

BGP

Interdomain Internet Routing

Quiz Review Notes 2007

Some terms to know

- AS (Autonomous System)- group of networks under a common administration and with common routing policies
- Wide-Area routing architecture is divided into Ases.
- IGP – Interior Gateway Protocol- Routing protocol *within* each AS
- EGP – Exterior Gateway Protocol (eg BGP)
- iBGP – internal BGP : used within an AS
- eBGP – external BGP: used between ASes

BGP

- Wide-Area Routing protocol – exchanges routing information between Ases
- Difference between BGP and IGP
 - IGPs concerned with loop-free routing and shortest path
 - EGPS such as BGP are more concerned about scalability and policy

eBGP and iBGP

- eBGP session between BGP routers in different Ases
- iBGP between BGP routers in the same AS
 - More than one router participates in eBGP sessions and learns a subset of the routes
 - Each such eBGP router disseminates this information to other routers in the AS
- Both use the same protocol
- iBGP is NOT IGP (iBGP messages routed via IGP)
- Why is iBGP needed?

BGP protocol

- Runs over TCP
 - Router sends OPEN message to start session
 - Two kinds of Updates
 - Announcements: Changes to existing routes or new routes
 - Withdrawls: For routes that no longer exist
- Keep Alive Messages

Inter-AS relationships

- Transit
 - Provider ->customer
 - Customer pays provider for internet access
- Peering
 - May not involve financial settlement
 - Two Ases provide limited mutual access to each other's routing tables
 - Tier1 ISPs sometimes peer with each other to obtain global reach

Exporting Routes

- Route Advertisement from B->A implies B will forward all packets sent via A to the advertised prefixes.
- Need route filtering. Why?
- Customer routes– advertise to everyone
- Provider routes- selective advertisement (Example)
- Peers- selective advertisement

Selecting Routes

- How does a router determine best route to a destination?
 - Use BGP attributes
 - LOCAL PREF (customer > peer > provider)
 - First criteria used to select routes.
 - Assigned when importing paths
 - ASPATH
 - Path vector of Ases that this route has been through
 - Used for loop avoidance
 - pick route with shortest AS path if LOCAL PREF not set.
 - MED
 - for choosing between multiple exits between two Ases
 - Normally ignored unless Ases have financial settlement
- Others attributes (eBGP > iBGP) etc

Interesting phenomenon

- Hot-Potato vs Cold- Potato Routing
- Example –

