

18. Addresses & Routes

HOST ROUTING LOGIC

A host creates an IP packet addressed to the final recipient and encapsulates it in an L2 frame addressed as follows:

- If the destination is local (part of our own subnet), send it directly. Use the L2 (MAC) address of the destination as the Ethernet frame destination address. If you don't already know it from your ARP cache, do a fresh ARP request to learn it.
- If the destination is outside of our own subnet, address the frame to the default gateway

ROUTER LOGIC

When receiving a frame, decide whether to process it

- Check for CRC errors in the FCS trailer of the frame
- Check that the destination MAC address belongs to the router (or is a relevant multicast/broadcast)

Perhaps Process It

- Strip the old L2 frame from the L3 packet. It already did its job of getting the packet as far as the router.
- Choose an exit interface (and maybe a next-hop address) by comparing the packet's destination address to the routing table
- Encapsulate the packet in a fresh frame of the right kind (Ethernet/HDLC/PPP) for the next hop, using ARP as necessary to address it.
- Send the new frame out that port.

Routers learn routes in 3 ways:

- **Connected Routes**—When you add an IP address & mask to an interface, you've also told the router how to reach that (sub)net. It'll automatically create a route for that net, using that exit interface, *if* the exit interface is usable (up/up).
- **Static Routes**—You can manually add a route to the routing table, using the "ip route ..." command.
- **Routing Protocols** can add routes they learn from other routers

ROUTER CONFIGURATION

Adding an interface address is just like with switch VLAN interfaces

```
interface gi0/0
  ip address 10.0.0.1 255.255.255.0
  no shut
```

Viewing a router's ARP Cache is a little different from the MAC address table on a switch

```
edge# show ip arp
Protocol Address           Age (min)  Hardware Addr  Type   Interface
Internet 69.12.233.1          42        0002.3b01.7ce6  ARPA   FastEthernet0/1
Internet 69.12.233.218        -         0018.bad1.a461  ARPA   FastEthernet0/1
Internet 192.168.22.1         -         0018.bad1.a460  ARPA   FastEthernet0/0
Internet 192.168.22.100      26        0016.cba6.9f58  ARPA   FastEthernet0/0
Internet 192.168.22.105      9         0016.cbad.3e8a  ARPA   FastEthernet0/0
Internet 192.168.22.210     189       0005.5e2f.40c0  ARPA   FastEthernet0/0
Internet 192.168.22.250     53        a418.753c.8dc2  ARPA   FastEthernet0/0
```

R O U T E R - O N - A - S T I C K

Router-on-a-Stick—Routing between VLANs on a trunk requires subinterfaces (one for each VLAN/subnet) and a little extra configuration to handle the 802.1Q encapsulation.

```
interface FastEthernet0/0
  no ip address
  !
interface FastEthernet0/0.1
  The subinterface number has no bearing on the VLAN #, that's in the encapsulation command
  description WILD vlan
  encapsulation dot1Q 1
  This subinterface will receive all frames encapsulated with 802.1Q for VLAN 2
  ip address 192.168.1.1 255.255.255.0
  !
interface FastEthernet0/0.5
  description CAGED vlan for DMZ Jails
  encapsulation dot1Q 5
  ip address 192.168.5.1 255.255.255.0
  !
interface FastEthernet0/0.29
  description TAME vlan
  encapsulation dot1Q 29
  ip address 192.168.29.1 255.255.255.0
```

There are two ways to handle the native VLAN

- Put an address on the underlying interface with no encapsulation

```
interface fa0/0
  ip address 192.168.1.1 255.255.255.0
  !
interface fa0/0.2
  encapsulation dot1Q 2
  ...
```
- declare one of the subinterfaces with 802.1Q encapsulation to be native

```
interface fa0/0.58
  encapsulation dot1Q 58 native
  ...
```

To see VLAN configuration and statistics, including which VLAN is native on which interface,

```
edge# show vlans
Virtual LAN ID: 1 (IEEE 802.1Q Encapsulation)

    VLAN Trunk Interface: FastEthernet0/0

This is configured as native Vlan for the following interface(s) :
FastEthernet0/0

    Protocols Configured: Address: Received: Transmitted:
        Other                0          1717982

    714653 packets, 141277256 bytes input
    1717982 packets, 133210465 bytes output

Virtual LAN ID: 2 (IEEE 802.1Q Encapsulation)

    VLAN Trunk Interface: FastEthernet0/0.2

    Protocols Configured: Address: Received: Transmitted:
        IP                192.168.2.1 170907068 327820853
        Other                0          56682

    170907312 packets, 17400604857 bytes input
    327877535 packets, 466664231516 bytes output
```

INTER - VLAN ROUTING ON A LAYER - 3 SWITCH

By using a Layer-3 switch, the packets never need to leave the switch to be routed between VLANs, eliminating a tiny lag and a possible bottleneck. SVIs (Switch Virtual Interfaces) become the default gateways for each VLAN. The configuration might look like:

```
ip routing
    Turn on L3 routing on the switch. This statement is already in router configs by default.
interface vlan 10
    ip address 10.0.10.1 255.255.255.0
interface vlan 20
    ip address 10.0.20.1 255.255.255.0
interface vlan 30
    ip address 10.0.30.1 255.255.255.0
```

You can even view the routing table, just like on a router. You'll probably need to add a default static route to a real router for any packets that need to go somewhere other than another local VLAN.

```
SW# show ip route
```

STATIC ROUTES

Static routes can be added with either an exit interface, a next-hop address, or both as the destination.

```
R1(config)# ip route 10.2.0.0 255.255.255.0 s0/0/0
    You can specify the exit interface if the recipient is unambiguous (serial lines only go to one destination).
R1(config)# ip route 10.2.1.0 255.255.255.0 10.12.0.2
    Or you can specify the next hop's address. The Cisco Academy calls this a "recursive static route" because the router will hit the routing table a second time to see which exit interface to use for that next-hop address. That term isn't used in the official cert guides.
R1(config)# ip route 10.2.2.0 255.255.255.0 s0/0/0 10.12.0.2
    To be completely unambiguous, you can specify both the exit interface and the next hop address.
```

You can limit the routing table display to just static routes:

```
R1# show ip route static
```

The router's operating system will automatically remove these routes from the routing table if the interface goes down, and put them back when it's available again. The same is true if the router doesn't have a route to the next-hop address. You can force the router to keep such routes in the routing table with the "permanent" keyword. Unfortunately, you can't force successful routing over broken links.

```
R1(config)# ip route 10.0.10.0 255.255.255.0 s0/0/0 permanent
```

A "static host route" is a static route where the mask is specific to an individual address (/32). It can exist alongside a more general route because the router will always choose the most precise match for any given packet's destination.

Bonus—Routing Table Details

From the routing table you can see how a static route was declared. The phrase "directly connected" tells us an exit interface was used with no next-hop address. Using the previous examples, we see:

```
R1# show ip route static
      <Stuff omitted>
S      10.0.0.0/8 is variably subnetted, 5 subnets, 3 masks
      10.2.0.0/24 is directly connected, Serial0/0/0
      Exit interface specified in the "ip route..." command
S      10.2.1.0/24 [1/0] via 10.12.0.2
      Next hop address specified
S      10.2.2.0/24 [1/0] via 10.12.0.2, Serial0/0/0
      Both specified. Notice that it's not "directly connected"
```

C O M P E T I N G R O U T E S

Competing Routes—the author's term for more than one route to the same destination (including mask). The router will need to choose which to put in the routing table, holding the other(s) in reserve in case the chosen one becomes unavailable.

AD (Administrative Distance)—how a router chooses between competing routes. The AD value is based on the source of the route. Lower is better.

Floating Static Route—A static route with a deliberately high AD so that it's treated as a backup for when some dynamic route (from a routing protocol like OSPF) is lost.

```
R1(config)# ip route 10.0.10.0 255.255.255.0 s0/0/0 180
      180 is the manually forced administrative distance
```

```
edge# show ip route
```

```
...
S*    0.0.0.0/0 [1/0] via 69.12.233.1, FastEthernet0/1
      The AD is 1 (fist of the two #s in the bracket)
```

```
edge# show ip route 0.0.0.0
```

```
Routing entry for 0.0.0.0/0, supernet
  Known via "static", distance 1, metric 0, candidate default path
  Routing Descriptor Blocks:
  * 69.12.233.1, via FastEthernet0/1
    Route metric is 0, traffic share count is 1
```

PROTOCOL	AD
Connected	0
Static Routes	1
EIGRP	90
OSPF	110
RIP (v1/v2)	120
DHCP Default Route	254
Unknown / Unbelievable	255

Static Default Route (AKA "gateway of last resort")—The previous two "show" commands happen to show a static default route that sends all traffic that doesn't match a more specific route out the default gateway to the internet at large. To create one, just use all zeroes for both the destination address and mask when specifying a static route. In the "show ip route" command, they're shown with both an "S" (static route) and a "*" (candidate default route).

```
edge(config)# ip route 0.0.0.0 0.0.0.0 FastEthernet0/1 192.51.100.1
```

TROUBLESHOOTING STATIC ROUTES

Good syntax, bad logic—check for subnet math errors that put your destination outside the range of the route.

Static Route Doesn't Show Up—routes won't be added to the routing table if

- Outgoing interface not up/up
- The next hop address doesn't have a route to reach *it*.
- Some other route to the same destination and mask has a lower AD

Static Route Works Poorly

- If you've got the "permanent" keyword, the route will stay in the routing table even when it's broken for other reasons.
- It may have nothing to do with routing—perhaps bad cables, etc.