Hardware Update

Arista EMEA Cloud Builders Technical Forum – Winter 2019

Confidential. Copyright © Arista 2019. All rights reserved.

ARISTA

Agenda

- Hardware Strategy
- 2019 Recap
- R3 series in details
- Optics:
  - Update
  - 400G optics
  - Breakout options for 400G
  - Interop options for 400G breakouts

Confidential. Copyright © Arista 2019. All rights reserved.

Hardware Update

Hardware Strategy

Confidential. Copyright © Arista 2019. All rights reserved.

Arista’s Hardware Strategy

- Use the best merchant silicon available
- Design for the datacenter / SP / Campus
- Build for scalable, high-performance networks
- Innovate to solve real world problems

Confidential. Copyright © Arista 2019. All rights reserved.

Rich Platform Portfolio

Single Arista EOS Image Across All Platforms

Confidential. Copyright © Arista 2019. All rights reserved.
ARISTA

Arista Silicon Landscape: 2016-2020

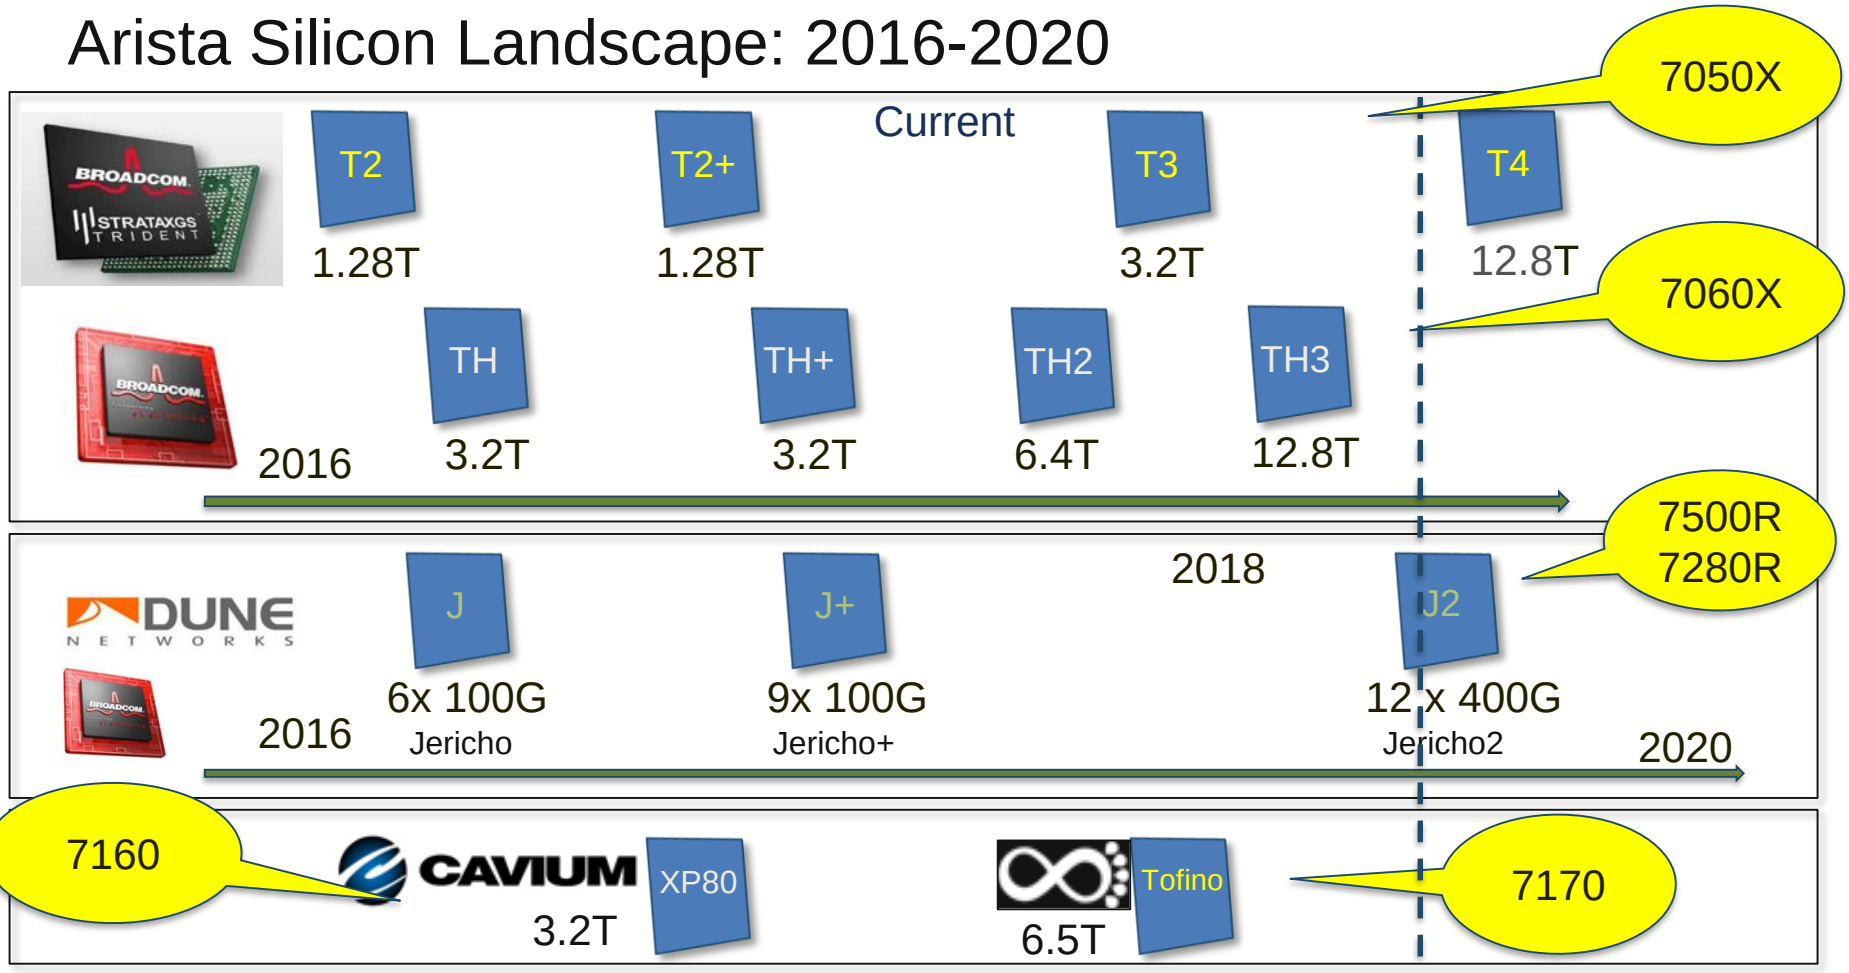

Confidential. Copyright © Arista 2019. All rights reserved.
ARISTA

Hardware Update

2019 Recap

Confidential. Copyright © Arista 2019. All rights reserved.

ARISTA

New Fixed DC Products

2019

Universal DC Leaf/Spine/Border

7050SX3-48YC8

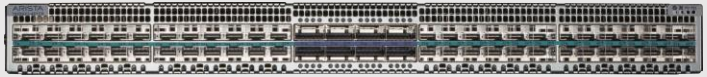

48x25G, 8x100G Ports

7050CX3M-32S

32x100G MACSEC Ports

Hyperscaler Leaf/Spine/Border

7060PX4-32

32x400G Ports (OSFP)

7060DX4-32

32x400G Ports (QSFP-DD)

Programmable L2/L3 Switch

7170-32CD

Double Pipeline, 32x100G Ports

Small deep-buffered Leaf & TapAgg

7020SR-32C2

32x10G, 2x100G Ports

7020SR/SRG-24C2

24x10G, 2x100G Ports
HW-based IPSec (G model)

Confidential. Copyright © Arista 2019. All rights reserved.

ARISTA

New Modular DC Products

7368X4

16xQSFP100

4xQSFP-DD

4xOSFP

High Network Radix Modular System

- Choice of port module configurations
- Improved power efficiency per bandwidth
- Upgradeable to next generation
- 128x 100G QSFP or 32x 400G in 4RU

2019

High Density DC Spine or DC/Campus Spline

- Two Chassis - 4 and 8 slot options
- Two Linecard options:
  - 32x100G
  - and new one

7300X3-48YC4 – 48x25G, 4x100G

32MB Buffer Linecard

Confidential. Copyright © Arista 2019. All rights reserved.

ARISTA

New 720XP Series: Campus Leaf with PoE

2019

Ready for Next Generation Campus

- Ready for Multi-Gig Ethernet transition: mGig down; up to 100G uplinks
- Endpoint behavior analytics with hardware-based flow telemetry
- Secure Segmentation with VXLAN overlay
- Compact, power efficient form factor
- Common “X-series” architecture

Simplicity through Consistent Software with EOS and CloudVision

Confidential. Copyright © Arista 2019. All rights reserved.

ARISTA

Hardware Update

R3 Series

Confidential. Copyright © Arista 2019. All rights reserved.

Proven Architecture - VoQ and Cell Fabric

- Packet buffers on Ingress to prevent incast
  - Large packet buffer for burst absorption
  - Egress buffer never fills up
- Credit based virtual queues on ingress
  - Egress ports pull traffic when ready
- Cell based fabric eliminates speed mismatch
  - Easy migration to next speed eg. 100G → 400G
  - Allows mix and match with previous generations

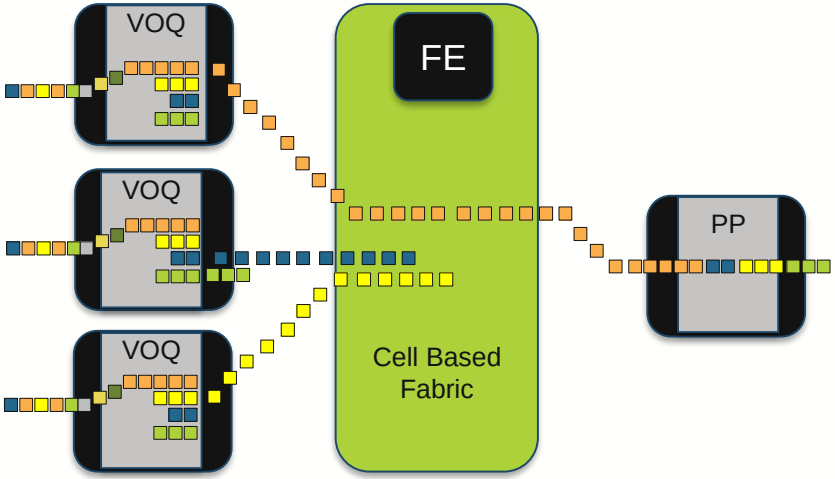

Best in Class - Distributed Buffer for Incast Prevention

Confidential. Copyright © Arista 2019. All rights reserved.

ARISTA

Proven Architecture - Distributed Buffer

- Large HBM2 and on On-Chip packet buffer for burst absorption
  - Integrated 8GB of HBM2 packet memory eliminates packet drops
  - 43% less power per bit than comparable GDDR5
  - Increases system density and allows higher port counts
- Flexible behavior when On Chip Buffer and HBM2 combine
  - Low latency, OCB-based under normal conditions
  - Seamless expansion to HBM2 memory for congested flows

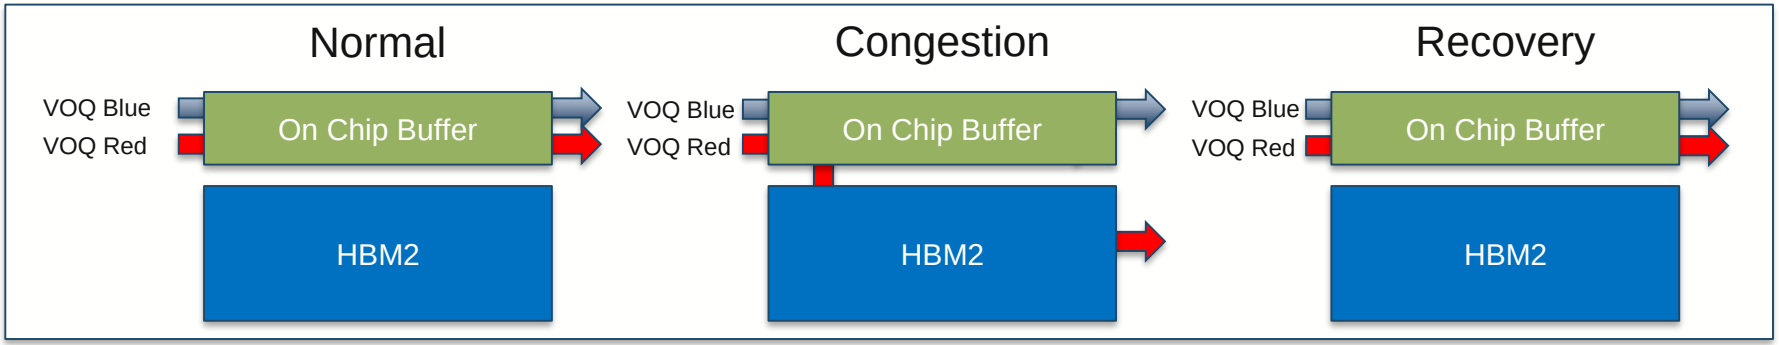

Optimized for High Performance Next Generation Routing at all Network Tiers

Confidential. Copyright © Arista 2019. All rights reserved.

ARISTA

Jericho2 – move to 50Gbps SerDes

		FCS	Process (nm)	Clock Rate (MHz)	Packet (Mpps)	Nif BW (Gbps)	SerDes (Gbps)
Petra	7048/7500	2010	65	250	125	80	6.25
Arad	7280/7500E	2013	40	600	300	240-280	12.5
Jericho	7280/7500R/R2	2016	28	720	720	600-1080	25.7
Jericho2	7280/7500R3	2019	16	1000	2000	4800	53.125

Increase over Jericho…

1.4x	2.7x	4.4x	2.1x

Confidential. Copyright © Arista 2019. All rights reserved.

Arista R3 Innovations and Continuity

Proven Architecture

Consistent VoQ non blocking fair delivery
Deep packet buffers with HBM2 with lower power
Power efficient with support from 10G to 400G
High bandwidth and high density

Investment Protection and EOS

Arista EOS Innovations and 10 Years of Experience
Investment Protection for current and next generation
Single EOS across portfolio
Feature rich and high quality on day-one

Programmable

Programmable packet lookups and forwarding
Flexible profiles for range of network use-cases
Add protocols and extensions with simple updates

Advanced Network Telemetry

Accelerated sFlow at 400G for traffic analysis and engineering
In-band Network Telemetry for flow and packet forensics
State Streaming, BGP FlowSpec and IPFIX support

Routing Scale and FlexRoute™

Internet Route Scale leveraging FlexRoute
Enhanced Segment Routing and EVPN
Rich Tunneling for v4/v6 overlays
DC Spine, DCI, MPLS, Peering and Edge networks

Security at Scale

New Algorithmic ACL for 4X capacity increase
Stateful ACLs for secure Edge networking
Rich counters to classify and filter in real-time
Network Address Translation

Confidential. Copyright © Arista 2019. All rights reserved.

Next Generation R-Series Portfolio

7280R3

7500R3

7800R3

115Tbps – 460Tbps

76Tbps – 230Tbps

4Tbps – 38Tbps

Confidential. Copyright © Arista 2019. All rights reserved.

7500R3 High Density 400G and 100G Spine Systems

High Performance 100G / 400G Spine:

- 230Tbps of throughput with choice of Chassis
- Consistent VOQ / Deep Buffers
- Backward compatible with 7500R and 7500R2

Chassis	400G OSFP	4 x 100G	100G QSFP
DCS-7512	288	1152	432
DCS-7508	192	768	288
DCS-7504	96	384	144

400G Spine:

- 24 x 400G OSFP linecards
- Supports range of optics and cables to ZR and ZR+
- Breakout to 4x100G and 2x200G

100G Spine:

- 36 ports of 100G with QSFP
- Supports copper cables, AOC, data center to DWDM optics
- 2.5M Route Scale Option

Confidential. Copyright © Arista 2019. All rights reserved.

7800R3 Series Next Generation 100G/400G

Cloud and Carrier Grade Networking

High Performance for next 10 years:

- Up to 576 x 400G wire speed ports
- Non-blocking up to 460 Tbps and 96Bpps
- 14.4 Tbps / slot with 36 x 400G linecards
- Upgradable to 800G (28.8Tbps /slot) for higher density

R-Series Architecture:

- VOQ architecture and deep buffers for lossless forwarding
- FlexRoute™ - 1.3M / 2.5 Million+ IPv4 & IPv6 Routes
- EOS for convergence and scale

Advanced Features:

- VXLAN Routing, Advanced Load Balancing
- Algorithmic ACLs, INT and Accelerated sFlow
- EVPN, MPLS, Segment Routing
- Dense 100G and 400G for SP, Cloud, Internet, HPC & CDN

Confidential. Copyright © Arista 2019. All rights reserved.

ARISTA

Highest Capacity 400G and 100G Spine System

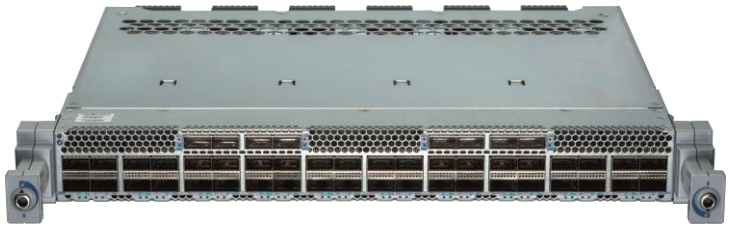

100G Spine:

- 48 ports of 100G with QSFP
- Supports copper cables, data center to DWDM optics
- 2.5 Million Routes with features

400G Spine:

- 36 ports of 400G OSFP – 14.4Tbps
- 6 Billion Packets per second of L2 and L3
- Range of optics and cables - ZR and ZR+
- Flexible 4x100G and 2 x 200G Modes

High Performance Spine:

- Choice of Chassis (4/8/16 slot)
- Future higher density and 800G

Chassis	Throughput	400G OSFP	4 x 100G	100G QSFP
DCS-7816	460Tbps	576	2304	768
DCS-7808	230Tbps	288	1152	384
DCS-7804	115Tbps	144	576	192

Confidential. Copyright © Arista 2019. All rights reserved.

7280R3 Series Fixed 100G/400G Switches

Wire Speed 100/400G with Deep Buffers

High Performance:

- Up to 48 x 400G wire speed ports
- Non-blocking up to 19.2 Tbps and 8Bpps
- FlexRoute™ - 1.3M / 2.5 Million+ IPv4 & IPv6 Routes

R-Series Architecture:

- VOQ architecture and deep buffers for lossless forwarding
- EOS for convergence and scale

Advanced Features:

- VXLAN Routing, Advanced Load Balancing
- Algorithmic ACLs and Accelerated sFlow
- EVPN, MPLS, Segment Routing

Cloud and Carrier Grade Networking:

- Dense 100G and 400G for SP, Cloud, Internet, HPC & CDN
- DC Optimized airflow and AC / DC power

7280PR3-48 - 48 x 400G

7280PR3-24 - 24 x 400G

7280CR3-96 - 96 x 100G

7280CR3-32P4 - 32 x 100G / 4 x 400G

7280CR3M-36 - 36 x 100G MACsec

Confidential. Copyright © Arista 2019. All rights reserved.

ARISTA

Arista 7280CR3-32P4 and 7280CR3K-32P4 Architecture

High Performance

- 32 wire speed 100G and 4 400G ports
- Non-blocking 4.8 Tbps and 2Bpps
- FlexRoute™ - 1.3 / 2.5 Million+ IPv4 & IPv6 Routes

R-Series Architecture

- VOQ architecture for lossless forwarding
- 8GB Deep packet buffers
- EOS for convergence and scale

Advanced Features

- VXLAN Routing, Advanced Load Balancing
- Algorithmic ACLs and Accelerated sFlow
- EVPN, MPLS, Segment Routing

Cloud and Carrier Grade Networking

- Dense 100G for SP, Cloud, Internet Routing, HPC & CDN
- Less than 12W per 100G port
- Front-to-rear or Rear-to-front Airflow

32x QSFP 100G and 4 x OSFP 400G

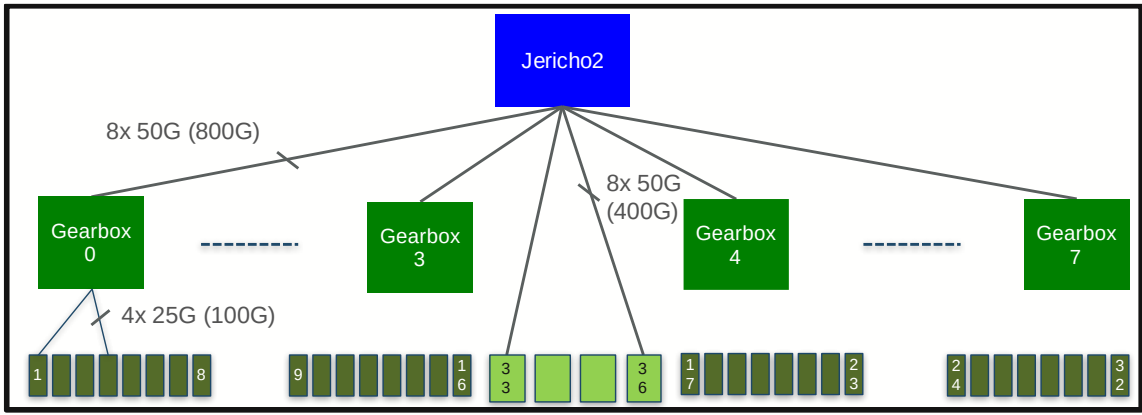

Confidential. Copyright © Arista 2019. All rights reserved.

Silicon Serdes Speeds and Optical Transceivers

Legacy 100G Optics

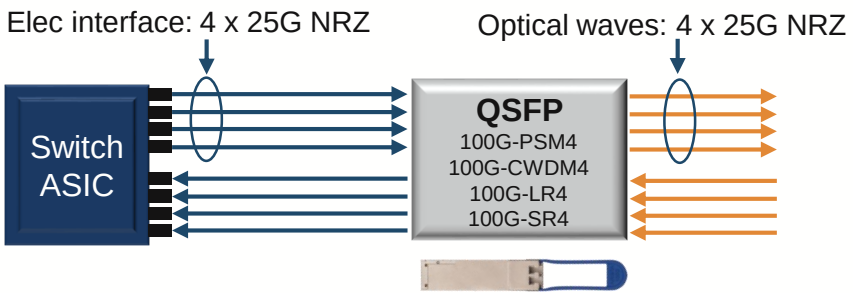

New 100G Optics (FR, DR)

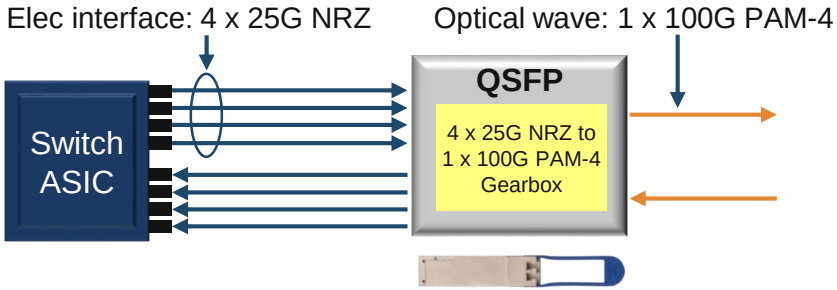

400G Optics (FR4, DR4, XDR4)

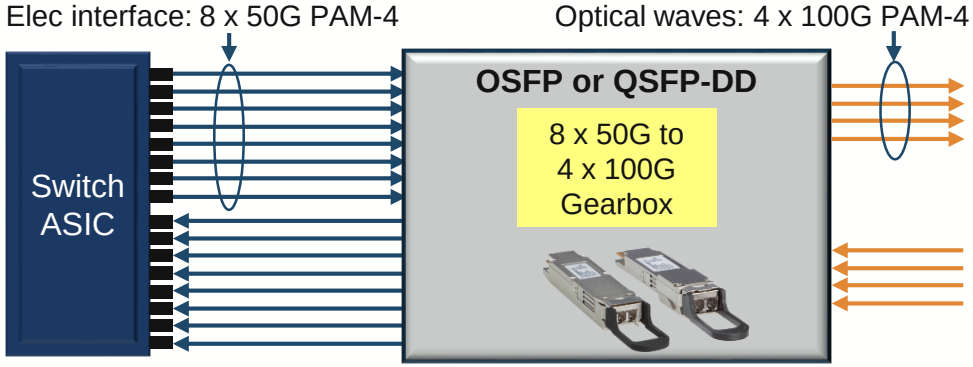

400G Optics (SR8, 2FR4)

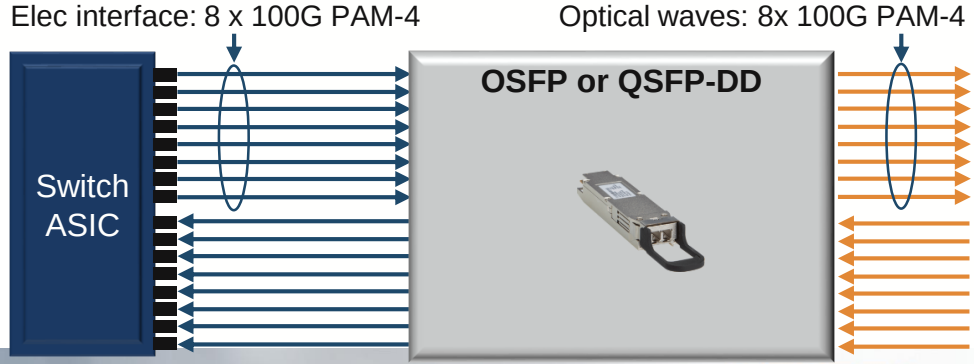

Confidential. Copyright © Arista 2019. All rights reserved.

Speed Groups / 25G serdes

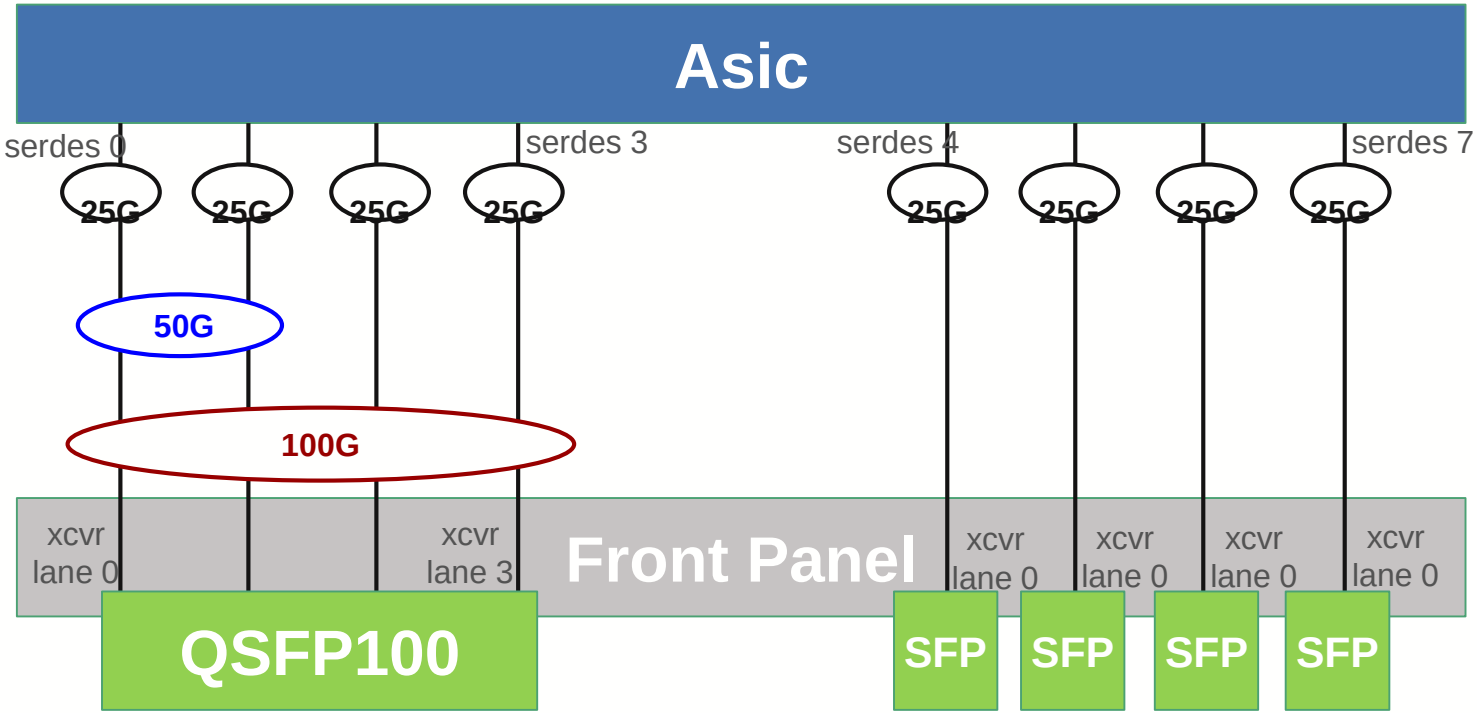

https://eos.arista.com/eos-4-22-1f/speed-group-management/

Confidential. Copyright © Arista 2019. All rights reserved.

50G Lane Challenge

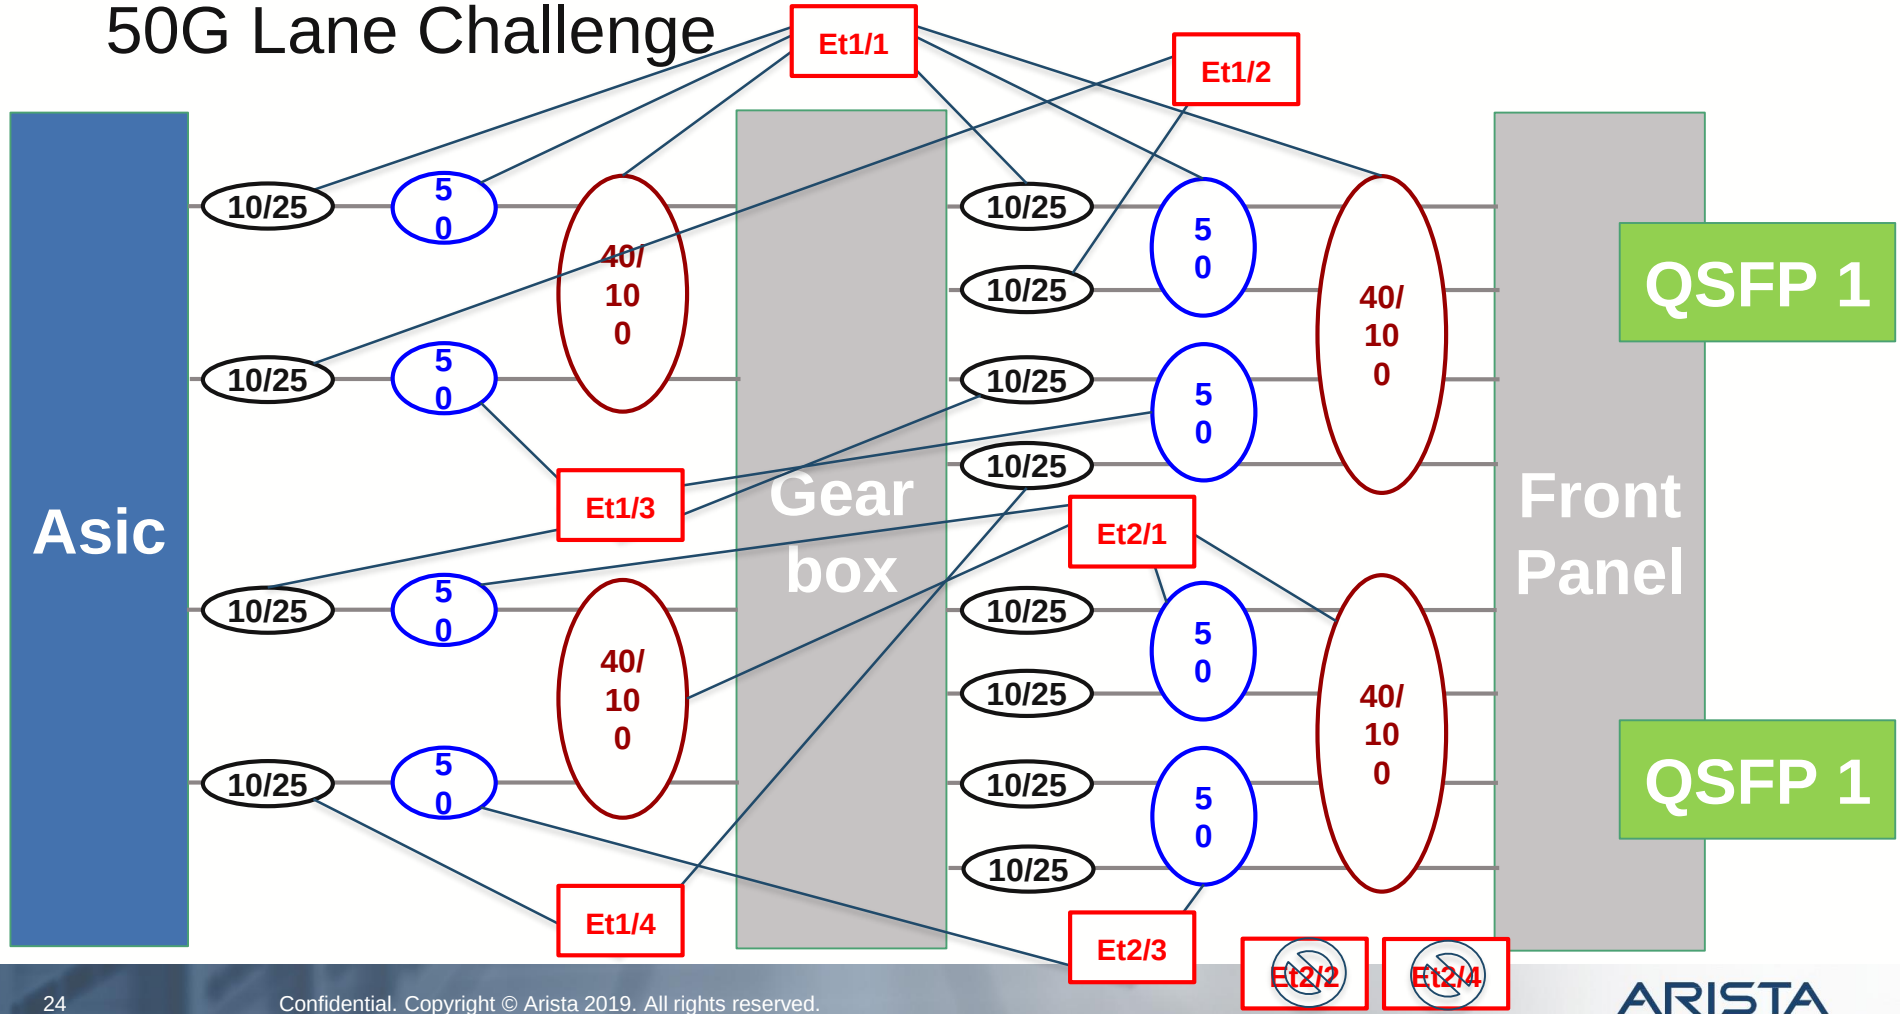

Confidential. Copyright © Arista 2019. All rights reserved.
ARISTA

Hardware Update

Optics / 400G / Breakouts

Confidential. Copyright © Arista 2019. All rights reserved.

Optical Transceiver Basics

Electrical interface	Form factor & data rate	Optical conn.	Fiber type	Wavelength	Max reach
1G & 10G	**SFP**: 1G & 10G	Dual LC	Duplex SMF	1310/1550nm	80km
			Duplex MMF	850nm	400m
4 x 10G (40G) or 4 x 25G (100G)	**QSFP**: 40G & 100G	Dual LC	Duplex SMF	4x ~1310/1550nm	40km
			Duplex MMF	4x or 2x ~850nm	40G: 150m, 100G: 100m
		MPO-12	Parallel SMF	1310nm	40G: 10km, 100G: 500m
			Parallel MMF	850nm	40G: 400m, 100G: 300m
8 X 50G	**OSFP**: 400G	Dual LC	Duplex SMF	4x ~1310nm	2km, 10km
		MPO-12	Parallel SMF	1310nm	2km
8 X 50G	**QSFP-DD**: 400G	MPO-16	Parallel MMF	850nm	100m
		Dual CS	2 x Duplex SMF	4x ~1310nm	2km

Confidential. Copyright © Arista 2019. All rights reserved.

OSFP and QSFP-DD 400G Optics

Arista have OSFP and QSFP-DD products

- Two connector standards: OSFP and QSFP-DD
- We prefer OSFP as it is technically superior
  - Higher power budget, easier to cool
  - More choices earlier in OSFP (not dependent on 7nm gearboxes)
  - More options for high-power optics (like ZR 120km 400g / ZR+ 1000km)
  - Supports 100G electrical which is the ***most*** cost-effective
- QSFP-DD is backwards-compatible with QSFP-100
- Solution: OSFP-to-QSFP Adapter for 100G compatibility
  - Inserts into an OSFP slot
  - Lets you deploy a 400G switch and run it at 100G!
  - Mechanical adaptor - purely passive

OSFP

QSFP-DD

OSFP-QSFP adapter

800G Optics (OSFP) - Future

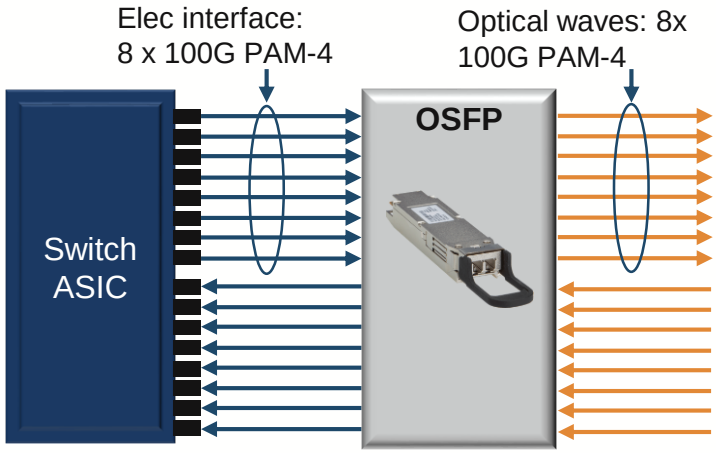

Confidential. Copyright © Arista 2019. All rights reserved.

ARISTA

400G Optical Transceiver & Cable Types

Application	Reach	Media Type	Solution
TOR to Server	3m	Copper	400G-CR8 (DACs)
TOR to Leaf	30m	AOC	400G-AOC
Switch to switch inside the datacenter	100m	Parallel MMF	400G-SR8
	100m	MMF	400G-BIDI
	500m / 2km	Parallel SMF	400G-DR4 / XDR4
	2km	SMF	400G-FR4
	2km	Dual SMF	400G-2FR4
DCI, Metro & Long-Haul	120km / 1000km+	SMF	400G-ZR / ZR+

Arista will support all available types

Confidential. Copyright © Arista 2019. All rights reserved.

400G Optical Modules for Single Mode Fiber

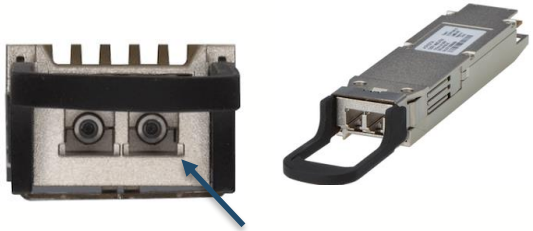

Duplex LC

400G-FR4, 2km reach

Same fiber plant as 100G-LR4 or 100G-CWDM4

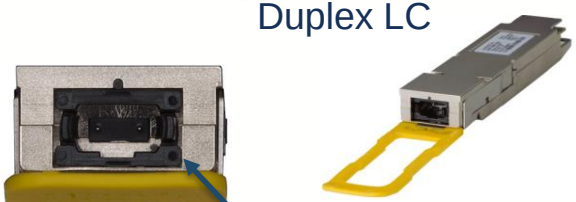

MPO-12

400G-DR4, 4 fiber pairs, 500m reach
400G-XDR4, 4 fiber pairs, 2km reach

Same fiber plant as 100G-PSM4
Breakout to 4 x 100G-DR or 100G-FR QSFP

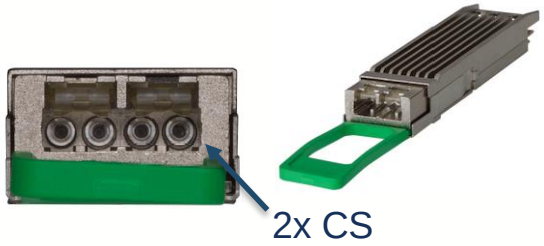

2x CS

400G-2FR4 (2 x 200G-FR4), 2km reach

Breakouts to 2 x 200G-FR4 QSFP, or
2 x 100G-CWDM4 QSFP (half rate)

Confidential. Copyright © Arista 2019. All rights reserved.

400G Optical Modules for Multi-Mode Fiber

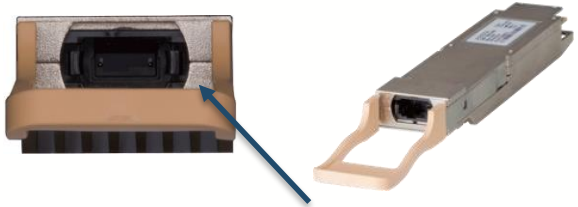

400G-SR8, 8 fiber pairs, 100m reach

Breakout to 2 x 200G-SR4 or 2 x 100G-SR4 QSFPs

MPO-16 APC (Angled MMF) - Requires new MPO-16 fiber connectivity

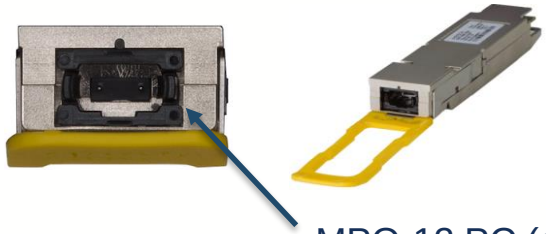

400G-BIDI, 4 fiber pairs, 100m reach

Breakout to 4 x 100G-BIDI QSFP

MPO-12 PC (std. unangled MMF)

OSFP and QSFP-DD versions

Lengths from 1m to 30m

AOC-O-O-400G-xM and AOC-D-D-400G-xM

Confidential. Copyright © Arista 2019. All rights reserved.

400G Platforms Offer Cost Effective and High Density 100G

Upgrade Spine to 400G Systems

32 x 100G

32 x 400G = 128 x 100G

Increase Network Radix 4x by Using 400G Spine Switches

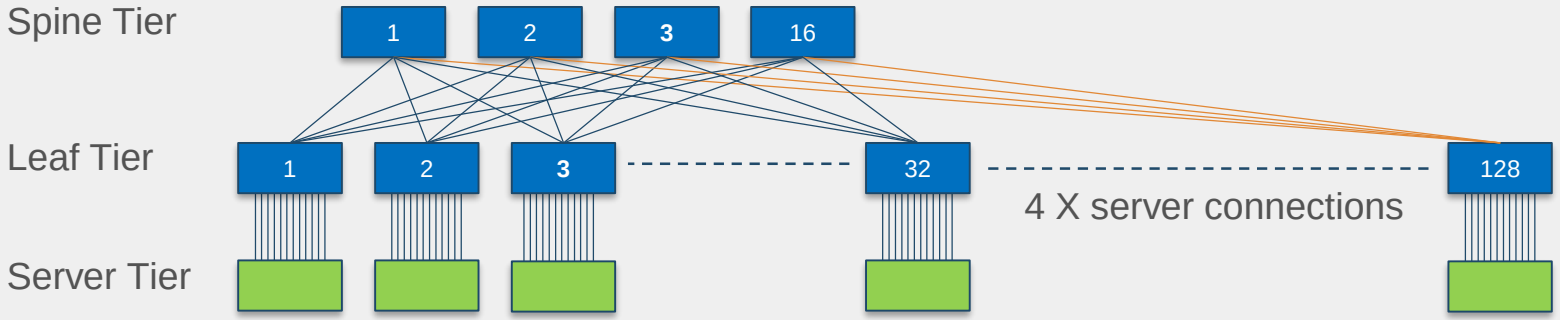

2048x100G / 8192x25G DC with no oversubscription

Confidential. Copyright © Arista 2019. All rights reserved.
ARISTA

400G Single Mode Breakout Options

400G Systems

400G-DR4 / XDR4

Parallel SMF (4 pairs)

500m / 2km

4 x QSFP-100G-DR / FR

100G QSFP Systems

100G-DR / FR QSFPs supported on ALL Arista 100G QSFP platforms

Gearbox: 2x50f ←→ 4x25g lane conversation

400G Systems

400G-2FR4
(2x200G or 2x100G mode)

duplex SMF

duplex SMF

2km

2 x 200G-FR4 QSFPs or
2 x 100G-CWDM4 QSFP

200G or 100G QSFP Systems

Confidential. Copyright © Arista 2019. All rights reserved.

400G Multi-Mode Breakout Options

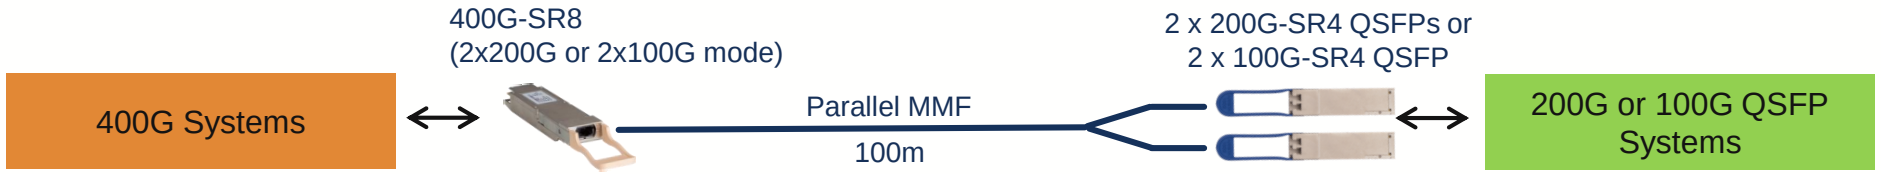

Confidential. Copyright © Arista 2019. All rights reserved.

400G Copper Cables for short links and breakout

400G to 400G cables (OSFP and QSFP-DD)

CAB-O-O-400G-xM and CAB-D-D-400G-xM

Breakout to 2 x 200G QSFPs, or 2 x 100G QSFP

CAB-O-2Q-400G-xM and CAB-D-2Q-400G-xM

CAB-O-2Q-200G-xM and CAB-D-2Q-200G-xM

Breakout to 4 x 100G-2 QSFPs, or 4 x 50G-2 QSFPs

CAB-O-4Q-400G-xM and CAB-D-4Q-400G-xM

CAB-O-4Q-200G-xM and CAB-D-4Q-200G-xM

Breakout to 8 x 50G SFPs (future), or 8 x 25G SFPs

CAB-O-8S-200G-xM and CAB-D-8S-200G-xM

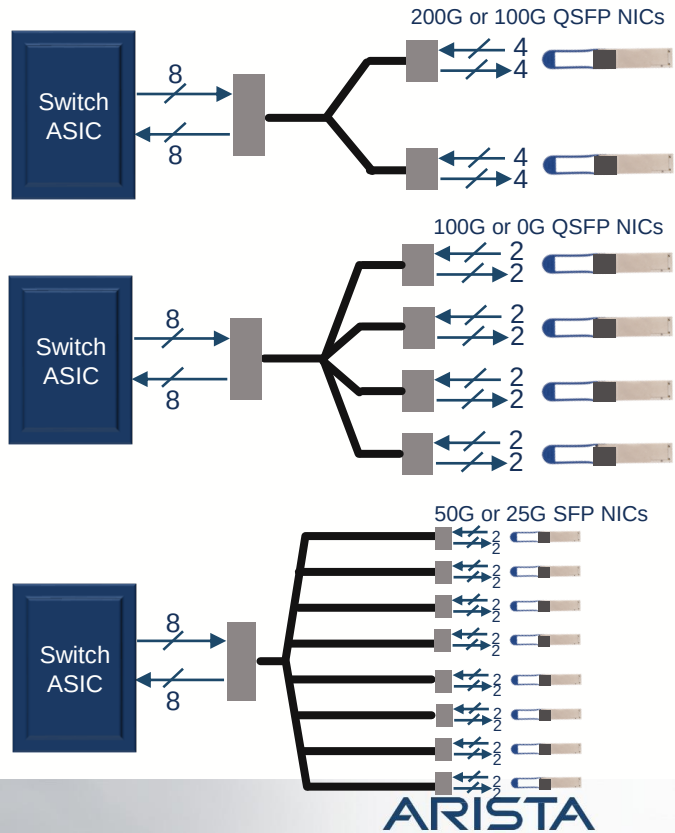

Confidential. Copyright © Arista 2019. All rights reserved.

Confidential. Copyright © Arista 2019. All rights reserved.