DATOS IO

Datos IO RecoverX 1.5

Overview

Enterprises are increasingly adopting scalable and non-relational databases (such as Apache Cassandra, MongoDB and more) rather than the traditional relational databases to handle the data requirements of their next-generation cloud applications (such as analytics, SaaS, IoT, ecommerce, and more). This fundamental shift raises critical issues in the lifecycle of data management. Traditional backup and recovery products were designed for small-scale databases, tape-based storage media, and on-premise deployments. This leaves the next-generation of cloud-native applications with a critical data protection gap.

Datos IO RecoverX is a purpose-built software that addresses the data protection requirements of cloud-native applications deployed on non-relational databases. RecoverX uses its industry-leading CODR™ architecture to address the following use cases - (1) Backup and Recovery (2) Test/Dev Refresh (3) Migrations and Upgrades. The CODR™ architecture does not dependent on media servers, and transfers data in parallel to and from file-based and object-based secondary storage. CODR™ delivers cluster-consistent backups that are highly space-efficient yet available in database native formats.

Key Benefits

Enterprises can use RecoverX for protecting their business-critical applications on-premise and in the cloud. RecoverX substantially reduces the risk of data loss and provides huge cost savings – both operational and capital investments.

- Reduce application downtime by ~6x through orchestrated recovery without any database repairs
- Reduce secondary storage requirements by ~70-95% through semantic deduplication and advanced support for compaction
- Reduce deployment infrastructure costs by ~80% through media-serverless architecture
- Increase DevOps efficiency to refresh test/dev environments by ~8x through flexible restore and data masking
- Increase operational resiliency through failure handling of source nodes/databases

---

Key Product Features

RecoverX is a scale-out software-only solution that may be deployed on any cloud platform (private or public) in a cluster configuration for high availability and performance. It provides an application-consistent versioning so no database repairs are required after recovery reducing the application downtime drastically. Its ability to restore to test/dev clusters with different topology configuration and mask personally identifiable information (PII) brings operational cost savings to DevOps teams in Continuous Development and Continuous / integration environment. Finally, RecoverX only stores a single logical copy of data using its Semantic Deduplication, an industry-first capability, which reduces secondary storage requirements multifold.

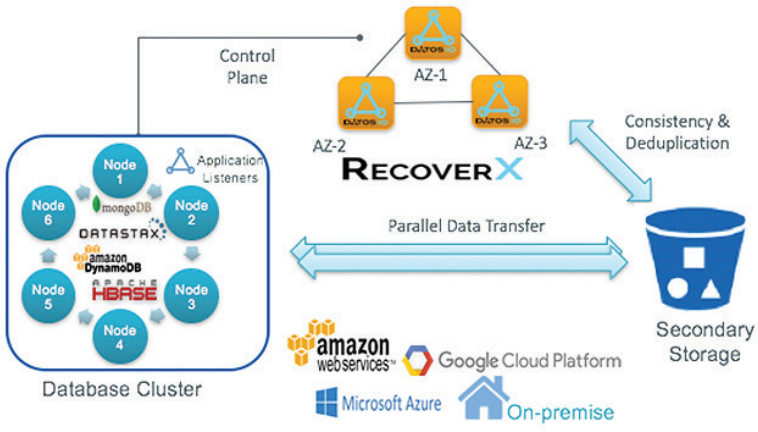

Figure 1 Reference Customer Deployment of Datos IO RecoverX

Features	Cassandra Database	MongoDB Database
Versioning Granularity	Table-level	Collection-level
Versioning Interval	15 mins and above	
Application consistency	Application-Consistent versioning; configurable per application needs	
Recovery Granularity	Table-level	Collection level
Recovery Options	• Fixed point-in-time • Any point-in time* • Node level recovery*	• Fixed point-in-time
Recovery Scenarios	• Recover to source cluster • Recover to target cluster that is smaller than source cluster • Recover to target cluster that is larger than source cluster	• Sharded › Sharded • Sharded › Unsharded
Secondary Storage	NFS, AWS S3, Google Cloud Storage	
Storage Efficiency	Incremental-forever, Semantic Deduplication (~67-90% savings)	
Test/Dev Support	• Restore to different topology • Data Masking*	• Restore to different topology
Advanced Features	• Compaction handling for storage efficiency • TTL data handling for restore	
RecoverX Deployment	Single node, Clustered (3-node) RecoverX	Single node RecoverX
User Interface	Graphical User Interface, CLI, Restful API	

* New features available in RecoverX 1.5

Compatibility Matrix

	Cassandra Database	MongoDB Database
Databases Supported	Apache Cassandra 2.0, 2.1 DataStax Enterprise 4.5, 4.6, 4.7, 4.8	MongoDB 3.0, 3.2
Deployment Profiles	On-premise (Physical Server, VM, NAS Storage) Amazon AWS Cloud (EC2 Instances, AWS S3) Google Cloud Platform (Google Compute Engine, Google Cloud Storage)	
RecoverX System Requirements (per node)	Compute – 8-core Local Storage – 128GB Local Memory – 32GB	Compute – 8-core Local Storage – 256GB Local Memory – 32GB
Operating System	RHEL/centos 6 (Python 2.6)	