Securing the Agentic Frontier:

A Defense in Depth Strategy with Cisco SASE

Agentic AI is quickly becoming part of the enterprise operating model. Across the business, employees are beginning to work alongside AI agents that can reason, take action, and interact with enterprise tools at machine speed. In fact, 87% of executives in a recent Cisco survey say Agentic AI has already impacted their strategic priorities.[1]

This shift creates new security and networking demands. Securing agentic workflows requires more than disconnected point solutions. It calls for a unified approach that brings together network intelligence, identity context, and real-time security controls to deliver visibility, policy consistency, and protection across AI-driven interactions.

Cisco helps organizations make this transition with an integrated, AI-ready architecture designed to reduce operational seams and strengthen security across users, devices, applications, and agents.

How Agentic AI Works

Three converging capabilities enable Agentic AI:

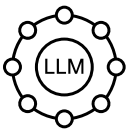

Large Language Model (LLM) Reasoning: Interprets context, plans steps, and generates responses/actions.

Tool Use: Agents can call APIs, applications, and services to create tasks.

Enterprise Integration: Agents can work across sources, SaaS Apps, email, and knowledge stores.

1 Cisco, The Race to Agentic AI: Why Infrastructure Will Make or Break Workforce Transformation, https://www.cisco.com/site/us/en/about/why-cisco/race-to-agentic-ai-report.html

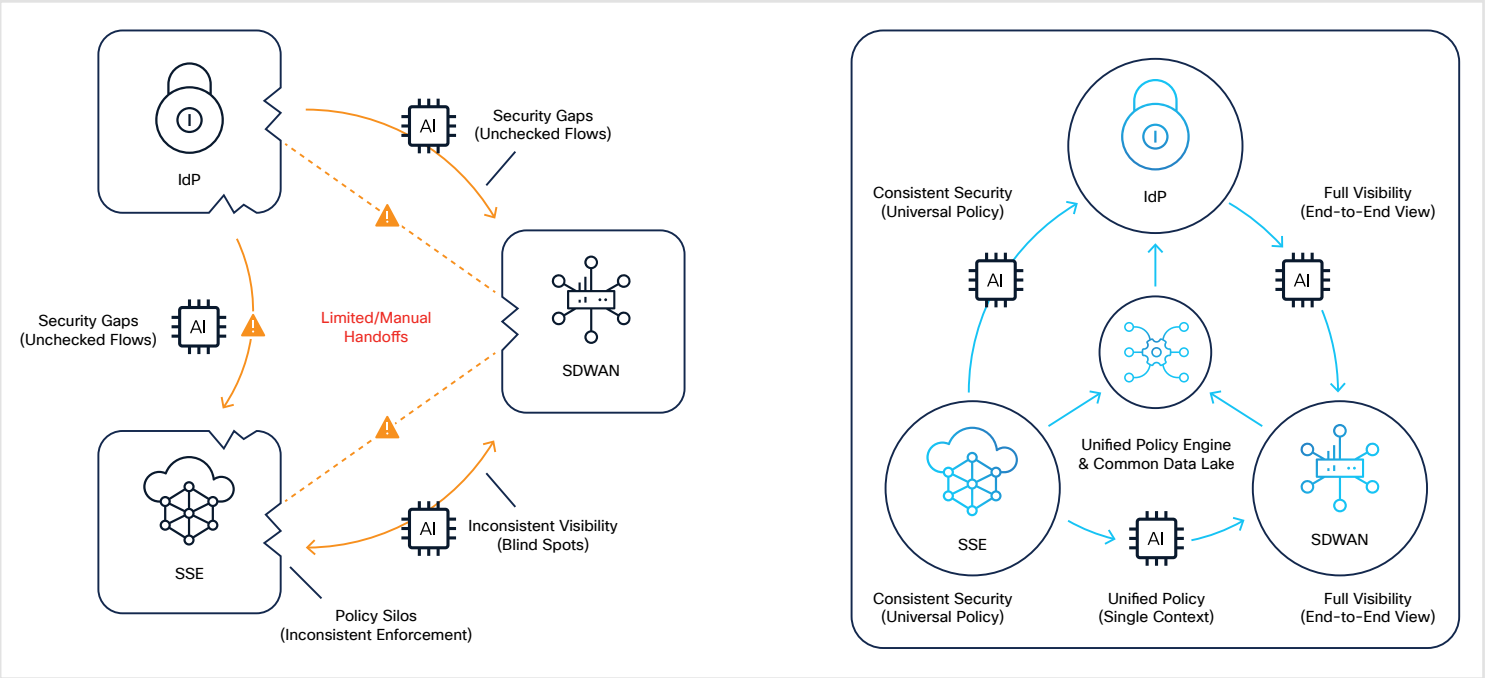

Single vendor SASE provides integrated guardrails for Agentic AI

A New Operational Reality

From experimentation to ubiquity. In less than 18 months, Generative AI (GenAI) and Agentic AI have transitioned from lab curiosities to essential business tools. This isn't just a technological milestone; it's a fundamental shift in how work is executed. Executives estimate that 55% of their workforce will regularly collaborate with AI agents within just 24 months.[2]

However, this shift has introduced intense pressures on the modern network:

1. **Unprecedented Traffic Demands:** Some AI workloads materially increase bandwidth demand and change traffic patterns, which can reveal the limits of traditional QoS policies.
2. **The "Shadow AI" Risk:** Employees are creating agents and using GenAI embedded in SaaS platforms and browser extensions, creating "Shadow AI" footprints that bypass traditional Data Loss Prevention (DLP) and compliance controls.
3. **New Exfiltration Vectors:** AI traffic patterns are fundamentally different. Traditional tools often fail to see the subtle ways data can be exfiltrated through prompt-based interactions.

The time to architect for this reality begins now.

2 Cisco, The Race to Agentic AI: Why Infrastructure Will Make or Break Workforce Transformation, https://www.cisco.com/site/us/en/about/why-cisco/race-to-agentic-ai-report.html

MCP: The Operating System for Agentic AI

To secure AI, you must understand the **Model Context Protocol (MCP)**. Think of MCP as a common protocol that standardizes how AI agents and applications connect to tools and context. It provides a standardized way for AI applications to connect to tools and context, while authentication and authorization must still be enforced by surrounding identity and security controls.

The Security Risk: MCP creates a "network of trust" that is often invisible to traditional security stacks. If an MCP server is compromised or an agent is over-privileged, it becomes a high-speed highway for unauthorized data access and lateral movement.

Here's an example.
A developer agent might use MCP to:

1. Access a GitHub repository (source code)
2. Query a customer database (sensitive PII)
3. Call a payment processing API (financial transaction)

If the MCP server is compromised—or if the agent's access isn't properly governed—it becomes a superhighway for data exfiltration or unauthorized actions.

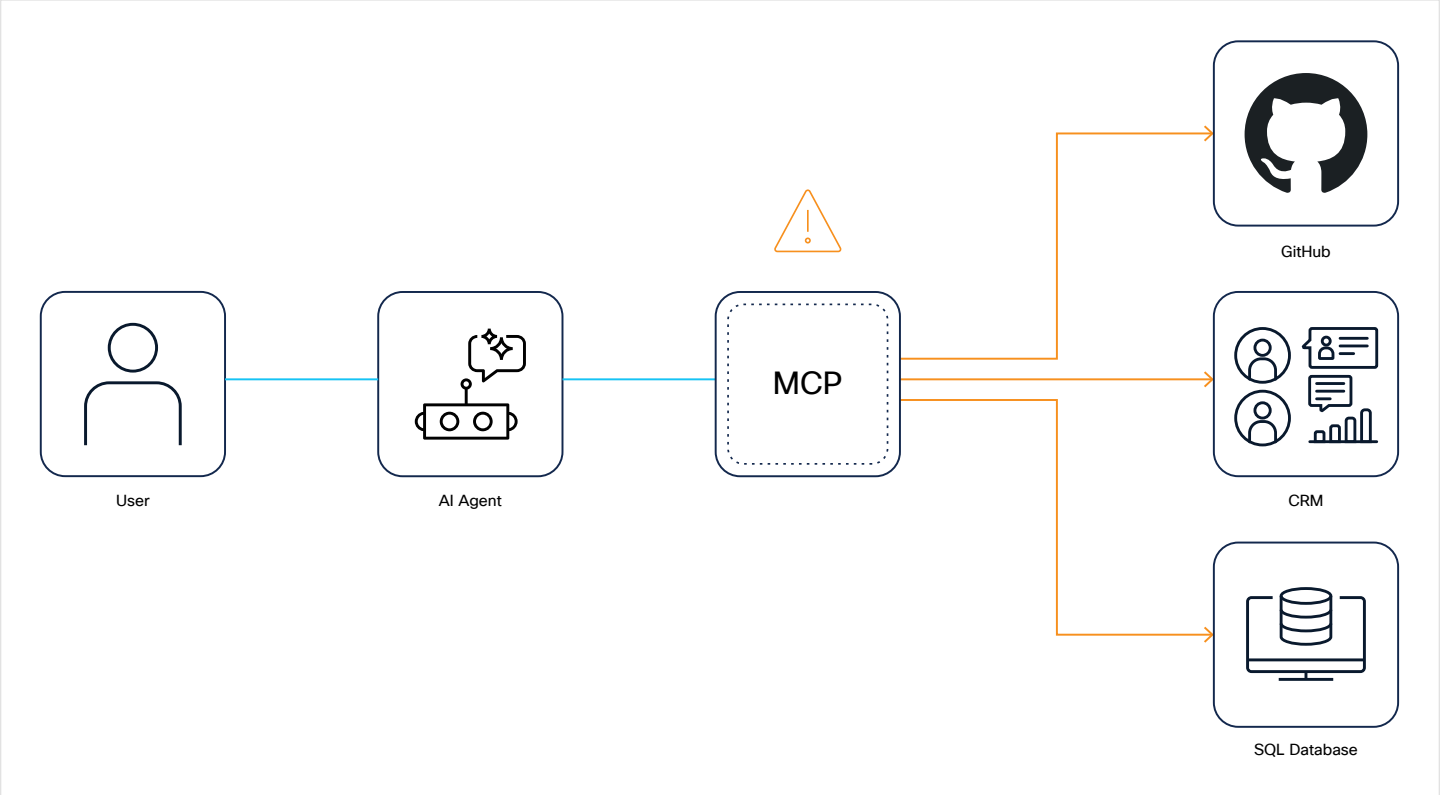

The Four Pillars of AI Defense-in-Depth

AI agents are non-determinstc, meaning their behavior is unpredictable, and operate across hybrid environments at machine speed. Security must be applied across four critical domains:

Know Every Agent in your organization. Visibility and identity with human accountability is required from onboarding, versioning to retirement.

Authorize Every Action with fine-grained authorization. Every API, tool, data or resource interaction is allowed just-in-time with least privilege access enforced.

Adapt to Risk in Real Time because agent behavior evolves rapidly. You need real-time detection for anomalies like “poisoned tools” or prompt injection.

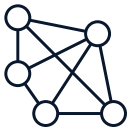

Secure Network Segmentation provides high performance with data insights into AI traffic with fast, low-friction deployment to enforce policy across and within clouds and datacenter.

According to a recent Cisco survey, 62% of executives struggle to protect networks from AI-driven attacks and manage agent identities.[3]

3 Cisco, The Race to Agentic AI: Why Infrastructure Will Make or Break Workforce Transformation, https://www.cisco.com/site/us/en/about/why-cisco/race-to-agentic-ai-report.html

Why Fragmented Security Falls Short

Agentic AI exposes the operational seams between disconnected security and networking tools. Identity systems, access controls, segmentation platforms, and inspection points may each provide value on their own, but when they operate without shared context and coordinated policy, gaps emerge.

For example:

- An identity system may verify who initiated a session, but not how an agent behaves once it begins interacting with tools and data.
- An SSE platform may enforce access decisions, but without deeper integration it may not share the same runtime context as identity, segmentation, and threat detection systems.
- Segmentation and firewall controls may restrict network movement, but they do not always reflect the identity, intent, or behavioral profile of an agent in real time.

The Result: A false sense of security due to fragmented enforcement, policy drift, and operational seams when identity, SSE, SD-WAN and segmentation are separate systems.

The Solution: Cisco SASE

Cisco SASE provides native, built-in integration across all four layers, managed through a single pane of glass: **Security Cloud Control**

- **Know Every Agent:** Cisco extends identity and access controls to non-human entities so organizations can verify and govern autonomous entities alongside human users.
- **Authorize Every Action:** Cisco applies zero trust principles to agent interactions, helping enforce least-privilege access to applications, APIs, and sensitive resources.
- **Adapt to Risk in Real Time:** Cisco provides continuous visibility and behavioral monitoring to help detect deviations from intended operations.
- **Secure Network Segmentation:** Cisco helps discover, prioritize, and segment critical AI traffic so organizations can enforce policy while maintaining performance.

With Cisco, organizations can expect these outcomes:

- More consistent policy enforcement across users, devices, applications, and agents
- Better end-to-end visibility into AI-driven traffic flows and access patterns
- Stronger identity-aware segmentation for hybrid and multi-cloud environments
- A modern zero trust access experience that can reduce reliance on legacy VPN architectures while improving performance and reliability

The Power of the Platform

A unified SASE approach helps organizations move beyond isolated controls and turn the network into a strategic layer of AI security. By reducing operational complexity and aligning policy across domains, organizations can better secure AI-driven workflows while improving resilience and manageability.

- **Operational efficiency:** Integrated visibility, policy, and automation help IT and security teams reduce manual coordination and respond faster to change.
- **Future-Proofing:** A platform-based approach makes it easier to adopt new capabilities over time, including agent-aware policy, advanced segmentation, and emerging cryptographic protections, without constantly stitching together new point products.

Cisco's Integrated AI Security and Safety Framework

Cisco takes a lifecycle-aware approach to AI security. The Integrated AI Security and Safety Framework organizes AI risks across models, agents, pipelines, and the broader ecosystem, aligning protections with industry standards such as NIST Adversarial Machine Learning, MITRE ATLAS, and OWASP guidance for generative AI and agentic applications.

Backed by Cisco Talos Intelligence

Cisco Security is powered by Talos, one of the world's largest commercial threat intelligence organizations. Talos continuously analyzes global threat activity to deliver real-time protection, research-driven insights, and rapid response across Cisco's security portfolio.

Proven at Global Scale

Cisco supports one of the largest installed bases in networking and security, with tens of thousands of enterprise customers worldwide. Our SASE technologies are widely deployed across industries to deliver secure access, resilient performance, and operational reliability at scale.

Embrace the AI Era with Confidence

Agentic AI is reshaping how work gets done and how enterprise traffic moves. Organizations need an architecture that delivers visibility, control, and performance across these new interactions.

Cisco SASE helps organizations secure agentic AI traffic with integrated policy, end-to-end visibility, and identity-aware protection.

Future-proof your infrastructure.
Visit www.cisco.com/go/sase.

P1506171712 04/26