

Answers to tutorial 1 ETSF05 and ETSF10

K&R Chapter 5

R7

Policy: Among ASs, policy issues dominate. It may well be important that traffic originating in a given AS not be able to pass through another specific AS. Similarly, a given AS may want to control what transit traffic it carries between other ASs. Within an AS, everything is nominally under the same administrative control and thus policy issues a much less important role in choosing routes within an AS.

Scale: The ability of a routing algorithm and its data structures to scale to handle routing to/ among large numbers of networks is a critical issue in inter-AS routing. Within an AS, scalability is less of a concern. For one thing, if a single administrative domain becomes too large, it is always possible to divide it into two ASs and perform inter-AS routing between the two new ASs.

Performance: Because inter-AS routing is so policy oriented, the quality (for example, performance) of the routes used is often of secondary concern (that is, a longer or more costly route that satisfies certain policy criteria may well be taken over a route that is shorter but does not meet that criteria). Indeed, we saw that among ASs, there is not even the notion of cost (other than AS hop count) associated with routes. Within a single AS, however, such policy concerns are of less importance, allowing routing to focus more on the level of performance realized on a route.

R8

False.

With OSPF, a router broadcasts its link-state information to all other routers in the autonomous system to which it belongs, not just to its neighboring routers. This is because with OSPF, each router needs to construct a complete topological map of the entire AS and then locally runs Dijkstra's shortest-path algorithm to determine its least-cost paths to all other nodes in the same AS.

R11

Routers use the AS-PATH attribute to detect and prevent looping advertisements; they also use it in choosing among multiple paths to the same prefix. The NEXT-HOP attribute indicates the IP address of the first router along an advertised path (outside of the AS receiving the advertisement) to a given prefix. When configuring its forwarding table, a router uses the NEXT-HOP attribute.

R12

A tier-1 ISP B may not carry transit traffic between two other tier-1 ISPs, say A and C, with which B has peering agreements. To implement this policy, ISP B would not advertise to A routes that pass through C; and would not advertise to C routes that pass through A.

R13

False.

A BGP router can choose not to add its own identity to the received path and then send that new path on to all of its neighbors, as BGP is a policy-based routing protocol. This can happen in the following scenario. The destination of the received path is some other AS, instead of the BGP router's AS, and the BGP router does not want to work as a transit router.

P9

No. loops involving three or more nodes will not be detected by the poisoned reverse technique.

P12

$O(N|E|)$, where N is the set of hosts and E is the set of edges or links.

P13

The chosen path is not necessarily the shortest AS-path. Recall that there are many issues to be considered in the route selection process. It is very likely that a longer loop-free path is preferred over a shorter loop-free path due to economic reason. For example, an AS might prefer to send traffic to one neighbor instead of another neighbor with shorter AS distance.

P16

One way for C to force B to hand over all of B's traffic to D on the east coast is for C to only advertise its route to D via its east coast peering point with C.

P18

BitTorrent file sharing and Skype P2P applications.

Consider a BitTorrent file sharing network in which peer 1, 2, and 3 are in stub networks W, X, and Y respectively. Due to the mechanism of BitTorrent's file sharing, it is quite possible that peer 2 gets data chunks from peer 1 and then forwards those data chunks to 3. This is equivalent to B forwarding data that is finally destined to stub network Y.

Answers to Extra questions

1

Diameter of the network is the longest of the shortest paths from every node to every other node. Assuming the hop count as the cost metric, in this network the diameter is 3 hops. Therefore the smallest value that can be used as infinity is 4 hops.

2

Before failure

- A. Net1-2hops-FromB
- B. Net1-1hop-FromB
- C. Net1-2hops-FromB

1 step after failure (Router B updates with infinity hop count to network 1.)

- A. Net1-2hops-FromB
- B. Net1-16hops-FromB
- C. Net1-2hops-FromB

2 steps after failure (A/C thinks: It is too expensive to get to network 1 from B, but C/A offers a cheaper way! Also split horizon prevents A and C from sending announcement with hop count 2 to B, because that is where they got it from.)

- A. Net1-3hops-FromC
- B. Net1-16hops-FromB
- C. Net1-3hops-FromA

3 steps after failure (Now split horizon doesn't prevent A and C from announcing their paths to B because they were not acquired from B. B thinks now: Oh I can now get to network 1 from either A or C! Also A and C increase the hop count metric by 1.)

- A. Net1-4hops-FromC
- B. Net1-4hops-FromA/C
- C. Net1-4hops-FromA

X steps after failure: All hops counts keep increasing by 1 and so on to infinity. Split horizon will not kick in because there is always another path to take than the one the announcement originally came from.

5

Simple net mask matching:

Network id	Net mask	Interface	
10.0.1.0	255.255.255.0	eth0	10.0.1.1 – 10.0.1.254
10.0.0.0	255.255.0.0	eth1	10.0.0.1 – 10.0.255.254
10.1.0.0	255.255.248.0	eth2	10.1.0.1 – 10.1.7.254
10.1.0.0	255.255.254.0	eth3	10.1.0.1 – 10.1.1.254
10.2.0.0	255.255.255.0	eth4	10.2.0.1 – 10.2.0.254
10.3.0.0	255.255.248.0	eth5	10.3.0.1 – 10.3.7.254


After considering overlapping routes and longest match:

Network id	Net mask	Interface	
10.0.1.0	255.255.255.0	eth0	10.0.1.1 – 10.0.1.254
10.0.0.0	255.255.0.0	eth1	10.0.0.1 – 10.0.0.255, 10.0.2.0 – 10.0.255.254
10.1.0.0	255.255.248.0	eth2	10.1.2.0 – 10.1.7.254
10.1.0.0	255.255.254.0	eth3	10.1.0.1 – 10.1.1.254
10.2.0.0	255.255.255.0	eth4	10.2.0.1 – 10.2.0.254
10.3.0.0	255.255.248.0	eth5	10.3.0.1 – 10.3.7.254

6

The address space for interface f1/0 would normally contain 254 usable IP addresses. (Actually 256 but .0 is the network identifier and .255 is the broadcast address). However, since the address space for f1/3 is overlapping with f1/0, and f1/3 has the longest mask, f1/3 is prioritized. Therefore there are 62 usable IP addresses on f1/3 (.0 is the network identifier and .63 is the broadcast address) and $254 - 62 - 1 = 191$ usable IP addresses on f1/0.

Address space for each interface:



7

The routing table at router X after the update should look like the table below. Since we have RIP the hop cost is 1.

Network id	Hops	Next hop
N2	2	C
N3	3	C
N4	5	A
N5	4	C
N6	4	C
N7	2	B

8

The routing table at router 1 after processing the SPF tree should look like the table below.

Network id	Cost	Next hop
NetA	1	-
NetB	7	Router 2
NetC	5	Router 3
NetD	5	Router 2
NetE	7	Router 3
NetF	8	Router 3

9

Net 4 at cost 1, Net 5 at cost 3 and Net 6 at cost 2.