

Introduction to Arista CHANGE in SP

Arkadiusz Gierdojc
Jaroslaw Grabowski
Arista SE

Arista Networks

15%,
27.8% for 100GE segment*
Market Share

10+ Million
Ports Shipped

4500+
Customers

\$1B+
Revenue

1 Operating System

*) IDC report from 9 June 2017



Agenda



Short introduction


CHANGE & Competitive advantage



Focused
Lean
Smart
Fast



Agenda




Architecture matters



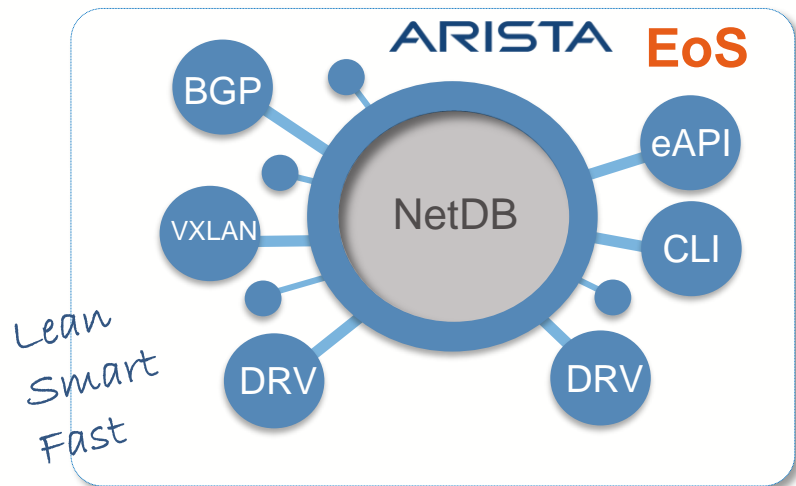
Algorithm /process



Data Structure /State




- Stateful Processes – All processes maintain their own state
- Enormously complex code interaction
- IPC with message queues and stores
- Impossible to coordinate development
- Serial Dependency (Christmas Lights)

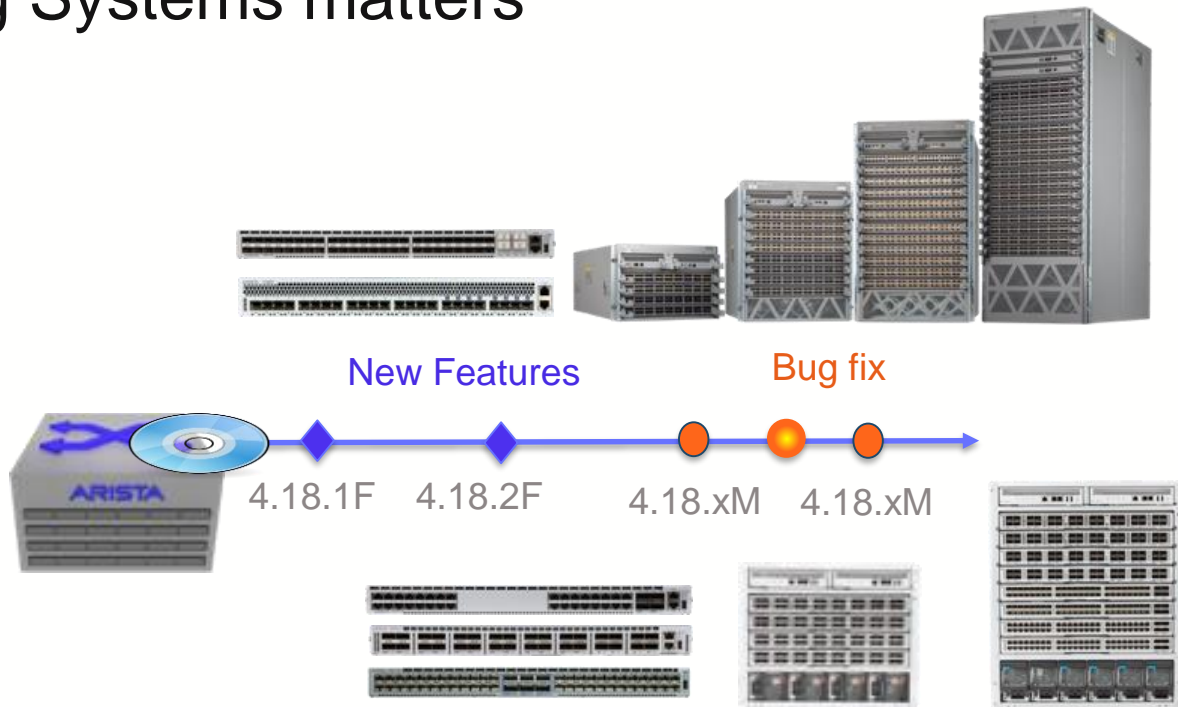


- Simple underlying architecture
- Fedora Core Linux Kernel
- Stateless Processes
- Publish/ Subscribe Model
- Self Healing

Number of Operating Systems matters



Multiple Operating Systems.
SW Image per device type
Complex feature and bug fix development



Single SW binary image
for all ARISTA products
Same CLI, API, bug fix, etc

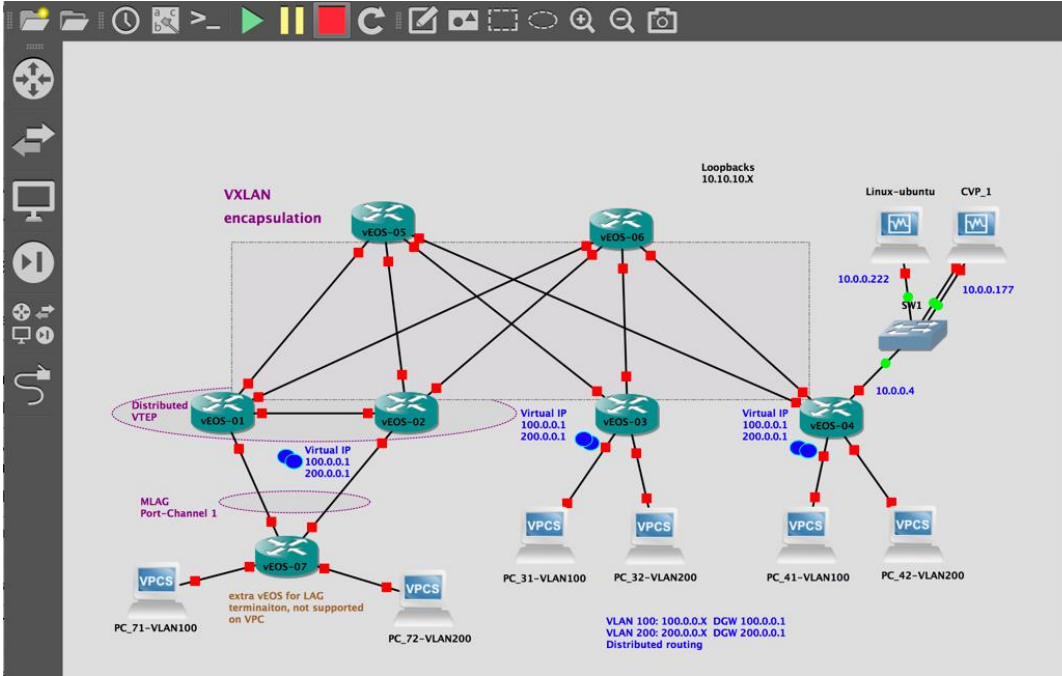
Software quality



20min Youtube video
„Arista Networks EOS Evolution and Quality with Ken Duda”
CTO of Arista Networks

*Quality
obsession...*

Free testing LAB DYI




Register at Arista

Download free
vEOS VM

Configure VMs or use free GNS3

L2, L3, VXLAN,
STP work!
Enjoy....

Final Note



Size matters...



45 RU

32 Tb/s



48 RU

16 Tb/s

ARISTA

7500R

7280R

Today



18 RU

57.6 Tb/s



2 RU

6 Tb/s

ARISTA

Plans for 2018



32x



4x100GE

1x400GE

100G-Lambda
Initiative

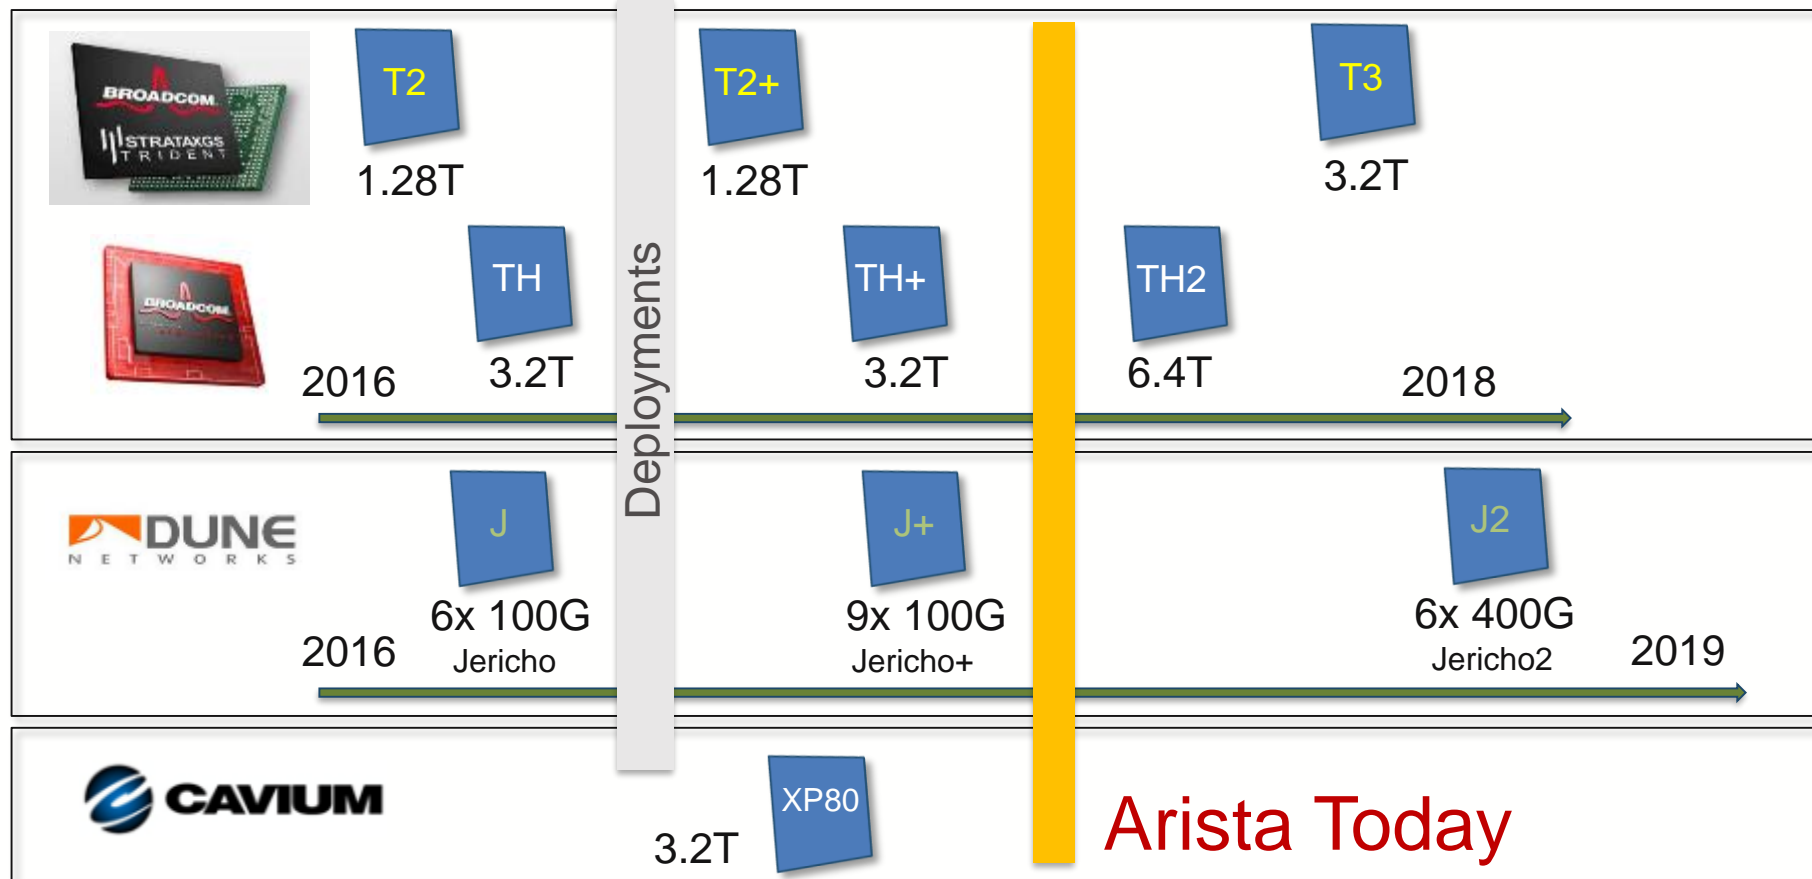
1 RU

12.8 Tb/s

High end routers: MPLS, Segment Routing, Internet routing,

Switching/Routing

Arista Silicon Landscape: 2016-2019



Arista Networks Portfolio

EOS: Extensible Operating System – Single Binary Image for all Platforms

EOS
CloudVision



7010

48-port Data Center
Class Gigabit
Ethernet Switch

7150S

Ultra Low Latency
24,52,64-port SFP+
1G-40GbE Switches

PTP High precisions
oscillator
Intelligent
Application Switch

7050X & 7200X

Dense Low Latency
32 & 64-port QSFP+
96xSFP+/8xQSFP+
48-port 10Gb w/
100Gb Uplinks

Advanced
Virtualization
Scale-out
Visibility



7280SE

Dense 48 Port 10G
switch with deep buffer
VOQ architecture with
choice of:
2 x 100G MXP
2 x QSFP100
4 x QSFP+

Advanced Virtualization
Scale-out
Visibility



7060CX

Dense Low Latency
32 & 64-port
CX32 – 32x40/100G,
64x50G or 128x10/25G
CX64 – 64x40/100G,
128x50G or
246x10/25G

Advanced Virtualization
Scale-out
Visibility



7280R

Deep buffer VOQ
architecture switch with
choice of:
48 x 10G(CU) + 6 x 100G,
48 x 10G(FX) + 6 x 100G,
36 x 40G (12 x 100G),
48 x 100G + 8 x 40G

All 100G can be used as
1 x 40G, 4 x 10G/25G or
2 x 50G

Advanced Virtualization
Scale-out
Visibility



7300X

High Density, Modular
System supporting up
to 512 40GbE

Cloud Scale
Leaf and Spine
10GbE-40GbE



7500R

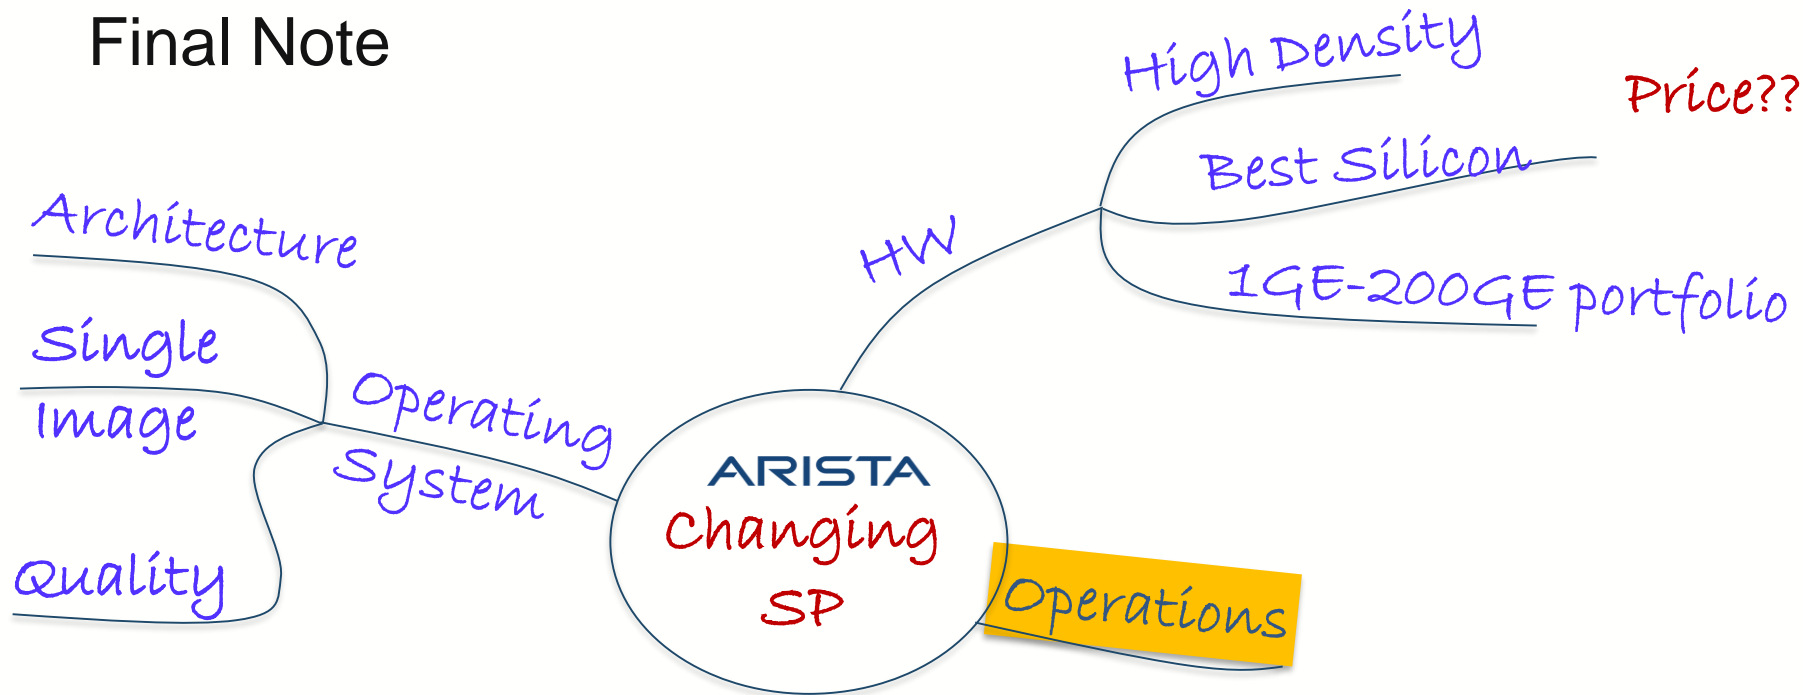
Lossless, High Density,
Modular Switching
System supporting up
to:

432 x 40G/100G ports
864 x 50G ports
1728 x 10G/25G ports
at wire speed.

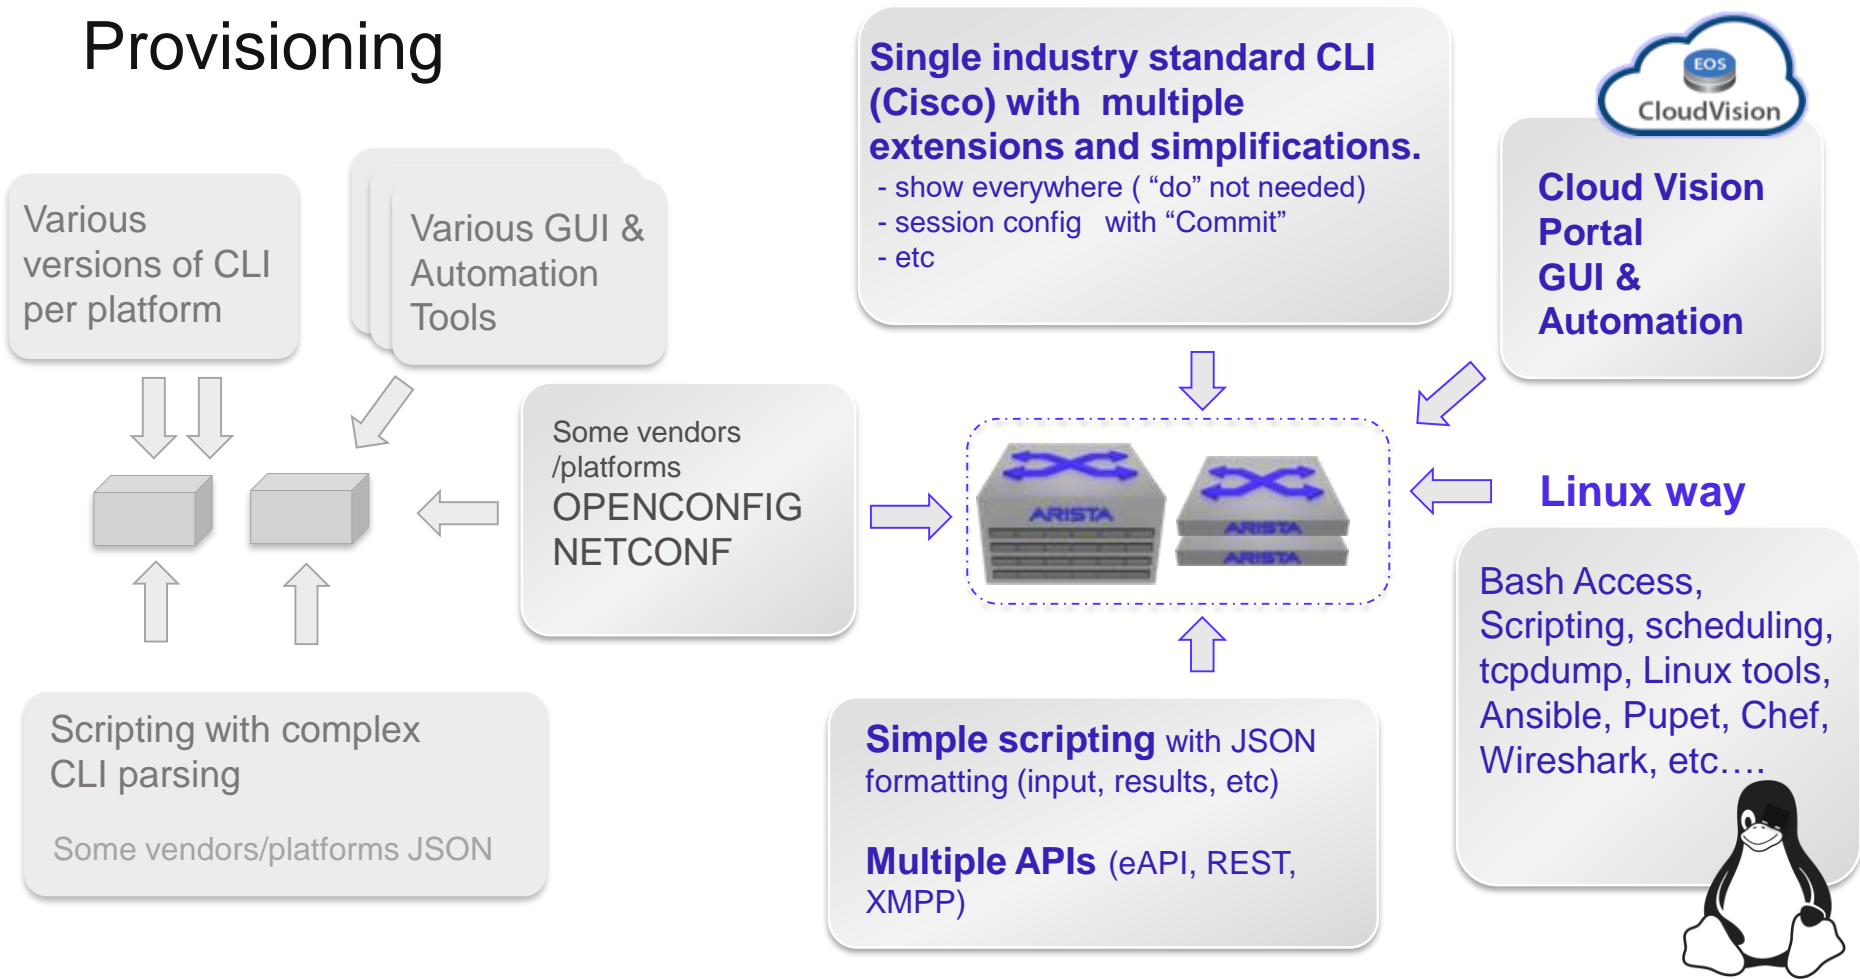
Spine Switch with Deep
buffer VOQ Architecture
10/25/40/50/100GbE
200G Coherent DWDM
with MACSec Encryption



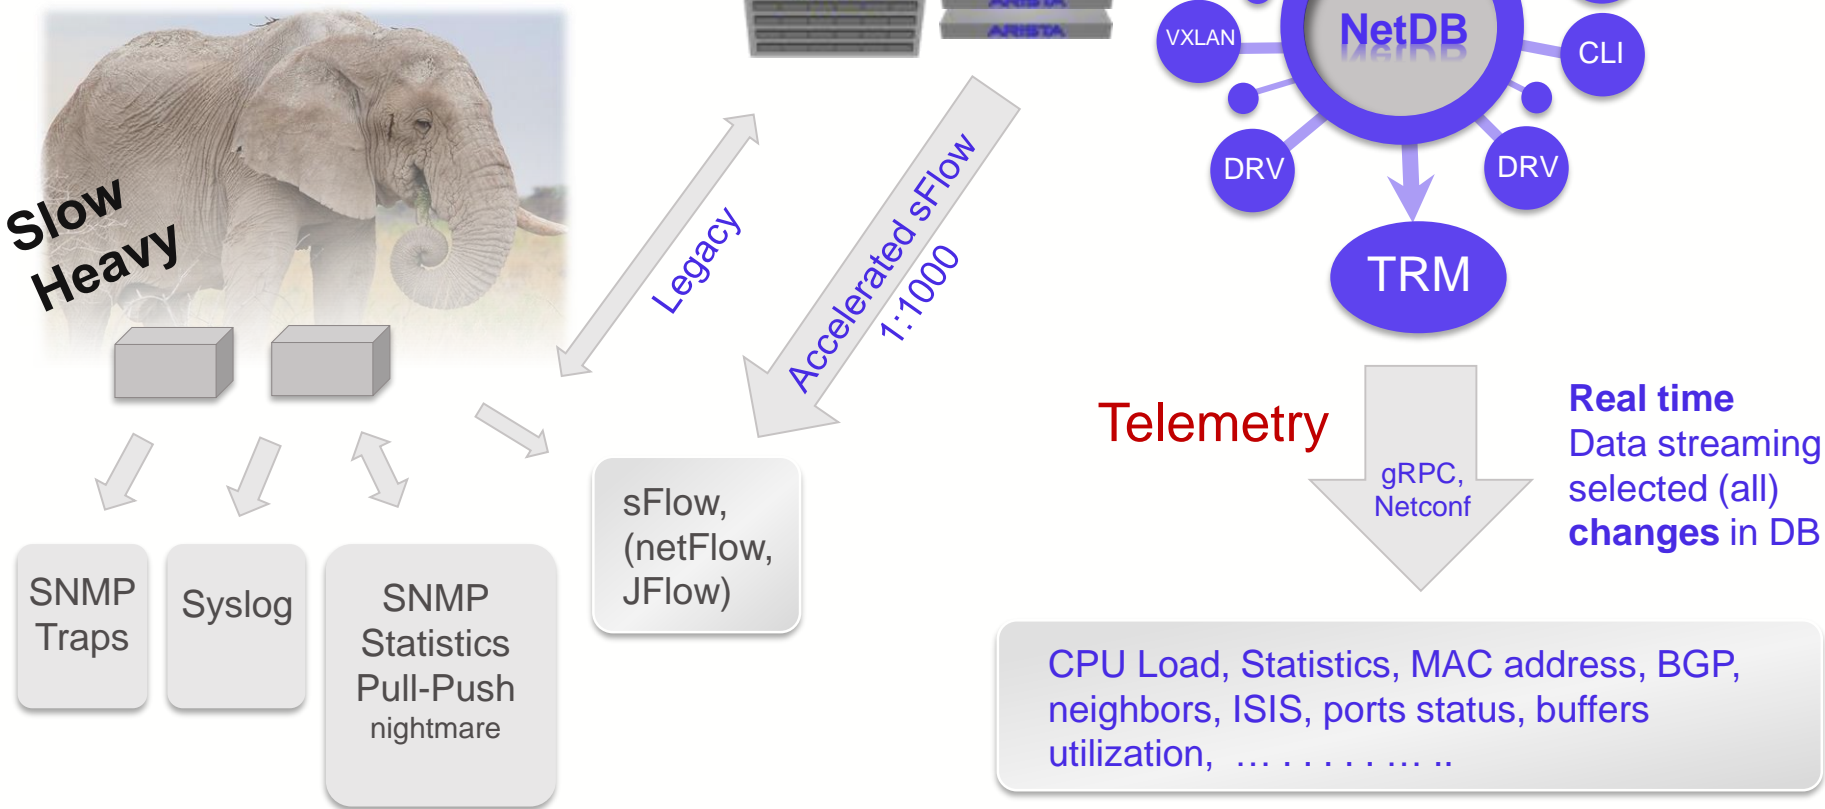
Final Note



Provisioning



Monitoring



Telemetry matters

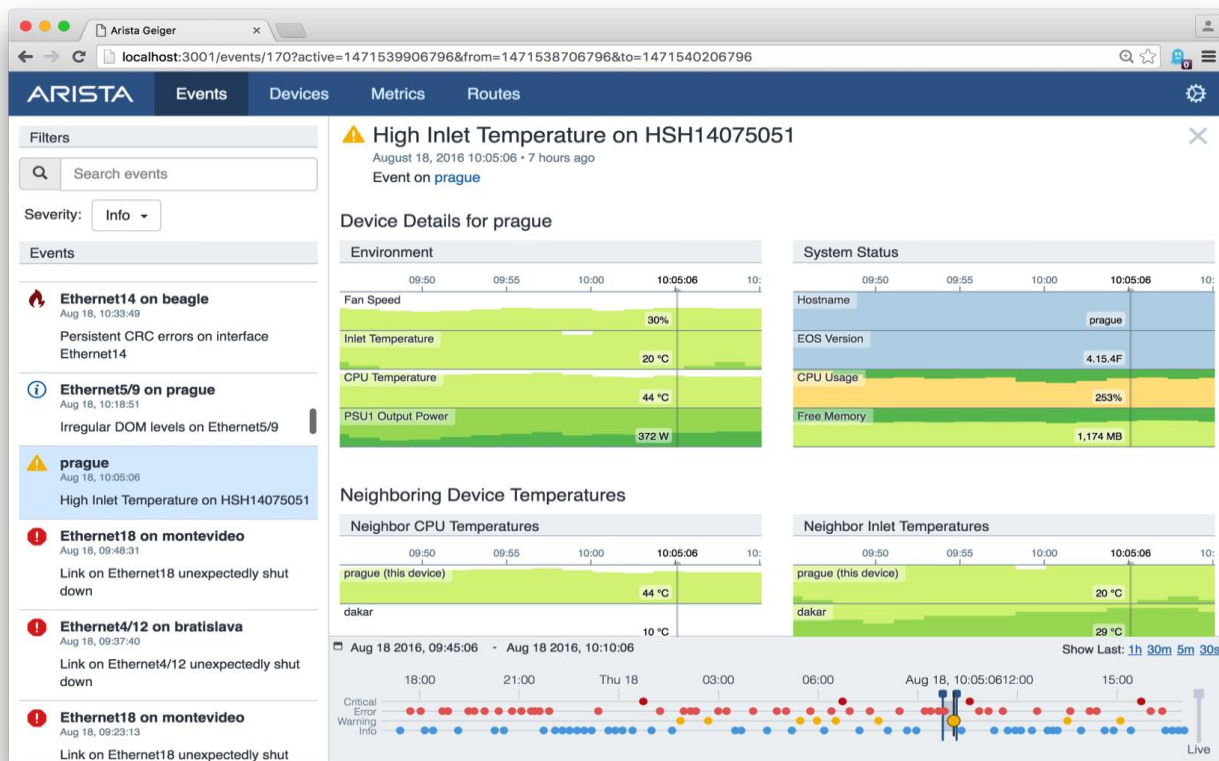
Select nodes

Select Metric

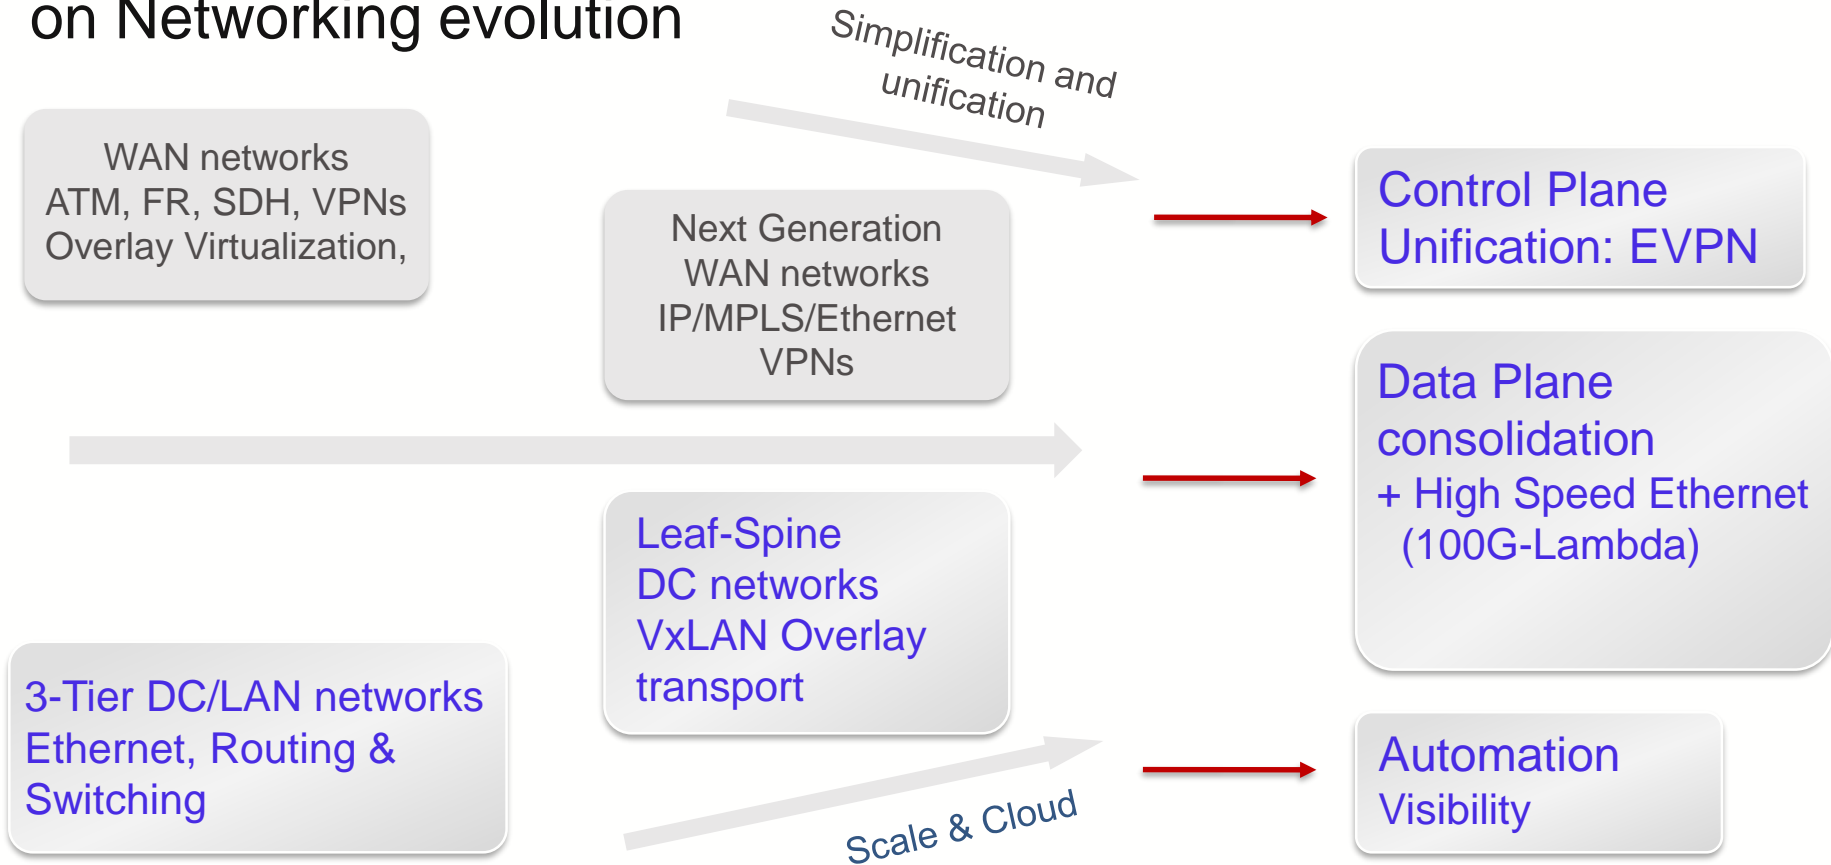
How can you answer when exactly specific MAC was learned?

Maybe 2 week ago in the middle of the night?

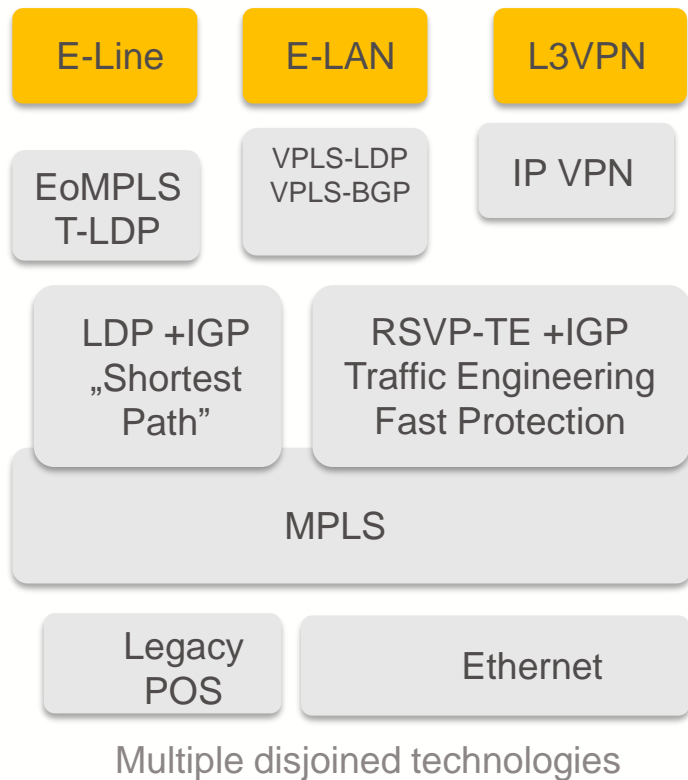
Select time range



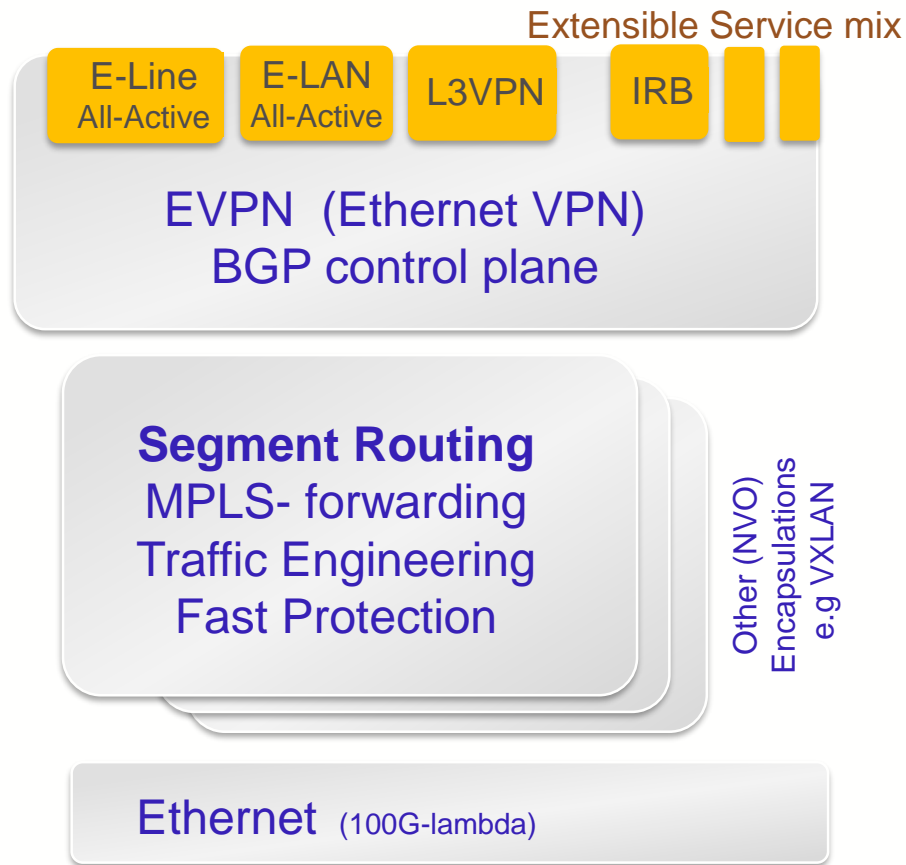
My personal view on Networking evolution



SP architecture change




New Generation Architecture



Segment routing in the nutshell

MPLS
Frame
Structure



MPLS forwarding

- Same paradigm
- Unique label per node
- IGP for routing and label distribution

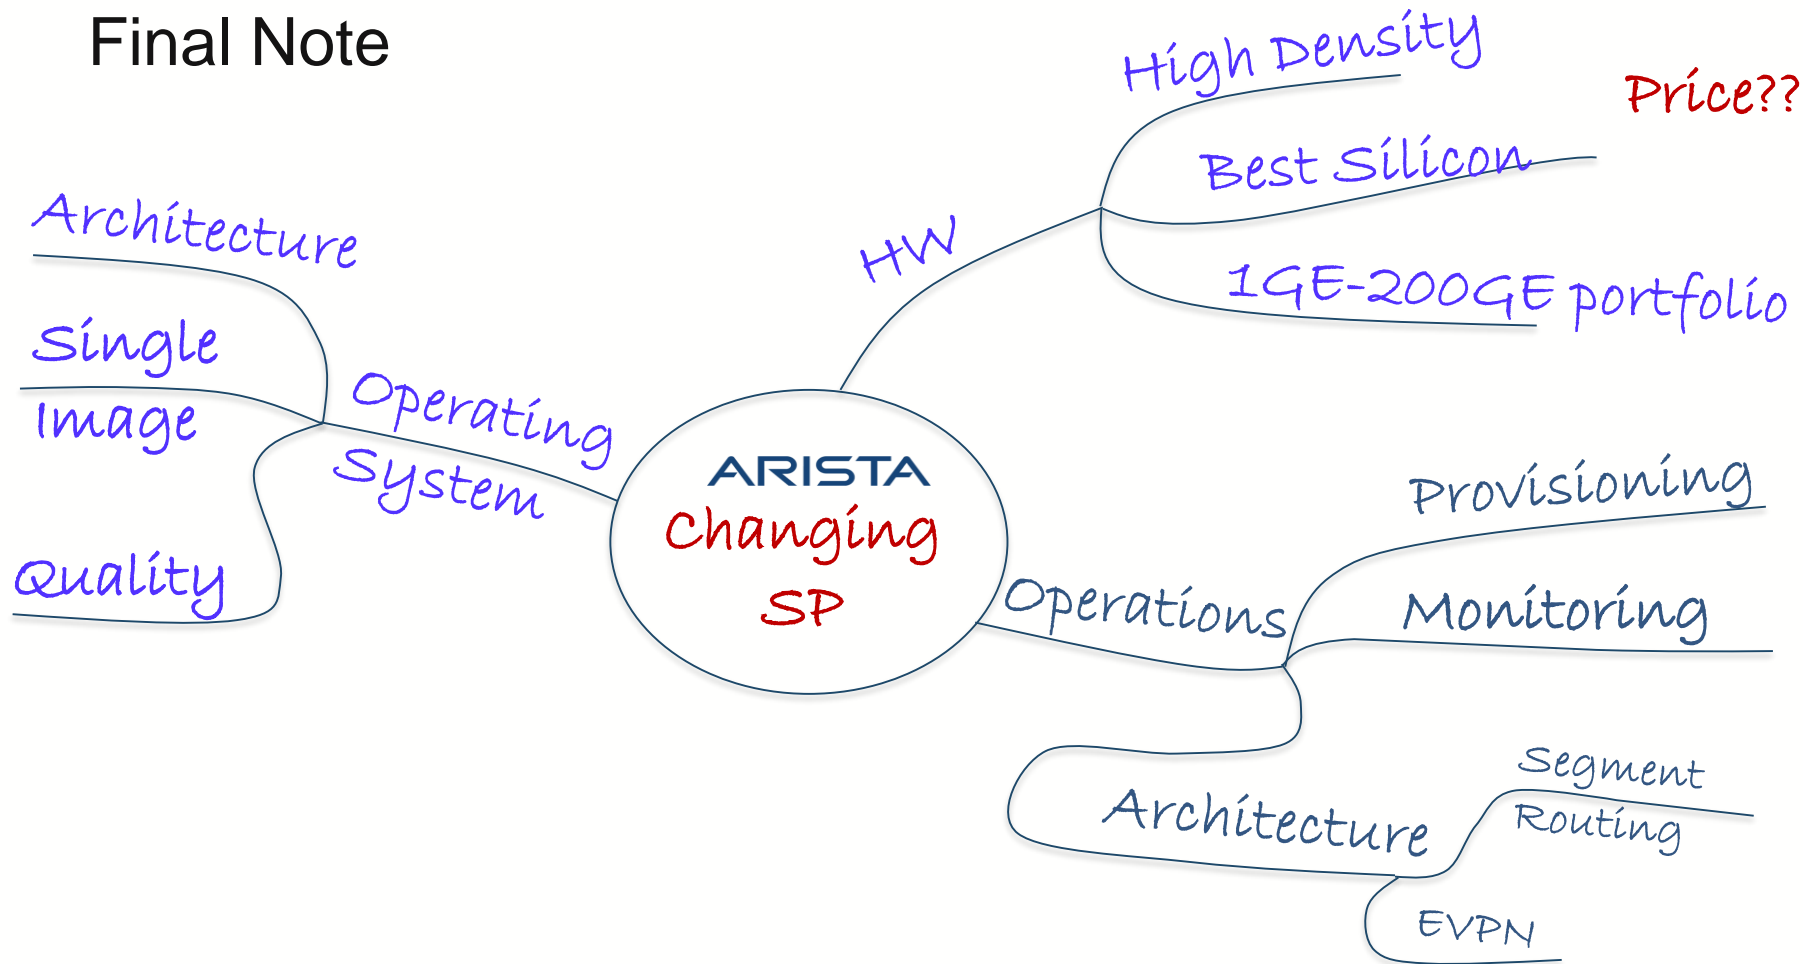
Traffic Engineering

- Programming through label stack
- No state kept by intermediate nodes (unlike RSVP-TE)

Fast protection

- Network calculates backup paths (TI-LFA)


Final Note




Use cases

DC Architectures


Spline™




Layer 2 / MLAG



Layer 3 / ECMP



L2 over Layer 3
VXLAN



Servers: 100 to 2,000

100 to 10,000

100 to 100,000+

100 to 100,000+

Customers

<p>Web</p> 	<p>Financial Services</p> 	<p>Service Provider / Cloud / Hosting</p> 	<p>High-Tech</p> 
<p>Government</p> 	<p>Media & Entertainment</p> 	<p>Education</p> 	<p>Oil and Gas</p> 
<p>Business Services</p> 	<p>Research Labs</p> 	<p>Healthcare</p> 	<p>Retail</p> 
			<p>Manufacturing</p> 

Facebook Luleå Datacenter




Arista 7300

“Look at these racks, the network devices, the cabling. Everything is like reference model!”
Max Zavyalov, Network Engineer in Edge & Network Services team
<https://www.facebook.com/zuck/posts/10103136694875121>

Internet peering

7280R – fixed, 1RU, 2RU
7500R – Modular Chassis



7280R



30x100GE (10G/40GE)



48x1/10/25GE + 6x100GE

BGP FIB for
multiple Internet
tables

1.3M (Standard)
2.0M (Extended)
5.0M (2018)

ECMP (Equal)
UCMP (Unequal)
Accelerated sFLOW
Deep Buffers

IXP L2 extensions

L2 emulation
Over VXLAN

ECMP

- All paths used
- No Spanning Tree


**Fast IP
convergence**

No VLAN limits

- Local VLAN
significance

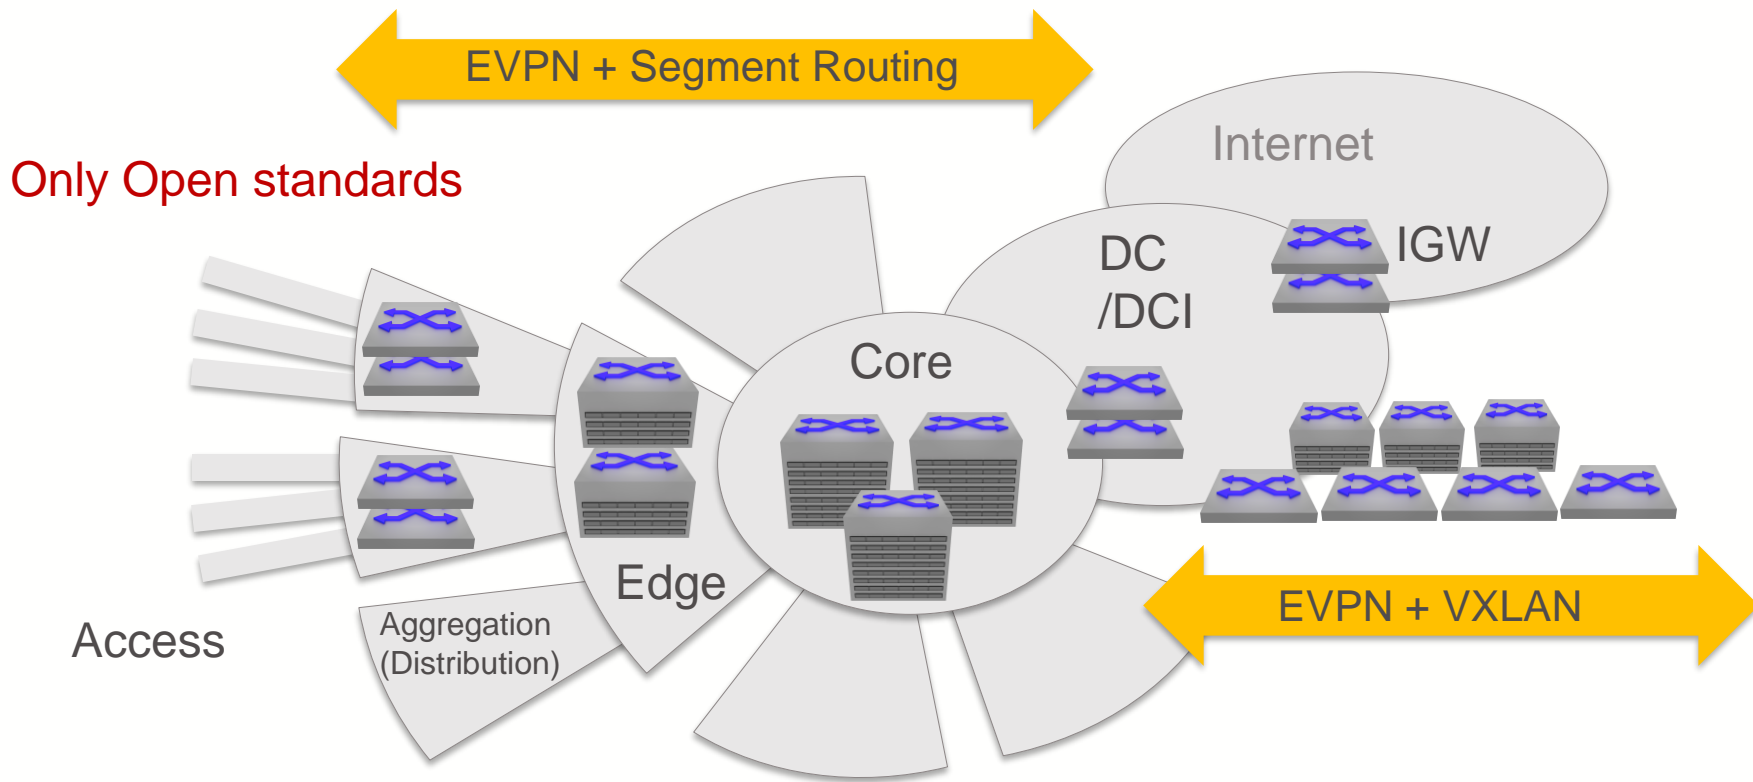
*No high end
routing needed!!*

*Lean and
fast*

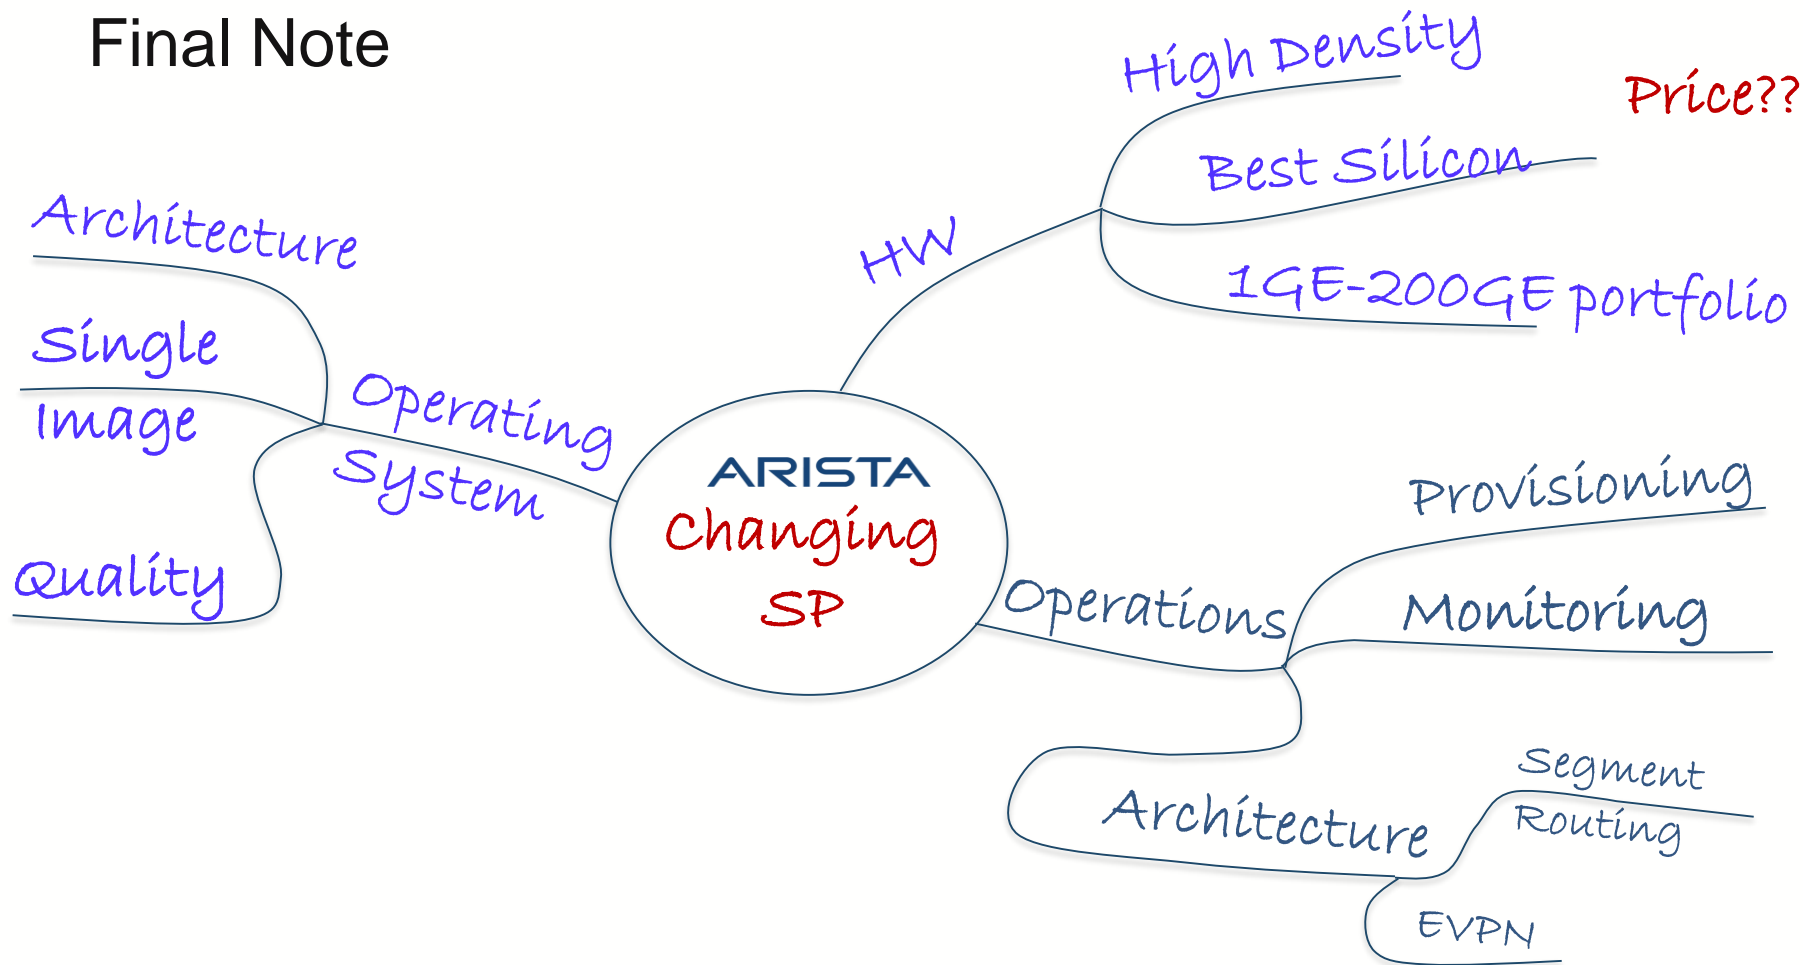


SP network coverage

Some features on roadmap but coming..



Final Note





Thank You

www.arista.com