

ECE 356/COMPSI 356

Computer Network Architecture

Inter-domain Networking.
Border Gateway Protocol (BGP).

Wednesday October 2nd, 2019

Recap


- Last lecture: link state routing
- Readings for this class: **PD 4.1.1, 4.1.2**

Lecture Outline

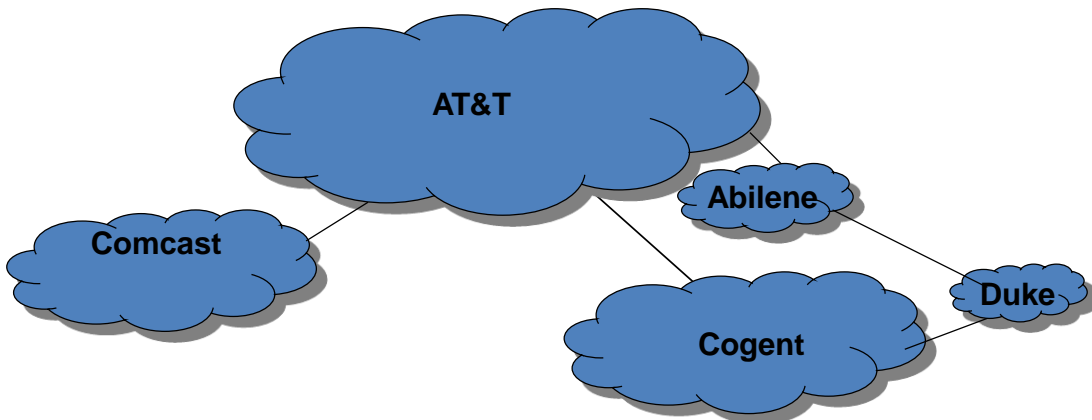
- Inter-domain routing
- BGP

3

The Internet

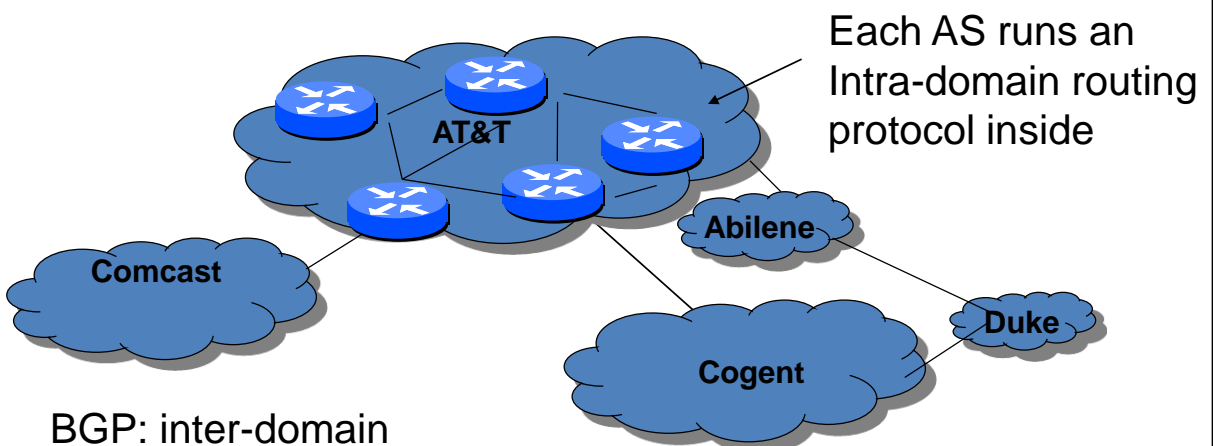


The Internet: Zooming In 2x



Duke UNIVERSITY

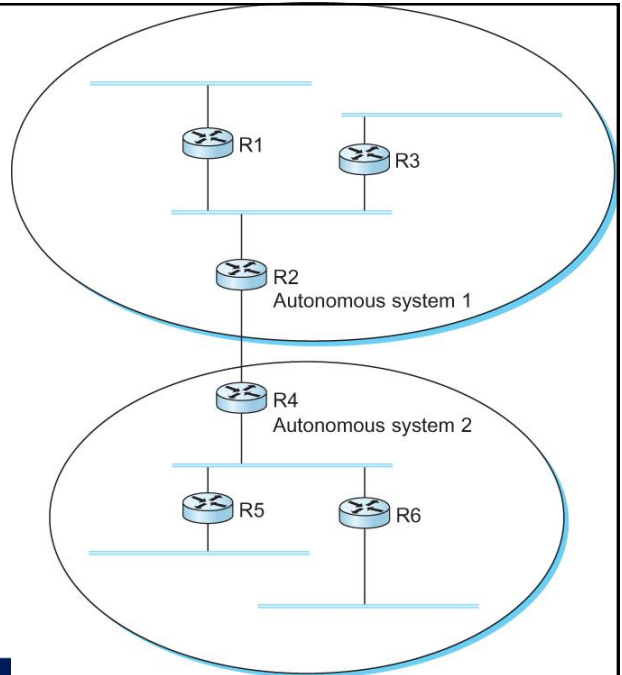
Intra-domain vs. Inter-domain Routing



Duke UNIVERSITY


Inter-domain Communications

- Always through the designated points
- E.g., if R3 and R6 are directly connected, they still need to go through designated *Area Border Routers (ABRs)*
 - Also called *gateways*



Duke UNIVERSITY


ASs Are Not Created Equal (1/4)



- **Local traffic:** traffic that originates at or terminates on nodes within an AS
- **Transit traffic:** traffic that passes through an AS


Duke UNIVERSITY

ASs Are Not Created Equal (2/4)




- *Stub AS*: an AS that has only a single connection to one other AS
 - Will only carry local traffic

ASs Are Not Created Equal (3/4)



- *Multihomed AS*: an AS that has connections to more than one other AS, but refuses to carry transit traffic

ASs Are Not Created Equal (4/4)

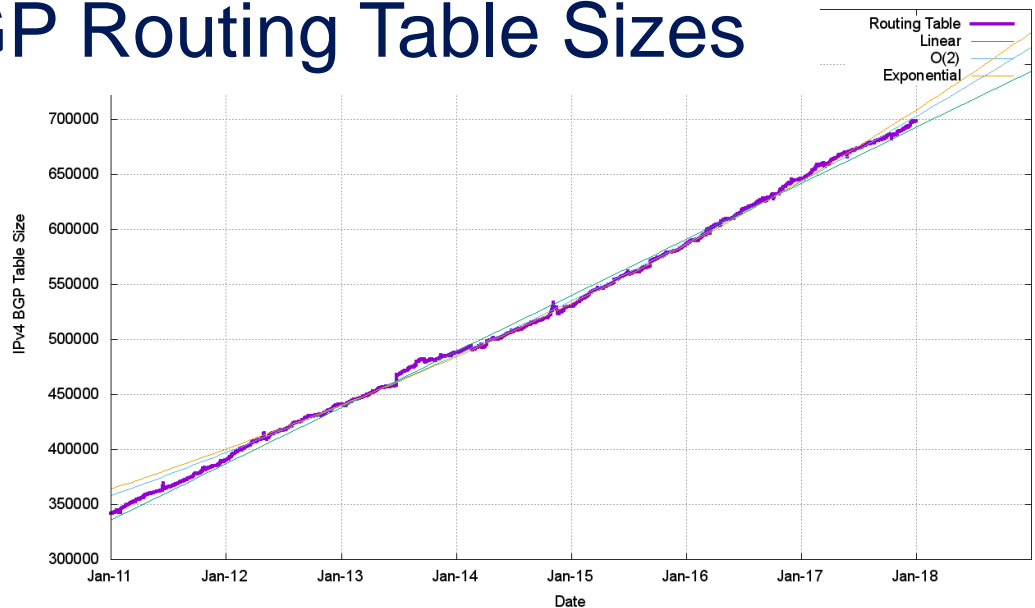


- *Transit AS*: an AS that has connections to more than one other AS, and is designed to carry both transit and local traffic

Challenges in Interdomain Routing: Scalability

- An Internet backbone router must be able to forward any packet destined anywhere in the Internet
 - Having a routing table that will provide a match for any valid IP address

BGP Routing Table Sizes



Duke UNIVERSITY

Challenges in Interdomain Routing

- Cannot find the true shortest path
 - Domains are autonomous and can use different routing metrics
- Find a path, advertise reachability
 - “Can reach this network through this AS”
- Paths must be policy-compliant, and policies can be complex
 - E.g., “*whenever possible, I prefer to send traffic via AS X than AS Y, but I’ll use AS Y if it is the only path, and I never want to carry traffic from AS X to AS Y or vice versa*”

Duke UNIVERSITY

Lecture Outline

- Inter-domain routing
- **BGP**

15

BGP (1/4)

- BGP = Border Gateway Protocol
 - Currently in version 4, specified in RFC 1771 (~ 60 pages)
- Inter-domain routing protocol for routing between autonomous systems
- Uses TCP to establish a BGP session and to send routing messages over the BGP session

BGP (2/4)

- BGP is a *path vector protocol*
 - Similar to distance vector routing, but routing messages in BGP contain complete paths
- Network administrators can specify routing policies
- BGP's goal is to find any path (not an optimal one)
 - Since the internals of the AS are never revealed, finding an optimal path is not feasible

BGP (3/4)

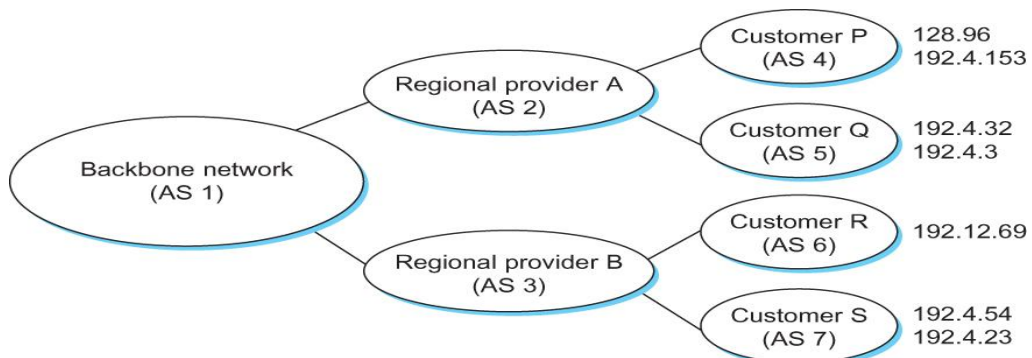
Each AS has:

- One BGP *speaker* that advertises:
 - Local networks
 - Other reachable networks (transit AS only)
 - Gives *path* information
- One or more border “gateways”
 - The routers through which packets enter and leave the AS
 - Need not be the same as the speakers

BGP (4/4)

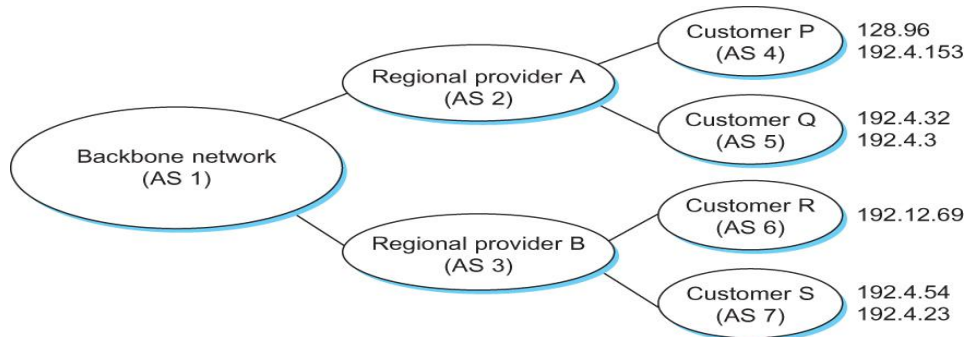
- BGP advertises *complete paths* as an enumerated lists of ASs to reach a particular network

BGP Example (1/2)



- Speaker for AS 2 advertises reachability to P and Q
 - Network 128.96, 192.4.153, 192.4.32, and 192.4.3, can be reached directly from AS 2

BGP Example (2/2)



- Speaker for backbone network then advertises
 - Networks 128.96, 192.4.153, 192.4.32, and 192.4.3 can be reached along the path <AS 1, AS 2>

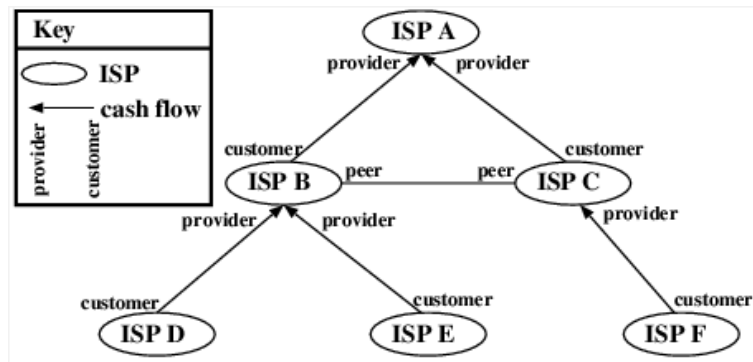
Loop Prevention in BGP

- ASs can see themselves in advertised paths
- AS numbers carried in BGP need to be unique
 - AS can only recognize itself in the AS path if no other AS identifies itself in the same way
- AS numbers are 16-bit numbers assigned by a central authority
 - Stub ASs do not need unique numbers

Different ASs Use Different Routing Policies


- E.g., BGP speaker does not have to advertise a route to a destination if it has one
 - This implements a policy of not providing transit

Common AS Relationships (1/3)




- Provider - customer
 - Advertise all routes I know about to my customer
 - Advertise routes I learn from my customer to everyone

Common AS Relationships (2/3)



- Customer - provider
 - Don't advertise routes learned from one provider to another provider

Common AS Relationships (3/3)



- Peer-peer
 - Advertise customer routes only
 - Don't advertise peer routes to a provider

Lecture Summary

- Inter-domain routing
- BGP

27

Next Lecture

- Miscellaneous IP topics
 - Host configuration (DHCP)
 - Network address translation (NAT)
 - IPv6
 - Tunneling

