ALTOROS™

Evaluation:

NetApp E-Series for Couchbase Deployments

This 11-page paper explores performance, scalability, and availability of the NetApp E-Series data storage and management system deployed with Couchbase Server.

Authors:

	Vladimir Starostenkov	R&D Engineer
	Alex Khizhnyak	Director of Tech Communications

Table of Contents

1. Abstract

This paper addresses the performance, scalability, reliability, and availability advantages of NetApp E-Series storage systems compared to commodity servers with internal direct-attached storage (DAS), while running Couchbase Server NoSQL database.

NetApp engaged with Altoros, a leading NoSQL consultancy, to evaluate, identify and validate how NetApp E-Series storage systems can better address the challenges encountered when deploying Couchbase Server at scale. The team evaluated Couchbase Server 3.0 running on an EF560 Flash Array and commodity servers with internal direct-attached storage (DAS). Test results cover Couchbase Server performance and reliability characteristics under using a number of workloads generated by a client-like benchmarking application.

2. Couchbase Server Overview

NoSQL technology was pioneered by leading internet companies to overcome the limitations of relational database technology for use with modern web applications. Today, enterprises are adopting NoSQL for a growing number of use cases, a choice that is driven by four interrelated megatrends: Big Users, Big Data, the Internet of Things, and Cloud Computing.

Much of this emerging Big Data is unstructured, and thus not suitable for loading and management by traditional relational database management systems (RDBMS) and their associated structured query language or SQL.

Couchbase Server is best described as an open-source, distributed NoSQL document-oriented database. It is being adopted by organizations that need a flexible data model to handle the new varieties of data, ease of scalability for data growth, consistent high-performance, and always-on 24x365 availability. As enterprises start expanding their Couchbase clusters, they may encounter deployment challenges maximizing performance, scalability, reliability, and availability.

3. NetApp E-Series Overview

Increasingly, organizations look for ways to drive greater speed and responsiveness from the applications that control their key business operations. Because the performance of these applications is tightly linked to time to market, revenue, and customer satisfaction, it is critical that they operate at maximum efficiency.

To achieve maximum performance, organizations have had to unnecessarily use partially filled hard disks to meet the required performance. This overprovisioning of hard disks wastes disk capacity and data center space and drives up power consumption. To eliminate overprovisioning and maximize return on investment from high-performance applications, companies now look to all-flash systems. However, as buyers consider these new systems for their tier 1 applications, they are often challenged to find an offering that is also enterprise proven and highly reliable.

The NetApp EF560 all-flash array is an all-SSD storage system designed for applications demanding the highest levels of performance, reliability, and availability. Requiring just 2U of rack space, the EF560 all-flash array combines extreme IOPS, ultralow response times, and up to 12 GB/sec of bandwidth with leading enterprise-proven availability features, including:

- Redundant components with automated failover
- Intuitive storage management with comprehensive tuning functions
- Advanced monitoring and diagnostics with proactive repair

Designed specifically for high-speed transactional applications that demand high IOPS and consistent low latency, the EF560 all-flash array delivers over 650,000 sustained IOPS and response times measured in microseconds. Bandwidth-oriented workloads also benefit from the EF560's capability to provide up to 12 GB/sec of throughput.

It is also important to note, that EF560 all-flash array continues NetApp's long heritage—leveraging expertise gained from nearly one million NetApp E-Series SANtricity-powered systems—of delivering powerful solutions to meet business needs.

4. Evaluation Methodology

Database performance can be defined by the speed at which a database will compute basic CRUD (create, read, update, delete) operations. In all Altoros benchmark tests, a basic operation is an action performed by the YCSB (Yahoo Cloud Serving Benchmark) workload executor, which drives multiple client threads. Each thread executes a sequential series of operations by making calls to the database interface layer to run the workload.

Below is a representation of the Couchbase Server Lab Infrastructure used by Altoros to execute the benchmark tests. The benchmark tests executed by Altoros scaled from 4–8 Couchbase Servers nodes.

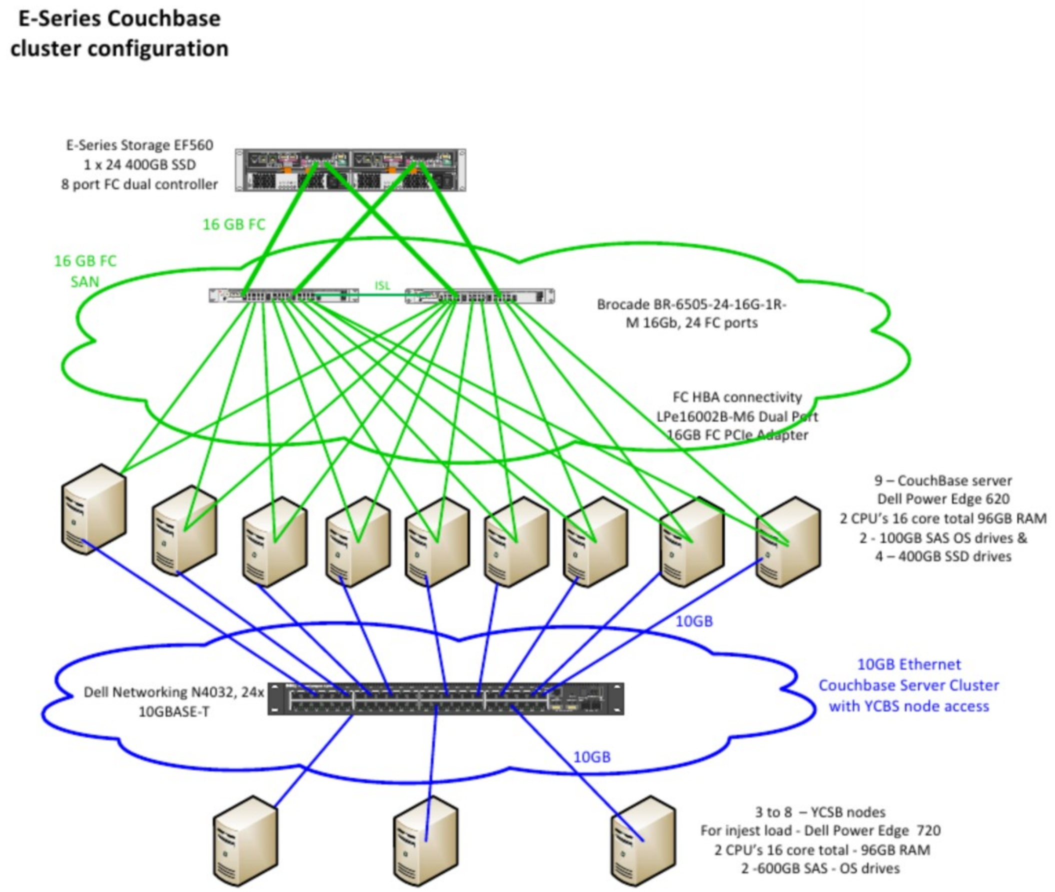

Figure 1. Single horizontal scaling unit of Couchbase Server Deployment over EF560

5. Results

Altoros's team ran multiple benchmark tests and found several advantages in deploying Couchbase Server over the NetApp EF-Series storage system compared to the traditional deployment of commodity server with direct-attached storage (DAS). The significance of these results becomes critical when enterprises need to deploy Couchbase at scale.

- **Consistent Performance Even During Failure**
  - Dynamic Disk Pools (DDP) in the EF-Series minimize the performance impact of a drive failure and can return the system to optimal condition up to eight times faster than traditional RAID found in commodity servers.
  - Shorter rebuild times significantly reduce exposure to multiple disk failures, offering a level of data protection that can't be achieved with traditional RAID.

- **High Availability and Enterprise Reliability**
  - EF560 all-flash array has fully redundant I/O paths with automated failover, extensive diagnostic and monitoring maximizes solution uptime.
  - All management tasks can be performed while the storage remains online with complete read and write access, enabling administrators to make configuration changes, perform maintenance, or expand storage capacity without impacting availability.
- **Flexibility and Ease of Use**
  - Simplified setup—DDP evenly distributes data, protection information, and spare capacity across the entire pool of drives increasing storage efficiency and utilization.
  - Flexibility for different requirement and workloads—the E-Series introduces the ability to mix different drive types with different storage capacities.
- **Online Scaling**
  - Decouples the Compute Node Layer from the Storage Layer, providing tremendous flexibility to scale each layer independently.
- **Consolidated Storage Management**
  - SANtricity can manage multiple E-Series systems. It provides a centralized management GUI for monitoring and managing important hardware events including drive failures and replacements. Manageability is also enhanced with predictive drive failure, data assurance, and advanced performance monitoring.
- **Data Encryption**
  - SANtricity full-disk encryption combines local key management with drive-level encryption for comprehensive security for data at rest with no impact on performance.

Couchbase Server deployment over NetApp E-Series

The storage system is shared across the Couchbase servers. NetApp E-Series supports a variety of RAID(s) like RAID 0, 1, 5, 6, 10, or NetApp Dynamic Disk Pool (DDP).

<table>
<tr><th colspan="4">Couchbase Server Deployment</th></tr>
<tr><th colspan="2">Direct-attached Storage</th><th colspan="2">E-Series Storage</th></tr>
<tr><th>RAID 0</th><th>RAID 5</th><th>RAID 5</th><th>DDP</th></tr>
<tr><td colspan="4">Feature: Performance</td></tr>
<tr><td colspan="4">Efficiency under disk I/O heavy workload (performance per used SSD drives)</td></tr>
<tr><td>Low</td><td>Low</td><td>High**</td><td>High</td></tr>
<tr><td colspan="4">Feature: Scalability</td></tr>
<tr><td colspan="4">Load scalability
(ease of expanding and contracting disk drive pool to accommodate heavier or lighter workloads)</td></tr>
<tr><td colspan="2">Low:
• Storage system requires homogeneous scalability—all the used disk drives must be equal in capacity and performance.
• Node restart could be required to scale storage capacity—time-consuming process.</td><td colspan="2">High:
• Heterogeneous scalability—disk drives of different capacities and types could be combined.
• Storage scalability on the fly.</td></tr>
<tr><td colspan="4">Growth planning and Flexibility</td></tr>
<tr><td colspan="2">Low. Complex and expensive process of changing the compute-storage balance. Once the system is planned various restrictions are applied:
• Number of disk drives to add or remove must be multiple of cluster size.***
• Number of hot spare drives must be multiple of cluster size.***</td><td colspan="2">High. E-Series decouples storage layer from compute layer providing tremendous flexibility to scale each layer independently.
• Arbitrary number of disk drives could be added or removed.
• Hot spare drives are shared across several cluster nodes—higher storage utilization.
• Single point of disk load monitoring.</td></tr>
<tr><td colspan="4">Feature: Availability</td></tr>
<tr><td colspan="4">Storage controller failure: Fault Tolerance</td></tr>
<tr><td colspan="2">Low. Partial service downtime for >30 seconds. Couchbase Server rebalance required.</td><td colspan="2">High. Fully redundant system (Redundant, hot-swappable storage controllers, disk drives, power supplies, and fans), partial service downtime for <2 seconds.</td></tr>
</table>

Disk drive failure: Fault Tolerance summary			
Zero: *Disk failure leads to node failover and cluster rebalance. Partial service downtime for >30 seconds.*	***Moderate***: *Disk failure leads to a significant performance degradation and time-consuming data reconstruction. No service downtime.*	***High***: *Negligible performance impact. No service downtime.*	***Highest***: *Tunable performance impact (against time to repair). No service downtime.*
Disk drive failure: time to repair			
Unrepairable. *Software level guarantees for data safety.*	***High***: *>10 hours to reconstruct 400 GB SSD.*	***Low***: *~1.5 hours to reconstruct 800 GB SSD.*	***Lowest***: *15 minutes to reconstruct 400 GB SSD. Tunable reconstruction time (against performance impact).*
Disk drive failure: system recovery automation			
Software layer only. Couchbase Server automated or manual failover.	***Low***: *Failed drive must be manually replaced or be backed by hot spare drive for each node.*	***High***: *Small number of hot spares is shared across several nodes. No manual intervention required.*	***Highest***: *Fully-automated. DDP evenly distributes data, protection information, and spare capacity across the EF560's entire pool of drives. Up to 2 disk drive failures could be handled by the storage system simultaneously.*
Additional Features			
Disk drives health monitoring			
Separate solution required.		*Built-in SANtricity proactive drive health monitoring identifies problem drives before they create issues.*	
Data encryption			
Separate proprietary solution required.		*SANtricity full-disk encryption.*	
(**) Expected on higher scale. DDP requires minimum of 11 drives in disk pool, but we were able to fulfill Couchbase Server storage I/O requirements with 8 drives in EF560 RAID 5. (***) All the nodes must be equal in underlying RAID configuration.			

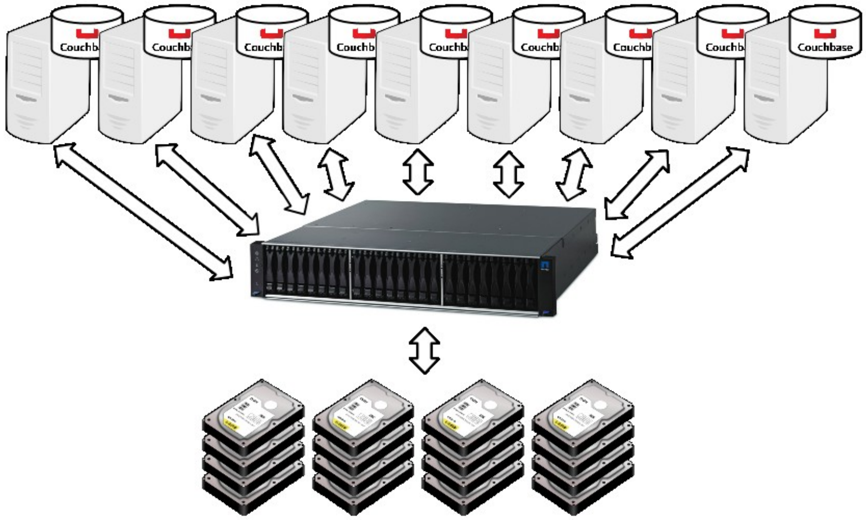

Figure 2. Couchbase Server deployment over NetApp E-Series

6. Conclusions

6.1 Performance

Our performance benchmarking evaluated workloads with the most intensive disk I/O and found significant price-performance advantages with the EF560 when compared to commodity servers with internal DAS. The benefits of using E-Series are even higher when Couchbase Server deployments are not capacity-intensive, allowing enterprise customers to safely deploy E-Series with fewer disks and still provide great performance. In every case, the EF560 solution proved to be more reliable and offered easier manageability.

6.2 Scalability

The standard RAID levels recommended for Couchbase server and implemented with traditional internal controllers for DAS do not allow administrators to combine different disk drive types and sizes, significantly reducing flexibility. Other limitations include the inability to add disk drives on the fly for RAID 10 and forcing basic RAID 10 configurations to double in size, at a time when increased storage is needed to accommodate data growth. It makes scaling storage a very challenging and costly task.

By decoupling storage from compute nodes, NetApp's Couchbase Server deployment, provides more flexibility in changing the compute-storage balance, and decreases cluster maintenance time and efforts. Using the SANtricity control application, the EF-series storage can easily add or remove disk drives while changing the nodes' volume size.

In contrast, using the traditional direct-attached storage drives results in more time consuming steps to change cluster capacity. Administrators must first exclude a node from the Couchbase Server cluster, then change the node's RAID configuration, before putting the node back and finally wait until the data is rebalanced on the software level and repeat the process for all the cluster nodes.

The rebalancing process can take days or even weeks. It is easier to just add more nodes to the Couchbase Server cluster, as this approach is recommended for storage scaling in the official documentation. This can potentially lead to unnecessary costs since added CPUs are being underutilized. The suggested Couchbase Server deployment over E-Series storage obviates this issue entirely as you can add storage capacity without needing to add compute nodes (CPUs).

The optimal number of compute nodes attached to a single disk array device varies, and generally depends on the particular type of production workload. A good starting point for a proof of concept deployment is 6–10 compute nodes per single storage array. In the current data-intensive world, successful POC clusters turn to production very quickly, especially with agile NoSQL workloads. For this use case, Couchbase Server deployment over E-Series provides superior flexibility to change the storage capacity and quality at runtime.

The EF560 is capable of handling up to 120 drives, fitting as many as 30 Couchbase Server nodes with relatively low disk I/O activity. During Couchbase Server tests on EF560, most of the frequently accessed data is cached by Couchbase Server in RAM while the whole data set is persisted inside the EF560. For additional fault-tolerance and cluster availability, enterprises could leverage 2 x EF560 in a rack-aware configuration supporting 20+ Couchbase Servers.

6.3 Availability

The EF560 has more powerful controllers and better reconstruction capabilities, resulting in significantly more uptime compared to the independent RAIDs used in direct-attached storage.

Even the standard RAID levels implemented by E-Series storage give the Couchbase Server administrator higher throughput and lower latency, better recovery in case of disk failures, much better availability and overall system resiliency. The combination of software level replication and reliable storage hardware is especially useful for deployments using Couchbase Server for mission-critical applications requiring High Availability.

NetApp's Dynamic Disk Pools showed unbeatable disk failure recovery time while minimizing the performance impact. Flexible disk pool sizing maximizes the available storage utilization and allows the Couchbase Server deployment to be backed by an arbitrary number of available disk drives.

More on NetApp E-Series: http://www.netapp.com/us/products/storage-systems/ef-series/index.aspx

7. About the Authors

Vladimir Starostenkov is a Senior R&D Engineer at Altoros. He is focused on implementing complex software architectures, including data-intensive systems and Hadoop-driven apps. Having background in computer science, Vladimir is passionate about artificial intelligence and machine learning algorithms. His NoSQL and Hadoop studies were published in NetworkWorld, CIO.com, and other industry media.

Alex Khizhnyak is Director of Tech Communications at Altoros and a co-founder of Belarus Java User Group. Since 1998, he has gained experience as a journalist, editor, IT blogger, and technical evangelist. Alex is mostly interested in big data, databases, data science, and high-load systems. His publications at TechRepublic, ebizQ, NetworkWorld, and DZone reached out to hundreds of thousands of readers. Follow him at @alxkh.

Altoros *brings Cloud Foundry-based "software factories" and NoSQL-driven "data lakes" into organizations through training, deployment, and integration. With 250+ employees across 9 countries in Europe and Americas, Altoros is the company behind some of the world's largest Cloud Foundry and NoSQL deployments. For more, please visit www.altoros.com.*

To download more NoSQL guides and tutorials:

- *check out our resources page,*
- *subscribe to the blog,*
- *or follow @altoros for daily updates.*