redundant topology. Therefore, not only are the remote routers required to respond to queries from the regional offices, but they also continue the search by reflecting the queries back toward the other regional office's router. This significantly complicates the convergence process on the network.

Figure 3-30 *Effect of the EIGRP Update and Query Process*

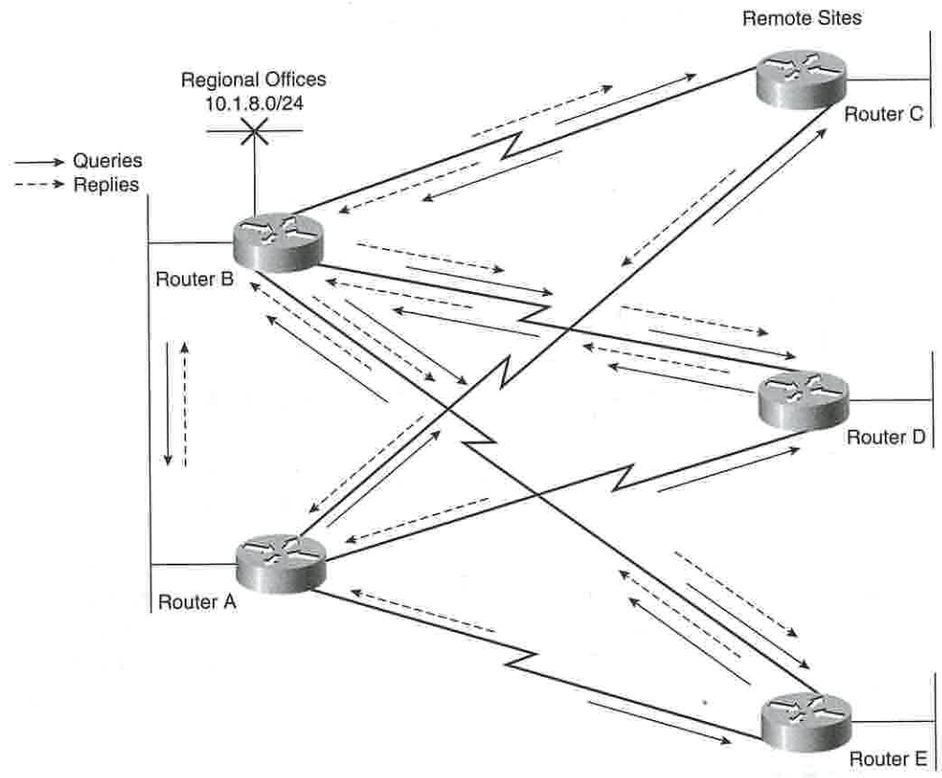

In this sample network with only two regional and three remote routers, the problem might not be very significant. In a network with hundreds of remote offices, the problem can be severe.

Examine the query process for the 10.1.8.0/24 subnet. Router B advertises 10.1.8.0/24 to all other routers. The best path for Router A to reach 10.1.8.0/24 is over the Ethernet link to Router B. The remote routers (C, D, and E) use the serial link to B as their preferred path to reach 10.1.8.0/24 but still learn about an alternative path through Router A. For this example, assume that the EIGRP metric for Ethernet is 1000 and the metric for a serial link is 100,000.

Table 3-5 shows the content of the IP EIGRP topology table on Routers C, D, and E for network 10.1.8.0/24. Table 3-6 shows the content of the IP EIGRP topology table on Router A for network 10.1.8.0/24.

Table 3-5 *IP EIGRP Topology Table for 10.1.8.0/24 on Routers C, D, and E in Figure 3-30*

Neighbor	FD	AD
Router A	102,000	2000
Router B	101,000	1000

Table 3-6 *IP EIGRP Topology Table for 10.1.8.0/24 on Router A in Figure 3-30*

Neighbor	FD	AD
Router B	2000	1000
Router C	201,000	101,000
Router D	201,000	101,000
Router E	201,000	101,000

Note that Routers C, D, and E determine that for network 10.1.8.0/24, Router B is the successor and Router A is an FS (because the AD is 2000 through Router A, which is less than the FD through Router B). Also, note that Router A does not have an FS, because all paths through the remote routers have an AD larger than the FD through Router B.

When Router B loses the path to network 10.1.8.0/24, it queries all four of its neighbors. When the remote sites receive this query, they automatically install the path through Router A in their routing tables and respond to Router B with their supposedly good path through Router A. They also remove the bad path through Router B from their topology tables.

Router B now has responses to three of its four queries, but it must wait until Router A responds as well.

When Router A receives the query from Router B for network 10.1.8.0/24, Router A creates a query and sends it to Routers C, D, and E, because Router A does not have an FS but knows that a path exists through each remote site to reach 10.1.8.0/24.

Routers C, D, and E receive the query from Router A; they now know that their path through Router A is not good, so they check their topology tables for alternative paths. However, none of these routers currently has another path, because Router B has just informed them that it does not have a path to this network. Because the remote routers do not have an answer to the query from Router A, Routers C, D, and E create a query and send it to all neighbors except the neighbor (interface) that these routers received the original query from. In this case, the remote routers send the query only to Router B.

Router B learns from these queries that none of the remote routers has a path to network 10.1.8.0/24, but it cannot respond that it does not know of a path, because Router B is waiting for Router A to