



# VCF TechCon

Powered by VMUG

## Troubleshooting EVPN in een NSX omgeving

- Robin van Altena

# Robin van Altena

- IT, SDDC, Private Cloud Consultant @ *itq*
- vExpert en Broadcom knight



# Agenda

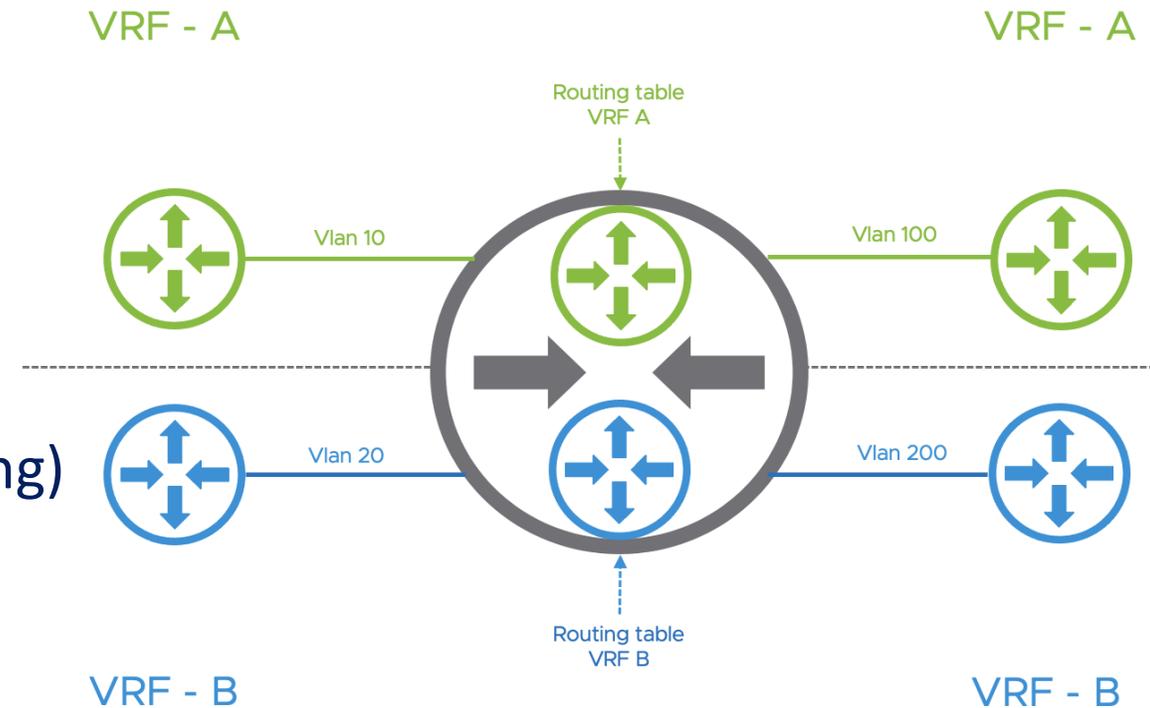
- Intro
  - NSX VRF-Lite en EVPN
  - Cisco Multi Site EVPN
- Omgeving
- Oeps.
- Troubleshooting..

**DISCLAIMER**

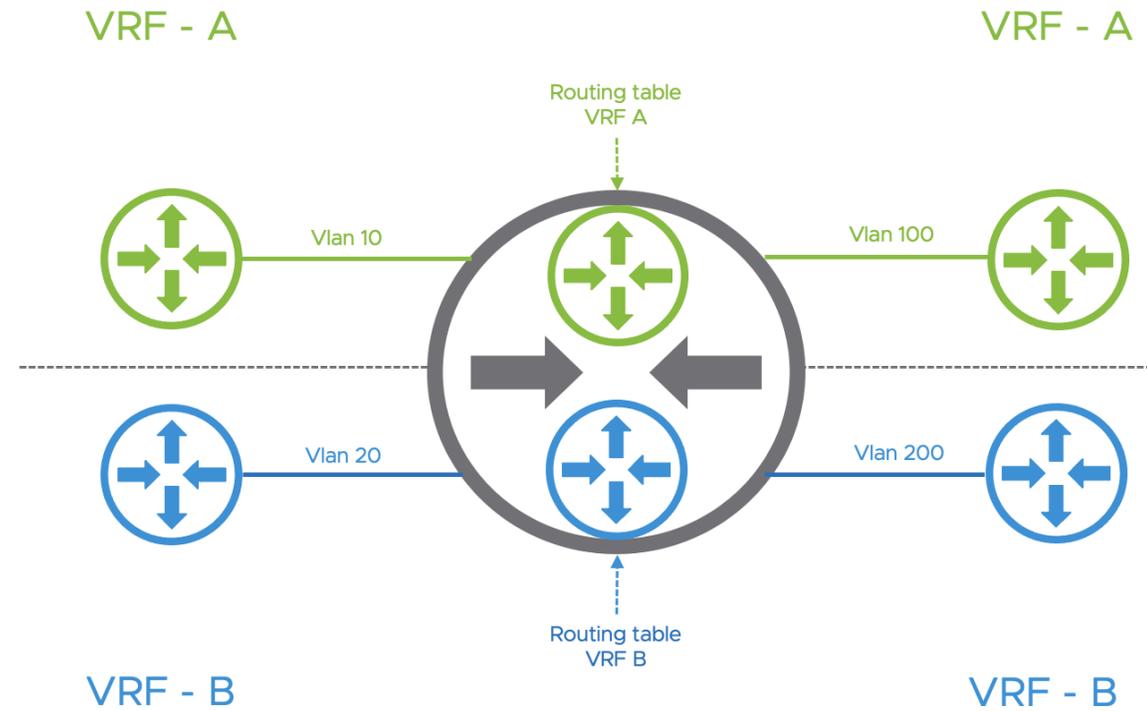
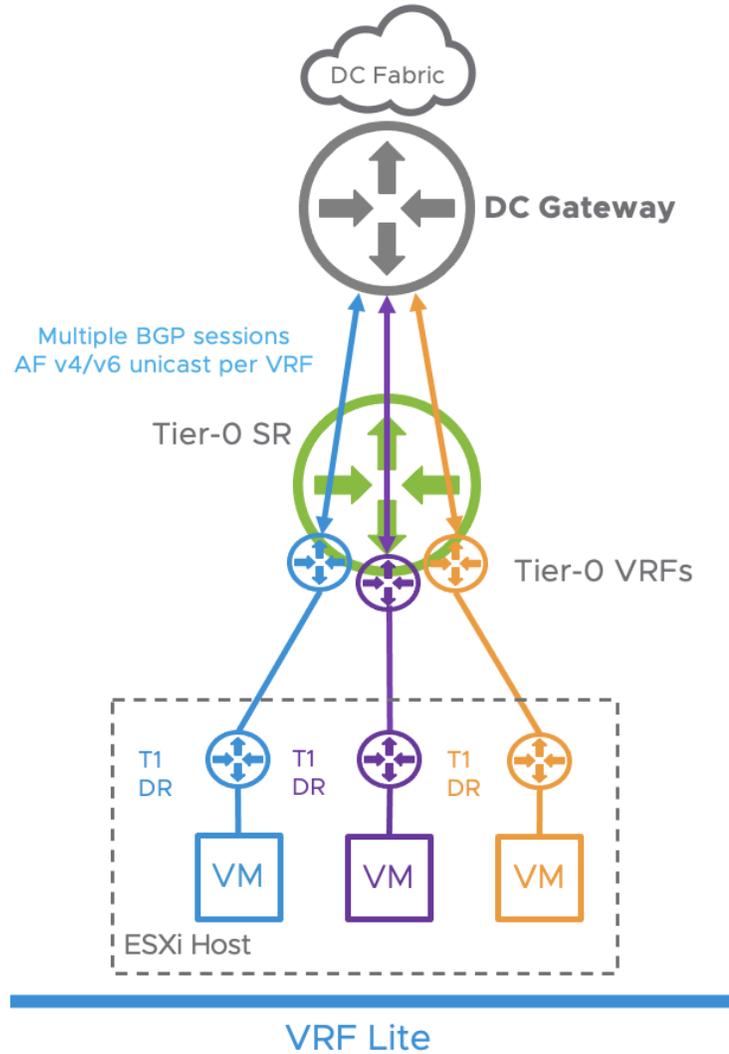


# Intro

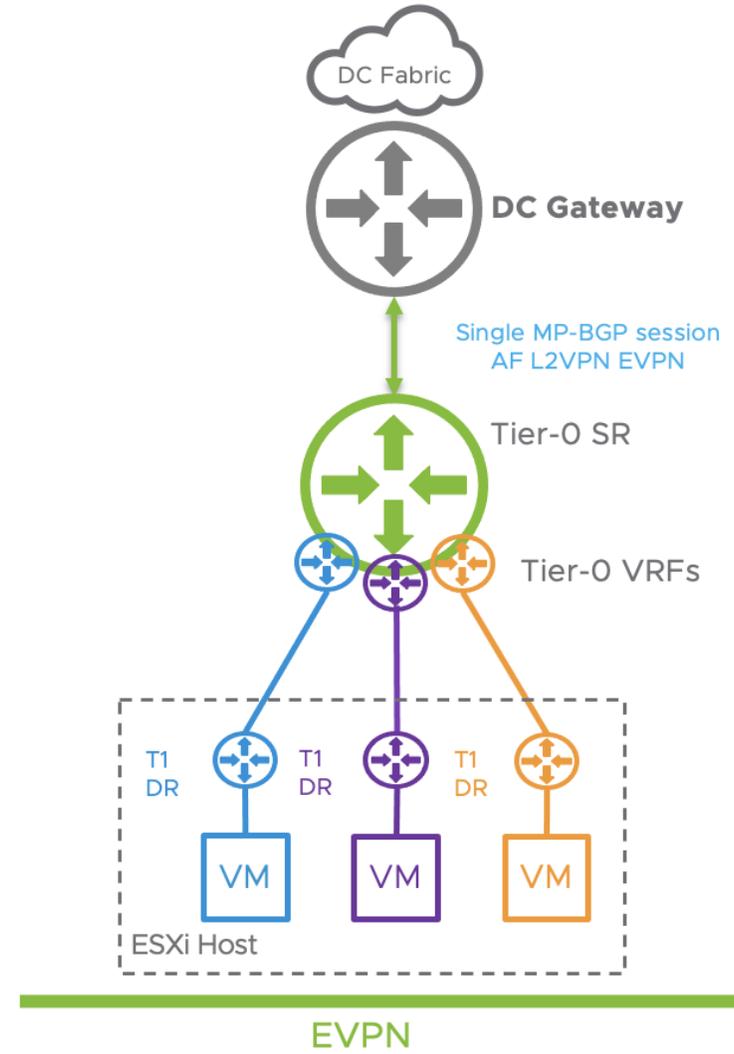
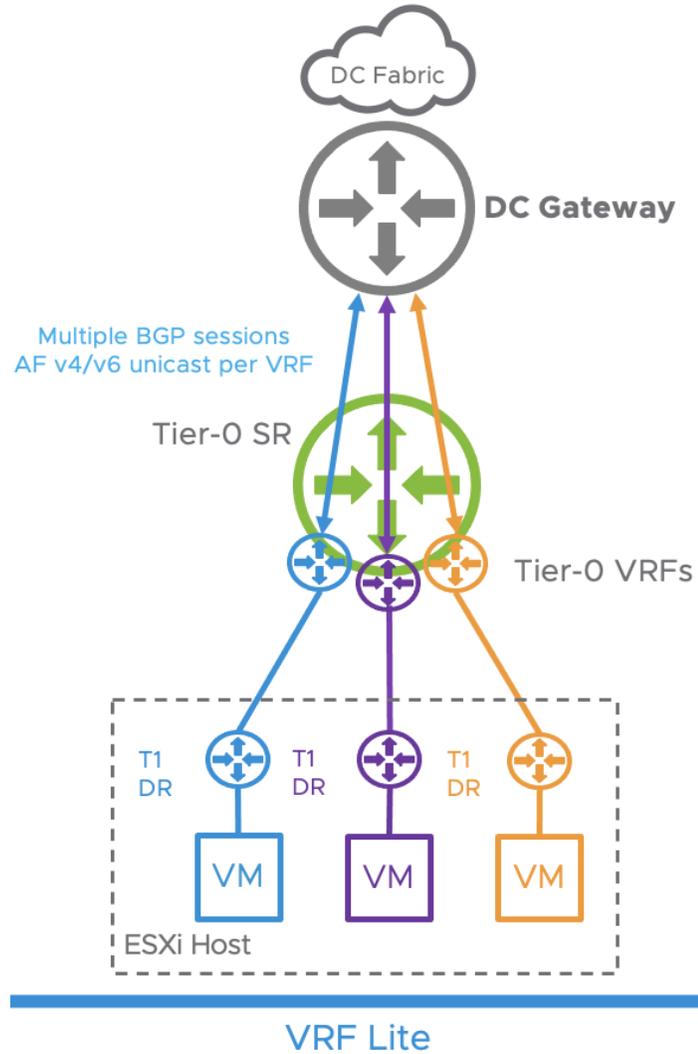
- Switching
  - VLAN,
  - Overlay (VXLAN / GENEVE)
- Routing
  - BGP
  - Multi Protocol BGP (MP-BGP)
  - VRF (Virtual Routing & Forwarding)
- Multi Tenant
- virtual Port Channel (vPC)



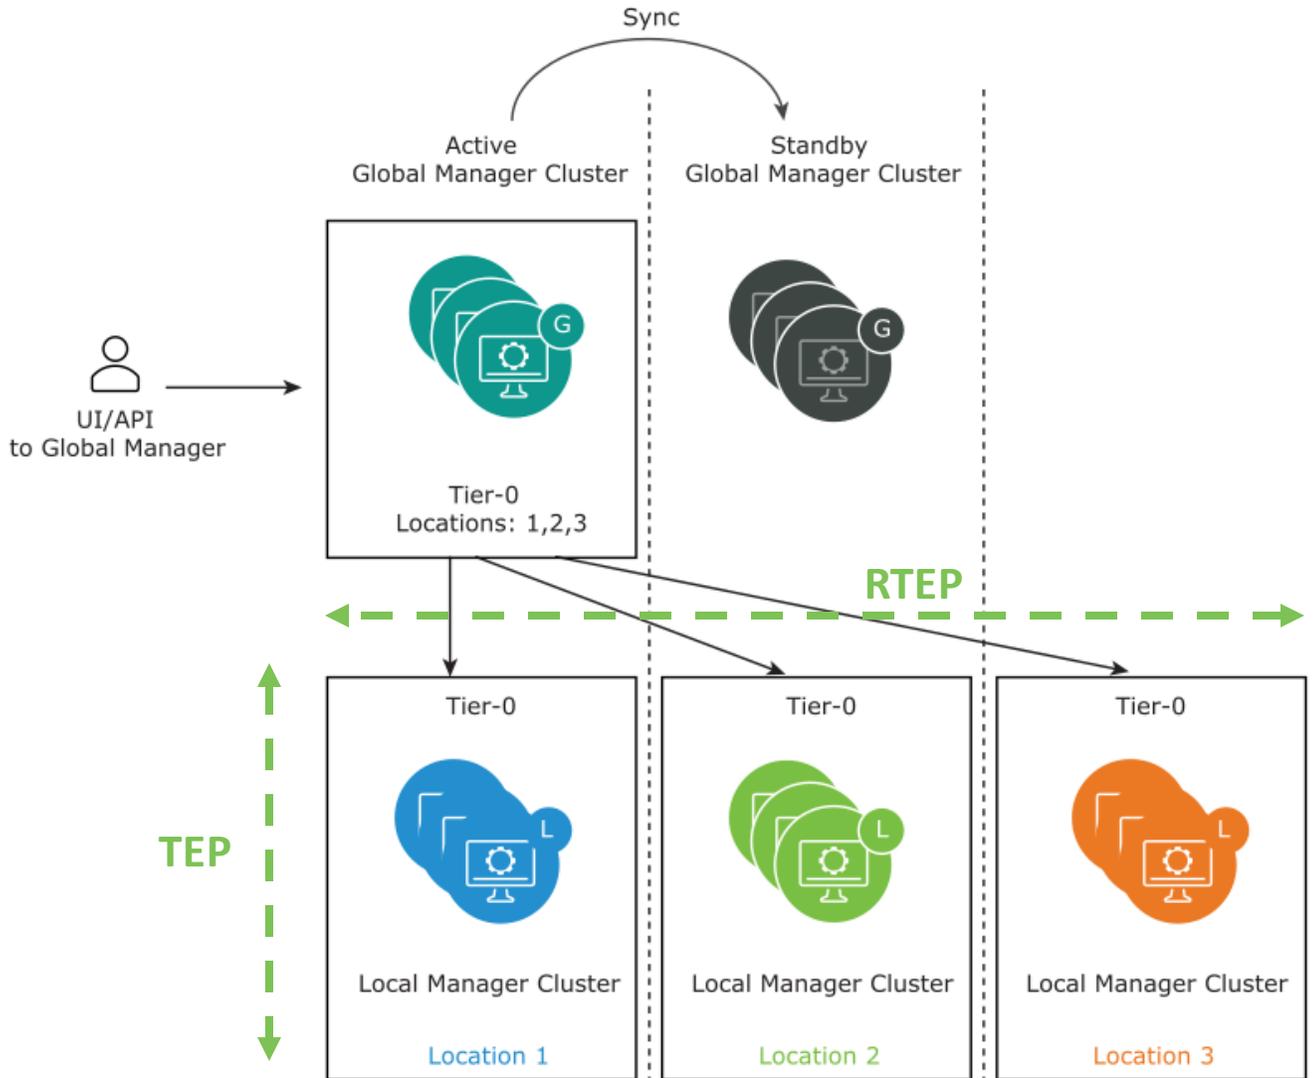
# NSX VRF-Lite



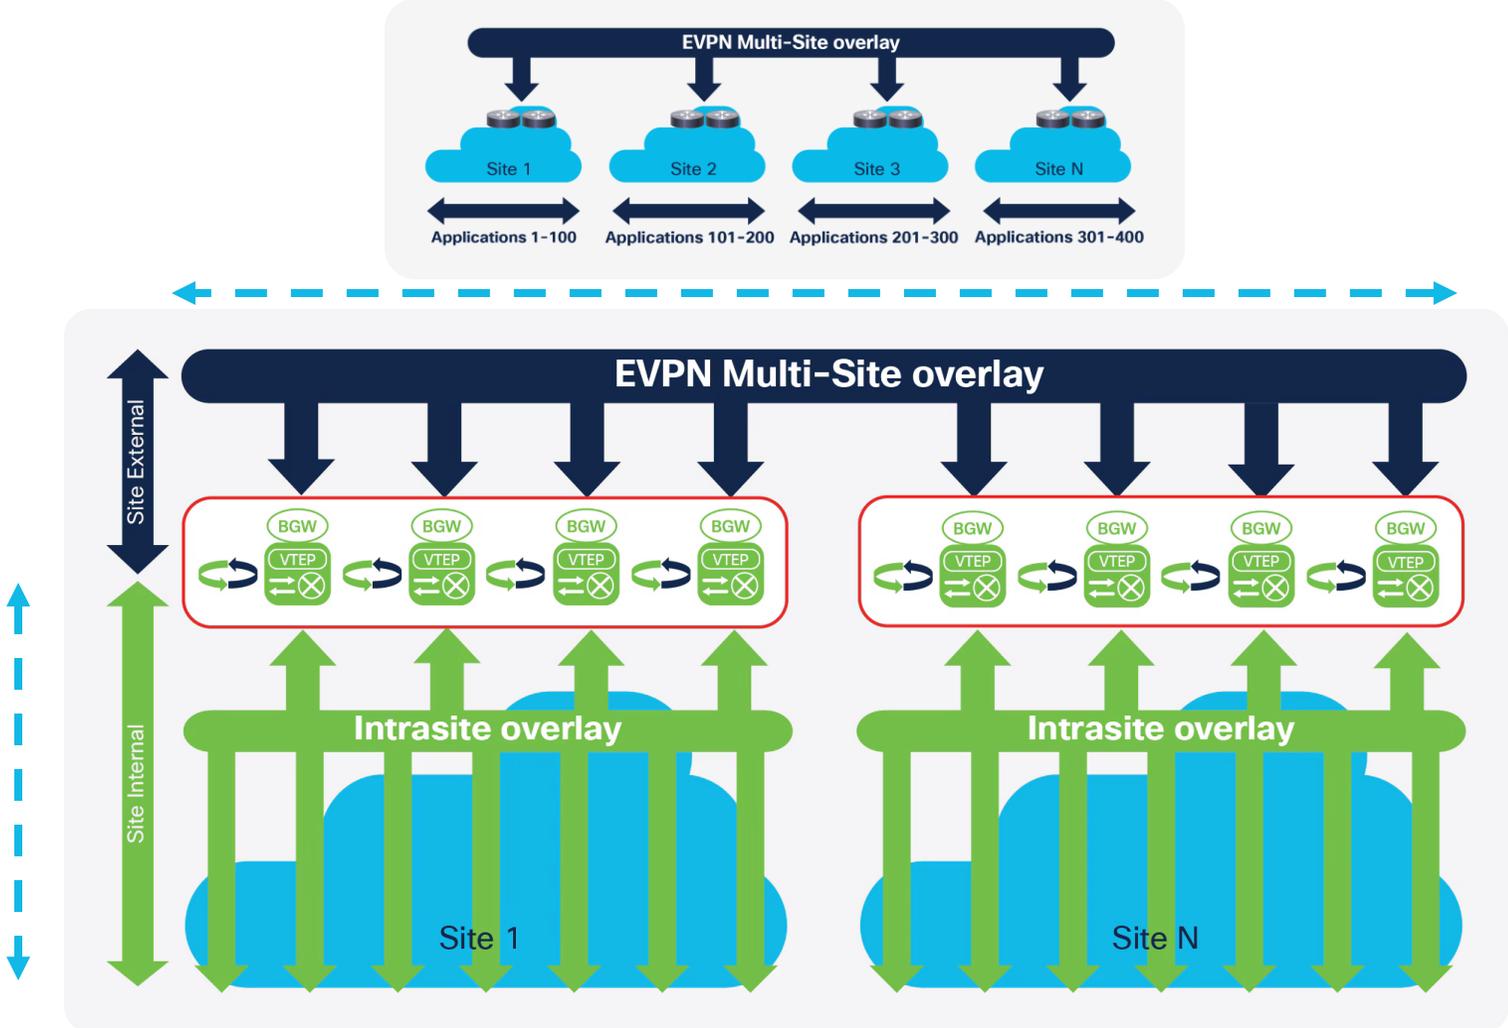
# NSX VRF-Lite en EVPN



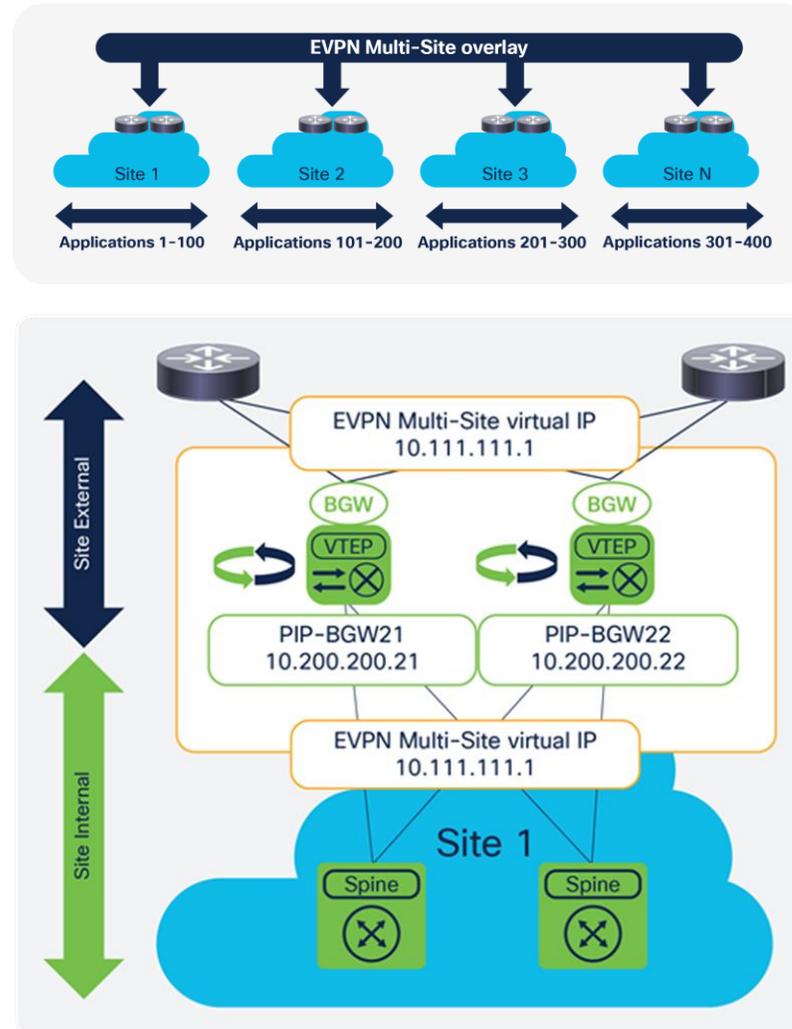
# NSX Federation



# Cisco Multi Site EVPN

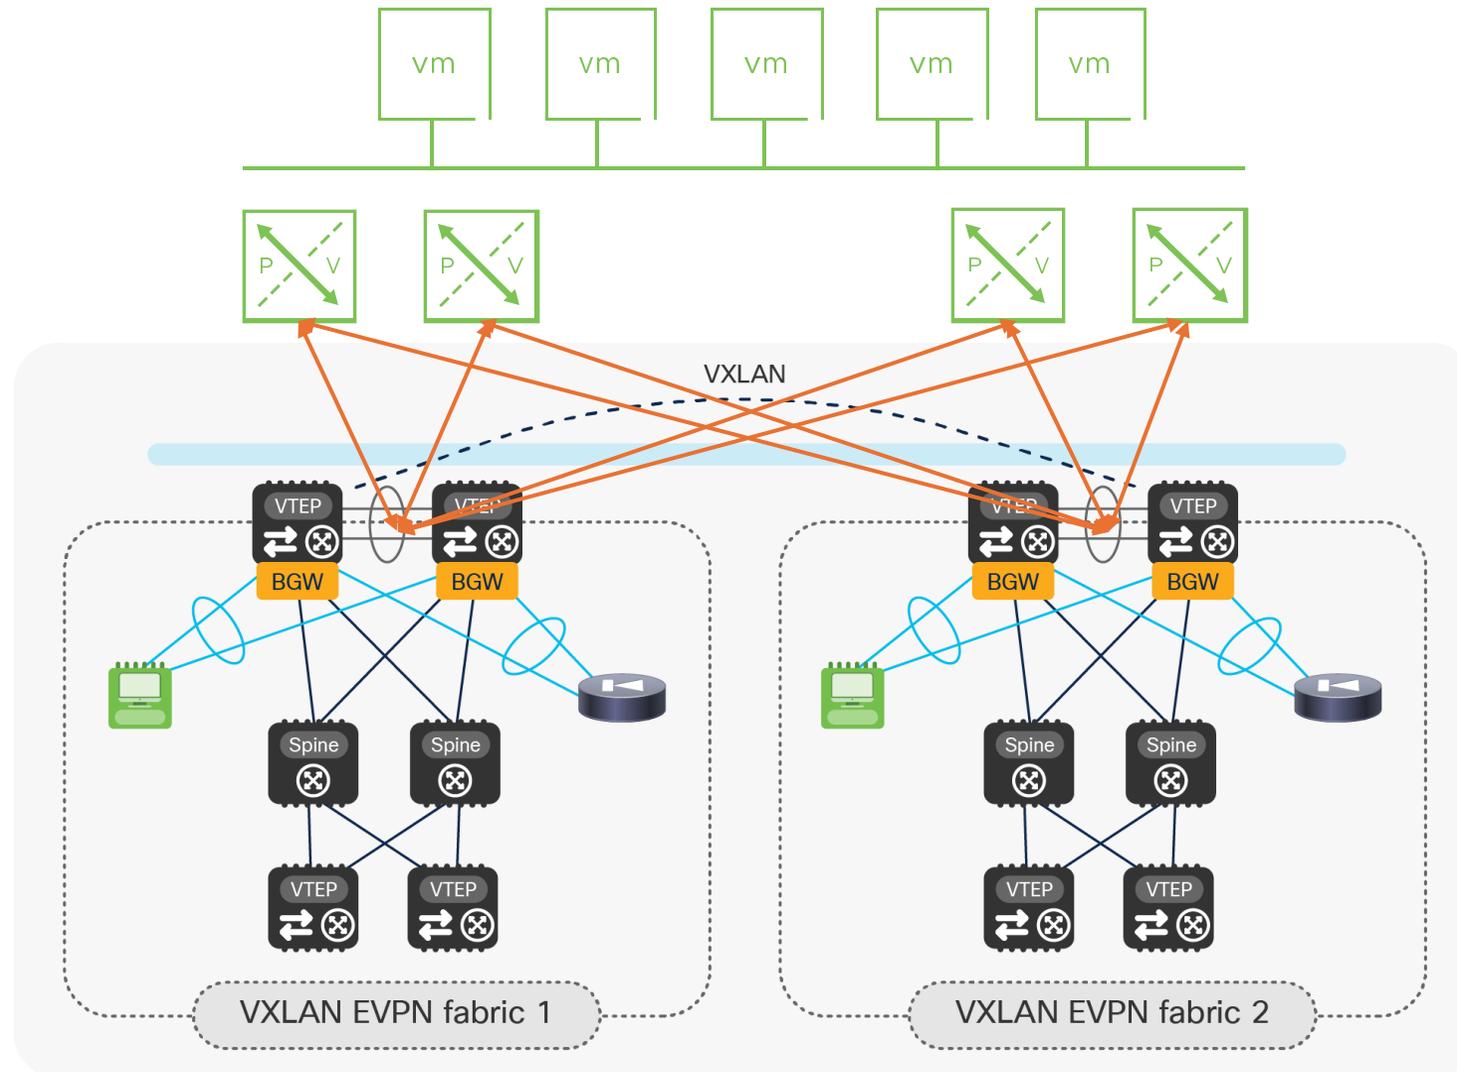


# Cisco Multi Site EVPN





# Cisco Multi Site EVPN en NSX EVPN



BGP sessies

# NSX EVPN Config

### Set Interfaces

Tier-0 Gateway | l-tier0 | #Interfaces 12

[ADD INTERFACE](#) [EXPAND ALL](#) Filter by Name, Path and more

Name	Type	IP Address / Mask	Connected To(Segment)	Status
> .edge-01-loopback	Loopback	46.126/32		Success
> .edge-01-uplink-01	External	46.97/30	.edge-upli... VLAN ⓘ	Success
> .edge-01-uplink-02	External	46.105/30	.edge-upli... VLAN ⓘ	Success

### Set BGP Neighbors

Tier-0 Gateway | tier0 | #Neighbors 12

[ADD BGP NEIGHBOR](#) [EXPAND ALL](#) Filter by Name, Path and more

IP Address	BFD	Remote AS number	Route Filter	Allowas-in	Status
> .138.110	On	11	1	No	Success
> .138.3	Off	11	2	No	Success
> .138.98	On	11	1	No	Success
> .138.6	Off	11	2	No	Success

### security

Connect to Tier-0 Gateway: -tier0

Edge Cluster: -edge-01

DHCP Config: Not Set

Additional Settings

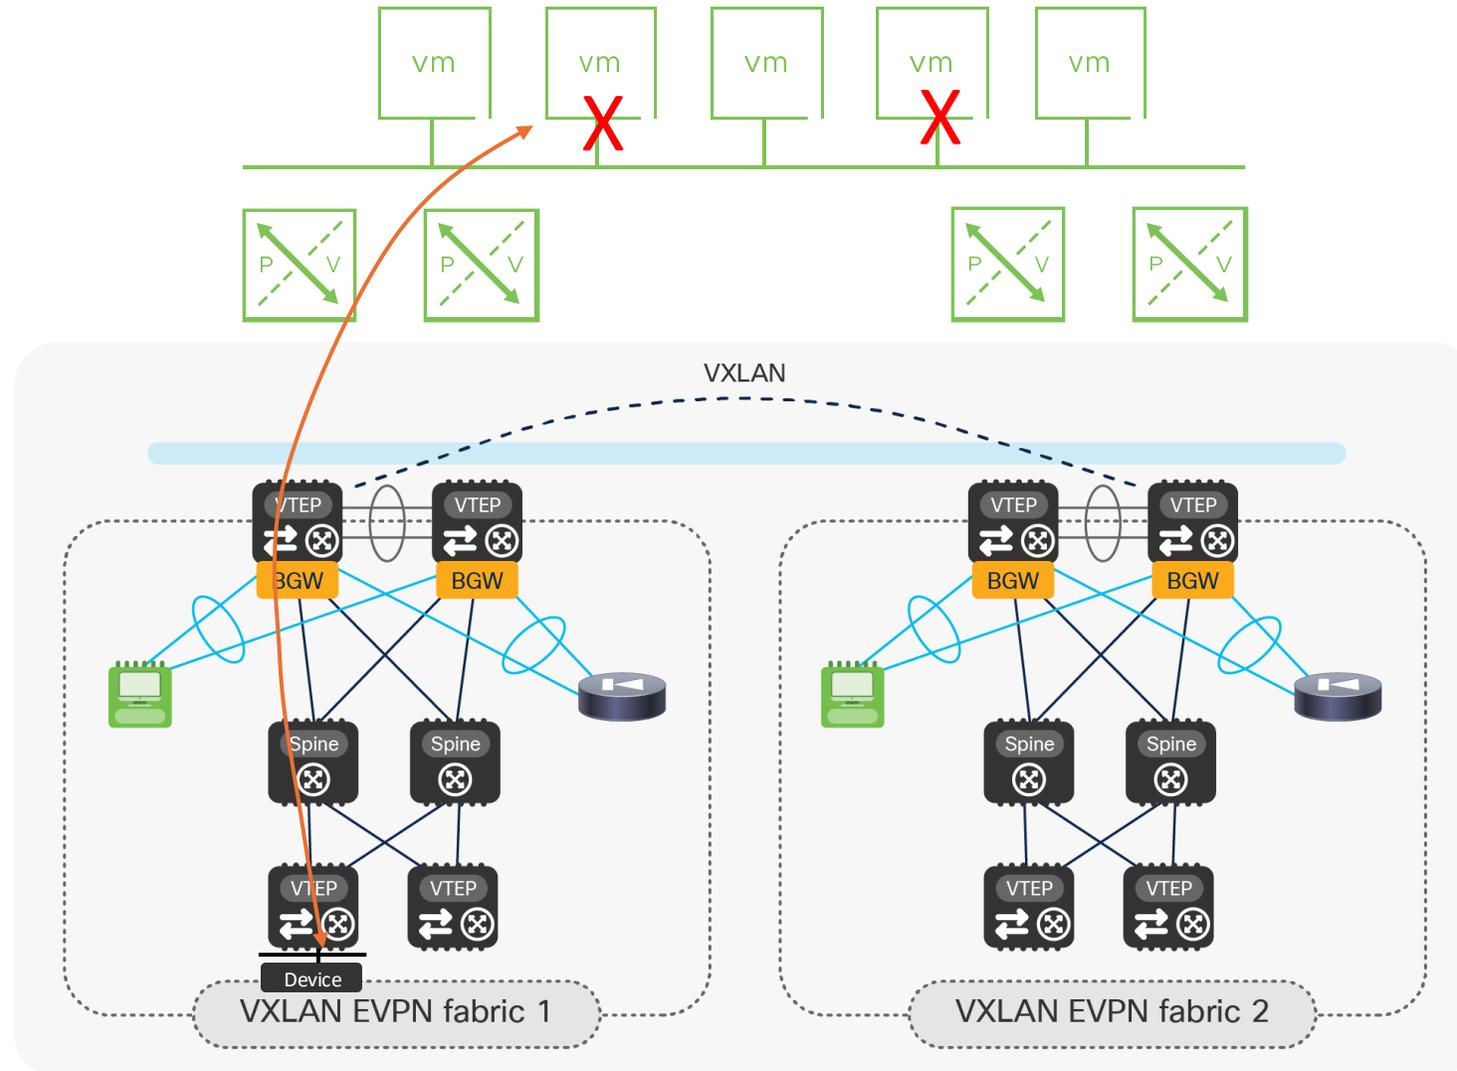
VRF Settings

L3 VRF Settings

Route Distinguisher: 44:75010

EVPN Transit VNI: 75010

# Oeps...



# NSX Config

```
Tue Nov 25 2025 CET 07:42:36.270
-01(tier0_sr)> get bgp
BGP IPv4 table version is 43, BGP IPv6 table version is 0, BGP EVPN table version is 9118
Local router ID is .126
Status flags: > - best, I - internal
Origin flags: i - IGP, e - EGP, ? - incomplete

EVPN type-2 prefix: [2]:[EthTag]:[MACLen]:[MAC]:[IPlen]:[IP]
EVPN type-3 prefix: [3]:[EthTag]:[IPlen]:[OrigIP]
EVPN type-4 prefix: [4]:[ESI]:[IPlen]:[OrigIP]
EVPN type-5 prefix: [5]:[EthTag]:[IPlen]:[IP]

Network                Next Hop                Metric    LocPrf  Weight  Path          RD
> 130.1/32              .146.106                0         100    0       22 11 ?
> 130.1/32              .146.98                 0         100    0       22 11 ?
> 130.2/32              .146.106                0         100    0       22 ?
> 130.2/32              .146.98                 0         100    0       22 ?
> 130.16/30             .146.106                0         100    0       22 ?
> 130.16/30             .146.98                 0         100    0       22 ?
> 130.20/30             .146.106                0         100    0       22 ?
> 130.20/30             .146.98                 0         100    0       22 ?
> 130.24/30             .146.106                0         100    0       22 ?
> 130.24/30             .146.98                 0         100    0       22 ?
> 130.28/30             .146.106                0         100    0       22 ?
```

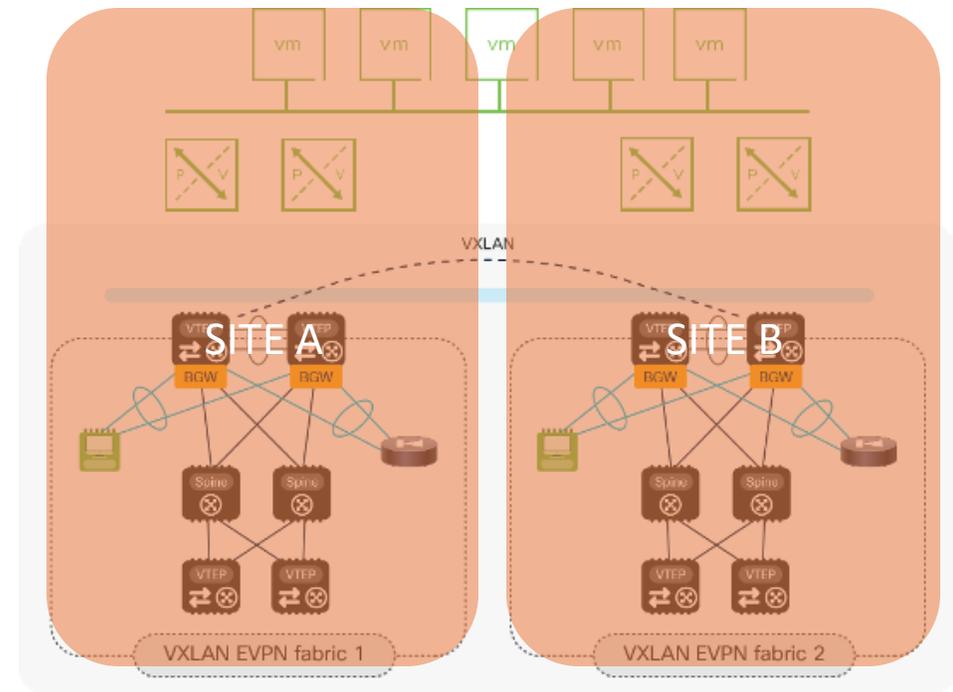
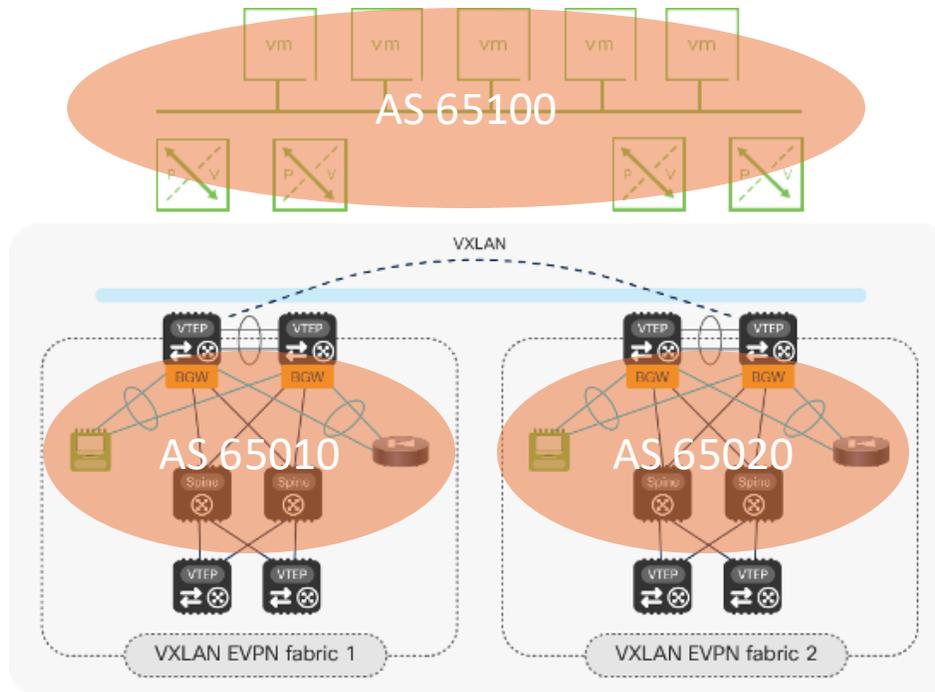
```
-edge-01> vrf 0
-01(tier0_sr)> get bgp evpn
BGP table version is 9118, local router ID is :
Status codes: s suppressed, d damped, h history, * valid,
Origin codes: i - IGP, e - EGP, ? - incomplete
EVPN type-1 prefix: [1]:[EthTag]:[ESI]:[IPlen]:[VTEP-IP]
EVPN type-2 prefix: [2]:[EthTag]:[MACLen]:[MAC]:[IPlen]:[IP]
EVPN type-3 prefix: [3]:[EthTag]:[IPlen]:[OrigIP]
EVPN type-4 prefix: [4]:[ESI]:[IPlen]:[OrigIP]
EVPN type-5 prefix: [5]:[EthTag]:[IPlen]:[IP]

Network                Next Hop                Metric LocPrf Wei
Extended Community
Route Distinguisher: 65111:20000
*> [2]:[0]:[48]:[00:50:56:52:f4:51]
.130.1 2000
RT:65111:20000 ET:8
* [2]:[0]:[48]:[00:50:56:52:f4:51]
.130.1 2000
RT:65111:20000 ET:8
*> [2]:[0]:[48]:[00:50:56:53:83:a6]
130.1 2000
RT:65111:20000 ET:8
```

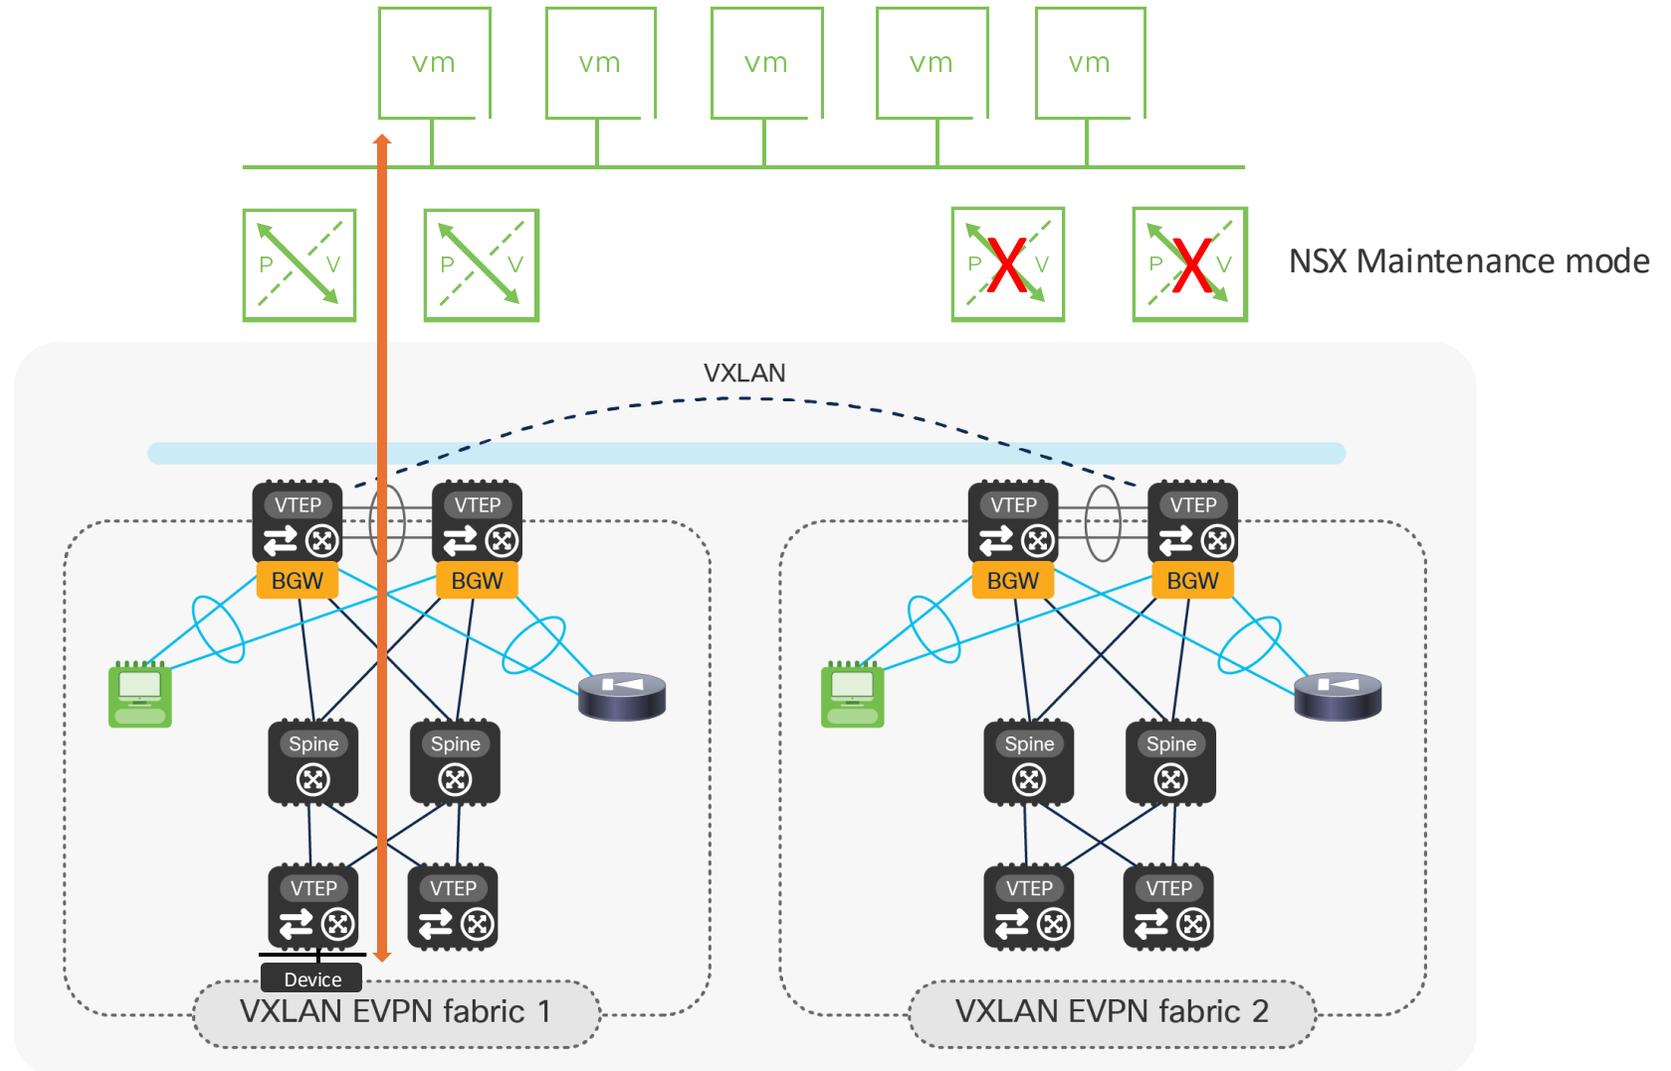
```
-01> vrf 5
-01(tier0_vrf_sr[5])> get bgp
BGP IPv4 table version is 13, BGP IPv6 table version is 0
Local router ID is .126
Status flags: > - best, I - internal
Origin flags: i - IGP, e - EGP, ? - incomplete

Network                Next Hop                Metric    LocPrf  Weight  Path
0.0.0.0/0              .138.194                0         100    0       65111 65500 ?
0.0.0.0/0              .138.196                0         100    0       65111 65500 ?
> 0.0.0.0/0            .146.198                0         100    0       65222 ?
> 0.0.0.0/0            .146.195                0         100    0       65222 ?
> .112.0/27            0.0.0.0                 0         100    32768  ?
> .112.32/27          0.0.0.0                 0         100    32768  ?
> 192.168.0.14/32     .138.198                0         100    0       65111 65500 i
192.168.0.14/32     .138.198                0         100    0       65111 65500 i
192.168.0.16/30     .138.194                0         100    0       65111 65500 ?
192.168.0.16/30     .138.196                0         100    0       65111 65500 ?
> 192.168.0.16/30     .146.198                0         100    0       65222 ?
> 192.168.0.16/30     .146.195                0         100    0       65222 ?
```

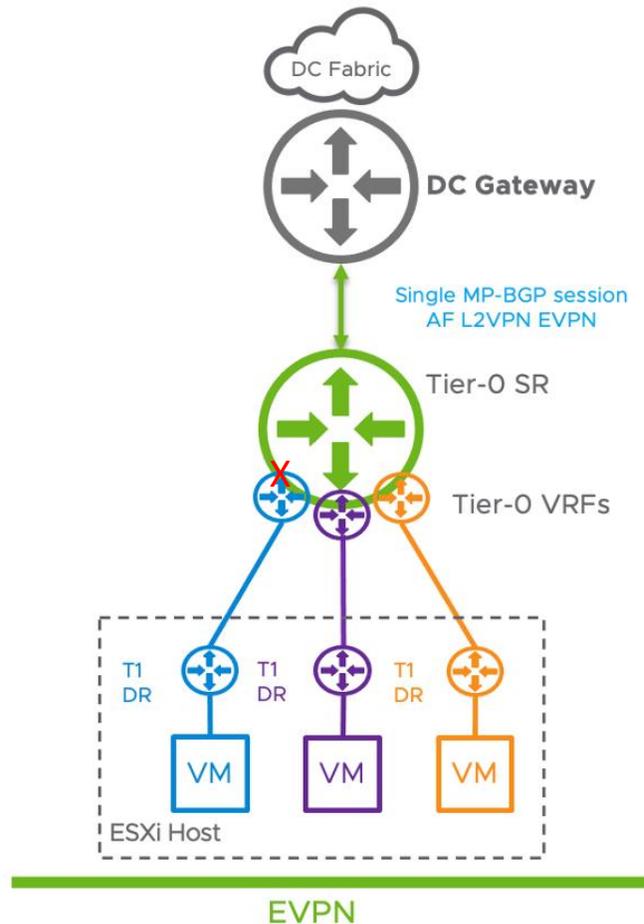
# Troubleshooting..



# Troubleshooting..



# Troubleshooting op de NSX edge



```
Interface      : 5fdf5a2d-4d66-4fdc-adac-e92214cef907
Ifuid          : 276
Name           : sr0-evpn-transit-port
Fwd-mode       : IPV4_ONLY
Mode           : lif
Port-type      : uplink
IP/Mask        :
MAC            : 02:50:56:00:48:00
VNI            : 75004
Access-VLAN    : untagged
LS port        : 1df9d0c4-7d74-47eb-874a-6f2c4deb970b
Urpf-mode      : PORT_CHECK
DAD-mode       : LOOSE
RA-mode        : RA_INVALID
Admin          : up
Op_state       : up
Enable-mcast   : False
MTU            : 8900
arp_proxy      :
```



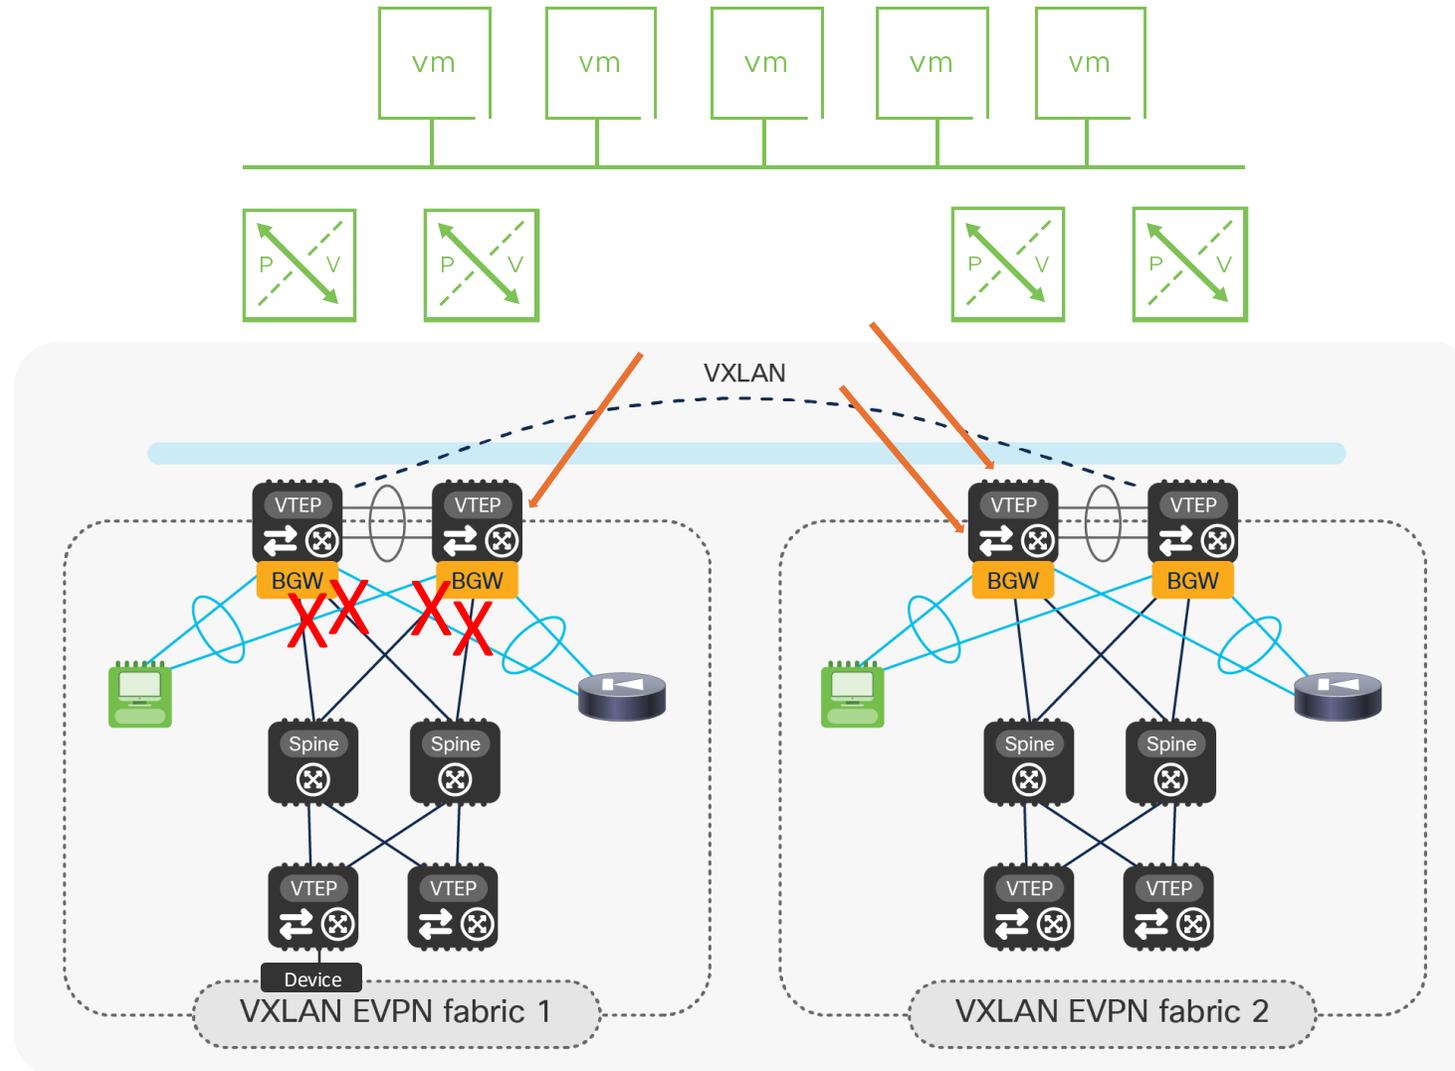
# Troubleshooting op de ESX host

- Net-stats -l
- Pktcap-uw --switchport <uuid> --capture <point> -o - | tcpdump-uw -r- -nne | grep ip-address

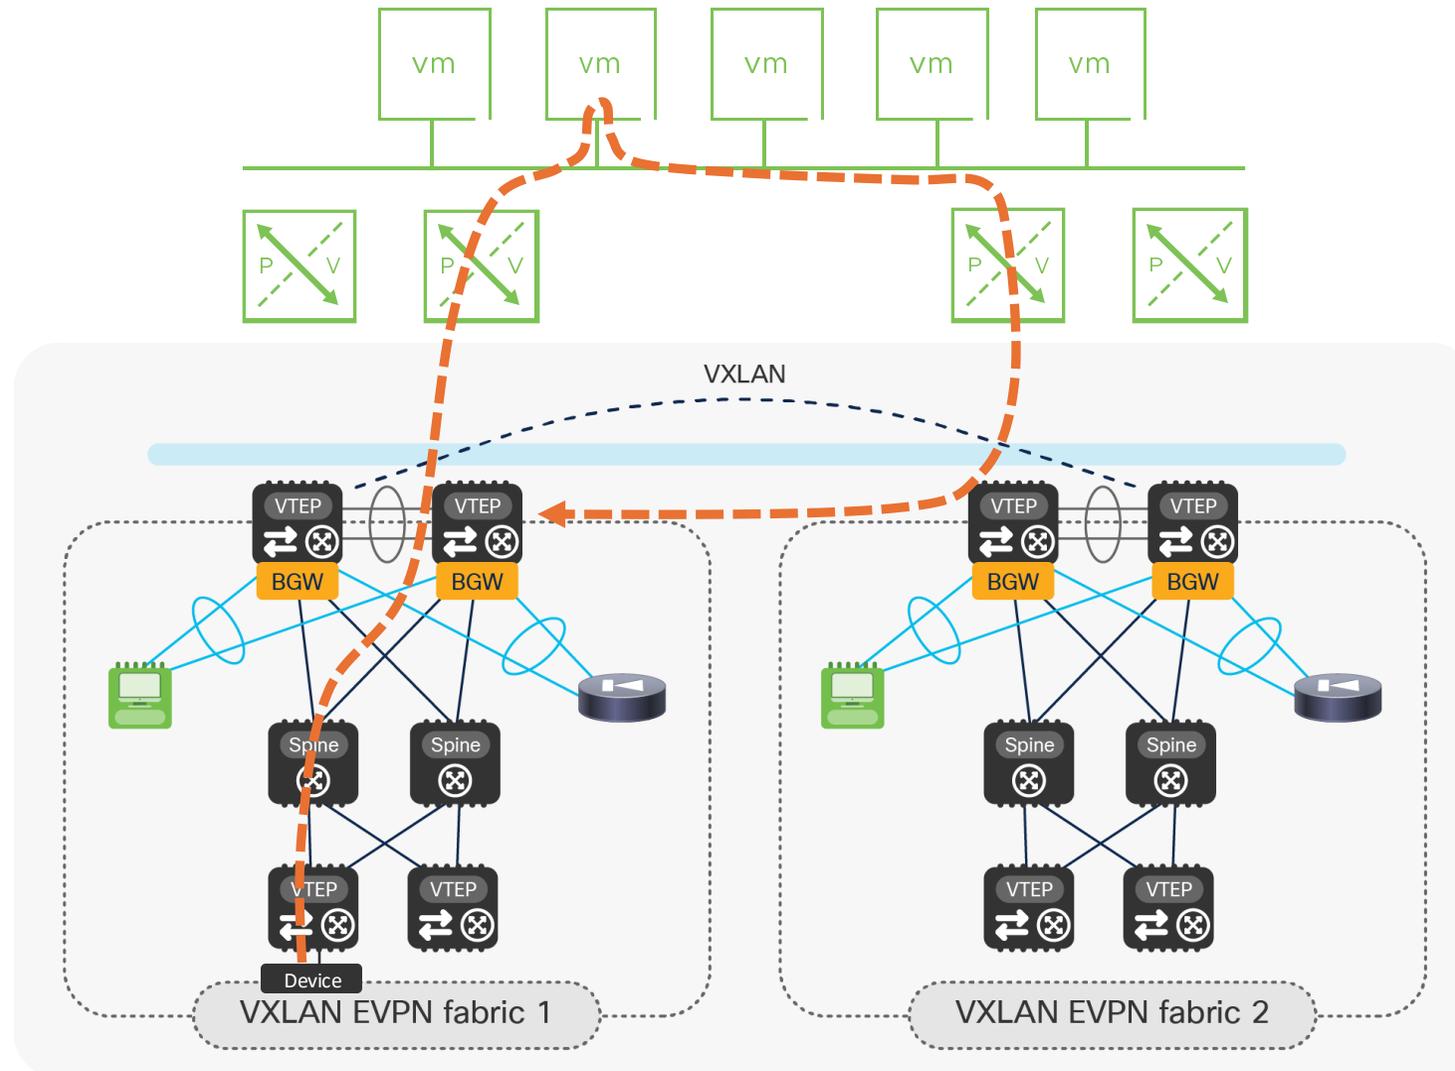
```
-esxi-01.od.intern - PuTTY
tcpdump-uw: pcap_loop: error reading dump file: Interrupted system call
pktcap: Receive thread exiting...
pktcap: error: Error writing 28 bytes of pkt header to file.
pktcap: Destroying session 1.
pktcap:
pktcap: Dumped 298 packet dropped 0 packets.
pktcap: Done.
[root@esxi-01:~] pktcap-uw --switchport 167772201 --capture VnicTx,VnicRx -o - | tcpdump-uw -r- -nne | grep .7.10
The switch port id is 0x0a000029.
The session capture point is VnicTx,VnicRx.
pktcap: The output file is -.
pktcap: No server port specified, select 1132 as the port.
pktcap: Local CID 2.
pktcap: Listen on port 1132.
pktcap: Main thread: 464756669632.
pktcap: Dump Thread: 464757200640.
pktcap: The output file format is pcapng.
pktcap: Recv Thread: 464757729024.
pktcap: Accept...
pktcap: Vsock connection from port 1028 cid 2.
reading from file -, link-type EN10MB (Ethernet), snapshot length 65535
02:50:56:00:70:00 > bc:8d:1f:07:12:1b, ethertype IPv4 (0x0800), length 74: .7.10 > .3.199: ICMP echo reply, id 10, seq 4613, length 40
02:50:56:00:70:00 > bc:8d:1f:07:12:1b, ethertype IPv4 (0x0800), length 74: .7.10 > .3.199: ICMP echo reply, id 10, seq 4614, length 40
02:50:56:00:70:00 > bc:8d:1f:07:12:1b, ethertype IPv4 (0x0800), length 74: .7.10 > .3.199: ICMP echo reply, id 10, seq 4615, length 40
tcpdump-uw: pcap_loop: error reading dump file: Interrupted system call
pktcap: Receive thread exiting...
pktcap: error: Error writing 28 bytes of pkt header to file.
pktcap: Destroying session 2.
pktcap:
pktcap: Dumped 24 packet dropped 0 packets.
pktcap: Done.
[root@esxi-01:~]
```

<code>--vlan VLAN_ID</code>	Capture or trace packets that belong to a VLAN.
<code>--srcip IP_address IP_address/subnet_range</code>	Capture or trace packets that have a specific source IPv4 address or subnet.
<code>--dstip IP_address IP_address/subnet_range</code>	Capture or trace packets that have a specific destination IPv4 address or subnet.
<code>--ip IP_address</code>	Capture or trace packets that have a specific source or destination IPv4 address.
<code>--proto 0xIP_protocol_number</code>	Capture or trace packets at Layer 3 according to the next level protocol that consumes the payload. For example, to monitor traffic for the UDP protocol, type <code>--proto 0x11</code> .
<code>--srcport source_port</code>	Capture or trace packets according to their source TCP port.

# Troubleshooting op de Cisco Switches

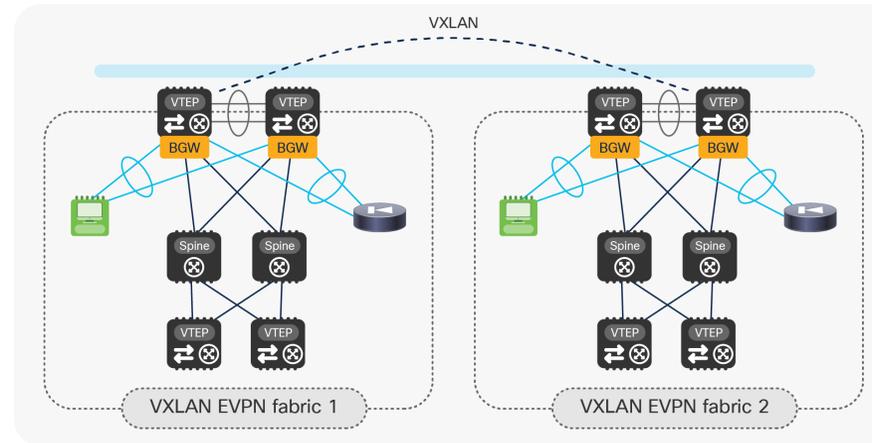


# Troubleshooting op de Cisco Switches

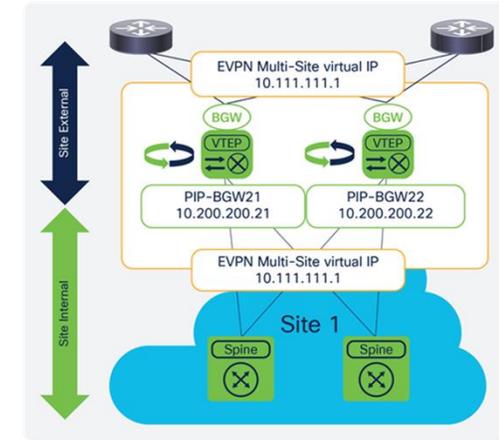


# Vermoeden en workaround

- vPC/PIP/VIP



- Workaround is device in ander VRF





# Vragen?