Reliability, Availability & Serviceability @ Cloud Scale

Ian Hood, P. Eng.
Chief Architect – Global Telco
June 2016

redhat.

The forces of change
Adapting to the networks of the future

The Business is Moving Fast - Downtime is Costly

Goal of Zero Service Downtime >>> Different Approaches

Laser Focus

- Deploy In Minutes
- Hot Deployments (DevOps)
- Optimize Resources / Costs
- Analytics / Self Healing

Stone Age Tools

- Deploy in Months / Years
- Complex Costly Service Environments
- Mandate 99.999% Availability
- Limited Service Visibility

Imagine if Your Network Was Completely Elastic

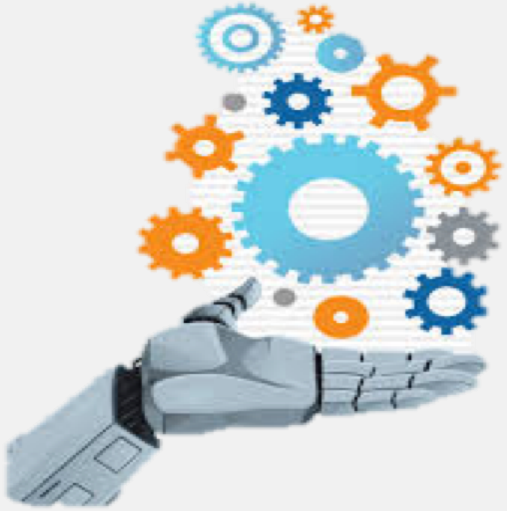

And Completely Automated!

We Have Done it Again ..

Architecture is More Complex!

Implementations are More Complicated..

TELCO CLOUD SERVICES

Pathway to New Services

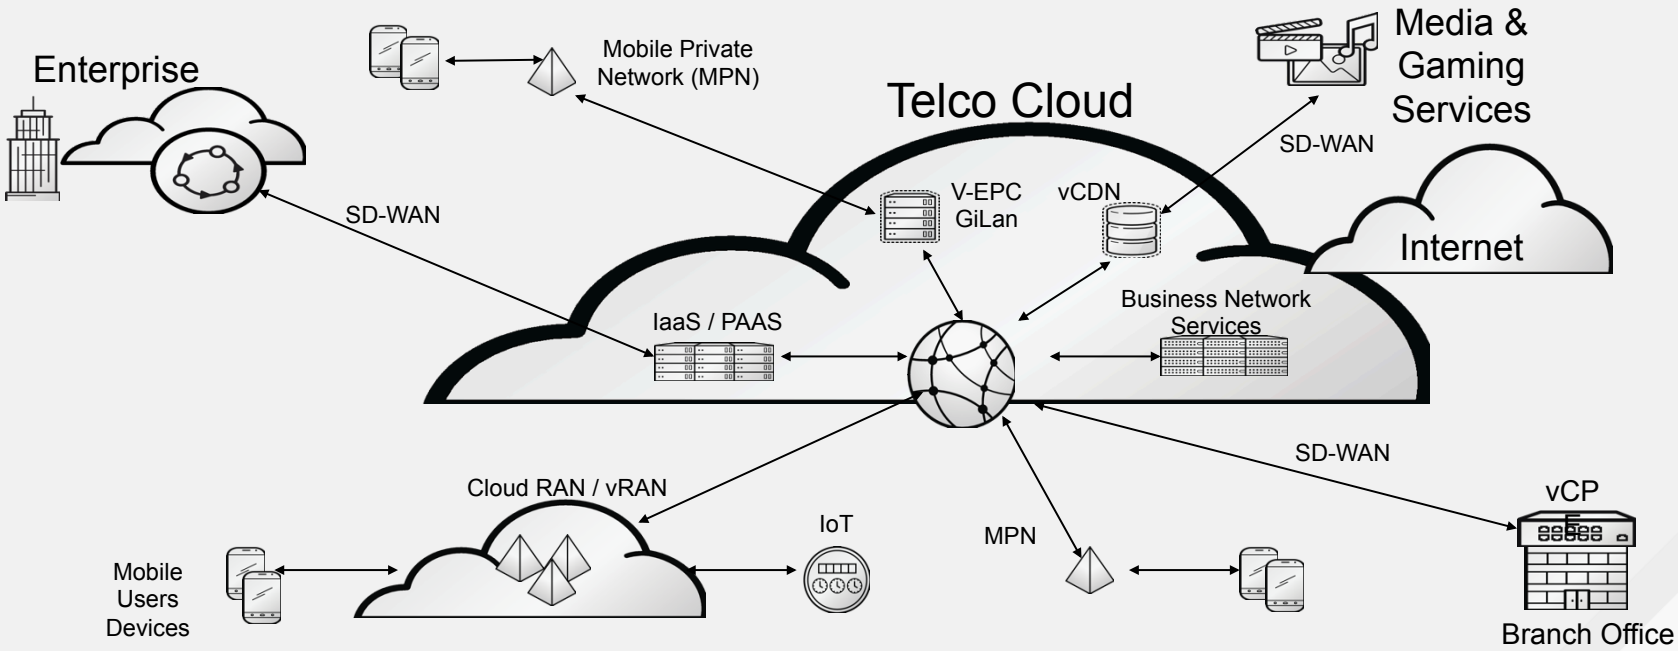

What Are Telco's Looking For?

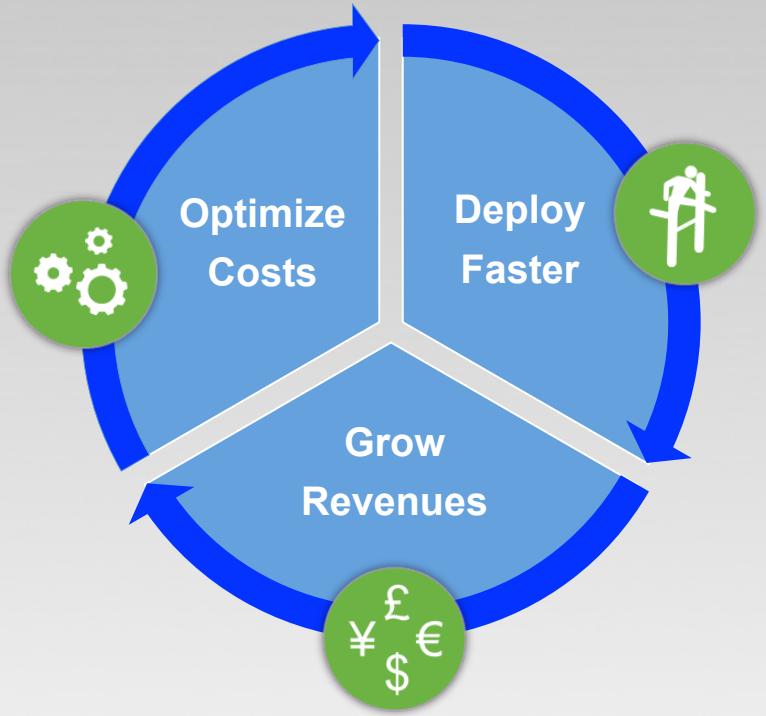

Reliability

Availability

Serviceability

"Automation is a key business imperative"

Telco NFV Service Delivery Framework

MULTI-VENDOR

COMMUNITY COLLABORATION

OPEN INTEROPERAPLE APIs

AUTOMATE EVERYTHING

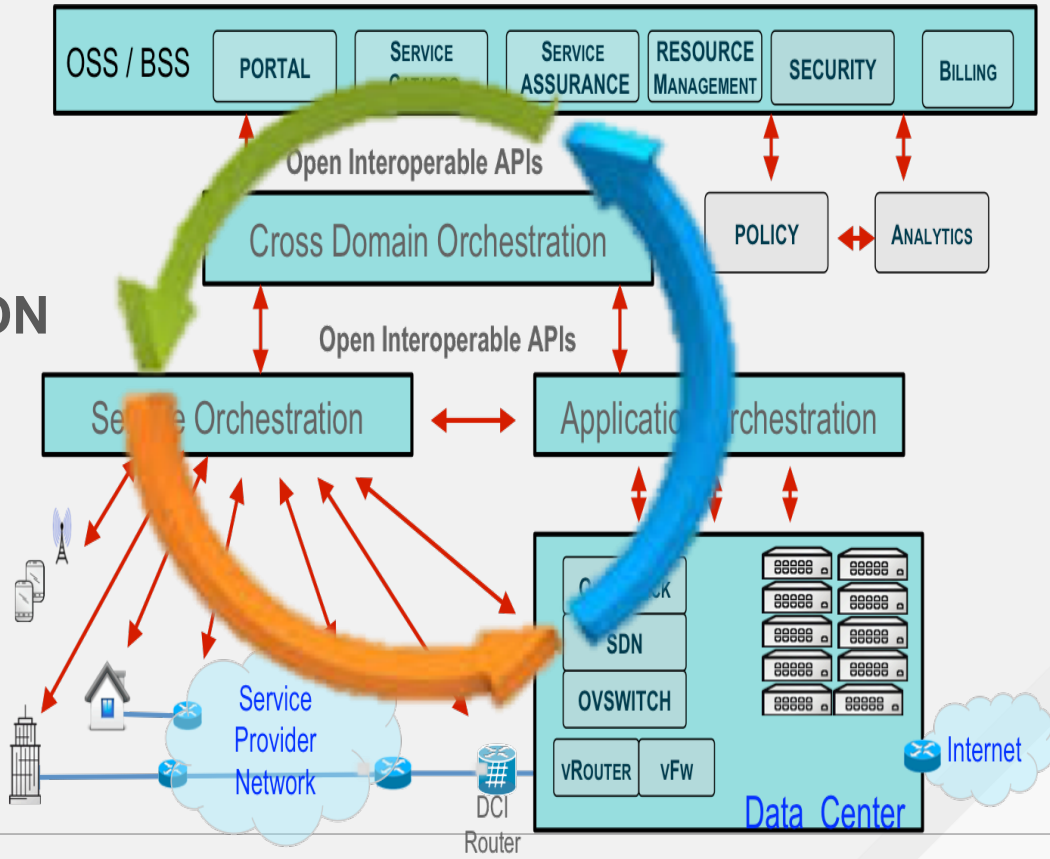

Modernizing Network Functions

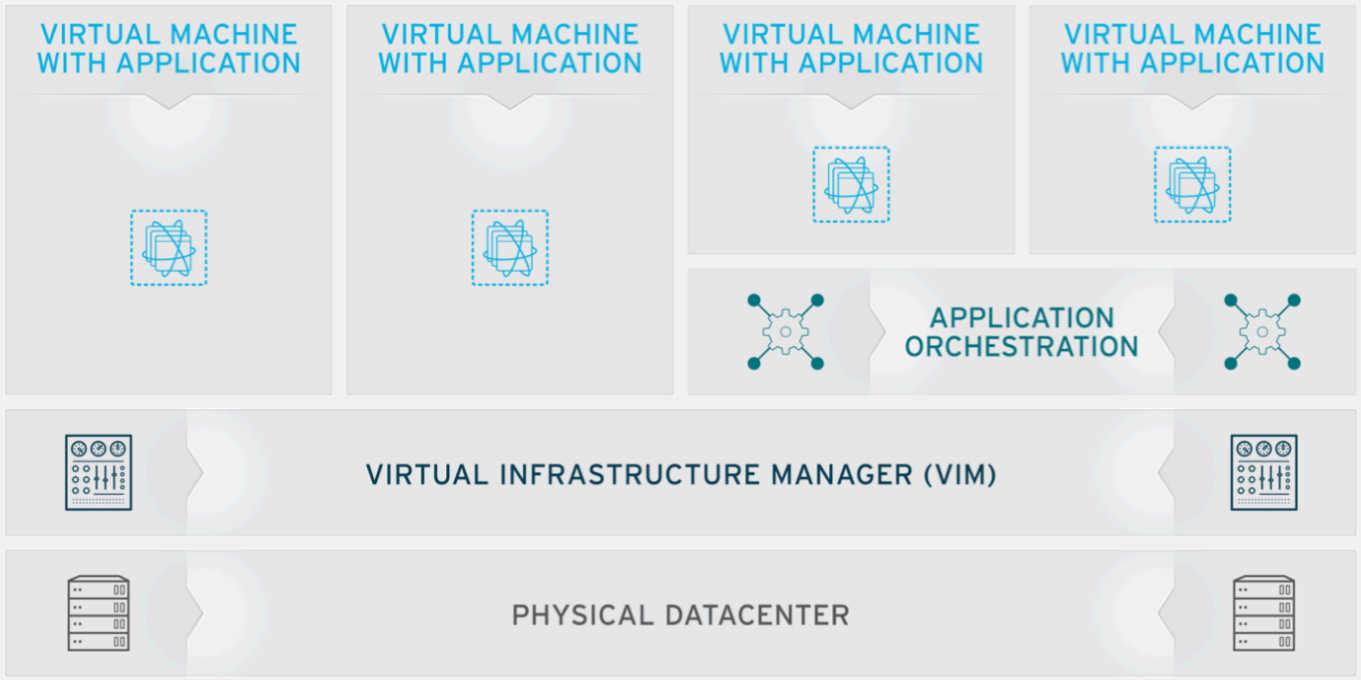

Flexible Redundancy / Availability Zones

MANAGEMENT PLANE

Active:Standby / 1:N / Geo-Redundant

CONTROL PLANE

Hierarchical / N + 1 / 1:N / Active / Active

DATA PLANE

Diversity / 1+1 / 1: N

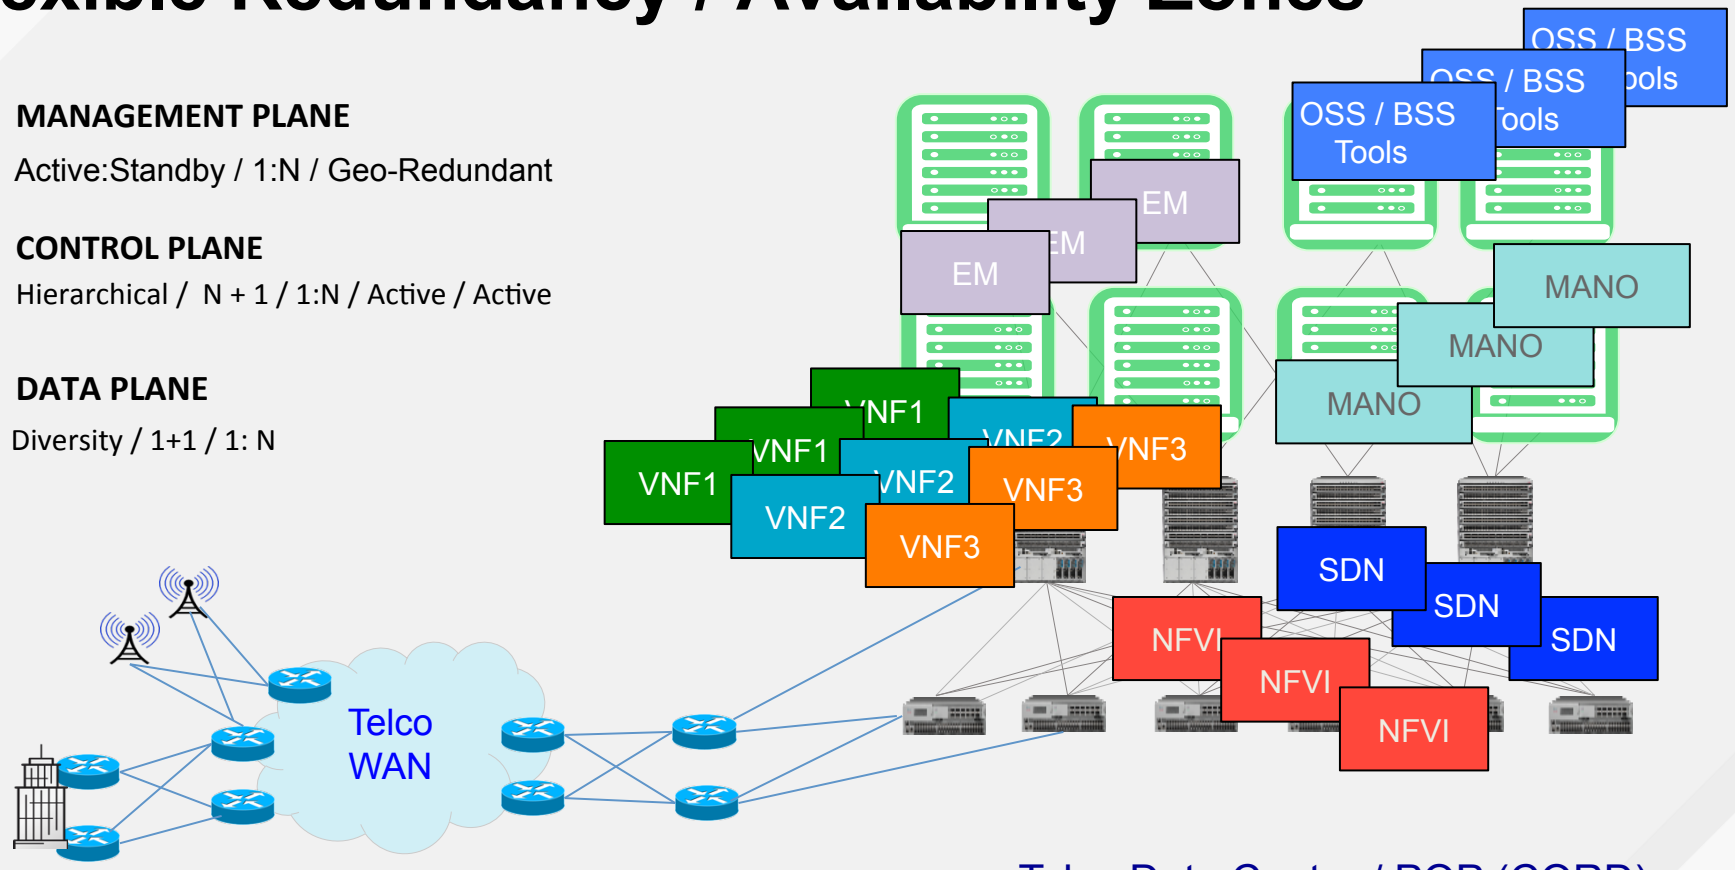

Telco Data Center / POP (CORD)

Serviceability / Manageability

Service Delivery at the Speed of Software

Automate -- Automate -- Automate (AAA)

- MANO - Performance, availability, and Life-Cycle Management
- Service Level Monitoring / Assurance Including Service Level OAM Telemetry
- Open Interoperable APIs and Data Models (Everywhere)
- Accelerate speed to Create / Modify / Delete Services -- Self-Service Portal
- Reduce Fault/Service Recovery Times – Automated Tools / Analytics

Aligning with Key Telco Priorities

	Telco Priorities	**Red Hat Open Telco**	**Key Telco Benefits**
1	Real-Time Low Latency Service Handling	Real-Time KVM DPDK Acceleration	Improve Ability to Deliver Per-Service SLAs Improve End Customer Experience Enhanced Service Throughput, Latency, Jitter
2	Cost Effective Infrastructure Scale	Hyperconverged Infrastructure (HCI)	Cost-Effective Scale Improve Customer Experience
3	Open Interoperable MANO Elements	Open Source Automation	Open Interoperable APIs in Multi-Vendor Environment Automate Network/Service Deployments Migrate workloads across physical / virtual environments
4	Security and Policy Management	RBAC with Audit SELinux + sVirt Validated Encryption	Role-Based Audit of Application / Access Controls Multi-Tenancy /Application Isolation Strong Intrusion Protection
5	In-Service Upgrades	High Availability OSP Reference Architecture	OSP Lifecycle Management Minimize Service Outages During Upgrades

WHAT IS NEEDED?

Resilient Services + Applications

Open Automation Software

Cost-Effective Robust Hardware

It's About Delivering Trusted Reliable Services … Anywhere

COLLABORATION IN A COMMON SPACE TO
SHARE INNOVATIONS AND SPEED SOLUTIONS.

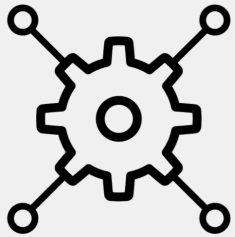

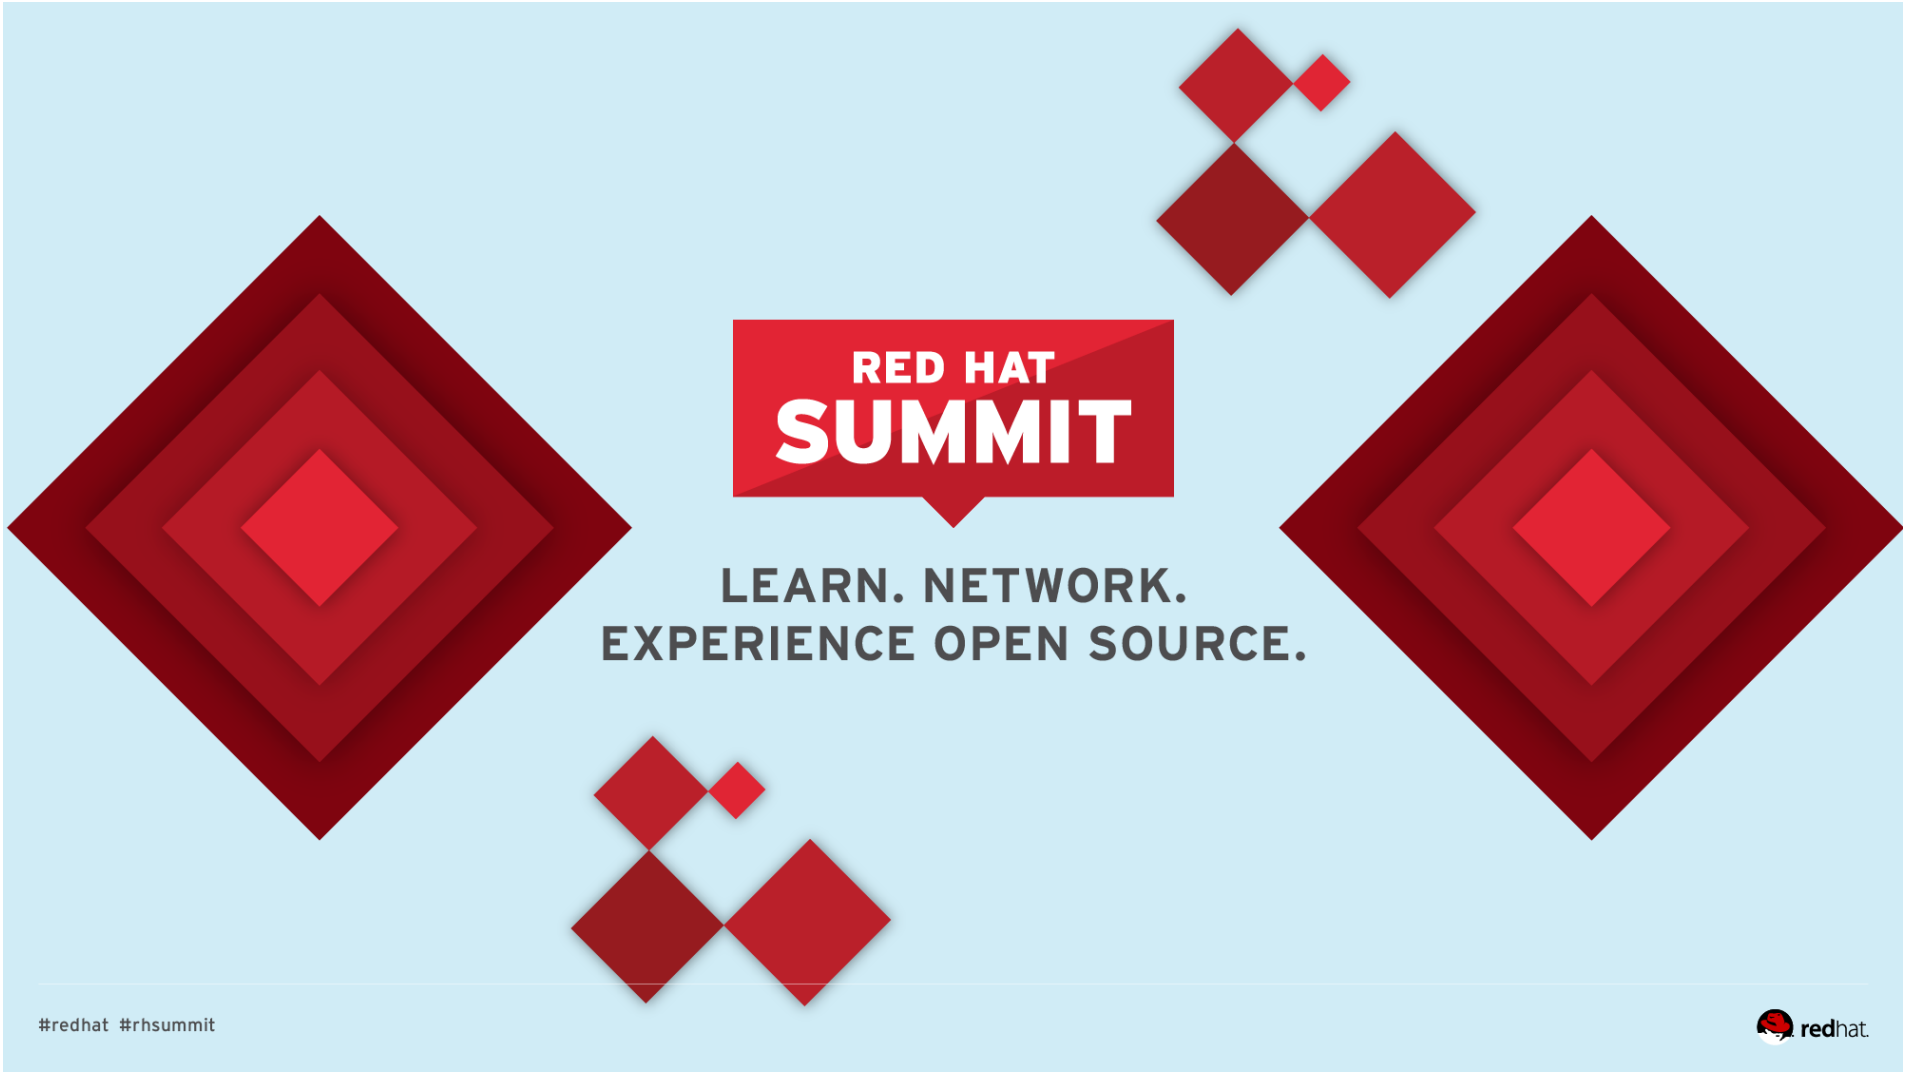
RED HAT
SUMMIT
LEARN. NETWORK.
EXPERIENCE OPEN SOURCE.
#redhat #rhsummit
redhat.

Story of an agile frog in a large company

RedHat Summit, telco breakout session
June 29, 2016

Cyrille Puget – PaaS Product Manager
Innovation, Marketing and Technology division

orange™

Orange at a glance

#53 in 2016 global brand ranking

263 millions
customers worldwide

1.8 million
Fibre customers (and 1 million in France)

450 000 kms
of undersea cables (enough to go around the earth 10 times!)

16 millions
Orange Money customers in 14 countries

726 millions
euros invested in research and innovation

6 930
patents in our R&D portfolio

4G
in 13 countries

+10 million
fans on Facebook

+130 000
followers on @orange on Twitter

+1 million
visits on Orange.com each month

Essentials2020

Group strategic ambition is focused on Customer Experience and based on 5 levers.

Orange aims to provide everyone with a unique customer experience every day, by designing digital services that help them make the most of what is essential to them – with complete peace of mind.

A digital, efficient and responsible company

Innovation, Marketing and Technology division

Our role is to design digital services that are simple, reliable and secured for all Orange countries we serve and their respective customers.

Communication services, IoT, SIM, API, TV, Cloud, Mobile Banking, Big Data, Enterprise services

7,900 employees worldwide
5,000 in innovation
€726 M invested

Situation a few years back

« Legacy » organization, infrastructures, tools and processes
From traditional infrastructure to virtualization and cloud:

- Mainly VMWare based but new cloud infrastructure with OpenStack
- Pure telco services still using traditional infrastructure

Tools and processes:

- Some old tools, manual processes
- Distributed teams, many actors

Impacts:

- Time to market
- Hard to achieve reliability, availability, and serviceability at cloud scale

How are we evolving

Group evolution on:

- Service oriented architecture
- Project organization
- Skills and tools
- Cloud and platform infrastructure
- DevOps transformation

IaaS is not enough, let's build a PaaS

Although improvments were realized, we believed we could do much more:

- We could go faster, no time to loose
- Provide modern cloud/platform solutions for projects
- Simplify projects' everyday life (for developers, testers, ops)
- Support the whole DevOps cycle: build, deploy & operate applications

Indeed, for many projects, IaaS is not enough or doesn't fit:

- It requires IT specialists/time that many project don't have
- Projects need more than just IaaS (e.g CI/CD, test, scaling, monitoring…)

We should have a PaaS, quickly, but how?

PaaS is the solution? Prove it… in less than 6 months

Mission: demonstrate that PaaS is
a good solution for our needs
Plan:

- Find a cute project name: Kermit
- Be agile, focus on serviceability, go fast
- Choose a hosting infrastructure
- Choose a PaaS solution
- Deploy a beta platform
- Get some users, support them
- Iterate
- Move to production

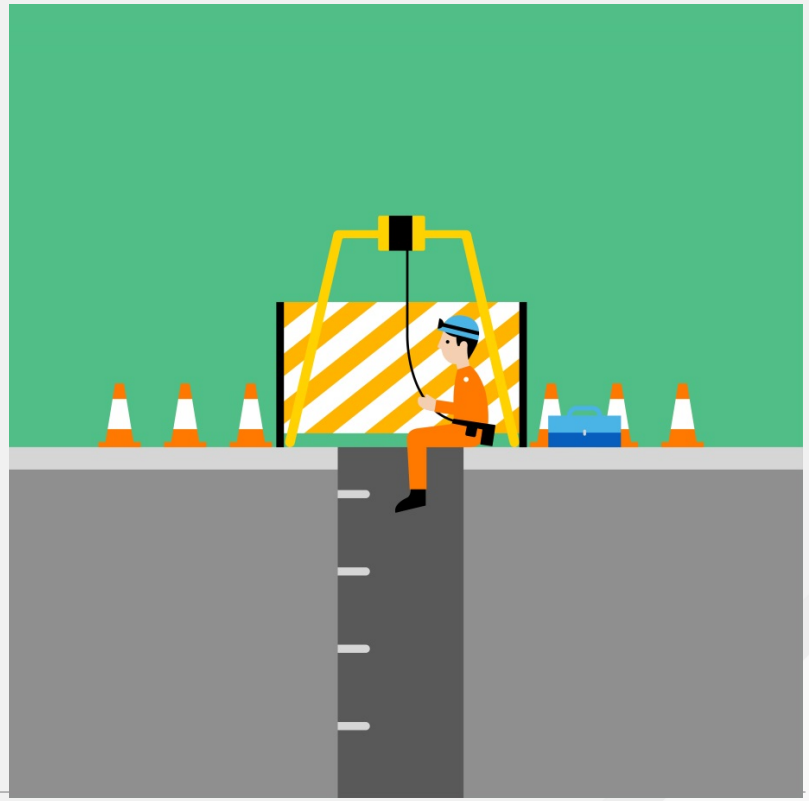

OpenShift adoption

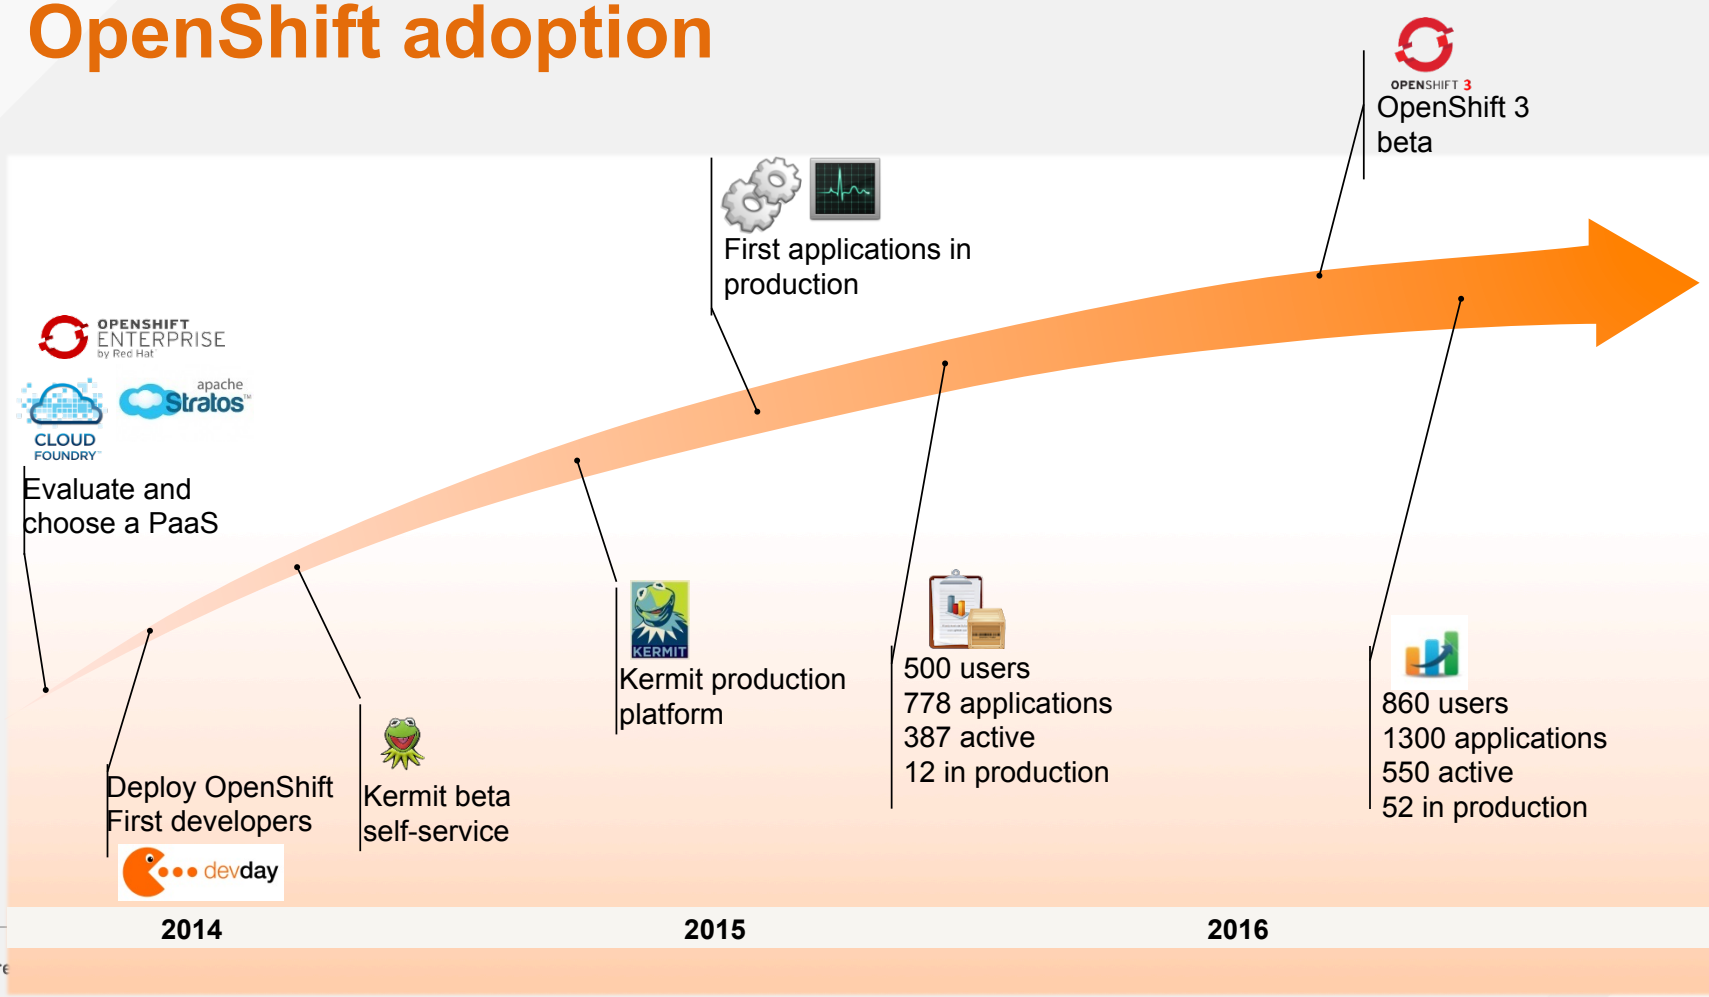

Examples of applications running on OpenShift

- Web application used in Orange shops
- Rich Communication Framework: infrastructure and client framework to ease development of communication and real-time services
- Backend of TV remote control mobile application providing an enriched experience by delivering information of on-going programs
- IoT data aggregation services targeting end users and service providers
- Ticketing tool used by some operations team
- Many backends for mobile applications

Outcomes

OpenShift platform:

- Easy to deploy and manage: only 2 persons doing ops/support
- Increadibly stable: >99.95% platform availability
- Valuable RedHat support

Users and projects:

- Strong adoption, numbers are speaking
- Enthusiast feedbacks, huge benefits
- 3 months saving on a 6 months project

Limitations:

- web app (http), scaling databases, resiliency
- Bronze SLA

What's next?

Support larger scope of applications:

- Complex application architecture (micro-services)
- Telco services, not just web apps

Provide higher service availability:

- Large scale, multi-DC deployment
- Improved resiliency at platform and application level

Hopefully, OpenShift 3 / Kubernetes is designed for this.

Thanks!

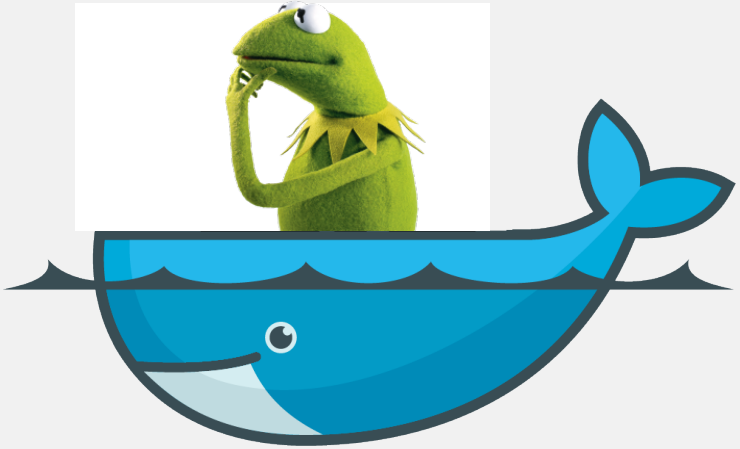

VIRTUALIZED SP NETWORKING

!= Enterprise Switching

Dave Ward
dward@cisco.com

Target Assumption:
Simple to deploy Apps and Services and Stack Does the Right thing To manage itself

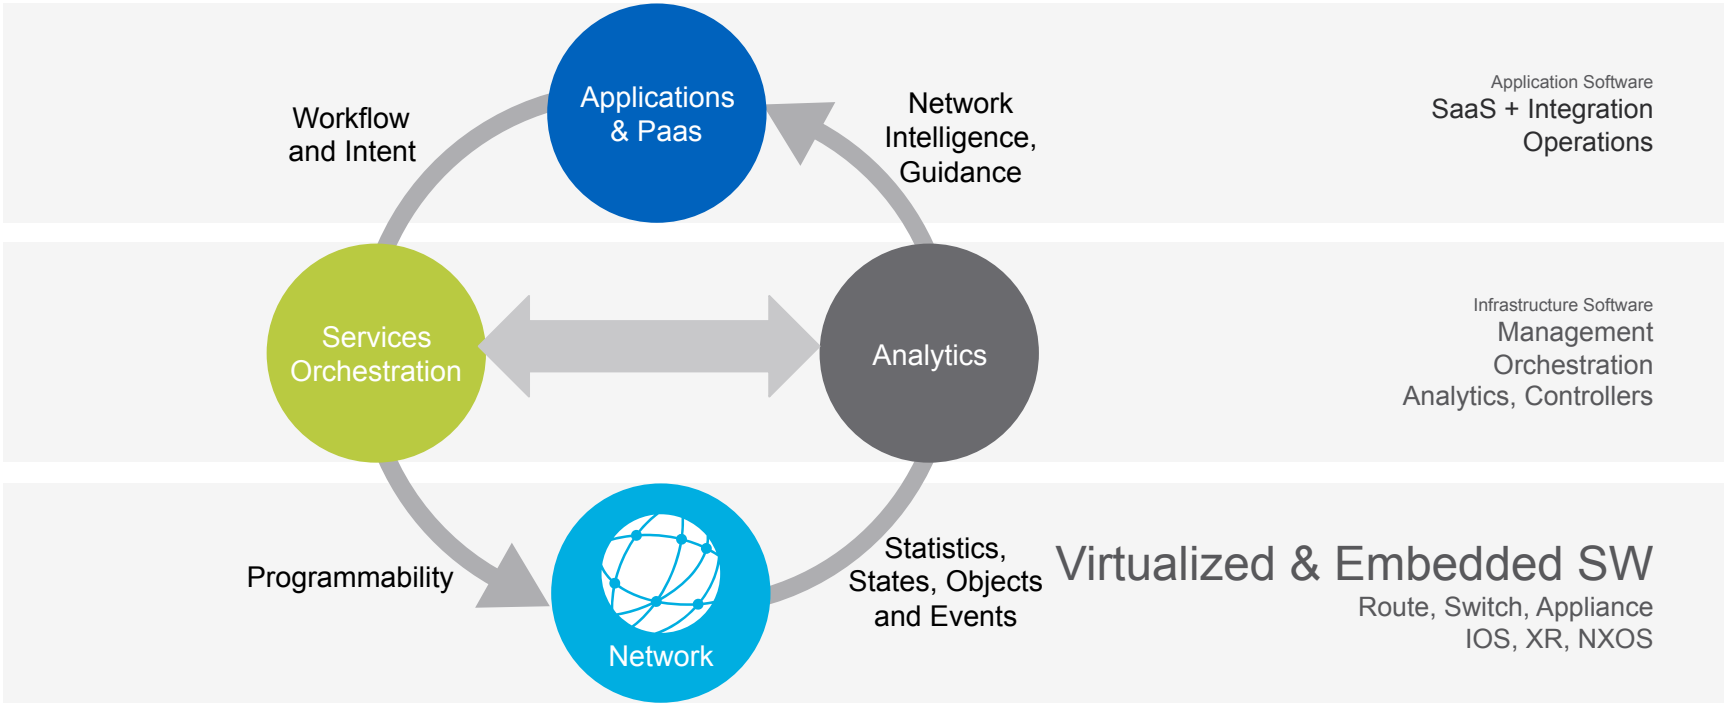

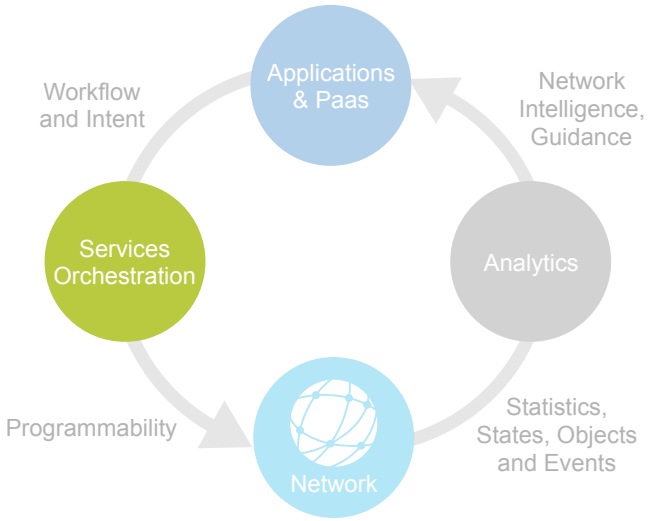

Services Orchestration Platform

Incorporating policy driven orchestration capabilities of complete services life-cycle manager that can implement Service Workflow topologies that span the network and data centers

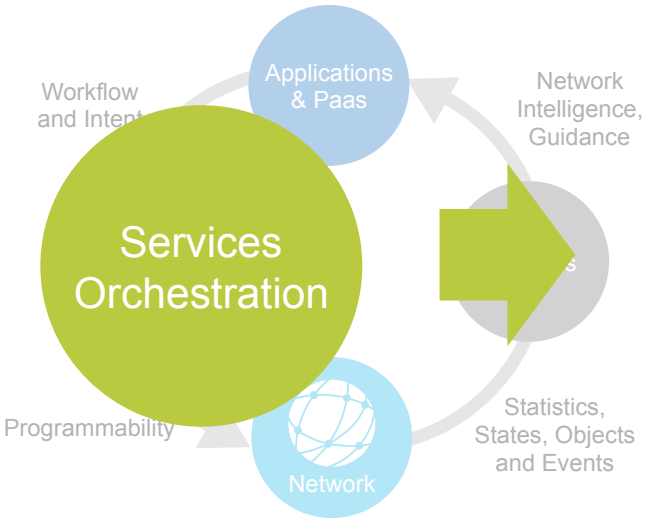

Business Impact

What business impact would you experience if you:

- Could repurpose your infrastructure in minutes by just moving software?
- Didn't have to physically rewire or reconfigure hardware?
- Could create broadcast pop-up channels in minutes instead of months?
- Had an infrastructure which automatically flexed to the different demands of your workloads?

Whole Stack Can't fit in a Everyone's Head?

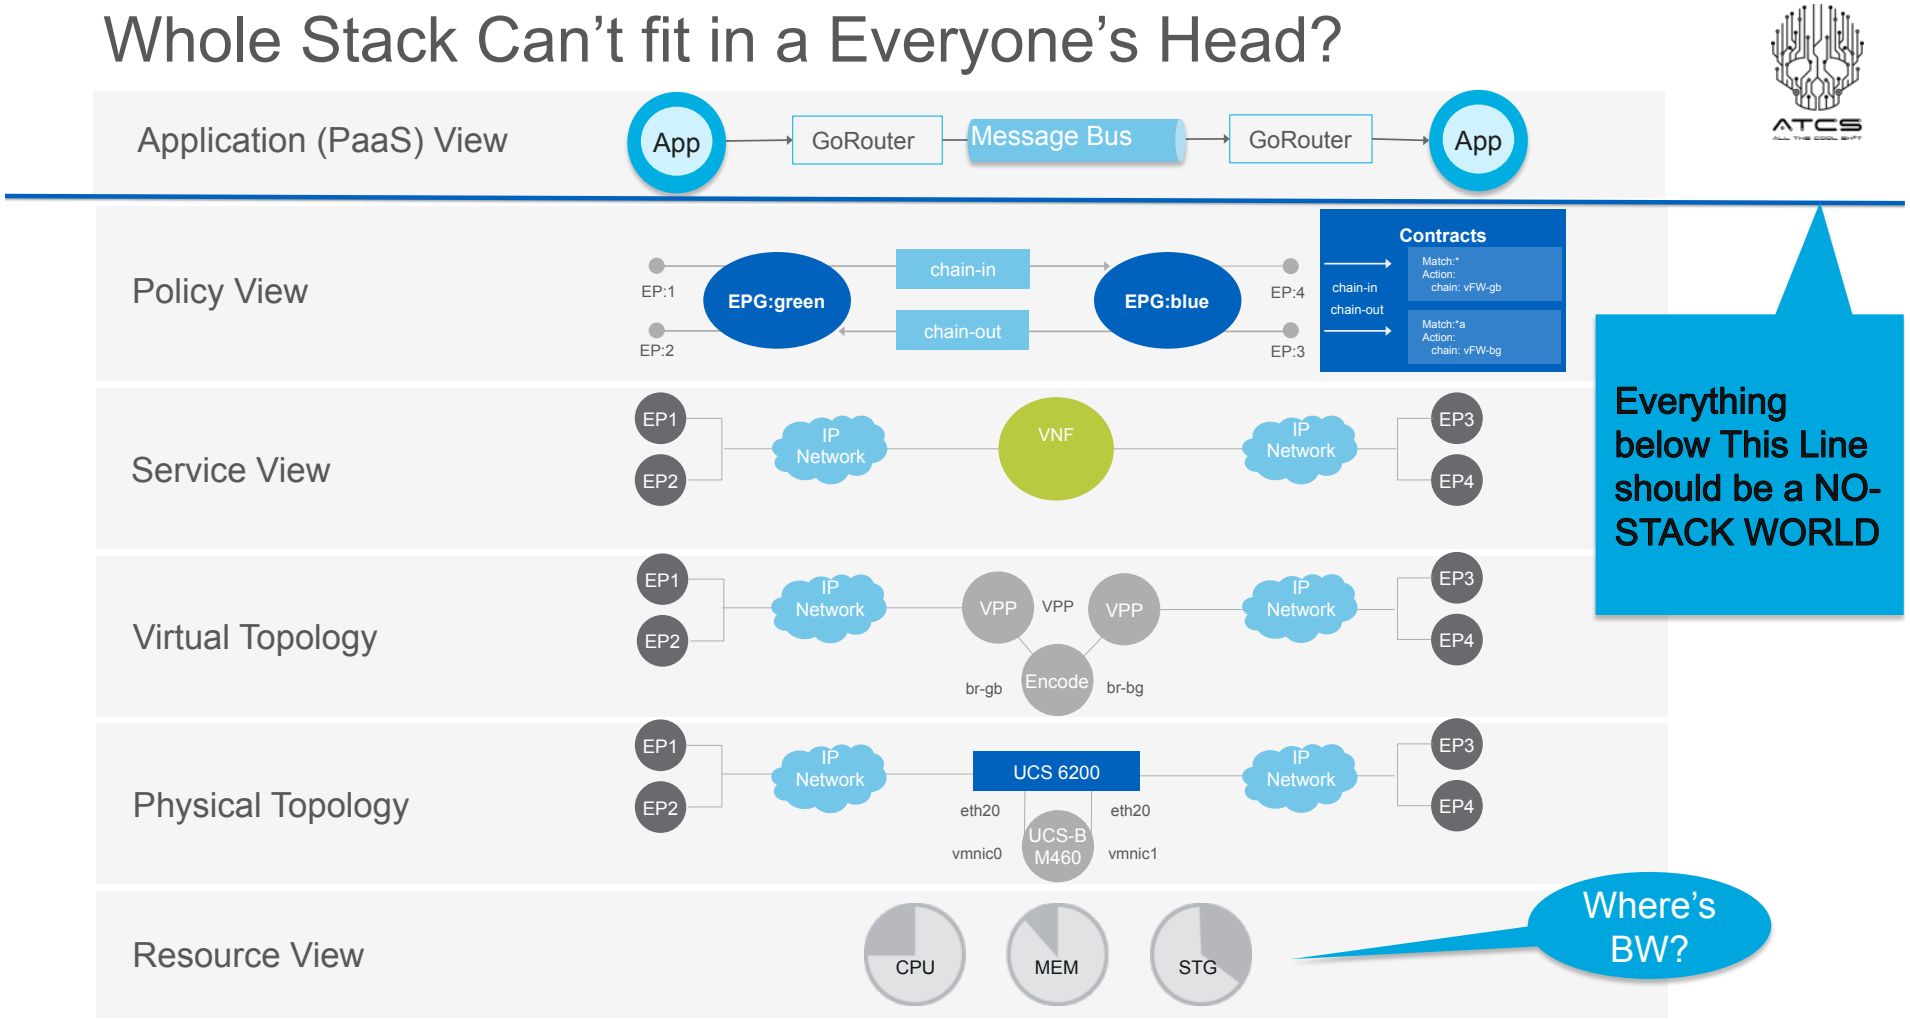

“Below the line”
Service WorkFlow orchestration: Just Make it happen

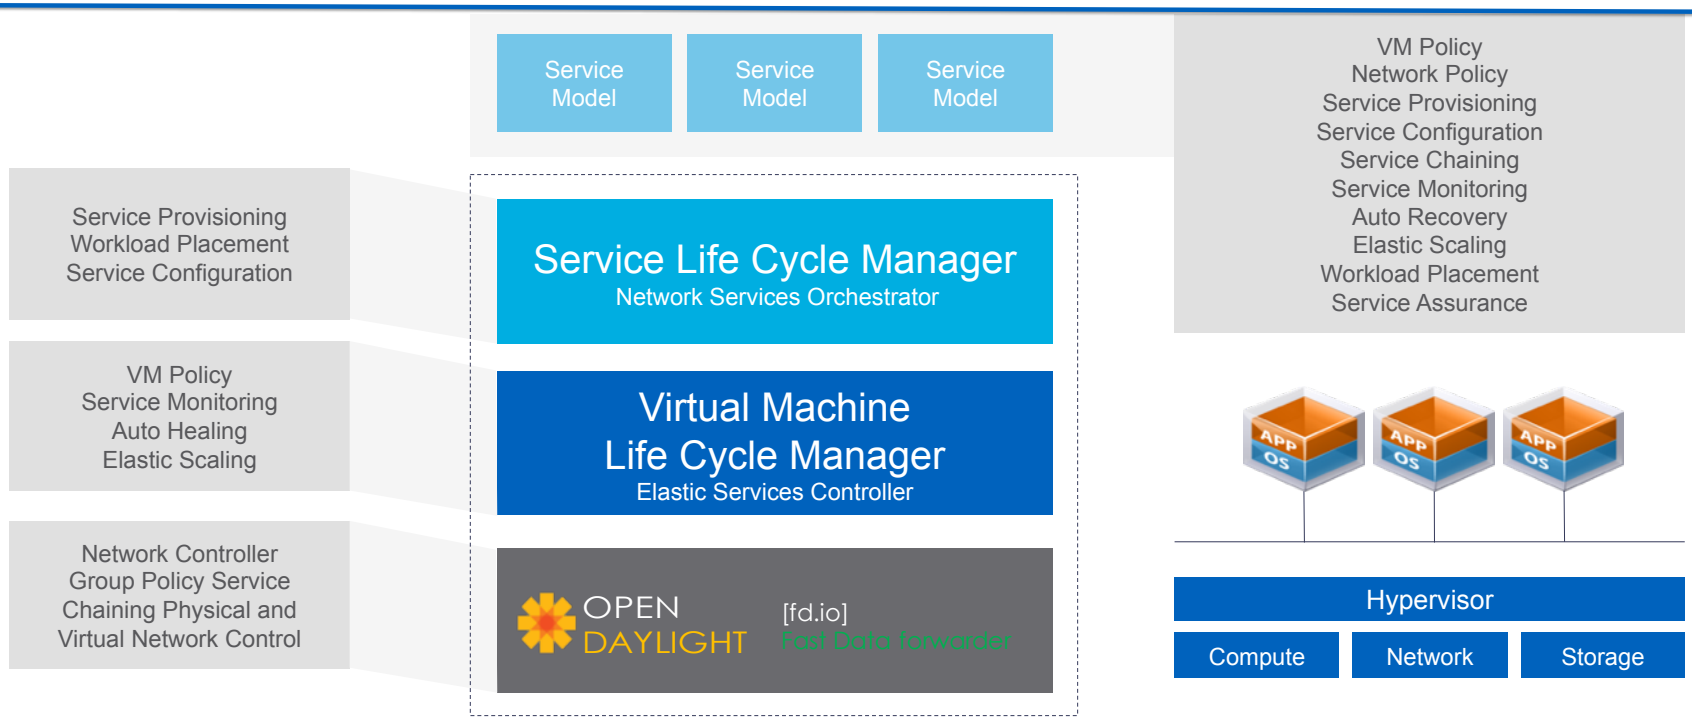

Containerized Virtualized Service Pipelines

Standalone Containerized Media Service with a REST Interface

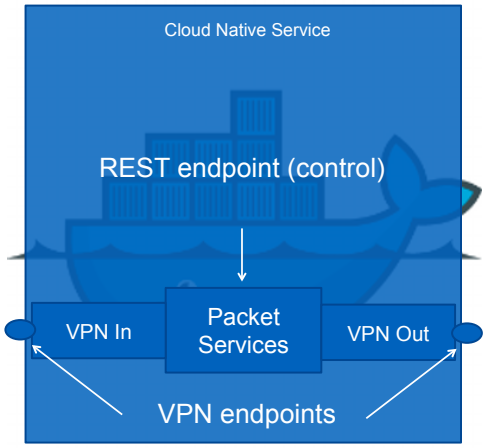

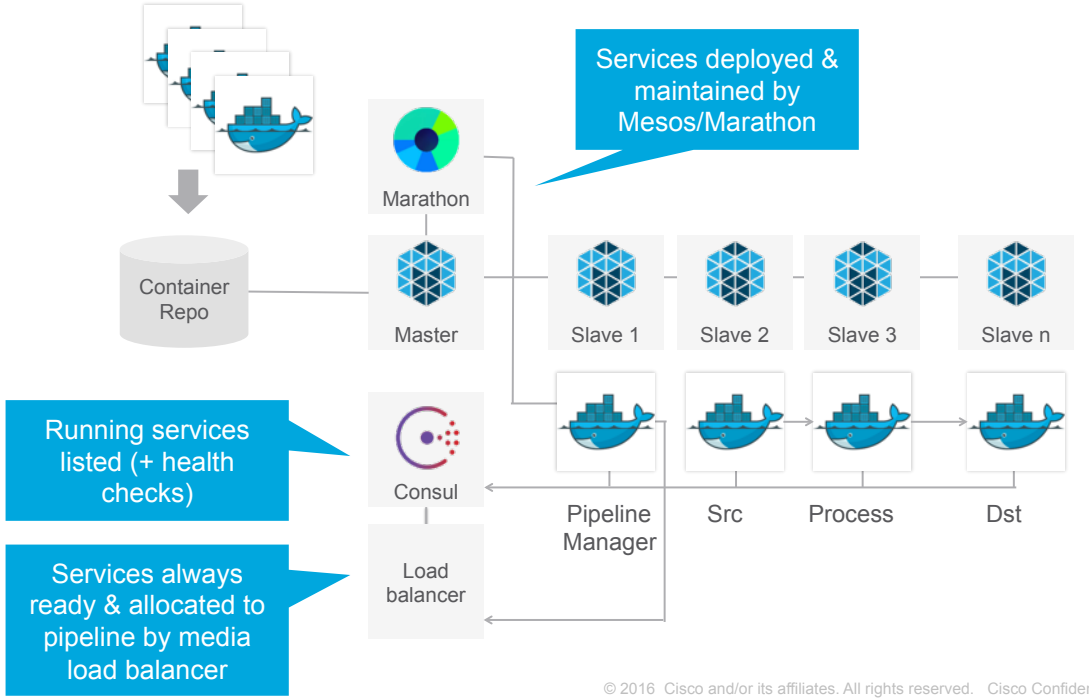

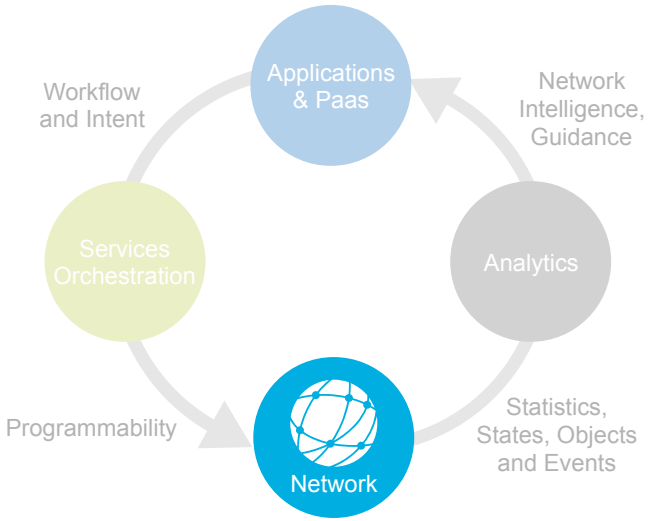

Network: IP Fabric for NFV

Creating the scalable foundation for virtualized services.

IPv4

Operational complexity: Underlays, Overlays, VLANs, NATs, ALGs,…

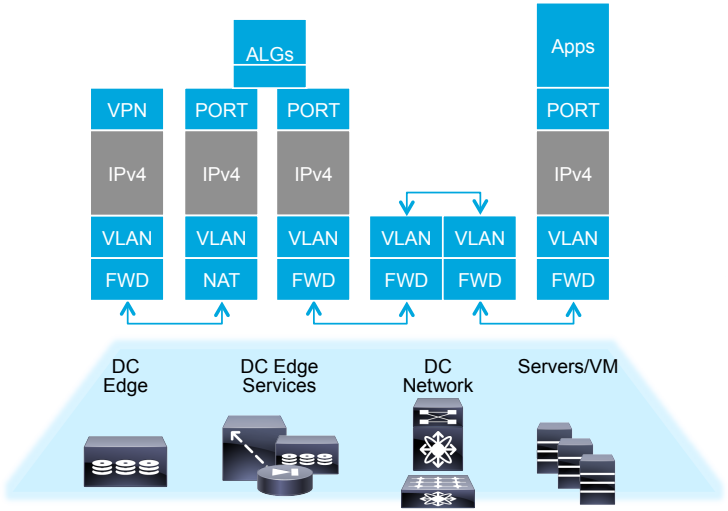

IPv6

IPv6-Centric Networking: Scale, Simplification and Services

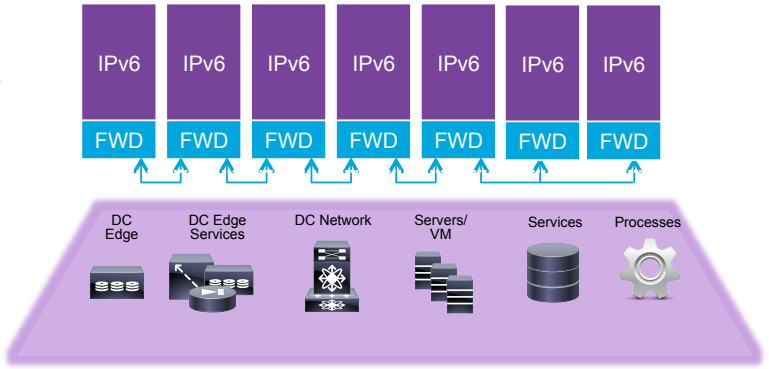

SP Media DC != Enterprise switching

- Live production media is well groomed: Load Engineering
- Learning Media flows and flow bandwidths
- End to End bandwidth reservation/Admission Control
- Quality of Service
- Flow security

Multivendor Network Topologies

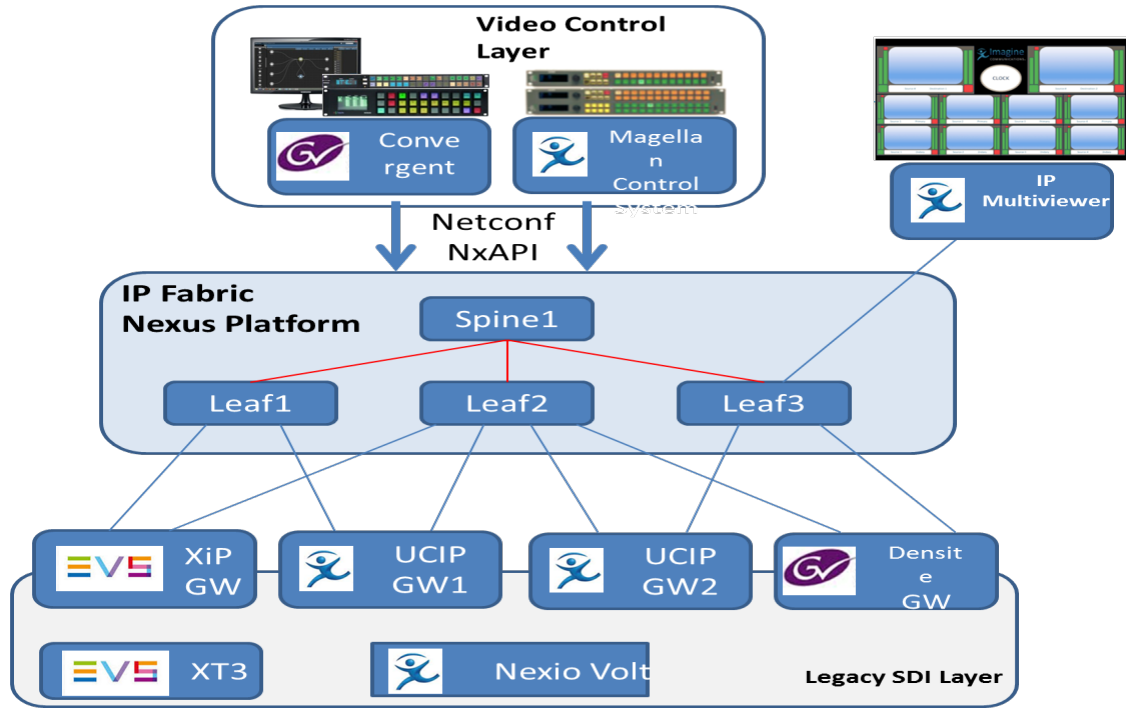

Highlights

- SMPTE 2022-6 Interop with 3 Vendors
- Non blocking 40G IP Fabric
- legacy control Panels to manage Video flows
- Mix of IGMP and SDN solutions

Service Brokered Networking: Guaranteeing WorkFlow Resources

- Rapid and transparent workload migration for VMs, containers, and content
- Dynamic, scalable, in-network workload discovery
- Data path, compute and storage management with integrated media load balancing
- Workflow service pipeline integrated with application engineered data paths
- Guaranteed high network throughput between workers

Uncompressed video data rates

Format	pixels/second	Pixel format	raw bitrate (Gbps)	+ audio & ancilliary data (Gbps)
480i60 (SD)	10 368 000	YUV 4:2:2 10 bit	0.21	0.27
1080p30	62 208 000	YUV 4:2:2 10 bit	1.24	1.49
1080p60 (HD)	12 4416 000	YUV 4:2:2 10 bit	2.49	2.97
4Kp60	497 664 000	YUV 4:2:2 10 bit	9.95	11.88
4Kp120	995 328 000	YUV 4:4:4 16 bit	47.78	
8Kp60	1 990 656 000	YUV 4:2:2 10 bit	39.81	47.52
8Kp120	3 981 312 000	YUV 4:4:4 16 bit	191.10	

SDI Bitrates

Greater bit depth for HDR

IO can keep up

IP Fabric for Media: Getting into Production

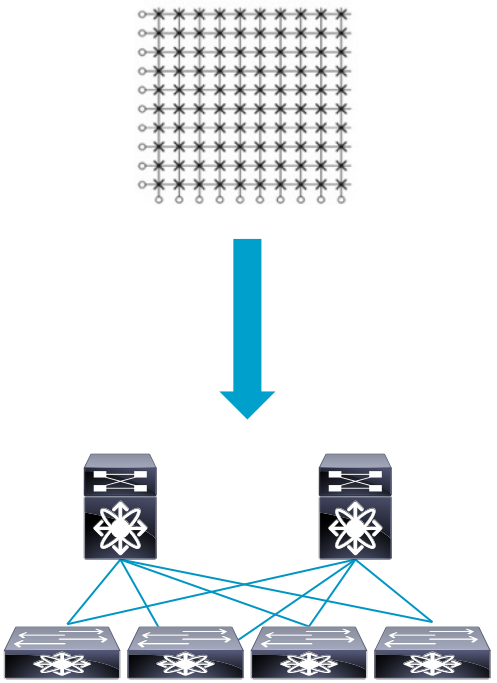

- Flexible Workflow
- Deterministic Low Latency and Jitter
  Deterministic Quality of Service
- Zero Packet Loss
  Reservation of network resources across
  redundant paths for zero congestion loss
- Video/ Audio End Point Sync and Lock with Micro-sec Accuracy
  Precision Timing and Synchronization
- Fast and Clean Switching
  Switching streams with minimal delay and on frame boundary
- System Availability
  Same or better than SDI-based system
- Network Security
  Protect network operations from any malicious attacks

VPP in the Overall Stack

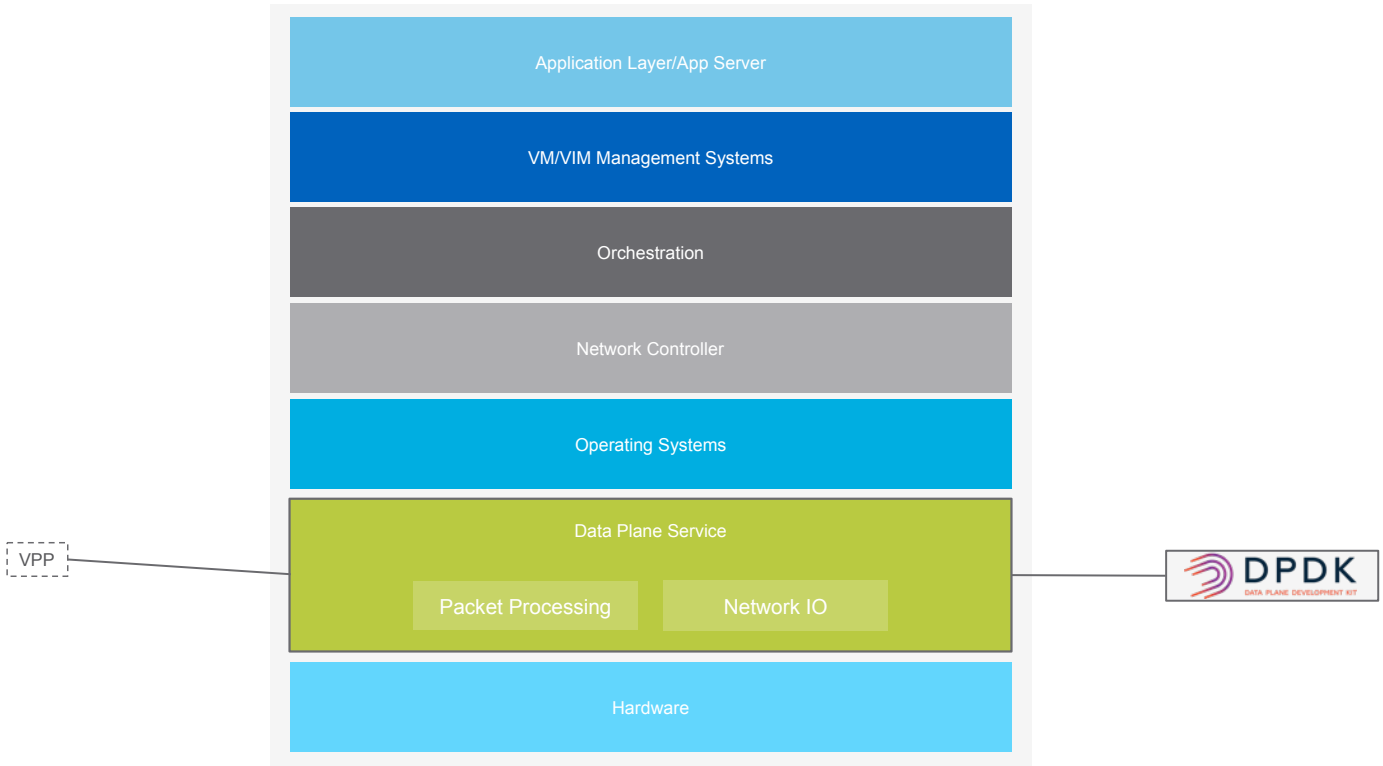

VPP TECHNOLOGY IN A NUTSHELL

- Throughput NDR = “No Drop Rate”
- “No Drop Rate” = Read zero frame loss
  - And not “ITU 10^-7 packet loss ratio”
- Simply put NDR here is better than ”ITU NDR” ☺

LightReading Networking the Communications Industry

EANTC

HOME NEWS & VIEWS FEATURED STORIES AUTHORS MESSAGES VIDEO AUDIO WEBINARS RESEARCH LR

TECHNOLOGY COMPONENTS GIGABIT/BROADBAND MOBILE CABLE OPTICAL ETHERNET/IP DATA

HOT TOPICS NEW IP IoT SECURITY WHITE BOX SERVICES BUSINESS WOMEN IN COMMS TELCO T

GIGABIT LTE" "ENEA ACQUIRES CENTERED LOGIC FOR NETWORK FUNCTIONS MANAGEMENT" "CELENO UNVEILS NE

NFV / NFV TESTS & TRIALS

Validating Cisco's NFV Infrastructure Pt. 1

As part of the New IP revolution, NFV (network functions virtualization) is going to change the world of communications networking and services. It's just a matter of when, not if.

Part of the journey towards the "cloudification" of wide area communications networks involves the development of coordinated hardware and software systems that will enable the creation, delivery, management and tear-down of virtual network functions (VNFs). Collectively, those systems are referred to as the NFV infrastructure (NFVi).

PRIME READING
LIGHT READING
10/12/2015

LINUX FOUNDATION
COLLABORATIVE PROJECTS

| NDR rates for 2p10GE, 1 core, L2 NIC-to-NIC

[IMIX Gbps]

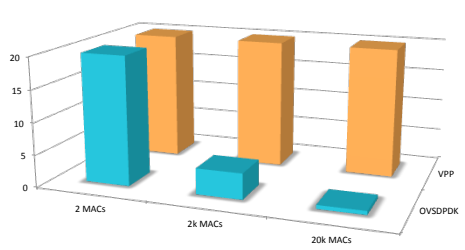

| NDR rates for 2p10GE, 1 core, L2 NIC-to-VM/VM-to-VM

[IMIX Gbps]

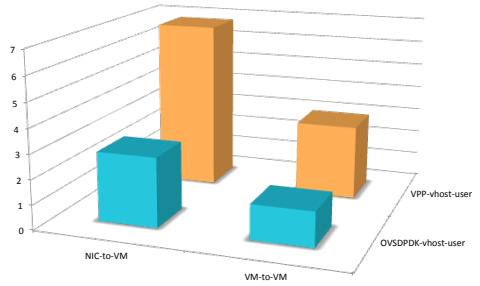

- Virtual network infrastructure – benchmarking the efficiency
- All tests with minimal connections only, single CPU core
- Higher performance with more connections, more CPU cores
- Latest SW at the time: OVSDPDK 2.4.0, VPP 09/2015

VNET-SLA BENCHMARKING AT SCALE: IPV6, V4, MAC

VPP-based vSwitch

Phy-VS-Phy

Zero-packet-loss Throughput for 12 port 40GE, 24 cores, IPv6

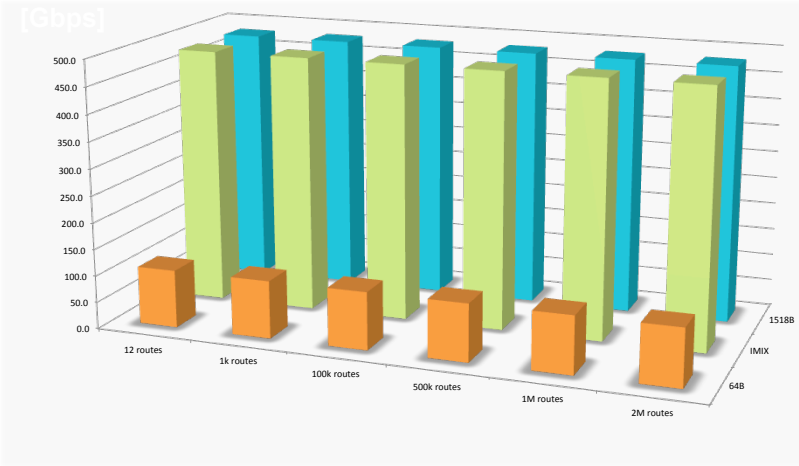

Zero-packet-loss Throughput for 12 port 40GE, 24 cores, IPv6

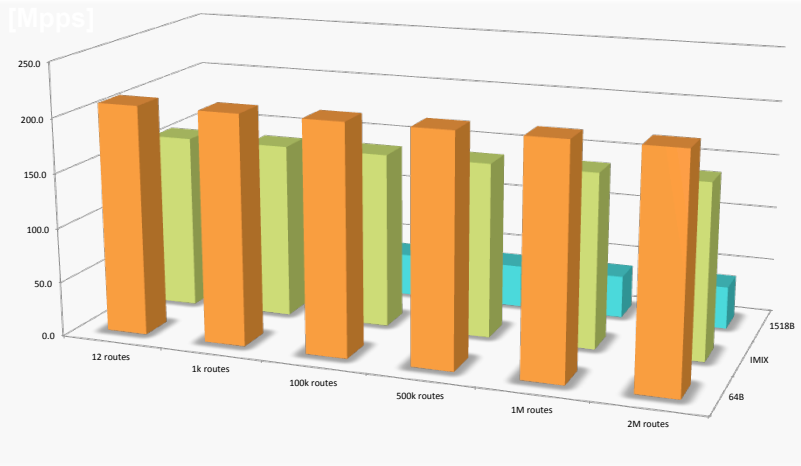

- FD.io VPP data plane throughput not impacted by large size of IPv6 FIB
- VPP tested on UCS 4-CPU-socket server with 4 of Intel “Haswell" x86-64 processors E7-8890v3 18C 2.5GHz
- 24 Cores used – Another 48 cores can be used for other network services!

VPP vSwitch IPv6 routed forwarding FIB with 2 milion IPv6 entries 12x40GE (480GE) 64B frames

200Mpps zero frame loss

NIC and PCIe is the limit not VPP

VPP vSwitch IPv6 routed forwarding FIB with 2 milion IPv6 entries

12x40GE (480GE) IMIX frames

480Gbps zero frame loss

“Sky” is the limit not VPP

VNET-SLA BENCHMARKING AT SCALE: IPV4 + SECURITY

VPP-based vSwitch

Zero-packet-loss Throughput for 18 port 40GE, 36 cores, IPv4

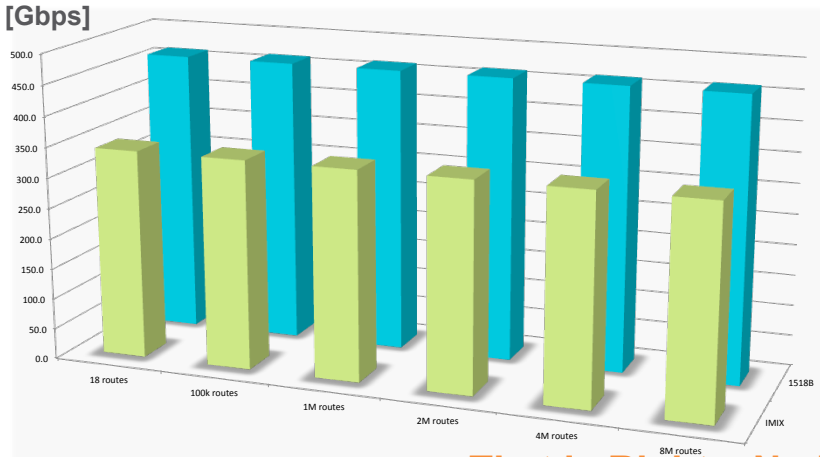

Zero-packet-loss Throughput for 18 port 40GE, 36 cores, IPv4

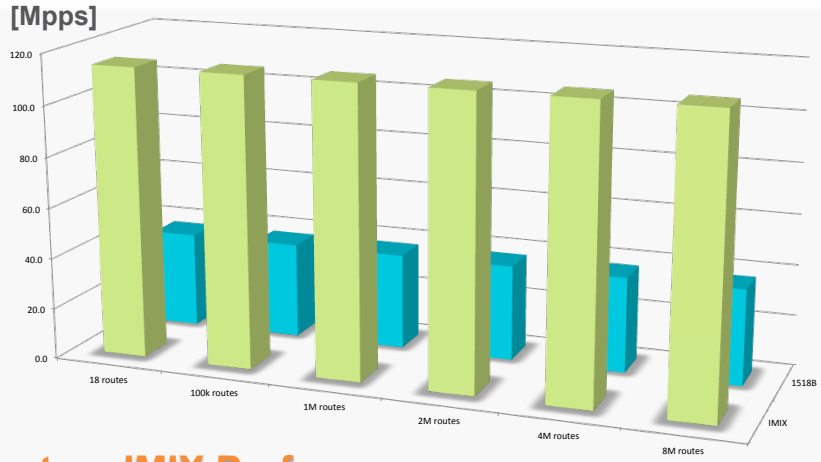

That is Right – No Impact on IMIX Performance

- FD.io VPP data plane throughput not impacted by large IPv4 FIB size, and stateless security input access lists and white-lists & black-lists*
- VPP tested on UCS 4-CPU-socket server with 4 of Intel "Haswell" x86-64 processors E7-8890v3 18C 2.5GHz
- **36 Cores used – Another 36cores can be used for other network services!**

VPP Cores Not Completely Busy
VPP Vectors Have Space For More Services and More Packets!!
PCIe 3.0 and NICs Are The Limit

VPP vSwitch IPv4 routed forwarding FIB with 2M IPv4 entries, 2k white-list* 18x40GE IMIX frame sequence

342Gbps zero frame loss

Sky" is the limit not VPP

VPP vSwitch IPv4 routed forwarding FIB with 2M IPv4 entries, 2k white-list* 18x40GE 1518B frames

462Gbps zero frame loss

Pcle 3.0 is the limit not VPP

* Test results extrapolated from 120GE 2-Socket machine. See next slides for why.

VPP Features

IPv4/IPv6

14+ MPPS, single core
Multimillion entry FIBs
Source RPF
Thousands of VRFs
Controlled cross-VRF lookups
Multipath – ECMP and Unequal Cost
Multiple million Classifiers –
Arbitrary N-tuple
VLAN Support – Single/Double tag

Counters for Everything
Mandatory Input Checks:
- TTL expiration
- header checksum
- L2 length < IP length
- ARP resolution/snooping
- ARP proxy

IPv4

GRE, MPLS-GRE, NSH-GRE,
VXLAN
IPSEC
DHCP client/proxy
CG NAT

IPv6

Neighbor discovery
Router Advertisement
DHCPv6 Proxy
L2TPv3
Segment Routing
MAP/LW46 – IPv4aas
iOAM

L2

VLAN Support
Single/ Double tag
L2 forwarding with EFP/Bridge Domain concepts
VTR – push/pop/Translate (1:1,1:2, 2:1,2:2)
Mac Learning – default limit of 50k addresses
Bridging – Split-horizon group support/EFP Filtering
Proxy Arp
Arp termination
IRB – BVI Support with Router Mac assignment
Flooding
Input ACLs
Interface cross-connect

MPLS

MPLS-o-Ethernet – Deep label stacks supported

LINUX FOUNDATION
COLLABORATIVE PROJECTS

cisco

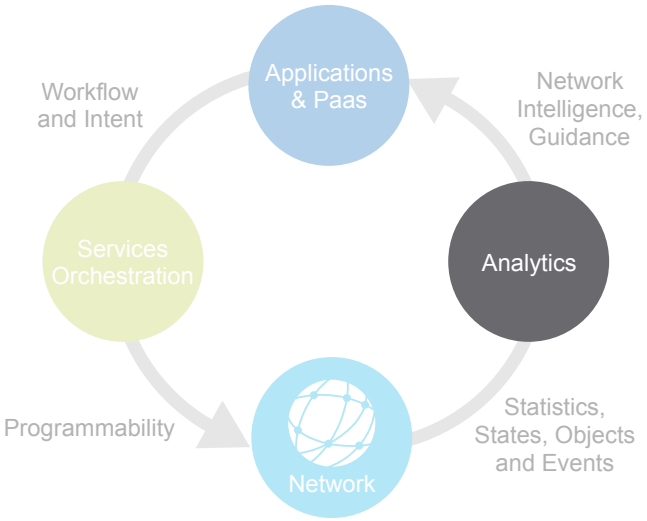

Incident Impact Analysis, Workflow Analytics, Application Engineered.

WorkFlow ANALYTICS IS A CF

- Analytics Today addressed by performing a query function against the entire data set = PROBLEM
  - Fault management = f(event data)
  - Performance management = f(metric data)
  - Service mediation = f(event data, metric data)
  - Capacity management = f(metric data)
  - Security analytics = f(metric data, route data)

Platform For Data Analytics

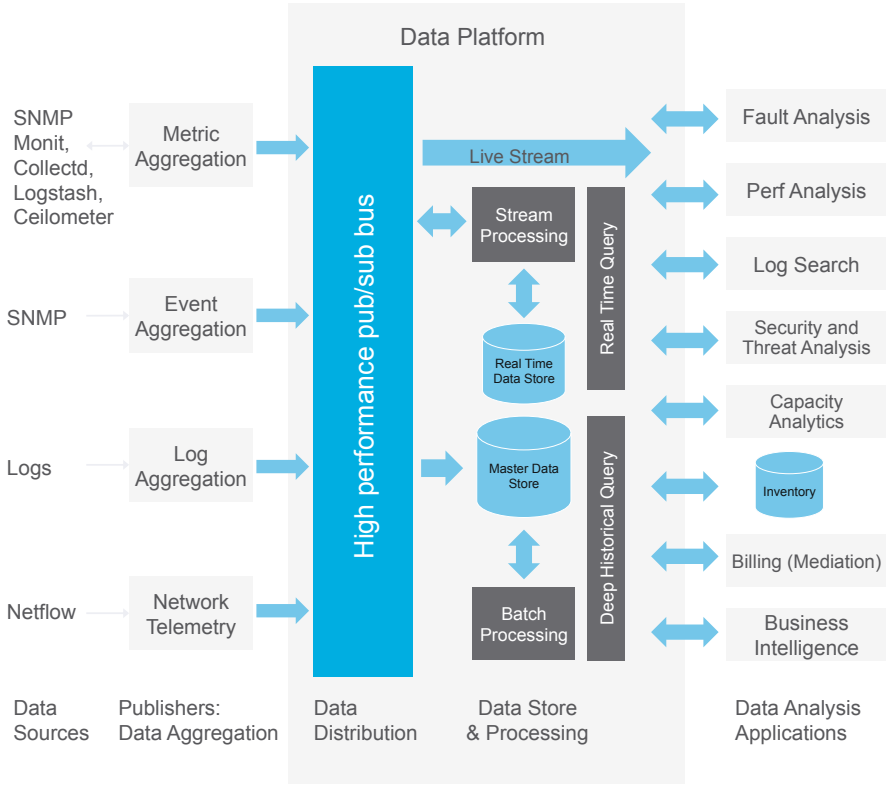

Principles

- Decouple data aggregation (publishers) from data analysis (consumers) – allow any OSS app the potential to access any data source
- Simple, scalable, open data distribution platform
  - Scale-out architecture with support for horizontal scale in all core components
  - Very highly available core platform
  - Low and predictable latency
- Immutable Dataset
  - All data stored raw
  - Minimal filtering/processing on ingress
- Minimal filtering/processing on ingress
- Analytics based approach to analysis functions
- Support for streaming apps, real-time queries and batch processing

Platform For Data Analytics

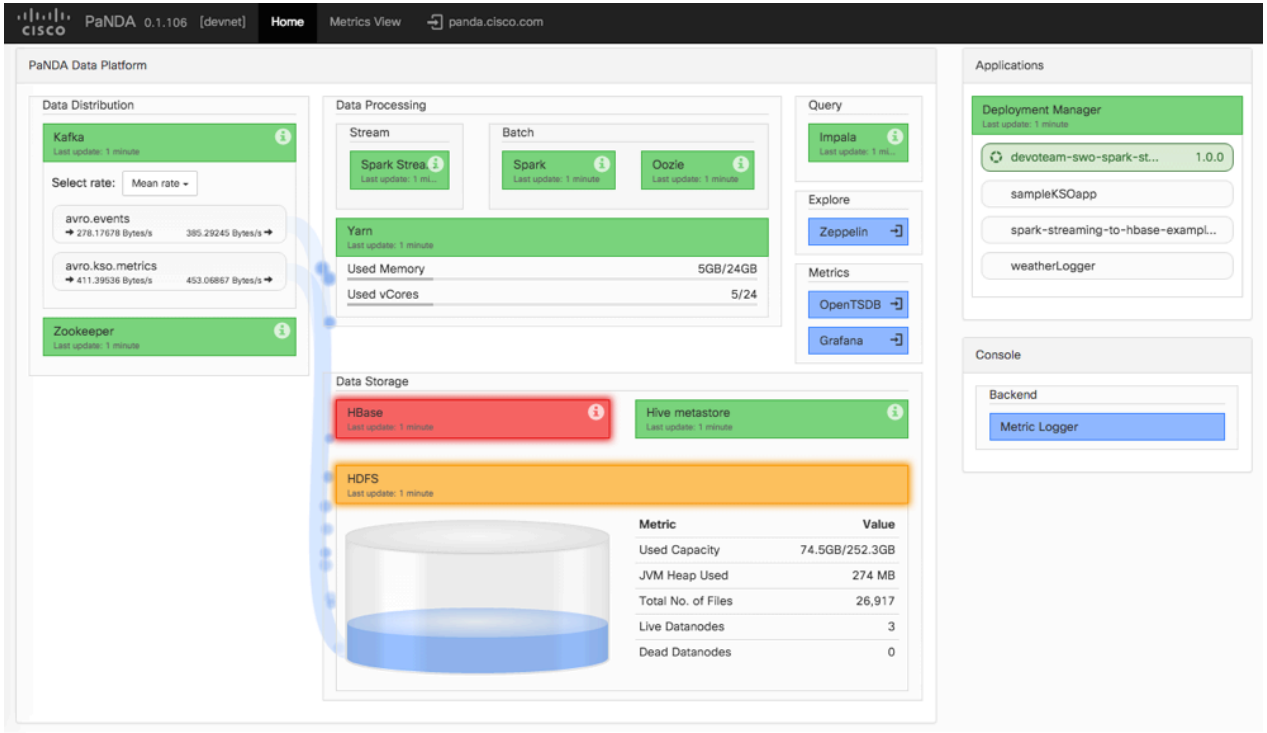

WorkFlow Service Impact Analysis

- Data analytics Platform
- Full-stack data instrumentation : BGP routing, telemetry, syslog, Netflow, Workflow
- Syslog analysis based on host | event frequency | event repetition
- Correlate data sources to identify root cause of service-impacting events
- Use this insight for platform service assurance e.g. optimal codec bitrate for max QoS and resource engineering
- Enable “Reactivity”

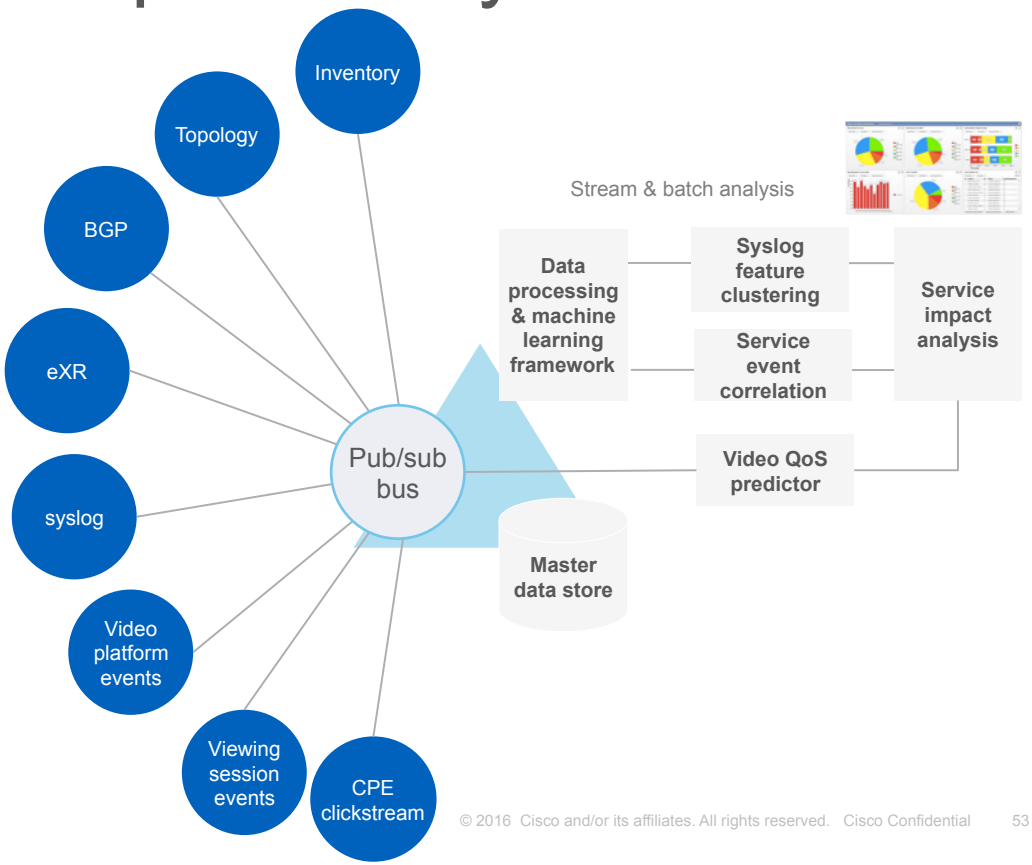

Putting the Pieces together

Applications/
Business Outcomes

Model Driven
Orchestration

Cloud
Platform

Compute, Network,
Storage Infrastructure

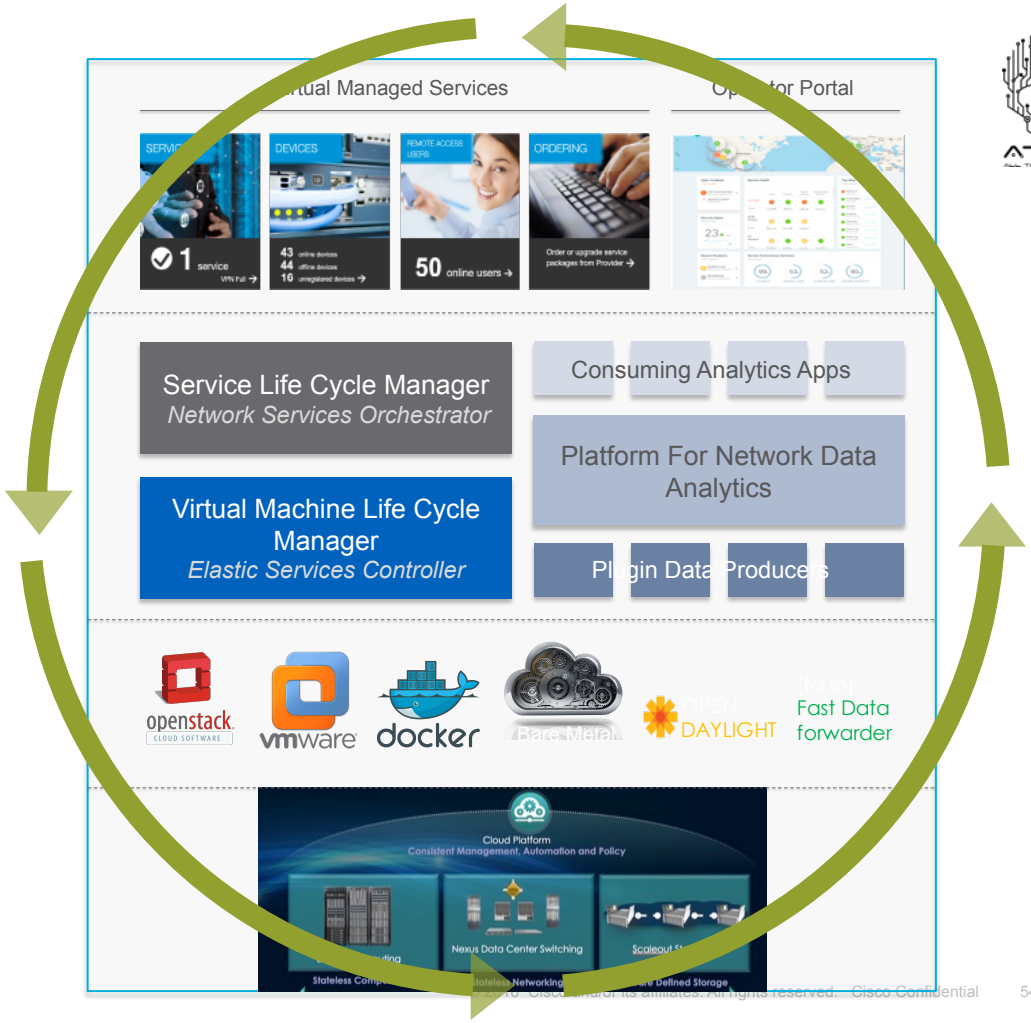

NETWORKING FOCUSED ON LIMITED DEV COMMUNITY AND APPLICABILITY

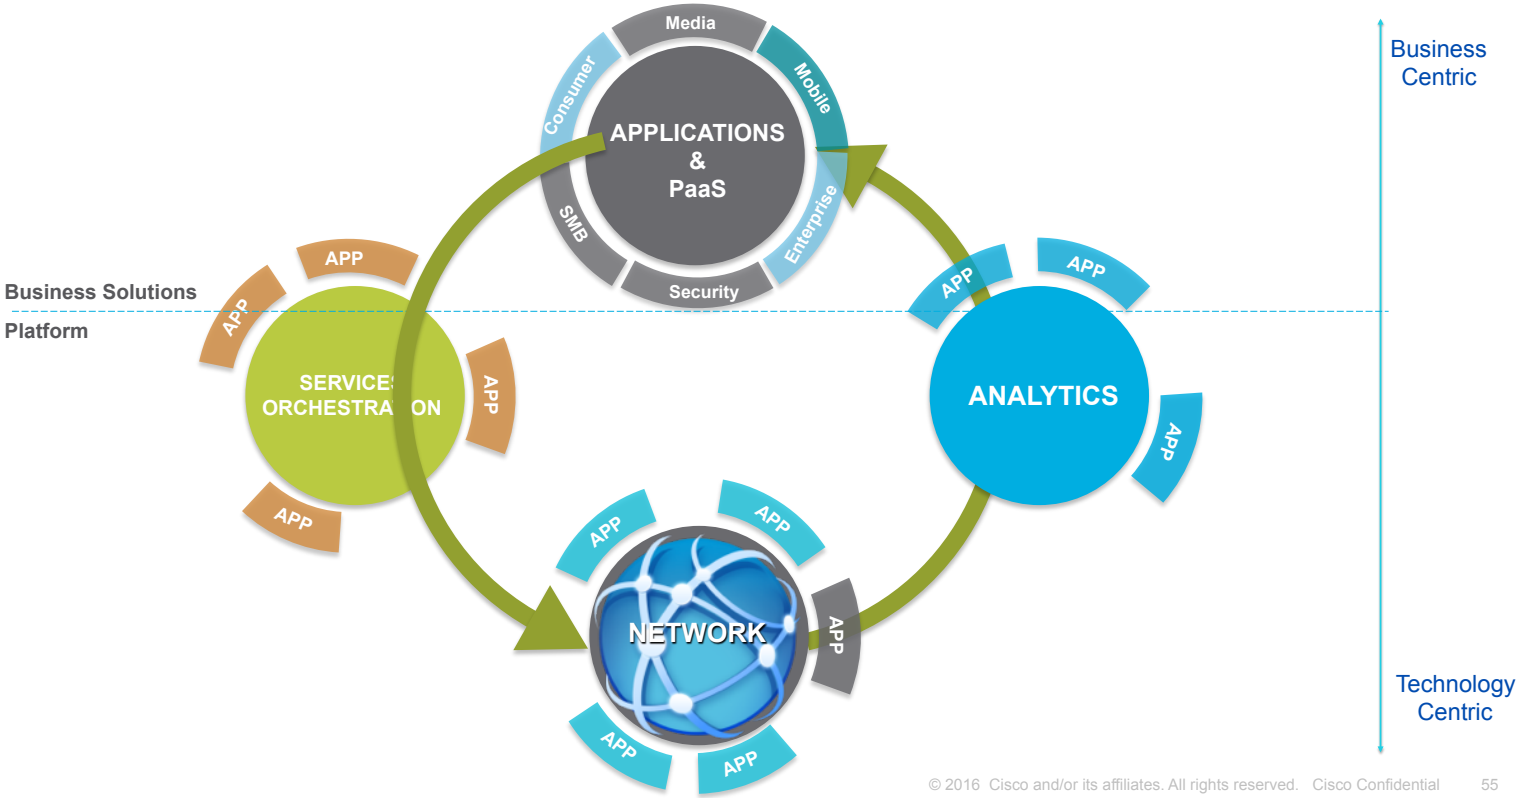

NEW COMMUNITY = NOSTACKDEVELOPER

| Turning data into information to drive intent

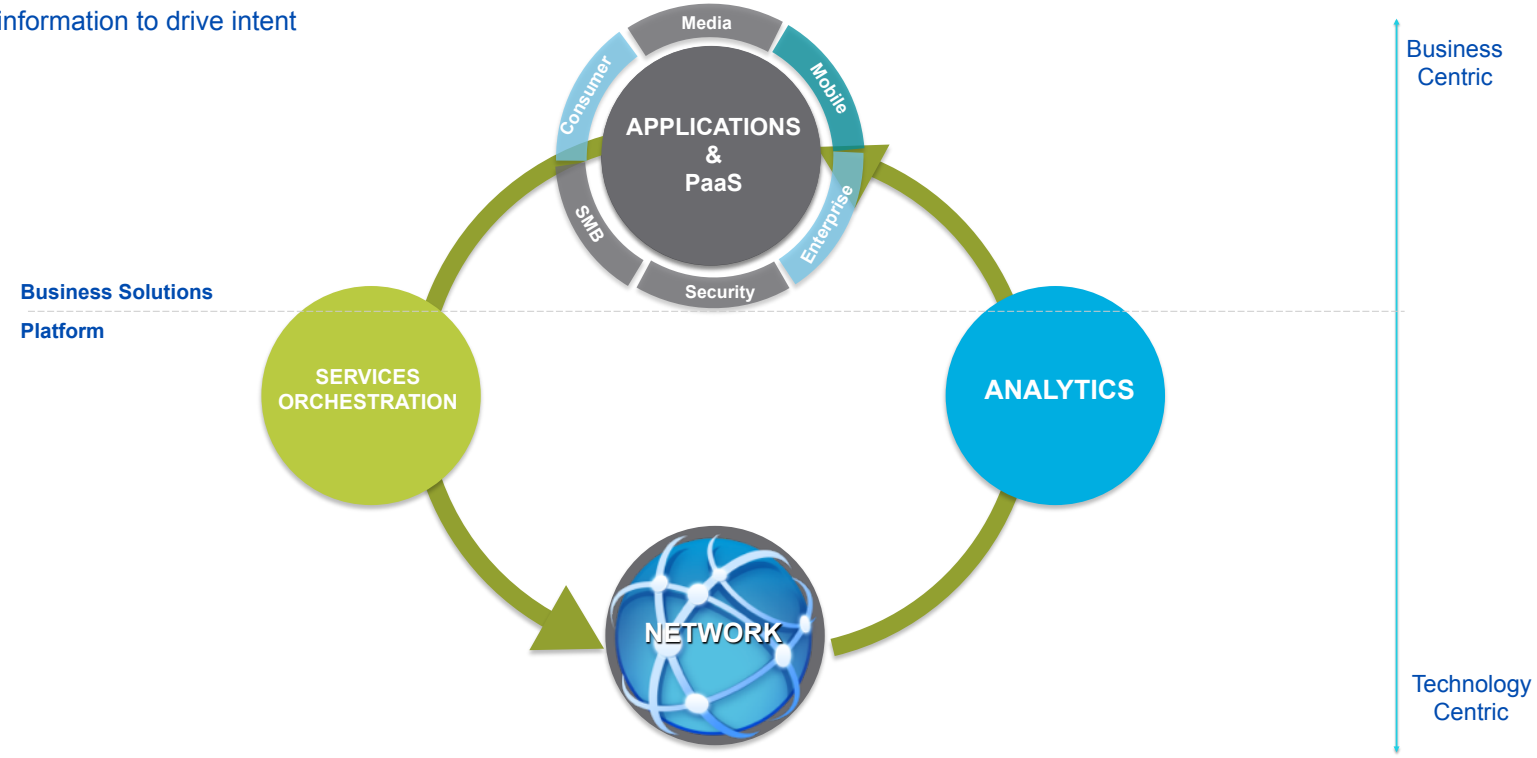

WORKFLOW -> REACTIVE SERVICE BROKER

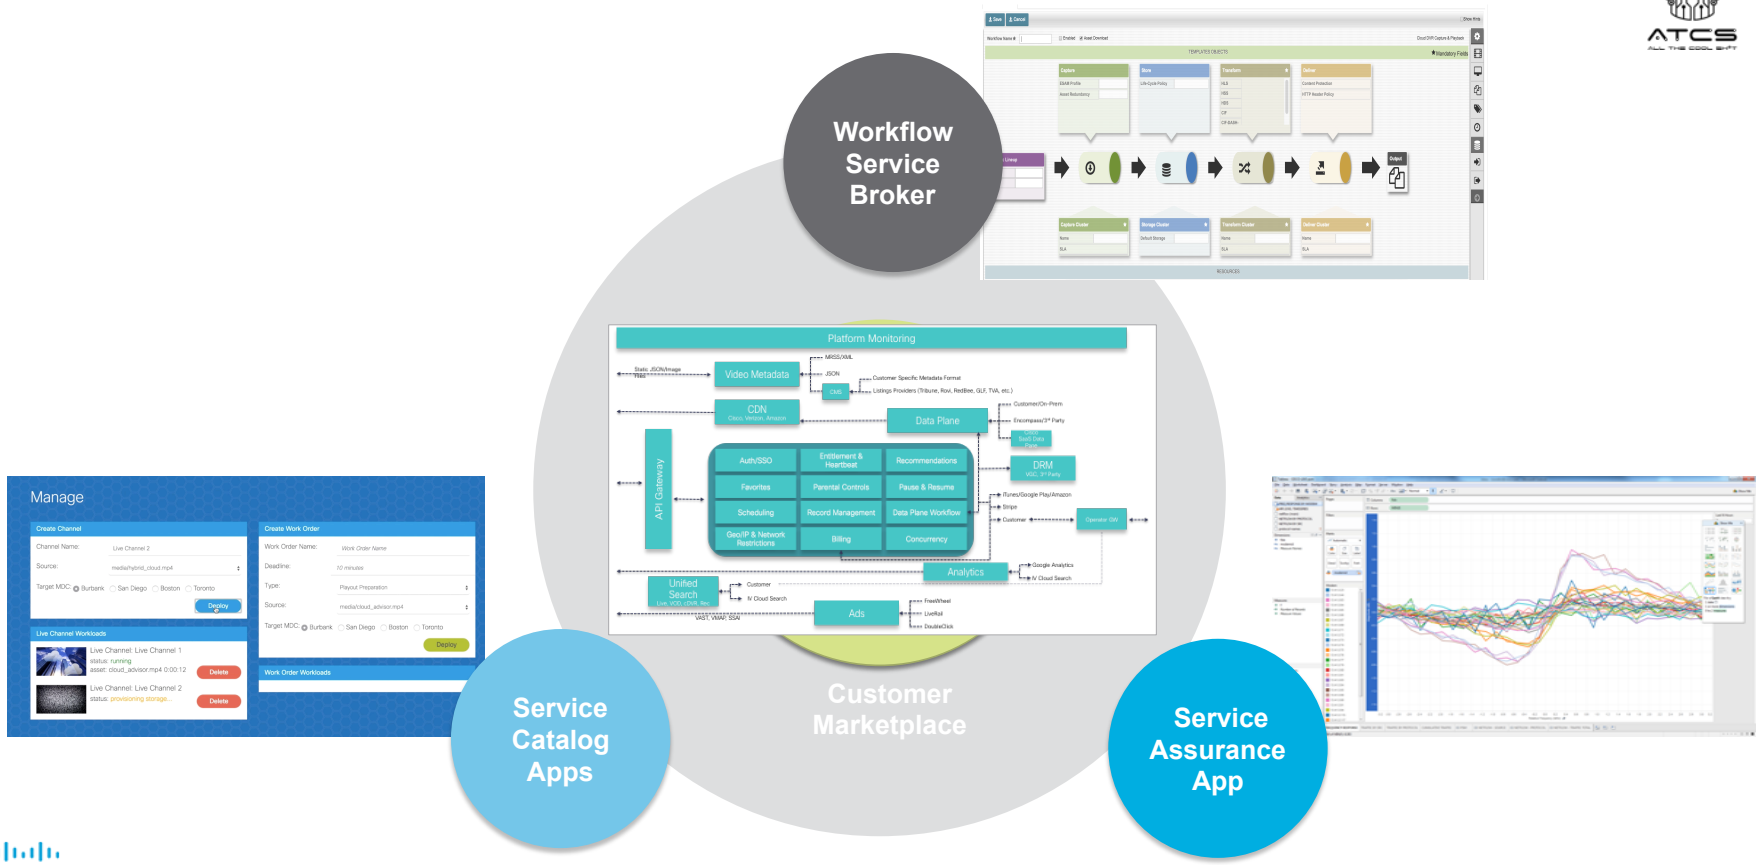

Enabling the SP Media Service Broker

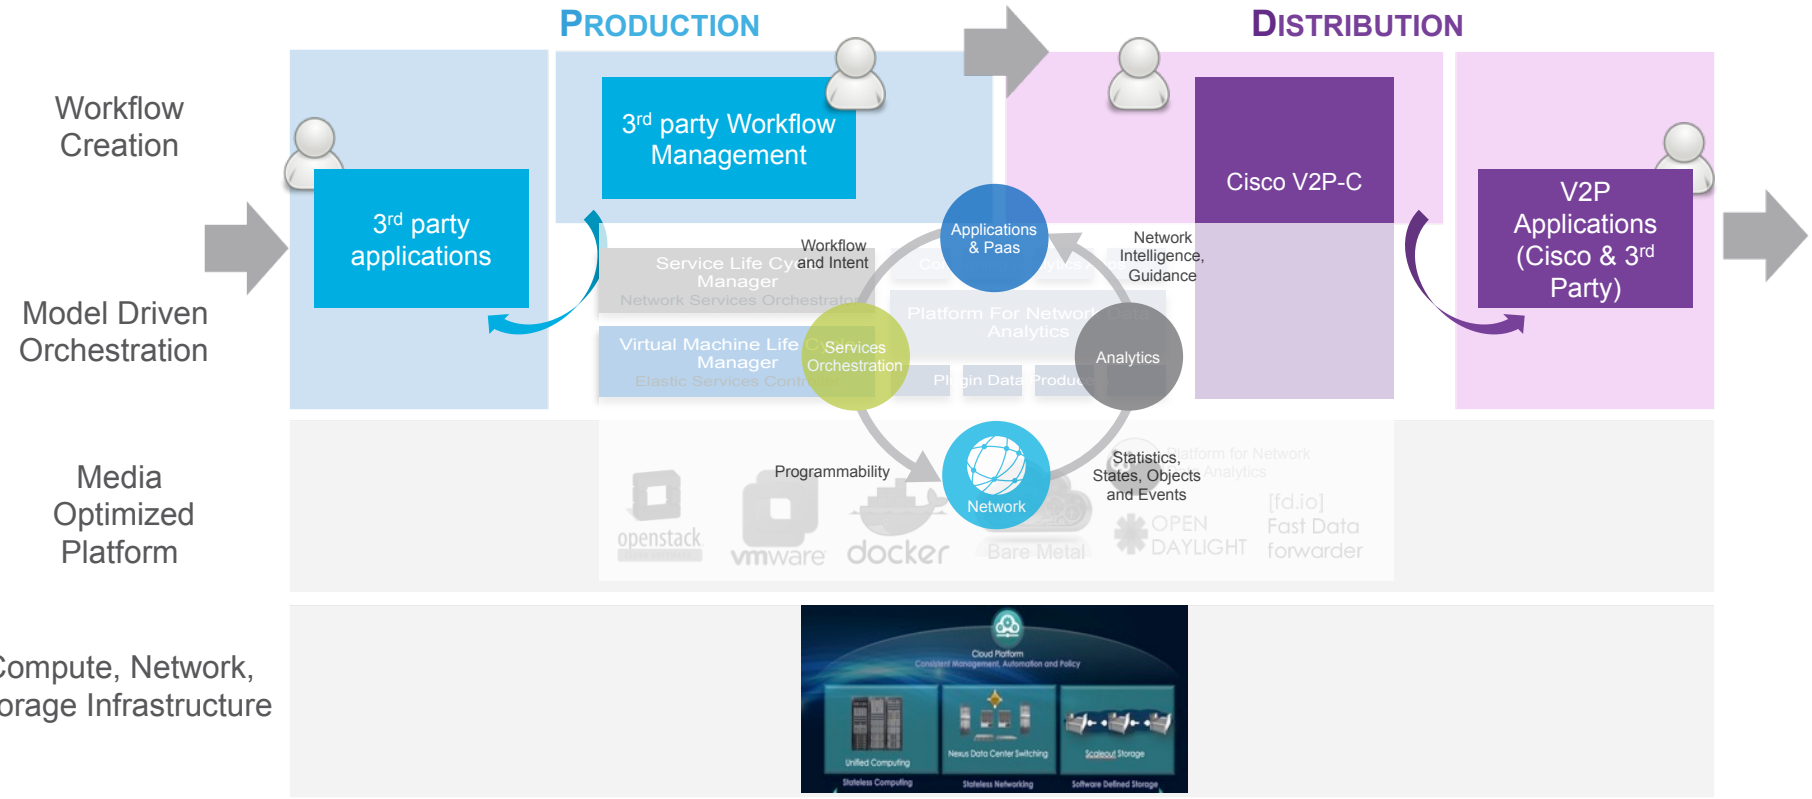

Networking Open Source Project Participation ~ 2 years

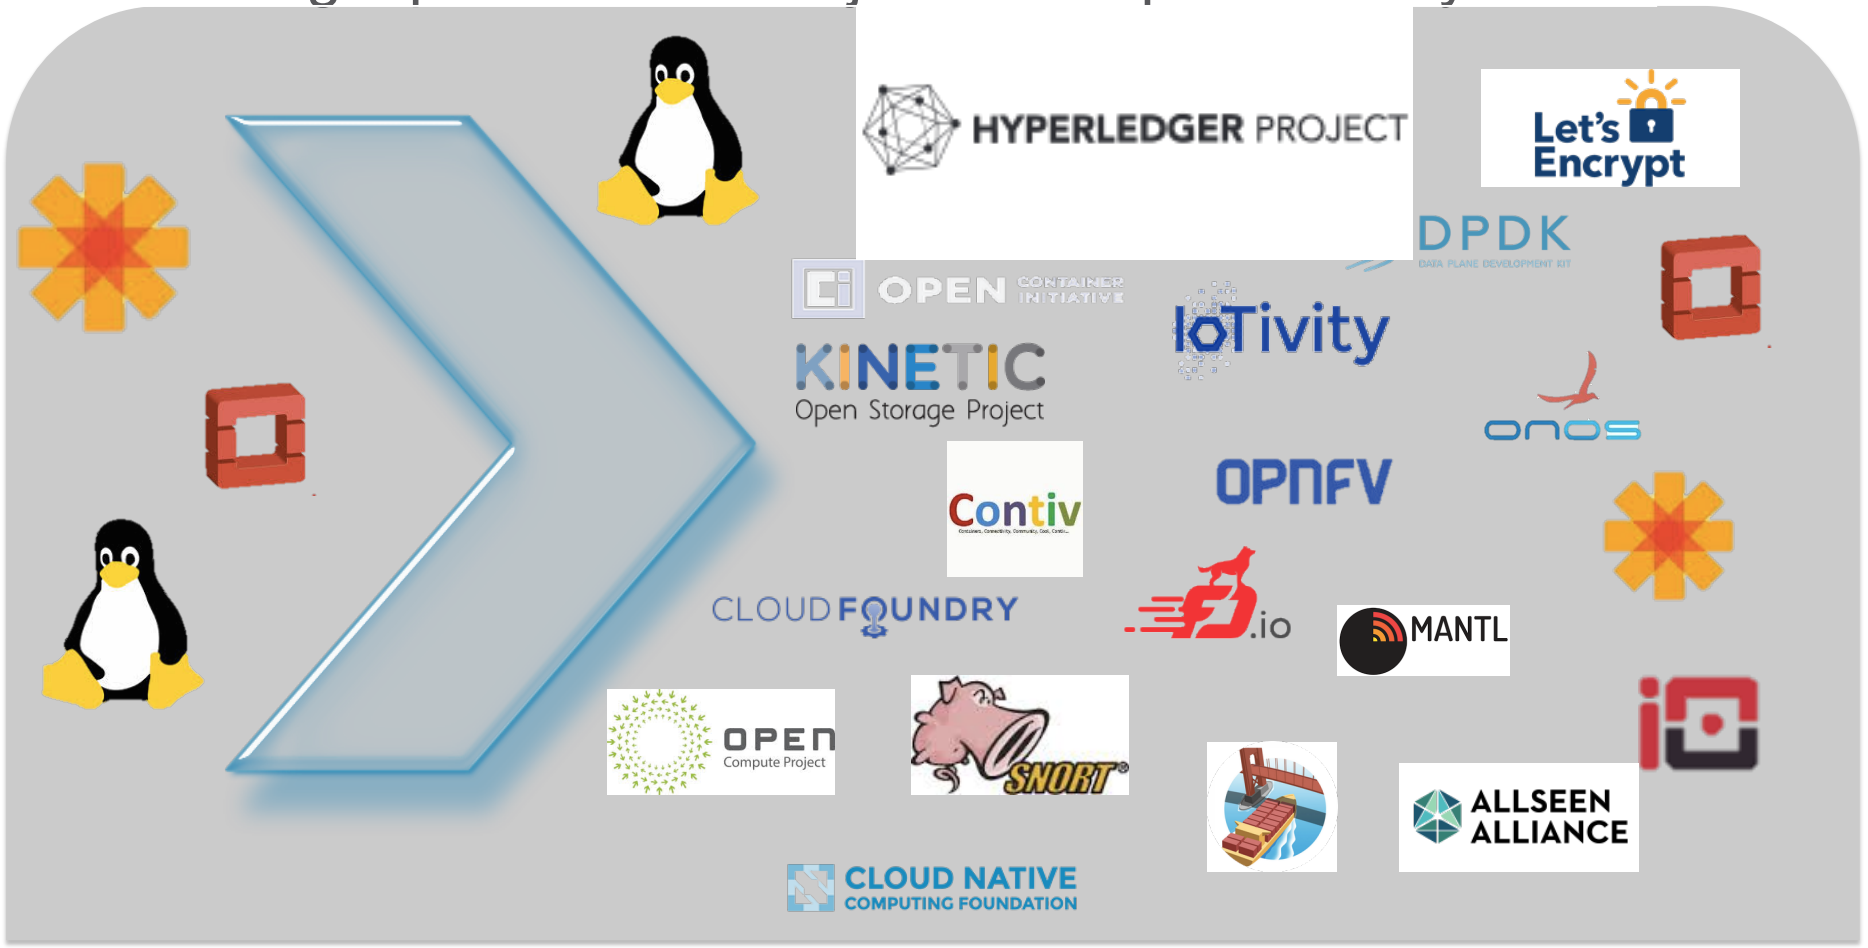

Measuring Success

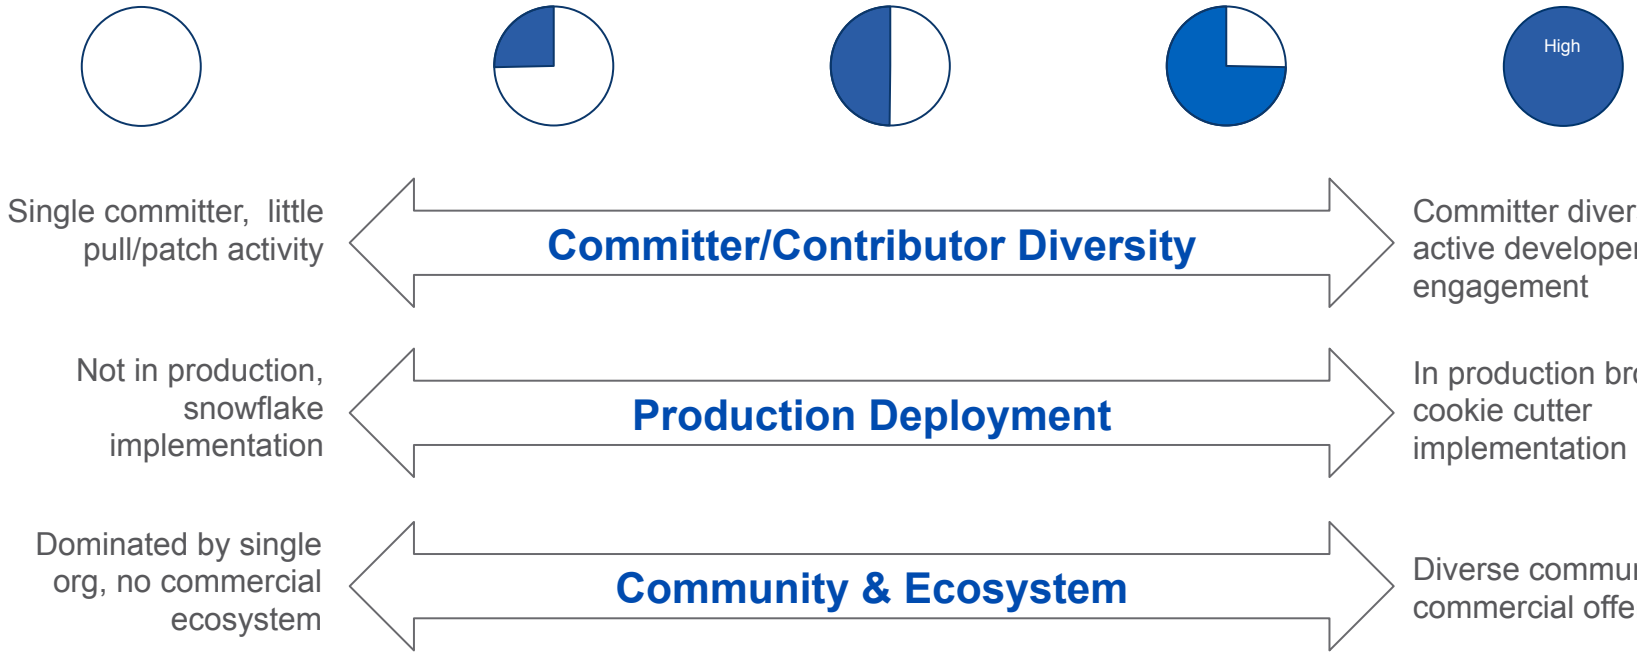

Massive Growth in Project Engagement

High

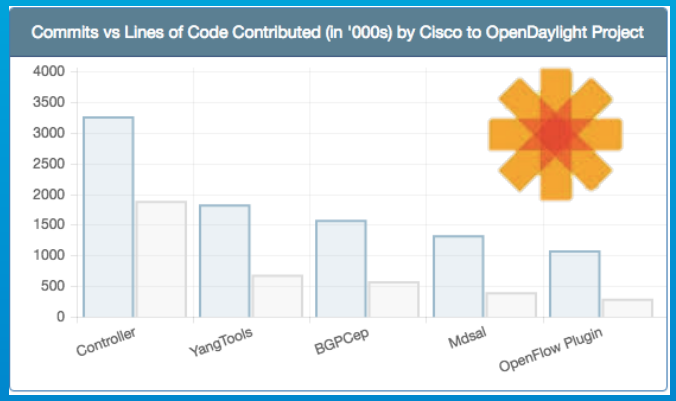

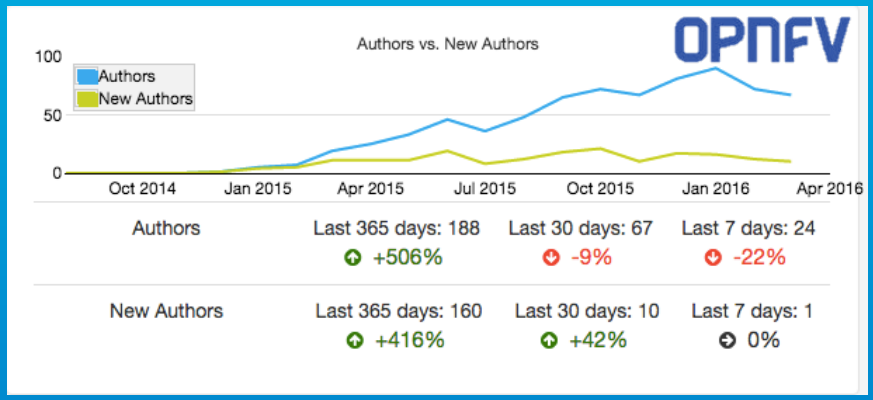

In a Nutshell, FD.io...

... has had 679 commits made by 60 contributors
representing 246,524 lines of code

... is mostly written in C
with an average number of source code comments

... has a codebase with a very short history
maintained by a very large development team
with stable Y-O-Y commits

... took an estimated 65 years of effort (COCOMO model)
starting with its first commit in November, 2015
ending with its most recent commit 3 days ago

In a Nutshell, OpenStack...

... has had 171,602 commits made by 4,009 contributors
representing 1,751,287 lines of code

... is mostly written in Python
with an average number of source code comments

... has a well established, mature codebase
maintained by a very large development team
with stable Y-O-Y commits

... took an estimated 499 years of effort (COCOMO model)
starting with its first commit in December, 2006
ending with its most recent commit 5 days ago

Early 2015: Open Source Umbrella Architecture

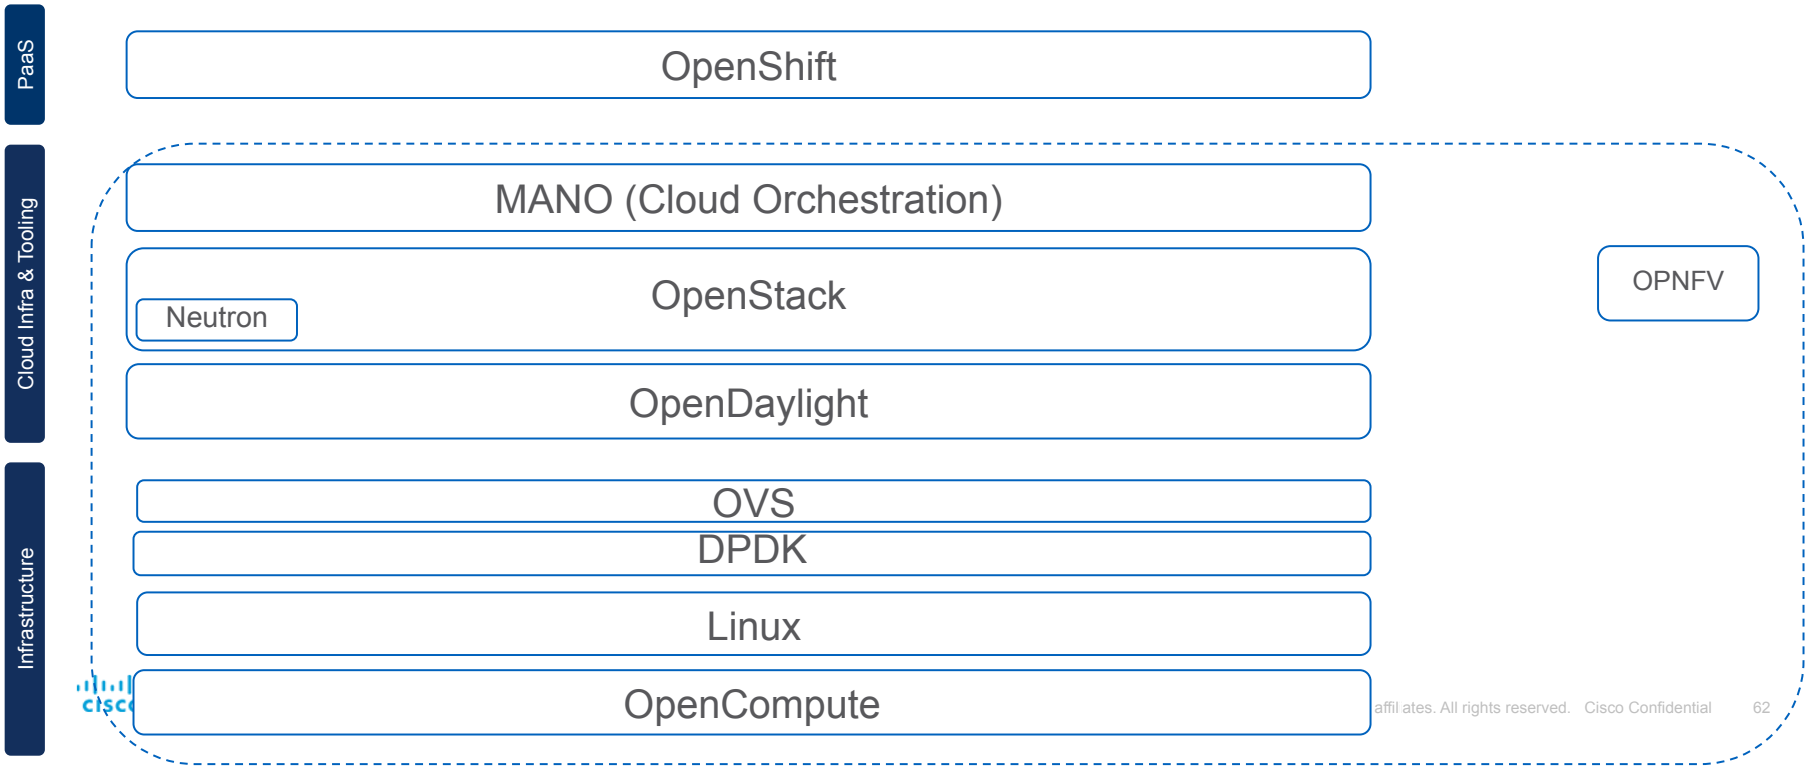

2015-2016: Open Source Umbrella Architecture

~ 8 New Open Source Projects Under the Linux Foundation

*New to LF in 2015/2016

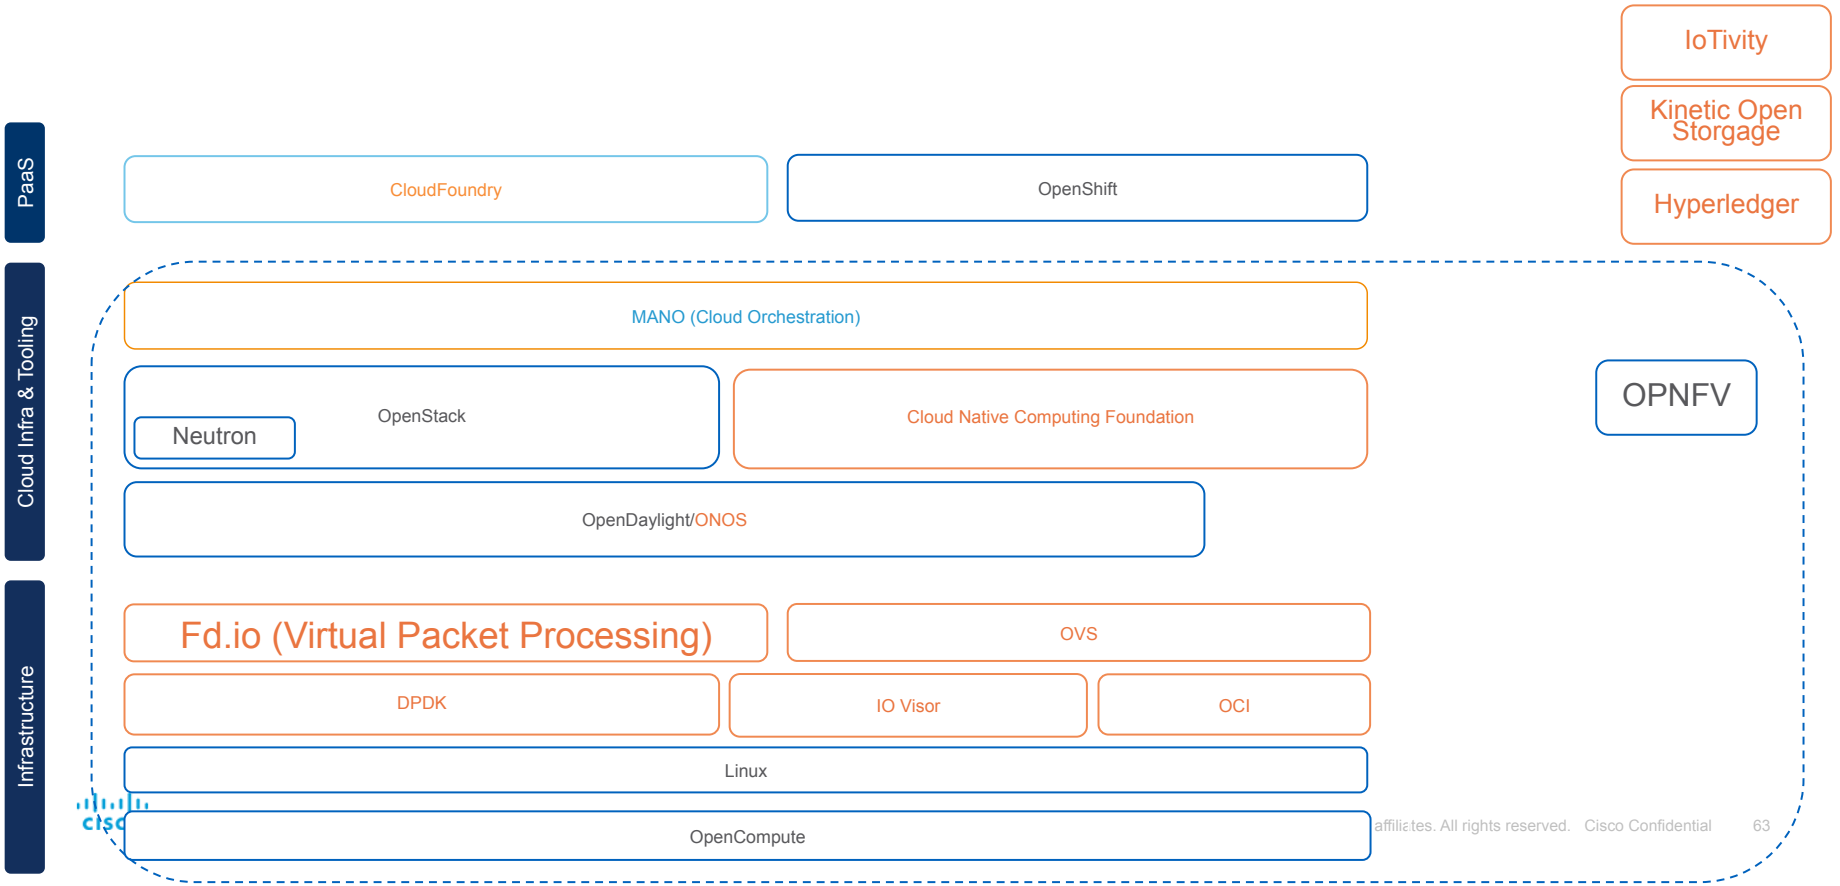

2016++: Open Source Umbrella Architecture

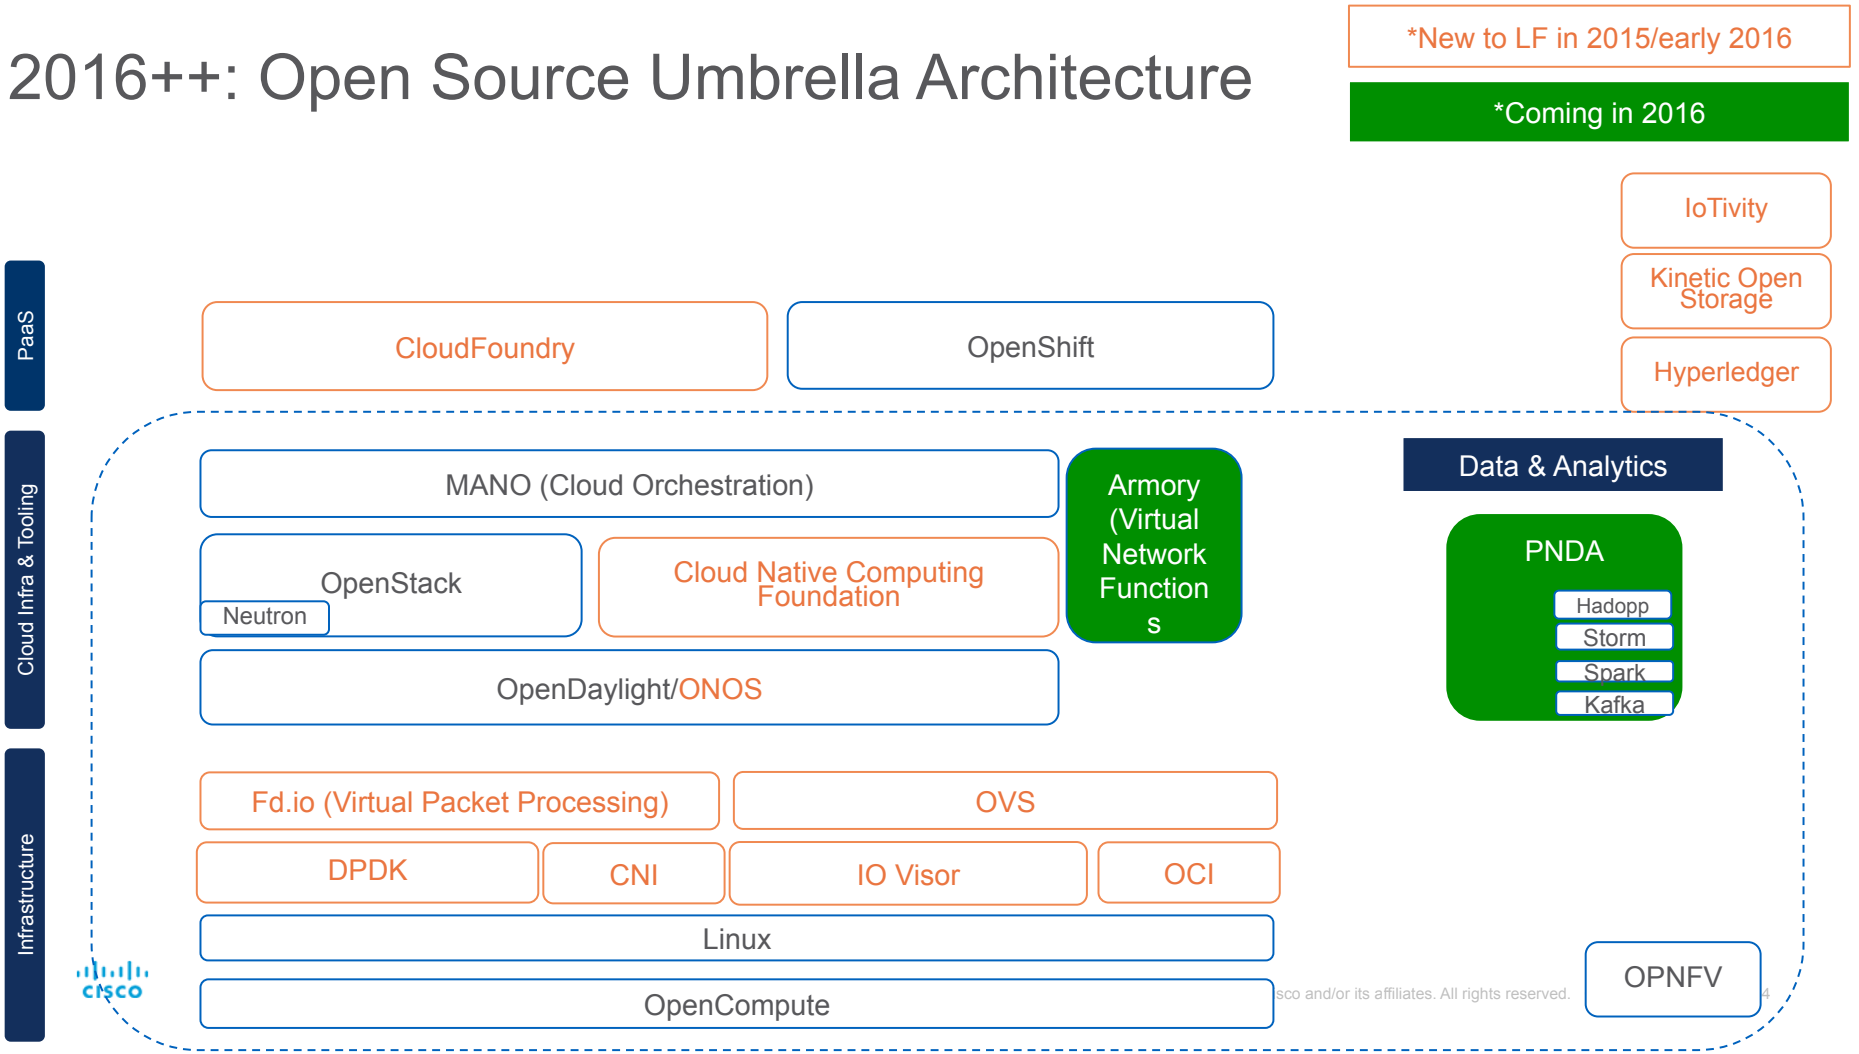

Open Networking Architecture

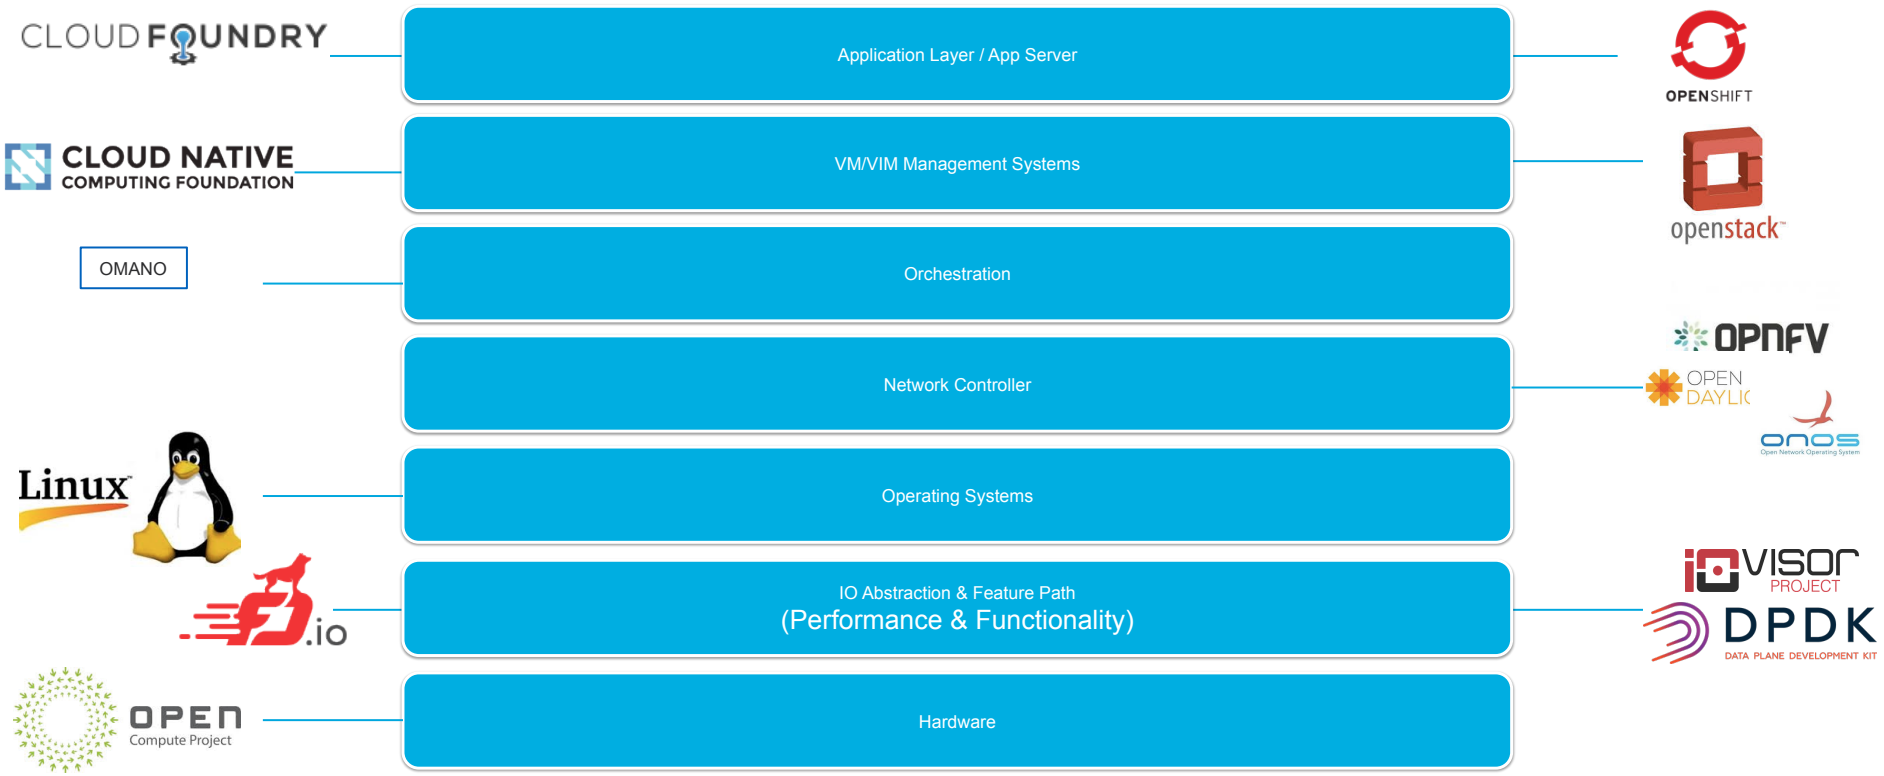

Enter The #NoStackDeveloper: “Dear Infra, Do what I need”

- Assume stable stack infra
  - Controllers, analytics, orchestration, service chaining
  - HV, Container management and lifecycle maintenance
- Simplify communication of intent w/ resources
- Visualize how Workflow actually performed and how it can work better
- Want to express Needs/Wants to layers down w/o every changing “Whole Stack” awareness/change in project

Applications/
Business Outcomes

Model Driven
Orchestration

Cloud
Platform

Compute, Network,
Storage
Infrastructure

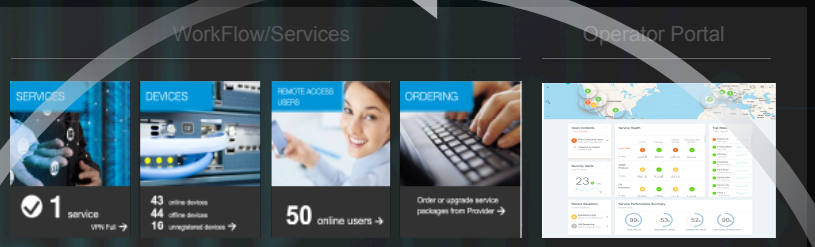

WF/Service Life Cycle Manager

VM/Container Life Cycle Manager

Consuming Analytics Apps

Platform For Network Data Analytics

Plugin Data Producers

openstack CLOUD SOFTWARE

vmware

docker

Bare Metal

OPEN DAYLIGHT

[fd.io]
Fast Data
forwarder

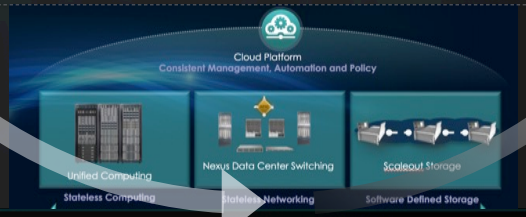

PANEL DISCUSSION