



F5XCの分散コントローラと データプレーンアーキテクチャについて SDNを利用したNaaSサービスの展開

Solution Architect

Daisuke Nakajima

Who am I ?



名前：中嶋 大輔

所属：F5ネットワークスジャパン

タイトル：ソリューションアーキテクト

過去の発表

- JANOG50: 集まれSRv6の森
- JANOG48: NaaS – Cloud Native時代のネットワークのあり方
- JANOG47: エッジコンピューティング時代のサービス運用の課題
- JANOG41: NFV+SDN標準化/実装動向と運用の課題

SDNとの関わり



<https://www.slideshare.net/slideshow/openvswitch-vps-20120429/12733200>



<https://www.slideshare.net/slideshow/open-contrail/30065485>

12年前！！！！

SDNとの関わり



<https://www.slideshare.net/slideshow/cloud-operator-days-tokyo-2020-237392286/237392286>




F5 Distributed Cloud Services とは？



Mesh – Networking & Security Services



Router



Firewall



ADC



DDOS



WAF



Signature/Signal
Bot Protection



AppStack – Application Platform Services



K8s Compute
Platform



K8s Cluster
Management



Service
Discovery



Secrets
Management




Identity



Console – Cloud-based Central Portal



Centralised
Operations



Visibility and
Analytics



Artificial Intelligence
Advanced Insights



ADN – Global app-to-app network



26 Global POPs



Public Clouds
Private peering



Multi-Tbps
Capacity



CONSOLIDATION



OPERATIONAL SIMPLICITY

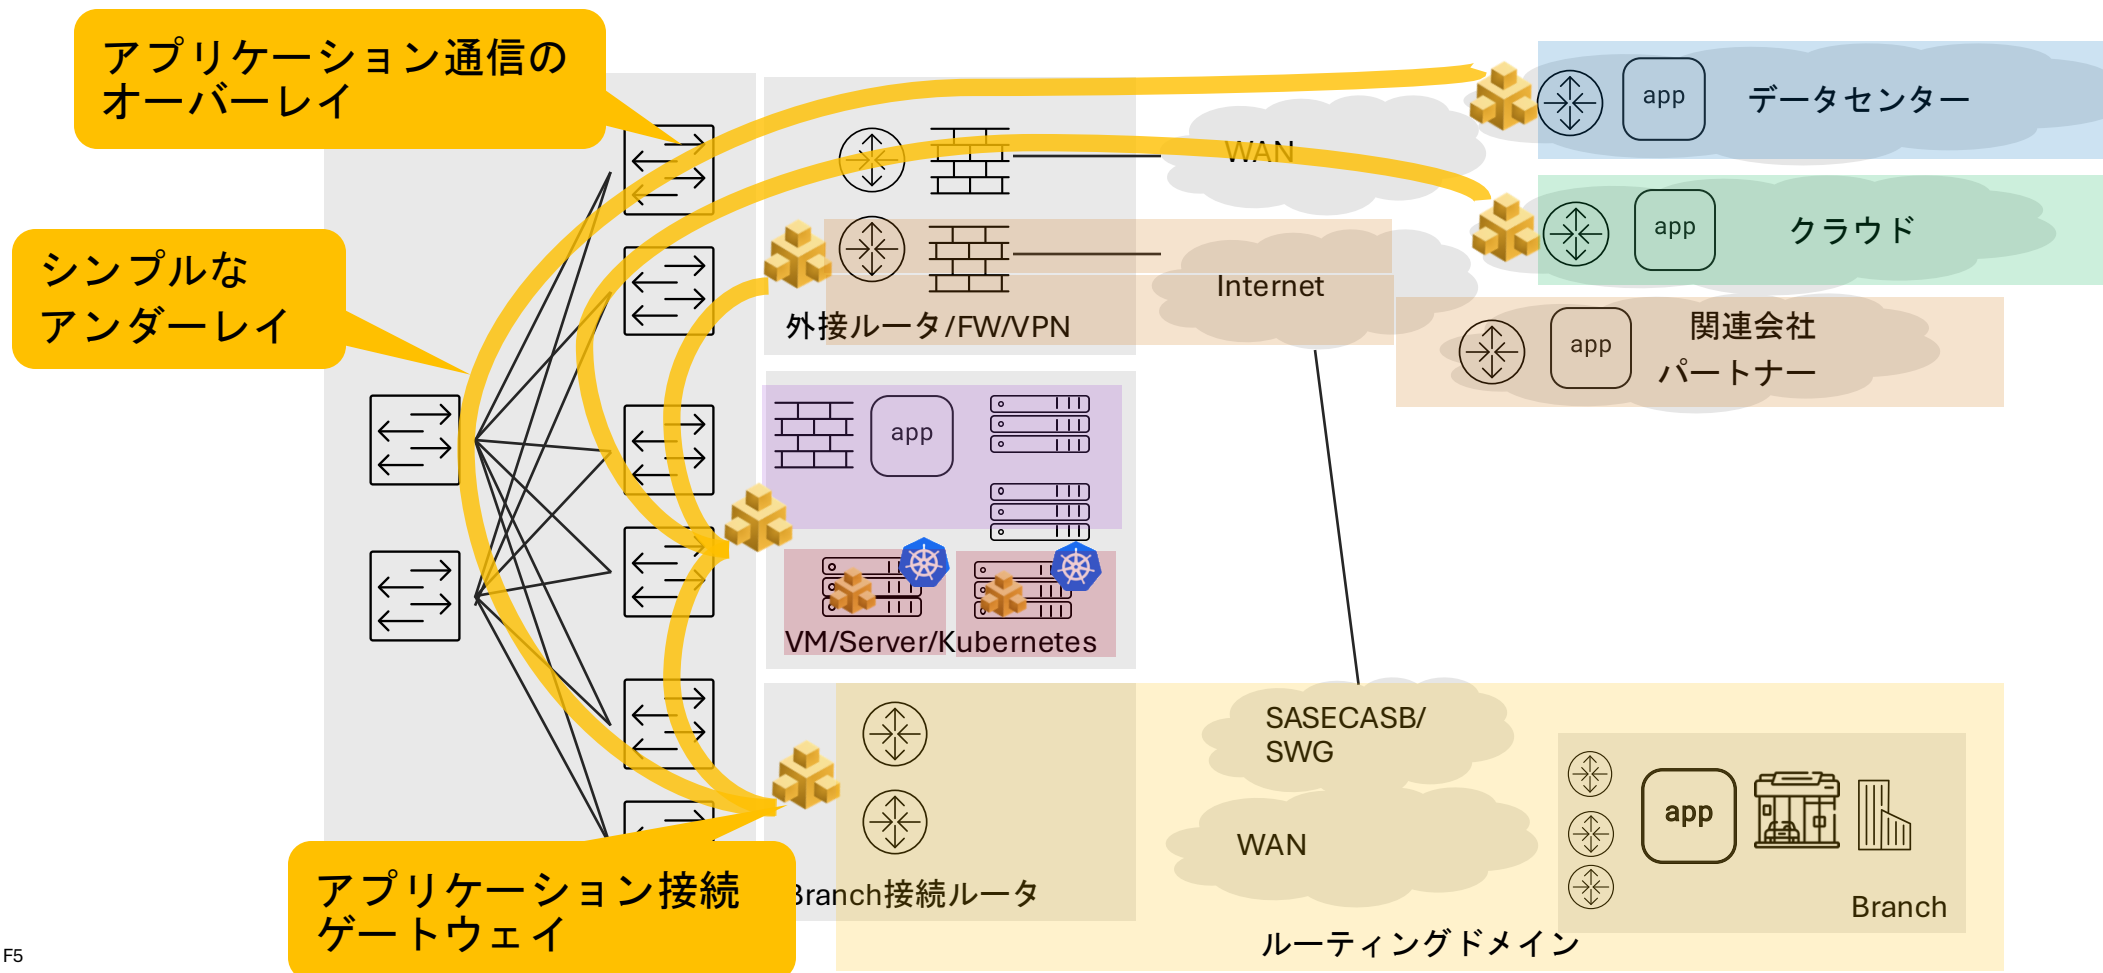


COMMON TOOLING

エンタープライズネットワーク再考

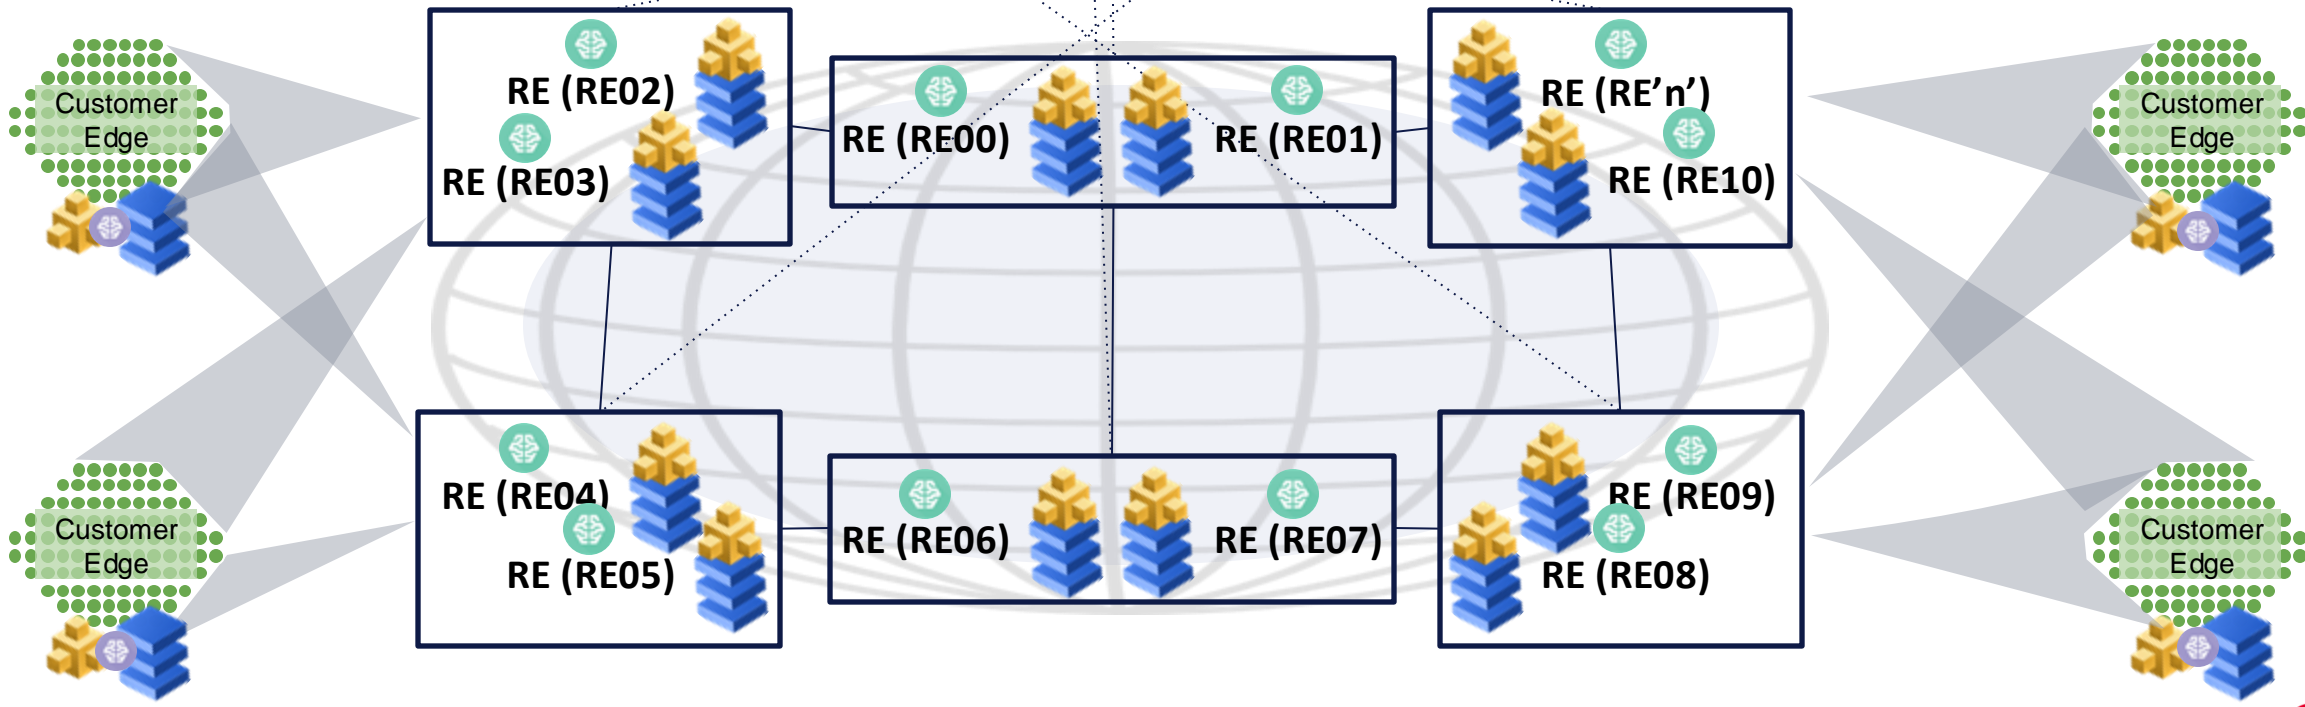
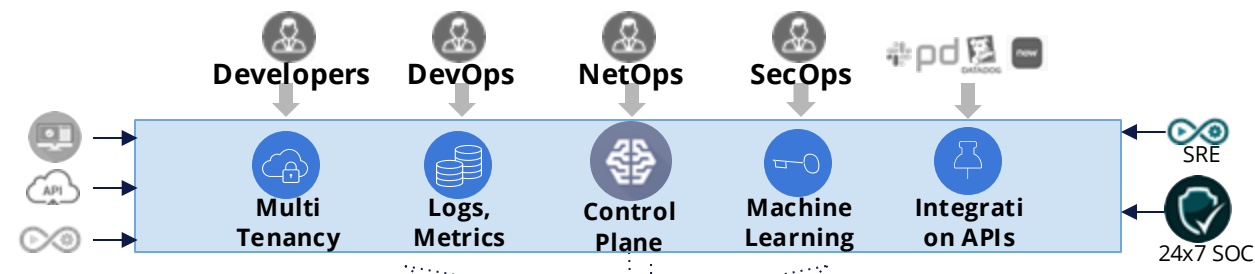
ルーティングドメインの分離と抽象化レイヤー

物理ネットワークのコンフィグをシンプル化し、L3レベルでのネットワーク相互接続を分離
アプリケーション接続をオーバーレイで行う



F5 Distributed Cloud Servicesコントロールプレーン

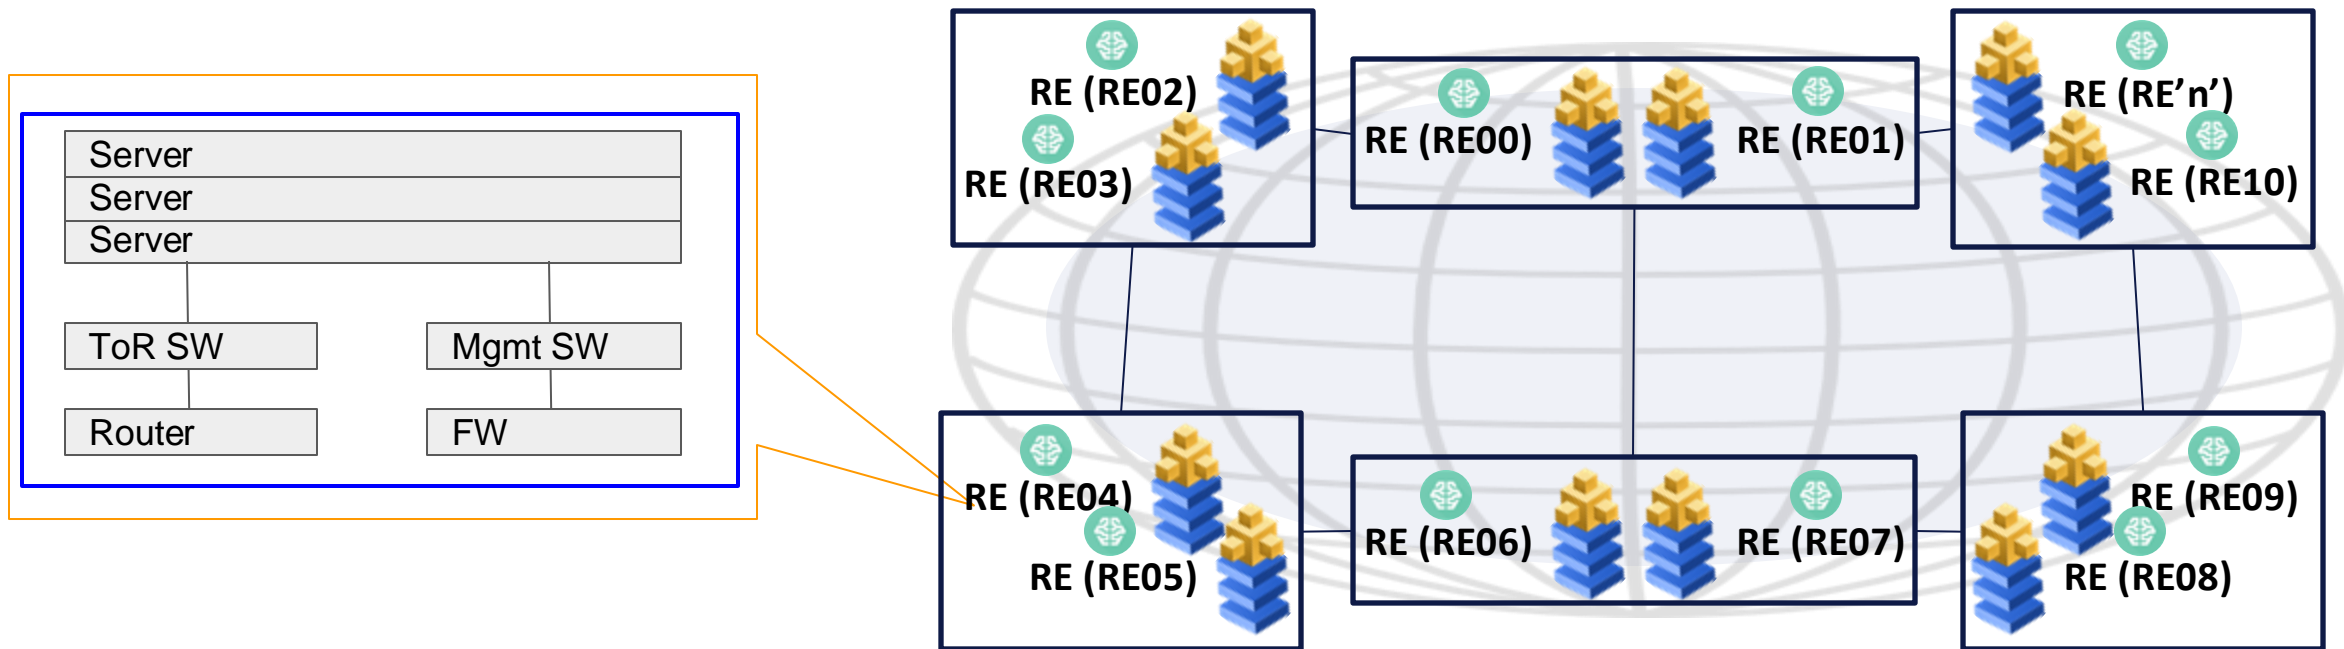
-  Global Control Plane
-  Regional Control Plane
-  Local Control Plane



Regional Edge (RE)

Regional Edgeはルータ/スイッチ/サーバで構成されており、スケールアウトが可能。
リモートよりOSインストールやコンフィグ変更を行うシステムを完備。現地保守はハードウェア交換やケーブリングのみで、OSのインストールやソフトウェアのアップグレードは全てリモートで完了。

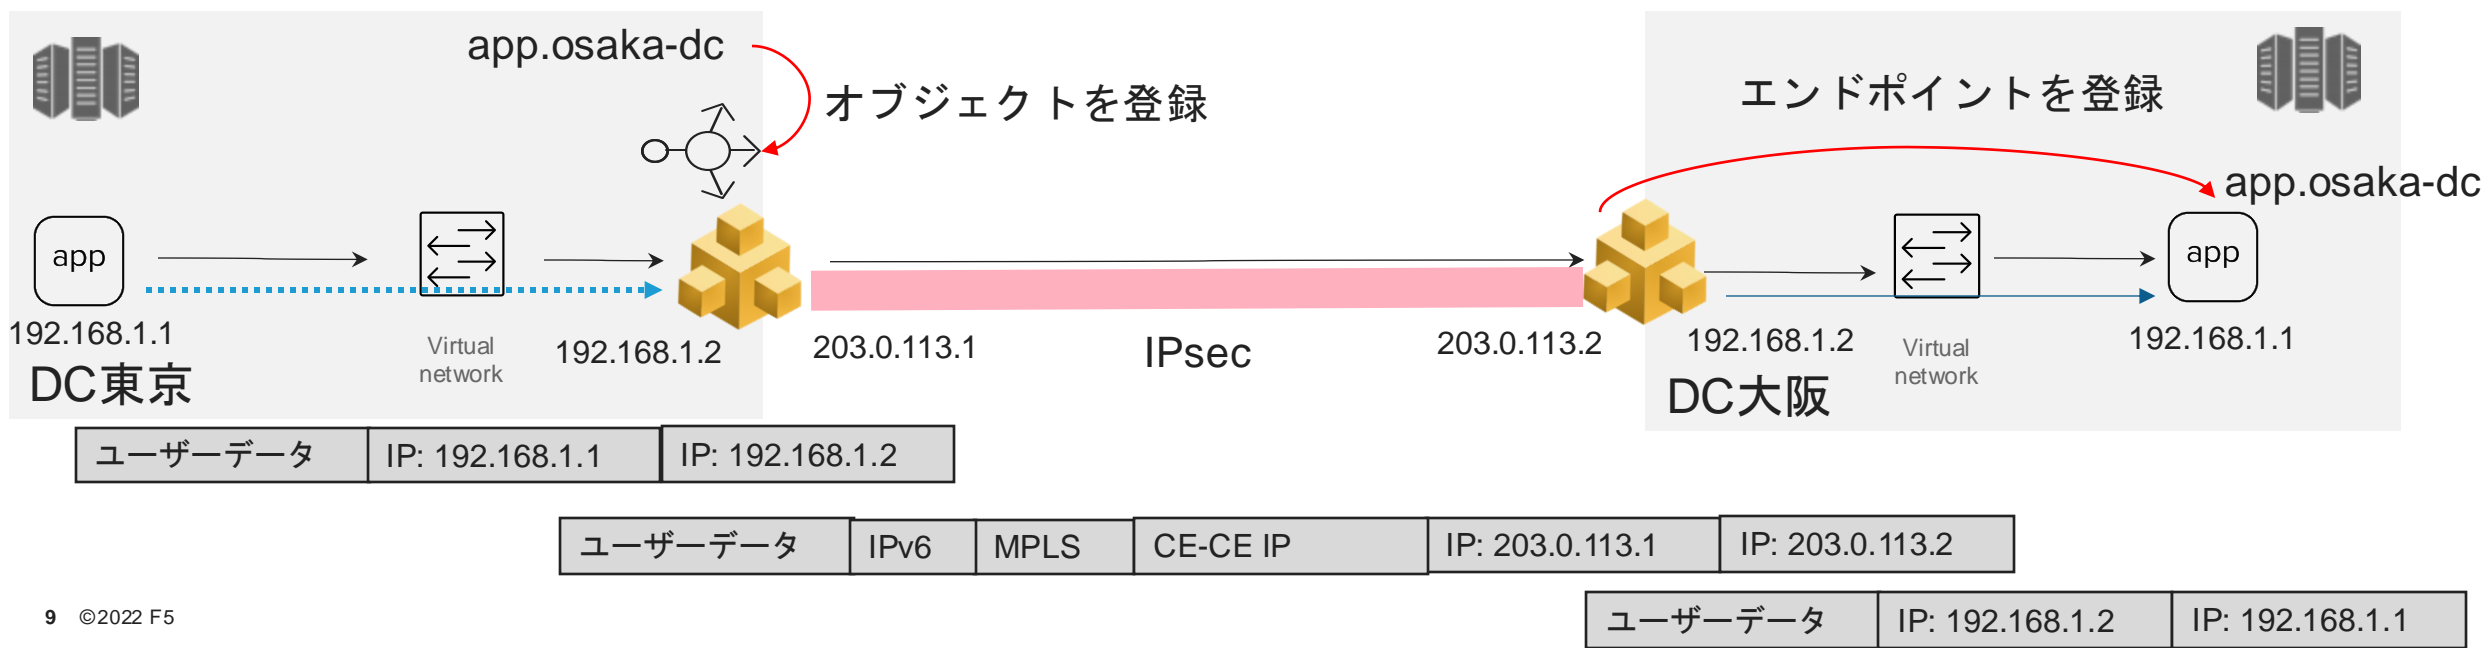
ルーターはバックボーンなどとBGPで接続。VRFで専用線やユーザーLANを収容することも可能

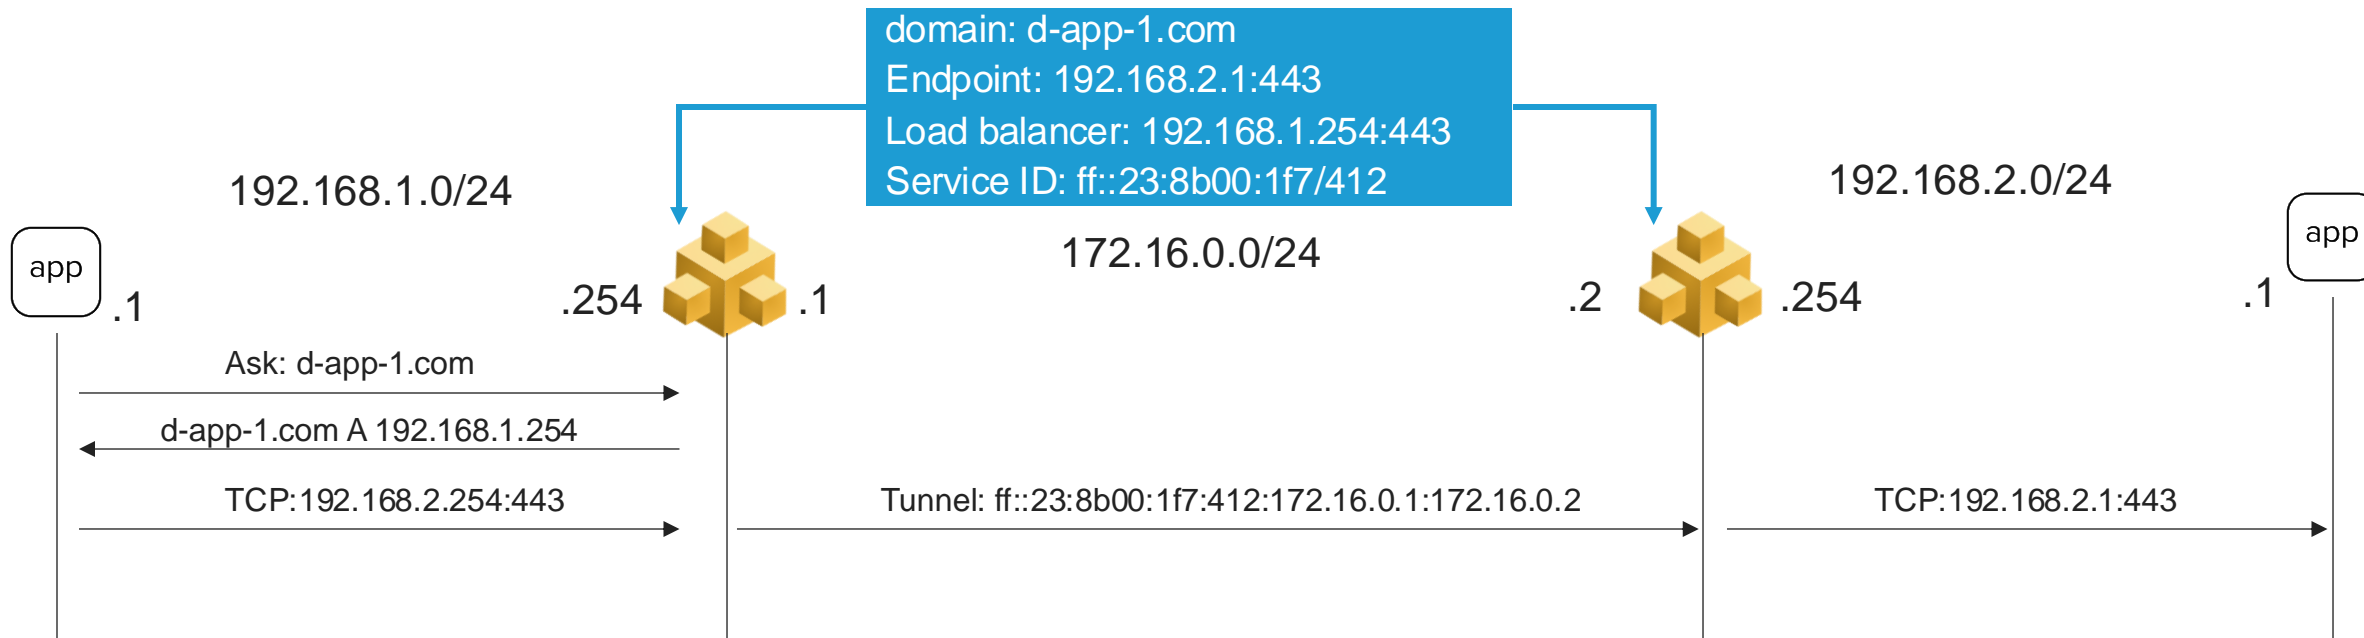


F5XCのネットワーク抽象化

F5XCでは拠点間ネットワークをL3 VPN (IPsec トンネル) も可能ですが、ネットワークを抽象化し、アプリケーションベースの制御を行うことでIPアドレス重複の排除も可能です。

例えば、大阪の app.osaka-dcアプリを東京のアプリケーションとの関係を行う場合、東京のF5XCノードにapp.osaka-dcへのアクセスを許可するオブジェクト(LB)を登録します。そのオブジェクトに対し、エンドポイントに実際のapp.osaka-dcを設定します。東京のアプリはDC東京のF5XCノードで一度終端されます。終端されたトラフィックはF5XCノード間のIPsecトンネルを通り、送信元IPをDC大阪のF5XCノードをセットしてアプリに通信します。



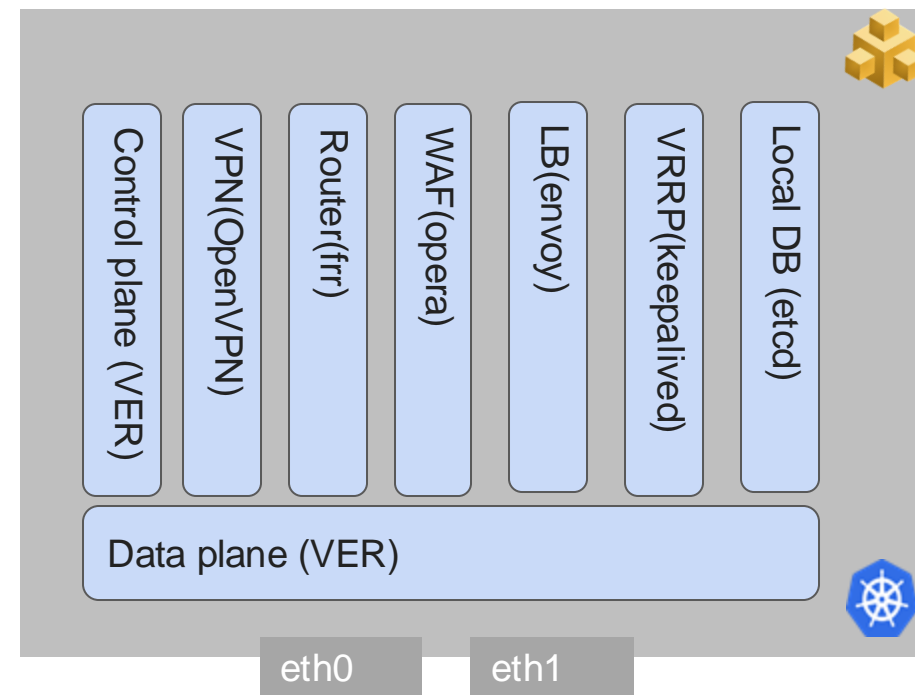


Load balancer オブジェクトが作成されると、エンドポイントに対して識別用のService IDが付与されます。

1. 作成されたロードバランサーのドメイン名をクライアントは名前解決します
2. CEはAレコードに自身のIPもしくはVIPをAレコードに設定して返信します。
3. クライアントはCEに対して接続します。
4. CEは受信したポートとSNI/HostnameなどからLoad balancerを特定し、エンドポイントのService IDをセットしてCEにデータを送信します。
5. 対抗のCEはService IDからエンドポイントを特定し、トラフィックを送信します。

Customer Edge (CE) ノードコンポーネント

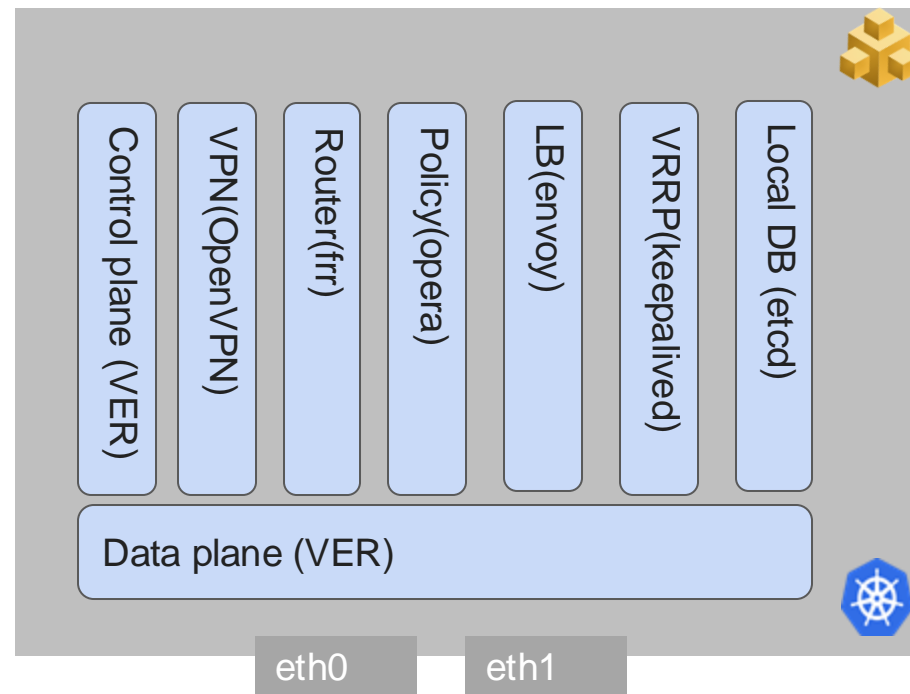
```
node-0 dnakajima-vlan-gw-ce-1 ~ kubectl get po
NAME                                READY   STATUS    RESTARTS   AGE
argo-rkfxw                          2/2    Running   0           20d
asaf-w5jcz                          2/2    Running   0           20d
aspen-6ljlc                         2/2    Running   0           20d
bdbewaf-7lrpc                       2/2    Running   0           20d
dnsmasq-5f698f8dff-m99kp            2/2    Running   0           20d
envoy-fpmnb                          2/2    Running   0           20d
etcd-0                               2/2    Running   0           20d
etcd-defrag-28841410-f9ppz          0/2    Completed 0           5d10h
frr-f97p8                           3/3    Running   0           20d
ganges-65ffb65877-gnwcf            2/2    Running   0           20d
gubernator-654f5786c9-7swqv        2/2    Running   0           20d
hellas-7b9f48bf47-z6d24           2/2    Running   0           20d
ike-dndbk                           2/2    Running   0           20d
keepalived-ktgb2                   2/2    Running   0           20d
obelix-5vfzm                        2/2    Running   0           20d
openvpn-5sgrx                       3/3    Running   0           20d
opera-4dr7r                         2/2    Running   0           20d
phobos1-d7cc45f-kt2qj              2/2    Running   0           20d
piku-5bb66b955d-5dhkd              2/2    Running   0           11h
pmtud-bzww6                        2/2    Running   0           20d
ver-6bd7975cdd-1cz9b               2/2    Running   0           20d
ver-jptwz                           2/2    Running   0           20d
webroot-9gkc4                      2/2    Running   0           20d
```




Customer Edge (CE) ノードコンポーネント

```
vif0/4 Ethernet: vethbfd0 Parent:vif0/-1
Type:Virtual HWaddr:0:0:5e:0:1:0 Driver: net_af_packet
Vrf:2 Mcast Vrf:2 VN:0x0 Flags:PL3ProxyTTL Ref:14
MTU: 1500 EncMSS: 130 NH: 13
RX port packets:1964185 errors:0
RX queue errors to lcore 0 0 0 0 0 0 0 0 0 0 0 0 0
RX packets:1964185 bytes:127759282 errors:0
TX packets:1962700 bytes:127658180 errors:0
Drops:753
TX port packets:1962700 errors:0
Ingress ACL: 65 67
Egress ACL: 65 67
Service ACL: 65 67

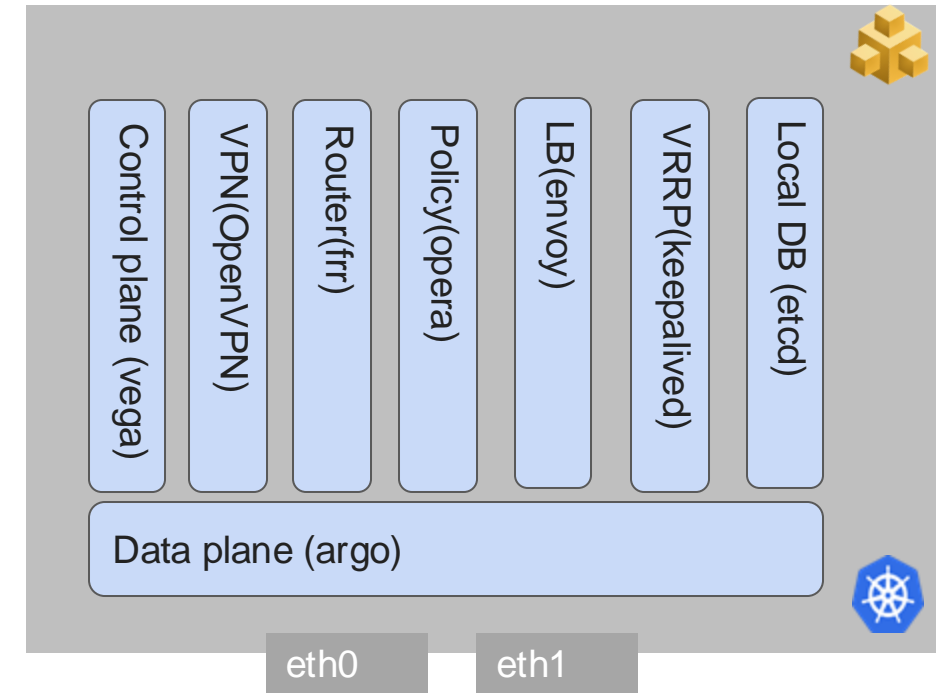
vif0/6 Ethernet: argo-envoy-0 Parent:vif0/-1
Type:Virtual HWaddr:0:0:5e:0:1:0 Driver: net_af_packet
Vrf:3 Mcast Vrf:3 VN:0x0 Flags:PL3MqProxy Ref:11
MTU: 1520 EncMSS: 130 NH: 14
VN offset:11
SRC VN Classifier Prefixes: SRC fe80:1::/48 fe80:3::/48 fe80:4::/48
RX port packets:749 errors:0
RX queue errors to lcore 0 0 0 0 0 0 0 0 0 0 0 0 0
RX packets:749 bytes:52726 errors:0
TX packets:1 bytes:86 errors:0
Drops:748
TX port packets:1 errors:0
Ingress ACL: 65 67
Egress ACL: 65 67
Service ACL: 65 67
```



Forwarding Planeの構造



Index	Source:Port/Destination:Port	Proto(V)
19524<=>229888	192.168.3.100:80 192.168.3.1:51219	6 (17->15)
(Gen: 17, K(Idx):4113, Action:N(SPsDPd), Flags:, TCP:SSrEEr, Stats:355/90737 E0:0, QOS:-1, S(nh):524, D(nh):525, D(L):278, SPort: 56013, TTL: 0, CPU: 12, Sinfo: 17.16.0.0, Flow label: 0 Created: 3089255928, Modified : 3089255928, LastSeen: 3089268274, Changed: 0, ACL: 65/1 Samples(Packets/Bytes):255/65037)		
149688<=>228780	192.168.3.100:80 192.168.3.1:51221	6 (17->15)
(Gen: 34, K(Idx):4113, Action:N(SPsDPd), Flags:, TCP:SSrEEr, Stats:265/67607 E0:0, QOS:-1, S(nh):524, D(nh):525, D(L):278, SPort: 57235, TTL: 0, CPU: 12, Sinfo: 17.16.0.0, Flow label: 0 Created: 3089256024, Modified : 3089256024, LastSeen: 3089268478, Changed: 0, ACL: 65/1)		



L3 VPN通信



L3 Tunnel (CE1)

```
root@node-0:~# vrfstats --dump
Vrf: 1
Discards 0, Resolves 0, Receives 0, L2 Receives 0, Vrf Translates 0, Unknown Unicast Floods 0
Ecmp Composites 0, L2 Mcast Composites 0, Fabric Composites 0, Encap Composites 0, Evpn Composites 0
Udp Tunnels 0, Udp Mpls Tunnels 0, Gre Mpls Tunnels 0, Vxlan Tunnels 0, Pbb Tunnels 0, IP Tunnels 0
L2 Encaps 0, Encaps 0, NAT 0
GROs 0, GRO merged 0, Diags 0
Arp Virtual Proxys 2, Arp Virtual Stitches 0, Arp Virtual Floods 0, Arp Physical Stitches 0, Arp Tor Proxys 0,
Arp Physical Floods 0
SRV6 tunnels 0

Vrf: 2
Discards 835, Resolves 0, Receives 0, L2 Receives 0, Vrf Translates 0, Unknown Unicast Floods 0
Ecmp Composites 0, L2 Mcast Composites 0, Fabric Composites 0, Encap Composites 0, Evpn Composites 0
Udp Tunnels 0, Udp Mpls Tunnels 0, Gre Mpls Tunnels 0, Vxlan Tunnels 0, Pbb Tunnels 0, IP Tunnels 0
L2 Encaps 0, Encaps 2105042, NAT 0
GROs 0, GRO merged 0, Diags 0
Arp Virtual Proxys 87391, Arp Virtual Stitches 0, Arp Virtual Floods 0, Arp Physical Stitches 0, Arp Tor Proxys
0, Arp Physical Floods 0
SRV6 tunnels 0

Vrf: 3
Discards 831, Resolves 0, Receives 0, L2 Receives 0, Vrf Translates 0, Unknown Unicast Floods 0
Ecmp Composites 0, L2 Mcast Composites 0, Fabric Composites 0, Encap Composites 0, Evpn Composites 0
Udp Tunnels 0, Udp Mpls Tunnels 0, Gre Mpls Tunnels 0, Vxlan Tunnels 0, Pbb Tunnels 0, IP Tunnels 0
L2 Encaps 0, Encaps 0, NAT 0
GROs 0, GRO merged 0, Diags 0
Arp Virtual Proxys 17, Arp Virtual Stitches 0, Arp Virtual Floods 0, Arp Physical Stitches 0, Arp Tor Proxys 0,
Arp Physical Floods 0
SRV6 tunnels 0
```

L3 Tunnel (CE1)

```
root@node-0:/# rt --dump 18
Flags: L=Label Valid, P=Proxy ARP, T=Trap ARP, F=Flood ARP, H=Proxy ARP for host interface, V=Preferred VRF
Labels/Misc: S=SiteId N=NetworkId
vRouter inet4 routing table 0/18/unicast
Destination          PPL      Flags      Label      Nextthop  Labels/Misc
                    Number of routes: 19

0.0.0.0/0            0         -          -          0
10.0.10.100/32      32        LV         270        489        (S,N):4972:48217 VN:65543 Stats UID:10
                    1:45 35:155
10.10.10.10/32     32        LV         270        584        (S,N):5049:48217 VN:65543 Stats UID:10
                    1:45 35:155
192.168.1.0/24     24        LV         280        489        (S,N):4972:48217 VN:65543 Stats UID:10
                    1:45 6:21 35:155
192.168.2.0/24     24        LV         280        529        (S,N):5446:48218 VN:65543 Stats UID:14
                    1:45 6:21 35:155
192.168.3.0/24     24        LV         279        489        (S,N):4972:48236 VN:65543
                    1:45 6:21 35:155
192.168.4.0/24     24        LV         279        529        (S,N):5446:48217 VN:65543 Stats UID:10
                    1:45 6:21 35:155
192.168.100.0/24   24        LV         279        501        (S,N):5049:48217 VN:65543 Stats UID:10
                    1:45 6:21 35:155
192.168.100.1/32  32        LPV        0          504        (S,N):5049:48217 VN:65543 Stats UID:10
                    1:45 6:6 6:21
                    35:155
```


L3 Tunnel (CE1)

```
root@node-0:/# nh --get 489
Id:489      Type:Tunnel      Fmly: AF_INET  Rid:0  Vrf:8
            Flags:Valid, MPLSoUDP, Indirect, Tunnel: None
            Direct NH(label): 532
            Oif:0  Len:14 Data:00 00 00 00 00 00 00 00 00 00 00 00 00 00
            Sip:10.101.10.39 Dip:10.101.9.245

Id:532      Type:Composite     Fmly: AF_INET  Rid:0  Vrf:0
            Flags:Valid, Ecmp, Tunnel: None
            Composite Flags:
            Sub NH(label): 85(0)(W:100) 86(0)(W:100)

Id:85       Type:Encap        Fmly: AF_INET  Rid:0  Vrf:8
            Flags:Valid, Tunnel: None
            Oif:4111 Len:14
            Encap Data: 00 00 5e 00 01 00 00 00 5e 00 01 00 08 00

Id:86       Type:Encap        Fmly: AF_INET  Rid:0  Vrf:8
            Flags:Valid, Tunnel: None
            Oif:4109 Len:14
            Encap Data: 00 00 5e 00 01 00 00 00 5e 00 01 00 08 00
```

L3 Tunnel (CE1)

```
root@node-0:/# vif --get 4111
Vrouter Interface Table
```

```
Flags: P=Policy, X=Cross Connect, S=Service Chain, Mr=Receive Mirror
Mt=Transmit Mirror, Tc=Transmit Checksum Offload, L3=Layer 3, L2=Layer 2
D=DHCP, Vp=Vhost Physical, Pr=Promiscuous, Vnt=Native Vlan Tagged
Mnp=No MAC Proxy, Dpdk=DPDK PMD Interface, Rfl=Receive Filtering Offload, Mon=Interface is Monitored
Nf=Dont Fragment, Vof=VLAN insert/strip offload, Mq=MultiQ, L=MAC Learning Enabled
Proxy=MAC Requests Proxied Always, Fr=Force Recv, Mn=Mirror without Vlan Tag, TTL=TTL not modified
Mn=Mirror without Vlan Tag, Tun=Tunnel, NoArp=No ARP interface, RawIp=RAW IP interface(No L2 header)
Npers=Non-Persistent, NoRPF=No RPF Checks, SynC=Syn Cookies, Gro=GRO Enabled
```

```
vif0/4111 Virtual: tunnel_ty8-tky_0 Parent:vif0/-1
Type:Tunnel HWaddr:0:1:2:3:4:5 Driver:
UserVisible
Vrf:8 Mcast Vrf:8 VN:0x0 Flags:L3Tun Ref:67
MTU: 1370 EncMSS: 0 NH: 85 FNH: 31
Outbound SA:SAId: 0xffff0060 Tunnel SRC:172.28.13.36 Tunnel DST:103.135.56.40
Tunnel Mode:Tunnel Window size:0 VIF: 0 State: Active
Inbound SA:SAId: 0xffff5f70 Tunnel SRC:103.135.56.40 Tunnel DST:172.28.13.36
Tunnel Mode:Tunnel Window size:4 VIF: 4111 State: Active
NAT Pool: 10.101.10.39
RX queue errors to lcore 0 0 0 0 0 0 0 0 0 0 0 0
RX packets:86228068 bytes:54213735085 errors:0
TX packets:83300681 bytes:44061827864 errors:18
Drops:39093
TCP reset ignored:36
Service ACL: 65 67
```

L3 Tunnel (CE2)

```
node-0 dnakajima-vlan-gw-ce-1 ~ shell argo
root@node-0:/# mpls --get 280
MPLS Input Label Map
```

Label	NextHop

280	513

```
root@node-0:/# nh --get 513
Id:513      Type:Vrf_Translate  Fmly: AF_INET  Rid:0  Vrf:18
           Flags:Valid, Policy, Tunnel: None
```

L3 Tunnel (CE2)

```
root@node-0:/# rt --dump 18
Flags: L=Label Valid, P=Proxy ARP, T=Trap ARP, F=Flood ARP, H=Proxy ARP for host interface, V=Preferred VRF
Labels/Misc: S=SiteId N=NetworkId
vRouter inet4 routing table 0/18/unicast
Destination      PPL      Flags      Label      Nexthop  Labels/Misc
                Number of routes: 13

0.0.0.0/0        0        -          -          0
10.0.10.100/32   32       LV         270        579      (S,N):4972:48217 VN:65543
                1:32 21:54
10.10.10.10/32   32       LV         270        525      (S,N):5049:48217 VN:65543
                1:32 21:54
192.168.1.0/24   24       LV         280        572      (S,N):4972:48217 VN:65543
                1:32 6:175
                21:54
192.168.1.1/32   32       LPV        0          575      (S,N):4972:48217 VN:65543
                1:32 6:6 6:175
                21:54
192.168.1.100/32 32       LV         0          577      (S,N):4972:48217 VN:65543
                1:32 6:175
                21:54
192.168.3.0/24   24       LV         279        499      (S,N):4972:48236 VN:65543 Stats UID:9
                1:32 6:175
                21:54
192.168.3.1/32   32       LPV        0          521      (S,N):4972:48236 VN:65543 Stats UID:9
                1:32 6:6 6:175
                21:54
```

L3 Tunnel (CE2)

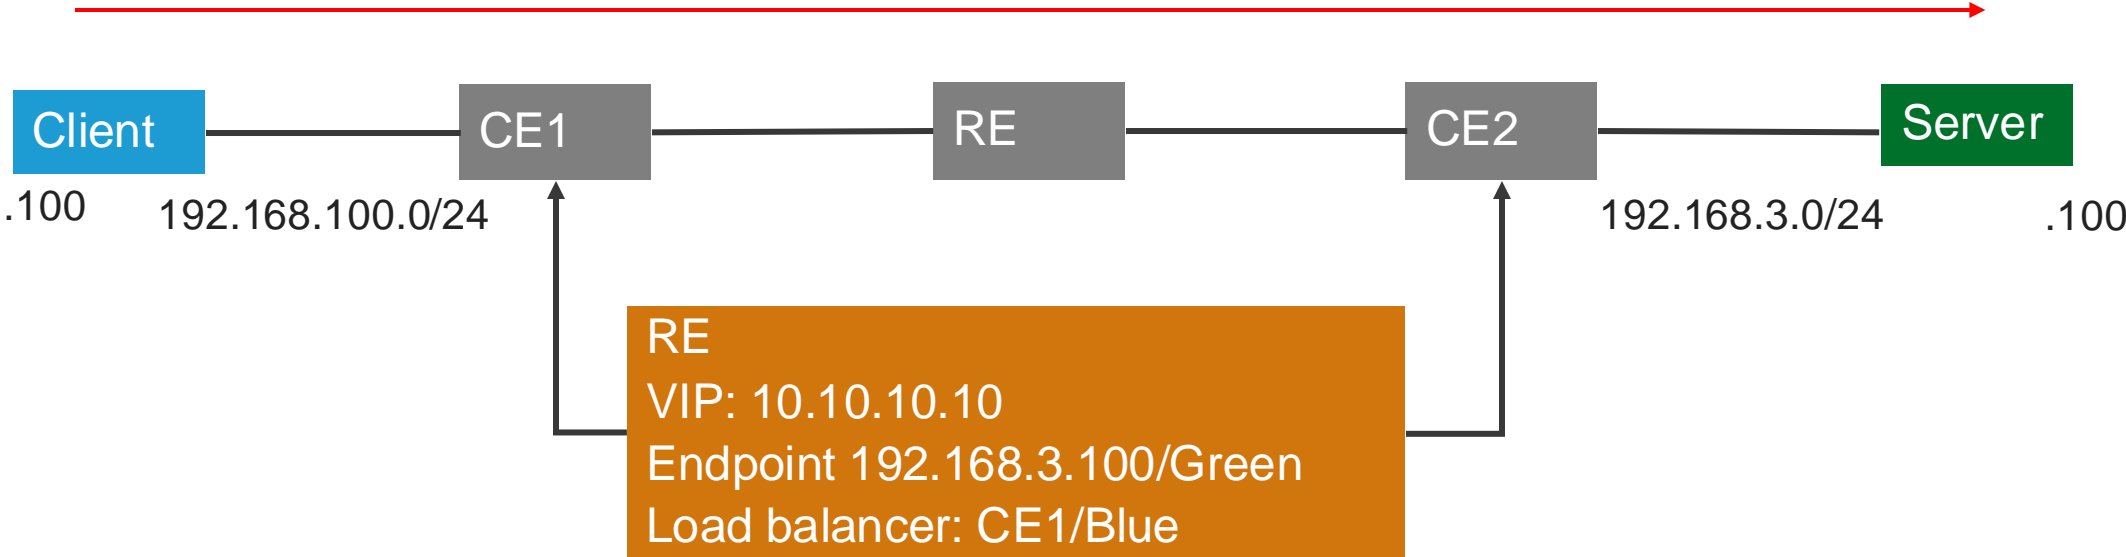
```
root@node-0:/# nh --get 577
Id:577          Type:Encap          Fmly: AF_INET  Rid:0  Vrf:18
                Flags:Valid, Tunnel: None
                Oif:4121  Len:18
                Encap Data: 00 0c 29 9b 69 69 00 0c 29 0f f4 a1 81 00 01 2c 08 00

root@node-0:/# vif --get 4121
Vrouter Interface Table

Flags: P=Policy, X=Cross Connect, S=Service Chain, Mr=Receive Mirror
Mt=Transmit Mirror, Tc=Transmit Checksum Offload, L3=Layer 3, L2=Layer 2
D=DHCP, Vp=Vhost Physical, Pr=Promiscuous, Vnt=Native Vlan Tagged
Mnp=No MAC Proxy, Dpdk=DPDK PMD Interface, Rfl=Receive Filtering Offload, Mon=Interface is Monitored
Nf=Dont Fragment, Vof=VLAN insert/strip offload, Mq=MultiQ, L=MAC Learning Enabled
Proxy=MAC Requests Proxied Always, Fr=Force Recv, Mn=Mirror without Vlan Tag, TTL=TTL not modified
Mn=Mirror without Vlan Tag, Tun=Tunnel, NoArp=No ARP interface, RawIp=RAW IP interface(No L2 header)
Npers=Non-Persistent, NoRPF=No RPF Checks, SynC=Syn Cookies, Gro=GRO Enabled

vif0/4121  Ethernet: ens224.300 Vlan(o/i)(,S): 300/300 (Speed 0, Duplex 0) Parent:vif0/24
Type:Physical(Vlan) HWaddr:0:c:29:f:f4:a1  Driver: vr_vlan
CrossConnectIdx: 28 UserVisible
Address: 192.168.1.1
Vrf:18 Mcast Vrf:18 VN:0x0 Flags:PL3 Ref:11
MTU: 1500 EncMSS: 130 NH: 571
NAT Pool: 192.168.1.1 192.168.1.1
RX queue errors to lcore 0 0 0 0 0 0 0 0 0 0 0 0 0
RX packets:597631 bytes:47659274 errors:0
TX packets:587832 bytes:49253887 errors:0
Drops:15
Ingress ACL: 65 67
Egress ACL: 65 67
Service ACL: 65 67
```

Load balancer



Load balancer

Metadata ⓘ


Name: ⓘ segment3-nginx

Origin Servers ^


Origin Servers ⓘ

1 item

Type	Name/IP	Site/VSite	Network/Segment	Labels
Private IP	192.168.3.100	dnakajima-vlan-gw-ce-1	demo-green	

 Origin server Port: ⓘ Port ⓘ

Port: ⓘ 80

 Port used for health check: ⓘ Endpoint port ⓘ

LoadBalancer Algorithm: ⓘ Load Balancer Override

Endpoint Selection: ⓘ Local Endpoints Preferred



Load balancer

View

Child objects
segment3-nginx

View x

Status
ves-io-origin-pool-segment3-nginx-757ccdb999

- Status Objects
- View Item
 - Metadata
 - Conditions
 - Ver status
 - Origin Span

Status ID: ⓘ 1620fcde-48cc-4734-975e-996675fd851d_VerCfgMgr

Publish: ⓘ Publish

VTRP ID: ⓘ mars2.gc01-cle.int.ves.io

VTRP Stale: ⓘ False

Conditions ⓘ

2 items

Type	Status	Reason	Last Updated
Validation	Success		2024-11-19 09:29:22
Operational	Installed		2024-11-19 09:29:22

Ver status ⓘ

1 item

Site	Health Status	Discovered IP Address	Allocated IP Address	Discovered port
dnakajima-vlan-gw-...		<u>192.168.3.100</u>	<u>ff::13:6c00:2</u>	80



Load balancer

Form Documentation JSON Search Edit Configuration

View

HTTP Load Balancer
nutanix-lb

Metadata

- Domains and LB Type
- Origins
- Web Application Firewall
- Bot Protection
- API Protection
- DoS Protection
- Client-Side Defense
- Common Security Controls
- Other Settings
- Virtual Host State
- Auto Cert State
- Cert State

1 item

nutanix-lb.test

- Load Balancer Type: HTTP
- Automatically Manage DNS Records: False
- HTTP Listen Port Choice: HTTP Listen Port
- HTTP Listen Port: 80

Origins

Origin Pools

1 item

Type	Name	Weight	Priority	Endpoint Subsets
Origin Pool	segment3-nginx	1	1	

Form Documentation JSON Search Edit Configuration

View

HTTP Load Balancer
nutanix-lb

- Other Settings
- VIP Advertisement : Custom
- Custom**
 - [Custom Advertise VIP Configuration](#)

Custom Advertise VIP Configuration

List of Sites to Advertise

1 item

Type	Site/VSite	Network/Segment	Listen Port
Segment on Site	dnakajima-nutanix-sms-1	demo-blue	0



Load balancer

The screenshot shows a configuration page for an HTTP Load Balancer. On the left, a sidebar contains a 'View' button and a list of settings: 'HTTP Load Balancer nutanix-lb', 'Other Settings', 'VIP Advertisement : Custom', 'Custom', 'Custom Advertise VIP Configuration', and 'List of Sites to Advertise'. The main area is titled 'Select Where to Advertise' and contains the following configuration items:

- Select Where to Advertise: Segment on Site
- Segment : system/demo-blue
- Site : system/dnakajima-nutanix-sms-1
- IP Address: 10.10.10.10
- TCP Listen Port Choice: Use Default TCP Listen Port

Load balancer (CE1)

```
root@node-0:/# rt --dump 16
Flags: L=Label Valid, P=Proxy ARP, T=Trap ARP, F=Flood ARP, H=Proxy ARP for host interface, V=Preferred VRF
Labels/Misc: S=SiteId N=NetworkId
vRouter inet4 routing table 0/16/unicast
Destination      PPL      Flags      Label      Nexthop  Labels/Misc
vRouter(Response): [No such file or directory (2)]
root@node-0:/# rt --dump 17
Flags: L=Label Valid, P=Proxy ARP, T=Trap ARP, F=Flood ARP, H=Proxy ARP for host interface, V=Preferred VRF
Labels/Misc: S=SiteId N=NetworkId
vRouter inet4 routing table 0/17/unicast
Destination      PPL      Flags      Label      Nexthop  Labels/Misc
Number of routes: 13

0.0.0.0/0         0        -          -          0
10.0.10.100/32   32       LV         270        489      (S,N):4972:48217 VN:65542
1:45 35:155
10.10.10.10/32   32       LV         270        584      (S,N):5049:48217 VN:65542
1:45 35:155
192.168.1.0/24   24       LV         280        489      (S,N):4972:48217 VN:65542
1:45 6:21 35:155
192.168.3.0/24   24       LV         279        489      (S,N):4972:48236 VN:65542 Stats UID:9
1:45 6:21 35:155
192.168.4.0/24   24       LV         279        529      (S,N):5446:48217 VN:65542
1:45 6:21 35:155
192.168.100.0/24 24       LV         279        501      (S,N):5049:48217 VN:65542
1:45 6:21 35:155
```

Load balancer (CE1)

```
root@node-0:/# nh --get 584
Id:584      Type:Receive      Fmly: AF_INET  Rid:0  Vrf:17
           Flags:Valid, Policy, Tunnel: None
           Oif:0
           Ports(Proto/Port/NH):
                   (1/8/34), (6/53/1), (6/80/585), (17/53/71),
           Ports(Proto/Port/NH):
                   (1/0-0/585), (1/8-8/34), (6/53-53/1), (6/80-80/585), (17/53-53/71), (58/0-
0/34),
           Default Port Match(Proto/NH):
                   (1/585), (58/34),

Id:1       Type:Drop          Fmly: AF_INET  Rid:0  Vrf:0
           Flags:Valid, Tunnel: None

Id:34     Type:NAT           Fmly: AF_INET  Rid:0  Vrf:7
           Flags:Valid, Policy, Tunnel: None
           NAT: SD VRF: 7 Family:2/0

           SIP: 169.254.254.4/32
           DIP: 172.28.13.36
           DIP1:

Id:585    Type:Encap         Fmly:AF_INET6  Rid:0  Vrf:17
           Flags:Valid, Policy, Tunnel: None
           Oif:4121 Len:14
           Encap Data: 00 00 5e 00 02 00 00 00 5e 00 01 00 86 dd
```

Load balancer (CE1)

```
root@node-0:/# vif --get 4121
Vrouter Interface Table
```

```
Flags: P=Policy, X=Cross Connect, S=Service Chain, Mr=Receive Mirror
Mt=Transmit Mirror, Tc=Transmit Checksum Offload, L3=Layer 3, L2=Layer 2
D=DHCP, Vp=Vhost Physical, Pr=Promiscuous, Vnt=Native Vlan Tagged
Mnp=No MAC Proxy, Dpdk=DPDK PMD Interface, Rfl=Receive Filtering Offload, Mon=Interface is Monitored
Nf=Dont Fragment, Vof=VLAN insert/strip offload, Mq=MultiQ, L=MAC Learning Enabled
Proxy=MAC Requests Proxied Always, Fr=Force Recv, Mn=Mirror without Vlan Tag, TTL=TTL not modified
Mn=Mirror without Vlan Tag, Tun=Tunnel, NoArp=No ARP interface, RawIp=RAW IP interface(No L2 header)
Npers=Non-Persistent, NoRPF=No RPF Checks, SynC=Syn Cookies, Gro=GRO Enabled
```

```
vif0/4121 Ethernet: ves-io-seg-volt-field-vhptnhxg-demo-blue_Envoy IP6Flow: 65542 Parent:vif0/6
Type:IP6Flow HWaddr:0:0:5e:0:1:0 Driver:
Vrf:17 Mcast Vrf:17 VN:0x10006 Flags:PL3Proxy Ref:2
MTU: 1520 EncMSS: 130 NH: 585
NAT Pool: 169.254.253.254
RX queue errors to lcore 0 0 0 0 0 0 0 0 0 0 0 0
RX packets:2968 bytes:566221 errors:0
TX packets:4438 bytes:356853 errors:0
Drops:0
Ingress ACL: 65 67
Egress ACL: 65 67
Service ACL: 65 67
IO Nat: 16
```

Load balancer (CE1)

```
"name": "ves-io-http-loadbalancer-nutanix-lb:volt-field-vhptnhxg:naas-demo",
  "domains": [
    "nutanix-lb.test",
    "nutanix-lb.test:80"
  ],
  "routes": [
    {
      "match": {
        "safe_regex": {
          "google_re2": {},
          "regex": "(.*?)"
        }
      },
      "route": {
        "cluster": "1fe31a5e-57d9-4a1e-81d3-2a779a77f3eb",
        "auto_host_rewrite": true,
        "timeout": "0s",
        "retry_policy": {
          "retry_on": "5xx",
          "num_retries": 1
        }
      },
      "name": "ves-io-http-loadbalancer-nutanix-lb-1:0"
    }
  ],
```

Load balancer (CE1)

```
"service_name": "1fe31a5e-57d9-4a1e-81d3-2a779a77f3eb"  
  },  
  "connect_timeout": "3.500s",  
  "per_connection_buffer_limit_bytes": 1048576,  
  "upstream_bind_config": {  
    "source_address": {  
      "address": "ff::100:13b9:100:0:1300",  
      "port_value": 0  
    },  
    "socket_options": [  
      {  
        "level": "1",  
        "name": "2",  
        "int_value": "1"  
      },  
      {  
        "level": "1",  
        "name": "15",  
        "int_value": "1"  
      }  
    ],  
    "flowlabel": 1  
  },  
}
```

Load balancer (CE1)

```
vif0/6      Ethernet: argo-envoy-0 Parent:vif0/-1
vtcpdump: verbose output suppressed, use -v[v]... for full protocol decode
listening on mon6, link-type EN10MB (Ethernet), snapshot length 262144 bytes
23:11:29.280696 IP6 fe80:2:0:1:6:0:c0a8:6464.45682 > fe80:1:0:1:6:0:a0a:a0a.http: Flags [S], seq 1523268449, win 29200, options [mss 1460,sackOK,TS
val 2077214626 ecr 0,nop,wscale 7], length 0
23:11:29.280705 IP6 fe80:1:0:1:6:0:a0a:a0a.http > fe80:2:0:1:6:0:c0a8:6464.45682: Flags [S.], seq 2126470129, ack 1523268450, win 31856, options
[mss 1460,sackOK,TS val 1273175143 ecr 2077214626,nop,wscale 7], length 0
23:11:29.281597 IP6 fe80:2:0:1:6:0:c0a8:6464.45682 > fe80:1:0:1:6:0:a0a:a0a.http: Flags [.], ack 1, win 229, options [nop,nop,TS val 2077214627 ecr
1273175143], length 0
23:11:29.281605 IP6 fe80:2:0:1:6:0:c0a8:6464.45682 > fe80:1:0:1:6:0:a0a:a0a.http: Flags [P.], seq 1:80, ack 1, win 229, options [nop,nop,TS val
2077214627 ecr 1273175143], length 79: HTTP: GET / HTTP/1.1
23:11:29.281608 IP6 fe80:1:0:1:6:0:a0a:a0a.http > fe80:2:0:1:6:0:c0a8:6464.45682: Flags [.], ack 80, win 249, options [nop,nop,TS val 1273175143 ecr
2077214627], length 0
23:11:29.282204 IP6 ff::100:13b9:100:0:1300.45477 > ff::13:6c00:2.http: Flags [S], seq 1231476591, win 32120, options [mss 1460,sackOK,TS val
1562631390 ecr 0,nop,wscale 7], length 0
23:11:29.286779 IP6 ff::13:6c00:2.http > ff::100:13b9:100:0:1300.45477: Flags [S.], seq 2256732437, ack 1231476592, win 28960, options [mss
1250,sackOK,TS val 3275469009 ecr 1562631390,nop,wscale 7], length 0
23:11:29.286783 IP6 ff::100:13b9:100:0:1300.45477 > ff::13:6c00:2.http: Flags [.], ack 1, win 251, options [nop,nop,TS val 1562631395 ecr
3275469009], length 0
23:11:29.286787 IP6 ff::100:13b9:100:0:1300.45477 > ff::13:6c00:2.http: Flags [P.], seq 1:215, ack 1, win 251, options [nop,nop,TS val 1562631395
ecr 3275469009], length 214: HTTP: GET / HTTP/1.1
23:11:29.292058 IP6 ff::13:6c00:2.http > ff::100:13b9:100:0:1300.45477: Flags [.], ack 215, win 235, options [nop,nop,TS val 3275469014 ecr
1562631395], length 0
23:11:29.292682 IP6 ff::13:6c00:2.http > ff::100:13b9:100:0:1300.45477: Flags [P.], seq 1:238, ack 215, win 235, options [nop,nop,TS val 3275469015
ecr 1562631395], length 237: HTTP: HTTP/1.1 200 OK
23:11:29.292686 IP6 ff::100:13b9:100:0:1300.45477 > ff::13:6c00:2.http: Flags [.], ack 238, win 250, options [nop,nop,TS val 1562631401 ecr
3275469015], length 0
23:11:29.293289 IP6 fe80:1:0:1:6:0:a0a:a0a.http > fe80:2:0:1:6:0:c0a8:6464.45682: Flags [P.], seq 1:245, ack 80, win 249, options [nop,nop,TS val
1273175155 ecr 2077214627], length 244: HTTP: HTTP/1.1 200 OK
```


Load balancer (CE1)

```
23:11:29.282204 IP6 ff::100:13b9:100:0:1300.45477 > ff::13:6c00:2.http: Flags [S], seq 1231476591,
win 32120, options [mss 1460,sackOK,TS val 1562631390 ecr 0,nop,wscale 7], length 0
23:11:29.286779 IP6 ff::13:6c00:2.http > ff::100:13b9:100:0:1300.45477: Flags [S.], seq 2256732437,
ack 1231476592, win 28960, options [mss 1250,sackOK,TS val 3275469009 ecr 1562631390,nop,wscale 7],
length 0
23:11:29.286783 IP6 ff::100:13b9:100:0:1300.45477 > ff::13:6c00:2.http: Flags [.], ack 1, win 251,
options [nop,nop,TS val 1562631395 ecr 3275469009], length 0
23:11:29.286787 IP6 ff::100:13b9:100:0:1300.45477 > ff::13:6c00:2.http: Flags [P.], seq 1:215, ack 1,
win 251, options [nop,nop,TS val 1562631395 ecr 3275469009], length 214: HTTP: GET / HTTP/1.1
23:11:29.292058 IP6 ff::13:6c00:2.http > ff::100:13b9:100:0:1300.45477: Flags [.], ack 215, win 235,
options [nop,nop,TS val 3275469014 ecr 1562631395], length 0
23:11:29.292682 IP6 ff::13:6c00:2.http > ff::100:13b9:100:0:1300.45477: Flags [P.], seq 1:238, ack
215, win 235, options [nop,nop,TS val 3275469015 ecr 1562631395], length 237: HTTP: HTTP/1.1 200 OK
```

Load balancer (CE1)

```
root@node-0:/# rt --dump 16 --family inet6
Flags: L=Label Valid, P=Proxy ARP, T=Trap ARP, F=Flood ARP, H=Proxy ARP for host interface, V=Preferred VRF
Labels/Misc: S=SiteId N=NetworkId
vRouter inet6 routing table 0/16/unicast
Destination      PPL      Flags      Label      Nexthop  Labels/Misc
                Number of routes: 31
::/0             0        L          278        11       (S,N):5049:19 VN:65541
                1:45 35:155
ff::13:6c00:2/128 128      L          284        489      (S,N):4972:19 VN:65541
                1:45 35:155
ff::13:b900:2/128 128      L          283        544      (S,N):5049:19 VN:65541
                1:45 35:155
```

Load balancer (CE1)

```
root@node-0:/# nh --get 489
Id:489      Type:Tunnel      Fmly: AF_INET  Rid:0  Vrf:8
           Flags:Valid, MPLSoUDP, Indirect, Tunnel: None
           Direct NH(label): 532
           Oif:0  Len:14  Data:00 00 00 00 00 00 00 00 00 00 00 00 00 00
           Sip:10.101.10.39 Dip:10.101.9.245

Id:532      Type:Composite     Fmly: AF_INET  Rid:0  Vrf:0
           Flags:Valid, Ecmp, Tunnel: None
           Composite Flags:
           Sub NH(label): 85(0)(W:100) 86(0)(W:100)

Id:85      Type:Encap        Fmly: AF_INET  Rid:0  Vrf:8
           Flags:Valid, Tunnel: None
           Oif:4111  Len:14
           Encap Data: 00 00 5e 00 01 00 00 00 5e 00 01 00 08 00

Id:86      Type:Encap        Fmly: AF_INET  Rid:0  Vrf:8
           Flags:Valid, Tunnel: None
           Oif:4109  Len:14
           Encap Data: 00 00 5e 00 01 00 00 00 5e 00 01 00 08 00
```

Load balancer (CE1)

```
root@node-0:/# vif --get 4111
Vrouter Interface Table
```

```
Flags: P=Policy, X=Cross Connect, S=Service Chain, Mr=Receive Mirror
Mt=Transmit Mirror, Tc=Transmit Checksum Offload, L3=Layer 3, L2=Layer 2
D=DHCP, Vp=Vhost Physical, Pr=Promiscuous, Vnt=Native Vlan Tagged
Mnp=No MAC Proxy, Dpdk=DPDK PMD Interface, Rfl=Receive Filtering Offload, Mon=Interface is Monitored
Nf=Dont Fragment, Vof=VLAN insert/strip offload, Mq=MultiQ, L=MAC Learning Enabled
Proxy=MAC Requests Proxied Always, Fr=Force Recv, Mn=Mirror without Vlan Tag, TTL=TTL not modified
Mn=Mirror without Vlan Tag, Tun=Tunnel, NoArp=No ARP interface, RawIp=RAW IP interface(No L2 header)
Npers=Non-Persistent, NoRPF=No RPF Checks, SynC=Syn Cookies, Gro=GRO Enabled
```

```
vif0/4111 Virtual: tunnel_ty8-tky_0 Parent:vif0/-1
Type:Tunnel HWaddr:0:1:2:3:4:5 Driver:
UserVisible
Vrf:8 Mcast Vrf:8 VN:0x0 Flags:L3Tun Ref:69
MTU: 1370 EncMSS: 0 NH: 85 FNH: 31
Outbound SA:SAId: 0xffff0060 Tunnel SRC:172.28.13.36 Tunnel DST:103.135.56.40
Tunnel Mode:Tunnel Window size:0 VIF: 0 State: Active
Inbound SA:SAId: 0xffff5f70 Tunnel SRC:103.135.56.40 Tunnel DST:172.28.13.36
Tunnel Mode:Tunnel Window size:4 VIF: 4111 State: Active
NAT Pool: 10.101.10.39
RX queue errors to lcore 0 0 0 0 0 0 0 0 0 0 0 0
RX packets:86249467 bytes:54224831979 errors:0
TX packets:83322562 bytes:44075751732 errors:18
Drops:39096
TCP reset ignored:36
Service ACL: 65 67
```

Load balancer (CE1)

```
root@node-0:/# vif --get 4111
Vrouter Interface Table
```

```
Flags: P=Policy, X=Cross Connect, S=Service Chain, Mr=Receive Mirror
Mt=Transmit Mirror, Tc=Transmit Checksum Offload, L3=Layer 3, L2=Layer 2
D=DHCP, Vp=Vhost Physical, Pr=Promiscuous, Vnt=Native Vlan Tagged
Mnp=No MAC Proxy, Dpdk=DPDK PMD Interface, Rfl=Receive Filtering Offload, Mon=Interface is Monitored
Nf=Dont Fragment, Vof=VLAN insert/strip offload, Mq=MultiQ, L=MAC Learning Enabled
Proxy=MAC Requests Proxied Always, Fr=Force Recv, Mn=Mirror without Vlan Tag, TTL=TTL not modified
Mn=Mirror without Vlan Tag, Tun=Tunnel, NoArp=No ARP interface, RawIp=RAW IP interface(No L2 header)
Npers=Non-Persistent, NoRPF=No RPF Checks, SynC=Syn Cookies, Gro=GRO Enabled
```

```
vif0/4111 Virtual: tunnel_ty8-tky_0 Parent:vif0/-1
Type:Tunnel HWaddr:0:1:2:3:4:5 Driver:
UserVisible
Vrf:8 Mcast Vrf:8 VN:0x0 Flags:L3Tun Ref:69
MTU: 1370 EncMSS: 0 NH: 85 FNH: 31
Outbound SA:SAId: 0xffff0060 Tunnel SRC:172.28.13.36 Tunnel DST:103.135.56.40
Tunnel Mode:Tunnel Window size:0 VIF: 0 State: Active
Inbound SA:SAId: 0xffff5f70 Tunnel SRC:103.135.56.40 Tunnel DST:172.28.13.36
Tunnel Mode:Tunnel Window size:4 VIF: 4111 State: Active
NAT Pool: 10.101.10.39
RX queue errors to lcore 0 0 0 0 0 0 0 0 0 0 0 0
RX packets:86249467 bytes:54224831979 errors:0
TX packets:83322562 bytes:44075751732 errors:18
Drops:39096
TCP reset ignored:36
Service ACL: 65 67
```

Load balancer (CE2)

```
root@node-0:/# mpls --get 284
MPLS Input Label Map
```

```
Label      NextHop
-----
284        605
root@node-0:/# nh --get 605
Id:605      Type:NAT          Fmly: AF_INET  Rid:0  Vrf:15
            Flags:Valid, Policy, Tunnel: None
            NAT: SSpDDp VRF: 17 Family:2/0

            SIP: 192.168.3.1/32
            SPORT: <8192, 65535>
            DIP: 192.168.3.100
            DIP1:
            DPortMap:
                6, 80->80
```

Load balancer (CE2)

```
root@node-0:/# rt --get 192.168.3.100/32 --vrf 17
Match 192.168.3.100/32 in vRouter inet4 table 0/17/unicast
```

Flags: L=Label Valid, P=Proxy ARP, T=Trap ARP, F=Flood ARP, H=Proxy ARP for host interface, V=Preferred VRF

Labels/Misc: S=SiteId N=NetworkId

vRouter inet4 routing table 0/17/unicast

Destination	PPL	Flags	Label	Nexthop	Labels/Misc
192.168.3.100/32	32	LV	0	524	(S,N):4972:48236 VN:65542 1:32 6:175 21:54

```
root@node-0:/# nh --get 524
```

```
Id:524      Type:Encap      Fmly: AF_INET  Rid:0  Vrf:17
Flags:Valid, Tunnel: None
Oif:4113  Len:18
Encap Data: 00 0c 29 b7 22 94 00 0c 29 0f f4 a1 81 00 01 2e 08 00
```

Load balancer (CE2)

```
root@node-0:/# vif --get 4113
Vrouter Interface Table
```

```
Flags: P=Policy, X=Cross Connect, S=Service Chain, Mr=Receive Mirror
Mt=Transmit Mirror, Tc=Transmit Checksum Offload, L3=Layer 3, L2=Layer 2
D=DHCP, Vp=Vhost Physical, Pr=Promiscuous, Vnt=Native Vlan Tagged
Mnp=No MAC Proxy, Dpdk=DPDK PMD Interface, Rfl=Receive Filtering Offload, Mon=Interface is Monitored
Nf=Dont Fragment, Vof=VLAN insert/strip offload, Mq=MultiQ, L=MAC Learning Enabled
Proxy=MAC Requests Proxied Always, Fr=Force Recv, Mn=Mirror without Vlan Tag, TTL=TTL not modified
Mn=Mirror without Vlan Tag, Tun=Tunnel, NoArp=No ARP interface, RawIp=RAW IP interface(No L2 header)
Npers=Non-Persistent, NoRPF=No RPF Checks, SynC=Syn Cookies, Gro=GRO Enabled
```

```
vif0/4113 Ethernet: ens224.302 Vlan(o/i)(,S): 302/302 (Speed 0, Duplex 0) Parent:vif0/24
Type:Physical(Vlan) HWaddr:0:c:29:f:f4:a1 Driver: vr_vlan
CrossConnectIdx: 25 UserVisible
Address: 192.168.3.1
Vrf:17 Mcast Vrf:17 VN:0x0 Flags:PL3 Ref:5
MTU: 1500 EncMSS: 130 NH: 498
NAT Pool: 192.168.3.1 192.168.3.1
RX queue errors to lcore 0 0 0 0 0 0 0 0 0 0 0 0
RX packets:19135 bytes:1545288 errors:0
TX packets:19748 bytes:1162588 errors:0
Drops:11
Ingress ACL: 65 67
Egress ACL: 65 67
Service ACL: 65 67
```


Load balancer (CE2)

```
root@node-0:/# flow --get 19524
Flow Index:                19524
Flow Generation ID:       17
Reverse Flow Index:       229888
VRF:                       17
Destination VRF:          15
Flow Source:               [192.168.3.100]:80
Flow Destination:         [192.168.3.1]:51219
Flow Protocol:             TCP
Flow Action:               NAT: SourceNAT, SourcePortNAT, DestinationNAT, DestinationPortNAT,
                           NAT(Source, Destination): [ff::13:6c00:2]:80, [ff::100:13b9:100:0:1300]:33181
```

