Network Implementation Tips

IP Address Facts

IP addresses allow hosts to participate on IP based networks. An IP address:

- Is a 32-bit binary number represented as four octets (four 8-bit numbers). Each octet is separated by a period.
- IP addresses can be represented in one of two ways:
    - Decimal (for example 131.107.2.200). In decimal notation, each octet must be between 0 and 255.
    - Binary (for example 10000011.01101011.00000010.11001000). In binary notation, each octet is an 8-character number.
- To convert from binary to decimal, memorize the decimal equivalent to the following binary numbers:

10000000	01000000	00100000	00010000	00001000	00000100	00000010	00000001
128	64	32	16	8	4	2	1

- For each bit position with a 1 value, add the decimal values for that bit together. For example, the decimal equivalent of 10010101 is:
  128 + 16 + 4 + 1 = 149
- The IP address includes both the network and the host address.
- The subnet mask is a 32-bit number that is associated with each IP address that identifies the network portion of the address. In binary form, the subnet mask is always a series of 1's followed by a series of 0's (1's and 0's are never mixed in sequence in the mask). A simple mask might be 255.255.255.0.
- IP addresses have a default *class*. The address class identifies the range of IP addresses and a default subnet mask used for the range. The following table shows the default address class for each IP address range.

Class	Address Range	First Octet Range	Default Subnet Mask
A	1.0.0.0 to 126.255.255.255	1-126 (00000001--01111110 binary)	255.0.0.0
B	128.0.0.0 to 191.255.255.255	128-191 (10000000--10111111 binary)	255.255.0.0
C	192.0.0.0 to 223.255.255.255	192-223 (11000000--11011111 binary)	255.255.255.0
D	224.0.0.0 to 239.255.255.255	224-239 (11100000--11101111 binary)	n/a
E	240.0.0.0 to 255.255.255.255	240-255 (11110000--11111111 binary)	n/a

- When using the default subnet mask for an IP address, you have the following number of subnet addresses and hosts per subnet:
    - There are only 126 Class A network IDs (most of these addresses are already assigned). Each class A address gives you 16,777,214 hosts per network.

  - There are 16,384 Class B network IDs. Each class B address gives you 65,534 hosts per network.
  - There are 2,097,152 Class C network IDs. Each class C address gives you 254 hosts per network.
  - Class D addresses are used for multicast groups rather than network and host IDs.
  - Class E addresses are reserved for experimental use.

As you are assigning IP addresses to hosts, be aware of the following special considerations:

Address	**Consideration**
Network	The first address in an address range is used to identify the network itself. For the network address, the host portion of the address contains all 0's. For example: • Class A network address: 115.0.0.0 • Class B network address: 154.90.0.0 • Class C network address: 221.65.244.0
Broadcast	The last address in the range is used as the broadcast address and is used to send messages to all hosts on the network. In binary form, the broadcast address has all 1's in the host portion of the address. For example, assuming the default subnet masks are used: • 115.255.255.255 is the broadcast address for network 115.0.0.0 • 154.90.255.255 is the broadcast address for network 154.90.0.0 • 221.65.244.255 is the broadcast address for network 221.65.244.0 **Note:** The broadcast address might also be designated by setting each of the network address bits to 0. For example, 0.0.255.255 is the broadcast address of a Class B address. This designation means "the broadcast address for this network."
Host Addresses	When you are assigning IP addresses to hosts, be aware of the following: • Each host must have a unique IP address. • Each host on the same network must have an IP address with a common network portion of the address. This means that you must use the same subnet mask when configuring addresses for hosts on the same network. The range of IP addresses available to be assigned to network hosts is identified by the subnet mask and/or the address class. When assigning IP addresses to hosts, be aware that you cannot use the first or last addresses in the range (these are reserved for the network and broadcast addresses respectively). For example: • For the class A network address 115.0.0.0, the host range is 115.0.0.1 to 115.255.255.254. • For the class B network address 154.90.0.0, the host range is 154.90.0.1 to 154.90.255.254. • For the class C network address 221.65.244.0, the host range is 221.65.244.1 to 221.65.244.254. **Note:** A special way to identify a host on a network is by setting the network portion of the address to all 0's. For example, the address 0.0.64.128 means "host 64.128 on this network."
Local Host	Addresses in the 127.0.0.0 range are reserved to refer to the local host (in other words "this"

	host or the host you're currently working at). The most commonly-used address is 127.0.0.1 which is the loopback address.

Because IP addresses assigned to hosts must be unique, the use of IP addresses on the Internet is controlled by organizations that ensure that no two organizations are given the same range of IP addresses to assign to hosts.

- The Internet Assigned Numbers Authority (IANA) manages the assignment of IP addresses on the Internet. IANA is operated by the Internet Corporation for Assigned Names and Numbers (ICANN).
- IANA allocates blocks of IP addresses to Regional Internet Registries (RIRs). An RIR has authority for IP addresses in a specific region of the world.
- An RIR assigns a block of addresses to Internet Service Providers (ISPs).
- An ISP assigns one or more IP addresses to individual computers or organizations connected to the Internet.

Subnetting Facts

Subnetting is the process of dividing a large network into smaller networks. When you subnet a network, each network segment (called a *subnet*) has a different network address (also called a *subnet address*). In practice, the terms *network* and *subnet* are used interchangeably to describe a physical network segment with a unique network address.

From a physical standpoint, subnetting is necessary because all network architectures have a limit on the number of hosts allowed on a single network segment. As your network grows, you will need to create subnets (physical networks) to:

- Increase the number of devices that can be added to the LAN (to overcome the architecture limits)
- Reduce the number of devices on a single subnet to reduce congestion and collisions
- Reduce the processing load placed on computers and routers
- Combine networks with different media types within the same internetwork (subnets can not be used to combine networks of different media type on to the same subnet)

Subnetting is also used to efficiently use the available IP addresses. For example, an organization with a class A network ID is allocated enough addresses for 16,777,214 hosts. If the organization actually uses only 10,000,000 host IDs, over 6 million IP addresses are not being used. Subnetting provides a way to break the single class A network ID into multiple network IDs.

- Subnetting uses *custom* rather than the default subnet masks. For example, instead of using 255.0.0.0 with a Class A address, you might use 255.255.0.0 instead.
- Using custom subnet masks is often called *classless* addressing because the subnet mask cannot be inferred simply from the class of a given IP address. The address class is ignored and the mask is always supplied to identify the network and host portions of the address.
- When you subnet a network by using a custom mask, you can divide the IP addresses between several subnets. However, you also reduce the number of hosts available on each network.

The following table shows how a Class B address can be subnetted to provide additional subnet addresses. Notice how by using a custom subnet mask the Class B address looks like a Class C address.

	Default Example	Custom Example
Network Address	188.50.0.0	188.50.0.0
Subnet Mask	255.255.0.0	255.255.255.0
# of Subnet Addresses	One	254
# of Hosts per Subnet	65,534	254 per subnet
Subnet Address(es)	188.50.0.0 (only one)	188.50.1.0 188.50.2.0 188.50.3.0 (and so on)
Host Address Range(s)	188.50.0.1 to 188.50.255.254	188.50.1.1 to 188.50.1.254 188.50.2.1 to 188.50.2.254 188.50.3.1 to 188.50.3.254 (and so on)

Note: It is possible to use subnet masks that do not use an entire octet. For example, the mask 255.255.252.0 uses six extra binary bits in the third octet. For the Network+ exam, you do not need to know how to work with such custom masks.

Be aware of the following additional facts about custom subnet masks:

- While subnetting divides a large address space into multiple subnets, *supernetting* combines multiple smaller network addresses into a single larger network. For example, this allows multiple Class C addresses to be combined into a single network.
- *Classful* addresses are IP addresses that use the default subnet mask. They are classful because the default subnet mask is used to identify the network and host portions of the address. *Classless* addresses are those that use a custom mask value to separate network and host portions of the IP address.
- Using classless addresses is made possible by a feature called Classless Inter-Domain Routing (CIDR). CIDR allows for non-default subnet masks (variable length subnet mask or VLSM). Routers use the following information to identify networks:
    - The beginning network address in the range
    - The number of bits used in the subnet mask

  For example, the subnet 199.70.0.0 with a mask of 255.255.0.0 is represented as 199.70.0.0/16 (with 16 being the number of 1 bits in the subnet mask).

Addressing Method Facts

The following table lists several options for assigning IP addresses.

Method	Uses
Dynamic Host Configuration Protocol (DHCP)	A DHCP server is a special server configured to pass out IP address and other IP configuration information to network clients. • When a client boots, it contacts the DHCP server for IP configuration information. • The DHCP server is configured with a range of IP addresses it can assign to

	hosts (Microsoft calls these ranges *scopes*). • The DHCP server can also be configured to pass out other IP configuration such as the default gateway and DNS server addresses. • The DHCP server ensures that each client has a unique IP address. • The DHCP server can be configured to not assign specific addresses in the range, or to assign a specific address to a specific host. • The DHCP server assigns the IP address and other information to the client. The assignment is called a *lease*, and includes a lease time that identifies how long the client can use the IP address. • Periodically and when the client reboots, it contacts the DHCP server to renew the lease on the IP address. • The DHCP lease process uses frame-level broadcasts. For this reason, DHCP requests typically do not pass through routers to other subnets. To enable DHCP across subnets: o Enable BootP (DHCP broadcast) requests through the router. o Configure a computer for BootP forwarding to request IP information on behalf of other clients. • You can configure a DHCP server to deliver the same address to a specific host each time it requests an address. Microsoft calls this configuration a *reservation*. • DHCP is a TCP/IP protocol. Any client configured to use DHCP can get an IP address from any server configured for DHCP, regardless of operating system. Use DHCP for small, medium, or large networks. DHCP requires a DHCP server and minimal configuration.
Automatic Private IP Addressing (APIPA)	APIPA is a Microsoft implementation of automatic IP address assignment without a DHCP server. Using APIPA, hosts assign themselves an IP address on the 169.254.0.0 network (mask of 255.255.0.0). With APIPA: • The host is configured to obtain IP information from a DHCP server (this is the default configuration). • If a DHCP server can't be contacted, the host uses APIPA to assign itself an IP address. • The host only configures the IP address and mask. It does not assign itself the default gateway and DNS server addresses. For this reason, APIPA can only be used on a single subnet. Use APIPA: • On small, single-subnet networks where you do not need to customize the IP address range. • As a fail safe for when a DHCP server is unavailable to provide limited communication capabilities.
Alternate IP configuration	With an alternate IP configuration, the system attempts to use DHCP for TCP/IP configuration information. If a DHCP server cannot be contacted, the static configuration values are used. When you configure an alternate IP address, APIPA is no longer used. Use an alternate configuration: • If you have a computer (such as a laptop) that connects to two networks: one with a DHCP server and another without a DHCP server.

	• If you want to provide values to properly configure the computer in case the DHCP server is unavailable.
Static (manual) assignment	Using static addressing, IP configuration information must be manually configured on each host. Use static addressing: • On networks with a very small number of hosts. • On networks that do not change often or that will not grow. • To permanently assign IP addresses to hosts that must have always have the same address (such as printers, servers, or routers). • For hosts that cannot accept an IP address from DHCP. • To reduce DHCP-related traffic. **Note:** Static addressing is very susceptible to configuration errors and duplicate IP address configuration errors (two hosts that have been assigned the same IP address). Static addressing also disables both APIPA and DHCP capabilities on the host.

DNS Facts

The Domain Name System (DNS) is a hierarchical, distributed database that maps logical host names to IP addresses. The DNS hierarchy is made up of the following components:

- . (dot) domain (also called the *root* domain)
- Top Level Domains (TLDs) such as .com, .edu, .gov
- Additional domains such as yahoo.com, microsoft.com, etc.
- Hosts

The fully-qualified domain name (FQDN) includes the host name and all domain names, separated by periods. The final period (for the root domain) is often omitted and implied.

DNS is a distributed database because no one server holds all of the DNS information. Instead, multiple servers hold portions of the data.

- Each division of the database is held in a *zone* database file.
- Zones typically contain one or more domains, although additional servers might hold information for child domains.
- DNS servers hold zone files and process name resolution requests from client systems.

When you use the host name of a computer (for example if you type a URL such as www.mydomain.com), your computer uses the following process to find the IP address.

1. The host looks in its local cache to see if it has recently resolved the host name.
2. If the information is not in the cache, it checks the Hosts file. The Hosts file is a static text file that contains hostname-to-IP address mappings.
3. If the IP address is not found, the host contacts its preferred DNS server. If the preferred DNS server can't be contacted, it continues contacting additional DNS servers until one responds.
4. The host sends the name information to the DNS server. The DNS server then checks its cache and Hosts file. If the information is not found, the DNS server checks any zone files that it holds for the requested name.

5. If the DNS server can't find the name in its zones, it forwards the request to a root zone name server. This server returns the IP address of a DNS server that has information for the corresponding top-level domain (such as .com).
6. The first DNS server then requests the information from the top-level domain server. This server returns the address of a DNS server with the information for the next highest domain. This process continues until a DNS server is contacted that holds the necessary information.
7. The DNS server places the information in its cache and returns the IP address to the client host. The client host also places the information in its cache and uses the IP address to contact the desired destination device.

You should know the following facts about DNS:

- A *forward* lookup finds the IP address for a given host name. A *reverse* lookup finds the host name from a given IP address.
- An *authoritative* server is a DNS server that has a full, complete copy of all the records for a particular domain.
- Zone files hold records that identify hosts.
  - A records map host names to IP addresses.
  - PTR (pointer) records map IP addresses to host names.
- *Recursion* is the process by which a DNS server or host uses root name servers and subsequent servers to perform name resolution. Most client computers do not perform recursion, rather they submit a DNS request to the DNS server and wait for a complete response. Many DNS servers will perform recursion.
- Some DNS servers might forward the name resolution request to another DNS server and wait for the final response rather than performing recursion.
- Root DNS servers hold information for the root zone (.). Root servers answer name resolution requests by supplying the address of the corresponding to top-level DNS server (servers authoritative for .com, .edu, and such domains).
- On very small networks, you could configure a HOSTS file with several entries to provide limited name resolution services. However, you would have to copy the HOSTS file to each client. The work involved in this solution is only suitable for temporary testing purposes or to override information that might be received from a DNS server.

Routing Facts

A *router* is a device that sends packets from one network to another network. Routers receive packets, read their headers to find addressing information, and send them on to their correct destination on the network or Internet. Routers can forward packets through an internetwork by maintaining routing information in a database called a *routing table*. The routing table typically contains the following information:

- The address of a known network
- The interface or next hop router used to reach the destination network
- A cost value (also called a *metric*) that identifies the desirability of the route to the destination network (using distance, delay, or cost)
- A timeout value that identifies when the route expires

Routers automatically have an entry in their routing tables for all directly-connected networks. Information about other networks can be placed in the routing table using one of two methods:

Method	Description
Static	Static routing requires that entries in the routing table are configured manually. • Network entries remain in the routing table until manually removed. • When changes to the network occur, static entries must be modified, added, or removed.
Dynamic	Routers can dynamically learn about networks by sharing routing information with other routers. The routing protocol defines how routers communicate with each other to share and learn about other networks. The routing protocol determines: • The information contained in the routing table • How messages are routed from one network to another • How topology changes (i.e. updates to the routing table) are communicated between routers Use a routing protocol to allow a router to automatically learn about other networks. The routing protocol generates some network traffic for the process of sharing routes, but has the advantage of being dynamic and automatic (i.e. changes in the network are propagated automatically to other routers).

Be aware of the following when managing routing tables:

- By default, routers know about directly-connected networks; you do not need to create static entries for directly-connected networks.
- You can use both dynamic and static routing together. You can add static routes to identify networks that are not learned about through the routing protocol.
- The most common reason to create a static routing table entry is to define a default route. The default route is similar to a default gateway setting on a workstation, it identifies a router that is used to forward packets to networks that do not appear in the routing table.
- When you configure a router for dynamic routing, you enable a routing protocol, then identify the interfaces that will participate in the exchange of routing information. Enabling a routing protocol on an interface does the following:
  - Configures the router to share information in its routing table with other routers accessible on that interface.
  - Configures the router to share information about that network with other routers.
- When using a routing protocol, changes in routing information take some time to be propagated to all routers on the network. The term *convergence* is used to describe the condition when all routers have the same (or correct) routing information.

Routing Protocol Characteristics Facts

Routers use a routing protocol to exchange information about known routes with other routers. The following table describes general characteristics of a routing protocol.

Characteristic	Description
Scope	Each organization that has been assigned a network address from an ISP is considered an Autonomous System (AS). That organization is free to create one large network, or divide the network into subnets. Each autonomous system is identified by an AS number. This number can be locally administered, or registered if the AS is connected to the Internet.

	Routing protocols can be classified based on whether they are routing traffic within or between autonomous systems. • An Interior Gateway Protocol (IGP) routes traffic *within* an AS. • An Exterior Gateway Protocol (EGP) routes traffic *between* ASs.
Metric	The *metric* is a value assigned to each route that identifies the distance or cost to the destination network. The metric is used by the routing protocol to identify and select the best route to the destination when multiple routes exist. A *lower* metric identifies a more preferred route. Common metrics include: • The *hop count* is a count of the number of routers between the current router and the destination network. • The metric might be based on an actual measure of the *bandwidth* or time it takes to reach the destination network (delay). For example, high speed links might be associated with a lower metric cost. • A *link cost* is a relative number that represents the cost for using the route. For example, the link cost could relate to the actual cost of using a link, such as an expensive WAN link, or it might identify the desirability of using a specific link. Be aware that the metric is used by the routing protocol to select the best or possibly alternate routes to the destination. Comparing route metrics used by different routing protocols is not useful. For example, a metric of 10 for a routing protocol that uses the bandwidth as the metric might indicate a better route than a metric of 4 for a protocol that uses hop count for the metric.
Routing update method	Routing protocols use different methods for sharing routing information and discovering networks. Three common methods are: • Using the *distance vector* method, routers share their entire routing table with their immediate neighbors. Routes learned from neighboring routers are added to the routing table, then shared with that router's neighbors. • Using the *link state* method, routers share only their directly-connected routes using special packets called link-state advertisements (LSAs) and link-state packets (LSPs). These route advertisements are *flooded* (forwarded) throughout the network. Routers use this information to build a topology database of the network. • A *hybrid* method combines characteristics of both the distance vector and link state methods. It shares its full routing table at startup, followed by partial updates when changes occur. In general, the different routing protocol methods have the following characteristics: • The distance vector method is simpler and requires less processing power for routers. Distance vector methods are best suited for small networks. • The link state method uses less network traffic for sending routing information, converges faster, and is less prone to errors. Link state methods are the best choice for large networks, or for sharing routes over WAN links. • Hybrid methods reduce the negative effects of a distance vector method while gaining many of the benefits of a link state method.

Classful or classless	Early routing protocols were not capable of variable length subnet masks (VLSM), and used only the default subnet masks to identify destination networks. Routing protocols can be classified based on their support for Classless Inter-Domain Routing (CIDR) features as follows: • A *classful* protocol uses the IP address class and the default subnet mask to identify network addresses. Classful protocols do not support CIDR or VLSM. • A *classless* protocol ignores the IP address class and requires that a subnet mask value be included in all route advertisements. Classless protocols support CIDR and VLSM.

Routing Protocol Facts

The following table lists the characteristics of specific routing protocols.

Protocol	**Description**
Routing Information Protocol (RIP)	RIP is a distance vector routing protocol used for routing within an autonomous system (i.e.an IGP). • RIP uses the hop count as the metric. • RIP networks are limited in size to a maximum of 15 hops between any two networks. A network with a hop count of 16 indicates an unreachable network. • RIP v1 is a classful protocol; RIP v2 is a classless protocol. RIP is best suited for small private networks.
Enhanced Interior Gateway Routing Protocol (EIGRP)	EIGRP is a hybrid routing protocol developed by Cisco for routing within an AS. • EIGRP uses a composite number for the metric that indicates bandwidth and delay for a link. The higher the bandwidth, the lower the metric. • EIGRP is a classless protocol. EIGRP is best suited for medium to large private networks.
Open Shortest Path First (OSPF)	OSPF is a link state routing protocol used for routing within an AS. • OSPF uses a relative link cost for the metric. • OSPF is a classless protocol. • OSPF divides a large network into areas. o Each autonomous system requires an area 0 that identifies the network backbone. o All areas are connected to area 0, either directly or indirectly through another area. o Routes between areas must pass through area 0. • Internal routers share routes *within* an area; area border routers share routes *between* areas; autonomous system boundary routers share routes *outside* of the AS. • A router is the boundary between one area and another area. OSPF is best suited for large private networks.

Intermediate System to Intermediate System (IS-IS)	IS-IS is a link state routing protocol used for routing within an AS. • IS-IS uses a relative link cost for the metric. • IS-IS is a classless protocol. • The original IS-IS protocol was not used for routing IP packets; use Integrated IS-IS to include IP routing support. • IS-IS divides a large network into areas. There is no area 0 requirement, and IS-IS provides greater flexibility than OSPF for creating and connecting areas. • L1 routers share routes *within* an area; L2 routers share routes *between* areas; an L1/L2 router can share routes with both L1 and L2 routers. • A network link is the boundary between one area and another area. IS-IS is best suited for large private networks, supporting larger networks than OSPF. IS-IS is typically used within an ISP, and easily supports IPv6 routing.
Border Gateway Protocol (BGP)	BGP is an advanced distance vector protocol (also called a *path vector* protocol). BGP is an exterior gateway protocol (EGP) used for routing between autonomous systems. • BGP uses paths, rules, and policies instead of a metric for making routing decisions. • BGP is a classless protocol. • Internal BGP (iBGP) is used *within* an autonomous system; External BGP (eBGP) is used *between* ASs. BGP is the protocol used on the Internet: ISPs use BGP to identify routes between ASs. Very large networks can use BGP internally, but typically only share routes on the Internet if the AS has two (or more) connections to the Internet through different ISPs.

NAT Facts

Network Address Translation (NAT) allows you to connect a private network to the Internet without obtaining registered addresses for every host. Private addresses are translated to the public address of the NAT router.

- Hosts on the private network share the IP address of the NAT router, or a pool of addresses assigned for the network.
- The NAT router maps port numbers to private IP addresses. Responses to Internet requests include the port number appended by the NAT router. This allows the NAT router to forward responses back to the correct private host.
- Technically speaking, NAT translates one address to another. Port address translation (PAT) associates a port number with the translated address.
    - With only NAT, you would have to have a public address for each private host. NAT would associate a single public address with a single private address.
    - Use PAT to allow multiple private hosts to share a single public address. Each private host is associated with a unique port number.

  Because virtually all NAT routers perform port address translation, most routers that are configured with NAT are really performing PAT. When you use a NAT router, you are normally using PAT and not just NAT. (NAT is typically used synonymously with PAT.)

- NAT supports a limit of 5,000 concurrent connections.
- NAT provides some security for the private network because it translates or hides the private addresses.
- A NAT router can act as a limited-function DHCP server, assigning addresses to private hosts.
- A NAT router can forward DNS requests to the Internet.
- There are three types of NAT implementation:

Type	Description
Dynamic NAT	Dynamic NAT automatically maps internal IP addresses with a dynamic port assignment. On the NAT device, the internal device is identified by the public IP address and the dynamic port number. Dynamic NAT allows internal (private) hosts to contact external (public) hosts but not vice versa. External hosts cannot initiate communications with internal hosts.
Static NAT (SNAT)	Static NAT maps an internal IP address to a static port assignment. Static NAT is typically used to take a server on the private network (such as a Web server) and make it available on the Internet. External hosts contact the internal server using the public IP address and the static port. Using a static mapping allows external hosts to contact internal hosts.
Dynamic and Static NAT	Dynamic and Static NAT, in which two IP addresses are given to the public NAT interface (one for dynamic NAT and one for static NAT), allows traffic to flow in both directions.

When connecting a private network to the Internet through NAT, assign IP addresses on the private network in several predefined private address ranges. These address ranges are guaranteed to not be in use on the Internet and do not need to be registered. The private IPv4 address ranges are:

- 10.0.0.1 to 10.255.255.254
- 172.16.0.1 to 172.31.255.254
- 192.168.0.1 to 192.168.255.254

The Internet Assigned Numbers Authority (IANA) is responsible for allocating IP addresses used on the Internet. When you want to obtain a public IP address, you would typically get the address from your ISP, which has received it from a Regional Internet Registry (RIR), which has been assigned a block of addresses from IANA. IANA is operated by the Internet Corporation for Assigned Names and Numbers (ICANN), so you might also see that ICANN is responsible for assigning public IP addresses.

ICS Facts

Internet Connection Sharing (ICS) is a service available on Windows systems that enables multiple computers on a single small network to access the Internet by sharing one computer's connection. With ICS, most configuration tasks are completed automatically. When using ICS:

- The ICS system is configured as a NAT router, a limited DHCP server, and a DNS proxy (name resolution requests from the private network are forwarded to DNS servers on the Internet).
- The IP address for the private interface is automatically changed to 192.168.0.1 with a mask of 255.255.255.0.
- The default gateway of the ICS system is set to point to the Internet connection.
- Hosts on the private network should use DHCP for address and DNS server information.
- The ICS system uses DHCP to deliver the following information to hosts on the private network:
    - IP address on the 192.168.0.0 subnet with a mask of 255.255.255.0.
    - DNS server address of 192.168.0.1 (the private interface of the ICS system).
    - Default gateway address of 192.168.0.1.

- When using ICS, do not use DHCP servers, DNS servers, or Active Directory on your private network.

IPv6 Facts

The current IP addressing standard, version 4, will eventually run out of unique addresses, so a new system is being developed. It is named IP version 6 or IPv6. The IPv6 address is a 128-bit binary number. A sample IPv6 IP address looks like: 35BC:FA77:4898:DAFC:200C:FBBC:A007:8973. The following list describes the features of an IPv6 address:

- The address is made up of 32 hexadecimal numbers, organized into 8 quartets.
- The quartets are separated by colons.
- Each quartet is represented as a hexadecimal number between 0 and FFFF. Each quartet represents 16-bits of data (FFFF = 1111 1111 1111 1111).
- Leading zeros can be omitted in each section. For example, the quartet 0284 could also be represented by 284.
- Addresses with consecutive zeros can be expressed more concisely by substituting a double-colon for the group of zeros. For example:
  - FEC0:0:0:0:78CD:1283:F398:23AB
  - FEC0::78CD:1283:F398:23AB (concise form)
- If an address has more than one consecutive location where one or more quartets are all zeros, only one location can be abbreviated. For example, FEC2:0:0:0:78CA:0:0:23AB could be abbreviated as:
  - FEC2::78CA:0:0:23AB or
  - FEC2:0:0:0:78CA::23AB

  But *not* FEC2::78CA::23AB

- The 128-bit address contains two parts:

Component	Description
Prefix	The first 64-bits is known as the *prefix*. ○ The 64-bit prefix can be divided into various parts, with each part having a specific meaning. Parts in the prefix can identify the geographic region, the ISP, the network, and the subnet. ○ The *prefix length* identifies the number of bits in the relevant portion of the prefix. To indicate the prefix length, add a slash (/) followed by the prefix length number. Full quartets with trailing 0's in the prefix address can be omitted (for example 2001:0DB8:4898:DAFC::/64). ○ Because addresses are allocated based on physical location, the prefix generally identifies the location of the host. The 64-bit prefix is often referred to as the *global routing* prefix.
Interface ID	The last 64-bits is the *interface ID*. This is the unique address assigned to an interface. ○ Addresses are assigned to interfaces (network connections), not to the host. Technically, the interface ID is *not* a host address. ○ In most cases, individual interface IDs are not assigned by ISPs, but are rather generated automatically or managed by site administrators. ○ Interface IDs must be unique within a subnet, but can be the same if the

	interface is on different subnets. ○ On Ethernet networks, the interface ID can be automatically derived from the MAC address. Using the automatic host ID simplifies administration.

IPv6 adds the following features which are not included in IPv4:

Feature	Description
Auto-configuration	Because hardware IDs are used for node IDs, IPv6 nodes simply need to discover their network ID. This can be done by communicating with a router.
Built-in Quality of Service	Built-in support for bandwidth reservations which make guaranteed data transfer rates possible. (Quality of service features are available as add-ons within an IPv4 environment, but are not part of the native protocol.)
Built-in Security Features	IPv6 has built-in support for security protocols such as IPSec. (IPSec security features are available as add-ons within an IPv4 environment.)
Source Intelligent Routing	IPv6 nodes have the option to include addresses that determine part or all of the route a packet will take through the network.

Although not yet widely adopted, you can implement IPv6 if your systems support it. As implementation of IPv6 proceeds, there will be cases when compatibility with IPv4 is required. Three strategies are recommended by IETF for IPv6 to IPv4 compatibility configuration:

Strategy	Description
Dual Stack	With a *dual stack* configuration, both the IPv4 and IPv6 protocol stacks run concurrently on a host. IPv4 is used to communicate with IPv4 hosts, and IPv6 is used to communicate with IPv6 hosts. When implemented on hosts, intermediate routers and switches must also run both protocol stacks. Use a dual stack configuration to enable a host to communicate with both IPv4 and IPv6 hosts.
Tunneling	*Tunneling* wraps an IPv6 packet within an IPv4 packet, allowing IPv6 hosts or sites to communicate over the existing IPv4 infrastructure. With tunneling, a device encapsulates IPv6 packets in IPv4 packets for transmission across an IPv4 network, and then the packets are de-encapsulated to their original IPv6 packets by another device at the other end. Tunneling solutions include: • Intra-site Automatic Tunnel Addressing Protocol (ISATAP) for implementations within a site • 6-to-4 tunneling for implementations across sites • Teredo for tunneling between two hosts Use tunneling to allow an IPv6 host to communicate with another IPv6 host through an IPv4 network.
Network Address Translation-Protocol Translation (NAT-PT)	NAT-PT is a protocol that converts the IPv6 packet header into an IPv4 packet header, and vice versa. This method is different than tunneling because the packet headers are converted between the IPv4 and IPv6, whereas tunneling wraps the IPv6 packet into an IPv4 packet. Use NAT-PT to allow IPv4 hosts to communicate with IPv6 hosts.

Multicast Facts

Multicasting creates logical groups of hosts--messages sent to the group are received by all group members. Multicasting is typically used for streaming video and audio such as Microsoft's NetMeeting or Windows Media viewer.

Without the ability to define multicast groups, messages that must be sent to a specific group could only use the following:

- With *unicasting*, messages are sent to a specific host address. Using unicasting to send messages to a group of computers would mean that the sending device must know the IP address of all recipients, and must create a separate packet for each destination device.
- With *broadcasting*, a single packet is sent to the broadcast address and is processed by all hosts. However, using broadcasting for sending data to a group would mean that all hosts, and not just group members, would receive the packet. In addition, broadcast packets are typically not forwarded by routers, so broadcast traffic is limited to within a single subnet.

The Internet Group Management Protocol (IGMP) is used to identify group members and to forward multicast packets onto the segments where group members reside. IGMP routers keep track of the attached subnets that have group members as follows:

- A router sends out a host membership query. This query is addressed to the IP address of 224.0.0.1.
- Hosts that are members of any group respond with a list of the groups to which the host belongs. Each group is identified with a multicast IP address in the range of 224.0.0.0 to 239.255.255.255.
- The router uses these responses to compile a list of the groups on that subnet that have group members. Routers do *not* keep track of individual hosts that are members of a group, rather they simply compile a list of groups on the subnet that have at least one member.
- When a host joins a new group, it automatically sends a join group message to the router. When the last host in a group leaves the group, it sends a leave group message to the router.
- The IGMP router reports to upstream routers that they have members of a specific group.
  - Upstream routers are the routers that exist between the router and the server that sends out the multicast data stream.
  - Upstream routers keep track of downstream routers that have group members.

The following process is used when sending a multicast stream:

1. The sending server sends packets addressed to the multicast group.
2. Routers receive the multicast packets and check their lists of group members.
   - If the router is connected to a subnet that has group members, or if the subnet includes a downstream router with group members, the multicast packet is sent on that subnet.
   - If a subnet does not have any group members, the packet is not forwarded on that subnet.
   - If a router does not have any subnets with group members, the packet is dropped and not forwarded.
3. Each intermediary router performs the same tasks until the data stream eventually reaches the multicast client.

Be aware of the following additional facts about multicasting:

- Hosts can be members of one or more groups.
- Hosts can join or leave groups at any time.
- A multicast group is identified by a multicast IP address.

    - Each group is identified by a different address.
    - Multicast addresses are in the range 224.0.0.0 to 239.255.255.255.
    - The address 224.0.0.1 is never assigned to a group because it is used for the query messages sent by routers.
- Frames that contain multicast traffic are sent to a special MAC address. The MAC address begins with 01-00-5E, with the last portion being a form of the IP multicast group address. Be aware that a single multicast MAC address could be shared by up to 5 other IP multicast addresses.
- Routers do not keep track of which hosts are part of the group, only that there is at least one group member on the subnet (or that the subnet is connected to a downstream router that has group members). A router sends multicast traffic it receives on all subnets with group members.
- A regular switch that receives multicast traffic sends the traffic out all ports, because the destination MAC address will be an unknown address. This means that a host might see multicast traffic on its segment, even if it isn't a member of the group. However, hosts that are not members of the group will not process the frame because they will not associate the multicast MAC address with their own address.
- IGMP snooping on a switch allows the switch to control which ports get IGMP traffic for a specific group. With IGMP snooping, the switch identifies which ports include members of a specific multicast group. When a message is received for a group, the message is sent only to the ports that have a group member connected.

VoIP Facts

Voice over IP (VoIP) is a protocol optimized for the transmission of voice (telephone calls) through a packet switched network. VoIP routes phone calls through an IP network, such as the Internet, instead of through the public telephone system (PSTN). However, VoIP solutions are typically integrated with the PSTN to allow VoIP customers to call any phone on the PSTN, and to allow phones on the PSTN to call phones connected to the VoIP network.

Obtaining VoIP service can be done in the following ways:

- Using an analog telephone adapter, the existing analog phone system is connected to a VoIP network. Analog signals are converted to digital signals, then encapsulated into IP packets for transmission on the VoIP network.
- A VoIP phone is a special phone capable of sending and receiving digital voice signals, already formatted for the VoIP network.
- When using VoIP phones, connect the phones to special switches with Power over Ethernet (PoE) capabilities. PoE supplies power to the VoIP phone through an Ethernet cable, the same cable that is used for transmitting data signals.
- With special software, you can use a computer with an existing broadband connection to send and receive VoIP calls. Software running on the computer converts the input from a microphone into a VoIP call.

VoIP uses regular IP datagrams for sending voice data over a network.

1. If using a regular phone, analog signals are converted to digital data.
2. Digital data is segmented and placed into IP packets.
3. Packets are sent through an IP network. A VoIP call consists of two data flows:
    - The voice carrier stream, consisting of Real-Time Transport Protocol (RTP) packets containing the actual voice samples.
    - The call control signaling, consisting of one of several protocols which set up, maintain, teardown, and redirect the call. Protocols used in call control include the following:
        - H.323
        - Session Initiation Protocol (SIP)
        - Media Gateway Control Protocol (MGCP)

4. At the receiving end, packets become segments which are re-assembled into the voice data stream. If necessary, digital data is converted back to analog for use on an analog phone or for final transmission onto the PSTN.

Using an IP network for voice has the following advantages:

- Administration is simplified because you maintain a single network for both data and voice instead of using a separate infrastructure for voice-only traffic.
- Costs for sending voice over an IP infrastructure are typically lower than long-distance costs over the PSTN.
- Adding additional phone lines is easier and cheaper than adding lines through the PSTN.
- Because packets are regular IP packets, encryption can be easily added to VoIP data--something that is difficult to accomplish for calls on the PSTN.

When you use the PSTN for a regular phone call, a dedicated circuit is established between the calling and the called phones. In addition, data is sent sequentially as it occurs, in a steady stream. Because VoIP uses data packets over a packet switched network, VoIP is susceptible to the following problems:

- *Delay* (or *latency*) occurs when data takes a long time to arrive at the receiving device. Delays cause long pauses between speaking and receiving, and might result in callers continually interrupting each other. International standards call for a delay of 150 ms or less.
- *Jitter* is a variation in the delay of individual packets. Jitter causes strange sound effects as the delay of packets fluctuates.
- Packet *loss* occurs when packets do not arrive. Packet loss causes drop-outs in the conversation.
    - Because voice traffic is time sensitive, lost packets do not need to be retransmitted.
    - Voice traffic is very sensitive to packet loss. Even a 1% loss of packets can be detected.
- *Echo* is hearing your own voice in the telephone receiver while you are talking. Excessive delay can cause unacceptable echo. When timed properly, echo is reassuring to the speaker.
- Power loss at your local facility or at any point in the IP network can disrupt phone service. With regular phone service, power to the phone line is supplied separately from the electrical power to the building. With regular phone service, you might still be able to make calls when the power is out. With VoIP, a normal power disruption typically affects VoIP calls as well.

With VoIP, specific measures have been implemented to reduce the negative impacts of using the IP network. Collectively, these measures are often referred to as Quality of Service (QoS). QoS measures ensure that voice data is given higher priority on the network to decrease the effects of delay, jitter, and packet loss.