SOLUTION BRIEF

# CASSANDRA DB: SLASH YOUR TCO BY CONSOLIDATING "DAS" ON THE FUNGIBLE STORAGE CLUSTER

## All the benefits of disaggregated storage at better than DAS performance

To meet the insatiable demands of real-time, high consumption-rate applications, organizations are employing NoSQL databases such as Apache Cassandra. However, with datasets growing in an unpredictable fashion, demanding more real-time processing, inefficiencies that may be tolerable in the past are now magnified a thousand-fold. With database operations running into the millions, compute and storage resources need to keep up. Specifically for Cassandra, the speed, scalability and availability of the solution is prioritized over consistency of data. The process of validating data for consistency slows down read operations. Thus, storage solutions must not only provide consistent reliable performance, but also increased performance to keep up with the performance and latency demands. Further, infrastructure provisioning needs to be agile to satisfy the continual and dynamic rate of change in requirements.

## THE CHALLENGES OF DIRECT-ATTACHED STORAGE (DAS)

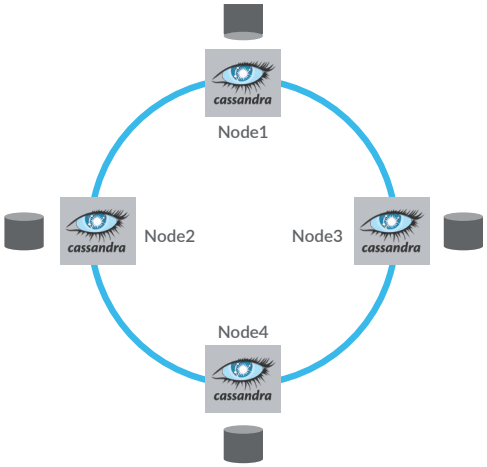


Figure 1. Traditional Apache Cassandra Deployment Utilizing Direct-Attached Storage (DAS)

While Apache Cassandra permits the latitude of deploying data center resources in a scale-out manner, it is costly as regards to storage and storage management. Today, storage for Cassandra is typically employed as direct-attached storage (DAS) in discrete servers. While DAS has served Cassandra well in the past, the winds of change in the data-centric era are exposing the flaws that IT leaders must address.

First, DAS storage cannot be shared efficiently since it is stranded in discrete servers and, as a result, commonly reduces storage utilization to levels as low as 30%. In addition, storage sizing determinations must be made at the time of infrastructure acquisition. Thus, the decisions pertaining to the storage capacity in each server must be finalized before the requirements of each Cassandra-based application are fully understood. This commonly leads to rigidity and progressively bloated expenses. Along the lines of shared storage, it is not uncommon for an organization to have multiple clusters of Cassandra that are managed independently. The fragmented usage of resources lead to underutilization of resources, specifically storage.

When scaling for either performance or capacity reasons, more server nodes need to be deployed to accommodate more directattached SSDs, effectively expanding the infrastructure unnecessarily.

While some Cassandra applications may offer built-in data protection mechanisms, the inherent Cassandra failure protection capability relies on basic data redundancy or multiple data replicas on server nodes. This causes greater storage inefficiencies and compounds the impact on data center infrastructure cost.

## THE FUNGIBLE STORAGE CLUSTER

Powered by the Fungible DPU, the Fungible Storage Cluster (FSC) is the highest performance, secure, scale-out disaggregated all-flash storage platform in the market today. The FSC comprises a cluster of two or more Fungible FS1600 storage target nodes and three Fungible Composer nodes. The control plane is managed by the Fungible Composer software, a centralized management solution that configures, manages, orchestrates, controls and deploys the Fungible Storage Cluster.

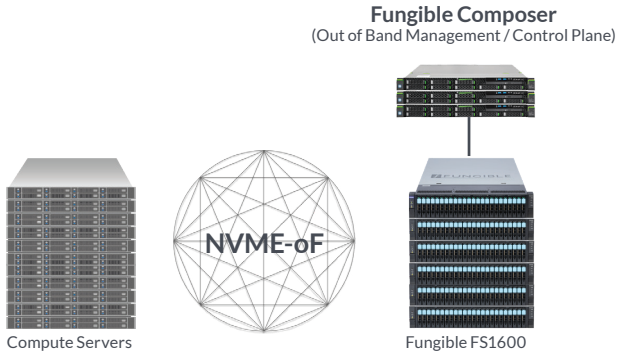


## THE SOLUTION - FITS LIKE A GLOVE FOR CASSANDRA STORAGE

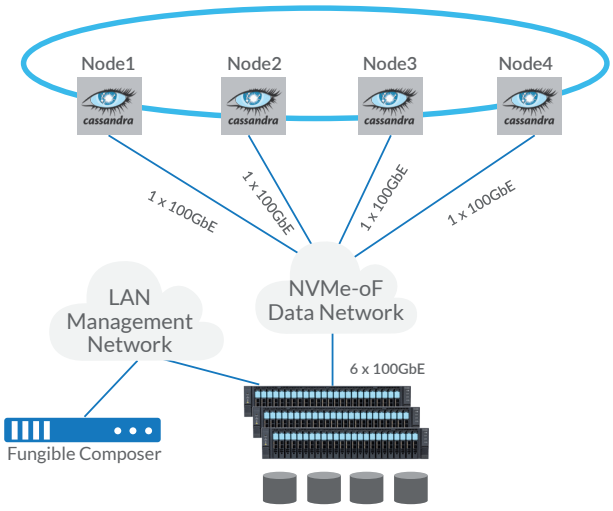


Figure 2. Apache Cassandra Deployment Using Fungible Storage Cluster Disaggregated Storage

The Fungible Storage Cluster (FSC) offers unprecedented storage performance and density that allows customers to deploy large-scale disaggregated flash storage for Cassandra DB. Consequently, data center administrators can now utilize logically shared storage to replace direct attached storage for cloud-level Cassandra environments.

The FSC is connected to Cassandra nodes via NVMe over Fabrics (NVMe-oF) Ethernet networks. Storage management functions such as volume creation are performed using the Fungible Composer. FSC supports storage volumes that are ephemeral or durable. Ephemera storage is great for Cassandra environments such as test/dev, staging etc. The FSC offers the world's fastest performance for storage, with each FS1600 node delivering 15MIOPS of 4K random read IOPS. The FSC also supports durable storage. Network erasure coding protects FSC data against drive, processor, system and rack failures and eliminates any single point of failure. Cassandra nodes can handle significantly higher levels of database workloads because the computationally intensive tasks of erasure coding, compression and encryption have all been offloaded to the FSC. FSC storage systems function together as massive, resilient, disaggregated pools of storage, and are linearly scalable in increments of 2RU/24 NVMe SSD nodes up to exabyte scale.

The FSC cluster provides a pool of highly available ephemeral or erasure coded, compressed and encrypted NVMe storage volumes that are shared across the Cassandra cluster. The FSC supports the diverse performance, resiliency and isolation needs of the full range of Cassandra bare metal, virtualized and containerized tenant applications. FSC permits scaling of storage capacity independent of compute capabilities to meet the challenges of the agile nature of Cassandra cloud-native applications. As a result, it enables extreme application scalability without causing resource over-provisioning or performance degradation.

## KEY BENEFITS

**Unmatched performance**. FSC delivers DAS-like NVMe performance without capacity and scale limitations to Cassandra applications running on any server. A single FSC provides up to 15 million IOPS of raw performance (4K read) at 60GB/sec by leveraging the performance of disaggregated NVMe storage supporting latencies that are within 10us of DAS using NVMe-oF/NVMe-TCP.

**Supports Modern Cassandra Deployment Methodologies**. The FSS supports bare metal, virtualized or containerized environments, offering both raw and durable block capabilities. It also includes a Container Storage Interface (CSI) Plugin for Kubernetes clusters.

**Cloud-scale Elasticity**. Organizations can provision storage on demand for Cassandra including adding linear scaling of capacity and/or performance. Any amount of storage can be dynamically allocated to any server on the network at any time from the FSC storage pool. And because compute and storage resources can be scaled independently, Capex inefficiencies can be significantly reduced.

**Hyper-scalability.** FSC supports scalability to an extremely large number of Cassandra workloads. The FSC can support thousands of servers running Cassandra instances with fine-grained QoS for IOPS and throughput scaling.

**Footprint and power efficiencies.** Real estate and power are important considerations for Cassandra deployments, whether at the edge or on-premises. The performance and density of FSC enables extremely large storage volumes to be consolidated into a small set of FSC storage systems. In addition, the FSC is the most power-efficient storage system on the market.

**Optimal use of compute resources.** FSC removes the sub-optimal use of server CPUs for running storage services such as data protection, compression and encryption, allowing data center administrators to ideally leverage CPU cores for Cassandra applications.

**Highly efficient data durability and management.** The FSC supports both within-system and across-system erasure coding for both hot and cold data at line rate, eliminating the need for inefficient replication in Cassandra DAS deployments. The FSC also supports snapshots and clones.

**Optimal data availability.** While storage is one component of the overall infrastructure, the FSC removes the need for data migration when a Cassandra DAS server is down. With a FSC disaggregated storage platform, data is not trapped in the DAS server and availability is handled at the storage system level. In addition, unlike the case of DAS, the FSC removes the need for data migration in the case of a planned or unplanned migration of a Cassandra-based virtual machine.

**Support for a highly secure, multi-tenant environment.** The FSC is architected with a secure hardware root of trust and supports data encryption, as well as per volume granularity for superior multi-tenant QoS. It also includes support for Active Directory user authentication, authorization and auditing.

## CONCLUSION

Fungible FSC provides scale-out Cassandra deployments with a new alternative to conventional storage that addresses severe infrastructure limitations to deliver flexibility, performance and value to unprecedented levels. The power, simplicity and density offered by the FSC for scale-out Cassandra applications, increases agility and lowers TCO. With the FSC, Cassandra cloud administrators can now convert these benefits to delivering the highest quality of service to their customer base.

## NEXT STEPS

For additional information and demo, contact sales@fungible.com. See more detailed performance results in this whitepaper.

## ABOUT FUNGIBLE

Silicon Valley-based Fungible is reimagining the performance, reliability and economics of today's data centers.

## CONTACT US

sales@fungible.com

FUNGIBLE, INC.

3201 Scott Blvd., Santa Clara, CA 95054, USA
669-292-5522

www.fungible.com