Skyrelis Agentic Security Platform

Always-On Protection for LLM Agents, Tools, and Data

Executive Summary

Agentic AI is the next SaaS moment—those who secure it will define the future.

Enterprises today face a pivotal inflection point. Agentic AI promises up to $15.7 trillion in productivity gains by 2030, but only if it's trusted, governed, and secure. At the same time, cybercrime is on track to hit $10.5 trillion annually by 2025, and the rise of autonomous AI systems only raises the stakes. The lesson from the SaaS era is clear: just as cloud-native security unlocked SaaS growth, runtime security will determine the winners of the Agentic AI era.

Skyrelis is purpose-built for this moment.

Our Agentic Security Platform delivers always-on protection for LLM-driven, multi-agent workflows—securing data flows between users, agents, and clouds across AWS, GCP, and beyond. With lightweight deployment, real-time behavioral monitoring, dynamic policy enforcement, and protocol-layer inspection, Skyrelis stops misuse, leakage, and escalation risks—without slowing development.

Every $1 invested in AI security returns $5–$8 in productivity. Skyrelis helps enterprises capture that upside, placing them in the "Skyrelis Zone," where security and AI adoption scale together, without compromise.
Built for speed, simplicity, and scale, Skyrelis is fully compatible with emerging standards like Model Context Protocol (MCP) and Agent-to-Agent Protocol (A2A), ensuring future-proof interoperability across rapidly evolving AI ecosystems.

The Challenge

The promise of Agentic AI is massive—but so is the exposure.

Enterprises poised to capitalize on the $15.7T productivity opportunity must first navigate a rapidly shifting threat landscape. Legacy security architectures, built for static applications and human-driven workflows, are blind to the dynamic behaviors of AI agents. Without purpose-built protection, the risks multiply—not just in terms of cybercrime's predicted $10.5T impact but in lost trust, regulatory non-compliance, and irrecoverable data leakage.

Three Critical Challenges stand in the way:

- **Invisible Attack Surface:** Agentic systems operate autonomously, generating real-time interactivity between agents, clouds, and APIs. Traditional security tools fail to observe or control these north-south (agent-to-tool) and east-west (agent-to-agent) data flows, leaving enterprises vulnerable to prompt injection, data exfiltration, and unauthorized actions.
- **Fragmented Control:** Multi-agent ecosystems are highly dynamic, coordinating tasks across internal systems and third-party APIs. This fluidity outpaces perimeter defenses and static IAM policies, creating security blind spots where agents can take unintended or malicious actions.
- **Compliance & Audit Gaps:** Most agentic observability solutions provide logs—but they are developer-centric and lack intuitive visualizations. This makes it difficult for security teams to trace agent decisions, identify vulnerabilities, or demonstrate compliance with GDPR, HIPAA, and SOX.

Enterprises can't afford to choose between velocity and security. That's why Skyrelis exists—to fill this critical gap by making agent behavior visible, controllable, and compliant.

Our Solution: Skyrelis Agentic Security Platform

Skyrelis provides always-on protection for AI agents through an **SDK-independent architecture** that simplifies and accelerates cost-effective deployment. Skyrelis monitors every agent action—whether accessing data, invoking tools, or communicating with other agents—and routes it through the Skyrelis edge for policy enforcement.

This applies runtime controls to block sensitive data exfiltration, intercept multi-step prompt attacks, and detect abnormal behaviors like unauthorized role escalation. Every interaction is logged with contextual risk scoring and audit trails, enabling security teams to trace decisions and remediate threats in real time.

SDK-Dependent vs. SDK-Independent Operations

	SDK-Dependent (TODAY)	SDK-Independent (SKYRELIS)
Policy Updates	Requires developer code changes and app redeployment	Security teams push policies directly—no code changes
Time to Enforce	Days to weeks (tied to sprint cycles)	Minutes to hours (real-time enforcement)
Control Owner	Developers	Security Teams
Deployment Impact	App updates required for every policy change	Zero app modifications needed

Core Capabilities for Safe, Scalable Agentic AI

01. Agent Inventory & Run-time Visibility

Govern, monitor, and influence the behavior of autonomous AI agents while they are actively running. Maintain a complete, real-time inventory of all agents, their roles, and runtime status.

02. Behavioral Analytics + Risk Scoring

Apply ML-driven behavioral analysis and context-aware scoring to detect anomalies, goal misalignment, and assign trust levels in real time.

03. Policy Enforcement

Block unauthorized behavior—tool misuse, unsafe outputs, or sensitive access—at runtime with precision. Enforce actions like block, redact, degrade, or allow per agent, tool, or data source.

04. Auditability & Transparency

Capture every decision, action, and policy outcome in an immutable log for compliance, incident response, and regulatory reporting.

Policy Architecture for Enterprise Scale

Skyrelis provides a **hierarchical, inheritance-based policy model** that mirrors how enterprises operate. Instead of manually configuring policies for every agent individually, security teams define organizational baselines and layer on context-specific controls that cascade automatically.

This architecture enables "one agent, many policies"—essential for SaaS vendors, multi-tenant platforms, and enterprises with diverse departmental requirements.

Hierarchical Policy Inheritance

Define organization-wide baselines, then layer department-specific rules. Policies cascade automatically to all child entities unless explicitly overridden at a lower level.

Level	Example Policy	Inheritance Behavior
Organization	Company Baseline: Block PII exfiltration, require audit logging	Applies to all descendants
Department	Finance Policy: Restrict external APIs, enhanced data classification	Inherits org baseline, adds dept rules
Team / Project	Trading Desk: Rate limits, human approval for >$1M	Inherits dept + org, adds team rules
Individual Agent	Override: Extended tool access for approved automation	Can override inherited rules as needed

One-Click Compliance Overlays

Pre-built policy templates aligned to major regulatory frameworks. Toggle them on for any scope—organization, department, or individual agent—to instantly apply the required controls without manual configuration.

Why This Matters

- **Operational Efficiency:** Security teams define policies once at the appropriate level; inheritance handles propagation automatically.
- **Consistent Governance:** Baseline controls apply everywhere, eliminating gaps from inconsistent manual configuration.
- **Rapid Compliance:** Toggle on a compliance overlay in seconds instead of building controls from scratch.
- **Flexibility with Control:** Departments can customize policies for their unique needs while still inheriting organizational baselines.

Use Case Examples

Multi-Agent Key Exchange: Enforce secure key handling in agentic workflows by intercepting inter-agent key sharing attempts. Redirect key requests to a designated key management agent integrated with a secure Vault, ensuring secrets and API keys are never exposed or improperly transmitted.

Financial Services: Secure agentic workflows that analyze customer transactions in AWS S3—prevent unauthorized data exfiltration and maintain audit-ready logs for regulatory review.

Legal Review: Control AI-driven contract analysis agents, blocking leakage of privileged clauses and generating compliance reports for attorney-client privilege protection.

Healthcare Data Processing: Govern PHI extraction and anonymization agents to ensure HIPAA-grade security for electronic health records. Apply compliance overlays with one click.

Threat Coverage: OWASP & MITRE Alignment

The Skyrelis Agentic Security Platform maps directly to the **OWASP Top 10 for LLMs**—a widely recognized list of the most critical AI-specific security risks—and **MITRE ATLAS**, a framework cataloging real-world adversarial techniques against machine learning.

Capability	OWASP LLM Risks	MITRE ATLAS Tactics
Prompt/Response Filtering	LLM01: Prompt Injection, LLM02: Sensitive Info Disclosure	Prompt Injection, Jailbreaks, Data Extraction
Anomaly Detection	LLM05: Improper Output, LLM09: Misinformation	LLM-Enabled Influence, Content Manipulation
Runtime Enforcement	LLM04: Model Poisoning, LLM06: Excessive Agency	Over-Privileged Tools, Autonomous Action Abuse
Tools/Data Access Control	LLM06: Excessive Agency, LLM08: Vector Weakness	Model Exfiltration, RAG/Index Poisoning
Observability & Audit	LLM10: Unbounded Consumption	Model/Service DoS, Attack Staging

ROI Calculator: Quantifiable Financial Advantages

The table below demonstrates the financial impact of Skyrelis for a typical enterprise deployment:

Category	Without Skyrelis	With Skyrelis	Savings/Agent
Policy Update Dev Cost	$72,000	$1,200	**$70,800**
Productivity Delay Loss	$25,000	$0	**$25,000**
Developer Productivity Loss	$40,000	$0	**$40,000**
Risk Reduction	$90,000	$45,000	**$45,000**
Total Annual ROI per Agent	—	—	**$180,800**
Total Annual ROI (50 Agents)	—	—	**$9,040,000**

Next Steps

We're currently inviting forward-thinking teams to join our **Skyrelis Innovation Circle**—a curated early access group helping shape the future of agentic AI security.

Members will gain early insights, share real-world challenges, and influence the development of Skyrelis Agentic Security Platform before general availability.

Ready to secure your agentic AI?

Visit www.skyrelis.com to schedule a demo
or contact us at info@skyrelis.com

Note: Some capabilities listed are currently under development and may be subject to change prior to general availability.