cisco Connect

Secure AI Factory

Integrating Compute, Network, Security, and Observability for AI-Ready Environments

Mastin Bailey
Account Executive, Cloud and AI Infrastructure

Andrew MacDonell
Solutions Engineer

February 5, 2026

The Journey So Far

1

COMPUTE

UCS + NVIDIA
AI PODs
Intersight

9:15 - 10:15 AM

2

NETWORK

Silicon One
Nexus Platforms
Hyperfabric AI

10:30 - 11:30 AM

3

SECURITY

Hypershield
Isovalent
Splunk

1:00 - 2:00 PM

4

INTEGRATION

How it all
works together

CURRENT SESSION

2:15 - 3:15 PM

Three acts. Three pillars. Now, the integration story.

INTEGRATION

The Three Pillars

FOUNDATION

COMPUTE

UCS Servers with NVIDIA GPUs

Intersight Management
AI PODs

CONNECTIVITY

NETWORK

Silicon One ASICs

Nexus Platforms
Hyperfabric AI

PROTECTION

SECURITY

Hypershield

Isovalent
Splunk Observability

But how do they work **together?**

INTEGRATION

The Integration Problem

TEAM 1

Server Team

Intersight • IPAM • Monitoring

- Compute utilization
- Provisioning time
- Uptime SLAs

TEAM 2

Network Team

Network Manager • Config • Flow

- Throughput
- Latency
- Packet loss

TEAM 3

Security Team

SIEM • Firewall • Compliance

- Zero breaches
- Policy compliance
- Mean time to detect

The result? 3 weeks to deploy. Endless ping-pong. "It works on my side."

Cisco Secure AI Factory with NVIDIA

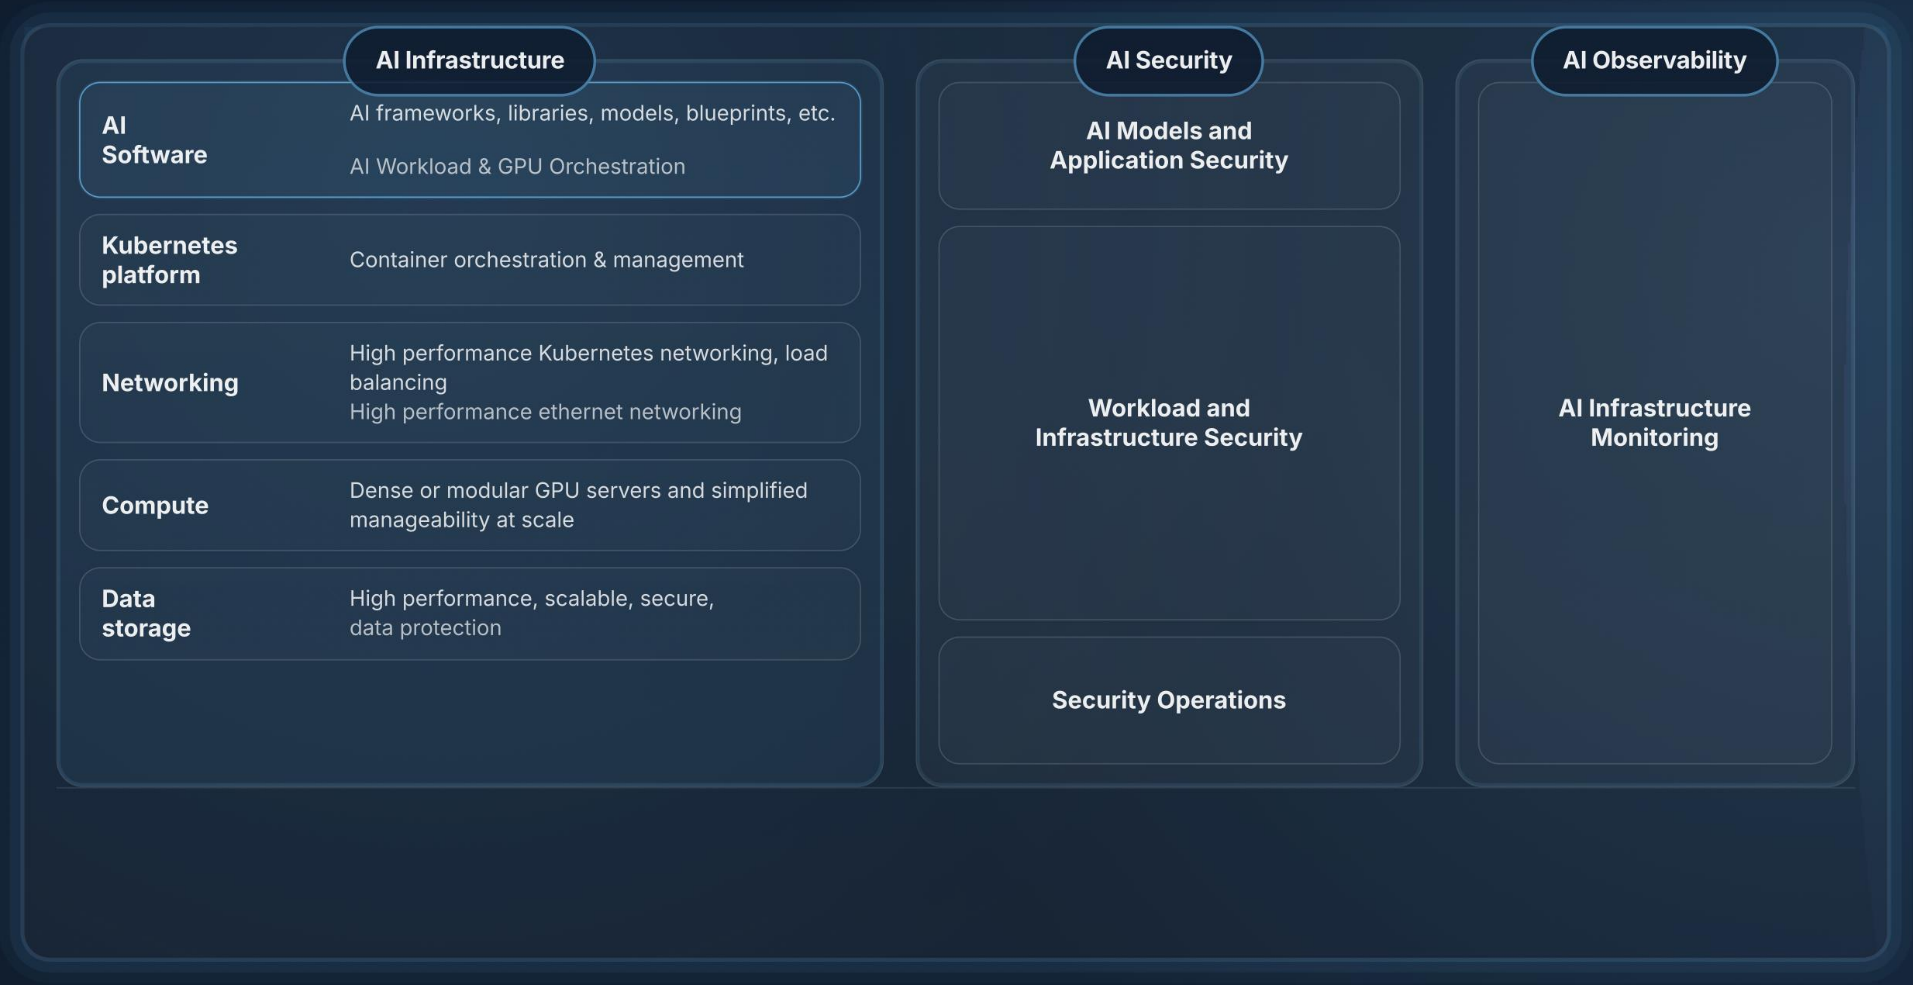

AI Infrastructure

AI Software	AI frameworks, libraries, models, blueprints, etc. AI Workload & GPU Orchestration
Kubernetes platform	Container orchestration & management
Networking	High performance Kubernetes networking, load balancing High performance ethernet networking
Compute	Dense or modular GPU servers and simplified manageability at scale
Data storage	High performance, scalable, secure, data protection

AI Security

AI Models and Application Security

Workload and Infrastructure Security

Security Operations

AI Observability

AI Infrastructure Monitoring

INTEGRATION

Traditional vs. Integrated Deployment

Traditional Approach

- **Day 1-5:** Server team provisions UCS
- **Day 6-7:** Email to network team
- **Day 8-14:** Network configures switches
- **Day 15-18:** Troubleshooting connectivity
- **Day 19-21:** Security review & approval

TOTAL TIME

3 WEEKS

VS

Integrated Approach

- **Hour 1:** Open Nexus Dashboard, select AI Fabric template
- **Hour 1-2:** Automated provisioning (Nexus ⇔ Intersight)
- **Hour 2-3:** Validation & telemetry to Splunk
- **Hour 3-4:** Security policies auto-applied

TOTAL TIME

2-4 HOURS

60-80% faster deployment. Not marketing. Reality.

Security-first architecture enables safe Enterprise AI

Security at all layers of the stack

Securing the Applications

Cisco AI Defense: Testing and runtime security of LLMs and GenAI applications, integrated with NVIDIA AI.

Securing the Workloads and Infrastructure

Cisco Hybrid Mesh Firewall: Unified management, consistent, pervasive policies.

- □ **Cisco Isovalent:** Enhanced visibility into cloud native interactions, consistent policy definition and enforcement.
- □ **Cisco Hypershield:** Protection against lateral movement, proactive vulnerability mitigation.
- □ **Cisco Secure Firewall:** Threat protection at scale without compromising performance.

Security Operations

Splunk Enterprise Security: Real-time threat detection, investigation, and response through analytics, automation, insights.

INTEGRATION

End-to-End Observability

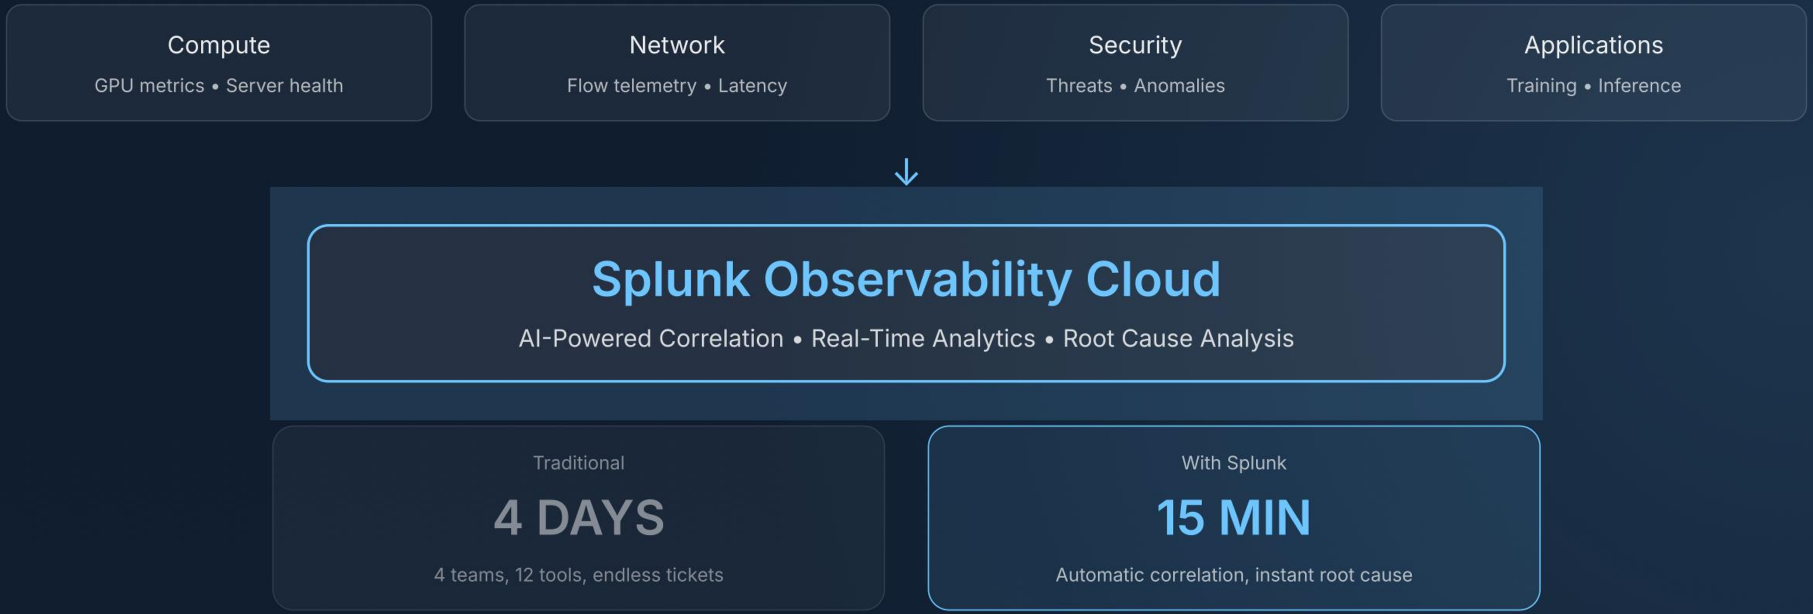

Compute

GPU metrics • Server health

Network

Flow telemetry • Latency

Security

Threats • Anomalies

Applications

Training • Inference

↓

Splunk Observability Cloud

AI-Powered Correlation • Real-Time Analytics • Root Cause Analysis

Traditional

4 DAYS

4 teams, 12 tools, endless tickets

With Splunk

15 MIN

Automatic correlation, instant root cause

Single pane of glass. Instant answers.

INTEGRATION

Cisco Secure AI Factory with NVIDIA

Secure

Security, Observability
and resilience

Scalable

High performance at any scale
enables faster delivery of
AI tokens and applications

Simple

Deployment simplicity and
flexibility helps improve AI
and IT team productivity

Secure. Scalable. Simple.

Customer Success Story

Deployment Time

~~12 weeks~~

↓

3 weeks

GPU Utilization

~~45%~~

↓

82%

> *"The integration isn't just about technology. It changed how our teams work together."*

Real results. Real transformation.

Your Next Steps

1

ASSESS

Take an honest look at your current state

→ How many management tools?
→ How many handoffs to deploy?
→ Can you trace issues end-to-end?
→ Do teams share platforms?

2

DEFINE

Set measurable integration goals

→ Faster time to production?
→ Better security posture?
→ Reduced operational costs?
→ Improved developer experience?

Pick 2-3 specific, measurable goals

3

PILOT

Start small, measure, expand

→ Choose one AI use case
→ Deploy on Cisco infrastructure
→ Measure vs. current process
→ Document lessons learned
→ Scale what works

Start with one team. One workload. Prove the value. Then expand.

cisco Connect

Thank you

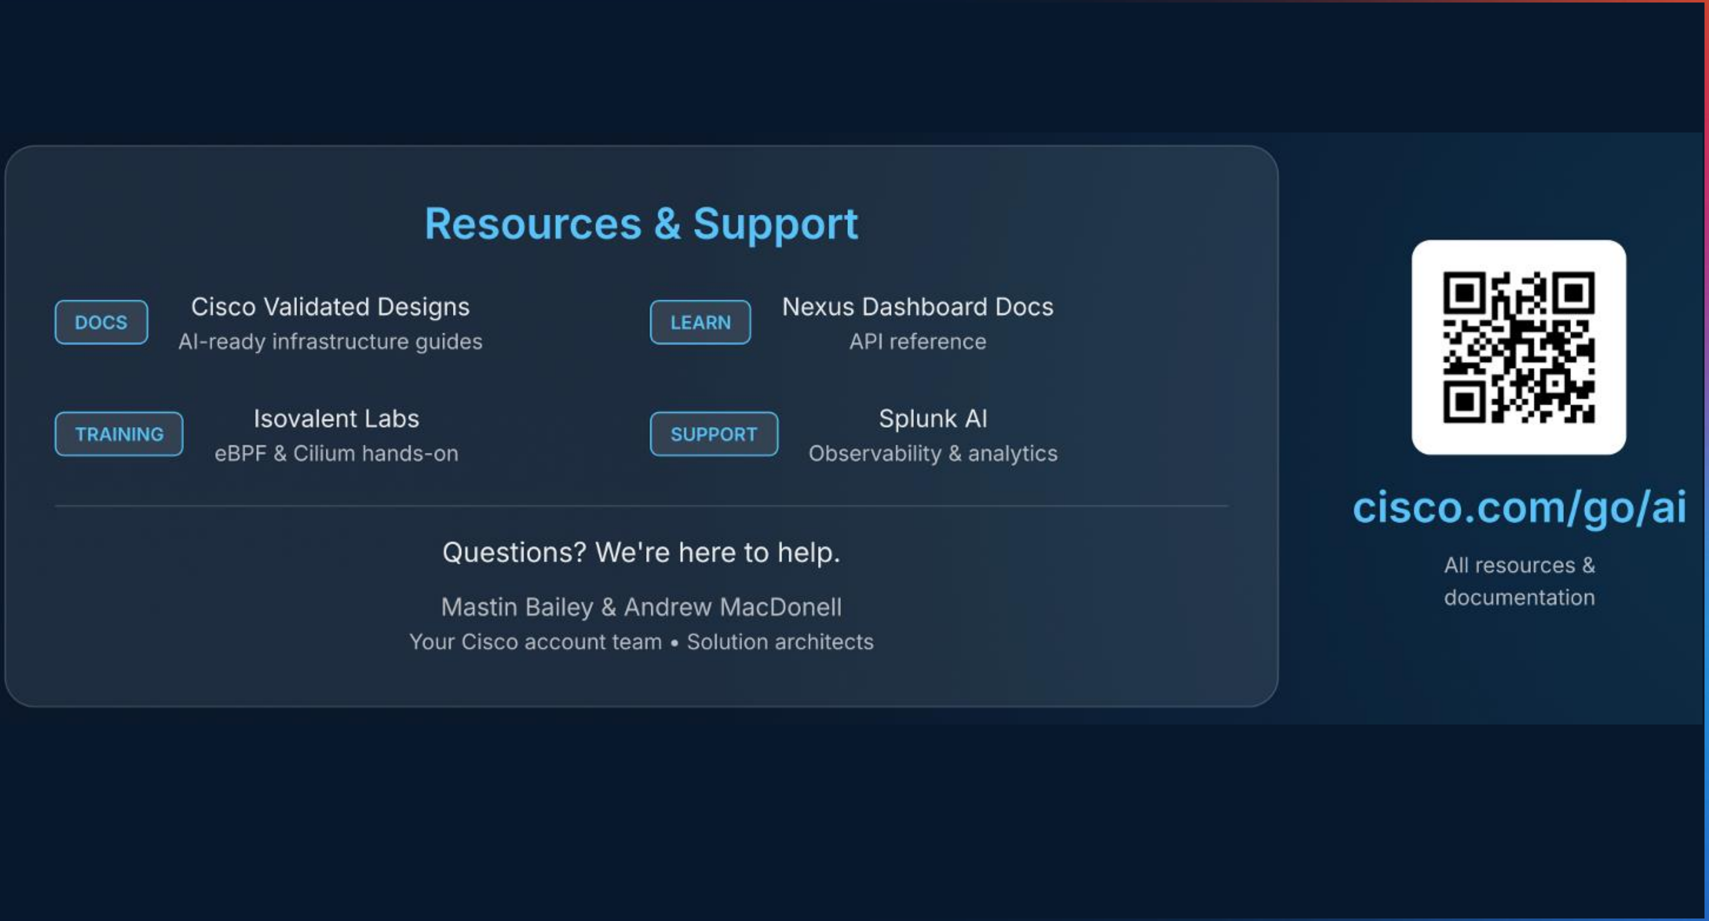

Resources & Support

DOCS Cisco Validated Designs
AI-ready infrastructure guides

LEARN Nexus Dashboard Docs
API reference

TRAINING Isovalent Labs
eBPF & Cilium hands-on

SUPPORT Splunk AI
Observability & analytics

Questions? We're here to help.

Mastin Bailey & Andrew MacDonell
Your Cisco account team • Solution architects

cisco.com/go/ai

All resources & documentation