basis, supporting true, secure multi-tenancy. Each FS node stores keys in a secure enclave served from a Key Management Server (KMS) whenever there is an IO operation. In addition, the FSC supports Bring Your Own Key (BYOK) and Keep Your Own Key (KYOK) cloud requirements for KMIP (v1.4 spec) compliant KMSs.

USABILITY AND SIMPLICITY REDEFINED

Designed for ease of adoption, the FSC fully supports industry-standard protocols, ensuring smooth deployments in today's data centers. Other features that embody usability and simplicity include:

API-First Platform with an Intent API: The FSC abstracts away implementation details, allowing organizations to focus on high-level business needs. For example, the Composer continuously monitors the health, capacity, utilization, and wear of the FSC, in addition to other critical metrics. This approach allows IT administrators to focus on making provisioning decisions and ensuring that the health of the cluster is constantly maintained. The FSC also supports standard storage orchestration APIs such as the Container Storage Interface (CSI).

Zero-Touch Provisioning (ZTP): The FSC supports ZTP, which greatly simplifies the process of adding nodes to a cluster. After a node is initialized, the management and data plane networking are configured by using ZTP. Data can be rebalanced at volume creation time or performed automatically via a scheduled redistribution batch operation. This feature makes the FSC a future-proof solution, allowing new storage nodes to be incrementally added with no downtime.

Centralized Control for Ease-of-Management: Similar to SDN control over networks, The Fungible Composer supports the centralized control of the storage cluster, allowing IT administrators to better deploy and maintain policies, manage operations, and monitor and act on security breaches.

¹ 4K random read on raw volumes.

POWER-EFFICIENT GREEN TECHNOLOGY

The FS1600 node's power consumption with 24 SSDs is 750 W (typical) and 850 W (peak). FSC is the most power-efficient storage cluster in the market today.

ECONOMICS REIMAGINED—UNLOCKING NEW POSSIBILITIES

The breakthrough performance of the FSC provides economic savings. For example, the FSC employs erasure coding to deliver data durability and resiliency. Erasure coding offers the same level of durability as replication but requires a much lower capacity overhead. The FSC also achieves media savings by (1) allowing storage to be pooled, improving capacity utilization (2) enabling parallel accesses to high demand content, obviating the need for duplication; and (3) attaining better data reduction factors and throughput with hardware accelerated compression. With these benefits, the Fungible solution offers 5x media savings compared to Hyperconverged-DAS solutions.

Other solutions require many processors in a given form factor to deliver performance, making them expensive and power-hungry. In contrast, the FS node delivers the highest "processor" efficiency with 6.5M IOPS or 38 GB/s per DPU. The FSC delivers 96.5% Performance Efficiency Percentage (PEP)².

PRODUCT USE CASES

Cloud-native storage for Hyper Disaggregation: The FSC offers cloud providers a compelling alternative to conventional storage. By disaggregating storage, the FSC enables independent scaling of compute and storage, increased utilization, reduced server SKU, reduced management complexity, and increased agility. Its performance characteristics are well suited for performance-hungry and latency-sensitive applications such as high performance databases, parallel file systems, AI/ML, and analytics, while its performance density offers consolidation, flexibility, and cost efficiencies.


The FSC also delivers, without performance penalties, efficient data durability, data reduction and security features that are required for reliable, secure multi-tenant environments of any size. Service providers can now deploy true elastic storage that can dynamically adapt to performance and capacity changes.

Artificial Intelligence / Machine Learning: Modern AI/ML workloads typically require massive parallelism in performance, low latency, and large capacity. The FSC, combined with highly scalable parallel file systems, eliminates storage bottlenecks to achieve unprecedented performance, latency, and efficiency for these modern workloads.

Cloud-native high performance databases: Many of today's high performance scale-out databases deploy DAS to meet latency requirements. These databases typically offer durability through clustered redundancy schemes such as replica sets or primary-secondary configurations. If a server fails, data is preserved on another server. The FSC preserves DAS-like latencies while offering improved storage utilization and clustered-redundancy, but at a lower capacity overhead. The FSC in-line line rate data reduction capabilities also significantly reduce capacity requirements.

Analytics: Data-centric organizations consider analytics a central component of their business strategy. Big Data platforms aim to offer quick and meaningful insights, but struggle to cope with the surge in data volume and the demand for real-time response. The FSC addresses the requirements of modern analytics by delivering a storage cluster that is highly performant and massively parallel, offering real-time responses and insights. With the flexibility of the FSC, organizations can add additional storage to new or existing analytics clusters without disrupting operations.

The Fungible Storage Cluster delivers a no-compromise high performance, secure storage solution that meets the most demanding application requirements, supports business agility and significantly improves ROI. To learn more about the Fungible Storage Cluster, visit the FSC [webpage](#) or contact sales@fungible.com.



FS1600 Node

PHYSICAL

FS1600 Node

- Support for up to 24x U.2 NVME SSD bays, with raw capacity options of 92 and 184 TB per FS1600 node
 - 3.84 TB high performance SSD: 92 TB raw capacity per FS1600 node
 - 7.68 TB high performance SSD: 184 TB raw capacity per FS1600 node
- H x W x D: 3.5" (88 mm) x 17.4" (440 mm) x 30.6" (776 mm)
- Weight: 43.4 lb (20 Kg) without SSDs
- Rack unit: 2 RU
- Power input/supply: (Redundant) 2 x 1500 W AC to DC power-factor-corrected (PFC), Universal Input Voltage Range: 90–246 VAC
- Power usage: 750 W/node (typical), 850 W/node (peak)
- Network: 12 x 100 Gb/s Ethernet ports
- Management: 1 x 1/10 Gb/s Ethernet port, 1 x 1Gb/s Ethernet IPMI/BMC port

Composer node

- H x W x D: 1.7" (43 mm) x 17.2" (437 mm) x 28.5" (724 mm)
- Weight: Net: 26 lb (11.8 kg) / gross: 41 lb (18.6 kg)
- Rack unit: 1 RU
- Power input/supply: (redundant) 2 x 750 W AC to 100–127 VAC, 50–60 Hz; 20–240 VAC, 50–60 Hz
- Power usage: 700 W (peak)
- Network: 4 x 1 G (4 x RJ-45 GbE ports), 2 x 10/25G NIC



² Learn more about [Performance Efficiency Percentage](#).