Routing loop

A **routing loop** is a common problem with various types of networks, particularly computer networks. They are formed when an error occurs in the operation of the routing algorithm, and as a result, in a group of nodes, the path to a particular destination forms a loop.

In the simplest version, a routing loop of size two, node A thinks that the path to some destination (call it C) is through its neighbouring node, node B. At the same time, node B thinks that the path to C starts at node A.

Thus, whenever traffic for C arrives at either A or B, it will loop endlessly between A and B, unless some mechanism exists to prevent that behaviour.

How a routing loop can form

For example, in this illustration, node A is transmitting data to node C via node B. If the link between nodes B and C goes down and B has not yet informed node A about the breakage, node A transmits the data to node B assuming that the link A-B-C is operational and of lowest cost. Node B knows of the broken link and tries to reach node C via node A, thus sending the original data back to node A. Furthermore, node A receives the data that it originated back from node B and consults its routing table. Node A's routing table will say that it can reach node C via node B (because it still has not been informed of the break) thus sending its data back to node B creating an infinite loop. This routing loop problem is also called a *two-node loop*.

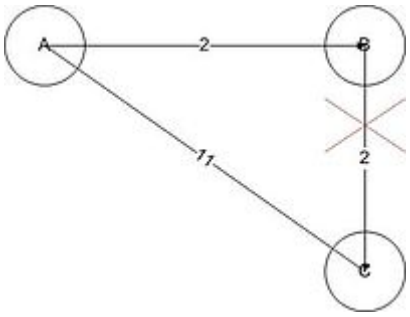

Broken network

How a routing loop can persist

Consider now what happens if both the link from A to C and the link from B to C vanish at the same time (this can happen if node C has crashed). A believes that C is still reachable through B, and B believes that C is reachable through A. In a simple reachability protocol, such as EGP, the routing loop will persist forever.

In a naive distance-vector protocol, such as the routing information protocol, the loop will persist until the metrics for C reach *infinity* (the maximum number of routers that a packet can traverse in RIP is 15. The value 16 is considered infinity and the packet is discarded).

Prevention and mitigations

In a link-state routing protocol, such as OSPF or IS-IS, a routing loop disappears as soon as the new network topology is flooded to all the routers within the routing area. Assuming a sufficiently reliable network, this happens within a few seconds.

Newer distance-vector routing protocols like BGP, EIGRP, DSDV, and Babel have built-in loop prevention: they use algorithms that assure that routing loops can never happen, not even transiently. Older routing protocols like RIP and IGRP do not implement the newest forms of loop prevention and only implement mitigations such as split horizon, route poisoning, and holddown timers.