

CSE/EE 461 – Lecture 11

Inter-domain Routing

David Wetherall
djw@cs.washington.edu


This Lecture

- Focus
 - How do we make routing scale?
- Inter-domain routing
 - ASes and BGP

Application
Presentation
Session
Transport
Network
Data Link
Physical

Structure of the Internet

- Inter-domain versus intra-domain routing




djw // CSE/EE 461, Autumn 2002

L11.3

Inter-Domain Routing

- Network comprised of many Autonomous Systems (ASes) or domains
- To scale, use hierarchy: separate inter-domain and intra-domain routing
- Also called interior vs exterior gateway protocols (IGP/EGP)
 - IGP = RIP, OSPF
 - EGP = EGP, BGP




djw // CSE/EE 461, Autumn 2002

L11.4

Inter-Domain Routing

- Border routers summarize and advertise internal routes to external neighbors and vice-versa
- Border routers apply policy
- Internal routers can use notion of default routes
- Core is “default-free”; routers must have a route to all networks in the world




djw // CSE/EE 461, Autumn 2002

L11.5

Exterior Gateway Protocol (EGP)

- First major inter-domain routing protocol
- Constrained Internet to tree structure; no longer in use



djw // CSE/EE 461, Autumn 2002

L11.6

Border Gateway Protocol (BGP-4)


- EGP used in the Internet backbone today
- Features:
 - Path vector routing
 - Application of policy
 - Operates over reliable transport (TCP)
 - Uses route aggregation (CIDR)

djw // CSE/EE 461, Autumn 2002

L11.7

Path Vectors

- Similar to distance vector, except send entire paths
 - e.g. 321 hears [7,12,44]
 - stronger avoidance of loops
 - supports policies (later)
- Modulo policy, shorter paths are chosen in preference to longer ones
- Reachability only – no metrics




djw // CSE/EE 461, Autumn 2002

L11.8

An Ironic Twist on Convergence

- Recently, it was realized that BGP convergence can undergo a process analogous to count-to-infinity!




- AS 4 uses path 4 1 X. A link fails and 1 withdraws 4 1 X.
- So 4 uses 4 2 1 X, which is soon withdrawn, then 4 3 2 1 X, ...
- Result is many invalid paths can be explored before convergence

djw // CSE/EE 461, Autumn 2002

L11.9

Policies

- Choice of routes may depend on owner, cost, AUP, ...
 - Business considerations
- Local policy dictates what route will be chosen and what routes will be advertised!
 - e.g., X doesn't provide transit for B, or A prefers not to use X



djw // CSE/EE 461, Autumn 2002


L11.10

Simplified Policy Roles

- Providers sell Transit to their customers
 - Customer announces path to their prefixes to providers in order for the rest of the Internet to reach their prefixes
 - Providers announces path to all other Internet prefixes to customer C in order for C to reach the rest of the Internet
- Additionally, parties Peer for mutual benefit
 - Peers A and B announce path to their customer's prefixes to each other but do not propagate announcements further
 - Peering relationships aren't transitive
 - Tier 1s peer to provide global reachability

Multi-Homing


- Connect to multiple providers for reliability, load sharing



- Choose the best outgoing path to P out of any of the announcements to P that we hear from our providers
 - Easy to control outgoing traffic, e.g, for load balancing
- Advertise the possible routes to P to our providers
 - Less control over what paths other parties will use to reach us

Impact of Policies – Example

- Early Exit / Hot Potato
 - “if it’s not for you, bail”
- Combination of best local policies not globally best
- Side-effect: asymmetry



djw // CSE/EE 461, Autumn 2002

L11.13

Operation over TCP

- Most routing protocols operate over UDP/IP
- BGP uses TCP
 - TCP handles error control; reacts to congestion
 - Allows for incremental updates
- Issue: Data vs. Control plane
 - Shouldn't routing messages be higher priority than data?

djw // CSE/EE 461, Autumn 2002

L11.14

Key Concepts

- Internet is a collection of Autonomous Systems (ASes)
 - Policy dominates routing at the AS level
- Structural hierarchy helps make routing scalable
 - BGP routes between autonomous systems (ASes)