Ethernet Fabrics for AI Clusters

Powered by Cisco SiliconOne or Nexus for Ultra High Performance, Scalable and Non-Blocking Ethernet Fabrics

S. John Banner - Technical Solutions Architect
Lieven Colman - Site Reliability Engineer

BRKCOC-3005

Webex App

Questions?

Use the Webex app to chat with the speaker after the session

How

1. Find this session in the Cisco Events mobile app
2. Click “Join the Discussion”
3. Install the Webex app or go directly to the Webex space
4. Enter messages/questions in the Webex space

Webex spaces will be moderated by the speaker until February 28, 2025.

About your Speaker

- S. John Banner
- Solutions Architect
- Based out of Porto, Portugal
- Cisco employee since 2000
- 35 years in IT industry
- Focus on DC Networking

How I spend my free time:

- Reading
- Exploring Portuguese Culture
- Spending time with Family and Friends

About your Speaker

- Lieven Colman
- Site Reliability Engineer
- Based out of Brussels, Belgium
- Cisco employee since 2021
- 17+ years with Cisco IT services
- Focus on DC Networking

How I spend my free time:

- Competitive padel fanatic and casual cycler
- Extensive culinary time with friends
- Amateur binge watcher and gamer
- Explore domestic creativity

Agenda

- Importance of AI to Cisco
- The Business Priority
- Architecture and Design
- Implementation
- Day 2 Support
- Lessons Learned
- What’s Next?

Why is AI Important?

- Changing the world
- It’s everywhere
- New market opportunities
- Improving our products
- Empower our services

Business Priority

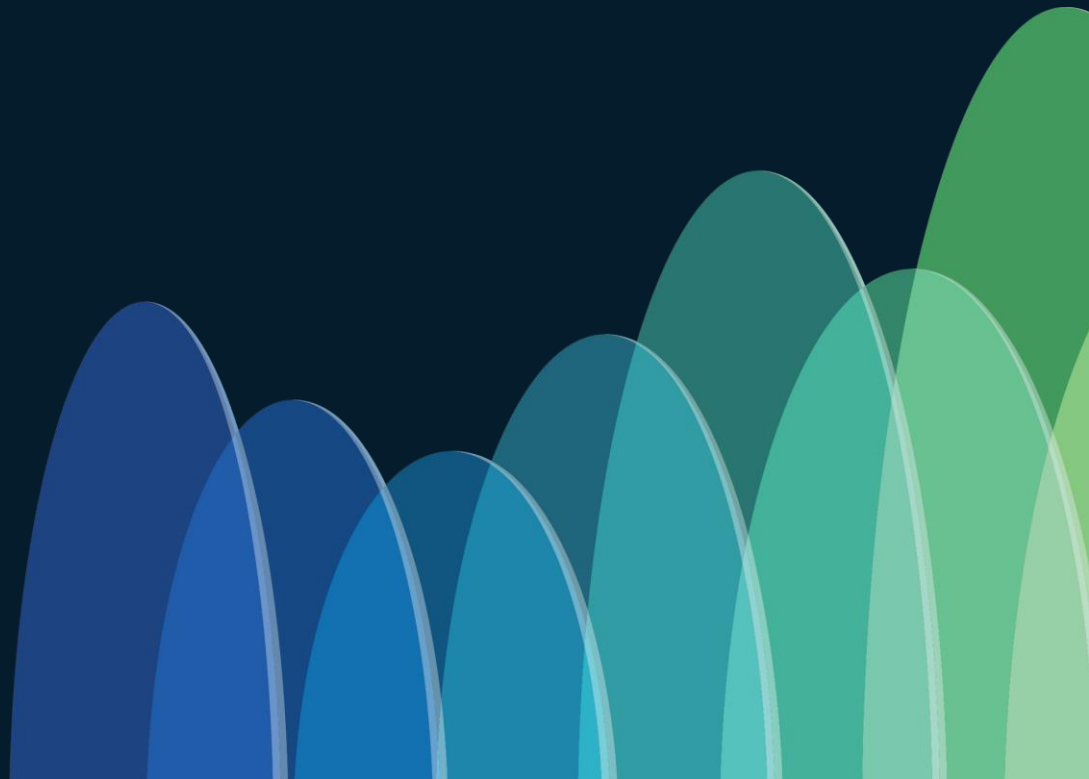

Ask from Leadership

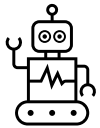

- Deployment of AI Cluster with 256 Nvidia H100 GPUs
- Blueprint for Large-Scale AI Cluster with Cisco Ethernet in Partnership with Common Hardware Group
- Deployment of AI Cluster with 512 H200 GPUs for Business Applications
- Deployment of AI Cluster with 92 H200 GPUs in UCS885s for Cisco Internal Cloud

Deployment Objectives

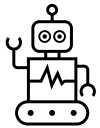

- Validate AI Infra design with Cisco Technology stack
- Support Engineering & Business use cases for AI model training
- Enable LLM use cases across Cisco products
- Enable AI Applications for Business Use Cases
- Offer GPU as a service

Timeline
Key Milestones

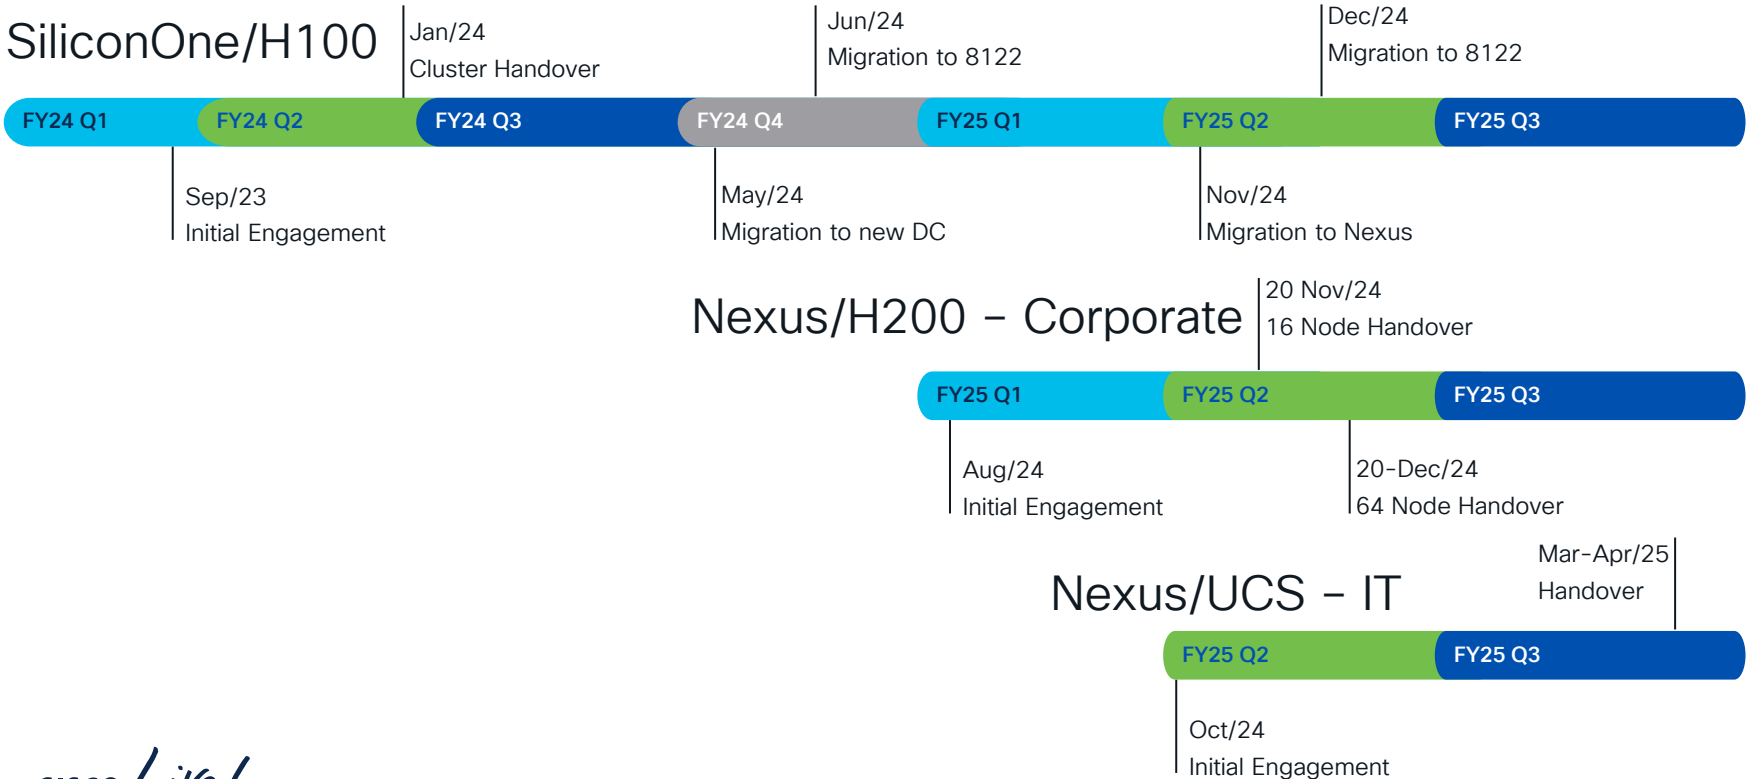

Getting Started
The clock is ticking

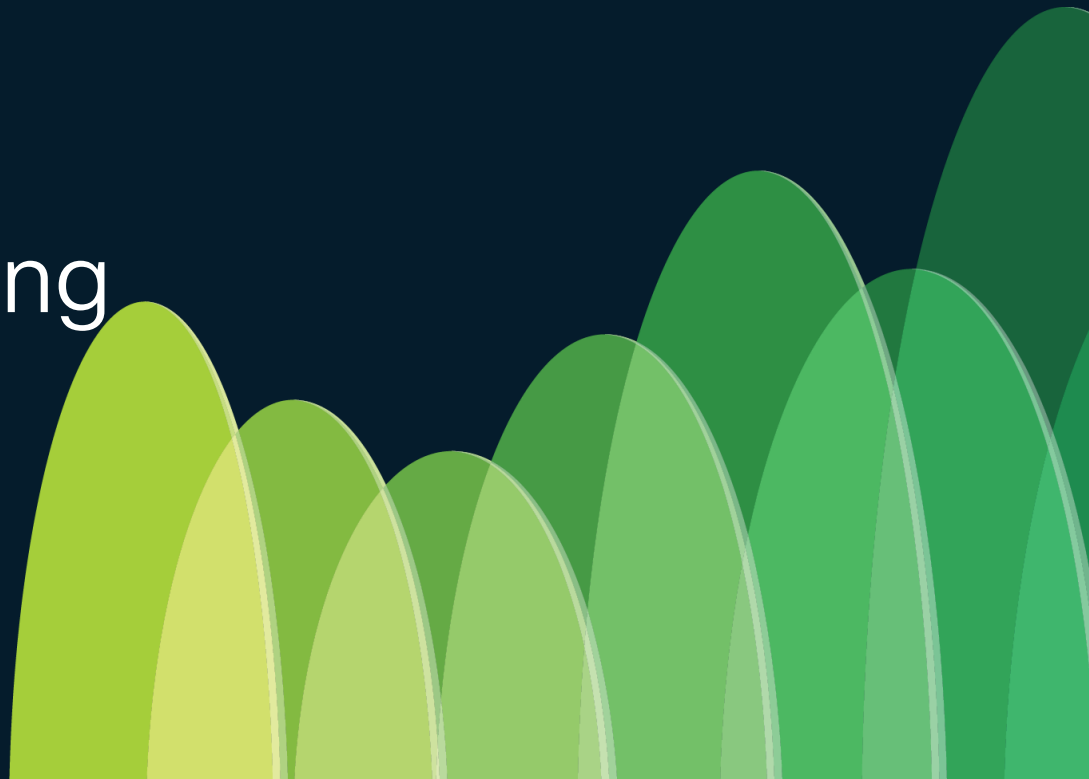

cisco Live!

Where We Began

Learning and Upskilling the Team

- Internal/External Research
- DevNet
- Training VODs/Classes
- Procured Lab Test Devices
- Partnerships with Cisco and NVIDIA Engineering Leadership
- Identified SMEs to Learn and Train on Design and Architecture Components

AI Cluster Fabrics
Scalable Dedicated Environments

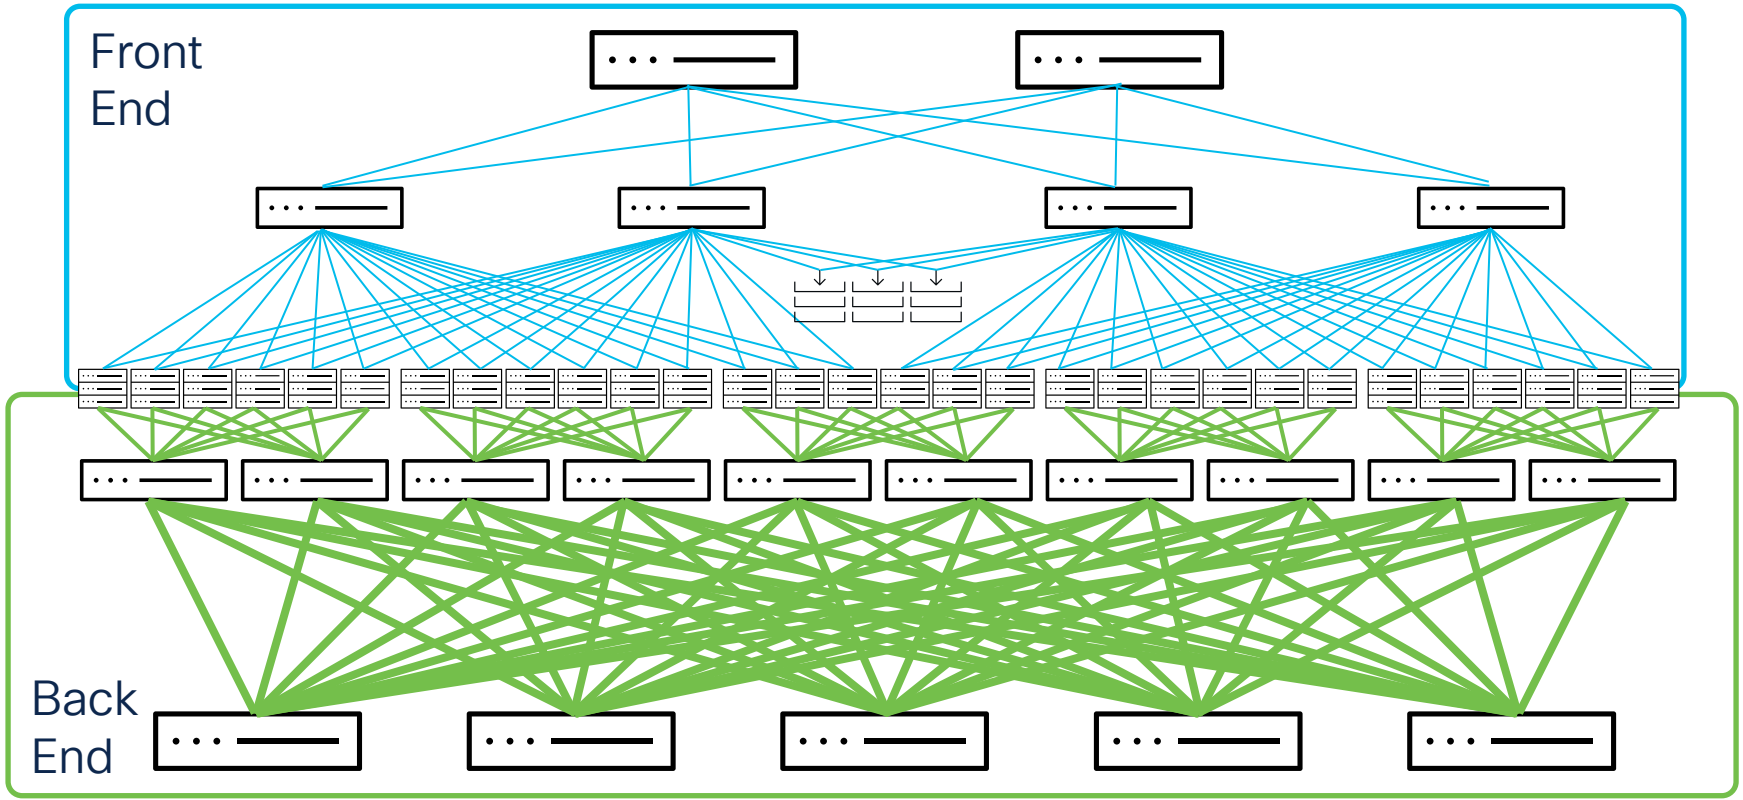

Non-Blocking Lossless Fabric
High Performance Network That Doesn't Drop Packets

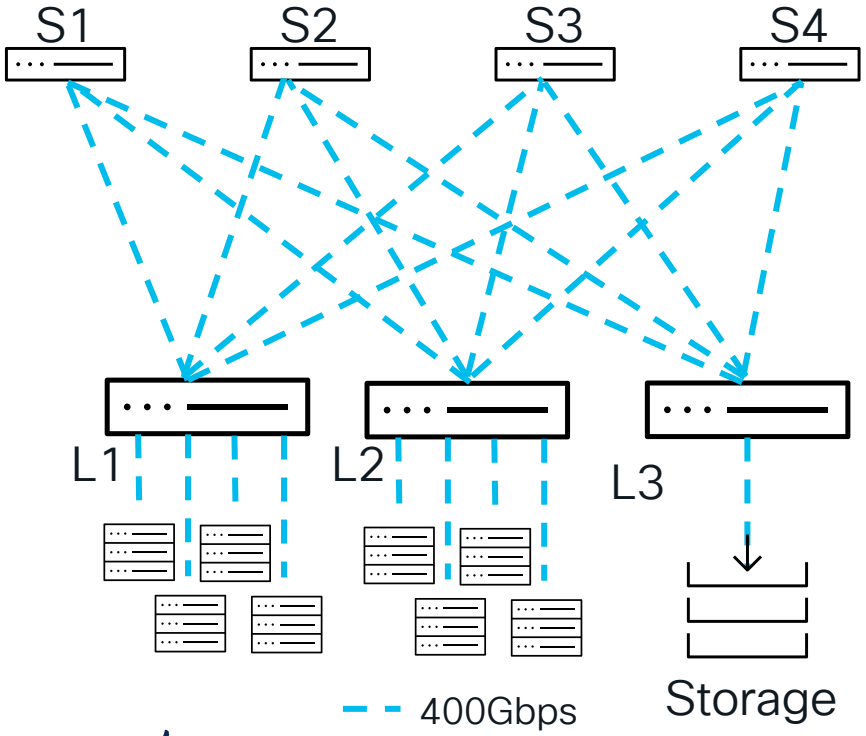

- Full bandwidth host communication
- Equal bandwidth from leaf to host and spine
- Congestion management to prevent packet loss

AI Cluster Architecture and Design

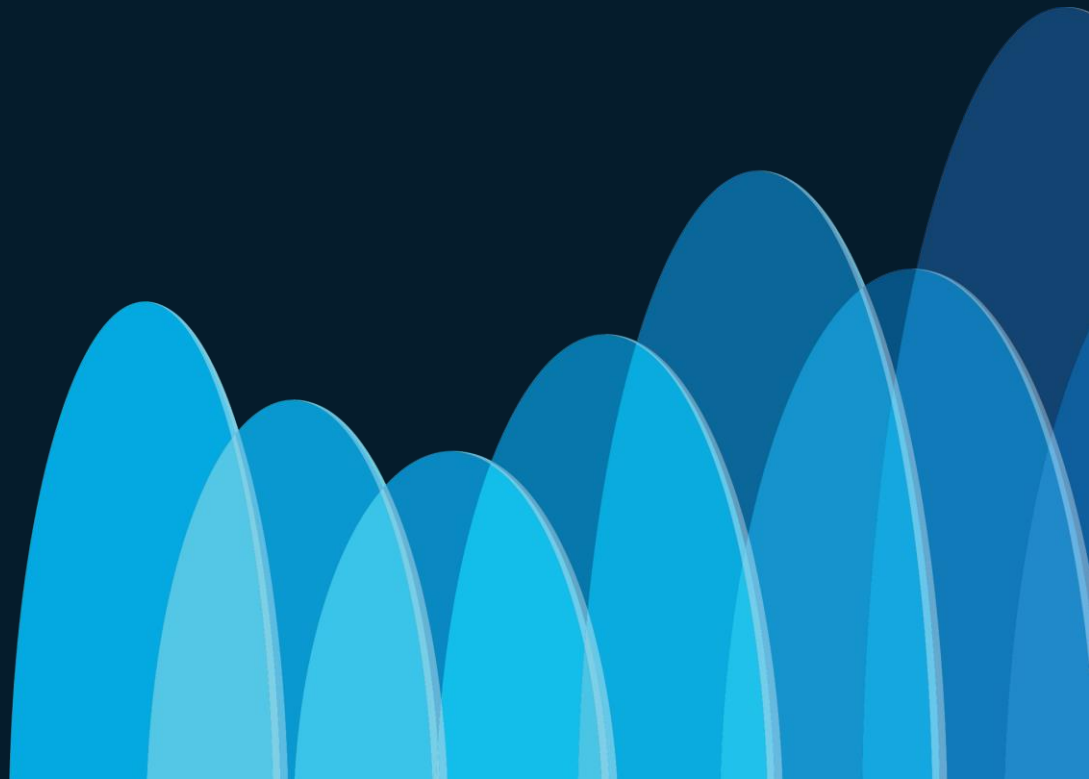

cisco Live!

Foundational Principles for our AI Clusters

- Focus on application performance
- Intelligent buffering, low latency, telemetry / visibility
- Dynamic congestion avoidance for various workloads
- Dedicated front-end / back-end networks. Non-blocking fabric
- Leverage Automation for Day 0 to 2 operations

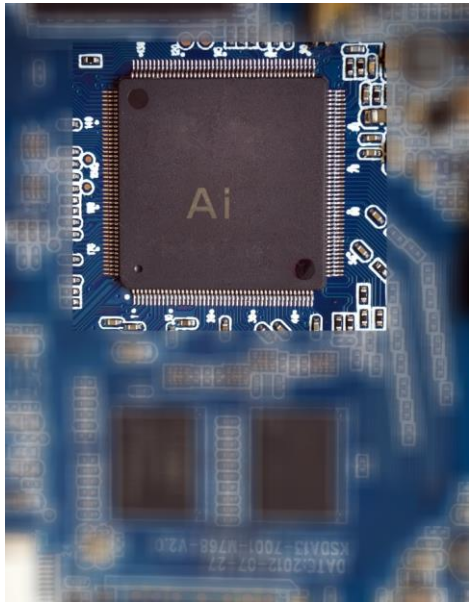

Network Zones

Original Design

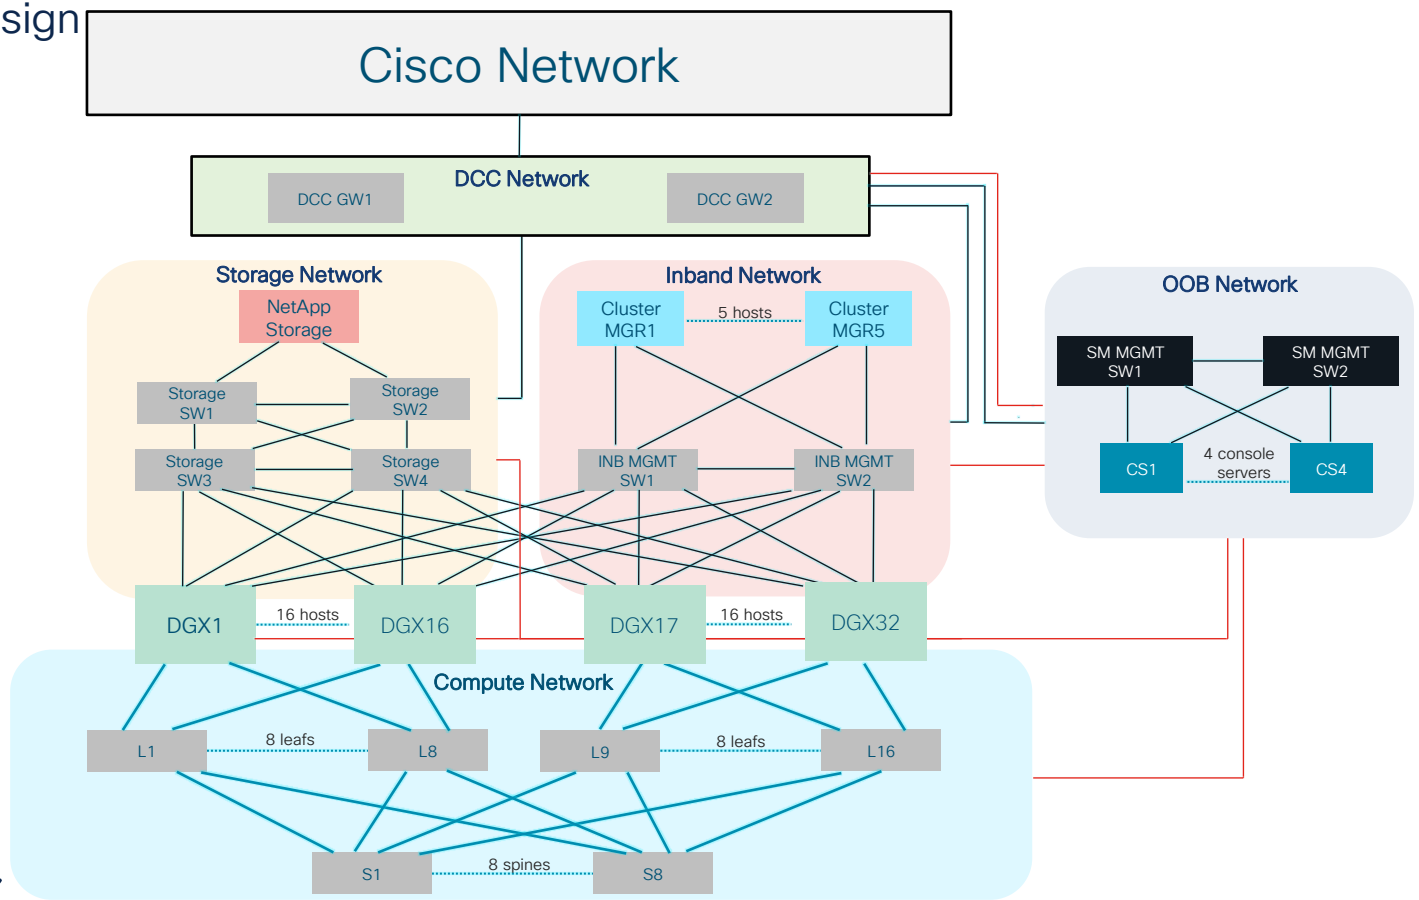

Inband Network

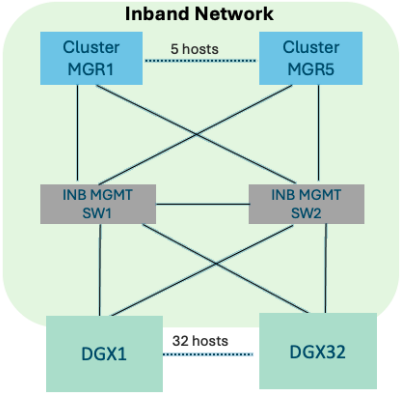

- Clients to access the AI cluster and schedule jobs
- 5 Management UCS hosts
- Simple Linux Utility for Resource Management (SLURM)
- Layer 2 between the DGX and UCS Managers
- Inband MGMT switches connected to Cisco Network
- Cisco 8102-64H-O switch,100Gig ports

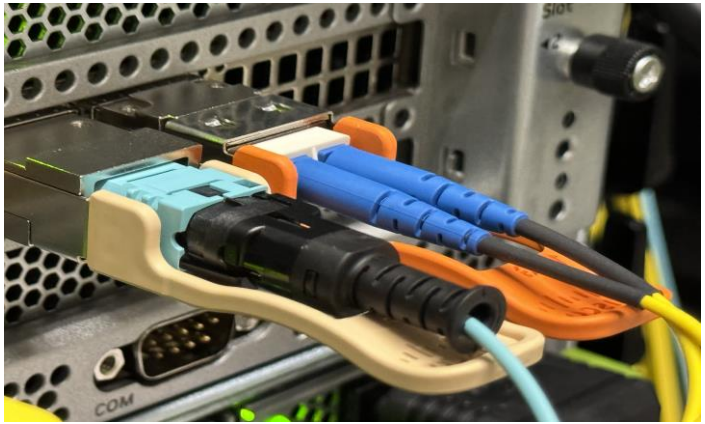

Storage Network

- Client access to periodic data snapshots
- Netapp Storage, A900 series
- Layer 2 network, 400Gig Links between storage switches
- 100Gig links to Netapp and DGX
- Cisco 8101-32FH-O

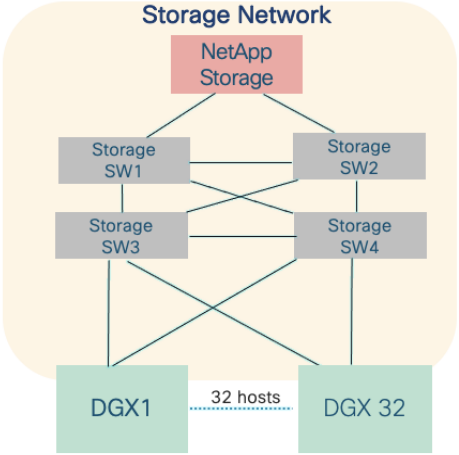

Future – Build L3 underlay with VxLAN overlay

Collapsed In-Band and Storage Network

New Design on Nexus

- Client and Data Access
- 8 Management UCS Hosts
- Netapp Storage, A900 series
- VXLan over BGP, 400Gig Links between Leaves and 4 Spines
- 100Gig links to Netapp, DGX and UCS
- Nexus N9K-C9364D-GX2A and Nexus N9K-C93600CD-GX
- Entirely built using NDFC

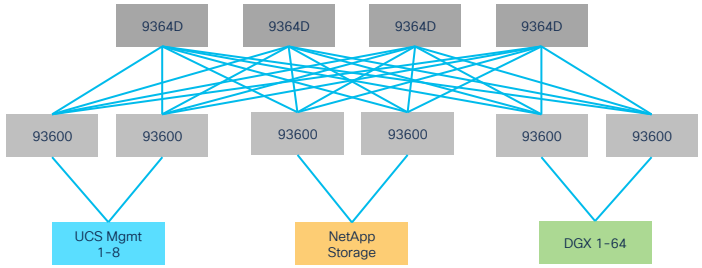

Compute Network

Original Design

- High-speed data transfer between GPUs
- RAIL optimized design with L3 connections
- 2x400Gig MLAG Leaf to Spine connectivity
- iBGP with Route Reflectors
- 400Gig non-blocking bandwidth per GPU
- Cisco 8101-32FH-O
  - Then 8122-64EH-O
  - Then N9K-C9364D-GX2A
- QoS with ECN and PFC
  - Then Distributed Scheduled Fabric
  - Then ECN, PFC and DLB

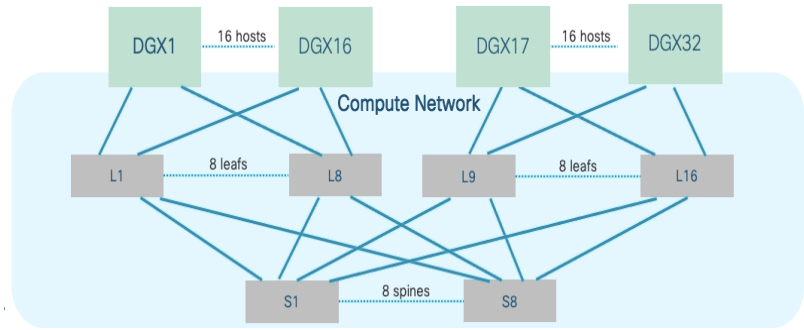

Cisco IT's Original AI Cluster

- First fabric built on Silicon One ASICs
- Utilizing SONiC as the NOS
- Ethernet-based architecture
- 400G non-blocking compute fabric
- 32 Nvidia compute nodes (8 GPUs each)
- 12.8T throughput per Compute leaf
- 32 PetaFLOPS AI performance per node
- 4 NetApp AFF A900 nodes

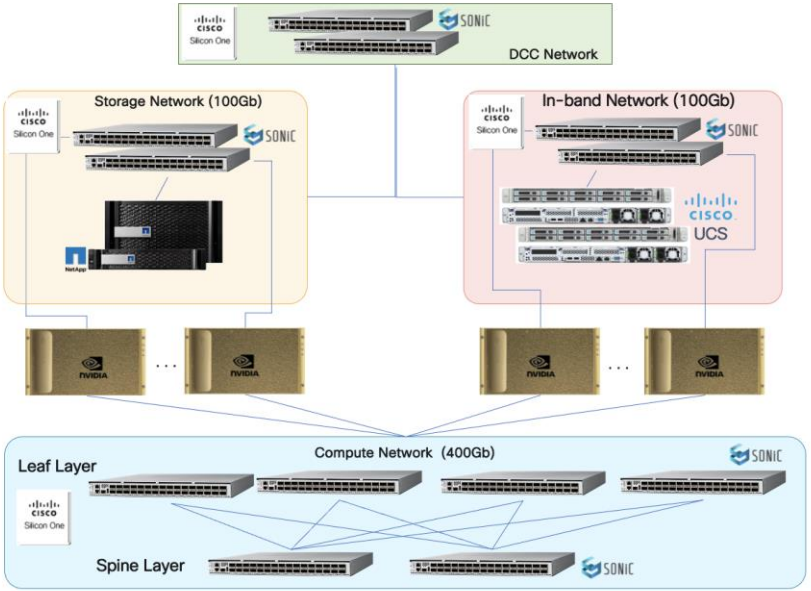

Cisco IT's Second AI Cluster

- Utilizing NX-OS as the NOS
- Our first VXLan DC Fabrics
- NDFC and NDI for Fabric Build and Management
- Ethernet-based architecture
- 400G non-blocking compute fabric
- 64 Nvidia compute nodes (8 GPUs each)
- 51.2T throughput per Back-End leaf
- 32 PetaFLOPS AI performance per node
- 4 NetApp AFF A900 nodes

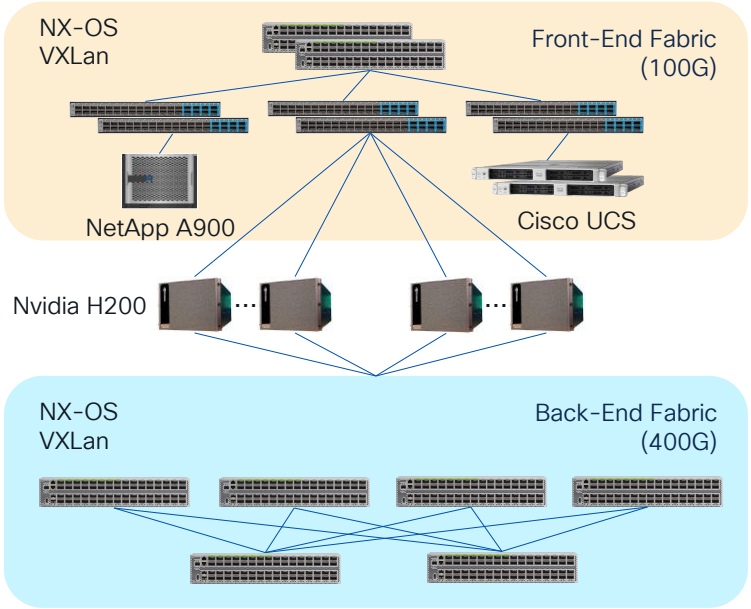

Original Infrastructure Build

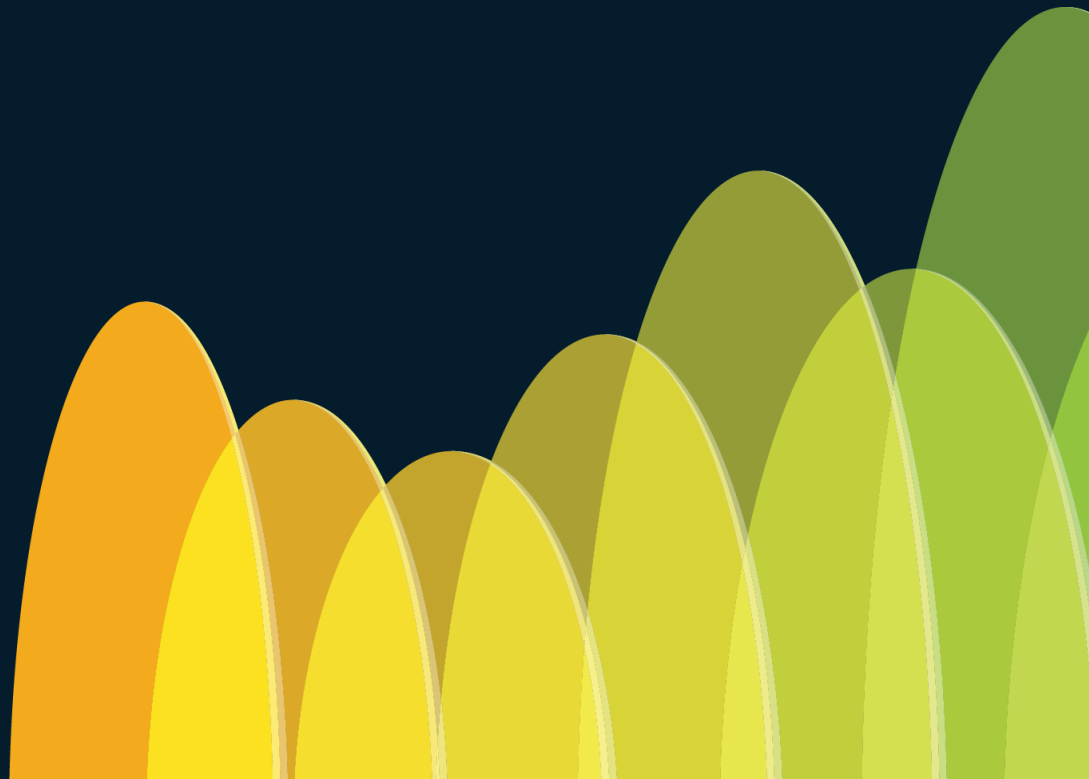

cisco Live!

Day1 Implementation

- Datacenter Layout
- Rack, Stack and Patch
- Bootup and Install SONiC
- Network Configuration

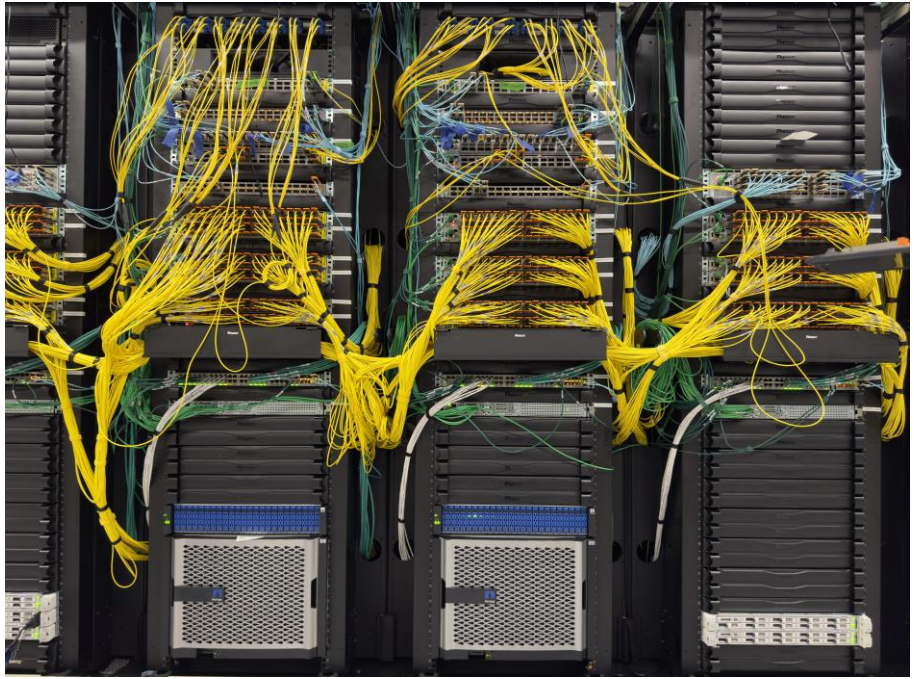

Data Center Cabinet Layout

- Utilized existing cabinets, power, and cooling
- 32 DGX Units distributed for high power and cooling (10.2 kW per DGX)
- Direct connections with long patch cords to save time and cost
- No major power or cooling upgrades needed

Data Center Cabinet Elevation Layout

Network Row Optimization

- Centralized network and storage hardware
- Reduced cabling and patching complexity
- Ensured failover and power redundancy

Compute Row Optimization

- Distributed DGX hardware (~10 kW per cabinet)
  - Newer DCs optimized for 25kW (2 Nvidia) per cabinet
- Maintained N+1 power and cooling redundancy
- Reduced cost and deployment time

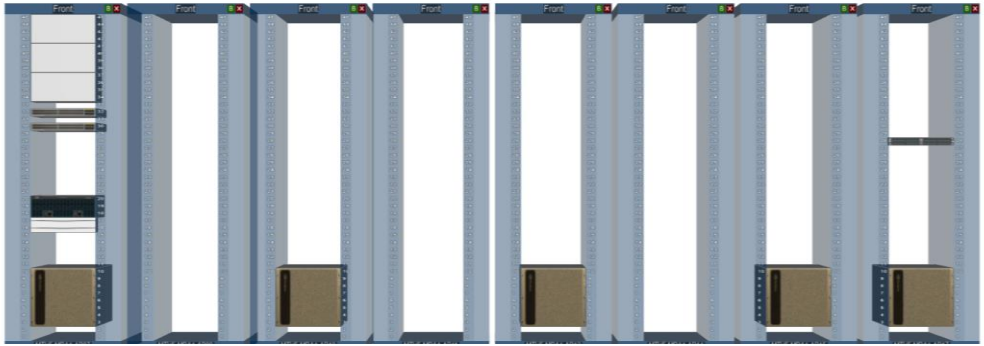

Optics for Original AI Cluster

	 QDD-400G-DR4-S	 QSFP-100G-SR4-S	 QSFP-100G-SM-SR	 MMS4X00-NS-FLT
Vendor	Cisco	Cisco	Cisco	Nvidia
Speed	400G	100G	100G	800G Duplex 400G-DR4
Fibre Mode	Single Mode	Multimode	Single Mode	Single Mode
Max Cable Distance	500m	70m	2000m	100m
Fibre Connector Type	MTP/MPO-12	MTP/MPO-12	Duplex LC	Duplex MPO-12
Use Case	Compute Fabric (Leafs/Spines)	Storage and Inband Mgmt	DCC	Compute Fabric

Second AI Cluster Added

	QDD-400G-FR4-S	QSFP-100G-DR-S	QSFP-100G-FR-S
Vendor	Cisco	Cisco	Cisco
Speed	400G	100G	100G
Fibre Mode	Single Mode	Single Mode	Single Mode
Max Cable Distance	2000m	500m	2000m
Fibre Connector Type	Duplex LC	Duplex LC	Duplex LC
Use Case	Leaf to Spine links	UCS Mgmt Nodes	Nvidia Storage and In-Band Links

Cisco Network Gear for Original AI Cluster

	 8101-32FH-O	 8102-64H-O	 93108TC-FX3H
ASIC	Silicon One	Silicon One	Cloud Scale
Rack Units	1RU	2RU	1RU
NOS	SONiC	SONiC	NX-OS
Ports	32xQDD	64xQSFP28	48x10GBASE-T
Total Throughput	12.8 Tbps	6.4 Tbps	1.8 Tbps
Use Case	Compute Fabric (leafs and spines) Storage Network	In-band Connectivity	OOB Mgmt
Port Speed	Four Hundred (FH)- 400Gb/s	Hundred(H)- 100 Gb/s	1/10Gb/s

Cisco Network Gear for Second AI Cluster

	 N9K-C9364D-GX2A	 N9K-C93600CD-GX	 N9K-C93108TC-FX3H
ASIC	Cloud Scale	Cloud Scale	Cloud Scale
Rack Units	2RU	1RU	1RU
NOS	NX-OS	NX-OS	NX-OS
Ports	64xQSFP-DD	28xQSFP28, 8xQSFP-DD	48x10GBASE-T
Total Throughput	25.6 Tbps	12 Tbps	1.8 Tbps
Use Case	Back-End Fabric (leafs and spines) Front-End Spines	Front-End Leaves OOB Mgmt GWs	OOB Mgmt SWs
Port Speed	400Gb/s	100 Gb/s	1/10Gb/s

Configuring the SONiC Network

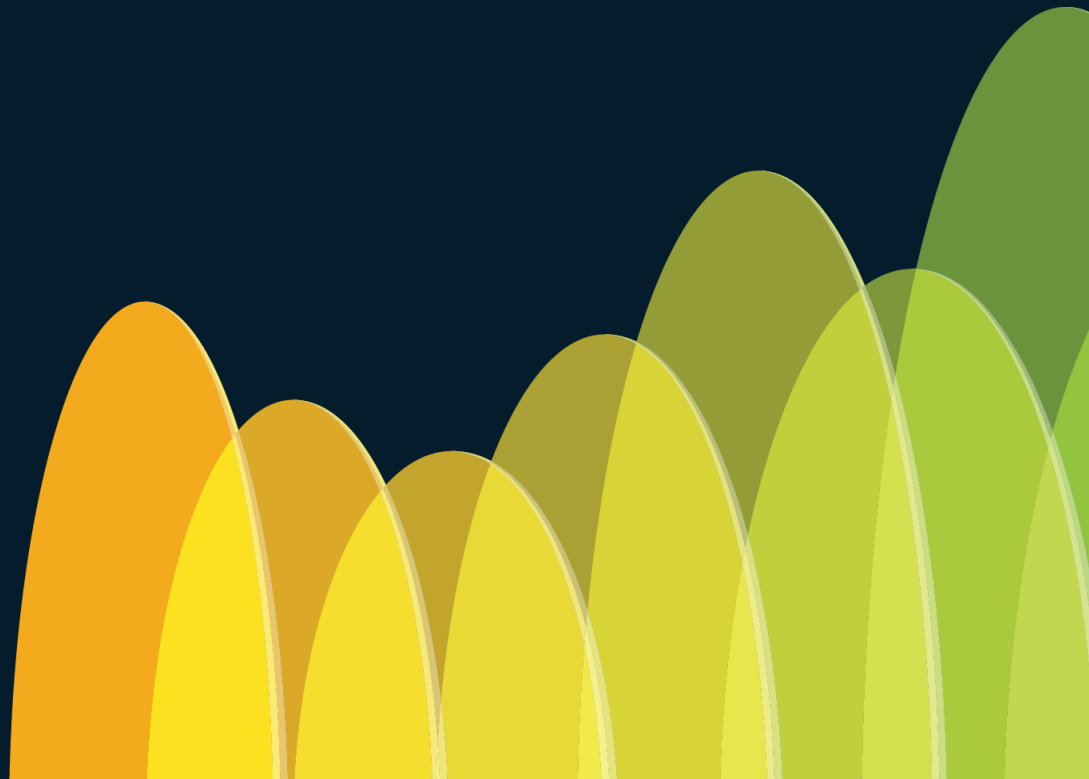

cisco Live!

Cisco 8000 Setup

- Bootup
  - Transfer SONiC image
  - Sonic Installer
  - Reload
- Upgrade
  - Backup Files
  - Sonic installer
  - Reload

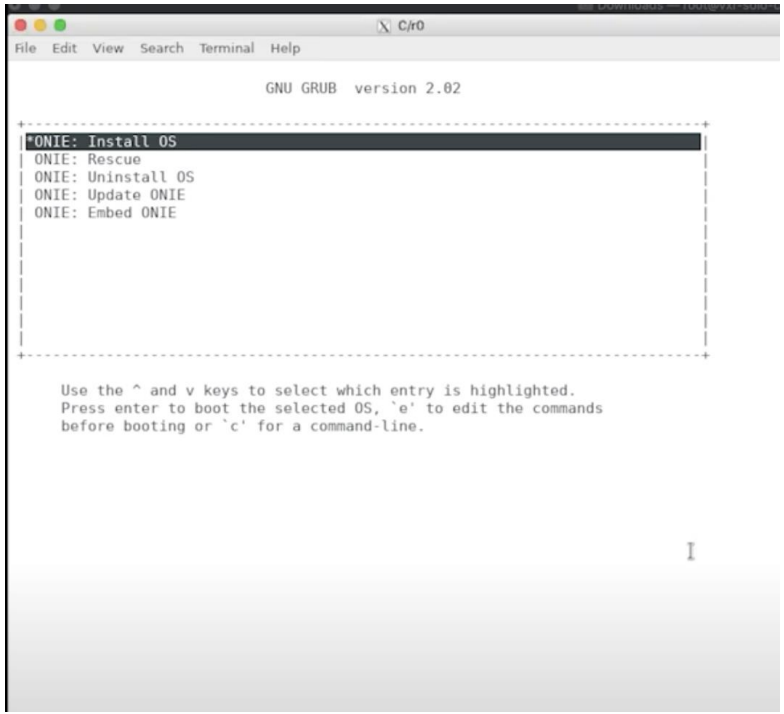

Network Configuration

- Native sonic cli

```
admin@svlngen4-sonic-8k-1:~$ sudo config interface ip add Ethernet16 172.20.12.1/30

admin@svlngen4-sonic-8k-1:~$ sudo config interface ip add Ethernet24 192.168.20.1/24

admin@svlngen4-sonic-8k-1:~$ sudo config save -y
admin@svlngen4-sonic-8k-1:~$



admin@svlngen4-sonic-8k-1:~$ show ip interfaces
Interface    Master    IPv4 address/mask    Admin/Oper    BGP Neighbor    Neighbor IP
-----------  --------  -----------------  ------------  -------------  ------------------------
Ethernet16             172.20.12.1/30       up/up         8101-2         ('default', '172.20.12.2')
Ethernet24             192.168.20.1/24      up/up         8101-2-vlan20  ('default', '192.168.20.2')
```

- Modify the config_db JSON File

```
admin@svlngen4-sonic-8k-1:/etc/sonic$ pwd
/etc/sonic
admin@svlngen4-sonic-8k-1:/etc/sonic$ ls | grep config_db.json
config_db.json

admin@svlngen4-sonic-8k-1:/etc/sonic$ grep -A 5 INTERFACE config_db.json
    "INTERFACE": {
        "Ethernet16": {},
        "Ethernet16|172.20.12.1/30": {},
        "Ethernet24": {},
        "Ethernet24|192.168.20.1/24": {}
    },

admin@svlngen4-sonic-8k-1:~$ sonic-cfggen -j /etc/sonic/config_db.json
admin@svlngen4-sonic-8k-1:~$
admin@svlngen4-sonic-8k-1:~$ sudo config save -y
Running command: /usr/local/bin/sonic-cfggen -d --print-data > /etc/sonic/config_db.json
admin@svlngen4-sonic-8k-1:~$
```

Network Configuration

```
sudo config portchannel add PortChannel0000 --min-links 2
sudo config portchannel member add PortChannel0000 Ethernet0
sudo config portchannel member add PortChannel0000 Ethernet8
sudo config portchannel add PortChannel0016 --min-links 2
sudo config portchannel member add PortChannel0016 Ethernet16
sudo config portchannel member add PortChannel0016 Ethernet24
sudo config portchannel add PortChannel0032 --min-links 2
sudo config portchannel member add PortChannel0032 Ethernet32
sudo config portchannel member a
sudo config portchannel add Por
sudo config portchannel member a
sudo config portchannel member a
sudo config portchannel add Por
sudo config portchannel member a
sudo config portchannel member a
sudo config portchannel add Por
sudo config portchannel member a
sudo config portchannel member a
sudo config portchannel add Por
sudo config portchannel member a
sudo config portchannel member a
sudo config portchannel add Por
sudo config portchannel member a
sudo config portchannel member a
```

INTERFACE.json

```
{
    "INTERFACE": {
        "PortChannel0000": {},
        "PortChannel0000|10.2.20.2/30": {},
        "PortChannel0016": {},
        "PortChannel0016|10.2.20.130/30": {},
        "PortChannel0032": {},
        "PortChannel0032|10.2.21.2/30": {},
        "PortChannel0048": {},
        "PortChannel0048|10.2.21.130/30": {},
        "PortChannel0064": {},
        "PortChannel0064|10.2.22.2/30": {},
        "PortChannel0080": {},
        "PortChannel0080|10.2.22.130/30": {},
        "PortChannel0096": {},
        "PortChannel0096|10.2.23.2/30": {},
        "PortChannel0112": {},
        "PortChannel0112|10.2.23.130/30": {},
```

BGP_NEIGHBOR.json

```
{
    "BGP_NEIGHBOR": {
        "default|10.2.20.33": {
            "admin_status": "up",
            "asn": "65200",
            "holdtime": "30000",
            "keepalive": "10000",
            "local_addr": "10.2.20.34",
            "name": "mtv5-ag10-sfab1-sw3901:PortChannel0064",
            "nhopself": "0",
            "rrclient": "0"
        },
        "default|10.2.20.161": {
            "admin_status": "up",
            "asn": "65200",
            "holdtime": "30000",
            "keepalive": "10000",
            "local_addr": "10.2.20.162",
            "name": "mtv5-ag11-sfab1-sw3902:PortChannel0064",
            "nhopself": "0",
            "rrclient": "0"
        },
        "default|10.2.21.33": {
            "admin_status": "up",
            "asn": "65200",
            "holdtime": "30000",
            "keepalive": "10000",
            "local_addr": "10.2.21.34",
            "name": "mtv5-ag13-sfab1-sw3903:PortChannel0064",
            "nhopself": "0",
            "rrclient": "0"
        },
```

QoS Configuration

PORT_QOS_MAP.json

```
{
    "PORT_QOS_MAP": {
        "Ethernet0": {
            "dscp_to_tc_map": "CISCO",
            "pfc_enable": "3,4",
            "pfc_to_queue_map": "CISCO",
            "pfcwd_sw_enable": "3,4",
            "tc_to_pg_map": "CISCO",
            "tc_to_queue_map": "CISCO"
        },
        "Ethernet8": {
            "dscp_to_tc_map": "CISCO",
            "pfc_enable": "3,4",
            "pfc_to_queue_map": "CISCO",
            "pfcwd_sw_enable": "3,4",
            "tc_to_pg_map": "CISCO",
            "tc_to_queue_map": "CISCO"
        },
        "Ethernet16": {
            "dscp_to_tc_map": "CISCO",
            "pfc_enable": "3,4",
            "pfc_to_queue_map": "CISCO",
            "pfcwd_sw_enable": "3,4",
            "tc_to_pg_map": "CISCO",
            "tc_to_queue_map": "CISCO"
        },
```

TC_TO_QUEUE_MAP.json

```
{
    "TC_TO_QUEUE_MAP": {
        "CISCO": {
            "0": "0",
            "1": "1",
            "2": "2",
            "3": "3",
            "4": "4",
            "5": "5",
            "6": "6",
            "7": "7"
        }
    }
}
```

WRED_PROFILE.json

```
{
    "WRED_PROFILE": {
        "CISCO_LOSSLESS": {
            "wred_green_enable": "true",
            "wred_yellow_enable": "true",
            "ecn": "ecn_green_yellow",
            "green_drop_probability": "0",
            "green_min_threshold": "0",
            "green_max_threshold": "0",
            "yellow_min_threshold": "0",
            "yellow_drop_probability": "0",
            "yellow_max_threshold": "6144000"
        }
    }
}
```

Generating Configurations

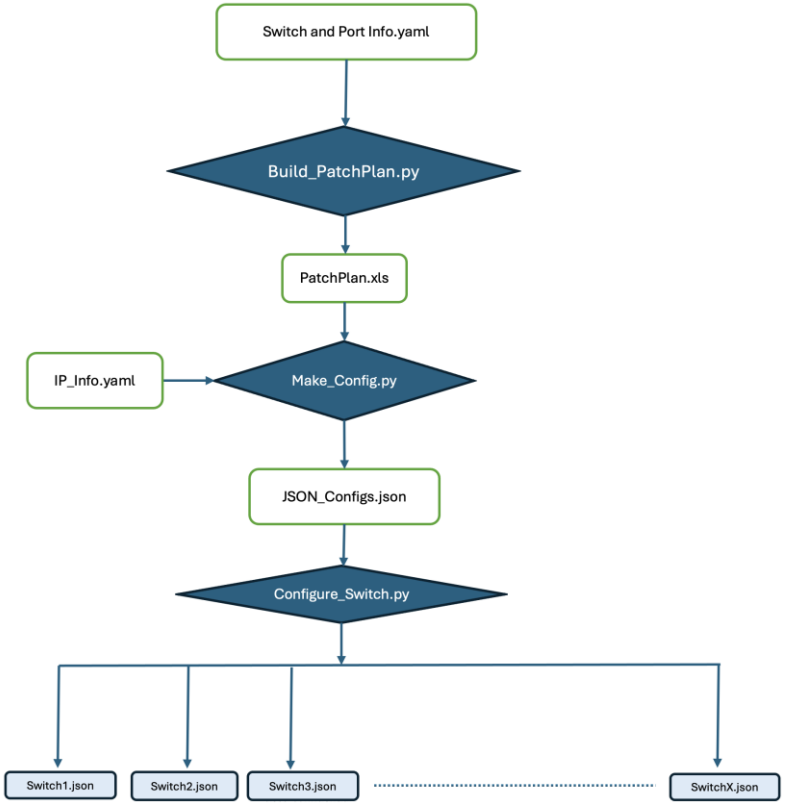

Configuring the Nexus Network

Nexus 9300 – NDFC Fabric Setup

- Fabric Setup
  1. Run Fabric Creation Wizard
  2. Select Fabric Template
  3. Define Fabric Parameters

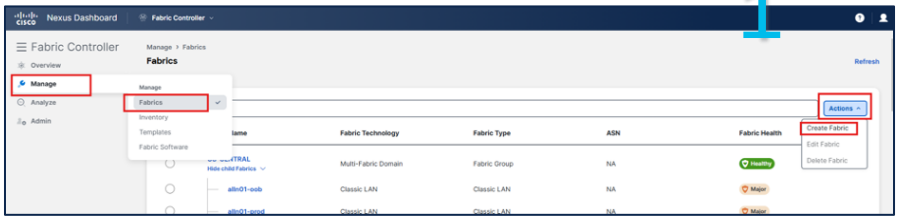

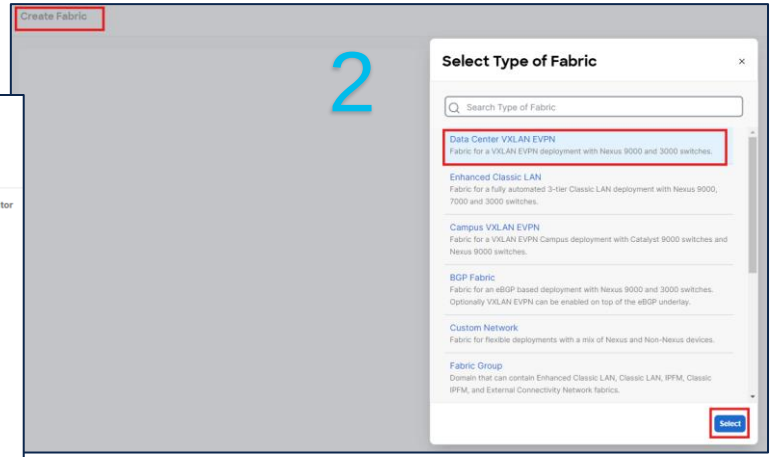

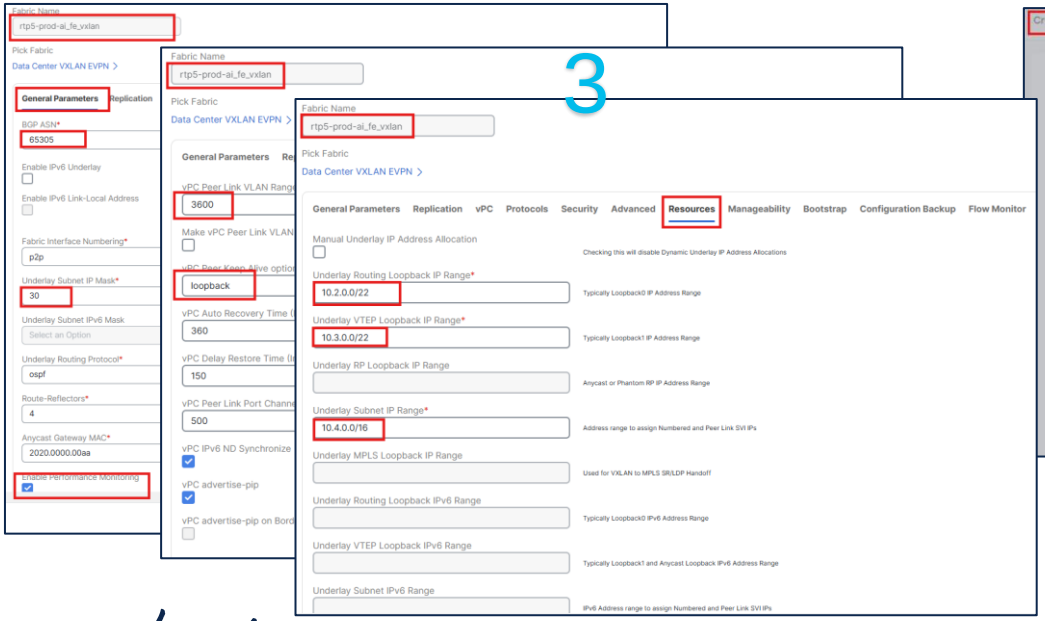

Nexus 9300 Setup

- POAP and Configure Network
  1. PreProvision Switches
  2. Import with POAP
  3. Deploy Network Configs

1

Host Name	IP Address	Role	serialNumber	Model	Software
rtp5-[redacted]-fefab1-sw1042	10.115.[redacted]	leaf	FDO2833034E	N9K-C93600CD-GX	10.3(6)
rtp5-[redacted]-fefab1-sw1041	10.115.[redacted]	leaf	FDO2833037Q	N9K-C93600CD-GX	10.3(6)
rtp5-[redacted]-fefab1-sw1032	10.115.[redacted]	leaf	FDO2833036J	N9K-C93600CD-GX	10.3(6)
rtp5-[redacted]-fefab1-sw1072	10.115.[redacted]	leaf	FDO2833031Z	N9K-C93600CD-GX	10.3(6)
rtp5-[redacted]-fefab1-sw3822	10.115.[redacted]	border	FDO2833033E	N9K-C93600CD-GX	10.3(6)
rtp5-[redacted]-fefab1-sw3904	10.115[redacted]	border spine	FDO283507P1	N9K-C9364D-GX2A	10.3(6)
rtp5-[redacted]-fefab1-sw1031	10.115.[redacted]	leaf	FDO28330VYS	N9K-C93600CD-GX	10.3(6)
rtp5-[redacted]-fefab1-sw1071	10.115.[redacted]	leaf	FDO28300FZF	N9K-C93600CD-GX	10.3(6)
rtp5-[redacted]-fefab1-sw3821	10.115.[redacted]	border	FDO28330313	N9K-C93600CD-GX	10.3(6)
rtp5-[redacted]-fefab1-sw3903	10.115.[redacted]	border spine	FDO283507NX	N9K-C9364D-GX2A	10.3(6)
rtp5-[redacted]-fefab1-sw1022	10.115.[redacted]	leaf	FDO28330335	N9K-C93600CD-GX	10.3(6)
rtp5-[redacted]-fefab1-sw1062	10.115.[redacted]	leaf	FDO2833031M	N9K-C93600CD-GX	10.3(6)

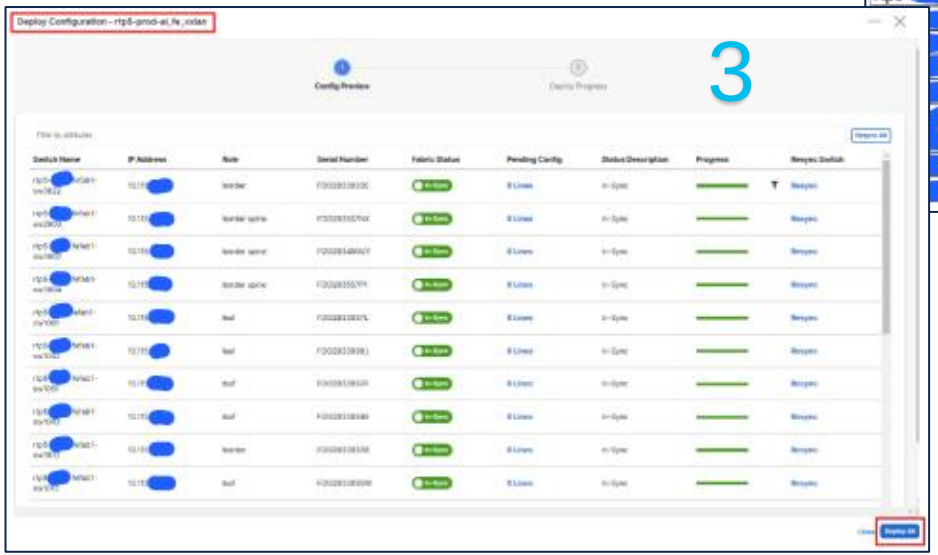

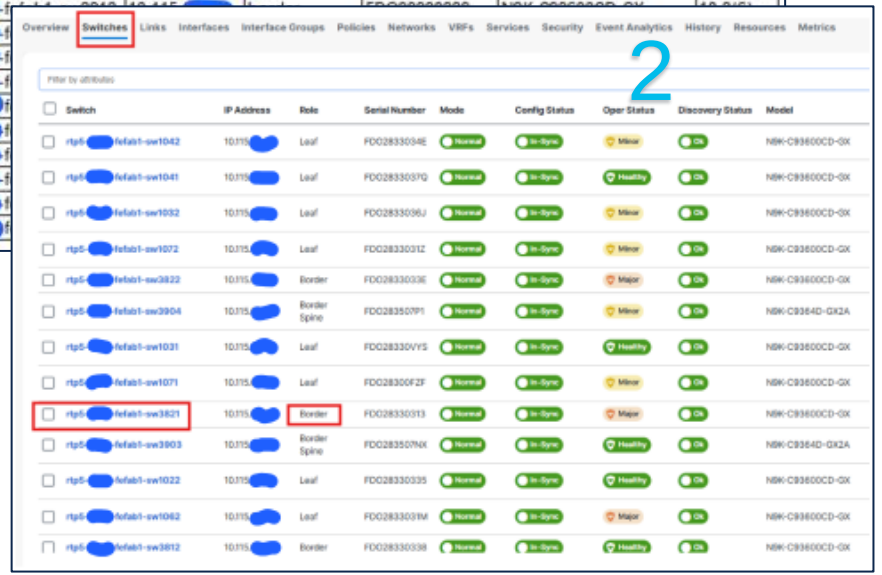

Interface and QoS setup

- Define VPC Pairs
- Define Interface configurations (Interface Wizard)
- Define QoS Policy (standard templates)
- Deploy policy

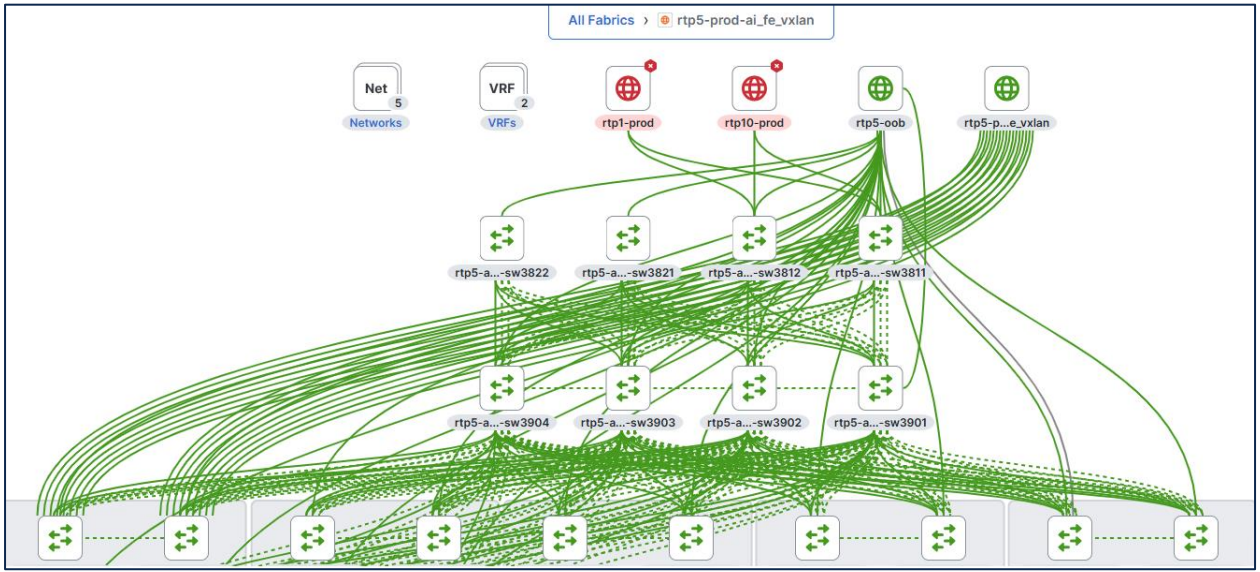

Why Ethernet?

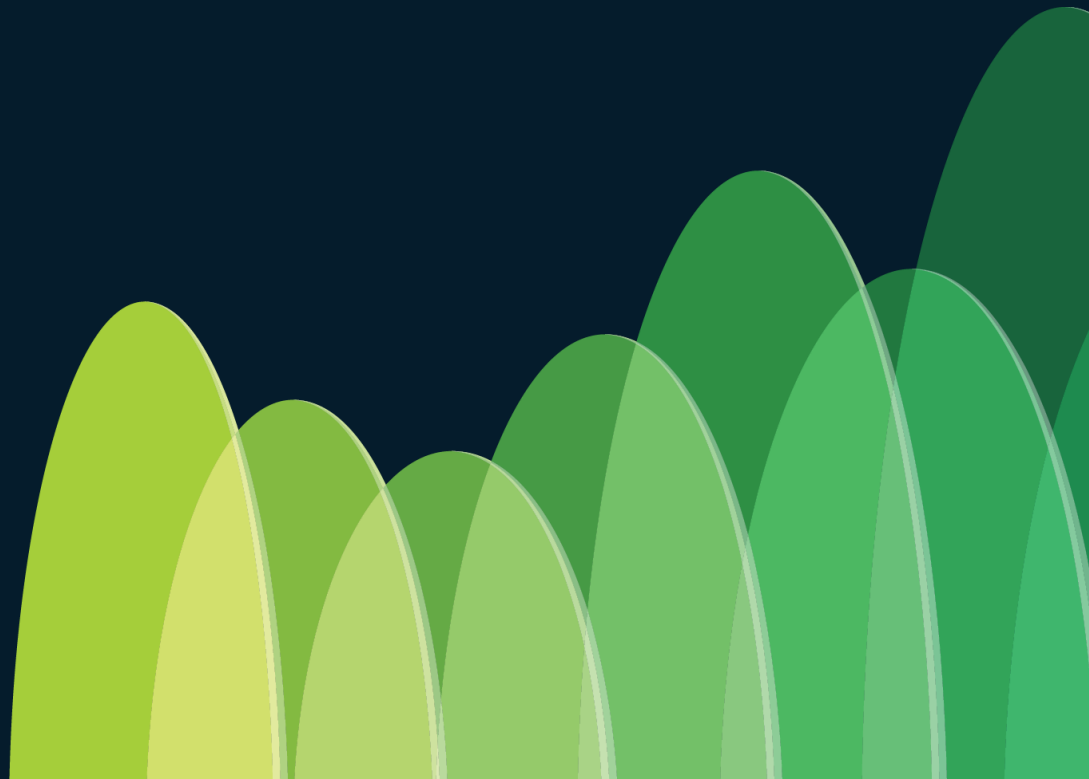

Ethernet vs Infiniband

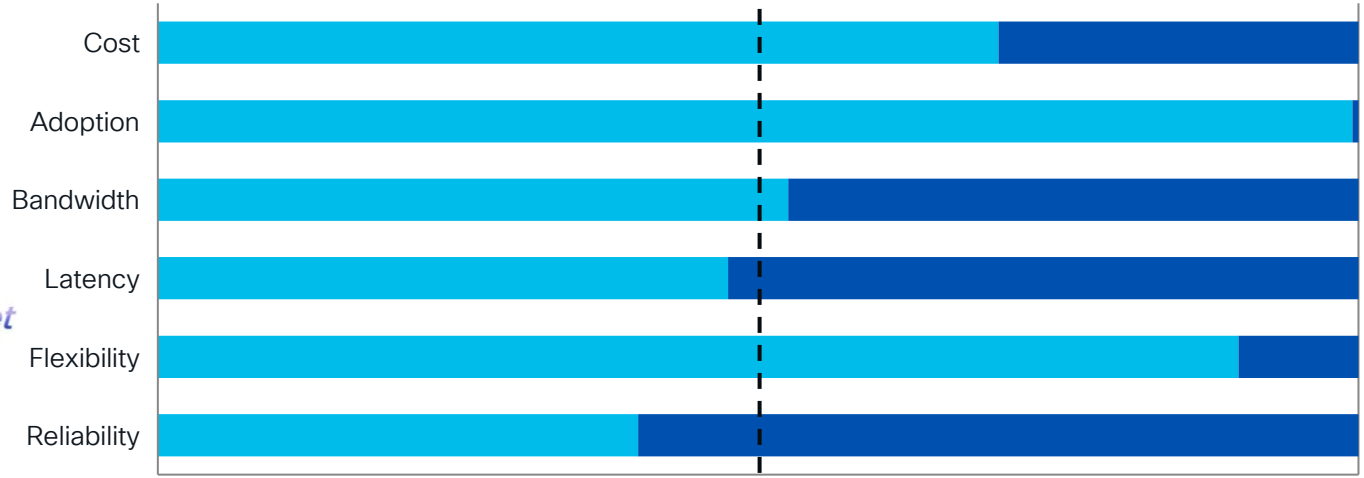

Key Takeaways

1. Ethernet trumps Infiniband due to its widespread use and cost advantages
2. Its flexibility allows much needed convergence in mixed use DCs
3. Infiniband will continue to prevail in some HPC and AI environments (Nvidia appliances)
4. Ultra Ethernet forum to deliver improvements to adapt Ethernet to AI

Note: Infiniband is not proprietary; open spec by IBTA

AI Network Market Trends

AI Back-End Networks - Ethernet Switch Port Growth

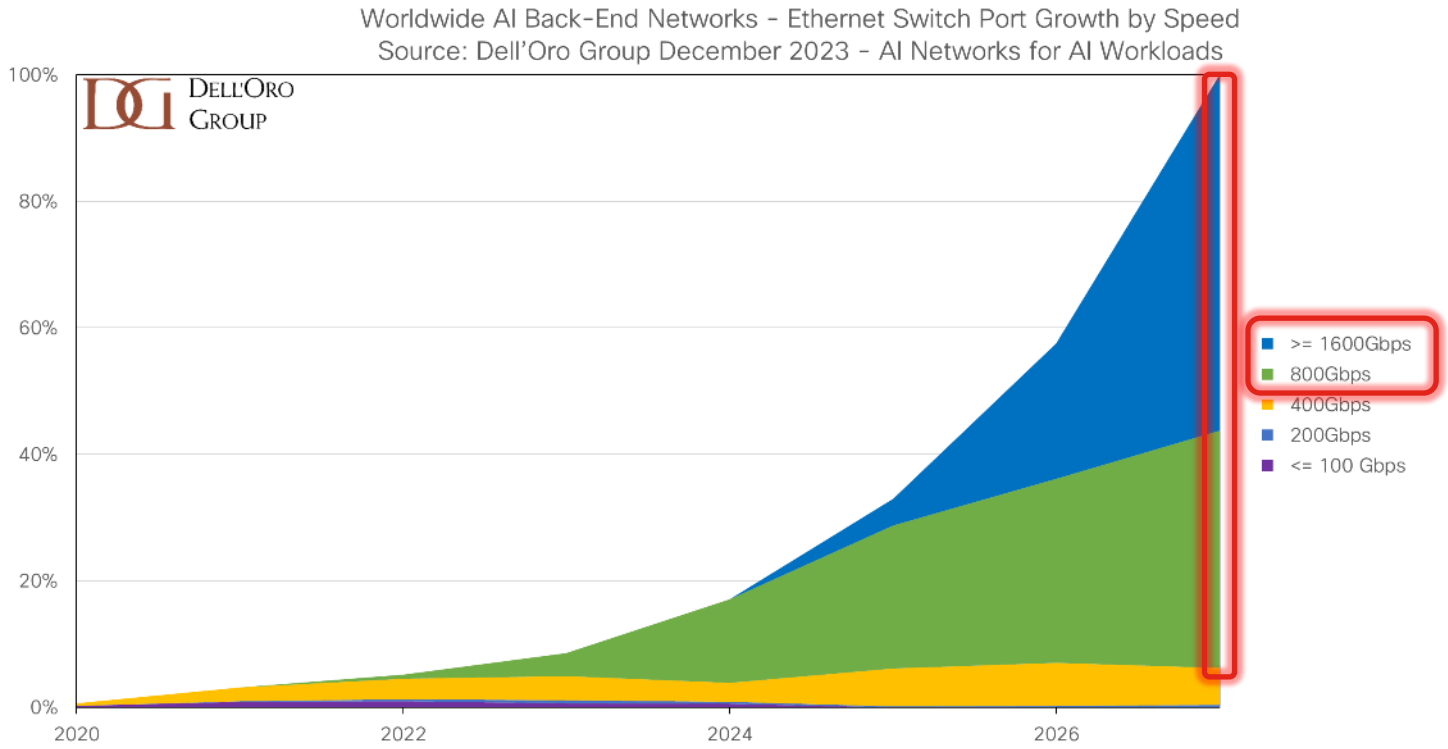

NOTE: The graph above shows percentage of 2028 total ports.

SONiC

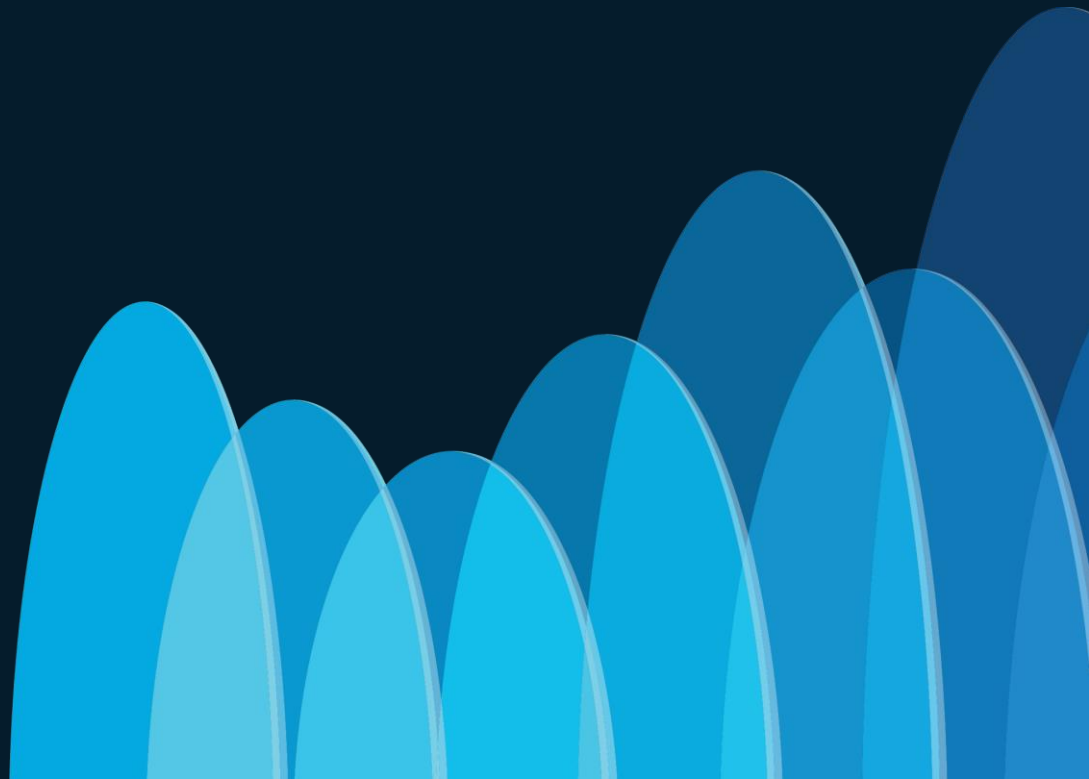

cisco Live!

What is SONiC?

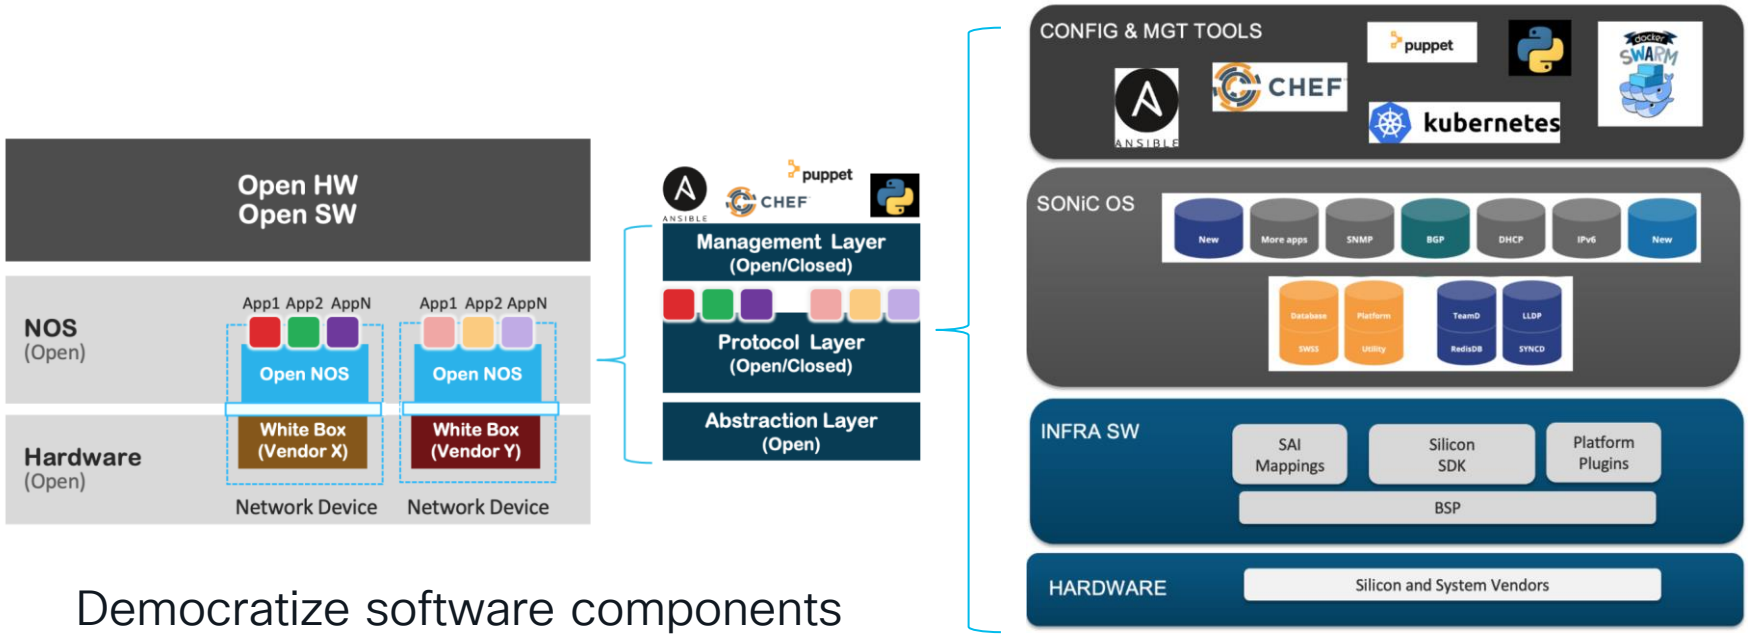

Democratize software components

SONiC – The functional stack

SONiC Feature Development

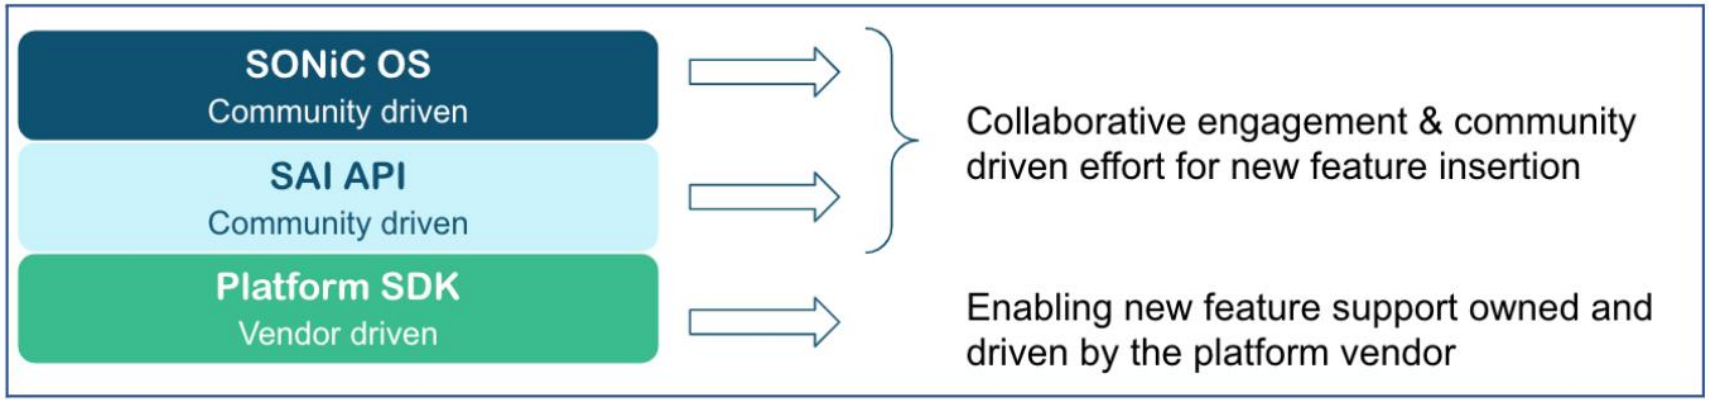

SONiC – Feature Set

SONiC Deployable – Technology Components

Management
- Telemetry (gRPC, gNMI)
- SNMP
- Syslog
- LLDP
- NTP
- TACACS+
- WRED
- QoS
- PFC / ECN
- Everflow
- IPinIP
- ZTP

Interfaces
- L2 & L3
- LAG
- SVI
- MTU setting

Layer 2
- VLAN
- MAC / ARP Aging

Layer 3
- IPv4/IPv6
- BGP
- BGP knobs (community, as-path-replace/relax, add-path, route aggregate, prefix-lists)
- ECMP
- BGP GR helper
- BGP MP
- ACL – IPv4 & IPv6
- COPP
- RPL (route-policies)
- VRF

Infrastructure
- SWSS Unit Test Framework
- ConfigDB framework

Overlay
- Static VxLAN

Security
- MACsec (8808)

High Availability
- Warm Boot
- Fast Reload

Use Cases
- DC Fabric
- DCI (L3/IP)
- CDN Fabric
- WAN LER/LSR*

Architectures
- IP Fabric – BGP
- Overlay Fabric – VxLAN
- Overlay – EVPN VxLAN*
- Core – MPLS/SR/SRv6*

Consumers
- Hyperscalers
- Service Providers
- Enterprises

Shipping

Under* Development

Congestion Management and Forwarding

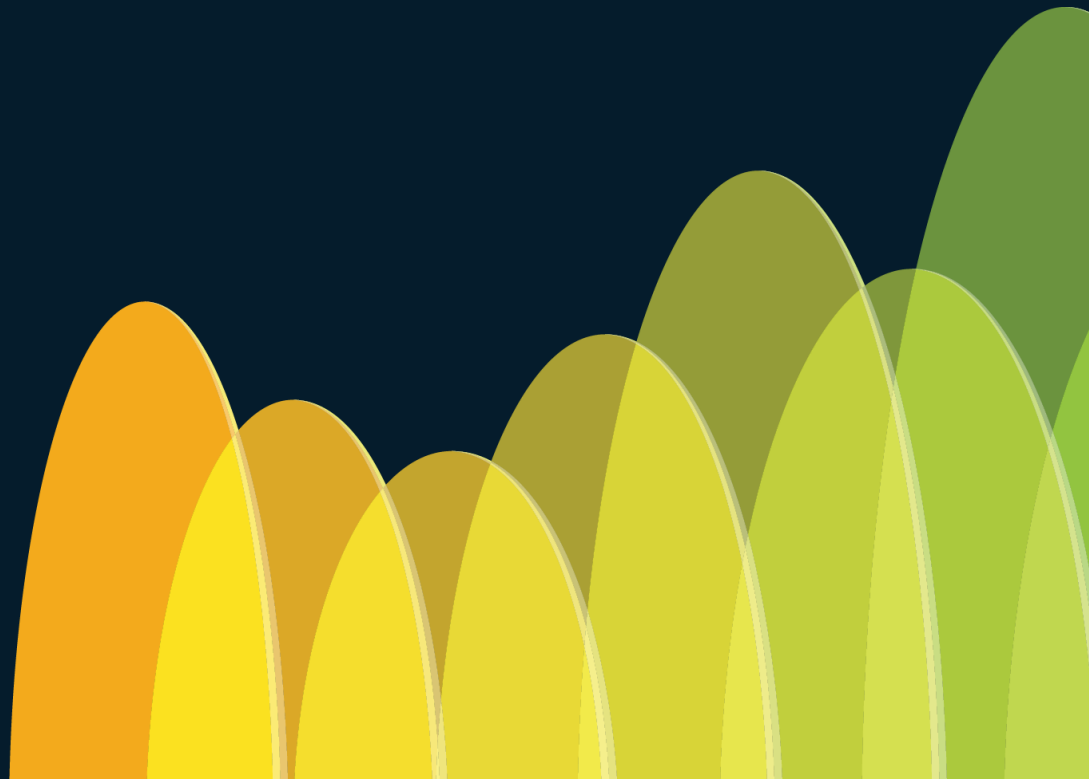

cisco Live!

Explicit Congestion Notification (ECN)

- Purpose: Provides congestion notification in IP networks
- Function: Enables end-to-end congestion notification between two endpoints
- Mechanism:
  - Utilizes 2 LSB of the Type of Service field in the IP header
  - Upon congestion, it triggers the transmitting device to reduce the transmission rate using a Congestion Notification Packet (CNP) without halting traffic

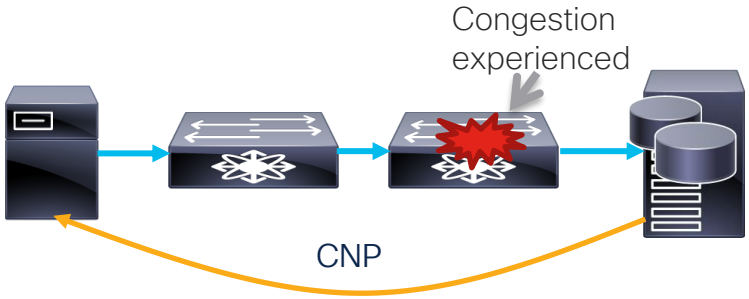

ECN	ECN Behavior
00	Non ECN Capable
10	ECN Capable Transport (0)
01	ECN Capable Transport (1)
11	Congestion Encountered

Priority Flow Control

Flow Control Mechanism – 802.1Qbb

- **Also Known As**: "Lossless Ethernet"
- **Function**: Enables Flow Control on a Per-Priority basis
- **Alternate Name:** Per-Priority-Pause
- **Key Benefits:**
  - Allows both lossless and lossy priorities on the same wire.
  - Ensures traffic can operate over a lossless priority independently.
  - Other traffic on different priorities continues to transmit, relying on upper layer protocols for retransmission.

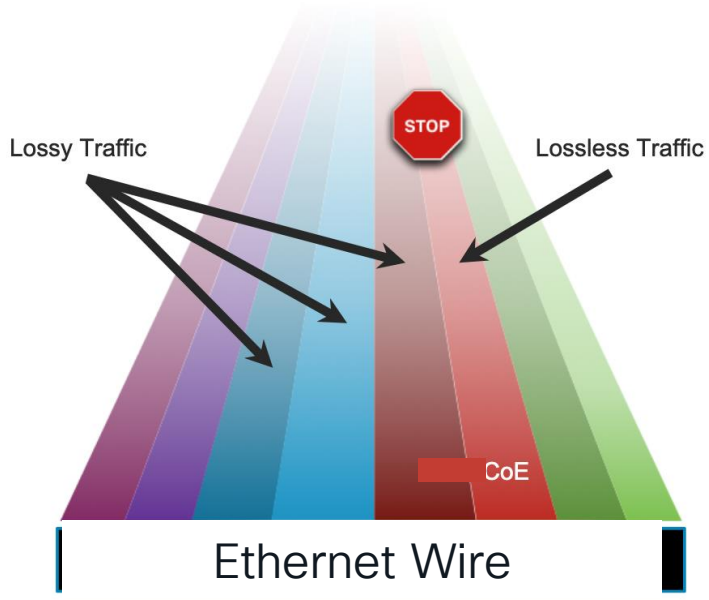

PFC Mechanism

- **Thresholds**: Set in no-drop queue with headroom for "in-flight" packets.
- **Buffering**: Traffic buffered in no-drop queue during congestion.
- **PFC Frames**: Sent to sender when queue utilization exceeds xoff threshold.
- **Queue Management**: Stops sending PFC frames when utilization drops below xon threshold.

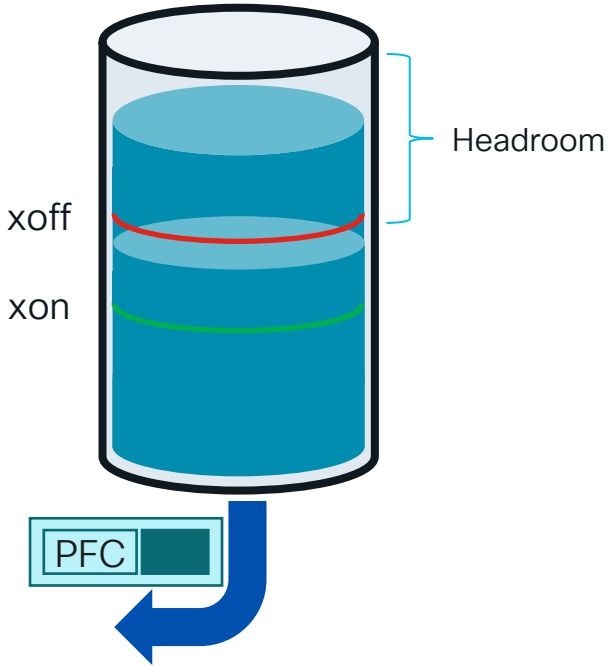

Fabric Forwarding
Operation with ECMP

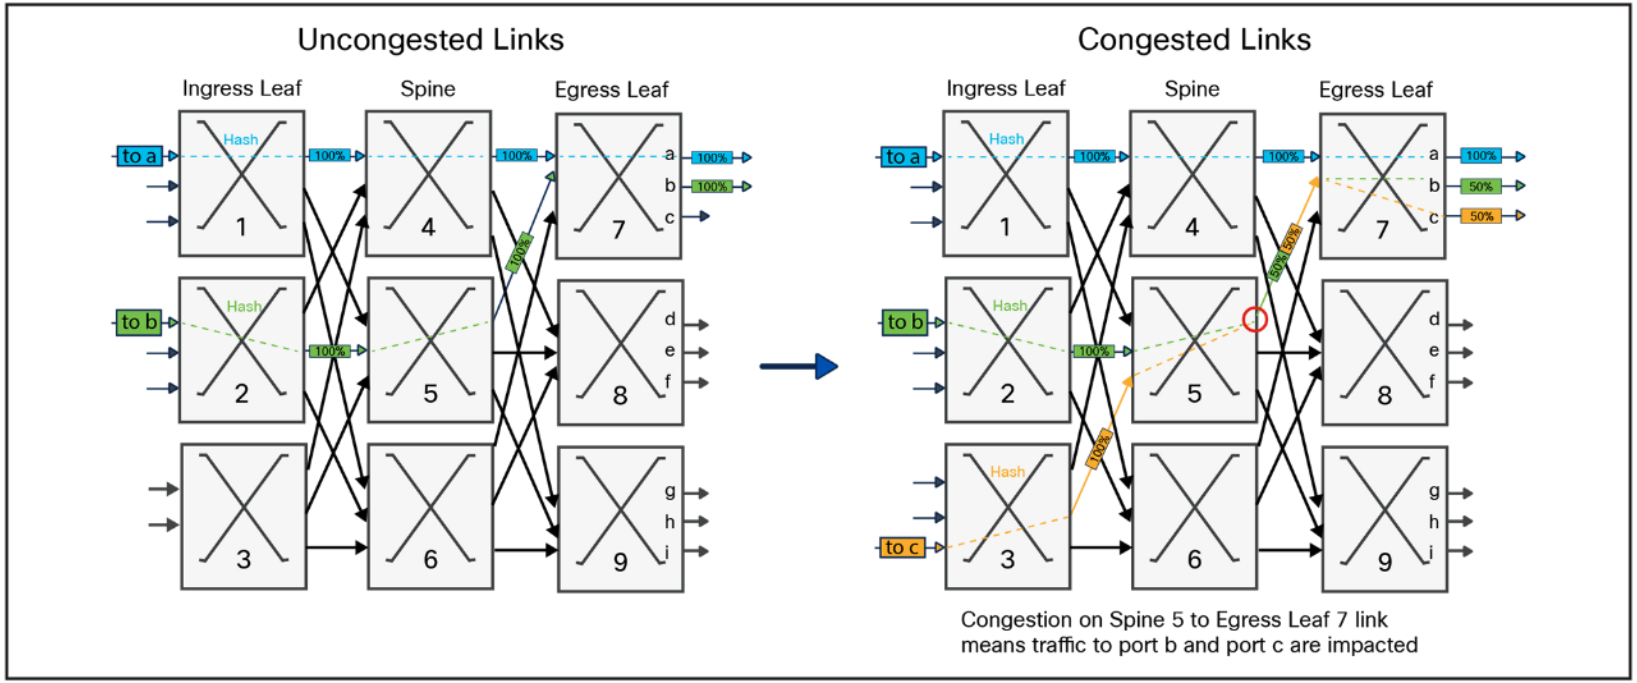

Day 2 Ops: SONiC

cisco Live!

SONiC Day 2 Ops – Grafana Dashboards

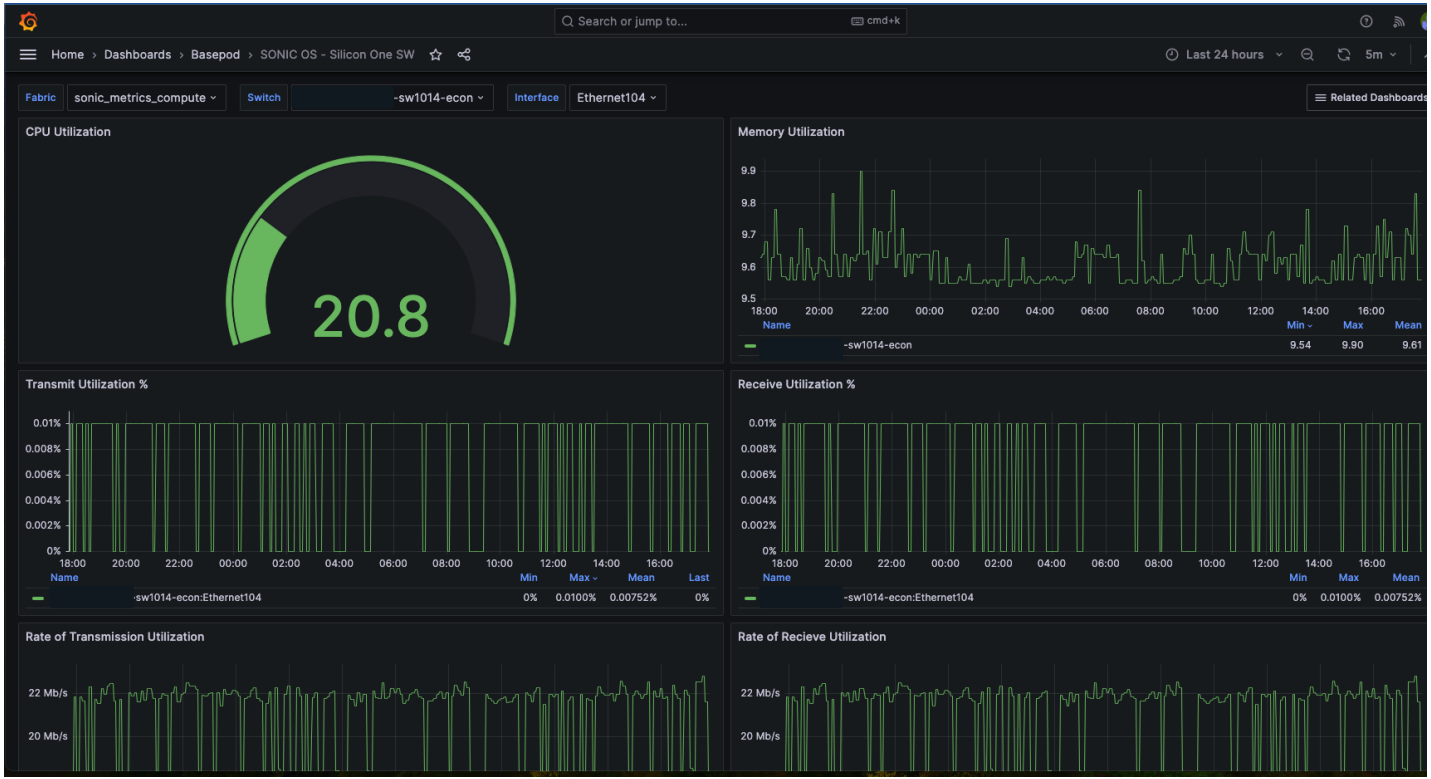

SONiC Day 2 Ops – Operational Runbooks

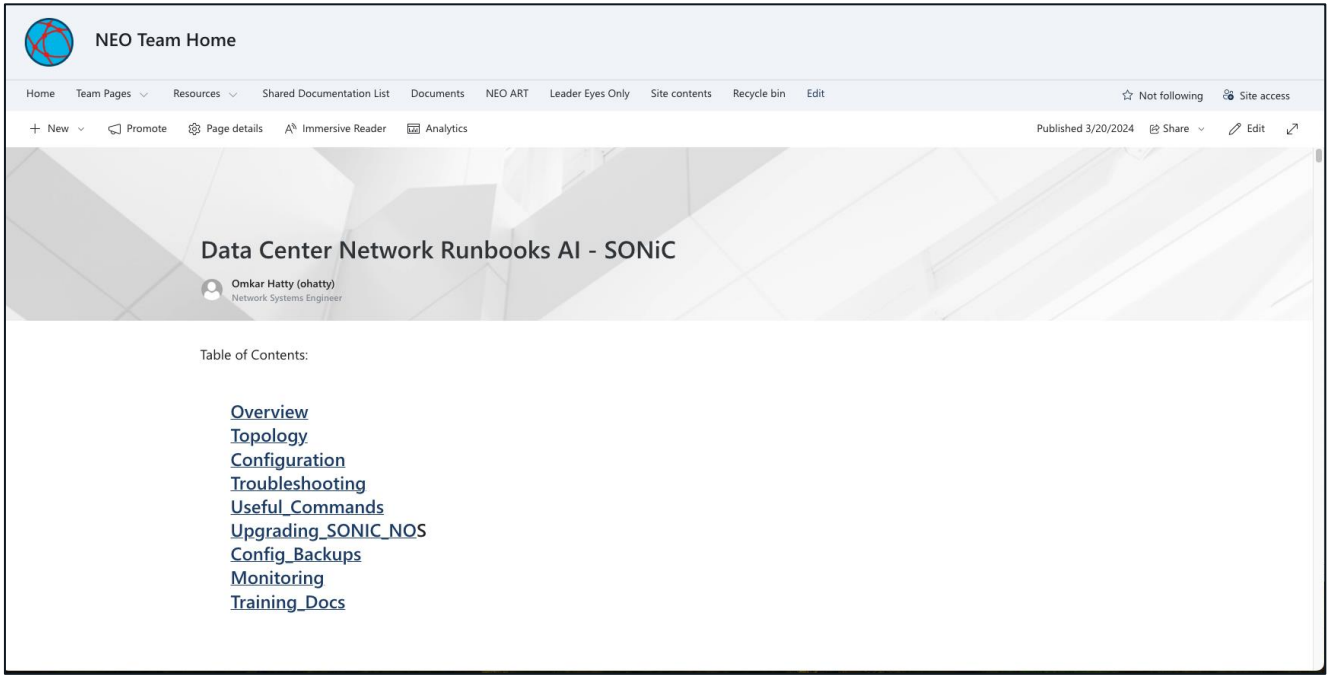

SONiC Day 2 Ops – Monitoring and Backup

- Zabbix

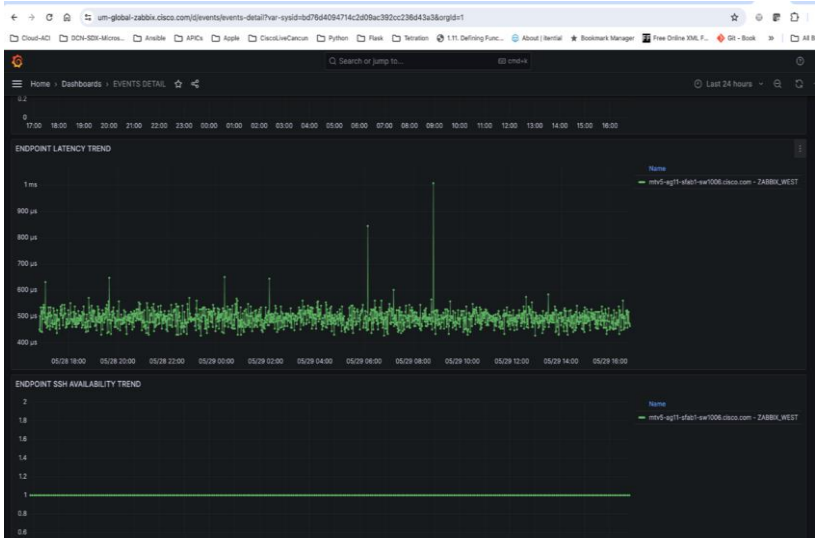

- Bitbucket

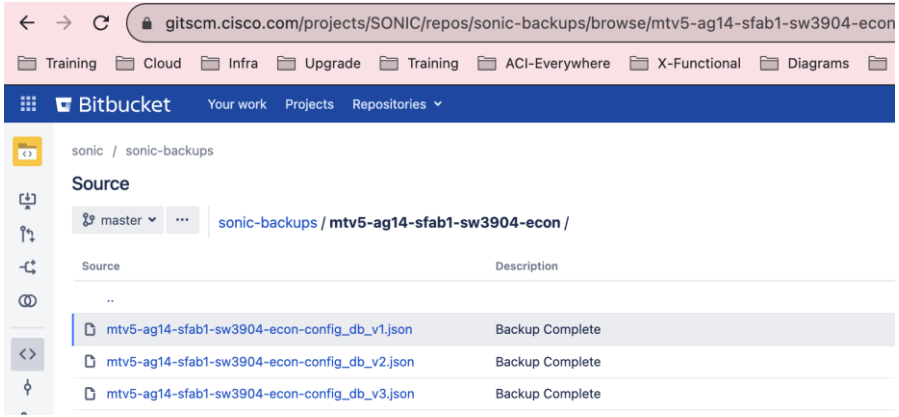

Day 2 Ops: Nexus

cisco Live!

Nexus Day 2 Ops – NDFC Inventory and Events

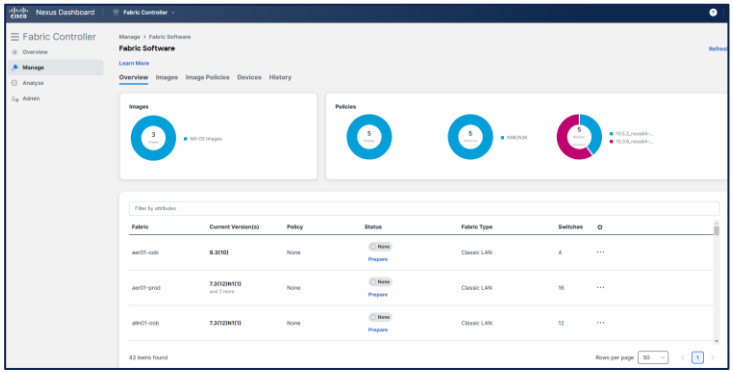

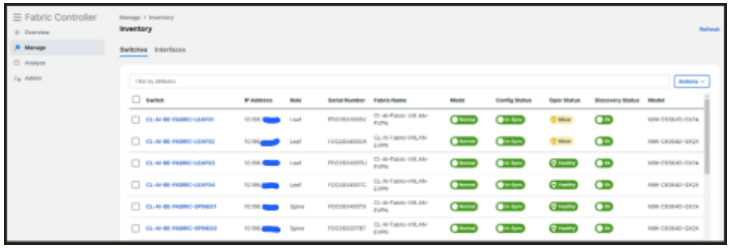

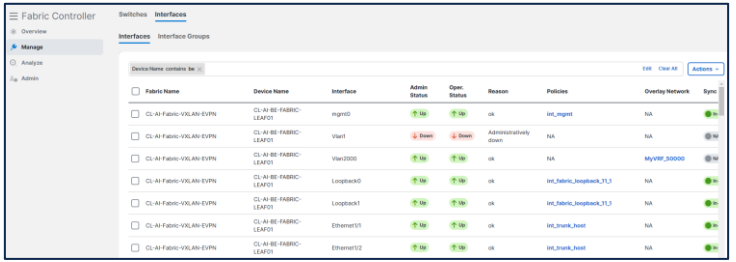

Nexus Day 2 Ops – NDI Performance Monitor

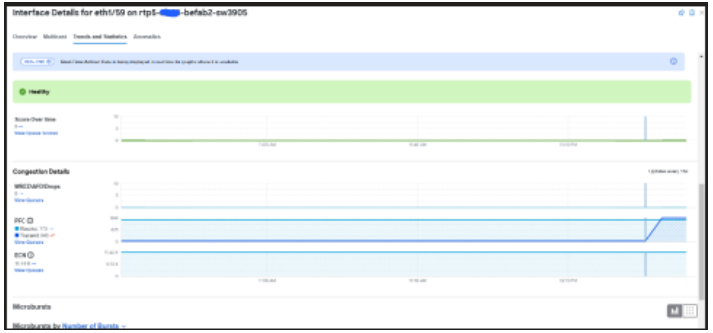

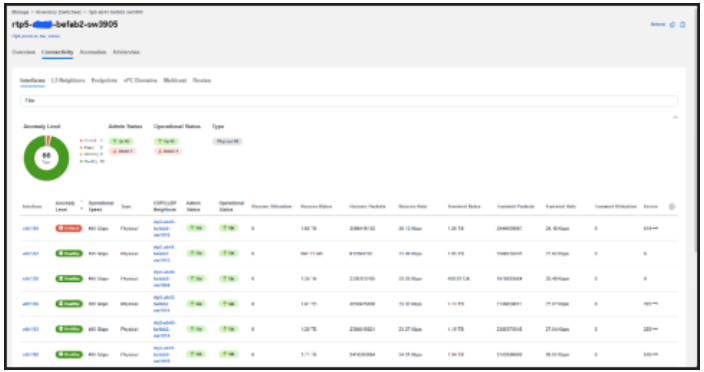

Nexus Day 2 Ops – NDI Advanced Capabilities

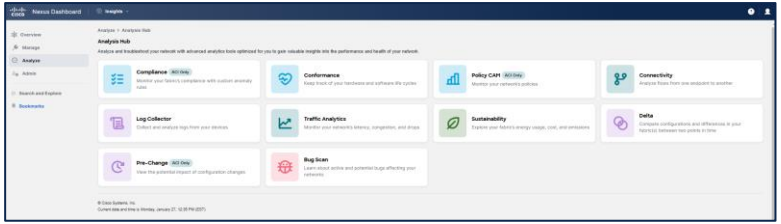

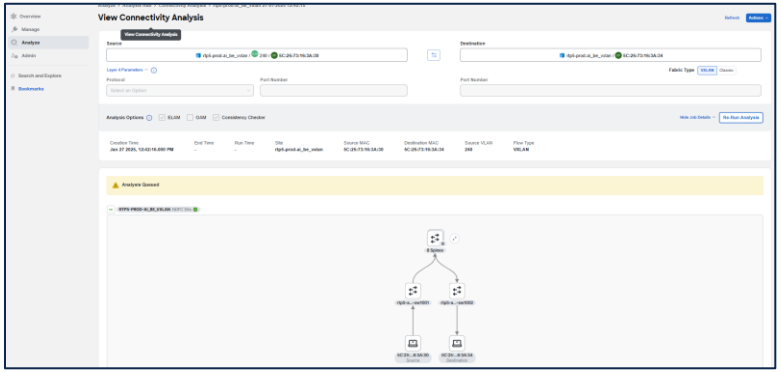

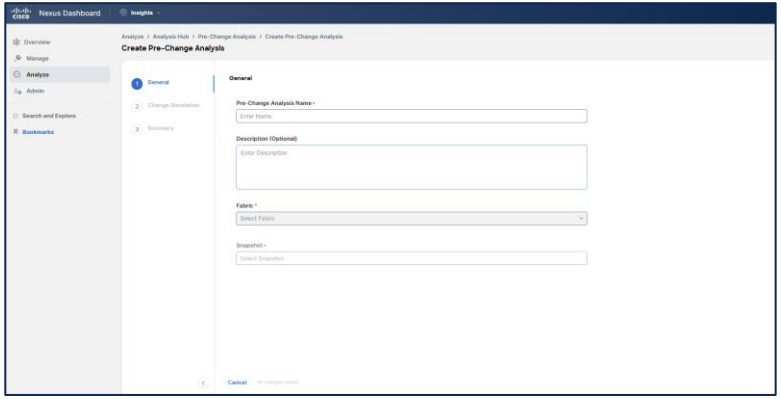

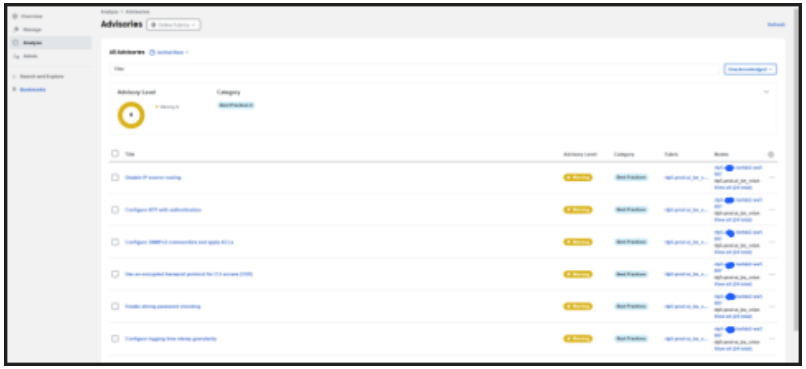

Use Cases

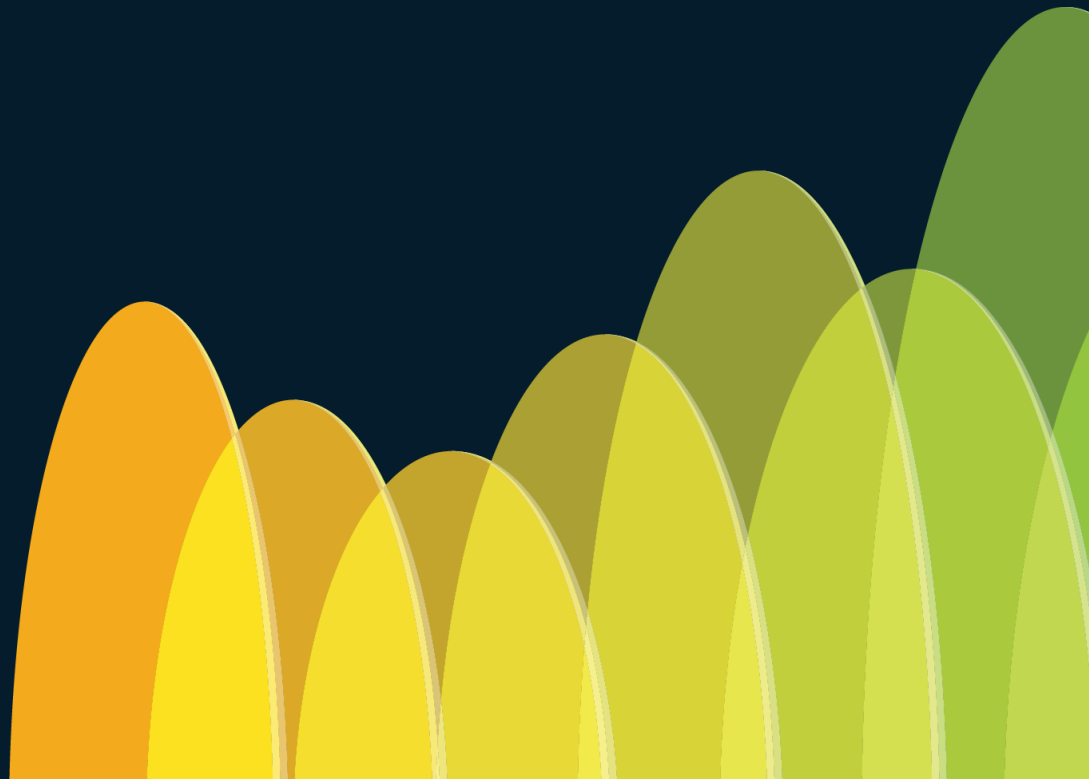

A Sample of Current Usage

Team	Use Case	Business Outcomes
Webex Audio (Babblelabs)	Improving codec development for noise cancellation and lower bandwidth data prediction	Improved Webex Teams feature performance
Webex Video	The team trains video AI models for Cisco, including background replacement, gesture recognition, and face landmarks.	Improved Webex Teams feature performance
Cisco SBG	Specialized code generation LLMs to reason over Cisco Security devices; Custom pretrained LLMs built for the cybersecurity domain.	Increase Cisco's security posture trough AI-driven service & application log analysis
I2C & Revenue (Core Finance)	Order Collection Optimization, Payment predictors, AI Ops predictors, Payment fraud analysis	Drive a reduction in late- and non-paying clients through the use of AI
Observability (SNOW)	Avoid business impacting incidents caused from IT changes	>80% accuracy in high-risk change detection leading to corrective action to avoid incidents Future Vision: GAI workflows to add approval paths, and suggestions to reduce Change Request scopes

Lessons Learned

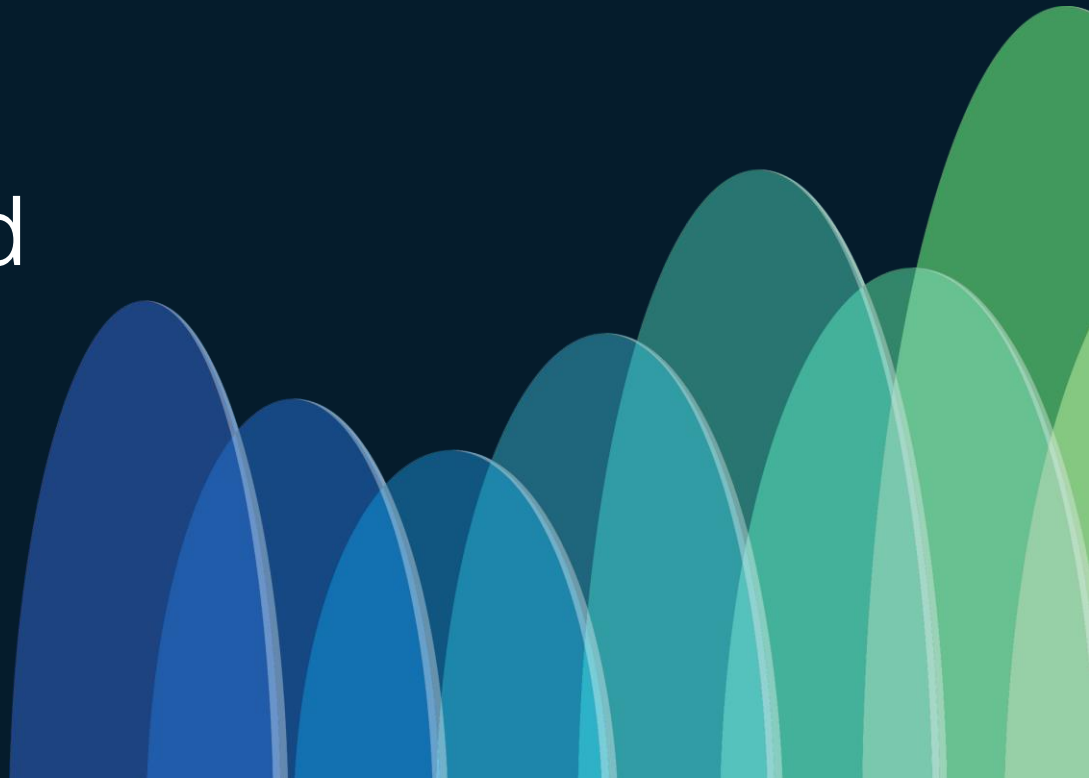

AI presents challenges for IT teams

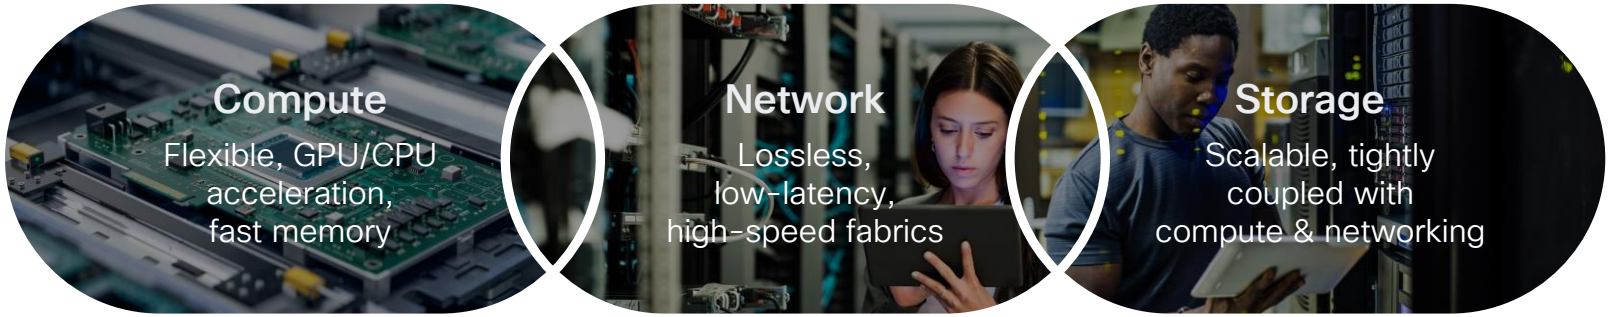

INFRASTRUCTURE DEMANDS

Rapid growth in data volume and variety	Unfamiliar application stacks and new, complex infrastructure patterns	Insufficient IT automation and observability	Greater cybersecurity threats
New operational silos	Shortage of technical expertise	Disorienting AI hype	High entry cost and lock-in issues

Key Takeaways

- **Power Efficiency**:
  - Prioritize power-efficient solutions.
- **Scalable CLOS Fabric Design**:
  - Focus on target CLOS Fabric Design.
  - Ensure a scalable two-tier design.
- **Fixed Switches**:
  - Utilize fixed switches for optimal performance.
  - Achieve lower latency with single ASIC designs.
- **NoS Choices:**
  - **SONiC** is targeted for Hyperscalers and SPs.
  - **Nexus** is targeted for Enterprises.
- **Congestion Management**:
  - Implement effective QoS is very important.
- **BGP for Control Plane**:
  - Utilize BGP for an efficient control plane.
- **Simplicity and Reusability**:
  - Aim for simplicity and reusability in designs.
- **Planning and Expectations**:
  - Cross-Team Planning and Double-check the details.
  - Set Clear expectations.
- **AI Evolution**:
  - Rapid advancements in AI technology.
  - Embrace a mindset of continuous learning.
- **Iteration and Evolution**:
  - Iterate and evolve with ongoing improvements.
  - Follow a Crawl/Walk/Run methodology for implementation.

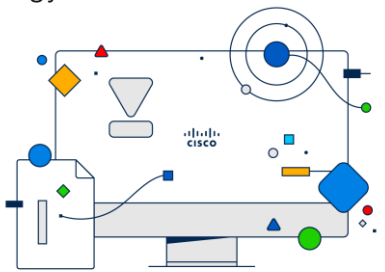

cisco Live!

Future Plans

Multitenant AIaaS	Design and deliver multitenancy capabilities
More AI clusters	Investing in more AI clusters leveraging Cisco UCS 885
SiliconOne Nexus	Deploy N9K-C9364E-SG2 based fabric
HyperFabric	Deploy HyperFabric for Improved Manageability

Webex App

Questions?

Use the Webex app to chat with the speaker after the session

How

1. Find this session in the Cisco Events mobile app
2. Click “Join the Discussion”
3. Install the Webex app or go directly to the Webex space
4. Enter messages/questions in the Webex space

Webex spaces will be moderated by the speaker until February 28, 2025.

Fill Out Your Session Surveys

Participants who fill out a minimum of 4 session surveys and the overall event survey will get a unique Cisco Live t-shirt.

(from 11:30 on Thursday, while supplies last)

All surveys can be taken in the Cisco Events mobile app or by logging in to the Session Catalog and clicking the ‘Participant Dashboard’

Content Catalog

Continue your education

- Visit the Cisco Showcase for related demos
- Book your one-on-one Meet the Engineer meeting
- Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs
- Visit the On-Demand Library for more sessions at ciscolive.com/on-demand. Sessions from this event will be available from March 3.


cisco
Thank you
cisco Live!

cisco Live!

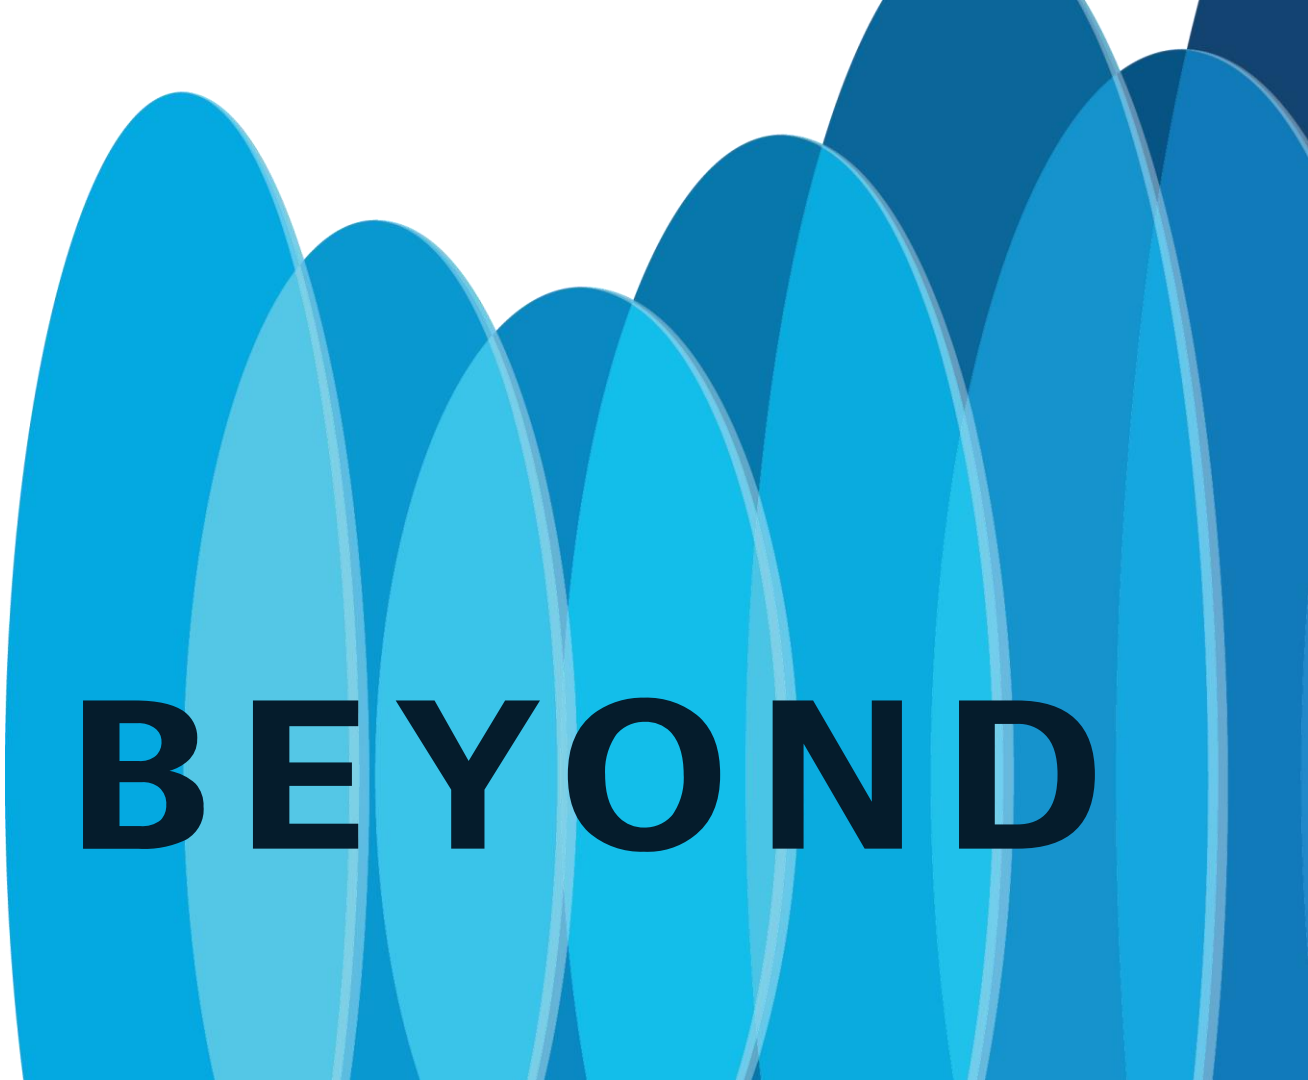

GO BEYOND