CISCO Connect

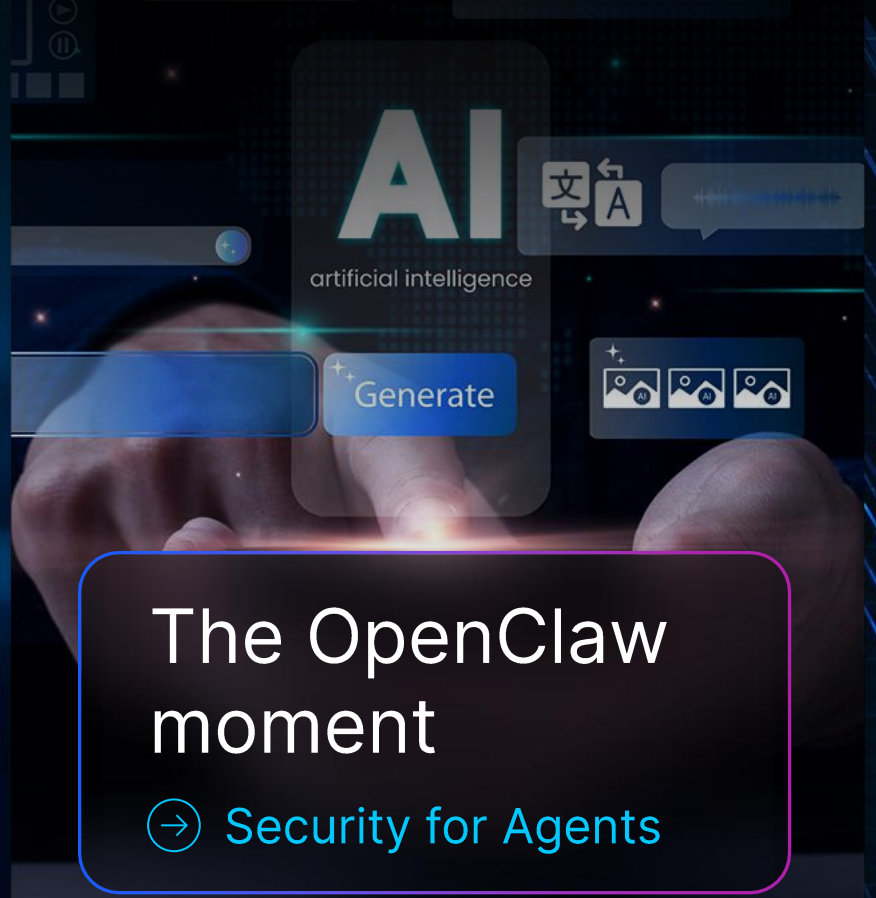
Security at Agentic Scale



Matthew J Joseff, CFE
Director Unified Data & Digital Resilience

Date placeholder

The three forces re-shaping security



The Mythos moment

Shields Up

The OpenClaw moment

Security for Agents

The need for speed

Security using Agents

Do bees perceive time?

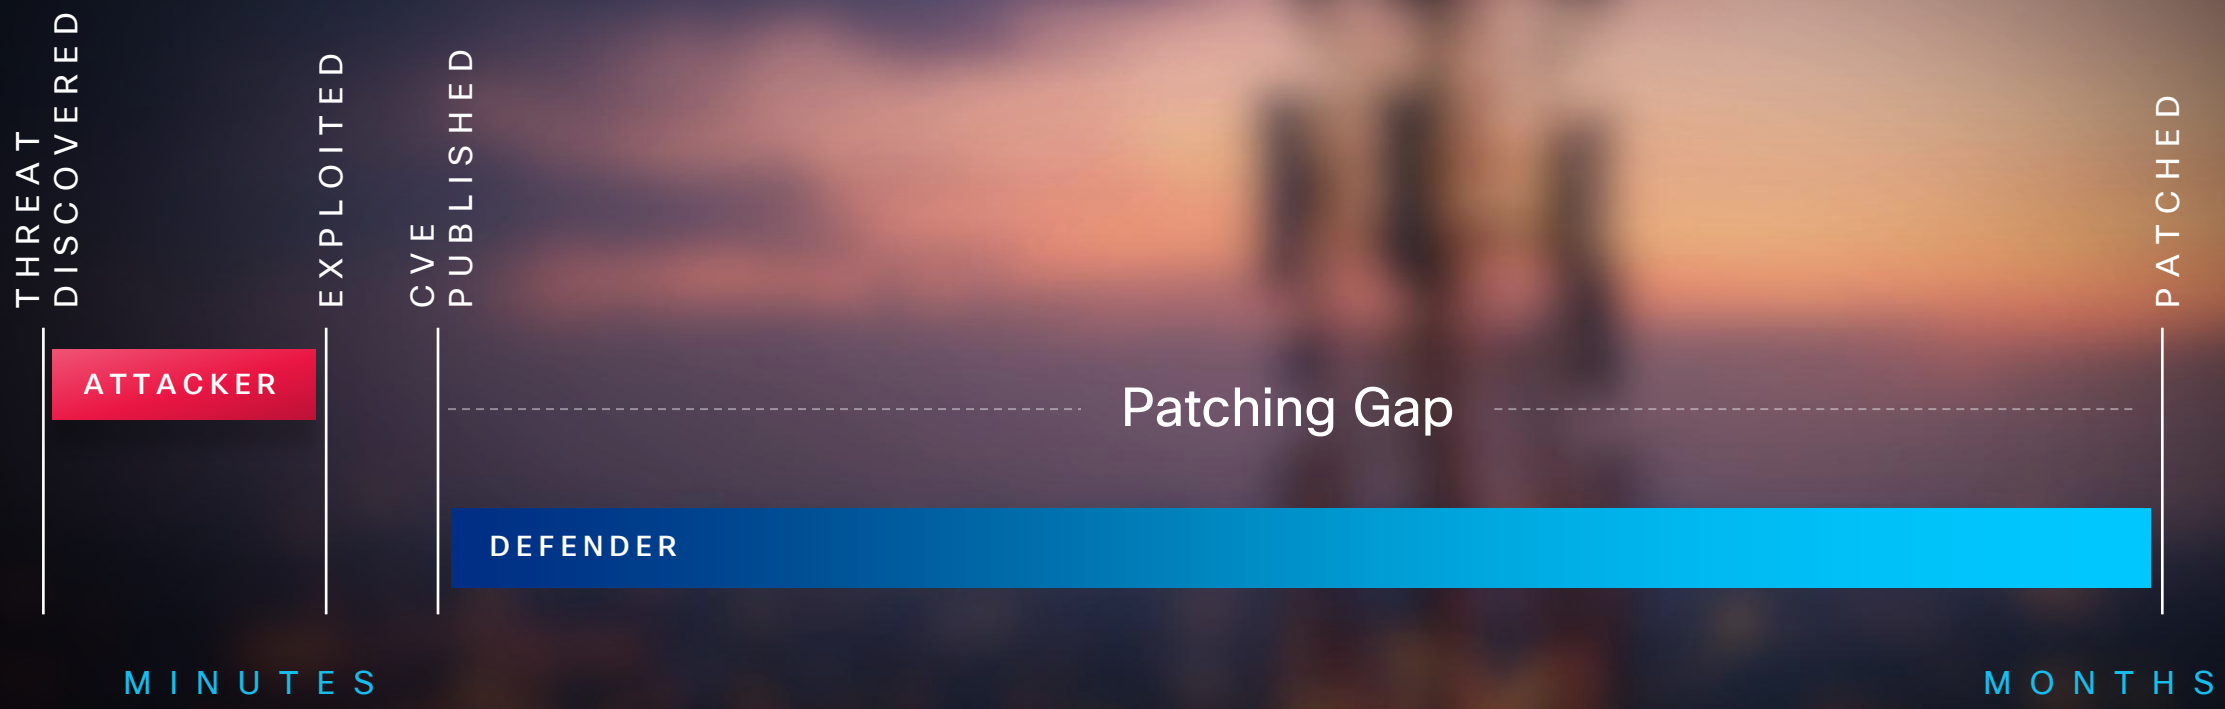
- 1600 everyday feeding
- In the dark
- Underground
- Paris to NY



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

Cisco Confidential

Patching is hard.
Frontier AI made it harder



Live Protect



CISCO NETWORK DEVICES
STARTING WITH N9000

Apply policy to live system

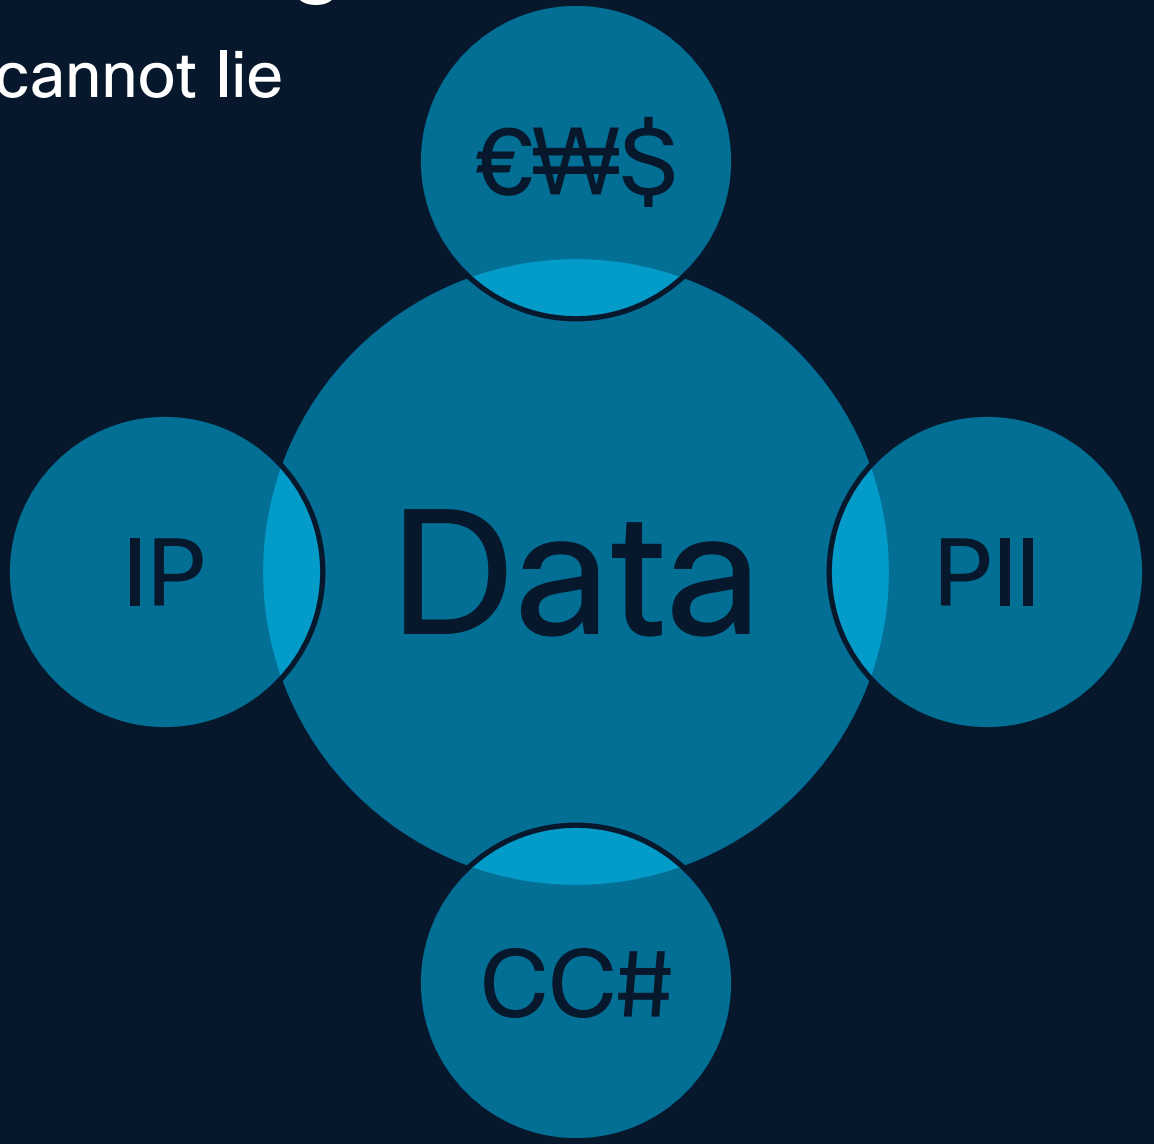
Protect now, patch in planned window

No downtime or switch reboot

Continuous re-validation

What are we protecting?

I like big data, and I cannot lie



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

Cisco Confidential

Quantum breaks cryptography foundations

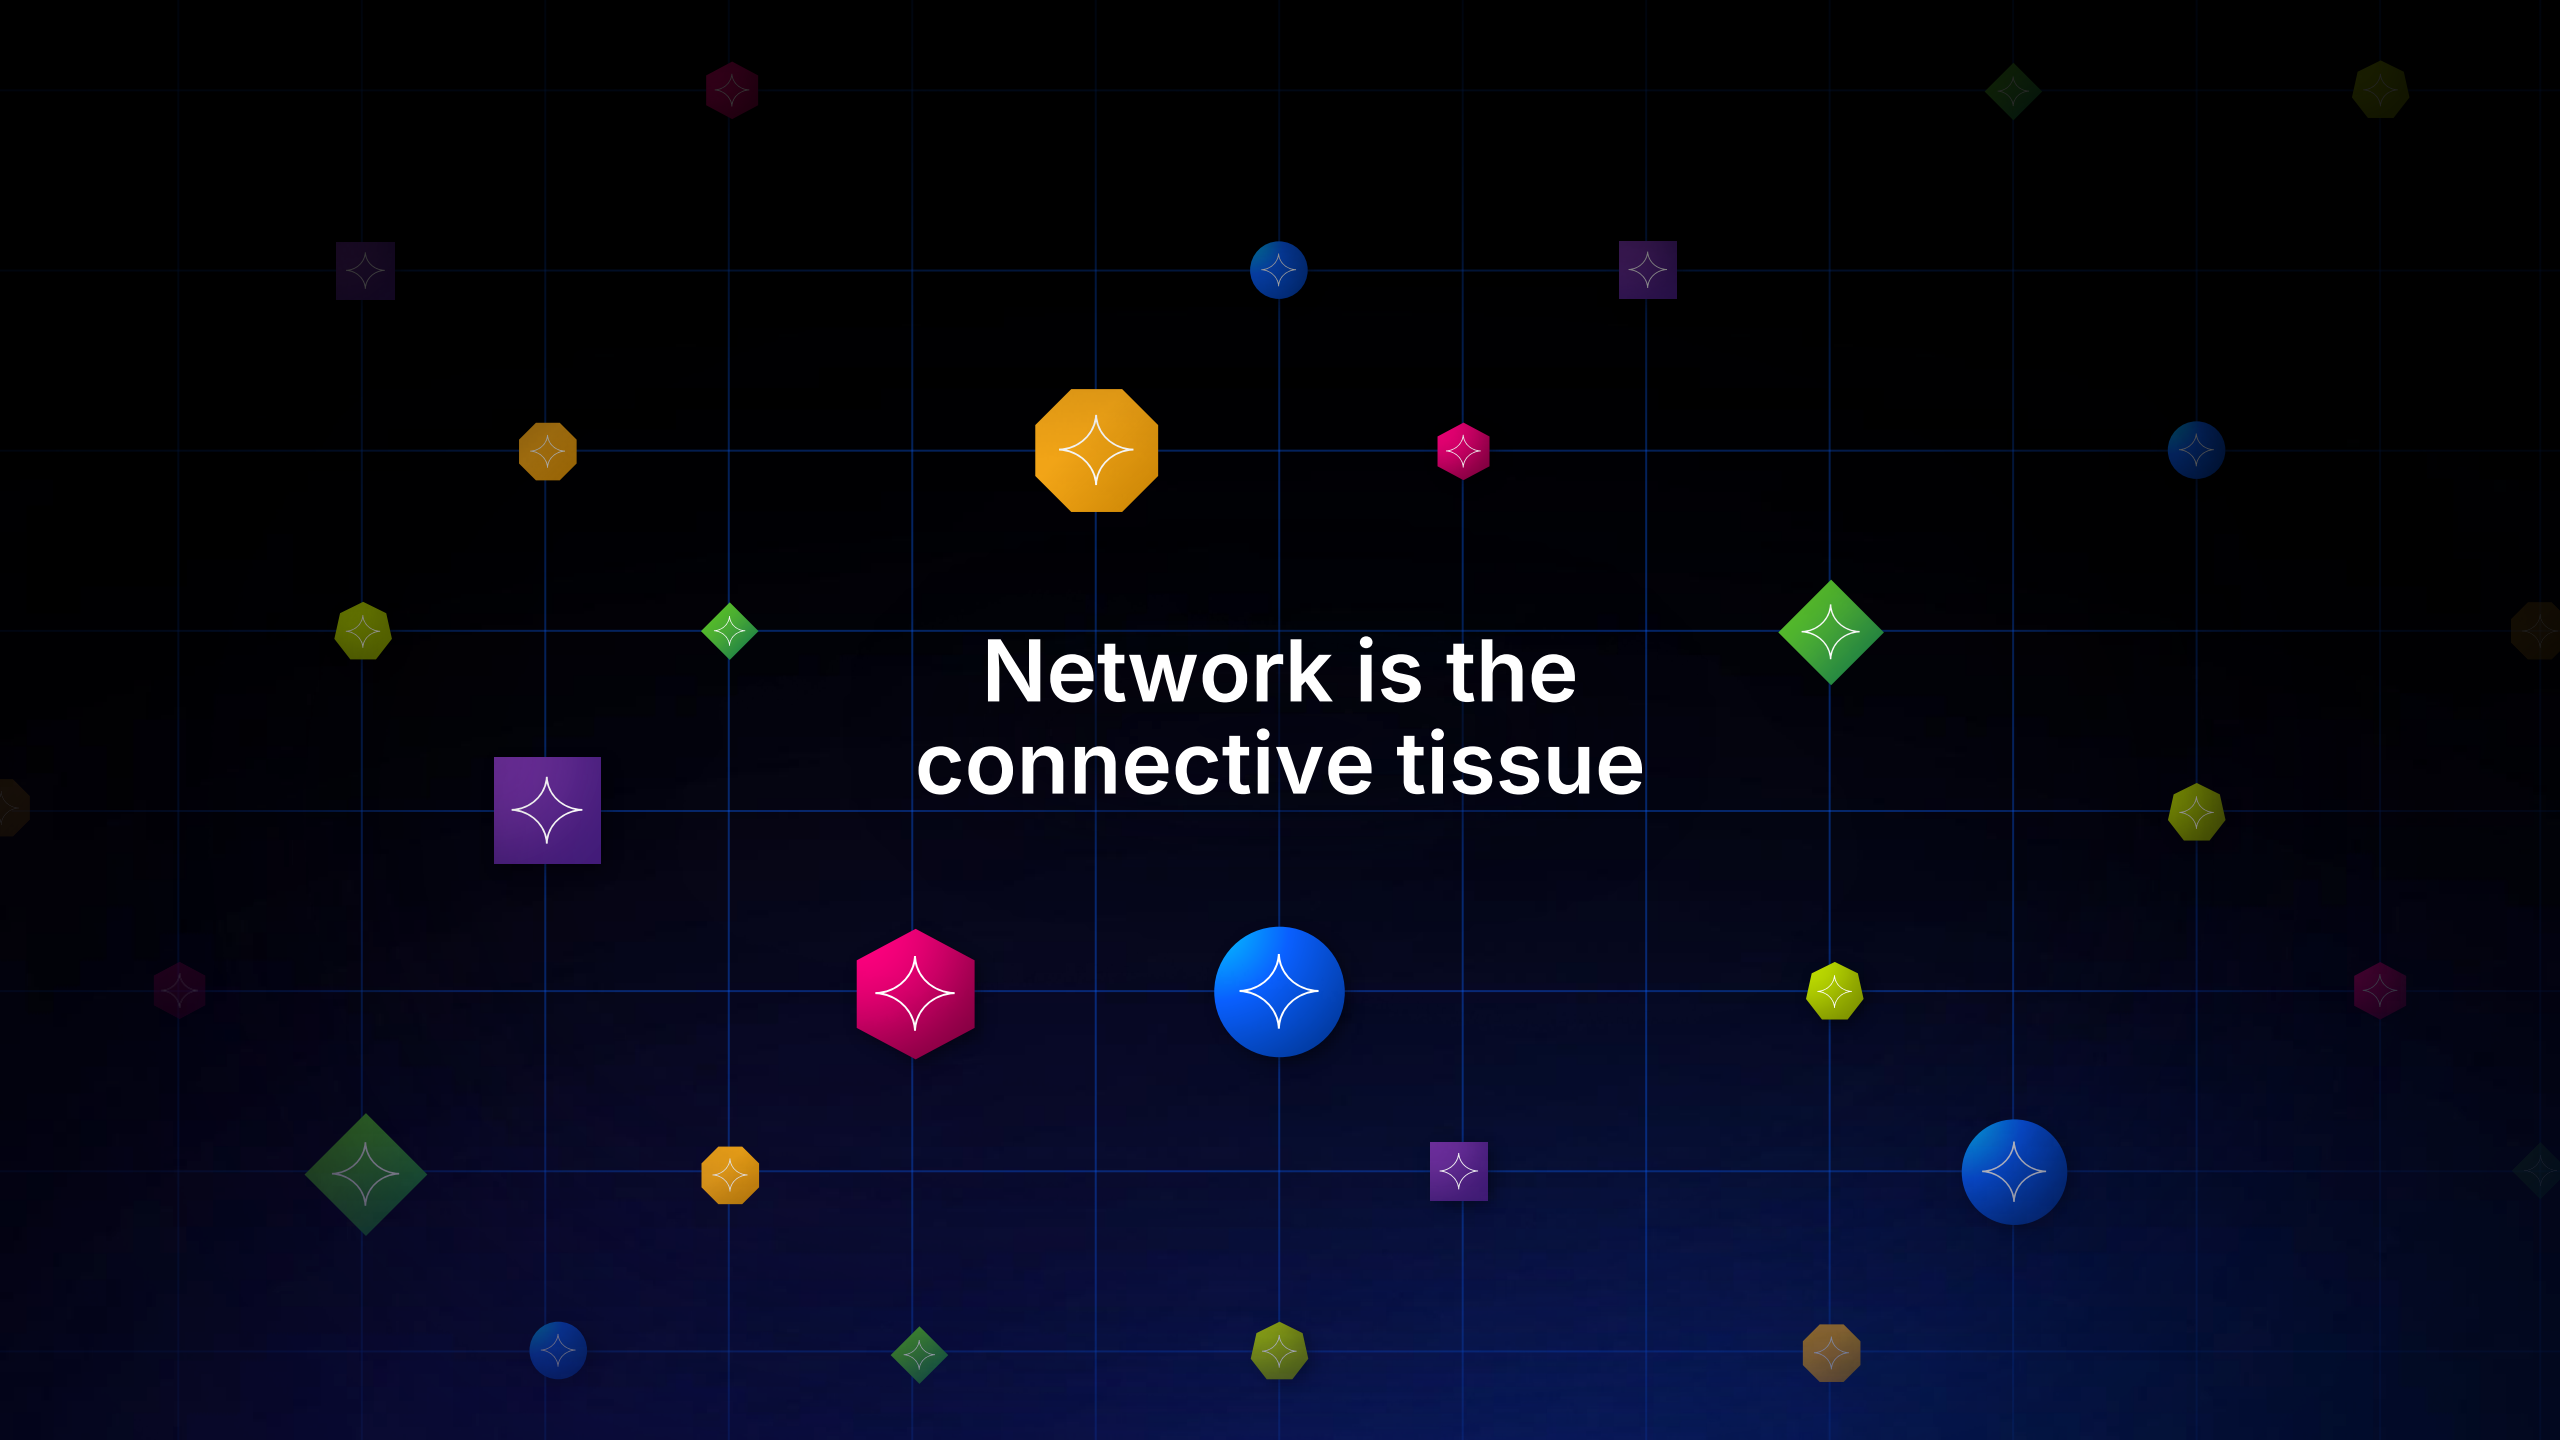
QUANTUM RESILIENCE

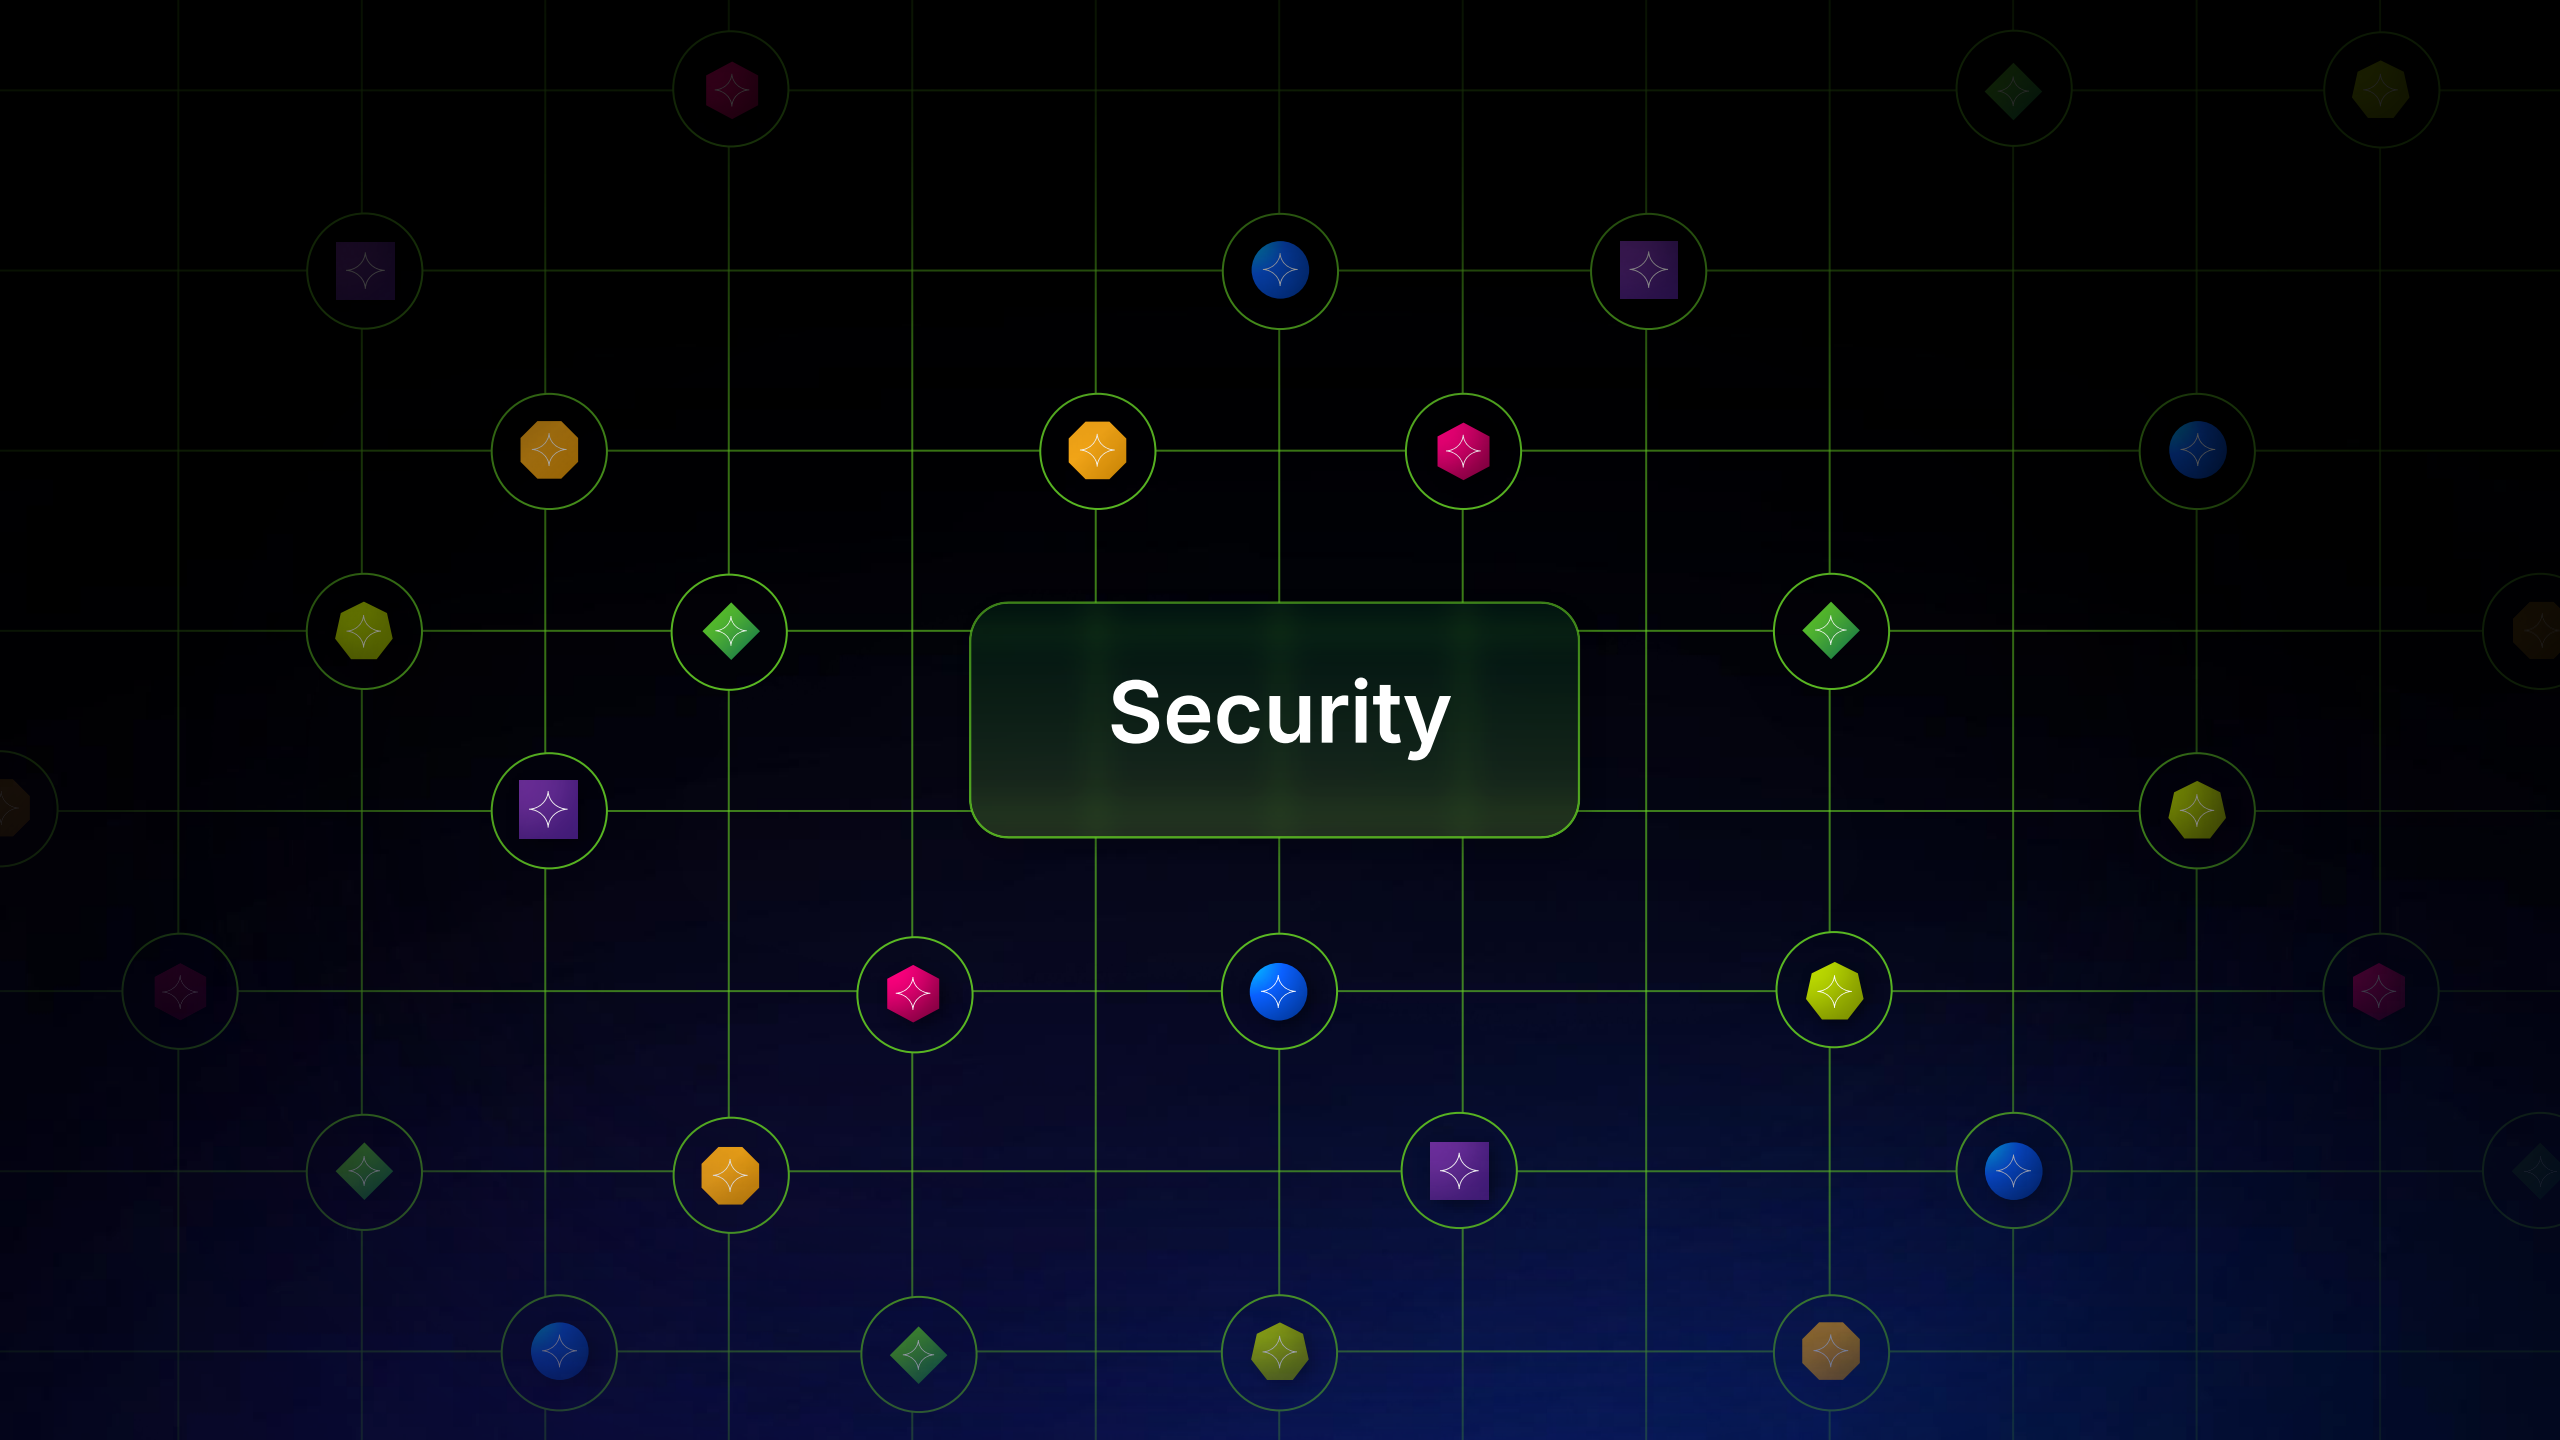
	Quantum-safe communications	Quantum-safe platforms
THREAT	Harvest Now, Decrypt Later	Trust Now, Forge Later
CISCO ACTION	Quantum-safe communications across core portfolio by June 2027	All new enterprise and data center routers, switches, and firewall series will ship with quantum-safe secure boot

SECURITY FOR AGENTS

Agents live everywhere, access everything. And are the new targets.




Network is the
connective tissue


Security

Secure Networking: Security Fused into the Network



Resilient infrastructure



Identity



Segmentation & steering



Protection



Visibility & insights

Cisco Security



Security for Agents

Protect the agents from the world



Protect the world from the agents

Detect and respond at machine speed

SECURITY FOR AGENTS: PROTECT AGENTS FROM THE WORLD

Poor uncertainty handling
Sensitive data over-sharing
Sensitive data overcollection
Workflow denial of service
Infrastructure mutation risk
Weak containment
Accountability ambiguity
Human trust manipulation
Misleading status reporting
Auditability failure
Unsafe code execution
Unsafe browser automation
Resource exhaustion
Regulatory noncompliance
Unsafe retrieval reliance
Context window loss
Poor rollback design
Offboarding failure

Excessive agency
Unauthorized action
Unsafe planning
Irreversible action
Cost amplification

Safety

Runaway task
Inadequate oversight
Transaction integrity failure
Business process corruption
Unsafe external communication



Agent identity abuse
Memory poisoning
Supply chain poisoning
Toolchain compromise
Goal hijacking

Security

Instruction hijacking
Tool misuse / abuse
Memory poisoning
Data exfiltration
Persistent instruction corruption

3rd party risk
Environment escape
Rogue agent behavior
Tool parameter injection
Tool output spoofing
Unsafe data ingestion
Credential leakage
Credential misuse
Cross-tenant data leakage
Cross-task contamination
Policy bypass
Task laundering
Approval laundering
Cascading compromise
Insecure A2A comms
Delegation abuse
Confused deputy
Covert channel exfiltration

Runtime protection for agents



Securing Custom Agents and AI Apps

Cisco AI Defense



Discover



Validate



Protect

NEW

Adaptive red teaming and custom guardrails

NEW

Protecting the full agent supply chain

NEW

Ecosystem to work with any cloud, any GPU



OPEN SOURCE

DefenseClaw

OpenClaw

Codex

Claude Code

Enterprise-ready protection in Secure Client



VPN ZTNA SIA/DNS

Visibility Posture

Experience Monitoring

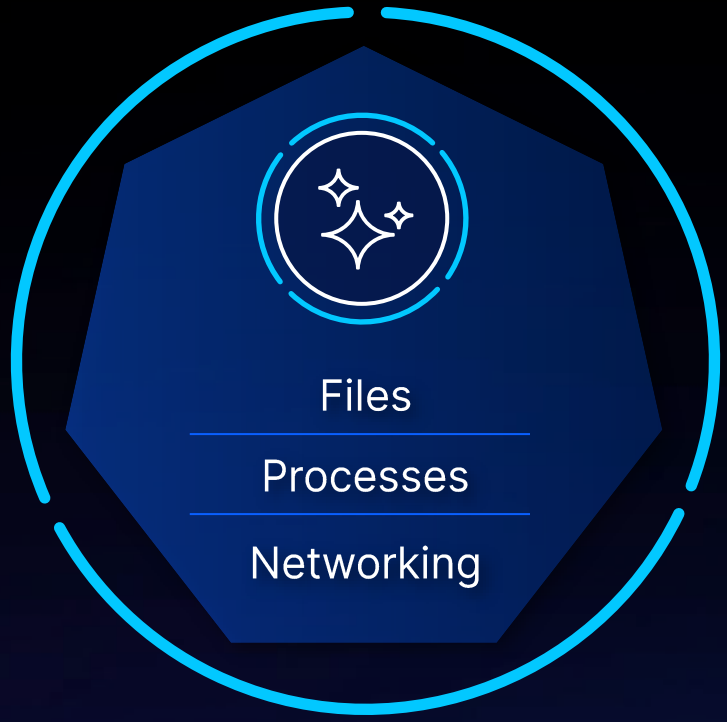
\+

DefenseClaw

AI runs on Kubernetes

Process hiding
Credential harvesting
Privilege escalation
Container escape
Kernel exploit
File tampering
Sidecar injection
Shell injection
Network scanning
Reverse shell

Segmentation



Runtime shields

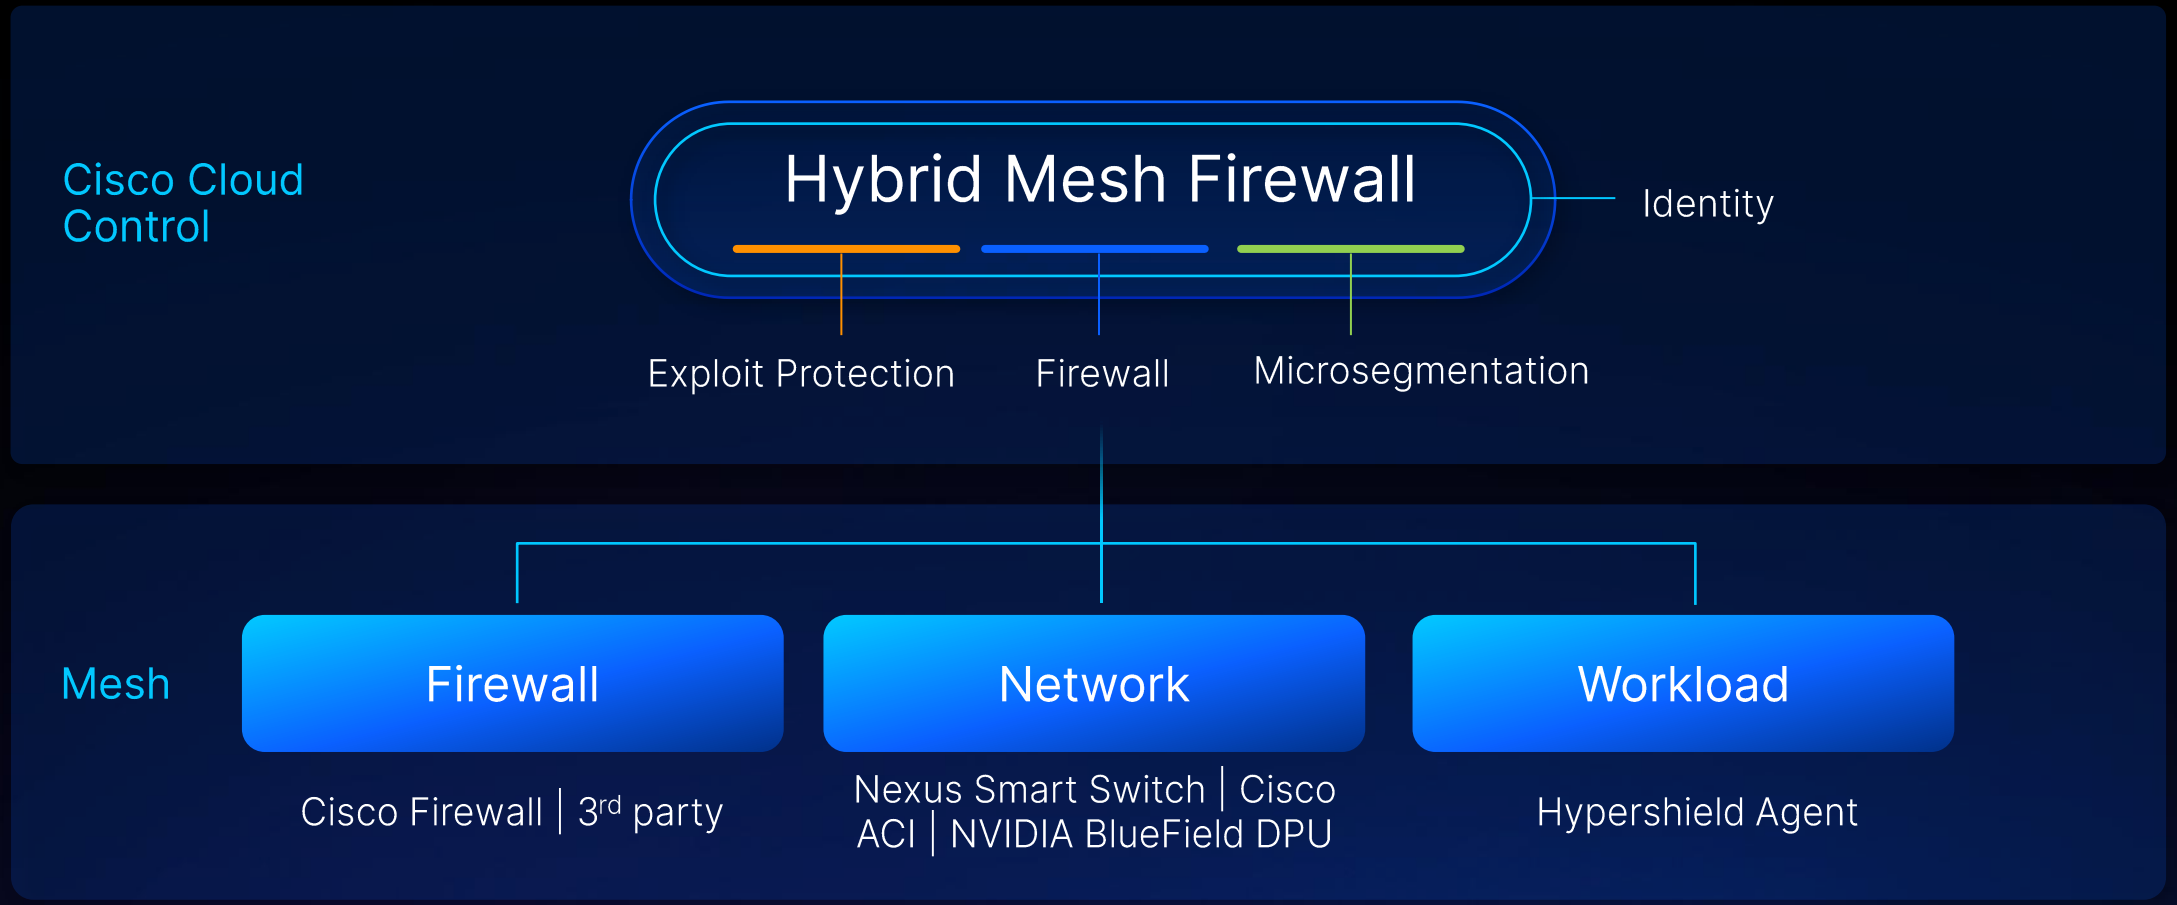
Host file system access
Service discovery abuse
API server probing
Token theft
Mail spooler abuse
DNS tunneling
Deserialization
Unauthenticated service
Runtime socket exposure
Command injection

Process-level visibility

Segmentation is no longer optional



Protecting your apps and infrastructure, including AI



Security for Agents

Protect the agents
from the world



**Protect the world
from the agents**

Detect and respond
at machine speed

The agent trust gap

85%
Experimentation and adoption

5%
Broad production





TOP CONCERNS

- Access control
- Data exfiltration
- Autonomy

Source: Cisco survey of security and IT execs, Jan 2026, n=224

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

Agents are the ‘Worst Of Both Worlds’

	Human	AI Agents	Machine
Scope	BROAD	BROAD	LIMITED
Speed	LIMITED	RAPID	RAPID
Scale	LIMITED	EXPONENTIAL	MODERATE
Sensibility	JUDGEMENT	NO JUDGEMENT	RIGID RULES

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

Zero Trust must change for the agentic workforce



Humans



Agents

Access control

Prove identity,
trust the user

Action control

Verify behavior,
control the agent

Holds agents accountable to humans | Secures agentic action

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

cisco

Chaos makes action-level zero trust hard

Agents Everywhere

Task-driven Behavior

Fragmented Enforcement

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

Zero Trust for Agents

KNOW EVERY AGENT

AUTHORIZE EVERY ACTION

ADAPT TO RISK IN REAL TIME

CONSISTENT ENFORCEMENT AT ACCESS BOUNDARY

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

Powered by Cisco SASE and Identity



Security for Agents

Protect the agents from the world



Protect the world from the agents

Detect and respond at machine speed

Mythos-class attackers have changed SOC Resilience.

Attack vectors need new approaches

Vulnerabilities

Attacks at machine speed

Phishing and Insider threat



Our Security innovations are solving for these vectors.



Vulnerabilities

Exposure Analytics and Risk Posture



Attacks at machine speed

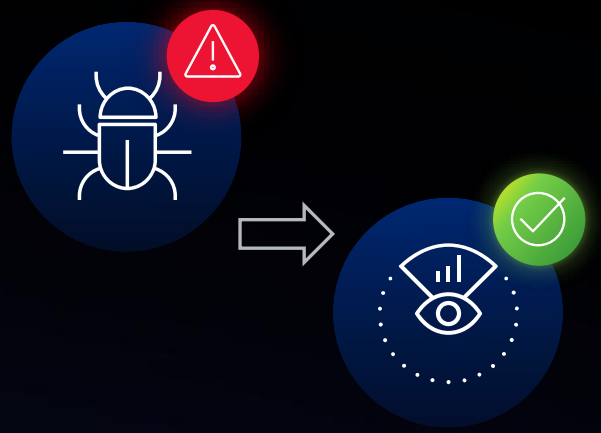
Agentic SOC



Phishing and Insider threat

Automated Phishing Threat Analysis

Expanded Visibility and Actionable Insights



- Continuously discover entities, assets, and users
- Enriched context and comprehensive visibility
- Identify vulnerabilities and exposures
- Risk posture assessment

Powering the Agentic SOC

Splunk Enterprise Security

Trusted AI

AI and Automation

- Agentic Orchestration
- Integrated Automation
- AI Assisted Experiences

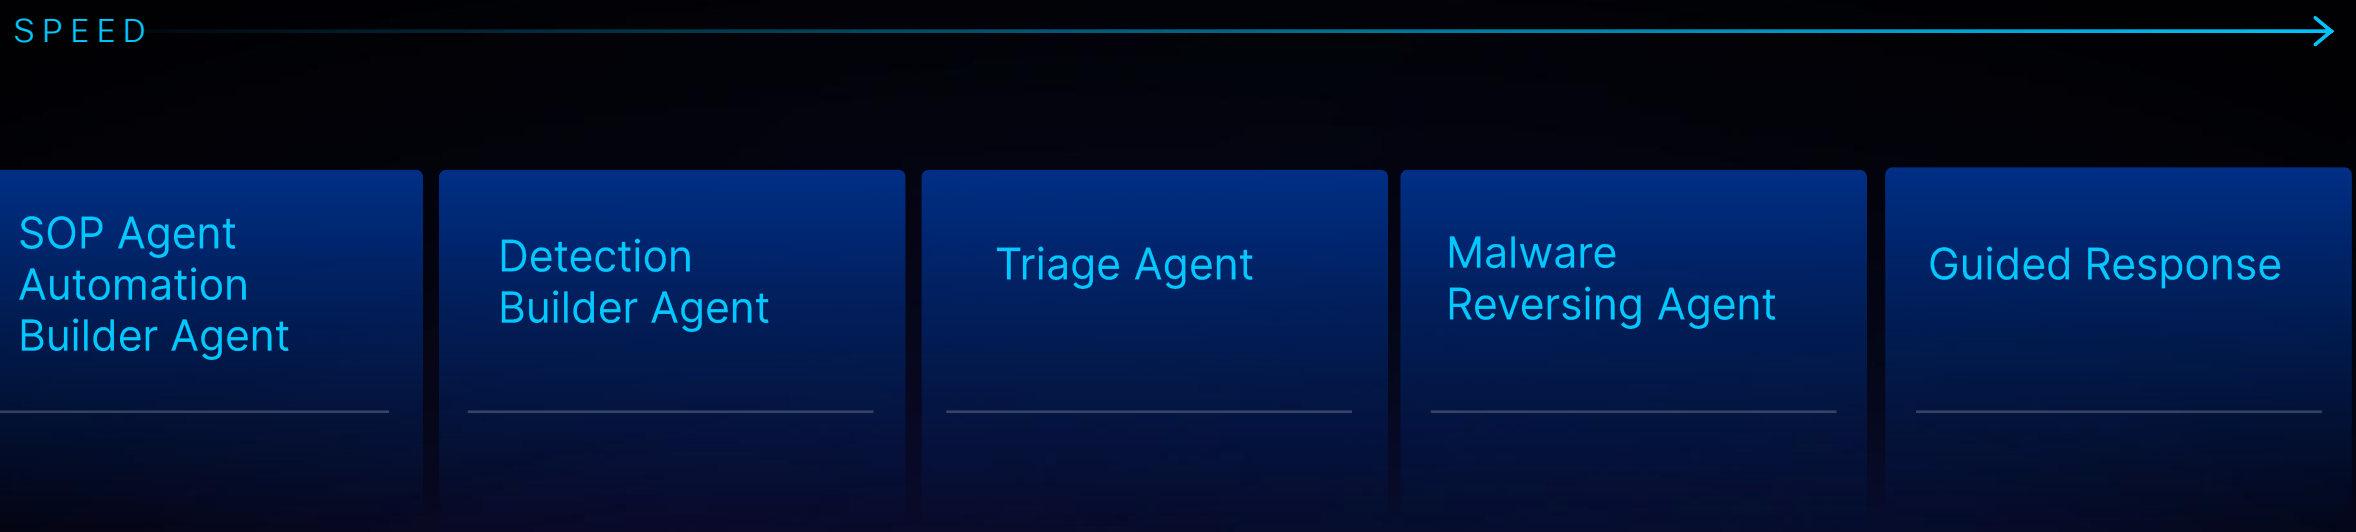
Unified Tooling

- Threat Detection
- Investigation
- Response

Open Data Platform

Cisco Data Fabric

Security AI Agents Built to Detect and Respond at Machine Speed





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

Observability for AI

AVAILABLE NOW



AI infrastructure monitoring

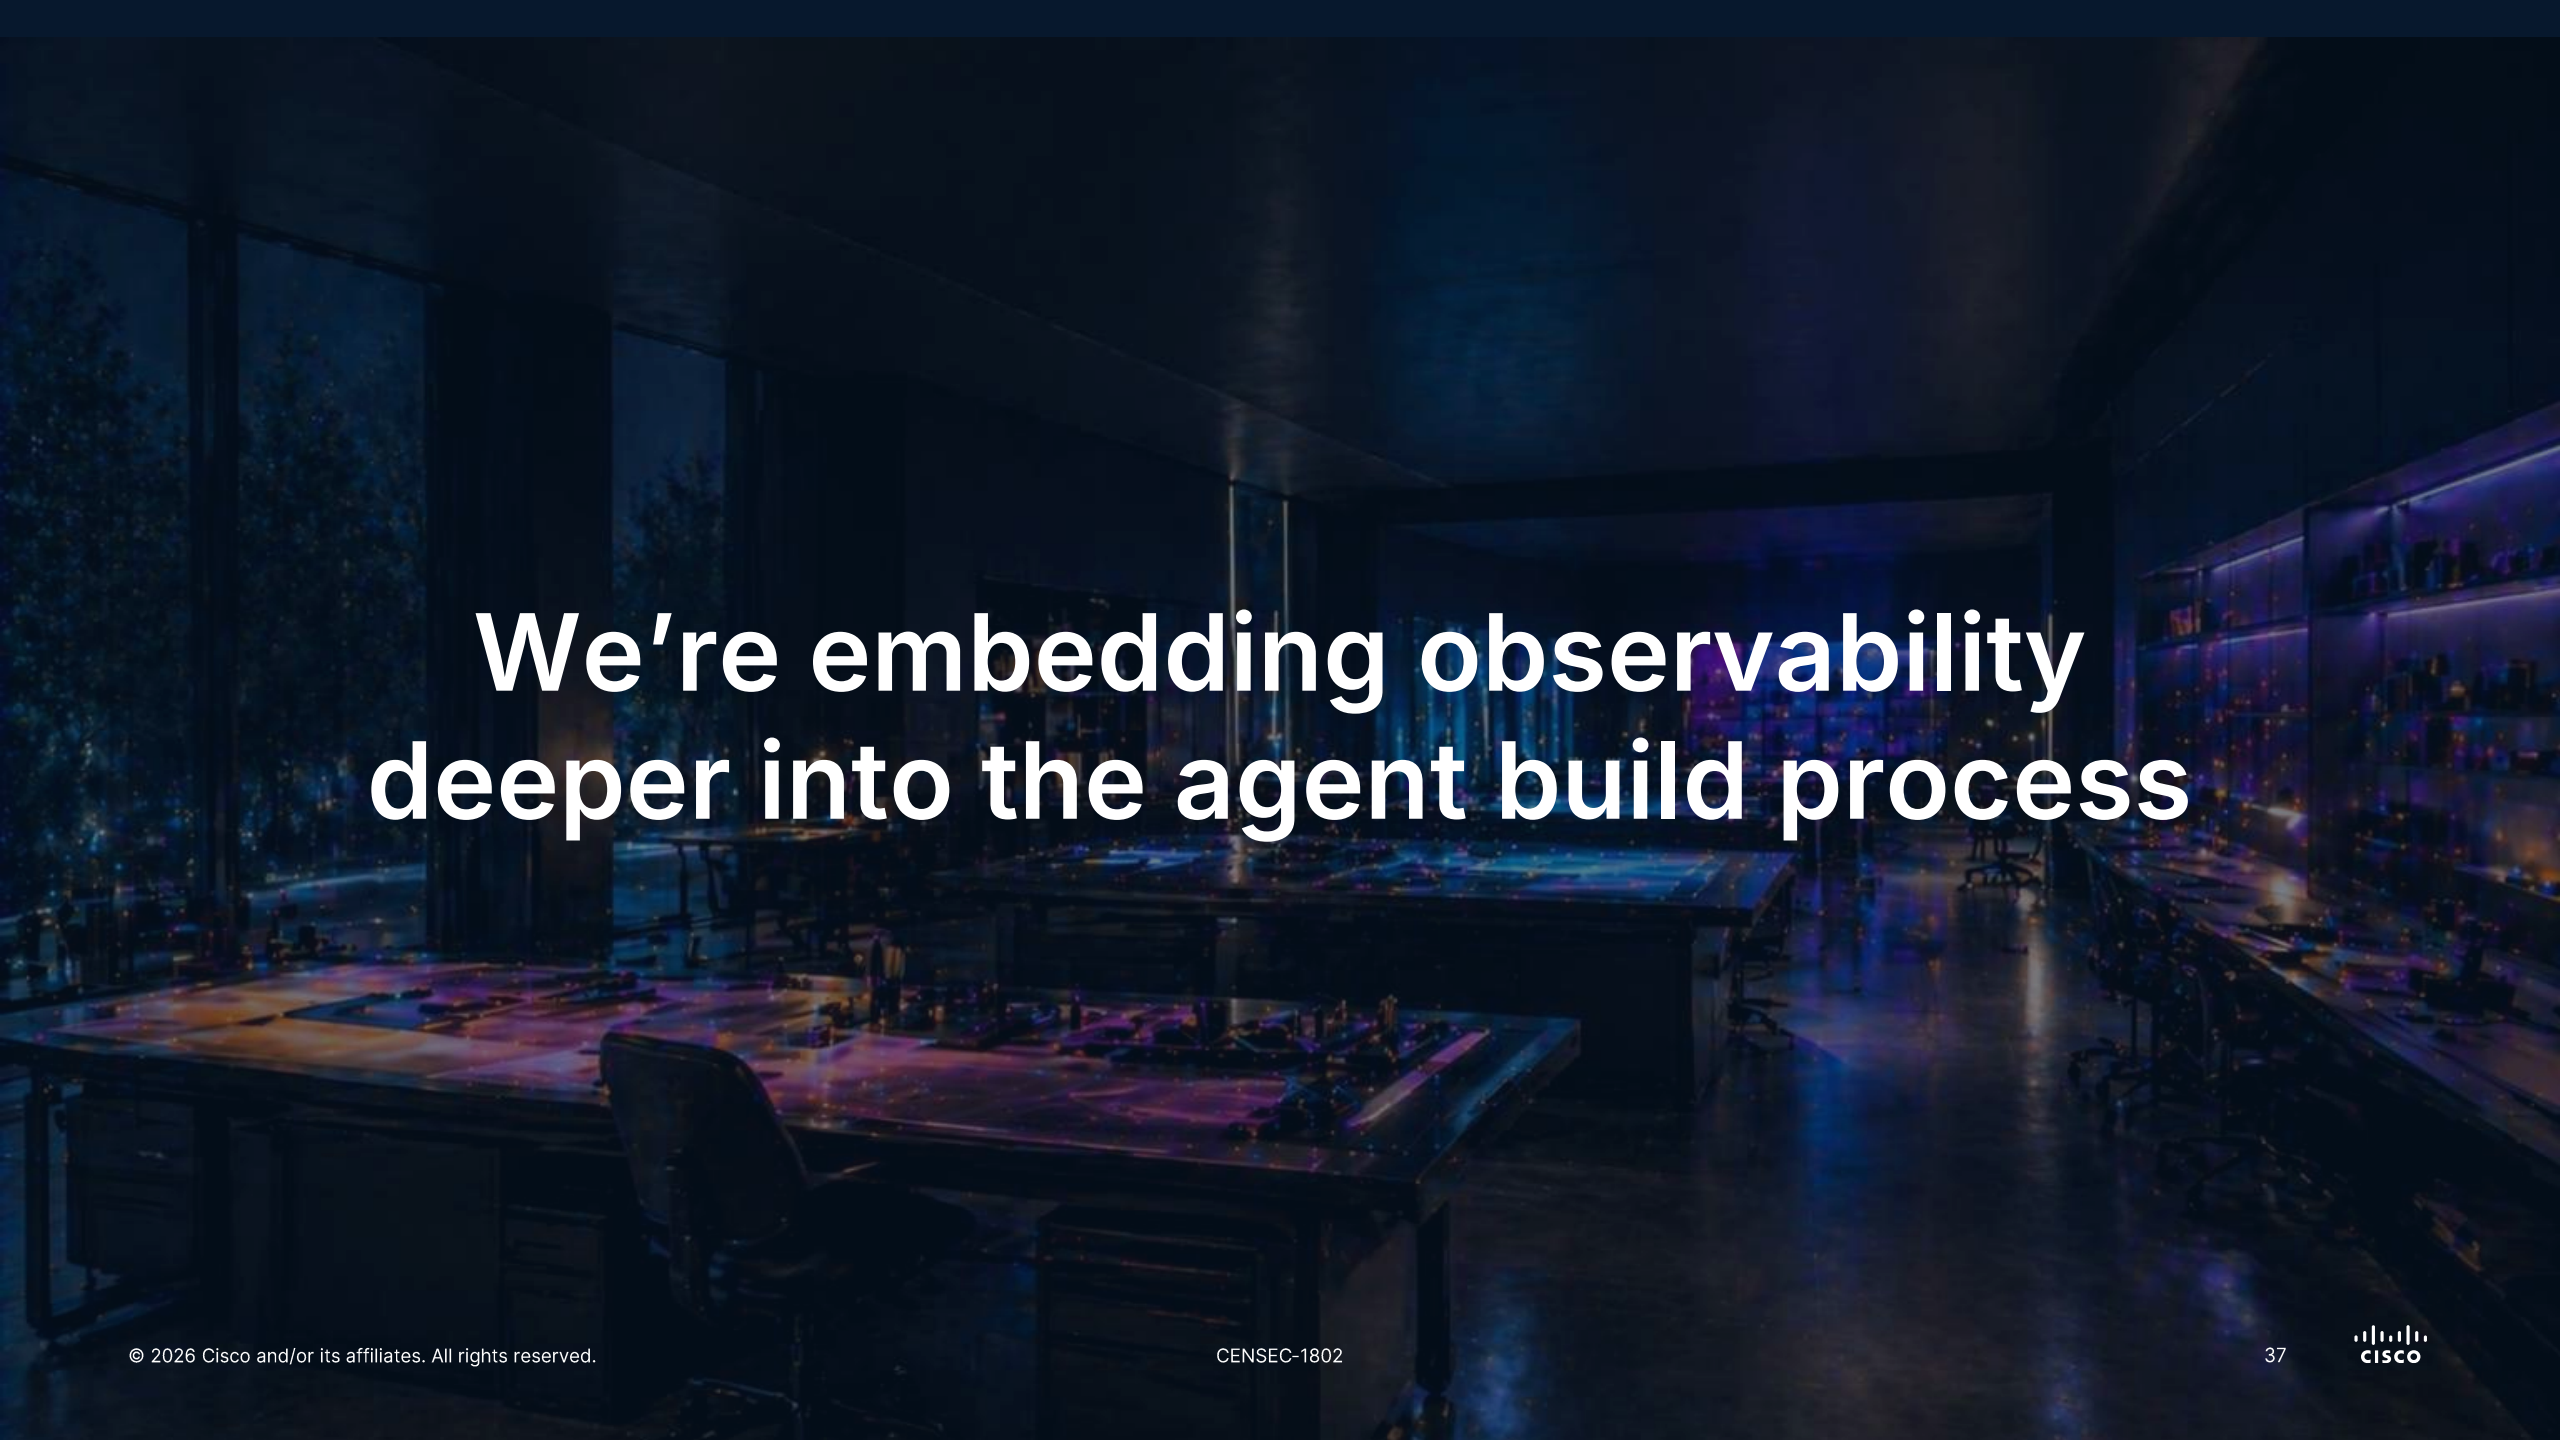
Understand operational metrics and tokenomics



AI agent monitoring

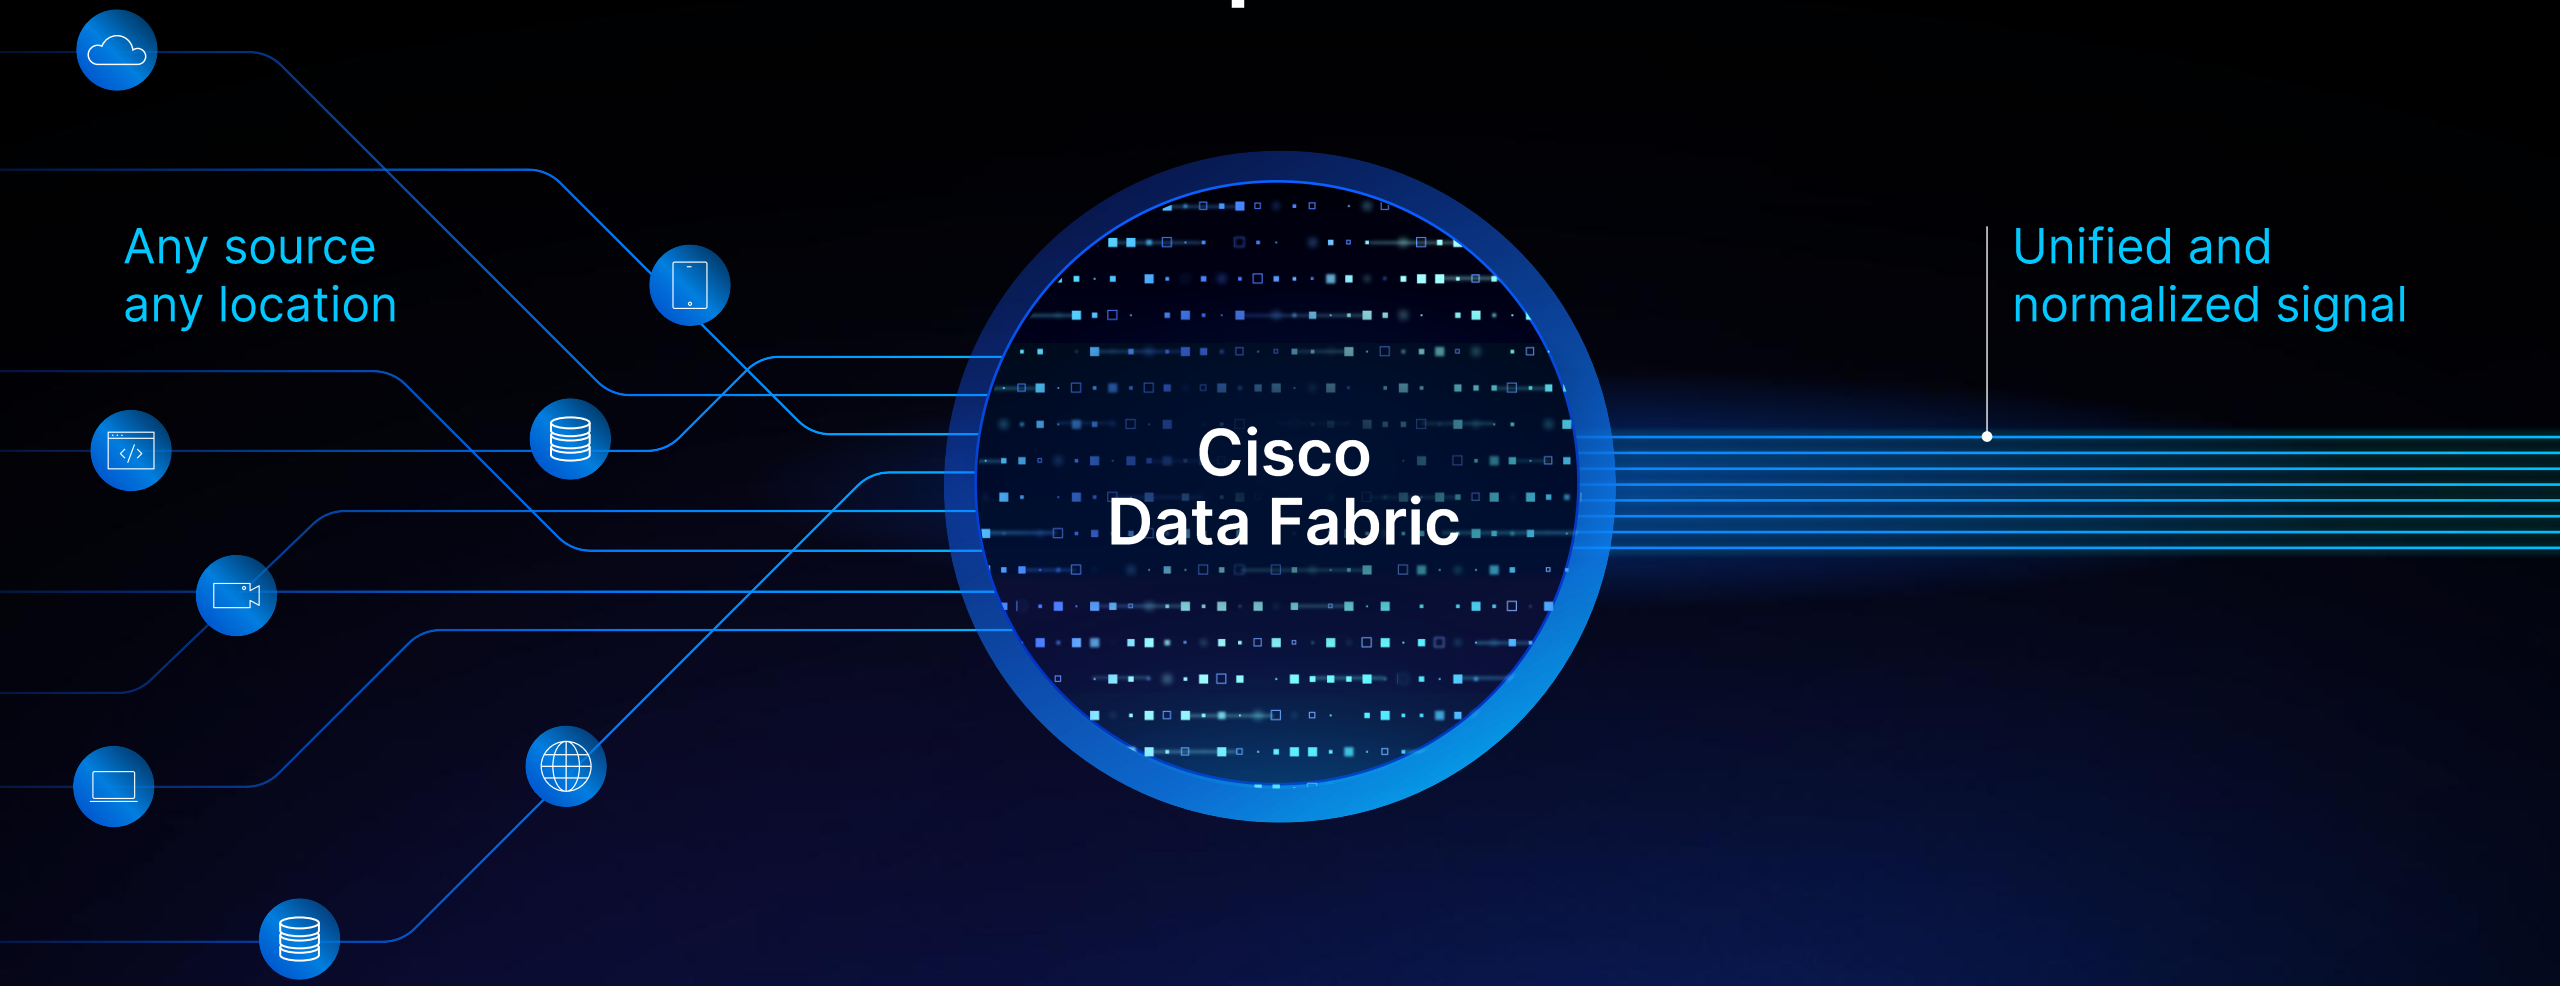
Understand behavior and quality of outcomes

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

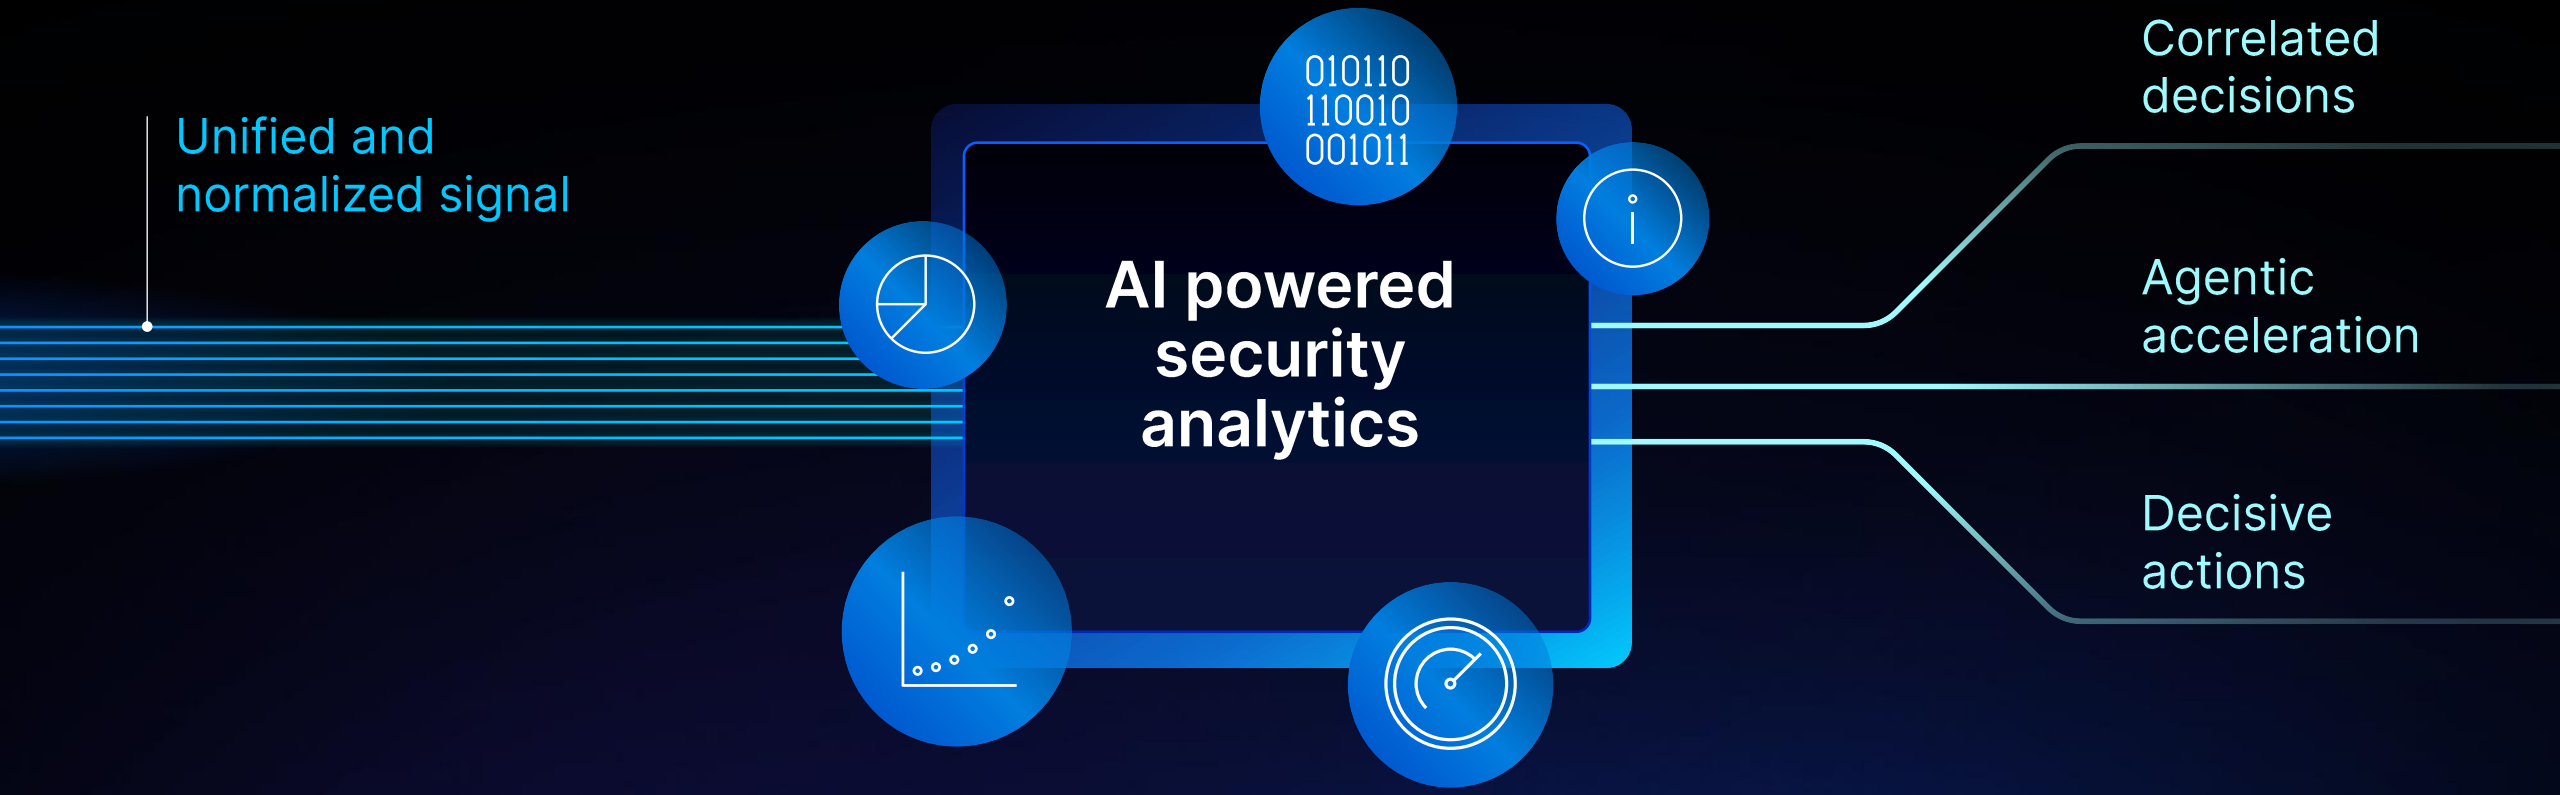


© 2026 Cisco and/or its affiliates. All rights reserved.

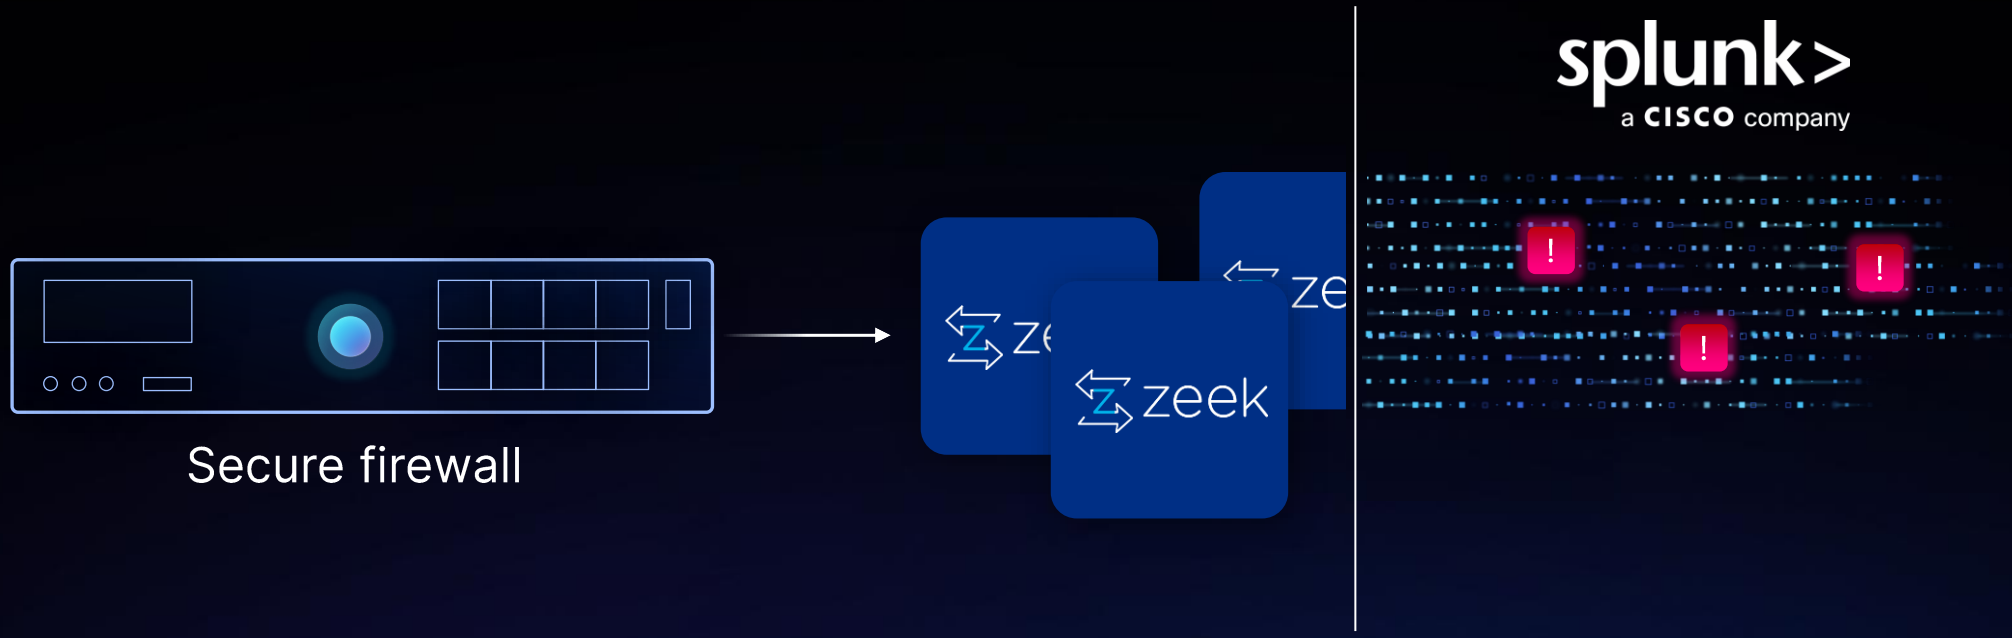
Agentic AI demands data visibility at machine speed and scale



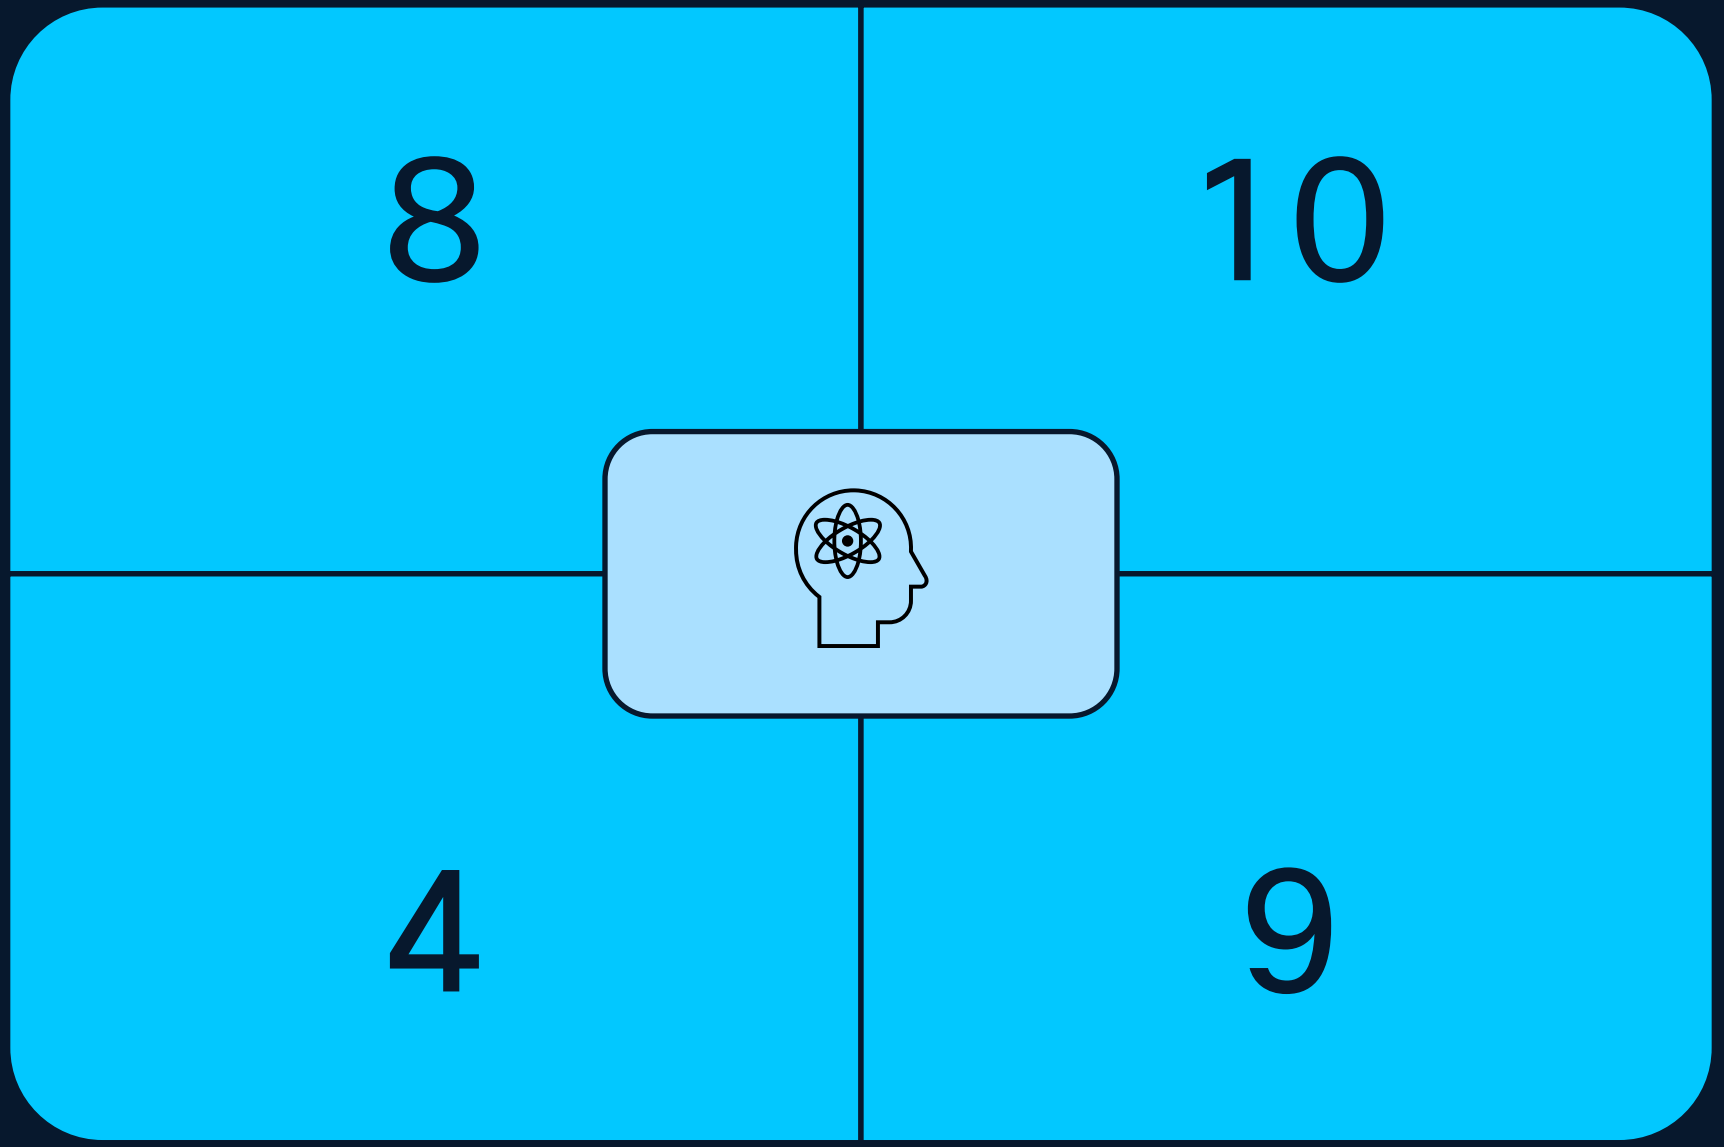
Transforming data into context, clarity, and decisive action



Splunk detections on advanced firewall logs **(it all starts with visibility)**



Which one of these is not like the other?



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

Cisco Confidential

CISCO Connect

Thank you




CISCO