

Securing AI Agents by Design



Overview

Agentic AI has opened new opportunities for re-imagining digital transformation initiatives. With Agents that can reason based on provided context and are capable of executing autonomous actions, it has opened new ways in which business processes/workflows can be redefined. HCLTech, in partnership with Google Cloud, has developed a suite of Agents using Agentspace, Agent Development Kit (ADK), Agent Builder and Agent to Agent Protocol (A2A) across industries, common enterprise functions and software engineering.

To enhance Model security and threat prevention, HCLTech and Palo Alto Networks have partnered to launch Prisma AIRS, a secure AI runtime platform designed to protect the agentic ecosystem. This collaborative solution, supported by expert consulting, reference architectures and managed services, enables enterprises to confidently and securely scale their autonomous AI initiatives by design.

Challenges



Expanding attack surface:

Autonomous agents operate across APIs, tools and external systems, increasing exposure to threats like prompt injection, exfiltration and identity spoofing.



Supply-Chain risks:

Dependencies across open-source models and tools create vulnerabilities like poisoned packages or stolen models, making runtime validation and software bills of materials (SBOM) essential.



Over-permissioned agents:

Unrestricted function bindings or misconfigured vector databases may result in agents performing unauthorized actions, such as issuing refunds or leaking proprietary data.



Emerging protocol threats:

With new interaction layers such as Agent-to-Agent (A2A) and Model-Context Protocol (MCP), agents can inadvertently expose prompts, identities and data without strict security protocols.

Our solution

Prisma AIRS integrates seamlessly into agent workflows to deliver defense-in-depth security capabilities:



Model scanning and SBOM

Validate LLM weights, licenses and supply-chain dependencies before deployment.



Posture management

Identify misconfigurations across data stores, embedding models and access controls.



AI Red teaming

Automatically simulate attack scenarios that evolve alongside agent tool graphs.



Runtime LLM and Agent firewall

Intercept prompts, responses and tool calls to prevent policy violations, data leakage, or malicious execution.



Agent security controls

Employ behavioral analytics to stop threats like spoofing, memory poisoning, or uncontrolled chaining.



A2A and MCP governance

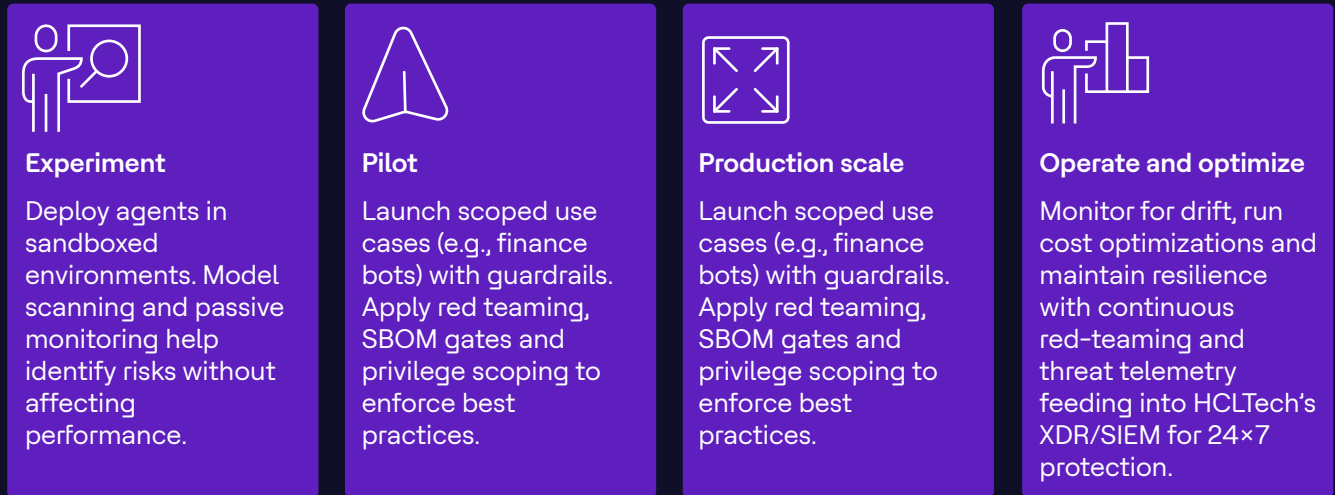
Identity-aware policies validate each A2A handshake, secure shared contexts and ensure auditability.

HCLTech's Prisma AIRS, supported by its consulting, integration and managed detection-and-response (MDR) services, ensures the safety and compliance of enterprise AI agents throughout their development, deployment and operational phases.

Value proposition

Secure-by-Design AI lifecycle

Enterprises can securely implement Agentic AI from experimental phases to full production through four stages.:



Key benefits



Next steps

Please feel free to reach out to ecosystem.marketing@hcltech.com to fast-track your Agentic adoption with enterprise guardrails

- Engage HCLTech Agentic AI COE Squad for a free AI security workshop
- Engage HCLTech Engineers to secure your first 2-3 agents for free (secured with AI Runtime security)
- Deploy a quick Proof of Concept (POC) within your environment - Rapid value demonstration on real-time threat prevention

HCLTech | Supercharging
Progress™

hcltech.com