CUMULUS

EVPN to the host

Host multitenancy

Attilla de Groot

Attilla de Groot | Sr. Systems Engineer, HCIE #3494 | Cumulus Networks

Agenda

EVPN to the Host

- Multi tenancy use cases
- Deployment issues
- Host integration with EVPN
- Caveats and future work

Multi tenancy use cases

Virtualization environments

VM environments

- MLAG to hypervisors
- Vxlan between hosts
- Double overlay
- Dedicated network nodes

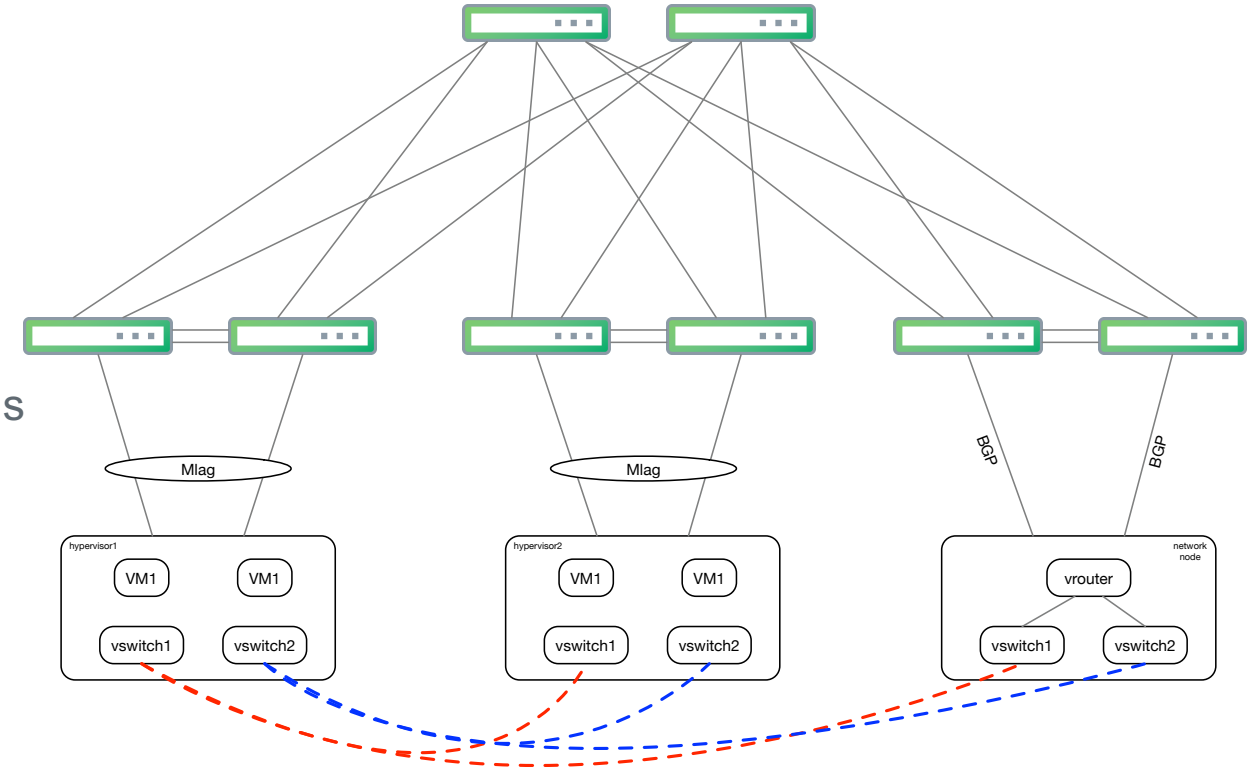

Multi tenancy use cases

VM deployment issues

Architecture issues

- MLAG to hypervisors
- Vxlan between hosts
- Double overlay
- Dedicated network nodes
- Orchestration problems

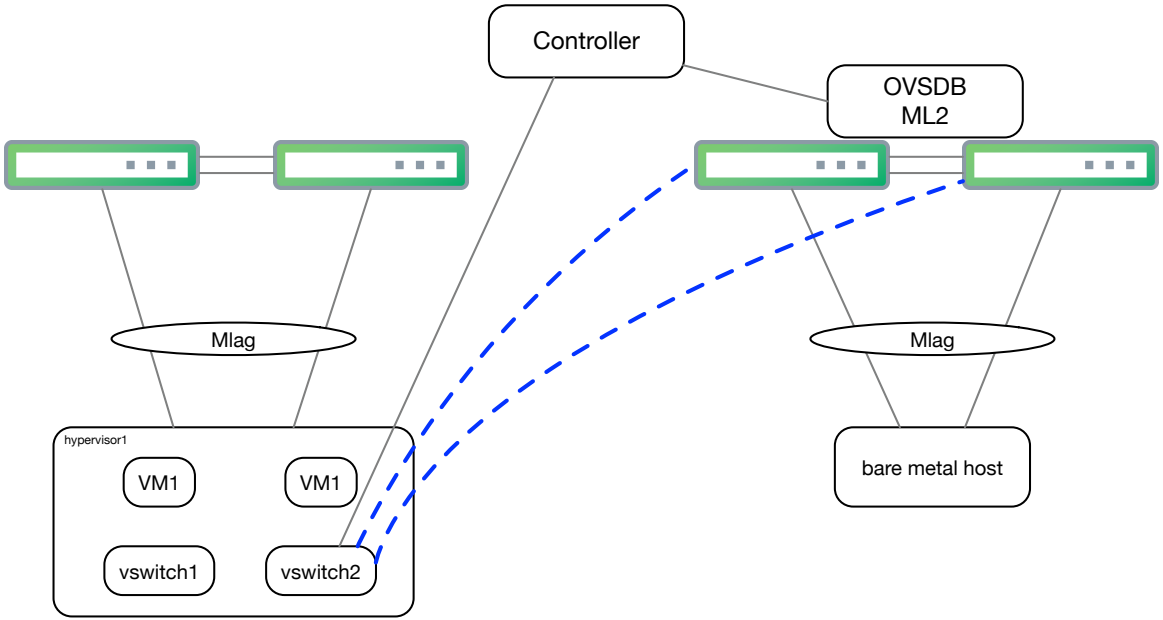

Multi tenancy use cases
Container environments

Container environments

- BGP to hosts
- Host route advertisement
- Docker management
- Container overlay

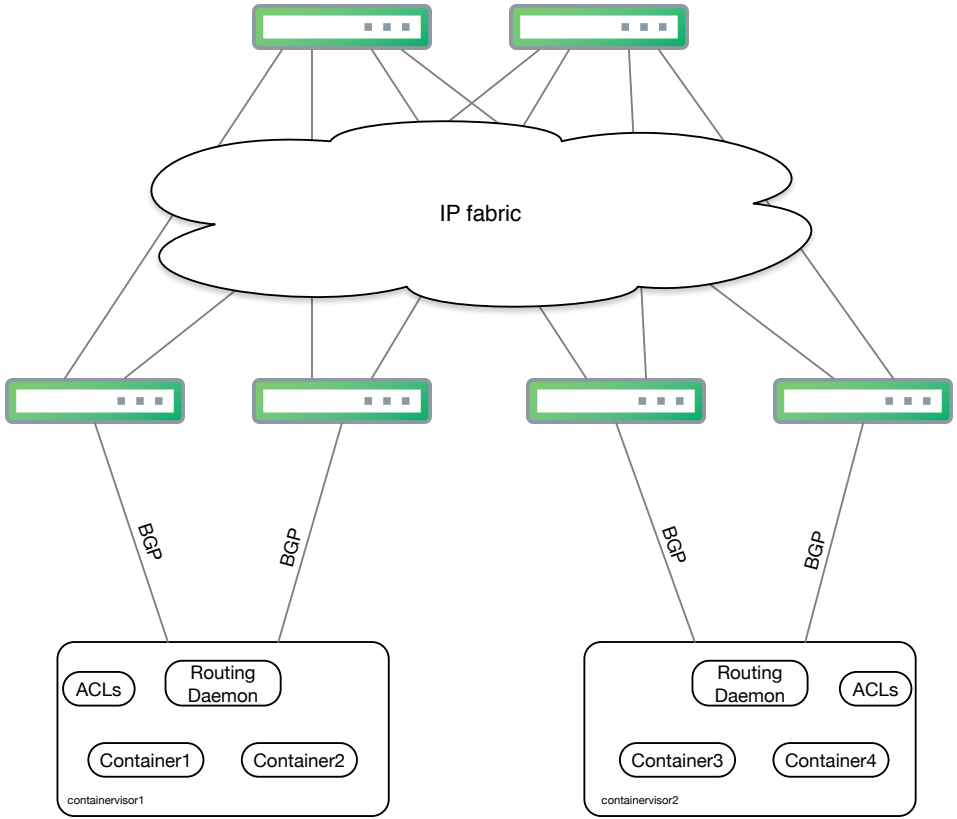

Multi tenancy use cases
Container deployment issues

Architecture issues

- Host networking
- Multi-tenancy
- IP prefix overlap
- ACL management

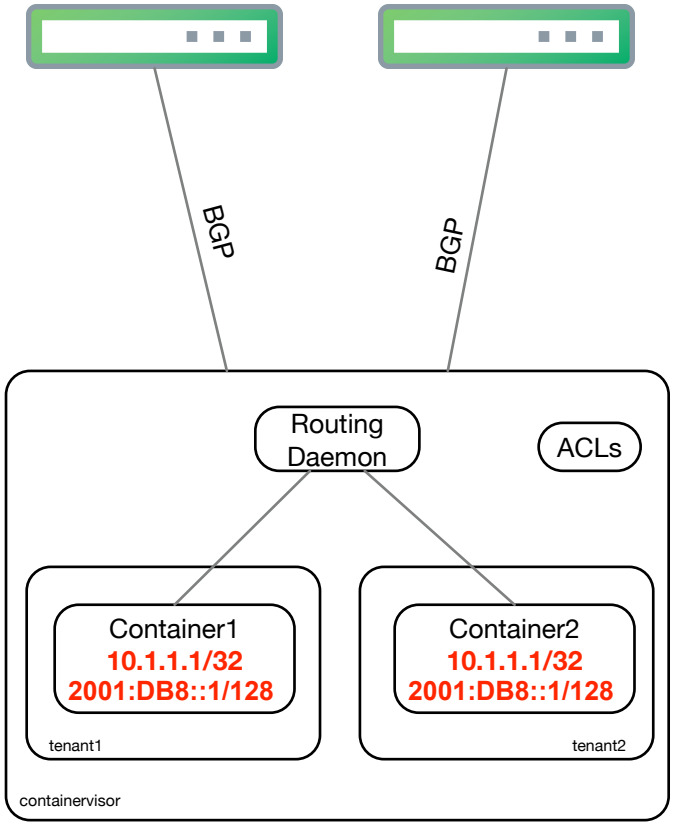

Host integration with EVPN
Open standards

EVPN on hosts

- Vlan aware bridge
- VRF in Linux kernel
- VxLAN
- Free Range Routing with EVPN
- Ifupdown2
- Iproute2

```
cumulus@server01:~$ uname -a
Linux server01 4.17.0-041700-generic #201806041953 SMP Mon
Jun 4 19:55:25 UTC 2018 x86_64 x86_64 x86_64 GNU/Linux

cumulus@server01:~$ vrf list

VRF            Table
---------------- -----
mgmt            1001
tenant3         1002
tenant2         1003
tenant1         1004
```

Host integration with EVPN
Host connectivity

Host connectivity

- Routing to hosts
- BGP unnumbered
- RFC5549
- Loopback advertisement
- EVPN address family

```
auto eth1
iface eth1

auto eth2
iface eth2
```

```
interface eth1
 ipv6 nd ra-interval 10
 no ipv6 nd suppress-ra
interface eth2
 ipv6 nd ra-interval 10
 no ipv6 nd suppress-ra
```

```
router bgp 65501
 bgp router-id 10.250.250.1
 bgp bestpath as-path multipath-relax
 neighbor FABRIC peer-group
 neighbor FABRIC remote-as external
 neighbor FABRIC timers 1 3
 neighbor eth1 interface peer-group FABRIC
 neighbor eth2 interface peer-group FABRIC
 !
 address-family ipv4 unicast
  redistribute connected route-map loopbacks
 exit-address-family
 !
 address-family l2vpn evpn
   neighbor FABRIC activate
   advertise-all-vni
   advertise ipv4 unicast
   advertise ipv6 unicast
 exit-address-family
!
route-map loopbacks permit 10
 match interface lo
```

L2 tenancy

- Vlan-aware bridge
- ARP / ND suppression
- L2VNIs
- EVPN Type-2

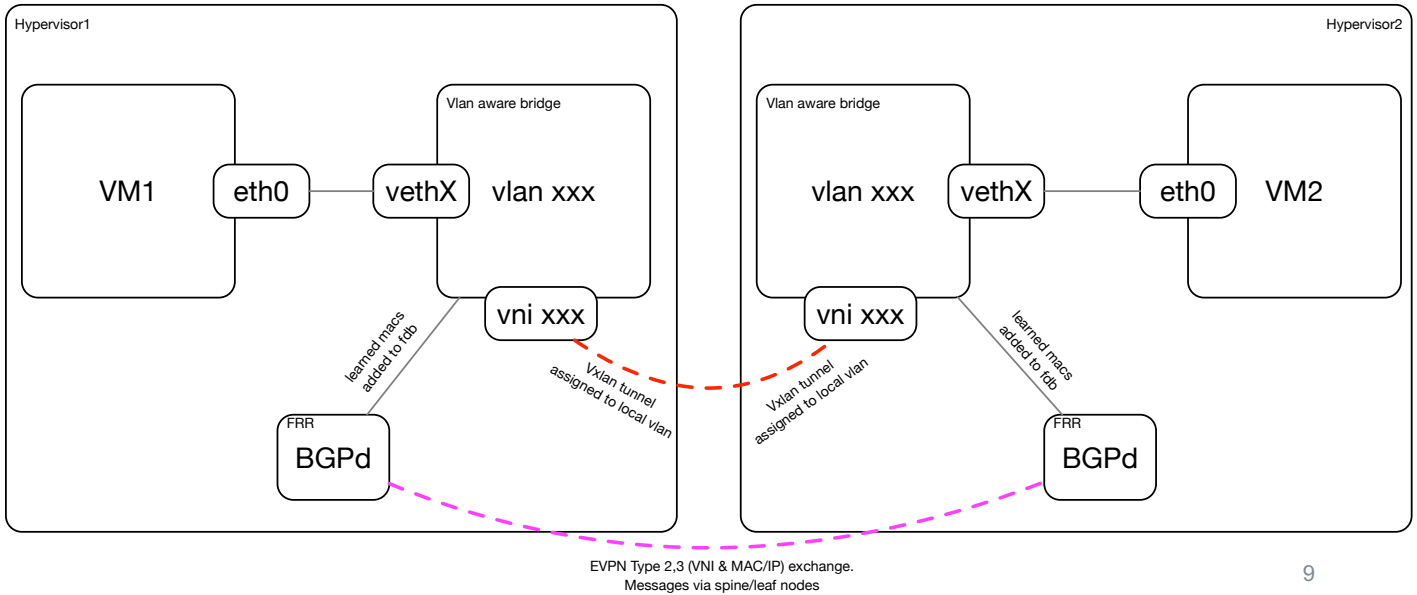

Host integration with EVPN
Distributed routing

Distributed routing

- Gateway location
- SVIs on host
- Anycast gateway
- Local host routing

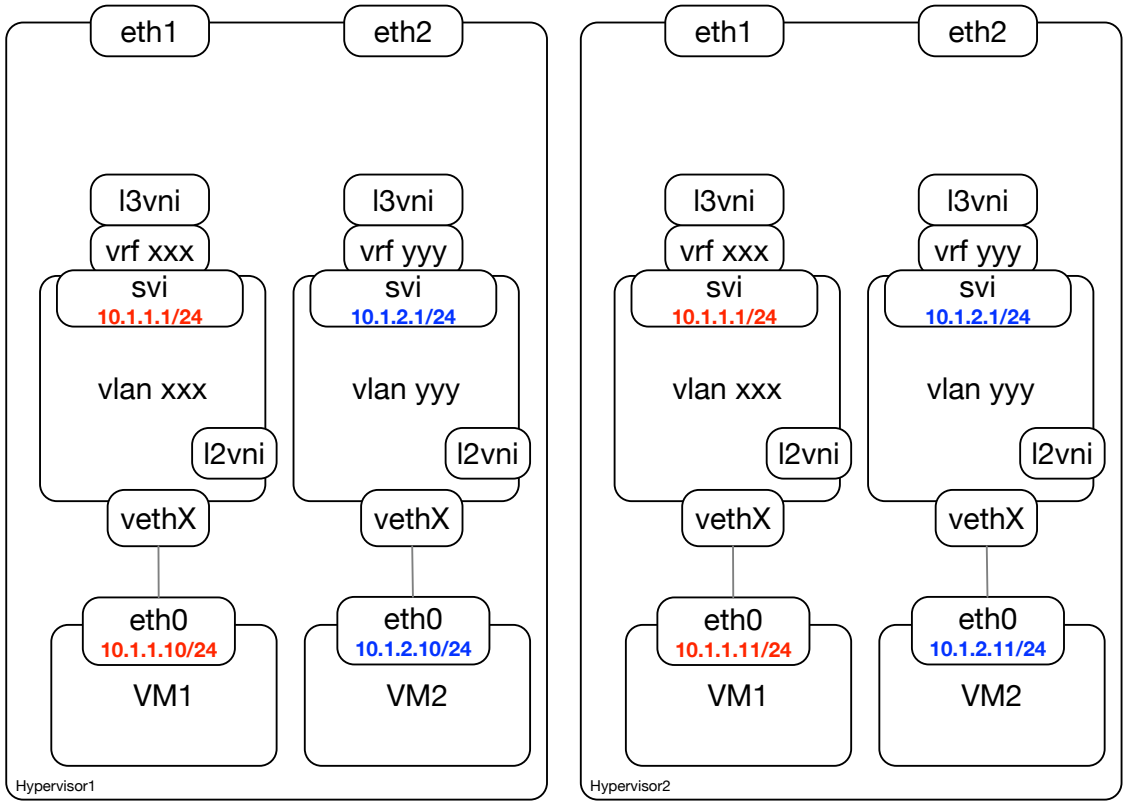

Host integration with EVPN
L3 tenancy

L3 tenancy

- VRFs on host
- Prefix advertisement
- L3VNIs
- EVPN Type-5

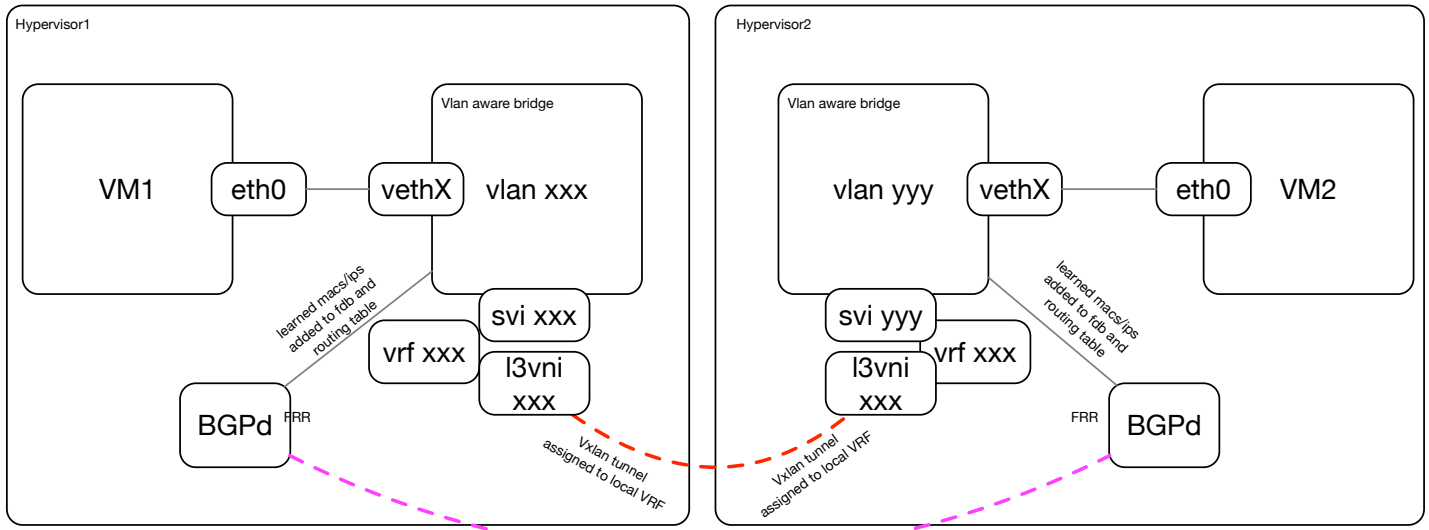

Host integration with EVPN
Container integration

Container integration

- Container IP redistribution
- Host route advertisement
- Prefix overlap
- No ACLs for tenant segregation

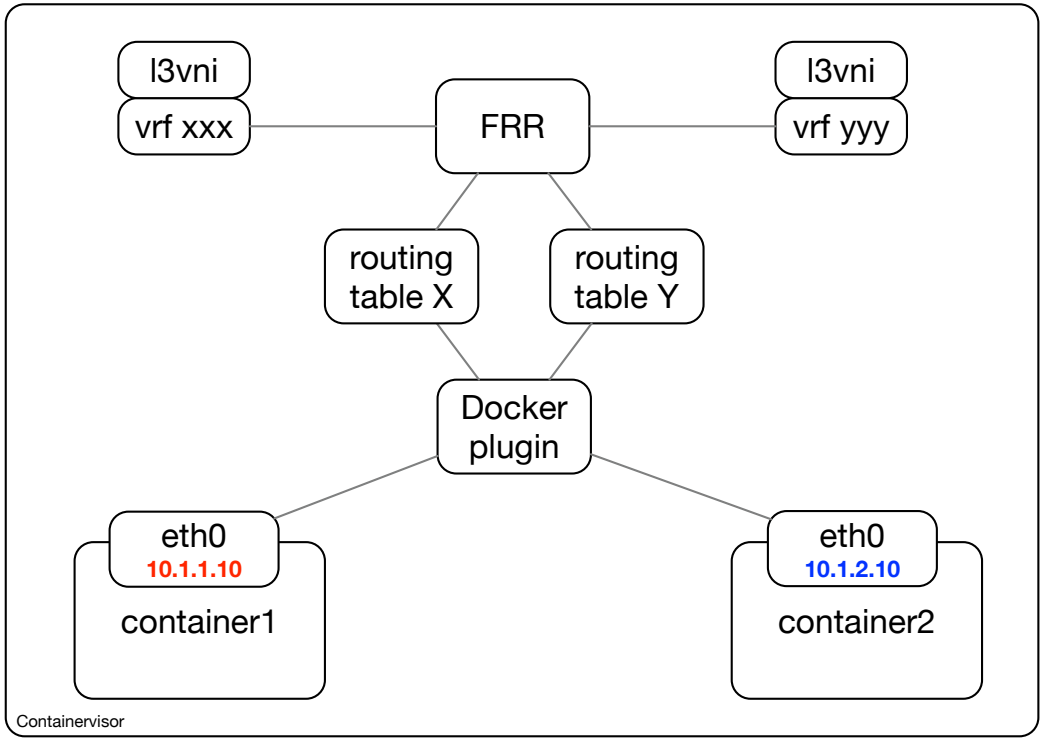

Bare metal integration

- Integration with RFC standard
- L2 stretching
- L3 tenancy
- Distributed routing

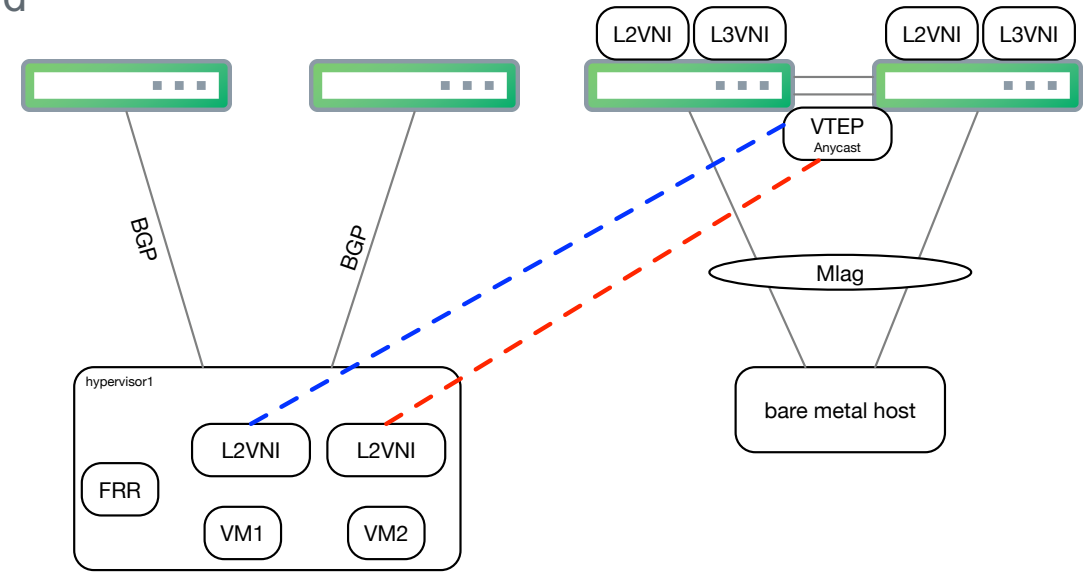

Proof of technology

- Commercial support
- Software / tools availability

  Linux kernel

  FRR
- Demo

```
cumulus@server01:~$ uname -a
Linux server01 4.17.0-041700-generic
#201806041953 SMP Mon Jun 4 19:55:25 UTC
2018 x86_64 x86_64 x86_64 GNU/Linux
```

```
cumulus@server01:~$ bridge vlan
port vlan ids
bridge 1 PVID Egress Untagged
1010
1011
1012
1013
1014
```

```
cumulus@server01:~$ vrf list


VRF            Table
-------------- -----
mgmt           1001
tenant3        1002
tenant2        1003
tenant1        1004
```

```
cumulus@server01:~$ sudo vtysh -c "show bgp evpn vni"
  VNI       Type RD                  Import RT           Export RT           Tenant VRF
* 101010    L2   10.250.250.1:5      65501:101010        65501:101010        tenant1
* 101012    L2   10.250.250.1:3      65501:101012        65501:101012        tenant2
* 101014    L2   10.250.250.1:7      65501:101014        65501:101014        tenant3
* 404001    L3   172.30.11.1:8       65501:404001        65501:404001        tenant1
* 404002    L3   172.30.13.1:9       65501:404002        65501:404002        tenant2
* 404003    L3   172.30.15.1:10      65501:404003        65501:404003        tenant3
```

https://gitlab.com/CumulusNetworks/evpn-to-the-host

Future work
Orchestration & Caveats

Orchestration

- Openstack neutron
- Kubernetes / Swarm
- Host / Network management

Head-end replication

- BUM anycast
- Merchant silicon limitations
- Multicast replication
- PIM-SM

Future work
Route leaking & Micro segmentation

Route-leaking

- Inter tenant traffic
- Service tenant
- FRR implementation

Micro segmentation

- Host traffic filtering
- Filtering with BPF
- Flowspec for ACL distribution

Questions ?

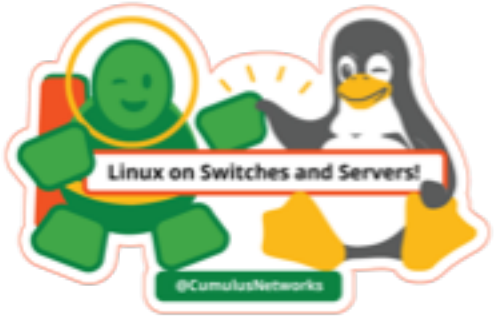

Thank you!

Visit us at cumulusnetworks.com or follow us @cumulusnetworks