



How Traffic Gets from Here to There At-A-Glance

Cisco Press

At-A-Glance—Routing

Why Should I Care About Routing?

Routing is one of the fundamental aspects of networking. The ability of routers to learn possible routes (rather than make you manually configure and constantly update the routes) is one of the primary reasons that ARPANET, which originally connected seven sites, scaled into the modern Internet in only a few short years.

What Are the Problems to Solve?

Routed networks are often large and complex, and it would be prohibitively difficult to manage and update network information on all routers all the time. Several algorithms have been developed to help address these difficulties. These algorithms allow the routers to learn about the network and then make decisions based on that information.

To learn paths (or routes) through a network, and make decisions on where to send packets, a router must know the following information:

- **Destination address**—Typically the Internet Protocol (IP) address of the data's (packet) destination.
- **Source address**—Where the information came from (typically an IP address).
- **Possible routes**—Routes that can get information from its present location or source to some other location (the destination or closest known point).
- **Best route**—The best path to the intended destination. ("Best" can mean many things—see below.)
- **Status of routes**—The current state of routes, which routers track to ensure timely delivery of information.

What Exactly Does "Best" Mean?

Routers often make decisions about the best possible path to get information from a source to a destination. "Best," however, is loosely defined, and it depends on what is valued by the network. These measurements of value are referred to as *metrics*. Which metrics are valued by the network is determined by the network administrator. Several metrics are listed here:

- **Hop count**—Number of times a packet goes through a router.
- **Delay time**—Time required to reach the destination.
- **Reliability**—Bit-error rate of each network link.
- **Maximum transmission unit (MTU)**—Maximum message length (or packet size) allowed on the path.
- **Cost**—Arbitrary value based on a network-administrator-determined value. Usually some combination of other metrics.

Static Versus Dynamic

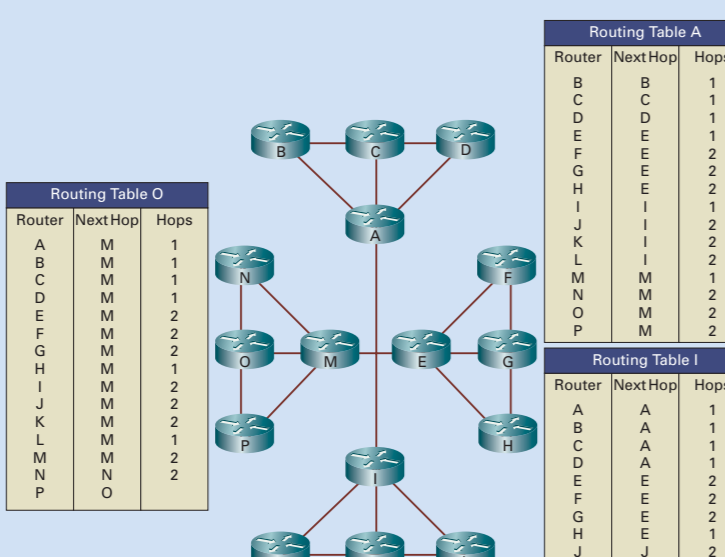
Routers must learn about the network around them to make determinations on where to send packets. This information can either be manually entered (static routes) or learned from other routers in the network (dynamic routes):

- **Static routes**—When a network administrator manually enters information about a route, it is considered a *static route*. Only a network administrator can change this information. (That is, the router does not learn from, or update, its routing tables based on network events.) Static routes allow for tight control of packets but are difficult to maintain and prone to human error.
- **Dynamic routes**—Routers on a network can learn about possible routes and current route status from other routers in the network. Routes learned in this way are called *dynamic routes*. Routers in dynamic routes learn about changes in the network without administrative intervention and automatically propagate them throughout the network.

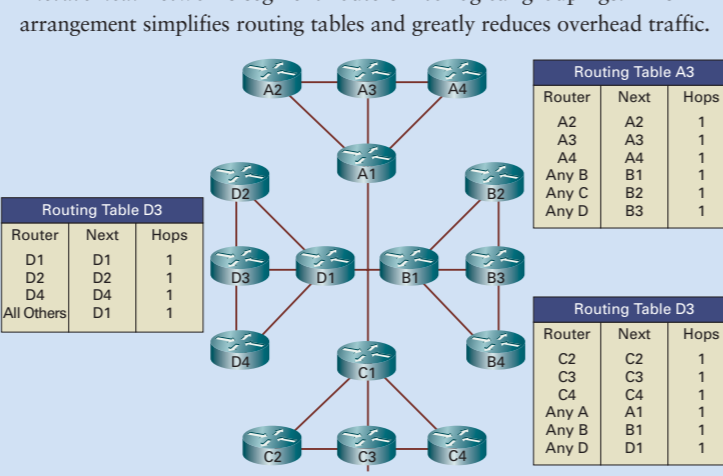
Flat Versus Hierarchical

With *flat* networks, all routers must keep track of all other routers on the network. As networks grow, the amount of information contained in the routing tables increases.

Although this method is simple, it can result in poor network performance because the number of routing updates traffic grows with each new router.




Hierarchical networks segment routers into logical groupings. This arrangement simplifies routing tables and greatly reduces overhead traffic.



Distance-Vector Versus Link-State Routing

The two main classes of routing are distant vector routing and link-state routing.

With *distance-vector routing*, routers share their routing table information with each other. Also referred to as "routing by rumor," each router forwards and receives updates from its direct neighbor. In the following figure, Router B shares information with Routers A and C. Router C shares routing information with Routers B and D. A distance vector describes the direction (port) and the distance (number of hops or other metric) to some other router. When a router receives information from another router, it increments whatever metric it is using. This process is called *distance accumulation*. Routers using this method know the distance between any two points in the network, but they do not know the exact topology of an internetwork.



Link-State Routing

With *link-state routing*, also known as shortest path first (SPF), each router maintains a database of topology information for the entire network.

Link-state routing provides better scaling than distance-vector routing because it only sends updates when there is a change in the network, and then it only sends information specific to the change that occurred. Distance vector uses regular updates and sends the whole routing table every time. Link-state routing also uses a hierarchical model, limiting the scope of route changes that occur.


Cisco Press and Cisco Systems work together to enhance classroom learning

Cisco Press works in conjunction with the Worldwide Education Group at Cisco Systems® to develop the only official books and classroom resources for the Cisco Networking Academy® Program. Together, Cisco Press and Cisco® provide an integrated learning environment that includes the online curriculum, trained instructors, and classroom textbooks to fulfill various learning styles.

Cisco Press Networking Academy Product Family Overview

There are three types of core Networking Academy textbooks that Cisco Press publishes: Companion Guides, Lab Companions, and Engineering Journals and Workbooks. These materials enhance the learning experience and lend support to all courses within the web-based curricula developed for the Cisco Networking Academy Program.

Companion Guides



Companion Guides are portable desk references of the course material to use anytime, anywhere. They are designed to reinforce course material, help focus on important concepts, and organize study time for exams. Key features include objectives, chapter summaries, margin notes, easy-to-understand figures and tables, well-defined key terms, "Check Your Understanding" review questions with answers, and Skill Builders.


"Overall, my experience with Cisco Press books has been very enjoyable. As a teacher I can use the textbook(s) as [either] the primary course materials or as a supplement to the online materials. I think the combination of text and online is put together very well."
—Richard B. Kirkland, CETsr, CCNA®, CCAI, Area Technical Trade Center

"I love to hear the student say 'Aha, so that's how it works'. Keep up the good work. I have never found a Cisco Press book I didn't enjoy."
—Santo Giabattari, Northwest Technical Institute, Networking Academy Instructor

"[What I like best about Cisco Press books is that] I have confidence in the accuracy of the information. It is clearly presented and easy to read."
—Bob Johnston, PTEC Clearwater, Networking Academy Instructor


"The Cisco Press books follow the Cisco Networking Academy Program very closely. [They] have numerous questions that have students verify their understanding of the material. The CDs that accompany the books have the e-Labs, which are valuable to students to practice commands."
—Jane Eckes, Central Virginia Community College, Networking Academy Instructor

Lab Companions



Convenient Lab companions provide the complete collection of lab exercises specifically written for Networking Academy courses. The labs are designed to give hands-on experience; each lab contains an introductory overview, a preparation/tools required section, explanations of commands, and step-by-step instructions to reinforce the concepts introduced in the online course and Companion Guide. Many labs also contain optional challenge exercises.


Engineering Journal and Workbooks



Engineering Journal and Workbooks are designed for some courses as a tool to provide additional exercises and activities, such as concept and focus questions and vocabulary exercises, to reinforce understanding of course topics.

For more information on Cisco Press titles, visit <http://www.pearson-books.com/ciscoacademy/>

Excerpts Taken from Cisco Networking Simplified



Available Now
Paul Della Maggiore and Jim Doherty
ISBN: 1-58720-074-0

This poster is a collection of excerpts from the Cisco Press title *Cisco Networking Simplified*. The book is an illustrated view of how networks operate, with an approach that answers, "How does it work?" This comprehensive review also covers topics like virtual private networks (VPNs), IP telephony, mobility, and storage area networks.

At-A-Glance—LAN Switching

Why Should I Care About Switching?

The advances in switching technology combined with the decrease in switch prices have made computer networks a common and increasingly important aspect of business today.

What Are the Problems to Solve?

Switches must learn about the network to make intelligent decisions. Due to the size and changing nature of networks, switches learned how to discover network address and keep track of network changes.

Switches must make decisions about what to do with traffic. The decisions are based on the switch's knowledge of the network.

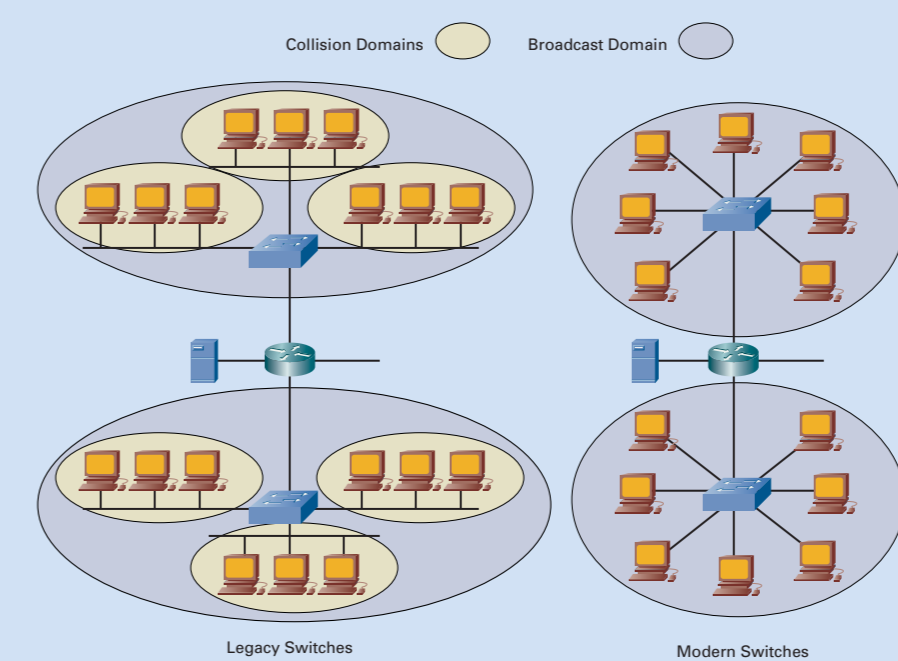
Switches must also have mechanisms for segmenting users into logical groupings to allow efficient provisioning of services.

Broadcast and Collision Domains

From time to time, a device on the network wants to communicate with all other "local" devices at the same time. Typically, this communication occurs when a device wants to query the network for an address, when a device is newly added to a network, or when there is a change in the network.

A group of devices that receive all broadcast messages from members within that group is called a *broadcast domain*. Network broadcast domains are typically segmented with Layer 3 devices (routers).

A group of devices that share a common access medium, and can therefore interfere with each other when transmitting simultaneously, define a *collision domain*. Traditionally, each broadcast domain had multiple collision domains. Modern switches, however, have a low price/performance ratio, making it feasible to dedicate a port to a single end device, effectively removing all collision domains.



Forwarding and Filtering

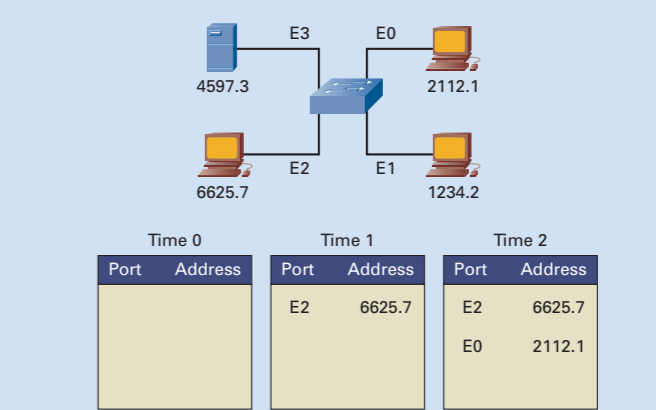
A switch always does something when it receives traffic. The preference is to send the traffic out a specific port (called *filtering*), but that only works when the location of the intended destination is known. When the destination address is not known, the switch forwards the traffic out every port, except the one on which the traffic was received. This process is called *flooding*.

From a network efficiency standpoint, it is much better for the network when the switch knows all the addresses on every port, but it is not always practical to enter this information manually. As the network grows and changes, all the port addresses are almost impossible to track.

Address Learning

A switch must therefore learn the addresses of the devices attached to it. It does so by inspecting the source address of all the traffic sent through it and then associates the port the traffic was received on with the Media Access Control (MAC) address listed. The following example illustrates this concept. (The MAC addresses, shown for clarity only, are not the correct format.)

- **Time 0**—The switch shown has an empty MAC address table.
- **Time 1**—The device attached to port 2 sends a message intended for the device on port 0. This message kicks off two actions within the switch. The switch now knows the address associated with the device on port 2, so it enters the information in its table; and because it does not have an association for the device the traffic is intended for (namely the computer on port 0), it floods the message out all ports except the one on which it was received.
- **Time 2**—The device on port 0 replies to the message. The switch now associates the source address of the message with port 0. This process happens all the time in every switch.




Virtual LANs

Virtual LANs (VLANs) provide the means to logically group several end stations with common sets of requirements. VLANs are independent of physical locations, meaning that two end stations connected to different switches on different floors can belong to the same VLAN. Typically, the logical grouping follows workgroup functions such as engineering or finance, but you can customize them as well.

With VLANs, it's easier to assign access rules and provision services to groups of users regardless of their physical location. For example, using VLANs, you can give all members of a project team access to project files by virtue of their VLAN membership. This ability also makes it easier to add or delete users without re-running cables or changing network addresses.

VLANs also create their own broadcast domains without the addition of Layer 3 devices.



At-A-Glance—Ethernet

Why Should I Care About Ethernet?

Ethernet was developed in 1972 as a way to connect newly invented computers to newly invented laser printers. Although recognized even at that time as a remarkable technology breakthrough, few people would have wagered that the ability to connect computers and devices would change communication on the same scale as the invention of the telephone and change business on the scale of the Industrial Revolution. Several competing protocols have emerged since 1972, but Ethernet remains the dominant standard for connecting computers into LANs.

What Are the Problems to Solve?

Ethernet is a shared resource where all end stations (computers, servers, etc.) all have access to the transmission medium at the same time. The result is that only one device can send information at a time. Given this limitation, there are two viable solutions:

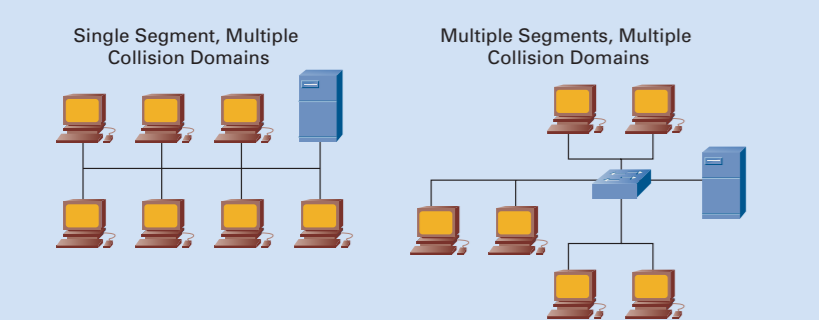
- Use a **sharing mechanism**—If all end stations are forced to share a common wire, then rules must ensure that each end station waits its turn before transmitting, or, in the event of simultaneous transmissions, follows rules for retransmitting.
- **Divide the shared segments and insulate them**—Another solution is to use devices that reduce the number of end stations sharing a resource at any given time.

Ethernet Collisions

In a traditional LAN, several users all share the same port on a network device and compete for resources (bandwidth). The main limitation of such a setup is that only one device can transmit at a time. Segments that share resources in this manner are called *collision domains* because if two or more devices transmit at the same time, the information collides and both endpoints must resend their information. Typically, the devices both begin a random countdown before attempting to retransmit. This method works well for a small number of users on a segment, each having relatively low bandwidth requirements. As the number of users increases, the efficiency of collision domains decreases sharply, to the point where overhead traffic (management and control) clogs the network.

Smaller Segments

You can divide segments to reduce the number of users and increase the bandwidth available to each user in the segment. Each new segment created results in a new collision domain. Traffic from one segment or collision domain does not interfere with other segments, thereby increasing the available bandwidth of each segment. In the following example, each segment has greater bandwidth, but all segments are still on a common backbone and must share the available bandwidth.




The basic tools for segmenting an Ethernet LAN into more collision domains follow:

- Bridges
- Routers
- Switches

This At-A-Glance sheet discusses segmenting using bridges and routers.

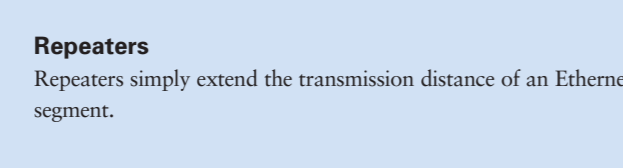
Ethernet Segment

A *segment* is the simplest form of network, where all devices are directly connected. In this type of arrangement, disconnecting or adding a computer disables the segment.




Hubs

Hubs enable you to add and remove computers without disabling the network but do not create additional collision domains.




Repeaters

Repeaters simply extend the transmission distance of an Ethernet segment.




Bridges

Bridges are simple Layer 2 devices that create new segments, resulting in fewer collisions. Bridges must learn the addresses of the computers on each segment to avoid forwarding traffic to the wrong port.



Switched Ethernet


You can think of a LAN switch as a high-speed, multipoint bridge with a brain. Switches not only give each end station a dedicated port (meaning there are no collisions), but they also allow end stations to transmit and receive at the same time, greatly increasing the efficiency of the LAN.



LAN Routers

LAN-based routers greatly extend the speed, distance, and intelligence of Ethernet LANs. Routers also allow traffic to travel along multiple paths.

Routers, however, do require a common protocol between the router and end stations.




At-A-Glance—Wireless LANs

Why Should I Care About Wireless LANs?

A wireless local-area network (WLAN) is a network of computers or terminals connected by radio frequencies. Unlike traditional LANs, WLAN users are free to move about while staying connected to the network.

Because of this mobility, WLANs offer business great flexibility when implementing a new network or when looking for new office space. You can implement a wireless LAN in a building not set up for traditional networking, saving the time and expense of making a new space business-ready. WLANs typically connect users to a corporate network, but they can also connect physically separated buildings. This implementation is referred to as a *building-to-building bridge system*.



The following are the advantages of WLANs over traditional LANs:

- With LANs, PCs must plug into Ethernet jacks. On WLANs, PCs can access the network from anywhere on the campus.
- With LANs, temporary networks are difficult to set up. With WLANs, temporary networks are easy to set up.
- Users on LANs typically share data files after work sessions due to a lack of connections. On WLANs, users can easily share data and files during work sessions.

What Are the Problems to Solve?

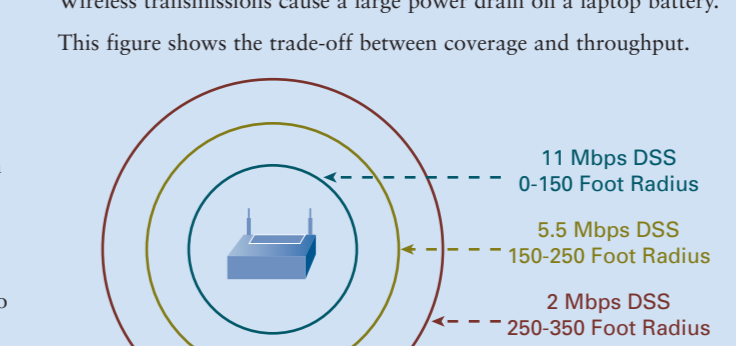
- Wireless LANs present network administrators with some new issues:
- Unlike fixed Ethernet, WLANs must trade off between throughput and the power consumption of mobile devices on battery power.
- One of the advantages of WLANs is mobility. Therefore, WLANs must employ schemes that allow users to remain connected as they move about a building or campus.
- WLANs present new security issues such as access control and data privacy.

In-Building Systems

In-building WLAN gives employees the flexibility to move about freely while staying connected to the network.


The number of access points (APs) required depends on the size of the building and the desired throughput. You must make trade-offs between power, battery life, and transmission throughput. Remember that in addition to receiving a signal, the PC must transmit a signal to the nearest AP.

Wireless transmissions cause a large power drain on a laptop battery. This figure shows the trade-off between coverage and throughput.



Direct Sequence Spread Spectrum (DSS) is a wireless multiplexing scheme for combining multiple signals into the same block of frequencies.


WLAN Roaming
Because wireless APs are relatively inexpensive, and the desire for bandwidth is high, most companies opt for deploying multiple APs with a reduced transmission radius and increased throughput. This solution introduces the need for a WLAN roaming scheme. *Roaming* describes switching from the control of one AP to another. You should position APs so there are no "dead spots." As a user moves away from one AP, the power and signal quality decrease. A good roaming plan ensures that as this happens, another AP signal is sufficiently strong enough to take control of the wireless connection. The network controls this "handoff" transparently to the user.



Wireless handoff can only occur on the same WLAN. If a user moves between two WLANs, connectivity is lost until the device authenticates on the new WLAN.

Building-to-Building Bridge Systems

Wireless bridges create a single LAN by linking remote networks together. For simple networks, the bridge connects to a hub or a switch on the LAN. If the network contains multiple subnetworks, the bridge is connected to a router. Wireless bridges are a convenient and cost-effective solution for rapidly growing companies or for users located in areas where a fixed connection is either expensive or impractical.



In some cases, building-to-building wireless bridges offer superior price and performance over the following competing technologies:

- **Direct cable connections:**
 - High installation costs
 - Difficulty overcoming physical barriers such as lakes, highways, and other buildings
 - Often require approval from local governments
- **Inflexible after deployment**
- **Telephone-line connections**
 - High monthly service fees
 - High installation and equipment costs
- **Microwave connections**
 - Expensive
 - Require licensing
 - Difficult to install

Security Issues and Options

Security is a major concern for WLANs. The two main security issues for WLANs are:

- **Access control**—Because WLANs use radio waves for access, any WLAN client in the area is capable of accessing the network. Some hackers access networks while sitting in a car outside of a building. Businesses should protect their networks with centralized user-authentication schemes to protect against unauthorized access.
- **Privacy**—Privacy is also an issue with WLANs. Unlike fixed connections, which send information point to point, WLANs broadcast information everywhere. Hackers can "snoop" this information out of the air. Therefore, it is essential to encrypt the data packets that transmit through the air.