

Seamless Operations and Observability in the Data Center

Brian Jacklin
Solutions Architect
Aaron Trompeter
Solutions Architect



Agenda

- Operational Changes in the AI Ready Data Center
- Visibility in the AI Ready Data Center
- Optimizing and Assure Network Infrastructure for AI Workloads

Cisco powers how people and technology work together across the physical and digital worlds

AI-ready data centers

Transform data centers to power AI workloads anywhere

Future-proofed workplaces

Modernize everywhere people and technology work and serve customers

Secure global connectivity

Digital resilience

Keep your organization securely up and running in the face of any disruption

Accelerated by Cisco AI

Transform data centers to power AI workloads anywhere

Public and private clouds, edge, on-premises



Comprehensive infrastructure

Power AI with networking, compute, and storage in fully-integrated, scalable, and modular systems for all workloads

Seamless operations and observability

Remove silos with unified management, observability, and assurance for traditional and AI workloads, across all environments

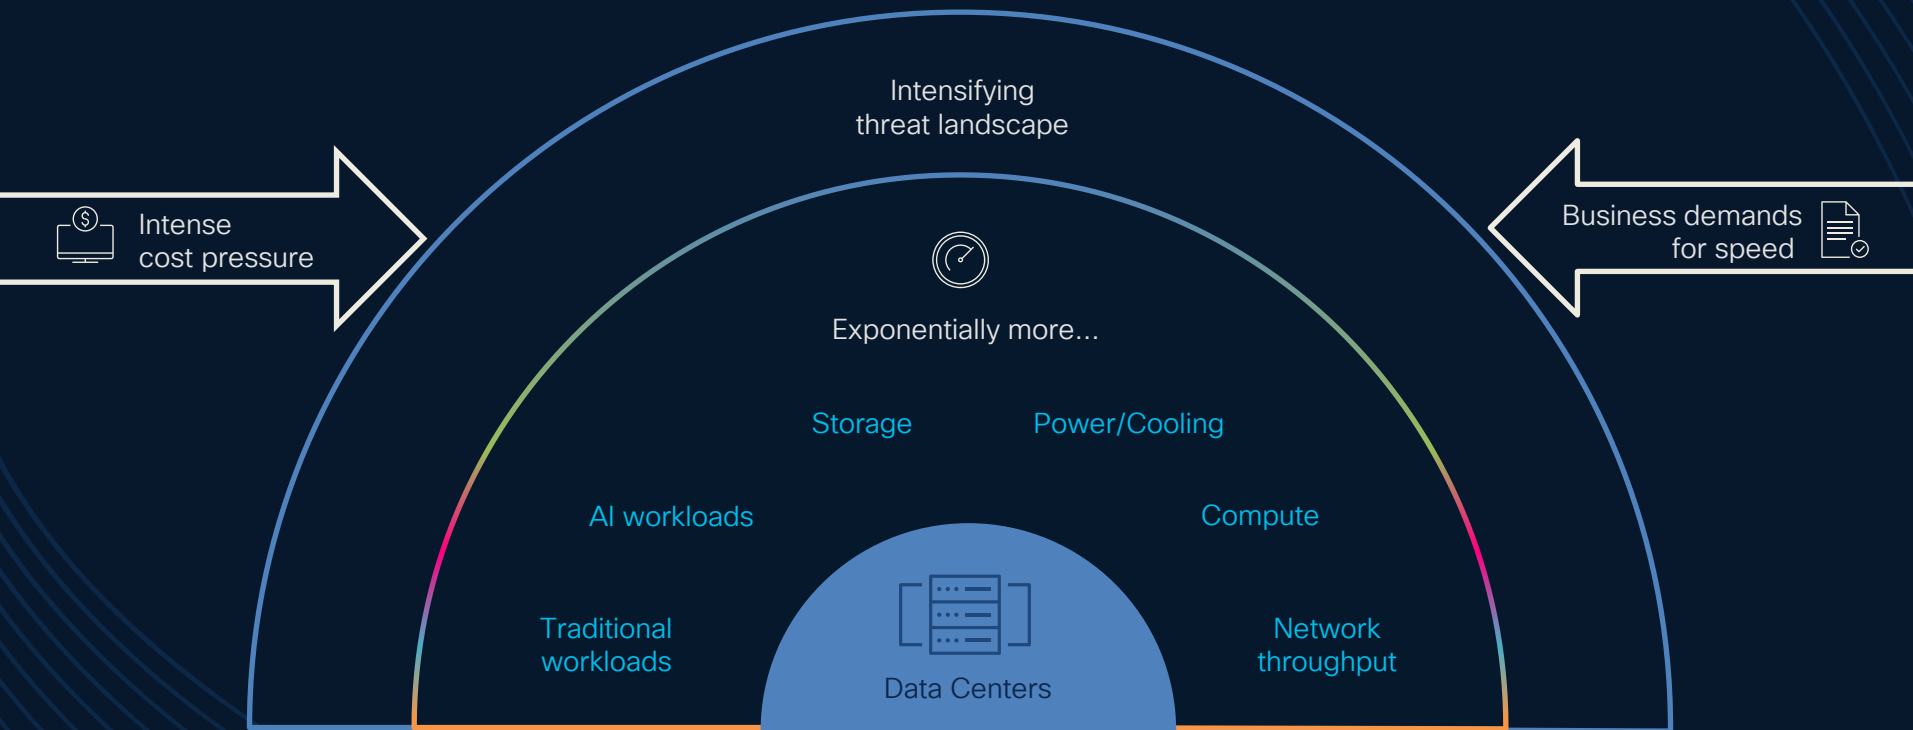
Security from ground to cloud

Protect hyper-distributed workloads by infusing security everywhere






Operations – in the AI Ready Data Center

We need to reimagine the data center for the AI era

While managing cost pressures, intensifying threats, and business demands



Cisco Data Center Networking Portfolio

	Nexus Hyperfabric		Nexus Dashboard	
			 	
Operating Model	Fabric-as-a-Service Cisco Cloud-Managed Controller		Customer Managed On-Prem Controller	
Flexibility & Customization	Prescriptive			Customizable
IT Staff Network Skillset	Generalist			Specialist
Deployment Type	Greenfield		Greenfield & Brownfield	

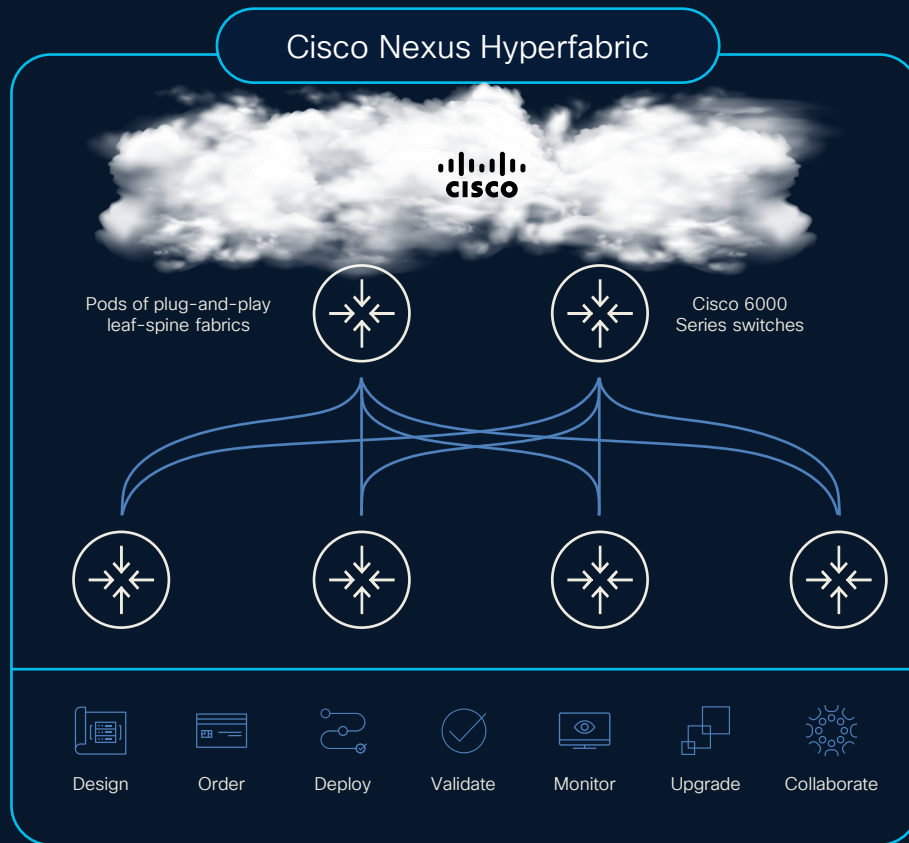
Greenfield: new fabrics not being managed
by Nexus Dashboard

Cisco Nexus Hyperfabric

✓ Design, deploy and operate on-premises fabrics located anywhere

✓ Easy enough for IT generalists, application and DevOps teams

✓ Outcome driven by a purpose-built vertical stack

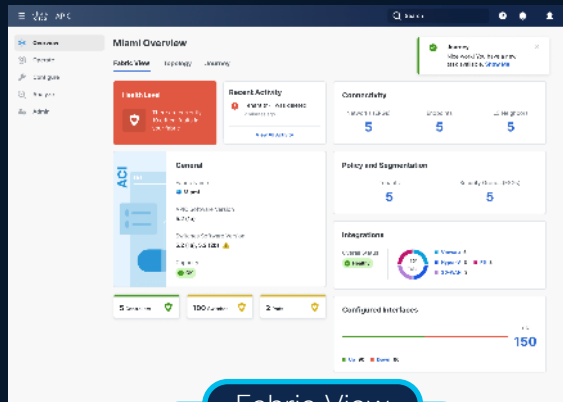
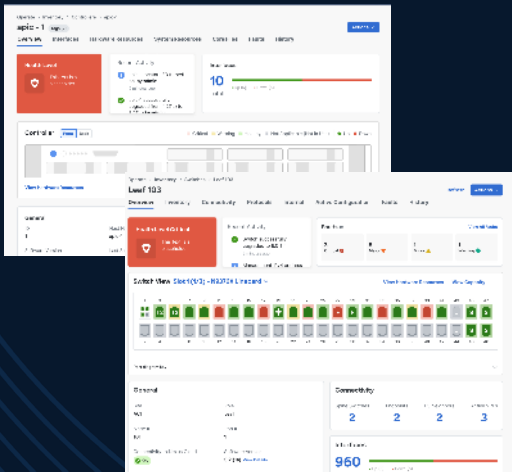


ACI simplification

A modern, faster UI for 2025 and beyond

Find the information you need with **ease**
 Edit configurations **faster** (Copy + Paste, Clone)
 Jointly **visualize** Policy Groups and EPGs
 configured on an interface

Switches and Controllers





Fabric View

- Quick access to main items
- In-context information
- Reduce jumping between sections

Quickly identify **who can talk to whom** and how
Associate **contracts** and **L4-L7** devices with ease
In-context security policy and/or domain association
for every EPG/ESG

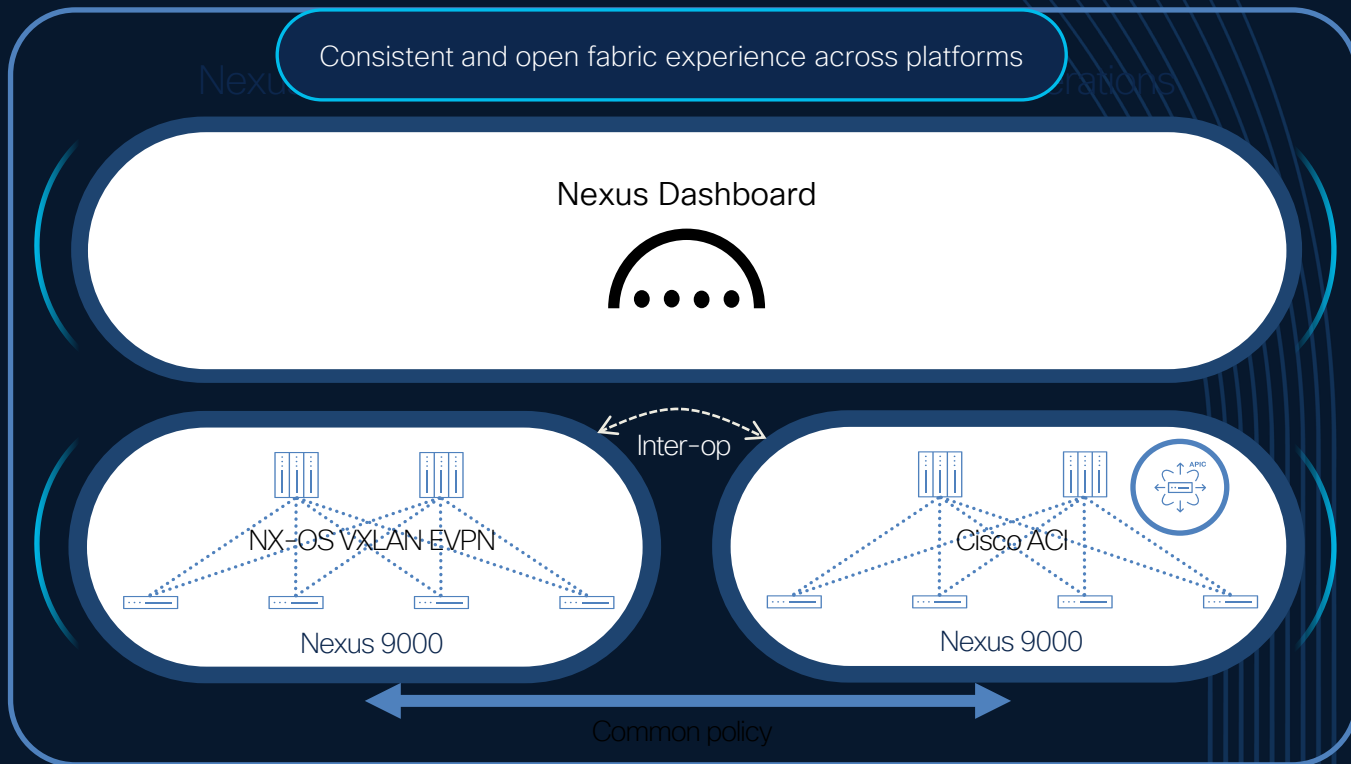
Tenants, EPGs & contracts



Introducing Cisco Nexus ONE

Solution Components

- Zero trust networking and micro-segmentation based on endpoint / security groups
- Advanced service chaining and service redirection
- Administrative multi-tenancy
- Standards based inter-op with 3rd party networks using open networking protocols
- Integration with CI/CD pipelines with DevOps ready APIs



Cisco Nexus Dashboard

Data Center Networking Management



Provision

Quickly Deploy any Data Center Network Fabrics
Deploy traditional LAN, EVPN VXLAN, AI/LM and SAN fabrics



Manage

Multisite fabric and device management in a single view
Device upgrades, Change control with roll-back, traffic steering



Analyze

Identify, troubleshoot and resolve issues quickly
Traffic Analytics, Delta and pre-change analysis



Segmentation

Unify data center networks across both ACI and NX-OS
Leveraging the capabilities of Cisco Nexus One

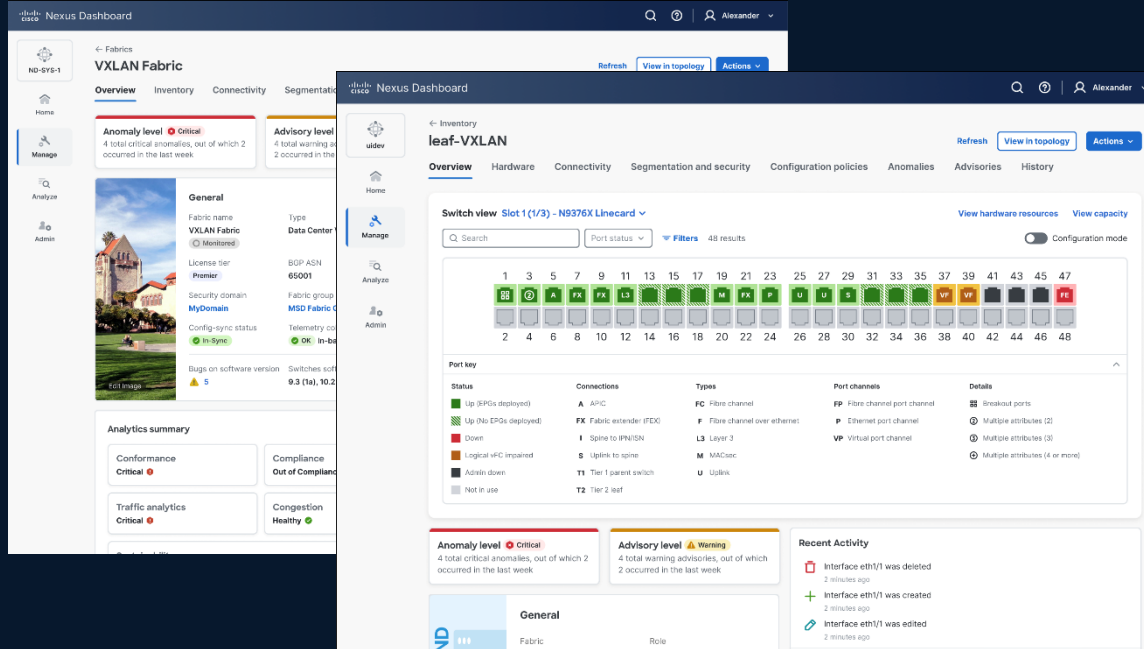
Simplified Network
Management for modern
data center architectures



Cisco Nexus Dashboard – Operations

ND 4.0 -> 4.1

Fabric and switch management



Consistent fabric management views

ACI, VXLAN, Classical, Routed, IPFM and more!

New AI/ML Fabric Type

Recent activity, security groups, top X interfaces

Switch configuration (NX-OS only) and monitoring

Anomalies, advisories, analytics summary, traffic

ACI two-way integration expected in 4.1 (future)


Benefits

Consistent operations for all switches

Nexus Dashboard - Operations

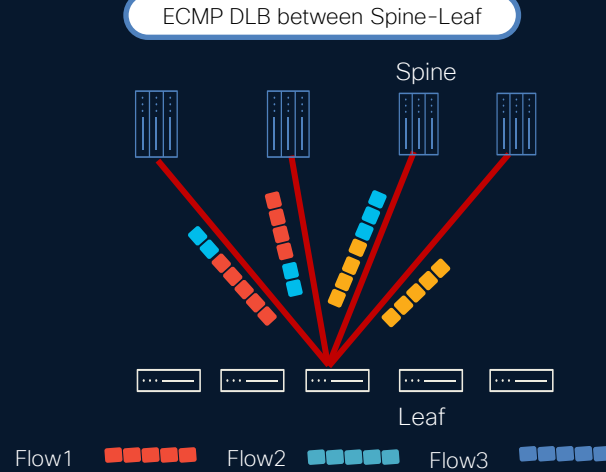
Automate your AI/ML Network Configurations

ND 4.0 -> 4.1



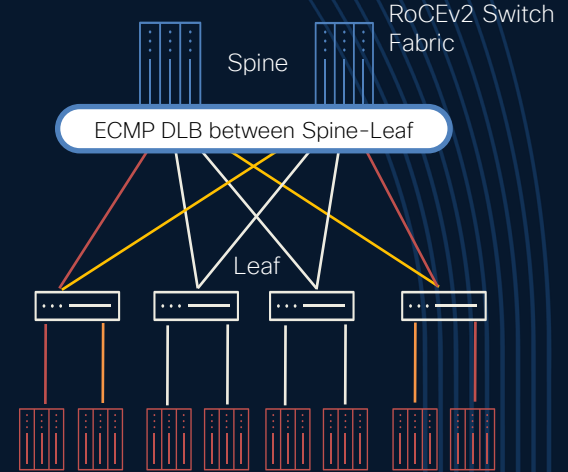
Manage network congestion
with Lossless Network (PFC + ECN)

RoCEv2 support on Nexus 9800
(Planning Q1CY25)



Load balance flows/flowlets
based on link utilization

Better hashing results in AI fabrics
with uniform flow size and header information



Traffic efficiency through pinning rules
Map traffic from each downlink to the desired uplink

Allows efficient selection of Spines for communication
between leaf and spines

Cisco Nexus Dashboard – Operations

ND 4.0 -> 4.1

Visibility with historical information



Fabric

Aggregate switches, interfaces, and endpoints (including VMs and containers). View anomalies, advisories, routing changes, and more



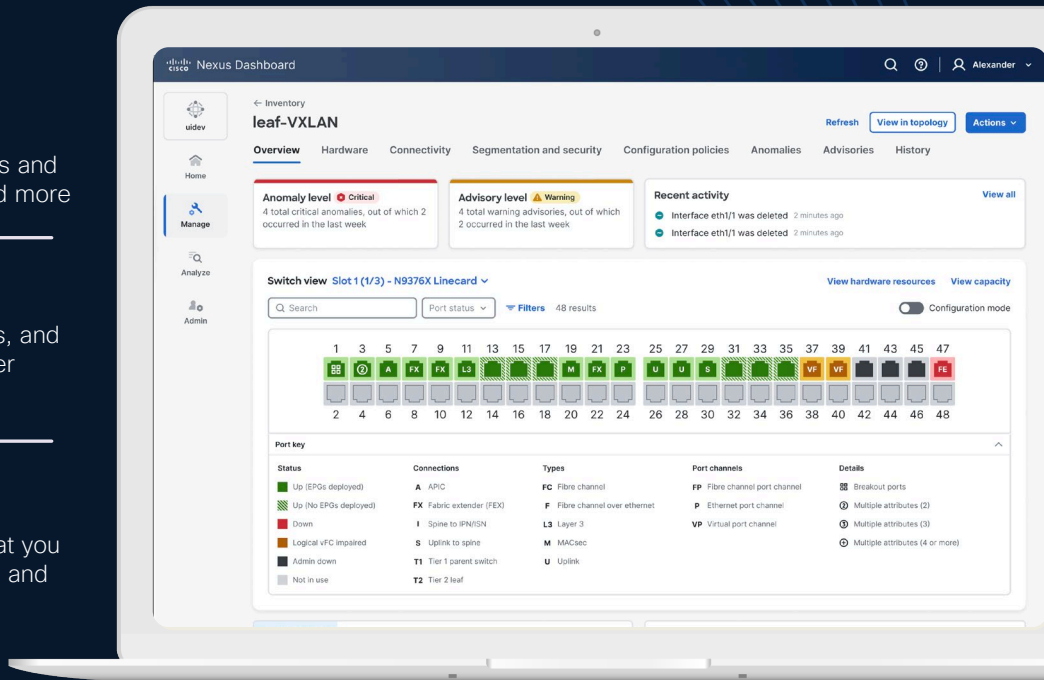
Switches and Interfaces

Learn in real time about traffic, statistics, hardware resources, and network events. Easily identify what is connected on the other side

1.1.1.1

Endpoints and VMs

Search for DNS names or IP and MAC addresses to find what you need faster. Extend beyond the switch port through VMware and Kubernetes integrations



Benefits

Consistent user experience

Easy navigation into an issue

Cisco Nexus Dashboard – Operations

ND 4.0 -> 4.1

Bug Scan



Continuous fabric assessment

Nexus Dashboard automatically performs a bug scan every 7 days and reports its findings. Optionally, it can be run on-demand



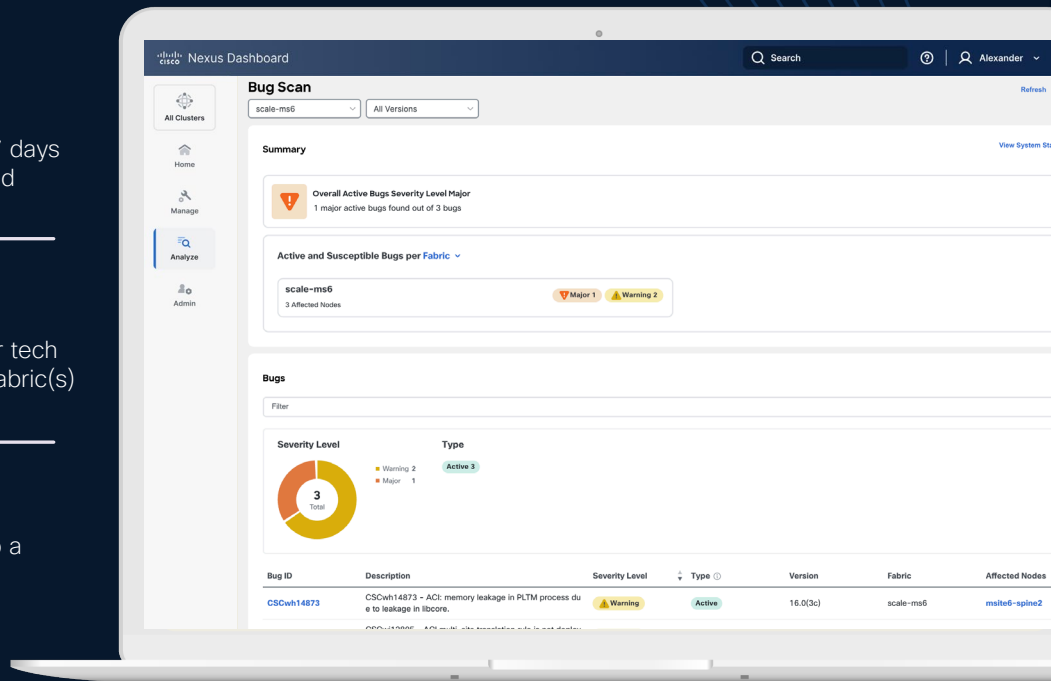
Active Bugs

Bugs that match Nexus Dashboard signatures based on your tech support files, will be reported as active bugs affecting your fabric(s)



Other known bugs

You may also search for specific bugs in a version related to a feature you have implemented or are planning to implement



Benefits

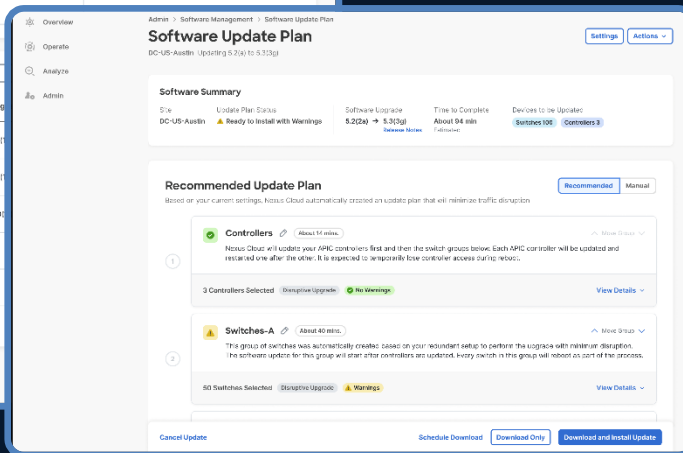
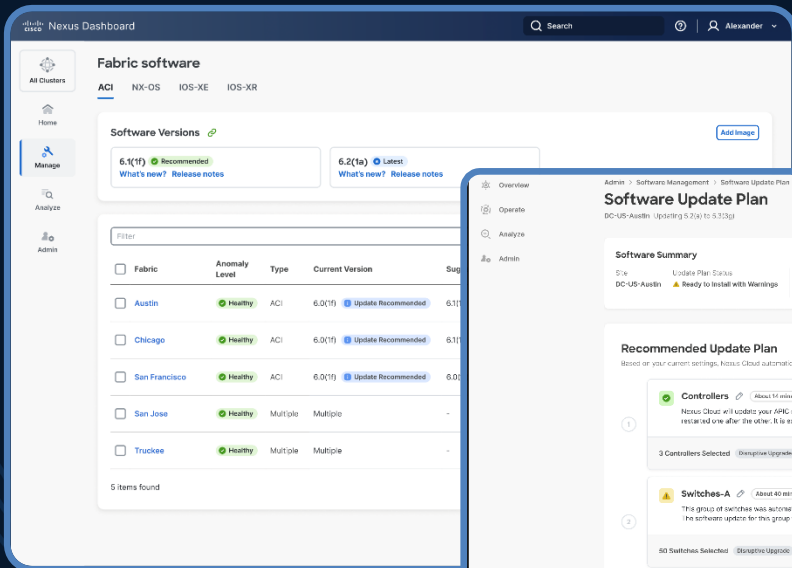
Continuous risk identification

Software recommendations

Cisco Nexus Dashboard – Operations

Software Upgrades – Workflows

ND 4.0 -> 4.1



Supports fabric upgrades for NX-OS, IOS-XE/XR

Recommended update plans to minimize disruption

Customize plans based on your preferences

ACI upgrade support is scheduled for 4.1 (future)

Pre and Post upgrade analysis for ACI, NX-OS

Direct image downloads from Cisco.com (future)*

Benefits

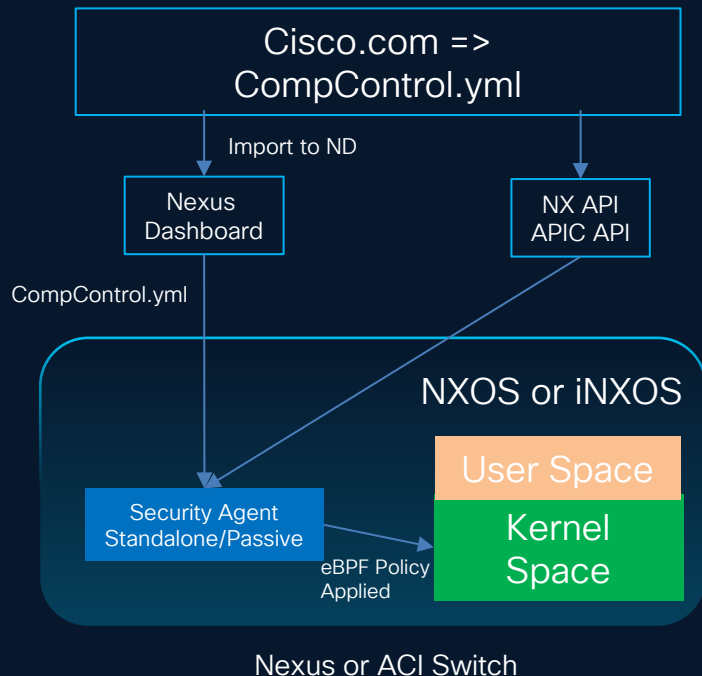
Minimize operational risks and keep software up to date

* Requires ND connectivity to Cisco

Securing Data Center Nexus and ACI Switches against Vulnerabilities

Problem: DCN customers face high Opex due to the need to upgrade 100s to 1000s of switches outside regular maintenance windows every 6 to 12 months to address CVEs

Requirement: A solution that offers compensating controls to mitigate the risk posed by CVEs without NXOS or iNXOS upgrades.



Workflow Proposal:

1. ND on prem with adv, premier license
2. NXOS runs security standalone agent in passive mode
 - Passive means no external connection cloud or controller
3. All CVE compensating controls (CompControl.yml) are published on CCO
4. ND imports CompControl.yml from CCO
5. ND pushes CompControl.yml to NX/iNX switch security agent
6. Security agent applies the fixes with eBPF policy

Packaging:

1. Capability included in DCN advantage and premier license tiers only
2. Security agent will be included in NXOS/iNXOS in a future version

Cisco Nexus Dashboard – Operations

ND 4.0 -> 4.1

Connected TAC – Log Collector



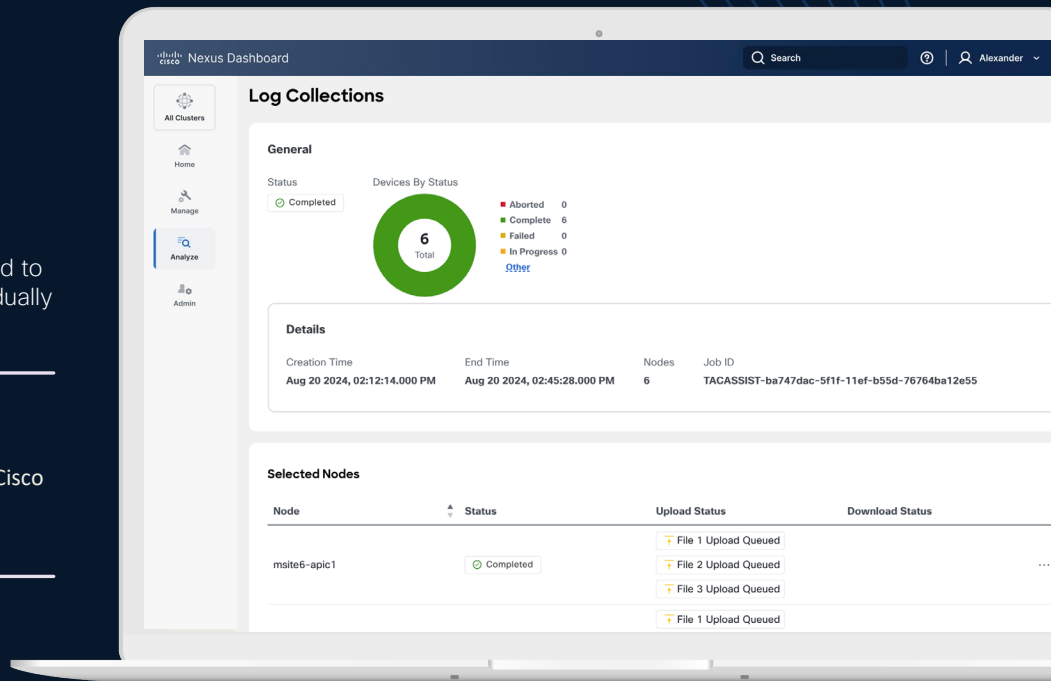
Avoid manual collection

Nexus Dashboard streamlines the process, removing the need to collect tech-support files for each switch in your fabric individually



Reduce time to resolution

By anticipating collection and establishing a direct connection with Cisco TAC for problem resolution



* Connected TAC requires Intersight connection

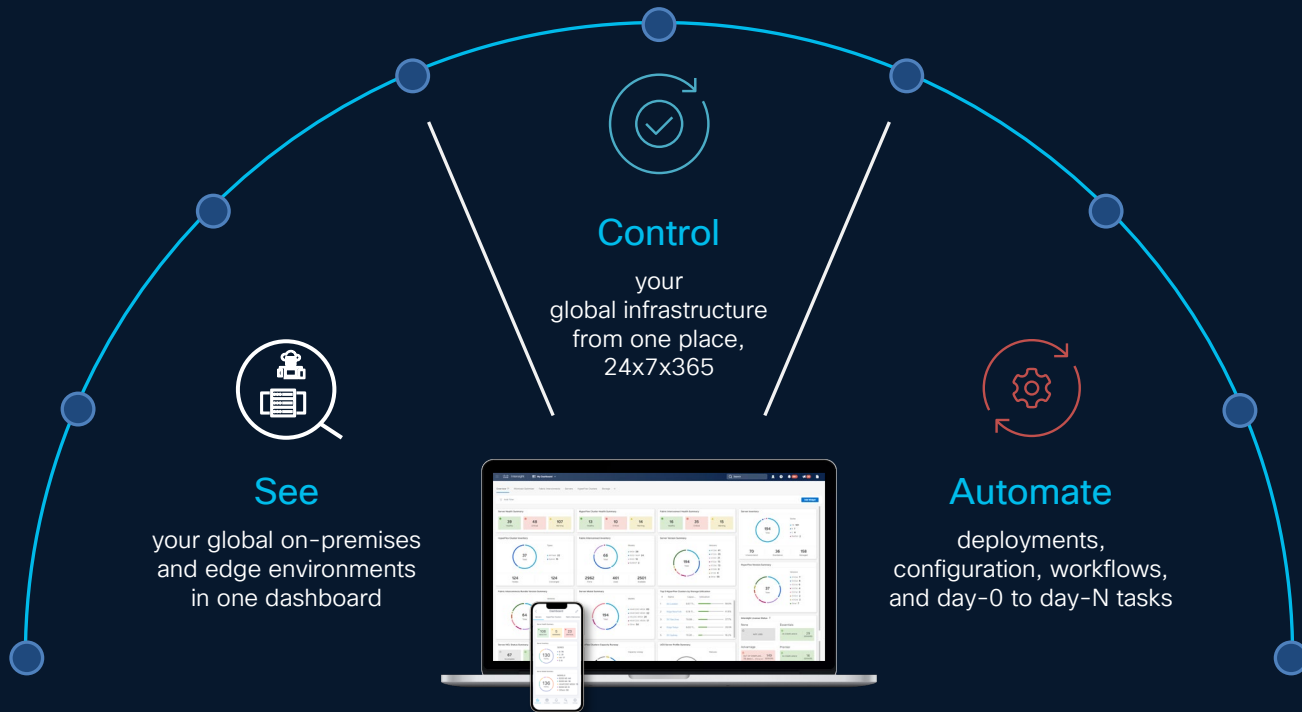
Benefits

Accelerate time to resolution

Save time through automation

Cisco Intersight Infrastructure Service

IT Operations. Simplified.



View in one place – Cisco Intersight

- Your global Cisco and third-party infrastructure
- Servers, virtual machines, third-party servers, storage, and networking
- Real-time status, security advisories, field notices, and hardware compatibility status
- Performance monitoring and metrics visibility



Customizable dashboard uses widgets for a single view of health and inventory status of managed devices

Prevent and Resolve – Cisco Intersight

Infrastructure issues with custom insights and automated support

- Reduce troubleshooting by 45-60 minutes*
- Close RMA cases up to two days faster*

Real-time and historical data

Security advisories and field notices

HCL baselining and notifications

Proactive RMAs

Connected support (Cisco® TAC)

Firmware upgrades

Compliance

Create automated workflows – Cisco Intersight

- Accelerate delivery of applications and infrastructure
- Reduce risk with standardization
- Simplify cross-domain orchestration

Two options

- Workflow designer
Drag-and-drop workflow authoring
- APIs/SDKs
For using your existing automation tools
(Ansible, Terraform, etc.)



Visibility in the AI Ready Data Center

Your Network Generates a LOT of Information

Software Telemetry

Config File

Tech-Support

Accounting Logs **RIB**

Debug Logs



Hardware Telemetry

ASIC

Statistics

Flow Events

Buffer

Utilization

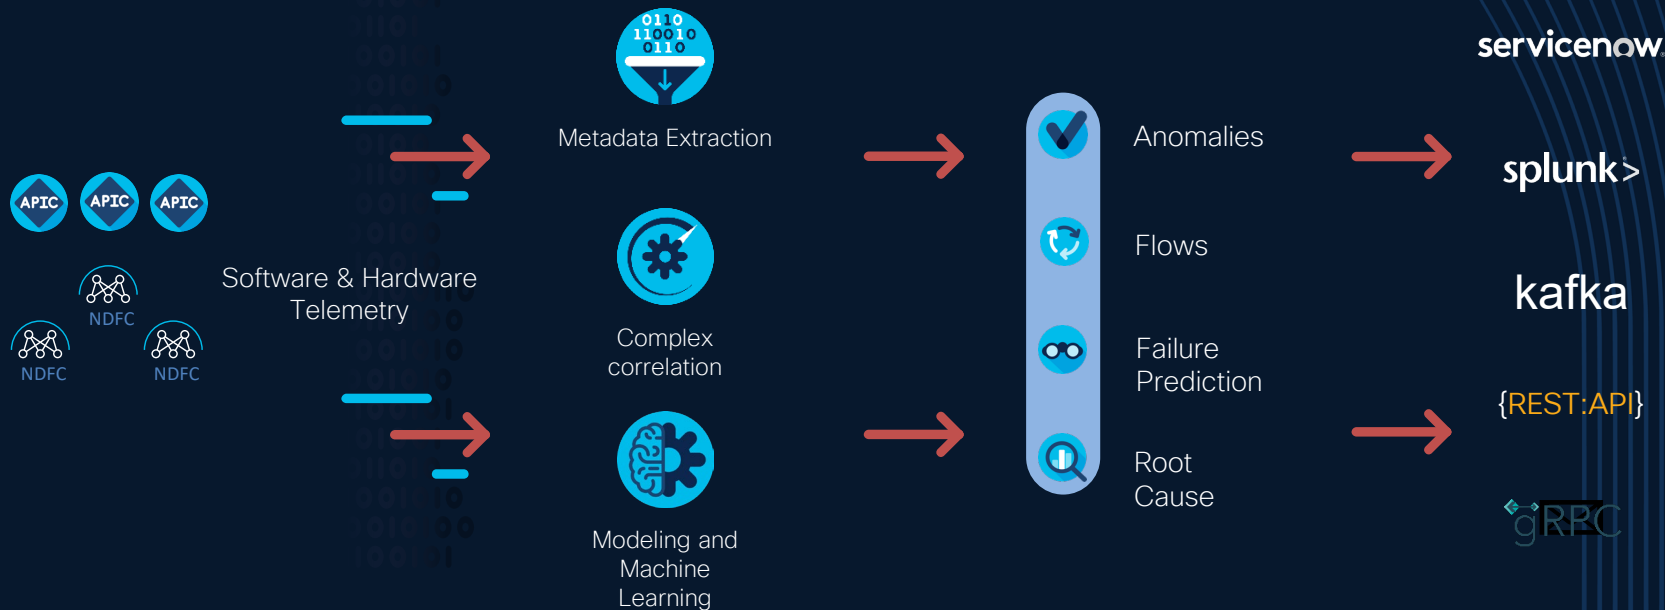
Advanced QoS Visibility



Having all the information is not the same as having all the answers

Nexus Dashboard - How it works

Streaming Telemetry + Actionable Insights



Proactive Operations: Increase Availability, Performance and Simplify Operation

Nexus Dashboard – SAN Traffic Visibility

ND 4.0 -> 4.1

Identify issues through SAN Insights



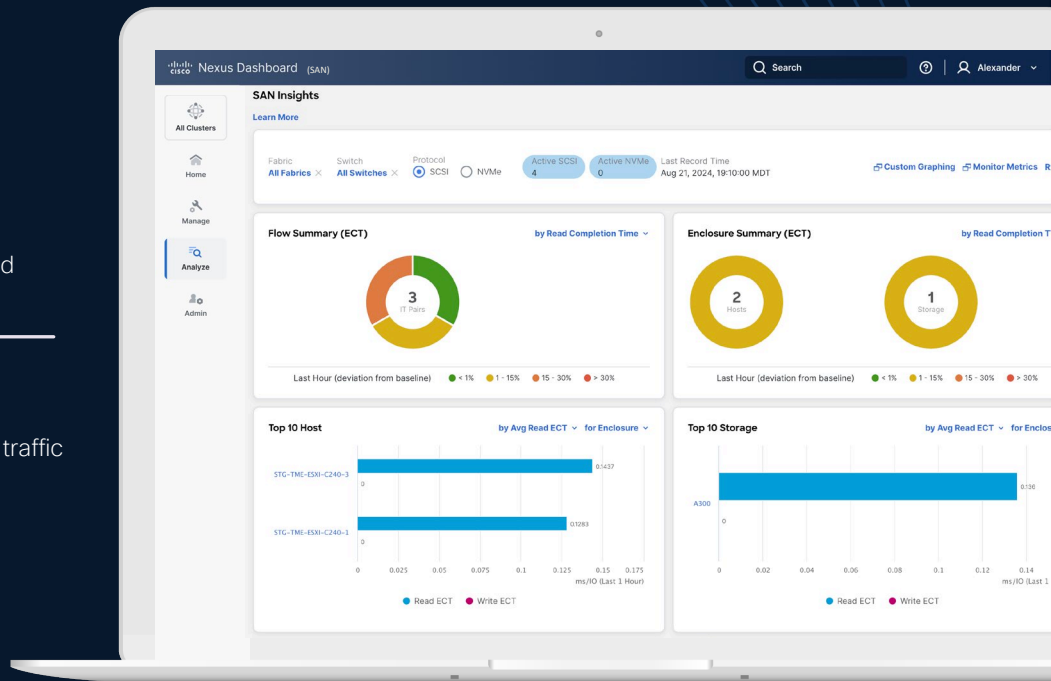
Maintain SLAs

Monitor SCSI and NVME metrics for your different Initiator and Target pairs, including ECT/DAL/read/write times and more



Identify deviations and anomalies

Based on active I/O, throughput, and IOPS, detect abnormal traffic patterns and spikes that may affect network performance



Benefits

Anticipate performance issues

Customize monitored services

Nexus Dashboard – LAN Flow Visibility

ND 4.0 -> 4.1



Automatically identify services in your network

Through well-known L4 ports (e.g. Web - TCP port 80) pre-loaded service categories are learned and monitored. Category customization is also allowed based on your own preferences.



Pervasive across switches and fabrics

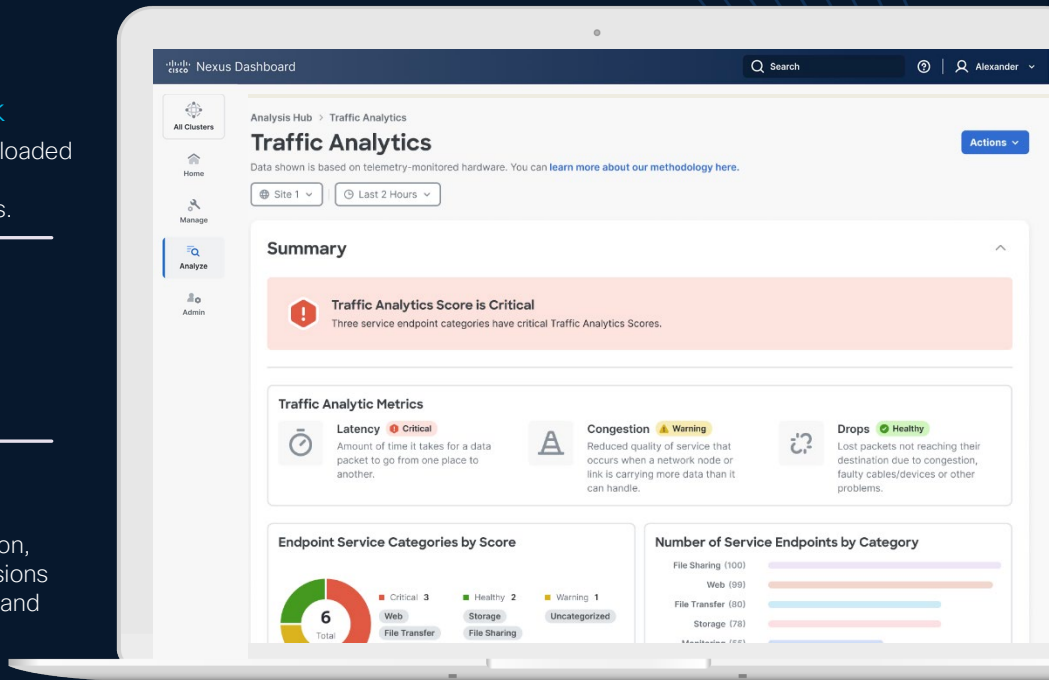
Learn about clients and services connecting across multiple fabrics without rules or any additional rule configuration.¹



Granular visibility for every connection

From overall fabric score, to category, service, and connection, Traffic Analytics can monitor individual client-to-service sessions and allows you to “tap-in” by capturing flow records on demand

Requires: ACI – 6.1.1 and NX-OS 10.4(2F)



* Fabrics must have PTP configured for timestamping

Benefits

Anticipate performance
issues

Customize monitored
services

Nexus Dashboard – Fabric Delta Analysis

ND 4.0 -> 4.1

More than comparing individual metrics



Anomalies

View new, unchanged, and cleared anomalies between two points in time to learn about changes in your network



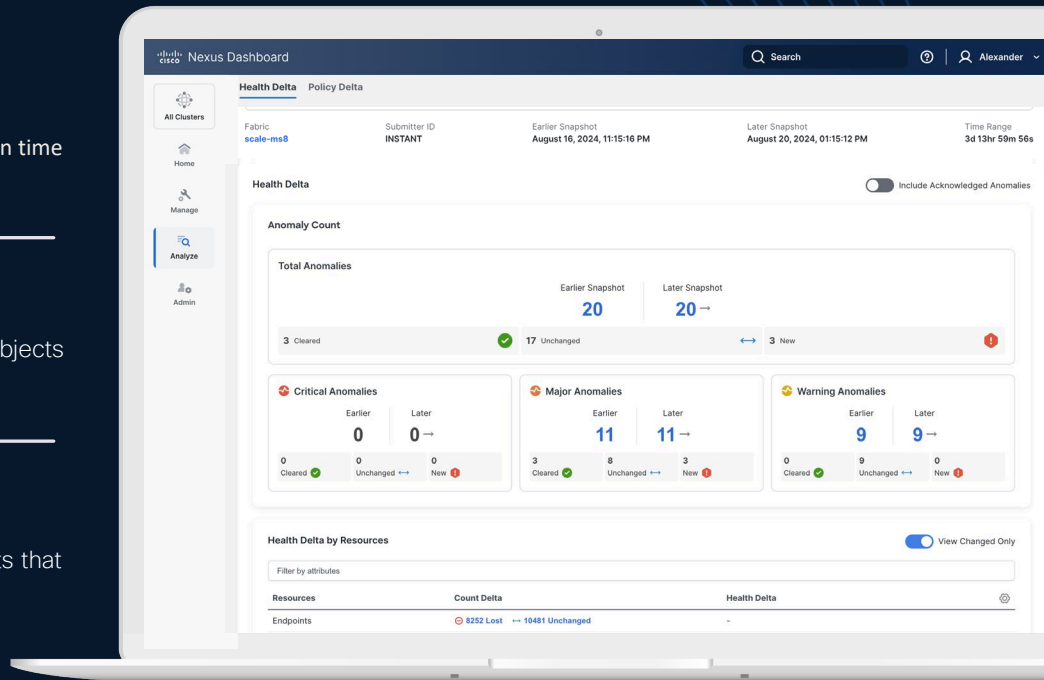
Resources

Learn about endpoints, interfaces, routes, VRFs, and other objects that may have changed their configuration or health status



Policies and configuration

Get an XML (ACI) or CLI (NX-OS) view of the specific objects that have changed, how they changed and who changed them



Benefits

Learn about any changes
in your network

Troubleshoot
faster

Cisco Nexus Dashboard – Operate

ND 4.0 -> 4.1

Pre-change analysis



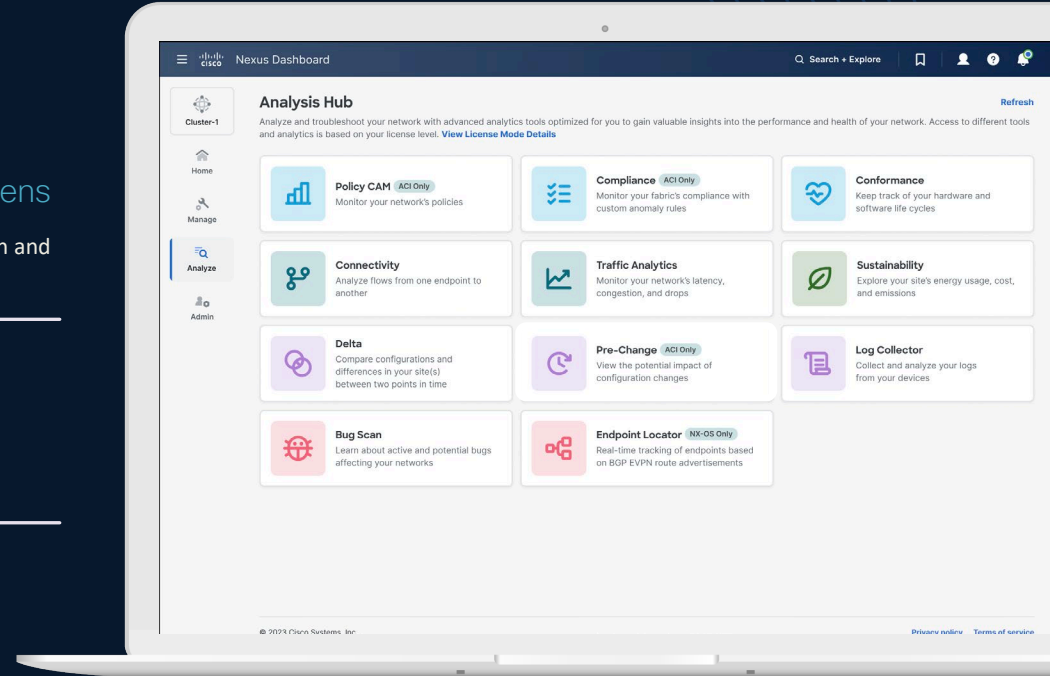
Simulate the impact of change before it happens

Simulate adding, removing, or modifying your fabrics' configuration and learn if existing anomalies will clear or if new ones will arise



Maintain compliance

Know beforehand if changes in your network will violate any configuration or communication compliance rules



Benefits

Execute changes
with confidence

Evaluate potential
change implications

¹ ACI only, NX-OS support expected with version 4.1 (not committed)

Nexus Dashboard – Connectivity Analysis

ND 4.0 -> 4.1

Endpoint to Endpoint Analysis



Evolve from traceroute

Add a source and destination IP or MAC address and let Nexus Dashboard do the rest. Optionally add L4 ports to be more specific



Increased visibility

Add forwarding (ELAM and OAM) and consistency checks, routing verification, traffic analytics and security¹

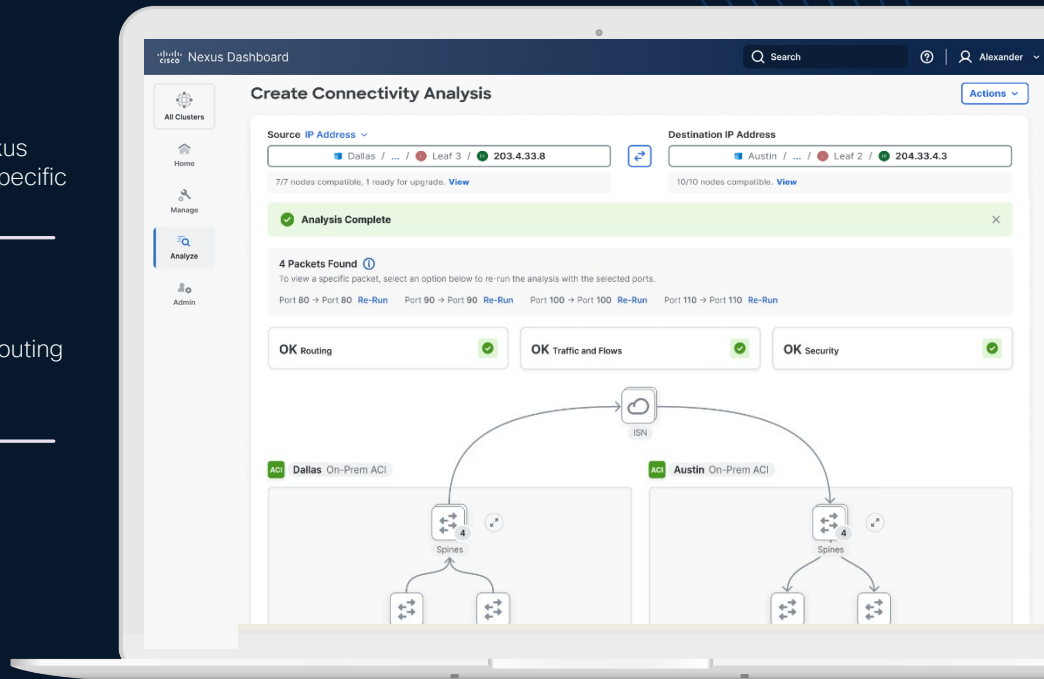


Single and multi-fabric²

Visualize the whole path even when endpoints communicate through different Nexus-Dashboard-managed fabrics

¹ Security and Traffic Analytics views are expected to be integrated in version 4.1 (not committed)

² Must run Nexus Dashboard 4.0 and above



Benefits

Learn about any
in-transit issues

Avoid isolated
troubleshooting

Cisco Nexus Dashboard – Search and Explore

ND 4.0 -> 4.1

Ask Questions -> Get Answers



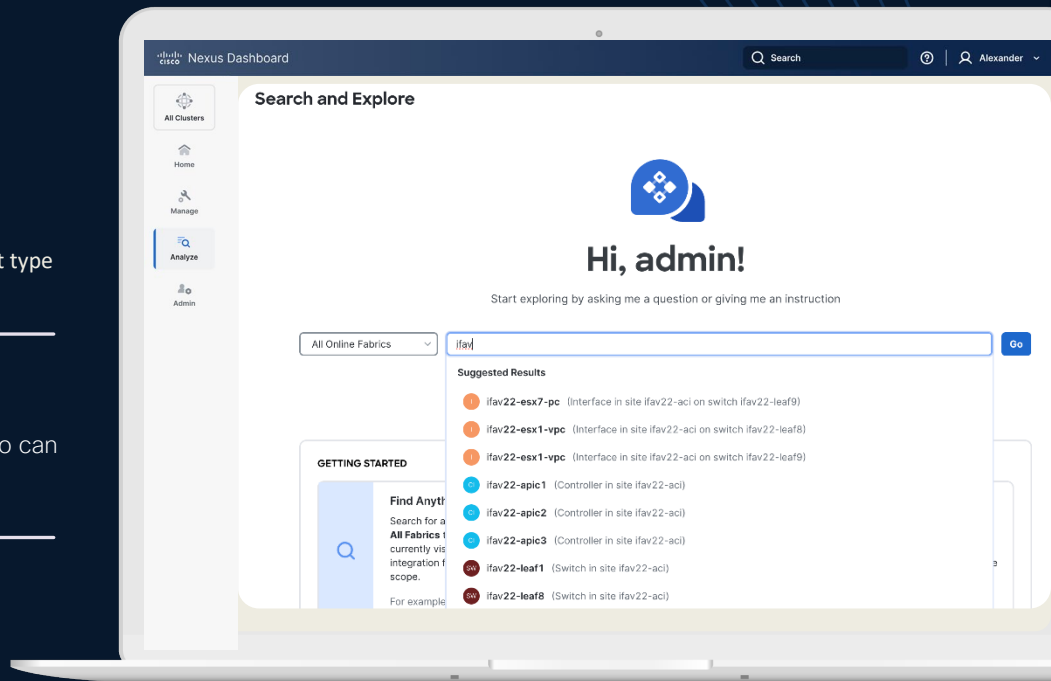
Search for anything across fabrics

Endpoints, IP addresses, switch names, configuration elements. Just type the first 3 characters and auto-suggestions will come up



Ask questions, get answers

Explore what interfaces are associated with a given VRF, who can talk to whom and how¹, and many others useful tasks.



* *Can* queries are only supported on ACI, expected NX-OS support by 4.1 (not committed)

Benefits

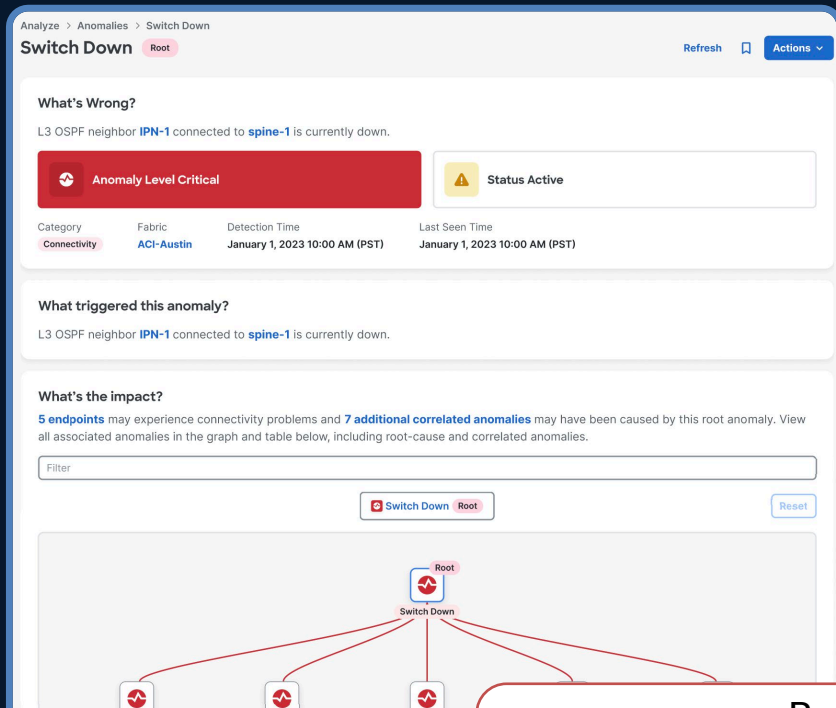
Find what you
need faster

Centralize
troubleshooting

Cisco Nexus Dashboard – Anomaly Correlation

Simplify path to Root Cause

ND 4.0 -> 4.1



Hardware process crashes

Fan tray, power supply failures

Interface utilization, microburst

DOM, CRC, FCS

Hardware temperature, memory

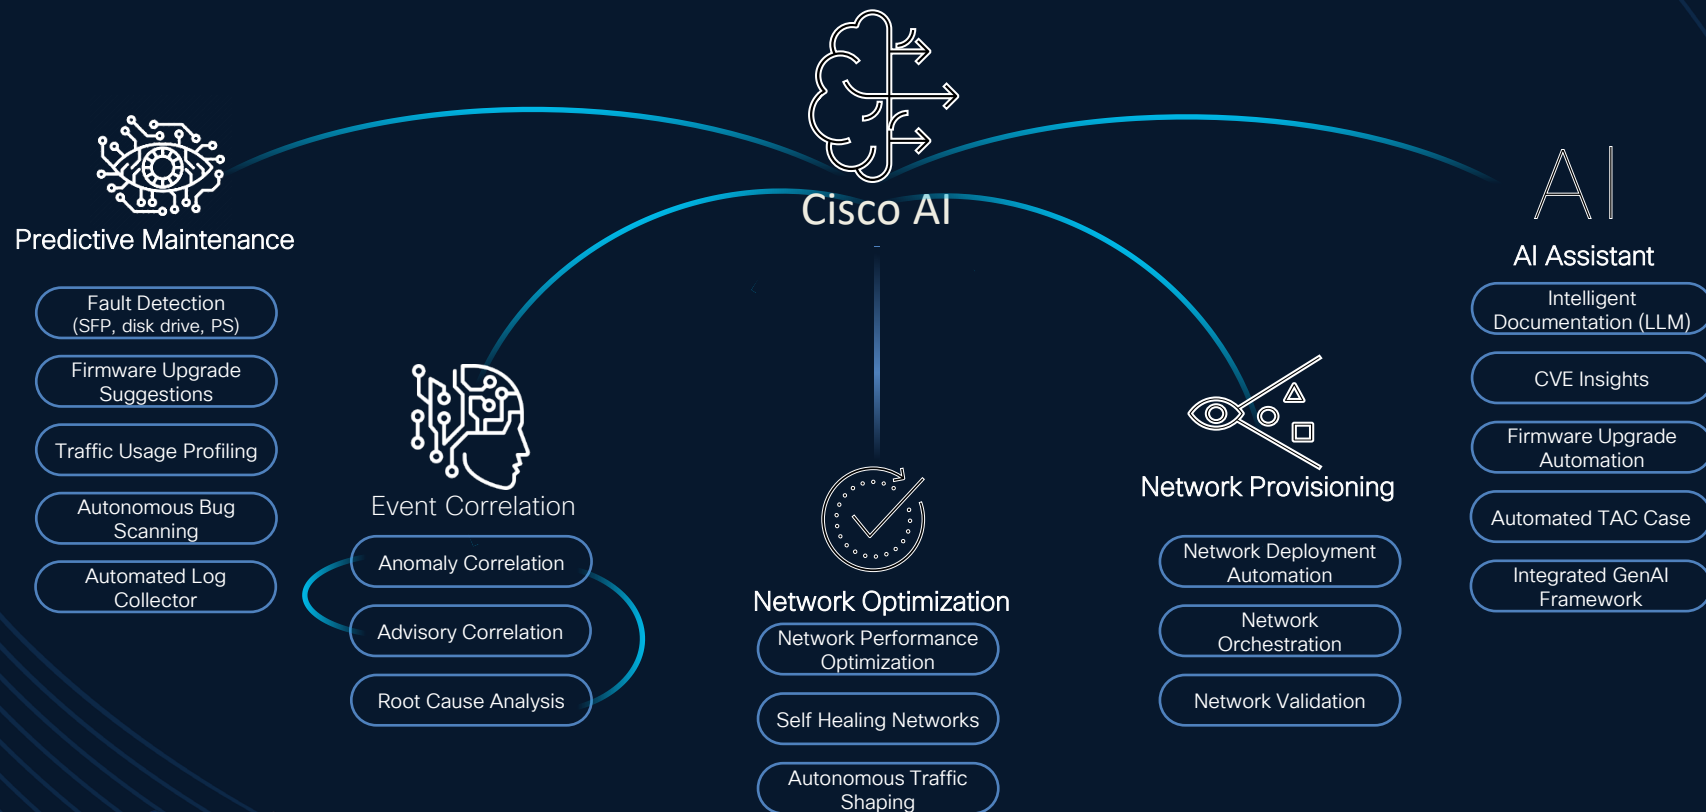
Benefits

Identify the root cause anomaly to reduce manual troubleshooting

Cisco Nexus Dashboard


Cisco AI

ND 4.1 -> Future



Optimizing and Assure Network Infrastructure for AI Workloads

Why optimize? AI traffic is moving the needle...



Meta, @Scale conference, September 2024

Meta backbone traffic grew >30% YoY for the last 4 years

AI-driven traffic grew faster and overtook non-AI traffic

GPT-4 has 1.76 trillion parameters, a massive increase from GPT-3's 175 billion parameters, requiring exponentially more compute and networking resources.

AI workloads require ultra-low latency (sub-100 microseconds) and high throughput for efficient processing.



Listening to the market

Gen AI is expected to be the

#1 driver

of infrastructure investments
over the next 18 months

McKinsey reports

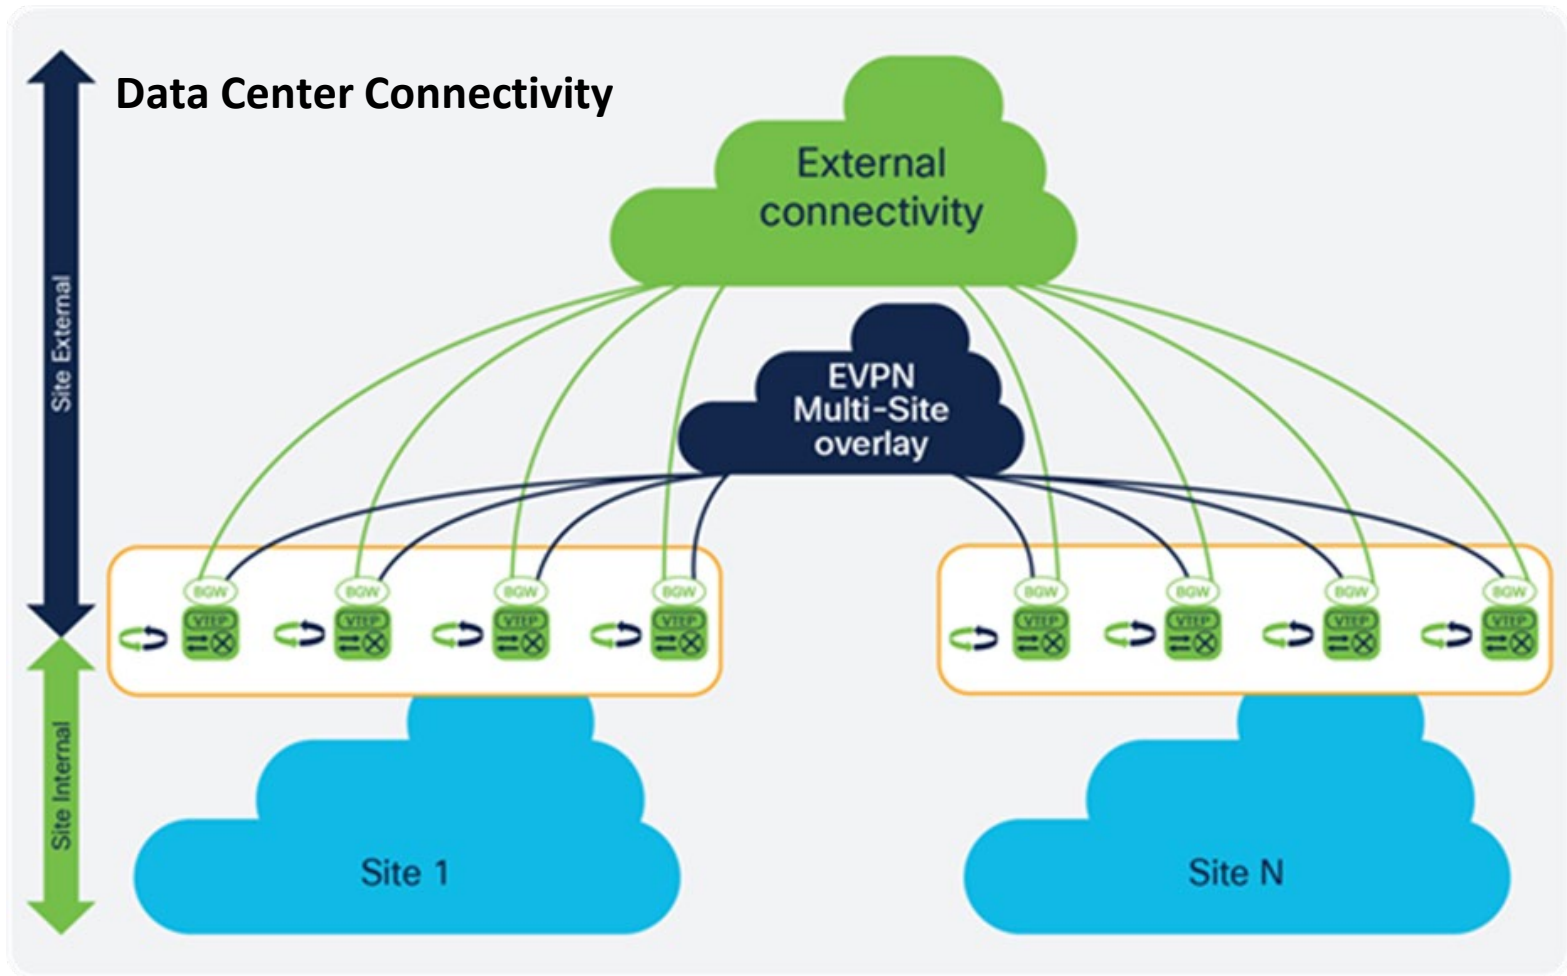
75%

of large enterprises will rely on
AI-infused processes by 2026

Lightreading says there was

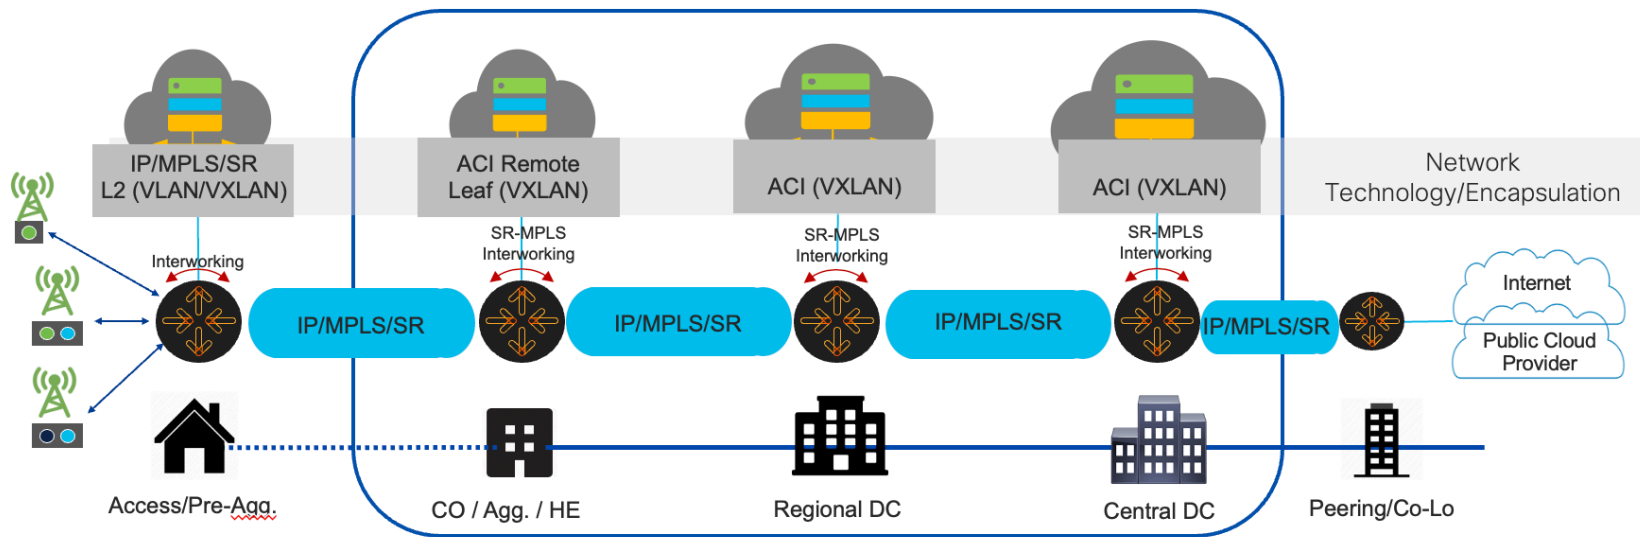
29%

Growth in DCI traffic in
2024



Monitoring DC to XYZ flows in transport network

- Using ACI or VXLAN EVPN to L2/L3/MPLS or SR handoff allows the transport team to monitor the DCs to XYZ flows using Provider Connectivity Assurance



The Need for Service- Centric Assurance

Customers expect a great experience as network complexity and costs increase.

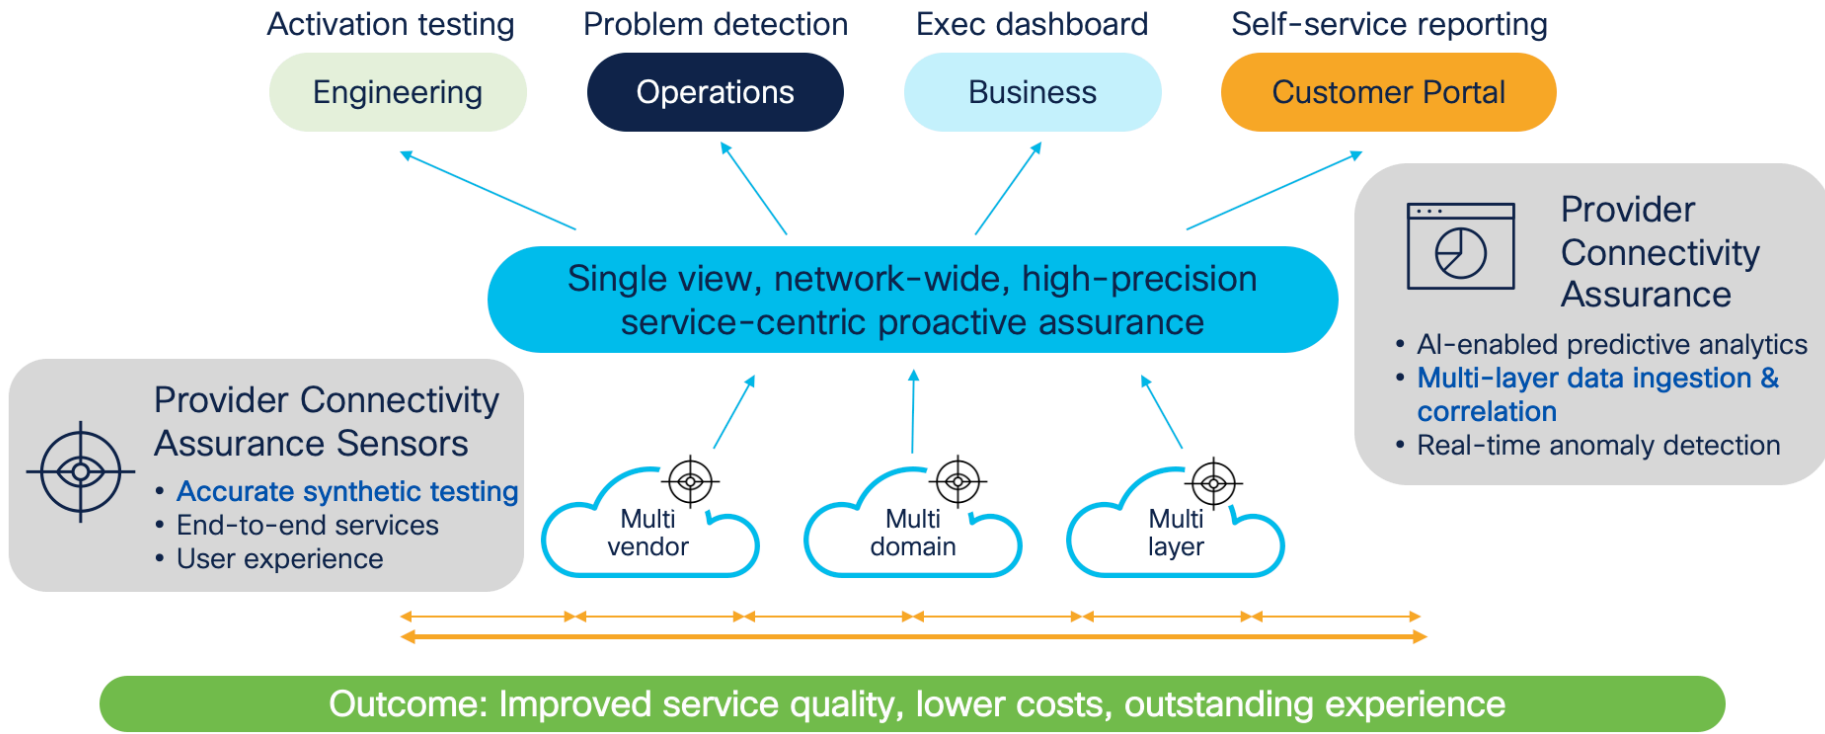
- Higher demands on networks and services

Complex and distributed network; fragmented tools

- Performance visibility silos

Introducing Cisco Provider Connectivity Assurance

Proactive assurance with real-time visibility and unparalleled accuracy for consistently great digital experiences



Assurance Sensor capabilities

Robust feature set enables critical network service quality measurements



FPGA powered hardware



Sensor Agents



Sensor Capture VM

End-to-end service and QoE assurance

Synthetic (active) multi-layer testing
Predictive/ultra-accurate
(e.g., jitter, bandwidth, latency)

Multi-layer real user experience
Service and application experience

Contextual performance assurance analytics

PM

Synthetic (active) multi-layer testing

- L2/L3: TWAMP, Y.1731, UDP/ICMP echo
μ-second precision, 1 second reporting
- L4/L7: TCP, HTTP, FTP, DNS, Trace

SAT

Service activation testing

RFC2544; ITU-T Y.1564, RFC6349

BM

Bandwidth metering

Microburst detection, 1ms sampling

RUM

Multi-layer real user experience (capture)

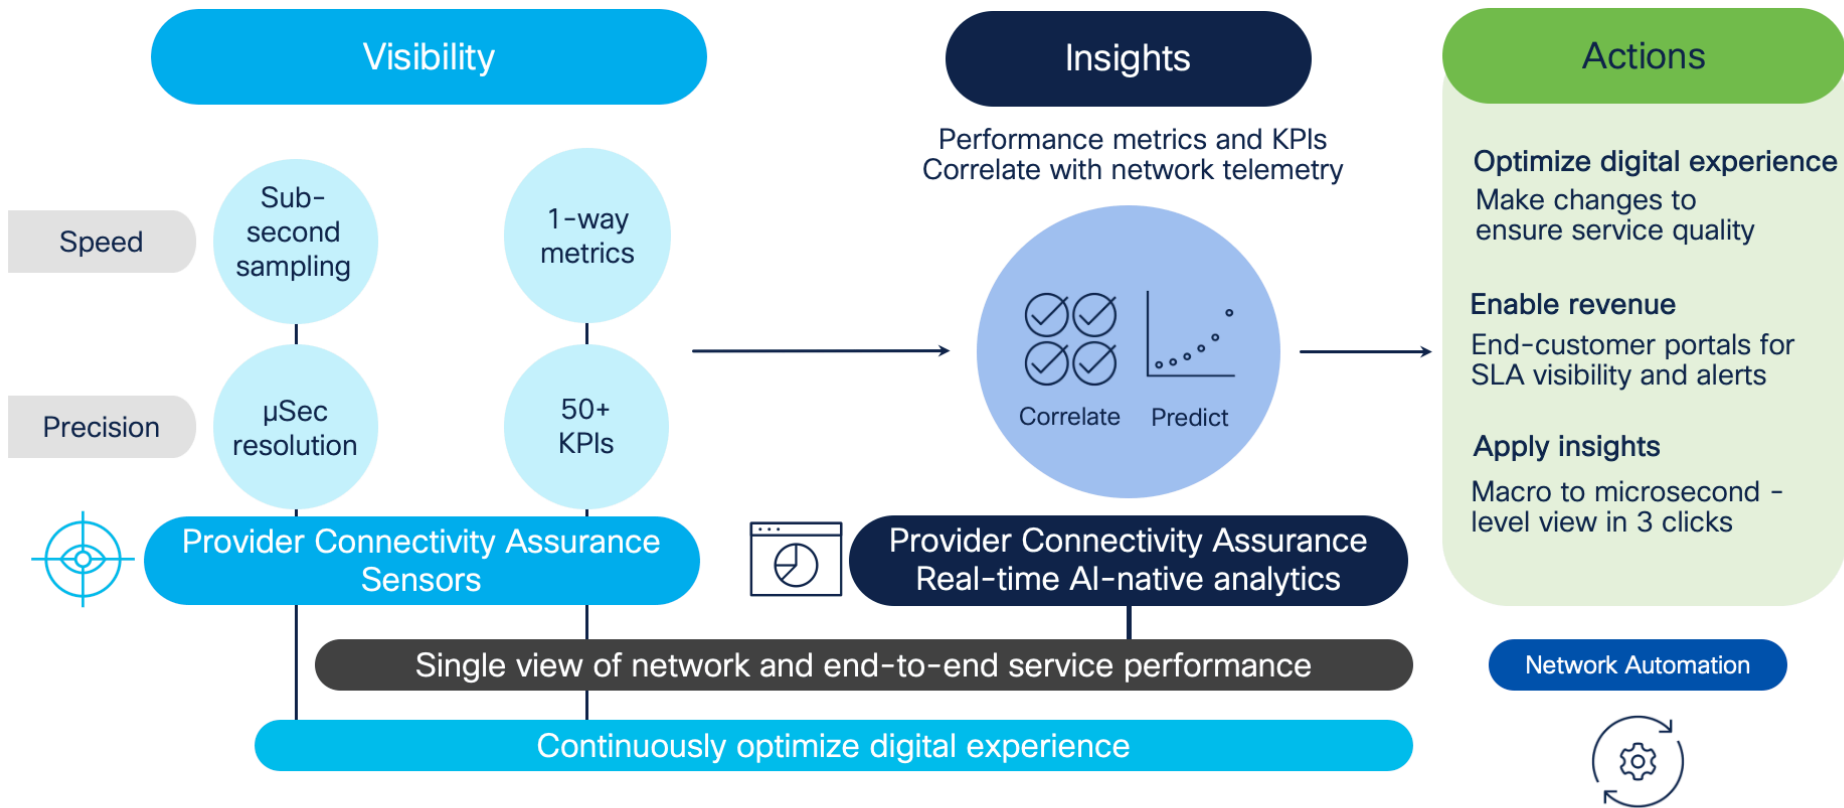
Service and application experience

ANA


Reporting, visualization and AI-native analytics

Multi-vendor, multi-layer, multi-persona

Service-centric visibility for proactive operations




Customer perception is measured in milliseconds, not minutes



Everything looks fine! Latency is below threshold

Perception

Service is performing according to SLA/SLO



But in reality

⚠ Customer is impacted!
SLA is breached

Service is impacted by intermittent latency issues

Data center network connectivity & service assurance

Challenge

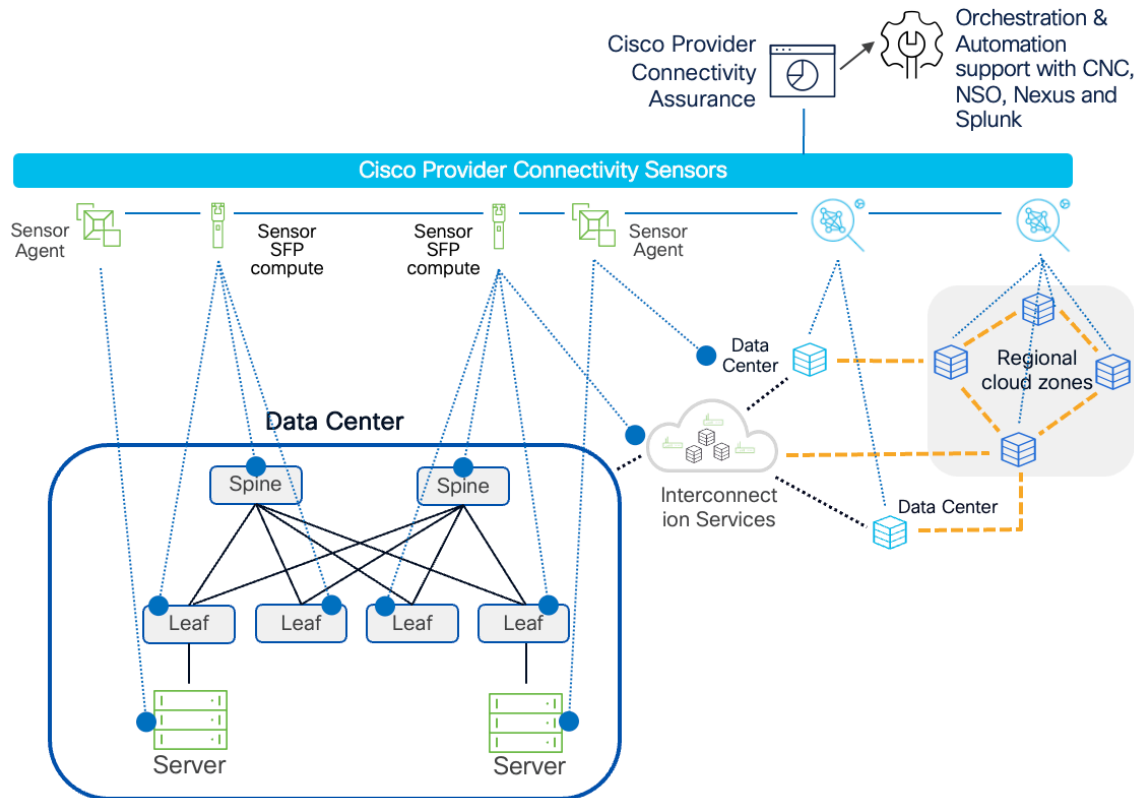
- Data centers are at the heart of today's digital infrastructures, supporting a myriad of applications and services. When problems arise, the impact can be crippling.
- The ability to detect and address performance bottlenecks, hardware failures, and network issues is key to safeguarding the functionality, reliability, and efficiency of DCs.

Solution

- Cisco Provider Connectivity Assurance enables data center providers to gain full visibility into performance across and/or within their owned or leased estate of DCs.
- Offers immediate benefits with proactive performance monitoring (latency, packet loss, and jitter) across inter DC, intra-DC, and cloud environments.
- Proof of innocence tool to determine if internal application performance issues are related to the network.

Outcomes

- Offers operators per-DC or end-to-end service
- Allows operators to monitor and identify issues before they can impact SLAs, either across the WAN or at the application layer.
- Enables operators to effectively troubleshoot and plan upgrades (or switching to alternative 3rd party providers in a leased scenario) - offering OpEx and CapEx benefits.



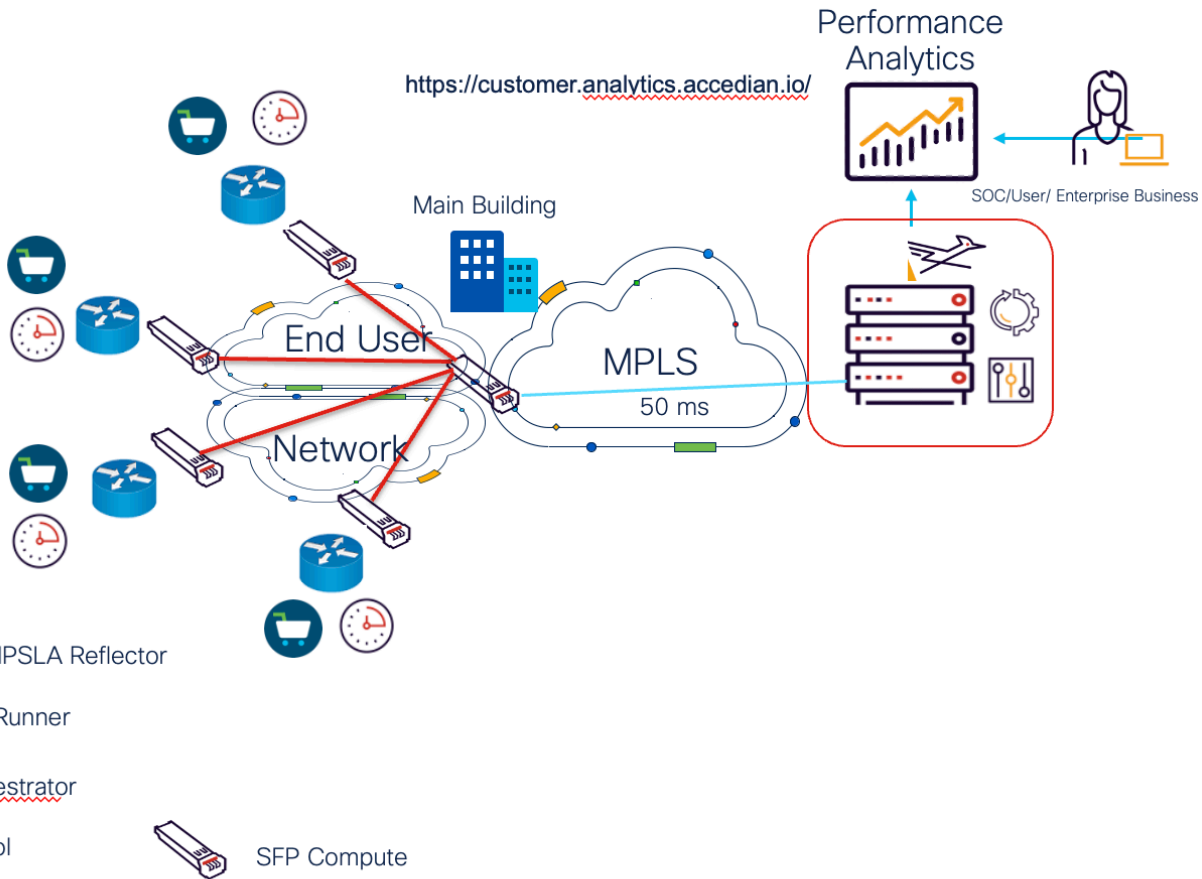
Performance Monitoring

Monitor

End-to-end end user customer network with metrics of Jitter, Delay, PL and BW with a central multi-user dashboard

Identify

Impact of BW consuption, PL, Jitter, Delay by branch
office, topology changes, service degradation, cor
map colors, macro to micro detail by circuit user



Bandwidth metering

Detecting microbursts

Benefits


- Detect microbursts: Throughput max, avg with 1 second reporting
- Validates capacity and traffic engineering
- Business benefits:
 - Usage reporting via end-customer portal
 - Upsell opportunities for extra bandwidth
- Proactively monitor, assess and diagnose issues when correlating with continuous PM

Measurements

- Continuous in-service monitoring of bandwidth utilization
- Supports L2-L4 filters enabling monitoring by IP, protocol, port number, CoS, VLAN, etc..
- Ultra granular real time monitoring with sampling every 1/10/100/1000ms and reporting per second
- Monitor up to 14 flows per direction (down/up)

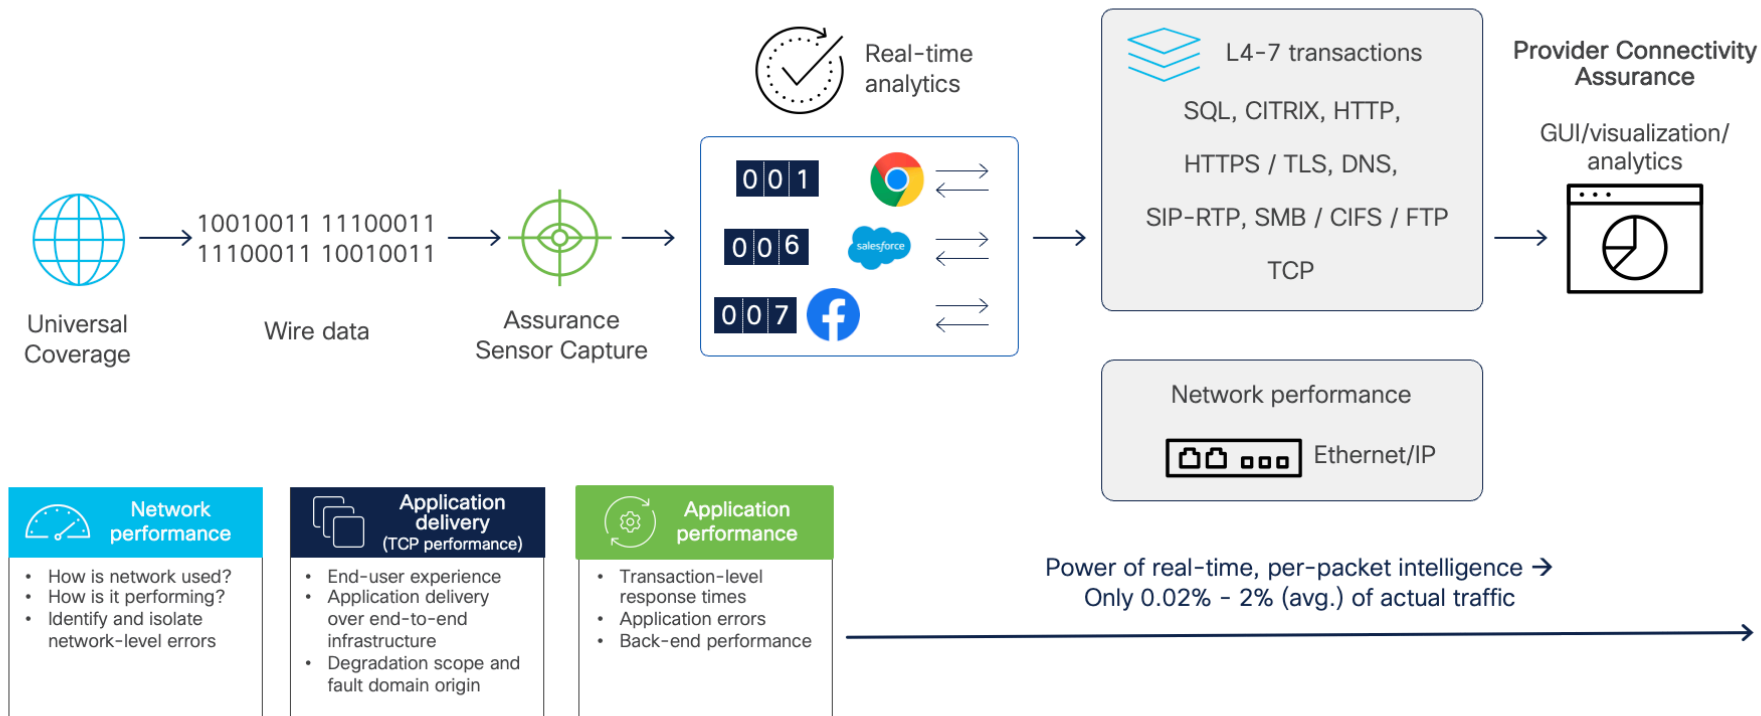


This information can lead to adequate buffer sizing which can result in either latency or packet loss when traffic increases



Sensor Capture: real user experience monitoring


Real-time, full stack visibility, down to each session and transaction



Recap




Nexus Dashboard and Insights driving to simplify Operations and Visibility




Streaming Telemetry + Insights leads to Actions – and Root Cause



PCA allows your team to monitor traffic from a passive and synthetic perspective



Analyze and forecast traffic patterns and behaviors



Granular msec monitoring



Thank you