

Understanding the IPv6 Routing Table

Joseph Davies

A routing table, which is present on all nodes running the IPv6 protocol for both the Windows Server 2003 family and Windows XP, stores information about IPv6 network prefixes and how they are reached either directly or indirectly. Before the IPv6 routing table is checked, the destination cache is checked for an entry that matches the destination address in the IPv6 packet that is being forwarded. The destination cache is a table that is maintained by IPv6 nodes. It maps a destination address to a next-hop address and stores the path MTU. If an entry for the destination address is not in the destination cache, the routing table is used to determine:

1. The interface to be used for forwarding (the next-hop interface)
The interface identifies the physical or logical interface that is used to forward the packet to either its destination or to the next router.
2. The next-hop address
For a direct delivery (the destination is on a local link), the next-hop address is the destination address in the packet. For an indirect delivery (the destination is not on a local link), the next-hop address is the address of a router.

After the next-hop interface and address are determined, the destination cache is updated. Subsequent packets that are forwarded to the destination use the destination cache entry instead of checking the routing table.

Contents of an IPv6 Routing Table

The following are the fields of a typical IPv6 routing table entry:

- **Destination Prefix:** The destination prefix is an IPv6 address prefix that can have a prefix length from 0 through 128. In the Windows XP and Windows Server 2003 family IPv6 routing table, this column is named Prefix.
- **Next-Hop Address:** The address to which the packet is forwarded. In the Windows XP and Windows Server 2003 family IPv6 routing table, this column is named Gateway/Interface Name.
- **Interface:** The network interface that is used to forward the packet. All the addresses defined by the Destination Prefix are reachable over the interface. In the Windows XP and Windows Server 2003 family IPv6 routing table, this column is named Idx (an abbreviation of Index) and stores the interface index of the forwarding interface.
- **Metric:** A number that is used to specify the cost of the route, so that the best route (potentially among multiple routes to the same destination) can be selected.

IPv6 routing table entries can be used to store the following types of routes:

- **Directly-attached Network Routes:** These routes are network prefixes for subnets that are directly attached and typically have a prefix length of 64.
- **Remote Network Routes:** These routes are network prefixes for subnets that are not directly attached but are available through other routers. Remote network routes can be subnet prefixes (typically with a prefix length of 64) or a prefix for an address space (typically with a prefix length less than 64).

Understanding the IPv6 Routing Table

Joseph Davies

- **Host Routes:** A host route is a route to a specific IPv6 address. Host routes allow routing to occur on a per-IPv6 address basis. For host routes, the route prefix is a specific IPv6 address with a prefix length of 128. In contrast, both types of network routes have prefixes that have a prefix length less than 128.
- **Default Route:** The default route is used when a more specific network or host route is not found. The default route prefix is `::/0`.

Route Determination Process

To determine which routing table entry is used to send or forward a packet, IPv6 uses the following process:

1. For a sending host, if the source address is specified by the sending application, the only routes that are checked are those that apply to the interface assigned the source address. For a sending host, if the source address is not specified by the sending application, or for a forwarding router, all the routes are checked. For each routing table entry that is checked, IPv6 compares the bits in the network prefix to those in the destination address for the number of bits specified in the prefix length of the route. For the number of bits in the prefix length, if all the bits in the network prefix match all the bits in the destination IPv6 address, the route is a match for the destination.
2. The list of matching routes is compiled. The route that has the largest prefix length (that is, the route that matched the most high-order bits with the destination address) is chosen. The longest matching route is the most specific route to the destination. If multiple entries with the longest match are found (for example, multiple routes to the same network prefix), the router uses the lowest metric to select the best route. If multiple entries exist that are the longest match and the lowest metric, IPv6 chooses which routing table entry to use.

For any given destination, the above procedure results in finding matching routes in the following order:

1. A host route that matches the entire destination address
2. A network route with the longest prefix length that matches the destination
3. The default route (the network prefix `::/0`)

The result of the route determination process is the selection of a single route in the routing table. The selected route yields a next-hop interface and address. The next-hop interface is the interface that is specified in the matching route. For remote traffic, the next-hop address is the address stored in the Next-Hop Address field (the address of a neighboring router). For traffic to neighbors on a directly-attached link, the next-hop address is the destination address of the packet. (In this case, an address is not stored in the Next-Hop Address field.)

If the route determination process on the sending host fails to locate a matching route, IPv6 treats the destination as locally reachable. If the route determination process on a router fails to locate a matching route, IPv6 sends an ICMPv6 Destination Unreachable-No Route to Destination message to the sending host and discards the packet.

Understanding the IPv6 Routing Table

Joseph Davies

Example IPv6 Routing Tables for Windows XP and the Windows Server 2003 Family

To view the IPv6 routing table on a computer that is running a member of the Windows Server 2003 family or Windows XP, type:

```
netsh interface ipv6 show routes
```

at a command prompt

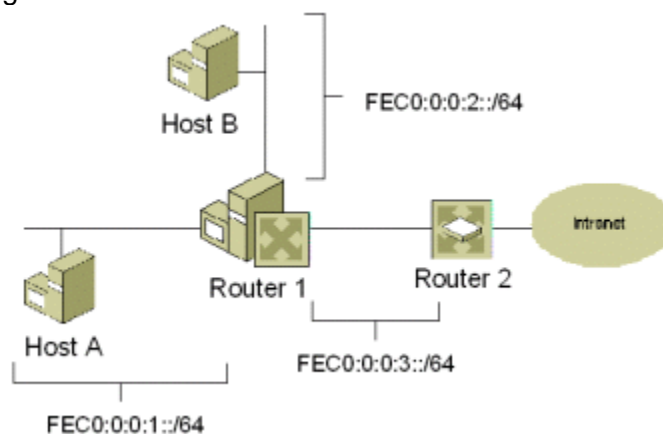
To view the entire routing table, which includes host routes and multicast routes and additional information for each route, use the:

```
netsh interface ipv6 show routes level=verbose
```

command.

The IPv6 routing table is built automatically, based on your IPv6 configuration. You can also add routes using the netsh interface ipv6 add route command. A route for the link-local prefix (FE80::/64) is never displayed in the IPv6 routing table.

To show IPv6 routing behavior for both a sending host and a router, we will examine the example network in the following figure.



Router 1 is computer using the IPv6 protocol for Windows XP and the Windows Server 2003 family that is configured as a forwarding and advertising router. Router 1 has three network adapters, is acting as a default router for two subnets configured with site-local address prefixes (FEC0:0:0:1::/64 and FEC0:0:0:2::/64), and has a default route that points to a default router on a third subnet (FEC0:0:0:3::/64). For more information about how to configure a computer running Windows XP or a member of the Windows Server 2003 family as an IPv6 router, see Manual Configuration for IPv6 (the September 2002 Cable Guy article).

The following is the display of the netsh interface ipv6 show routes command for Router 1:

```
Publish Type Met Prefix Idx Gateway/Interface Name -----
----- yes Autoconf 8 fec0:0:0:1::/64 3 Local Area
Connection yes Autoconf 8 fec0:0:0:2::/64 4 Local Area Connection 2 yes Autoconf 8
fec0:0:0:3::/64 5 Local Area Connection 3 yes Manual 256 ::/0 5
fe80::210:ffff:fed6:58c0
```

Understanding the IPv6 Routing Table

Joseph Davies

Each entry in the Windows XP and Windows Server 2003 family IPv6 routing table has the following fields:

- **Publish** specifies whether the route is published (advertised in a Routing Advertisement message).
- **Type** specifies the type of route.
- **Routes** configured by user applications have a route type of Manual. Routes configured by the IPv6 protocol have a route type of Autoconf.
- **Met** specifies the metric used to select between multiple routes with the same prefix. The lowest metric is the best matching route.
- **Prefix** specifies the address prefix.
- **Idx** specifies the interface index, which indicates the interface over which packets that match the address prefix are reachable.
- **Interface** indexes can be viewed from the display of the netsh interface ipv6 show interface command.
- **Gateway/Interface Name** specifies either a next-hop IPv6 address or an interface name.
- For remote network routes, a next-hop IPv6 address is listed. For directly-attached network routes, the name of the interface from which the address prefix is directly reachable is listed.

The first, second, and third routes are for the 64-bit site-local address prefixes of locally-attached subnets. An Ethernet network adapter named Local Area Connection (interface index 3) is connected to the subnet FEC0:0:0:1::/64. A second Ethernet network adapter named Local Area Connection 2 (interface index 4) is connected to the subnet FEC0:0:0:2::/64. The third Ethernet network adapter named Local Area Connection 3 (interface index 5) is connected to the subnet FEC0:0:0:3::/64.

The fourth route is the default route (prefix of ::/0), which matches all destinations. If the default route is the longest matching route for the destination, the packet is sent to Router 2's link-local address FE80::210:FFFF:FED6:58C0 by using the Ethernet network adapter named Local Area Connection 3 (interface index 5).

Because Router 1 is configured as an advertising router for the three attached subnets and the default route, it sends the following router advertisements:

- The router advertisement for the FEC0:0:0:1::/64 subnet contains an indication that Router 1 is a default router, a Prefix Information option for FEC0:0:0:1::/64, a Route Information option for FEC0:0:0:2::/64, and a Route Information option for FEC0:0:0:3::/64.
- The router advertisement for the FEC0:0:0:2::/64 subnet contains an indication that Router 1 is a default router, a Prefix Information option for FEC0:0:0:2::/64, a Route Information option for FEC0:0:0:1::/64, and a Route Information option for FEC0:0:0:3::/64.

Understanding the IPv6 Routing Table

Joseph Davies

- The router advertisement for the FEC0:0:0:3::/64 subnet contains an indication that Router 1 is a default router, a Prefix Information option for FEC0:0:0:3::/64, a Route Information option for FEC0:0:0:1::/64, and a Route Information option for FEC0:0:0:2::/64.

The Prefix Information option, specified in RFC 2461, is used to provide an address prefix from which autoconfigured addresses are derived. The Route Information option, specified in the Internet draft titled 'Default Router Preferences and More-Specific Routes, is used to populate the routing table of hosts with subnet prefixes for the subnets attached to the advertising router. This provides reachability between nodes on the subnets of the advertising router even when it is not acting as a default router.

The following is the display of the `netsh interface ipv6 show routes` command for Host A:

```
Publish Type Met Prefix Idx Gateway/Interface Name -----
----- yes Autoconf 8 fec0:0:0:1::/64 3 Local Area
Connection yes Autoconf 8 fec0:0:0:2::/64 3 fe80::210:ffff:fea4:211c
yes Autoconf 8 fec0:0:0:3::/64 3 fe80::210:ffff:fea4:211c
yes Autoconf 256 ::/0 3
fe80::210:ffff:fea4:211c
```

The first route is for the locally attached subnet. The second and third routes are for the other two subnets attached to Router 1. These routes are automatically added upon receipt of the router advertisement sent by Router 1 on the FEC0:0:0:1::/64 subnet. The fourth route is for the default route. FE80::210:FFFF:FEA4:211C is the link-local address of Router 1's interface on the FEC0:0:0:1::/64 subnet.

The next-hop interface for a route in the routing table is specified by the interface index in the `Idx` column. When determining the next-hop address from a route in the routing table:

- If the Gateway/Interface Name column of the route table entry specifies an interface name, then the destination is a neighbor and the next-hop address is set to the destination address of the IPv6 packet.
- If the Gateway/Interface Name column of the route table entry specifies an address (that is, the address of a neighboring router), then the destination is remote and the next-hop address is set to the address in the Gateway/Interface Name column.

For example, when Host A sends traffic to FEC0::1:2AA:FF:FE90:4D3C, the longest matching route in its routing table is the route for the directly-attached network FEC0:0:0:1::/64. The next-hop address is set to the destination address of FEC0::1:2AA:FF:FE90:4D3C and the interface is interface index 3, which corresponds to the Ethernet network adapter named Local Area Connection.

When Host A sends traffic to 3FFE:FFFF:2A90:17:2AA:FF:FE03:21A6, the longest matching route in its routing table is the default route (::/0). The next-hop address is set to the Router 1's address of FE80::210:FFFF:FEA4:211C and the interface is interface index 3, which corresponds to the Ethernet network adapter named Local Area Connection.

When Router 1 forwards the packet destined for 3FFE:FFFF:2A90:17:2AA:FF:FE03:21A6, the longest matching route in its routing table is the default route (::/0). The next-hop address is set to Router 2's link-local address of FE80::210:FFFF:FED6:58C0 on the FEC0:0:0:3::/64 subnet and the interface is interface index 5, which corresponds to the Ethernet network adapter named Local Area Connection 3.