

AI-ASSISTED VULNERABILITY DISCOVERY

COMMVAULT'S SECURITY EVALUATION APPROACH

Customers are not asking whether AI is a threat. They are asking whether their vendors are disciplined about how they use it — and that is a question about process, not technology.

Across the enterprise security industry, AI and large language models are being applied to vulnerability discovery — helping security teams evaluate more code, explore more attack paths, and identify exploitable conditions more efficiently. This is not a niche experiment: it is an industry-wide evolution in security engineering practice that is changing both the capability of offensive security research and the expectations customers place on software vendors.

Customers are asking questions that reflect genuine due diligence, not speculation. They want to know whether their software vendors are evaluating their own products against the same methods that adversaries may eventually use — and whether the processes for discovering, managing, and remediating those vulnerabilities are mature enough to scale with the pace of discovery.

These are the right questions to ask. Commvault's position is straightforward: we are already engaged, the work is structured and controlled, and our response processes are built on established security engineering practice — not on any single tool or model.

COMMVAULT'S SECURITY PHILOSOPHY AND OPERATIONAL APPROACH

Commvault's security posture is built on process maturity rather than dependency on any single tool, model, or vendor. Vulnerability management at Commvault follows an established, risk-based framework: findings are assessed for practical exploitability, prioritized by severity and exposure, and remediated through standard SDLC processes. That framework applies equally whether a finding originates from a static analysis tool, a penetration test, an external researcher, or an AI-assisted evaluation.

AI-assisted vulnerability discovery is integrated into this framework as an additional capability — not a separate program. Candidate findings generated through AI-assisted methods are treated as inputs that require human confirmation of exploitability before any remediation action is taken. This prevents both under-prioritization of genuine risk and over-investment in false positives.

© 2026 Commvault

GOVERNANCE AND CONTROL PRINCIPLES

- All AI-assisted evaluation activity is conducted in Commvault-controlled environments, either air-gapped local hardware or Commvault-managed cloud infrastructure
- Source code does not leave Commvault-controlled boundaries under any execution path
- Vendor and tooling access is governed by NDA and formal engagement terms
- Findings are processed through the same security engineering review pipeline applied to all vulnerability sources
- No AI-generated finding is actioned without human triage and exploitability confirmation

HOW AI-ASSISTED VULNERABILITY DISCOVERY IS EVOLVING

The application of large language models to security testing is an active area of development across the industry. Security engineering teams, specialist researchers, and dedicated vulnerability discovery programs are increasingly using AI models to assist with source code analysis, attack surface enumeration, and the identification of vulnerability classes that benefit from pattern recognition at scale.

This evolution does not make existing security practices obsolete. Static analysis, fuzz testing, manual code review, third-party penetration testing, and established scanning tooling remain essential components of a comprehensive security program. What AI-assisted discovery adds is coverage depth — the ability to evaluate a broader set of code paths, model more complex exploit conditions, and surface findings that require contextual understanding rather than simple pattern matching.

The practical implication for enterprise software vendors is that security evaluation programs need to expand to include AI-assisted methods alongside existing practices, not in place of them. Commvault is actively building that capability in a structured, governed way across multiple evaluation workstreams.

AI changes how thoroughly vulnerabilities can be found — not whether existing security processes matter. Commvault's evaluation program expands coverage without replacing the engineering disciplines that have always defined responsible vulnerability management.

CURRENT EVALUATION INITIATIVES

Commvault is currently executing structured AI-assisted security evaluation across two active workstreams, with a third in preparation. All operate under the same governance framework: Commvault-controlled environments, formal vendor engagement, and findings processed through standard security engineering review.

WORKSTREAM	STATUS	TARGET
Multi-model white box evaluation	Active — scan cycle in progress	Ongoing
Targeted supplementary scanning	Active — selected application components	Ongoing
Expanded model coverage (incl. Project Glasswing)	Pending	Coming Soon

© 2026 Commvault

STRUCTURED MULTI-MODEL WHITE BOX EVALUATION

Commvault is engaged with a specialist cybersecurity firm focused on vulnerability research, product evaluation, reverse engineering, and application security assessment, to conduct multi-model white box security evaluation against Commvault source code. The scope of this engagement is source-code-informed evaluation in controlled environments using a range of AI models, including both open-source and frontier commercial models.

Execution controls

- Evaluation conducted on isolated hardware and within Commvault-managed cloud infrastructure
- Air-gapped execution paths for open-source model evaluation
- Frontier cloud models accessed within Commvault-controlled VPC infrastructure
- Source code does not leave Commvault-controlled boundaries
- Vendor access governed by formal NDA and engagement terms

Models in scope

The current evaluation spans multiple open-source and frontier models, including Qwen, GPT OSS, Mistral, and NVIDIA-hosted models, with additional models being sequenced as access is confirmed. Commvault is also on the waiting list for Project Glasswing — a controlled-access defensive security initiative offering vetted organizations early access to advanced model-assisted vulnerability discovery. Validated findings will enter standard remediation workflows. Subsequent evaluation phases, including expanded model coverage, will be sequenced following completion of the current cycle.

TARGETED SUPPLEMENTARY SCANNING

In parallel, Commvault is conducting targeted model-assisted scanning against selected application components to evaluate coverage depth and testing behavior in specific areas. This work uses commercially available AI-assisted code analysis tooling integrated into existing security workflows, focused on identity management layers and user-facing application components.

The objective of this workstream is to identify coverage gaps not addressed by existing tooling and to evaluate where AI-assisted scanning can complement — and be sustainably integrated into — Commvault's standard security engineering practice.

Commvault's evaluation program is tool-agnostic and model-agnostic by design. No single vendor or model is a dependency; new approaches can be added as they prove out, without re-architecting how findings are governed or remediated.

HOW COMMVAULT EVALUATES AND MANAGES VULNERABILITIES

Commvault's vulnerability management process applies a consistent risk-based framework regardless of how a finding originates. AI-assisted evaluation feeds into the same triage and remediation pipeline as penetration testing, static analysis, bug bounty submissions, and external researcher disclosures.

© 2026 Commvault

TRIAGE AND PRIORITIZATION

Findings generated through AI-assisted evaluation are treated as candidates, not confirmed vulnerabilities. Each finding is assessed by the security engineering team for practical exploitability in realistic customer environments. Severity ratings are assigned based on exposure, exploitability, and impact — not on the novelty of the discovery method. This prevents both missed risk and the operational cost of chasing false positives.

REMEDIATION PROCESS

Confirmed vulnerabilities are remediated through standard SDLC processes with priority assigned based on severity rating and exposure profile. The same timelines and escalation paths apply as for any other vulnerability source. Findings are not disclosed externally until remediation is in place, at which point standard CVE and security advisory processes apply where applicable.

WHAT THE CURRENT EVALUATION HAS FOUND

The current evaluation cycle has identified real exploitable issues that were not surfaced by existing tooling. Those findings are being processed through standard remediation workflows. This outcome is consistent with the premise of the evaluation program: AI-assisted methods expand the coverage of vulnerability discovery, particularly for vulnerability classes that benefit from contextual code analysis.

WHY OPERATIONAL READINESS MATTERS MORE THAN SPECIFIC TOOLS

As AI-assisted vulnerability discovery becomes more widely used across the industry, the volume of potential findings that security engineering teams will need to evaluate is expected to increase. The critical question for enterprise software vendors is not which AI model they use — it is whether their vulnerability management process is mature enough, and scalable enough, to handle increased discovery throughput without creating a backlog that increases customer exposure.

Commvault's investment in AI-assisted evaluation is paired with an equivalent focus on the process infrastructure needed to act on findings: triage capacity, severity prioritization frameworks, remediation tracking, and coordinated disclosure practices. The evaluation capability is only as valuable as the response capability behind it.

Scalability and governance priorities

- Risk-based triage that scales with discovery volume without compromising prioritization accuracy
- Consistent remediation SLAs applied across all vulnerability sources, regardless of discovery method
- Governance controls that allow new AI models and evaluation methods to be added within a consistent operational framework
- Audit trails and documentation that support both internal accountability and external disclosure obligations

The advantage in AI-era security is not access to a specific model. It is having a vulnerability management process disciplined enough to respond effectively to what any Frontier AI models surface — at scale, consistently, and without operational gaps.

© 2026 Commvault

FAQ FOR FRONTIER AI CONVERSATIONS

When your auditors, management, board, and non-technical colleagues raise questions about AI-assisted security testing, they are rarely asking for a technical deep-dive on model architectures. They are performing due diligence on whether their software vendors are keeping pace with how the industry's threat and testing landscape is evolving, and whether the processes behind that effort are ones they can trust.

The most effective response is a process-grounded one. Commvault is not positioning this work as a marketing moment — it is describing how our security engineering practice is expanding to incorporate a new class of evaluation tooling within the same governance framework that governs all security testing.

IF ASKED: "WHAT IS COMMVAULT DOING AROUND AI-ASSISTED SECURITY TESTING?"

Commvault is actively evaluating our products using AI-assisted methods as part of our structured security engineering program. That work is conducted in Commvault-controlled environments, follows the same vulnerability management process as all other testing, and includes governance controls that prevent source code from leaving our infrastructure. Commvault is doing this now — it is not a roadmap item.

IF ASKED: "HOW IS COMMVAULT PREPARING FOR AI-DRIVEN VULNERABILITY DISCOVERY?"

Commvault built a program that is model-agnostic and tool-agnostic by design. Their goal is not to depend on a single vendor or approach — it is to ensure that their security engineering practice can incorporate the best available methods across a range of AI tooling within a consistent governance and risk management framework.

IF ASKED: "IS COMMVAULT USING THESE MODELS SAFELY?"

Yes. **All AI-assisted evaluation activity is conducted on isolated hardware or within Commvault-managed cloud infrastructure**. Source code does not leave Commvault-controlled environments under any execution path. Vendor and tool access is governed by formal NDA and engagement terms. AI-generated findings require human confirmation of exploitability before any remediation action.

IF ASKED: "HOW IS COMMVAULT SCALING VULNERABILITY MANAGEMENT?"

Commvault's focus is on ensuring the response process scales with discovery volume. AI-assisted evaluation will increase the number of potential findings our teams need to evaluate. They are investing in risk-based triage, consistent remediation SLAs, and the operational infrastructure needed to act on increased discovery throughput without creating exposure gaps.

Commvault is prepared, disciplined, and operationally mature. They are not reacting to a trend — they are extending a structured security engineering practice to incorporate new capabilities in a governed way.

© 2026 Commvault

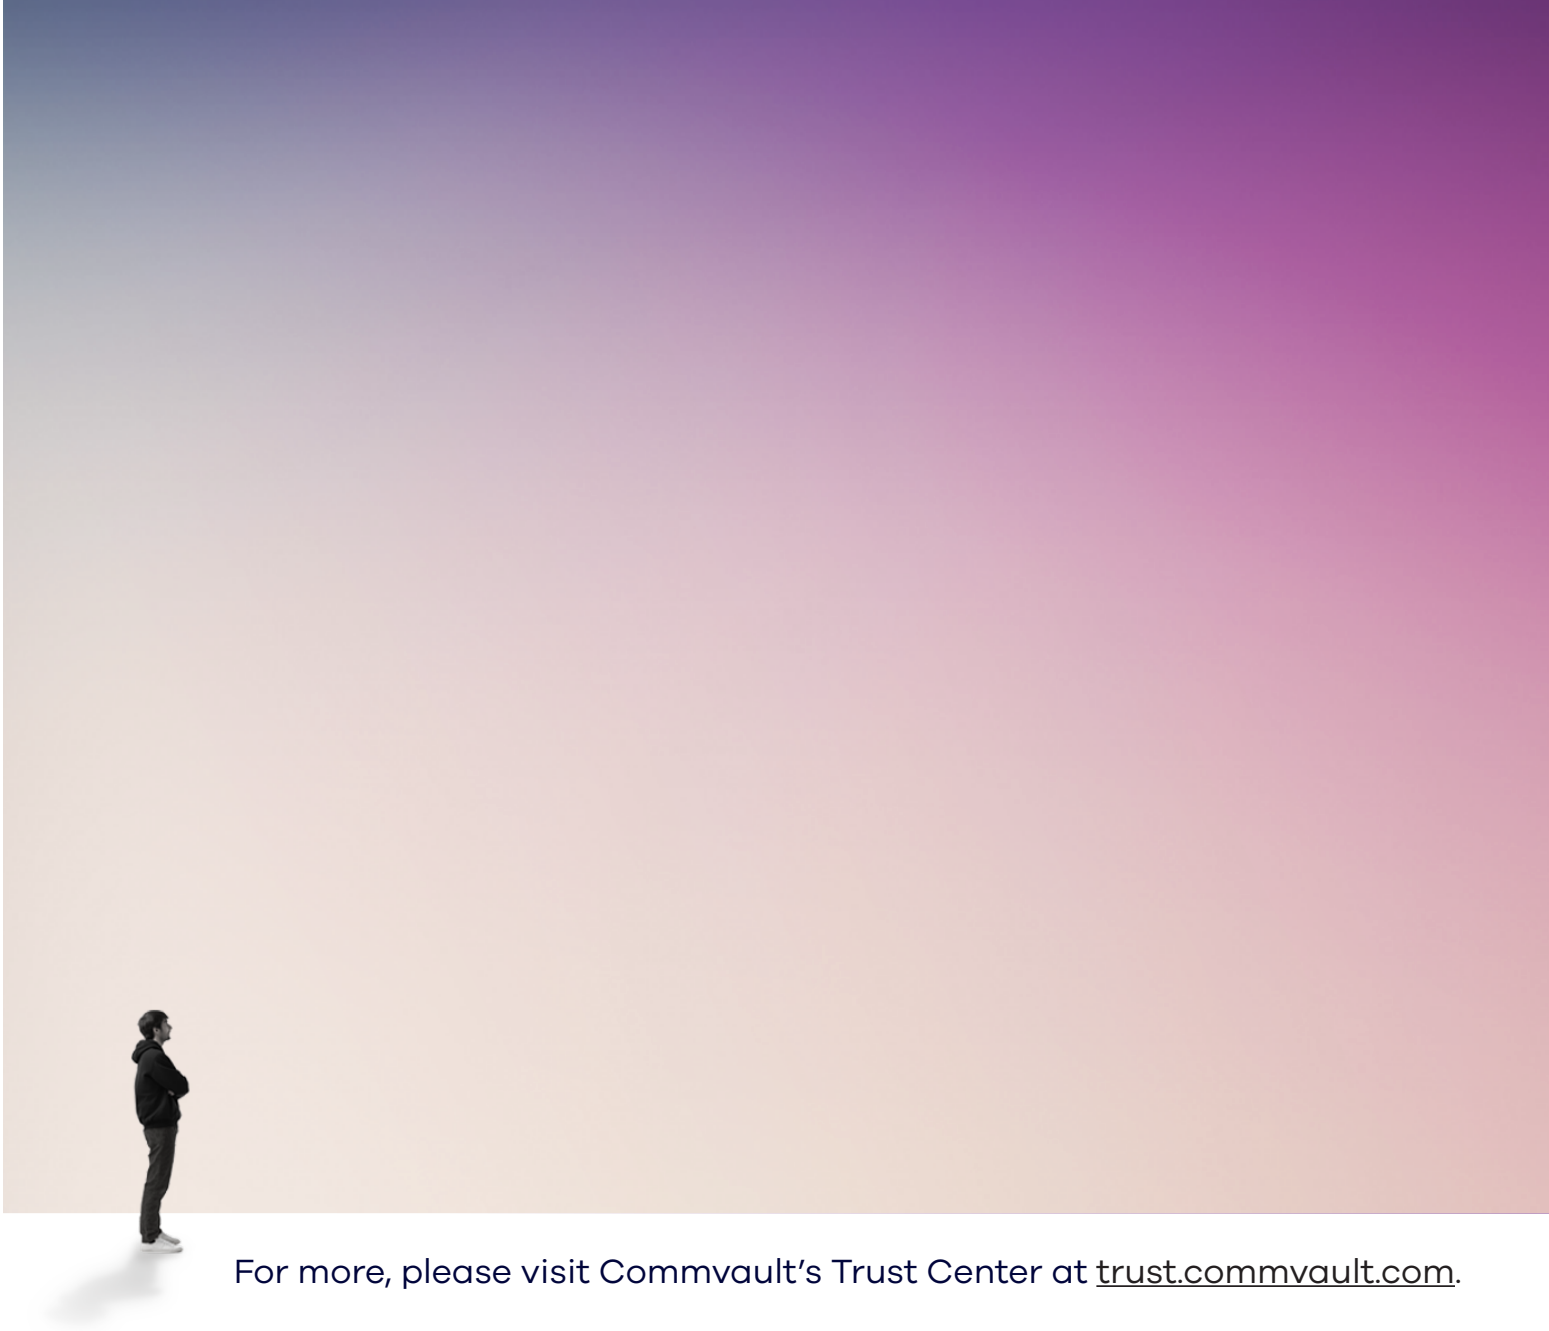

For more, please visit Commvault's Trust Center at trust.commvault.com.

Commvault®

commvault.com | 888.746.3849 | get-info@commvault.com

© 2026 Commvault. See here for information about our trademarks and patents. 06_26