

# ARISTA

## APAC Cloud Builders 2023

Friday, 29 September  
Hanoi



# ARISTA

APAC Cloud Builders 2023

Data Center Continuous Innovation

Hung Le



# Agenda


- ❑ Arista At A Glance
- ❑ EOS – Modern Networks Operating System
- ❑ Universal Cloud Networking Architectures
- ❑ Platform Overview

# ARISTA

## Arista At A Glance



# Arista At A Glance NYSE: ANET Market Cap \$60B



**IPO 2014**  
June 6th

**S&P 500**  
Added in 2018

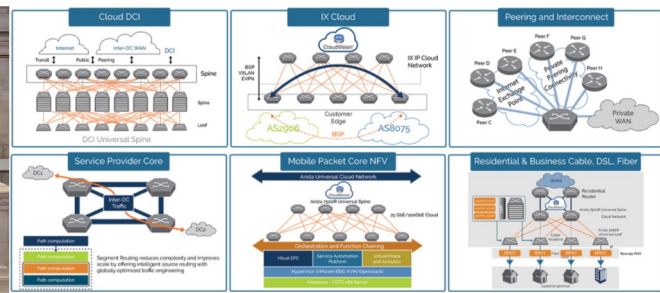
**Leader in 100/400G**

**9000+**  
Customers

**Gartner Leader**  
2020 Magic Quadrant for Data Center and Cloud Networking

**Forrester Wave Leader**  
Open Programmable Switches for Business Wide SDN

# Arista Networks



## 2004

Ken Duda and Andy Bechtolsheim (cofounder of Sun Microsystems) who had been with Granite Systems was acquired by Cisco. This would later become the Catalyst product. Went on to create **Arastra** (later renamed **Arista**)

With the vision of creating a Network OS which was built from the ground up to be highly scalable for the cloud. Customer immediately took notice.

## 2008

Jayshree Ullal who had been with Cisco for 15 years join Arista as the CEO. Before leaving Cisco she was the VP of Data Center & Switching which oversaw the Catalyst & Nexus product lines.

By this time Arista had customers from Financial Institutions, Hyperscalers, Enterprise and Service providers

## 2014

With all cylinders firing and customers such as Microsoft, Google, Facebook, Netflix, Twitter, eBay and many more Arista went public.

Arista continued to expand their teams but stayed true to a single OS across all platforms. As more features were added the lines between switches and routers blurred as every switch had full features

## 2020 & Beyond

While continuing to service enterprises, Arista continues to innovate and make strategic acquisitions to both expand the code base that make it so successful today - while adding strategic tools that are able to leverage the foundational network operating system of EOS.

2018 - Mojo Networks (Wireless)

2018 - MetaMako (Ultra Low Latency FPGA)

2020 - BigSwitch (Monitoring)

2020 - AWAKE (NDR)

2022 - Untangle

2022 - Pluribus Networks

# Arista by the numbers

**+80\***

Arista NPS  
2022

**#1**

100G  
400G

**6**

Consecutive Years  
Leader  
Gartner MQ

**1**

Operating  
System


**9000+**

Customers

## \* Net Promoter Score Definition

The Net Promoter Score is an index ranging from -100 to 100 that measures the **willingness of customers to recommend a company's products or services** to others. It is used as a proxy for gauging the customer's *overall satisfaction with a company's product or service* and the customer's loyalty to the brand.

# Data Center Ethernet Switch Analysis and Forecast

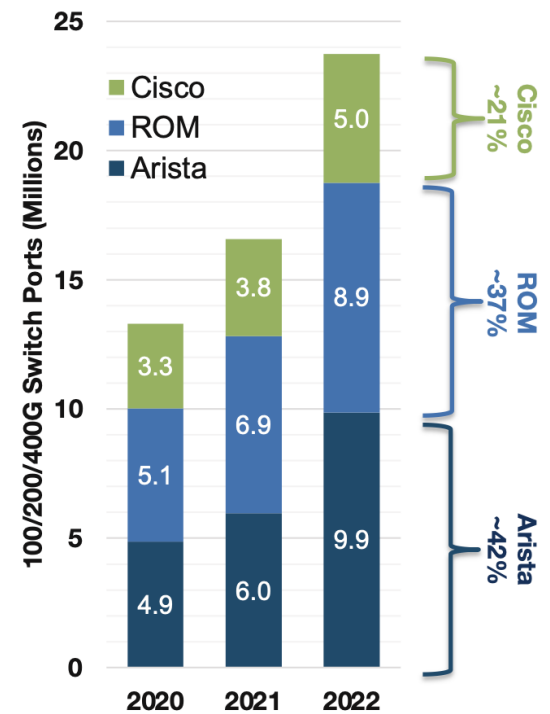
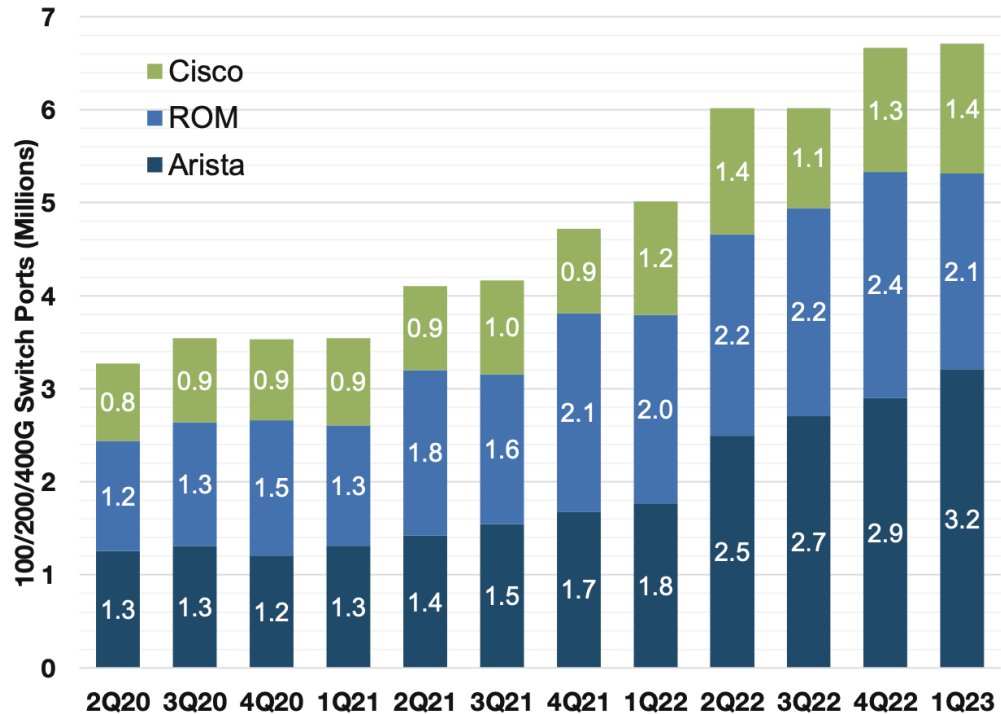


Source: Dell'Oro January 2023 - Long Term Ethernet Switch Forecast

Note - Initial 800 Gbps shipments will not be using 800 G Ethernet MAC and will be configured mostly as 2x400 Gbps or as 8x100 Gbps



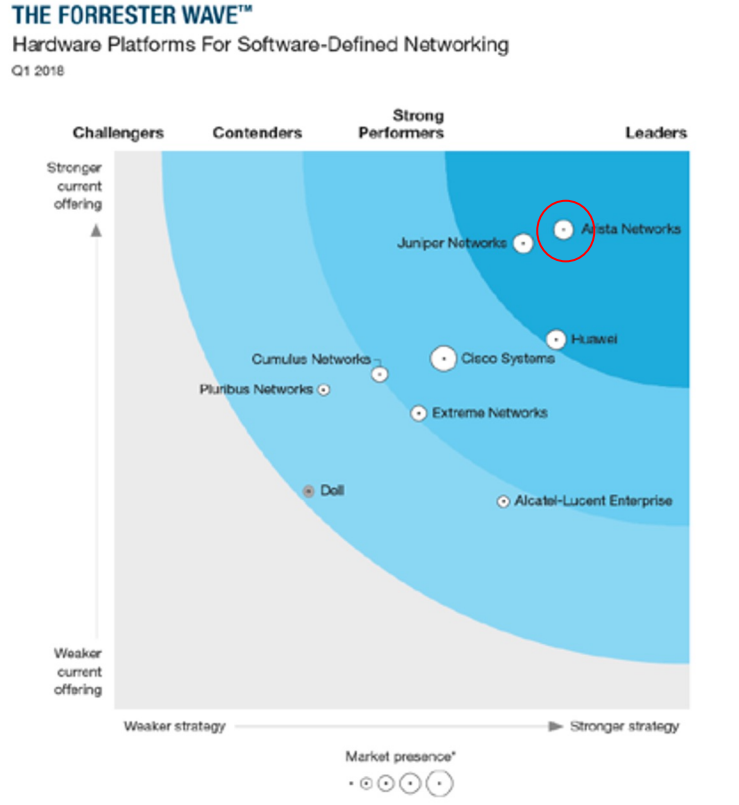
# Arista's Market Leadership in 100G/200G/400G Data Center High Speed Ethernet Port Analysis



Source: Crehan Ethernet Switch Data Center Total Vendor Tables – 1Q'23

# Arista Recognition

INDUSTRY RECOGNITION



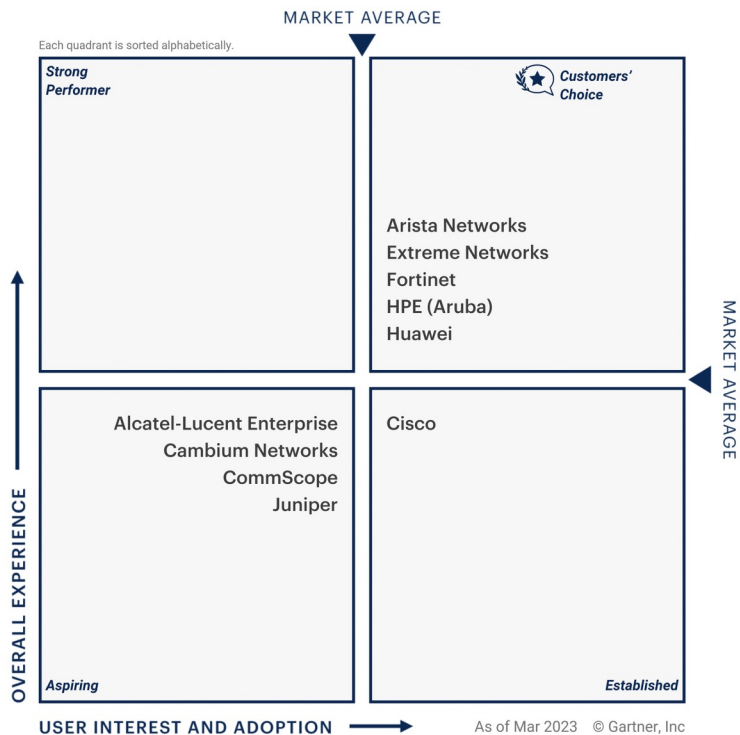
Arista recognized as a leader and visionary 6<sup>th</sup> year (2015-2020)



# Gartner Peer Insights “Voice of the Customer”

**CONGRATULATIONS!!**

Gartner Peer Insights “Voice of the Customer”  
Enterprise Wired and Wireless LAN Infrastructure



Gartner Peer Insights “Voice of the Customer”  
Enterprise Wired and Wireless LAN Infrastructure

Eligible Vendors	Overall Reviews	Willingness to Recommend
Alcatel-Lucent Enterprise	46	94%
Arista Networks	106	99%
Cambium Networks	49	96%
Cisco	67	96%
CommScope	23	98%
Extreme Networks	247	93%
Fortinet	100	99%
HPE (Aruba)	127	97%
Huawei	196	96%
Juniper	31	93%

Sorted alphabetically

As of Mar 2023 © Gartner, Inc

Gartner

# ARISTA

EOS – Modern Network  
Operating System



# Arista EOS Philosophy – It's all about QUALITY

Criteria	Total in ~15 years
# of EOS code versions across ~60+ shipping products	1
Software Field notices	54
Software Regression bugs	8
EOS Security Advisories	42
Security Advisories Requiring Downtime	2
Quality Control Testbeds	~1,000
Auto-tests Per Day	230,000+
EOS Defect Rate	< 1%

Quality comes from three places:

- ✓ Culture
- ✓ Architecture
- ✓ Testing



Quality is the most important attribute of our products. Not one of the most important, it's the most important.


Ken Duda - CTO

**Stability, reliability, and security starts with software quality**

# Our Aspiration from Day 1: Build a Better NOS

## Arista Extensible Operating System (EOS)

### A better architecture leads to a more reliable switch



- Linux Kernel - Standard and fully open
- Agent - completely isolated processes
- NetDB – contains all state
- Publish / Subscribe model
  - NetDB delivers state between all agents
- Hardware Abstraction
  - One binary for all hardware platforms

### Key Benefits:

Fault Isolation  
no other processes or data plane impacted

Fault Repair  
immediately restart a failed agent

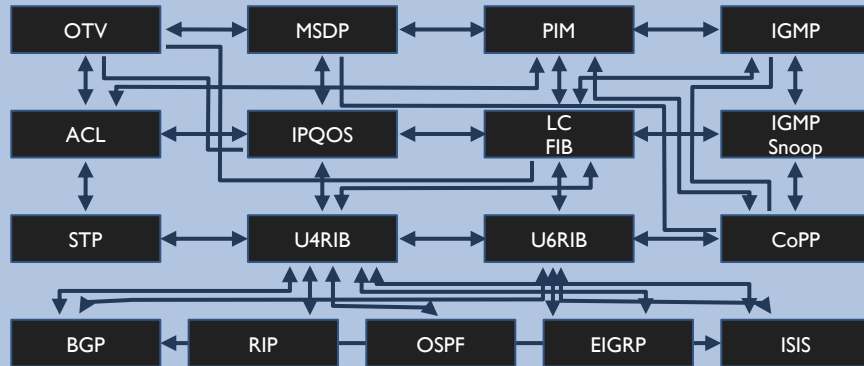
ISSU  
install RPMs or bug fixes or field upgrades with no  
downtime

3rd Party Integration  
seamless integration 3rd party

One, Single binary image for all platforms

# Resilient Hyperscale Architecture

## Legacy – Spaghetti Code

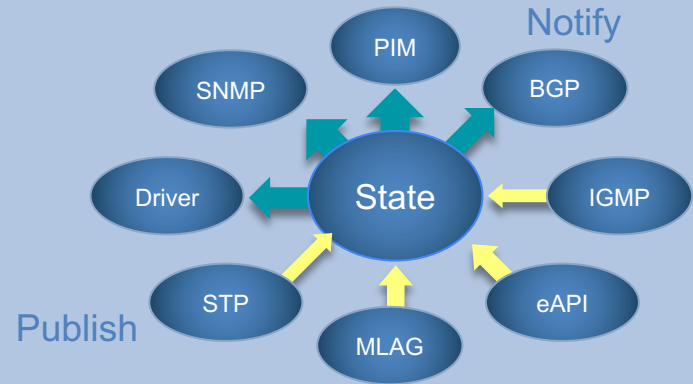


### High Inter-Process Communication

$N^2$  Messaging Passing  
Susceptible to Process Failure; Difficult Recovery

Difficult to Scale

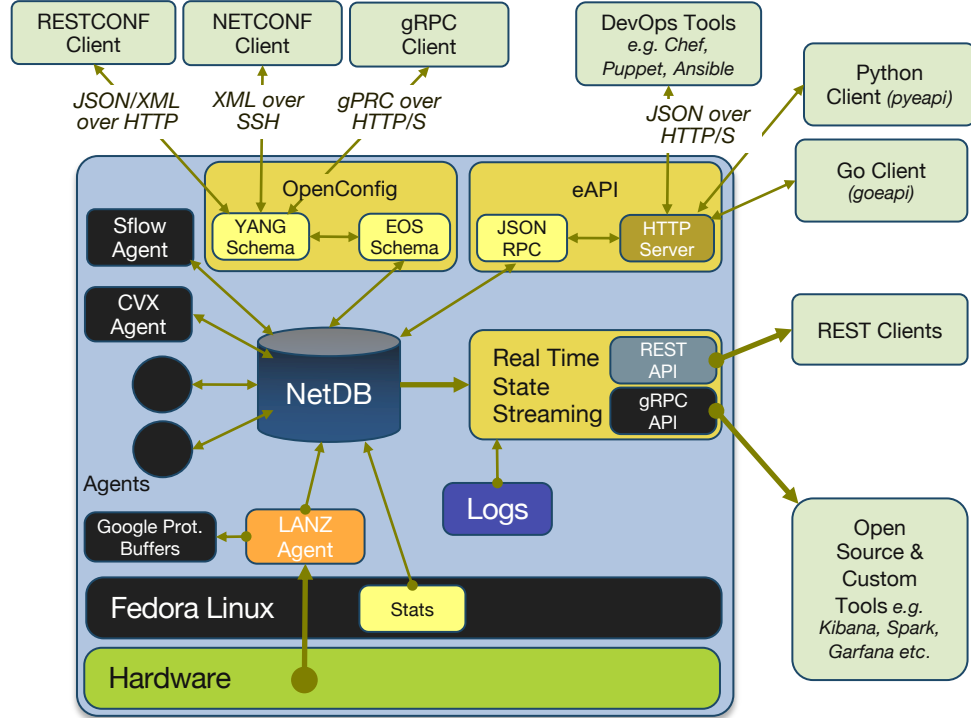
## Arista EOS – Publish / Subscribe



Efficient Publish/Subscribe

Linear Cloud Scale

# Arista EOS Architecture



## Benefits:

- Fault Isolation – no other processes or data plane impacted
- Fault Repair – immediately restart a failed agent
- Hitless Patching - install RPMs or bug fixes or field upgrades with no downtime; 91% hitless patch record
- 3<sup>rd</sup> Party Integration – seamless integration with 3<sup>rd</sup> party s/w and APIs. (i.e Splunk, Openconfig, Ansible)
- Fully Programmable Operating System

Better software architecture leads to a more reliable Carrier Grade Router



# If You Love Unix/Linux, You Will Love EOS!

## 1 Industry Standard Look and Feel:

```
Arista7124SX>
Arista7124SX >en
Arista7124SX #config t
Arista7124SX(config)#?
aaa          Authentication, Authorization, and Accounting
agent        Configure an agent
alias        Add a command alias
arp          Set a static ARP entry
banner       Configure system banners
boot         Modify system boot parameters
.....
```

## 2 To enter bash it, simply enter 'bash':

```
Arista7124SX#bash

Arista Networks EOS shell

[admin@tm225 ~]$
```

## 3 EOS commands are basically written in Python. 'Cli.py' resides in following directory :

```
[admin@tm225 site-packages]$ pwd

/usr/lib/python2.6/site-packages

[admin@tm225 site-packages]$ ls Cli*.py
Cli.py  CliExtensions.py  CliParser.py
CliAuth.py  CliInputWrapper.py  CliSave.py
```

## 4 You can view the processes running on switch:

```
[admin@tm225 flash]$ ps -ef
UID    PID  PPID  C STIME TTY          TIME CMD
root    1    0    0 May09 ?    00:00:00 /sbin/init
root    2    0    0 May09 ?    00:00:00 [kthreadd]
root    3    2    0 May09 ?    00:00:00 [migration/0]
root    4    2    0 May09 ?    00:00:00 [ksftirqd/0]
.....
```

## 5 You can easily do a tcpdump on an Arista switch.

```
tm225#bash sudo tcpdump
tcpdump: WARNING: fabric: no IPv4 address assigned
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on fabric, link-type EN10MB (Ethernet), capture size 65535 bytes
12:57:30.222171 00:1c:73:0f:80:a6 (oui Unknown) > 01:00:5e:00:00:05
(oui Unknown          ), ethertype 802.1Q (0x8100), length 82:
vlan 1, p 0, ethertype IPv4, 10.10.50.1          > OSPF-
ALL.MCAST.NET: OSPFv2, Hello, length 44
12:57:40.222687 00:1c:73:0f:80:a6 (oui Unknown) > 01:00:5e:00:00:05
(oui Unknown          ), ethertype 802.1Q (0x8100), length 82:
vlan 1, p 0, ethertype IPv4, 10.10.50.1          > OSPF-
ALL.MCAST.NET: OSPFv2, Hello, length 44
12:57:50.223172 00:1c:73:0f:80:a6 (oui Unknown) > 01:00:5e:00:00:05
(oui Unknown          ), ethertype 802.1Q (0x8100), length 82:
vlan 1, p 0, ethertype IPv4, 10.10.50.1          > OSPF-
ALL.MCAST.NET: OSPFv2, Hello, length 44
```

## 6 All images are saved in 'mnt/flash'. You can view, copy, remove, and the start-up config in bash:

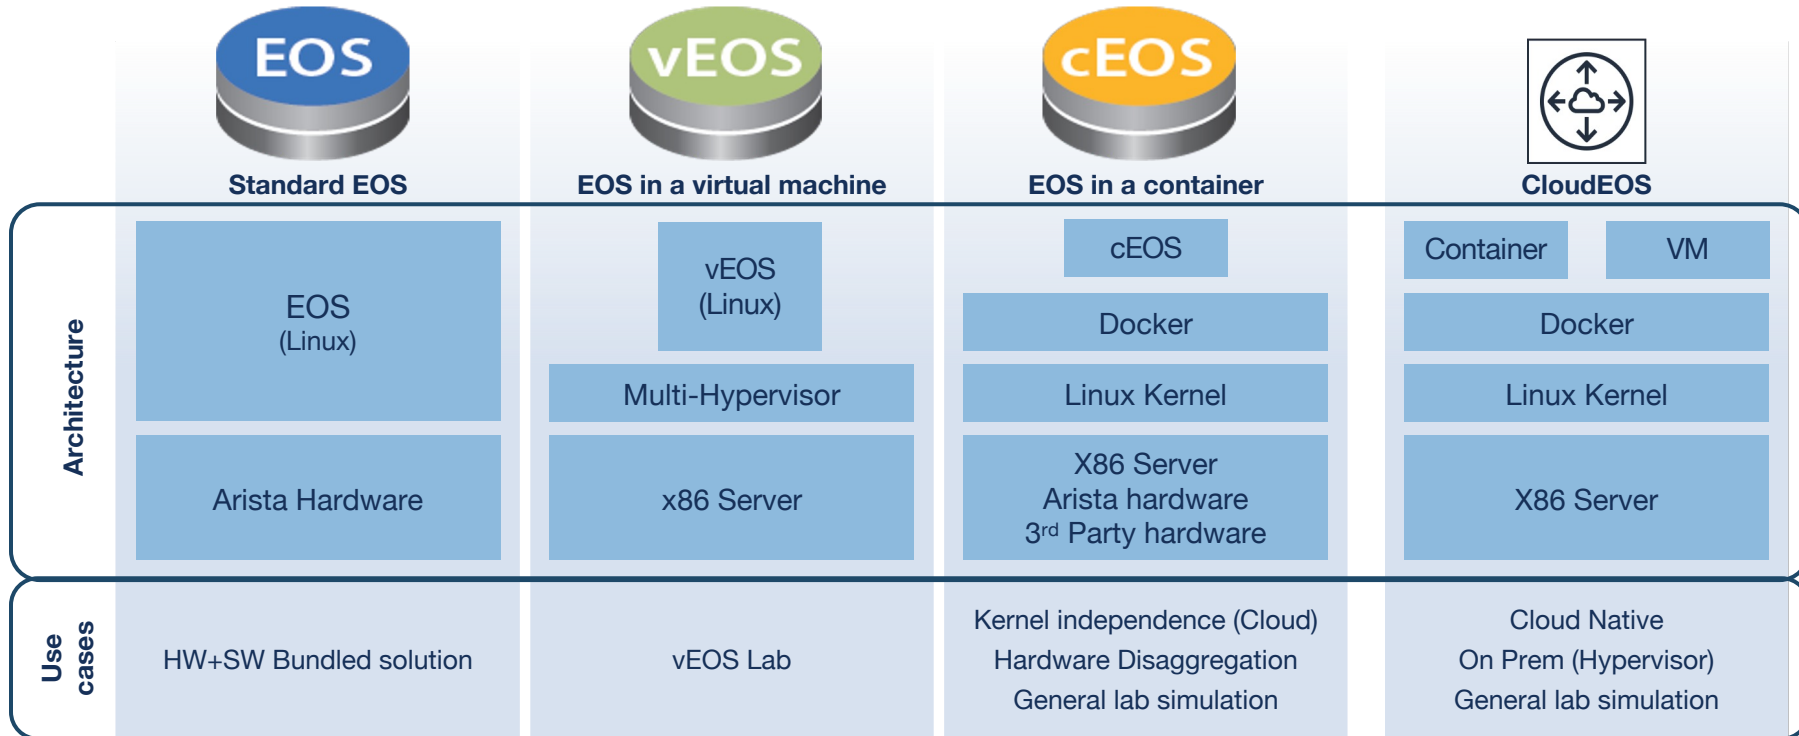
```
[admin@tm225 flash]$ pwd
/mnt/flash
[admin@tm225 flash]$ cat startup-config
! device: tm225 (DCS-7048T-A, EOS-4.6.3)
!
aaa root secret 5 $1$Vnh3khl/$d80eMfdcmHstP2gEou00/
aaa authentication policy local allow-nopassword-remote-login
!
prompt %H.%D{%H:%M:%S}%P
!
hostname tm225
ip name-server 172.22.22.10
ip name-server 172.22.22.40
ip domain-name aristanetworks.com
!
.....
```

## 7 You can create and add new commands by simply writing your own python script:


```
localhost# cat -n HelloCli.py
 1 import BasicCli, CliParser
 2 tokenHello = CliParser.KeywordRule('hello', helpdesc='Says hello')
 3 def doHello( mode ):
 4     print "World.\n"
 5 BasicCli.UnprivMode.addCommand( ( tokenHello, doHello ) )
localhost#show ?
...
hello Show your hello
...
localhost#show hello ?
| Output modifiers
<cr>
localhost#show hello
Hello World.
```

# EOS Platform Flexibility

## One EOS. Many packaging options



# EOS Security Track Record



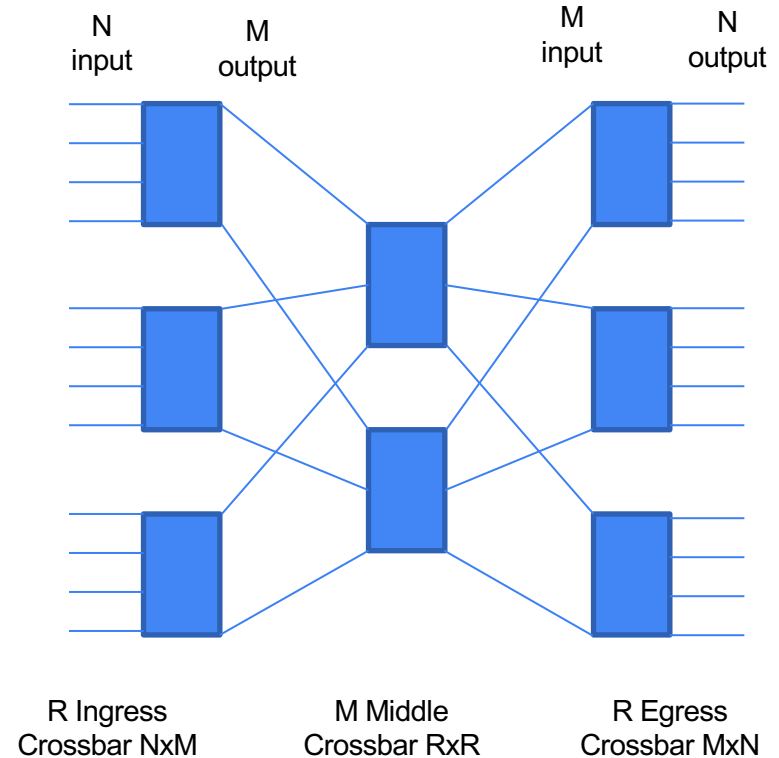
# ARISTA

Universal Cloud Network  
Architectures



# Principles

- A Clos network is a kind of multistage circuit-switching network and first formalized by Charles Clos in 1952.
- Non-blocking Clos network required  $M \geq N$

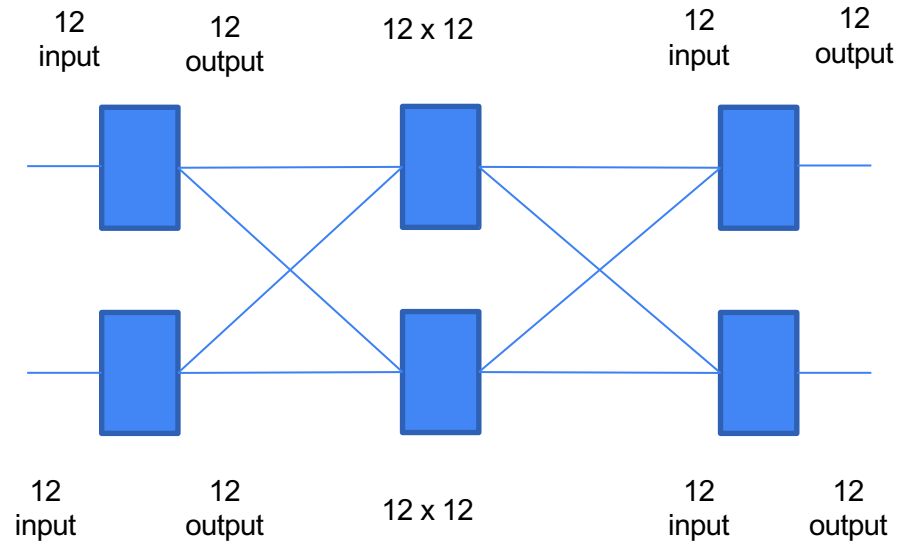


# Principles

- 48 Server with only 1G NIC
- Have only 24 port 1G Switch
- How many switches at least to build non-blocking network?

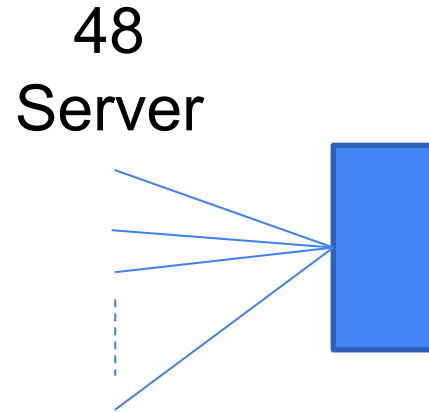
# Principles

- 48 Server with only 1G NIC
- Have only 24 port 1G Switch
- **6** switches at least to build non-blocking network.



# Principles

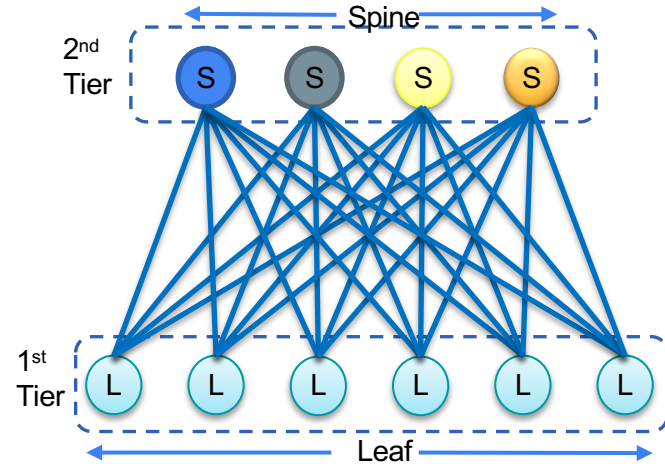
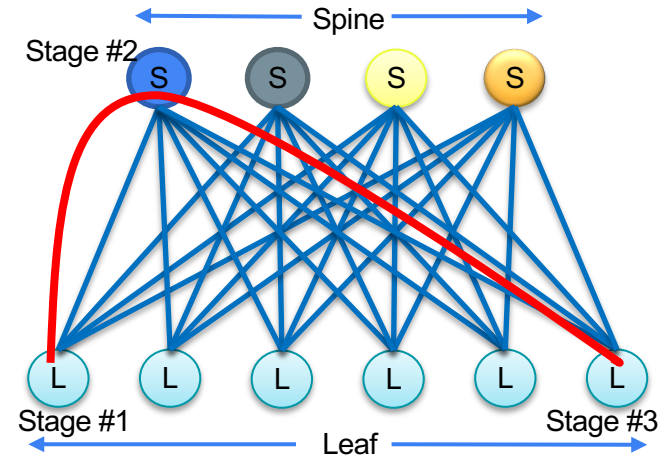
- 48 Server with only 1G NIC
- We wish we have 48 port 1G Switch
- And just need only **1** Switch





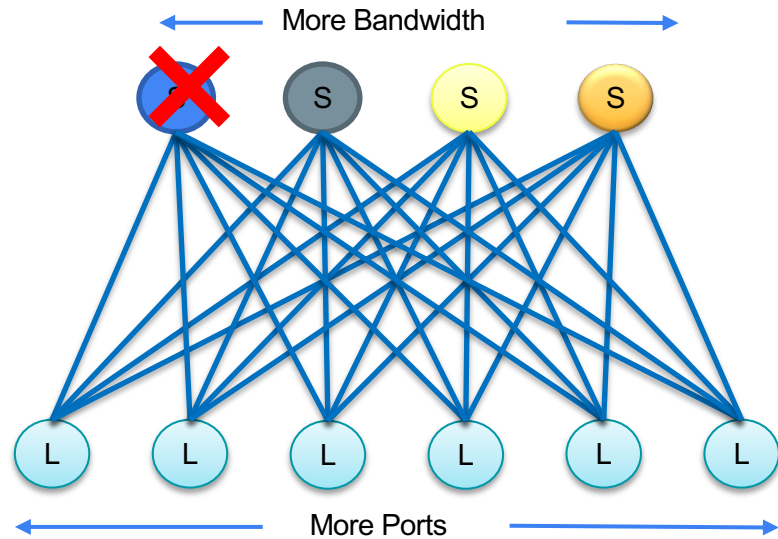
# Principles

- Leaf and Spine
- Folded Clos
- 3 Stage Clos
- 2 Tier Network
- Fat Tree






# Principles

- A Scale Out Architecture
  - More Leaf = More Ports
  - More Spine = More Bandwidth
- N +1 Redundancy
- On Spine failure
  - 4 Spine = 25% Impact
  - 8 Spine = 12.5% Impact



# Principles

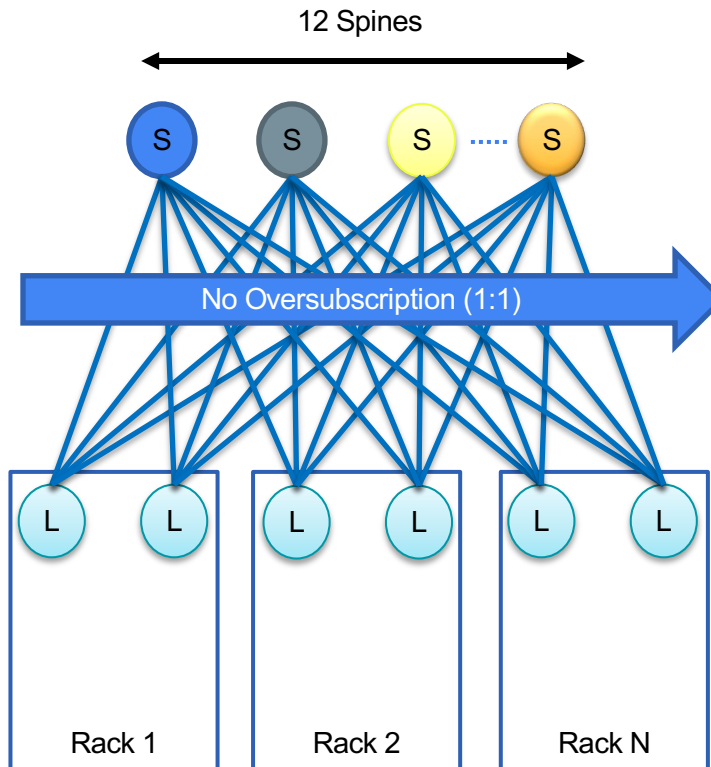
- West-East traffic account for more than 85% of traffic.
- Optimized for East to West
  - Consistent Latency from Leaf to Leaf
  - Wide ECMP
- Flexibility for North to South
  - External Connectivity at Leaf or Spine Layer



# The number of Spines?

## Server requirement

- 24 Server per Rack
- 2 x 25Gbps NIC per Server
- 1 x NIC per switch



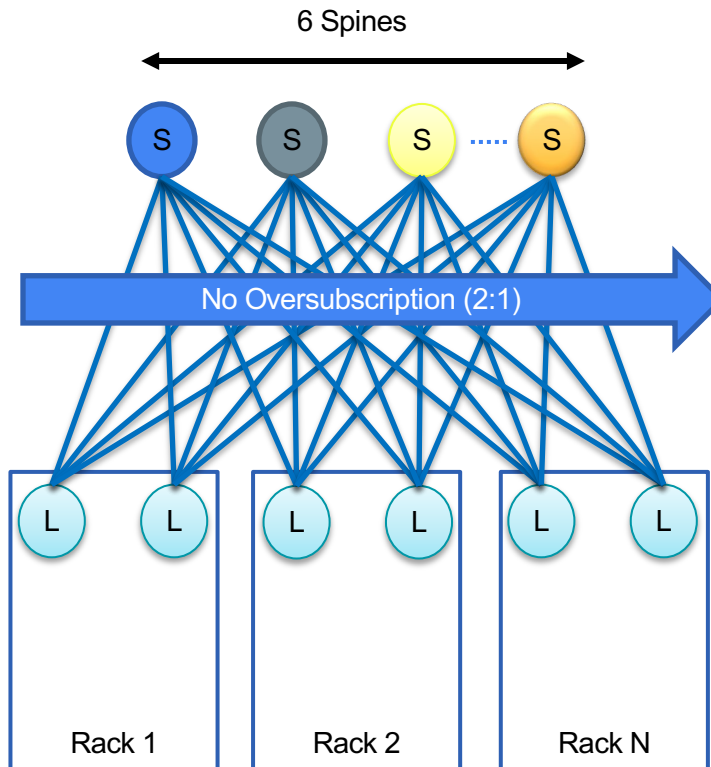
## Uplink Requirements

- 48 x 25 Gbps per Leaf
- 1.2 Tbps Uplink from Leaf to Spine
- **12 x 100 Gbps** toward Spine
- With 100G between Leaf-Spine, we need 12 Spines for 1:1 Oversubscription.

# The number of Spines?

## Server requirement

- 24 Server per Rack
- 2 x 25Gbps NIC per Server
- 1 x NIC per switch




## Uplink Requirements

- 48 x 25 Gbps per Leaf
- 1.2 Tbps Uplink from Leaf to Spine
- **12 x 100 Gbps** toward Spine
- With 100G between Leaf-Spine, we need 6 Spines for 2:1 Oversubscription.

# The number of Spines?

## Server requirement

- 24 Server per Rack
- 2 x 25Gbps NIC per Server
- 1 x NIC per switch




## Uplink Requirements

- 48 x 25 Gbps per Leaf
- 1.2 Tbps Uplink from Leaf to Spine
- **12 x 100 Gbps** toward Spine
- With 100G between Leaf-Spine, we need 4 Spines for 3:1 Oversubscription.

# The number of Spines?

## Server requirement

- 24 Server per Rack
- 2 x 25Gbps NIC per Server
- 1 x NIC per switch



## Uplink Requirements

- 48 x 25 Gbps per Leaf
- 1.2 Tbps Uplink from Leaf to Spine
- **12 x 100 Gbps** toward Spine
- With 100G between Leaf-Spine, we need 2 Spines for 6:1 Oversubscription.

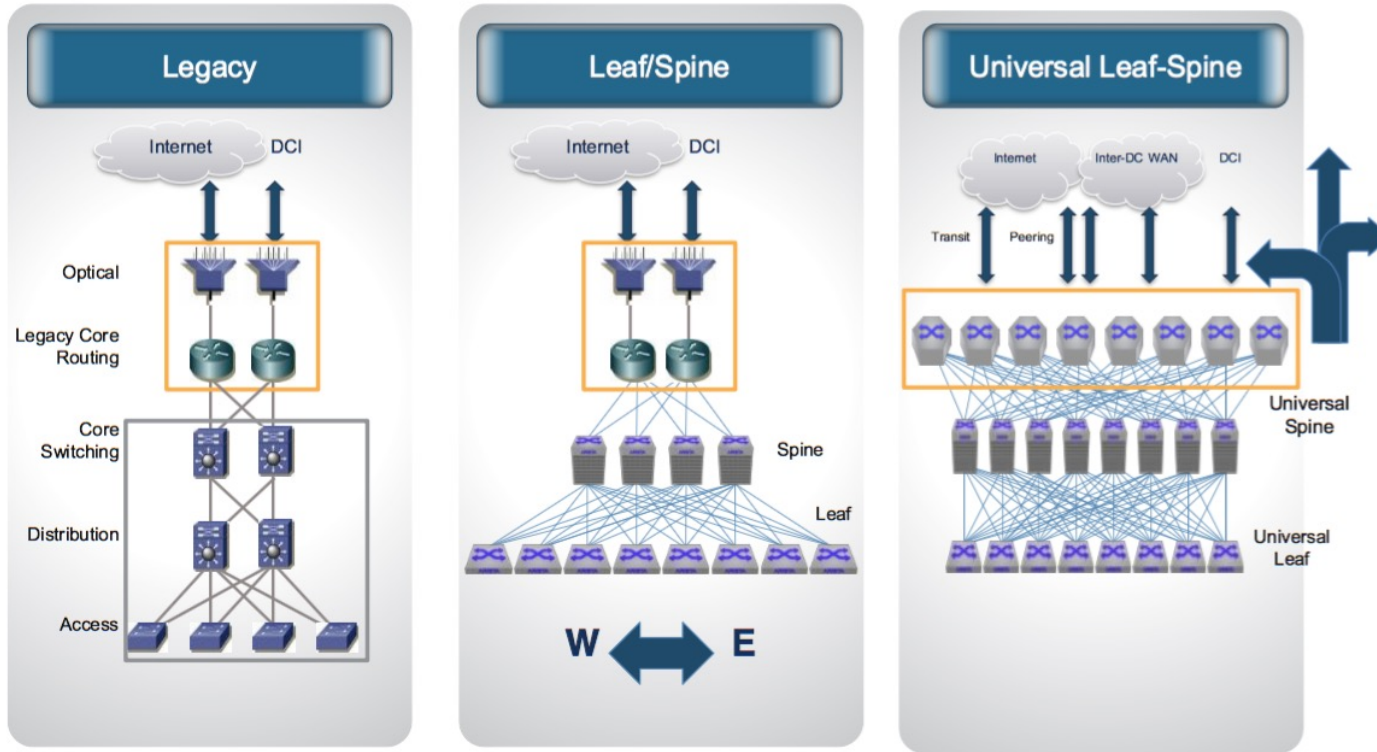
# Common DC Network: A system with Many Parameters

- ❑ Bandwidth & Capacity
  - Servers and services link speed, NICs & Optics
  - Switch ASICs, Platforms
  - Systems with mix of 25G to 400G speeds, also can support 800G in Core layer
- ❑ Scale and scalability
  - Topology and routing
  - MACsec and IPSec at high speeds
  - Lifecycle, Deployment and retrofits
- ❑ Cabling & Powering
  - Fiber infrastructure, high efficiency with 100G SerDes
  - Distance
  - Power budget
- ❑ Automation and management

Timelines: need by vs. technology availability and development



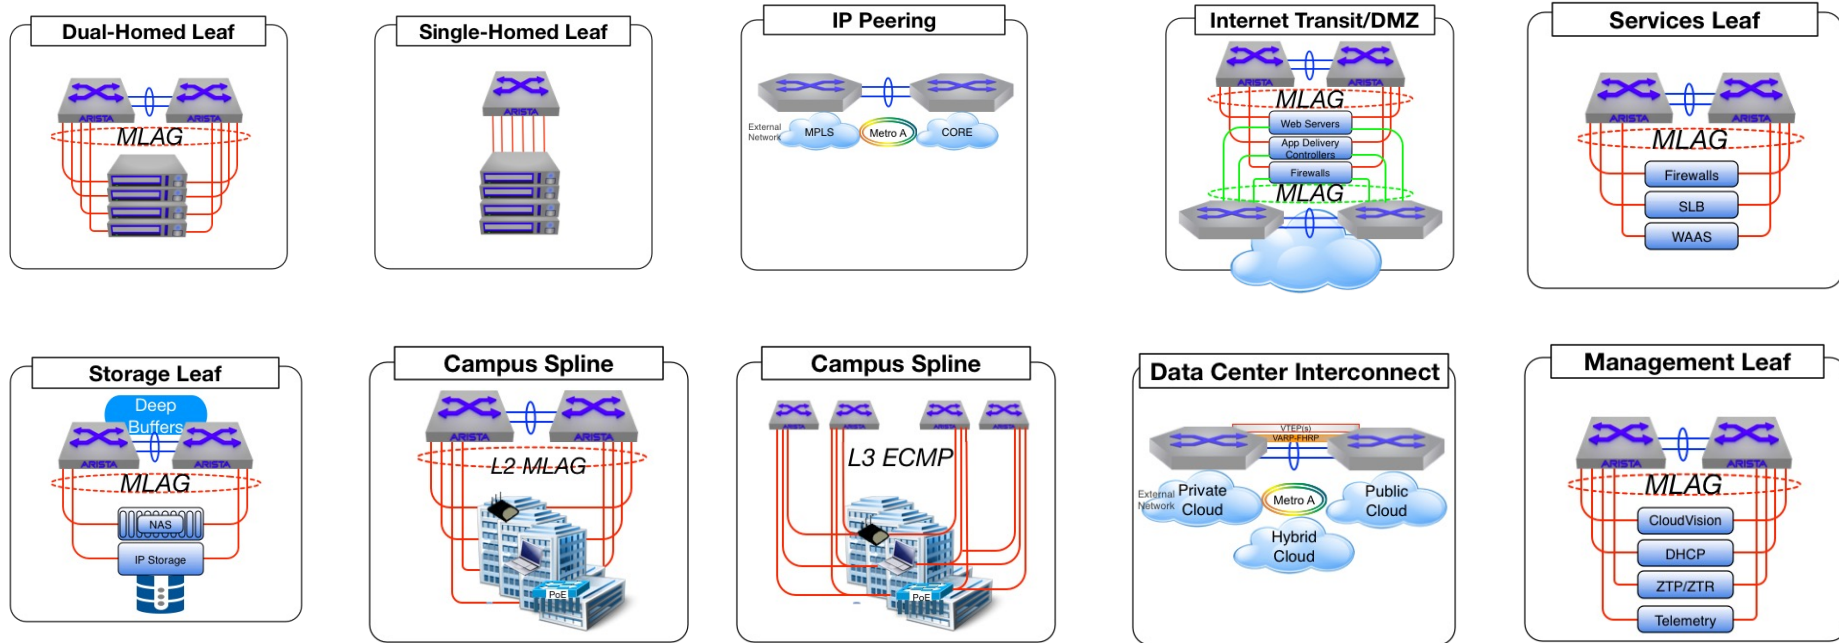
# Data Center Network Architecture Evolution



Routing is Integrated into the Universal Spine and Leaf

# Universal Cloud Network Design


Using a consistent architecture across all leaf switches and cabinets can easily support any portion of the network without having to relearn how each section of the topology is constructed.



# Universal Cloud Networking Architectures


## Layer 2 and Layer 3 Leaf/Spine network for data scale

Layer 2 / MLAG



100 to 10,000

Layer 3 / ECMP



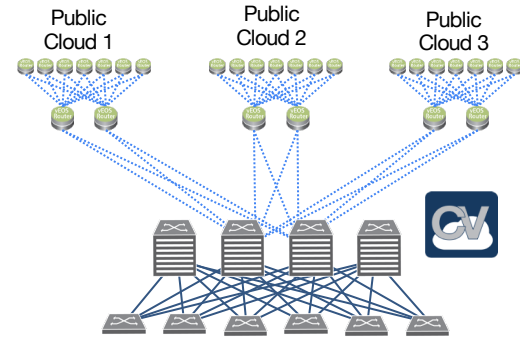
500 to 100,000+

VXLAN with EVPN



500 to 100,000+

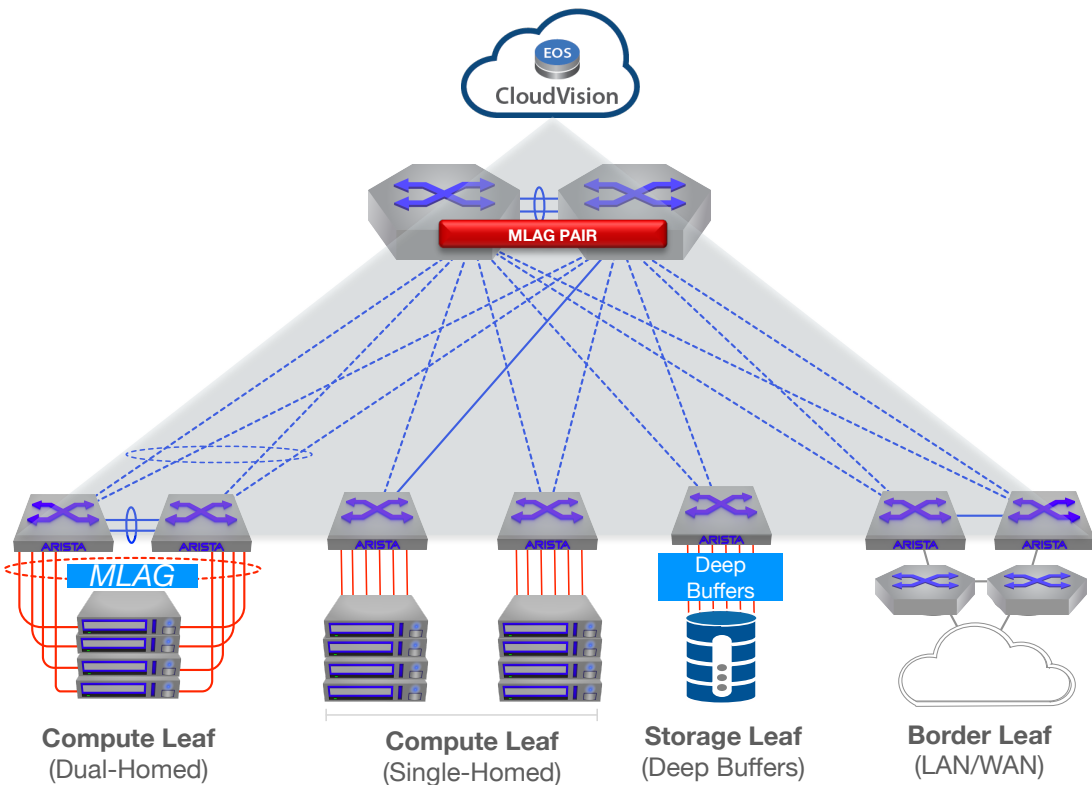
Hybrid Cloud



500 to 100,000+

Standard protocols - no proprietary fabrics

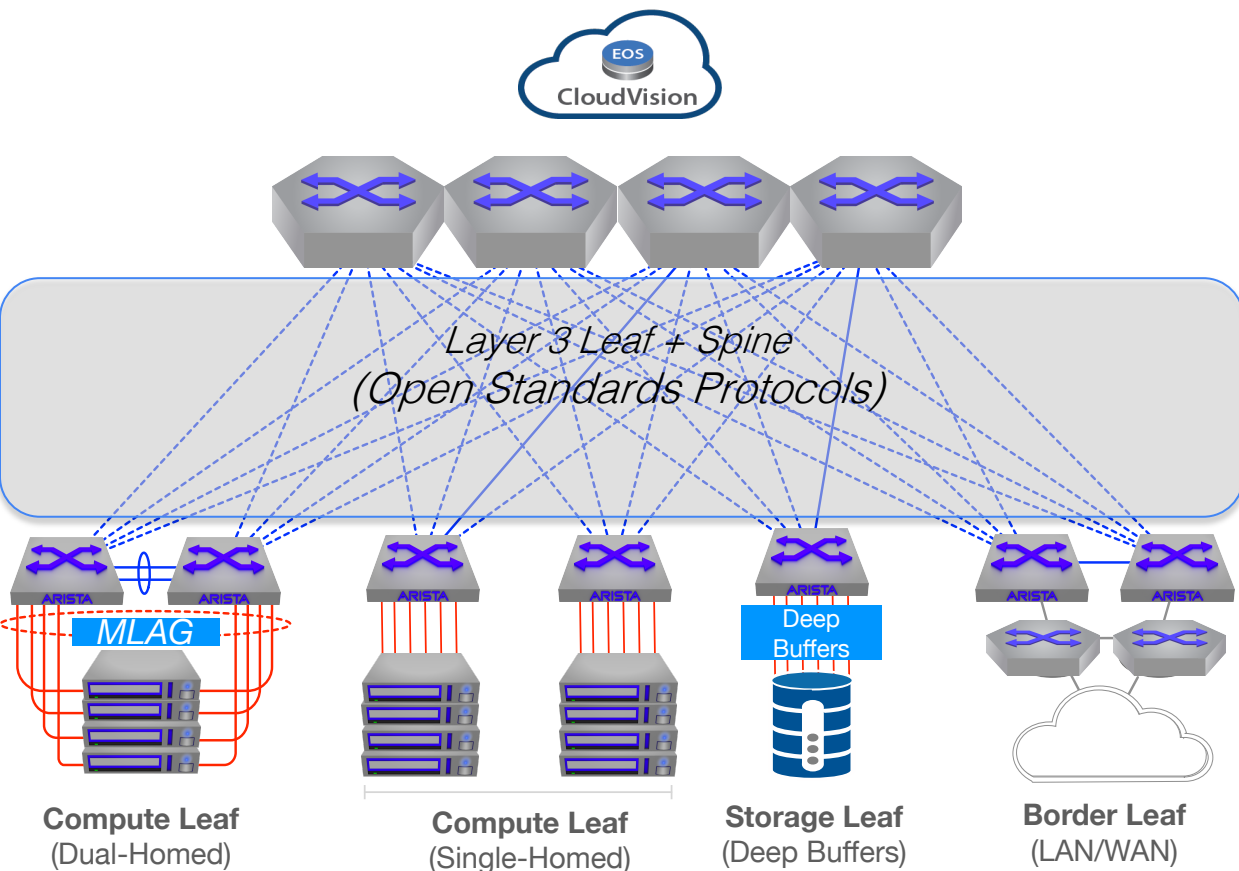
# Arista UCN – L2LS



## L2LS MLAG Spine Design

- Ease of Workload Mobility
- Segmentation via VLANs
- Simple design to provision and manage
- All Active links
- Leverage Standards based protocols
- Consistent Leaf designs (ToR)
- Spine to interconnect leafs
- Scale limited by MAC, VLAN and Spine density

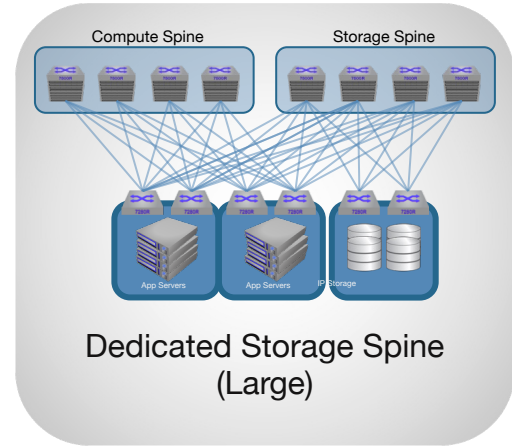

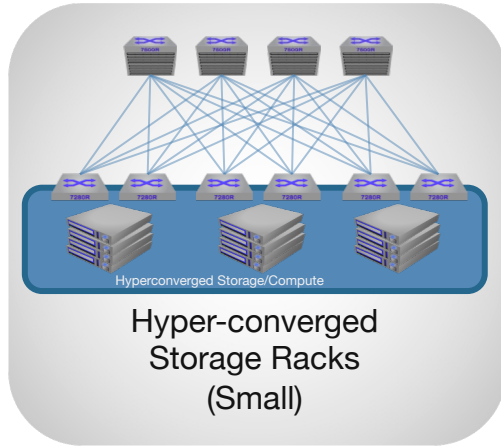
# Arista UCN – L3 LS



## L3LS ECMP Spine Design

- Spine redundancy
- Grow/scale from 2 to 128 way
- Minimal fault domains
- Deterministic failover and simpler troubleshooting
- (ToR) Leafs as FHR
- eBGP between Leaf & Spine
- Stability by reducing scope of Layer 2 domain
- Deployed and Managed by CloudVision

# Arista's IP Storage/Big Data UCN



UCN Architecture, Dense switching

Deep Buffering, VOQ, DCB/PFC



Smart System Upgrades, Self-healing OS

LANZ, VM Tracer, MapReduce, Tracer

CloudVision provisioning

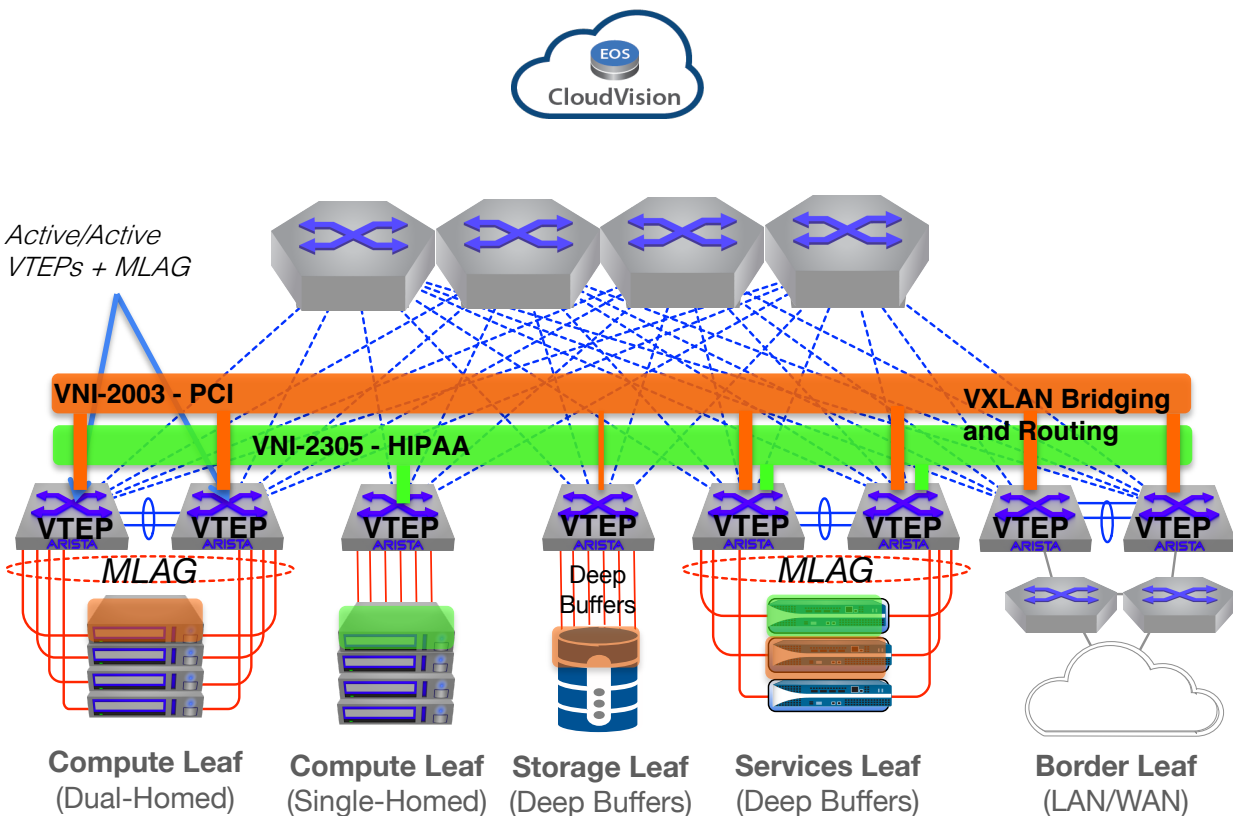
Low power consumption, compact form factors

# Arista Universal Leaf 7280: 100X More Buffer than traditional Leaf.



**Lossless: Deep packet buffer for Big Data, IP Storage, HPC & incast scenarios**

# Arista UCN – L3LS-With EVPN

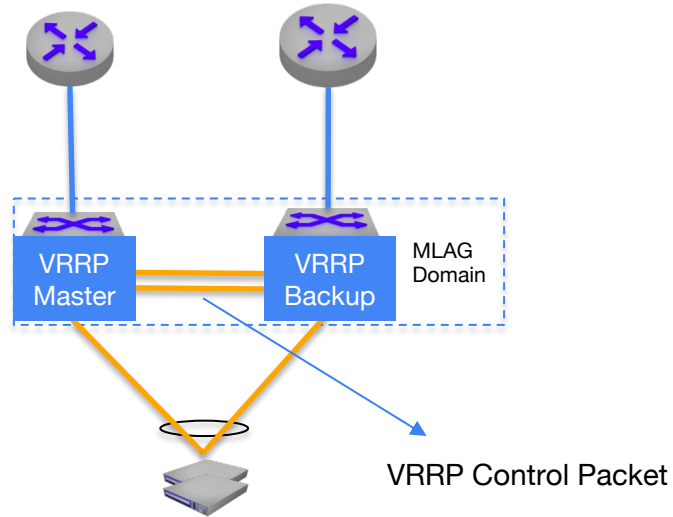


## L3LS-EVPN Use Case

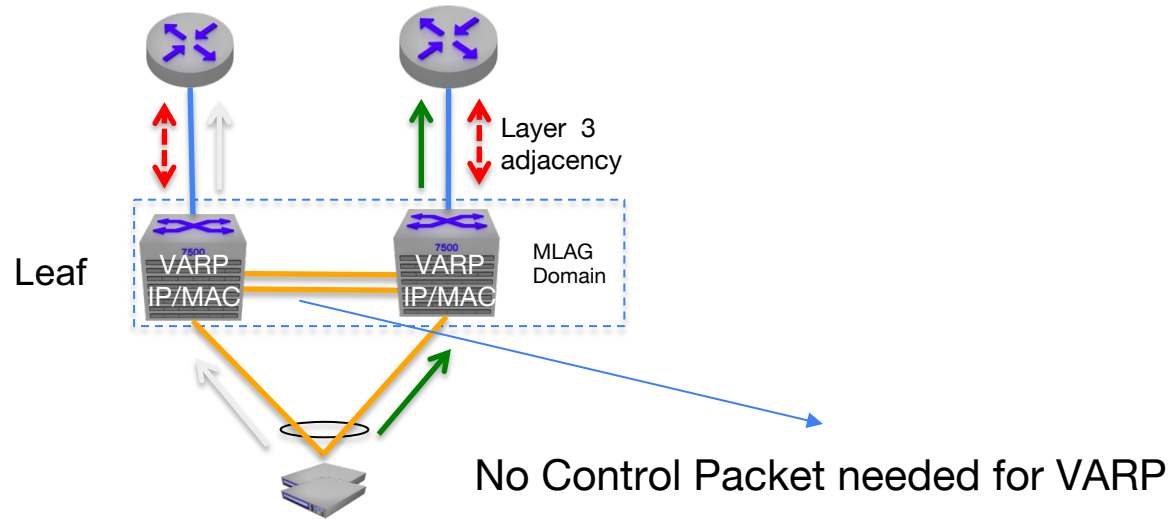
- Network Based Overlay
- Physical Virtual Tunnel End Points (VTEP's)
- BGP for underlay and overlay
- Enable Macro-Segmentation Services with VXLAN
- CloudVision Orchestration, Management, and Learning




# First Hop Redundancy Protocol – VRRP, HSRP,...



# First Hop Redundancy Protocol - VARP (Virtual ARP)

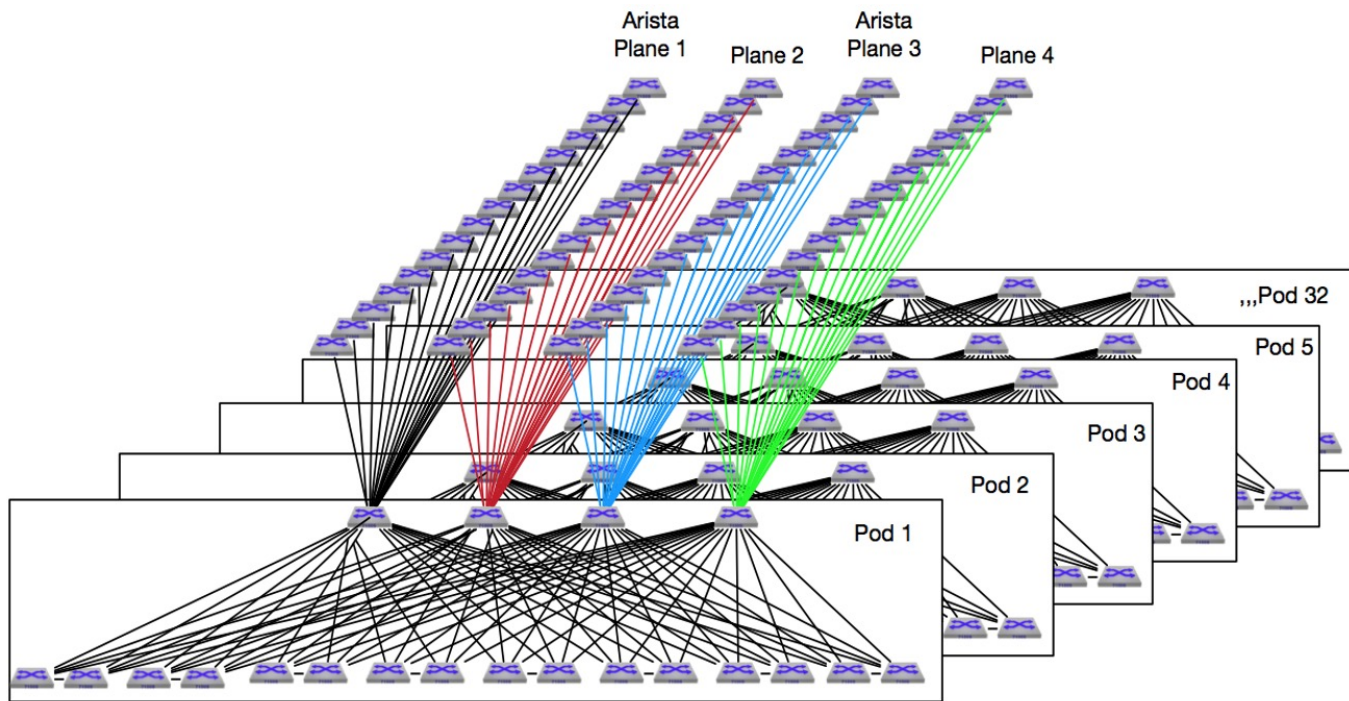


# Multi-Pod Architecture

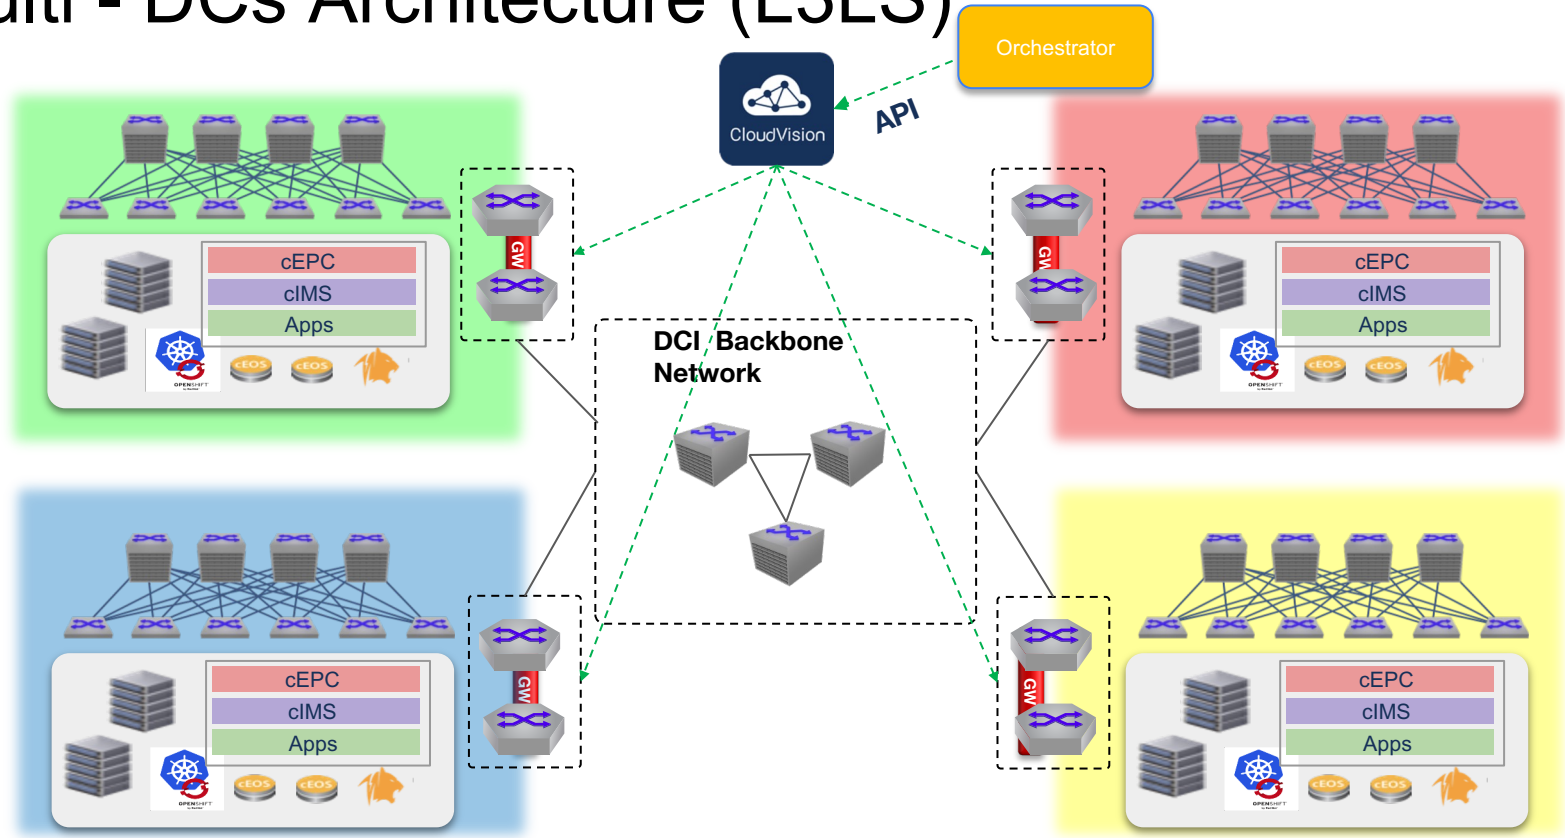


- Any Network Design (L2LS/L3LS/L3LS-V)
- Multiple Data Room Interconnect
- All with the same rich services (MSS/Virtualization/Tracers/DANZ/LANZ)
- CloudVision Orchestration, Management, and Learning

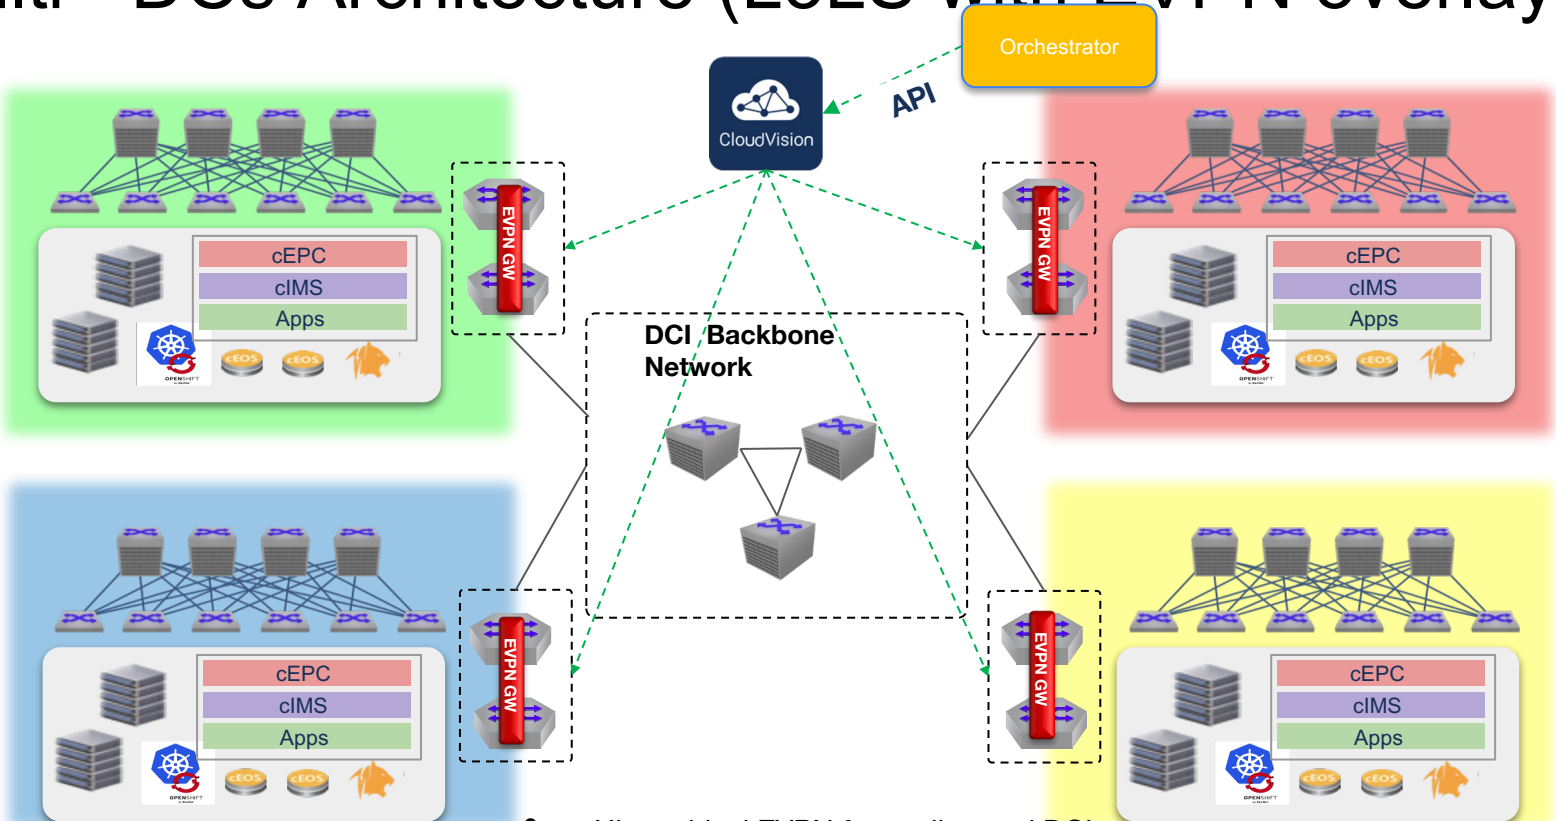
# Multi Plane Super Spine – Hyper Scale Architecture



# Multi - DCs Architecture (L3LS)



# Multi - DCs Architecture (L3LS with EVPN overlay)



- Hierarchical EVPN for scaling and DCI
- Interconnect EVPN domains geographically separate locations

# EVPN GW – Hierarchical EVPN for scaling and DCI

□ IETF BESS working group, number of RFCs/Drafts for EVPN GW behavior

- Support for both Layer 2 (type-2 & 3) and 3 (type-5) DCI solutions
- Interop across different BGP Address families and data-plane encapsulations (VXLAN, PBB, MPLS)

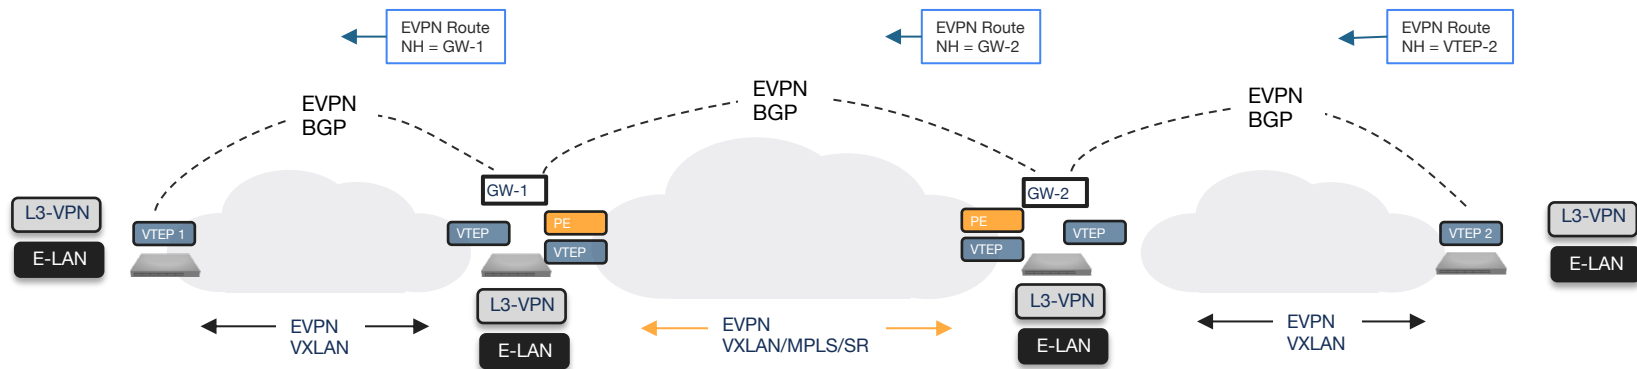
Draft	Overview
A Network Virtualization Overlay Solution using EVPN <b>RFC 8365</b>	EVPN control plane for L2 VPNs with an NVO environment with VXLAN, NVGRE and GENEVE encap- DCI using GWs and DCI using ASBRs
EVPN and IP-VPN Integrated Solution <b>draft-ietf-bess-evpn-ipvpn-interworking-07</b>	Layer 3 DCI interop between EVPN-VXLAN/MPLS and IP-VPN WAN for layer 3 DCI
Multi-site EVPN based VXLAN using Border Gateways <b>draft-sharma-bess-multi-site-evpn-01</b>	GW DCI solution focused only on EVPN-VXLAN, support for a single control planes (EVPN) and single data-plane (VXLAN)
Interconnect Solution for EVPN Overlay networks <b>RFC 9014</b>	EVPN GW solution for L2 interconnecting of multiple control planes (VPLS/EVPN) and data-planes (MPLS, VXLAN, PBB)
EVPN multicast forwarding for EVPN to EVPN GWs <b>draft-rabnic-bess-evpn-mcast-eeg-00</b>	EVPN GW solution for providing seamless multicast interconnect between EVPN domains, across VXLAN and MPLS data-planes
Domain Path (D-PATH) for Ethernet VPN (EVPN) Interconnect Networks <b>draft-sr-bess-evpn-dpath-02</b>	D-path community for EVPN routes to provide loop-free route advertisement between EVPN domains for layer 2.

L3 GW solution

Industry adopted L2 GW

# EVPN GW – EVPN VXLAN/MPLS GW

- How does EVPN GW work?



## EVPN GW behavior

- PE/VTEP nodes EVPN peer with their local GW node via eBGP or iBGP
- GW node EVPN peer with the GW nodes in the remote domain via eBGP or iBGP
- Import received type-2 & 5 routes based on RT policy
- Export type-2 & 5 routes between domains based on RT policy
- When exporting between domains, new Next-hop, encap and label

## Benefits

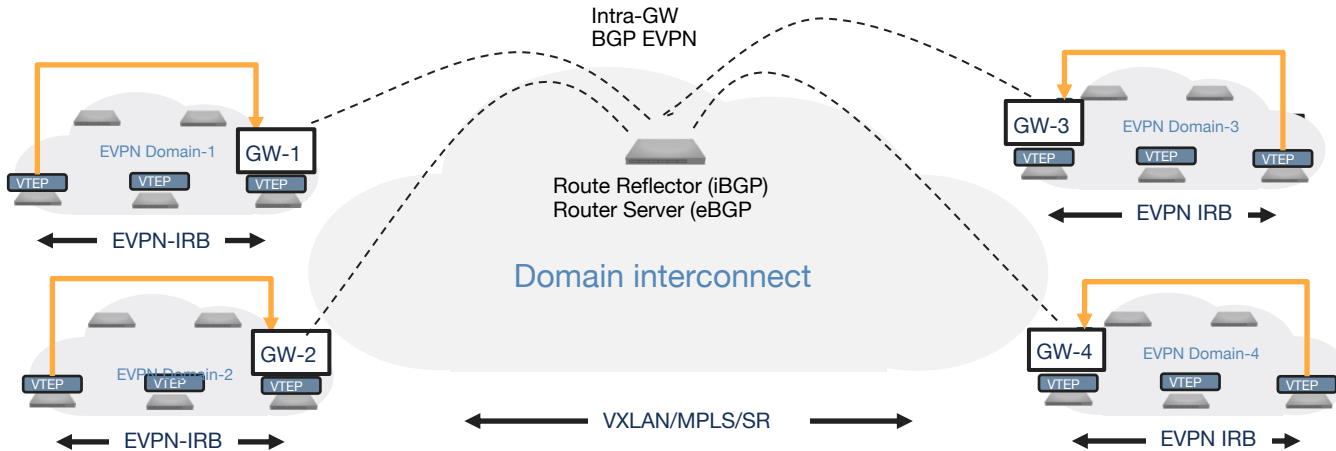
- End-to-End Layer 2 and 3 connectivity regardless of interdomain encap
- Support L2 and L3 VPN between VXLAN VTEPs and MPLS PE nodes
- EVPN A-A for GW redundancy for L2 interconnect across domains
- Hierarchical flood-list for BUM traffic forwarding
- Reduction in EVPN state churn across domains



# EVPN GW – EVPN VXLAN/MPLS GW

## □ EVPN Gateway Solution


- Not just a point-to-point DCI solution
- Support for multiple domains improved DC scale with inter-POD(s) or inter-Site(s) connectivity
- Support for all BGP topology variants in the Local and remote domains (eBGP and iBGP)



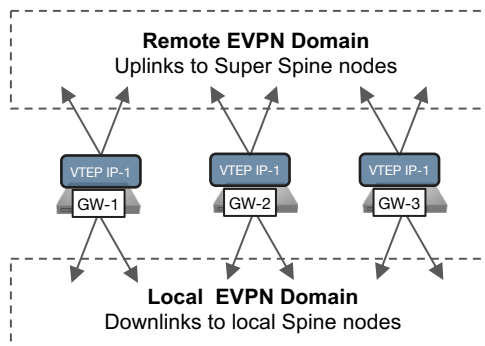
# EVPN Gateway Resiliency Models

- To provide resiliency, EVPN GW nodes can be deployed in choice of topologies
- The chosen approach depends on the level of resiliency required and how appliances/services will be attached to the GW nodes, if required

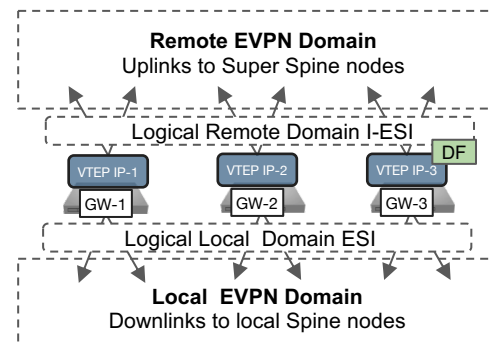
EVPN-GW with MLAG



EVPN-GW with Anycast IP



EVPN-GW with All-Active

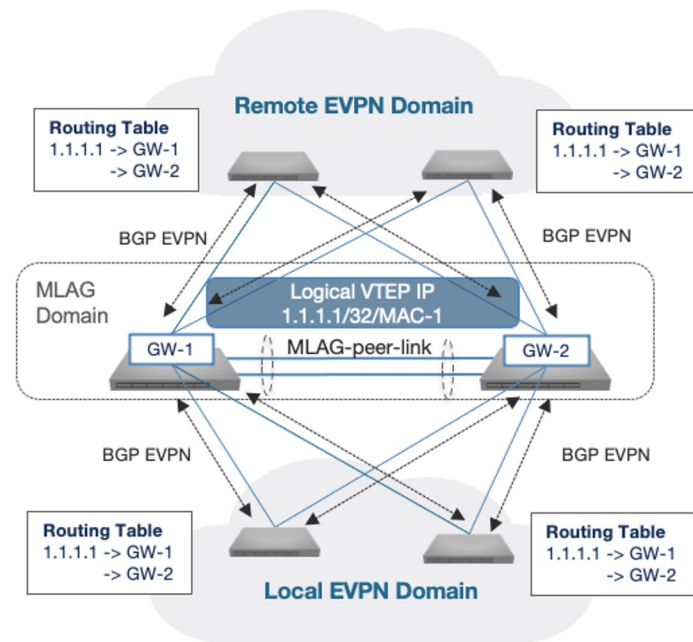


Note that R-Series products are recommended as EVPN Gateways

# MLAG Gateway for EVPN VXLAN

- **Pair of interconnected gateway nodes** within a domain are configured in an MLAG topology to also provide **host connectivity**.
- Each node has EVPN peering sessions with both the local and remote domains
- EVPN routes advertised between the domains
- The two GW nodes act as a single logical VTEP by using the same shared loopback IP and router MAC address

```
mLAG configuration → Standard MLAG configuration between GW-1 and GW-2
domain-id GW1-GW2-Domain
local-interface Vlan4094
peer-address 10.0.0.1
peer-link Po101
!
interface Loopback1
ip address 1.1.1.1/32 → Shared loopback IP address
!
interface Vxlan1
vxlan source-interface Loopback1 → Logical GW VTEP IP for the MLAG
vxlan virtual-router encapsulation mac-address mlag-system-id → Rtr-MAC for the GW VTEP
!
router bgp 64512
address-family evpn
neighbor EVPN_SPINE activate → EVPN Local domain peering with Spine
neighbor EVPN_RR activate
neighbor EVPN_RR domain remote → EVPN Remote domain peering with a RR
```



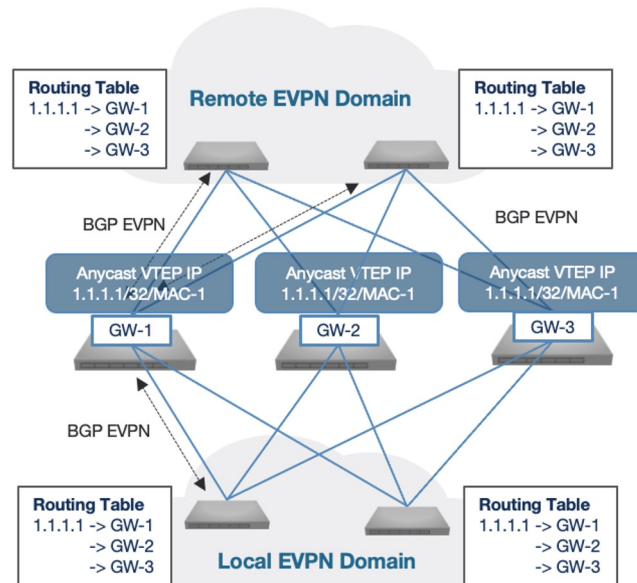
Note that MLAG is required for hosts

# Anycast Gateway for EVPN VXLAN

- **Gateway nodes** within a domain are **not interconnected**
- Each have EVPN peering sessions with both the local and remote domains
- For resiliency, nodes within a domain will share the same anycast VTEP IP and router MAC address, meaning the GW nodes within a domain act as a single anycast VTEP
- **No restriction on the number of gateway nodes** that can be deployed within a single domain

```
interface Loopback1
  ip address 1.1.1.1/32 → Anycast IP address
!
interface Vxlan1
  vxlan source-interface Loopback1 → Anycast GW VTEP IP
  vxlan virtual-router encapsulation mac-address 00:bb:bb:bb:bb:bb → Anycast GW Rtr-MAC
!
router bgp 64512
  address-family evpn
    neighbor EVPN_SPINE activate → EVPN Local domain peering with Spine
    neighbor EVPN_RR activate
    neighbor EVPN_RR domain remote → EVPN Remote domain peering with a RR
```

Same on All Gateway Nodes

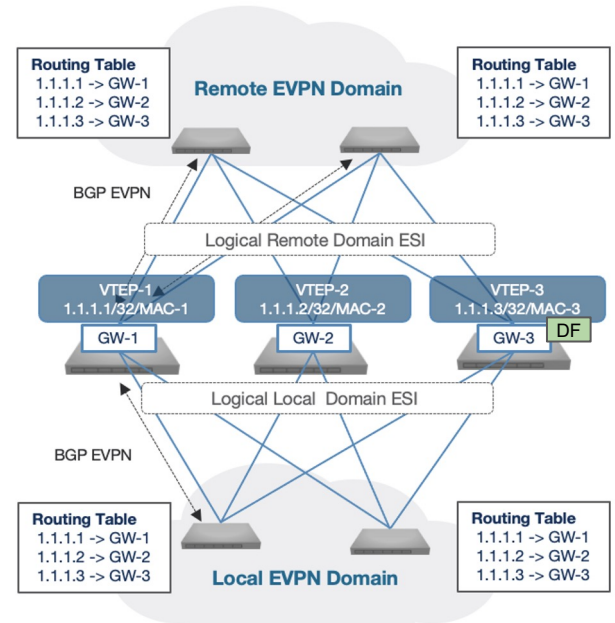


Note that front-panel ports are **NOT** supported

# All-Active Gateway for EVPN VXLAN

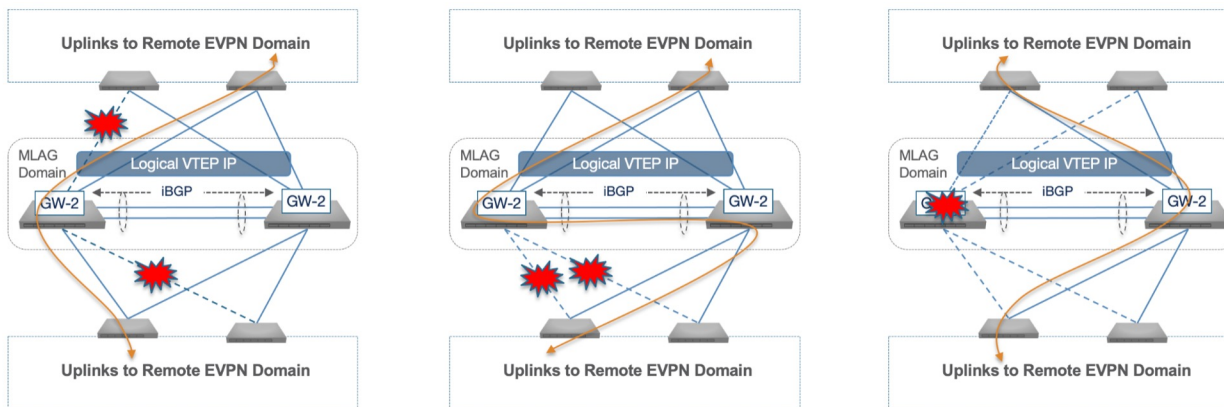
- **Gateway nodes** within a domain are **not interconnected**
- Each have EVPN peering sessions with both the local and remote domains
- To provide resiliency, GW nodes within each domain are deployed using EVPN multihoming as defined in RFC 7432.
- GW nodes use a unique Ethernet Segment (ES) for the domain, termed the I-ESI (Interconnect Ethernet Segment Identifier)
- Standard election for the Designated Forwarder (DF) for BUM forwarding
- Support for up to **16 nodes** in a single I-ESI

```
interface Loopback1
  ip address 1.1.1.1/32 → Unique IP address
!
interface Vxlan1
  vxlan source-interface Loopback1 → GW VTEP IP
!
router bgp 64512
  address-family evpn
    neighbor EVPN_SPINE activate → EVPN Local domain peering with Spine
    neighbor EVPN_RR activate
    neighbor EVPN_RR domain remote → EVPN Remote domain peering with a RR
  !
  evpn ethernet-segment domain remote
    identifier 0000:0001:0001:0000:0106 → I-ESI for the local and remote domain peerings
    designated-forwarder election algorithm preference 100 → preference based DF election
    route-target import 00:00:00:00:01:06
```



# EVPN Gateway Failure Behavior – MLAG

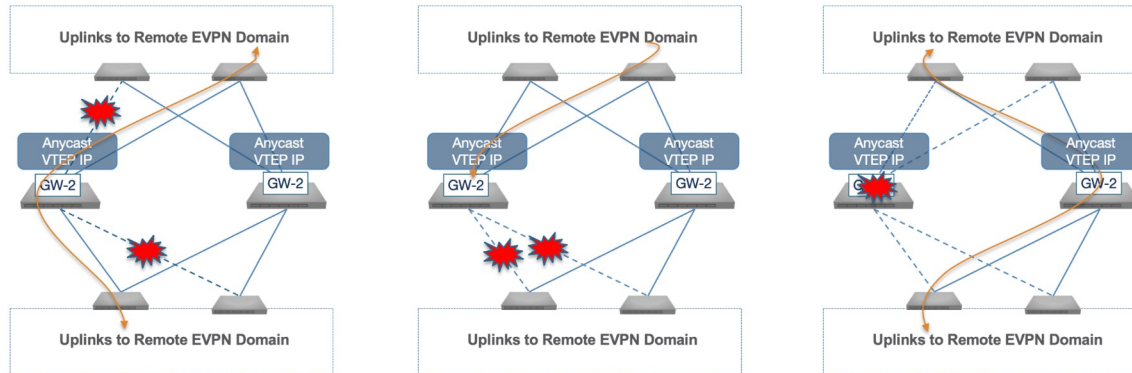
- Forwarding behavior during a failure within a domain is dependent on the resiliency model deployed
- Gateway nodes using an MLAG model are interconnected via a peer-link and IBGP session
  - Provides fast-failover in the event of single or multiple link failures
  - EVPN peering session of both GW nodes is retained
  - Traffic can take the path via the peer-link during a failure
- No additional configuration is required



**Recommendation: at least two links between MLAG peers**

# EVPN GW Failover Behaviour – Anycast and All-Active Gateway

- If a **node fails**, all EVPN sessions will go down and **traffic will route around** the failure
- If all **uplinks** to the remote domain **fail**, remote EVPN session will go down
  - **Local EVPN domain session will persist**, causing traffic to be **blackholed**
- Same issue occurs in the opposite scenario; local uplinks and EVPN session fail, remote uplinks and EVPN session remain up
- Note that Blackholing will only affect **Layer 2** traffic which is load-balanced across the GW members of an **All-Active I-ESI** which is still undesirable.



Loss of EVPN sessions to one of the domains will blackhole traffic

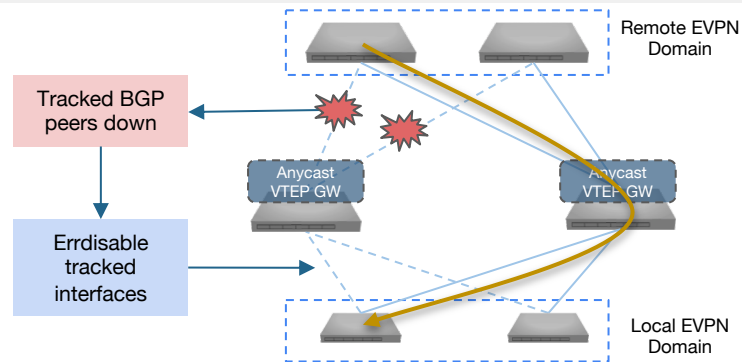
# BGP Session Tracker – Core Isolation

- Support for tracking static and dynamic BGP peers
- Tracked interface placed in “errdisabled” state
- Recovery time configurable per object
- Peer can only be a member of one tracker object
- Recovery on a single peer, not all peers
- Solution for ESI core-isolation
- **Solution for EVPN GW isolation**
  - Anycast-IP and All-Active

```
ANYCAST-GW1#show interface status errdisabled
```

Port	Name	Status	Reason
Et1/1		errdisabled	bgp-session-tracker

```
router bgp 100
  neighbor SPINE_EVPN session tracker EVPN_Local → Tracker for Local
  neighbor RR_EVPN session tracker EVPN_Remote → Tracker for Remote
  !
  session tracker EVPN_Local → BGP Session tracker definition
    recovery delay 300 seconds
  !
  session tracker EVPN_Remote → BGP Session tracker definition
    recovery delay 300 seconds
  !
  interface Ethernet1/1,2/1 -> Interface to errdisable
    bgp session tracker EVPN_Remote
  interface Ethernet3/1,4/1 -> Interface to errdisable
    bgp session tracker EVPN_Local
```



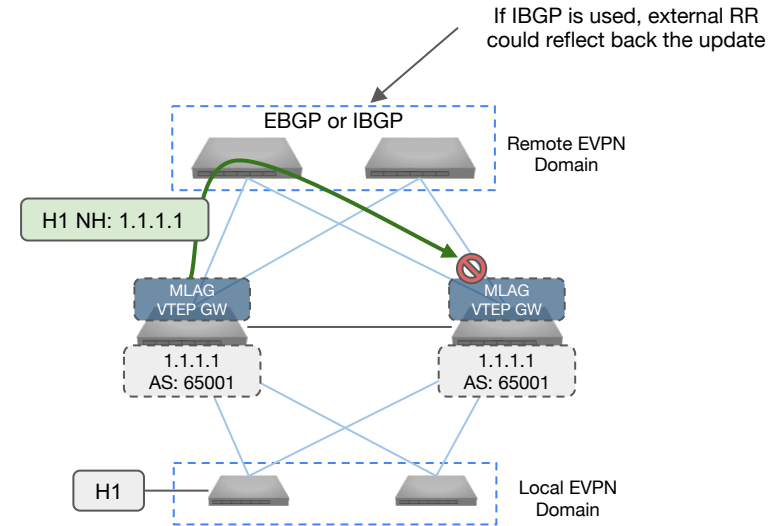
Helps with traffic blackholing by tracking BGP session state



# Gateway Loop Prevention Mechanisms

## MLAG & Anycast IP

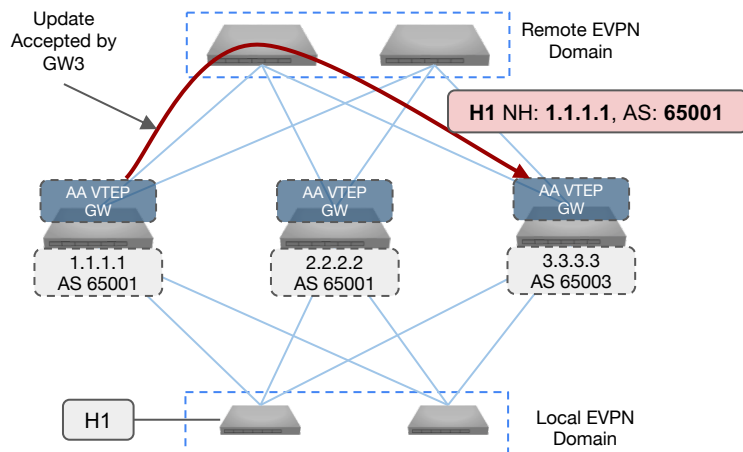
- EVPN Update Nexthop uses the shared logical VTEP loopback IP.
- Peers will mark the updates as invalid as the nexthop will match a local interface.
- Route-reflectors will re-advertise updates from GWs.
  - If a unique RD is configured per peer, both will be reflected.
  - Shared RD will cause only one update
  - But update will be ultimately marked invalid due to the local interface match



No Additional Configuration Required

# Gateway Loop Prevention Mechanisms

## Redundant All-Active Gateways

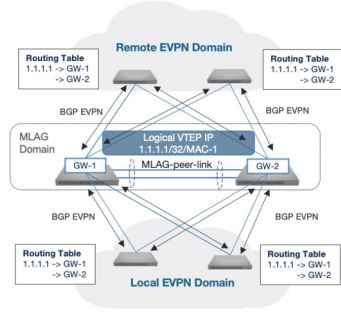


- Each peer uses a **unique VTEP IP/Nexthop & RD** (due to unique ASNs)
- Has the potential to be re-advertised back to another gateway peer and accepted causing a loop.
- A **route-map is required** to filter updates advertised by peer gateways in each domain.
- A future version of EOS will support **D-PATH** attribute as per *draft-sr-bess-evpn-dpath* to avoid loops

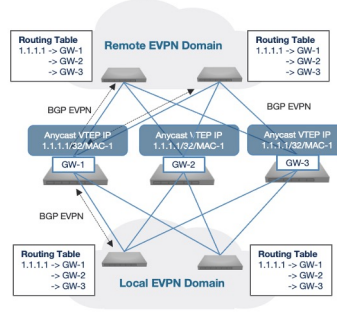
Filter Updates from other Gateways using a Route-Map

# EVPN Gateway Resiliency Model Summary

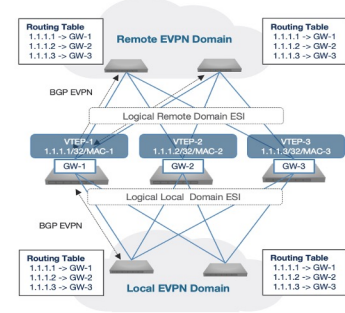
MLAG - 4.26.1F



Anycast-IP - 4.27.1F



All-Active - 4.29.0F




Max # of nodes	Limited to 2 nodes - Single MLAG domain site	No Restriction on # of GW per site	Up to 16 GW nodes
	Single Logical VTEP IP	Single Logical VTEP IP	VTEP IP per Node
r	No additional routes – MLAG Sync	No additional routes for state sync	Type-1 (AD per ES/EVI) – Type-4 (ES)
r	Next-hop Logical VTEP IP	Next-hop Anycast VTEP IP	VTEP IP of node, with ES on Type-2 routes
	Shared Logical IP – GW receive and forward a single copy	Shared Logical IP – GW receive and forward a single copy	Each GW receives a copy only DF forwards BUM out
	Underlay Load Balancing	Underlay Load Balancing	Overlay Load Balancing
	Supported	Not Supported	Supported
Loop Prevention	Native, no additional configuration	Native, no additional configuration	Should be considered and route-maps applied

# LS Topology with number of odd Spines?

Could we build the spine layer with the number of 3,5,.. spines?

Arista say **yes** with the high number ECMP path



# ARISTA

## Platform Overview



# Arista's Hardware Strategy

- Use the best merchant silicon available
- Design for datacenter, campus, service provider and special use cases
- Build for scalable, high-performance, ultra-low-latency networks and cloud computing networks
- Innovate to solve real world problems





# Arista Product Overview - 2023



Single Image Arista EOS Across All Platforms

# Arista Solution– Merchant Silicon, One EOS Image

Modular



Fixed



Standard EOS



EOS in a virtual machine



EOS in a container

General Data Centre

Scale Out & Richness

Single consistent EOS Image across all platforms across multiple applications

Telemetry & Analytics

Campus

Routing & Switching

IP Storage & Peering

OnPrem DC & DCI

M&E

TAP Agg



# Silicon Roadmap 2023 and beyond

- ❑ **Tomahawk2**, TH3, TH4, TH5 (speed, density, 100G, 400G,800G)

7060X

7368X

7388X

- ❑ **Trident3**, T3.\* T4.\* (DC features, 25G, 100G)

7050X

7300X

7358X

720X

- ❑ **Jericho2**, J2c, J2c+, J3 (buffers, features, tables, 100G, 400G, 800G)

7280R

7500R

7800R

- ❑ **Barefoot** Tofino, Tofino2 (flexibility, scale, L4)

7170

- ❑ **Xilinx** and **L1** (low latency switching)

7130

- ❑ \*: Alta, T+, H4, J, J+, Q, T2, T2+, TH, TH+, XP (still going supported)
- ❑ MACsec at 400G, 100G, 25G, 10G and more

# Arista 7000 Series products

To see full list of Arista 7000 Series products,  
Please visit:

<https://www.arista.com/en/products/platforms>

# Arista 7000 Series products

How to interpret Arista SKU:

## Majority Front Panel Ports

C = QSFP100

D = QSFP-DD

P = OSFP

Y = SFP28

S = SFP

7280CR3A-24D12

**Standard Scale**  
Standard Scaling  
No Encryption

7280R3A

7280CR3AM-24D12

7280CR3AK-24D12

## Chipset Family

R3 = Broadcom Jericho2

R3A = Broadcom Jericho2C+

## Encryption

Standard Scaling  
MACsec, IPsec,  
TunnelSec

## Large Scale

Large Scale Routing

# 2023: Next Generation Silicon for Customer Networks

## Feature Rich



### Trident4

4X Higher Performance  
12.8Tbps and 132MB Buffer  
Programmable Pipeline

## Highest Bandwidth



### Tomahawk4

2X Higher Performance  
25.6Tbps with 64 x 400G  
Scale Out and High Radix

## Carrier Grade Scale



### Jericho2C+

50% Higher Performance  
7.2Tbps and 2.7Bpps  
Deep Buffers and Extensible

Consistent High Performance and Extensible EOS

# Domain-Specific Products for Data Center Networks

## 7050X Series



Enterprise application stacks  
RoCEv2, EVPN, VXLAN  
Rich Telemetry for deep visibility  
Compute TOR for 10/25/50/100G  
Flexible traffic management

## 7060X Series



Cloud application stacks  
Highest switch performance  
Lowest Latency  
Scale Out & High Radix  
High density 400G Fixed Spines  
128 x 200G in 4RU

## 7280/7800R Series



Storage TOR with deep buffers  
EVPN, MPLS, SR  
High capacity scaling  
Metro & DCI with MACSec & ZR  
Fixed and Modular form factors

Consistent High Performance and Extensible EOS

# Key Silicon Families - Single EOS

	Trident 7050X	Tomahawk 7060X	Jericho 7280R/7800R
Table Scale	****	***	*****
Port speed	****	*****	****
Radix - fixed	****	*****	****
Radix - chassis	***	****	*****
Flexibility	****	***	*****
Buffering	****	***	*****
Variants	TD3*,TD4*	TH2,TH3,TH4,TH5	J2*,Q2*,J3,Q3*
EVPN	****	**	*****
Routing	****	**	*****

# Data Center and Routing Portfolio




## 7800R Series

100G and 400G, 800G ready  
Dynamic Deep Buffer



## 7500R Series

10G to 400G  
Dynamic Deep Buffer



## Arista 7280R Series

10G to 400G  
Dynamic Deep Buffer



## 7360X / 7388X Series

Cloudscale Data Center Networks  
25G to 400G Leaf and Spine



## Arista 7060X Series

Scale out applications and  
Cloudscale Data Center



## Arista 7350X / 7050X Series

Large Enterprise and Carrier Networks  
Rich features

# 10G to 400G : Industry Leading Platforms for Leaf and Spine Networks



## Arista R-Series

High Performance Routers for  
Universal Leaf / Spine and Carrier  
10G to 400G  
Dynamic Deep Buffers  
Optimized for Routing and IP Storage  
Feature rich for Edge, Core and DC



## Arista X-Series

High Performance Leaf and Spine  
Scale out applications and  
Data Center Networks  
Power Efficient 10G to 400G Spine  
Machine Virtualization and Containerization  
Rich features for wirespeed L2/L3 Networks



## Arista 7170 Series

Multi-function Programmable switch  
High performance 100G  
Fully programmable pipeline  
EOS and P4 profiles for Cloud,  
Enterprise, Service Provider and CDN  
Rich telemetry and visibility

Switches and Routers Optimized for high performance and next generation networking



# The Best Data Center Spine and Routing Portfolio



## Arista 7800R Series

High Performance Routers  
100G and 400G  
Dynamic Deep Buffer  
Universal Spine and Carrier Core  
Optimized for Routing and Storage

## Arista 7500R Series

High Performance Routers  
10G to 400G  
Dynamic Deep Buffer  
Universal Spine and IP Storage  
Optimized for Routing and Storage

## Arista 7300X3 Series

High Performance  
Enterprise and Carrier Networks  
Power Efficient 10G to 100G Spine  
Machine Virtualization and Containerization  
Rich features for L2/L3 Networks

## Arista 7368/7388X Series

Highest Performance  
Scale out applications and  
Cloudscale Datacenter Networks  
25G to 400G Leaf and Spine  
Optimized for throughput and density

# The Best Data Center Leaf and Routing Portfolio



## Arista 7280R Series

High Performance Routers  
10G to 400G  
Dynamic Deep Buffer  
Universal Leaf and Spine  
Optimized for Routing and Storage



## Arista 7060X Series

Highest Performance  
Scale out applications and  
Cloudscale Datacenter Networks  
25G to 400G Fixed Leaf and Spine  
Optimized for throughput and density



## Arista 7050X Series


High Performance  
Enterprise and Carrier Leaf Networks  
10G to 100G Fixed Leaf and Spine  
Machine Virtualization and Containerization  
Rich features for L2/L3 Networks

# Why Merchant Network Silicon is Winning

## Merchant Silicon Firsts

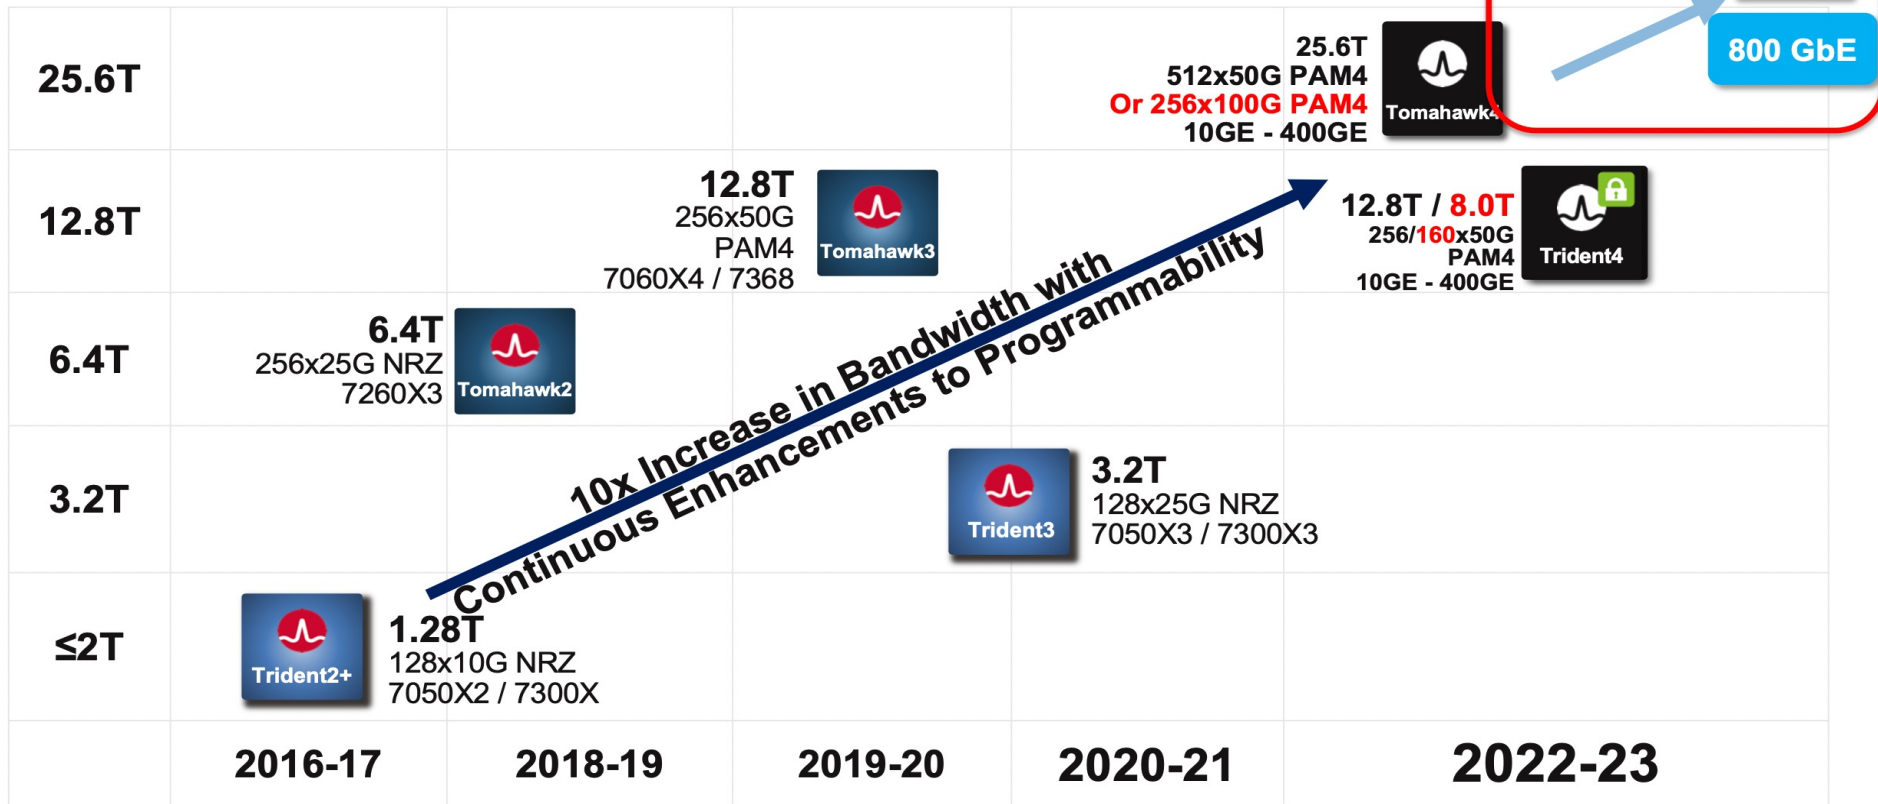
- 2008: First ultra-low latency 24-port 10G single chip
- 2010: First Large Buffer 10G Chip with VOQ Fabric
- 2011: First 64-port 10G single chip switch
- 2012: First 32-port 40G single chip
- 2013: First Large Buffer 40G Chip with VOQ Fabric
- 2015: First 32-port 100G single chip
- 2016: First Router 100G Chip with VOQ Fabric
- 2017: First 64-port 100G single chip
- 2018: First 32-port 400G single chip
- 2019: First Router 400G Chip with VOQ Fabric
- 2021: First 64-port 400G single chip

## Bandwidth Improvement

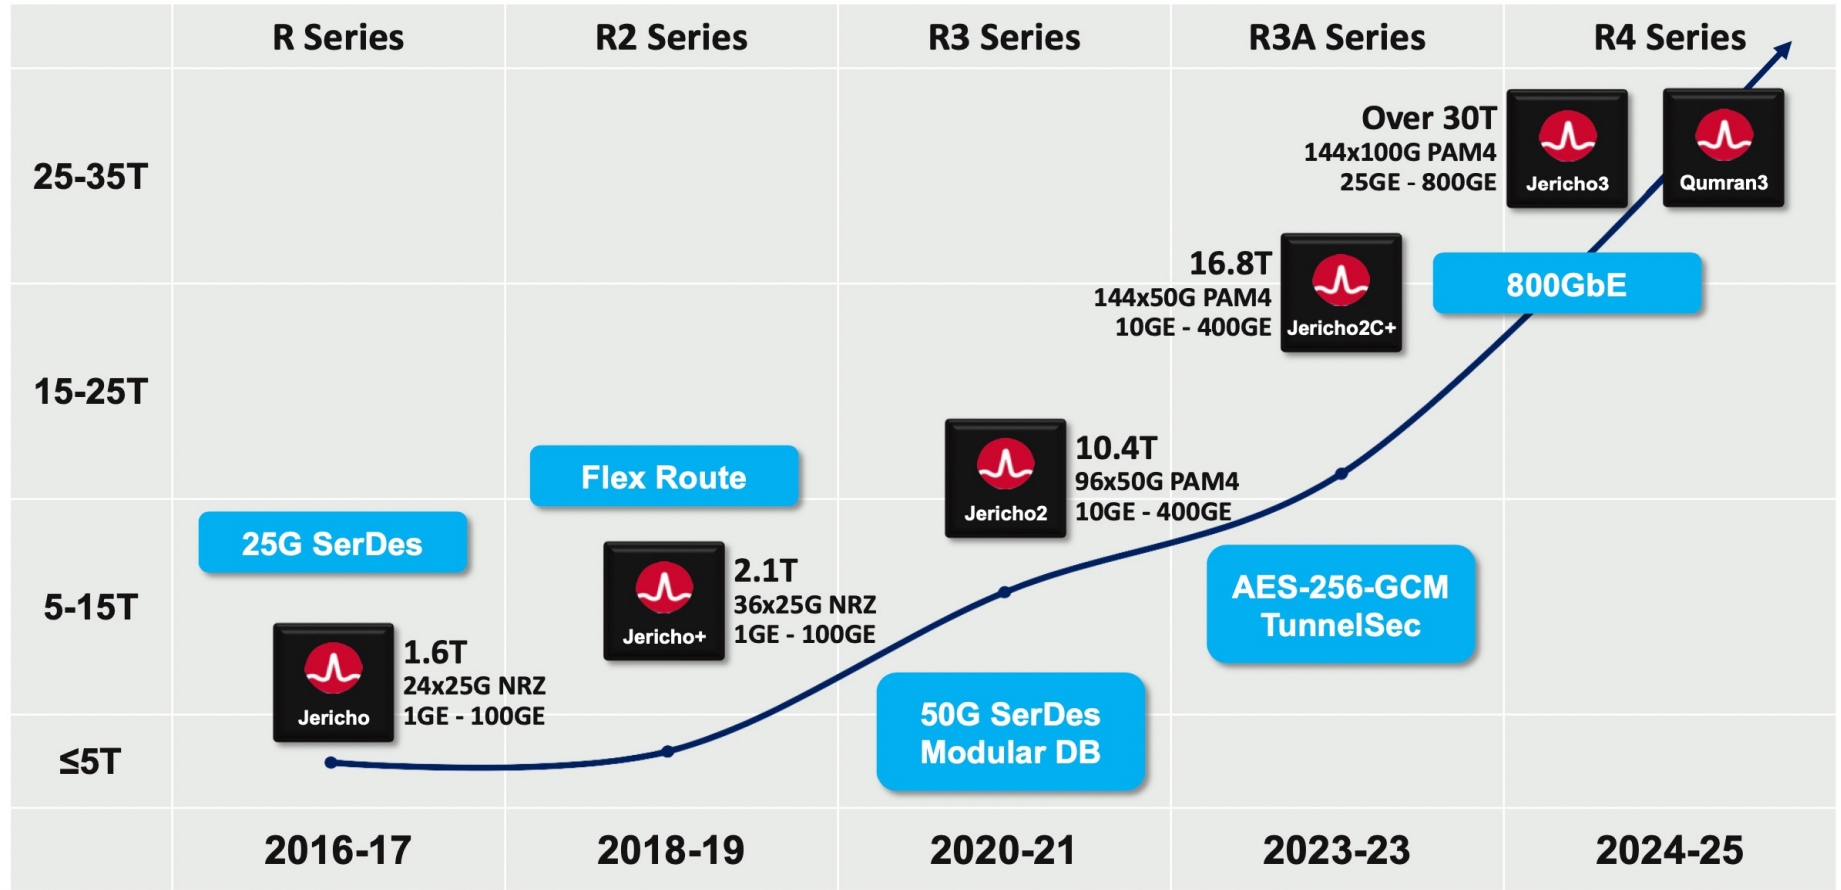


**Merchant Silicon: Faster Time-to-Market, Better Execution and Faster Innovation**


# Arista 7050X and 7060X Evolutions



# R Series Data Center Portfolio



# 5 Generations of 7050 Systems



# Arista 7050X4 Series 32 Ports 400G

## 1U High Density Fixed 400G Leaf

- Server rack optimized front to rear airflow
- High performance servers – 25G to 400G
- Full range of cables and optics:
  - 400G ports up to 16W per port
  - 4x100G and 2 x 200G modes on all ports
- Comprehensive L2 / L3 with forwarding enhancements
  - OSPF, BGP, Multicast & MLAG
  - Support for over 800K routes, 128-way ECMP, 64 way MLAG
  - Flexible and programmable pipeline
  - EVPN/VXLAN Routing in hardware
  - Rich instrumentation and in-band streaming telemetry
- Hot-swap / redundant power supplies
  - AC and DC Power options
  - Front to rear and rear to front
- Three hot-swap / redundant fan modules



32 400G Ports – 12.8Tbps

DCS-7050DX4-32S

DCS-7050PX4-32S

Agile deployment of 100/200/400G

- Flexible speeds in any OSFP/DD port
- Mix and Match speeds

# Next Generation 7050X4 Series – 8T Systems


## Flexible 100 and 400G Systems with Large Resources and Programmable Profiles

- 48 port 100G **SFP-DD** or **DSFP** and 8 port 400G in **1RU**
- Optimized for Next Generation PCIe4 Server
- 100G based on **2x50G** PAM4 lanes
- Consistent Trident Architecture


- 
- 48 port 100G **QSFP28** and 8 port 400G in **2RU**
  - Optimized for 50/100G NRZ Server Migration
  - 100G based on **4x25G** NRZ lanes
  - Consistent Trident Architecture
  - **MACSec** support on all ports

- 
- 40 port 200G **QSFP56** in 1RU
  - Optimized for Next Generation HPC 200G Server
  - 200G based on 4x50G PAM4 lanes
  - Consistent Trident Architecture

- 
- 24 port 200G QSFP56 and 8 port 400G in 1RU
  - Flexibility of native 200G & 400G ports
  - Optimized for Next Generation HPC 100/200G Server
  - Consistent Trident Architecture



7050SD/PX4-48D8



7050CX4M-48D8



7050CX4-40D



7050CX4-24D8



# Arista 7358X4 Series 100G/400G

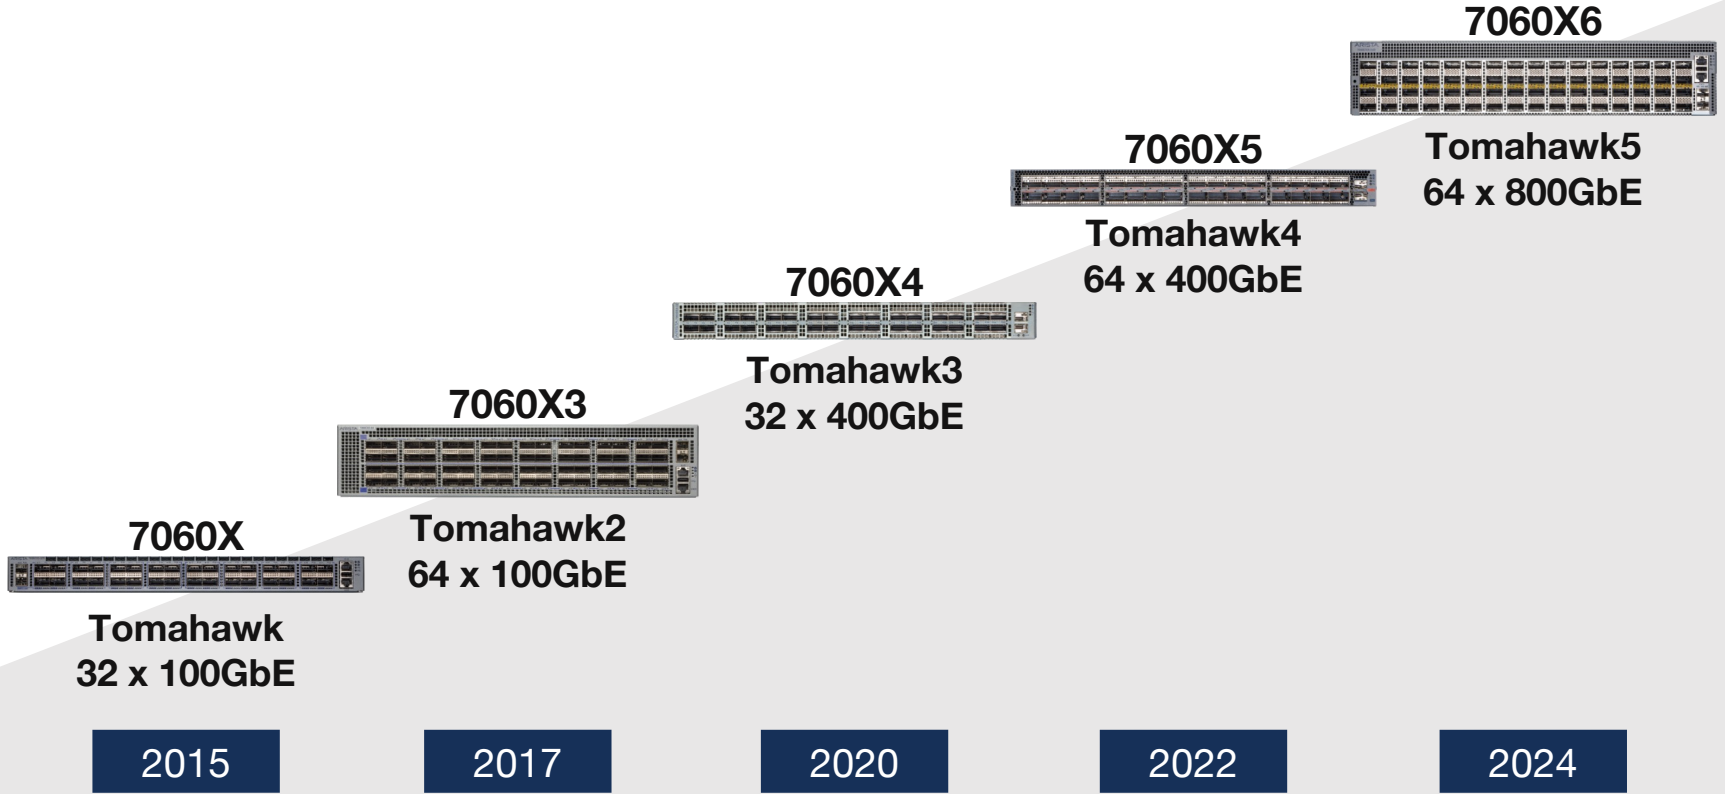
## 100/400G High Performance Modular System

- 7358X4-SC – Trident4 for 73xx Series (12.8T)
  - New switch card carries the Trident-4 switch chip
  - 132MB buffer, flexible scale and VXLAN
  - High density - 128 x 100G or 32 x 400G
  - Flexible 25G, 50G, 100G, 200G and 400G options



Common to 7368X4

# 5 Generations of 7060 Systems



# Next Generation 7060X5 Series Expanded Portfolio

High Density 100/200/400/800G Systems Optimized for Hyperscale Cloud

32 port **800G\*** (2x400G) OSFP or QSFP-DD in **1RU**

- 800G based on **256x 100G** PAM4 Serdes
- Each 800G port can support 2x400G based optics

64 port 400G QSFP-DD in 2RU

- 400G based on **512x 50G** PAM4 Serdes
- Consistent 7060X Series Architecture with addition of VXLAN
- Optimized for Hyperscale Cloud & Deep Learning AI Clusters

128 port 200G QSFP56 or 64 port 400G in 4RU

- Cloud Optimized – High Network Radix Modular Design
- 200/400G based on **512x 50G** PAM4 Serdes
- Flexible – 8 Hot Swap Removable IO Modules
- 16x 200G or 8x 400G MACsec or 8x 400G Non-MACsec
- Flexible speed and port breakouts – 256 100G ports

32 port 400G QSFP-DD in 1RU

- 400G based on **256x 50G** PAM4 Serdes
- Each 400G port can support 4x breakouts to 10/25/50/100G
- Power optimized with support for VXLAN



7060DX5-64 & 7060PX5-64



7060DX5-64S



7388X5



7060DX5-32

# Single Chip TH4: 7388X5 - 128x200G

- High Network Radix Modular System
- High performance 25.6Tbps switch card
- Removable Management Module
- Hot swap Fans and Power
- Choice of port module configurations
  - 8 vertical module slots:
  - 16 x 100G QSFP / 200G QSFP56
  - 8 x 400G QSFP-DD with MACSec
- Max Power Measured
  - Under 1.75KW for 128x200G
  - Under 3.5KW for 64x400G MACsec
- 4U Optimized for DC



# Arista 7280R3A Compact Modular System Overview

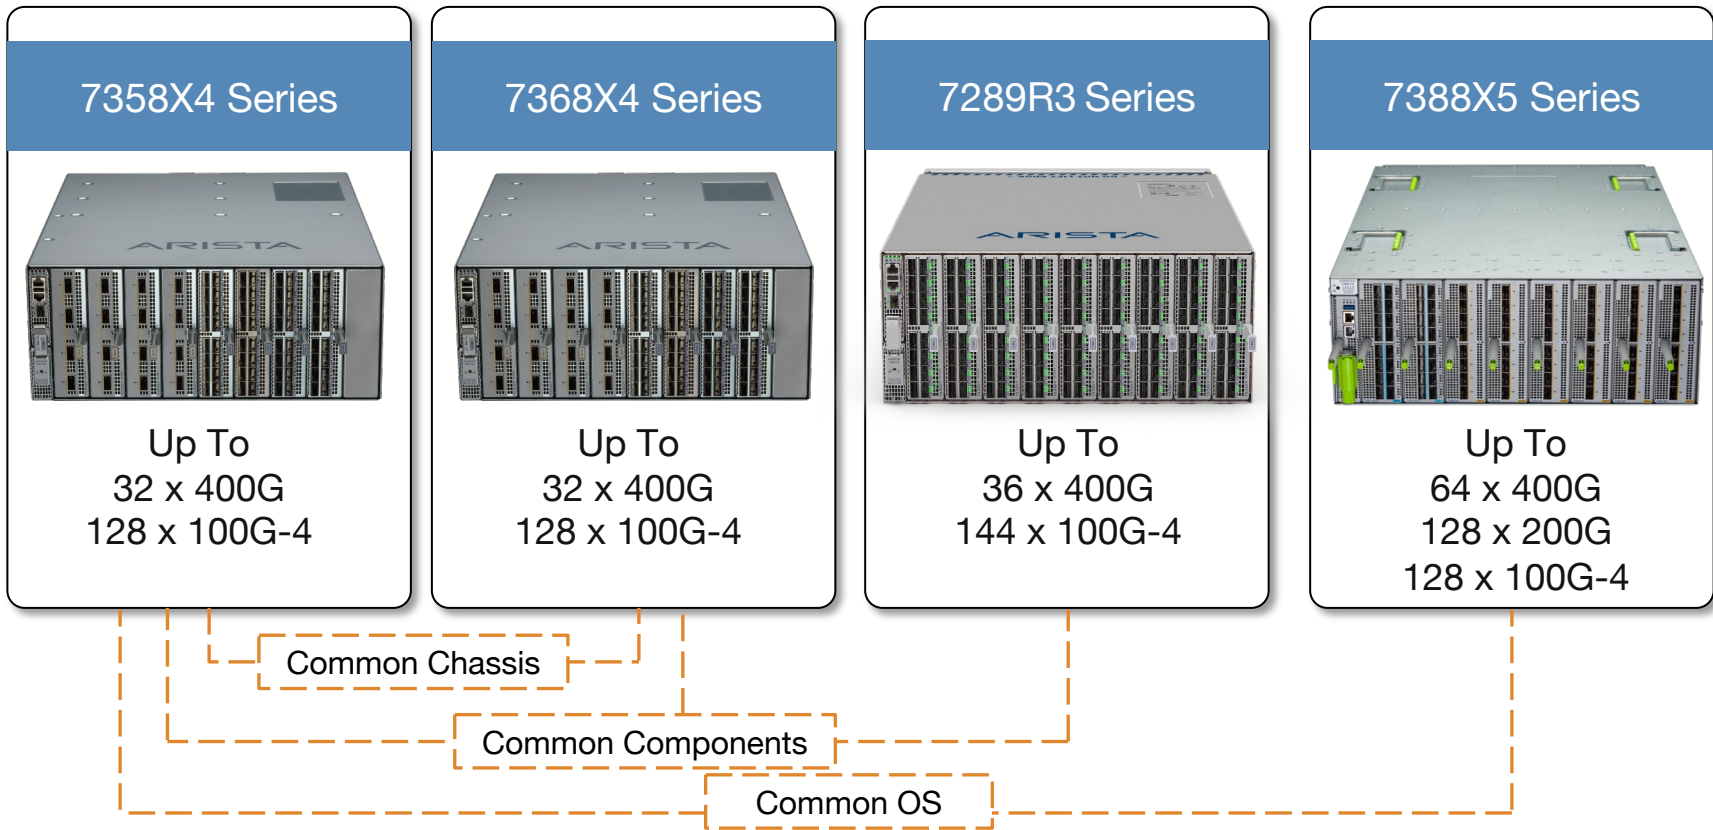
## 100/400G High Performance Modular System

- High Performance 100G/400G 2x J2C+ System
  - High Performance with 14.4Tbps and 5.6Bpps
  - Up to 36 x 400G or 144 x 100G
  - 16 GB Deep Packet Buffers
- Eight core x86 CPU
- 64GB System Memory
- 240GB M.2 SSD
- Console, Management and USB ports
- 5 Hot swap Fan Modules
- Up to 4 hot-swap / redundant power supplies
  - Front to rear airflow
  - HxWxD: 4U x 17.4" x 27.1" (17.9 x 44.2 x 68.83 cm)
  - Typical/Max Power: 1244 / 1927 W

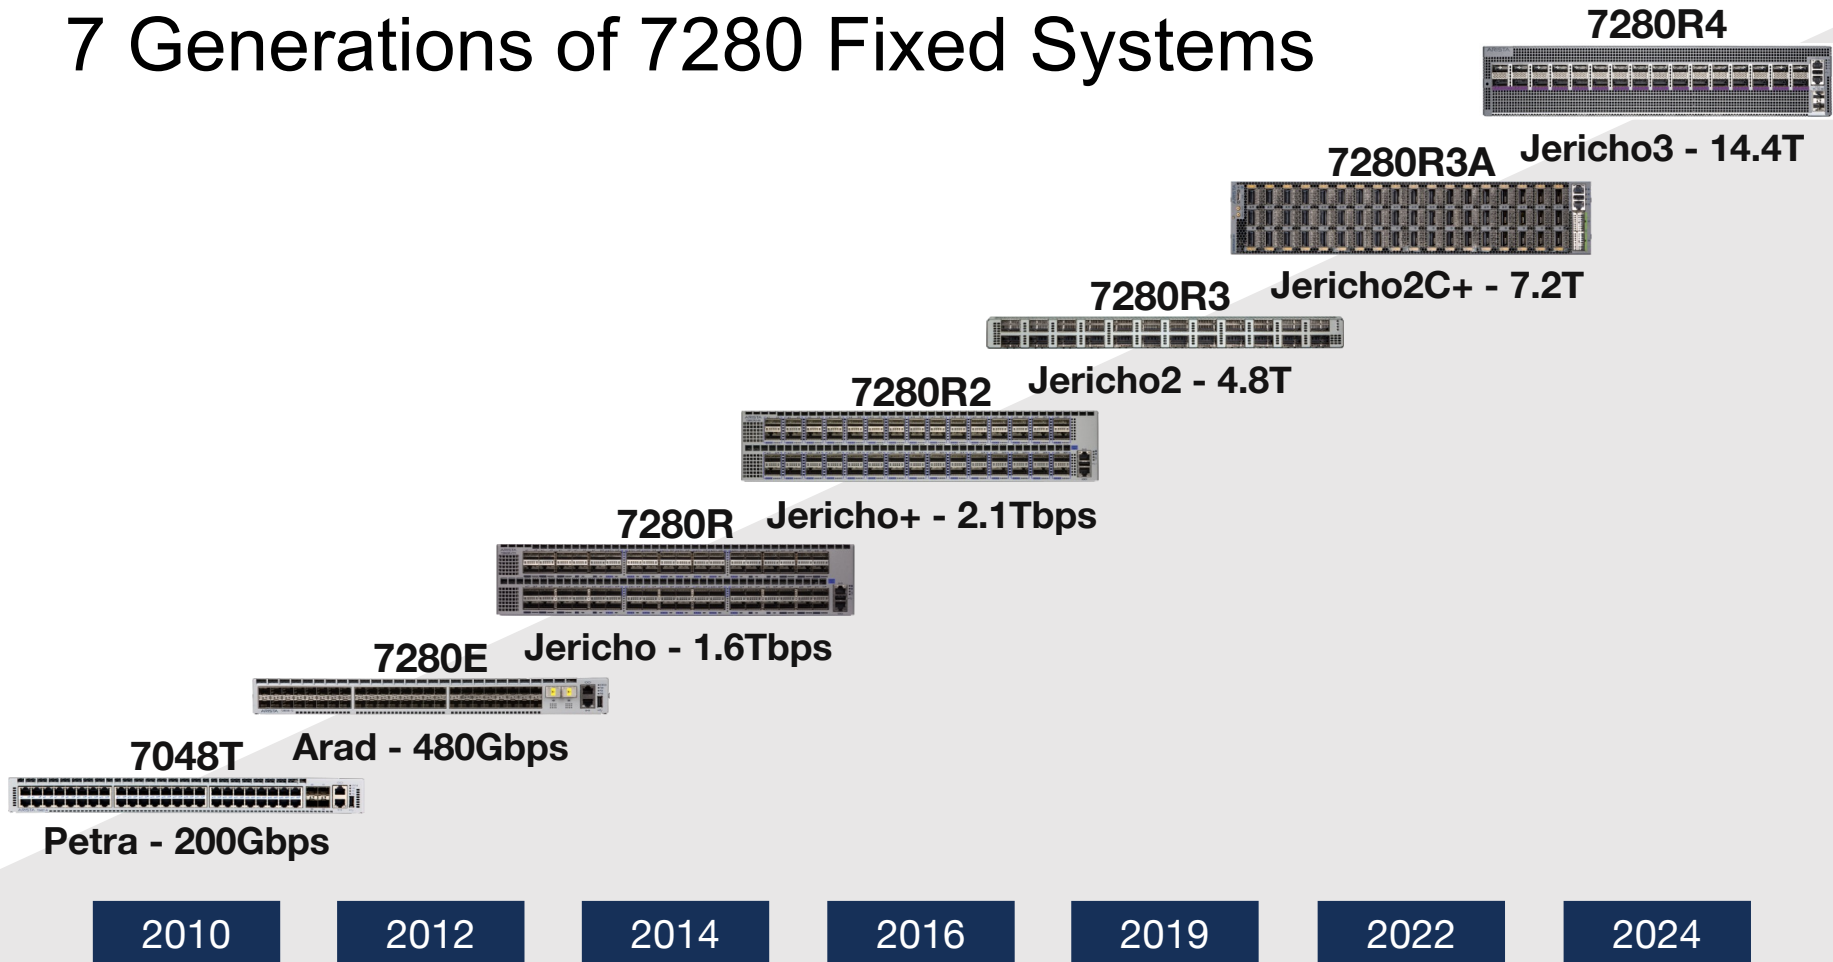


Consistent certification, knowledge, sparring, and architecture

# Streamlining Manufacturing – Common Components



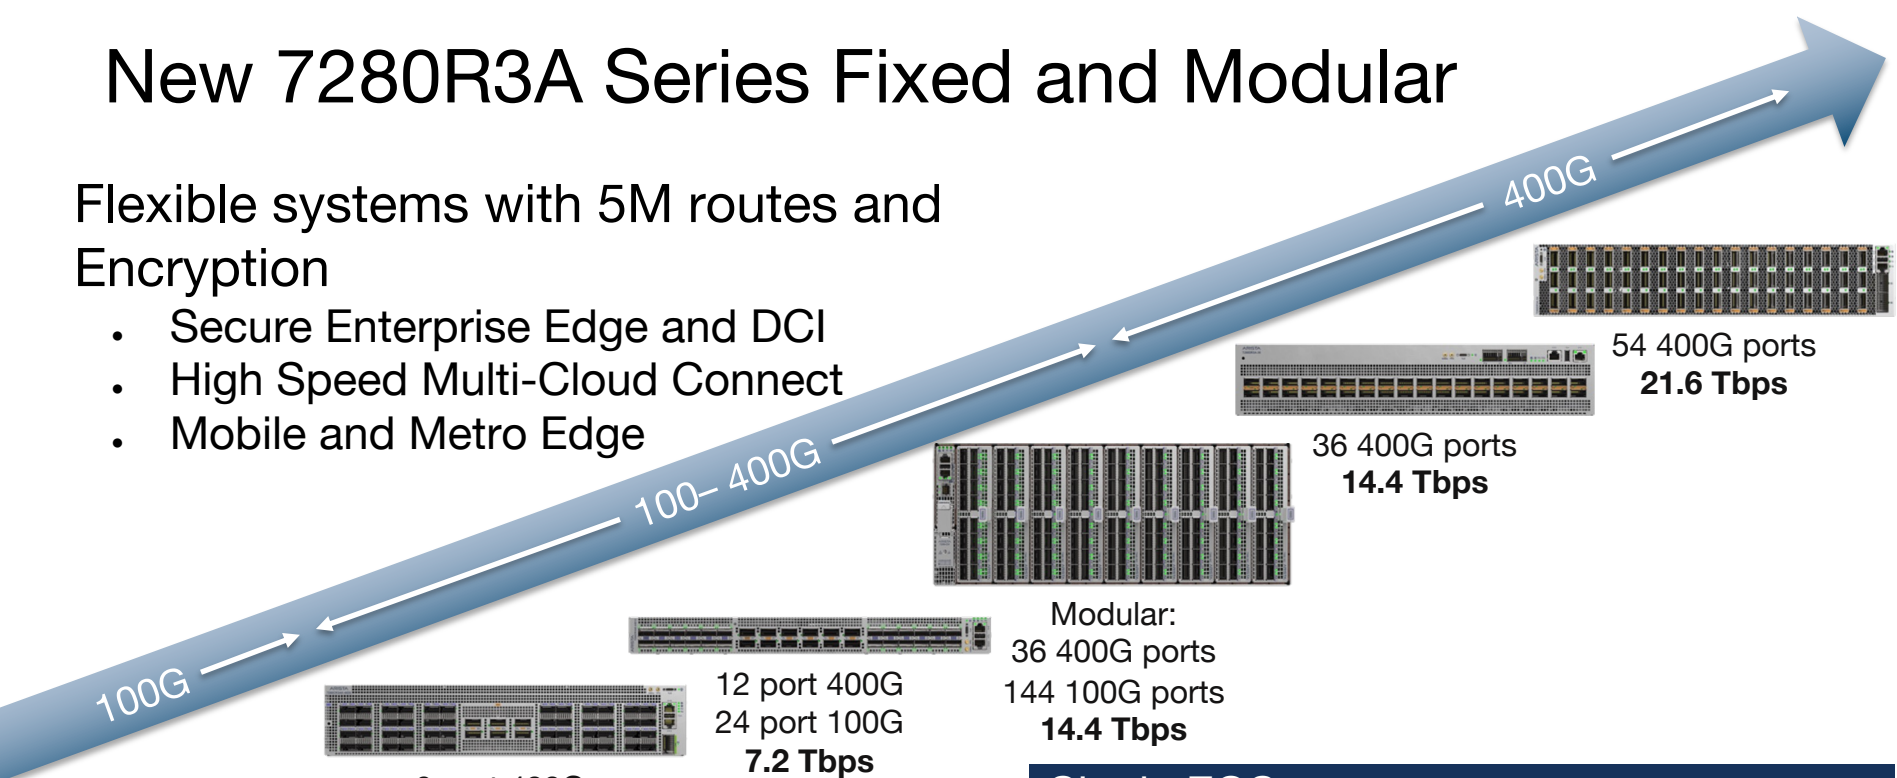
# 7 Generations of 7280 Fixed Systems



# New 7280R3A Series Fixed and Modular

Flexible systems with 5M routes and Encryption

- Secure Enterprise Edge and DCI
- High Speed Multi-Cloud Connect
- Mobile and Metro Edge



Single EOS  
Rich and Consistent R3 Features  
Integrated Arista TunnelSec™  
Base Encryption and Scale Options



# 7280R3A Series Fixed 100G/400G Switches

## Wire Speed 100/400G with Deep Buffers

### High Performance:

- Up to 54 x 400G wire speed ports
- Non-blocking up to 21.6 Tbps and 8.1 Bpps
- FlexRoute™ - 2 / 5 Million+ IP 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, Network Telemetry and Accelerated sFlow
- EVPN, MPLS, Segment Routing
- Wire speed AES-256 MACsec, IPsec and TunnelSec

### Cloud and Carrier Grade Networking:

- Dense 100G and 400G for SP, Cloud, Internet, HPC & CDN
- Dense 400G ZR and ZR+, Integrated Line-System Ports
- DC Optimized airflow and AC / DC power



7280DR3A-54 (Standard / Encryption / Scale)

3x J2C+  
54 400G  
21.6 Tbps



7280DR3A-36 (Standard / Encryption / Scale)

2x J2C+  
36 400G  
14.4 Tbps



7280CR3A-24D12 (Standard / Encryption / Scale)

1x J2C+  
24 100G + 12 400G  
7.2 Tbps



7280CR3A-48D6 (Standard / Encryption / Scale)

1x J2C+  
48 100G + 6 400G  
7.2 Tbps



7280CR3A-72 (Standard / Encryption / Scale)

1x J2C+  
72x 100G  
7.2 Tbps

# New 7800R3A – Systems and Line Cards

Cloud-Grade 460 Tbps with Encryption and 5M routes

- Secure Enterprise Edge, DCI and Peering
- High Speed Multi-Cloud Connect
- Service Provider, Mobile and Metro Edge
- 800G Ready




**7804R3**  
115 Tbps



**7808R3**  
230 Tbps

+50%



**7812R3**  
345 Tbps

+33%



**7816R3**  
460 Tbps




**R3A**  
Line Cards


>50%  
Power  
Reduction

Single EOS - Consistent R3 Features – Arista TunnelSec™


# 7800R3 Family 10G – 400G Line Cards




36 x 400G  
QSFP-DD



36 x 400G  
OSFP



48 x 100G  
QSFP




72 Port SFP28/56

Comprehensive Range of Line Cards for 10G-400G Deployments


# 7800R3 Family 10G – 400G Line Cards




36 x 400G  
QSFP-DD



36 x 400G  
OSFP



48 x 100G  
QSFP



72 Port SFP28/56

Comprehensive Range of Line Cards for 10G-400G Deployments

# ARISTA

APAC Cloud Builders 2023

Q&A



# ARISTA

Thank You!

We will see you in 2024!

