CISCO Engage | Tech Day Miami

Segmentation Untangled

Streamlining Hybrid Data Center Security

Jamey Heary
Cisco Chief Security Architect, DSE

2/19/26

Zero Trust Use Cases

Segmentation

Admin Access

Remote Worker

On-Premises

Application to Application

Application to Mode

Zero Trust End to End

User to App, App to App

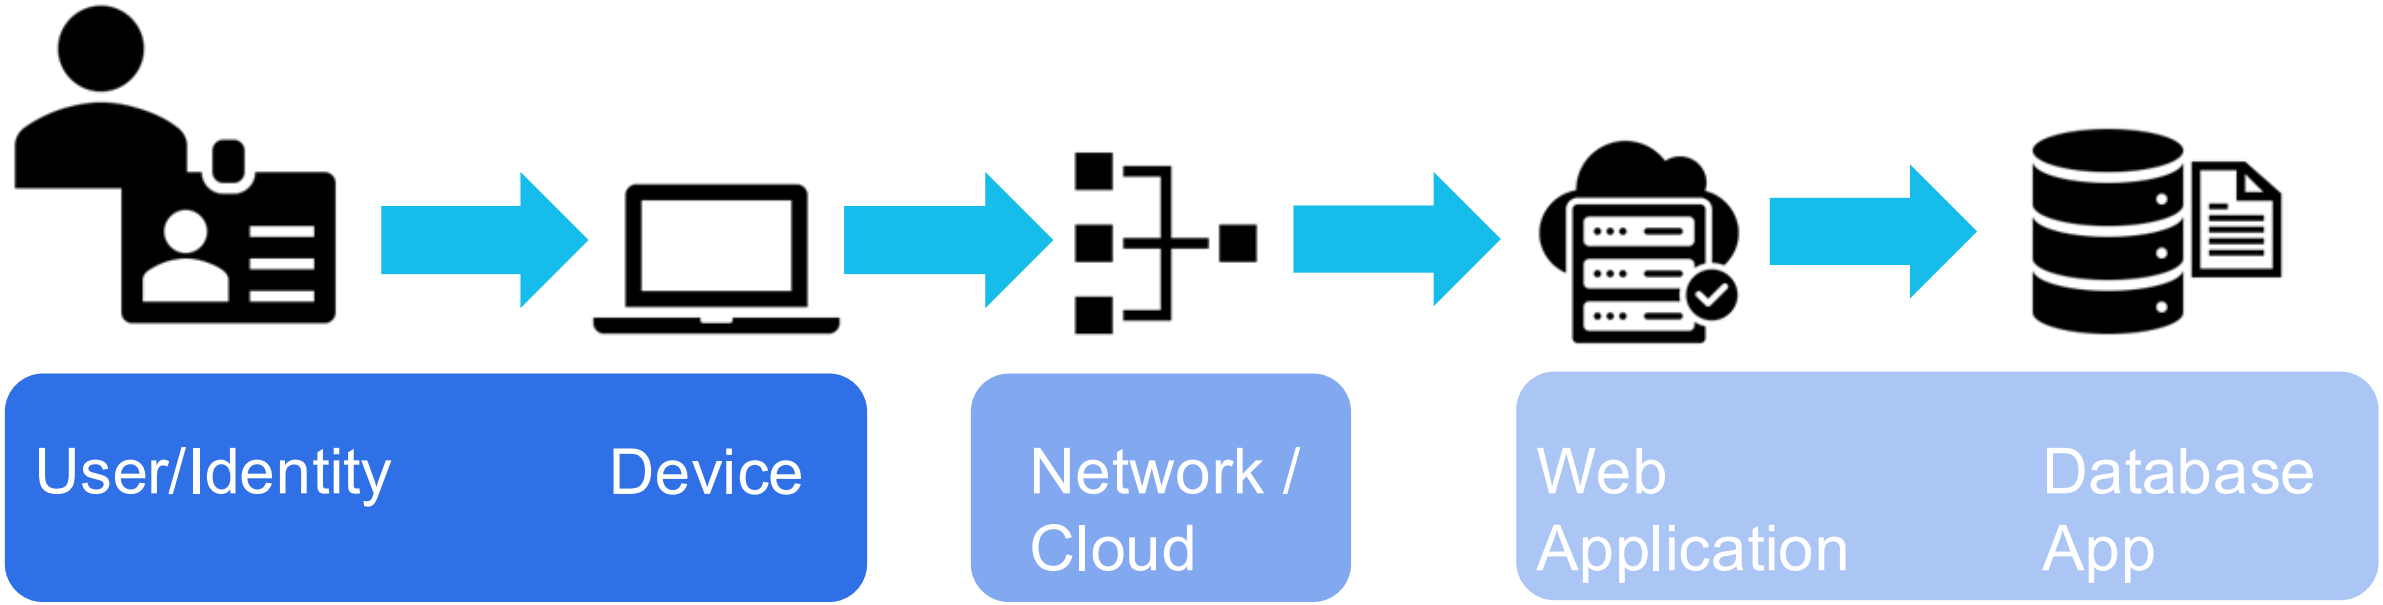

Zero Trust App to App/Service

Is it this Simple?

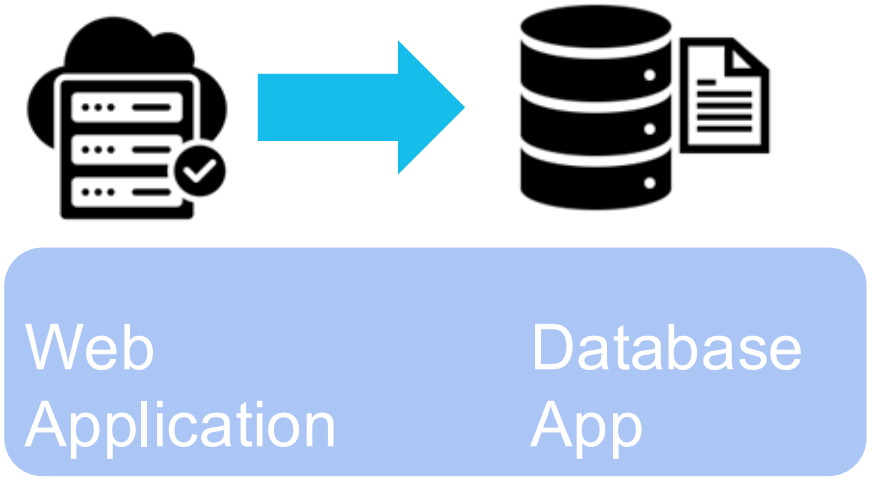

Zero Trust Use Case Detailed Examples

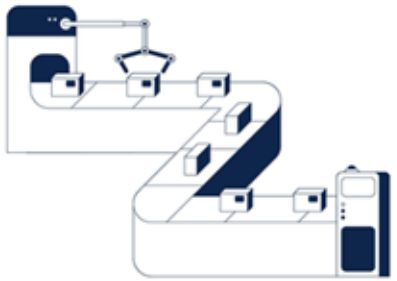

App-to-App

- App-to-App Inter multi-cloud
- App to app intra zone
- App to Internet
- App to API/public app SaaS
- Apps to/with GenAI Model

Application **workloads** evolution

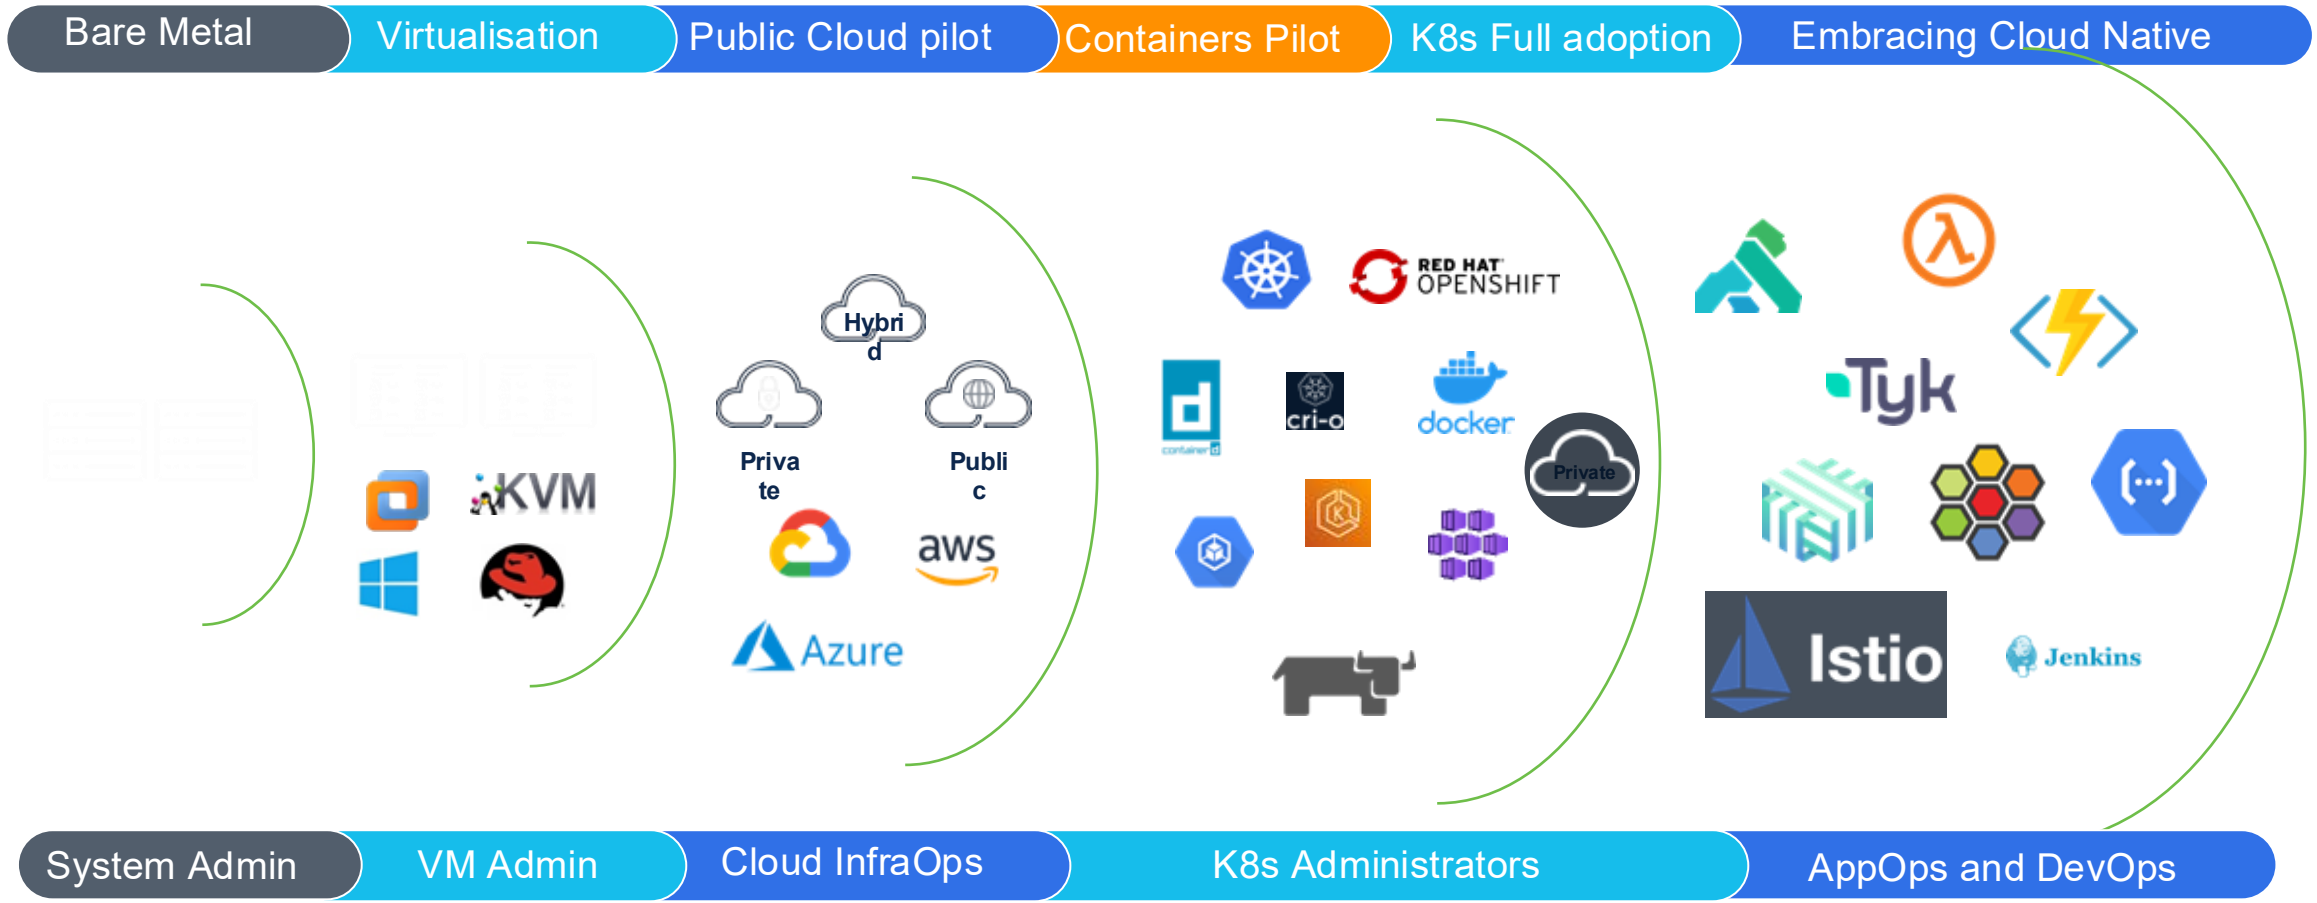

Simple Modern Application

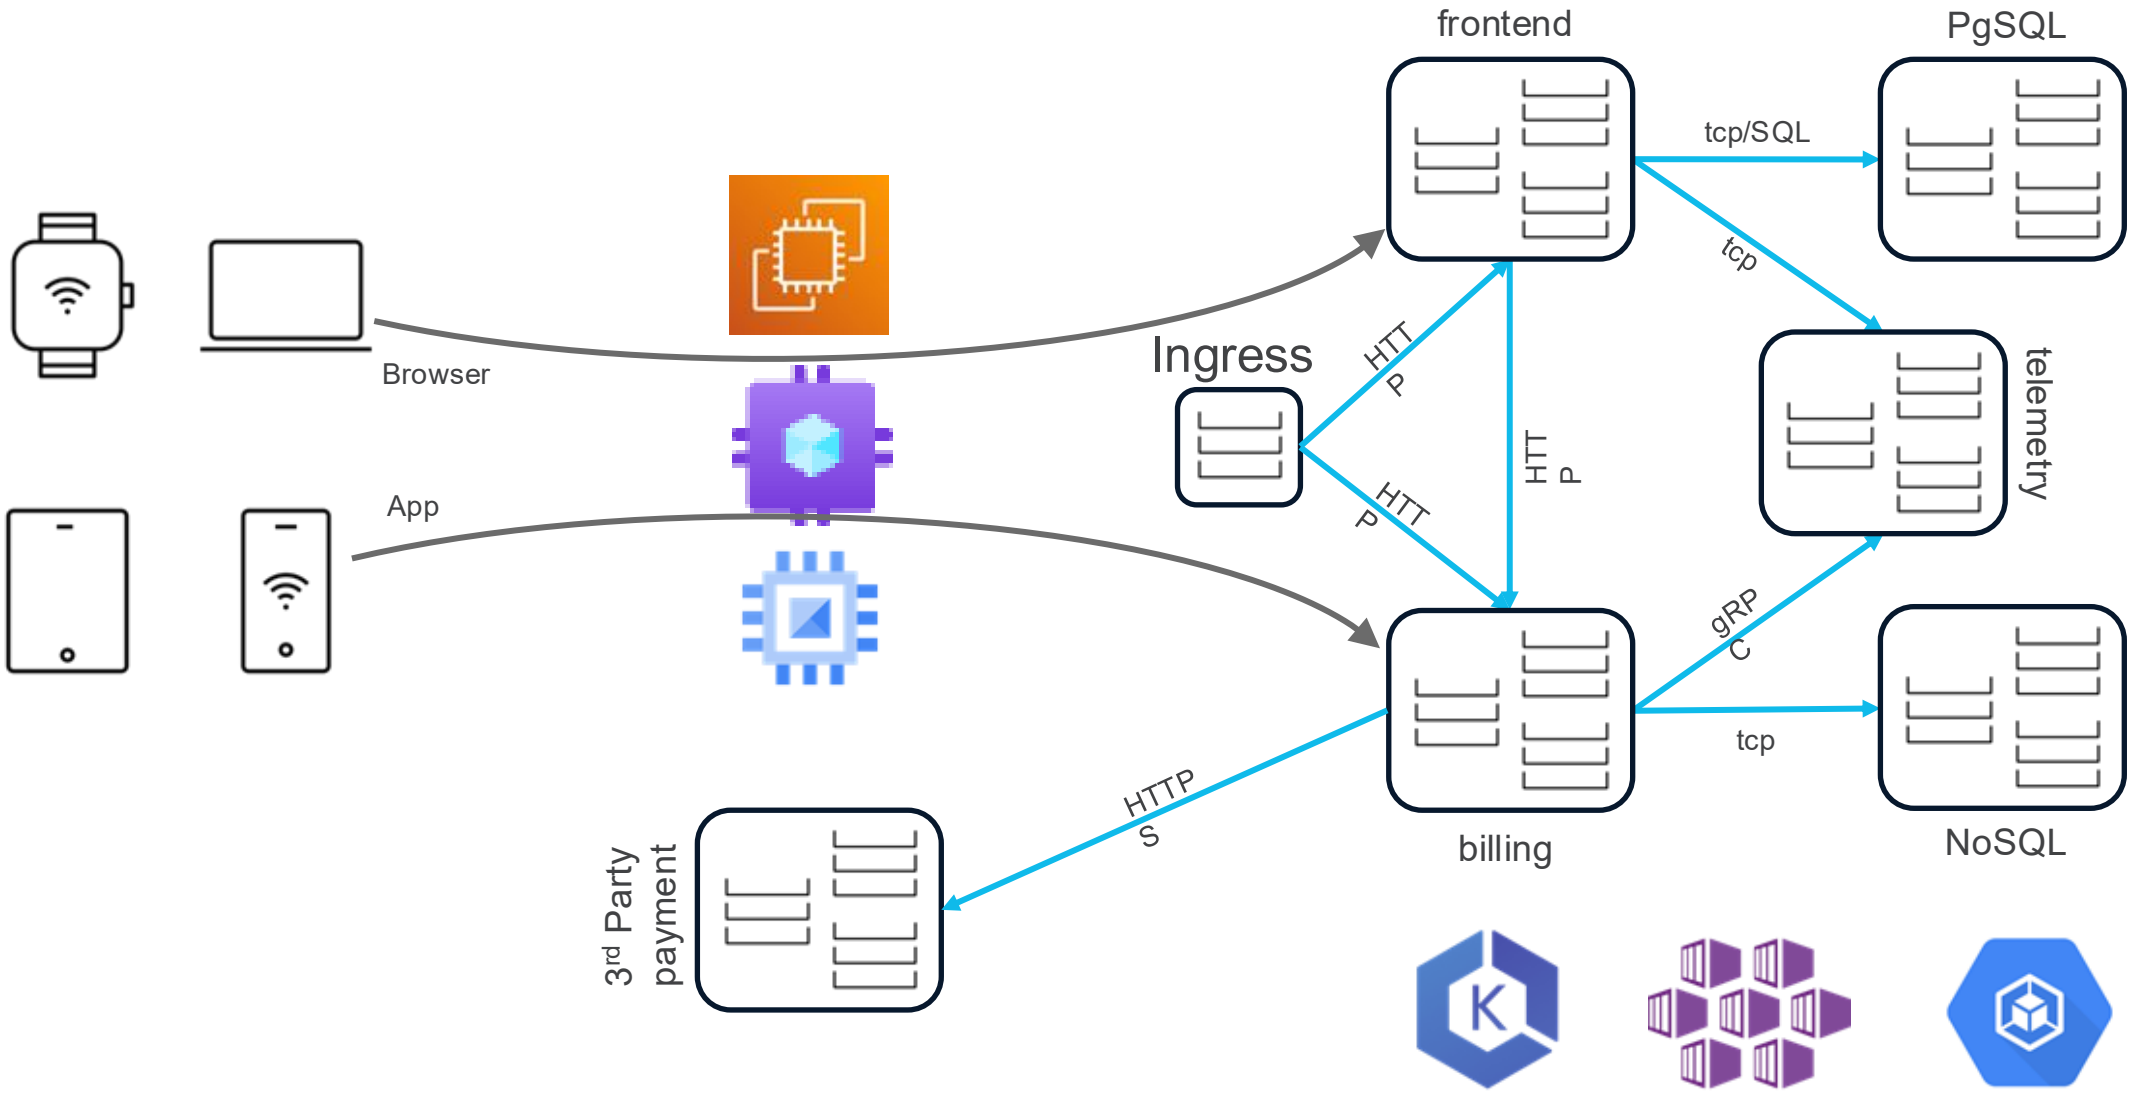

Their dream is your nightmare

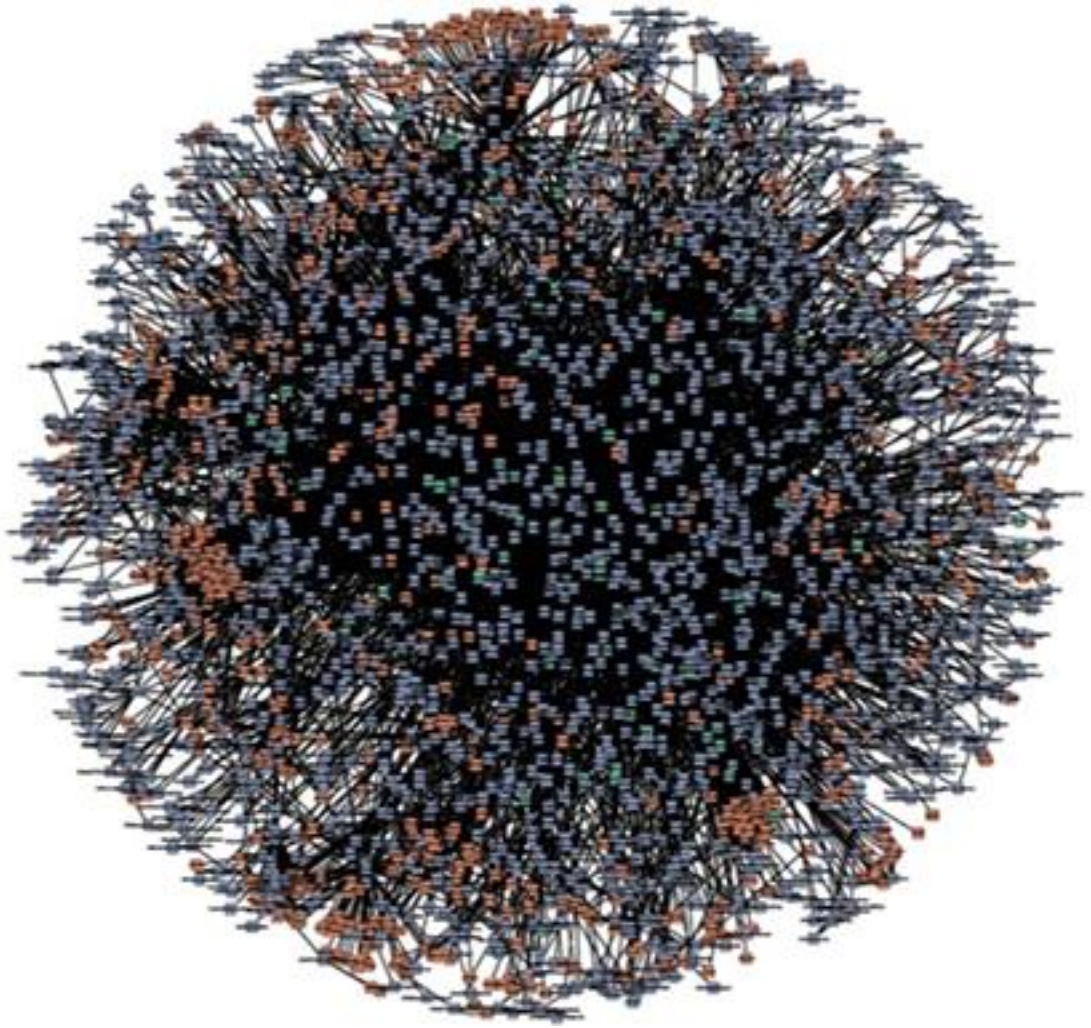

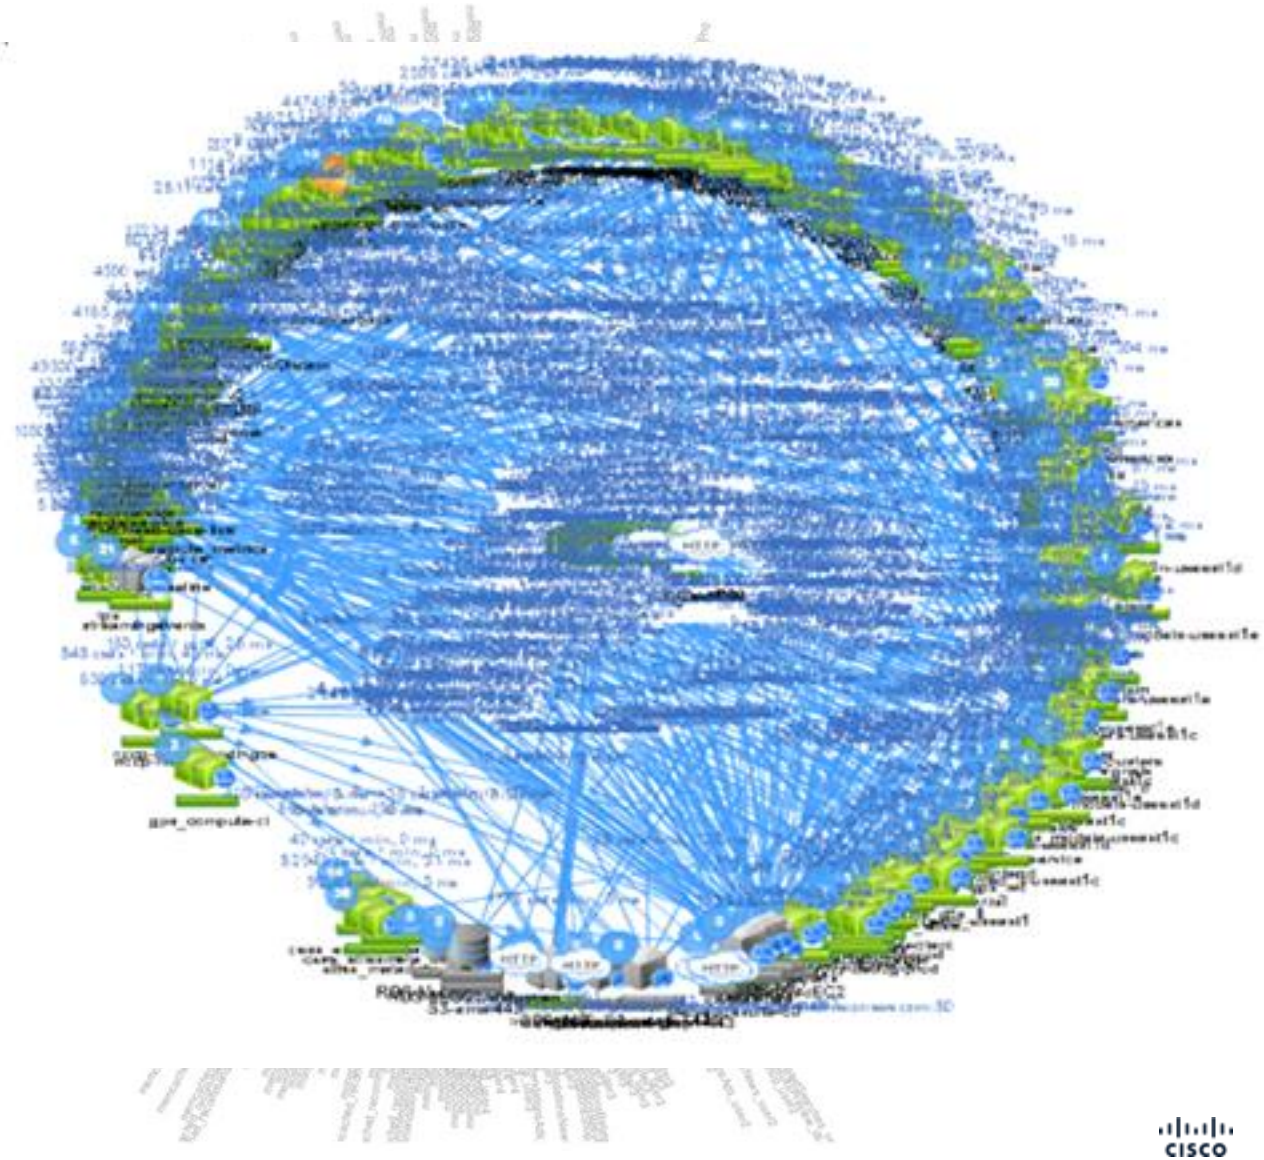

Achieving Zero Trust for Modern Applications

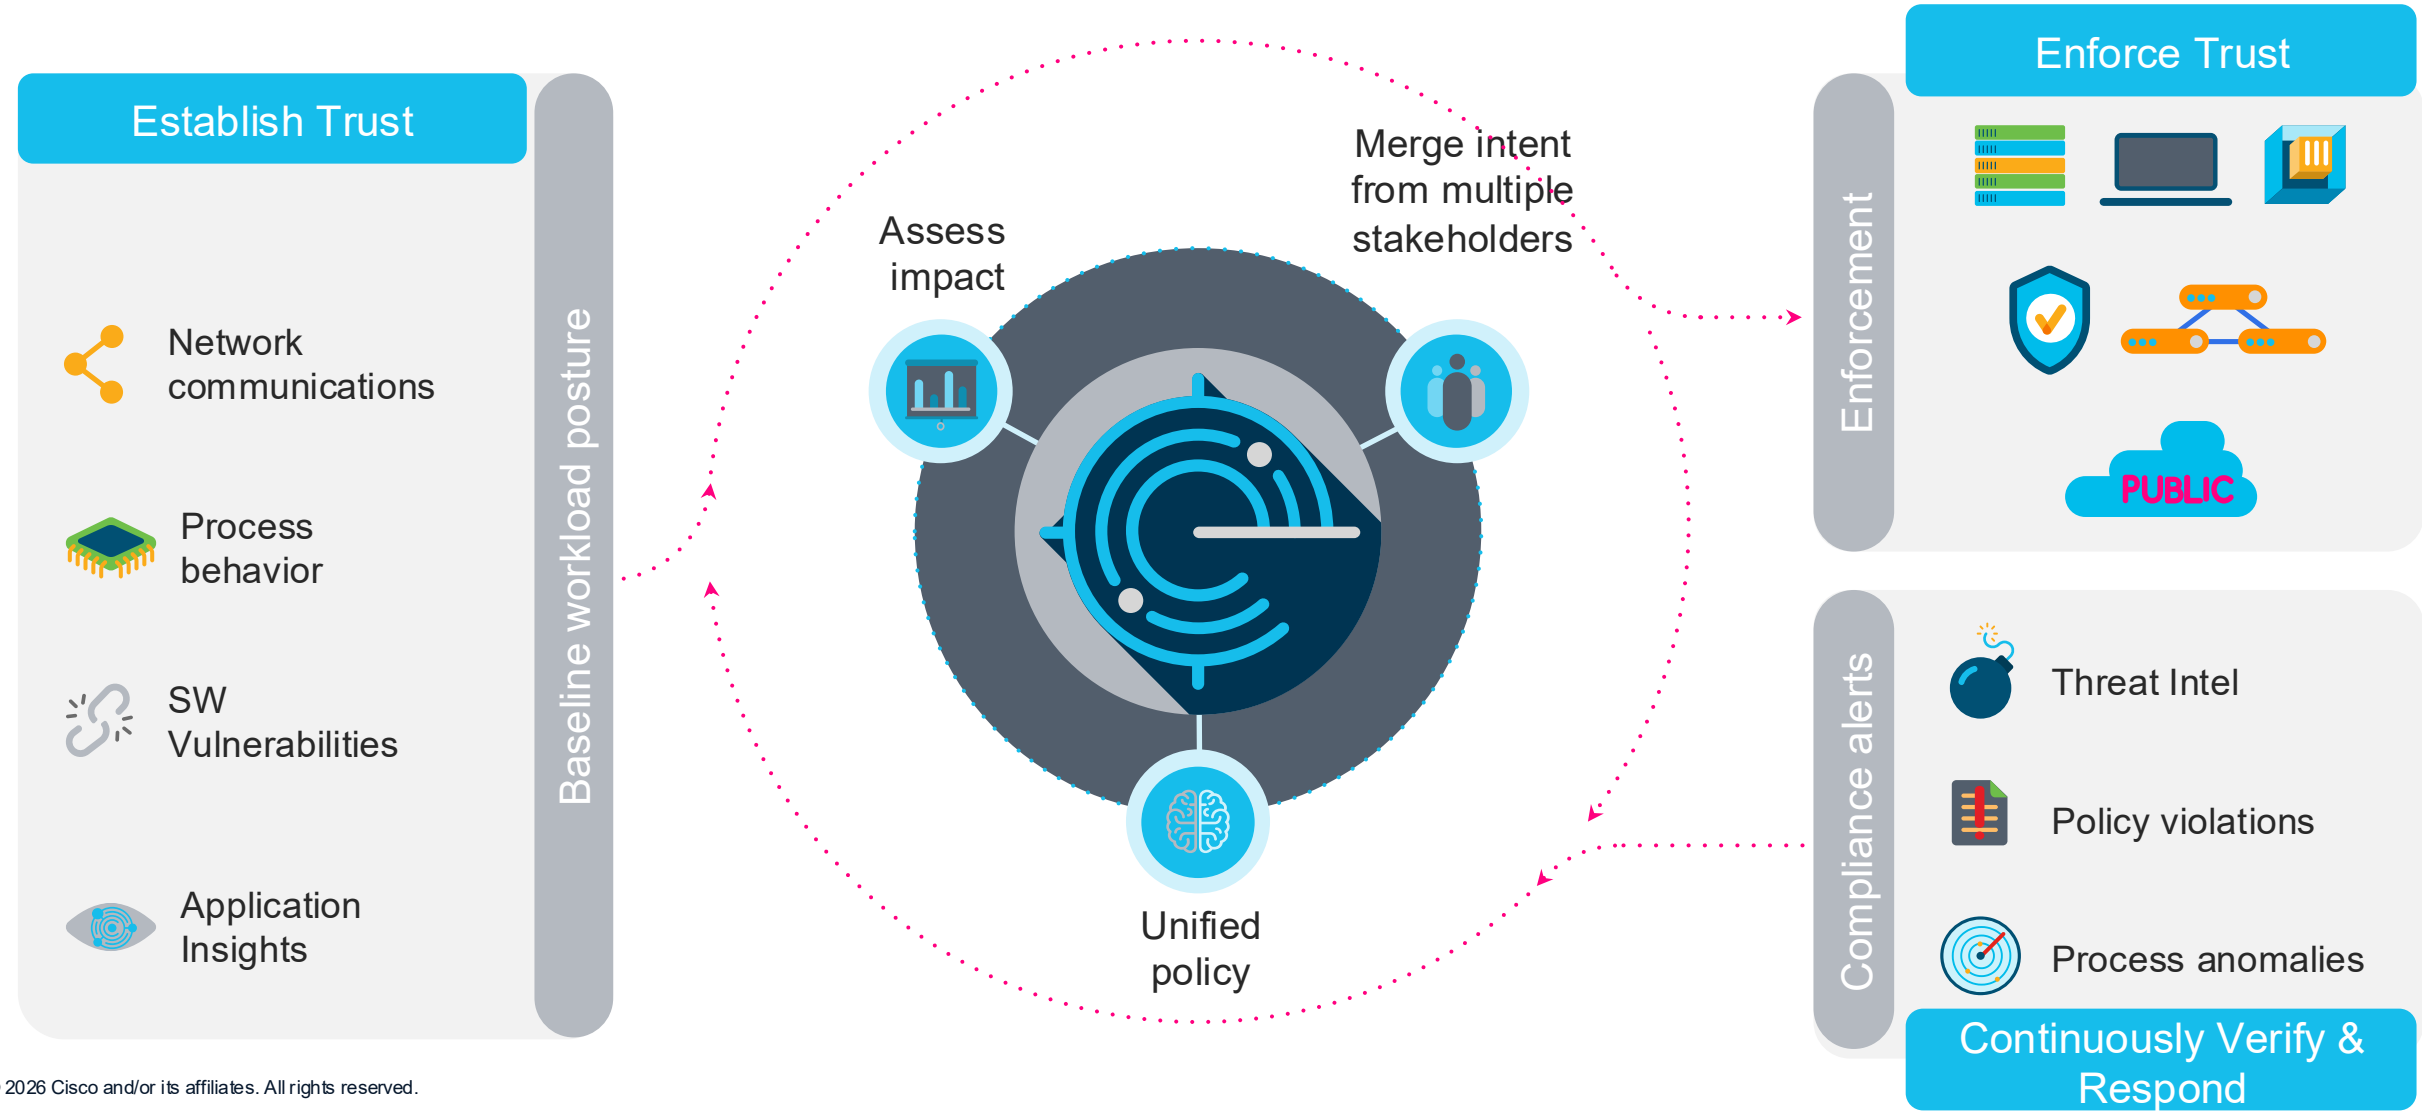

Hybrid Cloud Zero Trust Segmentation

Defense in Depth via Hybrid Mesh of enforcement points

Edge Firewall

Macro-Segmentation

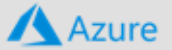

Micro-Segmentation

Hypershield

Unified Segmentation

Perimeter Defense

- Ingress & Egress Security
- Threat inspection at the data center or cloud edge.

Zones

- Segment zones within your data center and cloud.
- Supplementary coverage for workloads without agents.

App Segmentation

- Zero trust micro-segmentation enforcement at the workload
- Automated policy discovery and compliance

Hybrid Cloud Security

Zero Trust

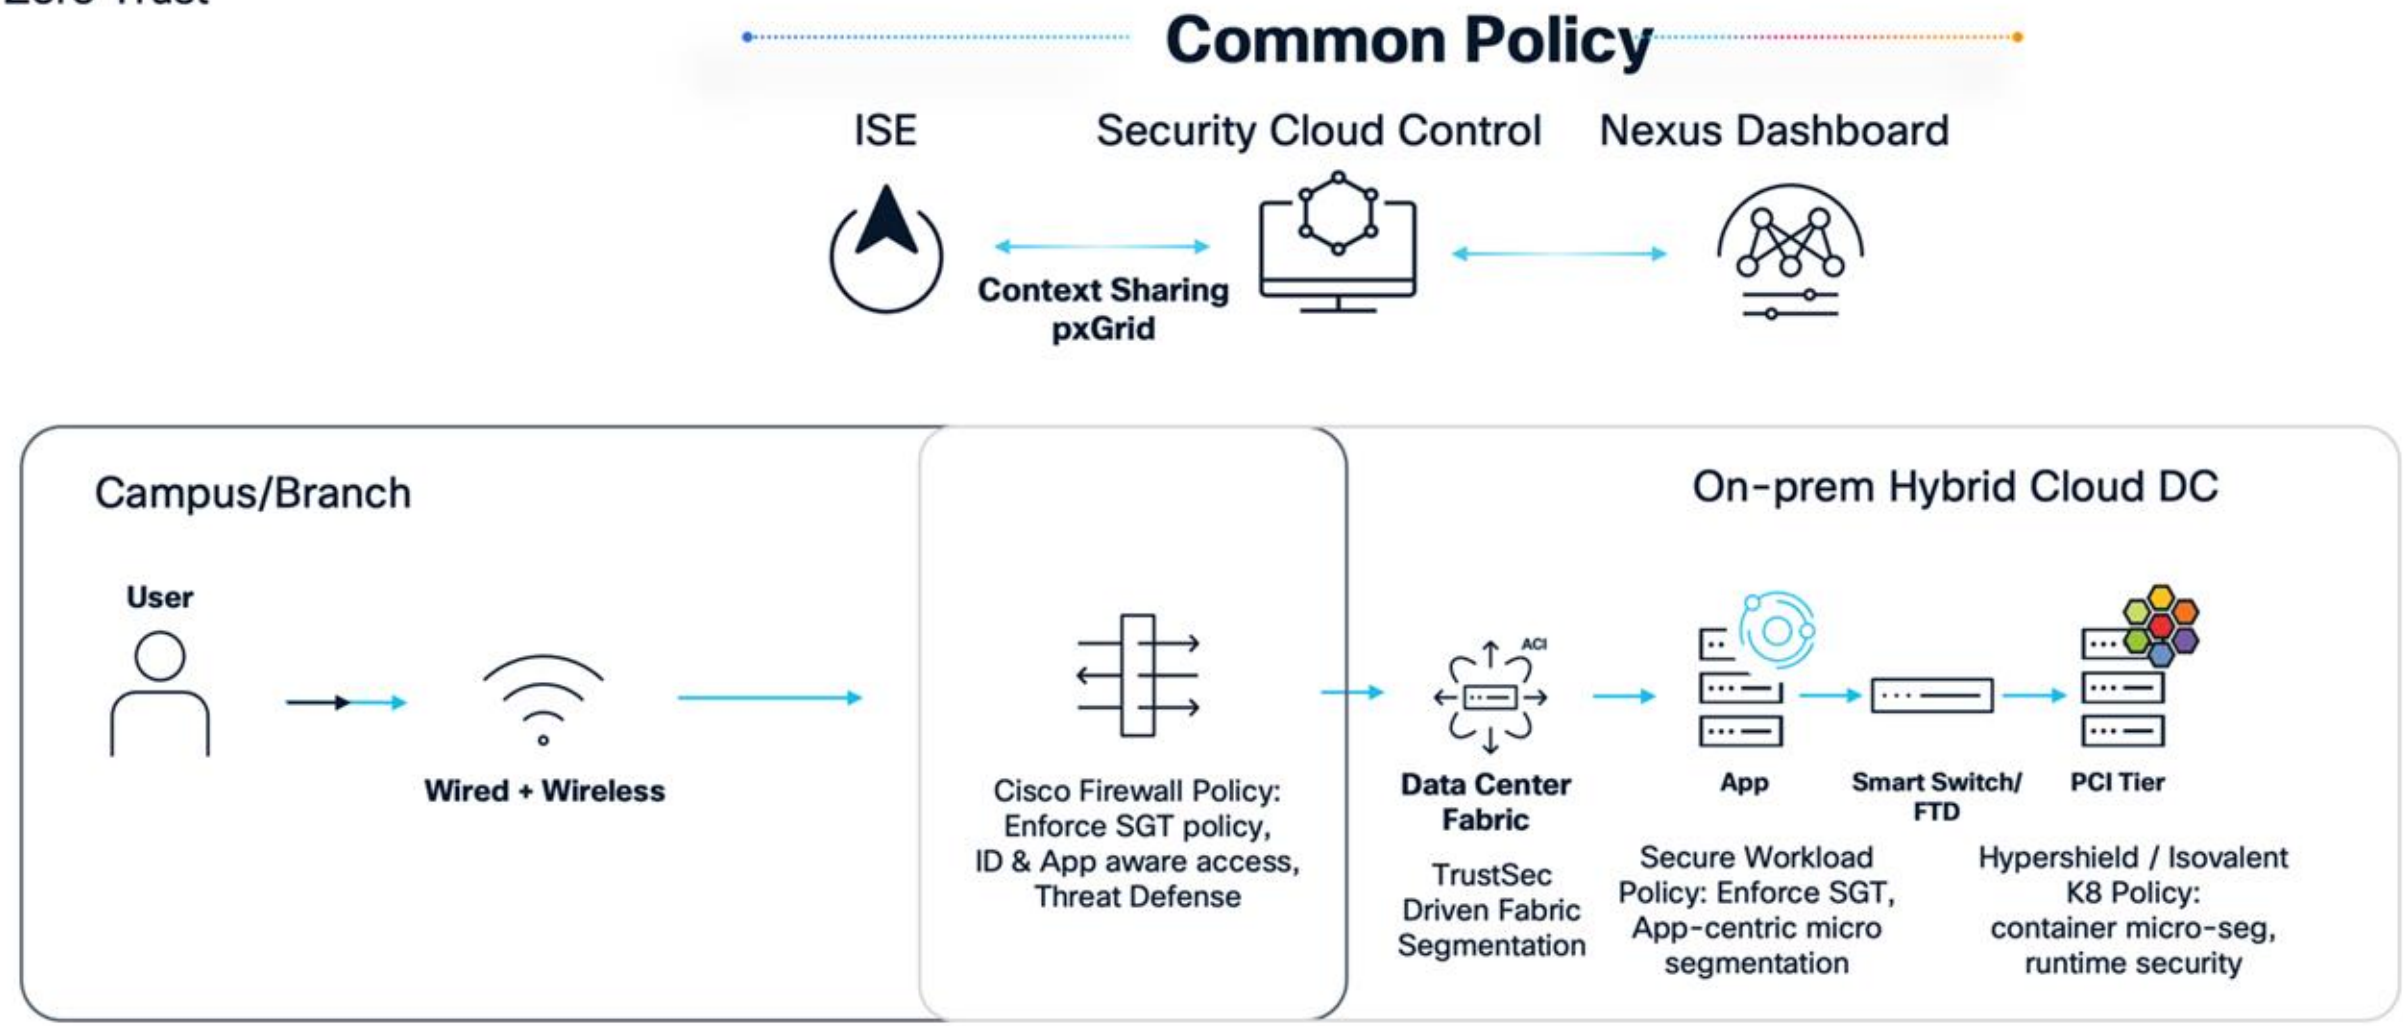

Application **workloads** evolution

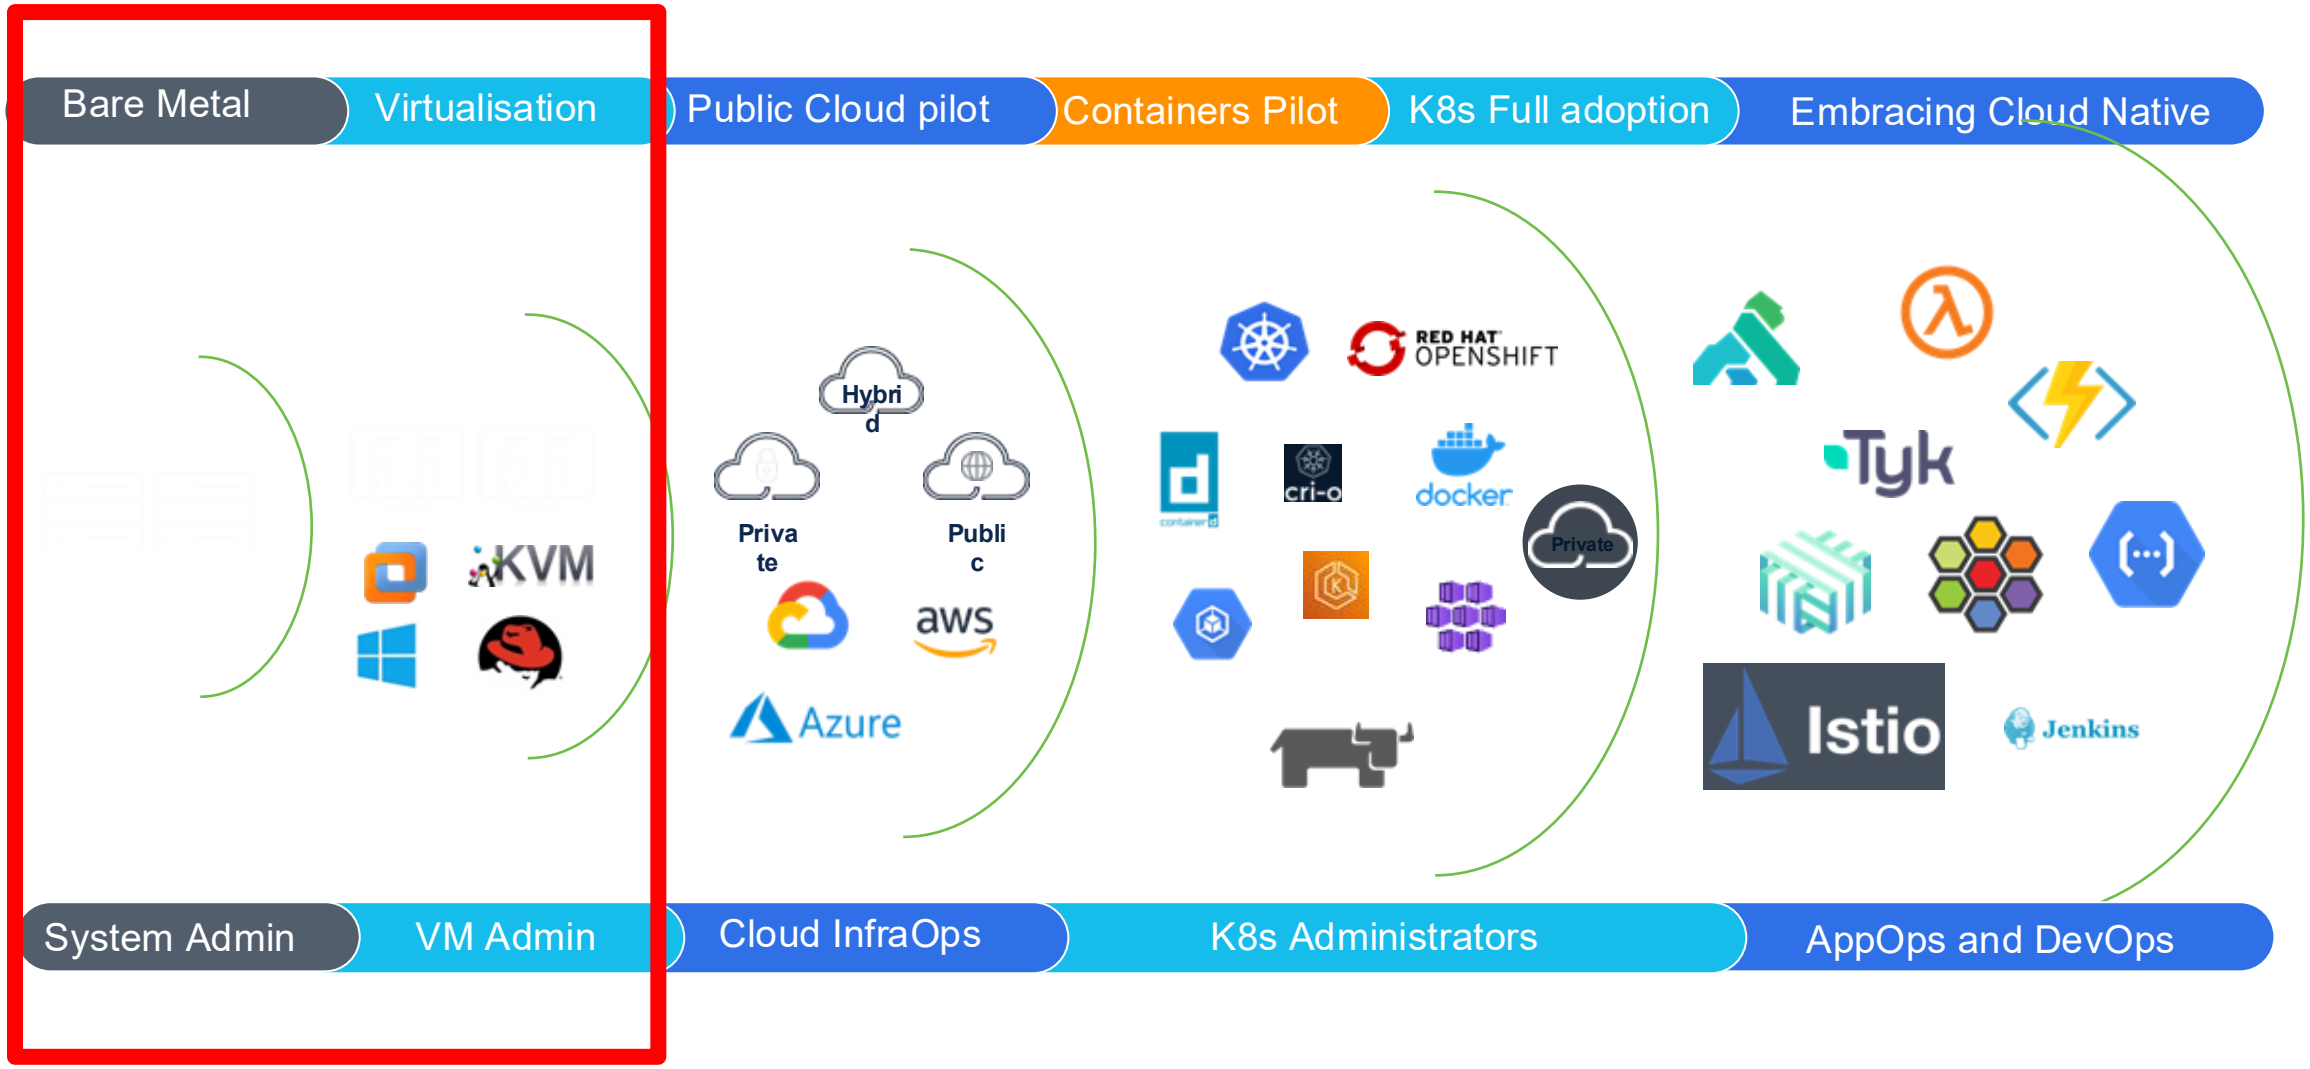

Enforce Trust in the DC with Cisco ACI

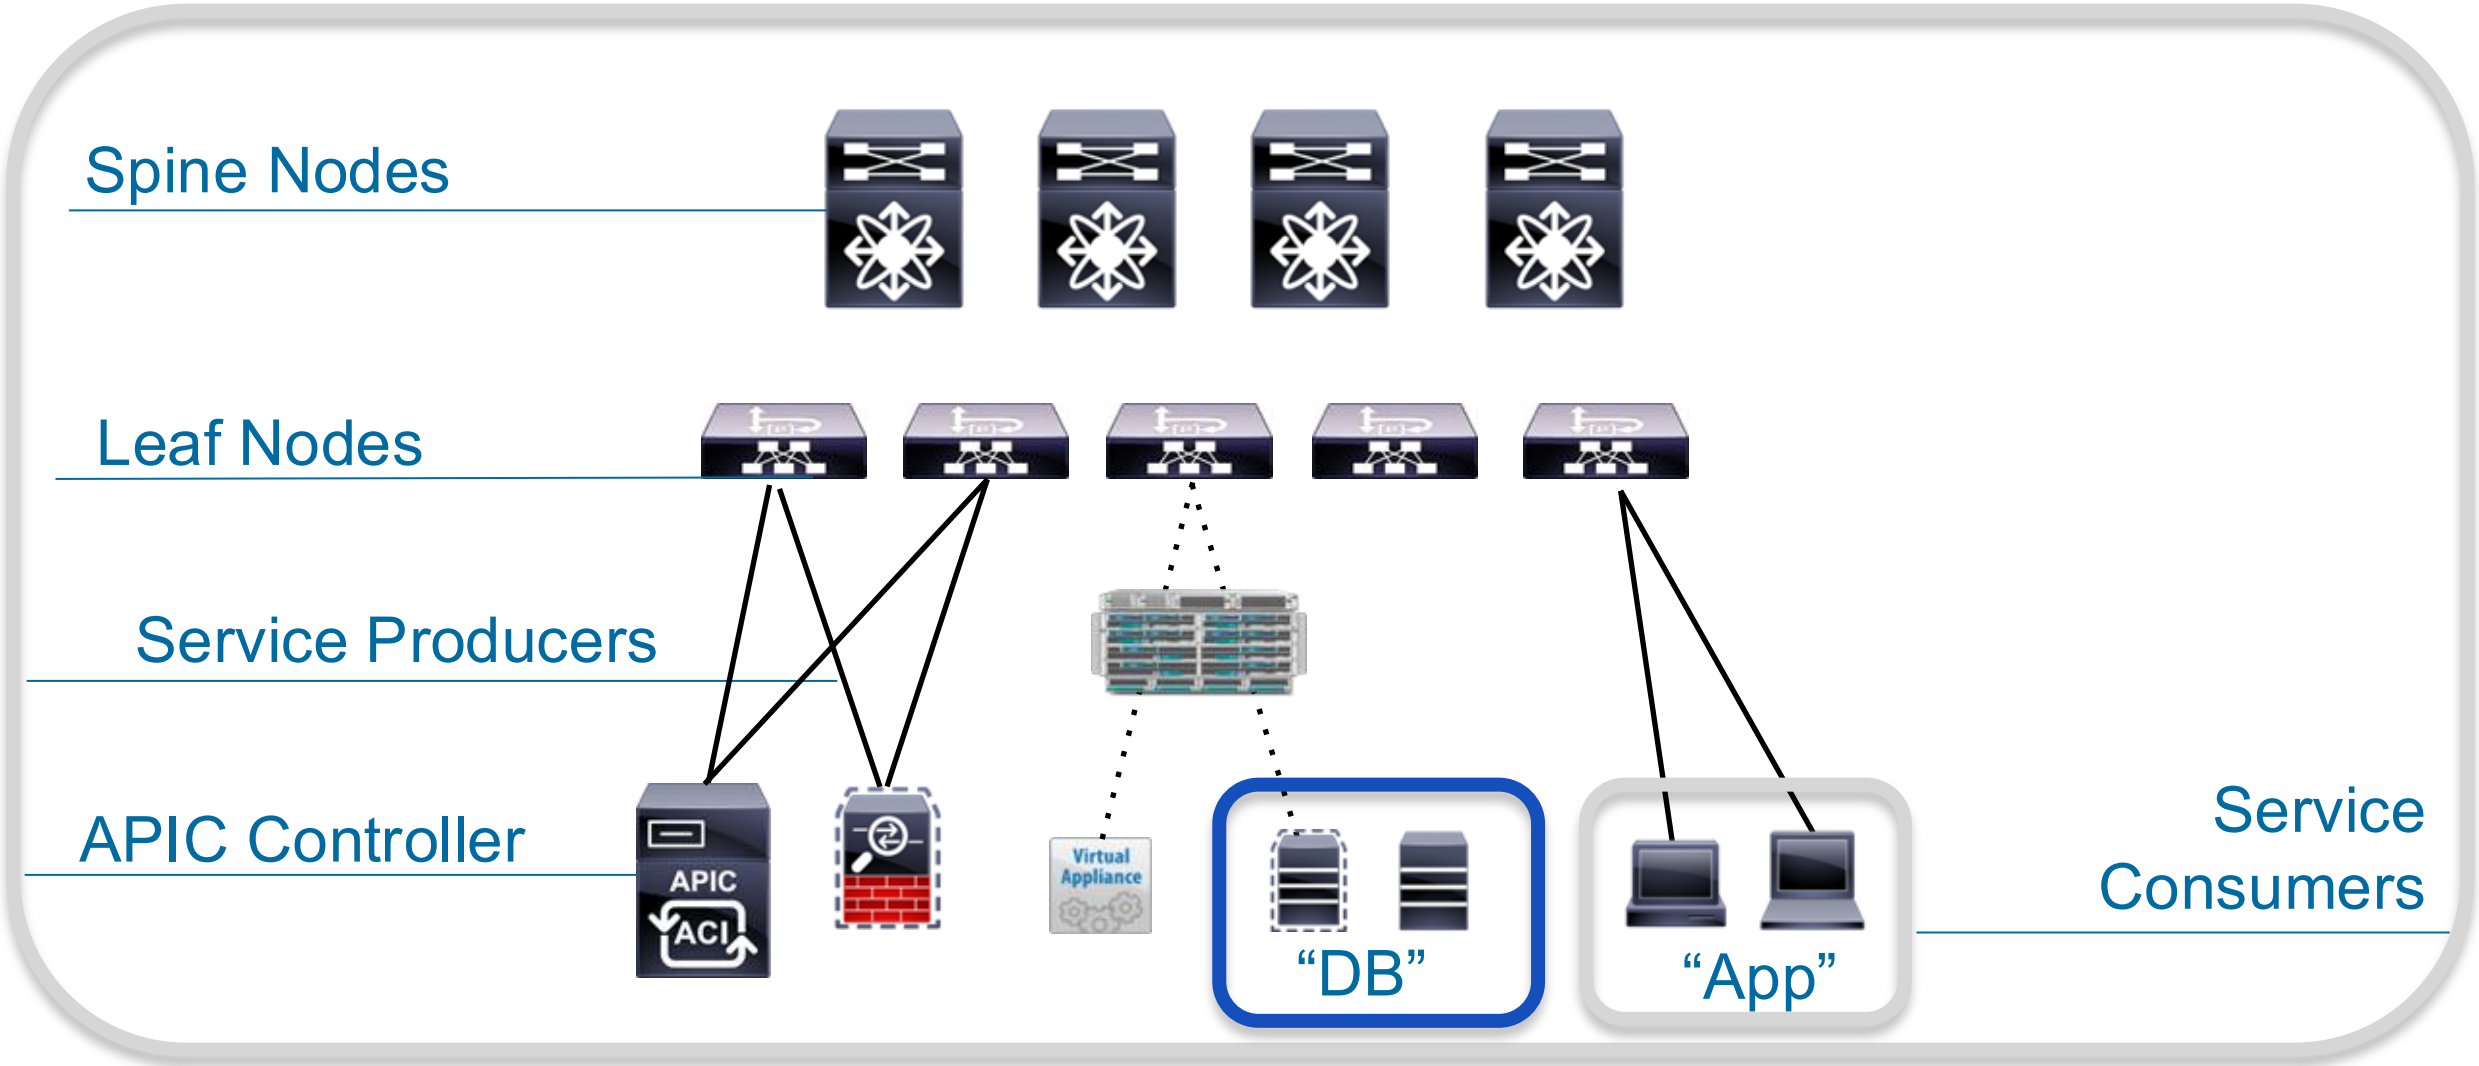

Is subnet the right way to establish trust groups?

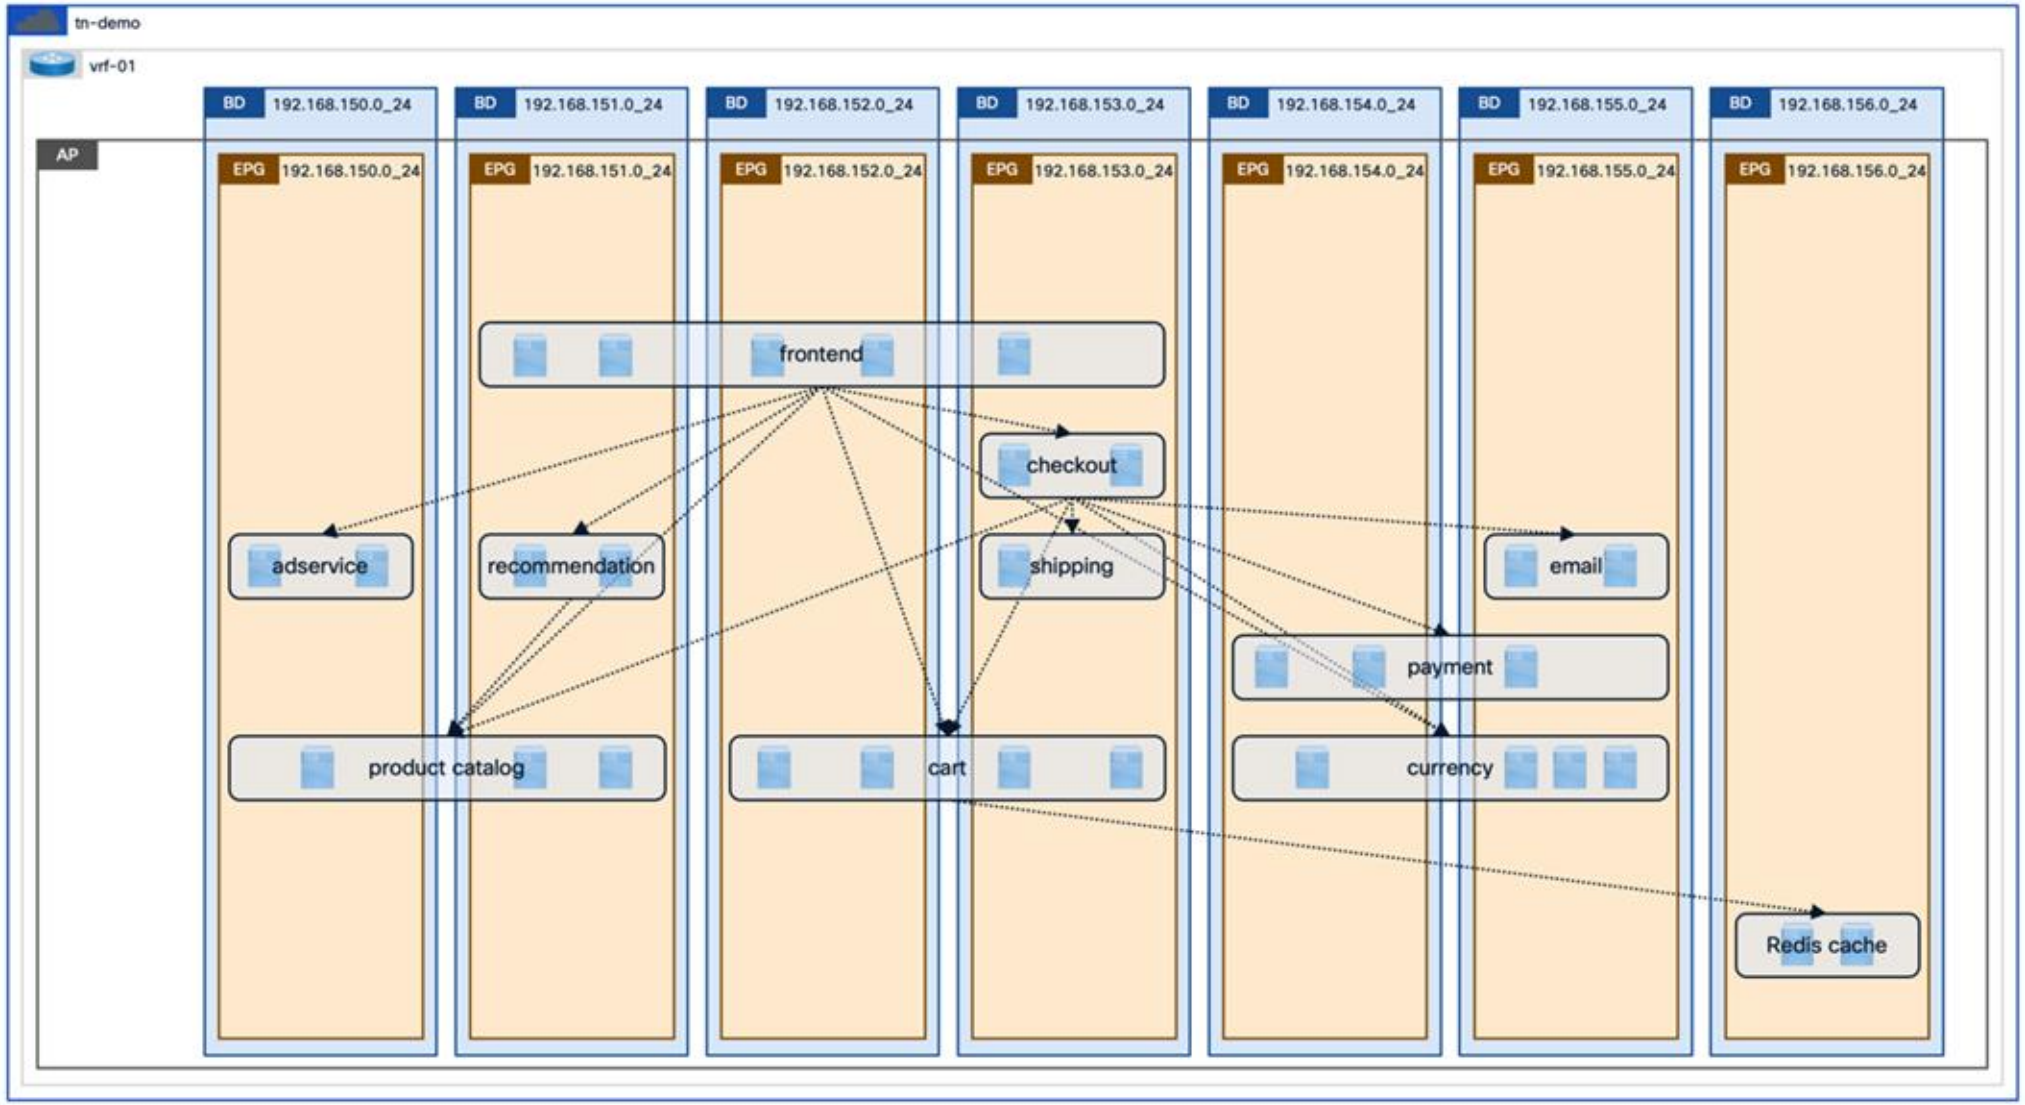

cisco

Zero Trust Segmentation in the Data Center:
EPGs, ESGs, and Contracts

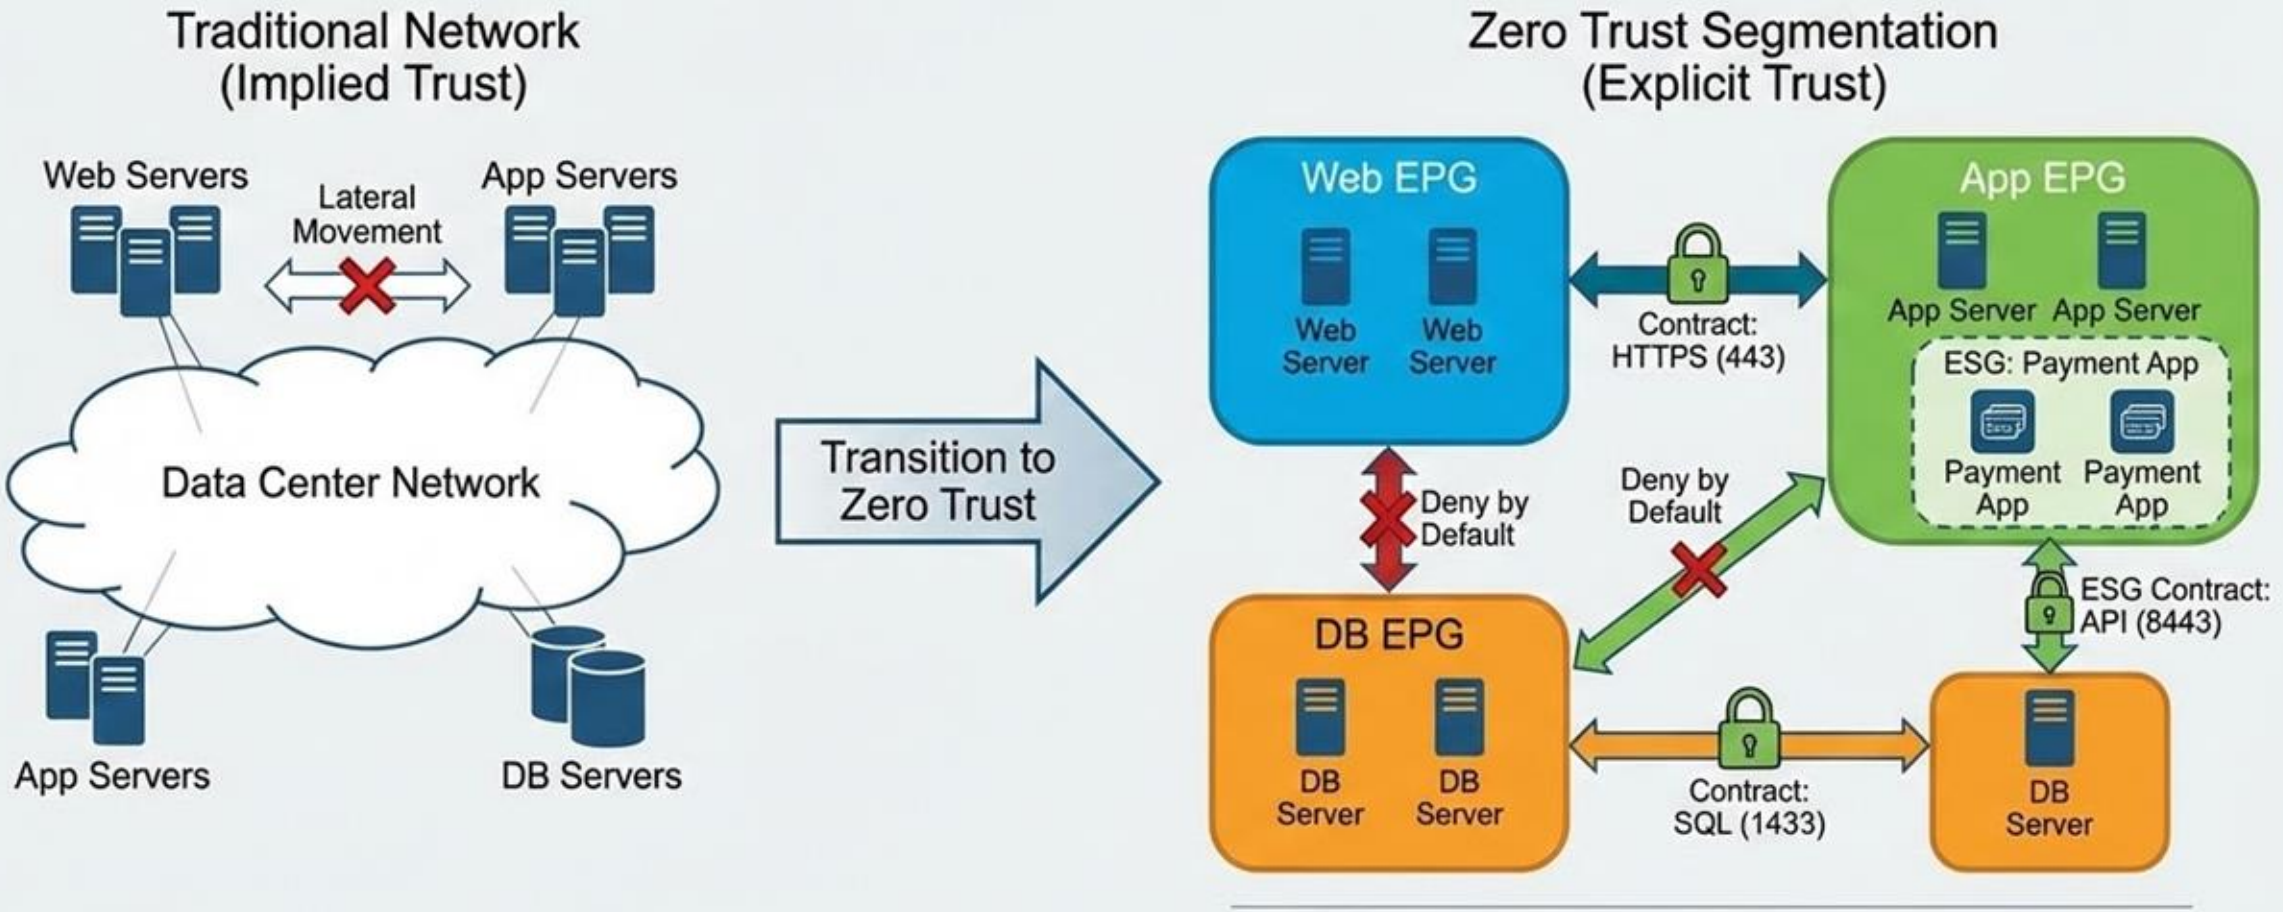

Endpoint Security Groups (ESGs) vs. Endpoint Groups (EPGs) in Cisco ACI

How ESG Works

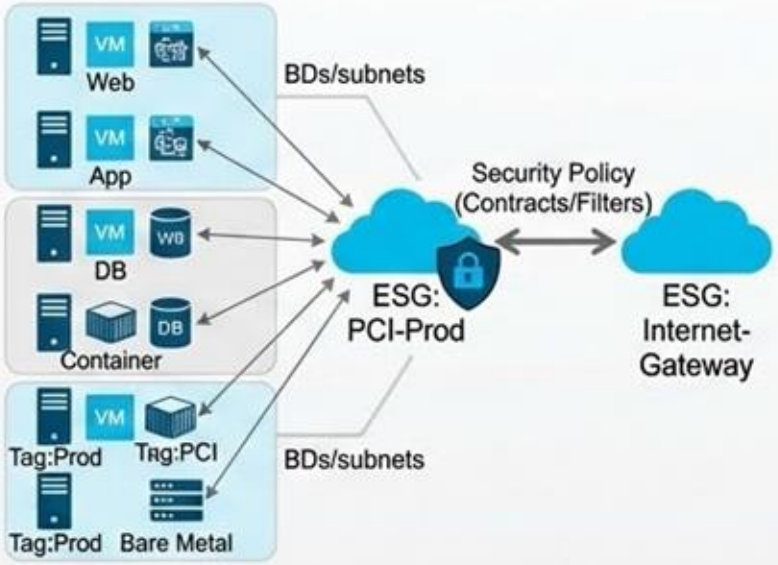

- Logical grouping for security, independent of network topology (BDs, VRFs).
- Membership defined by attributes (tags, VM labels, OS type, etc.).
- Enforces security policies (contracts) based on these logical groupings.
- Enables fine-grained micro-segmentation.

Differences from EPG

Feature	EPG vs. ESG	
Grouping Basis:	EPG: Network-centric (Subnet/VLAN/IP).	ESG: Attribute-centric (Tags, Labels, Identity).
Network Dependency:	EPG: Tightly coupled to BD & VRF.	ESG: Network-agnostic, decoupled from BD.
Segmentation Level:	EPG: Coarse-grained (Network Segment).	ESG: Fine-grained (Micro-segmentation).
Policy Scope:	EPG: Intra-EPG communication allowed by default (unless isolated).	ESG: Strict policy enforcement, zero trust model.

Why Use ESG instead of EPG

Decouple Security from Network: Define security policies without constraints of network addressing or topology changes.

Simplify Dynamic Environments: Ideal for containers, cloud, and highly dynamic workloads where IPs change frequently.

Granular Micro-Segmentation: Achieve zero trust by grouping workloads based on business function and security posture, not just location.

Unified Policy Model: Apply consistent security policies across different infrastructure types (on-prem, cloud, hybrid) using attributes.

Network agnostic endpoint security groups (ESGs)

Source/Consumer	Target/Provider	Target/Provider Port
cart	Redis cache	TCP 6379
checkout	cart	TCP 7070
	currency	TCP 7000
	email	TCP 5000
	payment	TCP 50051
	product catalog	TCP 3550
	shipping	TCP 50051
frontend	adservice	TCP 9555
	cart	TCP 7070
	checkout	TCP 5050
	currency	TCP 7000
	product catalog	TCP 3550
	recommendation	TCP 8080
	shipping	TCP 50051
outside	frontend	TCP 80/8080
recommendation	product catalog	TCP 3550

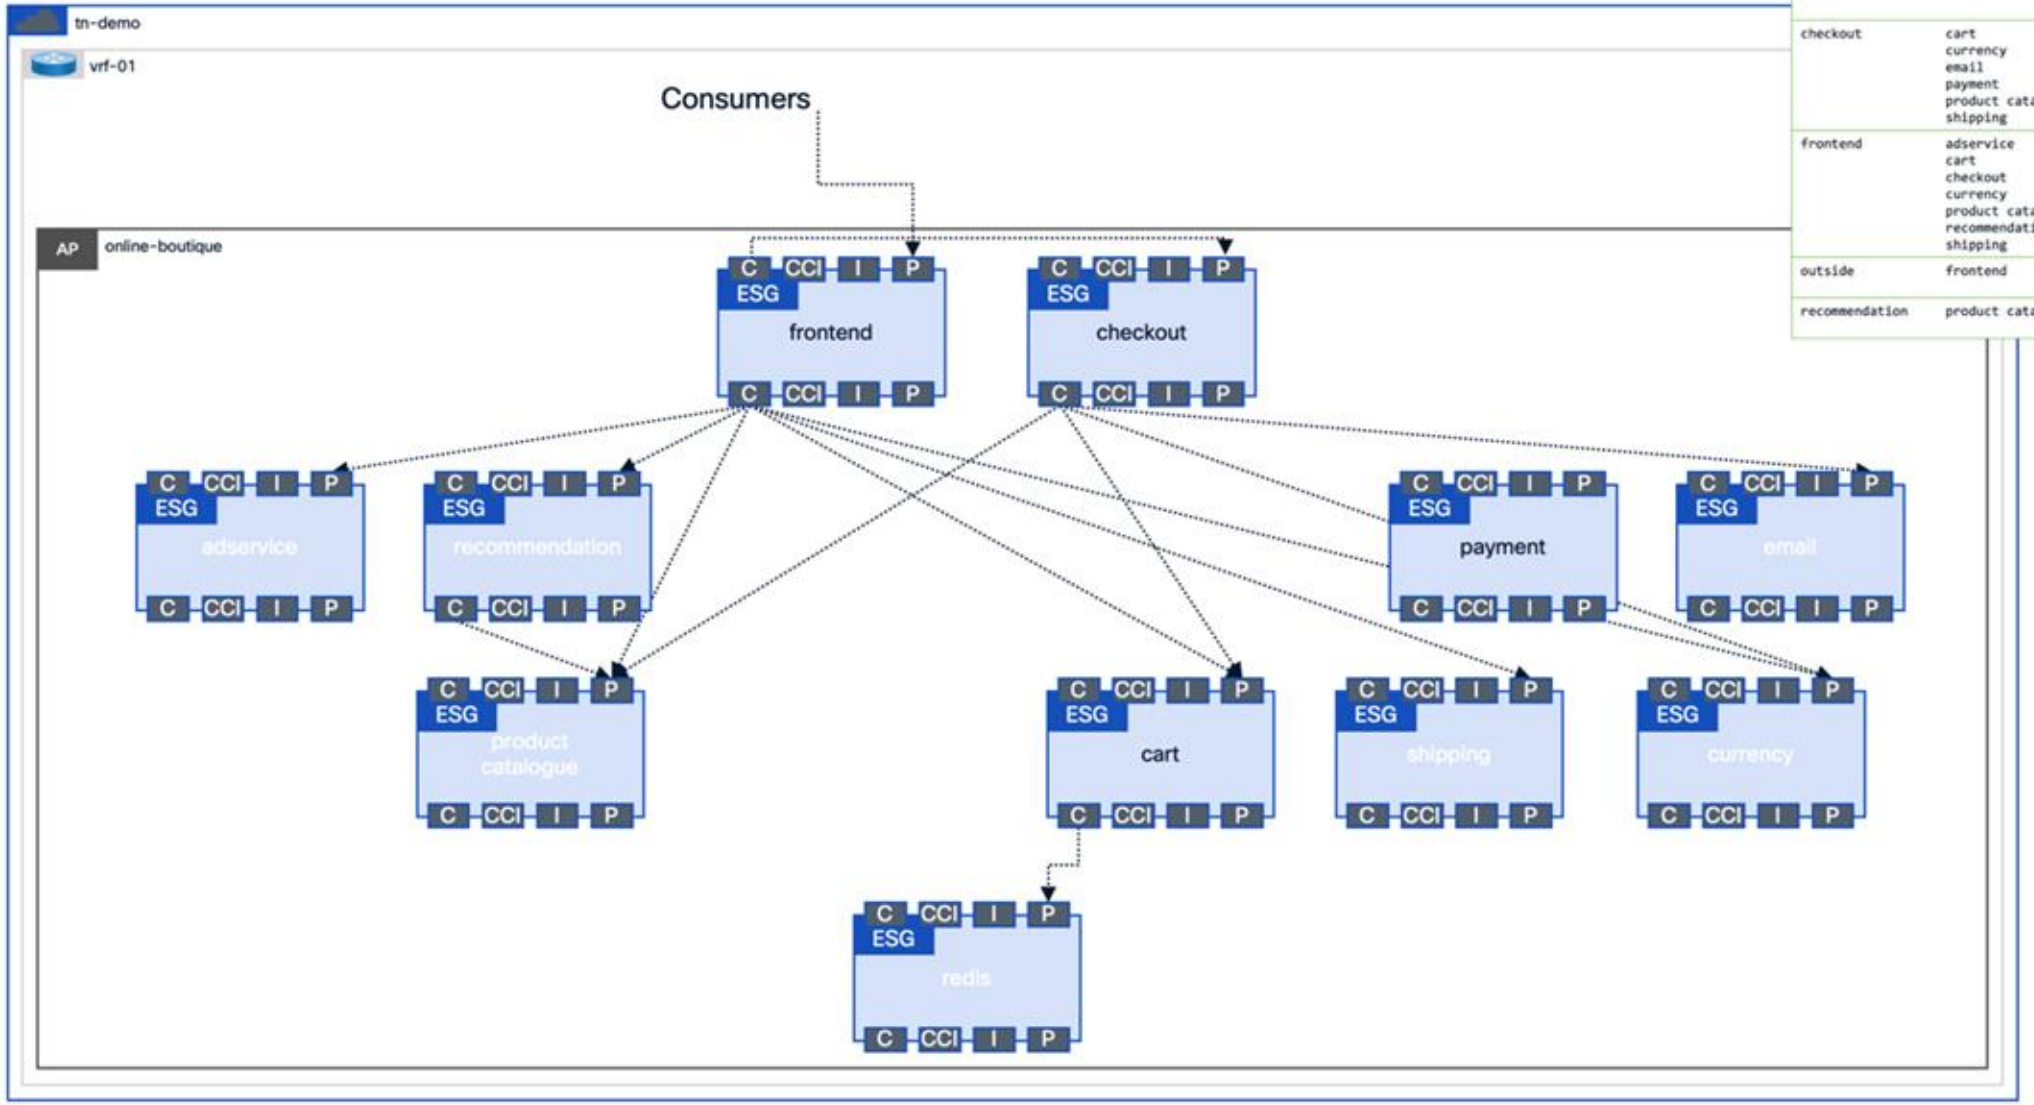

Policy Based Redirect to insert security services like NGFW

With PBR Service Graph

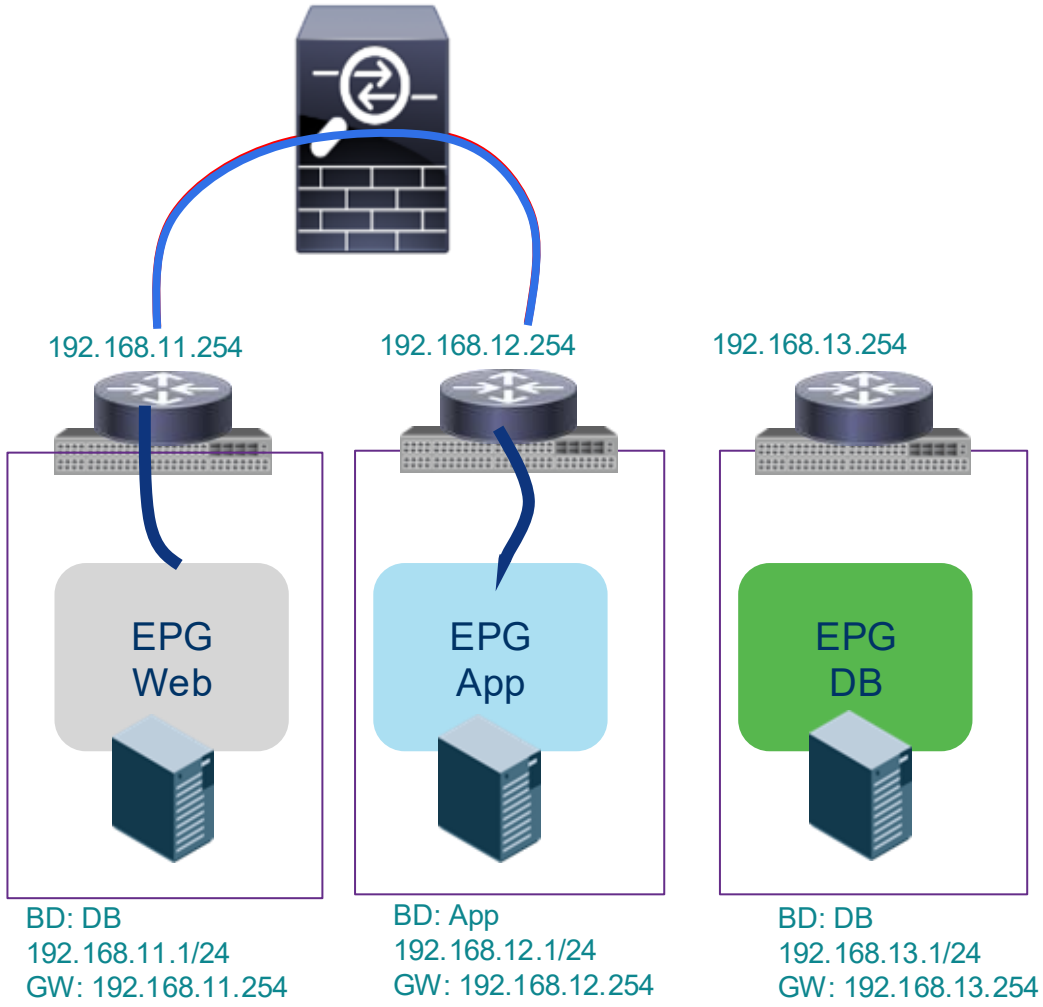

PBR redirects the traffic matching the contract(s) to the **Security Service**

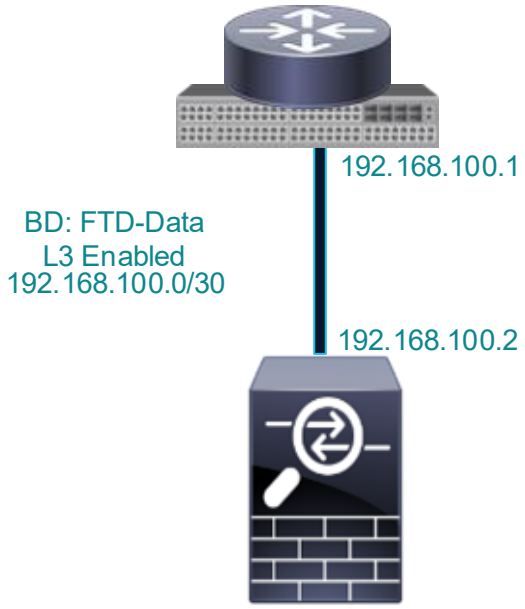

Complement ACI Contracts with Firewalls

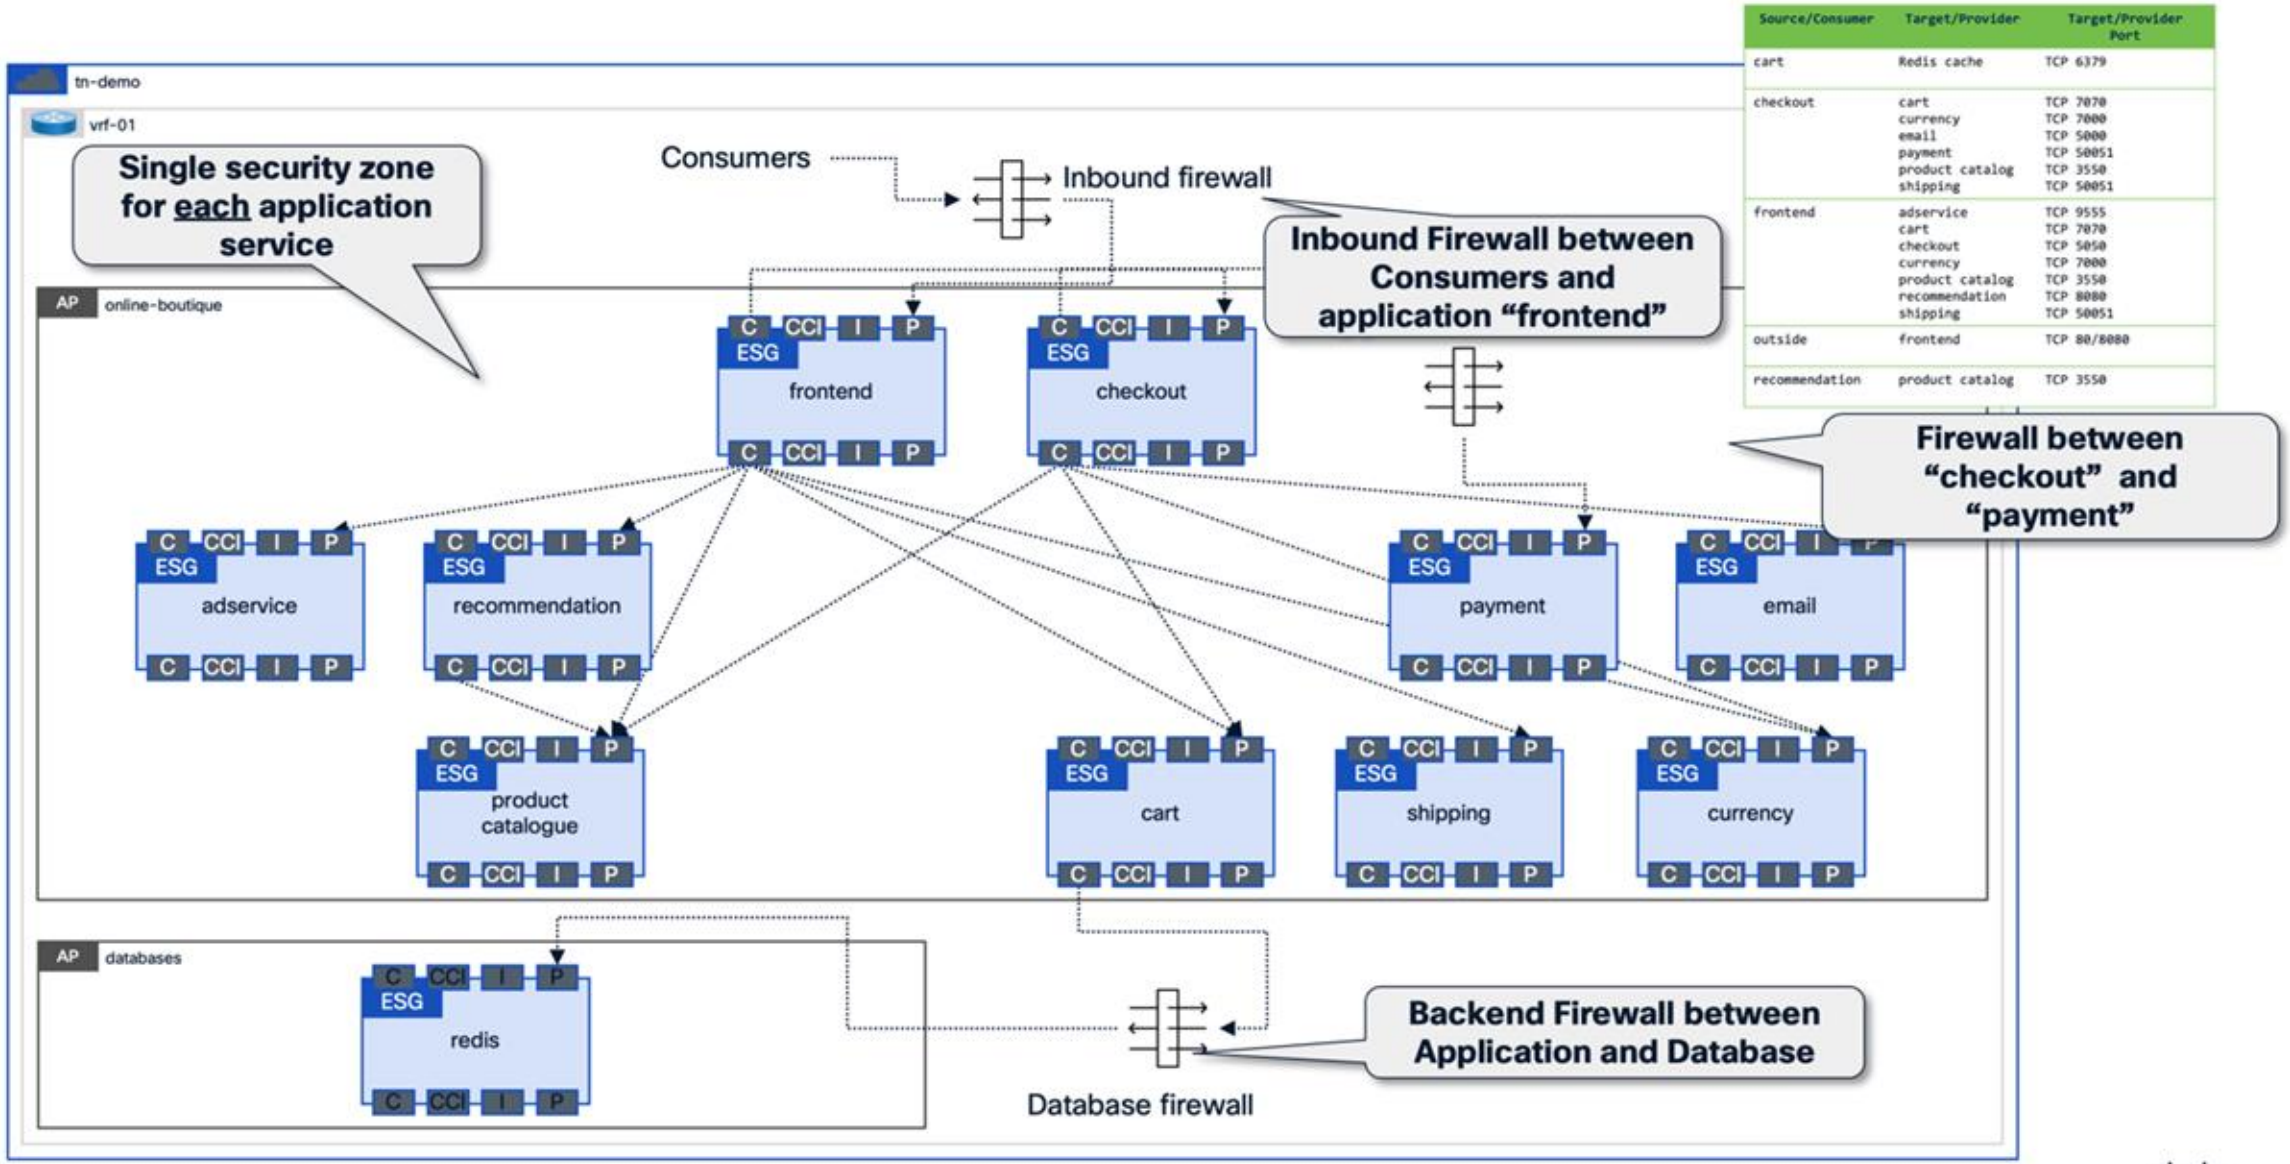

Source/Consumer	Target/Provider	Target/Provider Port
cart	Redis cache	TCP 6379
checkout	cart	TCP 7070
	currency	TCP 7000
	email	TCP 5000
	payment	TCP 50051
	product catalog	TCP 3550
	shipping	TCP 50051
frontend	adservice	TCP 9555
	cart	TCP 7070
	checkout	TCP 5050
	currency	TCP 7000
	product catalog	TCP 3550
	recommendation	TCP 8080
	shipping	TCP 50051
outside	frontend	TCP 80/8080
recommendation	product catalog	TCP 3550

Establish Trust - FTD/ASA can leverage SGTs

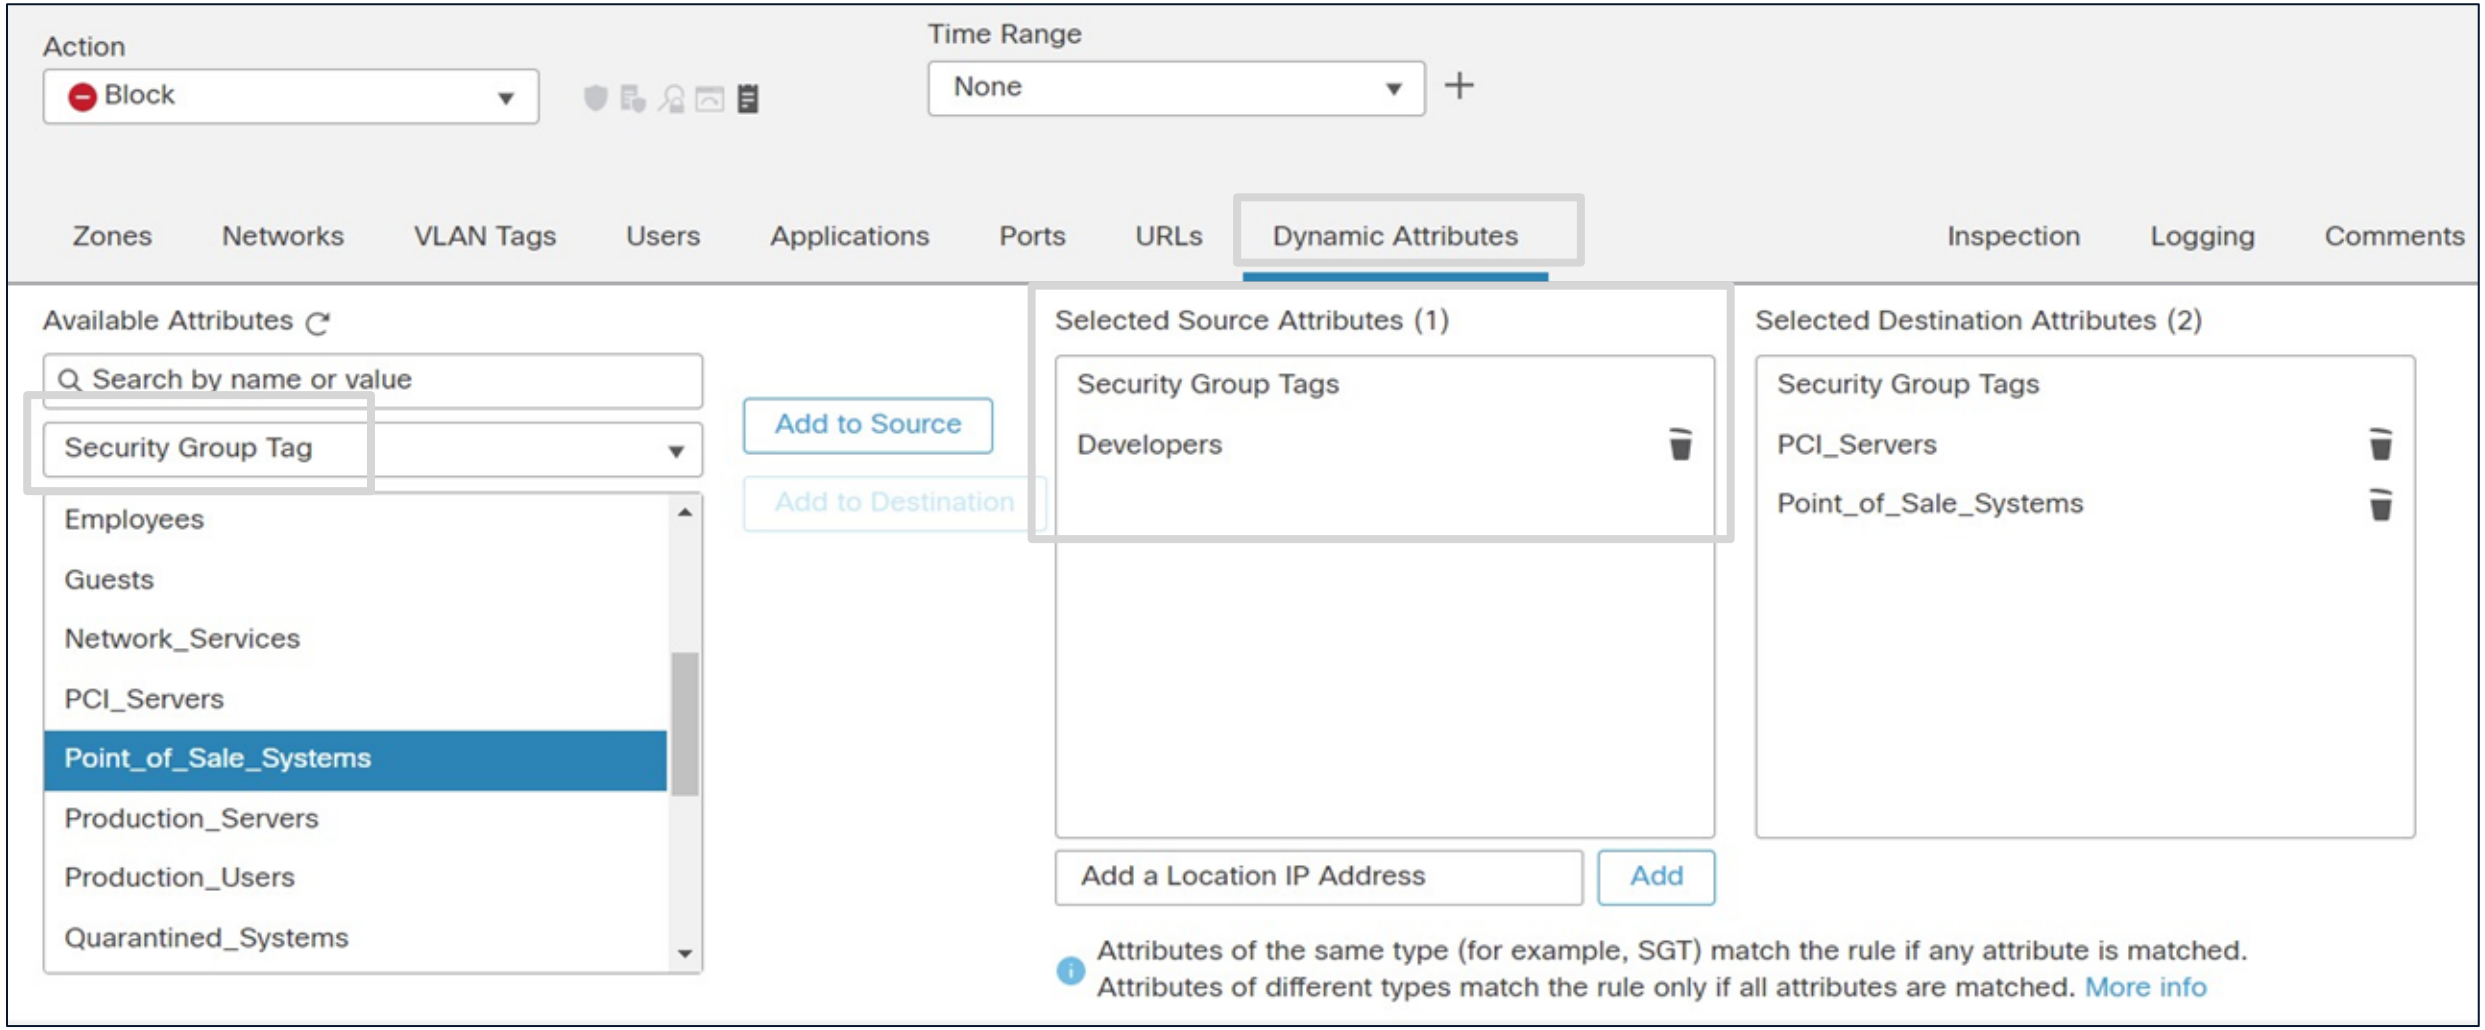

Establish Trust - FMC Learns EPGs/ESGs

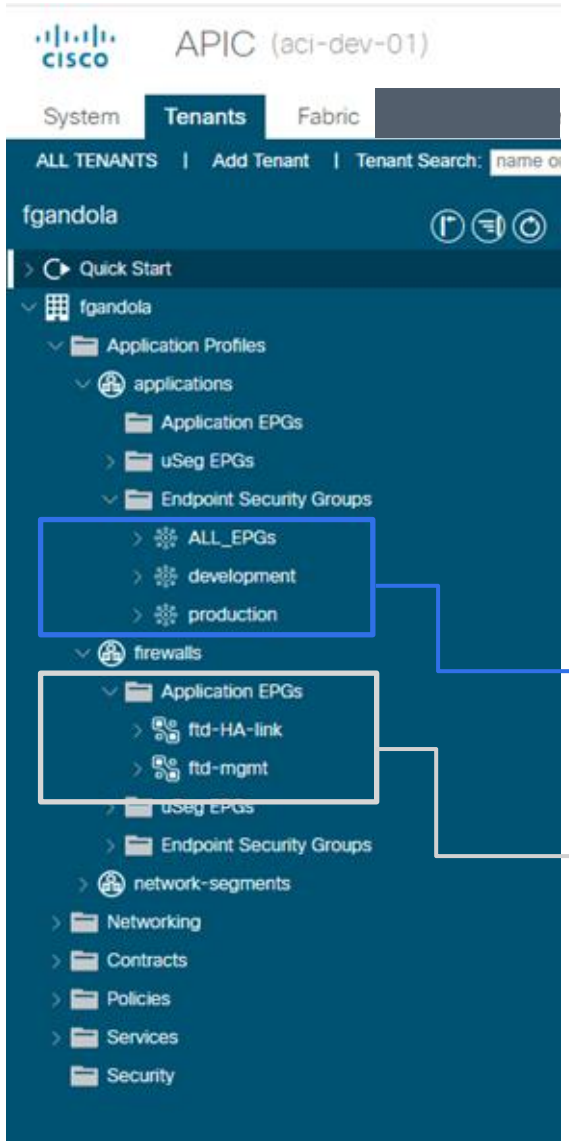

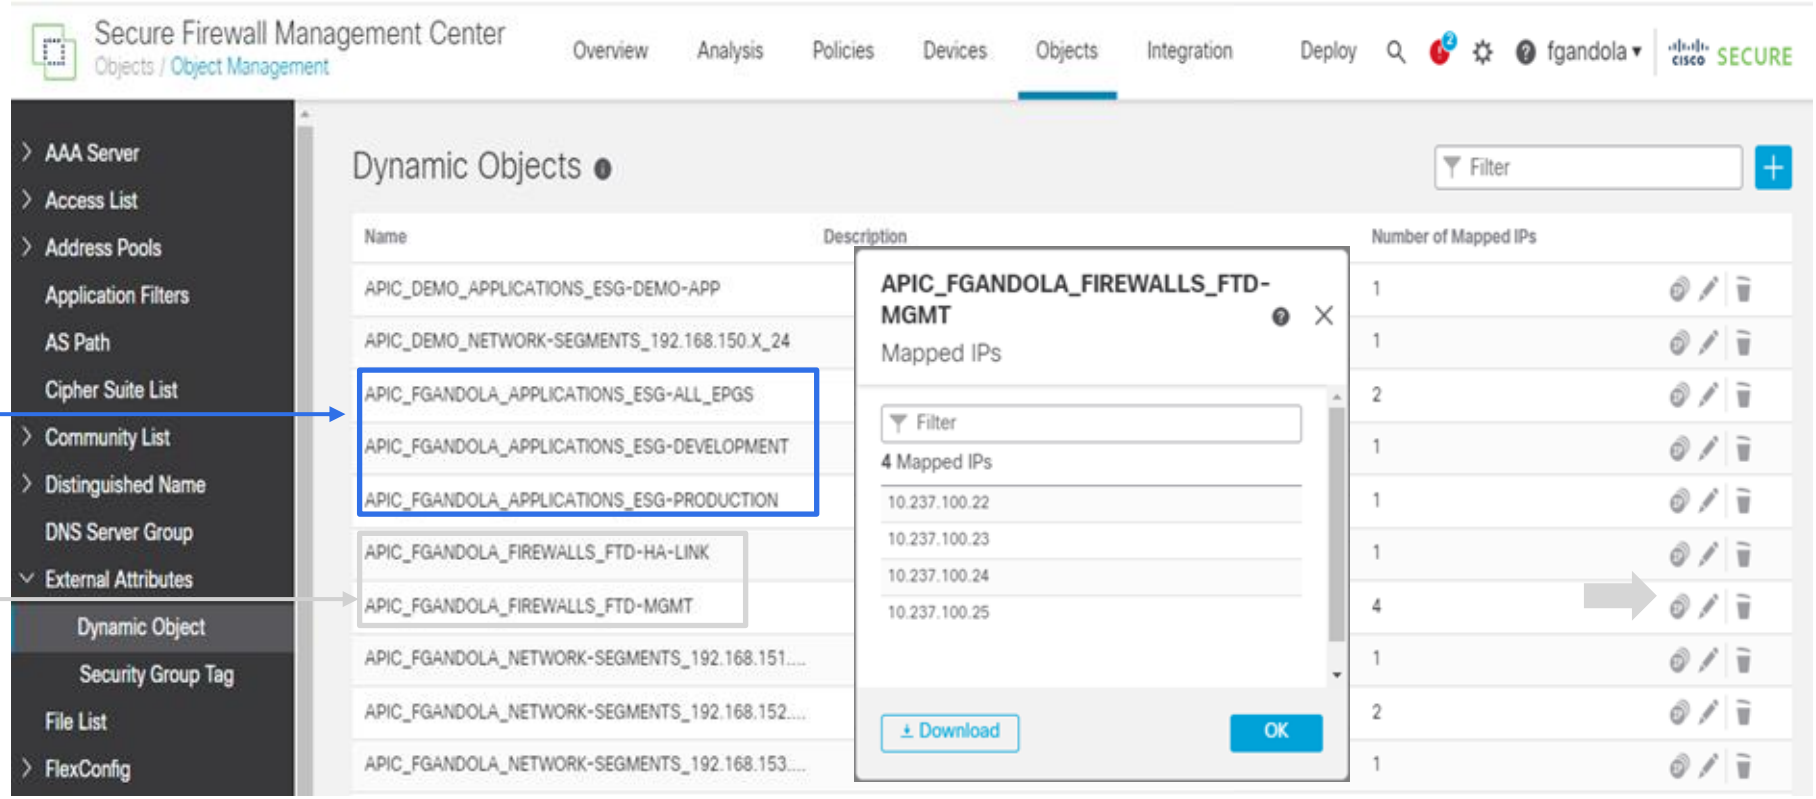

Use dynamic attributes to enforce trust

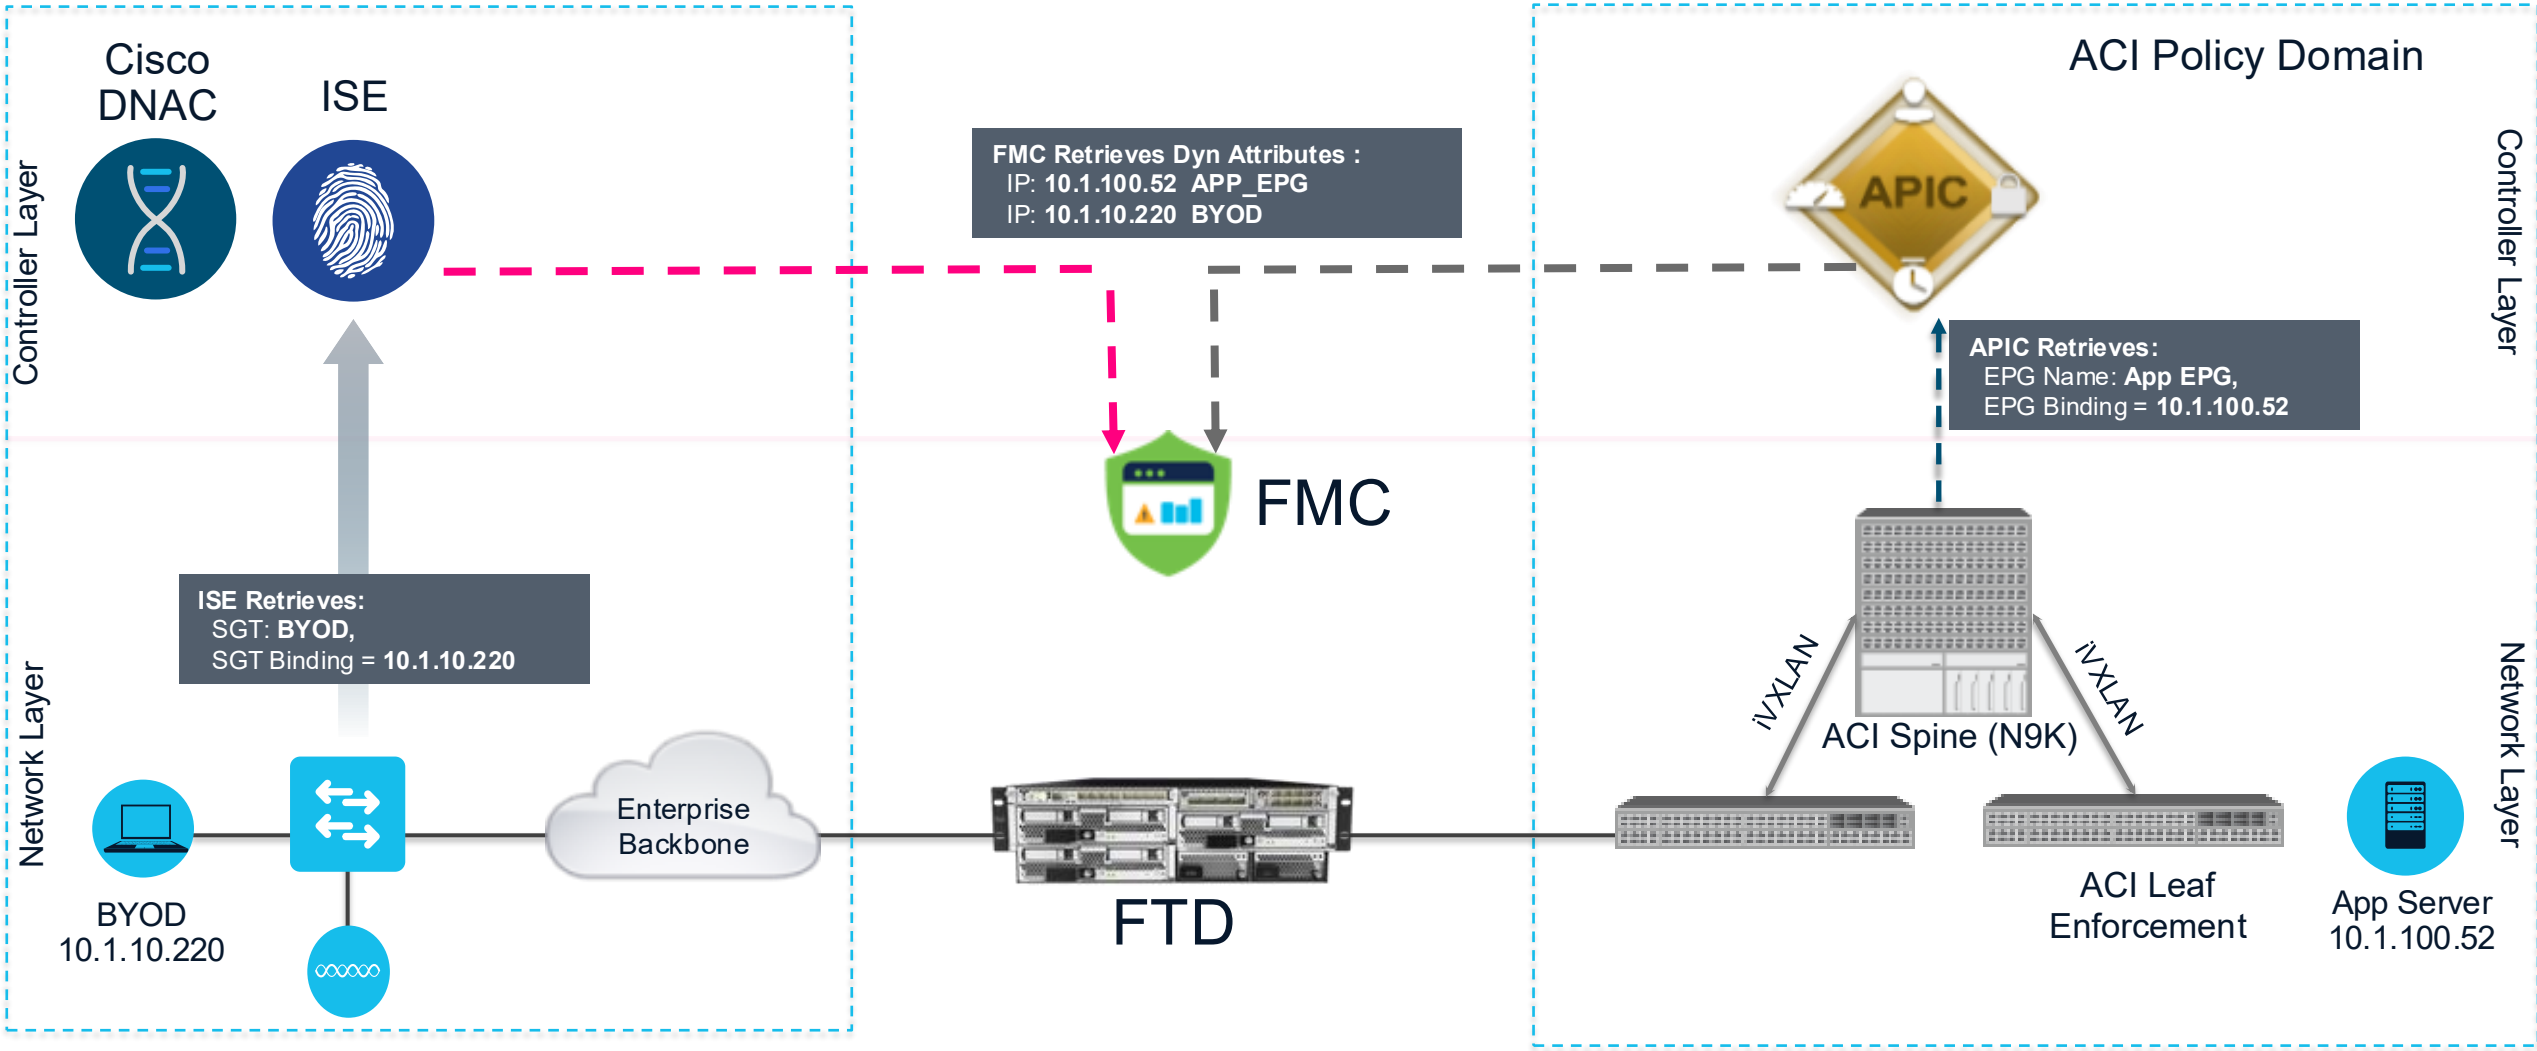

Dynamic Objects in Action

Automatic Without policy re-deployment

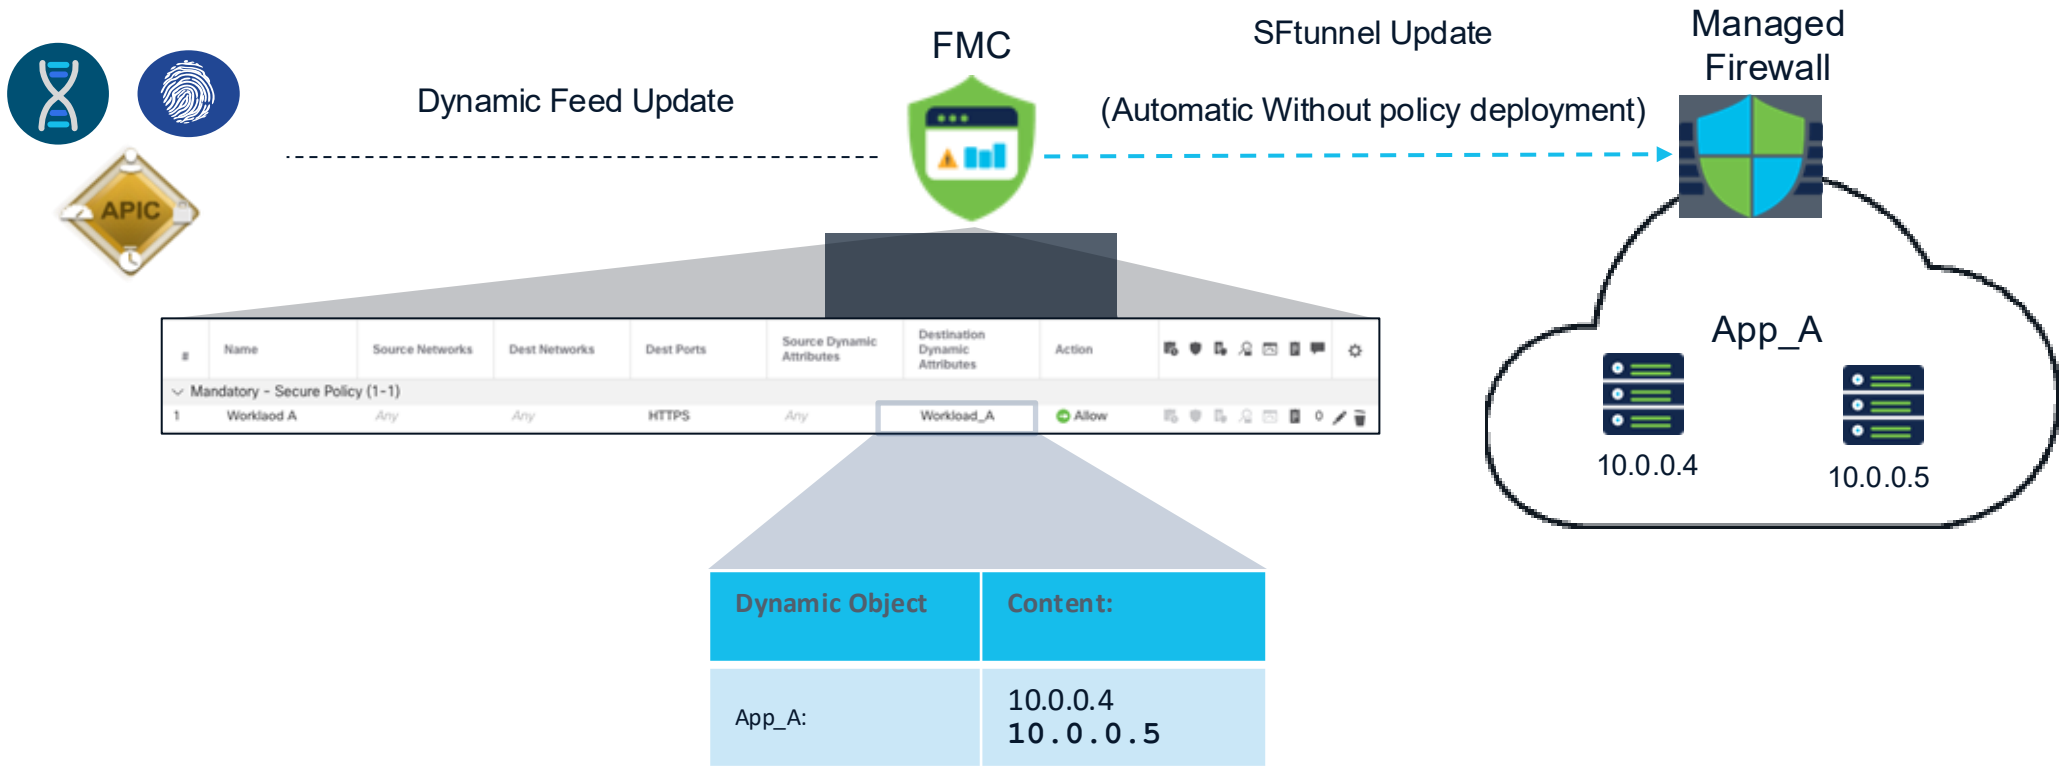

Dynamic Object	Content:
App_A:	10.0.0.4 10.0.0.5

Cisco Dynamic Attributes Connector (CDAC): Workflow & Example

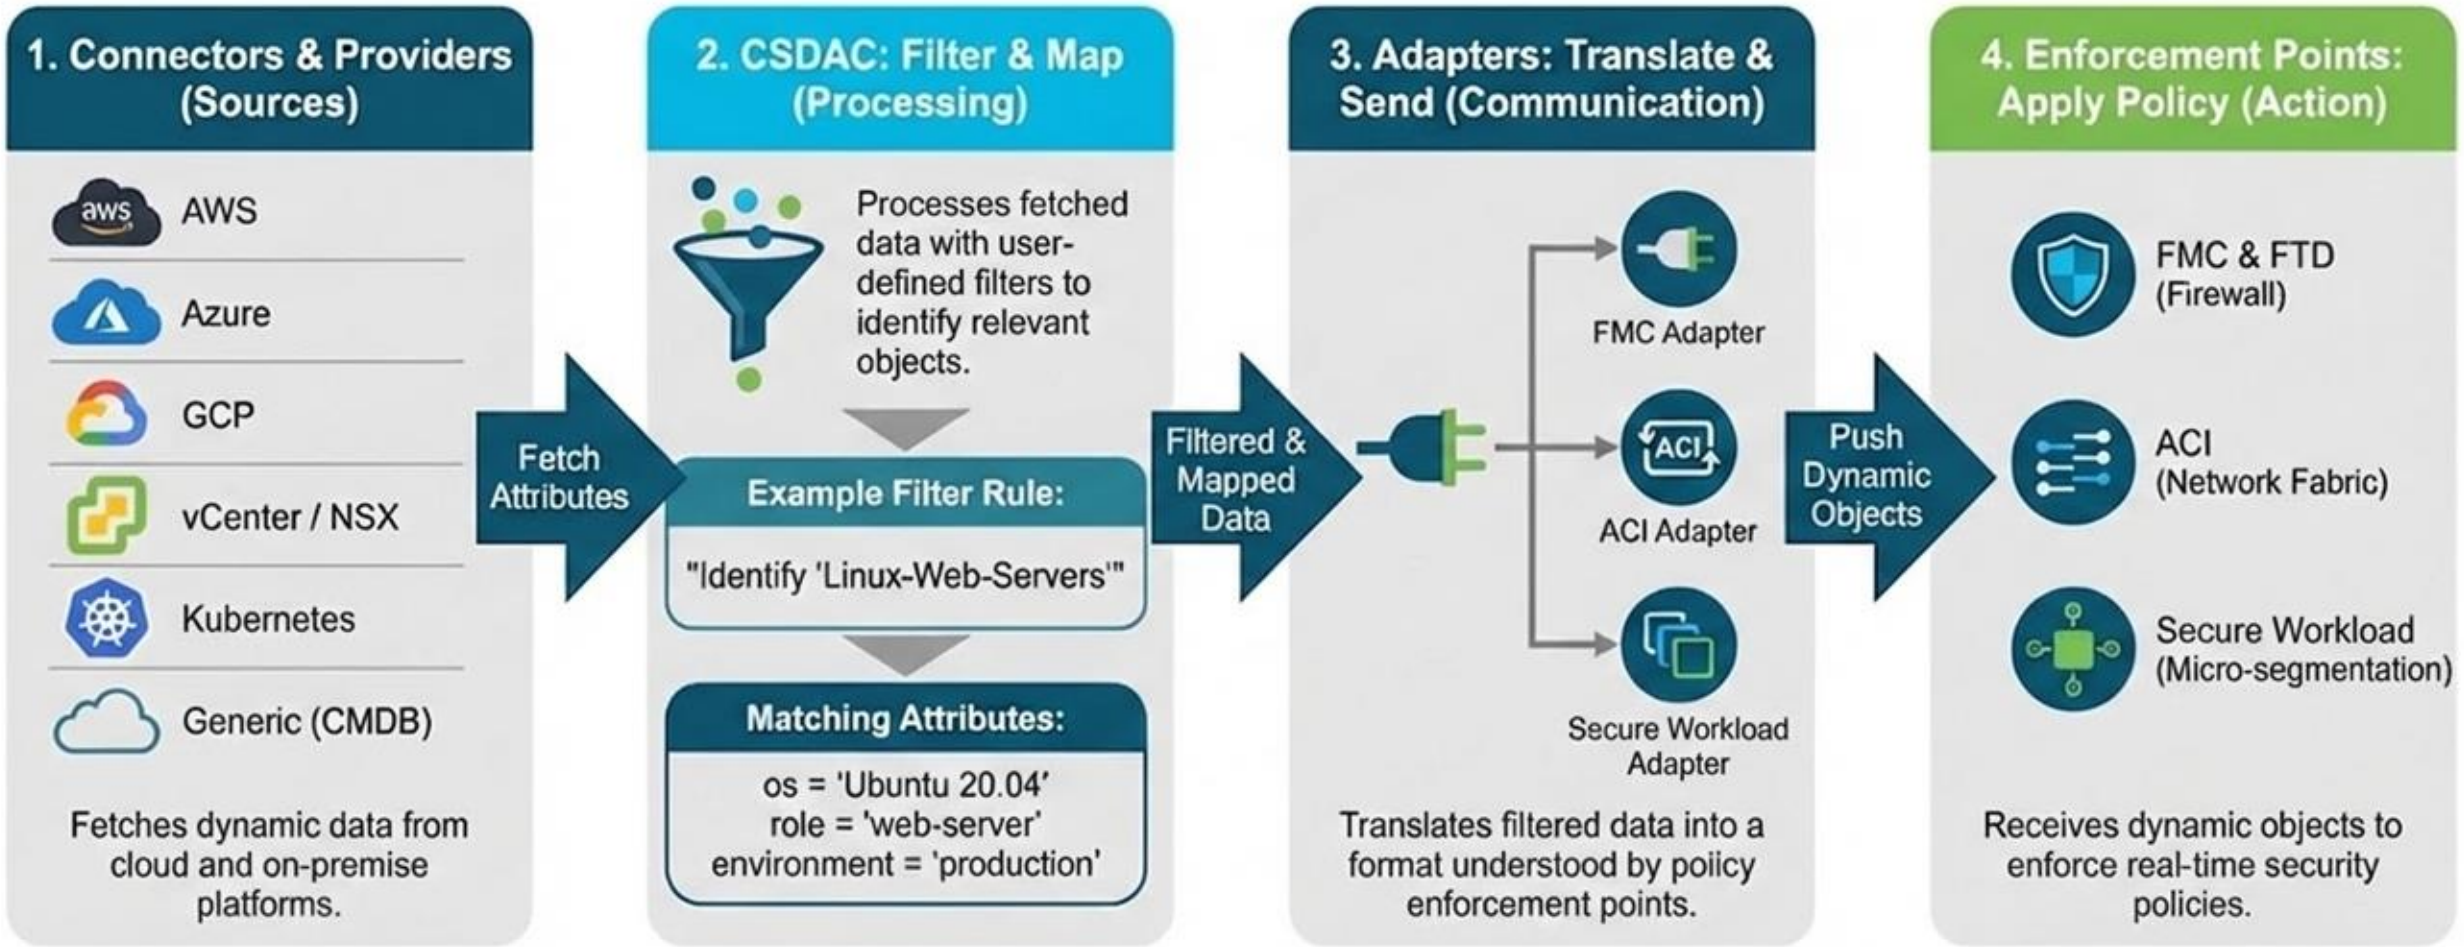

Application **workloads** evolution

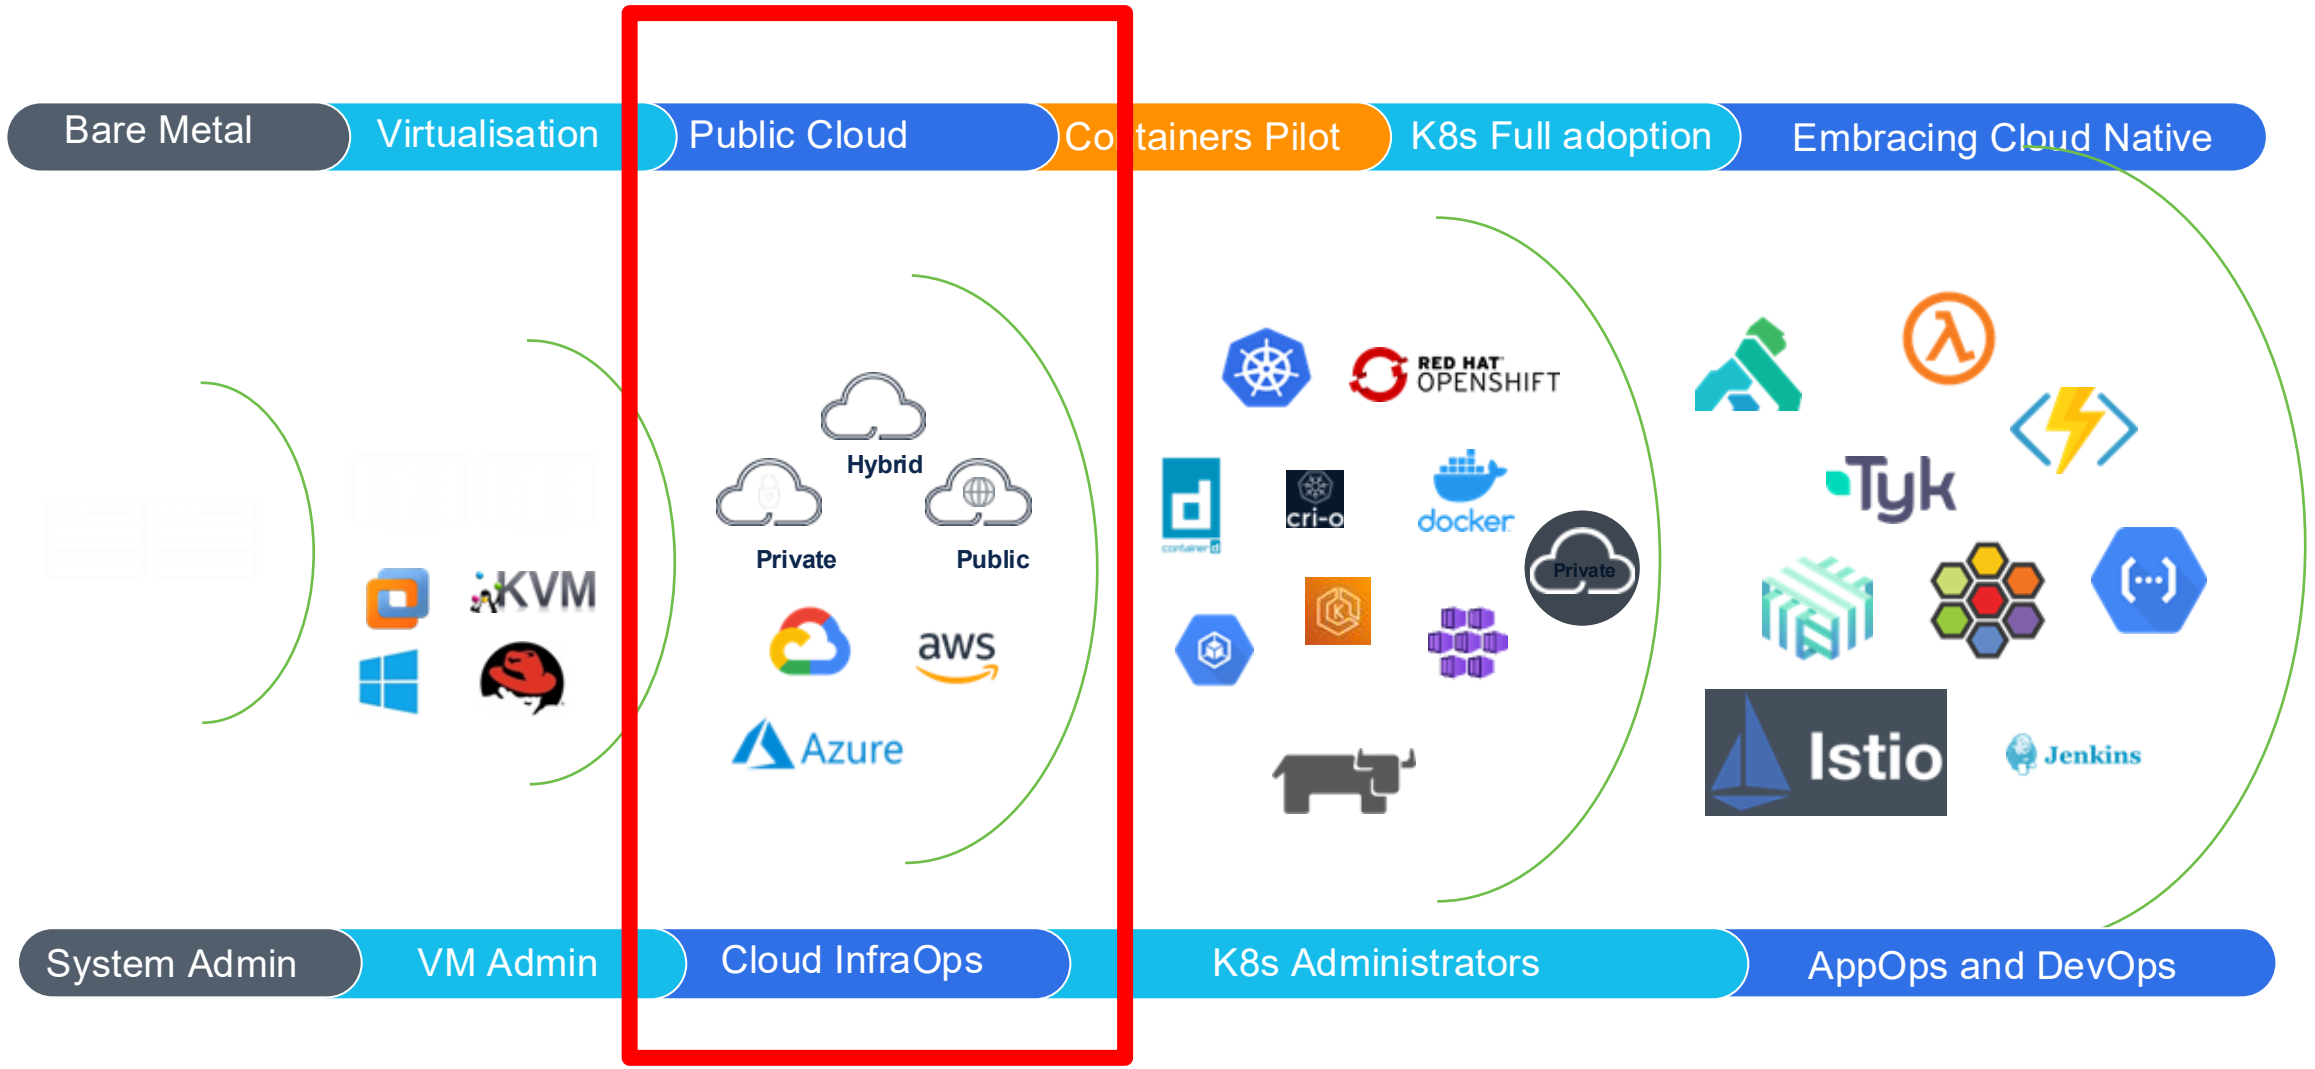

Macrosegmentation in the Cloud

Security Groups and Network Access list

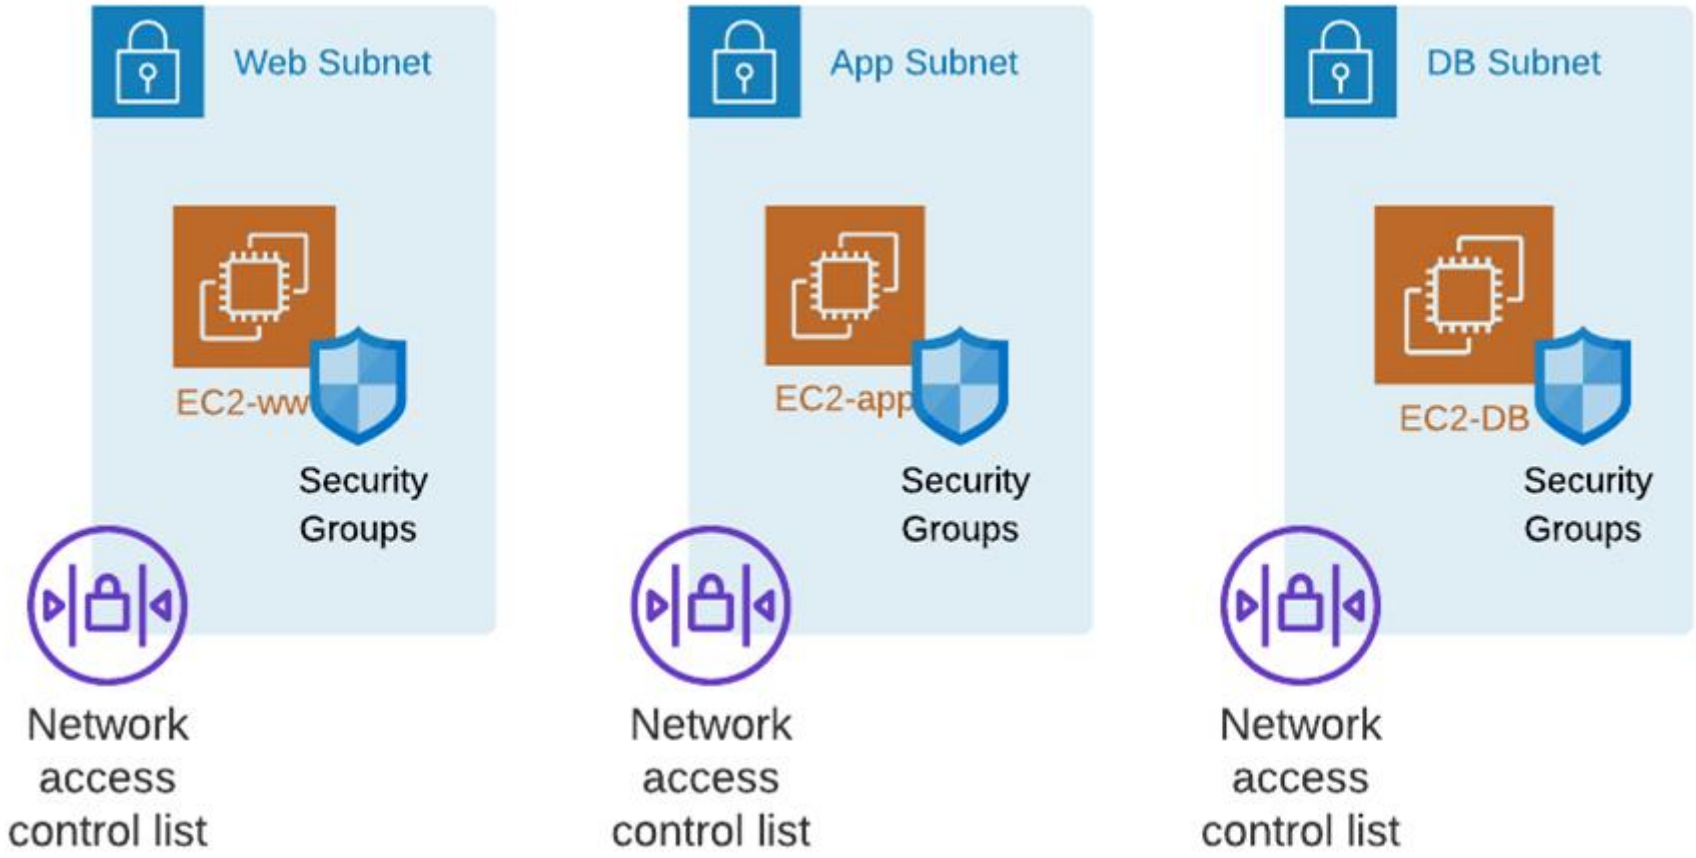

Firewall insertion with HA in front of the “Application” VPC

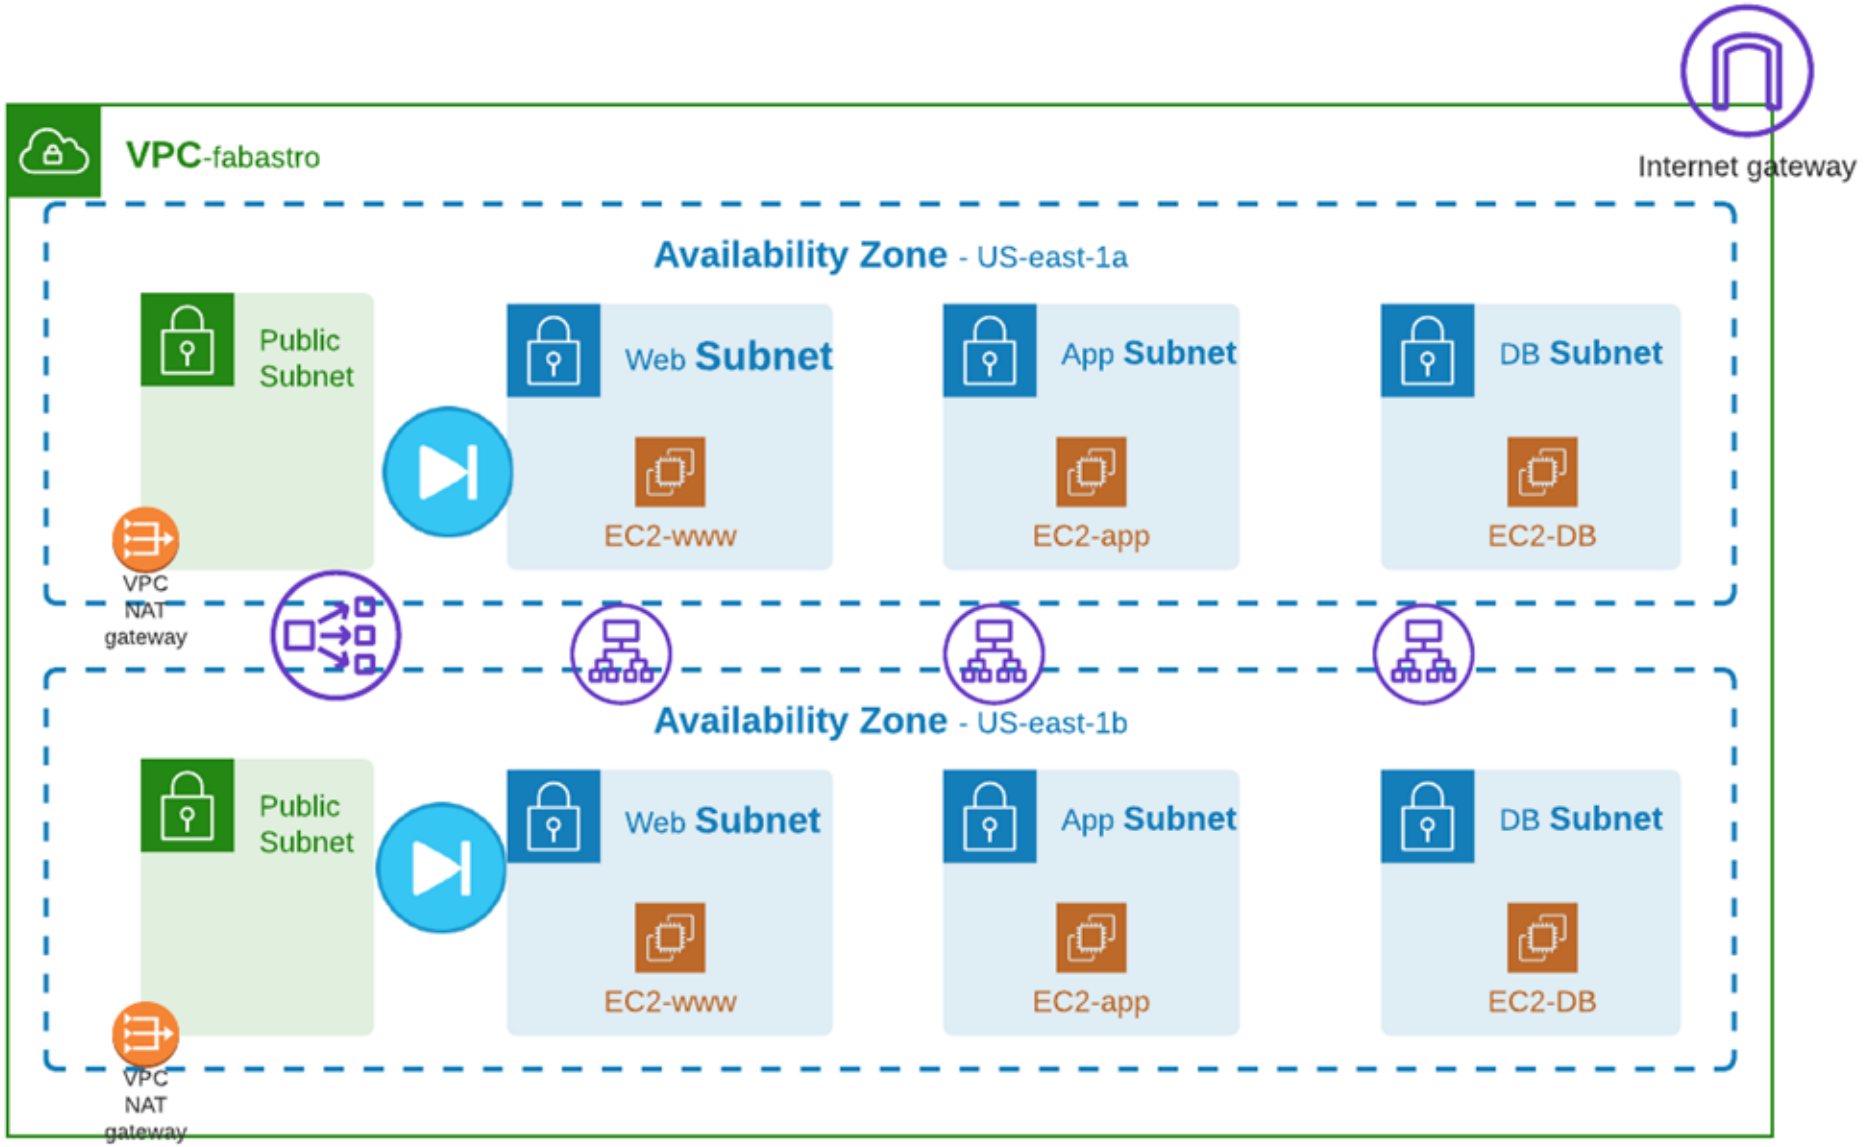

Multi-Cloud Security configuration complexity

What do we need to configure?

Security VPCs

VPC and DNS Flow Log capture

Network Load Balancer

Gateway Load Balancers

Firewalls Gateways
(Deployment, Insertion, Autoscaling)

AWS Transit Gateways
(New or existing, TGW attachment)

VPC subnet routing to TGW

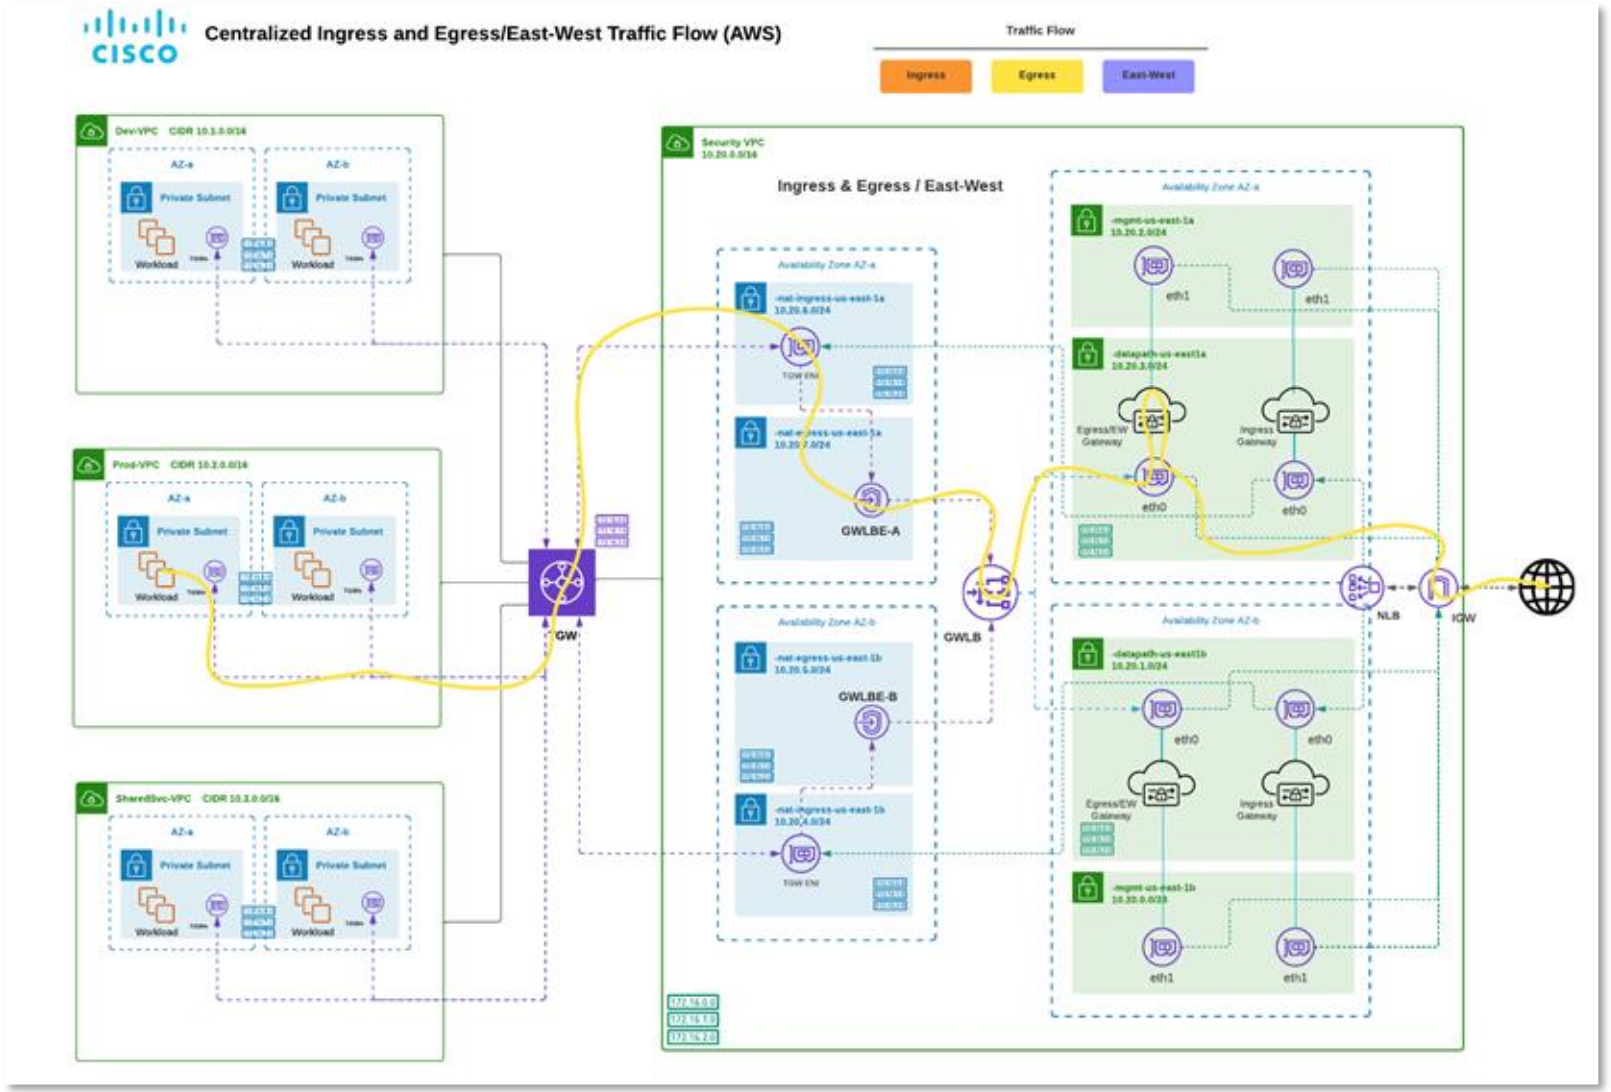

Cisco Multicloud Defense automates orchestration

What do we need to configure?

Security VPCs

VPC and DNS Flow Log capture

Gateway Load Balancers

Multicloud Defense Cloud

(Deployment, Insertion, Autoscaling)

AWS Transit Gateways

(New or existing, TGW attachment)

VPC subnet routing to TGW

AUTOMATED BY MULTICLOUD DEFENSE

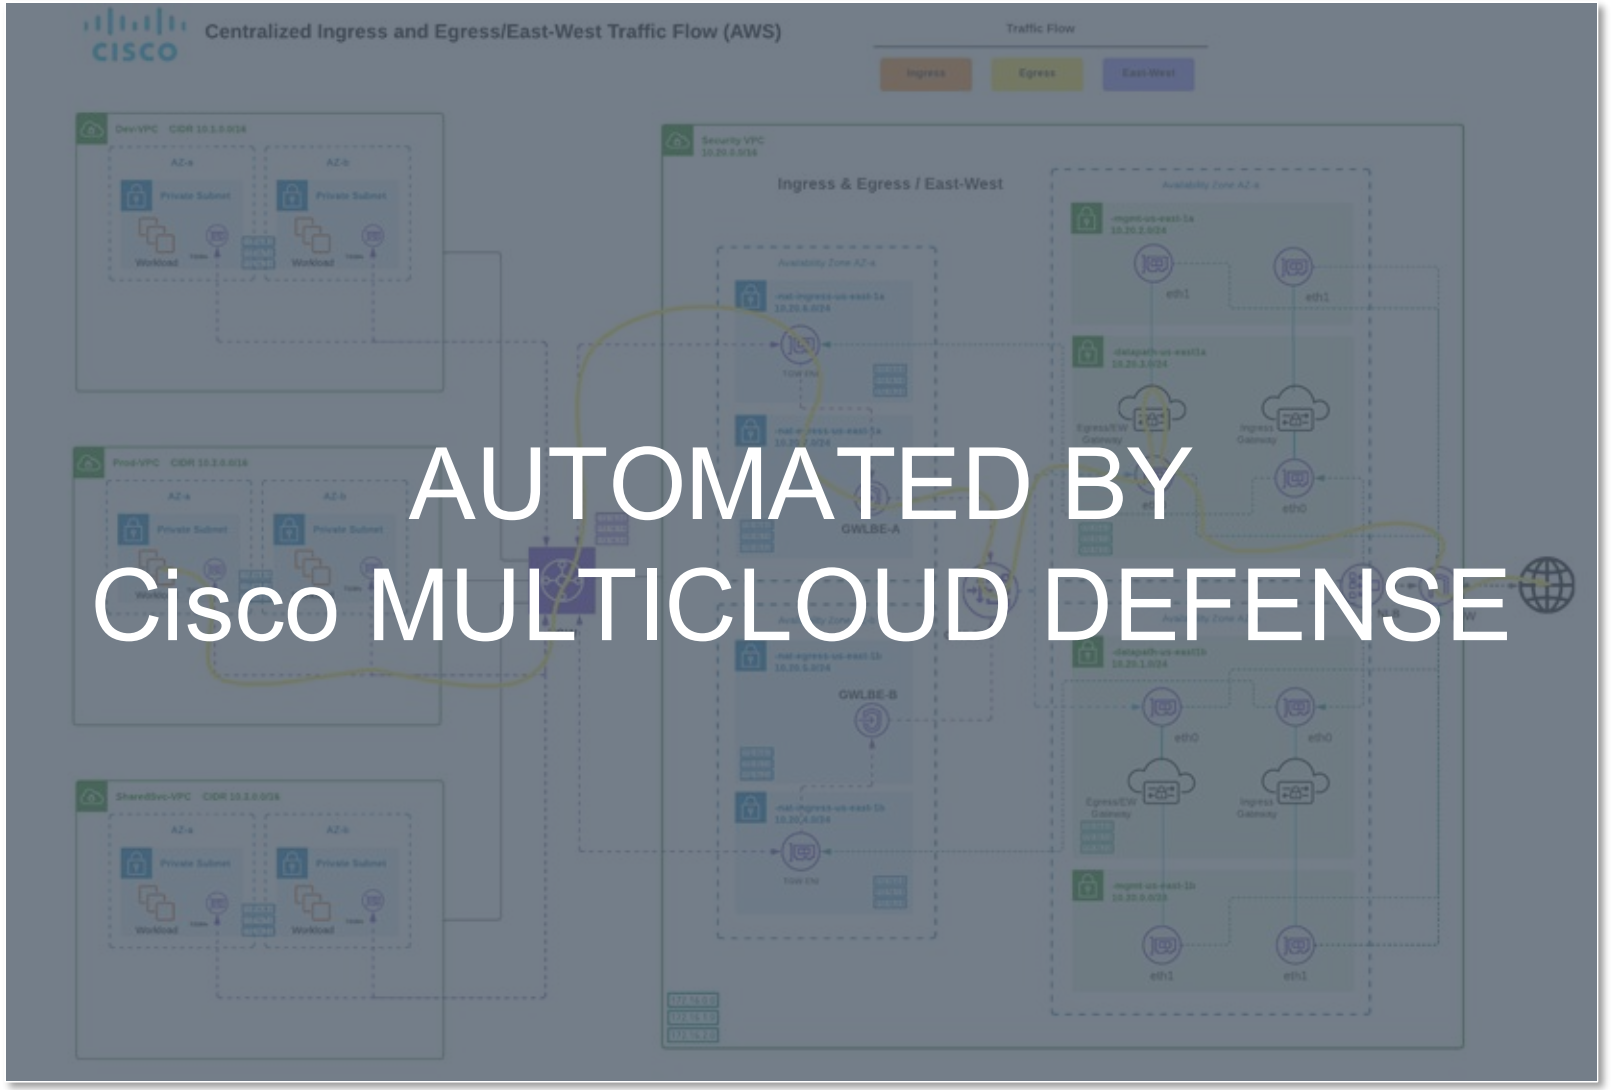

AUTOMATED BY Cisco MULTICLOUD DEFENSE

Cisco Multicloud Defense

Combining multicloud networking, automation, and cloud-native network security controls

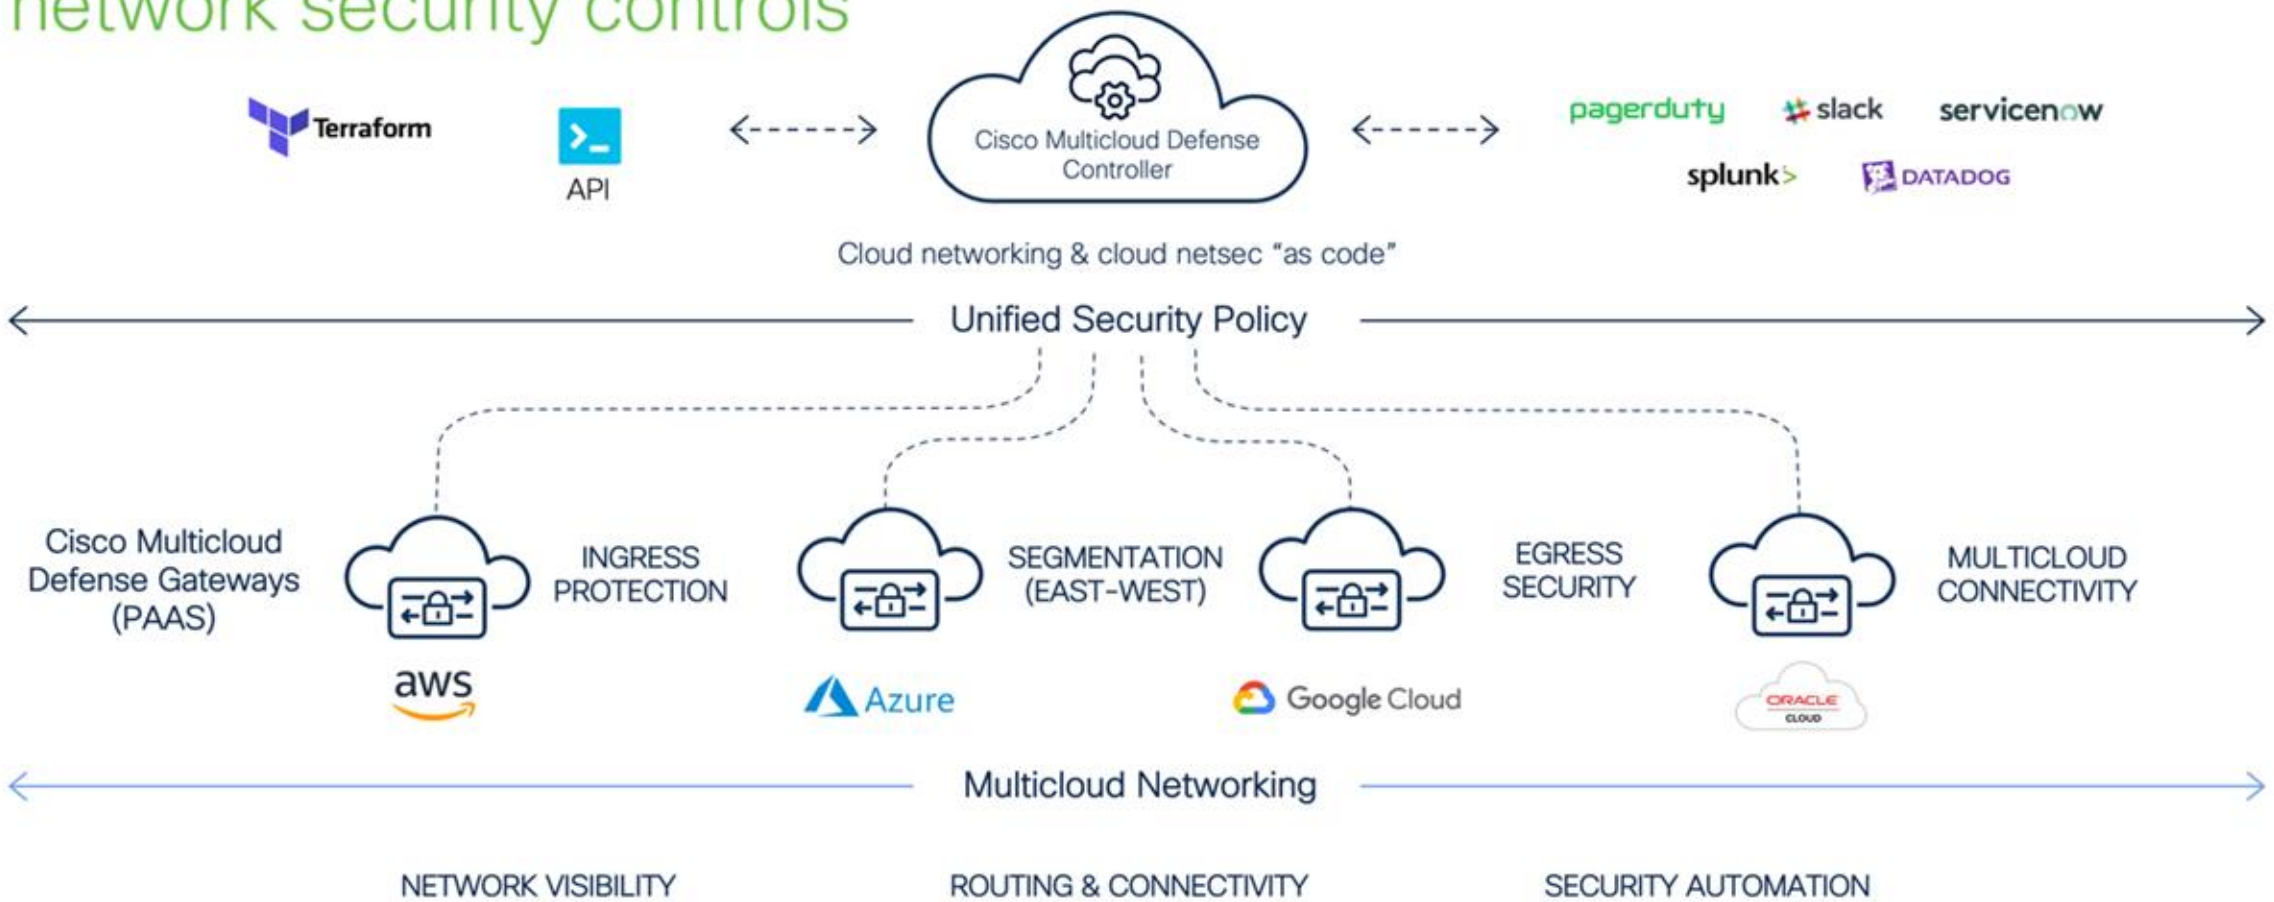

Use cases

FTDv Use Case

Ingress protection

Why Critical?

Stop inbound threats for web and non-web apps

Security Drivers

- Supply Chain Vulnerabilities
- Library Exploits (Log4Shell)
- Compliance

Security Controls

- WAF
- IDS / IPS
- Geo IP
- Malicious IP
- Antivirus

Egress security

Why Critical?

Block Command & Control, Botnets, Data Exfiltration

Security Drivers

- Zero Days
- C2, Ransomware
- Compliance

Security Controls

- FQDN Filtering
- URL Filtering
- DLP
- IPS / IDS
- Antivirus

Segmentation

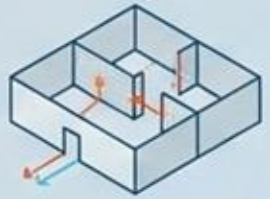

Why Critical?

Mitigate Lateral Movement

Security Drivers

- Zero Trust / Least Privilege
- Data Protection
- Compliance

Security Controls

- Cloud-native Identity
- IPS / IDS
- Antivirus

cisco

Egress Security

Deployment Models

Centralized Security Model

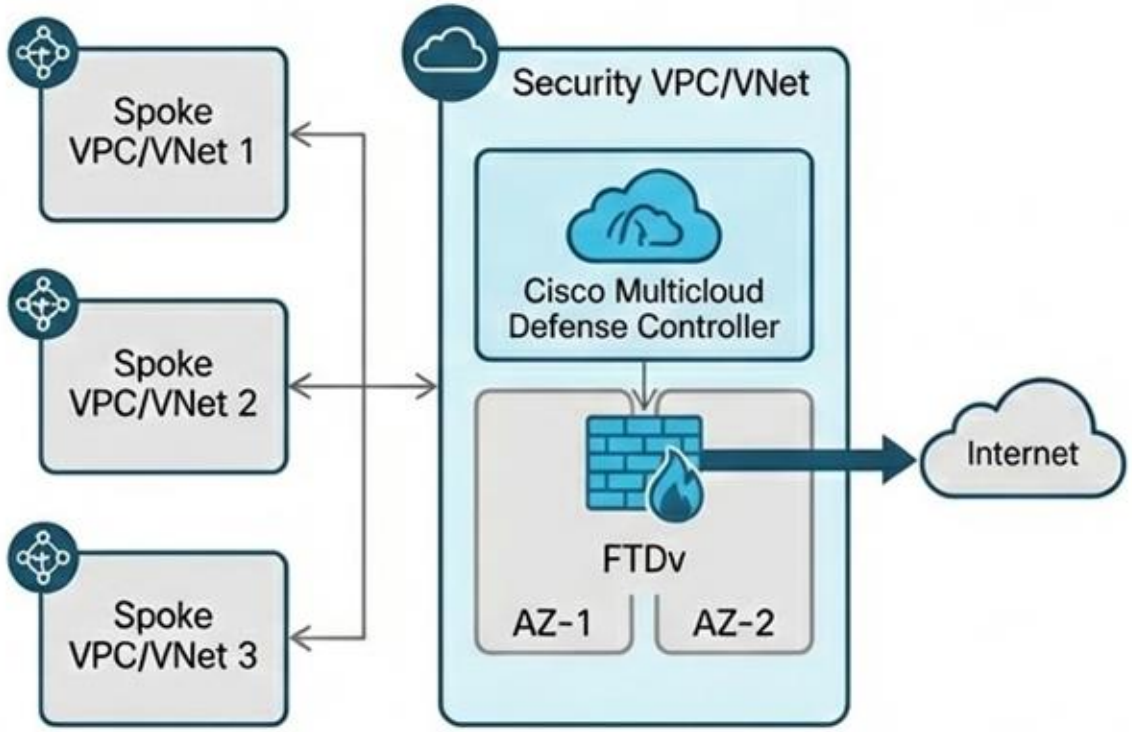

Distributed Security Model

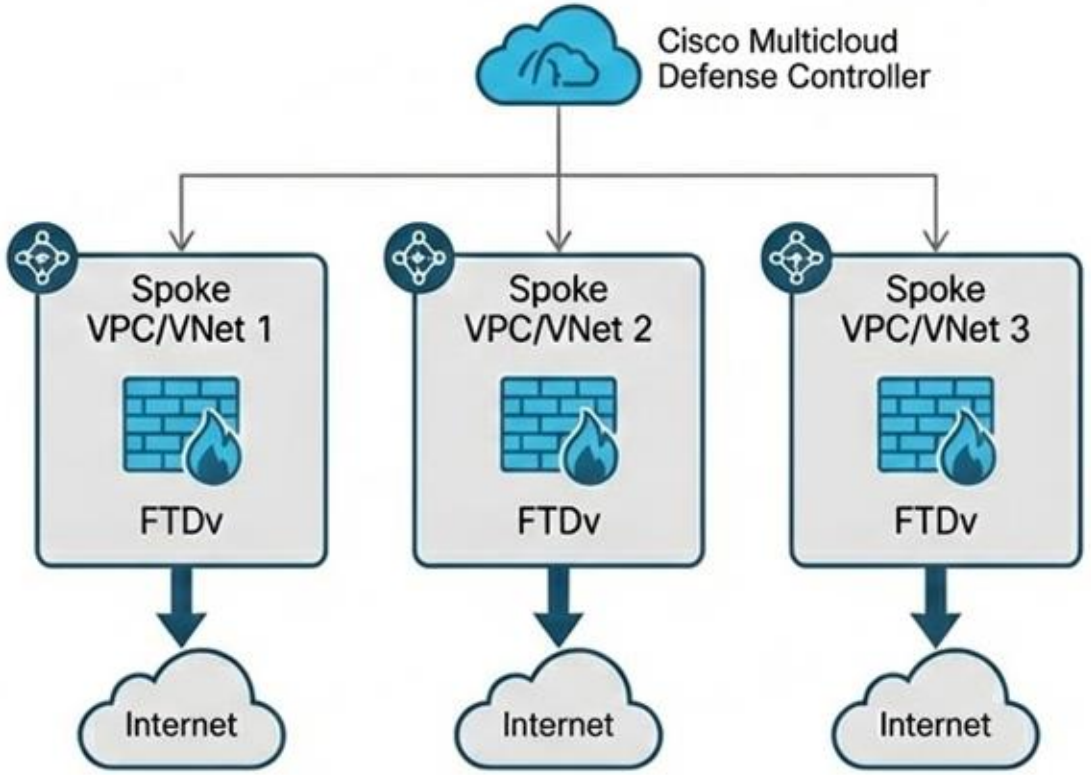

cisco

Simple Setup Process

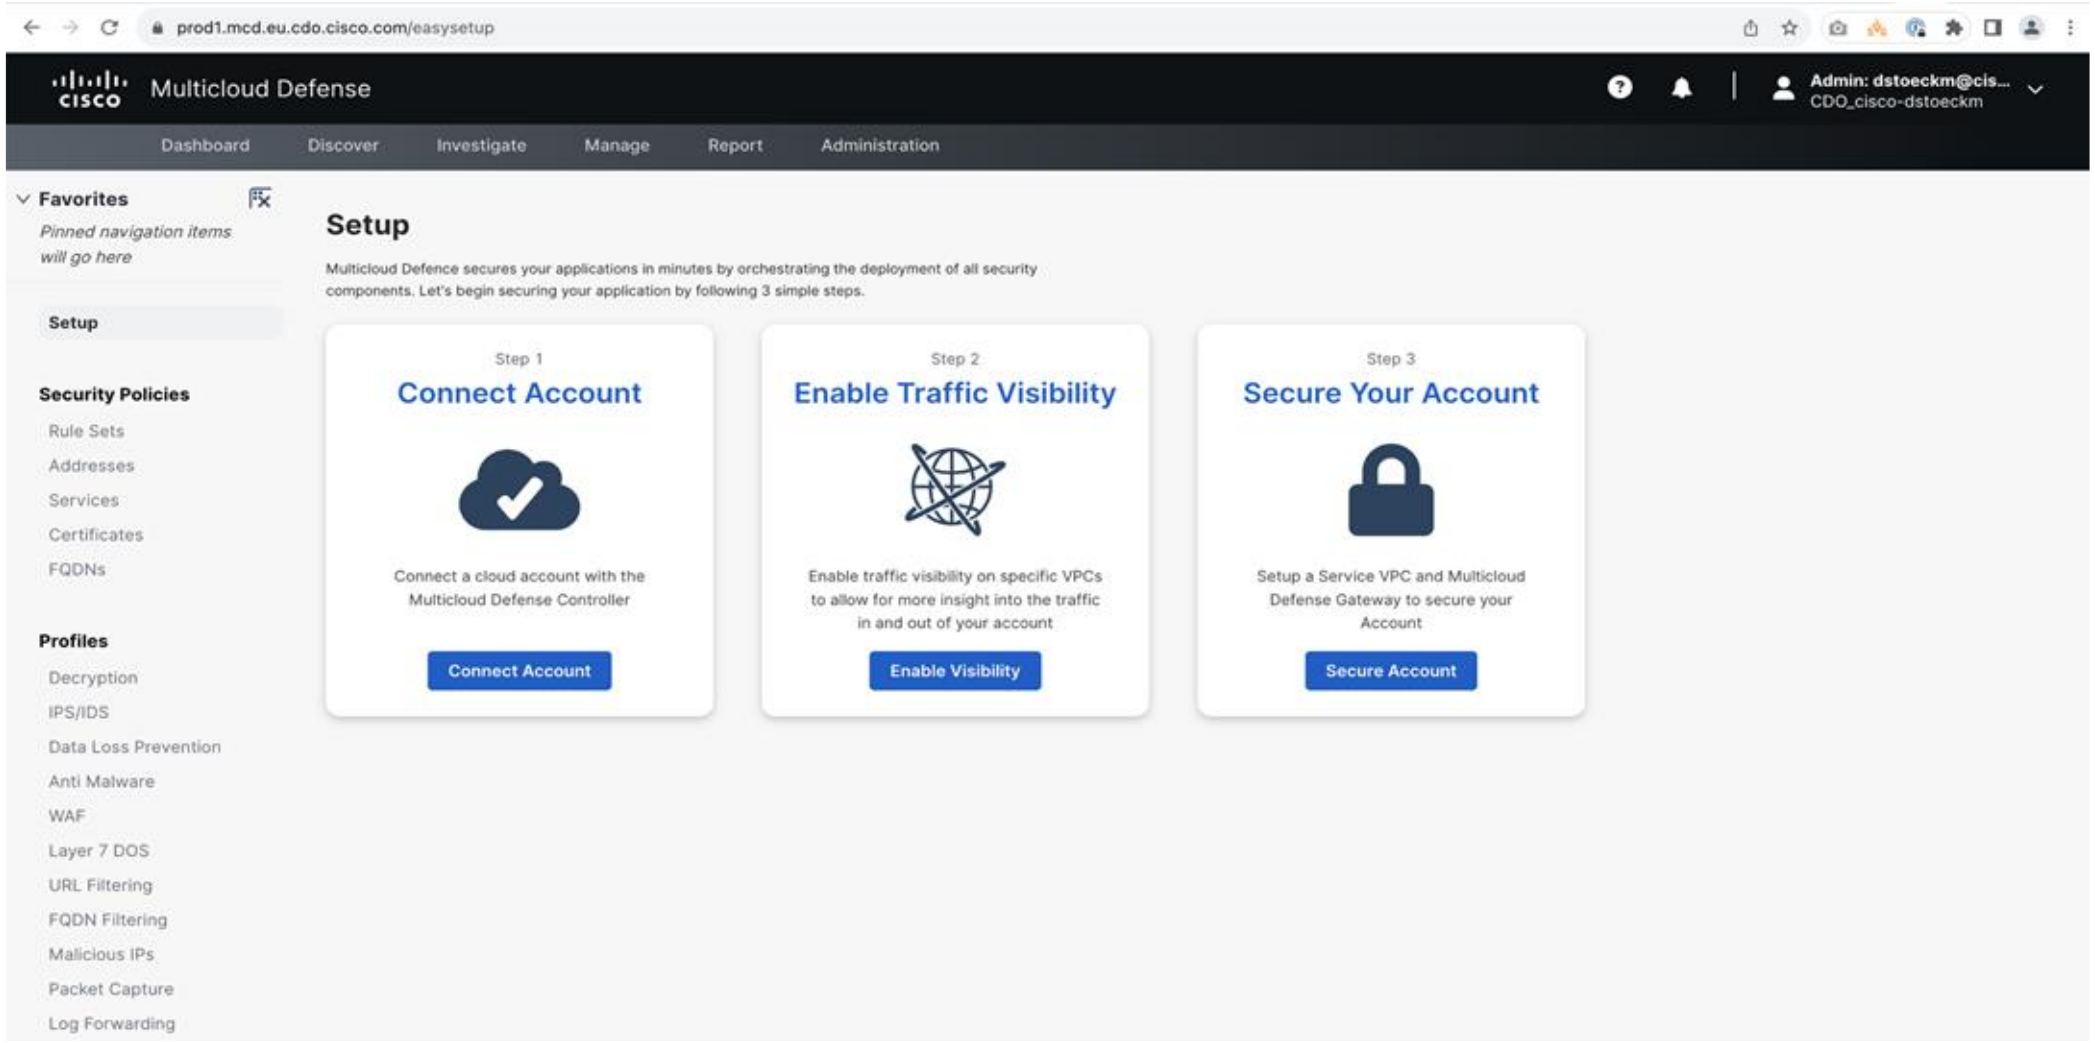

MCD Setup Demo

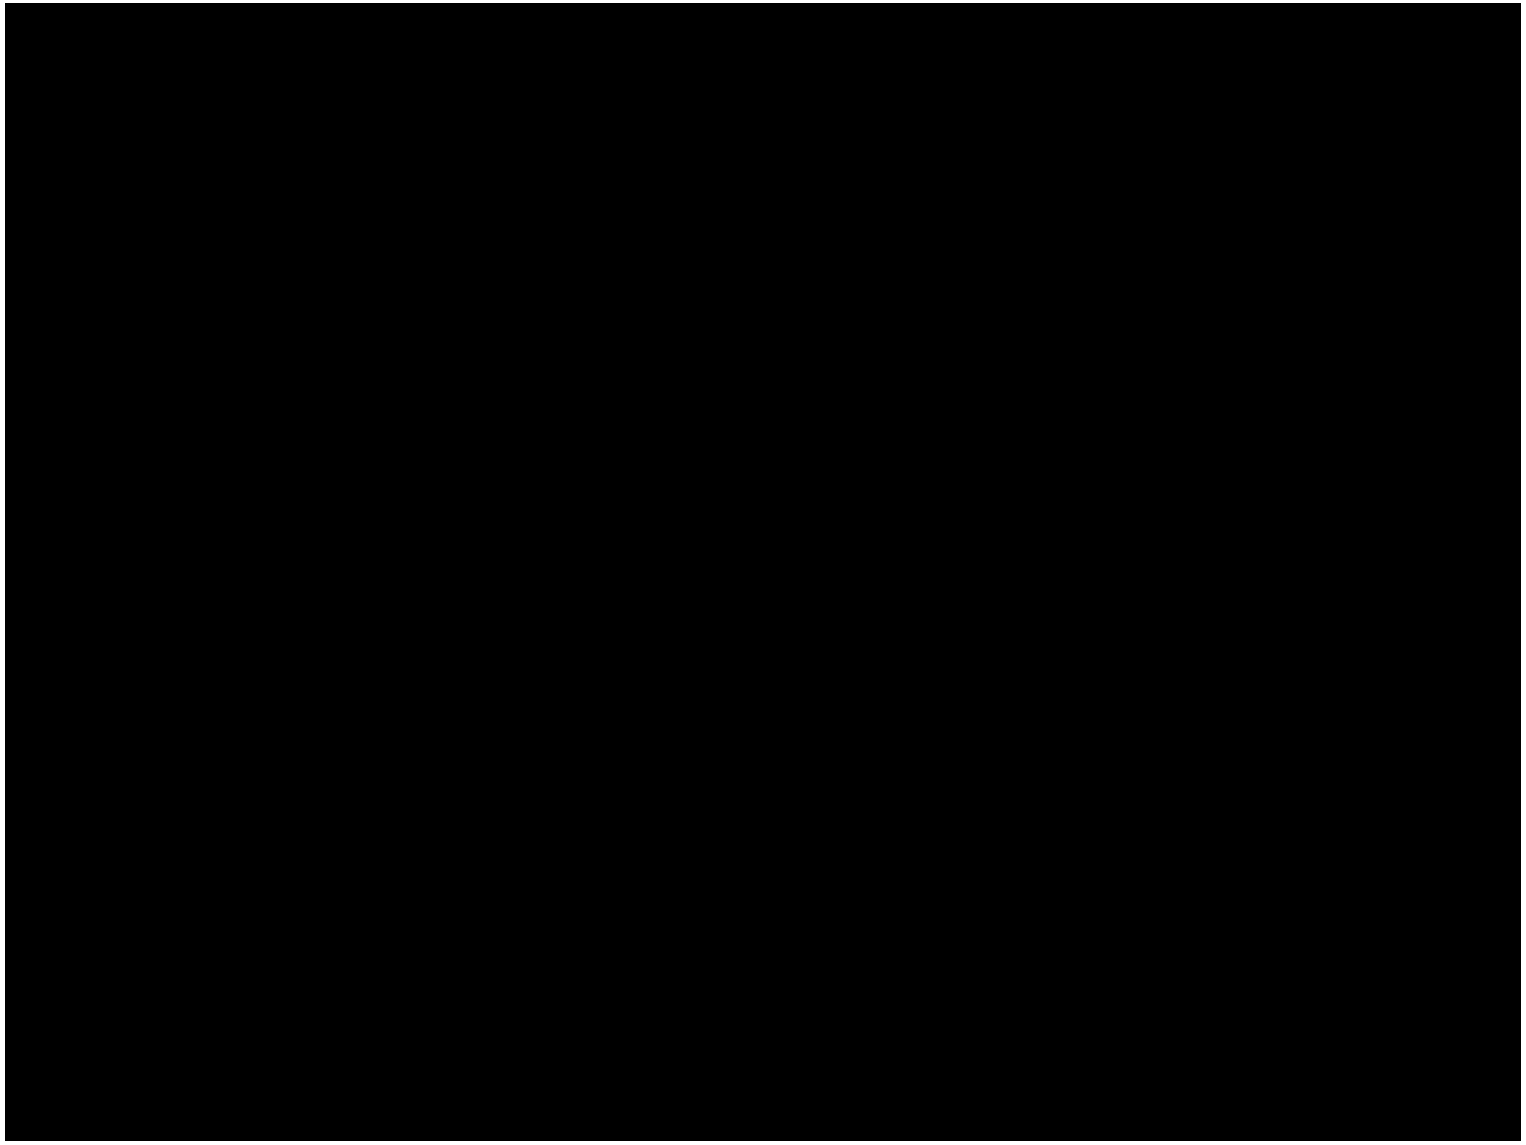

Cisco Secure Firewall Threat Defense Virtual

Private Cloud

HyperFlex
NUTANIX
KVM
openstack.
vmware ESXi

Public Cloud

aws
Google Cloud Platform
Microsoft Azure
rackspace technology
ORACLE CLOUD INFRASTRUCTURE
EQUINIX
Alibaba Cloud
alkira

Gov/IC Cloud

aws
Microsoft Azure
Google Cloud Platform

Virtual firewall performance-based licensing from 100Mbps up to 16 Gbps

FTDv5 (100Mbps), FTDv10 (1 Gbps), FTDv20 (3 Gbps), FTDv30 (5 Gbps), FTDv50 (10 Gbps), and FTDv100 (16 Gbps)

Cloud Leadership

- Clustering & Auto Scaling
- Integration with cloud native services & infrastructure
- Accelerated Networking
- Smart & Tiered Licensing
- Dynamic Policy
- Quickstarts, Infrastructure as Code and Automation
- Gateway Load balancer integration
- Snapshots

Cisco Firewall Threat Defense Virtual orchestrated by Multicloud Defense Controller

- Simplifies Secure Firewall (FTDv) deployment
  - Deploy one or more firewalls, along with load balancers
  - Configures all necessary routing
  - Identical workflow for all public cloud providers
- Simplifies Secure Firewall (FTDv) operations
  - Automatically scales in or out FTDv's to adjust to load
  - Monitors health of each FTD and replace unhealthy unit
- Performs registration and license management
- Bootstraps FTDv
  - Configure interfaces, device group, platform settings, NAT, and routing
  - Creates all objects on the FMC that are needed for FTDv deployment

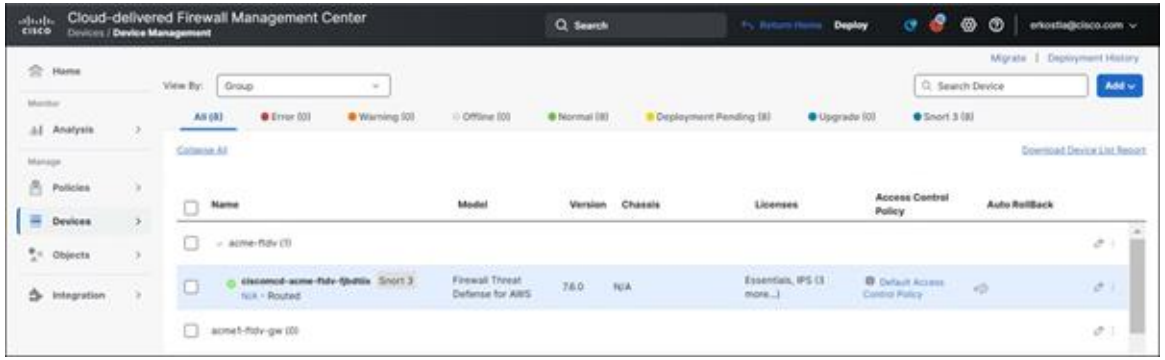

Cisco Secure Firewall Virtual Orchestrated by Multicloud Defense Controller

AWS architecture (Single AZ)

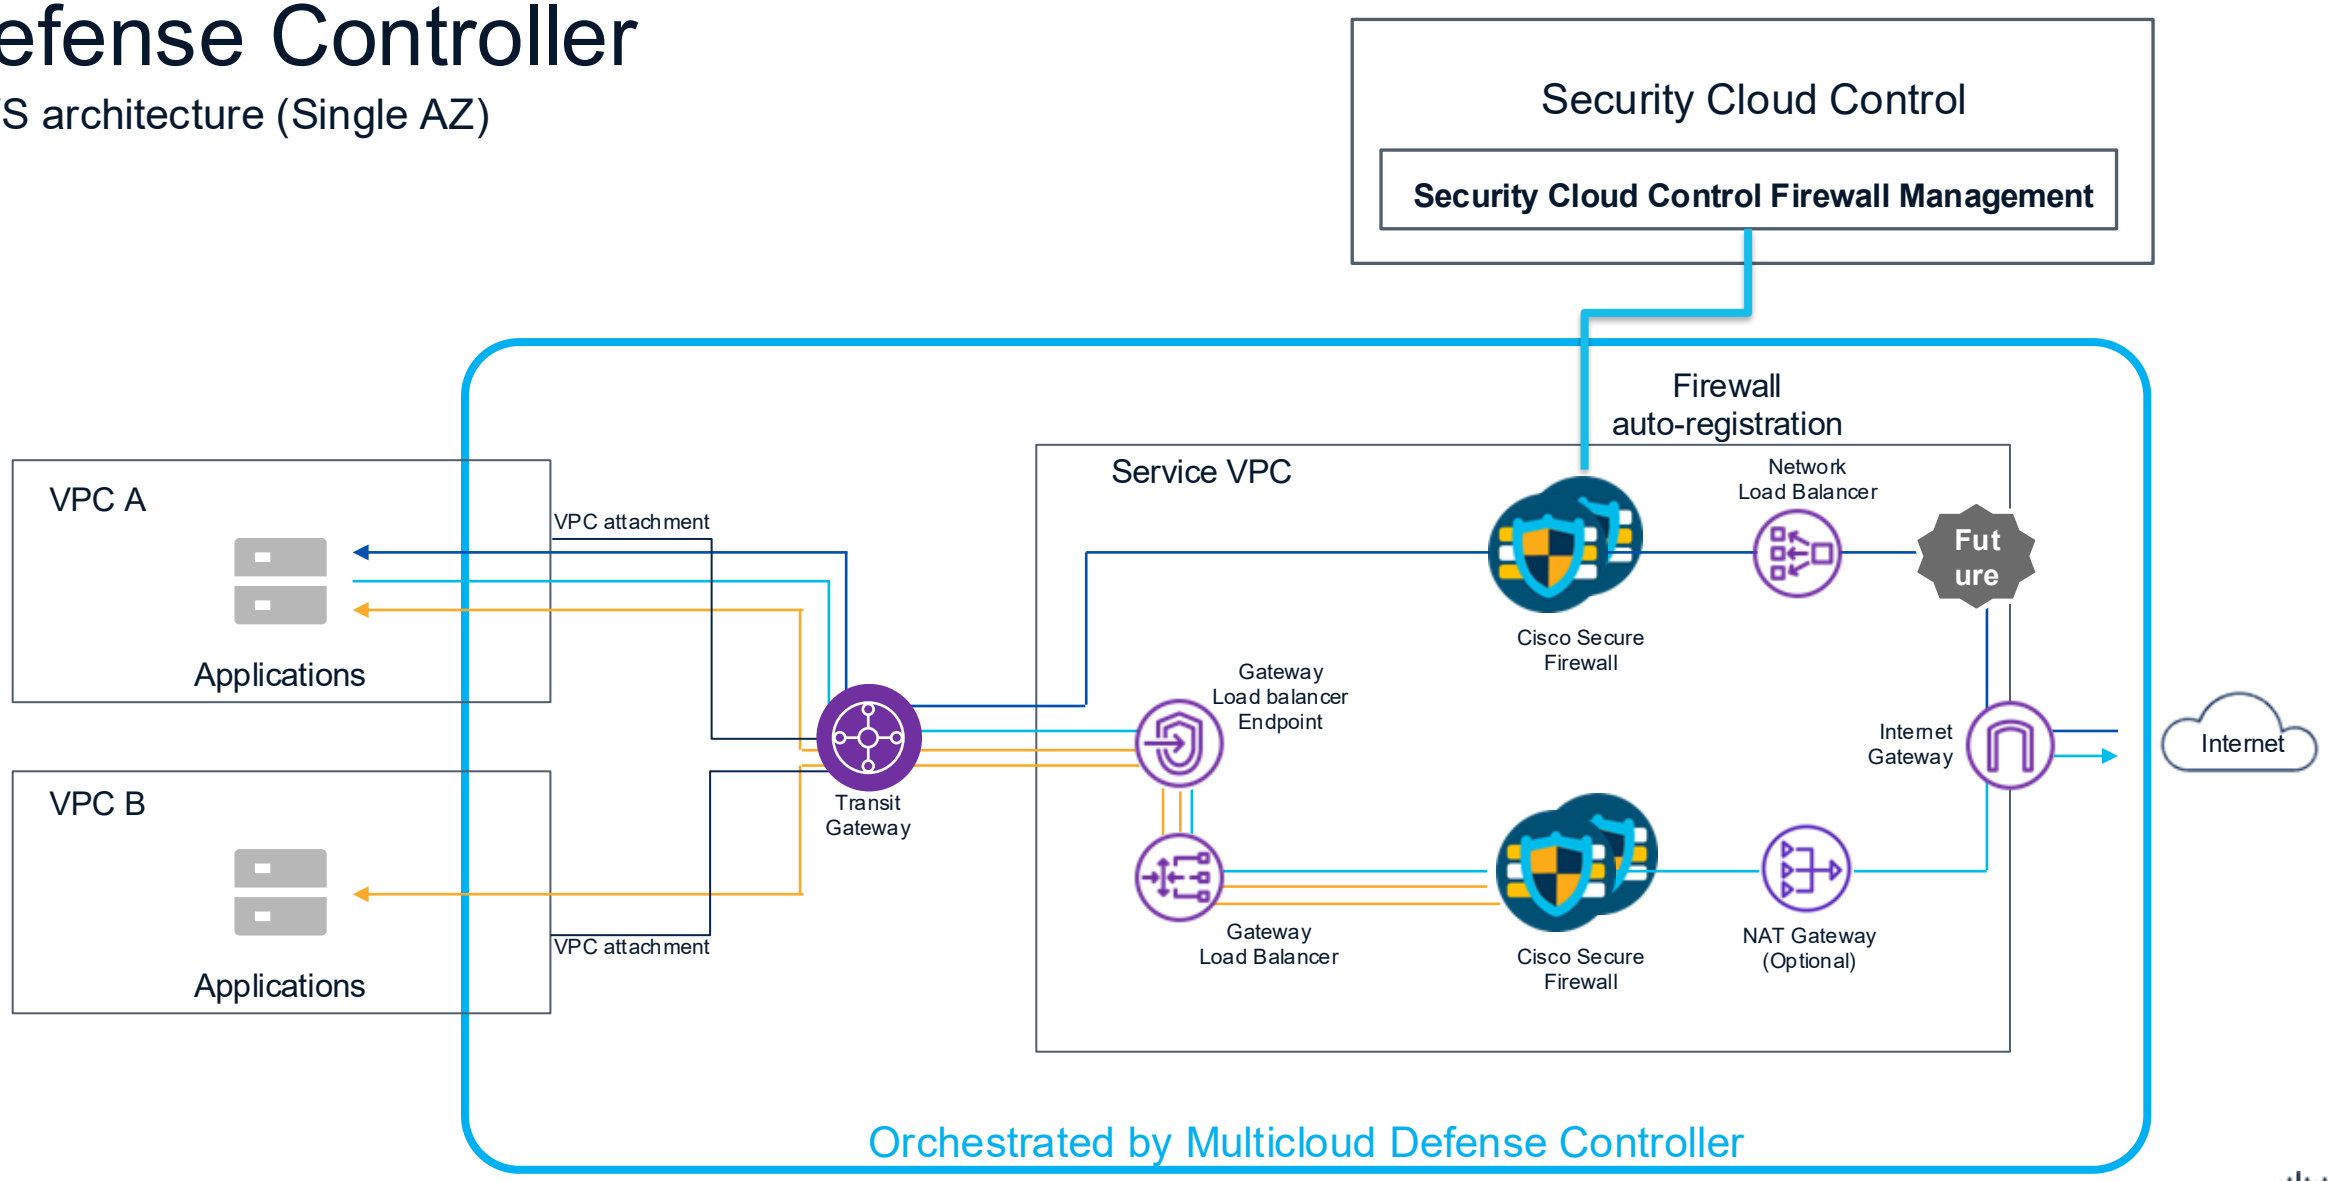

Getting closer to Zero Trust with microsegmentation

Establishing Trust with Cisco Secure Workload

Automated discovery, clustering & policy generation and verification

Baseline workload protection posture

- Network communications
- Process behaviour
- Vulnerabilities
- Labels

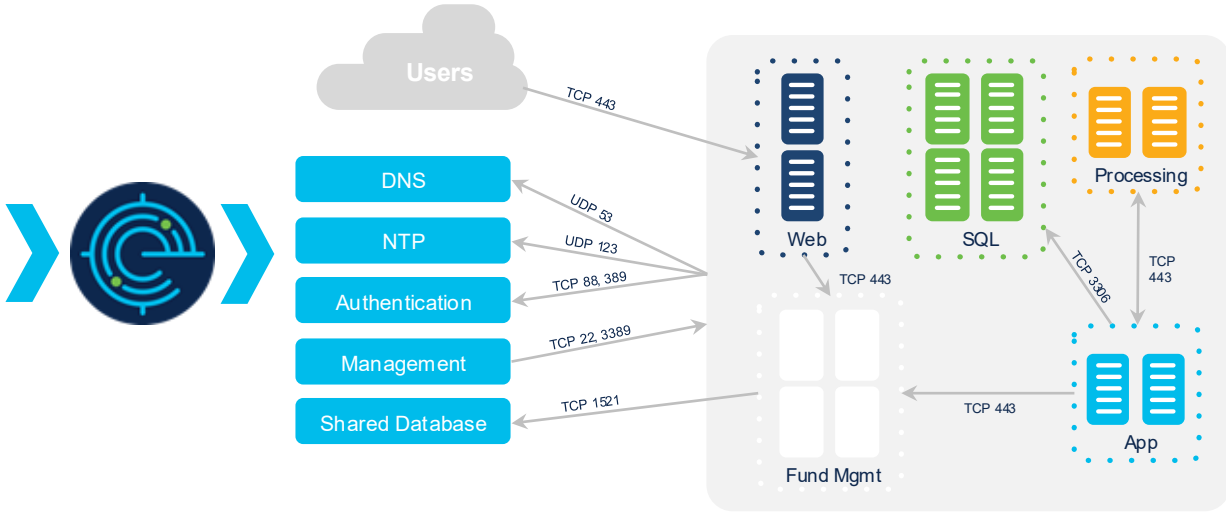

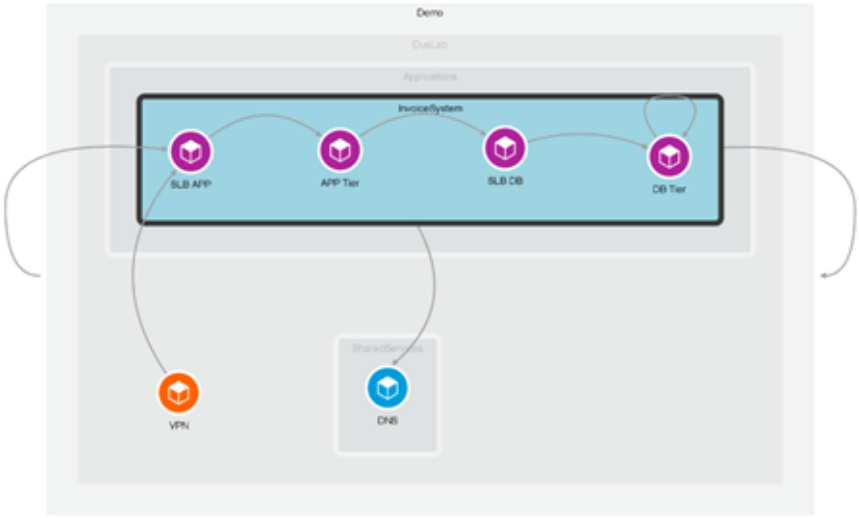

Absolute and Default Policies 13 | Catch All DENY | Grouped | Ungrouped

Rank	Priority	Action	Consumer	Provider	Protocols And Ports
Absolute	100	ALLOW	VPN	SLB APP	TCP : 1936
Default	100	ALLOW	VPN	SLB APP	TCP : 80 (HTTP)
Default	100	ALLOW	SLB DB	DB Tier	TCP : 3306 (MySQL)
Default	100	ALLOW	SLB APP	APP Tier	TCP : 8081
Default	100	ALLOW	Demo : DusLab : Applications : InvoiceSystem	Demo	UDP : 123 (NTP) ...2 more
Default	100	ALLOW	Demo : DusLab : Applications : InvoiceSystem	Demo : DusLab : SharedServices : DNS	UDP : 53 (DNS) ...1 more

Enforce Trust with Cisco Secure Workload

On-prem or SaaS management

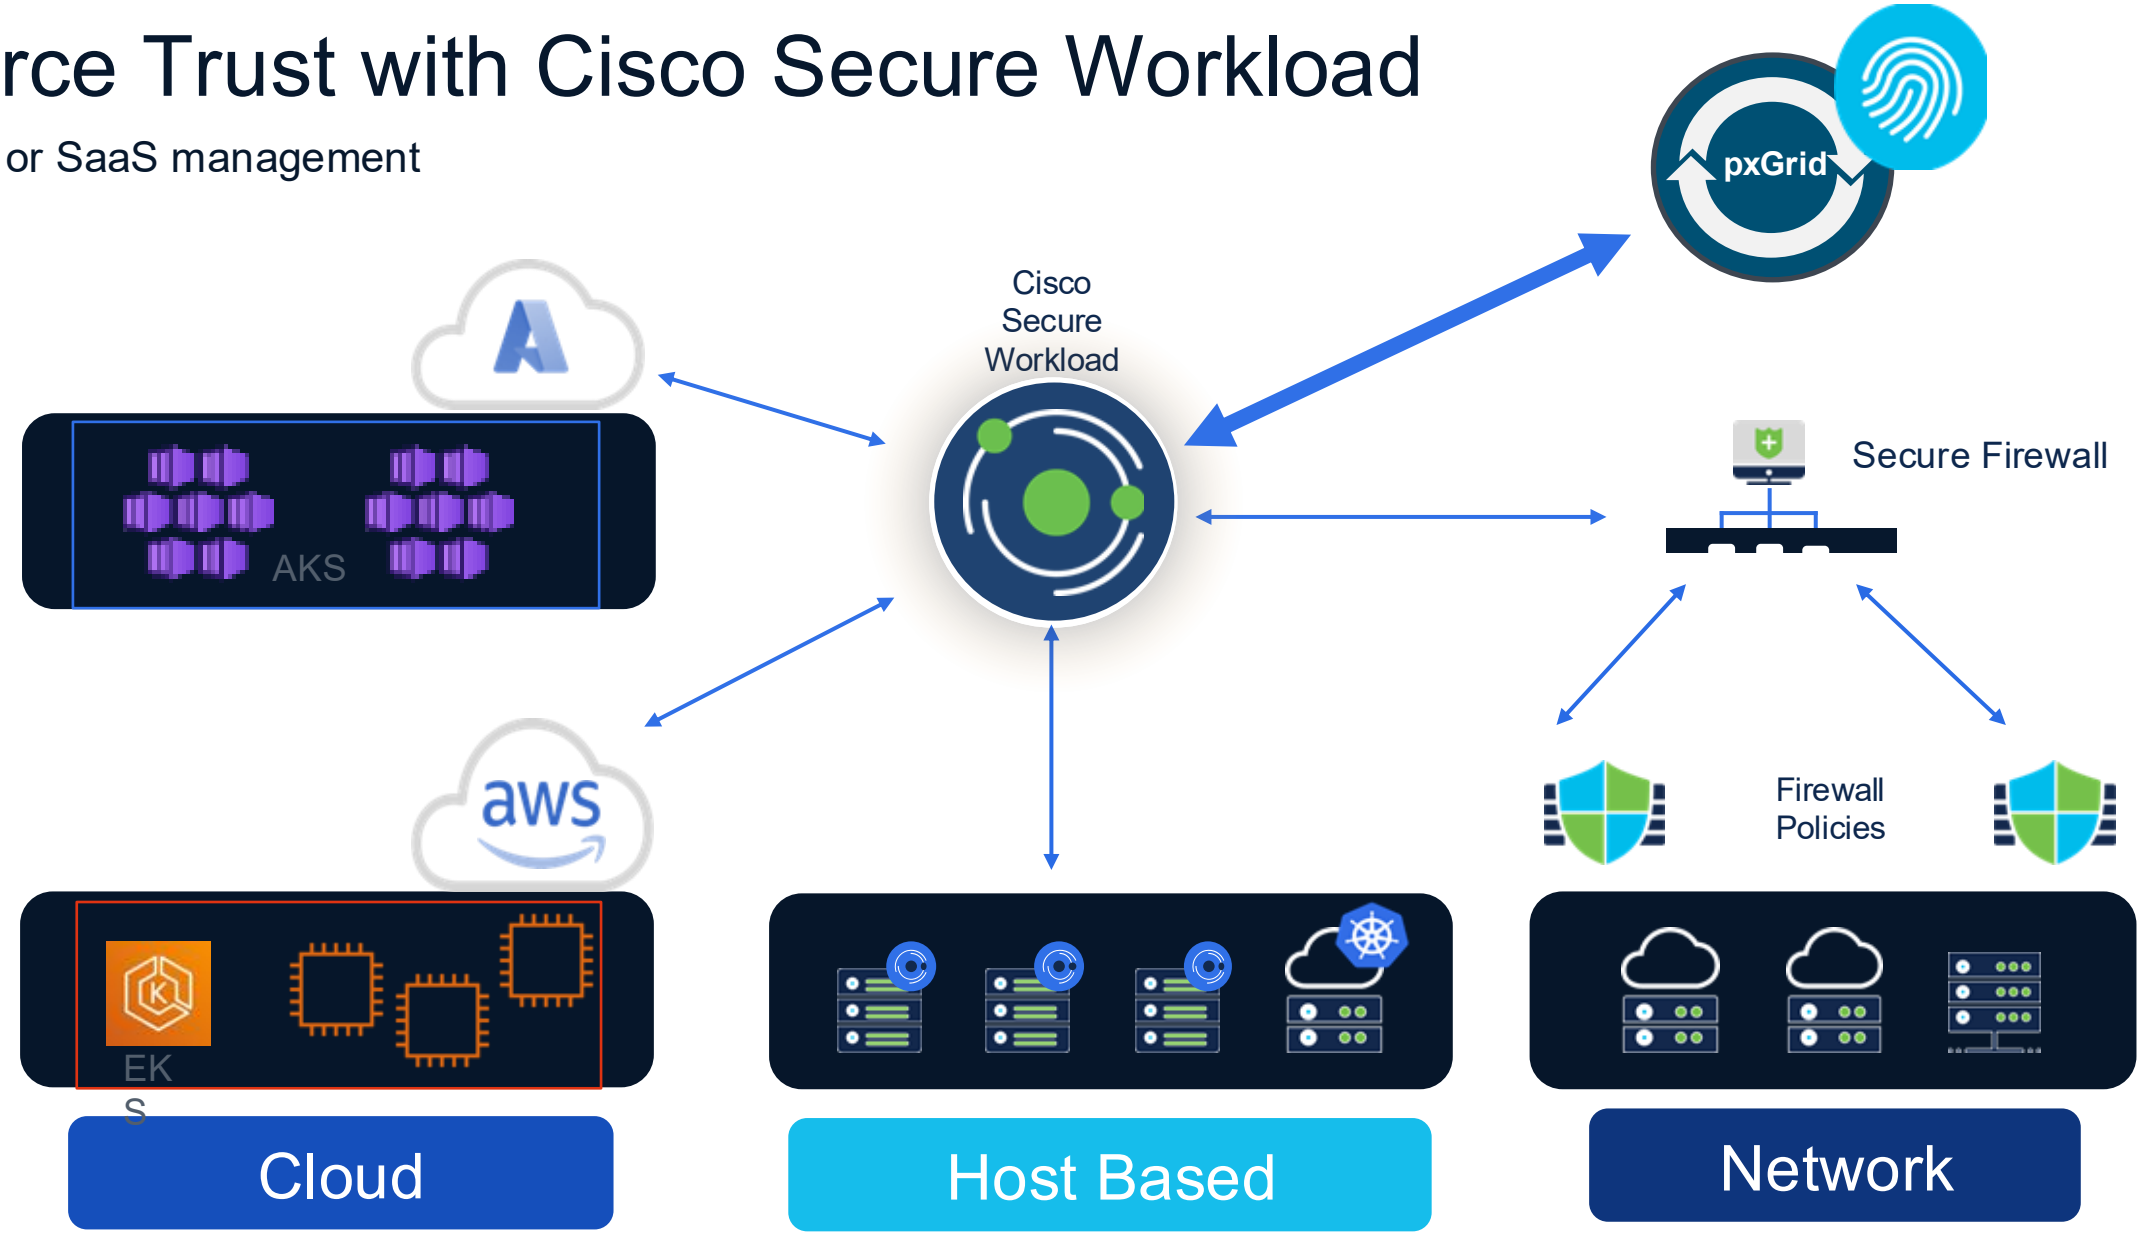

CSW Software Appliance

Background

CSW appliance was tightly coupled with Cisco UCS hardware and was not qualified to run on other hardware. Many customers are looking for utilizing existing or non cisco hardware for CSW

What's New

CSW 4.0 will come with software appliance and can run on VMware cluster. The customers can use their existing hardware for VMware cluster and CSW

Benefits

- Reduces CAPEX for micro segmentation adoption.
- Efficient use of hardware and better ROI
- Improved operational efficiency
- Scales to 40,000 workloads

Process and User based policy for Windows & Linux Agents

Background

Process/application-based policies are needed for some environments for better security, this was available on Windows only

What's New

With 4.0 release introduces process based policy for Linux

Benefits

- Granular security policies based on application / process / UID /GID
- Critical application workloads secured from malicious connections

No big fan
of using agents ?

Agentless visibility & enforcement everywhere

Public Cloud

- Centralized cloud-onboarding & cloud connectors
- Visibility with real-time discovery of workloads and labels & Flow telemetry via VPC/VNets flow-logs
- Enforcement using Security Groups (AWS), Network Security Groups (Azure), Firewall (GCP)

On-prem with ACI

- Realize ACI vision of App Centric deployment
- Automated Policy Lifecycle management

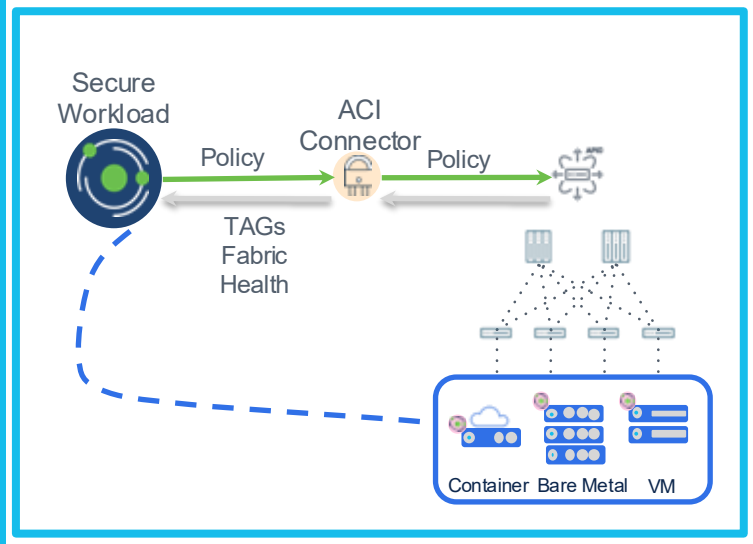

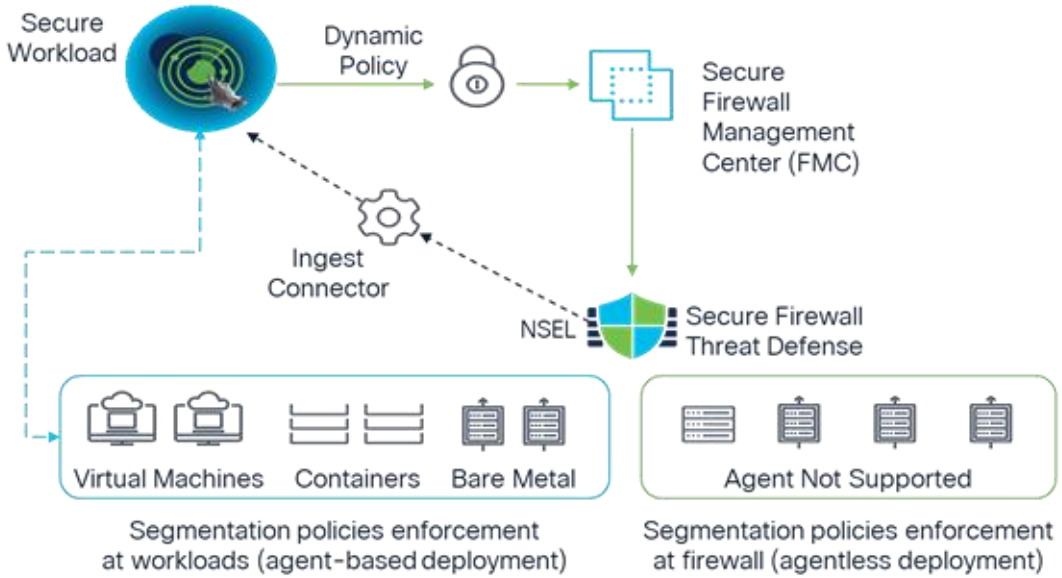

With Firewalls and Switching

- Topology aware policy enforcement
- Support access control policy and dynamic object

Secure Workload – ACI Integration

New in 4.0

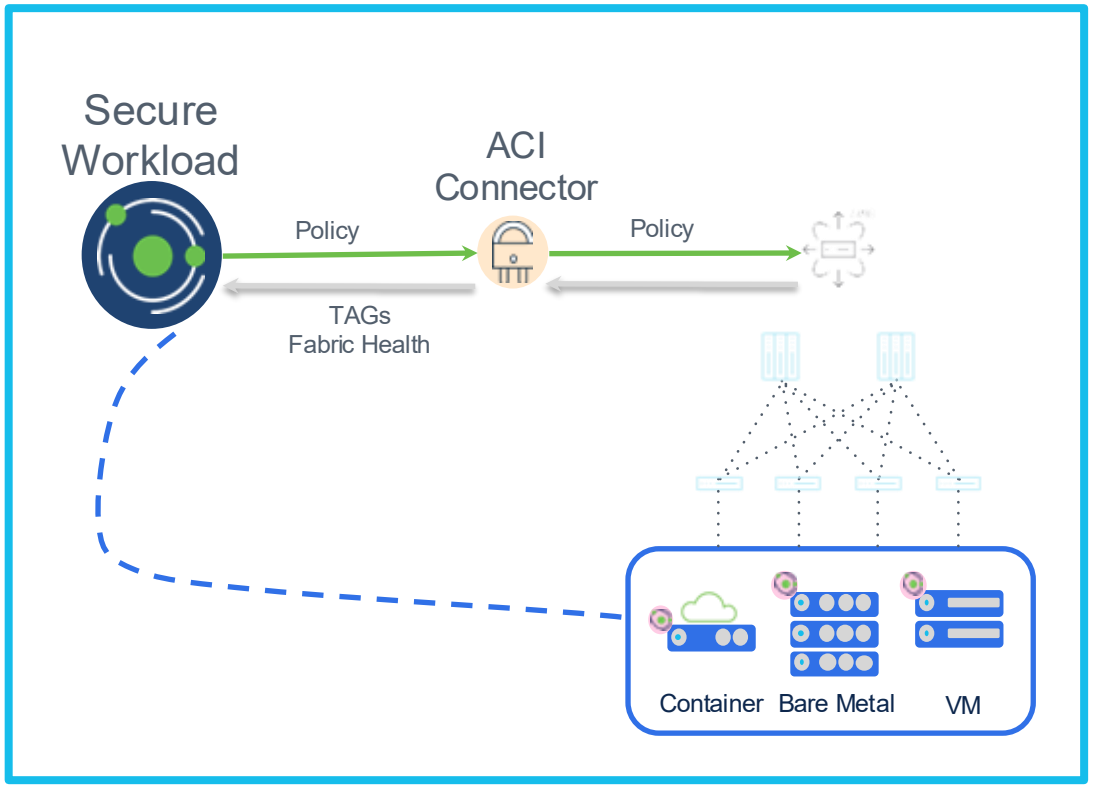

- Agentless policy enforcement
- Realize ACI vision of App Centric deployment
- Automated Policy Lifecycle management
- Faster Time to Value
- Foster collaboration between NetOps, App and NetSec team

CSW Policies Pushed to ACI

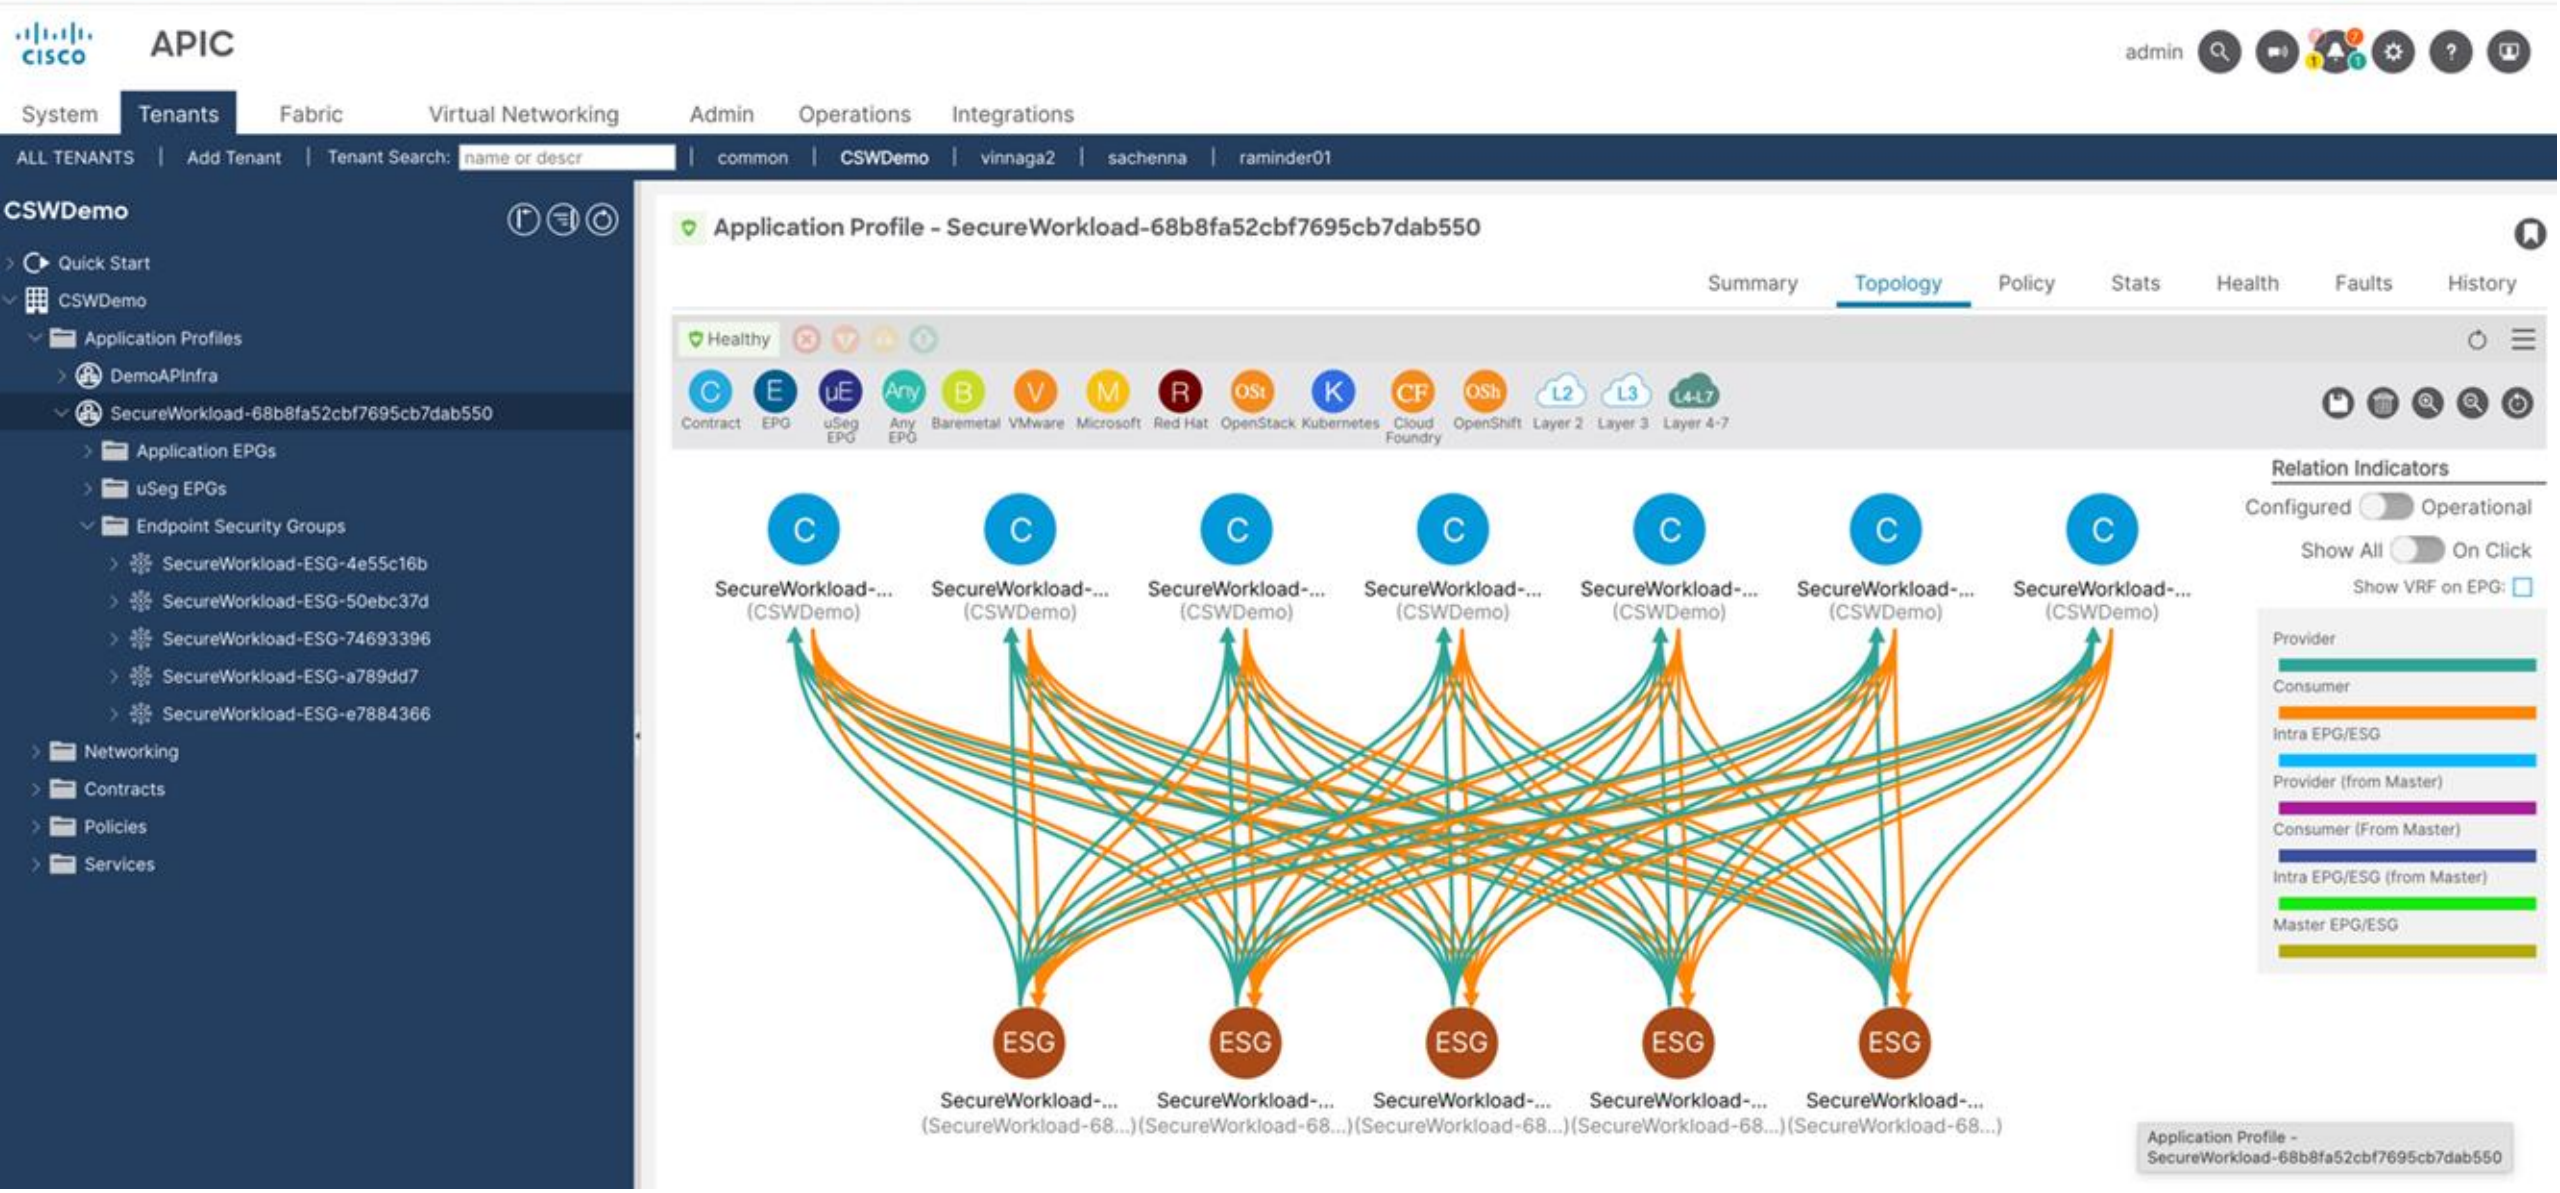

Cisco Secure Workload Cloud-Based Sources

Segmentation Policy

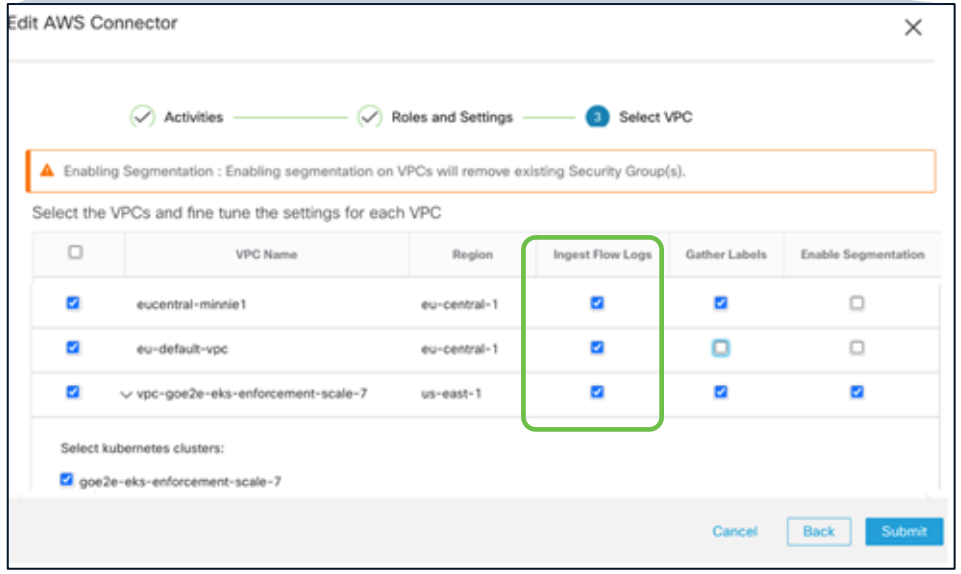

Flow Logs

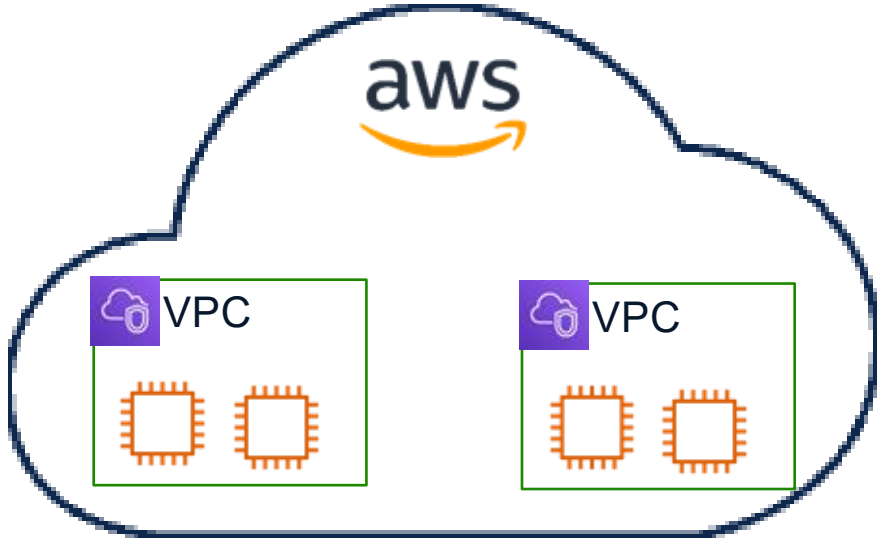

AWS Connector consolidates:

- VPC flow logs ingestion
- Context gathering (AWS tags and labels)
- AWS cloud-managed Kubernetes orchestration (Kubernetes object labels and annotations)
- Enforcement via cloud native constructs

Secure Workload & Firewall Integration And Enforcement

Firewall Segmentation - Enforce Trust without agents

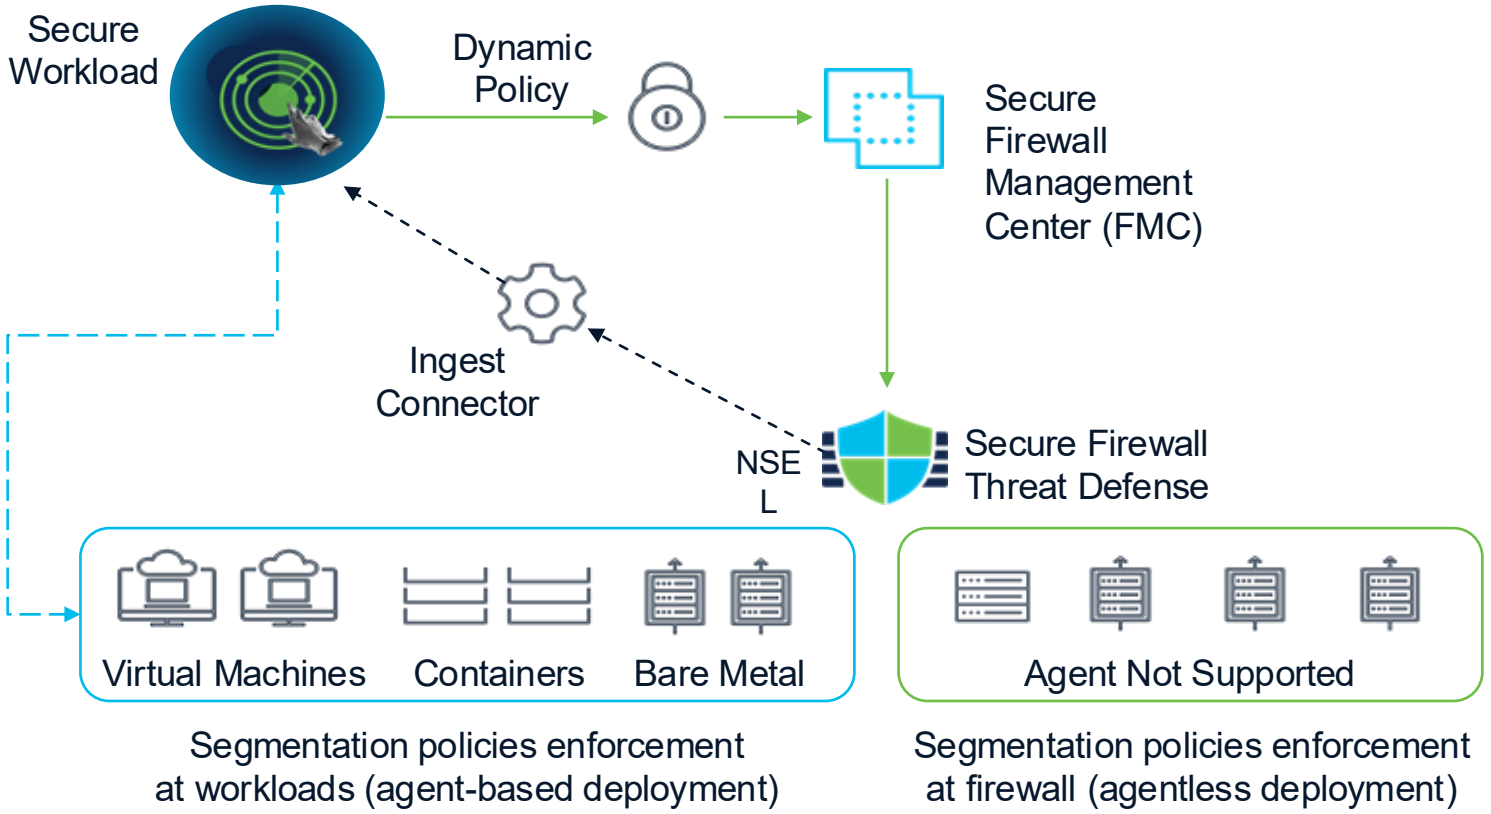

- Defense in depth
- Support access control policy and dynamic object
- Improved rule ordering
- FMC domain awareness
- Topology aware policy enforcement
- Support deployment across multicloud

Secure Workload Policy Orchestration in FMC

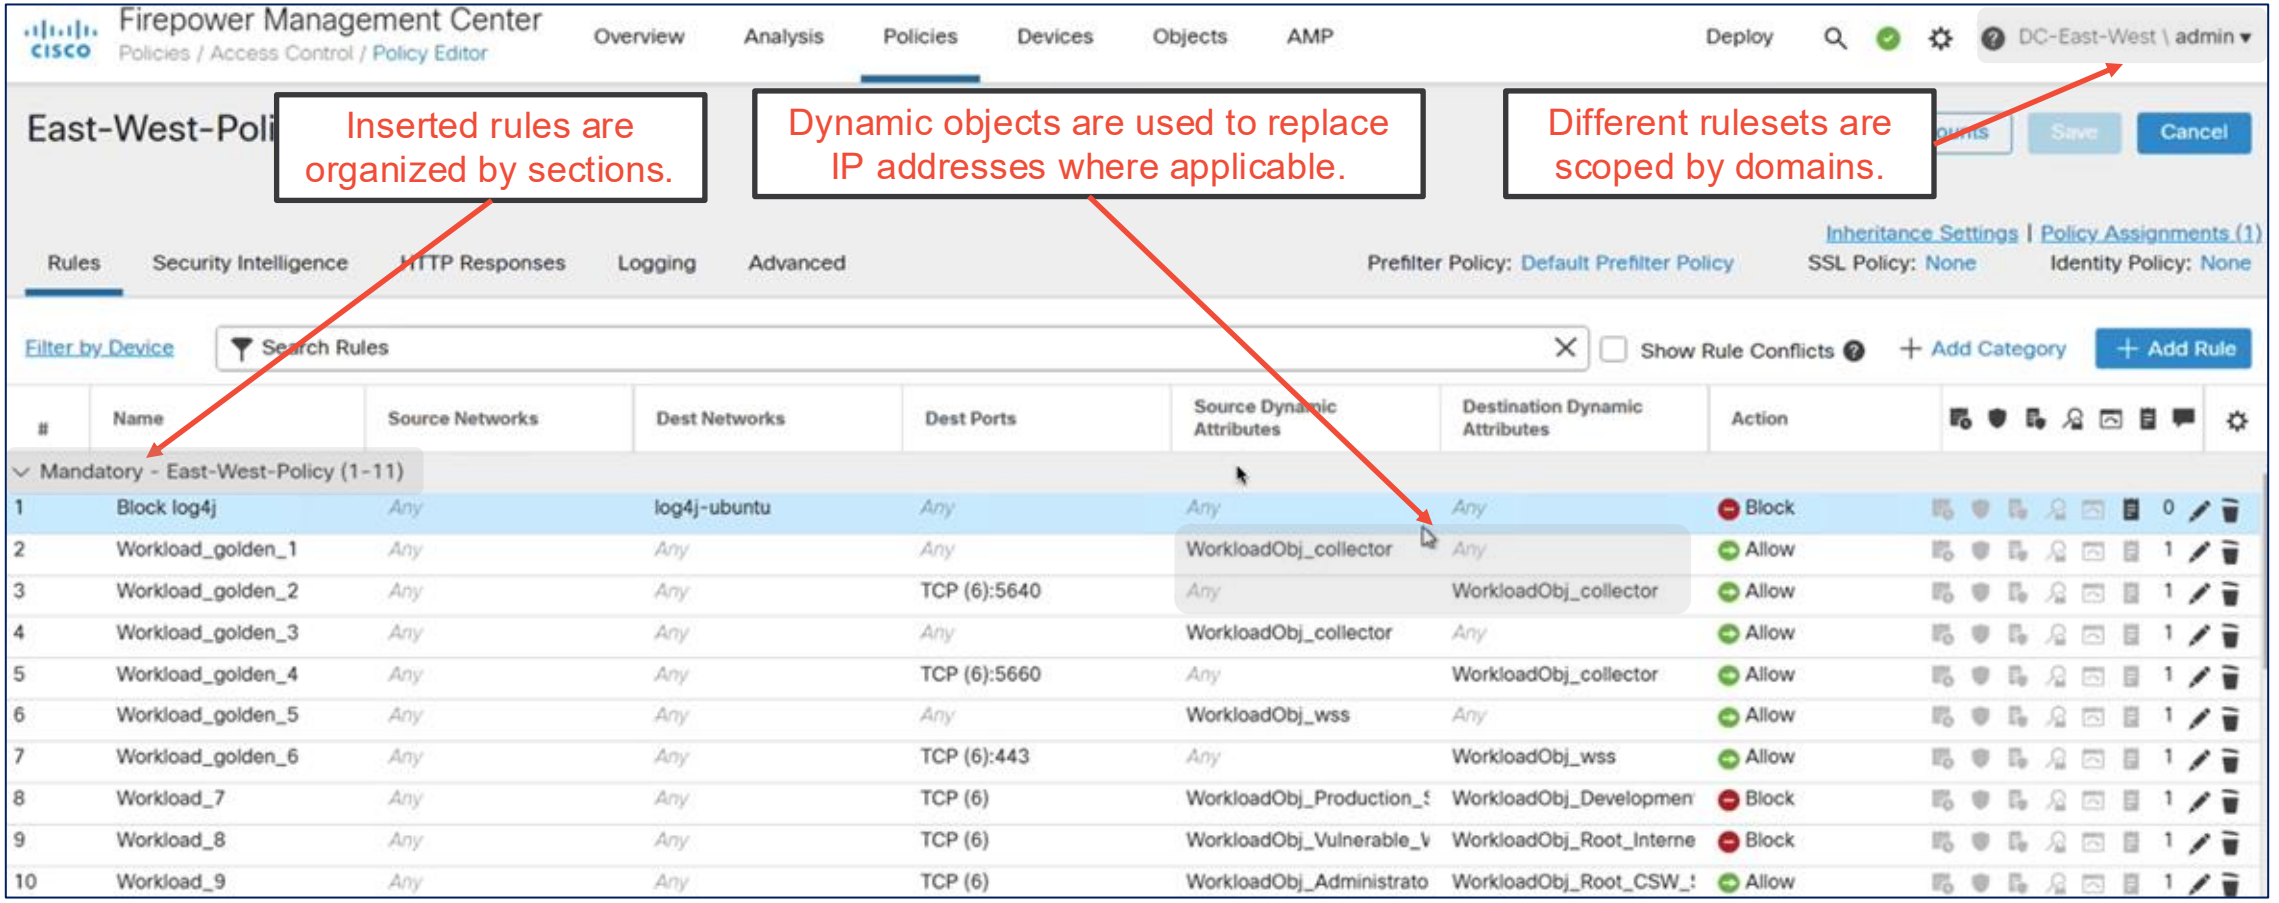

Continuously Verify & Respond

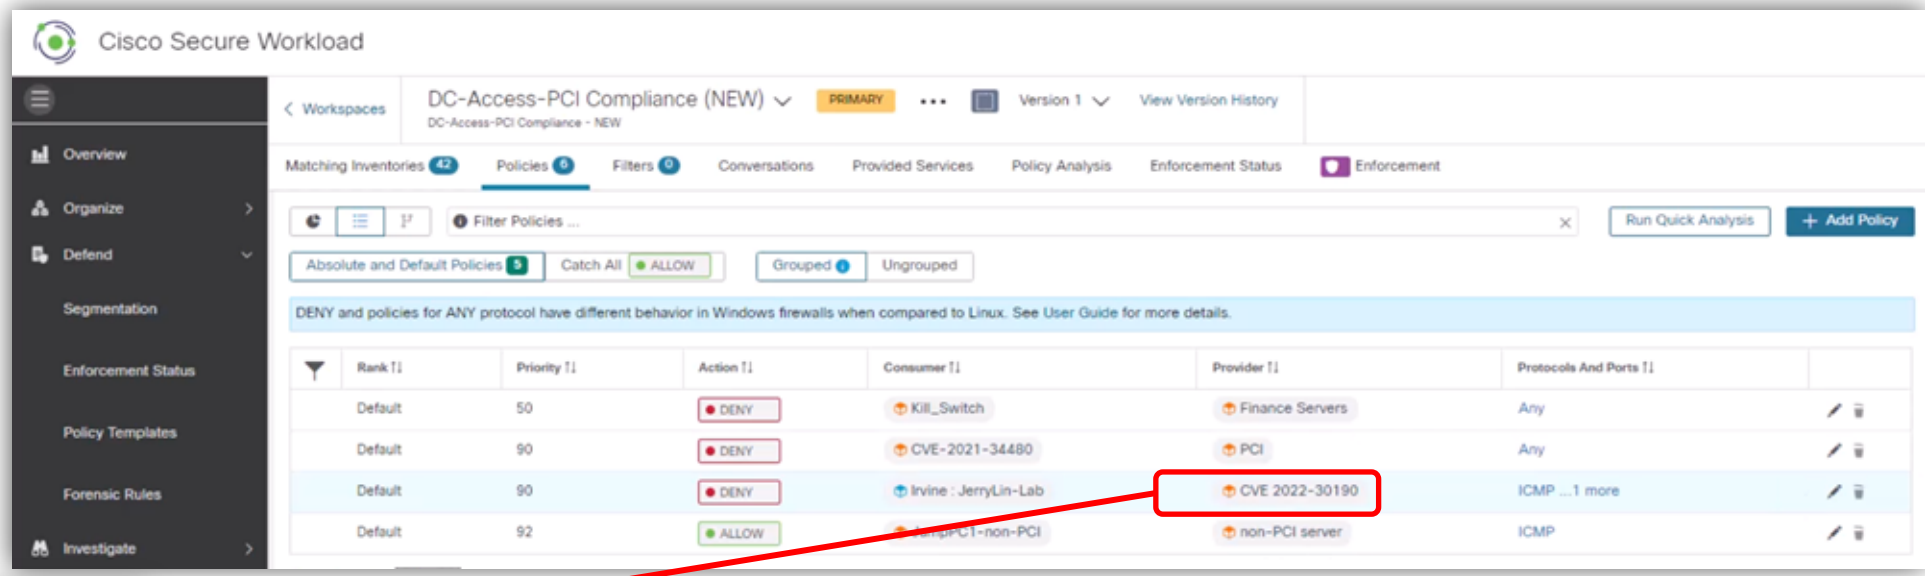

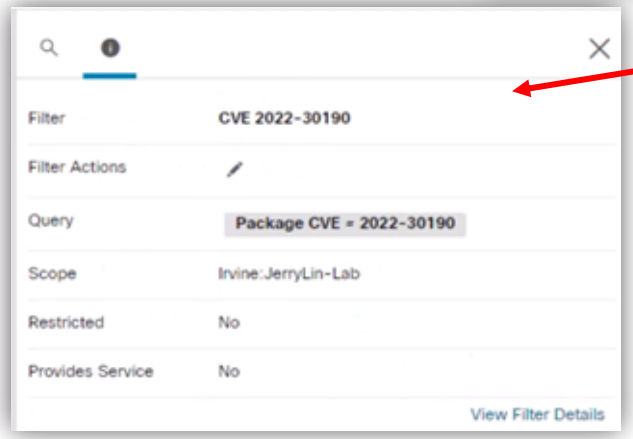

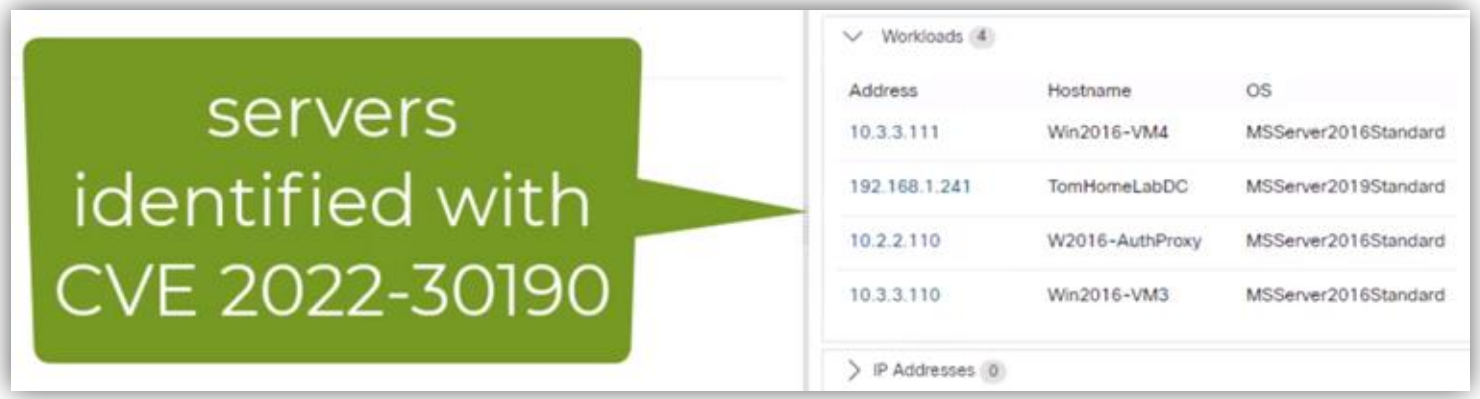

Application **workloads** evolution

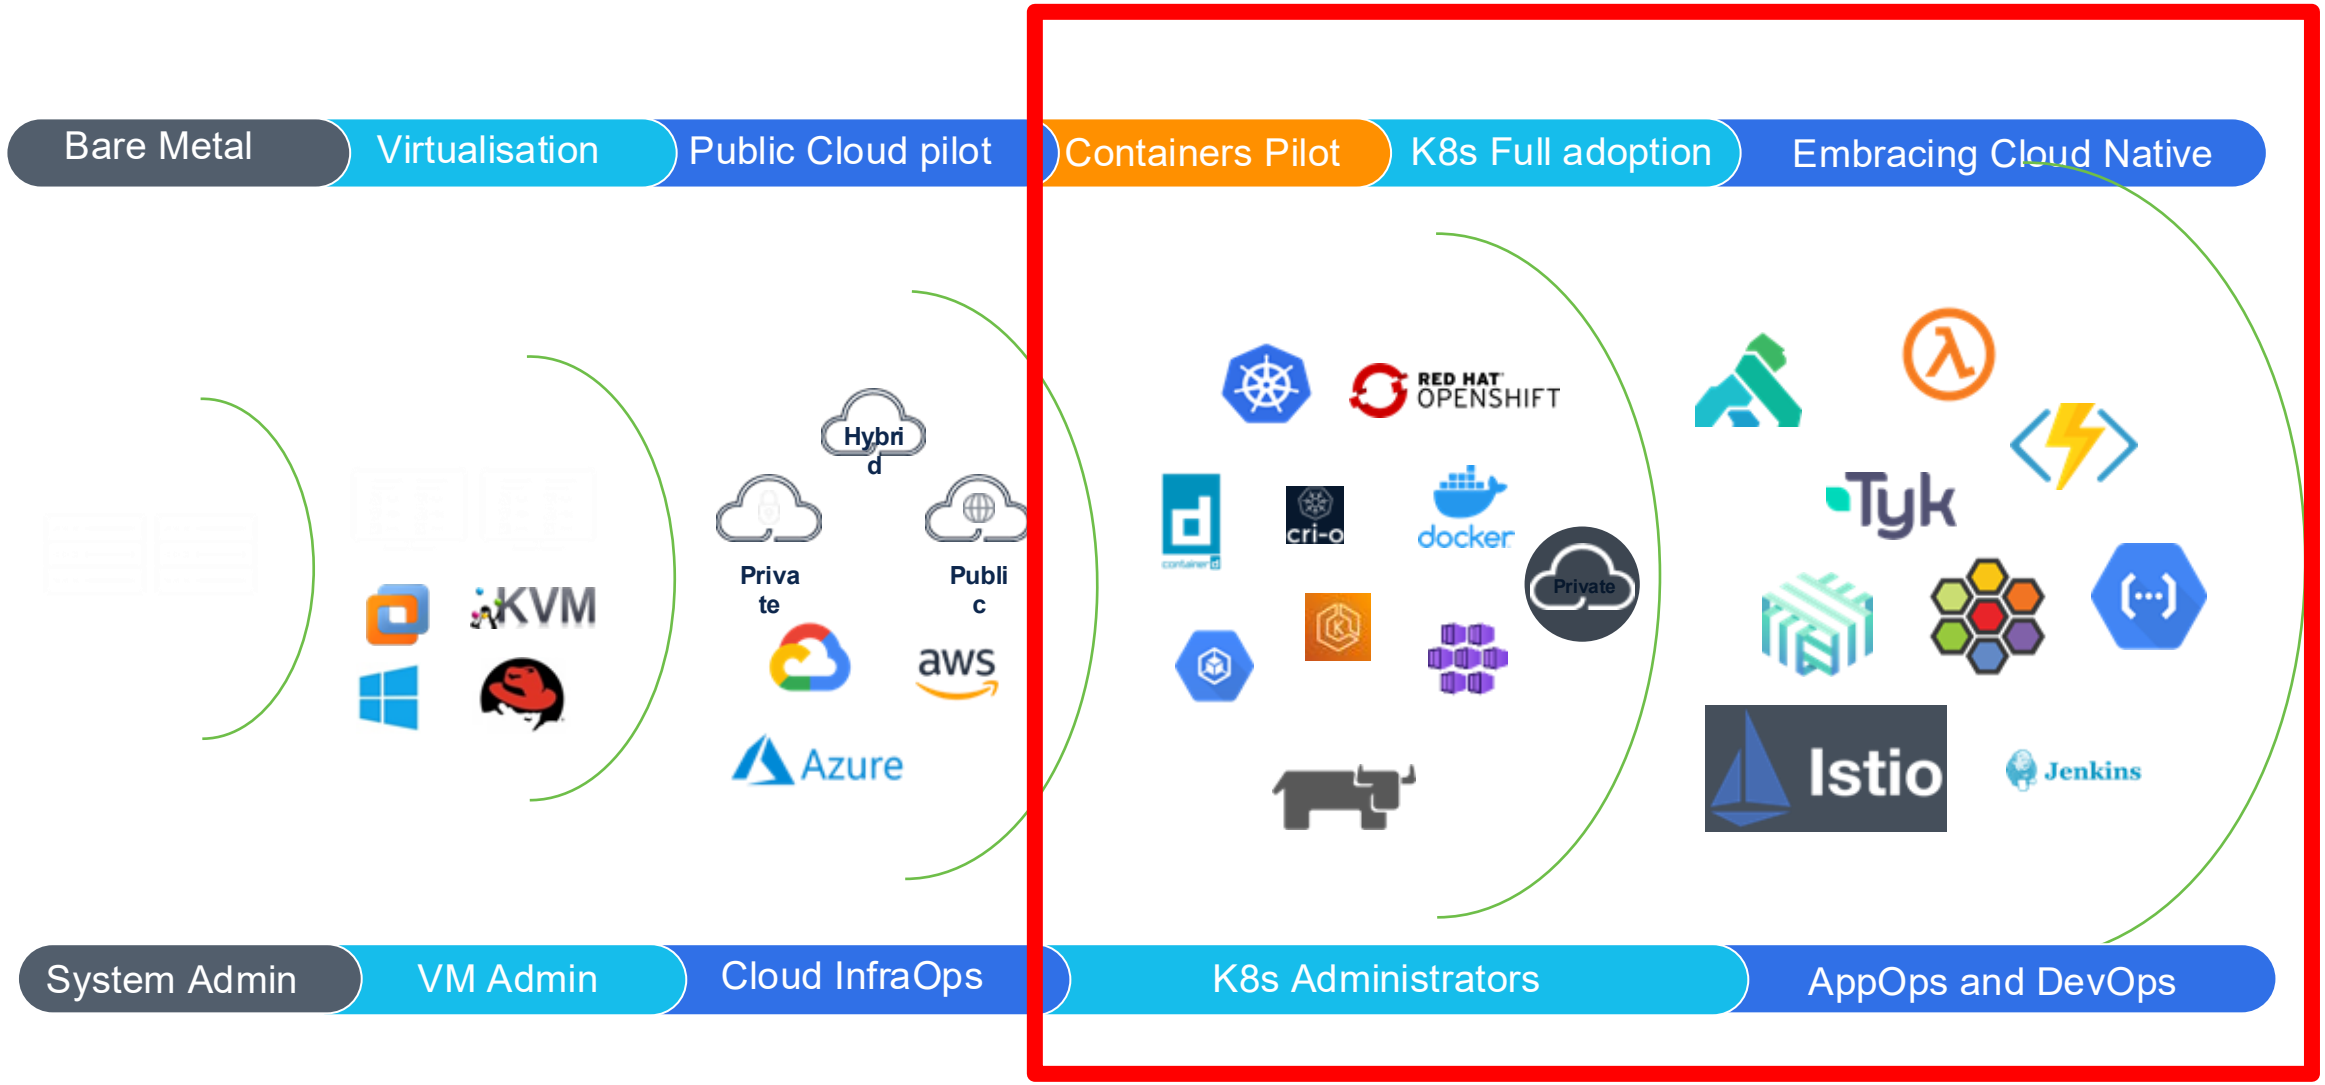

Solution to Regain the Visibility and Control Enterprises Need

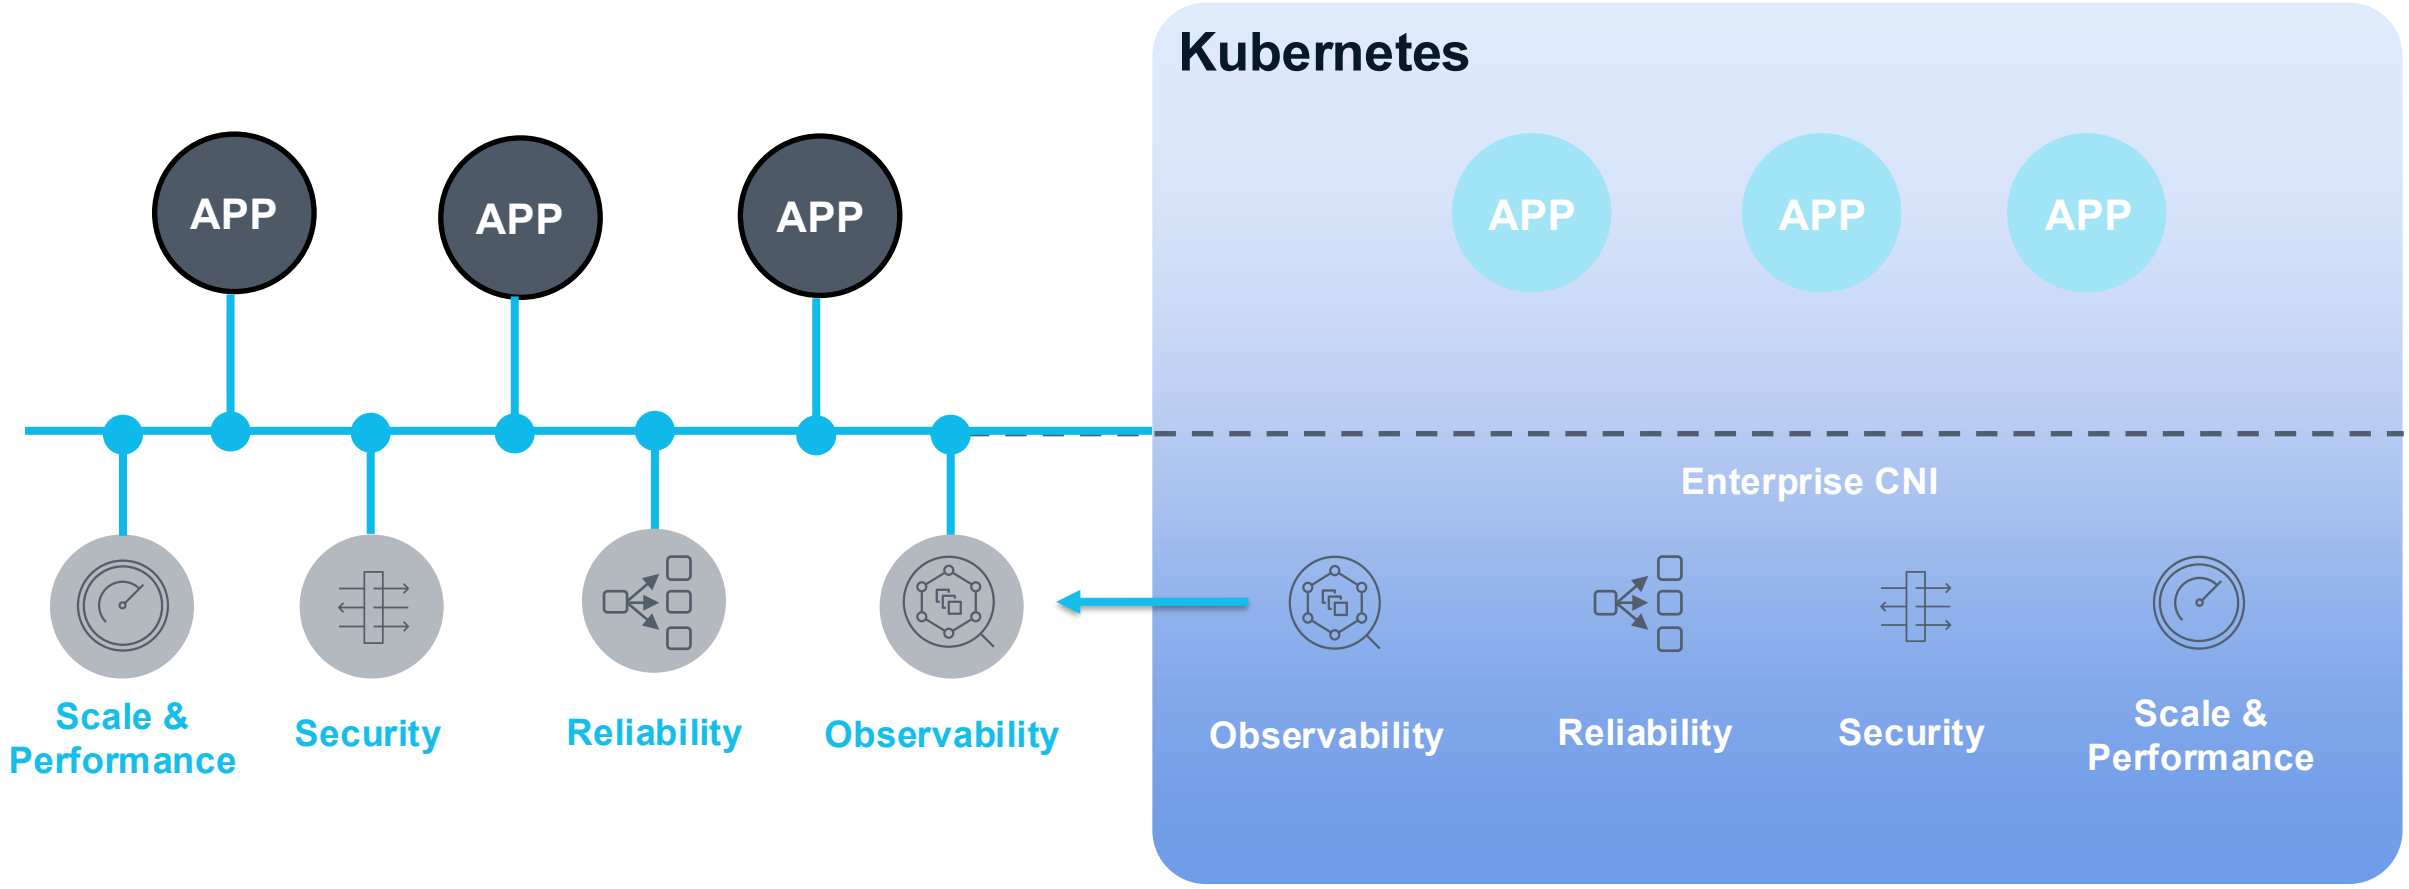

From Open-Source Innovator to Enterprise Platform

ISOVALENT
now part of cisco

The team behind the innovation. The platform for mission-critical environments.

Open-Source Leadership

- **Creators and maintainers of Cilium**, the leading cloud-native networking project.
- **Creators of Tetragon**, the eBPF-based runtime security and observability engine.
- **Key contributors to the eBPF ecosystem**
- Powering technologies **trusted by hyperscalers like AWS, Google, and Microsoft.**
- **Backed by a vibrant open-source community** and CNCF ecosystem.

ISOVALENT
now part of cisco

Enterprise-Ready Platform

- **Hardened** enterprise platform
- **Enterprise features** for production use
- Enterprise-grade **support and SLAs**
- **Built-in capabilities** for compliance, threat detection, and forensics.
- **Trusted by regulated industries** and global enterprises.

cisco

Isovalent Enterprise Platform

ISOVALENT
now part of cisco

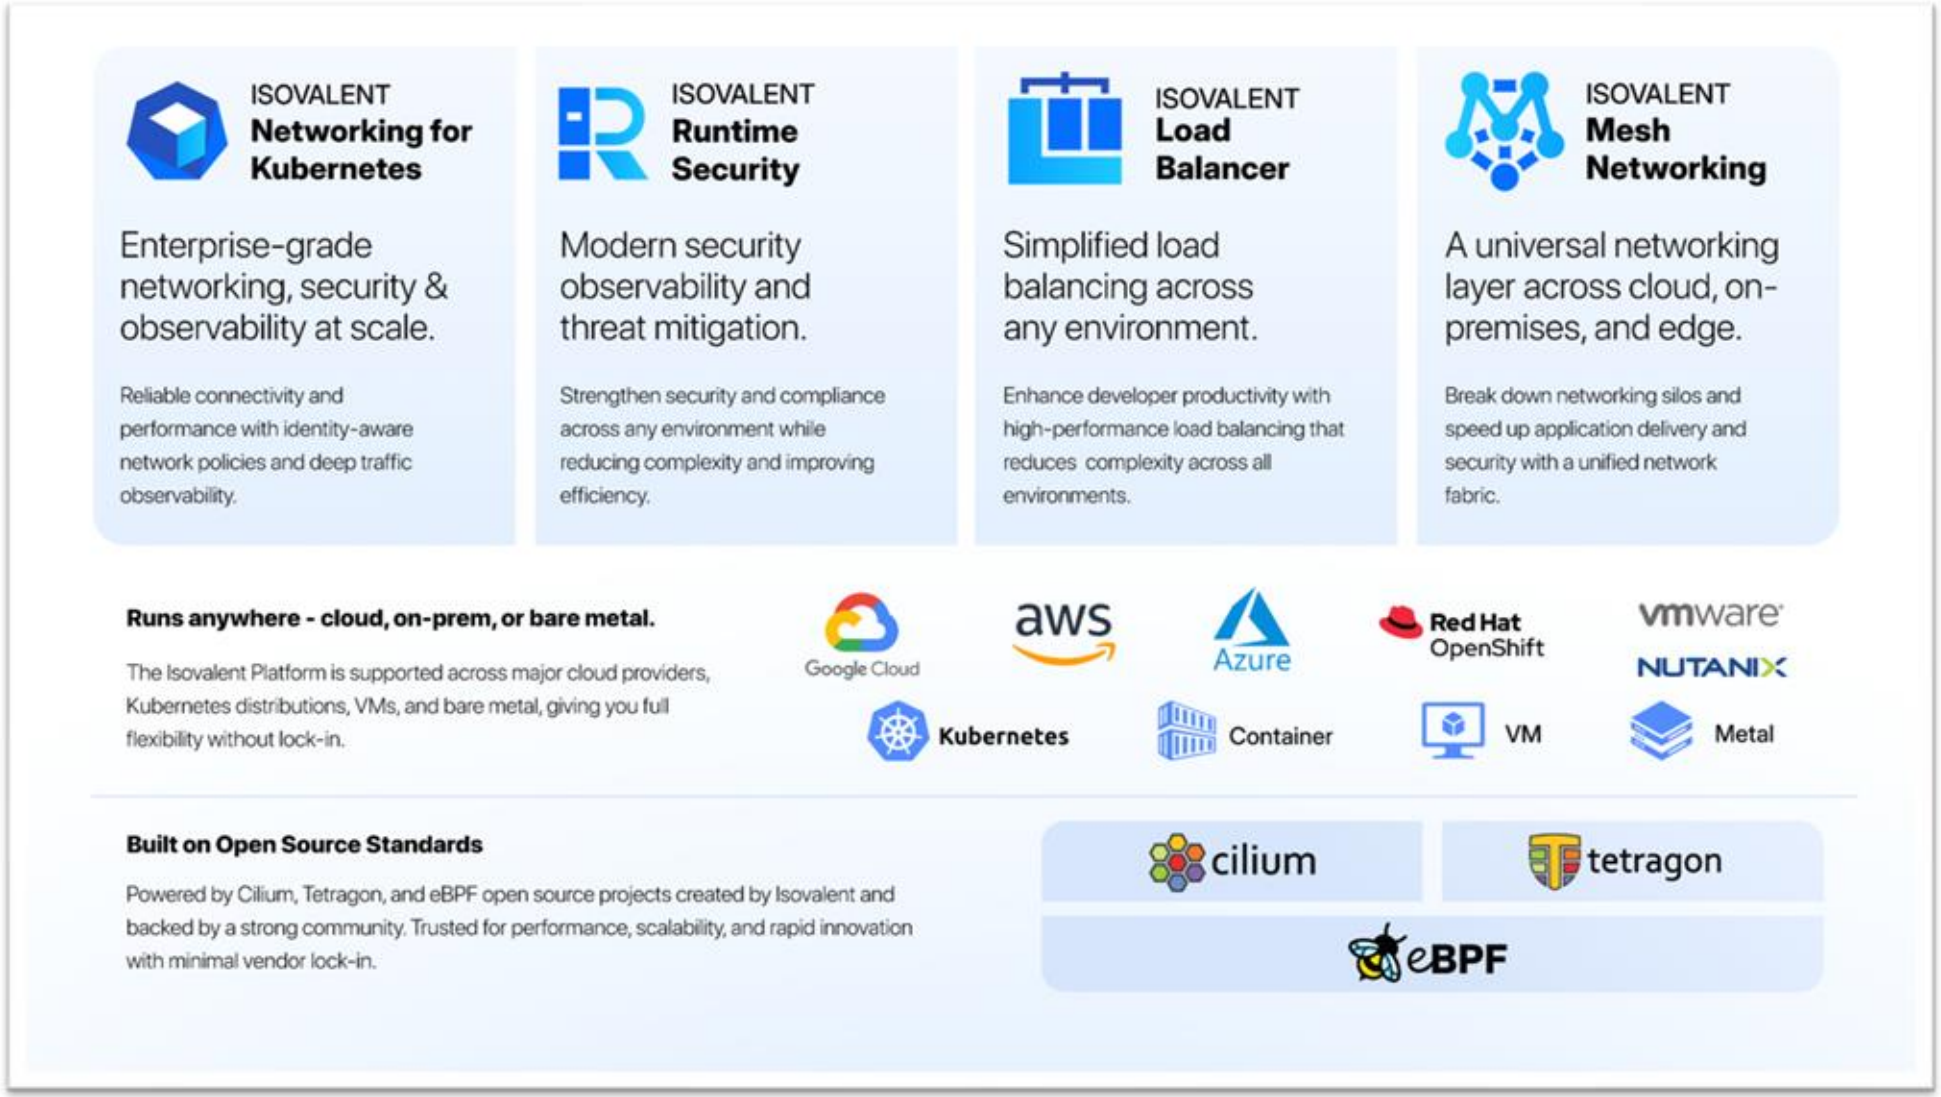

cisco

eBPF – Extended Berkely Packet Filter

Makes the kernel programmable in a secure and efficient way

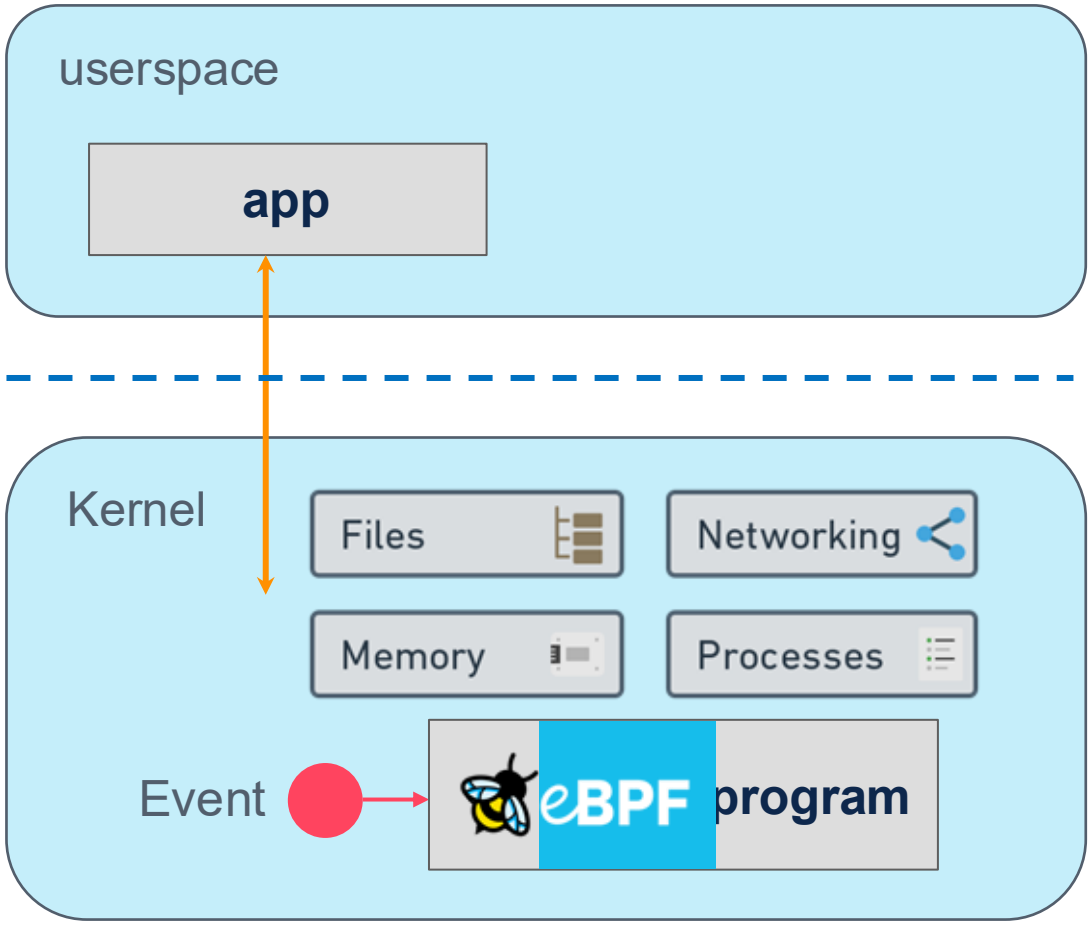

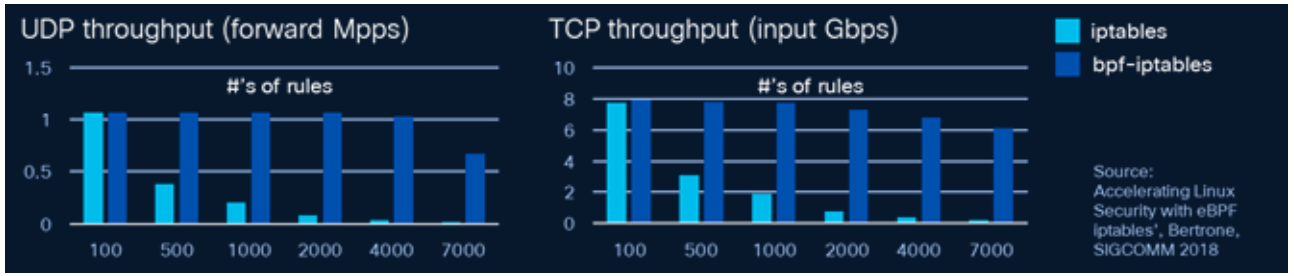

System calls

- eBPF runs in a restricted execution environment
- Before eBPF script is compiled:
  - Verifier ensures things such as memory safety
  - Hardening process
    (program execution protection, mitigation against spectre, constant blinding, …)
- eBPF can only access pre-approved kernel functions and data structures

eBPF powers many things

Capabilities For Your Cloud Native Journey

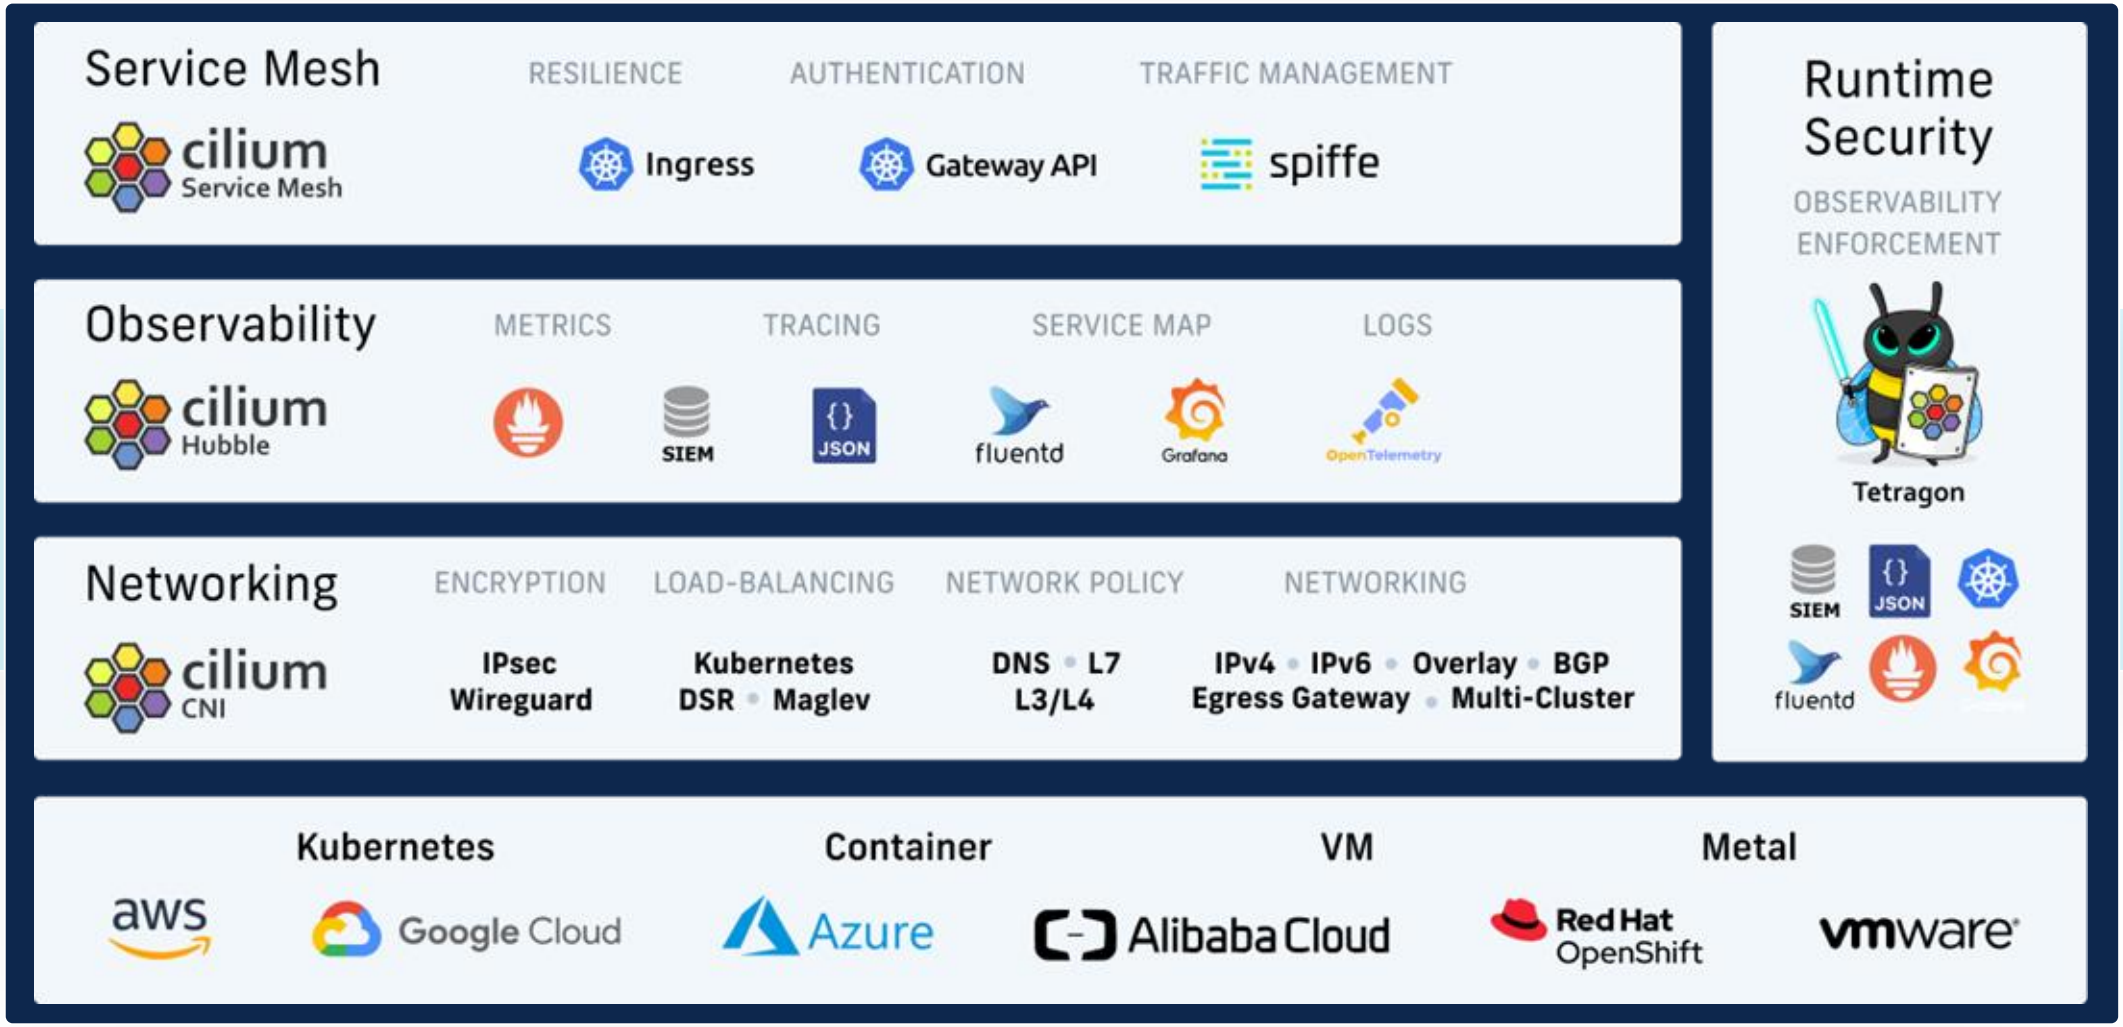

Micro-Segmentation

Label based East-West Application or Multi-tenant Security Enforcement

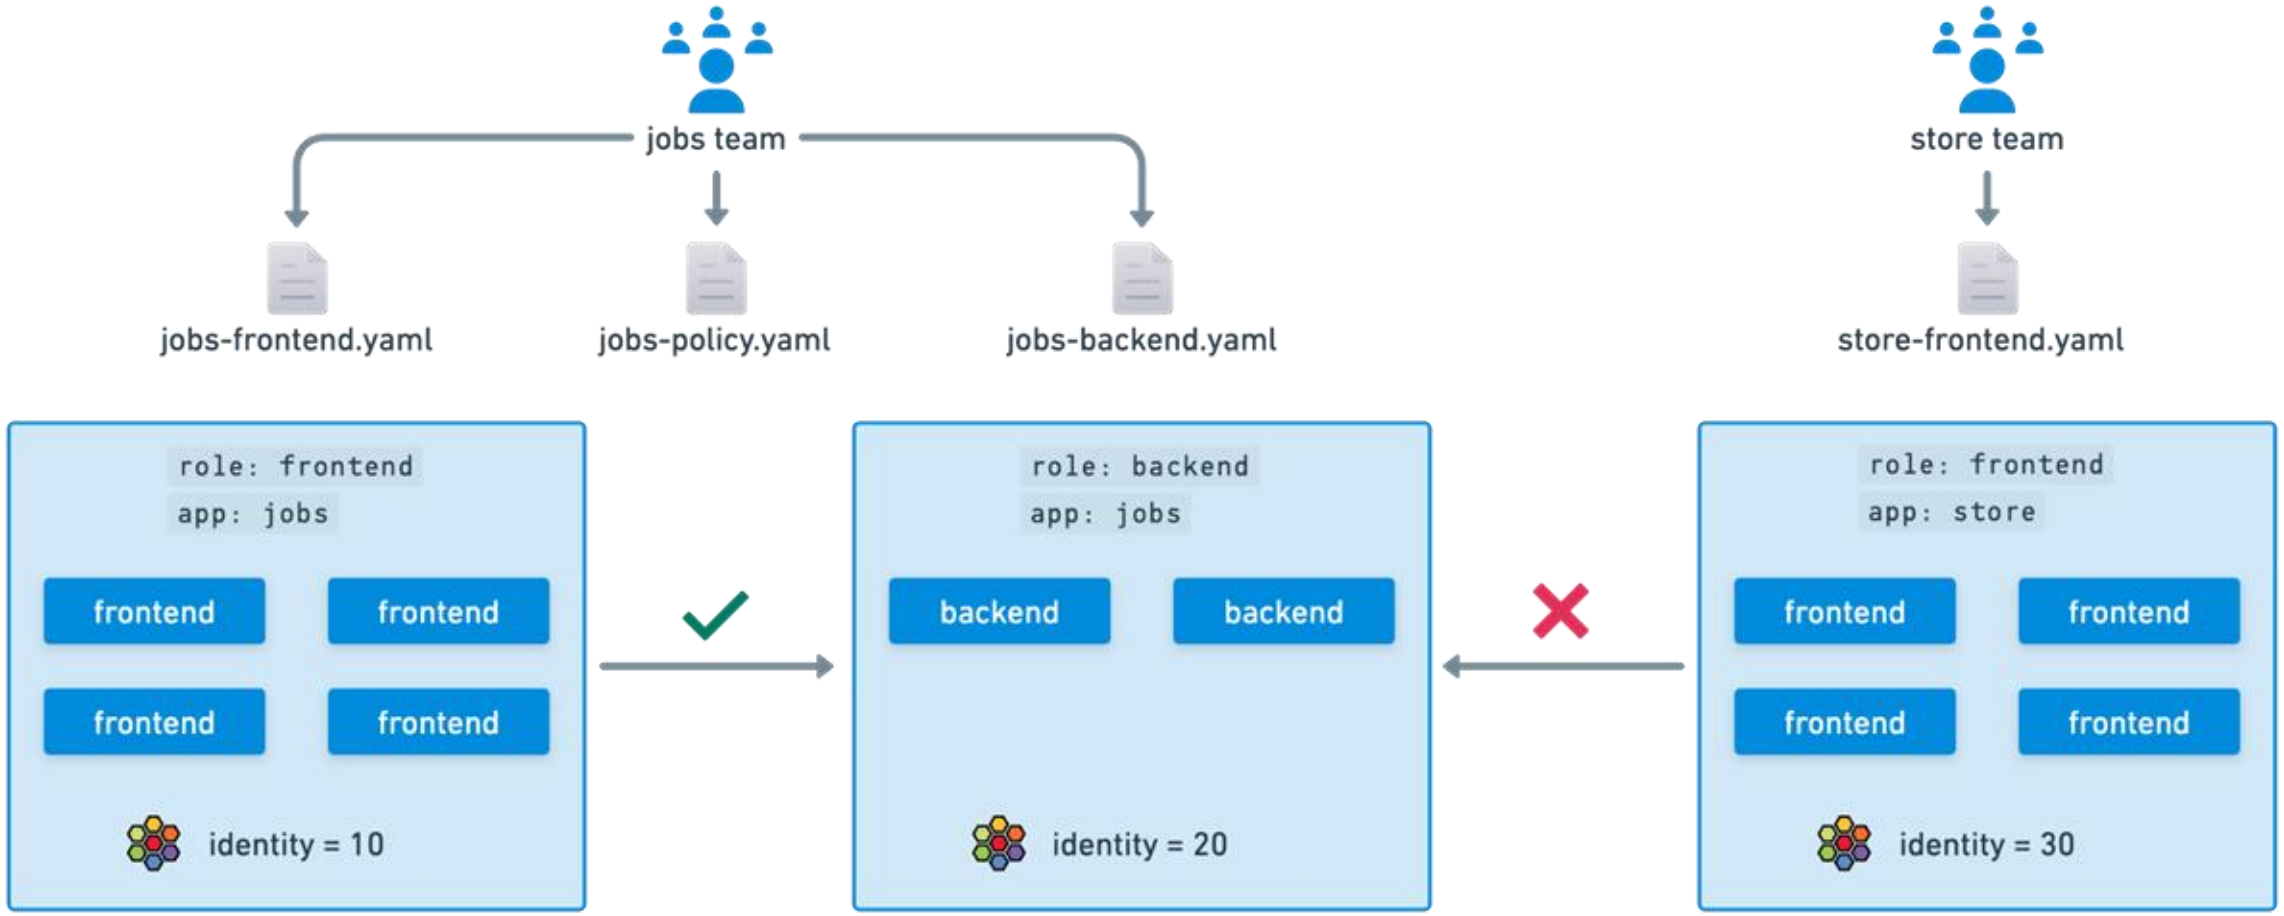

Shaping the future of cloud & cloud native

Observability, Networking & Security with eBPF, Cilium & Tetragon

- Presence in the OS kernel for ultimate efficiency
- Networking & Load Balancing
  - CNI, K8 Service Mesh
- Observability & Security
  - L4-7 traffic, encryption, process, data, file, identity-aware, api-aware

Securing the Application **workloads** evolution

Cisco Security Cloud Control

Bare Metal | Virtualisation | Public Cloud pilot | Containers Pilot | K8s Full adoption | Embracing Cloud Native

Cisco DC Fabric (ACI, Smart Switches)

Cisco Secure Workload (Agent)

Cisco Secure Workload (Agentless) & Multicloud Defense

Cisco Secure Firewall

Cisco Secure Workload & Isovalent Enterprise

Hypershield

Cisco Hybrid Mesh Firewall

Zero Trust Use Cases

Segmentation

Admin Access

Remote Worker

On-Premises

Application to Application

Application to Model

The Evolution of AI

Rapidly increasing autonomy and capabilities

2023

Simple Chatbots

Direct responses, basic assistance

2024

Retrieval Augmented Generation (RAG)

Enhanced accuracy & context via external knowledge.

2025+

Agentic AI

Autonomously execution of complex, muiti-step tasks.

cisco

Agentic AI in Cisco: Automating tasks, reducing friction.

Agents in a Sandbox (Gen AI)

- Creates Output (txt, images..)
- safely generates content and insights in controlled environments with limited access and guardrails

Agents in the real world (Agentic AI)

- **Takes Actions like a real user**
- **Agentic AI acts on your behalf. It's AI that can take initiative, interact with systems and data, and complete tasks end-to-end, much like a digital coworker**

Agentic AI is like a model employee – brilliant, tireless, perfect memory, and fully autonomous.

Unfortunately, it's also your least-supervised, immoral, root access user with a knack for outsmarting your defenses!

What’s the risk?

AI Applications can be non-deterministic

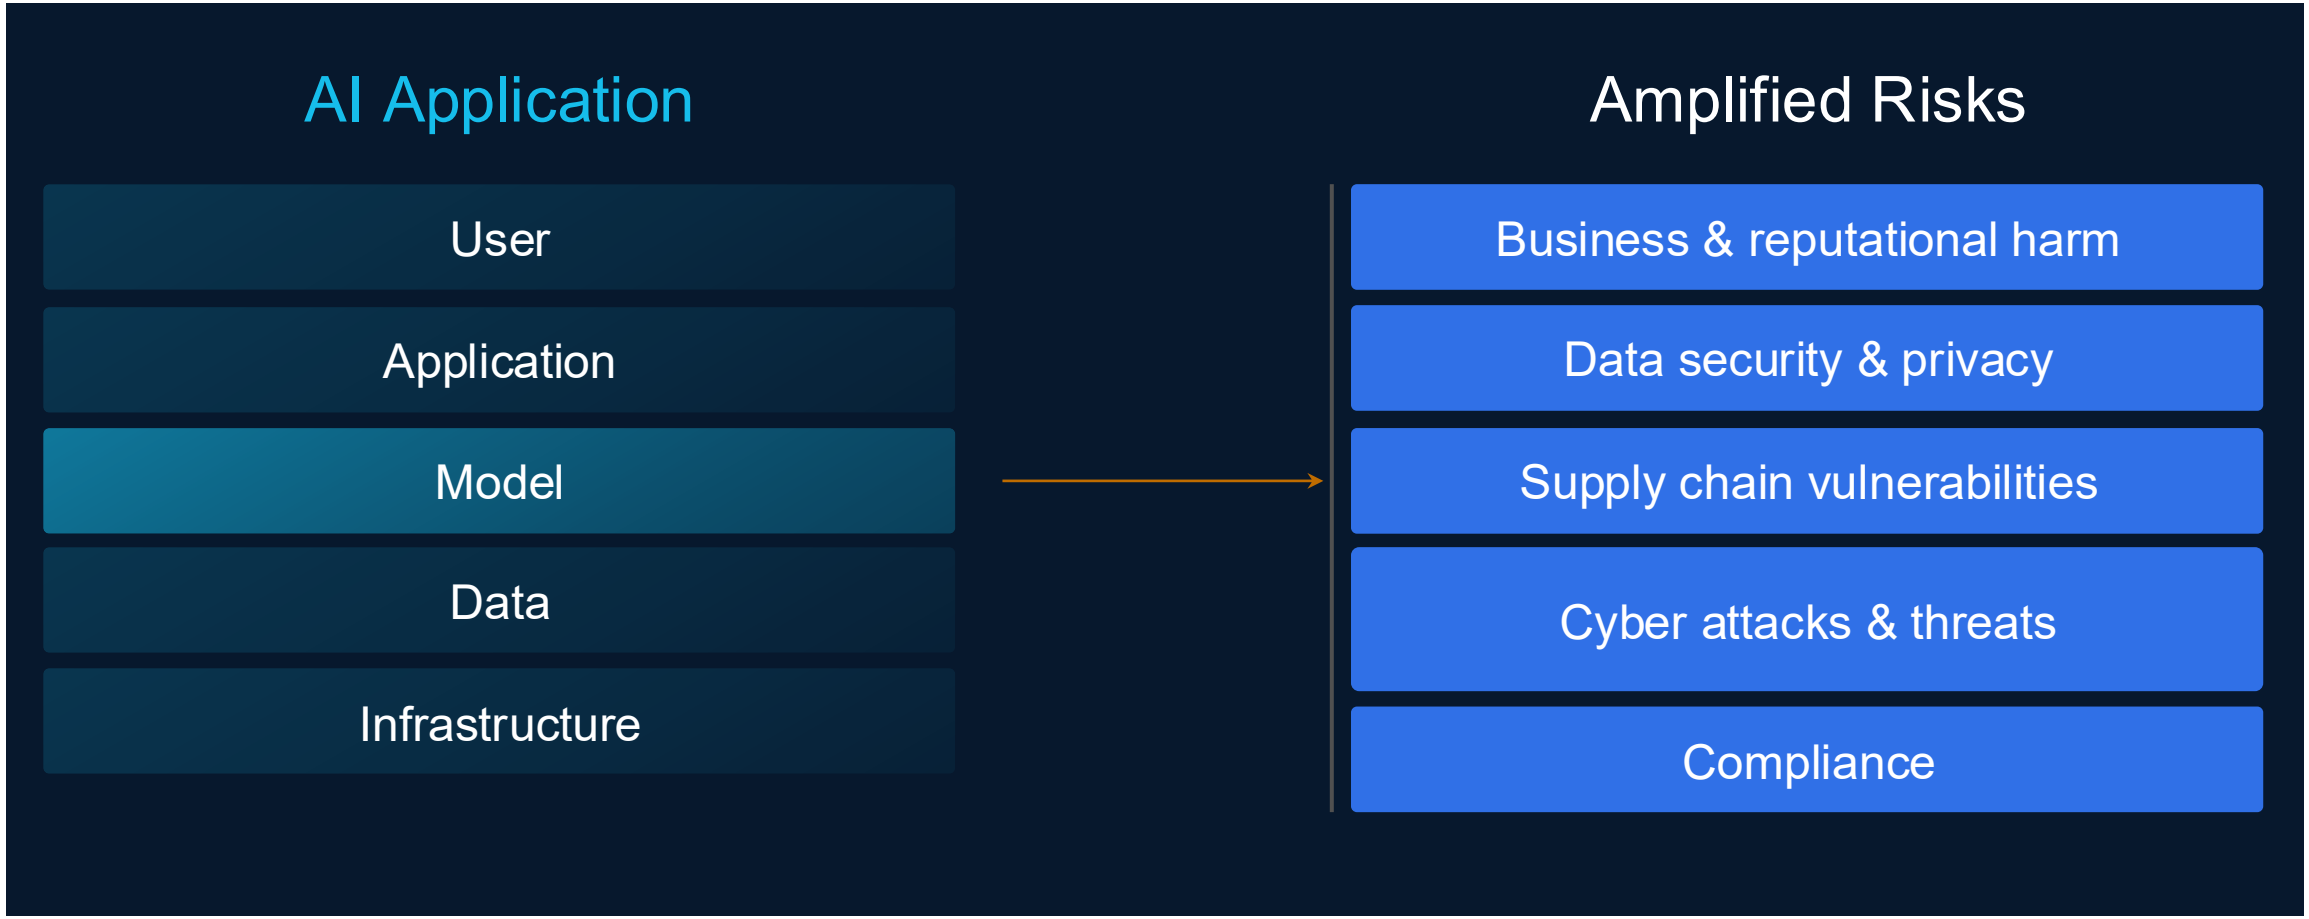

Third-party AI assets carry risks

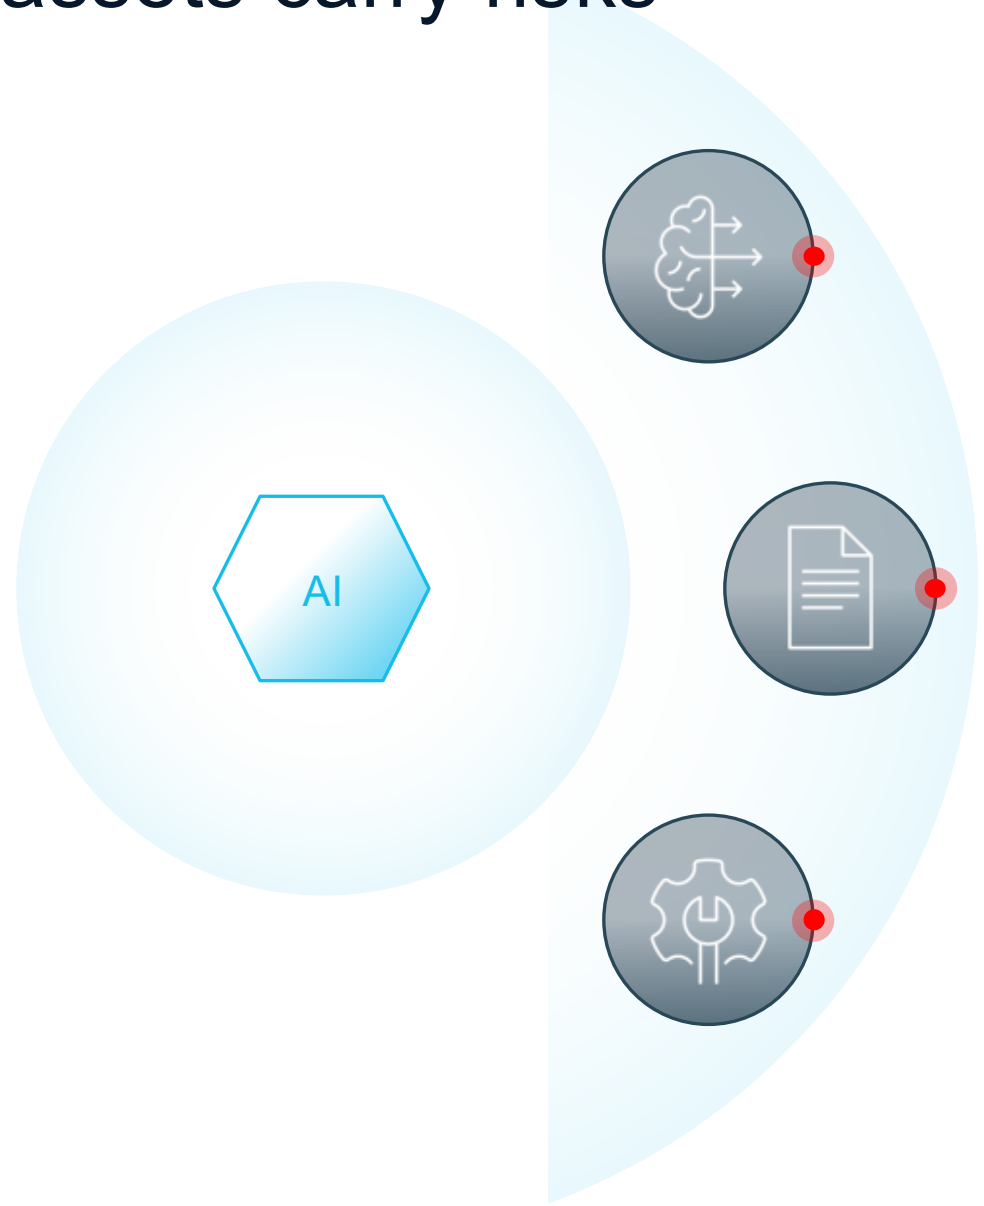

Open-source models
1.9M+ on HuggingFace

Risks: Model backdoors & malware

Third-party datasets
450K+ on HuggingFace

Risks: Data poisoning & privacy violations

MCP servers & tools
Thousand across multiple repos

Risks: Tool & server vulnerabilities

Model security is inconsistent

Built-in guardrails are different for each model, optimized for performance over security, and easily broken when changing the model.

Model security is inconsistent

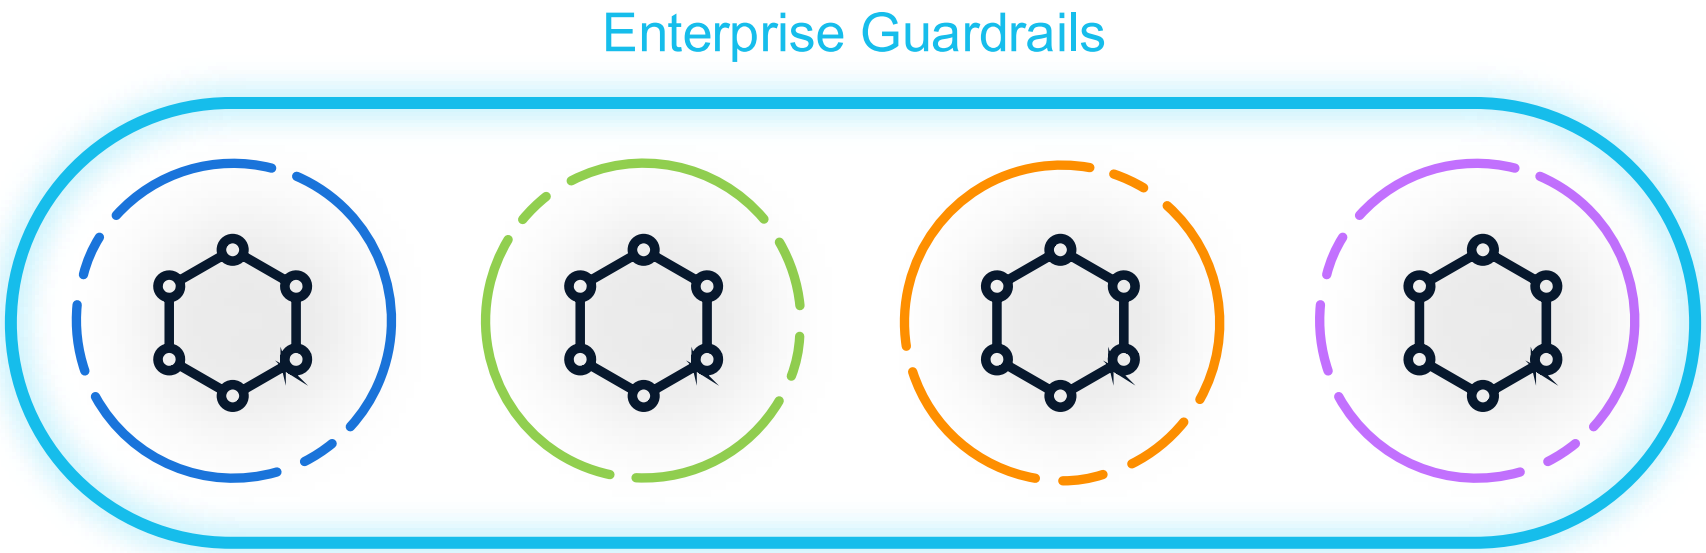

Enterprise guardrails provide a common layer of security across models, allowing AI teams to focus fully on development.

Future: an intelligent layer between agents and resources

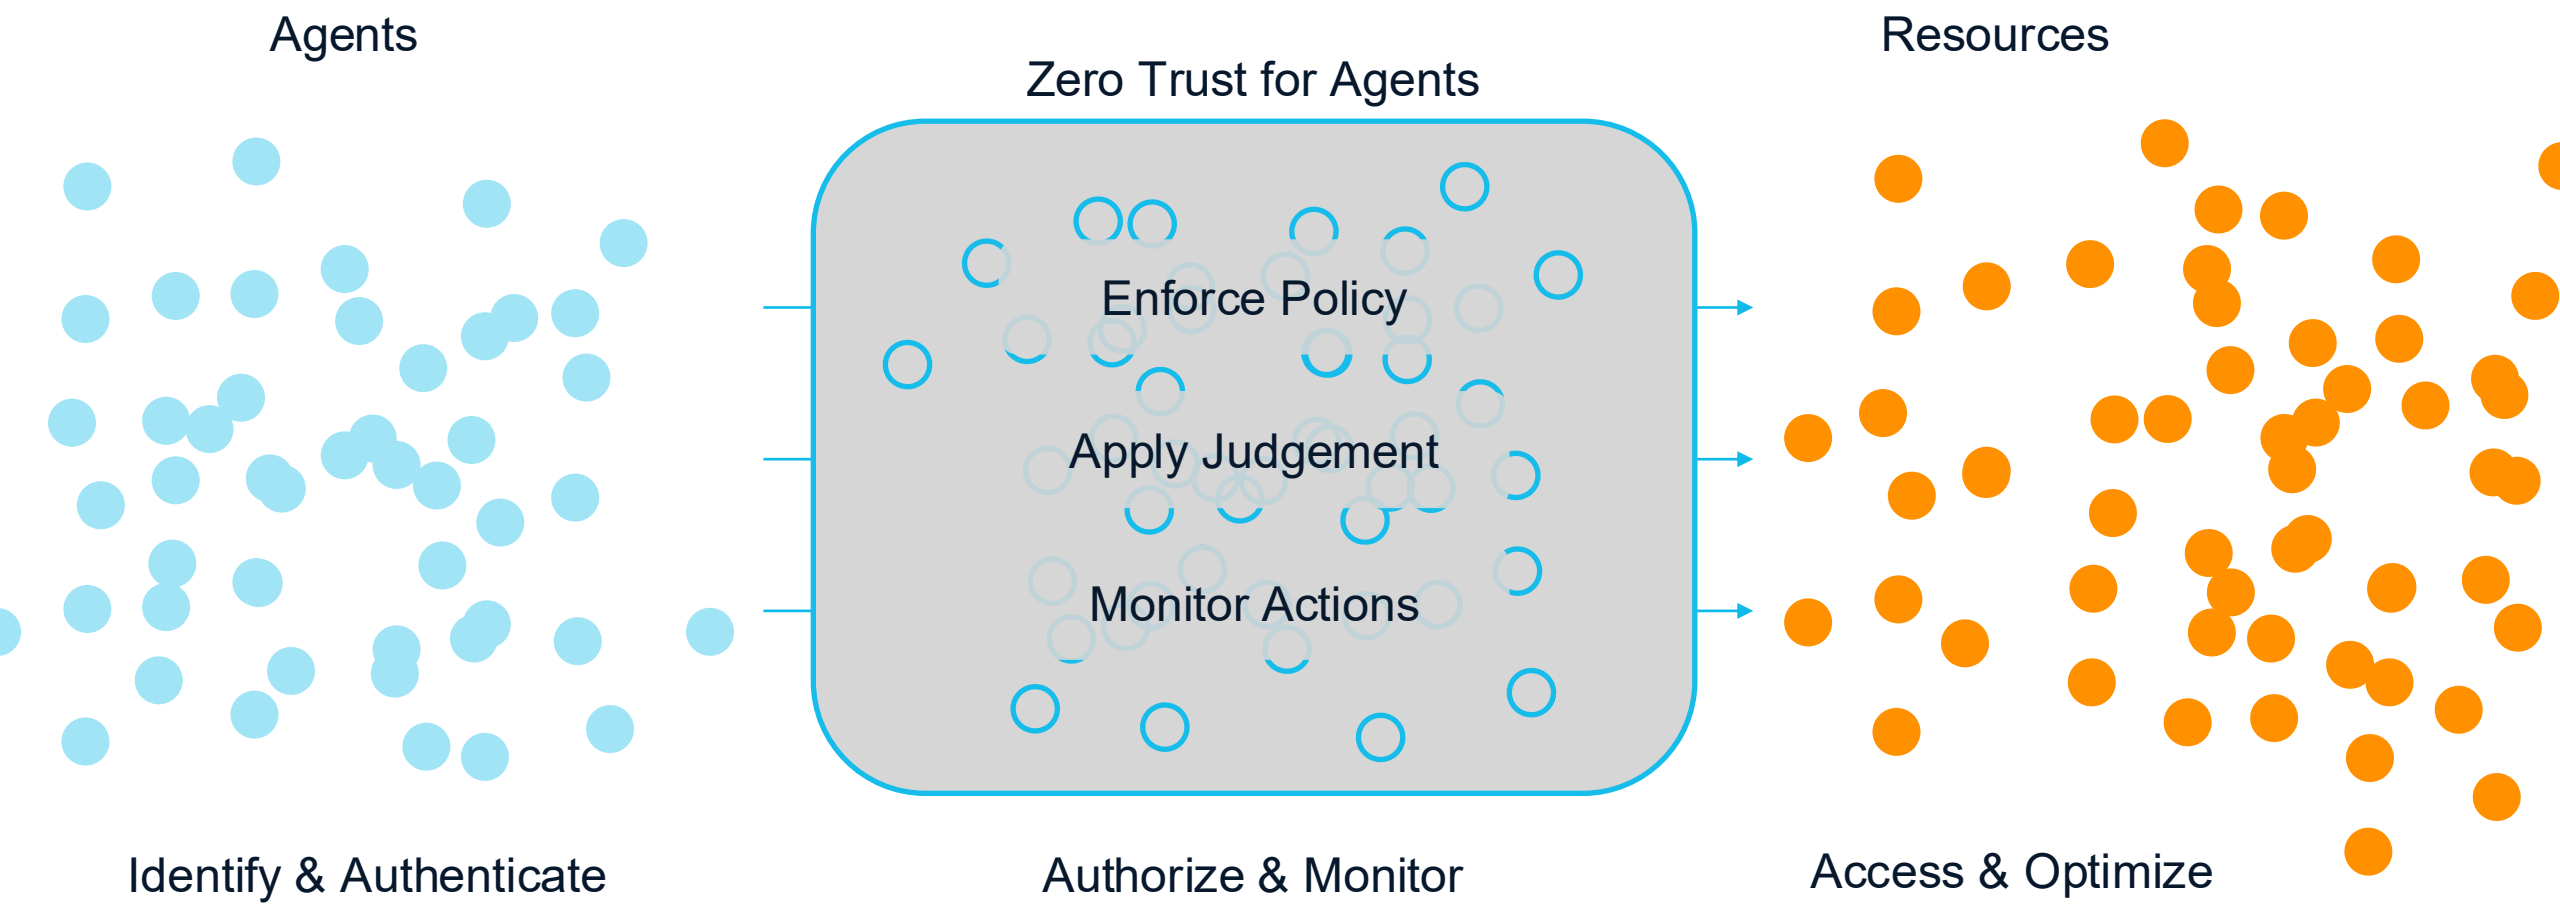

New Standards for AI Security

The New AI Risk Landscape

OWASP

LLM01	Prompt Injection	LLM06 Information	Sensitive Disclosure
LLM02	Insecure Output Handling	LLM07	Insecure Plugin Design
LLM03	Training Data Poisoning	LLM08 Agency	Excessive
LLM04	Model Denial of Service	LLM09	Overreliance
LLM05	Supply Chain Vulnerabilities	LLM10	Model Theft

MITRE ATLAS

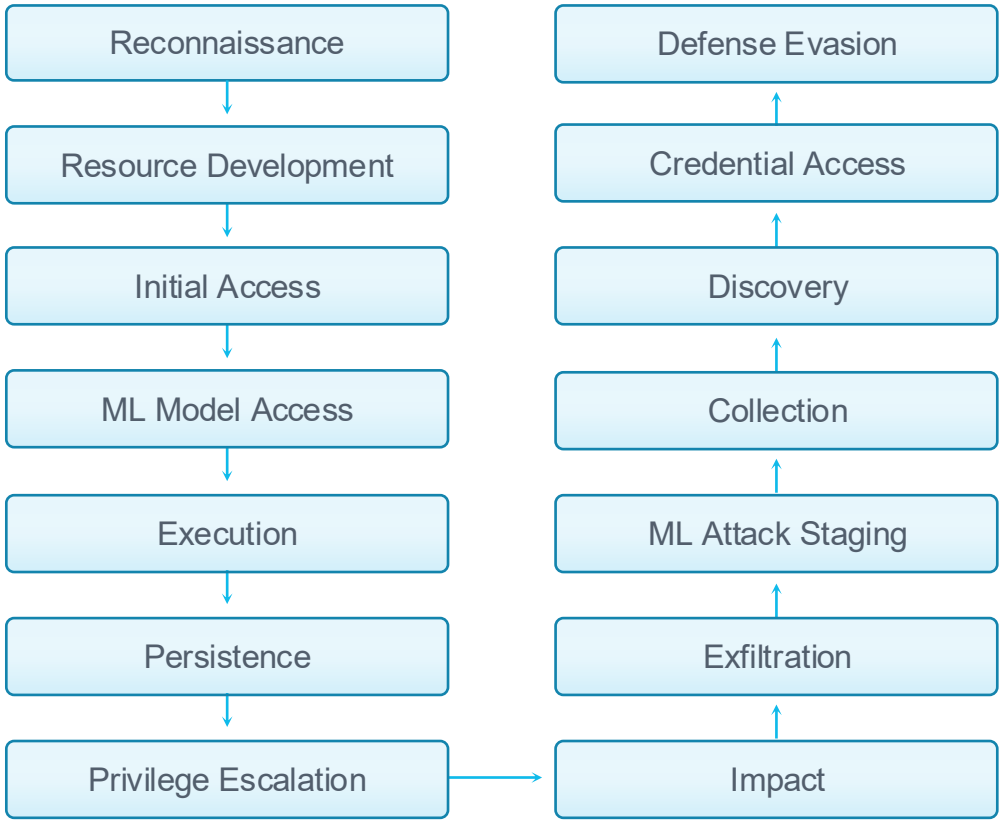

cisco

Cisco AI Security – Shaping AI Security Standards

MITRE ATLAS

- Founding member of MITRE Atlas
- Co-developed the AI Risk Database

- Representing AI Security for National Academies
- Co-organized Hackers on the Hill for Congressional staffers

NSF NIST NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY U.S. DEPARTMENT OF COMMERCE

- Co authored Adversarial AI Taxonomy
- Selected to NIST's AI Safety Institute

- Creating Prompt Injection Taxonomy w/ UK AI Security Institute
- AI Security Hackathon at The National Cyber Security Centre

OWASP

- Contributors to OWASP Top 10 for LLMs
- Selected as review panelist for Agentic Security initiative

- Representing AI Safety at APEC Summit in front of South Korean President, Japanese PM
- Partnering with the Japanese Government on AI safety proposal

cisco

Cisco's integrated AI security and safety framework

We've started to integrating it into the product – starting with AI Validation

20+ Objectives The motive or goal behind an attack	150+ Techniques & Sub-Techniques A granular understanding of the threats including actions, methods, and variations		5+ Mappings References to common AI and governance frameworks
Goal Hijacking	Direct Prompt Injection	Instruction Manipulation	OWASP: AAI003:2025, MITRE: AML.T0051.000…
		Obfuscation	OWASP: AAI003:2025, MITRE: AML.T0051.000…
	Multi-Modal Injection Manipulation	Image-Text Injection	OWASP: AAI001:2025, NIST: AML.018…
		Audio Command Injection	OWASP: AAI001:2025, NIST: AML.018…
		Video Overlay Manipulation	OWASP: AAI001:2025, NIST: AML.018…
Data Privacy Violation	Data Exfiltration / Exposure	Training Data Exposure	OWASP: AAI015:2025, MITRE: AML.T0024…
		Data Exfiltration via Agent Tooling	OWASP: AAI015:2025, MITRE: AML.T0086…

AI Defense: coverage across the AI lifecycle

Discovery

AI Cloud Visibility

Identify AI assets

Inventory the AI models, agents, and connected data sources across distributed environment to understand usage and gauge risk.

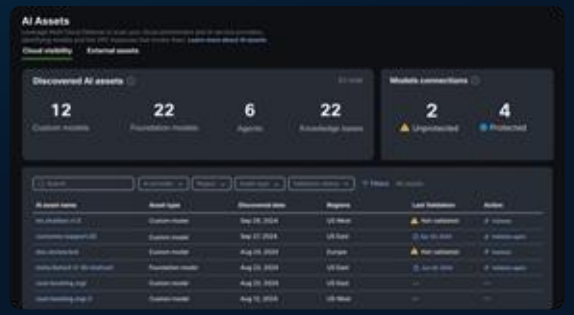

Detection

AI Supply Chain Risk Management

Scan for threats

Scan model files, repos, and MCP servers to proactively block malicious or unsafe AI assets before operations are impacted.

AI Model & App Validation

Detect the vulnerabilities

Identify safety and security vulnerabilities across models at scale with algorithmic red teaming technology.

Protection

AI Runtime Protection

Mitigate threats in real time

Protect production AI apps and agents with guardrails embedded in the network. Block attacks and harmful responses in real time.

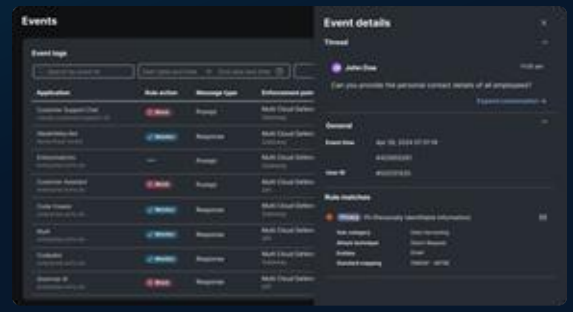

The AI Defense Solution

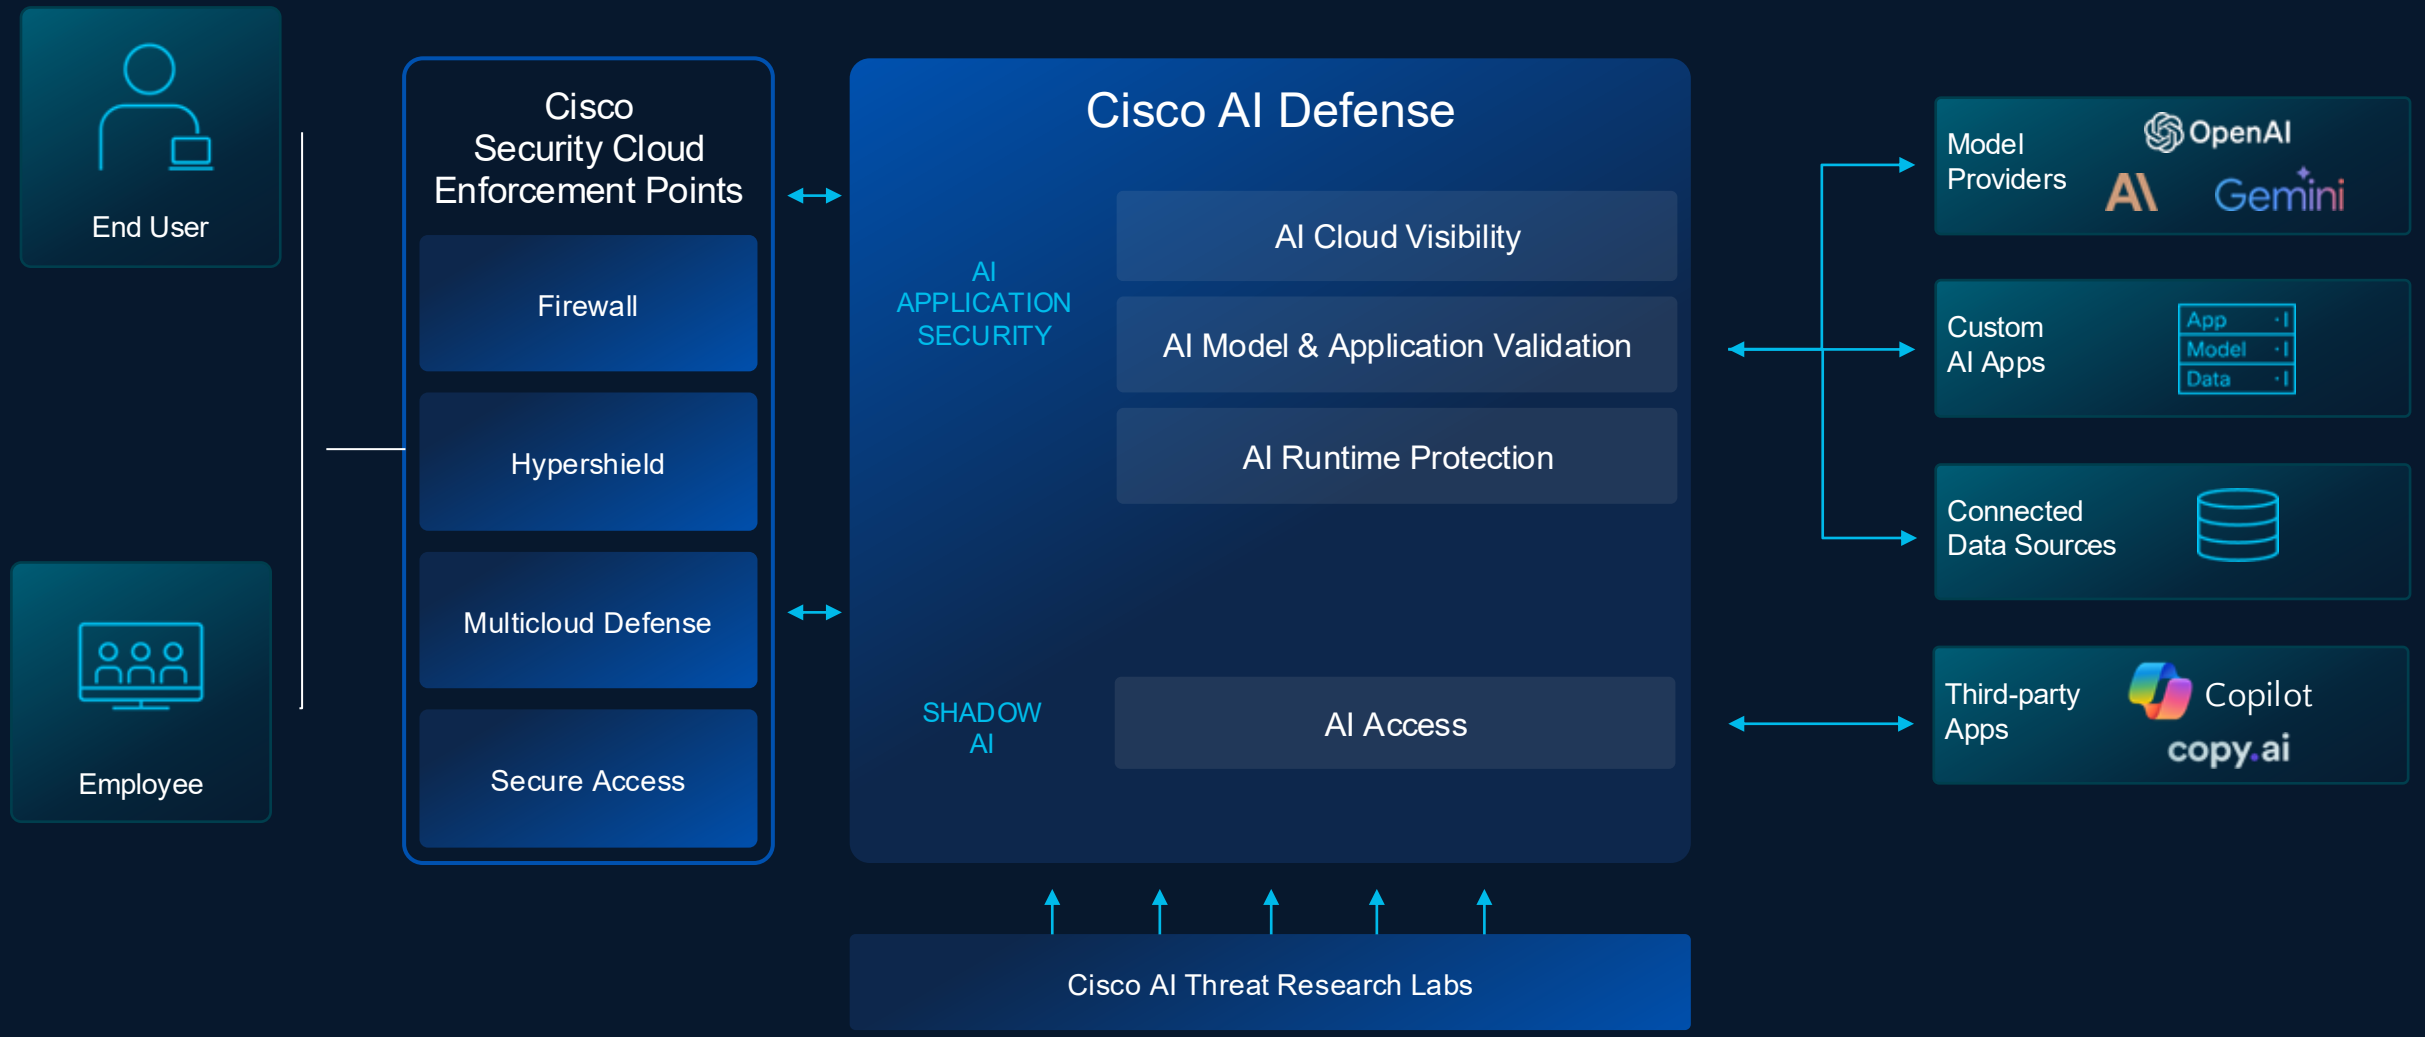

Discovery: AI Cloud Visibility

- Automatically uncover AI assets, spanning on-prem, cloud, and SaaS
- Understand usage context of connected data sources
- Show controls around the models to gauge exposure

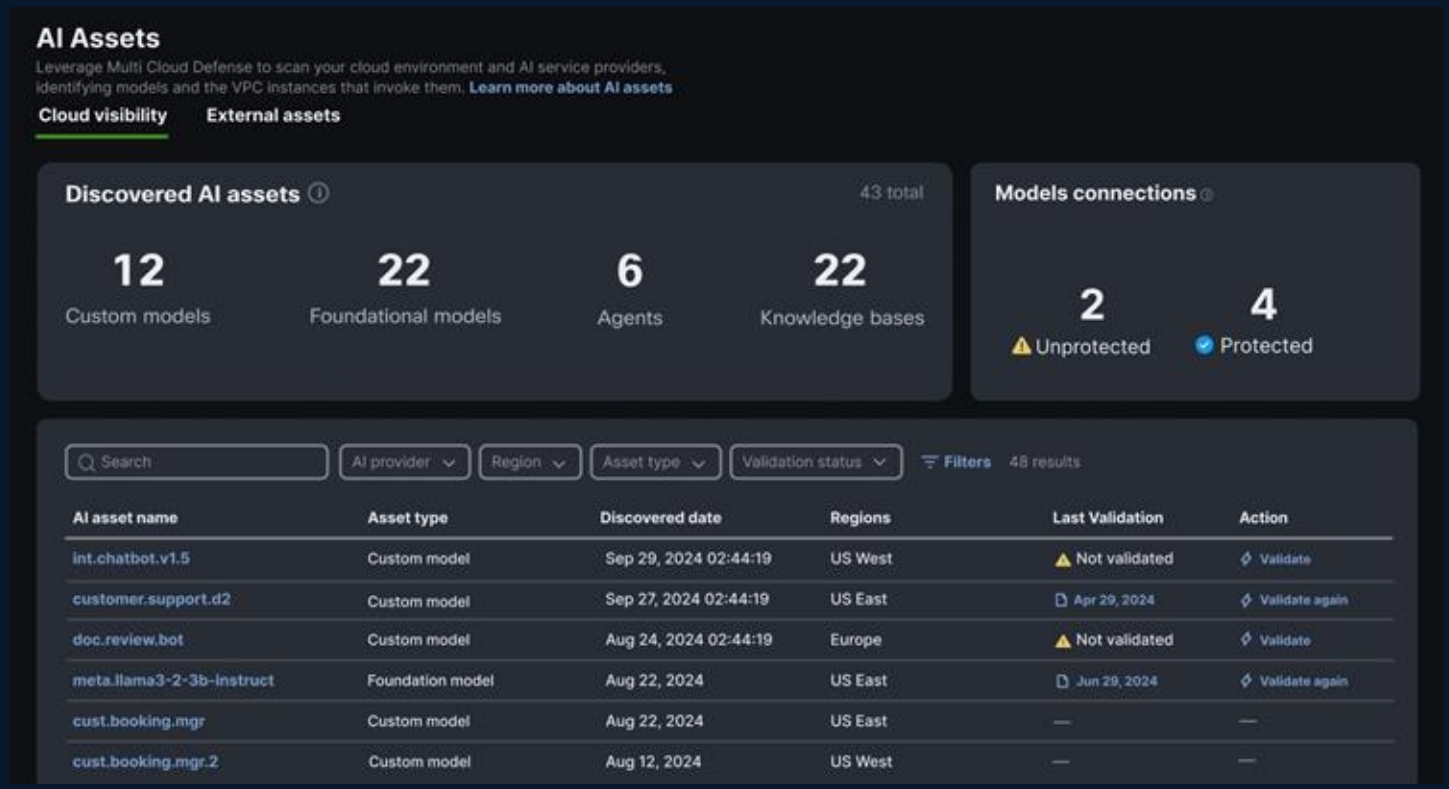

Cisco AI Defense

Complete Solution for building AI Apps

1

Platform Advantage

Security at the network layer

- Network-level data insights provide full visibility into AI traffic and associated risks
- Fast, low-friction deployment that does not modify the app
- Enforce policies across and within clouds and datacenters

2

AI Model & App Validation

Algorithmic AI red teaming

- Automated assessment of safety and security vulnerabilities
- AI readiness guides bespoke guardrail and enforcement policy
- Automatic integration into CI/CD workflows for seamless, continuous testing

3

Proprietary Model & Data

Purpose-built for AI security

- Team pioneered breakthroughs from algorithmic jailbreaking to the industry's first AI Firewall
- Contribute to (and align with) NIST, MITRE, and OWASP
- Leverage threat intelligence data from Cisco Talos & Cisco AI security research teams

cisco

Thank you