What You Make Possible

BUILT FOR
THE HUMAN
NETWORK
CISCO

Unified MPLS for Mobile Transport

BUILT FOR THE HUMAN NETWORK

Agenda

- Mobile Market Dynamics
- Unified MPLS for Mobile Transport (UMMT) system overview
- UMMT Functional Considerations
  - QoS
  - Resiliency
  - OAM and PM
  - Synchronization Distribution
  - Security
  - SDN Perspectives
- Summary and Key Takeaways

Mobile Market Dynamics

Mobile Service Evolution

- FROM:
- Voice/SMS
- Macro cell deployment
- 2G/3G
- MPC
- Physical Gi-LAN appliances
- Static data and application locations

- TO:
- Data and video
- Small cell deployment
- LTE
- MPC offload
- Network Function Virtualization
- Cloud based services

Mid 2000s SP Architecture

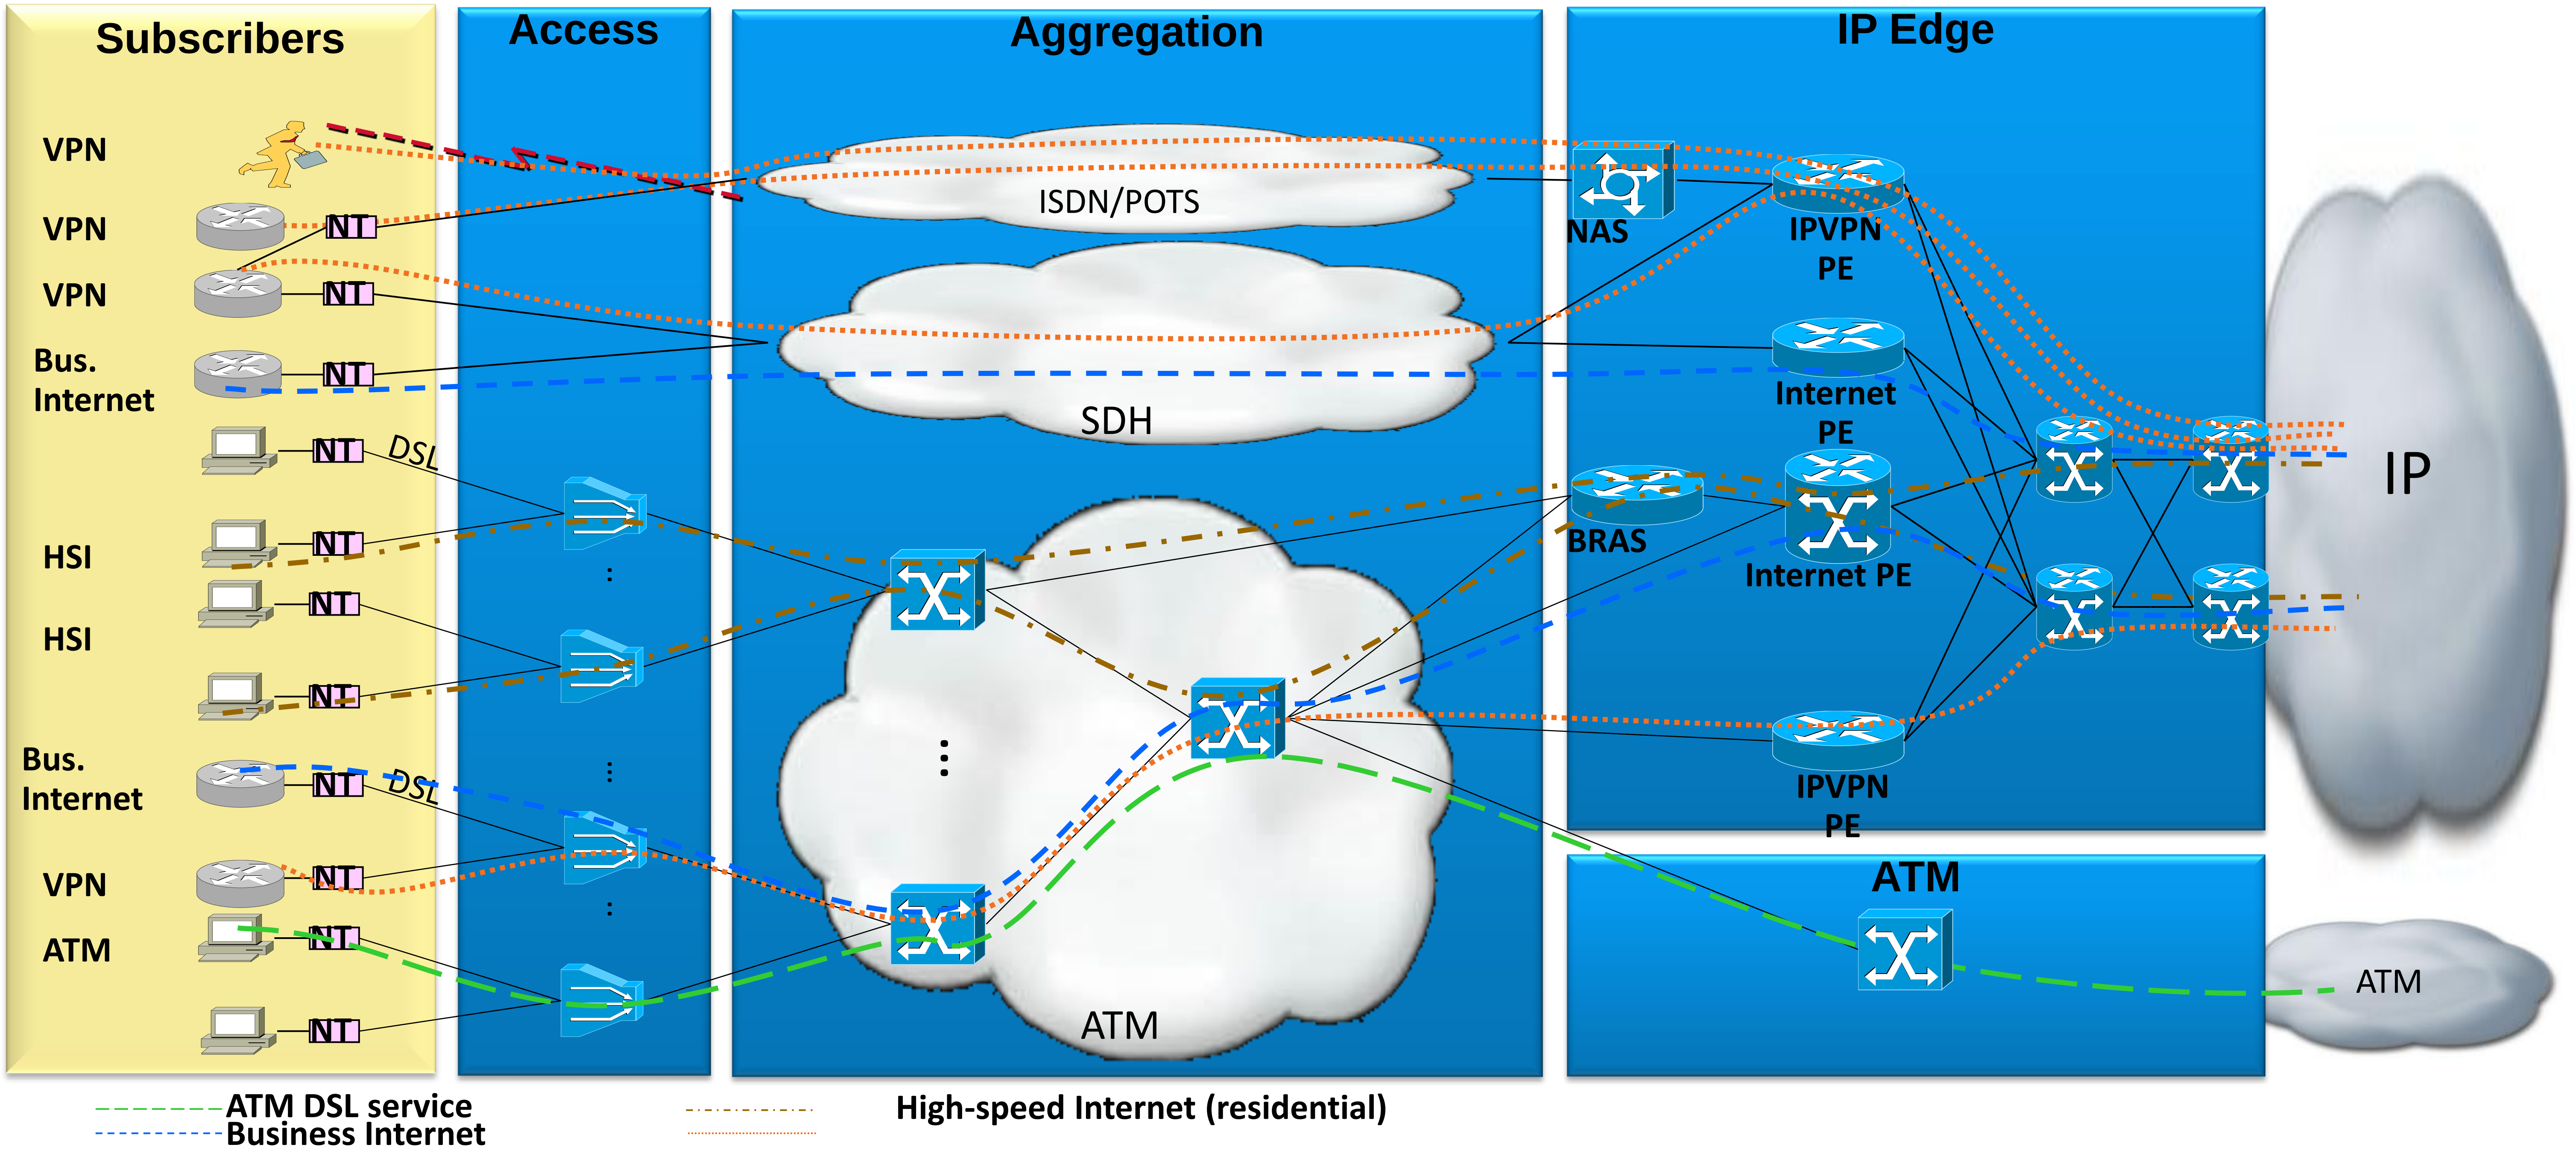

Global Mobile Data Traffic - 2011 to 2016

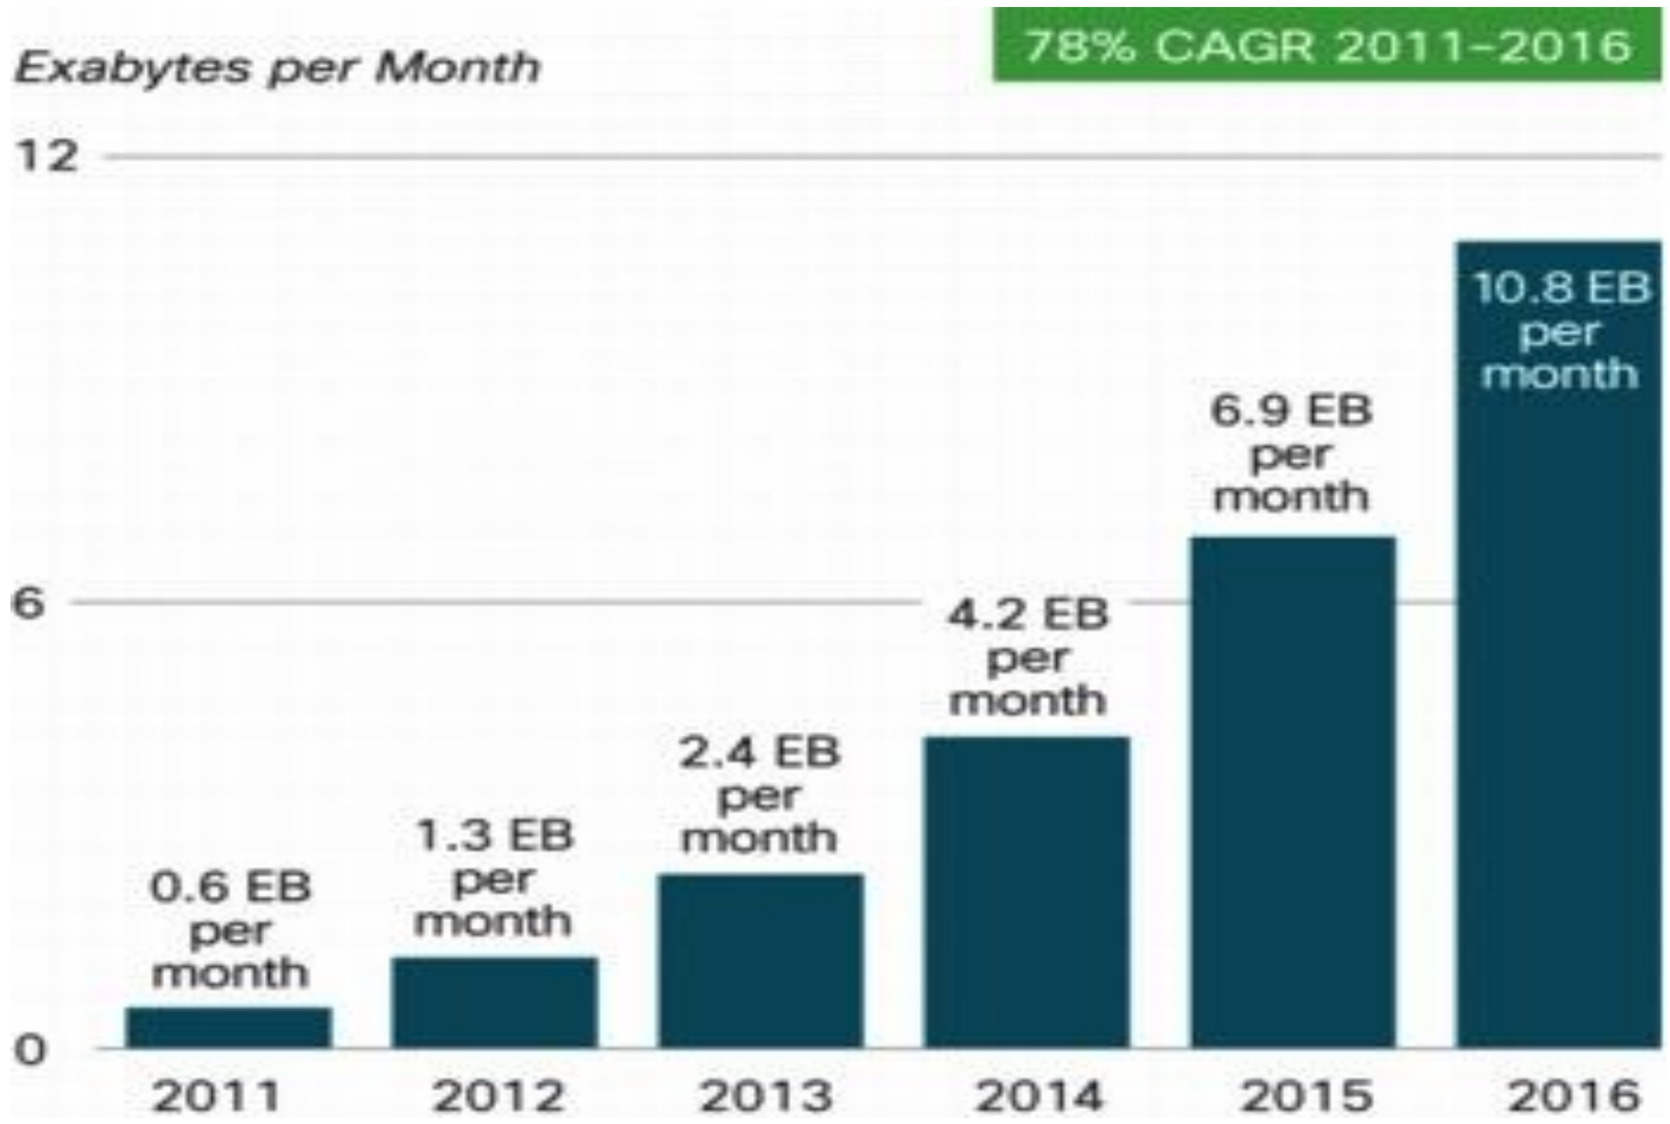

A 18x increase in mobile data traffic over the next 5 years

Video to comprise >70% of mobile data traffic by 2016

Machine-to-machine Traffic to increase 22x by 2016

Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2011-2016

Revenue Split and Traffic Predictions

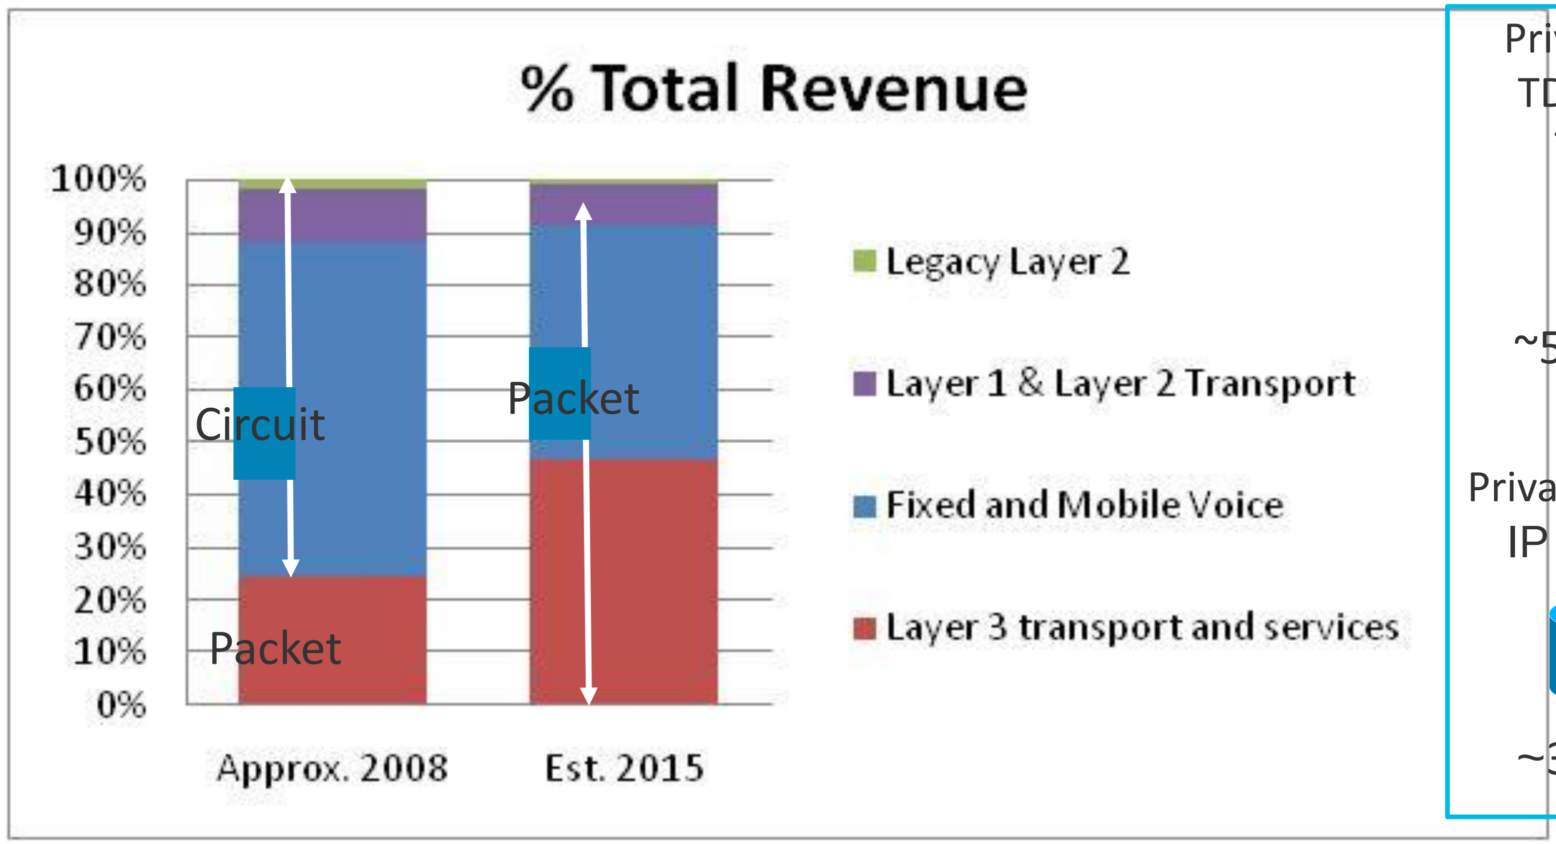

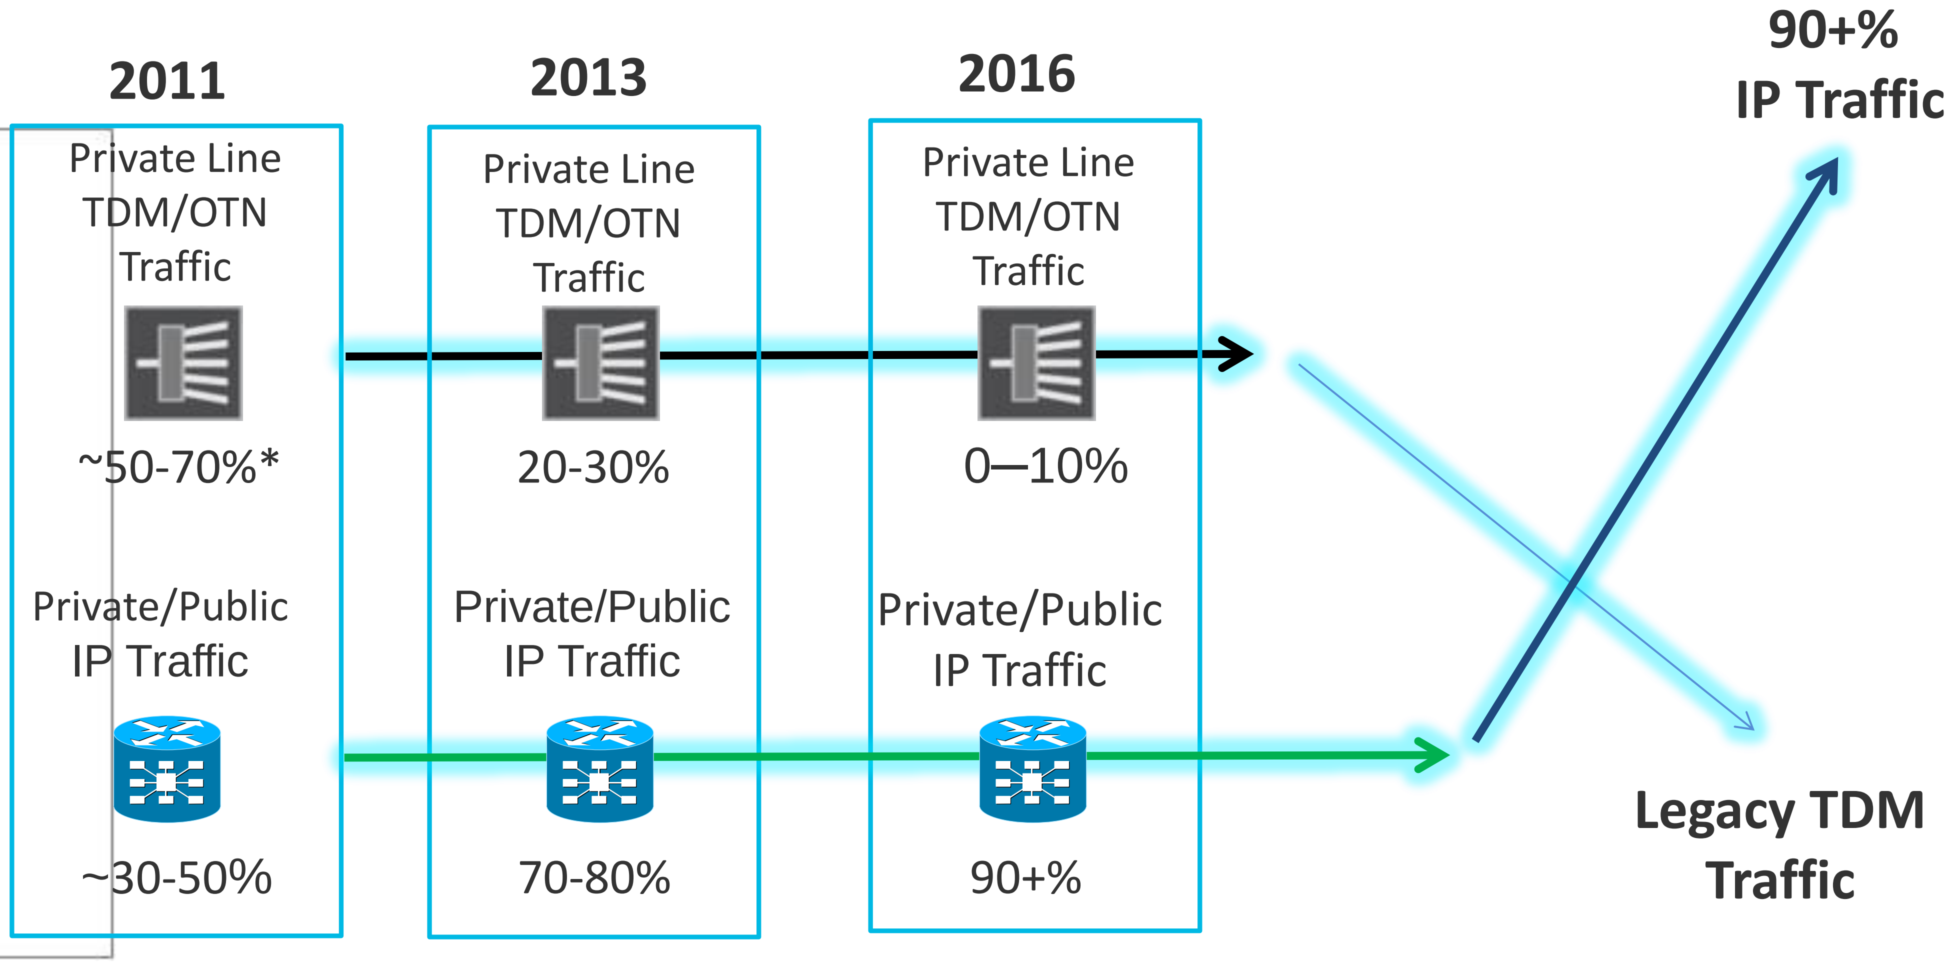

- SP revenue shifting from circuits to packet services**
  - 5 yrs → ~80% revenue derived from packet services
- Packet traffic increasing at 34% CAGR***
- Massive change in SP traffic make-up in next 5 years*

Economic Realities

Chart from Infonetics, Text from DT

Provisioning for Circuit/OTN and Packet/Router Aggregation

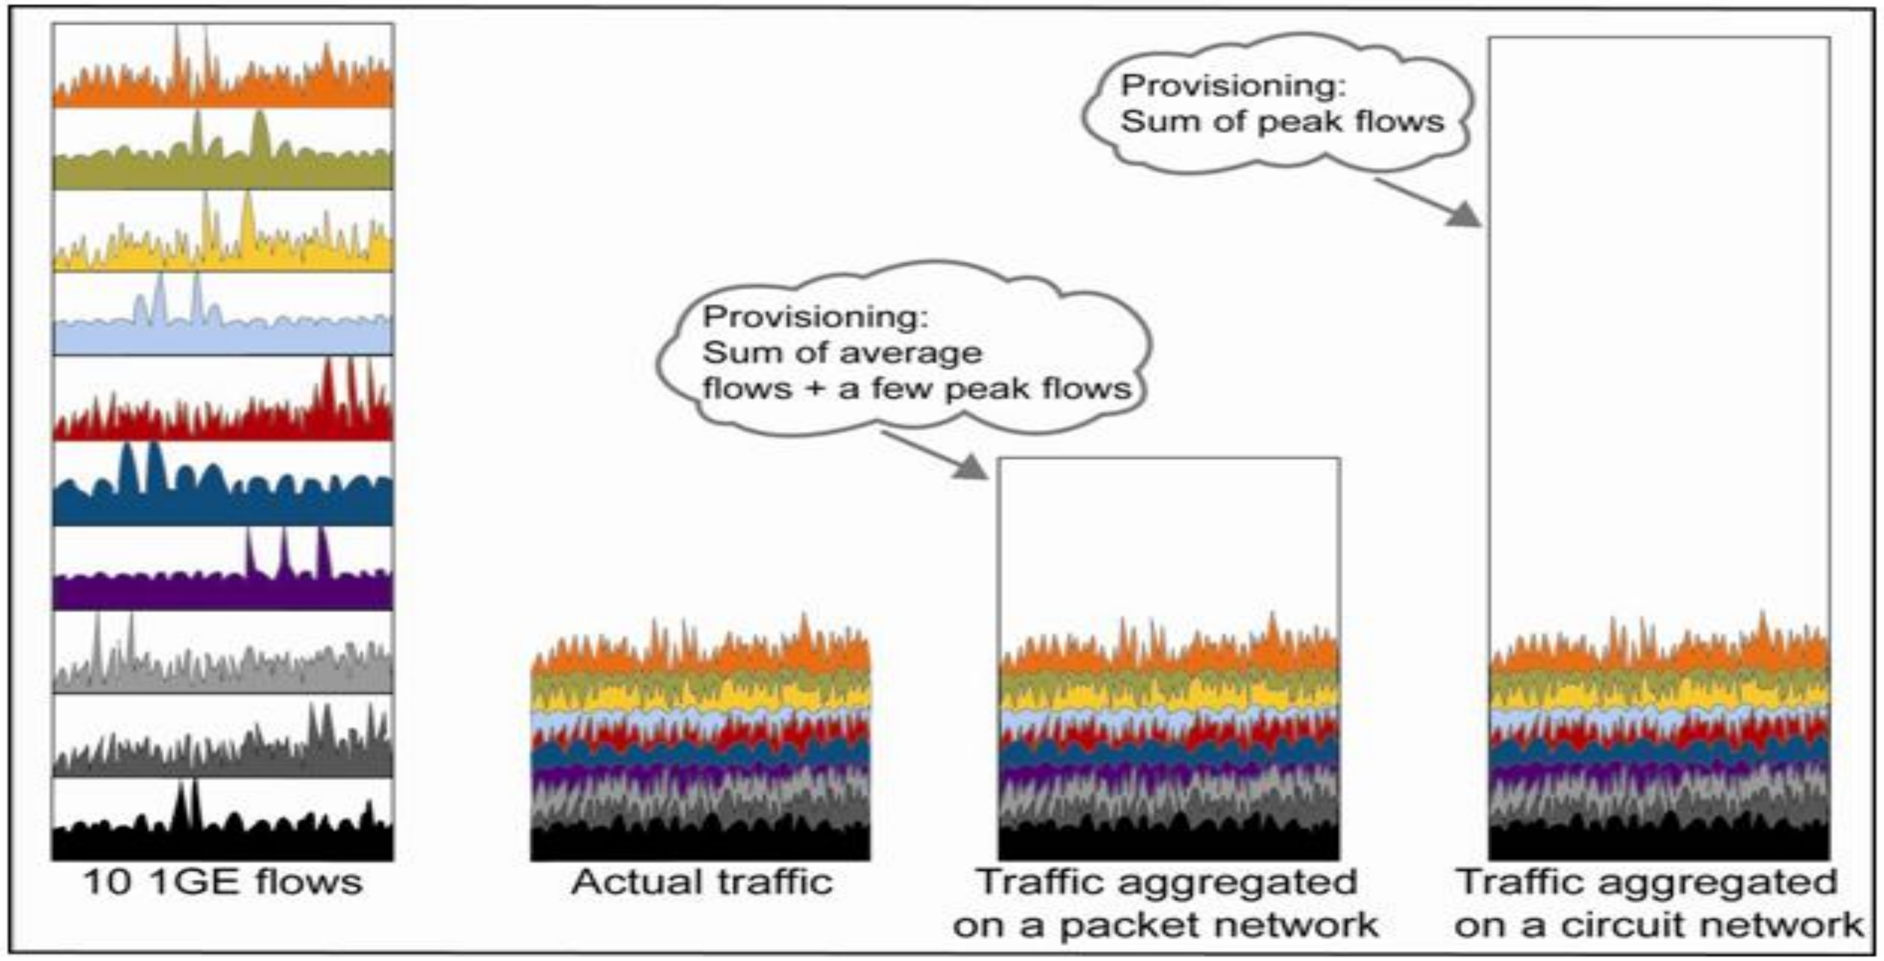

TDM transport of packets is no longer economically viable, lacks statistical multiplexing which makes it very expensive

Full transformation to NGN needs to occur from core to customer

Long term vision is critical, this will be the network for the next decade

What is the most effective technology choice that will:

- Minimize CapEx and OpEx?
- Provide carrier class service delivery?
- Maximize service agility?

Carriers want the deterministic attributes of transport networks with the flexibility of the internet

As we move towards 100G, Optical Interfaces NOT Router Port pricing dominate cost

Routers: 23% Cumulative Average $/Gbps Drop per year / fewer ASICs

Optics: $/G stays flat (best case) or increases from one technology to the next

Cisco Core Router Example

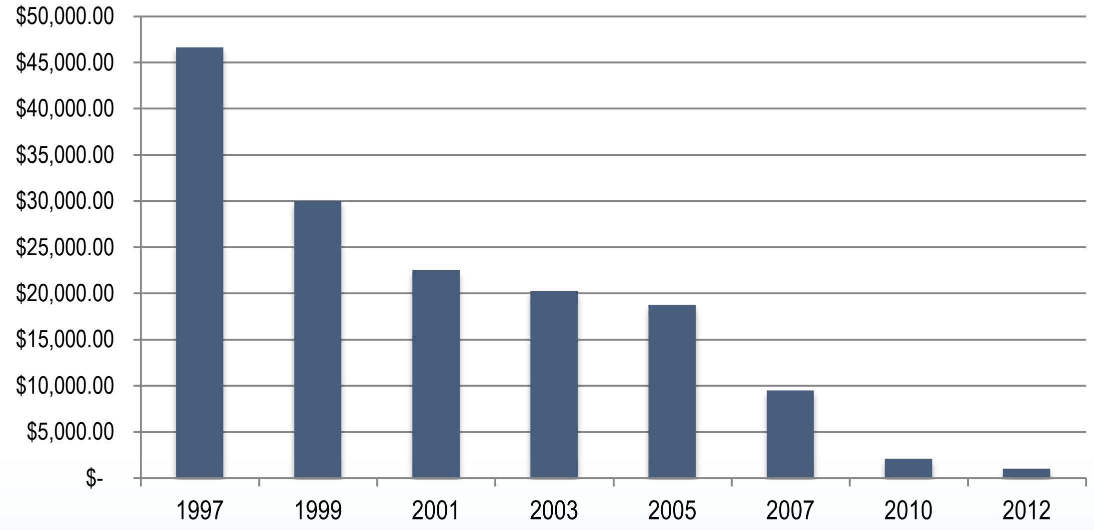

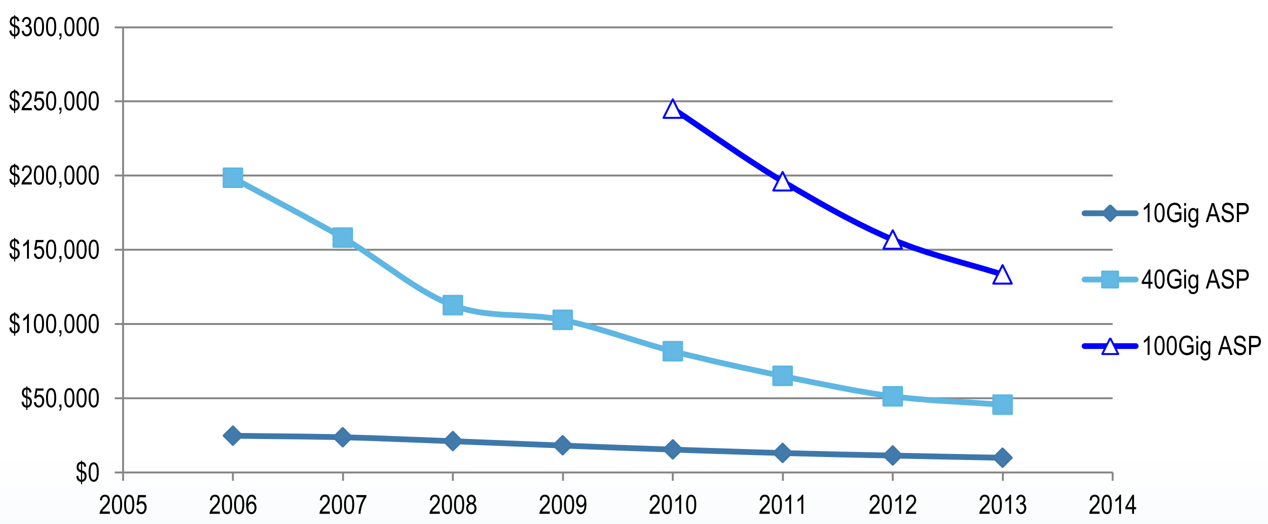

- Silicon has fundamentally followed Moore’s law
- Optics has fundamentally an Analog Problem

Mobile Transport Market Conditions

- High capacity requirements from edge to core:
  - 100Mbps eNB, 1Gbps Access, 10Gbps Aggregation, 100Gbps Core
- Higher scale as LTE drives ubiquitous mobile broadband
  - Tens- to hundred-of-thousands of LTE eNBs and associated CSGs
- Support for multiple and mixed topologies
  - Fiber and microwave rings in access, fiber rings, and hub-and-spoke in aggregation and core networks
- Need for graceful LTE introduction to existing 2G/3G networks
  - Coexistence with GSM Abis, TDM backhaul, and potentially UMTS IuB ATM
- Need to support transport for all services from all locations
  - Residential and business, retail and wholesale, L2 and L3 services from cell site where this is the most cost effective location for the customer
- Optimized operations with consistent packet transport

UMMT System Overview

Migrate With Simplicity; But is it Easy?

Simple is a pre-requisite for reliability

- Simple comes from Simplex, meaning single thread
- The opposite is Complex, meaning multi-threaded

Easy comes from adjacent, what is already known

- Opposite is what is difficult or unknown

Fig 1. represents each stovepipe networks. Each step by itself to bring universal connectivity is simple in the figure, but the end result is complexity

To create simplicity is initially hard, takes planning and understanding of all aspects of business and technology

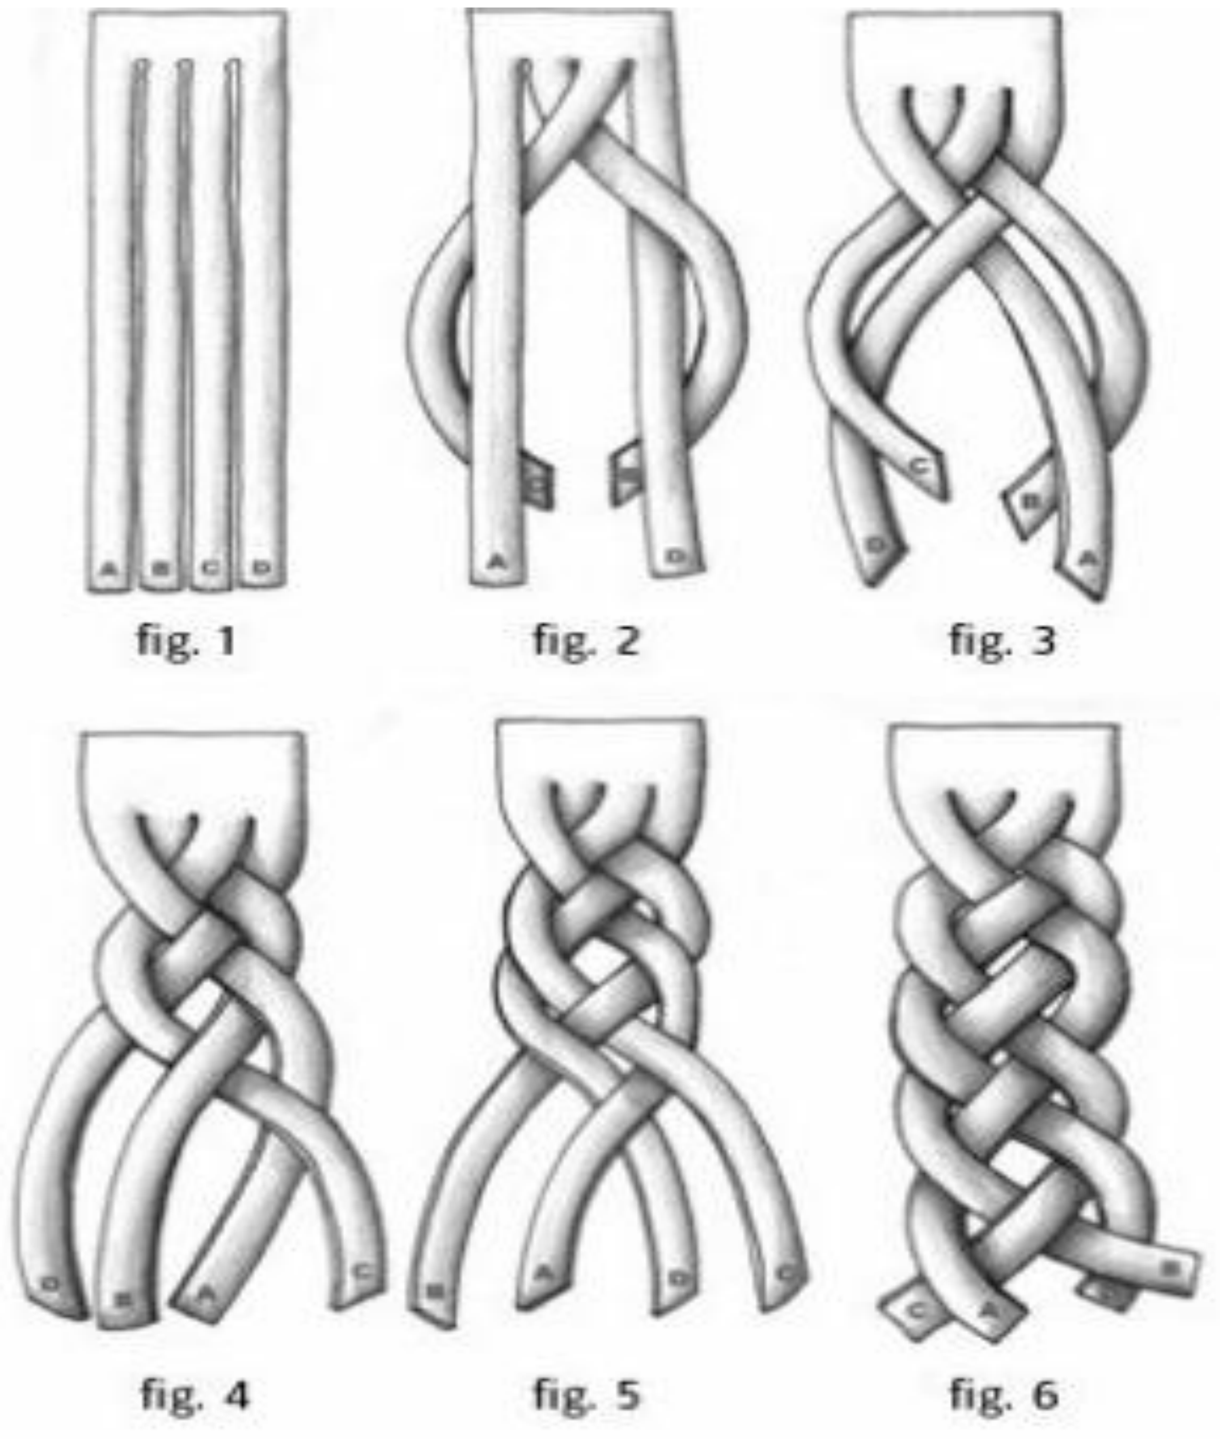

Unified MPLS Mobile Transport Concept

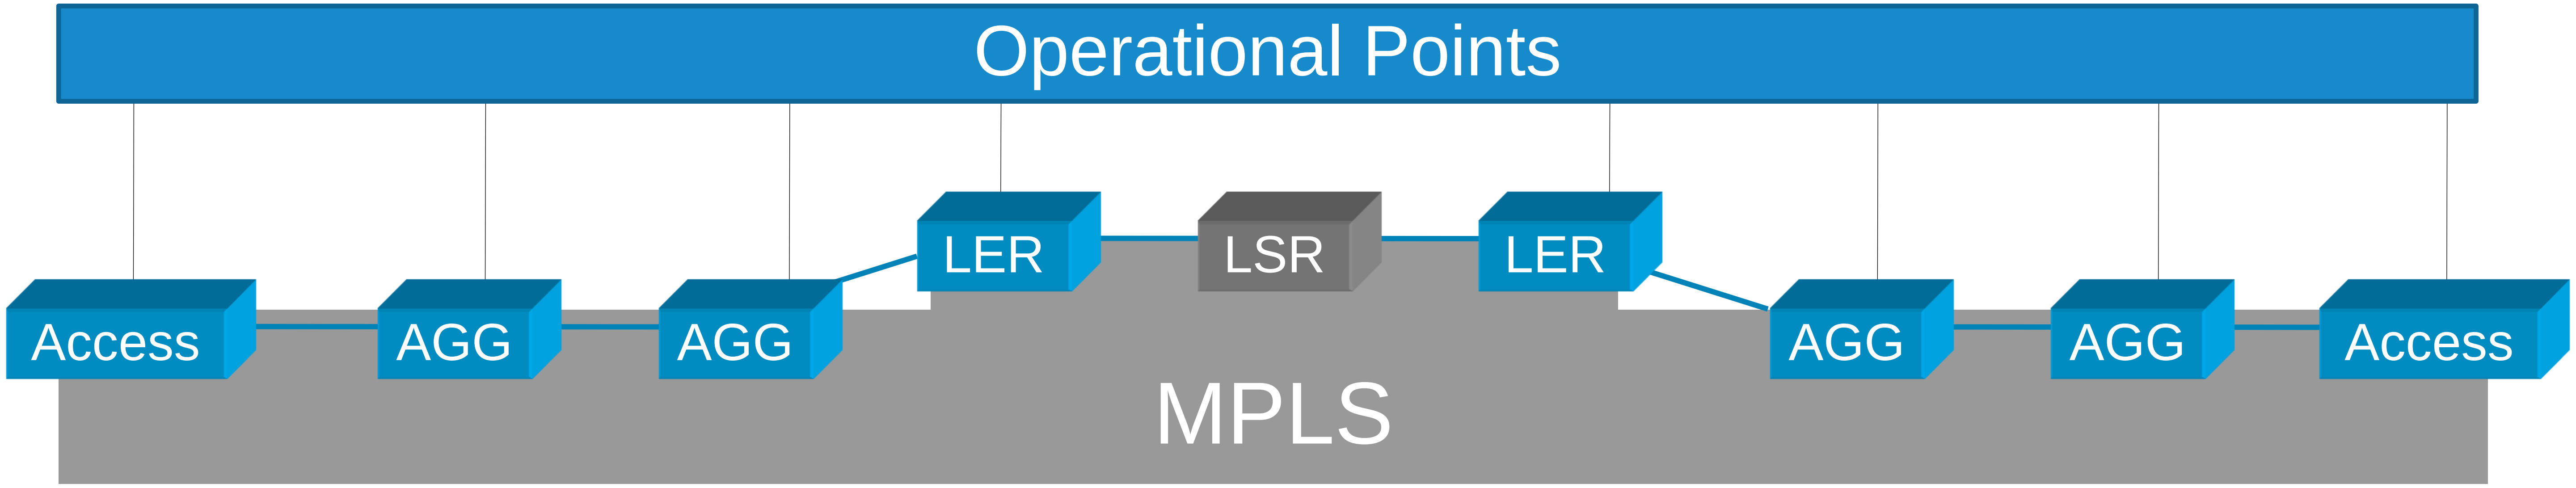

- **In general transport platforms, a service has to be configured on every network element via operational points. The management system has to know the topology.**
- **Goal is to minimize the number of operational points**
- **With the introduction of MPLS within the aggregation, some static configuration is avoided.**
- **Only with the integration of all MPLS islands, the minimum number of operational points is possible.**

Bringing it all together- “Show me the Labels”

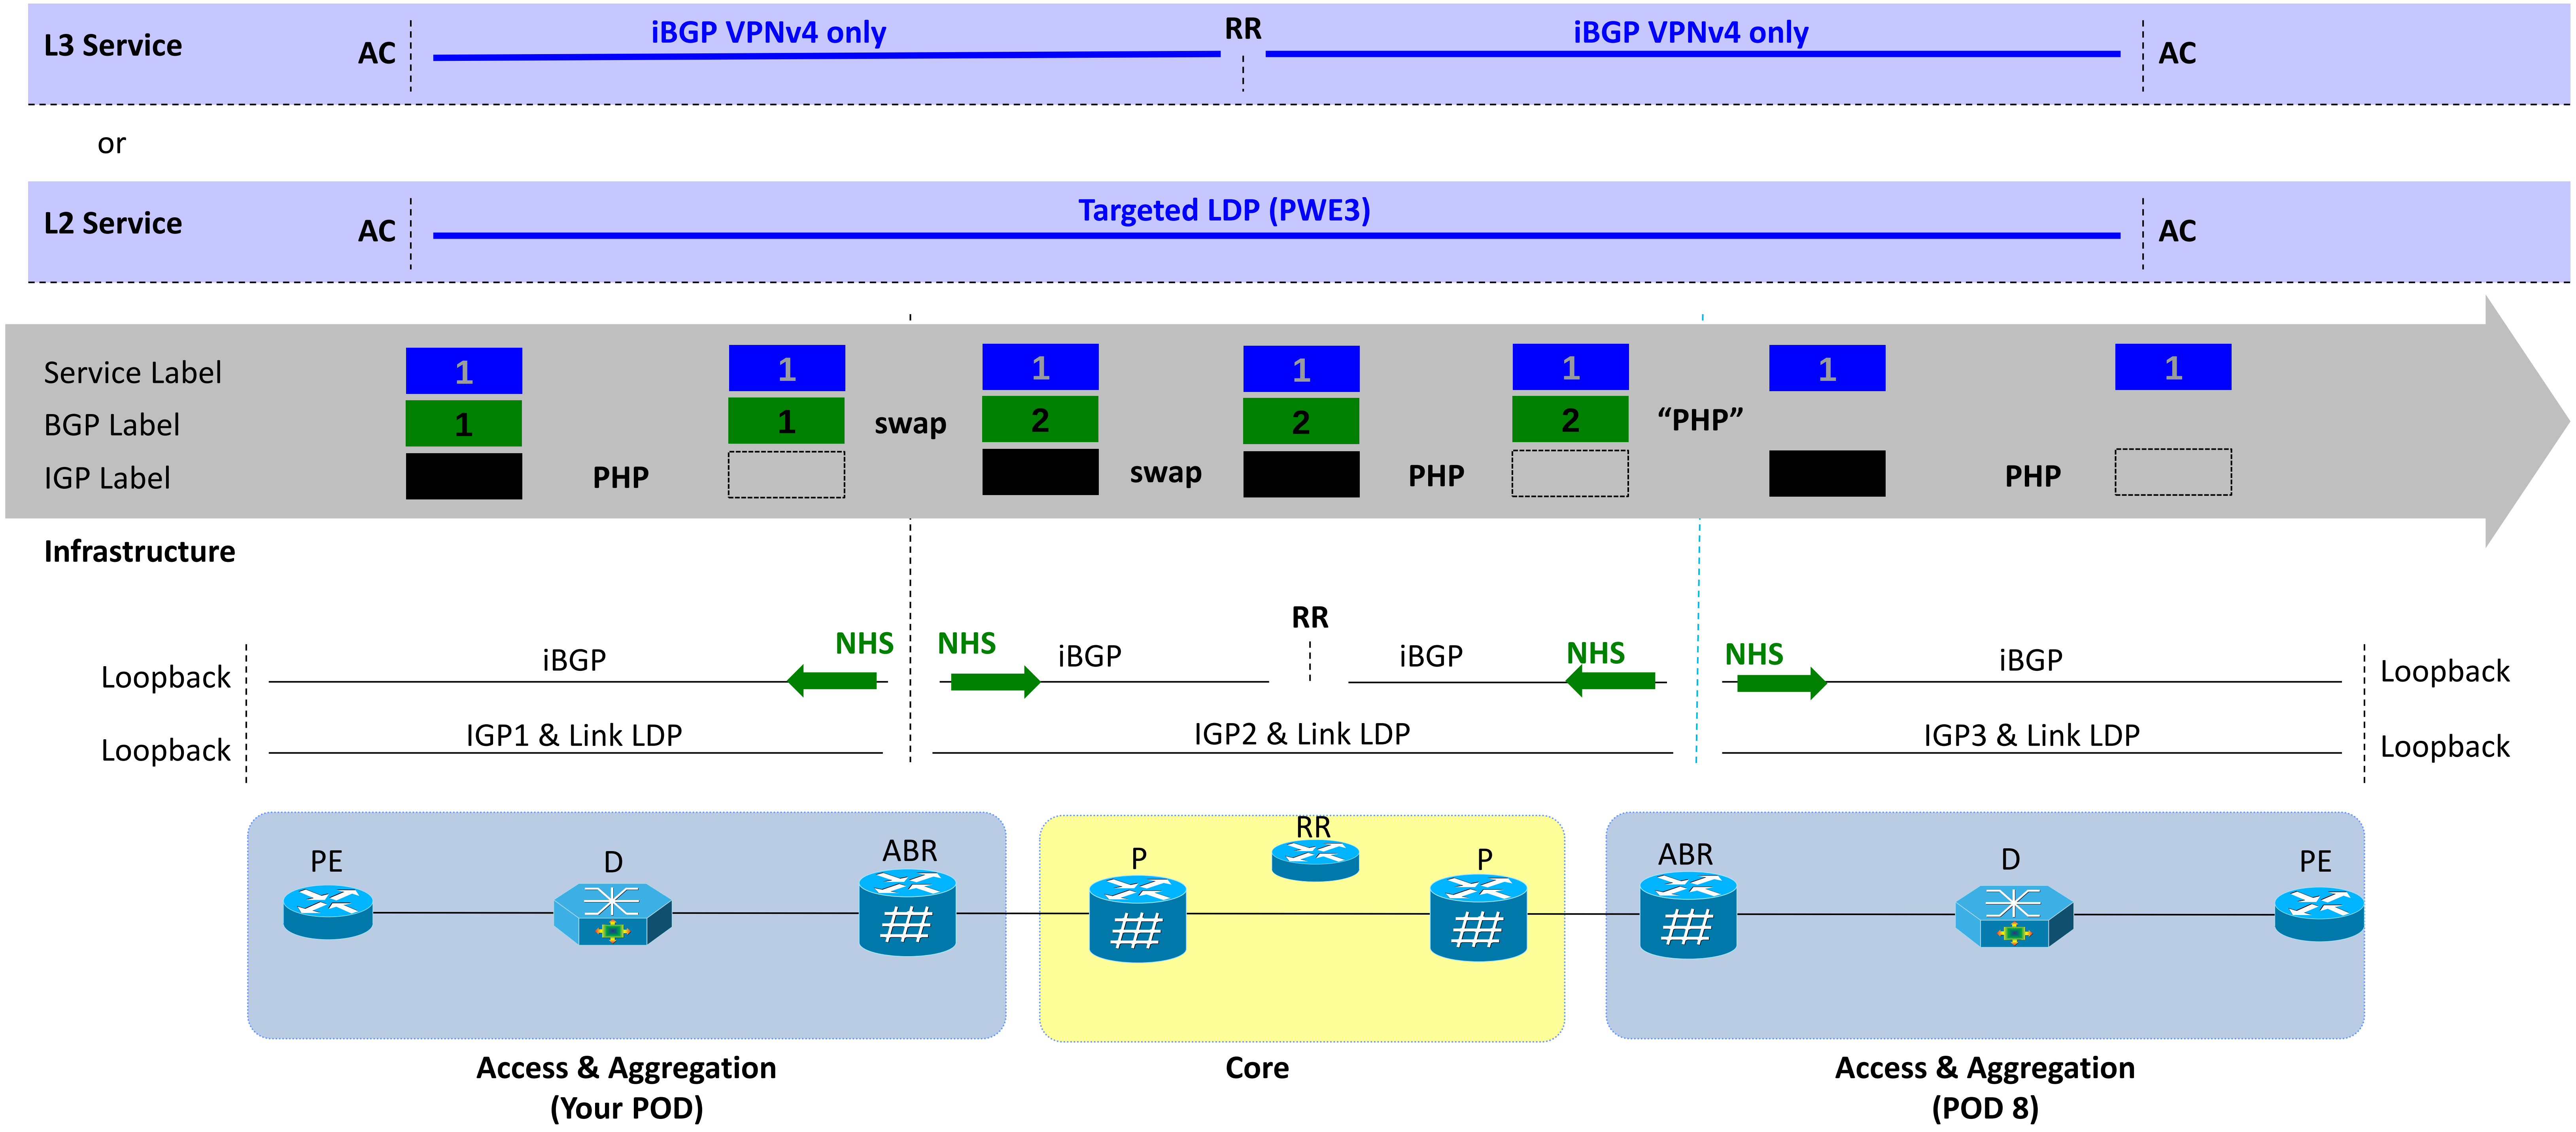

IGP1, IGP3, IGP3 … can be different OSPF/ISIS processes or ISIS L1/L2 or OSPF area 0 and X

Unified MPLS Mobile Transport Concept

- **A Unified MPLS transport simplifies the end-to-end architecture, eliminating the control and management plane translations inherent in legacy designs**
  - Seamless MPLS LSPs across Access, Aggregation & Core
  - Service provisioning done only at the edges
- **Flexible placement of L3 and L2 transport virtualization functions required to support retail and wholesale backhaul for GSM, UMTS, LTE**
  - Different options optimized for different topologies
- **Delivers a new level of scale for MPLS transport with RFC-3107 hierarchical labeled BGP LSPs**
- **Simplified carrier class operations with end-to-end OAM, performance monitoring, and LFA FRR fast convergence protection**
- **Extensible to wireline residential, business, retail/wholesale L2 and L3 VPNs, and IP services**

UMMT vs. LDP Control Plane Scalability

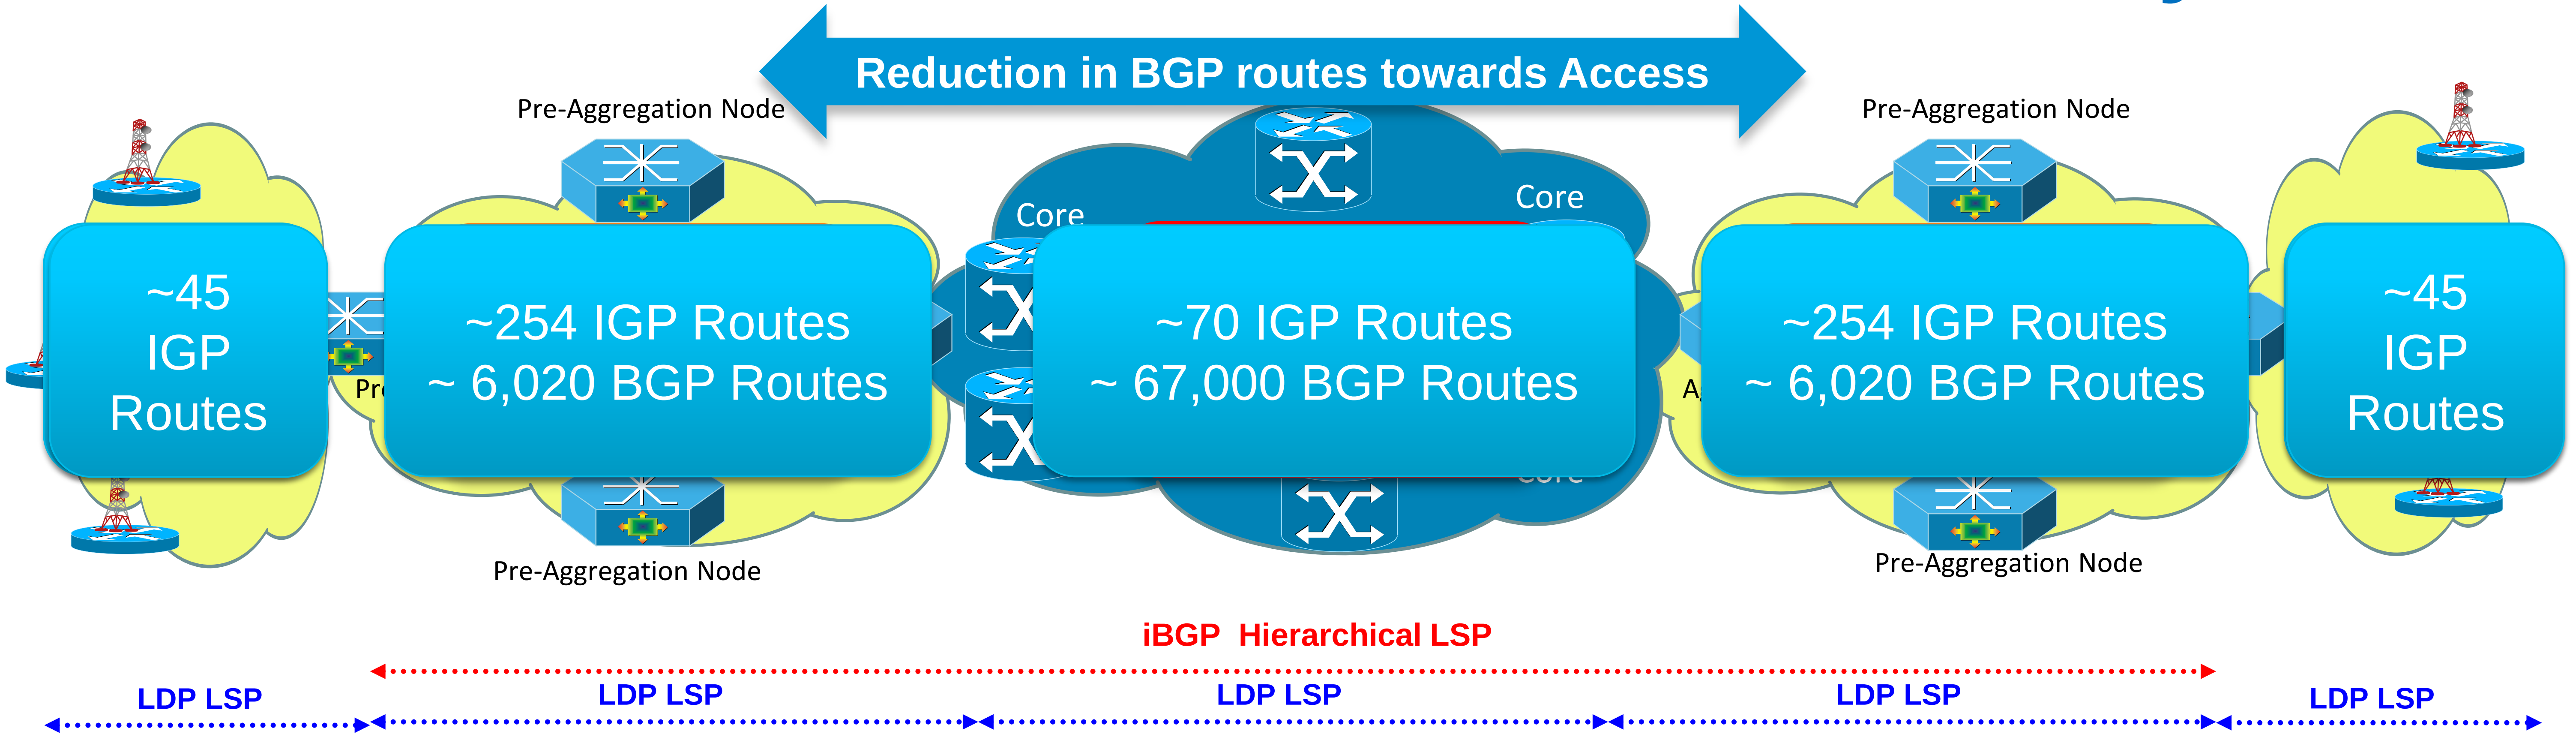

Node	Access Domain	Aggregation Domain	Network Wide
Cell Site Gateways	20	2,400	60,000
Pre-Aggregation Nodes	2	240	6,000
Aggregation Nodes	NA	12	300
Core ABRs	NA	2	50
Mobile transport Gateways	NA	NA	20

UMMT Architecture Overview

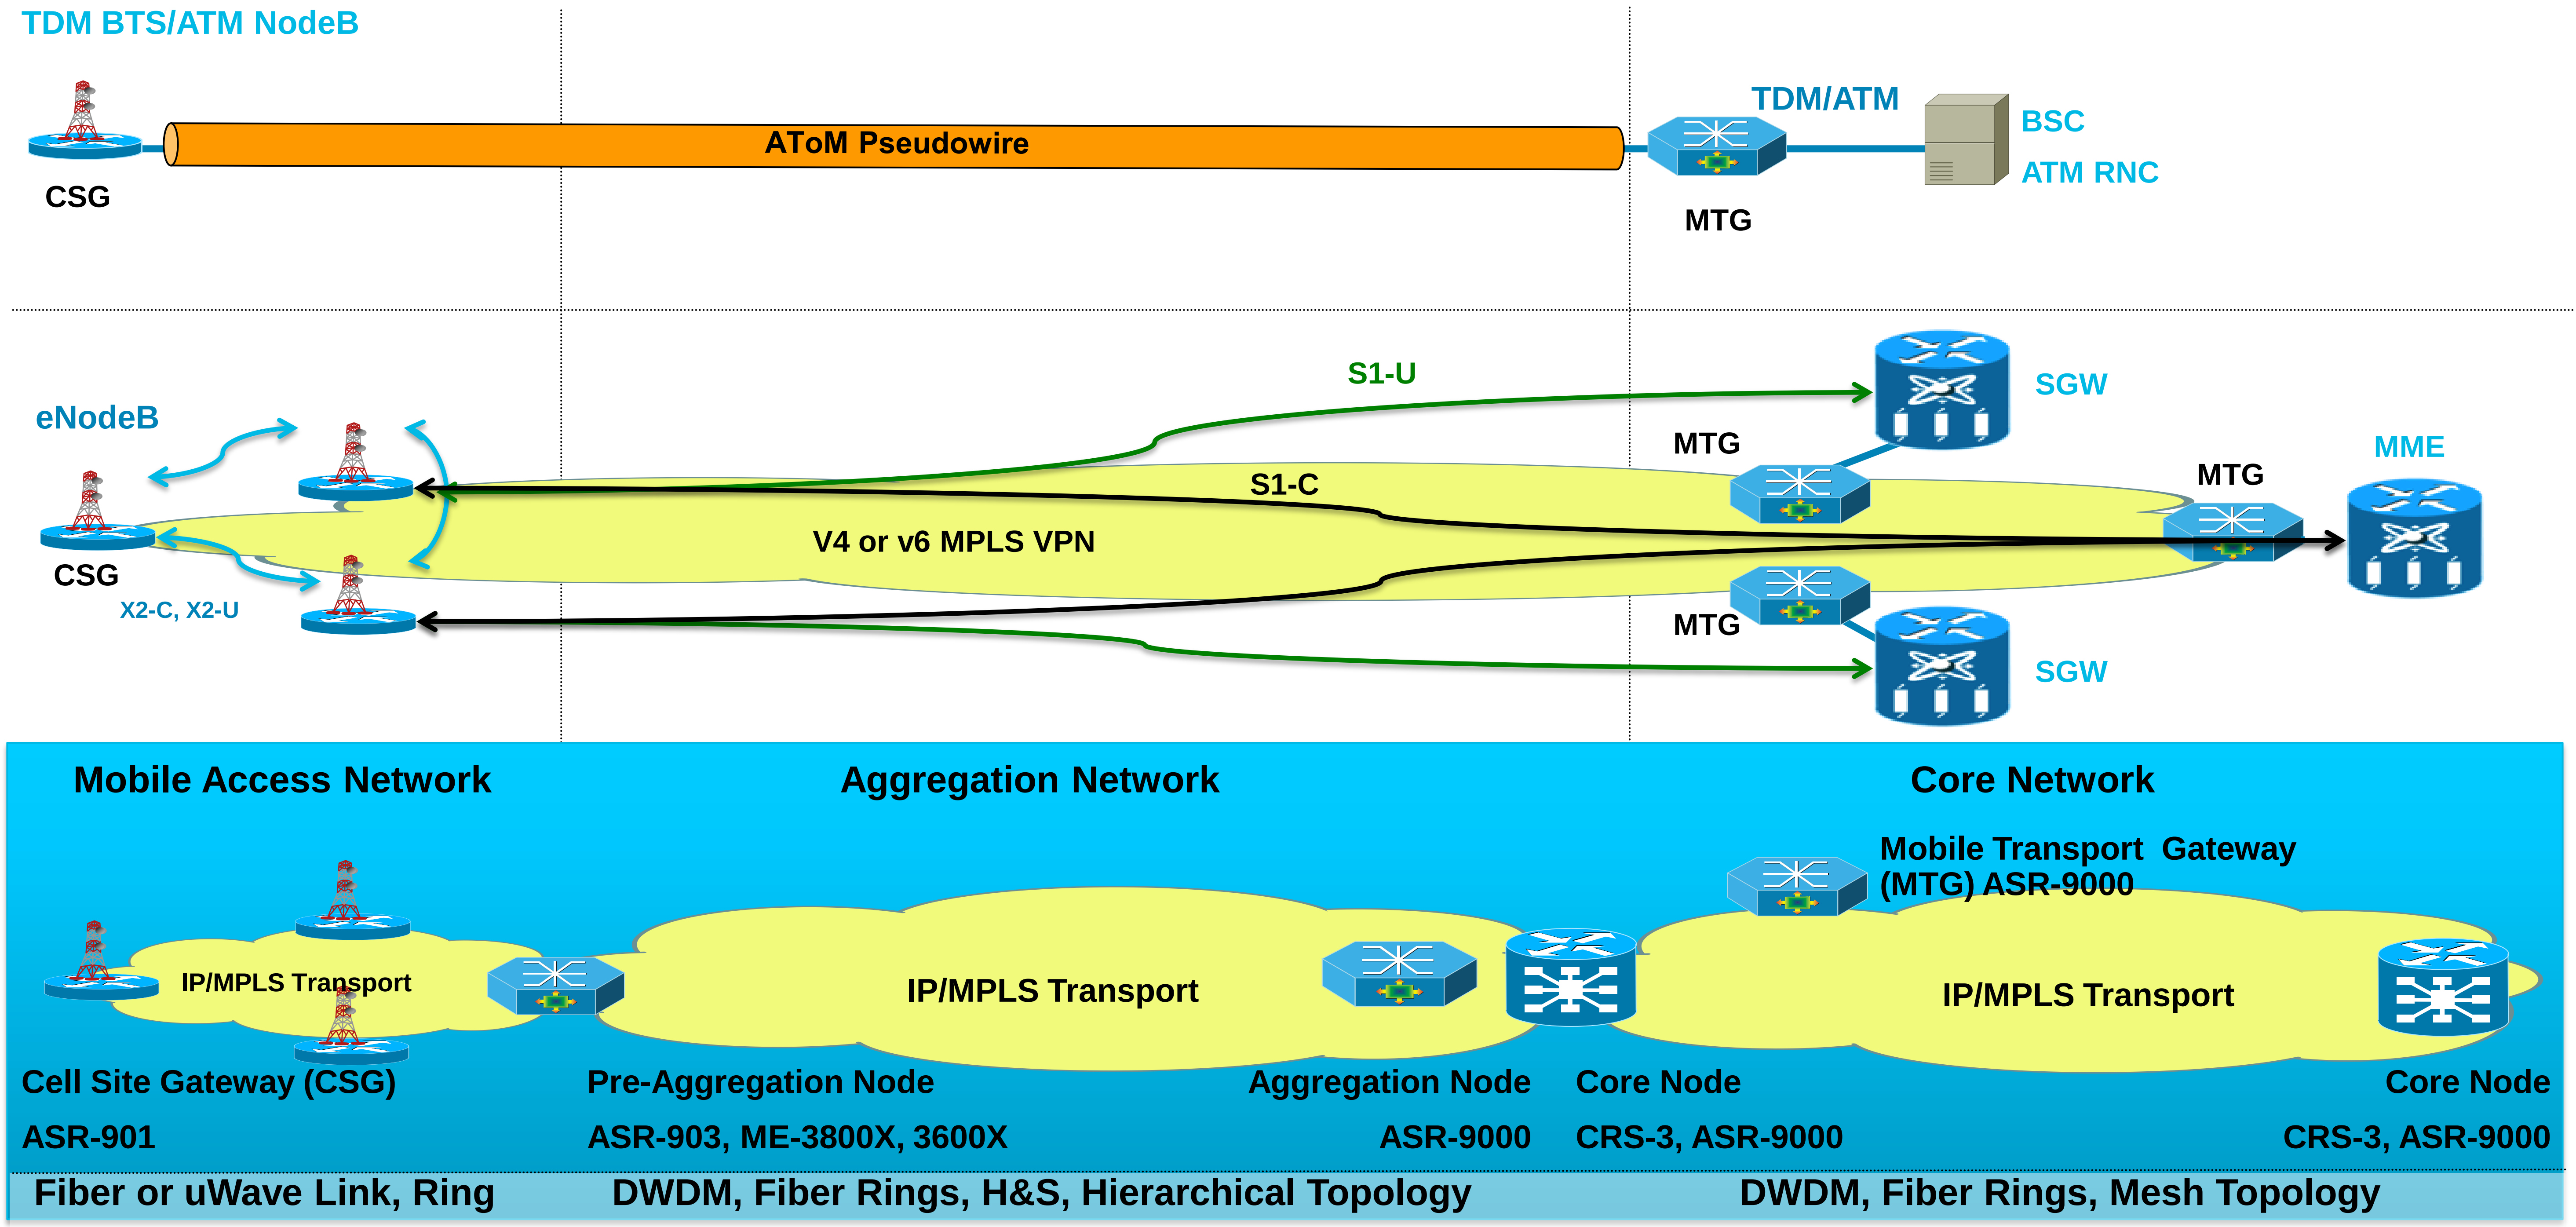

UMMT Platforms

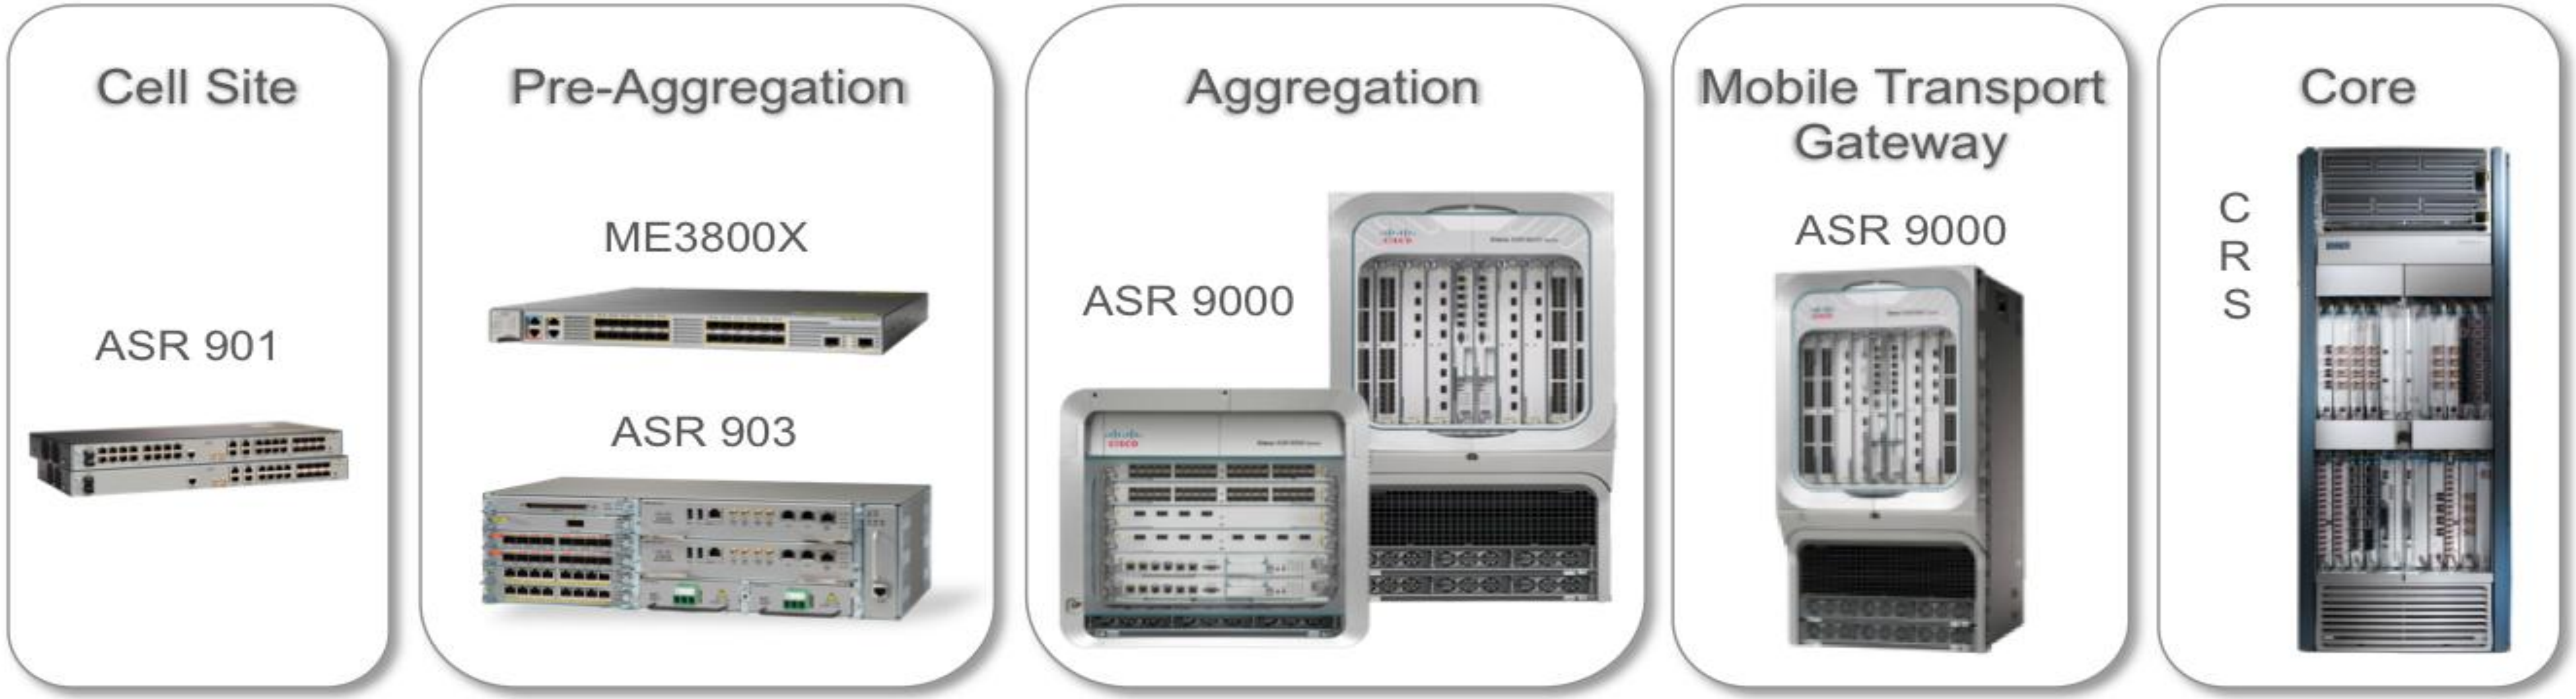

Unified MPLS Architecture Models

- Architecture Structuring based on
  - Access Type
  - Network Size

- Four ideal architecture models fitting various access options and network sizes

	Small Network (<1000 Nodes)	Large Network (>1000 Nodes)
Ethernet/TDM Access	**Model 1.1** Flat LDP Core and Aggregation Network	**Model 2.1** Hierarchical Labeled BGP Core and Aggregation Network
MPLS Access	**Model 1.2** Hierarchical Labeled BGP Access Network with 1.1	**Model 2.2** Hierarchical Labeled BGP Access Network with 2.1

1.1 Flat LDP LSP across Core and Aggregation

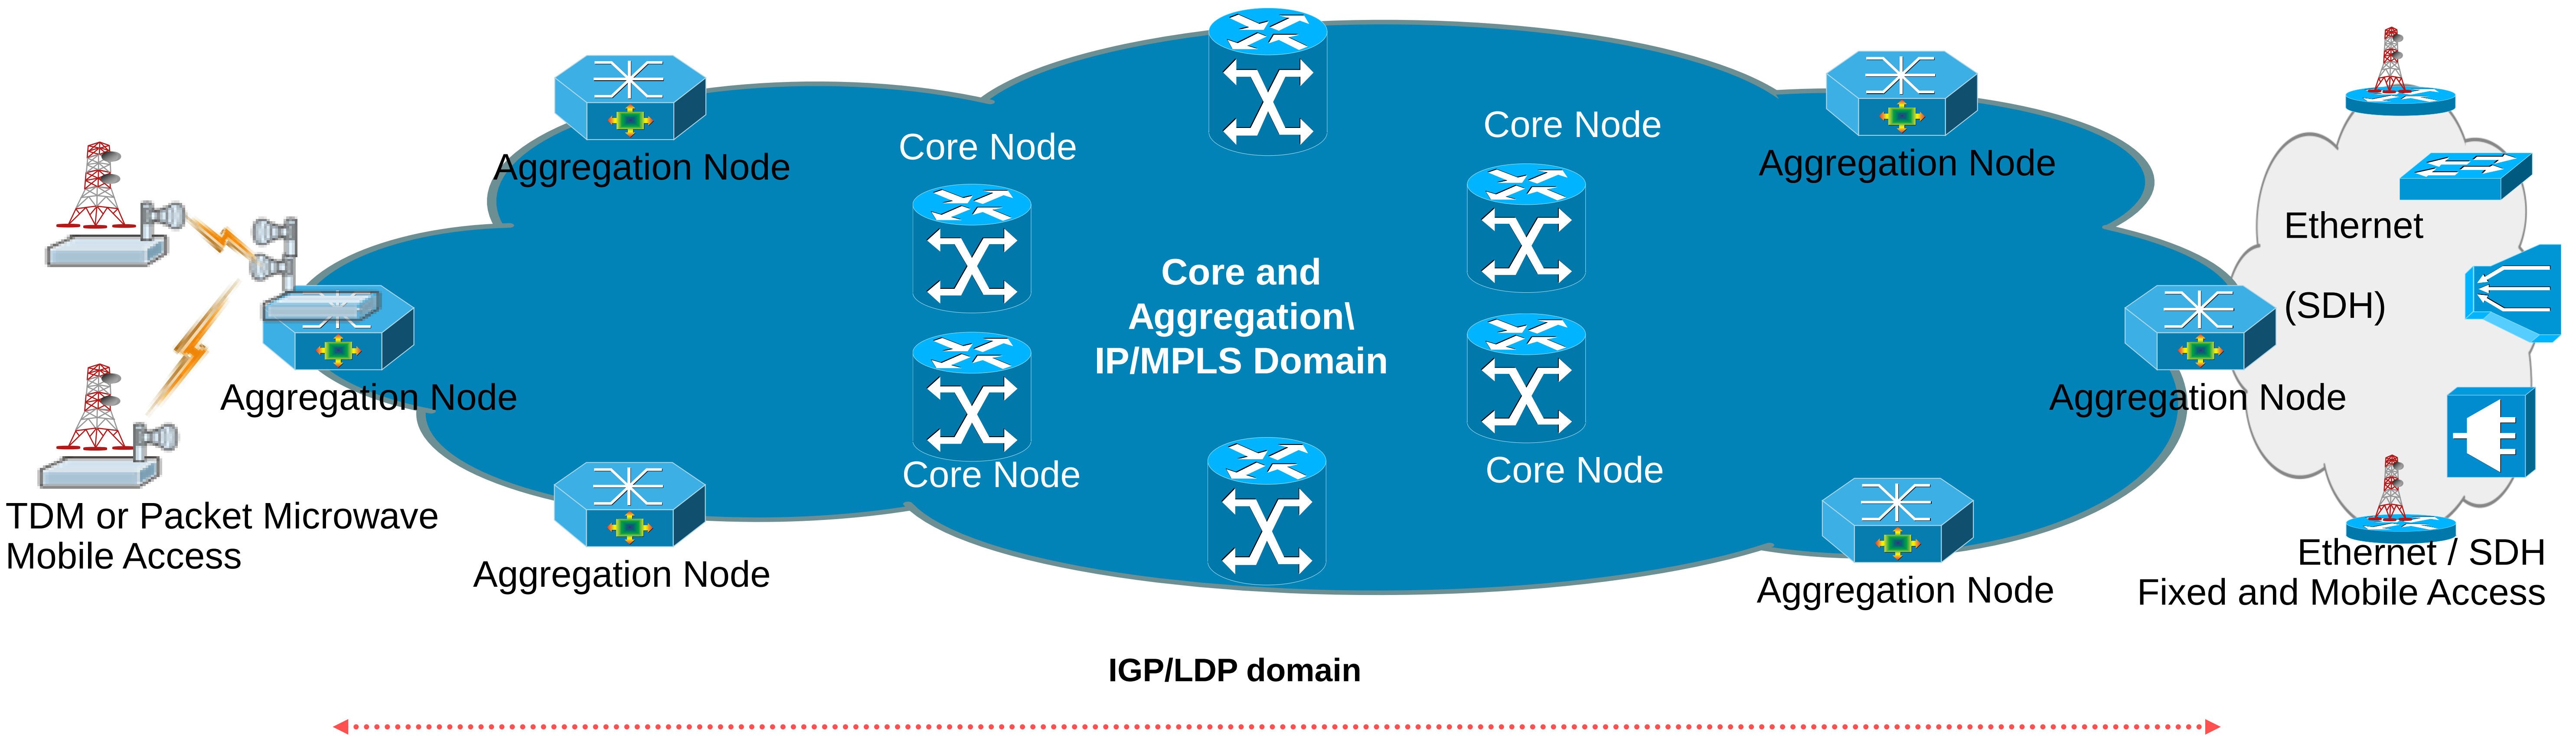

- Core and Aggregation Networks form one IGP and LDP domain.
  - Scale recommendation is less than 1000 IGP/LDP nodes for small aggregation platforms
- All Mobile and Wireline services enabled by the Aggregation Nodes. Access based on Ethernet, SDH or Packet/TDM Microwave links aggregated in Aggregation Nodes enabling TDM/ATM/Ethernet/IP interworking VPWS and MPLS VPN transport

1.2 Hierarchical BGP LSP Across Core + Aggregation and Access Networks

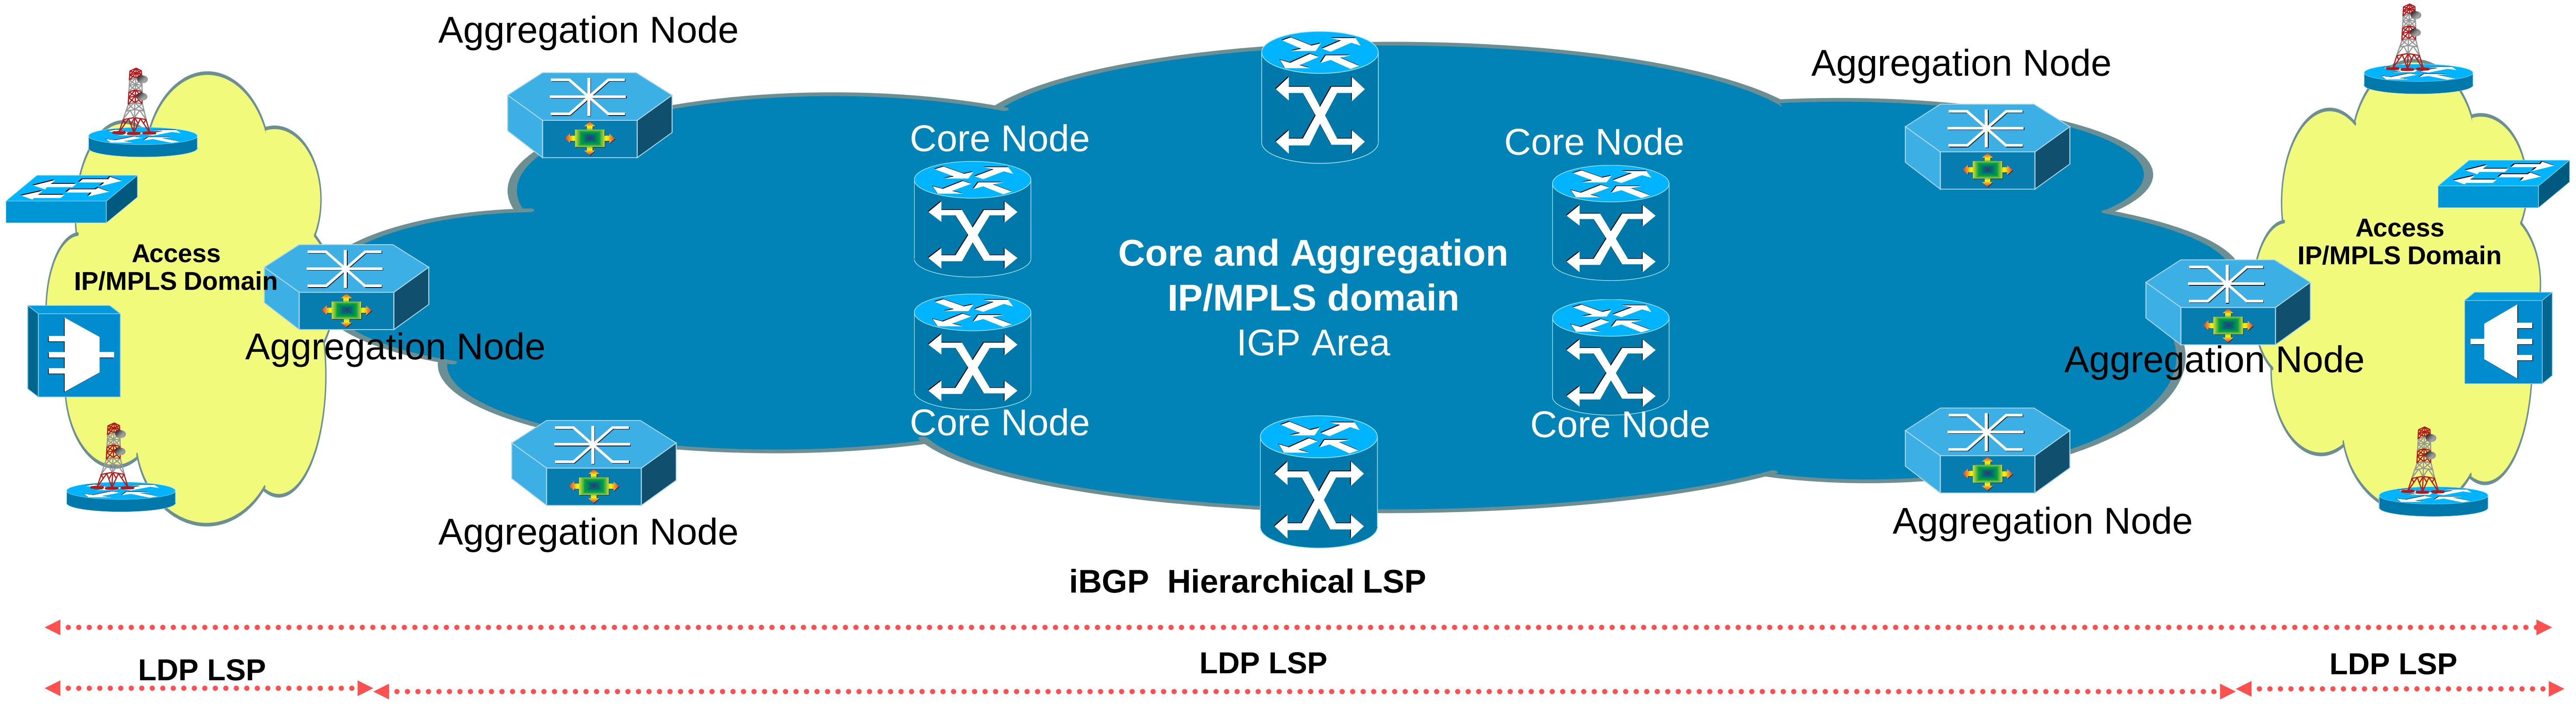

- Core and Aggregation form a relatively small IGP/LDP domain (1000 nodes)
- Access is MPLS enabled. Each Access Network forms a different IGP/LDP domain
- Core/Aggregation and RAN Access Networks are integrated with labeled BGP LSP
- Access Nodes learn only required service destinations based on inbound or outbound labeled BGP filtering (done in Access or Aggregation Node)

2.1 Hierarchical BGP LSP Across Core and Aggregation Networks Domains

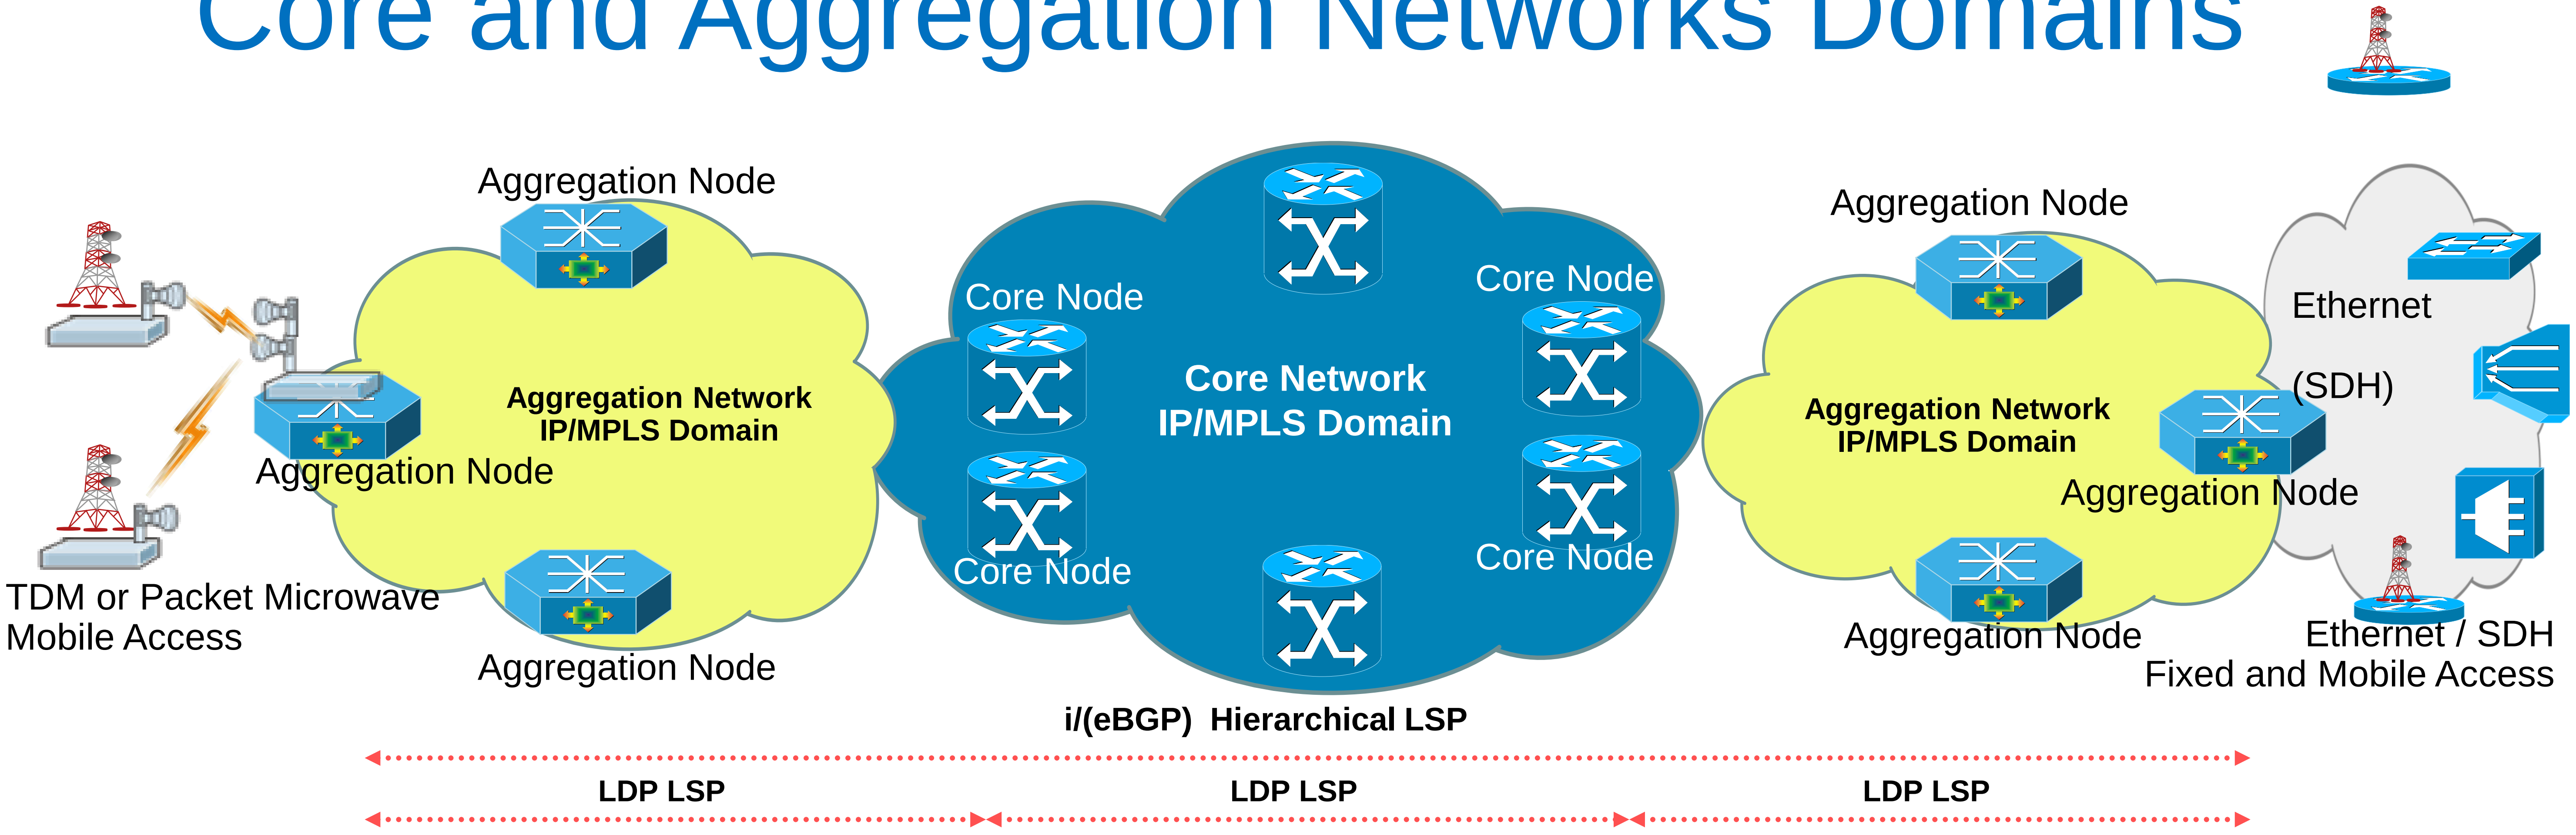

- Core and Aggregation Networks enable Unified MPLS Transport
- Core and Aggregation Networks organized as independent IGP/LDP domains
- Network domains interconnected by hierarchical LSPs based on RFC 3107, BGP IPv4+labels.
  - Intra domain connectivity is based on LDP LSPs
- Aggregation Nodes enable Mobile and Wireline Services. Access Network may be Ethernet, SDH, Packet/TDM Microwave

Model 2.2

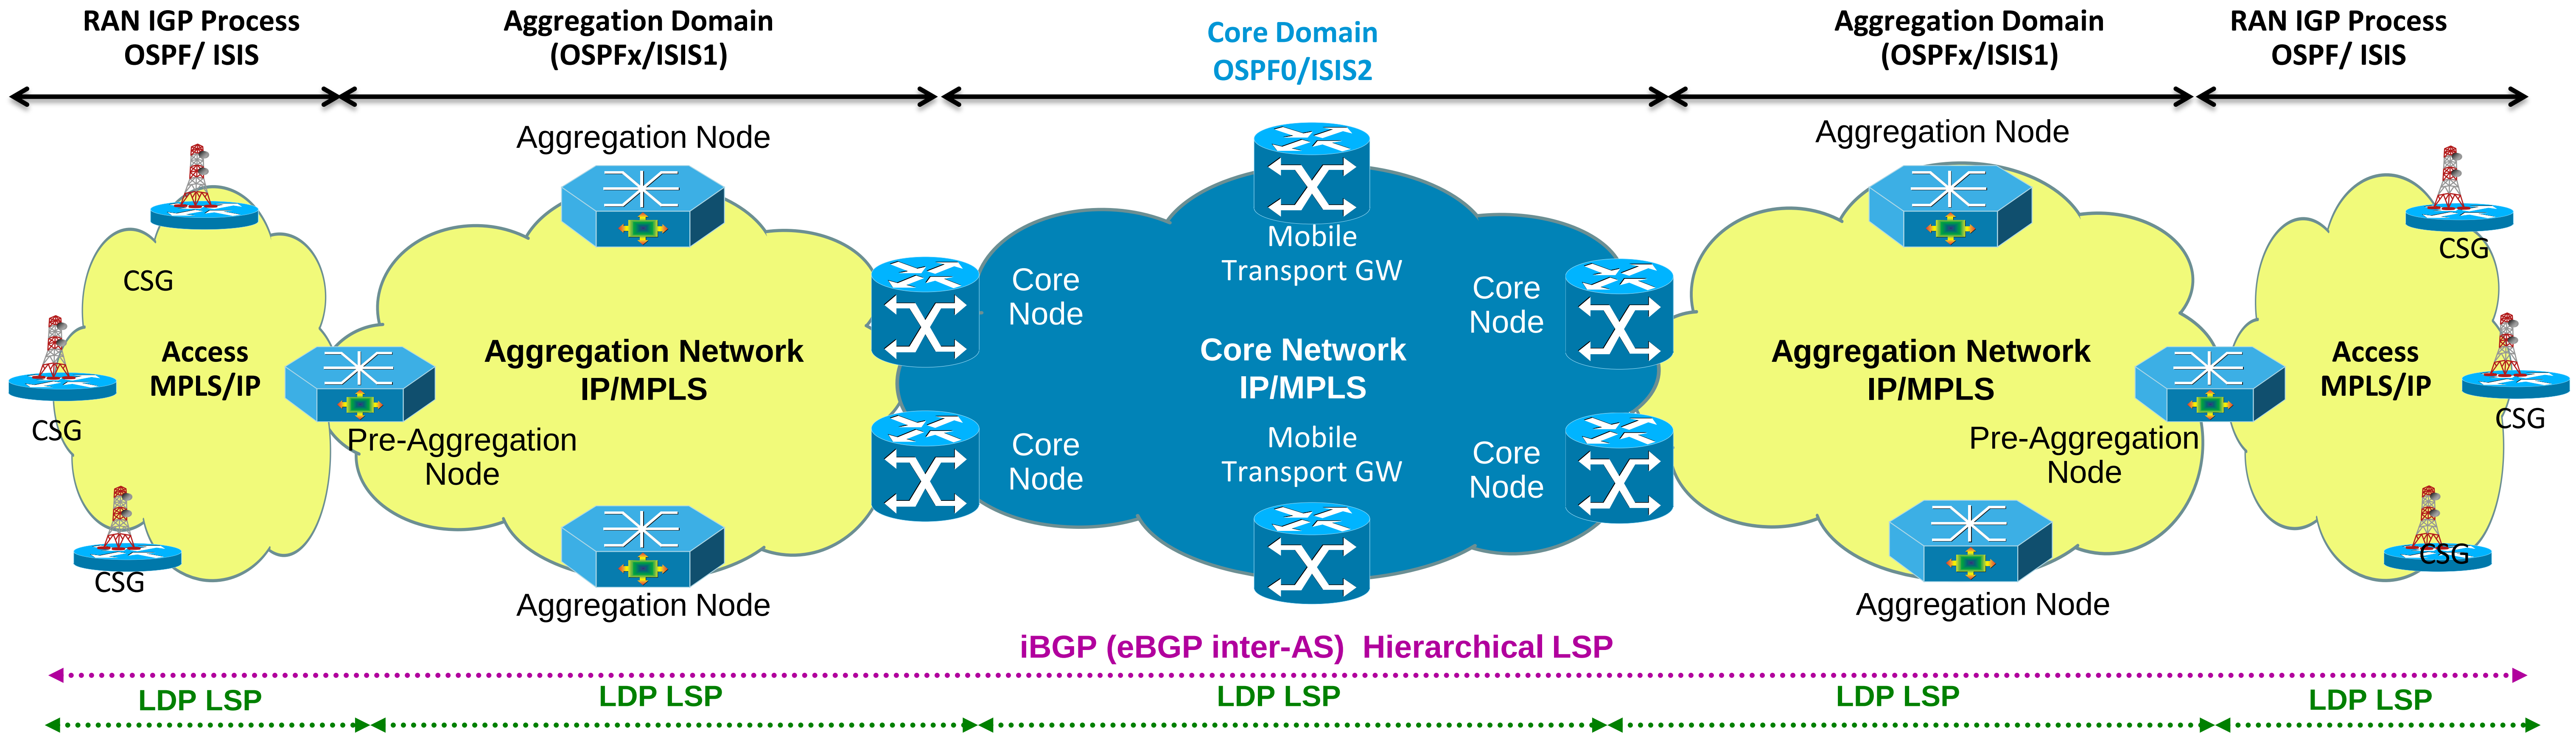

- Core, Aggregation, and Access partitioned as independent IGP/LDP domains.
  - Reduce size of routing & forwarding tables on routers to enable better stability & faster convergence.
  - LDP used to build intra-domain LSPs
- BGP labeled unicast (RFC 3107) used as inter-domain label distribution protocol to build hierarchical LSPs across domains
  - Single AS multi-area topologies use labeled iBGP to build inter-domain LSPs
  - Inter-AS topologies use labeled eBGP between ASes and labeled iBGP within the AS to build inter-domain LSPs
- Inter-domain Core/Aggregation LSPs are extended to Access with controlled redistribution between access and aggregation domains based on route targets.
  - Low route scale in RAN Access is preserved: Only local RAN IGP prefixes + few Mobile Packet Core loopbacks
- Intra-domain link & node failures protected by LFA FRR and Fast IGP convergence. ABR failures protected by BGP PIC

BGP Control Plane

Single AS, Multi Area IGP

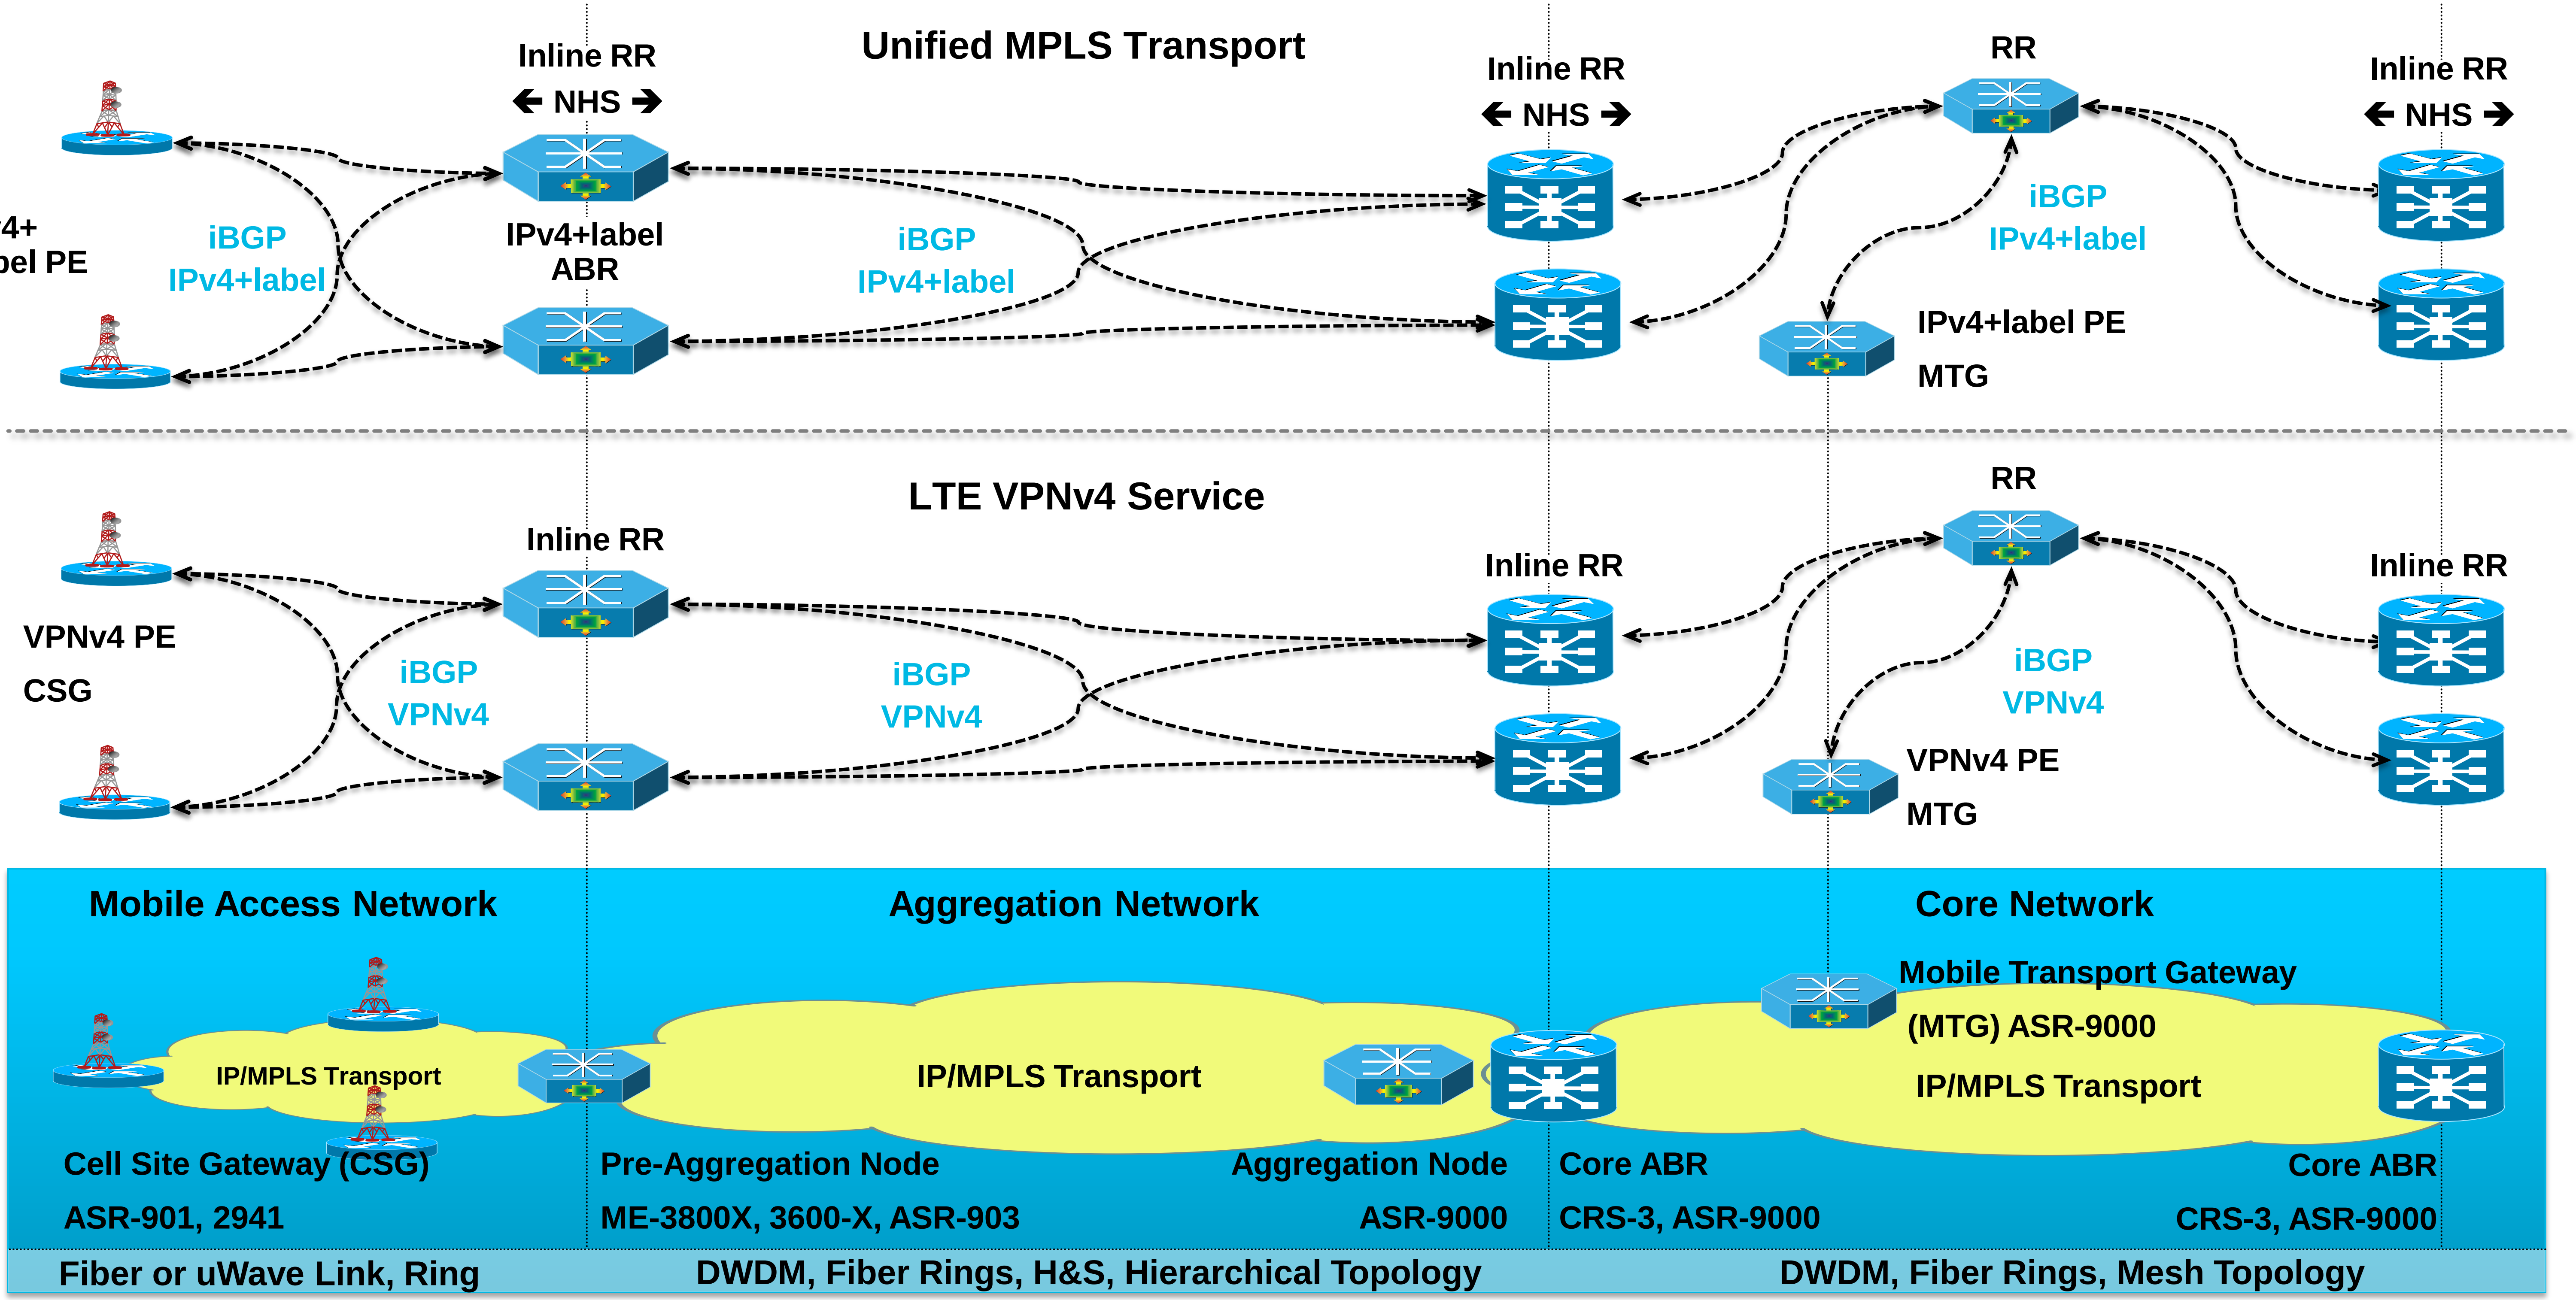

BGP Control Plane

Inter-AS with Centralized RR

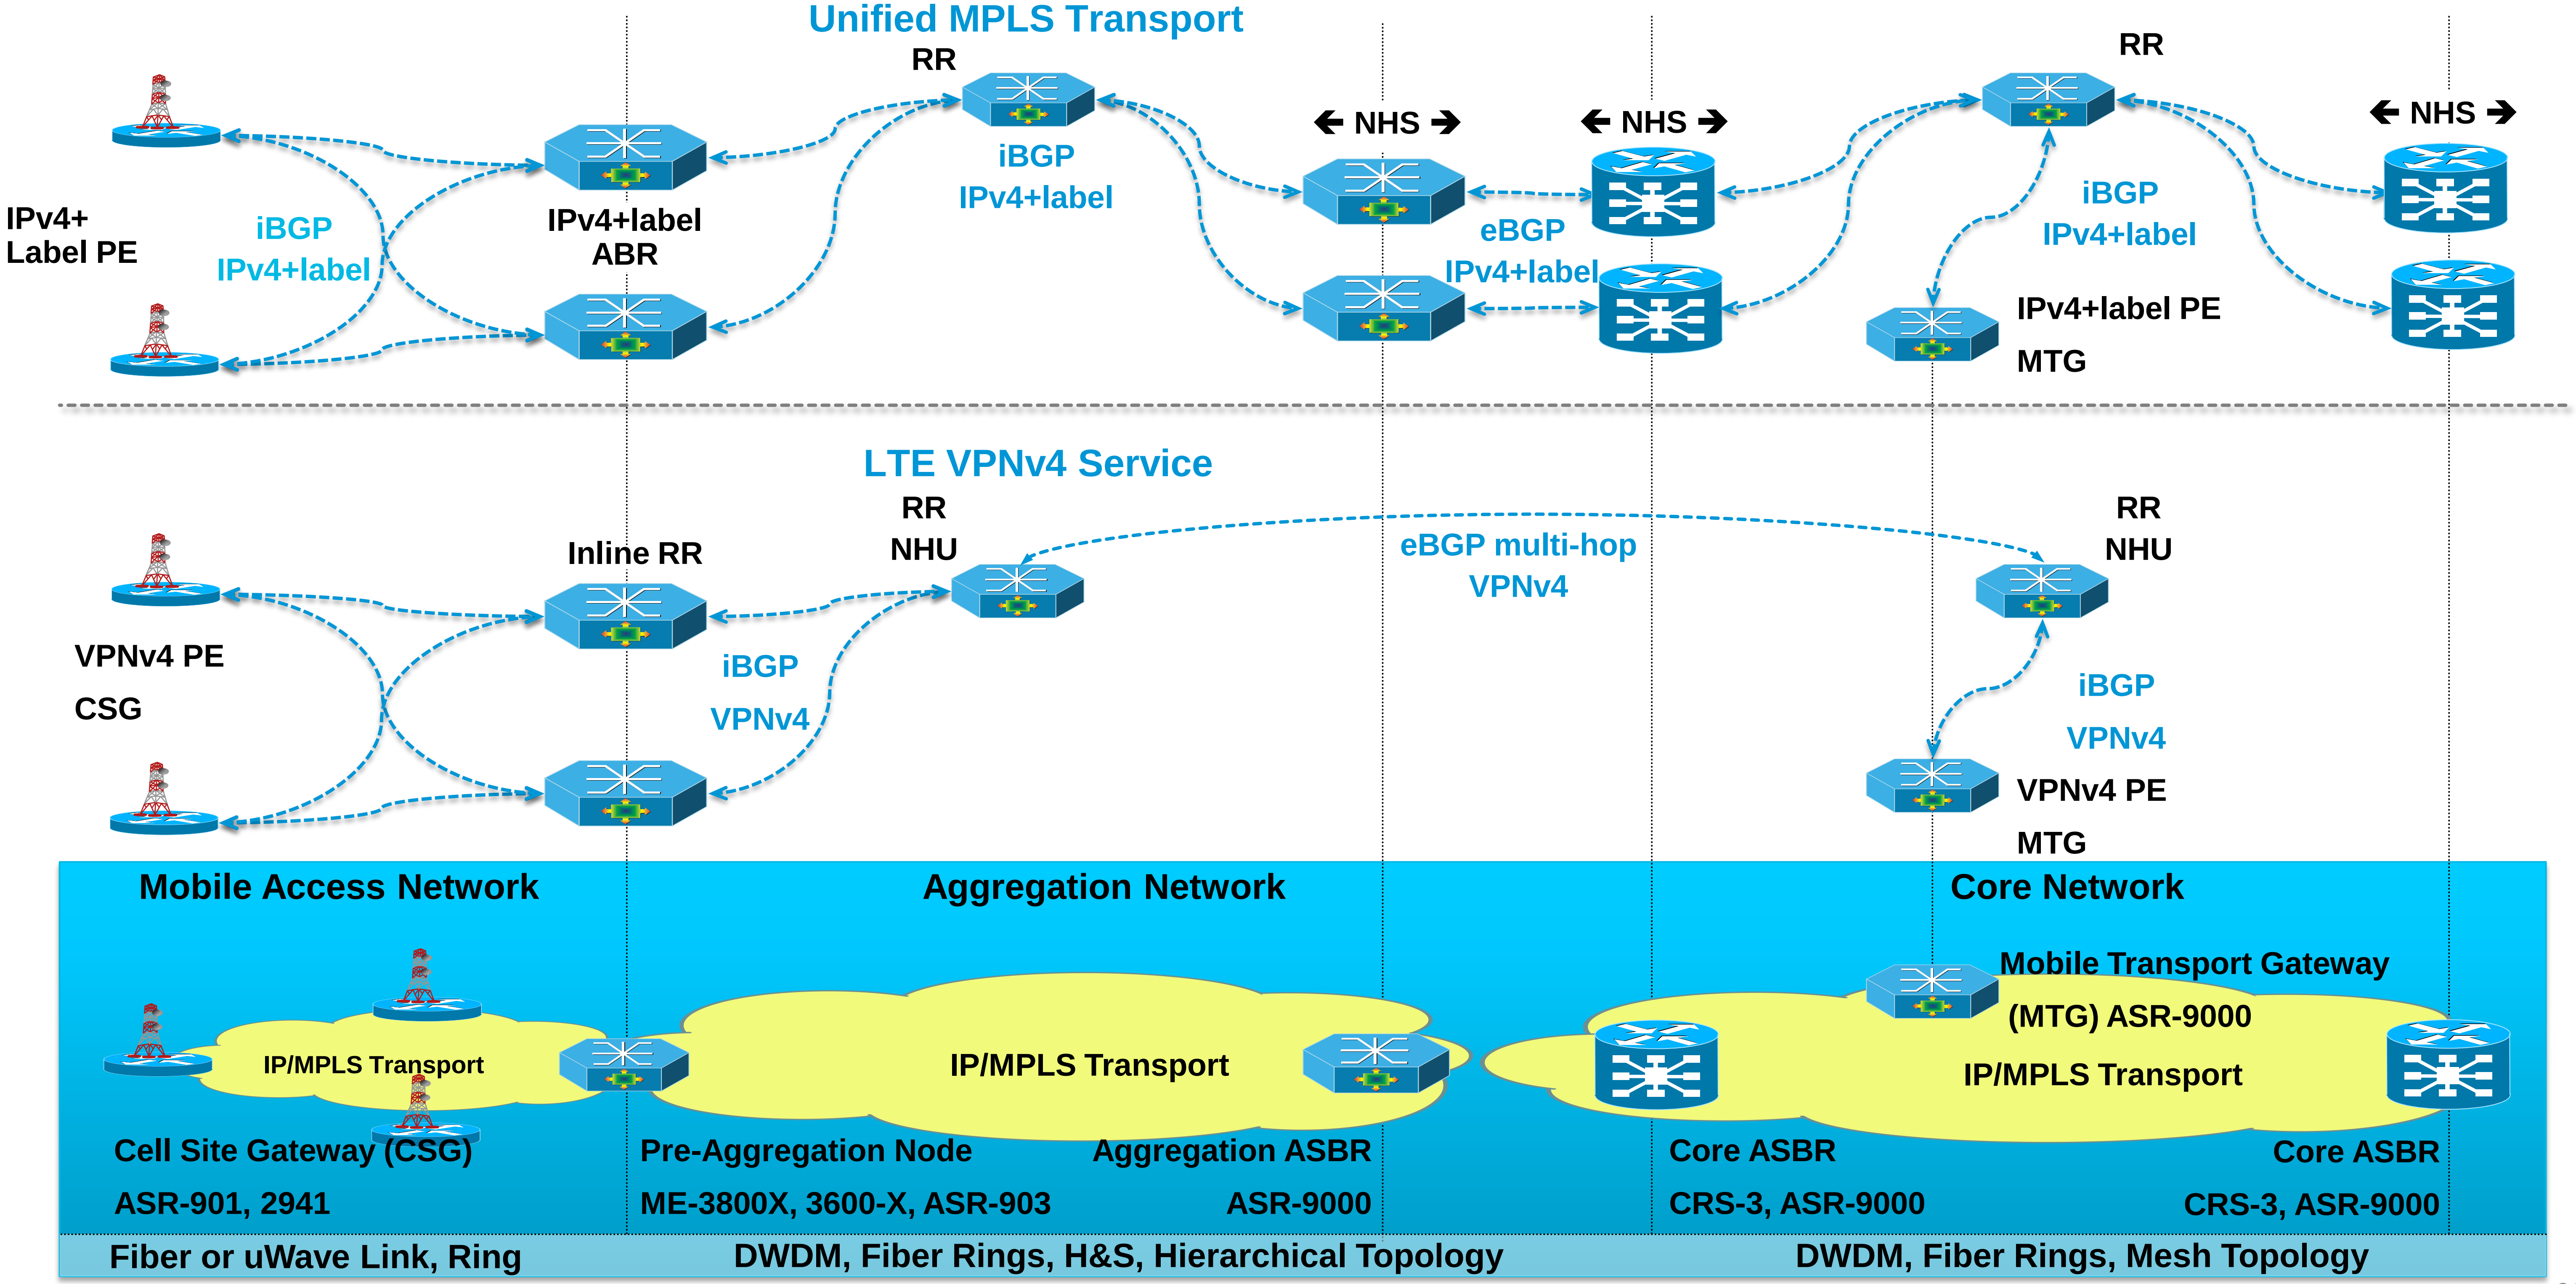

Mobile Service Scalability Control

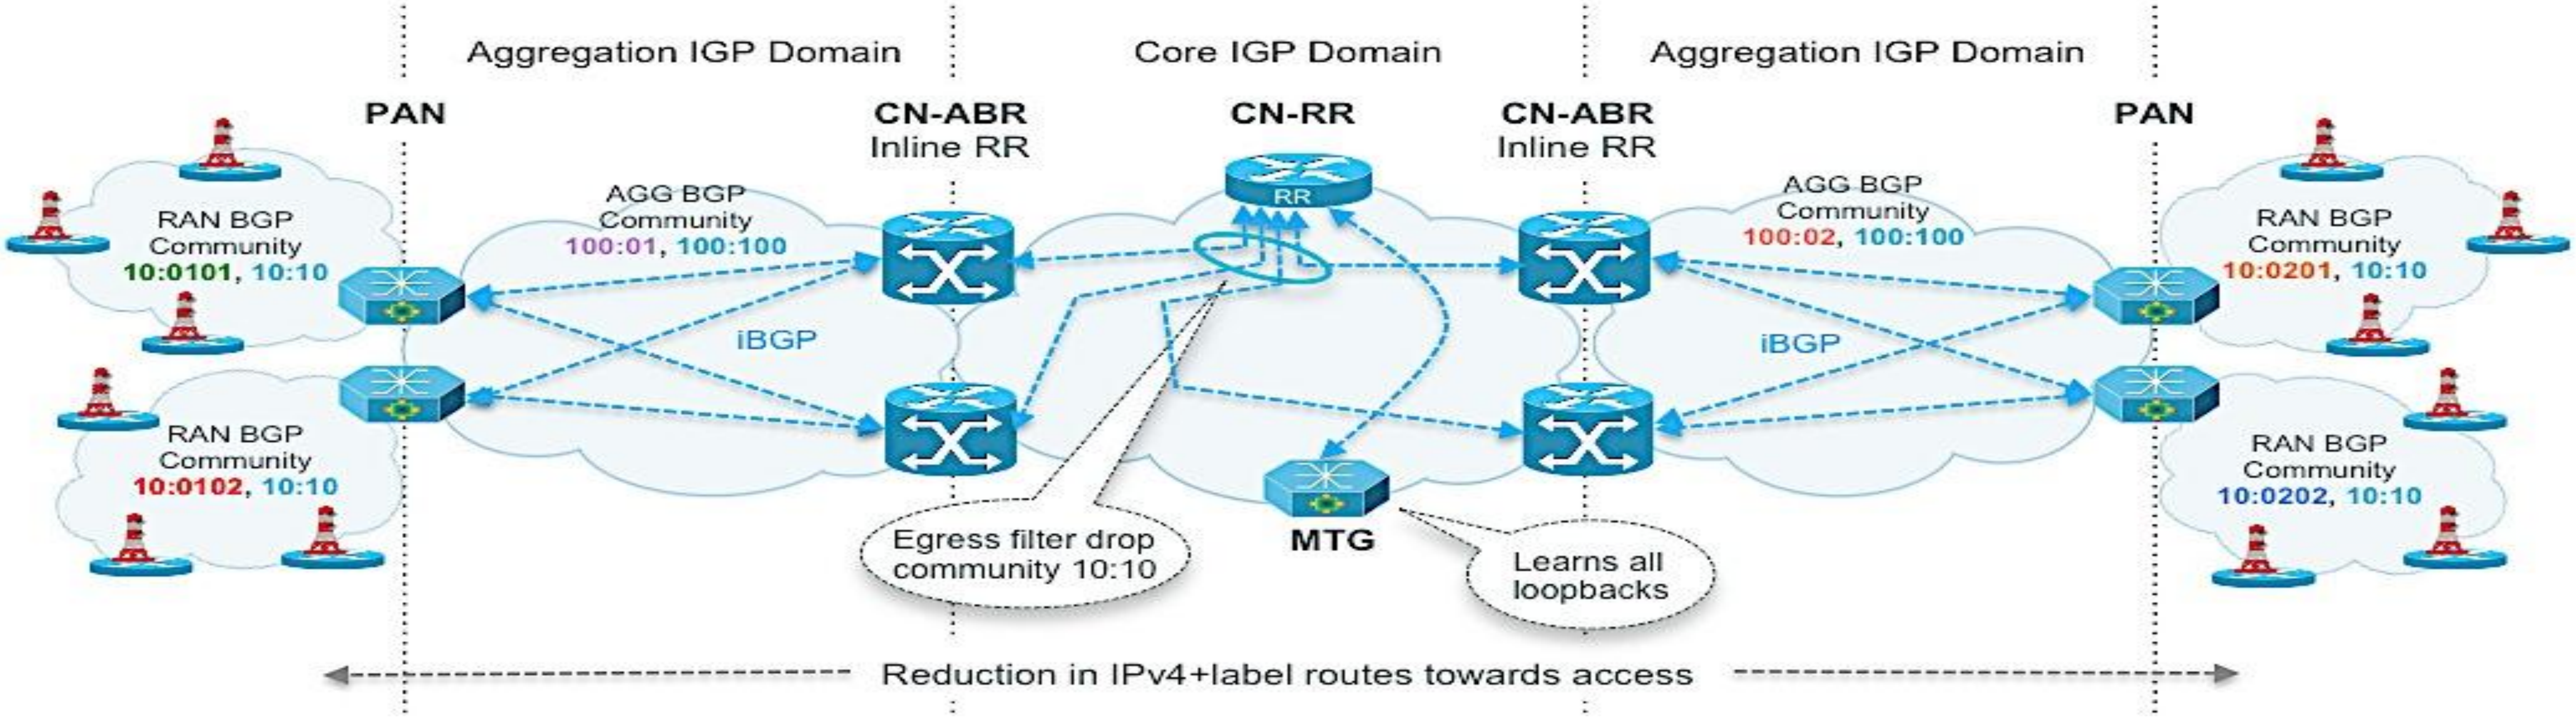

- Only the MPC community is distributed to RAN access
- The RAN Common Community is only distributed to MTGs

LTE S1 and X2
MPLS VPN Service and Scalability Control

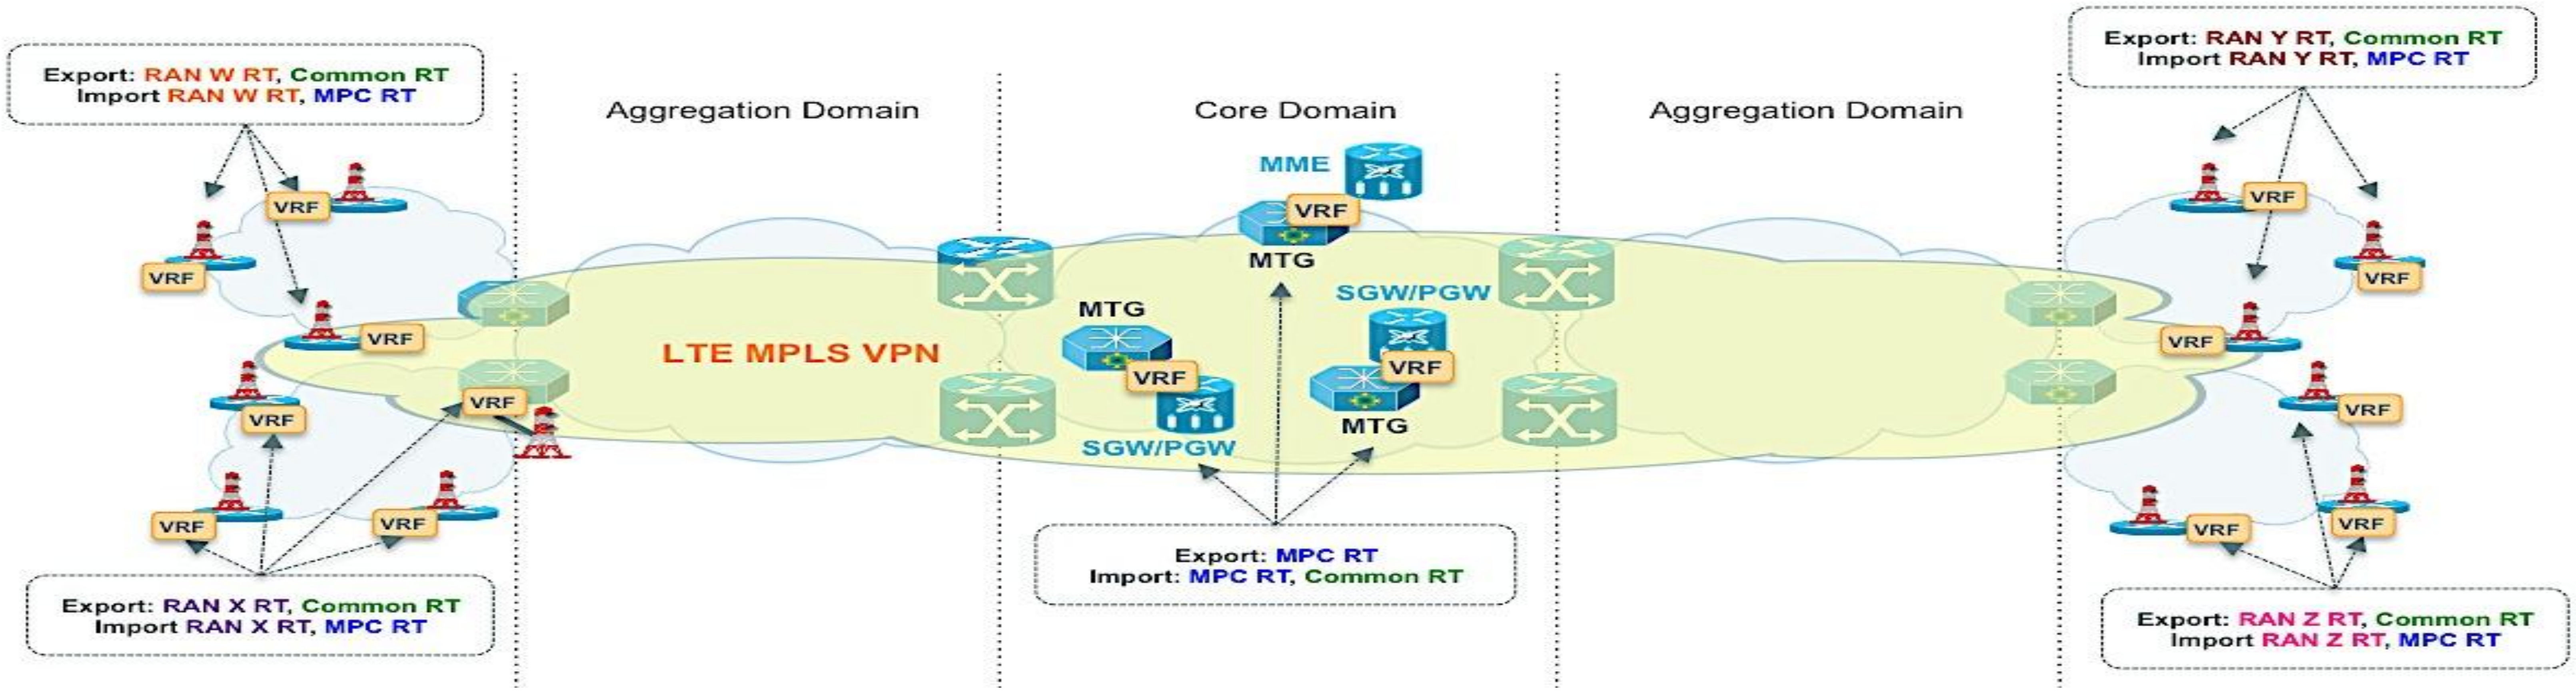

- Mobile Transport GWs import all RAN & MPC Route Targets, and export prefixes with MPC Route Target
- CSGs (and/or Pre-Aggregation Node) in a RAN region import the MPC and regional RAN Route Targets:
  - Enables S1 control and user plane with any MPC locations in the core
  - Enables X2 across CSGs in the RAN region
- MPLS VPN availability based on BGP PIC Edge and infrastructure LSP based LFA FRR
- Pre-Aggregation Nodes and Core POP Nodes form inline RR hierarchy for the MPLS VPN service
  - Core ABRs perform BGP community based Egress filtering to drop unwanted remote RAN VPNv4 prefixes
  - Pre-Aggregation Nodes implement RT Constrained Route Distribution towards CSR VPNv4 clients

LTE X2 Inter Access Network

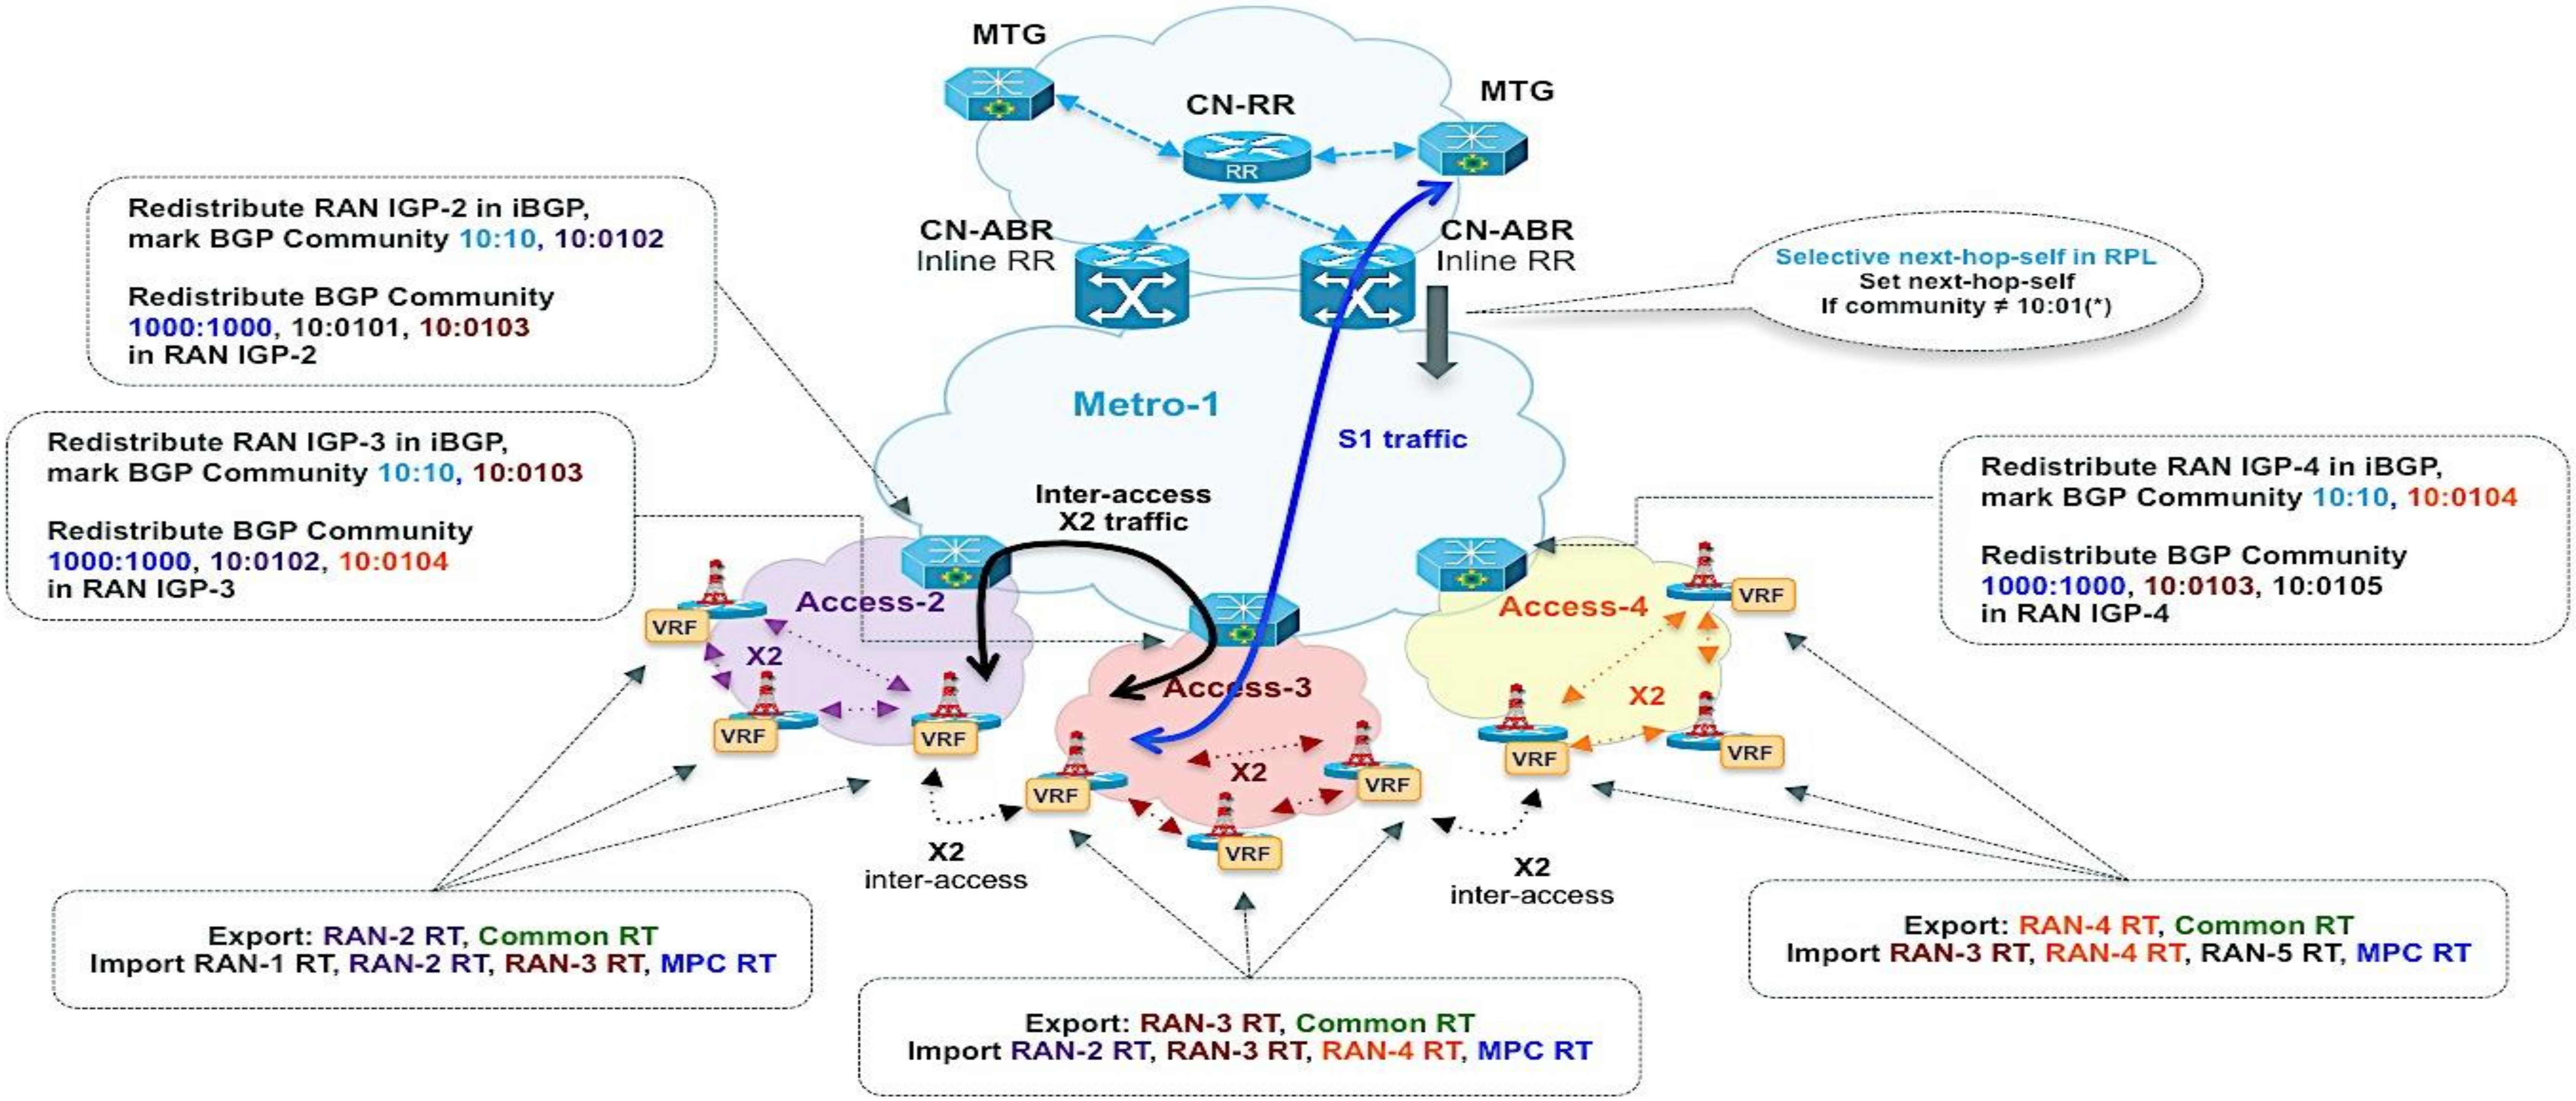

E-LINE Service and Scalability Control

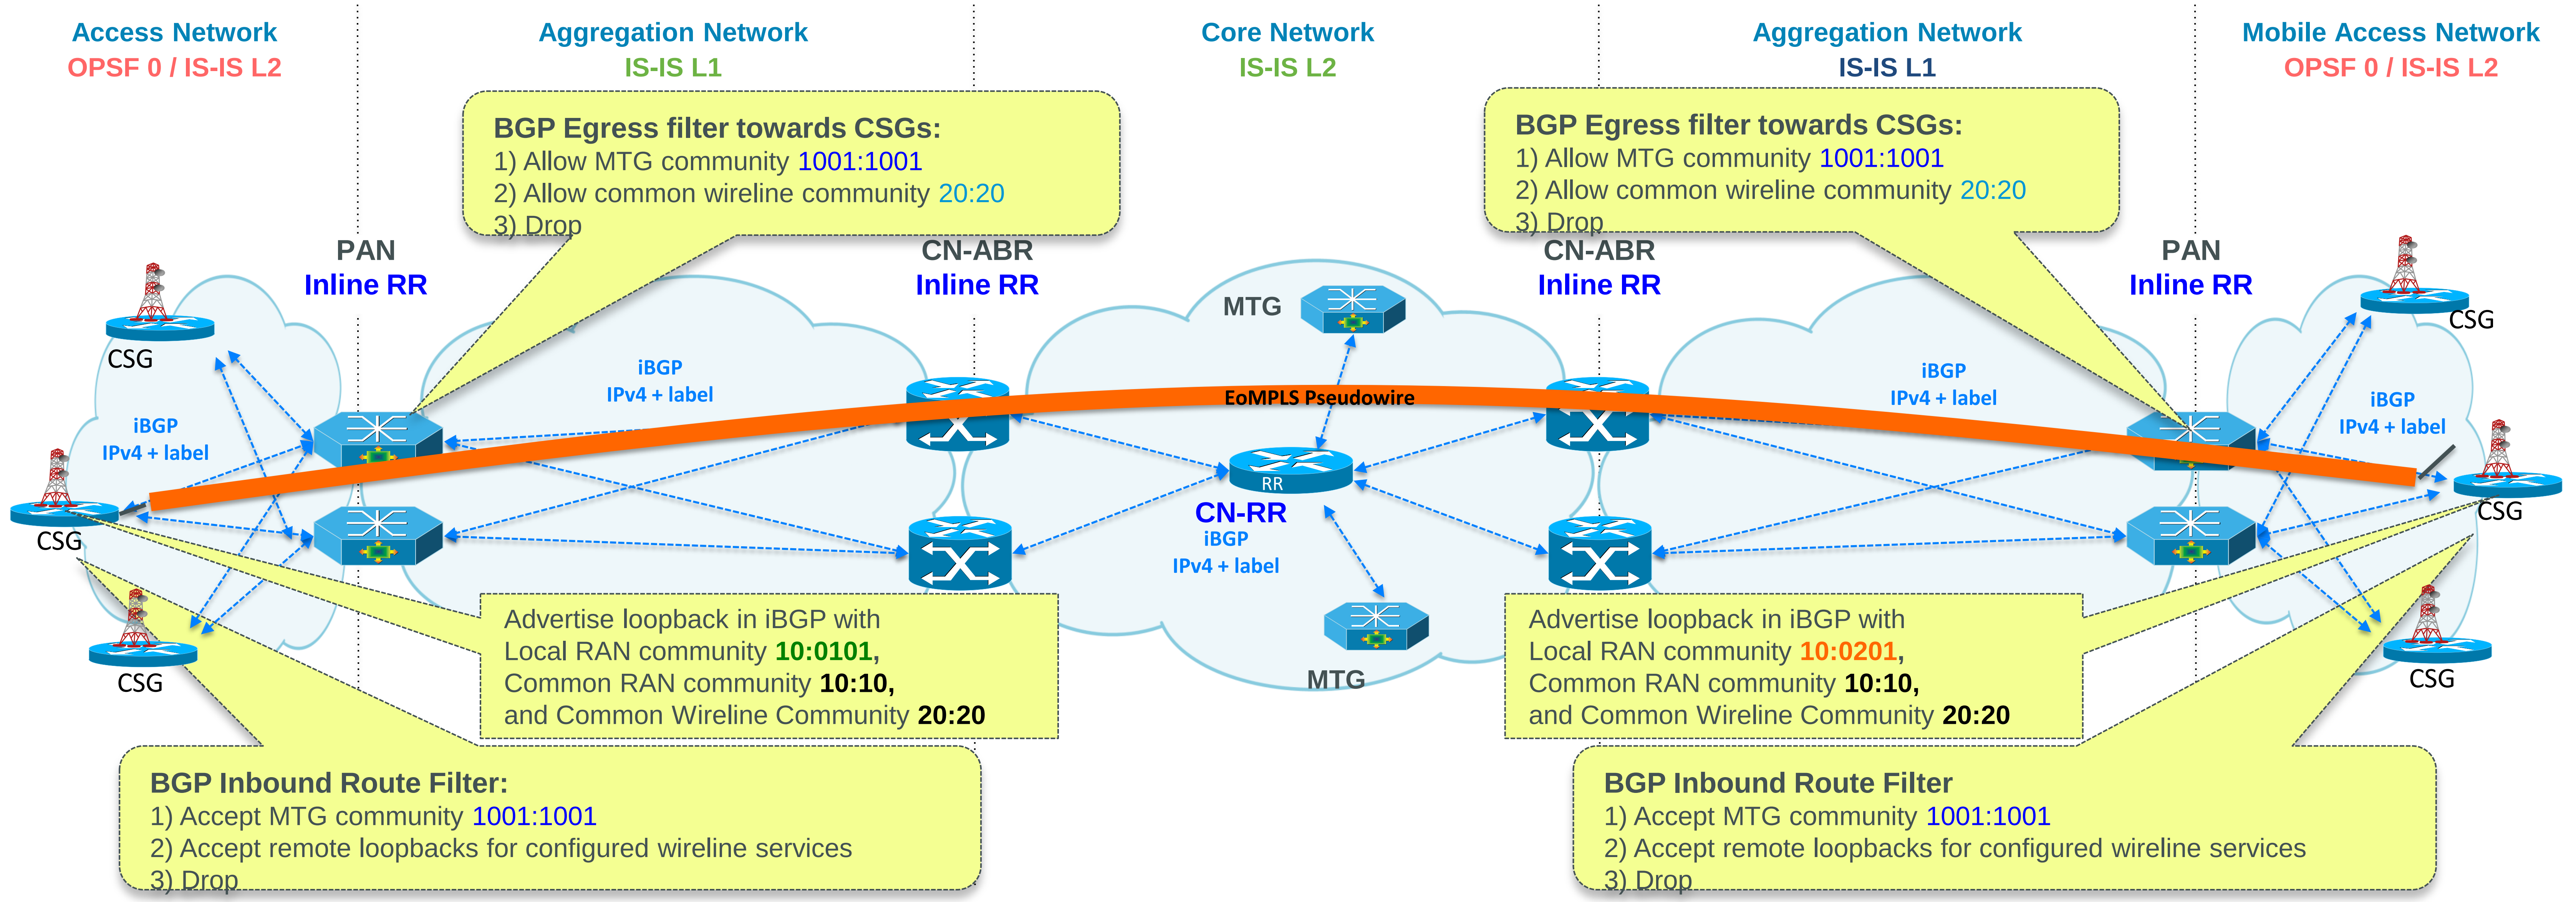

- MTG and Common wireline communities distributed to RAN access
- Common RAN Community is only distributed to MTGs
- CSG accepts MTG & remote loopbacks for configured wireline services

UMMT Functional Aspects
QoS

LTE QoS Concepts

- QoS Class Identifier (QCI): Scalar that controls bearer level QoS treatment
- Guaranteed Bit Rate (GBR): Bit rate that a GBR bearer is expected to provide
- Maximum Bit Rate (MBR): Limits the bit rate a GBR bearer is expected to provide
- Allocation and Retention Policy (ARP): Controls how a bearer establishment or modification request can be accepted when resources are constrained.
- IP pkts mapped to same EPS bearer receive same treatment. User IP packets are filtered onto appropriate bearer by TFTs.
- Default bearer always established with QoS values assigned by MME from data retrieved from HSS.

LTE QCI Mapping

QCI Value	Resource Type	Priority	Delay Budget [1]	Error Loss Rate [2]	Example Services
1 [3]		2	100 ms	10^{-2}	Conversational Voice
2 [3]	GBR	4	150 ms	10^{-3}	Conversational Video (Live Streaming)
3 [3]		3	50 ms	10^{-3}	Real Time Gaming
4 [3]		5	300 ms	10^{-6}	Non-Conversational Video (Buffered Streaming)
5 [3]		1	100 ms	10^{-6}	IMS Signalling
6 [4]		6	300 ms	10^{-6}	Video (Buffered Streaming) TCP-based (e.g., www, e-mail, chat, ftp, p2p file sharing, progressive video, etc.)
7 [3]	Non-GBR	7	100 ms	10^{-3}	Voice, Video (Live Streaming), Interactive Gaming
8 [5]		8	300 ms	10^{-6}	Video (Buffered Streaming) TCP-based (e.g., www, e-mail, chat, ftp, p2p sharing, progressive download, etc.)
9 [6]		9			

Quality of Service

- Core, Aggregation, Access
  – Differentiated Services QoS, MPLS EXP classification, diffserv queuing
- Microwave Access Networks
  – Hierarchical QOS with parent aggregate PIR/shape to the microwave speed, child differentiated services QoS, MPLS EXP classification, diffserv queuing
- Differentiated Services Class Mappings:
  – UNI classification based on IP DSCP or Ethernet CoS, or ATM, TDM mapping to a QoS group.
  – Generally the diffserv mapping is done in base stations, radio controllers and gateways

DiffServ QOS Domain

Traffic Class	LTE QCI	Resource	DiffServ PHB	Core, Aggregation, Access Network	Mobile Access UNI	
				MPLS/IP	IP NodeB, eNodeB	ATM NodeB
				MPLS EXP	DSCP	ATM
Network Management	7	Non-GBR	AF	7	56	VBR-nrt
Network Control Protocols	6	Non-GBR	AF	6	48	VBR-nrt
Network Sync (1588 PTP , ACR) Mobile Conversation (Voice & Video) Signaling (GSM Abis, UMTS Iub control, LTE S1c, X2c)	1 2 3	GBR	EF	5	46	CBR
Reserved	4	-	AF	4	32	VBR-nrt
Hosted Video	5	Non-GBR	AF	3	24	VBR-nrt
Reserved	8	-	AF	2 1	16 8	VBR-nrt
Internet Best Effort	9	Non-GBR	BE	0	0	UBR

UMMT Functional Aspects
Resiliency

High Availability

- Unified MPLS Transport:
  - Intra-domain LSP: IGP/LDP convergence, remote LFA FRR
  - Inter-domain LSP: BGP convergence, BGP PIC Core and Edge
- MPLS VPN Service (ASR-901, ASR-9000):
  - eNB UNI: static routes
  - MPC UNI: PE-CE dynamic routing with BFD keep-alive
  - Transport: BGP VPNv4 convergence, BGP PIC Core and Edge
- VPWS Service:
  - UNI: mLACP for Ethernet, MR-APS for TDM/ATM
  - Transport: PW redundancy, two-way PW redundancy
- Synchronization Distribution:
  - ESMC for SyncE, SSM for ring distribution.
  - 1588 BC with active/standby PTP streams from multiple 1588 OC masters

UMMT High Availability Overview

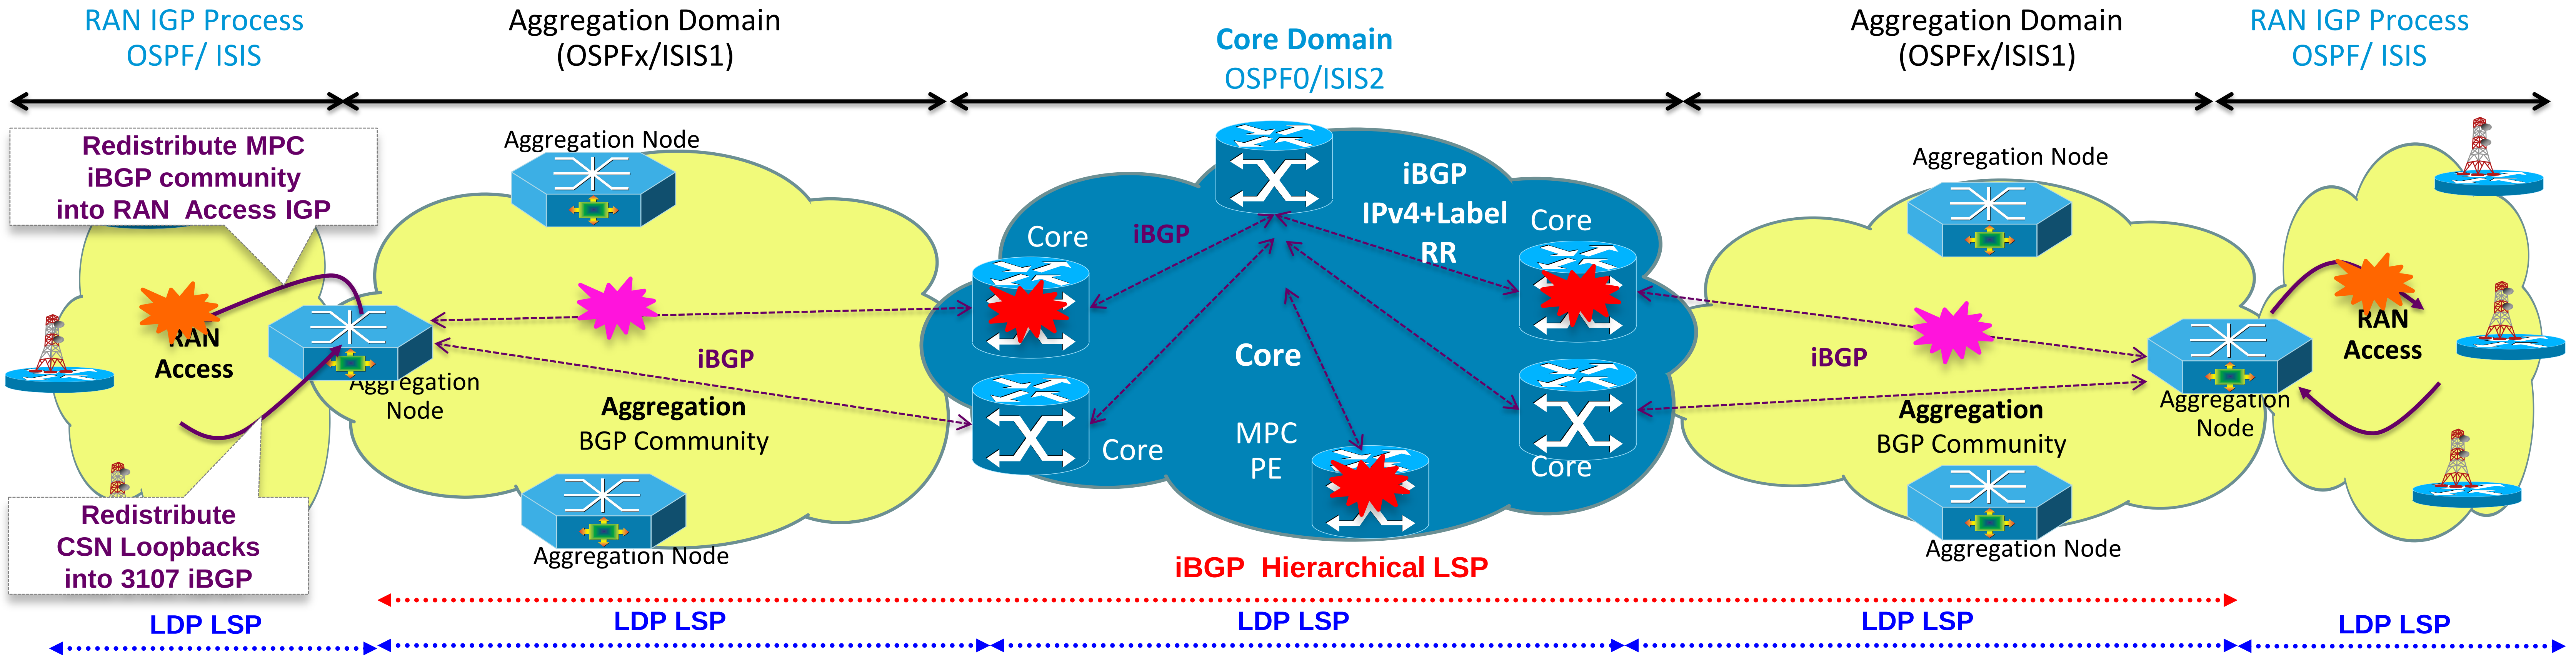

LFA L3 convergence < 50ms

BGP PIC Core L3 convergence < 100ms

BGP PIC Edge L3 convergence < 100ms

LFA FRR
Loop Free Alternate Fast Reroute

- What is LFA FRR?
  - Well known (RFC 5286) basic fast re-route mechanism to provide local protection for unicast traffic in pure IP and MPLS/LDP networks
  - Path computation done only at “source” node
  - Backup is Loop Free Alternate (C is an LFA, E is not)
- No directly connected Loop Free Alternates (LFA) in some topologies
- Ring topologies for example:
  - Consider C1-C2 link failure
  - If C2 sends a A1-destined packet to C3, C3 will send it back to C2
- However, a non-directly connected loop free alternate node (C5) exits

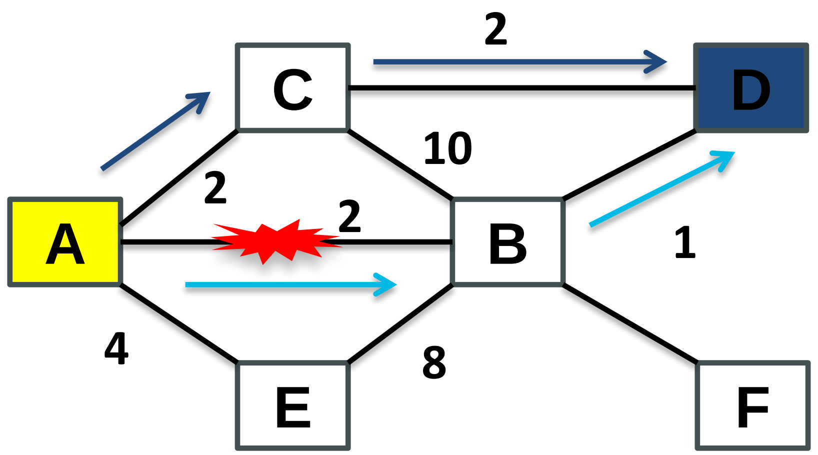

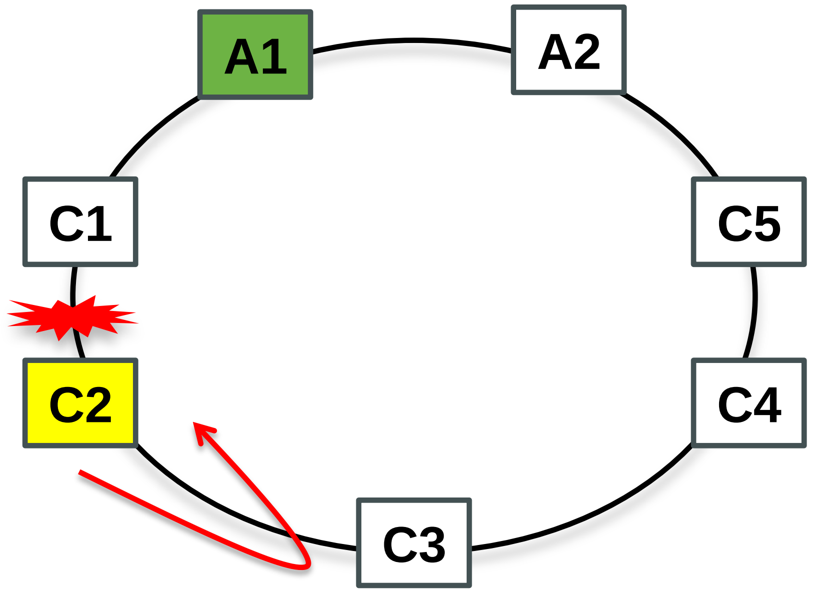

Remote LFA FRR

http://tools.ietf.org/html/draft-shand-remote-lfa

- Remote LFA uses automated IGP/LDP behavior to extend basic LFA FRR to arbitrary topologies
- A node dynamically computes its remote loop free alternate node(s)
  - Done during SFP calculations using PQ algorithm (see draft)
- Automatically establishes a directed LDP session to it
  - The directed LDP session is used to exchange labels for the FEC in question
- On failure, the node uses label stacking to tunnel traffic to the Remote LFA node, which in turn forwards it to the destination
- **Note**: The whole label exchange and tunneling mechanism is dynamic and does not involve any manual provisioning

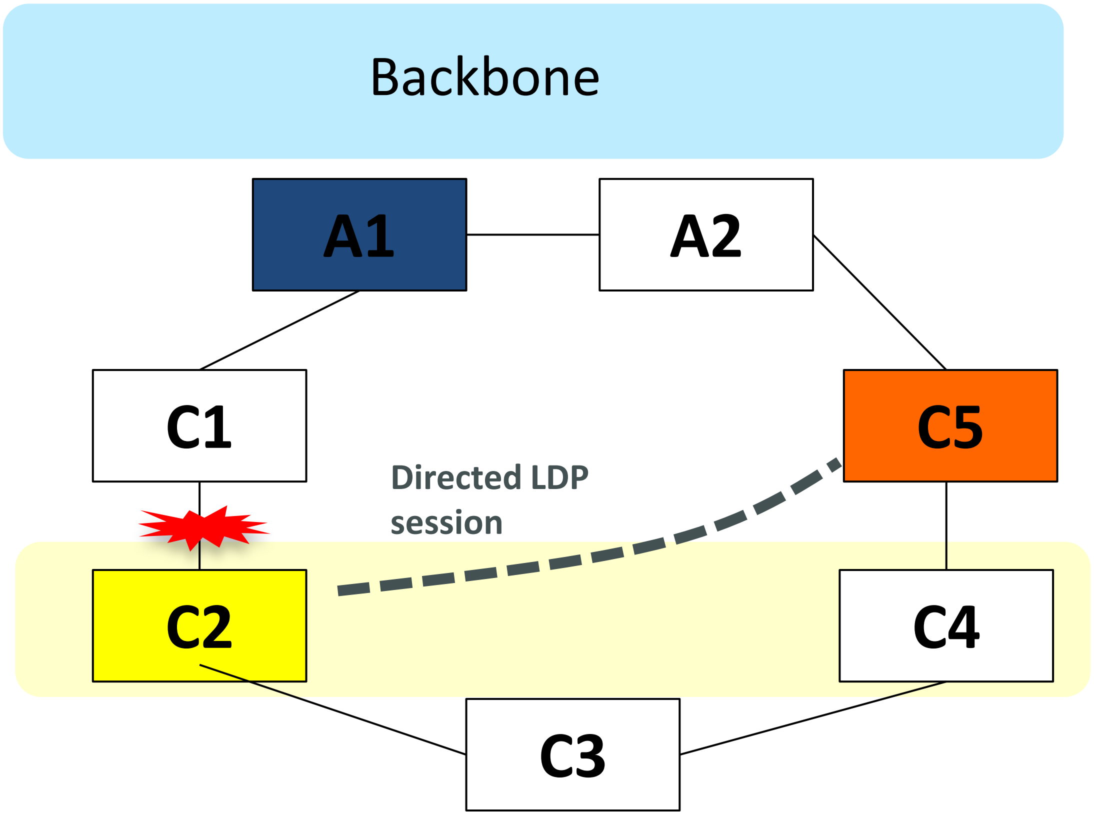

Without rLFA c3 does not offer a LFA to C2

Remote LFA FRR - Comparison

- MPLS-TE FRR 1-hop Link
  - 14 primary TE tunnels to operate
  - 14 backup TE tunnels to operate
  - No node protection
- MPLS-TE FRR Full-Mesh
  - 42 primary TE tunnels to operate
  - 14 backup TE tunnels to operate for Link protection
  - 28 backup TE tunnels to operate for Link & Node protection
- Remote LFA
  - Fully automated IGP/LDP behavior
  - tLDP session dynamically set up to Remote LFA Node

    Even ring involves 1 directed LDP sessions per node

    Odd ring involves 2 directed LDP sessions per node
  - No tunnels to operate

Odd Ring

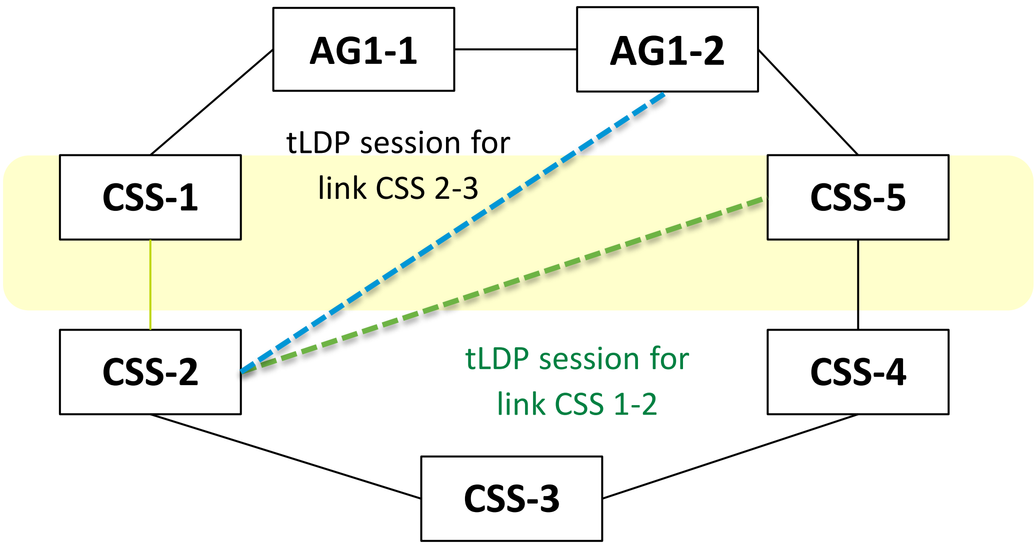

Even Ring

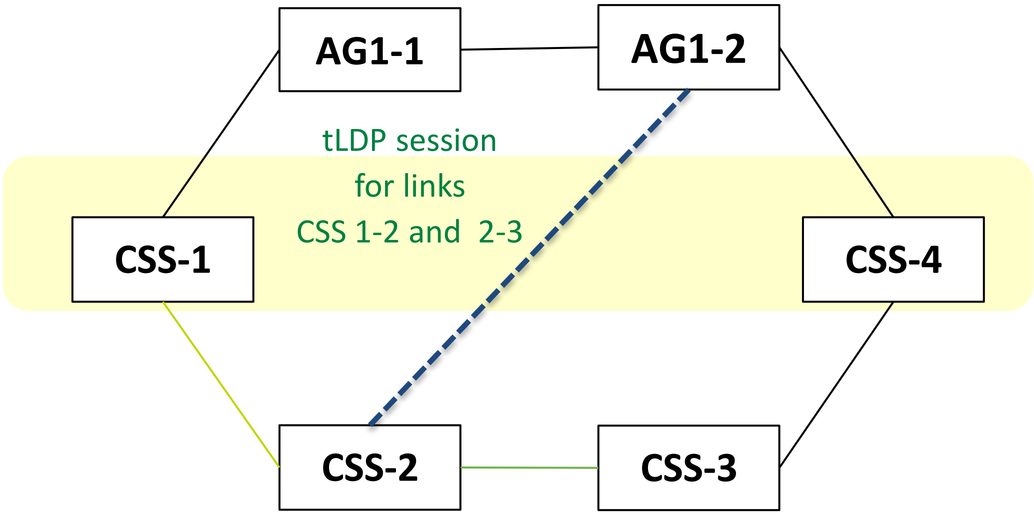

*For the count, account that TE tunnels are unidirectional

Remote LFA FRR Benefits

http://tools.ietf.org/html/draft-shand-remote-lfa

- Simple operation with minimal configuration
- No need to run an additional protocol (like RSVP-TE) in a IGP/LDP network just for FRR capability
  - Automated computation of PQ node and directed LDP session setup
  - Minimal signalling overhead
- Simpler capacity planning than TE-FRR
  - TE-FRR protected traffic hairpins through NH or NNH before being forwarded to the destination
  - Need to account for the doubling of traffic on links due to hairpinning during capacity planning
  - Remote-LFA traffic is forwarded on per-destination shortest-paths from PQ node

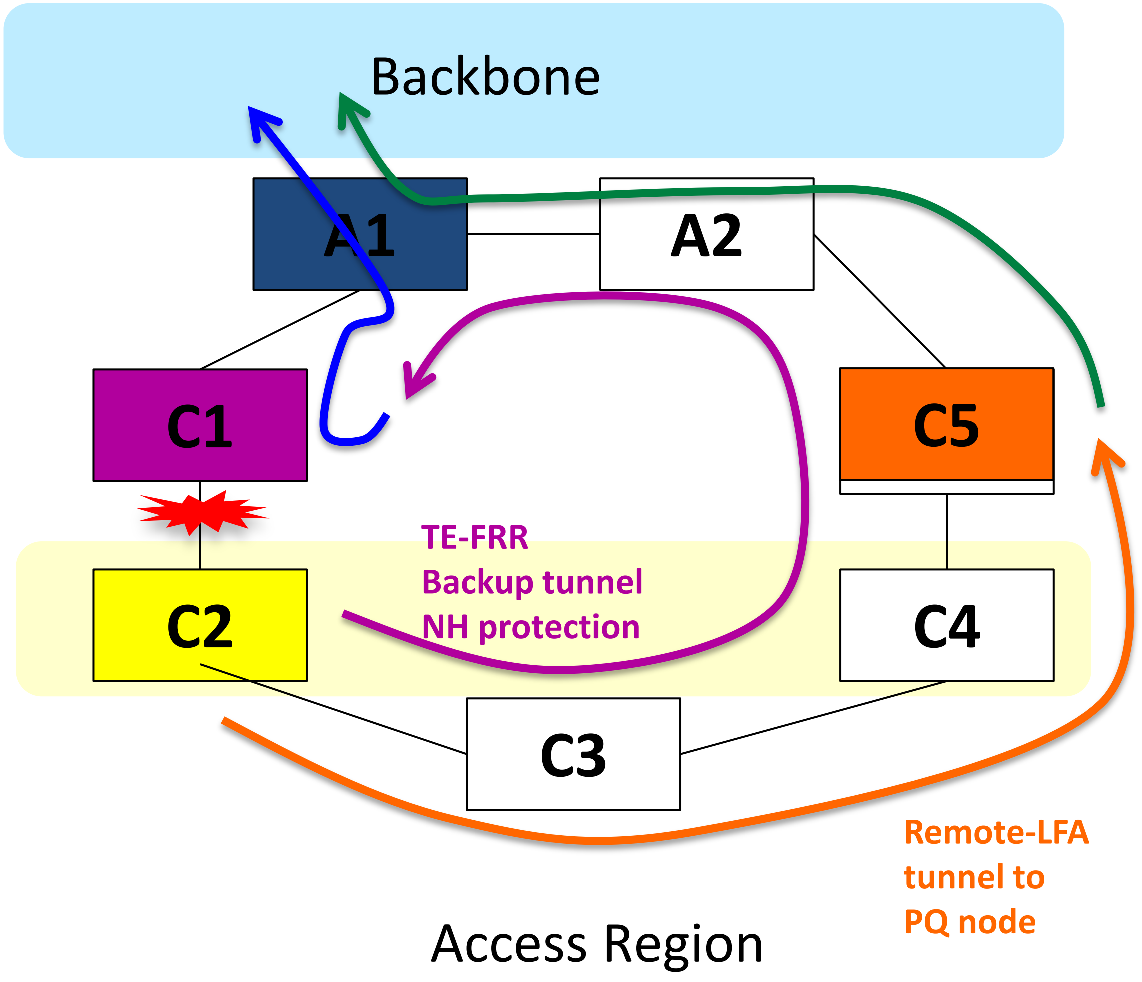

If you need Traffic Engineering then TE is the way to go.
But, if all you need is fast convergence, consider simpler options!

What is PIC or BGP FRR?

- **BGP Fast Reroute** (BGP FRR)—enables BGP to use alternate paths within sub-seconds after a path failure
- PIC or FRR dependent routing protocols (e.g. BGP) install backup paths
- Without backup paths
  - Convergence is driven from the routing protocols updating the RIB and FIB one prefix at a time - Convergence times directly proportional to the number of affected prefixes
- With backup paths
  - Paths in RIB/FIB available for immediate use
  - Predictable and constant convergence time independent of number of prefixes

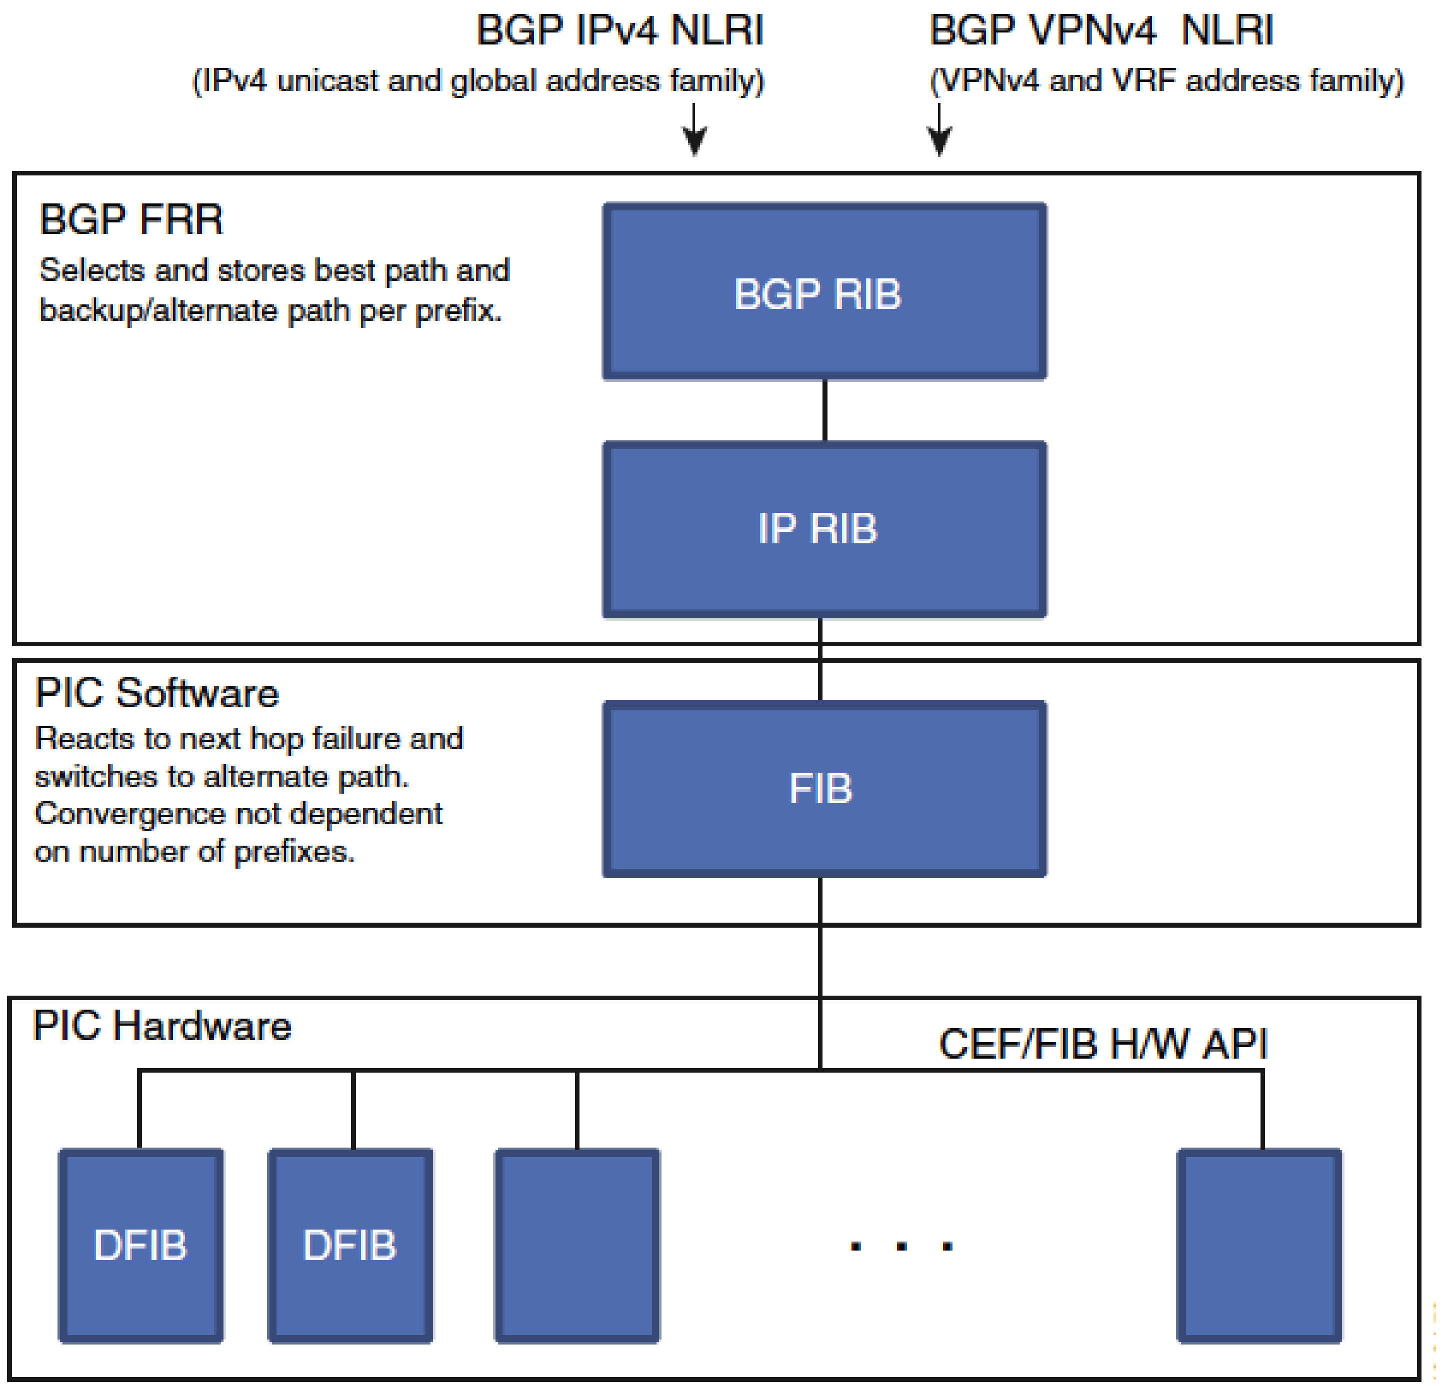

Convergence with and without PIC

BGP PIC Core and PIC Edge

PIC Core

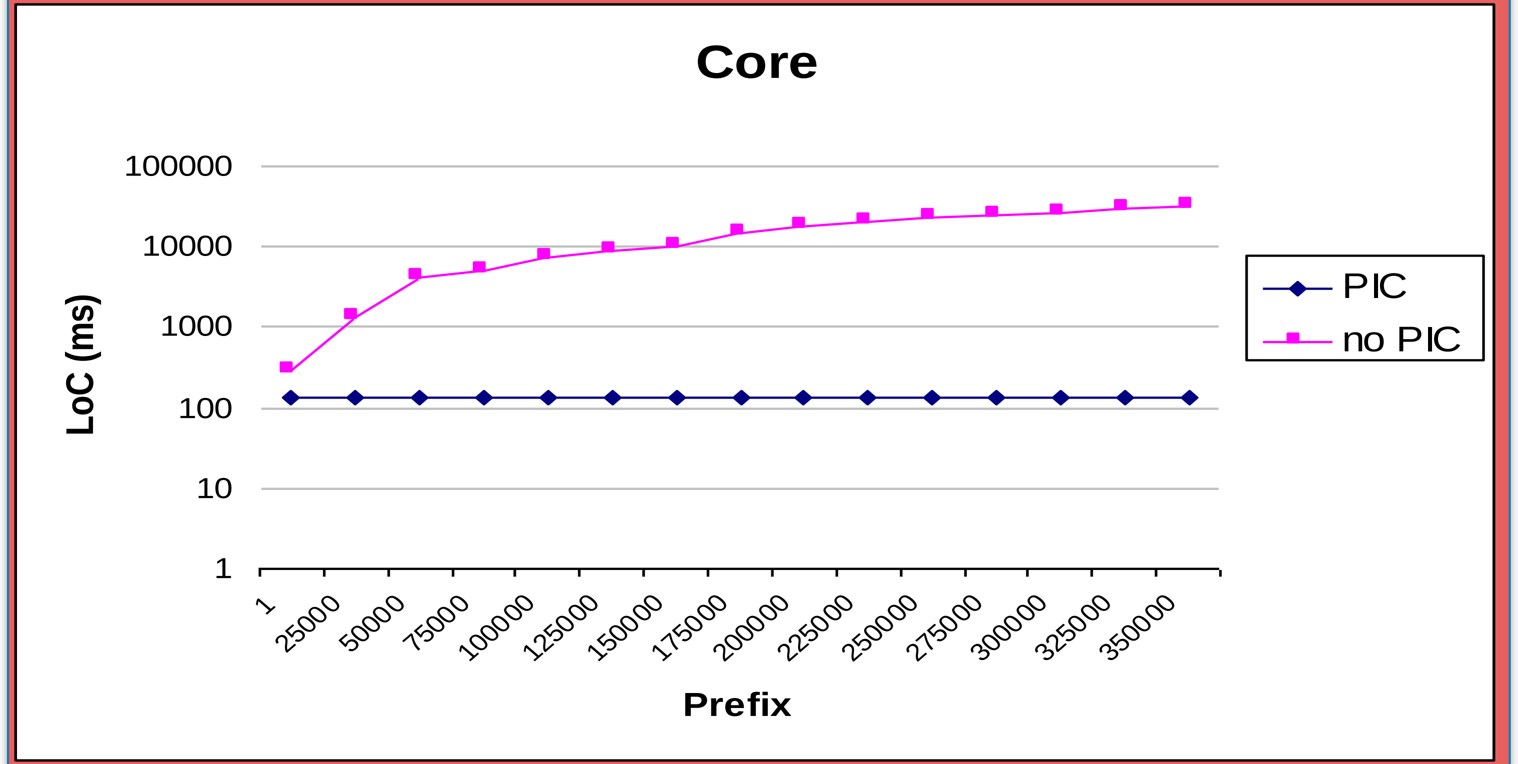

- **Upon failure in the core, without Core PIC, convergence function of number of affected prefixes**
- **With PIC, convergence predictable and remains constant independent of the number of prefixes**

PIC Edge

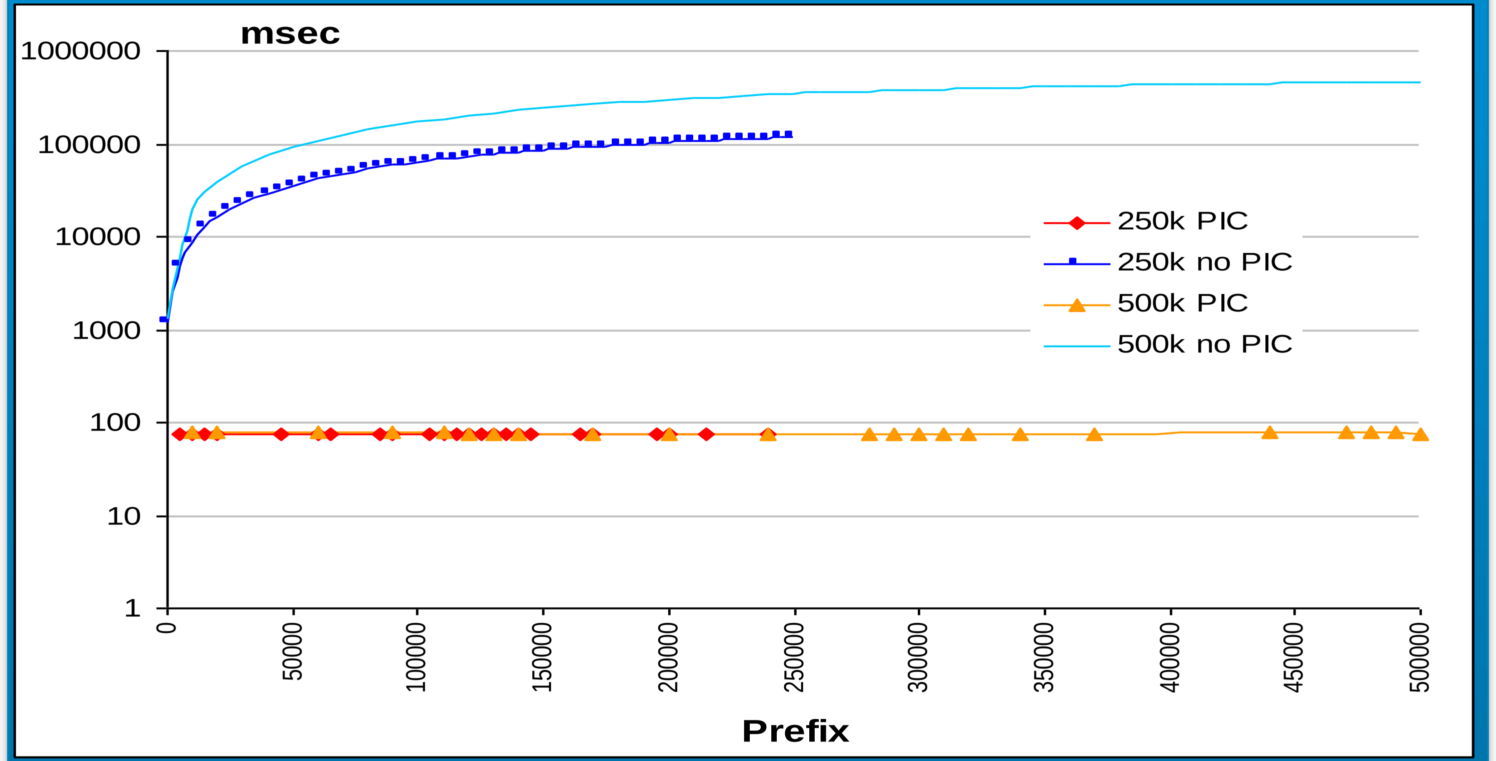

- **Upon failure at the edge, without edge PIC, convergence function of number of affected prefixes**
- **With PIC, convergence predictable and remains constant irrespective of the number of prefixes**

UMMT Functional Aspects
OAM and PM

Mobile Backhaul OAM and PM Drivers

- OAM benchmarks
  – Set by TDM and existing WAN technologies
- Operational efficiency
  – Reduce OPEX, avoid truck-rolls
  – Downtime cost
- Management complexity
  – Large Span Networks
  – Multiple constituent networks belong to disparate organizations/companies
- Performance management
  – Provides monitoring capabilities to ensure SLA compliance
  – Enables proactive troubleshooting of network issues

Transport and Service OAM and PM

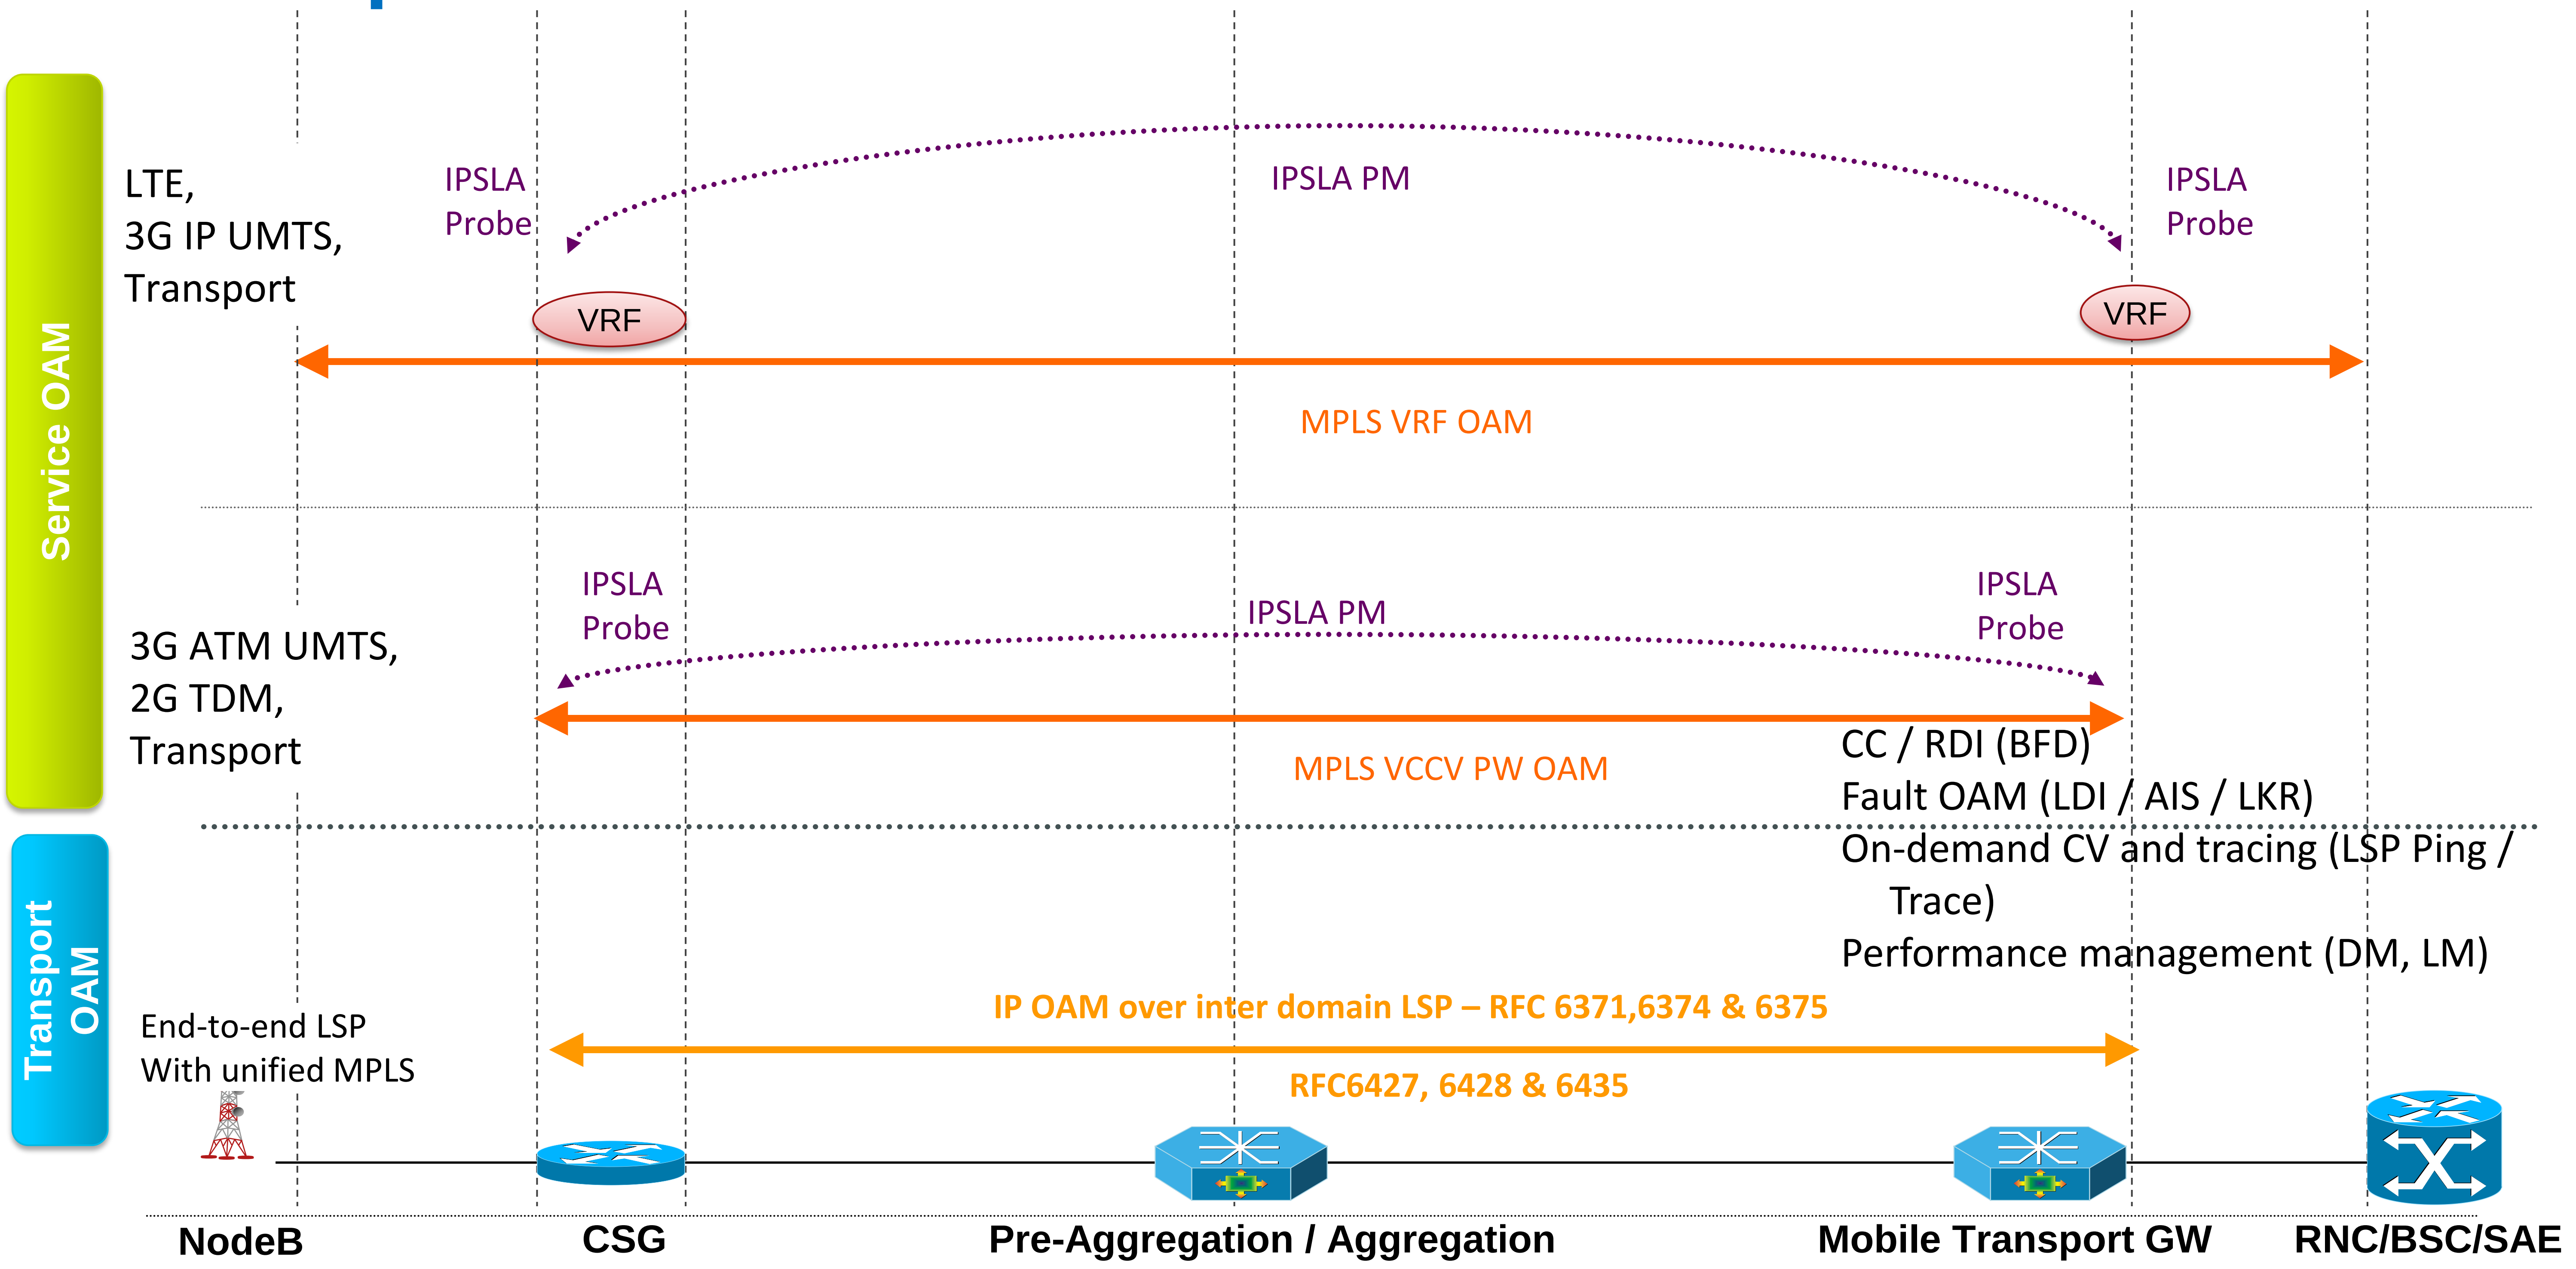

UMMT Functional Aspects
Synchronization Distribution

Mobile Backhaul Synchronization

- UMMT provides a single architecture for concurrent transport of different mobile backhaul technologies
- 2G TDM, 3G ATM/Ethernet, and 4G IP/Ethernet have different Synchronization requirements
  - CDMA, UMTS-TDD, LTE TDD require frequency and phase
  - GSM and UMTS-FDD require frequency only
- To accommodate, UMMT implements a combination of synchronization distribution methods:
  - TDM-based
  - Synchronous Ethernet
  - IEEE 1588v2 PTP

Base Station Synchronization Requirements

Mobile Technology	Frequency	Time / Phase
CDMA2000	±50 ppb	Goal: <3µs Must Meet: <10µs
GSM	50 ppb	N/A
WCDMA	50 ppb	N/A
TD-SCDMA	50 ppb	3µs inter-cell phase difference
LTE (FDD)	50 ppb	N/A
LTE (TDD)	50 ppb	**3µs inter-cell phase difference
LTE MBMS	50 ppb	**5µs inter-cell phase difference
FemtoCell	250 ppb	N/A
WiMAX (TDD)	2 ppm absolute ~50 ppb between base stations	Typically 1 – 1.5 µs
Backhaul Network	16 ppb	N/A

** Under Development

From: MEF-MBH_Synch_HaughHirdRam-Draft_101208_1725_1.pdf

UMMT Synchronization Distribution

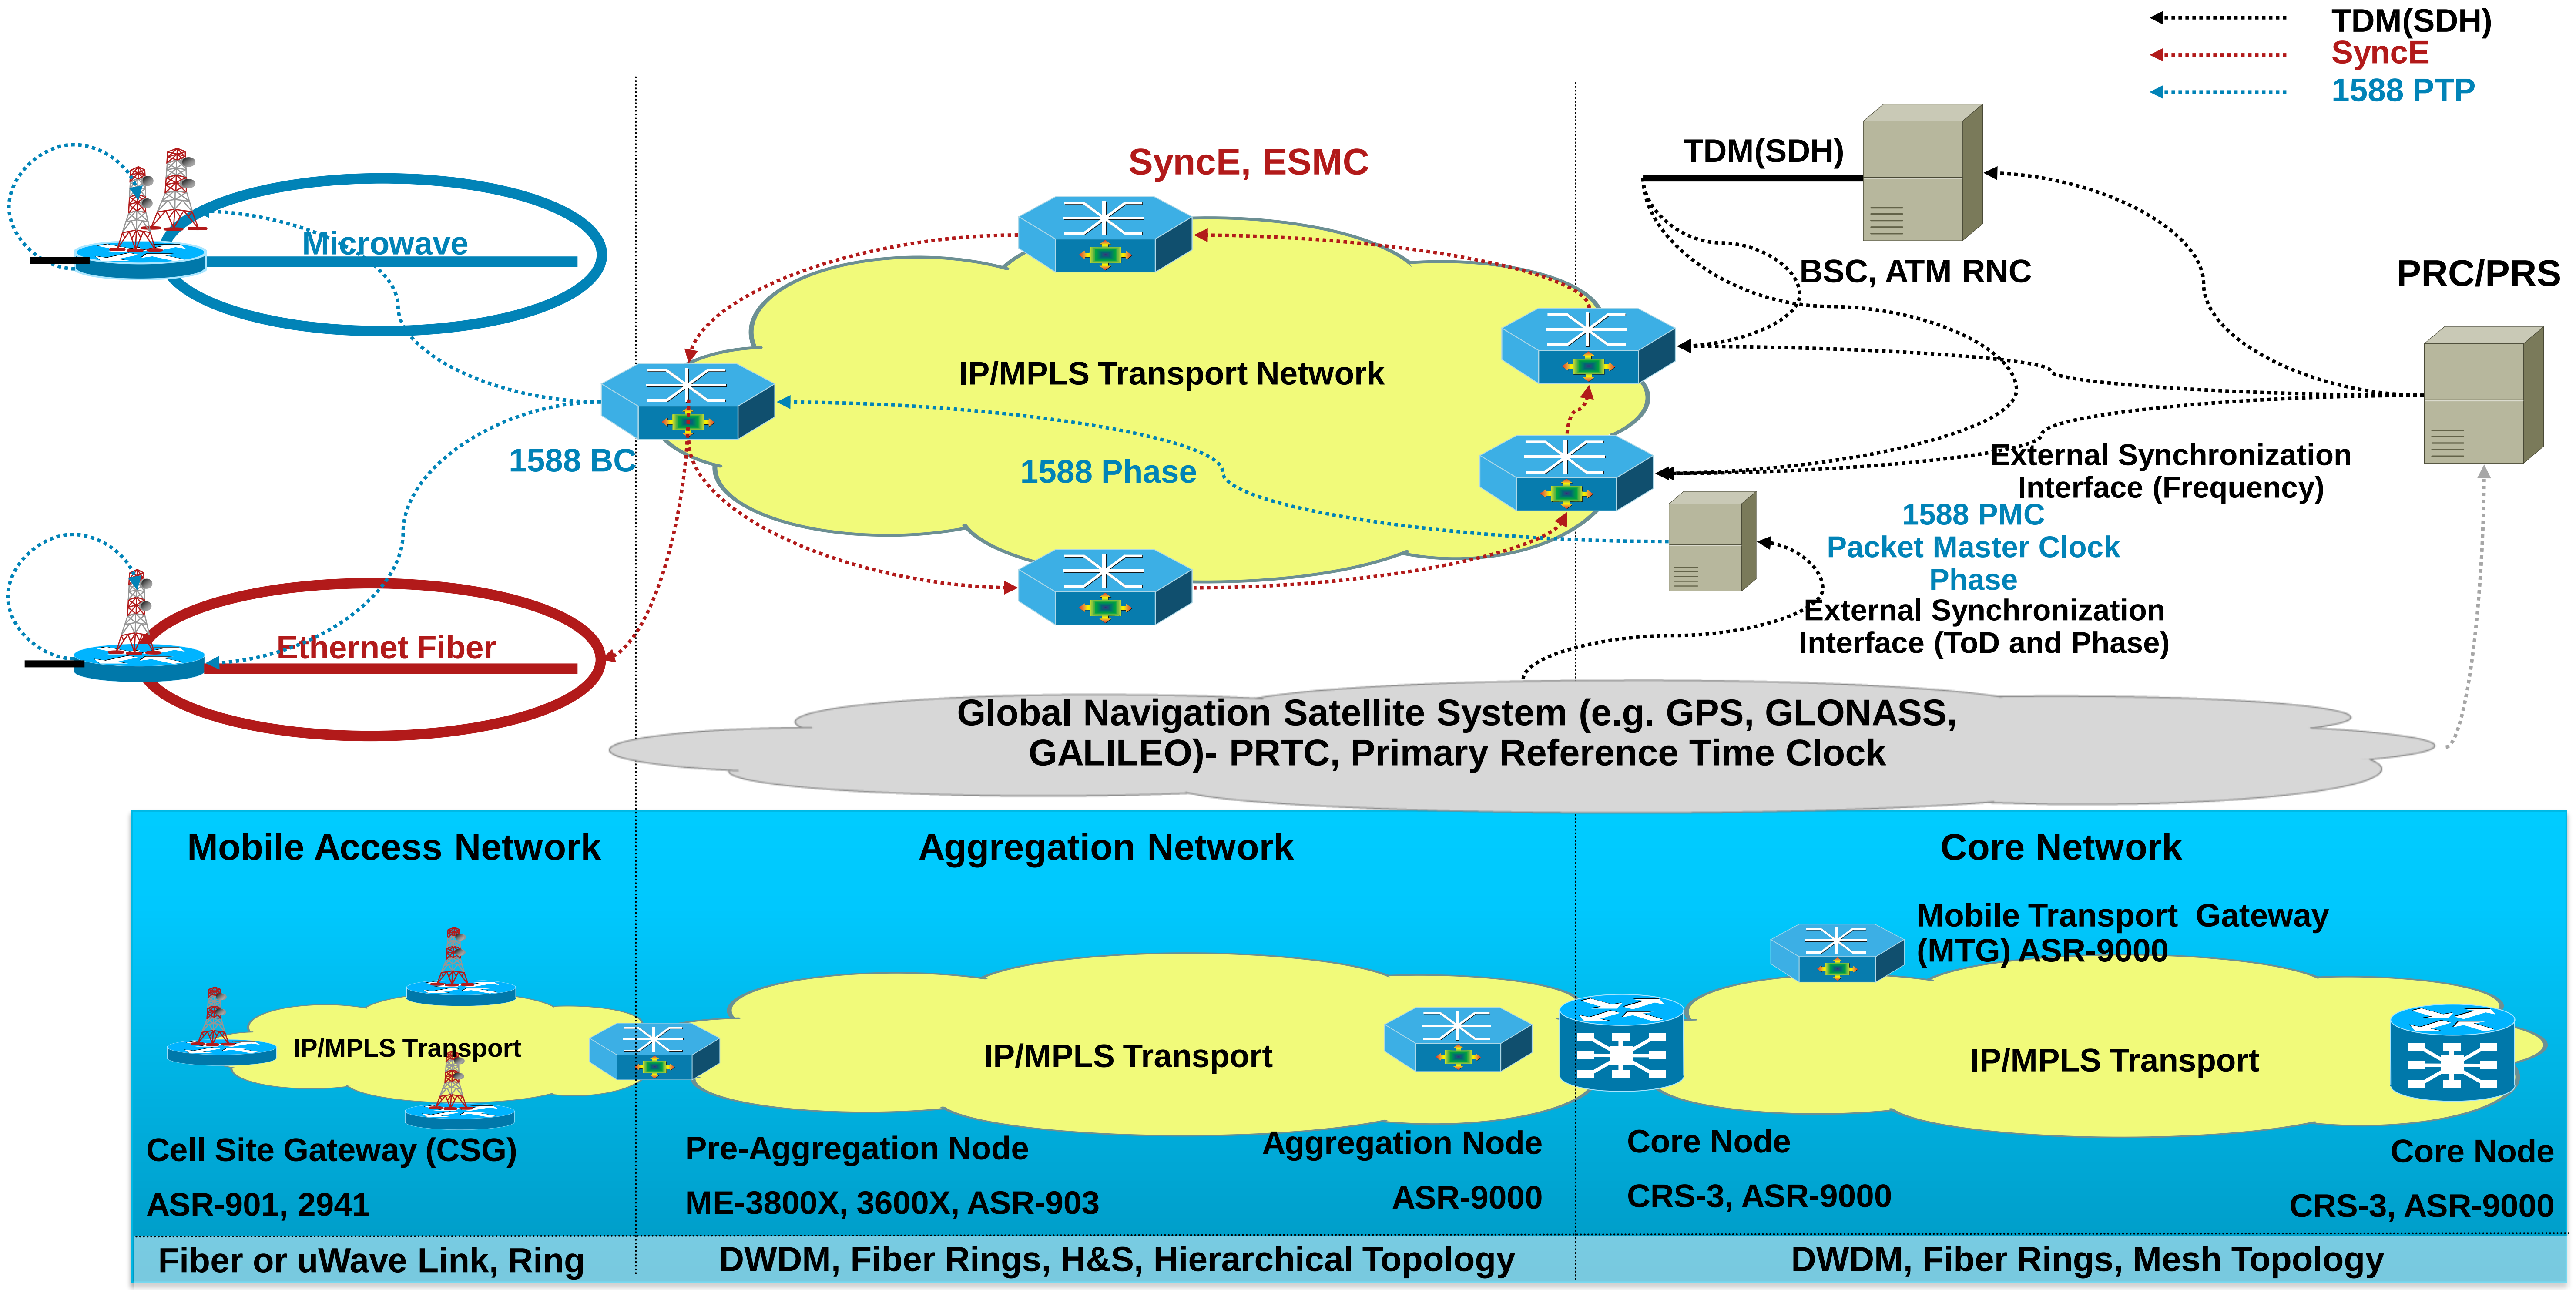

UMMT Functional Aspects
Security

UMMT Security Overview

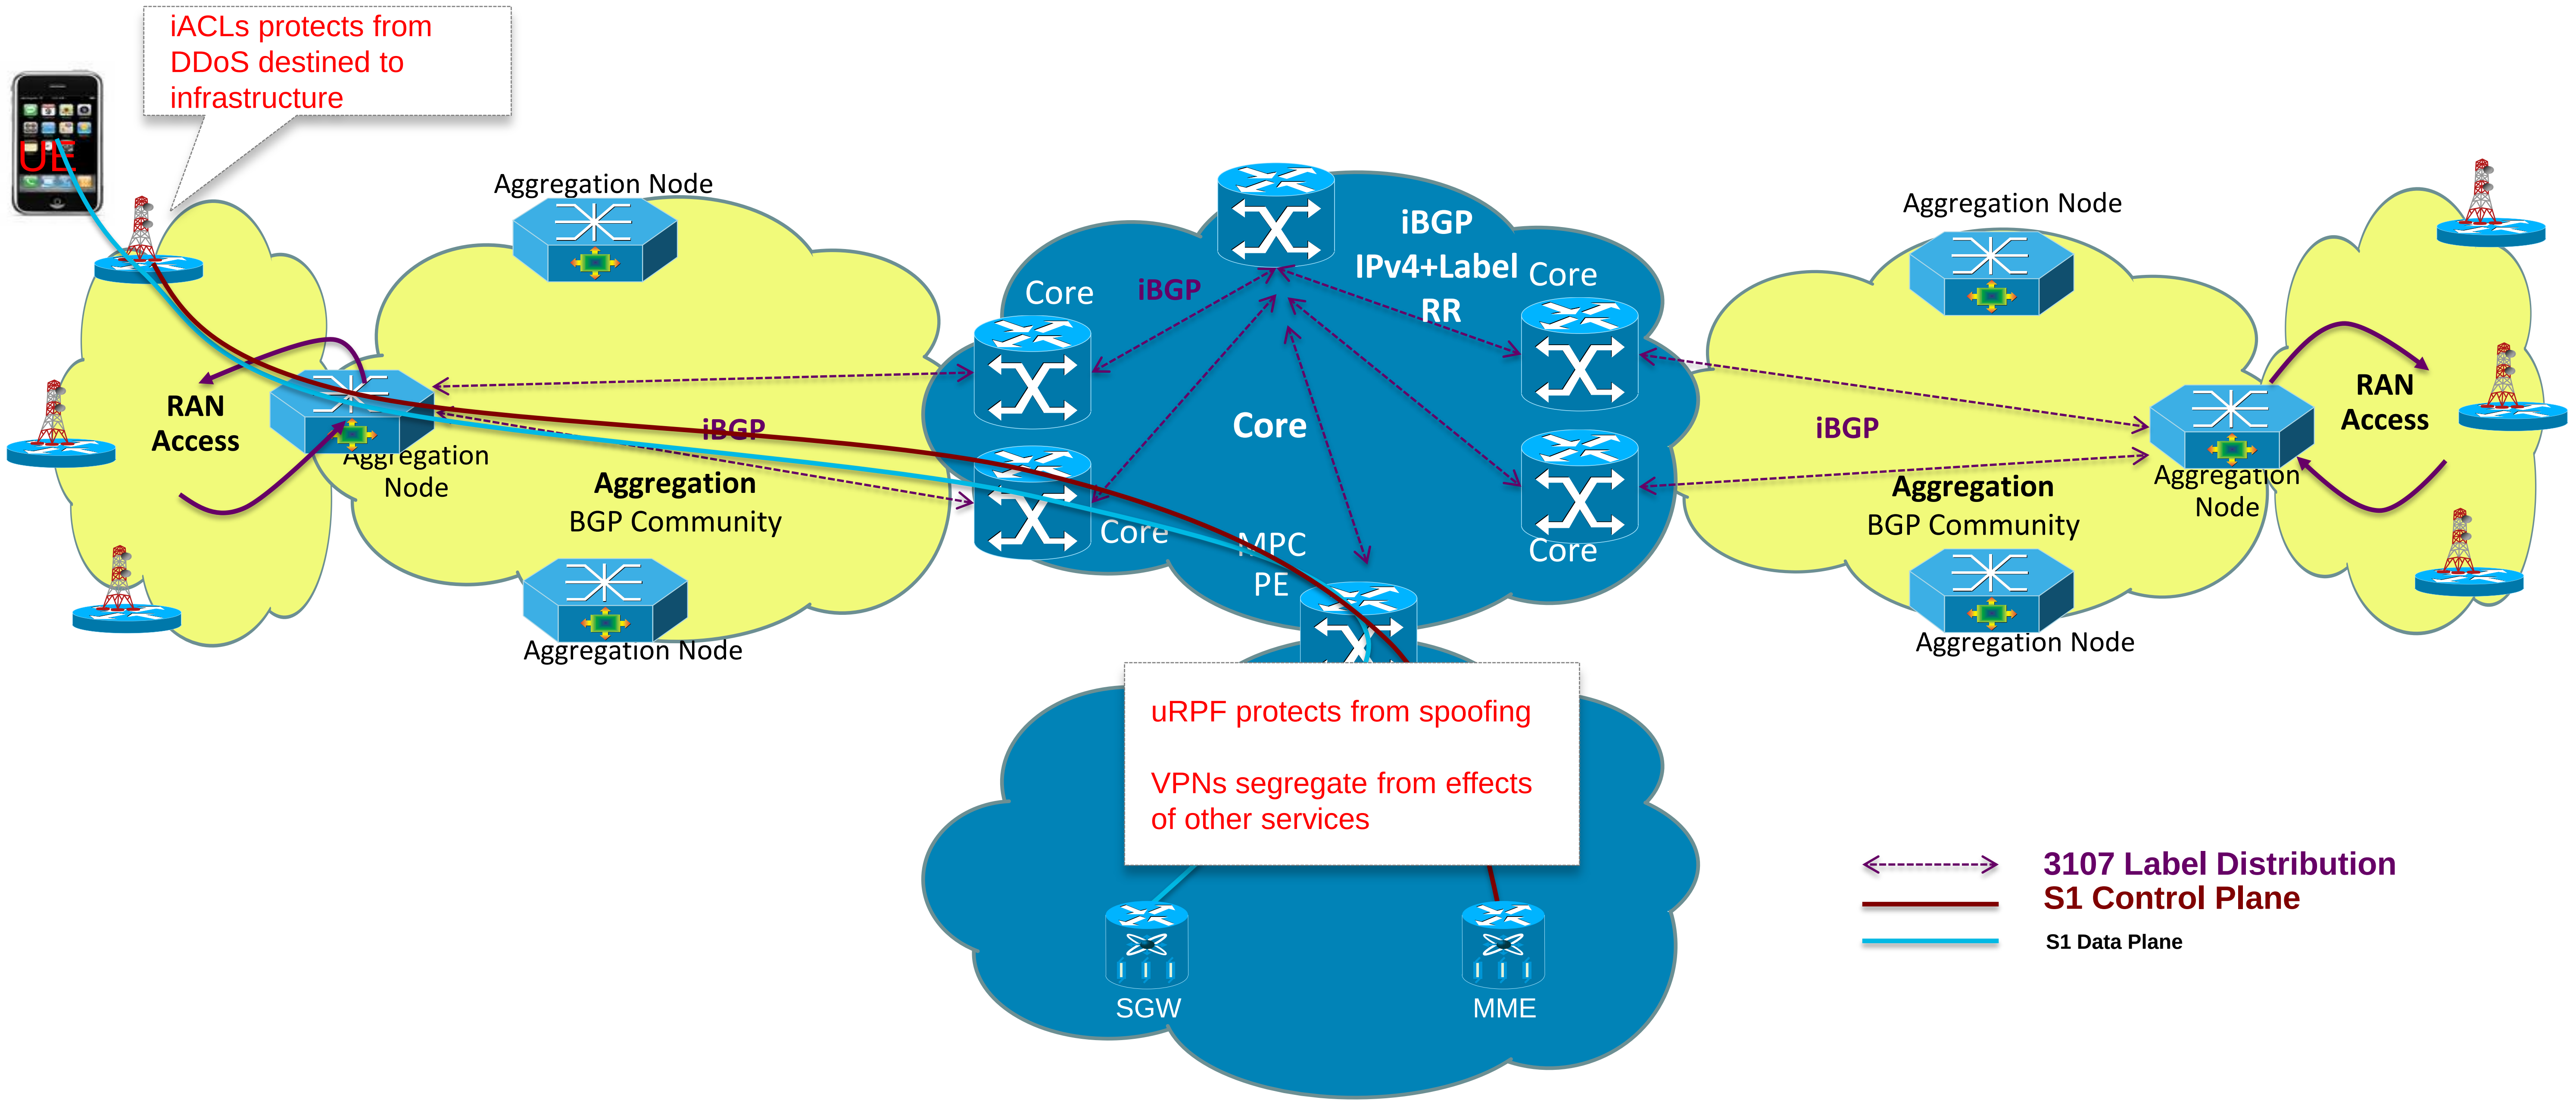

Mobile Backhaul Security

- IPSEC of S1-C can be offloaded to security gateways (S1-U encryption is optional depending on underlying transport).
- uRPF stops source spoofing.
- Firewalls stop radio “DoS” pings from internet.
- MPLS VPNs provide segregation of services across agg/core for FMC.
- Identifiable trails based on IP address – better LI visibility.
- iACLs and ABP (future) protect core from DDoS.
- Bogon filters protect BGP control plane.

SDN Perspectives

Infrastructure Evolution to SDN

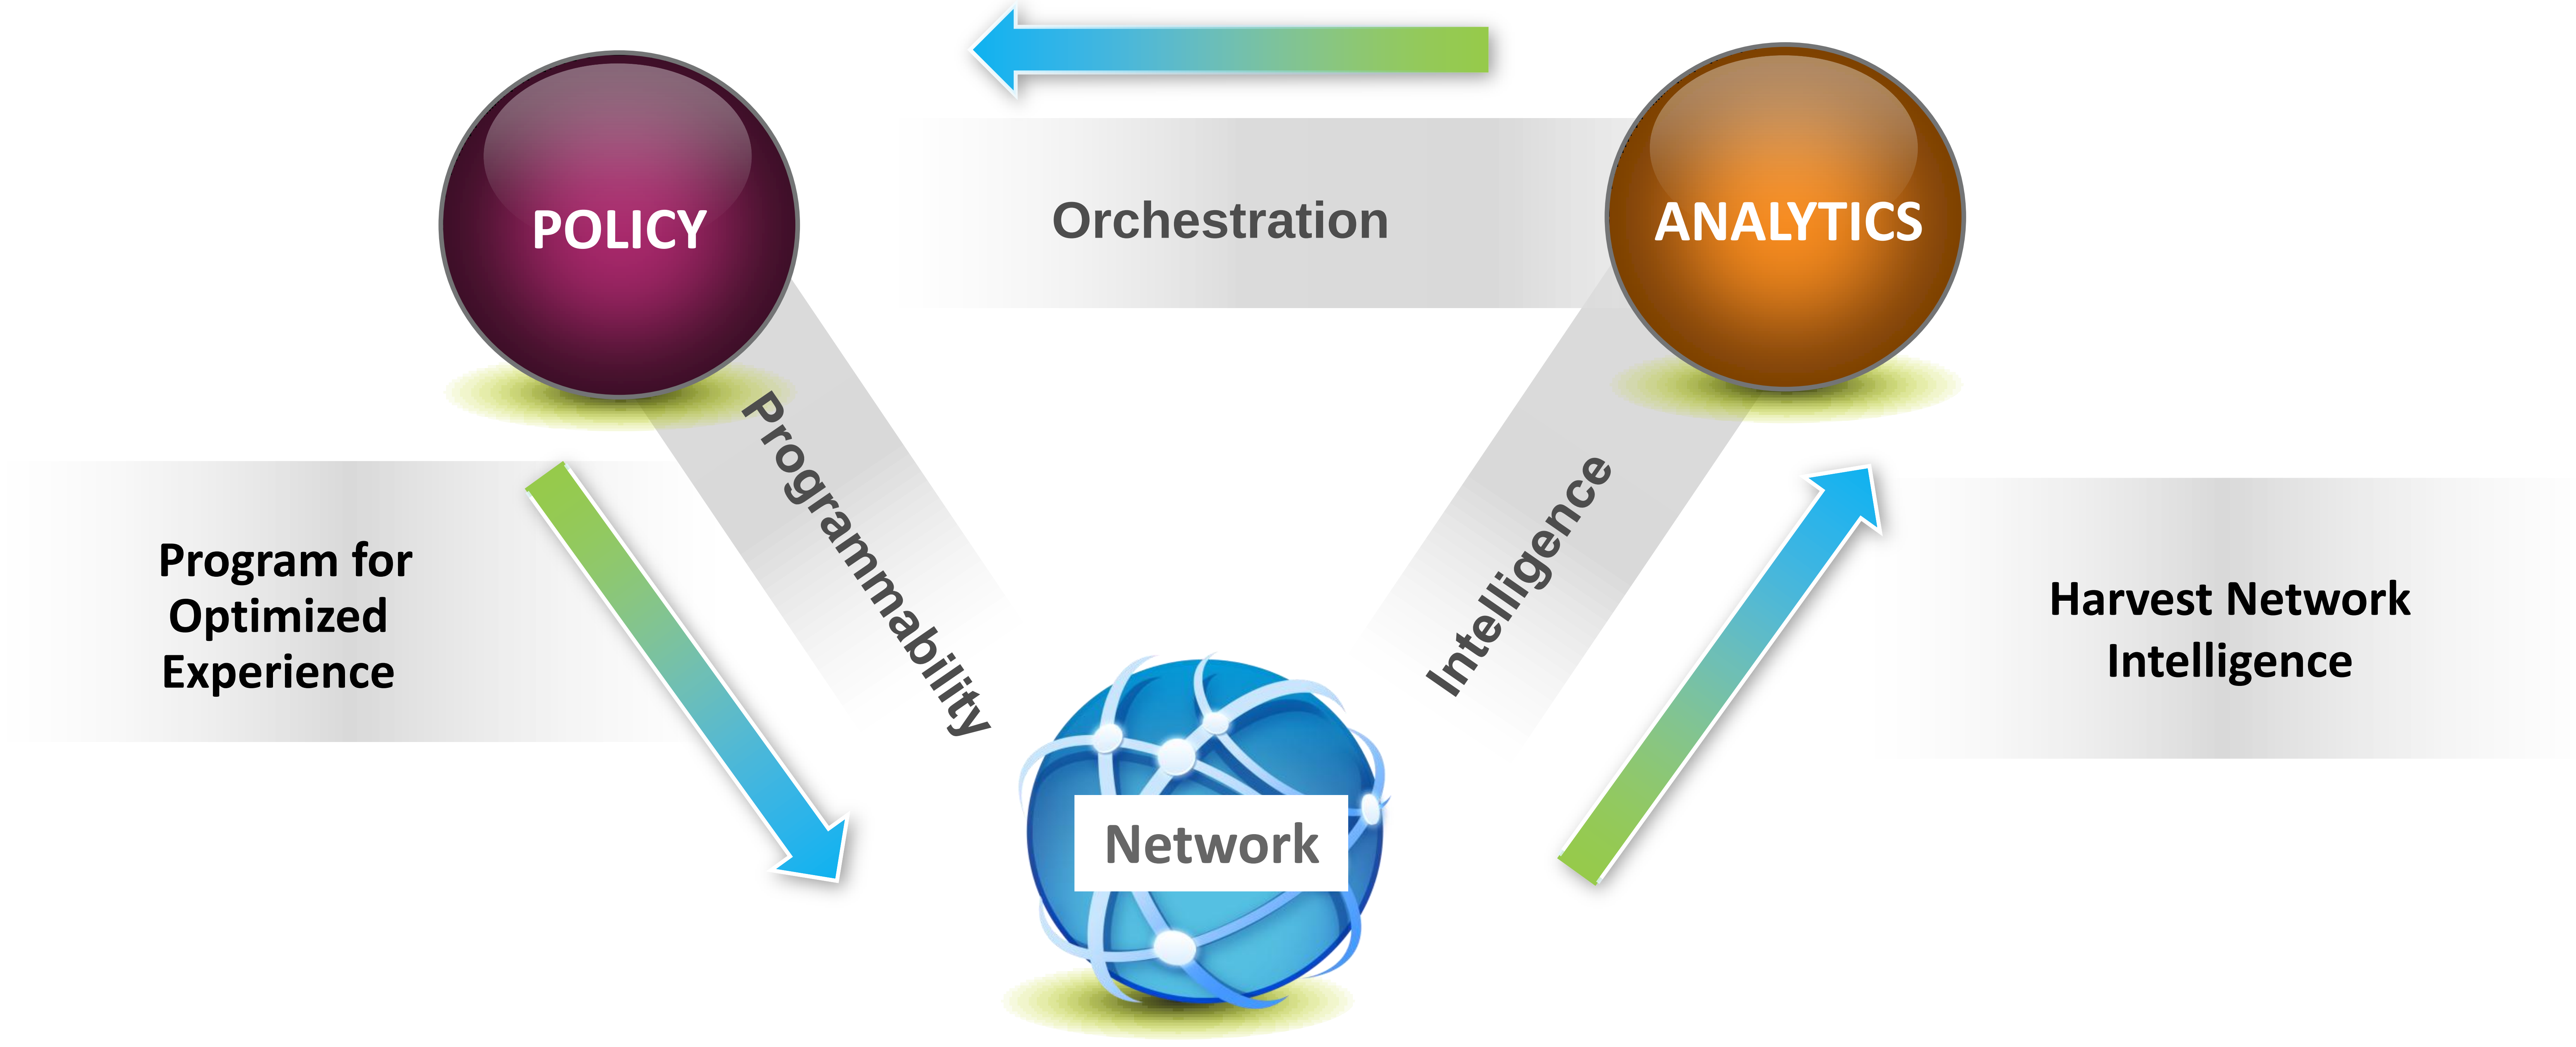

The Collaborative Control Plane

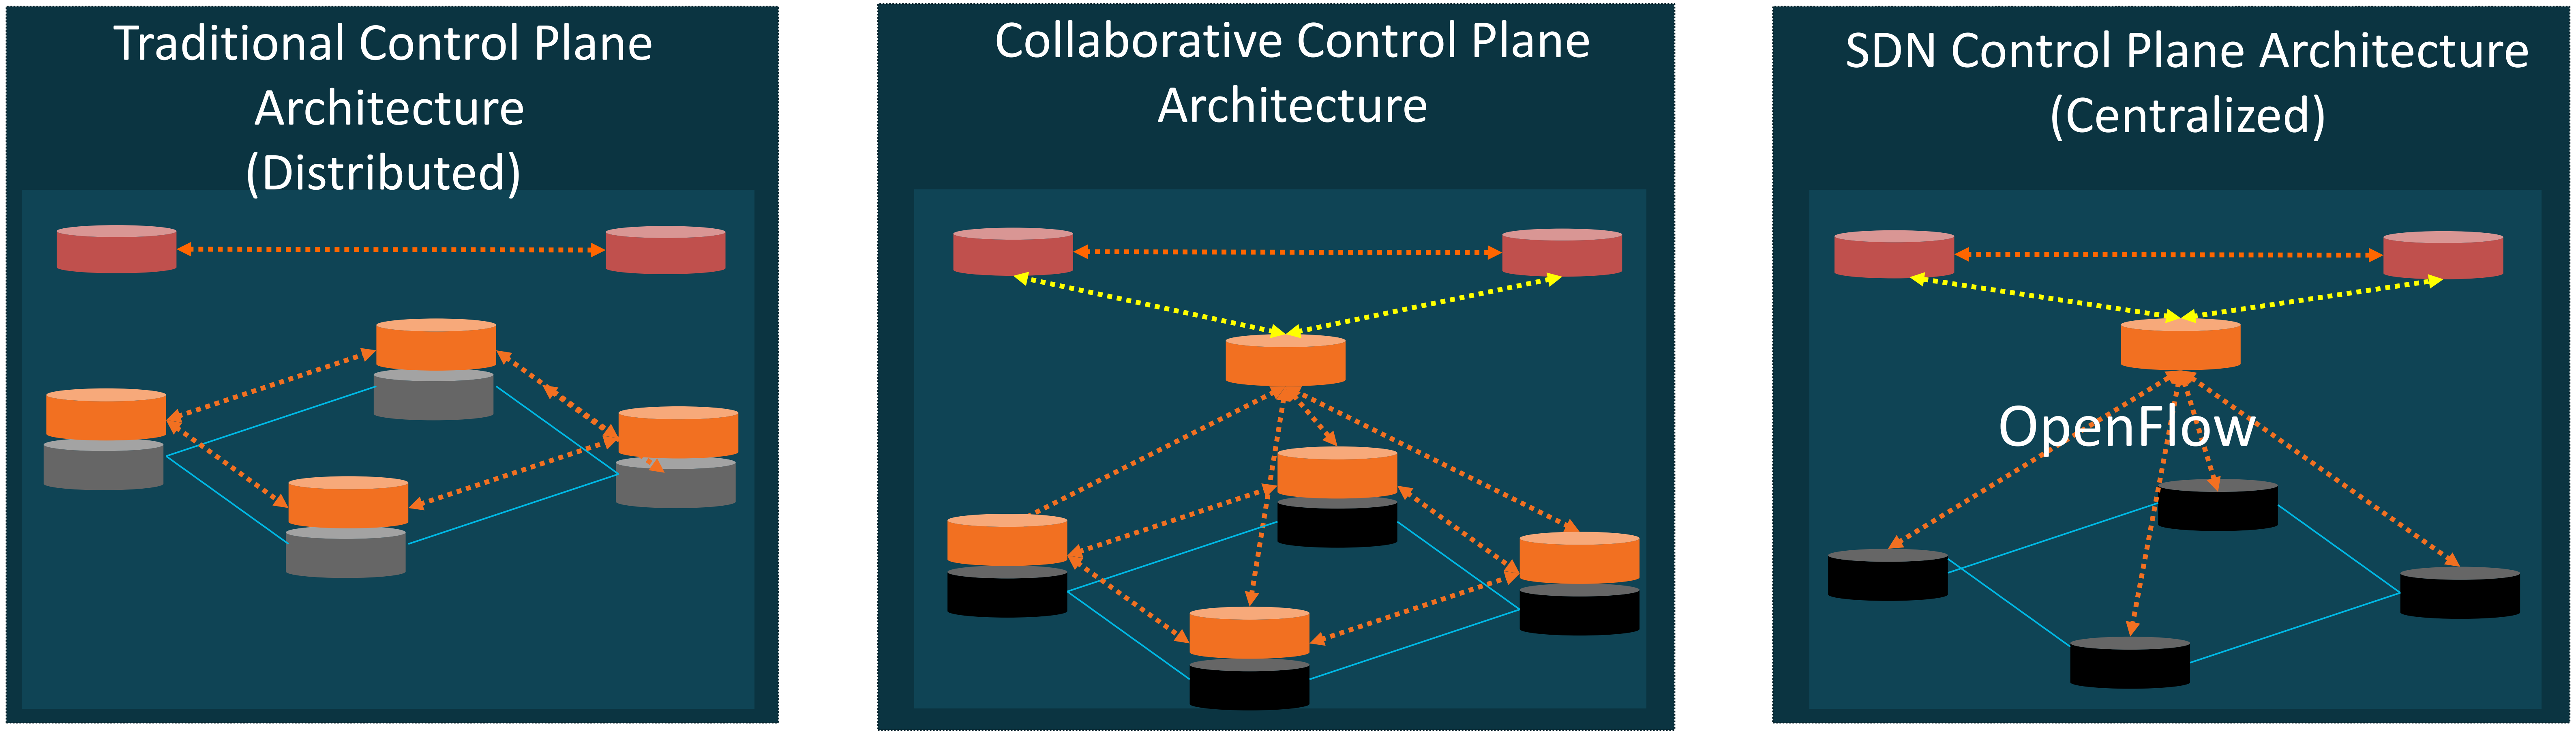

- Distributed Components –Functions tightly coupled to data plane
  - IGP convergence, OAM and physical link state driven protection
- Centralized Components –Functions where a holistic or abstracted view is required
  - Traffic Placement : 30% efficiency

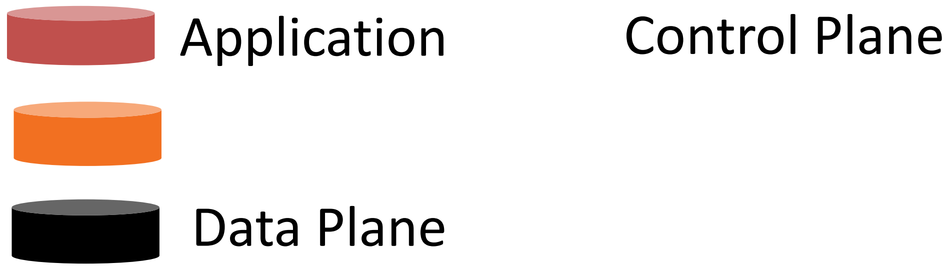

Cisco Perspective on SDN

- Cisco continues to pursue software defined networking
  - SDN includes (1) network overlay virtualization (2) programmatic device APIs (3) network functional abstractions
  - Cisco’s portfolio already includes several key components of an SDN solution
- OpenFlow is a protocol, not an architecture
  - OpenFlow primarily define a protocol for packet forwarding
  - OpenFlow is not complete for production (e.g. lacks: high availability, security, L3-forwarding model, management infrastructure, testing and certification framework, hybrid deployment capability)
- Migration to SDN will be evolutionary
  - Cisco will take a use-case driven approach that draws on several key elements of Cisco's product and technology portfolio
  - Cisco will in the near term engage with specific customers on OpenFlow as a prototype technology
  - There will likely be a need for more than one network controller

UMMT Summary

Key Takeaways

- Unified MPLS simplifies the transport and service architecture
- Seamless MPLS LSPs across network layers to any location in the network
- Flexible placement of L2/L3 transport virtualization functions to concurrently support 2G/3G/4G services
- Service provisioning only required at the edge of the network
- Divide-and-conquer strategy of small IGP domains and labeled BGP LSPs helps scale the network to hundred of thousands of LTE cell sites
- Simplified carrier-class operations with end-to-end OAM, performance monitoring, and LFA FRR fast convergence protection
- Easily extensible to support wholesale RAN sharing and wireline services

For More Information

- Email: nts-info@cisco.com
- Cisco SP Mobility Community: https://communities.cisco.com/community/solutions/sp/mobility
  - – The UMMT Design Guides mentioned in this presentation can be found at the SP Mobility Community

Complete Your Paper “Session Evaluation”

Give us your feedback and you could win 1 of 2 fabulous prizes in a random draw.

Complete and return your paper evaluation form to the room attendant as you leave this session.

Winners will be announced today.
You must be present to win!

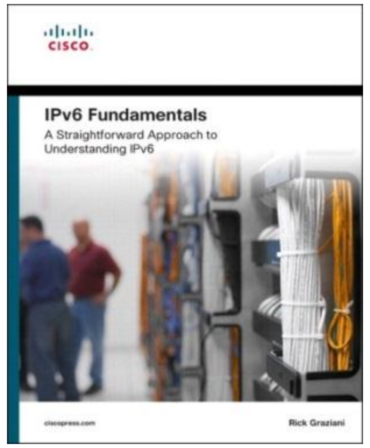

plantronics.

PEARSON ..visit them at BOOTH# 100

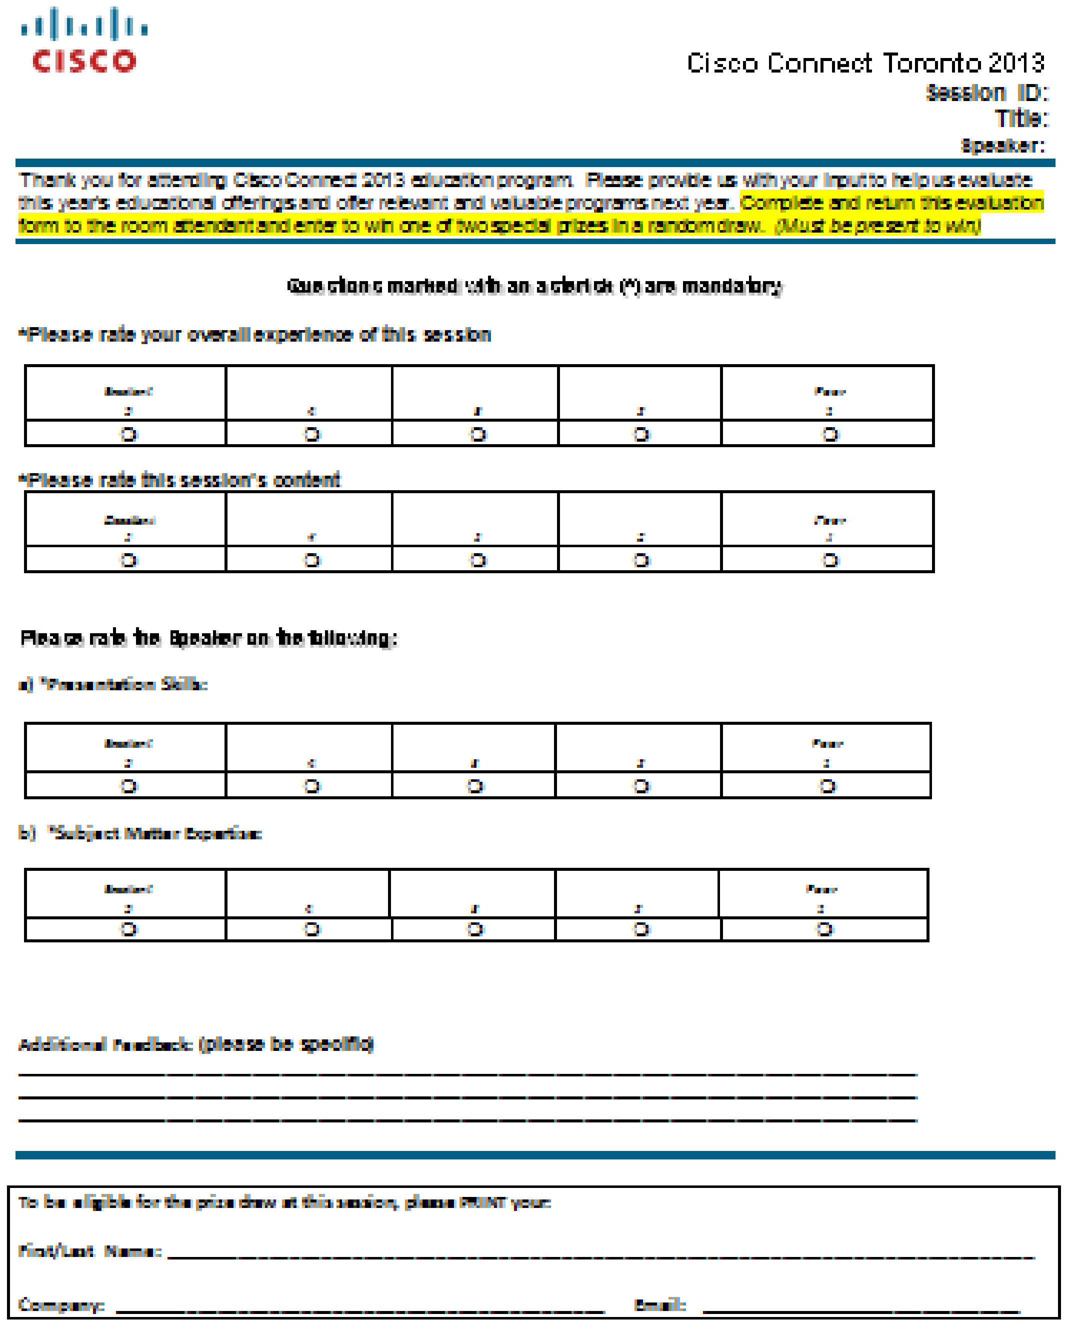

Cisco Connect Toronto 2013
Session ID:
Title:
Speaker:

Thank you for attending Cisco Connect 2013 education program. Please provide us with your input to help us evaluate this year's educational offerings and offer relevant and valuable programs next year. Complete and return this evaluation form to the room attendant and enter to win one of two special prizes in a random draw. (Must be present to win)

Questions marked with an asterisk (*) are mandatory

*Please rate your overall experience of this session

*Please rate this session's content

Please rate the Speaker on the following:

a) *Presentation Skills:

b) *Subject Matter Expertise:

Additional Feedback: (please be specific)

To be eligible for the prize draw at this session, please PRINT your:

First/Last Name:

Company: Email:

BUILT FOR
THE HUMAN
NETWORK

CISCO