

What's New in RecoverX 2.0

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

What's New in RecoverX 2.0

Enterprises are increasingly adopting next-generation applications and migrating traditional applications to multi-cloud environments. As a result, enterprise IT infrastructure now consists of multiple clouds (public, private and hybrid) in distributed geographies, all connected over multiple networking links. However, every cloud uses unique technology for infrastructure services. For example, private clouds are predominantly powered by the VMWare ESX virtualization suite, while AWS and Azure public clouds use Xen and Hyper-V for their virtualization technology. Accordingly, there are no ESX virtual machines and SCSI LUNs in the public cloud, and the only common denominator binding all the clouds together is the data itself, which does not change across clouds.

Because of these shifts, customers now need a comprehensive Cloud Data Management solution. That is why we are launching RecoverX 2.0 around three key dimensions: Cloud Mobility, Data Protection, and Platform Enhancements.

Cloud Mobility

- **Cloud Data Mobility For Relational Databases:** In multi-cloud environments, enterprises may run their production applications in one cloud (public or private) and test/dev instances in

another cloud. With RecoverX 2.0, customers can now move their relational databases (starting with Microsoft SQL Server) at a table-level granularity from on-premise to the cloud, across clouds, and from the cloud back on-premise.

- The **key product innovations** are grounded in semantic deduplication that drives 10x storage efficiency over traditional deduplication techniques, application listeners that perform next-generation source-side deduplication techniques, and our metadata catalog that is distributed across cloud boundaries to allow customers to backup anywhere, recover anywhere, and migrate anywhere.

Data Protection

- **Data Protection for Big Data Filesystems (HDFS):** Filesystem backup is the most prominent use case for the backup and recovery market. File system corruptions are the norm, and this is amplified in the triple replicated architecture of HDFS based file systems. To address the use case of operational recovery, we now support data protection for petabyte-scale Hadoop distributions to enable file-level recovery and archiving for long term retention. Both Cloudera and Hortonworks distributions

are supported. Backup or archive data can be stored either on-premise on S3 compatible object storage or in the cloud on cloud-native S3 storage -- giving customers freedom and choice.

- **Data Protection for Relational Databases:** Enterprises deploy Microsoft SQL Server on physical servers, virtual machines or hyper-converged infrastructure (HCI). Similarly, enterprises are deploying Microsoft SQL Server on Azure Cloud, Amazon Web Services (AWS), and Google Cloud. Regardless of infrastructure deployment and public cloud, customers are looking for application-centric data protection for Microsoft SQL Server that frees them from the underlying infrastructure (physical, VM, or HCI or cloud-native based). To address this, we have extended RecoverX to support data protection of relational databases, including Microsoft SQL server in private or public cloud environments to protect at an individual table level and recover only what you want.

RecoverX Platform Enhancements

- **RecoverX Platform Scale & Elasticity:** Enterprises can now deploy RecoverX in a 5-node cluster configuration to protect large scale database clusters. You may also dynamically scale up and scale down compute infrastructure to meet performance requirements depending on the application load.

- **Enterprise Policy Management:** Several enhancements to policy management have been made to simplify and automate data protection operations for DBAs, backup and recovery admins, and IT operations team. These include:
 - Adding/removing tables from a policy
 - Pausing/resuming a policy
 - Backup NOW
 - Rule-based addition for dynamically generated database tables
- **Operational Metrics:** New metrics are now available from the GUI to give a detailed operational overview to administrators. This allows administrators to monitor their data change patterns and storage footprint for troubleshooting and planning purposes.
- **Automated Deployments for Multi-Cloud Environments:** As enterprises adopt next generation infrastructure consumption models, such as containers and cloud-native deployments, RecoverX is available on Amazon Machine Image (AMI) and Docker image to simplify deployment and customer onboarding experience.
- **Support for Latest Database and OS versions:** RecoverX now supports the latest versions of databases - Cassandra 3.X, DSE 5.X and MongoDB 3.4. Customers may now also deploy RecoverX on RHEL/centos 7.X operating system.



About Datos IO

Datos IO is the application-centric data management company for the multi-cloud world. Our flagship Datos 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. Datos 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, Datos IO is headquartered in San Jose, California.