

Integrating Openstack with Tungsten Fabric

GÉANT 4th SIG-NGN meeting

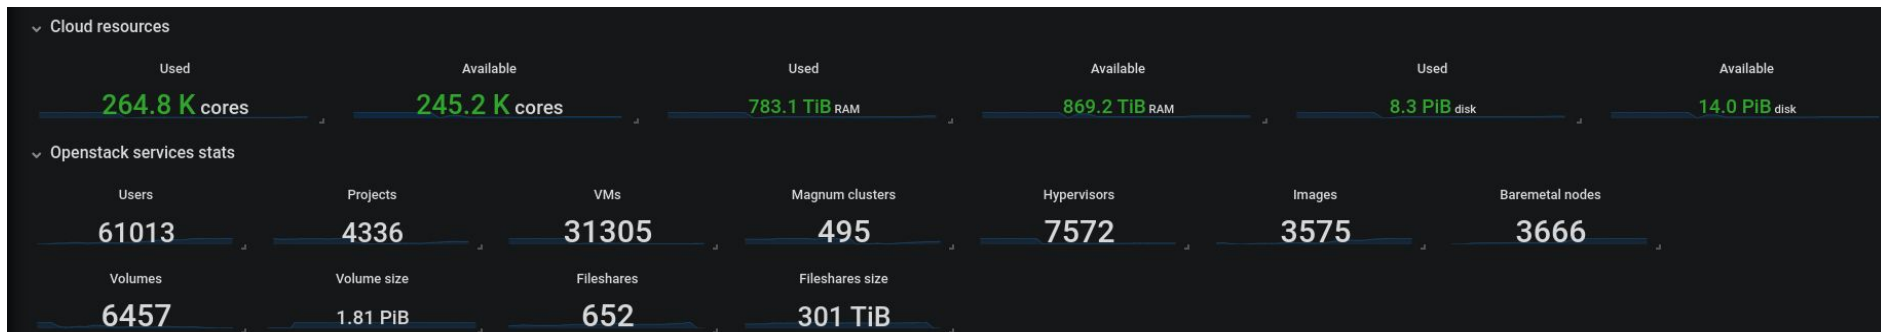
16.01.2020

Hamza Zafar, CERN IT-CM-RPS

hamza.zafar@cern.ch

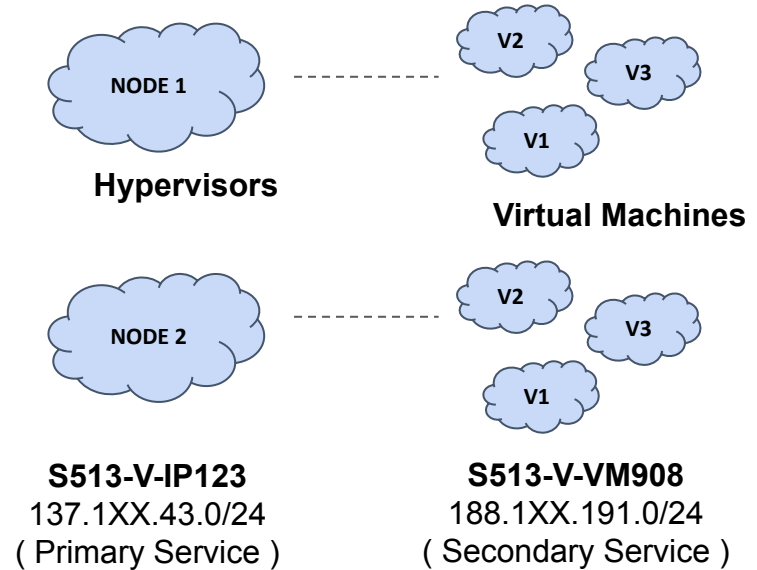
OpenStack @ CERN

- Platform to build and manage clouds
- Control compute, network and storage resources



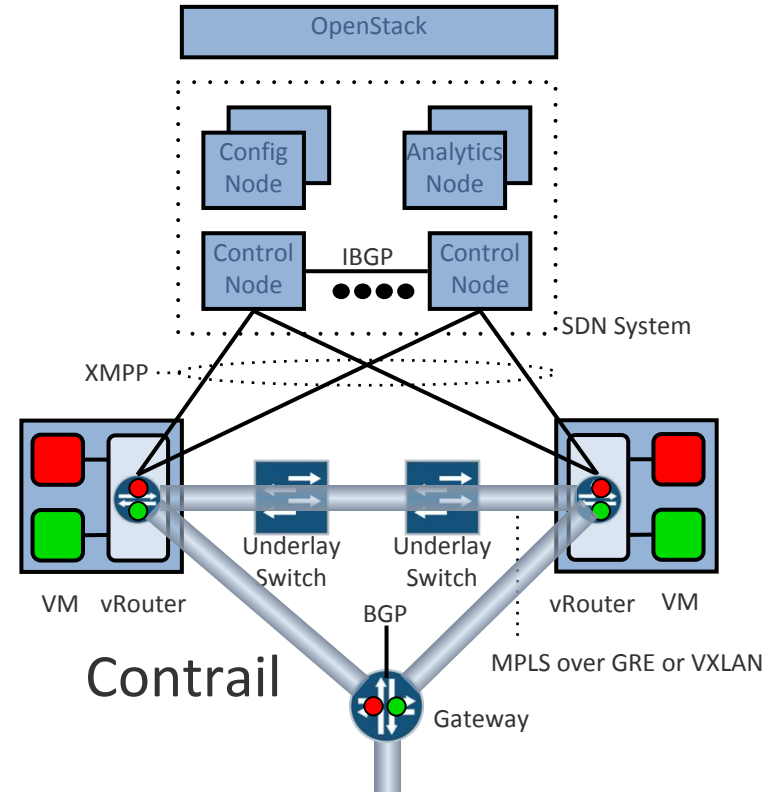
Exploring Tungsten Fabric

- CERN's network architecture
 - Flat but segmented network, with multiple broadcast domains
 - Limited network isolation
 - Limited IP mobility
- Overlay networks
- Load Balancer as a Service
 - DNS based load balancing
 - Need for L4/L7 load balancers
 - Ruled out Octavia load balancing solution



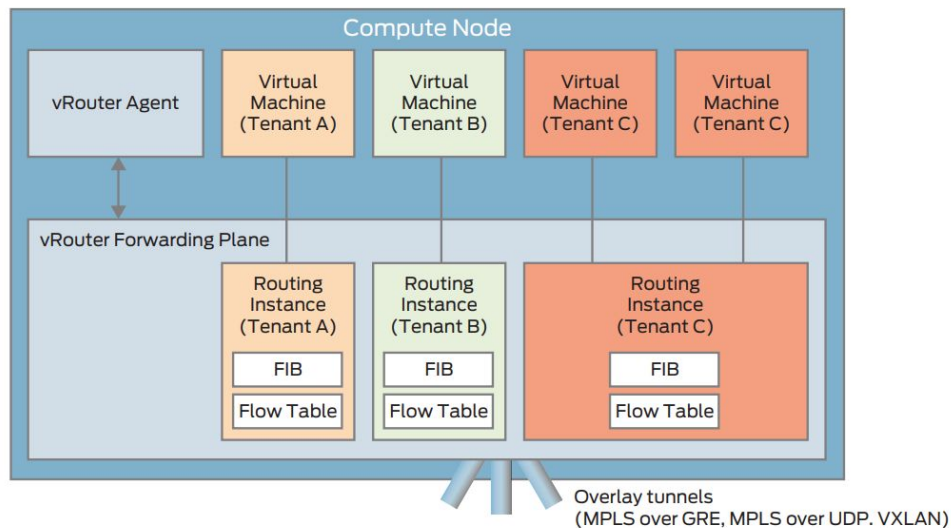
Tungsten Fabric in a nutshell

- Software Defined Networking (SDN) solution
- TungstenFabric aka. OpenContrail
- Network isolation (VxLAN, MPLSoGRE, MPLSoUDP)
- Virtual Network translates to a VRF
- Security and network policies
- Advanced features (load balancer, floating IPs, etc.)
- OpenStack and Kubernetes as orchestrators



Forwarding Plane - vRouter

- vRouter Agent
 - Receives routes from control plane
 - Installs forwarding state into the forwarding plane
 - Reports analytics state
 - Proxies DHCP, ARP and DNS
- vRouter Kernel Module
 - Encap/decap packets
 - Deliver packets to correct destination

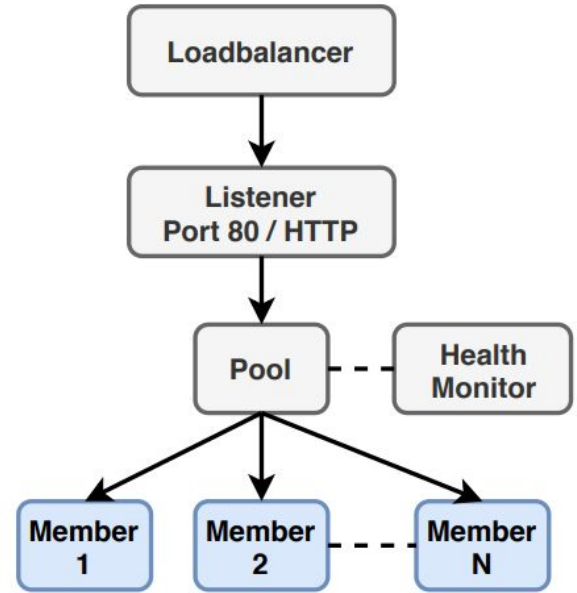


Tungsten Fabric Deployment @ CERN

- New SDN region in OpenStack
- Containerized control plane
- Helm for control plane deployment (OpenStack and Tungsten)
- 3 compute nodes
- QFX10k as SDN gateway
 - Auto-configured QFX10k by using tungsten's device manager
 - BGP peering between QFX10k and Tungsten Fabric controller
 - VxLAN tunnels between QFX10k and compute nodes
- **First service:** Load Balancing as a Service (LBaaS)

LBaaS Feature

- Integrates seamlessly with OpenStack
- HAProxy is spawned in a network namespace
- Offers a failover mechanism
- SSL Termination
- Choice of load balancing algorithm
- Health checks etc.
- L7 policies
- Documentation: http://clouddocs.web.cern.ch/clouddocs/networking/load_balancing.html



DEMO TIME

Step 1: Create a Load Balancer

```
$ openstack loadbalancer create --name mylb --vip-network-id public
```

Triggers creation of a Neutron port

HAProxy instances deployed in Active / Passive mode, different hypervisors

Step 2: Create a Listener

```
$ openstack loadbalancer listener create --name mylistener \  
    --protocol HTTP --protocol-port 80 mylb
```

Protocols supported

TCP, HTTP, HTTPS, TERMINATED_HTTPS, UDP

Step 3: Create a Pool

```
$ openstack loadbalancer pool create --name mypool \  
    --lb-algorithm ROUND_ROBIN --listener mylistener --protocol HTTP
```

Available LB Algorithms: ROUND_ROBIN, SOURCE_IP, LEAST_CONNECTIONS

Session Persistence using cookies

Step 4 (optional): Create a Health Monitor

```
$ openstack loadbalancer healthmonitor create --name tcp-monitor \  
    --delay 7 --timeout 5 --max-retries 3 --type TCP mypool
```

Available Types: PING, HTTP, HTTPS, TLS HELLO, TCP, UDP-CONNECT

Step 5: Add Members

```
$ openstack loadbalancer member create --address 137.138.53.95 \  
  --weight 2 --protocol-port 80 mypool
```

```
$ openstack loadbalancer member create --address 188.185.80.141 \  
  --weight 1 --protocol-port 80 mypool
```

Weights can be assigned

Enable/Disable

```
$ openstack loadbalancer member set mypool <member-id> --disable
```

Conclusion & Next Steps

- SDN Region up and running, LBaaS as the first available service
- Contributing upstream
- Deploy control plane in multiple AZs
- Disaster recovery?
- Ongoing Work
 - Integration of Virtual IPs with LanDB for DNS
 - Developing Tungsten provider driver for Octavia
- Coming Next
 - Full SDN Solution: Floating IPs, Overlay Networks, Security Groups, FWaaS, ...

Thank You!

Evaluation

	Neutron/OpenVSwitch	OpenDaylight	OVN
DHCP	Neutron	Neutron/Built-in	Built-in
Floating IPs	Yes	Yes	Yes
Distributed Routing	Only with DVR	Yes	Yes
Tunneling Protocols	vxlan / GRE / geneve	vxlan / GRE / geneve	vxlan / geneve
Security Groups	IPTables	OpenFlow Native	OpenFlow Native + Logging
Load Balancing	Octavia	Octavia	Octavia / OVN Native
Acceleration	Limited DPDK	DPDK	DPDK
Tracing	tcpdump	tcpdump	ovn-trace
Physical Switch Integr.	L2 / L3	L2 / L3	L2 / L3

Meet Tungsten Fabric

