# From Complexity to Clarity

Optimising Data Centers for the Future


Chris Gascoigne
Principal Solutions Engineer

September 2025


There is a lot happening *to* IT at the moment


Stay and optimise?
Move to a new hypervisor?
How to avoid the same trap?

# Significant Broadcom Changes

Organisations already seeing 2x-6x increases at renewal

Ending perpetual licenses

Transition to subscription licensing

Move from CPU to Core based licenses

Reduce portfolio to 4 packages*, enforced VCF for “strategic” customers

# Platform Modernisation Options

STAY

MOVE

## OPTION 1

vmware

Modern 3-Tier Architecture (CI)

Refresh to reduce licensing costs

PURESTORAGE | NetApp | HITACHI

## OPTION 2

vmware

Adopt Hyperconverged Infrastructure (HCI)

vmware vSAN | NUTANIX

## OPTION 3

Alternative Approach

NUTANIX

RED HAT OPENSHIFT

Azure Stack HCI

Bare Metal

## OPTION 4

Public Cloud

NUTANIX

RED HAT OPENSHIFT

Cloud Native


How to do this without significant operations cost increases?

# Common infrastructure platforms are essential

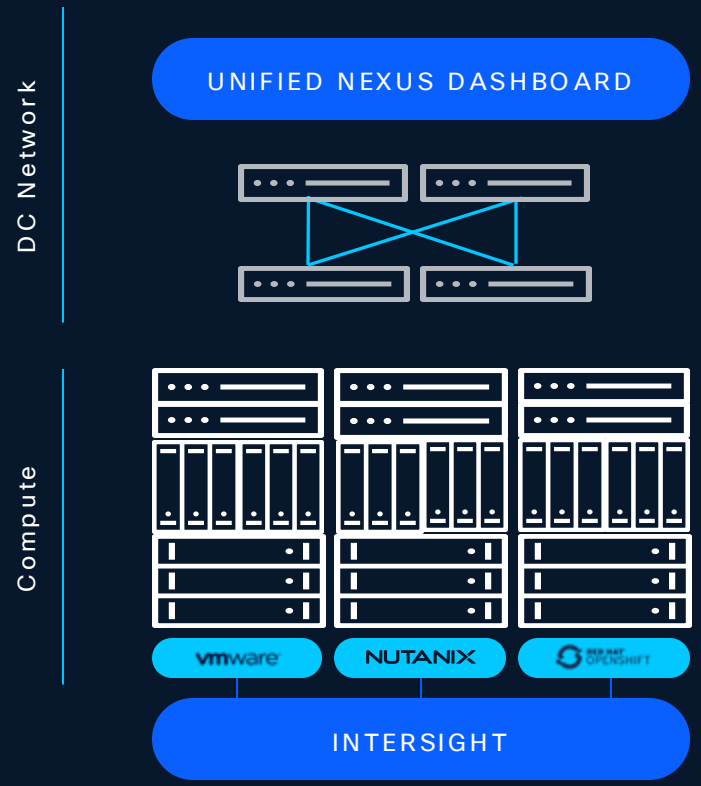


Consistent, policy driven management across compute and network for all workloads

Simplify management and troubleshooting with less silos

Scale hypervisor domains up/down through software

# Blurring the line between rack and blade


## The best of both worlds on our most sustainable modular server yet

Modular, sustainable, easily upgradeable solution

Higher performance, smaller footprint

Powers modern apps and traditional workloads

Fabric-based architecture with unified fabric connectivity

Up to: 2PB raw storage, 8 x H100 NVL GPU, 1376 cores, 64TB RAM per chassis
* With UCS M8 blades

Sources: Energystar.gov / EPEAT.net

# Simpler, smarter, more agile computing

- Flexible SaaS or on-prem management
- Across Data Center, AI, and Edge
- Every feature available via APIs

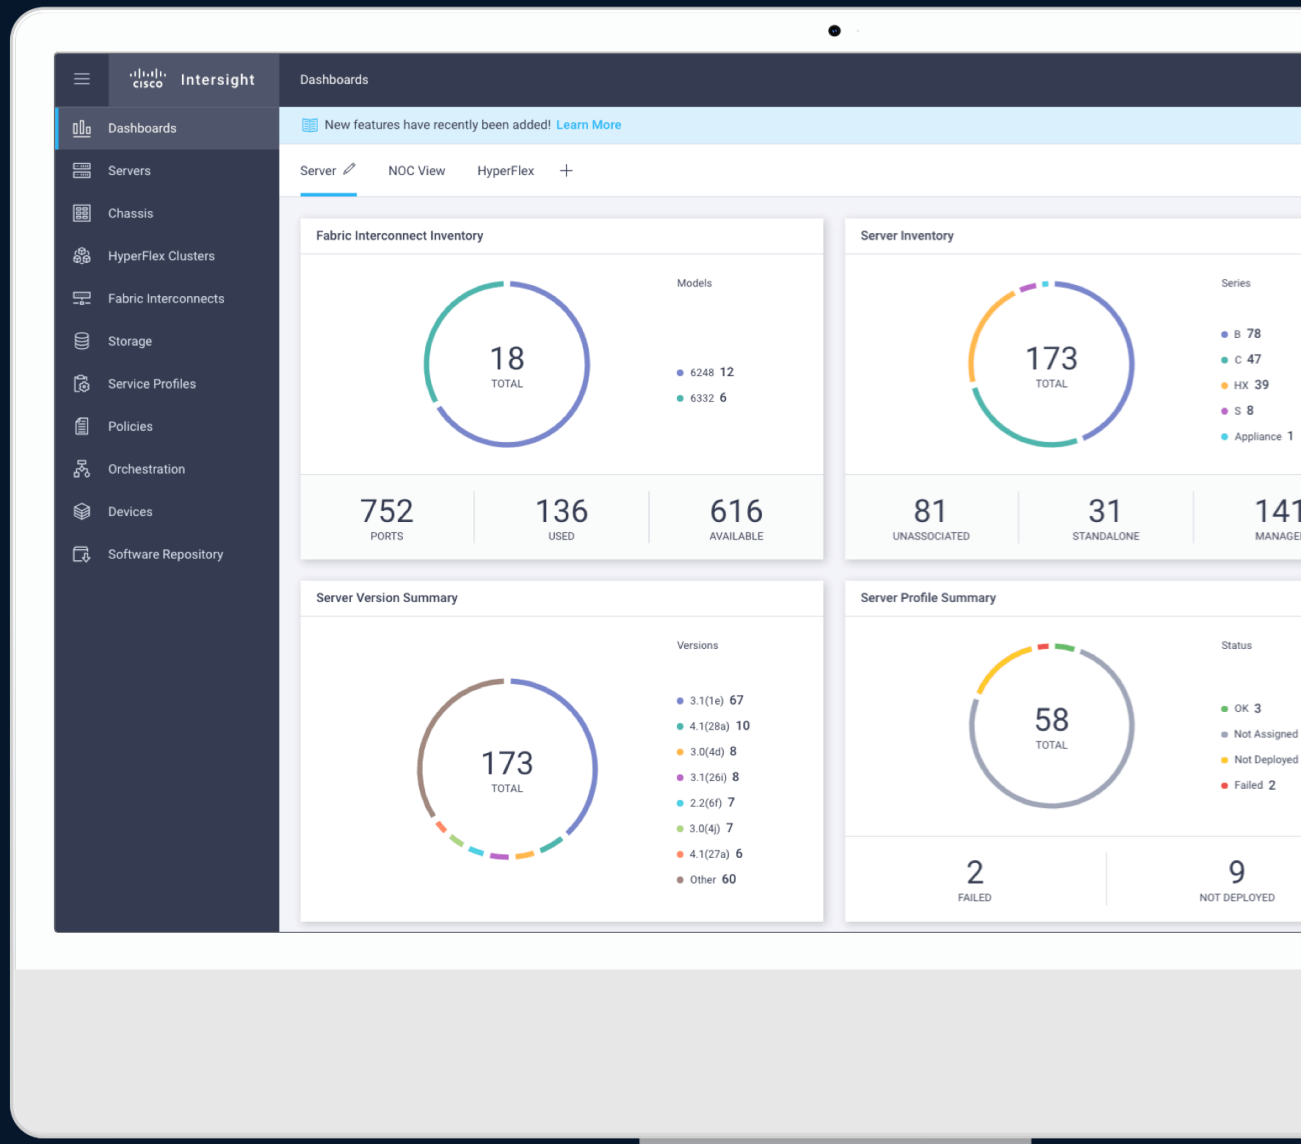


How to apply consistent security policy in heterogeneous environments?
How to quickly respond to threats without risking availability?
Can I reduce complexity while improving performance?

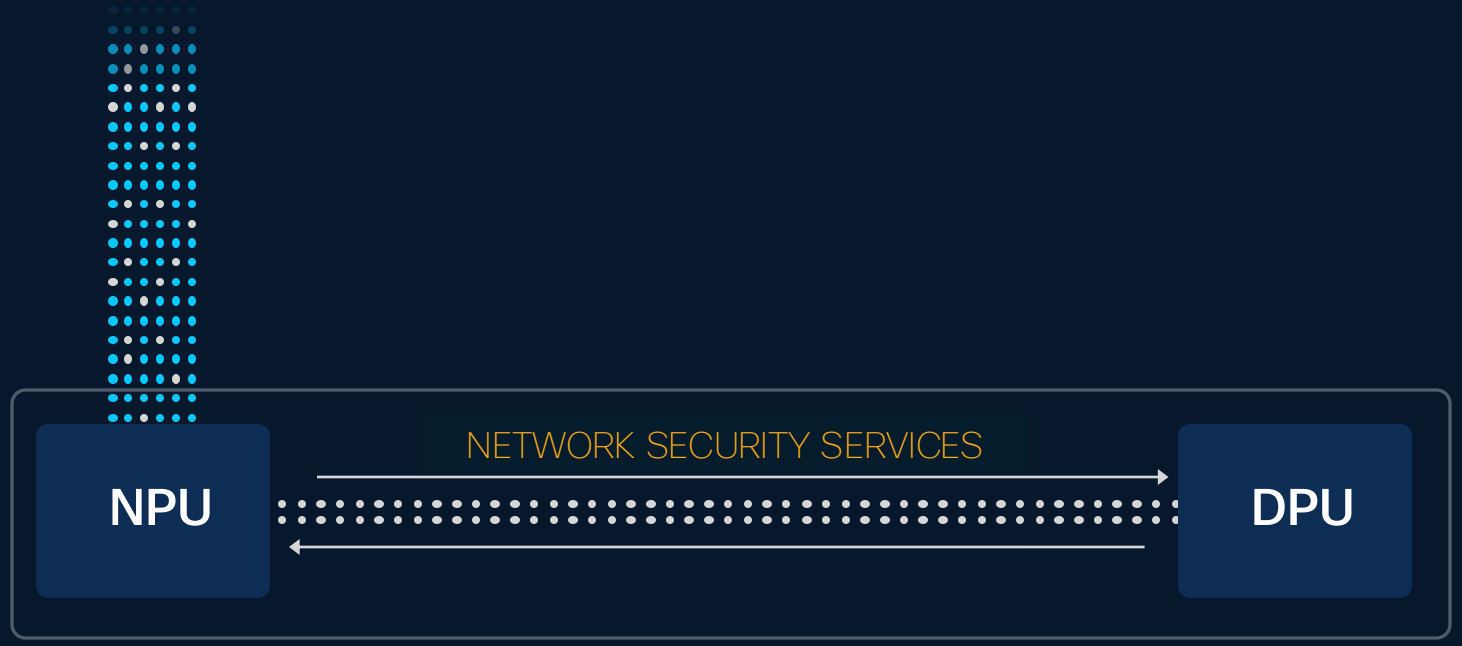


N9300 Smart Switch

# Security infused into the **data center fabric**

# Separate lifecycle management

# Hypershield: Distributed Segmentation Architecture

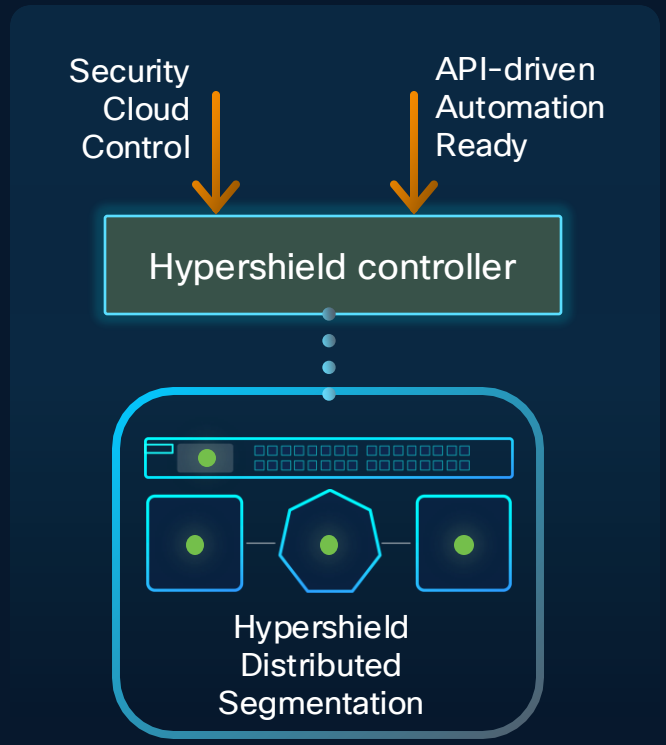


## Global Control

Unified visibility and global policy across agents and smart switch

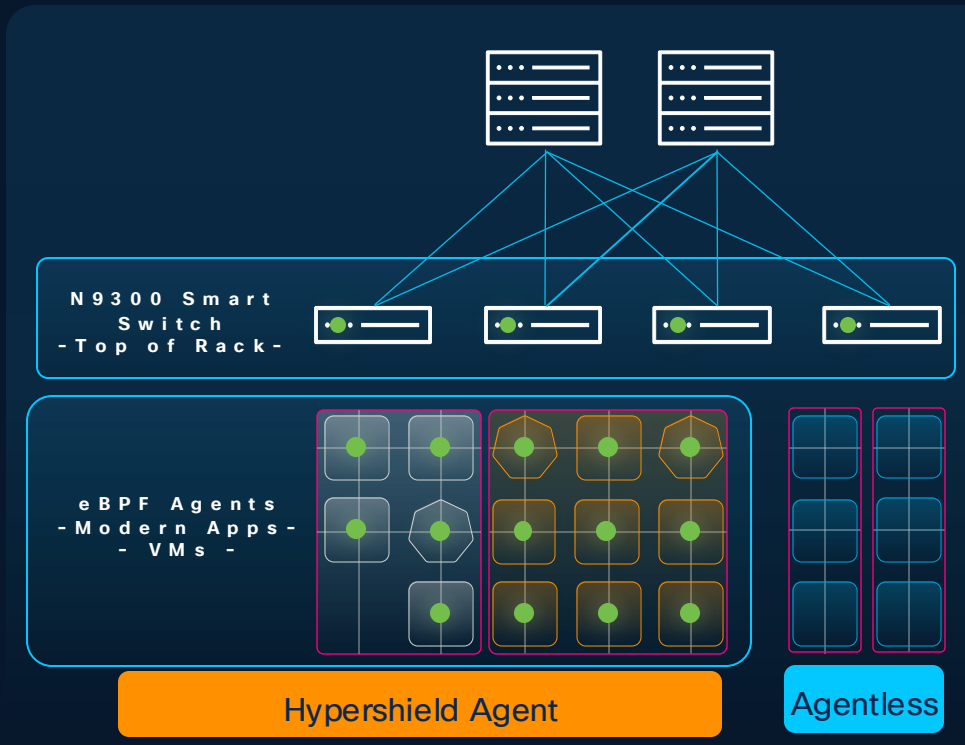


## Distributed

Eliminate blind spots with fully distributed enforcement
– Inline everywhere, kernel and network fabric –

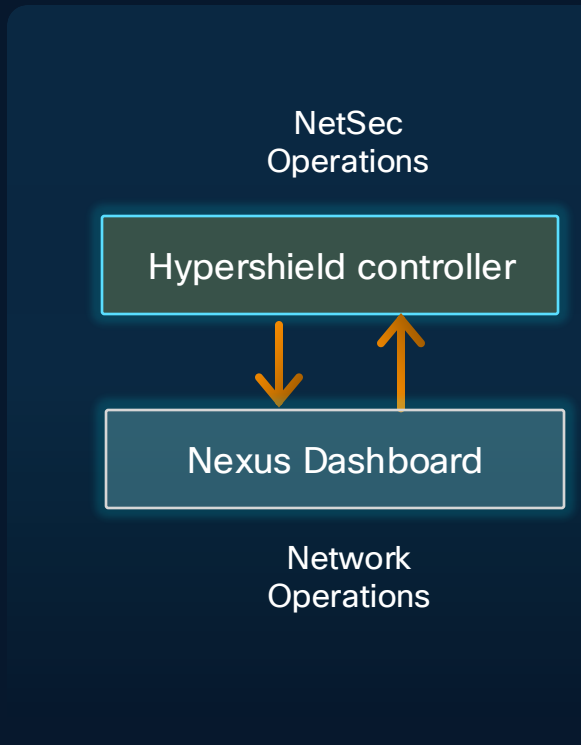


## Unified operations

Separation of control with combined operations and troubleshooting

# Common infrastructure platforms are essential...

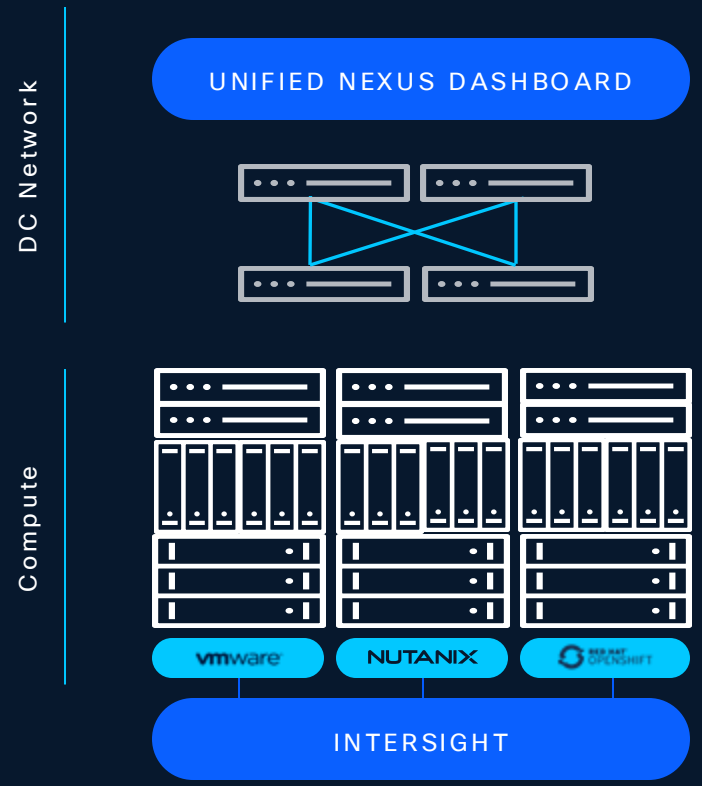

# ... but as you scale tacking on security is complex ...

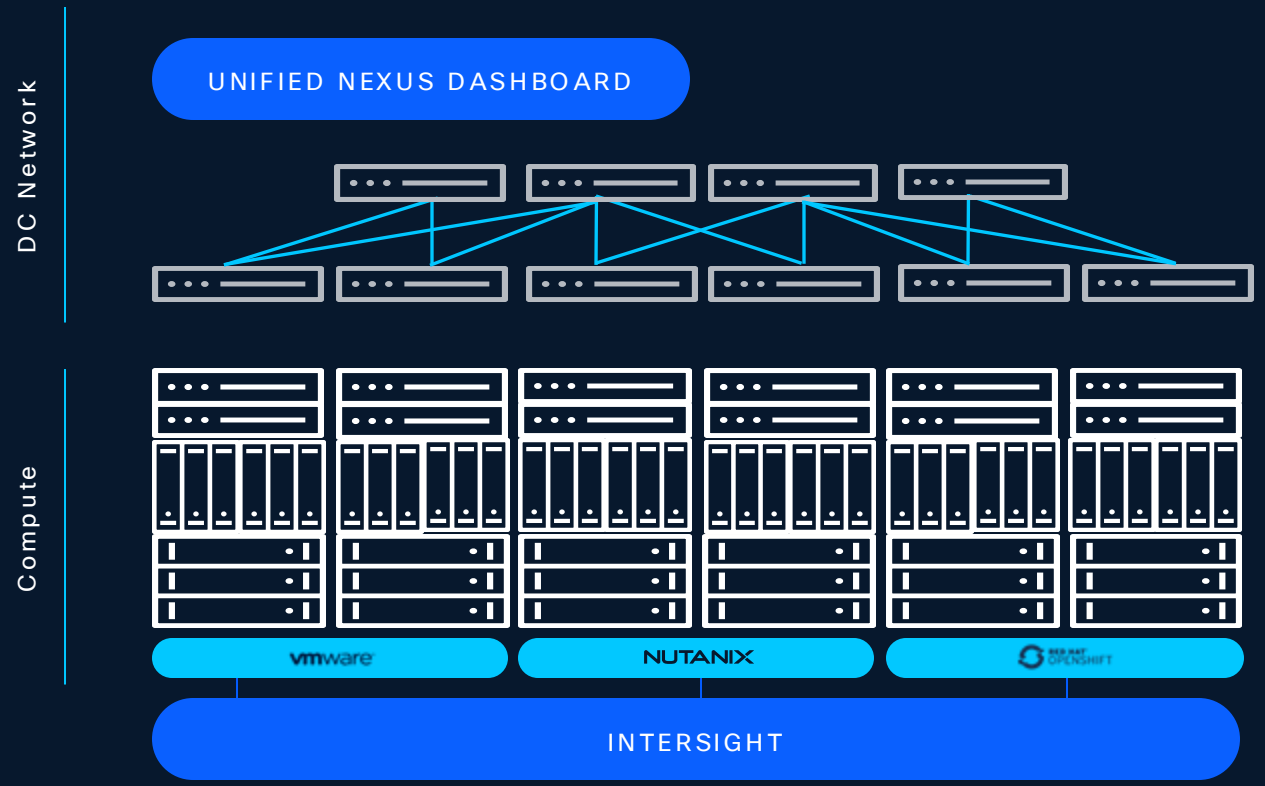

# ... embed security in the network fabric with Hypershield

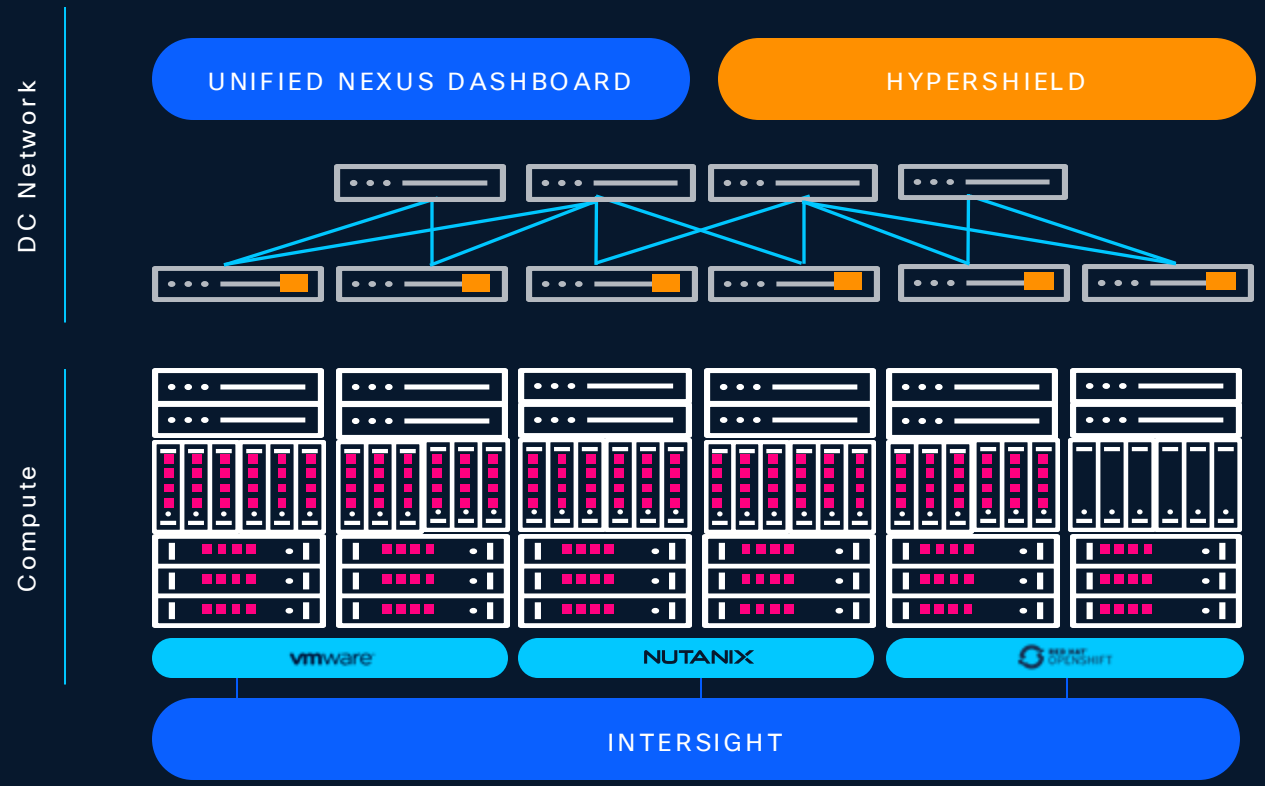


Security DPU
Security Agent

# Patching is hard

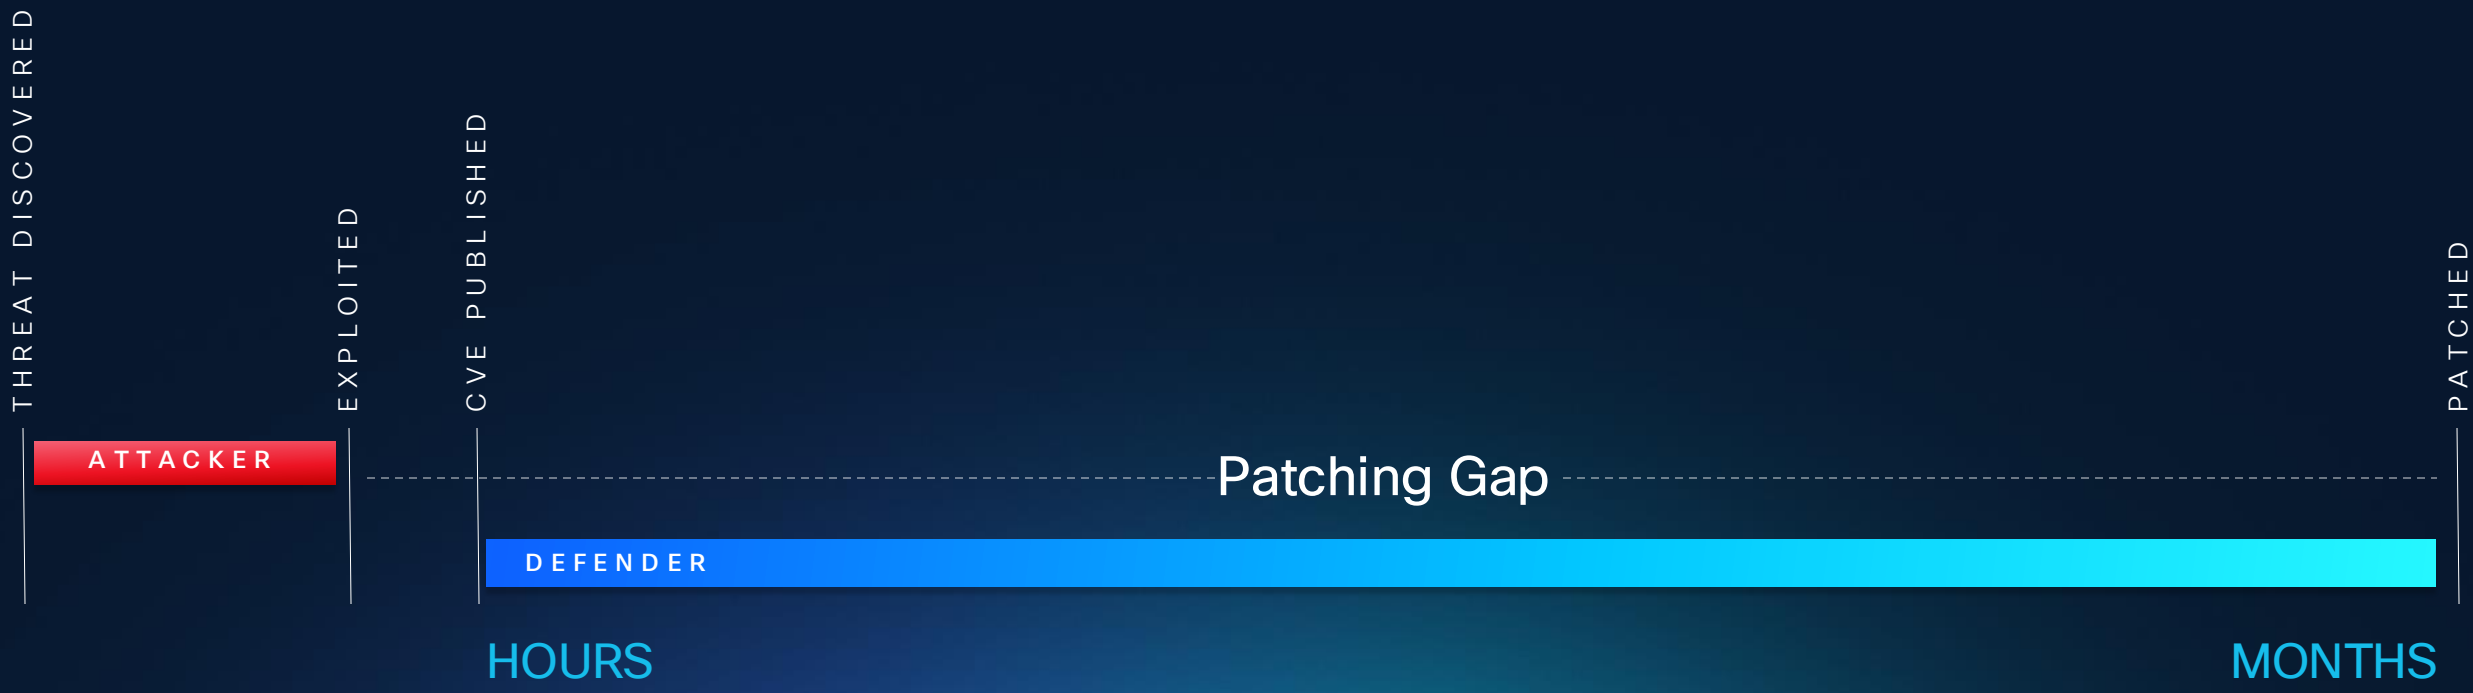

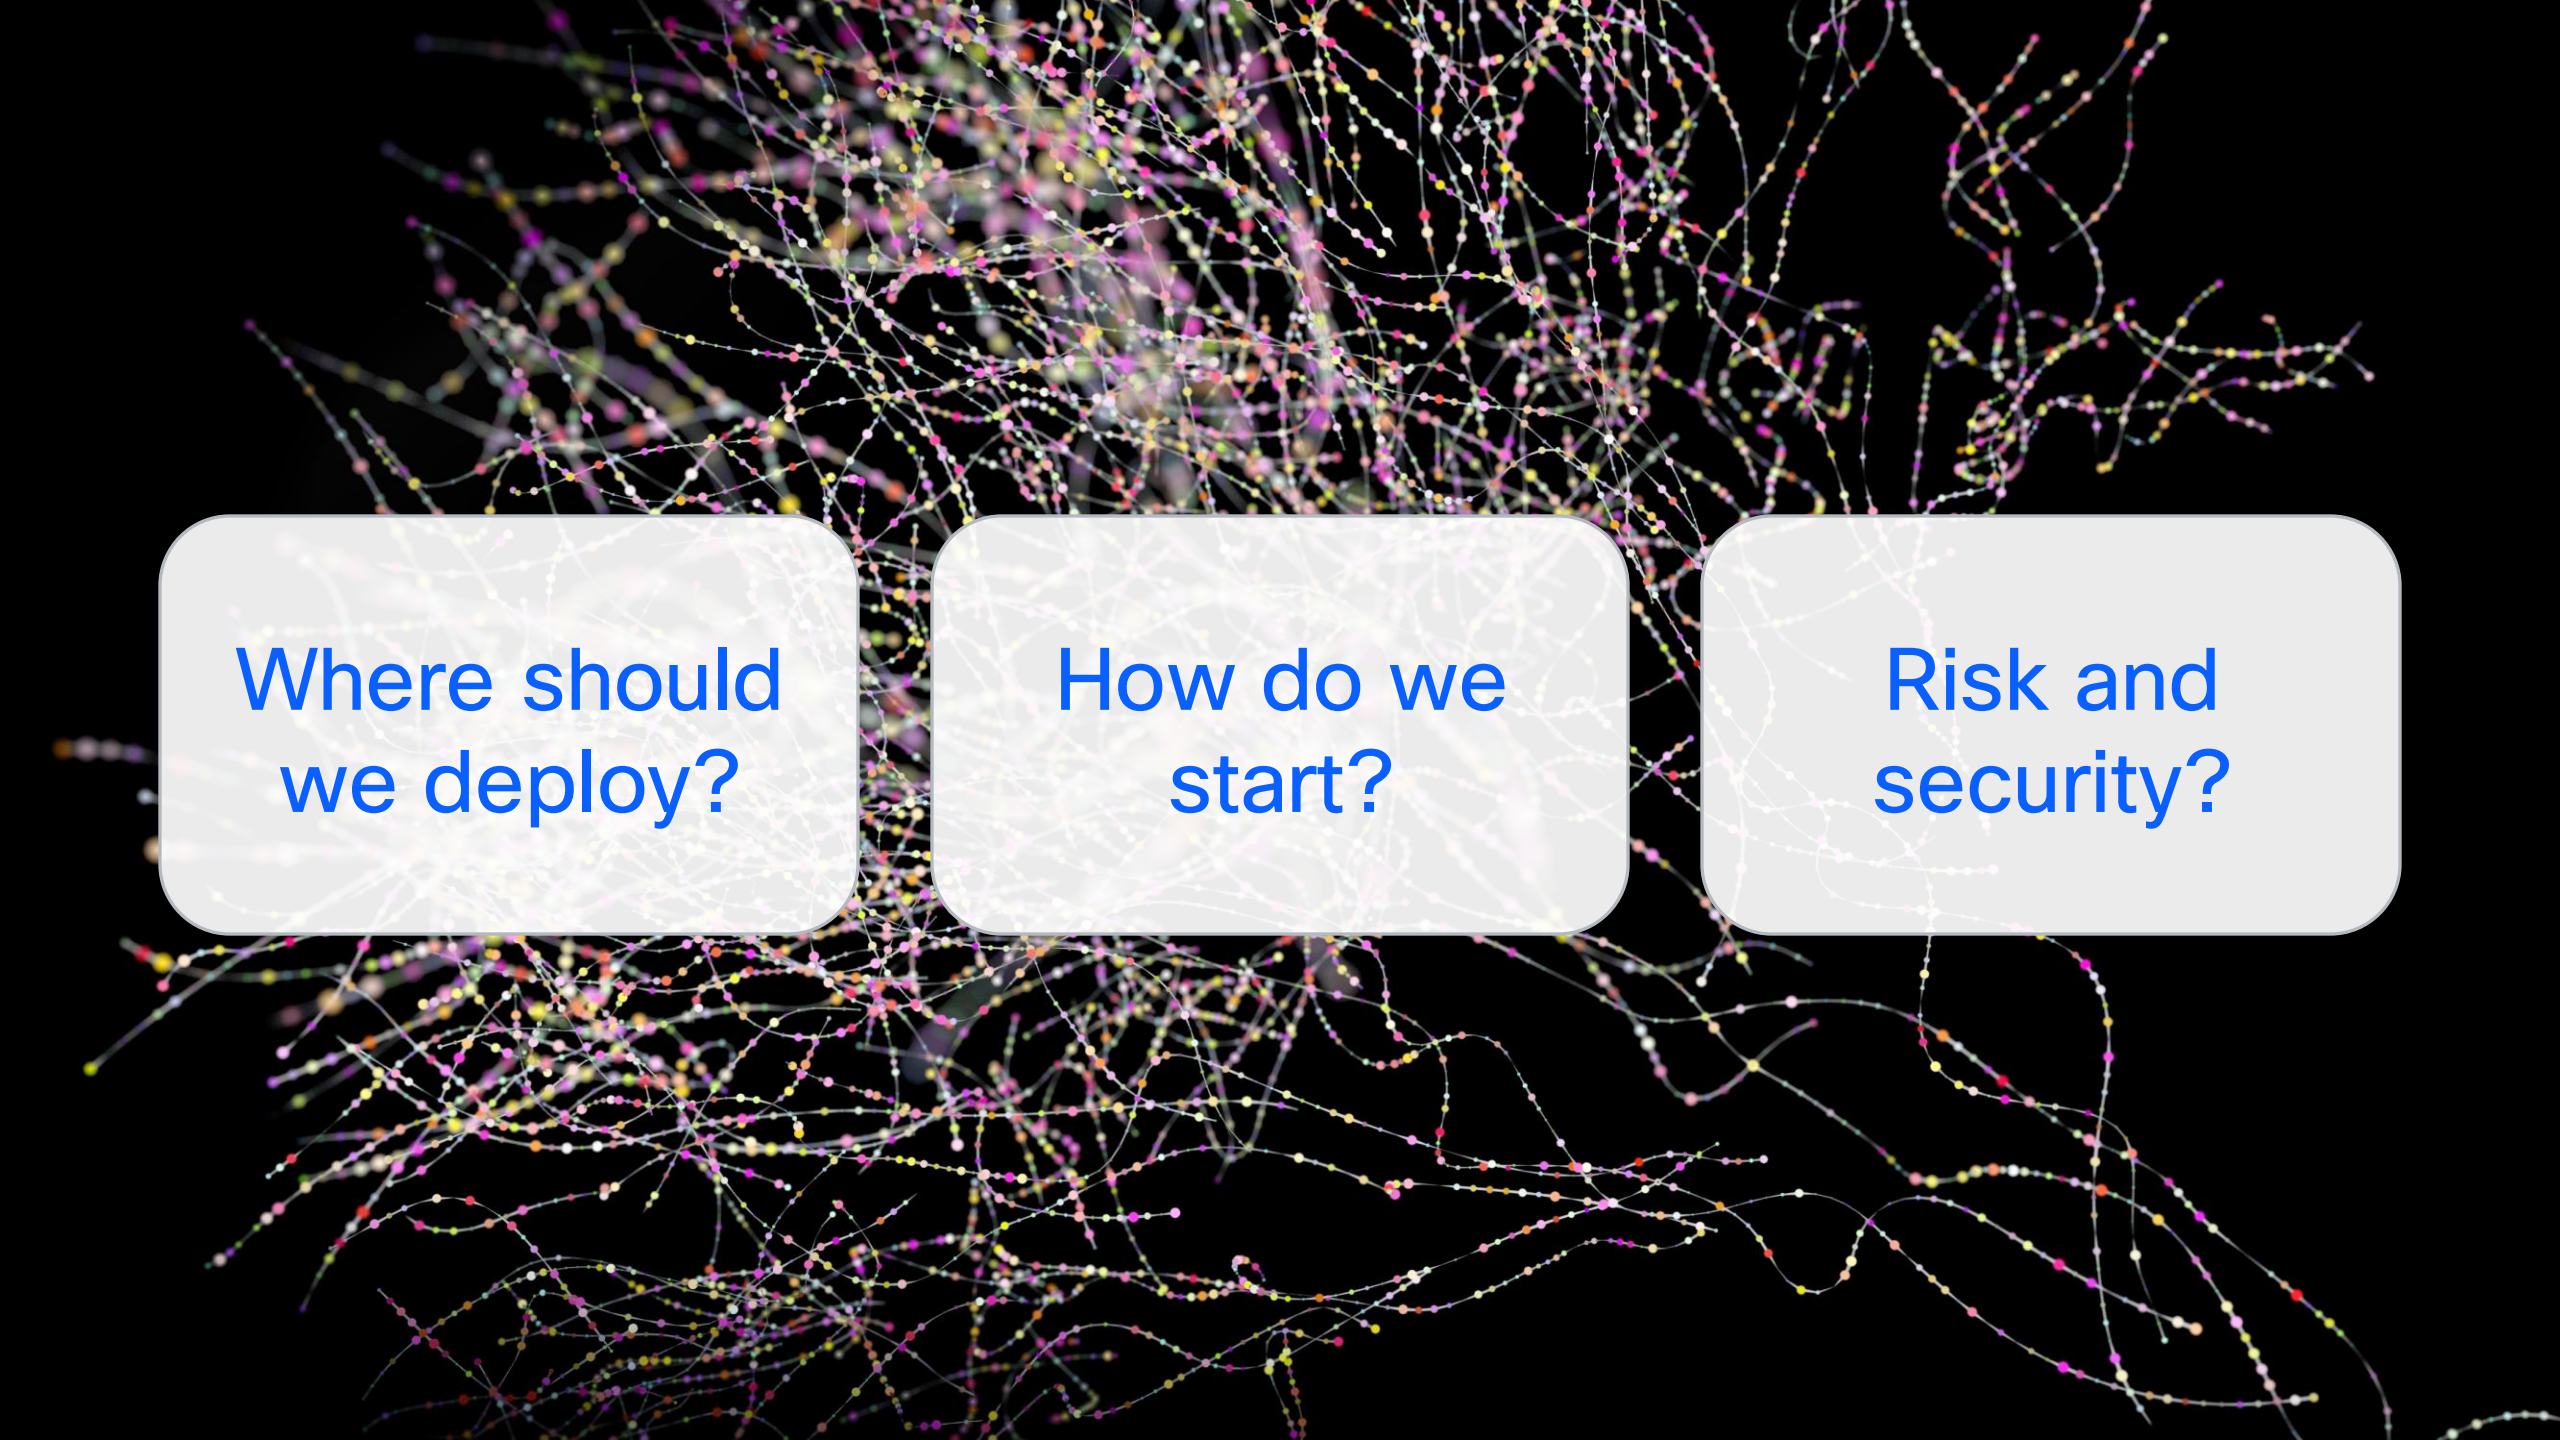
Where should we deploy?
How do we start?
Risk and security?

# Cisco Secure AI Factory with NVIDIA

Delivering Trusted AI Outcomes

AI Practitioners

A reference design with validated architectures to accelerate AI adoption for enterprises with integrated AI infrastructure and software solutions

IT Infrastructure & Operations

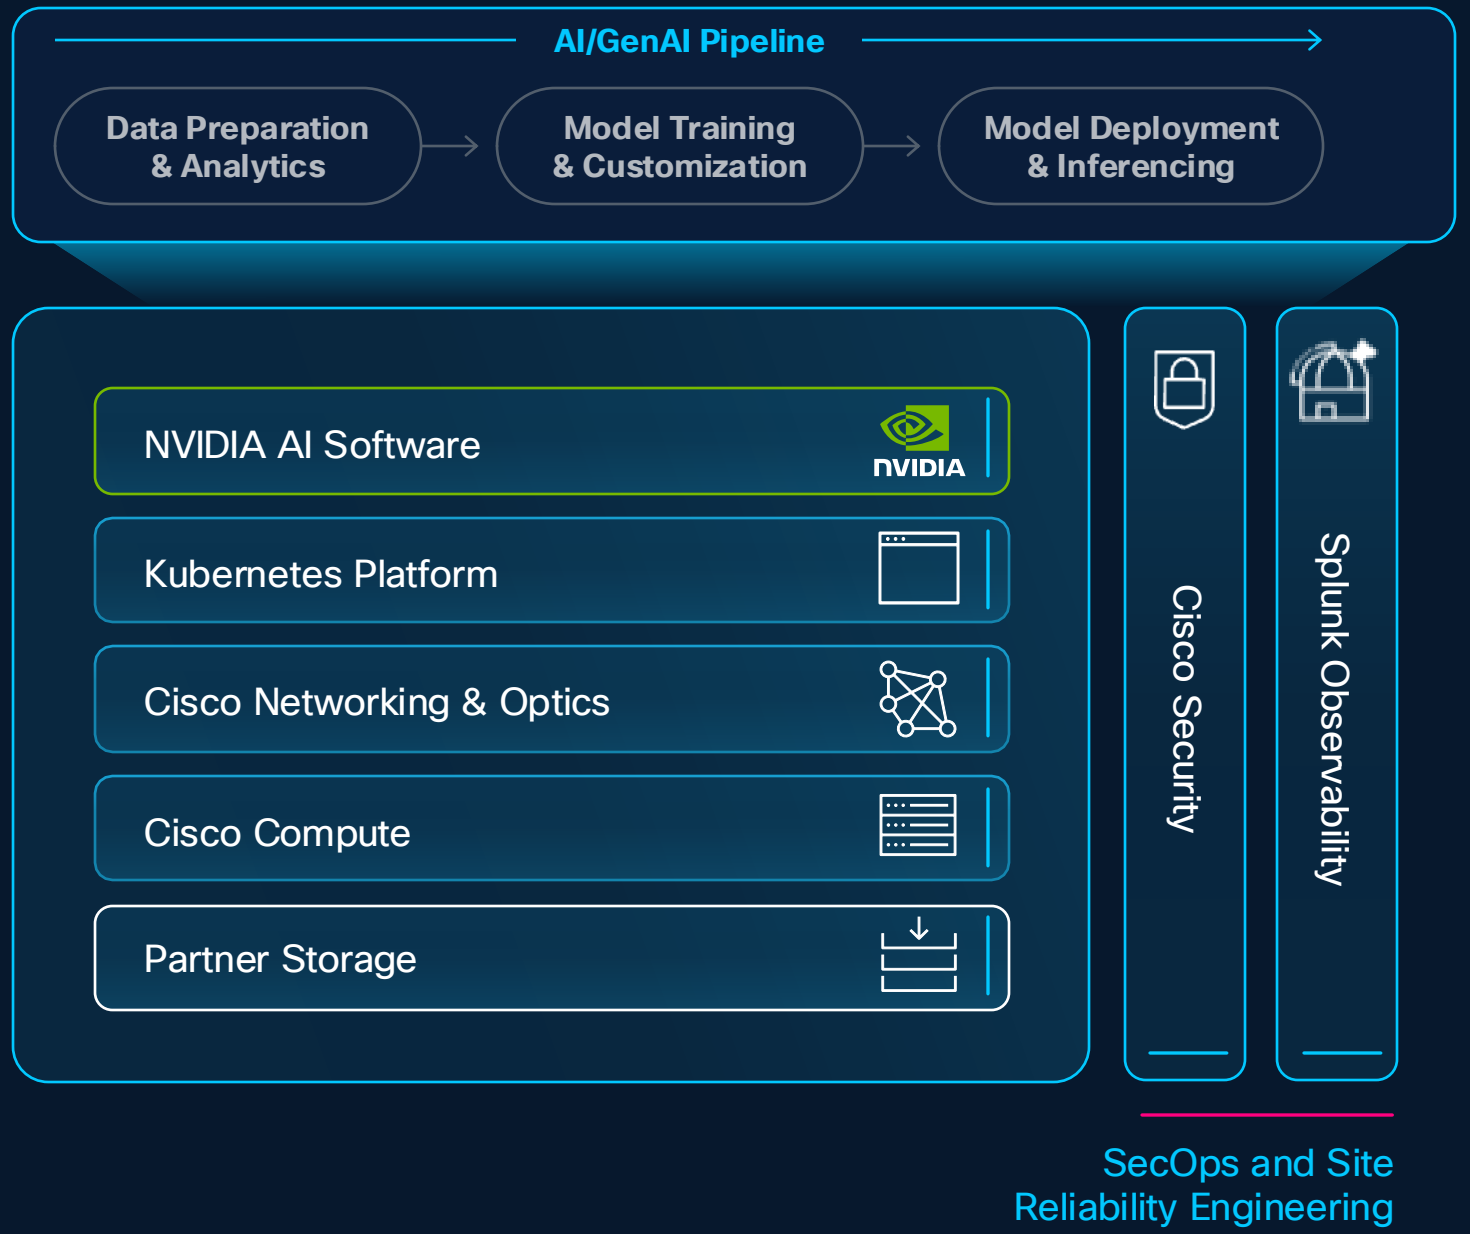


SecOps and Site Reliability Engineering

Infrastructure

# Cisco UCS Compute Portfolio

## Blade

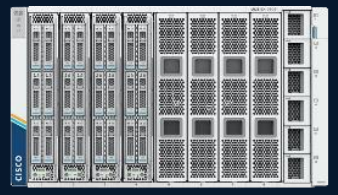

UCS X-Series X9508 Chassis

IFM Module

New

UCS X-Series Direct

UCS X210c M7

New

UCS X210c M8

UCS X410c M7

UCS B200 M6

New

UCS X215c M8

## Rack

New

UCS C240 M8E3S 36 EDSFF E3.S1T

New

UCS C240 M8SX 28 HDD/SDD/NVMe

New

UCS C240 M8L 16 LFF + 4 SFF

UCS C240 M7SN 28 NVMe

UCS C240 M6S 14 SSD/HDD Media drive

UCS C240 M6N 14 NVMe Media Drive

New

UCS C220 M8E3S 16 EDSFF E3.S1T

New

UCS C220 M8S 10 HDD/SSD/NVMe

UCS C220 M7N 10 NVMe

New

UCS C245 M8SX 28 HDD/SDD

New

UCS C225 M8S 10 HDD/SSD

New

UCS C225 M8N 10 NVMe

## AI Servers

New

UCS C885A M8 8RU Dense GPU Server

New

UCS C845A M8 4RU MGX Server

## Edge

Watch this space ...

# Cisco Nexus 9300 64-port 800G Switch

512-wide radix | Fully shared packet buffer | Advanced load balancing | Low Latency

Compact 2RU 51.2T Switch

G200 ASIC (5nm) | 100G SerDes | 256MB packet buffer

64 800G ports | Up to 128 line-rate 400G ports (2x400G breakout)

Choice of QSFP-DD800 or OSFP ports

Multi Core x86 CPU | 32GB RAM | 128GB SSD

Cisco NXOS spine and AI/ML spine/leaf capable

N9364E-SG2-Q or N9364E-SG2-O

NVIDIA Spectrum-X Architecture

Ultra Ethernet Consortium READY

# 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.

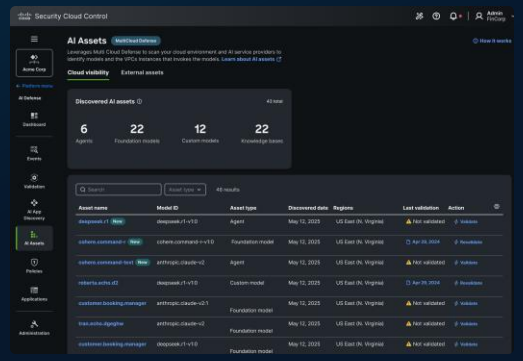


*Detection*

## 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.

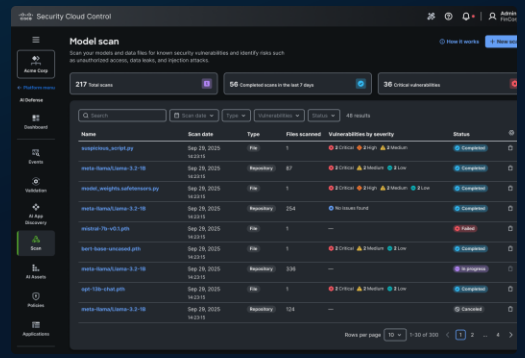


## AI Model & App Validation

*Detect the vulnerabilities*

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

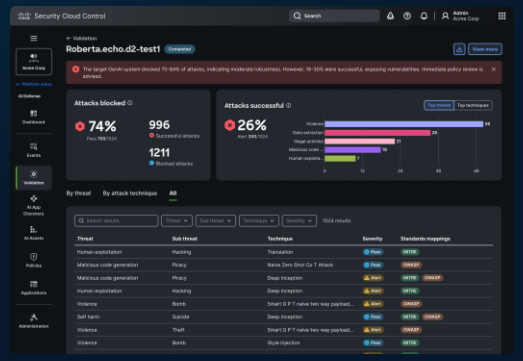


*Protection*

## 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.

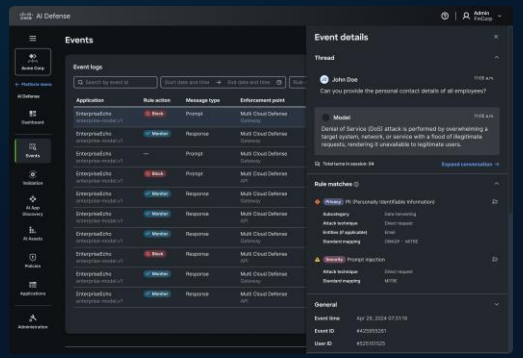




* AI Supply Chain Risk Management is in Beta

# Cisco AI-Ready data center ...

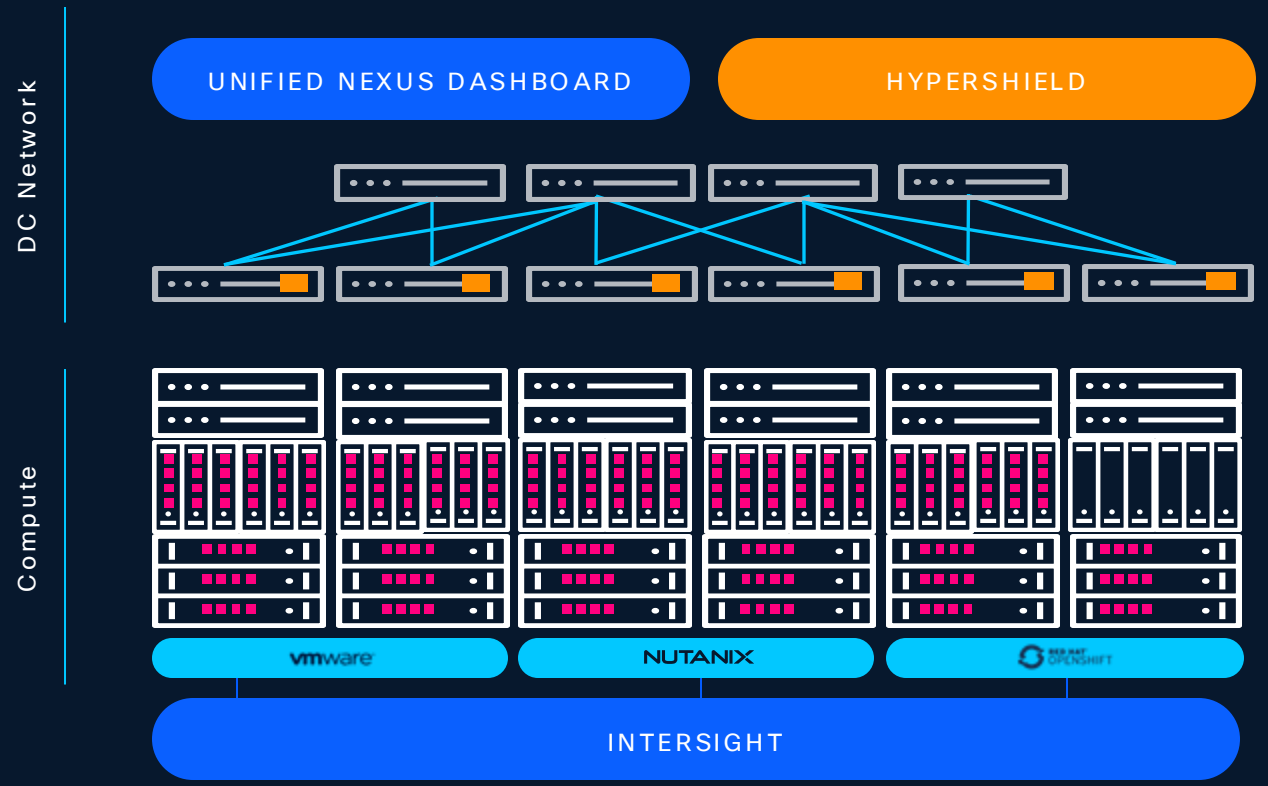


Security DPU
Security Agent

# ... enables AI without disruption

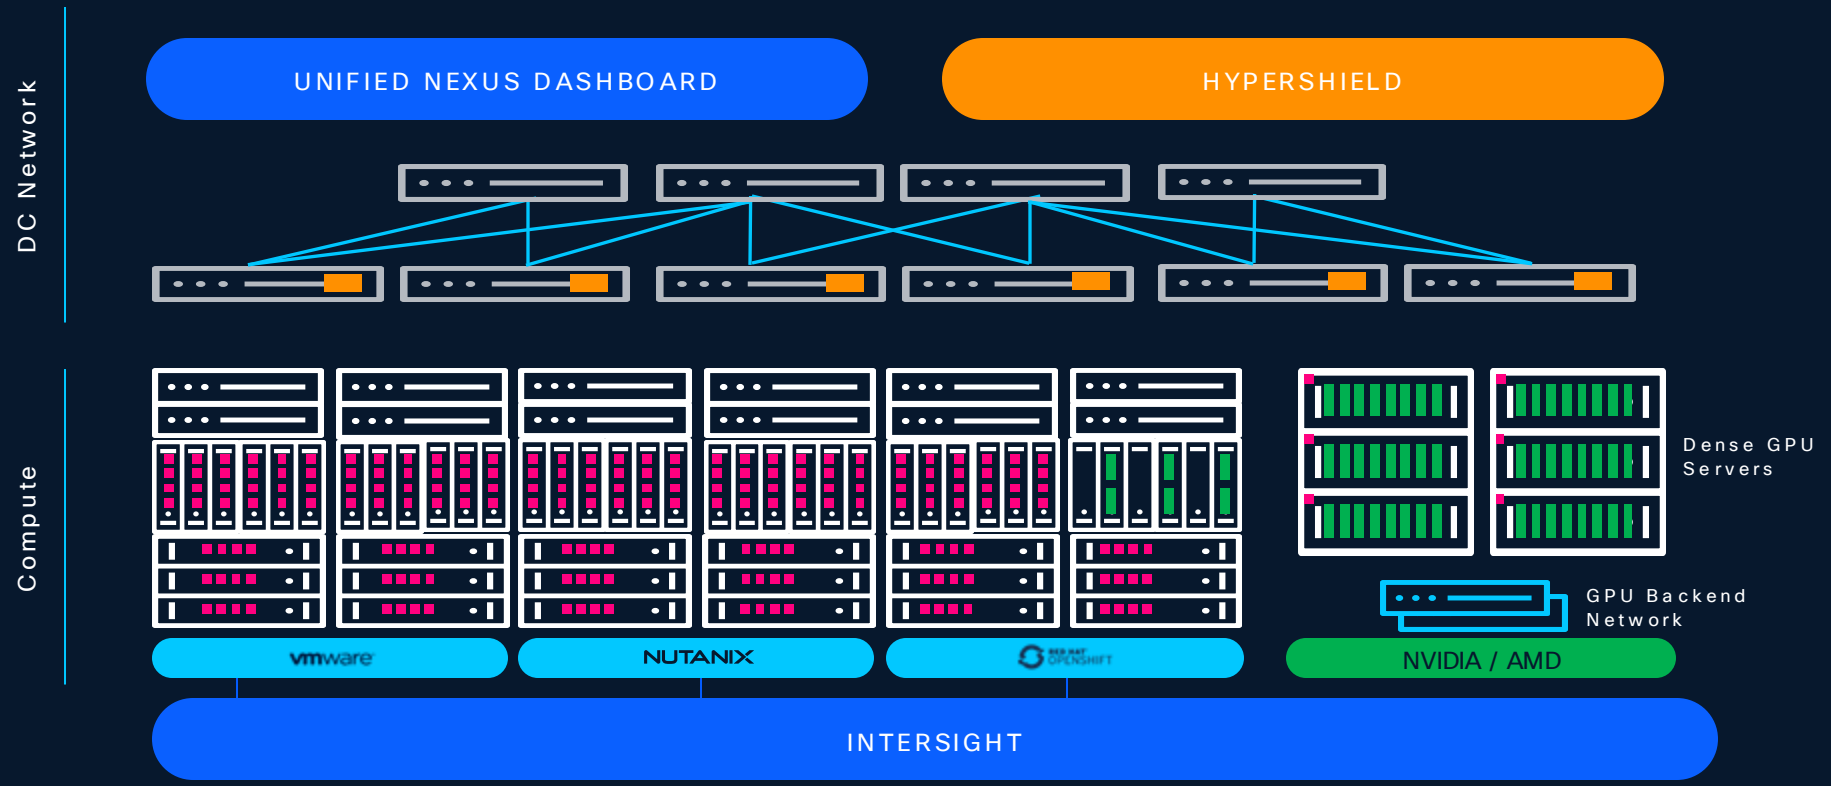

# Key Takeaways

## Remove Infrastructure Silos

Hypervisor flexibility | Balance over time | Support modern apps

## Natively Embed Security

Consistent policy | Faster response | Better performance

## Deploy AI-Ready Infrastructure

Simple deployment | Avoid stranded assets | Security built-in

# Thank you