You make **possible**

# Case study of SP customers running ACI based SDN for telecom datacenter

Abhishek Mande – Principle SE, SP
Sonu Khandelwal – TME, IBNG

BRKACI-3620

# Cisco Webex Teams

## Questions?

Use Cisco Webex Teams 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 Webex Teams or go directly to the team space
4. Enter messages/questions in the team space

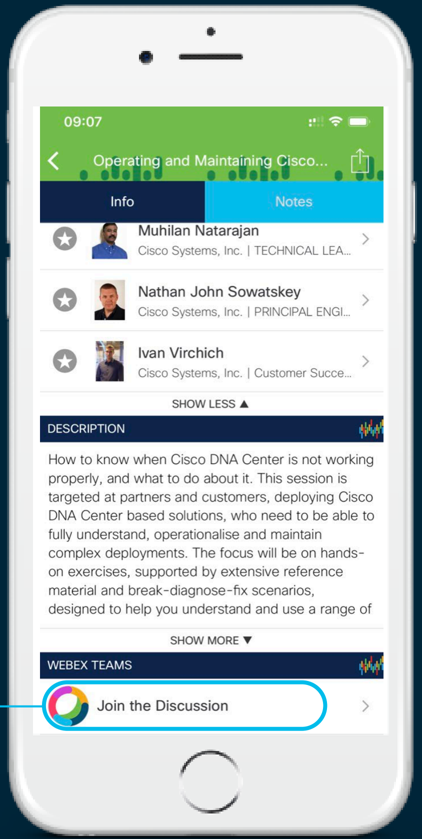

# Agenda

- **Evolution of Telco Datacenter**
- Customer Profiles & Technical Requirements
- How ACI met the requirements
- Looking Ahead
- Conclusion

  

# Session Objective

- This is a ”Case Study” session, highlighting how large Telecom Service Providers used ACI to meet their requirements.

Initial assumption:

The audience already has knowledge of ACI concepts (Tenant, VRF, BD, EPG, L3Out, etc.)

  

# Agenda

- Evolution of Telco Datacenter
- Customer Profiles & Technical Requirements
- How ACI met the requirements
- Looking Ahead
- Conclusion

# Evolution of Telco Datacenter

Connectivity Service: EPC CUPS

Edge Computing: CDN, IOT, Localized AR, Analytics

Connectivity Service: vRAN

Connectivity Service: vRAN (Packet layer), EPC/BNG CUPS, Edge services

Edge Computing: Analytics, IOT, AR/VR/Gaming, Public cloud hosting, CDN...

Connectivity Service: Control plane of CUPS gateways

Traditional IT apps

Connectivity Service: Packet core, Policy

Managed Services: CPE / SD-WAN, FW

IT Apps


  

# Agenda

- Evolution of Telco Datacenter
- **Customer Profiles & Technical Requirements**
  - Applications
  - Datacenter Fabric
- How ACI met the requirements
- Looking Ahead
- Conclusion

# Customer Profiles

- ✓ 15+ Top operators worldwide
- ✓ Largest deployment of 60+  fabrics serving 300mn+ subscribers
- ✓ Largest fabric in telco DC 200+ leafs
- ✓ Fabric for 3G, 4G services (Physical & Virtual)
- ✓ Mix of Cisco and Non-Cisco 3G, 4G services
- ✓ Readiness of 5G in some advanced markets (Leading with NSA deployments)

# Telco Applications Requirement

## IMS

- Faster Convergence
- Consistent low latency
- Multiple Hypervisor
- IPv4 & IPv6 connectivity

## SGW & PGW

- Virtualized deployment across multiple racks for control and data plane
- Failover & Redundancy
- Routing with fabric

## Gi-LAN

- TCP Optimizers
- Deep Packet Inspection (DPI)
- CG-NAT
- URL Filtering
- Service Chaining with multiple nodes

## CDN & OTT Caching

- Caching managed by OTT Provider
- Self created Media Content
- Driving 100G
- Migrating to 400G interfaces

## Mobile Networking Monitoring

- High performance data replication with scale

## PCRF

- IPv4 & IPv6 connectivity
- Multiple hypervisor

## Voice over Wifi

- Faster Convergence
- Consistent low latency
- Multiple hypervisor

## Assurance and insights

- Proactive troubleshooting
- Real time visibility
- Analytics
- Network assurance

# Datacenter Fabric Requirement

- Fabric Automation
- Distributed DC
- Security & segmentation

- North bound API
- Simplified Operation
- Troubleshooting

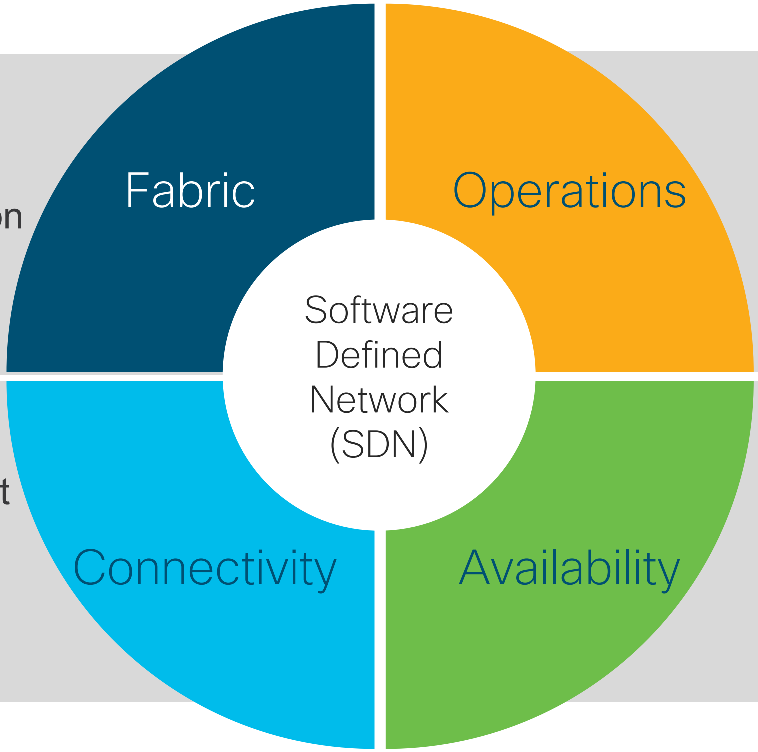


- Policy based deployment
- Service Chaining
- Multi-hypervisor

- Faster Convergence
- Carrier Grade Availability

# Datacenter Fabric Protocol & Scale Requirement

| | |
|---|---|
| Physical Fabric | • 20-200+ Leaf Per Physical Fabric<br>• Multi-Speed Interfaces on Same Leaf & Spine |
| Traffic Throughput | • 1.5-5 Tbps<br>• Linerate |
| Protocol | • BGPv4, BGPv6, Static route<br>• BFD<br>• IPv6 Multicast<br>• ERSPAN with filter, SPAN on drop<br>• Multi-Node PBR, PBR tracking, symmetric load-balancing, resilient hashing |
| Carrier Grade Scale | • 1500 Vlans<br>• 50 VRFs<br>• 1K Mac, ARP & ND per switch<br>• 10-50K end hosts<br>• 10-20K IPv4 & IPv6 LPM addresses<br>• 3K-15K BGP & BFD across fabric |

# Agenda

- Evolution of Telco Datacenter
- Customer Profiles & Technical Requirements
- **How ACI met the requirements**
- Looking Ahead
- Conclusion

  

# Application to Fabric Connectivity

## vEPC deployment with separate Gi-LAN Fabric

- Separated ACI fabric for L4-L7 Services.
- Isolated change domain
- Gi-LAN fabric for wireless and wireline customers

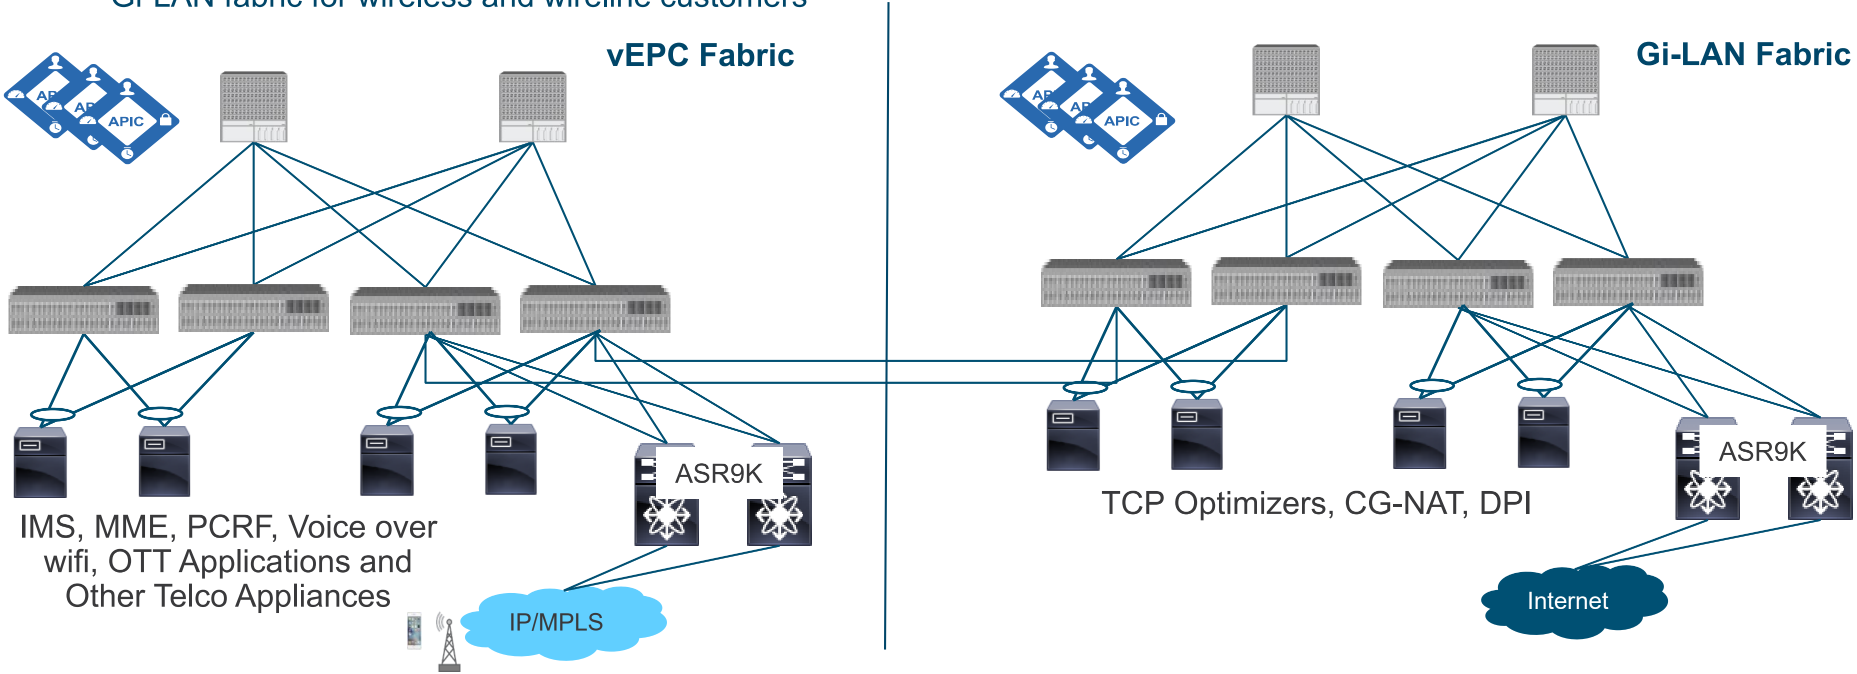


  

# Telco cloud deployment for 4G

Case-Study#2

## vEPC deployment with Cisco Ultra EPC and Cisco VIM

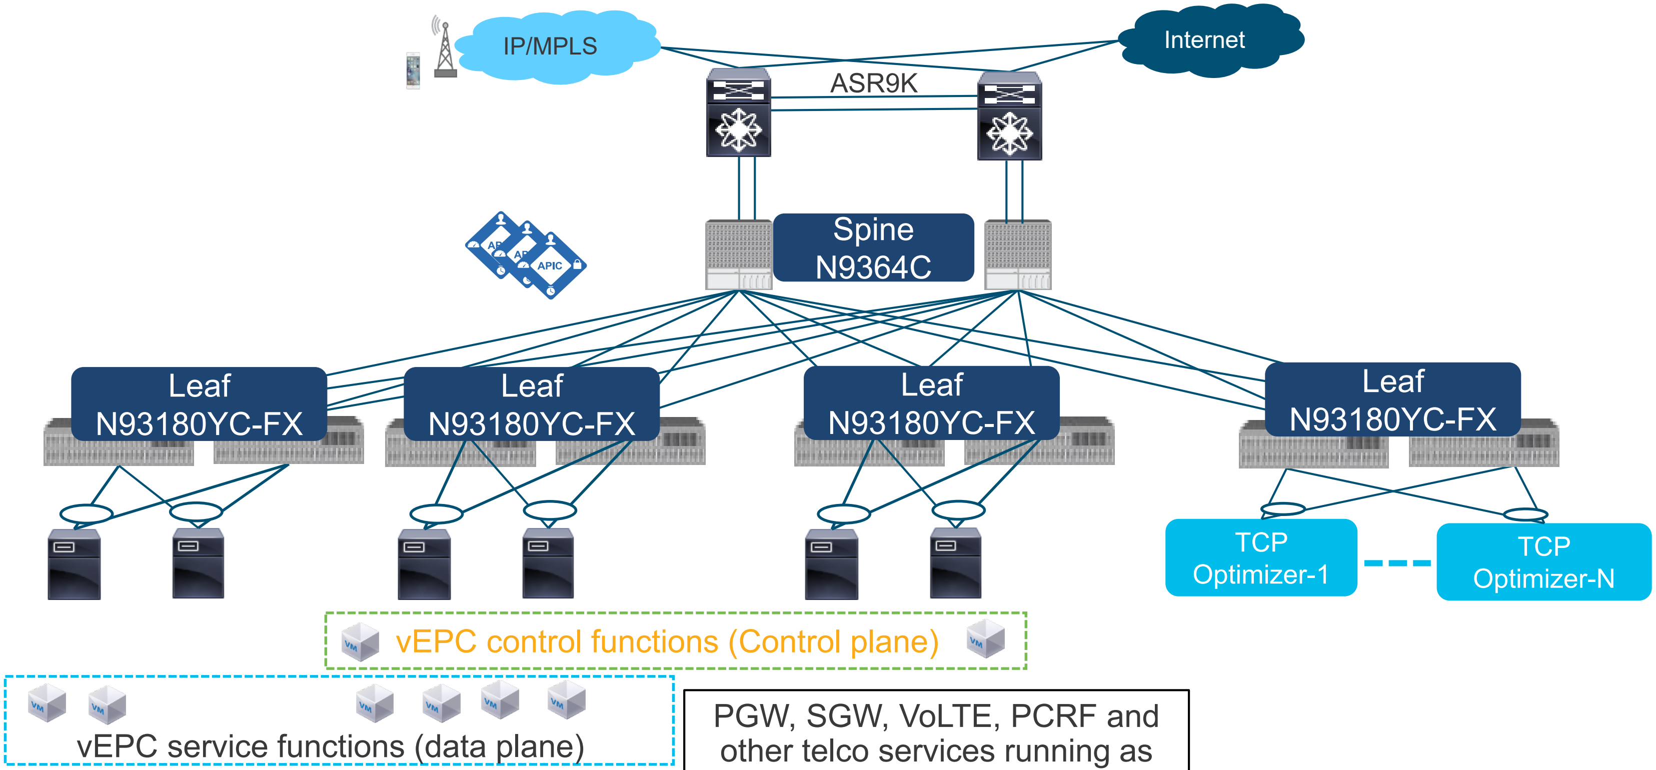


PGW, SGW, VoLTE, PCRF and other telco services running as virtual instances on Cisco VIM

# Distributed Edge DC with ACI Remote Leaf



- ✓ vEPC and Caching is deployed at Edge DC
- ✓ vEPC is connected to Remote Leaf
- ✓ APIC controller at Central site is managing Remote leaf deployed at Edge DC

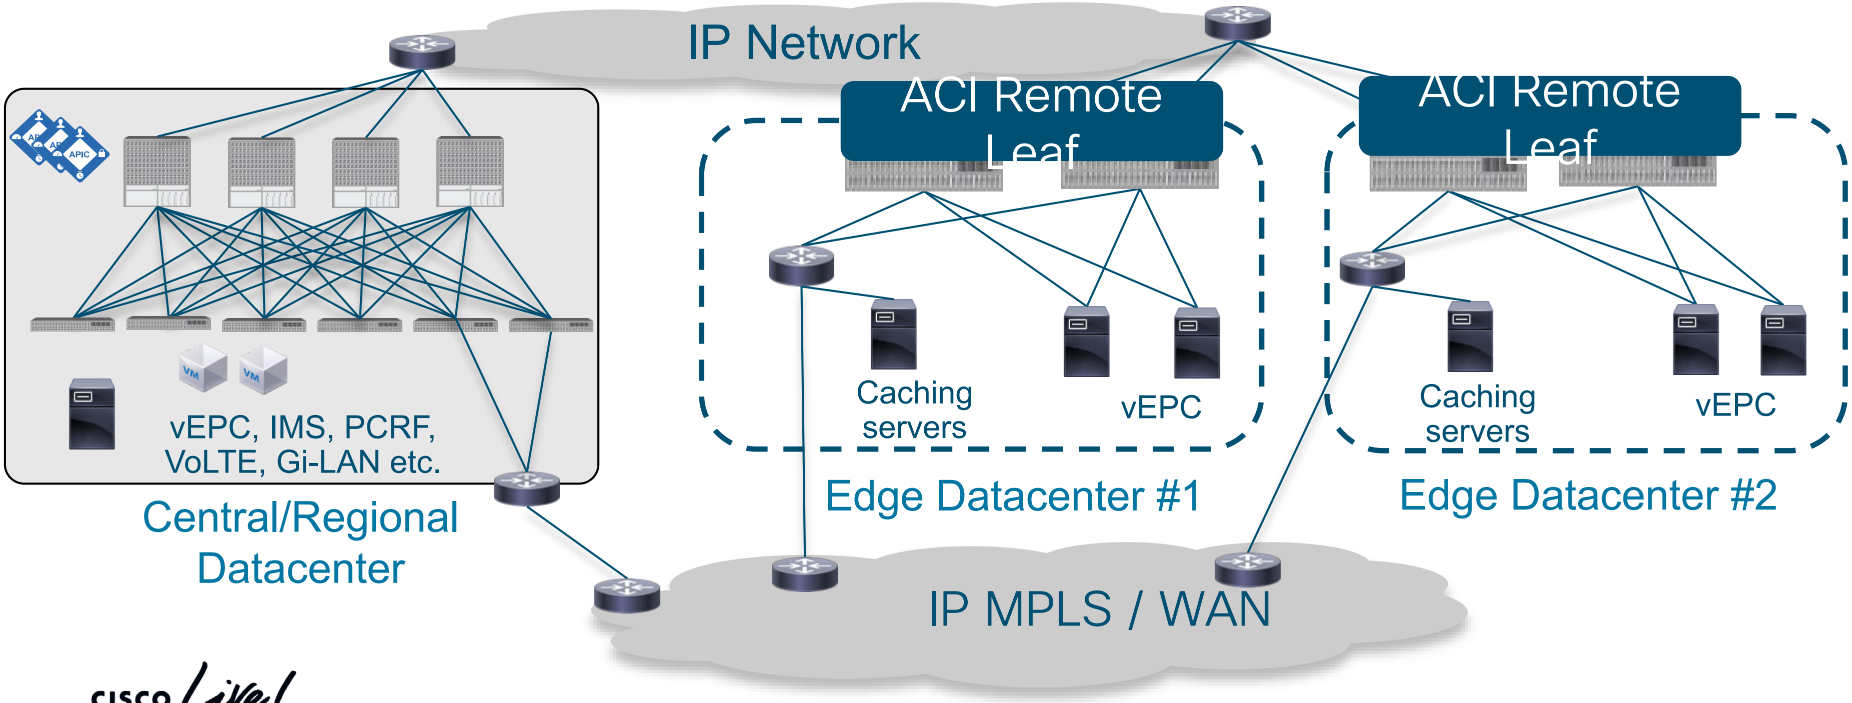


  

# Application to Fabric Connectivity

## 5G ready Central DC architecture

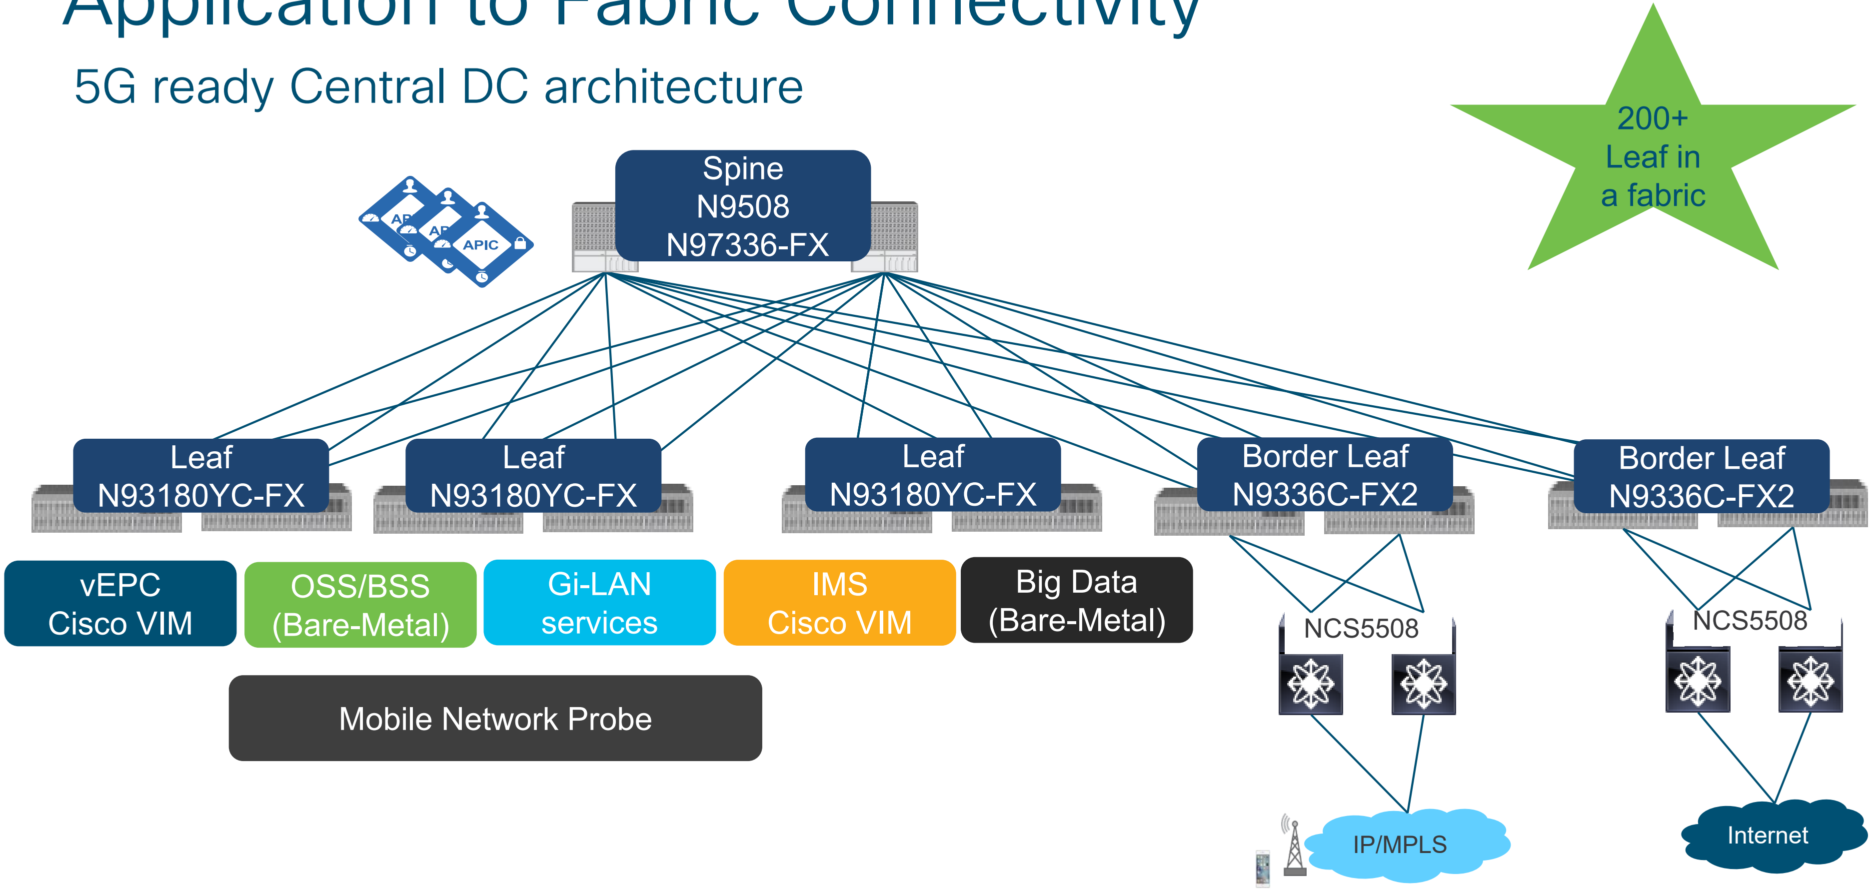


  

# Agenda

- Evolution of Telco Datacenter
- Customer Profiles & Technical Requirements
- **How ACI met the requirements**
  - **Automation**
  - Services Connectivity
  - Operational Simplification
  - Faster Convergence
  - Distributed Edge DC
- Looking Ahead
- Conclusion

# Fully Automated Provisioning of ACI Fabric

**75% reduction in time spent bringing up network**

- Fabric Provisioning
- Validation
- Inventory

Connect switches in Spine Leaf topology

Connect APICs to Leaf Pair

Power on APIC and switches

Input simple details like Fabric Subnet , APIC Out of Band Management IP & Login credential on APIC CIMC

Login to APIC and register switches

Fabric is up and running

# Topology View

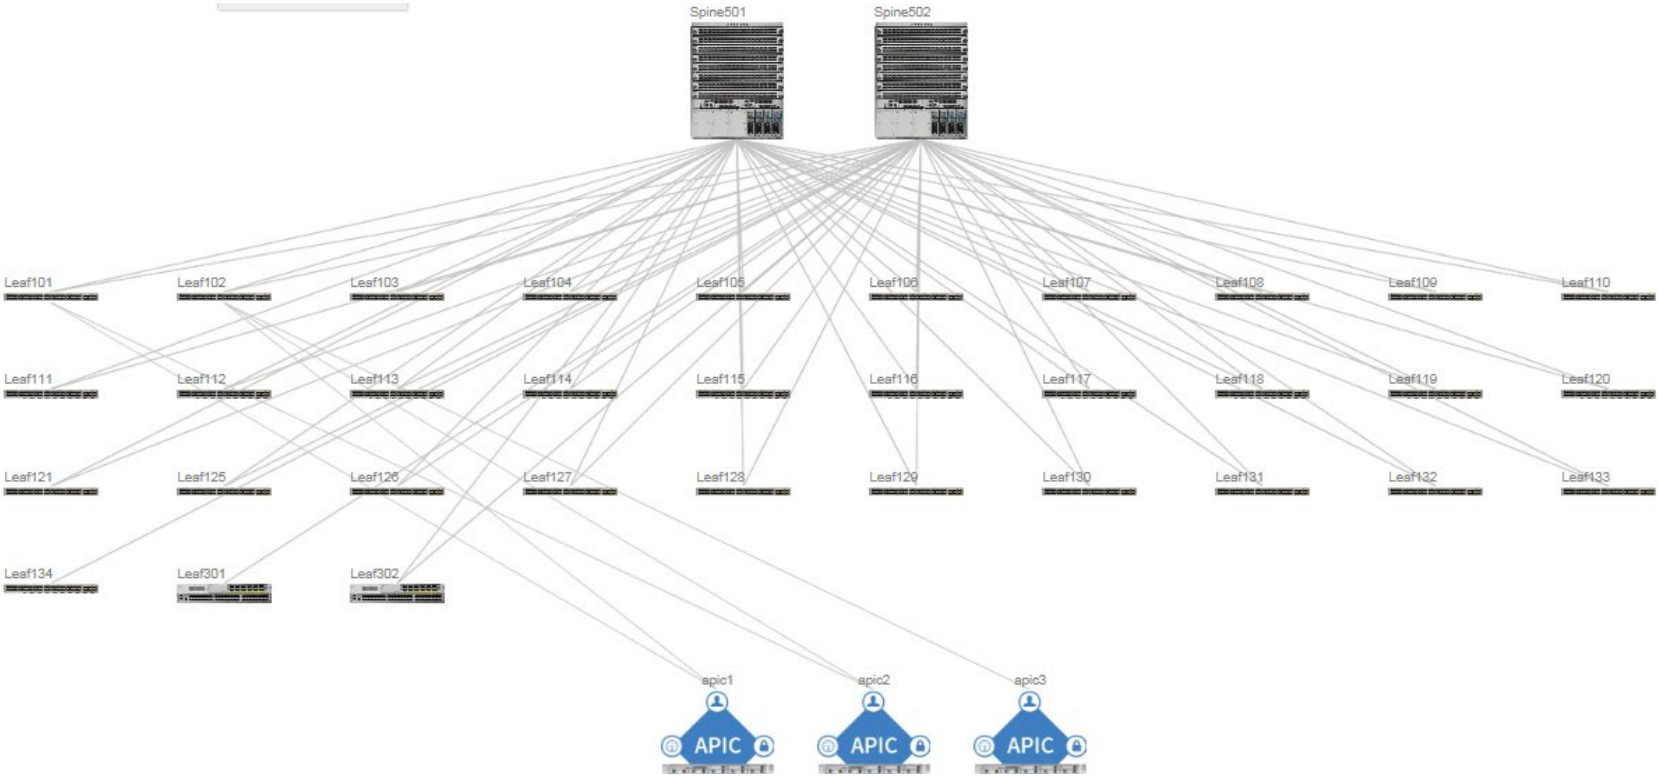


  

# Automatic Policy deployment

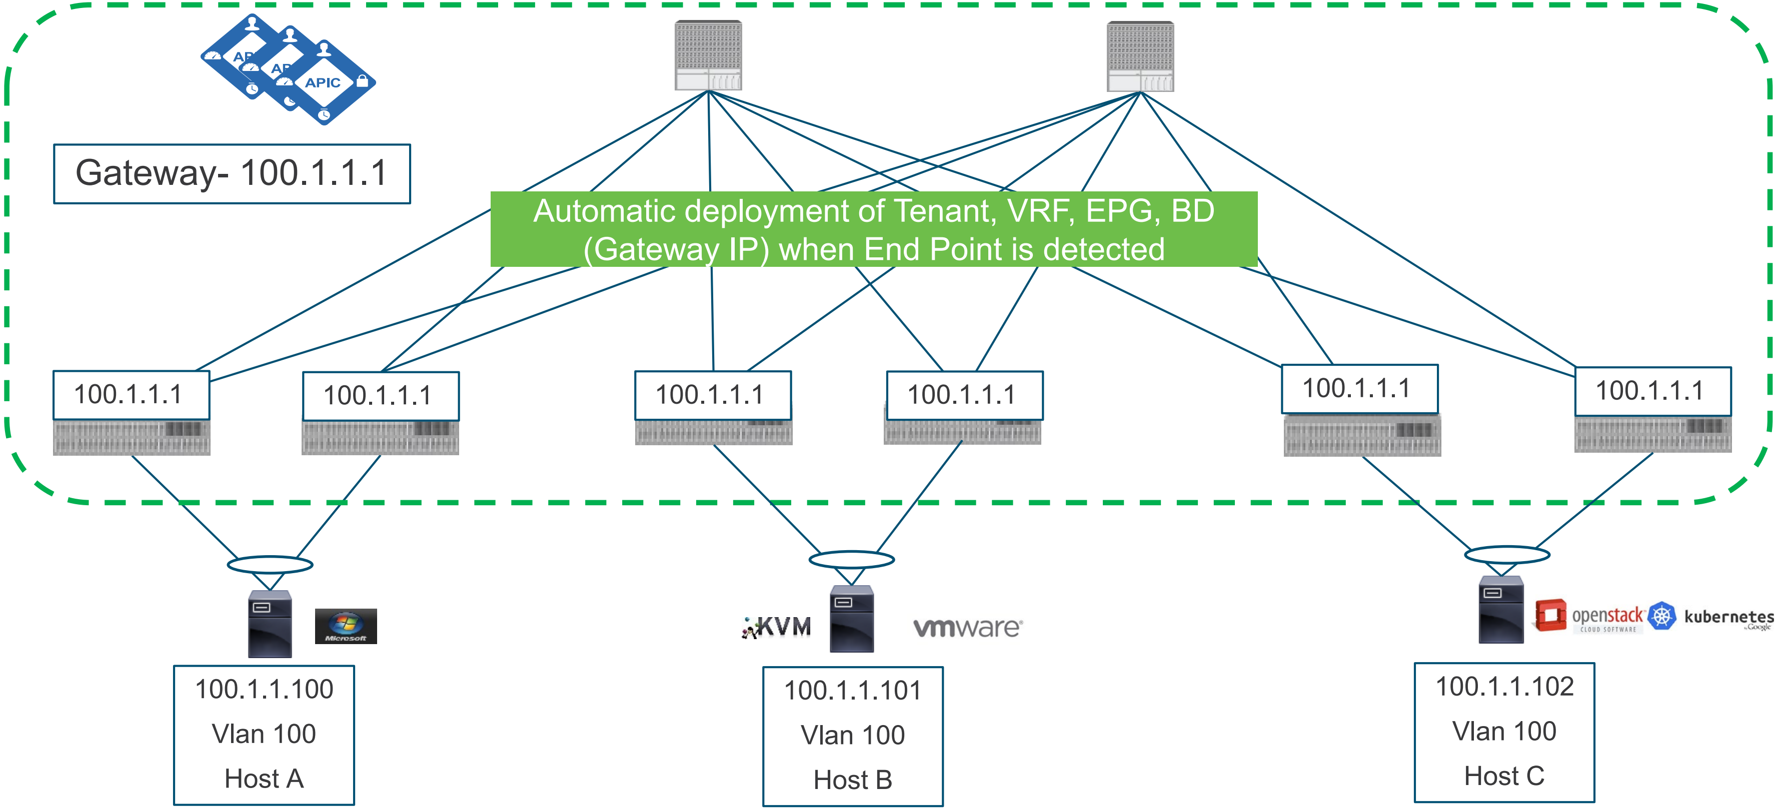


  

# Automation using APIs

Self developed Provisioning tool for pushing ACI policies as an alternative to UI based configuration

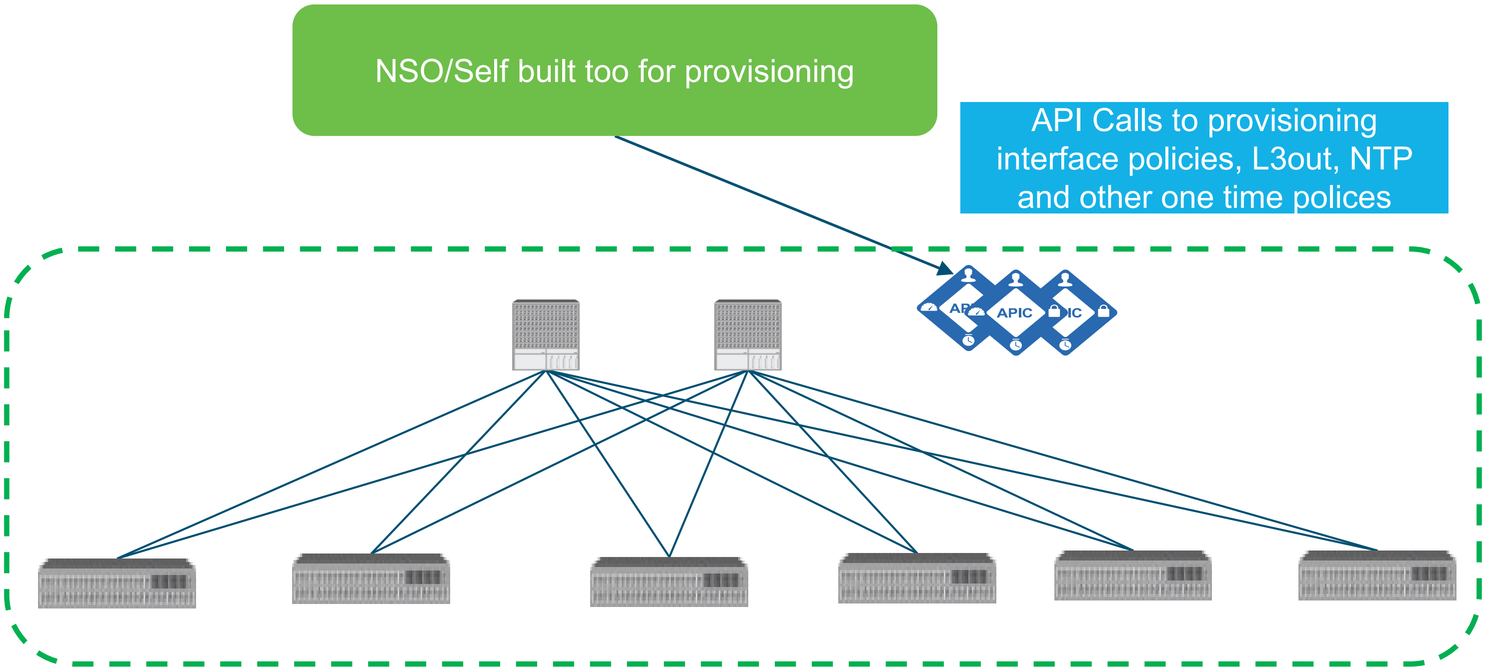

# Agenda

- Evolution of Telco Datacenter
- Customer Profiles & Technical Requirements
- **How ACI met the requirements**
  - Automation
  - **Services Connectivity**
  - Operational Simplification
  - Faster Convergence
  - Distributed Edge DC
- Looking Ahead
- Conclusion

  

# Evolved Packet Core (EPC) deployment

# SAE GW Connectivity to ACI Leaf

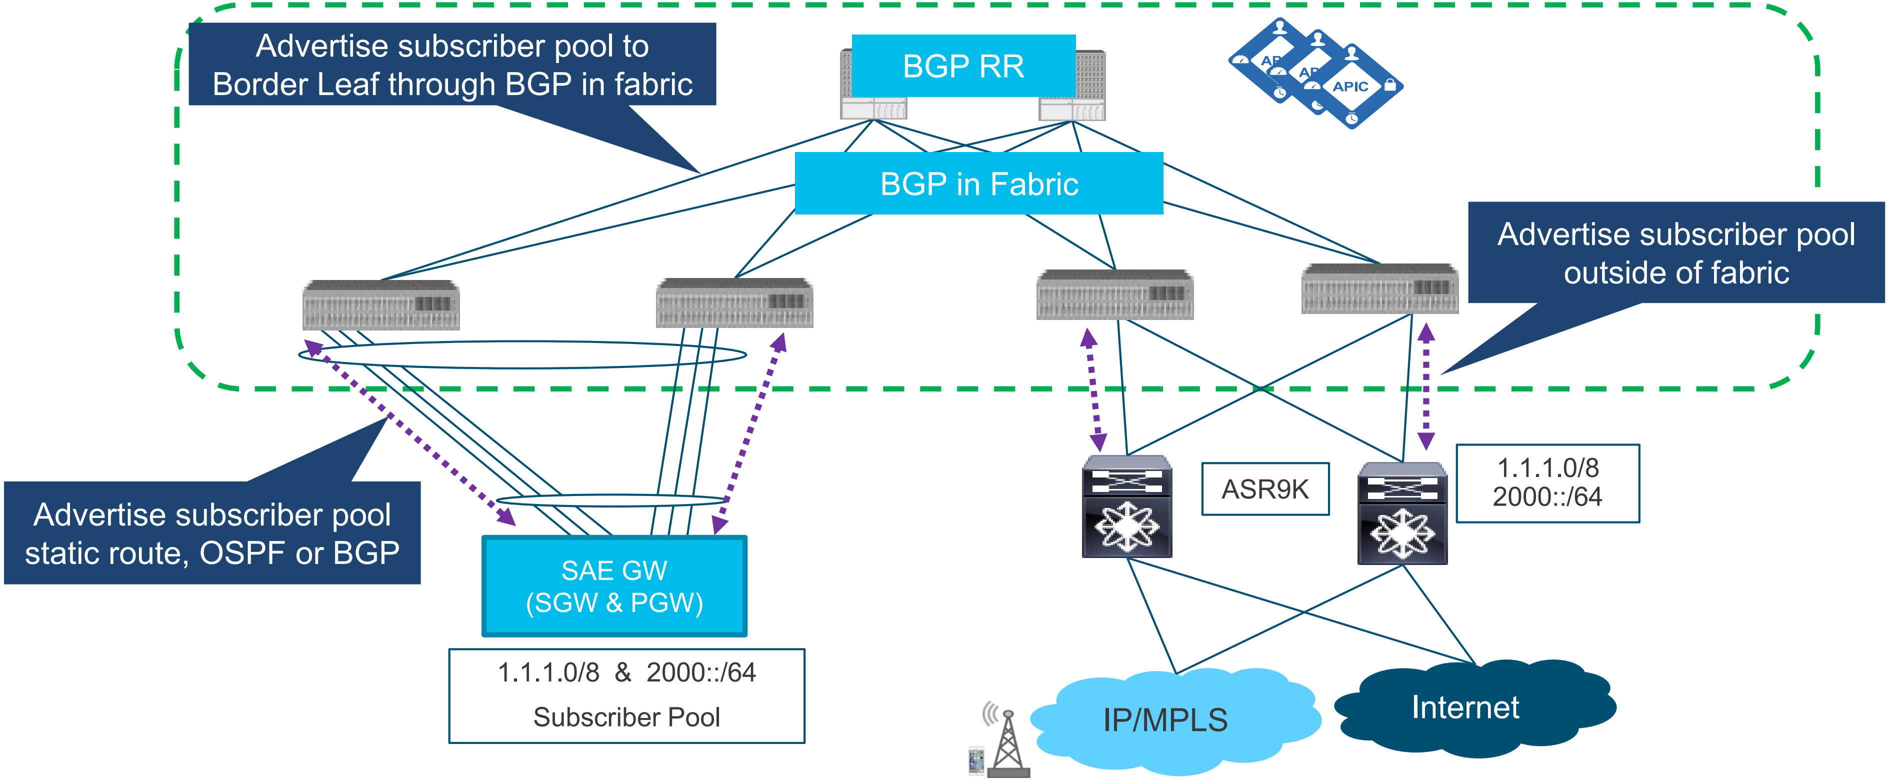

# vEPC characteristics

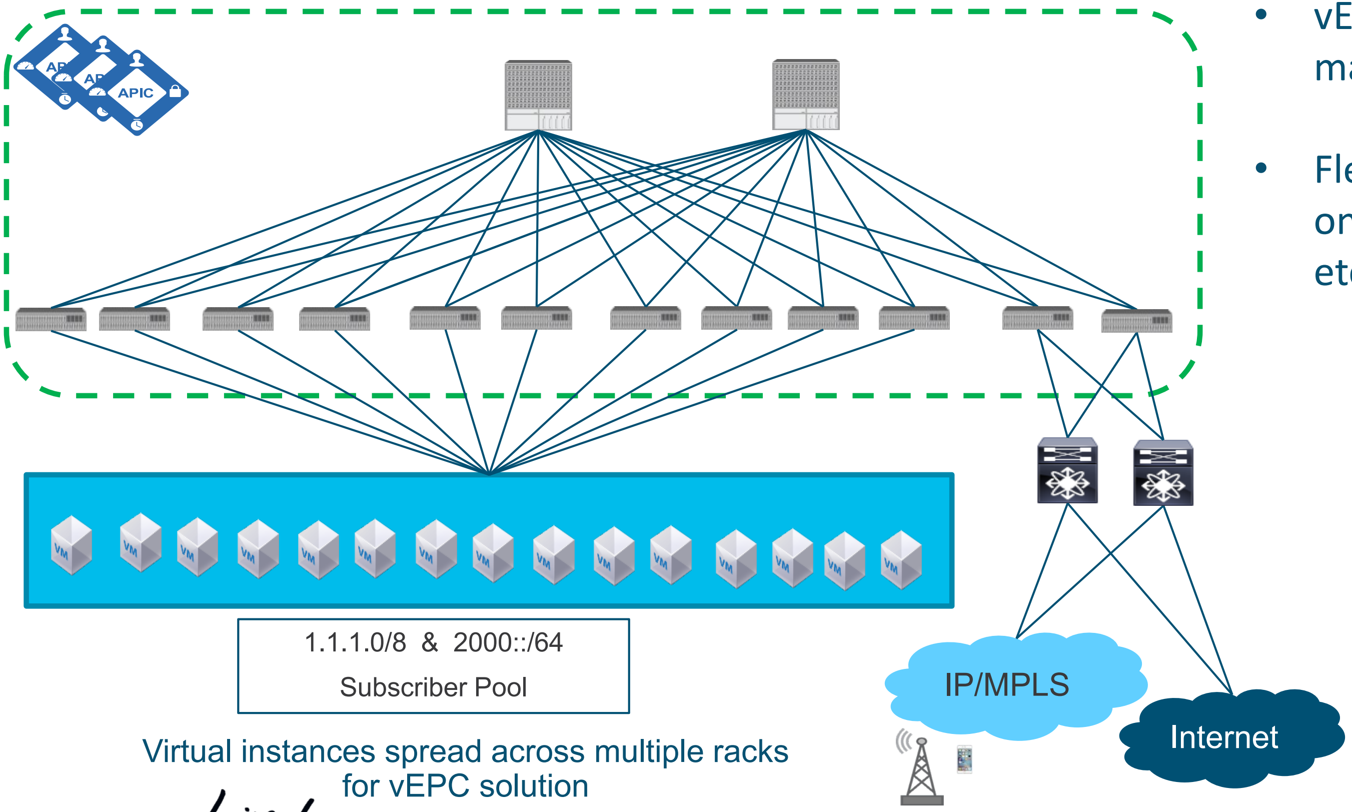


Virtual instances spread across multiple racks for vEPC solution

- vEPC VNF's distributed across many Rack, Leafs
- Flexibility to deploy VNF's based on capacity in Rack's, Servers etc.

# vEPC traffic forwarding requirements

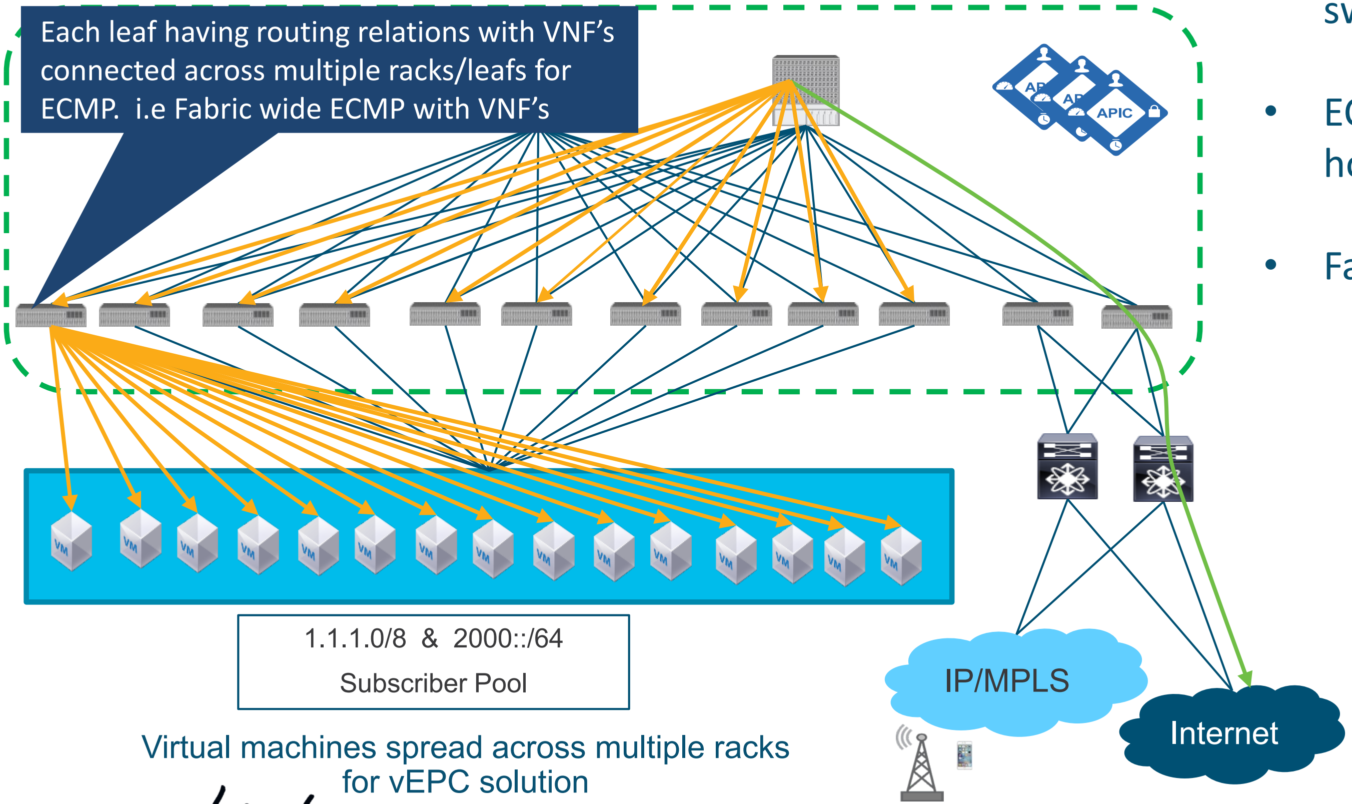


- ECMP within fabric towards leaf switches connected to VNF's.
- ECMP from each leaf to all VNF's hosted across multiple racks.
- Faster convergence using BFD

# Cisco vEPC design with ACI

- vEPC functionality is divided across service functions (SF) and control functions (CF).
- CF advertises subscriber pool information through BGP with the common next-hop of user-plane VNFs (SF)
- All the use-plane VNFs are sharing a common IP
- ACI leaf switches are configured with static route or dynamic routing protocol for provide reachability to common IP of SF
- 64-way ECMP from ACI Leaf to SF across racks.
- BFD is used between SF and ACI Leaf for resiliency

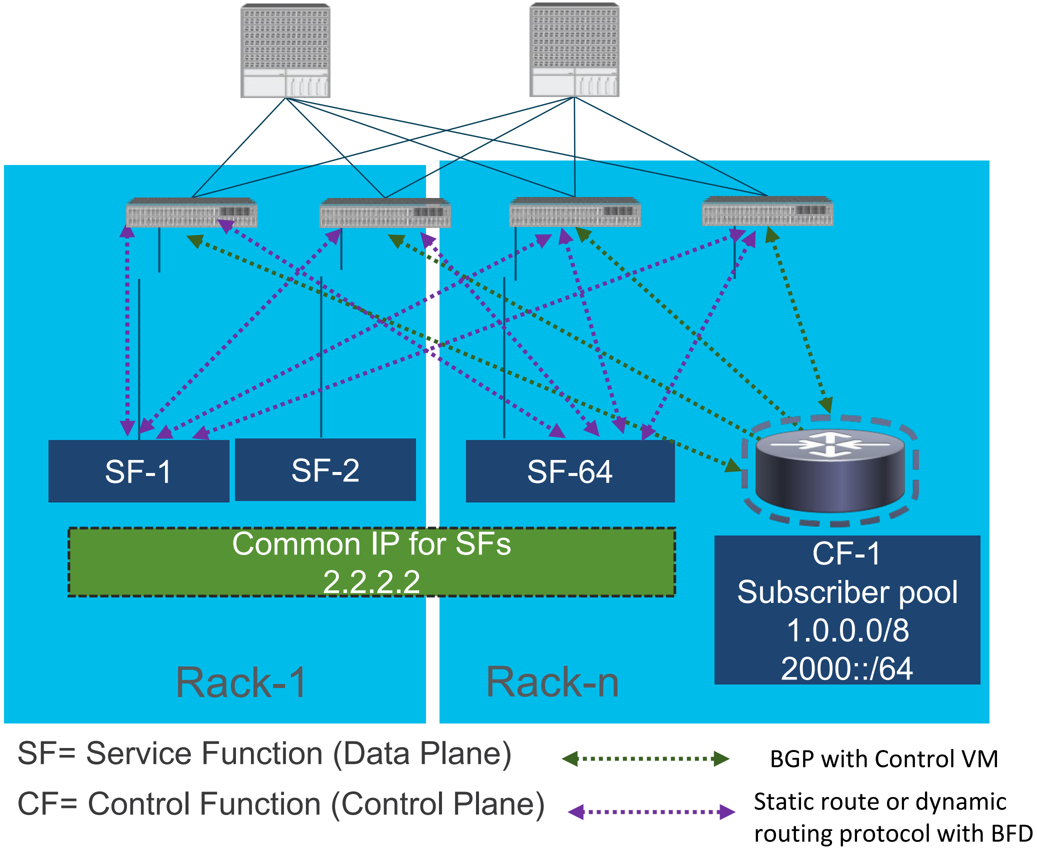




  

# vEPC design with ACI

## Design details

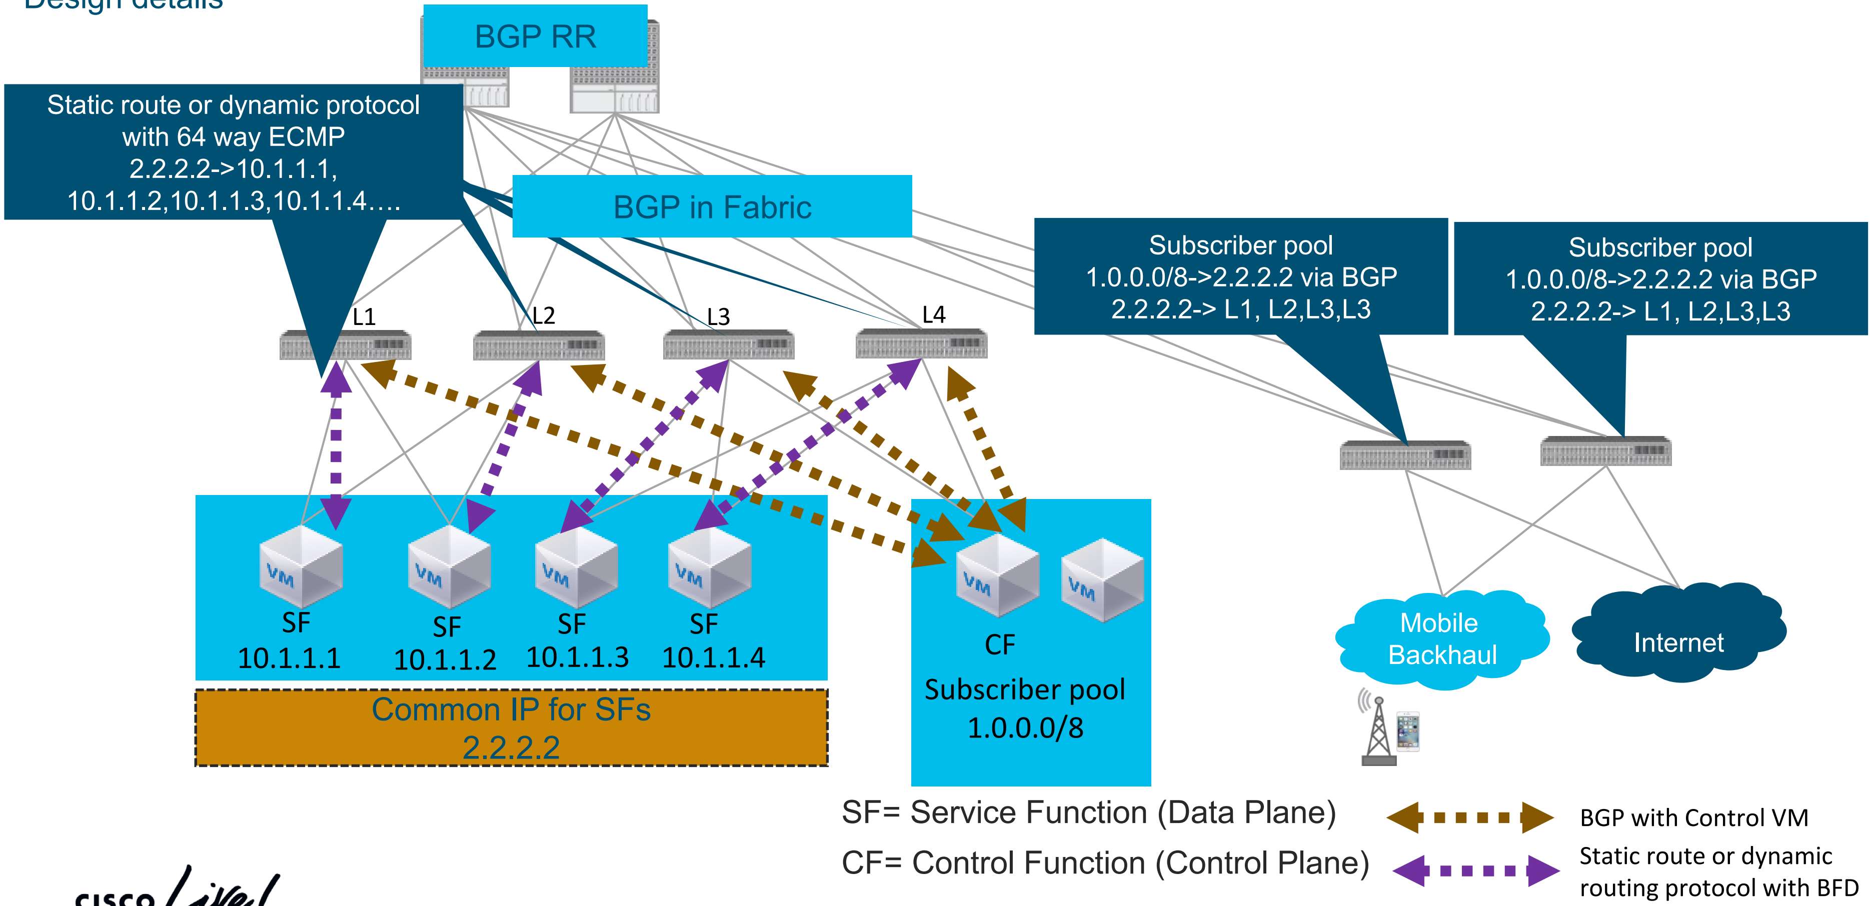


  

# vEPC design with ACI

Packet walk

Equal traffic distribution among VNFs irrespective of the amount of VNFs connected to each Leaf switch

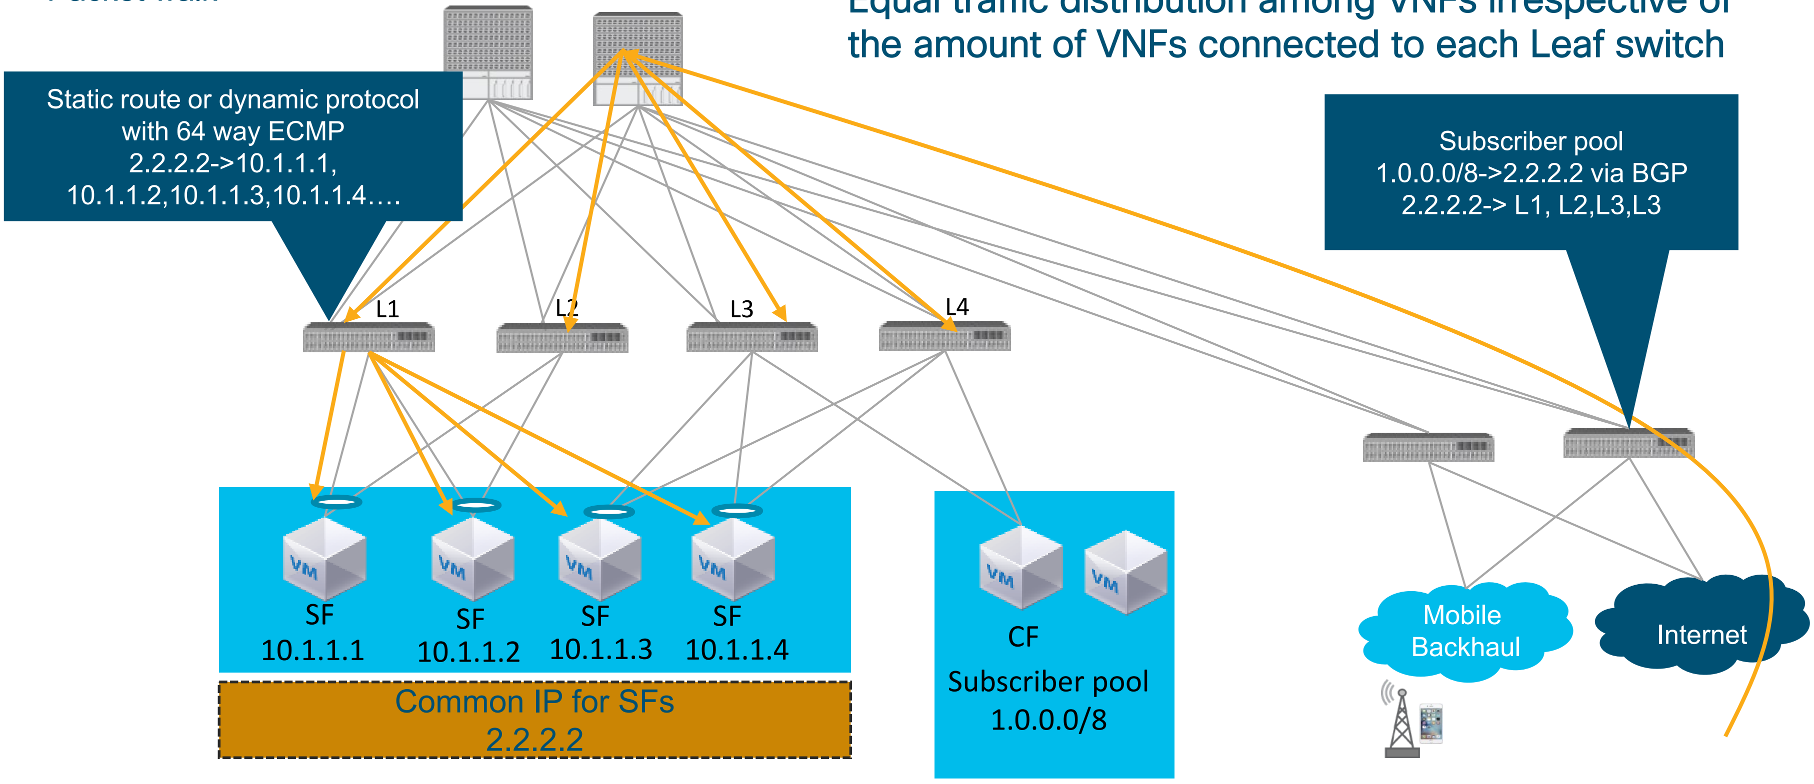


SF= Service Function (Data Plane)

CF= Control Function (Control Plane)

BGP with Control VM

Static route or dynamic routing protocol with BFD



  

# vEPC design with BGP from UPF with ACI

- ✓ 4G CUPS deployment for UPF
- ✓ Redundant and resilient design with BGP & BFD from each switch to SF
- ✓ Flexible Architecture, where UPF can be instantiated anywhere
- ✓ Each UPF is running BGP from UPF to all ACI leaf connected to vEPC cluster
- ✓ Increases overall BGP scale requirement on fabric
- ✓ 3000 BGP session across fabric is supported
- ✓ 15,000 BGP session across fabric (Q2CY20)

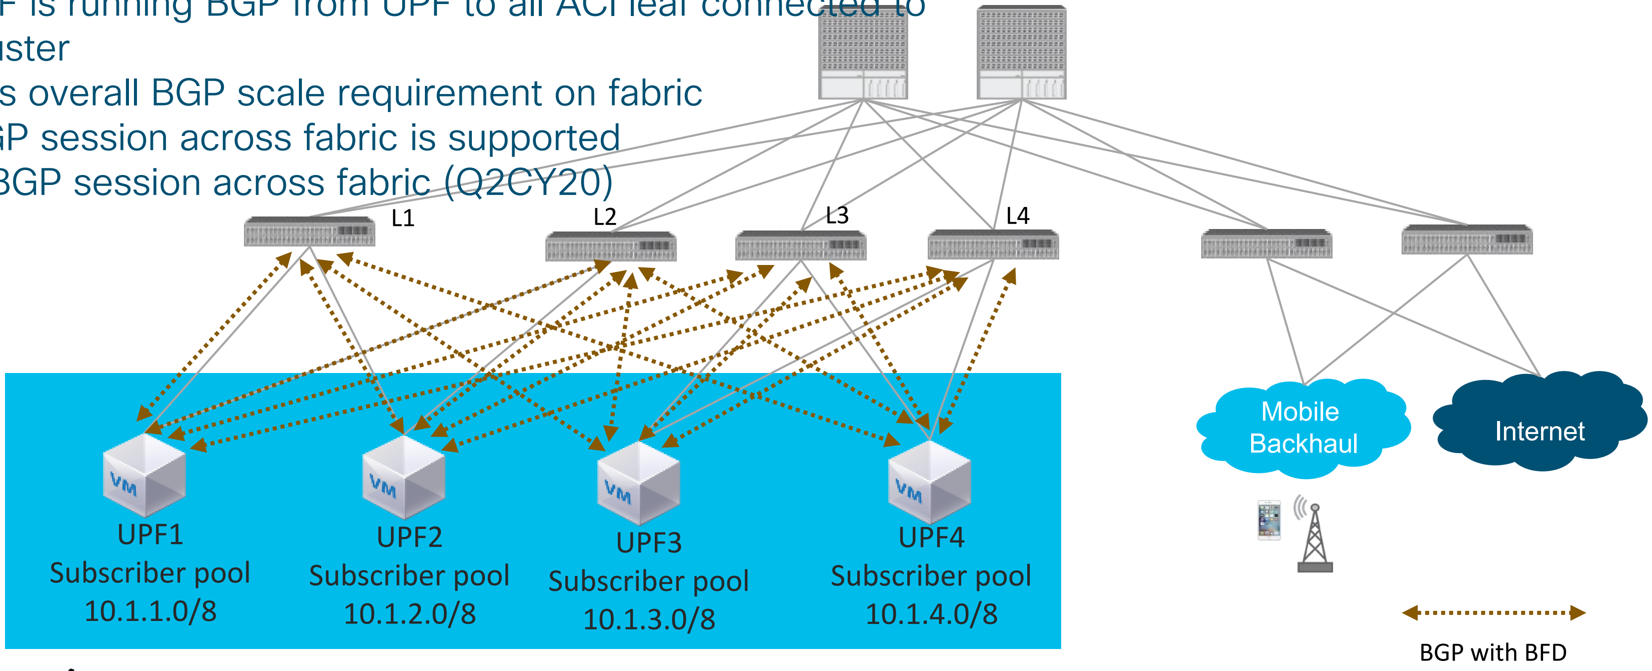


cisco Live!

  

# Load-balancing without GTP based load balancing

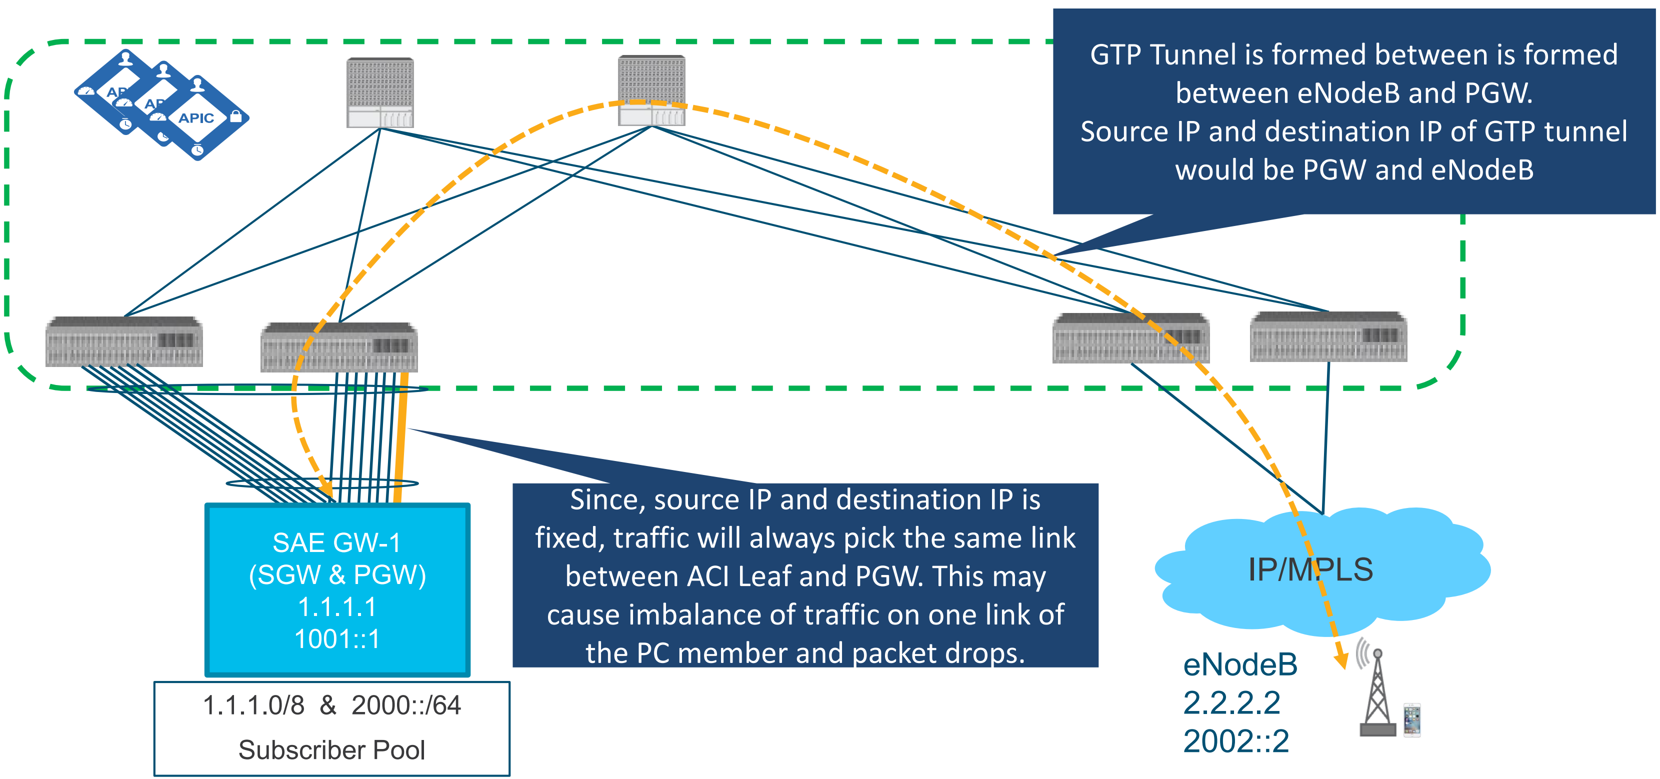

# Load-balancing without GTP based load balancing

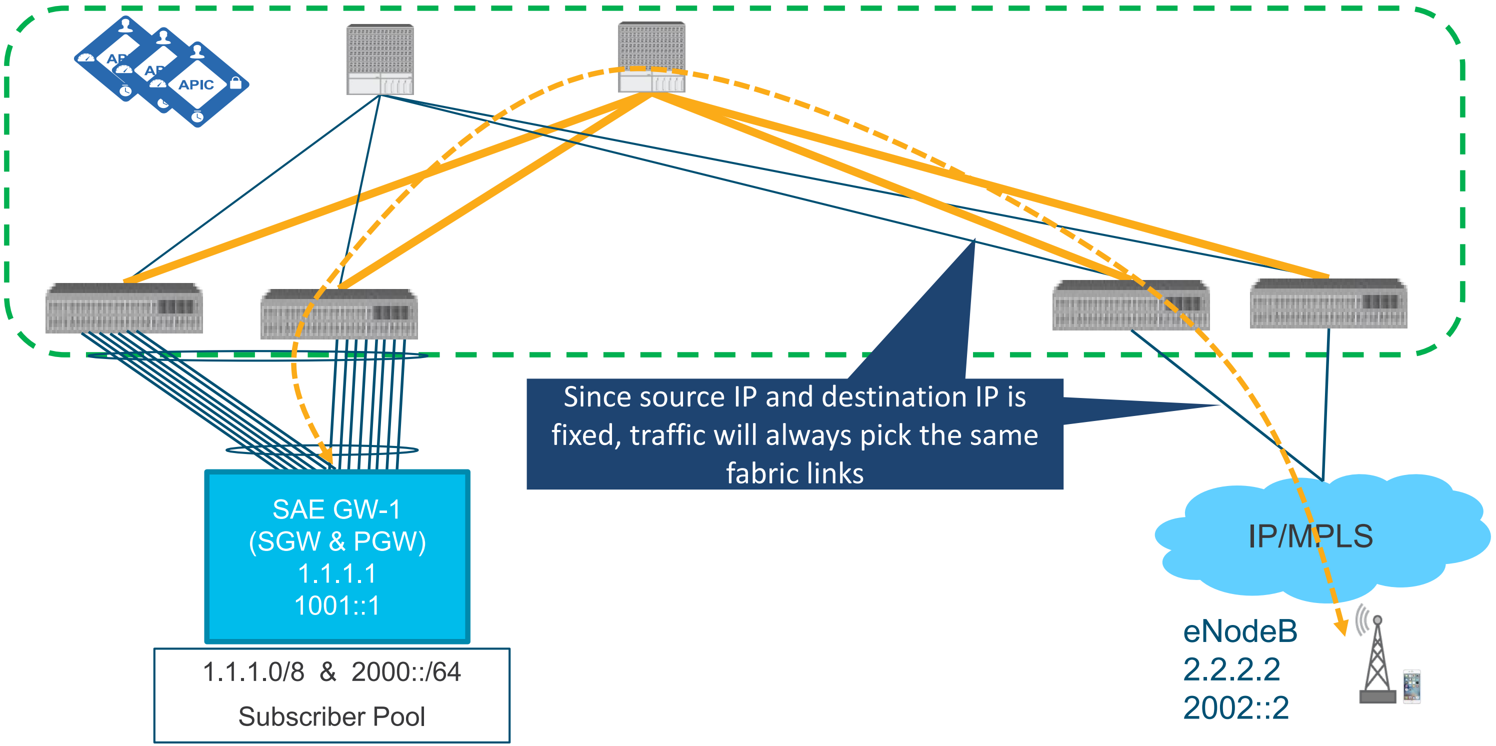


  

# GTP based load balancing

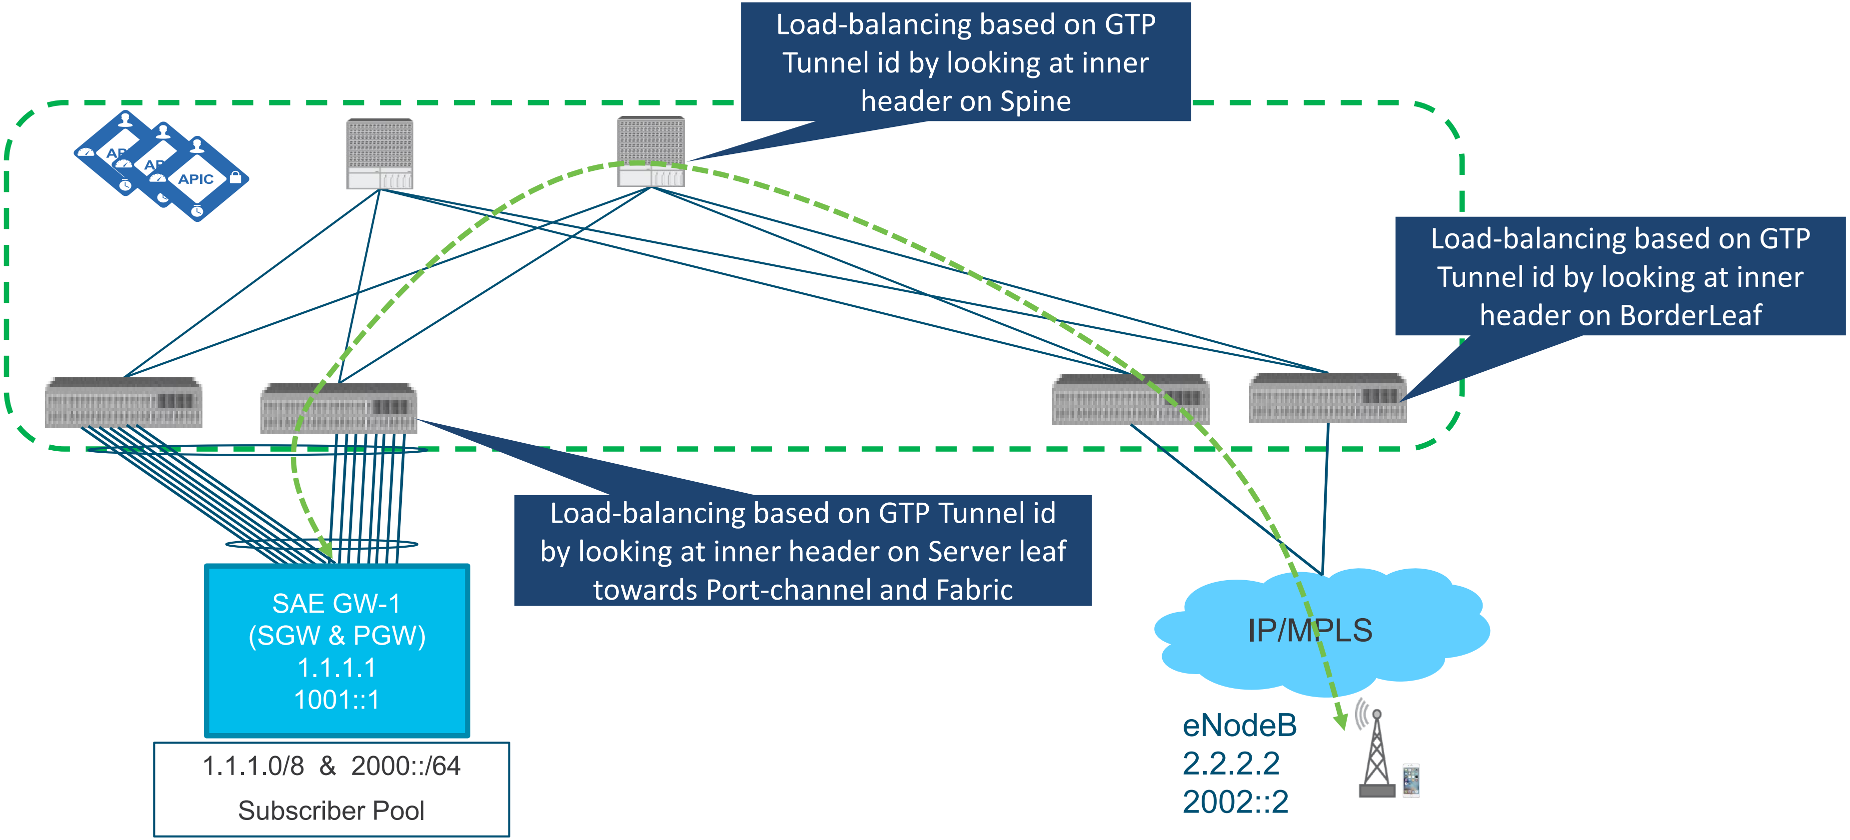

# Gi-LAN services

# Multi-Node Service chaining in Telco DC

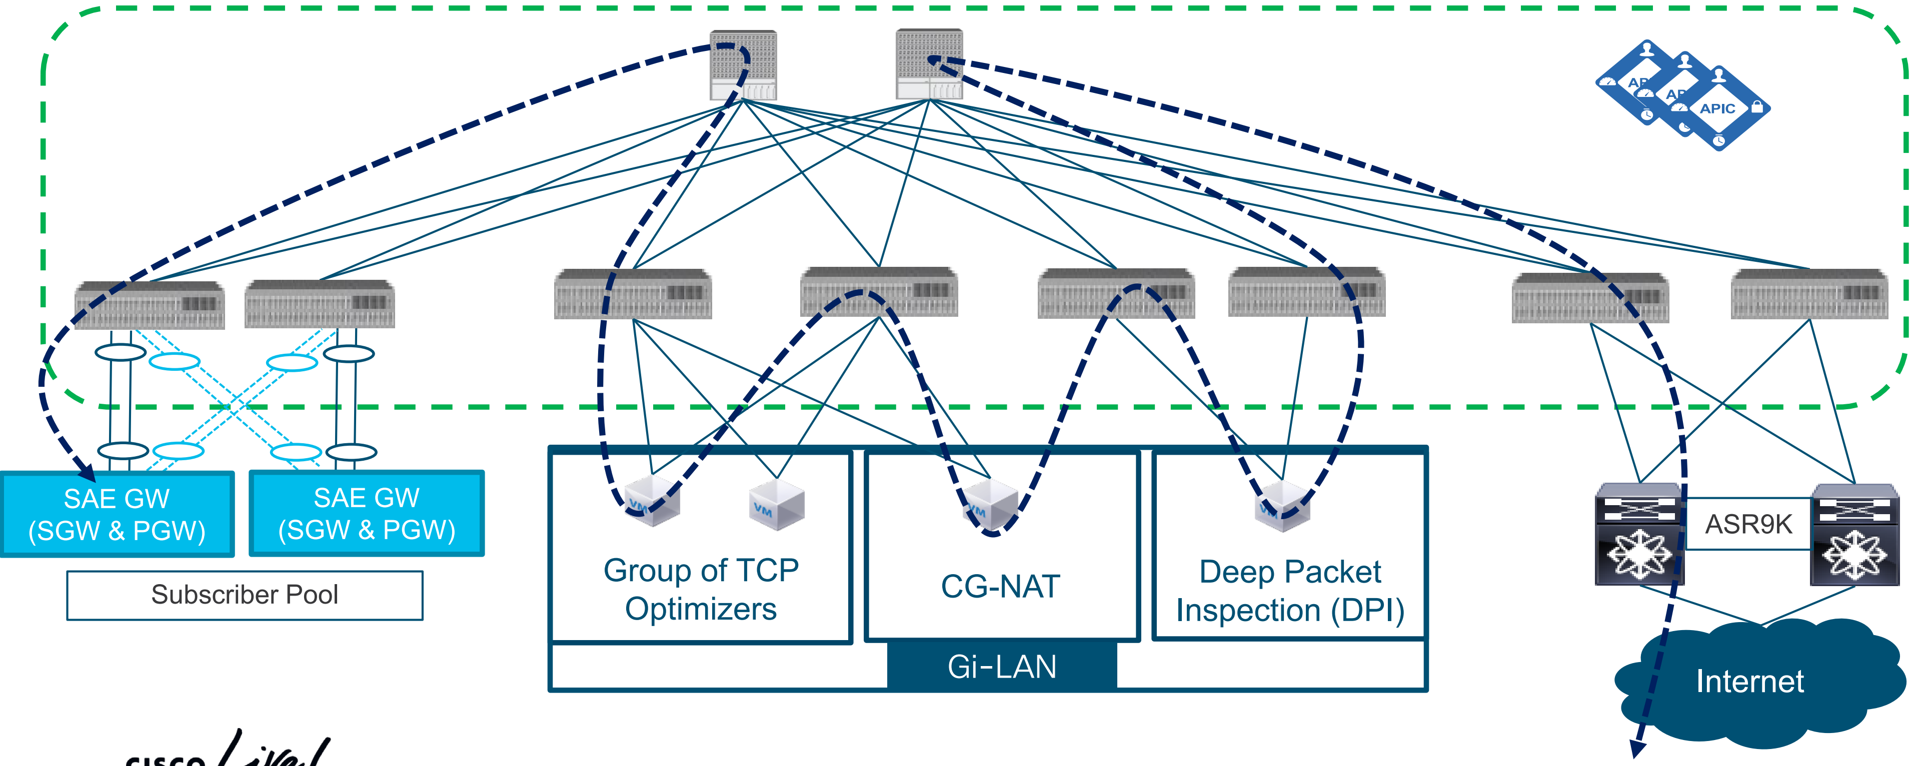

# Logical ACI Construct for service chaining

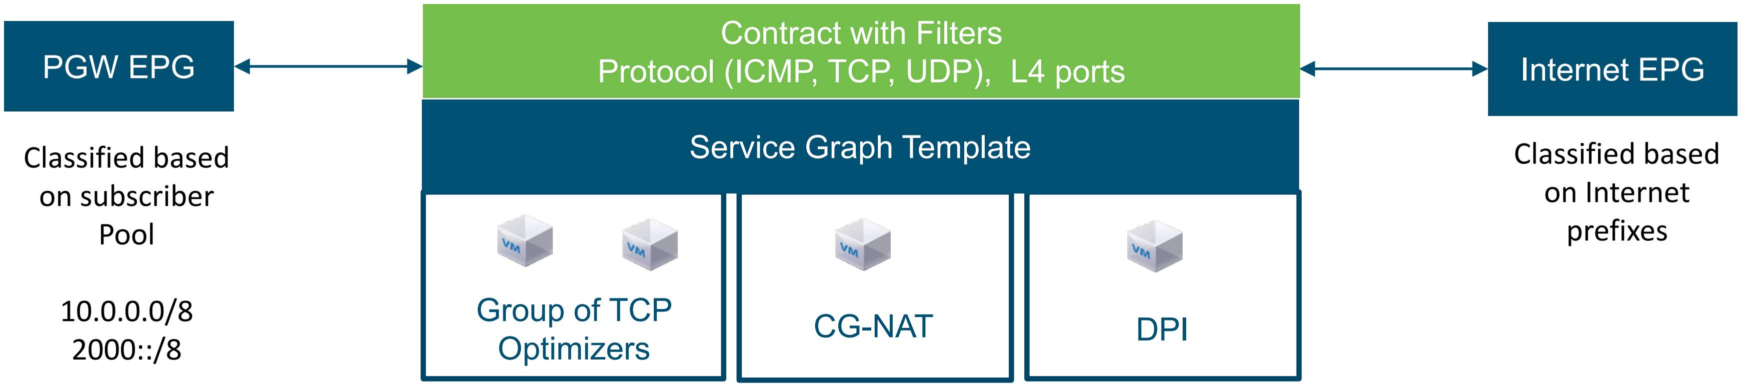


  

# Simplified Configuration

## L4-L7 Service Graph Template - service-chaining

Topology Policy Faults

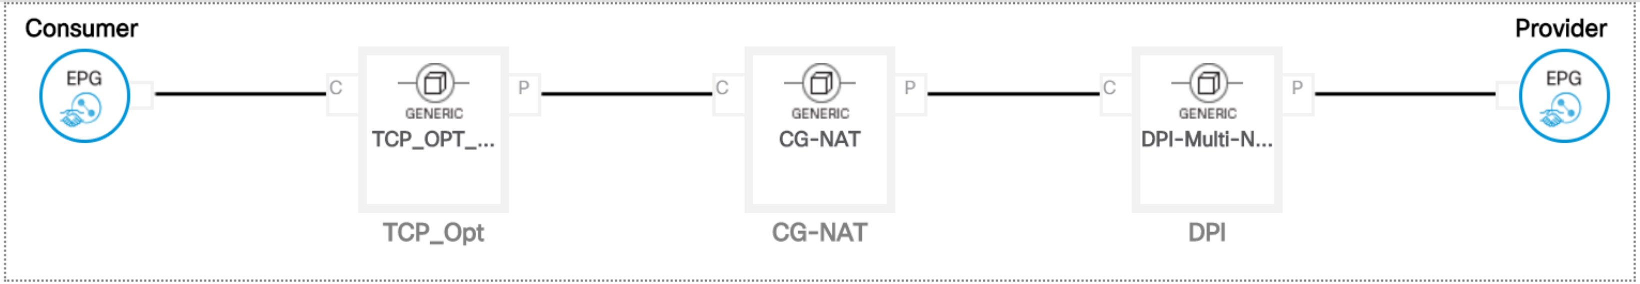


**TCP_OPT_Group Information**

Route Redirect: true

**CG-NAT Information**

Route Redirect: true

**DPI-Multi-Node Information**

Route Redirect: true

  

# PBR Bypass a node

ACI 4.1.2

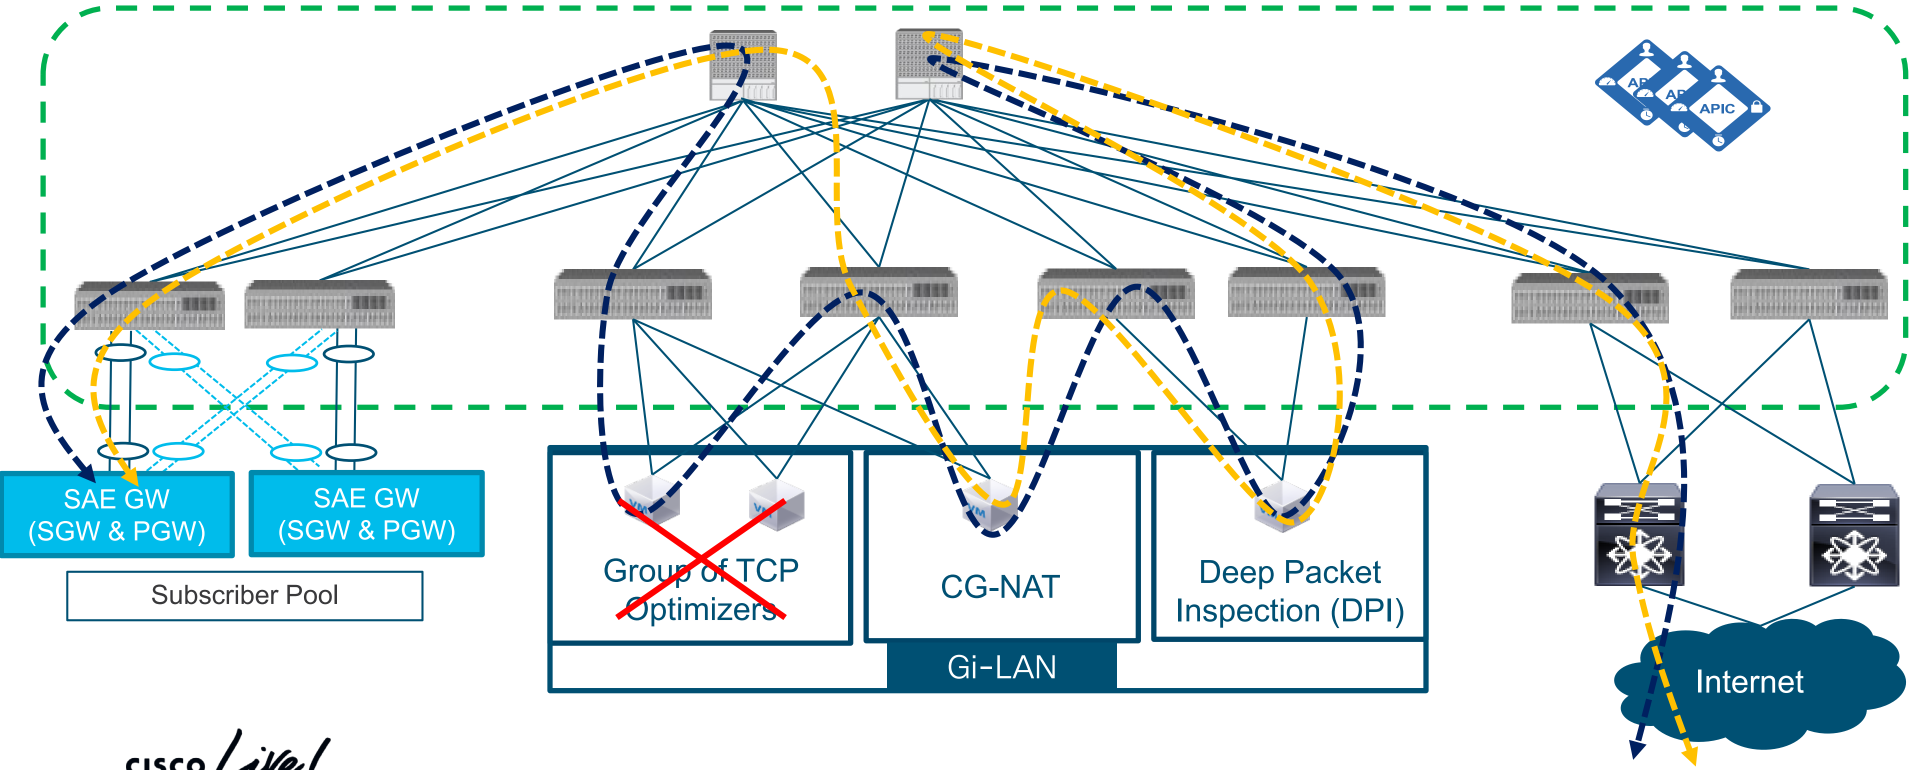

# Integration with TCP Optimizers

# TCP Optimizer Integration with ACI

## Automatic Load-balancing and Symmetry of traffic flow

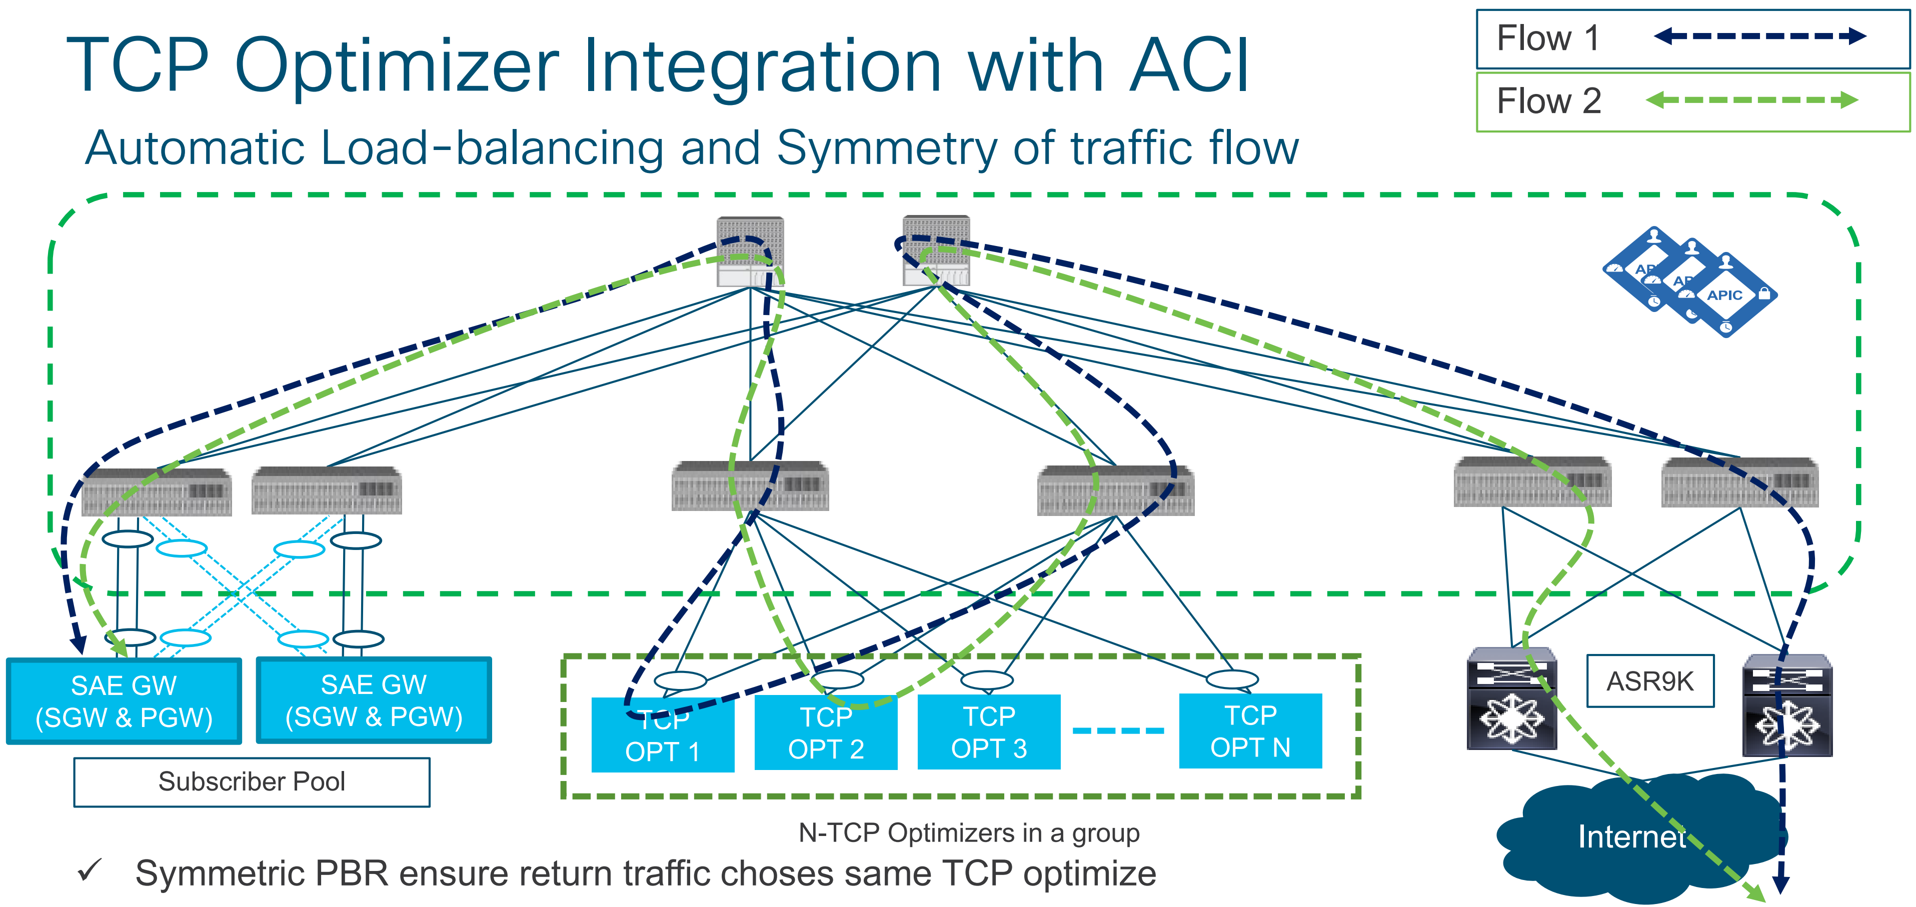


- ✓ Symmetric PBR ensure return traffic choses same TCP optimize
- ✓ Automatic load-balancing of traffic across different TCP optimizers based on forwarding table hash (Source IP, Destination IP, Source Port, Destination Port)

# TCP Optimizer Integration with ACI

## Tracking TCP Optimizer Liveliness

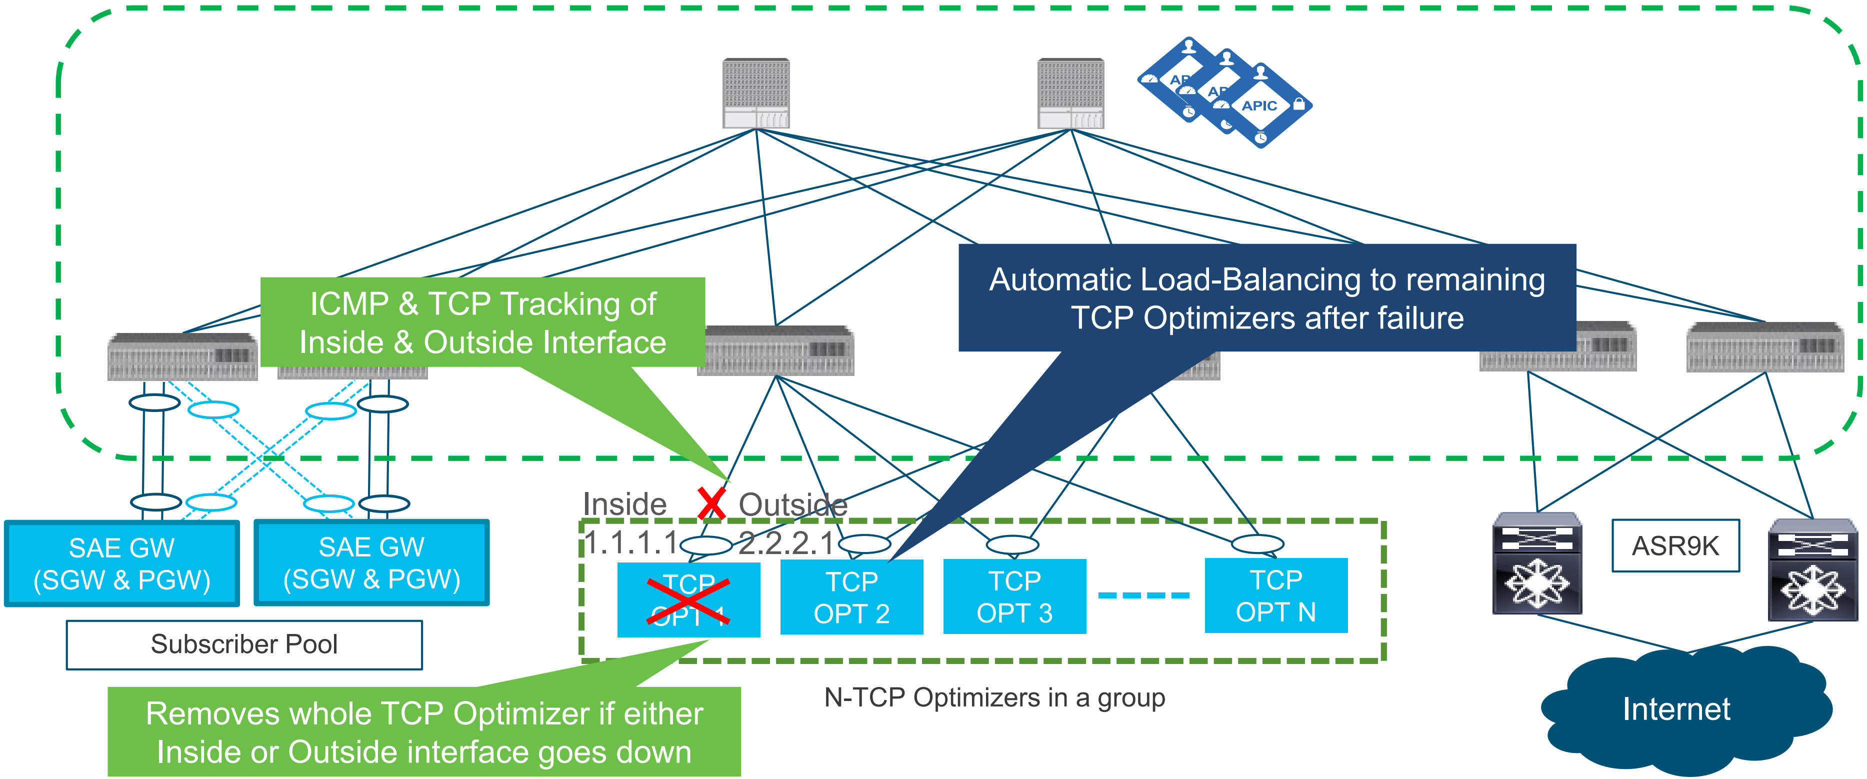

# TCP Optimizer Integration with ACI

## Bypassing TCP Optimizers to avoid congestion

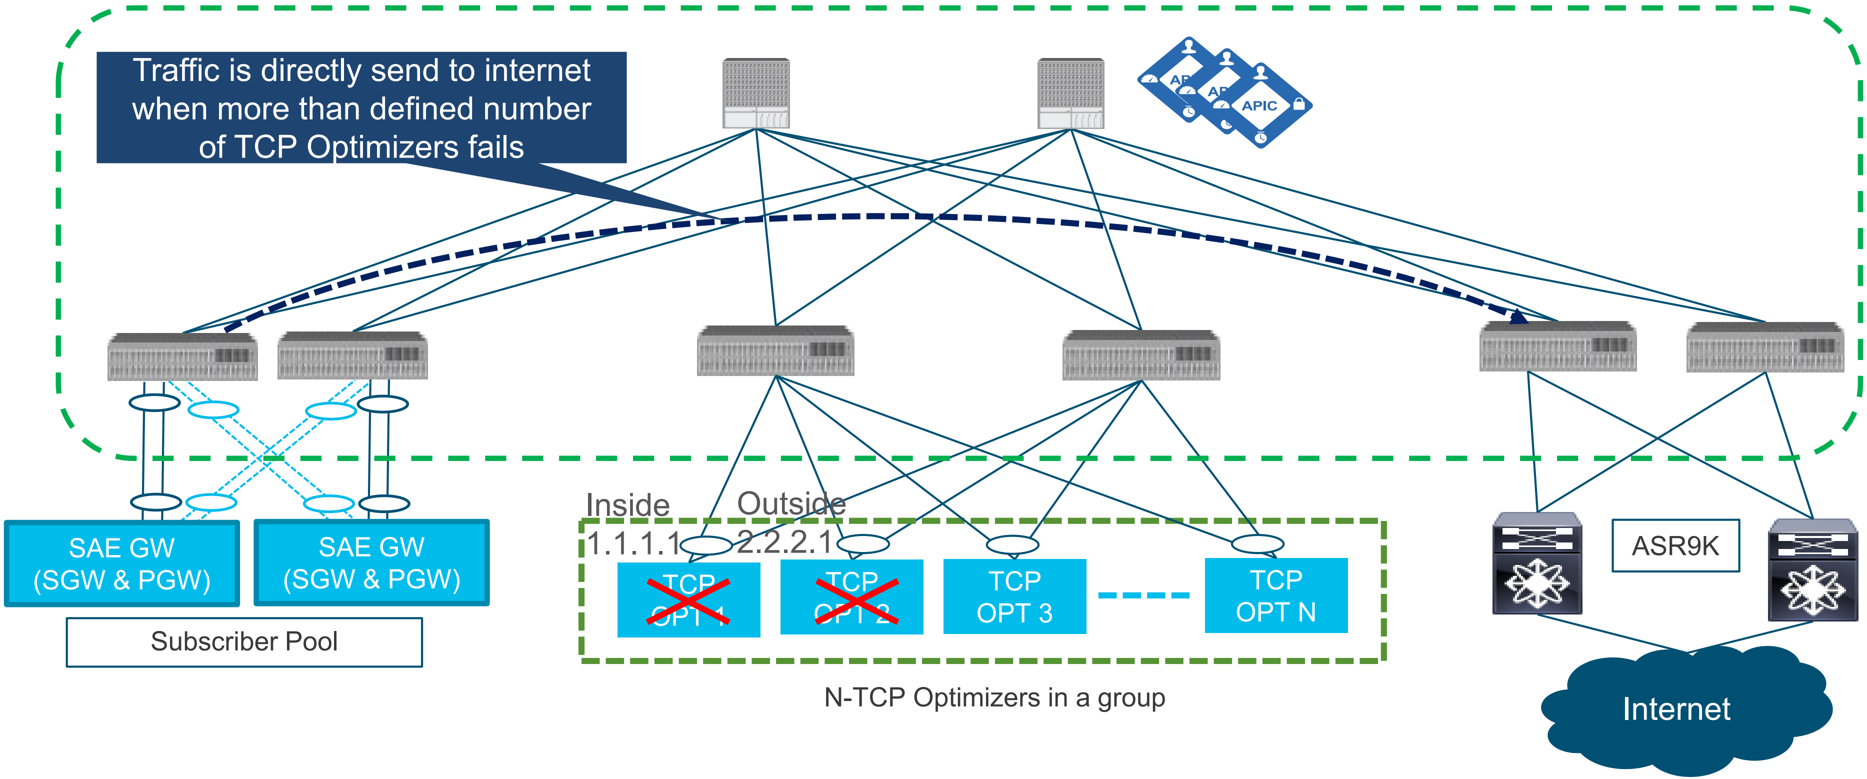

# TCP Optimizer Integration with ACI

## Non-Resilient Hashing

All flows get re-hashed on a PBR node failure, this can cause traffic drop for flows that lands on a PBR node that does not have a session information.

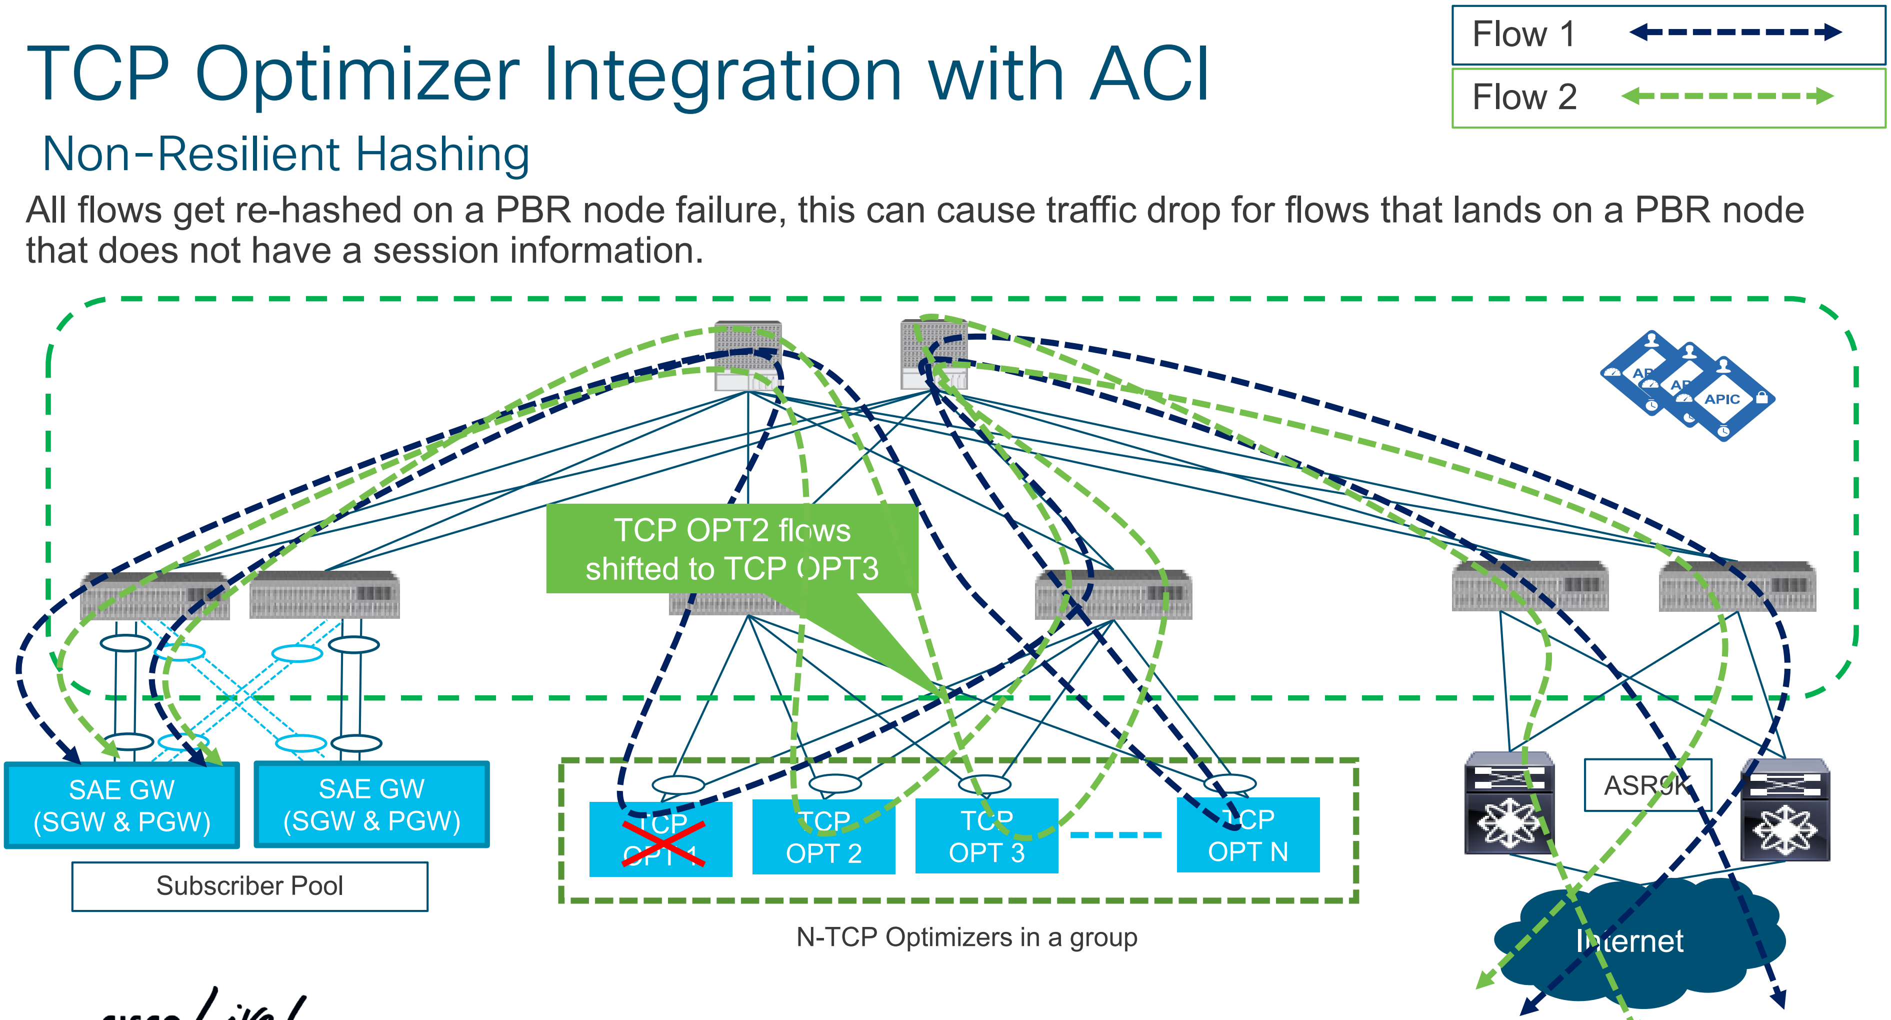


  

# TCP Optimizer Integration with ACI

## Resilient Hashing

Only the flows that were going through failed node gets re-hashed to one of the node

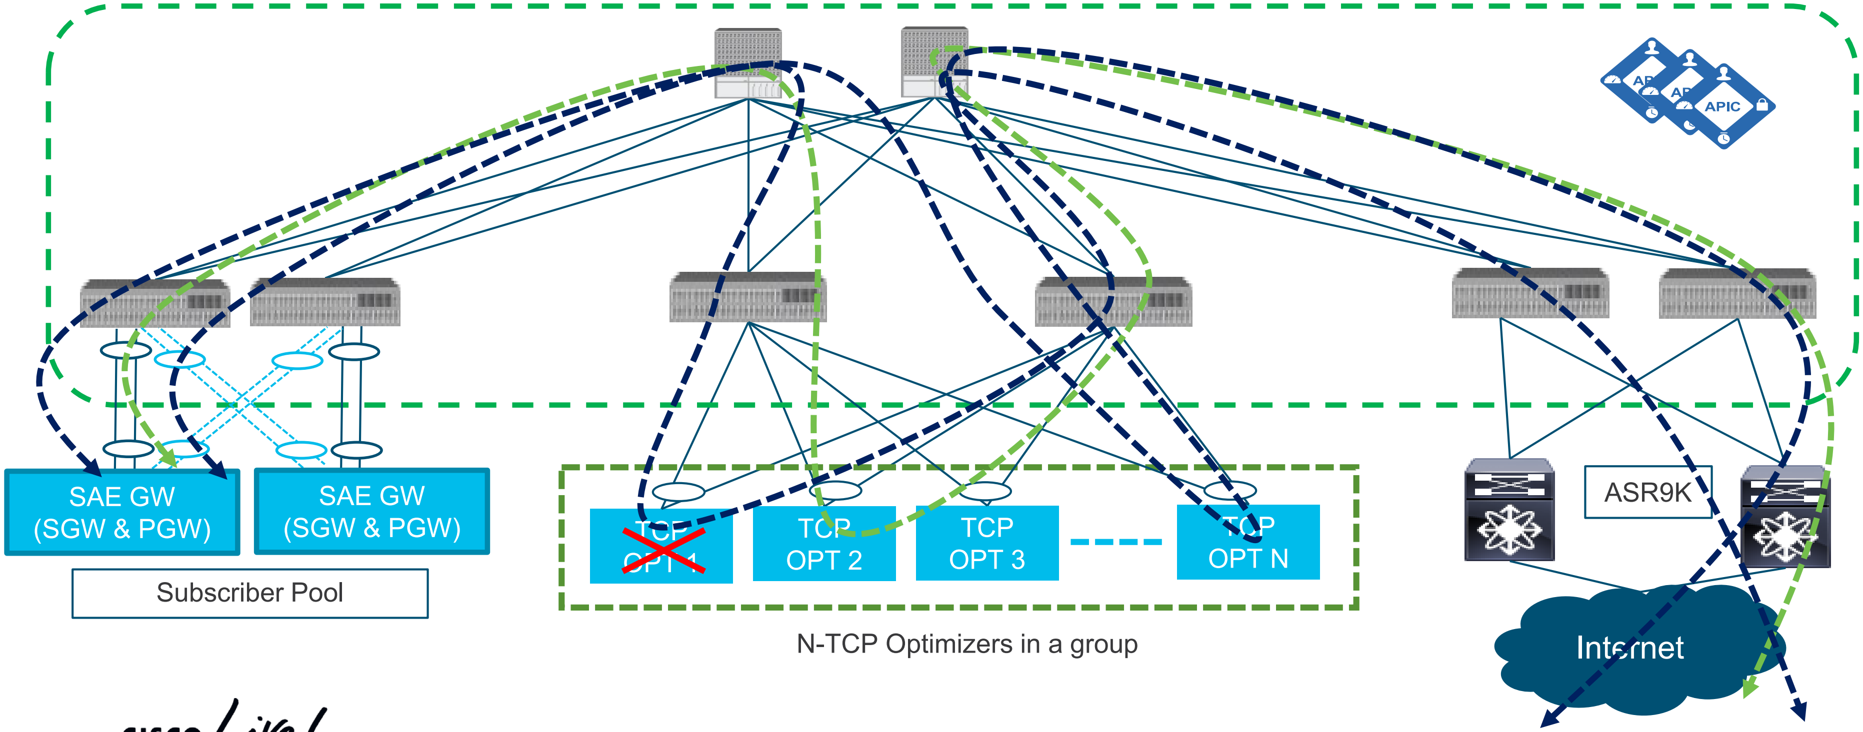



  

# TCP Optimizer Integration with ACI

ACI 4.2

## Resilient Hashing with standby node

Flow 1

Flow 2

Only the flows that were going through failed node gets re-hashed to standby node

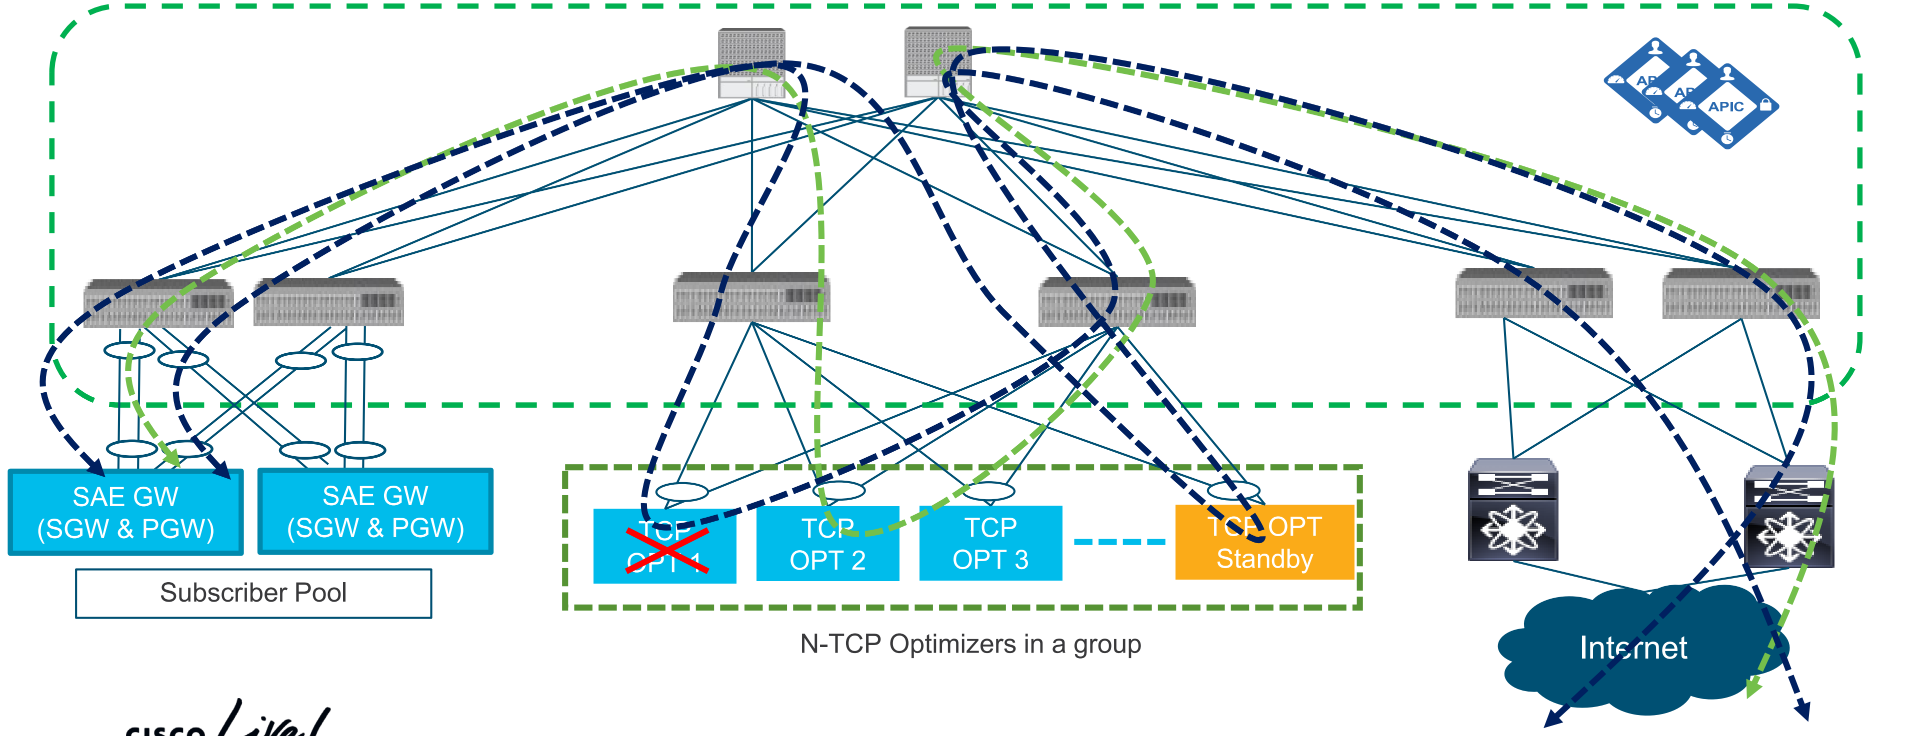

# TCP Optimizer Integration with ACI

## Simplified expansion

New TCP optimizers can be added anywhere in fabric

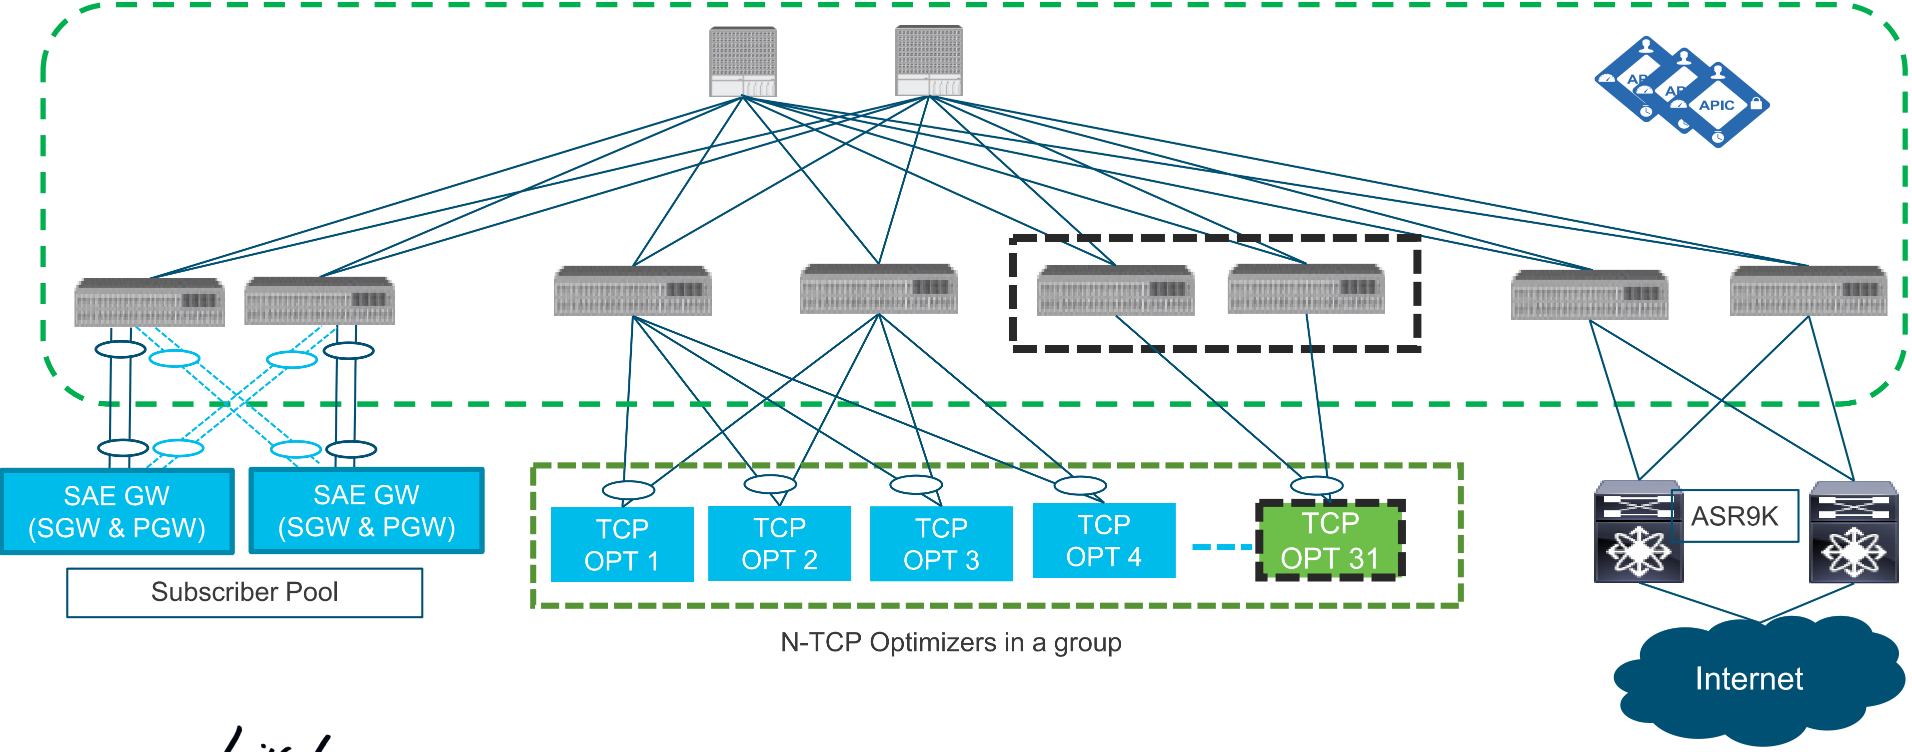

# Agenda

- Evolution of Telco Datacenter
- Customer Profiles & Technical Requirements
- **How ACI met the requirements**
  - Automation
  - Services Connectivity
  - **Operational Simplification**
  - Faster Convergence
  - Distributed Edge DC
- Looking Ahead
- Conclusion

# Operations Tools

## Topology Dashboard

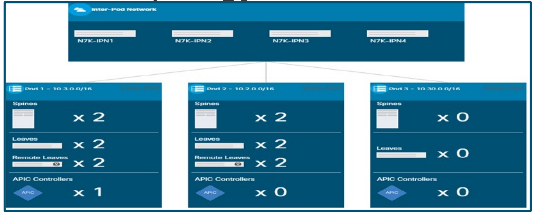

## Troubleshooting Wizard

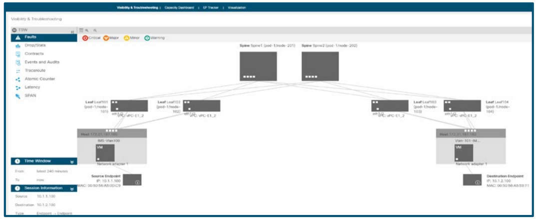

## End Point Tracker

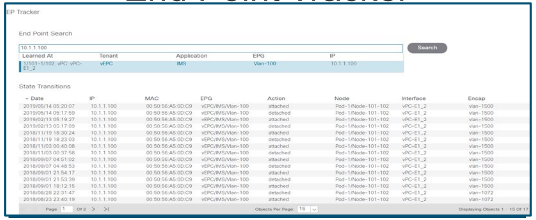

## Faults

## Capacity DashBoard

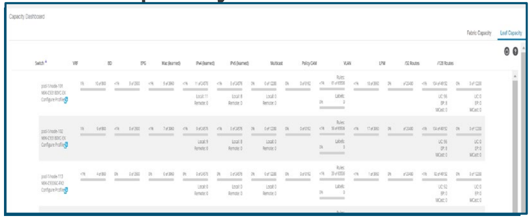

## Link Statistics

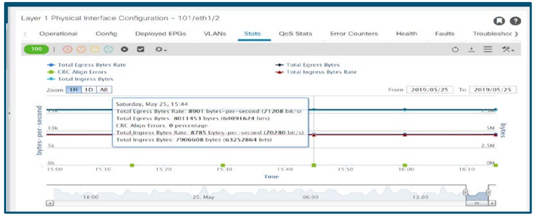

## Health Score Card

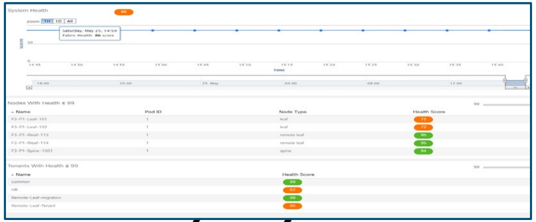

## Traffic Map

## Upgrade/Downgrade

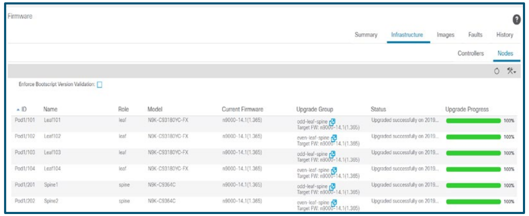

# Troubleshooting Wizard

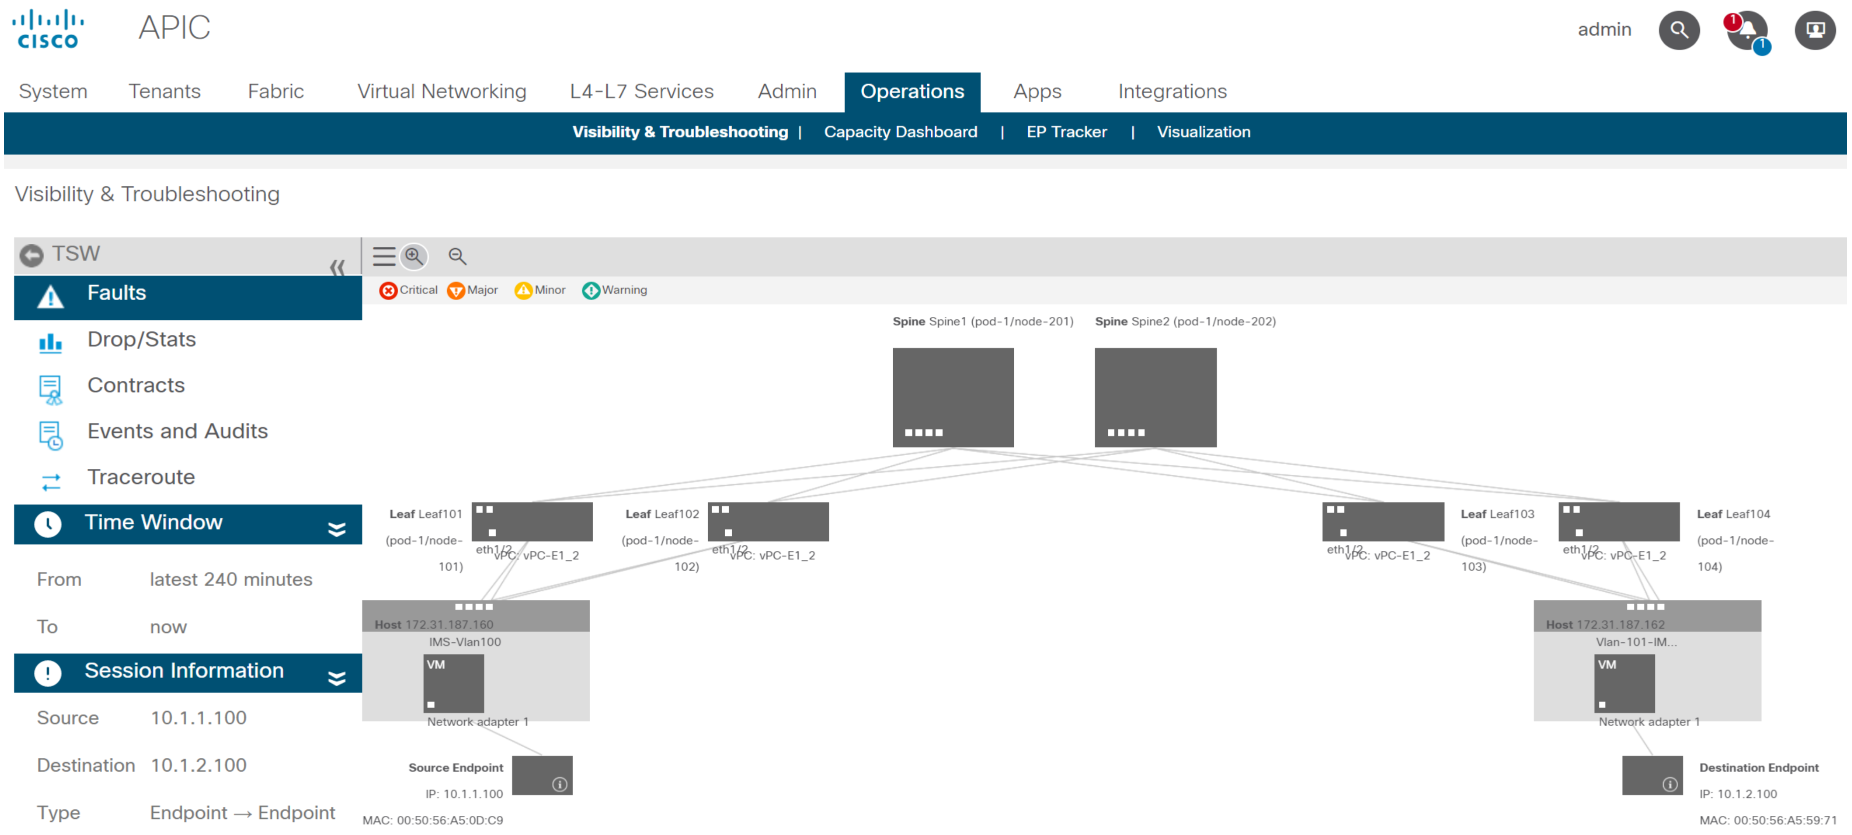




  

# Packet capture on APIC controller

- ✓ Easy to SPAN packets to APIC without dedicated server
- ✓ Packet policer to limit the number of packets sent to APIC
- ✓ Requires in band connectivity
- ✓ 10 PCAP files of 1MB each to limit the usage of APIC

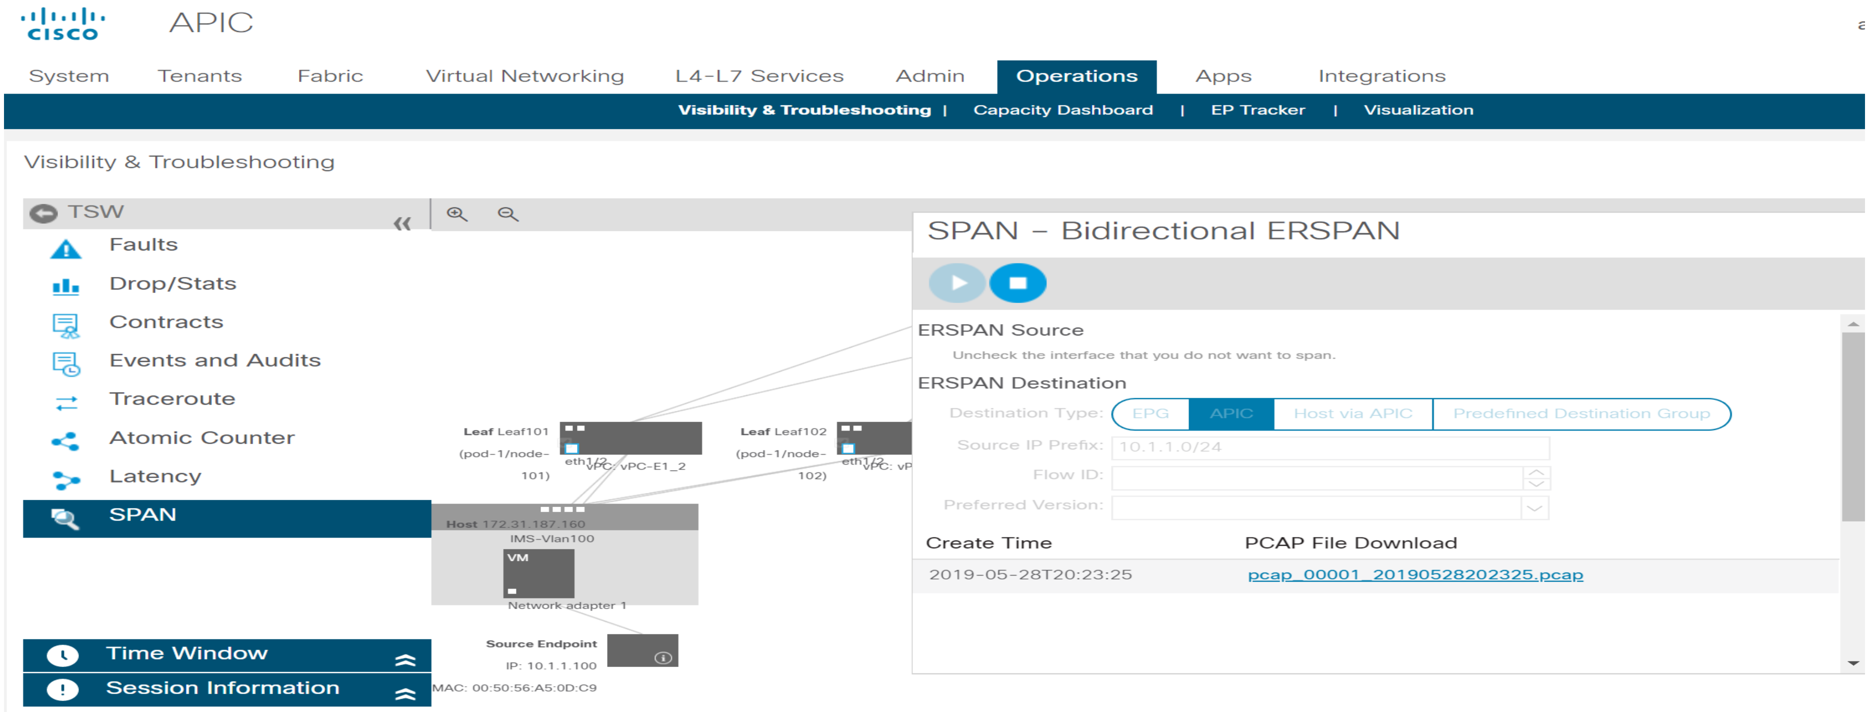



  

ACI 4.1

# SPAN based on filter

## SPAN session on ACI that sends SPANs packet matching the 5-tuples filter

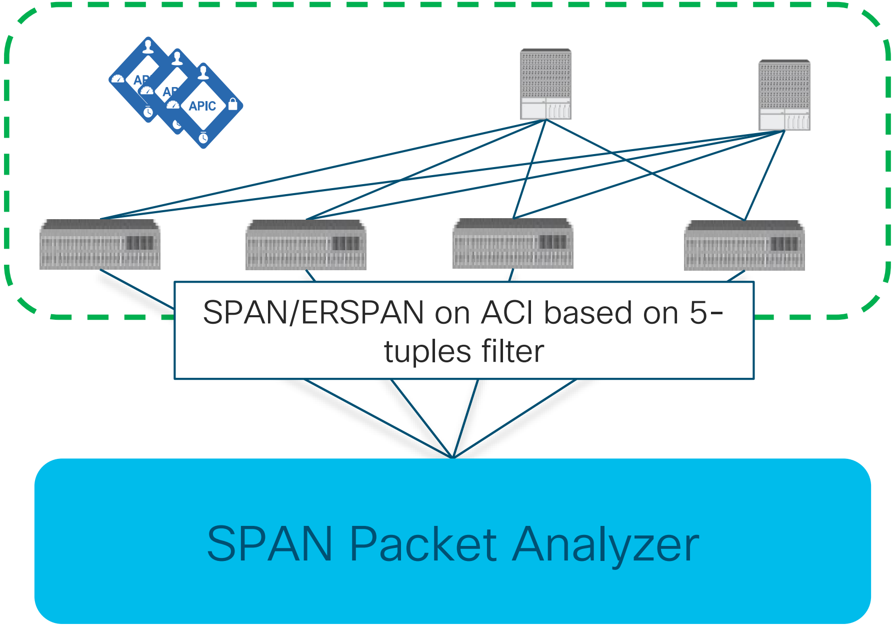


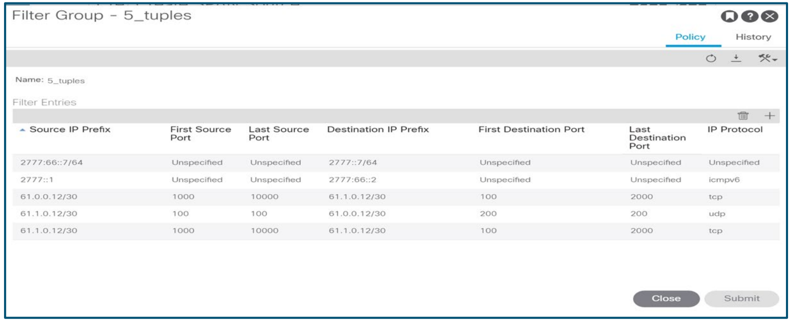


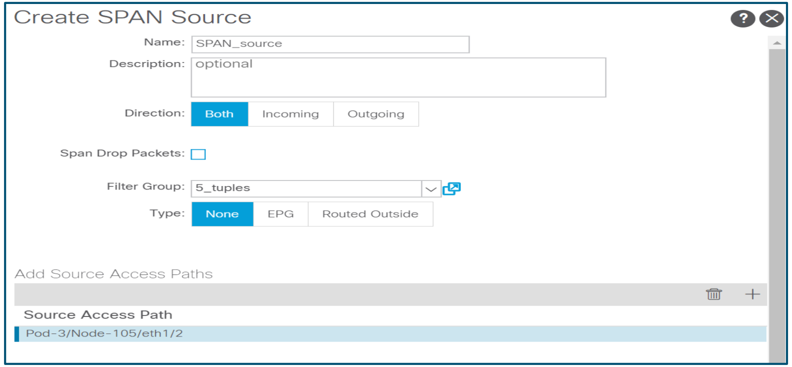


  

# SPAN on drop

## SPAN session on ACI that sends dropped packet to Analyzer

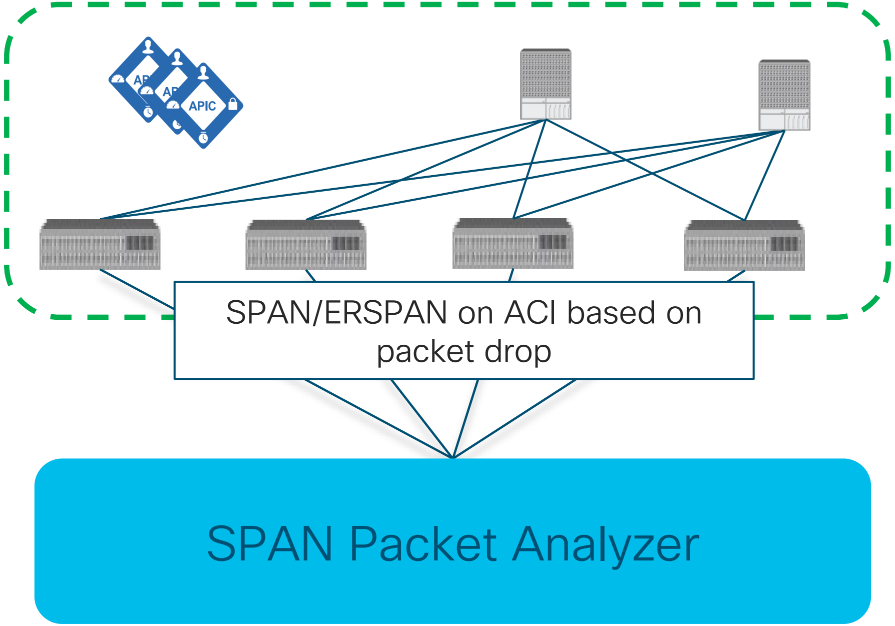


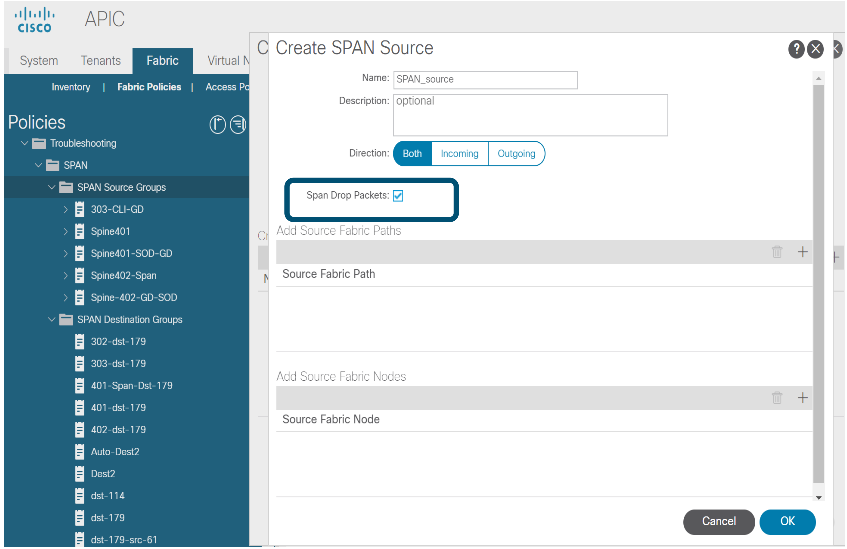


  

# ERSPAN from ACI fabric to Monitoring tools

- ✓ S1U traffic monitoring
- ✓ ERSPAN from each leaf to monitoring destination
- ✓ Dedicated rack for traffic monitoring tools
- ✓ Monitoring appliance could scale up to 10G traffic causing ERSPAN scale requirement on ACI fabric
- ✓ Increased ERSPAN scale from 4 bi-directional session to 16 bi-directional session per leaf

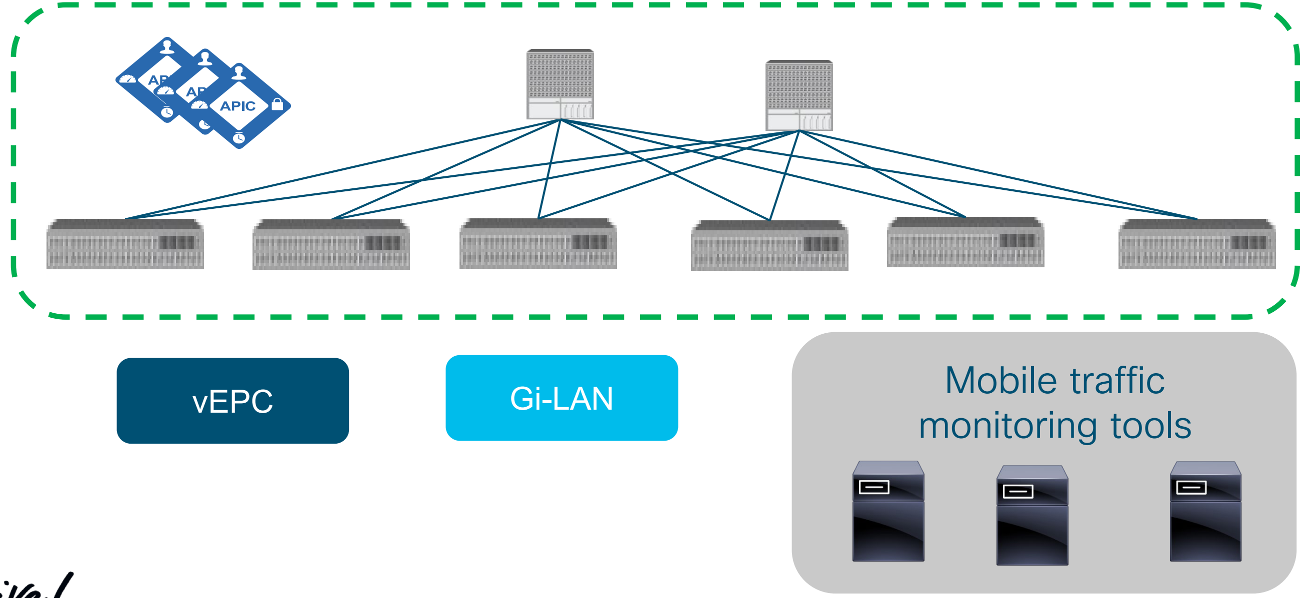


  

# Simple Return Merchandise Authorization (RMA)

Decommission the old switch by removing the controller

Register the new switch by provide same Node ID and Node Name

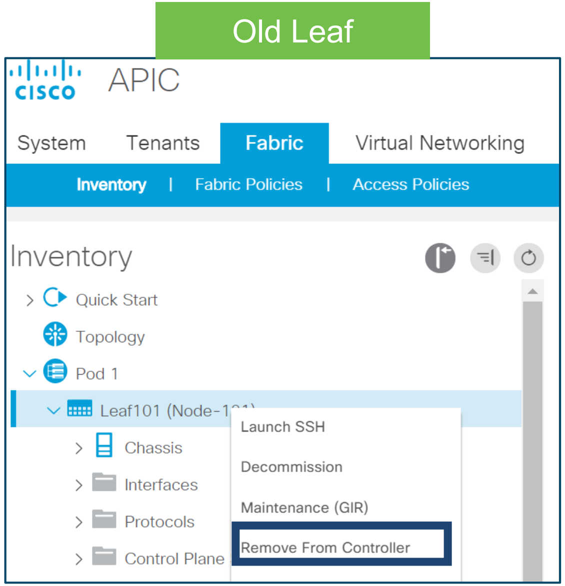


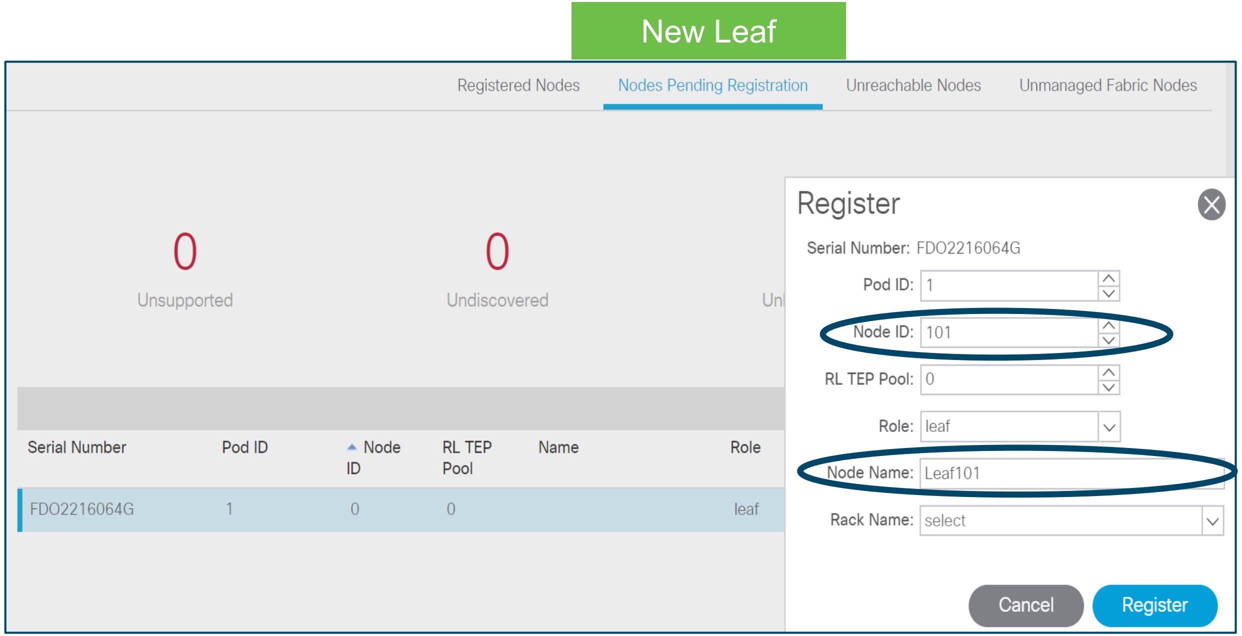


  

# Agenda

- Evolution of Telco Datacenter
- Customer Profiles & Technical Requirements
- **How ACI met the requirements**
  - Automation
  - Services Connectivity
  - Operational Simplification
  - **Faster Convergence**
  - Distributed Edge DC
- Looking Ahead
- Conclusion



  

# ACI Fabric Convergence

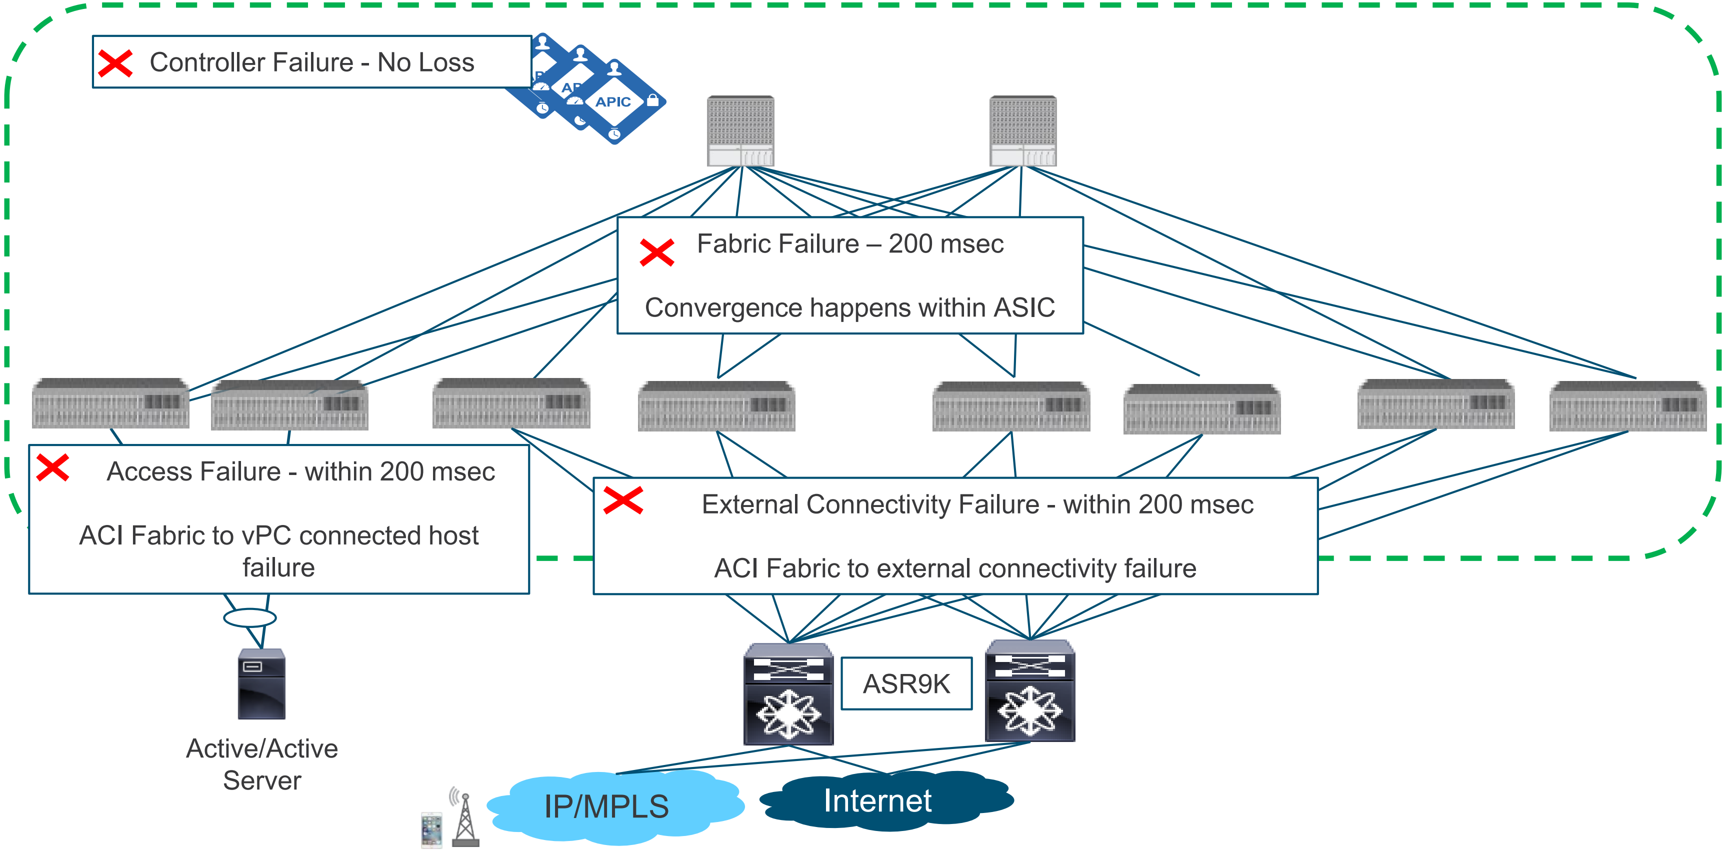

# Agenda

- Evolution of Telco Datacenter
- Customer Profiles & Technical Requirements
- **How ACI met the requirements**
  - Automation
  - Services Connectivity
  - Operational Simplification
  - Faster Convergence
  - **Distributed Edge DC**
- Looking Ahead
- Conclusion

# Architecture for Distributed Datacenters

## ACI Remote Physical Leaf

Remote Location contains Nexus 9300 connected to IP Network and fully managed by APIC cluster of Main DC

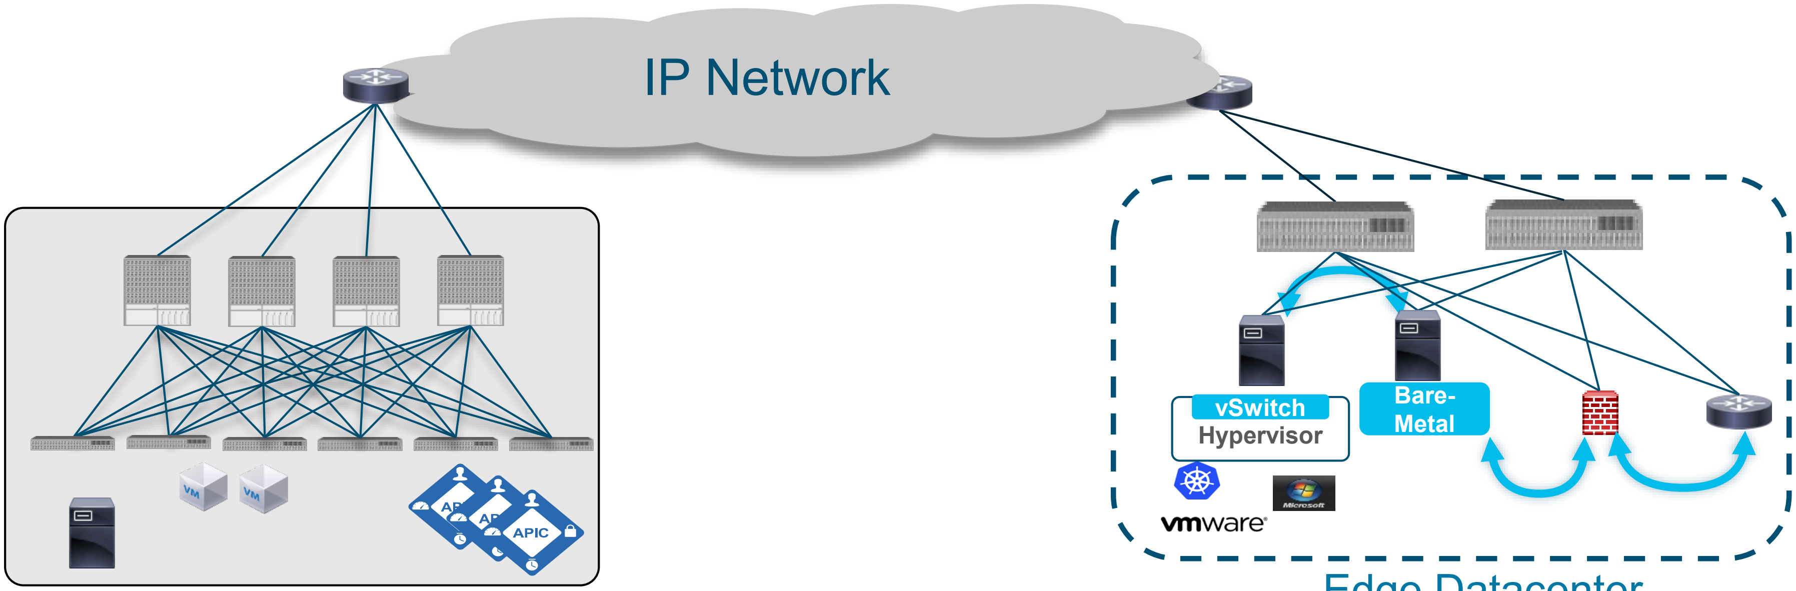


Central/Regional Datacenter

Edge Datacenter

All local traffic is switched directly between endpoints, both virtual and bare metal

  

# Architecture for Distributed Datacenters

Before ACI 4.1.2

## Traffic forwarding between pair of Remote Leaf switches

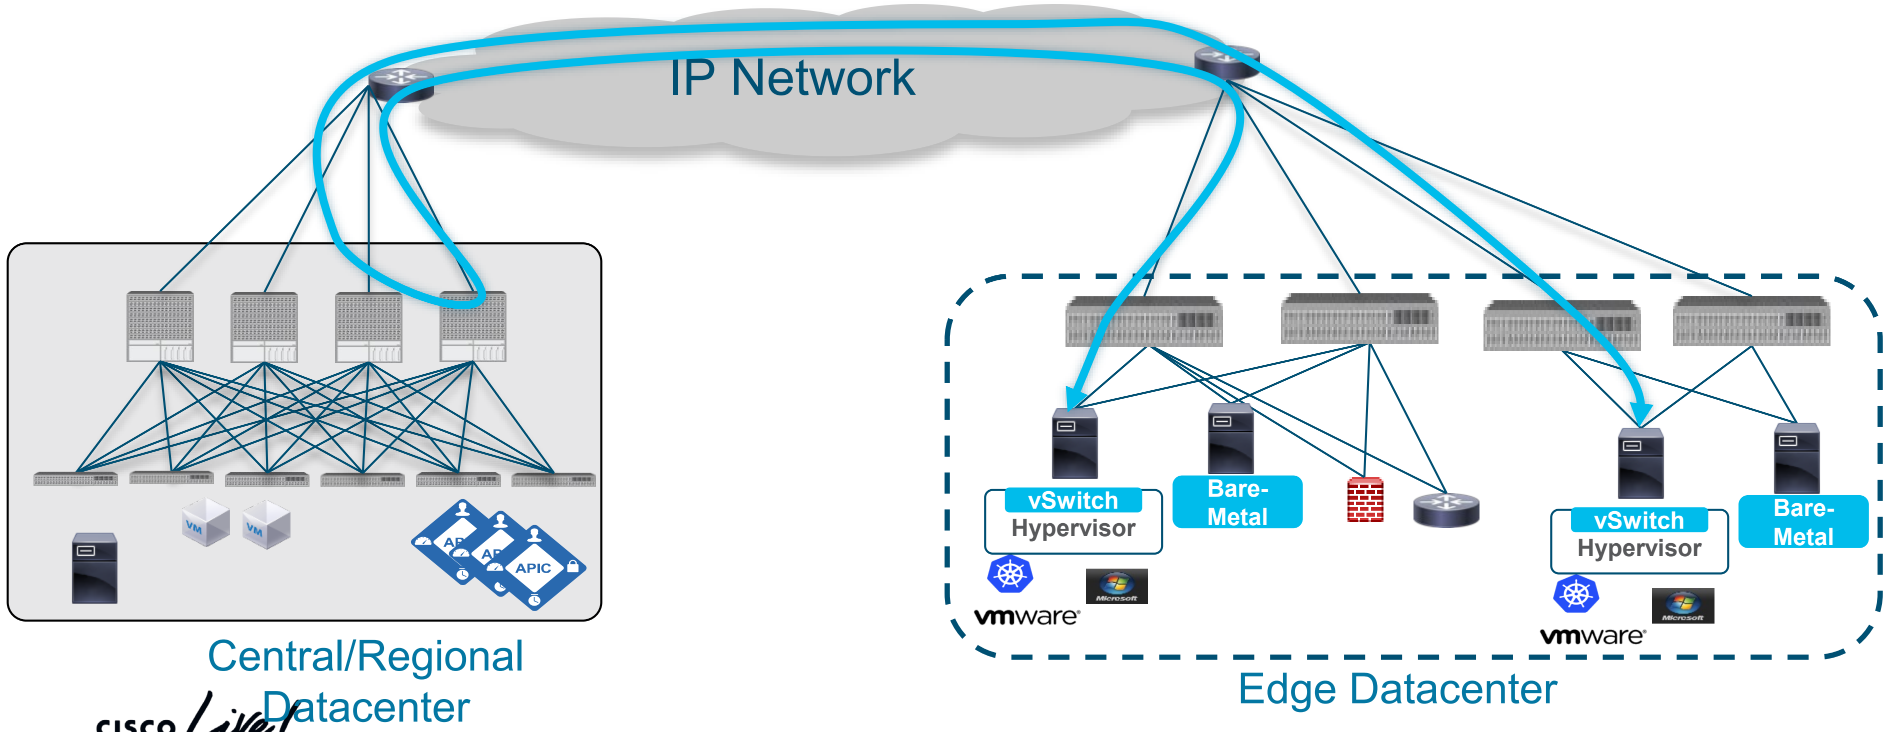


  

# Architecture for Distributed Datacenters

4.1.2

Direct traffic forwarding between pair of Remote Leaf switches

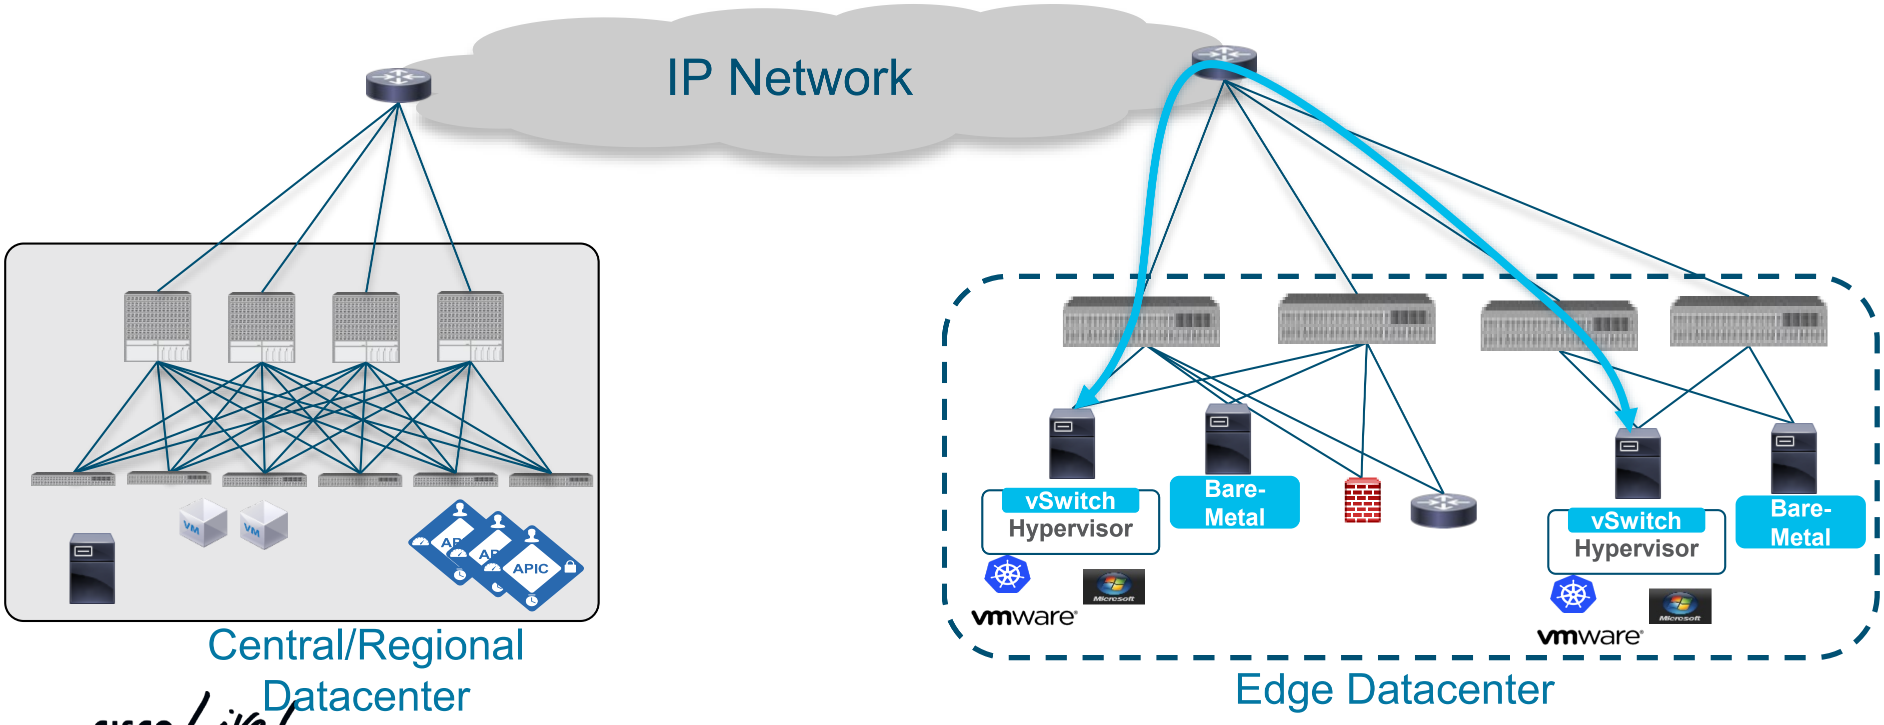


  

# Distributed Edge DC with ACI Remote Leaf

- ✓ IP Network is only for Remote Leaf management purpose
- ✓ Each DC has local VRF that maps to common VRF on DC-Edge router. This ensures traffic always take IP MPLS/ WAN path instead of IPN (VXLAN) path
- ✓ Non-VXLAN (IP MPLS/ WAN) path between DCs allows customer to use different policies in transport

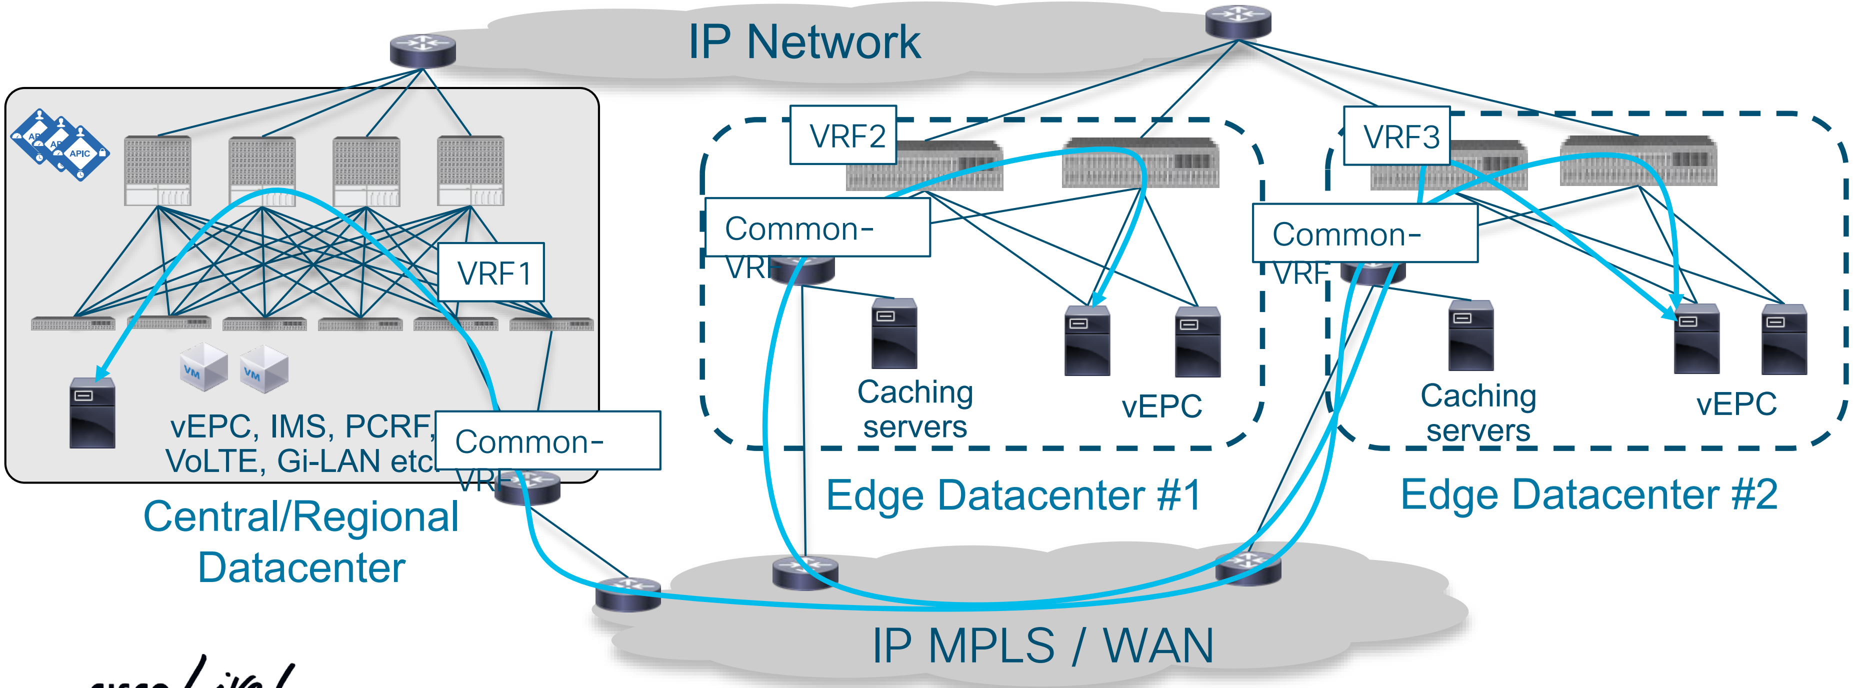

# Agenda

- Evolution of Telco Datacenter
- Customer Profiles & Technical Requirements
- How ACI met the requirements
- **Looking Ahead**
- Conclusion


  

# ACI Multi-Pod

## Single Telco DC Campus with multiple server halls

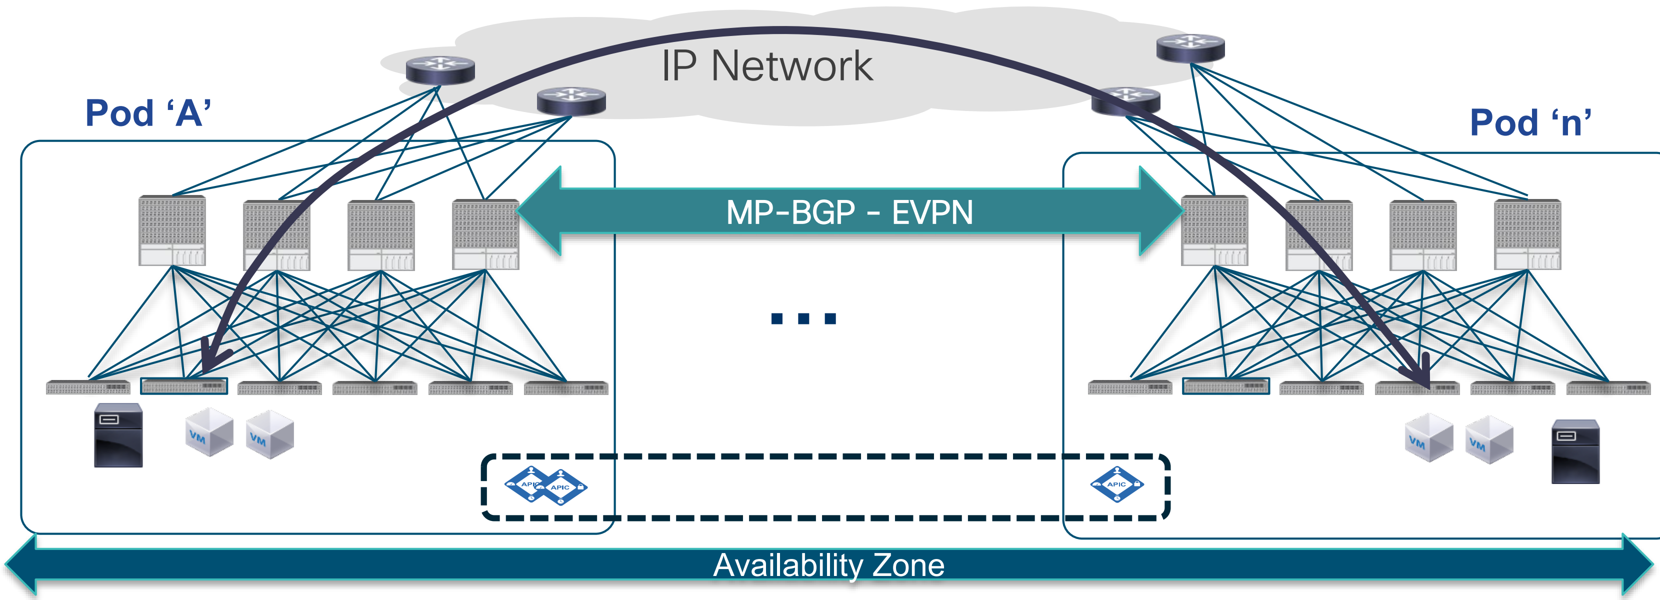


- Managed by a single APIC Cluster
- Single Management and Policy Domain
- End-to-end policy enforcement
- Control plane fault isolation

# ACI Multi-Pod with transport network slicing

- ✓ Multi-Pod with each Pod in different location, and with a requirement of network slicing in the transport
- ✓ ACI Multi-Pod solution is only used for management of multiple Pods in this scenario
- ✓ Traffic between Pod is forwarded through WAN and not through IPN

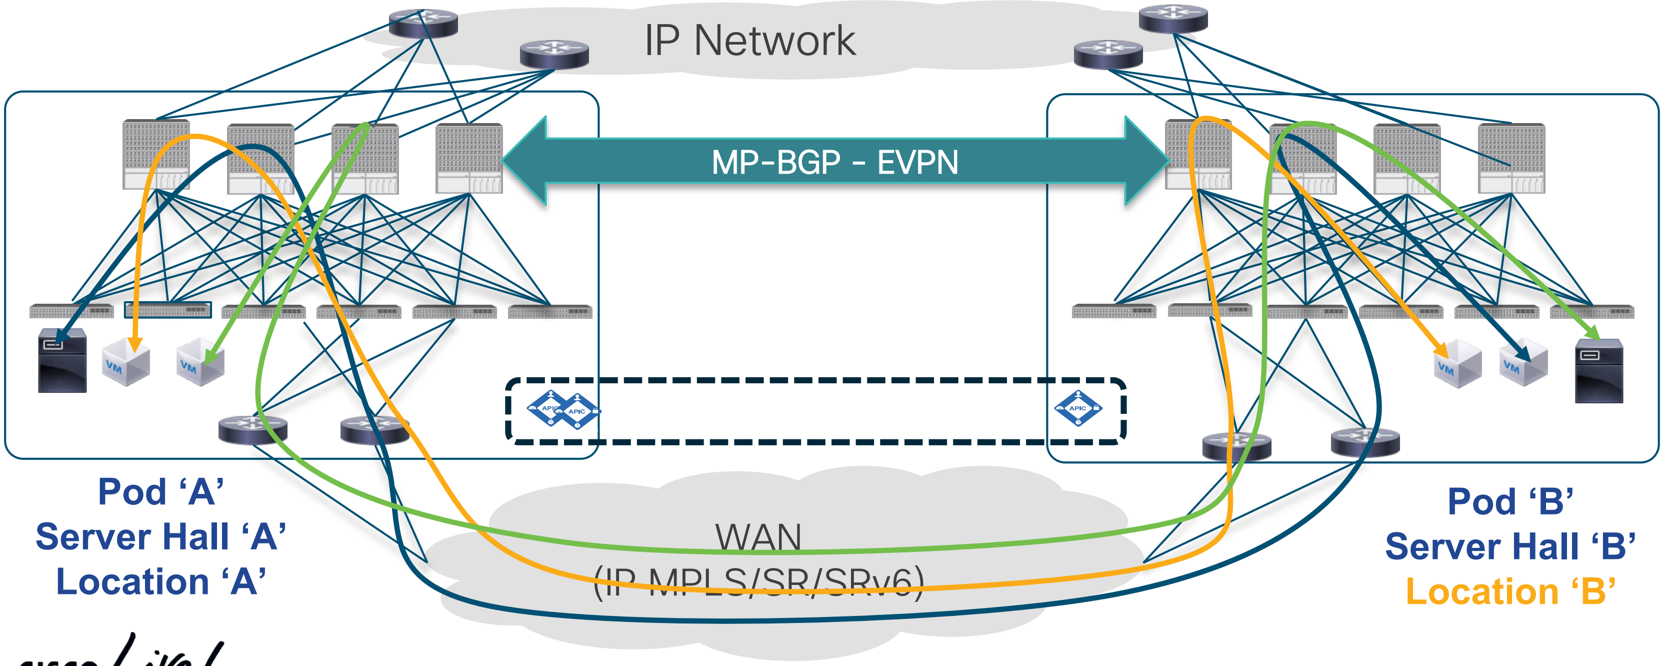

# Management of Multiple Sites with ACI Multi-Site Solution

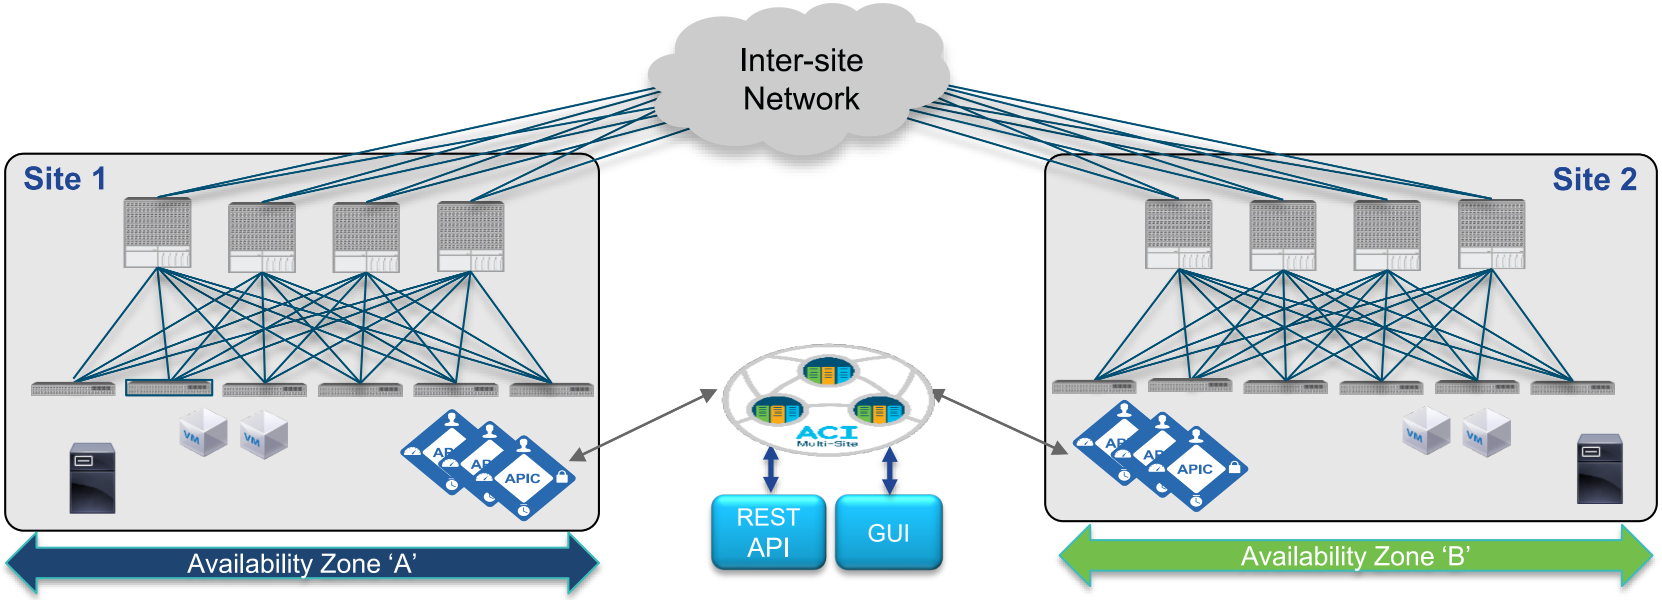


- Separate ACI Fabrics with independent APIC clusters
- ACI Multi-Site pushes cross-fabric configuration to multiple APIC clusters providing scoping of all configuration changes
- End-to-end policy definition and enforcement

# ACI Multisite with transport network slicing

Future

- ✓ Traffic between ACI site is forwarded through WAN
- ✓ No ISN and No need for VXLAN stretch between ACI Sites
- ✓ Multi-site Orchestrator is pushing policy configuration across sites along with day-2 operations

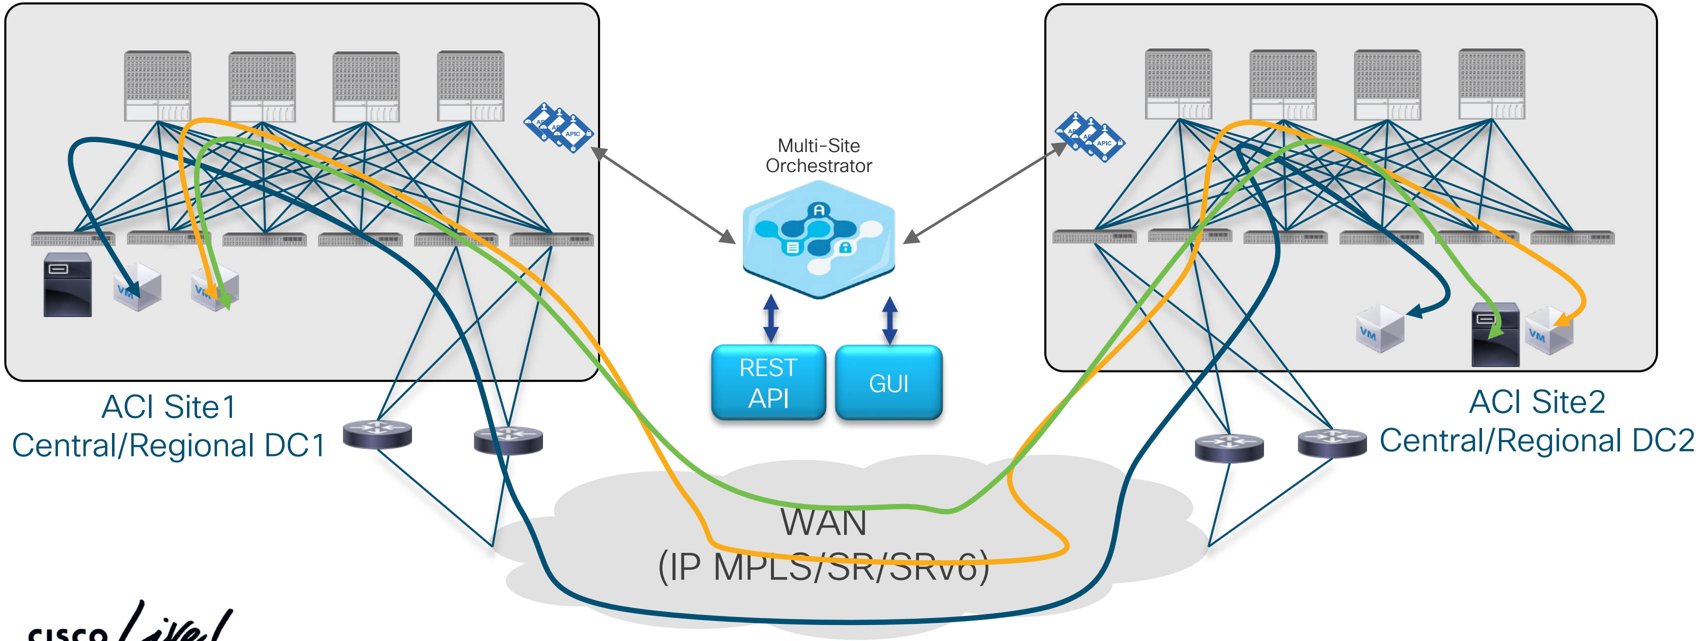

# Automation of Transport and DC using NSO

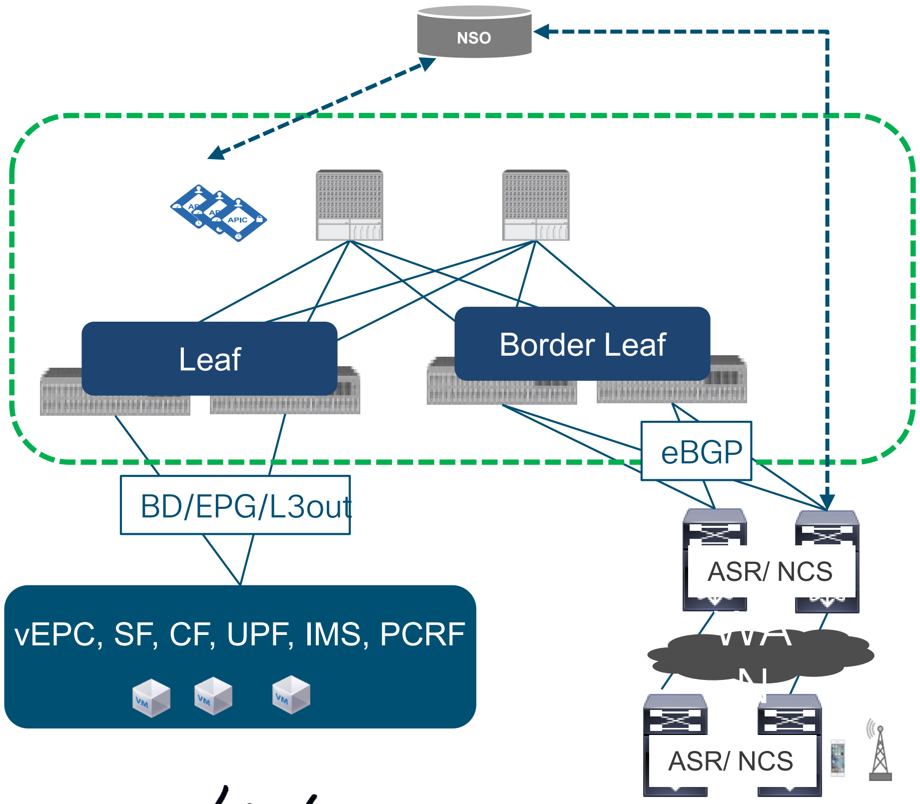


- NSO DC core function pack (DC-CFP) for DC to push policies on ACI
- NSO transport software defined network core function pack (T-SDN CFP) to push policies to transport devices
- Cross-domain integration between transport and DC to automate configuration such as VRFs, QOS/SR label mapping, route-target configs etc.

  

# Nexus Data broker for monitoring

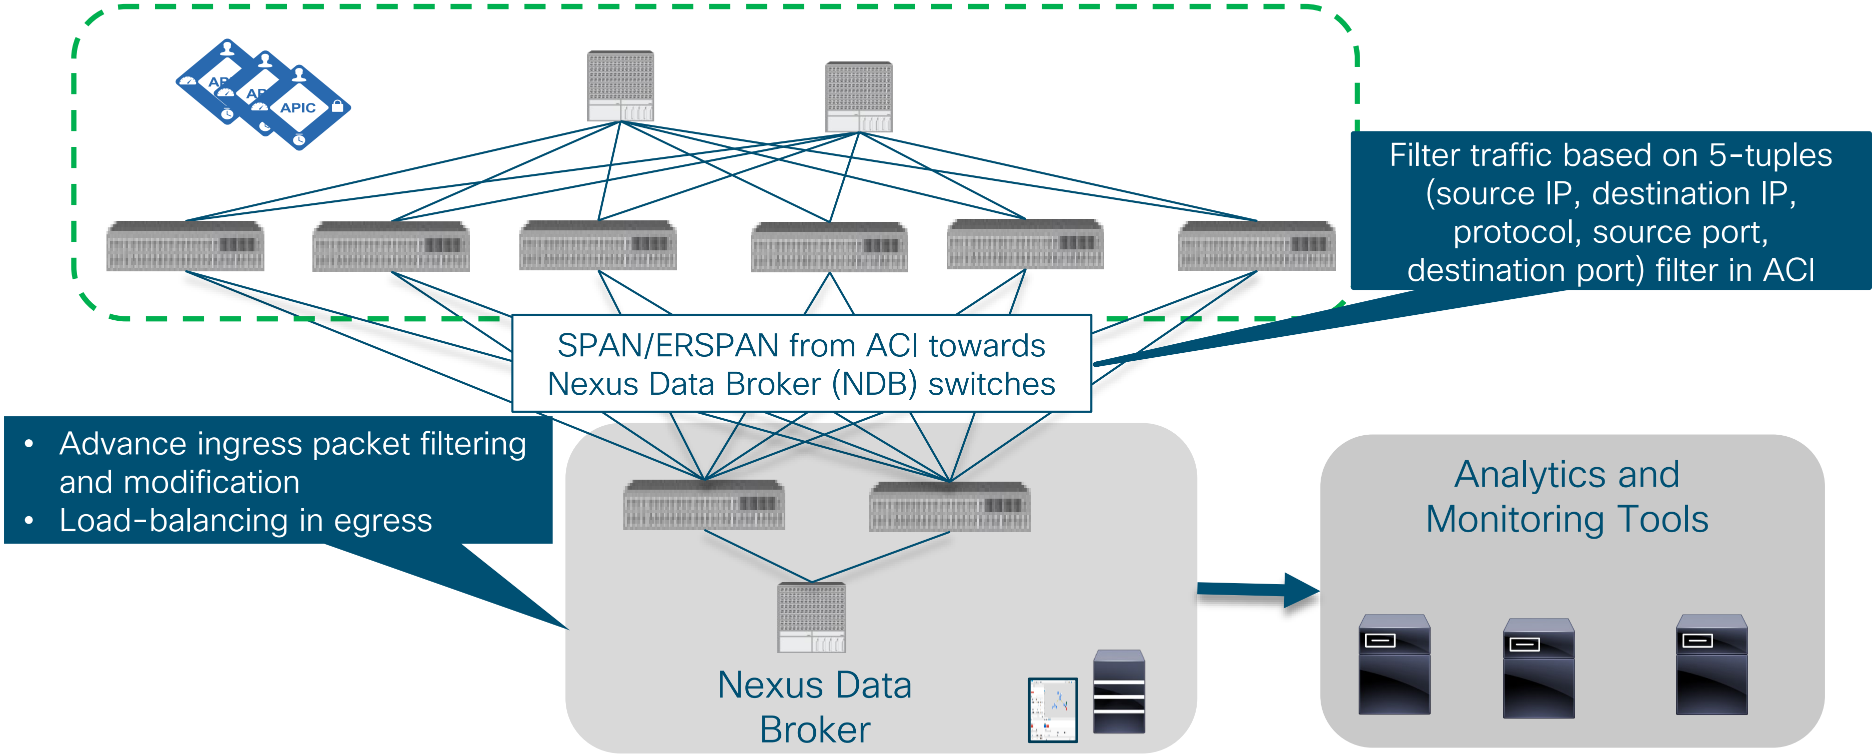


  

# ACI-SR Handoff at Border Leaf

- ACI to SR Interworking at ACI border leaf
- Remote Leaf can be border leaf and support same design
- ACI to SR MPLS Interworking with N9300-FX2
- APIC/MSC to WAN controller integration

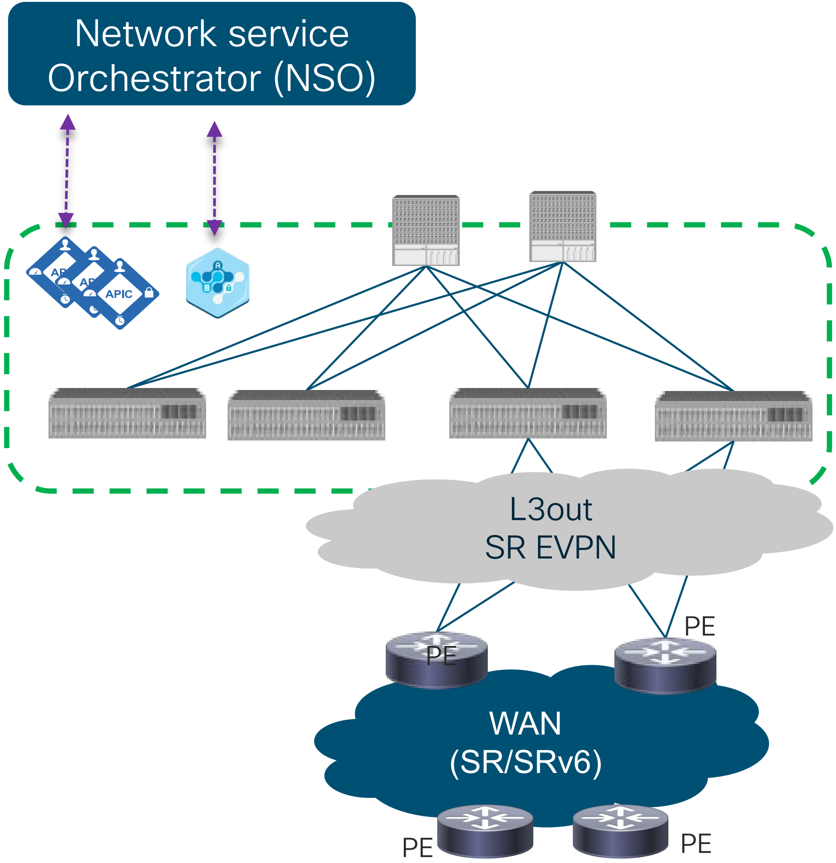

# Cisco Network Assurance and Insights

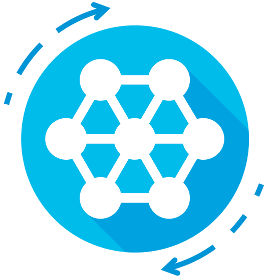

## Assure intent

"Ensure the business needs are consistently maintained"


## Guarantee Reliability

"Solve problems before they impact business"


## Troubleshoot intelligently

"Highlight the needle in the haystack"

# Network insights - Enable proactive action

## Increase availability, performance, and simplify operations

### Sources Of Telemetry Data

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

### Ingest And Process

- Metadata extraction
- Correlate against dBase
- Complex correlation

### Derive Insights

- Anomaly
- Flows
- Predictive failure
- Root cause

### Suggest Action

# Network Insights – Use Cases

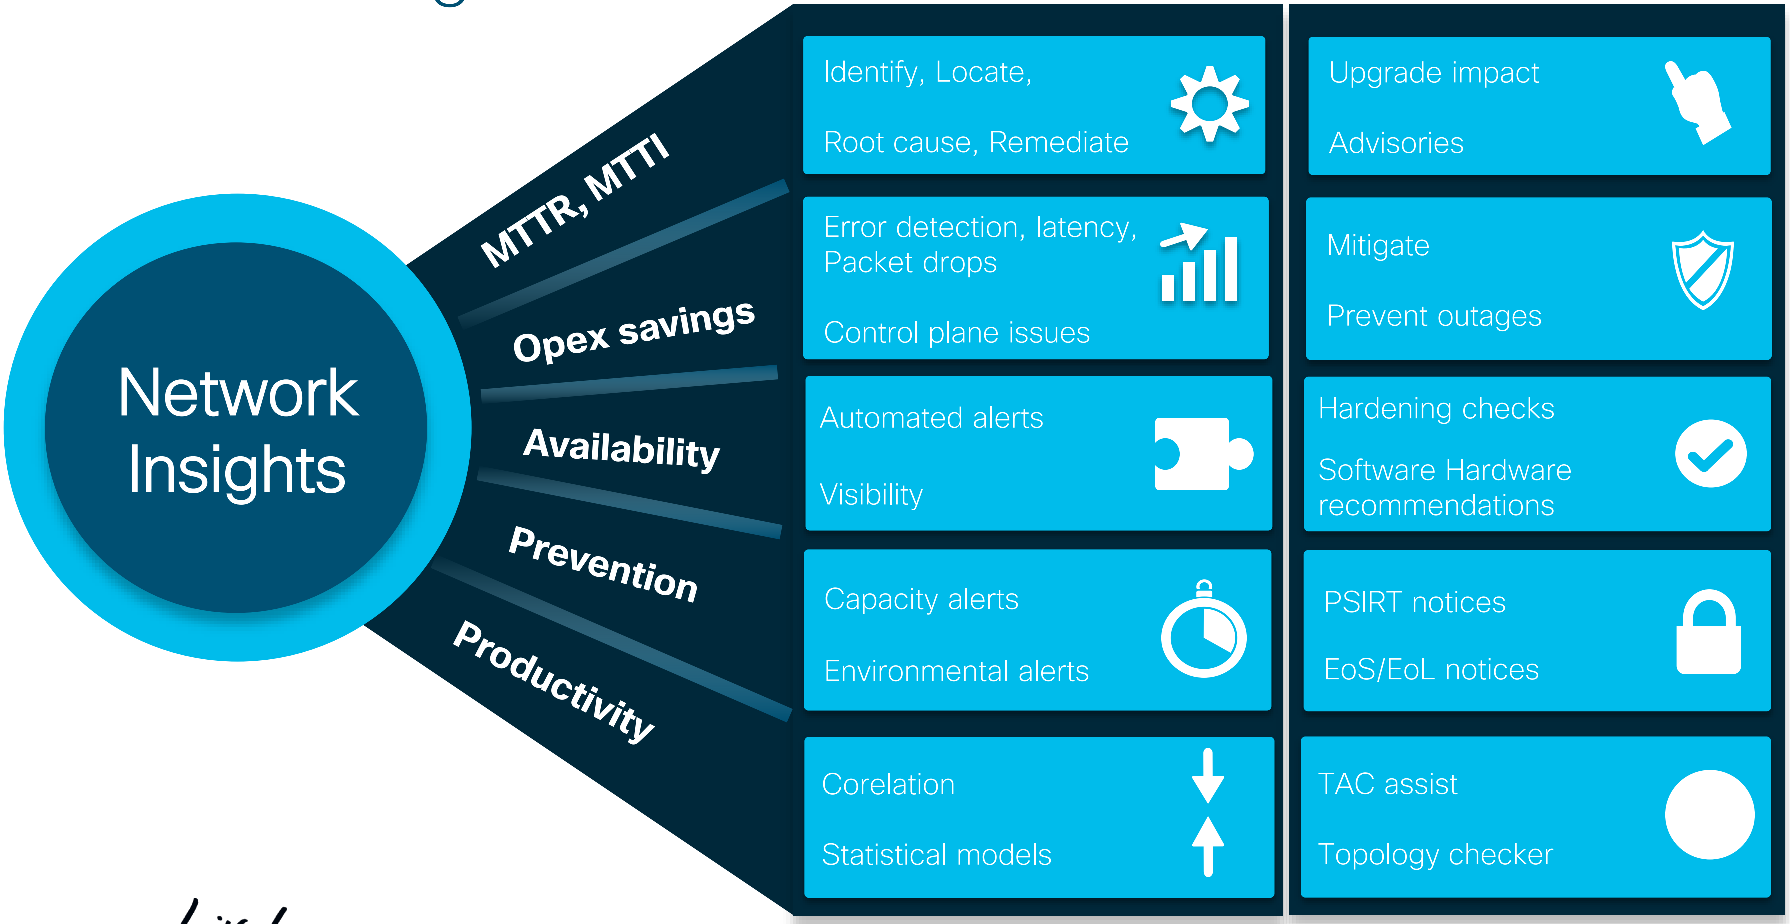


  

# Agenda

- Evolution of Telco Datacenter
- Customer Profiles & Technical Requirements
- How ACI met the requirements
- Looking Ahead
- **Conclusion**

# Conclusion

- Massive Scale
- Time to Market
- Operations
- Service chaining

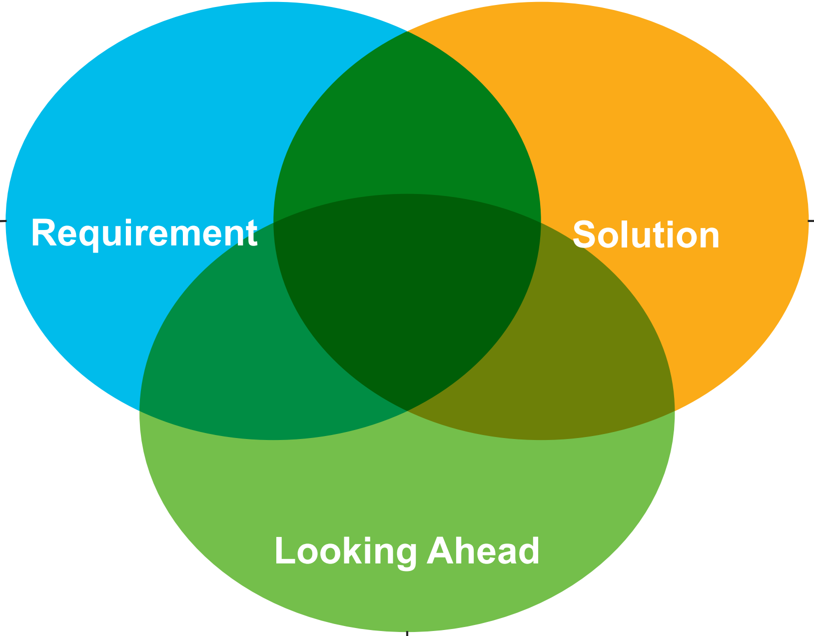


- Automation
- Distributed DC
- Scale
- Simplified Operation
- Integration with Tools

Consistent Policy & Management across Geography

Network Insight & Assurance

  

# Complete your online session survey

- Please complete your session survey after each session. Your feedback is very important.
- Complete a minimum of 4 session surveys and the Overall Conference survey (starting on Thursday) to receive your Cisco Live t-shirt.
- All surveys can be taken in the Cisco Events Mobile App or by logging in to the Content Catalog on ciscolive.com/emea.

Cisco Live sessions will be available for viewing on demand after the event at ciscolive.com.

  

# Cisco Webex Teams

## Questions?

Use Cisco Webex Teams 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 Webex Teams or go directly to the team space
4. Enter messages/questions in the team space

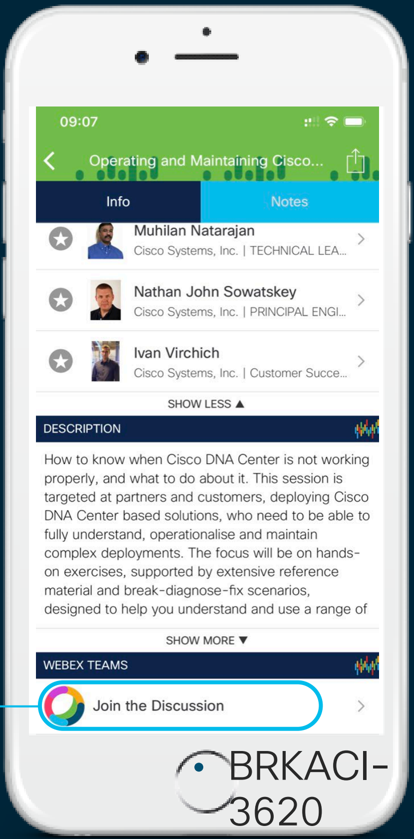

# Continue your education

- Demos in the Cisco Showcase
- Walk-In Labs
- Meet the Engineer 1:1 meetings
- Related sessions

# Thank you

You make **possible**