iproute2 and Advanced Linux Routing

Presenter William L. Thomson Jr.

What is iproute2

- A collection of utilities for controlling TCP/IP networking and traffic control in Linux
- Usually shipped in a package called iproute or iproute2 and conists of several tools, mainly ip and tc.
- ip controls IPv4 and IPv6 configuration
  - replaces arp, ifconfig, and route commands
- tc stands for traffic control (not discussed or covered)
- Available in most distributions
- Requires IP features to be enabled in kernel
  - Networking Options CONFIG_IP_*
  - Networking -> Networking Options -> IP:*

Why iproute2

- arp, ifconfig, and route commands are venerable
  - One command does it all, ip
- Consistent behaviour under Linux 2.2 and up
- Uses the redesigned network subsystem in 2.2 and above Linux Kernels
- Ability to do GRE tunnels
  - IP in IP tunneling
  - Transport multicast traffic and IPv6 through a GRE tunnel
- Robust features
  - Routing, filtering, and classifying
  - Rivals dedicated routers, firewalls, and traffic shaping products

Interface Management

- Displaying links, address, routes, and arp
  - ip link list
  - ip address show
  - ip route show
  - ip neighbour/neigh show
- Management
  - ip link set eth0 up/down
  - ip addr add/delete 192.168.0.161/27 brd 192.168.0.191 dev eth0
  - ip route add/change/delete/replace default via 192.168.0.192 dev eth0
  - ip neighbour add/change/delete/replace 192.168.0.1 lladdr 00:00:00:00:00:00 dev eth0 nud permanent
  - ip link/addr/route/neigh flush

Policy Routing

- Policy routing consists of rules and tables
- Rules are defined to match patterns, ip, interfaces, combo
- Rules are given priorities and are followed in order till a pattern is matched
- Rules point to tables that contain routes themselves
- Required for machines with multiple networks and gateways

Rules

- Default rules (be careful with flushing or modifying)

  # ip rule

  0: from all lookup local

  32766: from all lookup main

  32767: from all lookup default

- Custom rules

  ip rule add prio 5 table main

  ip rule add prio 10 from 192.168.0.0/24 table 10

  ip rule add prio 11 from 192.168.1.0/24 table 11

  ip rule add prio 12 from 192.168.2.0/24 table 12

  ip rule add prio 30 to 68.87.68.30 table 30

Tables

- Default tables (be careful with flushing or modifying)
  - local, main, default (ip route show table {local,main,default})
- Custom tables

  ip route del default table main

  ip route add default via 10.0.0.1 dev eth0 table proto static 10

  ip route add default via 10.0.1.1 dev eth1 table proto static 11

  ip route add default via 10.0.2.1 dev eth2 table proto static 12

  ip route add vid dev eth3 proto static table 30

Putting it all together, live example


```
ip link set lo up
ip link set eth0 up
ip link set eth1 up
ip link set eth2 up

ip addr add 127.0.0.1/8 brd 127.0.0.255 dev lo
ip addr add 192.168.1.250/24 brd 192.168.1.255 dev eth0
ip addr add 10.1.0.2/16 brd 10.1.255.255 dev eth1
ip addr add 10.1.0.97/16 brd 10.1.255.255 dev eth1
ip addr add 10.2.0.2/16 brd 10.2.255.255 dev eth2
ip addr add 10.2.0.57/16 brd 10.2.255.255 dev eth2

ip route add 127.0.0.0/8 dev lo

ip rule add prio 10 table main
ip route del default table main

ip rule add prio 20 from 10.1.0.0/16 table 20
ip route add default via 10.1.0.1 dev eth1 src 10.1.0.2 proto static table 20
ip route append prohibit default table 20 metric 1 proto static

ip rule add prio 30 from 10.2.0.0/16 table 30
ip route add default via 10.2.0.1 dev eth2 src 10.2.0.2 proto static table 30
ip route append prohibit default table 30 metric 1 proto static

# Set up load balancing gateways
ip rule add prio 50 table 50
ip route add default table 50 proto static \
        nexthop via 10.1.0.1 dev eth1 \
        nexthop via 10.2.0.1 dev eth2
```

Where to go from here

- help argument, ip help, ip route help, ip rule help
- Man page, man ip
- Distribution specific documentation
- General documentation and resources for this presentation
  - http://www.lartc.org/howto/
  - http://www.policyrouting.org/iproute2-toc.html
  - http://www.docum.org/docum.org/faq/cache/57.html
- Everyone's friend Google
- Local Users Group
- Ask the presenter via email
  - William L. Thomson Jr. wlt@obsidian-studios.com