SELF-HEALING CYBERSECURITY: AI IN ACTION

Autonomous and cost-aware cyber defense that links security investments to measurable reductions in risk, cost, and disruption.

By Matt Salmon - VP of Cybersecurity & AI Cyberhill Partners

Cyberhill Perspective

Enterprises generate more security telemetry than human teams can interpret, while adversaries automate, pivot, and monetize at machine speed. Cyberhill is building the next decisive advantage on two pillars: artificial intelligence that can reason and act, and cybersecurity ontologies that provide a shared, structured understanding of assets, identities, threats, and controls.

Together, they enable a self-healing cyber defense posture—a closed loop that senses threats, understands business impact, decides on the least-risk action, advises remediation, and learns to improve continuously.

The payoff is measurable: shorter detection and response times, lower loss expectancy, more transparent board communication, and a leaner, more efficient security spend.

The Case for Change

Most organizations rely on a patchwork of point tools—such as SIEM, EDR/XDR, IAM, CSPM, NDR, WAF, and vulnerability management—each with its own data model and competing for attention. Analysts spend precious time stitching together partial pictures, while CISOs struggle to translate technical metrics into business terms. Even mature dashboards remain primarily descriptive: they tell you what happened, but not what it means for revenue operations or the next best action at the lowest cost and risk.

This fragmentation also obscures financial clarity. Licenses accumulate, modules remain dormant, features overlap, and budgets grow without a clear narrative of value. Boards rightly ask whether the organization is overspending, whether it is secure, and where it can cut costs without increasing risk. Today's dashboards rarely address all three questions at once.

Why Ontologies Matter

An ontology is a formal, machine-readable model of a domain. In cybersecurity, it details the key entities—business processes, applications, data classes, endpoints and workloads, identities and privileges, networks and cloud services, vulnerabilities and misconfigurations, adversary behaviors and campaigns, controls and playbooks—and the relationships between them. Crucially, it also captures cost and utilization semantics: which tools and features protect which assets, at what cost, under which policies, and with what results.

When telemetry from different systems is standardized into a shared graph, AI can interpret it with context. A spike in failed logins is no longer an isolated alert; it is a potential credential-stuffing campaign that threatens the customer portal, which underpins revenue recognition and can have a quantifiable financial impact if degraded. An unpatched CVE isn’t just a score; it’s a vulnerability on a workload that handles regulated data, already targeted by the adversary's known tactics, partially mitigated by existing controls, and cost-effective to remediate via an approved playbook. Ontology transforms data into knowledge and knowledge into decisions that the enterprise can rely on.

From Visibility to Self-Healing

A self-healing posture is best viewed as a control loop. First, the environment senses logs, signals, and configuration states, which stream into the semantic layer and are resolved into consistent entities. Next, the system recognizes that AI enhances signals with business context, exposure data, and adversary tradecraft, creating a causal picture of current and future possibilities. Then, the system decides: a policy-aware engine evaluates candidate actions—such as blocking, isolating, rotating keys, stepping up authentication, segmenting networks, patching, or doing nothing—against potential disruption, residual risk, and cost. Afterward, it acts through automation, utilizing SOAR playbooks, identity workflows, infrastructure-as-code, and cloud policies to execute changes safely, with rollback paths and audit trails. Finally, it learns: outcomes are fed back into the models, improving detection accuracy, refining playbooks, and updating risk forecasts.

Self-healing does not imply removing humans. It means empowering them. Sensitive actions may require approval gates; change windows and kill switches remain in place; all decisions are transparent and explainable. Over time, the system gains autonomy by demonstrating reliability in situations where stakes are lower and consequences are reversible, allowing experts to focus on new threats and strategic planning.

"Chat With Your Data": Natural-Language Security Operations

Once knowledge is structured, language models become dependable assistants rather than creative guessers. CISOs and analysts can interrogate posture directly: Which critical vulnerabilities affect systems linked to revenue operations, and what is our patching latency? Which vendors present the largest blast radius if compromised? Where are we paying for features we aren't using, and what overlaps can be cut without reducing coverage?

Because questions are based on the ontology, answers are precise, traceable, and aligned with policy. Conversational access reduces the gap between curiosity and decision, turning the dashboard from a static reporting tool into a truly interactive command center.

Which critical vulnerabilities affect systems linked to revenue operations, and what is our patching latency?

Which vendors present the largest blast radius if compromised?

Where are we paying for features we aren't using, and what overlaps can be cut without reducing coverage?

Making Money Talk: Cost Analytics with Teeth

Security programs gain trust when they demonstrate value in both risk mitigation and financial savings. Ontologies make this possible by linking financial context to technical details. Each tool and module is represented by its license model, features, and deployment footprint; each feature is linked to the domains and assets it actually protects; and each control is associated with the incidents it prevents, the time it saves, and the disruptions avoided.

This structure supports key metrics that matter to executives: cost per tool and cost per feature, not in isolation but in relation to the protection provided; features in use versus features licensed, producing a genuine utilization rate by tool and across the stack; quarterly spend that aligns to budget cycles and forecasts; year-over-year trends that correlate investment to measurable reductions in MTTD, MTTR, and incident recurrence; overlap indices that reveal duplicative capabilities ripe for consolidation; and cost-to-risk reduction that ties dollars to decreases in annualized loss expectancy.

The result is a board-ready narrative: here is what we spend, here is how effectively we use it, and here is where optimization will lower cost without increasing risk.

A Pragmatic Path to Adoption

The journey begins with clarity, not code. Inventory the tools, licenses, and features you already fund; identify the business processes that truly matter; codify the policies that govern change. Establish a knowledge graph that aligns with widely used security taxonomies and your internal naming conventions. Normalize telemetry and configuration states into that model so that assets, identities, exposures, and controls are represented as durable entities.

Early wins should be visible and secure. Launch conversational queries that replace spreadsheet archaeology with instant answers—highlight utilization and overlaps so budget discussions are based on facts. Introduce assisted automation in low-risk playbooks—token revocations, IP blocks, ticket enrichment—and evaluate the impact on response times and analyst workload. As confidence increases, move on to medium-risk actions with rollback options—such as segmentation, WAF tuning, and secret rotation—and integrate policy-as-code so governance is automated alongside response.

Throughout, treat the platform as a product. Monitor model drift and data quality, maintain auditable decision logs, and run chaos experiments to validate fail-safes. The goal is not to achieve maximum automation on day one but to implement reliable automation that expands as it demonstrates its effectiveness.

Governance That Builds Trust

Autonomy must be earned. Sensitive workflows should include explicit approval gates and change windows. Every automated action should be explainable in terms that non-technical stakeholders can understand, and every decision should generate artifacts suitable for audit and review. Data minimization, access controls, and privacy constraints should be encoded in the ontology itself, rather than being bolted on after the fact. Finally, the program requires a model-risk discipline, including testing for adversarial inputs, monitoring performance, and retraining on a defined cadence.

Governance isn't a barrier to innovation; it's the system that allows innovation to be scaled.

Proving It Works

Boards and CEOs do not buy architecture diagrams; they buy outcomes. A concise scorecard should track improvements in detection and response, reductions in high-severity recurrence, increased feature utilization, retirement of overlapping tools, and demonstrable decreases in annualized loss expectancy. Track these metrics quarterly and annually, and pair the numbers with brief narratives—what the system learned, where it acted autonomously, where it requested approval, and what was saved in dollars and disruptions.

When the ontology serves as the backbone and AI functions as the brain, these stories write themselves from the evidence.

Conclusion

AI alone will not revolutionize cyber defense; it requires a solid foundation. Ontologies provide that base, converting fragmented telemetry into a unified, business-aware view of risk and control. With this foundation, automation becomes safe, explainable, and increasingly effective. The result is a self-healing cyber defense posture. This adaptive system senses, understands, decides, acts, and learns faster than adversaries can iterate, clearer than dashboards can describe, and leaner than budgets once believed possible.

Enterprises that move first will establish a new standard for resilience and cost discipline. Those who follow will adopt it as best practice. Those who wait will see it as hindsight.

See the Self-Healing Loop in Action

Experience how AI and ontologies enable autonomous, cost-aware cyber defense. In a Cyberhill strategy session, we'll walk you through the self-healing control loop — sensing, deciding, and acting at machine speed — and show how it links directly to business outcomes that matter to you.

BOOK A CYBERHILL SESSION