

Subnet Size



128	64	32	16	8	4	2	1	128	64	32	16	8	4	2	1	128	64	32	16	8	4	2	1	128	64	32	16	8	4	2	1
1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	

- The bigger the host portion of the network, the more hosts we can have
- If the subnet mask is /8, we have 24 bits available to allocate to hosts
- If the subnet mask is /24, we only have 8 bits available to allocate to hosts

How Internet Addressing Was Meant to Work

- The global coordination of Internet IPv4 addressing is performed by IANA (Internet Assigned Numbers Authority).
- This is the way it was originally supposed to work:
- When a company wants to communicate on the internet, they apply for a range of IP addresses.
- If they have 6000 hosts, they ask for a range of IP addresses big enough to cover that, plus room for growth.
- They then allocate their addresses to their hosts in their various offices.

How Internet Addressing Was Meant to Work

- Unfortunately, when IPv4 was created, the designers didn't realise how big the internet was going to get, and they didn't create a big enough address space – there's not enough addresses for everyone.
- The long term solution to this problem is IPv6 which has a much bigger address space.

How Internet Addressing Was Meant to Work

- Private IP addresses with NAT (Network Address Translation) are currently deployed in the majority of enterprise networks as a workaround.
- You'll learn all about private addresses, NAT and IPv6 in a later lecture.
- To understand the lectures until we get to that point, think about it from the context of the originally intended IPv4 design, where all hosts which can communicate on the Internet have a public IP address.

Class A



- The internet authorities split the IPv4 address space into separate classes.
- *Class A* addresses are assigned to networks with a very large number of hosts.
- The high-order (first) bit in a class A address is always set to zero.
- The default subnet mask is /8
- Valid network addresses range from 1.0.0.0 to 126.0.0.0 /8
- This allows for 126 networks and 16,777,214 hosts per network.

128	64	32	16	8	4	2	1	128	64	32	16	8	4	2	1	128	64	32	16	8	4	2	1	128	64	32	16	8	4	2	1
0	0	0	0	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	

- 15.0.0.0/8

Reserved Class A Addresses



- 0.0.0.0/8 is reserved and signifies 'this network'
- 0.0.0.1 to 0.255.255.255 are not valid host addresses

128	64	32	16	8	4	2	1	128	64	32	16	8	4	2	1	128	64	32	16	8	4	2	1	128	64	32	16	8	4	2	1
0	0	0	0	0	0	0	0	0	0	0	0	1	0	1	0	0	0	0	0	1	0	1	0	0	0	0	0	1	1	1	1

- 127.0.0.0/8 in the Class A space is reserved as the loopback address for testing the local computer
- 127.0.0.1 to 127.255.255.255 are not valid host addresses

128	64	32	16	8	4	2	1	128	64	32	16	8	4	2	1	128	64	32	16	8	4	2	1	128	64	32	16	8	4	2	1
0	1	1	1	1	1	1	1	0	0	0	0	1	0	1	0	0	0	0	0	1	0	1	0	0	0	0	0	1	1	1	1

- This wiped out 33,554,428 addresses from the global address pool – whoops!

Subnetting



- Obviously a company wouldn't put all 16,777,214 hosts into a single logical network, this would be terrible for performance and security.
- They would split their /8 address allocation into smaller subnets and allocate these to different offices and types of hosts
- For example if they received 15.0.0.0/8, they could allocate the subnet 15.0.1.0/24 to sales computers in New York, 15.0.2.0/24 to accounting PCs and 15.0.9.0/24 to sales computers in Boston.
- This is called *subnetting* and you'll master it later in this section.