

# ICS 451: Today's plan

- IP routing: domains, policy, BGP

# High-Level Structure of the Internet

- thousands of domains (Autonomous Systems)
- each operated independently
- some are for profit
  - some willing to carry transit traffic
  - others are not
- some are mostly to generate content, others mostly to consume content
- who pays?

# customer and peer relationships among domains

- stubs only connected to another domain
- multi-homed domains multiply connected
- transit: multi-homed and willing to carry data
- stub domains usually have to pay their provider
- providers at the core of the internet usually peer with each other without paying
  - can peer at an Internet Exchange Point (IXP)
  - or through a private peering link
- other providers negotiate case-by-case

# policy: where the traffic goes

- Autonomous System A peers with B and is a customer of C
  - should send to B all traffic for B
  - doesn't want to send B's traffic to C
- import: what routes we accept from domain
- export: what routes we give to domain
- ranking: which routes are preferred
  - based on non-technical factors

# BGP

- Border Gateway Protocol (v4)
- each AS has an AS number
- similar to distance vector, but each route lists the number of every AS on the route
  - path vector protocol
- policy preferences can be applied based on the list of AS numbers
- only “best” route (by policy) forwarded

# BGP properties

- runs over TCP (port number 179)
- BGP speakers in different domains establish BGP sessions with each other
  - sessions are configured by network admins
- information is only sent once
  - assuming reliable delivery
- updates only report changes from previously sent messages
- if a connection is inactive for 90s, all its routes are withdrawn

# BGP messages

- open: session initialization
- notification: session shutdown
- update
  - withdraw route (*withdraw* message)
  - new route (*update* message)
  - update existing route (*update* message)
- keepalive: avoid 90s timeout
  - sent if there is ever a 30s or longer gap in update messages

# RIB and FIB (and MIB)

- Routing Information Base contains all acceptable routes
  - may include multiple routes to same destination
- Forwarding Information Base contains at most one route to each destination
  - used to actually forward packets
- originally, Management Information Base contains description of managed systems



# BGP Route Management

- one import filter for each peer
- any routes that pass are added to RIB
- each best route from RIB placed into FIB
  - “best” determined by policy
- routes from FIB that pass the export filter for a peer, are sent to that peer
- new routes added to the FIB (new “best” routes) are re-advertised immediately
  - like Distance-Vector
  - but change routes advertised too

# Multiple BGP routers

- multiple BGP routers within a domain use BGP to exchange routing information
  - iBGP session used within domain
    - routers may not be directly connected
  - eBGP session used across domains
    - only between directly connected routers
- iBGP used to forward information learned over eBGP only
- need to forward packets between BGP routers
  - MPLS, IP-in-IP

# Routing Policy

- routes with higher local preference used first
  - customer routes have high local pref
  - provider (expensive) routes have low local pref
- shortest path used when local pref is same
- prefer routes with smallest Multi-Exit Discriminator (MED) among same AS routes
- prefer routes going outside AS
- prefer routes with closest next-hop
- break ties with lower router ID

# BGP summary

- used for interdomain routing
- policy more important than path optimality
- TCP connections allow focus on maintaining state rather than ensuring packet delivery
- connections set manually between trusted peers