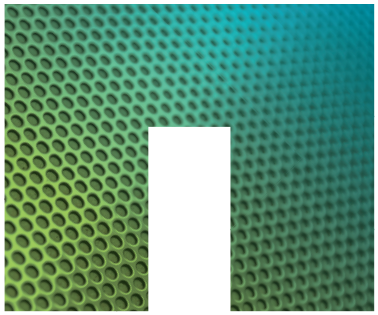

Solution Brief

NetApp Solutions for Cassandra

Get massive scalability, high performance, and continuous uptime—simply and securely

KEY BENEFITS

Scale Easily

- Add capacity with ease across clusters, data centers, and clouds
- Dynamically scale storage resources to meet business demands
- Scale-out without limitations

Accelerate Performance

- Deliver the lowest latency performance
- Dynamically adjust storage performance as needed
- Flexibly cluster flash and disk nodes to optimize performance and capacity

Maximize Availability and Data Protection

- Maximize uptime with >99.999% availability
- Encrypt data at rest
- Leverage unified data protection features

The Challenge

Modern applications operate in social, mobile, big data analytics, and cloud—i.e., third platform—IT environments that make complex demands on your data management. Apache Cassandra is a massively scalable, open-source, NoSQL database that offers continuous availability, linearly scalable performance, operational simplicity, and easy data distribution across multiple data centers and with cloud availability. Maintaining the highest levels of scalability, performance, and uptime for Cassandra workloads is critical.

The Solution

Scale your applications easily, while maintaining maximum performance and uptime, by deploying the Apache Cassandra NoSQL database on your choice of NetApp® data storage systems:

- On E-Series and EF-Series storage systems, including the all-flash E2800 and EF560, for submillisecond-latency workloads requiring exceptional read and write performance
- On FAS and All Flash FAS systems for moving high-performance application data securely across data centers and from flash to disk to cloud
- On SolidFire all-flash arrays for scale-out deployments with guaranteed application performance and capacity on demand

Scale Easily

Today's NoSQL database storage must keep up with continuous growth and meet the most demanding capacity requirements. The E-Series is purpose-built for capacity-intensive environments that require optimal space utilization and reduced power and cooling requirements. EF-Series all-flash storage array, supports for up to 384TB of raw capacity in modular 2U building blocks providing more data access with low-latency performance.

NetApp FAS and AFF makes it easy to optimize and accelerate your Cassandra environment as performance and capacity requirements change. Scale up by adding capacity, adding flash acceleration, and upgrading controllers. Scale out by growing from two nodes up to a 24 node cluster with 138PB of capacity, including combinations of different FAS and AFF models. With nondisruptive addition and replacement of storage systems and components, scaling occurs without maintenance windows or the challenge of coordinating downtime across teams.

With SolidFire you can linearly scale storage capacity and performance resources to meet the demands of rapid application data growth. The scale-out design complements the architecture of NoSQL datastores, and rapid volume cloning capabilities unlocks new deployment and scaling strategies.

NetApp®

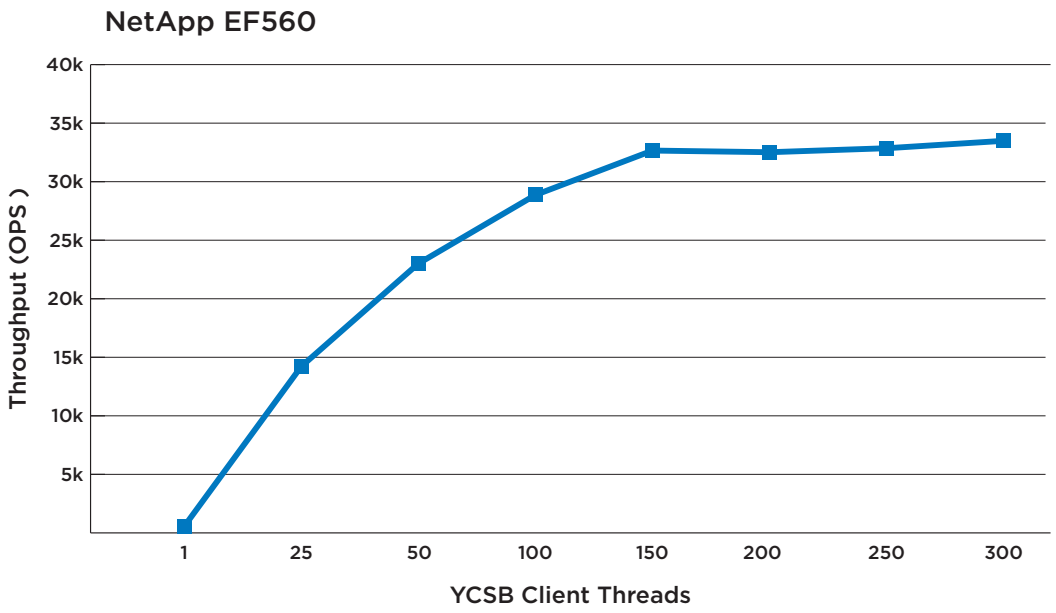

Figure 1) NetApp EF560 Cassandra 8-node cluster performance scaling, 95% reads, 5% updates.

Accelerate Performance

NetApp E-Series storage systems deliver up to 825,000 input/output operations per second (IOPS) and microsecond latency to help you complete operations more quickly. E-Series systems also offer bandwidth—12GB per second sustained to disk—to support the demanding performance and capacity needs of a Cassandra environment. Balanced performance features help support the requirements of different workloads. Dynamic Disk Pools (DDPs) dynamically rebalance data across all drives in the pool when new drives are added or old drives are removed, eliminating hot spots. The whole process is transparent to users, who experience little to no impact on performance. If a drive fails, DDPs help ensure that high performance within the Cassandra environment is maintained.

NetApp All Flash FAS systems with NetApp ONTAP® FlashEssentials enables up to 4 million IOPS with 24 nodes and meets the demands of Cassandra on the Data Fabric enabled by NetApp. FlashEssentials is what is behind the performance and efficiency of All Flash FAS. It encapsulates flash innovations and optimization technologies in ONTAP software. And with the ONTAP 9 release, performance can be further increased by up to 60%.

SolidFire systems allows you to dynamically adjust storage performance as needed to ensure your Cassandra environment delivers optimal throughput with minimal latency. Adjust database volume performance on the fly and control performance (IOPS) independent of capacity (GBs).

Maximize Availability and Data Protection

NetApp E-Series systems are built on a long heritage of serving diverse workloads to provide enterprise-class reliability. A fully redundant, building-block system with automated failover and advanced monitoring maximizes uptime. NetApp SANtricity® full disk encryption combines local key management with drive-level encryption for comprehensive security for data at rest that doesn't sacrifice performance or ease of use. SANtricity also supports FIPS-certified hard drives for security-sensitive customers.

NetApp FAS and AFF systems are designed to deliver high availability through a comprehensive approach that combines highly reliable hardware, innovative software, and sophisticated service analytics. With FAS and AFF systems, you can leverage the data fabric to maintain control of your data as you grow—on premises and in the cloud.. For simplified, standardized backup, restore, and disaster recovery, the Snap Creator framework offers unified data protection for physical, virtual, and cloud environments. This centralized solution for backing up critical data offers policy-based automation and integrates with existing application architectures to deliver data consistency and lower operating costs, as well as the ability to backup and restore point in time copies of an entire cluster, make copies for test/dev, disaster recovery with ONTAP Cloud and hybrid cloud deployments with NetApp Private Storage for Cloud.

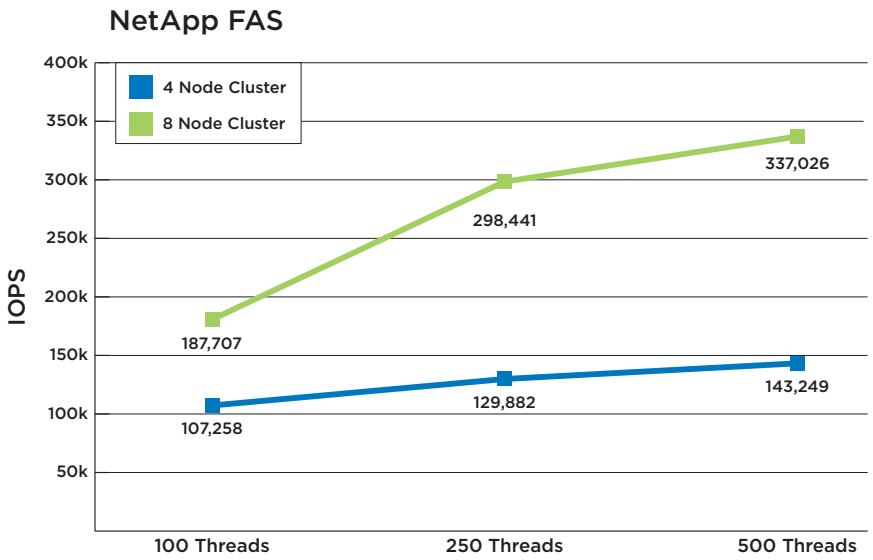

Figure 2) NetApp FAS Cassandra multi-node performance scaling, 100% reads.

Consolidate with confidence using SolidFire's Helix data protection—assurance for protecting your most critical information. Ensure complete data and performance availability regardless of system condition or application activity. Safeguard data with 256-bit encryption, and protect data with both synchronous and asynchronous replication, along with integrated backup/recovery.

About Cassandra

Apache Cassandra is a top-level Apache project born at Facebook and built on Amazon's Dynamo and Google's BigTable. Cassandra is the right choice when you need scalability and high availability without compromising performance. Linear scalability and proven fault-tolerance on commodity hardware or cloud infrastructure make it the perfect platform for mission-critical data. Cassandra's support for replicating across multiple datacenters is best-in-class, providing lower latency for your users and the peace of mind of knowing that you can survive regional outages. For more information, visit cassandra.apache.org.

About NetApp

Leading organizations worldwide count on NetApp for software, systems and services to manage and store their data. Customers value our teamwork, expertise and passion for helping them succeed now and into the future.

www.netapp.com