

The Practitioner's Guide to Cyber Resilience for Cloud Applications

Why Data Backup Alone Can't Save Your Business
(And What Comes Next)



**COMPROMISED
INFRASTRUCTURE**



**INSTANT
RECOVERY**



The Practitioner's Guide To Cyber Resilience For Cloud Applications

Why Data Backup Alone Can't Save Your Business (And What Comes Next)

Table of Contents

- 1 Why Traditional Backup Failed Cloud
- 2 Why Prevention Fails in the AI Era, and Why Cyber Resilience Is a Must
- 3 The Hidden Crisis in Cloud Disaster Recovery
- 4 Infrastructure Recovery: Beyond Data
- 5 The Business Case
- 6 Implementation Framework: From Backup to Recovery
- 7 The Future of Cloud Resilience

Executive Summary

In March 2026, a cyberattack shut down Stryker's global operations. It wasn't until nearly a month later that the company [confirmed in an SEC filing that it was fully operational again](#). This followed an AWS outage in late 2025 (among others) that had already taken down 1,000+ companies for 15 hours. And that was not because data was lost, but because the infrastructure collapsed, and no one had a plan for that.

Now add AI agents into the mix, and the blast radius gets bigger. Agents that provision, modify, and delete infrastructure autonomously can take down in seconds what used to take a human hours to break.


This is the backup industry's most disheartening secret: they protect your databases, but not the VPCs, IAM roles, load balancers, and dozens of additional resources from the configuration layer that make your applications run. When your infrastructure isn't available, there's no business continuity.

For years, the industry assumes that Disaster Recovery & Business Continuity is the same as backing up data. Modern cyber attacks, increases in cloud outage frequency, and the risk of AI Agents deleting and/or changing infrastructure makes it clear that backups alone aren't sufficient anymore

Resilience is not about back ups; it's about bringing your business applications back, quickly. This white paper explains what that looks like, and how leading organizations are getting there.

Why Traditional Backup Failed Cloud

Gartner's 2025 research on CAIRS, Cloud Application Infrastructure Recovery, shattered a comfortable illusion: that traditional data protection tools work in modern cloud environments. The reality is harder: cloud and third-party backup tools typically protect *only* the underlying data of an application, but not all dependent cloud services.



This insight from Gartner shouldn't come as a surprise, then.

"Gartner predicts that by 2030, 35% of organizations will utilize CAIRS solutions to complement infrastructure as code (IaC) disaster recovery orchestration, up from less than 5% in 2026: [a 10x jump]. (Top Trends in Backup and Data Protection for 2026, Michael Hoeck, March 12, 2026)

- Gartner

Modern applications are complex ecosystems of resources like auto-scaling groups, security policies, load balancers, and dozens of interconnected services. They all appear and disappear based on demand. When disaster strikes, a clean database backup is useless if the infrastructure that runs the application no longer exists. Over time, more business logic has shifted away from VMs running custom code and into managed infrastructure. This shift is even sharper in AI applications: instead of executing your own code, the app calls out to something like AWS Bedrock or Google Vertex AI, and the infrastructure becomes the business logic, not just the thing that hosts it.


Multi-cloud adoption has compounded the problem. Gartner identified the core issue:

"Organizations' use of multiple clouds has complicated efforts to consistently enforce and apply proper data protection strategies."

- Gartner

Using 3rd party tools on top of a multi-cloud estate, like Cloudflare for CDN, Datadog for Observability, and Okta for Identity - creates an even larger hole in many teams' Business Continuity and Cyber Recovery strategies.

Businesses spend millions recovering databases while ignoring the infrastructure required to run applications, only to discover the gap when it is too late. The question isn't "Do we have backups?", rather it is "Can we recover?". They are not the same.



Why Prevention Fails in the AI Era, and Why Cyber Resilience Is a Must

For two decades, security strategy has aimed at one goal: stop the attack before it lands. AI has made that goal unattainable, and the industry has largely conceded the point: measuring security success by breach prevention no longer reflects reality, because prevention at scale is no longer achievable.


The reason is structural. Prevention was always an unfair fight, as the attacker needs one working path, and the defender has to secure every one. AI widens that gap, letting attackers map, exploit, and hunt at a scale and speed no defensive team can match. Worse, it runs on a clock defenders can't keep: when an attack executes at machine speed, a defense waiting for a human to review and approve is simply too slow.

Security leaders have drawn the conclusion. Board-level investment is shifting from prevention alone toward response and recovery, driven by an "assume breach" mindset that treats a serious incident as a matter of when, not if. That shift has a name: cyber resilience, what NIST defines as the ability to anticipate, withstand, recover from, and adapt to adverse conditions, attacks, or compromises.

Prevention lowers the odds of an incident; resilience determines its cost. Because the modern business runs on cloud infrastructure, not just the data inside it, resilience comes down to one capability: rebuilding that infrastructure quickly, reliably, and on demand. That is the premise for everything that follows.

The Hidden Crisis in Disaster Recovery

Recent high-profile failures (ThinkL Microsoft Azure automation failures, Google Cloud deleting UniSuper's entire account, GitHub breaking critical workflows, an AI coding agent wiping **PocketOS's production database and its backups in under 10 seconds**, **Amazon's Kiro agent triggering a 13-hour AWS outage**, **a wiper attack disrupting Stryker's global operations**, and the latest AWS outage) all reveal the same problem: organizations cannot reliably recover complex cloud infrastructure when it matters most.



Crisis 1

IaC Adoption and Coverage $\neq$ Recovery Readiness

Enterprises now run on multi-cloud by default, and that default has a cost: fragmented environments, inconsistent controls, and infrastructure sprawl no single team fully understands. Most companies have adopted Infrastructure-as-Code to manage it, but few have the resources to use it as a real recovery tool. That gap between "we have IaC" and "we can actually rebuild from it" is the difference between a bad day and losing a business capability overnight.

Crisis 2

The Recovery Time Disconnect

There is a significant gap between what teams expect recovery to take and what it actually requires. Manual, undocumented rebuilds routinely stretch into weeks: one Accenture client faced a 26-day recovery timeline before automating the process, eventually cutting it to 16 hours. That's a 38x gap between where most teams start and where they need to be. The prerequisite: a trusted, executable blueprint of every asset, stored outside the blast radius. Infrastructure-as-Code is that blueprint.

Crisis 3

Ransomware Targets Infrastructure, Not Just Data

Modern ransomware attacks infrastructure directly. Effective response requires re-instantiating clean environments quickly, reliably, and securely. Traditional backup approaches, which are designed for data restoration, are completely inadequate for rebuilding entire infrastructures from scratch under time pressure.

Infrastructure Recovery: Beyond Data

Gartner defines Cloud Application Infrastructure Recovery as solutions that "automate the discovery of application blueprints to identify the dependencies of all cloud services, data, and configurations required to rebuild, replicate, and recover applications." This represents a fundamental shift: from data-centric to infrastructure-centric recovery.

The Four Pillars of Effective Infrastructure Recovery



Comprehensive discovery maps dependencies between services, applications, and data across multi-cloud environments, including continuous visibility into infrastructure relationships, real-time configuration drift tracking, and external dependencies traditional backup ignores.



Infrastructure codification converts live infrastructure into executable IaC rather than backing up ephemeral resources, generating clean, deployable code while maintaining version control across multiple IaC frameworks.

-  **Automated recovery** enables rapid rebuild through push-button deployment, cross-cloud recovery support, CI/CD integration, and comprehensive rollback verification.
-  **Governance and compliance** ensures recovery meets enterprise requirements through consistent policy enforcement, audit trails, and support for regulatory standards increasingly focused on infrastructure resilience.

The Business Case for Infrastructure Recovery

Cloud and cyber resilience treat disaster recovery as a software engineering problem, not a backup-and-restore task. That shift, from restoring data to rebuilding infrastructure, represents the most significant evolution in recovery thinking since automated backup systems were introduced. DR-as-Code is how that shift gets implemented in practice.

- 1 Revenue Protection and Business Continuity:** Cloud infrastructure failures eliminate entire business capabilities instantly, halting revenue streams. Years of infrastructure configuration can disappear when traditional backup lacks the scope to protect it.
- 2 Regulatory and Compliance Requirements:** Auditors now scrutinize cloud deployments for disaster recovery capabilities. DORA, NIST, SOX, GDPR, and evolving industry standards increasingly require infrastructure resilience that traditional backup cannot address.
- 3 Industry Requirements:** Most enterprises commit to 2-6 hour RTOs to customers and partners. Rebuilding to a "previous state" rarely gets there in time. Rebuilding to an optimal state, fast, is what actually closes the gap.
- 4 The Cost of Inaction:** Recovery times measured in weeks threaten customer relationships and competitive position. Clean data backups mean nothing without the ability to rebuild the infrastructure that runs the application.



Traditional DR recovery asks

"How do we restore our systems to their previous state?"

- Rebuilding to a previous state rarely meets a 2-6 hour RTO.

A cyber resilience approach asks

"How do we rebuild to an optimal state, and keep rebuilding faster every time?"

- Rebuilding to an optimal state, fast, is what actually closes the gap

What's on the line when infrastructure fails:

Revenue, compliance standing, and the trust of every customer counting on uptime

Implementation Framework: From Backup to Resilience



Phase 1: Infrastructure Discovery and Assessment

- **Comprehensive Asset Discovery** scans all multi-cloud resources continuously, capturing configuration details, security policies, and operational metadata.
- **Dependency Mapping** identifies relationships between infrastructure components and business applications, including both technical and business dependencies.
- **Application-Centric Organization** structures infrastructure around business applications rather than technical components, enabling recovery focused on restoring business capabilities.



Phase 2: Infrastructure Codification and Blueprint Creation

- **Automated Code Generation** converts existing infrastructure into clean, deployable IaC using Terraform, Pulumi, CloudFormation, and other frameworks: all with maintainable, modular components.
- **Blueprint Development** creates comprehensive deployment procedures for each application, including dependencies and validation, with regular automated testing across full and partial deployment scenarios.
- **Version Control Integration** manages infrastructure code through Git-based workflows with automated testing and deployment pipelines.



Phase 3: Recovery Automation and Integration

- **Automated Deployment Systems** provide push-button recovery with cross-account, cross-region, and cross-cloud capabilities, plus validation of deployed infrastructure.
- **Validation** of successful deployment to ensure recovery is working, trust worthy, and functioning as planned.
- **Workflow Integration** connects recovery with existing CI/CD pipelines and GitOps processes, enabling self-service recovery capabilities for development teams.
- **Monitoring Enhancement** implements continuous infrastructure state monitoring with automated alerting for configuration drift and policy violations.



Phase 4: Governance, Testing, and Continuous Improvement

- **Regular Testing** validates recovery capabilities through automated, scheduled deployment testing and application functionality verification.
- **Compliance Maintenance** ensures regulatory requirements are met through detailed logging of infrastructure changes, recovery procedures, and testing activities.
- **Continuous Improvement** establishes procedures for reviewing and updating recovery capabilities as infrastructure evolves.

From Backup to Resilience



Four phases to move from data backup to true, automated infrastructure recovery

Phase 1

Discovery and Assessment

- Continuous multi-cloud asset discovery
- Dependency mapping across services and apps
- Application-centric organization

Phase 2

Codification & Blueprint

- Automated IaC generation (Terraform, Pulumi, CFN)
- Executable recovery blueprints per app
- Git-based version control

Phase 3

Recovery Automation

- Push-button, cross-cloud deployment
- CI/CD and GitOps integration
- Automated deployment validation

Phase 4

Governance & Improvement

- Scheduled recovery testing
- Audit-ready compliance logging
- Continuous capability improvement

↑ Continuous Feedback Loop ↓

The Future of Cyber Resilience

Cloud and Cyber resilience are no longer a backup problem. They are an infrastructure problem, and solving it requires treating recovery with the same rigor, automation, and continuous improvement applied to the rest of the software development lifecycle.

The organizations that get there first will share a few things in common: they treat infrastructure as software, recovery as deployment, and every incident as an opportunity to strengthen their systems. Those still relying on traditional backup will find the gap widening: in recovery time, compliance posture, and resilience against the threat types that are already here.

Firefly is purpose-built for this shift. As one of only three vendors recognized by Gartner in the CAIRS category, Firefly's Autonomous platform automates cloud codification, drift remediation, and infrastructure recovery: so teams can rebuild in less than an hour, and stay ahead of the next failure before it happens.

Schedule a 24 Minute Demo

and See How to Build a Resilient Cloud with Firefly

Book Demo