The bridge to possible

Aktuální WAN portfolio

a progresivní transportní technologie

Peter Morvay
Systems Engineer– #55452
6. 6. 10111

A9K

Cisco ASR 9000 Series

Hardware portfolio

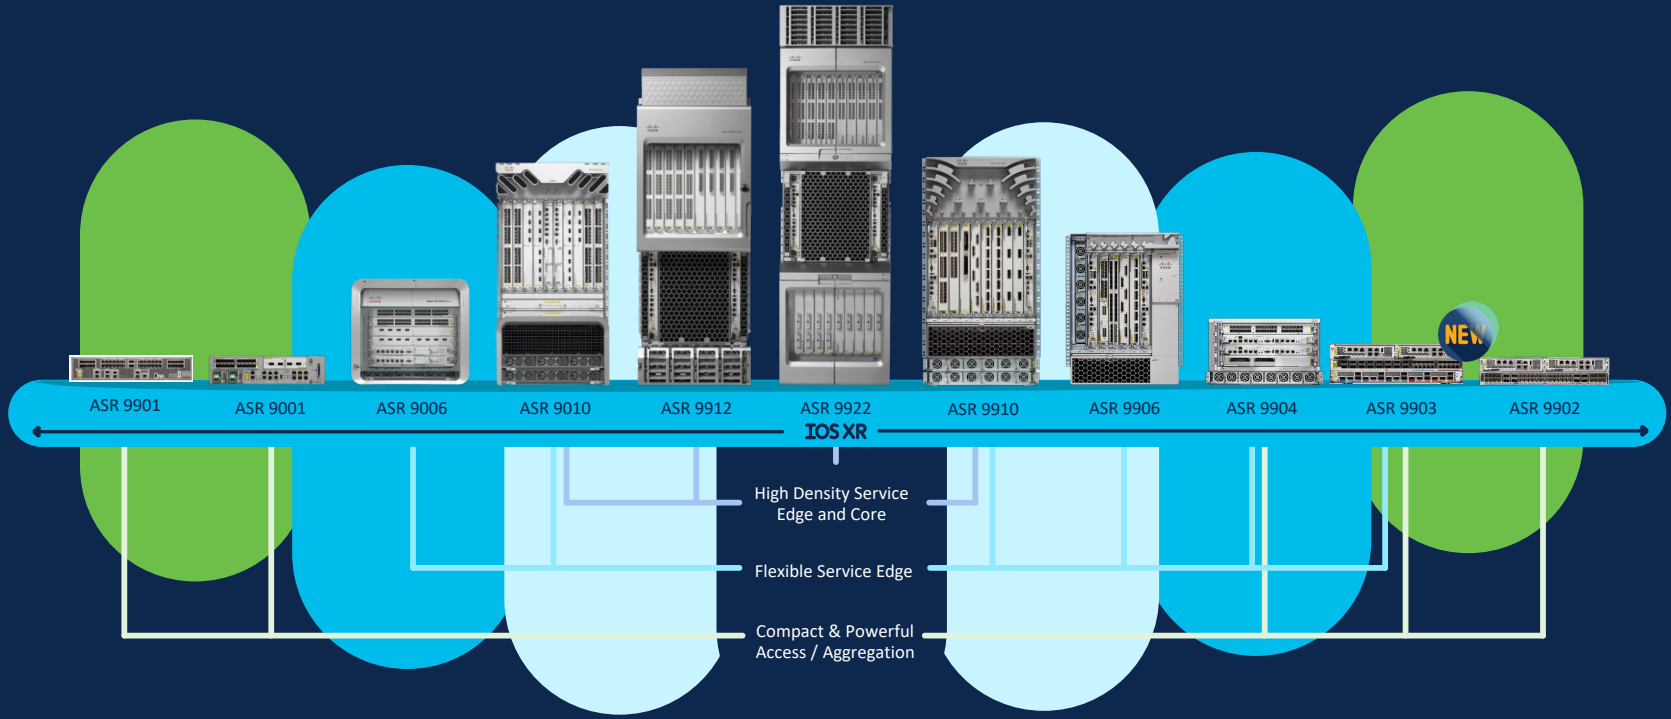

One Set of Commons for 3rd, 4th, 5th Generation

A99-RP3-SE/TR & A9K-RSP5-SE/TR

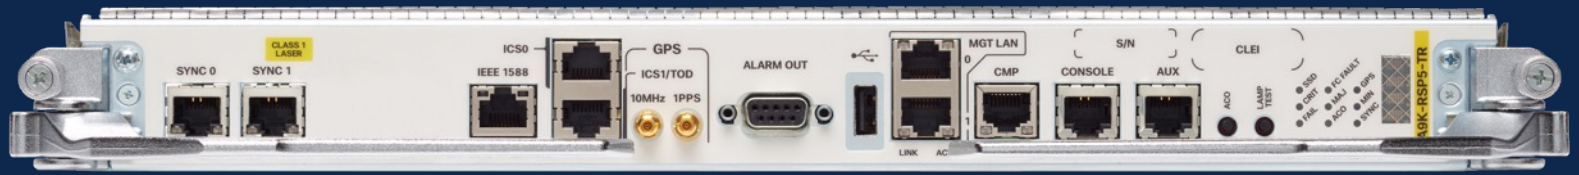

4th Generation Route Processor Card

- 9006/9010 - 1.2T/slot throughput (redundant configuration)
- 8 Core Intel CPU at 2G processor
- Available in both TR (24GB)/ SE (40GB) variants

Line Cards Supported

- 5th Generation
- 4th Generation
- 3rd Generation

RSP5 Supported Chassis

- 9006/9010/9910/9906/9904

RP3 Supported Chassis

- 9912 / 9922

Class C Complaint RSP5-X “+ Class C”

Use Case: 5G Packet Core & Converged wireline, wireless edge services

A9K-RSP5-X-SE/TR

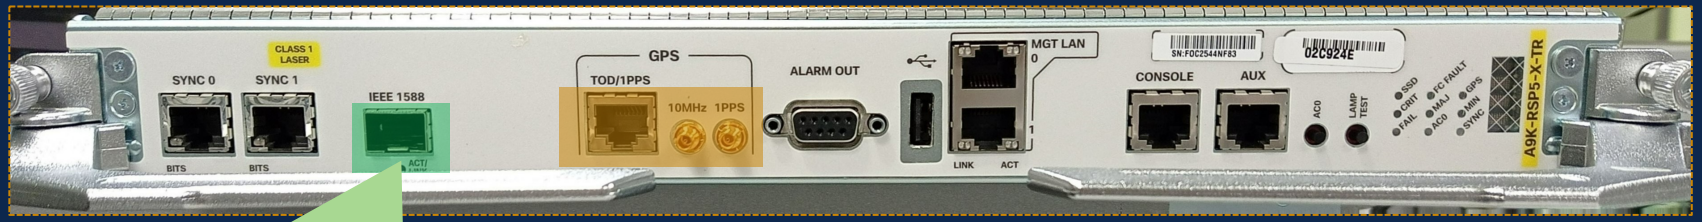

1588: Changed from RJ45 (RSP5) to 10/1G optical port (RSP5-X)

A9K-RSP5-SE/TR

ICS0/1 and CMP ports are removed from the new RSP5-X

ASR 9000 5th Generation Portfolio

400GE to 4T

	Ports	Bandwidth	Combo Ports	MACSec/ FlexE	Timing	RSP / RP
 A99-32X100GE-X	32 Ports of QSFP28	3.2 Tbps	No	No	Class B SyncE	RSP5, RP3
 A9K-20HG-FLEX	15 Ports QSFP28 5 Ports QSFP-DD	2 Tbps	Yes	MACSec + Flex-E	Class C SyncE	RSP5, RP3
 A9K-8HG-FLEX	6 Ports QSFP28 2 Ports QSFP-DD	800 Gbps	Yes	MACSec + Flex-E	Class C SyncE	RSP5, RSP880-LT, RSP880, RP3, RP2
 A99-10X400GE-X	10 Ports of QSFP-DD	4 Tbps	Yes	MACSec	Class B SyncE	RSP5, RP3
 A9K-4HG-FLEX	4 Ports QSFP28 16 Ports SFP28 24 Ports SFP+	400 Gbps	Yes	MACSec	Class C SyncE	RSP5, RSP880-LT, RSP880, RP3, RP2

A99-10X400GE-X-SE/TR (7-fabric) LC Architecture (when used in 9904)

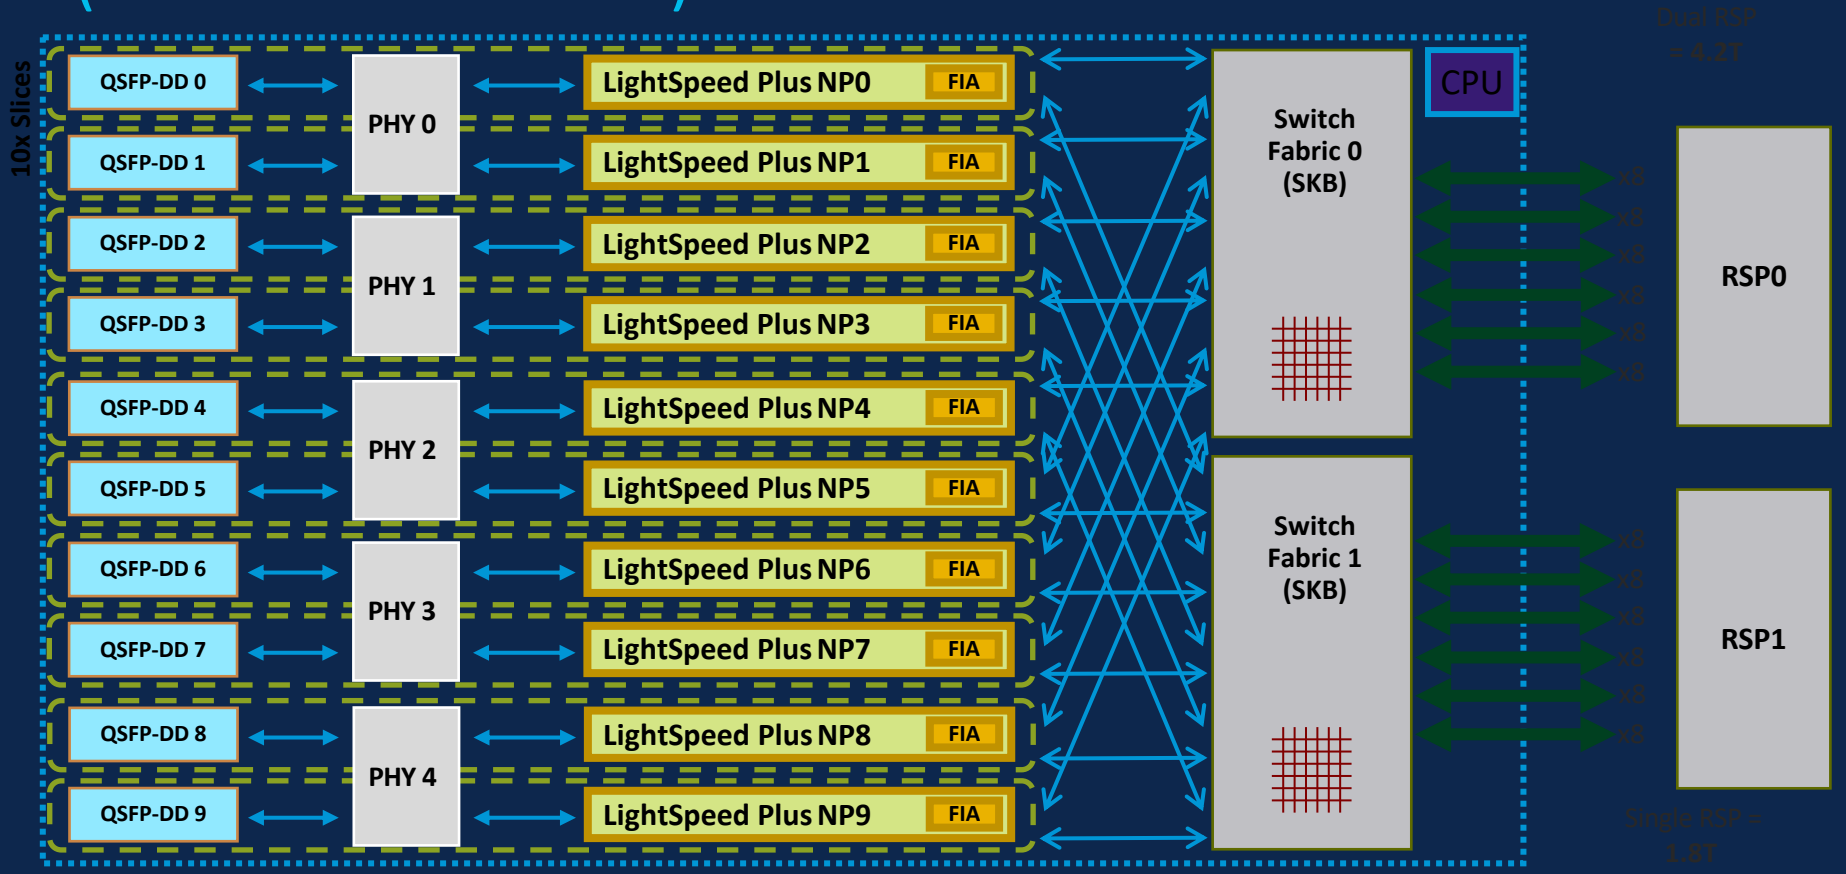

ASR 9000 5th Generation Compact Chassis

	Throughput	Ports	Multi-rate Ports	MACSec/OTN	Timing
 ASR-9902	800 Gbps	2 Ports QSFP-DD 6 Ports QSFP28 16 Ports SFP28 24 Ports SFP+	Yes	MACSec/ OTN	Class C
 ASR-9903 (Fixed Ports)	1.6 Tbps	16 Ports QSFP28 20 Ports SFP+	Yes	MACSec	Class C
 A9903-20HG-PEC	2 Tbps	15 Ports QSFP28 5 Ports QSFP-DD	Yes	MACSec	Class C
 A9903-8HG-PEC	800 Gbps	32 Ports SFP28 16 Ports SFP+	Yes	MACSec	Class C

Cisco ASR 9000 Compact Routers

	ASR 9901	ASR 9903	ASR 9902
			
RU Size	2 RU	3 RU	2 RU
Depth	600 mm	600 mm	483 mm
Air Flow	Front to Back	Front to Back	Front to Back
Capacity	Up to 456 Gbps	Up to 3.6 Tbps	Up to 800GE
Route Processor	Integrated RP	Redundant RP	Redundant RP
Ports/Slots	Fixed Ports: Total 42 • 2x100GE • 24 x 1/10GE (Linear tunable optics) • 16x1GE	Fixed Board: 14x100GE QSFP28 + 2x100G | 20x10GE SFP+ 1 Port Expansion Card: • 2T Port exp. card • 800G Port exp. Card	Fixed Ports: Total 48 • 8x 100GE • 16x 25G / 10GE • 24x 10GE w/ OTN mode ((Linear tunable optics)
MACSec	Yes	Yes	Yes
Applications	Access / Aggregation / Service Edge	Service Edge	Service Edge / Aggregation
OS	IOS XR (64 Bit)	IOS XR (64 Bit)	IOS XR (64 Bit)

10G to 1GE Rate Change SFP Optics (aka Smart 1G)

2 modules:

- **SFP-1G-SX**: 1km reach over OM3/OM4
- **SFP-1G-LH**: 10km reach over single-mode fiber

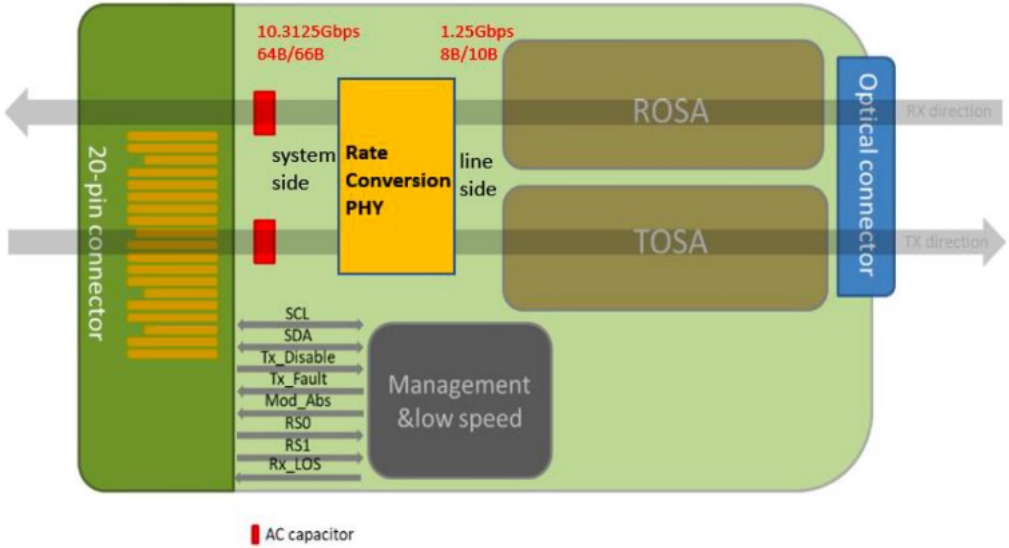

1G modules can work in 10G/25G ports

- designed to allow platforms without 1G Media Access Control to link up to 1G transceivers; works on 10G/25G ports
- provides a standard 10G electrical interface on the system side
- Contains PHY to convert 10G electrical interface to 1G optical interface
- In the ingress the module converts physical coding from 8B/10B to 64B/66B encoding
- In the egress the module converts physical coding from 64B/66B to 8B/10B encoding
- Due to higher thermal of these 10G-to-1GE SFPs **odd numbered** 10G/25G ports can run this optic (the upper row of SFP+/SFP28 ports) - It means that **max 20x 1GE ports can be supported.**
- The following port combinations are available:
  - **5x 1GE + 35x 10GE** or **10x 1GE + 30x 10GE** or **15x 1GE + 25x 10GE** or **20x 1GE + 20x 10GE**

Flexible Consumption Model (FCM)

2

Two simple suites: Essentials or Advantage

1

Hardware is separated out from new Software suites

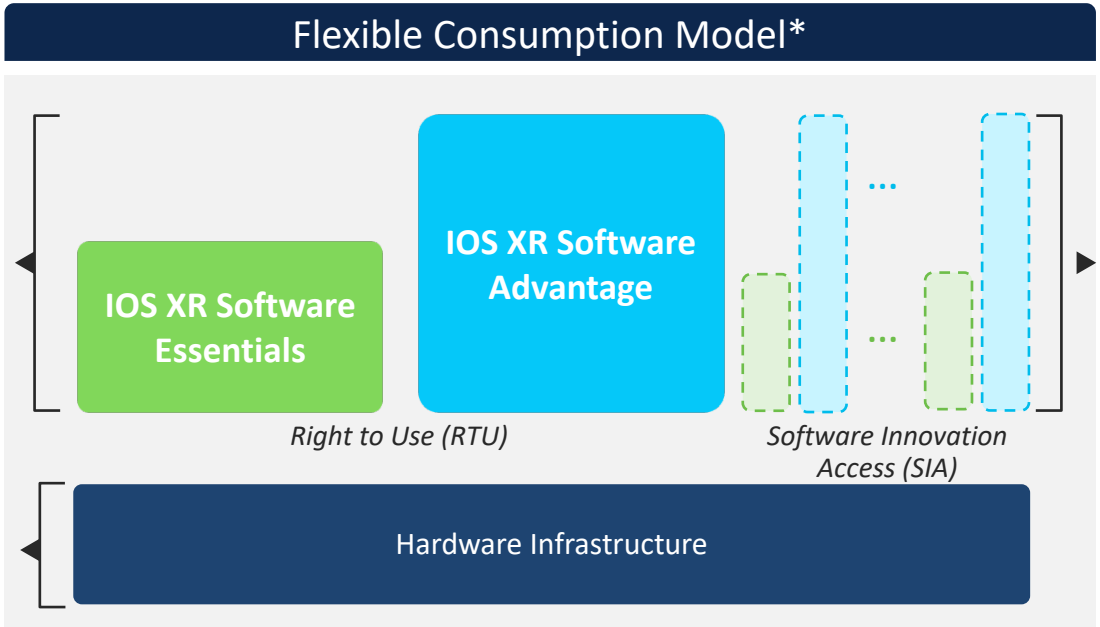

3

SIA provides pooling of licenses, continuous software innovation and investment protection.
3 years of mandatory SIA required

- Pricing Meter: 10G/100G/400G Capacity
- RTU: Perpetual license
- SIA: Subscription

* Smart Licensing registration & reporting required for FCM

Criteria for Advantage RTU License

Advanced Feature	System or Port Level	Check	Action
L2VPN	Per Port	Check all ports that are attachment circuits to an L2VPN or EVPN service (PW, VPWS, VPLS, EVPN)	All ports that match require an Advantage SW license
L3VPN	Per Port	Count the Number of VRFs that are Globally Configured	If <= 8, Essentials SW license only If > 8, Advantage SW license required
		Check all ports that have parent/subinterface configured with VRFs	All ports that match require an Advantage SW license
Peering Scale	Per System	Check FIB scale limits of IPv4: 512K routes, IPv6: 50K routes	Upon match, require an Advantage SW license for all ports in system
Traffic Engineering	Per System	Check if Traffic Engineering or Flex-Algo is enabled on the system	Upon match, require an Advantage SW license for all ports in system
Security / MACsec	Per Port	Check for all ports that are configured with MACSec	All ports that match require an Advantage SW license

NCS 540 FCM Advantage 10G RTU License

- PID: ADN-AC-10G-RTU-1 (includes Essentials) or ADV-AC-10G-RTU-1 (for upgrade from Essentials to Advantage)
- Not for actual bandwidth used but for “Up” interfaces in 10G increments
- Always maximum 1 Advantage license per 10G
- Rounding to 10G applies for 1GE or 25GE interfaces (15G→24G=2x 10G RTU, 25G→34G=3x 10G RTU, …)

Per-interface Advantage licenses:

- L3VPN: per access interface if router has >8 VRFs
- L2VPN: per access interface for: VPWS, VPLS, EVPN (for any interface in l2vpn before 7.4.2, 7.5.2)
- MACsec

Per-system Advantage licenses:

Needed for all ”Up” interfaces

- Traffic Engineering: In IGP, RSVP-TE, SRTE, Flex-Algo
- Peering scale: >512k v4 or >50k v6
- BGP features: EPE, FlowSpec*, QPBB, RPKI, BMP, Attribute Download
- Lawful Intercept, SL-API

Supporty v SP

Identify the right service level

Traditional Model to FCM Model Support Service Levels Translator

Equivalent Traditional Support Service GSPs

SW Support		SP SW Support	
Service	**Service Description**	**Service**	**Service Description**
SAS	24/7 SW Support	SPSAS	24/7 SW Support
ECMU	24X7 SW TAC with SW downloads	SPCMU	24X7 SW TAC with SW downloads

Smart Net Total Care		SP Base, HW Support	
Service	**Service Description**	**Service**	**Service Description**
-	-	SP RTF Svc(*)	SPRTF
SNT	SNTC-8X5XNBD	8X5XNBD AR Svc	SPAR1
SNTE	SNTC-8X5X4	8X5X4 AR Svc	SPAR2
SNTP	SNTC-24X7X4	24X7X4 AR Svc	SPAR3
S2P	SNTC-24X7X2	24X7X2 AR Svc	SPAR4
CS	SNTC-8X5XNBDOS	8X5XNBD Onsite Svc	SPCS
C4S	SNTC-8X5X4OS	8X5X4 Onsite Svc	SPC4S
C4P	SNTC-24X7X4OS	24X7X4 Onsite Svc	SPC4P
C2P	SNTC-24X7X2OS	24X7X2 Onsite Svc	SPC2P

FCM Support GSPs

What to Quote	
"SD"-Service	**SKU format**
SDSWK	SD-xxxx
SDSWK	SD-xxxx

What to Quote	
"SD"-Service	**SKU format**
SDRTFK(*)	SD-xxxx
SDAR1K	SD-xxxx
SDAR2K	SD-xxxx
SDAR3K	**SD-xxxx**
SDAR4K	SD-xxxx
SDCSK	SD-xxxx
SDC4SK	SD-xxxx
SDC4PK	SD-xxxx
SDC2PK	SD-xxxx

SDRTF service level is restricted for SP customers only

ASR 9000 5th Generation 2T Combo Card

A9K-20HG-FLEX-SE/TR

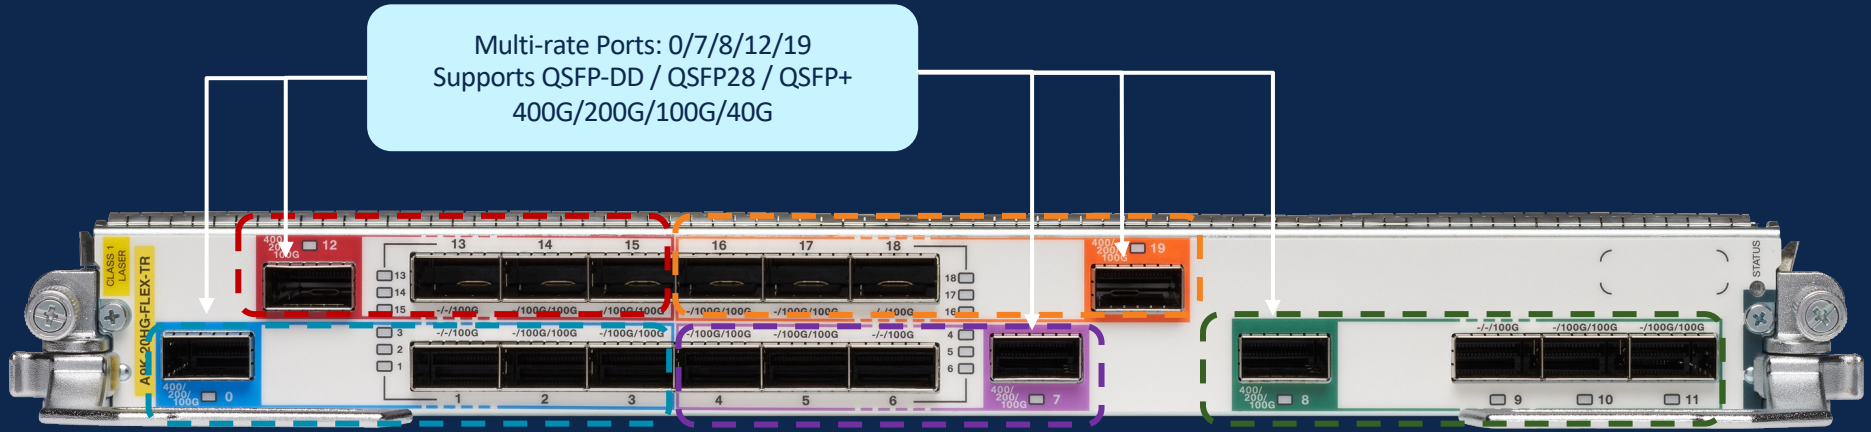

400G Ready

- 10G/25G/40G/100G/200G/400G Support

Each Slice Independently Configured as:

- 1x400G
- 1x200G + 2x100G or 2x40G
- 4x100G or 4x40G

Each 100G Breakout into 4x25G or 4x10G **Total 80x 10/25GE**

Cisco 8000

8000 Series Product Line

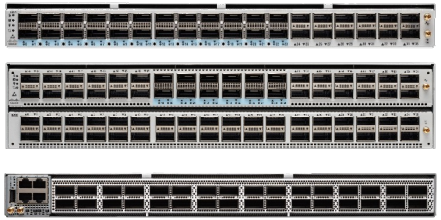

8200 SP/DC Fixed

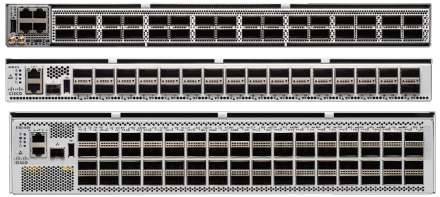

8100 DC Fixed

8800 Distributed Modular

Silicon One Device Lineup (mid-CY22)

8100 vs. 8200 Introduction

- 8200 has HBM for FIB expansion and deep buffers, 8100 does not
- Traffic passing through 8100 should stay in the building/site
- 8100 is for DC fabric (leaf, spine, and small-buffer ToR)

  Medium FIB (<1M v4) scale & smaller buffers (100s of microseconds)
- 8200 is for a wide range of roles

  SP core, aggregation, and peering

  DC leaf/spine/ToR and DCI

  Large FIB (~4M v4) and deep buffers (10s of milliseconds)

8200 Hardware Reference

	8212-48FH-M	8201-32FH	8202-32FH-M	8201-24H8FH	8201	8202
Bandwidth	19.2 Tbps	12.8 Tbps	12.8 Tbps	5.6Tbps	10.8 Tbps	10.8 Tbps
ASIC	P100	Q200	Q200	Q200	Q100	Q100
QSFP28	0	0	0	24	12	60
QSFP56-DD (400G)	48 (MACsec)	32	32 (MACsec)	8	24	12
Depth	23.6” / 600mm	23.6” / 600mm	23.6” / 600mm	23.6” / 600mm	20.1” / 511 mm	20.1” / 511 mm
Weight	42 lb / 19 kg	31 lb / 14.1 kg	42 lb / 19kg	31 lb / 14.1 kg	24 lb / 10.9 kg	42 lbs / 19 kg
CPU / Memory	Intel Broadwell 4-core with 32 GB DRAM & 128 GB SSD					
Fans	4	6	4	6	5	3
Airflow	PSI	Either	PSI / *PSE future*	Either	Either	Either
Typical/Max power	TBD	288/675W	700/1550W	200/525W	415/660W	700/1150W

Cisco Silicon One *P100*

- 19.2T routing device in 7nm
- Consistent architecture
- Highest routing performance
- SRAM and deep HBM buffers
- Lowest routing power per bit
- Increased scale vs. Q200
- 100G SerDes enable 800G

8800 Series Overview




	8804	8808	8812	8818
Rack Units	10 RU	16 RU	21 RU	33 RU
Line Cards	 48x100GbE (Q100, MACsec)  36x400GbE (Q100)	 36x400GbE (Q200, MACsec)  36x400GbE (Q200)	 34x100GbE & 14x400GbE (Q200,16x 100G MACsec)	 36x 2x400GbE (*P100*)
Capacity (28.8T LC)	115.2 Tbps	230.4 Tbps	345.6 Tbps	518.4 Tbps
Typical System Power (Q200 14.4T LC)	4.1 KW	8.0 KW	13.8 KW	17.4kW (18 LCs) 13.9kW (12 LCs)

800G vs. 800 GbE vs. 2x 400 & 8x 100 GbE

- P100 and G100 introduce 100G SerDes and 800G ports
- There is no single-fiber pair 800G in FY23
- 800G ports are currently 8x 100 or 2x 400 (maybe 4x 200)
- Interoperable 800G in next generation ASICs (~2024)

  Ethernet Technology Consortium standard first, then IEEE (very similar for initial PMDs)
- Use the terms carefully

  “800G” ports means 800G bandwidth, 8x100, 2x 400G, or proprietary 800G

  “2x 400 GbE” or “8x 100 GbE” IEEE standards shipping in first generation

  “800 GbE” means standardized single-channel (future – don’t use this yet)

8000 Series 800G Ports

- First shipped in mid 2022 – 32x 800G
- Second board will be 36x 800G line card
- First-gen 800G ports are 8x 100 GbE or 2x 400 GbE
  - 2x LC connector to two 400 GbE FR4
  - MPO12 to two 400 GbE DR4
  - MPO16 to eight 100 GbE

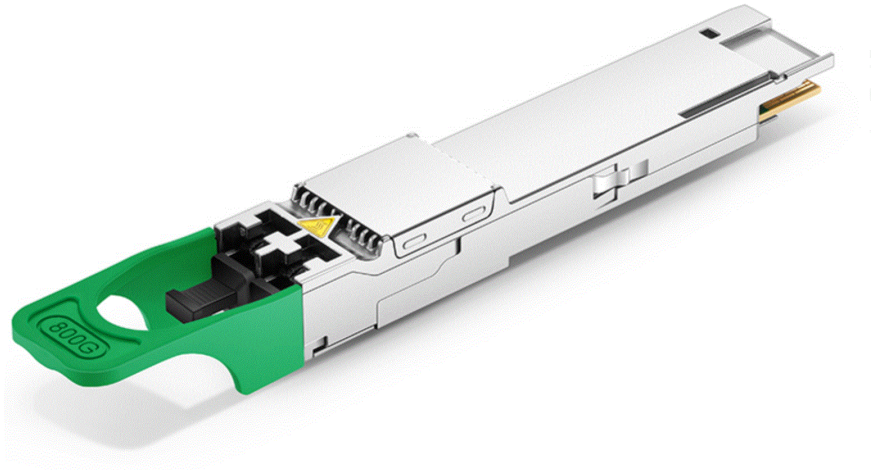

For more info: https://blogs.cisco.com/tag/qsfp-dd800

NCS 540

NCS 540 Family in 2022

Small	Medium	Large	Fronthaul
Darwin Fitzroy, Galapagos, Beagle Meerkat[7.5.2], Finches[7.8.1]	**Tortin & Big Bend** Tortin 16/32, Everglades, Crater Lake, Acadia, Denali, Olympic[7.5.2]	**Arches**	**Felidae** Lion, Jaguar
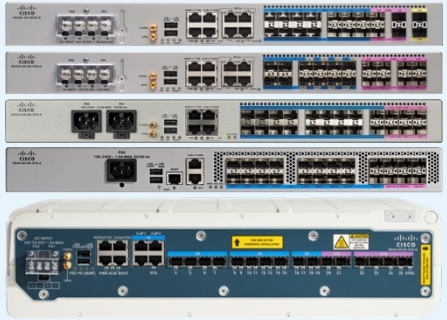	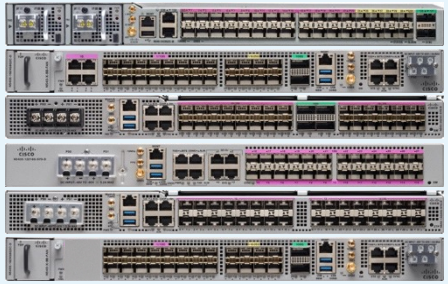		
64G or 104G QUX-64/120 1/10/25GE (2x SFP28) PSU: Fixed AC/DC 1+1 or non-redundant I-Temp S2S or F2L Airflow, Fixed Fans Depth: 23cm C-Temp PID: Single AC/DC PSU, F2L Airflow, no timing Passive Cooling (Fanless) PID: Fixed DC PSU 1+1, 2.5RU, depth 38cm **Shipping** 7.3.1/7.4.1/7.5.2/7.8.1	**136G to 300G** QAX-160/300 1/10/25/40/100GE (2/4x QSFP28) PSU: FRU 1+1 AC/DC or Fixed 1+1 DC/1 AC I-Temp or C-Temp F2B or S2S Airflow, Modular or Fixed Fans GNSS Receiver MACsec Depth: 25-28cm **Shipping** 6.3.2/6.5.2/7.0.1/7.5.2	**800G** Q2A 1/10/25/40/50/100/200/400GE (2x QSFP56-DD, 8x SFP56) PSU: FRU 1+1 AC/DC I-Temp F2B Airflow, Modular or Fixed Fans GNSS Receiver MACsec Depth: 30cm **Shipping** 7.4.1	**300G or 900G** QAX-300/J+ 1/10/25/40/100GE (2/4x QSFP28) 2/24x 10/25GE TSN 802.1Qbu 12/24x CPRI 3-8 PSU: FRU 1+1 AC/DC I-Temp or C-Temp F2B Airflow, Modular or Fixed Fans GNSS Receiver MACsec* Depth: 35/55cm **Shipping** 7.3.2

* Roadmap

NCS 540

				FCS	NPU	NPU Gbps	Max IO Gbps	PSUs	Fans	Airflow	Temp	Conf. Coat	Depth [cm]	GNSS	PTP Class	MACsec	CPRI	TSN	QDD56	QSFP28	SFP56	SFP28	SFP10	SFP	RJ45
Large	Arches		N540-24Q8L2DD-SYS	7.4.1	Q2A	800	1000	M	F	F2B	I		30	Y	C	Y			2		8	24			
Medium	Tortin	Tortin 16	N540-ACC-SYS	6.5.2	QAX	300	640	M	M	F2B	I		26	Y	B	Y				2		8	24		
		Tortin CC 16	N540X-ACC-SYS	6.5.2	QAX	300	640	M	M	F2B	I	Y	26	Y	B	Y				2		8	24		
		Tortin 32	N540-24Z8Q2C-SYS	6.3.2	QAX	300	640	M	M	F2B	I		26	Y	B	Y				2		8	24		
Medium	Big Bend	Everglades	N540X-16Z4G8Q2C-D/A	7.0.1	QAX	300	564	F[1]	F	S2S	I	Y	28	Y	C	Y				2		8	16		4
		Olympic	N540X-16Z8Q2C-D	7.5.2	QAX	300	560	F[D]	F	S2S	I	Y	28	Y	C	Y				2		8	16		
		Crater Lake	N540-28Z4C-SYS-D/A	7.0.1	QAX	300	680	F[1]	F	F2B	C		25		B					4			28		
		Acadia	N540X-12Z16G-SYS-D/A	7.0.1	QAX	160	136	F[1]	F	S2S	I	Y	25	Y	C								12	12	4
		Denali	N540-12Z20G-SYS-D/A	7.0.1	QAX	160	140	F[1]	F	F2B	C		25		B								12	20	
Small	Darwin	Fitzroy	N540X-4Z14G2Q-D/A	7.4.1	QUX	120	104	F[2]	F	S2S	I	Y	23		C							2	4	14	4[x]
		Galapagos	N540X-8Z16G-SYS-D/A	7.3.1	QUX	120	104	F[2]	F	S2S	I	Y	23		C								8	12+8[c]	4
		Beagle	N540X-6Z18G-SYS-D/A	7.3.1	QUX	64	78	F[2]	F	S2S	I	Y	23		C								6	18	
		Finches	N540-6Z18G-SYS-D/A	7.8.1	QUX	64	78	F[1]	F	F2L	C		23		-								6	18	
		Meerkat	N540-6Z14S-SYS-D	7.5.2	QUX	64	80	F[D]	-	-	I		38		C								6	10+6[c]	4
Fronthaul	Felidae	Lion	N540-FH-AGG-SYS	7.3.2	J+	900	1000	M	M	F2B	C		55	Y	C	Y[R]	Y	Y		4		24			
		Jaguar	N540-FH-CSR-SYS	7.3.2	QAX	300	550	M	F	F2B	I		35	Y	C	Y[R]	Y	Y		2		6	20		

M: modular F: fixed F[1]: dual DC or single AC F[2]: dual DC or dual AC F[D]: dual DC R: roadmap x: combo ports c: using cSFP

NCS 540	Interfaces	Throughput	Timing	IOS XR FCS	DRAM	Power
Tortin 32GB/16GB N540-24Z8Q2C-SYS N540(X)-ACC-SYS 	2x 100/40GE 8x 25/10/1GE 24x 10/1GE	300G Max Interfaces: 640G	GNSS Class B 1pps/10MHz/ToD	32GB: 6.3.2 Apr 2018 16GB: 6.5.2 Jan 2019	32GB 16GB	Modular: 1+1 AC/DC
Everglades N540X-16Z4G8Q2C-D/A 	2x 100/40GE 8x 25/10/1GE 16x 10/1GE 4x 1GE Copper	300G Max Interfaces: 564G	GNSS Class C 1pps/10MHz/ToD BITS	7.0.1 LNT August 2019	8GB	Fixed: 1 AC 1+1 DC
Crater Lake N540-28Z4C-SYS-D/A 	4x 100/40GE 28x 10/1GE	300G Max Interfaces: 680G	Class B* 1pps/10MHz/ToD BITS	7.0.1 LNT August 2019	8GB	Fixed: 1 AC 1+1 DC
Acadia N540X-12Z16G-SYS-D/A 	12x 10/1GE 12x 1GE 4x 1GE Copper	160G Max Interfaces: 136G	GNSS Class C 1pps/10MHz/ToD BITS	7.0.1 LNT August 2019	8GB	Fixed: 1 AC 1+1 DC
Denali N540-12Z20G-SYS-D/A 	12x 10/1GE 20x 1GE	160G Max Interfaces: 140G	Class B* 1pps/10MHz/ToD BITS	7.0.1 LNT August 2019	8GB	Fixed: 1 AC 1+1 DC
Fitzroy N540X-4Z14G2Q-D/A 	2x 25/10/1GE 4x 10/1GE 10x 1GE 4x 1GE Combo SFP/RJ45	120G Max Interfaces: 104G	Class C 1pps/10MHz/ToD	7.4.1 LNT August 2021	8GB	Fixed: 1+1 AC 1+1 DC
Galapagos N540X-8Z16G-SYS-D/A 	8x 10/1GE 4x 1GE SFP 4x 1GE RJ45 8x 1GE SFP or 16x 1GE cSFP	120G Max Interfaces: 104G	Class C 1pps/10MHz/ToD	7.3.1 LNT February 2021	8GB	Fixed: 1+1 AC 1+1 DC
Beagle N540X-6Z18G-SYS-D/A 	6x 10/1GE 18x 1GE	64G Max Interfaces: 78G	Class C 1pps/10MHz/ToD	7.3.1 LNT February 2021	8GB	Fixed: 1+1 AC 1+1 DC

* Class C validation on Crater Lake and Denali not planned now – please reach to PM team

NCS 540 Large

Arches

N540-24Q8L2DD-SYS

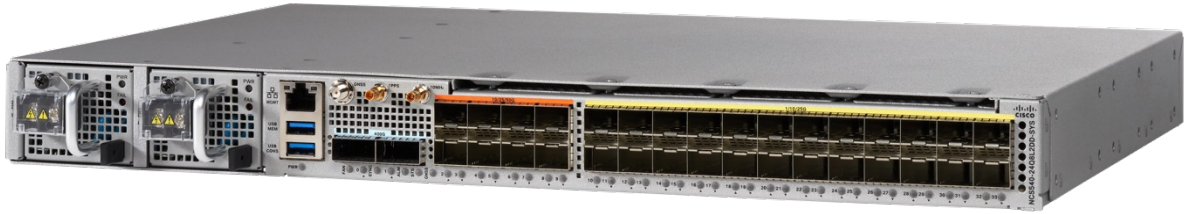

- Broadcom Qumran 2A (J2 family)
- 800Gbps/600Mpps, Max IO 1T, 2GB Packet Buffer
- CPU 4C Intel Xeon, 16GB DRAM, 64GB SATA
- 1RU, Depth 299mm
- Redundat Modular AC/DC 1+1 power supplies
- Front to Back airflow
- Fixed redundant fans
- I-Temp -40C to +65C
- G.8273.2 Class C & GNSS Receiver

2x 400GE QSFP56-DD

8x 50GE SFP56

24x 25GE SFP28

QDD-400G-ZR-S

QDD-400G-ZRP-S

NCS 540 Large
N540-24Q8L2DD-SYS

Q2A 800G @600Mpps	Intel 4C CPU 16GB DRAM	Power: Modular 1+1 DC/AC	F2B Airflow 5+1 Fans
IOS XR 7.4.1	MACsec on QDD56/SFP56	I-Temp -40/+65C	G.8273.2 Class C

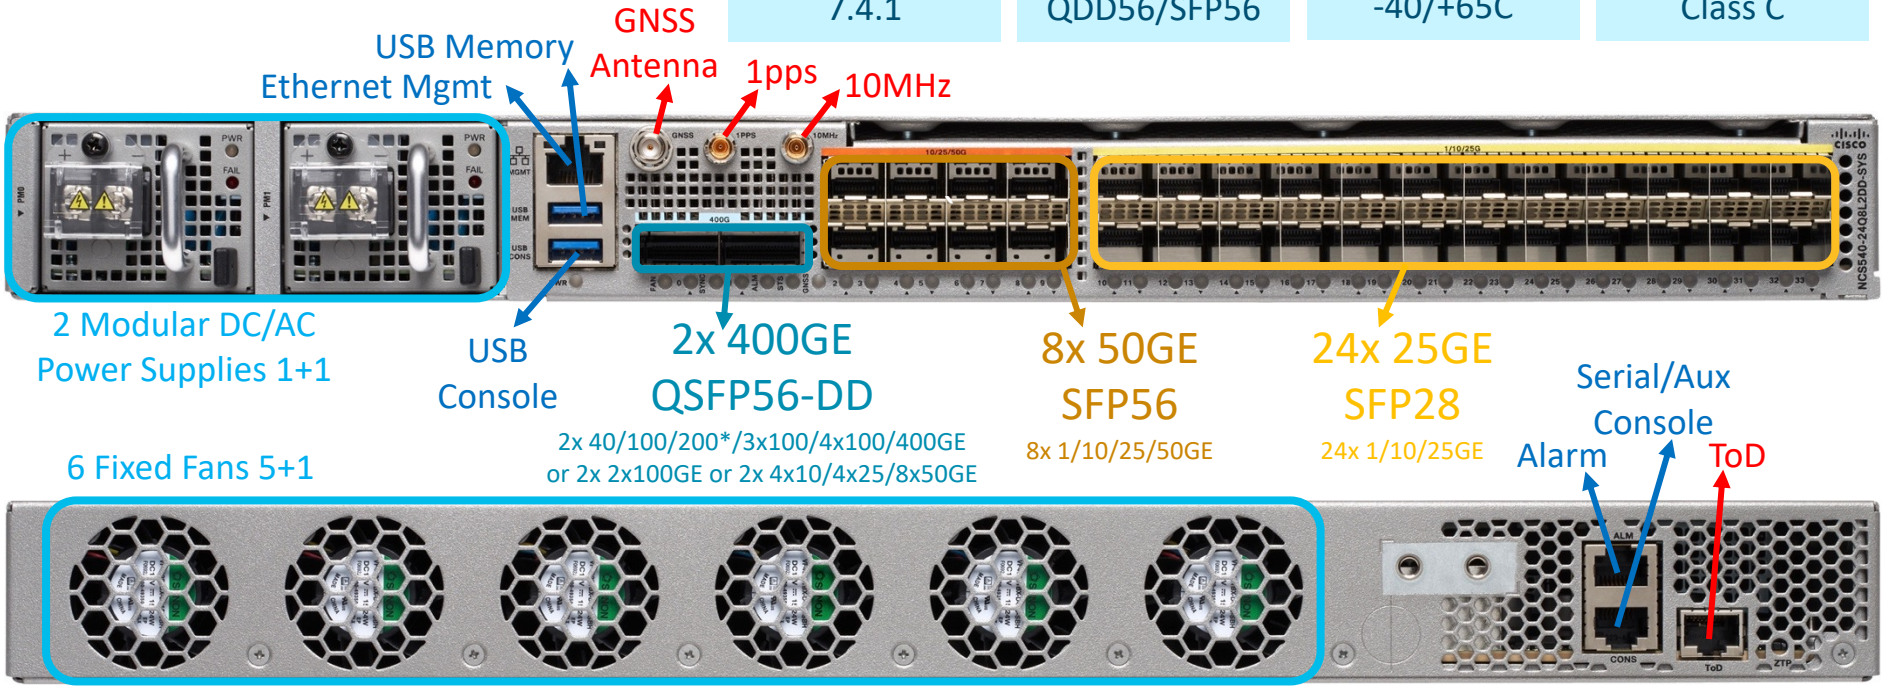

*roadmap

NCS 540 Olympic
N540X-16Z8Q2C-D

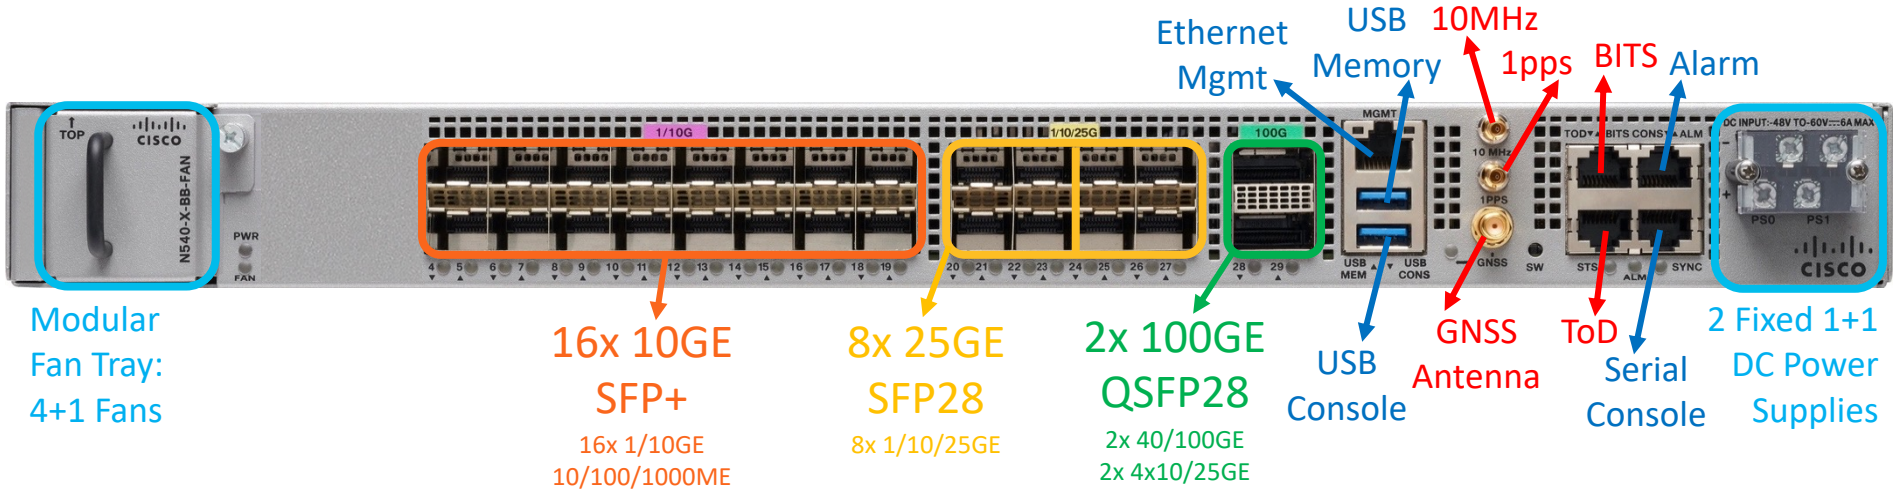

QAX 300G	I-Temp -40/+70C	Power: Fixed 1+1 DC	GNSS Receiver
8GB DRAM	Conformal Coated	S2S Airflow	Class C

Note for 25GE SFP28 ports: There are 2 Quads. All ports in one Quad operate in 25G (default) or 1/10G mode. Configurable.

NCS 540 Finches
N540-6Z18G-SYS-A/D

- “Beagle light”
- The same as Beagle but:
  - Single AC PSU or single dual-feed DC PSU
  - Front to Left airflow (Front to Back with air baffle)
  - No I-Temp = C-Temp 0°C to 55°C @300m
  - No Conformal Coating
  - No timing = no SyncE, no PTP
  - No additional surge protection for DC PSU

NCS 540 Finches
N540-6Z18G-SYS-A (-D)

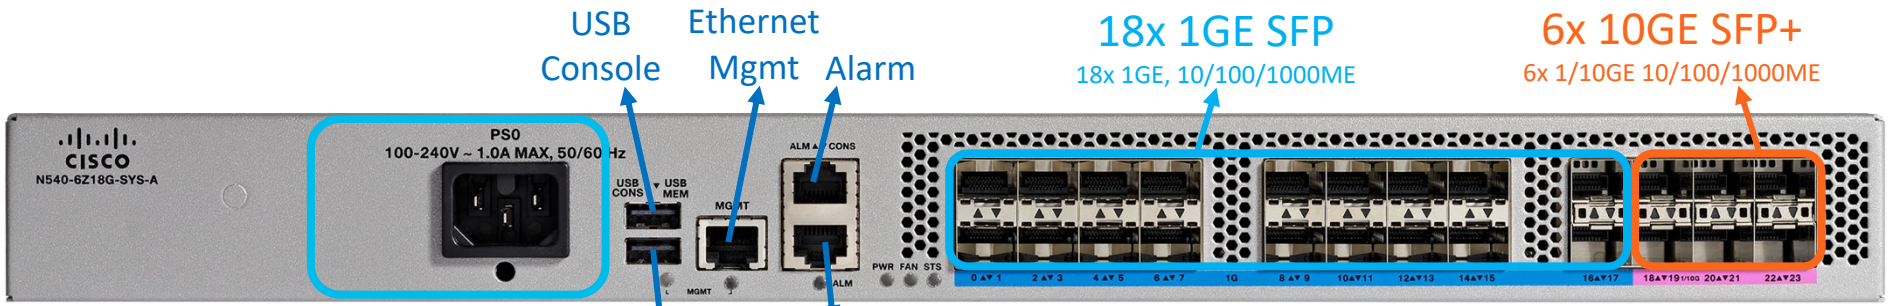

QUX 64G	C-Temp 0C/+55C	Power: 1x Fixed AC or 2-feed DC
8GB DRAM		F2L Airflow

FCM

Flexible Consumption Model

- licenses used per active port
- license pooled from smart account

NCS 5500, 5700

*(*BRKSPG-2397*)*

NCS 5500/5700 – Fixed Portfolio

High Scale Aggregation evolution

NCS5500 Products (Q-MX, J, J+)

1G | 10G | 25G

NCS 5501/SE

40G |100G

NCS 5502/SE

25G | 40G | 100G

NCS-55A1-48Q6H

NCS-55A1-24Q6H-S/SS

40G | 100G

NCS 55A1-36H-S/SE

NCS 55A1-24H

10G | 25G | 100G

NCS 55A2-MOD-S/SE

NCS5700 Products (J2/J2C/Q2C/J2C+)

NCS-57B1-6D24

- 400G ZR/ZR+
- 1RU; 4.8 Tbps throughput
- 24x100G + 6x400G
- MACSEC, Timing

NCS-57B1-5DSE

- 400G ZR/ZR+
- 1RU; 4.4 Tbps throughput
- 24x100G + 5x400G
- MACSEC, Timing
- External TCAM

NCS-57D2-18DD-S

- 400G ZR/ZR+
- 2RU; 7.2 Tbps throughput
- Flexible 66 ports 2x400G + 16x400G/64x100G
- MACSEC*, IPSEC*, Timing

NCS-57C3-MOD-S

- 400G ZR/ZR+
- 3RU; 2.4T throughput
- Fixed: 48x1/10/25G + 8x100G QSFP28
- 3 x MPA: 2x800G + 1x 400G
- MACSEC, Timing

NCS-57C3-MOD-SE

- 400G ZR/ZR+
- 3RU; 2.4T throughput
- Fixed: 48x1/10/25G + 4x100G QSFP28
- 3 x MPA: 2x800G + 1x 400G
- MACSEC, Timing
- External TCAM

NCS-57C1-48Q6D-S

- 400G ZR/ZR+
- 1RU; 2.4T throughput
- 32x1/10/25G + 16x1/10/25/50G + 6x400G
- MACSEC, Timing

Segment Routing EVPN MACSec Timing 400G ZR/ZRP

38

NCS 5500/5700 – Modular Portfolio

High Scale Aggregation evolution

NCS5500 Products (J, J+)

40G |100G

NC55-24H12F-SE

NC55-18H18F

100G

NC55-36X100G-S

NC55-24X100G-SE

NC55-6x200-DWDM-S

40G | 100G

NC55-36X100G-A-SE

Modular

NC55-MOD-A-S/SE

NC55-MOD-A-SE

10G | 25G | 100G

NC55-32T16Q4H-A

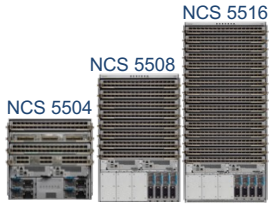

- Segment Routing
- EVPN
- MACSec
- Timing
- 400G ZR/ZRP

NCS5700 Products (J2)

NC57-24DD

- 400G ZR/ZR+
- 24x400G,
- Through put 9.6 Tbps
- No eTCAM

NC57-18DD-SE

- 400G ZR/ZR+
- 18x400G, 30x200G/100G
- Through put 7.2 Tbps
- External TCAM

NC57-36H6D-S

- 400G ZR/ZR+
- 100G, 400G
- Throughput 4.8 Tbps
- Timing, MACSEC,

NC57-36H-SE

- 400G ZR/ZR+ (1x100G mode)
- 100G
- Throughput 3.6 Tbps
- External TCAM

NC57-MOD-S

- 400G ZR/ZR+
- 10G, 25G, 50G, 100G, 400G
- Throughput 4.8 Tbps
- Timing, MACSEC, 800G-MPA

cisco Live!

NCS5500/5700 – NPU Evolution

	Jericho	Jericho +	Jericho2	Jericho2C	Jericho2C+
Bandwidth	720G	900G	4.8T	2.4T	7.2T
Power/100G	16.6W	16.6W	7.3W	5-6.7W	6.3W
Performance (pps)	720M	835M	2B	1B	2.83B
OCB	16MB	16MB	32MB	32MB	32MB
Buffer	4GB (GDDR)	4GB (GDDR)	8GB (HBM)	4GB (HBM)	8GB (HBM)
VOQ	96K	96K	64K per core	128K per core	256K per core
Counters	256K	256K	384K	192K	384K
Network IF	24x 25G+36x 12.5G	48x25G+24x12.5G	96x 50G	32x50G+96x25G	144x 50G
Fabric IF	36x 25G	48x 50G	112x 50G	48x 50G	192x 50G
MC Groups	-	128K	256K	256K	256K
Timing / Encryption	Class B / No	Class B / No	Class B / No	Class C / No	Class C / Yes

Database Overview – LEM/LPM

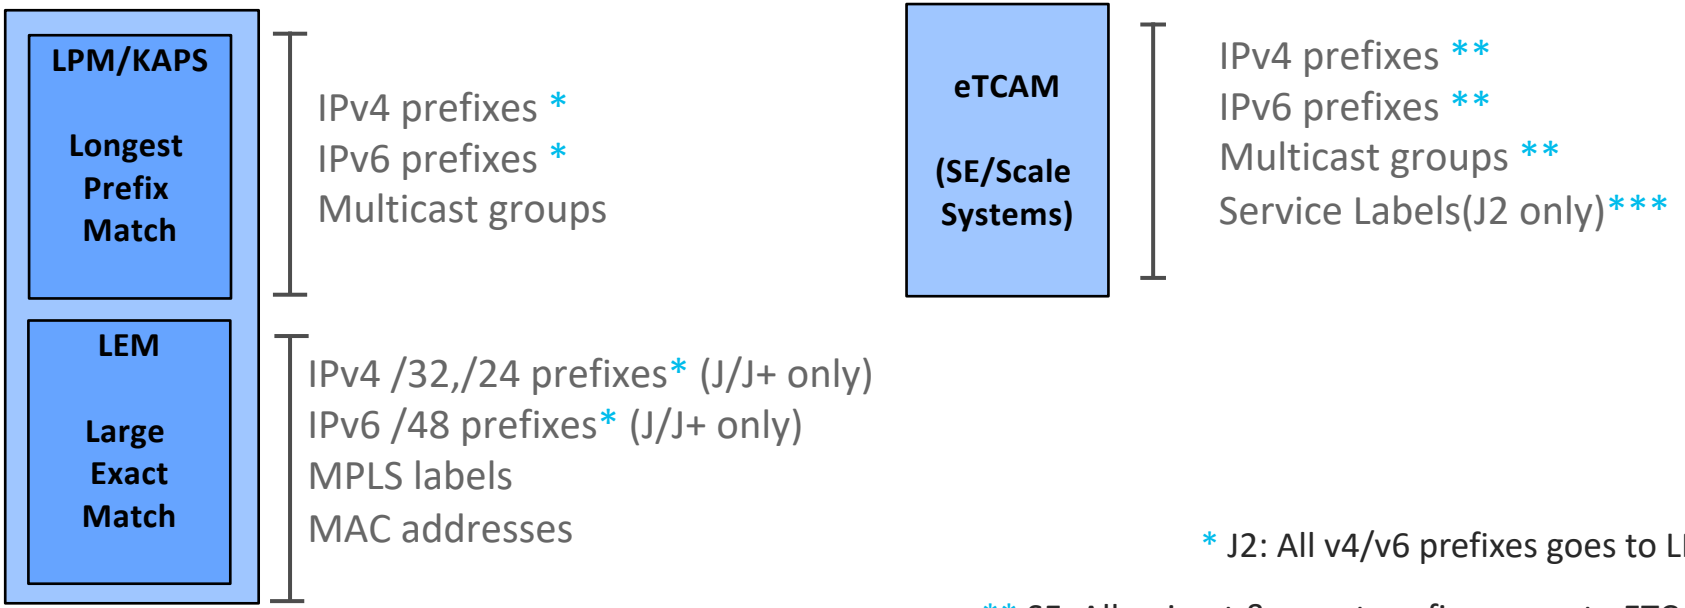

* J2: All v4/v6 prefixes goes to LPM

** SE: All unicast & mcast prefixes goes to ETCAM

*** J2-SE: Label in ETCAM for specific scenarios

J2

NCS-57B1-6D24 / NCS-57B1-5DSE

- 1RU: 24 ports QSFP + 6/5 ports QSFP-DD
  - 24x 40G/100G + (6 or 5)x 400G/200G
- 1x Jericho2 Forwarding ASIC (SoC) BCM88690
  - 2TPPS / 4,800 Gbps
- Base and Scale versions
  - Base: 6x QSFP-DD / Scale: 5x QSFP-DD ports
- IOS XR 7.3.1
  - XR7, Native only

NCS-57C1-48Q6D-S

- 1RU: 32 ports SFP28 + 16 SFP56+ 6 ports QSFP-DD
  48x 1G/10G/25G + 6 x 400G
- 1x Q2C Forwarding ASIC (SoC) BCM88820
  1B PPS / 2,400 Gbps
- **Base version**
- IOS XR 7.5.2
  XR7, Native only

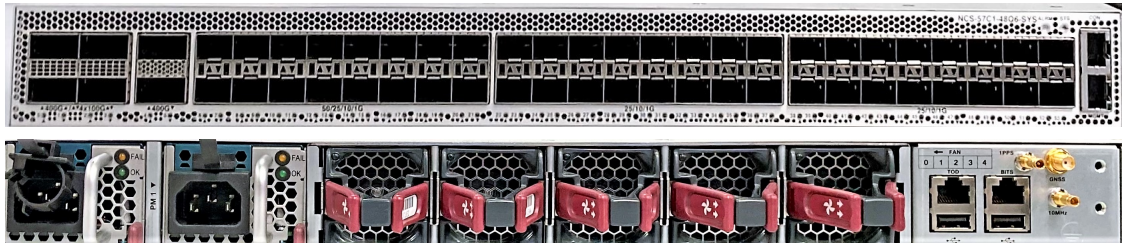

NCS-57C1-48Q6D-S

Front Panel

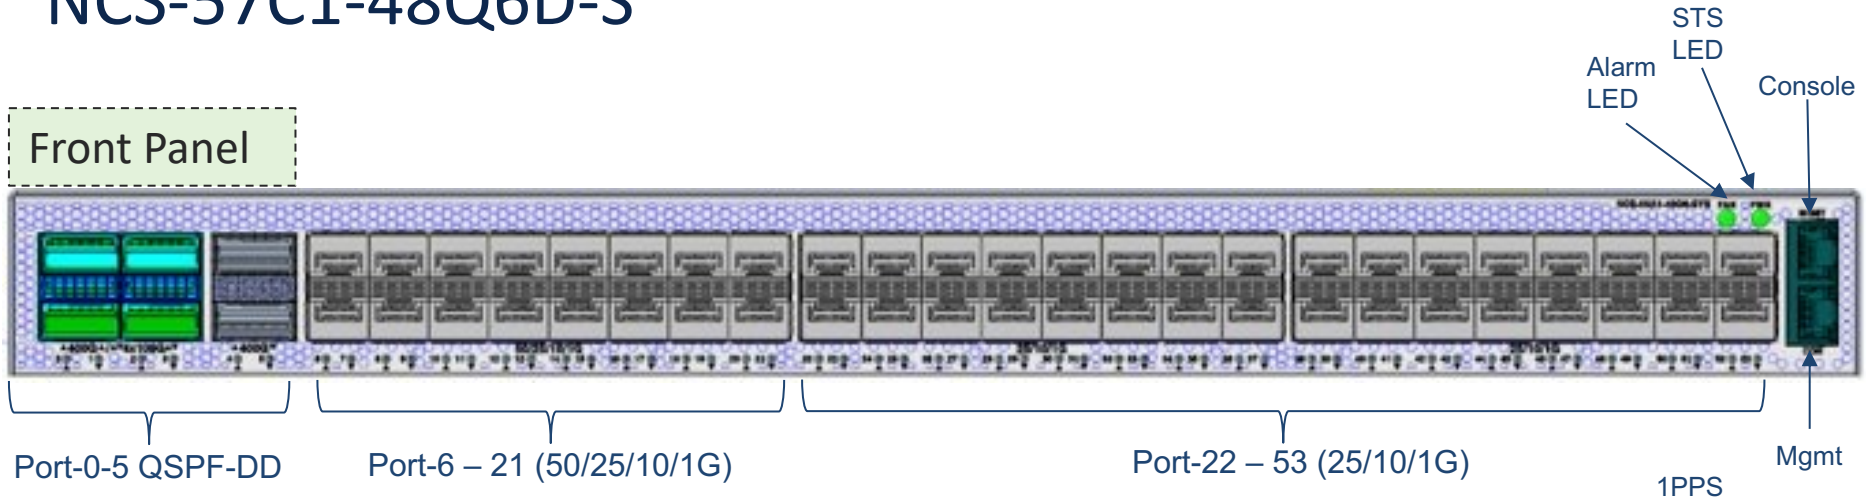

BackPanel

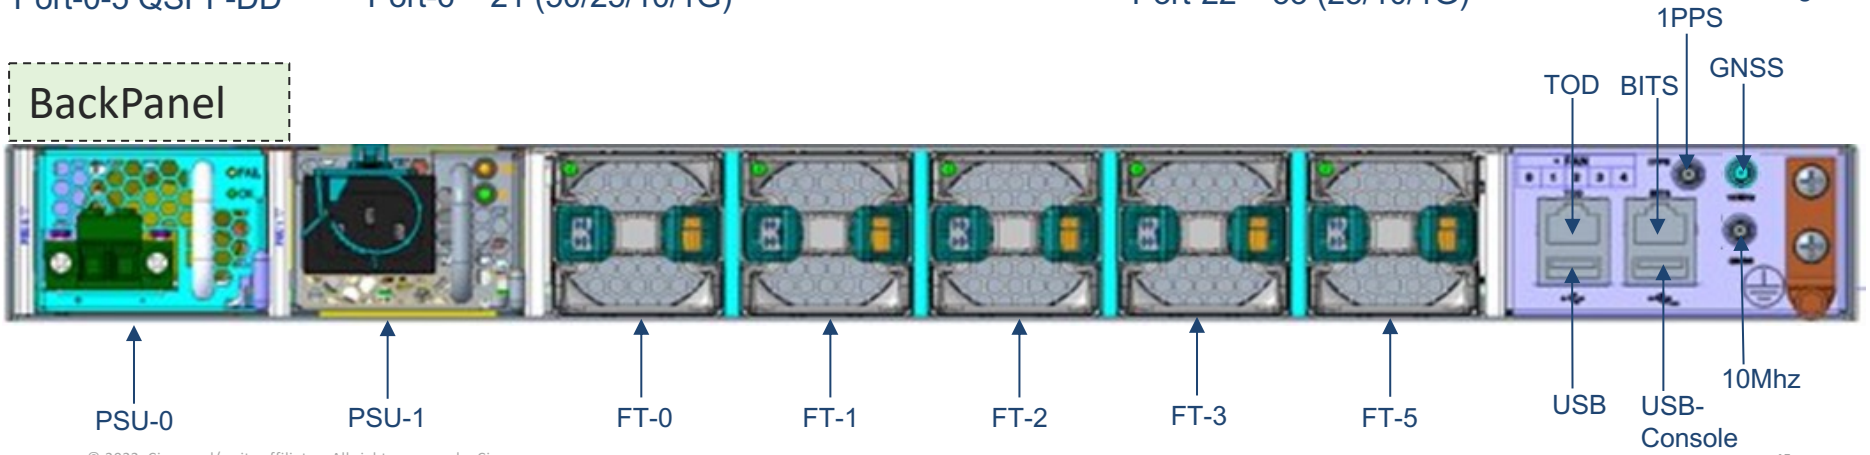

NCS-57C3-MOD

- **Code name: Eyrie**
- Compact 3RU, 284mm depth, F2B air-flow, dual RP, dual PSU, 6x fan trays
- Flexible with multiple interfaces support : 1G - 400G
- Ideal for network use cases like Mobile backhaul, Core/LSR, Peering etc.
- Platform will also support RON, PLE & cnBNG
- 400G ZR/ZRP, CFP2-DCO support via MPA
- MPA's supported:

4x QSFP-DD	**12x SFP56**	**1xCFP2 + 1xQSFP-DD**
PLE	**8x100G QSFP28**	**All existing 400G MPA's**

For more information, please refer to NCS 57C3 data sheet

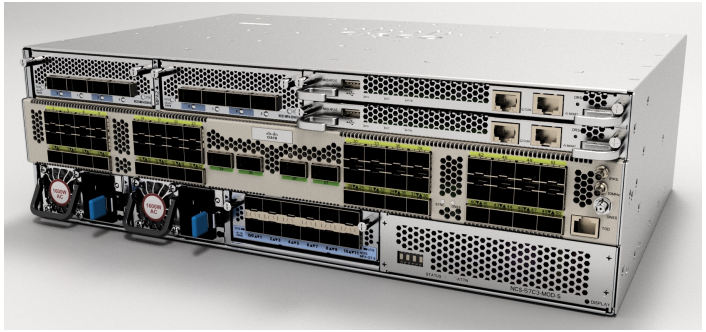

Quick Facts

Capacity	Base: 4.0T | Scale: 3.6T (Oversubscribed)	
NPU	1x Jericho2C (2.4 T)	
Port Configuration	2x MPAs (800G) + 1x MPA (400G) + 48x SFP28 + 8/4x QSFP28	
Power	Base Typical (25°C): 445W Maximum (50°C): 675W	Scale Typical (25°C): 485W Maximum (50°C): 735W
Temperature Support (at 1800m)	0-50°C 0-45°C (with NC57-MPA-2D4H-FC & low-powered 400G optics)	
SW Release	7.4.1 (Shipping)	
Hardware capabilities	MACSEC, Class C Timing, Built-in GNSS, Redundant RP	

Segment Routing

Simplicity Always Prevails

SR
Segment Routing

Segment Routing provides complete control over the forwarding paths by combining simple network instructions. It does not require any additional protocol. Indeed, in some cases it removes unnecessary protocols simplifying your network

Reduced Time to Deploy

Simplicity reduces time to deploy

- 60% reduction in internal testing (qualification) vs previous network design
- 4x improvement for software upgrade with fabric-style SP architectures

Better Productivity

Simplicity increases productivity

- 48% reduction in troubleshooting efforts vs previous network design

Reduced Capex

Low-End platforms also support SR

- 66% reduction in CapEx by optimizing the usage of feature-rich / higher-cost platforms only where it is needed, and using lighter platforms for simpler access / pre-aggregation / backhaul

* Actual business impact of Segment Routing deployment done by a large service provider in EMEA

Segment Routing

The path is in the packet

Segment Routing

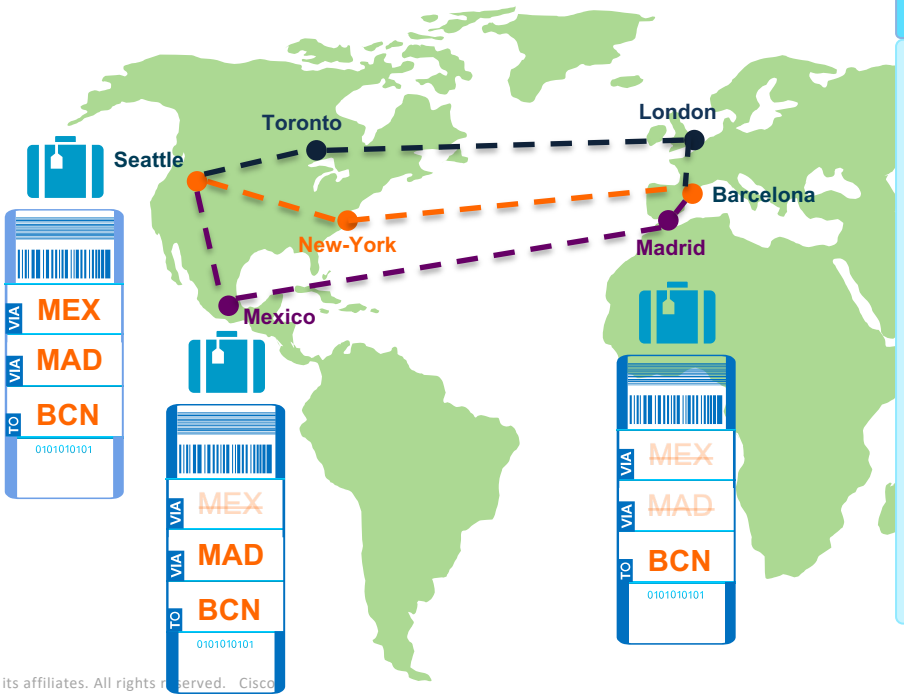

Intent – Route the luggage to Barcelona via Mexico and Madrid

1. A unique and global luggage tag is attached to the luggage with the list of stops to the final destination
2. At each stop, the luggage is simply routed to the next hop listed on the luggage tag

RESULT: **Path can be controlled**
Simple and scalable

Segment Routing

SR Prefix and Adjecency SIDs - Combining Segments

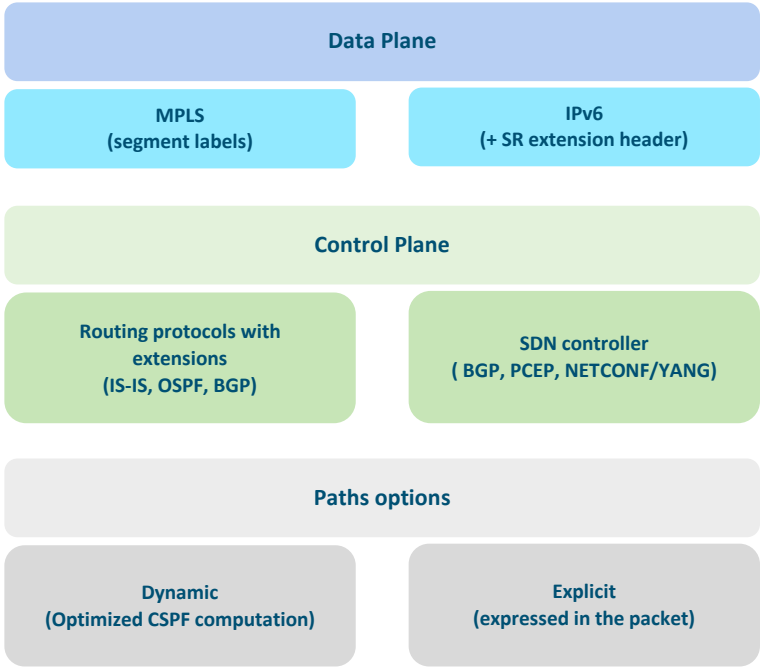

- Source Routing paradigm
  - Stateless IP fabric !!!

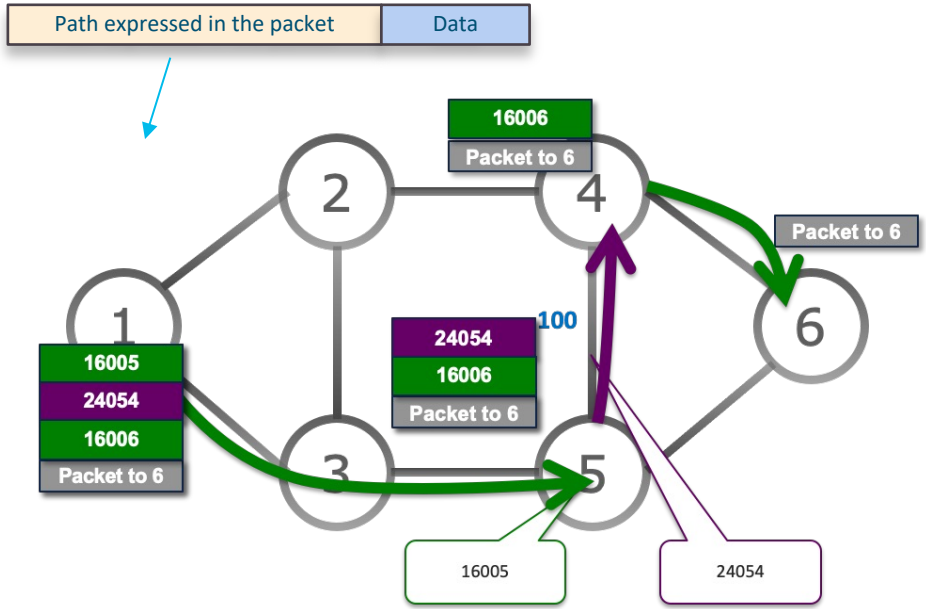

Understanding Today's Service Creation

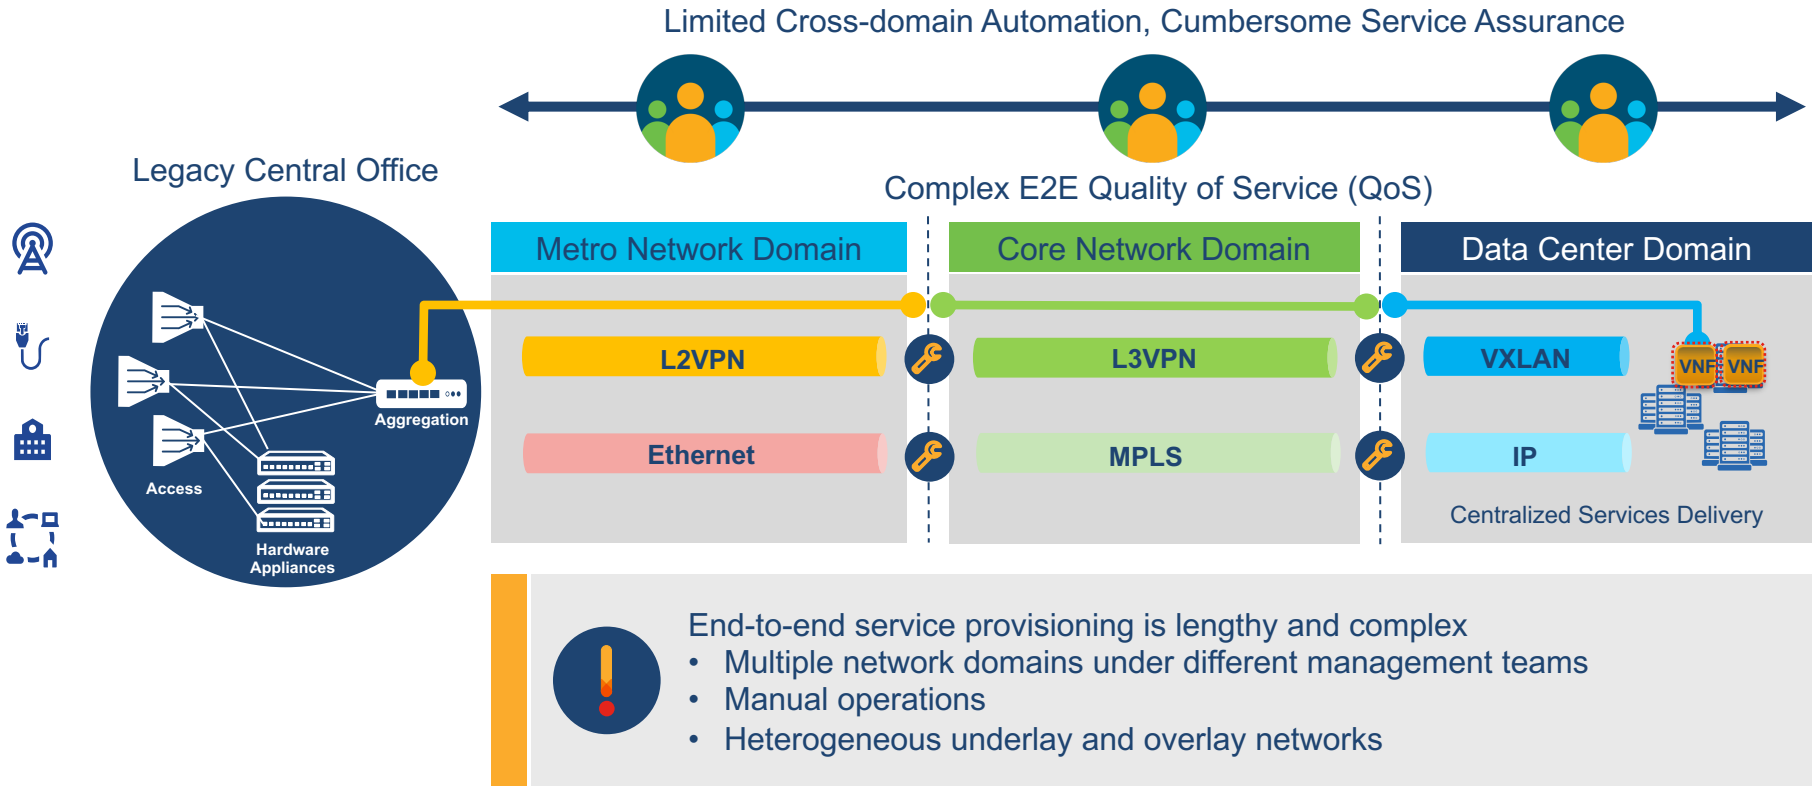

SR-MPLS: SDN ready
“Network as a Fabric” for Service Creation

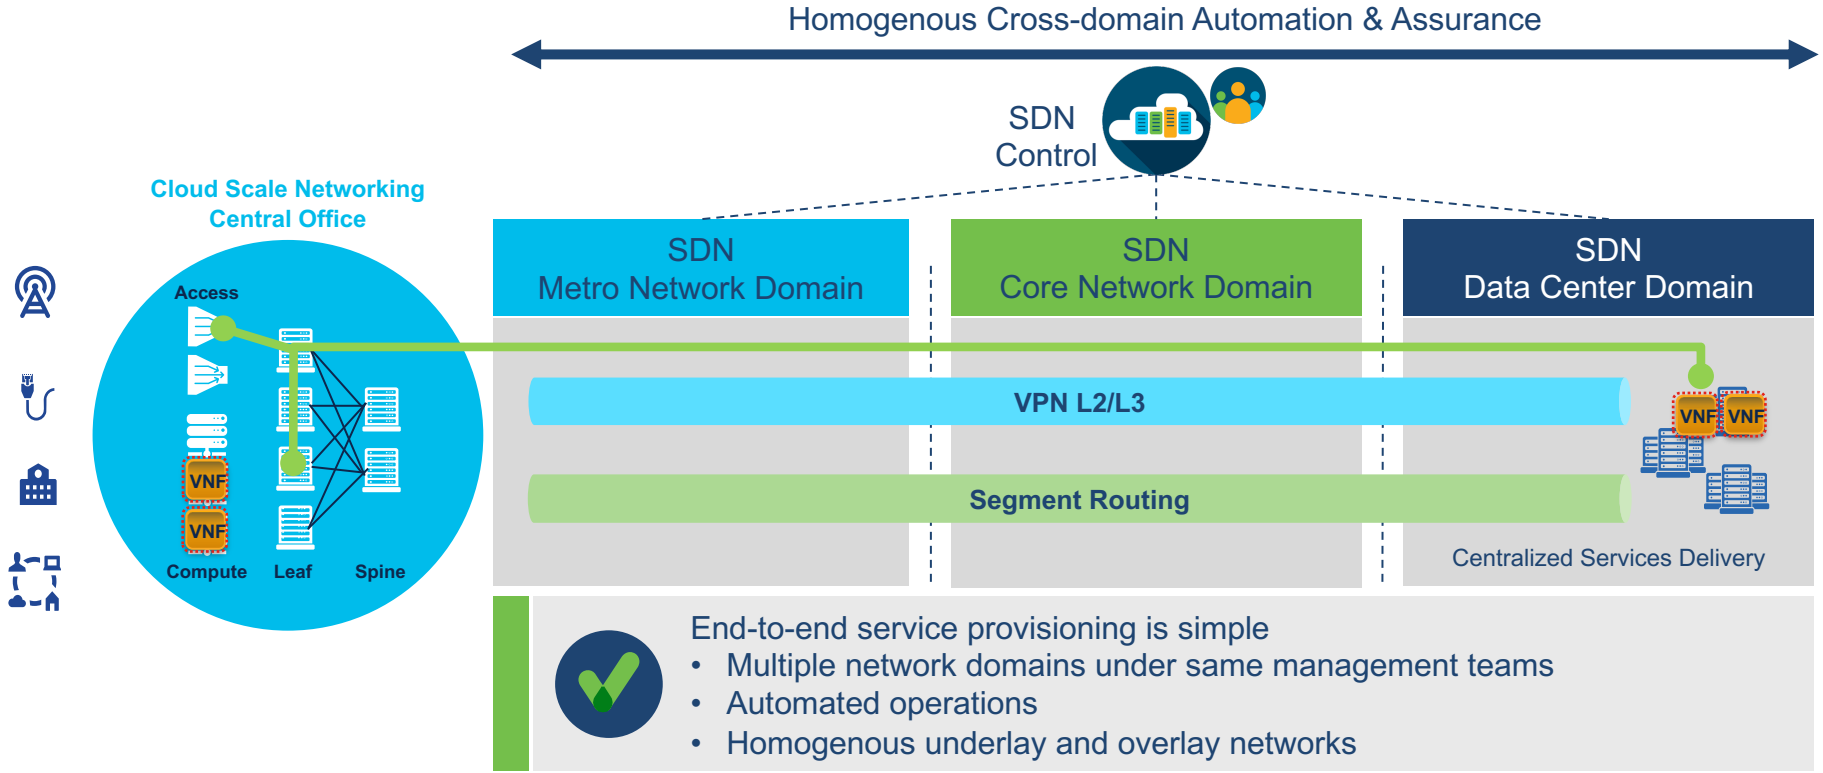

SRv6 Path to Simplicity

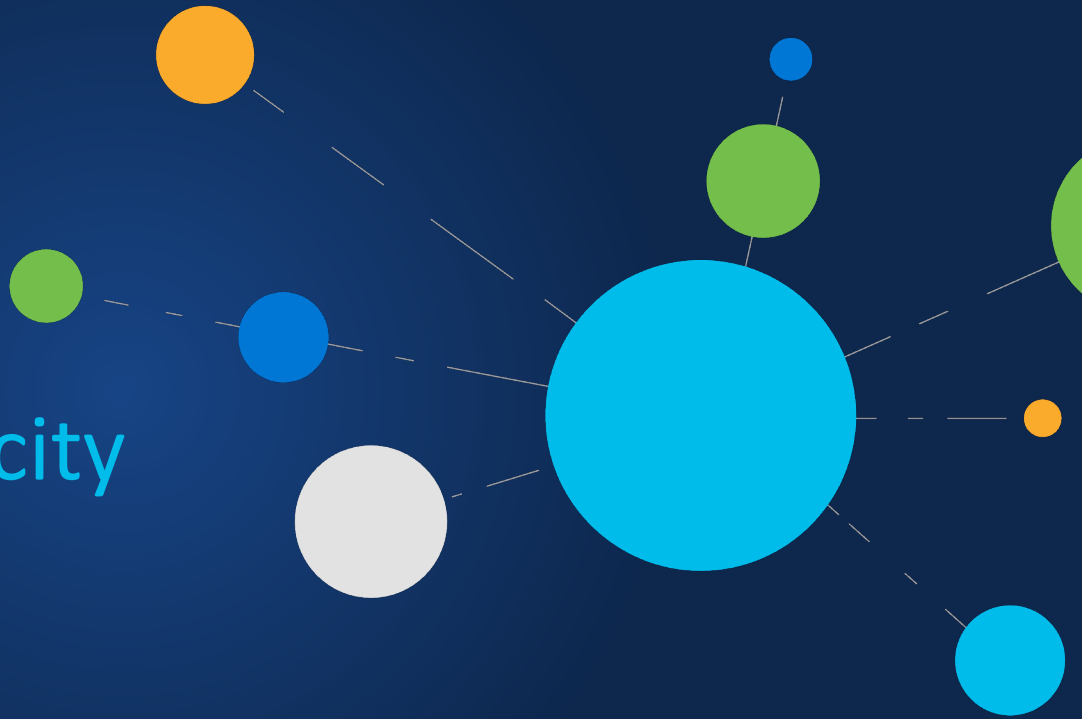

SRv6: SDN, NfV, 5G ready
“Network as an API” for Service Creation

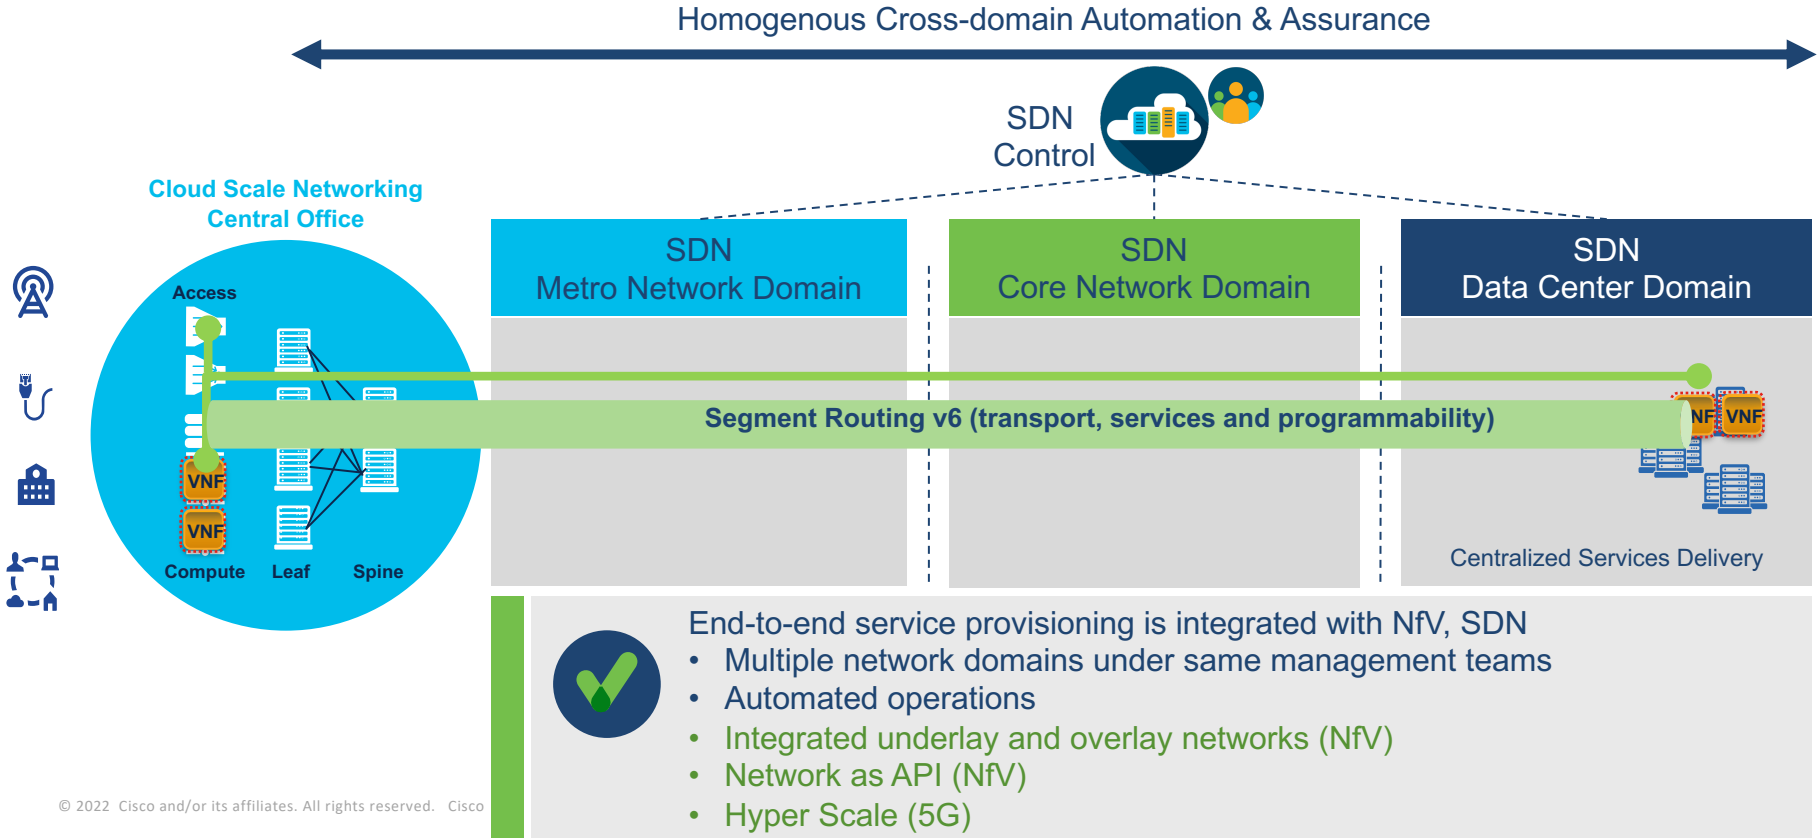

End-to-end service provisioning is integrated with NfV, SDN

- Multiple network domains under same management teams
- Automated operations
- Integrated underlay and overlay networks (NfV)
- Network as API (NfV)
- Hyper Scale (5G)

SRv6 Solution

Network Functions	IPv4	IPv6
Reachability	IPv4 Header	IPv6 Header
Engineered Load Balancing	MPLS Entropy Label, VxLAN UDP	IPv6 Header
VPN	MPLS VPN's, VxLAN	IPv6 Header
Traffic Engineering	RSVP-TE, SR-TE MPLS	IPv6 Header
Source Routing	SR-TE MPLS	IPv6 Header
Service Chaining	NSH	IPv6 Header

IPv6 Address 128bits
IPv6 Flow Header
- Engineered Flow optimization

SRv6 Header
- Source-Routing
- Traffic Engineering
- VPN
- Service Chaining

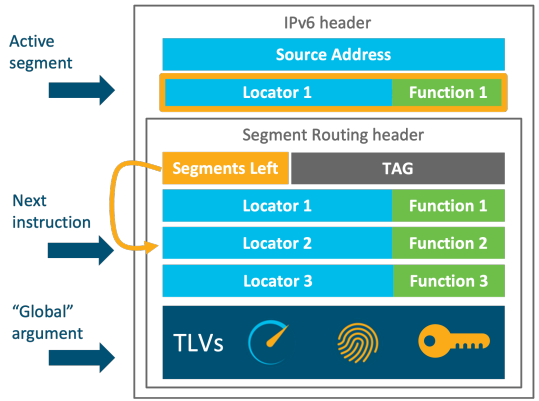

RFC 2460

RFC 8754
SR specific

SRv6 Micro-Program: Scale and MTU Efficiency

Applicability Examples

Optimized Transport

MTU Efficiency

5G Slicing

eMBB, mMTC, cMTC

Scalability

v6 Hyperscale

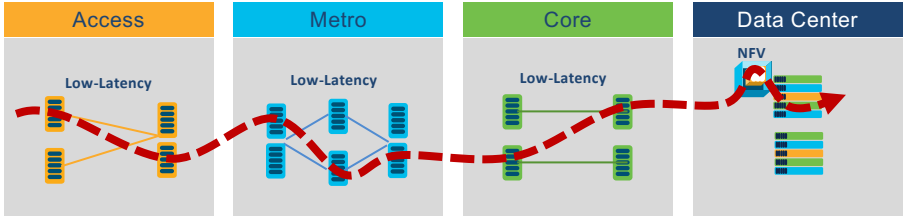

Solution

SRv6 Micro-Segments (uSID)

Ultra-Scale & MTU efficiency

New network instructions; **Shift & Forward**

SRv6, Net Prog and IPv6

uSID can be:

Node (uN), Adjacency (uA), Service (uDT4)

Up to 6 uSIDs in the outer DA + n x6 in SRH

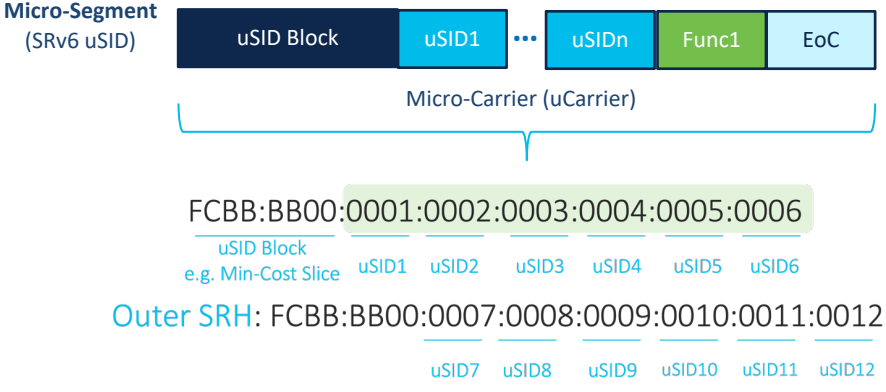

Scalable number of globally unique nodes in the domain

16-bit uSID: 65k uSIDs per domain block

32-bit uSID: 4.3M uSIDs per domain block

Benefits

Dataplane

- Best MTU efficiency (up to 6 uSIDs without SRH)
- Hyper-Scalable SR-TE (18 uSIDs with 40 bytes overhead)
- IP summarization and longest match is POWERFUL

Control Plane

- Scalable number of globally unique uSIDs per domain
- No new protocol extensions

Protect with automatic TI LFA FRR

Network Availability

Problem

Incomplete coverage, topology **dependent** coverage of classical LFA

Solution

Automated Topology **Independent** with guaranteed sub-50ms per-prefix protection

Benefits

Simple and Automated
IGP computed / No midpoint backup state

Optimal
Backup path following post-convergence path for 100% of the Topologies

Scalable
Cisco's TI-FLA algorithm – **optimized for scalability**
Post-convergence path computation and SID-list encoding

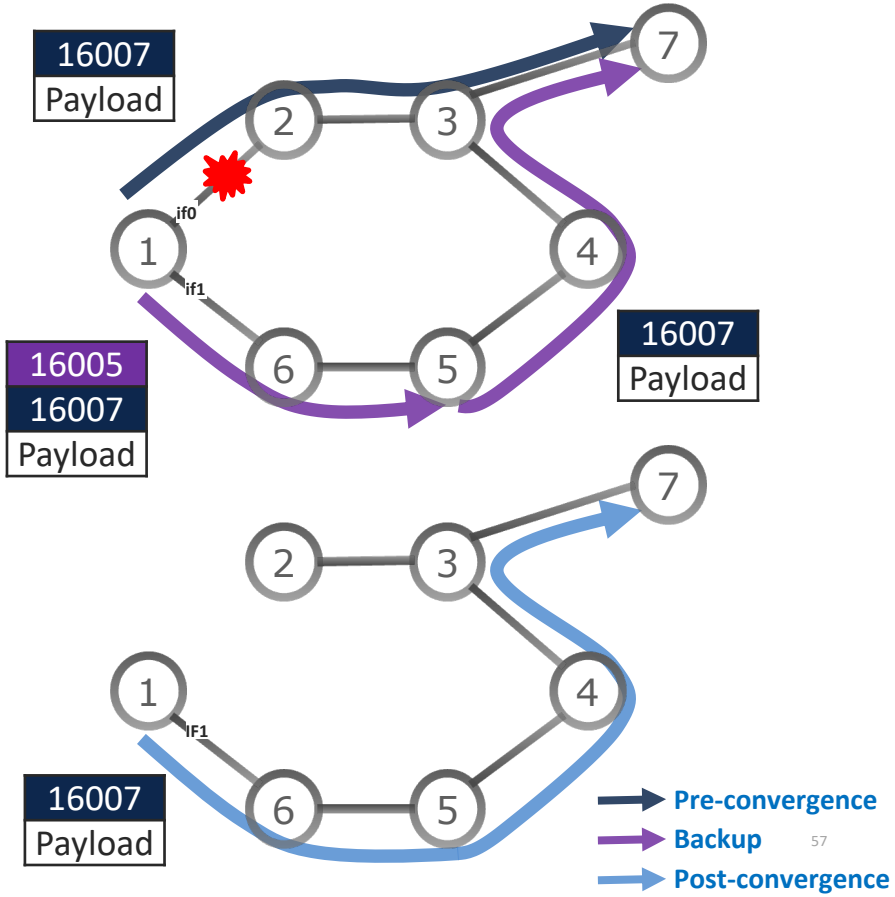

Latency Configuration

```
performance-measurement
 interface Gig0/0/0/0
  delay-measurement
   advertise-delay 7543
```

This will set latency of the link to 7543 **microseconds**

Performance Measurement

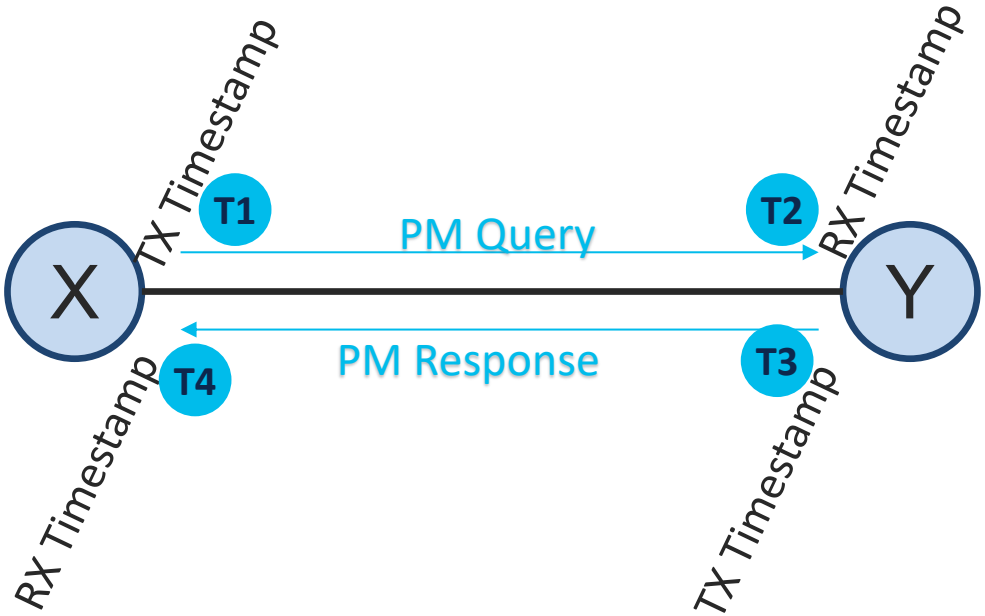

- TWAMP Ligth Protocol
- HW Level Timestamping
- ns precision!
- Link Latency Calculation:
  - One WAY Measurement

    $$Latency = T2 - T1$$
  - Two Way Measurement

    $$Latency = \frac{(T4 - T1) - (T3 - T2)}{2}$$

PM Configuration

```
performance-measurement
 interface Gig0/0/0/0
  delay-measurement
```

This will:

- Start PM probes on interface
- Provide Dynamic measurement values to IGP
- Both ends must be PM capable (provide HW based timestamping)

Service Programming

New Revenue Streams

Applicability Examples

Network Services
Firewall
Load-Balancer

Overlay
Multi-Could

As-a-Service
Network as a Service

Solution

Provide stateless NFV using Loc::Fun:Arg construct
Services are expressed with segments

SR aware Services
SR Proxies for legacy services

Integrated with Underlay intent/SLA and overlay

Benefits

Simplicity and automation at hyperscale
NSH creates per-chain states; SRv6 does not

Applications to control network behavior

Flexibility
New functions / Args can be easily added

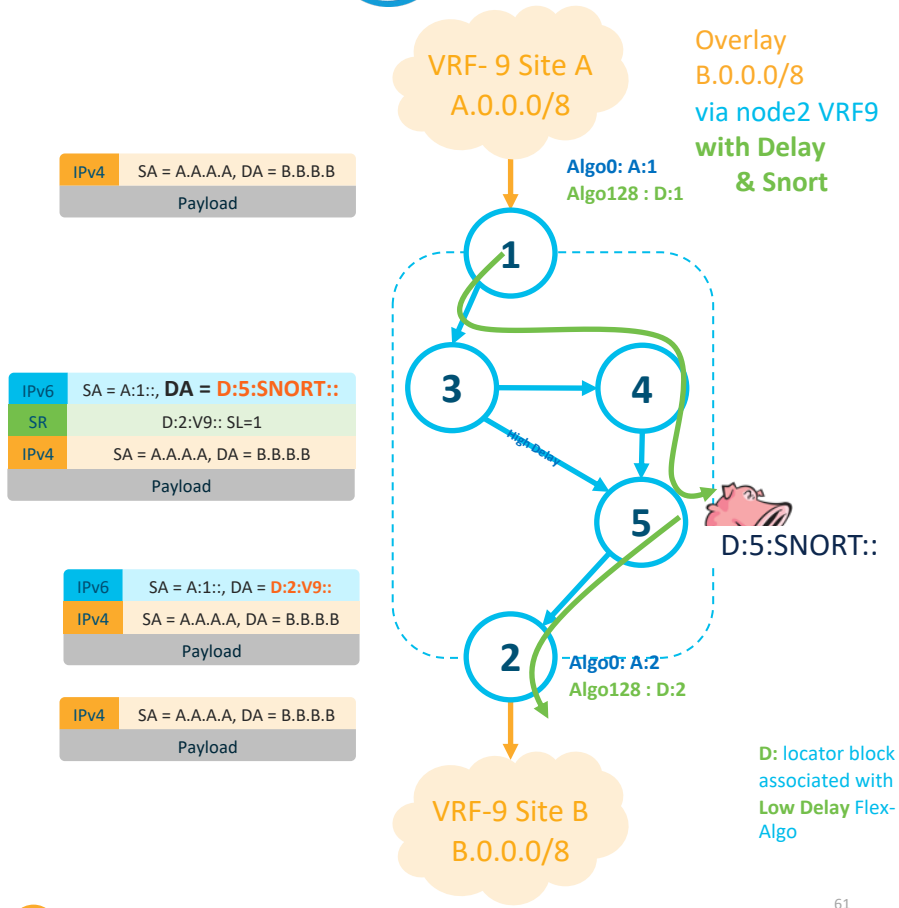

Interested? SRv6 Introduction on segment-routing.net


SRv6
Segment Routing
SR
Segment Routing
EVPN

Leadership in the Service Layer with EVPN

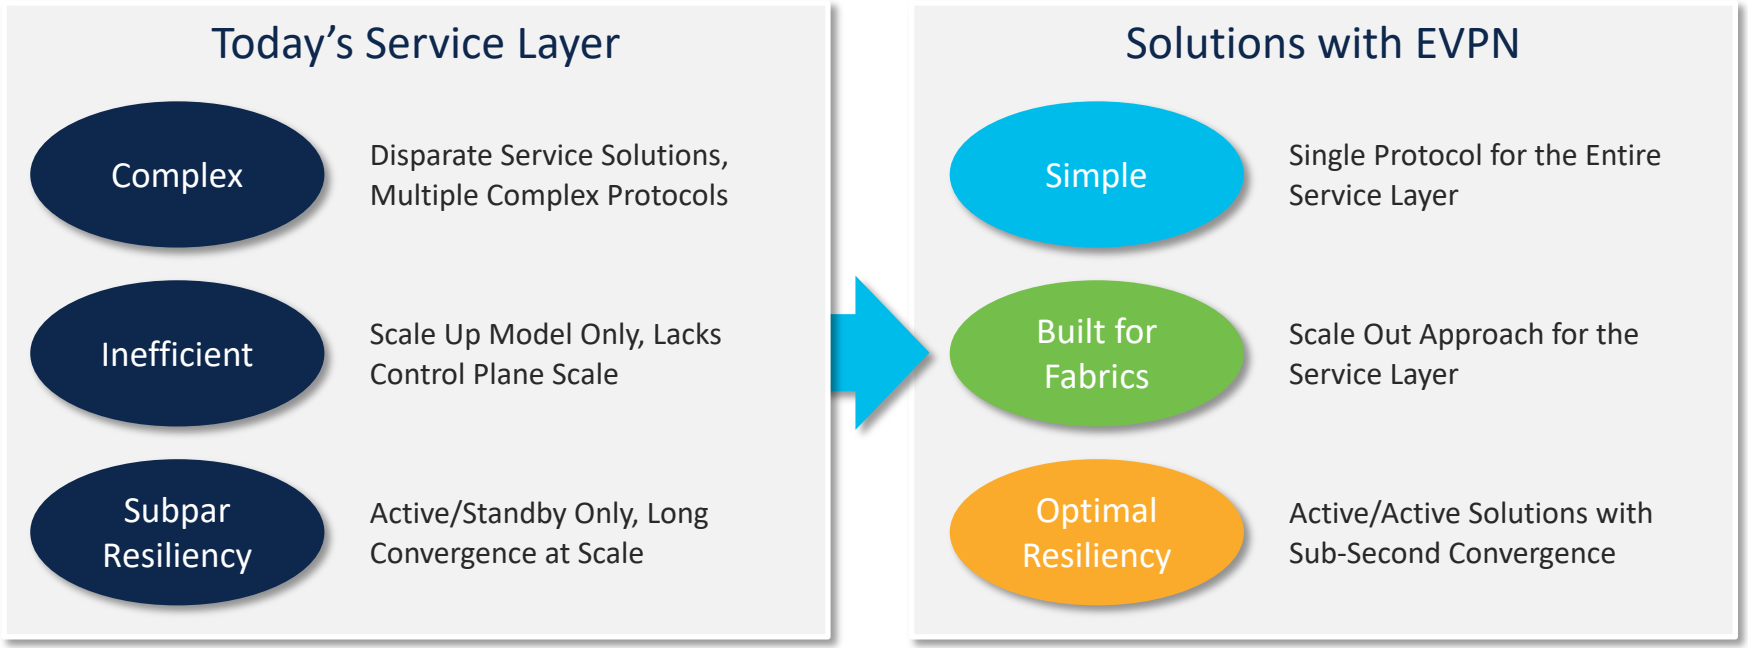

Unified services vision

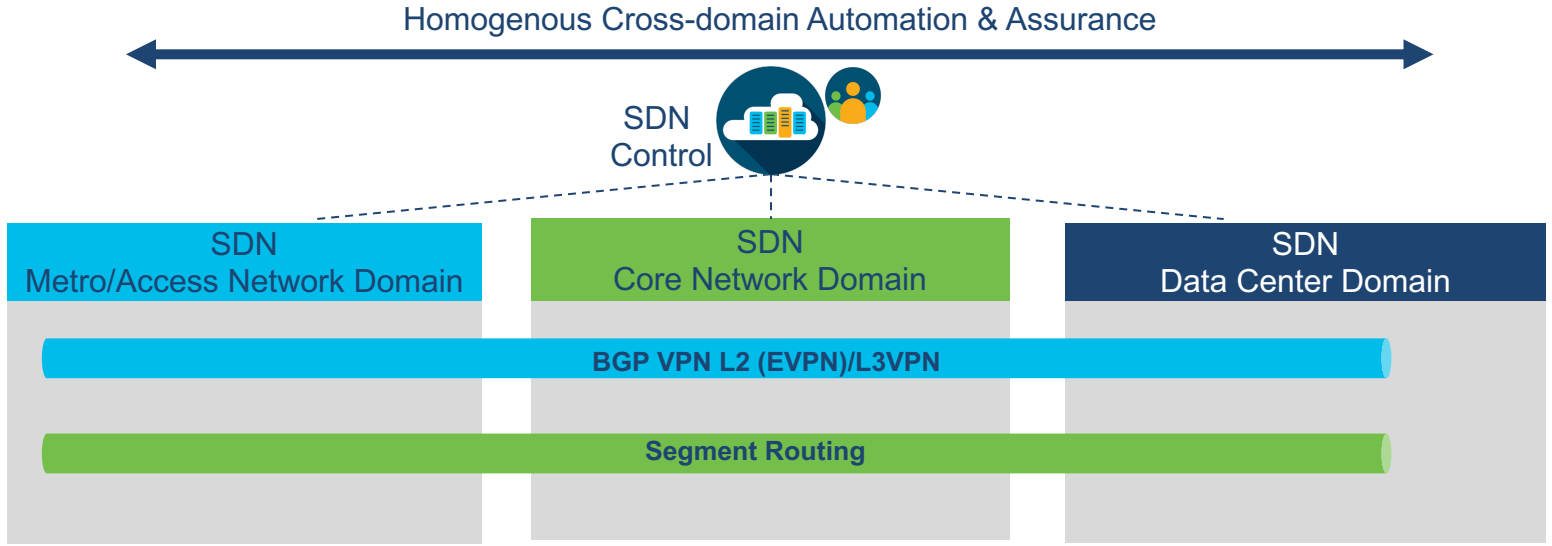

End-to-end service provisioning is simple and scalable

- Multiple network domains under same management teams
- Automated operations
- Homogenous underlay and overlay networks

Network Services Fabric: Value Proposition

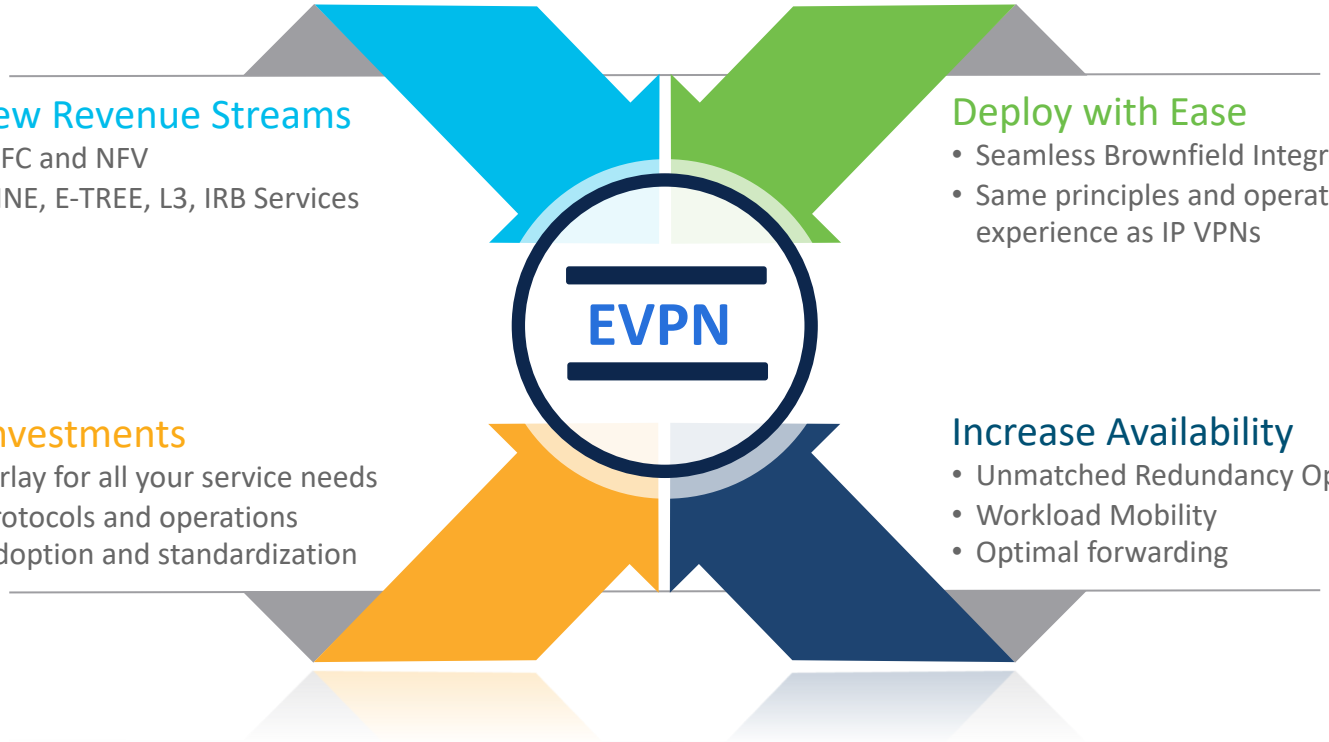

Create New Revenue Streams

- Stateless SFC and NFV
- E-LAN, E-LINE, E-TREE, L3, IRB Services

Deploy with Ease

- Seamless Brownfield Integration
- Same principles and operational experience as IP VPNs

Protect Investments

- Single Overlay for all your service needs
- Simplify protocols and operations
- Industry adoption and standardization

Increase Availability

- Unmatched Redundancy Options
- Workload Mobility
- Optimal forwarding

Fast, Resilient, Flexible Unified Services

EVPN Overview

Next-Generation Solutions for L2VPN
Solving VPLS challenges for per-flow Redundancy

- Existing VPLS solutions do not offer an All-Active per-flow redundancy
- Looping of Traffic Flooded from PE
- Duplicate Frames from Floods from the Core
- MAC Flip-Flopping over Pseudowire
  - E.g. Port-Channel Load-Balancing does not produce a consistent hash-value for a frame with the same source MAC (e.g. non MAC based Hash-Schemes)

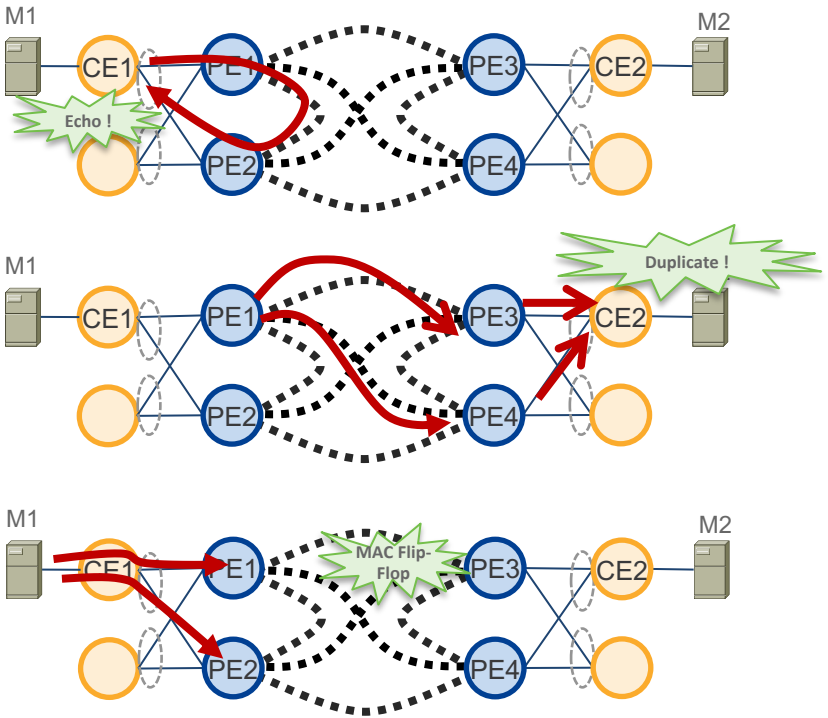

EVPN
Next generation network services

No technical benefit to replace them with EVPN L3!!

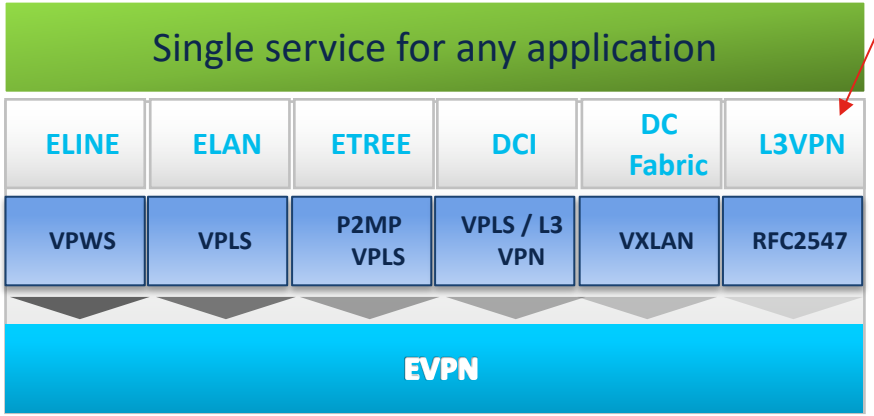

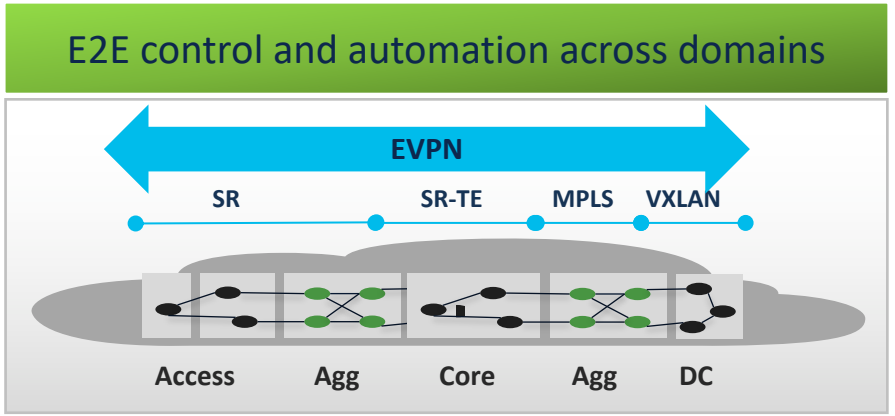

Optimized CapEx:
- Open Standards & Multi-vendor
- Active-Active multi-homing
- Enhanced load balancing

Reduced OpEx:
- Integrated L2 & L3 service, any application: faster time to market, certification
- E2E control and automation

Increased Customer Value
- Inter-domain SLA, faster convergence
- Better stability: no flood
- Granular policy control

EVPN - load-balancing modes

All-Active (per flow)

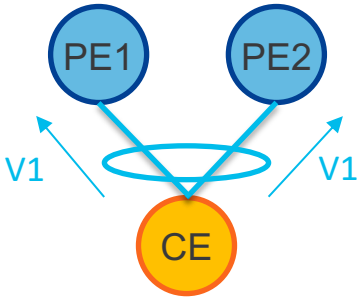

Single LAG at the CE
VLAN goes to both PE
Traffic hashed per flow
Benefits: Bandwidth, Convergence

Single-Active (per VLAN)

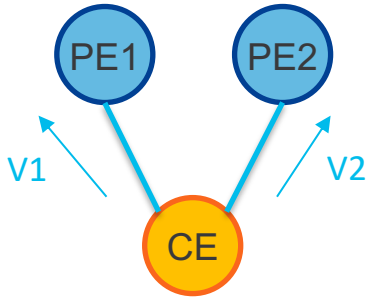

Multiple LAGs at the CE
VLAN active on single PE
Traffic hashed per VLAN
Benefits: Billing, Policing

Port-Active (per port)

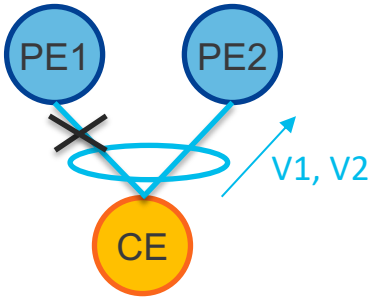

Single/Multiple LAGs at the CE
Port active on single PE
Traffic hashed per port
Benefits: Protocol Simplification

Single lambda optics

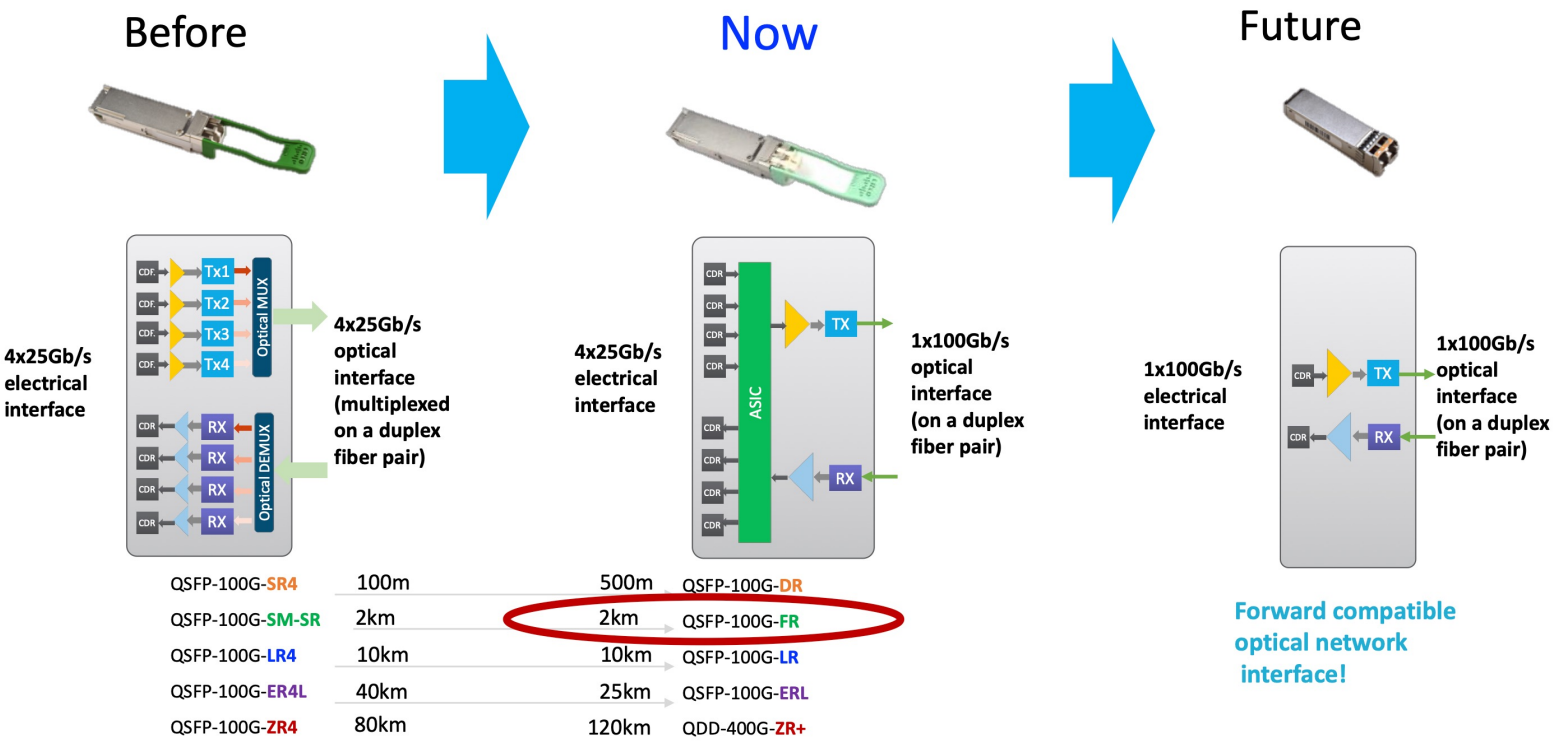

More about EVPN – the ultimate guide from Jiri Chaloupka

Service Overlay Cookbook - BRKSPG-2041

Jiri Chaloupka, Principal Technical Marketing Engineer, Cisco Systems, Inc. - **Distinguished Speaker**

Schedule **Monday, Jun 5 | 8:00 AM - 9:00 AM PDT | Level 3, South Seas C**

EVPN Deep Dive with IOS-XR Configuration examples for Service Provider Metro and Data Center - BRKMPL-2253

Jiri Chaloupka, Principal Technical Marketing Engineer, Cisco Systems, Inc. - **Distinguished Speaker**

Schedule **Wednesday, Jun 7 | 3:00 PM - 4:30 PM PDT | Level 2, Surf EF**

IOS XR EVPN Hands-On LAB - LTRSPG-2005

Jiri Chaloupka, Principal Technical Marketing Engineer, Cisco Systems, Inc. - **Distinguished Speaker**

David Jakl, Technical Marketing Engineer, Cisco Systems, Inc. - **Distinguished Speaker**

Schedule **Monday, Jun 5 | 1:00 PM - 5:00 PM PDT | Luxor - Level 1, Galleria DE**

Multicast with EVPN, Segment Routing & Traffic Engineering - BRKMPL-2123

Mankamana Mishra, Technical leader , Cisco Systems, Inc.

Schedule **Tuesday, Jun 6 | 10:30 AM - 12:00 PM PDT | Lower Level, Mariners AB**

Configure and Implement BGP-EVPN with Segment Routing using NCS 55xx/5xx platforms - LABSPG-3000

Tejas Lad, Technical Marketing Engineer, Technical Leader, Cisco Systems, Inc.

Paban Sarma, Technical Marketing Engineer, Cisco Systems, Inc.

The bridge to possible