CHAPTER 1

Unicast Routing

This chapter introduces the underlying concepts for Layer 3 unicast routing protocols in Cisco 1000 Series Connected Grid Routers (*hereafter* referred to as the Cisco CG-OS router) and WAN backhaul redundancy.The system software for the router is identified as the Cisco CG-OS software.

This chapter includes the following sections:

- Information About Layer 3 Unicast Routing, page 1-1
- Routing Algorithms, page 1-7
- Summary of Layer 3 Unicast Routing Features, page 1-9
- WAN Backhaul Redundancy, page 1-10

Information About Layer 3 Unicast Routing

Layer 3 unicast routing involves determining optimal routing paths. You can use routing algorithms to calculate the optimal path from the router to a destination. This calculation depends on the algorithm selected, route metrics, and other considerations such as load balancing and alternate path discovery.

This section includes the following topics:

- Routing Fundamentals, page 1-2
- Routing Metrics, page 1-3
- Router IDs, page 1-5
- Autonomous Systems, page 1-5
- Convergence, page 1-6
- Load Balancing and Equal Cost Multipath, page 1-6
- Administrative Distance, page 1-6
- Stub Routing, page 1-6

Routing Fundamentals

Routing protocols use a metric to evaluate the best path to the destination. A metric is a standard of measurement, such as a path bandwidth, that routing algorithms use to determine the optimal path to a destination. To aid path determination, routing algorithms initialize and maintain routing tables that contain route information such as the IP destination address, the address of the next router, or the next hop. Destination and next-hop associations tell a router that an IP destination can be reached optimally by sending the packet to a particular router that represents the next hop on the way to the final destination. When a router receives an incoming packet, it checks the destination address and attempts to associate this address with the next hop.

Routing tables can contain other information, such as the data about the desirability of a path. Routers compare metrics to determine optimal routes, and these metrics differ depending on the design of the routing algorithm used. See the Routing Metrics, page 1-3.

Routers communicate with one another and maintain their routing tables by transmitting a variety of messages. The routing update message is one such message that consists of all or a portion of a routing table. By analyzing routing updates from all other routers, a router can build a detailed picture of the network topology. A link-state advertisement, which is another example of a message sent between routers, informs other routers of the link state of the sending router. You can also use link information to enable routers to determine optimal routes to network destinations. For more information, see the Routing Algorithms, page 1-7.

Figure 1-1 Packet Header Updates Through a Network

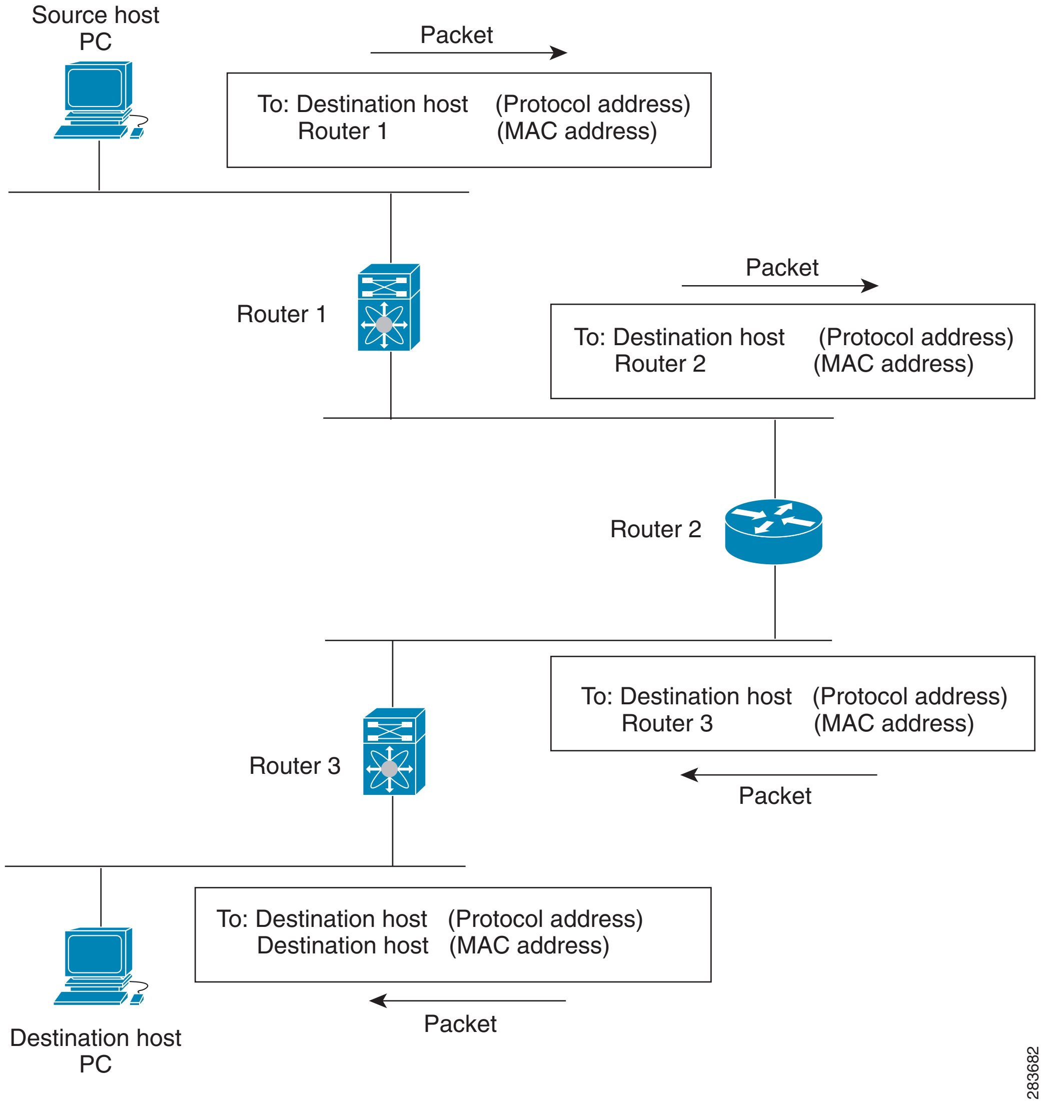
Source host
PC
Packet
To: Destination host (Protocol address)
Router 1 (MAC address)
Router 1
Packet
To: Destination host (Protocol address)
Router 2 (MAC address)
Router 2
To: Destination host (Protocol address)
Router 3 (MAC address)
Router 3
Packet
To: Destination host (Protocol address)
Destination host (MAC address)
Packet
Destination host
PC
283682

Routing Metrics

Routing algorithms use many different metrics to determine the best route. Sophisticated routing algorithms can base route selection on multiple metrics.

This section includes the following metrics:

- Path Length, page 1-4
- Reliability, page 1-4
- Routing Delay, page 1-4
- Bandwidth, page 1-4
- Load, page 1-4
- Communication Cost, page 1-4

Path Length

The path length is the most common routing metric. Some routing protocols allow you to assign arbitrary costs to each network link. In this case, the path length is the sum of the costs associated with each link traversed. Other routing protocols define the hop count, which is a metric that specifies the number of passes through internetworking products, such as routers, that a packet must take from a source to a destination.

Reliability

The reliability, in the context of routing algorithms, is the dependability (in terms of the bit-error rate) of each network link. Some network links might go down more often than others. After a network fails, certain network links might be repaired more easily or more quickly than other links. The reliability factors that you can take into account when assigning the reliability rating are arbitrary numeric values that you usually assign to network links.

Routing Delay

The routing delay is the length of time required to move a packet from a source to a destination through the internetwork. The delay depends on many factors, including the bandwidth of intermediate network links, the port queues at each router along the way, the network congestion on all intermediate network links, and the physical distance that the packet must travel. Because the routing delay is a combination of several important variables, it is a common and useful metric.

Bandwidth

The bandwidth is the available traffic capacity of a link. Although the bandwidth is the maximum attainable throughput on a link, routes through links with greater bandwidth do not necessarily provide better routes than routes through slower links. For example, if a faster link is busier, the actual time required to send a packet to the destination could be greater.

Load

The load is the degree to which a network resource, such as a router, is busy. You can calculate the load in a variety of ways, including CPU usage and packets processed per second. Monitoring these parameters on a continual basis can be resource intensive.

Communication Cost

The communication cost is a measure of the operating cost to route over a link. The communication cost is another important metric, especially if you do not care about performance as much as operating expenditures. For example, the line delay for a private line might be longer than a public line, but you can send packets over your private line rather than through the public lines that cost money for usage time.

Router IDs

Each routing process has an associated router ID. You can configure the router ID to any interface. The Cisco CG-OS router supports cellular, Ethernet (Fast Ethernet and Gigabit Ethernet), and WiMax interfaces. When you do not configure the router ID, the Cisco CG-OS router selects the router ID based on the following criteria:

- The Cisco CG-OS router prefers loopback0 over any other interface. When loopback0 does not exist, then the router prefers the first loopback interface over any other interface type.
- When you do not configure a loopback interface, the Cisco CG-OS router uses the first interface in the configuration file as the router ID. When you configure any loopback interface after the Cisco CG-OS software selects the router ID, the loopback interface becomes the router ID. When the loopback interface is not loopback0 and you configure loopback0 with an IP address, the router ID changes to the IP address of loopback0.
- When the interface that the router ID is based on changes, that new IP address becomes the router ID. When any other interface changes its IP address, there is no router ID change.

Related Topics

Cisco 1000 Series Connected Grid Routers WiFi Software Configuration Guide

Cisco 1240 Connected Grid Router Hardware Installation Guide

Cisco Connected Grid Cellular 3G Module for CGR1000 Series Installation and Configuration Guide

Cisco Connected Grid WiMAX Module for CGR1000 Series Installation and Configuration Guide

Autonomous Systems

An autonomous system (AS) is a network controlled by a single technical administration entity. Autonomous systems divide global external networks into individual routing domains, where local routing policies are applied. This organization simplifies routing domain administration and simplifies consistent policy configuration.

Each autonomous system can support multiple interior routing protocols that dynamically exchange routing information through route redistribution. The Regional Internet Registries assign a unique number to each public autonomous system that directly connects to the Internet. This autonomous system number (AS number) identifies both the routing process and the autonomous system.

The Cisco CG-OS router supports 4-byte AS numbers. Table 1-1 lists the AS number ranges.

Table 1-1 AS Numbers

2-Byte Numbers	4-Byte Numbers in AS.dot Notation	4-Byte Numbers in plaintext Notation	Purpose
1 to 64511	0.1 to 0.64511	1 to 64511	Public AS (assigned by RIR)[1]
64512 to 65534	0.64512 to 0.65534	64512 to 65534	Private AS (assigned by local administrator)
65535	0.65535	65535	Reserved
N/A	1.0 to 65535.65535	65536 to 4294967295	Public AS (assigned by RIR)

[1] RIR=Regional Internet Registries

Private autonomous system numbers are used for internal routing domains but must be translated by the Cisco CG-OS router for traffic that is routed out to the Internet. It is important not to configure routing protocols to advertise private autonomous system numbers to external networks. By default, the Cisco CG-OS router does not remove private autonomous system numbers from routing updates.

Note The autonomous system number assignment for public and private networks is governed by the Internet Assigned Number Authority (IANA). For information about autonomous system numbers, including the reserved number assignment, or to apply to register an autonomous system number, see this URL: http://www.iana.org/

Convergence

A key aspect to measure for any routing algorithm is how much time a router takes to react to network topology changes. When a part of the network changes for any reason, such as a link failure, the routing information in different routers might not match. Some routers will have updated information about the changed topology, while other routers will still have the old information. The convergence is the amount of time before all routers in the network have updated, matching routing information. The convergence time varies depending on the routing algorithm. Fast convergence minimizes the chance of lost packets caused by inaccurate routing information.

Load Balancing and Equal Cost Multipath

Routing protocols can use load balancing or equal cost multipath (ECMP) to share traffic across multiple paths.When a router learns multiple routes to a specific network, it installs the route with the lowest administrative distance in the routing table. If the router receives and installs multiple paths with the same administrative distance and cost to a destination, load balancing can occur. Load balancing distributes the traffic across all the paths, sharing the load. The number of paths used is limited by the number of entries that the routing protocol puts in the routing table.

Administrative Distance

An administrative distance is a rating of the trustworthiness of a routing information source. A higher value indicates a lower trust rating. Typically, a route can be learned through more than one protocol. Administrative distance is used to discriminate between routes learned from more than one protocol. The route with the lowest administrative distance is installed in the IP routing table.

Stub Routing

You can use stub routing in a hub-and-spoke network topology, where one or more end (stub) networks are connected to a remote router (the spoke) that is connected to one or more distribution routers (the hub). The remote router is adjacent only to one or more distribution routers. The only route for IP traffic to follow into the remote router is through a distribution router. This type of configuration is commonly used in WAN topologies in which the distribution router is directly connected to a WAN. The distribution router can be connected to many more remote routers. Often, the distribution router is connected to 100 or more remote routers. In a hub-and-spoke topology, the remote router must forward all nonlocal traffic to a distribution router, so it becomes unnecessary for the remote router to hold a complete routing table. Generally, the distribution router sends only a default route to the remote router.

Only specified routes are propagated from the remote (stub) router. The stub router responds to all queries for summaries, connected routes, redistributed static routes, external routes, and internal routes with the message “inaccessible.” A router that is configured as a stub sends a special peer information packet to all neighboring routers to report its status as a stub router.

Any neighbor that receives a packet that informs it of the stub status does not query the stub router for any routes, and a router that has a stub peer does not query that peer. The stub router depends on the distribution router to send the proper updates to all peers.

Figure 1-2 shows a simple hub-and-spoke configuration.

Figure 1-2 Simple Hub-and-Spoke Network

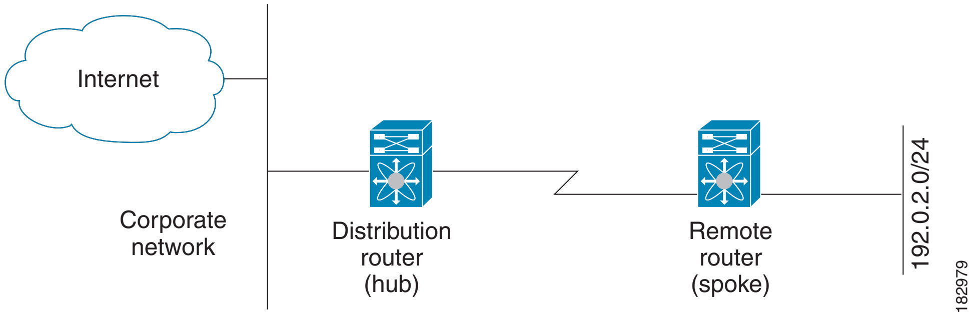

Stub routing does not prevent routes from being advertised to the remote router. Figure 1-2 shows that the remote router can access the corporate network and the Internet through the distribution router only. A full route table on the remote router, in this example, serves no functional purpose because the path to the corporate network and the Internet is always through the distribution router. A larger route table reduces only the amount of memory required by the remote router. The bandwidth and memory used can be lessened by summarizing and filtering routes in the distribution router. In this network topology, the remote router does not need to receive routes that have been learned from other networks because the remote router must send all nonlocal traffic, regardless of its destination, to the distribution router. To configure a true stub network, you should configure the distribution router to send only a default route to the remote router.

The Open Shortest Path First (OSPF) protocol supports stub areas. For more information on OSPF, see OSPF, page 1-9.

Routing Algorithms

Routing algorithms determine how a router gathers and reports reachability information, how it deals with topology changes, and how it determines the optimal route to a destination. Various types of routing algorithms exist, and each algorithm has a different impact on network and router resources. Routing algorithms use a variety of metrics that affect calculation of optimal routes. You can classify routing algorithms by type, such as static or dynamic, and interior or exterior.

This section includes the following topics:

- Static Routes and Dynamic Routing Protocols, page 1-8
- Interior and Exterior Gateway Protocols, page 1-8
- Link-State Protocols, page 1-8

Static Routes and Dynamic Routing Protocols

Static routes are route table entries that you manually configure. These static routes do not change unless you reconfigure them. Static routes are simple to design and work well in environments where network traffic is relatively predictable and where network design is relatively simple.

Because static routing systems cannot react to network changes, do not use them for large, constantly changing networks. Most routing protocols today use dynamic routing algorithms that adjust to changing network circumstances by analyzing incoming routing update messages. When the message indicates that a network change has occurred, the routing software recalculates routes and sends out new routing update messages. These messages permeate the network, triggering routers to rerun their algorithms and change their routing tables accordingly.

You can supplement dynamic routing algorithms with static routes where appropriate. For example, you can configure each subnetwork with a static route to the IP default gateway or router of last resort (the router to which all unrouteable packets are sent).

Interior and Exterior Gateway Protocols

You can separate networks into unique routing domains or autonomous systems. An autonomous system is a portion of an internetwork under common administrative authority that is regulated by a particular set of administrative guidelines. Routing protocols that route between autonomous systems are called exterior gateway protocols or interdomain protocols. Routing protocols used within an autonomous system are called interior gateway protocols or intradomain protocols. OSPF is an example of an interior gateway protocol that the Cisco CG-OS router supports.

Note The Cisco CG-OS router does not support any exterior gateway protocols.

Link-State Protocols

The link-state protocols, also known as shortest path first (SPF), share information with neighboring routers. Each router builds a link-state advertisement (LSA) that contains information about each link and directly connected neighbor router.

Each LSA has a sequence number. When a router receives an LSA and updates its link-state database, the LSA floods all adjacent neighbors. When a router receives two LSAs with the same sequence number (from the same router), the router does not flood its neighbors with the last LSA received because it wants to prevent an LSA update loop.

Discovering neighbors and establishing adjacency is an important part of a link state protocol. Neighbors use special Hello packets to discover on another. Hello packets also serve as keepalive notifications for each neighbor router. Adjacency establishes a common set of operating parameters for the link-state protocol between neighbor routers.

When a router receives an LSA, the router adds the LSA to its link-state database. Each entry consists of the following parameters:

- Router ID (for the router that originated the LSA)
- Neighbor ID
- Link cost
- Sequence number of the LSA

- Age of the LSA entry

The router runs the SPF algorithm on the link-state database, building the shortest path tree for that router. This SPF tree is used to populate the routing table.

In link-state algorithms, each router builds a picture of the entire network in its routing tables. The link-state algorithms send small updates everywhere, while distance vector algorithms send larger updates only to neighboring routers.

Because link-state algorithms converge more quickly, they are less likely to cause routing loops than distance vector algorithms. However, link-state algorithms require more CPU power and memory than distance vector algorithms and they can be more expensive to implement and support. Link-state protocols are generally more scalable than distance vector protocols.

OSPF is an example of a link-state protocol.

Summary of Layer 3 Unicast Routing Features

This section provides a brief introduction to the Layer 3 unicast features and protocols supported in the Cisco CG-OS router.

This section includes the following topics:

- IPv4 and IPv6, page 1-9
- IP Services, page 1-9
- OSPF, page 1-9
- Static Routing, page 1-10

IPv4 and IPv6

Layer 3 routing employs the IPv4 and/or the IPv6 protocol. IPv6 increases the number of network address bits to 128 bits from the 32 bits employed by IPv4.

The Cisco CG-OS router supports Address Resolution Protocol (ARP) and Internet Control Message Protocol (ICMP) when employing IPv4.

The Cisco CG-OS router supports Internet Control Message Protocol for IPv6 (ICMPv6) and Neighbor Discovery (ND) when employing IPv6.

For more information, see Chapter 2, "Configuring IPv4" and Chapter 3, "Configuring IPv6."

IP Services

IP Services addresses Domain Name System (DNS) clients. For more information, see Chapter 4, "Configuring IP Services."

OSPF

The Open Shortest Path First (OSPF) protocol is a link-state routing protocol that exchanges network reachability information within an autonomous system (AS). Each OSPF router advertises information about its active links to its neighbor routers. Link information consists of the link type, the link metric, and the neighbor router that is connected to the link. The advertisements that contain this link information are called link-state advertisements (LSAs). The Cisco CG-OS router supports both OSPFv2 for IPv4 networks and OSPFv3 for IPv6 networks.

For more information, see Chapter 5, "Configuring OSPFv2" and Chapter 6, "Configuring OSPFv3."

Static Routing

Static routing allows you to configure a fixed route to a destination. This feature is useful for small networks where the topology is simple. Static routing is also used with other routing protocols to control default routes. For more information, see Chapter 8, "Configuring Static Routing."

WAN Backhaul Redundancy

Redundant WAN backhauls can be configured on the Cisco CG-OS router within an Open Shortest Path First version 2 (OSPFv2) area by assigning link costs to cellular (3G) and WiMax interfaces. The interface with the lower assigned link cost remains the primary link until that link goes down; and, then traffic automatically goes to the secondary link with the next lowest cost. In cases where the link with the higher cost fails, no redirect of traffic occurs because the Cisco CG-OS router by default routes all traffic to the link with the lowest cost.

For more information, see Chapter 7, "Configuring WAN Backhaul Redundancy."