druva

White paper

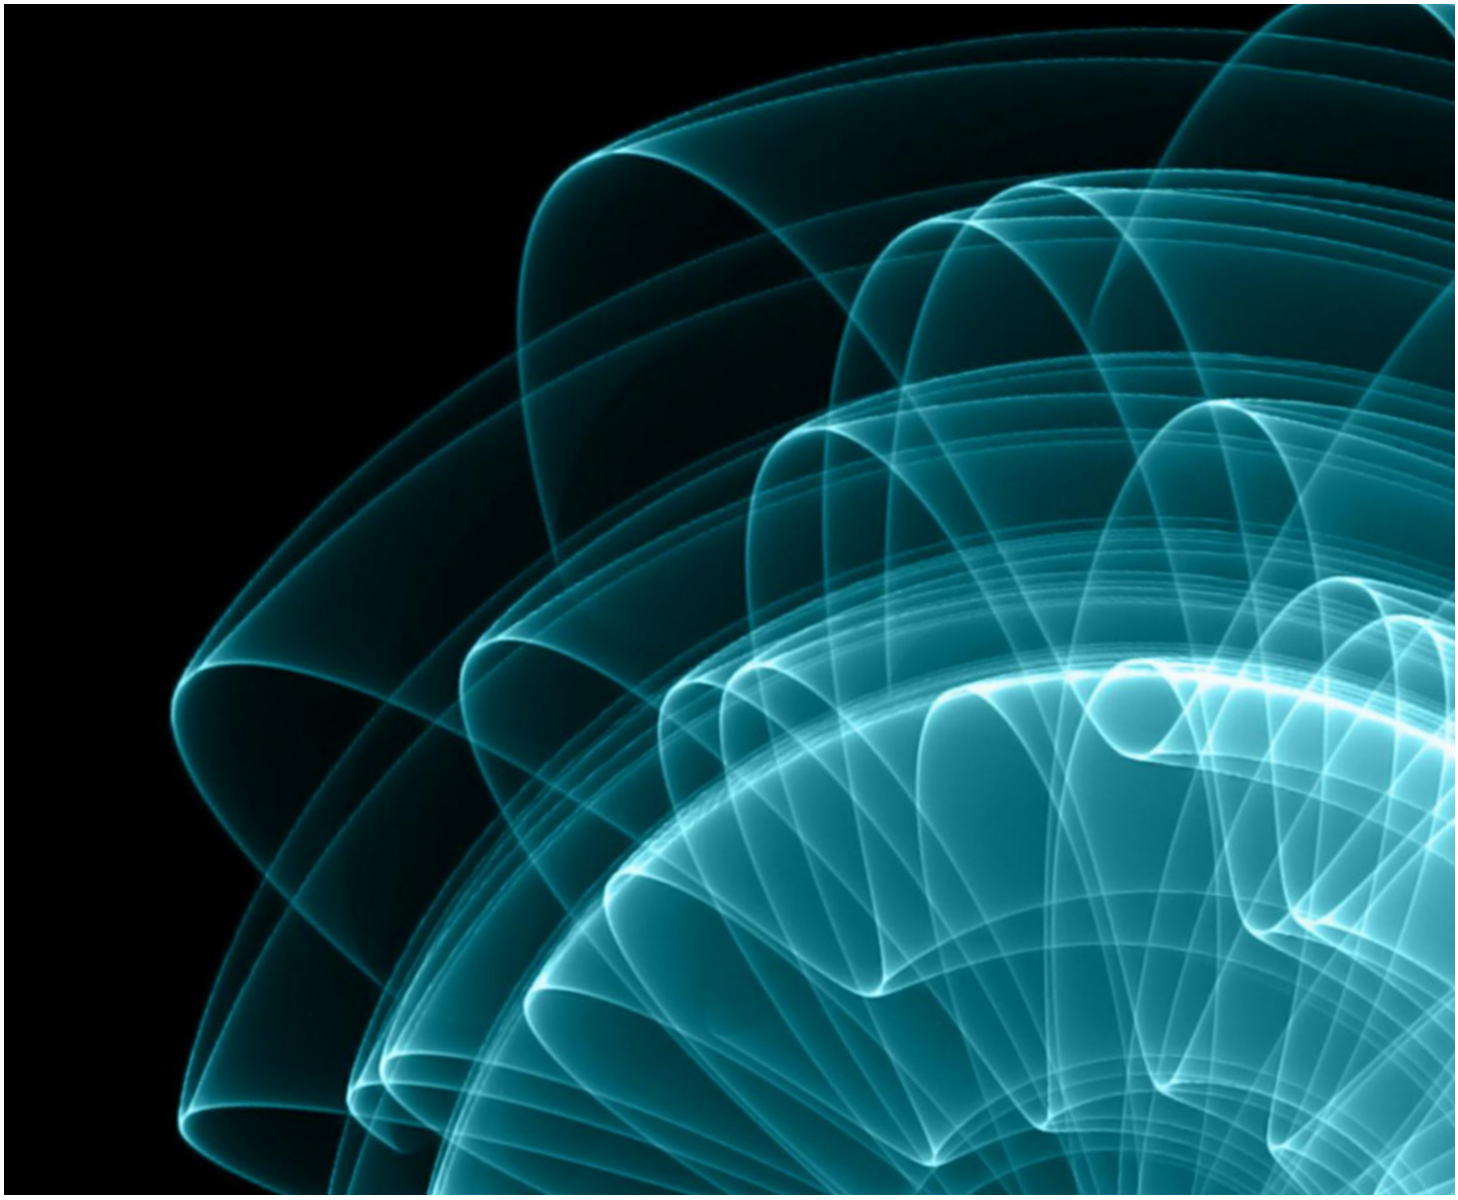

# The Four Critical Stages of Building Multi-Cloud Data Resilience



Q422-20330

## Executive Summary

As enterprises migrate their business processes to multi-cloud environments, they face numerous threats to data resilience. The sheer volume of priceless data has grown exponentially, and in the process, has often become siloed, poorly governed, and at risk of corruption, deletion, and ransomware. The solution is centralized management and automated protection processes made possible by robust, cloud-native SaaS apps.

This white paper discusses:

- Stages of the multi-cloud journey
- Data resilience in the cloud era
- Changing needs during the journey
- Critical requirements for a multi-cloud data protection solution
- How Druva protects data throughout the multi-cloud journey

## Stages of the multi-cloud journey

Adopting a multi-cloud strategy — leveraging multiple cloud computing and/or storage resources — is a journey. The starting point is where cloud apps and services are the exception, and the business is just experimenting with as-a-service efficiencies. The subsequent stages gradually expand usage, eventually leading to the adoption of multi-cloud architectures that encompass virtually all business processes.

The drivers for multi-cloud transformation are compelling:

- **Best-of-breed applicability —** Using multiple cloud platforms enables pairing a given workload with its most appropriate service. For example, an enterprise's SQL workloads may do best with one cloud, AI/ML apps on another, and SaaS productivity apps on proprietary clouds.
- **Greater business agility —** Different teams often have skills with different cloud tools; companies can avoid disruption during an acquisition by maintaining existing cloud architecture.
- **Lower costs —** A multi-cloud approach enables picking and choosing service providers based on how effectively, and affordably, they can meet specific enterprise needs.
- **Less vendor lock-in —** Managing cloud-native apps in production environments is complex. Relying on a single cloud service provider's proprietary tools to handle things is, for many enterprises, too much of a constraint.

This journey, of course, never ends. Cloud technologies have evolved so rapidly that most enterprises are in a constant state of simply keeping up. Indeed, IDC predicts over 70 percent of new applications will be based in the cloud by 2024[1]. Enterprises will find themselves running lift-and-shift legacy apps such as Oracle and SQL Server alongside SaaS apps such as Microsoft 365, Salesforce, and Google Workspace. And they'll all be hosted on both public and private resources, from hyperscalers like AWS, Azure, and Google Cloud to on-premises data centers and edge environments.

## Data resilience in the cloud era

Enterprises are generating and consuming unprecedented volumes of data that, too often, are distributed in difficult-to-access silos. These silos evolve for several reasons. Organizations rely on complex workloads with different types of data. The workloads operate across multiple clouds in diverse environments. This fragments data, increases the complexity of IT infrastructure, decreases business agility, and increases management, infrastructure, and resourcing costs.

1 IDC FutureScape: Worldwide Cloud 2020 Predictions, October 2019

Yet perhaps the greatest risks that accompany the journey to multi-cloud are those affecting data resilience: the ability to quickly recover business processes and data equity when outages or losses occur. Data is the lifeblood of the modern enterprise — any constraints are costly, and serious disruptions can be fatal.

A recent ESG study[2] found that 33 percent of production apps are mission critical and 58 percent of these apps are multi-cloud. And most enterprises tolerate less than an hour's outage to trigger disaster-recovery protocols. Clouds are secure and resilient, but the threats to business data are inescapable whether from accidental content deletion, corruption, ransomware, or malicious insider/external attacks.

### Factors that complicate multi-cloud data resilience

Protecting and managing multi-cloud data to ensure resilience involves significant challenges: shared responsibility models of cloud services (data, networks, applications, and operating systems) are often poorly understood; talent for securely moving to cloud-native architectures is rare and expensive; and, errors complying with regulatory mandates such as the GDPR, HIPAA, and CCPA risk significant fines and impact brand reputation. For example, in 2021, GDPR authorities fined Whatsapp €225M[3].

In addition:

- Enterprises that count on their legacy data protection solutions for data resilience through their multi-cloud journey are at particular risk. Legacy solutions weren't architected for the cloud, and they struggle to protect evolving cloud environments. Effective cloud-native apps are not rearchitected legacy apps, rather, they are wholly new and built specifically for the cloud.
- Maintaining data visibility across distributed infrastructure is hard-to-impossible when fragmented point solutions (multiple backup products) reinforce data silos. Any lack of visibility directly threatens data resilience.
- When a growing cloud footprint increases complexity, IT costs rise. Long-term data retention is typically cumbersome and expensive, as are storage expenses for rapidly expanding volumes of data.
- A persistent myth that SaaS productivity apps include adequate data protection significantly undermines multi-cloud data resilience. Users assume their data is forever safe in the cloud, when in fact it is always vulnerable to a considerable number of threats. Microsoft, Google, Salesforce, and other SaaS providers all stress the need for independent, third-party data protection solutions.

## Changing needs during the journey

Broadly speaking, protecting data during the journey to multi-cloud starts with a project stage, gets moving in a foundation stage, picks up steam in a migration stage, and matures at a reinvention stage. These stages have unique characteristics that affect how enterprises optimize data resilience. Starting out, enterprises recognize the value of a multi-cloud strategy and the need to adapt their data protection to cloud environments. They're concerned about:

- Long-term retention of certain types of data
- Ransomware/malware recovery
- Disaster recovery (DR)

Their target workloads are located in:

- Network attached storage (NAS)
- Virtual machines (VMs)
- File servers (FS)
- Databases (DBs)

At this point, they're evaluating different solutions and are committing to a change.

---

2 A Successful Multi-cloud Data Protection Blueprint, ESG, October 2021
3 BBC

Having successfully integrated some business processes in cloud environments, enterprises in the foundation and migration stages are implementing new data protection strategies that accommodate all their prior needs and add the ability to efficiently handle cloud-native apps. These apps use containers and microservices and require protection capabilities that legacy backup and DR solutions simply can't address effectively. Target workloads have expanded to include multi-cloud databases and public cloud file servers.

As SaaS apps and IaaS/PaaS solutions proliferate in the reinvention stage, enterprises face a much greater challenge to meet data governance requirements. For example, a SaaS workload may be subject to stringent data residence requirements. How does IT manage long-term storage within a specific availability zone? And the vulnerability of SaaS data in apps such as Microsoft 365, Google Workplace, and Salesforce persists.

## Critical requirements for a multi-cloud data protection solution

Enterprises at any stage of implementing a multi-cloud strategy have several basic needs for their data protection solution:

- To eliminate creating silos for data generated by cloud apps across distributed, multi-cloud infrastructure, centralized operations are a must. Enterprises need visibility into the status of all their data from one control plane — using four or five data protection solutions for different workloads isn't sustainable. This is the only way to minimize data silos and enable effective data governance that complies with ever-increasing regulations.
- Cloud-native, multi-cloud apps require cloud-native data protection solutions rather than legacy products that have been adapted for cloud environments. Containers and microservices represent a different and fast-evolving class of computing. Any app that can work with them efficiently must be built similarly.
- The hallmarks of SaaS and other cloud services are lower costs, simplicity, and scalability. These services eliminate capital expenditures for backup hardware as well as minimize operational expenditures for updating and maintaining software. They are easy to run with substantial automation and without specialized expertise. And they are infinitely scalable, immediately accommodating rapid shifts in data volumes.
- Minimizing exposure to ransomware and other cyberthreats requires an air-gapped solution that's been designed from the ground up with cloud-environment SecOps in mind.

## Druva protects data throughout the multi-cloud journey

Druva offers the first, at-scale SaaS platform that simplifies data resilience for modern multi-cloud environments. It centralizes data protection with multi-tenant, integrated workflows for comprehensive backup, disaster recovery, archival, compliance, and analytics. This helps reduce data silos and increase cyber and operational resilience without any hardware, software, or associated complexity.

> *"We needed our Microsoft 365 data to be backed up offsite, plus endpoint protection, robust disaster recovery (DR) capability, and ransomware protection — preferably all in the same place. With Veeam, it took 30 minutes to restore a file. We couldn't recover most of the files. With Druva, we can restore a lost file in 3 minutes. The service from Druva was spot on. We got in there, configured it, installed it. Done."*
>
> — Rob Ljunggren, IT Director, Vertrax

## Multi-cloud data resilience

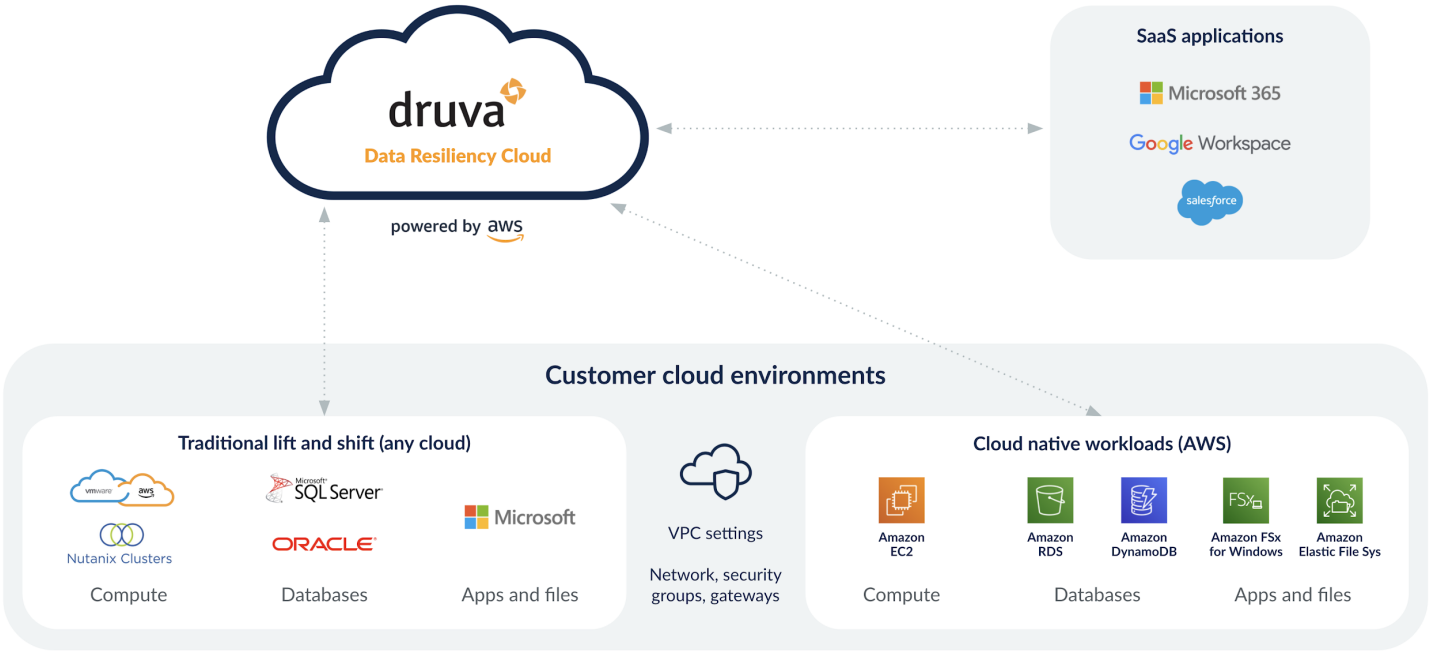


Druva provides the first SaaS defense-in-depth, zero-trust security architecture and immutable, air-gapped backups that ensure enterprises can always recover pristine data. Druva helps you:

- Drive cybersecurity, data governance, and analytics initiatives and leverage AI to transform your backup data into business intelligence.
- Enable eDiscovery and meet data governance and compliance mandates with best-in-class global deduplication, automated tiering, and archiving.
- Consolidate multiple point solutions into a single control pane to reduce infrastructure and administrative overhead.

## Next steps

Visit the Druva site to learn more about how we can accelerate and transform your multi-cloud deployments, and schedule a free demo to experience Druva for yourself.

druva Sales: +1 888-248-4976 | sales@druva.com

Americas: +1 888-248-4976
Europe: +44 (0) 20-3750-9440
India: +91 (0) 20 6726-3300
Japan: +81-3-6890-8667
Singapore: +65 3158-4985
Australia: +61 1300-312-729

Druva® delivers Data Protection and Management for the cloud era. Druva Data Resiliency Cloud is built on AWS and offered as-a-Service; customers drive down costs by up to 50 percent by freeing themselves from the burden of unnecessary hardware, capacity planning, and software management. Druva is trusted by thousands of companies worldwide, including over 50 of the Fortune 500. Druva is a privately held company headquartered in Sunnyvale, California, and is funded by Sequoia Capital, Viking Global Investors, CDPQ, Neuberger Berman, Tenaya Capital, Riverwood Capital, and Nexus Partners. Visit druva.com and follow us on LinkedIn, Twitter, and Facebook.