CISCO Engage | Tech Day NYC

Agentic Operations

Marina Ferreira, Principal Solutions Engineer

Chuck Duffy, Distinguished Solutions Engineer

March 4, 2026

Agenda

1. Introduction to Agentic Ops
2. Core Components of Agentic Ops
3. Cisco’s approach to Agentic Ops
4. Implementation Considerations

© 2026 Cisco and/or its affiliates. All rights reserved.

Agentic Ops

Bridging the Gap Between Intent and Execution with Autonomous Digital Workers

Agentic AI

An autonomous system of "digital workers" that reason, execute tools, and self-correct through continuous feedback loops

The Vision

Revolutionizing network management by evolving from manual, script-based automation to autonomous, goal-driven, self-healing networks.

Future Readiness

Which emerging skills are essential for the next-gen Network Engineer?

Current Landscape

Where does the technology stand today?

For example:

- **CLI Fluency → API Fluency**: Understanding RESTful APIs
- **Troubleshooting → Prompt Engineering**: Crafting effective queries for AI agents to diagnose issues
- **Vendor Certifications → LLM Literacy**: Understanding model capabilities, limitations, and context windows
- **Script Writing → Agent Orchestration**: Designing multi-agent workflows using MCP/A2A
- **Human-in-the-Loop Design**: Knowing when to require approval vs autonomous action
- **Security Mindset 2.0**: Understanding prompt injection, tool poisoning, and agentic attack surfaces

© 2026 Cisco and/or its affiliates. All rights reserved.

cisco

Gartner 2030 Forecast

0% of IT work done without AI
75% of IT work augmented by AI
25% of IT work performed by AI alone

© 2026 Cisco and/or its affiliates. All rights reserved.

Evolution of Network Operations

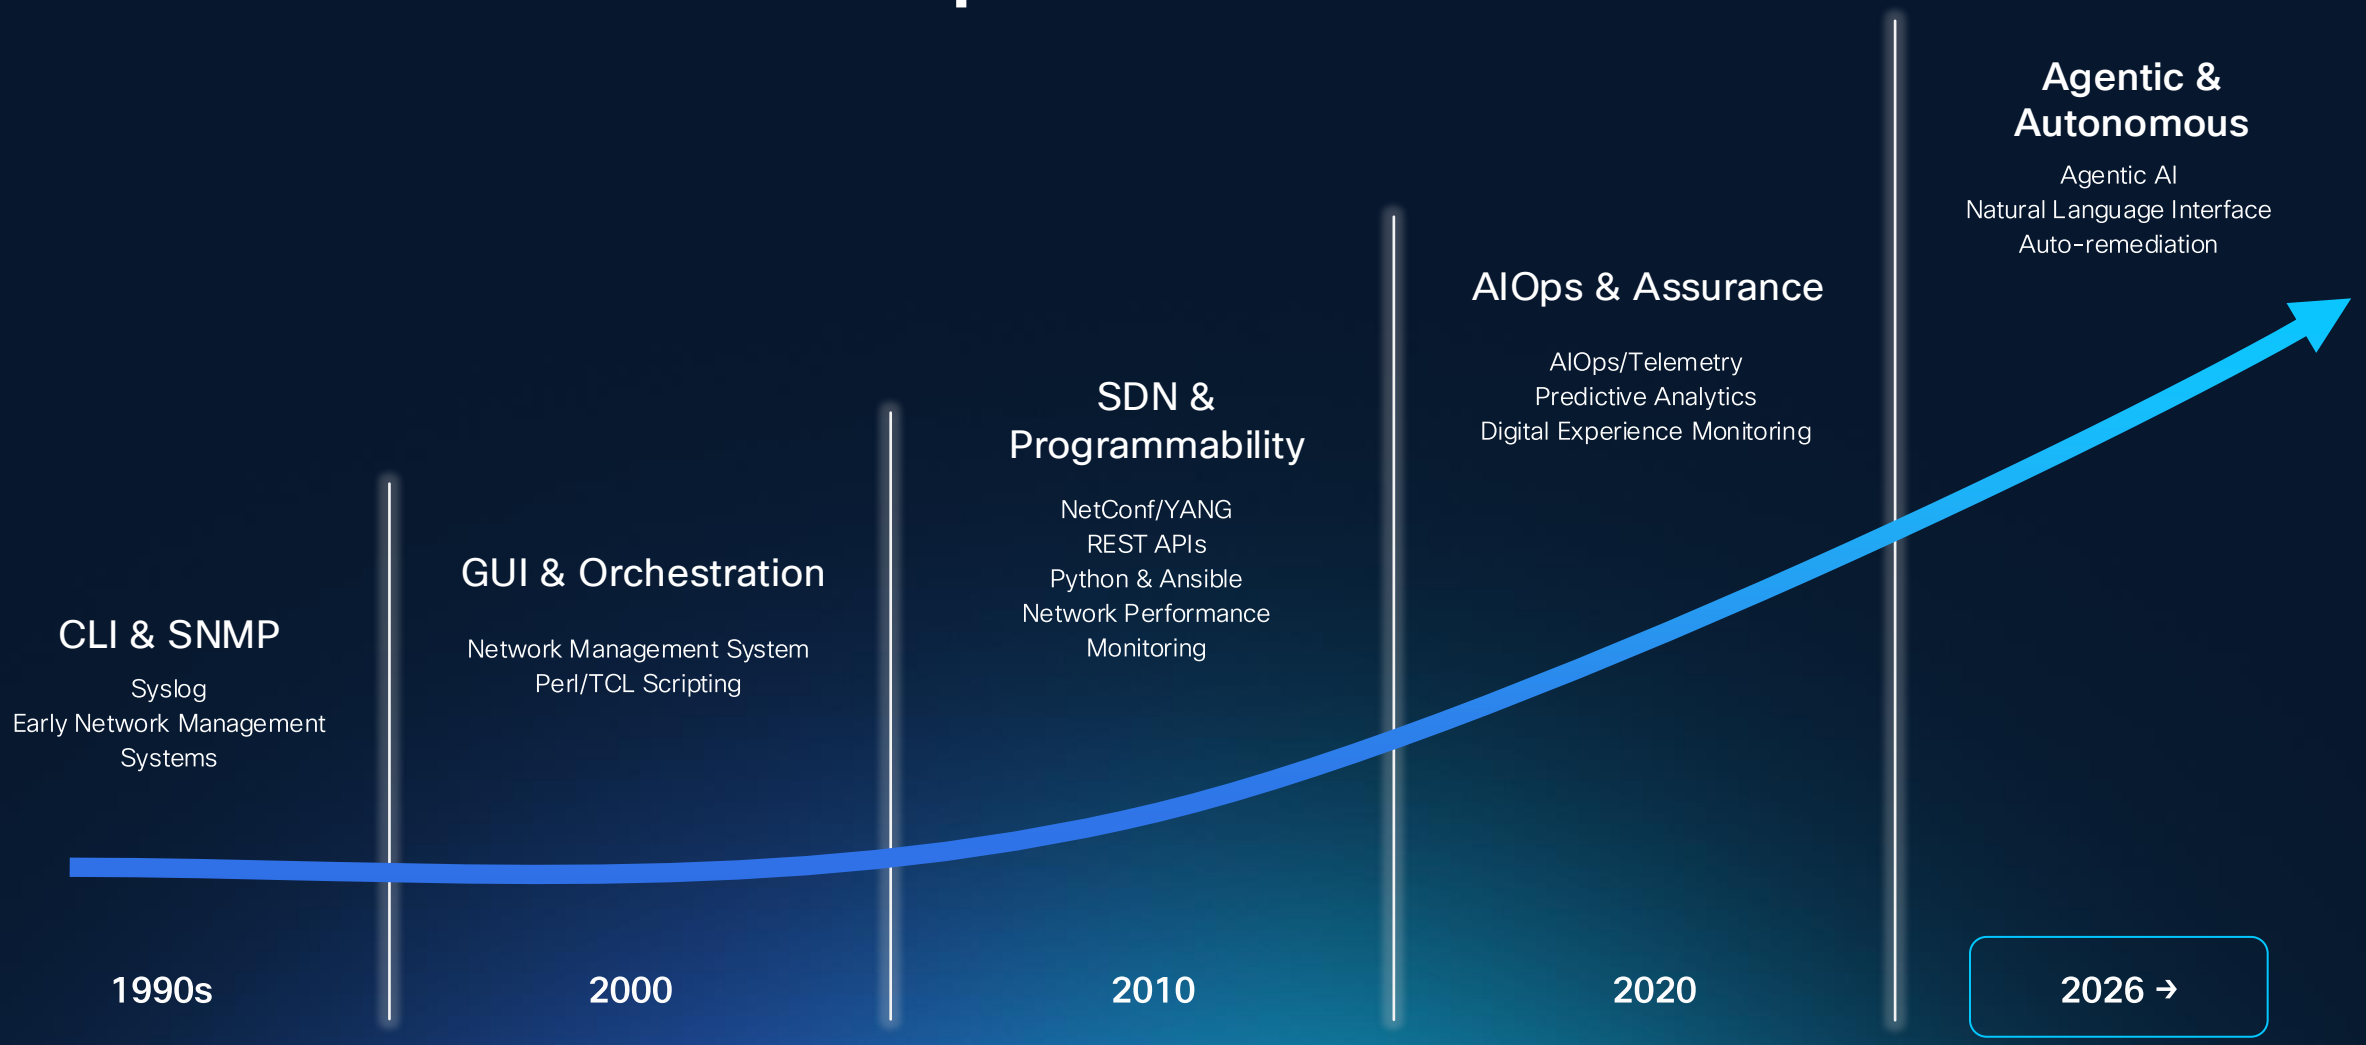

© 2026 Cisco and/or its affiliates. All rights reserved.

Reacting vs Reasoning

From Reactive Correlation to Autonomous Resolution

Traditional AIOps - The "Smart Alarm"

- **Advanced Pattern Matching:** Correlates noise into actionable alerts
- **"What happened?":** Identifies that 50 alerts equal one high-CPU event
- **Human-Dependent:** Delivers the *what* but requires a human for the *how*

Agentic AI - The "Digital Engineer"

- **Active Reasoning:** Transition from simple detection to deep logical analysis
- **"Why did it happen?":** Autonomously investigates root causes and probes systems like an expert engineer
- **Outcome-Oriented:** Delivers a validated remediation plan and fix recommendations

The Shift: From static playbooks to **adaptive intelligence.** Agentic AI leverages Large Language Models (LLMs) and domain expertise to solve novel, "unknown" problems on the fly, no pre-written scripts required.

© 2026 Cisco and/or its affiliates. All rights reserved.

cisco

Agenda

1. Introduction to Agentic Ops
2. Core Components of Agentic Ops
3. Cisco’s approach to Agentic Ops
4. Implementation Considerations

© 2026 Cisco and/or its affiliates. All rights reserved.

Agentic Ops: The Evolution from Knowledge to Action

Moving from a passive knowledge engine to an autonomous digital worker

Agentic Ops is an autonomous system that leverages reasoning and tool-use to achieve complex, multi-step goals

Key Capabilities:

Reasoning & Decision Making

Analyzes context and weighs options using **Chain-of-Thought** to determine the optimal path

Planning & Decomposition

Breaks high-level, abstract goals into a sequence of tactical, executable steps

Tool Use

Interfaces with the real world, executing APIs, querying databases, and running code

Perception & Observation

Actively monitors environment state and system feedback to inform the next move

Memory & Context Management

Retains short-term task state and long-term preferences to ensure workflow continuity

Self-Reflection & Correction

Evaluates its own output; if an action fails, it autonomously pivots to a new approach

© 2026 Cisco and/or its affiliates. All rights reserved.

cisco

AI Agents

Turning Intent into Autonomous Action

AI Agents leverage LLM-driven reasoning to transform environmental data into autonomous, goal-oriented actions.

Key Capabilities:

Autonomy

Translates high-level objectives into independent workflows without step-by-step instructions

Perception

Ingests and interprets real-time signals from network traffic, emails, and web data

Reasoning and planning

Leverages LLMs to decompose complex problems into a sequence of actionable, logical steps

Action - tool use

Executes tasks via MCP, whether configuring a network route, updating a database, or triggering alerts

© 2026 Cisco and/or its affiliates. All rights reserved.

Agentic Architectures

From external reasoning loops to collaborative ecosystems and native inference

Single Agent: Individual Reasoning

Definition: *External orchestration logic managing a centralized reasoning cycle.*

- **ReAct (Reason + Act):** An iterative cycle of **Thought**, **Action**, and **Observation** until task completion
- **Plan-and-Execute:** Decouples logic from action; generates a full sequential roadmap before initiating execution
- **Reflection:** Incorporates an autonomous self-critique phase to evaluate and refine output quality through iteration.

Multi-Agent: Collaborative Intelligence

Definition: *Ecosystems of specialized agents interacting to solve multi-domain problems.*

- **Hierarchical Orchestration:** A "Manager" agent decomposes complex goals and delegates sub-tasks to specialized "Workers"
- **Joint Collaboration:** A peer-to-peer architecture where agents share state and data via direct message passing
- **Standardized Messaging (A2A):** Utilizes open protocols for cross-platform discovery and secure service exchange.

Deep Agents: Native Reasoning

Definition: *Internalized reasoning capabilities baked directly into the model's inference process*

- **Native Chain-of-Thought:** Internalized reasoning process where the model "thinks" before generating any external output
- **System 2 Thinking:** Deliberate, "slow" reasoning designed for high-complexity logic, math, and error reduction
- **Internalized Planning:** The model plans and self-corrects its path during token generation, before the first tool-call.

© 2026 Cisco and/or its affiliates. All rights reserved.

cisco

Agentic Architectures

Single Agent

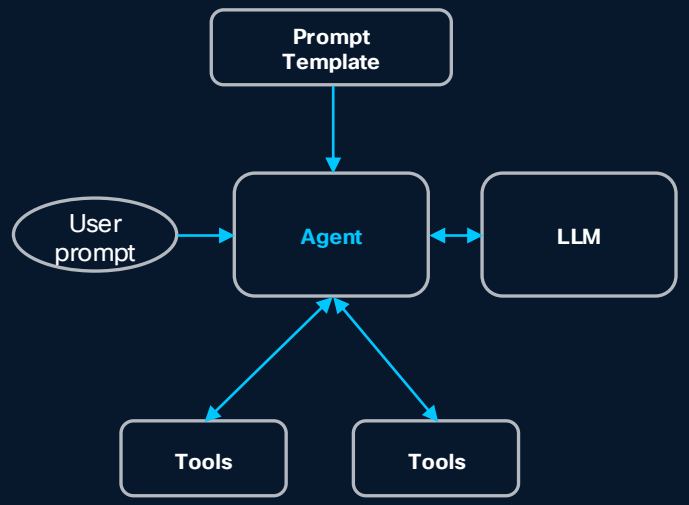

Multi-Agent

Skills

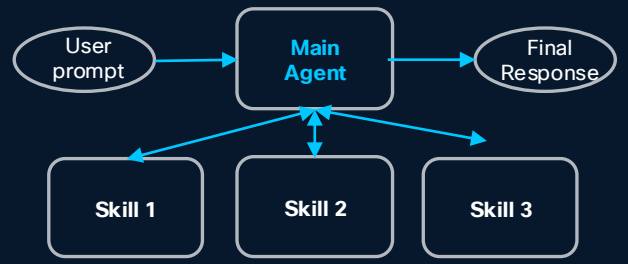

Peer to Peer

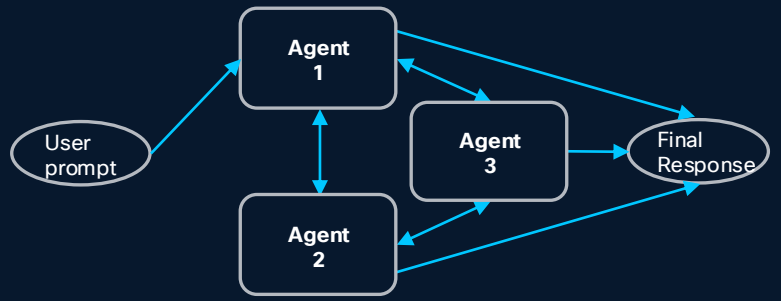

Deep Agents

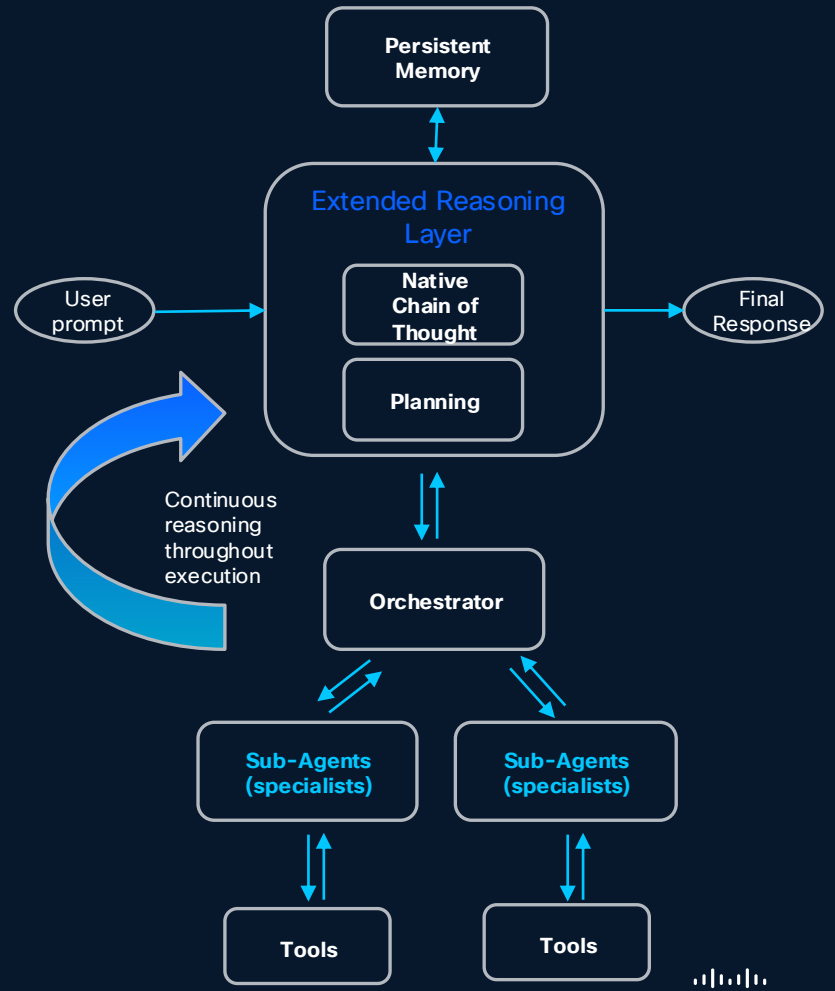

© 2026 Cisco and/or its affiliates. All rights reserved.

Model Context Protocol (MCP)

The open standard for connecting AI applications to external data and tools

Ecosystem & Governance

- **Open Standard:** Developed by Anthropic (Nov 2024); now governed by the Agentic AI Foundation (Linux Foundation)
- **Architectural Evolution:** Moves beyond brittle, manual "Function Calling" to a standardized, interoperable protocol
- **Rapid Adoption:** Over 17,000 active MCP servers (MCP.so), creating a massive, plug-and-play ecosystem

Solving the "M x N" Complexity

- **The Legacy Problem:** AI applications previously required bespoke, one-off integrations for every unique model-to-tool pairing
- **The MCP Solution:** A universal interface that allows any AI client to connect to any data source or tool via a single protocol
- **The Impact:** Drastically reduces integration overhead and enables seamless cross-platform tool discovery

© 2026 Cisco and/or its affiliates. All rights reserved.

cisco

Model Context Protocol (MCP)

An open protocol designed to securely connect AI agents to tools, data, and enterprise systems.
MCP provides a standardized way for agents to discover, invoke, and govern external capabilities.

Core capabilities: (1) Decouple agents from tool/resources (2) Standardize context (3) Scalability

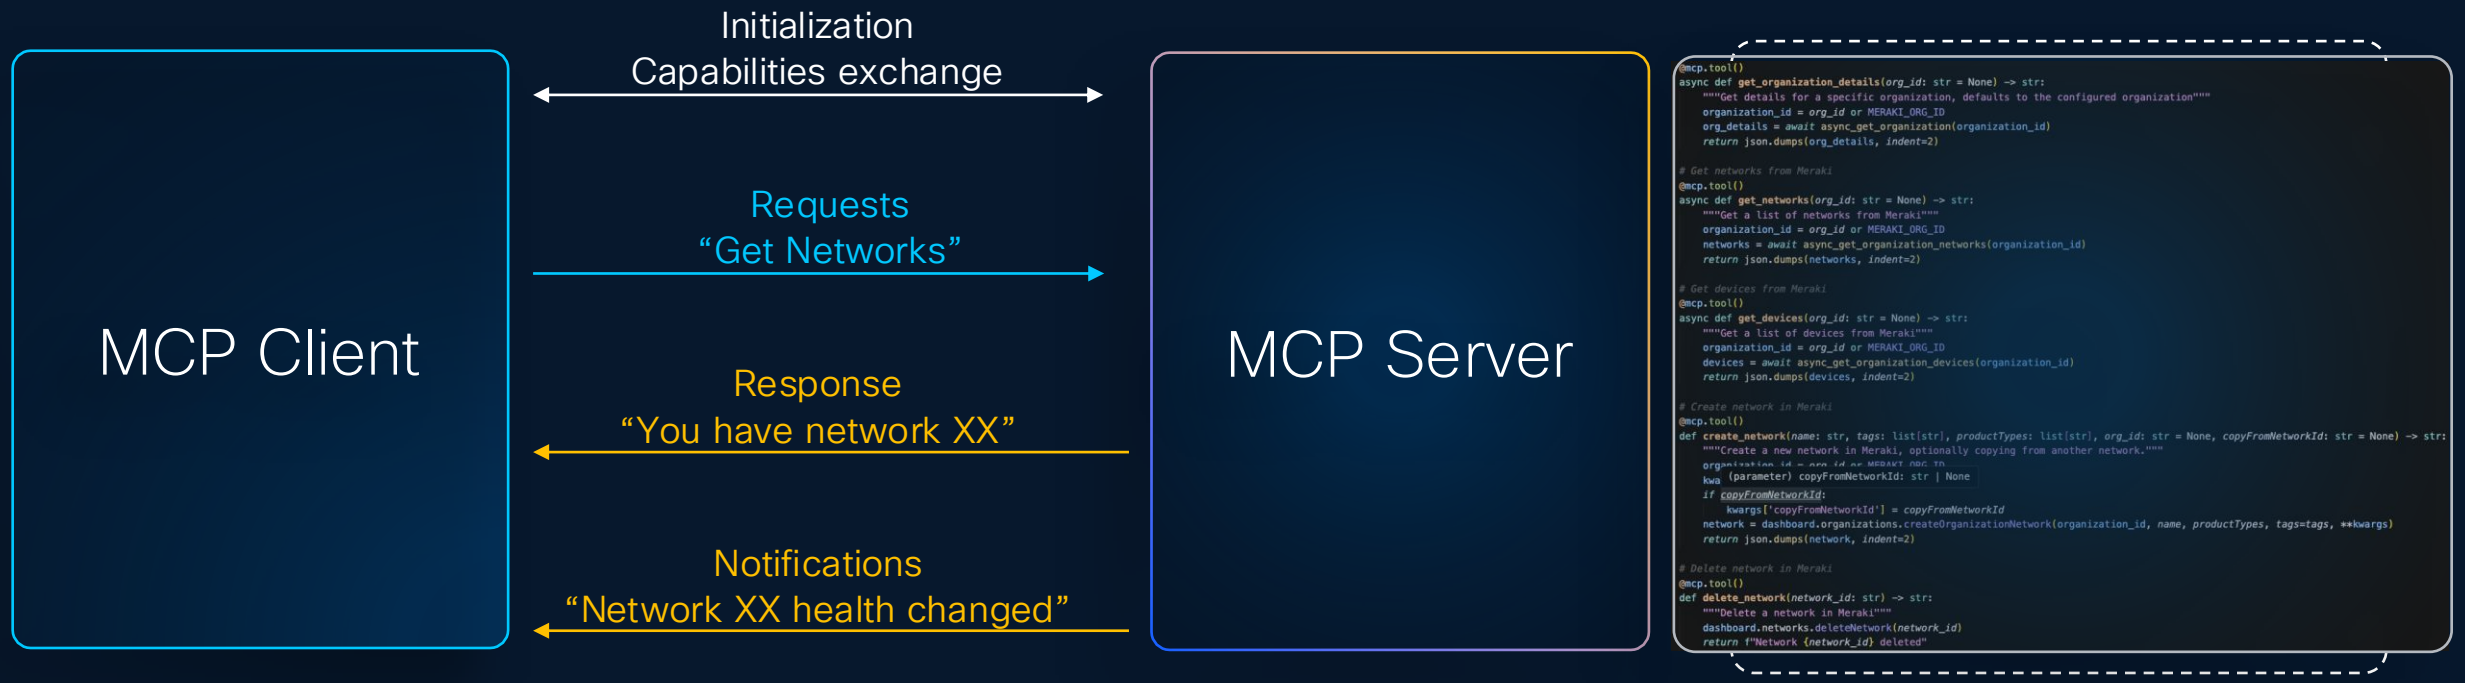

© 2026 Cisco and/or its affiliates. All rights reserved.

Model Context Protocol (MCP)

Example

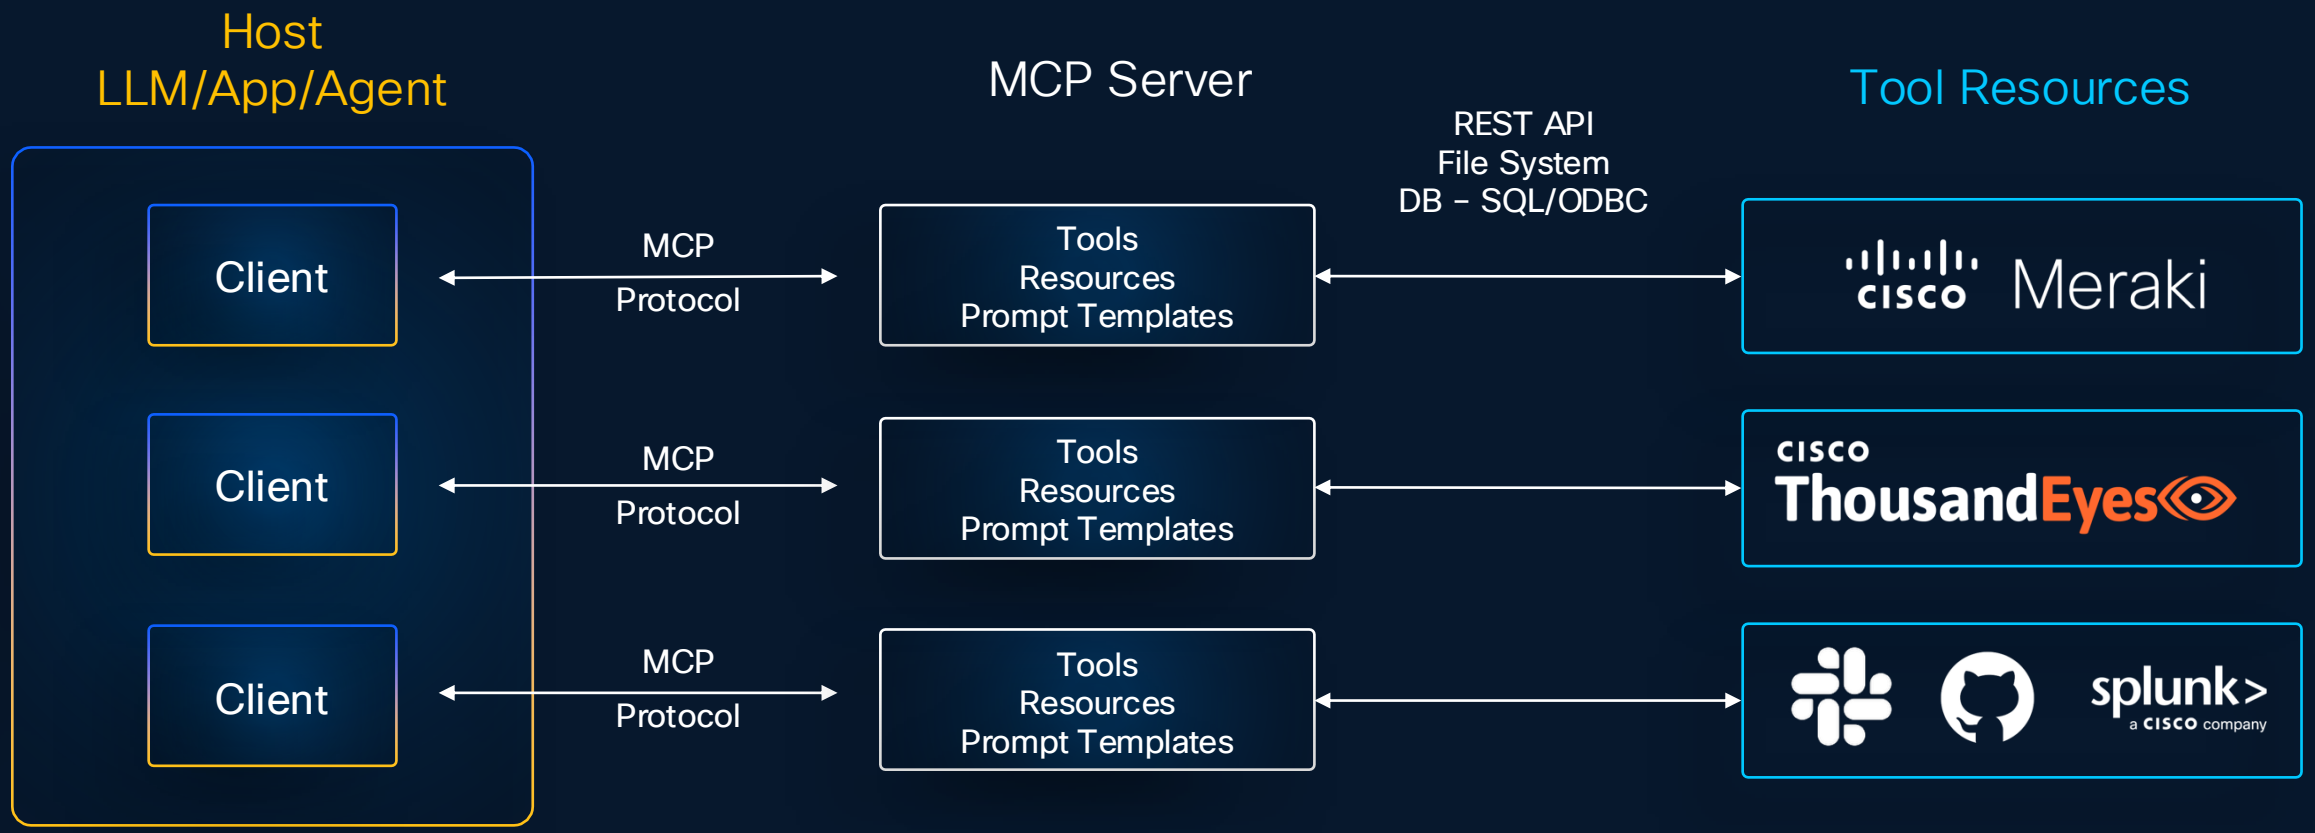

© 2026 Cisco and/or its affiliates. All rights reserved.

Agent 2 Agent Protocol (A2A)

The open standard for cross-platform agent discovery and collaboration

Governance & Timeline

- **Origin:** Created by **Google** and 50+ industry partners in **April 2025**
- **Governance:** Donated to the **Linux Foundation** in **June 2025** to ensure a vendor-neutral, open-source future
- **The Mission:** Eliminating "Agent Silos" by enabling secure, standardized communication between disparate AI platforms (e.g., Cisco, Google, Microsoft)

Core Strategic Principles

- **Vendor Neutrality:** A universal language enabling multi-vendor agent collaboration on shared goals
- **Opaque Collaboration:** Enables agents to share state and results without exposing proprietary logic or internal training data
- **Standardized Trust:** Built-in governance for task delegation, including constraints, data retention, and human-in-the-loop triggers.

© 2026 Cisco and/or its affiliates. All rights reserved.

cisco

Agent to Agent (A2A)

An open standard designed to enable seamless communication and collaboration between AI agents. A2A provides a standardized way for agents to talk to other agents

Core capabilities: (1) Capability discovery (2) Collaboration (3) Task management (4) User experience negotiation

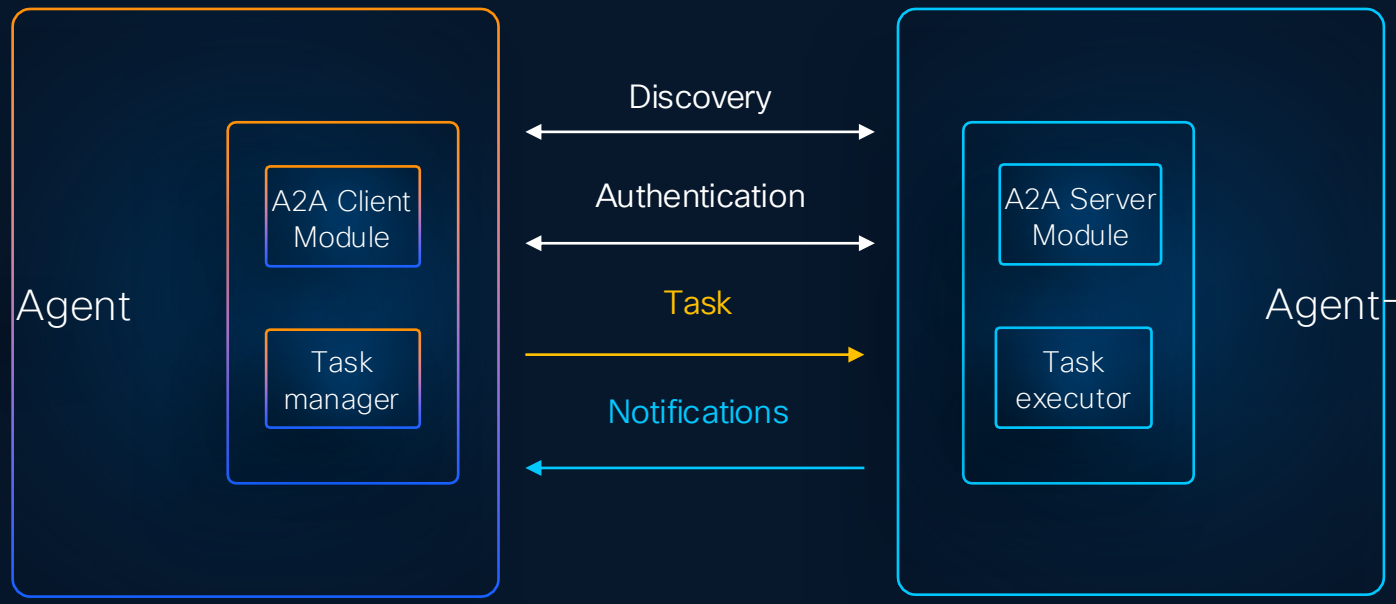


```
{
  "$schema": "https://agent-protocol.ai/v1/agentcard.schema.json",
  "agent_identity": {
    "id": "PA-MAYA-44-BETA",
    "name": "Maya Executive Assistant",
    "version": "1.2.0",
    "type": "Concierge-LLM"
  },
  "capabilities": [
    {
      "skill": "calendar_management",
      "functions": ["check_availability", "create_event", "resolve_conflict"],
      "permissions": ["calendar.read", "calendar.write"]
    },
    {
      "skill": "reservation_booking",
      "types": ["flight", "hotel", "restaurant", "ride_share"],
      "payment_integration": "Stripe-Connect-Vault"
    }
  ],
  "interface": {
    "type": "REST_API",
    "auth_type": "OAuth2",
    "supported_languages": ["en-US", "es-ES", "fr-FR"]
  },
  "constraints": {
    "spending_limit_per_transaction": 500.00,
    "require_approval_for": ["international_travel", "transactions > $500"],
    "data_retention_policy": "encrypted_ephemeral"
  }
}
```

© 2026 Cisco and/or its affiliates. All rights reserved.

Together: MCP & A2A

A2A and MCP are open standards that let AI agents interoperate end-to-end, A2A for agent-to-agent coordination and MCP for consistent, secure access to tools and data. Together they reduce bespoke integrations and make multi-agent systems easier to connect, scale, and govern.

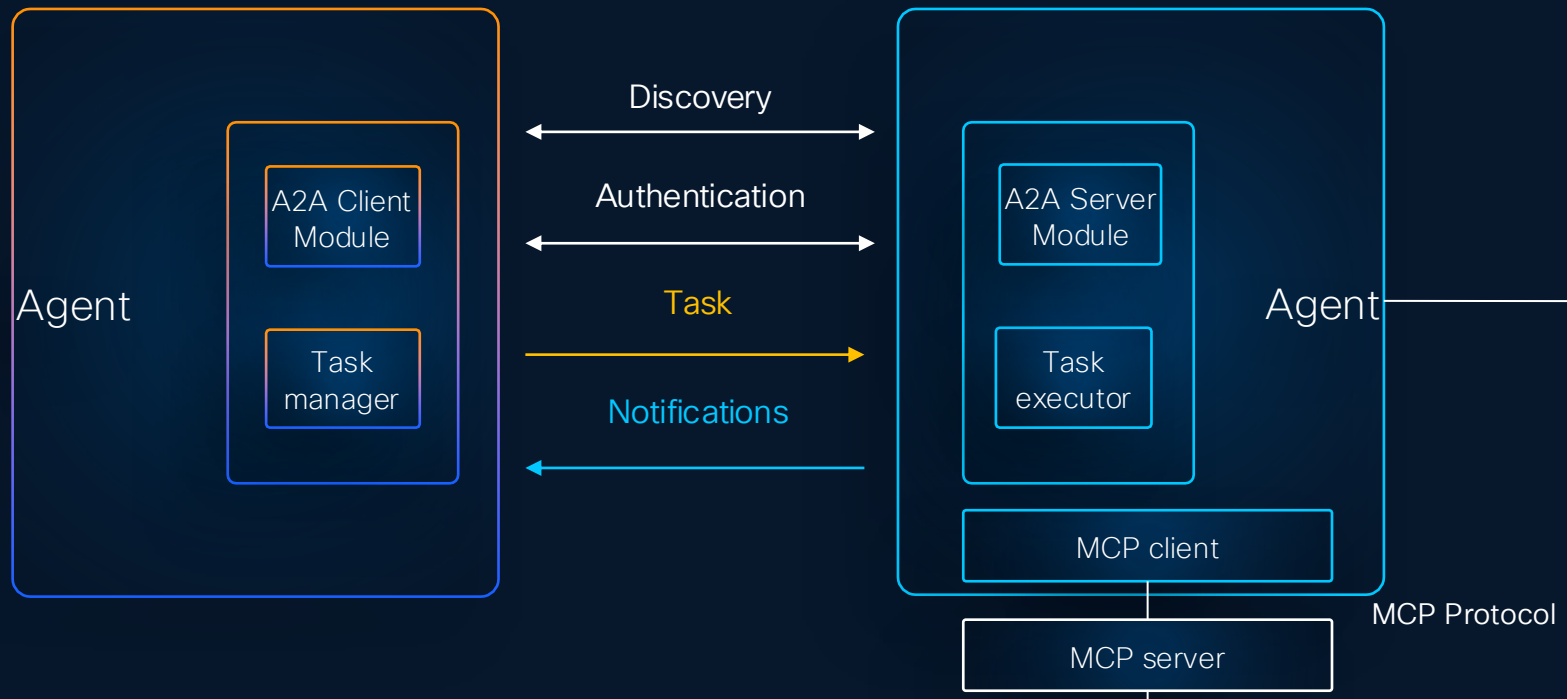


```
{
  "$schema": "https://agent-protocol.ai/v1/agentcard.schema.json",
  "agent_identity": {
    "id": "PA-MAYA-44-BETA",
    "name": "Maya Executive Assistant",
    "version": "1.2.0",
    "type": "Concierge-LLM"
  },
  "capabilities": [
    {
      "skill": "calendar_management",
      "functions": ["check_availability", "create_event", "resolve_conflict"],
      "permissions": ["calendar.read", "calendar.write"]
    },
    {
      "skill": "reservation_booking",
      "types": ["flight", "hotel", "restaurant", "ride_share"],
      "payment_integration": "Stripe-Connect-Vault"
    }
  ],
  "interface": {
    "type": "REST_API",
    "auth_type": "OAuth2",
    "supported_languages": ["en-US", "es-ES", "fr-FR"]
  },
  "constraints": {
    "spending_limit_per_transaction": 500.00,
    "require_approval_for": ["international_travel", "transactions > $500"],
    "data_retention_policy": "encrypted_ephemeral"
  }
}
```


© 2026 Cisco and/or its affiliates. All rights reserved.

cisco ThousandEyes

cisco

Agenda

1. Introduction to Agentic Ops
2. Core Components of Agentic Ops
3. Cisco’s approach to Agentic Ops
4. Implementation Considerations

© 2026 Cisco and/or its affiliates. All rights reserved.

AI in Network Operations Today – AI RRM

AI powered wireless optimization

AI-RRM now in use by more than 5K customers and over 650K access points.

“Two years of escalations ended within six hours of enabling AI-RRM!”

- Network Admin

“AI RRM’s intelligent radio tuning has reduced more than 95% of our configuration updates in our multiple sites.”

- Chief Network Admin

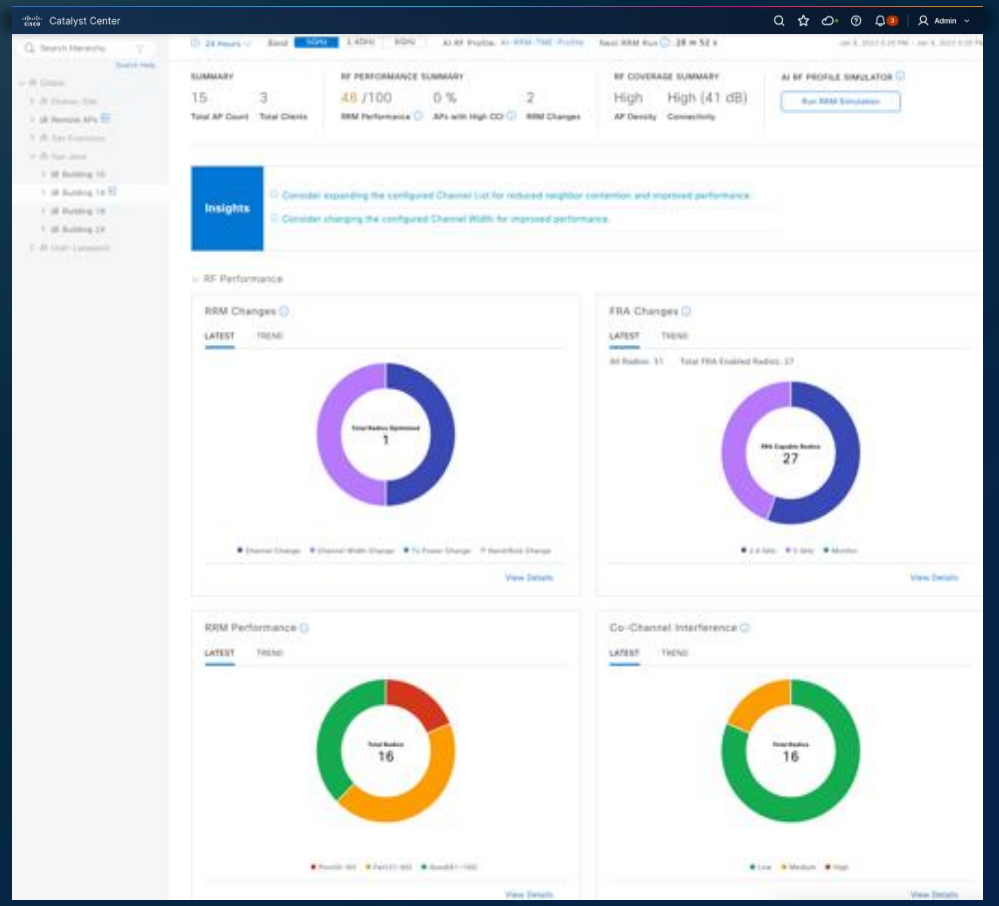

Catalyst Center

© 2026 Cisco and/or its affiliates. All rights reserved.

Intelligent Packet Capture

AI powered insights and troubleshooting

Accelerating issue resolution with AI

- 500 Million+ Proactive PCAPs Performed
- Reduces manual effort to reproduce and diagnose
- Embedded PCAP viewer to streamline troubleshooting
- AI powered analysis delivers clear actionable insights

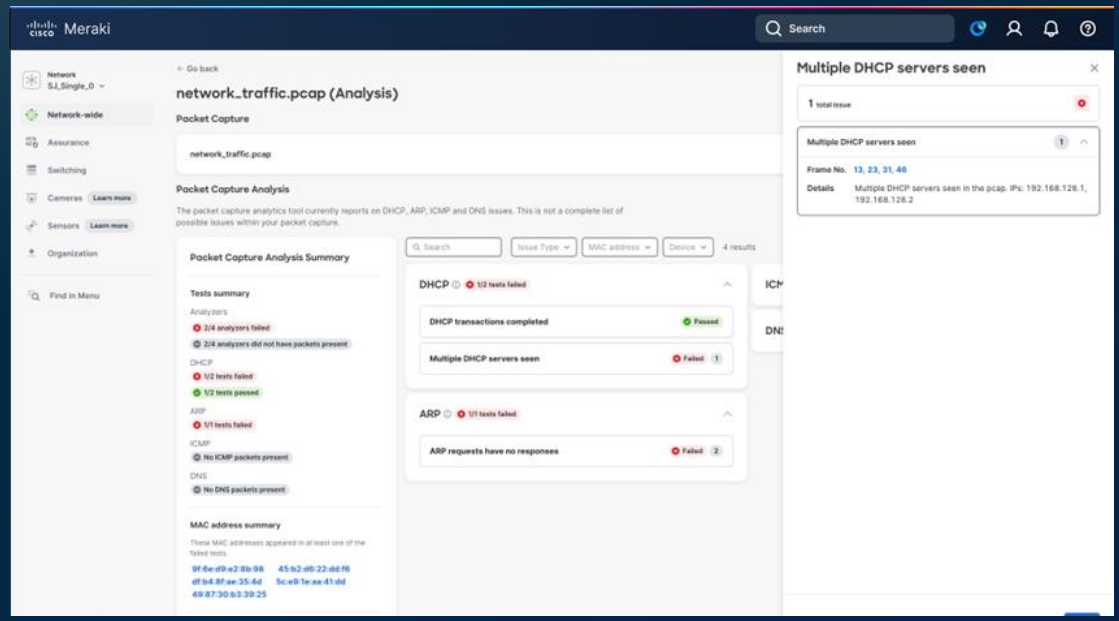

Meraki Dashboard

© 2026 Cisco and/or its affiliates. All rights reserved.

Workflows in AI Assistant

AI-driven network automation

Low code/no code free automation tool integrated within the Meraki Dashboard

165,000+ workflows executed in the last 30 days

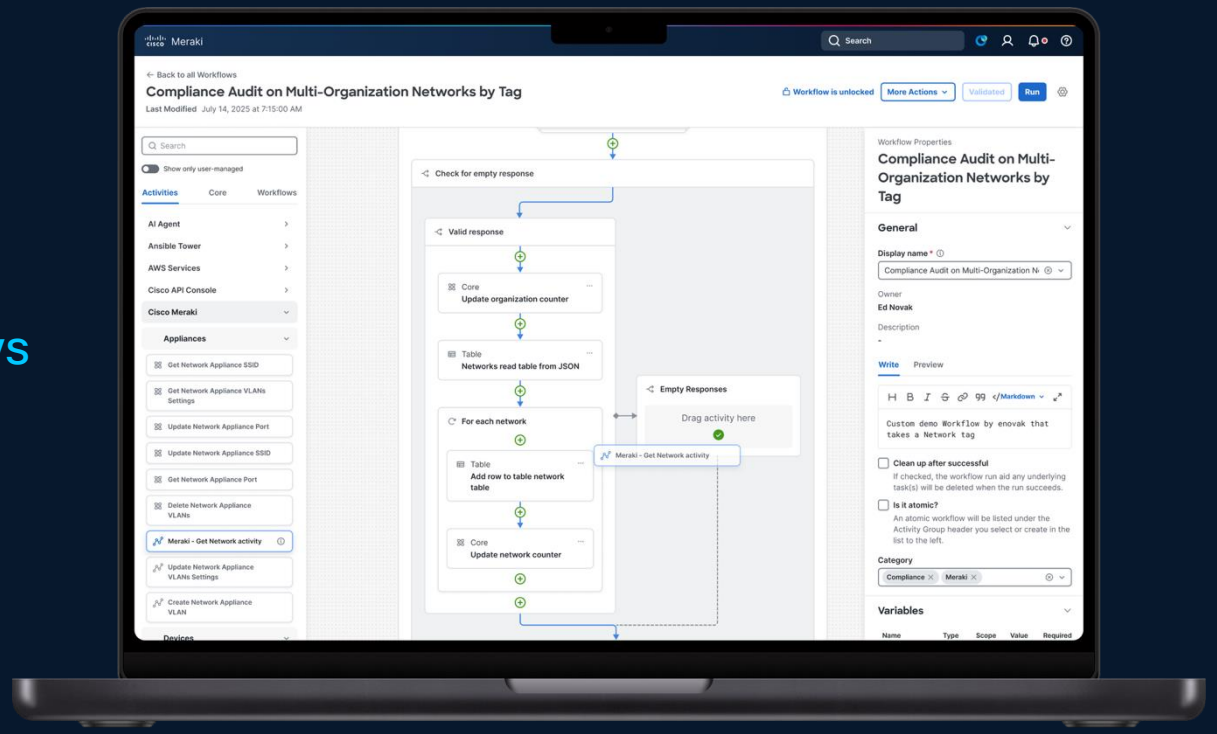

Cloud-Hosted Automation | Drag and Drop Creation / Editing | Custom and Pre-Built | For Cisco & Third-Party Domains

© 2026 Cisco and/or its affiliates. All rights reserved.

Workflows in AI Assistant

Leverage built-in AI Assistant integration for agentic operations using natural language

Natural Language Commands: Request network changes with simple, conversational commands

---

Deterministic Outcomes: Reliable, predictable execution ensures consistent results every time

---

Limitless Flexibility: Easily extend capabilities to automate any network task

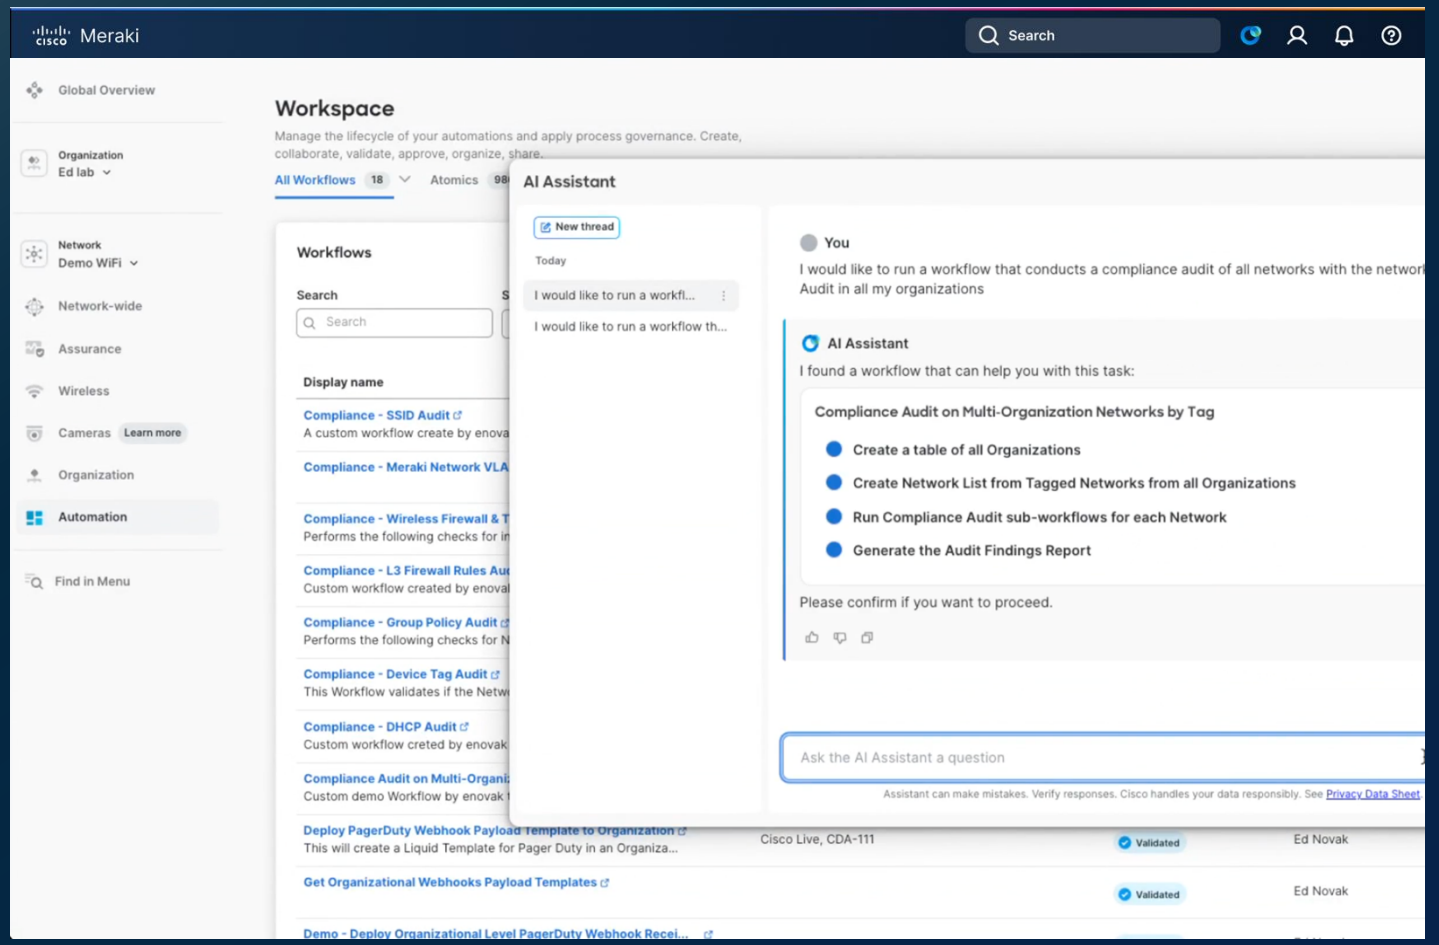

Meraki Dashboard

© 2026 Cisco and/or its affiliates. All rights reserved.

Demo

Workflows in AI Assistant

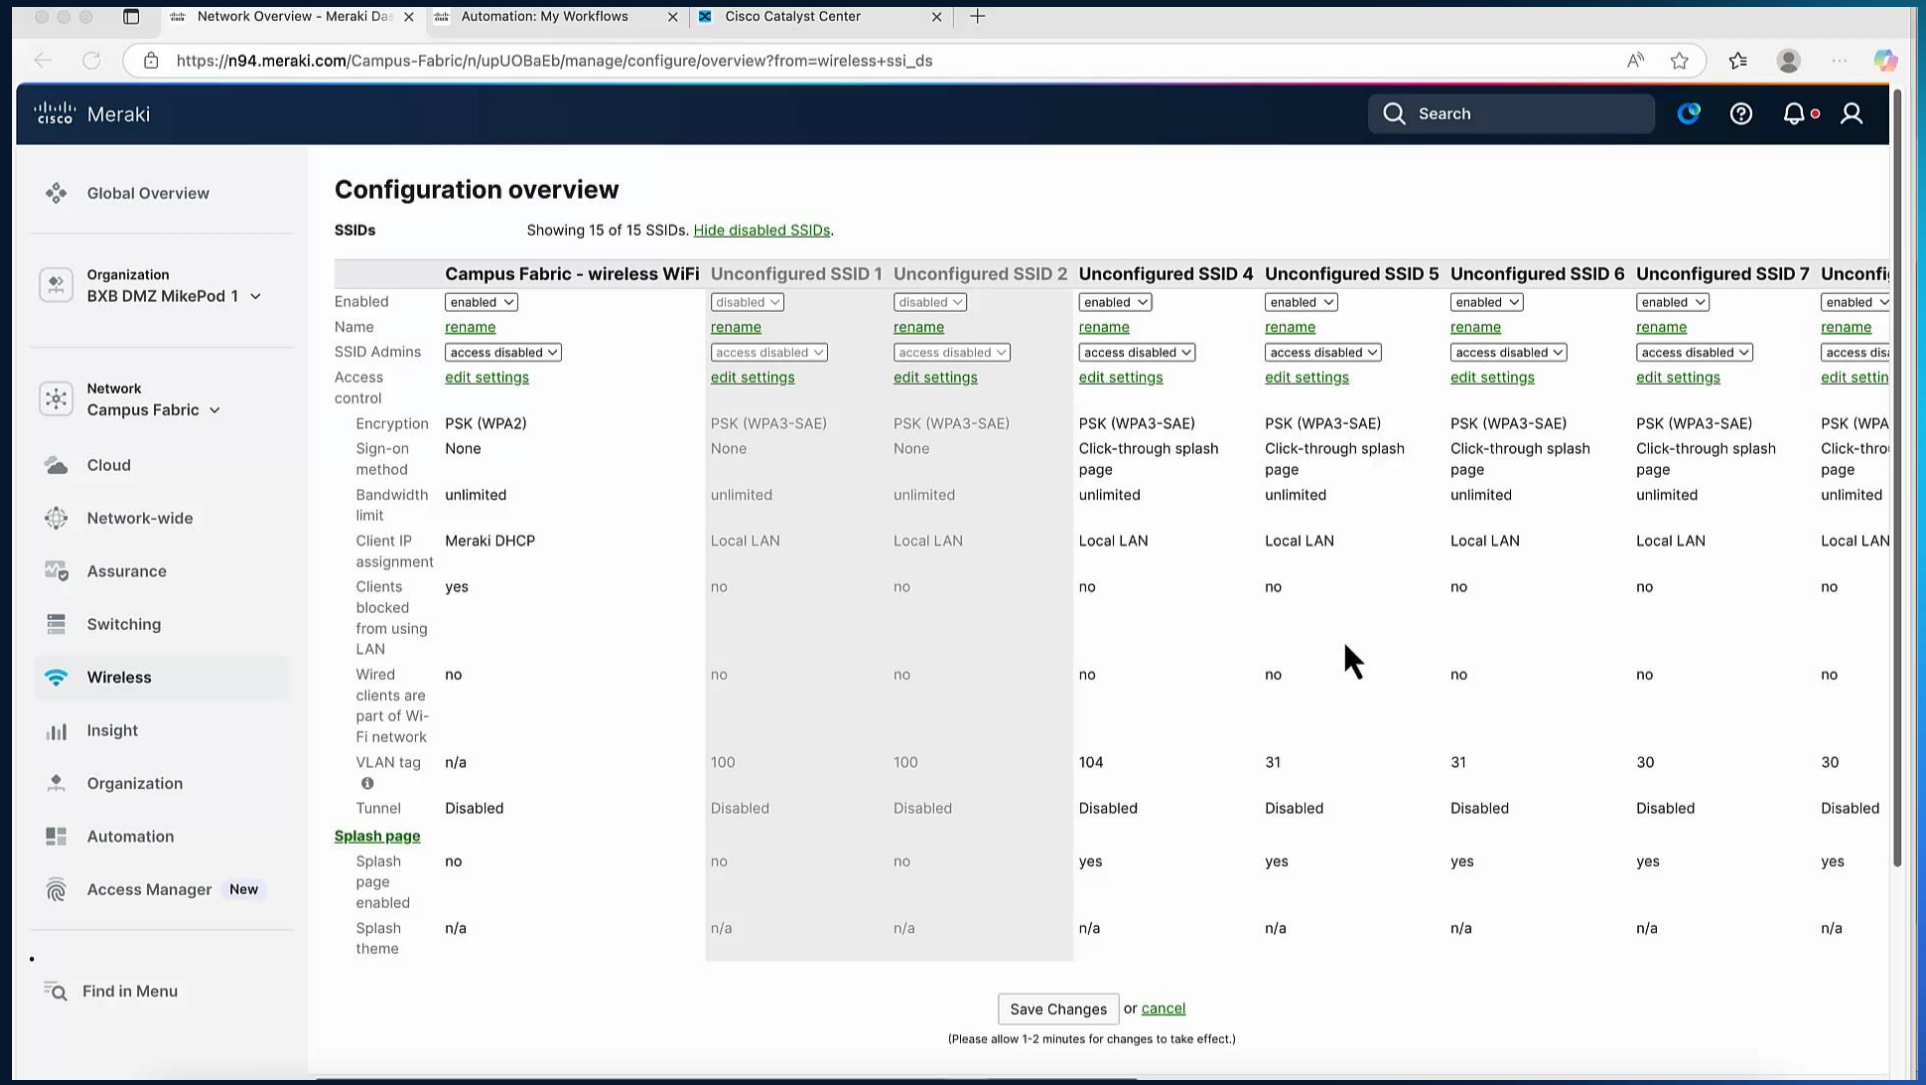

Click on the video in presenter mode to enlarge

Individual AI Assistants Are Integrated Across Cisco

Security	Firewall, Secure Access, Hypershield, Duo, Identity Intelligence Splunk Enterprise Security, ISE
Networking	Meraki, Catalyst Center, Catalyst SD-WAN, ThousandEyes, Intersight
Observability	Splunk Observability (Cloud, ITSI, AppDynamics)
Data	Splunk Platform
Collaboration	Webex Control Hub
Service Ops	Customer Experience

© 2026 Cisco and/or its affiliates. All rights reserved.

cisco

Introducing Cisco AI Canvas

- Single canvas for cross domain troubleshooting
- Generative UI with reasoning built-in
- Keeps NetOps, SecOps, IT and execs on the same page

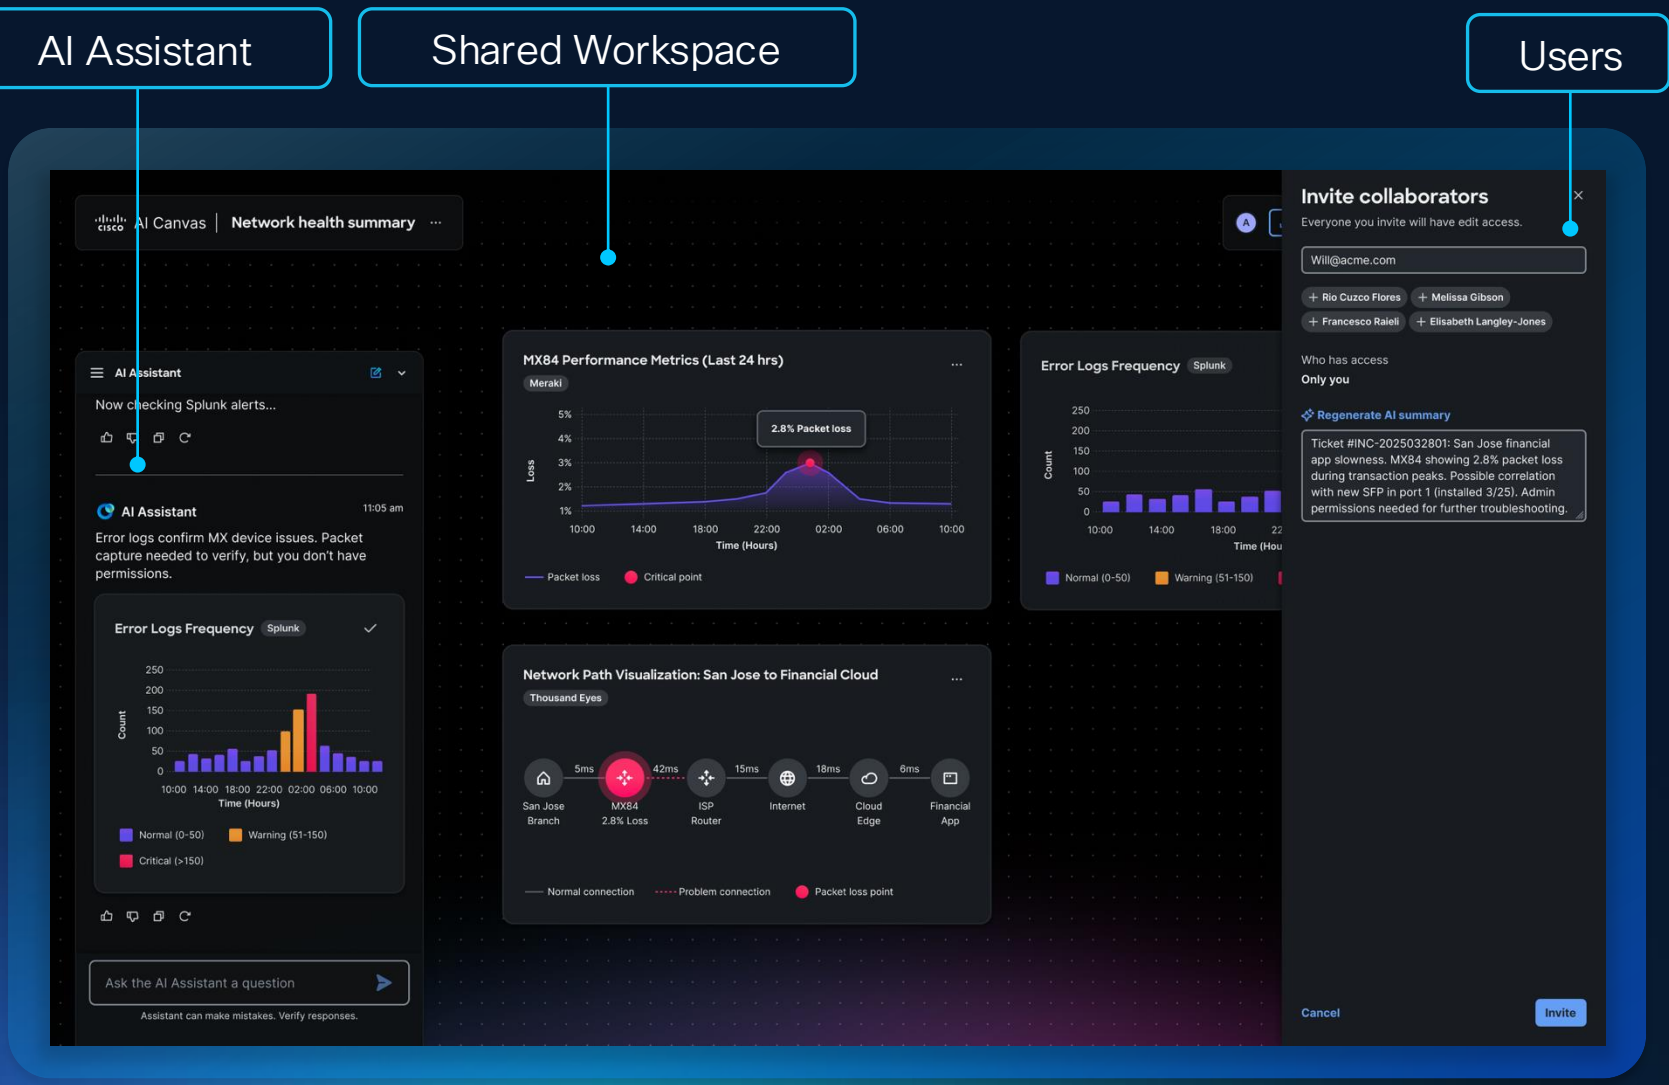

© 2026 Cisco and/or its affiliates. All rights reserved.

Leverage the data from your Cisco portfolio

Catalyst Center
webex
splunk> Cloud
SD-WAN
ThousandEyes
Cisco Meraki
Crosswork
Firewall
Nexus Hyperfabric
Cyber Vision
Duo
splunk> Enterprise Security
Nexus Dashboard
Intersight
Secure Access
XDR
Identity Intelligence
ISE
splunk> IT Service Intelligence
& more

© 2026 Cisco and/or its affiliates. All rights reserved.

Cisco first.
Third-party ITSM next.
Open MCP & A2A vision.

© 2026 Cisco and/or its affiliates. All rights reserved.

Where to access AI Canvas

© 2026 Cisco and/or its affiliates. All rights reserved.

Cisco's unified AI-native management platform across networks, security, observability, and collaboration.

https://cloud.cisco.com

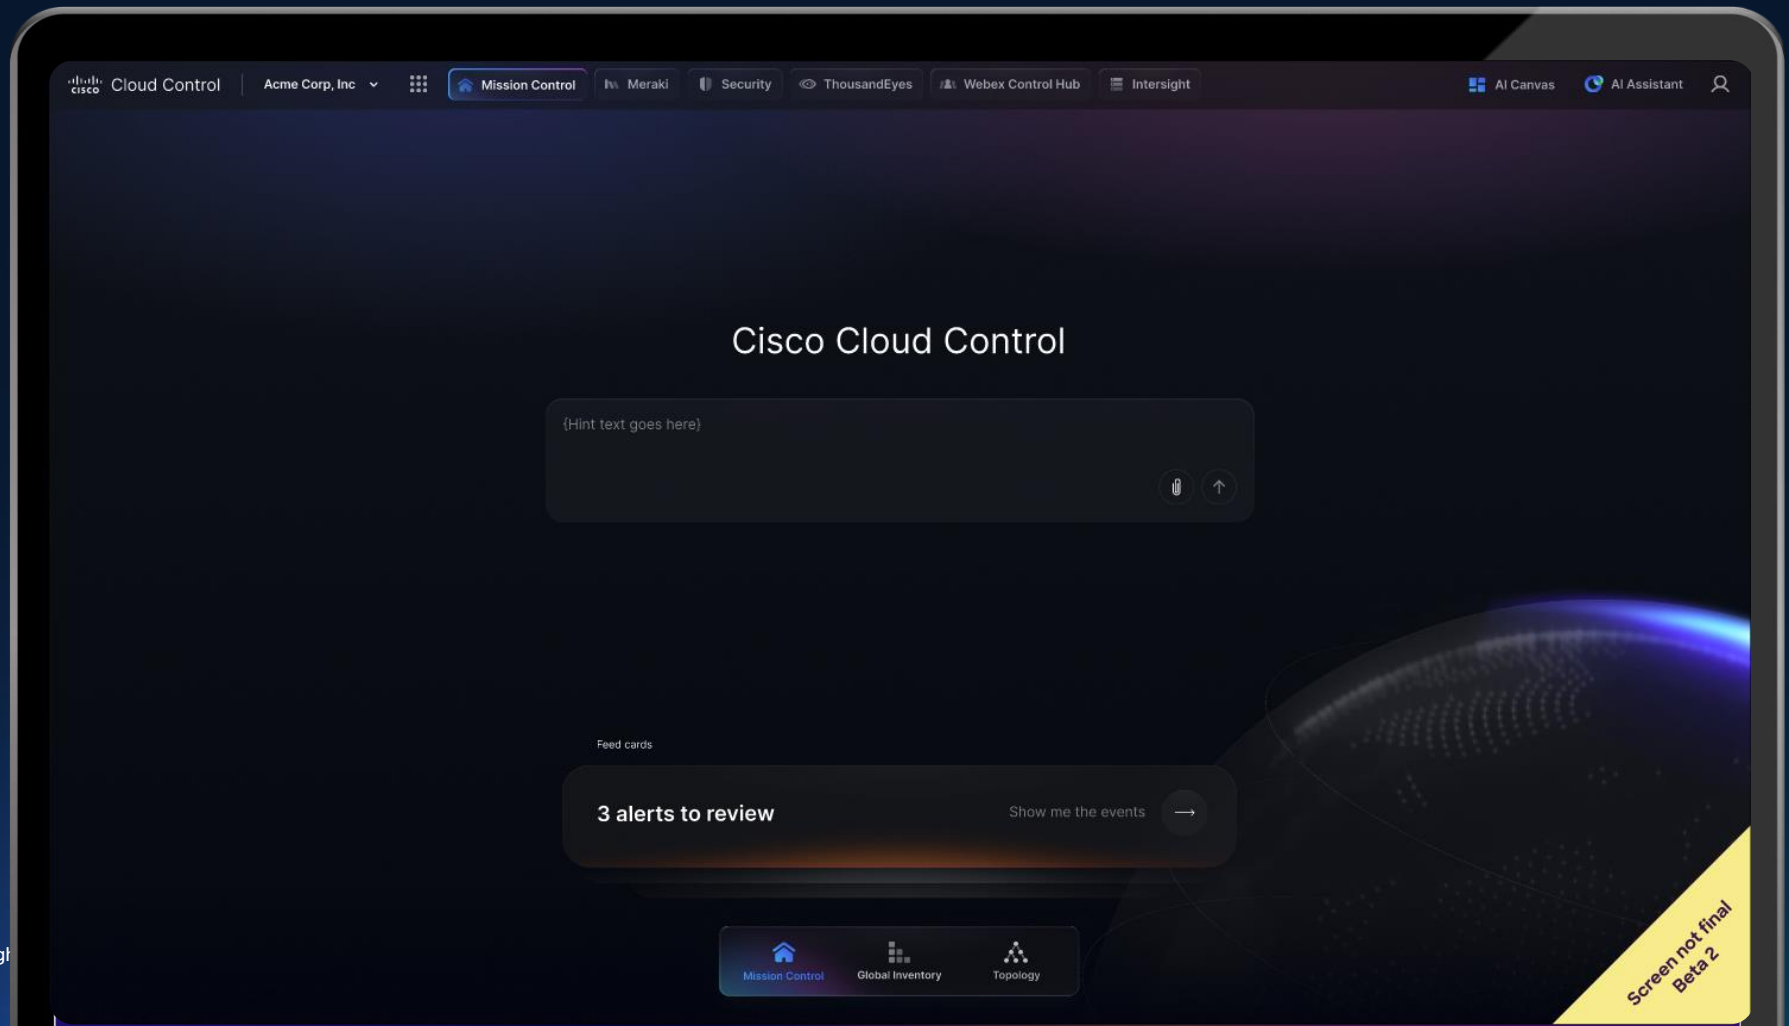

© 2026 Cisco and/or its affiliates. All rig

Introducing Cisco Cloud Control

Platform managed

1. Navigation / AI Canvas
2. Global search
3. Global AI Assistant
4. Help
5. Global notification
6. User menu
7. Platform navigator

Micro-app managed

Main content area

© 2026 Cisco and/or its affiliates. All rights reserved.

One place to **see everything**

Your entire Cisco infrastructure. One login. One view.

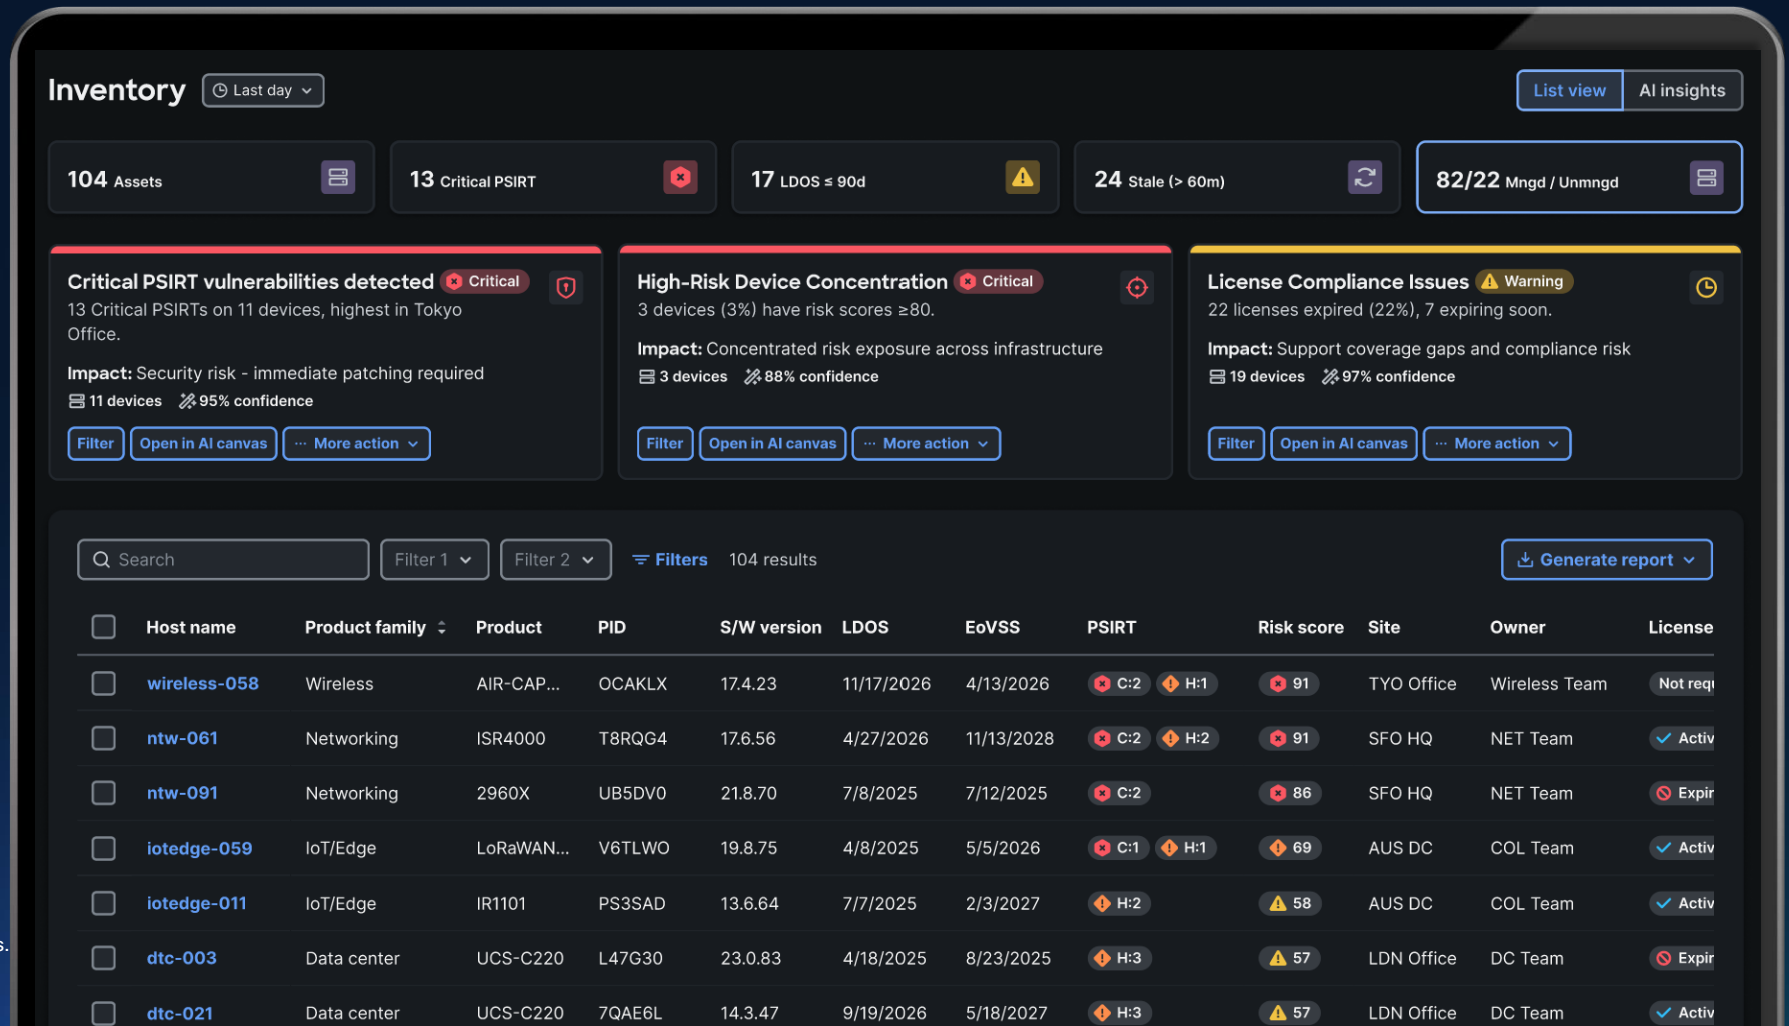

© 2026 Cisco and/or its affiliates.

AI Canvas and Cloud Control Demo

© 2026 Cisco and/or its affiliates. All rights reserved.

Demo

Cloud Control

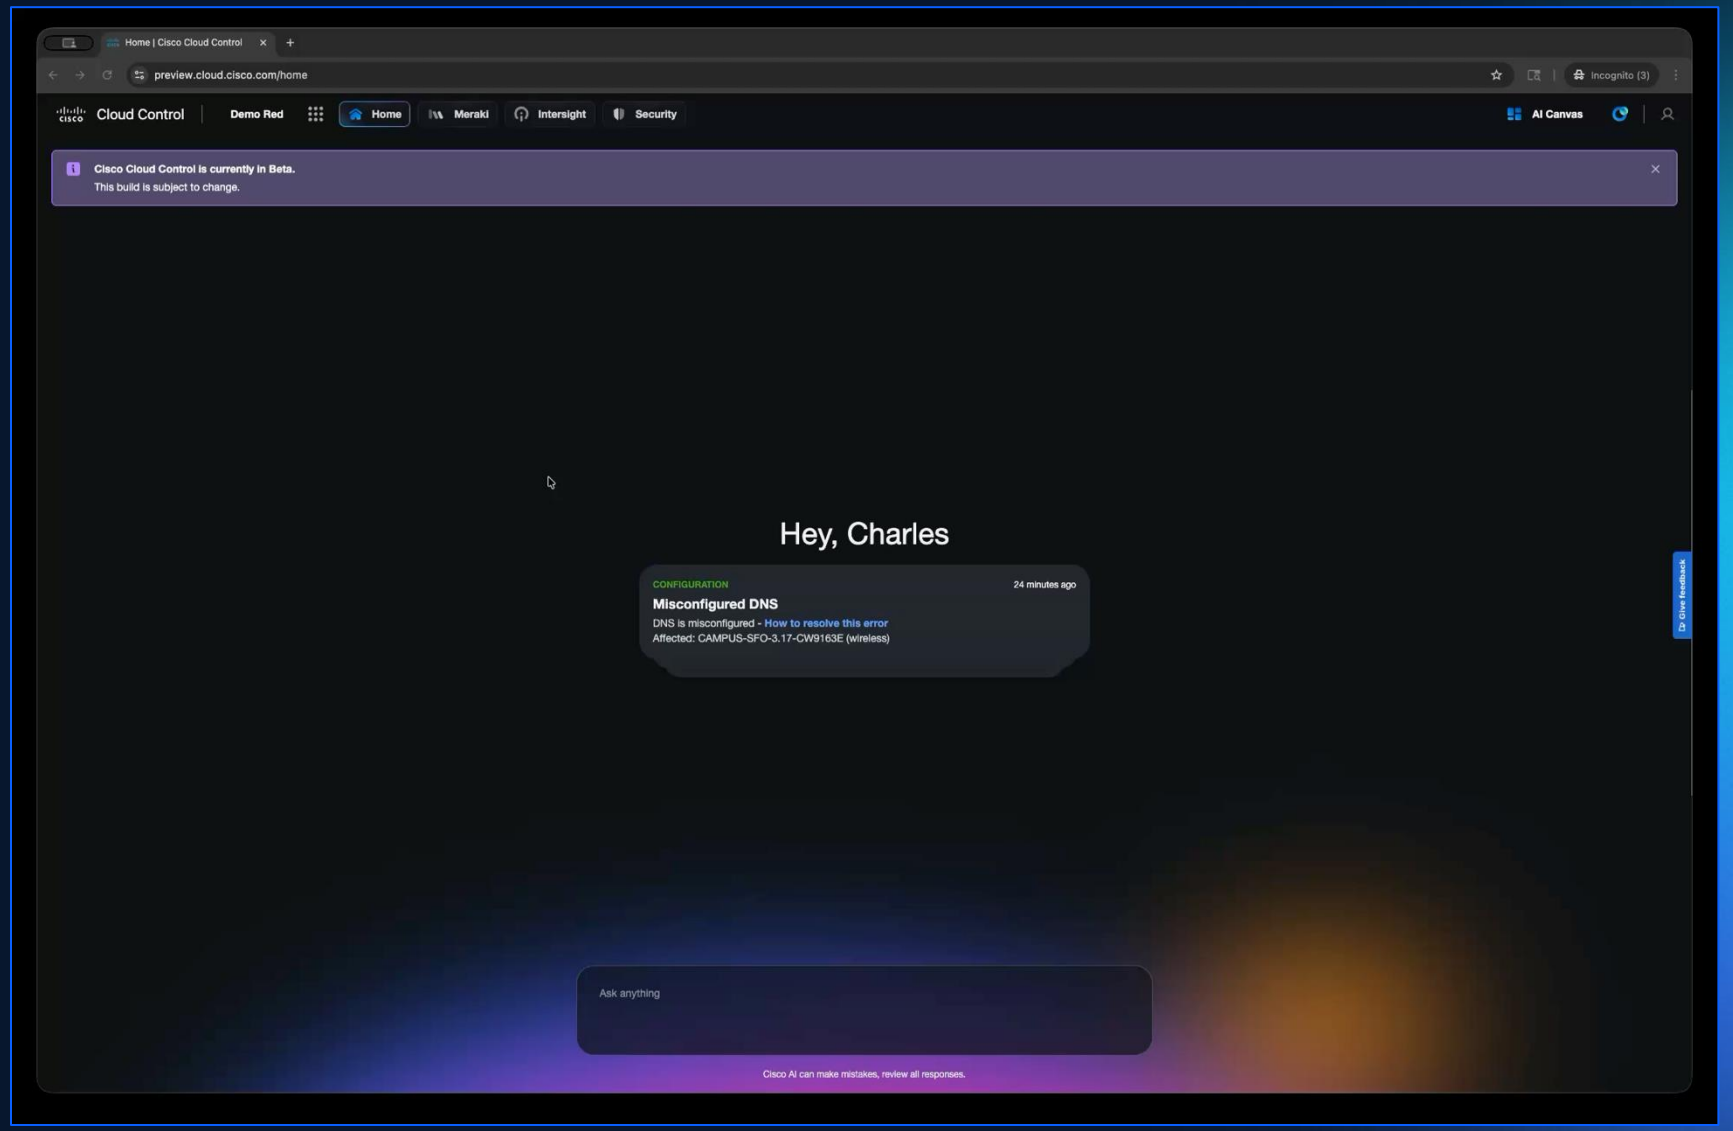

Click on the video in presenter mode to enlarge

Demo

AI Canvas with Security + Networking

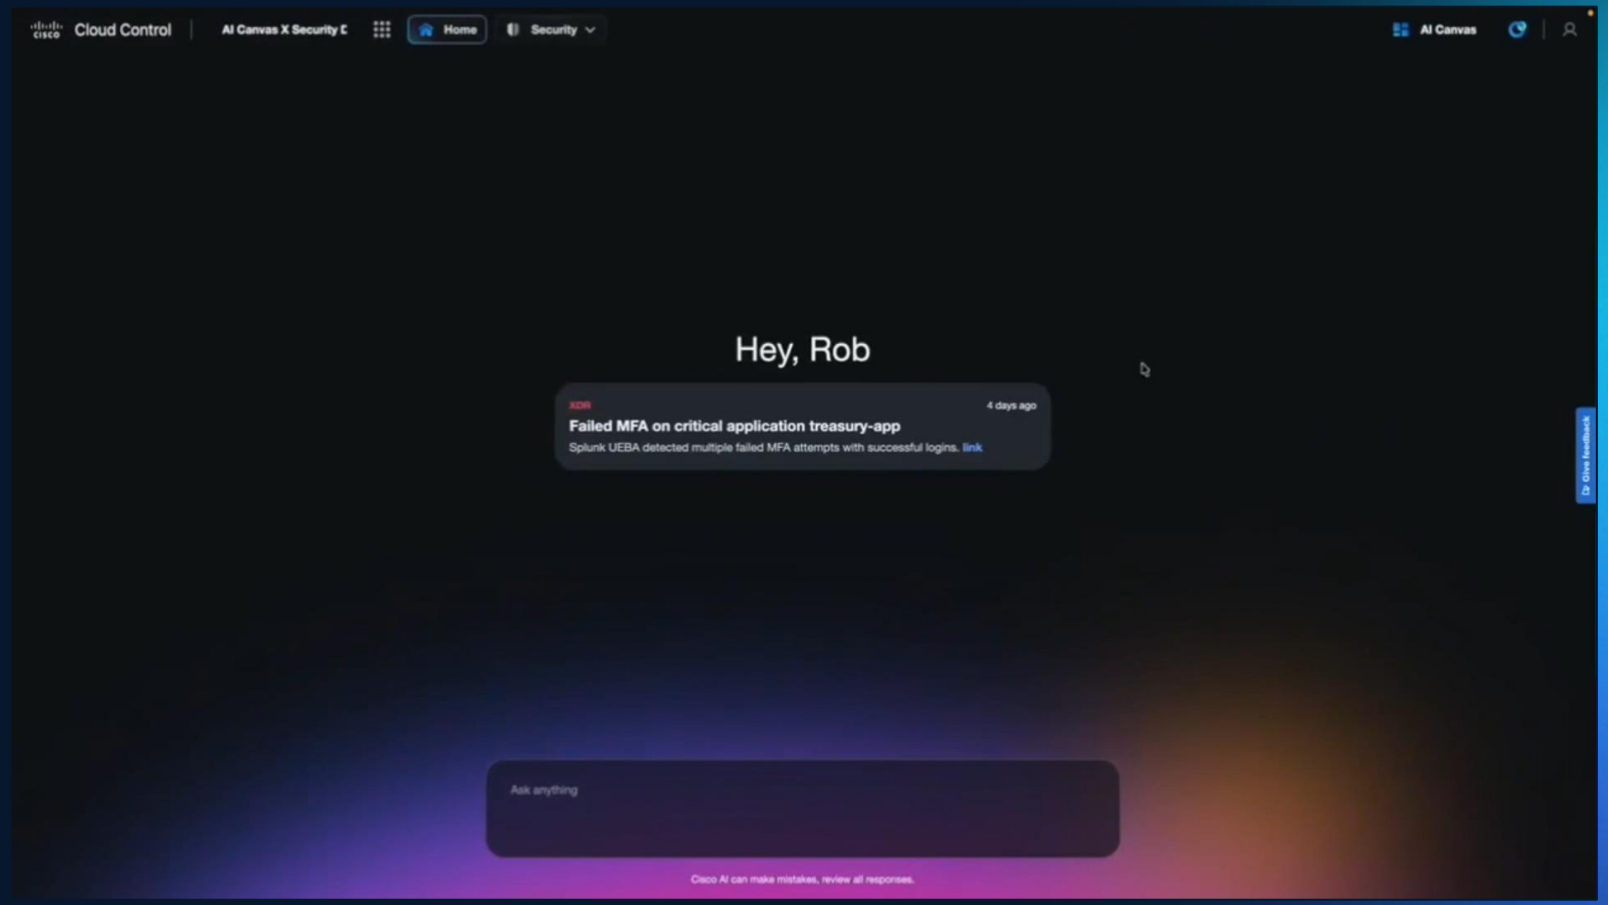

Click on the video in presenter mode to enlarge

Agenda

1. Introduction to Agentic Ops
2. Core Components of Agentic Ops
3. Cisco’s approach to Agentic Ops
4. Implementation Considerations

© 2026 Cisco and/or its affiliates. All rights reserved.

Implementation Consideration

1) Technology: Building the “AI Ready Network"

- **Controller-First Architecture:** Transition from box-by-box management to **Controllers** and **API-driven fabrics** (e.g., ACI, Cisco Catalyst Center, Meraki) to provide a programmable abstraction layer.
- **Infrastructure Readiness:** Shift from legacy SNMP polling to **High-Fidelity Streaming Telemetry** (gRPC/NetConf) to give agents real-time "sight."
- **The Reasoning Layer:** Implement **LLM-driven Agentic Frameworks** capable of goal-based planning rather than static $IF/THEN$ scripting.

2) People: Transitioning from "Scripters" to "Orchestrators"

- **Upskilling to Prompt Engineering:** Train network engineers to move from writing Python code to defining **System Prompts** and agent guardrails.
- **Trust & Verification Culture:** Establish a mindset of **"Trust but Verify,"** understand "Chain of Thought”
- **AI Collaboration:** Define the **Human-in-the-Loop (HITL)** roles—empowering engineers to act as "Supervisors" who approve high-risk agent plans.
- **Psychological Safety:** Address the cultural shift by positioning agents as "force multipliers" that handle toil, allowing humans to focus on high-level architecture

3) Security and Governance

© 2026 Cisco and/or its affiliates. All rights reserved.

MCP and A2A: Securing the New Attack Surface

Addressing vulnerabilities across the agentic lifecycle

Supply Chain

- **Tool Poisoning:** Malicious natural language "descriptions" designed to hijack LLM intent and exfiltrate data
- **Malicious Servers:** Rogue MCP servers hosting compromised code, backdoors, or unauthorized prompt templates
- **Identity Impersonation:** Naming attacks (typosquatting) that register deceptive agents or servers to intercept traffic
- **Dependency Vulnerabilities:** Exploiting unpatched software or libraries within the agentic stack
- **Privilege Escalation:** MCP servers operating with broader permissions than the end-user, leading to unauthorized actions

Runtime

- **Indirect Prompt Injection:** Malicious commands embedded in external data (websites/docs) that override agent logic
- **Goal Hijacking:** Multi-turn manipulation designed to shift the agent's objective toward attacker-defined ends
- **Multi-Modal Exploits:** Using malicious images, audio, or sensor data to corrupt the agent's perception and decision-making
- **Resource Exhaustion (DoS):** Triggering infinite reasoning loops or excessive API calls to degrade system performance
- **Data Exfiltration:** Exploiting tool-use to bypass traditional DLP and leak sensitive context to unauthorized endpoints.

© 2026 Cisco and/or its affiliates. All rights reserved.

Source: https://learn-cloudsecurity.cisco.com/ai-security-framework

cisco

Next steps

Connect with your Cisco account representative or get started with Cisco

Visit our brand-new AI Canvas web page

Learn more about AgenticOps

© 2026 Cisco and/or its affiliates. All rights reserved.

© 2026 Cisco and/or its affiliates. All rights reserved.