

Network Protocols

Internet Protocol (IP) Addresses (mostly about IPv4, IPv6 is later)

Why have addresses at layer 3?

Aren't there already globally unique addresses at L2?
If not, should there be?

IP address semantics

- Virtual, not specific to a hardware device
- 32-bit fixed address length (IPv4)
- Unique address for each interface (typically)
- Registry or upstream ISP assigns net bits (prefix)
- Local admin assigns subnet and host bits (suffix)
- Usually written in dotted decimal (dotted quad)
 - e.g. 140.192.5.1

IP address notation

*diagram courtesy of <http://www.netbook.cs.purdue.edu>

<u>32-bit Binary Number</u>	<u>Equivalent Dotted Decimal</u>
10000001 00110100 00000110 00000000	129 . 52 . 6 . 0
11000000 00000101 00110000 00000011	192 . 5 . 48 . 3
00001010 00000010 00000000 00100101	10 . 2 . 0 . 37
10000000 00001010 00000010 00000011	128 . 10 . 2 . 3
10000000 10000000 11111111 00000000	128 . 128 . 255 . 0

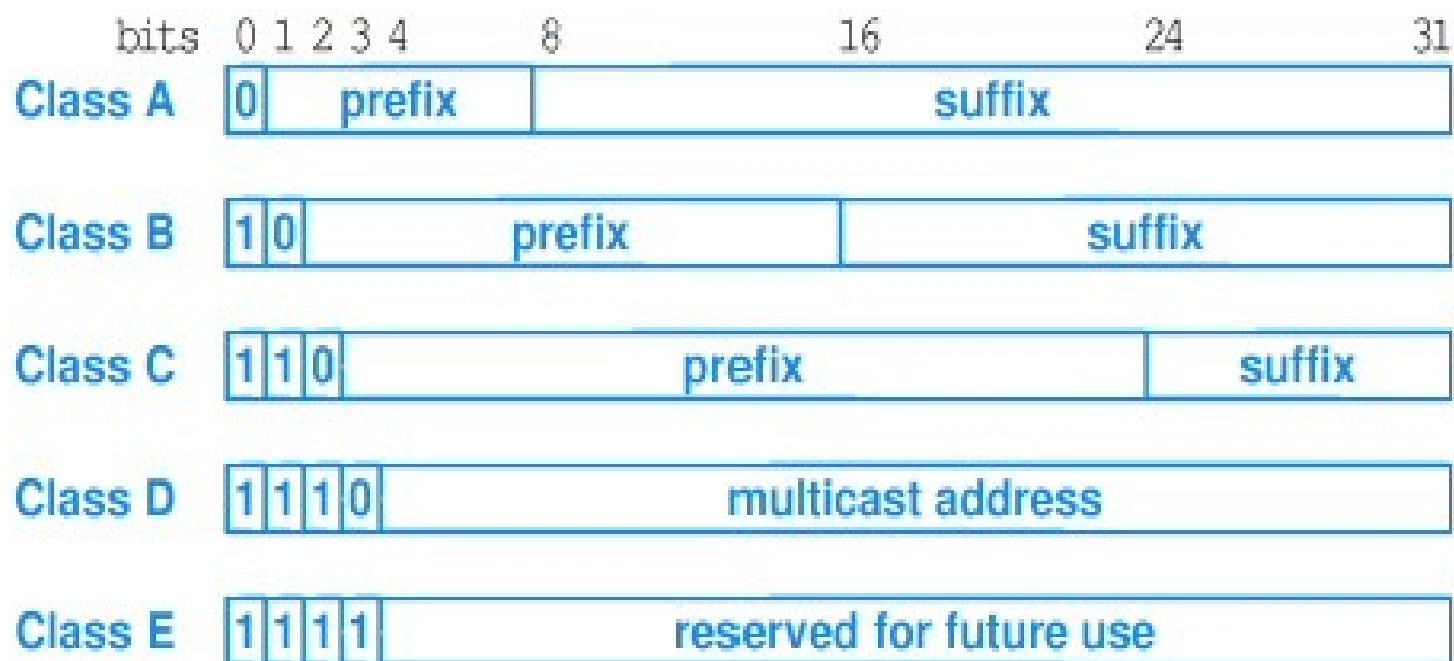
Special IP addresses

*diagram courtesy of <http://www.netbook.cs.purdue.edu>

Prefix	Suffix	Type Of Address	Purpose
all-0s	all-0s	this computer	used during bootstrap
network	all-0s	network	identifies a network
network	all-1s	directed broadcast	broadcast on specified net
all-1s	all-1s	limited broadcast	broadcast on local net
127	any	loopback	testing

Classful IP addressing

*diagram courtesy of <http://www.netbook.cs.purdue.edu>



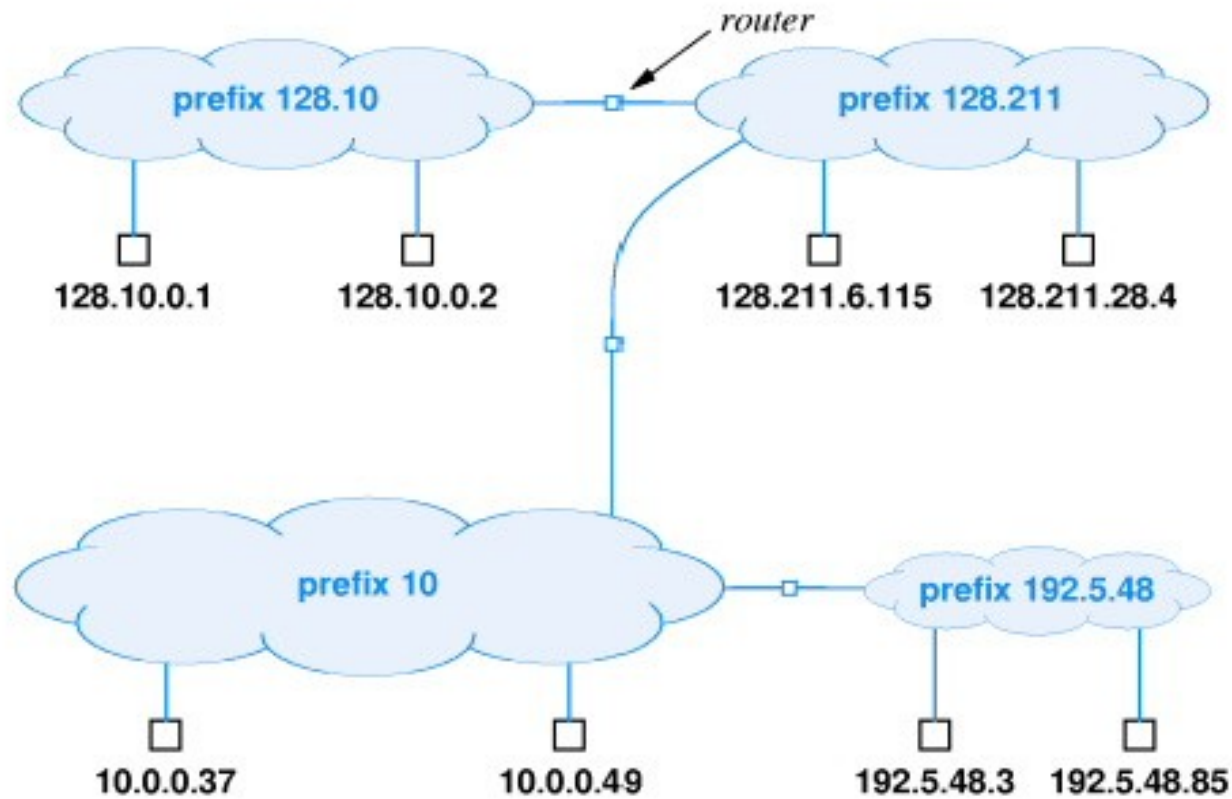
Classful address sizes

*diagram courtesy of <http://www.netbook.cs.purdue.edu>

Address Class	Bits In Prefix	Maximum Number of Networks	Bits In Suffix	Maximum Number Of Hosts Per Network
A	7	128	24	16777216
B	14	16384	16	65536
C	21	2097152	8	256

Example IP network

*diagram courtesy of <http://www.netbook.cs.purdue.edu>





Classful addressing limitations

- Internet growth and address depletion
- Route table size (potentially lots of class C nets)
- Misappropriation of addresses
- Lack of support for varying sized networks
 - Class B is often too big, Class C often too small

IP addressing solutions

- Subnetting
- Supernetting
- Classless interdomain routing (CIDR)
- Variable length subnet masks (VLSM)
- Temporary addresses (e.g. BOOTP, DHCP)
- NATs with port address translation (blech!)

Subnetting



Subnet masks

- The bit length of the prefix or the 'network' bits
- No more A, B or C class addresses ← important
- Use of the slash '/' notation to represent addresses
 - 140.192.5.1 with mask of 255.255.255.128 is:
 - 140.192.5.1/25
- A /25 mask in binary is:
 - 11111111.11111111.11111111.10000000

Subnet masks example

- Given 140.192.50.8/20 what is the...
 - subnet mask in dotted decimal notation?
 - directed broadcast address in dotted quad?
 - total number of hosts that can be addressed?

Supernetting

- Combine smaller address blocks into an aggregate
- If class B is too big and class C is too small...
 - Combine 199.63.0.0/24 to 199.63.15.0/24
 - To form 199.63.0.0/20

CIDR

classless inter-domain routing

- Routers “announce” prefixes
 - Maintain and announce millions of /24's or...?
- Aggregate
 - thank you, supernetting
- So instead of adding multiple class C blocks...
- Check out The Internet CIDR report:
 - <http://www.cidr-report.org>

CIDR example

- Given an ISP that announces:
 - 64.5.0.0/20
 - 64.5.16.0/20
 - 192.0.2.0/25
 - 192.0.2.192/26
 - 192.0.2.128/26
- What is the least number of CIDR announcements that can be made for this ISP?
- Why might address blocks not be aggregated?

VLSM

variable length subnet masks

- Many subnet sizes in an autonomous system (AS)
- Allows for efficient use of address space
- Can be used to build an internal hierarchy
- External view of the AS does not change
- An AS may be allocated 140.192.0.0/16, but...
- internally may use:
 - 140.192.0.0/17
 - 140.192.128.0/24
 - 140.192.129.0/25 and so on...

VLSM example

- Given an assignment of 140.192.0.0/16, create an addressing strategy to support:
 - 6 satellite offices and 1 large headquarter site
 - 6000 total hosts on all combined networks
 - headquarters needs about 50% of all addresses
 - satellite offices need 200 to 700 addresses
 - overall growth per year is 500 hosts

Obtaining IP addresses

- IANA has global authority for assignment
- RIRs delegate to ISPs and large nets
- ISPs assign addresses to end users and small nets
- RFC 1918 defines private address blocks
 - NOT globally unique
 - NOT for hosts attached directly to public Internet
 - 10.0.0.0/8, 172.16.0.0/12 and 192.168.0.0/16

IP address types

- Unicast (one-to-one)
 - source addresses should always be unicast
- Multicast (one-to-many)
 - receivers join/listen to group destination address
- Broadcast (one-to-all)
 - special case of multicast, usually unnecessary
- Anycast (one-to-one-of-many)
 - usually one-to-nearest, often used for reliability

NAT

network address translation

- Originally a solution to a shortage problem
 - Became a security monkey's best practice
- It can help alleviate renumbering problems
 - But if the net is that big, get your own allocation
- NAT back slaps the e2e argument in the face
- What has been a big motivating factor for Ipv6?
 - Is NAT for IPv6 for rational?
 - IETF, try as you might to “standardize” it away...

Final thoughts

- IP addressing is a pain, wait til you see IPv6
- IP addresses today are both a who and a where
- IP addresses make for poor trust relationships
- Private addresses and NATs, blech! Get a real net