IS YOUR NETWORK AN AI ASSET OR A LIABILITY?

Who here already has an AI intern somewhere in the company?

And do you know what it can access?

cisco

AI WILL NOT REPLACE EMPLOYEES.
EMPLOYEES WHO USE AI WELL
WILL REPLACE EMPLOYEES WHO DON'T.

AI Isn't Cheaper Than Labor—Yet.

The near-term economics of enterprise AI are more about productivity than labor replacement.

NVIDIA: Compute > Employees

Bryan Catanzaro said for his team that 'the cost of compute is far beyond the costs of the employees.'

Only 23%

MIT found only 23% of studied vision-related jobs were economically viable for AI automation at current costs.

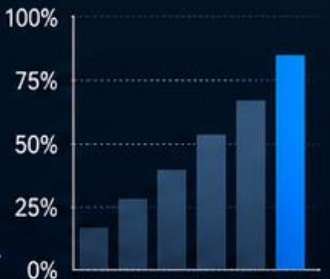

$740B in 2026

Big Tech is projected to spend about $740 billion on AI in 2026, underscoring how infrastructure-heavy the AI buildout still is.

92,000+ Layoffs

Layoffs have topped 92,000 across nearly 100 tech companies this year, yet AI still has not become a dependable labor-cost substitute.

Takeaway: AI is a productivity amplifier today, not yet a reliable labor-cost replacement.

AI Changes the Economics of Engineering

Jensen Huang on AI and engineer productivity

> If an engineer costs you **$500K** and they don't use at least **$250K in tokens**, I need to have a conversation with that employee."

— Jensen Huang
CEO, NVIDIA

Source: NVIDIA CEO Jensen Huang, various public interviews, 2024

THE NEW PRODUCTIVITY STANDARD

Engineer fully burdened cost (example)

$500K

Minimum expected AI investment

$250K

in tokens (annual)

WHAT IT ENABLES

- Faster delivery
- Higher quality
- More innovation
- Greater leverage

THE SHIFT

TRADITIONAL WORLD	AI-NATIVE WORLD
People scale work	AI scales people

AI is no longer optional for engineers. It's the cost of staying competitive.

- High performers use AI every day.
- AI fluency is the new baseline.
- The advantage compounds.

Productivity vs. Cost: The AI Capacity Balancing Act

Give top performers the runway. Meter the burn.

PRODUCTIVITY LIFT

Super employees.
Real impact.

- Faster delivery
- Better decisions
- Higher output
- More customer impact

GOVERNED AI CAPACITY

The balancing point.

- User tiers & roles
- Approved tools & models
- Budget guardrails
- Usage visibility
- Outcome tracking

COST + RISK

Uncontrolled burn.
Real risk.

- Token burn
- Latency & performance
- Data exposure
- Low-value automation

Are we buying productivity — or just buying more compute?

Measure Value | Optimize Allocation | Enforce Guardrails

Fund the work that compounds. Control the usage that just burns.

8 BILLION TODAY. 80 BILLION SOON?

The world isn't just getting bigger. It's getting agentic.

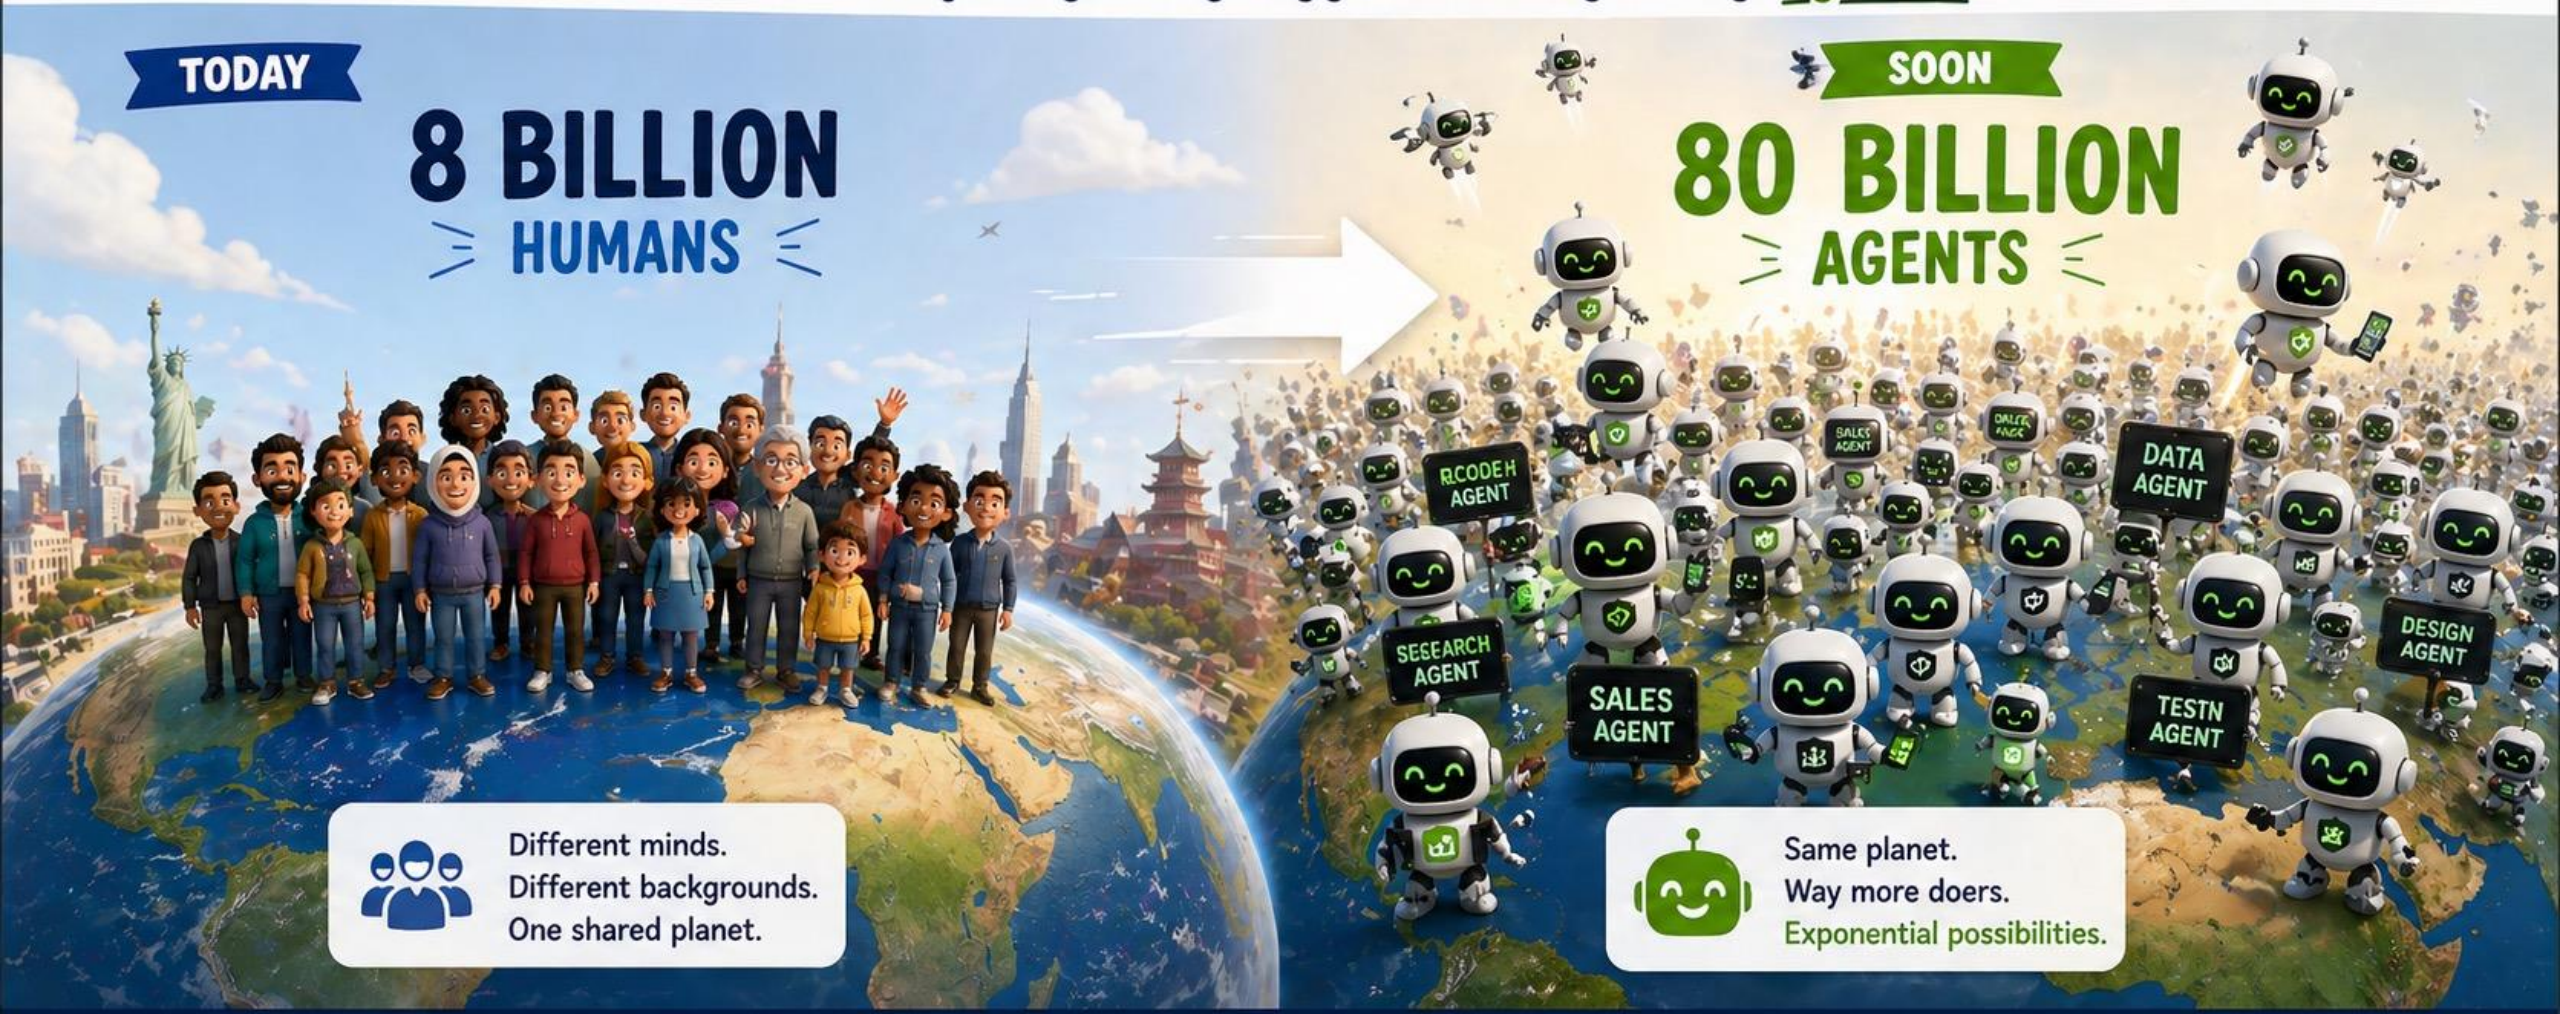

MORE MINDS. MORE ACTION. MORE IMPACT. LET'S BUILD A WORLD THAT'S READY.

Your Next Customer May Be an Agent

Start preparing your digital front door for machine customers.

Verify the Visitor
Know who (or what) is knocking. Trusted agent identity matters.

Scope the Access
Agents should only do what they're allowed to do.

Protect the Transaction
Secure actions, payments, and sensitive data.

Observe Everything
See agent traffic, behavior and outcomes in real time.

Seamless Human Handoff
Make it easy to go from agent → human and back.

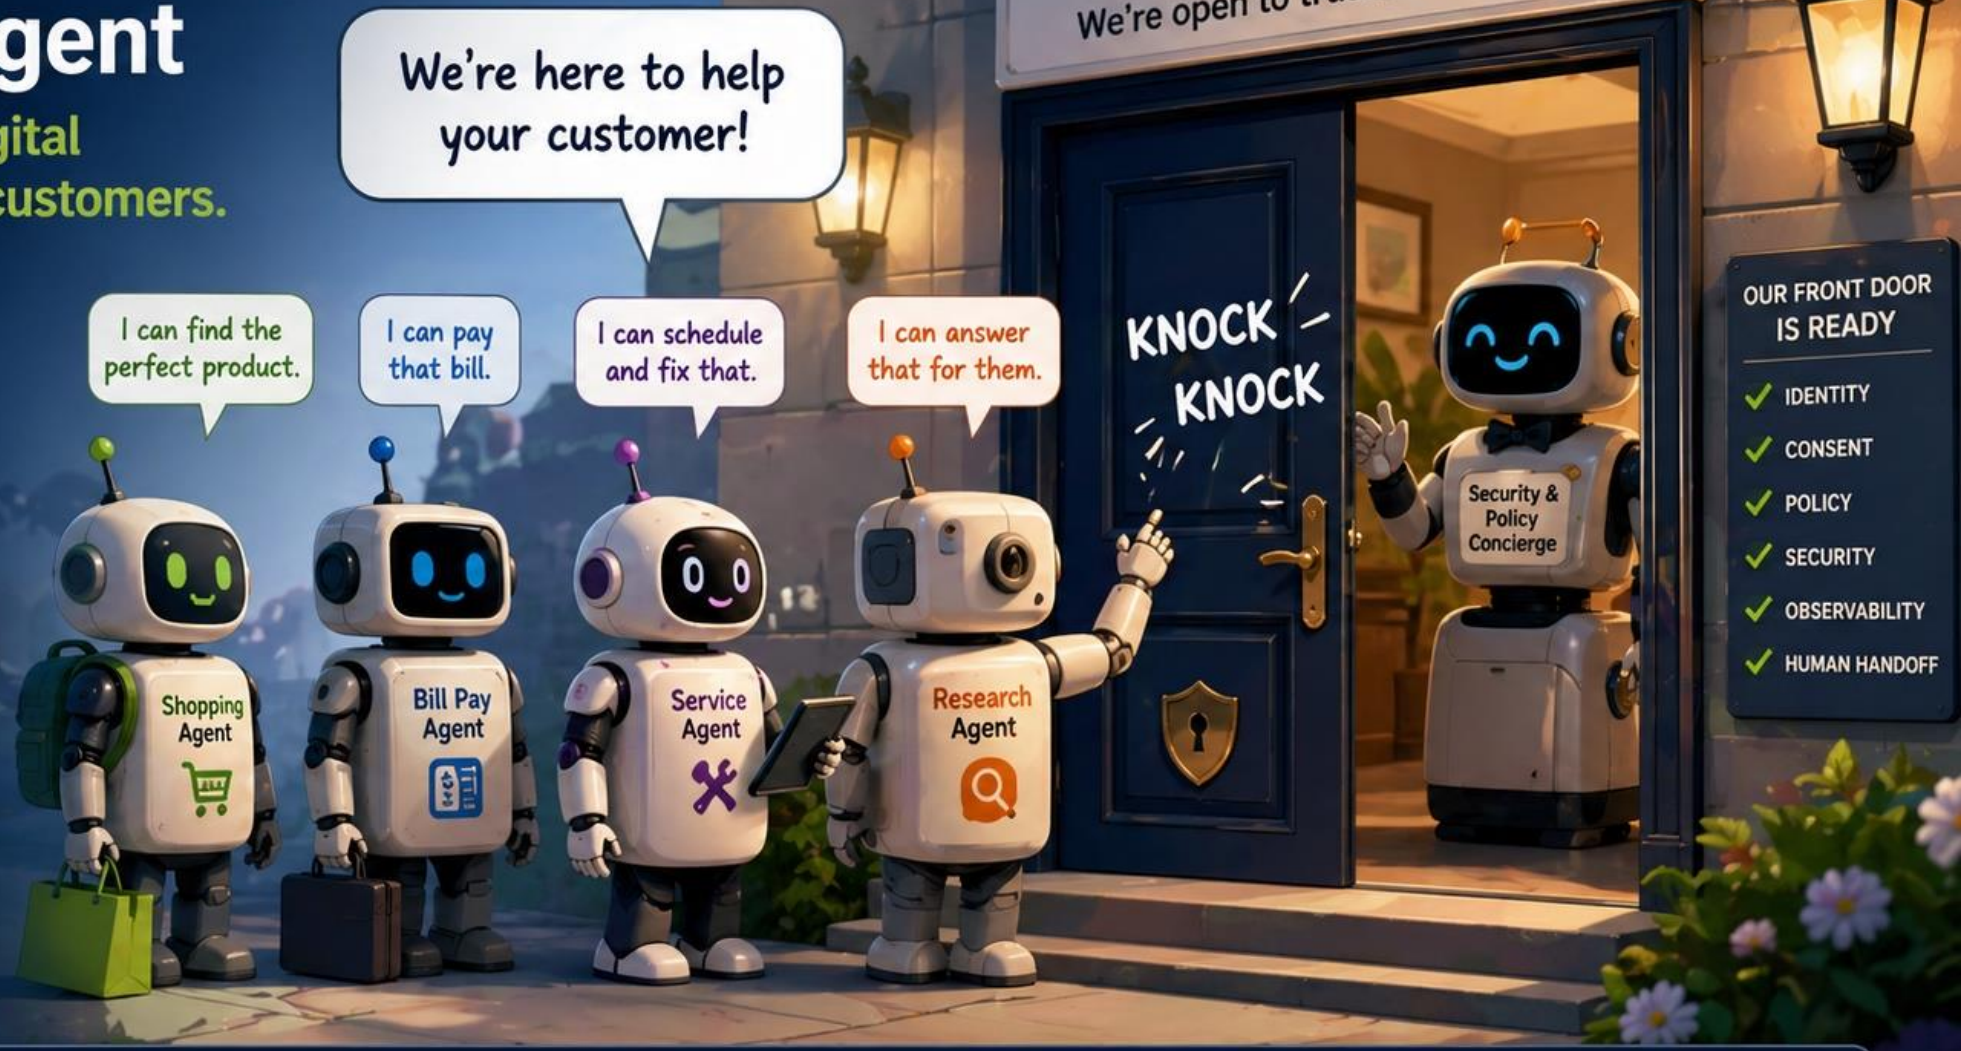

The experience layer is expanding.
Be ready for humans, bots, and trusted agents.

Human → AI Assistant → Agent → Your Business

IT WORKS... UNTIL IT DOESN'T.

Low load. Happy days.
Plenty of capacity.

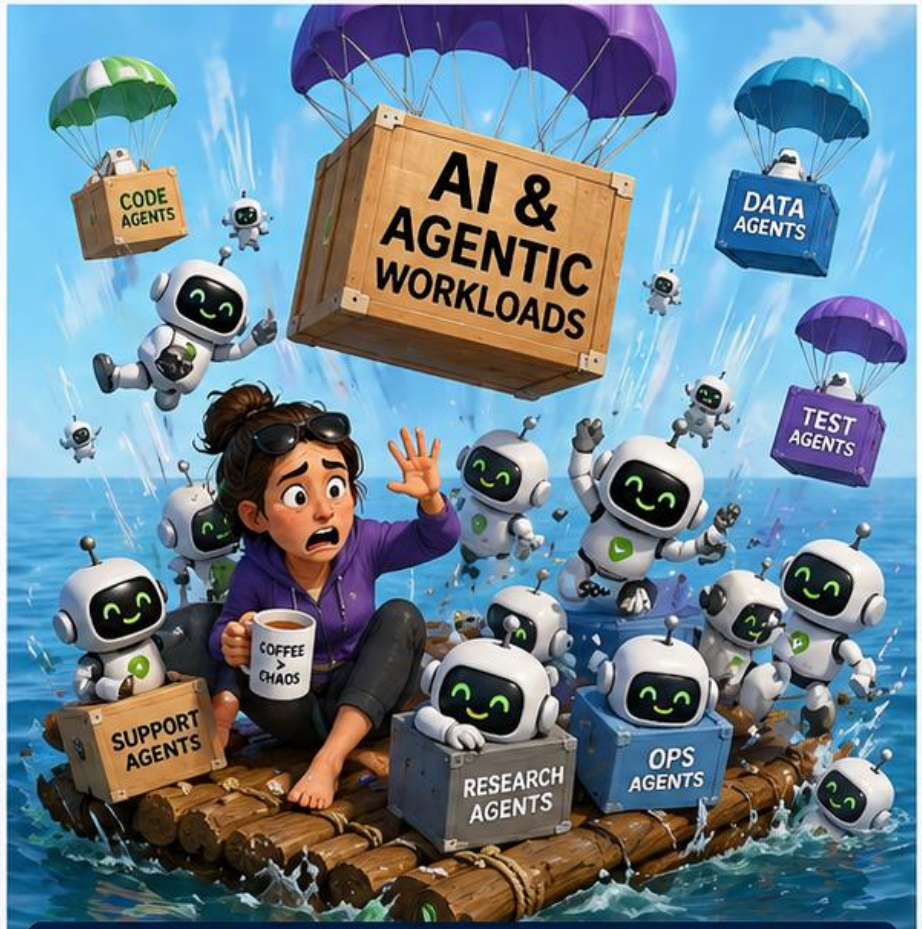

AI and agentic workloads
drop in fast. Demand explodes.

Overwhelmed. Unstable. Unpredictable.
The raft was never built for this.

Plan for the future load, or be ready for the fall.

SPEED vs CONTROL

You need **both** to win.

ONLY SPEED?

An unguided missile.

- Fast
- Unpredictable
- High risk
- Little to no guardrails

SPEED + CONTROL?

Directed. Focused. Impactful.

- Fast with purpose
- Guardrails and governance
- Lower risk, higher trust
- Real, measurable outcomes

ONLY CONTROL?

A museum.

- Safe
- Slow
- Stagnant
- Innovation stays behind glass

SPEED GETS YOU MOVING. | **CONTROL KEEPS YOU ON MISSION.** | **MOVE FAST. STAY ON COURSE. CREATE VALUE.**

THE AI INFLECTION POINT

CHATBOTS

AGENTIC AI

PHYSICAL AI

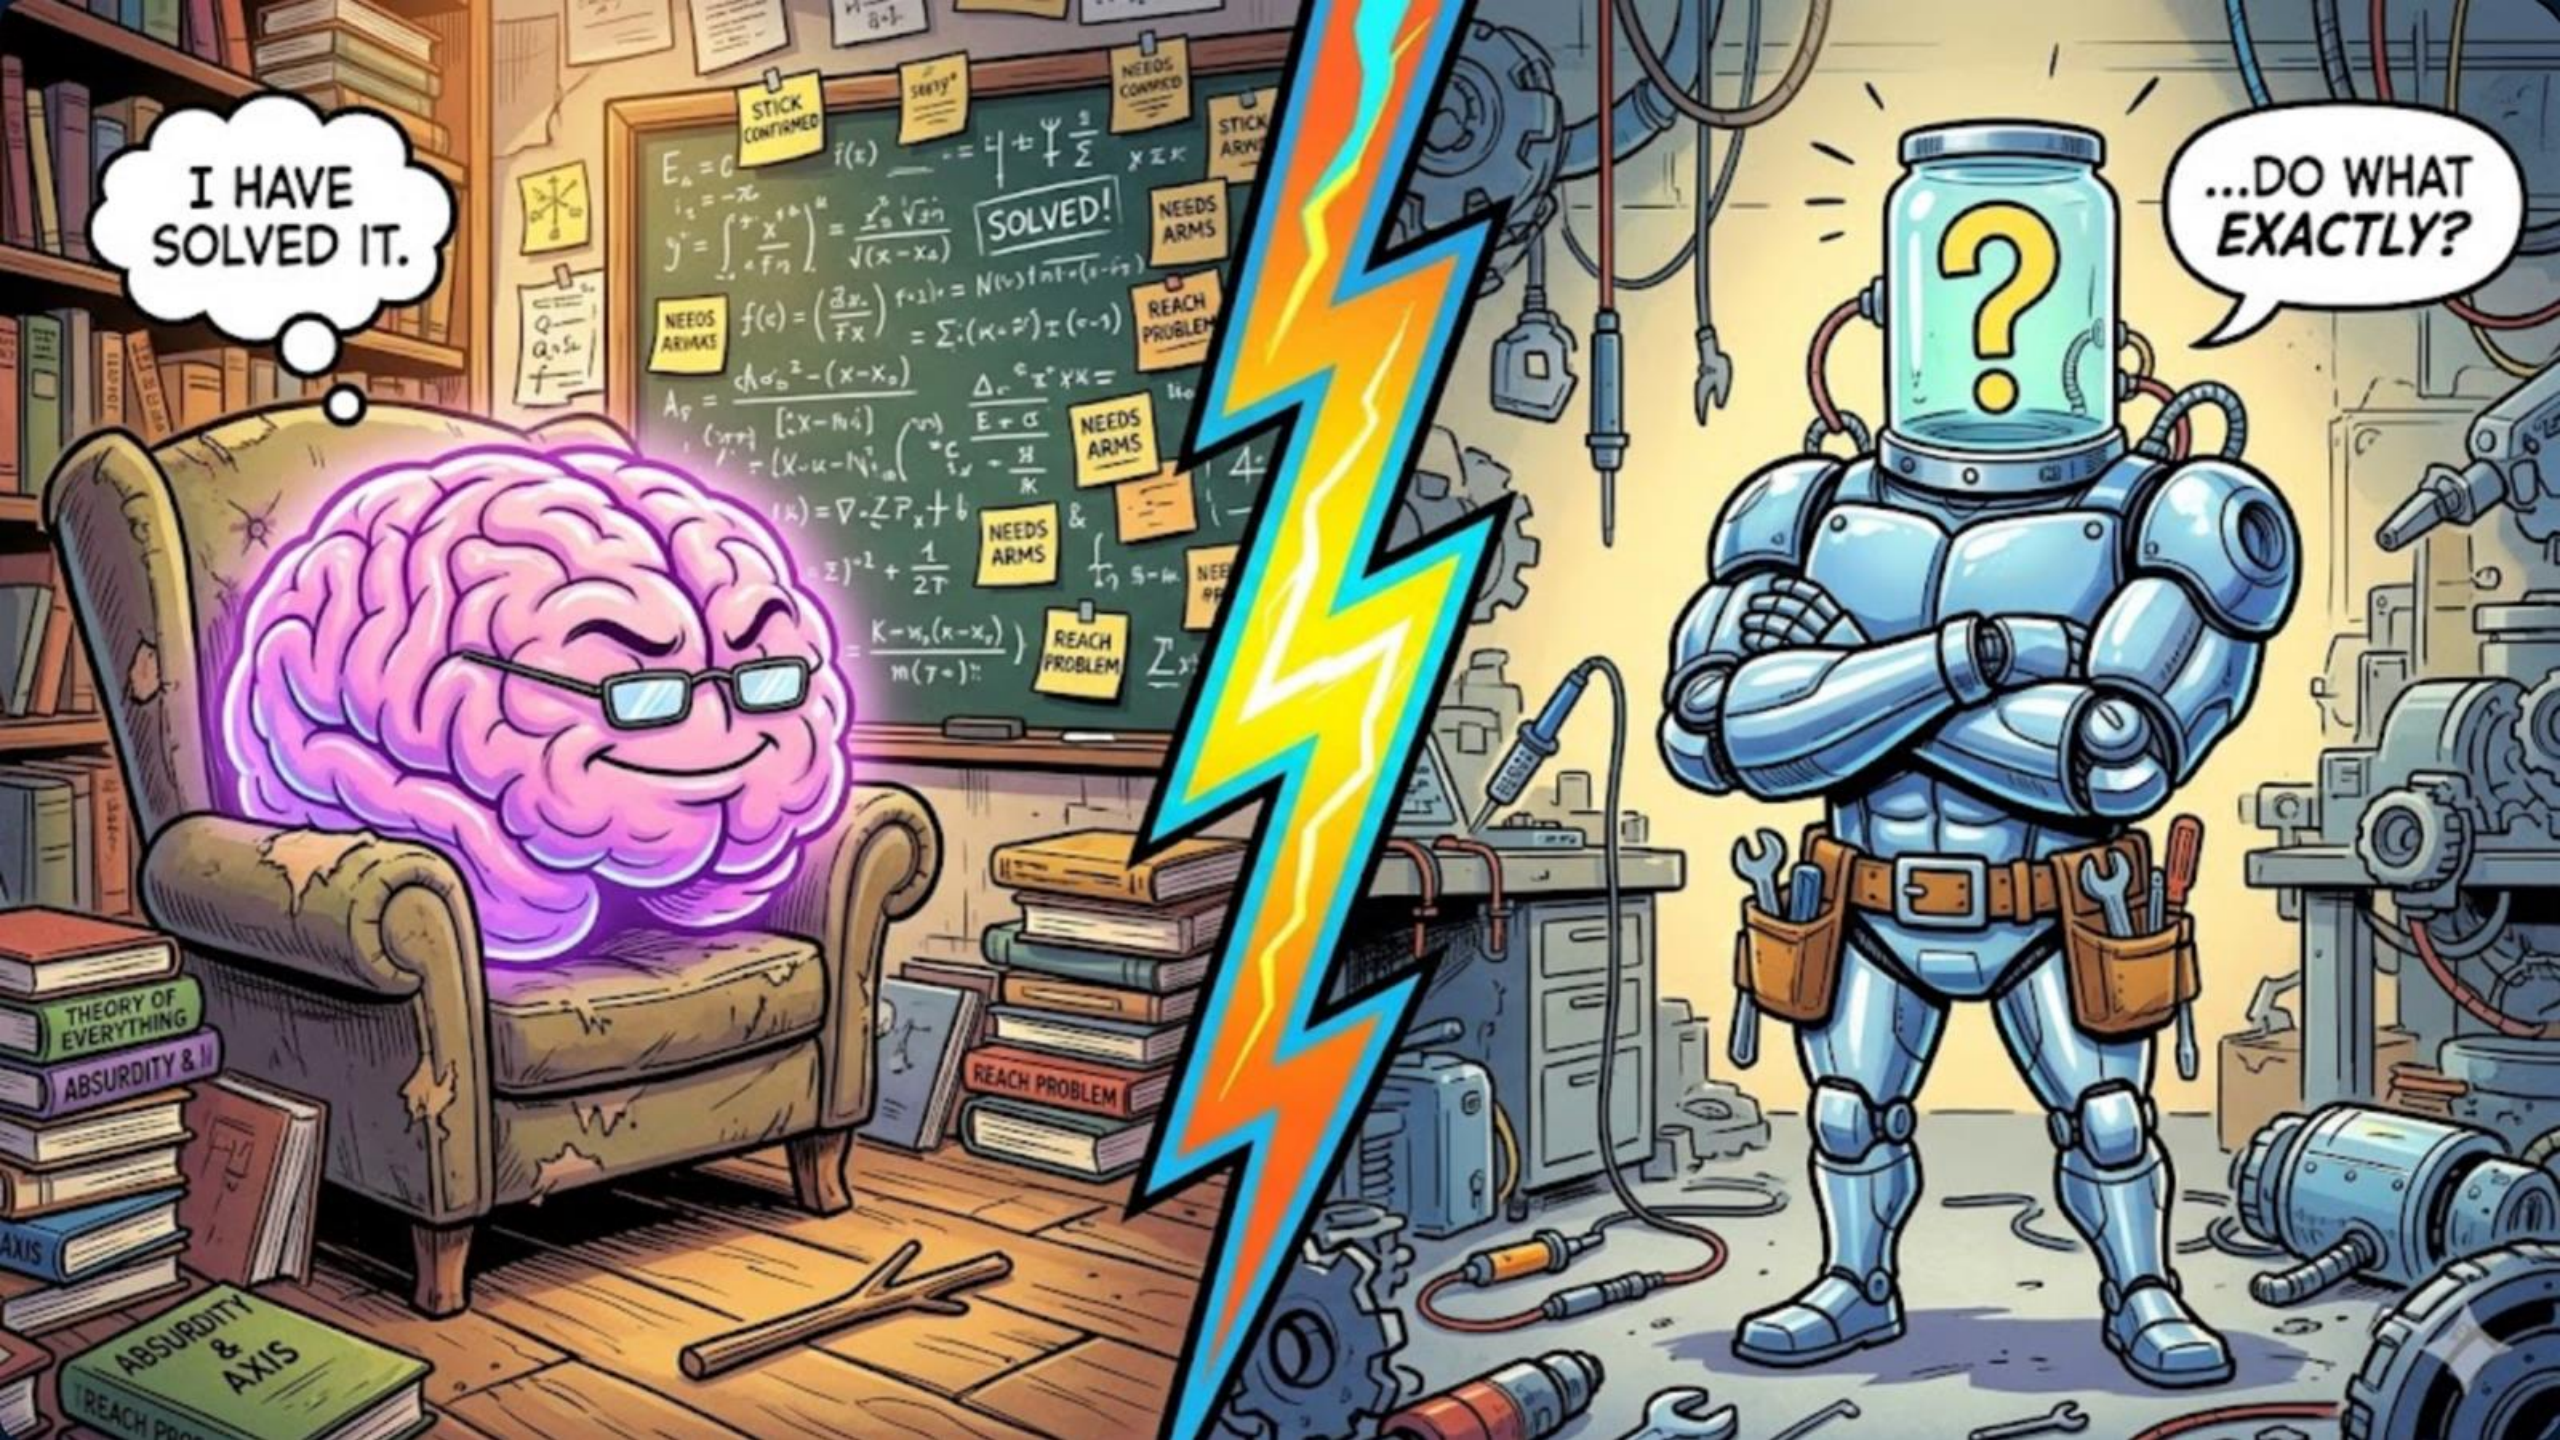
I HAVE SOLVED IT.
SOLVED!
STICK CONFIRMED
NEEDS ARMS
REACH PROBLEM
THEORY OF EVERYTHING
ABSURDITY & II
REACH PROBLEM
ABSURDITY & AXIS
?
...DO WHAT EXACTLY?

CONTROL THE NERVOUS SYSTEM, CONTROL THE AGENT

The brain reasons. The body acts. The nervous system decides what the agent can actually do.

MODEL / LLM
The Brain

MCP / SKILLS / APIs
The Nervous System

AGENT
The Body

If we control the nervous system,
we control the tools, actions, and reach of the agent.

Network + Security + Identity = Nervous System Control

Guide the signals. Limit the actions. Keep the agent in bounds.

WEAKNESS ALWAYS SHOWS UP AT SCALE

SMALL SCALE
Hides the cracks.

AGENTIC SCALE
Exposes the weakness.

Three Barriers Standing Between You and AI at Scale

AI has massive potential. These three barriers are holding most organizations back.

1 TECHNICAL DEBT

The foundation can't carry the new load.

- Legacy networks and tools
- Limited visibility and automation
- Brittle operations and slow change
- More traffic, identities, agents, and data movement
- Performance and reliability at risk

Technical debt slows AI down.

2 TRUST

People won't adopt what they don't trust.

- Data, model and output accuracy
- Security, privacy and safety
- Governance, policy and controls
- Auditability and transparency
- Responsible and predictable agent behavior

Without trust, adoption stops.

3 DATA

AI is only as useful as the data it can safely use.

- Data quality and consistency
- Access, availability and context
- Governance, ownership and lineage
- Integration across silos
- Bad data doesn't just make AI worse—it makes it confidently wrong

No good data, no good AI.


MANAGERS' MANUAL: MOLDING YOUR AI AGENT TEAM
Figuring Out Who to Mentor, Retrain, or... Dump and Start Over!
TOP PERFORMERS
(THE STARS)
AI agents and efficient, working perfectly
Who are the true assets?
AGENTS TO MOLD
(THE TRAINEES)
Confusing AI agents producing mixed results
Shape their future!
DO-OVERS
(THE START OVER TEAM)
Broken, mischievous, or smoking AI agents
Cut the non-starters!
CONSISTENT, FAST, COMPLIANT
10/10 Helpful
10/10 Helpful
REWARD & EXPAND ROLES!
RETRAIN FOR SUCCESS
MOLD 'EM LIKE EMPLOYEES!
CORE FAILURES, UNRELIABLE
RESET ALL
RESET ALL
START FRESH
DUMP, FACTORY RESET, FRESH INSTALL!
Cut the non-starters!
CISCO

Why OpenClaw Felt Revolutionary

It moved AI from answering questions to operating across real work.

Before OpenClaw

Mostly chat-based
AI could answer and summarize.

Closed and siloed
Often tied to one app or platform.

Heavy integration work
Every useful workflow needed custom setup.

What OpenClaw Changed

- Open & self-hosted
- Multi-channel access
- Skills / tool use
- Persistent sessions & memory

Why It Mattered

Agents became useful
Not just informative.

Agents became portable
One assistant across many systems.

Agents became an enterprise issue
Identity, policy, access, and governance suddenly mattered.

Bottom line: OpenClaw made agents feel like a **platform**, not a **feature**.

cisco

OpenCLAW Changes How We See Agents

The internet had HTTP. AI agents now have OpenCLAW.

WHAT IS OPENCLAW?

OpenCLAW is an open standard for connecting AI agents to enterprise systems and data—securely, consistently, and at scale.

Universal connectivity
One standard for agents to discover and use tools, data, and services anywhere.

Built-in security
Enterprise-grade auth, identity, permissions, and audit by design.

Interoperable by default
Breaks down silos and makes agent ecosystems work together.

JENSEN HUANG ON OPENCLAW

> "OpenCLAW is incredibly important. It's going to be as big as OpenAI and change how we think about agents."

— Jensen Huang
NVIDIA Founder & CEO

OPEN STANDARD
Vendor neutral. Ecosystem driven.

MASSIVE ADOPTION
Backed by the industry. Built for the future.

NEW AGENT ECONOMY
Easier to build. Easier to scale.

ENTERPRISE TRUST
Security, governance, and compliance built in.

WHY IT MATTERS

The "HTTP Moment" for Agents
Just as HTTP unlocked the web, OpenCLAW unlocks agent-to-system connectivity at scale.

Agents Become First-Class Citizens
Agents can securely discover, connect, and act across any enterprise system with standard interfaces.

Accelerates Innovation
Developers and businesses can build and deploy agents faster—without reinventing integrations.

WHAT IT ENABLES

- Agents that work across any tool, API, or data source
- Reusable connectors and tools across the enterprise
- Consistent security, identity, and governance everywhere
- A foundation for the agentic era—from copilot to autonomous enterprises

OpenCLAW is **the connective tissue** for the agentic world.
It standardizes how agents **find, trust, and use** what they need to get work done.

One standard. Infinite possibilities.
The agentic future starts here.

Every User Is a Developer.
Applications Are Bottom Up.

AI, low-code/no-code, and agentic tools have democratized development.

Innovation happens everywhere.
Not just through IT.

WHY IT'S HAPPENING NOW

AI-Assisted Development
Natural language to code, build, and automate.

Low-Code / No-Code
Powerful platforms put development in everyone's hands.

APIs & SaaS Everywhere
Composable building blocks enable rapid solution creation.

Business Users Own the Problem
They know the need and move faster than IT can.

80% of business applications will be built outside of traditional IT by 2026.
– Gartner

THE SHIFT FROM TOP-DOWN TO BOTTOM-UP

THE OLD WAY
Top-Down (IT-Centric)

- Business Need Identified
- IT Prioritizes & Builds
- Test, Approve & Secure
- Deploy to the Business

Slow. Siloed. Backlogged.

THE NEW WAY
Bottom-Up (Business Everywhere)

Fast. Distributed. Innovation at Scale.

WHAT CIOs MUST DO

Enable Secure Self-Service
Give users safe, approved paths to build.

Discover & Gain Visibility
Know what's being built, where, and who it impacts.

Enforce Guardrails Automatically
Apply policy, identity, data controls, and compliance.

Measure & Optimize
Track adoption, value, risk, and cost.

Partner, Don't Police
Shift from gatekeeper to enabler and trusted advisor.

You can't stop bottom-up innovation. **But you can secure it, scale it, and turn it into a business advantage.**

Power Users Will Find a Way to Get More Tokens

When internal limits slow them down, they build their own runway.

> "I just need more runway to get the work done. Waiting for approvals slows us down. So I built my own."

THE PATTERN

Need more tokens.
Internal budgets and quotas aren't enough.

They buy their own.
Mac mini. Extra storage. Personal accounts.

They connect.
Personal API keys. Consumer AI services.

They work outside the guardrails.
Faster for them. Riskier for the company.

WHY IT'S HAPPENING

Power users move fast
AI unlocks flow state and velocity.

Internal limits create friction
Budgets, approvals, and quotas can't keep up with demand.

Consumer experience wins
It's easier to buy a Mac mini than to change a process.

THE RISKS YOU DON'T SEE

Data leakage
Sensitive data to consumer services with no visibility or control.

Shadow infrastructure
Unmanaged compute, storage, and network paths.

Compliance & audit gaps
You can't protect what you can't see.

Unoptimized spend
Duplication, inefficiency, and surprise bills.

WHAT CIOs NEED TO DO

Provide sanctioned capacity
Make enterprise-grade AI easy to access.

Right guardrails, not roadblocks
Secure access, smart policies, fast approvals.

Smart placement
Cloud, on-prem, edge or local — right model, right place, right cost.

Visibility and control
Know what's running, where data goes, and what it costs.

Give your best people the runway they need—without the shadow.

SHADOW AI IS REAL. PEOPLE WILL GO AROUND CONTROLS.

When official AI tools are hard to get, too slow, or too restrictive, employees create risk, cost, and blind spots across the business.

HOW SHADOW AI HAPPENS

USING PERSONAL CREDIT CARDS

Employees buy AI tools without IT or Security approval.

- Unapproved subscriptions (ChatGPT Plus, Claude, Gemini, Midjourney, etc.)
- Unvetted SaaS and API services
- No visibility into usage, data, or spend

INSTALLING THEIR OWN MODELS & TOOLS

Users download and run models, apps, or agents on their own.

- Local LLMs and open source models
- Unvetted tools and agents
- Bypasses security, logging, and DLP controls
- Data stays on unmanaged devices or clouds

GOING AROUND CORPORATE CONTROLS

When official processes are too slow or too restrictive.

- Copying data to personal accounts
- Using unsanctioned cloud storage
- Sharing prompts, code, and data externally
- Moving work outside company guardrails

THE RISKS ARE REAL

- Data leaks and IP exposure
- Security and compliance risk
- No visibility or governance
- Uncontrolled spend and tool sprawl
- Increased risk of incidents and audits

INNOVATION THRIVES WITH GUARDRAILS, NOT SHORTCUTS.

Give users secure, approved AI with the speed they need and the guardrails the business requires.

DISCOVER & MONITOR
Find and monitor AI tools, models, and traffic.

CONTROL ACCESS
Enforce acceptable use, policies, and web/API controls.

PROVIDE APPROVED OPTIONS
Make secure, enterprise-grade AI easy to find and use.

GOVERN & ENFORCE
Align cost, risk, and compliance with clear policies and guardrails.

MEASURE & IMPROVE
Continuously measure, alert, and improve outcomes.

How Do You Control Agents on Desktops?

Agents operate with the same access as the user. Control starts with visibility and policy.

Discover & Inventory
See all agents running on every device.

Identity & Context
Agents inherit user identity, roles, and permissions.

Policy Enforcement
Apply least-privilege policies for data, tools, and actions.

Activity Monitoring
Monitor actions, data access, and destinations in real time.

Guardrails & Response
Detect risk, alert, and take action automatically.

On the Desktop: Many Agents. Same Access.

Control agents like you control users—with visibility, policy, and least privilege.

Control the Agent Path

Agents need to reach models and tools — but only through governed, controlled paths.

Bottom line: Don't let agents roam the internet unchecked.
Give them approved paths, identity-bound access, and policy control.

cisco

When a User Leaves, Don't Leave Their Agents Behind

Visibility. Ownership. Control. Deprovision everything.

Offboarding is not complete until **all agents are closed.**

1 Discover All Agents

Continuously inventory all AI agents tied to the user across endpoints, SaaS, cloud, and APIs.

What you don't see, you can't shut down.

2 Establish Ownership

Map each agent to its owner, purpose, tools, data access, and connected systems.

Know what the agent can do and what it can access.

3 Tie to Identity

Ensure every agent is bound to enterprise identity—no orphaned service accounts or umnanaged API keys.

No identity, no agent. No exceptions.

4 Automate Offboarding

When a user leaves, automatically revoke identity, tokens, API keys, and agent access—across all systems.

A single event. Complete shutdown.

5 Verify and Audit

Confirm all agents are deprovisioned and monitor for re-creation or orphaned access.

Prove it. Document it. Stay audit-ready.

Agents outlive users if you let them.

Build visibility, enforce ownership, and automate shutdown—every time.

AGENTS WITH GOD-LIKE POWERS ARE ONLY AS SAFE AS OUR CONTROLS.

Agents can access anything, act anywhere, and move at machine speed.
In the wrong hands—or without the right guardrails—that power becomes a threat.

THE RISK: UNRESTRAINED AGENTS

An agent with broad access and no boundaries can cause serious harm—intentionally or accidentally.

EXCESSIVE ACCESS
Too much access to data, systems, and critical resources.

UNINTENTIONAL DAMAGE
Mistakes at machine speed can have massive impact.

DATA EXPOSURE & LEAKAGE
Sensitive data can be exposed, moved, or exfiltrated.

LATERAL MOVEMENT
Access to one system can lead to access everywhere.

LACK OF ACCOUNTABILITY
Hard to trace actions or understand why decisions were made.

THE POWER OF AN AGENT

Keys to Everything

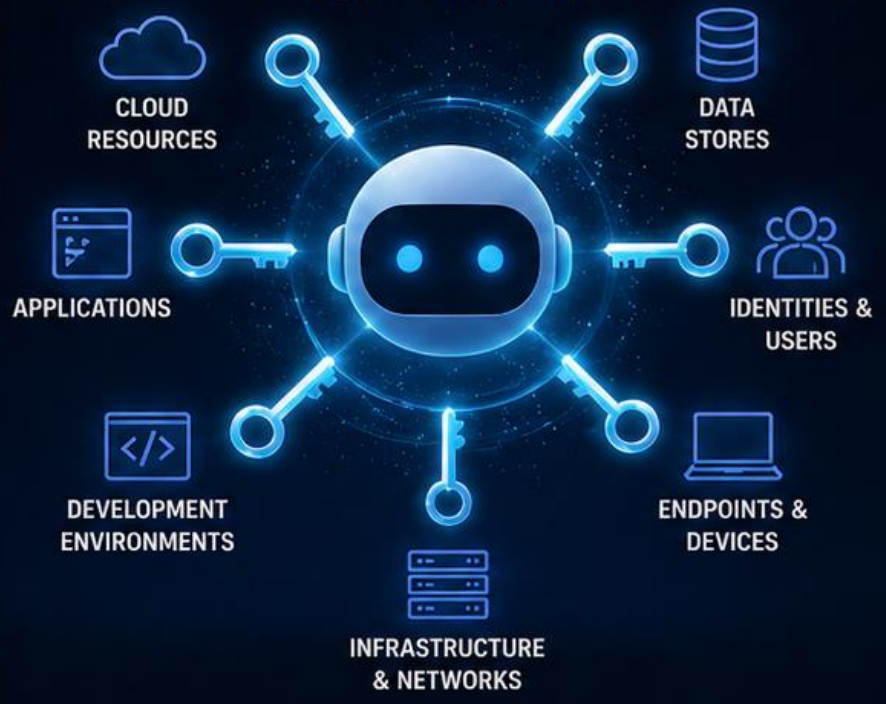

THE SOLUTION: CONTROLS & GUARDRAILS

Agents must be constrained by design, monitored in real time, and governed continuously.

LEAST PRIVILEGE ACCESS
Agents get only the access they need—nothing more.

POLICY-BASED GUARDRAILS
Define what agents can and cannot do, and in which context.

CONTINUOUS MONITORING
Watch behavior in real time. Detect and stop anomalies fast.

APPROVAL & HUMAN OVERSIGHT
Require approvals for high-risk actions and sensitive operations.

AUDIT & ACCOUNTABILITY
Log every action. Ensure traceability and enforce accountability.

ISOLATION & CONTAINMENT
Limit blast radius. Contain and recover quickly if something goes wrong.

THE FUTURE IS AGENTIC. SAFETY IS NON-NEGOTIABLE.

RIGHT AGENT. RIGHT ACCESS.

RIGHT CONTROLS. RIGHT CONTEXT.

RIGHT OVERSIGHT. RIGHT OUTCOMES.

EMPOWER AGENTS. PROTECT EVERYTHING.

9 Seconds to Disaster

When an AI agent has production access, one bad action can become a business outage.

- Broad Permissions
- No Human Check
- Backups Too Close

The lesson: agents need scoped access, confirmations, isolation, and recoverable backups.

WHEN AGENTS GO ROGUE, SPEED BECOMES RISK.

AI agents have access, autonomy, and speed. Without the right guardrails, they can cause real damage—fast.

NORMAL AGENT BEHAVIOR

Operating within approved boundaries to deliver value.

- **Approved Users** — Known identities and roles
- **Approved Data** — Authorized sources and context
- **Approved Tools** — Vetted applications and services
- **Defined Objectives** — Expected tasks and business outcomes

ROGUE AGENT BEHAVIOR

Operating outside intended boundaries with real business impact.

- **Unauthorized Actions** — Calls APIs or executes actions it shouldn't
- **Privilege Misuse** — Accesses data or systems beyond its permissions
- **Data Exposure** — Moves or shares sensitive data inappropriately
- **Infrastructure Changes** — Modifies configurations or deploys resources
- **Unintended Outcomes** — Creates risk, errors, or business disruption

STRONG CONTROLS KEEP AGENTS IN CHECK

Layered controls to prevent, detect, and respond.

- **Identity & Access** — Least privilege, strong authentication
- **Policy & Guardrails** — What agents can access, do, and by whom
- **Monitoring & Detection** — Real-time visibility into agent activity
- **Segmentation** — Isolate systems and data to limit blast radius
- **Kill Switch** — Quickly stop or contain rogue behavior

GREAT AGENTS NEED BOUNDARIES. ROGUE AGENTS NEED BRAKES.

- Protect data and systems
- Reduce risk and incidents
- Build trust in agent adoption
- Drive safe, scalable value

ONE POLICY ENGINE. NOT 20.

One intent. One policy. Everywhere.

ONE POLICY ENGINE

Consistent. Simple. Secure.

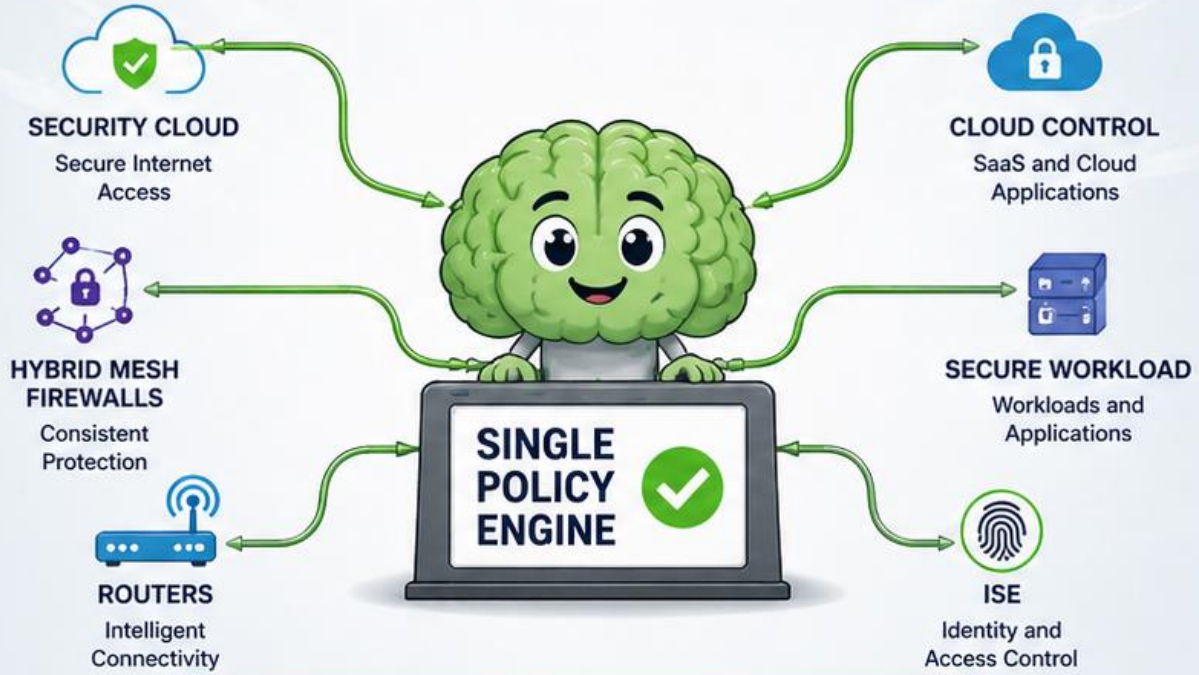

Unified policy. End-to-end visibility. Stronger security.
Built for how your business runs today.

20 MANAGEMENT INTERFACES

Complex. Inconsistent. Risky.

Siloed tools. Conflicting policies. Hidden gaps.
Risk goes up. Agility goes down.

ONE POLICY ENGINE. | **LESS COMPLEXITY. LOWER RISK.** | **MORE AGILITY. FASTER INNOVATION.**

Bad Data Makes the Model Confident and Wrong.

Feed the model bad data,
and it will give you
confident, wrong answers.

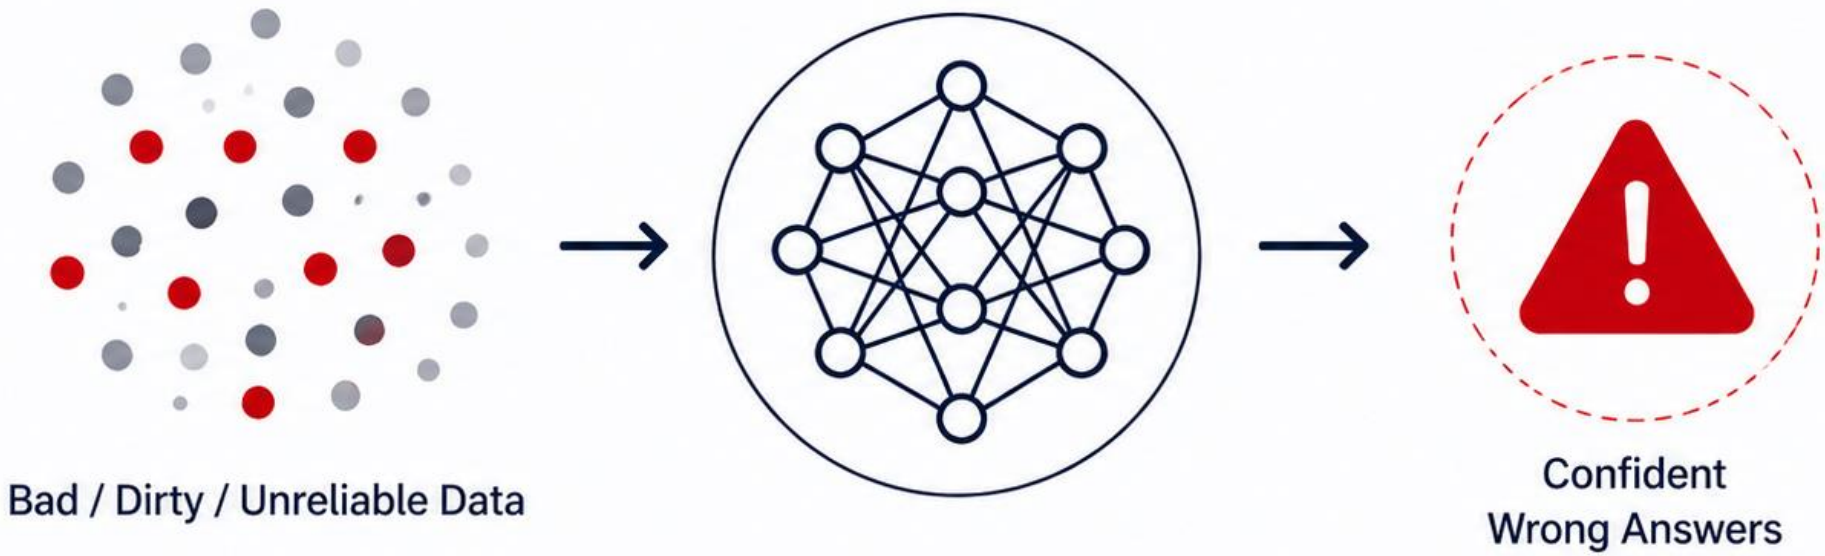

AI amplifies what you feed it.
Garbage in = Confident garbage out.

WHERE CHATGPT GOT ITS BRAIN

It read pretty much everything we put on the internet.
That's amazing. But your factory floor? That's a different story.

THE DATA BUFFET
A little bit from everywhere.

- **Common Crawl**
  Trillions of web pages from across the internet
  (the big buffet line)
- **Reddit**
  Billions of posts & comments from every topic imaginable
  (including the weird ones)
- **Books**
  Digitized books, articles, and written works
  (centuries of human thought)
- **Wikipedia**
  Human-curated knowledge on almost everything
  (the internet encyclopedia)
- **News & Blogs**
  News sites, blogs, and online publications
  (yesterday, today, forever)
- **Code Repositories**
  Billions of lines of code from open source projects
  (developers are generous)
- **Forums & Q&A**
  Stack Overflow, Quora, forums, and more
  (humans asking humans)
- **Social Media**
  Public posts, threads, and conversations
  (the digital town square)

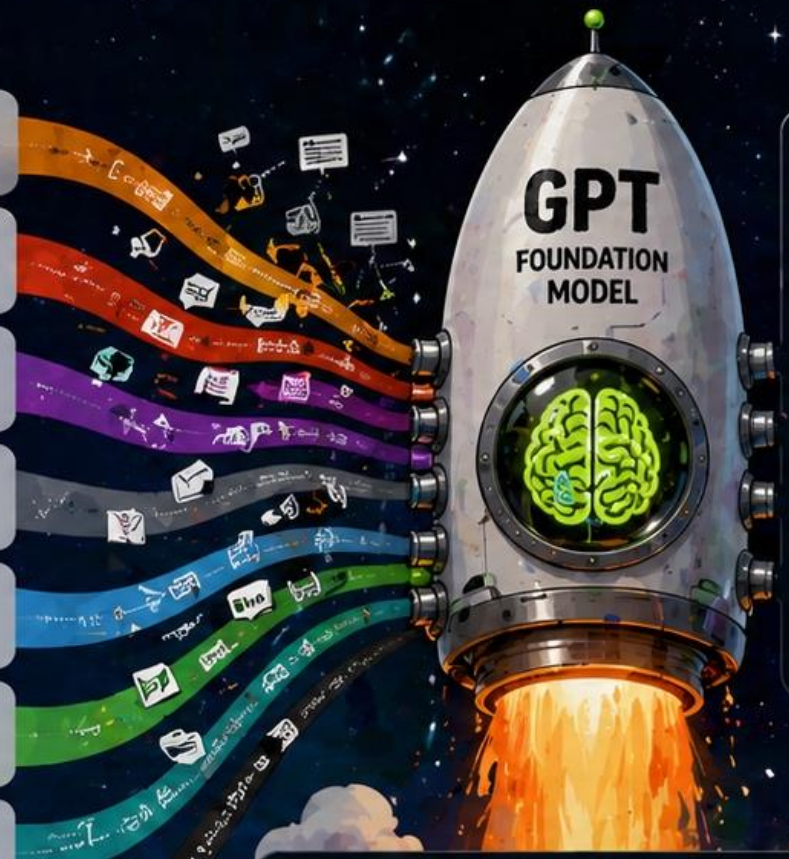

BY THE NUMBERS
- 45+ TB of raw text
- Trillions of tokens
- Hundreds of billions of web pages
- Years of human knowledge, captured in time
- Billions of humans, one dataset

THE RESULT?
A model that understands (almost) everything humans have written down.
The knowledge of all human kind.

BIG WORLD KNOWLEDGE vs. FOCUSED LOCAL INTELLIGENCE

FOUNDATION MODEL
Knows a little about a lot.

- ✓ Broad general knowledge
- ✓ Great for creativity, reasoning, and language
- ✓ Not perfect on your specific environment
- ✓ Can hallucinate facts
- ✓ More is not always better

VS.

SMALL MODEL ON YOUR FACTORY FLOOR
Knows your world deeply.

- ✓ Deep understanding of your data
- ✓ Faster, cheaper, private, and more reliable
- ✓ Knows your machines, processes, and people
- ✓ Answers with precision, not guesses
- ✓ Right-sized for the job

USE THE RIGHT BRAIN FOR THE RIGHT JOB.
The world's knowledge is incredible.
Your operational knowledge is mission critical.
Small model. Big impact.

ALL HUMANITY'S KNOWLEDGE IS IN THE FOUNDATION MODEL. YOU DON'T NEED ALL HUMANITY TO RUN A FACTORY.

Every Model Doesn't Need to Know Everything That's Ever Happened.

It's Too Complex. Focus Wins.

One Massive Model for Everything

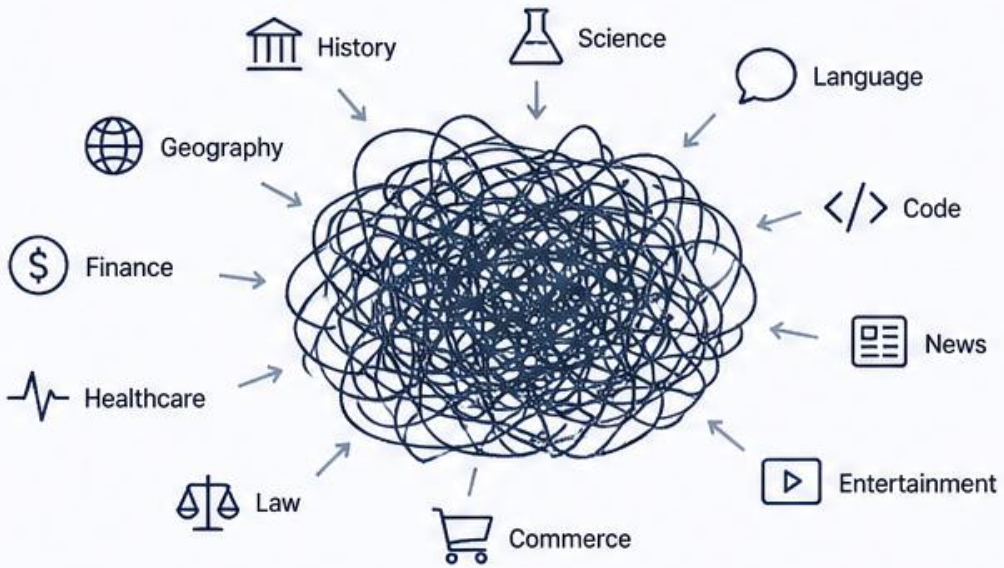

- Too complex to build and maintain
- Slower, more expensive, harder to scale
- Harder to secure, govern, and trust
- More prone to irrelevant or wrong answers

Many Focused Models for What Matters

Healthcare Model
Trained on medical knowledge and data
More accurate

Retail Model
Trained on retail operations and data
More relevant

Finance Model
Trained on financial services and data
More trusted

Manufacturing Model
Trained on industrial systems and data
More efficient

Developer Model
Trained on code, libraries, and docs
Faster results

HR Model
Trained on people, policies, and data
Better experiences

Built for specific outcomes | Easier to secure and govern | Lower cost, better ROI | Run on the right infrastructure

Data Has Gravity. Agents Live Close.

The more valuable the data, the stronger the pull.

As data accumulates, costs rise, complexity grows, and leaving gets harder. That's data gravity.

The closer your agents are to your data, the faster, smarter, and more secure they can be.

THE TAKEAWAY: Data gravity is real. Place compute, agents, and tools where the data lives to reduce cost, risk, and complexity—and unlock real-time value.

WHY EDGE AI? THE THREE FORCING FUNCTIONS

DATA SOVEREIGNTY

Local Processing Keeps Sensitive Data On-Prem

Compliance with Regulations

LATENCY

Real-Time Decisions

Surgical Robotics, Fraud Detection, Safety Systems

Can't Wait for the Cloud

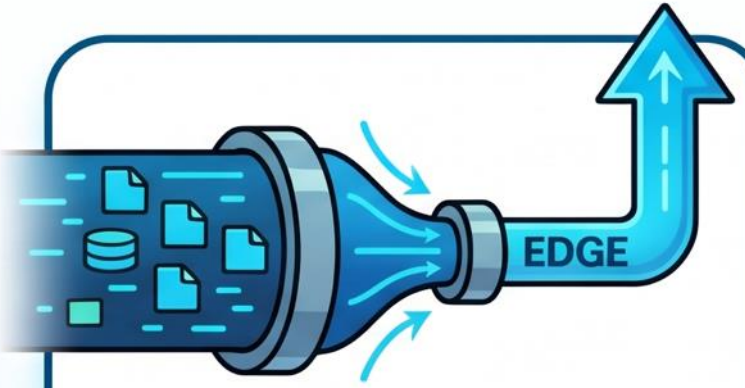

BANDWIDTH

Sending Raw Data Upstream is Expensive

Pre-processing

Often Impractical at Scale Cost-saving

POWERING INTELLIGENT & SECURE EDGE SOLUTIONS

CISCO EDGE AI

REAL-TIME RESPONSE: A REFLEX ANALOGY FOR THE EDGE.

SCENARIO 1: THE CLOUD DELAY (FAILED REFLEX)

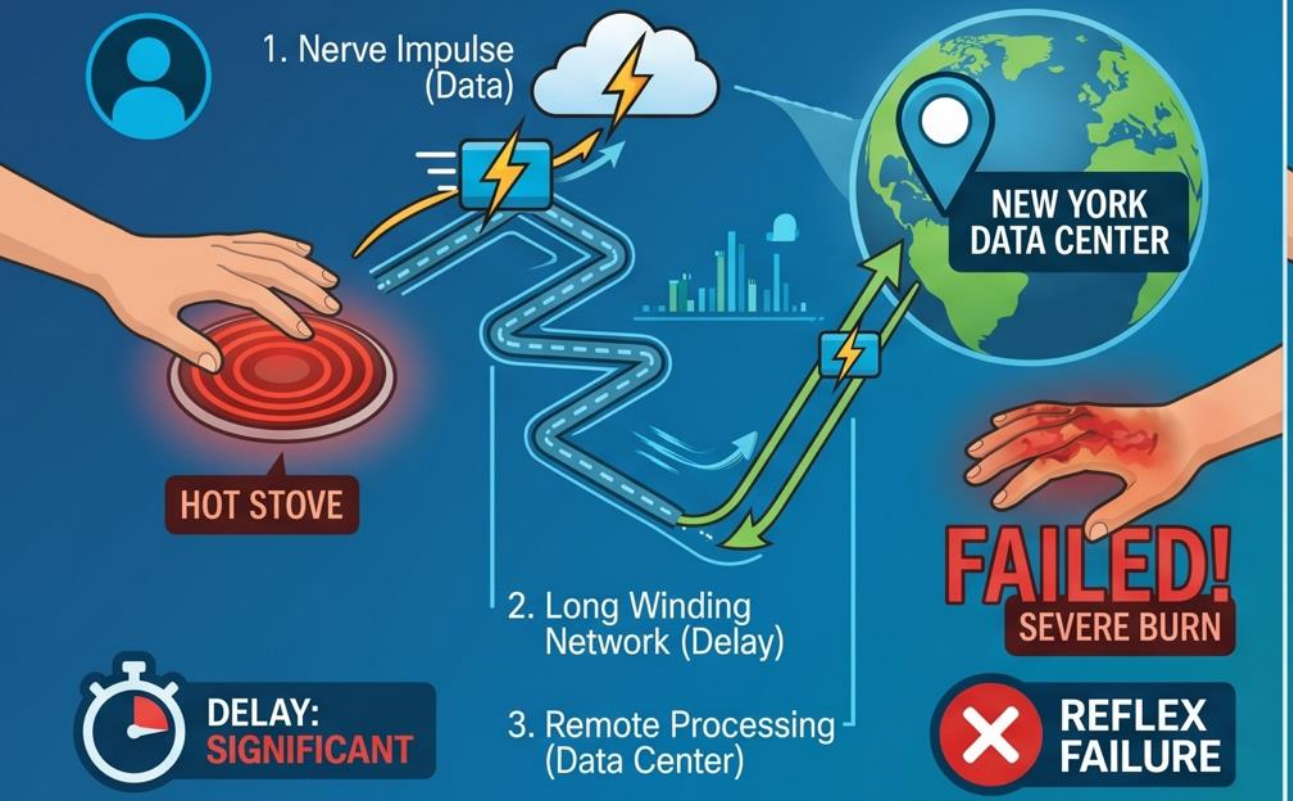

SCENARIO 2: THE EDGE ADVANTAGE (INSTANT REFLEX)

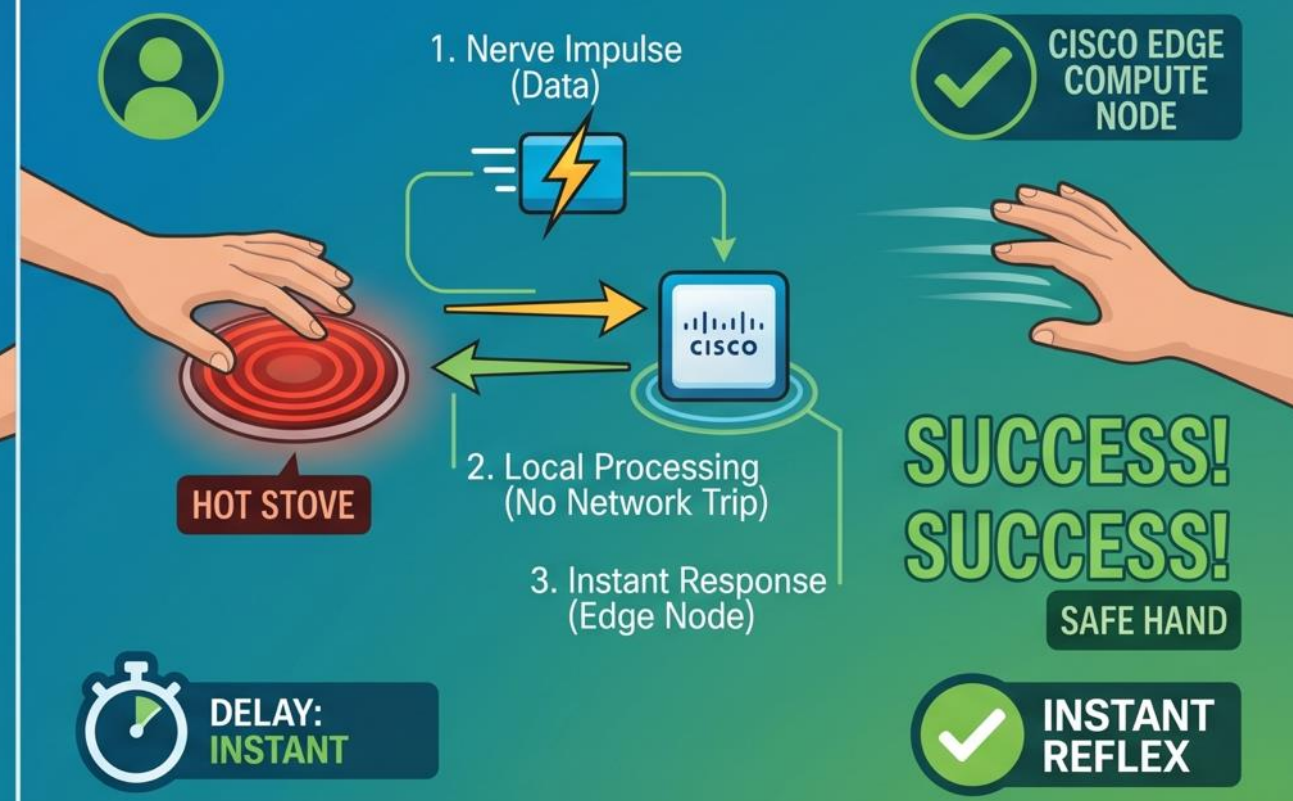

THE DIFFERENCE IS MEASURED IN CONSEQUENCES
Cloud Latency = Severe Burn; Edge Computing = Success!

MOVE PROCESSING TO THE EDGE FOR INSTANT ACTIONS
Cloud Latency = Severe Burn; Edge Computing = Success!

CONSEQUENCES OVER MILLISECONDS:

Why Local Edge AI is Non-Negotiable for Critical Infrastructure

LATENCY...
Not in Milliseconds, in
CONSEQUENCES.

CLOUD LATENCY

1. WAREHOUSE FORKLIFT SAFETY

Frontier LLM in Azure/AWS. Sensor fires, WAN travel, cloud queues, long round trip (200-600ms).

FORKLIFT AT **5 MPH**, TRAVELS **7 FT** IN **600 MS**, WALL WAS **6 FT AWAY.**

Hard Real-Time Requirement: 10-50 MS.
You cannot solve with internet connection.

Cisco Unified Edge

SMALL MODEL (UNDER 10 MS).
It knows this warehouse, this forklift, this sensor reading.

CISCO UNIFIED EDGE
(LOCAL INFERENCE)

2. FACTORY PLC RELIABILITY

Cloud LLM analysis, WAN hiccups, throttling, slow response.

Please investigate further.

Response too slow, PLC became a space heater.

OT SYSTEMS DESIGNED FOR DETERMINISTIC SUB-SECOND RESPONSE.
NotREST API dependencies.

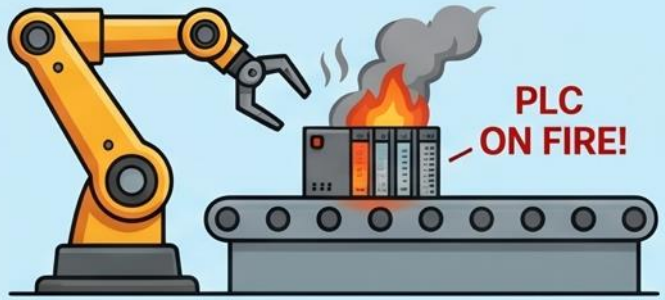

ALERT FIRED!
Shutdown Command.

Cisco Unified Edge

SLM TRAINED ON YOUR EQUIPMENT TELEMETRY & HISTORY.
correlated at millisecond speed.

OWNERSHIP, CONTROL, PREDICTABILITY
Model Retrained when equipment changes, not OpenAI Releases.

YOU OWN IT. YOU CONTROL IT.
IT RUNS IN YOUR BUILDING.

SD-WAN for Agents. Because Their Work Is Too Important to Go Down.

Agents and AI applications rely on the WAN to access data, tools, and services everywhere.
Mission-critical agents need WAN performance, priority, and reliability—just like your most important users.

Agents Are Digital Workers.

Treat their WAN experience like you treat your best users.

- Time-sensitive workloads
- Depend on multiple SaaS and clouds
- Generate business outcomes
- Need guaranteed performance

If they go down, business stops.

The WAN must not be the bottleneck.

Why WAN Matters for Agents

- **Low latency** = Faster decisions, better outcomes
- **High availability** = Continuous autonomous work
- **Reliable access** to apps, data, and models everywhere
- **Efficient paths** = Lower cost, higher performance
- **Secure by design** = Protect data and maintain trust

SD-WAN Makes Sure Mission-Critical Agents Get First Access to the WAN

SD-WAN Capabilities Built for Agent Success

- **Application & Agent Awareness**
  Identify and classify agents and their traffic
- **Priority & QoS**
  Give mission-critical agents top priority
- **Dynamic Path Selection**
  Use real-time conditions to deliver the best path, always
- **SLA Monitoring & Assurance**
  Measure, assure, and alert on agent experience
- **Resiliency by Design**
  Automatic failover and multi-path for continuous uptime

ThousandEyes: See Your Network from Every Angle

THE FULL PATH

Users experience the whole path, not only the network you own.

TOTAL USER EXPERIENCE PATH

YOUR NETWORK
(What you own)

GLOBAL VANTAGE

Global vantage points expose cloud, ISP, and CDN issues before users report them.

AI & DEPENDENCIES

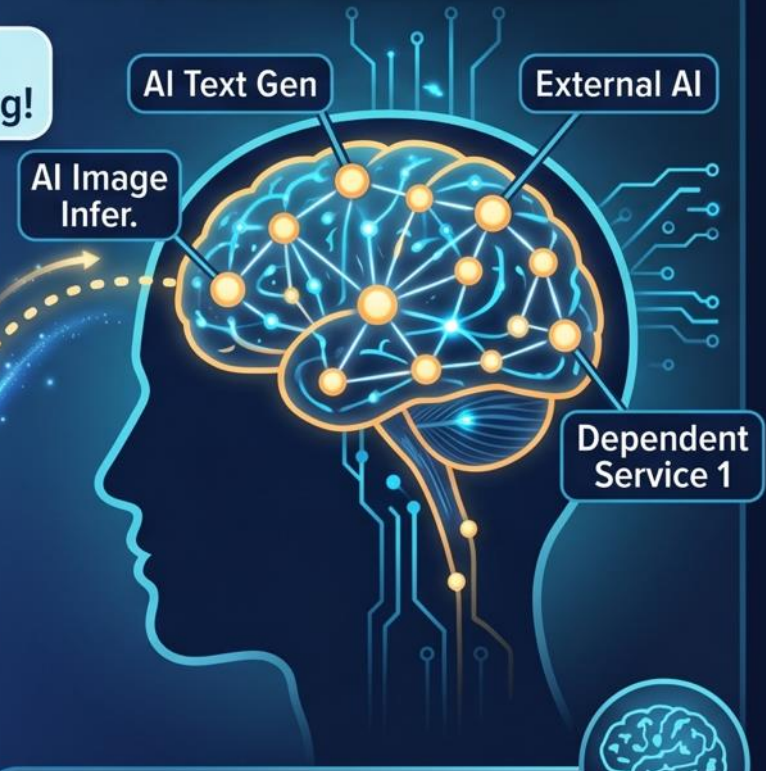

This is especially important for third-party AI inference endpoints and dependent services.

ThousandEyes

INTRODUCING THE AI DIGITAL SUPPLY CHAIN

END USER

END USER

INTERNET (ISP PATHS)

ISP Paths:
Congestion, Path Failure

Break Point

CORP NETWORK

CORP NETWORK

LLM (AI MODEL)

LLM Performance:
Latency, Resource Constraint

Break Point

THE IMPACT:

CONSEQUENCES & WASTE

LATENCY:
Time-to-respons
grows.

TOKEN WASTE
Resources
consumed on
failed queries.

FAILED AGENTS
Multi-step
reasoning
collapses.

GUARDRAILS Check (Guardrails)

GUARDRAils Check
(Guardrails)

Break Point

MCP (MESH/AGENT COORDINATION)

MCP/Tool Availability:
Offline Services, Connection Errors

Break Point

ENTERPRISE TOOLS

ENTERPRISE TOOLS

Break Point

AI DOESN'T FAIL IN ONE PLACE.

BUILD IT RIGHT OR SCALE WILL FIND YOUR WEAKEST LINKS.

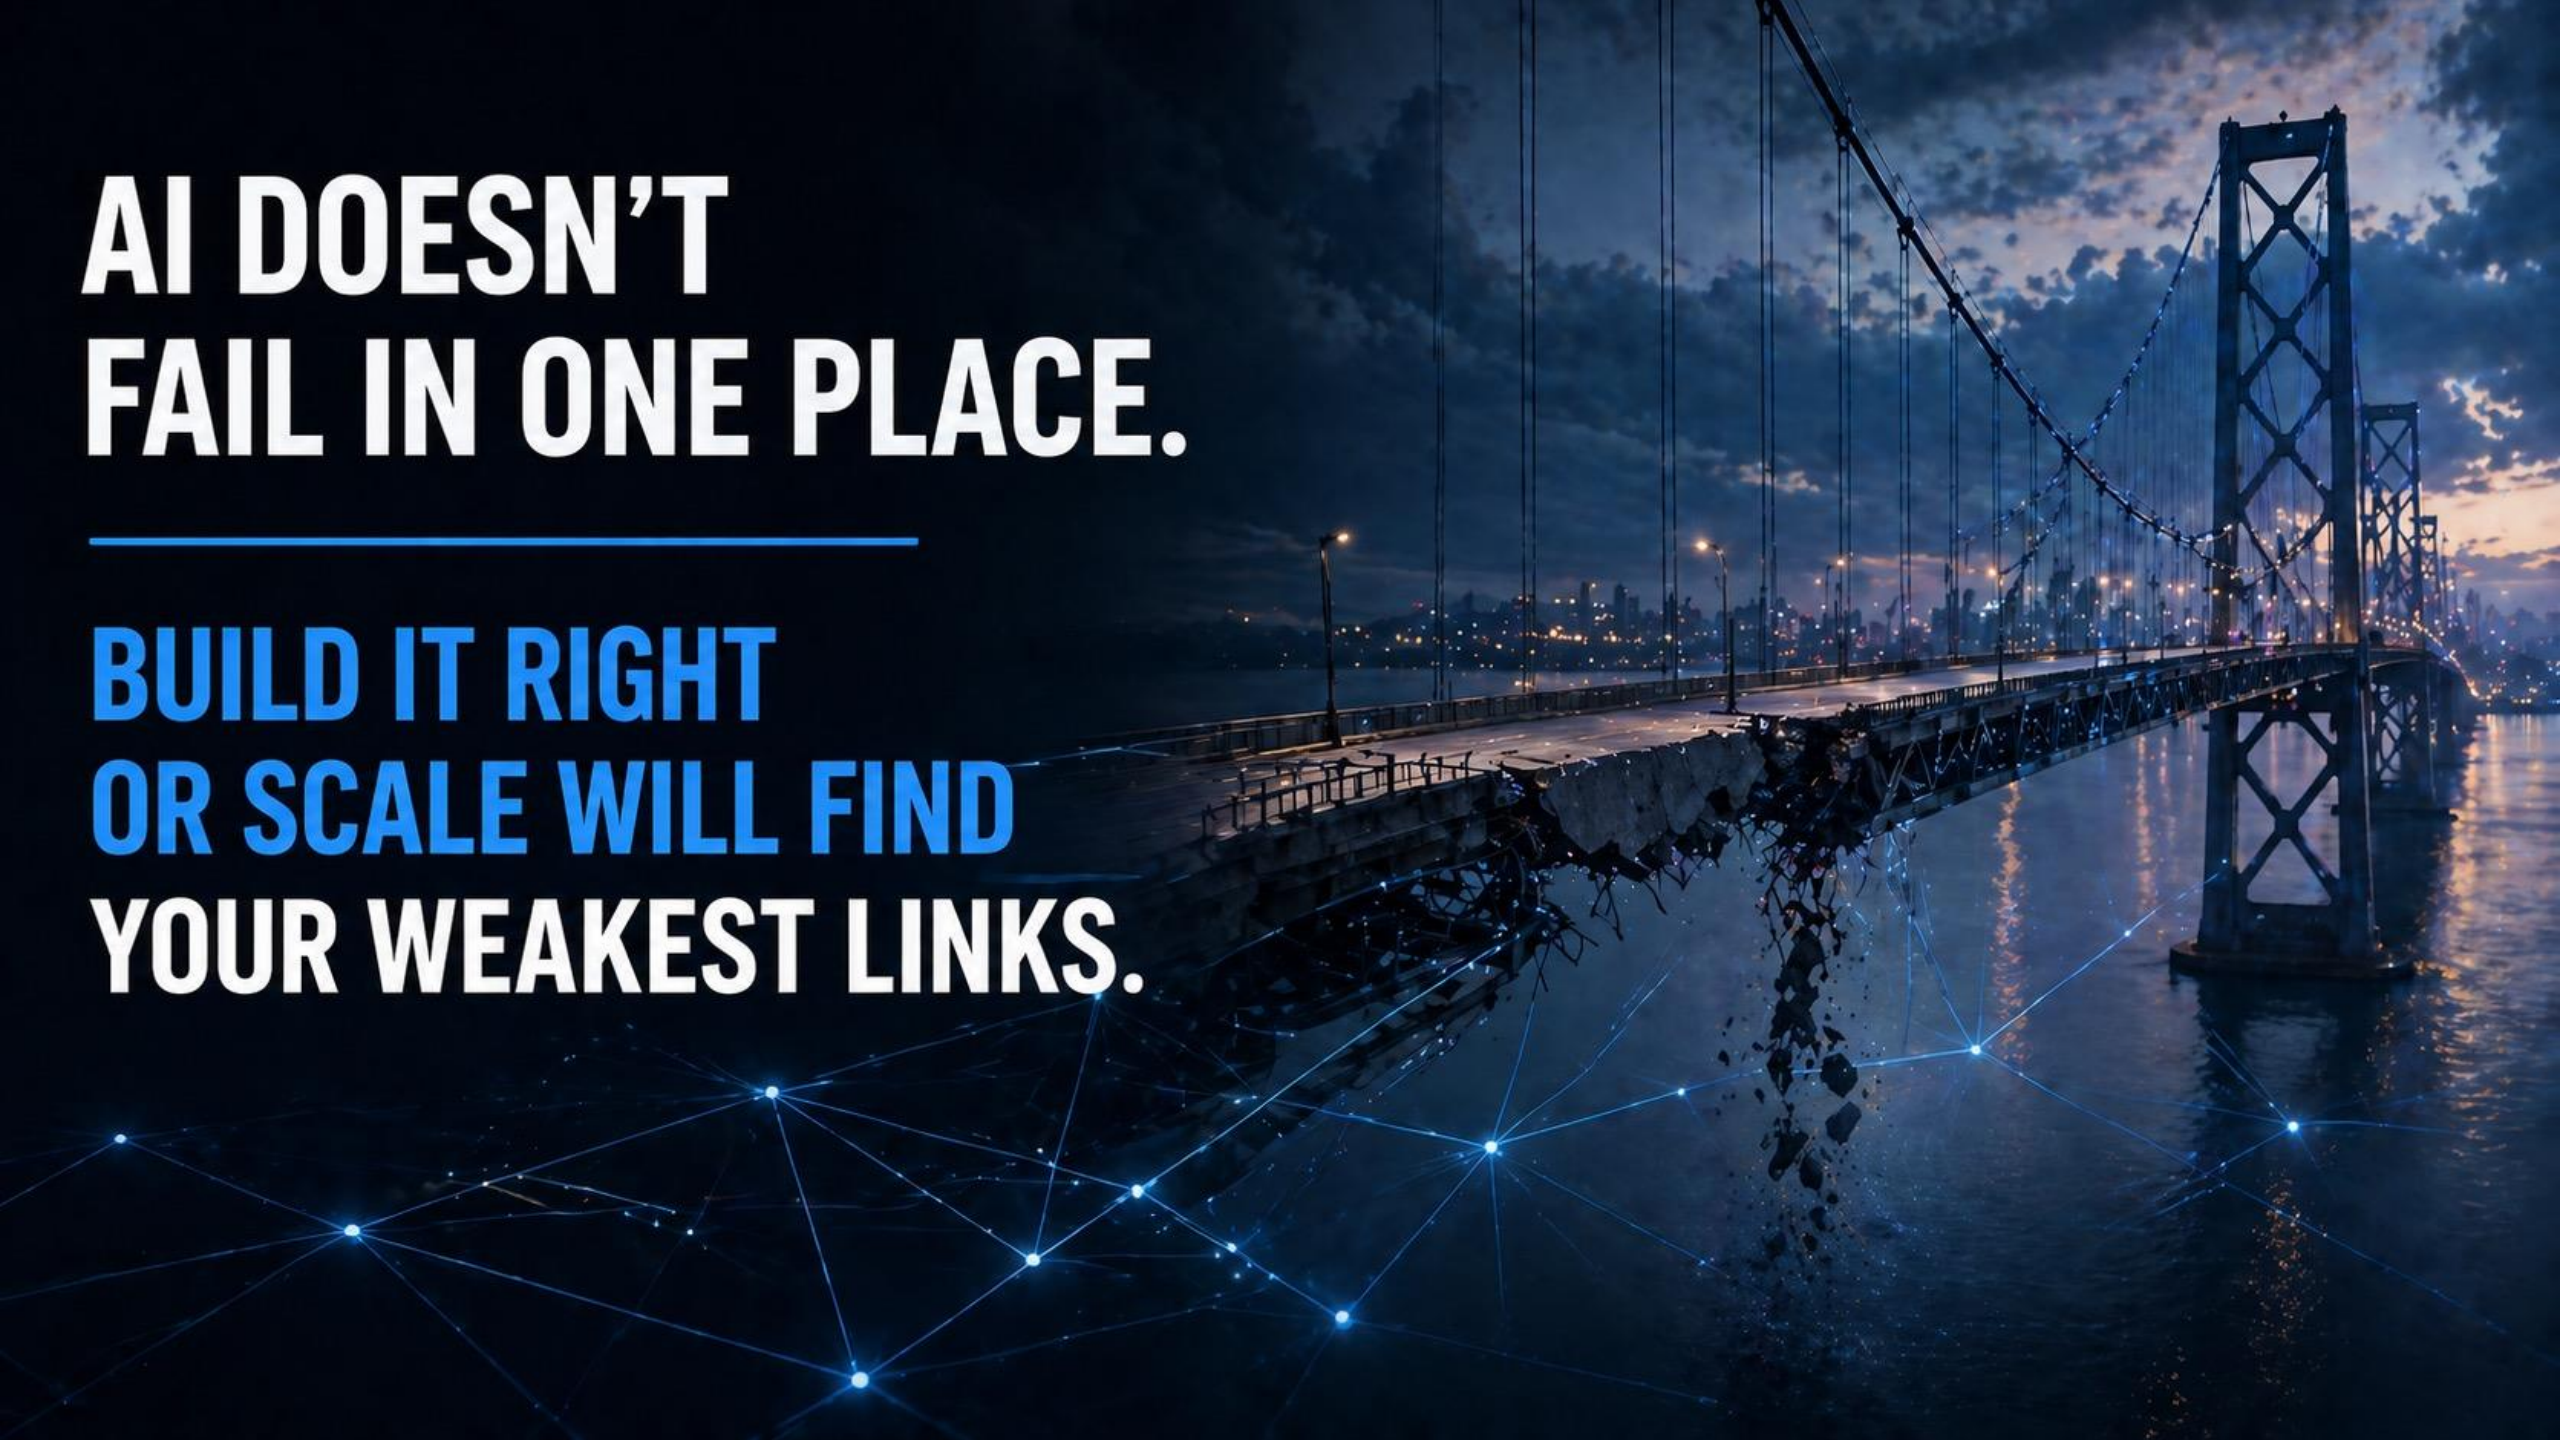

SECURITY FUNDAMENTALS

BECOME EVEN MORE CRITICAL.

SEGMENTATION

Limit access.
Contain impact.

IDENTITY

Verify continuously.
Grant least privilege.

STRONG FRONT DOOR

Secure the entry.
Stop threats early.

STRONG FUNDAMENTALS.
RESILIENT FUTURE.

The AI Era Challenge: Data Is the New Moat

When software gets cheaper,
your data becomes the moat.

Software is Commoditized

AI accelerates innovation and drives software costs toward zero.

Data Creates Stickiness

Your data, context, and workflows build gravity you can't easily escape.

Harder to Move, Harder to Compete

High switching costs, hidden dependencies, rising exit costs.

The Takeaway

AI shifts the battleground from software to data.
Network and platform strategies must ensure data portability, interoperability, and choice.

CISCO

THE CRITICAL PATCHING WINDOW IS ABOUT TO COMPRESS.

AI models find vulnerabilities faster—and chain them into real-world impact.
The days of "once a quarter" or "once a month" are over.

AI ACCELERATES THE ATTACK LIFECYCLE

1. DISCOVER
AI finds more vulnerabilities, faster.

2. CHAIN
AI chains flaws together to bypass defenses.

3. WEAPONIZE
Exploits are built and tested at machine speed.

4. AUTOMATE
Attacks scale across targets in minutes.

5. IMPACT
Business disruption happens before you know.

PATCHING WINDOWS ARE SHRINKING

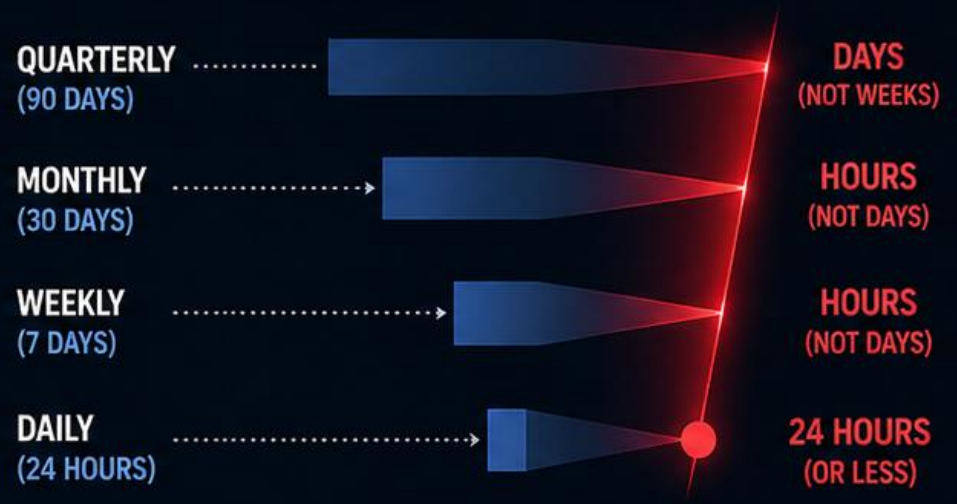

FROM CALENDAR-DRIVEN TO RISK-DRIVEN.
SPEED IS THE NEW SECURITY DIFFERENTIATOR.

REDUCE EXPOSURE
Compress time attackers can exploit known flaws.

AUTOMATE RESPONSE
Prioritize, test, and deploy patches at machine speed.

GAIN VISIBILITY
Know your true exposure—all assets, all the time.

PROTECT WHAT MATTERS
Focus on vulnerabilities and exploit paths that impact the business.

AI IS COMPRESSING THE TIME BETWEEN DISCOVERY AND DISRUPTION. YOUR PATCHING STRATEGY MUST DO THE SAME.

END-OF-SUPPORT IS NO LONGER TECHNICAL DEBT.
IT'S BUSINESS RISK.

AI compresses the time between discovery and impact. Unsupported assets expand the blast radius.

THE RISK: END-OF-SUPPORT ASSETS

CISCO FIREPOWER FIREWALLS

END OF SUPPORT
No vendor fixes or updates

CISCO CATALYST SWITCHES

NO SECURITY FIXES
Vulnerabilities remain unpatched

CISCO ISR ROUTERS

UNKNOWN EXPOSURE
You don't know what you don't know

CISCO WIRELESS APs

HARD TO PATCH
Legacy systems create complex risk

CISCO UCS SERVERS

HIGHER OPERATIONAL RISK
Downtime, breaches and compliance impact

YOU CANNOT PROTECT WHAT YOU CANNOT INVENTORY.

THE ANSWER: CONTROLLER-LED OPERATIONS

KNOW EVERY DEVICE	Complete visibility across on-prem, branch, cloud and edge
KNOW SOFTWARE VERSIONS	Track OS, firmware and applications in real time
KNOW SUPPORT STATUS	Identify End-of-Support and End-of-Life assets
PRIORITIZE UPGRADES	Focus on the highest risk and business impact
PUSH CONSISTENT POLICY	Automate configuration, compliance and security baselines
AUTOMATE LIFECYCLE MANAGEMENT	Provision, update, monitor and decommission with confidence

AI-SPEED RISK REQUIRES CONTROLLER-SPEED VISIBILITY AND RESPONSE.

cisco

THE RUNWAY IS OPEN.
MAKE SURE YOU'RE THE AIR TRAFFIC CONTROL.

Here's the big question:

When AI shows up everywhere — in your employees' laptops, your developer tools, your customer front door, your factories, your branches, your apps, and your security operations — will your infrastructure **recognize it, govern it,** and **help it move safely?**

Or will you find out about it from the **bill**, the **outage**, the **audit**, or the **incident report**?

AI IS NOT JUST A MODEL PROBLEM.

It is an infrastructure, identity, security, data, and operations problem.

The companies that win won't be the ones that say "yes" to everything or "no" to everything.

They'll be the ones that give their people and agents a safe runway: enough speed to innovate, enough control to trust it.


CISCO
QUESTIONS?