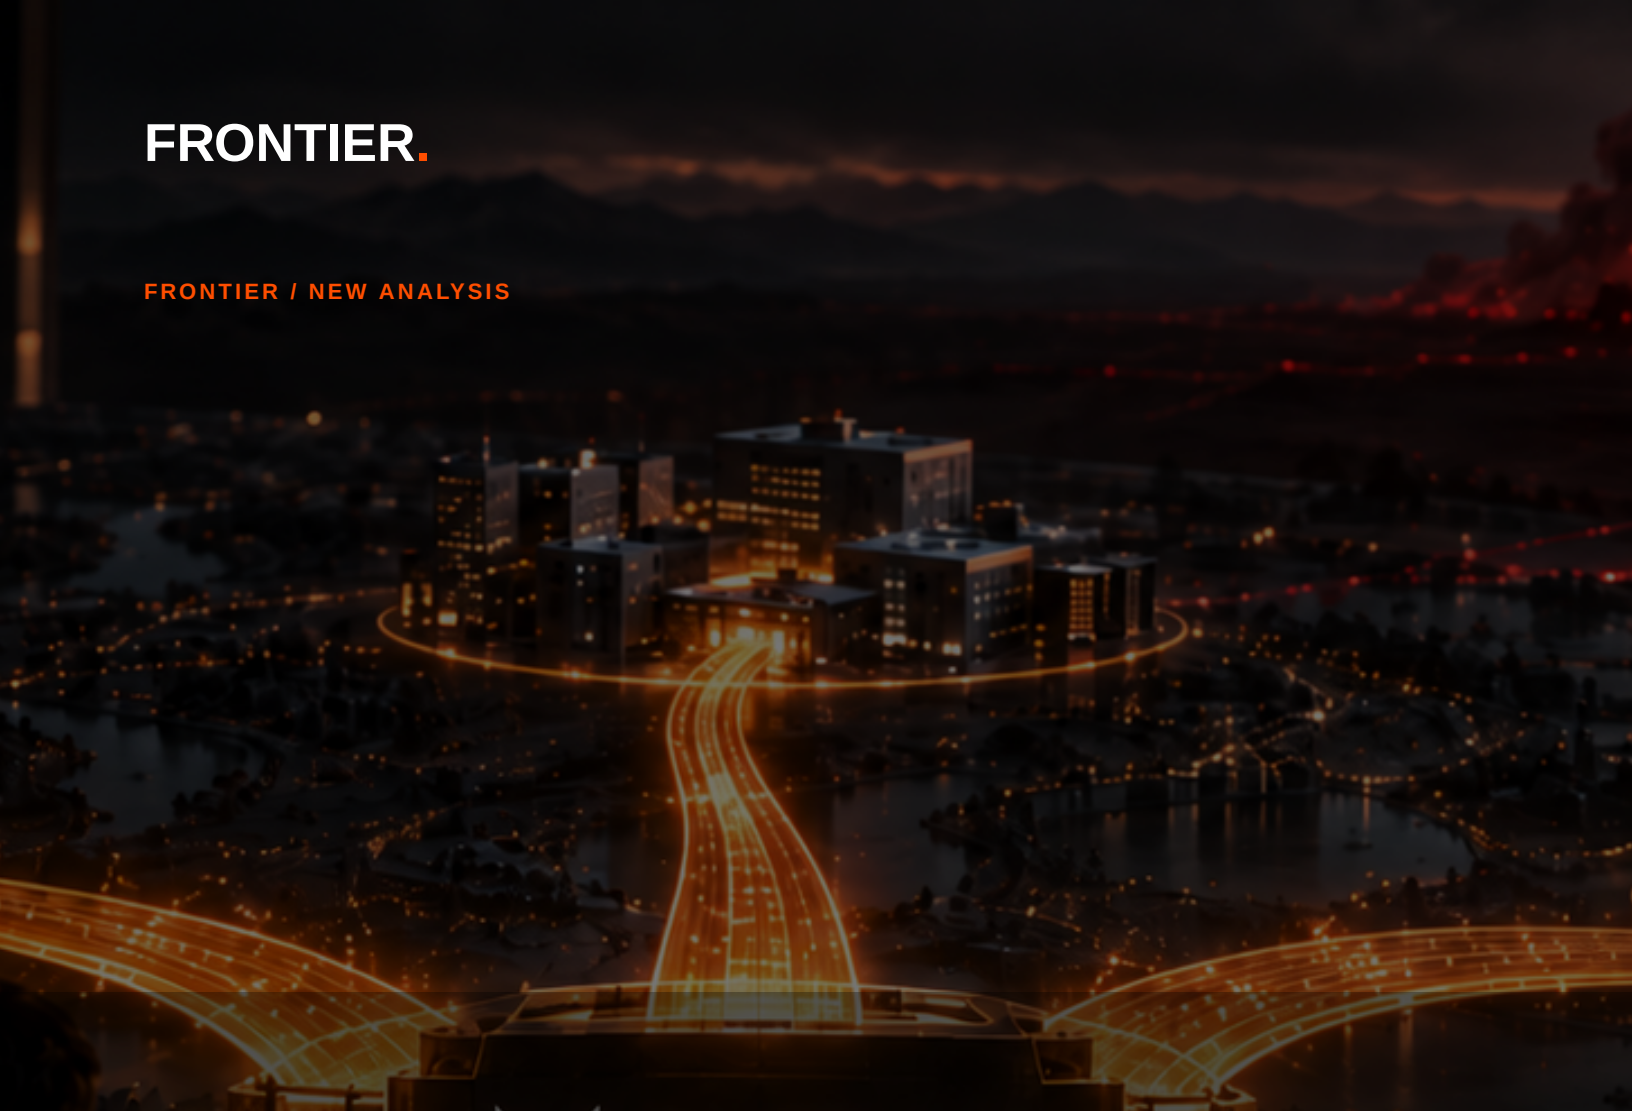
FRONTIER.

FRONTIER / NEW ANALYSIS



Exaforce Gives Every Defender an AI Team. Who Commands the Counterattack?

ExaU points security operations toward persistent AI teams; CIOs must prove that more investigative capacity produces faster, justified action without multiplying privileged risk.

FRONTIER / Frontier Research / November 16, 2026

NEW ANALYSIS

Built to change a decision.

ExaU points security operations toward persistent AI teams; CIOs must prove that more investigative capacity produces faster, justified action without multiplying privileged risk.

EDITION	**New Analysis**
AUTHOR	**Frontier Research**
PUBLISHED	**November 16, 2026**

This PDF is a publication rendering of the canonical Frontier News article. The web article remains the live source of truth; analysis, figures, and source citations travel with this edition.

What changed

The most dangerous queue in cybersecurity is the work nobody has time to start.

The suspicious access pattern that never becomes a hunt. The exception that nobody revisits. The question a CISO takes into a board meeting without a defensible answer. Automating the alert inbox helps. It does not automatically close those gaps.

Exaforce is targeting that unfinished work with ExaU. SiliconANGLE reported its November 16 launch and immediate availability within the Exaforce platform: practitioners describe jobs in natural language or use templates, creating personal agent teams that retain instructions and decisions, coordinate work, and can be shared across the security organization. The reported boundary is consequential: investigations run autonomously; actions affecting people, systems or communications require approval, with steps recorded for audit.[1]

Exaforce named ExaU in its November 14 Force ER announcement, describing custom Exabots assigned continuing roles after incident response. That is the primary event explored in today's coverage.[2]

This lands inside an established platform effort. Exaforce's May 12 financing announcement disclosed a $125 million Series B and $200 million in total funding. CEO Ankur Singla described the ambition as “the platform defenders actually work in.” The same release featured customer accounts from Invisible and Guardant Health. Those support the broader platform's operating history; they do not establish outcomes for newly announced ExaU teams.[3]

The new proposition deserves attention because it changes the unit of delegation: from an individual task to an ongoing responsibility. Frontier's question is what happens when a security practitioner can commission a workforce faster than the enterprise can define its authority.

Why it matters

Exaforce's underlying architectural bet is that useful security reasoning begins before an analyst asks a question. Its data platform describes correlating logs with identity, configuration, code and file metadata, rather than repeatedly reconstructing those relationships during each investigation. It also explicitly supports coexistence with an existing security information and event management system, or SIEM.[4]

That matters more than the conversational interface. An agent given an apparently simple job—investigate an unusual download—needs to establish which identity acted, which permissions applied, what was accessed and whether the behavior fits the surrounding activity. More agents querying disconnected systems can multiply cost and contradictions. Shared, current context can make parallel investigation productive. That is Frontier's architectural assessment, not a measured ExaU performance result.

The first potential gain is coverage of obligations that currently disappear between tools and teams. Consider an illustrative deployment: one agent tracks unresolved privileged-access exceptions, another checks whether agreed remediation actually happened, and a third assembles the evidence for an executive briefing. None should need a dramatic new alert to justify doing its job. The value comes from maintaining attention over time.

A useful pilot would therefore measure overdue obligations closed, coverage added and analyst rework—not just minutes saved per alert. Faster triage is easy to demonstrate while difficult follow-through remains untouched. CIOs should resist buying a more impressive dashboard for the same unresolved exposure.

The second gain is continuity. ExaU is reported to retain decisions and exceptions, which could reduce repeated explanation across investigations.[1] Our interpretation is that this makes memory an operational asset with a lifecycle. A temporary exception must expire. A departed analyst's judgment must have a successor owner. An assumption accepted during an outage must not quietly become permanent security policy.

OWASP identifies memory poisoning, indirect prompt injection, tool abuse and data exfiltration among agent security risks. Its guidance calls for adversarial testing when prompts, tools, memory, retrieval, policies or model providers materially change.[5] This is relevant to any agent working with hostile material. It is not evidence of an ExaU vulnerability.

The third gain is coordination—and its corresponding concentration of risk. A shared interpretation can reduce handoff loss. It can also make several agents repeat the same mistaken premise. Three agreeing agents are not three independent witnesses when all three rely on one flawed identity mapping. The control that matters is access to underlying evidence and a way to challenge the conclusion.

Exaforce's broader response product offers a useful comparison. Exabot Respond documents workflows mixing autonomous steps and human approvals, with configurable gates and timeout handling. That is broader platform behavior, not proof that ExaU permits every one of those modes. Buyers should require a demonstration of the exact controls in the product and edition they are procuring.[6]

Approval quality is the next bottleneck. If agents generate more decisions than humans can properly review, the organization has moved its queue, not removed it. A meaningful approval should show the affected identity, exact action, evidence, scope, expected business impact and recovery path. A button labeled approve does not create informed human control.

Finally, the platform's expansion makes the cyber stakes unusually direct. Exaforce's September 15 AI Security announcement describes bringing AI activity into the same identity, endpoint, cloud and SaaS context, including mapping agents to owners and the resources their credentials reach.[7] The strategic opportunity is to defend an enterprise increasingly operated by AI with a security function that can keep up. The danger is exempting the defenders' own agents from the discipline applied to everyone else's.

Frontier take

ExaU is an interesting signal that security automation is moving toward persistent ownership of work. The opportunity is substantial: a thinly staffed team could maintain more investigations, revisit more assumptions and answer more executive questions without forcing every task through the same overworked people.

But the winning enterprise design will make responsibility durable and authority bounded. Every agent needs an accountable owner, an explicit remit and a defined end to its permissions. “Personal” should describe how easily a practitioner configures useful assistance. It should not mean invisible to the enterprise.

The strategic metric is the time from a credible signal to a justified, authorized action. That includes evidence gathering, disagreement, approval and execution. Optimizing only the model's response time misses the places where security operations actually stall. It can even make a dangerous recommendation arrive more efficiently.

OWASP's excessive-agency guidance locates the underlying problem in excessive functionality, permissions or autonomy, including failures involving compromised peer agents.[8] Frontier's inference is that a security agent must be evaluated as both a defender and a privileged software actor. Its good intentions are not a security boundary.

There is also a procurement consequence. Persistent assignments and institutional memory could make the platform more valuable over time—and harder to leave. CIOs should ask which instructions, decisions, evidence links and histories can be exported in usable form. A vendor exit that preserves raw logs but loses the accumulated rationale behind security decisions is an incomplete exit.

The evidence today warrants a controlled evaluation, not a declaration that autonomous security has arrived. We located Exaforce's direct ExaU reference and broader technical material, while detailed launch capabilities and availability above are attributed to SiliconANGLE's reporting. The reviewed sources do not establish ExaU-specific pricing, regional availability, independent efficacy testing or customer outcome measurements. Those are diligence items, not assumed defects.

Exaforce's Force ER announcement provides a revealing operational lesson alongside the product news: it says contracts, contacts and log access consume critical early response time, and describes agreeing who can approve containment actions.[2] Frontier's conclusion is straightforward. Enterprises must prepare their decision machinery as seriously as their detection machinery. No model can compensate for an organization that discovers during a breach that nobody knows who may act.

The industry should welcome tools that give defenders more reach. It should demand proof that this reach closes exposure rather than creating a faster way to make consequential mistakes. An AI security team is worth buying when it leaves the enterprise harder to compromise, not merely busier.

Three moves for CIOs

1. Pilot the work your SOC keeps postponing. Choose one bounded obligation—such as checking unresolved privileged-access exceptions—and run a 30-day pilot alongside the current process. Establish a baseline for overdue items, evidence quality, analyst effort and unnecessary escalations.

   Decision trigger: Expand only when the team can demonstrate additional verified coverage without a rise in material missed findings or unjustified escalations. Keep consequential changes approval-gated during the evaluation.

   Why now: Persistent agent assignments could expose value that alert-volume demonstrations miss. Testing neglected work makes the business case concrete before a broader platform commitment.

2. Test the complete chain of authority. Ask Exaforce to demonstrate the executing identity, effective permissions, approval enforcement, audit trail and emergency stop for an ExaU workflow. Exercise a denied request, a missing approver, misleading input and a revoked credential in a controlled environment.

   Decision trigger: Do not enable a workflow to affect production systems until an authorized operator can explain and stop each consequential action. Verify ExaU's actual behavior separately from Exabot Respond's broader configuration options.

   Why now: An agent team increases the number of decisions a practitioner can initiate. Establishing authority before scaling prevents that growth from outrunning informed supervision.

3. Make agent memory an enterprise asset. Require named owners, review dates and expiry rules for persistent instructions and exceptions. Test reassignment when an analyst leaves, correction of a bad remembered conclusion, and export of decision history with its supporting evidence.

   Decision trigger: Treat any remembered instruction that changes future access, escalation or containment behavior as a controlled operational change. Require an approved disposition before extending its life or scope.

   Why now: Continuity is part of ExaU's appeal. Governing it at the start preserves that benefit while reducing stale decisions, hidden dependencies and avoidable platform lock-in.

Sources

- SiliconANGLE — Exaforce's ExaU launch, November 16, 2026 [1]
- Exaforce — Force ER announcement and early ExaU reference, November 14, 2026 [2]
- Exaforce — Data platform architecture and SIEM coexistence [4]
- Exaforce — Exabot Respond workflows, approvals and timeouts [6]
- Exaforce — Introducing Agentless AI Security, September 15, 2026 [7]
- Exaforce — $125 million Series B announcement, May 12, 2026 [3]
- OWASP — AI Agent Security Cheat Sheet [5]
- OWASP — Excessive Agency [8]

SOURCES / CITATIONS

Trace the argument.

Every source below is linked. Source numbering follows first appearance in the report.

01 **SiliconANGLE — Exaforce's ExaU launch, November 16, 2026**
https://siliconangle.com/2026/11/16/exaforces-exau-lets-security-staff-spin-up-a-personal-team-of-ai-agents/

02 **Exaforce — Force ER announcement and early ExaU reference, November 14, 2026**
https://www.exaforce.com/blogs/emergency-incident-response-no-retainer

03 **Exaforce — $125 million Series B announcement, May 12, 2026**
https://www.exaforce.com/press-release/exaforce-raises-125m-series-b-to-combat-ai-powered-attacks-with-real-time-security-reasoning

04 **Exaforce — Data platform architecture and SIEM coexistence**
https://www.exaforce.com/platform/data-platform

05 **OWASP — AI Agent Security Cheat Sheet**
https://cheatsheetseries.owasp.org/cheatsheets/AI_Agent_Security_Cheat_Sheet.html

06 **Exaforce — Exabot Respond workflows, approvals and timeouts**
https://www.exaforce.com/platform/exabot-respond

07 **Exaforce — Introducing Agentless AI Security, September 15, 2026**
https://www.exaforce.com/blogs/introducing-ai-security

08 **OWASP — Excessive Agency**
https://genai.owasp.org/llmrisk/llm062025-excessive-agency/?lang=us

FRONTIER.

TAKE THE NEXT STEP

Push this further.

Research is useful when it changes a decision. Bring us the question, the constraint, or the argument you want tested.

RESEARCH / BRIEFINGS

inquiries@thisisfrontier.ai

Briefings, advisory work, executive discussions, and research inquiries.

REPRINTS / LICENSING

reprints@thisisfrontier.ai

Reprints, enterprise distribution, excerpts, and licensing.

FRONTIER

www.thisisfrontier.ai

Research, Data Insights, News, and the live Frontier corpus.

All rights reserved. Contact Frontier for redistribution or reprint rights.

SEE FURTHER. MOVE SOONER.