The bridge to possible

Convergência na Rede de Transporte para 5G

Rúben Fonte
Network Consulting Engineer
Dezembro 2021

Agenda

1. Requisitos para nova Converged SDN Transport Network
2. Converged SDN Transport
3. Segment Routing
4. Converged Access & Aggregation Portfolio
5. Conclusão

Session Takeaway

- Understand the needs for Converged SDN Transport
- Learn the Technology & Platforms to build it

Requisitos para nova Converged SDN Transport Network

Mobile Operator Revenue Growth Opportunities

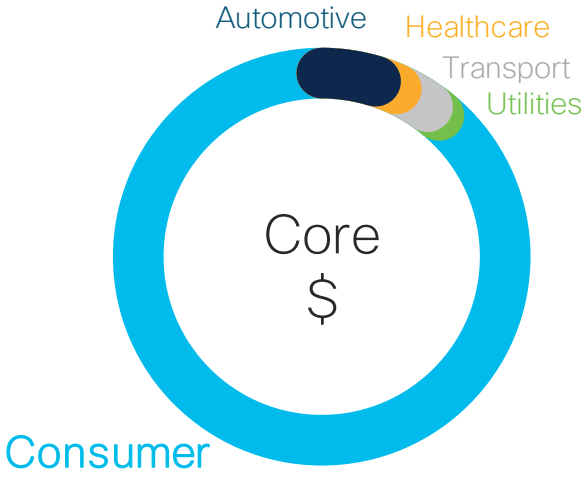

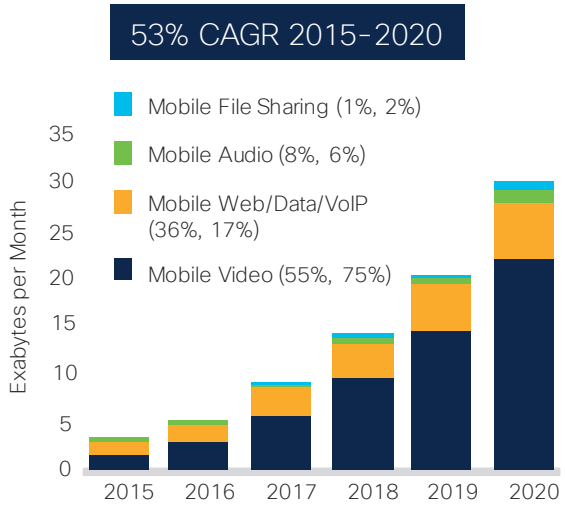

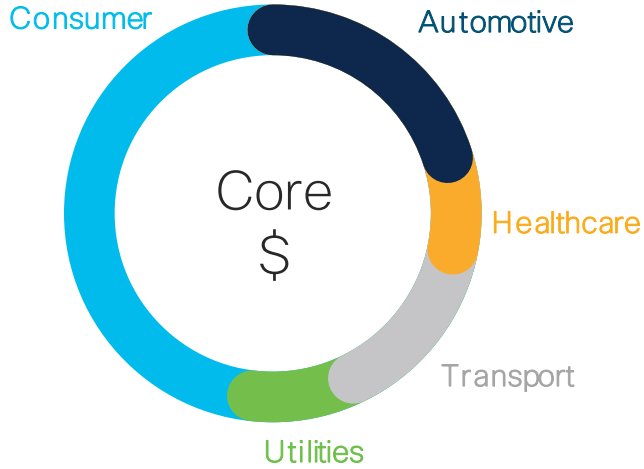

Today
Operator business mostly focused on the saturated consumer market

Video drives traffic… but not revenue

2025
Vertically targeted services will accelerate operator business growth

Source: European Commission Report, 2016

Evolving Transport Requirements

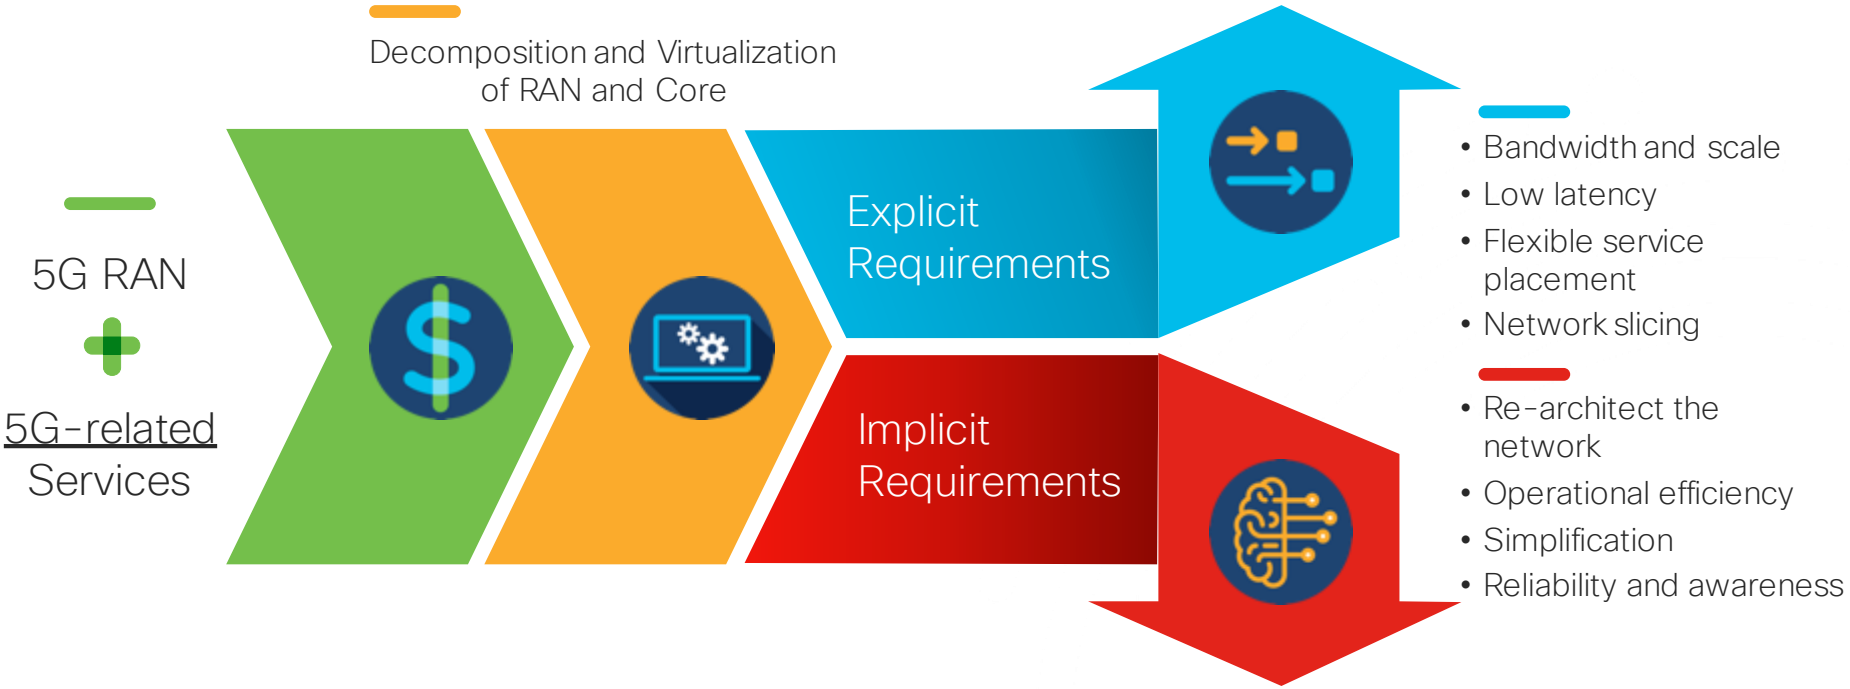

5G – Key Use Case Categories

Enhanced Mobile Broadband (inc. Fixed Wireless Access)

- Extra capacity delivered through new 5G frequency bands
- Not too concerned with connection density or latency

→ Increased bandwidth and capacity

Massive Machine-type Communication

- Focused on low power wide area NB-IoT with high connection density and energy efficiency

→ Scale, Reliability

Ultra-reliable, Low Latency Communication

- For mission critical use cases (self driving, Public safety, …)
- 1-25 msec latency

→ Push data plane to the edge, intelligence in network

Source: Recommendation ITU-R M.2083

Emerging – Low Latency

- Low latency applications, entertainment

→ Push data plane to the edge, Intelligence in Network

RAN Transport Changes

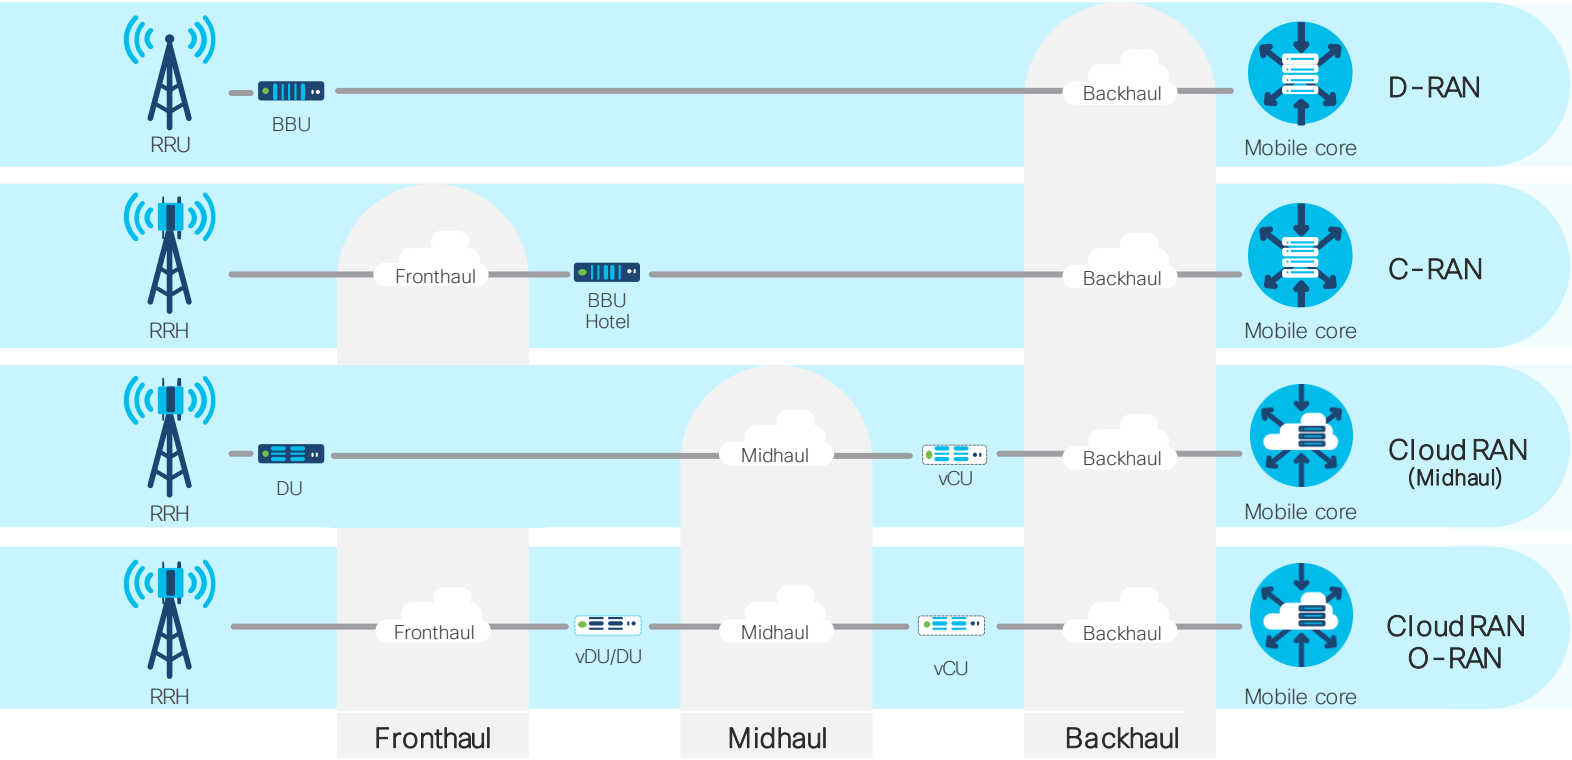

- Higher Speed Interfaces
- Lower Latency
- More Precise Timing & Synchronization
- Any-to-Any Connectivity

Benefits of Packet-Based Fronthaul vs xDWDM

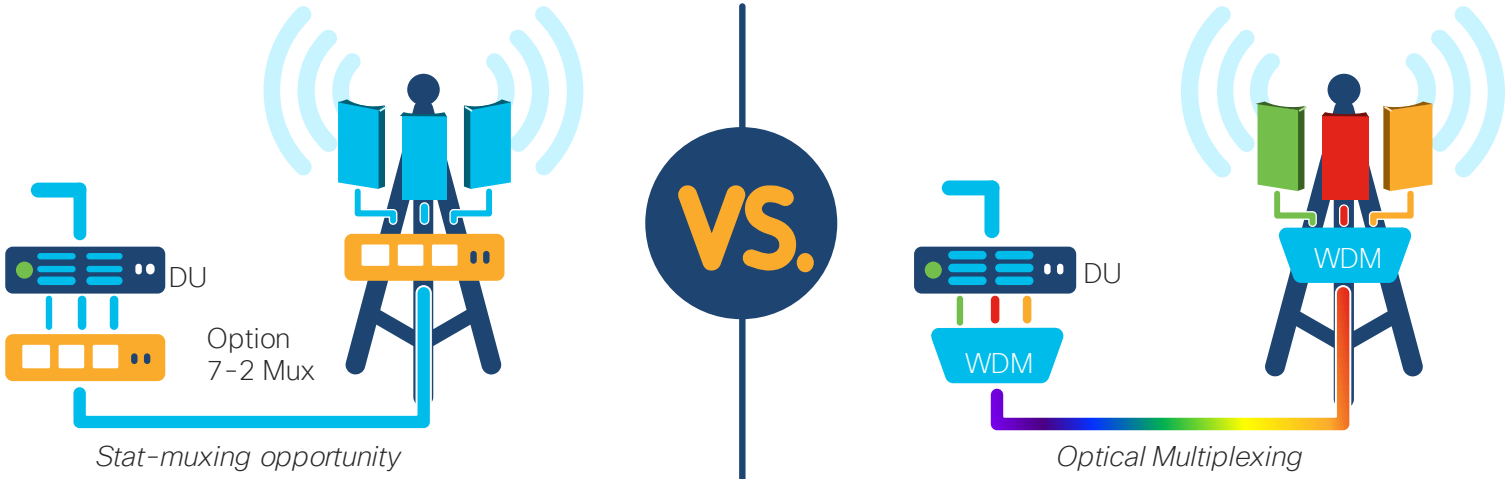

- ✔ Stat Mux Advantages
- ✔ Cost Effective
- ✔ Topology Independent
- ✔ Service Visibility & Transparency
- ✔ Scalable E2E Converged IP
- ✖ Optical multiplexing
- ✖ Non-scalable, architecturally rigid
- ✖ Point-to-point, topology dependent
- ✖ Limited service visibility
- ✖ Capex dependent scale

Comparing TCO for fronthaul

Packet vs optical fronthaul solutions

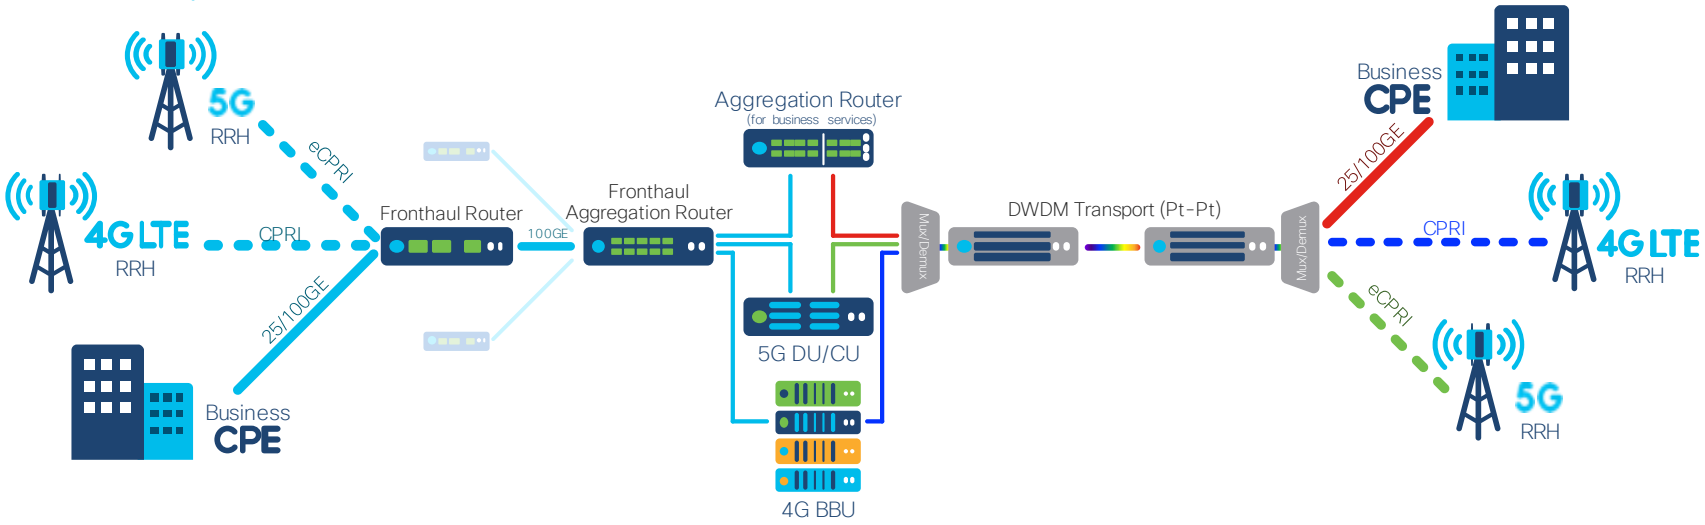

Packet-based fronthaul

Up to **65%** TCO Savings

Optical-based fronthaul

Savings								
+65%	ROADM	+46%	P2P Active DWDM	+40%	P2P Passive DWDM	+13%	P2P Passive CWDM

Source: ACG – An Economic Comparison of Fronthaul Architectures for 5G Networks

Requisitos para nova Converged SDN Transport Network Resumo

- Traffic Increase Transport
- 5G requirements
- Transport architecture challenges
- RAN architecture changes

Converged SDN Transport

Main Goals for New Architecture

- Converge multiple service on a single network infrastructure
- Simplify the infrastructure scaling to hundreds of thousands of nodes
- Support end to end constraint/SLA based services
- Flexible service placement with intelligence at the Edge
- Any to any connectivity between service endpoints
- End to end network resiliency and active/active services

Convergence Efficiency

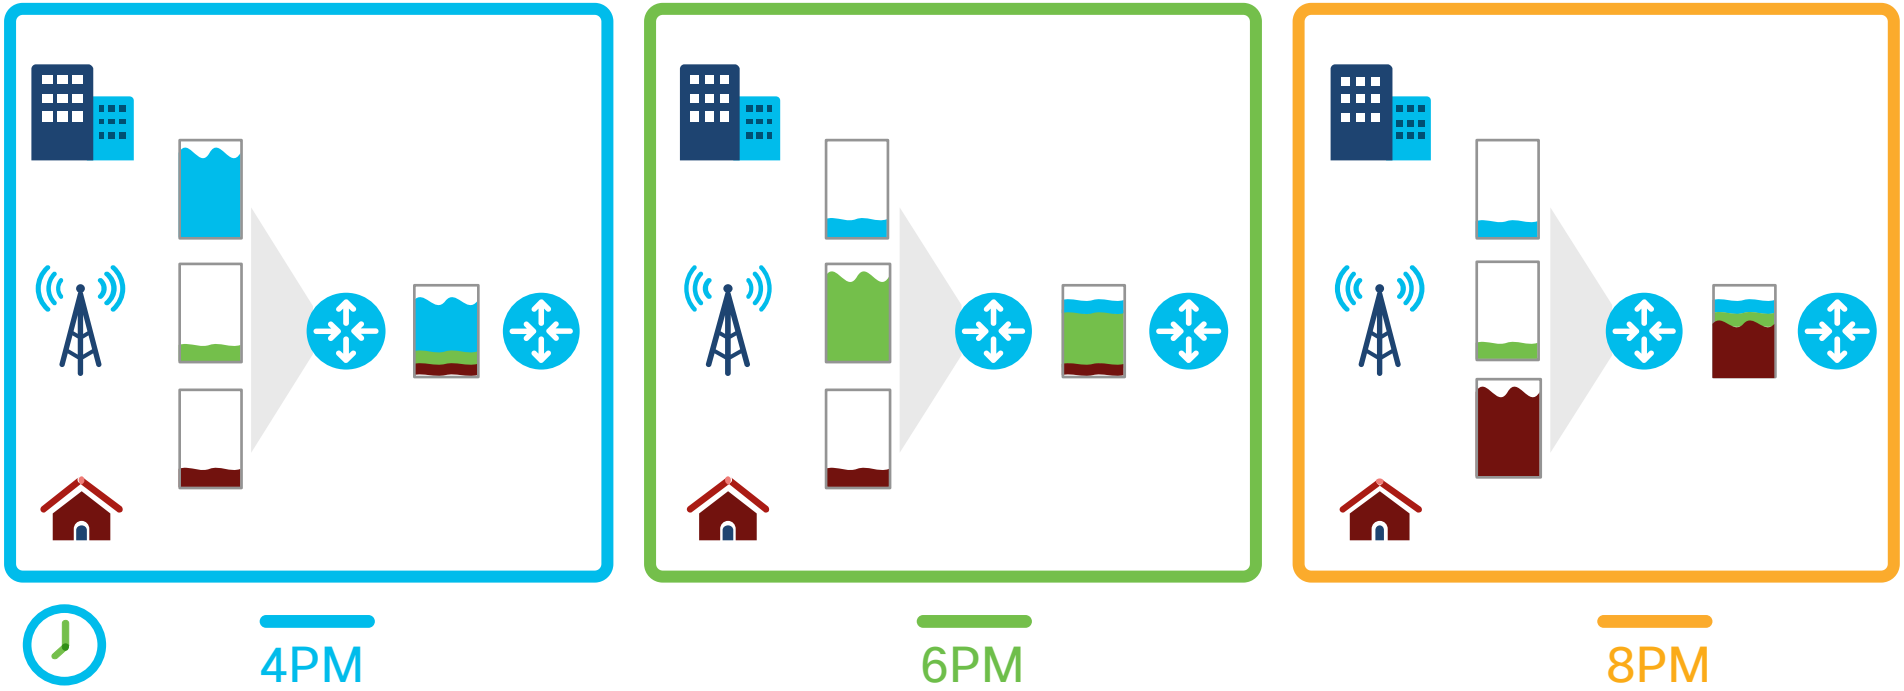

Slicing on Stat-Mux

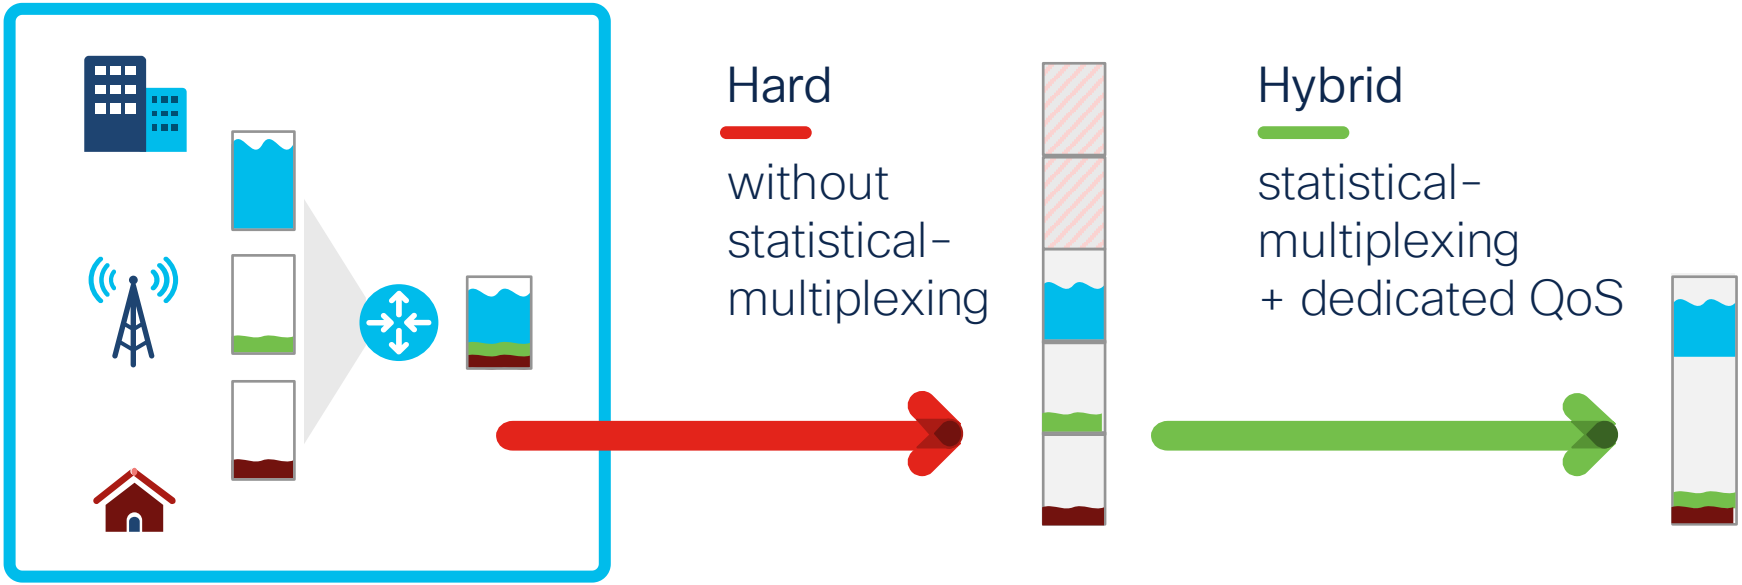

Transport Convergence

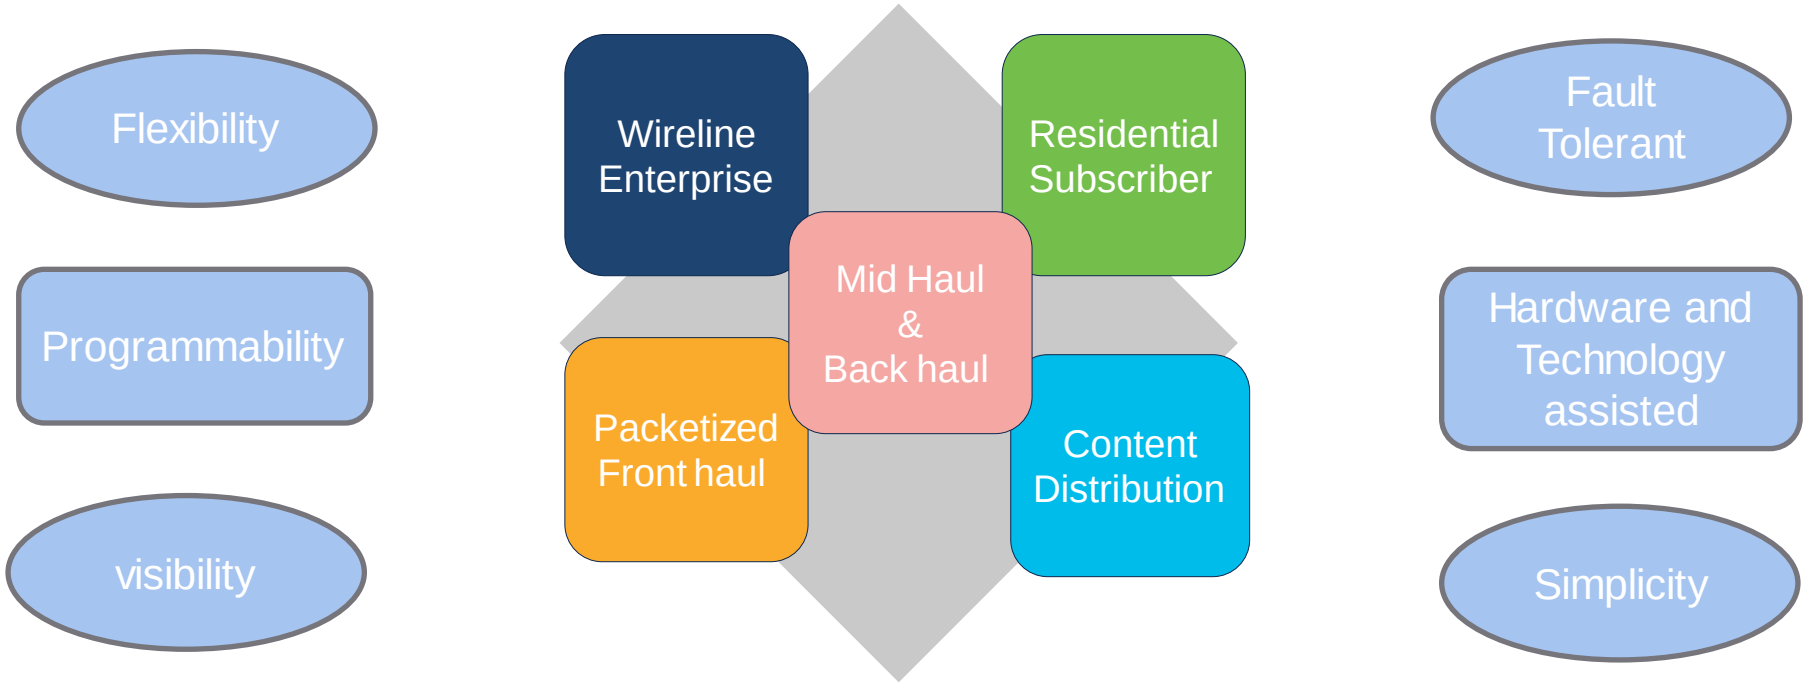

Differentiated Services
over Converged Transport

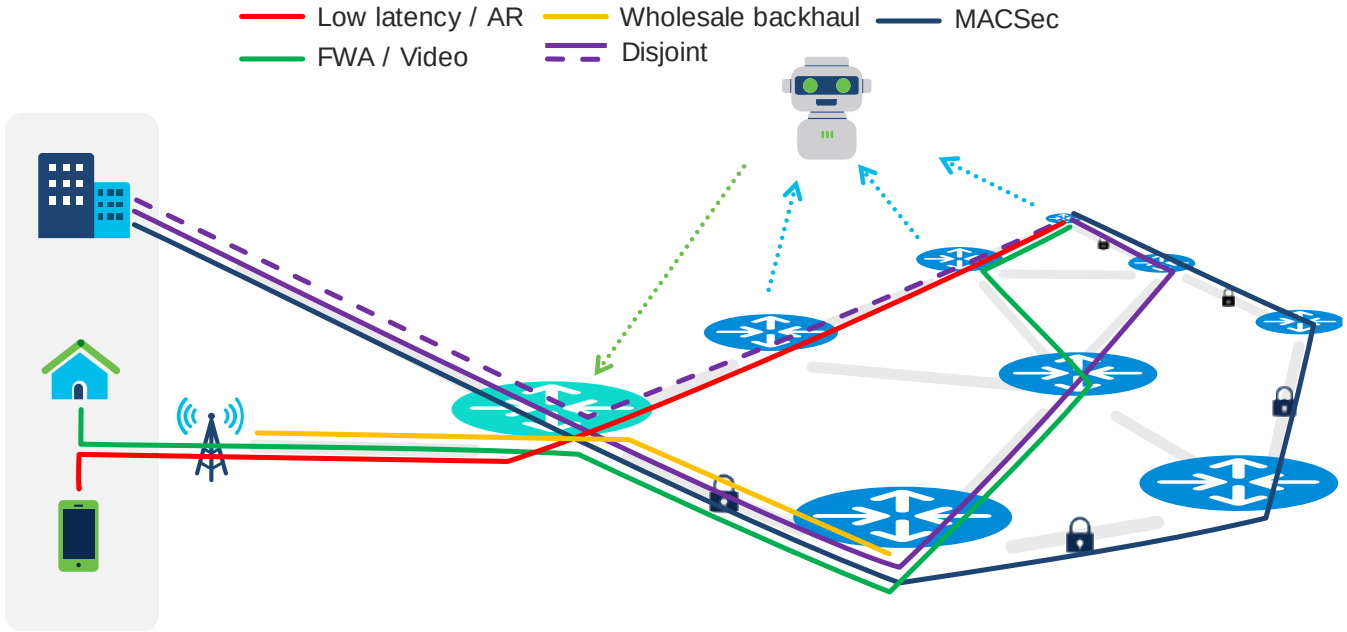

Intelligence in the Access

- Multi-service
- Network Slicing
- Synchronization
- Integrated Automation
- Telemetry
- Fronthaul Gateway
- Security

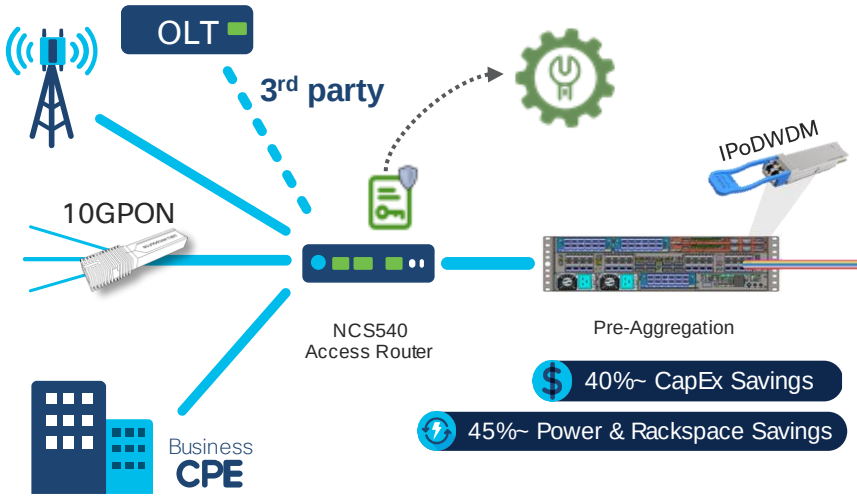

Benefits of Converging to IP

- Scalable any-to-any service connectivity
- Efficient load-sharing using statistical multiplexing
- Simplification by reducing the:
  - number of protocols and network touch-points
  - stitching points in service and transport paths
  - service-related state in the transport
- Consistent OAM and traffic flow control (Traffic Engineering)

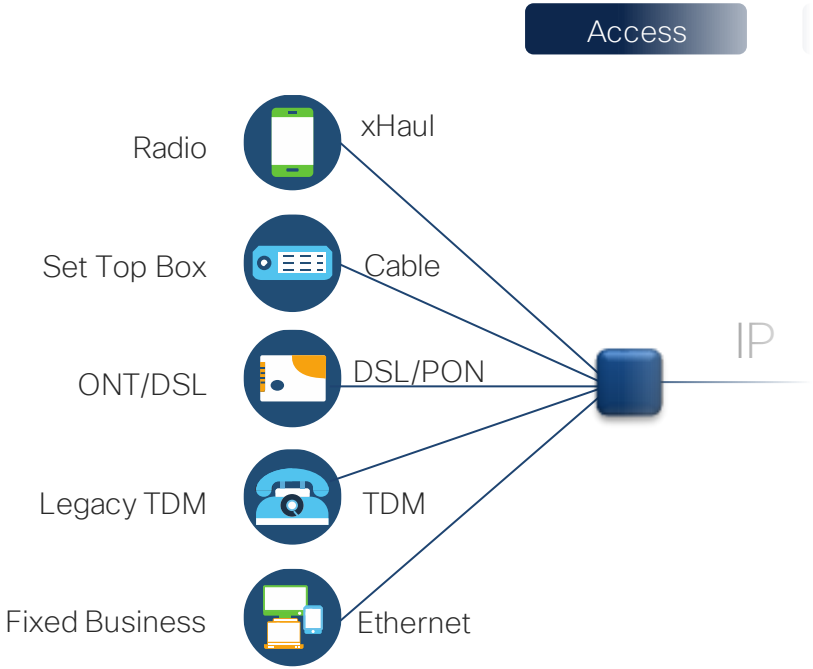

Converged SDN Transport Toolkit

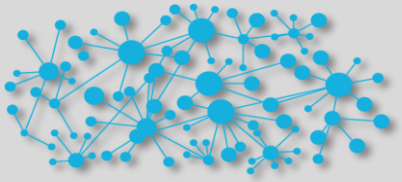

Mixed SLA Services

Unified Fabric – SR/ BGPVPN

Network Slicing & Virtualization

Automation & Programmability

Timing & Latency

Security & Visibility

Converged SDN Transport – End to End

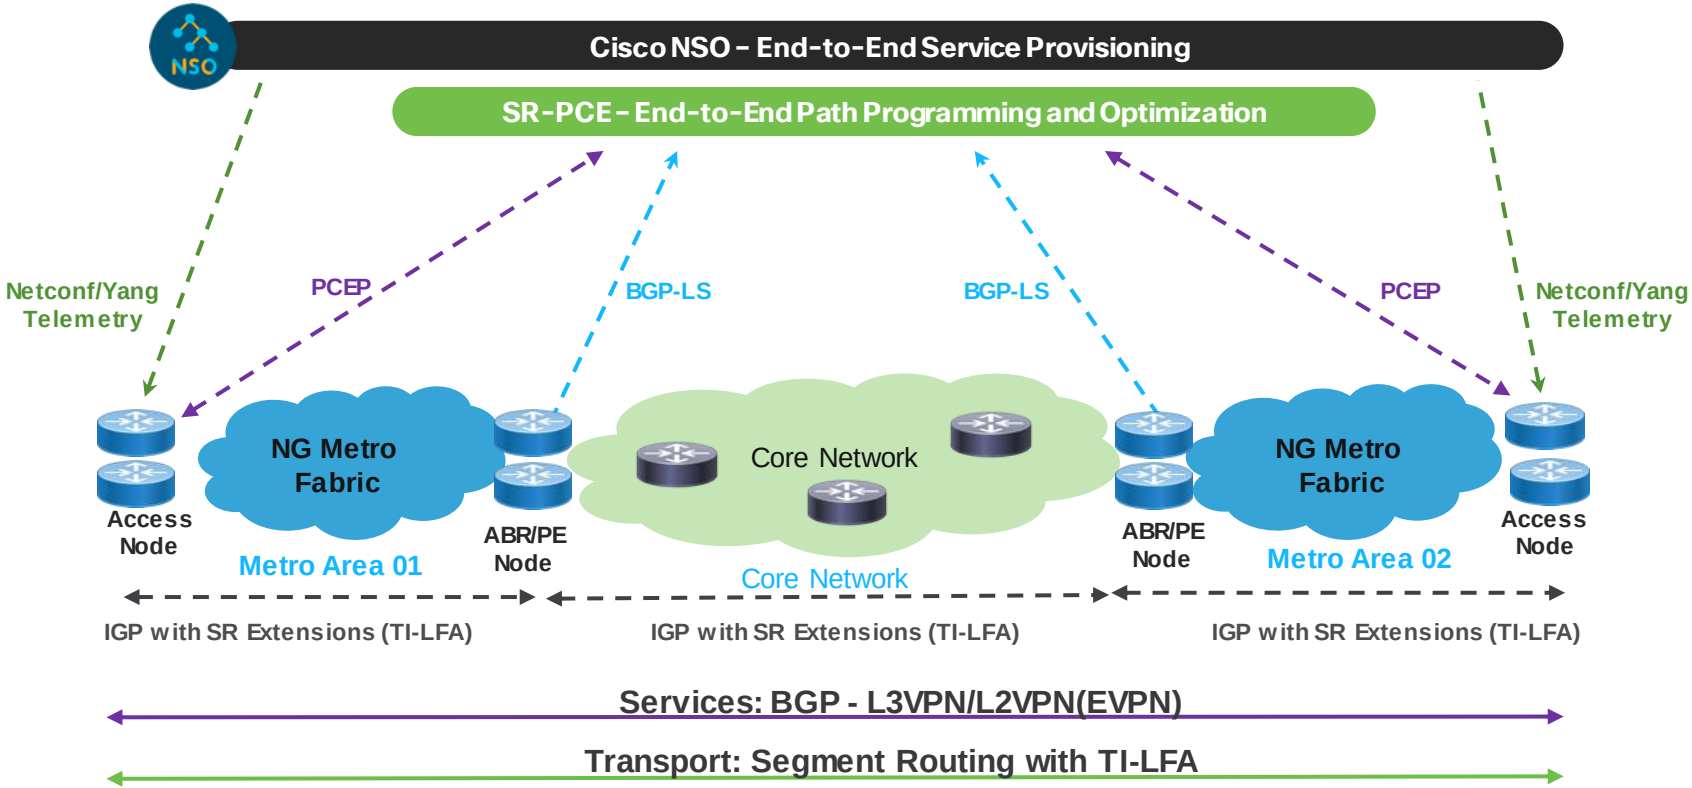

IP Transport Foundations “Done Right”

high-performance,
5G-ready portfolio

trusted
infrastructure

cross domain
automation

software-driven
simplification

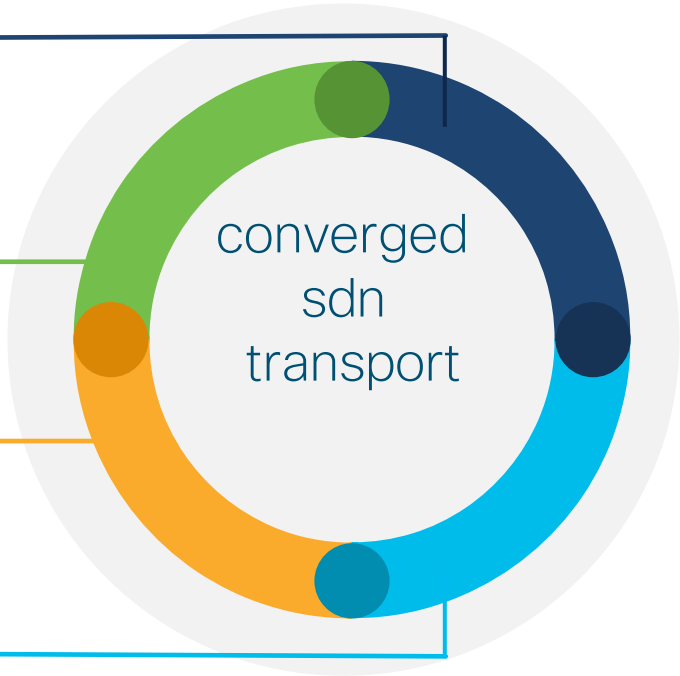

Lowering TCO by 62%

60% improved
capital efficiency

66% better
OPEX utilization

81% faster
time-to-service

Converged SDN Network Resumo

- Access Intelligence Changes
- Convergence requirements
- Transport Toolkit
- Architecture End to End

Segment Routing

Problem statement: Today's service creation

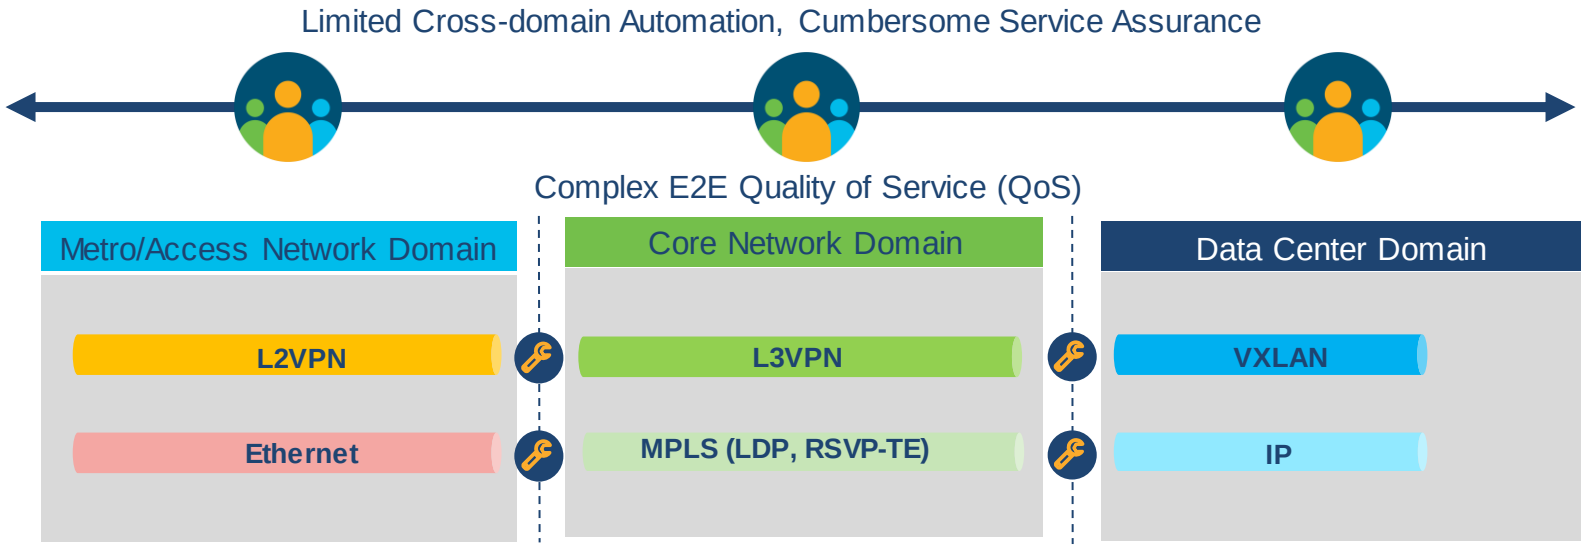

End-to-end service provisioning is lengthy and complex

- Routing protocols are blind to applications (Shortest-path for any traffic)
- Multiple network domains under different management teams
- Manual operations
- Heterogeneous underlay and overlay networks

Network Simplification with Converged Transport

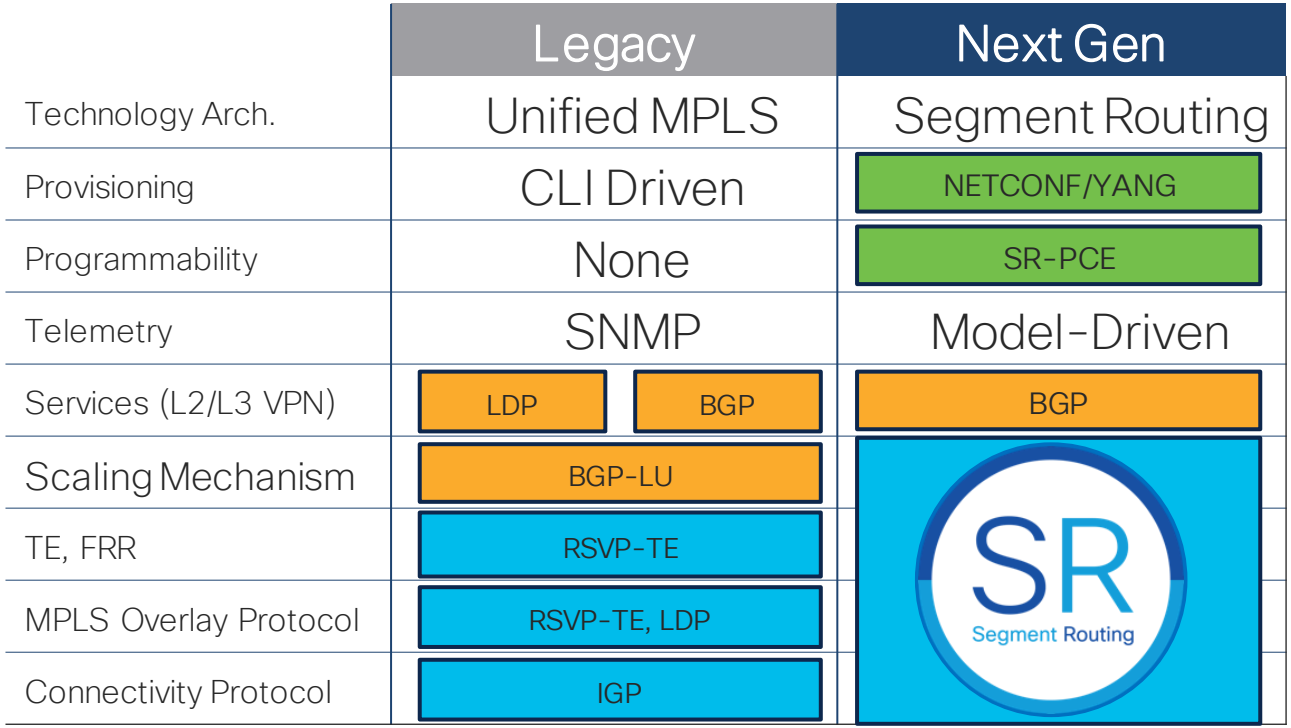

- End-to-End Segment Routing
  - IPv4 SR-MPLS
  - SRv6
- Programmable Transport
  - Utilizes Segment-Routing PCE
  - PCE eliminates need for BGP-LU by using SR On-Demand Next-Hop
- BGP Services Control Plane
  - EVPN for L2 Services
  - L3VPN for L3 Services

26

SP Networking – Simplification Journey

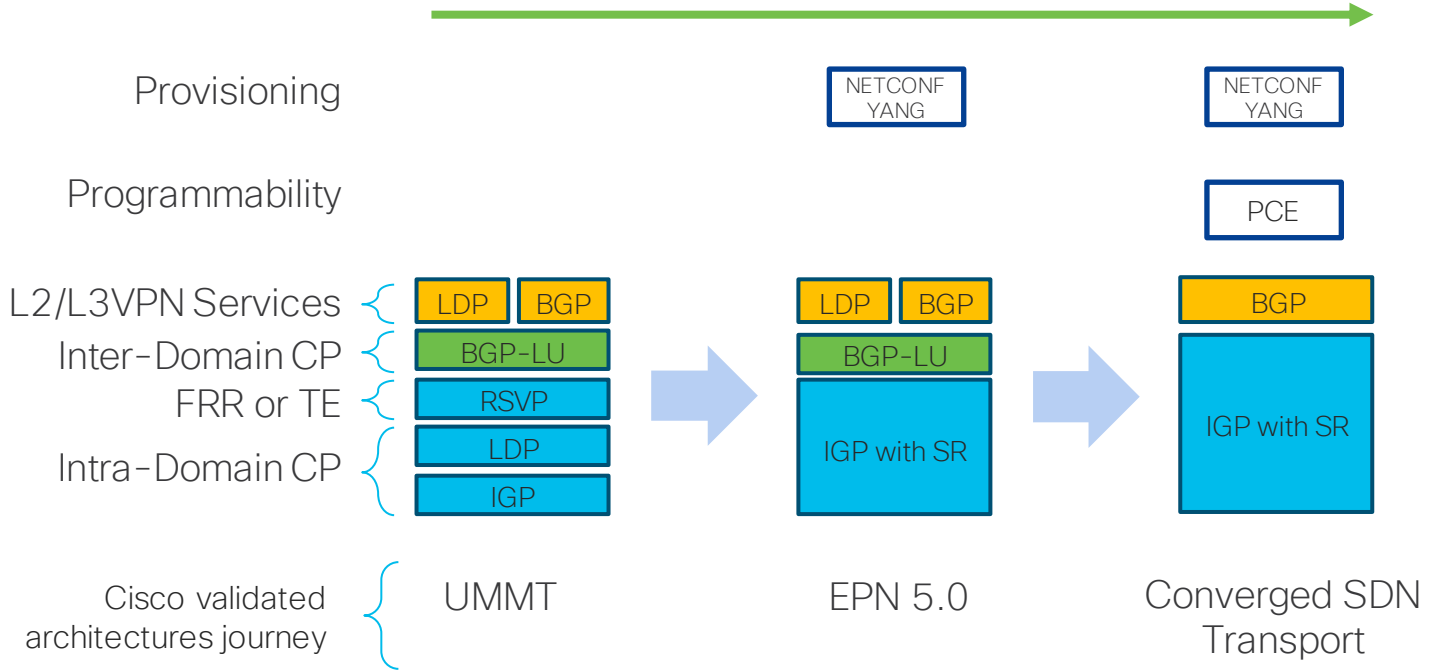

Why Segment Routing – MPLS

The 2 faces of segment routing

A LS IGP protocol extension bringing network simplification/optimization

- Eliminates LDP and RSVP-TE
  - Simple, less state and adjacencies
- No IGP to LDP synchronization
  - Eliminates delays in activating a path
- Topology independent FRR using post convergence back up path
  - 50 ms protection
  - Microloop avoidance
  - 100% coverage of network topologies

  Easy troubleshooting

An IP/MPLS architecture designed with SDN in mind

- Right balance between distributed intelligence and centralized optimization and programming
  - SR-TE
- Applicability across all network types
  - (SP, OTT/Web, GET) across (WAN, Metro/Agg, DC)
  - MPLS and IPv6 dataplanes
  - SDN controller

Segment Routing (SR)

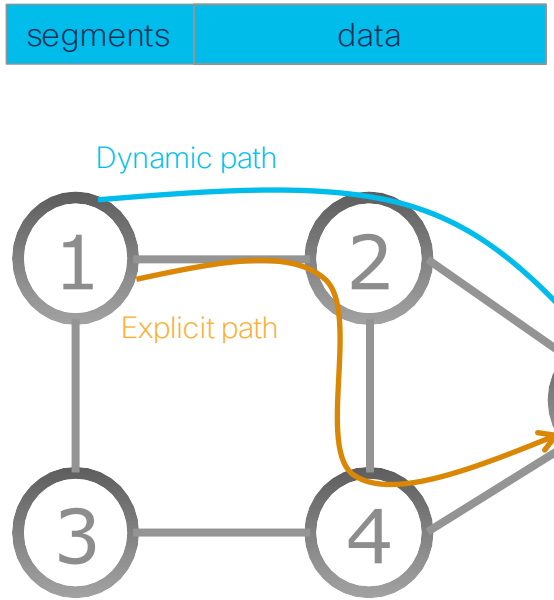

Data Plane

MPLS (segments)

IPv6

Control Plane

Routing protocols with SR extension (ISIS , OSPF, BGP)

SDN Controllers (BGP, PCEP, NC/YANG)

Path Construction

Dynamic (SPF computation)

Explicit (expressed in list)

SR Segments

IGP Prefix Segment

- Signalled by ISIS/OSPF
- Minor extensions to the existing link-state routing protocols (OSPF & IS-IS)
- Shortest-path to the IGP prefix
- Globally unique in SR domain
  - SRGB+ Index => 16000+5 = 16005

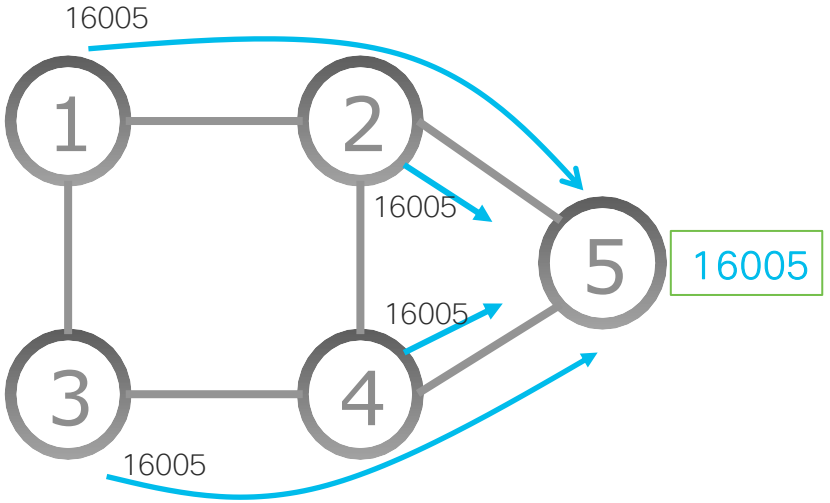

IGP Adjacency Segment

- Signalled by ISIS/OSPF
- Minor extensions to the existing link-state routing protocols (OSPF & IS-IS)
- Forwarded on IGP adjacency
- Local scope

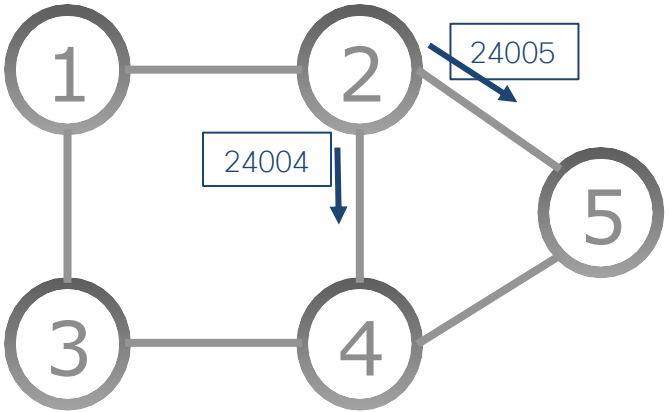

Combining IGP Segment
SR Policy

- Steer traffic on any path through the network
- Path is specified by list of segments in packet header, a stack of labels
- No path is signalled
- No per-flow state is created

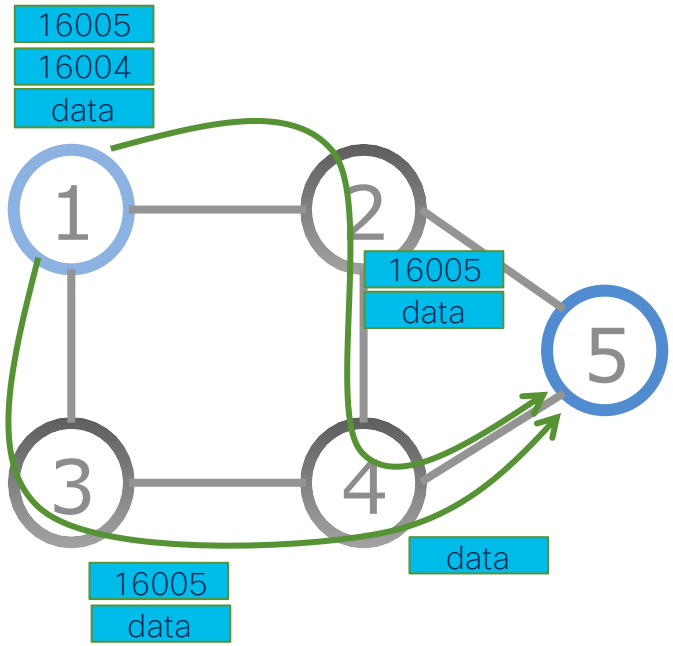

SR Policy Identification & Steering

- An SR Policy is uniquely identified by a tuple (head-end, color, end-point)
- At a given head-end, an SR Policy is uniquely identified by a tuple (color, end-point)
- Policy Color
  - BGP advertises as ext. community
  - Defines certain treatment
- Multiple candidate paths
  - explicit or dynamic

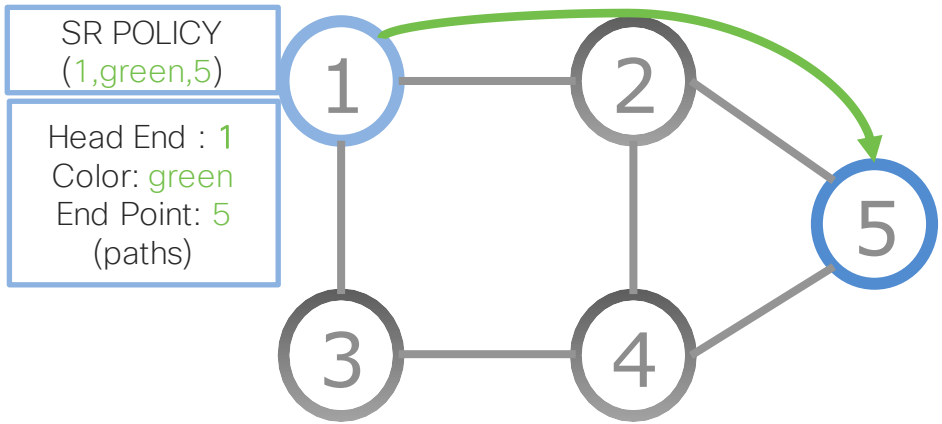

Programmable SDN Transport: SR PCE

- IOS XR-powered stateful PCE
  - Physical or virtual deployment
- Multi-domain topology collection
  - Real-time reactive feed via BGP-LS
  - Inter-area and Inter-domain with constraints and Anycast redundancy
- Computation
  - Shortest, Disjoint, Low Latency, SRLG avoidance
- SR-PCE is fundamentally distributed
  - Not a single all-overseeing entity, but distributed across the network; RR-like deployment
- Northbound APIs for fully programmatic operations

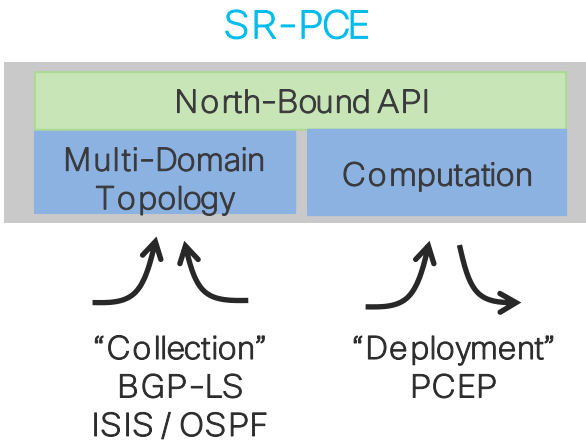

SR On Demand Next (ODN) Hop

- Head End requests PCE for colored path
- End point advertises color using BGP
- Path Initiates when certain end point advertises the color

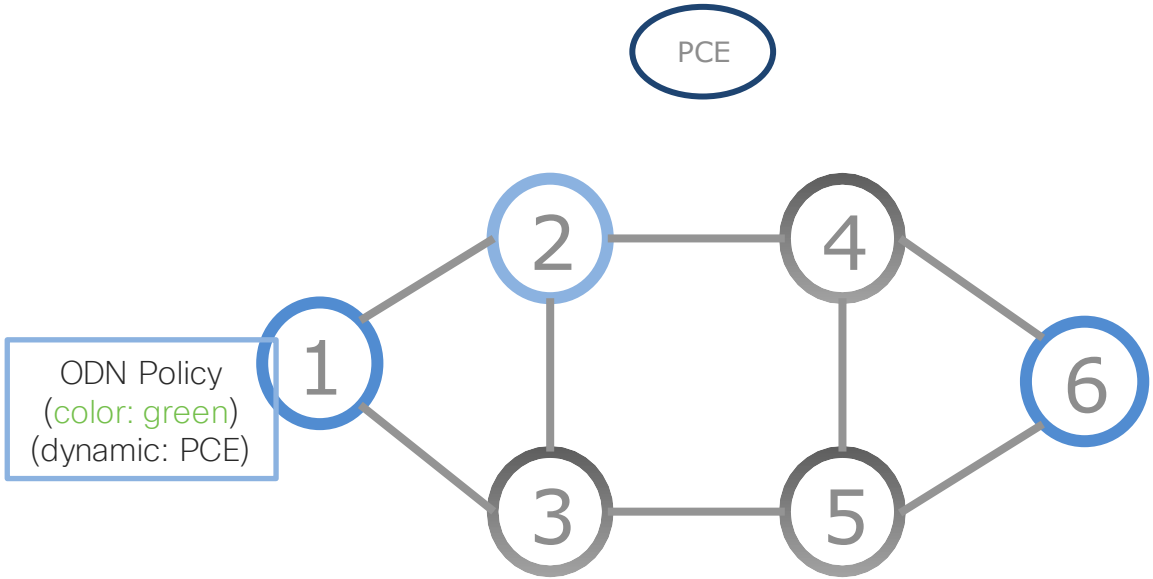

SR On Demand Next (ODN) auto steering

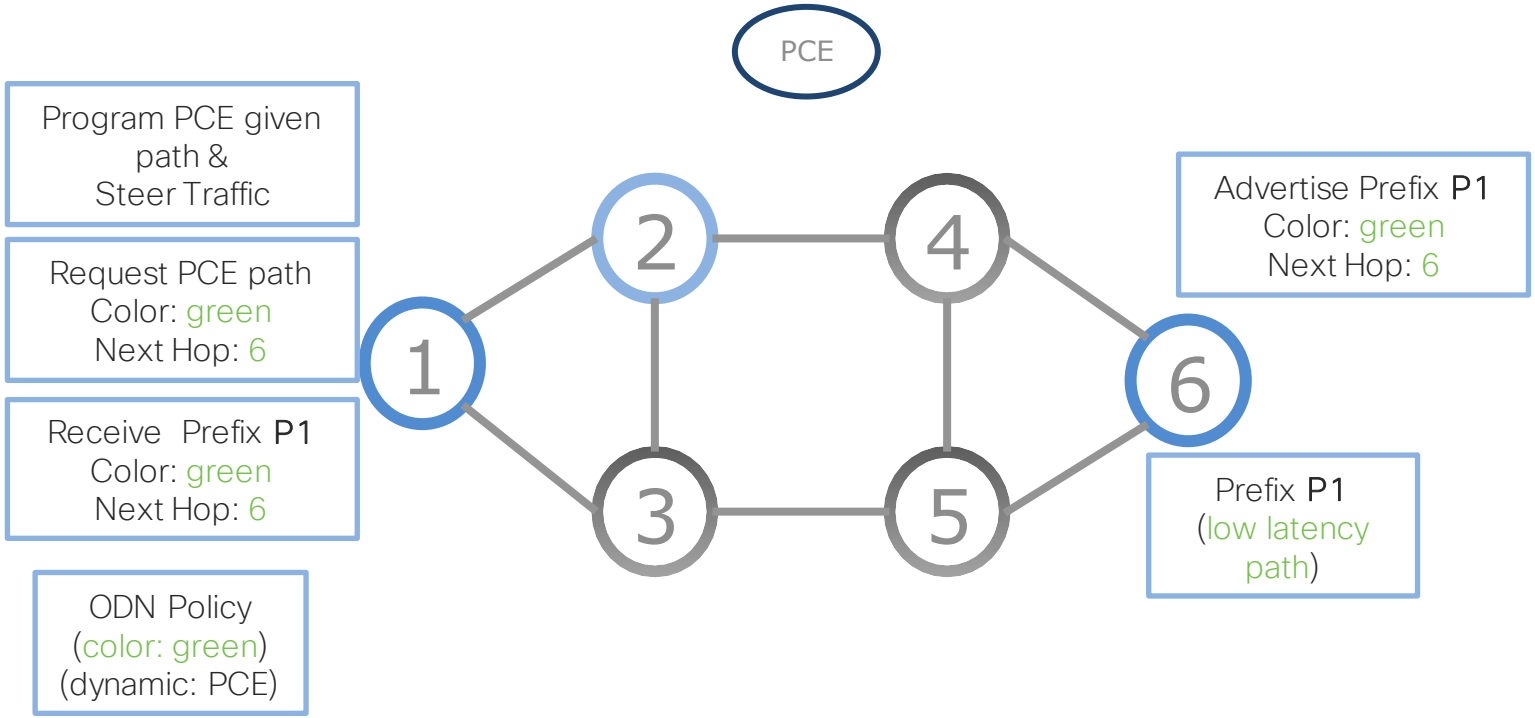

Inter-Domain SR-TE Policy via SR Policy

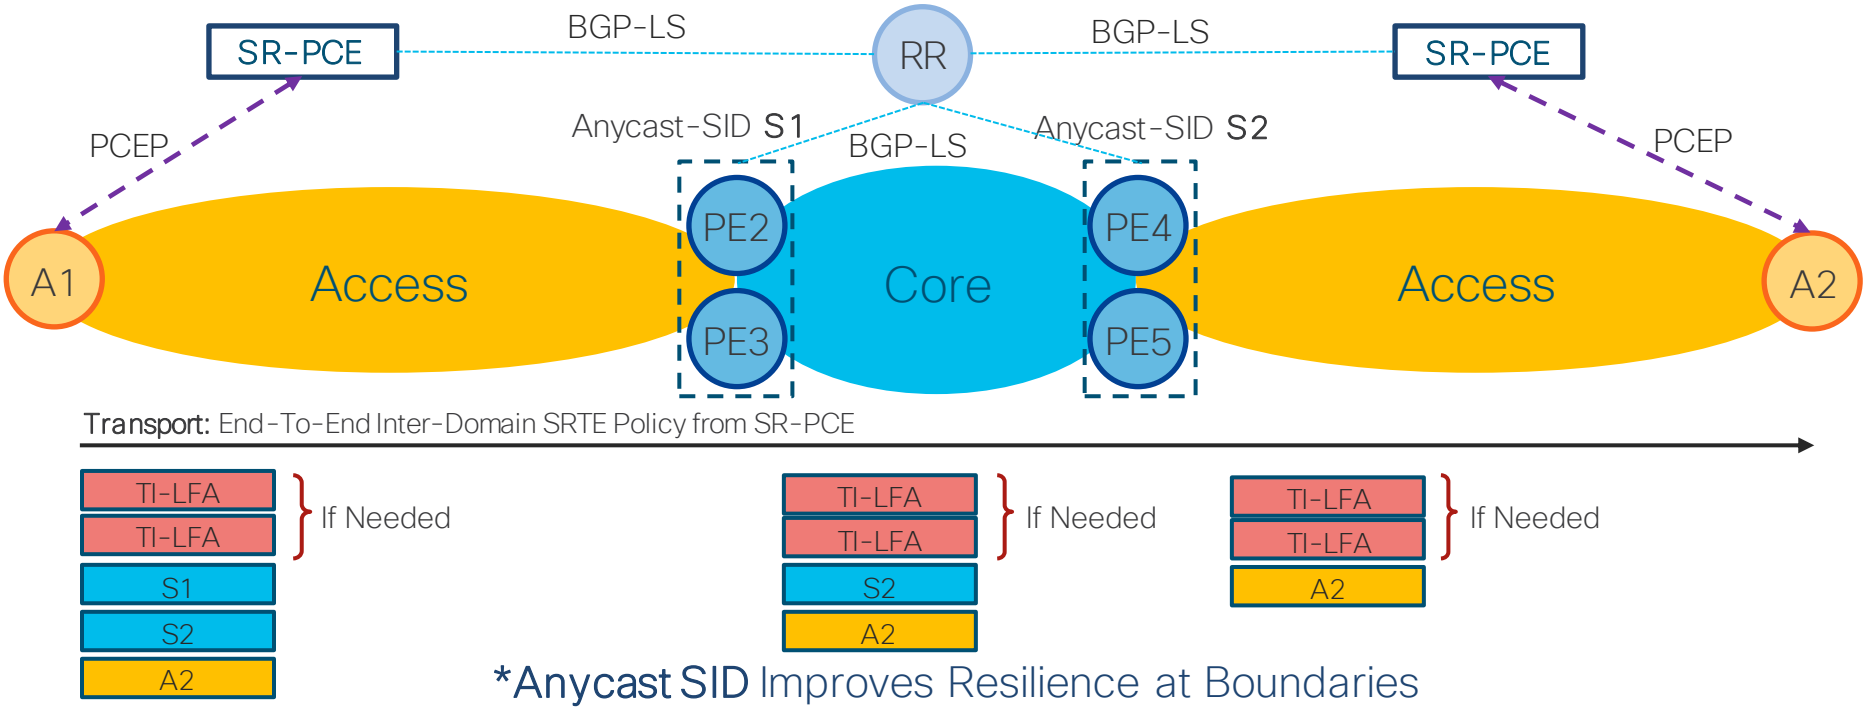

***Anycast SID** Improves Resilience at Boundaries

SR-PCE computes Inter-Domain paths with Anycast-SID

Performance Monitoring

PM Toolkit

- **Delay, Loss, Liveness Monitoring**
  - Per-Link Measurement
  - End-to-End (SR Policy) Measurement
- **Advertisement** of Link Metrics **in IGP**
- Standard Probe Packet **Encoding**
- **Intent-based** Network & Service **Monitoring**

https://datatracker.ietf.org/doc/html/rfc6374

Slicing: SR IGP Flexible Algorithms

Applicability Examples

Sensitive Data

Financial Transactions

Transport
Redundancy

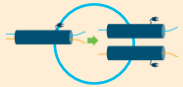

Different Fiber Conduits

Scalability

Low-End Platforms

Solution

Customized IGP algorithms defined by operator for **intent-based instantiation of traffic Engineering**

- Minimization of metrics: IGP, delay
- Exclusion of properties: link-affinity, SRLG

Benefits

Simplicity and Automation

- IGP-computed TE-path from anywhere to anywhere
- Sub-50msec protection (TILFA) optimized per Flex-Algorithm plane

Scalability

- Single SID (instead of label stack) to enforce TE path
- Single prefix segment can participate in many Flex-Algos

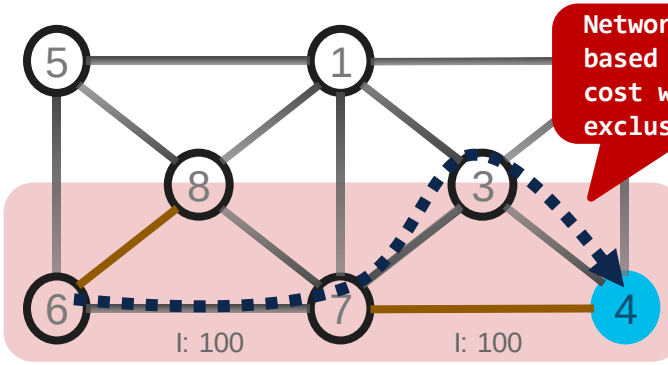

Default IGP
link metric: I:10

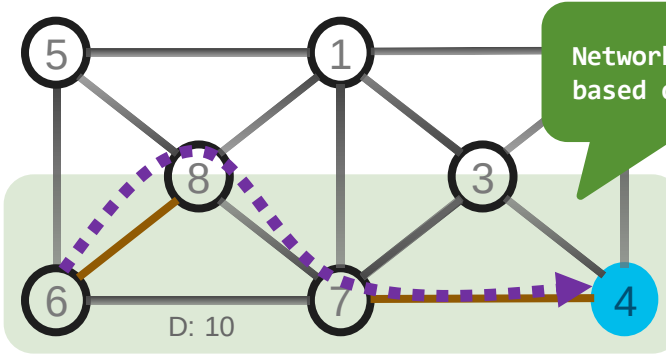

Measured link
Delay: D:1

Slicing: SR IGP Flexible Algorithms

- Each node MUST advertise Flex-Algo(s) that it is participating in

Nodes 0 and 9 participate to Algo 0 and 128 and 129

Nodes 1/2/3/4 participate to Algo 0 and 128

Nodes 5/6/7/8 participate to Algo 0 and 129

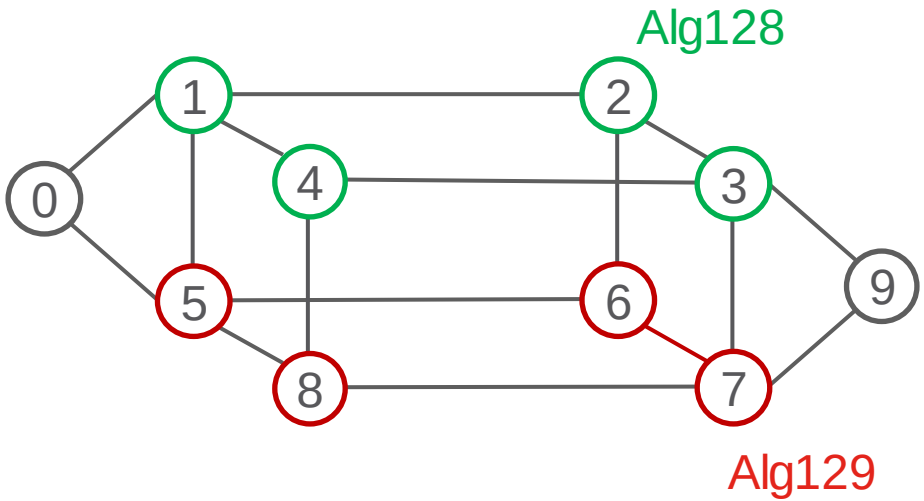

Slicing: SR IGP Flexible Algorithms

- Each node MUST advertise Flex-Algo(s) that it is participating in

Nodes 0 and 9 participate to Algo 0 and 128 and 129

Nodes 1/2/3/4 participate to Algo 0 and 128

Nodes 5/6/7/8 participate to Algo 0 and 129

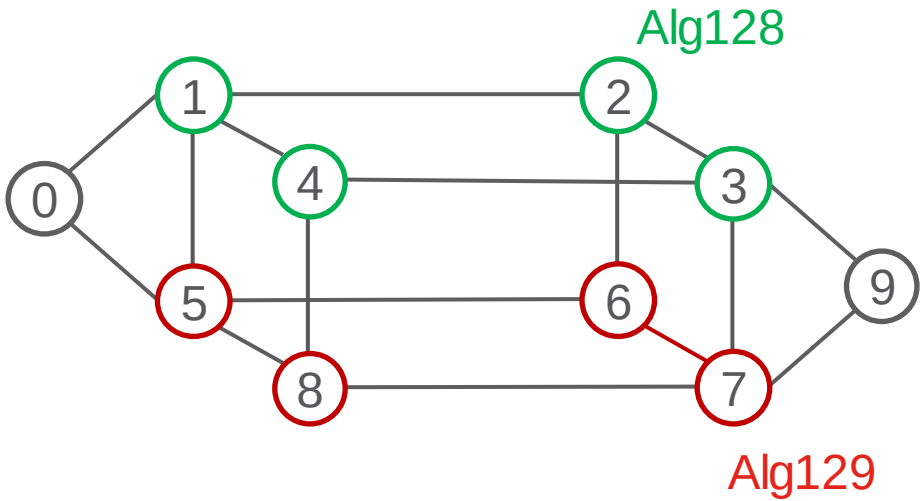

Slicing: SR IGP Flexible Algorithms

- If a node advertises participation in a Flex-Algo likely it also advertises a prefix SID for that Flex-Algo

Node 9 advertises

Prefix SID 16009 for ALGO 0

Prefix SID 16809 for ALGO 128

Prefix SID 16909 for ALGO 129

Node 2 advertises

Prefix SID 16002 for ALGO 0

Prefix SID 16802 for ALGO 128

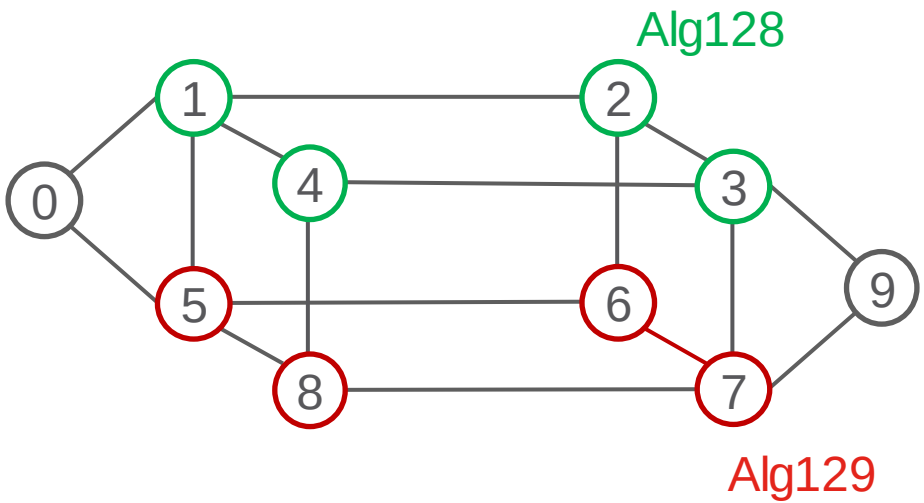

Slicing: SR IGP Flexible Algorithms

- Flex-Algo Prefix SID's can be advertised as additional prefix-SID's of the existing loopback address

Node 9 advertises loopback0 9.9.9.9/32 with

Prefix SID 16009 for ALGO 0

Prefix SID 16809 for ALGO 128

Prefix SID 16909 for ALGO 129

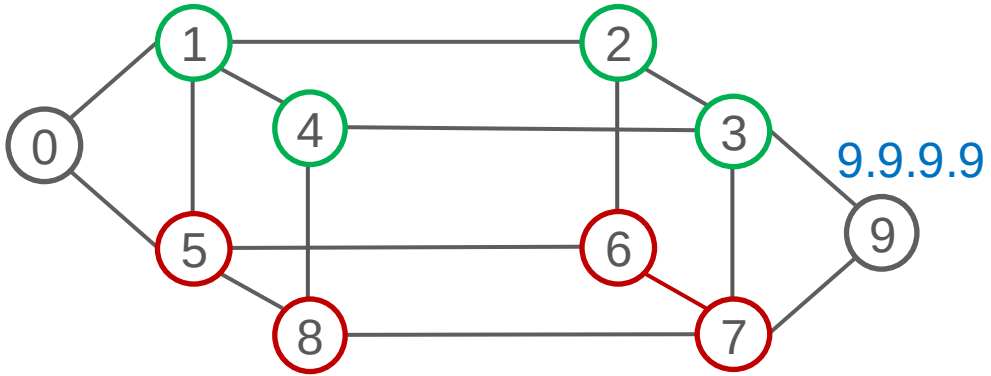

Slicing: SR IGP Flexible Algorithms

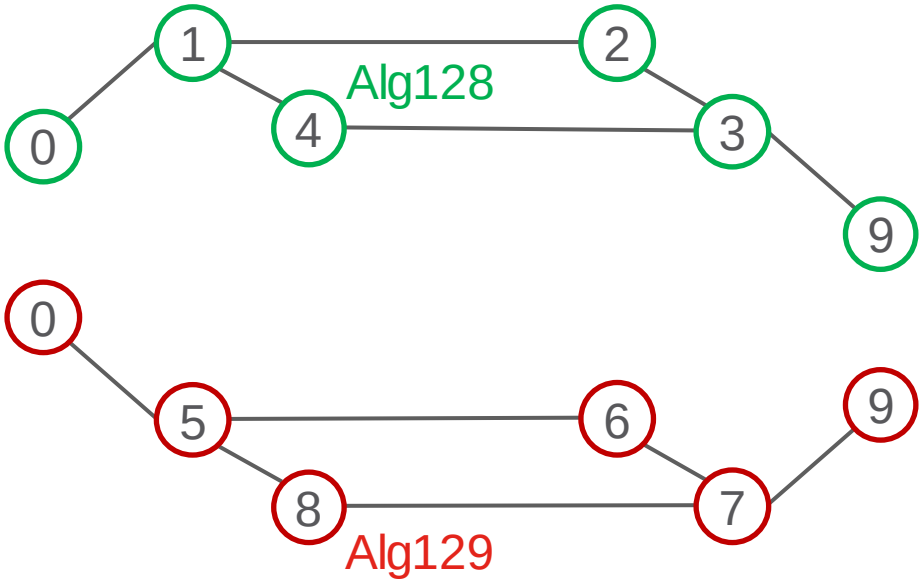

Topology Independent LFA (TI-LFA)

- 100%-coverage 50-msec link and node protection
- Prevents transient congestion and suboptimal routing
  - leverages the post-convergence path, planned to carry the traffic
- Simple to operate and understand
  - automatically computed by the IGP
- Incremental deployment
  - also protects LDP and IP traffic

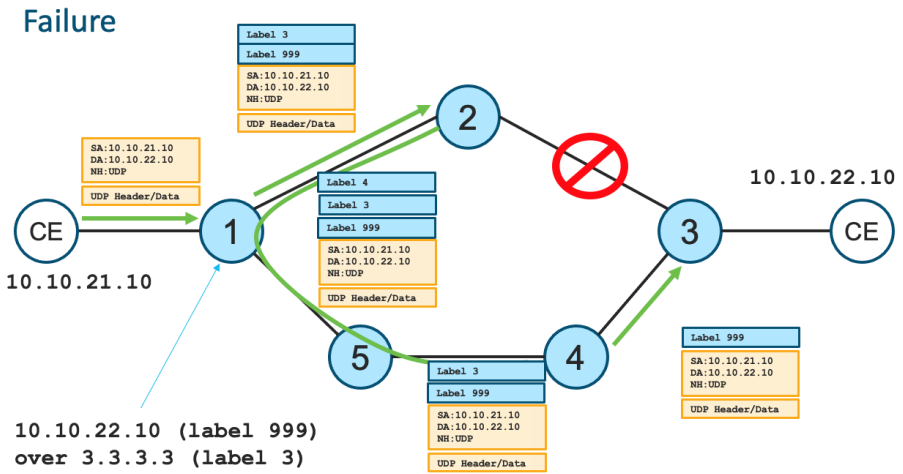

Transport Use-case Requirements

	5G xHaul	Cisco Solution
Bandwidth	✓	Silicon and Optics Innovation
Scalability	✓	Segment Routing
End to end Programmability	✓	Segment Routing
Synchronization	✓	Built-in
Security	✓	Built-in

Segment Routing Resumo

- Segment Routing Transport
- TI-LFA for improving High Availability
- Programable paths with SR-PCE
- Flex Algo for Network Slicing/Segmentation
- SR ODN & Auto Steering
- Performance Monitoring

Cisco Validated Network Designs

http://xrdocs.io/design

Converged SDN Transport High-Level Design

https://xrdocs.io/design/blogs/latest-converged-sdn-transport-hld

Converged SDN Transport Implementation Guide

https://xrdocs.io/design/blogs/latest-converged-sdn-transport-ig

Converged Portfolio

Cisco Access & Aggregation Portfolio

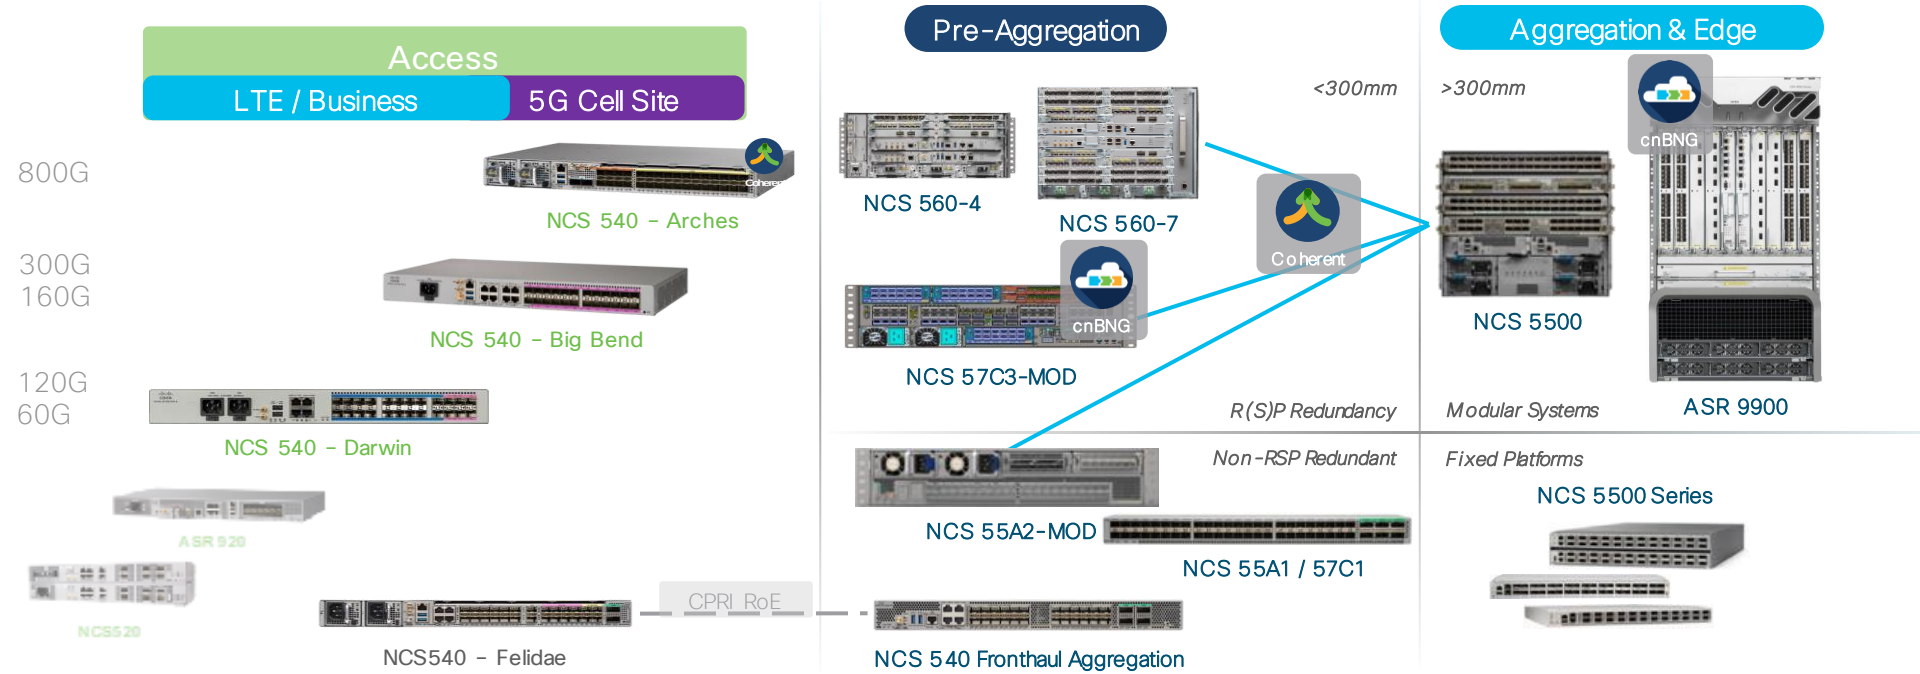

BGP-based Services with EVPN

Segment Routing

Embedded Security & Sync

Cell Site Router Use Cases and Examples

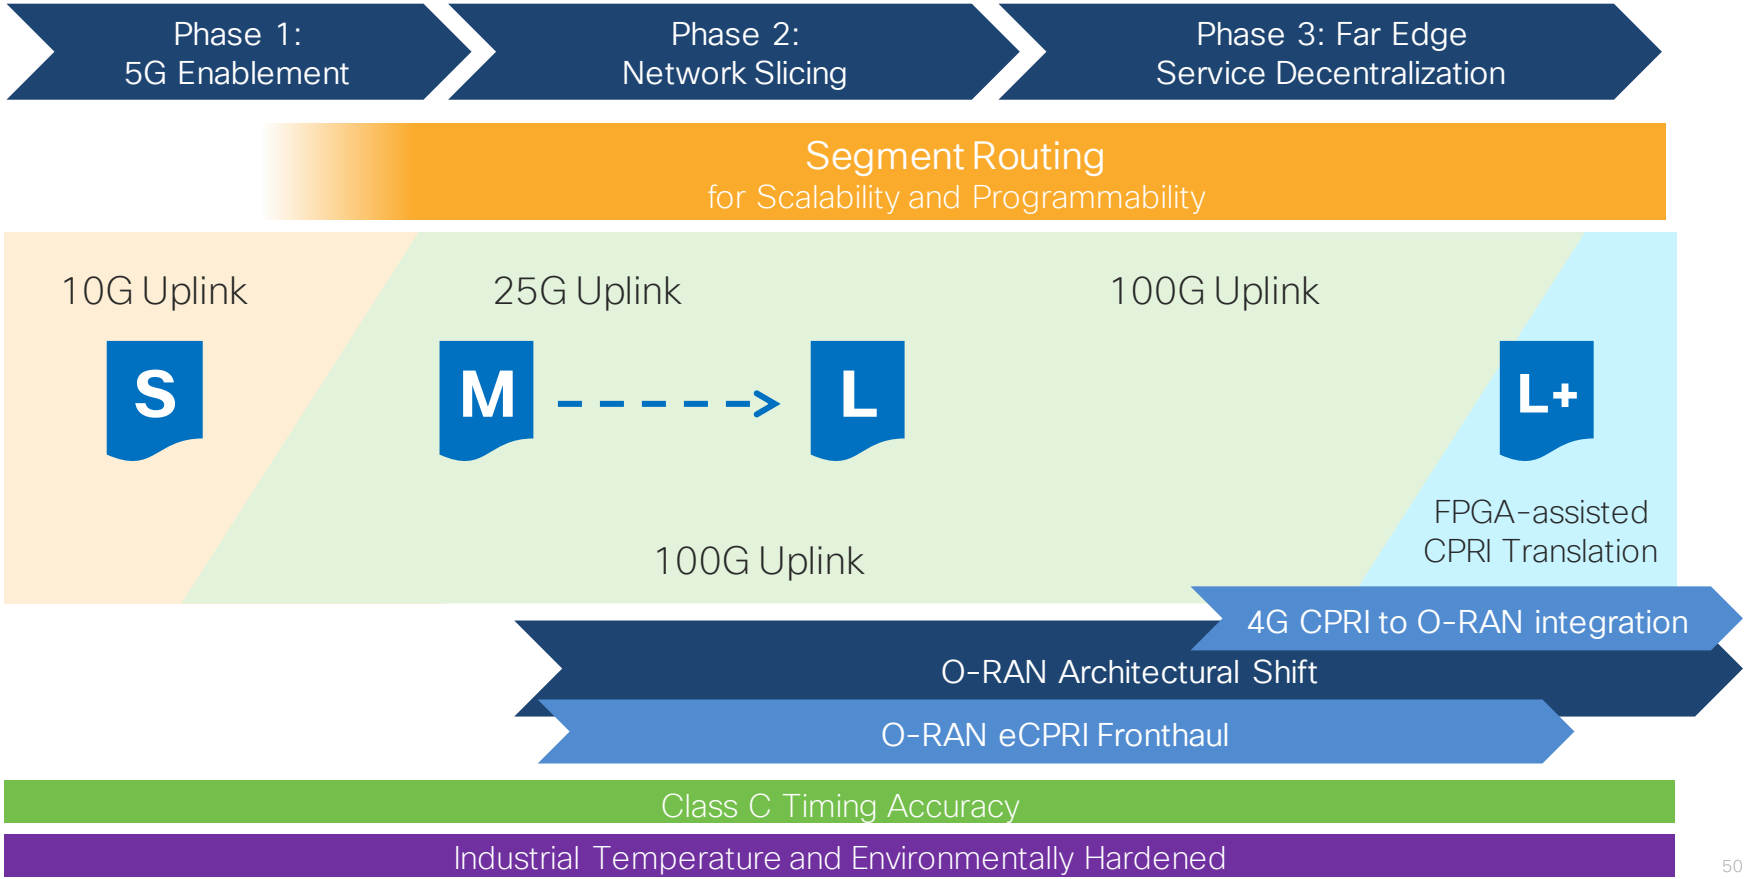

T-Shirt Sizes

		Specifics	*Use Cases*
S	Galapagos  Acadia 	10G Uplink	Distributed RAN sites (migration) Rural with Low-band Spectrum, no growth *Not a stop-gap*, these low-scale sites will be around for quite some time
M	Everglades 	25G Uplink	Distributed RAN sites with some 5G mid-band carriers and limited expected growth
L		100G Uplink	Multi-service sites with multliple 5G mid/high-band carriers and potential growth due to refarming or eCPRI fronthaul
L+	Felidae 	100G Uplink FPGA-assisted CPRI Translation	Multi-service sites with additional need for 4G CPRI translation to O-RAN compliant eCPRI fronthaul

Portfolio Highlights

Right Sized

1RU, <300mm – 21 RU fully modular

Speed & Density

1G/10G/25G/40G /50G/100G/200G/ 400G

ZTP

Secure & faster Deployment

Low Latency

Class-B, C portfolio, TSN*

Trust

Hardware anchored trusted

Single OS

Powered with Cisco IOS XR

Automation

Telemetry, Netconf-Yang

Simplified Transport & Service

SR/SRv6/EVPN

* Refer product documentation for exact capability

Portfolio Resumo

- NCS 5500 – Core
- NCS 560 – Aggregation
- NCS 540 – Access

540's Command Line Interface

NCS 540 Scale – IPv4/IPv6 Routes, MPLS Labels, MAC Addresses

LEM (Large Exact Match) Memory:

IPv4 Unicast /32 and /24
IPv6 Unicast /48
MPLS Labels (Local)
MAC Addresses (Max QAX: 128k, QUX: 64k)

LPM (Longest Prefix Match) or KAPS Memory:

IPv4 Unicast non-/32 and non-/24
IPv6 Unicast non-/48
IPv4 Multicast
IPv6 Multicast

LPM/KAPS is Algorithmic memory = Scale depends on prefixes and prefix length distribution
Numbers shown are official numbers from Broadcom for unicast prefixes (conservative – public Internet)

Route Distribution Tweaks for QAX and QUX

- Moves IPv4 /24 from LEM to LPM

```
hw-module fib ipv4 scale host-optimized-disable
```

- Moves IPv6 /48 from LEM to LPM

```
hw-module fib ipv6 scale internet-optimized-disable
```

- Choose which one of the IPv6 prefix lengths will be stored in LEM (since 7.1.1)

```
hw-module fib ipv6 scale custom-lem 40|44|48|52|56|60|64
```

DRAM & SRAM info

- TOTAL buffer utilization command

```
RP/0/RP0/CPU0:NCS540X_LAB#show controller fia diagshell all "diag ing" location all
Mon Aug 9 10:21:27.483 WEST

Node ID: 0/RP0/CPU0

R/S/I: 0/0/0

Ingress congestion statistics:

CORE 0:
-----------Free DRAM BDBs: 98302
Free minimum DRAM BDBs: 98301
Free SRAM free buffers: 16384
Free minimum SRAM free buffers: 16306
Free SRAM PDBs: 16380
Free minimum SRAM PDBs: 16368
RP/0/RP0/CPU0:NCS540X_LAB#
```

SRAM (16384*256B)
4Mb available
(100%)

DRAM 98302 BD available
Each BDB can points up to 8 consecutive DBs (DRAM Buffer Descriptor) that belong to a particular VoQ.
Buffer Descriptor: 4096 bytes

Examples

Command that provide information related to congested queue

```
RP/0/RP0/CPU0:NCS540X_LAB#show  controller fia diagshell all "diag cosq non" location all
Fri Aug 13 12:40:24.108 WEST

Node ID: 0/RP0/CPU0

R/S/I: 0/0/0

Core 0:
Ingress VOQs Sizes (format: [queue_id(queue_size)]):
[1288(93696B)]
RP/0/RP0/CPU0:NCS540X_LAB#
```

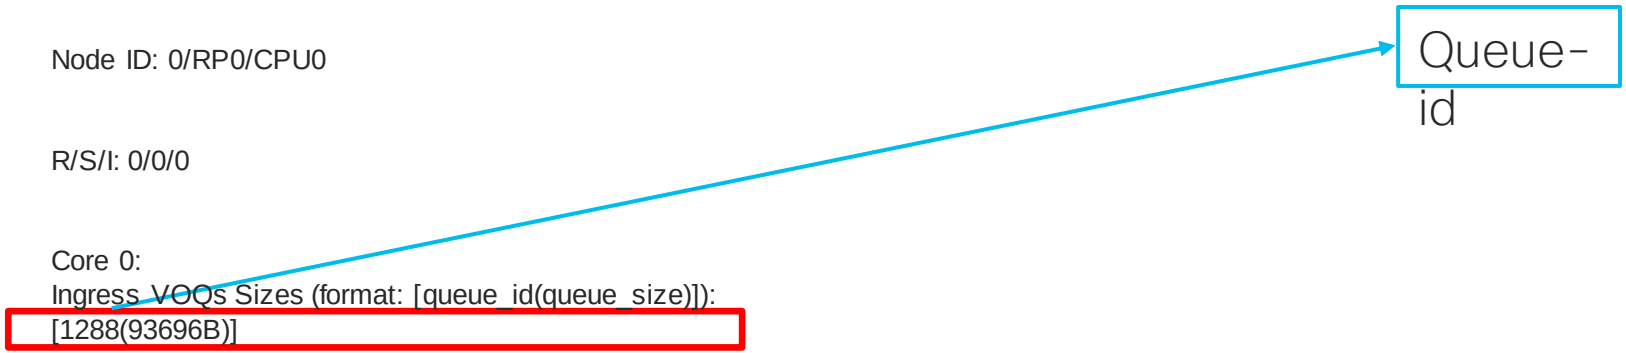

Conclusão

Segment Routing Unified Fabric vision

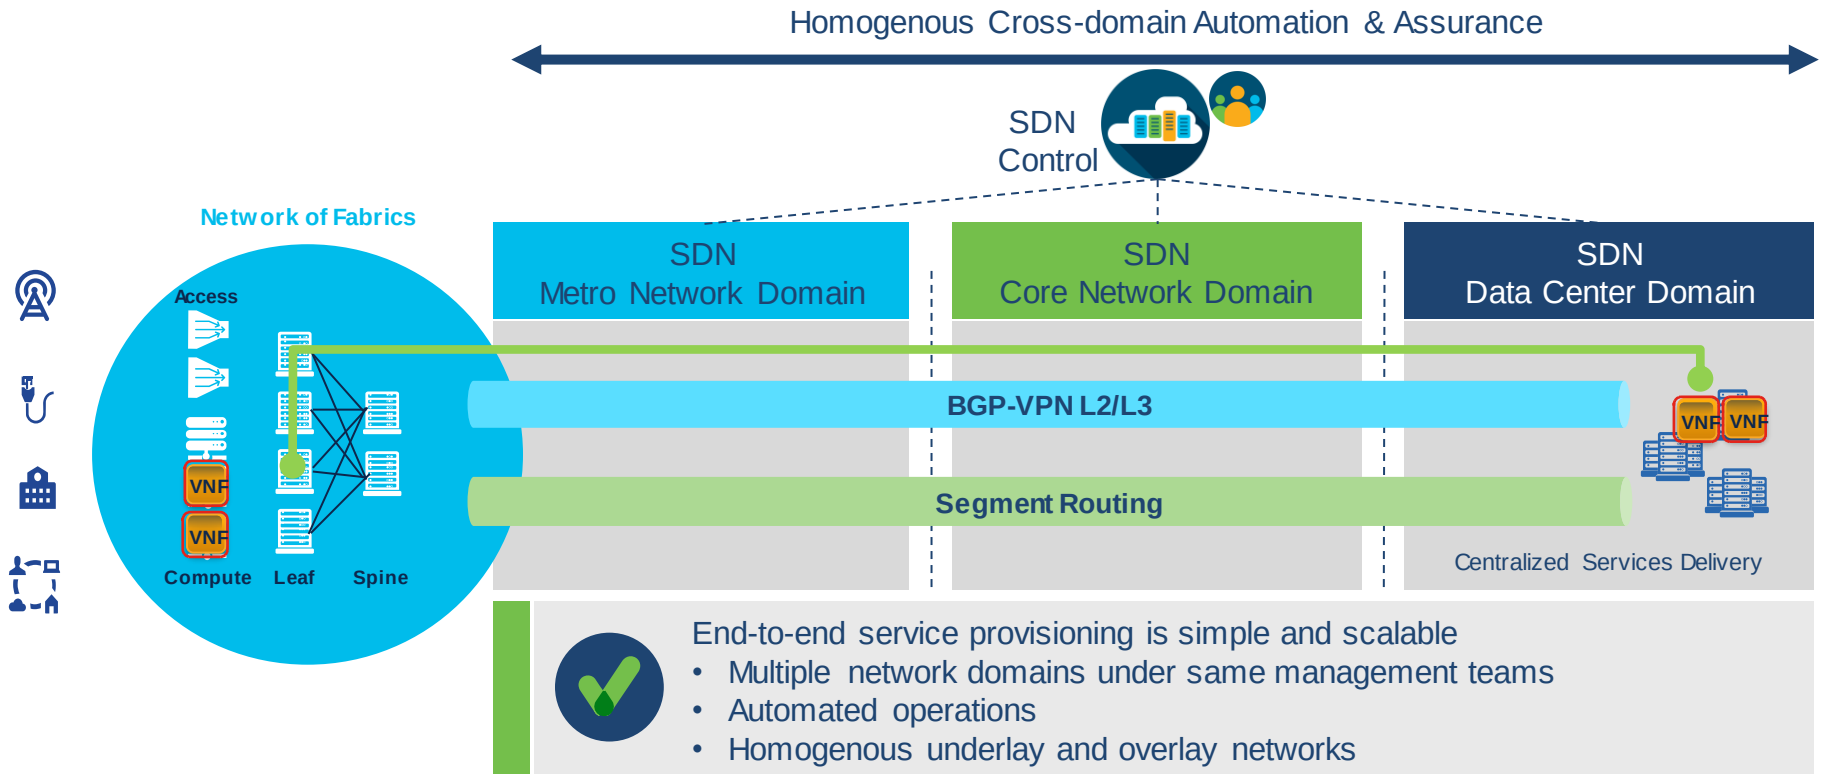

5G Needs from Transport

- ☑ Bandwidth (efficient usage)
- ☑ Latency awareness
- ☑ Network Slicing
- ☑ Reliability
- ☑ Simplicity = Programmability

The bridge to possible