7b. OSPF Concepts

BASICS

Link State protocols ensure that every router in the AS shares the same knowledge about the topology (flooding), then each router makes its own choices as to the best path for any given destination.

Router ID (RID)—A 32-bit number, displayed like an IP address, chosen on OSPF initialization. Order of preference:

- Explicit

```
R1(config)# router-id 10.0.0.1
```
Explicitly set the router ID
```
R1# clear ip ospf
```
Re-initialize OSPF so that a new RID is chosen (in this case the explicit one we just typed)

- Highest "up" loopback IP—These are stable in that they go up/down only by command.
- Highest "up" interface IP

Tables

- Neighbors
- Link State Database (LSDB)
- Also the routing table, but that's used by everything

```
R1# show ip ospf neighbors
R1# show ip ospf database
R1# show ip route
```

LSA (Link State Advertisement)—Data structure used within the LSDB and in LSU (Link State Update) messages. Not a message itself.

Djikstra SPF (Shortest Path First) Algorithm—used on each router to choose the best route to each destination, using information in the LSDB. The cost is calculated for every possible path to each destination subnet before the choice is made for that subnet.

Route Cost = sum of all outgoing interface costs along route. Ignore inbound interface costs.

Messages

- Hello—Multicast containing own RID and "seen" field for the RID of a potential neighbor, or null on the first hello. Put neighbor's ID in the "seen" field if parameters match (willing to be neighbors).
- DD (Database Description)—Lists available LSAs
- LSU (Link State Update)—Message containing LSA contents, perhaps for several LSAs

Exchange

- Use DD to send list of LSAs available
- From that list, ask for LSAs you don't already have.

States

- Init—After a router has received a hello with a null "seen" field and has responded with the other RID in its own "seen" field. The other router will be listed as neighbor before 2-way.
- 2-Way (Adjacent)—Router has received a hello with its own RID in a seen field and is willing to exchange LSAs. Obviously all checks have passed. Once both routers are 2-way (after the 3rd hello), database exchange can happen immediately.
- ExStart—Exchanging DDs
- Loading—Exchanging LSUs
- Full (Fully Adjacent)—LSDBs are converged (contents match). Fully adjacent.
  - Exchange Hellos every "Hello Interval"
  - If no hello received by "Dead interval" (default 4 x HELLO) neighbor considered down
  - Upon topo change, affected LSAs in the LSDB are changed and flooded. Each router reruns SPF to calculate routes.
  - Periodically (30 minutes), each router must re-flood LSAs that it originated. Each LSA has its own timer based on when it was created.

Designated Routers (DR)—Only on Ethernet with 2+ routers. Full state (and database exchanged) is achieved between DR and each other router. DROTHERs never achieve full state with each other. ("show ip ospf" shows them permanently stuck in 2-way.) Highest router-id wins. Change with interface priority.

```
R2(config-if)# ip ospf priority <0-255>
```
Default is 1. Higher wins.

BDR (Backup DR)—Also elected on Ethernet and reaches full status with all neighbors. Watches the DR and takes over on failure, with new BDR elected if that happens.

```
R2# show ip ospf neighbor
```
Tells neighbors and their state (e.g. full with DR and BDR)

MULTI-AREA OSPF

Best to break up large network into areas of ~50 routers. LSDB only lists routers & links inside its area. Subnets outside the area appear as if directly connected to the ABR (Area Border Router). Areas send summary LSAs to each other, which lack the detailed topology within the area.

Multi-Area Advantages

- Less memory for the LSDB
- Reduced CPU load because the SPF algorithm operates on a smaller LSDB—SPF load difference is exponential to changes in area size.
- Topology changes only require an SPF re-run within the one area where they occurred. It runs on each router in the area for each interface state-change.
- Reduced bandwidth inter-area because only summaries are sent outside each area

Area Design Guidelines:

- All interfaces in the same subnet should be in the same area (for efficient summarization)
- Areas contiguous
- All areas must connect to (or be) area 0

Terminology

Terminology	Definition
Area	All routers within have a matching LSDB
Backbone Router	Fully in area 0
Internal Router	Fully in a non-0 area
ABR (Area Border Router)	Straddles boundary, & participates in area 0 and another (some interfaces in one area, others in another, one area per interface)
Intra / Inter -area routes	To locations within / outside of a router's area

LSA TYPES (NOT COVERED IN 200-105)

Router LSA (Type 1)—One for each router in the area

- Router ID (LSA ID)
- Interfaces
- IPv4 Addresses & Masks
- Interface States
- Known Neighbors—and which interface they're on

Network LSA (Type 2)—One for each network that elects a DR where we have 1+ neighbors. Effectively, this means a non-LAN Ethernet with 2+ routers; a link. Not a point-to-point serial; they don't elect a DR. Router LSAs reference network LSAs if they exist.

- DR & BDR IP Addresses
- Subnet ID & Mask

Summary LSA (Type 3)—Tells of networks in external areas. Note: "Summary" doesn't imply summarization of the routes / networks; they remain separate. ¿¿¿ *Some kind of cost here too ???*

- Subnet ID & Mask
- RID of the ABR that created and advertises the summary

By moving to multi-area

- Fewer type 1 & 2 LSAs in each area
- Each ABR has a copy of the LSDB for each area they connect to
- Each area has a router LSA for their ABR
- Summary (Type 3) LSAs define subnets in other areas