Routing Overview

This chapter describes underlying concepts of how routing behaves within the ASA, and the routing protocols that are supported.

This chapter includes the following sections:

Information About Routing

Routing is the act of moving information across an internetwork from a source to a destination. Along the way, at least one intermediate node typically is encountered. Routing involves two basic activities: determining optimal routing paths and transporting information groups (typically called packets) through an internetwork. In the context of the routing process, the latter of these is referred to as packet switching. Although packet switching is relatively straightforward, path determination can be very complex.

This section includes the following topics:

Switching

Switching algorithms is relatively simple; it is the same for most routing protocols. In most cases, a host determines that it must send a packet to another host. Having acquired a router address by some means, the source host sends a packet addressed specifically to a router physical (Media Access Control [MAC]-layer) address, this time with the protocol (network layer) address of the destination host.

As it examines the packet destination protocol address, the router determines that it either knows or does not know how to forward the packet to the next hop. If the router does not know how to forward the packet, it typically drops the packet. If the router knows how to forward the packet, however, it changes the destination physical address to that of the next hop and transmits the packet.

The next hop may be the ultimate destination host. If not, the next hop is usually another router, which executes the same switching decision process. As the packet moves through the internetwork, its physical address changes, but its protocol address remains constant.

Path Determination

Routing protocols use metrics to evaluate what path will be the best for a packet to travel. A metric is a standard of measurement, such as path bandwidth, that is used by routing algorithms to determine the optimal path to a destination. To aid the process of path determination, routing algorithms initialize and maintain routing tables, which include route information. Route information varies depending on the routing algorithm used.

Routing algorithms fill routing tables with a variety of information. Destination or next hop associations tell a router that a particular destination can be reached optimally by sending the packet to a particular router representing 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 a next hop.

Routing tables also can include other information, such as 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.

Routers communicate with one another and maintain their routing tables through the transmission of a variety of messages. The routing update message is one such message that generally 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 network topology. A link-state advertisement, another example of a message sent between routers, informs other routers of the state of the sender links. Link information also can be used to build a complete picture of network topology to enable routers to determine optimal routes to network destinations.

Note Asymmetric routing is only supported for Active/Active failover in multiple context mode. For more information, see the "Configuring Active/Active Failover" section on page 1-9.

Supported Route Types

There are several route types that a router can use. The ASA uses the following route types:

Static Versus Dynamic

Static routing algorithms are hardly algorithms at all, but are table mappings established by the network administrator before the beginning of routing. These mappings do not change unless the network administrator alters them. Algorithms that use 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, they generally are considered unsuitable for large, constantly changing networks. Most of the dominant routing algorithms are dynamic routing algorithms, which adjust to changing network circumstances by analyzing incoming routing update messages. If 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, stimulating routers to rerun their algorithms and change their routing tables accordingly.

Dynamic routing algorithms can be supplemented with static routes where appropriate. A router of last resort (a router to which all unroutable packets are sent), for example, can be designated to act as a repository for all unroutable packets, ensuring that all messages are at least handled in some way.

Single-Path Versus Multipath

Some sophisticated routing protocols support multiple paths to the same destination. Unlike single-path algorithms, these multipath algorithms permit traffic multiplexing over multiple lines. The advantages of multipath algorithms are substantially better throughput and reliability, which is generally called load sharing.

Flat Versus Hierarchical

Some routing algorithms operate in a flat space, while others use routing hierarchies. In a flat routing system, the routers are peers of all others. In a hierarchical routing system, some routers form what amounts to a routing backbone. Packets from nonbackbone routers travel to the backbone routers, where they are sent through the backbone until they reach the general area of the destination. At this point, they travel from the last backbone router through one or more nonbackbone routers to the final destination.

Routing systems often designate logical groups of nodes, called domains, autonomous systems, or areas. In hierarchical systems, some routers in a domain can communicate with routers in other domains, while others can communicate only with routers within their domain. In very large networks, additional hierarchical levels may exist, with routers at the highest hierarchical level forming the routing backbone.

The primary advantage of hierarchical routing is that it mimics the organization of most companies and therefore supports their traffic patterns well. Most network communication occurs within small company groups (domains). Because intradomain routers need to know only about other routers within their domain, their routing algorithms can be simplified, and, depending on the routing algorithm being used, routing update traffic can be reduced accordingly.

Link-State Versus Distance Vector

Link-state algorithms (also known as shortest path first algorithms) flood routing information to all nodes in the internetwork. Each router, however, sends only the portion of the routing table that describes the state of its own links. In link-state algorithms, each router builds a picture of the entire network in its routing tables. Distance vector algorithms (also known as Bellman-Ford algorithms) call for each router to send all or some portion of its routing table, but only to its neighbors. In essence, link-state

algorithms send small updates everywhere, while distance vector algorithms send larger updates only to neighboring routers. Distance vector algorithms know only about their neighbors. Typically, link-state algorithms are used in conjunction with OSPF routing protocols.

How Routing Behaves Within the ASA

The ASA uses both routing table and XLATE tables for routing decisions. To handle destination IP translated traffic, that is, untranslated traffic, the ASA searches for existing XLATE, or static translation to select the egress interface.

This section includes the following topics:

Egress Interface Selection Process

The selection process follows these steps:

1. If a destination IP translating XLATE already exists, the egress interface for the packet is determined from the XLATE table, but not from the routing table.
2. If a destination IP translating XLATE does not exist, but a matching static translation exists, then the egress interface is determined from the static route and an XLATE is created, and the routing table is not used.
3. If a destination IP translating XLATE does not exist and no matching static translation exists, the packet is not destination IP translated. The ASA processes this packet by looking up the route to select the egress interface, then source IP translation is performed (if necessary).

   For regular dynamic outbound NAT, initial outgoing packets are routed using the route table and then creating the XLATE. Incoming return packets are forwarded using existing XLATE only. For static NAT, destination translated incoming packets are always forwarded using existing XLATE or static translation rules.

Next Hop Selection Process

After selecting the egress interface using any method described previously, an additional route lookup is performed to find out suitable next hop(s) that belong to a previously selected egress interface. If there are no routes in the routing table that explicitly belong to a selected interface, the packet is dropped with a level 6 syslog message 110001 generated (no route to host), even if there is another route for a given destination network that belongs to a different egress interface. If the route that belongs to a selected egress interface is found, the packet is forwarded to the corresponding next hop.

Load sharing on the ASA is possible only for multiple next hops available using a single egress interface. Load sharing cannot share multiple egress interfaces.

If dynamic routing is in use on the ASA and the route table changes after XLATE creation (for example, route flap), then destination translated traffic is still forwarded using the old XLATE, not via the route table, until XLATE times out. It may be either forwarded to the wrong interface or dropped with a level 6 syslog message 110001 generated (no route to host), if the old route was removed from the old interface and attached to another one by the routing process.

The same problem may happen when there are no route flaps on the ASA itself, but some routing process is flapping around it, sending source-translated packets that belong to the same flow through the ASA using different interfaces. Destination-translated return packets may be forwarded back using the wrong egress interface.

This issue has a high probability in some security traffic configurations, where virtually any traffic may be either source-translated or destination-translated, depending on the direction of the initial packet in the flow. When this issue occurs after a route flap, it can be resolved manually by using the **clear xlate** command, or automatically resolved by an XLATE timeout. The XLATE timeout may be decreased if necessary. To ensure that this issue rarely occurs, make sure that there are no route flaps on the ASA and around it. That is, ensure that destination-translated packets that belong to the same flow are always forwarded the same way through the ASA.

Supported Internet Protocols for Routing

The ASA supports several Internet protocols for routing. Each protocol is briefly described in this section.

- Enhanced Interior Gateway Routing Protocol (EIGRP)

  EIGRP is a Cisco proprietary protocol that provides compatibility and seamless interoperation with IGRP routers. An automatic-redistribution mechanism allows IGRP routes to be imported into Enhanced IGRP, and vice versa, so it is possible to add Enhanced IGRP gradually into an existing IGRP network.

  For more information about configuring EIGRP, see the "Configuring EIGRP" section on page 1-3.

- Open Shortest Path First (OSPF)

  OSPF is a routing protocol developed for Internet Protocol (IP) networks by the interior gateway protocol (IGP) working group of the Internet Engineering Task Force (IETF). OSPF uses a link-state algorithm to build and calculate the shortest path to all known destinations. Each router in an OSPF area includes an identical link-state database, which is a list of each of the router usable interfaces and reachable neighbors.

  For more information about configuring OSPF, see the "Configuring OSPFv2" section on page 1-5.

- Routing Information Protocol (RIP)

  RIP is a distance-vector protocol that uses hop count as its metric. RIP is widely used for routing traffic in the global Internet and is an interior gateway protocol (IGP), which means that it performs routing within a single autonomous system.

  For more information about configuring RIP, see the "Configuring RIP" section on page 1-4.

Information About the Routing Table

This section includes the following topics:

- Dynamic Routing in Multiple Context Mode, page 1-10

Displaying the Routing Table

To view the entries in the routing table, enter the following command:

```
hostname# show route

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is 10.86.194.1 to network 0.0.0.0

S    10.1.1.0 255.255.255.0 [3/0] via 10.86.194.1, outside
C    10.86.194.0 255.255.254.0 is directly connected, outside
S*   0.0.0.0 0.0.0.0 [1/0] via 10.86.194.1, outside
```

On the ASA 5505, the following route is also shown. It is the internal loopback interface, which is used by the VPN hardware client feature for individual user authentication.

```
C 127.1.0.0 255.255.0.0 is directly connected, _internal_loopback
```

How the Routing Table Is Populated

The ASA routing table can be populated by statically defined routes, directly connected routes, and routes discovered by the RIP, EIGRP, and OSPF routing protocols. Because the ASA can run multiple routing protocols in addition to having static and connected routes in the routing table, it is possible that the same route is discovered or entered in more than one manner. When two routes to the same destination are put into the routing table, the one that remains in the routing table is determined as follows:

- If the two routes have different network prefix lengths (network masks), then both routes are considered unique and are entered into the routing table. The packet forwarding logic then determines which of the two to use.

  For example, if the RIP and OSPF processes discovered the following routes:

  – RIP: 192.168.32.0/24

  – OSPF: 192.168.32.0/19

  Even though OSPF routes have the better administrative distance, both routes are installed in the routing table because each of these routes has a different prefix length (subnet mask). They are considered different destinations and the packet forwarding logic determines which route to use.

- If the ASA learns about multiple paths to the same destination from a single routing protocol, such as RIP, the route with the better metric (as determined by the routing protocol) is entered into the routing table.

Metrics are values associated with specific routes, ranking them from most preferred to least preferred. The parameters used to determine the metrics differ for different routing protocols. The path with the lowest metric is selected as the optimal path and installed in the routing table. If there are multiple paths to the same destination with equal metrics, load balancing is done on these equal cost paths.

- If the ASA learns about a destination from more than one routing protocol, the administrative distances of the routes are compared, and the routes with lower administrative distance are entered into the routing table.

Administrative Distances for Routes

You can change the administrative distances for routes discovered by or redistributed into a routing protocol. If two routes from two different routing protocols have the same administrative distance, then the route with the lower *default* administrative distance is entered into the routing table. In the case of EIGRP and OSPF routes, if the EIGRP route and the OSPF route have the same administrative distance, then the EIGRP route is chosen by default.

Administrative distance is a route parameter that the ASA uses to select the best path when there are two or more different routes to the same destination from two different routing protocols. Because the routing protocols have metrics based on algorithms that are different from the other protocols, it is not always possible to determine the best path for two routes to the same destination that were generated by different routing protocols.

Each routing protocol is prioritized using an administrative distance value. Table 1-1 shows the default administrative distance values for the routing protocols supported by the ASA.

Table 1-1 Default Administrative Distance for Supported Routing Protocols

Route Source	Default Administrative Distance
Connected interface	0
Static route	1
EIGRP Summary Route	5
Internal EIGRP	90
OSPF	110
RIP	120
EIGRP external route	170
Unknown	255

The smaller the administrative distance value, the more preference is given to the protocol. For example, if the ASA receives a route to a certain network from both an OSPF routing process (default administrative distance - 110) and a RIP routing process (default administrative distance - 120), the ASA chooses the OSPF route because OSPF has a higher preference. In this case, the router adds the OSPF version of the route to the routing table.

In this example, if the source of the OSPF-derived route was lost (for example, due to a power shutdown), the ASA would then use the RIP-derived route until the OSPF-derived route reappears.

The administrative distance is a local setting. For example, if you use the **distance-ospf** command to change the administrative distance of routes obtained through OSPF, that change would only affect the routing table for the ASA on which the command was entered. The administrative distance is not advertised in routing updates.

Administrative distance does not affect the routing process. The EIGRP, OSPF, and RIP routing processes only advertise the routes that have been discovered by the routing process or redistributed into the routing process. For example, the RIP routing process advertises RIP routes, even if routes discovered by the OSPF routing process are used in the ASA routing table.

Backup Routes

A backup route is registered when the initial attempt to install the route in the routing table fails because another route was installed instead. If the route that was installed in the routing table fails, the routing table maintenance process calls each routing protocol process that has registered a backup route and requests them to reinstall the route in the routing table. If there are multiple protocols with registered backup routes for the failed route, the preferred route is chosen based on administrative distance.

Because of this process, you can create floating static routes that are installed in the routing table when the route discovered by a dynamic routing protocol fails. A floating static route is simply a static route configured with a greater administrative distance than the dynamic routing protocols running on the ASA. When the corresponding route discovered by a dynamic routing process fails, the static route is installed in the routing table.

How Forwarding Decisions Are Made

Forwarding decisions are made as follows:

- If the destination does not match an entry in the routing table, the packet is forwarded through the interface specified for the default route. If a default route has not been configured, the packet is discarded.
- If the destination matches a single entry in the routing table, the packet is forwarded through the interface associated with that route.
- If the destination matches more than one entry in the routing table, and the entries all have the same network prefix length, the two entries with identical network prefixes and different interfaces cannot coexist in the routing table.
- If the destination matches more than one entry in the routing table, and the entries have different network prefix lengths, then the packet is forwarded out of the interface associated with the route that has the longer network prefix length.

For example, a packet destined for 192.168.32.1 arrives on an interface of an ASA with the following routes in the routing table:

```
hostname# show route
         ....
         R   192.168.32.0/24 [120/4] via 10.1.1.2
         O   192.168.32.0/19 [110/229840] via 10.1.1.3
         ....
```

In this case, a packet destined to 192.168.32.1 is directed toward 10.1.1.2, because 192.168.32.1 falls within the 192.168.32.0/24 network. It also falls within the other route in the routing table, but the 192.168.32.0/24 has the longest prefix within the routing table (24 bits verses 19 bits). Longer prefixes are always preferred over shorter ones when forwarding a packet.

Dynamic Routing and Failover

Because static routing systems cannot react to network changes, they generally are considered unsuitable for large, constantly changing networks. Most of the dominant routing algorithms are dynamic routing algorithms, which adjust to changing network circumstances by analyzing incoming routing update messages. If 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, stimulating routers to rerun their algorithms and change their routing tables accordingly.

Dynamic routing algorithms can be supplemented with static routes where appropriate. A router of last resort (a router to which all unroutable packets are sent), for example, can be designated to act as a repository for all unroutable packets, ensuring that all messages are at least handled in some way.

Dynamic routes are synchronized on the standby unit when the routing table changes on the active unit, which means that all additions, deletions, or changes on the active unit are immediately propagated to the standby unit. If the standby unit becomes active after the primary unit has been active for a period of time, routes become synchronized as a part of the failover bulk synchronization process, so the routing table on the active/standby failover pair should appear the same.

For more information about static routes and how to configure them, see the "Configuring Static and Default Routes" section on page 1-2.

Dynamic Routing and Clustering

Dynamic routing is fully integrated in a cluster, and routes are shared across units (up to eight units are allowed in a cluster). Routing table entries are also replicated across units in a cluster.

When one unit transitions from the slave to the master, the epoch number (32-bit sequence number) for the RIB table is incremented. After the transition, the new master unit initially has RIB table entries that are the mirror image of the previous master unit. In addition, the reconvergence timer starts on the new master unit. When the epoch number for the RIB table is incremented, all existing entries are considered stale. Forwarding of IP packets continues as normal. On the new master unit, dynamic routing protocols start to either update existing route entries or create new route entries with the new epoch number. These modified or new entries with the current epoch number indicate that they have been refreshed and are synchronized to all slave units. After the reconvergence timer has expired, old entries from the RIB table are removed. RIB table entries for OSPF routes, RIP routes, and EIGRP routes are synchronized to the slave units.

Bulk synchronization occurs only when a unit joins a cluster and is from the master unit to a joining unit.

For dynamic routing updates, when the master unit learns a new route through OSPF, RIP or EIGRP, the master unit sends those updates to all slave units through reliable message transmission. Slave units update their RIB tables after they receive a cluster route update message.

For the supported dynamic routing protocols (OSPF, RIP, and EIGRP), routing packets from layer 2 load balancing interfaces on the slave units are forwarded to the master unit. Only the master unit sees and processes dynamic routing protocol packets. When the slave unit requests a bulk synchronization, all routing entries learned through layer 2 load balancing interfaces are replicated.

When new routing entries are learned through layer 2 load balancing interfaces on the master unit, the new entries are broadcast to all slave units. When existing routing entries are modified because of a network topology change, the modified entries are also synchronized to all slave units. When existing routing entries are removed because of a network topology change, the removed entries are also synchronized to all slave units.

When a combination of layer 2 and layer 3 load balancing interfaces are deployed and configured for dynamic routing, the slave units only have partial topology and neighbor information (including details that were obtained through layer 3 load balancing interfaces) in the routing process because only RIB table entries are synchronized from the master unit for layer 2 load balancing interfaces. You must configure the network to have layer 2 and layer 3 belong to different routing processes and redistribute the load from each routing process.

Table 1-2 provides a summary of the supported configurations. Yes indicates that the combination of two processes (one process to layer 2 and one process to layer 3) work, No indicates that the combination of two processes does not work.

Table 1-2 Summary of Supported Configurations

Layer 2 or Layer 3	OSPF (Layer 3)	EIGRP (Layer 3)	RIP (Layer 3)
OSPF (layer 2)	Yes	Yes	Yes
EIGRP (layer 2)	Yes	No	Yes
RIP (layer 2)	Yes	Yes	No

All the units in a cluster must be in the same mode: either single or multiple context mode. In multiple context mode, the master-slave synchronization includes all the contexts and the RIB table entries of all the contexts in the synchronization message.

In clustering, if you have configured a layer 3 interface, you must also configure the router-id pool setting.

For more information about dynamic routing and clustering, see Chapter 1, “Configuring a Cluster of ASAs.”

Dynamic Routing in Multiple Context Mode

In multiple context mode, each context maintains a separate routing table and routing protocol databases. This enables you to configure OSPFv2 and EIGRP independently in each context. You can configure EIGRP in some contexts and OSPFv2 in the same or different contexts. In mixed context mode, you can enable any of the dynamic routing protocols in contexts that are in routed mode. RIP and OSPFv3 are not supported in multiple context mode.

The following table lists the attributes for EIGRP, OSPFv2, route maps used for distributing routes into OSPFv2 and EIGRP processes, and prefix lists used in OSPFv2 to filter the routing updates entering or leaving an area when they are used in multiple context mode:

EIGRP	OSPFv2	Route Maps and Prefix Lists
One instance is supported per context.	Two instances are supported per context.	N/A
It is disabled in the system context.		N/A
Two contexts may use the same or different autonomous system numbers.	Two contexts may use the same or different area IDs.	N/A

EIGRP (continued)	OSPFv2 (continued)	Route Maps and Prefix Lists (continued)
Shared interfaces in two contexts may have multiple EIGRP instances running on them.	Shared interfaces in two contexts may have multiple OSPF instances running on them.	N/A
The interaction of EIGRP instances across shared interfaces is supported.	The interaction of OSPFv2 instances across shared interfaces is supported.	N/A
All CLIs that are available in single mode are also available in multiple context mode.		
Each CLI has an effect only in the context in which it is used.		

Route Resource Management

A resource class called *routes* has been introduced, which specifies the maximum number of routing table entries that can exist in a context. This resolves the problem of one context affecting the available routing table entries in another context and also allows you greater control over the maximum route entries per context.

Because there is no definitive system limit, you can only specify an absolute value for this resource limit; you may not use a percentage limit. Also, there are no minimum and maximum limits per context, so the default class does not change. If you add a new route for any of the static or dynamic routing protocols (connected, static, OSPF, EIGRP, and RIP) in a context and the resource limit for that context is exhausted, then the route addition fails and a syslog message is generated.

Disabling Proxy ARPs

When a host sends IP traffic to another device on the same Ethernet network, the host needs to know the MAC address of the device. ARP is a Layer 2 protocol that resolves an IP address to a MAC address. A host sends an ARP request asking “Who is this IP address?” The device owning the IP address replies, “I own that IP address; here is my MAC address.”

Proxy ARP is used when a device responds to an ARP request with its own MAC address, even though the device does not own the IP address. The ASA uses proxy ARP when you configure NAT and specify a mapped address that is on the same network as the ASA interface. The only way traffic can reach the hosts is if the ASA uses proxy ARP to claim that the MAC address is assigned to destination mapped addresses.

Under rare circumstances, you might want to disable proxy ARP for NAT addresses.

If you have a VPN client address pool that overlaps with an existing network, the ASA by default sends proxy ARPs on all interfaces. If you have another interface that is on the same Layer 2 domain, it will see the ARP requests and will answer with the MAC address of its interface. The result of this is that the return traffic of the VPN clients towards the internal hosts will go to the wrong interface and will get dropped. In this case, you need to disable proxy ARPs for the interface on which you do not want proxy ARPs.

To disable proxy ARPs, enter the following command:

Command	Purpose
sysopt noproxyarp *interface* **Example:** `hostname(config)# sysopt noproxyarp exampleinterface`	Disables proxy ARPs.