keepit

Key takeaways

Secure by design

Rethinking data protection in the AI era

AI's promise is real — but today's gains are mostly speed and task-level augmentation.

The organizations that will benefit are the ones that start with a clear business outcome and keep tight control of data and systems as they innovate. As agentic and LLM (large language models) tools knit applications together, traditional permission boundaries blur, attack volume and sophistication rise, and shadow IT accelerates. This means governance, security guardrails, and human-directed oversight are non-negotiable.

This was the key focus of a recent Keepit webinar, where an expert panel led by Kim Larsen, CISO at Keepit, discussed how resilient data protection is the foundation for safe progress in AI. Immutable, sovereign backups with open, predictable access for analysis and precise recovery when things go wrong enable organizations to safely experiment.

Where AI is today and where it is going

Where AI is today

Productivity boosters (assistants, grammar, code acceleration). Helpful but not truly "intelligent"; limited novelty, mostly faster execution.

Where AI is going

Agentic systems making decisions, orchestrating across apps and data. This changes the engineering model from procedural steps to event-driven, outcome-driven operations.

Read our latest blog post: *AI today - Power, drift, and the discipline to stay in control* ⟶

What this means

With agentic AI, the security landscape shifts dramatically: Familiar attacks like phishing, fraud, and impersonation scale to massive volumes and become more convincing, while new vectors emerge through model jailbreaks — sometimes hidden in text or even images — along with remote manipulation of agent behavior and a stronger focus on data exfiltration over encryption.

At the same time, traditional control points erode as small agents link systems together, weakening predefined permissions unless organizations introduce brokers or overseers to enforce constraints and maintain security.

As AI systems act more autonomously, governance and control become more important, not less. Organizations must codify what "good" looks like and enforce it.

How to successfully adopt AI capabilities

Successful AI adoption follows a simple framework: Start with the outcome by defining the decision or process to improve, along with clear success metrics and acceptable risk; establish controls through policies and technical guardrails such as data boundaries, model scope, agent conditions, auditability, and rollback paths; and finally, drive enablement by providing approved tools, patterns, and sandboxes that meet compliance and sovereignty needs, allowing teams to innovate safely and at speed.

Data protection's evolving role

As systems grow more interdependent, data protection must evolve from simple point-in-time snapshots to restoring an entire coherent state. Achieving AI resilience requires immutability by design — so no agent or privileged user can tamper with backups — alongside vendor-independent storage that keeps data where policy dictates.

Open, cost-predictable access to the platform is essential for safe analytics and training, supported by anomaly detection and time-machine history to distinguish "good" from compromised objects. Recovery must also become more precise, with guided, fine-grained restores that roll back only what's necessary.

To stay resilient, organizations must demand immutability, sovereignty, anomaly detection, and granular recovery from their data-protection vendors. Ultimately, winning with AI isn't about adopting it everywhere, but about maintaining control. The companies that define outcomes, enforce guardrails, and ensure recoverability will be best positioned to innovate boldly, absorb missteps, and bounce back fast.

Keep an eye out for these emerging offerings…

- **Cross-workload intelligence.** Unifying SaaS silos in a single immutable platform enables anomaly detection, discovery, and selective access for model training.
- **Guided recovery.** Conversational restore assistants that leverage detailed data-change history to minimize user effort in complex incidents (e.g., restoring only the 5% of documents altered by ransomware across 1,000 users).

Intelligent data governance

Read our data governance report for more insight into how to best classify, protect, and recover your data so you can ensure operational continuity.

Download now

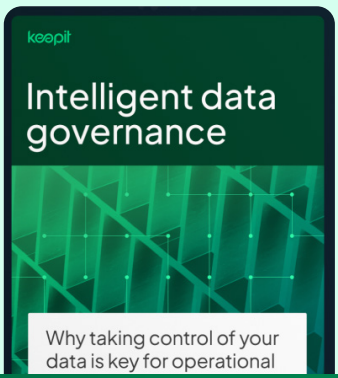

Keepit provides a next-level SaaS data protection platform purpose-built for the cloud by securing data in a vendor-independent cloud to safeguard essential business applications, boost cyber resilience, and future-proof data protection.

For more information visit www.keepit.com or follow Keepit on Linkedin