

## Layer 2 Networking and IP Addressing

**File Location:** /etc/network/interfaces

### inet Types:

- loopback
- dhcp
- *Other, Non-specified* (L2/L3 ports)

```
# loopback interface
auto lo
iface lo inet loopback
address 10.2.1.1/32

# oob interface
auto eth0
iface eth0 inet dhcp


# front panel port interface
auto swp1
iface swp1
address 172.16.1.2/24
mtu 9216
link-speed 10000
link-duplex full
link-autoneg off
```

Linux Bridge in vlan-aware mode:

```
auto bridge
iface bridge
bridge-vlan-aware yes
bridge-ports swp2 swp3 swp4
bridge-vids 1-200
bridge-pvid 1

#Switch VLAN Interface for VLAN10
auto bridge.10
iface bridge.10
address 172.16.10.2/24
```

## MLAG (Multichassis Link Aggregation)



```
auto peerlink.4094
iface peerlink.4094
address 169.254.1.1/30
clagd-peer-ip 169.254.1.2
clagd-backup-ip 10.0.1.50
clagd-sys-mac 44:39:39:FF:40:94

# ToR pair #1
auto downlink1
iface downlink1
bond-slaves swp29 swp30
clag-id 1
```

## Routing

**File Location:** /etc/quagga/Quagga.conf

### Supported Routing Protocols:

- BGP
- OSPF

### Configure OSPF for swp1 and put it into area 0

```
!
service integrated-vtysh-config
!
interface swp1
ip ospf area 0.0.0.0
!
router ospf
ospf router-id 10.2.1.1
!
```

**Dameon Location:** /etc/quagga/daemons

```
root@leaf1:/etc/cumulus# cat /etc/quagga/daemons
zebra=yes
bgpd=yes
ospfd=yes
ospf6d=yes
```

## Security (CoPP, iptables, acl)

**File Location:** /etc/cumulus/acl/policy.d/\*.rules

```
[iptables]
-A FORWARD -o swp1 -p tcp --dport 200 -j DROP
```

### Tables:

- **iptables** corresponds to IPv4
- **ip6tables** corresponds to IPv6
- **eatables** corresponds to layer2

### Chains:

- **FORWARD** corresponds to traffic through the switch
- **INPUT** corresponds to traffic with destination on the Switch (e.g. loopback)
- **OUTPUT** corresponds to traffic being generated by switch

### Applying Rules to Multiple Swps

```
-A OUTPUT,FORWARD -o swp+ -p tcp --sport 123 -j DROP
```

### Applying Policing to Swp


```
-A FORWARD --in-interface swp1 -j POLICE --set-mode KB
--set-rate 125000 --set-burst 2000
```

## VRR (Virtual Router Redundancy)

VRR for VLAN 10 configured under /etc/network/interfaces

```
auto bridge.10
iface bridge.10
address 172.16.10.2/24
address-virtual 00:00:5e:00:01:01 172.16.10.1/24
```

# BGP Unnumbered



**File Location:** /etc/quagga/Quagga.conf

```
interface swp1
  ipv6 nd ra-interval 3
  no ipv6 nd suppress-ra

router bgp 65000
  bgp router-id 10.2.1.10
  bgp default show-hostname
  bgp bestpath as-path multipath-relax no-as-set
  bgp bestpath compare-routerid
  network 10.2.1.10/32
  network 10.2.1.12/32
  neighbor PEER_GROUP peer-group
  neighbor PEER_GROUP remote-as external
  neighbor PEER_GROUP description Remote Peer
  neighbor PEER_GROUP advertisement-interval 0
  neighbor PEER_GROUP timers 1 3
  neighbor PEER_GROUP timers connect 3
  neighbor PEER_GROUP capability extended-nexthop
  neighbor swp1 interface
  neighbor swp1 remote-as external
  neighbor swp1 peer-group PEER_GROUP
  maximum-paths 64
!
address-family ipv6
  network fd::10/128
  neighbor PEER_GROUP activate
  neighbor swp1 peer-group PEER_GROUP
exit-address-family
!
```

**Relevant BGP IETF Info:**

- <https://tools.ietf.org/html/rfc5549>
- <https://www.ietf.org/id/draft-walton-bgp-hostname-capability-02>
- <https://tools.ietf.org/html/draft-lapukhov-bgp-routing-large-dc-02>

**BGP Configuration Guide:**

<https://docs.cumulusnetworks.com/display/DOCS/Configuring+Border+Gateway+Protocol+-+BGP>

**Troubleshooting Commands:**

show ip bgp summary	show neighbor adjacency
show ip bgp	show IPv4 networks
show ip bgp neighbors <swp#> advertised-routes	show IP routes advertised
show ipv6 bgp	Show IPv6 networks
Show ipv6 bgp neighbors <swp#> advertised-routes	show IPv6 routes advertised

**Explanation of Commands:**

- > ipv6 nd ra-interval 3  
Router Advertisement interval set to 3 seconds
- > no ipv6 nd suppress-ra  
Disable Suppression of Router Advertisements
- > router bgp 65000  
Switch is running in AS 65000
- > bgp router-id 10.2.1.10  
Router ID matches loopback IP 10.2.1.10
- > bgp default show-hostname  
Hostname for troubleshooting enabled
- > bgp bestpath as-path multipath-relax no-as-set  
multipath-relax allows ECMP across multiple paths
- > bgp bestpath compare-routerid  
Use router-id of adjacent devices to consider routes
- > network 10.2.1.10/32 & network 10.2.1.12/32  
Implicitly advertise loopback and the attached VM
- > neighbor PEER\_GROUP peer-group  
Create peer group called PEER\_GROUP
- > neighbor PEER\_GROUP remote-as external  
Peers using this peer group will be in another AS (not 65000)
- > neighbor PEER\_GROUP description Remote Peer  
Human readable plain text description for this peer group
- > neighbor PEER\_GROUP advertisement-interval 0  
Send new advertisements immediately
- > neighbor PEER\_GROUP timers 1 3  
Set keepalive timer to 1 second and hold timer to 3 seconds
- > neighbor PEER\_GROUP timers connect 3  
Try to reconnect to down neighbor in 3 seconds
- > neighbor PEER\_GROUP capability extended-nexthop  
Allows RFC5549 to work (use IPv6 link-locals for IPv4 Adjacency)
- > neighbor swp1 interface  
Enable swp1 interface for BGP
- > neighbor swp1 remote-as external  
Remote neighbor on swp1 on different AS (not 65000)
- > neighbor swp1 peer-group PEER\_GROUP  
Use peer-group PEER\_GROUP parameters for swp1 peer
- > maximum-paths 64  
If multiple paths, allow switch to use up to 64 ECMP routes
- > network fd::10/128  
Advertise fd::10/128 into IPv6 BGP routing
- > neighbor PEER\_GROUP activate  
Activate peer\_group PEER\_GROUP for IPv6
- > neighbor swp1 peer-group PEER\_GROUP  
Enable swp1 for IPv6, and to use peer\_group PEER\_GROUP

**Logging:**

```
Create Log File
switch(config)# log file /var/log/quagga/bgplot

Use Debugs
switch# debug bgp neighbor-events
BGP neighbor-events debugging is on
```