

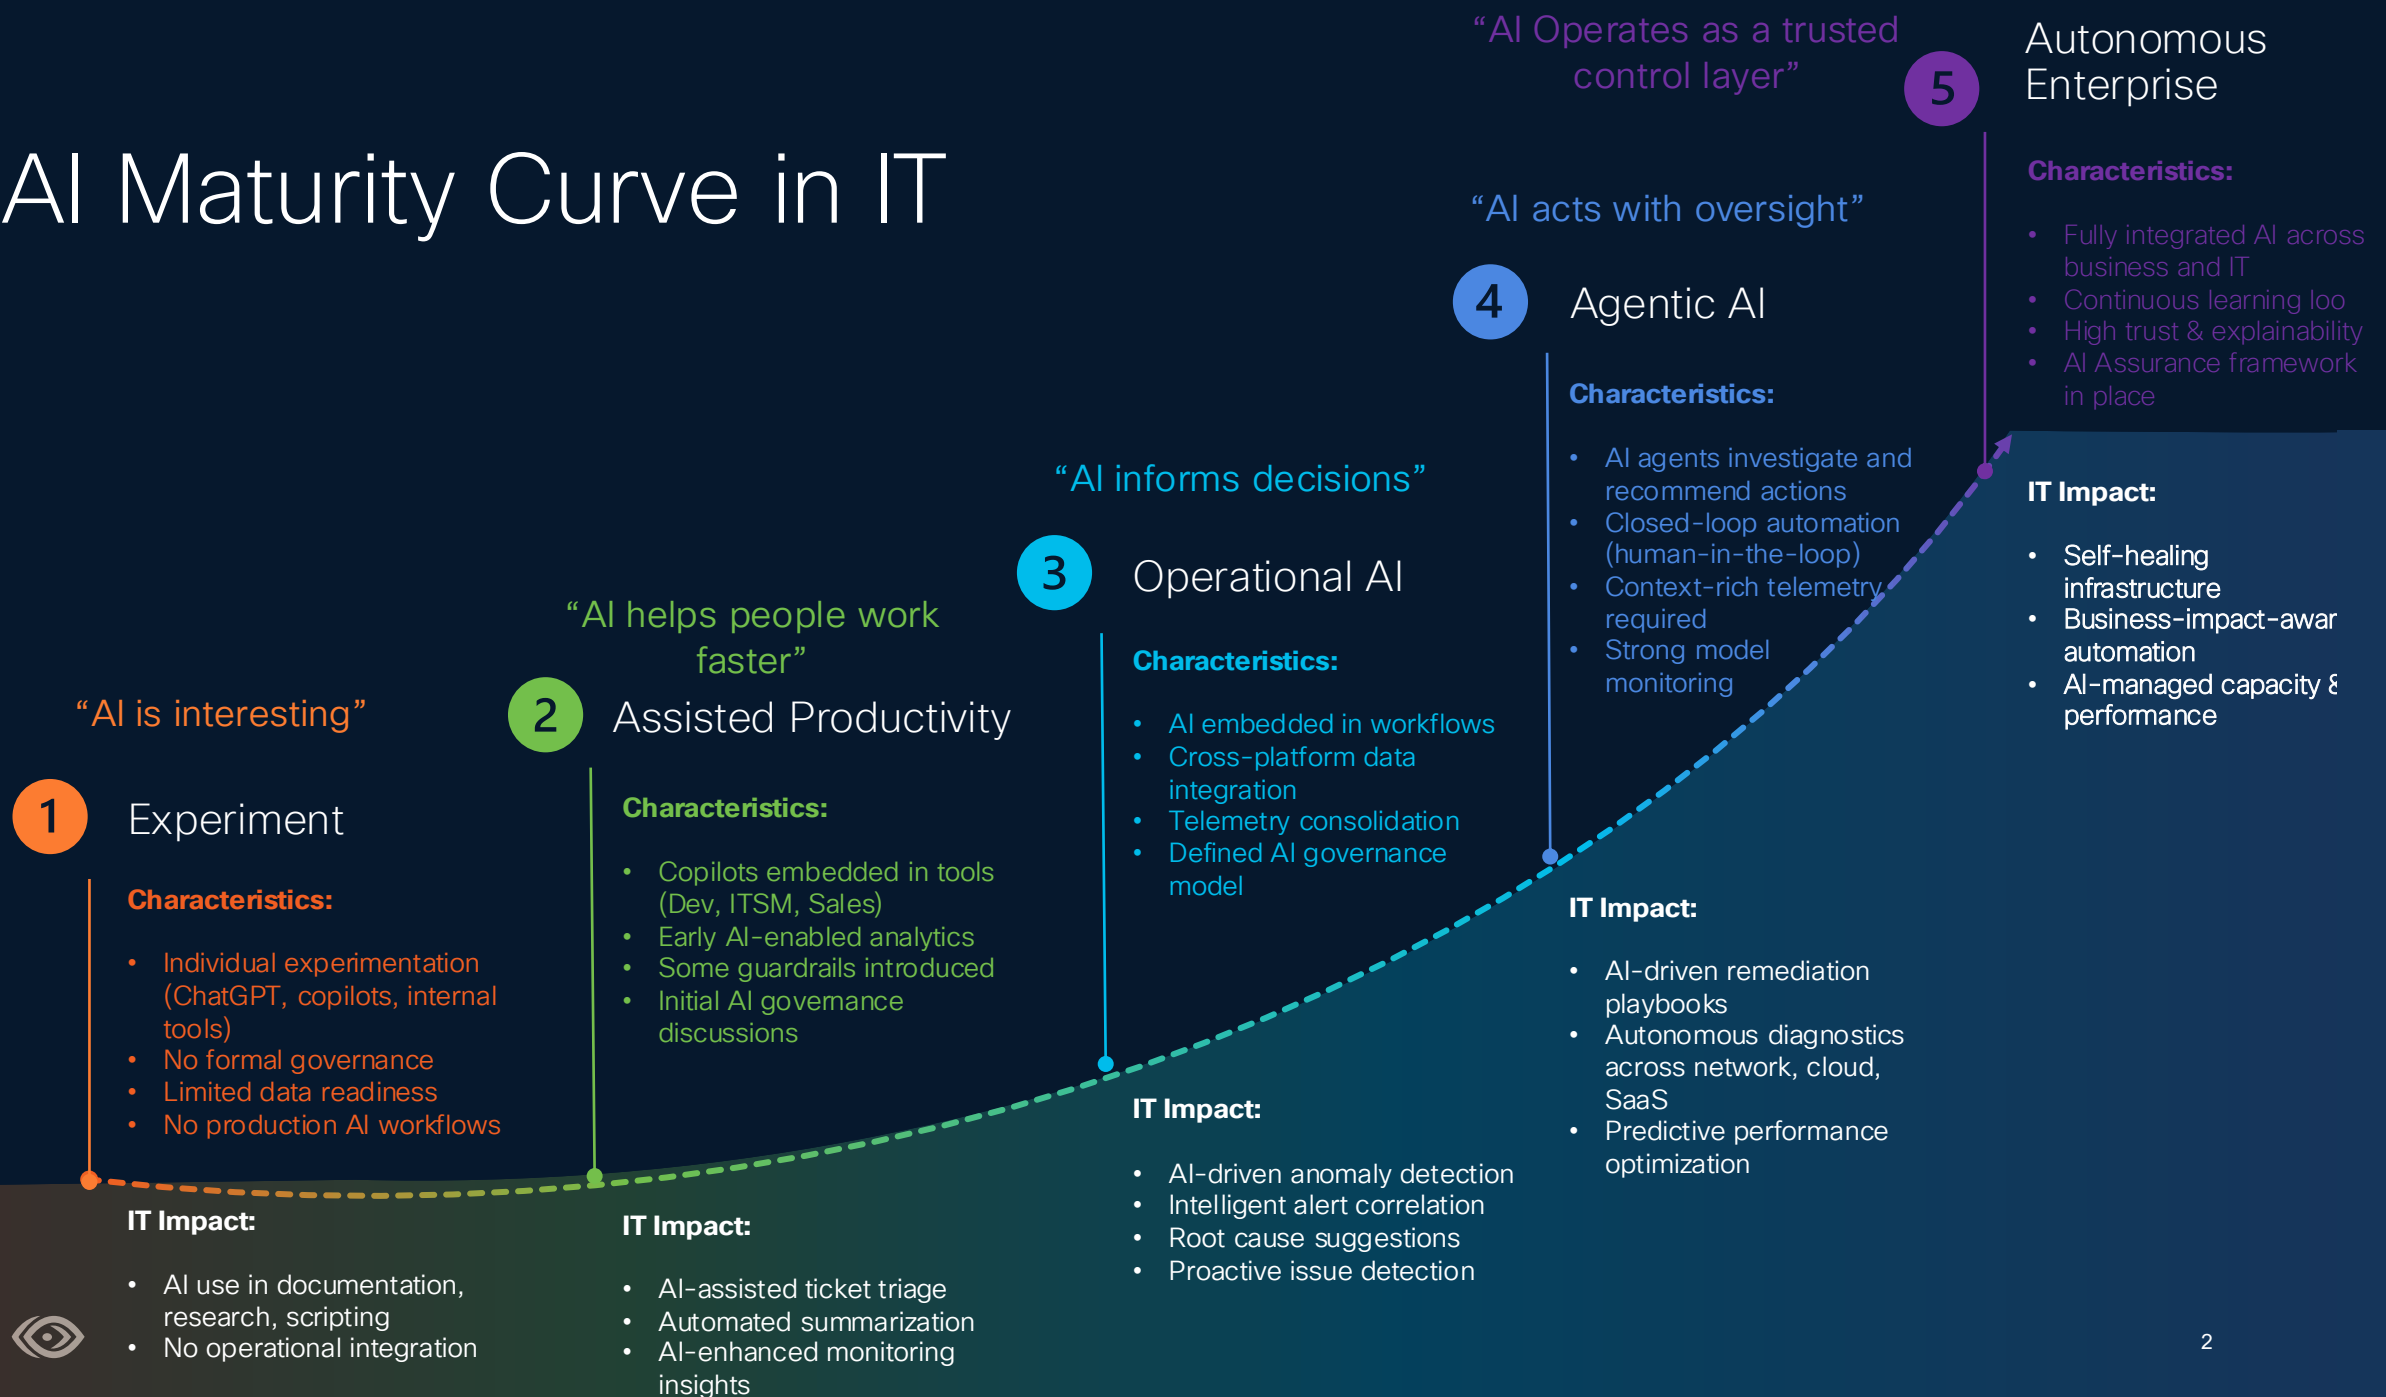
ThousandEyes and AgenticOps

Bill Donoghue
ThousandEyes Solutions Engineering Leader

March 31, 2026



AI Maturity Curve in IT



Cisco ThousandEyes Platform

Experience & Service Assurance



Digital experience

End user | Mobile
IIoT | VoIP



Proactive monitoring

Web | Network
BGP | DNS



Global visibility

On-Prem | Cisco
3rd party



Domain insights

Traffic | WAN
Internet | Cloud



AI-powered platform

Automation | Integrations | APIs



Path
visualization



Cross-layer
correlation



Anomaly detection &
alerting



Event summarization



Predictive
recommendations



Natural language
query

How It Works

CLOUD AGENT



- 400+ ThousandEyes maintained POPs
- Global scale
- T1/2 DCs, Cloud and Broadband providers
- Outside-in visibility
- Public facing sites and APIs
- Customer experience

ENTERPRISE AGENT

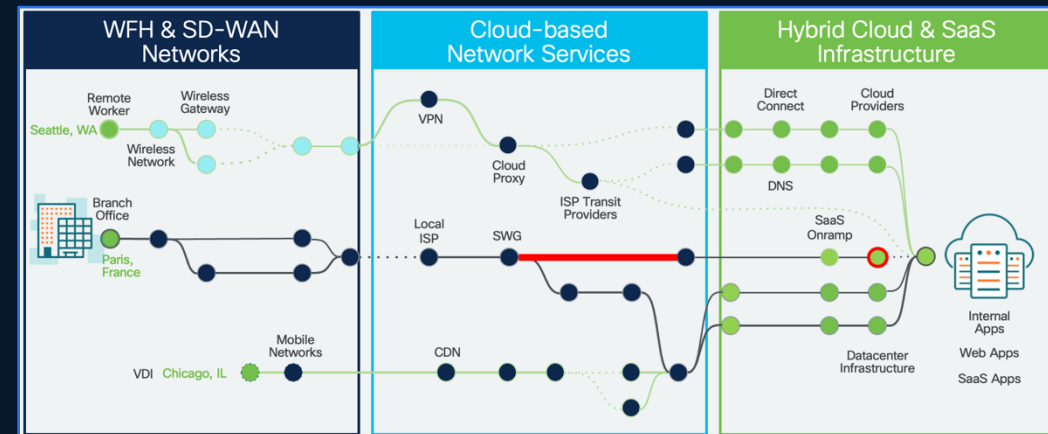
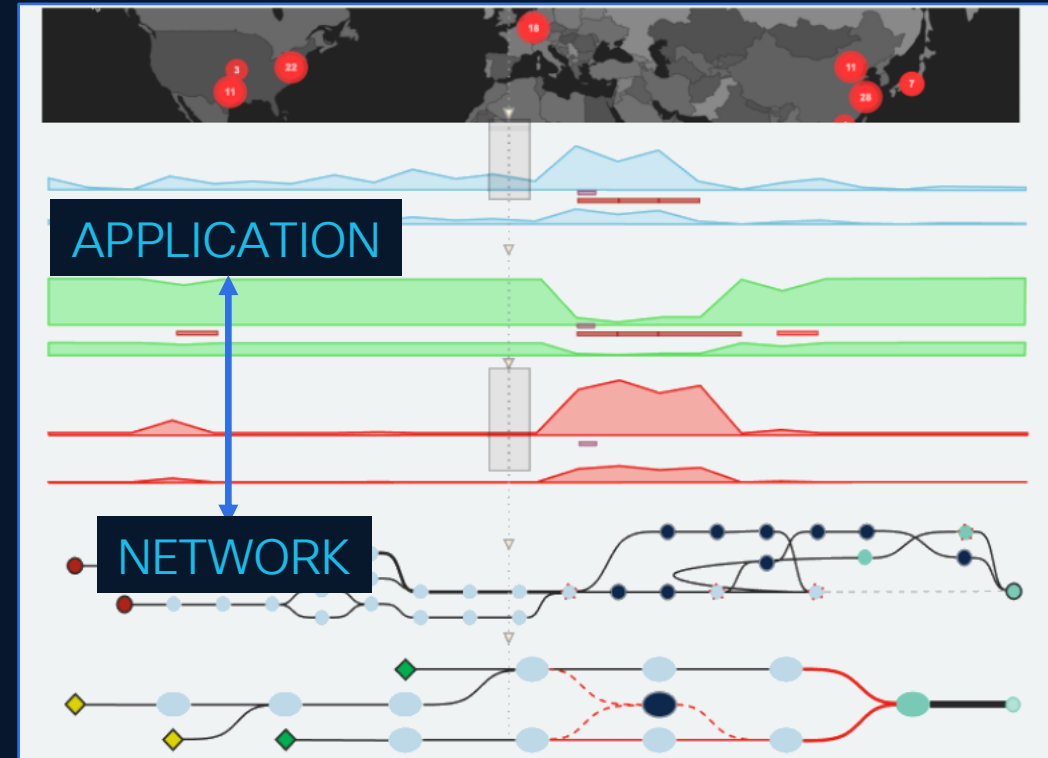


- Deployed in YOUR environment
- DCs, sites, offices, branches, stores...
- VMs, Servers, Containers, Cisco HW
- Inside-out, inside-inside
- Internal apps, SaaS, network
- Employee / network experience

ENDPOINT AGENT



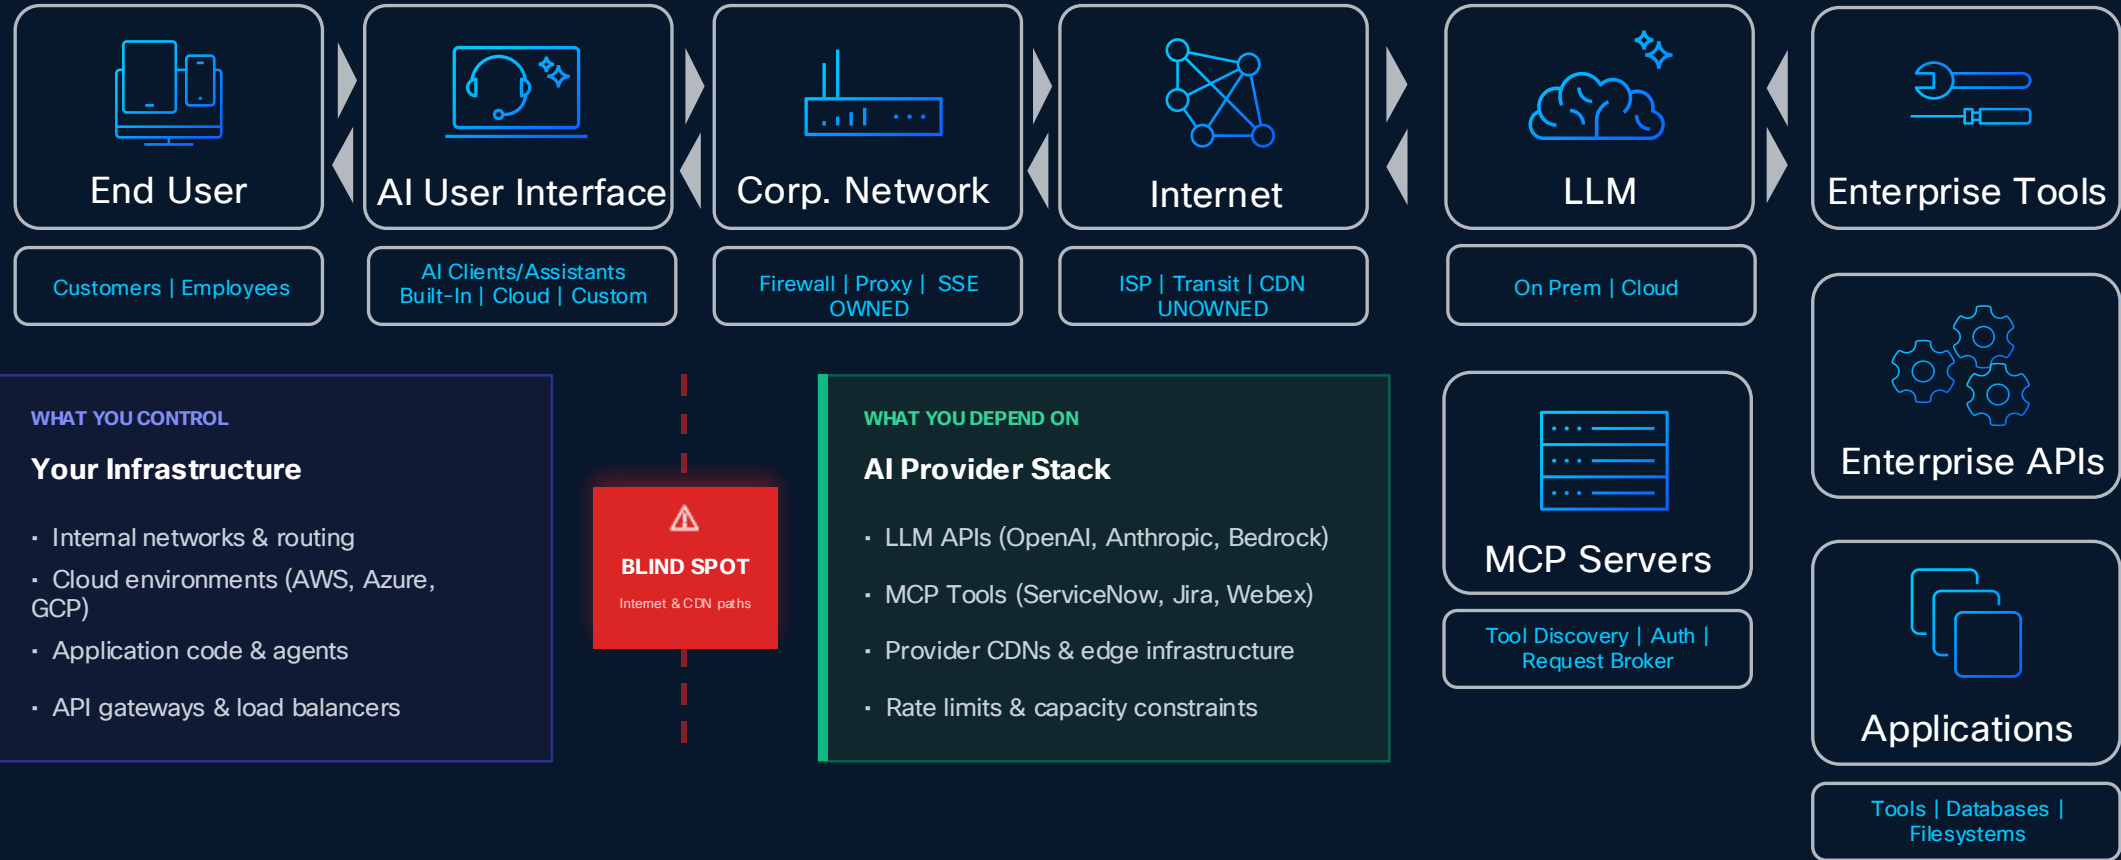
- Deployed on your employees' devices
- Home, office, anywhere...
- Laptops, RoomOS, Secure Access, Mobile
- Last mile visibility
- Internal/external apps, SaaS, network
- Wi-fi, VPN, ISP, any app



ThousandEyes and AI

Monitor and Feed Your AI Ecosystem

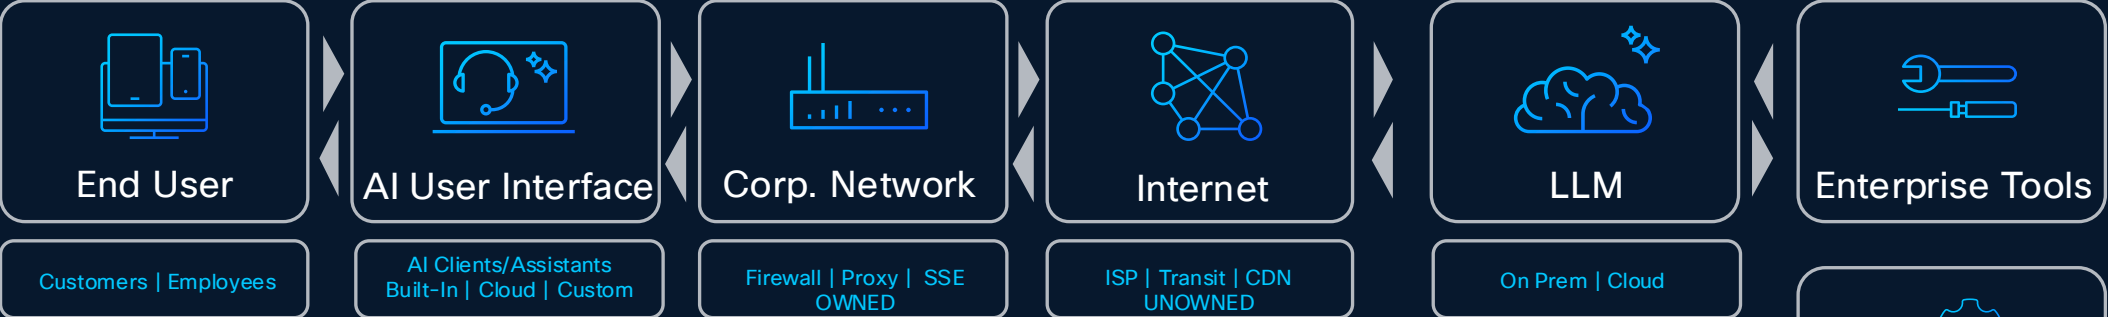
A New Generation of Complexity



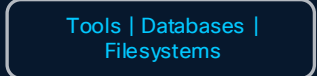
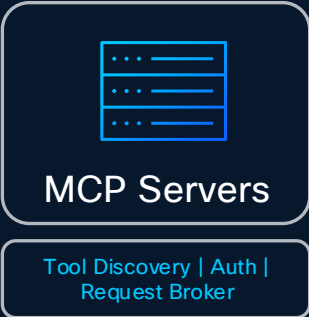
ThousandEyes and AI

Monitor and Feed Your AI Ecosystem

- CHALLENGE**
- Multiple models, systems and environments
 - On-prem and cloud based hosted components
 - Embedded AI Assistant reliability
 - Adoption, scale and security challenges
 - Many potential bottle necks across owned and unowned infrastructures



- PROMISE**
- Correlation of multiple systems without integration maintenance/upkeep
 - Rapid RCA and intelligent analysis
 - Dramatically reduce MTTI/MTTR
 - Advanced data correlation and visualization across platforms
 - Improve efficiency and satisfaction



ThousandEyes and AI

Monitor and Feed Your AI Ecosystem

Contribution



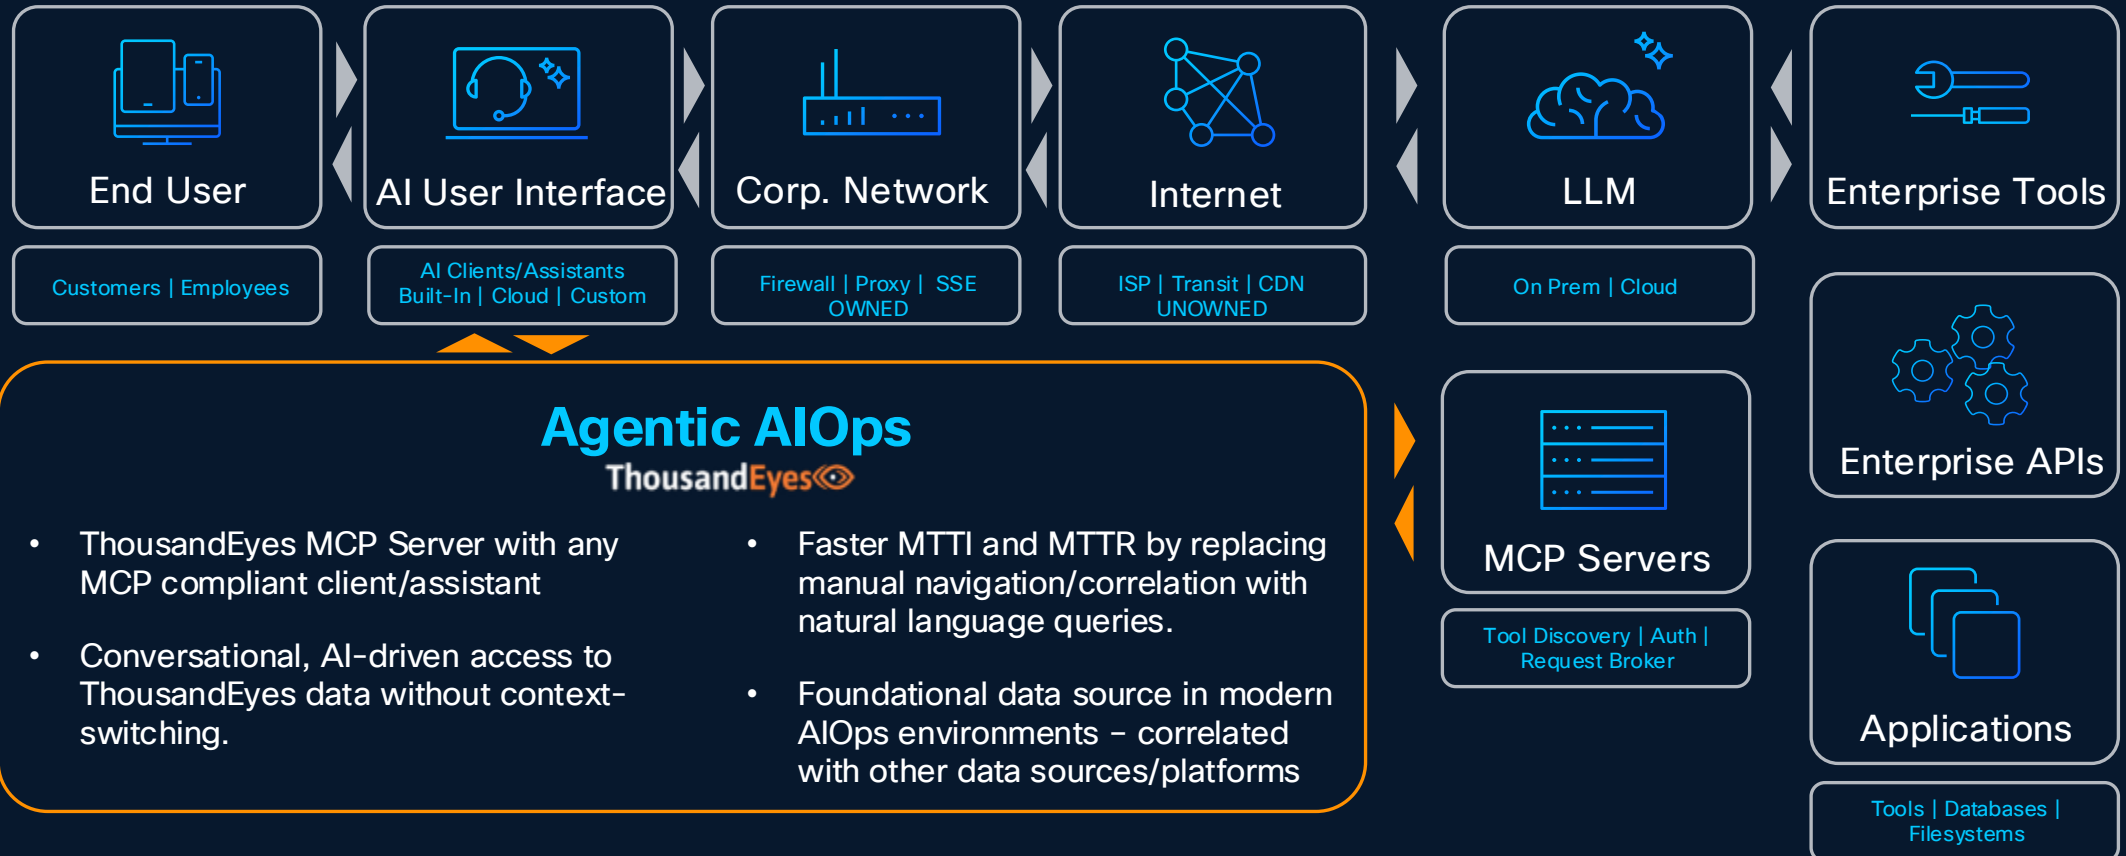
Assurance

- Visibility into latency, response times, token efficiency, and service availability of LLMs
- Inspect MCP resources validating the state of available tools and their configurations

ThousandEyes

Assurance for Agentic Apps

- Validate that responses maintain accuracy and consistency
- Assure connectivity across multiple owned and unowned environments



ThousandEyes MCP Server

Contributing to the App/Model

What is Agentic Operations

Autonomous AI systems that observe application, system and network conditions, make intelligent decisions, and take actions to optimize operations—without requiring human intervention for routine tasks.

Autonomous Monitoring

Continuously observes network performance, application health, and user experience across all digital infrastructure.

Intelligent Analysis

Uses AI to correlate data, identify root causes, predict issues, and understand context beyond simple threshold violations.

Contextual Decision-Making

Evaluates impact, urgency, and business context to prioritize actions and determine appropriate responses.

Automated Response

Takes corrective actions, creates tickets, notifies teams, and remediates issues while learning from outcomes.

ThousandEyes Role: Provides the comprehensive visibility that agents need to make informed decisions for assurance.

Why Agent Operations Matters

Business Impact

Speed to Resolution

Minutes instead of hours. Agents detect, analyze, and remediate issues faster than manual workflows.

Reduced Operational Costs

Automate tier-1 triage and routine troubleshooting, allowing engineers to focus on strategic initiatives.

24/7 Coverage

Continuous monitoring and response without staffing constraints or alert fatigue.

Technical Advantages for Operations

Correlation Across Data Sources

Combine ThousandEyes path visualization with BGP, DNS, and endpoint data for complete context.

Predictive Problem Detection

Identify degradation trends before they impact users or trigger SLA violations.

Intelligent Alert Filtering

Reduce noise by understanding which alerts represent true business impact vs. transient issues.

Continuous Learning

Improve accuracy over time by learning from past incidents and remediation outcomes.

Critical for Modern Hybrid Networks: As networks become more complex (cloud, SaaS, SD-WAN, multi-cloud), the volume and complexity of monitoring data exceeds human capacity. Agentic operations transform this data into actionable intelligence at machine speed.

ThousandEyes MCP Server

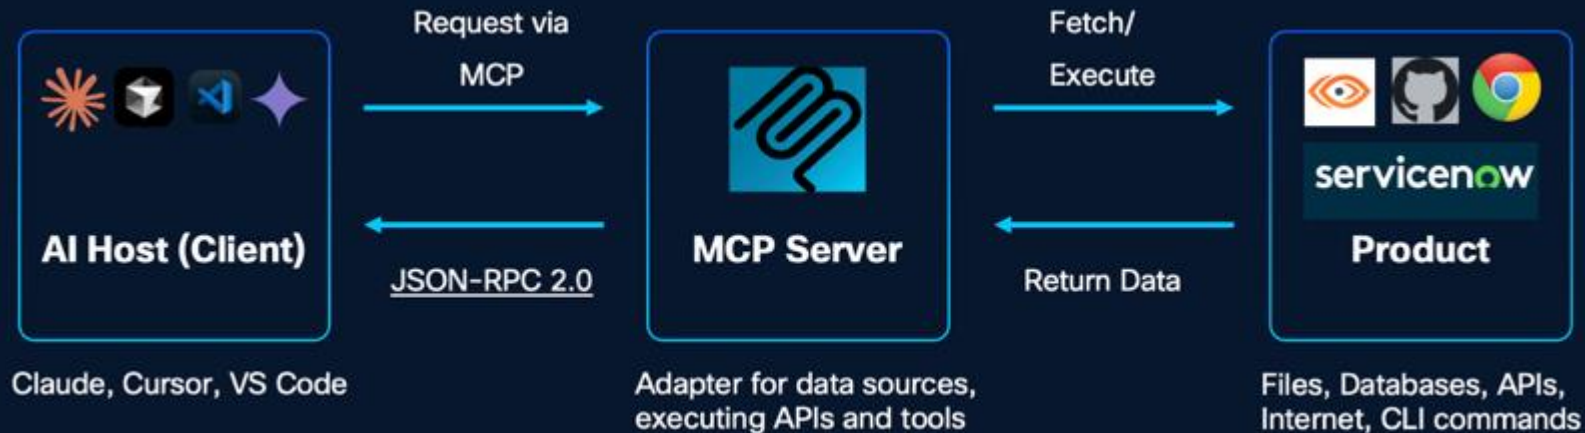
What is Model Context Procol (MCP)?

Lightweight protocol connecting AI agents to tools, APIs, databases, and filesystems

Secure, scalable client/server architecture for AI-powered workflows

Simplifies integration by providing a uniform interface to diverse infrastructure and operational tools

Critical component for customers building Agentic AI frameworks, tools and platforms



ThousandEyes MCP Server

Value

1

Conversational, AI-driven access to ThousandEyes data without context-switching.

2

Faster MTTI and MTTR by replacing manual navigation/correlation with natural language queries.

3

Positions ThousandEyes as a foundational data source in modern AIOps environments.

4

Expands user base from dedicated NetOps engineers to any technical stakeholder using an AI assistant.

Tools

Core Monitoring & Investigation

- List: Tests, Events, Alerts
- Get Details: Tests, Events, Alerts
- Search Outages

Network Path Analysis

- Get
 - Path Visualization
 - BGP Test Results
 - BGP Route Details

Endpoint Monitoring

- List
 - Endpoint Agents
- Get
 - Endpoint Agent Metrics

AI-Powered Skills

- ThousandEyes Assistance
- Troubleshooting
- Views Explanation

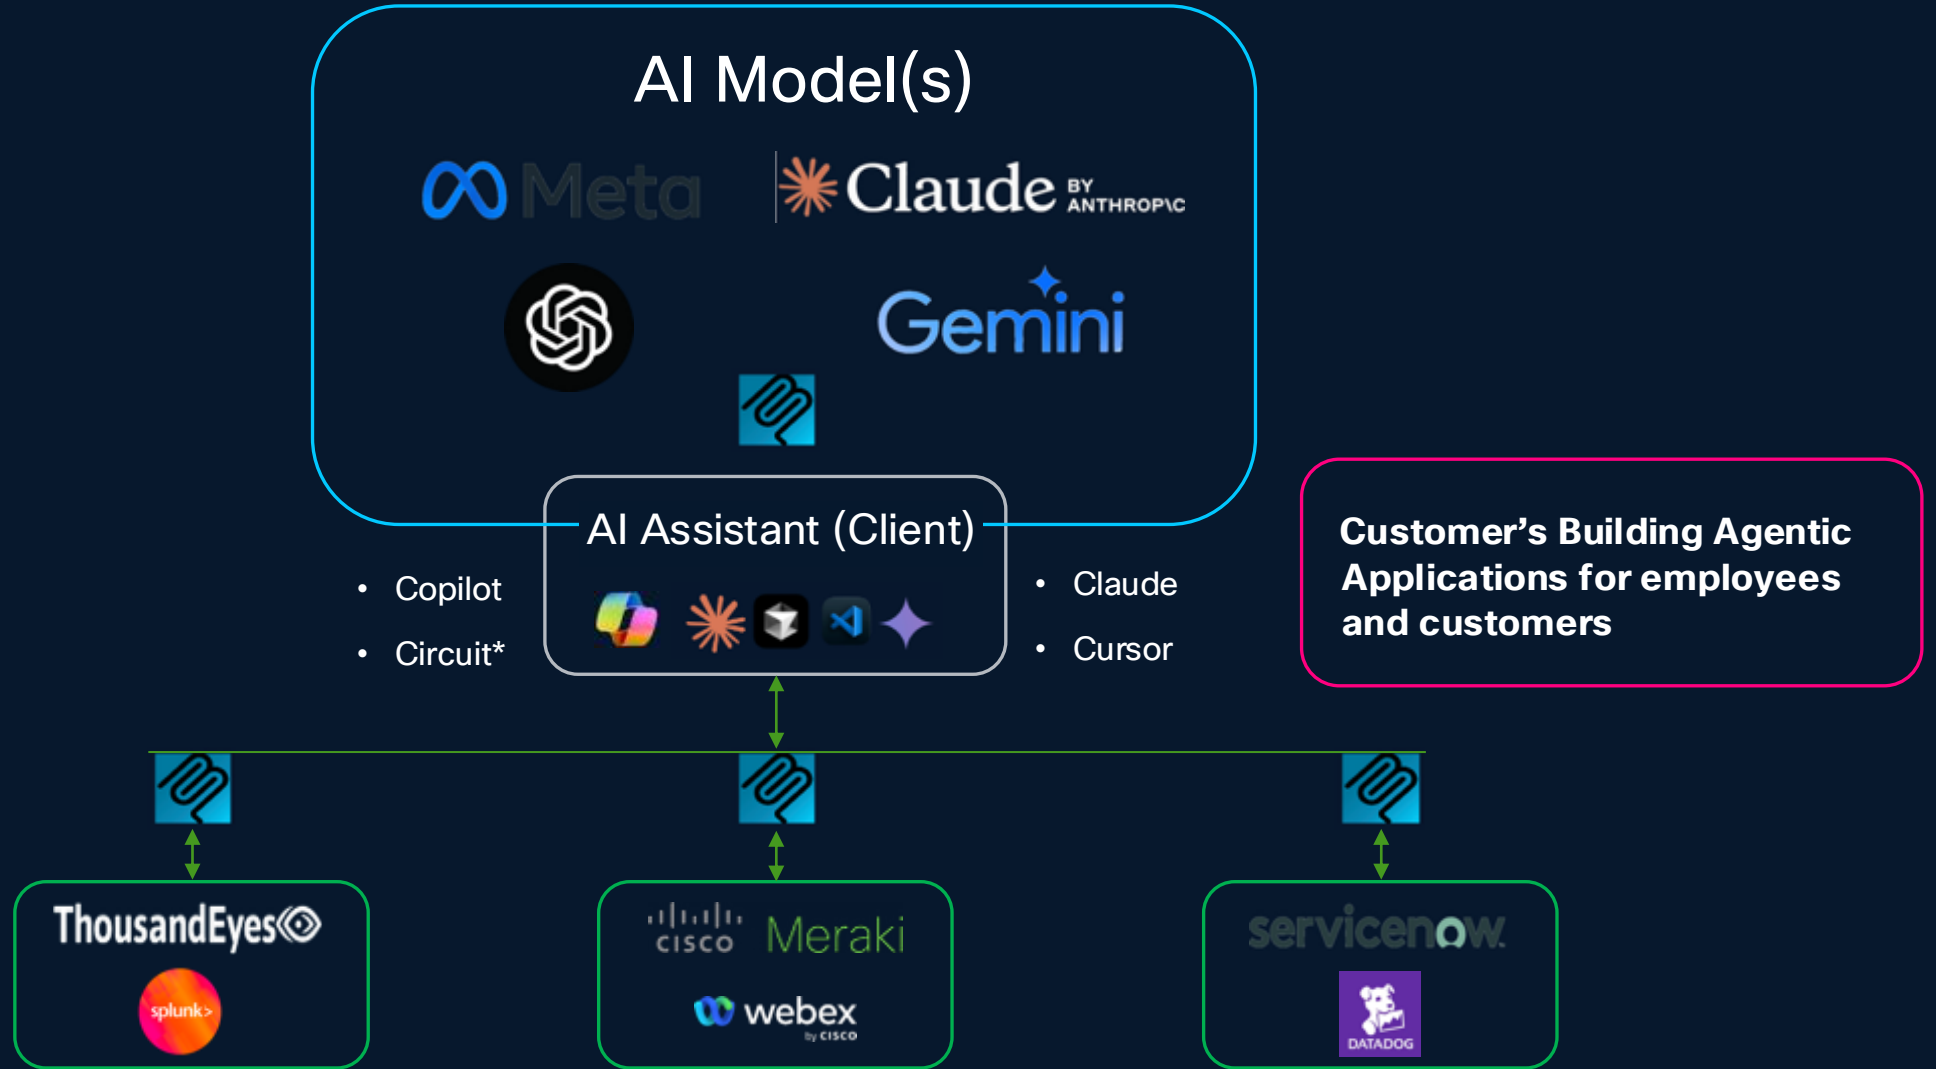
Account Management

- Get Account Groups

Advanced Analysis

- Get Anomalies

ThousandEyes MCP Server



ThousandEyes MCP Server

Current Customer Personas



ThousandEyes
Power User

- Leverage advanced data analysis, incident and RCA reporting
- They LOVE what MCP does for them
- **Interact with the data the way you want**



IT Operations

- Isn't a day-to-day ThousandEyes user
- Allows them to leverage the DATA – without needing to be in the UI
- **More value across the organization**



Leadership

- Thinking about agentic AI and how it can help their teams
- MCP allows us to be part of your strategy
- **Strategic Vision, QUICK AI WIN**

ThousandEyes MCP Server

What Does this Mean

1 Data is (still) Valuable

ThousandEyes uniquely collects data, now customers can leverage AI to correlate the data and derive conclusions much faster than ever before with little experience.

2 Integrations and APIs

“The decoupling of integration and correlation”

Say goodbye to “can you integrate with X”. Integration with a customer IT environment has always been a challenging goal. Data correlation and normalization across platforms can now be done fast and more accurately with AI will minimal effort from the customer.

3 Time to Value Decreases

The time it takes to arrive at conclusions using ThousandEyes data goes down significantly. The AI client can correlate the data and provide root cause and remediation recommendations.

4 Exposure/Adoption Increases

No longer will customers be required to understand our UI, or how to analyze the data to get value. This allows non-technical or general IT Ops personas to get value from ThousandEyes.

Agentic AIOps

ThousandEyes MCP Server



Incident Report

Root Cause Analysis

Alert ID: 9949709-WiFi-444-2024-08-14C-0001a

Test Name: Epic - ShareWeb Page/L1: Epic (ShareWeb app com)

Alert Started: 2024-02-09 04:18 UTC

Alert Duration: 03 minutes

Report Date: February 09, 2024 at 04:02 PM CST

Account ID: 542749913621986

Report Type: Root Cause Analysis

ThousandEyes Link: [https://cloud.thousandeyes.com/Incidents/Details/9949709-WiFi-444-2024-08-14C-0001a](#)

1. Executive Summary

A minor severity adaptive page-load alert triggered at 2024-02-09T04:18:02Z (04:18 PM CST) impacting the Epic ShareWeb test (https://shareweb.epic.com) across 3 geographic distributed cloud agents. Users experienced page load timeouts with the entire pageLoad() method true and pageLoad() false preventing successful page completion. The incident self-resolved at 2024-02-09T04:21:00Z (04:21 PM CST) with a total duration of 3 minutes. Root cause: Simultaneous page load timeout failures at multiple US locations due to server-side responsiveness degradation at the Epic application by an underlying network path issues detected. This alert represents a recurring pattern occurring at approximately the same time.

2. Incident Timeline

- 2024-02-09T04:18:02Z (04:18 PM CST): Baseline period: 0 violations out of 11 agents (0% probability: 0.25%)
- 2024-02-09T04:19:00Z (04:19 PM CST): Violation count increased: 2 violations out of 11 agents (probability: 20%)
- 2024-02-09T04:21:00Z (04:21 PM CST): Alert triggered: 3 agents violated threshold, severity increased: 3 violations out of 11 agents (probability: 25%)

Done. I've created the incident:

Incident Number: INC0080991

Field	Value
Short Description	Investigate wireless connectivity issues - MGC1DC7QFYWY on Mustang network
Category	Network
Subcategory	Wireless
Priority	1 - Critical
Urgency	1 - High

The description includes all the ThousandEyes findings: the poor RSSI readings (-67 to -72 dBm), low signal quality (54-66%), elevated roaming events indicating AP hunting behavior,

ThousandEyes Global Demos NETWORK & APPLICATION ASSURANCE **LIVE** Auto-refresh every 15 min Data as of 3/24/26, 10:03 AM Refresh

99.7% SCORE

147 MONITORED SERVICES

58 INTERNET OUTAGES

ASSURANCE SCORE
Composite availability across all active synthetic tests in the selected window. Use business service tiles below to prioritize follow-up.

OVERALL AVAILABILITY: 99.97% (vs 24h avg)

AVG RESPONSE TIME: 148ms (P95: 312ms)

PACKET LOSS: 0.02% (12 affected paths)

ACTIVE ALERTS: 7 (3 critical, 4 warning)

ACTIVE AGENTS: 342 (CE: 48, EP: 294)

INTERNET OUTAGES: 3 (2 providers affected)

GLOBAL FACILITY HEALTH - THOUSANDEYES MONITORING (22 / 24 Healthy)

Facility Drill-Down - Arlington, VA

Availability	Response Time	Packet Loss	Latency
99.96%	67ms	0%	2.3ms

Endpoint Agent Overview (294 agents active)

Agent	Latency	Signal	CPU	Issues
NYC-FIN-Trader-01	14ms	-52 dBm	44%	4
NYC-FIN-Trader-02	18ms	-	-	-
STL-ENG-Dev-14	23ms	-	-	-
LON-OPS-Analyst-07	112ms	-	-	-
MUM-DEV-QA-03	245ms	-	-	-
SYD-PM-Mgr-02	188ms	-	-	-

Microsoft Office 365 Performance Dashboard

Real-time monitoring powered by ThousandEyes | Account Group: 100376

CURRENT LATENCY: 47.9 ms | **AVERAGE LOSS: 0.31 %** | **AVAILABILITY: 99.97 %** | **RESPONSE TIME: 144 ms**

Network Latency Timeline

Packet Loss Timeline

Path Trace - Office 365

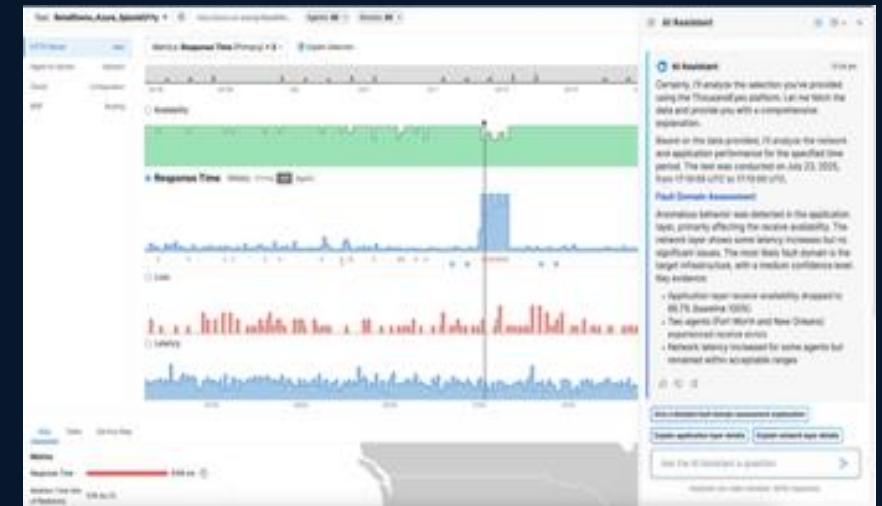
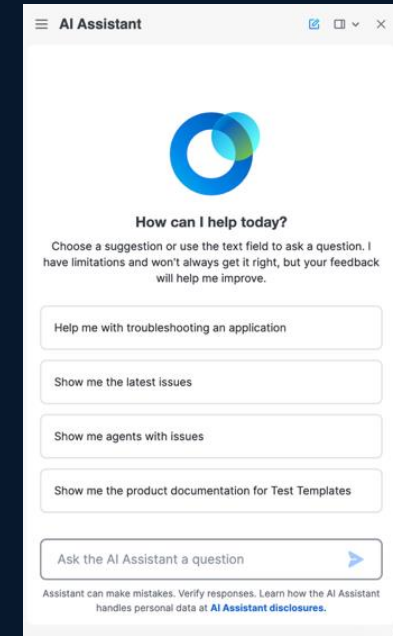
Step	Latency	Loss	Signal
1	2 ms	0.0%	-
2	12.0ms	0.0%	-
3	18.0ms	0.0%	-
4	26.0ms	0.0%	-
5	26.0ms	0.0%	-
6	26.0ms	0.0%	-
7	26.0ms	0.0%	-
8	30.0ms	0.0%	-



AI-Driven Intelligence

Low Touch, No Touch Assurance

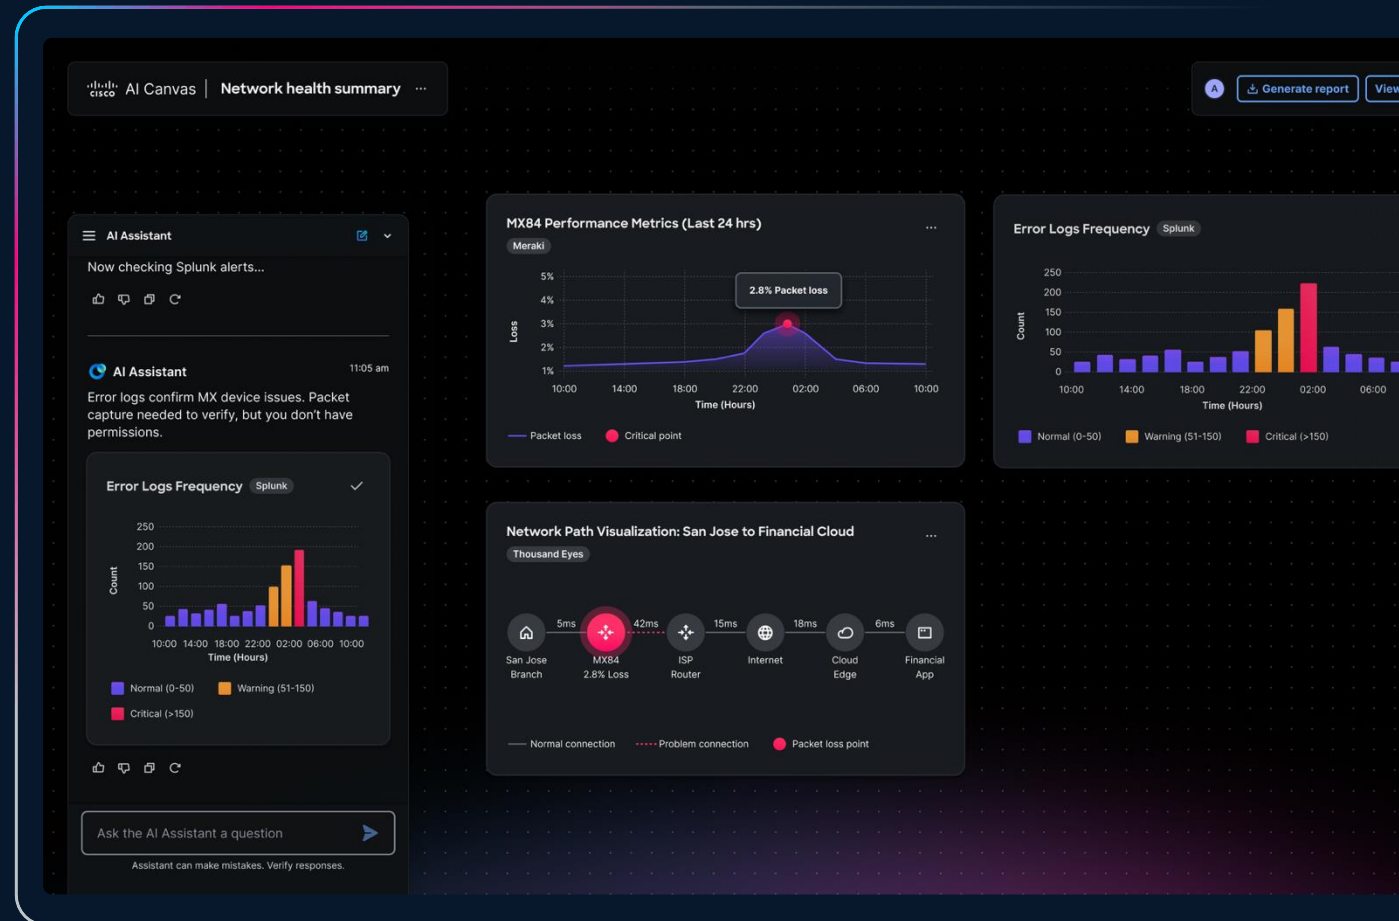
- Dynamics Baselines
- Views Explanations
 - Automatically identifies fault domain reducing the time to identify root cause
- Predictive Analytics & Events
 - Forecasts potential disruptions, allowing teams to act before users are impacted
- Cisco AI Assistant integrated into ThousandEyes
 - Summarizes, explains, and provides troubleshooting steps for network issues using natural language



Agentic AIOps

AI Canvas

- Single canvas for cross domain troubleshooting
- Generative UI dashboard with AI reasoning built-in
- Support data from Meraki, ThousandEyes, Splunk, and more



THOUSANDEYES · INCIDENT REPORT RESOLVED

Microsoft Teams Service Outage Event ID 4ac8bd67-73b5-4f4f-b35b-ee0c248e4440

March 3, 2026 07:30 – 07:44 UTC (~15 min) Global HIGH SEVERITY

Summary

On March 3, 2026, Microsoft Teams experienced a global service disruption caused by **HTTP 5xx server errors** originating from Microsoft's infrastructure. ThousandEyes detected the incident at **07:30 UTC** across 3 monitored tests and 55 agents spanning 6 continents. Service was fully restored by **07:44 UTC**. No internal network or DNS issues were identified — this was entirely a third-party provider failure.

KEY METRICS

DURATION 15m 07:30 – 07:44 UTC	AFFECTED AGENTS 55 Cloud & Enterprise	LOCATIONS HIT 419 Provider-wide	TESTS TRIGGERED 3 HTTP Server tests	SERVERS AFFECTED 745 Provider-wide	TRANSIT HOPS AVG 14 2–23 hops observed
--	---	---	---	--	--

INCIDENT TIMELINE

[file:///Users/bdonoghru/Downloads/Microsoft%20Teams%20Outage%20v2\(2\).html](file:///Users/bdonoghru/Downloads/Microsoft%20Teams%20Outage%20v2(2).html)

Assurance for Agentic Applications

Beyond Availability

The Visibility Gap

AI applications span your infrastructure and external providers. Current tools only see half the picture.

WHAT YOU CONTROL

Your Infrastructure

- Internal networks & routing
- Cloud environments (AWS, Azure, GCP)
- Application code & agents
- API gateways & load balancers



BLIND SPOT
Internet & CDN
paths

WHAT YOU DEPEND ON

AI Provider Stack

- LLM APIs (OpenAI, Anthropic, Bedrock)
- MCP Tools (ServiceNow, Jira, Webex)
- Provider CDNs & edge infrastructure
- Rate limits & capacity constraints

APM Tools

PARTIAL

See application traces but blind to network path issues

Visibility:

App traces Code spans **Network path** **Internet hops**

LLM Observability

PARTIAL

Track tokens and prompts but assume the network works

Visibility:

Tokens Latency **Root cause** **Path issues**

Provider Dashboards

PARTIAL

Show usage stats but no diagnostic capability

Visibility:

Usage stats Status page **Your path** **Diagnostics**

When your AI agent fails or slows down, where do you look?

API gateway? Provider infrastructure? Internet routing? CDN issues? Regional capacity?

Key Monitoring Considerations for AI

See across the entire AI delivery path—from your infrastructure to every provider endpoint.



LLM APIs

- OpenAI, Anthropic, Bedrock
- Response latency & availability
- Token throughput monitoring
- Regional endpoint health



MCP Servers

- ServiceNow, Jira, Webex
- Tool call latency
- Integration availability
- Error rate tracking



Vector DBs

- Pinecone, Weaviate, Chroma
- Query response times
- Embedding retrieval
- Database connectivity



Network Path

- Internet routing analysis
- CDN performance
- Path visualization
- Hop-by-hop latency

WHY THOUSANDEYES



Vantage Points

Test from 400+ global cloud agents + your enterprise agents



Path Visualization

See every network hop between you and AI providers



Proactive Detection

Identify issues before users experience them

Three Gaps Nobody Talks About

Where AI monitoring falls short today



Availability Without Quality

Your pipeline returns 200 OK, but the answer is fabricated. Uptime metrics say everything is fine. Your customer says otherwise.



Latency Without Context

A 4-second response looks normal. But if one pipeline stage consumed 3.8 of those seconds, you have a ticking time bomb. No tool surfaces this.



Observability Without Assurance

Token-level telemetry tells you what the model did. It doesn't tell you whether the answer was correct, grounded, or safe.

Your Version of This

Three questions to sit with



How many stages are in your AI pipeline?

Most leaders can't answer this. If you can – how many of those stages are monitored?



When was the last time you validated output quality?

Not model accuracy on a benchmark. Your production AI, answering real questions, today.

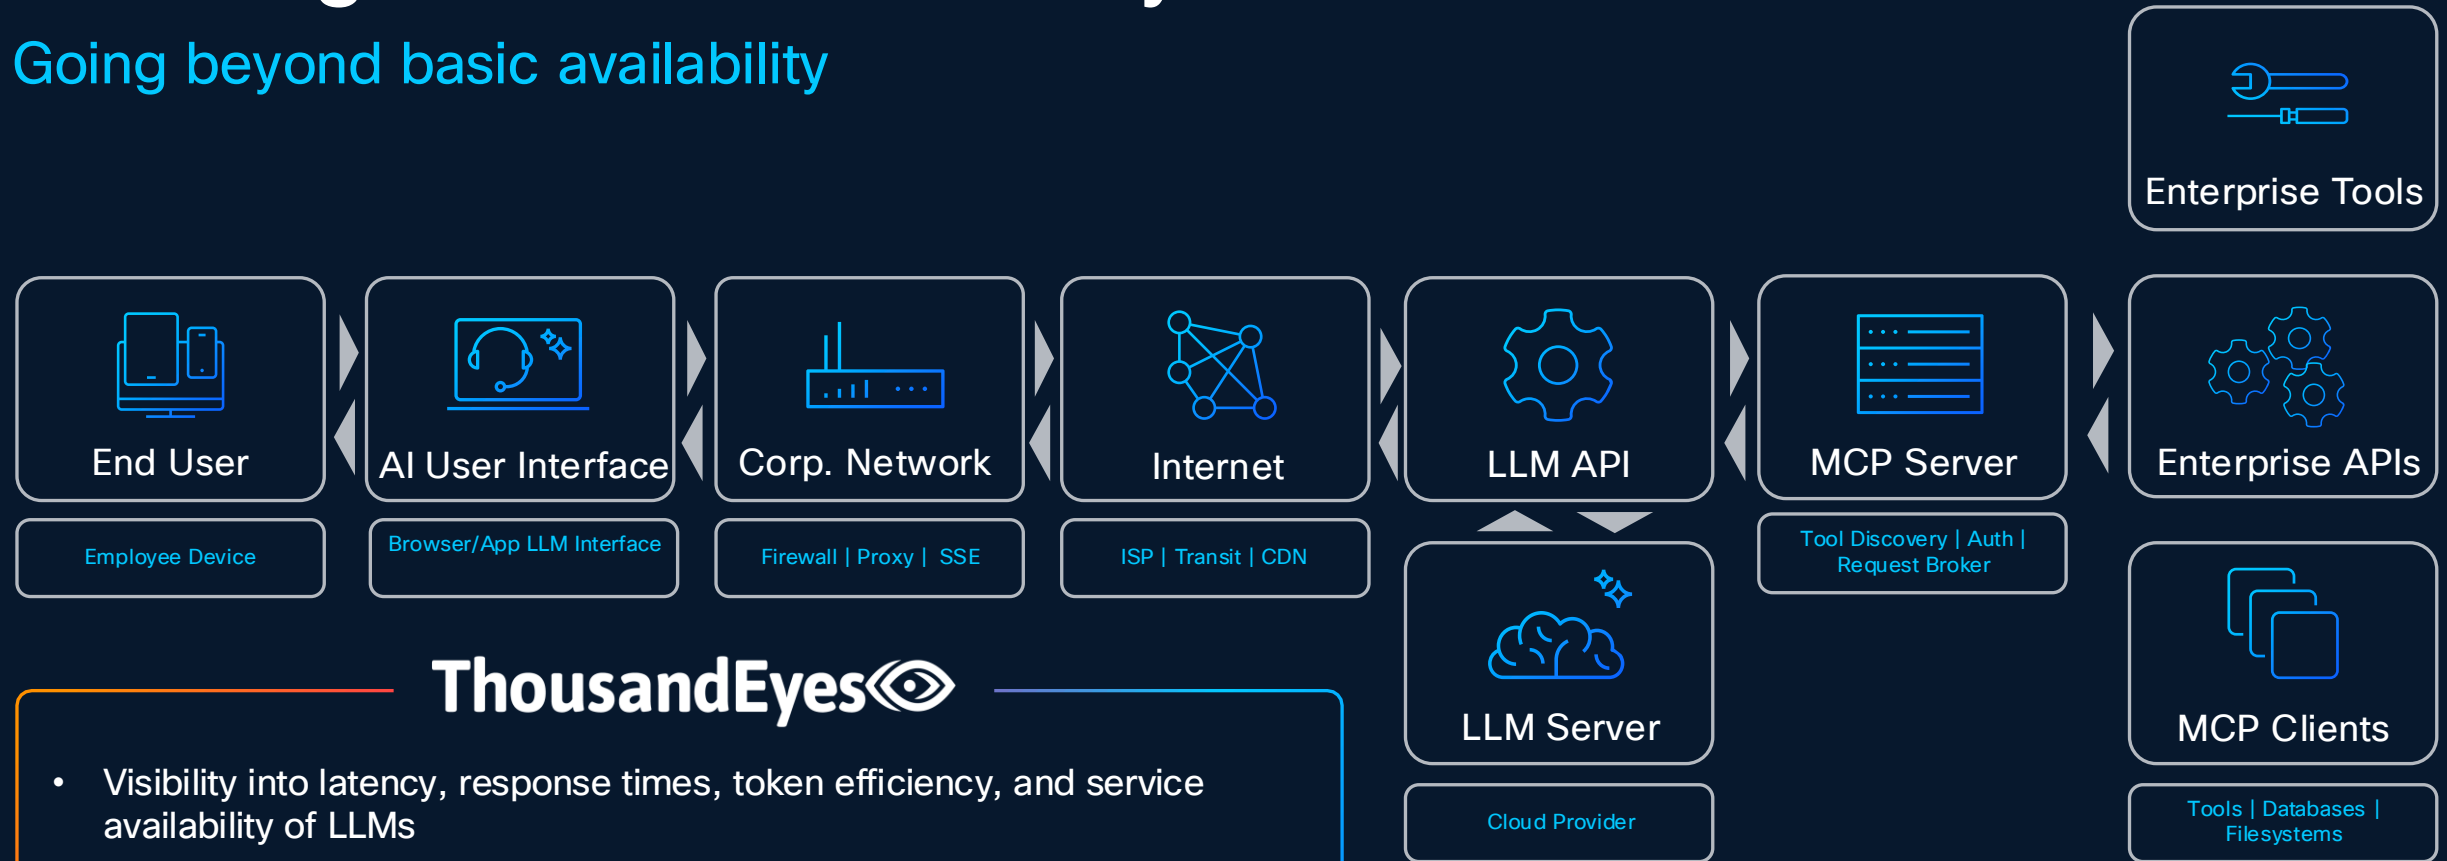


If your AI hallucinated right now – how long until you'd know?

Not 'the model went down.' A confident, well-formatted, completely wrong answer. How many hours? Days?

Assuring AI with ThousandEyes

Going beyond basic availability



ThousandEyes

- Visibility into latency, response times, token efficiency, and service availability of LLMs
- Inspect MCP resources validating the state of available tools and their configurations
- Validate that responses maintain accuracy and consistency
- Comparing performance of multiple models
- Assure connectivity across multiple owned and unowned environments

End-to-End Test Architecture

1

API/Transaction Test

Application Logic

- Custom token utilization checks
- Rate limit threshold alerts (80/90%)
- API response validation
- Login -> Token -> Prompt flow
- Vector Database Monitoring

2

HTTP Server Test

Connection Layer

- DNS resolution time
- TCP connect latency
- SSL/TLS handshake
- Time to first byte (TTFB)

3

Network Test

Path Analysis

- Hop-by-hop visualization
- Packet loss per segment
- Latency breakdown
- ISP/provider identification

4

BGP Monitoring

Routing Layer

- Route hijack detection
- Path change alerts
- Upstream provider issues
- Traffic blackhole prevention

We Can Provide AI Assurance TODAY

The screenshot displays the Cisco ThousandEyes Network & App Synthetics dashboard. The left sidebar contains navigation options: Dashboards, Event Detection, Alerts (8), Network & App Synthetics (selected), Endpoint Experience, Routing, Traffic Insights, Devices, Cloud Insights, Internet Insights, and Manage. The main content area is titled "Start with ThousandEyes based recommendations" and includes a "View all recommendations" link. Below this, there is a section for "Associated Service Recommendations" with a "Monitored 0% (0 of 15)" status. The "Start with templates" section features a grid of monitoring templates for various AI services: Anthropic, ChatGPT, OpenAI, Google Cloud Vision AI, Google Gemini, Azure AI Foundry, AWS Bedrock, and Custom MCP Server. Each template card includes the service logo, name, and a brief description of the template's purpose.



Agent to Server

Network

BGP

Routing

Metrics: Latency

Explain Selection

Wed, Mar 4 02:00 - 02:05 UTC (1 day ago)

24h 7d 14d



Latency Overlay Off Agents



Path Visualization

Map

Table

Show

Agents: All

IP Address Labels

3 hops

3 hops

Group

Agents by Agent

Interfaces by IP Address

Destinations by No Grouping

Highlight

All

Search Network, Country, IP Address, Prefix or Title


0 matches

Forwarding Loss > 10% (0 Nodes)

Link Delay (Avg) > 100ms (0 Links)

Select

Click a node or link



Assurance for Agentic Apps

The Power of Cisco

ThousandEyes Assurance for Agentic Apps is essential as Cisco helps you build your AI Ecosystem



AI Networking

Deliver AI everywhere with intelligent networking purpose-built for high performance, simplified operations, and security at scale.



Secure AI Factory

Accelerate AI deployment with a modular and validated design that combines infrastructure, security, and observability.



AI Defense

Protect AI development and usage with visibility, guardrails and real-time threat intelligence designed for emerging AI-specific risks.

“AI won’t replace humans. But humans who use AI will replace those who don’t.” – *Sam Altman*

Thank you

