druva

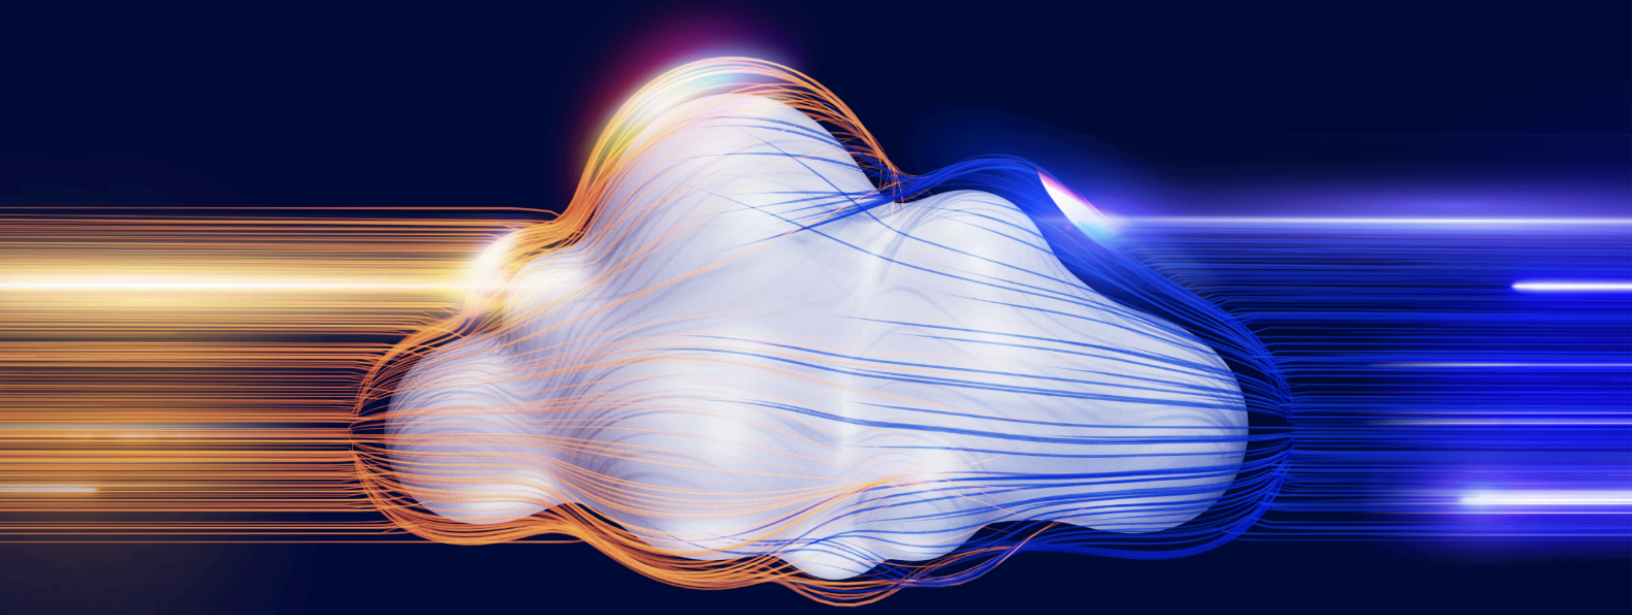

Redefining Cyber Resilience for Cloud-Native Workloads

A Framework for Cloud-First Organizations Using AWS and Azure

Whitepaper
Q226-20569

druva.com

Executive Summary

Enterprises are deepening their investments in cloud-native infrastructure—shifting critical applications and workloads to platforms like AWS and Microsoft Azure. While this transition unlocks agility, elasticity, and scalability, it also marks a decisive change in the threat landscape and the nature of cyber risk.

In this cloud-first paradigm, traditional backup and disaster recovery models fall short. Attackers increasingly exploit identity access, misconfigurations, and platform-native tools to destroy recovery data before it can be used. Cloud providers secure the underlying infrastructure, but data protection and recoverability are squarely the customer's responsibility. This responsibility spans architecture, identity, policy enforcement, and threat response.

Cyber resilience—defined not just by the ability to back up data, but to recover it safely and reliably under attack—must become an operational cornerstone. This whitepaper offers a framework for understanding, evaluating, and implementing cyber resilience for cloud-native workloads, drawing from best practices in the AWS and Azure well-architected frameworks and the latest threat research. It also calls out where existing approaches—custom scripts, native services, and snapshot-based third-party tools—fail to deliver.

The New Threat Model: Complex, Cloud-Native, and Targeting Recovery

Cloud computing has fundamentally reshaped enterprise architecture. Workloads are now globally distributed, ephemeral, and highly dynamic. Yet this architectural shift has introduced new threat vectors, especially as security boundaries blur across shared services, accounts, and regions.

Modern adversaries have adapted. They exploit the complexity of cloud environments by:

- Targeting backup and recovery infrastructure directly to eliminate fallback options.
- Using stolen credentials or insider access to modify, delete, or encrypt recovery data.
- Leveraging automation gaps to disable protections silently or over time.
- Exploiting weak configurations, shared control planes, and unmonitored storage tiers.

These are not theoretical risks. In multiple recent incidents, attackers breached environments and intentionally destroyed snapshots and backup vaults before deploying ransomware. In such cases, the failure wasn't a lack of backup—it was the absence of isolated, verifiable, immutable recovery.

The Shared Responsibility Model—Its Misunderstood Consequences

AWS and Azure operate under a shared responsibility model. While they ensure physical infrastructure security, network integrity, and platform availability, customers retain full responsibility for securing access to their data, managing protection policies, and enabling recovery.

AWS Shared Responsibility Model

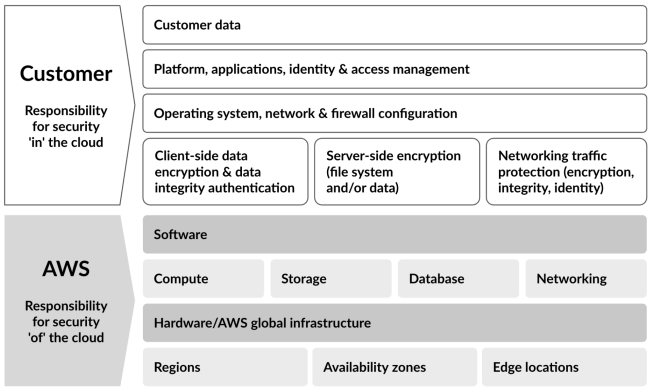

Azure Shared Responsibility Model

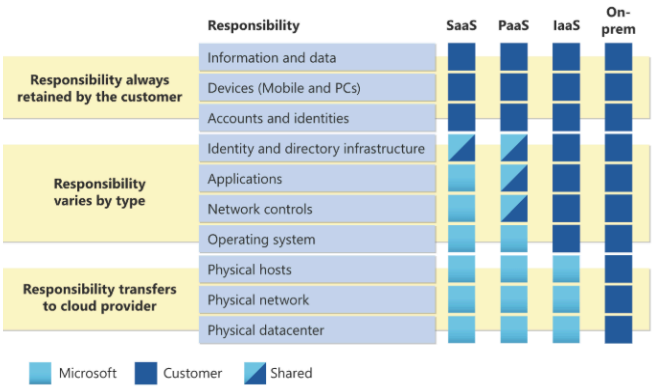

Many organizations assume that enabling native backup tools satisfies this requirement. However, these tools:

- Often operate in the same identity and access management (IAM) boundary as production.
- Do not separate control planes between backup and recovery operations.
- Lack built-in threat detection, rollback, or forensic insight.
- Provide point-in-time recovery without assurances of integrity or cleanliness.

This illusion of coverage is dangerous. If an attacker gains access to the cloud account, they can delete production data and the backups alongside it. In these cases, the recovery mechanism becomes a single point of failure.

Cloud-Native Workloads Demand a New Backup Mindset

Despite the shift to cloud-native environments, many organizations continue to use outdated backup models. These include:

Custom Scripting

Scripting provides granular control, but it is brittle and labor-intensive. It introduces risks such as:

- Silent failures due to syntax errors or missed updates.
- No policy consistency across accounts or clouds.
- No automated alerts or recovery orchestration.
- Lack of auditability or centralized oversight.

Native Cloud Tools

Services like AWS Backup and Azure Backup are easy to activate but were designed for operational recovery, not cyber resilience. Common issues include:

- No immutability or air-gap by default.
- Backups are often stored in the same account or subscription.
- Limited policy visibility across multi-region or multi-account deployments.
- No built-in threat hunting, rollback, or recovery verification.

Other Third Party Backup Vendors

Many third-party solutions replicate snapshots across regions or accounts but rely heavily on basic cloud-native capabilities. They rarely offer:

- Full data separation from the cloud vendor control plane.
- Zero-trust recovery vaults.
- Pre-recovery threat scans.
- Integrated compliance and governance controls.

All of these approaches may check boxes for continuity—but fail when the recovery environment is compromised or corrupted alongside production.

The Operational and Financial Burden of Incomplete Strategies

The indirect costs of suboptimal cloud backup strategies are substantial:

- **Operational Drag:** Maintaining scripts or navigating multiple native consoles diverts engineering time from innovation.
- **Audit Complexity:** Proving retention, immutability, and recoverability requires manual correlation of logs and snapshots across clouds.
- **Unpredictable Costs:** Egress fees, cross-region replication, and storage redundancy in native tools can create budget overruns.
- **False Confidence:** Many teams discover backup gaps only during a crisis—when the data is needed most.

As environments scale across workloads, accounts, subscriptions, regions, and clouds, these issues multiply. A sustainable strategy must prioritize simplicity, automation, and security.

Operationalizing Resilience Across AWS and Azure

AWS Best Practices

(as defined by the AWS Well-Architected Framework):

- Replicate snapshots across regions and accounts to defend against outages and account compromise.
- Maintain immutable, time-separated backups for recovery from corruption or ransomware.
- Isolate backups using separate IAM roles, accounts, and storage vaults.

Azure Best Practices

(as defined by the Azure Well-Architected Framework):

- Use Zone-Redundant and Geo-Redundant Storage (ZRS/GRS) for high availability.
- Store backups in vaults with soft-delete and immutability policies.
- Separate backup control from production IAM paths and accounts.

However, executing these practices manually or through native tools introduces cost, complexity, and risk. A centralized, cloud-agnostic architecture is essential to make these principles actionable.

Data Protection and Cyber Resilience, Simplified for AWS and Azure

Druva is purpose-built for this moment—a fully-managed, cloud-native SaaS platform based on zero trust principles, eliminating infrastructure overhead and delivering unmatched cyber resilience for AWS and Azure workloads.

As cyber threats grow more sophisticated and cloud environments become increasingly complex, you need a data protection solution that simplifies operations, reduces costs, and strengthens your security posture by design.

With air-gapped, immutable backups, built-in threat hunting, anomaly detection, and instant rollback, you get protection that goes far beyond basic recovery. No infrastructure to manage. No egress surprises. Just clean, fast, and reliable recovery—when it matters most. This is cyber resilience, done right—for AWS and Azure.

Zero Infrastructure. Zero Egress. Zero Headaches.

Support also extends to additional AWS and Azure workloads

Why Druva?

Simplified Operations

No infrastructure. No agents. A single control plane to manage protection across all your AWS accounts, Azure subscriptions regions, and workloads — no silos, no complexity.

Comprehensive Data Security

Immutable, air-gapped backups with built-in threat hunting, data lock, rollback actions, cyber recovery, anomaly detection, and quarantine — built for ransomware recovery and cyber resilience.

Proven Savings

Lower Total Cost of Ownership (TCO) than native AWS snapshot replication and Azure storage redundancy strategies with no egress fees, deduplication, and no hidden maintenance costs.

Redefining Cyber Resilience for Cloud-Native Workloads

A Framework for Cloud-First Organizations

Druva's Cyber Resilience Maturity Model provides a practical framework for assessing readiness and advancing through five maturity levels. This model, built on Druva's extensive customer insights, helps organizations boost resilience, enhance data protection, and mitigate risks.

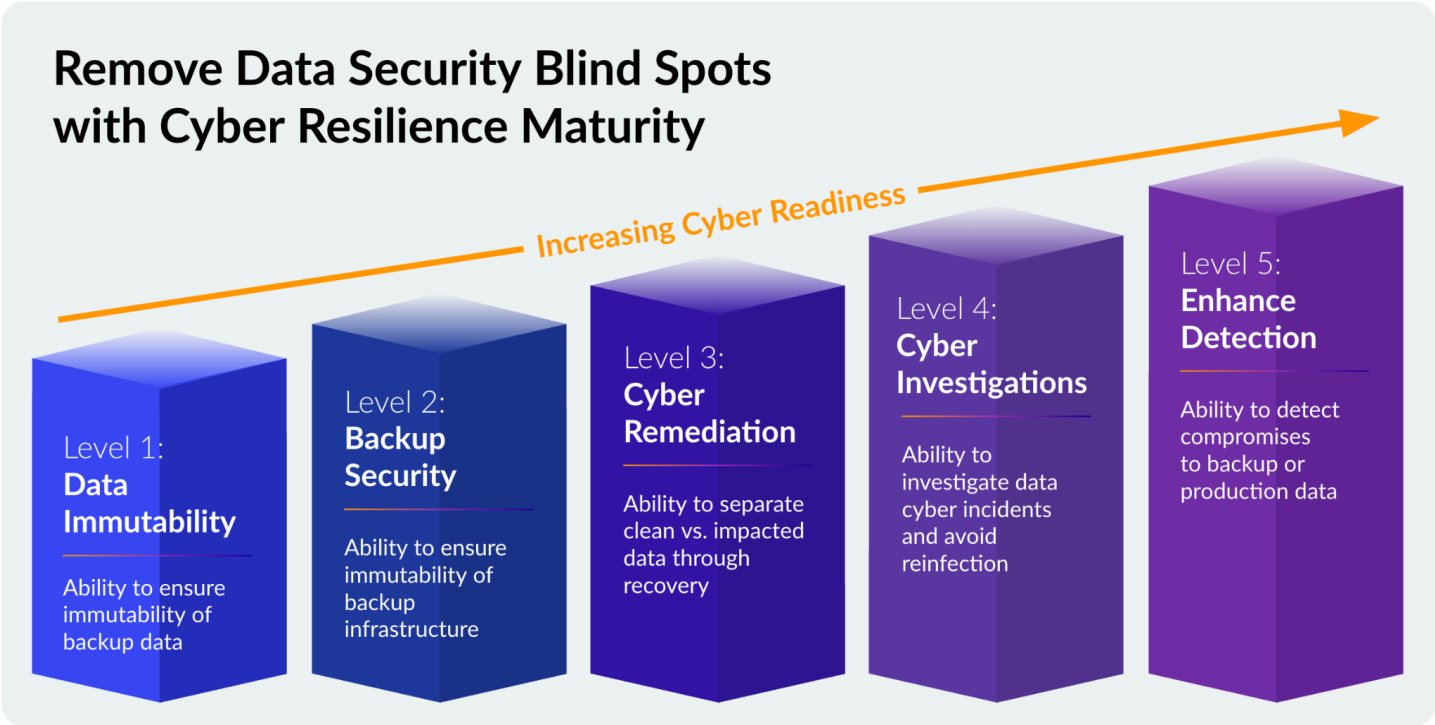

Level 1:
Data Immutability

Protect your backup data from tampering, deletion, or encryption—even by internal threats—with air-gapped, immutable storage.

Immutable, Air-Gapped Backups

Guarantees tamper-proof data storage for assured recovery after cyberattacks.

Data Lock

Prevents backup data from being overwritten or erased before a specified retention period.

Encrypted Data

Data is encrypted in-flight and at-rest to prevent unauthorized access and maintain confidentiality.

Level 2:
Backup Security

Secure your backup environment with a zero trust foundation that separates access, locks down data, and prevents compromise.

Zero Trust Platform

Drastically reduces the attack surface and insider threat risks.

Access Separation

Separation of control and data planes ensures administrative controls don't compromise backup integrity.

Secure Air-gapped Cloud Vault

An isolated cloud recovery vault prevents attackers from tampering with or deleting backup data, ensuring a clean recovery path.

Level 3:
Cyber Remediation

Recover faster and safer from ransomware with clean, quarantined backups and one-click rollback to a known-good state.

Cyber Recovery

Makes every recovery a cyber recovery. Ensures only clean data is restored by scanning backup sets for malware or IOCs.

Rollback Actions

Easily revert systems to known-good states while preserving forensic evidence of the incident.

Quarantine Impacted Snapshots

Automatically isolate suspicious or compromised backups to prevent further contamination.

Level 4:
Cyber Investigations

Accelerate incident response and root-cause analysis with built-in threat intelligence and AI-powered investigations

Threat Hunting

Identifies threats before they spread—minimizing risk and accelerating response by analyzing backup metadata and behavioral patterns for early signs of compromise.

Gen AI Cyber Investigation

Uses generative AI to assist in root cause analysis, threat correlation, and recovery point identification.

Level 5:
Enhanced Detection

Detect threats early and respond faster with 24x7x365 expert monitoring and managed threat detection across your backup data

24x7x365 Monitoring with Managed Data Detection & Response (MDDR)

Druva's expert threat analysts continuously monitor for ransomware behaviors within your backup environment.

Adds a dedicated layer of human-led threat detection and remediation guidance—around the clock.

Data Anomaly Detection

Machine learning establishes baselines and flags suspicious deviations, such as mass deletions or encryptions.

Conclusion: Every Recovery Is a Cyber Recovery

Cloud-native workloads demand a modern approach to data protection—one that sees backup not as an operational afterthought, but as a frontline security control. Organizations must move beyond traditional backup toward architectures that deliver cyber resilience and cyber recovery by default.

This requires rethinking backup as an intelligent, secure, and scalable capability:

- One that decouples backups from the primary environment.
- One that enforces immutability, visibility, and compliance.
- One that enables clean, verified, and rapid recovery—even under active attack.

Cyber resilience is no longer a theoretical ideal. It's a practical, achievable outcome—when recovery is treated not as a fallback, but as a strategic design goal.

Take the Next Step

Ready to modernize your cloud data protection and cyber resilience strategy?

- Start a free trial
- Take a guided product tour
- Request a TCO comparison based on your environment

druva Sales: +1-800-375-0160 | sales@druva.com

Americas: +1-800-375-0160
Europe: +44 (0) 20-3750-9440
India: +91 (0) 20 6726-3300
Japan: japan-sales@druva.com
Singapore: asean-sales@druva.com
Australia: anz-sales@druva.com