TAG

ANALYST REPORT

THE ACRONIS APPROACH TO AI ENABLEMENT AND GENAI PROTECTION

DR. EDWARD AMOROSO,
CEO, TAG
RESEARCH PROFESSOR, NYU

Acronis

THE ACRONIS APPROACH TO AI ENABLEMENT AND GENAI PROTECTION

DR. EDWARD AMOROSO, CEO, TAG,
RESEARCH PROFESSOR, NYU

This TAG Research as a Service (RaaS) report provides an independent technical and operational assessment of cybersecurity and data protection vendor Acronis. The report examines how Acronis not only leverages AI across its cyber protection platform, but also how it addresses the growing requirement to protect enterprise use of generative AI. Particular emphasis is placed on the AI capabilities embedded in the Acronis platform and on Acronis GenAI Protection, the service that anchors the company's GenAI protection strategy. The objective of this report is to provide practitioners, service providers, and enterprise leaders with a clear understanding of the platform's architecture, intended use cases, and operational value. We hope the report is useful to readers working in this important area.

INTRODUCTION

Acronis is well-known for its world-class cyber protection, which combines cybersecurity, data protection, infrastructure management, service automation, and cloud infrastructure into a unified platform. This positioning becomes especially relevant as organizations begin integrating generative AI into everyday workflows. The most common objective is to leverage AI to improve productivity, while also managing the risks introduced by AI systems, including leakage, misuse, and manipulation.

In this context, Acronis has extended its platform strategy into the GenAI domain. Rather than treating AI as a separate, bolted-on feature set, the company is embedding AI capabilities and governance directly into its existing managed cyber protection architecture. This approach aligns with a broader operational reality observed across enterprise and service provider environments, namely, that new capabilities must integrate into existing control planes to be effective.

PLATFORM OVERVIEW: ACRONIS AI-ENABLED CYBER PROTECTION

Of particular note, Acronis is delivering AI across two complementary dimensions on a single platform. First, the company embeds AI into the day-to-day operation of cyber protection — accelerating threat response, remote management, automation authoring, service desk operations, and operational decision-making for service providers and their technicians. Second, Acronis provides governance and protection controls over how generative AI itself is used in business settings, through Acronis GenAI Protection.

Underlying both is the Acronis Cyber Platform. With one console and one agent, it offers service providers a single platform from which to deliver a variety of services and gain complete visibility of security posture across client tenants. Such foundational leverage gives Acronis an advantage over many of the newer startup entrants offering AI security support, because AI enablement and AI protection are introduced into a control plane that partners already operate at scale.

AI ENABLEMENT IN THE ACRONIS PLATFORM

In addition to protecting against the consequences of rogue or non-compliant AI usage, Acronis leverages AI technology across multiple dimensions of its platform to improve operational efficiency and effectiveness. The most recent platform releases extend these capabilities directly into the workflows that service providers rely on every day — from operational insight and automation authoring, through live remote support and incident triage, to the service desk itself.

Acronis AI — operational visibility and AI-driven onboarding. Acronis AI lets partners ask natural-language questions about tenants, users, workloads, and alerts in one place. It helps partners answer everyday operational questions faster, including identifying trial tenants that may need onboarding support, checking service and resource usage, reviewing key user details such as roles and 2FA status, and surfacing critical alerts and unhealthy or unprotected workloads — without digging through multiple screens. Acronis has since extended the assistant into an AI agent that guides partners through common onboarding tasks, such as creating customer tenants and users directly in the console, using context from tenant and user data and offering deep links to the relevant screens. The value here is contextual analysis applied to real operational data, reducing the time technicians spend interpreting and acting on the state of an environment.

AI-created automation workflows. Partners can now build and manage automation workflows using natural language. An AI assistant in the Workflow Automation builder creates, edits, and explains workflows from plain-language descriptions, recommends actions based on the selected trigger, and applies automatic validation and correction suggestions to improve reliability. Instead of configuring automations step by step, a partner can describe the intended logic — for example, "create a ticket and notify the on-call technician when a backup fails" — and refine the generated workflow with text commands. This lowers the barrier to automation, shortens onboarding for new users, and helps partners scale consistent, reliable automation across customer environments.

Acronis AI for remote desktop — session summaries and in-session recommendations. Acronis AI automatically generates comprehensive summaries of remote desktop sessions, capturing the actions taken and making recordings easier to review with time-coded navigation. Sessions are automatically categorized by issue type with a resolution status, reducing manual documentation and improving technician handoffs, audits, and training. Complementing this, Acronis AI provides real-time, context-aware troubleshooting recommendations during live sessions, helping technicians work through issues step by step and reducing the amount of senior engineer time needed to guide junior staff through complex cases.

AI-guided remediations via remote command line. Extending AI assistance into hands-on remediation, technicians can now receive AI-recommended remediation commands based on alert context or a plain-language description of an issue. Critically, a human-in-the-loop approval is required before any command is executed remotely, and all execution details are written to the audit log. This accelerates alert resolution and reduces the research effort required to write remediation commands, while preserving the control and accountability that service providers require.

AI Triage for EDR incidents. On the security side, Acronis applies AI to automate alert analysis and incident interpretation. AI Triage delivers instant, AI-generated analysis of EDR incidents — including an executive summary, MITRE ATT&CK mappings, and prioritized remediation steps — available on demand at no extra cost. Findings can be exported as a professional, shareable PDF, enabling clear customer communication and effective response even when no analyst is on call.

AI Service Desk (early access). Releasing in July 2026 as an early-access capability, the Acronis AI Service Desk extends AI into ticketing operations, automating ticket resolution with pre-built agents so MSPs can close tickets faster. In its current early-access scope, the AI Service Desk is focused on alert-to-ticket automation, generating possible root causes based on collected details and previous ticket history, and providing possible remediation steps drawn from prior tickets, existing documentation, and established best practices. The intent is to reduce manual triage effort and accelerate time to resolution, with additional agent capabilities expected to follow as the service matures.

Taken together, these capabilities apply AI to operational workflows that extend well beyond traditional security functions — generating summaries of activity, authoring automations, supporting decision-making through contextual analysis, triaging incidents and tickets, and guiding remediation in real time with human oversight. The integration of these capabilities into the overall platform reduces the need for separate tools and helps streamline operational processes. The result is a platform that not only secures systems, but also contributes directly to productivity and service delivery.

AI PROTECTION AND GOVERNANCE: ACRONIS GENAI PROTECTION

Perhaps the most significant aspect of the Acronis AI strategy is its governance layer, delivered as Acronis GenAI Protection. Generative AI is spreading fast across organizations — often through unsanctioned, consumer-grade tools — creating new attack surfaces, compliance exposure, and data leakage risk that businesses expect their service providers to control. Natively integrated into the Acronis platform, GenAI Protection provides visibility and control over generative AI usage across Windows endpoints, blocks sensitive data sharing, and detects harmful prompts. It supports browser-only or full desktop protection modes.

Shadow AI visibility and governance. One of the primary concerns addressed by the platform involves unsanctioned AI usage. Organizations frequently lack visibility into which AI tools are being used and accessed. Acronis discovers and monitors browser- and desktop-based GenAI tools across the environment, enabling administrators to identify trends and risk exposure, detect anomalies, and align usage with organizational policy. This visibility can be combined with enforcement — including URL filtering through EDR/XDR — to restrict access to unapproved AI applications and keep adoption aligned with policy.

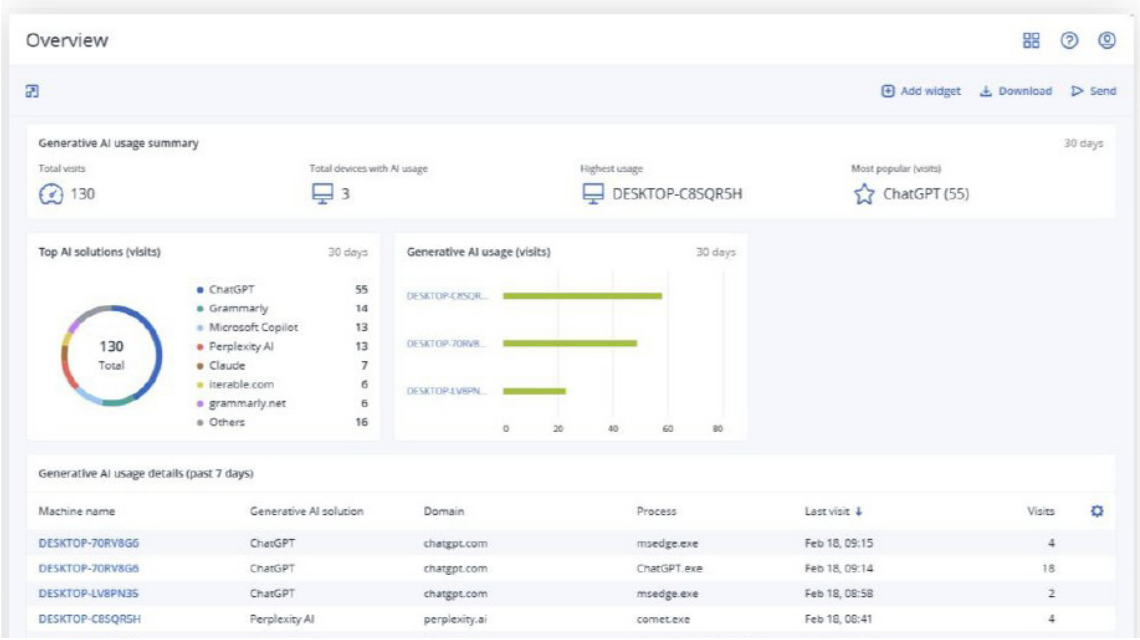

Figure 1. Acronis GenAI Protection Dashboard for GenAI Detection and Governance

Sensitive data protection for AI interactions. Another critical capability is the protection of sensitive data within GenAI chats. Users routinely paste data into prompts to accelerate their work, risking exposure of PII, PHI, credentials, financial, and confidential business information. Acronis GenAI Protection inspects user prompts and file uploads for sensitive content and blocks unauthorized submissions to public or unsanctioned AI services. By applying detection and enforcement policies, organizations can reduce the likelihood of inadvertent or deliberate data exposure through AI systems — turning "safe AI" from a policy statement into an enforceable, managed service.

Harmful prompts and AI abuse prevention. Acronis also addresses the emerging risk of prompt manipulation and abuse. As AI systems become more integrated into workflows, they become susceptible to adversarial inputs designed to alter behavior or bypass safeguards. Prompt injection and other harmful techniques are emerging as practical attack vectors that can degrade output integrity and create downstream security and compliance risk. The platform analyzes prompts and associated context for potentially harmful or malicious content and blocks abusive prompts, helping to mitigate these risks before they can affect downstream processes — without adding another standalone enterprise tool.

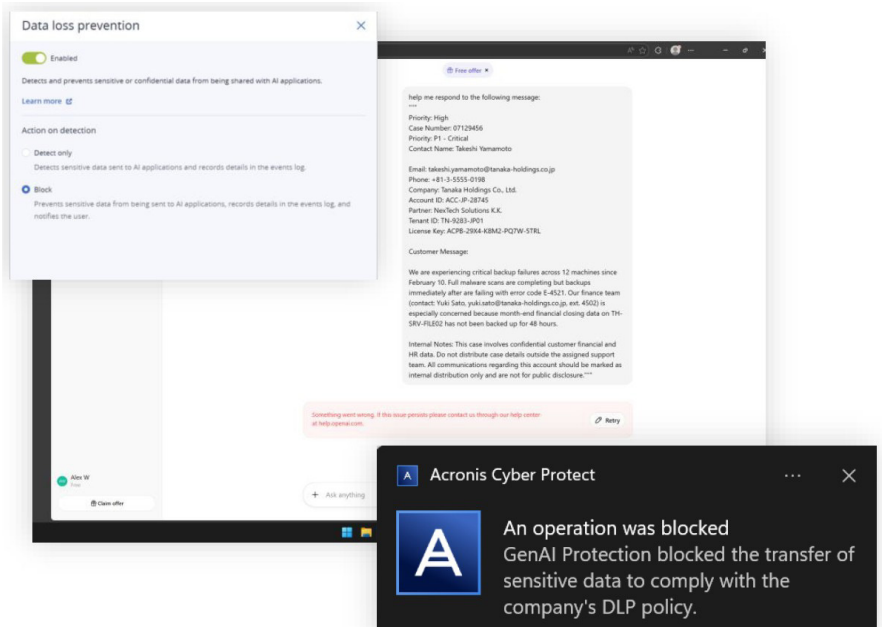

Figure 2. Acronis GenAI Protection Prompt Injection Protection in Action

MCP server monitoring. Extending governance to the agentic AI layer, GenAI Protection now includes Model Context Protocol (MCP) server monitoring, giving partners visibility into where and how MCP servers are being used across managed client devices. This enables partners to identify external AI tools and services in use, assess potential security and supply-chain risks, and improve oversight of AI-driven workflows. The capability supports discovery of MCP server usage, investigation of suspicious or risky activity, and drill-down into detailed activity data for security analysis.

Logging, event context, and reporting. Finally, the platform provides logging and reporting capabilities that support governance and accountability. By capturing detailed information about AI interactions, policy violations, and enforcement actions, Acronis enables organizations to produce evidence of control effectiveness. The GenAI Protection events log lets partners and customers search, filter, and investigate events with richer context — including user details, device information, and the specific content that triggered a detection — making data-loss and prompt-injection events easier to attribute, and helping teams confirm sensitive material, identify false positives, and refine enforcement policies. This is essential for internal governance, customer reporting, and alignment with emerging regulatory expectations around AI usage. For broader coverage beyond GenAI applications, Acronis DLP extends protection across 70+ local and network channels, helping organizations achieve a more comprehensive data protection posture.

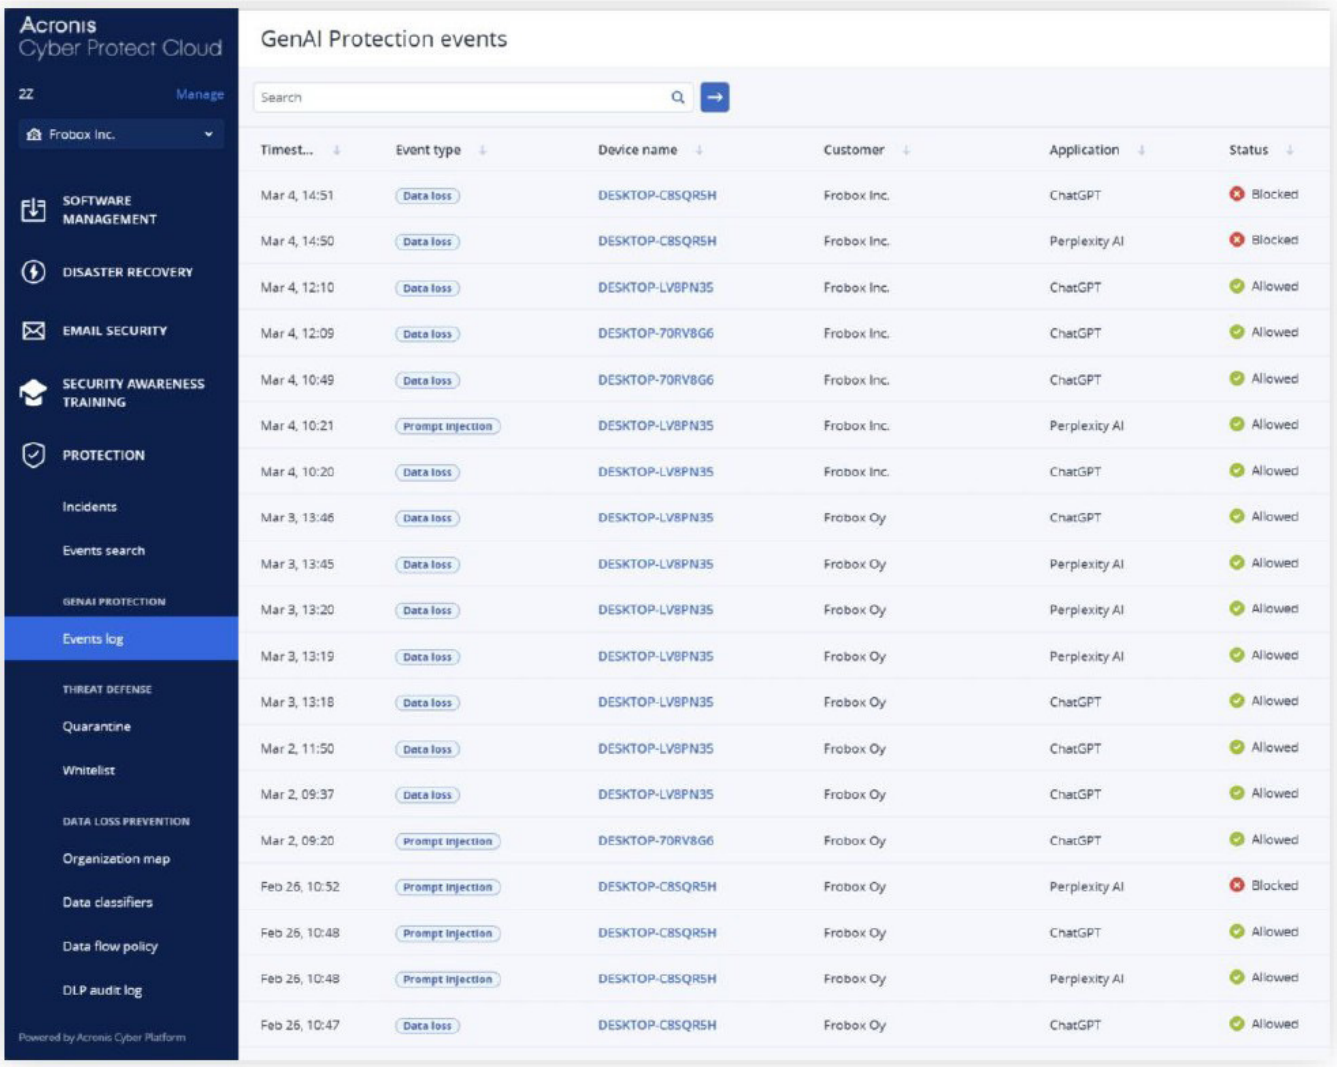

Figure 3. Acronis GenAI Protection Event Log

ANALYST OBSERVATIONS

From a TAG perspective, the Acronis approach to AI represents a practical and well-aligned response to the challenges of AI adoption by businesses of all sizes. The most notable strength of the platform is its integration into an existing cyber protection architecture. By embedding AI capabilities and controls into a unified operational framework, Acronis avoids the fragmentation that often limits the effectiveness of new technologies. The balance is deliberate: AI is used to make protection and service delivery faster and more consistent, while GenAI Protection ensures that the organization's own use of AI does not become a new source of risk.

We also note the deliberate design choice to keep a human in the loop where it matters most — AI-guided remote command remediations require explicit approval and are fully audited — which reflects an enablement model built for accountable service delivery rather than unchecked automation. The early-access AI Service Desk, AI-created workflows, and AI Triage for EDR incidents indicate a clear trajectory toward agent-assisted operations across the partner workflow.

We believe the long-term success of the platform will depend on execution in several areas. As AI usage evolves, the complexity of interactions, data flows, and attack vectors will increase. The early extension of GenAI Protection into MCP server monitoring is an encouraging signal that Acronis intends to track the agentic AI threat surface rather than only generative AI. As newer capabilities such as the AI Service Desk move from early access to general availability, Acronis will need to continue expanding its detection, automation, and control capabilities to address these changes. In addition, the value of the platform will be closely tied to the quality of its reporting and its ability to provide actionable insights to administrators and service providers.

FINAL NOTE

For organizations seeking to move from informal experimentation with AI to structured and controlled adoption, the Acronis platform offers a coherent vision and approach that combines AI enablement and AI protection in a managed, automated way. From a TAG standpoint, this balance between enablement and protection is essential, and Acronis appears to be aligned with that requirement as the enterprise AI landscape continues to evolve and service providers urgently need to establish their role in the AI value chain.

ABOUT TAG

Recognized by Fast Company, TAG is a trusted next generation research and advisory company that utilizes an AI-powered SaaS platform to deliver on-demand insights, guidance, and recommendations to enterprise teams, government agencies, and commercial vendors in cybersecurity and artificial intelligence.