

AgentMinder

Run Your Business with AI You Can Trust

Thank you for your interest in exploring this content.

Carahsoft is the **Trusted Government IT Solutions Provider®** supporting a broad portfolio of industry-leading technologies through OH MSCA, NASPO ValuePoint, Texas DIR-CPO-5683 and a wide range of other contract vehicles.

As the **Master Government Aggregator®**, Carahsoft connects government agencies, industry partners, and technology providers to deliver innovative, mission-focused solutions.

In partnership with Broadcom, we provide technology solutions that drive modernization, strengthen operations, and ensure compliance with evolving government standards.



To learn more about how Carahsoft can support your technology needs, please visit carahsoft.com



Explore More Resources:
carah.io/BroadcomResources



Join Events & Webinars:
carah.io/BroadcomEvents



Discover Technology Solutions:
carah.io/BroadcomSolutions



Learn About Procurement:
carah.io/BroadcomContracts



Connect With Our Team:
Broadcom@carahsoft.com
(571) 662-3260

Product Brief

At A Glance

Turn autonomous AI agents into governed, production-grade digital employees.

Key Benefits

- Comprehensive security
- Cost optimization
- Flexible convenience
- Breadth of coverage

Key Features

- Agentic fabric
- Action-based governance
- Agentic AI enforcement point
- Delegation and mission constraints
- Unified observability and audit
- Pluggable policy decision point
- Cloud-native architecture

AgentMinder

Run Your Business with AI You Can Trust

Overview

AI has moved beyond the pilot phase and is now an infrastructure priority. According to Foundry's 2026 AI Priorities Study, nearly two-thirds of IT decision-makers report having a dedicated technology budget for AI this year. Within that spending, generative AI—and, increasingly, autonomous AI agents—is driving the next wave of enterprise adoption. This shift is not simply about deploying large language models; it is about operationalizing autonomous agents that can act across systems, APIs, and data. Yet most enterprise security models were not designed for that level of autonomy. As Dario Amodei wrote in *The Adolescence of Technology* (January 2026), "We are past the event horizon; the takeoff has started."

Business Challenges

Agentic AI is an artificial intelligence designed to act and pursue a goal to completion, rather than merely answer a prompt. This represents a fundamental shift in the enterprise, as software moves from acting as a consultant to operating as an autonomous digital employee. An agentic system's core capabilities include autonomy (making independent decisions), reasoning (planning its steps and self-correcting errors), and action (using digital tools to execute tasks like sending emails or manipulating files).

This autonomy creates a significant operational trust gap: agents are non-deterministic, often unpredictable, and lack inherent boundaries. Executives struggle to answer basic questions: Which agent did what, why, and who allowed it? Agentic AI security provides the control plane that turns these high-risk agents into governed, production-grade assets.

Specifically, agentic AI security addresses three urgent security challenges and business risks:

- **Regulatory compliance:** Organizations must prove that AI agents comply with evolving regulations and internal policies, with clear audit trails of who did what, when, and why.
- **Agent behaviors:** As agents gain autonomy, executives worry about unsafe or uncontrolled actions and need reliable guardrails that keep behavior within approved bounds.
- **Interoperability and integration:** Security teams need a consistent way to identify agents, enforce policy over actions, and monitor their activity without slowing innovation.

Together, these pressures are driving demand for dedicated agentic AI security platforms that can make autonomous agents governable, observable, and trustworthy at enterprise scale.

Solution Overview

Broadcom® AgentMinder™ software is an agentic security and governance platform that turns autonomous AI agents into governed, production-grade digital employees:

- **Identity and intent:** AgentMinder treats agents as first-class identities and introduces intent as a control object, so every agent must declare what it is trying to do, not just who it is, before it can touch your systems.
- **Runtime enforcement:** A cloud-native AI gateway enforces least-privileged policies on every tool call, authenticates tokens, and routes requests only to authorized back ends, while a pluggable policy engine evaluates rich context (identity, tool, intent, resource) in real time.
- **Observability and audit:** An observability layer built on OpenTelemetry gives security, risk, and platform teams full, compliance-grade visibility into every agent session and action, delivering chain of custody, anomaly detection, and operational insight at machine speed.

This combination of intent governance, runtime enforcement, and deep auditability lets enterprises safely scale agentic AI into high-risk workflows without rewriting their existing identity or authorization stack.

Key Features

- **Agentic fabric:** Verifiable identity for agents, with explicit mission and intent objects that govern what each agent is allowed to do, and when it must, hand decisions back to humans.
- **Action-based governance:** Policies that evaluate who, what tool, which resource, and why (intent) before permitting any action.
- **Agentic AI enforcement point:** Protocol-aware AI gateway that authenticates tokens, applies per-tool and per-back-end policy, and prevents direct agent access to back ends.
- **Delegation and mission constraints:** Support for agent-to-agent delegation, where child agents inherit and cannot exceed the parent's mission boundaries, so helpers that ask for help never gain more authority than the original agent.
- **Unified observability and audit:** OpenTelemetry traces, metrics, and structured logs for every agent session and tool invocation, with full context and immutable audit trails.
- **Pluggable policy decision point (PDP):** Integration with existing authorization stacks leveraging the AuthZEN standard, so organizations can reuse current policy enforcement endpoints.
- **Cloud-native architecture:** Runs alongside LLMs on-premises, in virtual private clouds, or in public clouds without forcing traffic through a single SaaS chokepoint.

Key Benefits

- **Comprehensive security:** Each agent gets intent-aware authorization, meaning an evaluation of not only what it's asking to do, but also why. Delegation is explicit and auditable, so you always know which human or system authorized an agent's actions, and an end-to-end chain of custody is maintained.
- **Cost optimization:** Runtime guardrails are in place to prevent unauthorized charges and tool abuse before they happen. Resource telemetry tracks token consumption and API calls in real time, so finance and engineering teams get predictable, accountable AI costs instead of surprise invoices.
- **Flexible convenience:** AgentMinder is purpose built and cloud native, so it runs alongside the LLMs teams are already using. This means no platform lock-in, and no requirement to route all traffic through a single SaaS checkpoint.
- **Breadth of coverage:** AgentMinder extends the Broadcom reach into agentic AI governance, so customers aren't buying a point solution from a startup; they're adding a capability from the company that already dominates AI infrastructure end to end.

AgentMinder secures agents with real identities and auditable controls, makes AI costs visible and predictable, and deploys wherever your models already run. No other vendor in this space brings all three benefits together in a single, integrated solution; that is the full agentic fabric AgentMinder delivers.

Summary

AgentMinder is the AI-native control plane for autonomous agents. Our agentic fabric gives every agent a verifiable identity, enforces least-privileged policies on each tool call, and continuously audits behavior for unsafe or noncompliant actions. A cloud-native architecture lets security and governance run wherever your LLM lives, on-premises, in virtual private clouds, or in public clouds, rather than forcing traffic through a single SaaS chokepoint. Security, risk, and AI teams gain unified visibility, rich audit trails, and runtime guardrails that turn high-risk, non-deterministic workflows into governed digital employees, so the business can safely deploy its most sensitive workloads on agentic AI. In short, AgentMinder is how you stop AI quietly running your business in the background, and start running your business with AI you can trust.



For more information, visit our website at: www.broadcom.com

Copyright © 2026 Broadcom. All Rights Reserved. The term "Broadcom" refers to Broadcom Inc. and/or its subsidiaries. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.
AM-PB100 June 9, 2026