



Routing

“It’s how the bits get there!”

CIS 68C2-01

UNIX Network Administration

Routing

- *What is Routing?*
 - ✗ TCP/IP's directional decision maker for packet flow
 - ✗ Routing is ultimately just passing a packet from gateway to gateway until it reaches its destination host
- *Why are packets routed?*
 - ✗ TCP/IP is a switching network
 - ✗ Packets must be switched across physical network boundaries
 - ✗ IP can only deliver packets through known networks and gateways
- *Who does routing?*
 - ✗ IP Protocol is responsible for performing packet switching
 - ✗ IP decides where to send a packet next

Routing

- *When* are packets routed?
 - ✗ Every time a packet is received by IP!
- *Where* are packets routed?
 - ✗ Up the TCP/IP stack
 - ✗ To an attached physical network
 - ✗ To the local loopback network
 - ✗ Nowhere – just dropped on the floor
- *How* are packets routed?
 - ✗ IP routes packets with the help of a table of instructions

Routing

- What is a route?
 - ✗ Simple instructions specifying how to forward a packet
 - ✗ No route ... no packet delivery!
- Routing consists of two distinct components:
 - ✗ Packet Routing
 - ✗ Forwards packets based on a routing table
 - ✗ Performed in all TCP/IP networking systems
 - ✗ Routing Protocols
 - ✗ Define the exchange of routing table information
 - ✗ Used by systems whose routing tables are updated dynamically
 - ✗ Optional

Routing

□ Routing Configurations

✗ Minimal – Single Route

- ✗ Single route created by **ifconfig**

- ✗ **ifconfig up** typically installs a single route to the local network

✗ Static Routes

- ✗ Routing table built manually by system administrator

✗ Dynamic Routes

- ✗ Routing table built dynamically by routing table exchange software

- ✗ Protocols help define and determine “best” routes

Routing

- A Packet's Journey
 - ✗ Transmitting NIC places packet on wire
 - ✗ Packet headers contain:
 - ✗ Source and destination Ethernet addresses (MAC addresses)
 - ✗ Source and destination IP addresses
 - ✗ Receiving NIC picks up packet as it passes on the wire
 - ✗ NIC passes packet up the protocol stack to the IP layer
 - ✗ IP makes routing decision based on destination IP address
 - ✗ If address matches host's, packet is passed up to Transport Layer
 - ✗ Otherwise, it is considered for *forwarding* to the next *hop*

Routing

- A Packet's Journey - *continued*
 - ✗ Forwarding occurs when...
 - ✗ The system is enabled for packet forwarding, *and*
 - ✗ The routing table contains either
 - ✗ a specific entry for a directly connected network, or
 - ✗ a specific entry for a router on a directly connected network, or
 - ✗ a default route (a place to go when no specific route exists)
 - ✗ When IP forwards a packet...
 - ✗ Original **datagram** (IP header plus payload) is intact
 - ✗ Packet is re-framed with a MAC header for transmitting interface
 - ✗ Source address is interface's MAC address
 - ✗ Destination address is next hop's MAC address

Routing

□ Building Static Routes

- ✗ Use **route** command to add/delete routes
- ✗ Linux **route** syntax is somewhat unique (and bizarre)
- ✗ Commands to view the route table:
 - ✗ **netstat -rn** and **route**
- ✗ Set default route to gateway that accesses maximum number of networks
- ✗ Routes in Linux can be added at boot
 - ✗ `/etc/sysconfig/static-routes`
 - ✗ Format: *interface args*
 - ✗ Where *args* are the arguments to **route** that follow the **add** keyword

Routing

□ Reading the Route Table

1) For each row in the route table, apply the network mask in the Genmask column to the packet's destination IP address and if the resulting network address matches the Destination address here ...

2) ... then IP can forward the packet onto the network or to the gateway listed for that row...

3) ... by transmitting the packet through this interface.

```
$ netstat -rn
Destination      Gateway          Genmask         Flags   MSS Window  irtt  Iface
10.0.0.0         0.0.0.0         255.255.255.0  U       40  0        0    eth0
11.0.0.0         0.0.0.0         255.255.255.0  U       40  0        0    eth1
127.0.0.0        0.0.0.0         255.0.0.0      U       40  0        0    lo
0.0.0.0          11.0.0.1        0.0.0.0        UG      40  0        0    eth1
```

4) Finally, if there is no previous match, then the 0.0.0.0 entry here *will always match* since applying this row's Genmask of 0.0.0.0 to any IP address results in a 0.0.0.0 network address. This entry is called the Default Route.

Routing

□ ICMP Redirects

- ✗ Router A instructing host to use Router B
- ✗ Causes new static **host route** to be added to route table
 - ✗ Only removed via reboot or administrator
 - ✗ Use is generally not recommended for security concerns
- ✗ `/proc/sys/net/ipv4/conf/*/accept_redirects`

Routing

- Additional Information

- ✗ Linux Networking HOWTO

- ✗ Routing Section

- ✗ <http://www.linux.org/docs/ldp/howto/Net-HOWTO/x552.html>