Routing Protocols

"Helping systems to help themselves"

CIS 68C2-01

UNIX Network Administration

Routing Protocols

- Overview
  - Routers use routing protocols to dynamically exchange information used to build route tables
  - Two groups of protocols: **interior** and **exterior**
    - Interior Protocols
      - Used *within* an autonomous system
      - Determine and distribute *best* route information to systems
      - Administrator chooses the protocol
      - Protocols: RIP, RIP-2, OSPF, IS-IS, Hello
    - Exterior Protocols
      - Exchange *reachability information* between autonomous systems
      - ISP or administrator of autonomous system chooses the protocol
      - Protocols: EGP, BGP

Routing Protocols

- Two implementation approaches:
  - **Distance Vector** and **Link State**
- Distance Vector
  - A "gossipy" protocol
  - A router *advertises* its distance to neighboring networks
    - *"How many routers (hops) away is a given network?"*
    - Distance metric is measured in terms of *hops*
      - Each router typically adds one hop (router metric is 1)
  - Other routers listen for *route advertisements*
    - Replace their expensive routes with advertised cheaper routes
    - More expensive routes advertisements are ignored

Routing Protocols

- Distance Vector
  - Limitations
    - Infinite loops are a challenge
      - A receives from B, sends to C, which sends back to B
    - Slow convergence
      - Takes time before routes settle down
    - Requires arbitrary maximum hop count
      - Prevents infinite loops or delays
    - Frequent, bulky updates
      - All information exchanged very frequently
        - RIP: every 30 seconds
        - EGRP/EIGRP: every 90 seconds

Routing Protocols

- Link State
  - Exchange the *state* of adjacent routers
  - *"Router A is connected to router B and link is up"*
  - Requires connectivity map of entire network
    - At most, an entire autonomous system
    - Map is a (possibly very large) *directed graph*
  - Yields an end-to-end overview of the entire network (AS)
    - Cost to reach any end point is known
      - Better than only having cost to next hop
  - Are more complicated than distance vector
  - Requires more CPU and RAM

Interior Routing Protocols

- RIP – Routing Information Protocol
  - Simple Distance Vector protocol
  - Implemented in **routed** daemon
  - Uses hop count as its metric
  - Limitations
    - Bandwidth hog – broadcasts all information
    - Slow convergence after network changes
      - Heuristic solutions: split horizon, poison reverse, triggered updates
    - Beyond 15 hops is considered unreachable
    - No support for CIDR (no way to network masks)

Interior Routing Protocols

- RIP-2
  - RIP w/extensions to overcome RIP's limitations
  - Backwards compatible with RIP
  - Enhancements
    - Supports CIDR (distributes network masks)
    - Minor security enhancements (not widely used)
    - Next hop updates allow advertising routes not your own
    - Route tags to allow propagation of externally discovered routes
    - Uses multicast to reduce network load

Interior Routing Protocols

- OSPF – Open Shortest Path First
  - Link-state protocol
  - Implemented in **gated** daemon
    - Widely-used and most popular multi-protocol routing daemon
  - Pros
    - Works well in large, complex environments
    - Much less chatty than RIP
    - Handles multiple paths to a destination
    - Allows arbitrary cost metrics (hops are default w/gated & Cisco)
    - Allows sectioning networks into Areas, Backbone, and Stub Areas
    - Supports authentication (clear text and MD5)

Interior Routing Protocols

- OSPF
  - Cons
    - Very complex (see appendix B, page 449-496)
  - Determines neighbors via HELLO packets
  - Link State Announcements are flooded
    - Neighbors + cost are advertised via each interface
  - Neighbors are added to graph
  - Map is pruned of more expensive routes
  - Use of designated routers improves efficiency
    - Treats all others as neighbors
    - Reduces map size

Interior Routing Protocols

- OSPF Directed Graph
  - **A** sends HELLO packet to neighbors and neighbors respond with HELLO packets
  - **A** floods LSA to neighbors **B**, **C**, **F**, **H**, and **I**
  - Neighbors flood LSAs in turn to their neighbors
  - **A** adds paths to its map
  - **A** prunes map of more expensive routes

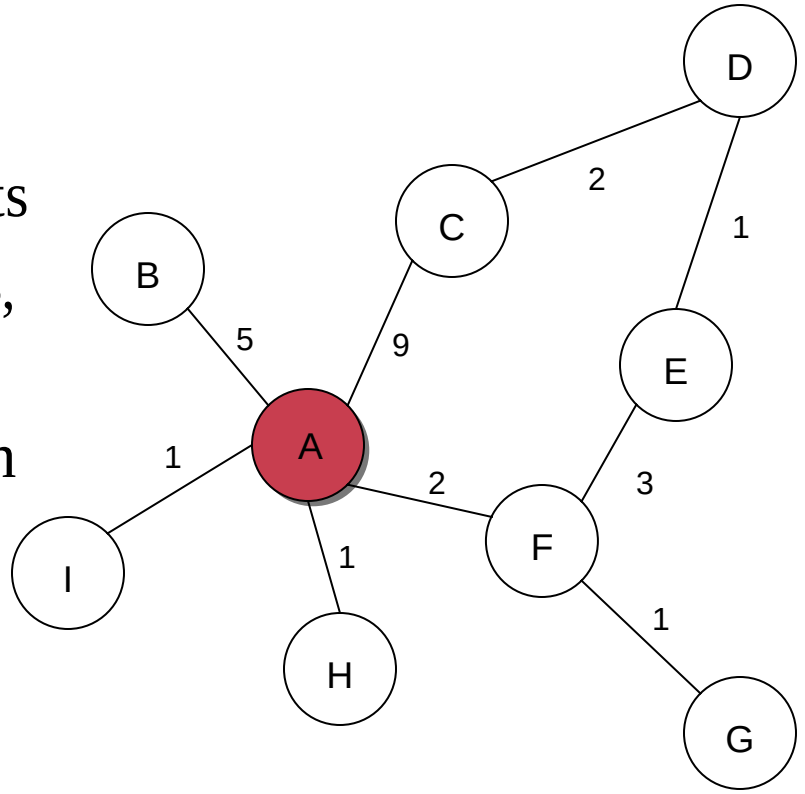

Interior Routing Protocols

- Others
  - IGRP / EIGRP
    - Proprietary protocols used in Cisco routers
  - IS-IS – Intermediate to Intermediate System
    - Rarely used ISO attempt at OSPF-like functionality
  - Hello
    - Was used in 56k NFSNET
  - MOSPF / DVMRP / PIM
    - Multicast IP routing protocols
  - Router Discovery Protocol
    - Not (yet) supported by all routers

Exterior Routing Protocols

- EGP – Exterior Gateway Protocol
  - Only announces its own networks
  - An EGP system *acquires a neighbor*
    - Via *Hello* and *I-Heard-You* messages
  - Does not choose best route
    - Distance vector information is not interpreted, just passed
      - Each autonomous system uses its own metric
    - Trusted core gateways evaluated and redistributed *best* routes
  - No longer popular

Exterior Routing Protocols

- BGP – Border Gateway Protocol
  - Leading exterior protocol
  - Path vector protocol
    - Provides entire end-to-end path
  - Supports *Policy-based Routing*
    - Allows non-technical policies to control routing
  - Does not require central routing authority
  - Very efficient

Routing Daemons

- The **routed** Daemon
  - Long time standard in UNIX
  - Supports RIP only
    - Some implementations support RIP-2
  - Options
    - **-s** Run in server mode advertises routes
    - **-q** Run in quiet mode only listens
    - **-g** Advertises a single gateway
  - Configuration file: /etc/gateways
    - For static routes

Routing Daemons

- The **gated** Daemon
  - Generally preferred over **routed**
  - Supported many protocols
    - Interior: RIP, RIP-2, OSPF, IS-IS
    - Exterior: EGP, BGP (v3, v4), HELLO
  - Configuration file: /etc/gated.conf
  - Administration via the **gdc** command