



Broadcom Unveils AgentMinder, An Enterprise Solution for AI Agent Governance and Runtime Control

August 31, 2026

New Solution Makes It Possible for Enterprises to Use AI Agents While Ensuring They Follow Company Rules and Security Standards

Broadcom Utilizes AgentMinder to Scale Enterprise Grade Agentic AI

LAS VEGAS, Aug. 31, 2026 (GLOBE NEWSWIRE) -- **VMware Explore 2026** -- Broadcom Inc. (NASDAQ: AVGO), a global technology leader that designs, develops, and supplies semiconductor and infrastructure software solutions, today introduced AgentMinder, a new solution that acts like a traffic controller for autonomous artificial intelligence (AI) agents. As AI agents move from generating content to accessing data, invoking tools, and completing business processes, organizations need more than model guardrails or static permissions. AgentMinder independently verifies agent identity and authorizes each action against its declared mission, intent, context, and current risk before the action reaches an enterprise resource.

"AI has moved from the pilot stage to becoming a critical part of enterprises' daily business operations," said Clayton Donley, vice president and general manager, Identity Management Security Division, Broadcom. "However, existing security models were designed for human users, not for autonomous 'digital employees' that can make decisions and take actions on their own. AgentMinder acts as a traffic controller, verifying exactly what these AI agents are doing, providing guardrails, and tracking their work. It helps businesses stop worrying about AI running unchecked, and start running their operations with confidence."

According to IDC Perspective, *Securing the Agentic Enterprise: Part 2 — Assessing the Impact of Agents on Key Security Domains* "This shift demands unified governance, continuous authorization, and real-time telemetry to ensure every agent action is observable, attributable, and reversible." (1) AgentMinder addresses these requirements by providing a consistent framework for managing AI agents across models, tools, and deployment environments. By establishing policy-based controls, continuous authorization, and full telemetry, AgentMinder helps enterprises scale agentic AI while maintaining security, compliance, and operational trust.

Closing the Operational Trust Gap in Agentic AI

Agentic AI differs from traditional software because it doesn't follow a set of rigid instructions. Instead, it observes, thinks, makes decisions, and takes actions to reach a goal independently. While this makes AI powerful, it also introduces risks: it can accidentally access restricted data, make unauthorized changes or overstep authority. AgentMinder enables enterprises to turn autonomous AI agents into governed digital employees via the following capabilities:

- **Identity and intent:** AgentMinder treats agents as enterprise-grade identities and binds their authority to a declared mission, permitted intents, approved tools, and authorized resources. Every agent must declare what it is trying to do, not just who it is, before it can touch enterprise systems.
- **Runtime enforcement:** AgentMinder's cloud-native AI gateway secures every tool call at runtime, authenticating tokens and directing traffic exclusively to authorized backends. Driven by a dynamic policy engine, the solution instantly evaluates context—from user identity to intent—to help enable total policy compliance
- **Observability and audit:** With an observability layer built on OpenTelemetry, AgentMinder provides 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 enables enterprises to safely scale agentic AI into high-risk workflows such as Finance, HR, or IT without rewriting their existing identity or authorization stack.

With a cloud-native architecture, AgentMinder is deployed alongside existing large language models (LLMs)—whether on-premises, in virtual private clouds (VPCs) or across public cloud environments. By integrating with existing authorization stacks via the AuthZEN standard, organizations can reuse current policy enforcement endpoints without routing traffic through a single software as a service (SaaS) chokepoint.

AgentMinder benefits enterprises by helping secure agents with real identities and auditable controls, making AI costs visible and predictable, and deploying wherever their models run—all in a single, integrated solution.

Broadcom Uses AgentMinder for Agentic Pipeline at Scale

AgentMinder provides a highly flexible deployment model across platforms like VMware vSphere Kubernetes Service (VKS), Google Cloud Platform, and other standards-based, cloud-native Kubernetes platforms. AgentMinder's multi-region, active-active architecture supports continuous uptime, effortlessly supporting peak loads of nearly 36 million customer-related and seven million workforce-related API calls every single day.

"With AgentMinder powering our platform, we've achieved massive global scale paired with zero downtime, even during maintenance and upgrades," said Alan Davidson, CIO, Broadcom. "It has enabled company wide deployment of our AI agentic pipeline by providing 'chain of custody' capabilities between developers and multiple agents/skills. The secure, multi-region redundancy, delivers high-availability, low-latency connections for over 20 million customer identities and 72,000 workforce identities."

AgentMinder is generally available today. For more information, visit <https://www.broadcom.com/agentminder>

About VMware Explore

VMware Explore is the established cloud event for IT professionals to advance their skills and credentials. VMware Explore 2026 will connect

attendees with technology experts and solution architects who are solving real-world challenges. Focused technical sessions, Hands-on Labs, and certification opportunities will enable attendees to apply their learnings and deliver greater value back to their organizations. Attendees will leave with the training, tools, and strategies to build and operate a modern private cloud that is AI-native—running any workload while keeping data local and secure. To learn more about VMware Explore, please visit: <https://www.vmware.com/explore>

About Broadcom

Broadcom Inc. (NASDAQ: AVGO) is a technology leader that designs, develops, and supplies semiconductors and infrastructure software for global organizations' complex, mission-critical needs. Broadcom combines long-term R&D investment with superb execution to deliver the best technology, at scale. Broadcom is a Delaware corporation headquartered in Palo Alto, CA. For more information, visit www.broadcom.com.

1) IDC Perspective: Securing the Agentic Enterprise: Part 2 — Assessing the Impact of Agents on Key Security Domains, #EUR154585326, June 2026

Broadcom, the pulse logo, and Connecting everything are among the trademarks of Broadcom. The term "Broadcom" refers to Broadcom Inc., and/or its subsidiaries. Other trademarks are the property of their respective owners.

Media Contact:

Eloy Ontiveros
Broadcom Global Communications
+1-650-427-6145
eloy.ontiveros@broadcom.com