Lecture 4

CS625: Advanced Computer Networks
Fall 2004

Monday, 04 August 2003

Bhaskaran Raman
CSE, IIT-Kanpur

http://www.cse.iitk.ac.in/users/braman/courses/cs625-fall2004/outline.html

Outline for Today

- Internet routing architecture
- Border Gateway Protocol (BGP)
-
- Class timing change?
- Course projects:
  - Two kinds
  - Sign-up sheet, Appointment sheet
- *Scribe for today?*

Routing (recall)

- Routing protocols
  - Link-State
  - Distance-Vector
  - Path-Vector
- Hierarchical routing
- Address and Network mask

Internet Service Providers (ISPs)

- Internet Service Provider (ISP):
  - A business entity which provides *Internet connectivity* to its customers
  - Example: VSNL
  - Customer gets address(es) and connectivity
  - How is *connectivity* ensured?

ISP Relationships

- Business relationships between ISPs:
  - Peer-Peer
  - Provider-Customer (Transit relationship)
- ISP relationships are NOT transitive
  - Connectivity is not transitive
- Reasons to peer, and to not peer [Nor00]

The ISP or AS Hierarchy

- Autonomous System
  - An ISP in most cases
  - A large ISP may consist of many ASes
- Hierarchy of ISP relationships [SARK02]
  - Tier-1, Tier-2, Tier-3...
  - Dense-core, Transit-core, Outer-core, Regional-ISPs, Customers

Border Gateway Protocol (BGP)

- BGP is the routing protocol in the AS graph
  - Path-vector based protocol
  - Data path is reverse of control path
  - Allows for policy-based routing
- Each AS has border routers which speak BGP
- BGP is implemented on top of TCP!

BGP Messages

- OPEN
  - For initialization, exchange of AS numbers
- UPDATE
  - Main routing messages
  - Route announcements and withdrawals
- NOTIFICATION
  - Before closing a connection
- KEEP-ALIVE
  - Periodic heart-beat to keep BGP session alive

BGP Path Attributes

- ORIGIN
  - How the route to the prefix was learnt
- AS-PATH
  - To implement the path vector protocol
- NEXT-HOP

Path Attributes (continued)

- MULTI-EXIT-DISCRIMINATOR (MED)
- AS2 announces about AS3 to AS1
  - With MED=10 on Link A, MED=50 on Link B

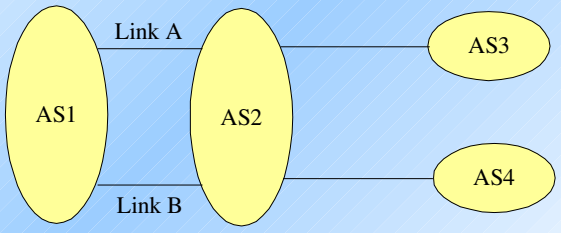

Path Attributes (continued)

- LOCAL-PREF
- AS4 may prefer route via AS3

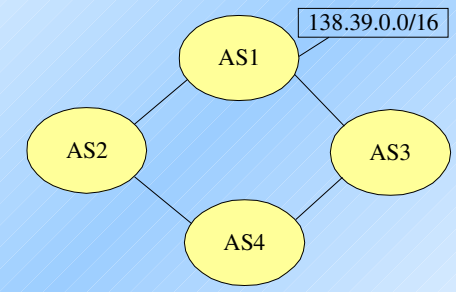

Some Remarks

- I-BGP versus E-BGP
  - I-BGP used between routers of the same AS
  - This is NOT the same as IGP
- Route selection: order of rules
  - LOCAL-PREF
  - AS-PATH
  - MED
  - E-BGP > I-BGP
  - IGP cost

More Remarks

- BGP not guaranteed to converge!
  - Has poor convergence properties in practice
  - Can take up to a few min. to recover from failure
  - Overlay networks (later in course)
- About 100K BGP routers
  - Largest distributed system to date
  - About 150,000 routing entries per router!
    - http://bgp.potaroo.net/
  - Its a surprise that it works!

Later this Week

- OSPF: Open Shortest Path First
  - Guest lecture by Dr. Mukul Goyal
- Border Gateway Protocol (BGP)
  - Policy control, convergence, and other issues