3 slide aside: details about IP addresses

Evolution of IPv4 addresses

Originally: Each address was 32b and of a certain class:

Class A: 1b type, 7b network, 24b host	`0nnnnnnnhhhhhhhhhhhhhhhhhhhhhhh`
Class B: 2b type, 14b network, 16b host	`10nnnnnnnnnnnnnnhhhhhhhhhhhhhhhh`
Class C: 3b type, 21b network, 8b host	`110nnnnnnnnnnnnnnnnnnnnhhhhhhhh`

Classification is simple, e.g. 3 tables, one for each class

Problems:

- × Free addresses become scarce as Internet becomes popular => larger (128b) addresses for IPv6
- × Poor utilisation: Organisation *O* with 2^8+1 hosts on one network needs Class B space and uses only $1/256^{th}$
  - Could use 2 separate Class C networks, but expands router database [13W]…
- × Router tables become large (2M Class C nets)
  - Router chatter (to maintain tables) uses appreciable transmission capacity

Solution: Classless InterDomain Routing (CIDR)

Classless InterDomain Routing

Allow network prefixes of arbitrary length

Different way of indicating which bits identify network, and which bits identify host.

- No longer indicated by address class
- Instead, indicated by a separate network mask (=1 in network ID bits)

could have been class B => 16b host

Address: 10010101.10101011.00100100.00110000 (149.171.36.48)

Mask: 11111111 11111111 11111110 00000000 used with 9b for host

- Usually the mask covers contiguous bits, starting with the first bit, so it can be defined by a prefix length (23 in this case)

=> represent addresses as dotted.decimal/prefix_length e.g.

- 149.171.36.48/23 = 149.171.36 (net) + 0.48 (host)
- 149.171.37.48/23 = 149.171.36 (net) + 1.48 (host)

$37_{10} = 00100101_2$

bit 23 bit 24

Prefix length => how many trailing 0s identify network, e.g.

192.16.32.0/20=192.16.0010hhhh.hh... (h=bit identifying host),
192.16.32.0/22=192.16.001000hh.hh...

- Routing protocols exchange network addresses + prefix lengths

e.g. assign one organisation ("*O*"):

- two "Class C" spaces (192.16.32 & 192.16.33)
- but only one router entry of 192.16.32/23

$192_{10}.16_{10}.00100000_2$

$192_{10}.16_{10}.00100001_2$

23 bits

Address aggregation

Router has moderate number of ports (e.g. 4-64) & different networks may be reachable through same port

e.g. Organisations

O with network 192.16.32/23 through port P, and

Q with network 192.16.34/23 through port P

=> forward aggregate 192.16.32/22 through port P

Aggregation reduces number of router table entries.

Ideally assign topologically-adjacent organisations numerically-adjacent addresses; e.g. Asia-Pacific: APNIC

Organisations may move (e.g. change network service provider) => exceptions

e.g. 192.16.32/19 through port *P*, except 192.16.34/23 through port *Q*

Search for **longest prefix match** (rules with longer prefixes have priority)

… end of aside.