Application Note

IP Addressing
A Simplified Tutorial

July 2002
COMPAS ID 92962

Avaya Labs

AVAYA

All information in this document is subject to change without notice. Although the information is believed to be accurate, it is provided without guarantee of complete accuracy and without warranty of any kind. It is the user's responsibility to verify and test all information in this document. Avaya shall not be liable for any adverse outcomes resulting from the application of this document; the user must take full responsibility.

Companion document

- LANs and VLANs: A Simplified Tutorial
  http://www1.avaya.com/enterprise/whitepapers/vlan-tutorial.pdf

Introduction

The purpose of this tutorial is to give the newcomer to data networking a basic understanding of IP addressing. The following topics are covered.

- IP addressing fundamentals
- Classful IP addressing
- Subnet masks
- Variable length subnet masks (VLSM)
- Classless inter-domain routing (CIDR)
- Routing and routing protocols

AVAYA

IP Addressing Fundamentals

OSI and TCP/IP

OSI Reference Model	TCP/IP	Terms used in this tutorial
7 – Application	Application	
6 – Presentation		
5 – Session		
4 – Transport	Host – to – Host (TCP/UDP)	TCP port, UDP port
3 – Network	Internet (IP)	**IP address**
2 – Data Link	Network Interface	MAC address
1 – Physical		

- This table is presented for reference purposes.
  - The first column shows the 7-layer OSI Reference Model, which is a model used to design protocols that make networking possible.
  - The second column shows the TCP/IP protocol stack in reference to the OSI model. TCP/IP is the prevalent protocol stack for data networking.
  - The third column shows that an IP address is a layer 3 (L3) address, as well as its relationship to the MAC address and TCP/UDP port, which are not covered in this tutorial.

AVAYA

Anatomy of an IP address

- The IP address is a 32-bit address that consists of two components.
- One component is the network portion of the address, consisting of the network bits.
  - The network bits make up the left portion of the address.
  - They consist of the first bit up to some boundary, to be discussed later.
- The second component is the host portion of the address, consisting of the host bits.
  - The host bits make up the right portion of the address.
  - They consist of the remaining bits not included with the network bits.

AVAYA

The mask

- The network portion of the address is separated from the host portion of the address by a mask.
- The mask simply indicates how many bits are used for the network portion, leaving the remaining bits for the host portion.
- A 24-bit mask indicates that the first 24 bits of the address are network bits, and the remaining 8 bits are host bits.
- A 16-bit mask indicates that the first 16 bits of the address are network bits, and the remaining 16 bits are host bits.
- And so forth…
- The difference between a network mask and a subnet mask will be explained as this tutorial progresses.

AVAYA

Quick lesson in binary math

- Binary math is based on powers of 2, as opposed to powers of 10 for decimal math.
  - Whereas decimal math has a 1s place, 10s place, 100s place, and so forth…
  - Binary math has a 1s place, 2s place, 4s place, 8s place, and so forth.
- Given an octet (8 bits), when a bit in the octet is set (1) its value is…
  - 128 = left-most bit (most significant bit) = 2^7
  - 64 = next bit = 2^6
  - 32 = next bit = 2^5
  - 16 = next bit = 2^4
  - 8 = next bit = 2^3
  - 4 = next bit = 2^2
  - 2 = next bit = 2^1
  - 1 = right-most bit (least significant bit) = 2^0
- When a bit in an octet is not set (0) its value is zero.
- The decimal value of an octet is the sum of each set bit's value.
  - 11000000 = 128 + 64 = 192
  - 10101000 = 128 + 32 + 8 = 168
  - 11111111 = 128 + 64 + 32 + 16 + 8 + 4 + 2 + 1 = 255

AVAYA

Dotted decimal notation

- Machines read the IP address as a stream of 32 bits.
- However, for human consumption, the IP address is written in dotted decimal notation.
  - The 32-bit address is divided into 4 groups of 8 bits (an octet or a byte).
  - Each octet is written as a decimal number ranging from 0 to 255.
  - The decimal numbers are separated by periods, or dots.

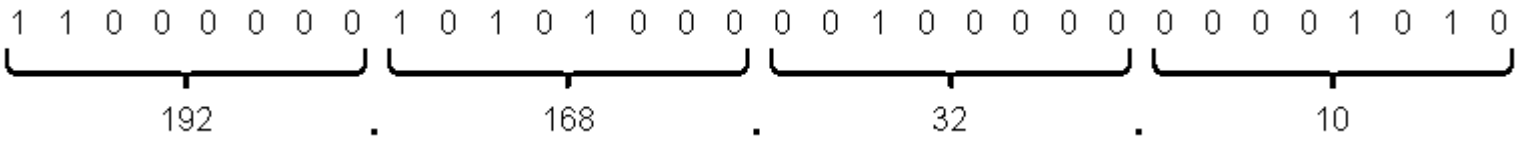

AVAYA

Network, host, and broadcast addresses

- For a given IP network…
  - the network bits remain fixed and the host bits vary.
  - the network address is the one that results when all the host bits are not set (the result of performing an AND operation on the address and its mask).
  - the broadcast address is the one that results when all the host bits are set.
  - host addresses are those that result with all remaining combinations of the host bits.
- The next two slides show examples of how to determine the various addresses for two networks.

AVAYA

24-bit mask (255.255.255.0)

16-bit mask (255.255.0.0)

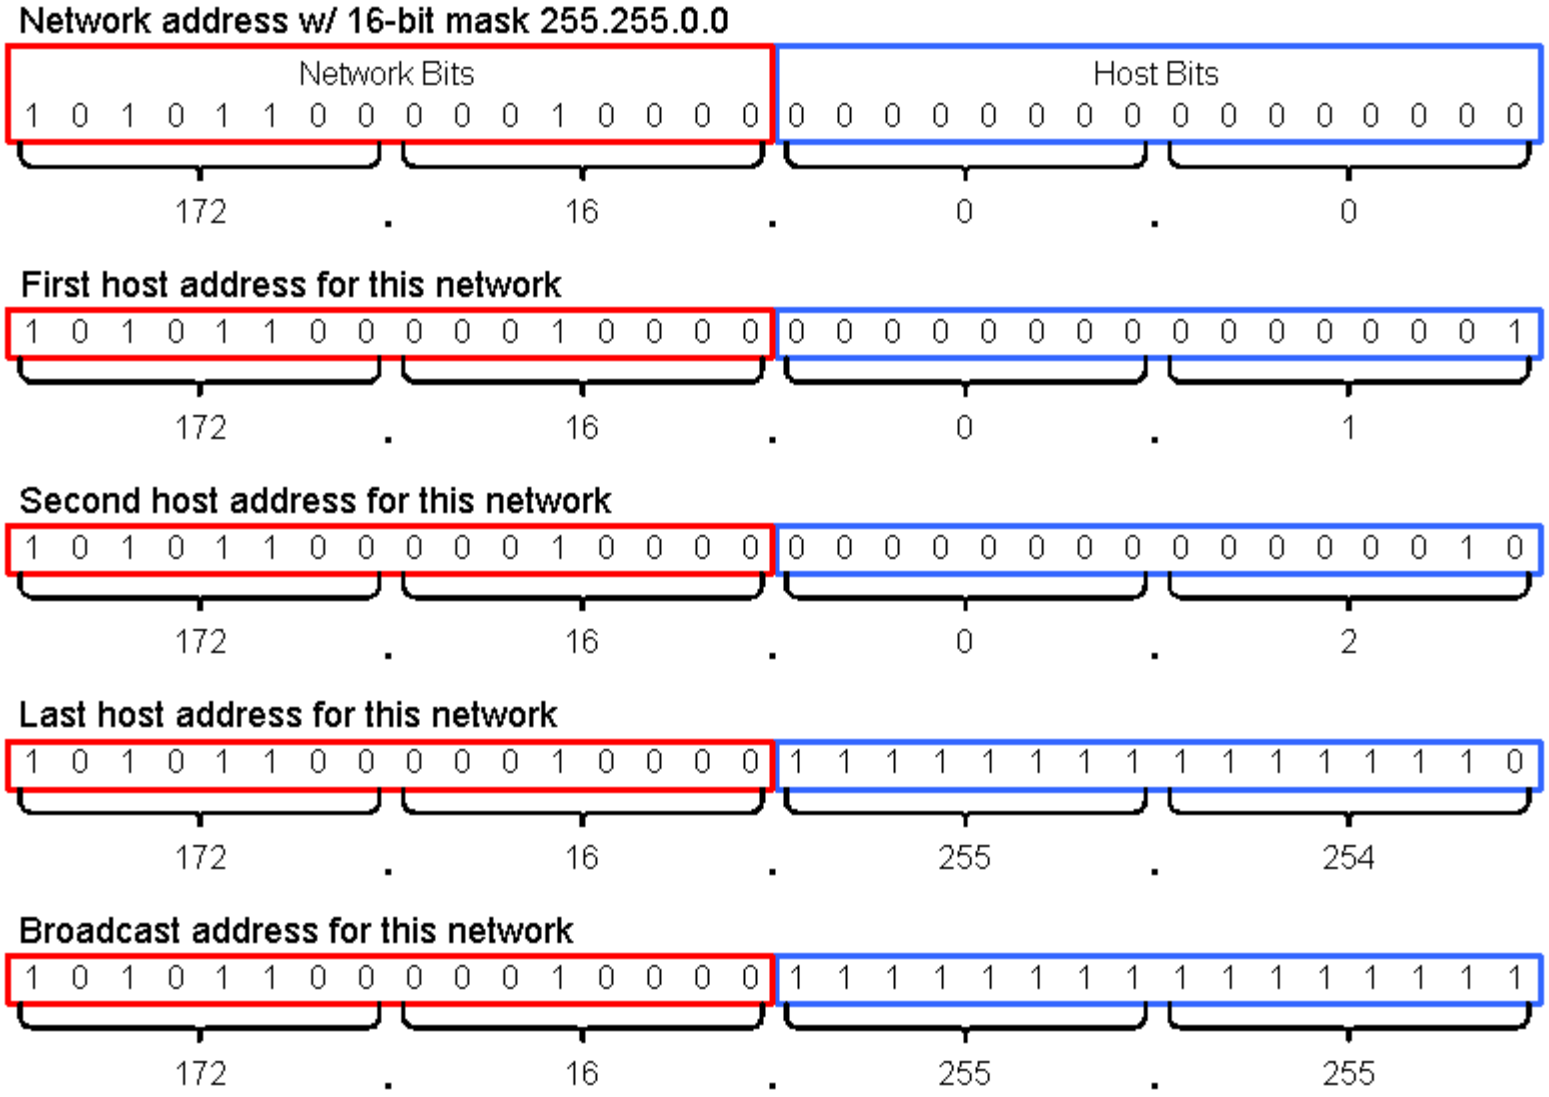

Significance of IP networks and hosts

- An IP host is any device with an IP address, such as a PC.
- Multiple hosts reside on a given IP network or subnet (short for subnetwork). Subnets will be discussed later.
- A group of IP networks is an internetwork, with the largest internetwork being the Internet.
- What is typically called a “data network” is technically an internetwork, because multiple IP networks are connected together by routers.

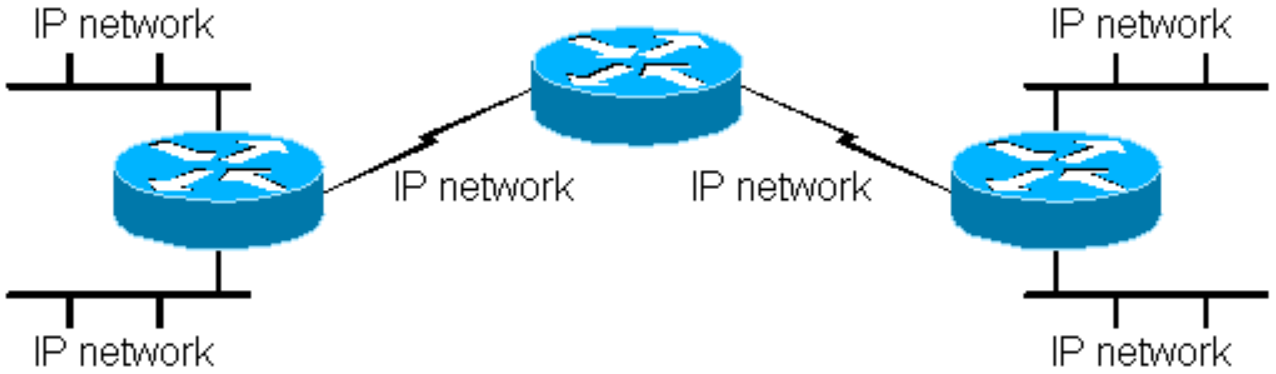

- This internetwork contains 6 IP networks.
- Note that even a link between routers is a network.

AVAYA

Physical network vs. IP network

- In terms of physical connectivity, it is common to see these terms used. The term “network” here describes as a whole the connected devices that make up this data network.

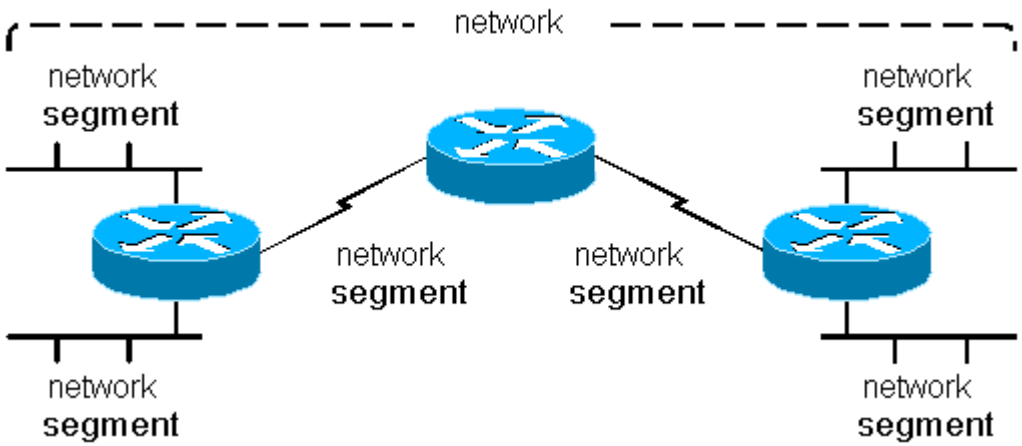

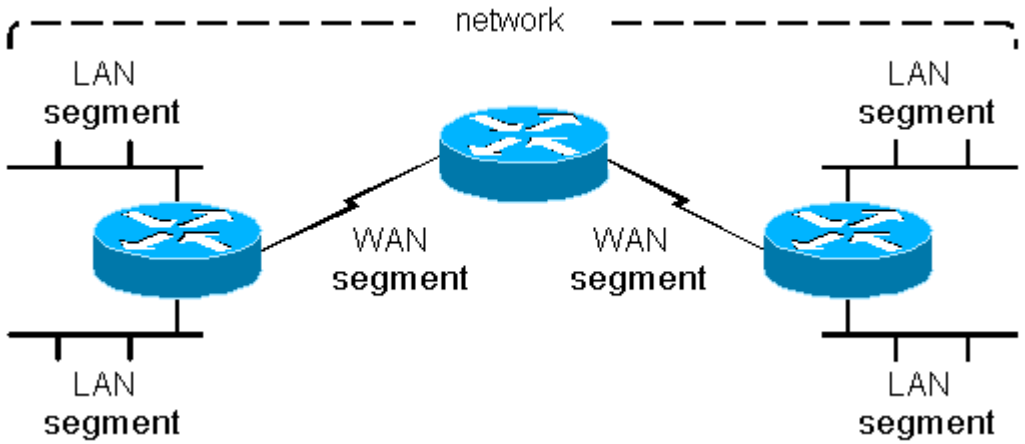

- To be more precise, these terms are used to describe physical connectivity.
- Each physical segment has a separate logical IP network or subnet.

AVAYA

Formula to determine number of hosts on a given network

- Given that there are N host bits in an address, the number of hosts for that network is $2^N - 2$. Two addresses are subtracted for the network address and the broadcast address.
- 8 host bits: $2^8 - 2 = 254$ hosts
- 16 host bits: $2^{16} - 2 = 65534$ hosts
- 24 host bits: $2^{24} - 2 = 16777214$ hosts
- As this tutorial progresses, it will become apparent how networks are typically sized so that there is a manageable number of hosts.

AVAYA

Public addresses

- Most IP addresses are public addresses. Public addresses are registered as belonging to a specific organization.
- Internet Service Providers (ISP) and extremely large organizations in the U.S. obtain blocks of public addresses from the American Registry for Internet Numbers (ARIN http://www.arin.net). Other organizations obtain public addresses from their ISPs.
- There are ARIN counterparts in other parts of the world, and all of these regional registration authorities are subject to the global Internet Assigned Numbers Authority (IANA http://www.iana.org).
- Public IP addresses are routed across the Internet, so that hosts with public addresses may freely communicate with one another globally.
- No organization is permitted use public addresses that are not registered with that organization!

AVAYA

Private addresses

- RFC 1918 designates the following as private addresses.
  - Class A range: 10.0.0.0 through 10.255.255.255.
  - Class B range: 172.16.0.0 through 172.31.255.255.
  - Class C range: 192.168.0.0 through 192.168.255.255.
- Private addresses may be used by any organization, without any requirement for registration.
- Because private addresses are ambiguous - can't tell where they're coming from or going to because anyone can use them - private addresses are not permitted to be routed across the Internet.
- ISPs block private addresses from being routed across their infrastructure.
- Note: The use of private addresses, network address translation (NAT), and proxy servers solved the IP address shortage problem for the short and medium terms. The projected long-term solution is IPv6. These topics will not be discussed here.

AVAYA

Classful IP Addressing and its Shortcomings

Three main classes

- Class A networks
  - First octet values range from **1 through 126**.
  - First octet starts with bit **0**.
  - Network mask is **8 bits**, written **/8** or **255.0.0.0**.
  - **1.0.0.0 through 126.0.0.0** are class A networks with **16777214** hosts each.

- Class B networks
  - First octet values range from **128 through 191**.
  - First octet starts with binary pattern **10**.
  - Network mask is **16 bits**, written **/16** or **255.255.0.0**.
  - **128.0.0.0 through 191.255.0.0** are class B networks, with **65534** hosts each.

- Class C networks
  - First octet values range from **192 through 223**.
  - First octet starts with binary pattern **110**.
  - Network mask is **24 bits**, written **/24** or **255.255.255.0**.
  - **192.0.0.0 through 223.255.255.0** are class C networks, with **254** hosts each.

AVAYA

Two additional classes, and reserved addresses

- Class D addresses
  - First octet values range from **224 through 239**.
  - First octet starts with binary pattern **1110**.
  - Class D addresses are multicast addresses, which will not be discussed in this tutorial.

- Class E addresses
  - Essentially everything that's left.
  - Experimental class, which will not be discussed in this tutorial.

- Reserved addresses
  - 0.0.0.0 is the default IP address, and it is used to specify a default route. The default route will be discussed later.
  - Addresses beginning with 127 are reserved for internal loopback addresses. It is common to see 127.0.0.1 used as the internal loopback address on many devices. Try pinging this address on a PC or Unix station.

AVAYA

The need to improve IP addressing efficiency

- As IP networking and internetworking progressed, it became very apparent that class A and B networks were simply too large.
- 254 hosts on one network segment is manageable, but 65534 hosts or more on a single network segment is difficult to manage.
  - This would result in class A and B networks not being fully utilized, meaning that not all the host addresses would get used.
  - Or it would result in more hosts being put onto a single network segment than could reasonably be managed.
- For these and other reasons, there was a need to improve the efficiency of IP addressing. That is, to provide a way to limit the number of host addresses per network segment to what is actually needed, regardless of the network class.
- This need was met progressively through the conceptions of subnet masks, variable-length subnet masks, and classless inter-domain routing.

AVAYA

Subnet Masks

Extending the classful network mask

- Subnet masks are used to make classful networks more manageable and efficient, by creating smaller subnets and reducing the number of host addresses per subnet to what is actually required.
- Subnet masks were first used on class boundaries.
- Example
  - Take class A network 10.0.0.0 with network mask 255.0.0.0.
  - Add additional 8 subnet bits to network mask.
  - New subnet mask is 255.255.0.0.
  - New subnets are 10.0.0.0, 10.1.0.0, 10.2.0.0, and so on with 65534 host addresses per subnet. Still too many hosts per subnet.
- Example
  - Take class A network 10.0.0.0 with network mask 255.0.0.0.
  - Add additional 16 subnet bits to network mask.
  - New subnet mask is 255.255.255.0
  - New subnets are 10.0.0.0, 10.0.1.0, 10.0.2.0, …, 10.1.0.0, 10.1.1.0, 10.1.2.0, …, 10.2.0.0, 10.2.1.0, 10.2.2.0, and so on with 254 host addresses per subnet.

AVAYA

Subnetting continued

- Example
  - Take class B network 172.16.0.0 with network mask 255.255.0.0.
  - Add additional 8 subnet bits to network mask.
  - New subnet mask is 255.255.255.0
  - New subnets are 172.16.0.0, 172.16.1.0, 172.16.2.0, and so on with 254 host addresses per subnet.

- As shown in these examples…
  - A class A network can be subnetted to create 256 (2^8) /16 subnets.
  - A class A network can be subnetted to create 65536 (2^{16}) /24 subnets.
  - A class B network can be subnetted to create 256 (2^8) /24 subnets.

- Note: Technically there really is no such thing as a classful subnet or classful subnet mask. However, terms such as “class C subnet” and “class C subnet mask” are used routinely to describe a class A or B network that has been subnetted with a 24-bit mask.

AVAYA

Terminology check

- By now it should be apparent that the terms network and subnet technically mean two different things.
- However, it is both common and somewhat accepted to use the terms interchangeably in casual communication.
- If there is a more general term of the two, it is network.

- It should also be apparent by now that the terms network mask and subnet mask technically mean two different things.
- But again, it is both common and somewhat accepted to use the terms interchangeably in casual communication.
- The term mask is a general term that is commonly used because of its ambiguity.

- The proper terms are used more consistently in formal communication.

AVAYA

Variable-length Subnet Masks (VLSM)

Improving the efficiency of subnet masks

- VLSM removed the class boundary restriction of traditional subnet masks.
- With VLSM a network of any class can be subnetted to almost any size.
- The next series of slides shows examples of VLSM.

AVAYA

Successive subnets w/ a 26-bit subnet mask

- The next four slides show a class A network subnetted with a 26-bit subnet mask.
- Successive subnets are shown; that is, the network addresses are shown in sequence.
- These slides show that in the same address space as a /24 network or subnet with 254 host addresses, four smaller /26 subnets can be created, each with 62 host addresses.

AVAYA

1. 26-bit subnet mask (255.255.255.192)

Network address of a subnet w/ 26-bit subnet mask 255.255.255.192

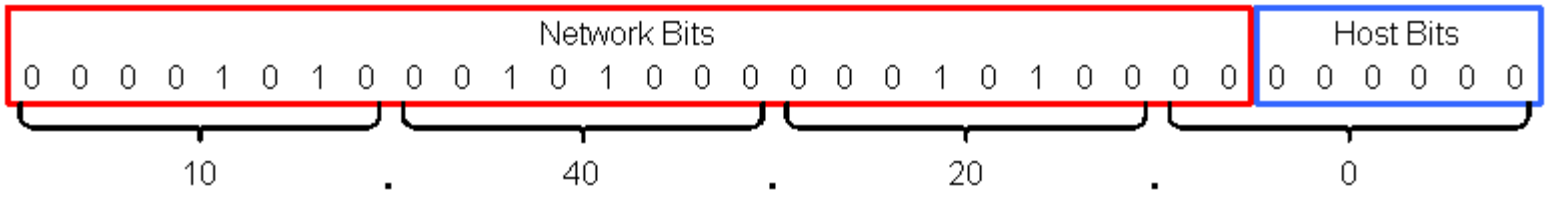

First host address for this subnet

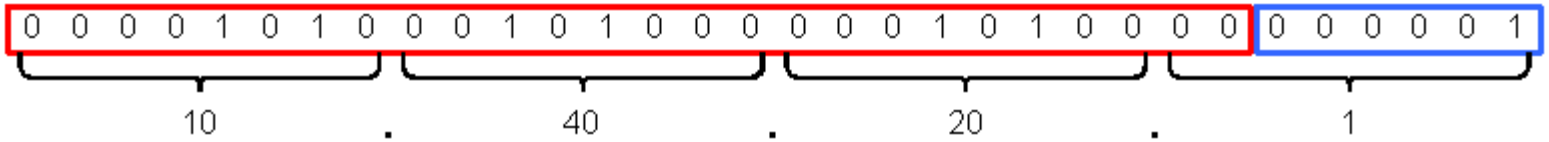

Second host address for this subnet

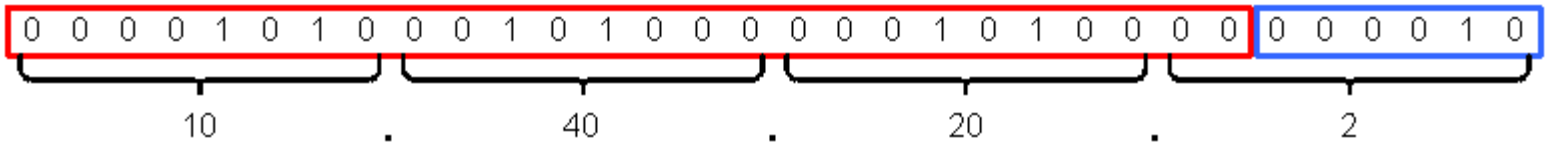

Last host address for this subnet

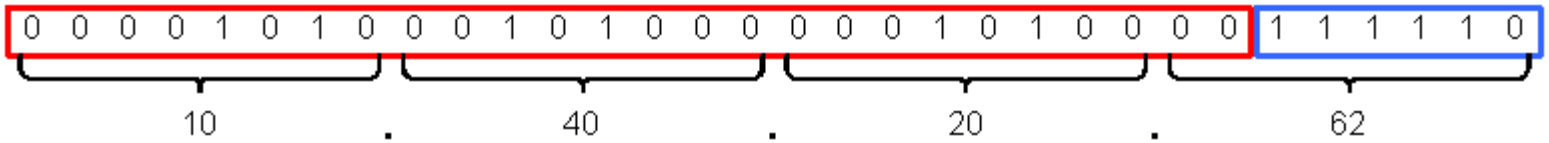

Broadcast address for this subnet

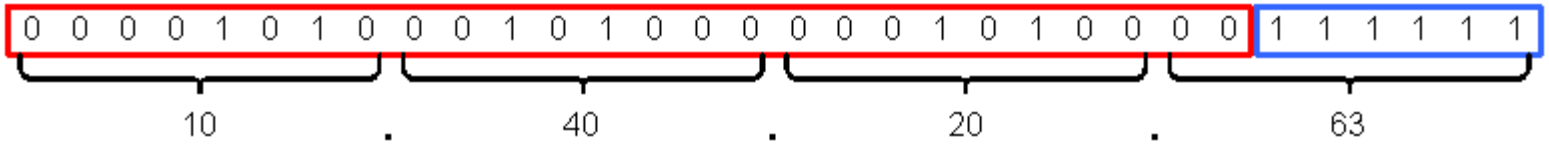

AVAYA

2. 26-bit subnet mask (255.255.255.192)

Network address of a subnet w/ 26-bit subnet mask 255.255.255.192

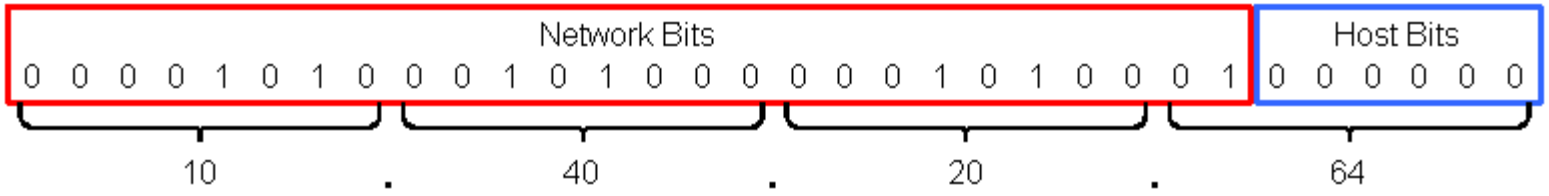

First host address for this subnet

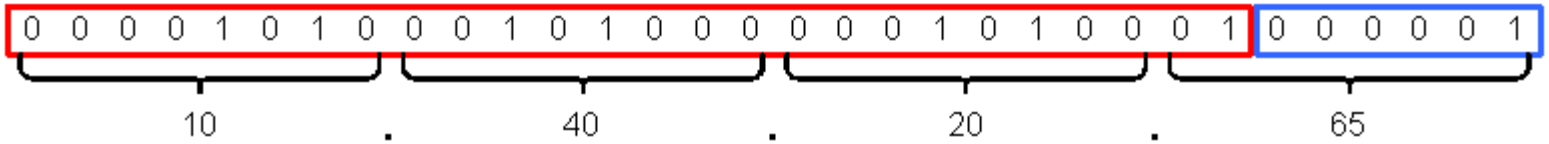

Second host address for this subnet

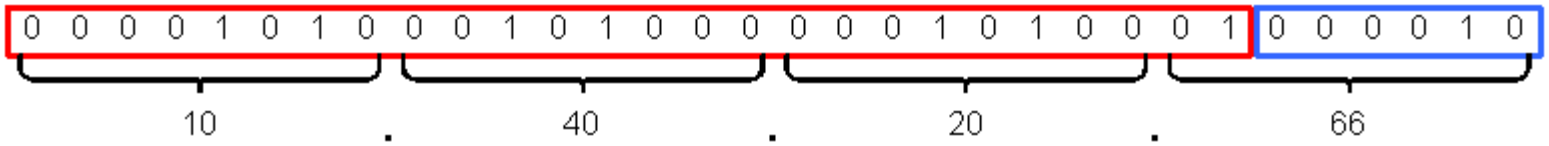

Last host address for this subnet

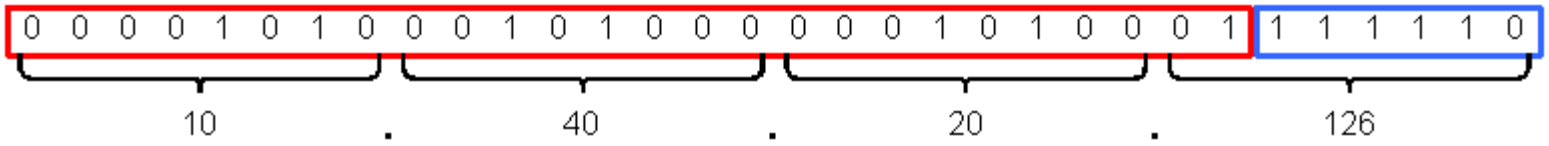

Broadcast address for this subnet

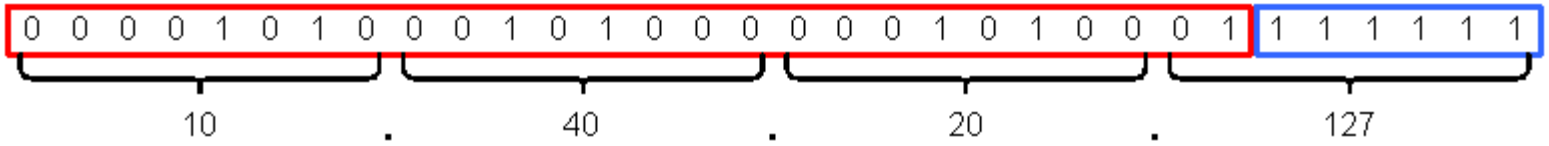

3. 26-bit subnet mask (255.255.255.192)

Network address of a subnet w/ 26-bit subnet mask 255.255.255.192

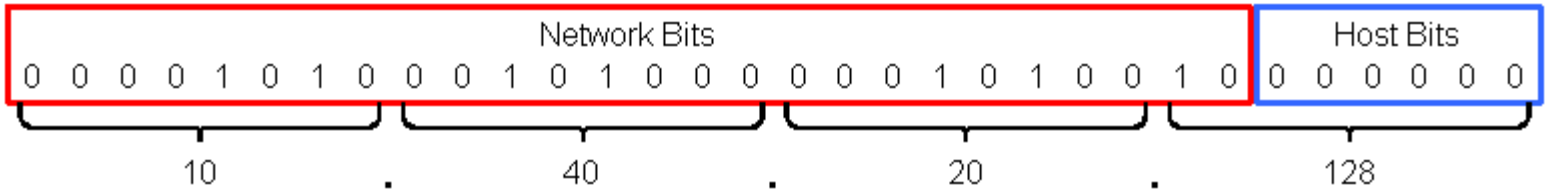

First host address for this subnet

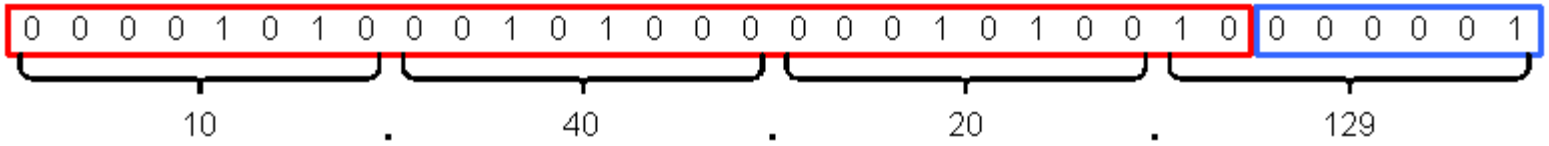

Second host address for this subnet

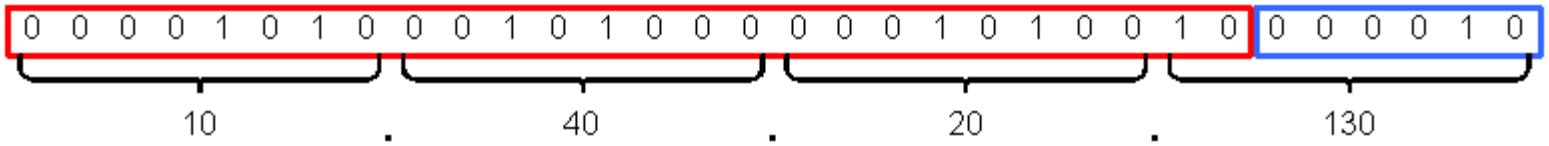

Last host address for this subnet

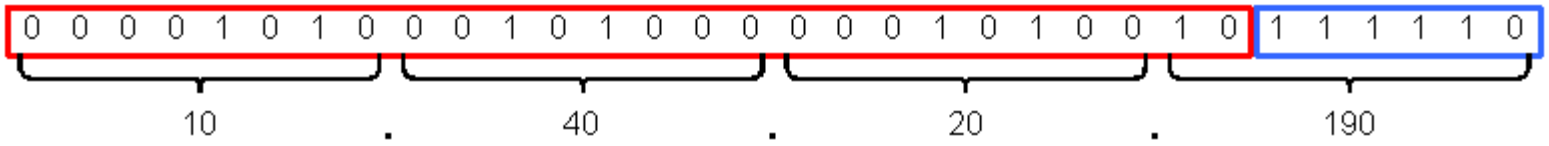

Broadcast address for this subnet

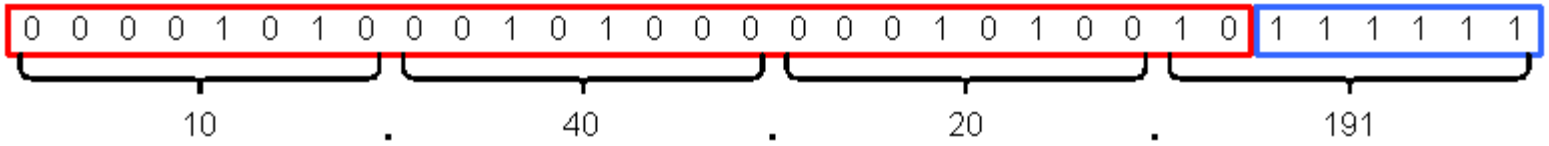

AVAYA

4. 26-bit subnet mask (255.255.255.192)

Network address of a subnet w/ 26-bit subnet mask 255.255.255.192

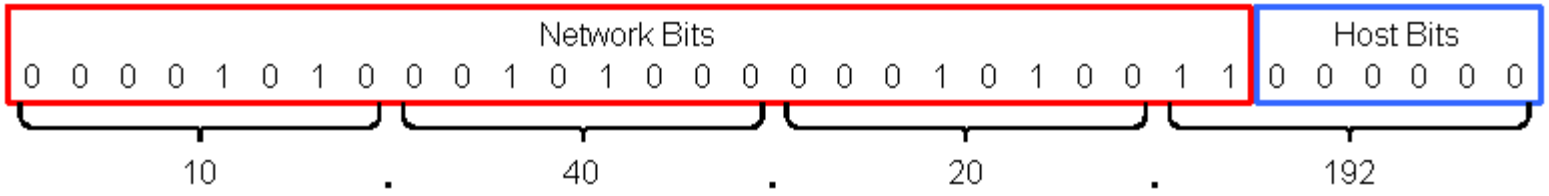

First host address for this subnet

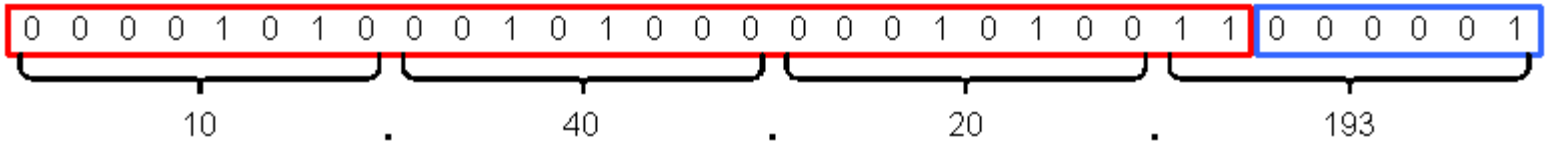

Second host address for this subnet

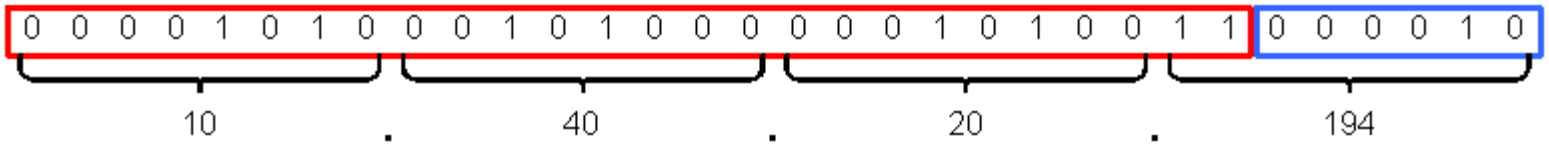

Last host address for this subnet

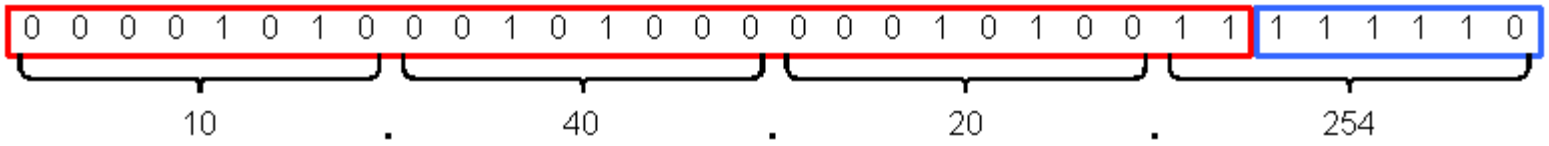

Broadcast address for this subnet

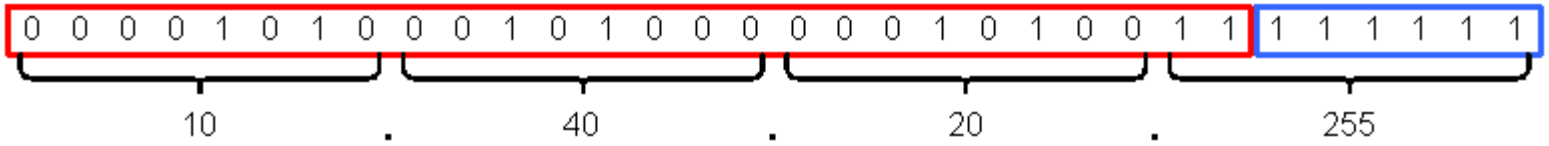

Successive subnets w/ a 23-bit subnet mask

- The next four slides show the same class A network as in the previous slides, but with a 23-bit subnet mask.
- Successive subnets are shown; that is, the network addresses are shown in sequence.
- A /23 subnet has 510 host addresses, and this is arguably the largest number of hosts to practically put on a LAN segment.
  - The next larger subnet (22-bit subnet mask) has 1022 host addresses.
  - As explained in the companion document, “LANs and VLANs: A Simplified Tutorial,” all IP hosts must transmit ARP broadcasts.
  - In addition, IP hosts participate in other exchanges that require broadcasts, and some applications are very broadcast-intensive.
  - A thousand or so hosts transmitting frequent broadcasts to all other hosts on a LAN segment can be very taxing on network and host devices.
  - There are also other reasons, in addition to broadcasts, why a thousand endpoints on one LAN segment would not be desirable.

1. 23-bit subnet mask (255.255.254.0)

Network address of a subnet w/ 23-bit subnet mask 255.255.254.0

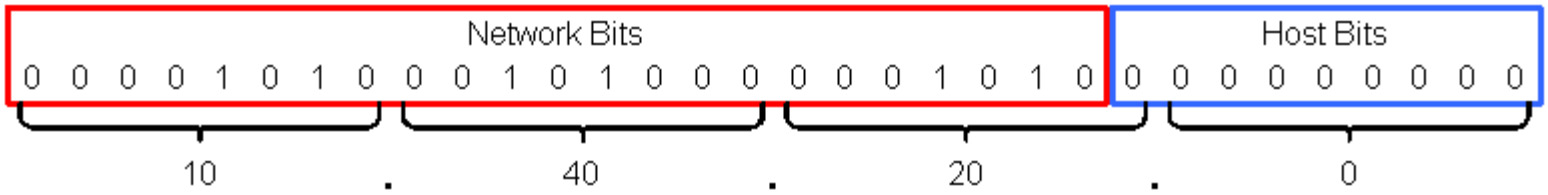

First host address for this subnet

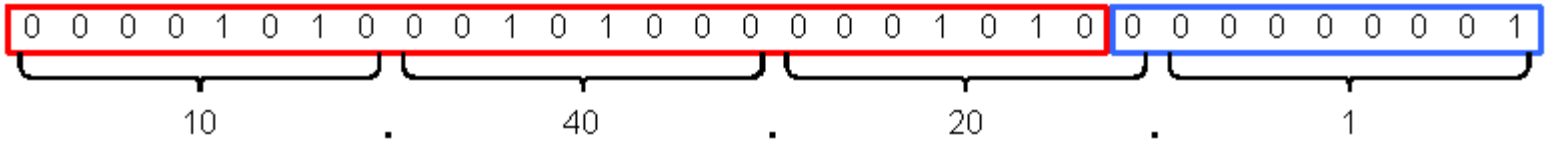

Second host address for this subnet

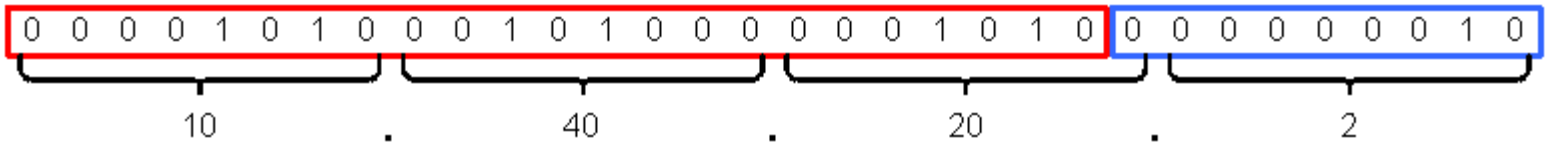

Last host address for this subnet

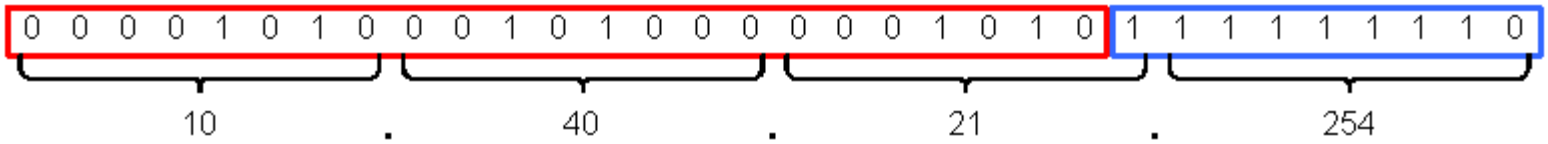

Broadcast address for this subnet

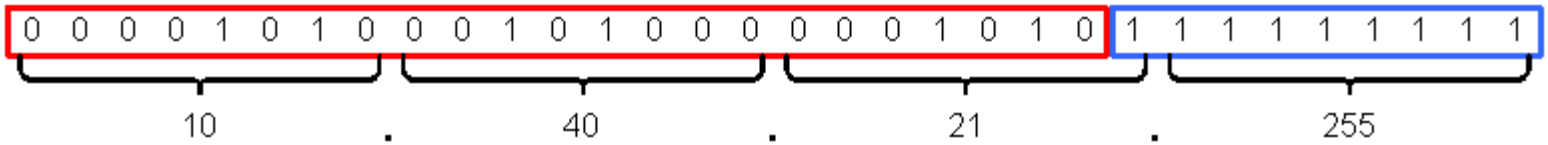

AVAYA

2. 23-bit subnet mask (255.255.254.0)

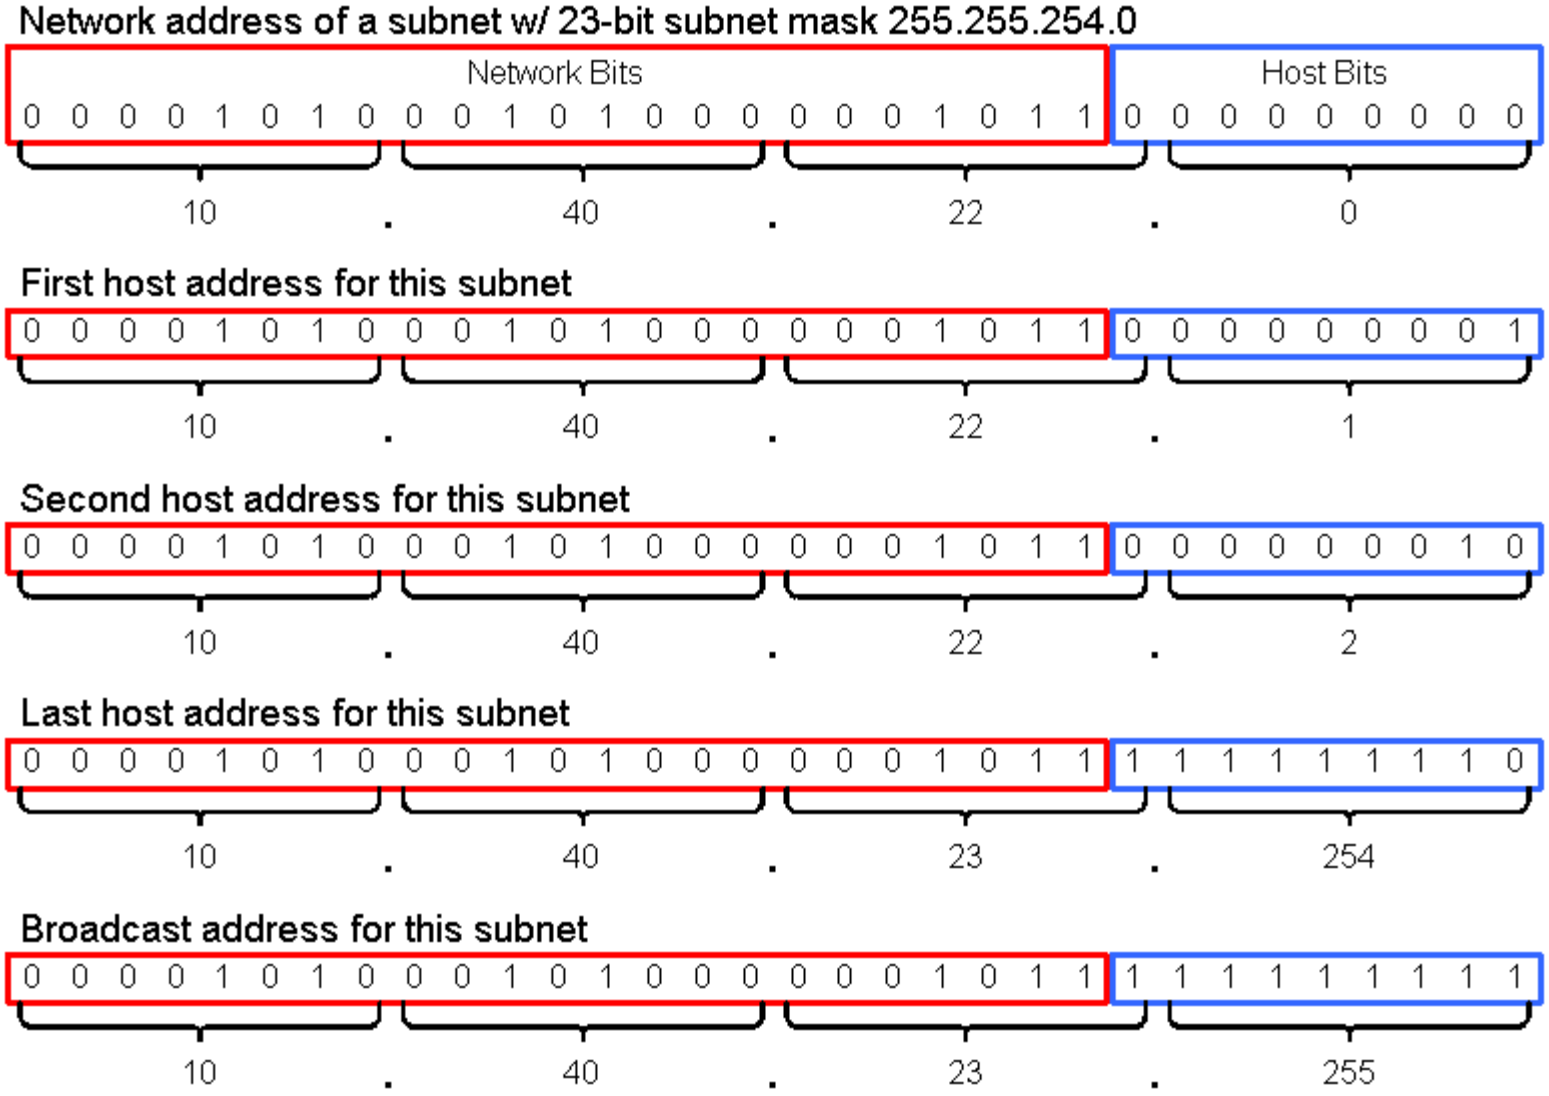

3. 23-bit subnet mask (255.255.254.0)

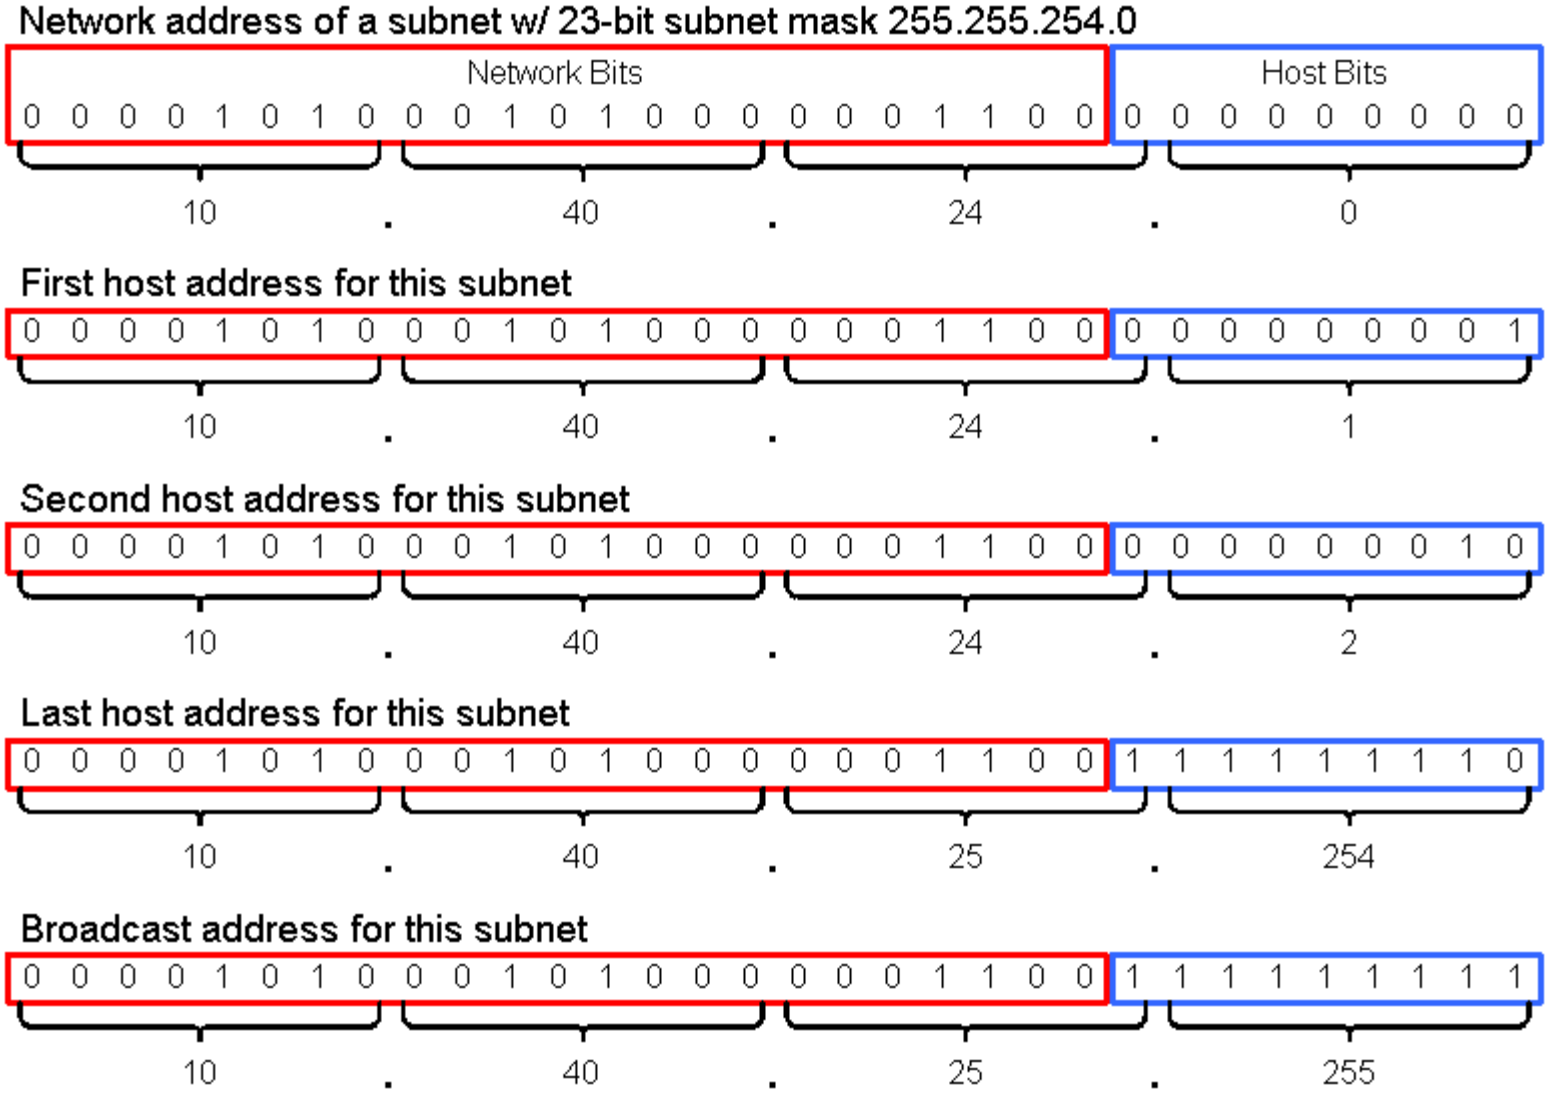

4. 23-bit subnet mask (255.255.254.0)

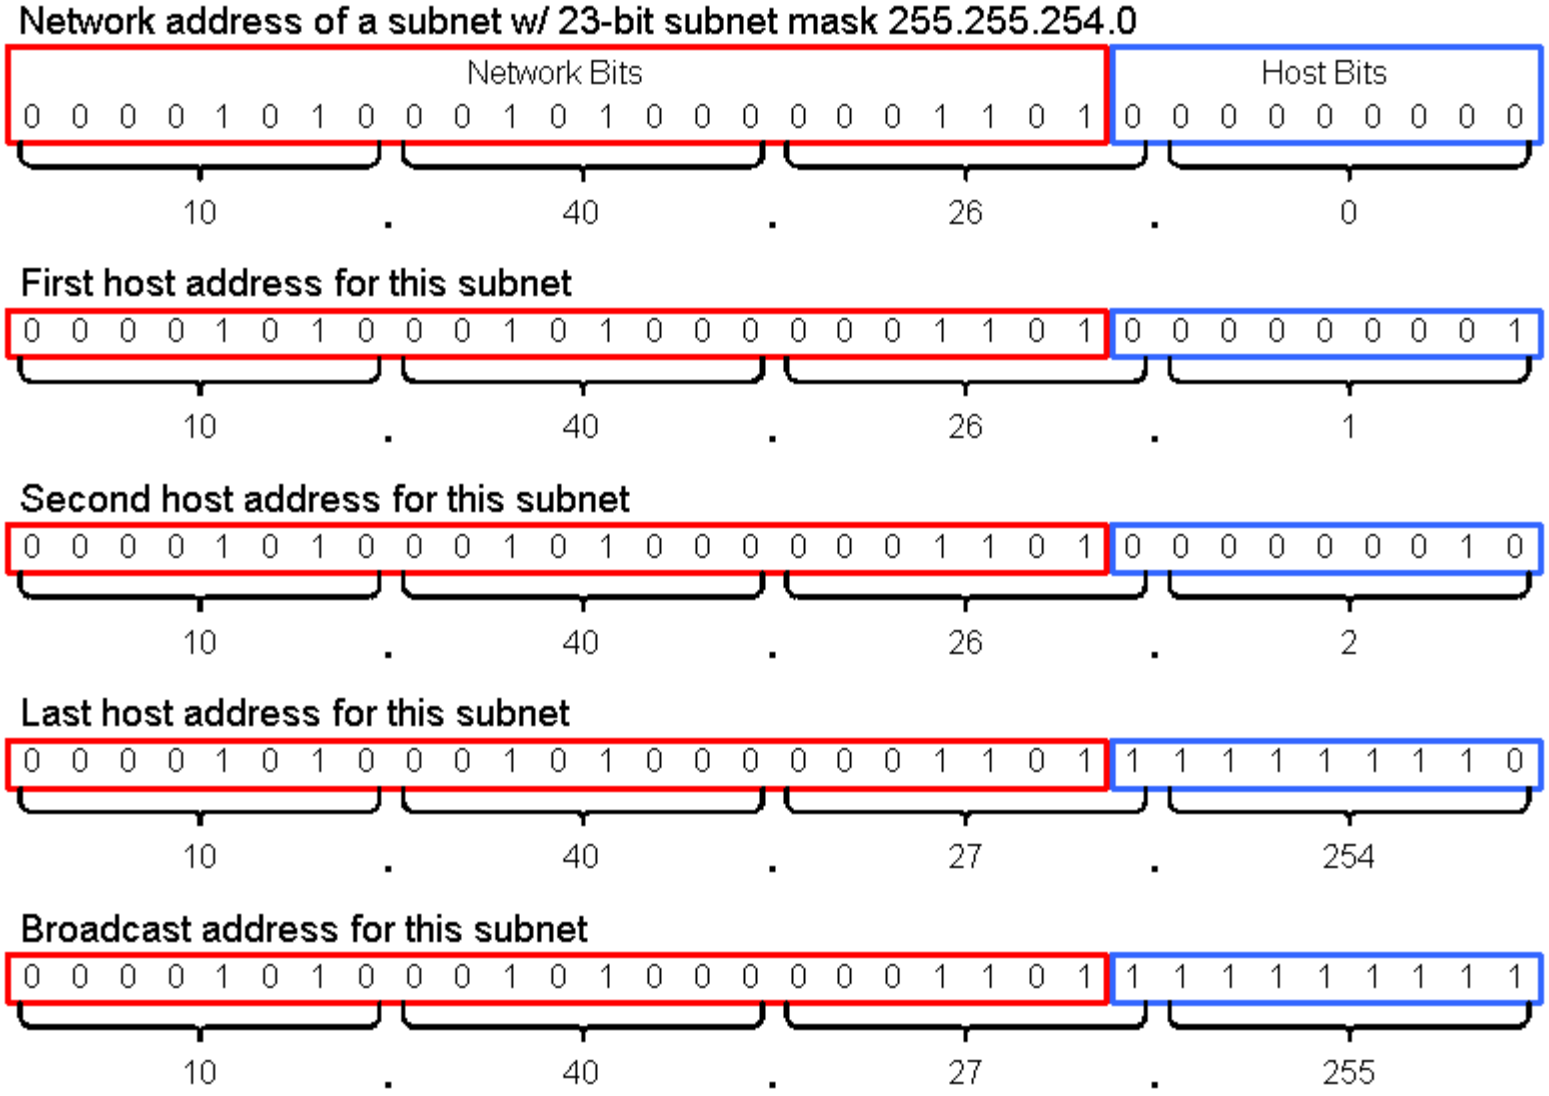

Successive subnets w/ a 30-bit subnet mask

- The next four slides show the same class A network as in the previous slides, but with a 30-bit subnet mask.
- Successive subnets are shown; that is, the network addresses are shown in sequence.
- A /30 subnet has two host addresses, and this is the smallest permissible subnet.
  - A /31 subnet has only two addresses because there is only one host bit.
  - This would allow for a network address and a broadcast address and no host addresses.
- /30 subnets are commonly used for point-to-point links, such as WAN links, because there are only two hosts on the network segment.

AVAYA

1. 30-bit subnet mask (255.255.255.252)

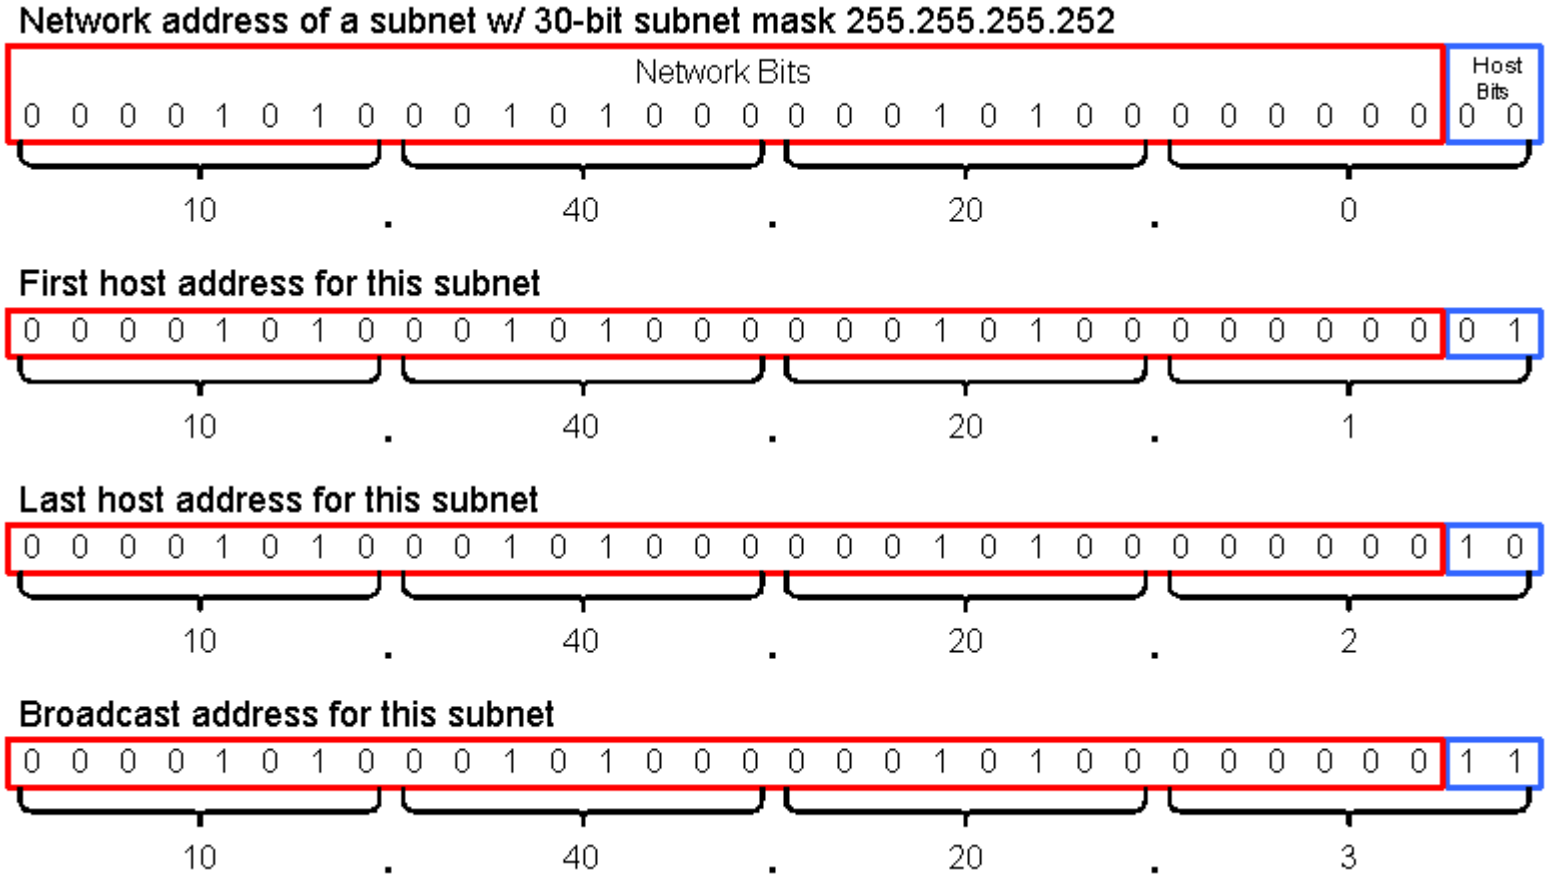

2. 30-bit subnet mask (255.255.255.252)

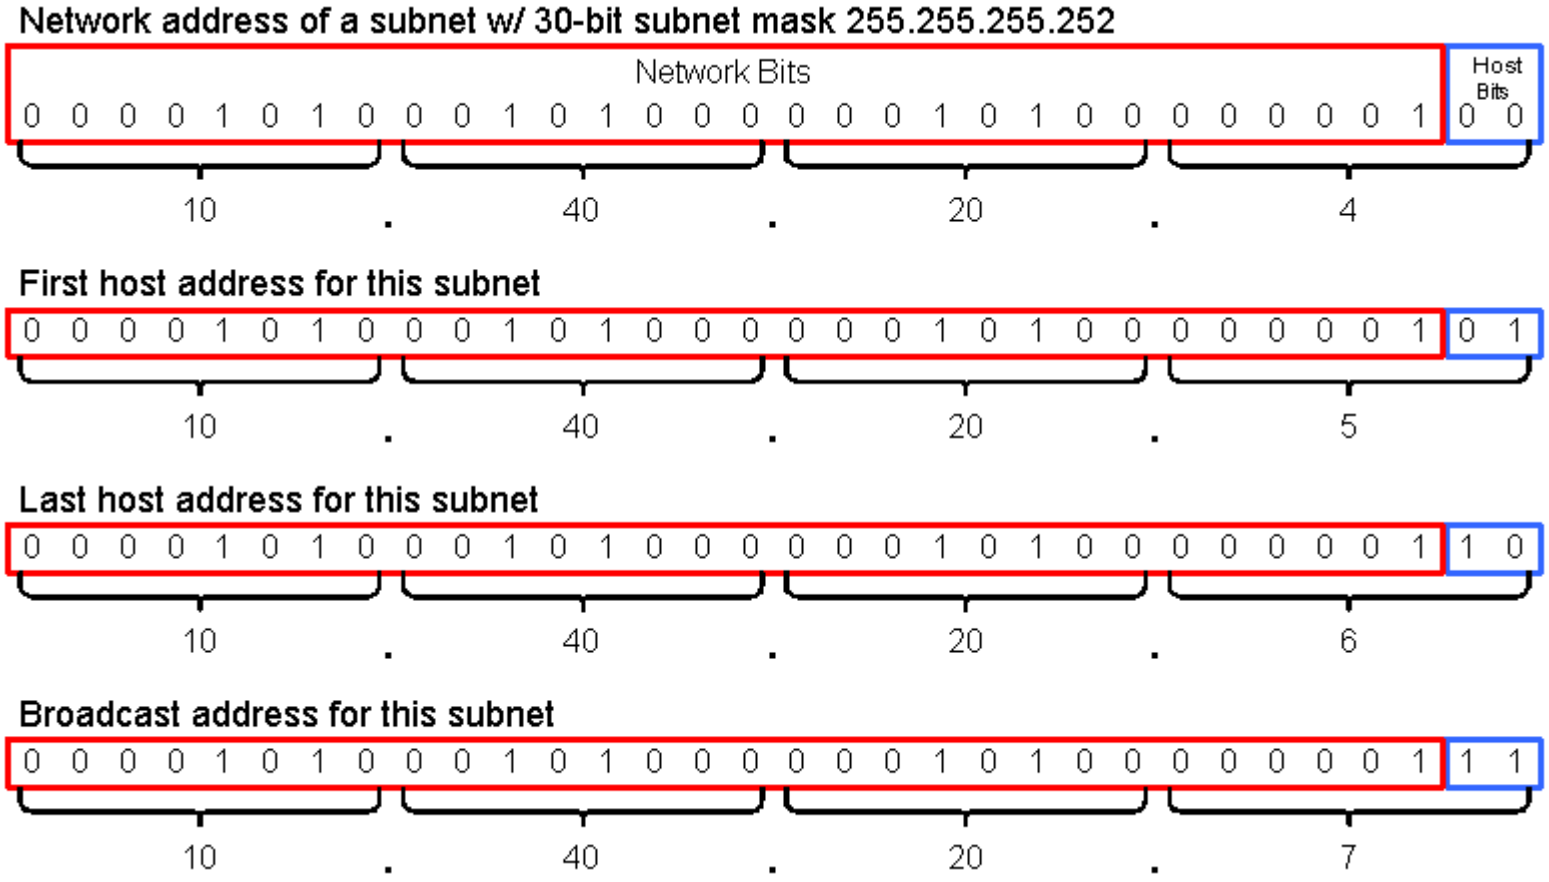

3. 30-bit subnet mask (255.255.255.252)

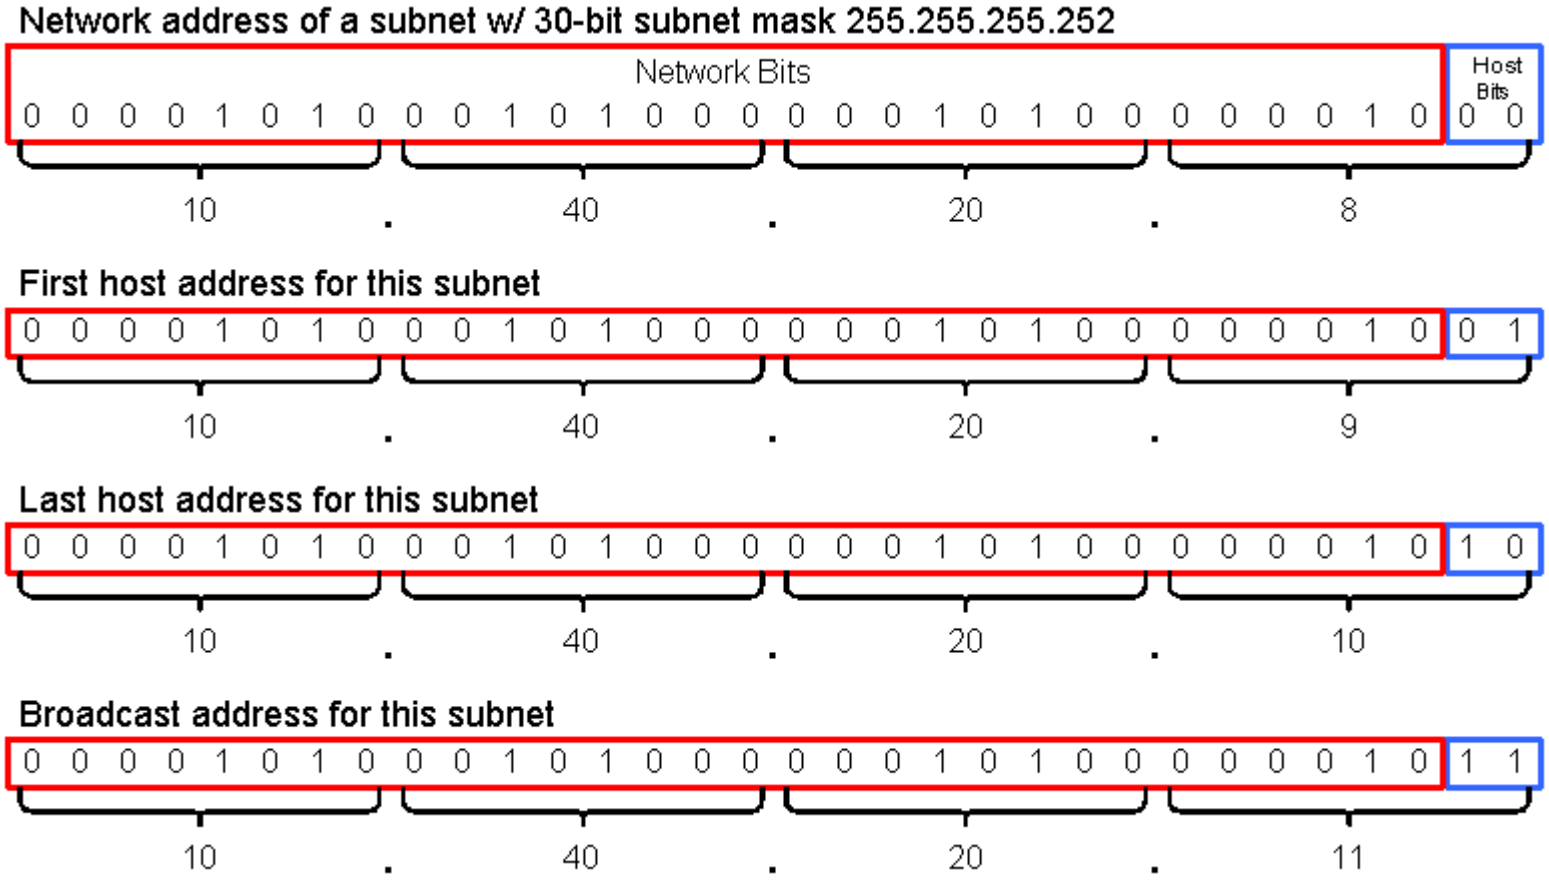

4. 30-bit subnet mask (255.255.255.252)

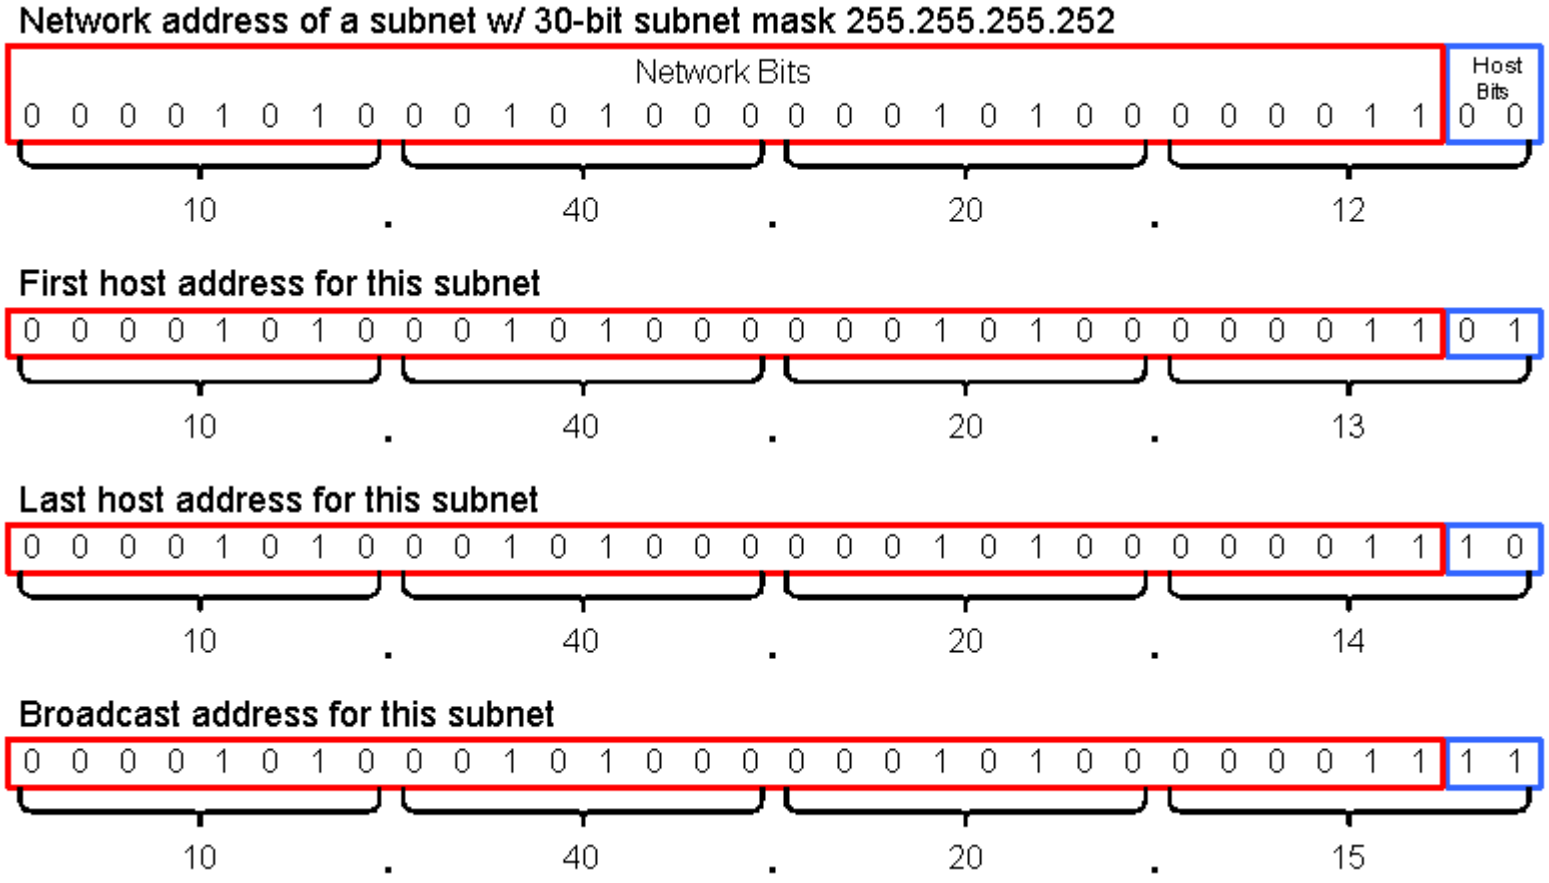

Classless Inter-domain Routing (CIDR)

Supernetting

- Very simply stated, CIDR is combining two or more classful networks to create a supernetwork.
- The most common use of CIDR for actual addressing is to combine two or more class C networks to create a /23 or /22 supernet.
- For example, the class C networks 192.168.32.0 and 192.168.33.0 could be combined to create 192.168.32.0/23.
- The class C networks 192.168.34.0 and 192.168.35.0 could be combined to create 192.168.34.0/23.
- These two examples are illustrated in the following slides.

AVAYA

1. Supernet w/ 23-bit mask (255.255.254.0)

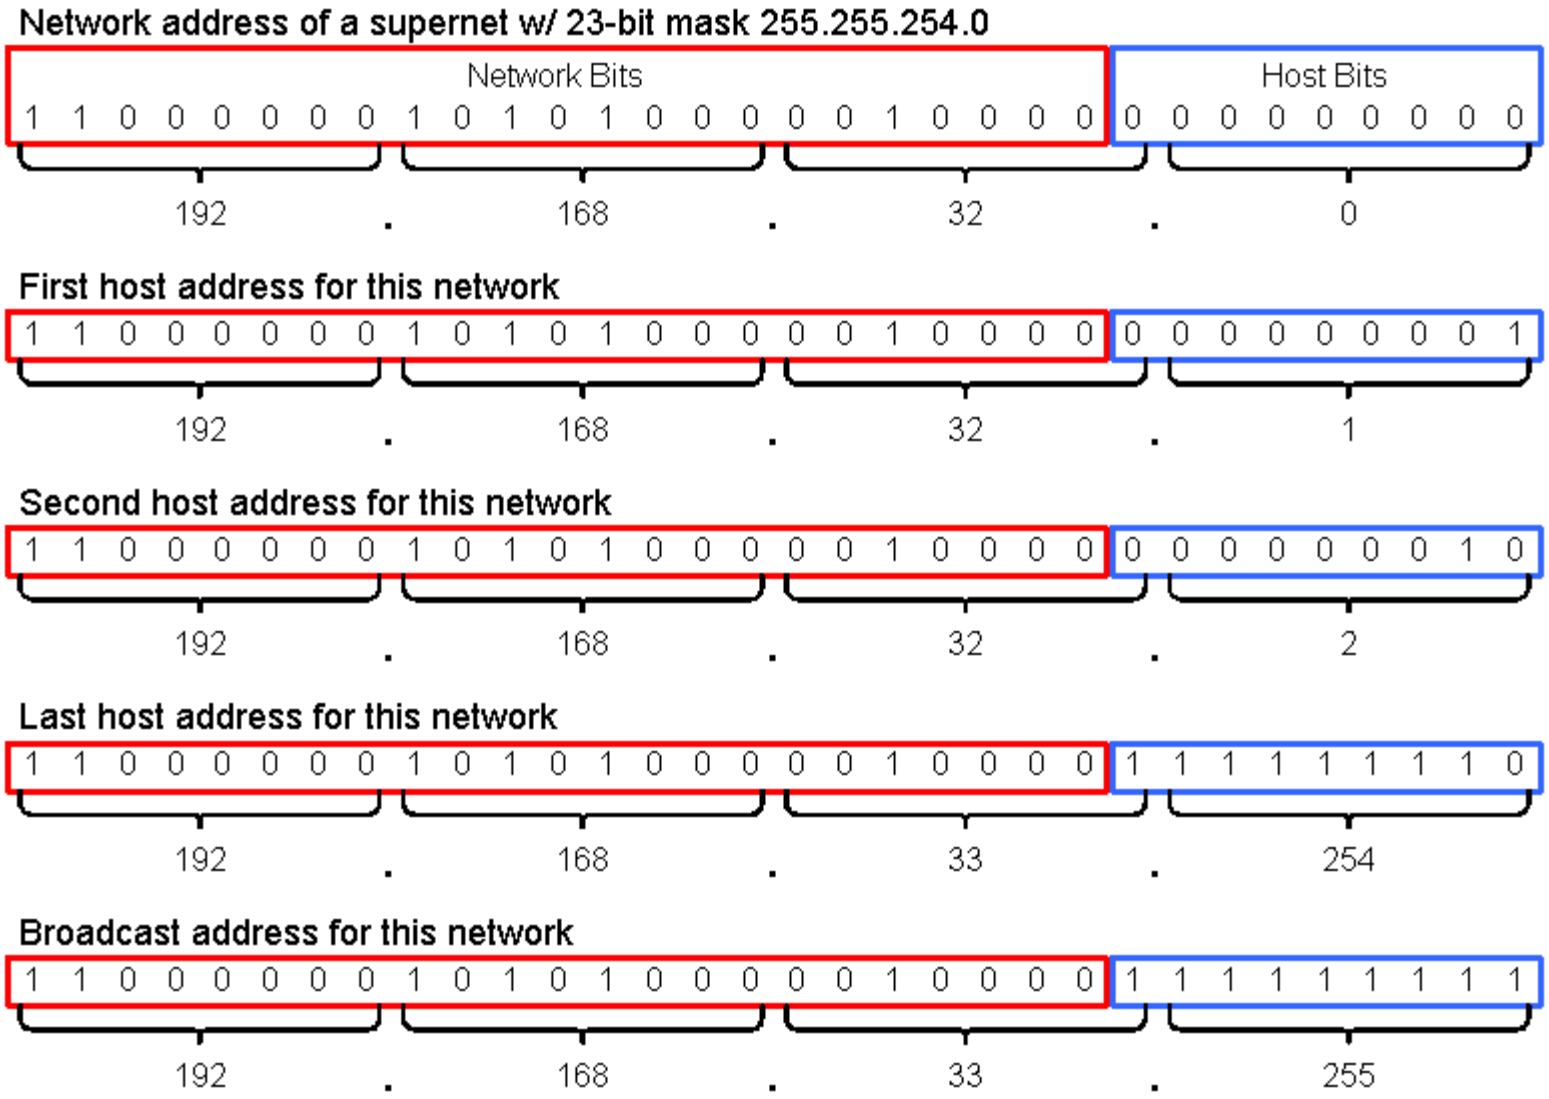

2. Supernet w/ 23-bit mask (255.255.254.0)

Network address of a supernet w/ 23-bit mask 255.255.254.0

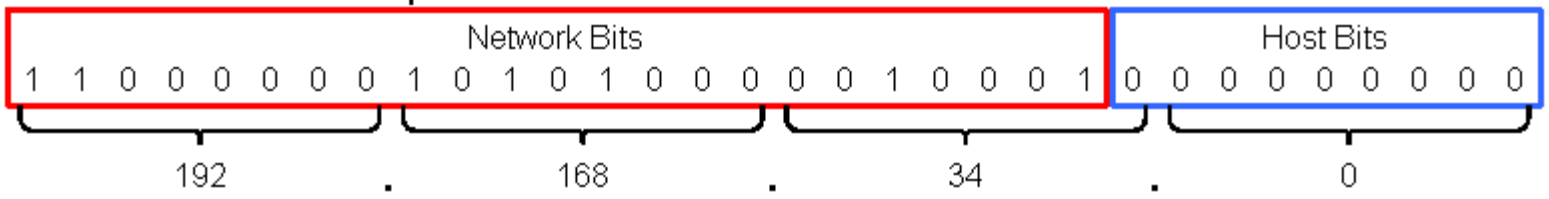

First host address for this network

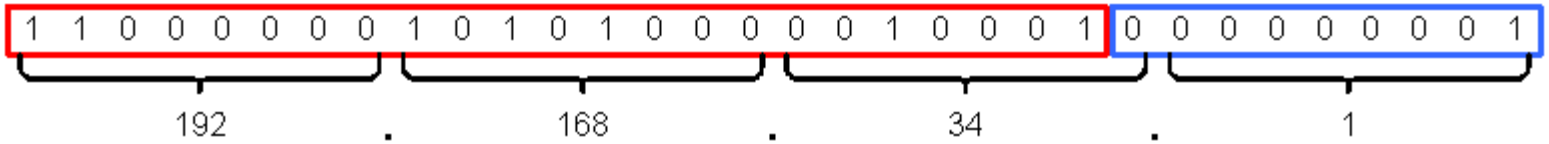

Second host address for this network

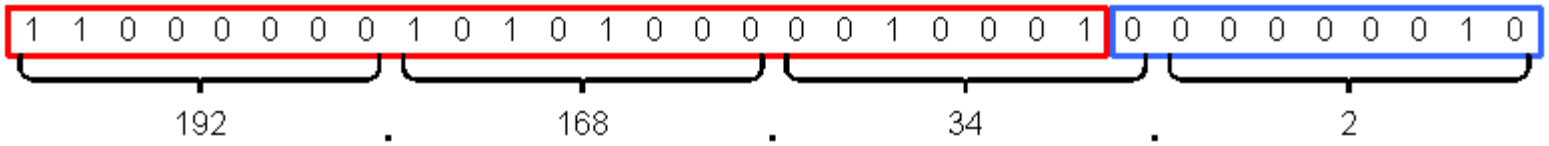

Last host address for this network

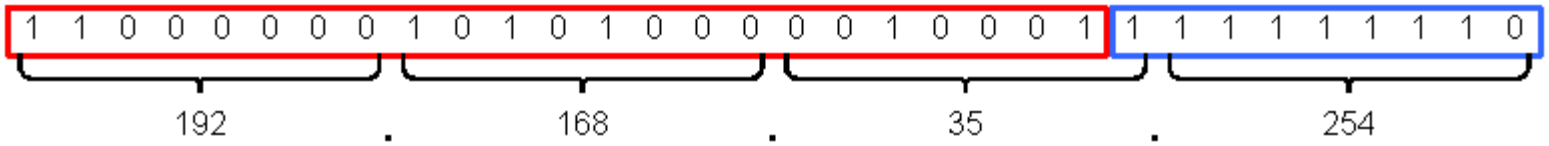

Broadcast address for this network

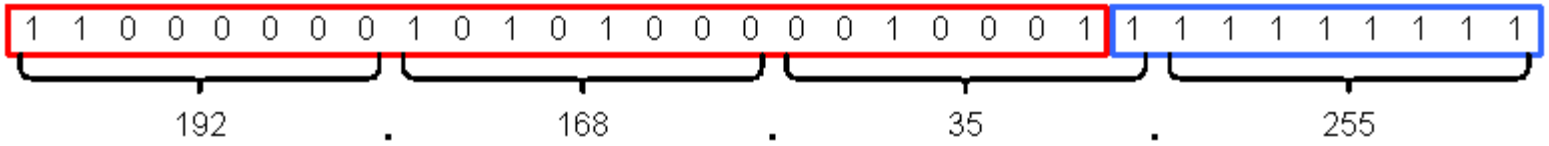

AVAYA

CIDR observations

- There really is no significant difference between the /23 supernet examples just shown and the /23 subnet examples shown previously.
  - Both networks utilize the same 23-bit mask.
  - Both networks have 510 host addresses.
  - The only difference is that a class A network was subnetted in one set of examples, and two class C networks were supernetted in the other.
- As with all things in IP addressing, supernets must fall on bit boundaries.
  - There must be a power of 2 number (2, 4, 8, …) of networks in a supernet.
    - A /23 supernet must include two class C networks.
    - A /22 supernet must include four class C networks.
  - For example, it is not possible to create a supernet that includes just 192.168.1.0 and 192.168.2.0. To include both these networks the mask must be 22 bits, and a 22-bit mask must also include the networks 192.168.0.0 and 192.168.3.0.
- Although the term supernet is used in this tutorial to explain CIDR, this term is not commonly used in casual communication. Instead, most will simply use the term network.

AVAYA

Routing and Routing Protocols

The IP header

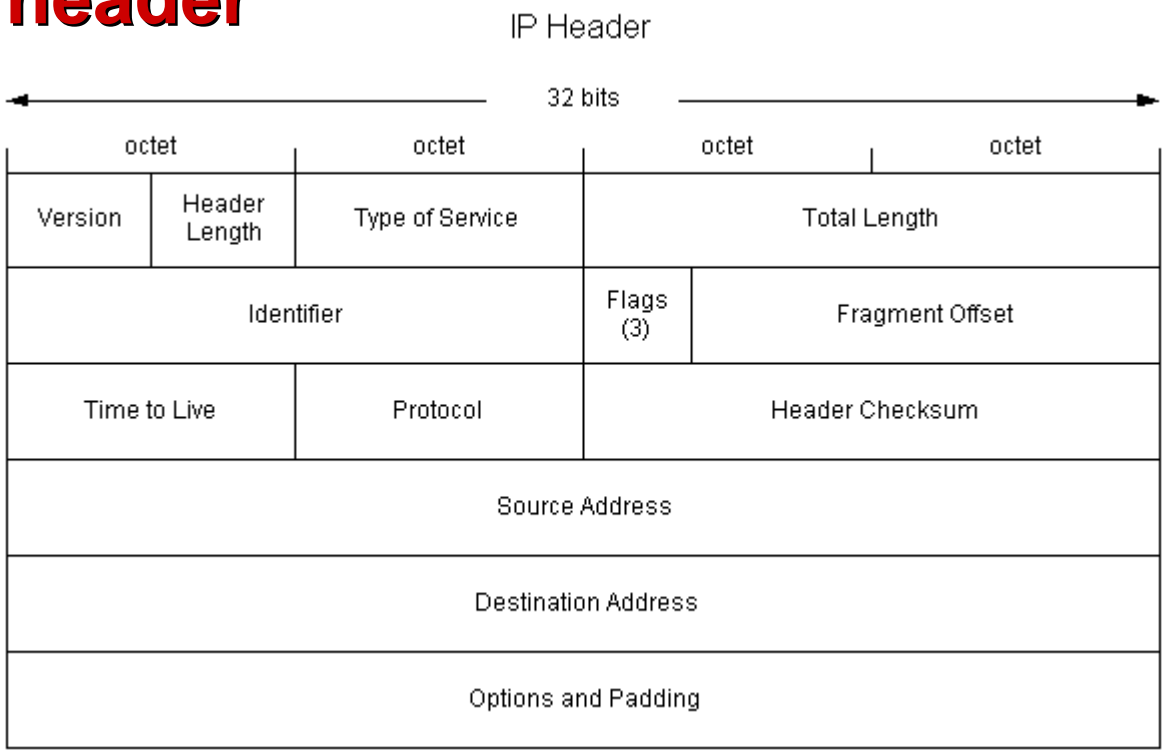

- This tutorial will not cover the entire IP header, but it is presented here for reference.
- For this section it is important to note that every IP packet has an IP header, and therefore has a source and destination IP address.
  - The source address belongs to the host that originated the IP packet.
  - The destination address belongs to the host, or hosts if the packet is a broadcast, that is to receive the IP packet.

AVAYA

IP routing defined

- IP routing is the process of delivering IP packets from one IP network to another.

- Routing is the core function of a router.
  - To forward an IP packet from one IP network to another, a router must know how to route that packet to its destination.
  - If a router is directly connected to a packet's destination network, the router can independently forward the packet to that network.
  - If a router is not directly connected to a packet's destination network, the router must know the next-hop router to which the packet must be forwarded.

- IP hosts participate in routing to a much lesser degree.
  - Hosts on any given IP network can independently communicate with one another.
  - Hosts on different IP networks, however, must communicate through an IP gateway, or router.
  - A host must be aware of its gateway or gateways.

AVAYA

Anatomy of a route

- In general, a route consists of a destination and a gateway.

- If a host or router is only capable of classful routing, its route table will have two fields.

  destination gateway

  – Such a device can only route to classful networks, using the classful network mask.
  – Because there can be only one classful network mask based on the destination IP address, there is no need to explicitly specify the mask.

- If a host or router is capable of classless routing, its route table will have three fields.

  destination mask gateway

  – Such a device can route to both classful networks and classless subnets.
  – This requires that a mask be explicitly specified.

- Most devices are capable of classless routing.

AVAYA

Default address, route, and gateway

- The default address is 0.0.0.0. This address encompasses all IP addresses.
- The default address is used to specify a default route, which is the route to be taken when no other more specific route is available.
- For example:

destination	gateway
0.0.0.0	address-of-router-on-this-network

destination	mask	gateway
0.0.0.0	0.0.0.0	addr-of-router-on-this-net

- As shown here, the mask for a default route is 0.0.0.0.
- The gateway used for the default route is the default gateway.

AVAYA

Static routing

- Static routes are manually entered into a router or host.
- An administrator must know the internetwork layout and the paths that exist between networks.
- Then the administrator must program each router in the internetwork with the proper routes to get from any given network to any other network.
- The hosts obtain their routes manually or via DHCP.

AVAYA

Dynamic routing

- Dynamic routes are routes learned via one or more routing protocols.
- Routing protocols are used by routers to inform one another of the IP networks accessible to them.
- There are classful routing protocols, such as RIPv1, that do not transmit masks in their routing updates - the classful network mask is implied.
- There are also classless routing protocols, such as OSPF, that do transmit masks in their routing updates.
- Dynamic routing is much too complex a subject to be covered in this tutorial. Suffice it to say that OSPF is the most prevalent standard routing protocol today, and likely the most prevalent routing protocol overall.
- Routing protocols typically do not apply to hosts. Hosts obtain their routes by manual configuration or by DHCP.

AVAYA

Static vs. dynamic routing

- One difference between static routes and dynamic routes is that one is entered in manually and the other is learned and/or calculated dynamically.
- The big differentiator, however, is in how the routers adapt to sporadic changes in network topology caused by outages.
- A statically routed network has almost no way of adapting to temporary topology changes. But routing protocols are designed for this purpose.
- A key factor in designing networks and choosing a routing protocol is convergence time, which is the time it takes for the network as a whole to discover its topology and reach a steady state.
- In general, the shorter the convergence time the better. A network that converges quickly can better compensate for unexpected outages.

AVAYA

One last note about route tables

- In addition to destination, mask, and gateway, most routers have one other field in their route tables, which was not previously mentioned.
- This field is a source field, which indicates how a route was obtained, whether statically or by a routing protocol.
- Although the source of a route is very significant for network administration, it is not the focus of this tutorial.
- This tutorial is more concerned with the route itself, and not necessarily how it was obtained.

AVAYA

Routing example

- This example applies regardless of how the routes were learned.

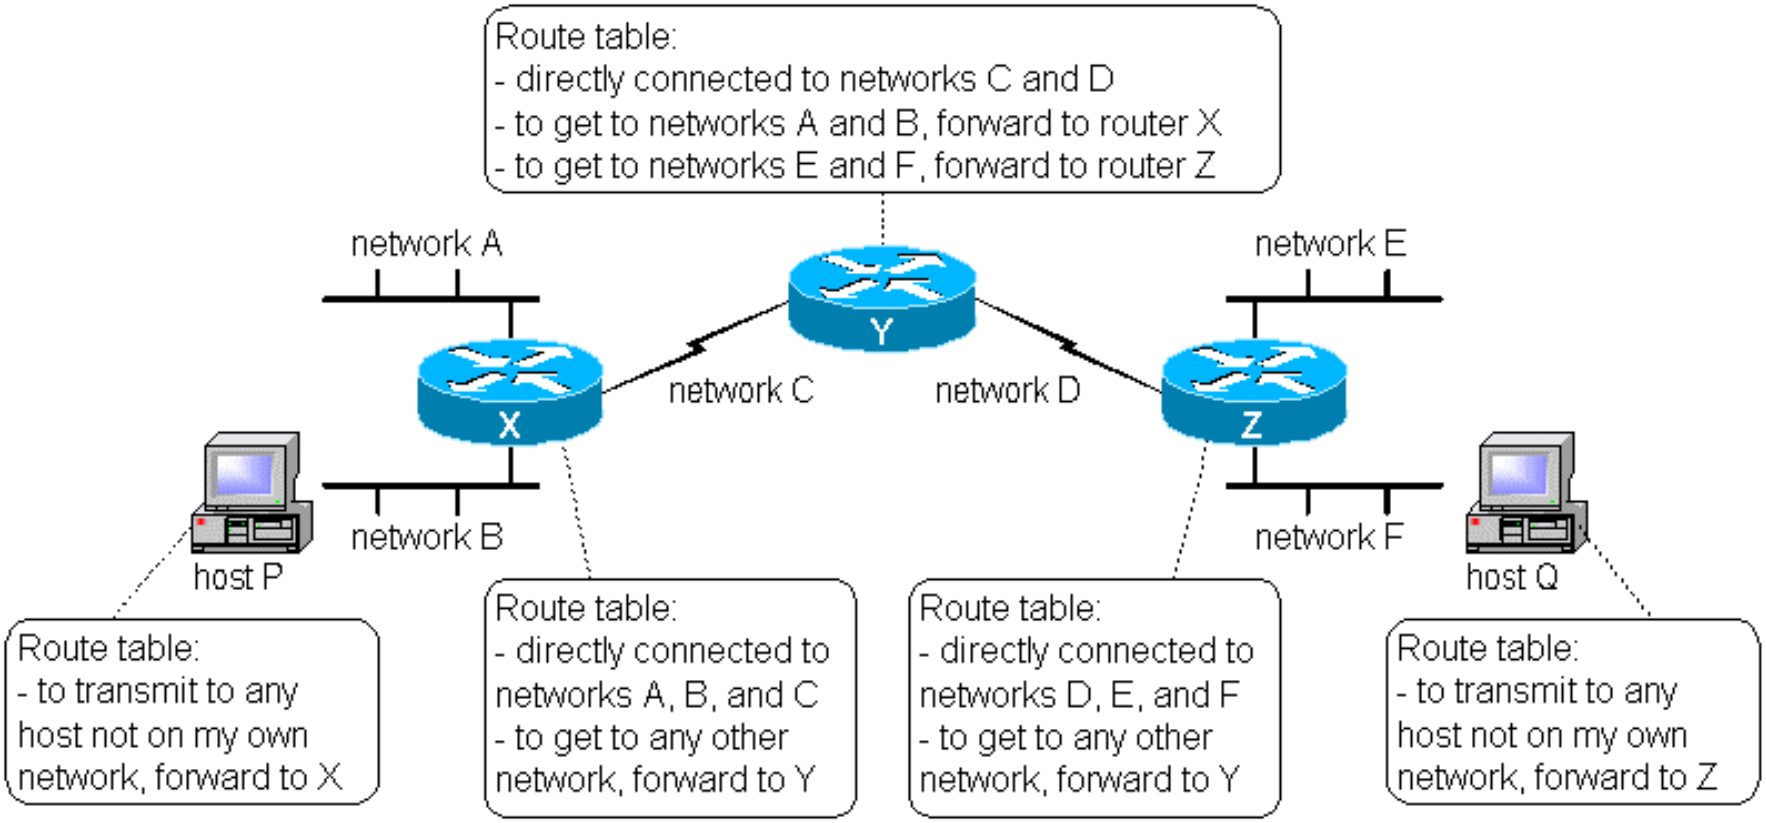

Routing analysis

- Hosts P and Q
  - Hosts P and Q have only one way to get off their networks, and that is to forward traffic to the router connected to their respective networks.
  - These hosts each have a default route to a default gateway - router X or Z.
- Router X
  - Router X has interfaces directly connected to networks A, B, and C, so it can route traffic between these networks w/o additional configurations.
  - This router has only one option to get to other networks, and that is to forward traffic to router Y. So it has a default route to router Y.
- Router Z
  - Router Z has interfaces directly connected to networks D, E, and F, so it can route traffic between these networks w/o additional configurations.
  - This router has only one option to get to other networks, and that is to forward traffic to router Y. So it has a default route to router Y.
- Router Y
  - Router Y automatically knows how to route traffic between networks C and D.
  - To get to networks A and B, this router forwards traffic to router X.
  - To get to networks E and F, this router forwards traffic to router Z.
  - One of these routes could be made the default route.

AVAYA

Routing example with IP addresses added

- The routing analysis in the previous slide still applies.

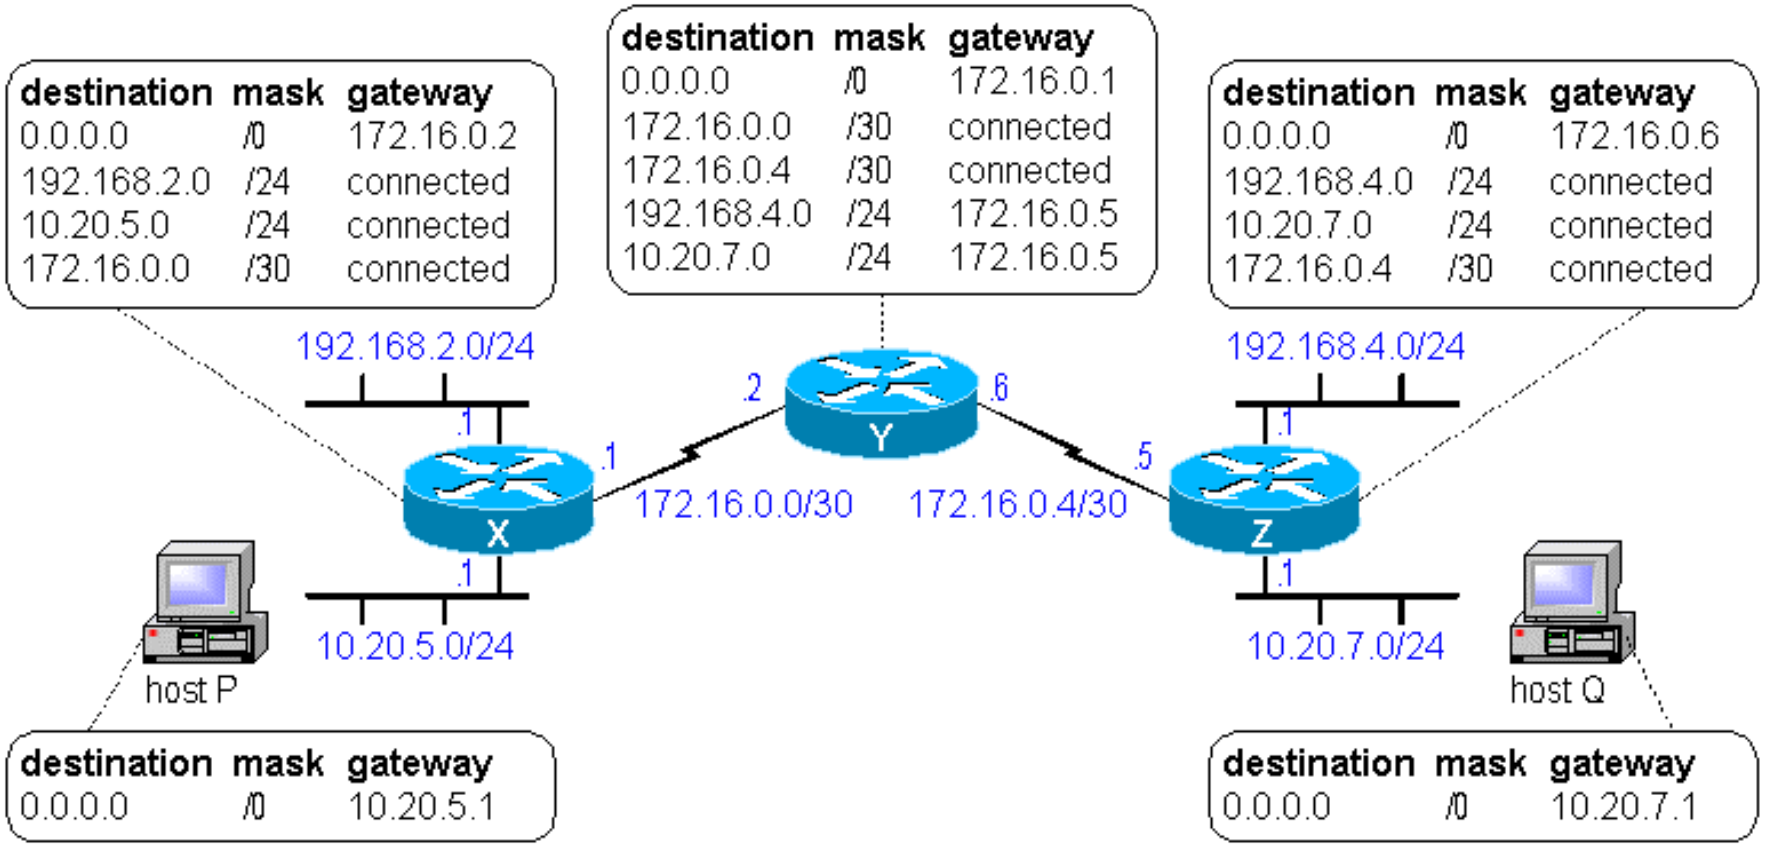

Routing observations

- In the preceding diagram, the network could not have worked with classful routing.

- Without subnet masks, router Y would have a route to class A network 10.0.0.0 via gateway 172.16.0.5.
  - This would result in all 10.x.x.x packets, that traverse router Y, being forwarded to router Z.
  - For example, router Z would forward destination 10.20.5.x packets to router Y, and router Y would return them to router Z.
  - This is because router Y has no way to distinguish between 10.20.5.0 and 10.20.7.0 without a subnet mask.
  - Destination 10.20.5.x packets within router X - that is, those sourced from 192.168.2.x - would not be affected by this phenomenon.

AVAYA

Route summarization

- Where VLSM and CIDR become truly powerful is in route summarization.
- This example shows 32 individual routes being summarized to a single route using a 19-bit mask, and 16 individual routes being summarized to a single route using a 20-bit mask.
- Route summarization results in greater routing efficiency, as each route adds a processor burden to the router.

destination	mask	gateway
0.0.0.0	/0	172.16.0.6
192.168.0.0	/24	172.16.0.2
192.168.1.0	/24	172.16.0.2
- - -		
192.168.31.0	/24	172.16.0.2
10.30.0.0	/24	172.16.0.2
10.30.1.0	/24	172.16.0.2
- - -		
10.30.15.0	/24	172.16.0.2

destination	mask	gateway
0.0.0.0	/0	172.16.0.6
192.168.0.0	/19	172.16.0.2
10.30.0.0	/20	172.16.0.2

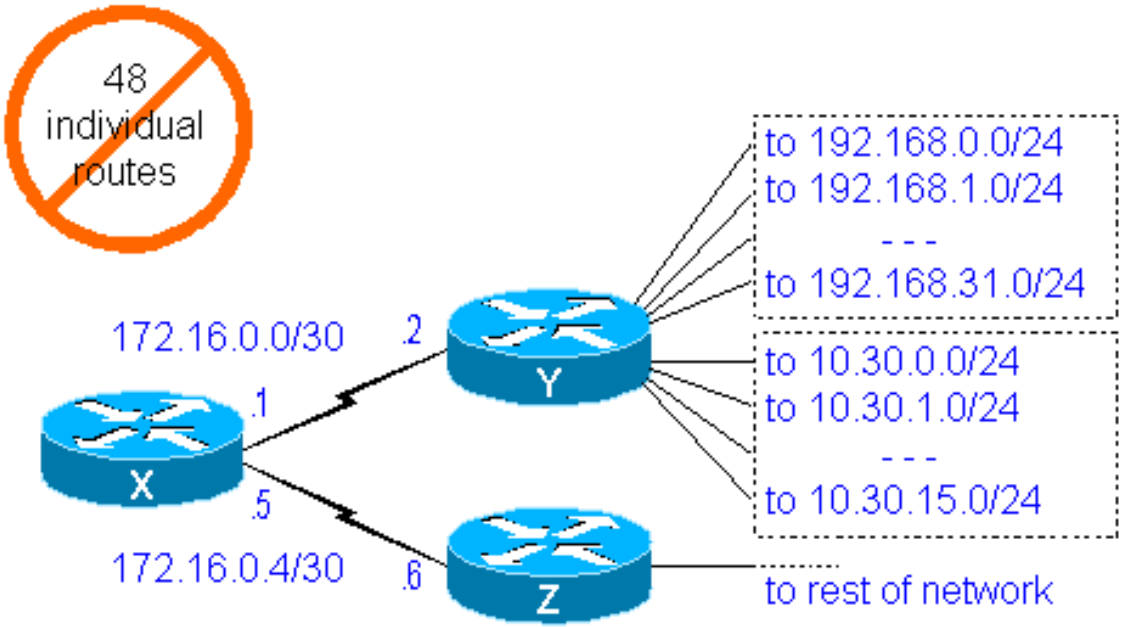

2 summarized routes

Route summarization analysis

- The supernet 192.168.0.0/19 (mask 255.255.224.0) covers addresses

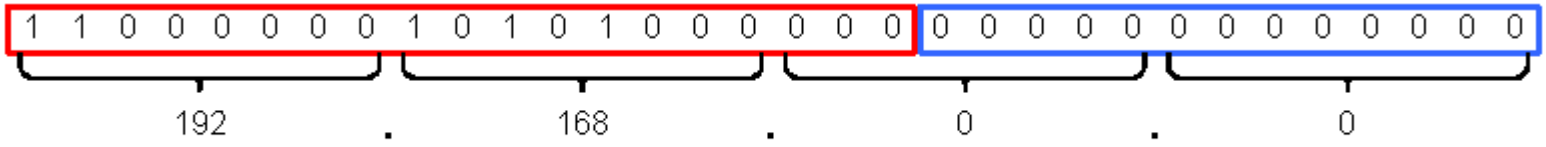

through

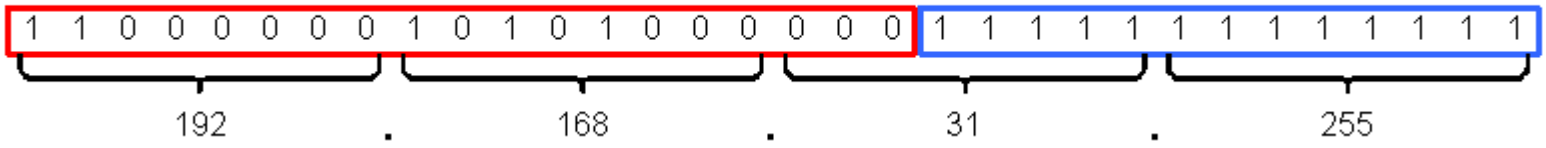

- While such a supernet would not be practical to service a single LAN segment, a summarized route to this supernet covers all possible addresses in the 32 individual class C networks 192.168.0.0/24 through 192.168.31.0/24.

AVAYA

Route summarization analysis continued

- The subnet 10.30.0.0/20 (mask 255.255.240.0) covers addresses

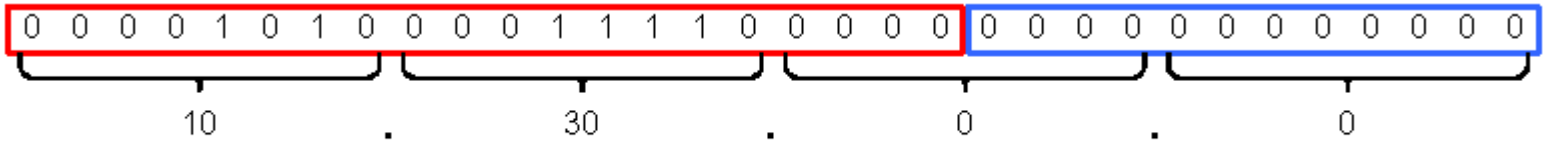

through

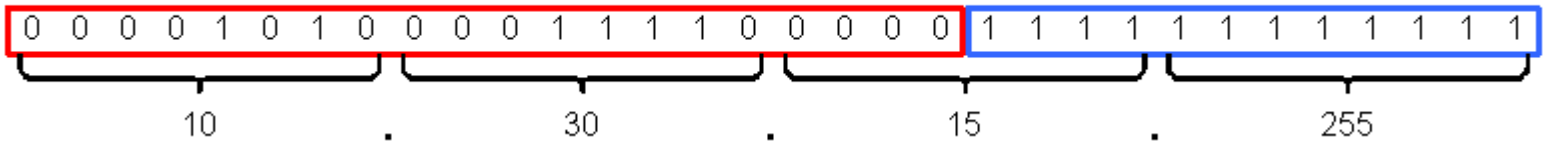

- While such a subnet would not be practical to service a single LAN segment, a summarized route to this subnet covers all possible addresses in the 16 individual subnets 10.30.0.0/24 through 10.30.15.0/24.

AVAYA

Routes w/ multiple IP gateways

- This is an example of a LAN segment that has two IP gateways to two different parts of the data network.

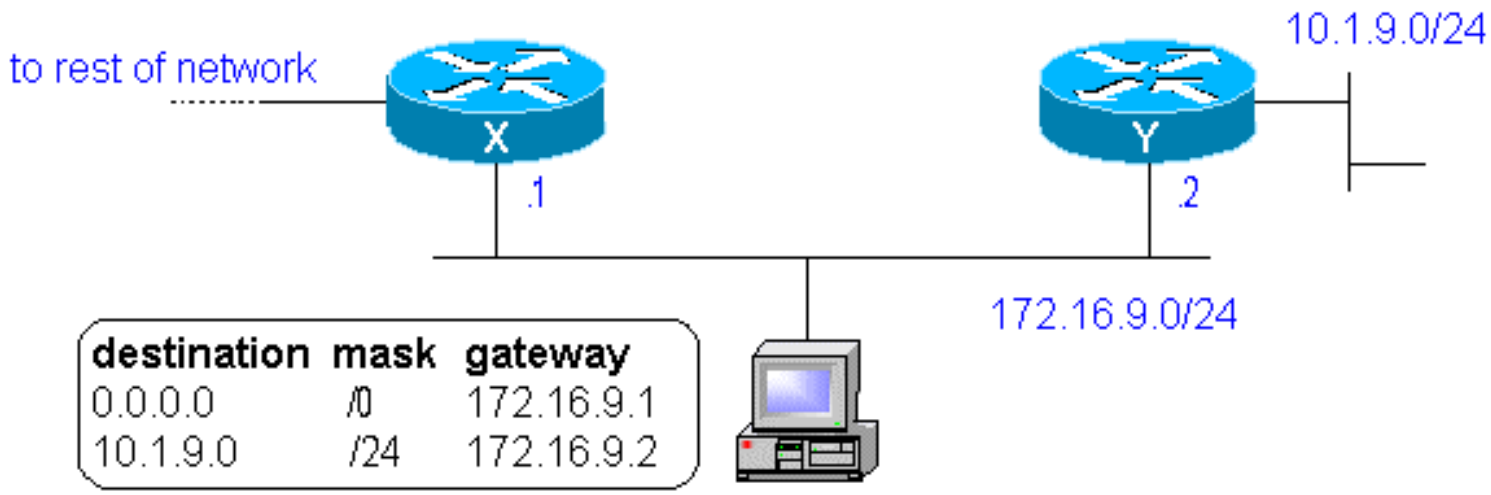

- In this configuration it is preferable that each host have two different routes - a default route to the default gateway, and a network route to the second gateway.

ICMP Redirects

- If the host does not have a network route to the second gateway, it can learn host routes from ICMP Redirect messages.

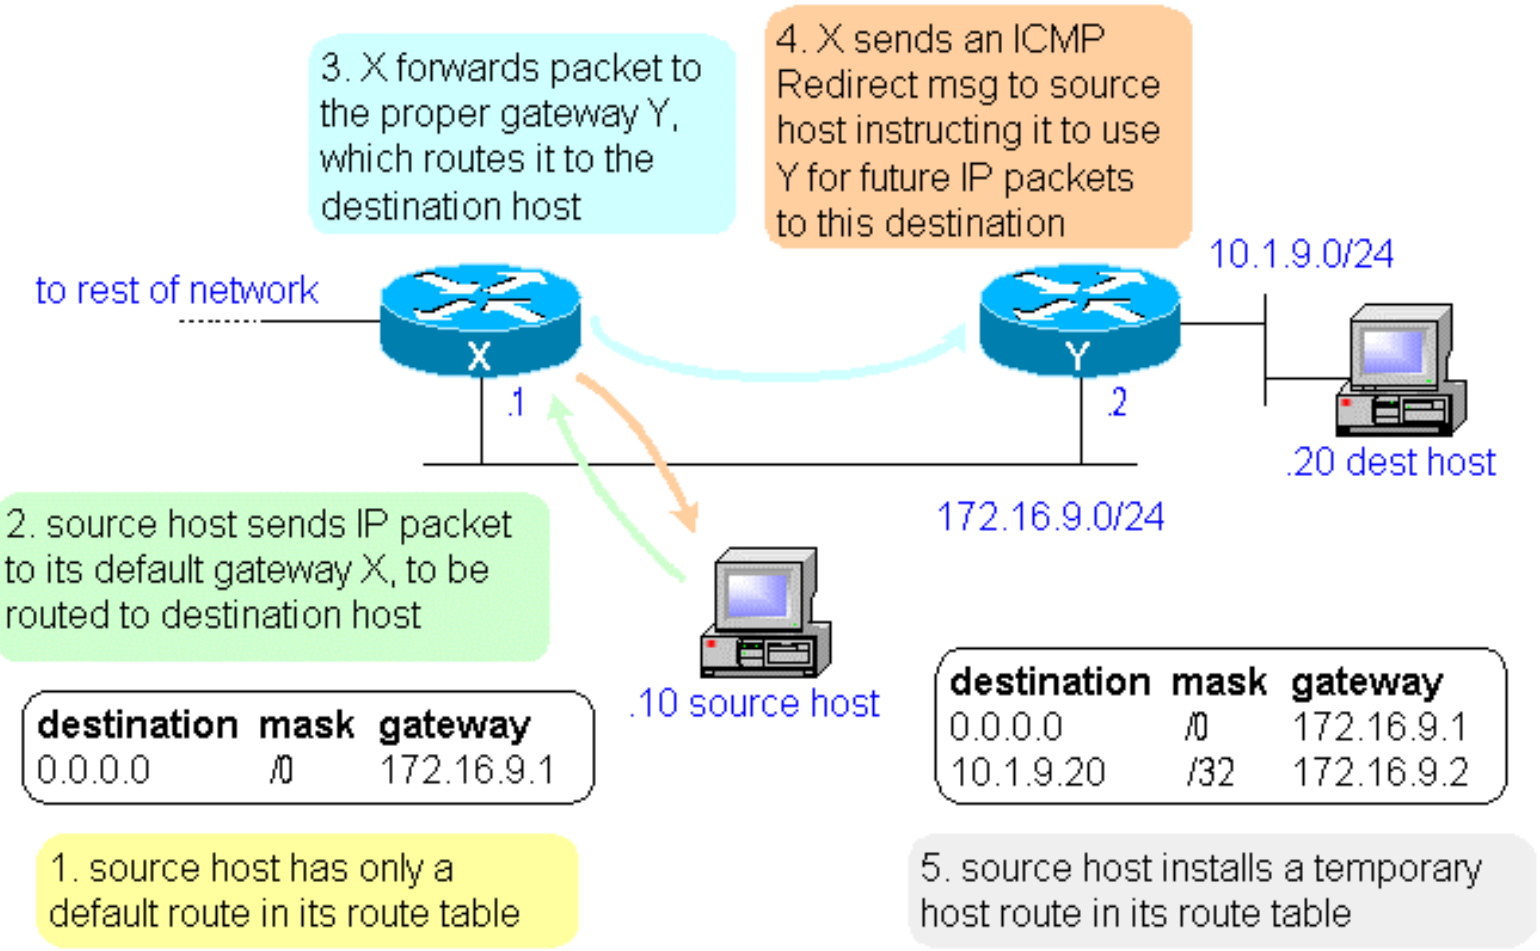

ICMP Redirect analysis

- The route learned via an ICMP Redirect message is a host route, which is a route to an individual host rather than to a network or subnet.
- A host route has a 32-bit mask (255.255.255.255), which means that all the bits in the IP address are fixed.
  - This makes sense because a typical network or subnet mask fixes the network bits and allows the host bits to vary.
  - With a 32-bit mask there are no bits left to vary, so the only valid address with this mask is the single address.
  - Thinking along these lines the 0.0.0.0 mask for the default route also makes sense, because this mask allows all the bits to vary, meaning that every address is included with this mask.
- In most implementations the host route learned via an ICMP Redirect message is a temporary route. It remains in the route table for a few minutes and then times out.

AVAYA

Recommended reading

- RFC 950 - Internet Standard Subnetting Procedure
- RFC 1519 - Classless Inter-Domain Routing (CIDR)
- RFC 1918 - Address Allocation for Private Internets

AVAYA