Dynamic Routing and OSPF

Static and Dynamic Routing

- Static routing is a simplistic approach
- Shortcomings:
  - Cumbersome to configure
  - Cannot adapt to link/node failures, addition of new nodes and links
  - Doesn't scale to large networks
- Solution: Dynamic Routing

Desirable Characteristics

- Automatically detect and adapt to network topology changes
- Optimal routing
- Scalability
- Robustness
- Simplicity
- Speed of convergence
- Some control of routing choices (e.g. which links we prefer to use)

Convergence - Why do I care?

- Convergence is when all the routers have the same routing information
- When a network is not converged, there is network downtime
  - Packets don't get to where they are supposed to be going: routing loops, black holes
  - Occurs when there is a change in the status of a router or link

Other Interior Gateway Protocols (IGPs)

- RIP
  - Lots of scaling problems
  - RIPv1 is classful and officially obsolete
- EIGRP
  - Proprietry (Cisco only)
- IS/IS
  - The forerunner of OSPF
  - Multiprotocol (OSPF is IP only)

Why not use RIP?

- Distance Vector algorithm
  - Listen to neighboring routes
  - Install all routes in table, lowest hop-count wins
  - Advertise all routes in table
  - Very simple
  - Very stupid
- Broadcasts everything (not scalable)
- Metric is hop-count only
- Infinity of 16 (not large enough)
- Slow convergence (routing loops)
- Poor robustness

OSPF

- Open Shortest Path First
- Dynamic IGP (Interior Gateway Protocol)
  - Use within your own network
- Link state algorithm

Shortest Path First

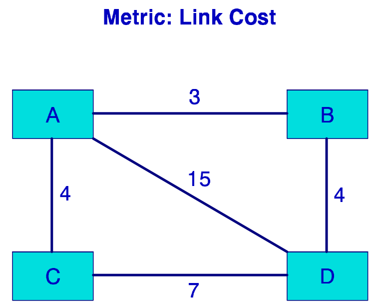

Link State Algorithm

- Each router maintains a database containing map of the whole topology
  - Links
  - State (including cost)
- All routers have the same information
- All routers calculate the best path to every destination
- Any link state changes are flooded across the network
  - "Global spread of local knowledge"

Note: Routing is not the same as Forwarding

- Forwarding: passing packets along to the next hop
  - There is only one forwarding table
  - Just has prefix and next-hop info
- Routing: populating the forwarding table
  - You might have multiple routing databases - e.g. both OSPF and BGP
  - Routing databases have more information

Routing and Forwarding

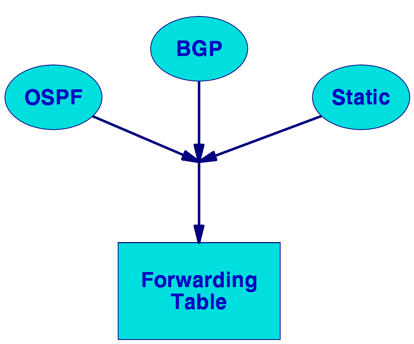

On Cisco, if the same prefix is received from multiple protocols, the "administrative distance" is used to choose between them

OSPF: How it works (1)

- "Hello" packets sent periodically on all OSPF-enabled interfaces
  - become "neighbors"
  - establishes that link can carry data
- Adjacencies (virtual point-to-point links) formed between *some* neighbors

How it works (2)

- Once an adjacency is established, trade information with your neighbor
- Topology information is packaged in a "link state announcement"
- Announcements are sent ONCE, and only updated if there's a change (or every 30 minutes)

How it works (3)

- Each router sends Link State Announcements (LSAs) over all adjacencies
  - LSAs describe router's links, interfaces and state
- Each router receives LSAs, adds them into its database, and passes the information along to its neighbors

How it works (4)

- Each router builds identical link-state database
- Runs SPF algorithm on the database to build SPF tree
- Forwarding table built from SPF tree

How it works (5)

- When change occurs:
  - Broadcast change
  - All routers run SPF algorithm
  - Install output into forwarding table

HELLO

- Broadcast* HELLO on network segment
- Receive ACK
- Establishes 2-way communication
- Repeat periodically
  - Default: HELLO sent every 10 seconds
  - Default: if no HELLO heard for 40 seconds, link is assumed to be dead
- Now establish adjacencies

* Actually uses Multicast addresses (224.0.0.5, 224.0.0.6) so that non-OSPF devices can ignore the packets

The HELLO packet

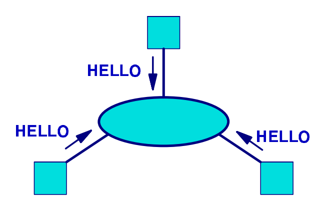

- Router priority
- Hello interval
- Router dead interval
- Network mask
- List of neighbors

These must match (Hello interval, Router dead interval, Network mask)

Neighbors

- Bi-directional communication
- Result of OSPF hello packets
- Need not exchange routing information

Who is adjacent?

- "Adjacent" neighbors exchange routing information
- Not all neighbors are adjacent
- On a point-to-point link
  - everyone
- On broadcast medium
  - not everyone
  - why?

Broadcast neighbors

Order of N^2 adjacencies

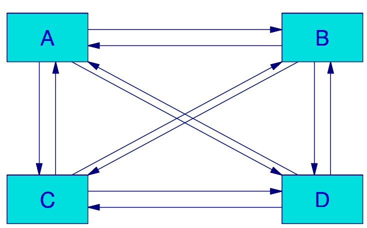

Broadcast medium

- Select a neighbor: Designated Router (DR)
- All routers become adjacent to DR
- Exchange routing information with the DR
- DR updates all the other neighbors
- Scales
  - Adjacencies reduced from N^2 to 2N
- Backup Designated Router (BDR)

LSAs propagate along adjacencies

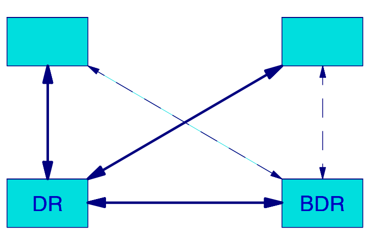

Other nice features of OSPF

- Authentication (optional)
- Equal-cost multipath
  - more than one "best" path - share traffic
- Proper classless support (CIDR)
- Multiple areas
  - For very large networks (>150 routers)
  - Aggregate routes across area boundaries
  - Keep route flaps within an area
  - Proper use of areas reduce bandwidth and CPU utilisation
  - Backbone is Area 0