

Simplifying Network Operations with Cisco Agentic AIOps

Tony Davitt

Head of Technical Strategy, Cisco Ireland & Scotland

Brian O'Donoghue

Solution Engineer, Cisco Ireland

Last login: Mon Apr 13 10:30:05 on ttys000

The default interactive shell is now zsh.

To update your account to use zsh, please run `chsh -s /bin/zsh`.

For more details, please visit <https://support.apple.com/kb/HT208050>.

[TODAVITT-M-6H2M:~ todavitt\$ ping 10.3.1.6

PING 10.3.1.6 (10.3.1.6): 56 data bytes

Request timeout for icmp_seq 0

Request timeout for icmp_seq 1

Request timeout for icmp_seq 2

Request timeout for icmp_seq 3

Request timeout for icmp_seq 4

Request timeout for icmp_seq 5

Request timeout for icmp_seq 6

^C

--- 10.3.1.6 ping statistics ---

8 packets transmitted, 0 packets received, 100.0% packet loss

TODAVITT-M-6H2M:~ todavitt\$

Last login: Mon Apr 13 10:30:13 on ttys000

The default interactive shell is now zsh.

To update your account to use zsh, please run `chsh -s /bin/zsh`.

For more details, please visit <https://support.apple.com/kb/HT208050>.


[TODAVITT-M-6H2M:~ todavitt\$ traceroute 10.3.1.6

traceroute to 10.3.1.6 (10.3.1.6), 64 hops max, 40 byte packets

1 skyrouter.home (192.168.0.1) 14.657 ms 2.981 ms 3.975 ms

2 * * *

3 *^C *







Apply a display filter ... <Ctrl-/>

Expression... +

No.	Time	Source	Destination	Protocol	Length	Info
10	9.041865	192.168.200.21	192.168.200.135	TCP	66	cisco-sccp(2000) → 7876 [SYN, ACK] Seq=...
11	9.047489	192.168.200.135	192.168.200.21	TCP	60	7876 → cisco-sccp(2000) [ACK] Seq=1 Ac...
12	9.047526	192.168.200.135	192.168.200.21	TCP	15...	7876 → cisco-sccp(2000) [ACK] Seq=1 Ac...
13	9.047543	192.168.200.21	192.168.200.135	TCP	54	cisco-sccp(2000) → 7876 [ACK] Seq=1 Ac...
14	9.047559	192.168.200.135	192.168.200.21	TCP	15...	7876 → cisco-sccp(2000) [ACK] Seq=1461
15	9.047567	192.168.200.21	192.168.200.135	TCP	54	cisco-sccp(2000) → 7876 [ACK] Seq=1 Ac...
16	9.047570	192.168.200.135	192.168.200.21	TCP	15...	7876 → cisco-sccp(2000) [ACK] Seq=2921
17	9.047574	192.168.200.21	192.168.200.135	TCP	54	cisco-sccp(2000) → 7876 [ACK] Seq=1 Ac...
18	9.047577	192.168.200.135	192.168.200.21	TCP	15	7876 → cisco-sccp(2000) [ACK] Seq=4381

- Frame 12: 1514 bytes on wire (12112 bits), 1514 bytes captured (12112 bits) on interface 0
- Ethernet II, Src: Dell_96:12:0e (ec:f4:bb:96:12:0e), Dst: Vmware_b4:90:14 (00:0c:29:b4:90:14)
 - Destination: Vmware_b4:90:14 (00:0c:29:b4:90:14)
 - Source: Dell_96:12:0e (ec:f4:bb:96:12:0e)
 - Type: IPv4 (0x0800)
 - Internet Protocol Version 4, Src: 192.168.200.135 (192.168.200.135), Dst: 192.168.200.21 (192.168.200.21)
 - Transmission Control Protocol, Src Port: 7876 (7876), Dst Port: cisco-sccp (2000), Seq: 1, Ack: 1, Len: 1460
 - Data (1460 bytes)

```

0000  00 0c 29 b4 90 14 ec f4 bb 96 12 0e 08 00 45 00  ..)....E.
0010  05 dc 1d 1f 40 00 80 06 c6 0e c0 a8 c8 87 c0 a8  ....@.....
0020  c8 15 1e c4 07 d0 6a f0 7c f6 6f 9b 26 e0 50 10  ....j.|o&P.

```

Evolution of Network Operations



Simplifying Network Operations with Cisco Agentic AIOps

Brian O'Donoghue

Solution Engineer, Cisco Ireland

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

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?

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

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.

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

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

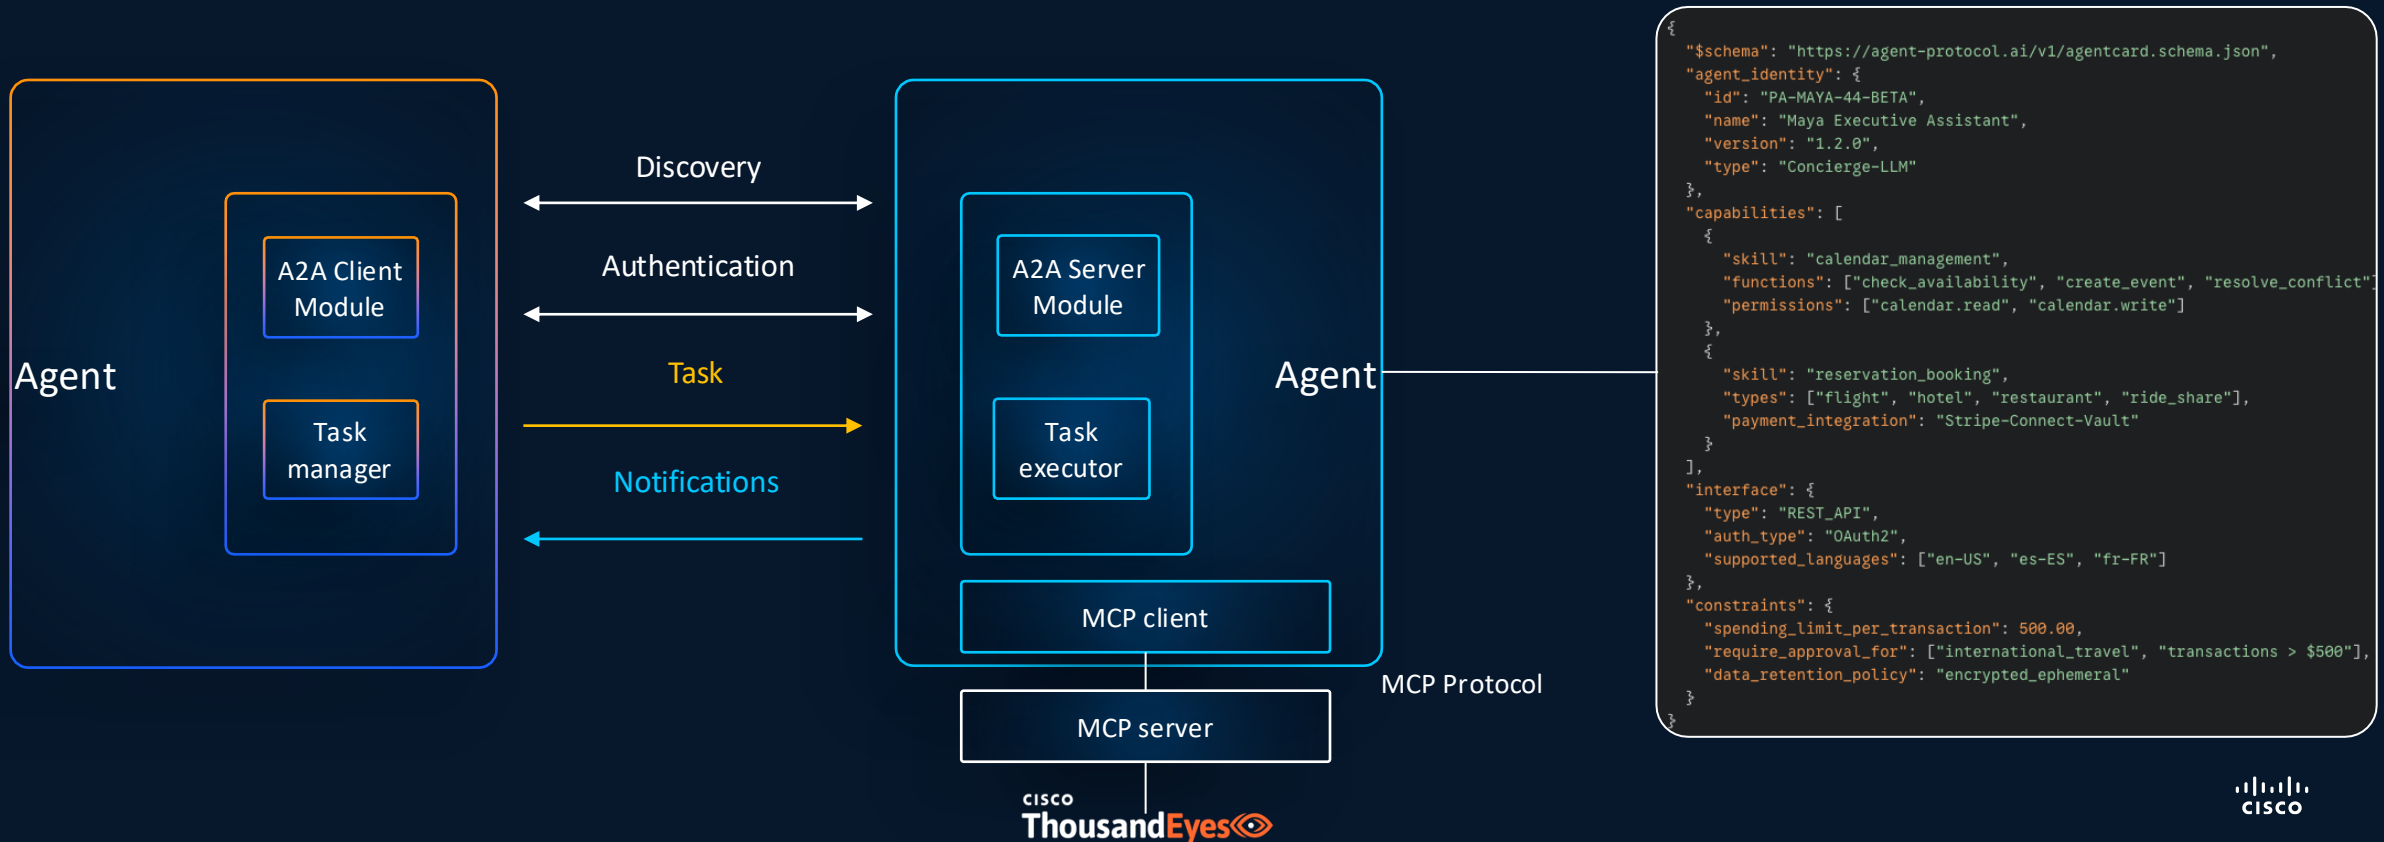
Leverages LLMs to decompose complex problems into a sequence of actionable, logical steps

Action - tool use

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

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.



Introducing AI Packet Analyzer

Accelerating issue resolution with AI

Automatically seizes the moments of client failure and uploads PCAP files to the AI cloud

Quickly analyzes PCAPs with AI summary

Isolates fault with reasoning view or visually with packet flow view

Actionable recommendation provided to resolve issue

500 Million+ Proactive PCAPs Performed

The screenshot displays the Meraki AI Packet Analyzer interface. At the top, there are two callout boxes: "AI Summary" and "Reasoning". The interface is divided into several sections:

- AI Summary:** A panel titled "Summary AI-generated" with a "Complete PCAP analysis" status. It contains a text description of a client connection issue related to an invalid PMKID during a "Welcome2Helz 6G" network. Below the text are two buttons: "Packet Flow" and "Reasoning".
- Reasoning:** A panel titled "Reasoning AI-generated" with a "Complete PCAP analysis" status. It provides detailed analysis of frames:
 - Frames 21-35 : Action Frames:** Describes radio management and optimization actions like neighbor discovery and Block ACK setup.
 - Frames 28-33 : ARP:** Describes ARP requests for network resource discovery, noting improved signal strength around -60 dBm.
 - Frames 36-37 : DNS:** Describes successful DNS queries and the beginning of application-level network usage, with signal strength around -58 dBm.Below the analysis, it lists suggested actions to solve the issue, such as checking for fast roaming misconfiguration or stale key cache entries.

- Packet Capture:** A table at the bottom showing captured packets. The table has columns for No., Time, Source, Destination, and Transmission. The first three rows are highlighted in blue.

No.	Time	Source	Destination	Transmission
1	0.000000000	ce:0e:f4:59:6a:2b	Broadcast	ce:0e:f4:59:6a:2b
2	0.009151000	ce:d6:66:66:c9:80	ce:0e:f4:59:6a:2b	ce:d6:66:66:c9:80
3	0.591894000	ce:0e:f4:59:6a:2b	Broadcast	ce:0e:f4:59:6a:2b

Packet Capture

Remediation suggestion

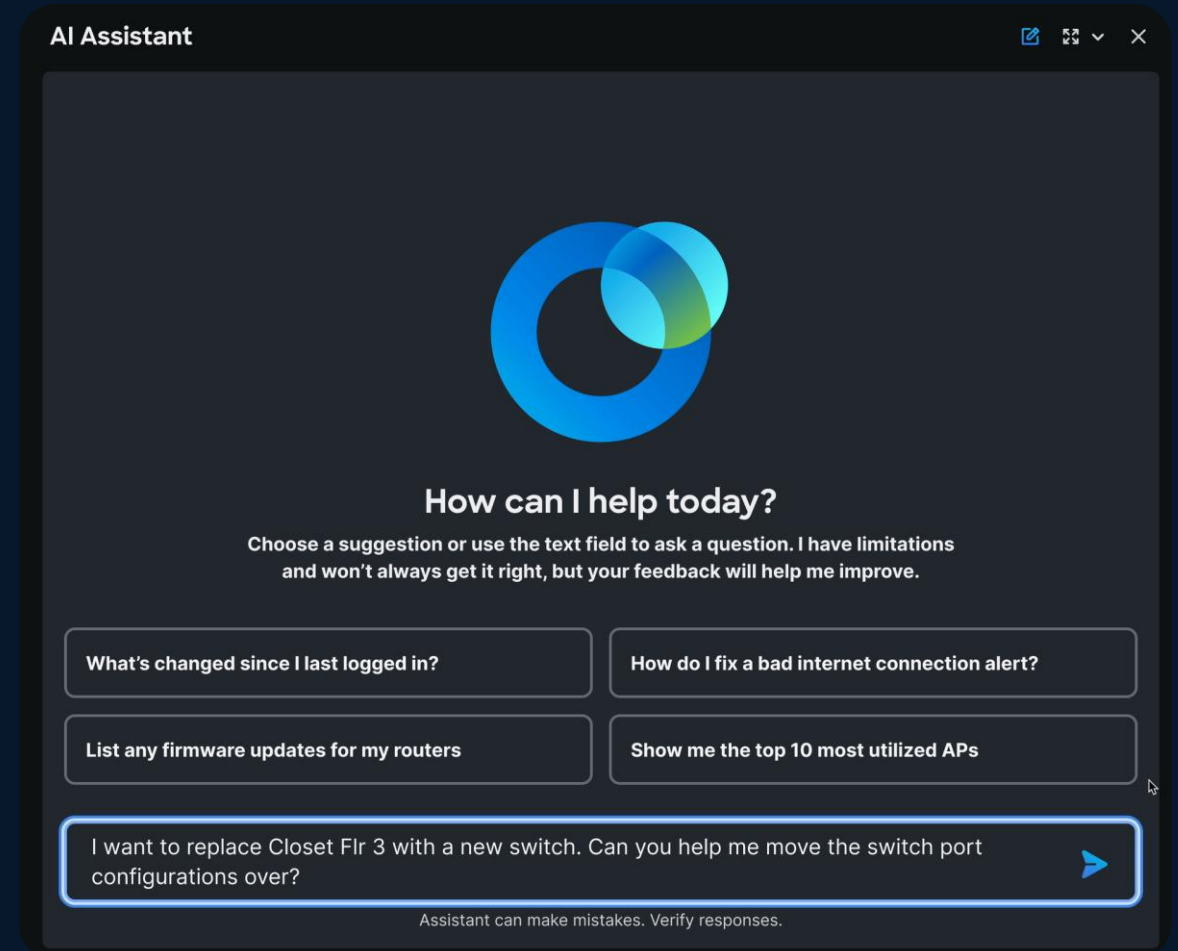


New Networking Skills for AI Assistant

Ask, explore and act in natural language

New automated workflows for configs changes and switch migrations

New integrated assurance capabilities



Implementation Considerations

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

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.

Cisco Networking with AgenticOps

From:

- Manual configuration and ticket triage
- Limited visibility and context beyond the network
- Manual network optimization
- Fragmented tools and processes to ID and solve issues

To:

- Agentic AI-powered configuration to reduce manual errors
- Predictive, end-to-end visibility
- Automated AI-driven network optimization
- Rapid, predictive detection and remediation

Powered by a unified platform for both on-premises and cloud environments

Demo Time

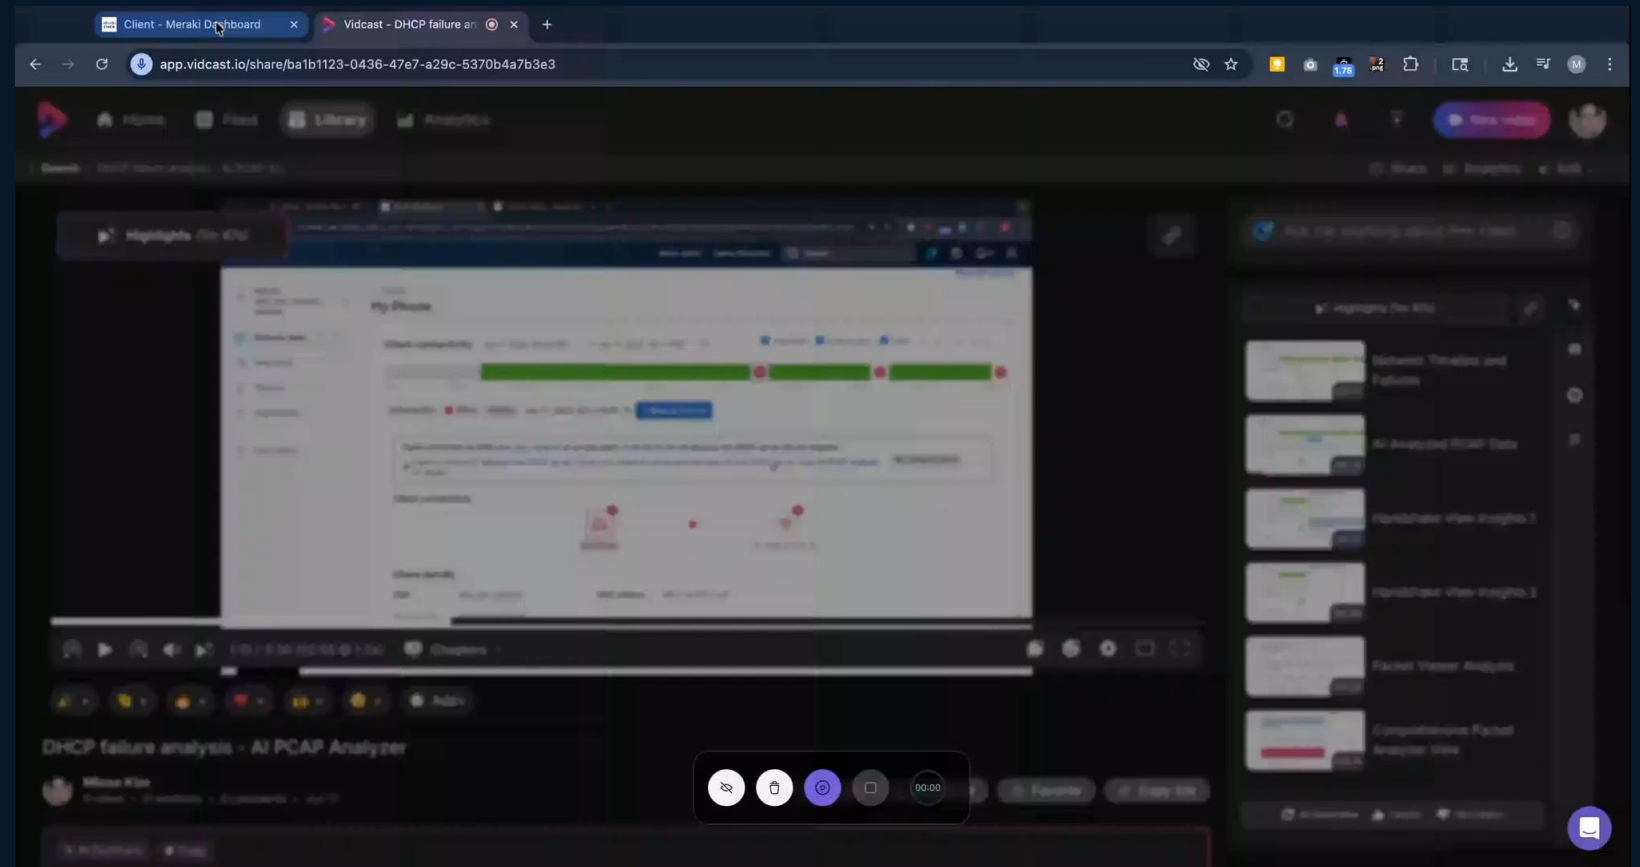
Proactive AI Packet Capture and Analysis

AI Config Recommendations

AgenticOps with Cisco AI Canvas

Demo

Proactive AI Packet Capture and Analysis



Demo

AI Config Recommendations

The screenshot displays the Cisco Meraki Alerts dashboard. At the top, the navigation bar includes the Cisco Meraki logo, 'Show admin', 'Demo Networks', and a search bar. The left sidebar lists various network management categories: Network (City Hall), Network-wide, Assurance, Wireless, Cameras, Sensors, and Organization. The main content area is titled 'Alerts' and shows a filter for 'Last week'. Below this, there are tabs for 'Alerts' (0) and 'Optimizations' (10). A bar chart titled 'Alerts triggered over time' shows the number of alerts per hour from June 10 to June 17, 2025. A pop-up window titled 'Alert Counts' for the period 2025-06-13 05:47 - 2025-06-13 10:48 shows 8 Critical alerts (Unreachable device) and 0 Warning alerts. Below the chart, there are summary cards for '0 Critical' and '0 Warning' alerts. At the bottom, there are sections for 'Top alert counts' by network and by alert type.

Alert Type	Current Period	Previous Week
Critical	8	0
Warning	0	0

AgenticOps building on AIOps:

Product operators must work across domains and with each other to enable great end-user experiences



Cross-Product Troubleshooting is Complex

Individual products don't scale to triage across domains effectively.



Cross-Team Collaboration Has High Overhead

Effort is required across siloed teams and tools, which is time consuming.



Cross-Team Handoffs Lack Full Context

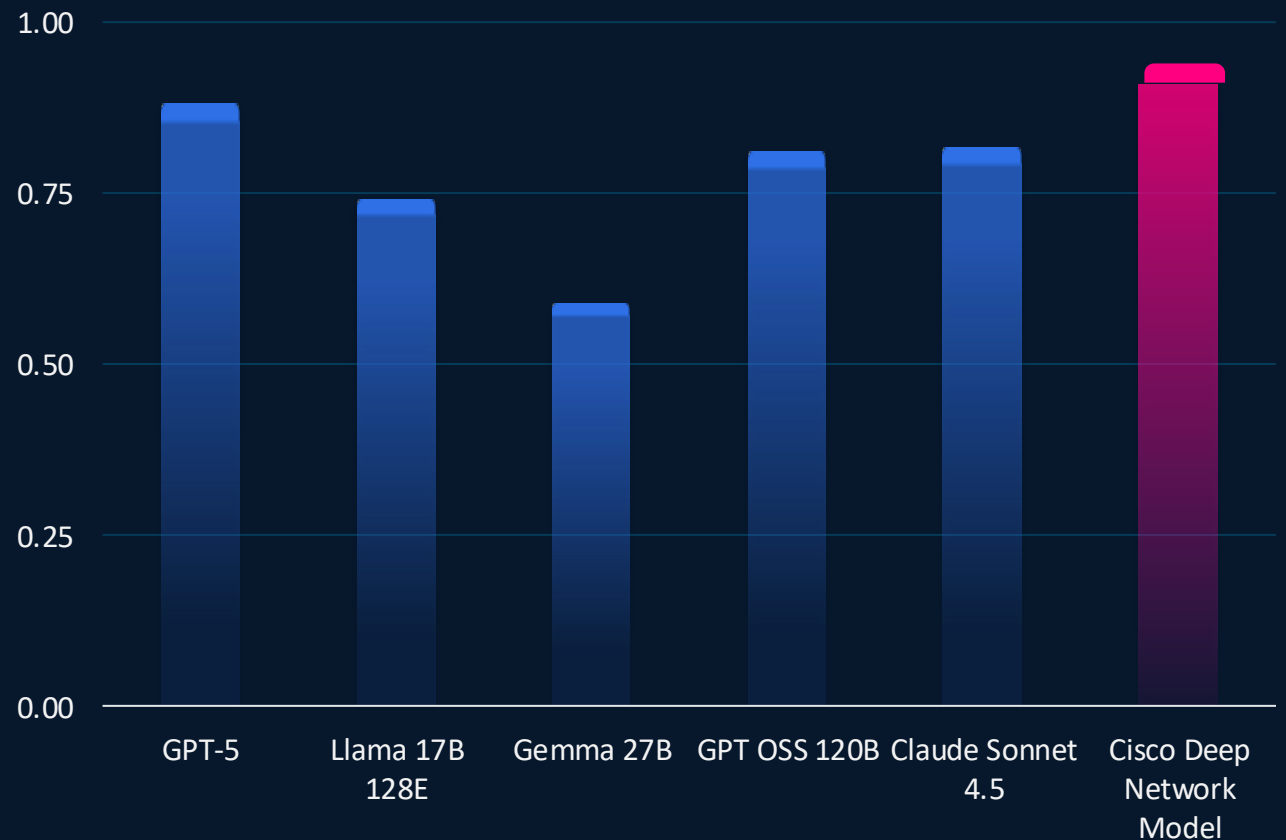
Root cause analysis is a challenge because of insufficiently detailed handoffs.

Leveraging the power of Cisco Deep Network Model

Purpose-built for networking, expert accuracy

- Fine-tuned on 40+ years of expertise and expert-vetted for accuracy
- More precise reasoning for troubleshooting, configuration, and automation
- Up to 5x fewer tool-calls for troubleshooting tasks; up to 3x lower latency for networking Q&A

Outperforms general-purpose models by ~20%



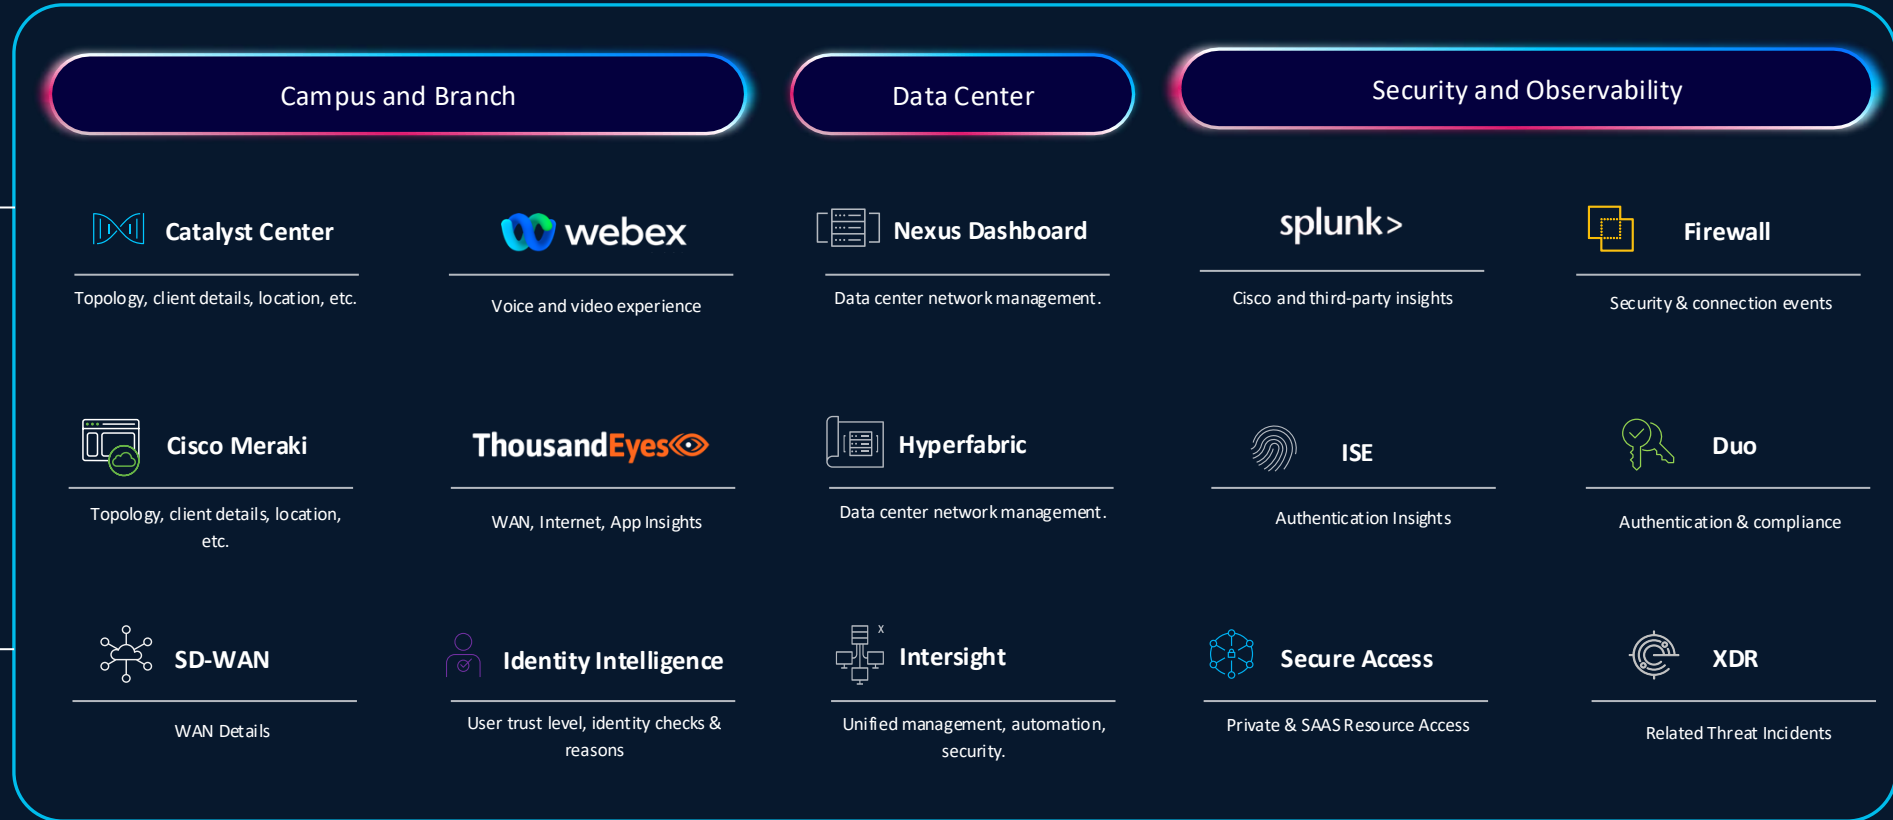
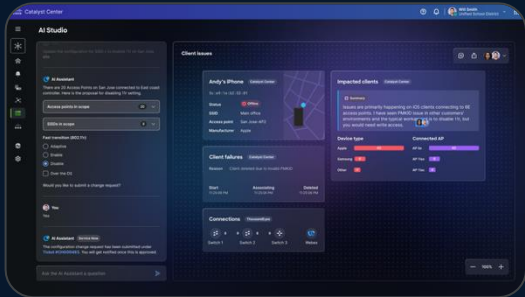
Accuracy on CCIE-style multiple choice questions (590-question benchmark), October 2025

AgenticOps with cross-product skills and unified data

AI Assistant

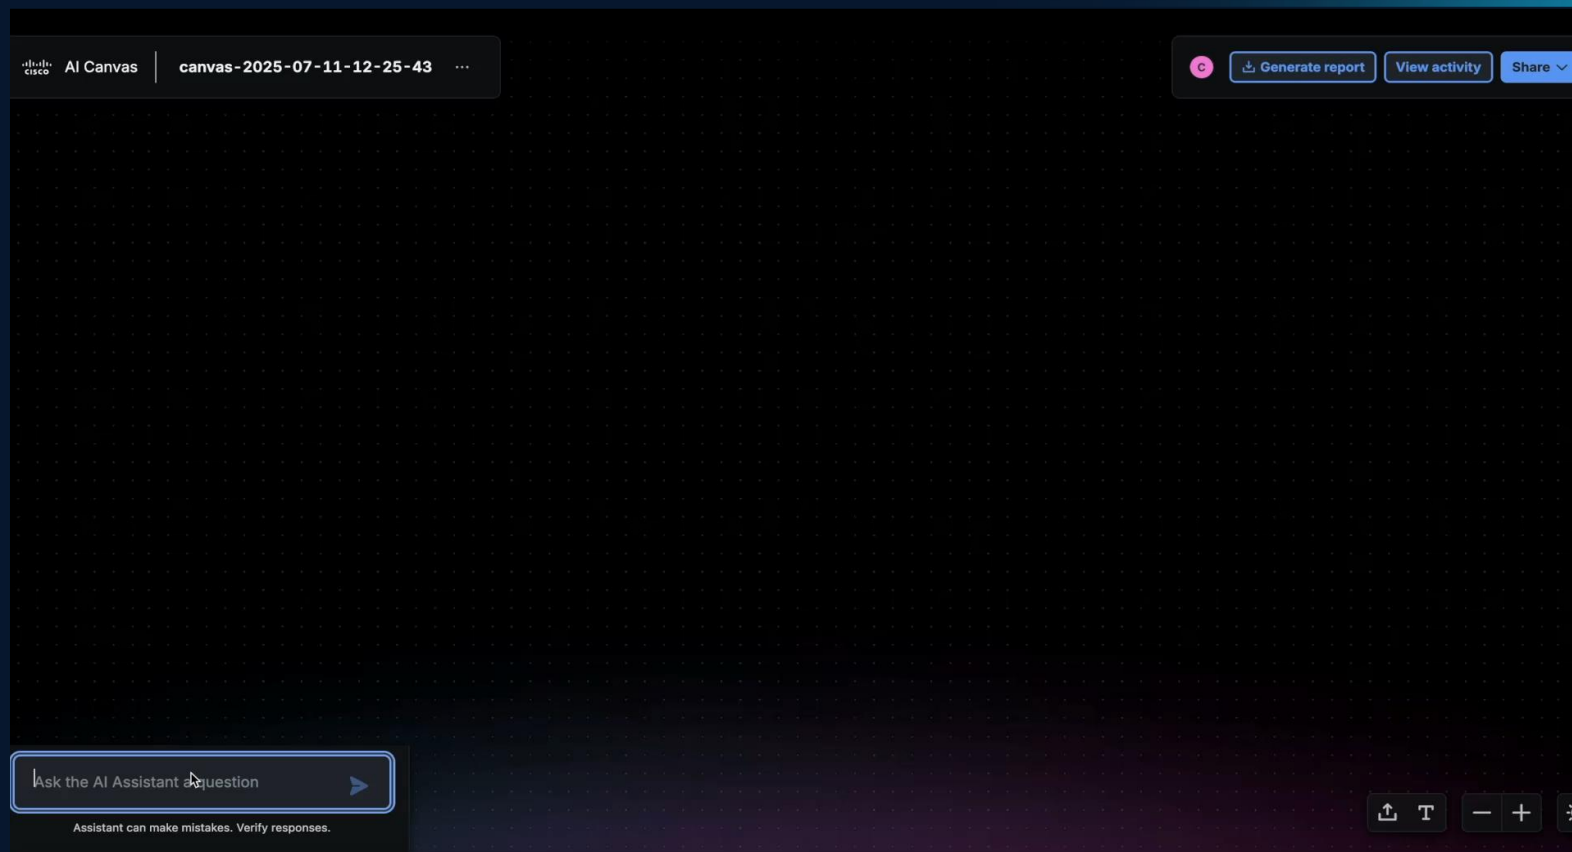


AI Canvas



Demo

AI Canvas with Meraki, ThousandEyes and Splunk



Customer zero

Quoted from a live feedback session

“ ...data isn't actionable alone. You need to interpret it, which takes a lot of human time. **AI Canvas gives us the potential to step back as humans - step out of the toil - and into a more privileged position where we can make decisions.** AI generates the insights. We decide. ”



Chris Tomazic

Technical Leader, Tech Systems Engineering

What you reduce

DECREASED

Mean time to resolution

Faster isolation leads directly to faster fixes.

DECREASED

Manual toil

Less dashboard hopping, log digging, and ad-hoc querying.

DECREASED

Cognitive load

Operators consume a single investigative narrative instead of fragmented views.

DECREASED

Reliance on tribal knowledge

Fewer escalations to “the one expert who knows this system.”

What you gain

INCREASED
Operator leverage

One human can supervise more systems and incidents because AI performs correlation and reasoning.

INCREASED
Decision consistency

Investigations follow the same logic every time, independent of who is on call or on shift.

INCREASED
Confidence in actions

Causal reasoning and explainability reduce guesswork and second-guessing.

IT operations are entering the AgenticOps era
where humans and AI agents work side-by-side to drive faster and safer outcomes.

Cisco's AI Canvas is the tool that brings it to life



