

# NetApp Plus Mosaic: Continuous Data Protection for MongoDB

## KEY BENEFITS

### Minimize Application Downtime

- Near-zero recovery point and recovery time objectives (RPOs and RTOs)
- Single-click and orchestrated recovery
- Granular collection-level recovery

### Increase Operational Efficiency

- Flexible policy-based management
- Recovery to different topology database clusters (testing and development use case)
- Automated failure handling

### Simplify Deployment

- On-premises or public cloud use
- Compatibility with NetApp® ONTAP® 9 software
- Performance and availability benefits of NetApp arrays

## THE DATA PROTECTION CHALLENGE

Rapid proliferation in social, mobile, cloud, and Internet of Things is driving enterprises to deploy hyperscale, distributed, data-centric applications such as customer analytics, e-commerce, security, surveillance, and business intelligence. To handle the data requirements of these high-volume, high-ingestion rate, and real-time applications, enterprises are rapidly adopting massively scalable and nonrelational databases such as MongoDB.

Like any business-critical application, these databases require data protection, including application-consistent backup; near-zero RPO and RTO; granular, repair-free recovery; failure handling; and backup storage efficiency. However, given the hyperscale, distributed nature of these databases, traditional backup and recovery products don't support these requirements, leaving a critical data protection gap.

## THE SOLUTION: APPLICATION-CENTRIC DATA PROTECTION FROM RUBRIK MOSAIC

NetApp and Rubrik have partnered to extend the NetApp Data Fabric by providing application-centric data protection for next-generation applications. Enterprises can now leverage the performance, reliability, and flexibility of ONTAP 9 with the continuous data protection capability of Rubrik Mosaic. With this solution, organizations can scale business-critical applications on MongoDB and be confident in the recoverability of data and the ability to maintain high application uptime. Mosaic is an industry-first scale-out data protection software-only product that is purpose-built for applications running on MongoDB.

### Scale-out architecture

Rubrik Mosaic is founded upon Consistent Orchestrated Distributed Recovery (CODR), Mosaic's cloud-first, scale-out data management architecture that enables customers to meet data protection requirements for MongoDB. CODR uses elastic compute services that can be auto-scaled based on load and removes the dependency on media servers. CODR also transfers data in parallel to and from file-based and object-based secondary storage for multiple use cases, including data protection and testing and development. To simplify the data recovery process and to avoid vendor lock-in, the protected data is stored in the database native format.

### Continuous backups

By using native application intelligence, Mosaic creates a true point-in-time consistent backup copy of MongoDB collections (both sharded and unsharded) at user-specified intervals, a concept called cluster-consistent versioning. Mosaic produces these cluster-consistent versions across all shards without quiescing the MongoDB database. Backups can be generated at a user-specified time interval and at any granularity (collection-level

or entire database), providing operational ease of use to database administrators. And with Mosaic, backup operations are resilient to failovers (primary switch) and failures (node and so on).

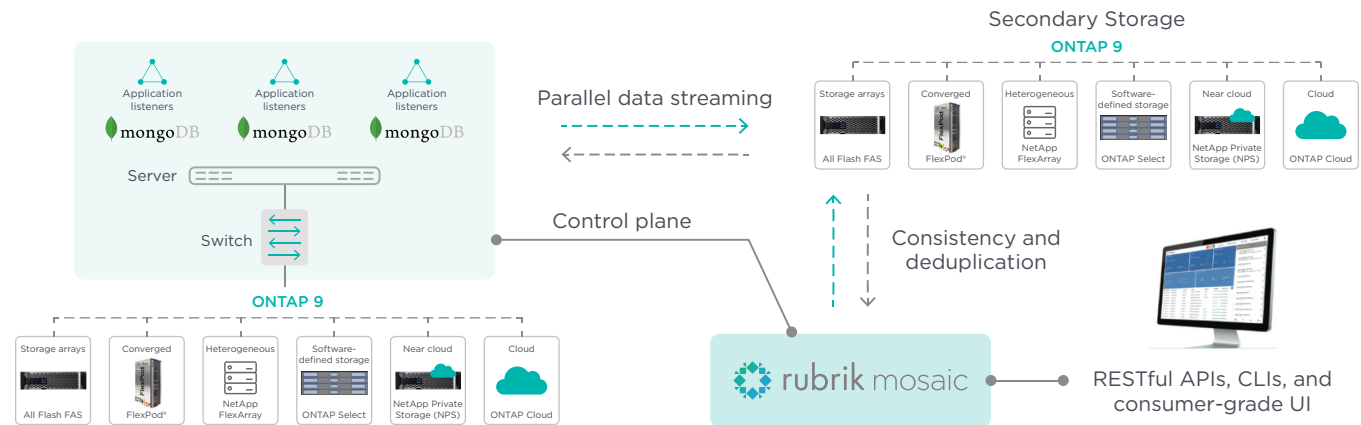


Figure 1 - Reference deployment of MongoDB with NetApp and Rubrik Mosaic.

### Fully orchestrated and granular recovery

Rubrik Mosaic provides fully orchestrated, any-point-in-time recovery. MongoDB collections can be recovered directly back into the same MongoDB database (operational recovery). They can also be recovered to a different MongoDB database instance (testing and development refresh) with a different topology (the number of nodes on the destination cluster differs from the node count of the source cluster).

Mosaic supports all combinations of recovery—sharded to unsharded, unsharded to sharded, sharded to sharded—thus reducing the operational burden of refreshing testing and development clusters for continuous development DevOps environments. Further, the recovery process deals only with the logical data, making it three times faster than with traditional approaches. During recovery, the data is directly transferred from secondary storage into target databases, resulting in a very low RTO.

### WHY NETAPP PLUS MOSAIC?

NetApp big data solutions deliver an open, scalable storage system for building big data applications. Customers gain business insights—and value—more quickly. The NetApp All Flash FAS system offers predictable and high performance with consistent and low latency, resulting in a very fast response time to the most

demanding applications that are deployed on MongoDB. And to deliver maximum uptime and high availability, it also offers non-disruptive operations and integrated data protection across applications, virtual infrastructures, and cloud architectures. Together with Rubrik Mosaic, NetApp offers an industry-leading enterprise data management solution for MongoDB that is based on ONTAP 9 and Rubrik.

Rubrik Mosaic is a software-only product that can be deployed on a physical server, a virtual machine, or any cloud compute instance (for example, Amazon EC2). It communicates with the MongoDB database through a Secure Shell (SSH) connection that forms a control plane to orchestrate data movement. The data can be backed up to a secondary FAS array or to a NetApp E-Series system. In addition to CLIs and RESTful APIs, customers can use the Mosaic consumer-grade UI to manage their data protection environment.

### ABOUT NETAPP

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

[www.netapp.com](http://www.netapp.com)



#### Global HQ

1001 Page Mill Rd., Building 2  
Palo Alto, CA 94304  
United States

1-844-4RUBRIK  
[inquiries@rubrik.com](mailto:inquiries@rubrik.com)  
[www.rubrik.com](http://www.rubrik.com)

Rubrik delivers a single platform to manage and protect data in the cloud, at the edge, and on-premises. Enterprises choose Rubrik's Cloud Data Management software to simplify backup and recovery, accelerate cloud adoption, and enable automation at scale. As organizations of all sizes adopt cloud-first policies, they rely on Rubrik's Polaris SaaS platform to unify data for security, governance, and compliance. For more information, visit [www.rubrik.com](http://www.rubrik.com) and follow [@rubrikinc](https://twitter.com/rubrikinc) on Twitter.

20190614\_v1