

Linux Networking: IP internetworks

David Morgan

Routing: local perspective vs global perspective

- it's **local** as viewed by
 - any individual machine in the path
- it's **global** as viewed by
 - the collective of machines
 - a network administrator

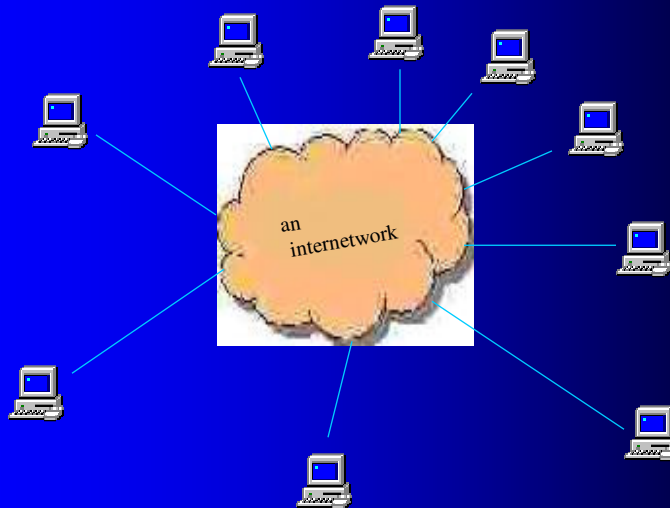
Local view

- the “boundary” is the local machine
- the “destination” is one of the machine’s interfaces
- path to IP drives choice of interface


Global view

- the “boundary” is the internetwork
- the “destination” is the other machine
- interface choice drives path to IP


An Internetwork:
looks like a network...



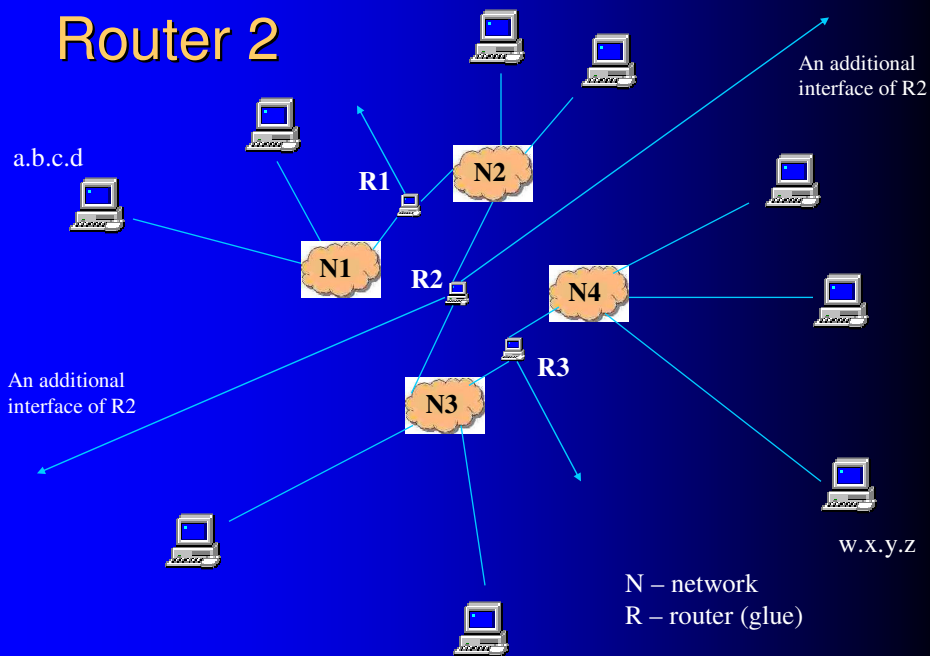
...but it's a bunch of 'em...



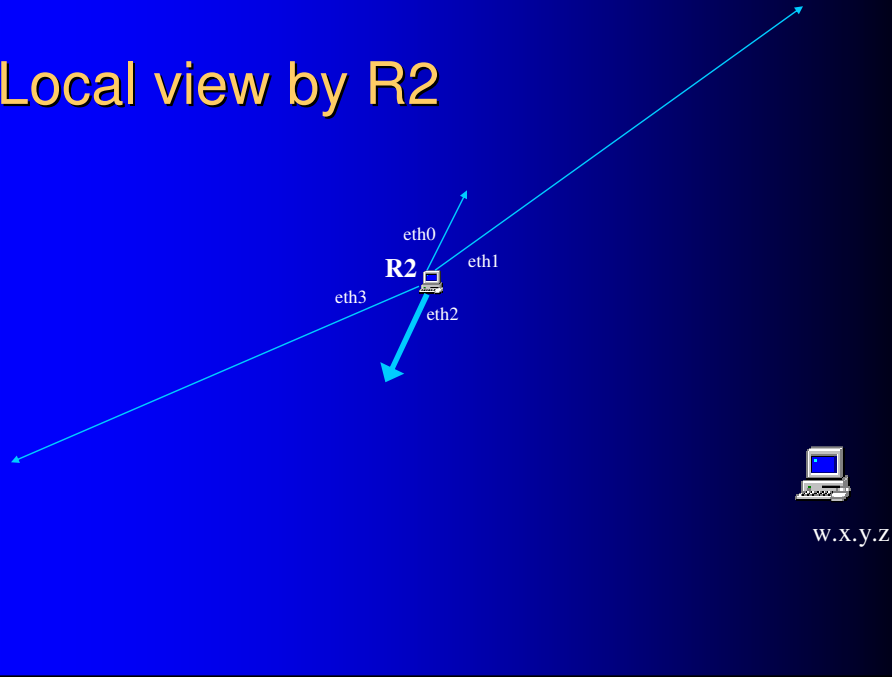
...with glue (routers).



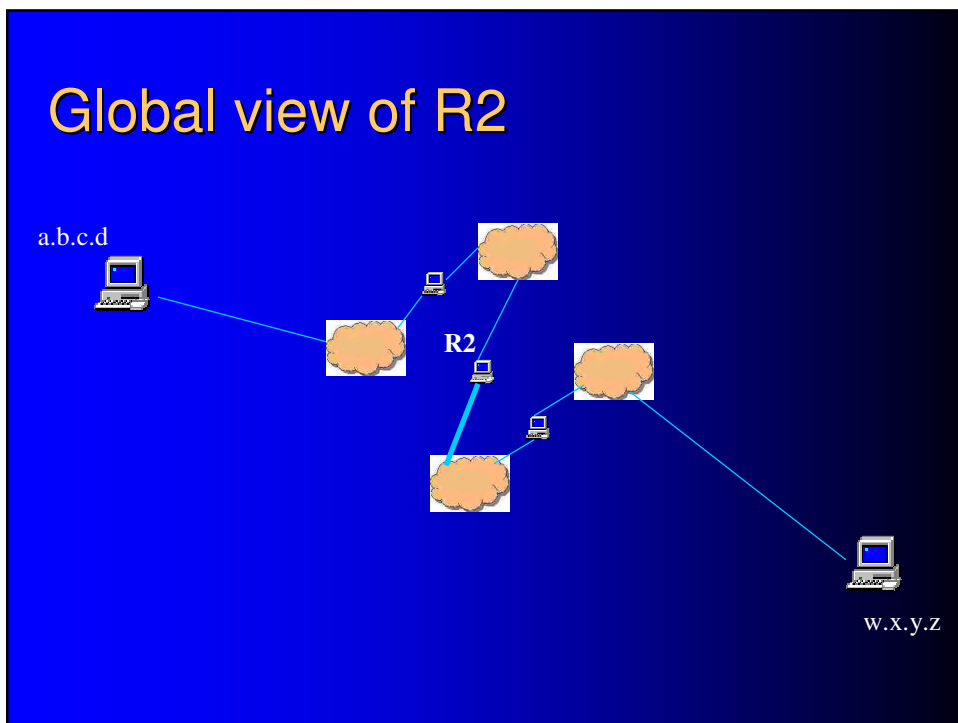
Router 2




Local view by R2



Global view of R2



Routers are key



Routers are key



Routers' route tables

– 2 types of entries



| Destination | Next hop (gateway entry) | Direct delivery (interface entry) |
|-------------|-----------------------------|--------------------------------------|
| N1 | | |
| N2 | | |
| N3 | | |
| N4 | | |


“Direct delivery” type – linux syntax

- host route - to a single machine
 - route add ~~–host~~ 192.168.4.2 eth0
- network route, local - to a group of machines
 - route add ~~–net~~ 192.168.4.0 netmask 255.255.255.0 eth0
- network route, thru gateway - to a group of machines
 - route add ~~–net~~ 192.168.5.0 netmask 255.255.255.0 gw 192.168.4.1
- default route - to “any and all” else
 - route add default gw 192.168.4.1


“Next hop” type – linux syntax

- host route - to a single machine
 - route add –host 192.168.4.2 eth0
- network route, local - to a group of machines
 - route add –net 192.168.4.0 netmask 255.255.255.0 eth0
- network route, thru gateway - to a group of machines
 - route add –net 192.168.5.0 netmask 255.255.255.0 gw 192.168.4.1
- default route - to “any and all” else
 - route add default gw 192.168.4.1

Network and host addresses




Routers' route tables




| Destination | Next hop (gateway entry) | Direct delivery (interface entry) |
|-------------|-----------------------------|--------------------------------------|
| N1 | | |
| N2 | | |
| N3 | | |
| N4 | | |

R1's route table




| Destination | Next hop (gateway entry) | Direct delivery (interface entry) |
|-------------|-----------------------------|--------------------------------------|
| N1 | | eth0 |
| N2 | | eth1 |
| N3 | 192.168.3.2 | |
| N4 | 192.168.3.2 | |

R2's route table



| Destination | Next hop (gateway entry) | Direct delivery (interface entry) |
|-------------|-----------------------------|--------------------------------------|
| N1 | 192.168.3.1 | |
| N2 | | eth0 |
| N3 | | eth1 |
| N4 | 192.168.4.2 | |

R3's route table



| Destination | Next hop (gateway entry) | Direct delivery (interface entry) |
|-------------|-----------------------------|--------------------------------------|
| N1 | 192.168.4.1 | |
| N2 | 192.168.4.1 | |
| N3 | | eth0 |
| N4 | | eth1 |

Addresses put into frames/packets

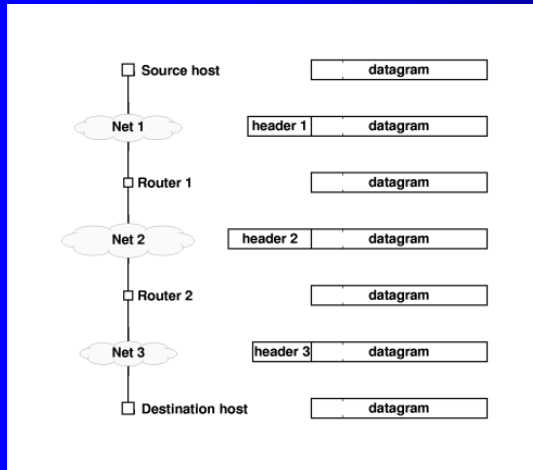
| | Direct delivery (target in same subnet) | Next hop (target in different subnet) |
|--------------|---|---------------------------------------|
| IP address: | target's | target's |
| MAC address: | target's | router's |

Because...

To transmit to a host that is in the **same** Broadcast Domain, IP software will ARP **directly** for the **destination host** IP address.

To transmit to a host that is in a **different** Broadcast Domain, IP software will ARP for the **default router** IP address.

One frame header at a time



- frame headers do not accumulate
- removed at each router, where then new header affixed
- networks of heterogeneous types (ethernet, point-to-point, frame relay, ATM,...) headers differ

Figure: Douglas Comer textbook

Now they look like a network...

