

# Inter-Domain Routing (Continued)

CS640, 2015-03-03

## Announcements


- Assignment #2 due at 11pm today
- Quiz #3 on Thursday

## Overview

- Inter-Domain Routing Issues


## Other Inter-Domain Routing Issues

- **\*\*What are the benefits of path vector routing (i.e. BGP)?**
  - Loop free -- if receive a path that already contains your ASN, discard the advertisement
  - No count to infinity
  - Path selection based on local policy -- no need to tell neighbors all possible paths
    - Important for AS to choose, because certain path selections cost AS money!
  - Scalable -- there is no flooding
- Convergence
  - Internet is always in flux -- convergence is important
  - Go through path exploration (series of temporary routes) before convergence
  - Example



- To get to AS 0
  - AS 1 has possible routes: (0), (2, 0), (3, 0)
  - AS 2 has possible routes: (0), (1, 0), (3, 0)
  - AS 3 has possible routes: (0), (1, 0), (2, 0)
- If AS 0 fails, each AS will withdraw first route and resort to alternate route
  - AS 1 -> AS 2: withdraw (1, 0)
  - AS 1 -> AS 3: withdraw (1, 0), announce (1, 2, 0)
  - AS 2 -> AS 1: withdraw (2, 0)
  - AS 2 -> AS 3: withdraw (2, 0), announce (2, 1, 0)
  - AS 3 -> AS 1: withdraw (3, 0)
  - AS 3 -> AS 2: withdraw (3, 0), announce (3, 1, 0)
- Depending on message ordering, it is possible
  - AS 3 -> AS 2: new route (3, 1, 2, 0) -- AS 2 discards
  - AS 2 -> AS 1: new route (2, 3, 1, 0) -- AS 1 discards

- Eventually all routes get withdrawn
- Route flapping
  - Faulty links can cause series of announce/withdraw advertisements
  - Convergence is on the order of minutes => flapping is bad!
  - Causes lots of BGP messages, and high router load to deal with these
  - Solution: add damping mechanism
- Hot potato routing
  - Transit AS wants to carry traffic through their network as little as possible
  - Prefer to hand off to another AS as soon as possible
  - Example



- **\*\*How does AS 2 convince AS 1 to handoff traffic in New York?**
  - Only AS 2's BGP speaker in New York advertises 2.2.2.0/24 to AS 1
    - AS 1 will never use the link to AS 2 in Los Angeles
    - If the link in New York fails, no traffic can reach 2.2.2.0/24 until AS 2 detects the failure and advertises from the BGP speaker in Los Angeles
  - Multi-exit discriminator (MED)
    - Optional attribute included in a BGP advertisement to tell neighboring AS which link is preferred
    - It's still up to AS 1 to obey AS 2's hint
  - Prepending
    - Generally, ASes prefer advertisements with shorter paths
    - Add your own AS number to the path a few times so the path appears longer
    - AS 2 will advertise path (2,2,2) in Los Angeles and (2) in New York
    - It's still up to AS 1 to prefer the shorter path

**\*\*Routing worksheet**