

# Highly Scalable Data Management for Cassandra and DataStax Enterprise

The number of mission-critical applications built on top of Cassandra and DataStax Enterprise is rising rapidly. Although modern data platforms like Cassandra have basic data redundancy (multiple replicas) built into them, redundancy alone isn't enough to protect companies from large-scale loss of mission-critical data resulting from user errors, ransomware, or application corruption. Loss of data and increased downtime can severely impact your revenue and business reputation.

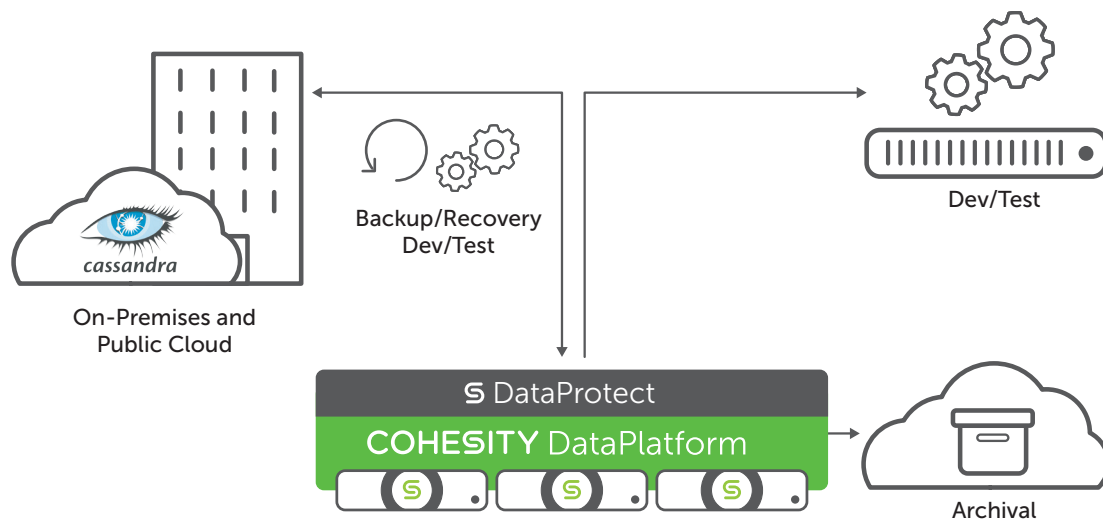
Platform tools or homegrown solutions cannot handle the large volume of data, so enterprises need a new approach to manage their Cassandra applications. They need intelligent systems to effectively manage and quickly clone or recover data so they can get up and running faster with minimal downtime. More intelligent, rapid data management is a must-have in today's modern data world.

## Simplified Management and Protection

Cohesity provides the industry's fastest data management software to handle massive data sets — from a few terabytes to petabytes and beyond — residing on modern data platforms such as Cassandra, DataStax Enterprise (DSE), MongoDB, Couchbase and, Hadoop HBase/ Hive/HDFS. Cohesity ensures data resiliency in the event of disasters or corruption, enabling companies to get back online faster. Cohesity backs up, clones, and recovers terabyte and petabyte-sized data sets and beyond much faster than any other solution on the market, minimizing the impact of data loss associated with human and application errors and reducing downtime to minutes and hours, as opposed to days and weeks. Cohesity offers extreme scale, rapid recovery, and smart storage optimization, one of the many reasons why Cohesity is used by the leading Fortune 500 businesses in retail, financial services, and healthcare.

## KEY BENEFITS

- *Enterprise-class data management and protection*
- *Rapid recovery*
- *Limitless scalability for petabytes and beyond*
- *Make backup data productive*
- *Reduce data footprint with Content-aware storage optimization*



## Key Features

- **Flexible Backup and Recovery** allows backup and recovery of individual tables, keyspaces, or the whole database to the original or alternate location of a different size or configuration.
- **Extremely Scalable** uses a scale-out architecture that can support a few terabytes to petabytes of data.
- **Rapid Recovery** leverages the scale-out architecture, fully materialized restore points, and Imanis Data FastFind™ so enterprises can quickly locate the backed-up data and restore rapidly.
- **Any Point-in-time Recovery** allows recovery of data to a specific date and time further reducing data loss.
- **Cloud Native architecture** supports deployment on various Cloud platforms, including AWS, Azure, and Google Cloud. Older backups can be transparently migrated and recalled from object storage to reduce costs.
- **Enterprise-class solution**, including key security capabilities (Kerberos, LDAP integration, granular RBAC, encryption), auditing, alerting, high availability with erasure coding, and call-home technologies.
- **Content-aware Storage Optimization** uses specialized compression, de-duplication, and erasure coding, saving significant resources and cost.
- **Network Efficient** reduces network costs by employing techniques such as throttling, incremental-forever transfers, compression, and multi-datacenter awareness.
- **Heterogeneous Backup** enables backup of multiple applications and databases using a single user interface
- **Consolidate Backup and Dev/Test** on a single platform to simplify data management.