Configure Route Selection for Routers

Contents

Introduction

Prerequisites

- Requirements
- Components Used
- Conventions

Background Information

Processes Involved

Build the Routing Table

- Backup Routes
- Adjust the Administrative Distance
- How Metrics Determine the Route Selection Process
- Prefix Lengths

Make Forwarding Decisions

- IP Classless

Summary

Related Information

Introduction

This document describes how routers work, are configured, and how to select a route for them.

Prerequisites

Requirements

There are no specific prerequisites for this document.

Components Used

This document is not restricted to specific software and hardware versions.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Conventions

For more information on document conventions, see the Cisco Technical Tips Conventions.

Background Information

One aspect of Cisco routers is how the router chooses the best route among those presented by protocols, manual configuration, and various other means. Route selection requires some knowledge about the way Cisco routers work.

Processes Involved

There are three processes involved to build and maintain the routing table in a Cisco router:

- Various routing processes, which actually run a network (or routing) protocol, such as Enhanced Interior Gateway Routing Protocol (EIGRP), Border Gateway Protocol (BGP), Intermediate System-to-Intermediate System (IS-IS), and Open Shortest Path First (OSPF).
- The routing table itself, which accepts information from the routing processes and also replies to requests for information from the forwarding process.
- The forwarding process, which requests information from the routing table to make a packet forwarding decision.

You need to examine the interaction between the routing protocols and the routing table to understand how the routing table is built.

Build the Routing Table

The main considerations when you build the routing table are:

- **Administrative distance**- This is the measure of trustworthiness of the source of the route. If a router learns about a destination from more than one routing protocol, the administrative distance is compared and the preference is given to the routes with lower administrative distance.
- **Metrics**- This is a measure used by the routing protocol to calculate the best path to a given destination, if it learns multiple paths to the same destination. Each routing protocol uses a different metric.
- **Prefix length**

As each routing process receives updates and other information, it chooses the best path to any given destination and attempts to install this path into the routing table.

For instance, if EIGRP learns of a path toward 10.1.1.0/24, and decides this particular path is the best EIGRP path to this destination, it tries to install the path it has learned into the routing table.

The router decides whether or not to install the routes presented by the routing processes based on the administrative distance of the route in question. If this path has the lowest administrative distance to this destination (when compared to the other routes in the table), it is installed in the routing table. If this route is not the route with the best administrative distance, then the route is rejected.

For example, assume a router runs four routing processes: EIGRP, OSPF, RIP, and IGRP. Now, all four of these processes have learned of various routes to the 192.168.24.0/24 network, and each has chosen its best path to that network through its internal metrics and processes.

Each of these four processes attempts to install their route toward 192.168.24.0/24 into the routing table. The routing processes are each assigned an administrative distance, which is used to decide which route to install.

Default Administrative Distances

Connected	0
Static	1
eBGP	20
EIGRP (internal)	**90**
IGRP	100
OSPF	110
IS-IS	115
RIP	120
EIGRP (external)	170
iBGP	200
EIGRP summary route	5

Because the internal EIGRP route has the best administrative distance (the smaller the administrative distance, the higher the preference), it is installed in the routing table.

Backup Routes

- What do the other protocols, RIP, IGRP, and OSPF, do with the routes that were not installed?
- What if the most preferred route, learned from EIGRP, fails?

Cisco IOS® software uses two approaches to solve this problem.

The first is to have each routing process attempt to install its best routes periodically. If the most preferred route fails, the next best route (determined by the administrative distance) succeeds on the next attempt.

The other solution is for the routing protocol that failed to install its route in the table to hang on to the route and tell the routing table process to report if the best path fails.

For protocols that do not have their own routing information tables, such as IGRP, the first method is used.

Every time IGRP receives an update about a route, it attempts to install the updated information in the routing table.

If there is already a route to this same destination in the routing table, the installation attempt fails.

For protocols that have their own database of routing information, such as EIGRP, IS-IS, OSPF, BGP, and RIP, a backup route is registered when the initial attempt to install the route fails.

If the route installed in the routing table fails for some reason, the routing table maintenance process calls each routing protocol process that has registered a backup route, and asks them to reinstall the route in the routing table. If there are multiple protocols with registered backup routes the preferred route is chosen based on administrative distance.

Adjust the Administrative Distance

The default administrative distance is not always right for your network; you can adjust it so that RIP routes are preferred over IGRP routes. But, first, look at the implications if you change the administrative distance.

It is very dangerous to change the administrative distance on routing protocols. It can lead to routing loops and other oddities in your network. Therefore, always change the administrative distance with caution. Ensure that you plan the change and know the consequences before you do this.

For entire protocols, it is easy to change the distance. Just use the **distance** command in the routing process sub-configuration mode. You can also change the distance for routes learned from one source only in some

protocols, and you can change the distance on just some routes. For more information, refer to Adjust Administrative Distance for Route Selection in Cisco IOS Routers Configuration Example.

For static routes, to change the distance of each route enter a distance after the **ip route** command:

```
ip route network subnet mask next hop distance
```

You cannot change the administrative distance for all the static routes at the same time.

How Metrics Determine the Route Selection Process

Routes are chosen and built in the routing table based on the administrative distance of the routing protocol. The routes learned from the routing protocol with the lowest administrative distance are installed in the routing table. If there are multiple paths to the same destination from a single routing protocol, then the multiple paths would have the same administrative distance and the best path is selected based on the metrics. Metrics are values associated with specific routes that rank 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. For more information on load balancing see How Does Load Balancing Work?

Prefix Lengths

Look at another scenario to see how the router handles another common situation: varying prefix lengths. Assume, again, that a router runs has four routing processes, and each process has received these routes:

- EIGRP (internal): 192.168.32.0/26
- RIP: 192.168.32.0/24
- OSPF: 192.168.32.0/19

Which of these routes can be installed in the routing table?

Because EIGRP internal routes have the best administrative distance, you can assume the first one can be installed.

However, because each of these routes has a different prefix length (subnet mask), they are considered different destinations, and they can all be installed in the routing table.

The next section provides the information from the routing table to make forwarding decisions.

Make Forwarding Decisions

Look at the three routes that were installed in the routing table and see how they look on the router.

```
<#root>

router#

show ip route

     ....
     D   192.168.32.0/26 [90/25789217] via 10.1.1.1
     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
....
```

- If a packet arrives on a router interface destined for 192.168.32.1, which route would the router choose?
- It depends on the prefix length, or the number of bits set in the subnet mask.

Longer prefixes are always preferred over shorter ones when forwarding a packet.

In this case, a packet destined to 192.168.32.1 is directed toward 10.1.1.1, because 192.168.32.1 falls within the 192.168.32.0/26 network (192.168.32.0 to 192.168.32.63).

It also falls within the other two routes available, but the 192.168.32.0/26 has the longest prefix within the routing table (26 bits verses 24 or 19 bits).

Likewise, if a packet destined for 192.168.32.100 arrives on one of the router interfaces, it is forwarded to 10.1.1.2,

because 192.168.32.100 does not fall within 192.168.32.0/26 (192.168.32.0 through 192.168.32.63), but it does fall within the 192.168.32.0/24 destination (192.168.32.0 through 192.168.32.255).

Again, it also falls into the range covered by 192.168.32.0/19, but 192.168.32.0/24 has a longer prefix length.

IP Classless

Where the **ip classless** configuration command falls within the routing and forwarding processes is often confusing.

In reality, IP classless only affects the operation of the forwarding processes in Cisco IOS; it does not affect the way the routing table is built.

If IP classless is not configured (with the **no ip classless** command), the router cannot forward packets to supernets.

As an example, again place three routes in the routing table and route packets through the router.



Note: If the supernet or default route is learned via IS-IS or OSPF, the **no ip classless** configuration command is ignored. In this case, packet switching behavior works as though **ip classless** were configured

---

<#root>

router#

show ip route

```
....
     172.30.0.0/16 is variably  subnetted, 2 subnets, 2 masks
D        172.30.32.0/20 [90/4879540] via  10.1.1.2
D       172.30.32.0/24  [90/25789217] via 10.1.1.1
S*   0.0.0.0/0 [1/0] via 10.1.1.3
```

The 172.30.32.0/24 network includes the addresses 172.30.32.0 through 172.30.32.255, and the 172.30.32.0/20 network includes the addresses 172.30.32.0 through 172.30.47.255, therefore, you can then try switching three packets through this routing table and see what the results are.

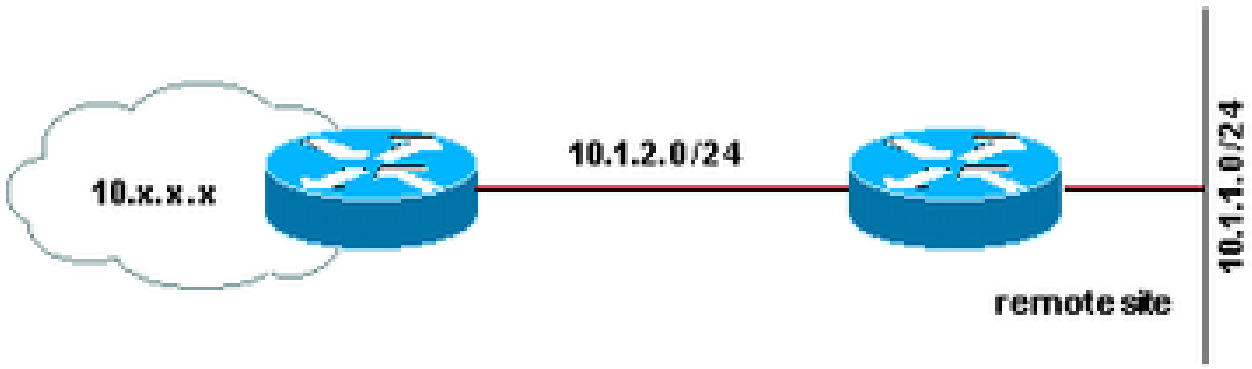
- A packet destined to 172.30.32.1 is forwarded to 10.1.1.1, because this is the longest prefix match.
- A packet destined to 172.30.33.1 is forwarded to 10.1.1.2, because this is the longest prefix match.
- A packet destined to 192.168.10.1 is forwarded to 10.1.1.3; because this network does not exist in the routing table, this packet is forwarded to the default route.
- A packet destined to 172.30.254.1 is dropped.

The answer out of these four is the last packet, which is dropped. It is dropped because its destination, 172.30.254.1, is within a known major network, 172.30.0.0/16, but the router does not know about this particular subnet within that major network.

This is the essence of classful routing: If one part of a major network is known, but the subnet toward which the packet is destined within that major network is unknown, the packet is dropped.

The most confusing aspect of this rule is that the router only uses the default route if the destination major network does not exist in the routing table at all.

This can cause problems in a network where a remote site, with one connection back to the rest of the network, runs no routing protocols, as illustrated.



Runs No Routing Protocol

The remote site router is configured like this:

```
interface Serial 0
     ip address 10.1.2.2 255.255.255.0
   !
   interface Ethernet 0
     ip address 10.1.1.1 255.255.255.0
   !
   ip route 0.0.0.0 0.0.0.0 10.1.2.1
   !
   no ip classless
```

With this configuration, the hosts at the remote site can reach destinations on the Internet (through the 10.x.x.x cloud), but not destinations within the 10.x.x.x cloud, which is the corporate network.

Because the remote router knows about some part of the 10.0.0.0/8 network, the two directly connected subnets, and no other subnet of 10.x.x.x, it assumes these other subnets do not exist and drops any packets destined for them.

Traffic destined to the Internet, however, does not ever have a destination in the 10.x.x.x range of addresses, and is therefore correctly routed through the default route.

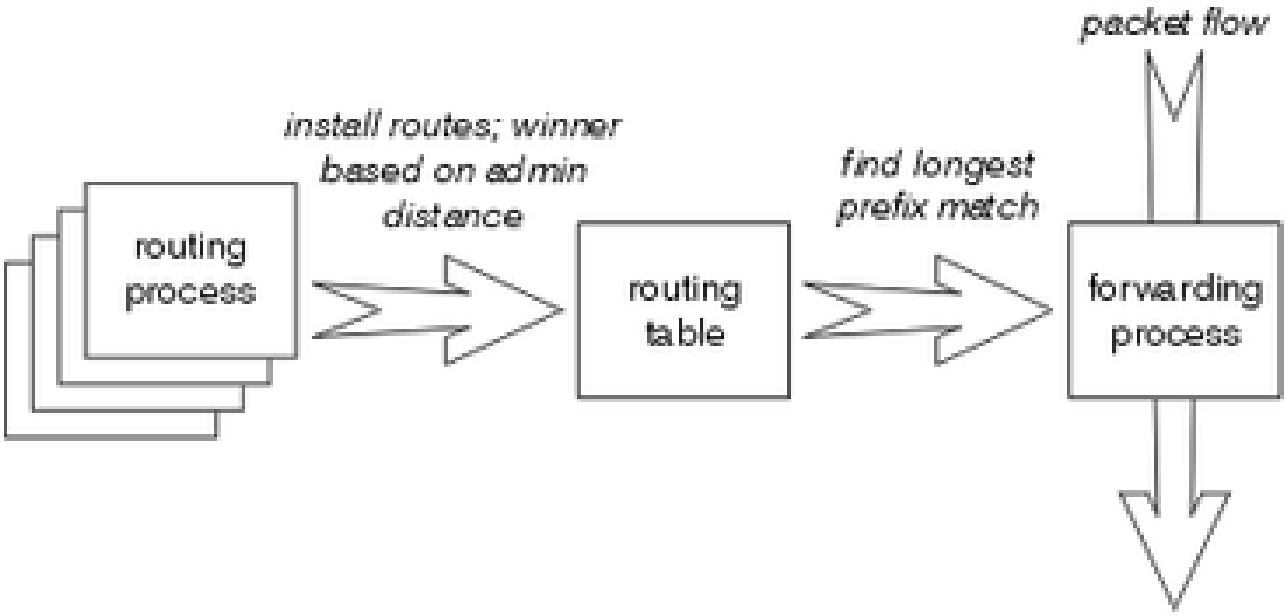
If you configure **ip classless** on the remote router, this problem resolves because it allows the router to ignore the classful boundaries of the networks in its routing table and simply route to the longest prefix match it can find.

Summary

In summary, to make a forwarding decision consists of three sets of processes:

- the routing protocols,
- the routing table, and
- the actual process which makes a forwarding decision and switches packets.

These three sets of processes are illustrated, along with their relationship, in the next image:



Three Sets of Routing Processes

The longest prefix match always wins among the routes installed in the routing table, while the routing protocol with the lowest administrative distance always wins when the routes are installed into the routing table.

Related Information

- **How Does Load Balancing Work?**
- **What is Administrative Distance?**

- **IP Routing Page**
- **IP Routed Protocols Page**
- **Cisco Technical Support & Downloads**