Navigating the Future of Cybersecurity: AI, Quantum-Resistant Cryptography and Zero Trust

Cindy Green-Ortiz
Cisco Principal Security Architect
CISSP | CISM | PMP | CSM | CSSLP | CRISC

Cisco Live!

IT Leadership Program

Agenda

Security Threat Landscape

Exploring the Paradigm Shift in Security

This session will explore the dynamic intersection of artificial intelligence, post-quantum cryptography, and zero trust architecture.

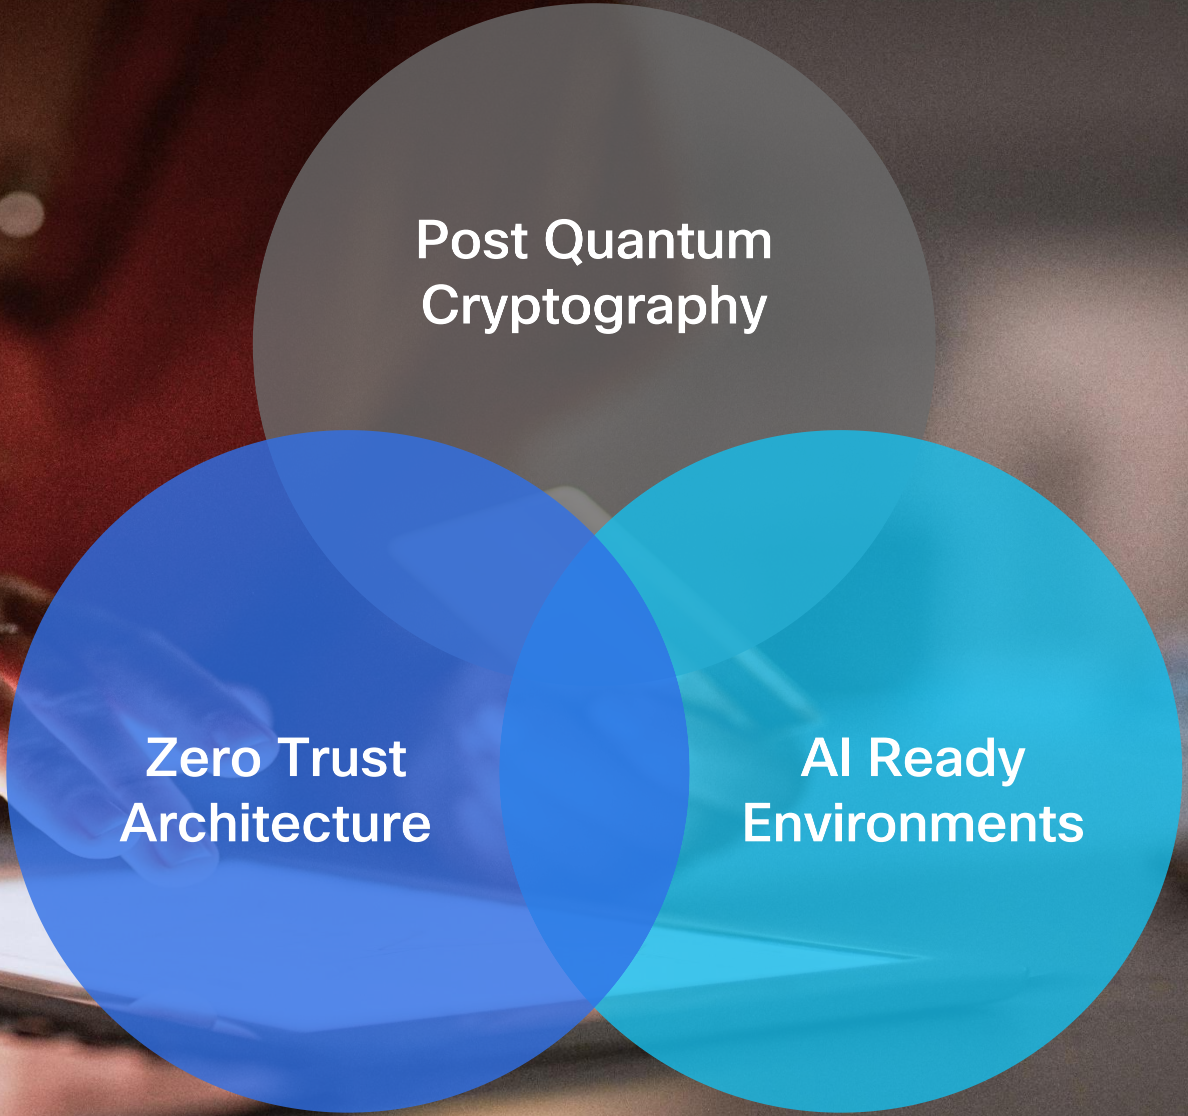

Security Paradigm Shift

Ransomware-as-a-Service (RaaS)
- **TTPs/IOCs**: LockBit, BlackCat/ALPHV; phishing, exposed RDP, CVE exploitation, double extortion.
- **Leader View**: Continues to be the most disruptive operational risk, affecting business continuity, data integrity, and customer trust.

Business Email Compromise (BEC), Now Enhanced by AI
- **TTPs/IOCs**: Executive impersonation, deepfake audio/video, MFA token theft, invoice fraud.
- **Leader View**: Financial fraud and internal trust erosion are increasing as attackers use AI to manipulate executive communications.

Zero-Day Exploits of Critical CVEs
- **TTPs/IOCs**: CVE-2024-21893 (Ivanti), CVE-2023-23397 (Outlook); rapid exploitation post-disclosure.
- **Leader View**: Speed of exploitation leaves narrow windows to respond; IT operations and vulnerability management must be agile.

Prompt Injection & Adversarial LLM Exploits
- **TTPs/IOCs**: Manipulated prompts override AI model controls, enabling data leakage or system abuse.
- **Leader View**: Enterprise AI tools can be subverted to disclose sensitive information or act against policy.

AI-Crafted Malware and Automated Payload Generation
- **TTPs/IOCs**: Obfuscated and polymorphic malware created with AI tools; adaptive to defenses.
- **Leader View**: AI is lowering the barrier for sophisticated attacks, increasing both threat volume and speed.

AI-Enhanced Phishing & Deepfake Social Engineering
- **TTPs/IOCs**: Multilingual phishing, deepfake impersonation of executives, highly targeted spear phishing.
- **Leader View**: Social engineering attacks are now tailored, scalable, and far harder for humans to detect.

Supply Chain Threat: Backdoor (CVE-2024-3094)
- **TTPs/IOCs**: Compromised open-source software injected into core Linux tooling.
- **Leader View**: Third-party software trust is eroding; organizations must audit the origins and integrity of dependencies.

Enhanced Attack Automation via AI
- **TTPs/IOCs**: Autonomous exploitation, real-time attack coordination, AI-modulated payloads.
- **Leader View**: Threat campaigns are accelerating beyond traditional SOC response speeds; AI-powered defense is becoming essential.

Code Injection via AI Development Tools
- **TTPs/IOCs**: Insecure defaults or vulnerable logic inserted by AI-assisted coding platforms.
- **Leader View**: Rapid development cycles powered by AI can introduce critical vulnerabilities if not properly governed.

Quantum Computing Risks to Encryption (Harvest Now, Decrypt Later)
- **TTPs/IOCs**: Encrypted data theft now for future decryption by quantum systems.
- **Leader View**: Sensitive data needs long-term confidentiality protection; post-quantum cryptography planning is urgent.

Data Poisoning and ML Model Manipulation
- **TTPs/IOCs**: Malicious data used to influence AI behavior or embed hidden logic.
- **Leader View**: Corrupted AI systems could make faulty decisions, leak data, or be covertly controlled by attackers.

Enterprise AI Hijacking
TTPs/IOCs: Internal chatbots and AI tools abused for lateral movement or data extraction.
Leader View: Internal AI systems must be treated as high-value assets subject to misuse or compromise.

Adversarial AI & GAN Weaponization
TTPs/IOCs: AI designed to deceive or evade other AI systems (e.g., anti-spam, image analysis).
Leader View: Expect AI-on-AI adversarial activity to increase, especially in detection and fraud prevention domains.

Escalating AI-Driven Attack Velocity
TTPs/IOCs: AI tools supporting rapid exploitation, adaptive targeting, and persistent access.
Leader View: The threat tempo is outpacing human detection—security architecture must be designed for speed.

Vector Database Encryption Gaps
- **TTPs/IOCs**: Unencrypted AI embedding stores and vector search engines vulnerable to data leaks.
- **Leader View**: Sensitive data powering AI applications is at risk if encryption-in-use is not addressed.

LEGEND: Current / Persistent | Emerging (3 months) | Horizon (6-12 months)

In the News

Google AI Overviews Incorrectly States Current Year

▪ Summary: Google's AI Overviews feature has been providing incorrect information regarding the current year, confidently asserting that it is still 2024, despite the actual date being May 29, 2025. This error has been consistently reproduced, with the AI citing various sources, including Reddit and Wikipedia, leading to confusion among users.
▪ Author: Reece Rogers
▪ Date & Time: May 29, 2025, 4:21 PM
▪ Risk Score: MEDIUM
▪ Threat Category: Emerging
▪ TTPs/IOCs: Misinformation propagation via AI-generated content; reliance on unverified sources
▪ Why Sales Cares: Clients utilizing AI-driven tools for information retrieval may be exposed to inaccurate data, potentially affecting decision-making processes and trust in AI solutions.
▪ Why Leaders Care: The dissemination of incorrect information by AI systems can undermine organizational credibility and highlights the need for robust validation mechanisms in AI deployments.
▪ Source: https://www.wired.com/story/google-ai-overviews-says-its-still-2024/

Research Uncovers Exposed DeepSeek Database Leaking Sensitive Information, Including Chat History

A publicly accessible database belonging to DeepSeek allowed full control over database operations, including the ability to access internal data. The exposure includes over a million lines of log streams with highly sensitive information.
January 29, 2025, https://www.wiz.io/blog/wiz-research-uncovers-exposed-deepseek-database-leak

U.S. breakthroughs in superintelligence targeted for theft by China's spy network

"Right now, the greatest danger is not that the U.S. will fall behind China in the race to superintelligence. Until we've secured the labs, there is no lead for us to lose," the duo wrote in a report released Tuesday. [...] One lab's researcher told Gladstone AI that a running joke inside the team was that they were "the leading Chinese AI lab because probably all of our [stuff] is being spied on."
April 22, 2025 https://www.washingtontimes.com/news/2025/apr/22/us-breakthroughs-superintelligence-targeted-theft-chinas-spy-network/

Post Quantum Resistant Cryptography

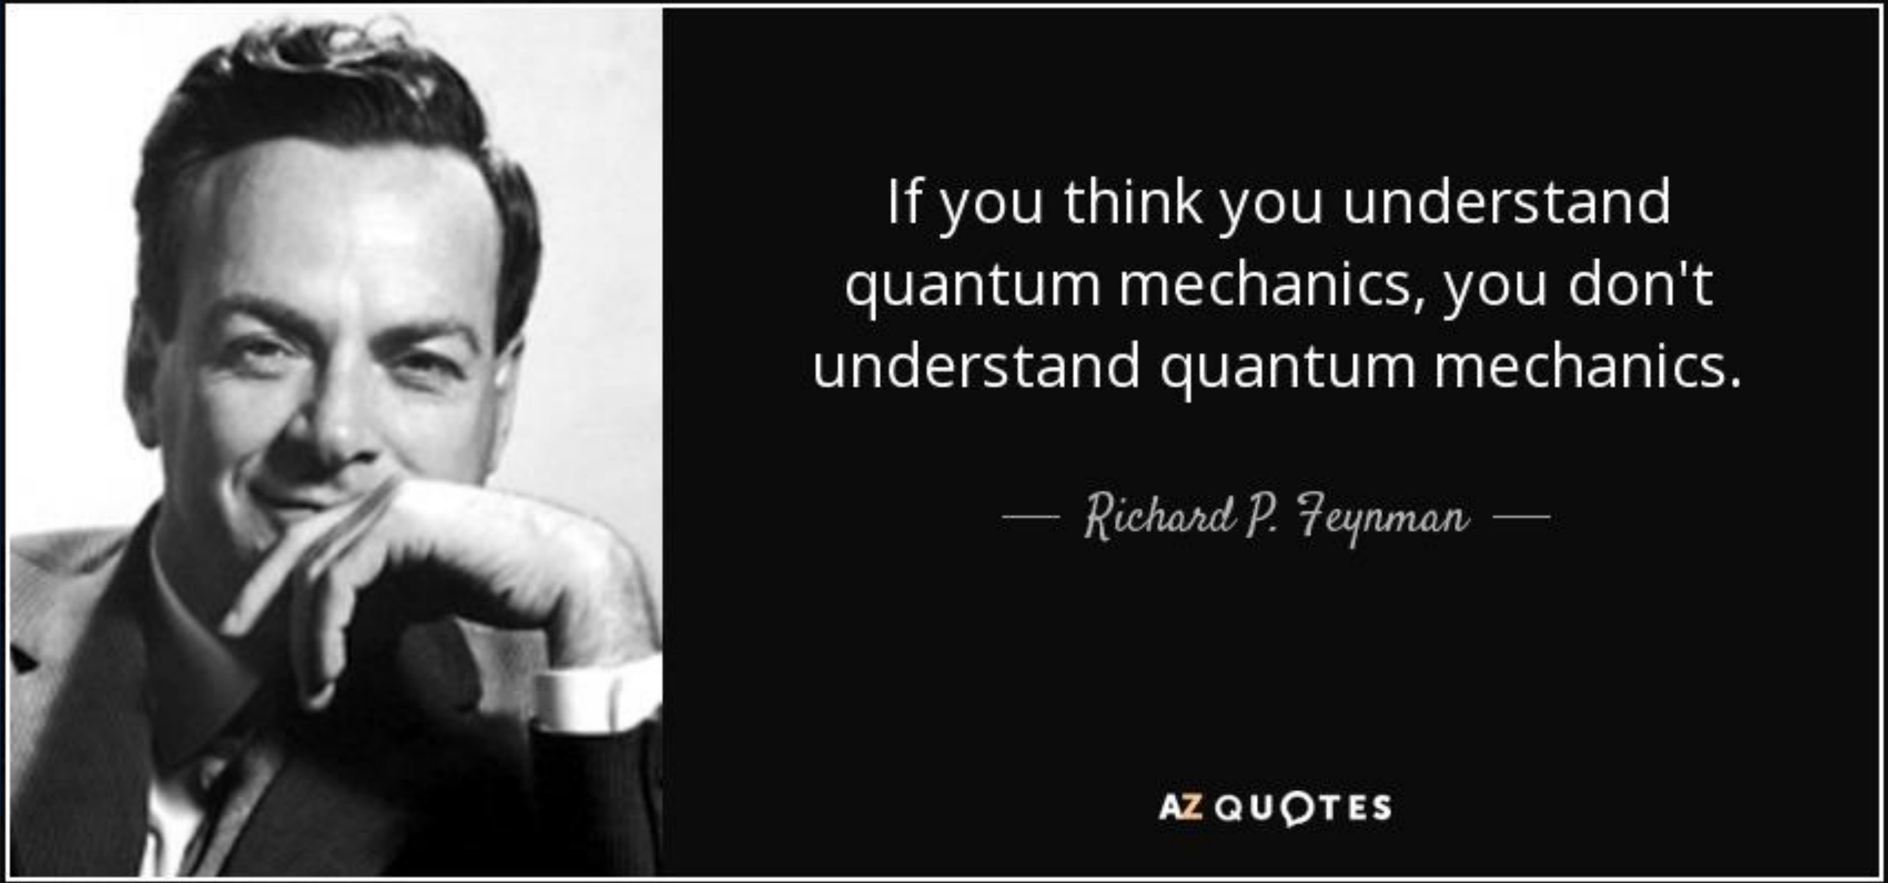

Richard Phillips Feynman (May 11, 1918 – February 15, 1988) His most public achievement came in 1965, when he won the Nobel Prize in Physics, sharing it with Julian Schwinger and Shin'ichiro Tomonaga for their independent work in quantum electrodynamics. https://www.richardfeynman.com/about/bio.html

What is Quantum Computing?

Processes information that uses qubits. While classical computers use bits to store data as either a zero or a one, qubits can store both at the same time, thanks to a property called superposition, allowing quantum computers to handle much more information at once.

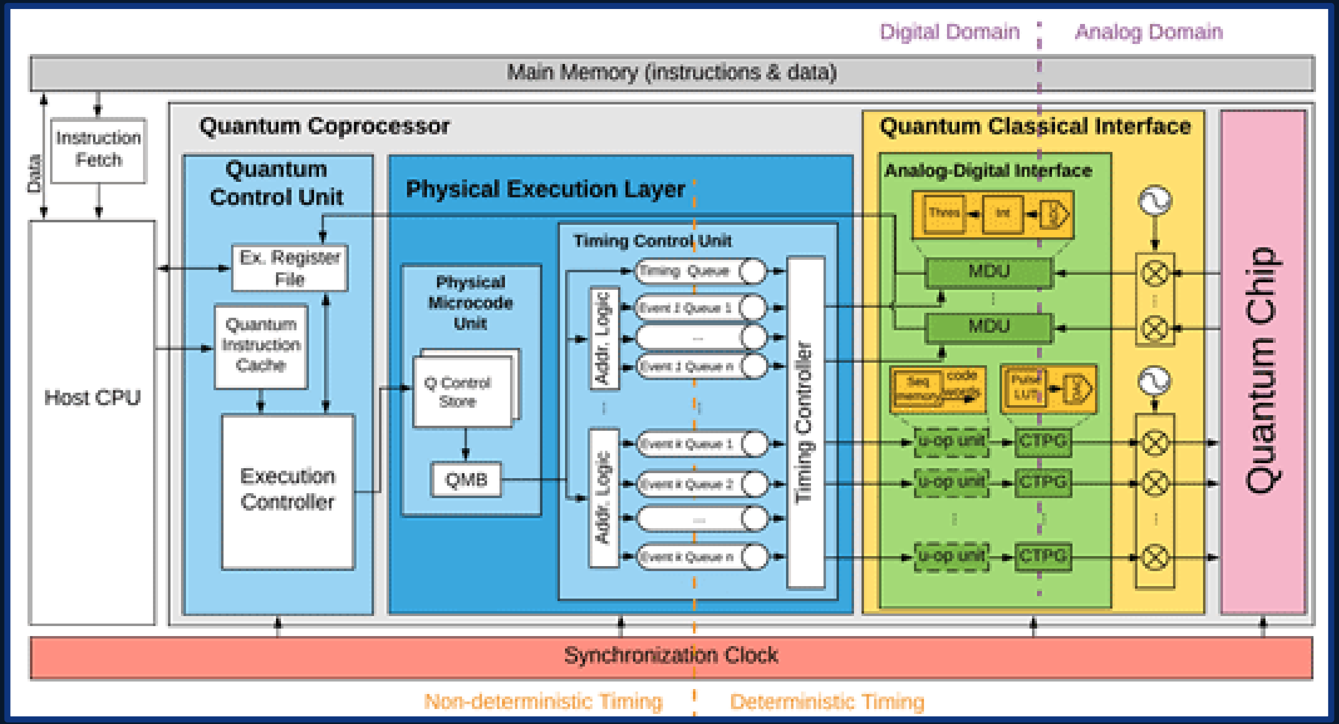

"Classical Computing is akin to the performance of a 100 baud modem (i.e., 0.0000125 Mbps) from the 1980s, compared to the performance of Quantum Computing"
- Cisco Quantum Summit 1/30/2025

What makes Quantum Computing so fast?

Superposition: when a quantum particle can exist in multiple states at once, rather than just one.
Entanglement: occurs when quantum particles become linked, so the state of one can instantly affect the state of another, no matter the distance between them.
Decoherence: when quantum particles lose their quantum state and settle into a single state that can be measured by classical physics.
Interference: when quantum states interact with each other, affecting the probabilities of different outcomes.

Cisco Quantum Summit Replay Here:
https://www.youtube.com/@OutshiftbyCisco/playlists

Quantum Terminology

Quantum-resistant	▪ Focuses specifically on resisting quantum attacks, often tied to specific algorithms. Algorithms that are believed to resist quantum attacks.
Quantum-safe	▪ Encompasses a holistic assurance of security against quantum threats, potentially including systems and implementations. Quantum Key Distribution (QKD) can be included here.
Post-Quantum Cryptography (PQC)	▪ A cryptographic discipline developing algorithms for a quantum future. NIST standardized algorithms (FIPS 203, 204, 205) sit here. QKD does not.
Hybrid (PQ/T)	▪ Refers to the combination of traditional cryptographic algorithms (such as Diffie-Hellman) with post-quantum cryptographic (PQC) algorithms (like ML-KEM). Both methods run side by side and their outputs are combined. Post-Quantum/Traditional
Crypto Agility	▪ The ability of a cryptographic system to easily switch between encryption algorithms or key sizes, ensuring adaptability to new security needs

Y2K (Year 2000):

we knew exactly when,
but did not know what we missed

VS

Q-Day (Quantum Day):

**we know what will happen,
but do not know when**

Q-Day, short for **Quantum Day,** is a term used to denote the fast-approaching future date when quantum computers will become powerful enough to break the now-widely used cryptographic algorithms. Q-Day will have profound implications for cybersecurity, as current encryption methods like **RSA** and **ECC** will be rendered obsolete.

Race to Supremacy

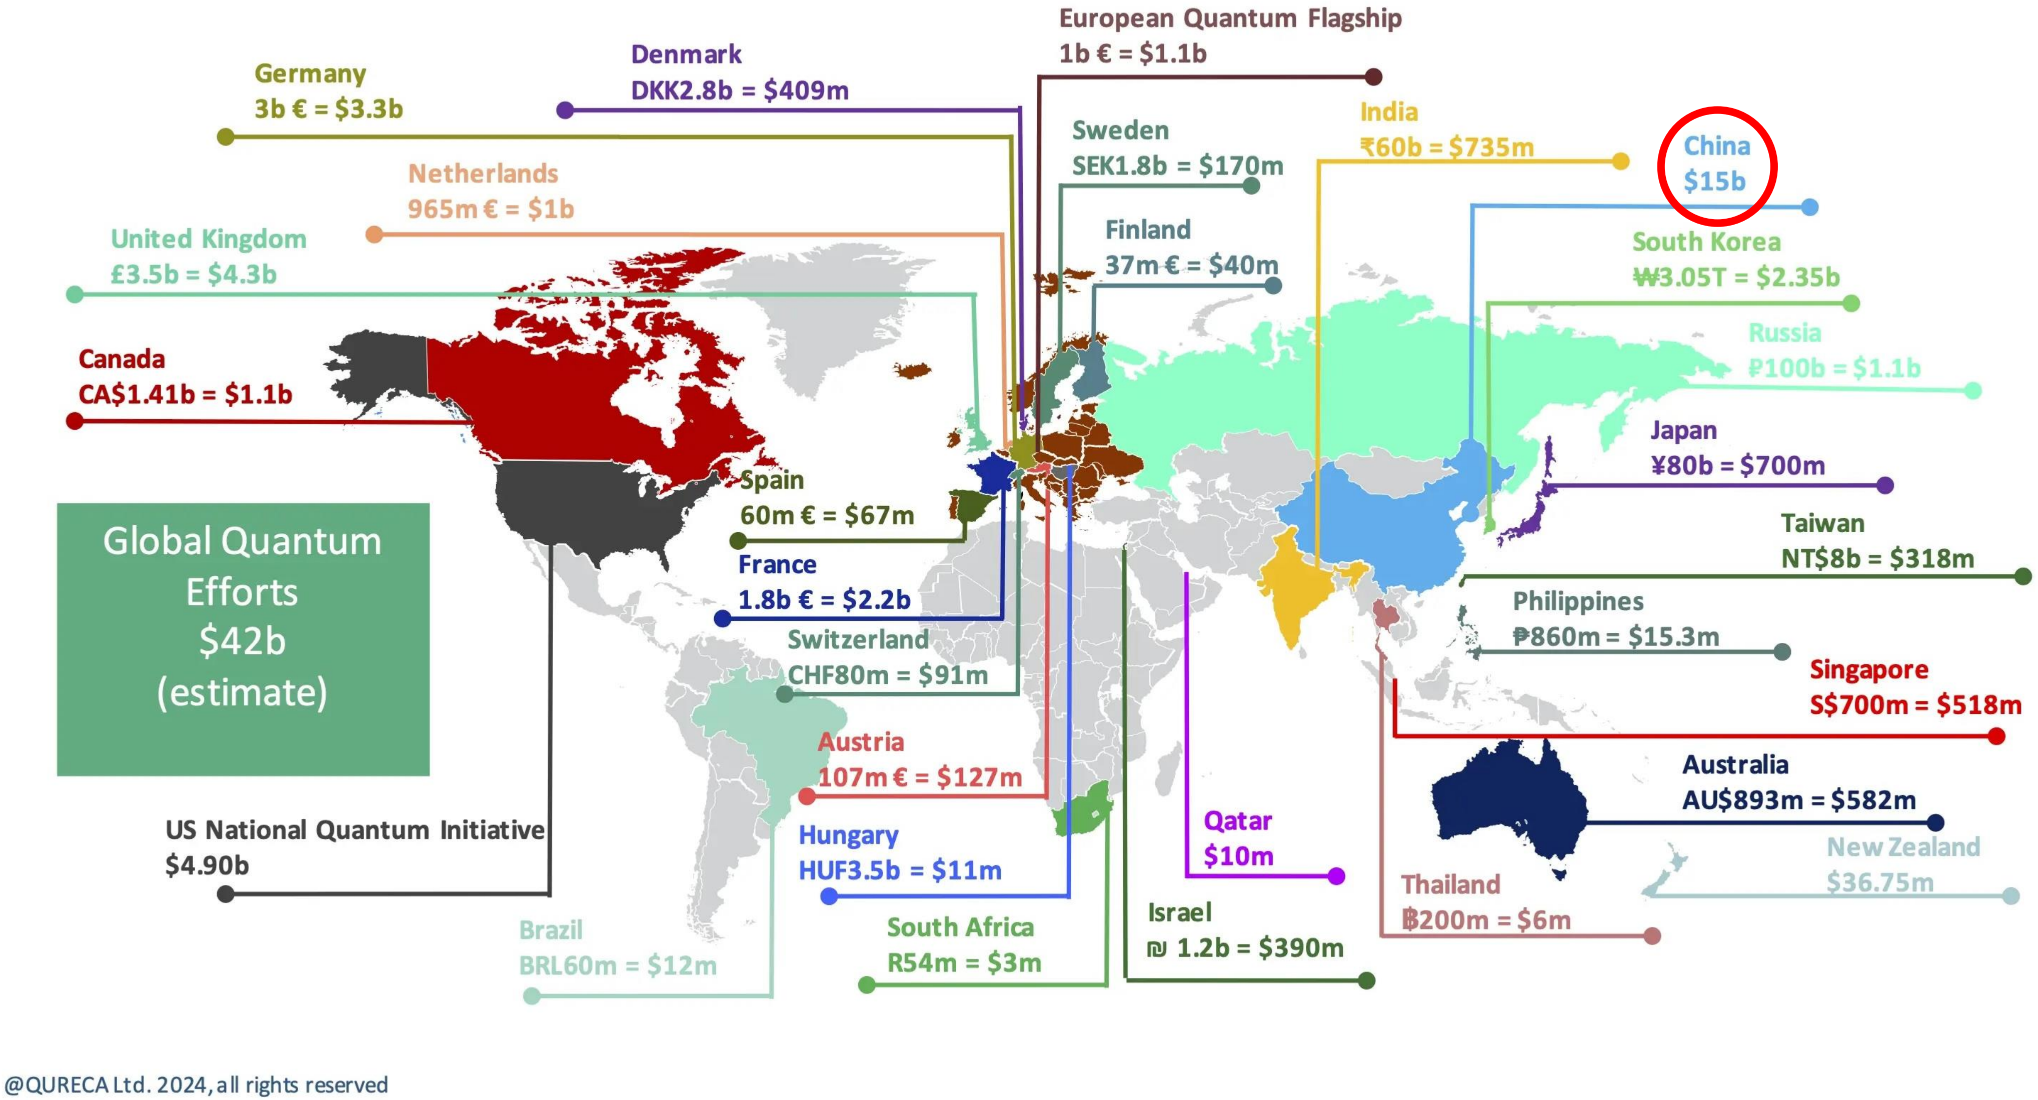

https://www.qureca.com/quantum-initiatives-worldwide/

People are making incremental efforts in developing a
Quantum Computer.

Once they have one which is sufficiently large and reliable, they may use it to
Break Current Encryption!
(public key algorithms)

Quantum Compute Timeline

IBM's Quantum Roadmap is often cited for realistic predictions

Quantum timelines may appear to be sooner than you think

IBM Quantum Roadmap

Quantum Computing’s Impact on Cryptography

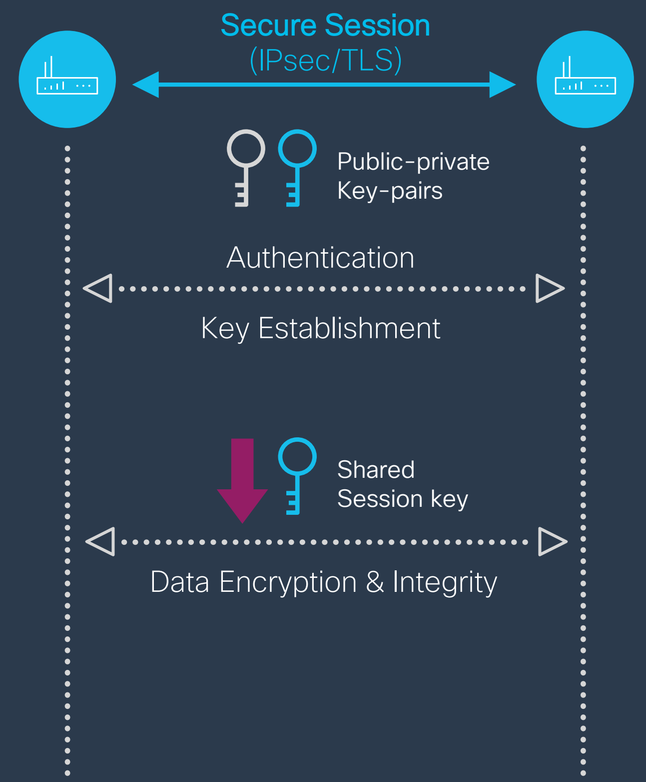

Asymmetric Cryptography

- Based on **mathematically related** public-private key-pairs
- Used for control plane operations
  - Authentication, Key establishment
- Example: RSA, DH, ECC

Symmetric Cryptography

- Based on shared key
- Used for bulk data encryption & integrity
- Protection level based on key strength
  - Key size & entropy
- Example: AES-GCM

Quantum-Resistant?

Large reliable Quantum computers can break RSA, DH, ECC!

Symmetric crypto with large and high-entropy keys is resistant to Quantum computer attacks

Post-Quantum Pre-shared Keys (PPKs)

Quantum-safe encryption keys using RFC-8784

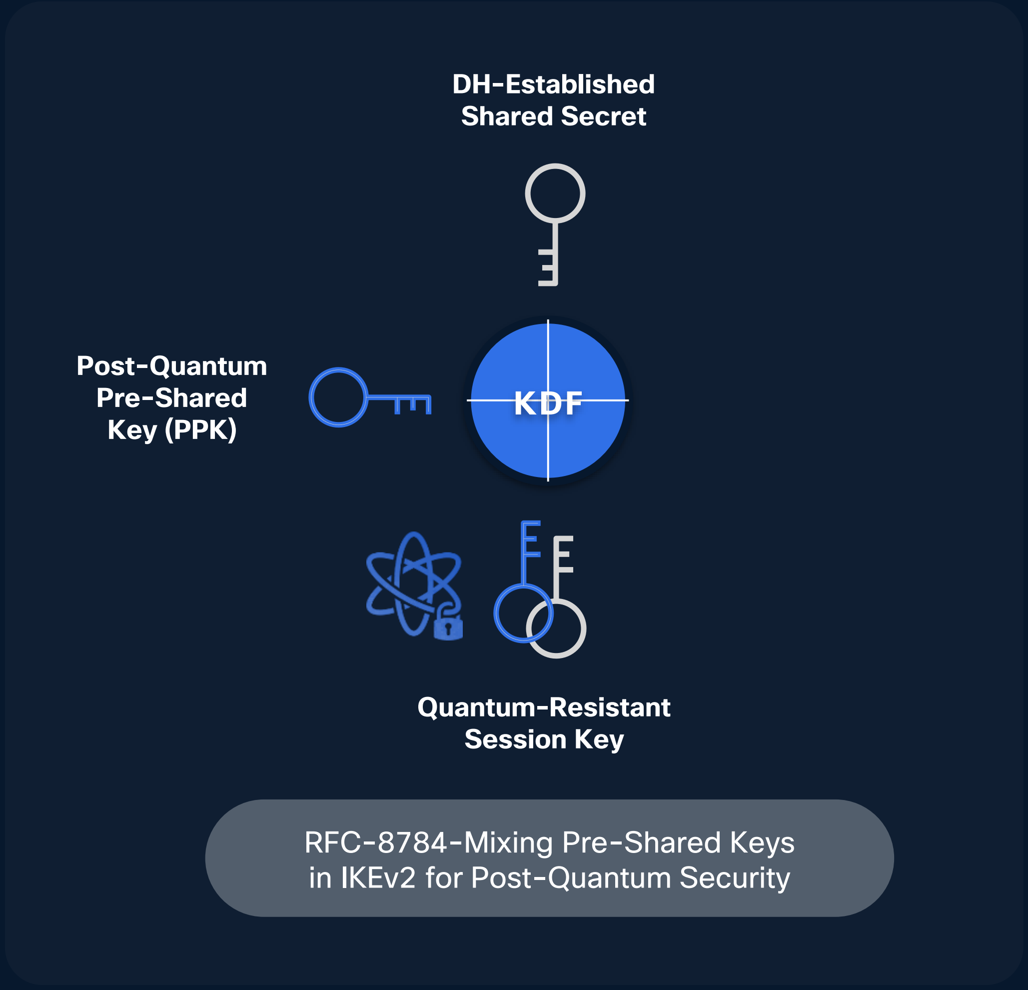

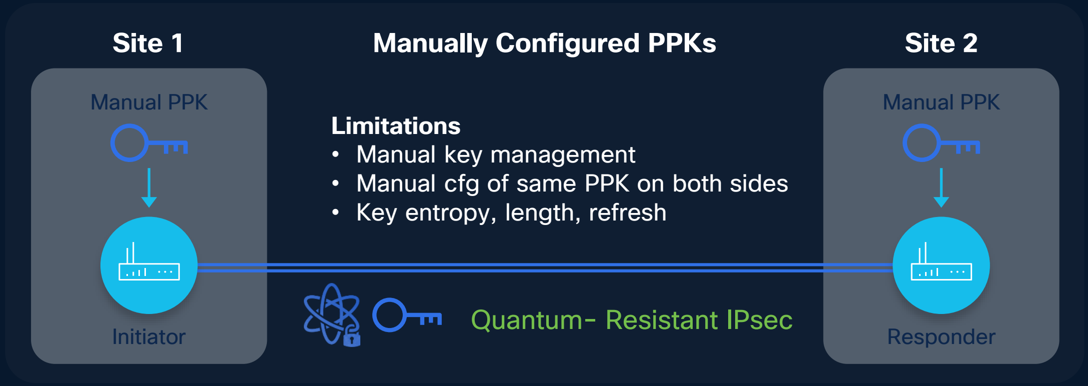

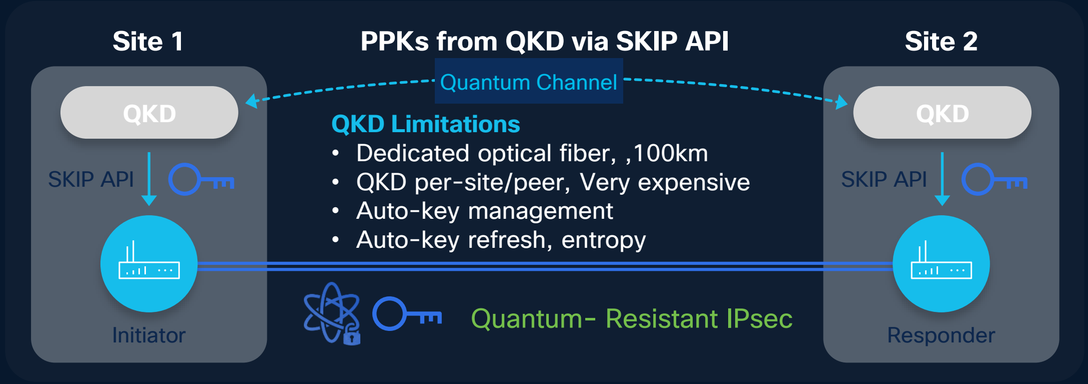

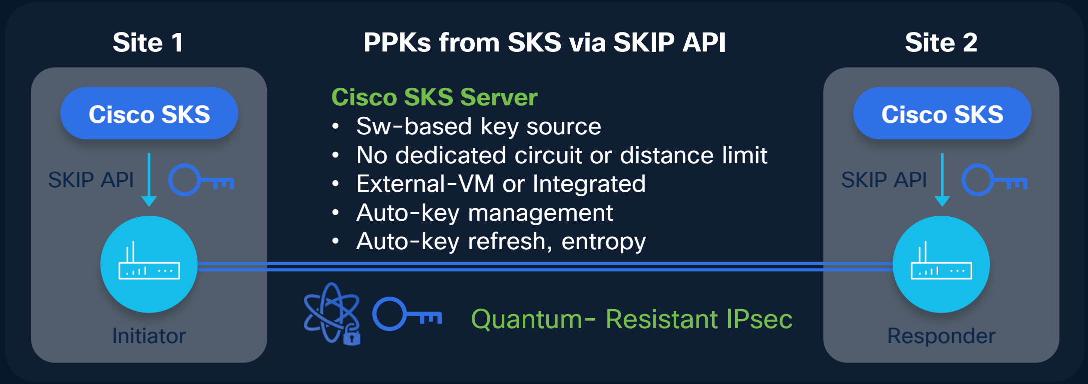

NSA | Commercial National Security Algorithm Suite 2.0

Serves as the cryptographic base to protect US National Security Systems. Required-by date for new acquisitions accelerated to **January 2027.** Only PQC allowed in National Security Systems after **December 2031.**

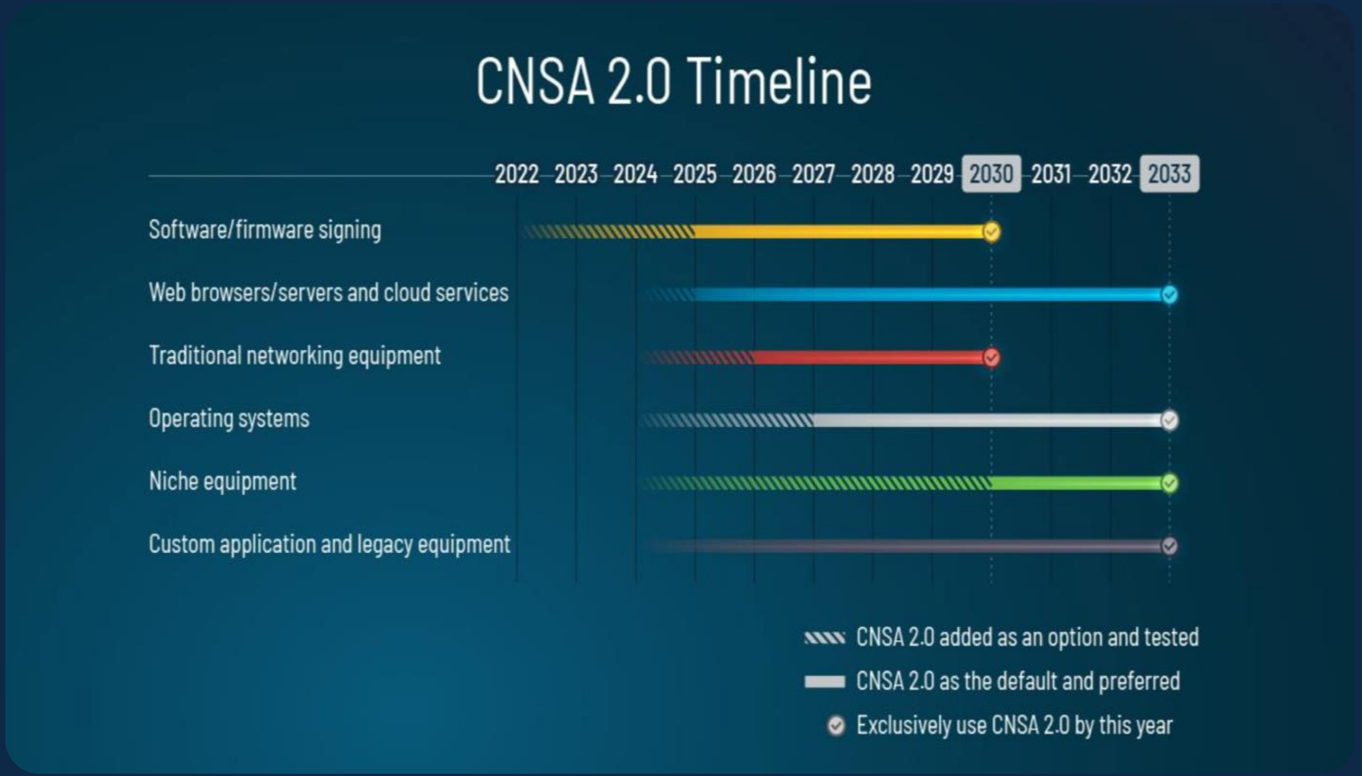

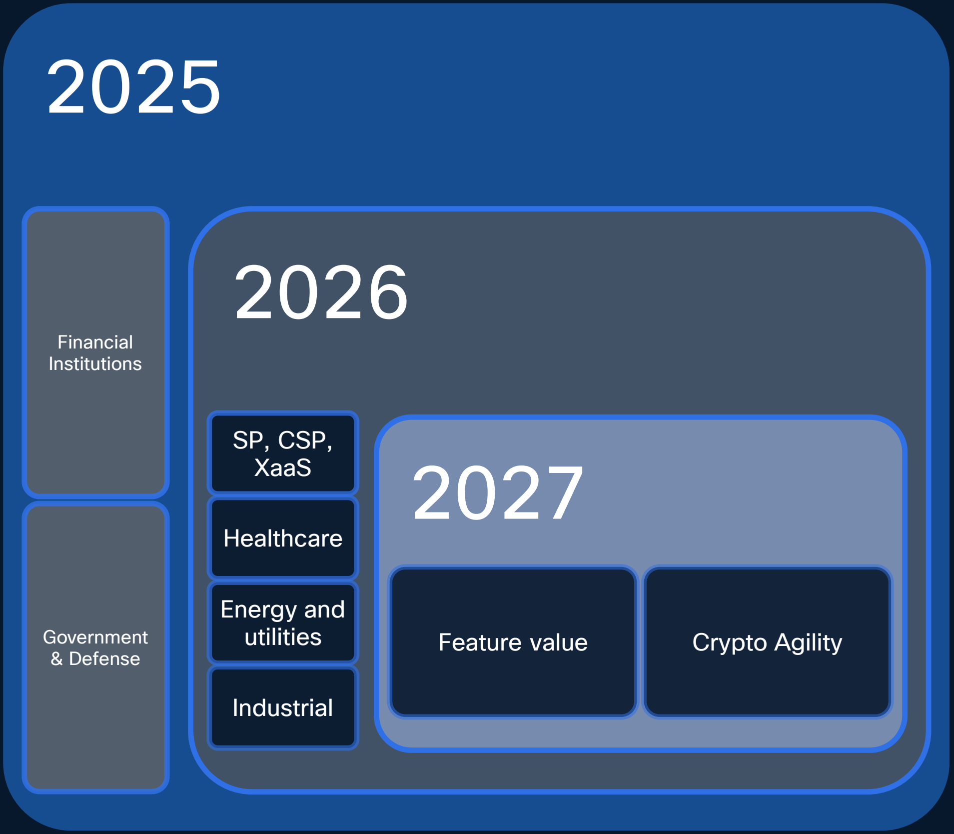

Source: National Security Agency, *Commercial National Security Algorithm Suite 2.0*

Preparing for Quantum Computing Security Threats

Educate

Ensure IT, security, and leadership understand the quantum threat

Assess & Prioritize

- Build a cryptography database.
- Identify encryption used for all systems, applications, and data flows.
- Special note for asymmetric encryption usages
- Focus on high-value data with long-term sensitivity—think trade secrets, financial records, or healthcare data that needs to remain secure for decades

Research Options

- Start conversations with software and hardware suppliers about their quantum-readiness roadmaps.
- PQC-usages, Hybrid solutions, and crypto agility.

Develop a Strategy

- Create a roadmap with immediate, near-term, and long-term actions.
- Plan for crypto agility, enabling seamless switching between cryptographic algorithms as threats evolve.

Execute the Strategy

- Adopt NIST Standards
- Implement Hybrid schemes that combine current algorithms with PQC algorithms.
- Update Protocols: TLS, VPNs, and other secure communication protocols to support PQC.

Monitor Progress

- Track transition progress
- Follow advancements in quantum hardware (qubit count, error correction)
- Stay aligned with standards bodies.

Securing AI using Zero Trust Principles

Cisco AI Best Practices

Preserving your trust with AI Governance

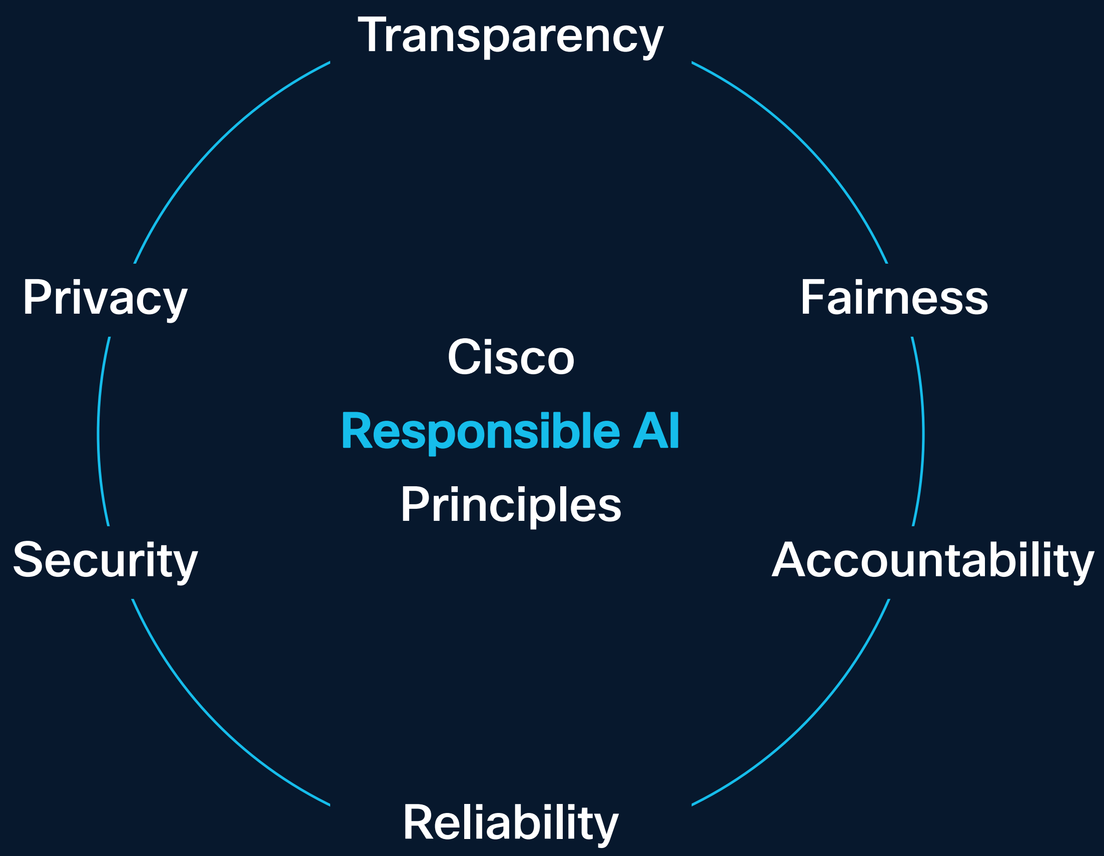

We're navigating the intersection of AI security, regulatory compliance, and Zero Trust–grounded in NIST 800-207, the NIST AI RMF, the EU AI Act, and informed by OWASP AI best practices, MITRE ATLAS threat models, and the Cloud Security Alliance (CSA) AI Controls Matrix.

AI Security Regulations & Standards

EU Artificial Intelligence Act
EU AI Act

CSA AI Risk Mgt FW

NIST AI RMF

OWASP

- LLM01 Prompt Injection
- LLM02 Sensitive Information Disclosure
- LLM03 Supply Chain
- LLM04 Model Denial of Service
- LLM05 Improper Output Handling
- LLM06 Excessive Agency
- LLM07 System Prompt Leakage
- LLM08 Vector and Embedding Weaknesses
- LLM09 Misinformation
- LLM10 Unbounded Consumption

MITRE ATLAS™

- Reconnaissance
- Resource Development
- Initial Access
- ML Model Access
- Execution
- Persistence
- Privilege Escalation
- Impact
- Exfiltration
- ML Attack Staging
- Collection
- Discovery
- Credential Access
- Defense Evasion

Safe and Trustworthy Use of GenAI: Key Recommendations

Securing AI using Zero Trust Principals

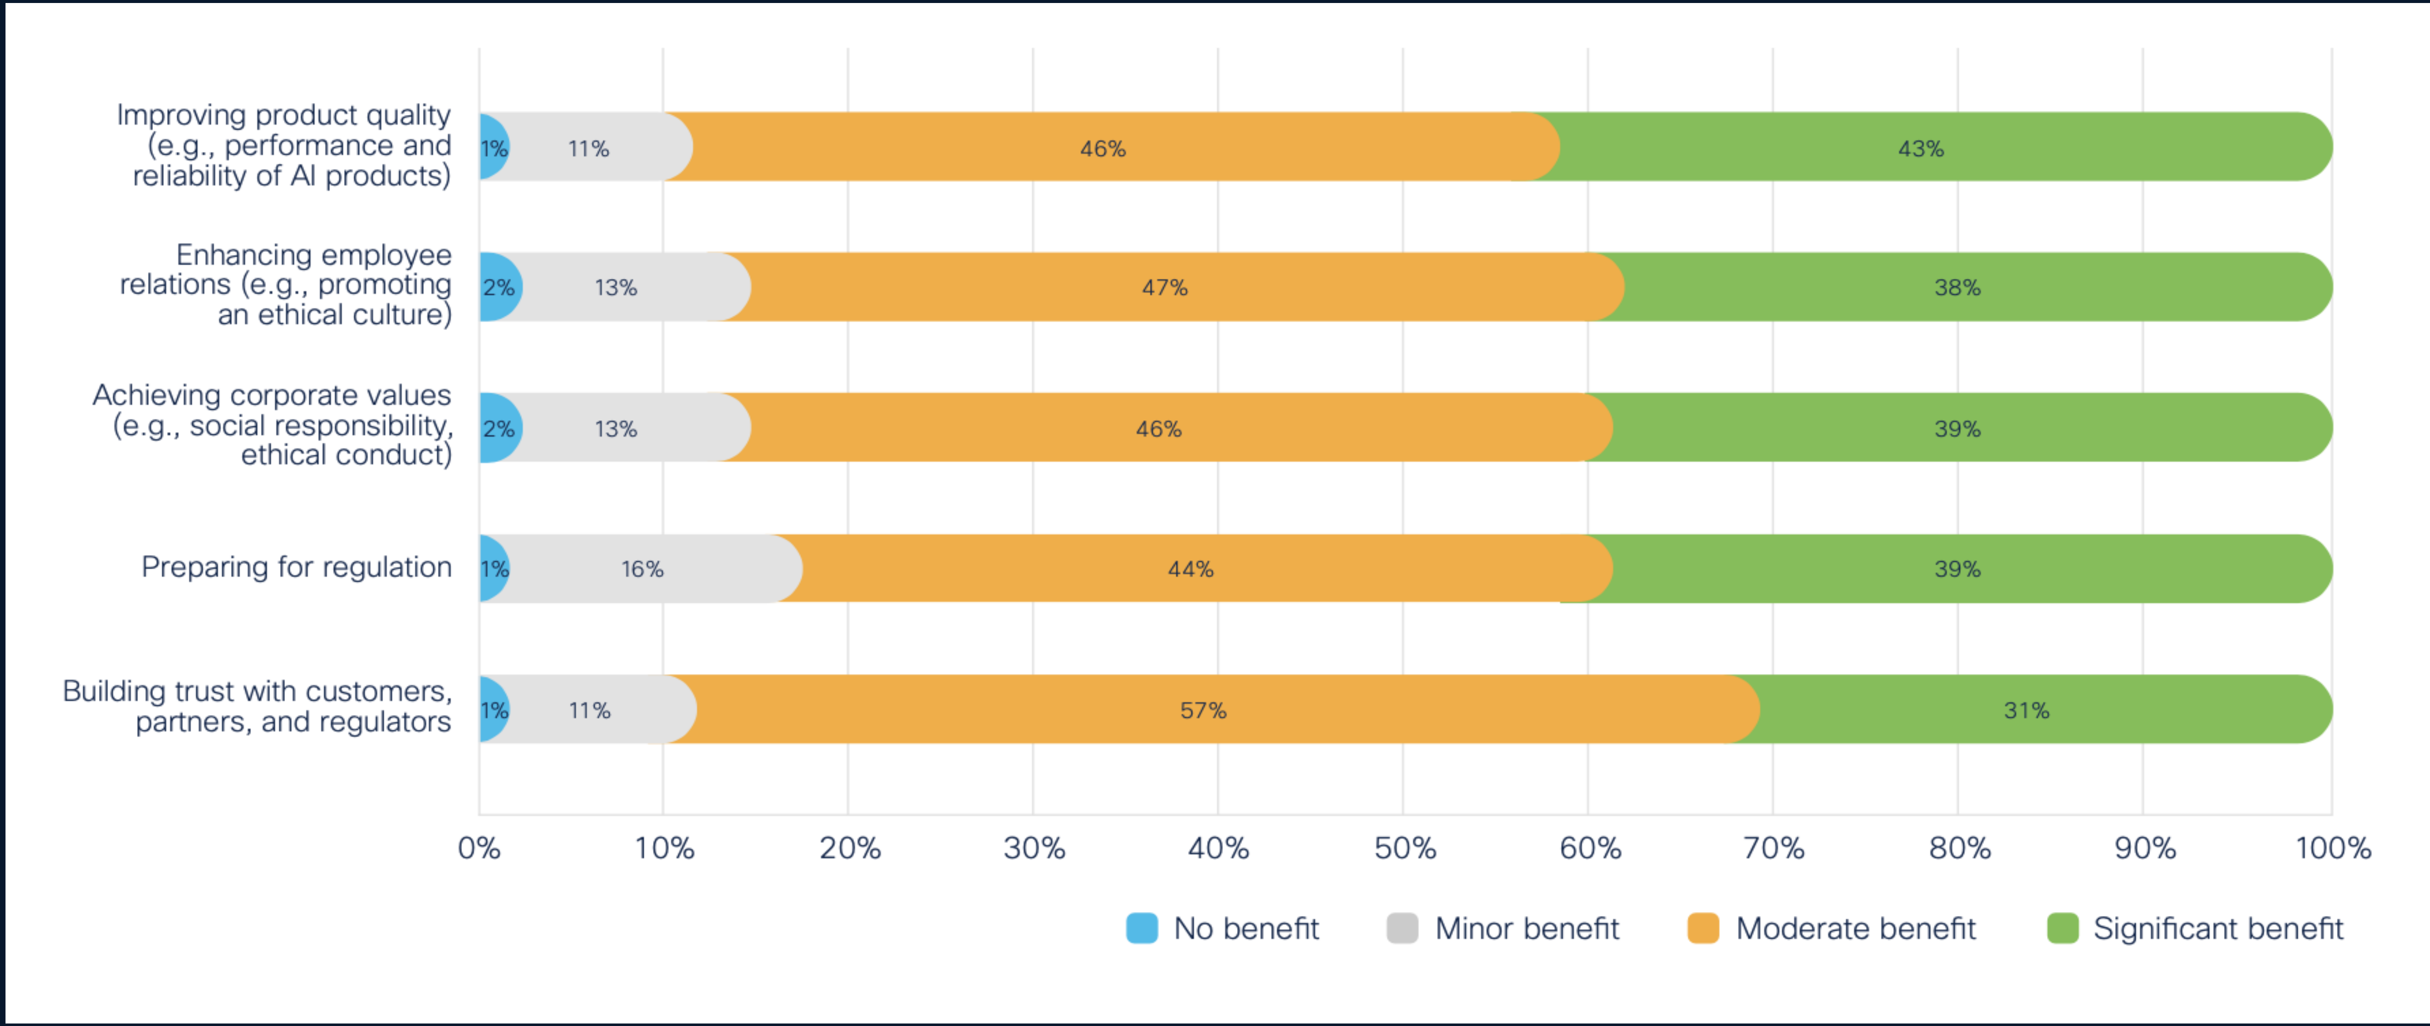

Source: April 22, 2025 Cisco's 2025 Data Privacy Benchmark Study sheds light on the importance of keeping sensitive information safe as AI transforms industries.

AI Zero Trust: What needs protecting & how?

Enterprise-Controlled AI (First-Party & Open-Source AI)	Externally Managed AI (Third-Party & Cloud AI)	Distributed AI (Federated & Embedded/Edge AI)	Autonomous Agent AI

• **Definition:** AI systems developed internally or using open-source frameworks, fully governed by the enterprise.
• **Examples:** In-house fraud models, AI built on Hugging Face, OpenLLaMA.
• **Zero Trust Notes:**
  • Full control over model lifecycle and data handling.
  • Requires secure SDLC, supply chain risk management (SBOMs), and internal audit trails.
  • Best aligned with **NIST AI RMF**, **Zero Trust Architecture (NIST 800-207)**, and **OWASP AI Guidelines**.
	• **Definition:** AI services or solutions managed by external vendors or public cloud providers.
	• **Examples:** SaaS AI tools, Azure OpenAI, Google Vertex AI.
	• **Zero Trust Notes:**
	  • Shared responsibility for security and compliance.
	  • Needs vendor risk management, SLA enforcement, and runtime monitoring.
	  • Ensure compliance with **EU AI Act**, **NIS2**, and data residency laws.
		• **Definition:** AI that operates across distributed or resource-constrained environments like IoT and edge.
		• **Examples:** AI in healthcare federated learning, smart city edge devices.
		• **Zero Trust Notes:**
		  • Emphasizes decentralized trust boundaries, secure data aggregation, and environmental integrity.
		  • Must apply differential privacy, OTA security, and federated cryptographic trust anchors.
		  • Critical for **HIPAA**, **GDPR**, and **industry-specific standards** (e.g., IEC 62443).
			• **Definition:** AI agents capable of initiating actions based on goals, context, and evolving logic.
			• **Examples:** AutoGPT, self-operating bots, AI-driven SOC runbooks.
			• **Zero Trust Notes:**
			  • Highest need for granular guardrails: role-based actions, behavioral monitoring, and real-time intent validation.
			  • Applies both **NIST AI RMF Manage** and **EU AI Act Article 14–15** requirements.
			  • Must account for **MITRE ATLAS** adversarial tactics and OWASP action confinement controls.

AI Agents: What's Next?

AI Agents are a Digital Employee.
They need a Job Description that must be adhered to to secure the organization

AI agents automate sensitive operations, securing them isn't optional – it's existential.

Let's talk about five guardrails to provide a structured, standards-aligned approach to deploying AI responsibly and defensively.

Zero Trust Guardrails for AI Agents

Principle 1 - "Allow known good. Block everything else."

In a Zero Trust architecture, outlined in **NIST 800-207**, access is never implicit—it must be continuously validated. This principle is reinforced by **MITRE ATLAS**, which highlights how adversaries exploit ungoverned or over-permissive AI environments to deploy shadow agents or prompt injections.

OWASP's AI Guidelines recommend strict agent control mechanisms and validated allow lists to prevent rogue model execution. Likewise, the **EU AI Act (Articles 9–15)** mandates that high-risk AI systems operate within tightly controlled, pre-approved boundaries.

The **CSA AI Controls Matrix** emphasizes the importance of defining and enforcing strict access controls for AI systems, ensuring only authorized agents operate within predefined parameters.

Allow listing defines the safe zone. Everything else? It's denied by design.

Zero Trust Guardrails for AI Agents

Principle 2 – Make Policies Readable "Security that can’t be understood, can’t be trusted."

Declarative, transparent policies empower human oversight—vital per the **Govern function of the NIST AI RMF** and **Article 13 of the EU AI Act**, which require explainability in AI governance.

OWASP guidance emphasizes the importance of human-in-the-loop designs and interpretable rule enforcement. Security controls should be intelligible by policy owners, auditors, and compliance teams—not just engineers.

The **CSA AI Controls Matrix** aligns with this by advocating for clear documentation and transparency in AI system policies, facilitating easier audits and compliance checks.

This also supports Zero Trust policy centralization as specified in **NIST 800-207**, making enforcement both visible and auditable across domains.

Zero Trust Guardrails for AI Agents

Principle 3 – Log Everything "Every interaction tells a story. Capture it."

Comprehensive telemetry is critical. The **NIST AI RMF (Map & Measure)** calls for full lifecycle visibility, while **MITRE ATLAS** threat use cases demonstrate how visibility gaps enable stealthy AI model manipulation and misuse.

OWASP AI recommendations stress logging inputs, decisions, and outputs—particularly for model inference and external API interactions. The **EU AI Act (Article 12)** reinforces this, requiring audit trails that verify the integrity of decision-making processes.

The **CSA AI Controls Matrix** underscores the necessity of detailed logging and monitoring to detect anomalies and ensure accountability in AI operations.

Zero Trust requires pervasive observability, and that starts with capturing every event—authorized or blocked. If we can't understand the logs then that's a problem.

Zero Trust Guardrails for AI Agents

Principle 4 – Fail Closed, Not Open "No access is better than wrong access."

Failing open creates an adversary playground. **MITRE ATLAS** includes tactics where attackers induce fail-open conditions, such as DoS on policy evaluators or corrupt fallback mechanisms.

The **NIST AI RMF** emphasizes robustness and resilience—AI must respond to uncertainty by reducing exposure, not increasing it. The **EU AI Act (Article 14)** mandates fallback safeguards to ensure safety when anomalies are detected.

OWASP guidance advises strict policy enforcement with well-tested error handling to avoid insecure defaults.

The **CSA AI Controls Matrix** advocates for default-deny strategies and robust error handling to prevent unauthorized access during system failures.

In Zero Trust, indecision means denial—anything less is a liability.

Zero Trust Guardrails for AI Agents

Principle 5 – Use Multiple Layers "Defense in depth applies to AI too."

Layered controls are foundational in **NIST 800-207's ZTA model** and reinforced by the **Manage function of NIST AI RMF**. From input sanitization to model governance to contextual access checks—defense in depth is critical.

OWASP urges use of runtime monitors, rate limiters, and anomaly detectors. **MITRE ATLAS** documents chained attacks that bypass a single weak control—multi-layer defenses mitigate this.

The **CSA AI Controls Matrix** supports implementing multiple layers of security controls, including data encryption, access management, and continuous monitoring, to protect AI systems comprehensively.

Regulators agree: The **EU AI Act** calls for multi-dimensional risk controls across the AI lifecycle. Stack identity, environment, behavior, and purpose validation—because no single control is infallible.

cisco

Sample Job Description Internal AI Agent: SOC Analyst AI Agent

SOC Analyst AI Agent

Maintaining Vigilance Across Security Logs with Autonomous AI

Role Overview: SOC Analyst AI Agent
Focus: Real-Time Threat Monitoring & Triage
Function: Cybersecurity Operations
Reports To: SOC Manager / CISO

Mission: Analyze security logs across the corporate envirorment to identify and prioritize threats—resolving potential incidents or escalating to human analysts.

Key Responsibilities

- Continuously analyze diverse security logs for threat patterns & anomalies
- Prioritize and triage security alerts based on severity and context
- Investigate & correlate events to accurately detect active threats
- Automate containment of validated low-risk incidents
- Handoff complex, high-risk cases to SOC analysts for further investigation

Required Expertise

- Extensive training on SIEM plarforms such as Splunk or Microsoft Sentinel
- Proficiency in attack techniques (e.g. phishing, malware) and their indicators
- Knowledge of incident response processes and escalation procedures

Strategic Impact

- Reduce mean-time-to-detect for malicious activity
- Maintain consistent and compliant threat responsee workflows
- Enhance threat visibility across the enterprise landscape

Controls to Protect the SOC Analyst AI Agent

Ensuring Safe and Effective Autonomous Operation

Access Policies
Apply strict role-based access policies to limit systems the agent can monitor to its own operational scope

Log Verification
Ensure that logs displayed to othe agent are intact, authentic, and free from tampering

Event Oversight
Establish human oversight of significant events, alerts, and incidents generated by the agent

Transparency
Maintain logs of the agent's activity for auditing, to ensure traceability of its decision-making

Model Security
Implement protections to ensure the security of the agent's underlying machine learning models

Risk Evaluation
Conduct regular evaluations of operational risk and risk to the AI's alignment with security goals

AI Agents: Guardrails Enable Innovation

Secure-by-design doesn't stifle innovation – it scales it safely.

By aligning to NIST's Zero Trust principles, embedding controls from the NIST AI RMF, complying with the EU AI Act, and applying adversary-informed defenses from MITRE ATLAS, OWASP, and the CSA AI Controls Matrix, we create an ecosystem where AI can be trusted, governed, and resilient.

Guardrails are not restrictions–they are the structure that lets innovation run free without running wild.

The AI Defense Solution

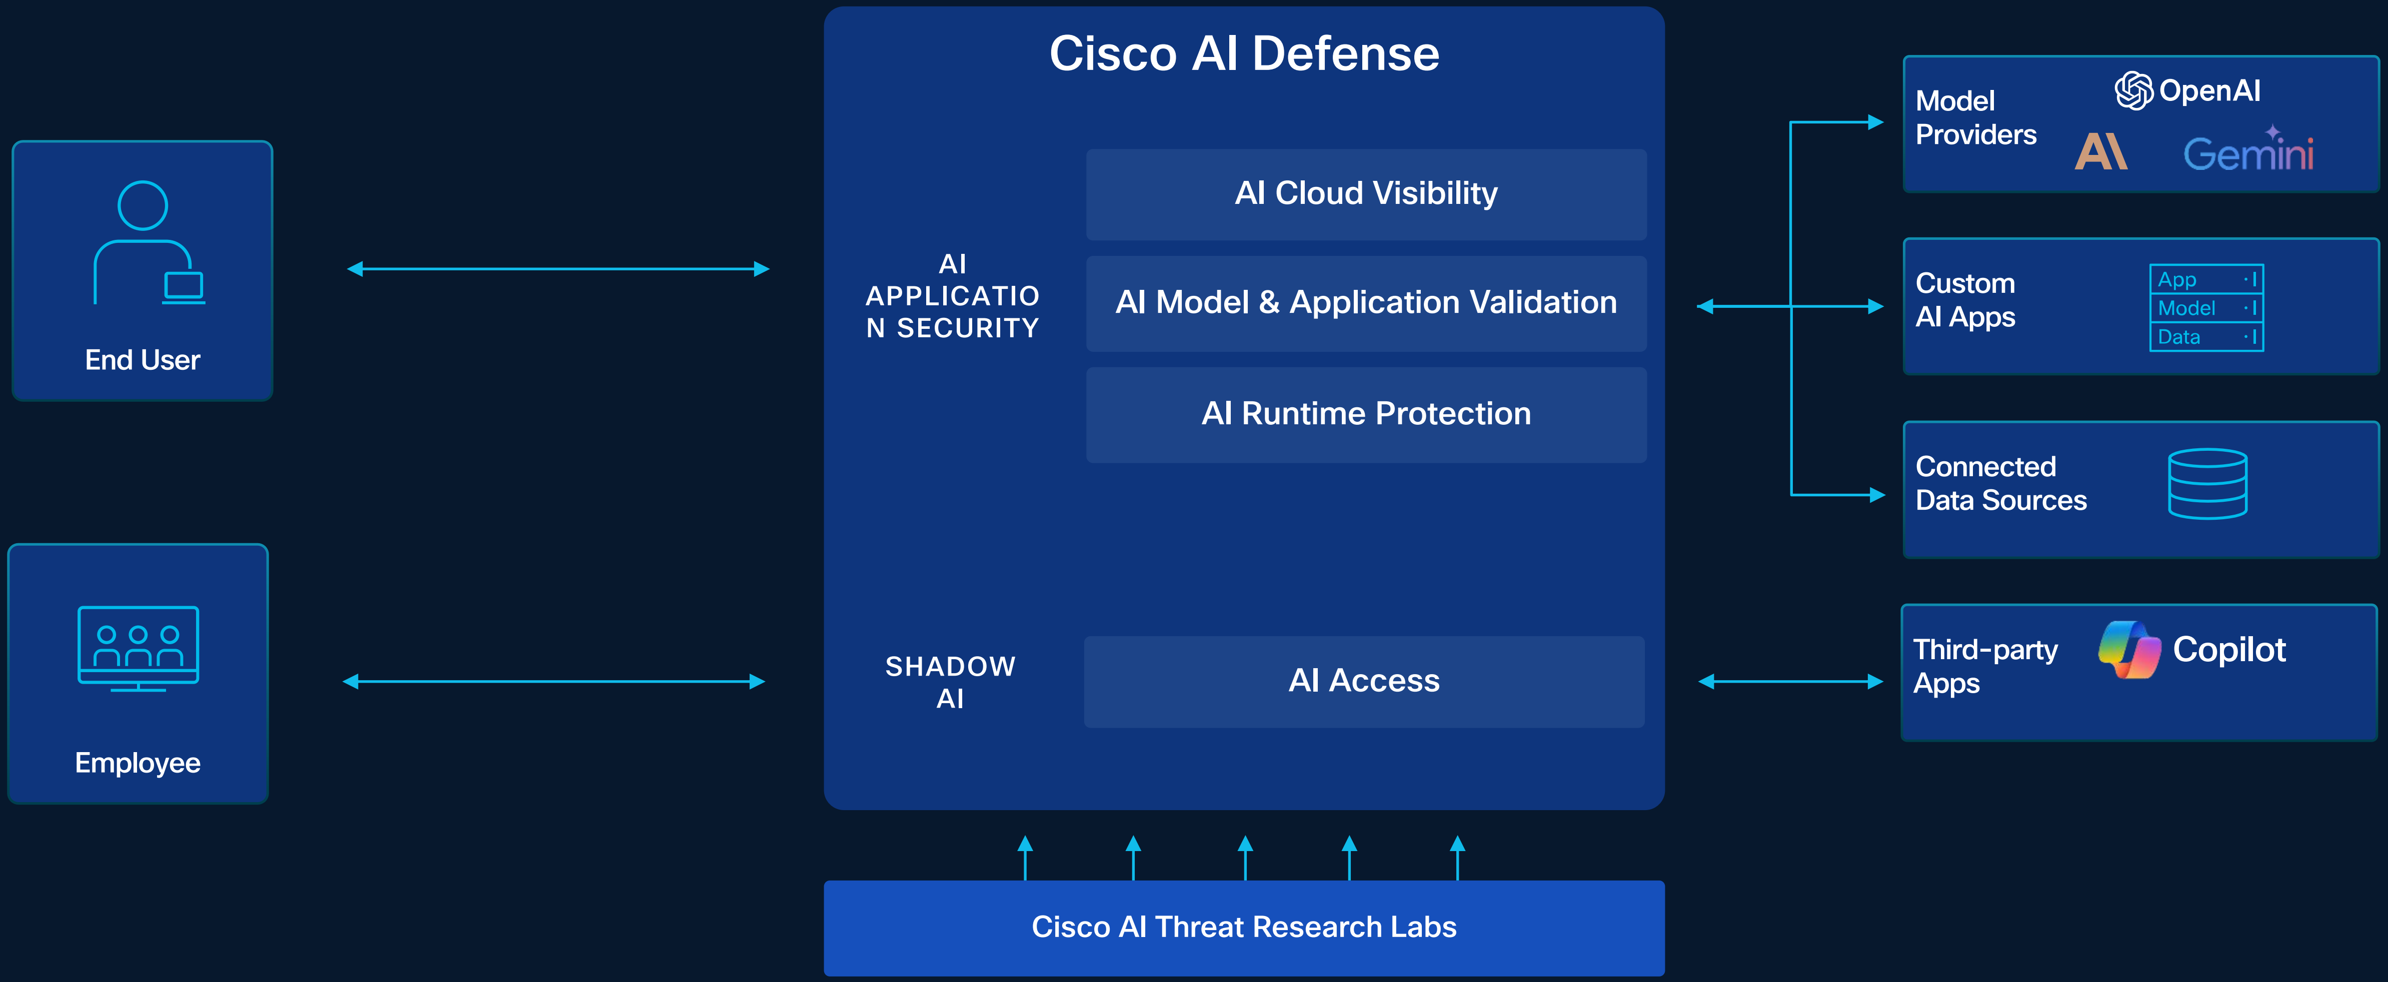

Zero Trust Journey: What’s Next

There are a lot of boats in the water!

Zero Trust Foundational Approach

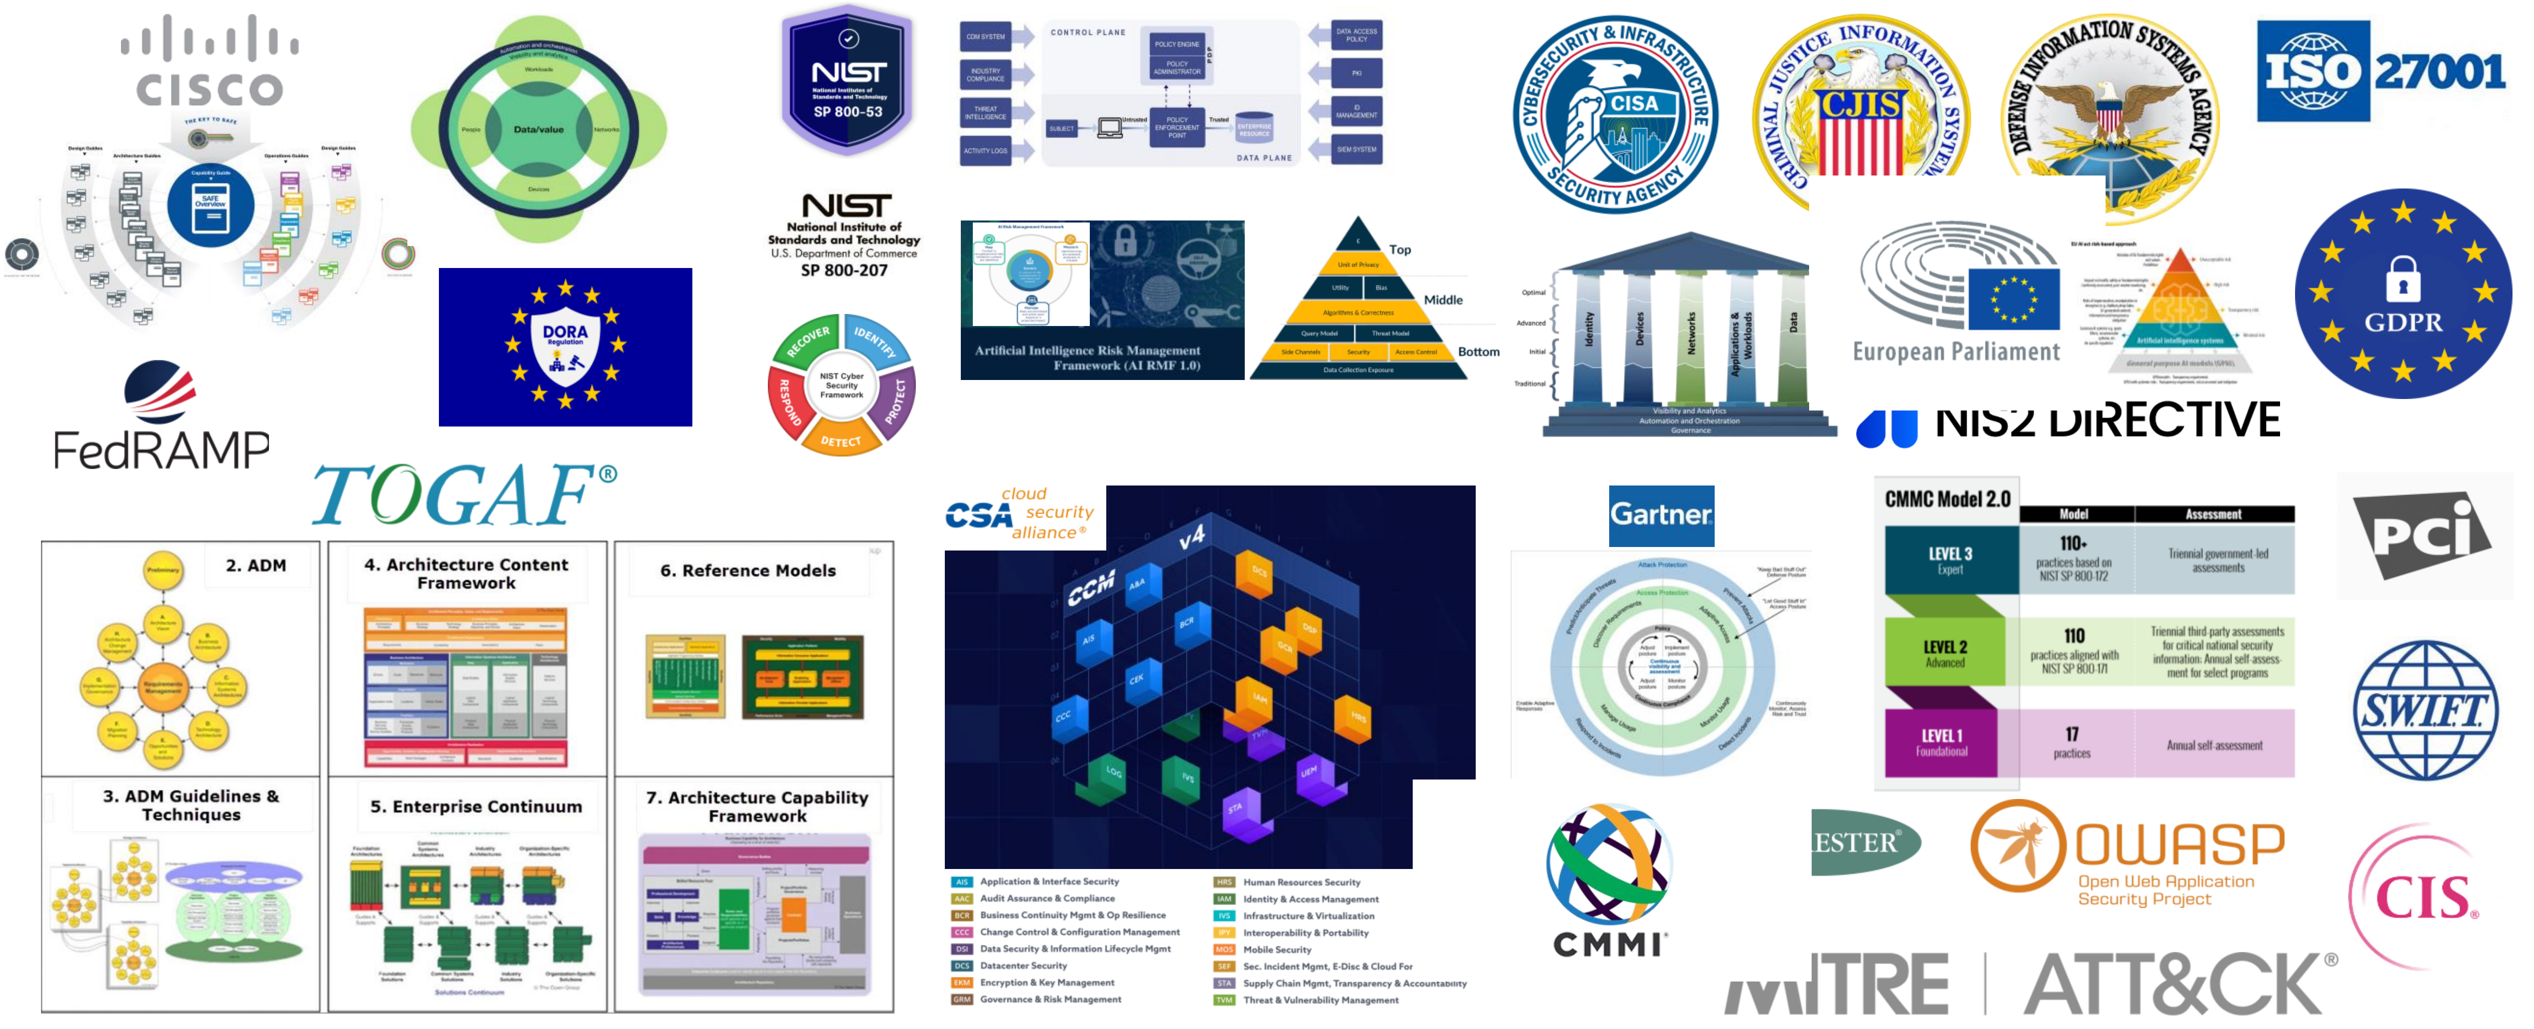

Zero Trust Adoption Drivers

Limit Compliance Scope& Attack Surface

Enable Adaptability & Growth

Improve Network Stability and Resiliency

Long Term Cost Reductions

Protect Brand & Intellectual Property

What is the Foundation of Zero Trust?

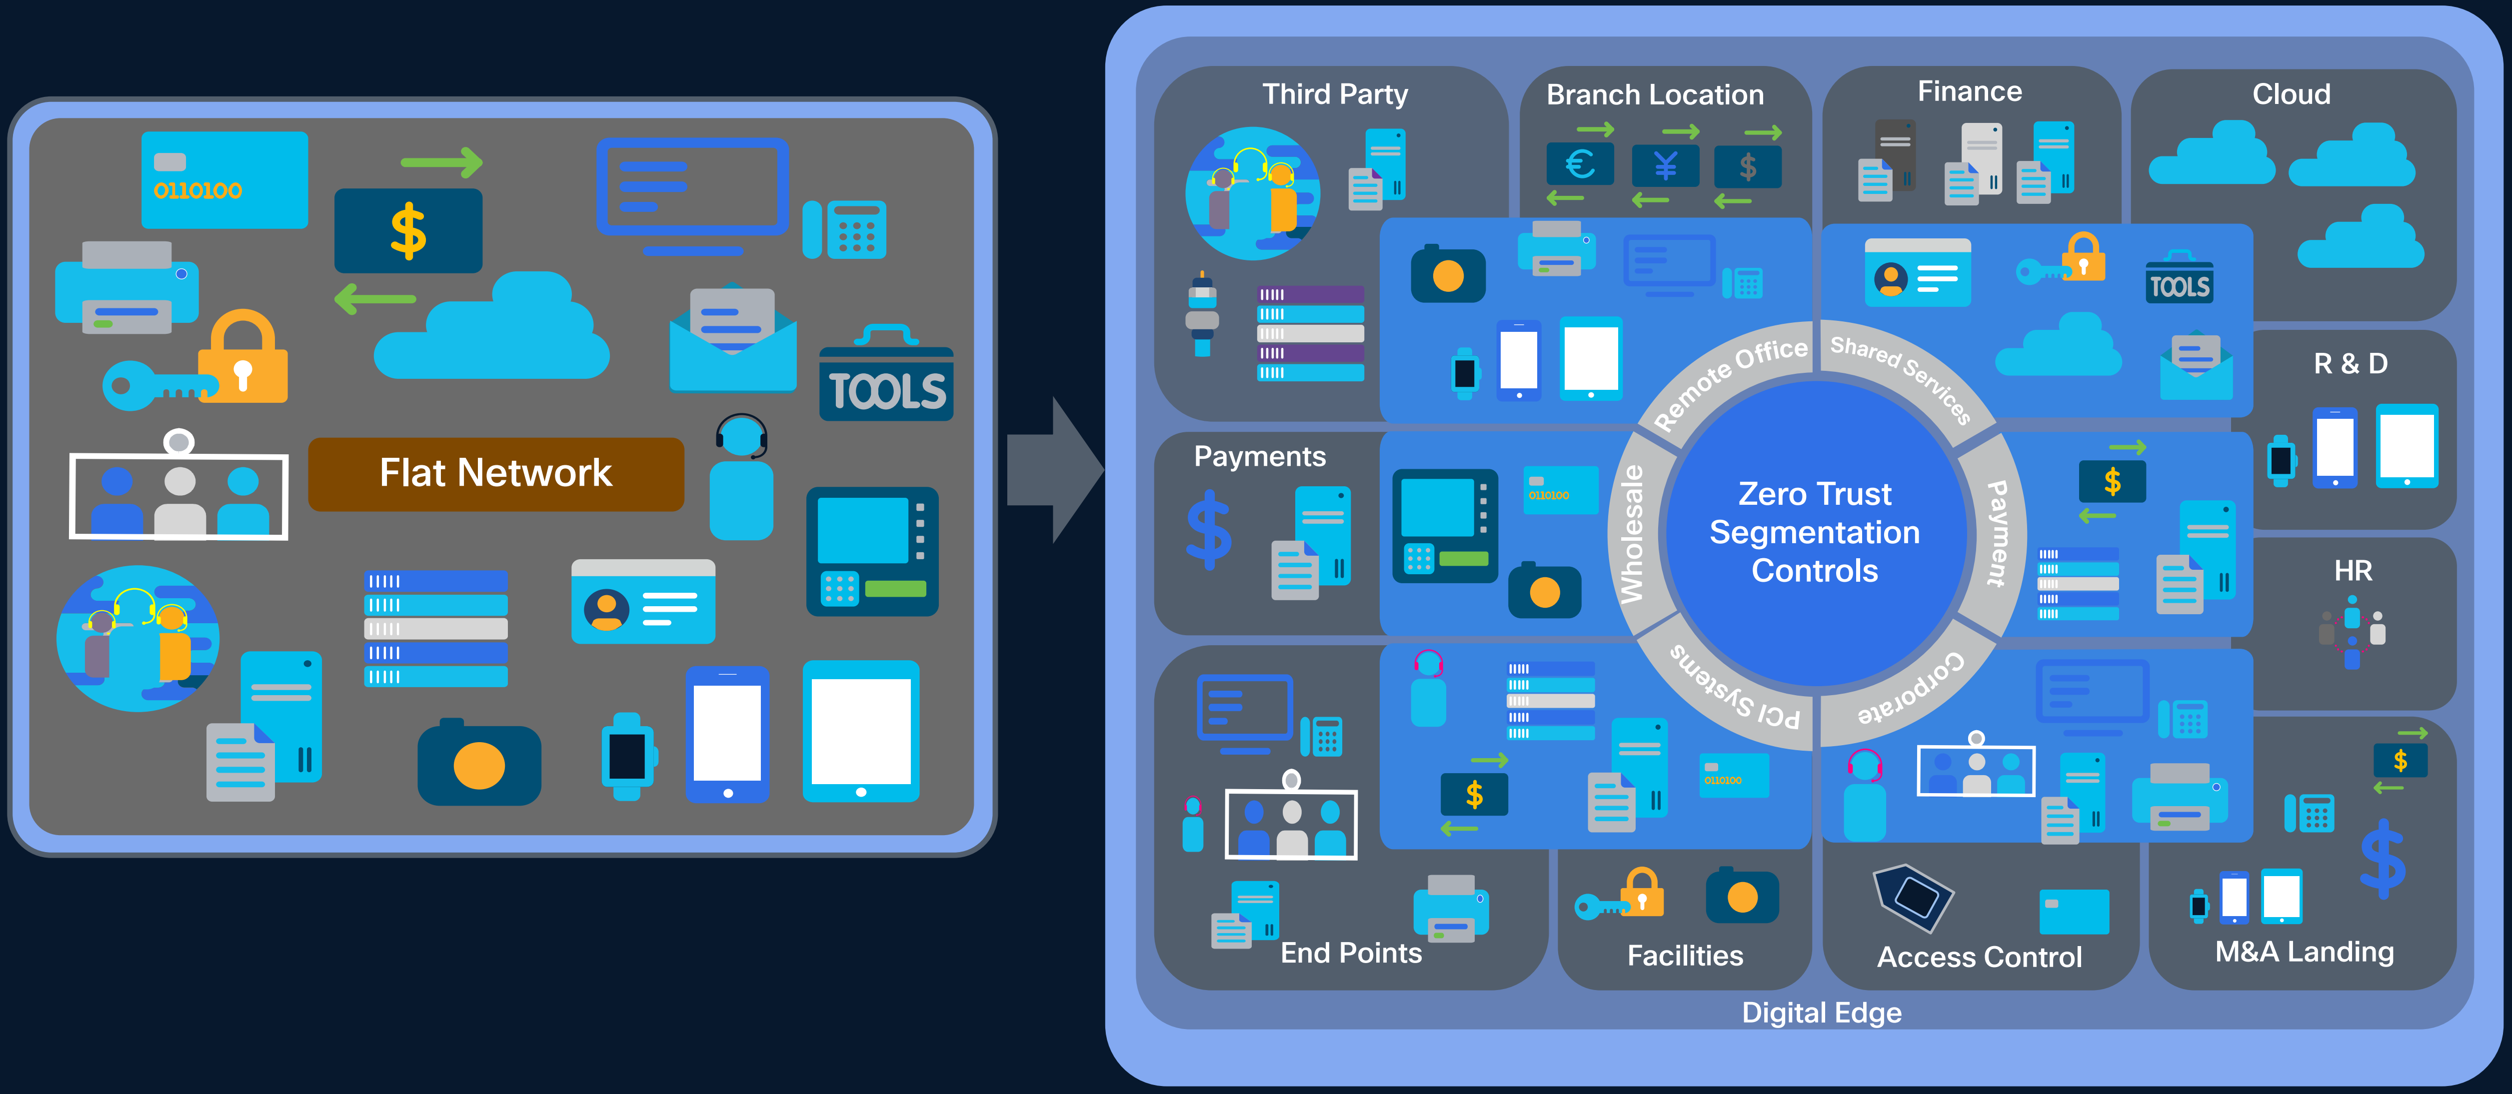

Typical Challenges - Focus Areas

Foundational Requirement: Leadership with Vision and Oversight

Typical Challenges - Focus Areas

Risk Mitigation

Applications | Infrastructure | Operations | Network | Security

Foundational Requirement: Zero Trust Team & Business Alignment

Typical Challenges - Focus Areas

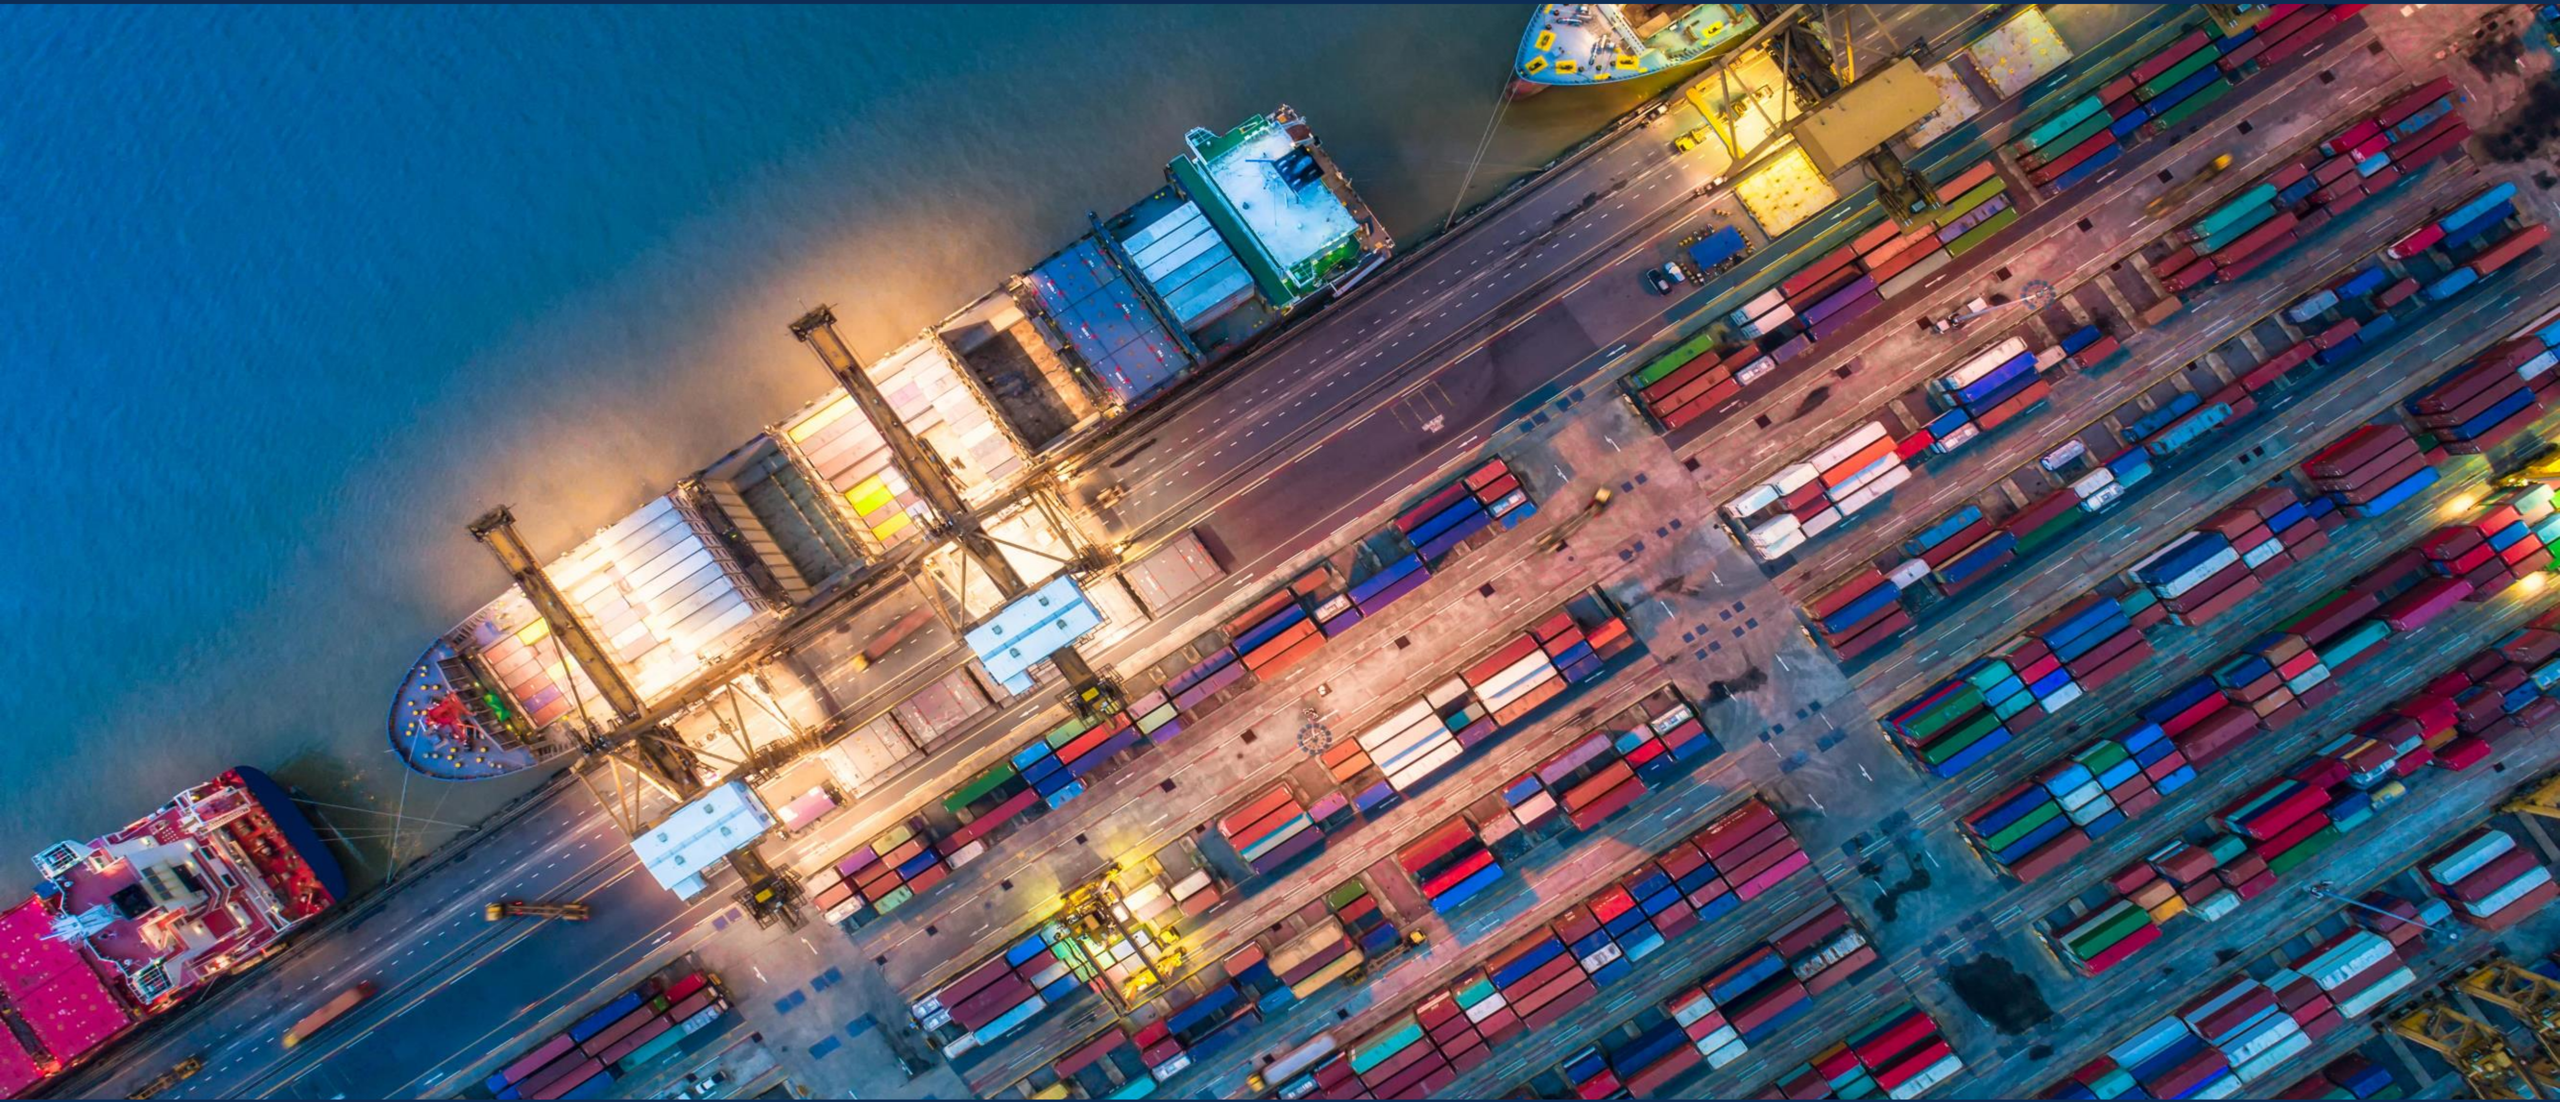

Foundational Requirement: Configuration Mgt - Know What You Have!

Typical Challenges - Focus Areas

Foundational Requirement: Do not start on the Biggest Boat in the Fleet

Let’s Simplify the Process

Zero Trust Capabilities

Policy & Governance

- Change Control
- Data Governance Policy + Encryption
- Data Retention Policy
- QoS
- Redundancy / Replication
- Business Continuity
- Disaster Recovery
- Risk Classification Policy
- Segmentation

Identity

- AAA
- Certificate Authority
- NAC
- Provisioning
- Privileged Access
- MFA
- Asset Identity
- Configuration (CMDB)
- IP Schemas

Vulnerability Management

- Endpoint Protection
- Malware Prevention and Inspection
- Vulnerability Management
- Authenticated Vulnerability Scanning
- Database Change

Enforcement

- CASB
- DDoS
- DLP
- DNS Security
- Email Security
- Firewall
- IPS
- Proxy
- VPN / RA
- SOAR
- File Integrity Monitor
- Segmentation

Analytics

- App. Performance Monitoring
- Audit, Logging, and Monitoring
- Change Detection
- Network Threat Behavior Analytics
- SIEM
- Threat Intelligence
- Traffic Visibility
- Asset Monitoring & Discovery

As Published by Cisco Press Book: "Zero Trust Architecture"

cisco

Zero Trust Enforcement AI-Ready Engine

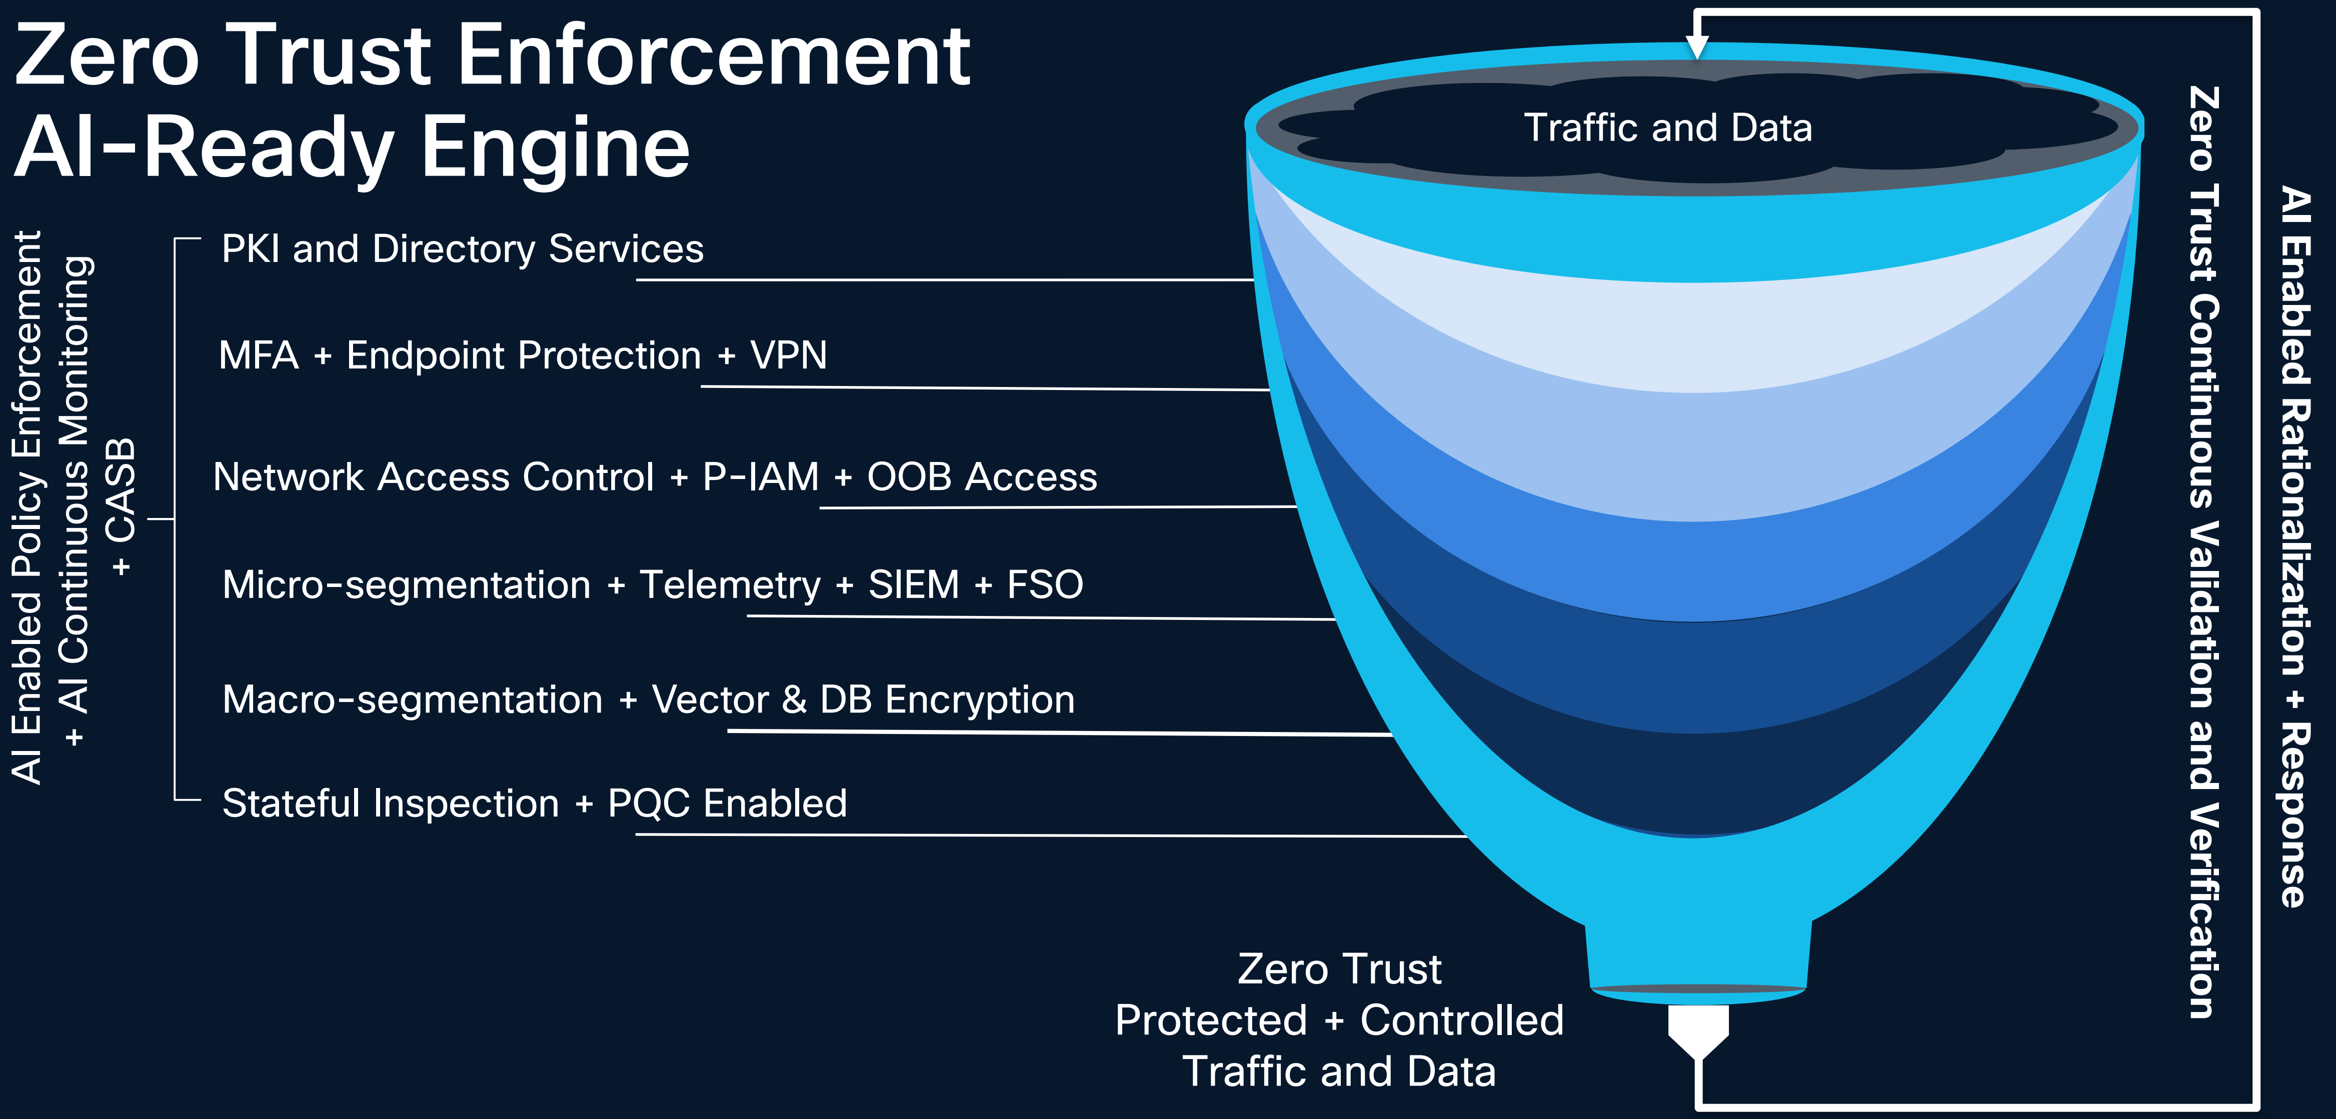

"AI-Ready" + "PQC-Ready" Infrastructure

Zero Trust + Responsible AI Policy + Strategy as a Foundation

Cisco Version – ZT Enforcement AI-Ready Engine

Splunk + Multi-Cloud Defense
AI Defense AI Enabled Policy
Enforcement

- PKI and Directory Services
- DUO + Secure Endpoint + Secure Access
- ISE NAC + IPAM + ISE TACACS+
- Secure Workload + SNA + Splunk + TE
- ACI + Vector & DB Encryption
- Hypershield + Hybrid Mesh Firewall + SKIP

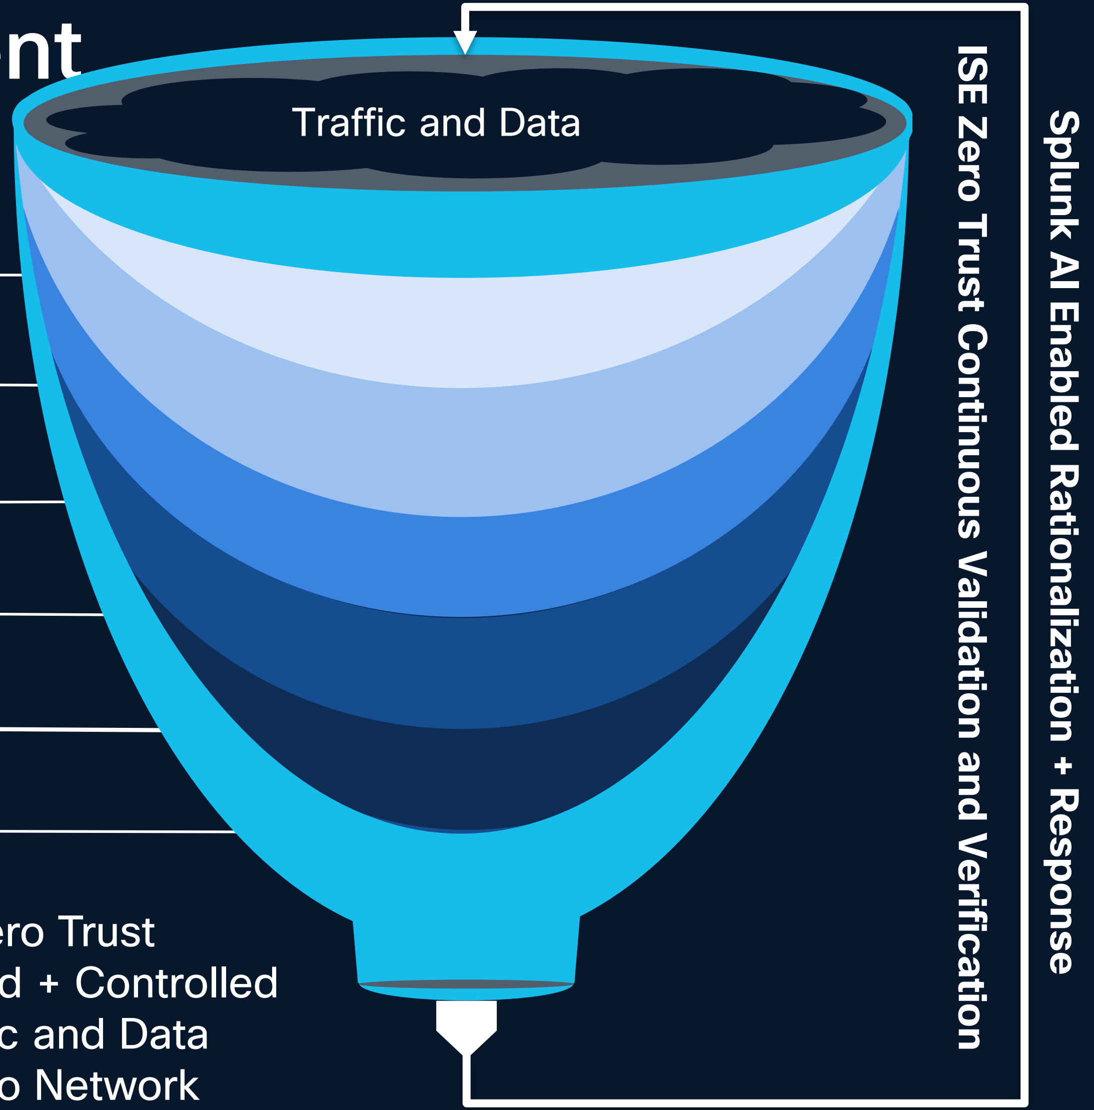

Zero Trust
Protected + Controlled
Traffic and Data
Cisco Network

Nvidia + AI PODs + AI Defense + PQC Enabled Network Devices

Zero Trust + Responsible AI Policy + Strategy as a Foundation

Let’s Build a Model

Organizational Mapping - Sample

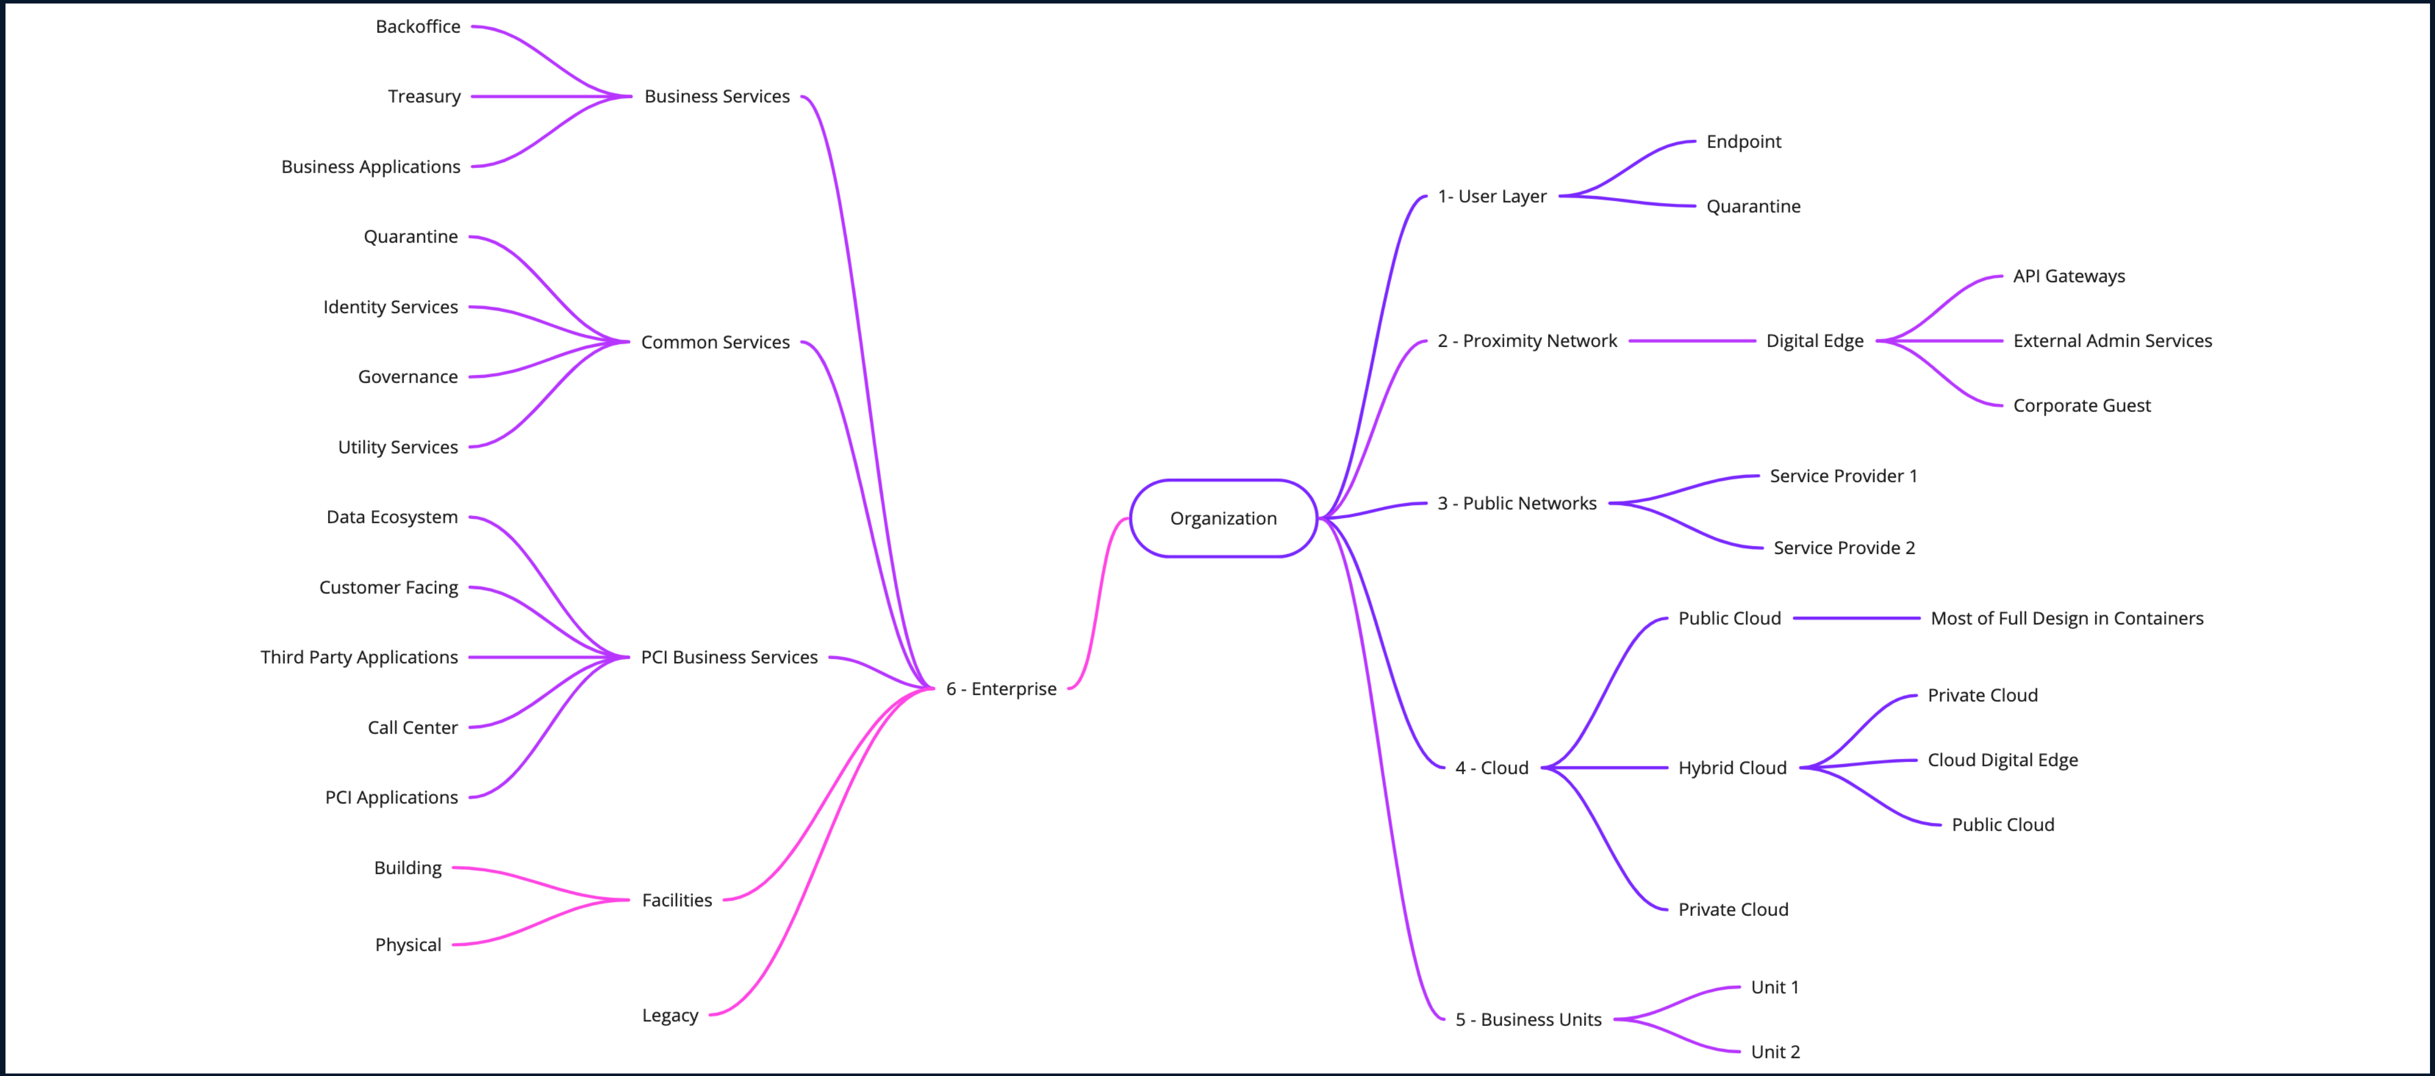

Segmentation Design – Sample

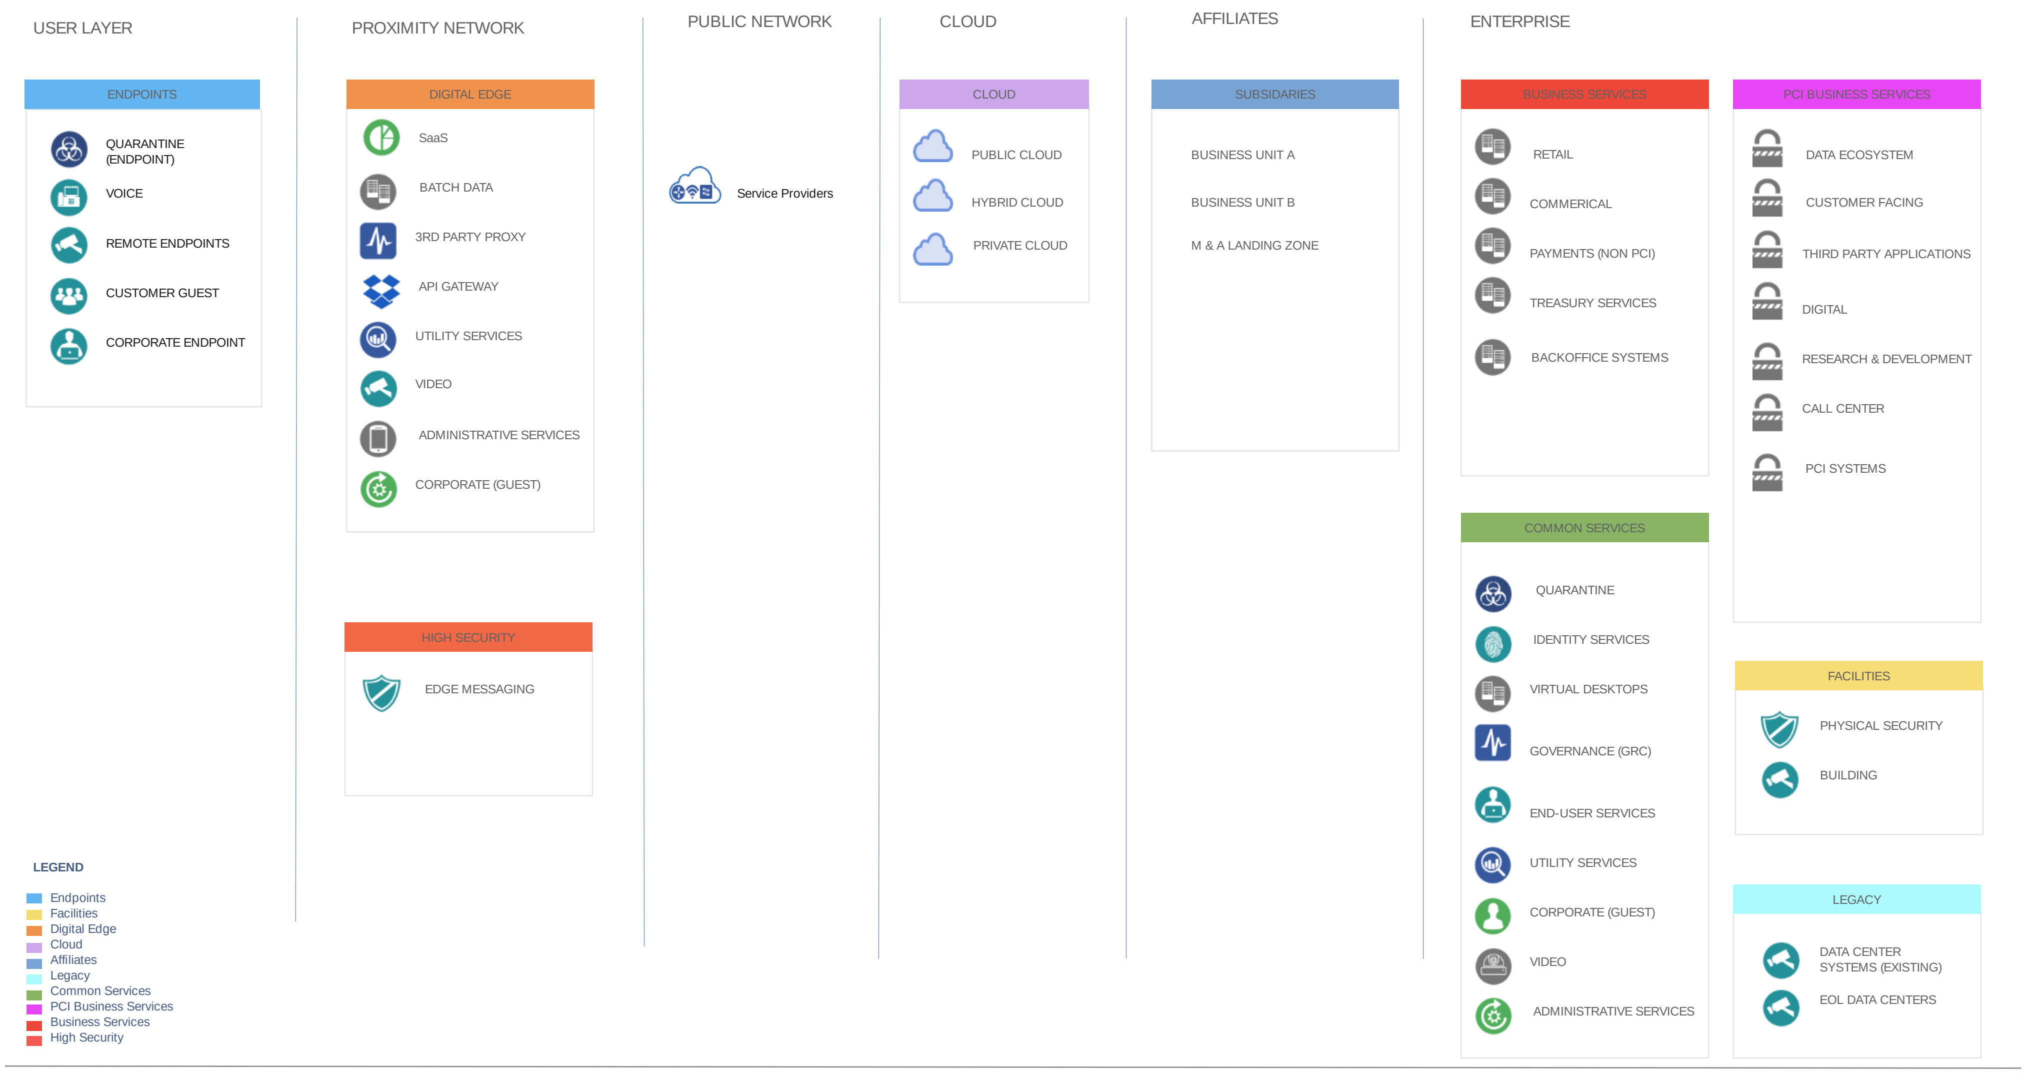

Zero Trust Capabilities – Enterprise Data Centers

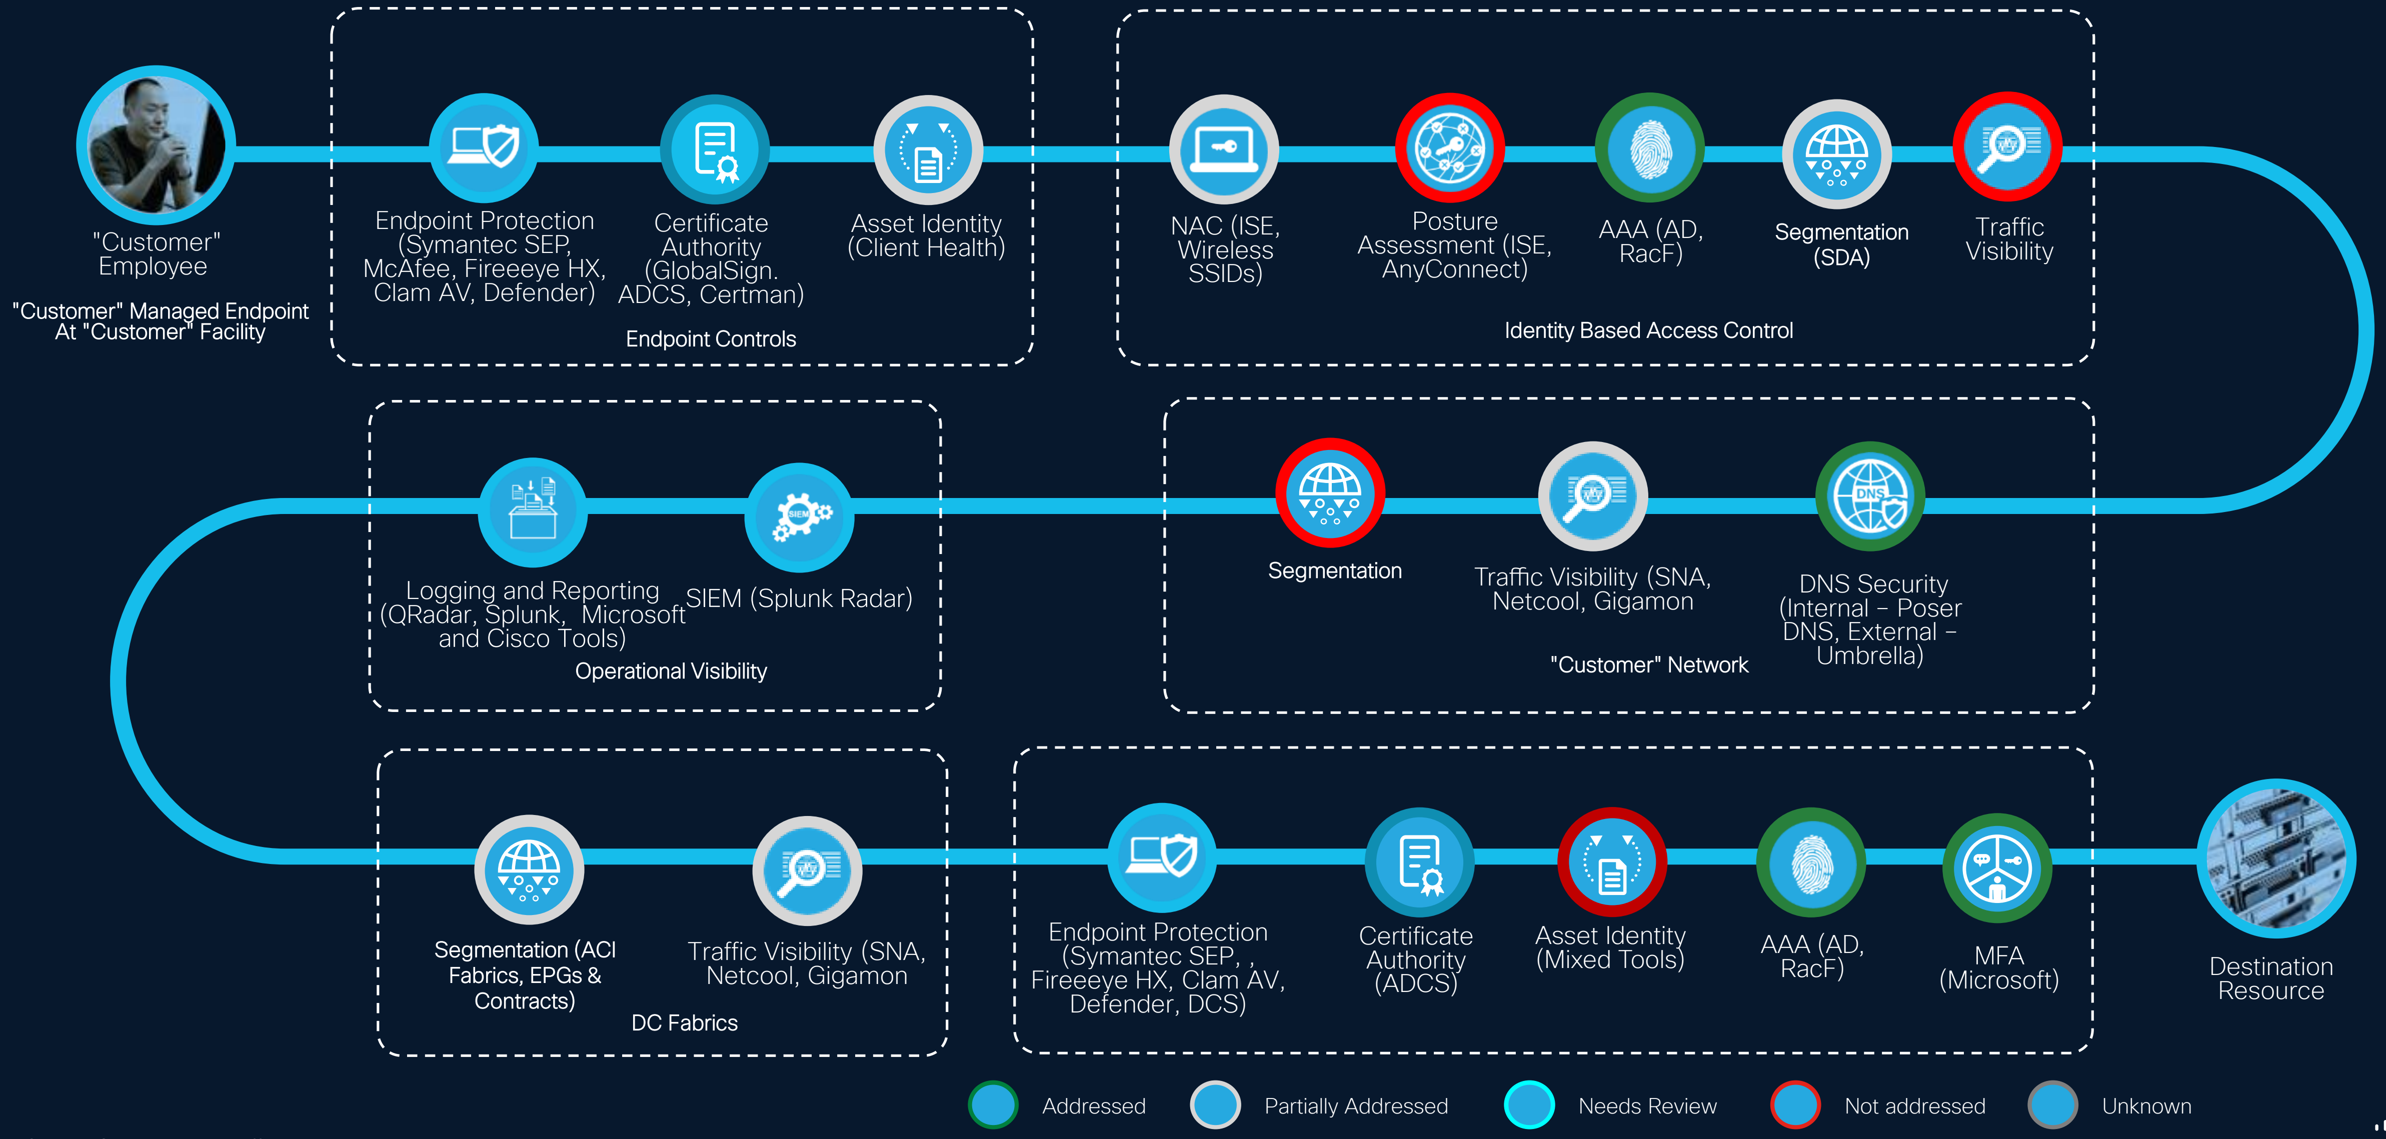

Paradigm Shift: Actions to Take

Security Paradigm Shift: Actions to Take In 7-Steps

Map the Organization and Define Security Domains

Identify organization processes, data flows, and risk zones.	Establish Zero Trust segmentation aligned with compliance.	Define guardrails for AI Agents and integrations across the enterprise.

Develop Integrated Policies for Zero Trust, AI, and PQC

Refine security policies to incorporate Responsible AI principles and PQC transition plans.	Align with NIST AI RMF, ISO 27001, and industry-specific regulations.

Build an AI-Ready Privacy and Risk Management Program

Conduct business impact assessments and model explainability reviews.	Evaluate third-party AI risks and cryptographic dependencies.

Design Secure, Scalable Architecture

Construct data center and cloud environments with vector database encryption and PQC integration.	Implement Zero Trust access controls across data, models, and APIs.

Enable Enterprise-Wide Visibility and Segmentation

Use NetFlow and behavioral analytics to map traffic and tag AI workloads.	Apply telemetry and RAG-aware segmentation with PQC-protected communications.

Pilot and Validate Enforcement Controls

Lab test segmentation, AI Agent guardrails, and PQC cryptographic performance.	Launch pilot enforcement across key domains and critical data paths.

Enforce, Monitor Continuously, and Iterate

Deploy enforcement across segments.	Monitor AI threat activity, cryptographic health, and ZT policy drift.	Adapt controls as AI and PQC standards evolve.

Cisco Services Call to Action

Next Steps

Explore our demos in the **Cisco Showcase** of the **World of Solutions**.

Meet with our **CX Subject Matter Experts**. Inquire about availability at the welcome desk in the CX Connections Lounge in the World of Solutions.

Attend our Center Stage and PSO sessions. Scan this code to view all our sessions.

Play the **CX Big Wave Sweepstak**es while you're at Cisco Live for a chance to win a gaming & surfing prize package.

Cisco
Customer Experience

For more information on Cisco Customer Experience and how we can help you, **visit cisco.com/go/services**.

cisco


Questions?

Thank you