SUBNETTING

Subnet mask.
What the heck is that about?

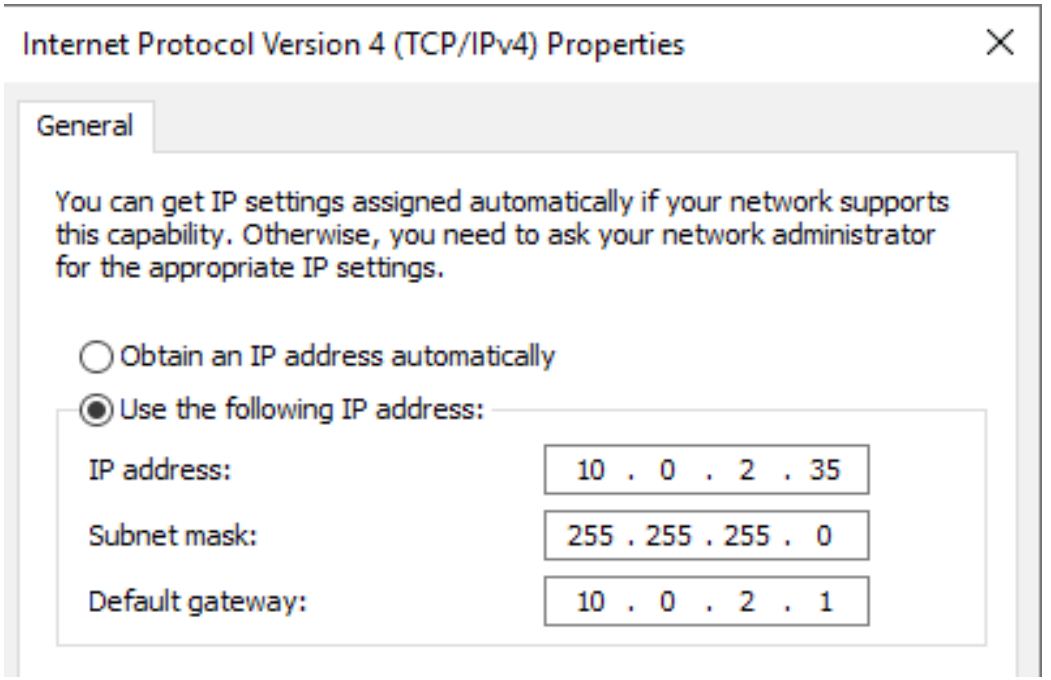

SERIOUSLY

This thing! What does it mean?

What's going on? What's with all the 255's?

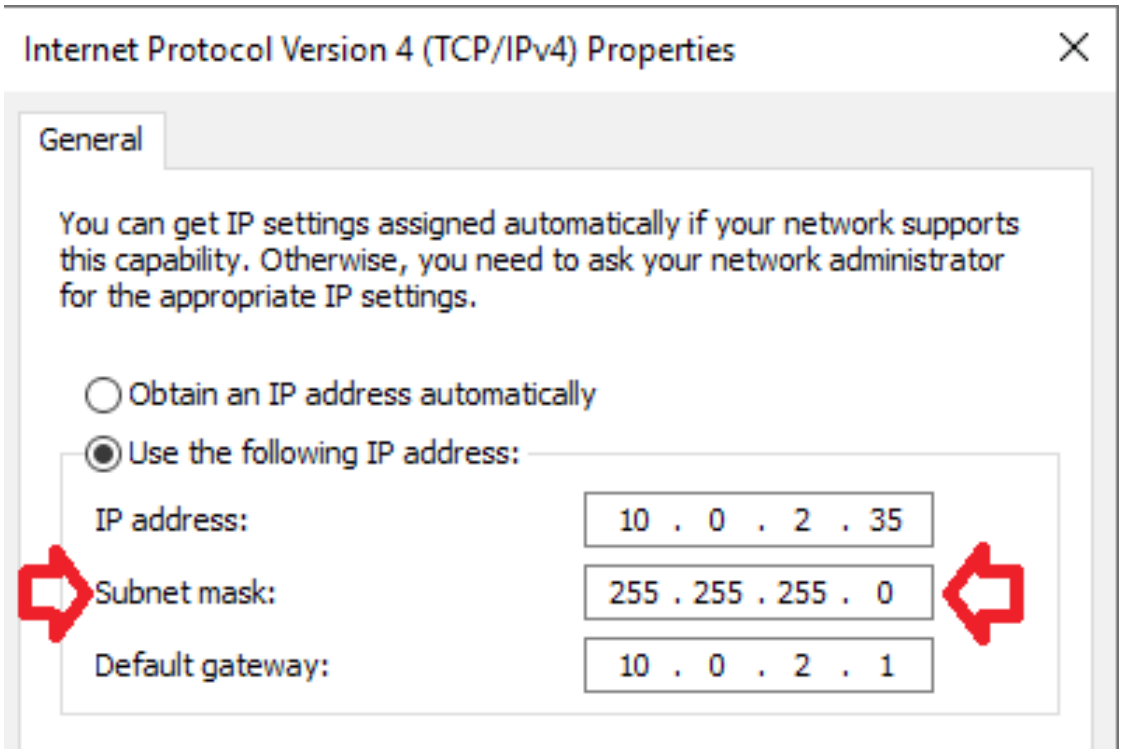

Wait, let me back up. Some basics.

- Every machine that talks on the Internet has an Internet Protocol (IP) address.
- There are two kinds, or protocols: IPv4 or IPv6.
- IPv4 was first deployed in 1983 – most Internet traffic uses this protocol.
- IPv6 was first deployed in the mid 1990's, but is still only a fraction of Internet traffic.
- We're talking about IPv4 for now, here.

IP Addressing

(IPv4)

- An IP address is a 32 bit binary number.
- Humans find decimal easier than binary, so we write IP address in base 10.
- “Dotted decimal” format is 4 binary “octets” (groups of 8), with each 8 digit binary number translated to decimal.

Binary	Decimal
00000001	= 1
00000010	= 2
00010011	= 19
11100000	= 224
11111111	= 255

IP Addressing (cont)

Here's an IP address in dotted decimal and binary:

```
172.16.100.10
 172 -> 10101100
        16 -> 00010000
              100 -> 01100100
                     10 ->00001010
10101100 00010000 01100100 00001010
```

Meanwhile, on your computer

- Things your computer knows that enable it to talk to the network:
  - Its IP address
  - Its “Default gateway”: This is the next hop traffic has to take to get off of the local network and out to where the rest of the Internet is
  - Its “Subnet mask”: This is what lets your computer figure out if another IP address is part of its local network or not. (We'll explain more about this soon.)

Local network vs Remote network

- When your computer is talking to another computer:
  - If the other computer is on the same local network as your computer, your computer sends its traffic directly to the other computer.
  - If the other computer is not on the same local network, your computer sends its traffic to the default gateway and the default gateway ships it off towards the other computer.
    - Other gateways in the path from here to there help out like it's some kind of bucket brigade and each one sends it another link along the route.

So...

How does my computer know if the other computer is on the local network or not?

Meet the Subnet

- Large networks are broken up into smaller network pieces, called “subnets”
- Can be for technical reasons, eg
  - The distance is too far to have some machines on the same network with the media you're using.
  - Vendor recommendations for max number of hosts on a network.
- or administrative reasons, eg
  - HIPAA or PCI might require segmentation at a network layer in some situations.
  - One set of users might be on a privileged network that has access that another network segment doesn't have.

Breaking an IP address up into “host” and “network” portions

- The 32 bits of the IP address are broken into a “network” segment and a “host” segment.
- All hosts on a network have the same “network” address segment.
- Hosts on the same network expect to be able to talk to each other without going through a router/gateway.
- IRL metaphor: neighbors in the same floor of an apartment building being able to visit each other's apartments without using stairs/elevator.

Splitting an IP address into “Host” and “Network”

- Network segment starts from the 'left' (higher bits). Whatever isn't in the network part of the address is the host portion.
- If the first 24 bits of our “172.16.100.10” IP are the network identifier, the bits in red are the network portion:

  10101100 00010000 01100100 00001010
- The bits in black are the host identifier.

Two hosts on the same network

- Let's look at the decimal and binary addresses for two hosts if we have a 24 bit network identifier (the bits in red):

  172.16.100.10

  10101100 00010000 01100100 00001010

  and

  172.16.100.1

  10101100 00010000 01100100 00000001

  The portions in red are the same, so the two hosts are on the same network.

Two hosts on different networks

- Let's look at the decimal and binary addresses for two other hosts if we have a 24 bit network identifier (the bits in red):

  172.16.100.7

  10101100 00010000 01100100 00000111

  and

  10.2.2.1

  00001010 00000010 00000010 00000001

  The portions in red are different, so the two hosts are not part of the same network.

Different sizes of networks

- The size of your network segment determines how many hosts you can have on a network.
- For a 32 bit address with an N bit network identifier, this means that you have [32 – N] bits left over for host numbering.
- 23 bit network identifier leaves 9 bits for host ID
- 24 bit network identifier leaves 8 bits for host ID
- 27 bit network identifier leaves 5 bits for host ID

Special host IDs

- A host portion of all 0's is reserved and should not be used for a host ID.

  172.16.100.0

  10101100 00010000 01100100 00000000

- A host portion of all 1's is used as the broadcast IP address – this is an address that all hosts should respond to.

  172.16.100.255

  10101100 00010000 01100100 11111111

Number of hosts on a network

- Because the network ID of all 1's is reserved for broadcast, you can only number host identifiers up to 1 less than the maximum number of the host ID portion.
- 27 bit network identifier leaves 5 bits for host ID
- 5 bits = 11111 binary = 31 decimal
- Largest host ID is 30 (11110)
- Maximum number of hosts on a 27 bit network is 30.

Number of hosts on a network (cont)

- 23 bit network identifier leaves 9 bits for host ID
  - 111111111 binary = 511 decimal: 510 hosts
- 24 bit network identifier leaves 8 bits for host ID
  - 11111111 binary = 255 decimal: 254 hosts
- 26 bit network identifier leaves 6 bits for host ID
  - 111111 binary = 63 decimal: 62 hosts

So... Breaking a network into smaller subnets doesn't give me more host addresses?

- No. It does the opposite since you can't assign the “all 0” and “all 1” IPs to hosts. This means you lose 2 IPs on each network when you break a network into smaller subnets.

Describing a subnet, including size

- Networks are often described by listing just the network portion (with all 0's for the host portion), followed by a / and the number of bits in the network ID
- The network we used before with 24 bits in the network identifier:

  172.16.100.0

  10101100 00010000 01100100 00000000

  would be described as

  172.16.100.0/24

More examples of network descriptions

10.7.246.0/23

00001010 00000111 11110110 00000000

172.16.200.192/28

0101100 00010000 11001000 11000000

172.16.100.0/24

10101100 00010000 01100100 00000000

What's a “mask”?

- In programming, a “mask” can be used to look at only some digits of a binary number.
- If the mask has a “1” in a position, that bit is kept intact. If the mask has a “0” in a position, that bit is set to “0” (whatever it was before).
- In binary math, this is an “and” function.
- Number AND Mask:
  - 1 AND 0 = 0 (bit is always set to 0 if mask bit is 0)
  - 0 AND 0 = 0 (bit is always set to 0 if mask bit is 0)
  - 1 AND 1 = 1 (bit is kept as 1 if mask bit is 1)
  - 0 AND 1 = 0 (bit is kept as 0 if mask bit is 1)

So, what's a “subnet mask”?

- A subnet mask is a 32 bit number with the “network ID” portion set to all 1's
- A subnet mask for a network with 28 bits in the network portion would have 28 “1” bits followed by 4 “0” bits
- Examples of network masks:
  - /25 = 11111111 11111111 11111111 10000000
  - /16 = 11111111 11111111 00000000 00000000
  - /12 = 11111111 11110000 00000000 00000000

Formatting subnet masks

- Since humans find decimal easier than binary, subnet masks are often printed the same way as IP addresses – dotted decimal format:
- **/25 = 11111111 11111111 11111111 10000000**
  - **255.255.255.128**
- **/16 = 11111111 11111111 00000000 00000000**
  - **255.255.0.0**
- **/12 = 11111111 11110000 00000000 00000000**
  - **255.240.0.0**

Looking back to the beginning of the slide deck...

A subnet mask of “255.255.255.0” means 24 bits of this IP address are the network portion.

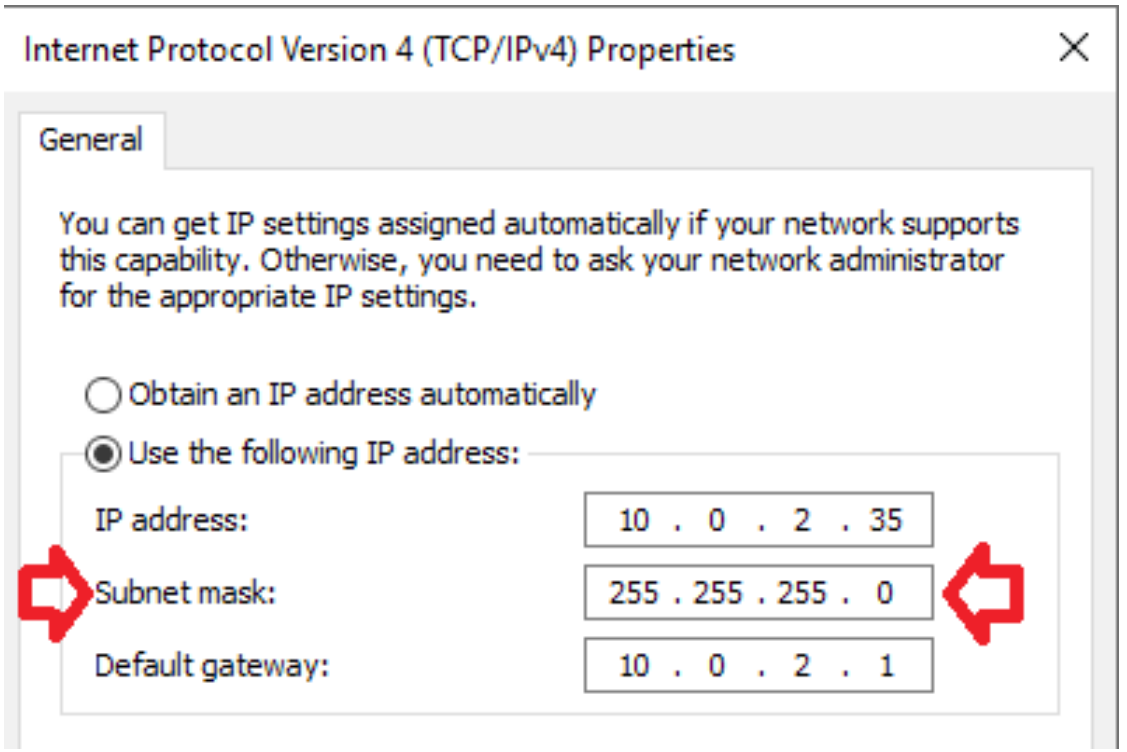

Hex formatting of subnet masks

- Another way to format a subnet mask is to use hex for each octet instead of decimal.
- Each 8 bit octet becomes 2 hex digits.
- 11111111 = 255 = FF
- /24 subnet = 255.255.255.0 = FFFFFF00

Some special cases

- A subnet mask of 32 bits, 255.255.255.255, is used to describe a single host.
- A subnet mask of 31 bits, 255.255.255.254, is not normally usable since the network portion is either “0” or “1”.
- A subnet mask of 0 bits, 0.0.0.0, is used to describe all hosts on the internet.
- Some – but not all – hosts will handle a host ID of all 0's as another broadcast IP. This is why you shouldn't number a host with a host identifier of all 0's.

“Class” addressing

- You may encounter the terms “Class A”, “Class B”, or “Class C”. These are an obsolete way of describing network allocation.
- Class A: first bit of the IP is 0, a /8 subnet
- Class B: first bit of the IP is 1, second bit is 0, a /16 subnet
- Class C: first two bits of the IP are 1, a /24 subnet
- In common usage, “Class C” refers to any /24 network, “Class B” to any /16, etc.

We have no class...

- “Classful” network allocation (Class A, B, C) was superceded by “CIDR” addressing, Classless Inter-Domain Routing.
- CIDR format uses the “/” followed by the number of bits in the network portion.
- Allowed allocation of networks in sizes other than 8, 16, and 24 bit networks.

Network masks vs wildcard masks

- Another way to look at IP addresses is with a “wildcard mask” rather than a “network mask”.
- These are not used for subnetting, but to select certain IPs or ranges of IPs.
- Cisco IOS access control lists (ACLs) are often written with “wildcard” masks.
- Network masks use a “1” to indicate what bits to pay attention to.
- Wildcard masks use a “1” to indicate what bits to ignore.

Network masks vs wildcard masks

(cont)

- Network mask 255.255.255.0 says to ignore the last 8 bits.
- Wildcard mask of 0.0.0.255 says to ignore the last 8 bits.
- Wildcard masks don't have to be contiguous:
  - Wildcard mask “255.0.0.255” says to ignore first and last octets of an IP address.

“Well, actually ...”

- Under normal circumstances, you cannot create a /31 subnet, since there are no usable host entries. (31 bit network identifier leaves 1 bit for host identifier, “0” for network ID, “1” for broadcast.)
- HOWEVER, for point-to-point links, some network equipment will allow a /31, since there are only ever 2 hosts on a point-to-point link and no need for broadcast traffic.
- If you must use it, use with great caution!

“Okay, so, technically...”

- Be careful using leading “0”s in your IP addresses.
- A weird artifact of “dotted decimal” notation is that some programs on some computers will parse a leading “0” in an IP address as an indicator that the number is octal instead of decimal. (Not even all programs on all computers – it's dependent on what libraries are used by the program.)
- Example: octal 010 is 8 in decimal, so 010.010.010.010 Could be parsed as 8.8.8.8:
  - $ ping 010.010.010.010

    PING 010.010.010.010 (8.8.8.8) 56(84) bytes of data.

    64 bytes from 8.8.8.8: icmp_seq=1 ttl=58 time=1.86 ms
- Similarly, some browsers will parse an IP even if it's translated into a giant decimal number:
  - http://3627729134 -> http://216.58.192.238 (google.com)

MOAR INFORMATION

- Wikipedia page on subnets: https://en.wikipedia.org/wiki/Subnetwork
- Cisco doc about subnetting: https://www.cisco.com/c/en/us/support/docs/ip/routi
  - If that link is too badly formatted, try this: https://preview.tinyurl.com/h3rlj3l
- There are many subnet calculators online. http://www.subnet-calculator.com/ is one of them.

WHOSE FAULT IS THIS SLIDE DECK?

- Regis Donovan – Senior Network Engineer
- 20+ years experience with large and small networks.
- Routing, switching, firewalls, load balancers, wireless, security, ops, etc. (Sort of a “utility infielder” kind of network engineer)
- Twitter: rmd1023
- regis at regisdonovan.org