TECHNICAL SUPPORT NOTE
Understanding OSPF
Featuring ADTRAN OS

Introduction

OSPF or Open Shortest Path First is an Internet Protocol (IP) routing protocol defined in the Request for Standards (RFC) 2328 dated April, 1998. The purpose of this and all other routing protocols is to share the paths to destination networks among peer gateways or routers. The scope of operation for OSPF is designed to be within one autonomous system. A single autonomous system includes all routers normally under a single administrative control as compared with the Internet, which consists of many autonomous systems. OSPF meets the definition of an Interior Gateway Protocol (IGP).

OSPF Overview

Historically, OSPF followed a class of routing protocols known as distance vector or Bellman-Ford protocols. Perhaps the best known of these is the Routing Information Protocol or RIP. RIP shared routes by advertising all the routes that the router knew about with a single or multiple RIP update packets that were sent periodically (about every 30 seconds). Very large networks could take a long time to propagate new information concerning links that had failed.

OSPF was developed by the OSPF working group of the Internet Engineering Task Force (IETF) to utilize the Shortest Path First (also know as the Dijkstra) algorithm. This link-state algorithm normally only forwards updates about changes in member links in the form of Link-State Advertisements (LSAs). All routers learn about the autonomous system (AS) they are a part of by sharing LSAs and develop a database of their neighbors.

The OSPF database can be compared to a map that has all nearby towns and the conditions of the roads leading to them. Based on this map, a traveler (the IP packet) can be quickly sent on a route that results in the shortest amount of time traveled. RIP also does not allow for bandwidth or delay considerations when choosing a route. OSPF allows for these and other parameters in calculating the “best path, or paths.” In the instance of equal cost paths, OSPF permits load balancing across them to optimize data flow.

OSPF Architecture

OSPF is a hierarchical routing protocol. Simply, this means that the protocol was developed to take advantage of preplanning for the large network in a top down fashion.

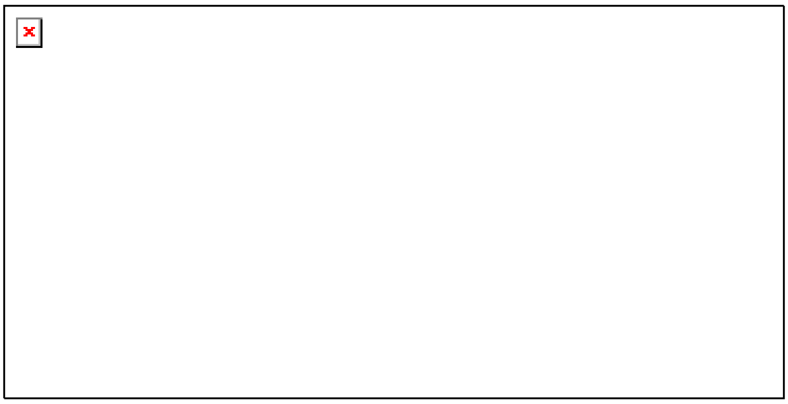

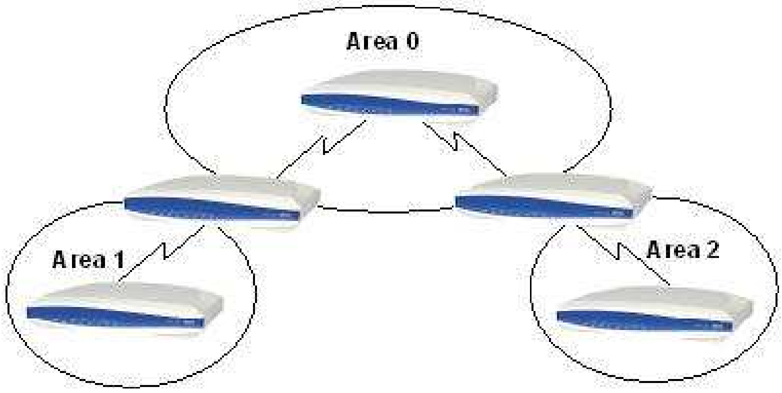

Figure 1 – OSPF Topology

As shown in Figure 1, a large internetwork can be subdivided into smaller “Areas.” These areas branch from a central or “core” area zero. If all routers reside in a single area, then that area normally is area 0.

The benefits of hierarchical topologies include:

- Summarization – or the ability to represent an entire area with a single network address
- Control – Sensitive areas can be restricted from general access
- Smaller routing tables – Only the summary routes get included
- Reduced Link-State exchanges – Updates only reflect an area
- Fewer SPF calculations – a single interface failure in the domain need not force a recalculation, reducing the load on area routers.
- Larger internetworks – OSPF is designed to support large (greater than 15 hop) networks

With the advent of Variable Length Subnet Masks, the Internetwork can be designed in such a way that summarization becomes a natural result. Suppose that the private network scheme 10.0.0.0 is used for the entire AS or domain (collection of all areas). Each area could then be provided with a single sixteen-bit subnet such as 10.1.0.0, 10.2.0.0, etc. Within a single area, the sixteen-bit subnet could then be further divided into twenty-four bit subnets (e.g. 10.1.1.0/24, 10.1.2.0/24, etc.) The routers between areas are known as Area Border Routers or ABRs. At these routers, the summarization and control between areas are performed. For connections to the Internet or to another AS, the Autonomous System Border Router (ASBR) is used to NAT addresses or enforce policies between domains.

OSPF Network Types

OSPF separates three major network types (see Fig. 2): broadcast, point-to-point, and Non-Broadcast Multi-Access or (NBMA). In general, broadcast refers to networks that have the ability for routers to broadcast requests to each other such as those joined on a single Ethernet segment. To reduce the amount of updates on the network, the concept of a Designated Router (DR) and Backup Designated Router (BDR) was developed.

A DR is used to reduce routing traffic. Imagine that 10 backbone routers are joined on a single one hundred-megabit Ethernet segment at the core or area 0. If a single router had to send every update to all other routers, it would open connections to nine other routers. If the other routers needed to do the same thing (at startup, for instance), a large amount of routing traffic would occur. Now consider that all routers can elect a single "spokesrouter" that has the job of notifying all other routers about changes to the network (the DR). That single router could broadcast or flood the update to members of the group (multicast group) as it is told about any change. As a side benefit, all routers have the same routing table reducing the opportunity for errors.

A special "Hello" packet is used to learn about each router. The Hello protocols allows each router to establish its OSPF neighbors and whether they are active (adjacencies). The Hello packet is used to share information about neighbors in the DR election process. The router with the highest priority becomes the DR, next highest the BDR.

OSPF Network Types

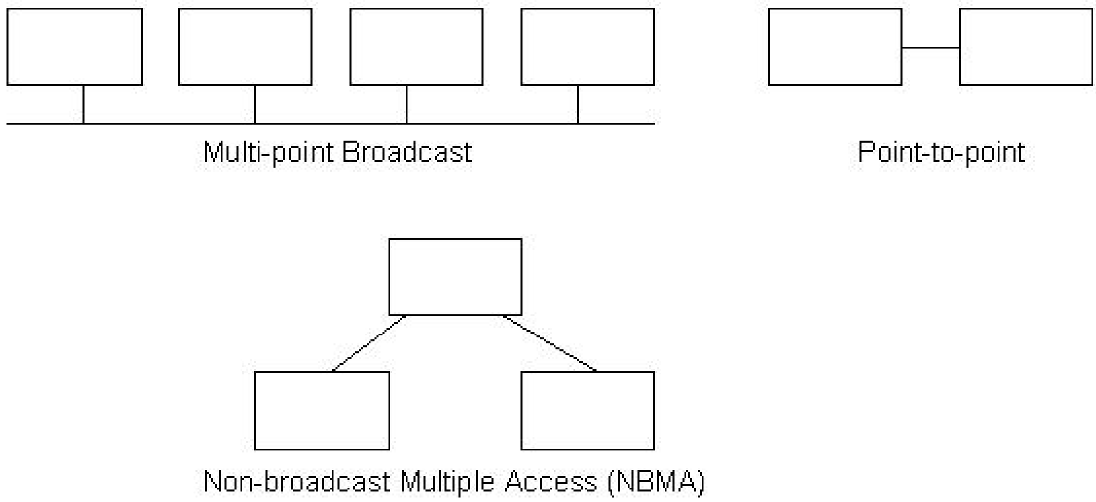

Figure 2 – OSPF Network Types

The Point-to-point network uses the Hello packet to establish an adjacency, but no DR or BDR is necessary. In general, Hello packets are sent by default at 10-second intervals for broadcast networks, and 30-second intervals for point-to-point networks.

Finally, the Non-Broadcast Multi-Access or NMBA network type is used in partially meshed network types such as Frame-relay point-to-multipoint topologies. In general, Frame Relay networks employ a hub and spoke design. In some rare instances, a partially meshed or fully meshed topology could be used, but costs rise as the number of links increase. Fully meshed networks links increase exponentially for every new node added! A fully meshed 20-site network requires 190 links. The NetVanta access routers are specially designed to function in the point-to-point mode in their role as remote routers in a point-to-multipoint network. In this topology, all routers are in the same subnet in a manner similar to broadcast networks. The same multipoint architecture could also use pure point-to-point configurations since many commercial routers can utilize sub-interface or virtual interface technologies in Frame Relay routing networks. In the case where point-to-point links are used, each subinterface pair has their own subnet. Both topologies are supported and allow the network designer flexibility for his design.

Summary

The key benefits of OSPF include:
Multivendor interoperable routing solution
Automatic routing in large (radius > 15 routers) internetworks
Fast convergence
Low bandwidth routing related traffic
Hierarchical logical network designs
Summarization of multiple subnets
Multiple load sharing links
Highly configurable network architecture security
Ease of configuration

If you experience any problems using your ADTRAN product, please contact ADTRAN Technical Support.