CS 168 Introduction to the Internet: Architecture and Protocols
Fall 2022 Sylvia Ratnasamy Discussion 6:BGP and IP

1 True/False

1. iBGP is used for intradomain routing.

   Solution: False, IGP, such as OSPF or RIP, is used for intradomain routing. iBGP is used to distribute externally learned routes internally

2. Avoiding loops is one reason why BGP uses path vector.

   Solution: True, also helps ASes adopt certain policies

3. BGP always advertises a shortest path.

   Solution: False, the path relies on other policies (like money) instead of purely distances

4. BGP route advertisements use classless addressing.

   Solution: True

2 Interdomain vs Intradomain

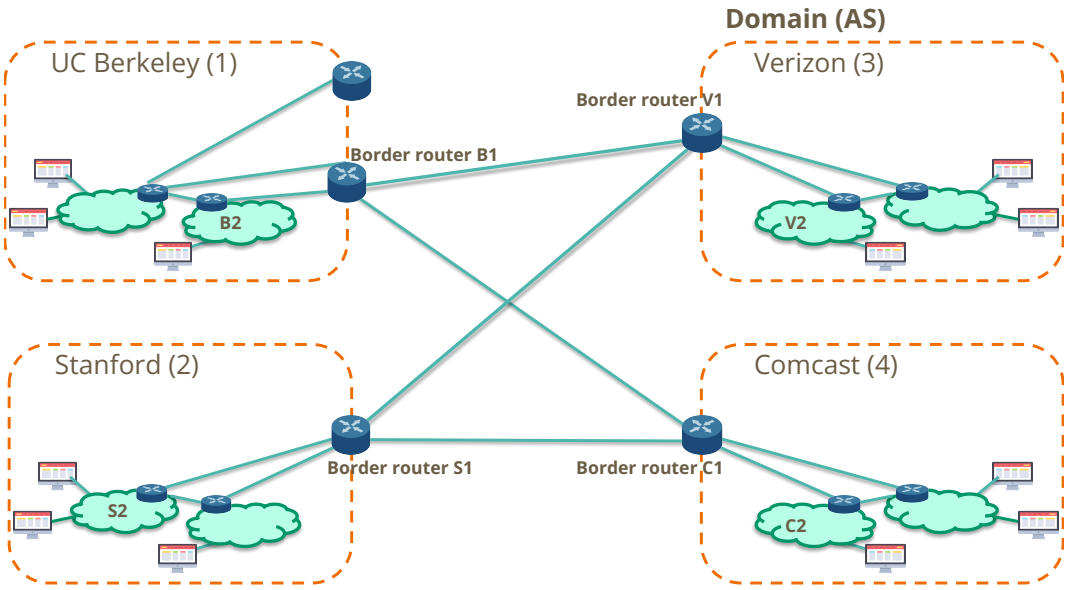

Consider the four ASes in the diagram above. ASes Berkeley, Verizon, Stanford and Comcast have border routers B1, V1, S1 and C1 respectively, and internal routers B2, V2, S2 and C2 respectively.

Berkeley and Stanford both use Comcast's and Verizon's services. The (fake) cost metrics are 10/MB for using Comcast's bandwidth and 20/MB for using Verizon's bandwidth. Please answer following questions with the assumption made in lecture: a border router establishes iBGP sessions to all other routers within its AS.

1. Which one of eBGP, iBGP and IGP distributes externally learned routes internally, and which routers, if any, speak it?

   Solution: iBGP. V1, C1, B1, S1, V2, C2, B2, S2. All routers speak iBGP.

2. Which one of eBGP, iBGP and IGP learn routes to external destinations, and which routers, if any, speak it?

   Solution: eBGP. V1, C1, B1, S1. Only border routers speak eBGP.

3. Which one of eBGP, iBGP and IGP provides internal reachability, and which routers, if any, speak it?

   Solution: IGP. V1, C1, B1, S1, V2, C2, B2, S2. All routers speak IGP

4. Which AS would Berkeley use to reach Stanford, in terms of cost effectiveness?

   Solution: Comcast. It costs less amount of money.

5. Given now Comcast knows Berkeley and Stanford don't get along with each other, it doesn't advertise its route of Berkeley to Stanford, or the other way around. However, Verizon still remains neutral. Which AS would Berkeley use to reach Stanford now?

   Solution: Verizon.

3 IP Fragmentation

Maximum Transmission Unit (MTU) is the size of the largest packet that a link can carry. Host A sends an **600 byte** IP packet (including header) to Host B, which is fragmented along the way. Assume the typical IP header length of 20 bytes.

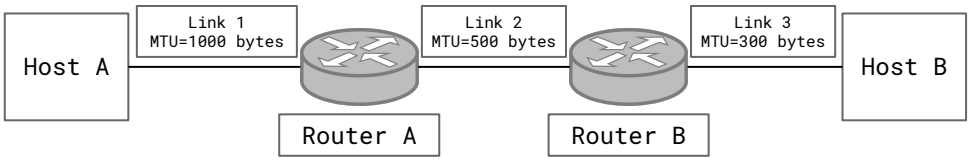

1. The packet fits within the MTU of Link 1 and arrives at Router A. What are the resulting fragments that traverse Link 2? For each fragment, identify the total length (including header), flags, and offset.

   Solution: Original payload = 600 byte total length – 20 byte header = 580 bytes
   F1: total length = (20 byte header + 480 byte payload) = 500 bytes, flags = 001, offset = 0
   F2: total length = (20 byte header + 100 byte payload) = 120 bytes, flags = 000, offset = 480/8 = 60

2. The fragments arrive at Router B. What are the resulting fragments that traverse Link 3?

   Solution: F1a: total length = (20 byte header + 280 byte payload) = 300 bytes, flags = 001, offset = 0
   F1b: total length = (20 byte header + 200 byte payload) = 220 bytes, flags = 001, offset = 280/8 = 35
   F2: total length = (20 byte header + 100 byte payload) = 120 bytes, flags = 000, offset = 480/8 = 60
   Observe that fragmentation offsets in F2 stays the same.

3. Why is the MF flag needed?

   Solution: Packets can arrive out of order; MF flag tells the end host which fragment is the last one.

4. Why can't we just number our fragments instead of keeping track of fragmentation offsets?

   Solution: How else would you order fragments of fragments? (Consider the further fragmentation of F1)

4 Longest Prefix Matching

Your routing table contains the following entries:

Address	Port
01*	Port 1
000	Port 2
001	Port 1
1**	Port 2
101	Port 1

Mark the entries on the following tree. Then, find a more concise representation of the table.

Solution:

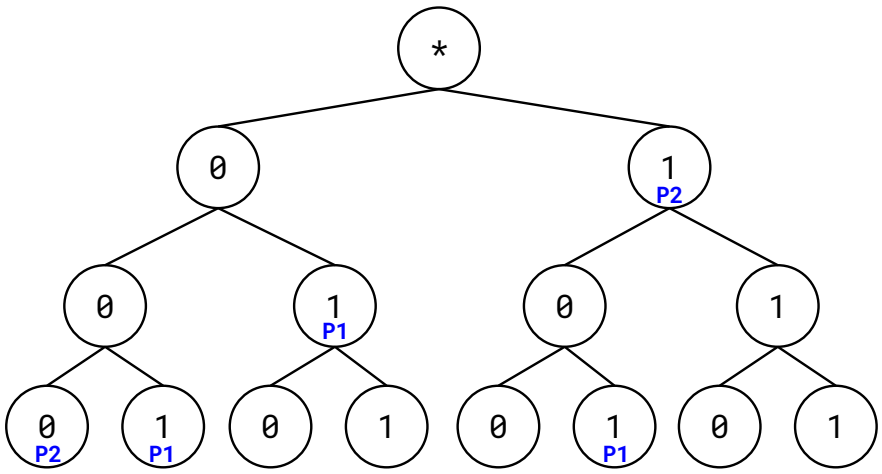

Address	Port
***	Port 1 (or 0** → Port 1)
000	Port 2
1**	Port 2
101	Port 1