

Lecture 16:

Border Gateway Protocol

CSE 123: Computer Networks
Alex C. Snoeren



Some figures courtesy Mike Freedman

Lecture 16 Overview




- Border Gateway Protocol (BGP)
 - The canonical path vector protocol
 - How routing gets done on the Internet today



Path-vector Routing


- Extension of distance-vector routing
 - ◆ Support flexible routing policies
 - ◆ Avoid count-to-infinity problem
- Key idea: advertise the entire path
 - ◆ Distance vector: send *distance metric* per destination
 - ◆ Path vector: send the *entire path* for each destination





Loop Detection



- Node can easily detect a loop
 - ◆ Look for its own node identifier in the path
 - ◆ E.g., node 1 sees itself in the path “3, 2, 1”
- Node can simply discard paths with loops
 - ◆ E.g., node 1 simply discards the advertisement



Policy Support



- Each node can apply local policies
 - ◆ Path selection: Which path to use?
 - ◆ Path export: Which paths to advertise?
- Examples
 - ◆ Node 2 may prefer the path “2, 3, 1” over “2, 1”
 - ◆ Node 1 may not let node 3 hear the path “1, 2”




Border Gateway Protocol



- Interdomain routing protocol for the Internet
 - ◆ Prefix-based path-vector protocol
 - ◆ Policy-based routing based on AS Paths
 - ◆ Evolved during the past 28 years

- **1989 : BGP-1 [RFC 1105], replacement for EGP**
- **1990 : BGP-2 [RFC 1163]**
- **1991 : BGP-3 [RFC 1267]**
- **1995 : BGP-4 [RFC 1771], support for CIDR**
- **2006 : BGP-4 [RFC 4271], update**

Basic BGP Operation



Step-by-Step




- A node learns multiple paths to destination
 - ◆ Stores all of the routes in a routing table
 - ◆ Applies policy to select a single active route
 - ◆ ... and may advertise the route to its neighbors
- Incremental updates
 - ◆ Announcement
 - » Upon selecting a new active route, add own node id to path
 - » ... and (optionally) advertise to each neighbor
 - ◆ Withdrawal
 - » If the active route is no longer available
 - » ... send a withdrawal message to the neighbors



A Simple BGP Route

- Destination prefix (e.g., 128.112.0.0/16)
- Route attributes, including
 - ◆ AS path (e.g., “7018 88”)
 - ◆ Next-hop IP address (e.g., 12.127.0.121)



(some) BGP Attributes




- **AS path:** ASs the announcement traversed
- **Next-hop:** where the route was heard from
- **Origin:** Route came from IGP or EGP
- **Local pref:** Statically configured ranking of routes within AS
- **Multi Exit Discriminator:** preference for where to *exit* network
- **Community:** opaque data used for inter-ISP policy

Export Active Routes



- In conventional path vector routing, a node has one **ranking function**, which reflects its routing policy




BGP Decision Process



- Default decision for route selection
 - ◆ Highest local pref, shortest AS path, lowest MED, prefer eBGP over iBGP, lowest IGP cost, router id
- Many policies built on default decision process, but...
 - ◆ Possible to create arbitrary policies in principal
 - » Any criteria: BGP attributes, source address, prime number of bytes in message, ...
 - » Can have separate policy for inbound routes, installed routes and outbound routes
 - ◆ Limited only by power of vendor-specific routing language




Example: Local Pref






Example: Short AS Path

Shorter AS Paths are more preferred



AS Paths vs. Router Paths



BGP Has Lots of Problems



- Instability
 - ◆ Route flapping (network x.y/z goes down... tell everyone)
 - ◆ Long AS-path decision criteria defaults to DV-like behavior (bouncing)
 - ◆ Not guaranteed to converge, NP-hard to tell if it does
- Scalability still a problem
 - ◆ >500,000 network prefixes in default-free table today
 - ◆ Tension: Want to manage traffic to very specific networks (eg. multihomed content providers) but also want to aggregate information.
- Performance
 - ◆ Non-optimal, doesn't balance load across paths

For Next Time



- Read P&D 3.4
- Project 2 introduced in discussion tomorrow