

## Datos IO RecoverX

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 Features

- Elastic scale-out architecture
- Application-consistent backups
- Parallel streaming data movement
- Any point-in-time recovery
- Sub-table, query-able and incremental recovery
- Globally distributed metadata catalog
- Industry-first semantic deduplication
- Enterprise policy management
- Enterprise security



### Data Sources Supported

#### Non-relational Databases

- MongoDB
- Apache Cassandra
- DataStax Enterprise

#### Big Data Filesystems

- Apache HDFS
- Commercial Distributions (Cloudera, Hortonworks)



### Overview

There is a fundamental, once in a generation, technology shift occurring in IT. Modern applications deployed in public cloud environments and across multiple clouds (private, hybrid) are emerging as the de facto enterprise architecture. According to IDC, deployment of these applications will double in the next 24 months! This modern IT architecture is comprised of micro-services based applications being deployed on the elastic database, elastic compute and elastic storage services of cloud native environments. This shift, from the traditional monolithic data center centric IT stack to the modern IT stack is fundamentally disrupting the core tenets of data management, specifically backup and recovery, one of the most sacred pillars of IT.

Datos IO addresses this challenge with the industry's first cloud-scale, application-centric, data management platform, RecoverX, enabling organizations to protect, mobilize, and monetize their application data across private cloud, hybrid cloud and public cloud environments. RecoverX is built on top of our seminal data management architecture called Consistent Orchestrated Distributed Recovery (CODR™) engine, which is not dependent on media servers and transfers data in parallel directly between application sources and file-based or object-based secondary storage. The architecture is fully distributed in nature, resulting in high availability in failure scenarios and uses elastic compute resources for scalable performance and resiliency. Using this architecture, RecoverX

delivers application-aware data management that allows massive storage efficiency and query-able sub-table level recovery/mobility solutions at scale for traditional and modern applications. The key use cases that RecoverX addresses are: point-in-time backup & recovery, test/dev/QA refresh, cross-cloud data mobility, and archival/long-term retention — either in the cloud, on-premises/private clouds, or for hybrid cloud environments.



## Benefits

Unlike legacy solutions, RecoverX is built from the ground-up to address the requirements of cloud-first applications. This allows RecoverX to deliver massive benefits for enterprises, including:

- Software-only elastic data management architecture allows private cloud, public cloud and multi-cloud deployment
- Fully automated, application-consistent backups for fast backup and repair free recovery
- Flexible recovery including any point-in-time, and queryable recovery enables selective recovery based upon time or query
- Enterprise-level security enables LDAP authorization, Kerberos authentication, and TLS/SSL encryption
- Industry-first semantic deduplication results in 80-90% storage cost savings
- Native failure handling provides operational resiliency when there are node or database failures
- Application-awareness enables automated refresh of test/dev/QA environments and results in ~8x operational efficiency
- Application-awareness enables universal data mobility across private and public cloud



## Product Features

Datos IO RecoverX provides application consistent backup and orchestrated recovery for modern applications and databases deployed across hybrid cloud environments. Key Features:

- Datas IO RecoverX is an elastic cloud data management software platform
- Requires no dedicated media server hardware
- Provides 10x greater backup storage efficiency for cloud-native applications and workloads
- RecoverX supports modern data sources from databases to filesystems, to cloud infrastructures of the digital enterprise
- RecoverX software can be deployed on physical servers, virtual machines or any compute instance in the cloud as a single node or clustered (3 nodes or 5 nodes) configuration
- Clustered configuration provides high-availability resiliency and high-performance scalability to handle massive scale environments (100TB +)
- RecoverX performance also scales elastically with the underlying infrastructure (compute and memory) that is provided

Features	Cassandra Database	MongoDB Database	Hadoop Filesystem
Backup Granularity	Column Family	Collection	Directory
Backup Interval	15 mins and above		
Recovery Granularity	Sub-Column Family	Collection	File
Recovery Options	Fixed point-in-time Any point-in time Query-able	Fixed point-in-time Any point-in-time	Fixed point-in-time
Recovery to Different Topology	Yes; for test/dev use cases		
Cross-Cloud Mobility	Yes; for test/dev, compliance, cloning, migration use cases		
Secondary Storage	NFS, AWS S3, Google Cloud Storage		
Advanced Features	Compaction handling for storage efficiency TTL data handling for restore	Continuous data protection	File-level deduplication
RecoverX Deployment	Single node or clustered (3 or 5 nodes)		
User Interface	Graphical User Interface, CLI, Restful API		
Supported Versions	Apache Cassandra 2.1, 3.0, 3.10 DSE 4.8, 5.0, 5.1	MongoDB 3.0, 3.2, 3.4	CDH 5.8+ HDP 2.5+

Multi-cloud is the new normal. To keep pace with this cloud transformation, enterprises must adopt a cloud-first data management strategy. Datas IO's mission is to help organizations accelerate their adoption of multi-cloud by enabling them to protect, mobilize, and monetize their traditional and next-generation applications. Datas IO has leapfrogged the data management industry, creating the world's first hyper-scale, distributed data management platform — Datas IO RecoverX. Datas IO RecoverX is elastic, scale-out software for multi-cloud environments, delivering efficient data protection and data management services for traditional and next-generation applications.



## About Datas IO

Datas IO is the application-centric data management company for the multi-cloud world. Our flagship Datas IO RecoverX delivers a radically novel approach to data management helping organizations embrace the cloud with confidence by delivering solutions that protect, mobilize, and monetize their data — at scale. Datas IO was recently awarded Product of the Year by Storage Magazine, and was recognized by Gartner in the 2016 Hype Cycle for Storage Technologies. Backed by Lightspeed Venture Partners and True Ventures, Datas IO is headquartered in San Jose, California.