
cisco Live!
Let's go
#CiscoLive

The bridge to possible

Design distributed Telco DC for Open-RAN and 5G with Nexus Dashboard, ACI and NX-OS

Sonu Khandelwal
Principal Technical Marketing Engineer

BRKDCN-2972

Abstract

5G brings in many interesting use-cases such as self-driving cars, multi-player gaming, factory robots etc. With Open/Virtual RAN, SPs can move from appliance-based RAN to cloud native open/virtual RAN architecture to take advantage of scale, innovation and reduced costs. In this session, we are going to explain the impact of 5G on Telco DCs. Attendees will learn how to build 5G ready distributed DC with automation, service chaining, network analytics, assurance and telemetry using Cisco Nexus Dashboard, Application Centric Infrastructure (ACI) and NX-OS based on real life experience. We will cover design and architecture that spans across DC and transport for far edge, edge, and central/regional DC for open/virtual RAN deployment.

Agenda

- Telco DC trends and challenges
- Solution overview for automation and operations of distributed DCs
- Building distributed Telco DC
- Cloud native packet core integration
- DC to transport handoff
- Cross-domain automation
- Day2ops and integration with OSS/BSS systems
- Customer references
- Conclusion

Cisco Webex App

Questions?

Use Cisco Webex App to chat
with the speaker after the session

How

1. Find this session in the Cisco Live 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 June 9, 2023.

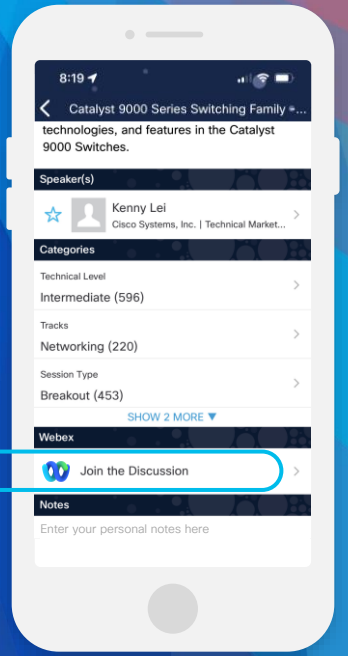

https://ciscolive.ciscoevents.com/ciscolivebot/#BRKDCN-2972

Drivers for distributed Telco DC deployments


Open RAN


Gaming


Edge CDN


Connected cars


Video ingest & analysis

Manufacturing

Healthcare

Retail

Smart cities

IOT

Telco DC Trends

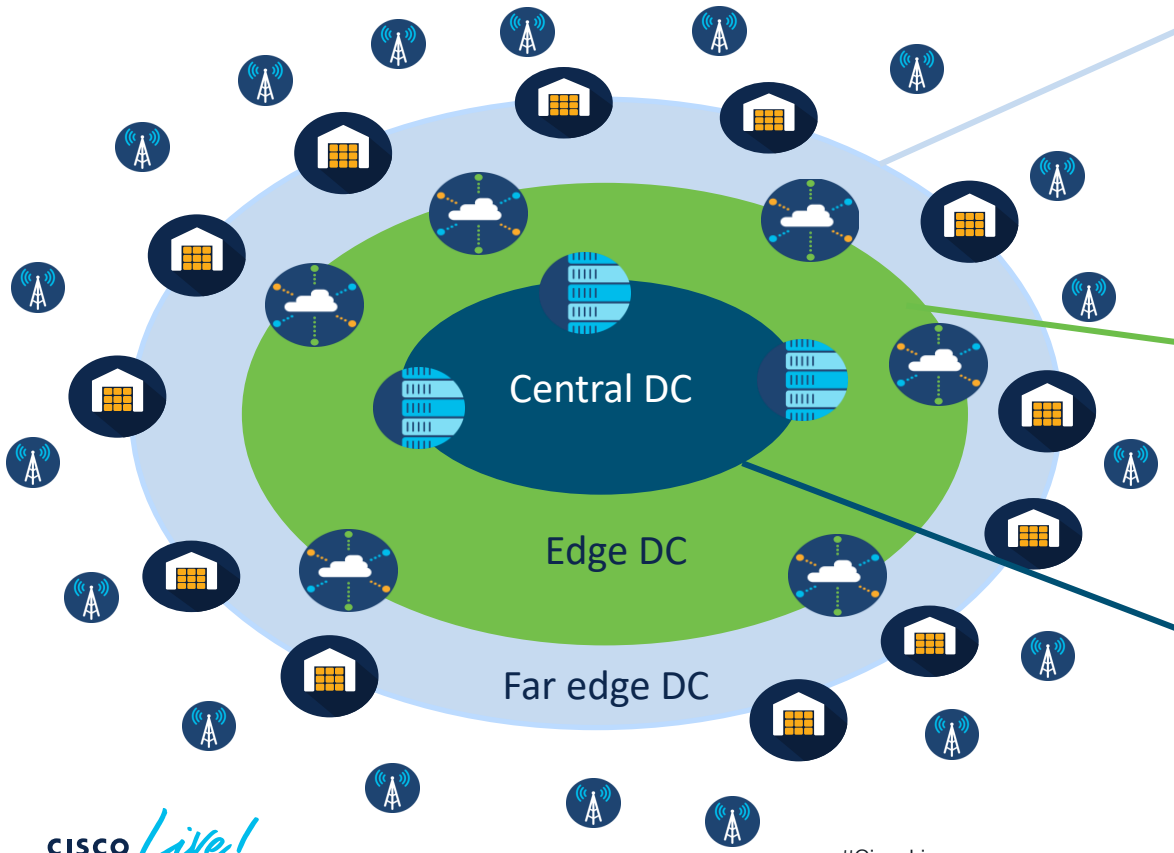

Far Edge DC
1000+ location

- **Use Case:** SP Network Edge
- **Apps:** O-RAN DU
- **Characteristics:** Micro-second latency

Edge DC
100+ locations

- **Use Case:** Network Edge, Service Edge
- **Apps:** O-RAN CU, UPF, Gi-LAN, BNG-U, Caching, Gaming, MEC, IOT
- **Characteristics:** Milli-second latency

Central DC
2-10+ locations

- **Apps :** 5G Control Plane, IMS, Gi-LAN
- **Characteristics:** Non-latency sensitive

Challenges due to distributed Telco DC deployments

Automation

- 100s of fabrics and 1000s of switches
- Single tool to automate distributed DC configuration
- Cross-domain orchestration across transport and DC for use-cases like network slicing

Cloud native apps scale and features

- End point and route scale with fast convergence
- NF movement and placement
- Cloud native packet core integration

Proactive Day2 Ops

- Alarms, KPIs, automated correlations of faults and suggested remedies

Scalable DC to transport handoff

- Automated and scalable handoff for any type of DC fabric to IP, MPLS-LDP, SR-MPLS or SRv6 handoff

Agenda

- Telco DC trends and challenges
- Solution overview for automation and operations of distributed DCs
- Building distributed Telco DC
- Cloud native packet core integration
- DC to transport handoff
- Cross-domain automation
- Day2ops and integration with OSS/BSS systems
- Customer references
- Conclusion

Automation and operations for distributed Telco DCs

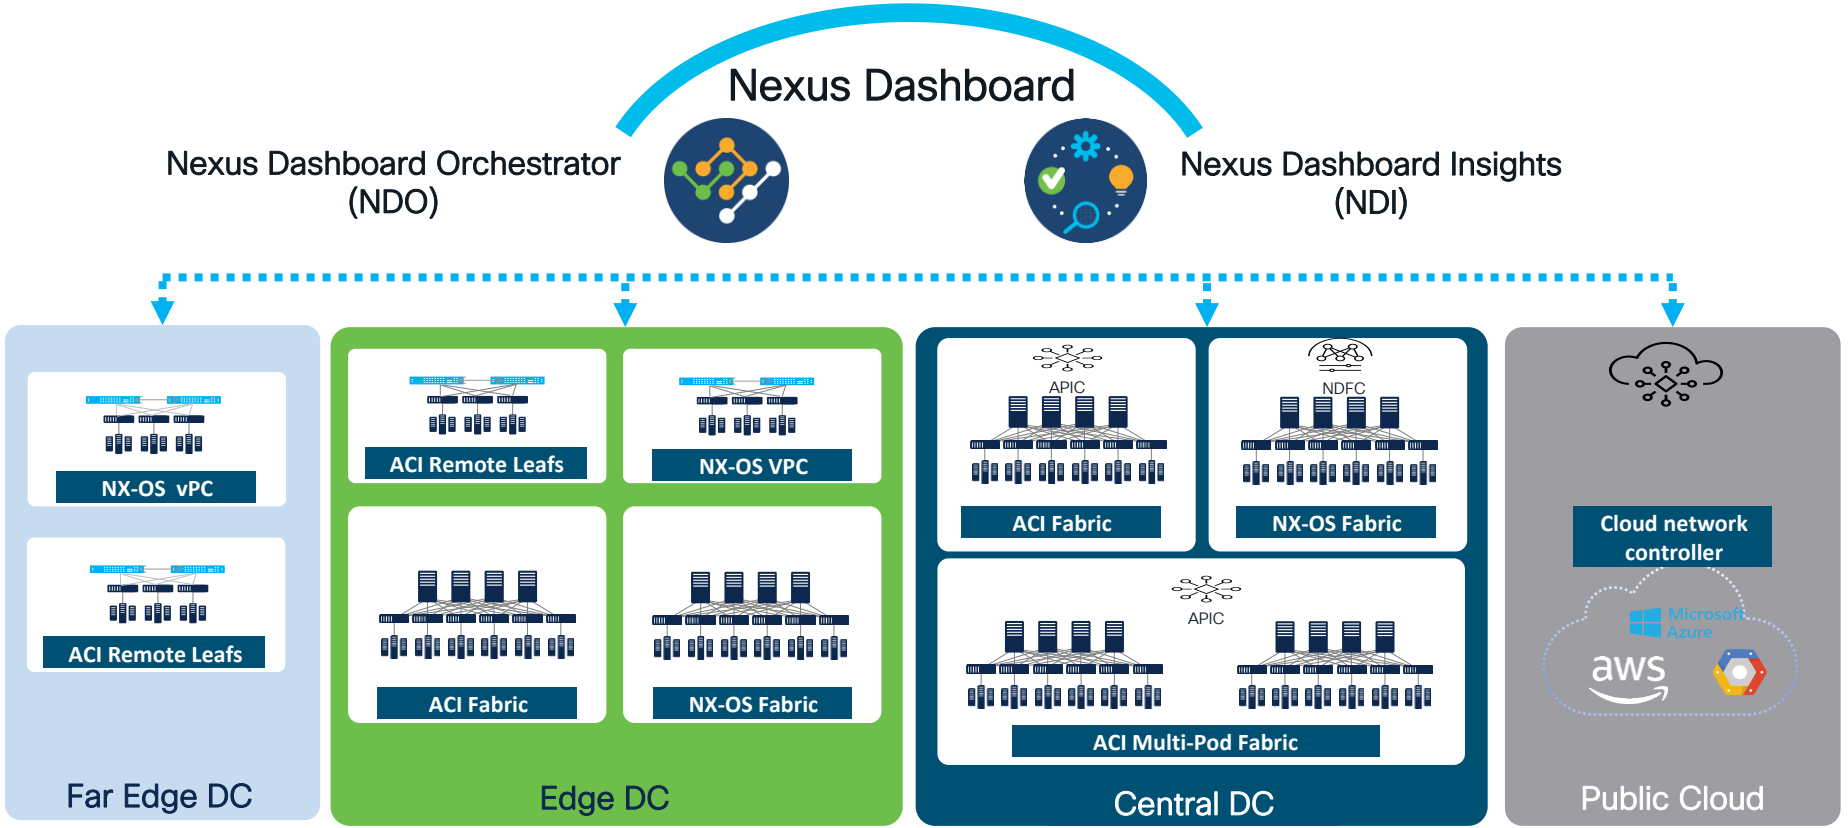

Agenda

- Telco DC trends and challenges
- Solution overview for automation and operations of distributed DCs
- Building distributed Telco DC
- Cloud native packet core integration
- DC to transport handoff
- Cross-domain automation
- Day2ops and integration with OSS/BSS systems
- Customer references
- Conclusion

Building distributed Telco DC with NX-OS fabric

cisco Live!

Automate distributed Telco DCs with NDFC

- Full lifecycle automation with best practices
- Centralized visibility & maintenance
- Support for both brownfield and greenfield deployments

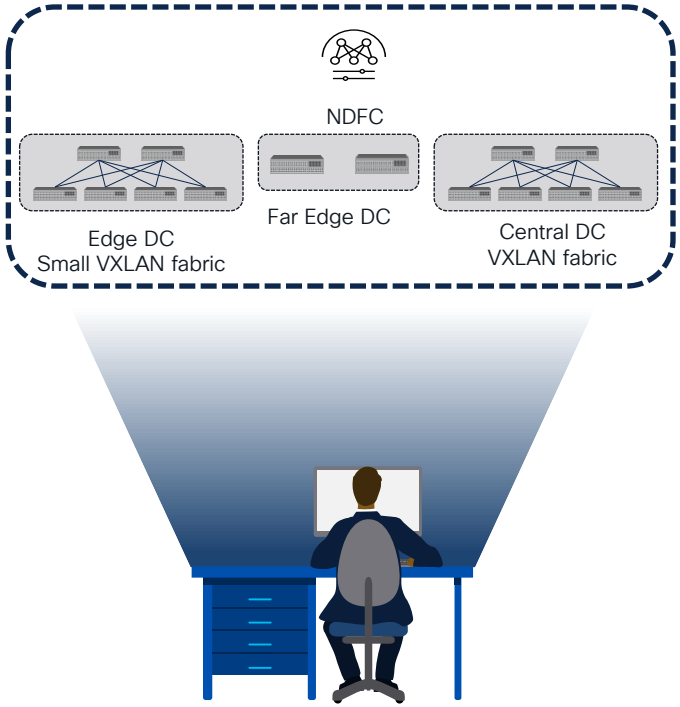

Benefit

Simplify deployment time, reduce chances of errors

Site connectivity: Far Edge DC

- Build small edge sites with just a pair of switches in each site without Spine.
- Far edge sites don't need VXLAN unless L2 extensions are required between far edge sites
- Only IP reachability is required for NDFC to bring up switches, and put relevant network configuration
- NDFC centrally manages whole life cycle of all far edge switches from day-0/day-1 configuration, and basic day-2 ops including software upgrades, performance monitoring, brownfield import

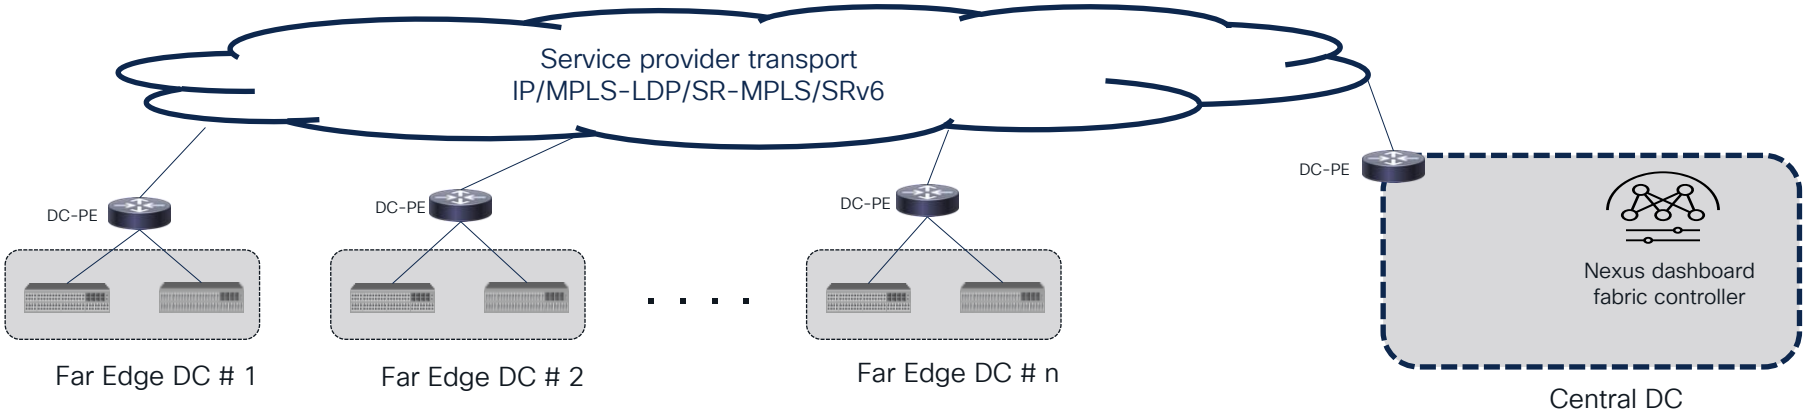

Site connectivity: Edge DC

- A smaller VXLAN fabric is required for relative bigger Edge DC sites
- VXLAN is only used within edge DCs
- DC to transport handoff is required for all external communication
- NDFC is centrally hosted in one of the edge fabric or in a central DC can manage all edge DC sites

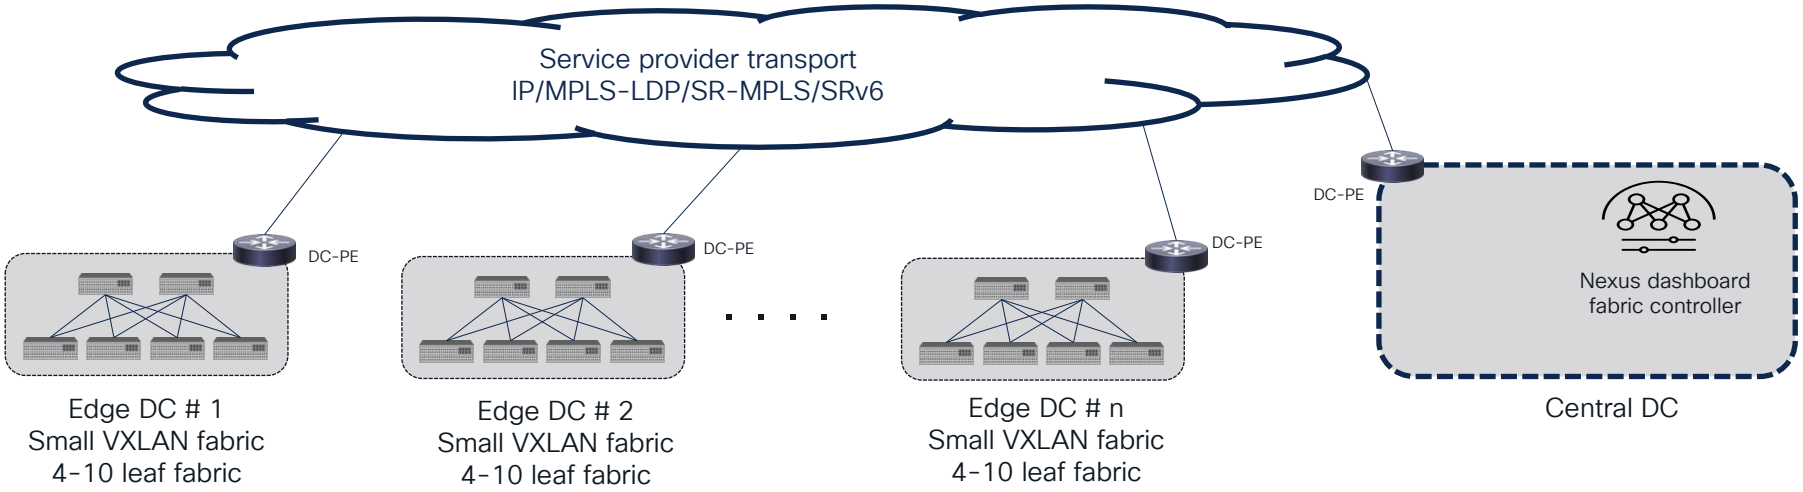

Site connectivity: Far Edge, Edge and central DCs

NDFC supports multiple fabric types (VXLAN, External fabrics, and Enhanced classic LAN etc.) in a single NDFC cluster, this capability allows multiple different DC fabrics to be managed by single NDFC cluster

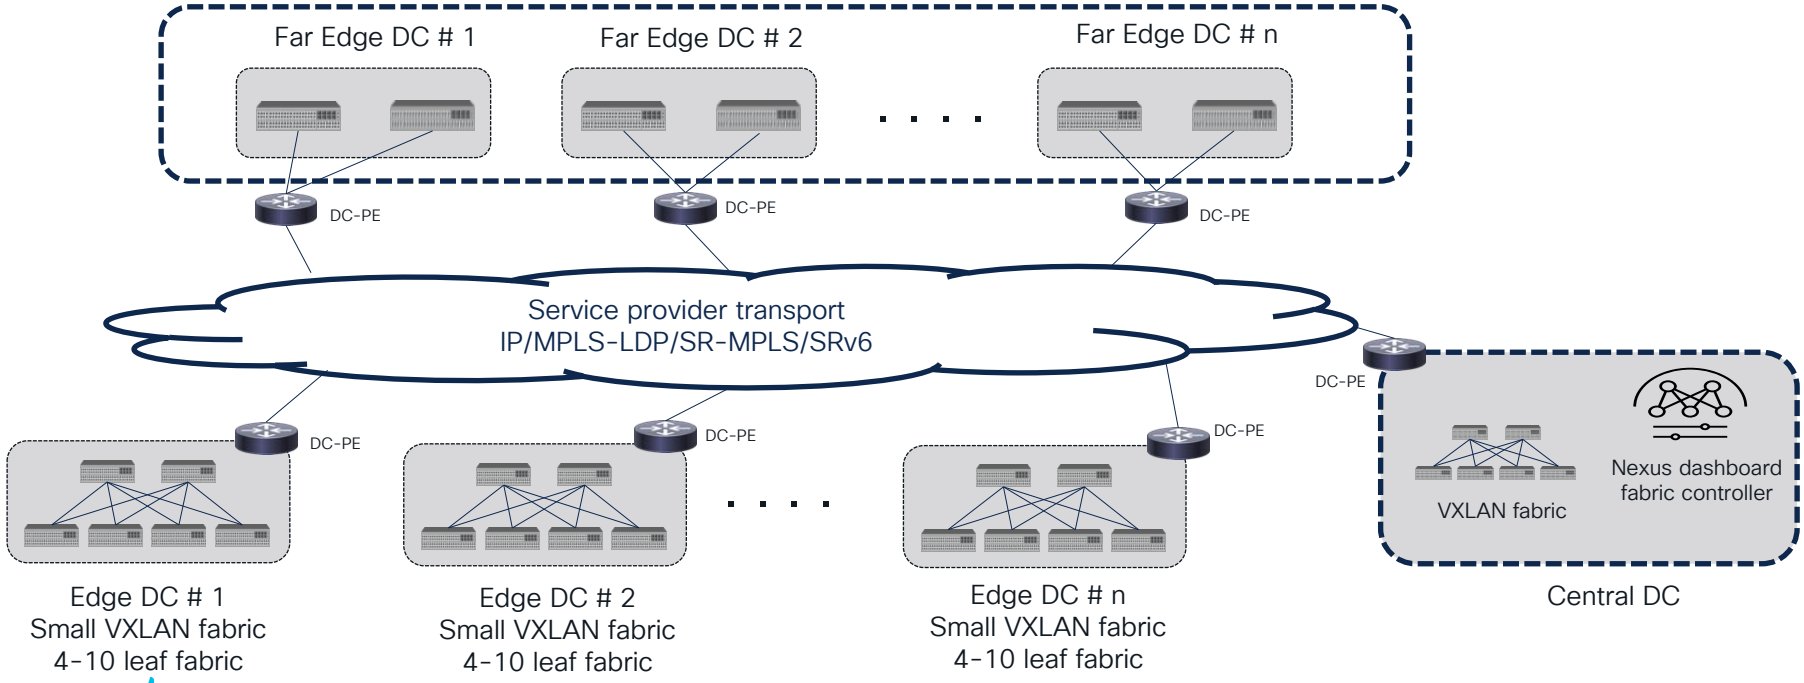

On-prem network extension and scale out with NDO

- Centralized orchestrator for VXLAN multisite capabilities across multiple NDFC instances
- Automates VXLAN-EVPN Day-1 configuration like add/delete/modify VRF and networks
- Provides config deployment visibility

- Config change control - Designer, Approver, Deployer
- Configuration compliance for VXLAN VRF/Network/Attachments using drift/reconcile feature
- Config rollback and versioning

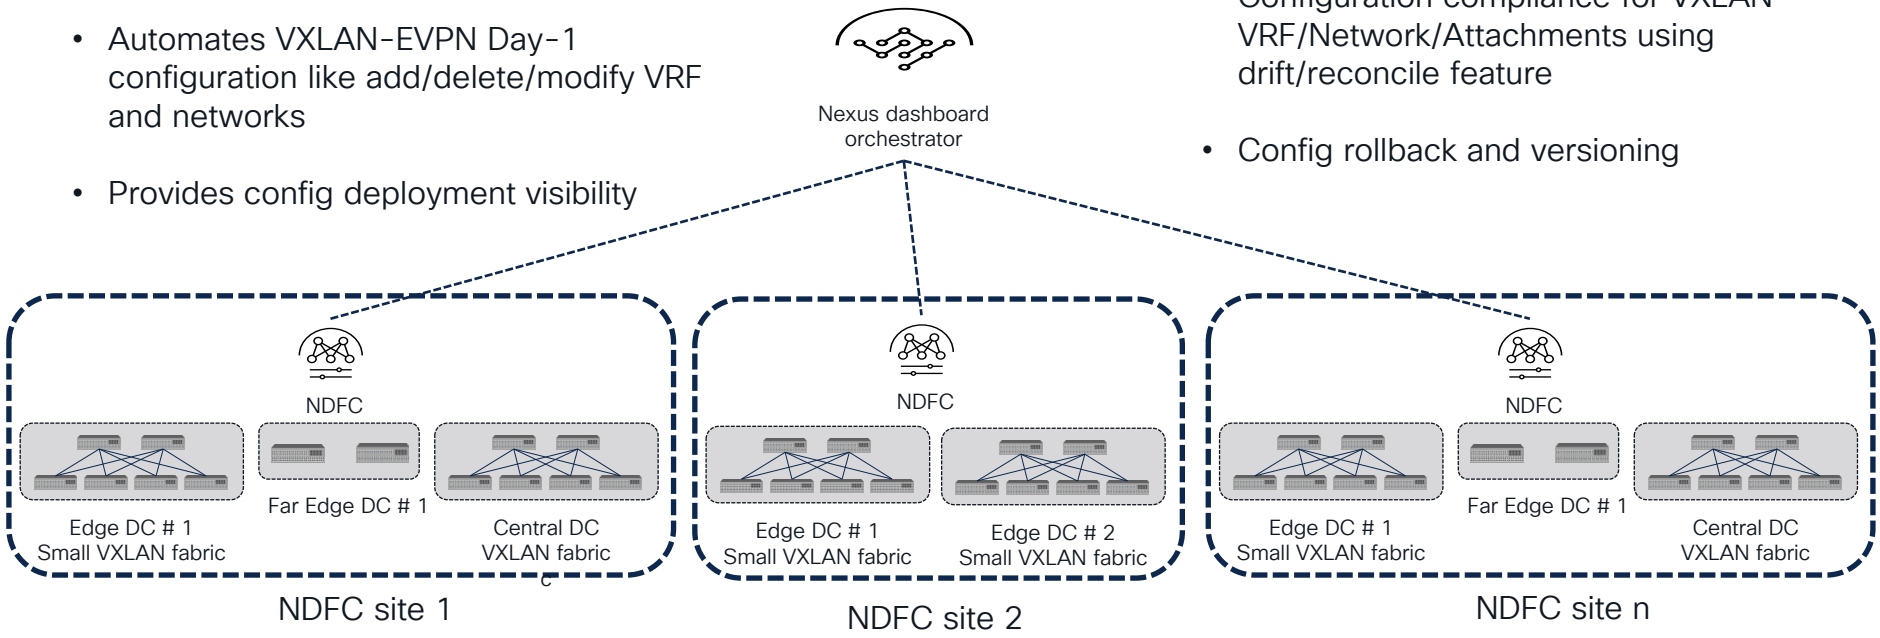

Service chaining with NX-OS ePBR

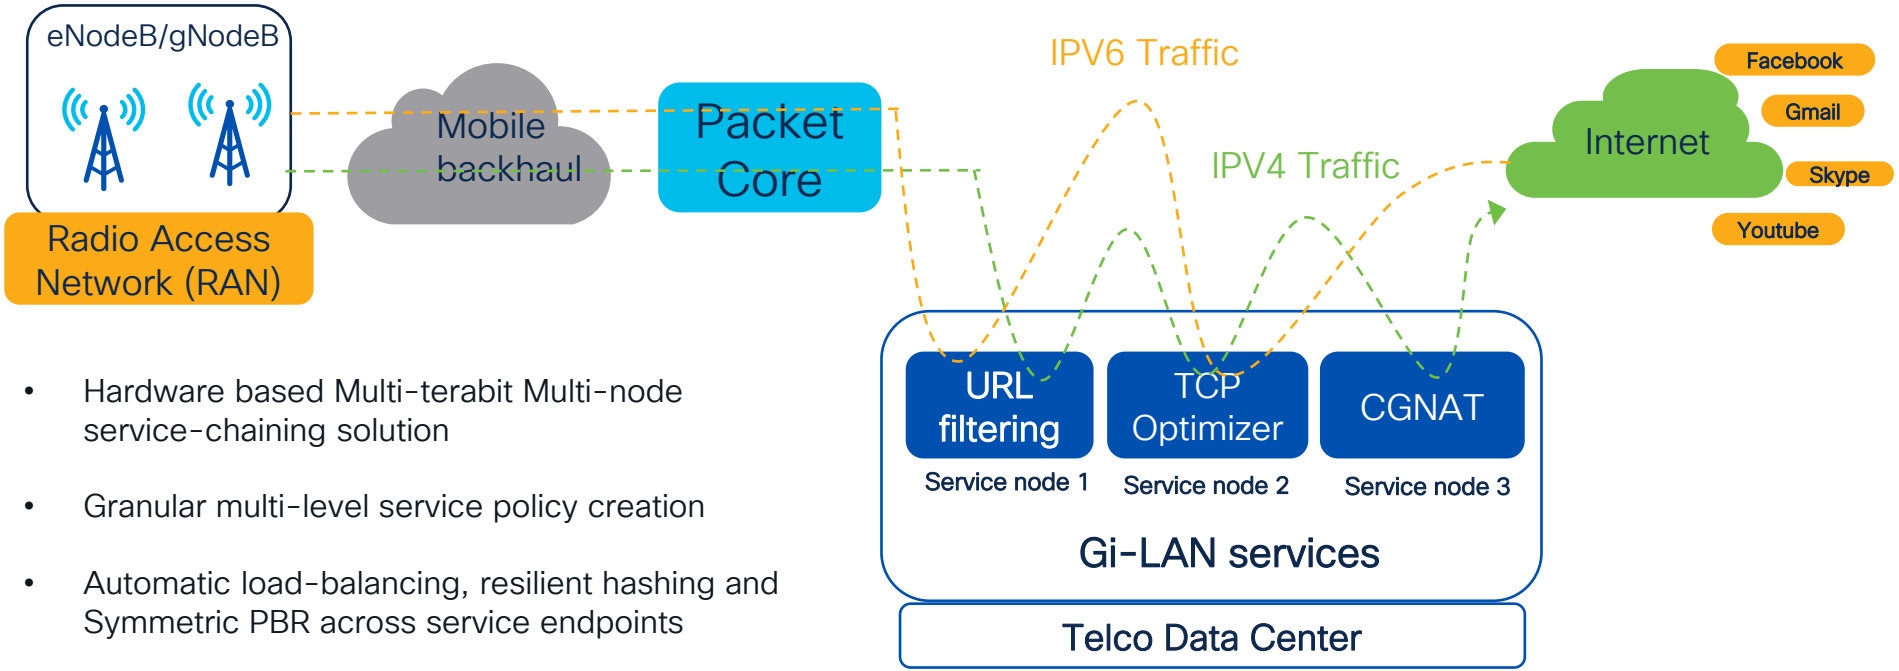

- Hardware based Multi-terabit Multi-node service-chaining solution
- Granular multi-level service policy creation
- Automatic load-balancing, resilient hashing and Symmetric PBR across service endpoints
- Non-disruptive in-service config updates

Latency for deploying ND, NDFC, NDI and NDO

Service	Connectivity	Maximum RTT
Nexus Dashboard cluster	Between nodes	150 ms
Nexus Dashboard Orchestrator	Between nodes	150 ms
	To sites	For APIC sites: 500 ms For NDFC sites: 150 ms
Nexus Dashboard Insights	Between nodes	50 ms
	To switches	50 ms
Nexus Dashboard Fabric Controller	Between nodes	50 ms
	To switches	50 ms
Nexus Dashboard Data Broker	Between nodes	150 ms
	To switches	500 ms

Latency requirement details

Building distributed Telco DC with ACI

cisco Live!

Site connectivity: Far Edge DC

- Deploy ACI Remote leaf (RL) at the far edge DCs
- APIC controllers in central DC would manage discovery of switches, policy push, service chaining, fault/performance monitoring, upgrade/downgrade etc.
- IPN is only used for discovery and control plane traffic, while all external traffic is forwarded via L3out (IP or SR/MPLS)
- Up to 200 RLs are supported per fabric

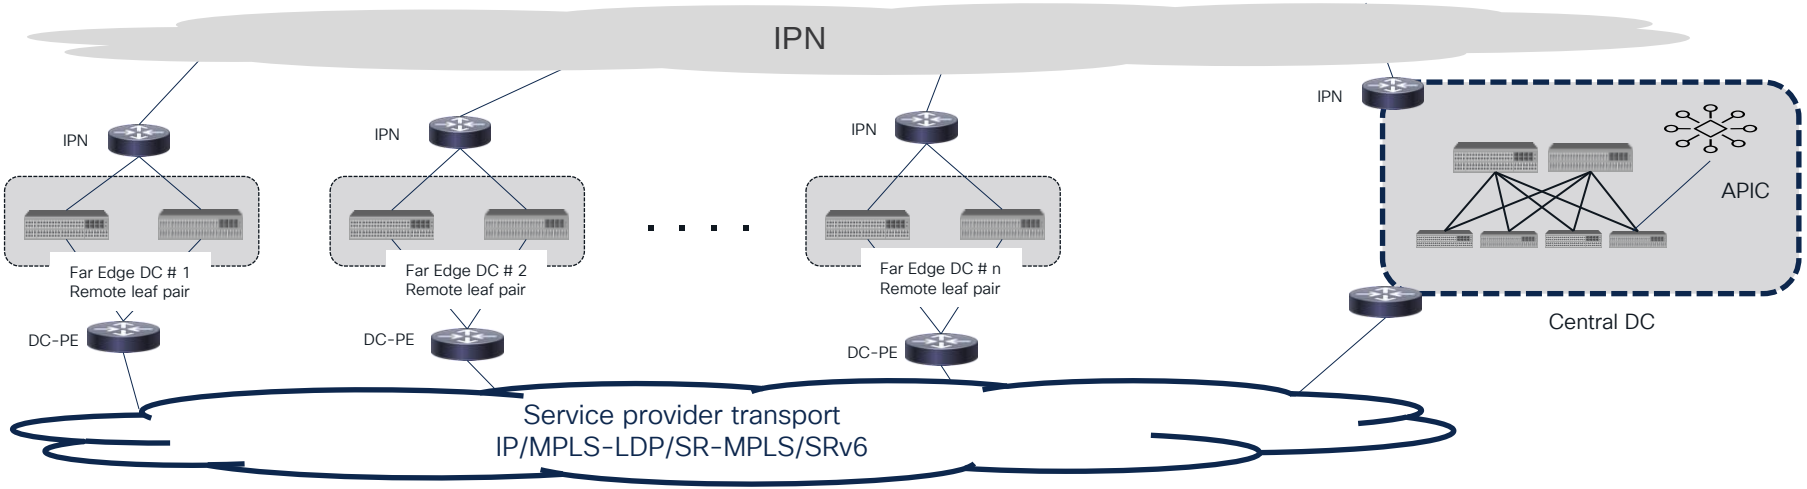

Site connectivity: Edge DC

- Build small Edge DC Pods for small DC edge Pods. Use ACI Multipod solution to manage multiple Pods.
- All Pods don’t need APIC controllers. APIC controllers can be placed in central DCs or be distributed across edge Pods.
- 25 Pods per fabric (APIC cluster) is supported with 500 leaf across all Pods

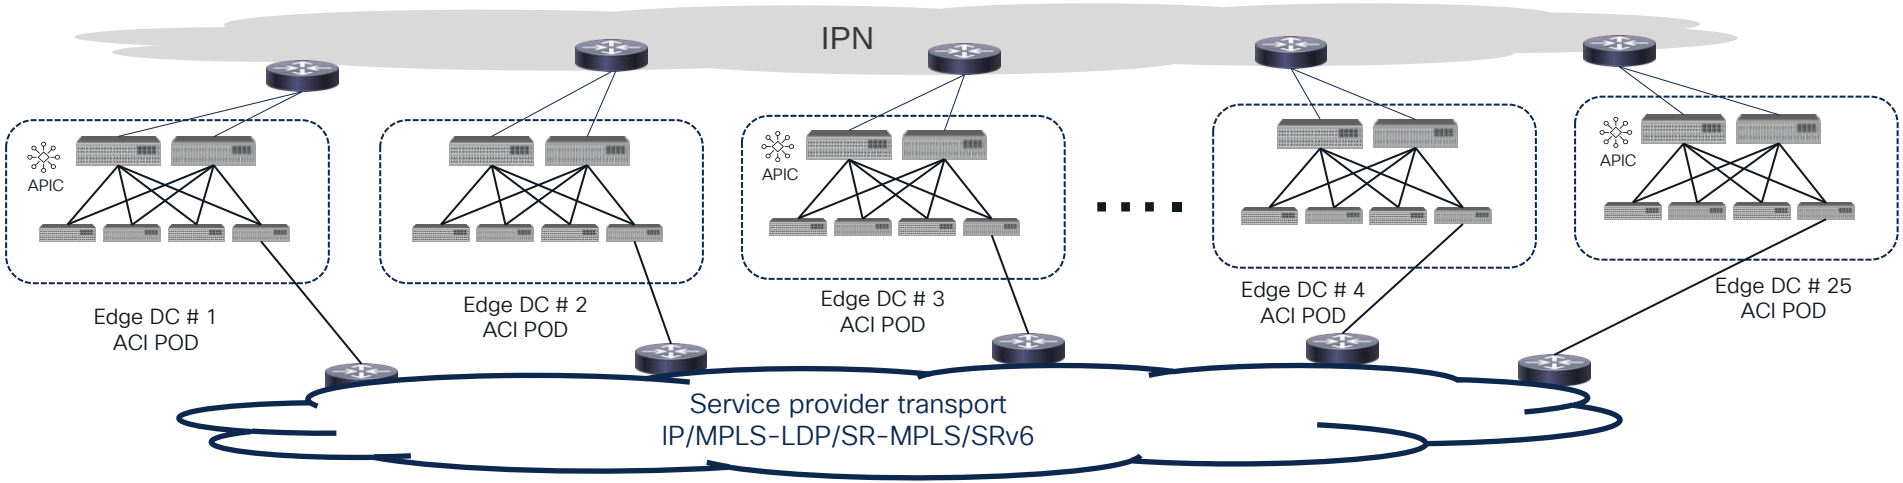

Site connectivity: Far Edge, Edge & Central DC

APIC in central DC managing all DCs – Central, Regional and Edge

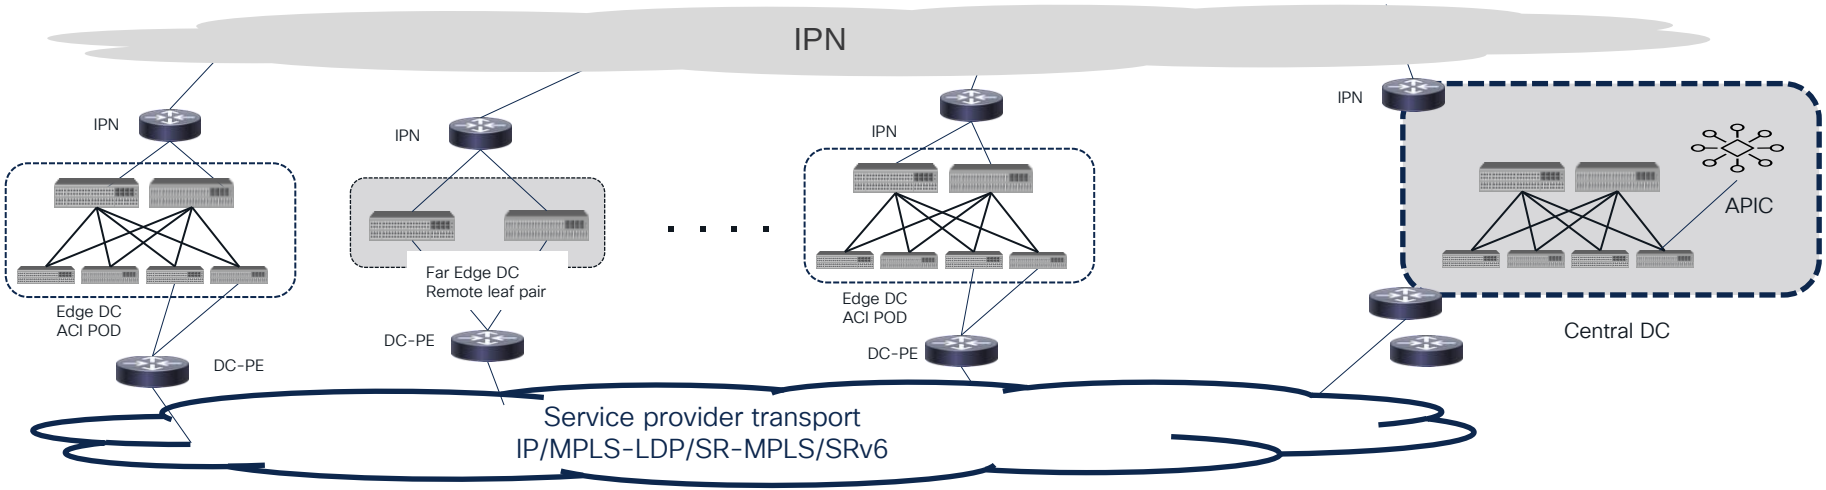

Site connectivity: APIC over L3 network

- Some central DC may not have ACI leaf switches. APIC over L3 enables use-cases to deploy APIC controllers without physically connecting APIC controllers to leaf. In this architecture, APIC can be connected to any standard L2/L3 networking device.
- APIC over L3 supports mix of Remote Leaf and ACI multi-Pod architecture for Edge and Far Edge deployment

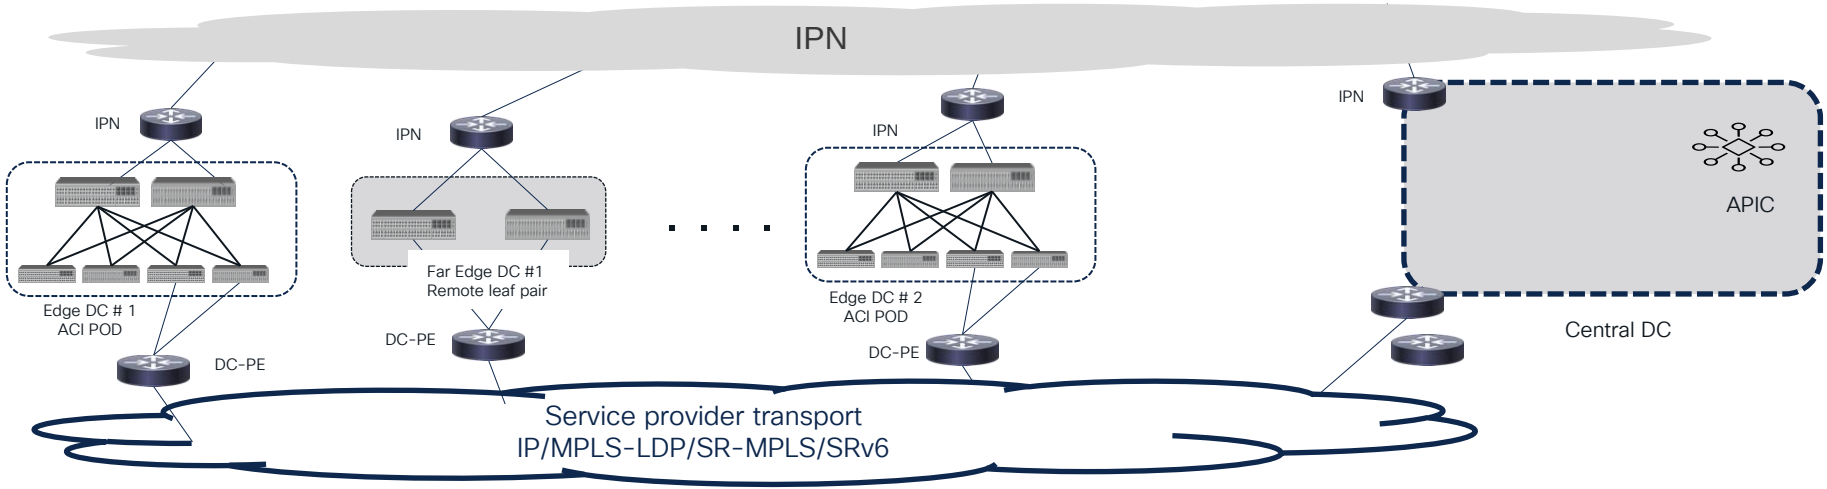

Site connectivity: Virtual APIC

- Virtual APIC controllers may be preferred instead of Physical APICs due to multiple reasons - supply chain constraints or due to a preference for standard virtualized platforms for app hosting etc.
- Stating ACI 6.0.2, virtual APIC is supported for all types of deployment Remote Leaf, Multi-Pod and Multi-site
- Virtual APIC on ESXi/AWS are supported

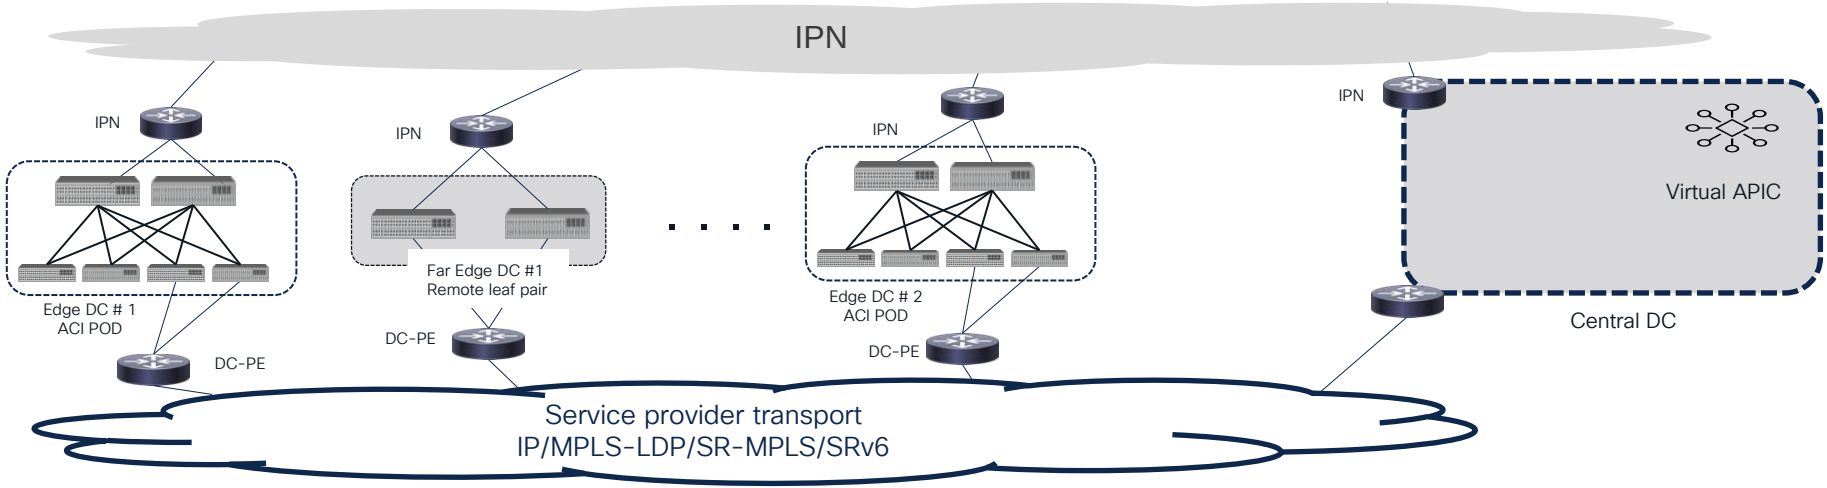

ACI NDO for massively scalable configuration

- Supports 100 autonomous sites with each site supporting 500 leaf, effectively supporting automation across 50K leaf
- Automates Complete far edge and edge configuration – PTP, Synce, SPAN, interface, L3out and QOS
- Greenfield and brownfield support
- Roadmap to support service graph and common services to complete end-to-end telco DC orchestration
- Supports L2/L3 extensions (EPG, BD, VRF, Tenant etc.) for IT use-cases
- Configuration template versioning and drift support

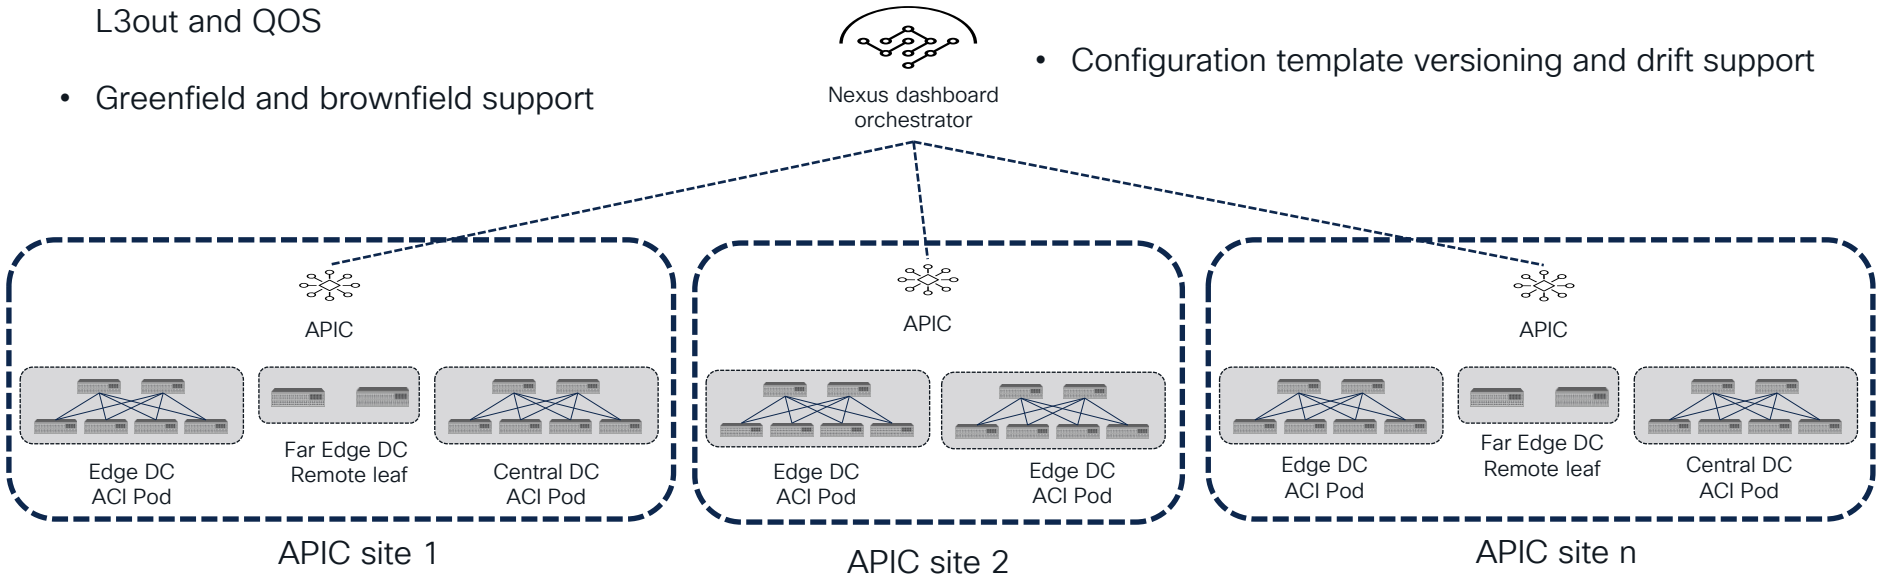

Service chaining with ACI

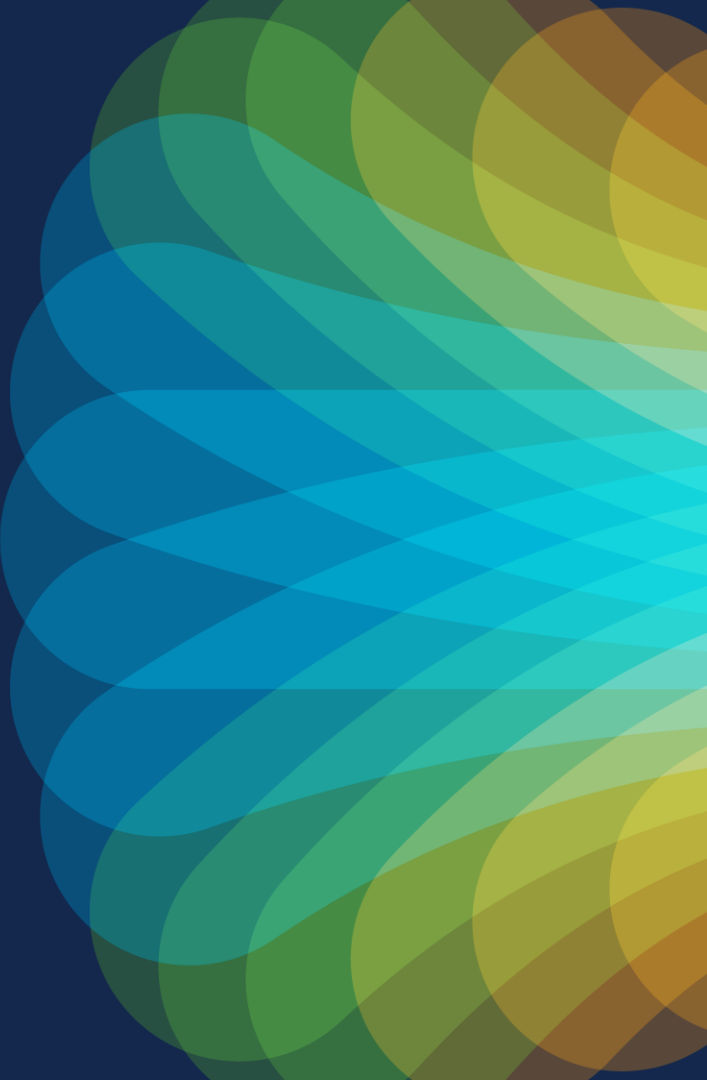

Multi-node service chaining in ACI

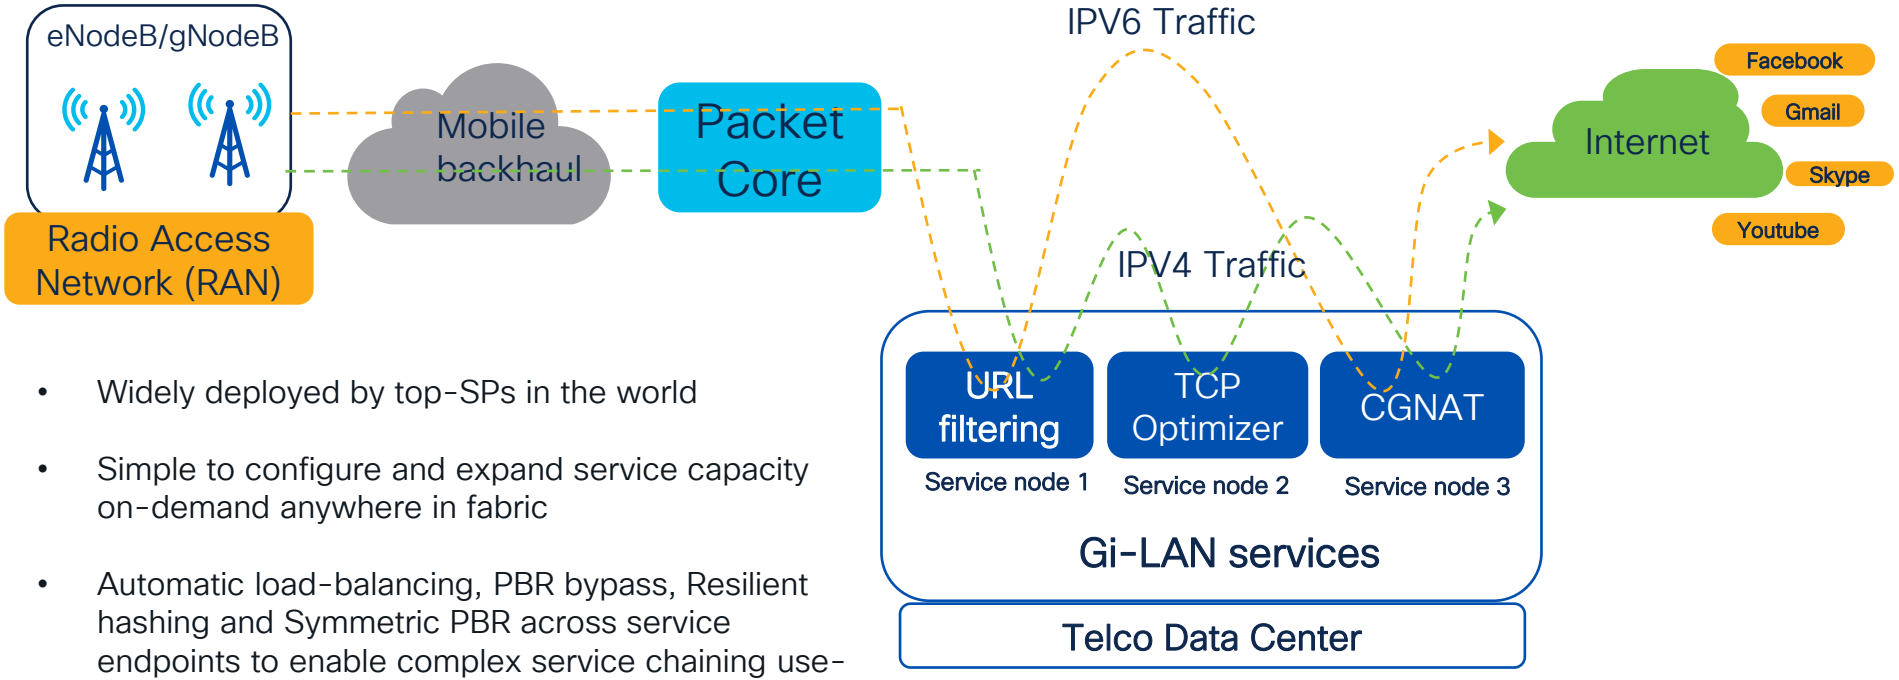

- Widely deployed by top-SPs in the world
- Simple to configure and expand service capacity on-demand anywhere in fabric
- Automatic load-balancing, PBR bypass, Resilient hashing and Symmetric PBR across service endpoints to enable complex service chaining use-cases

Service chaining configuration with ACI

- Create source and destination EPG for Packet core subscriber pool and internet prefixes
- Create a service chain with different nodes based on application requirements
- Use different service chains for different use-cases

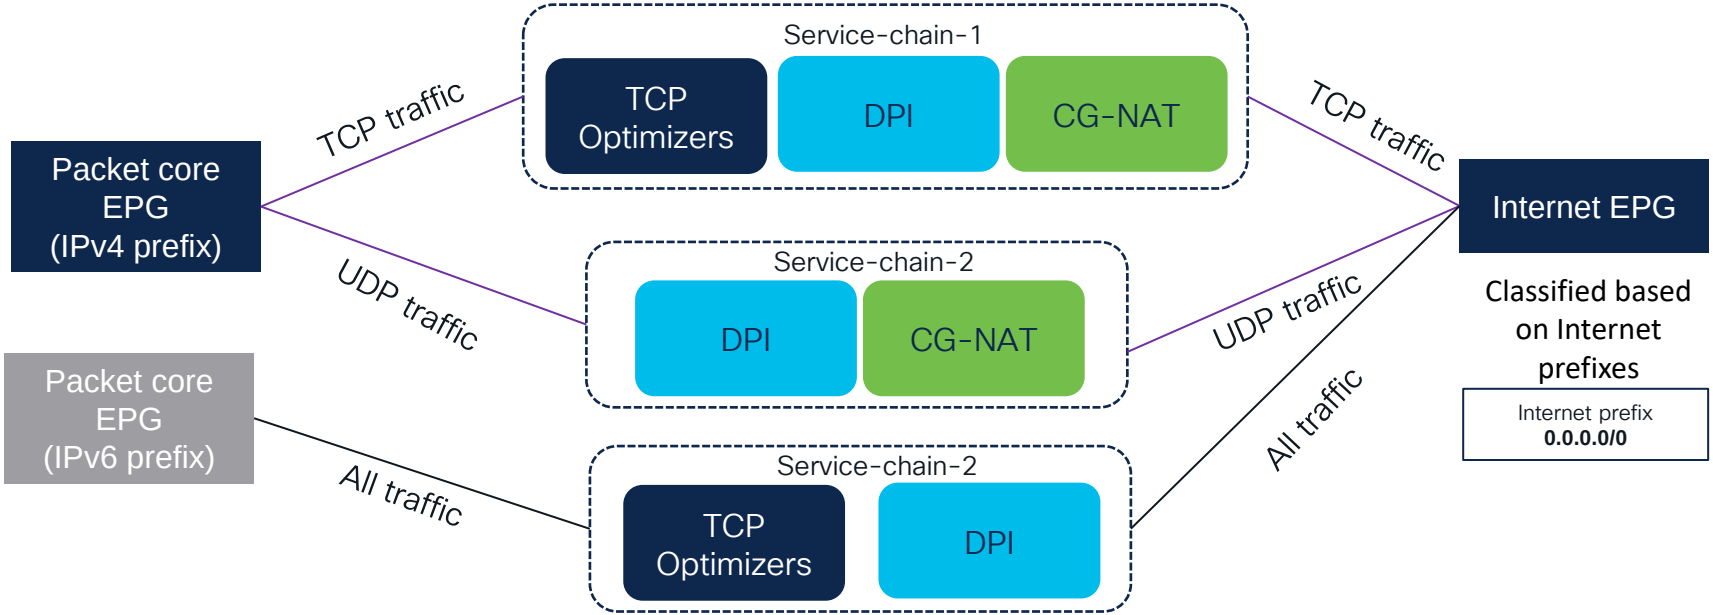

Simplified Configuration

L4-L7 Service Graph Template - service-chaining

Topology | Policy | Faults

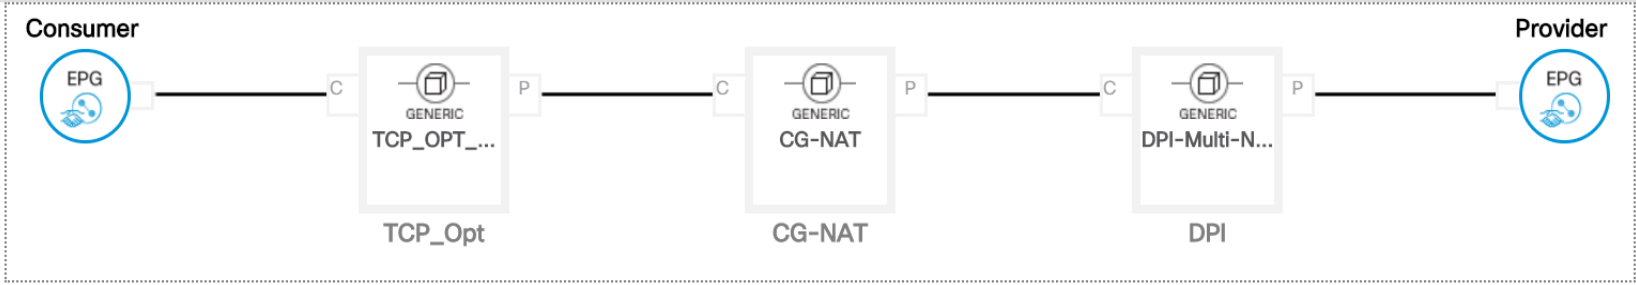

TCP_OPT_Group Information

Route Redirect: true

CG-NAT Information

Route Redirect: true

DPI-Multi-Node Information

Route Redirect: true

Agenda

- Telco DC trends and challenges
- Solution overview for automation and operations of distributed DCs
- Building distributed Telco DC
- Cloud native packet core integration
- DC to transport handoff
- Cross-domain automation
- Day2ops and integration with OSS/BSS systems
- Customer references
- Conclusion

Telco grade fabric for cloud native packet core

- **Cisco DC fabric is massively deployed in production for 4G and 5G services with top packet core vendors**
- Cloud native packet core requires following features since packet core is deployed with several small containers instead of single appliance. Cisco DC fabrics already support these features
  - Huge route scale (512K v4 and v6 prefixes)
  - High ECMP (128-way ECMP)
  - 200msec convergence with L3 fabric to support hitless operations
  - 2K Routing and BFD peering per leaf
  - Flexible Architecture, where NF can be instantiated anywhere

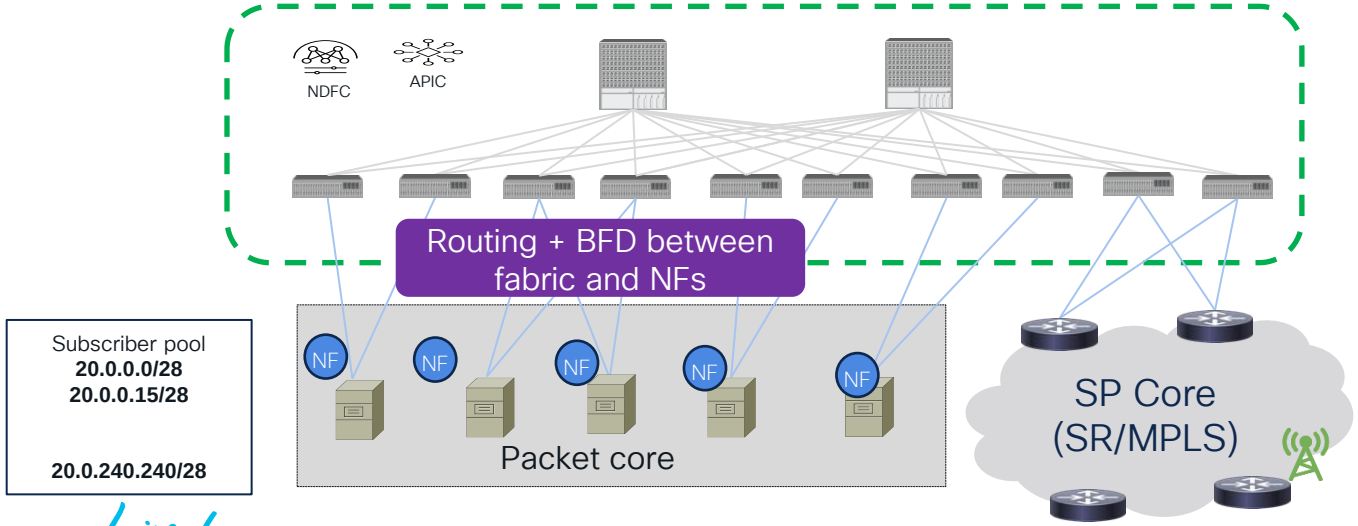

Why do I need to pay special attention for packet core fabric?

Typical Enterprise application deployment

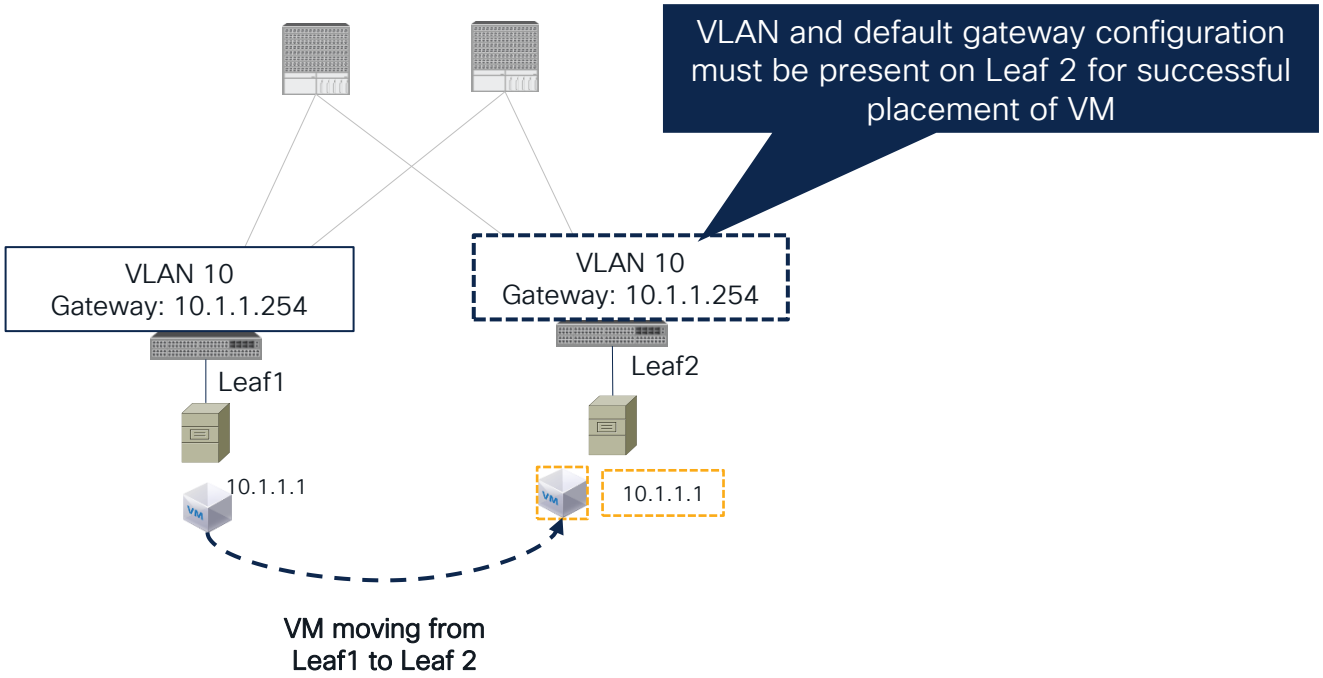

Cloud native packet core deployment

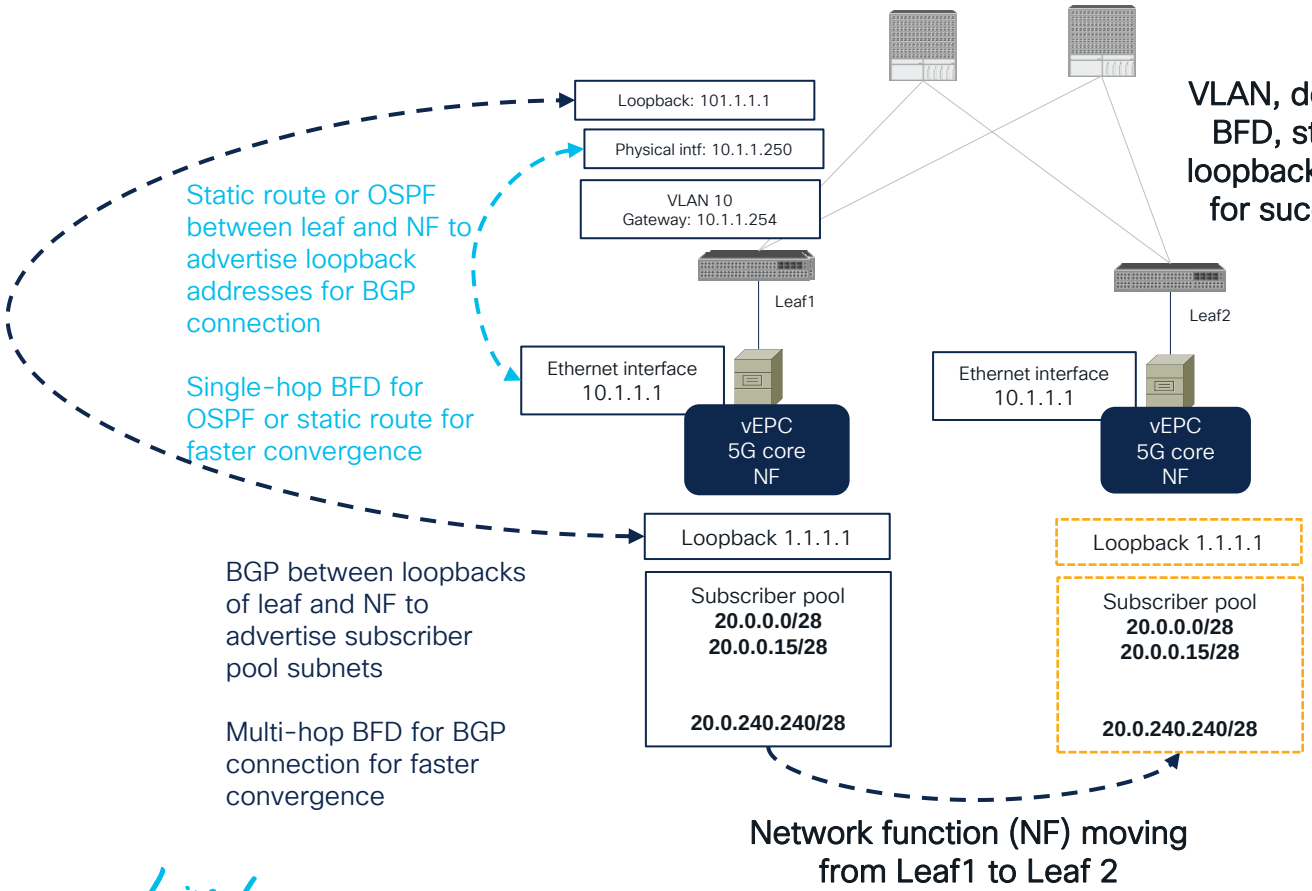

Cisco innovations in building telco grade fabric for packet core deployments

Centralized route peering for packet core deployment in NX-OS

Centralized Routing Peering

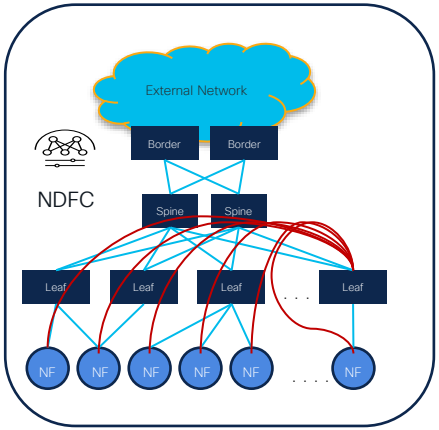

NF integration :
Local & Centralized route peering options
with BGP and BFD

NF Mobility

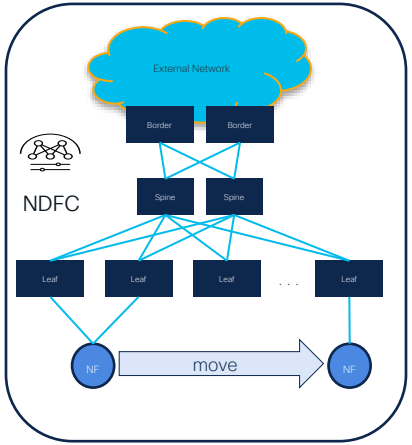

Seamless NF mobility to new leaf
VNF, CNF or PNF connectivity to fabric.

Proportional load-balaning

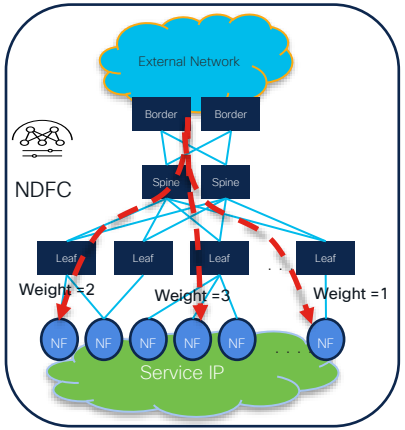

Scale NFs with capability to optimally load-share traffic

Packet core design with ACI Floating L3out

- Widely deployed by leading service providers
- Local traffic forwarding – Even though routing relationship is with Anchor leaf, both incoming and outgoing traffic is locally forwarded by each Leaf

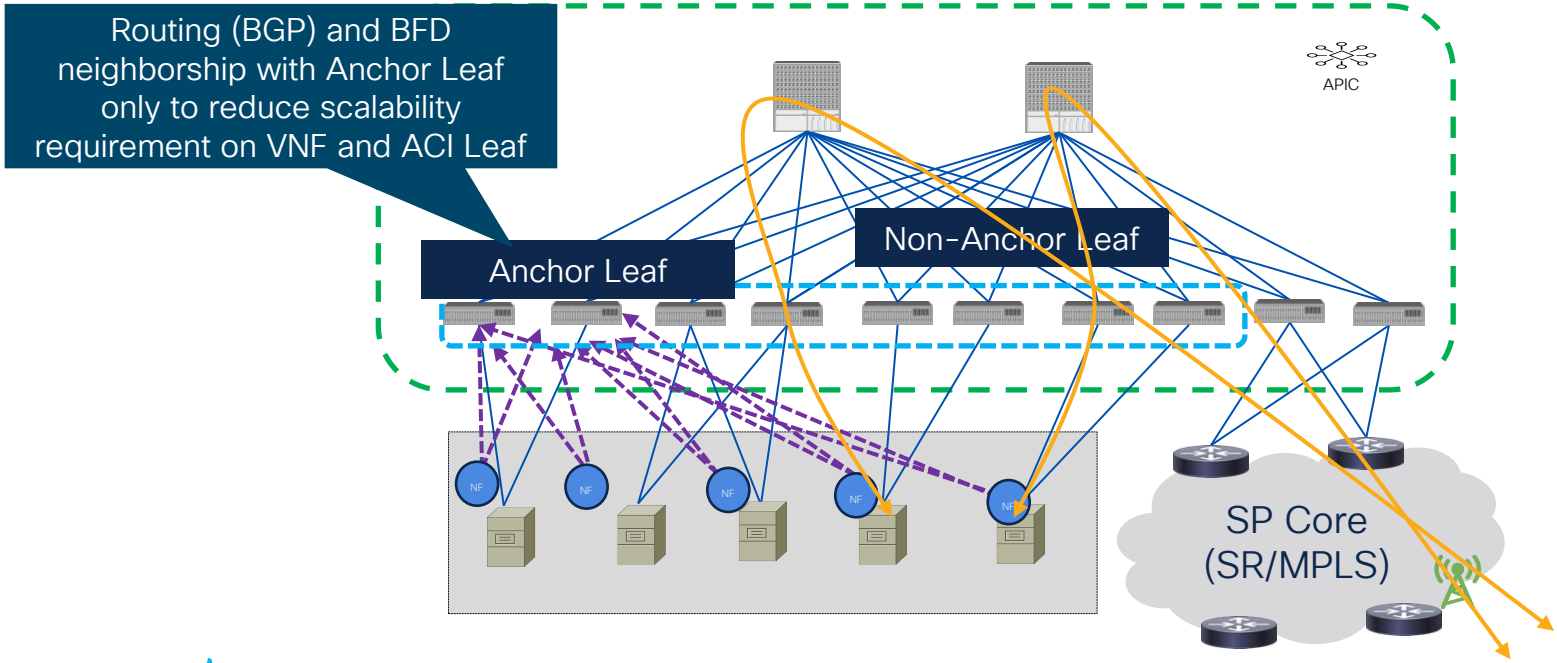

Why should I build L3 fabric for packet core?

Who asked for L2 fabric and why?

- Packet core application vendors pushed for L2 design in the past to keeping fabric simple
- Pushed service providers to deploy L2 fabric solution through validated designs specially for 5G packet core (UPF)

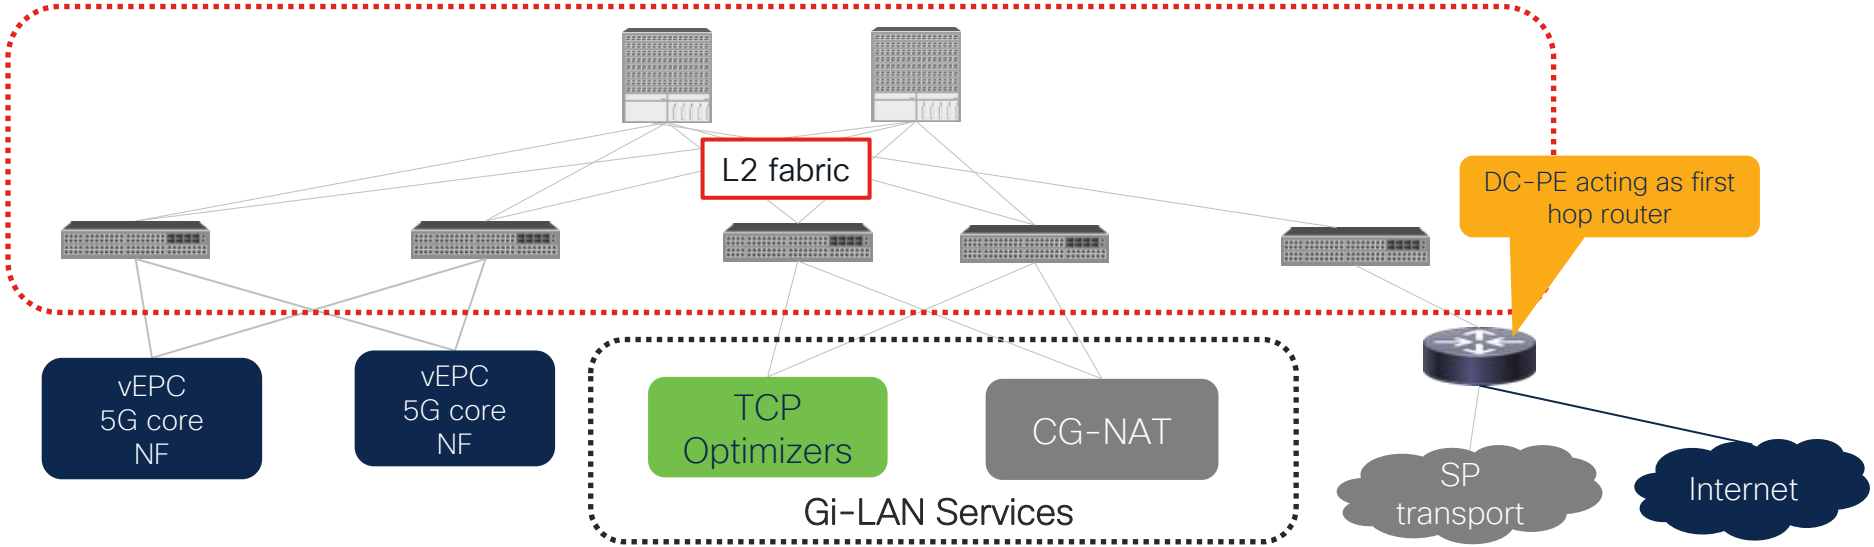

Scale challenges with L2 fabric

L2 fabric creates huge scale requirement for bandwidth, routing, and forwarding scale on DC-PE since all routing sessions are terminated on it and it's acting as a first hop router.

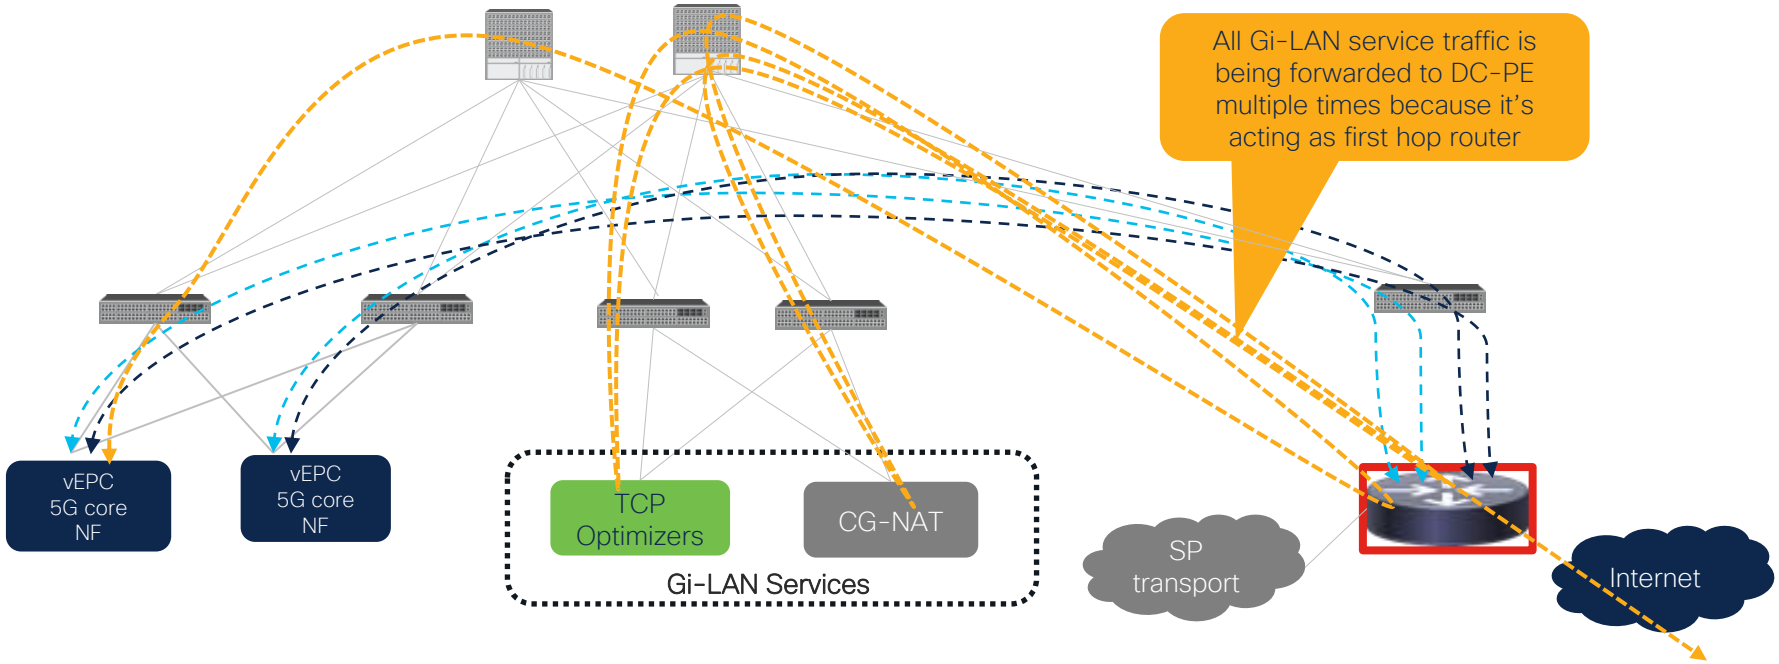

Horizontal scale with L3 fabric

Recommended solution for packet core deployment

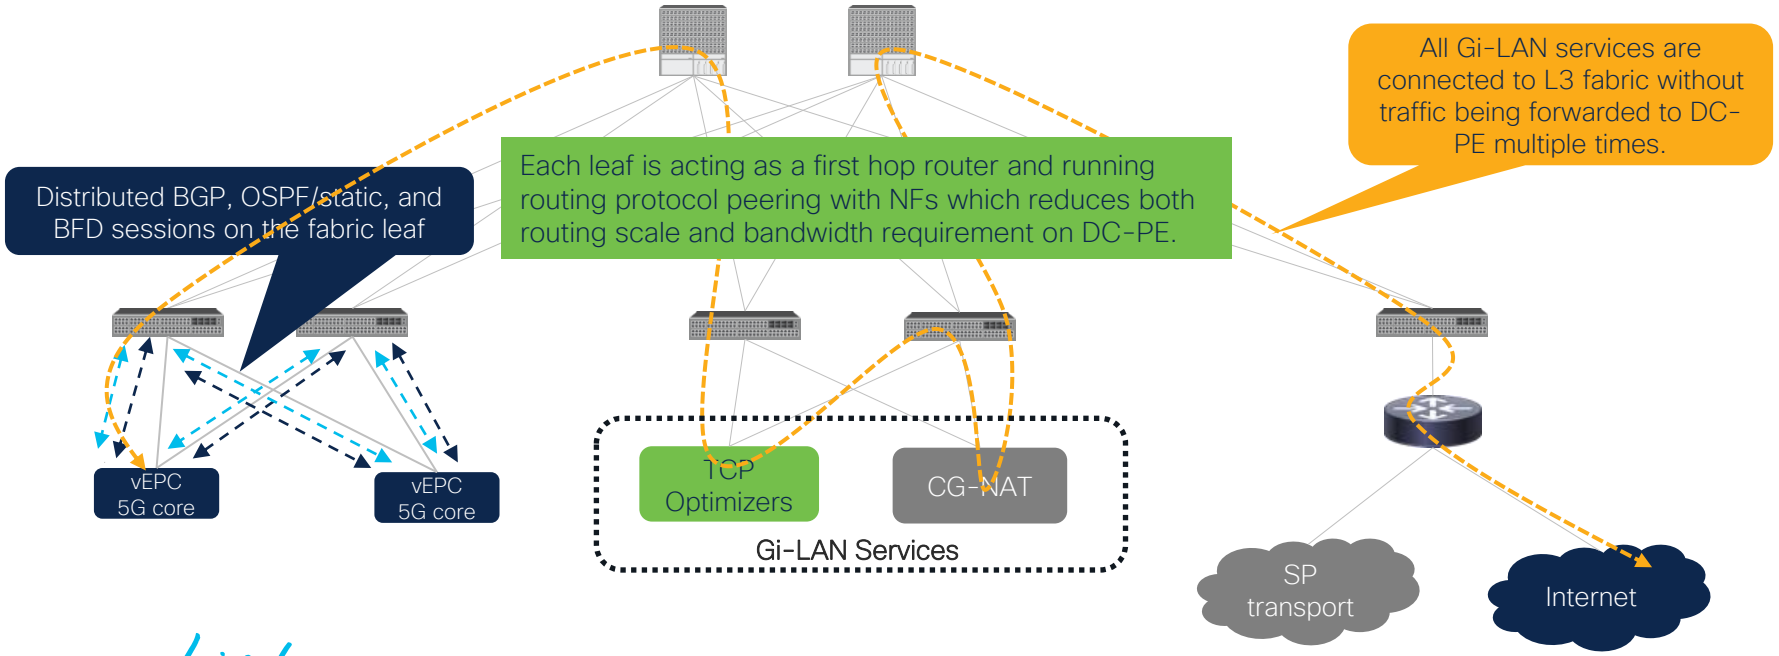

Agenda

- Telco DC trends and challenges
- Solution overview for automation and operations of distributed DCs
- Building distributed Telco DC
- Cloud native packet core integration
- DC to transport handoff
- Cross-domain automation
- Day2ops and integration with OSS/BSS systems
- Customer references
- Conclusion

Practical reasons for DC handoff to transport

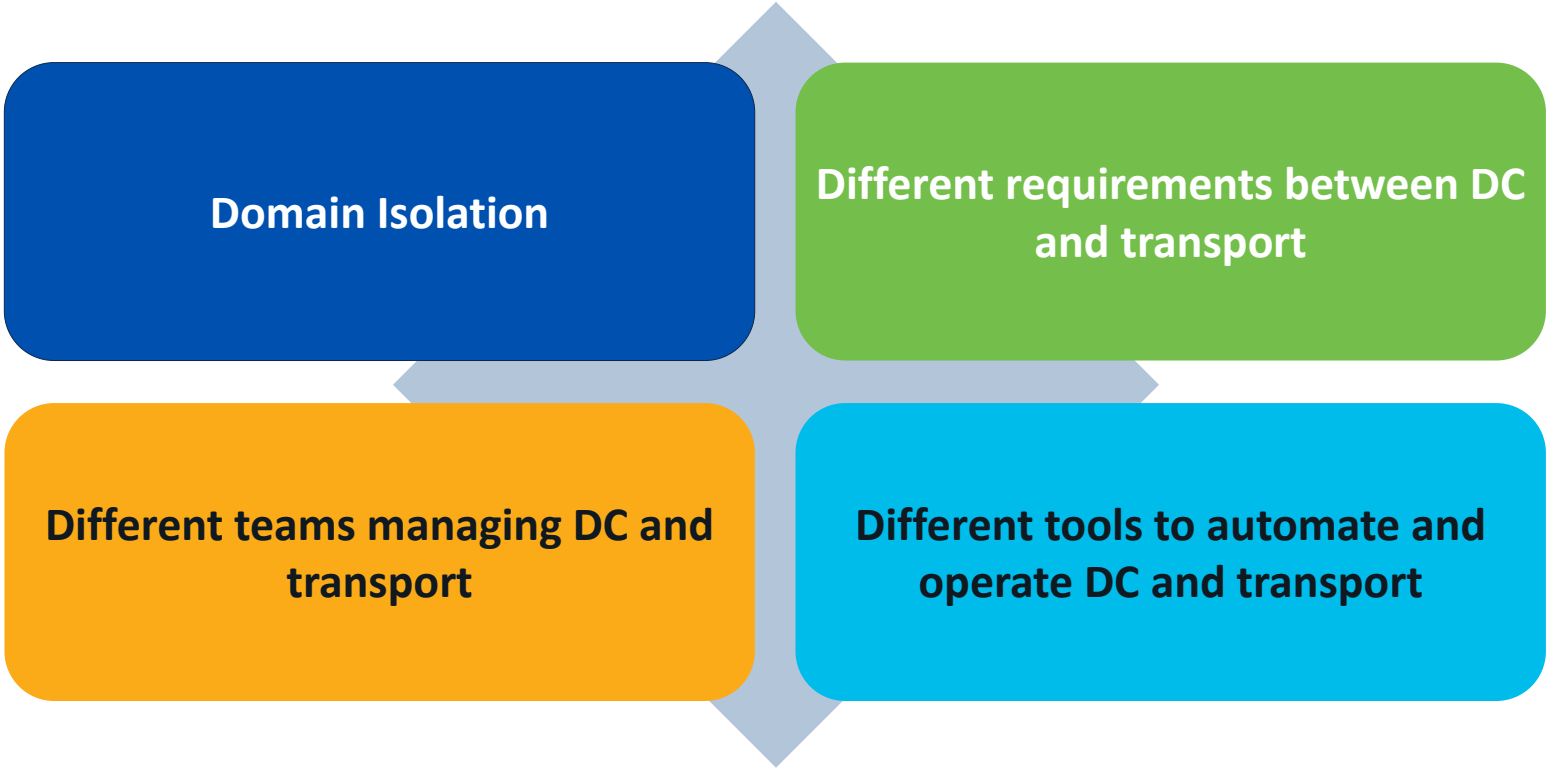

DC requirements

NF movement and flexible placement	Automation of distributed DC configuration across 100s of fabric and 1000s of switches	E<->W communication between application endpoints	ECMP scale since single service is being handled by multiple smaller containers or VNFs
Service chaining	Horizontal scale for endpoints, bandwidth, routing etc.	Integration with south bound VMM domains	End point visibility into fabrics

Transport requirements

Traffic engineering	Quality of service	Vertical routing scale
Fast reroute (50 msec convergence)	Multi-vendor SP core network	Multi-vendor automation and operation tool to enable end-to-end network slicing

Handoff for best of both worlds

Datacenter
VXLAN-EVPN / ACI

- ✓ Default standard for DC deployments
- ✓ Mature automation and operations tools for DC use-cases
- ✓ Scale and DR solutions- with Single Pod (availability zone), Multiple Pod (Multi-AZ) and Multi-site (Multiple regions) with L2/L3 extensions
- ✓ Mature and proven service chaining use-cases
- ✓ Distributed routing, forwarding and horizontal scale for CNFs & VNFs
- ✓ Flexible placement and Mobility of Workloads

Transport
SR-MPLS/MPLS-LDP/SRv6

- ✓ Network slicing
- ✓ Traffic engineering
- ✓ Quality of service
- ✓ Vertical routing scale
- ✓ Fast reroute (50 msec convergence)
- ✓ Multi-vendor SP core network
- ✓ Multi-vendor automation and operation tool to enable end-to-end network slicing

DC to transport Handoff

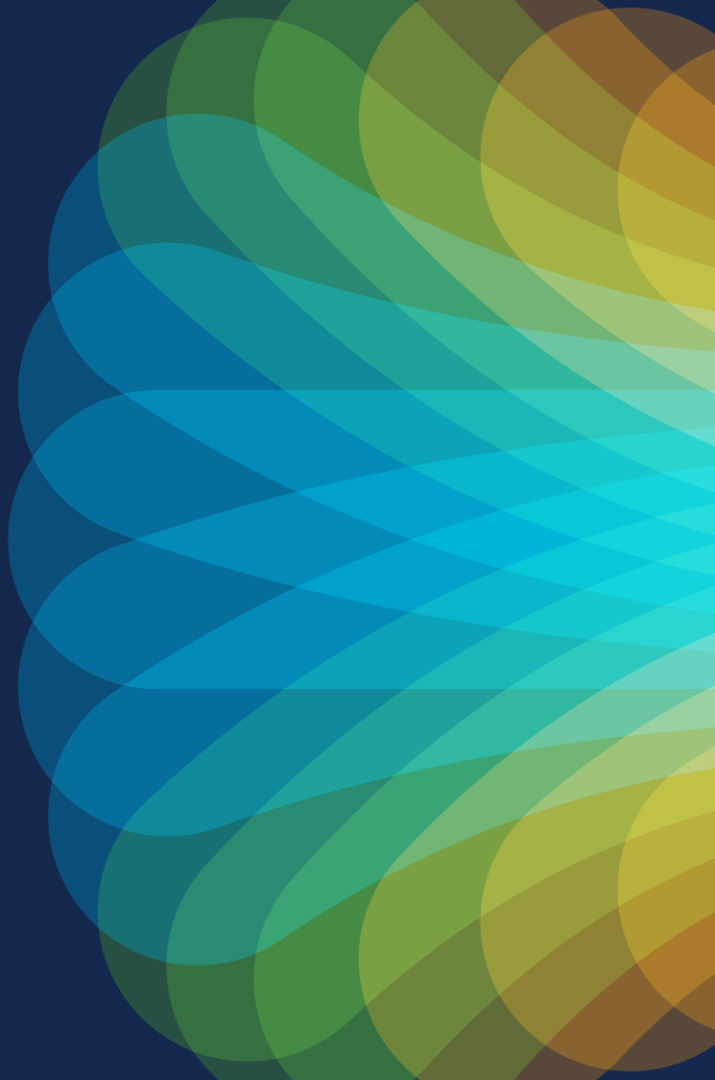

DC to transport handoff with VRF-lite

- Interface and routing protocol session per VRF between BL and DC-PE
- Simple solution to connect DC and transport that allows any type of transport datapath encapsulation (SR-MPLS, LDP or SRv6)
- Automation and scalability are key challenges in this solution due to per VRF routing protocol and sub-interface configuration
- Supported on all hardware platforms

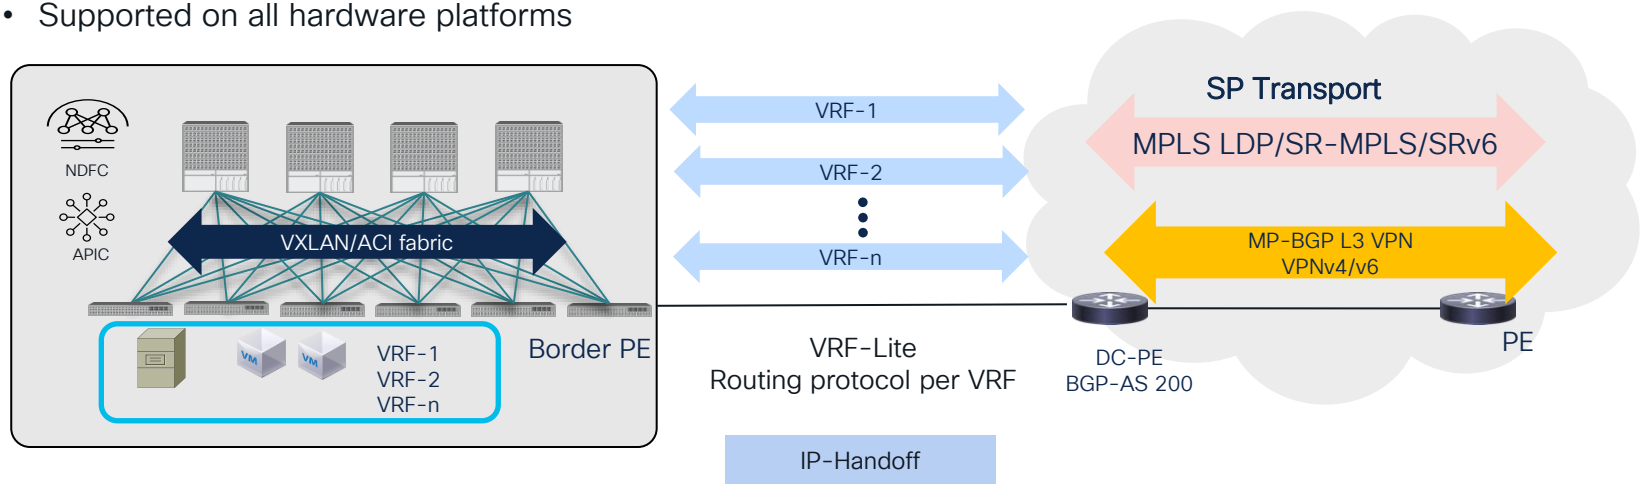

VXLAN EVPN to transport handoff with SR-MPLS

- Better scale with single control plane and data plane session instead of per VRF route peering and sub-interface
- Unified SR-MPLS transport network
- Border PE encapsulates packet with SR-MPLS VPN label based on incoming VXLAN L3VNI
- Border PE encapsulates packet with VXLAN L3 VNI based on incoming SR-MPLS VPN label

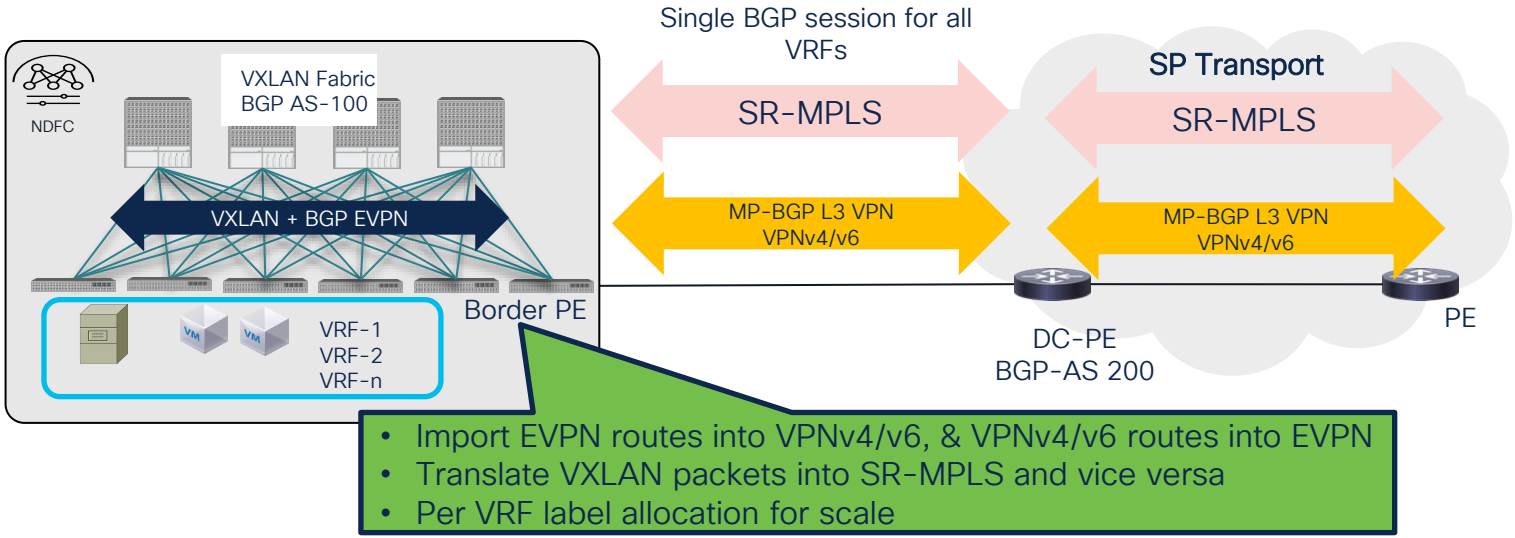

VXLAN EVPN to transport handoff with SRv6

- Better scale with single control plane and data plane session instead of per VRF route peering and sub-interface
- Unified SRv6 transport network
- Border PE encapsulates packet with SRv6 SID based on incoming VXLAN L3VNI
- Border PE encapsulates packet with VXLAN L3 VNI based on incoming SRv6 SID

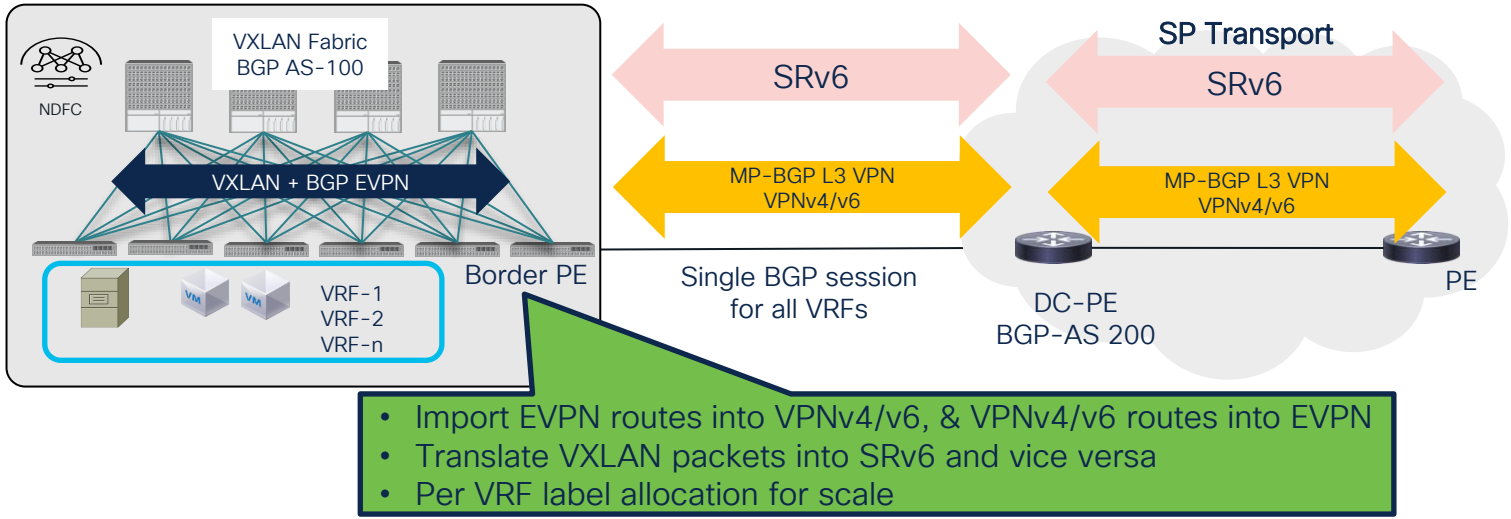

ACI to transport handoff using SR-MPLS handoff

- Better scale with single control plane and data plane session instead of per VRF route peering and sub-interface
- Unified SR-MPLS transport network
- Border Leaf encapsulates packet with SR-MPLS VPN label for each VRF in ACI fabric
- Border Leaf encapsulates packet with ACI VXLAN header based on incoming SR-MPLS VPN label

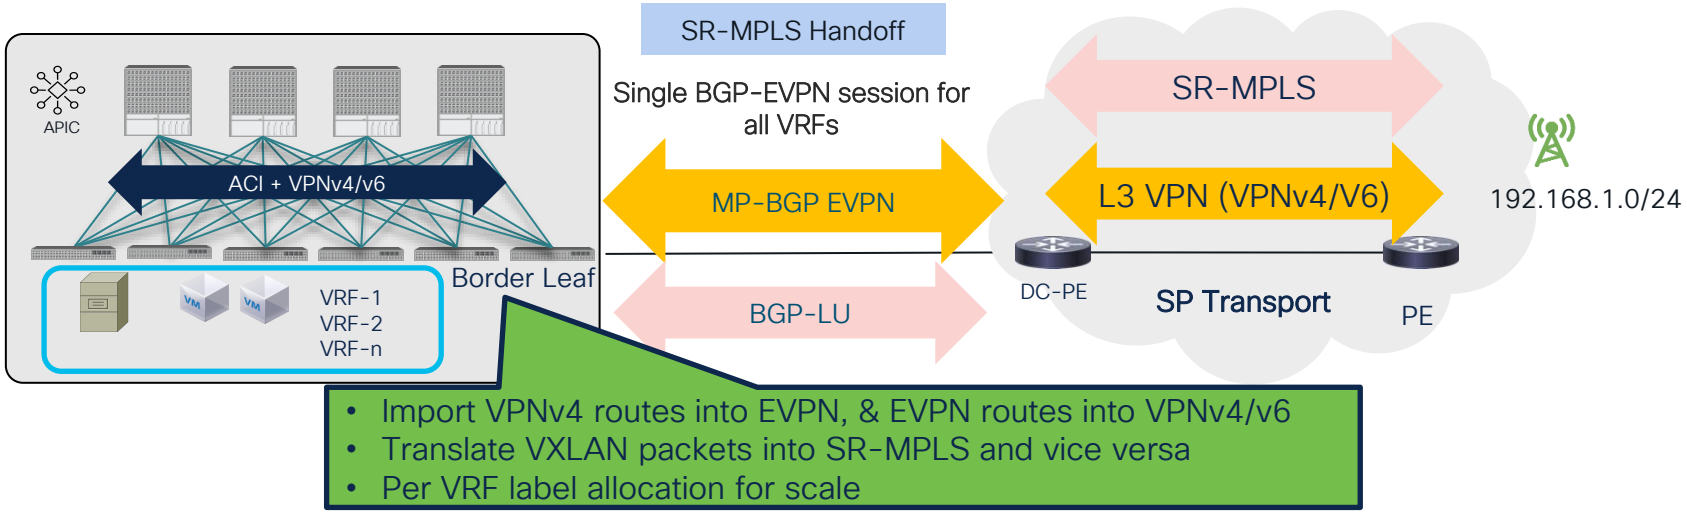

ACI to transport handoff for SRv6 based network

- Solution utilizes ACI to SR-MPLS handoff option
- BL does ACI to SR-MPLS handoff, and DC-PE does SR-MPLS to SRv6 handoff.
- DC-PE encapsulates the packet with SRV6 SID based on incoming SR-MPLS VPN label in the packet.
- DC-PE encapsulate the packet with SR-MPLS VPN label based on SRv6 SID in the packet

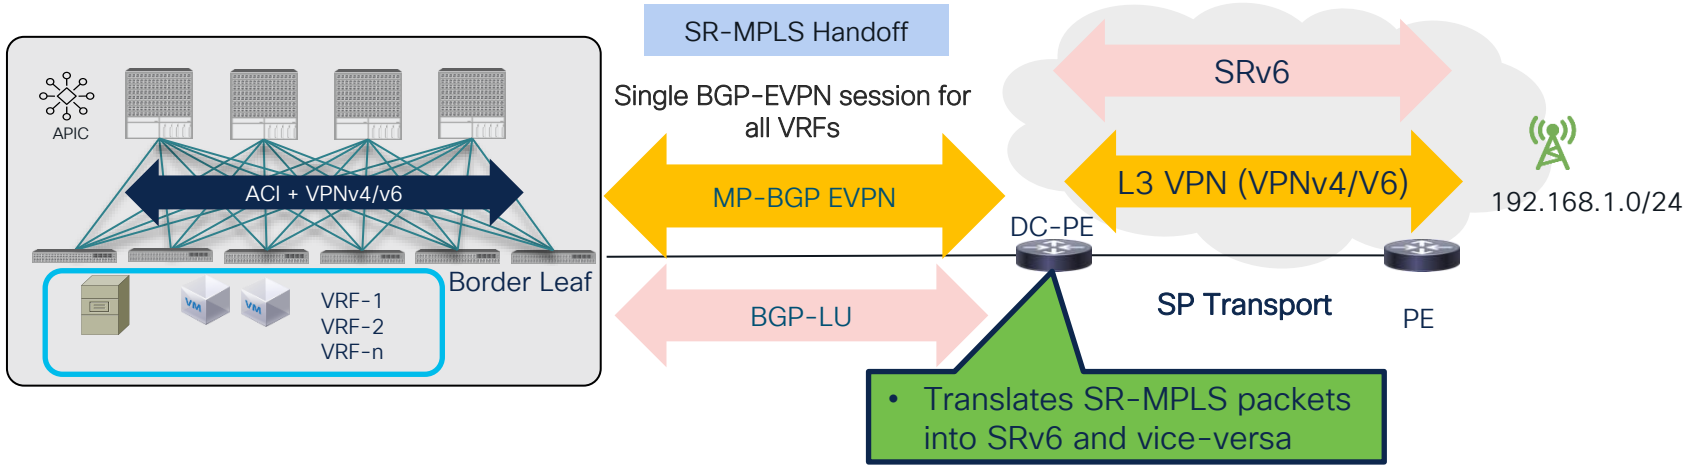

Agenda

- Telco DC trends and challenges
- Solution overview for automation and operations of distributed DCs
- Building distributed Telco DC
- Cloud native packet core integration
- DC to transport handoff
- Cross-domain automation
- Day2ops and integration with OSS/BSS systems
- Customer references
- Conclusion

What is cross-domain automation?

Applications:
Open RAN, 5G, IOT, Edge

- ✓ Applications are hosted in the DC
- ✓ Application to application and user communication is through SP transport network

Cross-domain automation is essentially a synchronized configuration across different domains for faster time to deployment and provision/remove network resources on-demand

Cross-domain automation with NSO

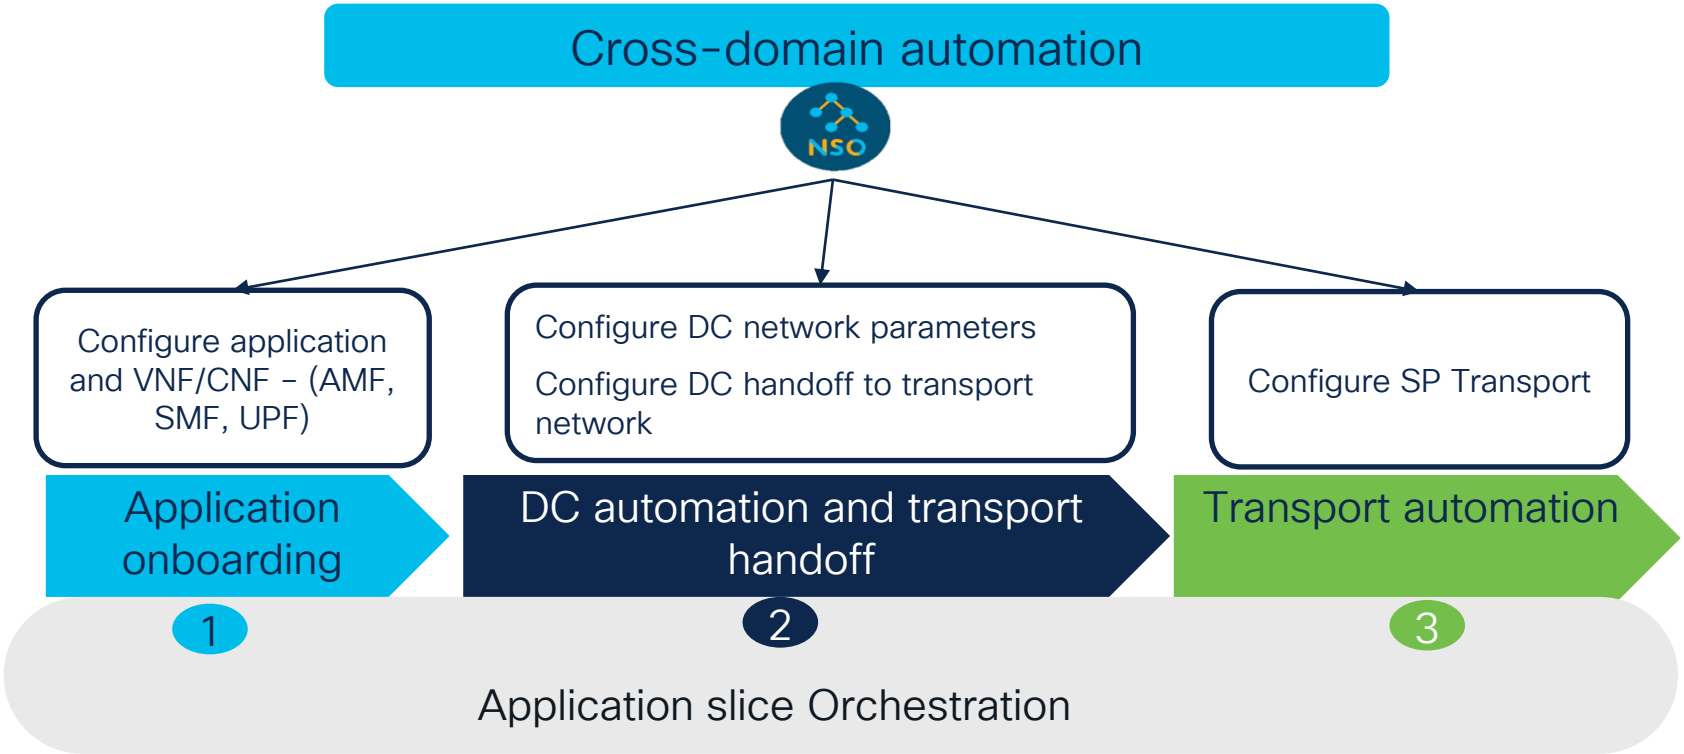

Cross-domain (DC and transport) automation using NSO

- Multi-Domain automation across Transport and DC
- Telco DC, and DC to transport handoff provisioning for both IP and SR handoff

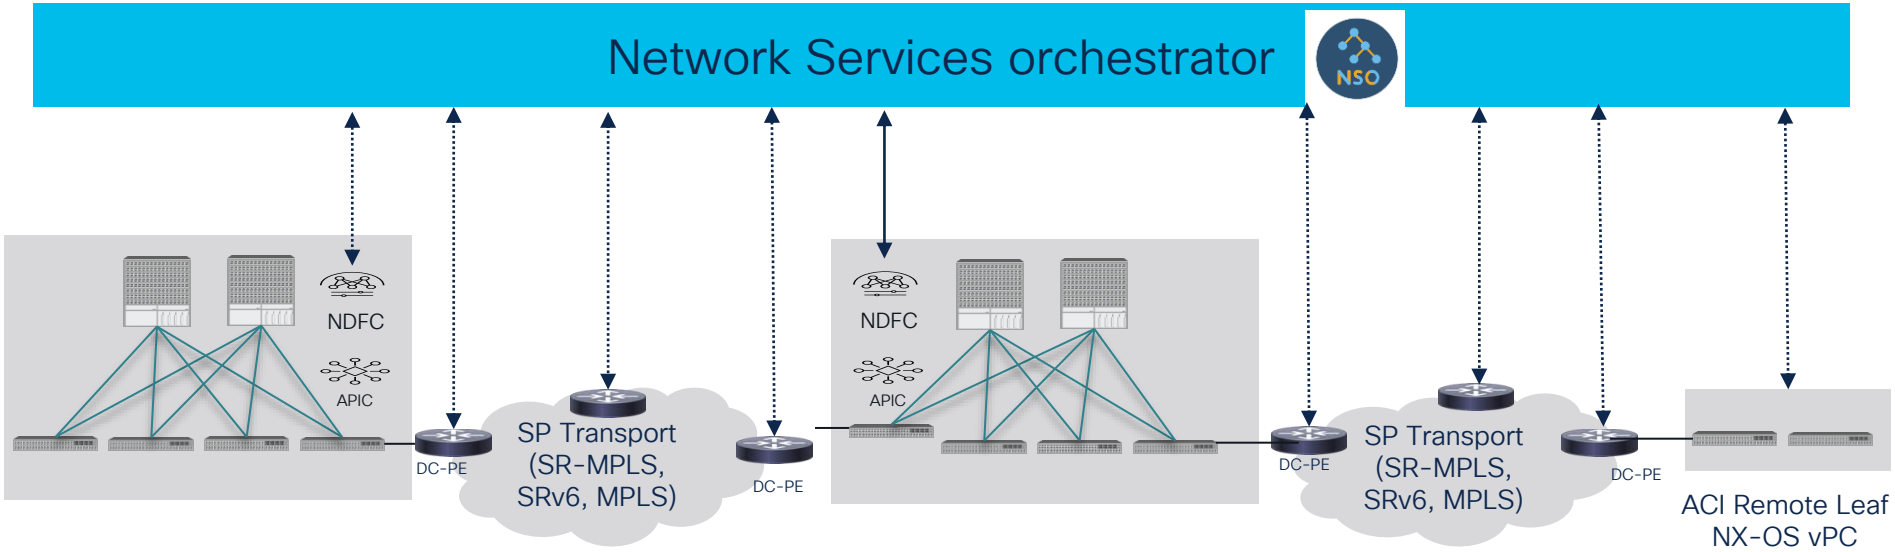

Architecture for cross-domain automation

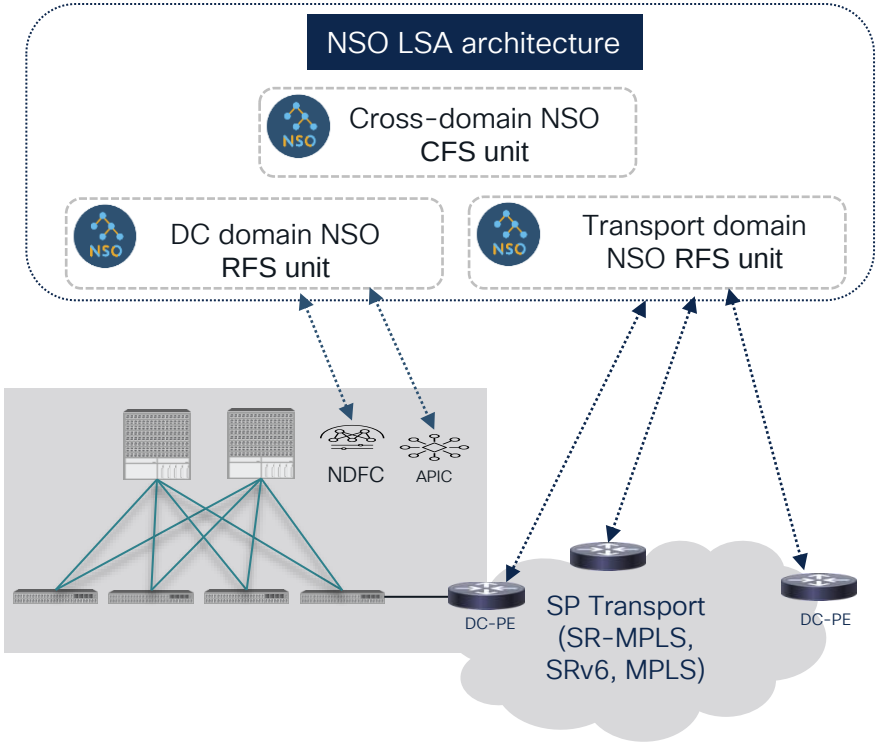

- Utilize NSO LSA (Layered service architecture) for cross-domain automation for scalable deployment by separating different NSO for each domain
  - NSO CFS (consumer facing service) unit is used to deploy cross-domain services across DC and transport
  - NSO transport RFS (resource facing service) for automating transport network – SR, L3VPN etc.
  - NSO DC RFS for automating DC networks by utilizing domain level controllers – APIC and NDFC

Agenda

- Telco DC trends and challenges
- Solution overview for automation and operations of distributed DCs
- Building distributed Telco DC
- Cloud native packet core integration
- DC to transport handoff
- Cross-domain automation
- Day2ops and integration with OSS/BSS systems
- Customer references
- Conclusion

Application visibility and troubleshooting requirements

- Network functions are getting virtualized/containerized, hence static mapping of application endpoints to network nodes doesn't work any more.
- Network teams want application (NF) visibility into network operations for faster troubleshooting or proactive alerting.

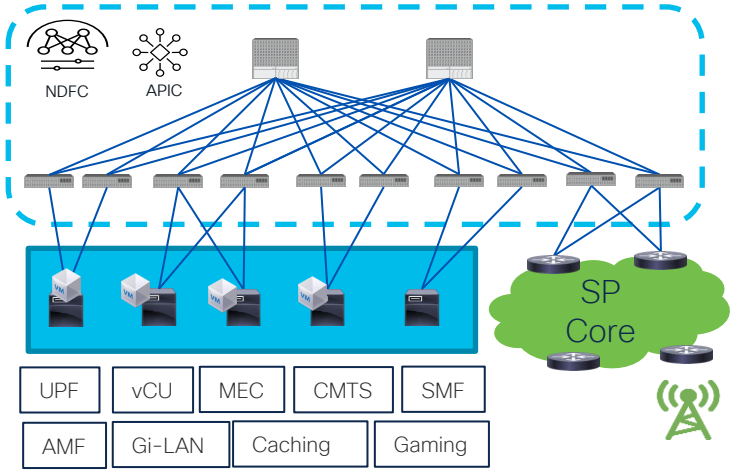

Network operator

- UPF is connected to which leaf and ports in the fabric?
- What is causing high latency and traffic drop?
- What are the IP addresses used by 5G NFs?
- Are application endpoints flapping?

Nexus Dashboard Insights

Proactive Operations : Increase Availability, Performance, And Simplify Operations

Sources Of Telemetry Data	Ingest And Process	Derive Insights	Suggest Action

Config file
Syslog
Tech-support
RIB
FIB
Accounting logs
Debug logs
Streaming telemetry
Environmental
Event history
Cores
Consistency checkers

Metadata extraction

Correlate against dBase

Complex correlation

Proactive Action

Nexus Dashboard Insights – Proactive operations

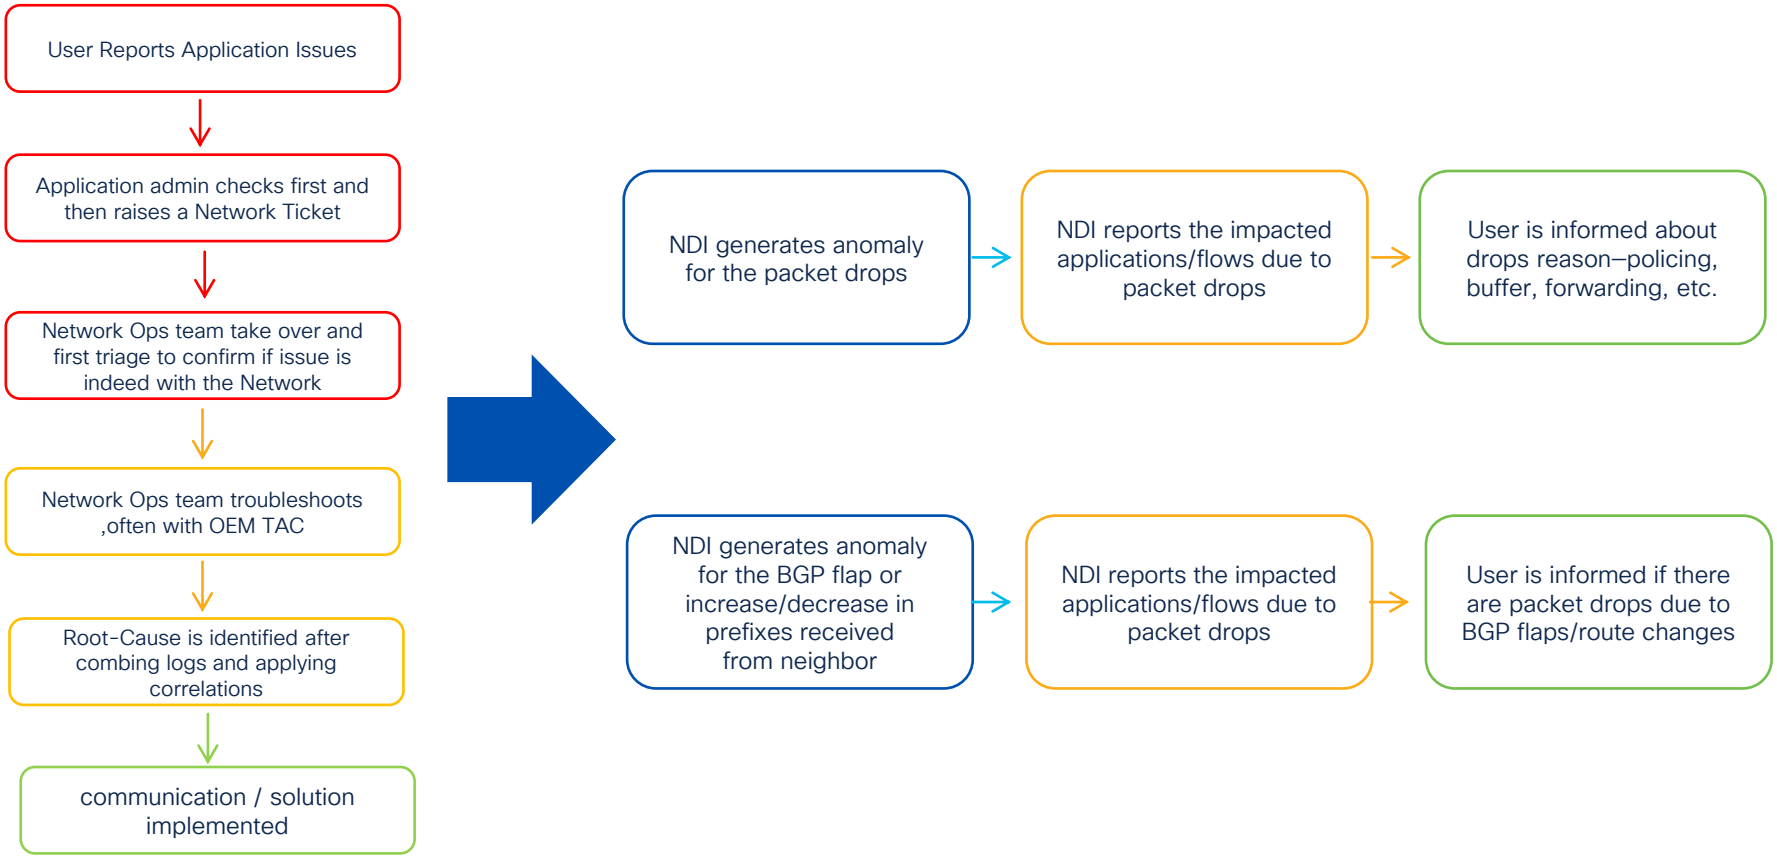

Seamless integration with existing OSS/BSS tools

- Integrate ND Insights with 3rd party OSS/BSS systems using Kafka/APIs
- Customers can use mix of NX-OS and ACI environment. For example - NX-OS at Far Edge can be used for ORAN, and ACI at central/regional DCs can be used for packet core

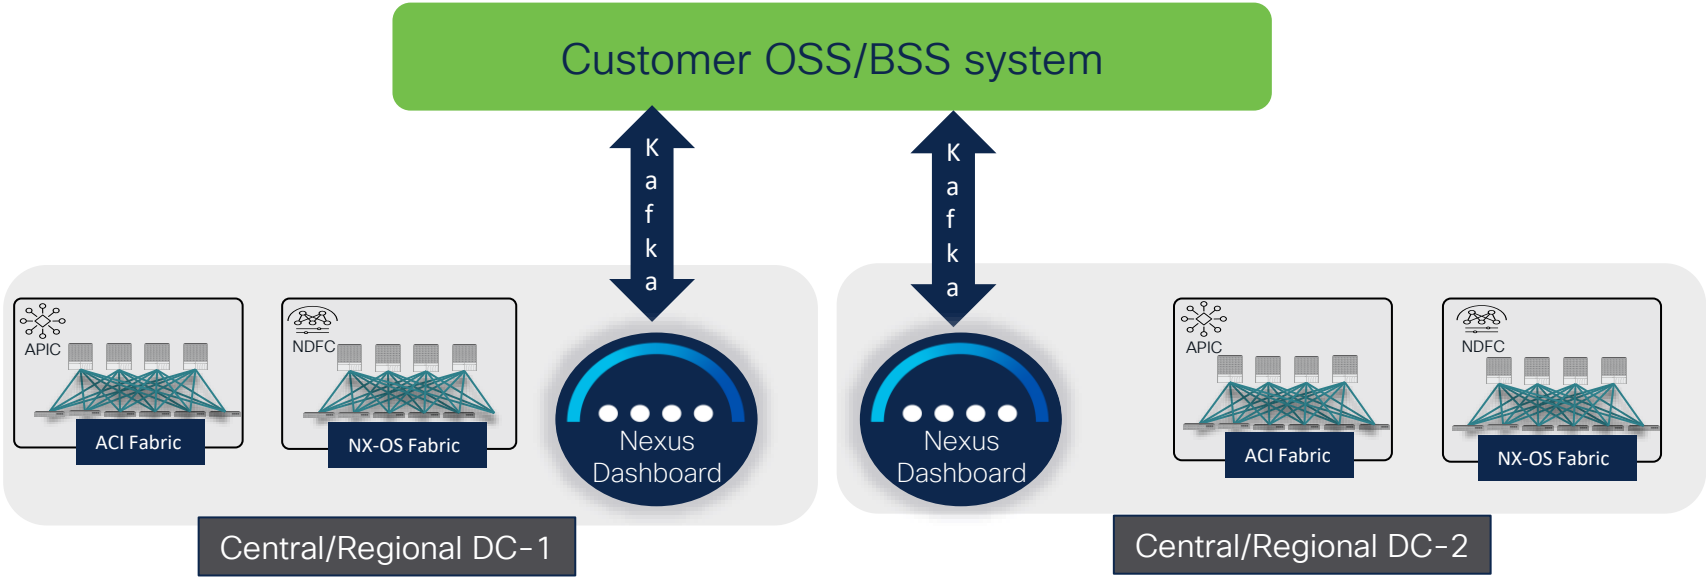

Agenda

- Telco DC trends and challenges
- Solution overview for automation and operations of distributed DCs
- Building distributed Telco DC
- Cloud native packet core integration
- DC to transport handoff
- Cross-domain automation
- Day2ops and integration with OSS/BSS systems
- Customer references
- Conclusion

Large SP deployments for 5G and ORAN

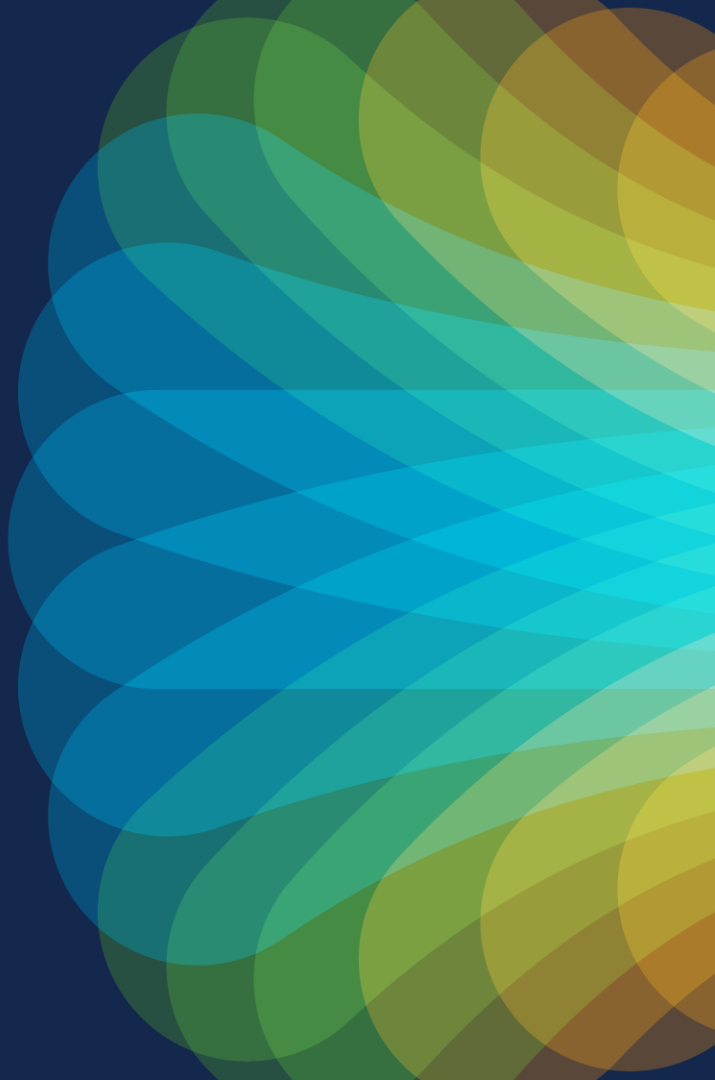

Customer profile and use-cases

Greenfield deployment for ORAN and 5G

Automation first approach for end-to-end orchestration across DC, transport, application and NFs

Deployed Cisco NSO for end-to-end automation, Cisco APIC/DCNM for DC automation, ESC for VNF management, and Cisco VIM (openstack platform) for hosting VNFs

Zero-Touch deployment of 12 Central DC, 58 Regional DC and 3500+ far edge DC

Distributed Telco DC deployment

Far edge DCs

- About 9000 switches deployed across 3500+ far edge locations
- 16 DCNM instances for all 9000 far edge switches
- Evaluating the migration from DCNM to NDFC

Central DCs

- 12 central DC & 58 regional DC with ACI fabric to provide end-to-end automation, service chaining and integration with CVIM (hosting platform for packet core and other mobility services)

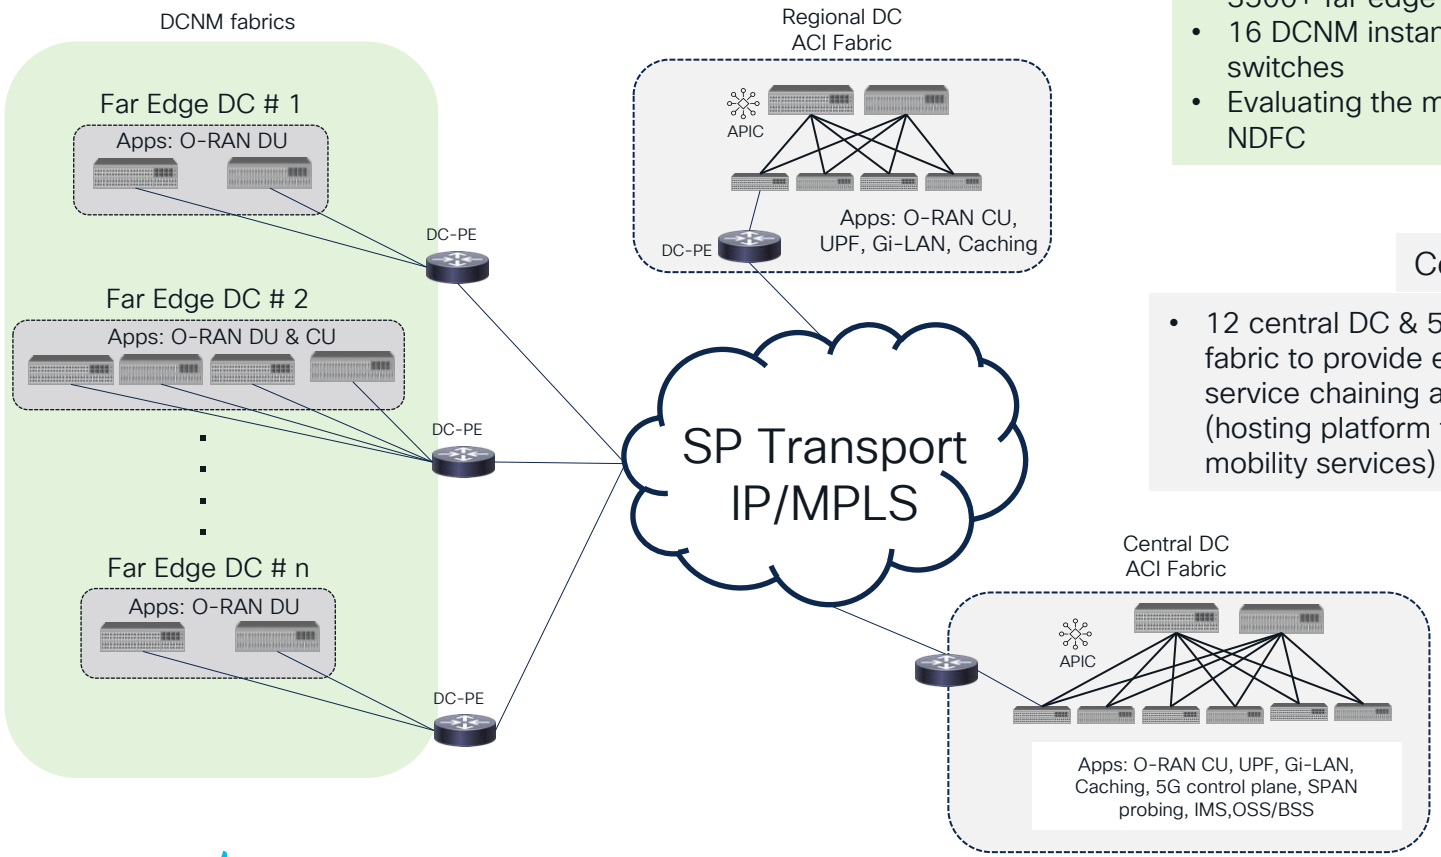

Large APJ SP customer

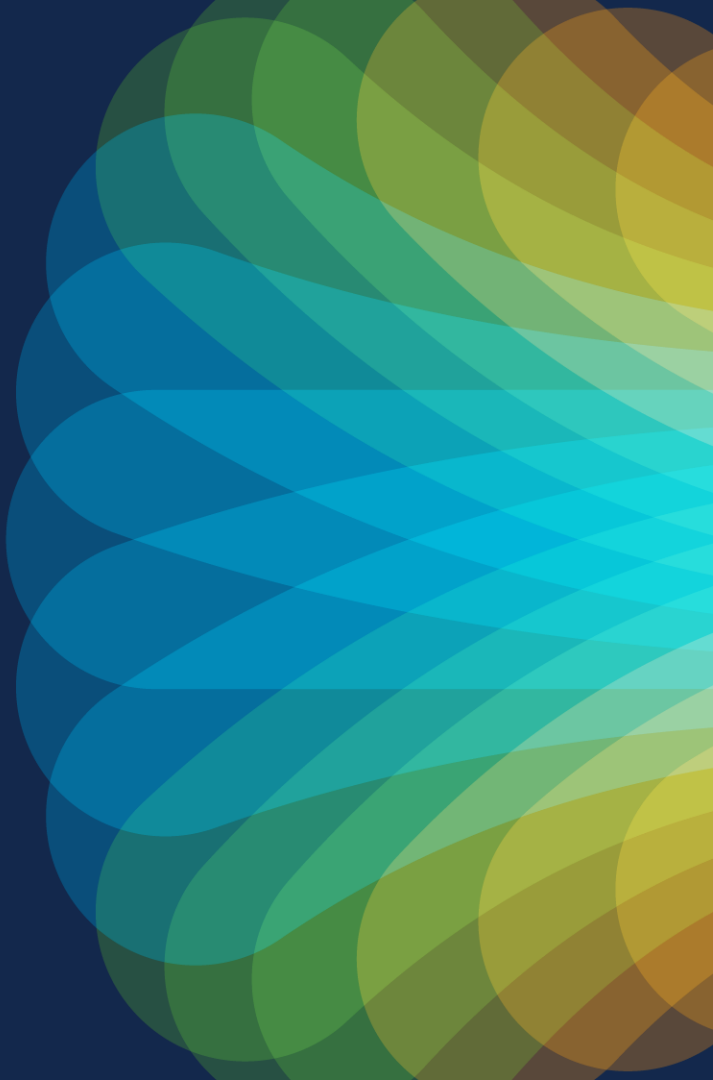

Journey of a large customer deployment in APJ

Nexus 9K/7K Deployment

- Physical Packet core deployment
- Central/Regional deployment
- Simple Gi-LAN services such as FW, LB
- **30+ Fabrics with DCNM**

Expansion of N9K/ACI Fabrics

- 25+ additional fabric for vEPC
- Separate fabrics for wireline and wireless services
- Additional Gi-LAN services such as CG-NAT
- Multi-Pod Deployment

Future

- 1300+ Edge DCs and 500+ Regional/Central DCs
- NDO & Day2ops deployment
- SRv6 handoff

N9K/ACI Fabric deployment

- Physical Packet core deployment
- Central/ Regional Deplyment
- **40+ ACI Fabrics deployed in less than 3 months**
- Additional Gi-LAN services such as TCP optimizer deployed with fabric as load-balancer

5G & Edge deployment

- 5G containerized packet core deployment with 400G fabric
- IOT applications(smart meters, car parking, industrial IOT etc.) are being deployed at new and existing fabrics
- SR MPLS handoff
- 160+ ACI fabric deployment in production

5G deployment overview

Deploying 5G for Enhanced Mobile Broadband (eMBB) service and Internet of Things (IoT) use cases

Distribution of network functionality (UPF, Gi-LAN) to edge locations

Separate fabrics for control plane and data plane services

400G fabrics for the data plane services

5G workload placement

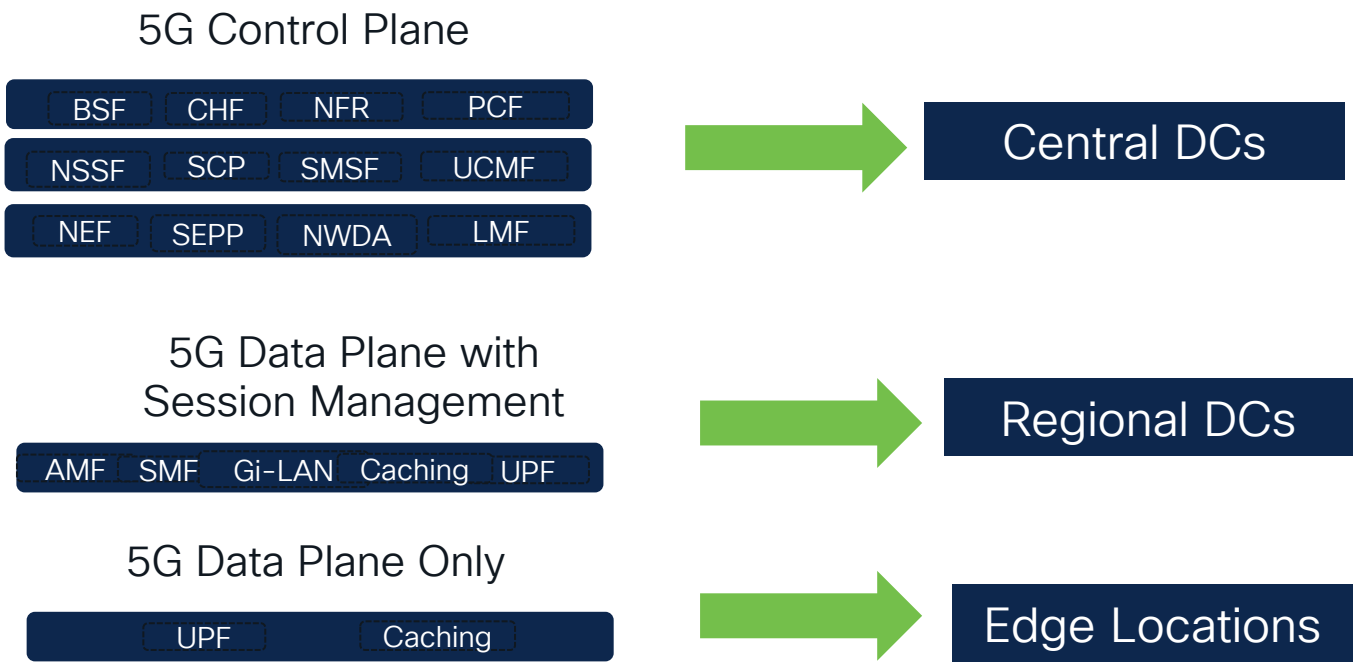

Small DC Fabric deployment at Edge DCs

UPF, Caching and Gi-LAN for low latency and high bandwidth use-case

- Small fabric deployments at the edge
- UPF connected with 100G links
- Gi-LAN connected with 10/25G links
- Routers connected with 100G links

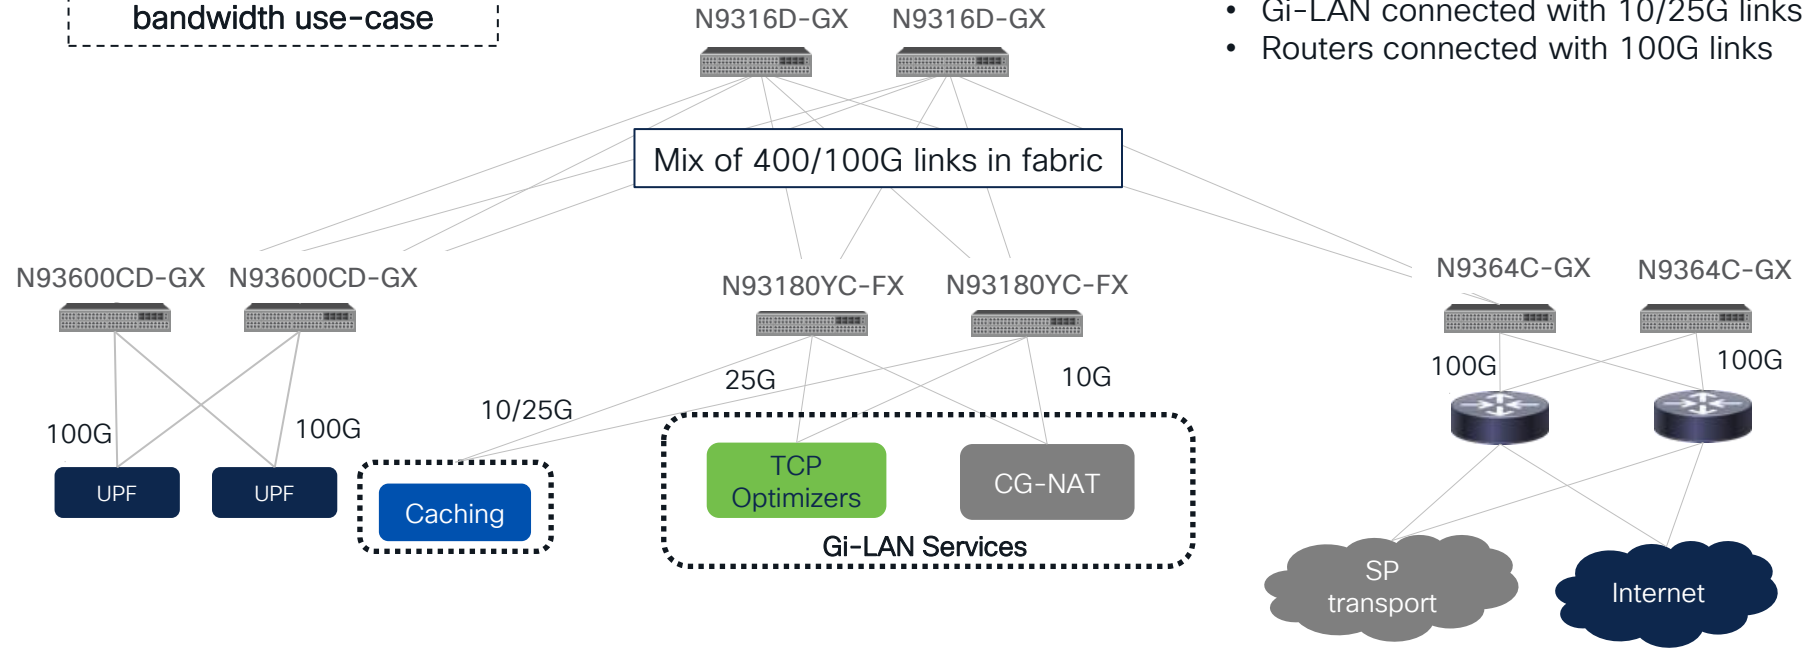

DC Fabric deployment at Regional DCs

AMF, SMF deployed in regional DCs for session management along with UPF, Caching and Gi-LAN

- UPF connected with 100G links
- Caching, SMF, AMF, Gi-LAN connected with 10/25G links
- Routers connected with 100G links

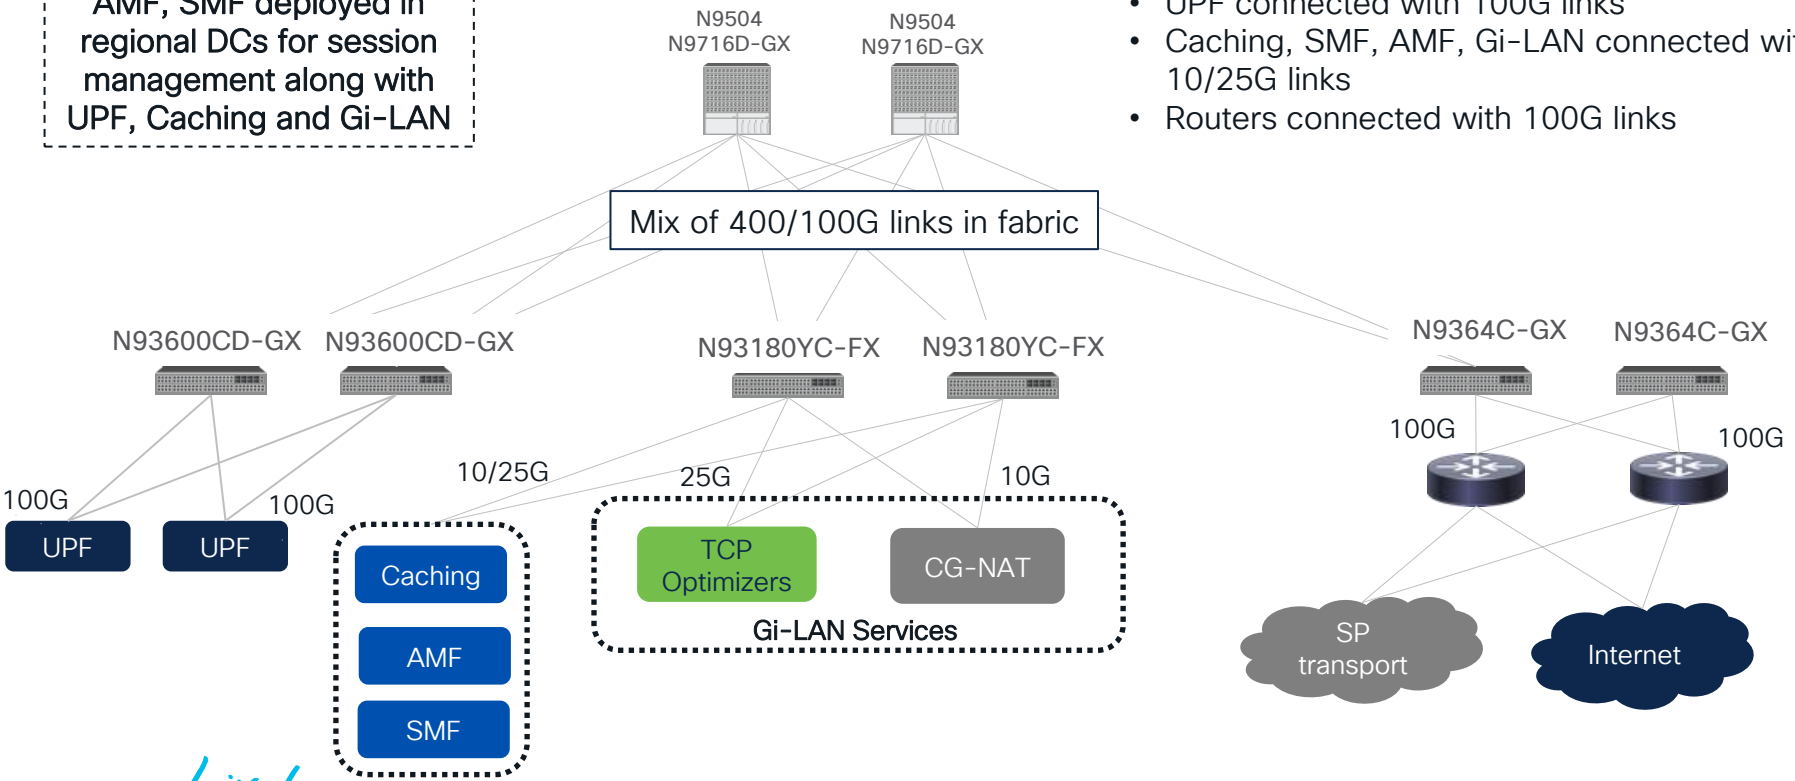

DC Fabrics deployment at Central DCs

Separate fabric for control and data plane for isolation

Outage in control plane won't impact data plane and vice versa

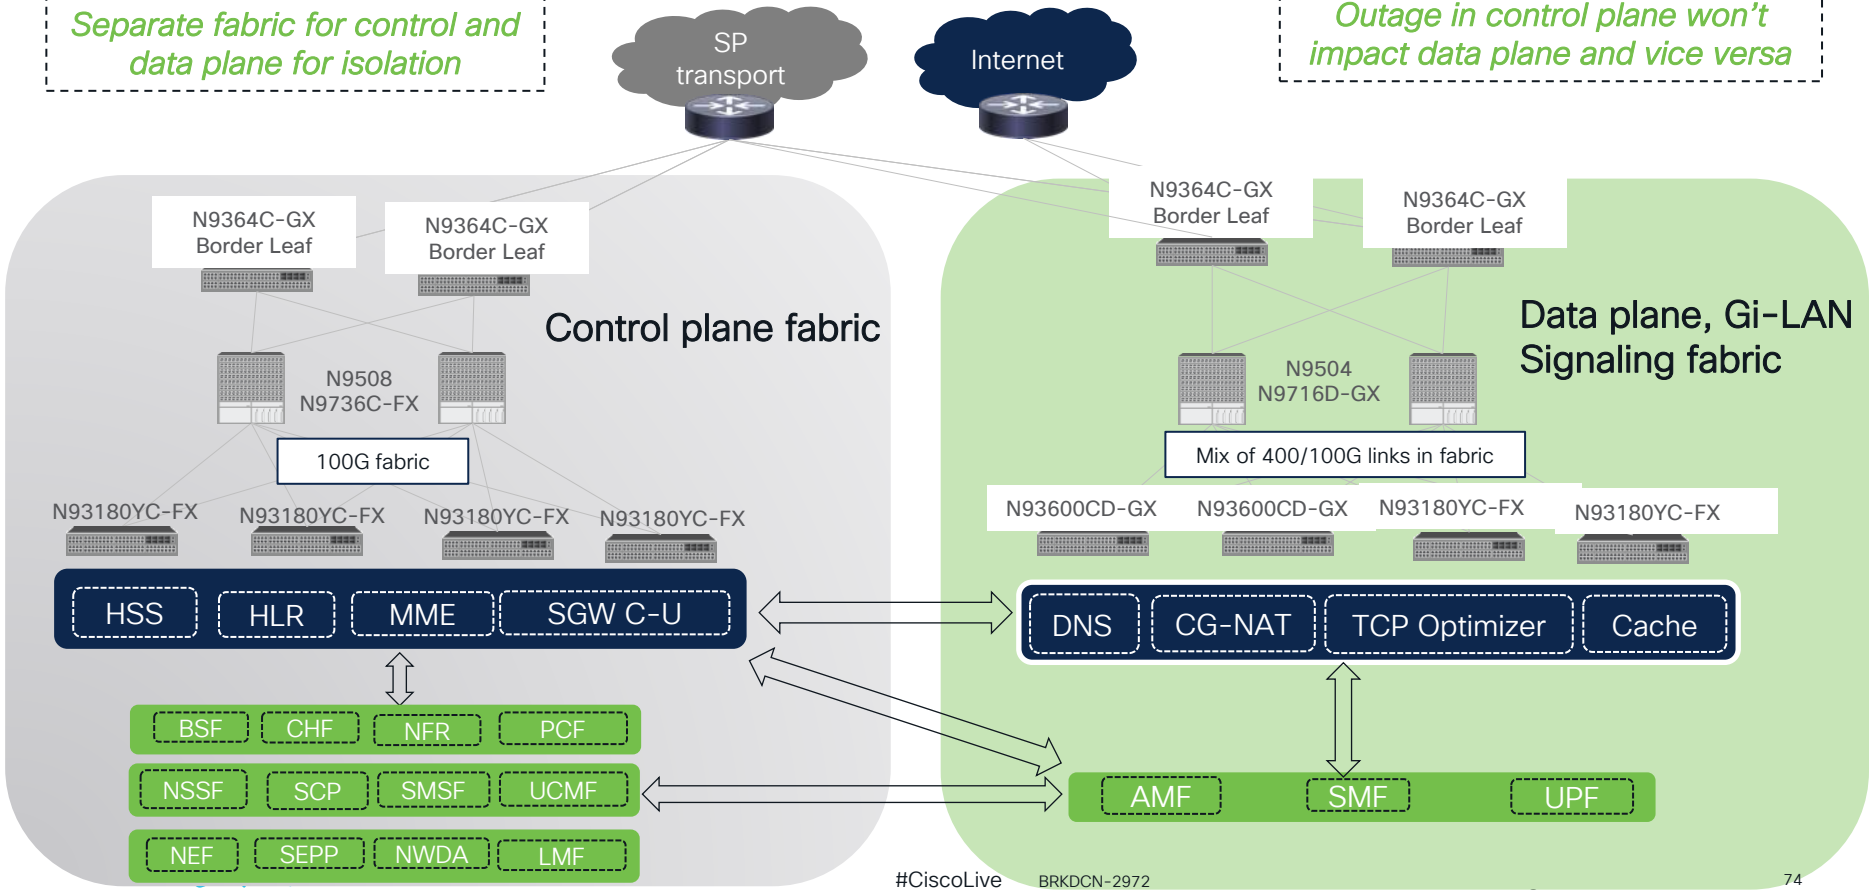

Agenda

- Telco DC trends and challenges
- Solution overview for automation and operations of distributed DCs
- Building distributed Telco DC
- Cloud native packet core integration
- DC to transport handoff
- Cross-domain automation
- Day2ops and integration with OSS/BSS systems
- Customer references
- Conclusion

Summary

Automation

Automation of infrastructure is critical for deployment of 5G and ORAN. Cisco has built end-to-end automation across distributed DC using NDO, NDFC and ACI.

Cross-domain integration

Seamless integration for data plane – VXLAN-EVPN/ACI to SR-MPLS/SRv6 handoff

Cross-domain automation using NSO

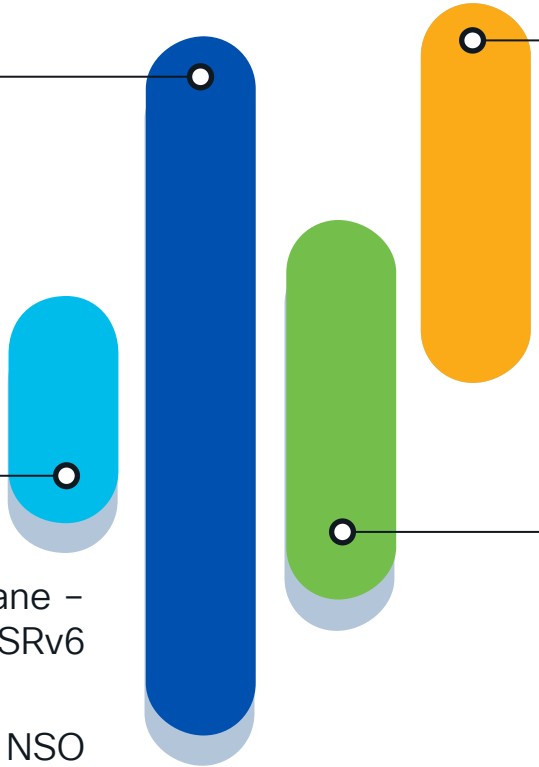

Operation

Proactive operation is must for any edge deployment including 5G and ORAN. Nexus Dashboard insights is a tool for operation team to proactively troubleshoot and fix network problems.

Proven deployment

Cisco has several successful deployment of distributed Telco DC for 4G, 5G and Open RAN across the globe

Fill out your session surveys!

Attendees who fill out a minimum of four session surveys and the overall event survey will get **Cisco Live-branded socks** (while supplies last)!

Attendees will also earn 100 points in the **Cisco Live Challenge** for every survey completed.

These points help you get on the leaderboard and increase your chances of winning daily and grand prizes

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 www.CiscoLive.com/on-demand

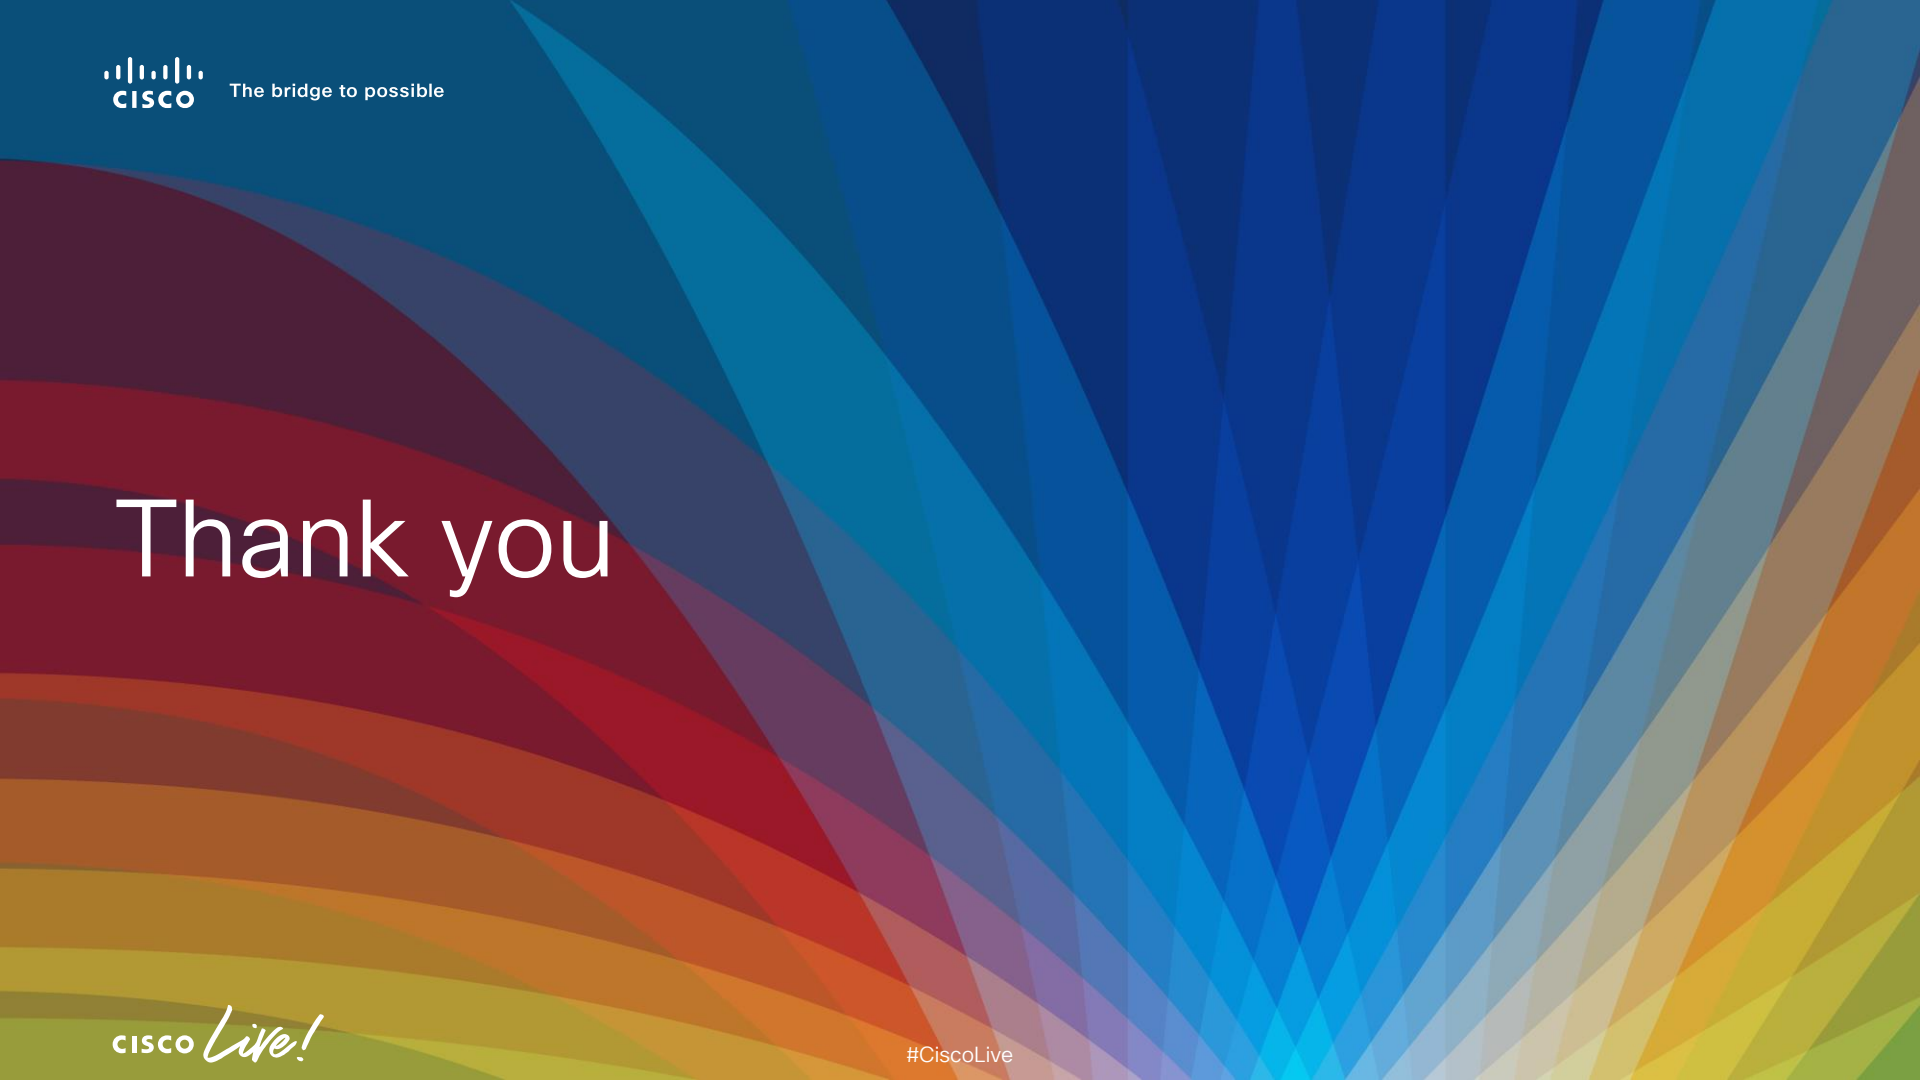
cisco
The bridge to possible
Thank you
cisco Live!
#CiscoLive

Cisco Live
Challenge

Gamify your Cisco Live experience!

Get points for attending this session!

How:

1. Open the Cisco Events App.
2. Click on 'Cisco Live Challenge' in the side menu.
3. Click on View Your Badges at the top.
4. Click the + at the bottom of the screen and scan the QR code:

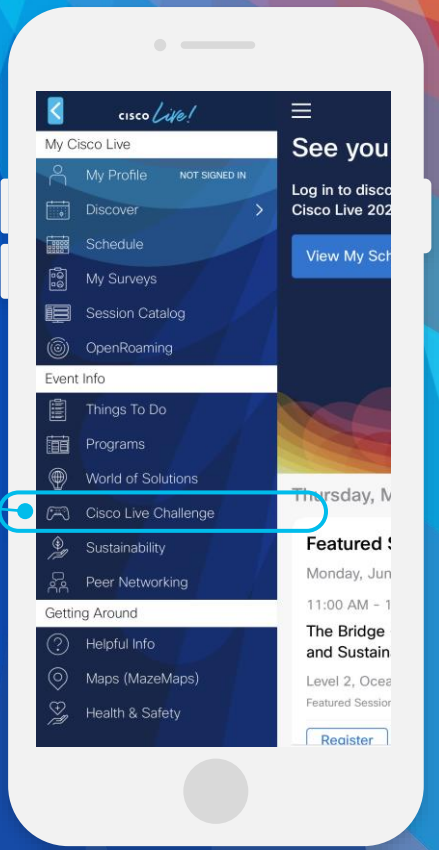

cisco Live!


cisco Live!
Let's go
#CiscoLive