CSE461 Section #6

Anran Wang

Routing

Distance Vector Routing vs. Link State Routing

BGP

Practice

Distance Vector Routing

What to exchange: distance vector for all nodes

How to route: directly from distance vector

Loop protection, partition detection

Link State Routing

What to exchange: local link states

How to route: calculate shortest path

Computational expensive, more data to be stored

	Link State	Distance Vector	Path Vector
Dissemi-nation	Flood link state advertisements to all routers	Update distances from neighbors' distances	Update paths based on neighbors' paths
Algorithm	Dijsktra's shortest path	Bellman-Ford shortest path	Local policy to rank paths
Converge	Fast due to flooding	Slow, due to count-to-infinity	Slow, due to path exploration
Protocols	OSPF, IS-IS	RIP, EIGRP	BGP

BGP - The interdomain routing problem

Each AS determines its own routing policies

◦ One AS only wants to send and receive packets from the internet

◦ One AS can carry transit traffic for others if you pay this service

Political considerations

◦ Never send traffic from the Pentagon on a route through Iraq

Security considerations

◦ Traffic starting or ending at Apple should not transit Google Economic considerations

◦ Use cheaper service

Interconnection Relationships

Local

Transit

Peering

BGP Basics

Types of routers:

•Border router: packets enter and leave the AS

•BGP Speaker: handles advertisements, usually the same as border routers

Path-vector protocol

•Not distance vector or link-state

•AS Path: list of autonomous systems to reach a particular network

•Built on TCP

Loop Detection

Assign each AS a unique number

◦ BGP current version: 16 bits

Route selection

Routes via peered networks are favored over routes via transit providers

- Free!

Shorter AS paths are better

Prefer the route that has the lowest cost within the ISP

Only advertise routes that are good enough for you

Allow route withdrawal

One example

Consider the following network with 6 Ases

- AS1 is the provider for AS2, AS3, and AS4
- AS2 is the provider for AS5
- AS3 is the provider for AS5 and AS6
- AS5 and AS6 have a peer agreement

Practices

Consider the network shown in Figure 4.28, in which horizontal lines represent transit providers and numbered vertical lines are interprovider links.

(a) How many routes to P could provider Q's BGP speakers receive?

(b) Suppose Q and P adopt the policy that outbound traffic is routed to the closest link to the destination's provider, thus minimizing their own cost. What paths will traffic from host A to host B and from host B to host A take?

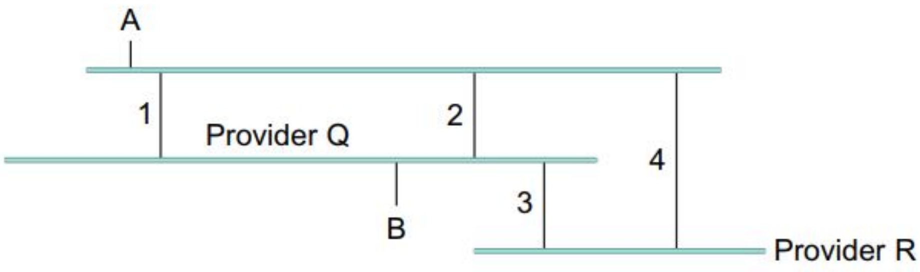

Practices

Give an example of an arrangement of routers grouped into autonomous systems so that the path with the fewest hops from a point A to another point B crosses the same AS twice. Explain what BGP would do with this situation.

Practice

Let A be the number of autonomous systems on the Internet, and let D (for diameter) be the maximum AS path length.

Assuming each AS number is 2 bytes and each network number is 4 bytes, give an estimate for the amount of data a BGP speaker must receive to keep track of the AS path to every network. Express your answer in terms of A, D, and the number of networks N.