

EVPN edge routing (incl. type 5 routes) best practises

March 2019

written by Johannes Resch

Below a collection of highly recommended config knobs to use EVPN VXLAN edge routing on QFX5110 leave devices.

For edge-routing on QFX5110 leaves, JunOS 18.1R3-S3 (or newer service releases) should be used for the time being.

The listed config snippets apply for the following design:

- edge routing on QFX5110 leaves
- EVPN model used is “VLAN aware service” (using single default-switch instance), n:1 mapping of VLAN to EVI
- use of “virtual-gateway address” and a unique IP per leaf switch IRB interface
- eBGP is used for underlay
- iBGP is used for overlay
- Underlay ASN: one ASN for all spines, each leaf uses its own ASN

Configure `ecmp-fast-reroute` and `composite NH`, aside with standard `ecmp` policy.

```
routing-options {
  forwarding-table {
    export POL_ECMP;
    ecmp-fast-reroute;
    chained-composite-next-hop {
      ingress {
        evpn;
      }
    }
  }
}
```

```
policy-options {
  policy-statement POL_ECMP {
    term T_ECMP {
      then {
        load-balance per-packet;
      }
    }
  }
}
```

Note: Without chained-composite-next-hop, PFE would not program the tunnel NH.

Configure overlay-ecmp:

```
forwarding-options {
  vxlan-routing {
    overlay-ecmp;
  }
}
```

Note: Configuring this statement will restart PFE.

Configure “multipath multiple-as” for underlay and overlay BGP sessions.

Routing-instances vs. loopback-interfaces

Make sure every routing-instance used is configured with a loopback interface. Even if not required otherwise, assign a “dummy” loopback logical interface. No configuration aside from “family inet” is required on the dummy loopback.

```
interfaces {
  lo0 {
    unit 1 {
      family inet;
    }
  }
}
```

```
routing-instances F00 {
  interface lo0.1;
```

```
}
```

Note: if you are using control-plane protection firewall filters (which you should), keep in mind that filters will have to be configured on all loopback interfaces to also filter access from VRFs to the device control plane.

IRB configuration

Make sure the local/unique IRB IP is configured with the “preferred” knob. It is mandatory when VGA IP address is lower than one assigned to the IRB interface, as this will break ARP processing. IRB IP address is needed to process ARP requests and by default lower IP is preferred.

```
interfaces {
  irb {
    unit 1234 {
      virtual-gateway-accept-data;
      description IRB_1234;
      family inet {
        address 192.168.1.254/24 {
          preferred;
          virtual-gateway-address 192.168.1.1;
        }
      }
    }
  }
}
```

Config statements NOT to use for this topology

Make sure the following configuration statements are NOT applied:

- IRB proxy-macip-advertisement
- default-gateway do-not-advertise
- default-gateway no-gateway-community

Other topics to be aware of

With 18.1R3, the following features are enabled by default, which might influence forwarding behavior in unexpected ways:

- Duplicate MAC detection
(https://www.juniper.net/documentation/en_US/junos/topics/reference/configuration-statements/duplicate-mac-detection.html)
- ARP/NDP suppression
(https://www.juniper.net/documentation/en_US/junos/topics/concept/evpn-proxy-arp-support.html)

On QFX5110, 8'000 nexthops for EVPN-VXLAN are allocated by default. In case more entries are needed, this can be configured (change requires PFE restart):

https://www.juniper.net/documentation/en_US/junos/topics/reference/configuration-statement/next-hop-edit-forwarding-options-vxlan-routing.html

QFX5110 does not support routing between VXLAN and non-VXLAN interfaces.