Chapter 6
General Routed Design Theory

Routed network design is at the heart of the CCDE. Understanding the concepts and principles of Layer 3 design, focused on routing, is critical to achieving the CCDE certification. This chapter discusses the general principles of routed network design. Included are sections on route aggregation, fate sharing, redundancy and resiliency, convergence tuning, configuration complexity, and multicast design principles. Then Chapter 7, “Topology Design Theory,” discusses specific topologies and how they interact with the routing protocols.

Route Aggregation

Route aggregation serves two specific purposes in a network:

- Breaking the network into multiple failure domains
- Reducing the amount of information the routing protocol must deal with when converging

Figure 6-1 will be used to illustrate both of these principles.

FIGURE 6-1
Aggregation Principles

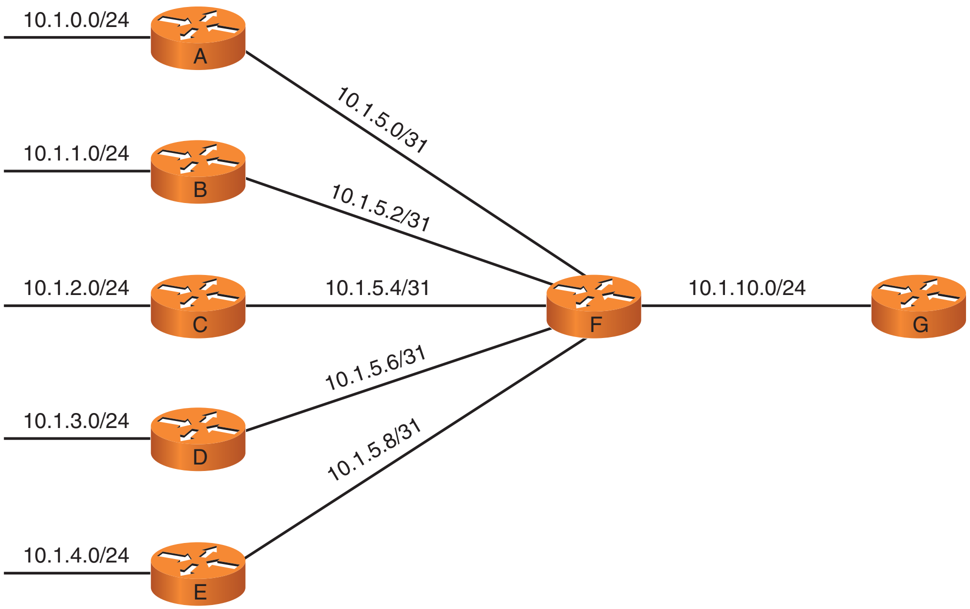

Based on the IP addressing within this network, it is possible to aggregate at Router F toward Router G to 10.1.0.0/22, which would provide reachability to 10.1.0.0 through 10.1.7.255. Aggregation at Router F breaks up the failure domain by blocking notifications or updates about individual link failures between Routers F and A, B, C, D, or E from being transmitted to Router G. For instance, if the link between Routers A and F fails, the aggregate Router F is sending to Router G does not change.

Aggregation at Router F can also improve network convergence by decreasing the amount of information routers in the network must process. If a new router is attached to Router G, it needs to send just 2 routes, rather than 11. There is a large multiplier effect as a network grows.

Most routers build a *discard route* when advertising an aggregate. The discard route causes the traffic falling within the aggregate route, but not within one of the more specific routes in the aggregating routers routing table, to be discarded. For instance, if Router G transmits a packet toward 10.1.6.1, the only route in Router F's routing table that will match this destination is the discard route built off the aggregate Router F is advertising toward Router G.

Network engineers should always be aware that the longest prefix within the local routing table that contains the destination address will be used for forwarding. For instance, if a packet is received with a destination address of 10.1.1.1, and there are two routes in the local routing table (10.1.1.0/24 and 10.1.1.0/25), which both contain, or could provide a valid route, to the destination, the route with the longer prefix length will always be preferred over the route with the shorter prefix length. This might appear to be a simple rule, but this simple rule will often cause unexpected misrouting of traffic.

The longest prefix rule can also be used as an advantage. For instance, in the network in Figure 6-2, you could use the longest-match rule to optimally route traffic to the correct destination, while allowing for a backup path that is always in the local routing table, for faster convergence.

FIGURE 6-2
Longer Prefixes and Optimal Routing

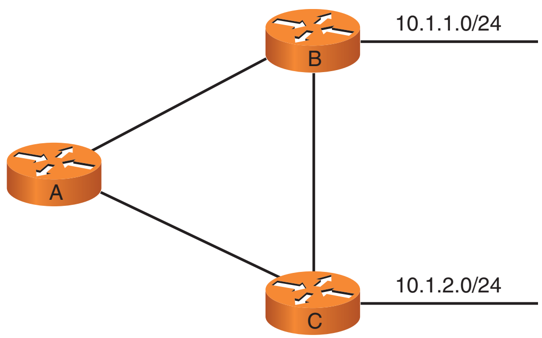

If the following advertisements are configured in the network

- Router B advertises 10.1.0.0/16 and 10.1.1.0/24 toward Router A.
- Router C advertises 10.1.0.0/16 and 10.1.2.0/24 toward Router A.

For packets destined to 10.1.1.1, Router A will choose the path through Router B, whereas for packets destined to 10.1.2.1, Router A will choose the path through Router C. If the Router A to Router B link fails, however, Router A will choose the path through Router C to reach 10.1.1.1, because that is the only path available to this destination.

In many situations, the discard route built through the configuration of an aggregate route will create a black hole in the event of some specific link failures, as shown in Figure 6-3.

FIGURE 6-3
Aggregate Black Holes

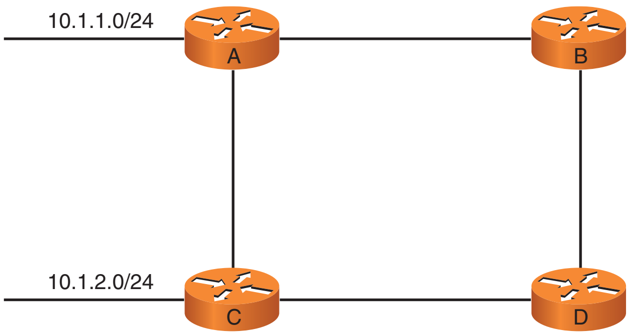

If the following routes are being advertised in this network

- Router A is advertising 10.1.0.0/16 to Routers B and C.
- Router C is advertising 10.1.0.0/16 to routers A and D.

If the link from Router A to Router B fails, Router D will forward traffic destined to 10.1.1.1 to Router C. Because Router C has an aggregate route to 10.1.0.0/16, and no more specific route to the destination, it will forward the traffic to the discard route created when the aggregate was configured. Although a physical path to the destination is available, the traffic is discarded because of the way the aggregation is configured. In general, you should always have a link between any set of routers that are configured with the same aggregates on which aggregation is not configured.

Network engineers should also be careful of the following when configuring or using aggregated routing information:

- Aggregate routes can bring traffic further into the network than you might want, which could result in security issues. For instance, packets resulting from an attack that are directed to a destination that does not exist in your network could pass deep into your network before they are discarded at the router with a discard route because of an aggregate.
- Most routing protocols take the metric for an aggregate route from the component routes, or the routes that are blocked by the aggregate route. Routing protocols take the metric from the component route with the lowest metric or the highest metric, and advertise the aggregate route with the same metric. If the route from which the aggregate's metric is taken fails, or is withdrawn, the aggregate's metric will change, too, communicating the change to the routers beyond the aggregate route. This defeats the purpose of the aggregate route. Network engineers should consider this when designing route aggregation, using techniques to keep the aggregate's metric from changing where possible.

Fate Sharing

Network engineers generally consider fate sharing only when there are multiple logical signals on a single physical wire, but the concept of fate sharing, and its application in network design, is much broader than this specific case. Any time a network is virtualized, fate sharing will be the result. For instance:

- Multiple wavelengths over a single fiber, using dense wavelength-division multiplexing (DWDM)
- Multiple circuits multiplexed through add/drop multiplexers on a SONET link
- Multiple circuits running over a single Frame Relay circuit
- Multiple VLANs running over a single physical Ethernet cable
- Multiple VPNs running over a single Layer 3 infrastructure