ARMIS®
from ServiceNow

SOLUTION BRIEF

Defending at machine speed

The path to Shift Zero with Armis Centrix™

Executive summary

The software development landscape has fundamentally shifted. With the rapid adoption of agentic AI, engineers are writing and deploying code faster than ever before. But this blistering speed comes with a dark side: a mountain of AI-generated technical debt and a brand-new class of security risks. We are no longer just defending against human hackers operating on human schedules.

The reality of today's AI arms race includes advanced models deployed by threat actors, autonomous 24/7 bots capable of scanning, weaponizing, and exploiting vulnerabilities at superhuman speeds. While the security industry has historically focused on finding flaws and managing backlogs, that paradigm is now obsolete.

The core challenge for the modern enterprise is not merely the existence of vulnerabilities, but the ability to prioritize and remediate them in near real-time. Armis is rewriting the playbook, offering a platform where security is so deeply embedded into workflows that vulnerabilities are prevented before they ever manifest.

Discovery is not enough

Recent advancements in AI have revolutionized vulnerability discovery, turning it into a double-edged sword. Breakthroughs from leading AI organizations highlight both the defensive potential and the looming offensive threat:

> **Anthropic Mythos:** An advanced reasoning model capable of autonomously identifying vulnerabilities and engineering functional exploits. Under Project Glasswing, Anthropic used Mythos to discover multiple zero-days at an alarming speed, demonstrating that AI can bridge the gap from a raw finding to being exploit ready.

> **OpenAI Daybreak:** An AI native scanner combining the intelligence of models like GPT-5.5 with Codex. Daybreak excels at secure code review, threat modeling, and patch validation by deeply scanning software repositories.

While these tools are highly effective at finding risks, they inadvertently highlight the critical failure of the modern cybersecurity apparatus: finding risks is one thing; the real focus should be on fixing them.

The data deluge and alert fatigue

Models like Mythos represent the hunter-gatherer stage of security, diving deep into code and logic to uncover exhaustive lists of flaws. However, without context, these AI native scanners simply hand security teams a staggering list of potentially 10,000 broken items. For Vulnerability Management teams already drowning in backlogs, this capability adds to an insurmountable heap of "prioritized" fixes. A scanner running separately to discover risk is no longer fit for daily use.

Remediation and the ideal fix

Codex agents are adept at generating patches within repositories, but generating a generic patch raises critical questions: Does it actually fix the issue? Does it introduce a new flaw? Crucially, does it fit into the organization's unique architecture and best practices?

This is where Armis shines. Armis recognizes that even if an AI can generate a technically viable fix, a generic patch is insufficient. Armis delivers Knowledge-Driven Remediation to provide the "Ideal Fix".

Contextual AI reasoning: Rather than offering boilerplate code, the Armis platform ingests an organization's historical data, unique architectural patterns, frameworks, and coding standards.

Tailored solutions: The Armis-generated fix perfectly mirrors the way your company expects software to function, resolving issues exactly as a senior developer would.

Self-healing workflows: When a flaw is identified, Armis automates the entire remediation lifecycle. It auto-assigns the issue to the exact developer responsible for that specific line of code and automatically generates Pull Requests directly into GitHub or your preferred Git platform for easy review and auto-approval.

"The speed of the attacker is now the speed of light. The speed of the defender must now be the speed of context."

Welcome to Shift Zero

Rather than operating "right of boom", solutions like Armis Centrix™ for Early Warning already leverage AI-driven intelligence from dark web forums, dynamic deception technologies, and expert human intelligence to identify emerging threats before threat actors can weaponize them. By delivering highly contextualized, actionable alerts days, weeks, or even months before vulnerabilities hit public registries like the CISA KEV list, Armis gives security teams a critical head start. This vital time advantage allows organizations to patch systems, reconfigure networks, and harden their environments "left of boom," effectively preempting attacks and neutralizing risks before any operational damage or disruption can occur.

For years, the gold standard in AppSec was also to "shift left", moving security testing earlier into the development pipeline. But in the age of AI-driven agentic development, even shifting left is too slow. If you're waiting for code to be pushed to a repository to scan it, you've already lost the time advantage. Armis Centrix™ introduces Shift Zero: enforcement that happens before the code is even finished.

Enterprise Enforcement → Machine-Time Workflows → Knowledge-Driven Remediation

Enterprise enforcement and the Judge Model

With hundreds of developers leveraging AI coding agents, Armis acts as the ultimate enterprise enforcement layer. Armis automatically detects every AI agent installed and writing code across the environment. Every agent is paired with Armis's proprietary Judge Model. Because an AI model cannot reliably audit its own mistakes, the Armis Judge Model independently monitors its behavior, validating and countering the generated code in real-time.

Machine-time workflows

Armis applies pre-commit guardrails directly in the IDE. It scans code for vulnerabilities, supply chain risks, and poor architectural practices as the code is being generated, much before the code commit, build, and deployment phases. Problems are found, tailored fixes are applied within seconds, and new code comes out secure without adding to the backlog.

Knowledge-driven remediation

Apply the ideal fix within seconds through knowledge-driven remediation, significantly streamlining your security workflow. This approach drives AI coding agents to remediate vulnerabilities using your standards - your frameworks, your libraries, your approved patterns - instead of whatever a generic model would produce. Conformant fixes, at machine speed.

Why it matters now

A new generation of capable AI models can chain low-severity bugs into real exploits, autonomously. Critical and high vulnerabilities have always been triaged; lows and mediums rot in the backlog because no human can justify the time. That backlog is the soft underbelly, and it's about to be reachable at machine speed. The only viable defense closes the backlog at the same speed.

Generic AI agents can write fixes, but their patches don't match your standards - that's not remediation; it's a new review queue. Knowledge-Driven Remediation turns those agents into something you can actually ship: fixes that conform, traceable to a pattern you authored.

The code to device continuum

Tools like Daybreak operate strictly within the boundaries of software supply chains and code repositories. However, a vulnerability in code often manifests as an exploit on a physical device or a piece of critical infrastructure. A flaw discovered in a Linux kernel means very little until you know if that kernel is running on a non-critical guest Wi-Fi router or a life-support system in a surgical suite. Armis is the only solution that connects the dots from a vulnerability in a software repository directly to the physical IT, IoT, or OT asset running that code.

Furthermore, you cannot push an automated GitHub patch to a 15-year-old Siemens PLC on a manufacturing floor or a smart HVAC system. For the millions of unmanaged, physical, and legacy assets that AI scanners simply cannot remediate, Armis acts as the discovery engine for rogue repositories and shadow IT. We also prioritize reachability analysis, taking into account any compensating controls (like a web application firewall or segmented VLAN) that remove the immediate need for patching or mitigation without a single line of code being changed.

From backlog to breakthrough

The transition from merely identifying vulnerabilities to implementing actionable, knowledge-driven remediation represents a pivotal evolution in modern cybersecurity. The emergence of advanced AI systems, such as Anthropic's Mythos and OpenAI's Daybreak, serves as a stark reminder of the escalating risks associated with autonomous exploitation.

To fight this, Armis offers an AI-scale defense tailored for an AI-driven landscape. By integrating Shift Zero visibility, comprehensive business intelligence, and automated remediation, Armis actively clears security backlogs rather than simply managing them.

This decisive approach ensures that organizations can enable their teams to innovate rapidly and securely.