Link-State Routing

Link-state routing differs from distance-vector routing in that each router knows the exact topology of the network. This reduces the number of bad routing decisions that can be made because every router in the process has an identical view of the network. Each router in the network will report on its state, the directly connected links, and the state of each link. The router will then propagate this information to all routers in the network. Each router that receives this information will take a snapshot of the information. This ensures all routers in the process have the same view of the network, allowing each router to make its own routing decisions based upon the same information.

In addition, link-state routing protocols generate routing updates only when there is a change in the network topology. When a link, i.e., a point on a route, changes state, a link-state advertisement (LSA) concerning that link is created by the device that detected the change and propagated to all neighboring devices using a multicast address. Each routing device takes a copy of the LSA, updates its topological database and forwards the LSA to all neighboring devices. An LSA is generated for each link on a router. Each LSA will include an identifier for the link, the state of the link, and a metric for the link. With the use of LSAs, linkstate protocols reduces routing bandwidth usage.

Examples of link-state routing protocols are: Open Shortest Path First (OSPF) and Integrated Intermediate System to Intermediate System (IS-IS). Another protocol, Enhanced Interior Gateway Routing Protocol (EIGRP) is considered a hybrid protocol because it contains traits of both distance-vector and link-state routing protocols. Most link-state routing protocols require a hierarchical design, especially to support proper address summarization. The hierarchical approach, such as creating multiple logical areas for OSPF, reduces the need to flood an LSA to all devices in the routing domain. The use of areas restricts the flooding to the logical boundary of the area rather than to all devices in the OSPF domain. In other words, a change in one area should only cause routing table recalculation in that area, not in the entire domain.

Classful Routing

Classful routing is used in routing packets based upon the class of IP address. IP addresses are divided into five classes: Class A, Class B, Class C, Class D, and Class E. Class A, Class B and Class C are used to private and public network addressing; Class D is used for multicast broadcasting; and Class E is reserved by the Internet Assigned Numbers Authority (IANA) for future use.

Classful routing is a consequence of the fact that routing masks are not advertised in the periodic, routine, routing advertisements generated by distance vector routing protocols. In a classful environment, the receiving device must know the routing mask associated with any advertised subnets or those subnets cannot be advertised to it. There are two ways this information can be gained:

. Share the same routing mask as the advertising device

. If the routing mask does not match, this device must summarize the received route a classful boundary and send the default routing mask in its own advertisements.

Classful routing protocols, such as Routing Information Protocol version 1 (RIPv1) and Interior Gateway Routing Protocol (IGRP), exchange routes to subnetworks within the same network if network administrator configured all of the subntworks in the major network have the same routing mask. When routes are exchanged with foreign networks, subnetwork information from this network cannot be included because the routing mask of the other network is not known. As a result, the subnetwork information from this network must be summarized to a classful boundary using a default routing mask prior to inclusion in the routing update. The creation of a classful summary route at major network boundaries is handled automatically by classful routing protocols. However, summarization at other points within the major network address is not allowed by classful routing protocols.

Classless Routing

One of the most serious limitations in a classful network environment is that the routing mask is not exchanged during the routing update process. This requires the same routing mask be used on all subnetworks. The classless approach advertises the routing mask for each route and therefore a more precise lookup can be performed in the routing table. Classless routing, which is also known as Classless Interdomain Routing (CIDR), is thus not dependent on IP address classes but, instead, allows a variablelength subnet mask

(VLSM), which extends IP addressing beyond the limitations of using fixed-length subnet masks (FLSM),to be sent in the routing update with the route. This allows you to conserve IP addresses, extending the use of IP addresses. Classless routing protocols also addressed the need to summarize to a classful network with a default routing mask at major network boundaries. In the classless environment, the summarization process is manually controlled and can be invoked at any point within the network.

The routing protocols that support classless routing protocols are: Routing Information Protocol version 2 (RIPv2); Enhanced Interior Gateway Routing Protocol (EIGRP); Open Shortest Path First (OSPF); and Integrated Intermediate System to Intermediate System (IS-IS).

RIP and IGRP Convergence

Convergence time is one of the problems associated with distance-vector protocols, such as RIPv1 and IGRP. When a router detects a link failure between itself and a neighbor, it sends a flash update with a poisoned route to it other neighbors. These neighbors in turn create a new flash update and send it to all of its neighbors, and so on. The Router that detected the link failure purges the entry for the failed link and removes all routes associated with that link from the routing table. The router then sends a query to its neighbors for the routs that have been removed. If a neighbor responds with a route, it is immediately installed in the routing table. The router does not go into hold-down because the entry was already purged. However, its neighbors are in hold-down for the failed route, thus ignoring periodic advertisement for that route. As the other routers come out of hold-down, the new route announced by the router that detected the failed link will cause their routing table entries to be updated.

EIGRP Convergence

Enhanced IGRP (EIGRP) convergence differs slightly. If a router detects a link failure between itself and a neighbor, it checks the network topology table for a feasible alternate route. If it does not find a qualifying alternate route, it enters in an active convergence state and sends a Query out all interfaces for other routes to the failed link. If a neighbor replies to the Query with a route to the failed link, the router accepts the new path and metric information, places it in the topology table, and creates an entry for the routing table. It then sends an update about the new route out all interfaces. All neighbors acknowledge the update and send updates of their own back to the sender. These bi-directional updates ensure the routing tables are synchronized and validate the neighbor's awareness of the new topology. Convergence time in this event is the total of detection time, plus Query and Reply times and Update times.

Link-State Convergence

The convergence cycle used in Link-State Routing Protocols, such as OSFP and IS-IS, differs from that of the distance-vector protocols. When a router detects a link failure between itself and a neighbor, it tries to perform a Designated Router (DR) election process on the LAN interface, but fails to reach any neighbors. It then deletes the route from the routing table, builds a link-state advertisement (LSA) for OSFP or a link-state PDU (LSP) for IS-IS, and sends it out all other interfaces. Upon receipt of the LSA, the other neighbors that are up copy the advertisement and forward the LSA packet out all interfaces other than the one upon which it arrived. All routers, including the router that detected the failure, wait five seconds after receiving the LSA and run the shortest path first (SPF) algorithm. There after the router that detected the failure adds the new route to the routing table, while its neighbors update the metric in their routing table. After approximately 30 seconds, the failed router sends an LSA after aging out the topology entry from router that detected the failure. After five seconds, all routers run the SPF algorithm again and update their routing tables to the path to the failed link. Convergence time is the total of detection time, plus LSA flooding time, plus the five seconds wait before the second SPF algorithm is run.