

Architecting for Rapid Recovery at Scale

by DCIG Principal Storage Analyst & Partner, Ken Clipperton



COMPANY

Arcitecta
 LOUISVILLE OFFICE
 357 South McCaslin Blvd.
 Suite 200
 Louisville CO 80027
 USA
 +1 720 583 4006
 talk@arcitecta.com

MELBOURNE OFFICE
 15 Little Bakers Lane
 Northcote, Victoria 3070
 Australia
 +61 3 9029 3437
 talk@arcitecta.com

arcitecta.com

INDUSTRY

Data Management Software

SOLUTION

- Mediaflux
- Point in Time Ransomware Rapid Recovery Solution

This paper discusses the concept of architecting for rapid recovery at scale, explores the key principles and technologies involved, and introduces the key capabilities that enable rapid recovery. Organizations adopting this approach will minimize the impact of inevitable system failures.

Rapid Recovery is Challenging Yet Essential

Many core business operations require consistent access to business applications and the rapidly growing amounts of business data associated with those operations. Any disruption to that access can have significant impacts on businesses. Thus, organizations have a growing need for rapid recovery from system failures, no matter the cause of the failures.

Minimizing the impact of system failures requires an architectural approach that focuses on rapid recovery at scale. As the amount of data an organization manages grows beyond several hundred terabytes, existing backup solutions and manual procedures do not meet the organization’s recovery requirements. Thus, they discover they need a next-generation enterprise data platform that is architected for rapid recovery.



The importance of making rapid recovery the design center for a next-generation enterprise data platform cannot be overstated, as orga-

nizations that fail to prioritize rapid recovery risk significant financial losses, reputational damage, and even legal consequences.

Beyond rapid recovery, organizations need a data platform that enables them to derive more value from data, while optimizing their growing data estates for access, performance, and cost.

Foundational Data Platform Capabilities for Rapid Recovery at Scale

A global namespace provides the essential foundation for intelligent data management and efficient data sharing. It enables a unified view of the organization’s data regardless of location or underlying storage system, with identity and policy-based access controls.



Intelligent data placement and movement so that the data is in the right place at the right time and at the right cost.

Data provenance refers to the metadata that identifies how and when data was created, who has accessed or modified it, and what has been done to it. This metadata can help organizations ensure data quality and traceability, aid in compliance with data regulations, and support data analytics and management.

“Minimizing the impact of system failures requires an architectural approach that focuses on rapid recovery at scale.”

Key Principles of Architecting for Rapid Recovery at Scale

Resilience. The first key principle of architecting for rapid recovery is resilience. Resilient designs minimize downtime through redundancy and fault tolerance so that there is no single point of failure in the system.

Designing for recoverability. Recoverability includes efficient and effective recovery capabilities that address all the major sources of downtime and data loss, and does so inside the recovery point objective (RPO) and recovery time objective (RTO) limits required to sustain the business. Financial constraints generally result in organizations accepting some level of downtime and data loss as part of their recovery planning. But the ideal for both RPO and RTO is zero—zero data loss and zero downtime.

Immutable data, not backup sets. Protecting multiple petabytes of data is best achieved by making data immutable in place, rather than by making multiple copies of data in multiple places.

**How would you backup 100 petabytes of data?
You don't.**

While backup sets have traditionally been the primary method for protecting data, creating and managing multiple copies of data becomes ever more time-consuming, costly, and challenging as data volumes grow.

In contrast, immutability offers a more efficient and reliable way to protect data. By making data immutable in place, organizations can ensure that their data remains protected against accidental or malicious modification, corruption, or deletion. This approach eliminates the need for multiple backup sets, simplifies data management, facilitates compliance with data regulations, and can enable rapid recovery since massive data movement is not required to restore systems to an operational state.

Software-defined infrastructure. Software-defined infrastructure facilitates recovery by abstracting hardware resources and enabling the quick deployment of IT services, including virtualized servers and storage in the primary data center, in a secondary data center, or in the cloud.

Recovery anywhere. A well-architected solution can leverage private or public clouds to enable organizations to recover individual data sets and applications in the cloud, or even full Disaster Recovery as a Service (DRaaS) capabilities in the cloud.

Key Capabilities for Rapid Recovery at Scale

Policy-based automation and orchestration. Automation and orchestration tools benefit routine procedures and are critical for efficient recoveries. Automation minimizes the delays and the opportunities for errors associated with manual processes and interventions. Policy-based automation enables consistent data management while reducing ongoing administrative overhead and is critical to managing data at scale.

Self-service tools. Accidental deletion or overwriting of files and other human errors account for many data loss incidents. Robust and easy-to-use self-service tools enable end users, application administrators, and developers to recover from these incidents without the delays associated with involving IT staff in the recovery process. This approach saves time and minimizes disruption for all involved.

“Every business must develop and implement a strategy for timely recovery from successful cyberattack attacks and other system failures”

Guidance for Business and IT Leaders

Every business must develop and implement a strategy for timely recovery from successful cyberattack attacks and other system failures. Enterprising professionals will leverage this necessity to implement a next-generation enterprise data platform that creates new opportunities for their organizations while mitigating the risks posed by ransomware.

Beyond meeting data recovery time requirements, a truly scalable solution must enable effective self-service for its users, automated data placement for optimal business results, and enterprise-class data governance.



Arcitecta’s Mediaflux Data Fabric Enables Rapid Recovery at Scale

Arcitecta is a data management company, not a storage company.

Its Mediaflux data fabric virtualizes multiple underlying storage systems to create an integrated data environment that enables continuous inline data protection at scale. Its Point in Time solution solves the rapid recovery problem by providing IT and end users with self-service tools that rewind any file or file system to a specific point in time without performing a restore operation. Its intelligent search facilities help to reduce the RTO to near zero, even at scale. ■

About Arcitecta

Arcitecta is a creative and innovative data management software company. Founded in 1998, Arcitecta builds the world’s best data management platforms, enabling thousands of users worldwide in some of the most demanding data-driven environments. Arcitecta’s flagship Mediaflux platform began with the vision to provide organizations with extraordinary technology for handling all forms of data, from small to very large and complex. Today, it forms the foundation for managing the simplest and the most complex data for all sizes of organizations and global enterprises, empowering them to simplify data-intensive workflows and accelerate time to insight from their data to improve business and research outcomes.

About DCIG

The Data Center Intelligence Group (DCIG) empowers the IT industry with actionable analysis. DCIG analysts provide informed third-party analysis of various cloud, data protection, and data storage technologies. DCIG independently develops licensed content in the form of TOP 5 Reports and Solution Profiles. Please visit www.dcig.com.