TCP/IP Addressing and Subnetting

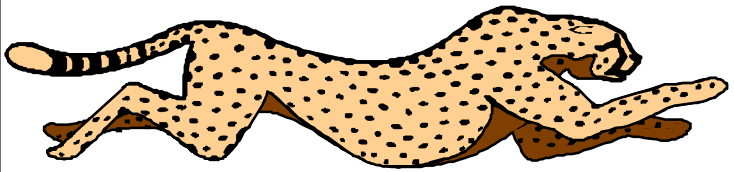

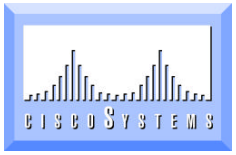

an excerpt from:

A Technical Introduction to TCP/IP Internals

IP Addressing Roadmap

- Format of IP Addresses
- Traditional Class Networks
- Network Masks
- Subnetting
- Supernetting
- Special IP Addresses

IP Addresses

- All IP interfaces have IP addresses
- Each IP interface must have its own unique IP address
- Internally, this is represented as a 32-bit number of 0's and 1's
- IP addresses consist of two parts
  - network identification
  - host identification

Breaking it up into network number and host is key

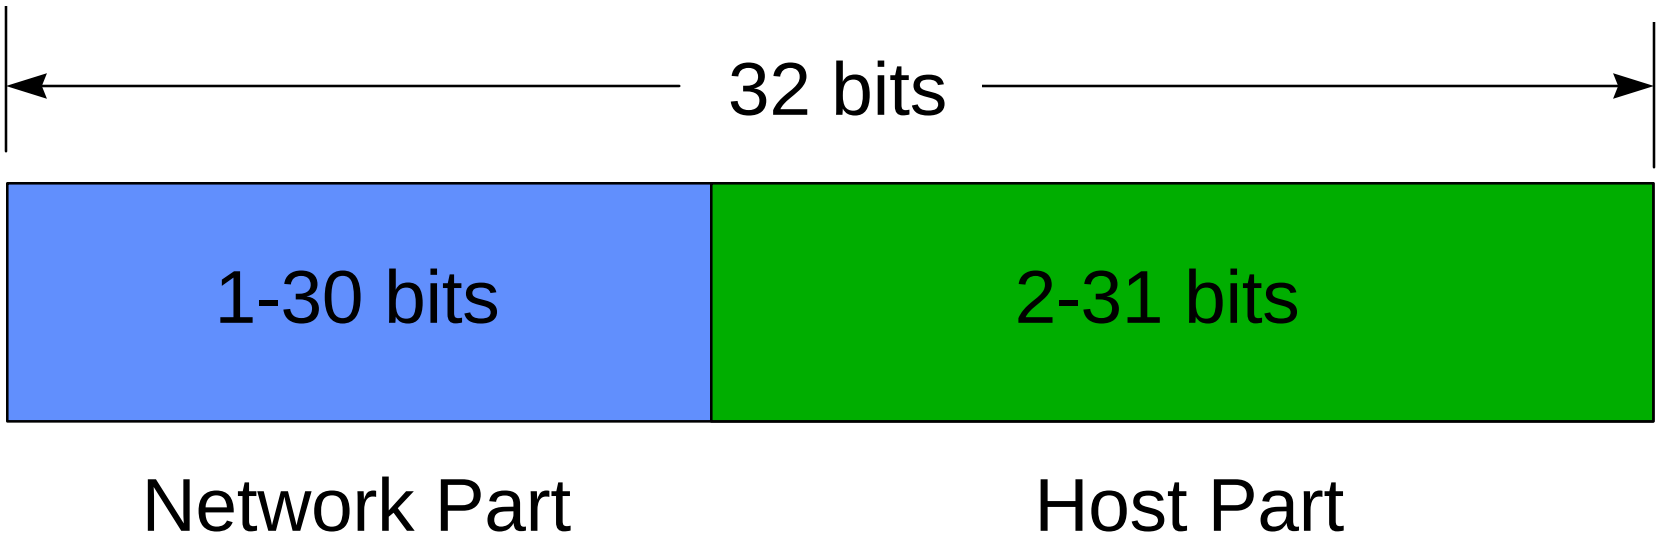

We care because that's how we do routing

- IP routing is based on a simple "next hop" model:
  - Is the destination address ON my network or NOT?
    - If it is ON my network, send it directly
    - If it is NOT on my network, send it via a router
- To match network numbers, you must know what part is network and what part is host

Representing IP Addresses

- There are several ways the IP address can be represented
  - 32 bit number of 0's and 1's
    - 10100001 00101100 11000000 00000001
  - four decimal numbers separated by dots
    - 161.44.192.1
  - hexadecimal representation
    - 9D.2C.BC.01

Traditional Network Class Addresses

- The first dotted quad value identifies the network class and how much of the IP address is the network identifier
  - Class A Networks (first number between 1-127)
  - Class B Networks (first number between 128-191)
  - Class C Networks (first number between 192-223)
- There are also some special IP addresses which are defined in a different way
  - Class D Networks (first number between 224-239)
    - for IP multicast
  - Class E Networks (first number between 240-255)
    - for Landmark routing

This is due to a clever hack by the IP authors

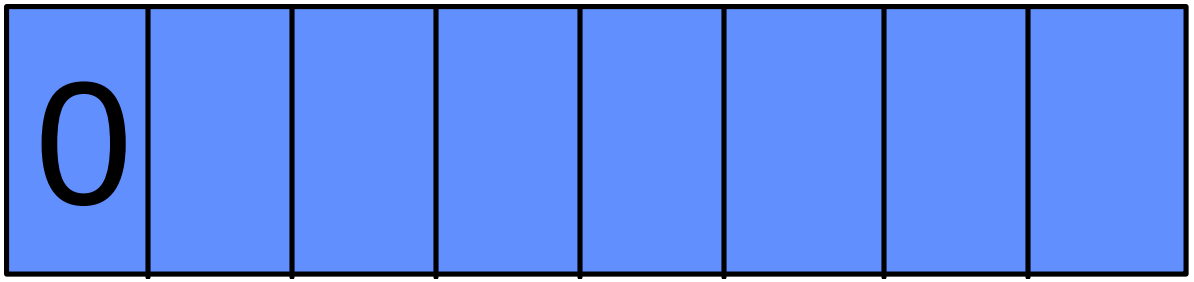

Class A Address

Network = 8 bits

Host = 24 bits

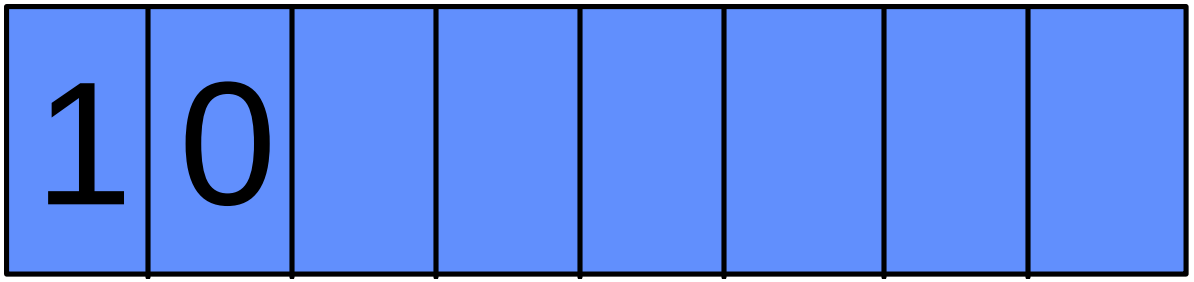

Class B Address

Network = 16 bits

Host = 16 bits

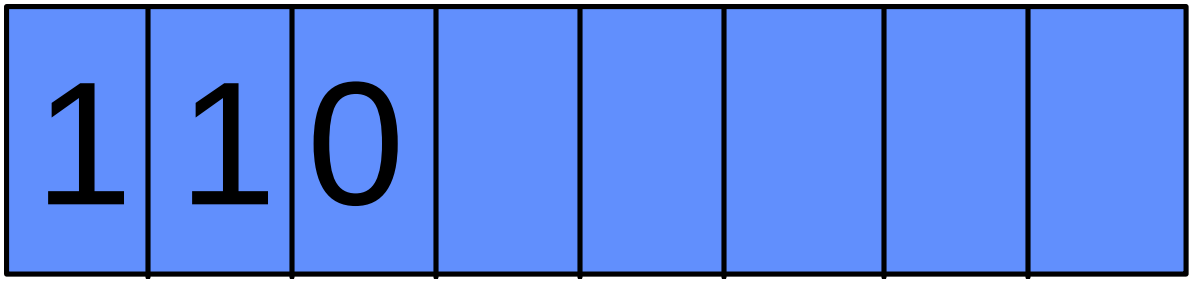

Class C Address

Network = 24 bits

Host = 8 bits

Assigning Network Numbers

- Network numbers imply some space for hosts
  - Network numbers are assigned by your Internet Service Provider, who got them from the InterNIC (Network Information Center)
- Network numbers are written as a full 32-bit quantity (and an implied network mask)
- Networks end with some number of contiguous zero-bits on the right
- These zero-bits are where customers can use one bits for host addresses

You can also use RFC 1597 addresses

- For “local” use, although your provider may reserve some of them

- 10.0.0.0 - 10.255.255.255 (10/8)
- 172.16.0.0 - 172.31.255.255 (172.16/12)
- 192.168.0.0 - 192.168.255.255 (192.168/16)

- (see also RFC 1918 and RFC 1627)

Network Mask

- Identifies how many bits of the IP address the host may use
- The mask contains a 1 bit for every bit in the “network portion” of the address
- The mask contains a 0 bit for every bit in the “host portion” of the address
- Every host on a network must have the same network mask
- May also be called the Subnet Mask

Default Network Masks

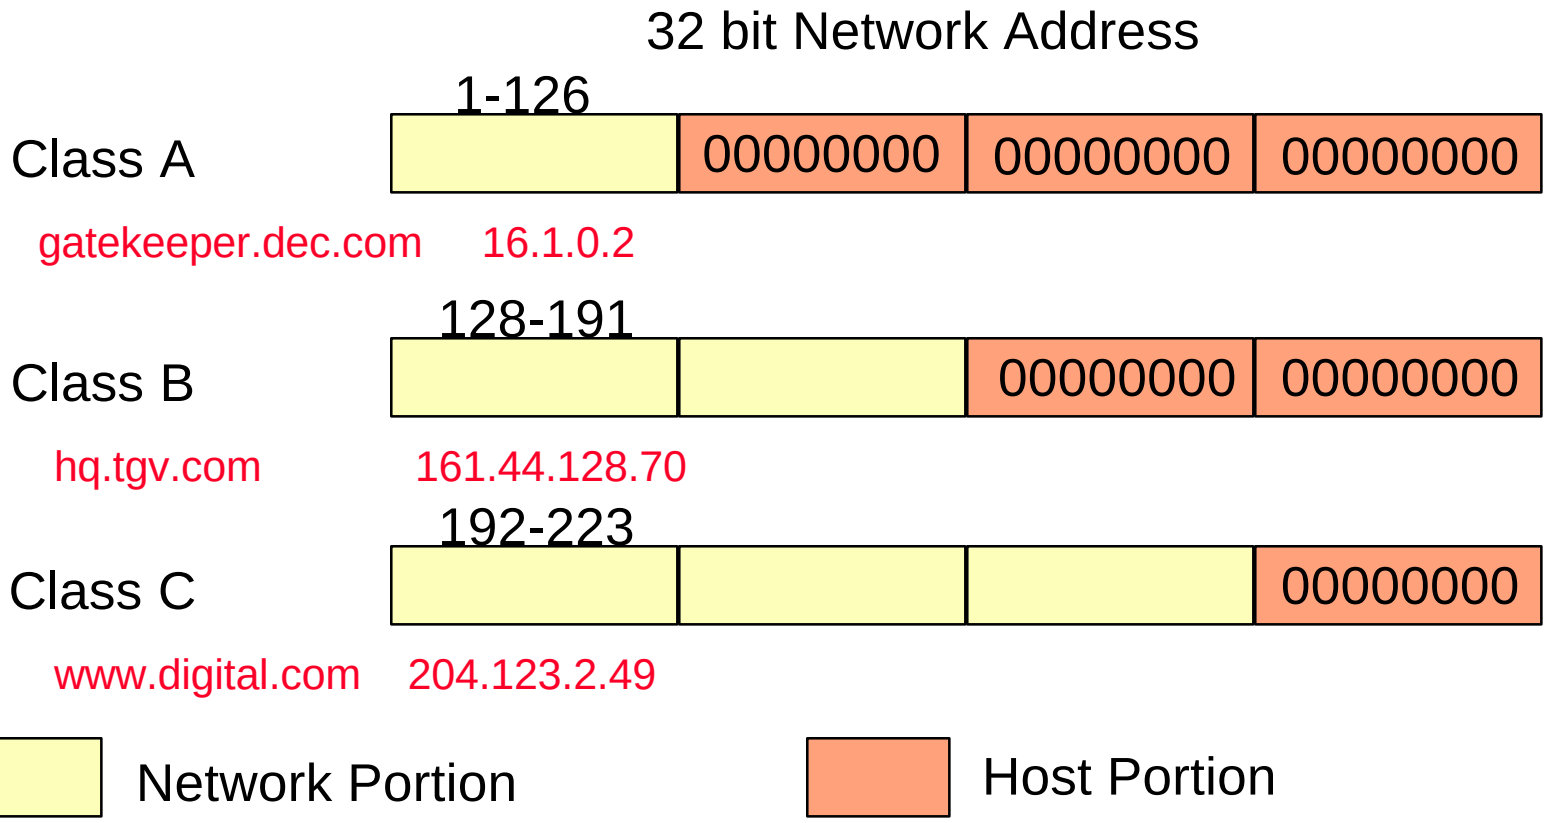

Network Masks are now shown with slash notation

- Class A network number
  - 8 bits of network, 24 bits of host
  - 10.0.0.0/8
- Class B network number
  - 16 bits of network, 16 bits of host
  - 128.196.0.0/16
- Class C network number
  - 24 bits of network, 8 bits of host
  - 192.245.12.0/24

technically, Class D numbers have a 28-bit prefix, but this is never used in practice.

Prefixes and Network Masks almost the same

- A network mask can represent an arbitrary set of bits:
  - 11111111 11110111 10101010 00000000
- A prefix can only represent contiguous ones bits:
  - 11111111 11111111 11111100 00000000
  - is the same as /22
- “Subnet numbers SHOULD be contiguous...” (RFC 1812)

Translating between the two is easy

255.255.0.0	/16
255.255.128.0	/17
255.255.192.0	/18
255.255.224.0	/19
255.255.240.0	/20
255.255.248.0	/21
255.255.252.0	/22
255.255.254.0	/23
255.255.255.0	/24
255.255.255.128	/25
255.255.255.192	/26
255.255.255.224	/27
255.255.255.240	/28
255.255.255.248	/29
255.255.255.252	/30

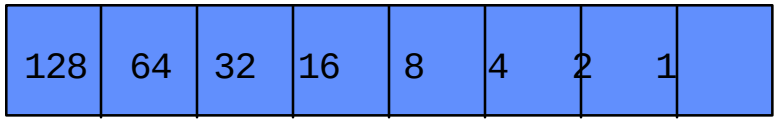

10000000 = 128 = 128

11000000 = 192 = 128+64

11100000 = 224 = 128+64+32

11110000 = 240 = 128+64+32+16

11111000 = 248 = 128+64+32+16+8

11111100 = 252 = 128+64+32+16+8+4

11111110 = 254 = 128+64+32+16+8+4+2

11111111 = 255 = 128+64+32+16+8+4+2+1

Simple Network Example

- Network address 192.195.240.0
- Network mask 255.255.255.0 or /24
- Host numbers
  - 192.195.240.1 - 192.195.240.254
- First 24 bits identify the network
- Last 8 bits are for the host EXCEPT:
  - Can’t use all 0’s (.0, assigned network)
  - Can’t use all 1’s (.255, broadcast address)

Two addresses in every network are special

- Host part all ones (usually “255-ish”)
  - This is defined as the broadcast address, and means “all systems on the current network”
- Host part all zeros (usually “0-ish”)
  - This is defined as the network number and cannot be used
- Example:
  - 192.245.12.0/24 is a network with 8 bits
  - 192.245.12.255 is the broadcast address
  - 192.245.12.0 is the network number
  - 192.245.12.1 through 192.245.12.254 are hosts

Network Mask Usage

- Host address: 192.195.240.4
- Network Mask: 255.255.255.0 (/24)
  - Logical AND yields network 192.195.240.0
- Destination host: 192.195.241.4
  - Logical AND yields network 192.195.241.0
- Since the network 24 bits of the local host and destination host are unequal, the destination host is not on local net

IP Subnetworks

- Allows the “host” part of IP address to be further split
- Arbitrary bit position divides subnet and host
- Transparent outside of local network
- Must be agreed upon by all hosts in local network
- Allows additional layer of hierarchy to be built into a single IP network number
- Helps reduce address space waste

Originally used to break up Class B networks

- Organization would get a Class B network number (e.g., 128.196.0.0)
- Organization would start to buy routers
- Organization would want to break up that network into smaller pieces

“Subnet a B into Cs”

- Original network number was 128.196.0.0
- Original network mask was 255.255.0.0

- Subnet with network mask 255.255.255.0
- This gives 256 networks of 254 hosts each
  - 128.196.0.1 through 128.196.0.254
  - 128.196.1.1 through 128.196.1.254
  - 128.196.255.1 through 128.196.255.254

Finishing “Subnet a B into Cs”

- The world (everyone outside) knows of the network as 128.196.0.0 (no subnet)
- Everyone inside must agree that the network mask is 255.255.255.0

That's where we used to end the class...

- You can't get a class B network number any more
- You probably get a block of class C network numbers which you need to break up yourself
- Address "space" is scare
  - Class B addresses are very scarce
  - Class C addresses are common, but routing table space is very scarce
- Major ISPs are filtering "inefficient" blocks

Subnets and Supernets

- In the old Internet the default network mask was based on the first few bits of the first octet
- In the new Internet network masks are defined for all networks
  - a network subdivided into smaller subnets uses subnet masks
  - a network comprised of a consecutive range of network numbers uses supernet masks (CIDR)

Example of Subnetting

- Physical topology of two physical LANs (ethernets) separated by a router
- The router (host) must know which interface to select
  - Each interface must be on a different IP network

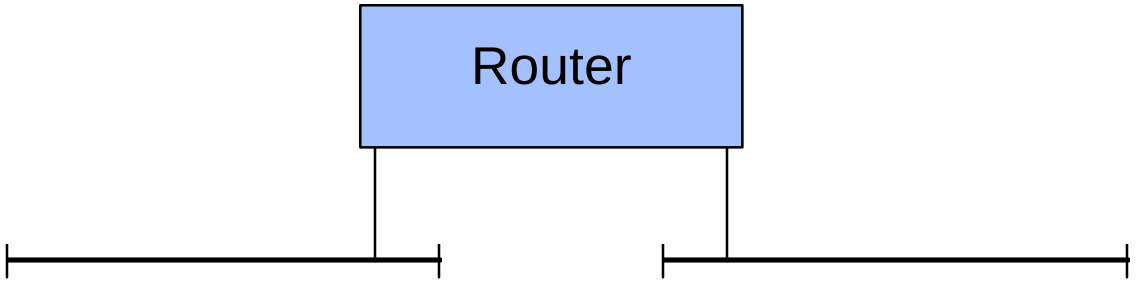

Subnet Example

- We could assign each its own, like 192.195.240.0 and 192.195.241.0
  - wastes lots of IP addresses if < 510 hosts
- We can take our /24, and split it into /25 networks:
  - 192.195.240.[0][7 host bits]
    - 192.195.240.1 - 192.195.240.126
  - 192.195.240.[1][7 host bits]
    - 192.195.240.128 - 192.195.240.255
- This gives us two subnetworks of 2**7 hosts each (minus 2 per subnet, of course)

However

- We can't use a subnetwork of all 0 bits
  - some routers can't handle that
    - can't distinguish between route to both nets and route to subnet 0
- Therefore we can't use a one-bit network mask, such as in the previous example, because it's either all zeroes or all ones

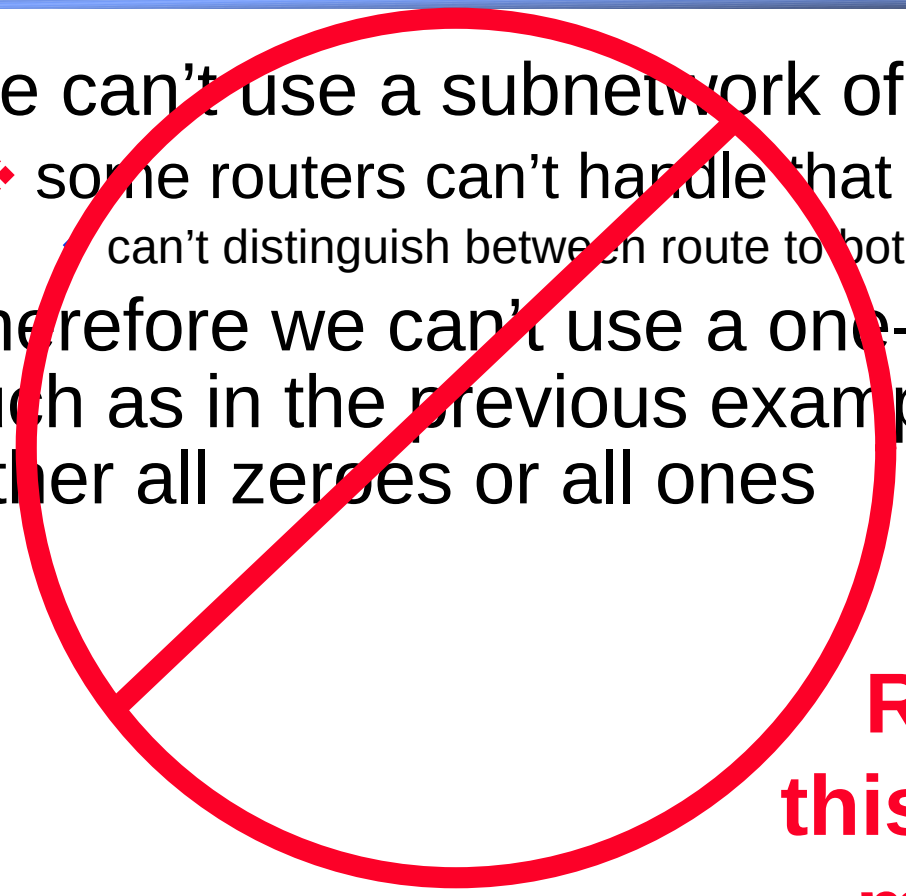

RFC 1812 changed this! Get your router manufacturer to fix their software!

Let's do Two Subnets anyway

- If we assign two bits:
  - 192.195.240.[00][6 host bits]
  - 191.195.240.[01][6 host bits]
  - 192.195.240.[10][6 host bits]
  - 192.195.240.[11][6 host bits]

The Subnet Mask

- The subnet mask in this case must represent the part the IP kernel needs to compare when checking for whether this is on the local network
  - 255.255.255.192 includes those extra two bits at the end
    - 192 = 11000000
    - mask = 11111111.11111111.11111111.11000000
    - prefix = /26 (/24 + 2 bits)

Subnetting 192.195.240.0

First three octets are: 192.195.240.xxx

net num	net num in binary	num in decimal	b-cast address	host range
0	00 000000	.0	.63	.1 through .62
1	01 000000	.64	.127	.65 through .126
2	10 000000	.128	.191	.129 through .190
3	11 000000	.192	.255	.193 through .254

Network Mask = /26 = 255.255.255.192

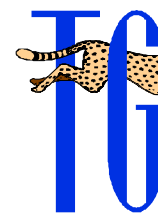

Further Subnetting

- Let’s say we have need for multiple physical networks, like 10 or so, each of which will have a few systems on it

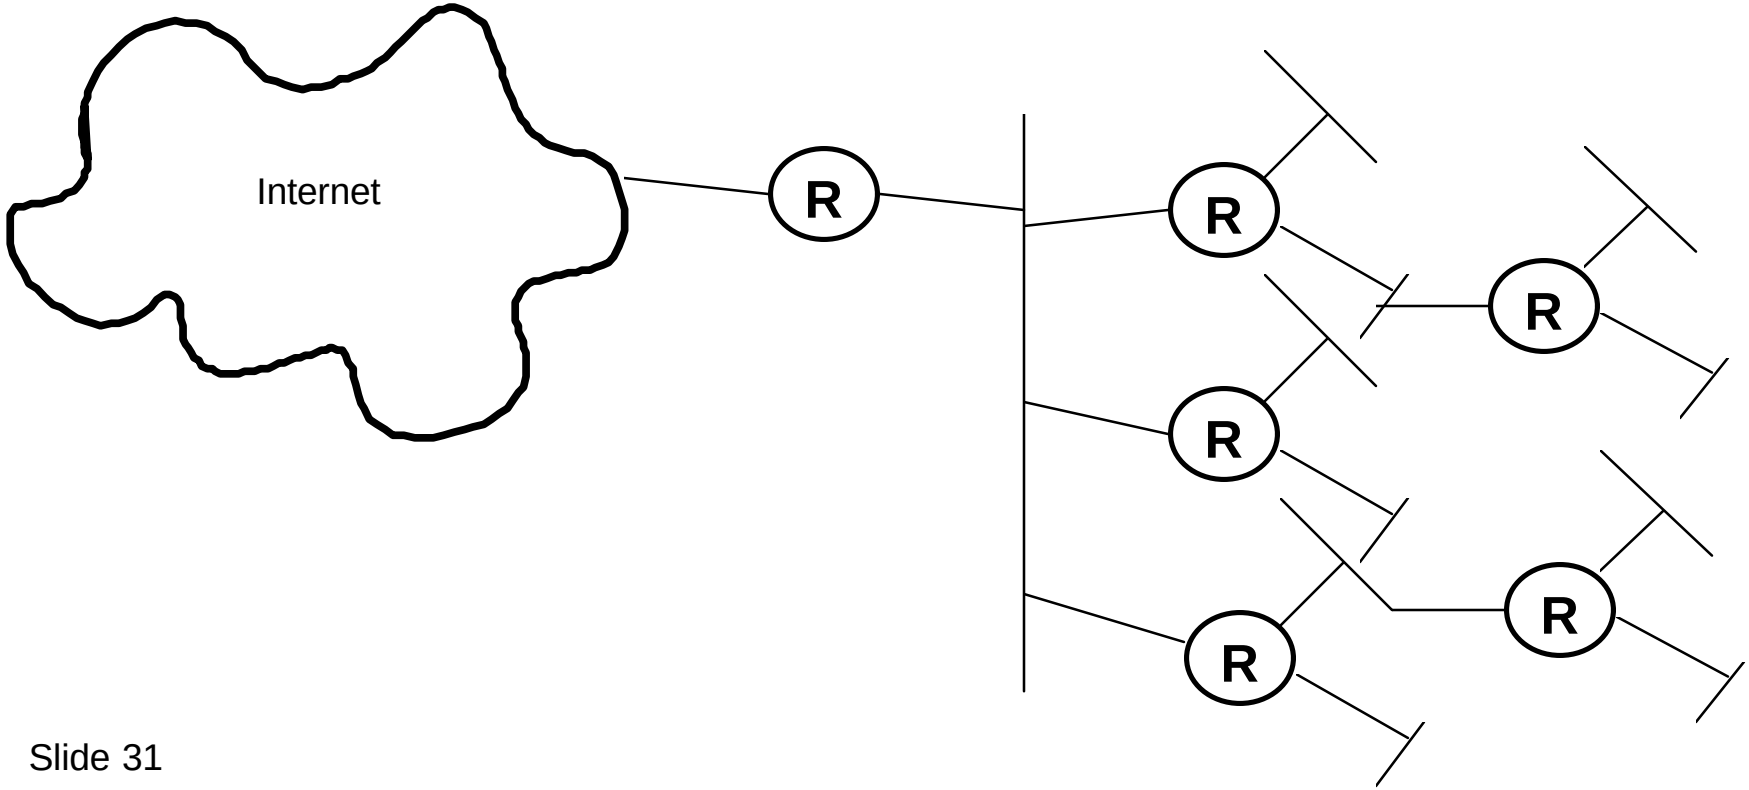

Find the lowest power of 2 that fits

2**8	256	(not very useful)
2**7	128	
2**6	64	
2**5	32	
2**4	16	
2**3	8	
2**2	4	
2**1	2	
2**0	1	(not very useful)

"32 is too many, and 8 is too few, so 16 must be just right"

- little Red Riding Hood

Four additional bits for network number works

- The original network had a 24 bit netmask
  - /24 prefix
  - 255.255.255.0 mask
- Subnet as a /28 (/24 + /4)
  - /28 prefix
  - 11111111 11111111 11111111 11110000
  - 255.255.255.240 mask
- This will leave us 16 host addresses per subnet, minus one for the network address and one for the broadcast address = 14

Example of subnetting a network to a /27

Original network number: 192.245.12.0/24

First three octets of everything: 192.245.12.xxx

net num	net num in binary	num in decimal	b-cast address	host range
0	000 00000	.0	.31	.1 to .30
1	001 00000			
2	010 00000			
3	011 00000			
4				
5				
6				
7	111 00000	.224	.255	.225 to .254

A little binary-to-decimal conversion table

0000 0000	0	0010 0000	32	0100 0000	64	0110 0000	96	1000 0000	128	1010 0000	160	1100 0000	192	1110 0000	224
0000 0001	1	0010 0001	33	0100 0001	65	0110 0001	97	1000 0001	129	1010 0001	161	1100 0001	193	1110 0001	225
0000 0010	2	0010 0010	34	0100 0010	66	0110 0010	98	1000 0010	130	1010 0010	162	1100 0010	194	1110 0010	226
0000 0011	3	0010 0011	35	0100 0011	67	0110 0011	99	1000 0011	131	1010 0011	163	1100 0011	195	1110 0011	227
0000 0100	4	0010 0100	36	0100 0100	68	0110 0100	100	1000 0100	132	1010 0100	164	1100 0100	196	1110 0100	228
0000 0101	5	0010 0101	37	0100 0101	69	0110 0101	101	1000 0101	133	1010 0101	165	1100 0101	197	1110 0101	229
0000 0110	6	0010 0110	38	0100 0110	70	0110 0110	102	1000 0110	134	1010 0110	166	1100 0110	198	1110 0110	230
0000 0111	7	0010 0111	39	0100 0111	71	0110 0111	103	1000 0111	135	1010 0111	167	1100 0111	199	1110 0111	231
0000 1000	8	0010 1000	40	0100 1000	72	0110 1000	104	1000 1000	136	1010 1000	168	1100 1000	200	1110 1000	232
0000 1001	9	0010 1001	41	0100 1001	73	0110 1001	105	1000 1001	137	1010 1001	169	1100 1001	201	1110 1001	233
0000 1010	10	0010 1010	42	0100 1010	74	0110 1010	106	1000 1010	138	1010 1010	170	1100 1010	202	1110 1010	234
0000 1011	11	0010 1011	43	0100 1011	75	0110 1011	107	1000 1011	139	1010 1011	171	1100 1011	203	1110 1011	235
0000 1100	12	0010 1100	44	0100 1100	76	0110 1100	108	1000 1100	140	1010 1100	172	1100 1100	204	1110 1100	236
0000 1101	13	0010 1101	45	0100 1101	77	0110 1101	109	1000 1101	141	1010 1101	173	1100 1101	205	1110 1101	237
0000 1110	14	0010 1110	46	0100 1110	78	0110 1110	110	1000 1110	142	1010 1110	174	1100 1110	206	1110 1110	238
0000 1111	15	0010 1111	47	0100 1111	79	0110 1111	111	1000 1111	143	1010 1111	175	1100 1111	207	1110 1111	239
0001 0000	16	0011 0000	48	0101 0000	80	0111 0000	112	1001 0000	144	1011 0000	176	1101 0000	208	1111 0000	240
0001 0001	17	0011 0001	49	0101 0001	81	0111 0001	113	1001 0001	145	1011 0001	177	1101 0001	209	1111 0001	241
0001 0010	18	0011 0010	50	0101 0010	82	0111 0010	114	1001 0010	146	1011 0010	178	1101 0010	210	1111 0010	242
0001 0011	19	0011 0011	51	0101 0011	83	0111 0011	115	1001 0011	147	1011 0011	179	1101 0011	211	1111 0011	243
0001 0100	20	0011 0100	52	0101 0100	84	0111 0100	116	1001 0100	148	1011 0100	180	1101 0100	212	1111 0100	244
0001 0101	21	0011 0101	53	0101 0101	85	0111 0101	117	1001 0101	149	1011 0101	181	1101 0101	213	1111 0101	245
0001 0110	22	0011 0110	54	0101 0110	86	0111 0110	118	1001 0110	150	1011 0110	182	1101 0110	214	1111 0110	246
0001 0111	23	0011 0111	55	0101 0111	87	0111 0111	119	1001 0111	151	1011 0111	183	1101 0111	215	1111 0111	247
0001 1000	24	0011 1000	56	0101 1000	88	0111 1000	120	1001 1000	152	1011 1000	184	1101 1000	216	1111 1000	248
0001 1001	25	0011 1001	57	0101 1001	89	0111 1001	121	1001 1001	153	1011 1001	185	1101 1001	217	1111 1001	249
0001 1010	26	0011 1010	58	0101 1010	90	0111 1010	122	1001 1010	154	1011 1010	186	1101 1010	218	1111 1010	250
0001 1011	27	0011 1011	59	0101 1011	91	0111 1011	123	1001 1011	155	1011 1011	187	1101 1011	219	1111 1011	251
0001 1100	28	0011 1100	60	0101 1100	92	0111 1100	124	1001 1100	156	1011 1100	188	1101 1100	220	1111 1100	252
0001 1101	29	0011 1101	61	0101 1101	93	0111 1101	125	1001 1101	157	1011 1101	189	1101 1101	221	1111 1101	253
0001 1110	30	0011 1110	62	0101 1110	94	0111 1110	126	1001 1110	158	1011 1110	190	1101 1110	222	1111 1110	254
0001 1111	31	0011 1111	63	0101 1111	95	0111 1111	127	1001 1111	159	1011 1111	191	1101 1111	223	1111 1111	255

Example of subnetting a network to a /27

Original network number: `192.245.12.0/24`

First three octets of everything: `192.245.12.xxx`

net num	net num in binary	num in decimal	b-cast address	host range
0	000 00000	.0	.31	.1 to .30
1	001 00000	.32	.63	.33 to .62
2	010 00000	.64	.95	.65 to .94
3	011 00000	.96	.127	.97 to .126
4	100 00000	.128	.159	.129 to .158
5	101 00000	.160	.191	.161 to .190
6	110 00000	.192	.223	.193 to .222
7	111 00000	.224	.255	.225 to .254

Subnet Mask Summary

- A network can be split into multiple smaller logical networks
- Network mask or prefix indicates which bits to compare when making routing decisions
  - 255.255.255.0 is the same as /24
- Host part: All 1s and all 0s cannot be used
  - All host bits ones are broadcast address
  - All host bits zero are network address
- Network part: All 0s can be a problem
  - With non-RFC 1812 compliant routers

Supernets

- Supernetting takes multiple logical networks and makes one new logical network
  - Combine multiple Class-C networks for one physical network
  - More than 256 hosts on a cable
- Supernetting makes the network mask less specific than the default mask

Supernet Example

- Assigned network numbers of 204.17.32.0 and 204.17.33.0
- A supernet mask of 255.255.254.0 would address both nets on the same physical wire

204.17.32.0	=	11001100.00010001.0010000**0**.00000000
204.17.33.0	=	11001100.00010001.0010000**1**.00000000
255.255.254.0	=	11111111.11111111.1111111**0**.00000000

Special IP Addresses

- A number of IP addresses are considered “special” by the RFCs and most implementations
- These are mostly for broadcast and loopback purposes
- We’ll use the notation { xxx, yyy } to indicate the network and host part
  - xxx = network part
  - yyy = host part

{ 0 , 0 } and { 0 , <host> }

- {0,0} means “this host, on this network”
- Written also as 0.0.0.0
- Never used except in testing or booting
  - BOOTP uses 0.0.0.0 to indicate “me”
- {0,<host>} means “this host, on this network” as well.
- Reserved
  - but I’ve never seen it used

{ -1 , -1 } and { <this net> , -1 }

- {-1 , -1} is the “everywhere” broadcast address
- Usually written as 255.255.255.255
- Does not go outside of your local network
-
- {<this net>, -1} is the broadcast to all hosts in your local net
- Very commonly used
  - For example, 192.245.12.0/24 broadcast is 192.245.12.255

{ 127 , <anything> }

- Any address with the first octet 127
- Typically used as 127.0.0.1
- A Class A network number which is reserved for loopback purposes
- You may never use Net 127, even if you want to

Special Address Summary

- 0.0.0.0 means “me”
- network.0 means “this network”
- network.255 means “broadcast”
- 255.255.255.255 means “broadcast everywhere”
- 127.0.0.1 means “loopback”
  - (actually: 127.anything)

IP Addressing
Key Concepts

- IP Addresses are 32 bit numbers represented as a “dotted quad”
- Network numbers are assigned by the Internic or Internet access provider
- Host numbers are assigned by the network manager
- Network masks identify which part of the IP address is the network portion

TCP/IP References

- TCP/IP Illustrated, Volume 1, The Protocols, W. Richard Stevens, Addison-Wesley Publishing Company, 1994
- Interconnections: Bridges and Routers, Radia Perlman, Addison-Wesley Publishing Company, 1992
- The Simple Book, An Introduction to Internet Management, Marshall T. Rose, PTR Prentice-Hall, Inc, 1994

IP Addressing

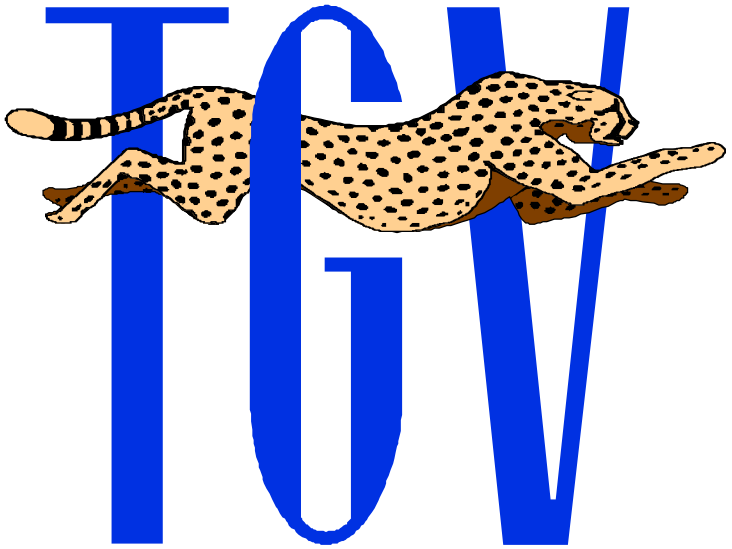

Questions ?

TCP/IP Addressing and Subnetting

Mike Sullenberger
Cisco Systems, Inc.
mls@cisco.com
+1 800 553 NETS