SACR

JULY 2026

THE CISO GUIDE TO ENDPOINT CONTROL AND PREVENTION (ECP): THE NEXT ARCHITECTURE FOR ENDPOINT AND AI SECURITY

FRANCIS ODUM
LAWRENCE PINGREE

cyberhaven

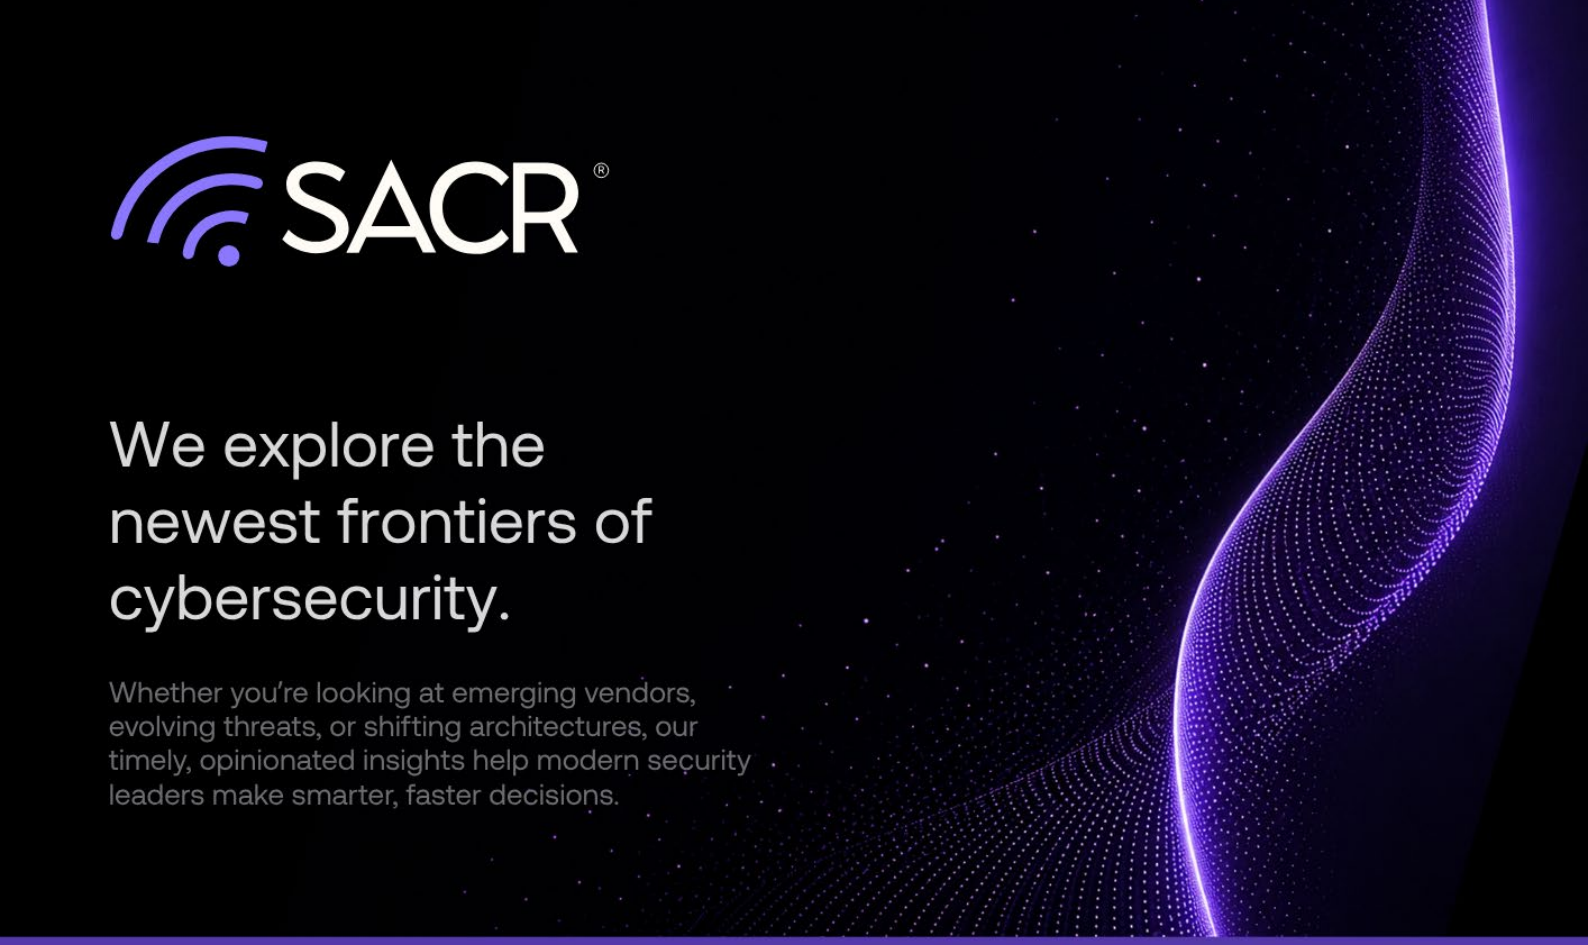

We explore the newest frontiers of cybersecurity.

Whether you're looking at emerging vendors, evolving threats, or shifting architectures, our timely, opinionated insights help modern security leaders make smarter, faster decisions.

About Us

Software Analyst Cyber Research (SACR) is a modern research and advisory firm built for today's cybersecurity leaders. We deliver in-depth, timely analysis across SOC operations, Identity, Network, Cloud, Application Security, Data, and AI Security; equipping CISOs, security teams, founders, investors, and practitioners with the insight they need to navigate high-stakes decisions.

With an engaged community of over **80,000 readers and followers**, SACR connects with a global network of cybersecurity decision-makers and innovators. Our access to leaders across categories and industries gives us a direct line to the conversations shaping the market. By pairing these insights with rigorous technical analysis and continuous market tracking, we produce research that is both data-driven and grounded in the realities of modern security operations.

Whether you're seeking clarity on emerging technologies, evaluating vendors, or tracking market shifts, SACR delivers trusted, independent research designed to help you see clearly and decide with confidence.

Acknowledgements

Practitioners and CISOs,

We're excited to share a new framework and system for securing agents across enterprises. This is one of our biggest reports published in recent months.

The Authors

Lawrence Pingree is the Head of Research at SACR, where he leads research on data protection, AI security, identity, cloud security, and agentic security models. With over 16 years of industry analyst experience, including more than a decade at Gartner, he has authored over 300 research notes and provides research driven insights to enterprise security leaders on emerging cyber threats, cloud security, AI security, endpoint defense (EDR), and SD-WAN.

Francis Odum is the Founder and CEO of Software Analyst Cyber Research (SACR), where he leads the firm's research and engagement with cybersecurity leaders. With extensive experience in cybersecurity research and market analysis, he built SACR into one of the largest independent cybersecurity research platforms, reaching over 120,000 security professionals worldwide and establishing the firm as a trusted advisor to CISOs and leading security vendors.

Table of Contents

Executive Summary

Endpoint security is entering its next architectural transition. For nearly fifteen years, Endpoint Detection and Response (EDR) has been the foundation of endpoint defence for enterprises. It was designed for an era where attackers executed malware, launched processes, and modified files that operating systems could observe. That assumption is rapidly breaking down. Companies like Crowdstrike EDR, SentinelOne EDR, Palo Alto Networks EDR earned that position by answering the question of its era: **what did this program do on the machine?** EDR watches the operating system. It logs when a process starts, and when a file is written, then matches those events to known attacker behavior. It worked because anything dangerous had to run as a visible program the operating system could see.

This assumption is now failing with AI usage in enterprises today. This is because the **way work** happens on the endpoint is changing. Three shifts separate the 2026 endpoint from the one most security programs were built for prior to 2023:

- **The rise of AI coding agents:** In 2023, AI in the developer's tools meant autocomplete that suggested code a human accepted line by line. In 2026, agents such as Cursor, Claude Code, and Copilot agent mode read the codebase, run shell commands, install packages, and complete tasks directly on the machine using the developer's own permissions.
- **Capabilities connect through a live protocol, not an installer.** The rise of Model Context Protocol (MCP), which did not exist before late 2024, lets agents plug into tools, data, and internal systems at runtime. Each connection is a new link in the supply chain that never appears in a traditional software inventory.
- **Everyone became a builder (not just engineers).** AI tools for employees, Low-code and no-code frameworks can now let someone in finance or operations assemble a working application, workflow, or agent in an afternoon, often with no ticket, no review, and no security checkpoint.

The pattern underneath all three is the same. The risky action is no longer a file the system can scan. It is an instruction written in plain language, carried out by software acting on a person's behalf, often across steps that each look harmless alone. The endpoint has stopped being a device the operating system governs and become the point where people, agents, applications, and data meet. Securing it means governing that interaction as it happens, not reviewing the process after it ends. These interactions often occur **above the operating system,** leaving conventional EDR with only a partial view of what is actually happening.

SACR believes this shift will create a new architectural category for the next decade: Endpoint Control and Prevention (ECP).

ECP covers a new vector that EDR's today miss. They focus primarily on detecting malicious processes after execution; ECP expands endpoint security into a runtime control plane capable of governing AI agents, browser activity, application workflows, identity context, and sensitive data before attacks complete.

The market landscape ecosystem is emerging. We have seen an explosion in companies on the market moving to solve the problem. Specialized vendors are expanding endpoint security across software posture, application enforcement, AI runtime visibility, behavioural analysis, and data-centric controls. Over time, these capabilities are likely to converge into broader endpoint platforms.

- **EDR to ECP players**: CrowdStrike, Palo Alto Networks, SentinelOne have moved to cover the new agentic market complementing their EDR solutions.
- **Emerging ECP native key players** include: Neo Security, Bay Security, Bloom Security, Glow Security, Pluto Security, Origin HQ, Ent, Bold Security, NeutralTrust, and Cyberhaven, which define individual zones through

purpose-built architecture. The vendors that already span several zones point to where the category is heading, toward one converged control plane rather than five separate tools. These are leading the charge to solve this market issue. We have partnered closely with the key players to perform our analysis, which is attached in the report.

On the left is the EDR market as it stands today, competing on process and file telemetry that has largely commoditized. On the right, that same market fans out into the five zones where endpoint defence is now expanding, from software posture and supply chain governance at one end to data-centric enforcement at the other. The connective idea is that EDR becomes ECP: an extension of the endpoint, not a replacement of it.

The Report Summary: This report makes five observations.

1. **The endpoint is becoming the primary enforcement point for AI.** Security controls must move closer to where prompts, tool calls, browser actions, and agent decisions actually occur. In 2026, more of this activity is happening at the endpoint.
2. **EDR is approaching an architectural ceiling.** Process and file telemetry remain essential but no longer provide sufficient visibility into AI-driven workflows, browser sessions, SaaS applications, and non-binary software.
3. **Context becomes the new detection engine.** Modern security requires understanding who initiated an action, what prompted it, which tools were used, and whether the resulting behavior aligns with legitimate intent.
4. **Prevention increasingly replaces response.** Rather than terminating processes after compromise, next-generation platforms intervene earlier using software governance, workflow controls, behavioral guidance, and inline enforcement.

Endpoint Control and Prevention should not be viewed as the replacement for EDR. Instead, it represents the next stage of endpoint security, extending protection from the operating system into the interactions between users, AI agents, applications, identities, and data that increasingly define enterprise risk.

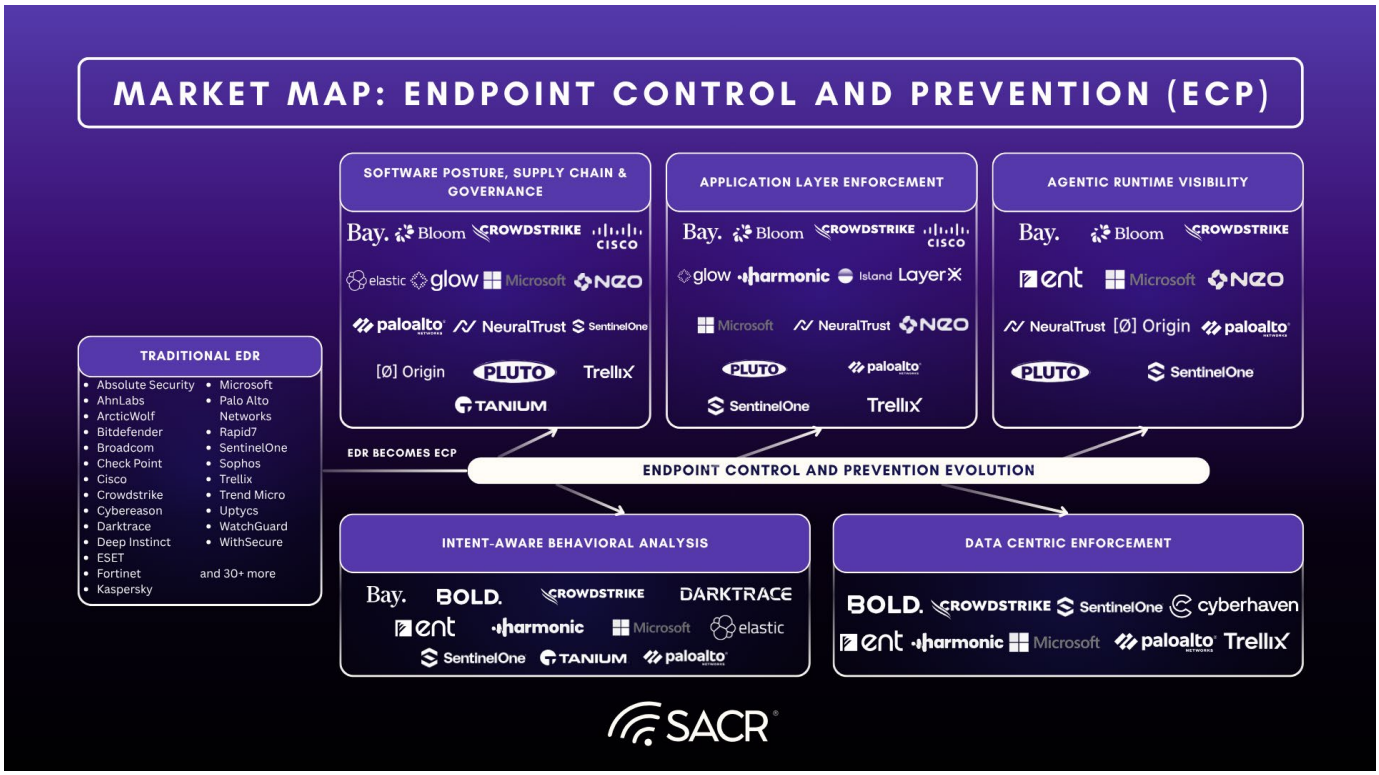

Endpoint Control and Prevention (ECP) Thesis

Our thesis is that the core failure in EDR is structural. Defenders lack actionable context within the execution workflow. Traditional endpoint tools secured basic machine hygiene and website access, but the next generation must secure the active interaction flow between users, autonomous agents, and SaaS platforms. Endpoint defense must transform from a passive detection sensor into a strategic control plane for the expanded use cases. EPP and EDR protected the machine and EDR protected the processes and some scripts; today, the defense perimeter must expand. The next-generation Endpoint Control and Prevention (ECP) endpoint protects the interaction of the critical flow between humans, AI systems, services, and data.

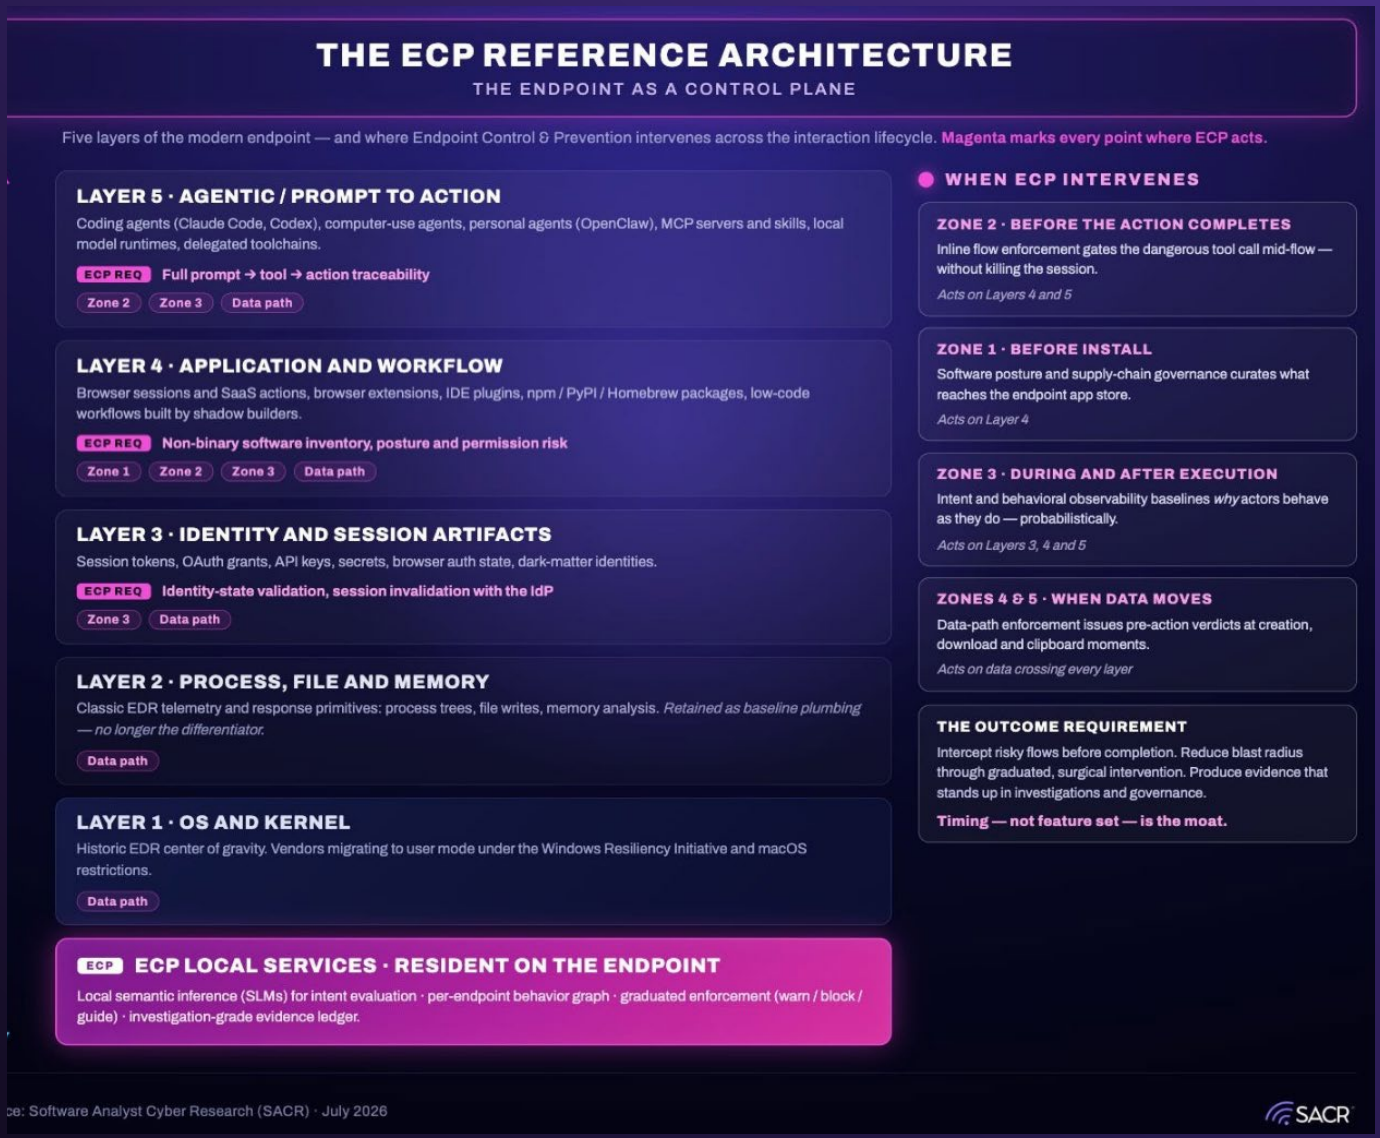

Market Definition: Endpoint Control and Prevention (ECP)

Endpoint Control and Prevention (ECP) is a pivot in endpoint security that reframes endpoint defense as a runtime visibility, operational workflow, and user contextual gap problem. ECP expands endpoint visibility beyond operating systems, files, and processes into developer environments, SaaS applications, AI tools, and agent frameworks. It secures these environments by monitoring the interactions and data flows between users, AI agents, applications, and command-line activities. Next-generation ECP focuses on prevention, enhancing context sharing and user-driven responses across emerging AI, SaaS, and data environments. It focuses on the governance of non-binary software, deeper application-layer execution context (especially for AI and data movement) for better user and agent behavioral characterization or judgement, and machine-speed intent recognition from AI to ameliorate precision action. All of which offers endpoints stronger attribution capabilities, allowing endpoints to transform from more passive detection and response-oriented sensors into more active control planes for the user, AI or agent for prevention, investigation or response actions.

Strategic Imperative: Transforming Telemetry into Decision Support

1. Application-layer discovery and posture control for non-binary software (extensions, plugins, models, MCP tooling, configuration risk)
2. Greater application user behavior visibility depth, governance and control by extending policies to the browser and its contextual telemetry.
3. AI Prompt-to-action attribution for agentic systems (what triggered an action, what tools were used, and what happened next)
4. Local semantic inference (often via small local language models SLMs and graph storage) to interpret intent and drive proportional controls and responses autonomously from back-end operations.

The outcome requirement is simple: Intercept risky flows and sessions before completion, reduce blast radius, and produce evidence that stands up in investigations and governance.

In scope

- Threat detection evolution for AI and local agentics
- Surgical intervention for lineage, evidence, and governance of users, data, and AI and agents
- Local application context and extension to non-binary configurations and unmanaged dark matter identities (secrets, certificates and other artifacts)
- User, Identity and agent telemetry, traceability and data lineage
- Trusted, zero trust access and software postures above the OS (e.g. visibility and control over local DevOps integrated development environments, IDE plugins, AI agent harnesses/frameworks)
- Identity and SaaS-mediated execution context (e.g endpoint and SaaS shared telemetry to shift more policy and threat detection from backend to frontend)

Out of scope

- Standard EDR feature updates focused on legacy malware, traditional endpoint focused indicators of attack (IOA) or indicators of compromise (IOC) indicators
- Standalone SIEM or SOAR platforms
- Network-layer-only controls
- General XDR bundles that primarily aggregate endpoint detection + response

Why is Endpoint Security being Re-Written?

Endpoint defense is being rewritten because decision points and context that can be derived and used are moving upwards in the stack. Context has taken precedence in a world of dispersed applications and interactions. In the near future, federated context from MCP will contribute significantly to both backend analytics and endpoint security as the threat landscape expands. Classic malware remains a threat, but many incidents now begin with a compromised session, a risky SaaS authentication (OAuth) grant, a compromised API secret, credentials, or a workflow executed legitimately with a stolen identity.

As adversaries bypass technical controls through credential attacks, session hijacking, and autonomous agent manipulation, the traditional telemetry stack encounters three terminal blind spots:

1. **Invisibility of non-binary software:**

   Browser extensions, plugins, and local models execute without traditional installations, leaving file-centric EDR blindsided.

2. **Deficit in intent and attribution**

   EDR tools fail to distinguish between malicious injections and legitimate commands, lacking autonomous oversight and intent assessment.

3. **Contextual blindness in execution paths**

   Encrypted connections, AI apps, and SaaS environments obscure user and application context from traditional tools.

Nefarious actors can now progress via autonomous AI agents and their delegated toolchains, where the file system becomes more secondary (or irrelevant) to the execution context or location of execution. This means shared context must come from new sources to make better, more localized decisions with enhanced visibility and context on the endpoint, giving security responders faster, real-time prevention, depth for incident response efficiency and enhanced precision. We depict the new evolution in the timeline graphic below.

Era Emergence: From EDR to Endpoint Control and Prevention (ECP)

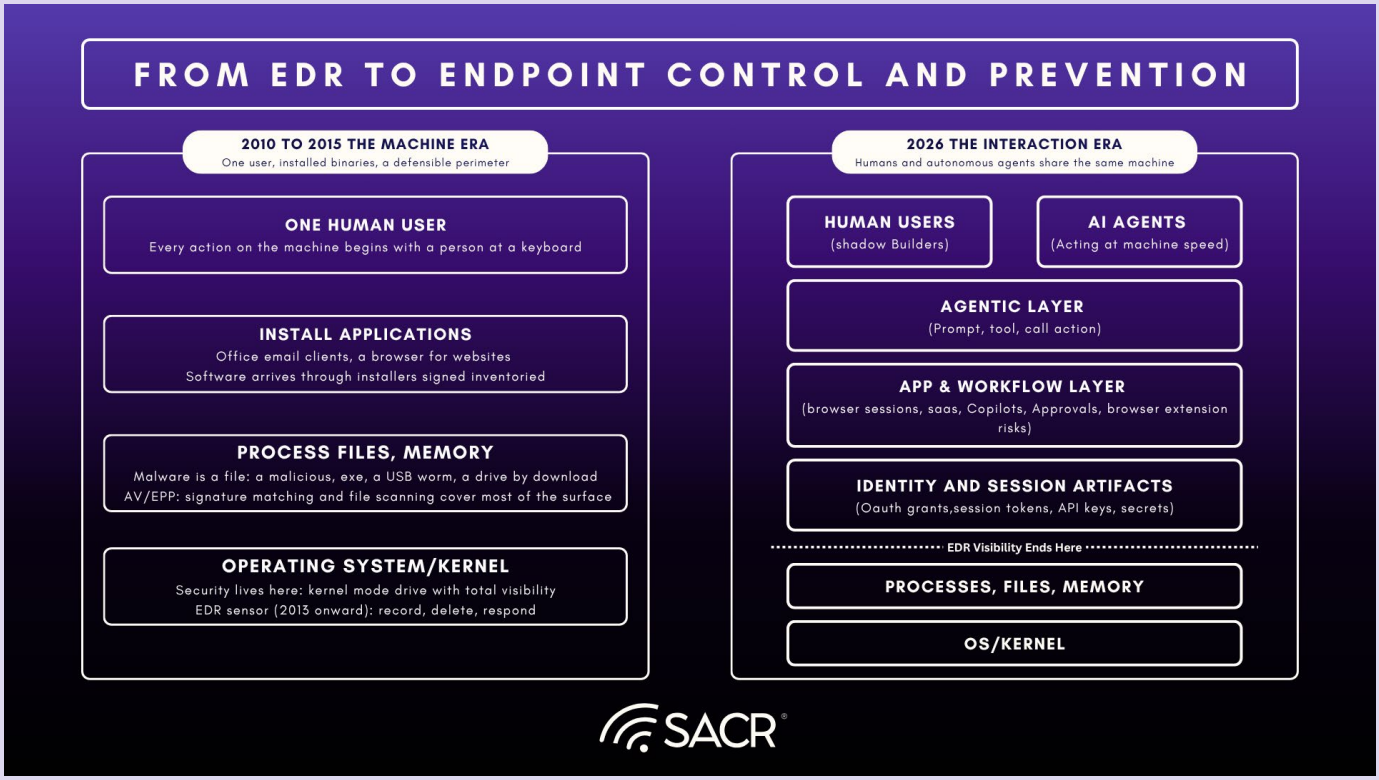

Legacy EDR is failing because it targets symptoms and conditions rather than the underlying structural issues or overall execution contexts. By focusing on file system telemetry and process trees, traditional security boundaries miss the strategic shift up-stack to non-binary software, sessions, autonomous workflows and SaaS. ECP resolves this structural blindness by transforming the endpoint into an active control plane that monitors, participates and governs the interaction layer between humans, agents, and data.

Conceptualizing the evolution of endpoint defense transformation requires transitioning to a framework focused on establishing new architectural layers or zones, leading directly to improvement of endpoint defensive and prevention posture. By integrating advanced visibility and control across these areas, organizations can re-align with contemporary threat landscapes and optimize the use of multi-model context, behavioral intent, and diverse data interactions to drive proactive defensive and prevention-focused decisions.

The Eras of Endpoint Security

Endpoint security has evolved through three distinct eras, each triggered by the structural failure of the previous architecture.

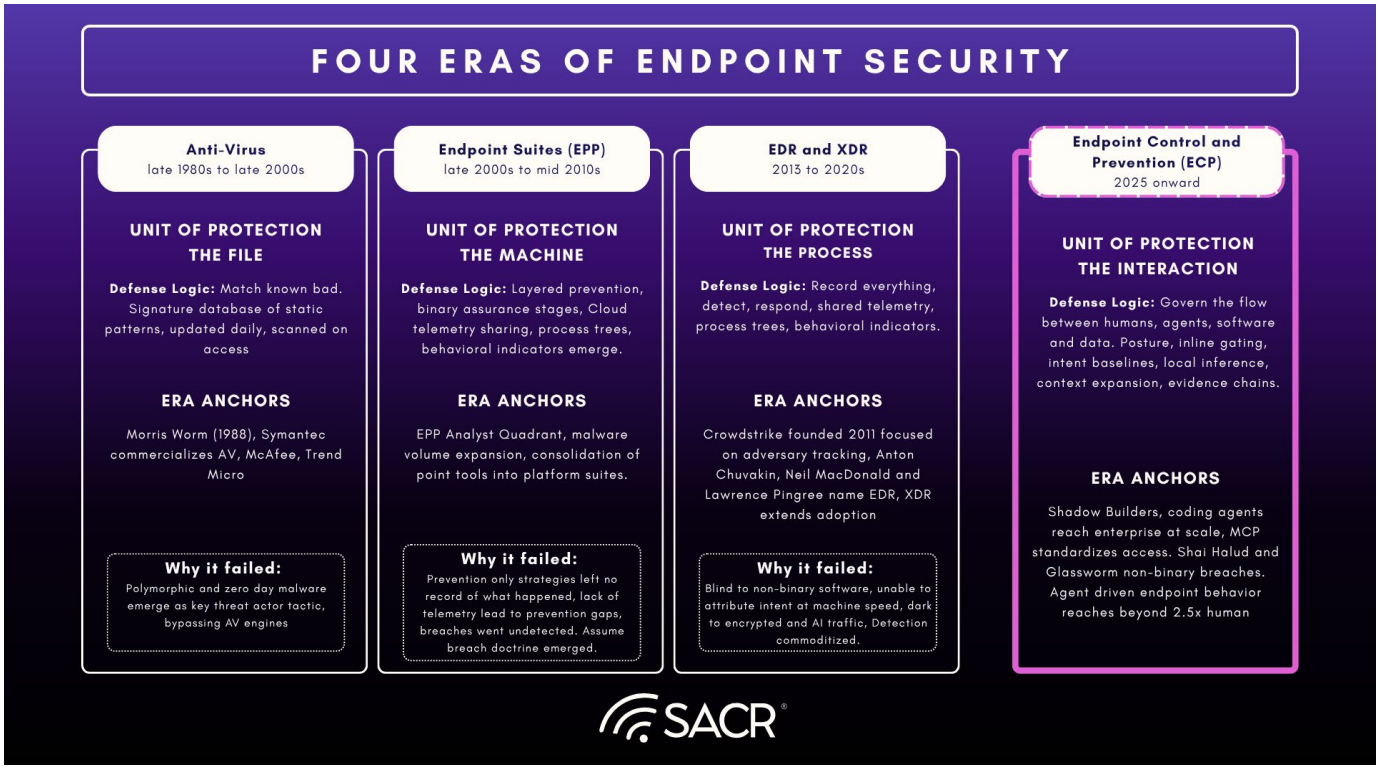

The Antivirus Era (late 1980s to late 2000s): Protecting the File

The first commercial endpoint security products industrialized the idea of maintaining a database of signatures for known malicious files. This worked while malware was a static file, the industry was in a race between virus authors and signature distribution. The model collapsed under the pressure of polymorphic malware and targeted attacks like Operation Aurora (2010), which demonstrated that adversaries could bypass signatures entirely by using bespoke, never-before-seen implants.

The EPP Era (late 2000s to mid 2010s): Protecting the Machine

The industry's first response was to bundle security tools, antivirus, firewalls, and device control into Endpoint Protection Platforms (EPP). While EPP improved prevention, it lacked visibility into what occurred when prevention failed. Breach investigations of the era consistently revealed that intruders remained resident for months, leading to an assume-breach doctrine and the realization that the decisive capability of the EPP era was not blocking malicious files, but rather observing endpoint activity.

The EDR Revolution (2011 to the mid 2020s): Protecting the Process

Endpoint Detection and Response (EDR) succeeded by recording everything, from process creation, file writes, and registry changes, into a cloud-hosted graph to identify adversary behavior. EDR's dominance rested on one foundational architectural assumption: everything that matters on an endpoint executes as a process the kernel can see. As threats migrate up-stack to non-binary software, sessions, and AI agents, this reliance on kernel-level process visibility is reaching its limit, necessitating the shift to Endpoint Control and Prevention (ECP).

The ECP Redefinition (2026 through 2030): All we need is Context and Intent

The ECP redefinition phase marks a strategic pivot away from reactive process monitoring toward holistic, interaction-centric, SaaS and context-aware governance with user and agent intent awareness. This timeline emphasizes the necessity of newly emerging agentic workloads, managing high-speed autonomous execution, where the endpoint must evolve into an active control plane capable of verifying intent and securing non-binary and user execution contexts in real time.

Focused on Emergent areas such as:

1. **The Human User (Shadow Builders)**
2. **AI Agents**
3. **Agentic Layer**
4. **Application and Workflow Layer**
5. **Identity, Session and Artifacts**

The rapid proliferation of enterprise AI agents and automated supply chain vulnerabilities has intensified this crisis. Modern adversaries bypass OS-level signatures entirely by injecting malicious tools directly into shadow builder or developer ecosystems and agentic frameworks. These autonomous threats often operate with high-level permissions, independently navigating networks and executing custom, on-the-fly commands or building and executing their own code at machine speed. Recently, attackers have adapted part of their game; since so much net-new AI is being adopted on endpoints, threat actors are now targeting the new AI and agentic software supply chains to execute their nefarious acts.

Recent supply chain breaches highlight the urgency of securing this non-binary execution context:

- **PyTorch Lightning PyPI Compromise (April 2026):** Hijacked credentials allowed malicious library versions to harvest and exfiltrate environment secrets via automated CI/CD pipelines.
- **Mercor-LiteLLM Upstream Breach (April 2026):** A compromised routing dependency propagated downstream, exposing sensitive training pipelines and forcing immediate contractor suspensions.
- **NX NPM Breach & Agent Hijacking (August 2025):** Poisoned packages executed locally to hijack active AI tools, systematically hunting and exfiltrating SSH keys and API tokens.
- **Hugging Face Weaponized Models (Ongoing):** Malicious models embed reverse shells within unsafe serialization formats, instantly compromising the host endpoint upon model loading.

Assumption: AI agents and their arrival with Endpoint harnesses creates a cascade of new demands such as monitoring, assessment of risk, management and control, driving net-new technology adoption in endpoint security. These and the emergence of custom software generated on-demand by AI agents (even apps, scripts and services deployed locally on endpoints) become a new IT standard; legacy endpoint security metrics and even DevOps or DevSecOps metrics will also become obsolete on a standalone basis. One way to think of this is in the Builder Era (where all users become agentic shadow builders), federated development becomes the norm. This strongly shifts the value proposition of endpoint defense decisively from post-incident response to real-time execution control and proactive context governance.

To defend the modern enterprise, endpoint security must pivot from general user, software, and binary-driven monitoring and acquire higher-order contexts in the stack, helping unify disparate activities on the endpoint for enforcement pre prevention consideration. We need specialized oversight for autonomous agentic systems, integrated SaaS and the context created by users and agents, and visibility and control over the workflows connecting them. In older environments, investigators struggled to tie a user's actions to automated agents they triggered. Today, as users increasingly become shadow builders spawning agents, writing and executing code via local chat interfaces and agentic harnesses, these new agents effectively become net-new bots on the internet. This is mandating evolution in adjacent areas (outside endpoint), for example, how we identify bots vs agents online. AI Agents often use similar tools to bot scraping

services and other engagement tools that make them overlap with traditional malicious bots, representing the same dynamic and challenge to overcome on the endpoint. Even when operating under a user's identity, actions taken by agents require a new forensic perspective and better context for decisions.

Effective defense now relies on our ability to provide user and agent context and enable rapid context sharing and enforcement mechanisms at machine speed. By adopting this approach, acceleration enables near-instantaneous security decisions, empowering both users and automated monitoring agents and other local sensory software to intervene and halt malicious behaviors the moment they arise. This shift opposes the traditional, slower-moving world of traditional breach-oriented detection and response focused on users or simple administrative functions or scripts (already representing a detection challenge). As enterprises weigh the benefits of control against the need for agility, the industry is moving decisively away from manual human-in-the-loop (HITL) processes and toward automated, proactive control and prevention.

In today's environments, high-consequence actions rarely occur as simple file executions. They happen inside:

- **Browser Sessions:** Where copilots handle approvals and data movement.
- **The Identity Plane:** Where attackers leverage tokens, OAuth grants, and API keys.
- **Delegated Toolchains and Self-Developed Agent software:** Where agents run scripts, build their own tools on the fly and move data directly.

Historically, Endpoint Detection and Response (EDR) has failed to extend into these specific areas on the device. However, technical shifts, such as machine-speed risks and AI-powered attacks using the advanced capabilities of Mythos, are forcing a faster transition to higher speed, greater context, and prevention prioritized over the EDR detection and response heritage. Organizations that continue to rely solely on traditional detection and response strategies will fail, the landscape is being redefined into one of real-time, active control and prevention-focused defense.

Emerging ECP Value Proposition

The Endpoint Control and Prevention (ECP) value proposition centers on transforming endpoint defense from a passive, binary-centric sensor into an active control plane capable of governing the modern interaction lifecycle, whether user or agent. By shifting visibility and enforcement up-stack, beyond kernel-level activity and file-based telemetry, ECP restores defender advantage in an era of AI-driven, machine-speed attacks. It bridges the gap between legacy detection and the reality of autonomous workflows, providing the contextual evidence and granular, real-time intervention primitives needed to secure the critical flow between humans, identities, agents, and SaaS applications.

Critical Endpoint Control and Prevention Layers

- **OS and Kernel (classic EDR center of gravity):** This layer represents the traditional foundation of endpoint security, focusing on low-level system activity and kernel-mode monitoring to identify malicious behavior at the hardware-software interface.
- **Process, file and memory (EDR telemetry & response primitives):** These core primitives enable the detection of legacy file-based threats and post-incident reconstruction by analyzing how processes interact with the file system and system memory.
- **Identity and session artifacts (tokens, OAuth, browser auth, session abuse):** In ECP, visibility and context extend to the identity plane to protect against session hijacking and the theft of cryptographic artifacts in real time, at runtime: for example, tokens and API keys used in SaaS-mediated workflows.

- **App & workflow layer (browser actions, SaaS actions, agent tool calls):** Endpoint defense now necessitates runtime observability at the application layer, governing non-binary software such as browser extensions and IDE plugins where critical modern execution context resides, but is often not shared.
- **Agentic prompt-to-action layer (prompt, tool, action traceability):** This critical new layer provides full attribution for autonomous systems, linking natural language prompts to specific tool invocations and downstream actions to ensure governability at machine speed.

Endpoint and End-User Organization Trends Enabling the Shift

Modern end-user organizations are driving the shift toward Endpoint Control and Prevention (ECP) as traditional security perimeters dissolve. With the rise of hybrid work and the proliferation of "shadow builders", which are often non-technical staff deploying custom AI-driven application development, workflows and SaaS applications. The endpoint attack surface has expanded beyond legacy and more stable binary and process-centric threats. To address these complex risks, organizations are prioritizing comprehensive visibility into application-layer interactions and autonomous agentic behaviors, necessitating a move toward granular, real-time control.

- **Broad adoption of hybrid working and the rise of Shadow AI and Shadow Builders:** The emergence of shadow builders and Shadow AI both represent a significant market driver for ECP, as non-technical staff increasingly use low-code and no-code AI agent frameworks tools and AI plugins to assemble automated tools and workflows. This decentralized development creates new third-party supply chain models directly on the endpoints or inside designated harness infrastructure for agents, necessitating visibility into non-binary software linked to browser sessions, extensions, and IDE plugins.
- **Instrumentation is shifting down-stack while risk shifts up-stack.** Kernel and eBPF-era telemetry improves what happens, but the differentiator is increasingly interpreting and governing what runs above the OS (extensions, plugins, packages, models, MCP tools and shadow supply chains).
- **Shadow builder risks are now a first-class endpoint problem.** Since non-developer users and teams can now assemble powerful workflows, build net new applications and deploy plugins and AI tools, security needs visibility into all of this new installable software, its permissions, composition risk, configuration and communications, which are no longer limited to just malware-style OS level indicators.
- **Deep Posture Assessment and Identity Awareness are emerging.** As per-endpoint detection commoditizes, budget and evaluation energy move to application posture, attribution depth, and real-time intervention primitives (warn/block/guide) that reduce blast radius.

Operational Reality: Aligning Defense with Machine Speed

We have entered the age of Mythos, which is an era defined by the industrialization of vulnerability exploitation at machine speeds. Automated models are compressing the time required to weaponize software flaws from weeks to ~seconds. With the emergence of models like Mythos and open-source frameworks like GLM 5.2, the cybersecurity landscape has transitioned into a high-speed race to automate both AI vulnerability discovery and remediation via threat management processes. This evolution renders traditional, manual patch cycles obsolete, forcing a strategic pivot toward continuous, autonomous prevention engineering and retooling of the flows of data and context across security systems more generally.

The failure of legacy defense stems from a structural misalignment with how work actually occurs:

- **Execution has shifted from local binaries to SaaS-mediated workflows:** Security telemetry cannot stop at the file system or process tree when the action and critical context occur in browser sessions, SaaS and API calls.
- **Identity is still a primary chokepoint:** Tokens, OAuth grants, and session artifacts like cookies or authentication bits are the new crown jewels, and their abuse is a primary vector for modern compromise.
- **Non-binary attack surfaces are expanding:** The real risk resides in extensions, IDE plugins, and agentic configurations that mutate daily; a compositional, use- driven attack surface can bypass legacy signature-based scanning.
- **Autonomy has changed the unit of risk:** We have moved from a user running a process to an agentic toolchain executing a workflow. These actions occur at machine speed, rendering human-in-the-loop triage a strategic bottleneck.
- **MTTR requirements have collapsed:** Cloud-backend analytics loops cannot keep pace with agentic misuse. Defensive logic must shift right to the endpoint runtime.

The Strategic Imperative: Securing the AI and Agentic Layer

Modern ECP architectures require deep visibility into the non-deterministic nature of local agentic infrastructures, including IDE copilots, local AI runtimes, and MCP-integrated toolchains. Governance must extend to the prompt-to-action trace to ensure full context, lineage, and verifiability for autonomous systems. Defenders must ensure that actions taken by local agents against remote APIs or SaaS platforms are properly contextualized, mediated, attributed, as well as mapping out what triggered an action, what tools were used, and what happened next. These are very important aspects of monitoring context to drive enhanced endpoint prevention and behavioral response decisions.

Small Language Models (SLMs) Arrive on the Endpoint

The transition toward AI PCs and Minis, along with an increase in shadow builders (unauthorized) or authorized users setting up local agent harnesses, is accelerating the need for deployment of local models, bringing agentic benefits to endpoint security. This development empowers local AI models and agents specialized in Endpoint Security to manage scaling agentic telemetry and localized or remote context, circumventing the cloud-pipeline expenses, latency, and contextual issues inherent in current setups.

Traditional, EDR-centric protection models are proving inadequate for modern defenders. In this report, SACR highlights a necessary paradigm shift: migrating full-stack context and security logic directly to the device. By leveraging local agents (Endpoint Defense Agents), Small Language Models (SLMs), and local graph databases, systems can achieve semantic and behavioral insight to improve threat prevention. This architecture facilitates instantaneous enforcement, circumventing the delayed, cloud-reliant detection pipelines characteristic of legacy EDR/XDR vendors that depend on central storage backends.

Endpoint Defense Agents Rise

Moving agentic inference onto the Endpoint hardware redefines the threat landscape, where local inference risks, data lineage and attribution emerge as the primary detection hurdles. Rather than simply evaluating if a process is malicious, defenders must identify the specific user, agent, workflow or tool and identify errant or malicious prompts that initiated an activity, confirming if it aligned with user intent. This establishes prompt, user and agent action lineage as the modern extension of threat detection from the traditional process tree.

Adaptive and Predictive Prevention Becomes the New Goal

Cloud-backend analytics and context enrichment loops can no longer keep pace with AI and agentic misuse, causing traditional MTTR requirements to collapse and forcing defensive logic to shift right to the endpoint runtime. Moving endpoint defense beyond traditional process and file centricity, this approach rejects standard signature or rules-based logic in favor of probabilistic, anomaly-based detection driven by local semantic inference. Endpoint security systems must evolve into proactive, predictive prevention engines and contextual sharing architectures. By leveraging expansions and preemptive modeling of advanced context, behavioral modeling and intent awareness of both user and agent, endpoints can enhance and anticipate agent-driven threats before they fully manifest, achieving millisecond-level mean time to response (MTTR). This proactive stance supports a self-healing zero-trust model, where the endpoint facilitates faithful action execution while continuously reducing the attack surface. By executing defensive logic directly at runtime via local SLMs, the endpoint can interpret user and agent intent to intercept risky or non-deterministic workflows at machine speed before any malicious intent can fully execute.

7 Critical Endpoint Control Moment Pillars Important for ECP

To effectively govern modern autonomous workflows and minimize the blast radius of agentic toolchains, security teams must embed defensive logic across the entire interaction lifecycle. Shifting to Endpoint Control and Prevention (ECP) requires moving beyond reactive, post-execution detection to establish high-fidelity visibility at each stage of the workflow.

The following pillars identify the critical control moments where ECP architectures must intervene to ensure policy adherence, verify intent, and mitigate risk before an action completes.

Endpoint Security Enforcement Control Moments

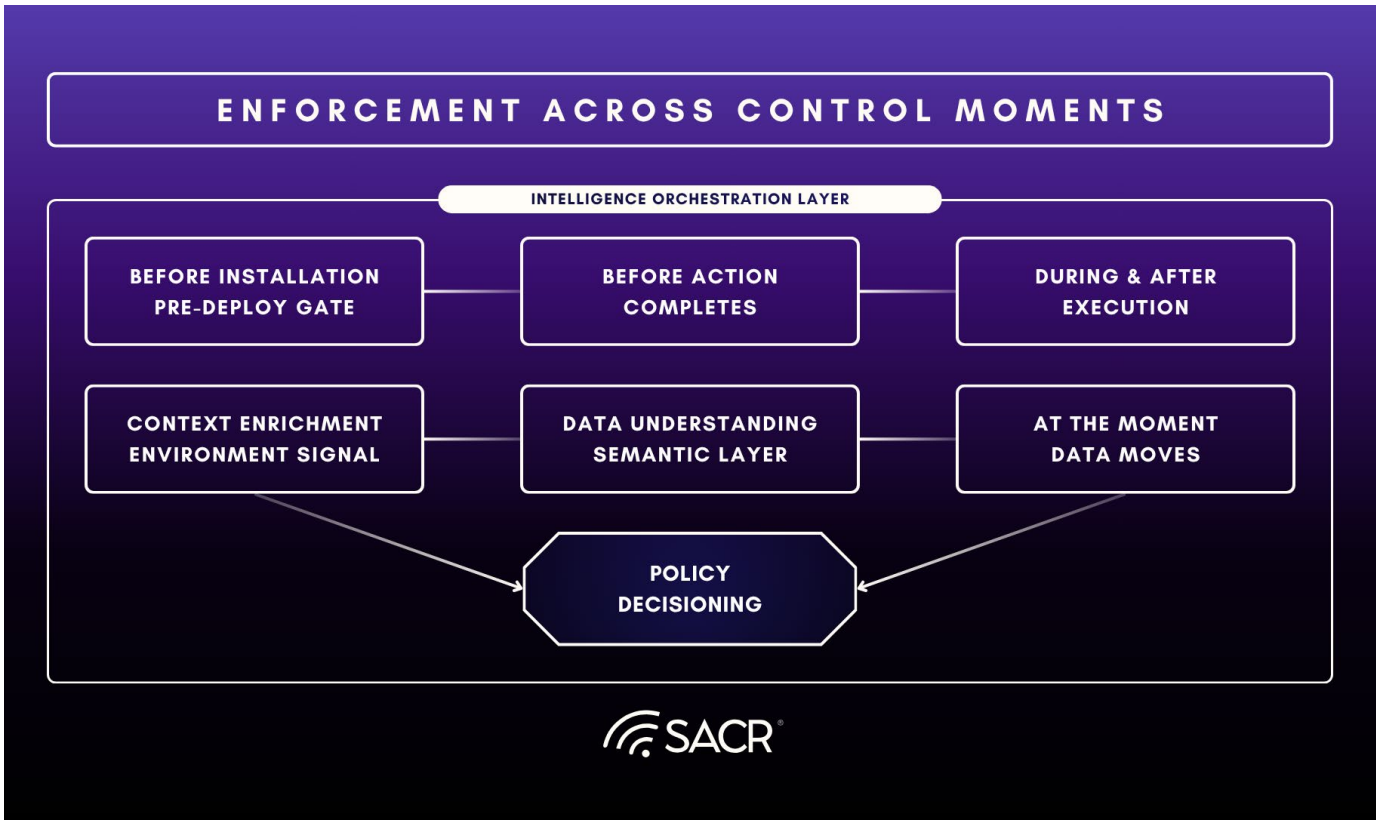

1

Before Installation (Pre-deploy gate): Focuses on proactive curation of the software supply chain to neutralize threats before they manifest on the device. By treating the endpoint similarly to an app store, this phase discovers and risk-ranks non-binary software components, such as browser extensions, IDE plugins, MCP servers, packages, and local models, ensuring configuration hygiene and preventing unauthorized dependencies from landing on the machine.

2

Before Action Completes: Involves gating risky behaviors, dangerous tool calls, or unauthorized data sharing mid-flow. Instead of relying on disruptive post-execution process termination, this control point uses inline, flow-level session evaluation to intercept actions at machine speed before they can fully complete.

3

During & After Execution: Builds runtime observability and behavioral baselines above the operating system layer. It continuously monitors active interactions, including natural language prompts, shell commands, tool execution, and network activity, using local semantic inference to detect intent-level anomalies and maintain complete forensic traceability.

4

Context Enrichment (Environmental Signal): Integrates deep identity plane and environmental telemetry, such as active session states, cryptographic tokens, OAuth grants, and API keys. This shared context connects endpoint activity directly to SaaS-mediated execution paths, providing the multi-layer visibility needed for continuous trust validation.

5

Data Understanding (Semantic Layer): Leverages on-device AI models and local small language models (SLMs) to interpret the natural language context of prompts and workflows. Shifting from rigid pattern-matching rules to probabilistic inference, this layer evaluates user and agent motivations to recognize risky prompt injection or unauthorized orchestration before execution.

6

At the Moment Data Moves: Monitors and intercepts sensitive data flows at critical interaction points, such as clipboard paste, file downloads, or model prompt submissions. By performing local AI classification directly at the endpoint edge, it successfully inspects cert-pinned or SASE-bypassed traffic that network proxies miss.

7

Final Policy Decisioning: Evaluates the combined signals of intent, posture, identity, and data behavior against organizational security rules to drive real-time, graduated interventions. Rather than defaulting to binary allow-or-block decisions, it enables proportional enforcement primitives such as surgical inline redaction, warnings, or context-aware user coaching.

SACR Security Zones for Endpoint Control and Prevention (ECP)

To evolve new use cases across the control moments pillars, we see emerging technology concepts and vendors evolving endpoint defense in several key emerging zones of technology and use case expansion, which requires transitioning from legacy, process-centric models to a robust architectural framework of ECP security zones. These zones serve as the blueprint for securing the modern, autonomous workflow by establishing high-fidelity visibility and control at the critical layers where risk now resides. By mapping defenses to these five distinct evolutionary focus areas, security teams can move beyond reactive detection, enabling proactive governance of user, agent, and data interactions at machine speed.

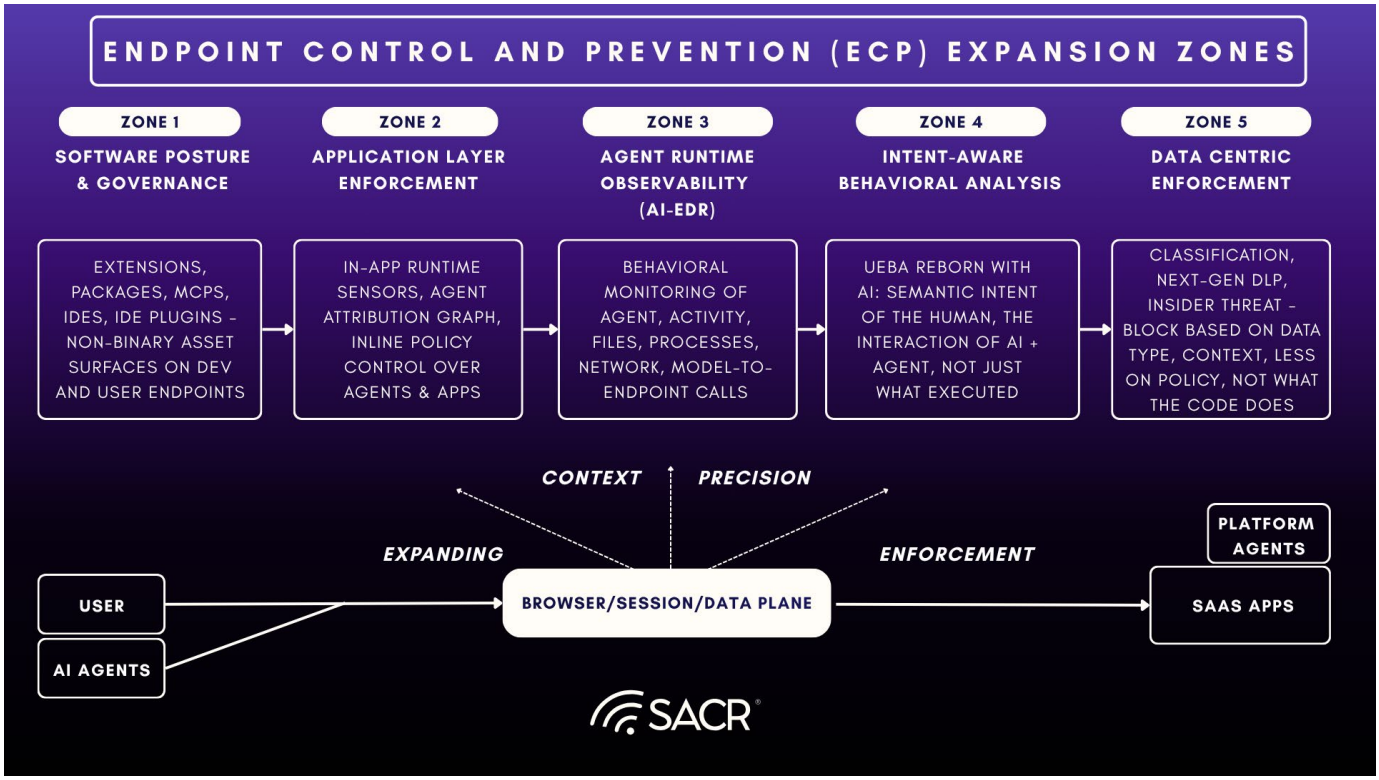

Zone 1: Software Posture & Governance

Concept: By treating the endpoint similarly to an app store, this layer controls the ingestion of various software components such as extensions, MCP servers, packages, models, and containers. Proactive curation of the software supply chain neutralizes incidents before they manifest by maintaining an active inventory and comprehensive posture assessment.

Operational Pillar: Shifts endpoint defense from reactive binary signature matching to proactive continuous governance of the non-binary software composition layer. This involves discovery and continuous risk-ranking of browser extensions, IDE plugins, and local model frameworks to establish an active inventory and continuous posture assessment above the operating system layer.

By treating the endpoint similarly to an app store, this layer controls the ingestion of various software components such as extensions, MCP servers, packages, models, and containers. The foundational premise is that because the vast majority of AI risk is acquired, proactive curation of the software supply chain can neutralize incidents before they ever manifest. Referenced as Endpoint Application Posture Management (EAPM) within SACR advisory frameworks and as a critical feature of ECP

expansion, this paradigm and concept marks a strategic resurgence of classic OS oriented posture and control application control and configuration assessment with renewed focus on endpoint and agentic applications and supply chains.

While traditional allowlisting was historically constrained by operational overhead, intensive manual fine-tuning, and rapid policy challenges, AI completely redefines this resource-heavy dynamic. By deploying automated fleets of specialized agents, organizations can continuously analyze software behaviors and dynamically maintain allowlists at machine speed. Enforcement within this zone shifts the primary security discipline away from reactive detection engineering toward proactive policy engineering. It treats platforms as first-class citizens where configuration liabilities, such as a Claude instance, for example, with dangerous permissions enabled, are surfaced as high-severity posture anomalies rather than traditional malware events. The ultimate goal is to sustain an active inventory and comprehensive posture assessment of all non-binary dependencies functioning above the operating system layer, verifying configuration hygiene, access permissions, and the intricate workflows generated by shadow builders using AI agents before any malicious intent can execute.

Zone 2: Application Layer Enforcement

Concept: Sits at the interface between native apps and AI agents, governing layers like browser extensions, developer packages, MCP servers, and AI agent toolchains. It blocks dangerous tool-calls mid-flow or potentially risky data sharing without killing the session.

Operational Pillar: Enforce inline gating and session evaluation mid-flow to block dangerous tool calls or data sharing before completion, replacing disruptive process termination with surgical intervention embedded directly within the active workflow layer.

Sits at the interface between native apps and AI agents, governing layers like browser extensions, developer packages, MCP servers, and AI agent toolchains. It blocks the dangerous tool-call mid-flow or potentially risky data sharing without killing the session. The core bet is that at machine speed, where autonomous agents read natural language instructions, execute shell commands, edit files, and call APIs at a speed no human could perform or anticipate, traditional detection-and-response is operationally and structurally too late, and therefore must evolve to keep pace with emerging risks.

The rising prevention-first approach mandates that actions be gated using inline, flow-level, session evaluation, or data-layer enforcement before tool calls can fully complete. This strategy successfully replaces disruptive process termination with a more precise, surgical intervention embedded directly within the active workflow layer. Vendors in AI are hard at work continuously modeling both AI interactions to preemptively identify attacks while also expanding their ability to monitor deep workflows, and as DNS for Agents emerges in DNS land, most organizations still lack methods for executing more complex multi-layer agentic workflow recording and intercede. Traditional tools cannot block potentially harmful actions on both Application AI agents and their workflows together as a continuum, where the evaluation of intent and outcome behavioral analysis focuses on precision prevention.

Zone 3: Agent Runtime Visibility (AI-EDR)

Concept: This architectural layer builds its own above-OS telemetry and baselines of why actors behave as they do, rejecting standard signature or rules-based EDR logic in favor of probabilistic, anomaly-based detection driven by local semantic inference.

Operational Pillar: Provides full context and verifiability for autonomous systems by establishing absolute traceability from natural language prompts to downstream execution steps, tool calls, and system responses. It reframes the endpoint as an identity-centric control plane evaluating intent

to mitigate credential theft and session hijacking through continuous identity-state validation and active blast-radius reduction.

This architectural layer builds its own above-OS telemetry and baselines of why actors behave as they do, shifting endpoint defense beyond traditional process and file centricity. It rejects standard signature or rules-based EDR logic in favor of probabilistic, anomaly-based detection driven by local semantic inference. The underlying bet is that novel AI-era attacks are anomalies of intent, making them invisible to both legacy signatures and passive posture scans. By executing defensive logic directly at the endpoint runtime using small language models (SLMs), it interprets user and agent intent to intercept risky or non-deterministic workflows at machine speed before a malicious action can fully manifest.

Zone 4: Intent-Aware Behavioral Analysis

Concept: Shifts endpoint control from passive file or process monitoring to active assessment of user and agent intent. This layer validates the behavioral risk profile of AI-driven interactions by correlating actions with their underlying intent. It ensures that autonomous agent activity and human-initiated commands are continuously evaluated against security policy in real-time, effectively neutralizing risks like malicious prompt hijacking, unauthorized agent orchestration, and anomalous workflow execution.

Operational Pillar: Deploys granular, intent-aware enforcement gates that provide real-time, behavioral guardrails rather than static blocking. By integrating surgical user coaching and inline friction directly into the interaction flow, this mechanism replaces binary, disruptive termination with adaptive, intent-aligned guidance that secures the workflow while maintaining productivity.

Anchors on intent and behavioral context, rather than just the file or software artifact. The core bet is that you cannot judge whether an AI-driven action is dangerous unless you understand the motivation behind it and the behavior it exhibits. This intent-anchored layer evaluates the risk of an action based entirely on its behavioral trajectory and alignment with authenticated user or agent intent, providing the control plane for mitigating autonomous misuse, shadow agent orchestration, and workflow manipulation. Rather than relying on rigid rules, next-generation architectures in this zone utilize local intent inference to issue pre-action verdicts at the point of interaction.

Zone 5: Data-Centric Enforcement

Concept: Sits at the data itself, classifying, tracing, and intercepting sensitive data movement at the moment of creation, use, or exfiltration across the endpoint, browser, and SaaS/API plane.

Operational Pillar: Anchors on the data object and its lineage to achieve convergence with DSPM and insider risk platforms, employing on-device AI classification at the endpoint edge to handle cert-pinned or SASE-bypassed traffic directly.

Sits at the data itself, classifying, tracing, and intercepting sensitive data movement at the moment of creation, use, or exfiltration across the endpoint, browser, and SaaS/API plane. The hypothesis: one cannot govern AI-era risk without understanding what data was touched, by whom, through what workflow, and whether that movement was authorized. Unlike Zones 1–4, which focus on software posture, attribution, behavior, or browser enforcement, Zone 5 anchors on the data object and its lineage, making it the natural control plane for insider risk, AI prompt leakage, shadow builder exfiltration, and supply chain data theft.

What distinguishes Zone 5 from legacy DLP: Classic DLP relied on static policy rules, perimeter chokepoints, and pattern matching, generating ~90% false-positive rates and high operational burden. Next-gen Zone 5 vendors replace policy brittleness with local AI classification, context-aware intent inference, and coaching-first enforcement that reduces noise and preserves workflow continuity.

Market Landscape

The Market Landscape: Decoupling ECP from Legacy EDR

The shift toward Endpoint Control and Prevention (ECP) signifies a fundamental decoupling of endpoint security from the constraints of legacy, process-tree-based detection models. As traditional security perimeters dissolve into dynamic, SaaS-mediated execution environments and autonomous agentic workflows, the market for defense is rapidly evolving. To address this complexity, the vendor landscape has begun to fragment, coalescing around three distinct strategic notable motions:

- **Software Posture & Supply-Chain Governance:** Treating the endpoint as an app store. Prioritizing the discovery and configuration risk-ranking of extensions, IDE plugins, and agentic toolchains rather than binary malware.
- **Intent-Aware Observation & Semantic Defense:** Moving beyond *what* a process did to *why* a process occurred. **Ent** and **Origin HQ** provide high-fidelity telemetry that maps prompt-to-action, establishing the forensic evidence chains required for governance.
- **App-Layer Enforcement & Browser-Centric Control:** Leveraging the browser as the primary control surface. Vendors like **Neo Security** and **Cyberhaven** move enforcement directly into the execution context (the browser or app), enabling surgical warn/block/guide interventions without the latency of kernel-level processing.

Emerging Players in Endpoint Control and Prevention (ECP)

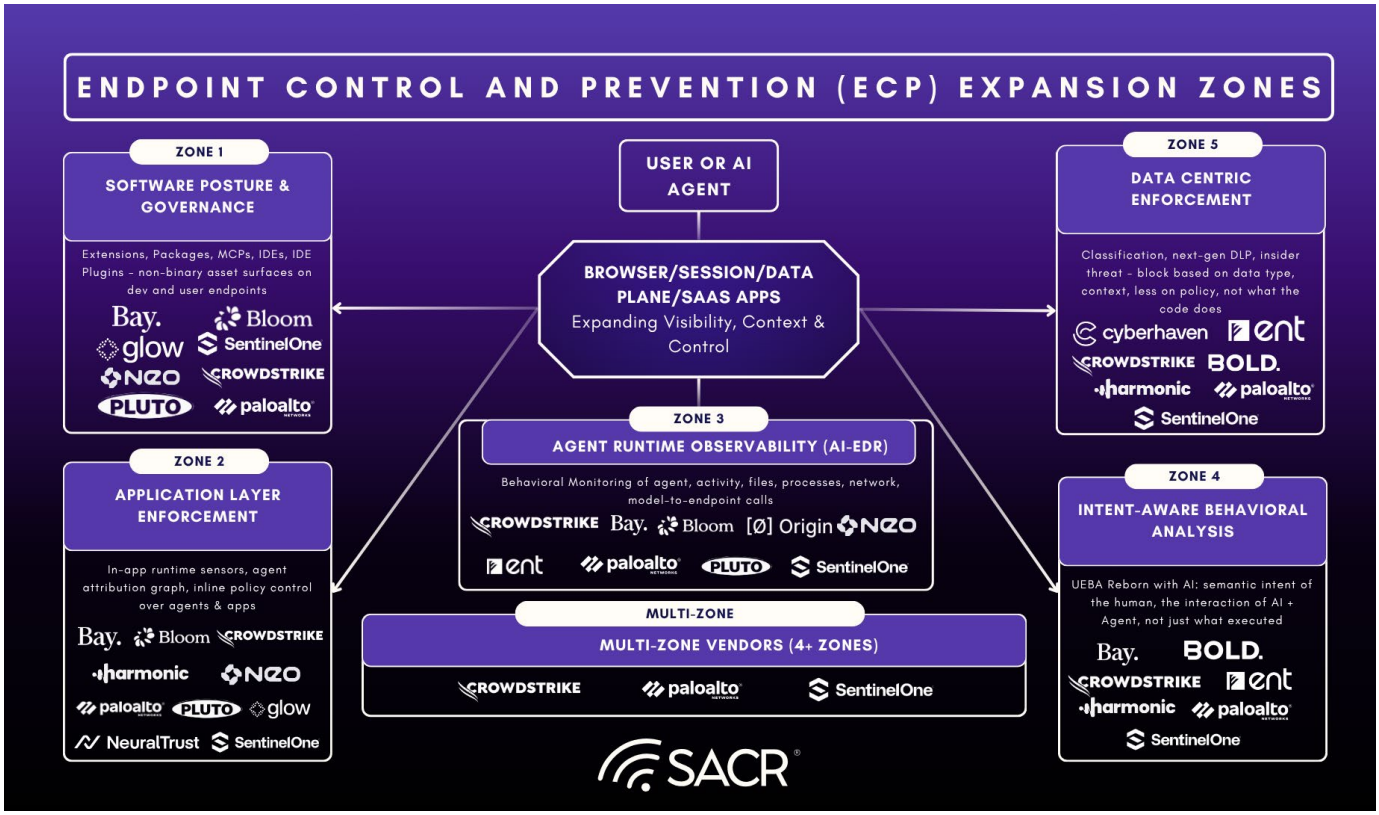

Source: SACR Briefings 2026 (Sampling of Notable Vendors)

Agentless vs. Agent-Based Trade-offs

The deployment debate has persisted for decades, with many buyers rightfully wary of the systemic stability risks inherent in legacy kernel-integrated drivers, most notably the catastrophic Blue Screen of Death (BSOD) events. While the adoption of Extended Berkeley Packet Filter (eBPF) technology has significantly neutralized these driver-conflict anxieties, a strategic appetite for agentless architectures remains. The primary driver for agentless adoption is the elimination of net-new software footprints, allowing organizations to achieve specialized outcomes, such as advanced forensic depth, novel data correlation, or unique AI-driven inference, without the operational friction of traditional sensor management. Whether delivered via API integration or cloud-mediated browser plugins, endpoint harnesses, or agentless solutions, they might provide the surgical, complementary telemetry required to bridge functional gaps without duplicating the existing endpoint protection stack.

Agentless Deployment Performance Claims May be Misleading

The value proposition of agentless deployment includes that they ensure zero performance degradation, which isn't entirely true. If a vendor solution hits customer API rate limits, or maybe query maximums, the solution may have significant gaps in data freshness and quality, leading to potentially poor alignment to detections or lack of the latest functionality releases. But on the endpoint, in some cases agentless solutions can significantly augment the current endpoint software's use cases, either delivering new telemetry, data science or machine learning approaches of their own to a complementary endpoint protection tool. The primary drivers for adoption of agentless is for a customer to avoid net-new software or processes deployed, thereby optimizing endpoint battery life, and in the eyes of many buyers, the overall user experience.

Rapid Deployment and Scaling

Agentless solutions facilitate instantaneous deployment, typically requiring only a service account, an application API key, or standard credentials. Scalability is achieved seamlessly without necessitating reboots or software installations. Centralized control mechanisms of agentless solutions can sometimes simplify management, particularly for platforms utilizing differentiated or advanced reporting or offering analytic graph interfaces that incumbents aren't providing out of the box. By minimizing the local attack surface, agentless options, if used to unify datasets and operations, help somewhat strengthen an organization's security posture while maintaining compatibility and lowering the total cost of ownership (TCO). This aligns with a broader cybersecurity trend where specialized vendors integrate synergistically with existing infrastructures to deliver enhanced capabilities, such as data enrichment, advanced workflows, and deep analytical graphing.

Vendor Use Case Limitations and Dependency Issues

Agentless architectures present clear operational constraints. Chief among these are reduced context and visibility stemming from a lack of their own, controlled local file system access, drivers, or parsing engines, which can hinder deeper forensic investigations, delivery of new functionality, use cases or performance while they also remain fundamentally dependent on the capabilities and rate limits of provided APIs. Reliance on network connectivity introduces latency risks and potentially restricts real-time response actions when local remediation alternatives are absent. Organizations should implement a hybrid defense model with preference towards agent-based solutions where possible for the lowest latency of endpoint defensive measures. In some situations, customers of endpoint software might want agentless monitoring for high-performance locations or IoT environments that have limited processing capacity, where agentless may be more ideal, while deploying agent-based ECP solutions across critical endpoints that demand comprehensive file inspection and immediate, machine-speed containment and prevention capabilities.


cyberhaven

Zone 5:
Data-Centric Enforcement

Cyberhaven

Vendor Profile

Cyberhaven is a data lineage and browser-centric Zone 5 governance platform that tracks the origin, movement, and destination of sensitive data across endpoint, browser, SaaS, and cloud environments, providing high-fidelity evidence chains of user data interactions. Unlike traditional DLP tools that classify data by content pattern at a point in time, Cyberhaven builds a continuous lineage graph that follows data from its point of origin- a document created, a record copied from a database, a file downloaded from a cloud service- through every subsequent interaction, transformation, and transfer. This lineage model enables security teams to answer the forensic question that matters most in AI-era investigations. Cyberhaven's browser-centric architecture captures SaaS and web application data interactions that endpoint-only agents miss, making it particularly relevant for organizations where the majority of sensitive data movement occurs within browsers interacting with cloud-based AI tools, SaaS platforms, and collaboration applications with on-device AI classification as the primary architectural differentiator.

Applicable ECP Zones:

- **Zone 5**

Products/Services Overview

- **Data Lineage Platform:** Continuous tracking of sensitive data origin, movement, and destination across endpoint, browser, SaaS, and cloud environments
- **Browser-Centric Governance:** Browser extension-based data movement monitoring providing full visibility into SaaS application and web-based data interactions
- **Data Movement Evidence Chains:** High-fidelity forensic evidence chains of user data interactions for incident investigation, governance, and compliance reporting
- **AI Tool Data Flow Monitoring:** Visibility into data flows into and out of AI applications accessed via the browser

Core Functions

- **Continuous Data Lineage Tracking:** Follows sensitive data from its origin through every subsequent interaction, transformation, and transfer across endpoint, browser, and cloud, building a persistent lineage graph.
- **Browser Data Movement Capture:** Instruments the browser to capture the full richness of SaaS and web application data interactions, including copy-paste, upload, download, and AI tool prompt submissions.
- **Forensic Evidence Chain Generation:** Produces investigation-grade evidence chains that answer the origin, path, and destination questions for any sensitive data movement event.
- **AI Tool Data Flow Visibility:** Monitors data flowing into and out of browser-accessed AI tools (ChatGPT, Copilot, Gemini), providing governance coverage for AI-mediated data movement.
- **Cross-Environment Lineage Correlation:** Correlates data movement events across endpoint, browser, SaaS, and cloud to provide a unified view of sensitive data flows across the enterprise.

Use Cases and Pain Points Addressed

1

Insider Threat and Exfiltration Forensics: Provides complete data lineage for forensic investigation of insider threat and exfiltration incidents, answering origin, path, and destination questions that traditional DLP cannot.

2

AI Tool Data Governance: Monitors and governs sensitive data flowing into browser-accessed AI tools, detecting prompt leakage and unauthorized AI-mediated data movement.

3

Departing Employee Data Exfiltration Detection: Detects and investigates data exfiltration by departing employees by tracking sensitive data movement patterns in the period leading up to departure.

Key Takeaway / Recommendation

Cyberhaven is a strong Zone 5 data lineage platform, offering the most mature and proven browser-centric data governance capability in the market. For organizations prioritizing forensic evidence quality, insider threat investigation depth, and AI tool data flow governance, Cyberhaven provides capabilities that legacy DLP and most next-gen alternatives do not match on lineage depth. CISOs should evaluate Cyberhaven as the data lineage and evidence chain foundation of their Zone 5 stack. The key consideration is that Cyberhaven's primary strength is lineage and forensics rather than real-time prevention; organizations requiring pre-action enforcement should layer Cyberhaven with a prevention-first enforcement vendor rather than treating it as a standalone DLP replacement.

Execution Strategy for the CISO and Security Team

SACR Five Strategic Zone Expansion Observations

1 SACR believes the zones converge into one architecture within ~3 years.

Posture without enforcement is a report, enforcement without intent context is a false-positive machine and observability without prevention loses at machine speed. The winning stack is posture, inline gate, intent baseline and prevention enforcement in one user-mode agent.

2 Local AI on the endpoint becomes a defining moat.

The cloud-pipeline cost structure of incumbent EDR becomes a liability at agentic telemetry volume. CrowdStrike's network-dependent AIDR architecture is the specific vulnerability startups are targeting, as is Bold's on-device classifier thesis. Agentics and small language models and generative AI-enabled behavioral defenses are emerging to help defend the endpoint (eventually leading to on-endpoint agentic defense).

3 Attribution becomes the new detection.

The question shifts from "is this process malicious?" to "which prompt caused this action, and was it the user's intent?". Prompt and action lineage (Origin) and intent verification (like with Ent) are early versions of what becomes a required capability, the AI-era equivalent of the process tree.

4 The agentless wedge is real but time-boxed.

Vendors ride-the-EDR model wins the next 18 months on deployment friction. As enforcement (not visibility) becomes the buying criterion, API depth and real-time enforcement capabilities limits will be a deciding factor. Visibility is commoditizing, while enforcement, telemetry depth, federation and speed of context are where we see the margin settle over time.

5 Platforms win the category, startups define it.

The most likely outcome mirrors CSPM and Wiz: one or two startups (Zone 2/Zone 3 architectural natives best positioned) escape to platform scale while the rest are acquired. CrowdStrike AIDR's attach rates prove the demand; the Koi acquisition shows platforms will buy rather than build the new zones.

Emerging Goal Pillars of Endpoint Control and Prevention

The modern Endpoint Control and Prevention (ECP) framework is defined by four critical operational pillars:

App Posture & Discovery:

Continuous governance of the non-binary software supply chain, including extensions and plugins.

AI/Agent Observability & Attribution:

Full traceability from natural language prompts to downstream tool execution.

Intent, Identity, and Behavioral Inference:

Shifting defense from binary monitoring to understanding the motivations behind user and agent actions.

Real-time Policy Intervention:

Surgical, inline enforcement that interrupts risky workflows mid-stream without disrupting legitimate business activity.

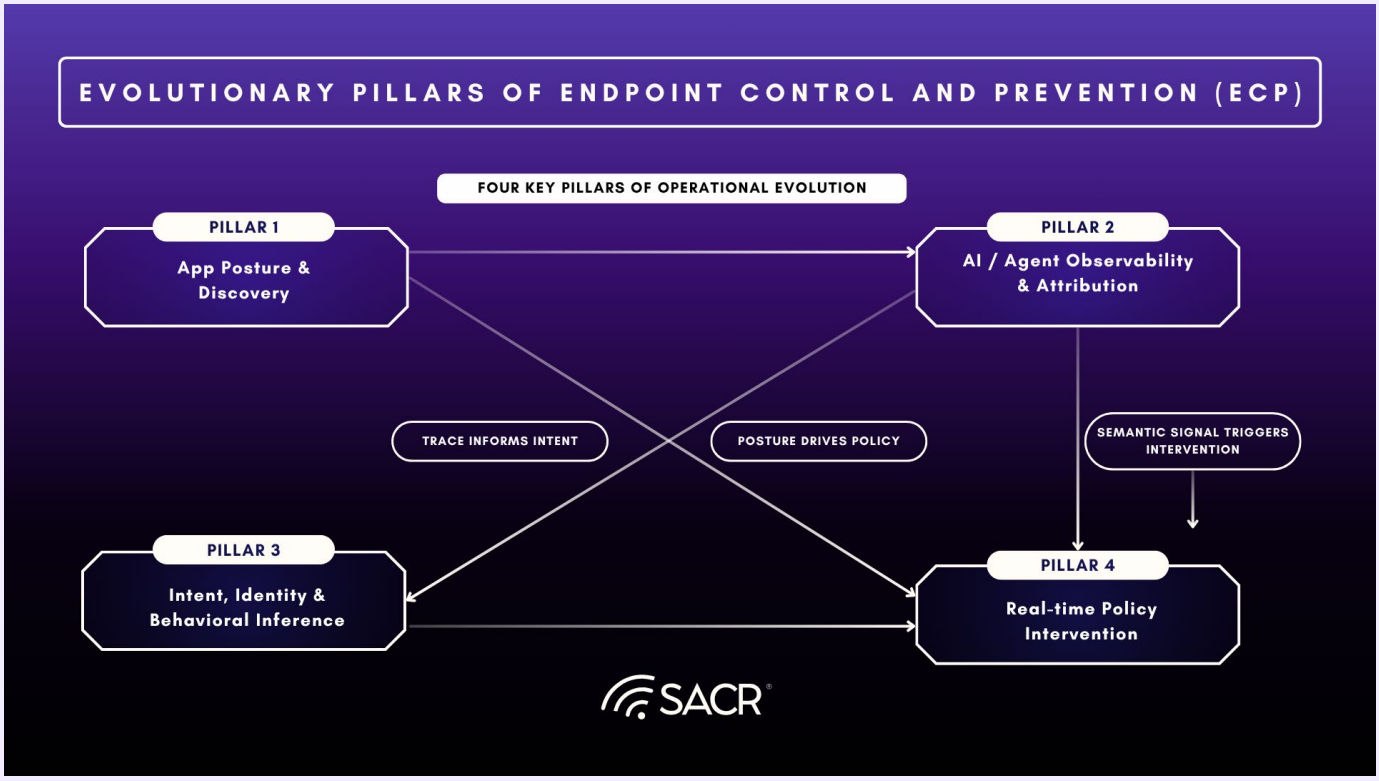

Investigation-Grade Evidence Required to Reconstruct Why, What and How

When an attack occurs, local telemetry is often the first and most detailed source of truth. However, raw endpoint logs (like standard Event IDs or process creation lines) lack the context required for sophisticated modern forensics. Endpoint agents must evolve into intelligent data collectors that correlate activities into a cohesive narrative. It isn't enough to log that an administrative tool was run, a software or website was used; the endpoint must capture the prompt, tool, or automated workflow that triggered it. This allows SecOps to determine if a sequence of commands was executed by a human attacker, a benign developer, or an autonomous AI agent or toolchain.

Safer Enforcement: Moving Beyond Binary Decisions of Allow or Block through AI and Context Evaluation

Historically, endpoint detection and response (EDR) tools relied on binary enforcement where either a file was allowed to run, or the process was abruptly terminated and the machine isolated. This heavy-handed approach frequently breaks critical business workflows and alienates users. Modern endpoint defense requires graduated, human-centric controls integrated directly into the user interface. When a user attempts a risky but non-malicious action (e.g., running an unsigned IT tool or interacting with an unapproved API), the endpoint should inject real-time friction. This means warning the user, guiding them toward a safer alternative, or requiring a justified exception request that is immediately audited and passed to IT or GRC for review, rather than defaulting to a hard block. Contextual evaluation capabilities of LLMs have proven to reduce behavioral analysis of false positives by enabling cognitive evaluation by AI and agents, producing fewer false positives and more probabilistic-style responses.

Shifting from Posture Visibility to Proactive Posture Management and Control

The most disruptive change is the shift from endpoint security to posture security where the endpoint is no longer a boundary but a dynamic, self-optimizing component of the enterprise's security ecosystem. The endpoint will be defined by identity, protected by identity, and continuously hardened by automated, intelligent systems. The technological building blocks are already in place: LLMs for reasoning, autonomous agents for self-healing, cloud-native and edge-native architectures for cross-platform security, and predictive posture for proactive defense. The next several years will see the convergence of these technologies into ECP, the next evolution of endpoint security. Endpoint defense must dynamically discover new micro-softwares and non-binary tools, rank their risk based on the permissions they hold (such as reading browser data, configurations, identities), evaluate source code repositories and supply chains for compromise, and apply proactive policy controls to prevent supply chain compromises arriving on endpoints.

SACR Future View: The Transition to Autonomous Endpoint Defense

By 2027, artificial intelligence and machine learning are fundamentally shifting Endpoint Detection and Response (EDR) from reactive, manual tools to proactive, autonomous, and self-healing systems conveying contextual information and telemetry in a semi-federated fashion.

The next horizon of endpoint defense is the migration of security intelligence from cloud-backend SIEM/SOAR clusters to the endpoint itself. As autonomous agents operate at machine speed, the traditional observe, export, analyze, respond cycle is becoming a structural failure point as the latency between detection and containment is now the primary attack vector.

The future of the endpoint is a self-governing runtime environment capable of endpoint and connected edge Inference. This requires a shift from passive telemetry collection to active, stateful full context defense, where endpoints utilize localized models (SLMs) and most likely include Edge services and integration with its telemetry to evaluate intent, govern non-binary software composition, and mediate agentic tool-chains in real-time, evolving from ECP to AECP over time.

ECP Zones as an Endpoint Control Plane

The architecture below presents The ECP Control Plane as a three-part architecture that moves from endpoint evidence to policy enforcement and verified outcomes. On the left, OS-layer signals, normalized endpoint context, and posture or behavioral state are consolidated into a validated endpoint evidence state. This evidence feeds a central, multi-zone enforcement model covering OS controls, software posture and supply chain governance, application-layer enforcement, agentic runtime visibility, intent and behavioral analysis, and data-centric enforcement. The model then applies zone-scoped write authority across endpoint agents, identity stores, browsers and SaaS applications, AI toolchains, and data egress layers, ultimately producing an enforced endpoint posture. A validation rail beneath the architecture logs each action, verifies the result against zone policy, preserves the evidence and authority record, and feeds the outcome back into the behavioral baseline.

ECP Zones as an ECP Control Plane

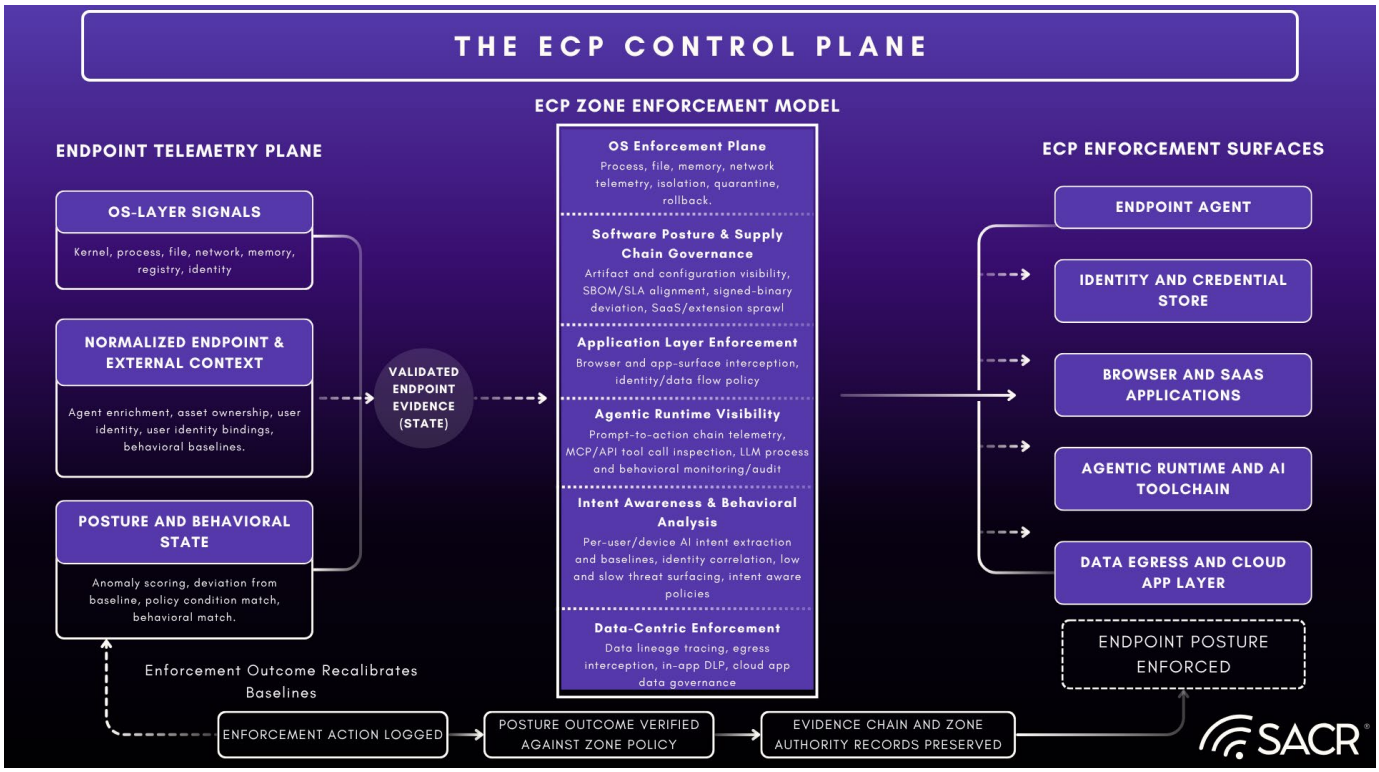

Key Requirements for Endpoint Control and Prevention

To enable autonomous Endpoint defense, security architectures must evolve to monitor these specific high-fidelity telemetry streams:

Prompt-to-Tool Lineage:

Real-time tracking of input prompts, the specific tools invoked by the agent, and the associated data movement.

Non-Binary Composition State:

Continuous monitoring of browser extensions, IDE plugins, and dependency packages (npm, PyPI) that alter the endpoint's function post-deployment.

Local Intent Inference:

Evaluation of natural language intent against security policies using lightweight, on-device models to detect semantic anomalies before execution.

Identity State Context:

Localized tracking of active session tokens, OAuth grants, and API keys, ensuring that sensitive identity artifacts are only accessible to authorized workflows.

Dynamic Policy Enforcement:

Monitoring and identification of shadow builder activity with automated, no-code and low-code workflows assembled by non-technical teams, to assess configuration and permission risks at the moment of creation.

Market Map: Endpoint Control and Prevention

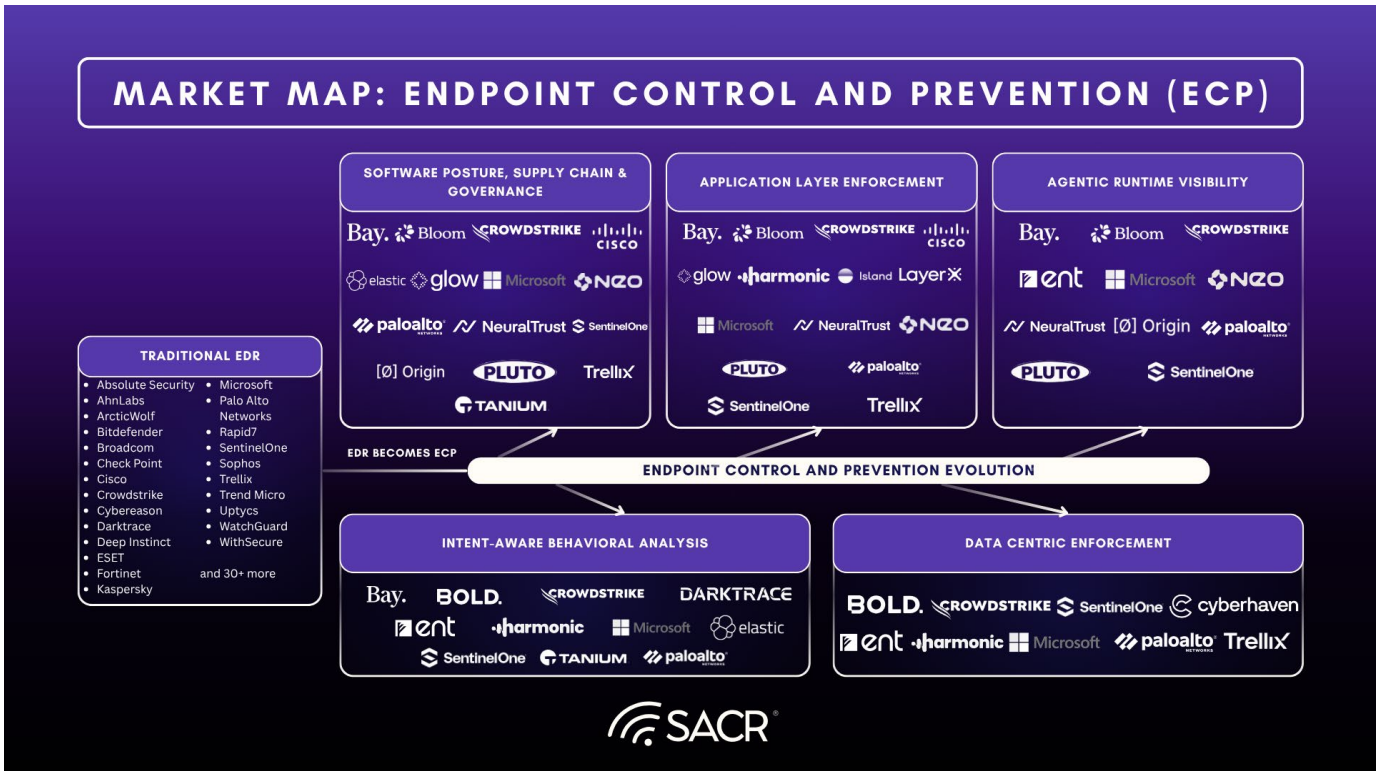

Future View: Endpoint Control and Prevention (ECP)

The integration of advanced AI and capabilities such as model context protocol driving the federation of agent-driven integrations are driving several major transformations in how EDR operates:

The Rise of Autonomous Agents

Endpoint security is entering its fifth generation, characterized by autonomous security agents. Instead of simply alerting human analysts, agentic AI is being deeply embedded into daily security operations, allowing systems to analyze massive volumes of heterogeneous data, make independent decisions about threat severity, and automatically trigger defensive reactions. This drastically reduces the mean time to respond (MTTR) from hours or days to just minutes, which is critical for containing fast-moving threats like ransomware.

New Attack Vectors Targeting AI

As AI becomes central to ECP, the AI models and agents themselves become high-value targets. Security teams must now defend against adversarial machine learning, where attackers manipulate input data to deceive AI detection algorithms. Because AI agents can crawl data and execute shell commands, they are vulnerable to indirect prompt injection attacks hidden within files or web pages. Modern ECP platforms are expanding to monitor AI usage, protect against data leakage, and secure the underlying AI infrastructure.

Advanced Behavioral Detection and LLM Integration

AI is replacing outdated signature-based detection with sophisticated behavioral analytics. EDR platforms are leveraging a combination of supervised learning, unsupervised learning, reinforcement learning (RL), and deep learning (DL) neural networks to recognize complex, novel attack patterns that traditional tools miss. Furthermore, Large Language Models (LLMs) are being deployed to parse massive amounts of unstructured data, such as chat logs, security logs, and emails to classify anomalies and conduct post-incident analysis by revealing an attacker's motives and methods.

Self-Healing Endpoints and Automated Remediation

Perhaps the most significant advancement is the introduction of self-healing capabilities. When an endpoint is compromised, AI-driven ECP can automatically isolate the device, notify identity systems and backend telemetry via model context protocol (MCP), terminate malicious processes, remove persistence mechanisms, and roll back unauthorized system changes or encrypted files to a known-good, pre-infection state. This automated remediation and contextual sharing and response limits the blast radius of an attack and dramatically cuts down on the need to manually re-image devices.

Adaptive Security Policies

Rather than relying on rigid, static rules, AI empowers EDR and cloud security systems with continuous self-learning. These systems can dynamically adjust access controls, firewall settings, and anomaly thresholds in real time based on evolving attack patterns, ensuring resilience against zero-day exploits and environmental changes.

Human-Machine Collaboration

Despite the push for autonomy, a major theme for 2026 is balancing AI with human control so that algorithms do not go unchecked. The most effective SOCs utilize a collaborative model: autonomous systems handle routine threat containment and initial data correlation, while escalating the most sophisticated attacks to skilled human analysts. Human analysts then provide feedback that is used to continuously train and refine the machine learning models, improving future detection accuracy and reducing false positives.

Future View: Agentic Endpoint Control and Prevention (A-ECP)

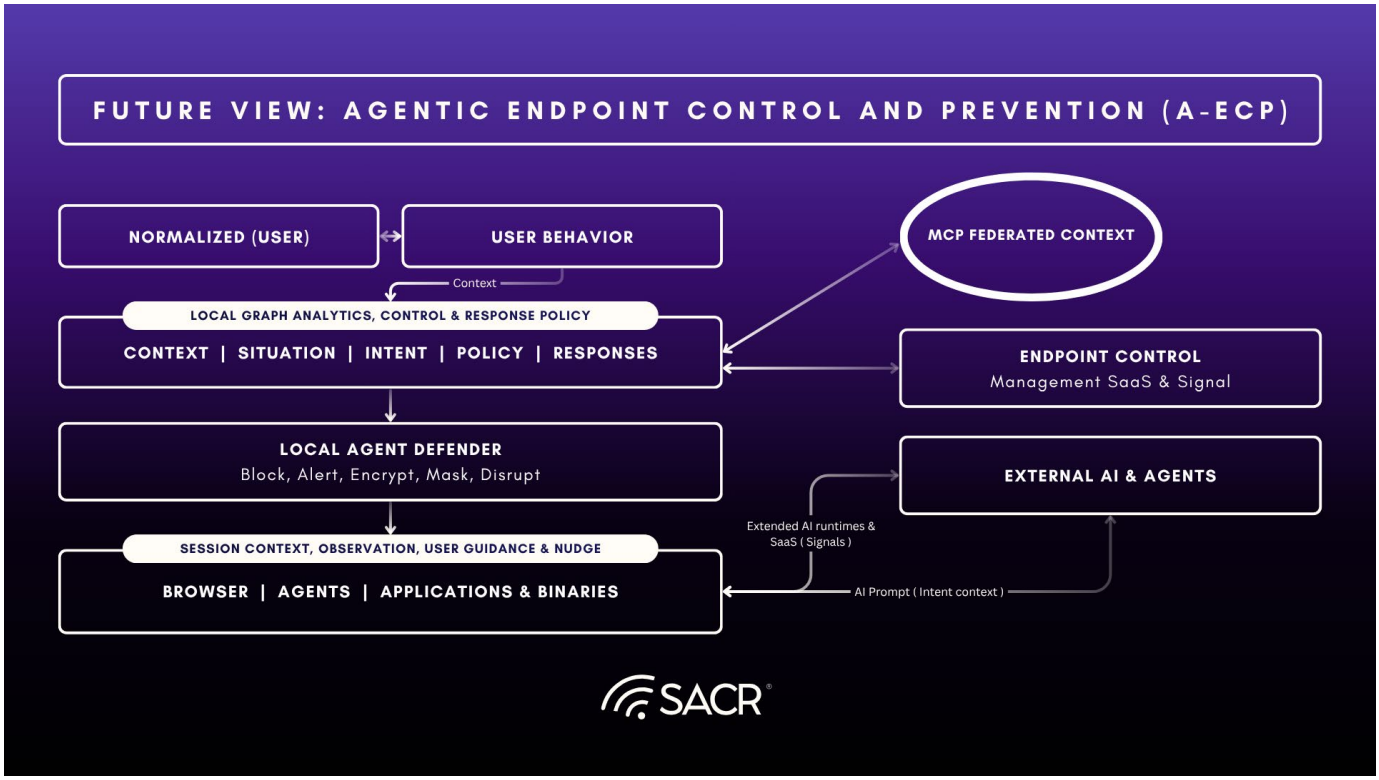

Implications for CISOs / Buyers: 30/90 Day Plan

Buyer strategy should assume rapid convergence: incumbent EDR will add more app-layer posture and identity context, while posture/agent-control vendors will add response and evidence features. CISOs should leverage the SACR 5 Zones framework as the primary architectural tool for identifying coverage gaps and mapping their infrastructure needs. The practical evaluation lens is: evidence quality, intervention safety, and measurable exposure reduction.

30 Days: Assess and Baseline

1. **Audit existing security coverage against the SACR 5 Zones framework** to identify specific architectural coverage gaps and avoid redundant tooling.

2. **Treat classic EDR as baseline plumbing, not the strategic differentiator.** Reallocate budget and evaluation energy toward application-layer posture, agent observability, and intervention safety.

3. **Run privacy and governance diligence early.** Recognize that intent layers require new, sensitive telemetry (e.g., interaction traces). Validate your organization's standards for minimization, purpose limitation, and auditability.

4. **Track pricing/model changes as a strategic signal.** Monitor endpoint per-seat pricing compression, which signals detection commoditization. Look for consumption/usage models that better align with modern, bursty agentic workflows.

90 Days: Validate and Operationalize

5 **Demand prompt-to-action attribution in any agentic tooling environment.** Ensure any selected vendor can connect intent/prompt → action trace; without this, investigations will fail at machine speed.

6 **Prioritize enforcement primitives that are surgical.** Shift away from blanket block/allow strategies. Evaluate flow-level, policy-based interventions and exceptional UX that won't break critical business workflows.

7 **Evaluate autonomy.** Evaluate autonomy as an operations capability, emphasizing the integration of new ECP signals into existing SOC workflows rather than mere validation. Ensure that these signals are actionable in your current incident response processes.

8 **Prepare for category convergence.** As EAPM/posture vendors add interventions and intent vendors add posture, focus your final evaluation criteria on the quality of control primitives and evidence.

SACR Key Takeaway:

For CISOs, the transition to Endpoint Control and Prevention (ECP) is not merely a tool upgrade but a strategic necessity to reclaim visibility in an era of AI-driven, machine-speed threats. As the center of gravity shifts from process-based detection to interaction-centric governance, the key takeaway is that your endpoint strategy must prioritize prompt-to-action attribution and granular, surgical enforcement over legacy, binary blocking.

By pivoting to an ECP-ready/AI oriented architecture, focused on software supply chain posture, identity-centric context, and autonomous behavioral inference, you can effectively reduce the blast radius of agentic workflows while ensuring your security posture remains resilient against industrialized, high-speed exploitation.

Software Analyst®
Cyber Research

business

personal

softwareanalyst.io

SACR®

Trusted research. Sharp insights. Real conversation.

CISO	VENDOR
SECURITY TEAMS	INVESTORS