A Telco Story of OpenStack Success

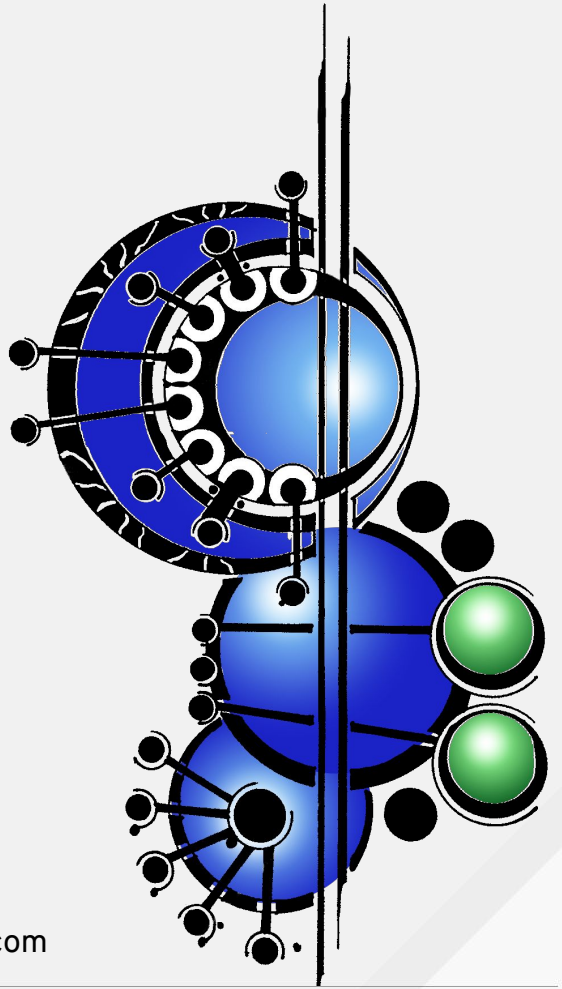

Chris Janiszewski
Solutions Architect
chris.openstack@redhat.com

Darin Sorrentino
Solutions Architect
darin@redhat.com

Aaron Hinkle
Systems Architect
Aaron.Hinkle@sprint.com

redhat.

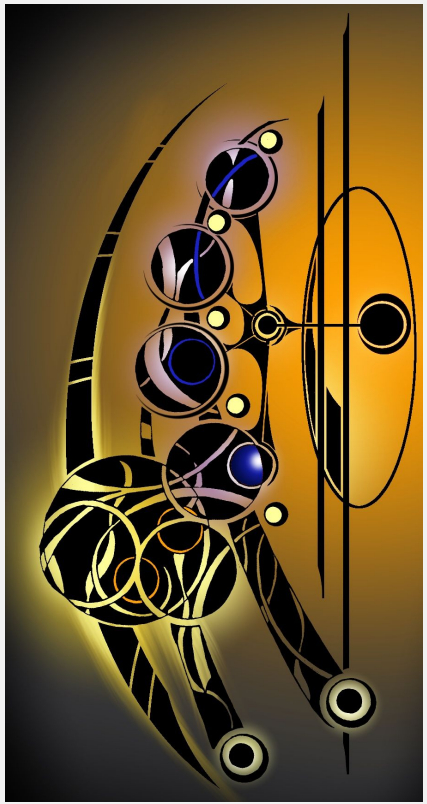

- Why the Telco’s are moving to the cloud
- How the Telco’s are implementing cloud
- What’s changing in Network Architecture

“Demand” versus “cost”

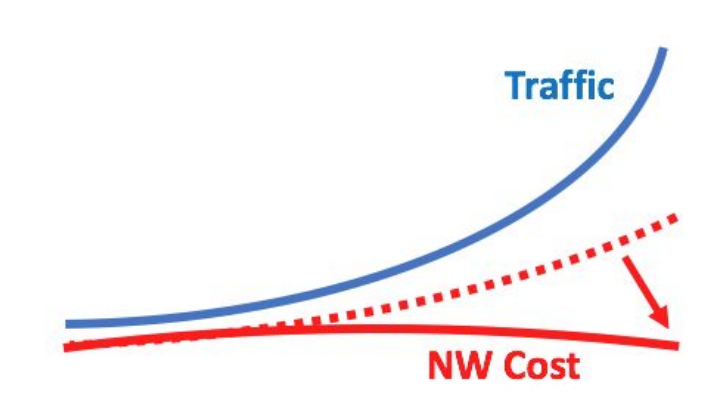

Carefully balancing the two is difficult at best…

- Traffic growth is a hockey stick
- Unlimited use for a flat fixed fee
- Cost of delivering network service needs to plateau and slowly decline

Time to Market

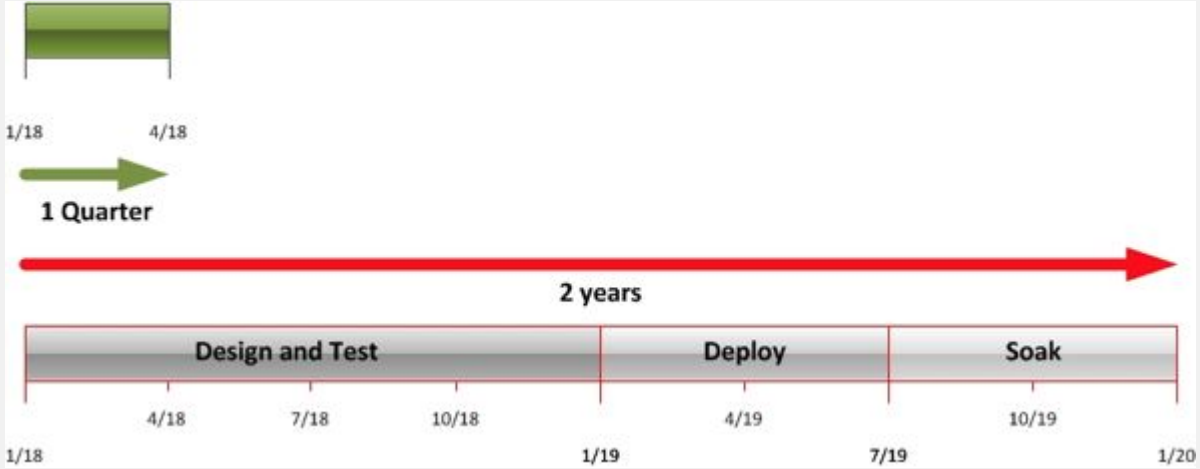

For services to be competitive

Time from ideation to customer realization needs to shrink

Business Needs Drive the Path Forward

Speed and Cost

- Rapid deployment for customers in less than a day
- Lower cost shared infrastructure

Distribution

- IoT is BIG!
- Backhaul is expensive
- Need to distribute closer to the markets

Availability

- Emergency services must ALWAYS be reachable!

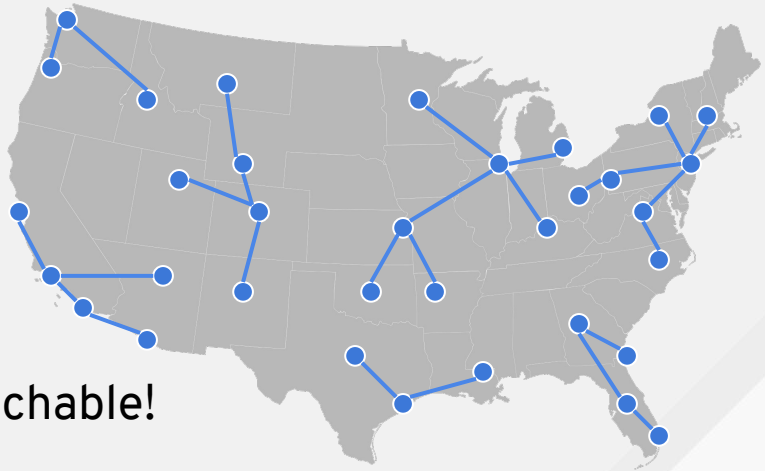

How are the Telco's deploying Openstack?

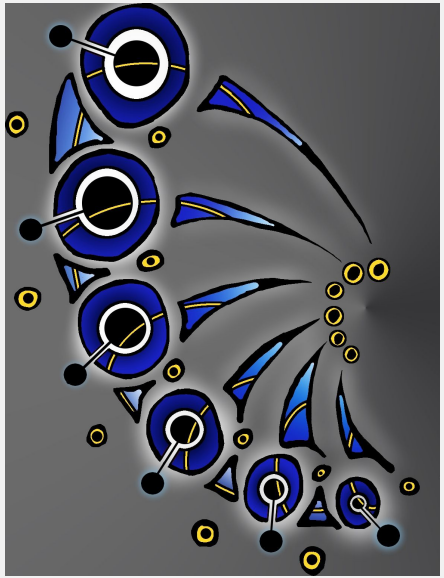

Survey

- **What** - anonymous survey for Telcos in North America
- **Target** - system and infrastructure architects, technical decision makers
- **Goal** - learn what drives OpenStack adoption at Telco and discover best practices and utilize experience to overcome common challenges

Why are telco's moving to the cloud?

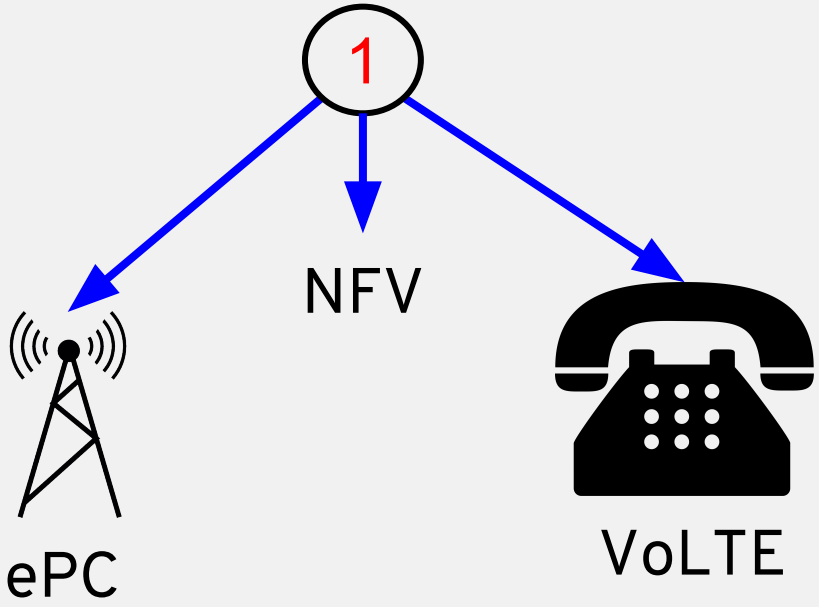

Why Openstack?

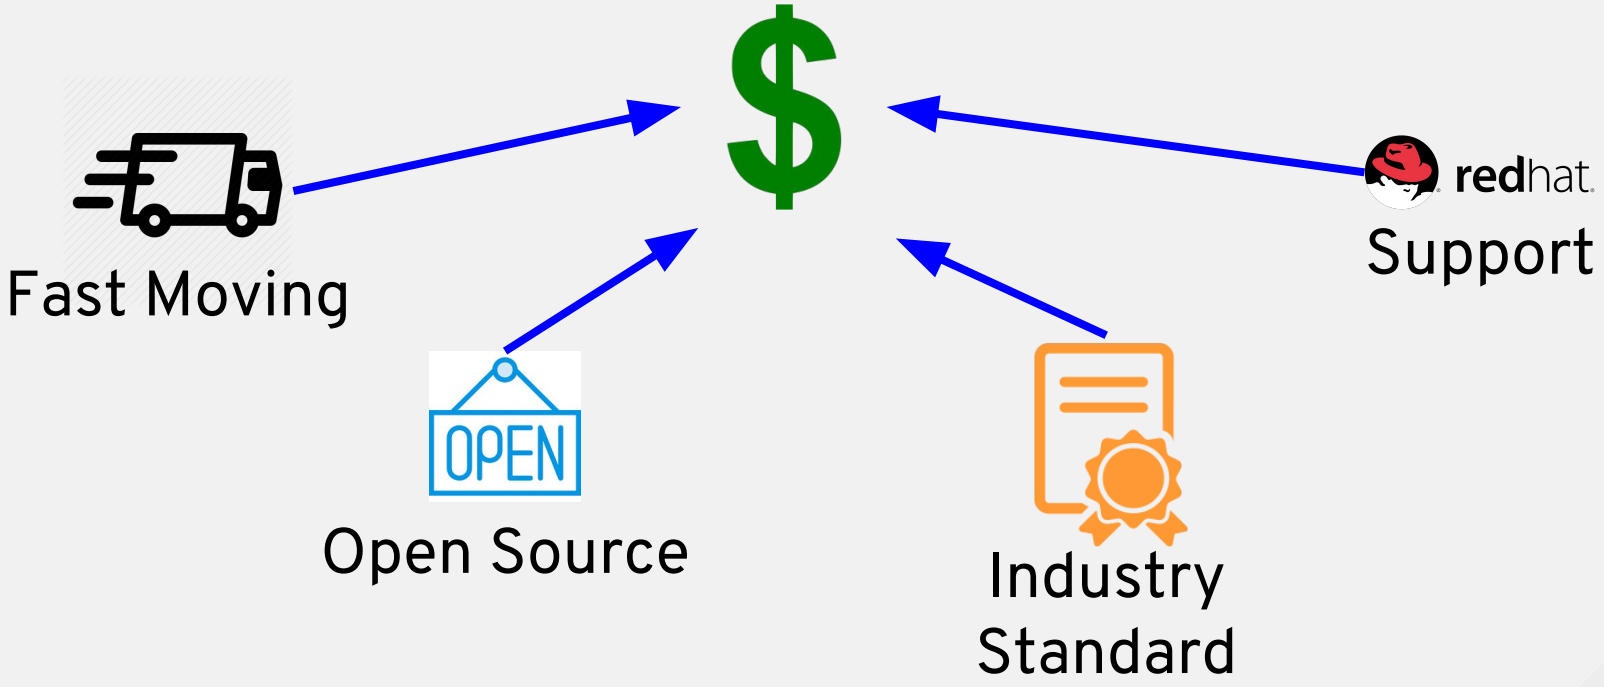

Which infrastructure?

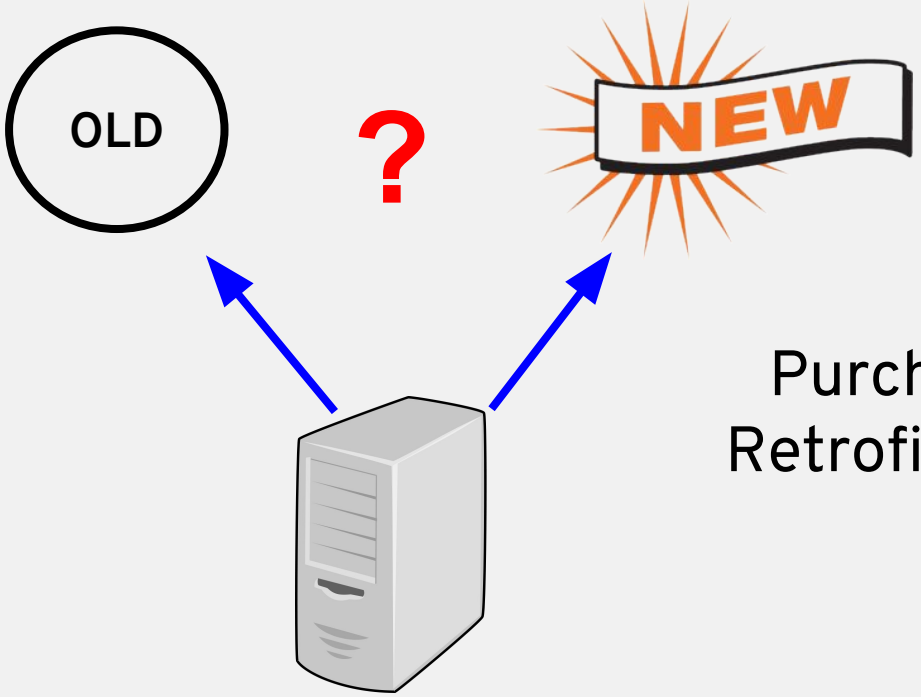

Purchasing cycle takes too long
Retrofit old hardware is too painful

Go New

Networking?

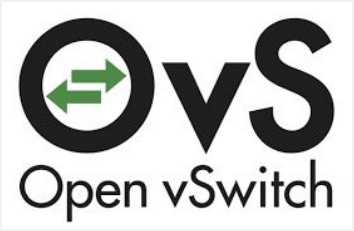

?

- **Open vSwitch** is the initial chose due to cost and open source
- **SDN** products offer a lot of potential and are currently under feature evaluation

What is the Telco footprint for Openstack?

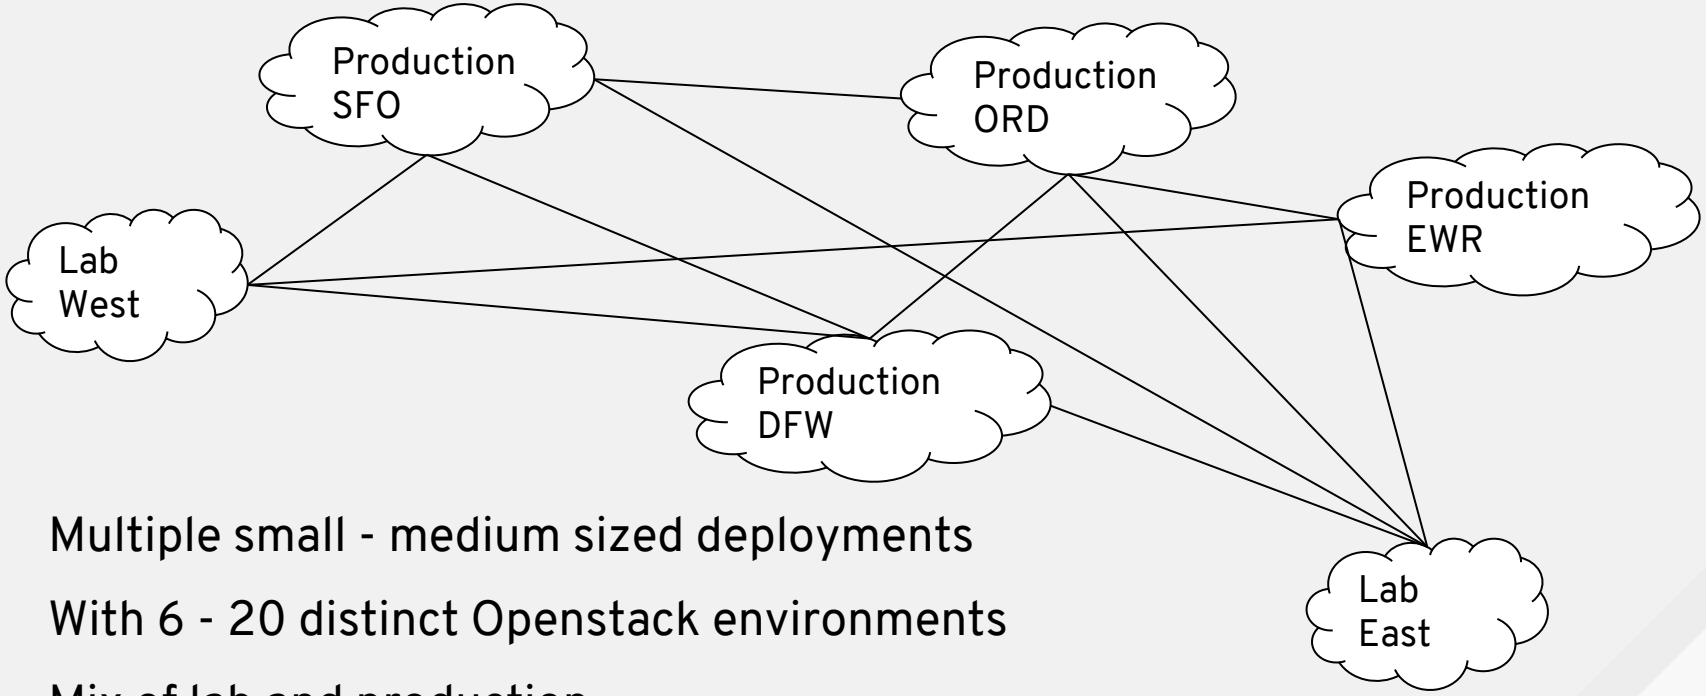

- Multiple small - medium sized deployments
- With 6 - 20 distinct Openstack environments
- Mix of lab and production

What are some of the hurdles to deploying Openstack?

Champions educate the rest

- Skills gap running the platform
- Changing people's paradigm

  ➡ Clouds are very different to traditional telco!
- Getting silo'd groups to work together
- Getting a consistent deployment process and post deployment customizations

Lessons Learned

- Automation is king
- Evaluate emerging technologies like OVS+DPDK and SR-IOV
- Focus on the why and the use cases

  Don’t build the “field of dreams” …

  … they are not coming
- VNF vendors are still learning, vet them well
- Don’t allow the “way we have always done it” hinder forward progress

Be Agile.
Pick the right partners.
Learn from the best.

Why think about storage in the cloud?

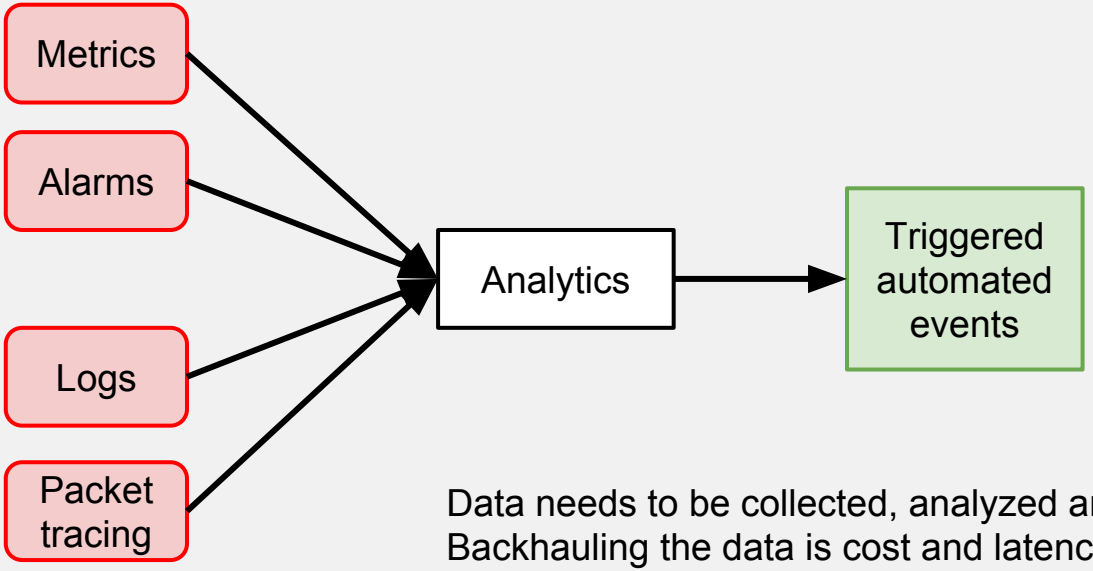

Data needs to be collected, analyzed and events triggered locally
Backhauling the data is cost and latency prohibitive
That's a lot of data that needs to be stored

Storage - Architectural evaluation

Cost:

- Local - free
- SDS solution - $$$ per node
- Traditional storage - $$$ per system

Features:

- Unified delivery (block, object, filesystem)
- Capacity and scalability
- Link Clones
- Disaster Recovery
- Management

Performance:

- Multi-tier
- Add performance on demand

Storage - Architectural evaluation

Local storage

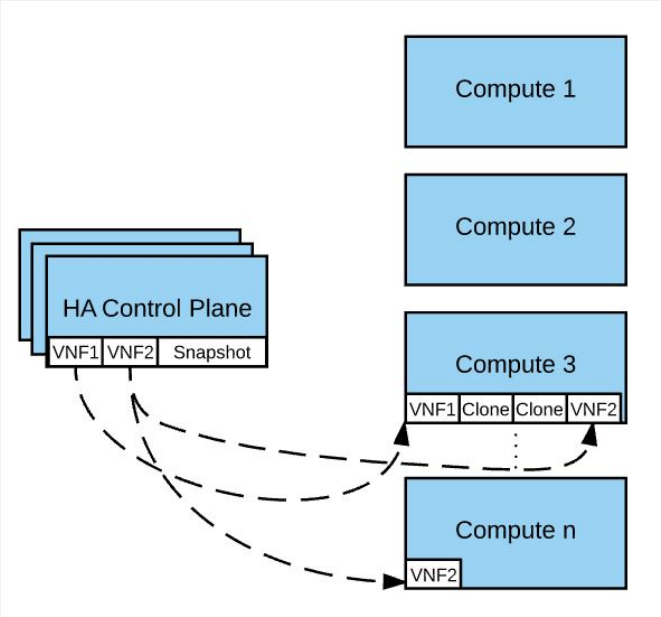

SDS / Centralized

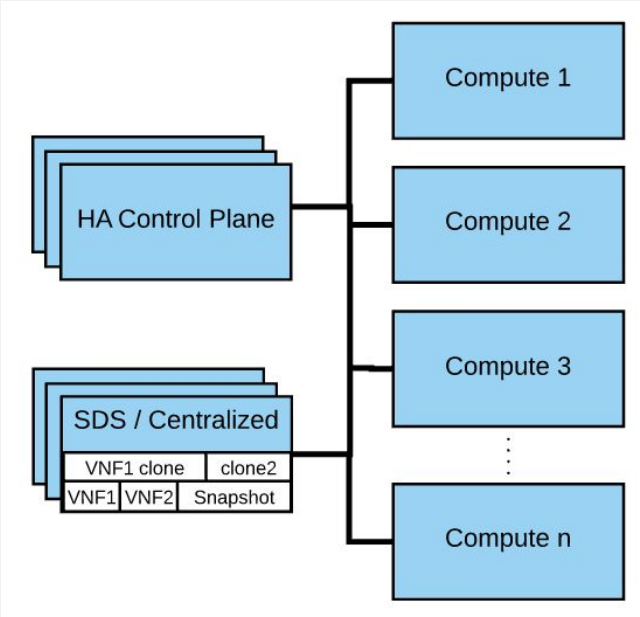

redhat.

Storage - Architectural evaluation

Deploying VNFs on unoptimized storage can cost you 5x deployment penalty.

SDS		Traditional	
Pros	**Cons**	**Pros**	**Cons**
Unified Delivery (block, object and filesystem)	Requires Linux skills	Storage team knows it	Siloed
Scale-out architecture		Management and Monitoring	Slow moving
Fast moving		Mature	
Live migrations, better performance and HA!			

Network Optimizations

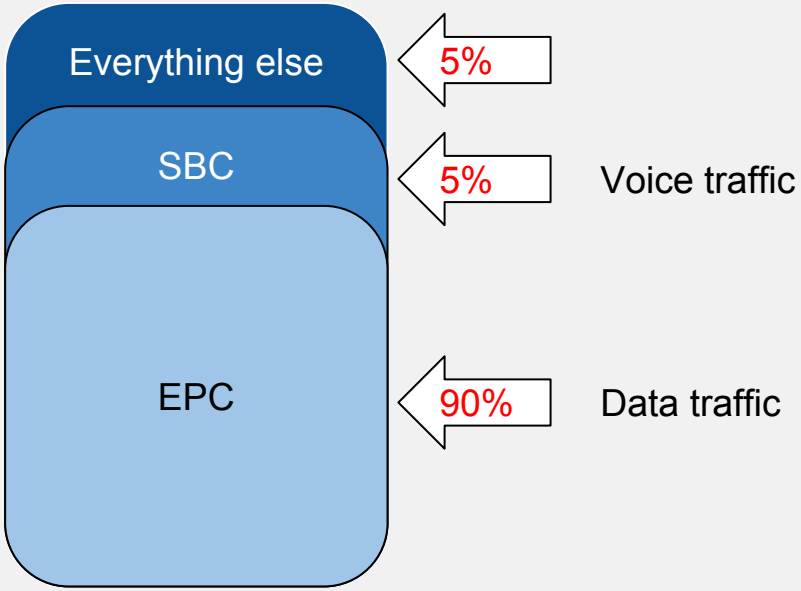

Need line rate throughput

Bearer path functions can move 10 gbps or greater on a fraction of the CPU cores available.

Network - Architectural evaluation

Cost:

- OVS - free
- Third-party SDN solution - $$$ per node

Features:

- Stretching L2/L3 between OpenStack PODs
- Integration with Service Provider networks
- Service function chaining
- Federated overlay networks
- SR-IOV
- DPDK

Performance:

- Workflow management
- Traffic shaping

Network - Architectural evaluation

Don't *interrupt* your speed to a successful implementation!

SR-IOV	
Pros	**Cons**
Near line Speed	Loss of flexibility - Live migrations
NIC sharing	Loss of functionality - Disable Neutron FW

DPDK	
Pros	**Cons**
~90% line speed	CPU overhead - 100% util of PMD
Software based	Complex "tuning" - Framework for dev

This is not an either/or situation, choose the technology that best suits the NFV!!

Network - Architectural evaluation

Options - Evaluating vanilla OVS & Third party SDN vendors

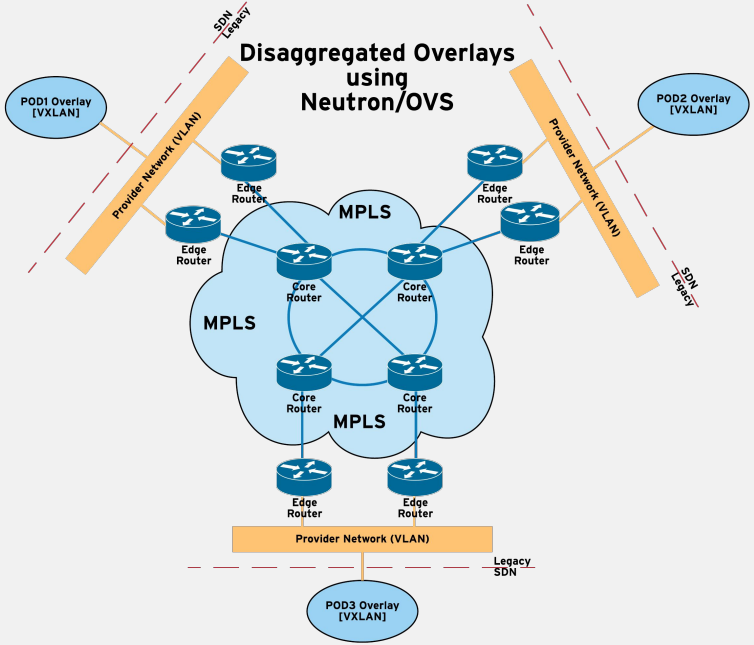

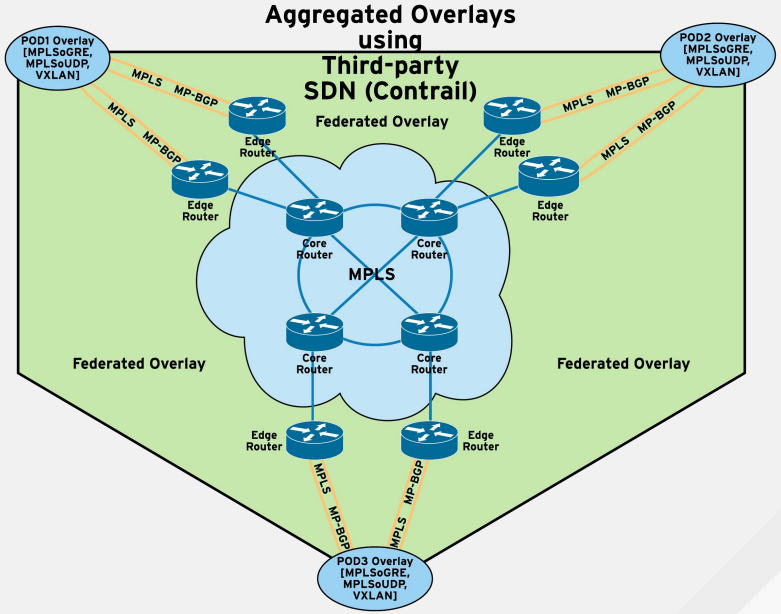

Network - Architectural evaluation

Options - Evaluating Vanilla OVS

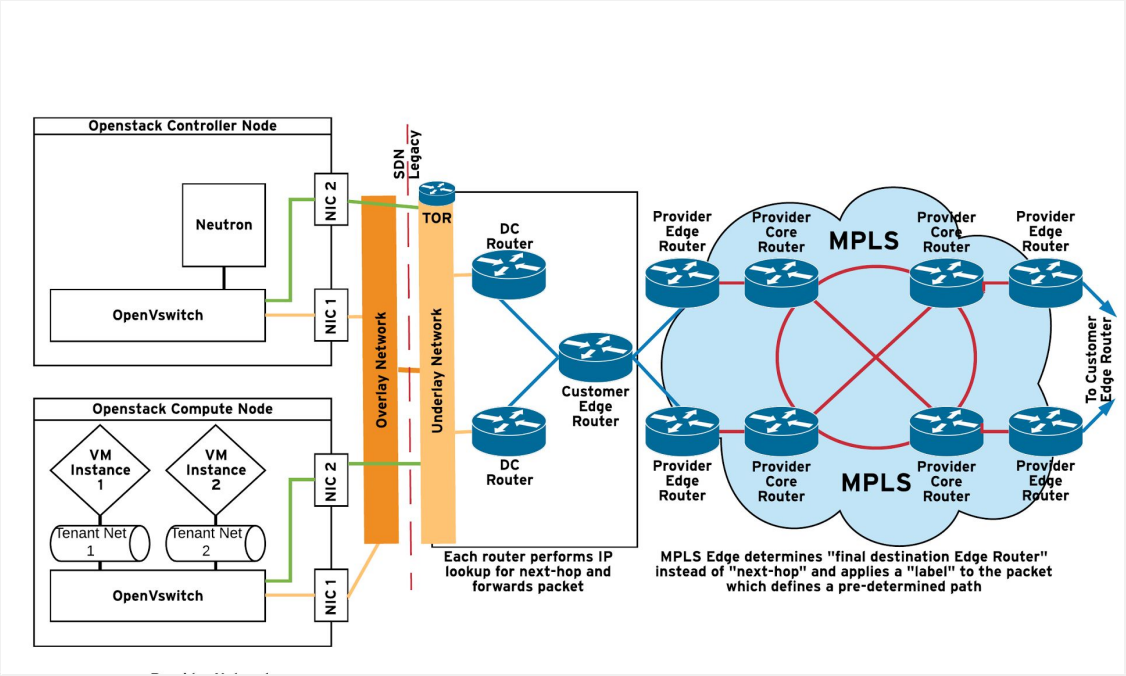

- Included as part of Openstack
- Provider networks required to remove Neutron from the Data Path (Performance)
- Decreased complexity for implementation, increased complexity for Pod-to-Pod routing
- MPLS Edge is outside of each Openstack Pod:
  - Increased workload in the Data Center Core using per-packet “next-hop” routing
  - Isolated L3 from one pod to the next
  - Lack of control to manage capacity, prioritize different services, and prevent congestion within DC

Network - Architectural evaluation

Options - Third party SDN (such as Juniper Contrails)

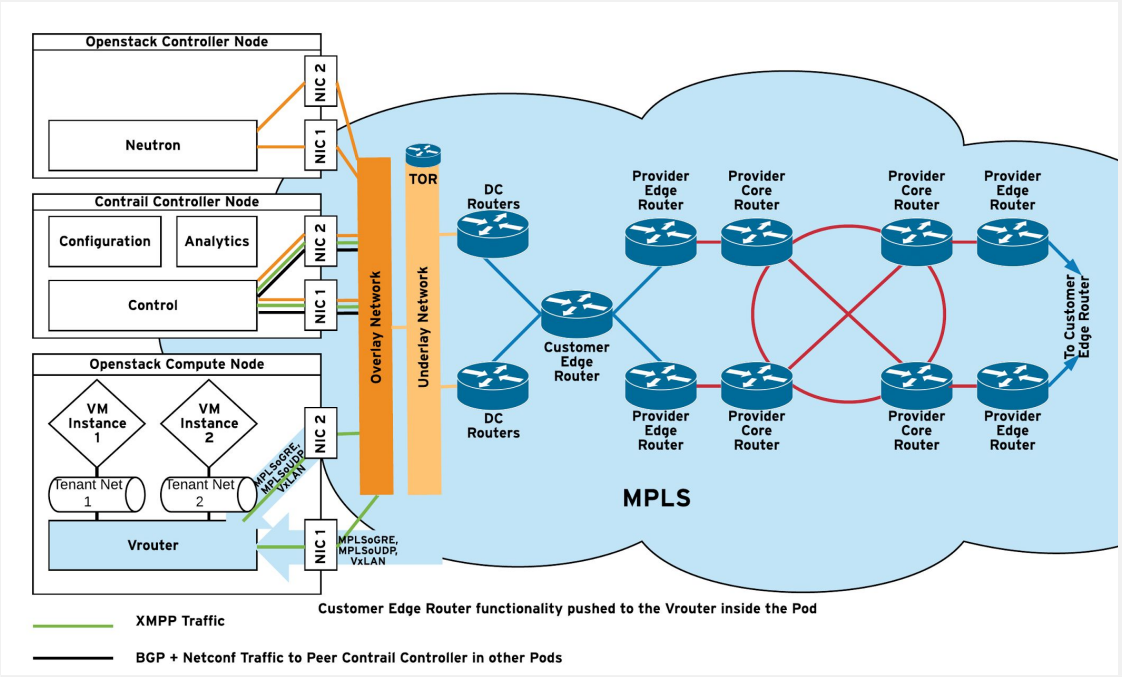

- Additional expense outside of Openstack
- Routing capabilities (Vrouter) replace Switch capabilities (OpenVswitch)
- Increased complexity of implementation, decreased complexity of Pod-to-Pod routing
- Third-party SDN extends MPLS closer to the Virtual Instance
    - Decreased workload on DC Core Routers
    - Stretch L3 using MPLS labels
    - Increased control to manage capacity, prioritize different services, and prevent congestion

3rd Party SDN Vendors

Additional considerations:

- Monolithic versus ML2 plug-in implementation
- Self sufficient pod using federated deployment versus centralized external deployment
- Licensing models

BREAKOUT SESSIONS - Tuesday, November 7th

Neutron-based networking in Kubernetes using Kuryr – a hands-on lab	Sudhir Kethamakka, Geetika Batra, and Amol Chobe (JP Morgan Chase)	10:50am - 12:20pm
A Telco story of OpenStack success	Krzysztof Janiszewski, Darin Sorrentino, and Aaron Hinkle (Sprint)	1:50pm - 2:30pm
Turbo-charging OpenStack for NFV workloads	Ajay Simha, Vinay Rao, and Ian Wells (Cisco)	3:20pm - 3:30pm
Windmill 101: Ansible-based deployments for Zuul / Nodepool	Paul Belanger and Ricardo Carrillo Cruz	3:20pm - 4:50pm
Simpler encrypted volume management with Tang	Nathaniel McCallum and Ade Lee	3:50pm - 4:00pm
Deploying multi-container applications with Ansible service broker	Eric Dube and Todd Sanders	5:00pm - 5:40pm

BREAKOUT SESSIONS - Wednesday, November 8th

OpenStack: the perfect virtual infrastructure manager (VIM) for a virtual evolved packet core (vEPC)	Julio Villarreal Pelegrino and Rimma Iontel	9:00am - 9:40am
Questions to make your storage vendor squirm	Gregory Farnum	9:50am - 10:30am
Bringing worlds together: designing and deploying Kubernetes on an OpenStack multi-site environment	Roger Lopez and Julio Villarreal Pelegrino	10:20am - 10:30am
DMA (distributed monitoring and analysis): monitoring practice and lifecycle management for Telecom	Tomofumi Hayashi, Yuki Kasuya (KDDI) and Toshiaki Takahashi (NEC)	1:50pm - 2:00pm
Standing up and operating a container service on top of OpenStack using OpenShift	Dan McPherson, Ata Turk (MOC), and Robert Baron (Boston University)	1:50pm - 2:30pm

BREAKOUT SESSIONS - Wednesday, November 8th

Why are you not a mentor in the OpenStack community yet?	Rodrigo Duarte Sousa, Raildo Mascena, and Telles Nobrega	1:50pm - 2:30pm
What the heck are DHSS driver modes in OpenStack Manila?	Tom Barron, Rodrigo Barbieri, and Goutham Pacha Ravi (NetApp)	1:50pm - 2:30pm
SD-WAN – the open source way	Azhar Sayeed and Jaffer Derwish	2:40pm - 3:20pm
Adding Cellsv2 to your existing Nova deployment	Dan Smith	3:30pm - 4:10pm
What's your workflow?	Daniel Mellado and David Paterson (Dell)	3:30pm - 4:10pm
Glance image import is here…now it's time to start using it!	Erno Kuvaja and Brian Rosmaita (Verizon)	4:30pm - 5:10pm

THANK YOU

plus.google.com/+RedHat

facebook.com/redhatinc

linkedin.com/company/red-hat

twitter.com/RedHatNews

youtube.com/user/RedHatVideos