



Campus Networks: Layer 3 routing

Routing intro / Routing security / Routing protocols

Author:
Sami Ait Ali Oulahcen


Nouakchott, Mauritania
17-22 February 2025

Intro

What is a router? / Routing & Forwarding / Control plane & Data plane

? What is a router

- Layer 3 equipment
- Used to interconnect networks
- Looks for packet destinations then forwards accordingly



Routing & Forwarding

- Routing = building maps and giving directions
- Forwarding = moving packets between interfaces according to the directions



Control plane & Data plane


- Key concepts in SDN (Software Defined Networks) and Cloud computing
- Control plane: Controls how packets are moved. This includes routing/forwarding tables and policies.
- Data plane (aka Forwarding Plane): Uses data from the control plane to do the actual forwarding.

Routing table & Forwarding table

Connected & Static routes / Routing protocols / RIB & FIB


Connected routes

- **Connected routes:** refer to routes that are directly associated with an interface



Connected routes

- **Connected routes:** refer to routes that are directly associated with an interface



Static routes

- **Static routes:** are routes explicitly defined by the network administrator

ip	route	destination	prefix	next hop
ip	route	203.0.108.64	255.255.255.192	10.0.2.254

ipv6	route	destination/prefix	next hop
ipv6	route	2001:db8:d1:B203::/64	2001:db8:d1:2::254

Routing protocols

- **Routing protocols:** standardized methods to determine and communicate best paths for packets across a network

* 203.0.108.0/26 [20/0] via 192.0.2.1, BGP, 00:05:12

203.0.108.0/26 [115/20] via 10.1.1.2, IS-IS, 00:02:30

203.0.108.0/26 [110/30] via 10.2.2.1, OSPF, 00:01:45

More on routing protocols in the next chapter.

Routing Table

- Each routing protocol builds its own routing table (Local RIB)
- Routing table (Global RIB) is built from connected routes, static routes and all routing protocols' local RIBs
- Updated periodically or as topology changes (event driven)

```
sw00.wan.noc.rabat#sh ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is 196.200.160.1 to network 0.0.0.0


S*    0.0.0.0/0 [1/0] via 196.200.160.1
      1.0.0.0/8 is variably subnetted, 5 subnets, 5 masks
B      1.10.10.0/24 [200/0] via 10.0.0.1, 4d13h
B      1.51.0.0/16 [200/0] via 10.0.0.1, 00:01:36
B      1.51.32.0/20 [200/0] via 10.0.0.1, 00:01:36
```

Forwarding Table

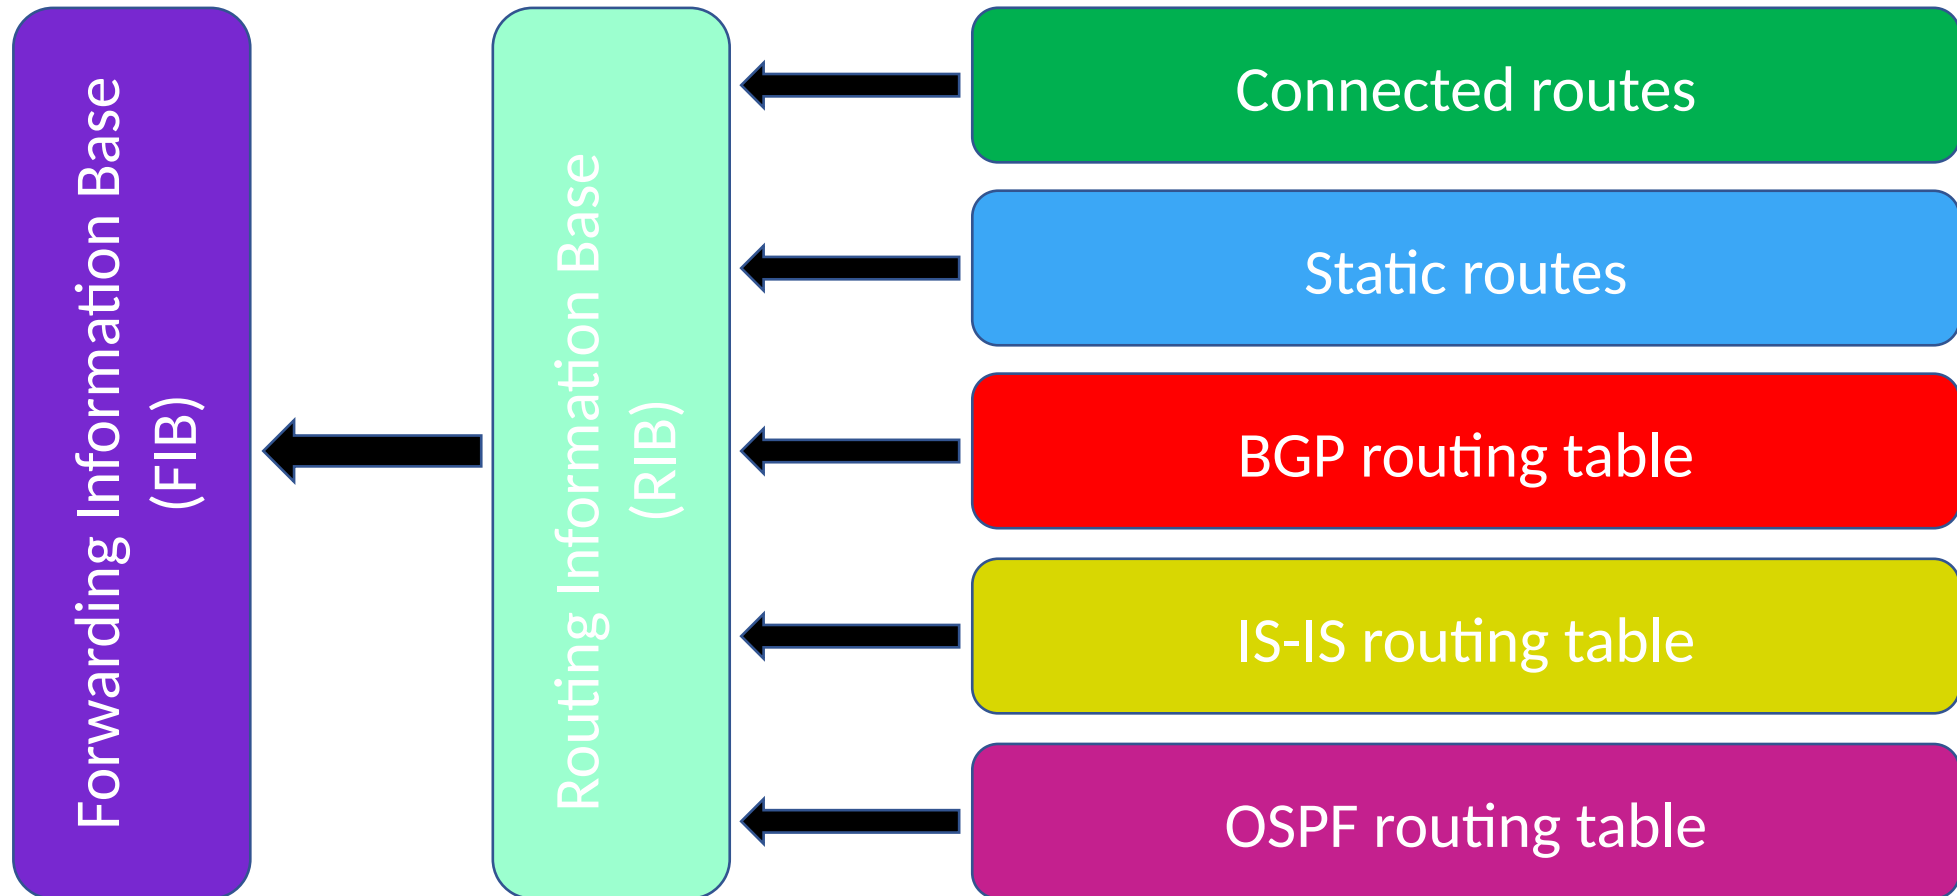
- Built from the routing table by selecting the best route for each destination network.
- Contains prefix, next hop, and corresponding outgoing interface
- Sometimes we call it the routing table

```
sw00.wan.noc.rabat#sh ip cef
```

Prefix	Next Hop	Interface
0.0.0.0/0	196.200.160.1	Vlan160
0.0.0.0/8	drop	
0.0.0.0/32	receive	
1.10.10.0/24	10.0.0.1	Vlan161
1.51.0.0/16	10.0.0.1	Vlan161
1.51.32.0/20	10.0.0.1	Vlan161



Routing Table & Forwarding Table summary



What does the router do ?

- Router looks at the packet's **destination address** then forwards the packet to the correct interface
- Router uses **longest match routing**: the more specific prefix is always preferred over the less specific prefix

Example: where do we route 203.0.108.102 ?

203.0.108.64/26 [20/0] via 10.0.0.2

203.0.108.0/24 [1/0] via 203.0.108.254

- When the address doesn't match any entry in the forwarding table:

Default route is used if it exists

OR Router **drops** packet and sends ICMP destination unreachable

What does the router do ?

- Router looks at the packet's **destination address** then forwards the packet to the correct interface
- Router uses **longest match routing**: the more specific prefix is always preferred over the less specific prefix

Example: where do we route 203.0.108.102 ?

203.0.108.64/26 [20/0] via 10.0.0.2

203.0.108.0/24 [1/0] via 203.0.108.254

- When the address doesn't match any entry in the forwarding table:

Default route is used if it exists

OR Router **drops** packet and sends ICMP destination unreachable

Border routers for small sites ($\leq 1\text{Gbps}$ WAN)

MikroTik

Model: CCR2116-12G-4S+

Price: \$1,000



Cisco

Model: Catalyst C8200-1N-4T

Price: \$4,300 (+yearly license)



Huawei

Model: AR2204XE


Price: \$1000



Fortinet

Model: Fortigate 60F

Price: \$1,000



Border routers for medium sites ($\leq 10\text{Gbps}$ WAN)

MikroTik

Model: CCR2216-1G-12XS-2XQ

Price: \$3,000



Cisco

Model: ASR1001X-10G-K9

Price: \$19,000



Juniper

Model: MX204-HW-BASE

Price: \$23,000



Fortinet

Model: Fortigate 600F

Price: \$19,000



Routing protocols

IGP / EGP

Routing protocols

- Standardized methods to determine and communicate best paths for packets across a network. Two types:
 - IGP (Interior Gateway Protocol) is used to designate the process running on routers inside a network.
 - EGP (Exterior Gateway Protocol) is used to designate the process running between routers bordering directly connected networks.

IGP

- Within an autonomous system
- Transmits information on internal network prefixes
- Two widely used IGPs:
 - OSPF (Open Shortest Path First)
 - IS-IS (Intermediate System to Intermediate System)

EGP

- Used to transmit routing information between autonomous systems
- Decoupled from the IGP
- The only EGP is the Border Gateway Protocol (BGP)

Autonomous System (AS)

- A set of IP networks and routers that are under the control of a single administrative entity and share a common routing policy.
- Used as part of the Border Gateway Protocol (BGP) to exchange routing information between autonomous systems on the Internet.


AS characteristics

- **AS number:** Each autonomous system is assigned a unique number, known as an Autonomous System Number (**ASN**), which is used for identification in BGP
- **Routing policy:** ASes can have different routing policies, which means they can choose how and when to exchange routing information with other ASes
- **Interconnection:** ASes can be connected to each other by direct links or via Internet Exchange Points (IXPs).

Common terms

- Neighbors : neighboring ASes that directly exchange routing information
- Announce: send routing information to a neighbor
- Accept: receive and use routing information sent by a neighbor
- Originate: insert routing information into external announcements (usually following the IGP)
- Peers: routers in neighboring ASes or within an AS that exchange routing and policy information

? How does BGP work



- For AS1 and AS2 networks to communicate :
 - AS1 must announce to AS2
 - AS2 must accept AS1's announcement
 - AS2 must announce to AS1
 - AS1 must accept AS2's announcement

Default Administrative Distances

Route Source	Cisco	Juniper	Huawei	Dell	Nokia	MikroTik
Connected Interface	0	0	0	0	0	0
Static Route	1	5	60	1	1	1
External BGP	20	170	255	20	170	20
Internal BGP	200	170	255	200	130	200
IS-IS	115	18	15	115	18	N/A
OSPF	110	10	10	110	10	110
Unknown	255	255	?	255	?	


Routing security

Filtering / RPKI / MANRS

Useful tips for border routers


- Route unused subnets to NULL0
- Use uRPF to prevent spoofing from your network (BCP10)
- Use ROV to filter invalid RPKI prefixes
- Sign your prefixes using RPKI to secure them against hijacking or mis-origination
- Check MANRS (Mutually agreed norms for routing security) for network operators for current best practices

? What is uRPF



- The router compares the packet's source address with the FIB entry:
 - If the FIB entry corresponds to the incoming interface, the packet is transferred
 - If the FIB entry doesn't correspond to the incoming interface, the packet is dropped

? What is uRPF



- The router compares the packet's source address with the FIB entry:
 - If the FIB entry corresponds to the incoming interface, the packet is transferred
 - If the FIB entry doesn't correspond to the incoming interface, the packet is dropped

URPF Configuration

```
!  
interface VLAN 230  
description Students in Building 2  
ip address 203.0.230.1 255.255.255.0  
ip verify unicast reverse-path  
ipv6 address 2001:DB8:2:30::1/64  
ipv6 verify unicast reverse-path  
!  
interface VLAN 302  
description Labs in Building 3  
ip address 203.0.302.1 255.255.255.0  
ip verify unicast reverse-path  
ipv6 address 2001:DB8:3:2::1/64  
ipv6 verify unicast reverse-path  
!
```

RPKI

- A framework designed to secure the routing infrastructure of the Internet
- Provides cryptographic evidence for the association between IP address blocks and their legitimate owners
- Mitigates route hijacking and misconfigurations by verifying route origins

Key components of RPKI

- Trust anchors: Root certification authorities (RIRs) that issue certificates
- Route Origin Authorizations (ROAs): Cryptographic objects that certify which Autonomous System (AS) is authorized to create routes for specific IP prefixes
- Route Origin Validators (ROVs): Infrastructure elements that publish and verify the validity of certificates and ROAs.

ROV Configuration

- Point the router to an RPKI cache (10.0.0.34 for example):

```
router bgp 64512
  bgp rpki server tcp 10.0.0.34 port 43779 refresh 60
```

- BGP table after using ROV:

```
RPKI validation codes: V valid, I invalid, N Not found

Network      Metric LocPrf Path
N*> 1.0.4.0/24      0      37100 6939 4637 1221 38803 56203 i
N*> 1.0.5.0/24      0      37100 6939 4637 1221 38803 56203 i
...
V*> 1.9.0.0/16      0      37100 4788 i
N*> 1.10.8.0/24      0      37100 10026 18046 17408 58730 i
N*> 1.10.64.0/24     0      37100 6453 3491 133741 i
...
V*> 1.37.0.0/16     0      37100 4766 4775 i
N*> 1.38.0.0/23     0      37100 6453 1273 55410 38266 i
N*> 1.38.0.0/17     0      37100 6453 1273 55410 38266 {38266} i
...
I* 5.8.240.0/23     0      37100 44217 3178 i
I* 5.8.241.0/24     0      37100 44217 3178 i
I* 5.8.242.0/23     0      37100 44217 3178 i
I* 5.8.244.0/23     0      37100 44217 3178 i
...
```

Courtesy of SEACOM: <http://as37100.net>

Questions ?