

# Next-Generation Backup and Recovery for Modern Databases in Virtualized Environments

Datos IO provides the industry's first cloud-scale, application-centric, data management platform enabling organizations to protect, mobilize, and monetize all their application data across private cloud, hybrid cloud and public cloud environments.

To learn more, visit [www.datos.io](http://www.datos.io)



## Key Benefits

### Application-Centric Backup and Recovery

- Application-consistent backup of non-relational NoSQL databases and big data file systems
- Single-click, fast recovery leading to reduced application downtime
- Industry-first semantic deduplication for 80% reduction in backup storage capacity

### Minimize Application Downtime

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

### Increase Storage Efficiency

- Industry-first semantic deduplication

### Simplify Deployment

- Software based solution easily deployed on virtual machines
- VMware Ready certified solution
- Complete API set for integration with VMware orchestration



## The Data Protection Challenge

Rapid proliferation in social, mobile, cloud, and Internet-of-Things are driving enterprises to deploy hyper-scale, distributed, data-centric applications on modern, NoSQL databases.

These applications and uses cases include analytics, e-commerce, security, surveillance, and business intelligence. Enterprises are rapidly adopting massively scalable and non-relational databases to handle the data requirements of these high-volume, high-ingestion-rate, and real-time applications. Likewise enterprise utilize VMware based Converged and Hyper-converged infrastructures to increase their agility in deploying scalable applications.

Like any business-critical application, these databases require robust data protection requirements that are standard for enterprise applications. These include application-consistent backup, near-zero RPO and RTO, granular, repair-free recovery, failure handling, backup storage efficiency, and software-only deployment for cloud-first environments. However, given the hyper-scale, distributed nature of these databases traditional backup and recovery products (are infrastructure centric or VM centric) don't support these requirements, leaving a critical data protection gap. Native VMware data protection methods don't deliver the tools or methods to meet these protection requirements for distributed and NoSQL databases.



## The Solution: Datas IO RecoverX for Application-Consistent Backup and Recovery

VMware has certified Datas IO RecoverX as a VMware Ready solution to provide application-centric data protection for cloud native applications and distributed NoSQL databases. Enterprises can now leverage the agility, flexibility and scalability of deploying cloud-native applications on VMware with the continuous data protection capability of Datas IO RecoverX. With this solution, enterprises can scale business-critical applications on a Cassandra and MongoDB databases and be confident in the recoverability of data and the ability to maintain high application uptime.

### Scale-Out Architecture

Datas IO RecoverX is based on Consistent Orchestrated Distributed Recovery (CODR), engine, Datas IO's cloud-first, scale-out data management architecture that enables customers to harness the cloud for both next-generation data protection and data management.

### Scalable Versioning

By using native application intelligence, RecoverX creates application-consistent point-in-time backups at user-specified intervals. Because the backups are application-consistent no repair is needed when a version is restored thus reducing RTO.

### Fully Orchestrated and Granular Recovery

RecoverX simplifies the operational burden of refreshing QA/Test/Dev in continuous-development DevOps environments. Further, the recovery process deals only with the logical data, making it 3 times faster than with traditional approaches. During recovery, data is directly transferred from secondary storage into target databases, an orchestrated media server-less architecture, resulting in very low RTO.

### Industry-First Semantic Deduplication

Semantic deduplication is an industry-first capability developed by Datas IO to reduce the cost of storing backups. As part of versioning Datas IO RecoverX deduplicates all the replicas of a primary dataset, thus greatly increasing storage savings without losing native formats.

**VMWare Deployment:** Datas IO RecoverX is a software-only data protection solution designed to be deployed cloud environments. Hosting RecoverX on VMware enables rapid and flexible deployment.

- Because memory, CPU and disk resources are virtual it is simple to provision and configure VMs with machine resources appropriate for a RecoverX deployment.
- Virtual networking makes it simple to assign networking to RecoverX and enable connectivity to the Data Sources it will be used to protect.
- VMware monitoring makes it easy to monitor resource use of RecoverX and plan for growth.

