

Cloudify the Enterprise with BYOC (Bring Your Own Controller)

김창민 수석부장

Korea SE Lead

charles.kim@arista.com

ARISTA

Gartner 2013 to 2014 Data Center Networking Magic Quadrant



2015: The Year of the Cloud




The BVP Cloudscape

Top 300 Privately Held Cloud Companies

Business Users

SaaS

Marketing	Sales	Service & Support	Finance	HR	Vertical
action HubsSpot bizo AGILONE gigya ensighten main street hub Optimizely BRIGHT EDGE INVESTARS CLICOTALE™ Infusionsoft KENBIHQ kapost SALTTRU yodle Owlter RAVEN GINZA Metrics CORALTE unbounce hootsuite Simply Measured moz jonrain DELIVERYAGENT SpreadFast DYNAMIC YIELD TEALIUM sprinklr infer Zuberance MailChimp doubledutch	clearslide insidesales.com xactly toutapp ConnectAndSell clari hearSay social hoopla Q RelateIQ WorKinBox nimble SW Steelwedge instightly InsideView Yesware BOXER SUGARCRM pipedrive selligy CIRBUSPATH	Intercom™ TOA NEWVOICEMEDIA gainsight Lithium TOTANGO liveops MEDALLIA Preact BLUENOSE satisfaction serviceMAX uservice	Adaptive Insights zenpayroll anaplan FRESHBOOKS aria hostanalytics kyriba Avalara Chargify Intacct Bill.com Recurly coupa tidemark SAASU zuora EXPENSIFY wave FINANCIAL FORCE.COM	HireVue entelo SmartRecruiters Parklet LEVER ZipRecruiter greenhouse evolv bamboohr workable Jobvite sumtotal SilkRoad PEOPLE MATTER REPLICON jobscience Resumator Avature Maxwell Health.	Healthcare DocuTAP ClearCare Wellsco inLum omodo comprehend WebPT AMERICAN WELL KINNSER CoreCloud tigertext evariant PointClickCare GRAND ROUNDS medica PROPELLER HEALTH doximity naviCare practicefusion Real Estate HIGHTOWER Floors rentlytics YARDI PROCORE properly VIEW THE SPACE RAINMAKER ppfolio propertybase Geary Technologies
Collaboration box EVERNOTE 37signals Alfresco DESKAWAY DocuSign slack moxie Wrike TeamLab asana Dropbox clarizen Atlassian AtTask HIGHTAIL huddle Redbooth sendthisfile EGN*TE LiquidPlanner SugarSync Zoho WatchDox					Education KNEWTON INSTRUCTURE edemy BETTERLESSON Clever Verificient Desire2Learn edmodo topschool BrightBytes HistoryConnect U usacity dreambox edmentum schoolbag courseera
BI / Analytics DOMO Chartbeat FLURRY birst INSIGHTSQUARED UPSIGHT sumAll pentaho INFO data GoodData mixpanel Datemeer Localytics KISSmetrics kognitio Lattice bime PivotLink SpatialKey visier Rosslyn Analytics					



PaaS

IaaS

Developers

IT Ops

Security

twilio zapier SendGrid PRGRDUTY Skytap RIGHTSCALE Cloudius ravello sumologic IID CloudLock defense.net apprenda CrowdFlower Iron.io cloudera VictorOps APPERIAN boundary MarkLogic wandera mojave networks CipherCloud NITROUS PLCloud Acquia Hortonworks Bluejeans APPTIO cloudshare snapLogic MuleSoft simplified onelogin CLOUDFLARE WhiteHat CloudBees stripe Engine Yard CRITERICISM catchpoint Abiquo EUCALYPTUS Zerto GIGASPACE Vaultive RAPID7 VERACODE Lookout CloudPassage AUTOMATTIC DigitalOcean linode.com nebula SCALEXTREME thinking share networks SPICEWORKS MobileIron loggly ONECLOUD SilverSky okta
--

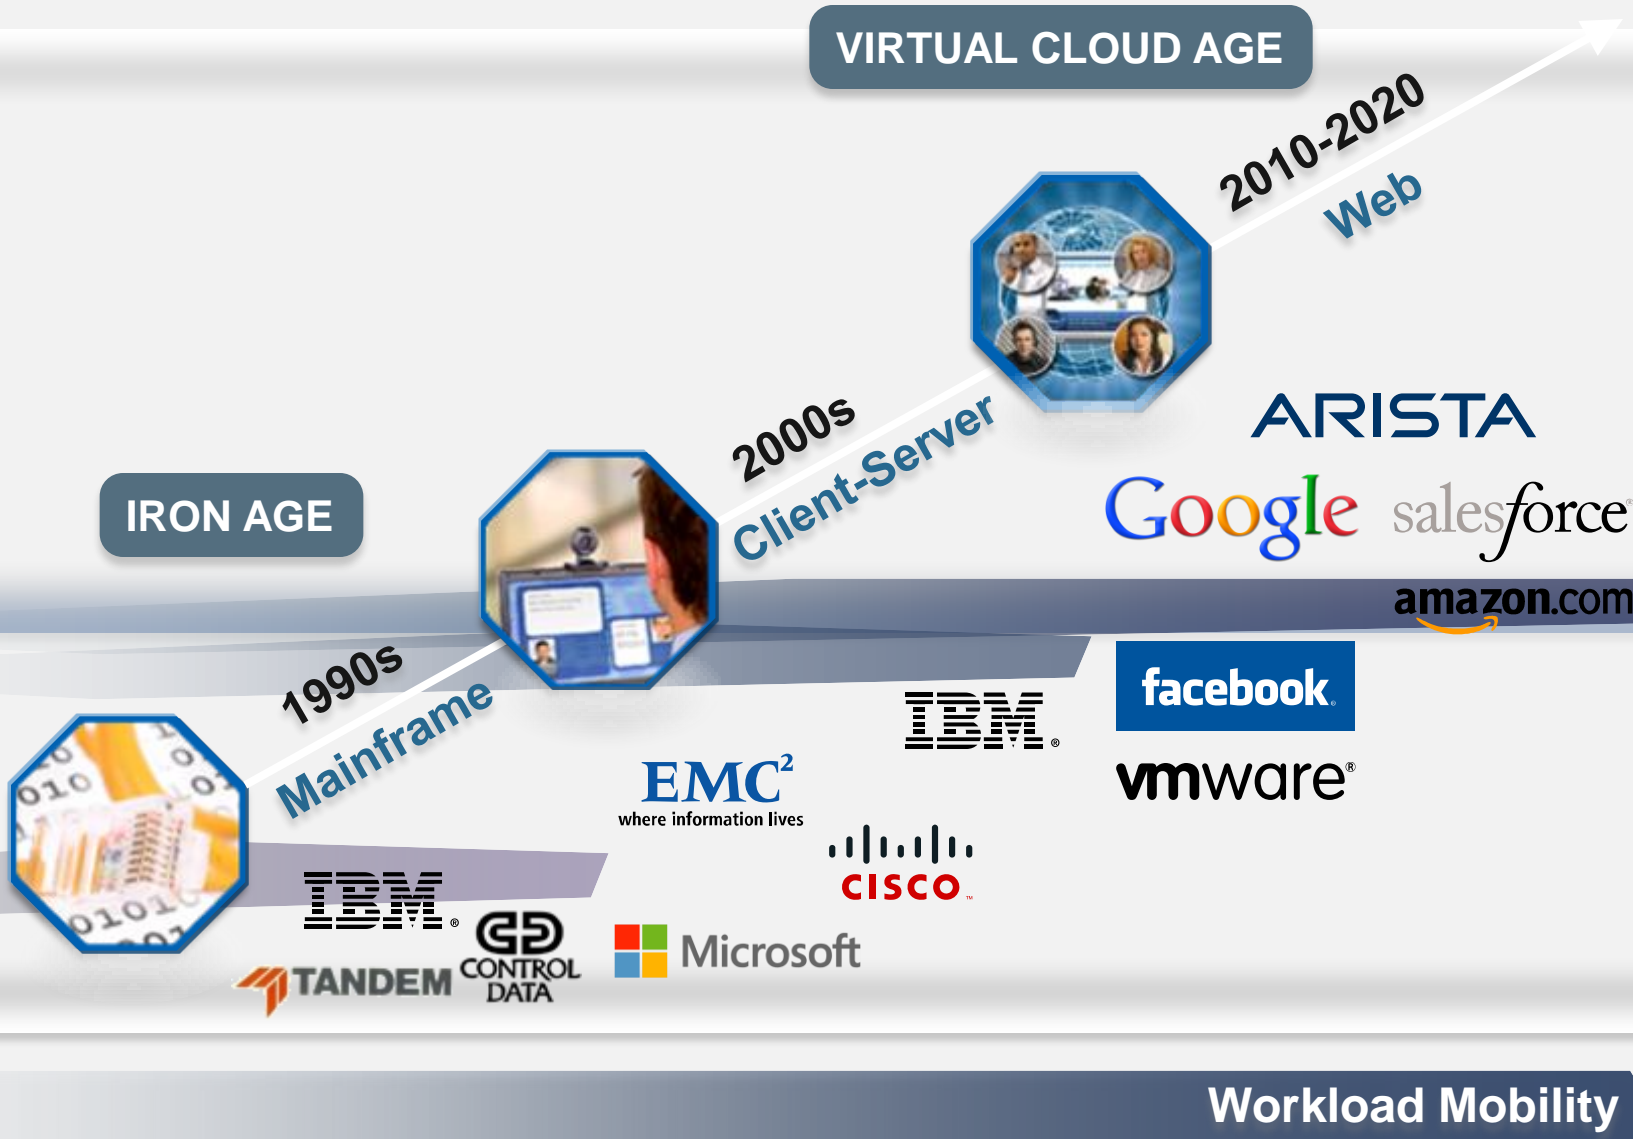
© Bessemer Venture Partners 2014 v4.0

Download a digital copy or nominate your company: bvp.com/cloud

The Next Disruptive Frontier in IT

ARISTA


Application Mobility



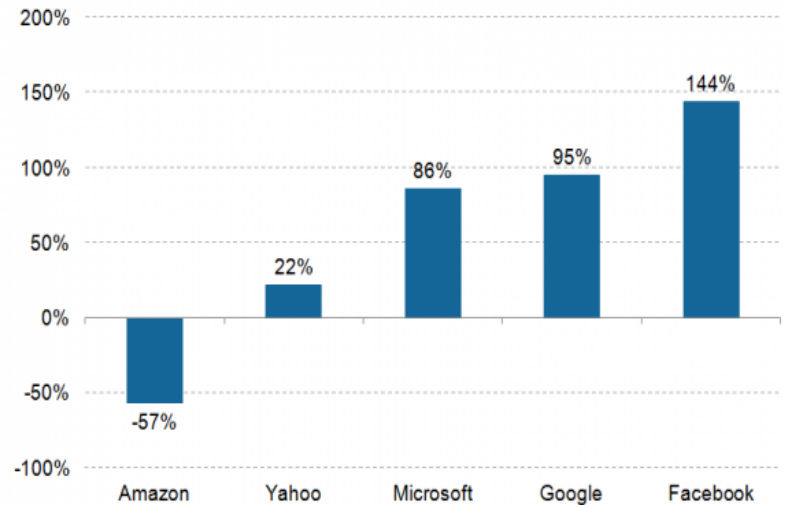
IDC Report for Worldwide Public IT Cloud Service and Cloud CAPEX Spend



Worldwide Public IT Cloud Services Spending by Segment (in \$ billions)




CapEx Spending Growth, Latest Quarter's YoY Growth



Market Realist

Source: Company Filings



Web/Cloud:


- Computing *is the business*
- ~25% of servers in 2014
- ~50% of servers in 2018(E)
- Small number of large-scale web/cloud service providers

Enterprise:

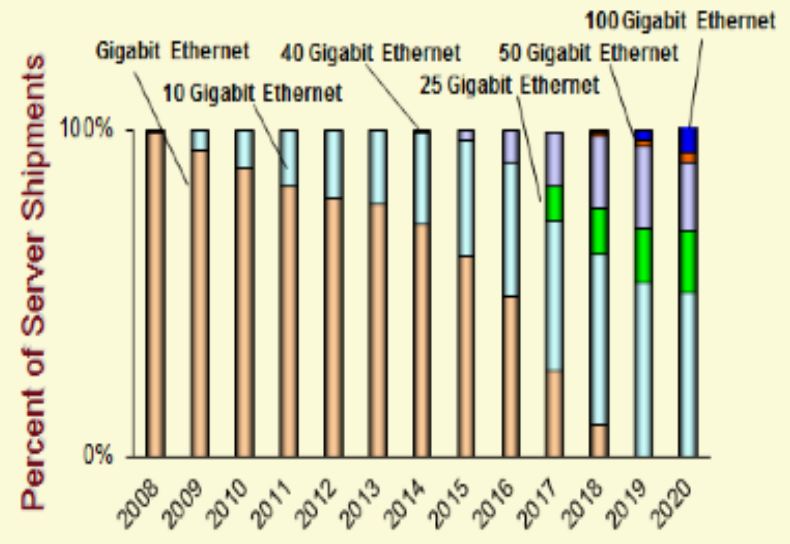
- Computing *supports the business*
- Thousands of companies with 100-1,000s of servers
- Hundreds of companies with 1,000-10,000s of servers

Server Ethernet Speed Adoption Cloud vs Enterprise

**Speed Migration on Cloud Servers
(Included in Dell'Oro Group's Server Report)**



**Speed Migration on Premise Servers
(Included in Dell'Oro Group's Server Report)**



In 2015, Cloud server interface is 90% 10G or above, Enterprise < 50%

Why do we need another speed?



- 40G and 100G ports need more SerDes, consume more power and reduce port density
- The Cloud needs to hit the sweet-spot of lowest price per Gigabit vs optimal performance

25G

50G: 2x25G

- 25G is a single lane specification, just like 10G
- Leverages IEEE 802.3 Ethernet framing
- Offers 2.5X the speed at a cost structure closer to 10G
- Same port density & connectors as 10G SFP+
- 50G is dual-lane
- Offers 1.25X the speed of 40G
- Cost structure is closer to 2X of 10G
- 2X the port density as 40G using splitter cables from QSFP

July 2nd 2014 - 25G and 50G Ethernet Consortium

- Founding members:
 - Arista, Broadcom, Google, Mellanox & Microsoft
 - 25gEthernet.org consortium website
- An open specification for the new speeds
- Consortium open to everyone in the industry





Cloud and New Enterprise:

- Open/standard cloud architectures
- Large Scale Deployments
- Automation, Resiliency, Programmability
- Rapid adoption of new technologies
- Competitive, multi-vendor

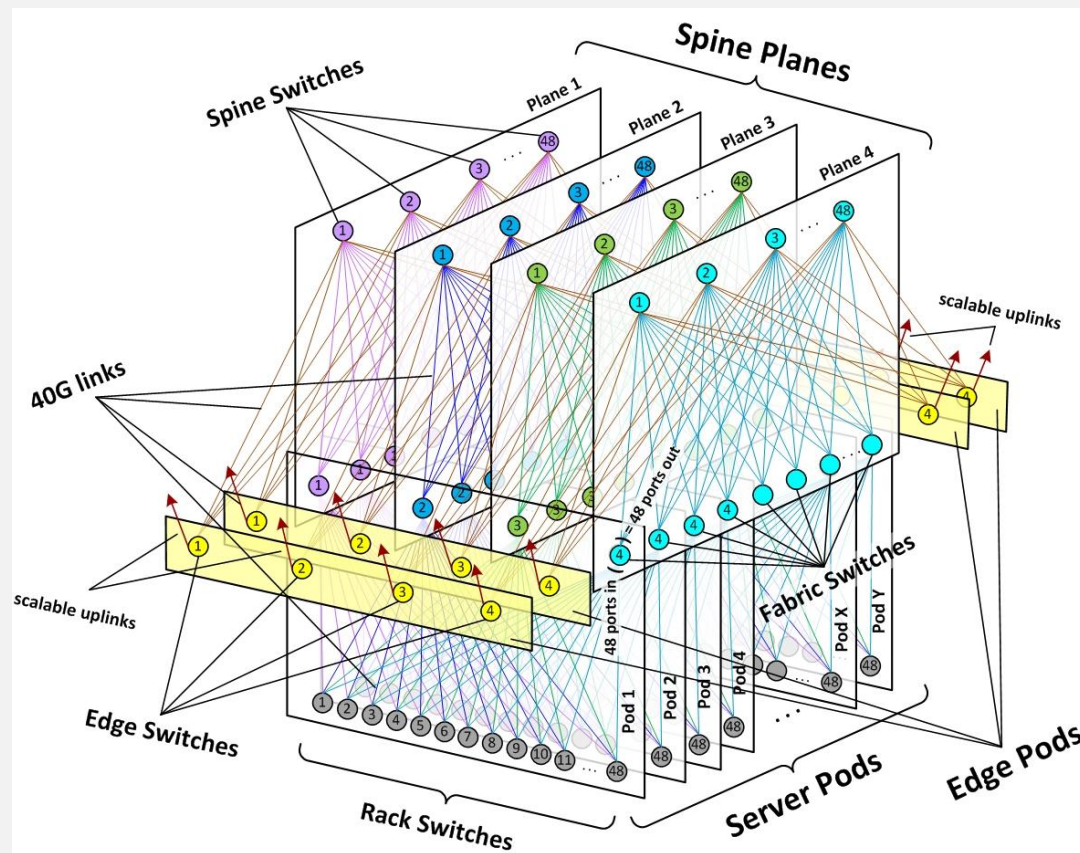


Legacy Enterprise:

- Complex/legacy network architectures
- Smaller incremental deployments
- Long time to deploy
- Slower adoption of new technologies
- Often single-vendor

Open Scalable Architecture chosen by Google, Facebook, Microsoft etc.

- Generally Available HW
- Merchant Silicon
- Clos Fabric
- Shifting to BGP Only



<https://code.facebook.com/posts/360346274145943/introducing-data-center-fabric-the-next-generation-facebook-data-center-network/>
<https://www.youtube.com/watch?v=mLEawo6OzFM>



What If Your Network Was Smarter Than You?

Jeremy Pollard



Who Am I?

Jeremy Pollard

Network Engineer @ Box.com

SIGGRAPH2015 GraphicsNet
Committee Chair

Automator

Lindy-Hop and Blues Dancer



<http://www.slideshare.net/DevOps4Networks/infrastructure-api-lightning-talk-by-jeremy-pollard-of-boxcom>

https://www.youtube.com/watch?v=2_1M3E57iio&feature=youtu.be

BGP Controller Extensibility

- The BGP controller by default only computes top prefixes and passes all the information used and the results to 'plugins'.
- Plugins can do with this information whatever they want:
 - Build reports
 - Build a prefix list and send it to a router.
 - Compare possible next-hops, AS PATH... with a monitoring tool to choose peers based on reliability, latency, company policies, etc...

```
# cat etc/config.yaml
max_age: 48
csv_delimiter: ";"
max_routes: 30000
min_bytes: 0
packet_sampling: 10000

... (output omitted)

plugins:
- 'prefix_data.SavePrefixData'
- 'statistics.RouteStatistics'
- 'statistics.OffloadedBytes'
- 'bird.Bird'

... (output omitted)
```

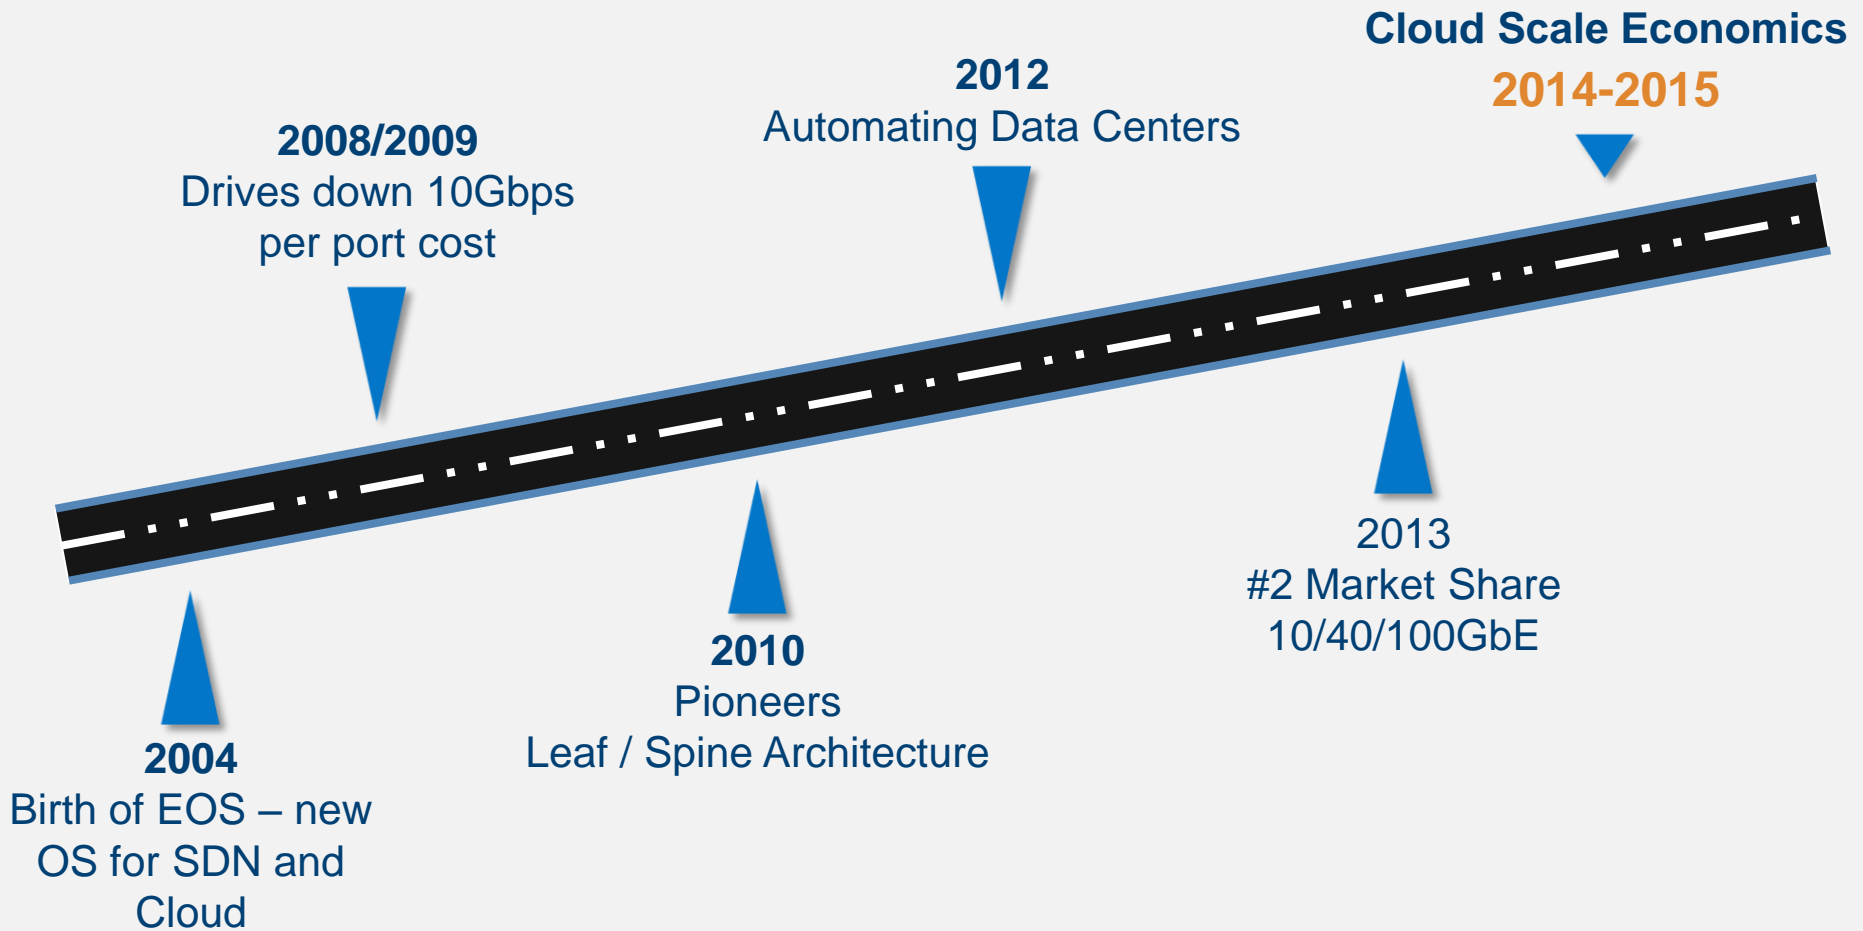
<https://github.com/dbarrosop/sir>

<https://www.youtube.com/watch?v=o1njanXhQqM>

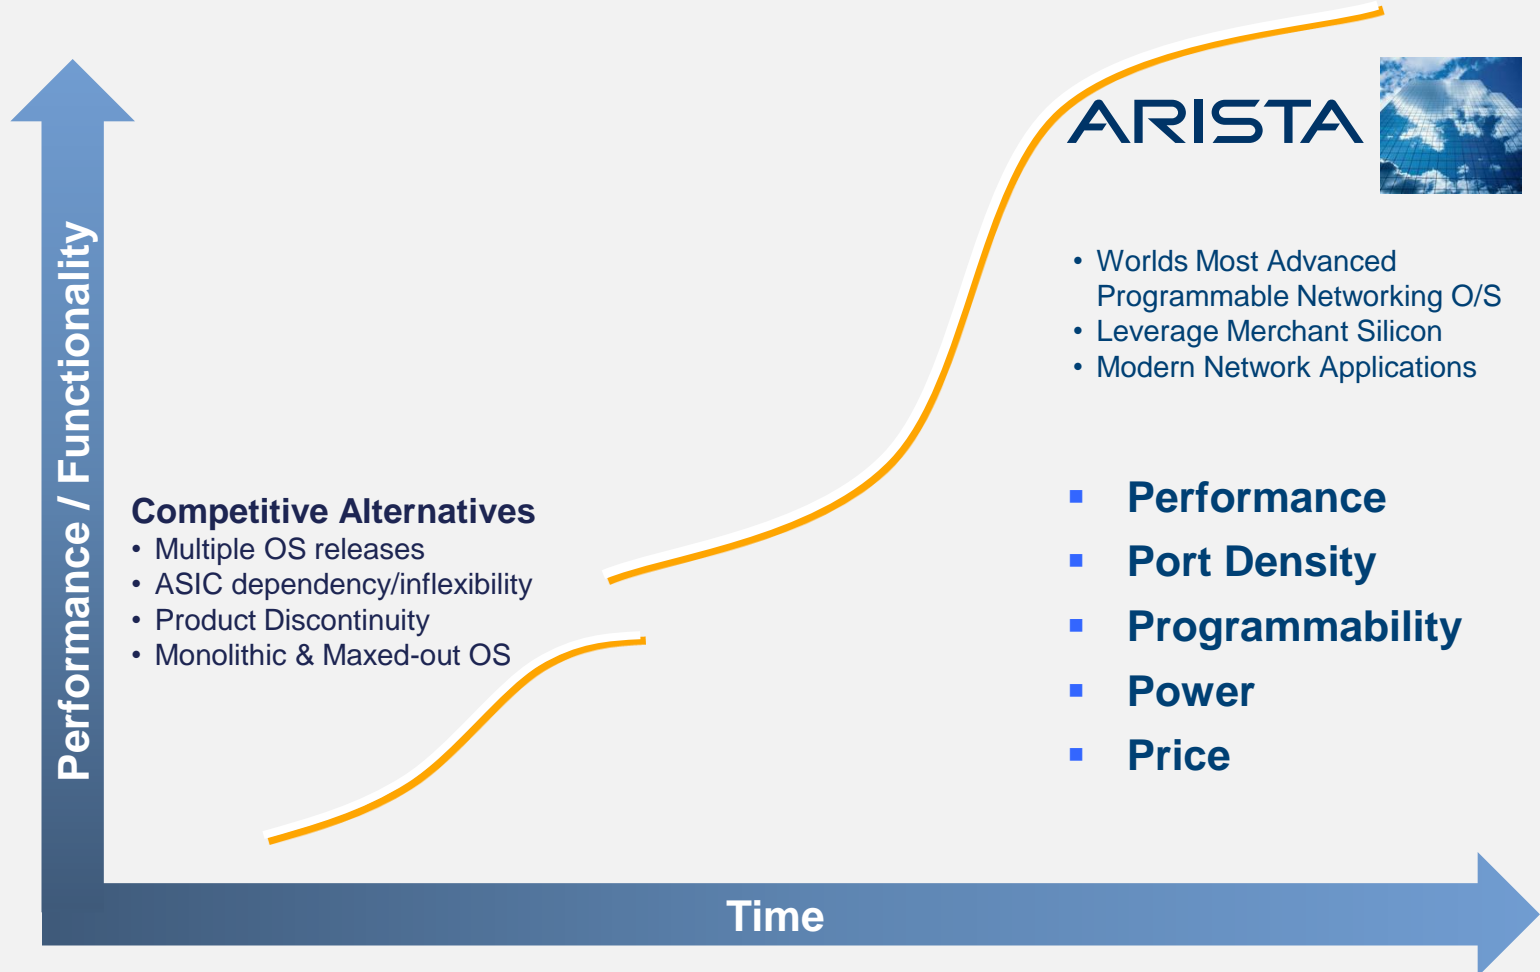
[v=o1njanXhQqM&index=5&list=PLXSSXAe33jl2IIWtfnnEj5J7B7KoixKCe](https://www.youtube.com/watch?v=o1njanXhQqM&index=5&list=PLXSSXAe33jl2IIWtfnnEj5J7B7KoixKCe)

Arista's Ongoing Innovation Path

ARISTA



Disruption causes Structural Disadvantage



Merchant Silicones based on your applications



Intel FM4000:

10G Focus
Ultra Low Latency
Onboard Telemetry

ULL Applications

Broadcom Trident+

10G/40G Focus
High Density
Low Power

Low cost leaf/spine

Dune Petra:

10G Modular Focus
Focus
Ultra Deep Buffer

High Load /Application proof

Intel FM6000:

10/40G SDN Focus
Programmable HW
Virtualization

High function applications

Broadcom Trident2:

10/40G Focus
48 – 2048 ports

General purpose leaf/spine/spline

Broadcom Arad:

40/100G Focus
Ultra Deep Buffer
True fabric design

Multi workload / Future proof

Real World Data

Customer	Real Buffer Utilization Observations	Max Buffer Used
HPC	Storage Cluster – Medium	33 MB
Animation	Storage Filer (NFS)	6.2 MB
Software vendor	Engineering Build Servers (Perforce)	14.9 MB
Business Analytics	Hadoop 2K servers – Big Data	52.3 MB
Enterprise	Data Center Virtualization	52.4 MB


ARISTA

Extensible Operating System (EOS)


Cloud Purposed Operating System



Unmodified Linux



Shared State



Modular / Resilient



Open / Programmable

High Performance, Merchant Data Plane

Physical Machines



Virtual Machines



Storage



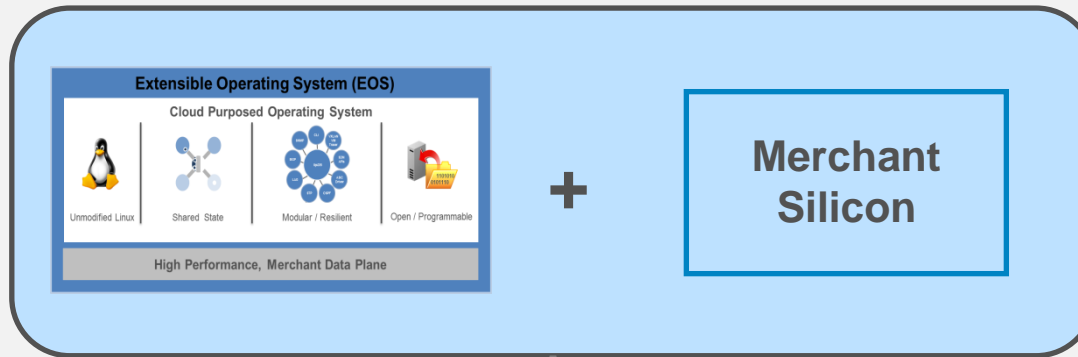
Big Data Nodes



Cloud



Software Driven Networking



Cloud Scale Architecture

Availability

Workflow Visibility

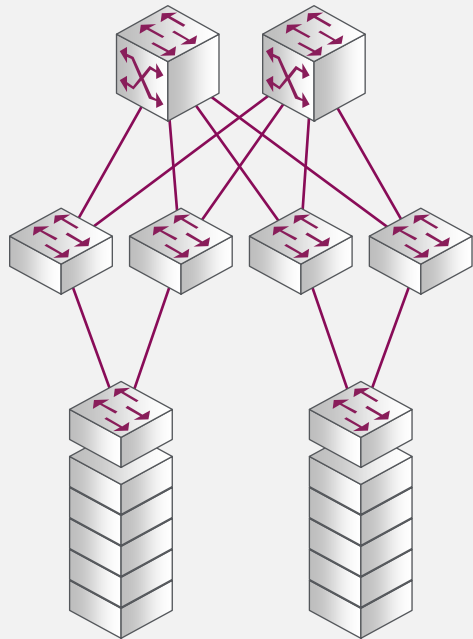
Open and Programmable

Workload Automation

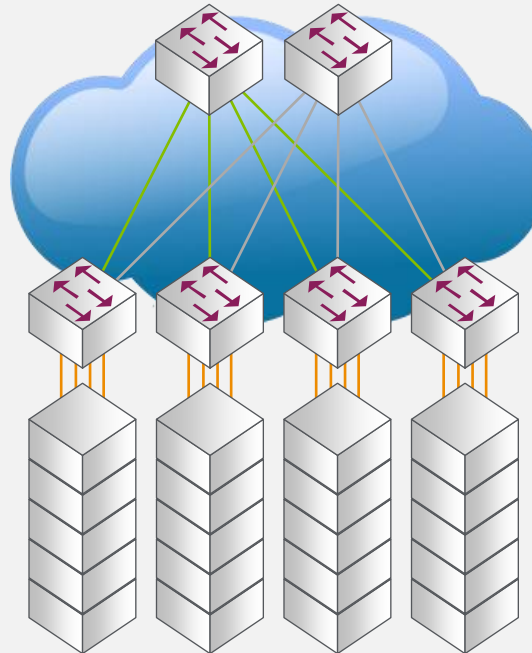
Accelerates Time to Service and Reduces TCO

The Next Generation of Scale in Cloud Networking

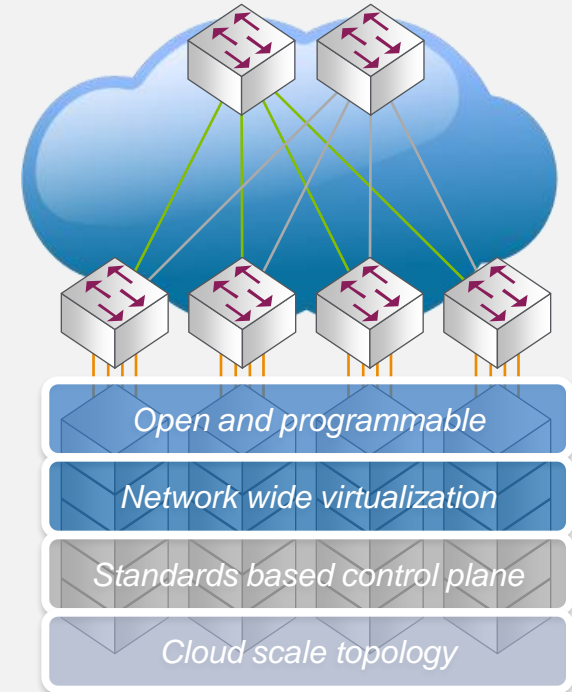
Traditional Network Architecture



Spine Leaf Or Spline™



Software Driven Cloud Networking



Optimized to transport data in/out of the data center

Optimized to handle east-west traffic within the data center

A smart, flexible and manageable network at cloud scale

Savings via total cost of operational (TCO)

ARISTA

Alternatives

Savings

Single Image

V
S

Multiple Images

3x-6x

Virtual Environment

V
S

None/Limited

10x

Automated Deployment

V
S

Manual Provisioning

60x

Proactive Notifications

V
S

Reactive responses

100x

Ecosystem Integration

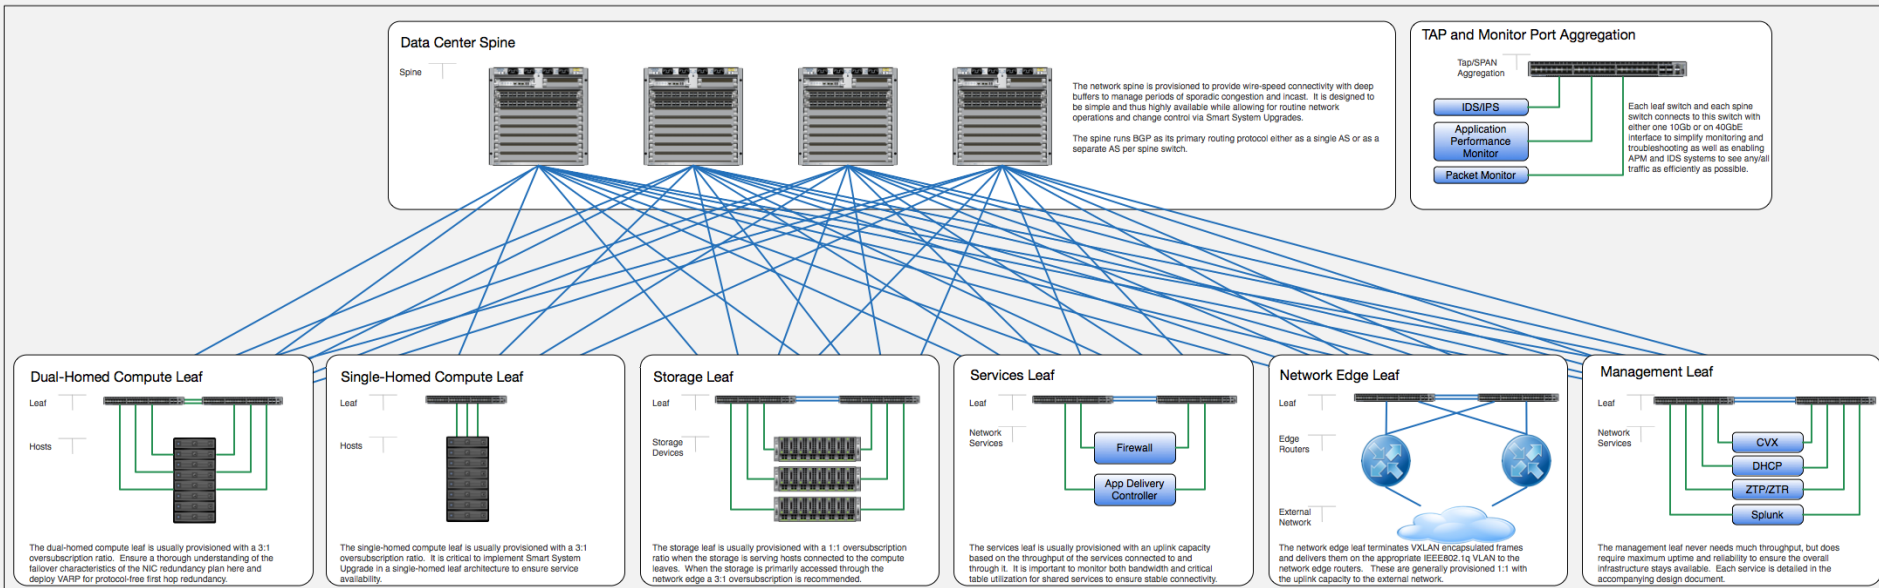
V
S

Proprietary Stack

Silo IT to SDCN

Arista UCN(Universal Cloud Network) Architecture

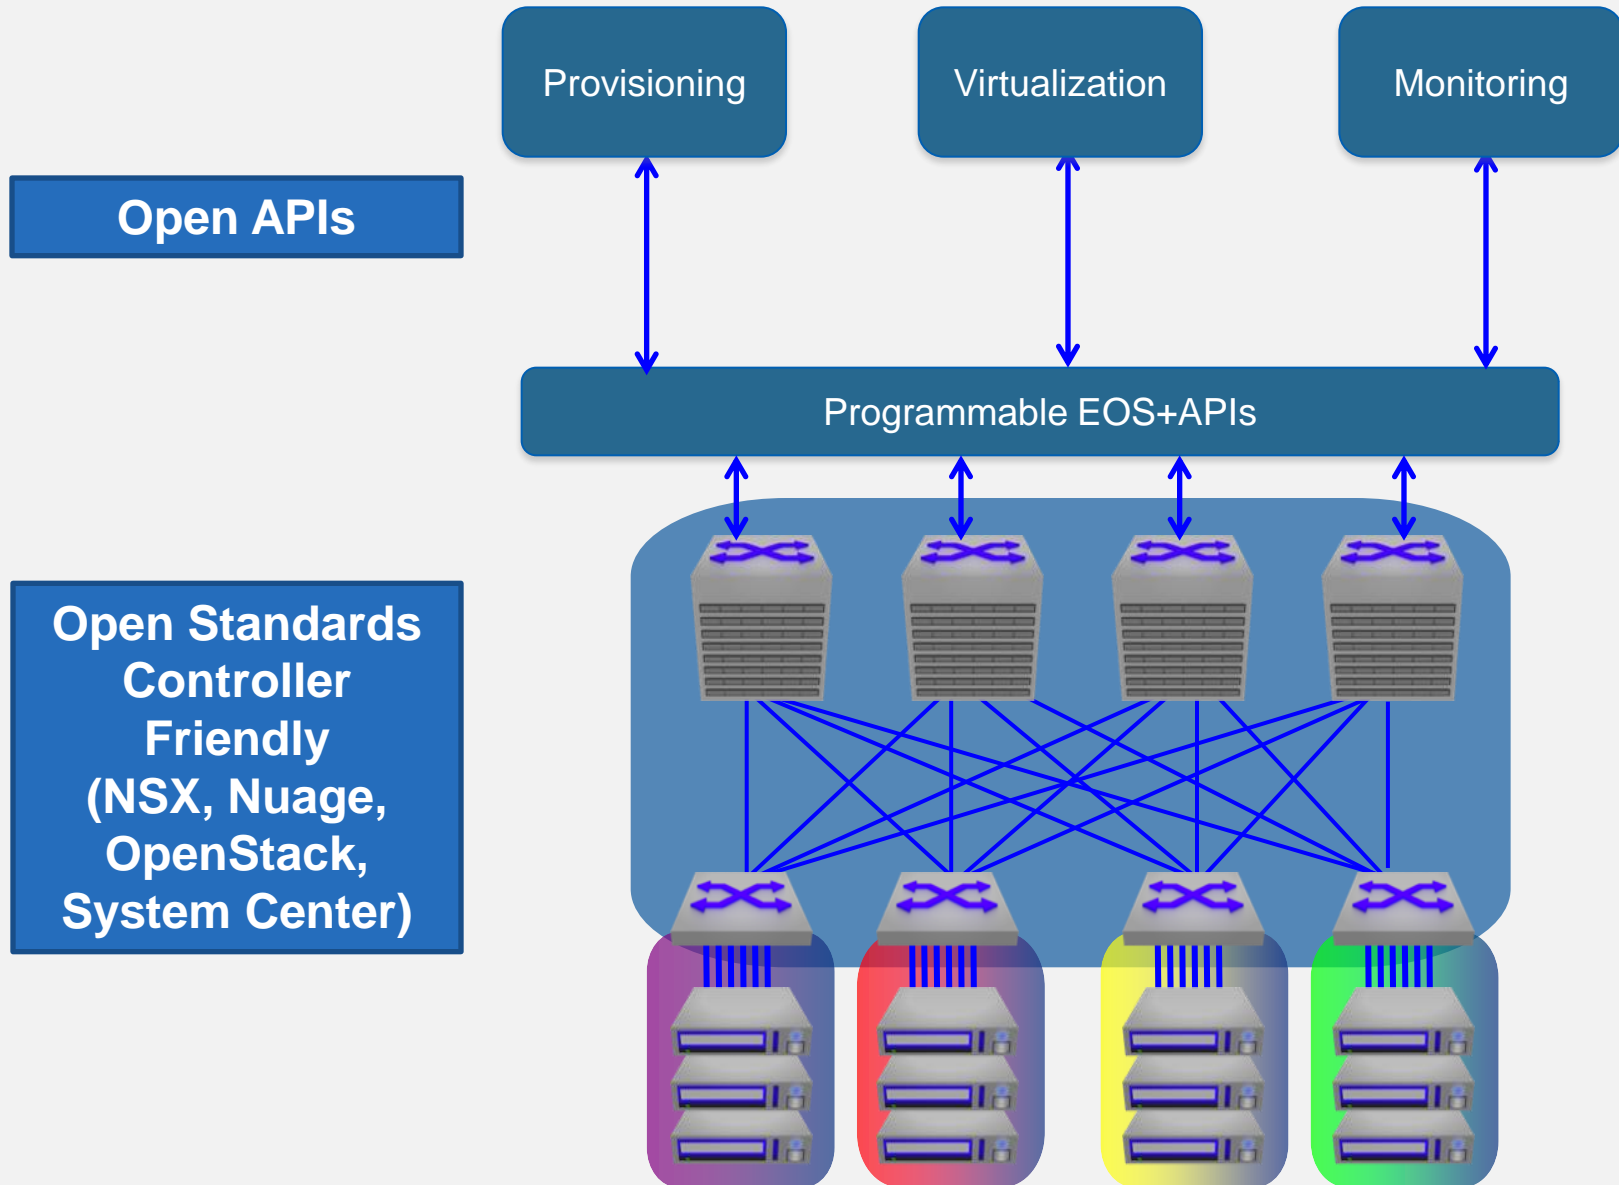
ARISTA




Universal Cloud Network Design Principles

- Highly Available and Resilient Network Topology
- Simplicity
- Automation of network management functions
- Open
- Scalable
- Network Telemetry

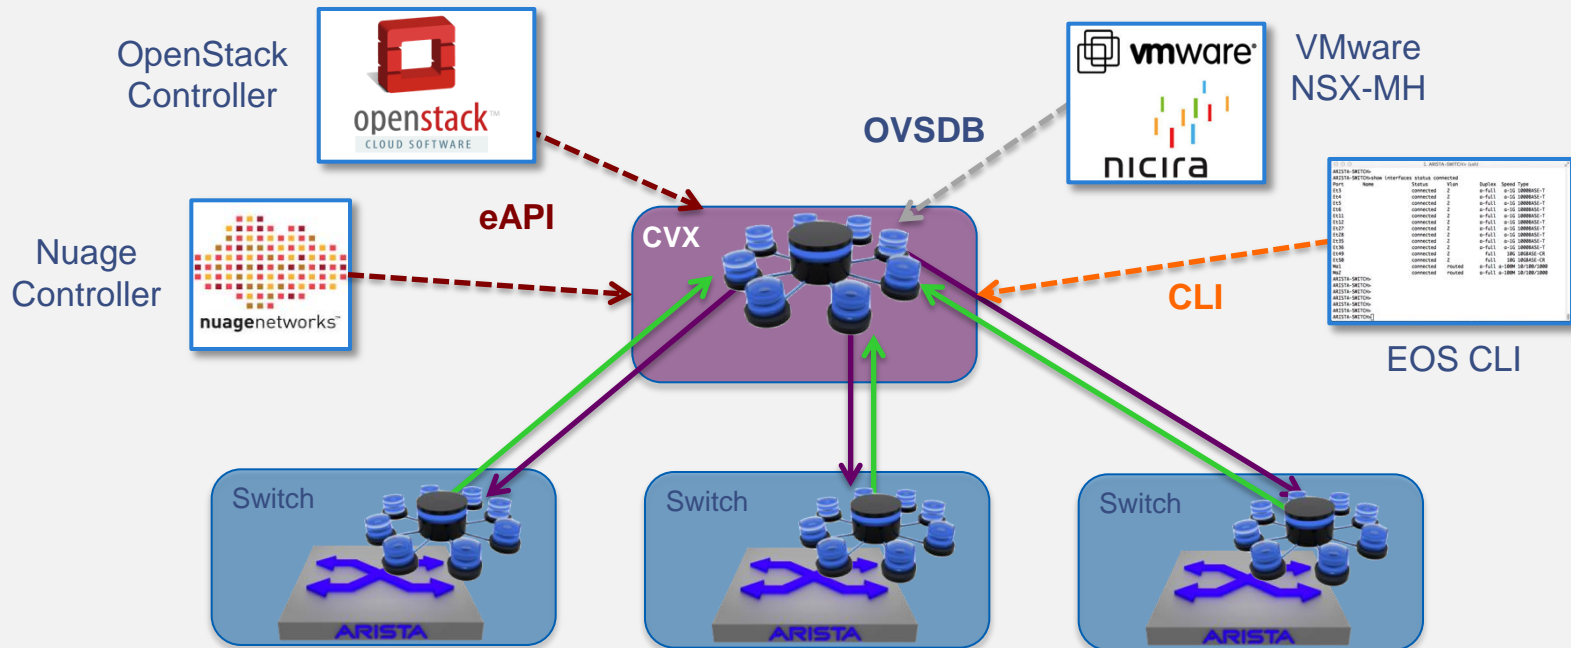
Cloud Vision: The Network View is Dramatically Different than a Decade Ago





- **CVX allows for EOS SysDB to aggregate and synchronize a subset of state information**
- **CVX can run on physical switch or VM with vEOS**
- **Communication can be in-band or over the management network**

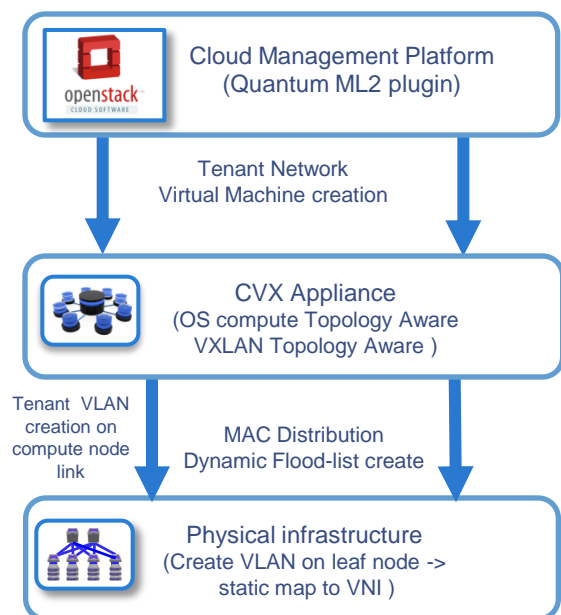
Arista CVX with BYOC(Bring Your Own Controller)



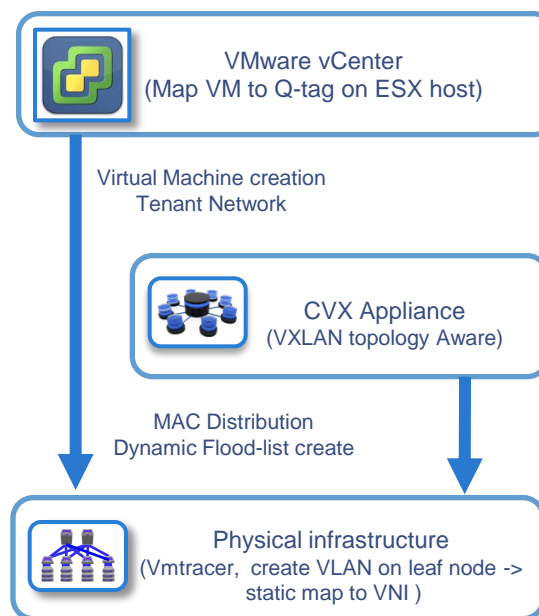
- CVX Serves as integration point for controllers, orchestration systems and NMS
- Native support for for NSX for Multi-Hypervisor with OVSDB
- eAPI drivers for OpenStack and Nuage (others coming)
- Provides integrated provision of VLANs and VXLANs across physical and virtual switches

▪ Models for Orchestration of Performance Based Hardware only VTEPs

Arista - OpenStack Integration



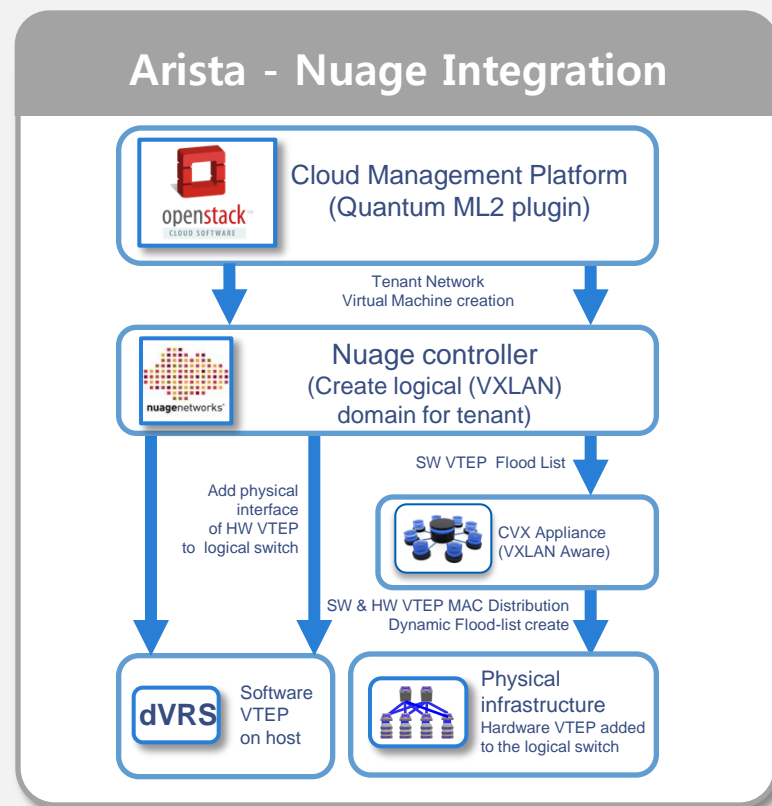
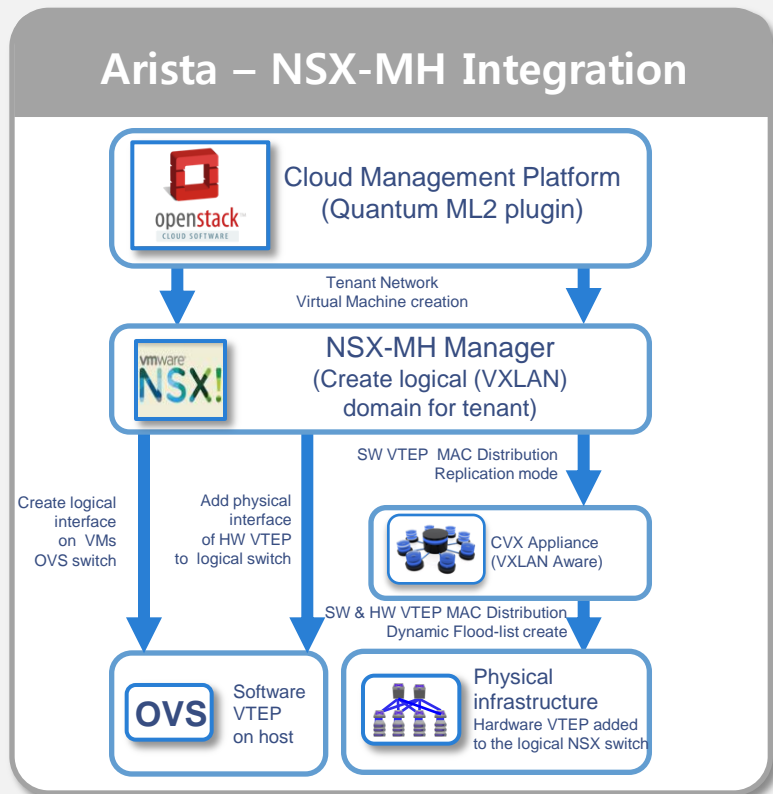
Arista - vSphere Integration



- Dynamic creation and deletion of tenant networks (VLAN) on the physical infrastructure driven by OpenStack.
- VXLAN awareness via eAPI created mapping on the physical switches

- Dynamic creation and deletion of tenant networks (VLAN) on the physical infrastructure driven by VMTracer
- VXLAN awareness via eAPI/CLI created mapping on the physical switches

Models for Orchestration of Hardware & Software only VTEPs



- Integration of Arista HW VTEP into the NSX logical switch for VTEP performance, BMS and storage
- Dynamic distribution of SW MAC addresses between to HW VTEPs

- Integration of Arista HW VTEP into the Nuage logical switch for VTEP performance, BMS and storage
- Dynamic distribution of SW VTEP flood list

Summary – Delivering World Class Software Driven Cloud Networking

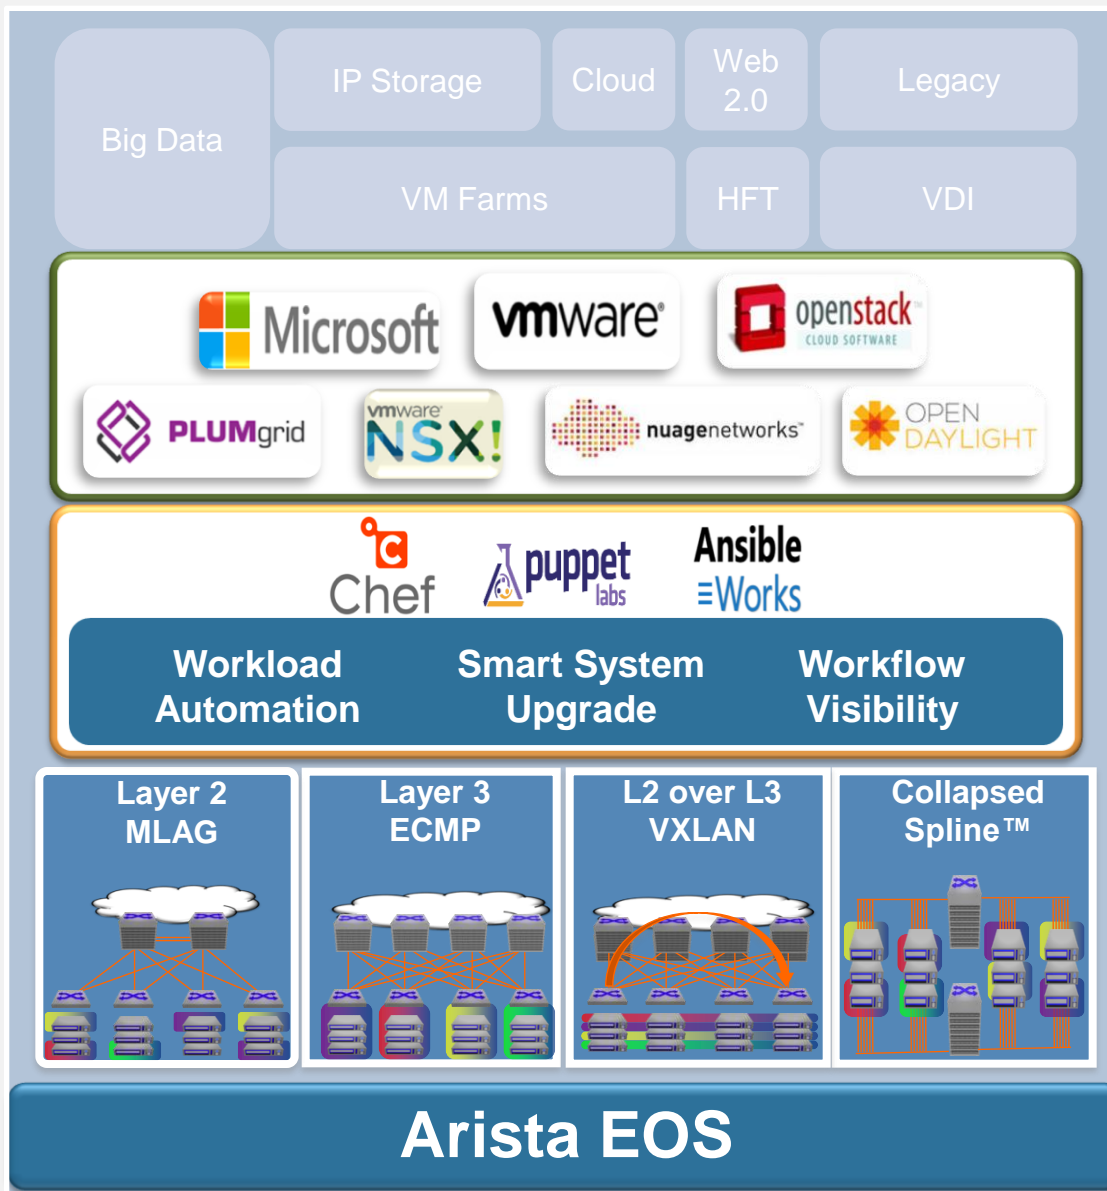
ARISTA

Use Cases

Open Network Apps with BYOC

Operations Deployment

Universal Cloud Network



Thank You!