Introduction to Routers

TEL-335

Week 6 – Interior Routing with OSPF Part 2

Brief Review: Link State vs. Distance Vector

- **Distance Vector** routing protocols **send their entire routing table to their nearest neighbors at regular intervals**. Those neighbors then filter through the received routing table to learn which direction to send traffic to reach a given network. Distance Vector routing protocols use a distance measurement (metric) to determine the best path.

Brief Review: Link State vs. Distance Vector

- **RIP:** Uses a hop count to determine which route is best. If one router says that a network is directly connected to it, it will have a smaller hop count than a router two links away. In RIP, routing decisions are based entirely on this hop count.
- Other Distance Vector routing protocols, such as IGRP, improve on this by using such factors as bandwidth, current load, dollar cost or reliability as well as hop count to determine the best path. *Advance Course Material*

Brief Review: Link State vs. Distance Vector

- **Link State** routing protocols (like OSPF) **send only the parts of the routing table which have changed to every router within their area**. Routers using Link State routing protocols maintain a picture of the entire network, and are aware of outages several hops away. The router can then use this knowledge to determine the best route for traffic, and can avoid sending data across the network if that network is down.

Brief Review: OSPF multicasting and “Hello”

- In OSPF, two multicast addresses are used. When an OSPF area is started, one router is elected the *Designated Router* (DR), and another is made the *Backup Designated Router* (BDR), there may be several BDRs.
- The Designated Router tells all the other routers about changes in the network by sending out Link State Advertisements (LSAs) on multicast address **224.0.0.5**. When a router notices a change in the network, it sends this information out on **224.0.0.6**, the multicast address reserved for the DR and the BDR.

Brief Review: OSPF multicasting and “Hello”

- When a new router is added to the net, it announces its presence to the world by sending out a *hello packet* on 224.0.0.5. If it doesn't get a response, it becomes the new Designated Router. If a Designated Router responds to the hello message, the new router will use 224.0.0.6 from then on to send out LSAs.

Brief Review: OSPF multicasting and “Hello”

- OSPF routers check the status of other routers on the network by sending a small *hello packet* at regular intervals.
- If a router does not respond to the hello packet, it is assumed dead, and routing updates are sent to every other router by using the multicast address.
- In the case where there are no network changes, OSPF will use very little bandwidth (only sending hello packets).

Brief Review: OSPF multicasting and “Hello”

- As soon as there is an outage, however, OSPF will flood the network as the change is sent to every router (and then every router notifies every other router about the change).
- This system of near silence when possible and flooding when necessary ensures that routing information gets propagated throughout the network as quickly as possible.

OSPF: Area Types

- **Backbone area**
  - The backbone area (also known as *area zero*) forms the core of an OSPF network. All other areas are connected to it, and inter-area routing happens via a router connected to the backbone area.
- **Stub area**
  - A stub area is an area which doesn't receive *external* routes. External routes are defined as routes which were distributed in OSPF via another routing protocol. Therefore, stub areas typically need to rely on a default route to send traffic to routes outside the present domain.

OSPF: Area Types

- **Totally stubby area**
  - A totally stubby area is similar to a stub area, however this area does not allow *summary* routes in addition to the *external* routes, i.e., inter-area (IA) routes are not summarized into totally stubby areas. The only way for traffic to get routed outside of the area is a default route which is the only Type-3 LSA advertised into the area. When there is only one route out of the area, fewer routing decisions have to be made by the route processor, which lowers system resource utilization.

OSPF: Area Types

- **Not-so-stubby area**
  - Also referred to as NSSA, a not-so-stubby area is a type of stub area that can import external routes and send them to the backbone, but cannot receive external routes from the backbone or other areas.
  - Used when you want a stub area to be able to inject external route information
  - *Advanced Routers Topic*

OSPF: Area Types

- Every area has to directly connect to the backbone area
  - The backbone area is Area 0
  - This is because of the way OSPF operates when transmitting
  - Backbone area routers can connect to another area by having an interface in that area
    - These routers are called Area Border Routers (ABRs)
  - The type 4 LSA is used to advertise information about an ASBR – *Advanced Class*

OSPF: LSA Details

- In the last lab we saw two types of LSAs
  - Type 1 – router LSA
  - Type 2 – network LSA
- This lab will introduce the type 3 LSA
  - The network summary LSA
    - Used to advertise networks between areas
- The type 5 LSA is used for advertising external route sources

LAB!