Spotlight, brought to you by:

Dreeze

Beyond Data Restores: Recovering Cloud Applications for Operational Success in the Agentic AI Era

WRITTEN BY:

Archana Venkatraman

Senior Research Director, Cloud Data Management, IDC Europe

Johnny Yu

Research Manager, Storage and Data Management, Enterprise Infrastructure, IDC

PUBLISHED:

November 2026

Cyberattacks, AI agents, sovereignty rules and data growth are making resilience harder. IDC examines why recovering whole cloud applications, not just data, needs to become the focus point for truly resilient and consistent digital operations.

I. Introduction

In the digital and agentic era, resilience has increasingly shifted from an IT concern to a boardroom and C-suite priority. Digital businesses are characterized by data growth, deep interconnection between applications, infrastructure sprawl, and ever-challenging cyberattacks, and each of these has intensified in the past two years. Regulators have raised expectations, geopolitical tension has made supply and sovereignty risks tangible, and AI has added both a new class of workload and a new class of risk.

At a glance

Key stats

- 95% of organizations run AI agent workflows in production.
- Only 3 in every 10 organizations are very confident of a full ransomware recovery with minimal disruption.
- 16% of organizations have tiered workloads for recovery enterprise-wide.

Key takeaways

- Regulation, geopolitics, data growth, and AI compound one another and challenge resilience.
- Business resilience must be judged by recovery of the application rather than recovery of the data.
- Tested, repeatable, dependable, and consistent recovery is the evidence boards and regulators seek.

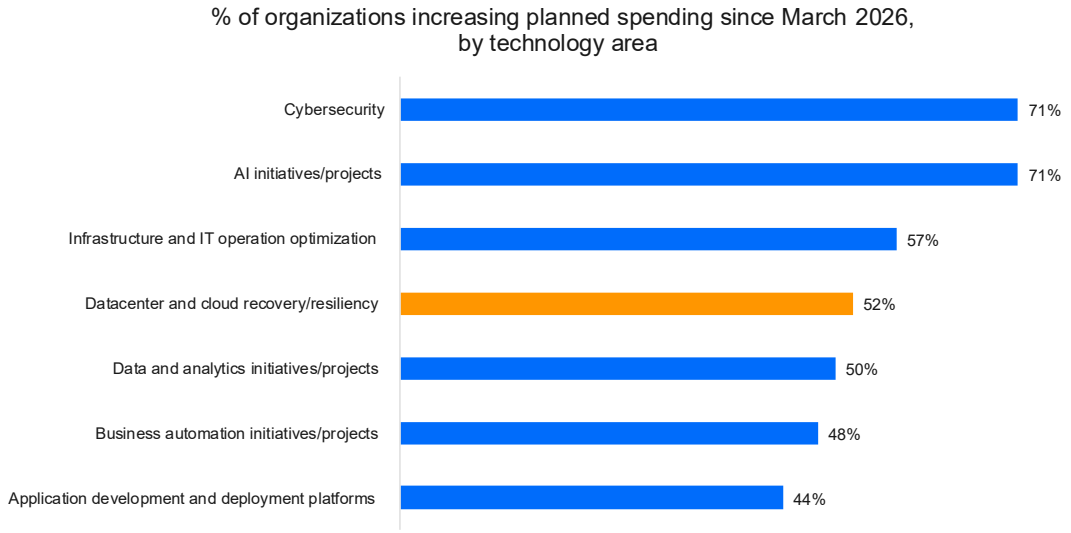
Technology leaders' spending plans reflect this. IDC's July 2026 *Future Enterprise Resiliency and Spending Survey* revealed that 71% of organizations have increased planned cybersecurity spending since March and the same share have increased AI spending. A further 52% had raised funding for datacenter and cloud recovery and resiliency (see Figure 1). Money is moving toward resilience even as budgets tighten.

There is growing scrutiny of whether that spending is delivering desired resilience, digital continuity, and recovery outcomes. IDC believes resilience is both harder and more pressing today for four reasons that compound one another: a denser regulatory landscape that asks for evidence of recovery rather than just intent; geopolitical and sovereignty pressures that instantly threaten workload and data operations; end-to-end digitization raising the cost and reputational damage of downtime; and cloud-native and agentic architectures. These reasons overthrow traditional protection policies and paradigms that are data and infrastructure centric rather than workload- or operations centric.

In addition, recovery success has traditionally been judged by data-centric metrics. Most organizations still report recovery in terms of how quickly a database or storage bucket can be restored. The business, however, experiences recovery as an application that works. Business leaders need to evaluate the impact of the forces making resilience harder, what better looks like, and how to run business-critical applications in the cloud in a stable and risk-mitigated manner.

Figure 1

Planned spending is rising for security and AI; recovery spend is catching up



n = 1,835

Source: IDC's *Future Enterprise Resiliency and Spending Survey (FERS Survey)* Wave 4, July 2026

II. Definitions

Broadly, C-suites want to balance resilience and AI-driven innovation. Successful digital resilience also provides a stable and sustainable foundation for AI adoption. This requires a pivot in how recovery is planned and measured.

- **Digital resilience:** A strategic, data-driven, and technology-led approach to keeping the business operating through disruption, whether from a cyberattack, an outage, an AI failure, or a regulatory change (It extends business continuity and cyber-resilience into how digital services are designed, operated, and tested. It embraces a proactive and intelligent approach to business continuity rather than a reactive one.)
- **Cyber-recovery:** The ability to return systems and data to a known-good, trusted state after a malicious attack, typically from isolated and immutable copies
- **Data recovery:** Restoring the data layer, such as databases, object stores, and file systems; necessary, but on its own not enough for a cloud-native application to run
- **Application recovery:** Restoring everything an application needs to serve users, including compute and orchestration, networking, identity roles, secrets and keys, configuration, and third-party connections, in the correct sequence alongside data (Some suppliers call this application revival. IDC treats it as a core part of digital resilience, distinct from the multizone or multiregion design that is sometimes also called application resilience. Savvy organizations are focusing on recovering the integrity and functionality of their full digital business services as a measure of resilience because data without applications will be of little use for operations.)

III. Benefits

Organizations that can recover all application artifacts rather than only data, and can prove it, gain in five areas:

- **Shorter, more predictable time to restore business services.** Recovery is judged by when customers and employees can transact again, not when a database comes back online or when data becomes accessible again. Automating sequencing and rebuilding removes the trial and error that stretches an incident from hours to days.
- **Confidence built on evidence.** In IDC's 2025 *Cloud Data Logistics and Protection Survey,* only 34% of organizations were very confident, or had no doubt, that they could recover from a major ransomware attack without paying a ransom, without significant disruption, and with minimal data loss. Frequent, nondisruptive recovery tests turn that assumption into a measured recovery time.
- **Less reliance on scarce skills and individual memory.** In the same survey, 45% found it difficult to recruit or retain cyber-recovery preparedness skills. Recovery plans that are generated and maintained automatically depend less on the one engineer who built the environment.
- **Regulatory and audit readiness.** Documented dependencies, defined recovery objectives, and test results are what EMEA regulations such as DORA, NIS2, and similar rules require, and they shorten preparation for each review.
- **Safer speed for the business.** When coverage follows change and rollback has been rehearsed, development and AI teams can deploy and experiment faster, because the cost of a mistake, human or machine, is bounded. It also protects AI investment, since models and agents are only as dependable as the applications and data they run on.

IV. Trends

Regulation is shifting the burden from intent to proof

Regulation is the clearest signal that expectations have changed. DORA, applicable since January 2025, requires financial entities to test their ICT resilience, manage third-party risk, and report incidents, and it brings critical technology providers under direct oversight. NIS2 extends comparable duties, including early warning of significant incidents within 24 hours, to energy, health, transport, digital infrastructure, and other sectors. The GDPR has long required the ability to restore availability of, and access to, personal data in a timely manner. Since September 2026, the Cyber Resilience Act has also required manufacturers to report actively exploited vulnerabilities, which moves some of that discipline into the software supply chain.

The pattern extends beyond Europe. NIST's Cybersecurity Framework 2.0 treats recovery as one of its six core functions and is widely used outside the United States. SEC rules require United States–listed companies to disclose material cyberincidents within four

business days of determining materiality. The United Kingdom's Cyber Security and Resilience Bill, still moving through Parliament, would bring managed service providers and datacenters into scope. Multinationals tend to design to the strictest regime they face, and in IDC's July 2026 *Future Enterprise Resiliency and Spending Survey,* respondents already expect to spend 21% of their 2026 IT budgets on meeting regulatory compliance mandates.

Geopolitics and sovereignty constrain where recovery can happen

Geopolitical risk has moved from scenario to planning assumption. According to *FERS Survey* data, 50% of organizations placed the impact of the Middle East War on prices, profits, and operations among the tech risks that concern them most over the next 12 months. Among those concerned, 50% are lowering business assumptions yet raising IT spending to offset the effect. IDC's *Cloud Data Logistics and Protection Survey* found geopolitical conflicts and tensions to be the risk expected to be most disruptive to infrastructure procurement and consumption over the following 12 months (33%), with 23% citing cyberespionage or warfare that causes infrastructure outages.

Sovereignty is the regulatory expression of the same worry. In IDC's *EMEA Multicloud Survey,* 25% of European organizations said geopolitical uncertainty had increased their interest in digital sovereignty compared with 2024; 22% already use sovereign public cloud and a further 37% plan to within 12 months. IDC predicts that by 2028, 60% of organizations with digital sovereignty requirements will migrate sensitive workloads to new cloud environments to reduce risk and increase autonomy. For resilience, this means recovery plans have to be portable. A runbook written for one provider, region, or account structure does not survive a move to another, and the ability to rebuild an application elsewhere from a trusted copy is itself a sovereignty control. IDC also hears the discussion spreading beyond Europe, notably among Canadian, Middle Eastern, and Asia/Pacific organizations.

Data growth and AI raise the cost of slow recovery

Data volumes keep climbing: 42% of European organizations in IDC's *EMEA CloudOps Survey* reported public cloud data growth above 20% over the previous 12 months. AI changes what that data is for. Agentic systems need real-time, contextual access to data, and IDC predicts that by 2027, 80% of agentic AI use cases will require it (see *IDC FutureScape: Worldwide Data and Analytics 2026 Predictions,* IDC #US53860225, October 2025). When an incident freezes the data and the applications around it, the loss is no longer confined to a customer-facing service. Internal and automated decisions stop as well.

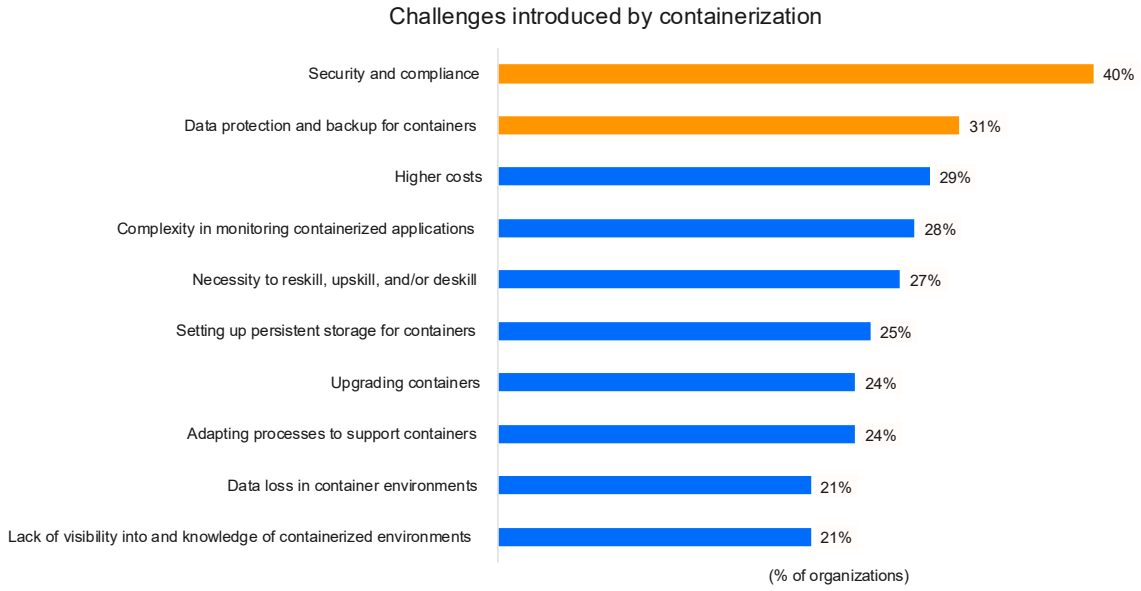
Data quality is the other half. Among organizations using AI, 31% named access to quality data as a top data challenge in moving projects to production, and 29% each named data preparation effort and pipeline scalability. Asked for the single biggest risk to their data platform, 16% chose AI model bias and data poisoning, second only to data security, sovereignty, and privacy compliance (23%). Recovery therefore can't simply restore to the latest copy. Instead, organizations must identify and restore the last trustworthy state quickly enough that AI workloads are not left idle.

Cloud-native design and AI agents change what has to be recovered

In IDC's *Cloud Data Logistics and Protection Survey,* the average organization runs 43% of its production applications in containers, and 48% run at least half. Respondents classify, on average, 58% of those containerized applications as business critical or above. Cloud-native applications are assembled from containers, functions, managed databases, identities, keys, network rules, and external services, defined as code and changed continuously. In IDC's December 2025 *CloudOps Survey,* 31% of organizations identified data protection and backup for containers as a challenge introduced by containerization, and 40% pointed to security and compliance (see Figure 2). A recovery approach built around fixed groups of virtual machines sees little of this, and protection policies tied to manual tagging drift out of date as environments change.

Figure 2

Cloud-native design introduces security, data protection, and cost challenges



Source: IDC's *CloudOps Survey*, December 2025

AI agents speed up that change and add a failure mode of their own. IDC's *FERS Survey* found that 95% of organizations already run at least one AI agent workflow in production, and 72% run five or more. Agents provision, configure, and delete resources at machine speed, so a misconfigured or compromised agent can cause damage in minutes that would take a human insider far longer to inflict.

According to IDC's 2025 *Cloud Pulse Survey,* the top reasons cloud buyers gave for IT operations outages are:

- Internal IT operations failures: 40% of respondents (single most cited reason)
- Downtime across networks: 35%

- Unscheduled maintenance at cloud provider: 33%
- Datacenter downtime: 30%
- Cloud provider downtime: 28%
- Internal security breach: 27%
- External security breach: 24%

IDC predicts that by 2030, up to 20% of G1000 organizations will face lawsuits, substantial fines, and CIO dismissals due to high-profile disruptions stemming from inadequate controls and governance of AI agents. Resilience planning that covers cyberattack but not agent error leaves a gap.

Preparedness has not kept pace with concern

Budgets are rising, but preparedness is uneven. In IDC's *Cloud Data Logistics and Protection Survey,* 79% of organizations expected cyber-recovery budgets to increase, and AI data protection showed the highest expected growth (22% on average). Yet the average organization devotes about 5% of its IT budget to data protection infrastructure, and only 16% have completed a data and workload classification effort across the enterprise to tier recovery; 44% have done so on a limited scale. The leading challenges to cyberpreparedness are staff knowledge and skills (34%), budget constraints (34%), and keeping up with attack types (33%). Coordination between IT operations and security is a top 3 challenge for 29% of organizations, rising to 33% in North America, while 26% are challenged by testing and validation and 23% know they have insufficient recovery capabilities.

These are organizational problems as much as technical ones. Detection, containment, restoration, and forensics often sit with different teams and tools, each with its own view of the environment. Collapsing those boundaries, with shared data on what is running and what is protected, is what separates a rehearsed recovery from an improvised one.

V. Vendor profile: Dreeze

Dreeze is a cloud application backup and recovery company founded in 2025 by a team with a long background in backup and disaster recovery, including at Zerto, now part of Hewlett Packard Enterprise (HPE), and earlier at Kashya, whose technology became Dell RecoverPoint. The company has raised seed funding from investors including Norwest and Bessemer, employs more than 50 people, most of them developers, and emerged from stealth in November 2026. It reports more than 30 design partners, many of them large enterprises, and its platform became generally available in the same month.

Dreeze’s focus is end-to-end recovery of cloud-native applications rather than recovery of the data alone, or ad hoc recovery of data along with minimal discrete resources. Its premise, drawn from calls with many Fortune 100 and Fortune 500 enterprises before it began building, is that many organizations have backups but little confidence in recovering applications in the public cloud.

The platform is designed around four capabilities:

- **Discovery and mapping.** Automated discovery of the cloud environment, classifying hundreds of resource types that include nondata components, such as infrastructure, networking, identity roles, and secrets, and grouping them into applications through deterministic and AI-assisted dependency mapping (Where they exist, tags and infrastructure-as-code definitions can inform the mapping, but the approach is designed not to depend on them. This matters for organizations whose tagging is inconsistent or whose infrastructure code has drifted from what is running. The application dependency map drives the rest of the Dreeze platform, including what gets automatically protected as apps evolve as well as the order in which resources come back during recovery.)
- **Protection.** Backup of data and nondata resources into an isolated, customer-managed vault held in a clean cloud account and protected by immutability, with policies assigned automatically to resources as they appear, intended to reduce silent gaps in coverage (With every resource captured, the state of the application is maintained and can be restored during the recovery process.)
- **Recovery planning.** A dependency-aware plan that sets the order in which resources return and is updated as the application changes, in place of manually maintained runbooks
- **Orchestrated recovery and testing.** Full-stack recovery, including Kubernetes cluster state, that runs either for a live incident or as a nondisruptive test drill, with the option to connect automated application tests to validate the outcome

Dreeze says that applications its customers describe in 5–10 objects often expand to 100+ once dependencies are mapped. The initial platform supports AWS, with Azure support in Early Access at the November company launch and Google Cloud Platform support slated for the road map. Dreeze's early traction is with enterprises building and running cloud-native applications built from IaaS and PaaS services such as Kubernetes clusters, managed databases, object storage, and serverless functions, along with the identity, key management, and networking resources that connect them.

The company focuses on recovery rather than prevention or detection and uses the term *application revival* to describe the whole-app recovery of both data and nondata resources in the right dependency order. Dreeze positions its platform for incidents caused by cyberattacks, misconfigured or compromised AI agents, and human error when building or maintaining complex cloud apps.

IDC considers the design relevant to the gaps described previously: Automatic coverage and generated recovery plans address protection gaps and stale runbooks, while test drills address the infrequent end-to-end testing that organizations report. How well these capabilities hold at enterprise scale will become clear as production deployments grow.

Challenges

Like many new entrants, Dreeze faces some hurdles. As a new vendor, it has a limited production track record and few public references, and buyers of recovery tools typically look for proof under real incidents, along with support depth and staying power, before trusting

critical applications to them. IDC believes its strategy of beta testing, pilots, and POCs alongside reliance on a proven channel ecosystem can help it establish production credibility faster.

Dreeze needs to be mindful that established backup, cyber-recovery, and cloud platform providers are starting to move toward application-level recovery, so the company will need to explain its differentiators. *Application revival* is a new term, and application resilience already carries a different meaning in cloud architecture, so category education will take effort, but it can capitalize on the established vendors' tailwinds. Communicating with a wider range of personas beyond traditional IT and data protection specialists, including security, application developers, and platform engineers, can help drive investments. These personas are closer to application and operational challenges and will value faster time to full application recovery over rebuilding systems from data, runbooks, and templates. IDC predicts that by 2028, 50% of CFOs will take on COO roles, driving system investments with a sharper focus on ROI, risk, and operational efficiency for agile, data-driven enterprises.

Budget is another hurdle. IDC's *FERS Survey* research shows 71% of organizations raising cybersecurity spend against 52% for datacenter and cloud recovery, and security and identity products often claim resilience funding first. Recovery responsibility is spread across infrastructure, security, and application teams, so Dreeze must identify who owns the budget.

Coverage will be tested too. AWS comes first, leaving multicloud, hybrid, and sovereign environments to follow, and organizations with state-heavy AI workloads such as GPU clusters will look for support that the company says it could add later. Finally, value depends on customers committing to regular recovery drills, which requires process change as well as tooling.

VI. Conclusion

Resilience has become harder and more urgent, and organizations must evolve their recovery and restore strategies from data and infrastructure centric to application centric to find success. Regulators ask for proof of recovery, geopolitics constrains where workloads can run, data growth and AI raise the price of every hour of downtime, and cloud-native estates that people and agents change outpace policies written for virtual machines. IDC believes resilience will increasingly be measured by the time taken to return a business service, tested and evidenced, rather than by the existence of a protected copy of data.

For C-suite leaders, the practical questions are whether recovery is measured at the application level, how recently it was tested end to end, and who owns it across IT, security, and development. To the extent that Dreeze can build customer references, extend coverage beyond its initial platform, and show buyers where application recovery fits alongside security spending, IDC believes the company has a credible opportunity in a market where the need is growing.

About the IDC analysts



Archana Venkatraman

Senior Research Director, Cloud Data Management

Archana Venkatraman is based in London and has covered data management and data protection for more than 10 years, with a focus on regulation, resilience, and digital sovereignty in Europe. Her current research also covers agentic AI governance and observability.

More about Archana Venkatraman



Johnny Yu

Research Manager, Storage and Data Management, Enterprise Infrastructure

Johnny Yu is based in Boston and covers data protection and the intersection of cyber-resilience and AI. Before joining IDC, he covered data protection as a technology journalist.

More about Johnny Yu

Message from the sponsor

Dreeze is pioneering Application Revival, going beyond traditional backup and recovery to bring complete cloud-native applications back, not just their data. The platform autonomously maps every resource and dependency to the application it serves, protects everything the app needs to run, then orchestrates whole-app recovery. Dreeze gives cloud teams the confidence of validated recovery after cyber-attacks, rogue AI agents, or human error. Backup veterans Avi Raichel and brothers Ziv and Oded Kedem, founders of Zerto (acquired by HPE) and Kashya (acquired by EMC), formed Dreeze in 2025. Learn more at www.dreeze.com and get a demo to see what Application Revival can mean for your business.

IDC Custom Solutions

IDC Custom Solutions produced this publication. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis that IDC independently conducted and published, unless specific vendor sponsorship is noted. IDC Custom Solutions makes IDC content available in a wide range of formats for distribution by various companies.

This IDC material is licensed for external use, and in no way does the use or publication of IDC research indicate IDC's endorsement of the sponsor's or licensee's products or strategies.



IDC Research, Inc.
One Beacon Street, Suite 33100, Boston, MA 02108, USA
T +1 508 872 8200

blogs.idc.com | LinkedIn @IDC | www.idc.com

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications, and consumer technology markets. With more than 1,300 analysts worldwide, IDC offers global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries.