



ITA DATA CENTRE ARCHITECTURE

WITH OPEN NETWORKING





Schumacher Neto

Gestor Core



CORE MANAGER

- www.paratus.ao
- www.paratus.africa
- www.peeringdb.com/asn/33763
- AS33763



AGENDA



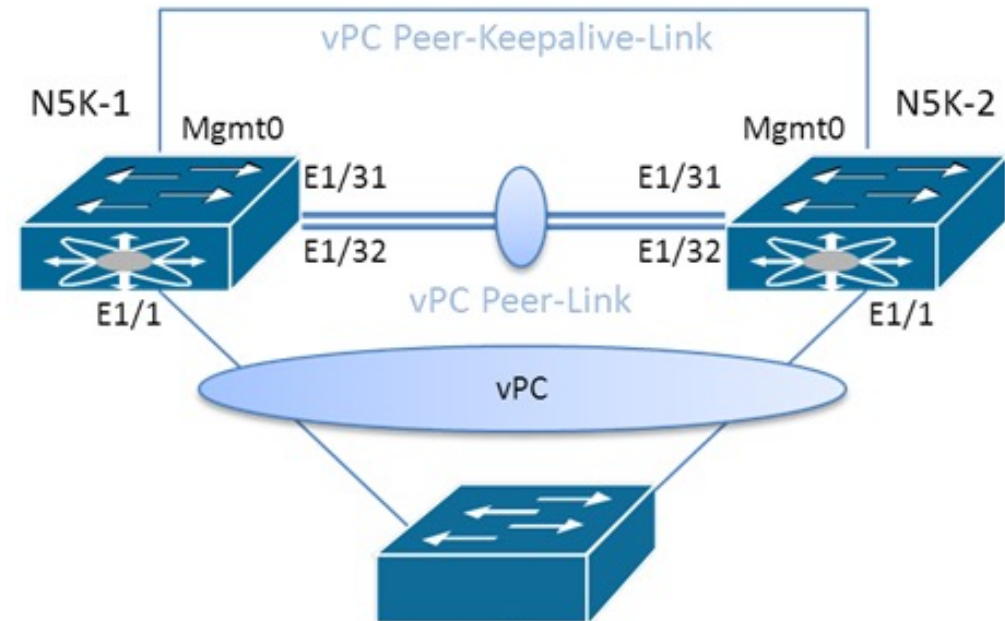
1. MOTIVATION
2. ARCHITECTURE (2018)
3. CLOUD SERVICES FUTURE-PROOF
4. RATIONALE
5. ROLLOUT
6. UNDESIRE NEWS & MOVING FORWARD
7. DEMO

MOTIVATION



ACCESS – AGGREGATION – CORE (2018)

- Cisco Nexus 5000
- Fabric Extenders
- Layer 2 Fabric (Bridging)
- Rapid Spanning-Tree Protocol
- Virtual Port-Channel
- 4094 VLANs



CLOUD SERVICES FUTURE-PROOF



Scalability and cost



Flexibility



Cloud services availability Zones



Complexity and technical debts



Automation



PARATUS
Always Prepared

OPEN NETWORKING JOURNEY




RATIONALE

- Right Architecture
- Network Disaggregation
- Features
- Automation




CLOS TOPOLOGY




Charles Clos

- Spine and Leaf
- Economical to build large networks
- Simple building blocks
- Networks become more scalable, flexible and agile
- Fine-grained failure domain




Network As Legoland

- Construction of a large infrastructure in small block
- Network operations become simpler
- Failures become simpler
- Networks become more scalable, flexible and agile



HARDWARE & NOS

- Hardware investment tailored to the needs
- Cost effective Hardware CAPEX



Hardware & NOS

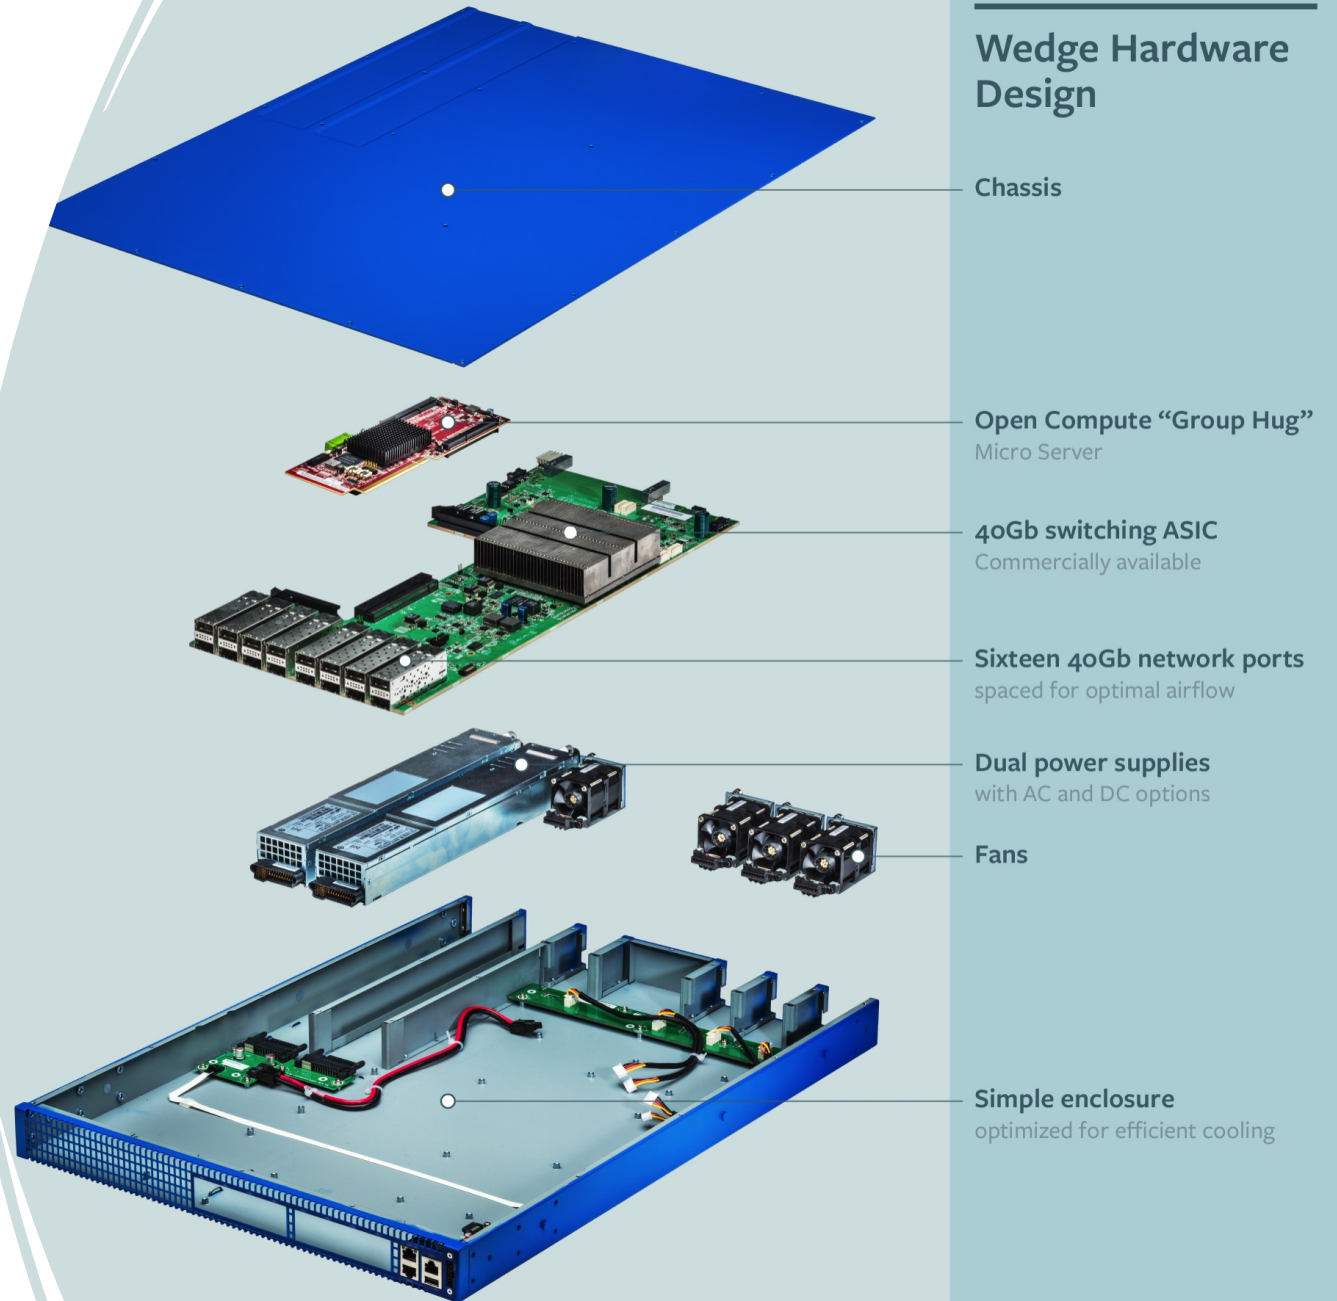
- Network Disaggregation
- Open Compute Switches
- Application-specific Integrated Circuits (ASICs)

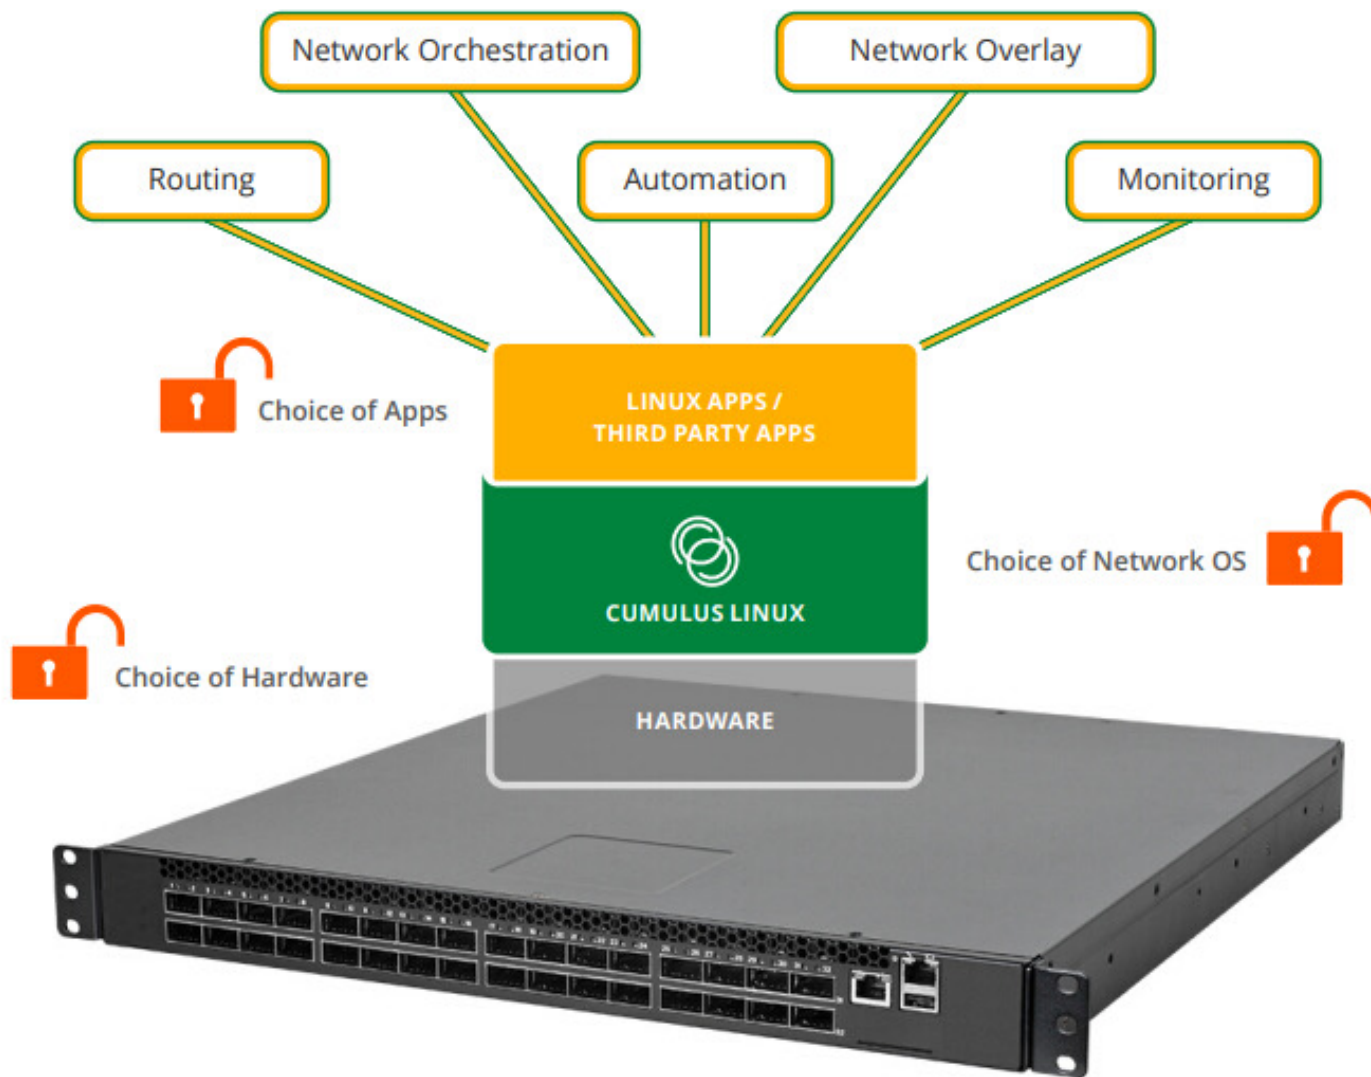
CUMULUS 

ipinfusion™

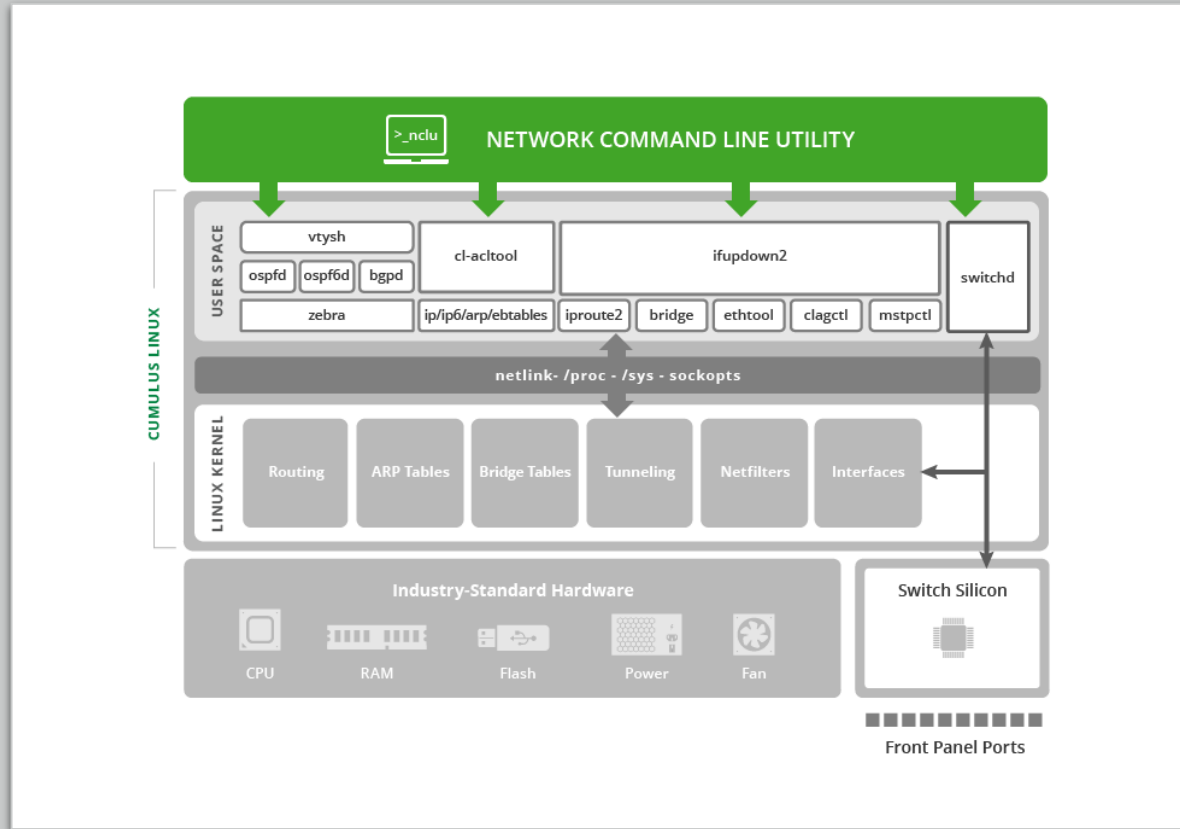
 SONiC

MELLANOX
ONYX
Shape Your Network





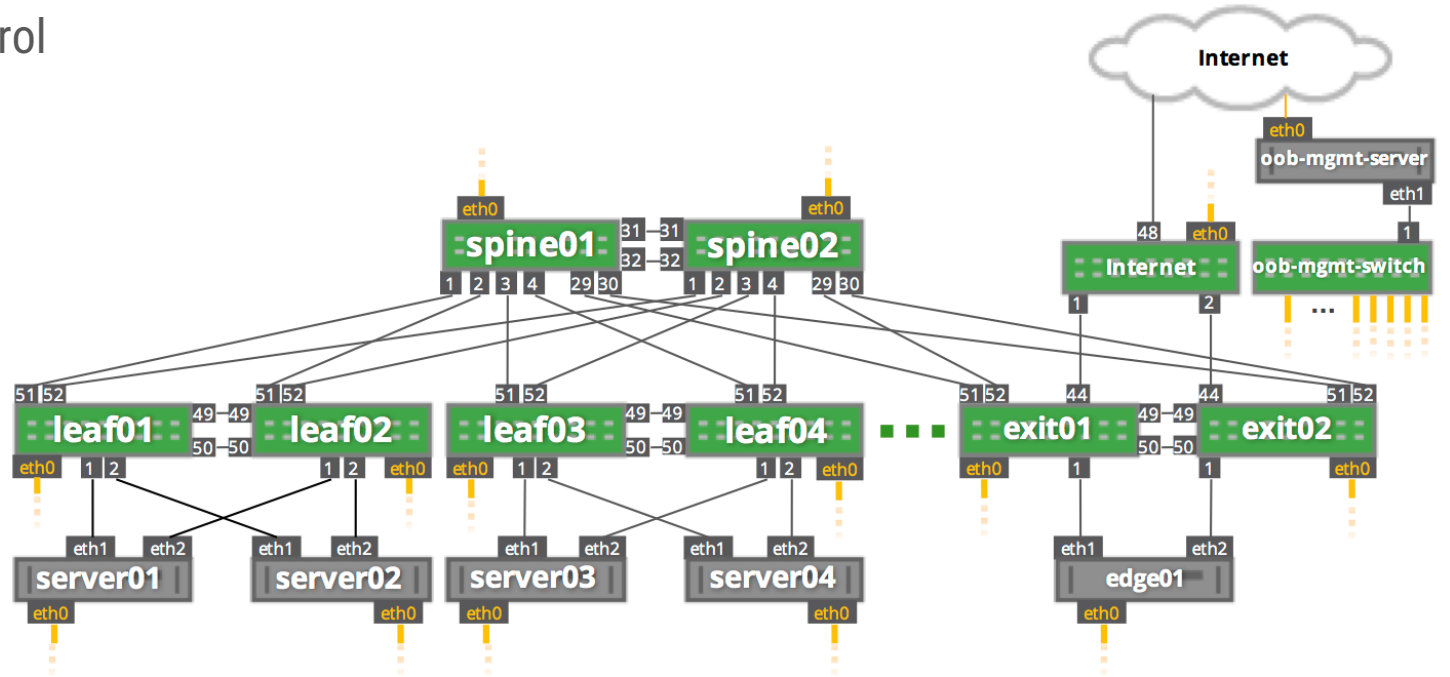
Cumulus Linux NOS



L2 FABRIC – 2019 ARCHITECTURE

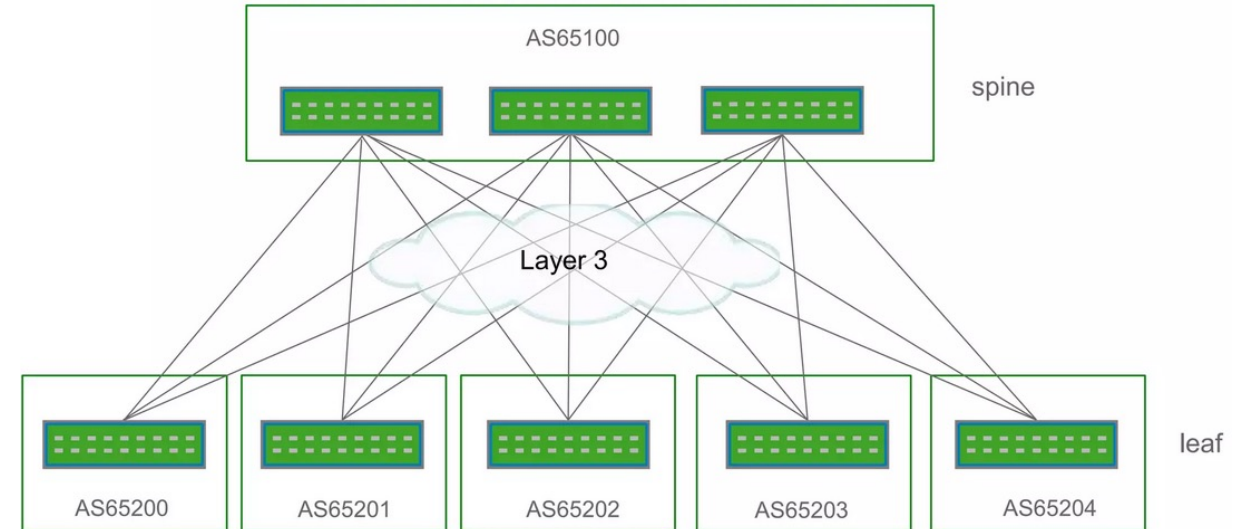


- 2 Tier Clos
- L3 Underlay (BGP)
- EVPN network virtualization and control plane
- VXLAN Overlay
- ECMP
- Ansible Playbooks



LAYER 3 UNDERLAY WITH eBGP

- BGP with private ASNs
- RFC 5549 BGP unnumbered
- eBGP easier to understand path selection
- BGP optimized timers

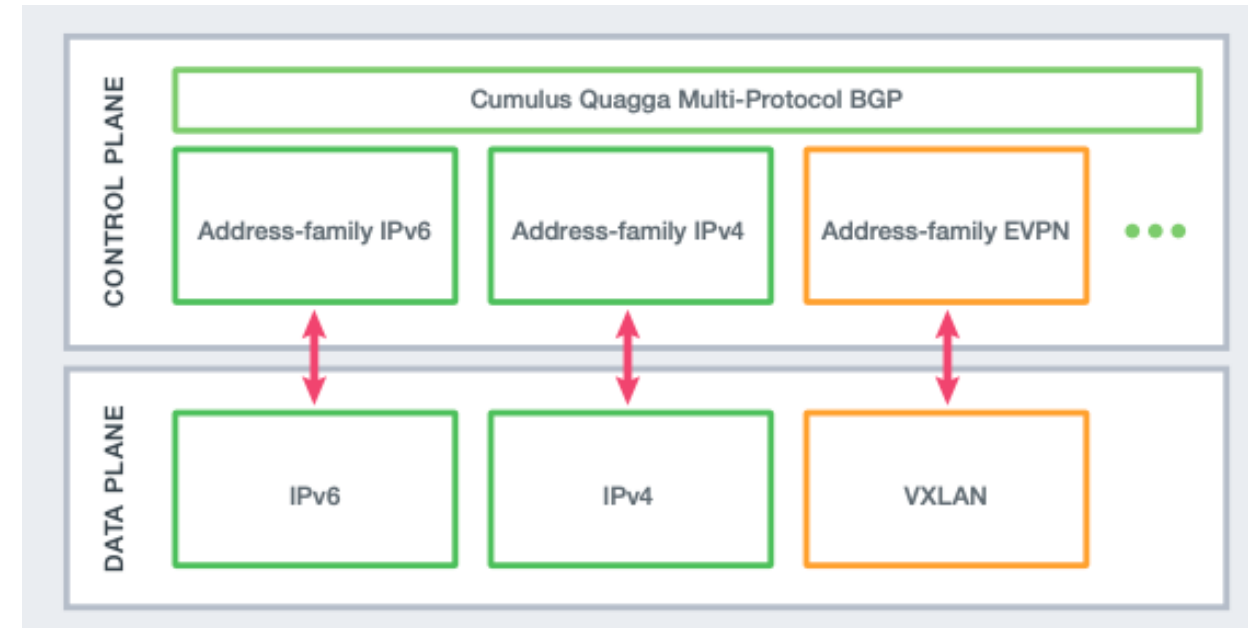


```
router bgp 65200
  bgp router-id 172.19.1.1
  neighbor CUMULUS peer-group
  neighbor CUMULUS remote-as external
  neighbor CUMULUS timers 1 3
  neighbor swp1 interface peer-group CUMULUS
  neighbor swp2 interface peer-group CUMULUS
  neighbor swp3 interface peer-group CUMULUS
!
address-family ipv4 unicast
  network 172.19.1.1/32
!
```


EVPN NETWORK VIRTUALIZATION

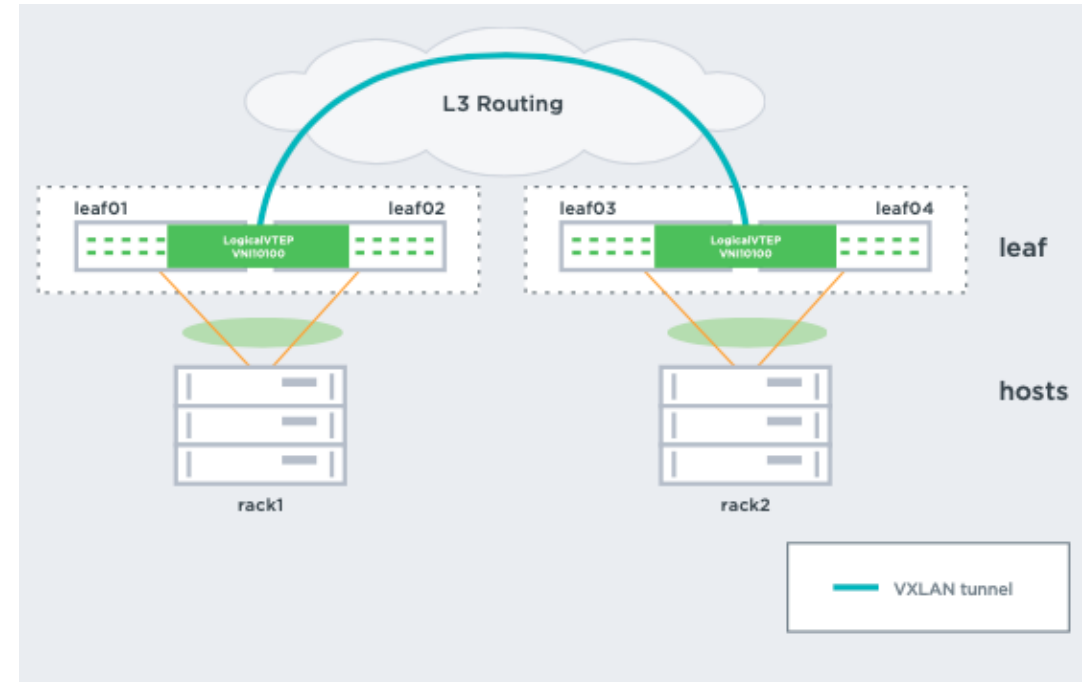
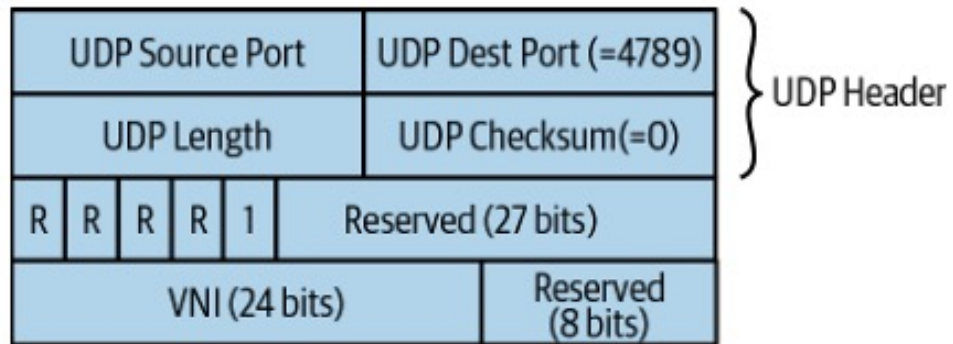


- Control plane for network virtualization
- Offers fast convergence VM moves
- Controller-less VXLAN tunnels
- Scale and robustness
- Interoperability between vendors
- MAC address learning



VXLAN

- Virtual Extensible LAN
- Overlay virtual network solution
- L2 virtual network over a L3 network infrastructure



Equal-Cost Multipath



- ECMP routing traffic forwarded to any available paths
- Traffic load balanced across links
- Equal distribution of flows across links

- Flows and Congestion

Elephant-mice problem

Traffic polarization

Automation



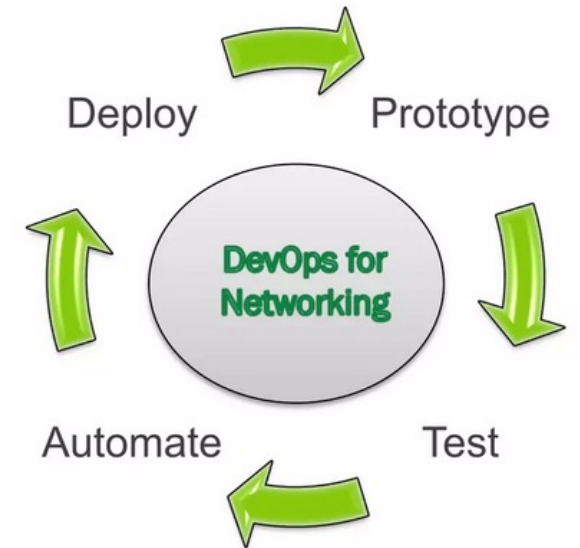
- Subversions repository
- Ansible
- Playbooks validation & PoC using containerlab



PoC & Validation



- **Cumulus VX**
<https://cumulusnetworks.com/cumulus-vx/>
- **Cumulus on the cloud**
<https://www.nvidia.com/en-us/networking/ethernet-switching/cumulus-linux/>
- **Containerlab (Docker container)**
<https://containerlab.dev/manual/kinds/cvx/>



UNDESIRE NEWS & MOVING ON



- NVIDIA Acquires Cumulus Networks (NOS)

End of support for non Spectrum ASICs up to version 4.x

Hardware cost constraints

- NVIDIA Acquires Mellanox Technologies

(Spectrum ASIC and Onyx NOS)

- Moving forward

Move for Spectrum line products

Potentially integrate Sonic NOS to remain on Trident ASIC





References



Nvidia Cumulus Linux:

<https://www.nvidia.com/en-us/networking/ethernet-switching/cumulus-linux/>

Containerlab:

<https://containerlab.dev/manual/kinds/cvx/>

Cloud Native Data Center Networkin, Dinesh G. Dutt





Thank You

