Extending IP

The University Of Alabama in Huntsville Computer Science

Internet Structure

cs570

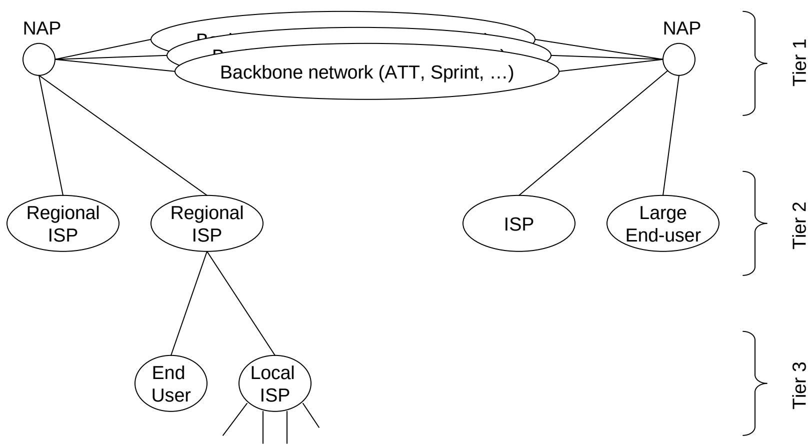

IP must accommodate a global-scale internet

Scaling problems

cs570

The University Of Alabama in Huntsville Computer Science

- Huge number of physical networks and nodes
  - IP v4 can address only about 2M networks (and most of those must be small)
  - Historically, IPv4 address space has not been efficiently used

G. W. Cox -- Fall 2007

IP part B - 3

cs570

The University Of Alabama in Huntsville Computer Science

Autonomous Systems

G. W. Cox -- Fall 2007

IP part B - 4

Problems with our ideas of a homogenous Internet

The University Of Alabama in Huntsville Computer Science

cs570

- The scale of the Internet makes it prohibitive to do routing in the ways we've discussed
  - Example: LS flooding approach
- Individual organizations want to run and administer their internal networks in any way they wish, and still connect to "outside" networks

G. W. Cox -- Fall 2007

IP part B - 5

Autonomous Systems (AS)

The University Of Alabama in Huntsville Computer Science

cs570

- aka "Routing Domains"
- Large networks are divided into AS's, usually along administrative boundaries

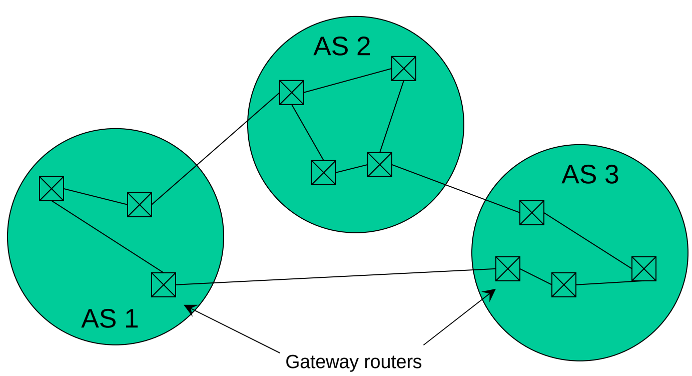

G. W. Cox -- Fall 2007

IP part B - 6

AS concepts

- Within an AS, all routers run the same routing protocol and share the same information
- Gateway routers ("Boundary routers"):
  - Run their local AS's protocol and,
  - Run a separate routing protocol to route between the Gateway routers ("Inter-AS Routing protocol")

The University Of Alabama in Huntsville Computer Science

Inter-AS routing

cs570

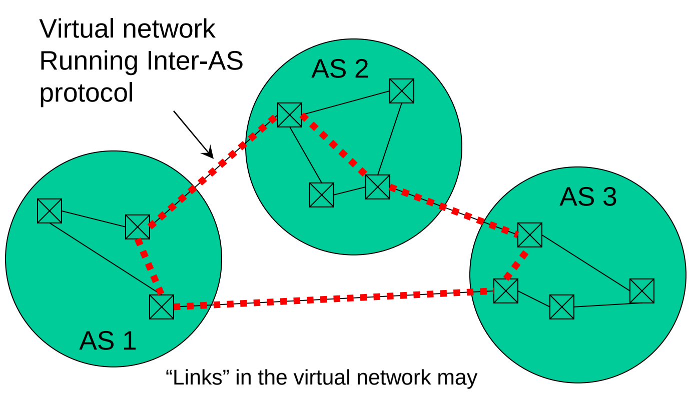

"Links" in the virtual network may correspond to physical links or to multi-hop paths

The University Of Alabama in Huntsville Computer Science

How routing packets flow over the virtual Inter-AS network

cs570

Inside AS1:

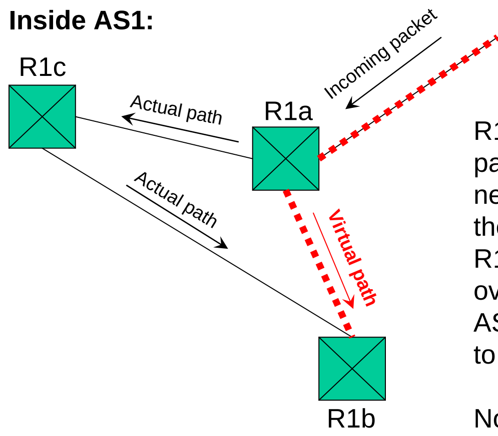

R1a receives a routing packet on the Inter-AS net. The next node on the virtual net is R1b. R1a hands the packet over to the AS1 Intra-AS protocols to send it to R1b.

Note that routers outside of AS1 require no knowledge of AS1's internals

The University Of Alabama in Huntsville Computer Science

Keeping track of AS's

cs570

- Each AS is identified by a globally-unique AS number assigned by ICANN (RFC 1930*)

- Internet Engineering Task Force (IETF) Request for Comments http://www.ietf.org/rfc.html

This helps, but…

cs570

The University Of Alabama in Huntsville Computer Science

- The AS concept helps by breaking the Internet up into manageable-size pieces
- But it does not help make efficient use of the address space

Why is there an efficiency problem?

cs570

The University Of Alabama in Huntsville Computer Science

You own a growing small business and you need a network for about 220 computers. You tell your net admin to get a network address space. He gets 220.126.44.x.

Do you thank him or fire him?

Subnetting

cs570

The University Of Alabama in Huntsville Computer Science

- The idea:
  - Introduce a new level of hierarchy by using part of the Class B Node field as a “subnet” number
  - This lets us share a single Network # across several physical nets

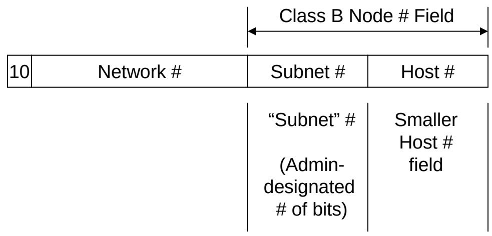

Addressing with Subnetting

cs570

The University Of Alabama in Huntsville Computer Science

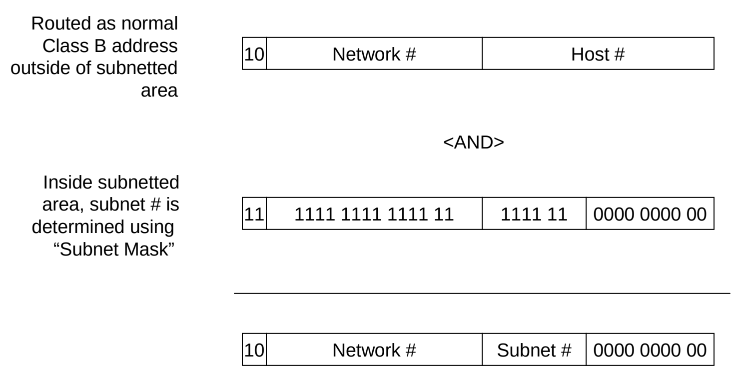

Note: Host number is obtained by ANDing address with Complemented Subnet Mask

Example subnetting

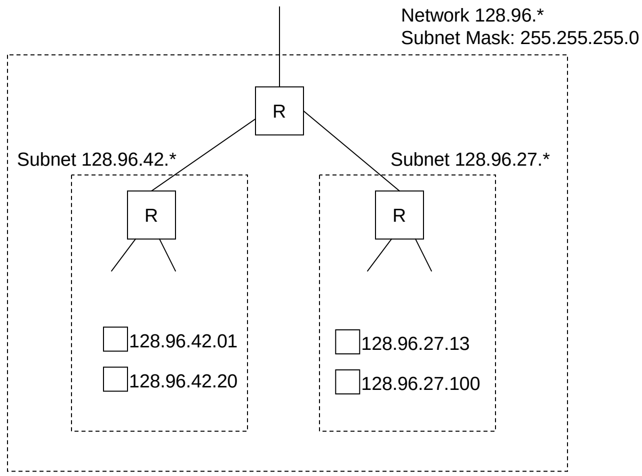

Some notes on subnetting

- Subnets will usually be physically near to each other, since all their traffic will be routed to the same router
- Subnet masks are often described by the number of 1's (e.g, 128.96.*/24)
- Subnet masks don't necessarily have contiguous 1's, but anything else is confusing

The University Of Alabama in Huntsville Computer Science

Another approach to extending IPv4

cs570

- Subnetting subdivides Class B address spaces to form subnets that lie in between Class C and Class B in the hierarchy
- Another way we could achieve the same end is to combine contiguous Class C address spaces

The University Of Alabama in Huntsville Computer Science

Classless addressing ("Supernetting")

cs570

Classless
Inter-
Domain
Routing

CIDR example

cs570

- You have a network with 16*254 hosts.
- To conserve Class B space, assign 16 contiguous Class C networks (e.g, 192.4.6.* -> 192.4.21.*)
- Some number of the high-order bits will be the same (for the example, all addresses start with 1100 0000 0000 0100 000 – 19 bits are the same)
- We can think of this as a new type of network with a 19-bit network number

High-order 19 bits	Low-order 13 bits
1100 0000 0000 0100 000	Node # (or subnetting)

- Anywhere from 4 to 30 bits could be used as determined by the number of Class C's that are combined

The University Of Alabama in Huntsville Computer Science

A complication with CIDR

cs570

- Since network numbers do not occupy fixed fields, backbone routers must be able to interpret the CIDR encoding
- This can be confusing since we may have the same high-order bits for two different CIDR nets:
  - 171.69/16 and 171.69.10/24 can be two different networks
- For routing, use the "longest match" principle – choose the network that matches the most high-order bits of the IP address:
  - 171.69.10.5 -> 171.69.10/24
  - 171.69.20.5 -> 171.69/16

The University Of Alabama in Huntsville Computer Science

The University Of Alabama in Huntsville Computer Science

Another type of approach: Network Address Translation (NAT)

cs570

- The idea:
  - Conserve address space by assigning each organization just a few IP addresses.
  - Internal to the organization, every PC has a unique address in the ranges:
    - 10.0.0.0 to 10.255.255.255/8
    - 172.16.0.0 to 172.31.255.255/12
    - 192.168.0.0 to 192.168.255.255/16
  - When an IP packet exits the organization, it passes through a NAT Box that translates it to one of the “real” IP addresses.
- Tannenbaum says NAT is an “ugly hack.” But many give it partial credit for extending the life of IPv4.

The University Of Alabama in Huntsville Computer Science

cs570

A long-term solution: IPv6

The University Of Alabama in Huntsville Computer Science

IPv6 ("IPng")

cs570

- Intended to be a long-term fix
- Goals:
  1. Extend the address space
  2. Some additional feature (QoS support, Security support…)
  3. Improved performance
- Characteristics:
  – 128-bit addresses (3.4 x 10^{38} hosts max)
  – Classless addressing
  – Less complicated packet format than v4 (7 vs 13 header fields)

The University Of Alabama in Huntsville Computer Science

IPv6 packet header format

cs570

<table>
<tr><td>IP version</td><td>Traffic class</td><td colspan="2">Flow label</td></tr>
<tr><td colspan="2">Payload length</td><td>Next header</td><td>Hop limit</td></tr>
<tr><td colspan="4">Source address</td></tr>
<tr><td colspan="4">Source address</td></tr>
<tr><td colspan="4">Source address</td></tr>
<tr><td colspan="4">Source address</td></tr>
<tr><td colspan="4">Destination address</td></tr>
<tr><td colspan="4">Destination address</td></tr>
<tr><td colspan="4">Destination address</td></tr>
<tr><td colspan="4">Destination address</td></tr>
<tr><td colspan="4">Extension header(s) -- optional</td></tr>
</table>

Changes in header fields

cs570

- New
  - Traffic Class – ID's special delivery req'ts (e.g, real time delivery)
  - Flow label – supports circuit-oriented channels
  - Payload length -- # data bytes (not including header as in v4)
  - Next header – ID's extension header type, if any or Transport protocol (TCP or UDP)
  - Extension headers – various optional header fields to modify basic format (e.g, over-length payloads, authentication,…)
- Gone
  - IHL – basic and extension headers are now fixed-length
  - Protocol – function replaced by Next Header
  - Fragmentation support fields – fragmentation handled differently
  - Checksum – to improve performance (redundant with L2, L4)

Fragmentation in IPv6

cs570

- IPv6 senders are expected to know the max packet size for their path
- Oversize packets are treated as errors – error message returned to sender specifying max size for future transmissions

The University Of Alabama in Huntsville Computer Science

The state of IPv6

cs570

- Hasn't caught on as fast as expected
  - Huge cost of changing hardware
  - Success of IPv4 extension measures
- Many IPv6 "islands" now in operation on the Internet
  - Communicate by tunneling through IPv4
- Eventually, islands will merge, "take over"

The University Of Alabama in Huntsville Computer Science

Shorthand for IPv6 addresses

cs570

- No longer using dotted decimal. Changed to hex format:
  - Ex: 8000:0000:0000:0000:0123:4567:89AB:CDEF
- Simplifications:
  - Omit leading zeros in a group ('0123' -> '123')
  - Replace zero groups with pair of colons when not ambiguous
  - Ex: 8000::123:4567:89AB:CDEF
- IPv4 addresses written as pair of colons followed by dotted decimal
  - ::192.31.20.46