What Linux can do for you

- Throttle bandwidth TO certain computers
- Help you to fairly share your bandwidth
- Protect your network from DoS attacks
- Protect the Internet from your customers
- Multiplex several servers as one, for load balancing or enhanced availability
- Restrict access to your computers
- Limit access of your users to other hosts
- Do routing based on user id (yes!), MAC address, source IP address, port, type of service, time of day, or content

iproute2

- Linux has a sophisticated system for bandwidth provisioning called Traffic Control.
- This system supports various method for classifying, prioritizing, sharing, and limiting both inbound and outbound traffic.

Requirement

- Any Linux distribution such as RedHat or Debian
- Iproute2 package

source policy routing

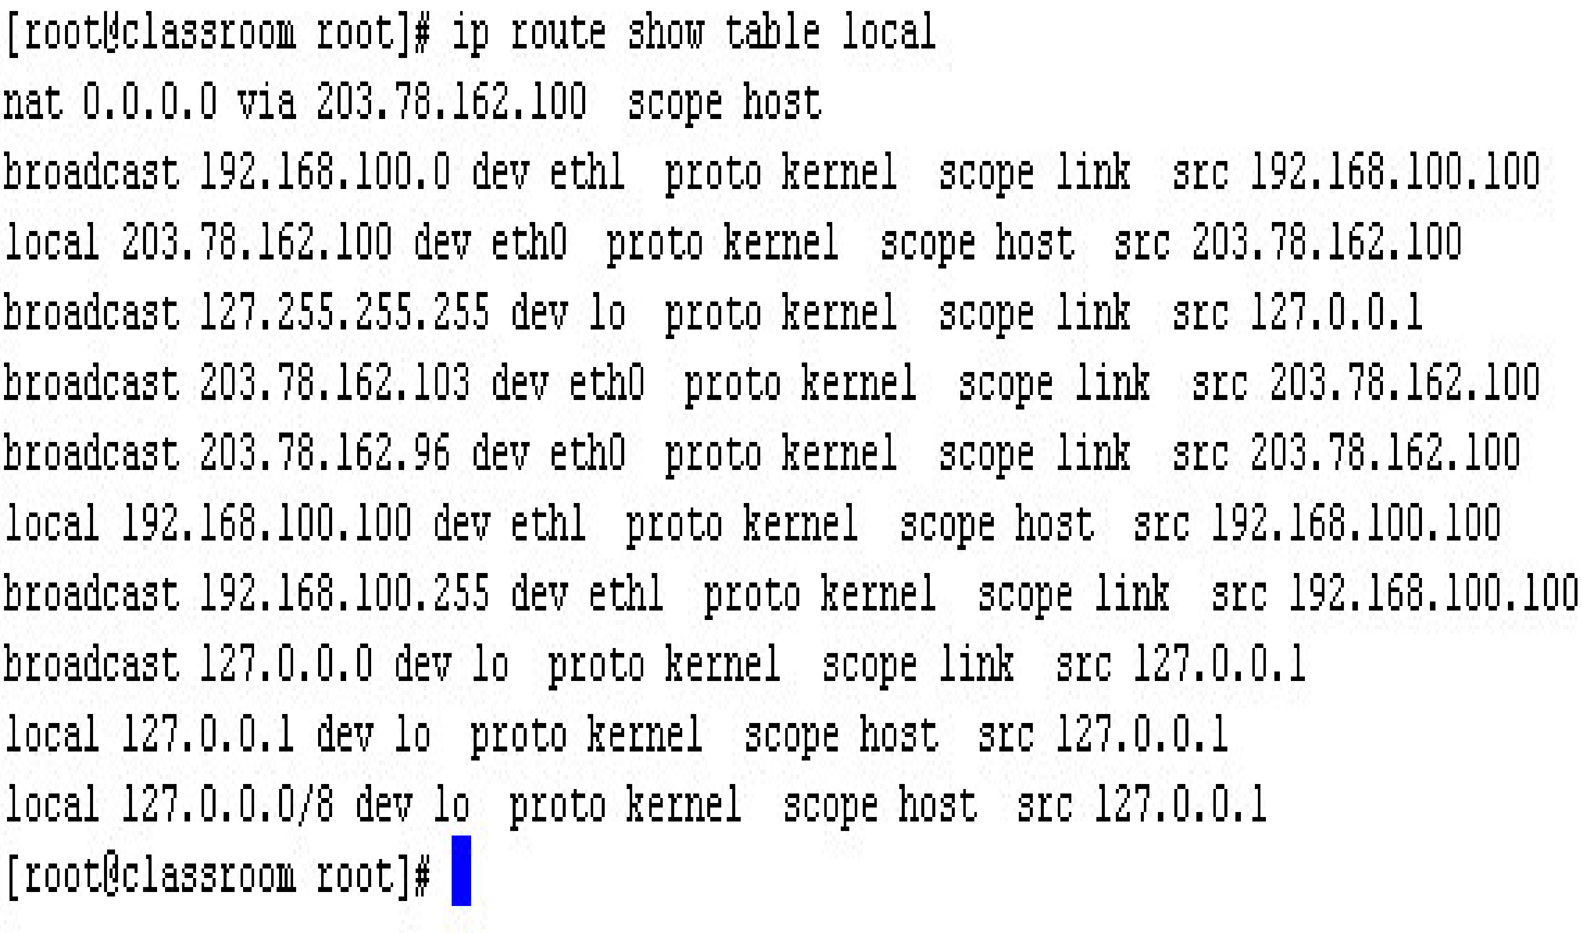

```
[root@classroom root]# ip route show table local
nat 0.0.0.0 via 203.78.162.100  scope host
broadcast 192.168.100.0 dev eth1  proto kernel  scope link  src 192.168.100.100
local 203.78.162.100 dev eth0  proto kernel  scope host  src 203.78.162.100
broadcast 127.255.255.255 dev lo  proto kernel  scope link  src 127.0.0.1
broadcast 203.78.162.103 dev eth0  proto kernel  scope link  src 203.78.162.100
broadcast 203.78.162.96 dev eth0  proto kernel  scope link  src 203.78.162.100
local 192.168.100.100 dev eth1  proto kernel  scope host  src 192.168.100.100
broadcast 192.168.100.255 dev eth1  proto kernel  scope link  src 192.168.100.100
broadcast 127.0.0.0 dev lo  proto kernel  scope link  src 127.0.0.1
local 127.0.0.1 dev lo  proto kernel  scope host  src 127.0.0.1
local 127.0.0.0/8 dev lo  proto kernel  scope host  src 127.0.0.1
[root@classroom root]#
```

Routing for multiple providers

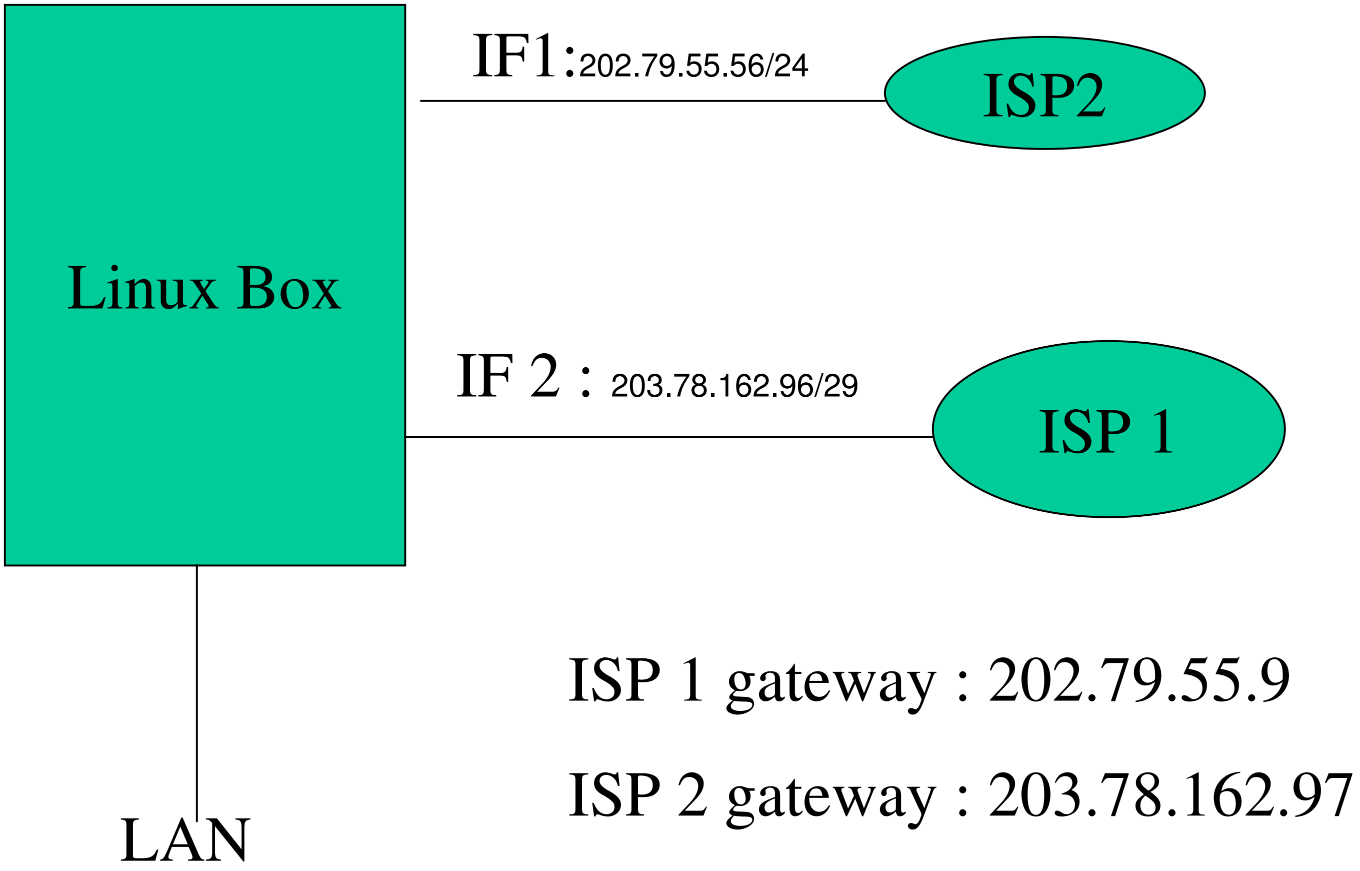

Linux Box

IF1:202.79.55.56/24

ISP2

IF 2 : 203.78.162.96/29

ISP 1

LAN

ISP 1 gateway : 202.79.55.9

ISP 2 gateway : 203.78.162.97

how to route answers to packets coming in over a particular provider, say ISP 1, back out again over that same ISP.

- creates two additional routing tables, say **ISP1** and **ISP2**.

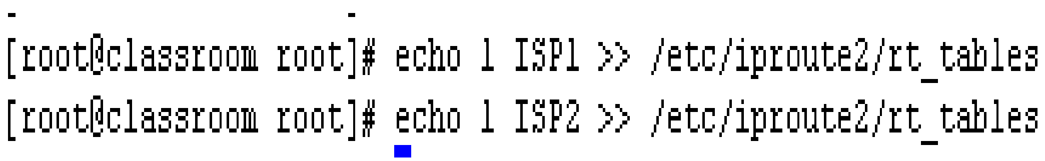

```
[root@classroom root]# echo 1 ISP1 >> /etc/iproute2/rt_tables
[root@classroom root]# echo 1 ISP2 >> /etc/iproute2/rt_tables
```

Then you set up routing in these tables as follows:

```
ip route add 202.79.55.0 dev eth1 src 202.79.55.56 table ISP1
ip route add default via 202.79.55.9 table ISP1
ip route add 203.78.162.96 dev eth2 src 203.78.162.100 table ISP2
ip route add default via 203.78.162.97 table ISP2
```

set up the main routing table.

```
ip route add 202.79.55.0 dev eth1 src 202.79.55.56
ip route add 203.78.162.96 dev eth2 src 203.78.162.100
```

preference for default route:

```
ip route add default via 203.78.162.97
```

set up the routing rules.

```
ip rule add from 202.79.55.0 table ISP1
ip rule add from 203.78.162.100 table ISP2
```

Load balancing

- Instead of choosing one of the two providers as your default route, you now set up the default route to be a multipath route. In the default kernel this will balance routes over the two providers.

- `ip route add default scope global nexthop via 203.78.162.97 dev eth2 weight 1`
- `nexthop via 202.79.55.9 dev eth1 weight 1`
- balancing will not be perfect, as it is route based, and routes are cached. This means that routes to often-used sites will always be over the same provider.