

Communication Networks


Prof. Laurent Vanbever

Exercise 8 – BGP - advanced concepts

8.1 Putting Everything Together (Exam Question 2016)

Consider the ISP network composed of 5 routers (A, B, C, D, E) depicted in the Figure below. Three of these routers, A, E and D , are connected to routers located in neighboring ASes via eBGP. These neighboring routers are indicated by X, Y and Z . Each of them advertises the same three distinct IP prefixes $p1, p2$ and $p3$.

The three tables in the Figure indicate the Local-Preference (LP) associated to each external prefix by A, E and D along with their corresponding AS-PATH length. For instance, A learns a route to $p1$ from X with an AS-PATH length of 10 to which it associates a LP of 200. Internally, the ISP uses an iBGP full-mesh to distribute the BGP routes and OSPF as intra-domain routing protocol. The weight of each internal link is indicated next to it.



An ISP network which receives BGP routes for 3 external prefixes ($p1, p2, p3$) from 3 routers (X, Y, Z) in neighboring ASes.

For each router in the ISP, indicate the router ID of the selected egress (A, E, D) along with the router ID of the internal next-hop (A, B, C, D, E or *direct*) used to reach it. For that you can use the tables on the next page. You can assume that A, E and D use the next-hop-self configuration.

A

prefix egress internal NH

<i>p1</i>		
<i>p2</i>		
<i>p3</i>		

B

prefix egress internal NH

<i>p1</i>		
<i>p2</i>		
<i>p3</i>		

C

prefix egress internal NH

<i>p1</i>		
<i>p2</i>		
<i>p3</i>		

D

prefix egress internal NH

<i>p1</i>		
<i>p2</i>		
<i>p3</i>		

E

prefix egress internal NH

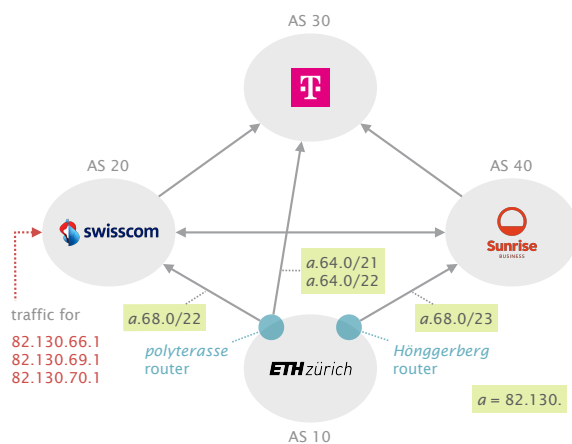
<i>p1</i>		
<i>p2</i>		
<i>p3</i>		

Fill in the following tables with the selected egress and internal next-hop.

8.2 Traffic (not so much) Engineered

After passing the Communication Networks exam with flying colors, ETH hires you as a junior network engineer. *Congrats!*

Your first mission is to analyze their BGP configuration. They indeed suspect that something might be wrong, especially since they installed this box from Cisco Systems that automatically configure BGP announcements according to Traffic Engineering objectives. For the sake of simplicity, assume again that ETH has only one prefix: 82.130.64.0/21 and three providers: Swisscom, Deutsche Telekom and Sunrise. The actual announcements are depicted on the left. Customers are drawn below their providers (Swisscom is a customer of Deutsche Telekom), while peers are drawn next to each other (Swisscom is a peer of Sunrise).



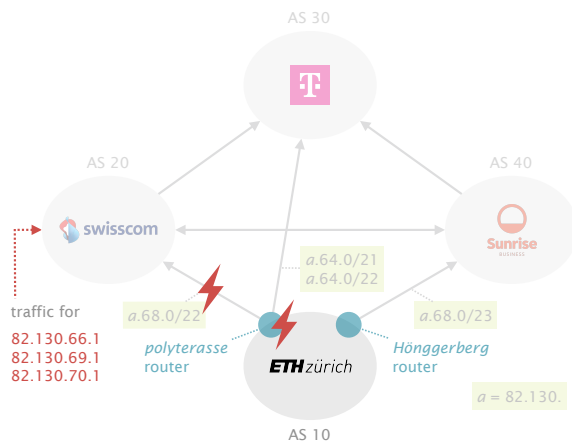
Where are my packets going?

Consider the incoming traffic from Swisscom. What path is taken for packets destined to:

- 82.130.66.1?
- 82.130.69.1?
- 82.130.70.1?

Are Swisscom, Deutsche Telekom and Sunrise happy about these announcements? Can they do anything about that? Explain briefly.

As the *polyterasse* router is getting older, its reliability starts to suffer. In theory, this should not be a big problem as ETH is *triple*-homed! Yet, in practice, the ETH engineers observe regular network connectivity upon failures.



Is this network as redundant as it looks?

Assuming the same announcements, what path ends up being taken by the packets destined to the above three IP addresses when:

- the link between *polyterasse* and Swisscom goes down?
 - 82.130.66.1?
 - 82.130.69.1?
 - 82.130.70.1?
- the *polyterasse* router dies?
 - 82.130.66.1?
 - 82.130.69.1?
 - 82.130.70.1?


What would you change in the ETH announcements to improve reliability, without disturbing the inbound Traffic Engineering performed by the Cisco box in the steady case (without failures)? Explain briefly.

8.3 BGP Hijack (Exam Question 2018)

Consider the Internet topology consisting of 11 Autonomous Systems (ASes) in the Figure below. Single-headed plain arrows point from providers to their customers (AS *A* is the provider of AS *E*), while double-headed dashed arrows connect peers (AS *A* and AS *B* are peers). Each AS is made up of a single BGP router and applies the default selection and exportation BGP policies based on their customers, peers and providers.

In this task, the routers break ties using the AS number of the neighbor: in case multiple routes are equally good, the router selects the route of the neighbor with the lowest AS number (alphabetical order).

AS *G* is the origin of prefix 20.0.0.0/22 and advertises it to its neighbors. Independently of what the external advertisements are, AS *G* **always** prefers its internal route to reach any IP destination in 20.0.0.0/22.



An Internet topology with 11 ASes. AS *K* aims at hijacking traffic destined to AS *G*.

- AS *K* wants to hijack all the traffic going to AS *G* for 20.0.0.0/22. It starts advertising the exact same prefix. From which ASes is it able to attract the traffic?

b) AS *K* is not satisfied by the result. What can it do to attract traffic destined to AS *G* from more of the ASes? List the ASes from which it is able to attract the traffic and explain why this works.

c) The ASes from which AS *K* manages to attract the traffic realize what is happening as all their traffic to 20.0.0.0/22 goes to a dead-end (AS *K*).

Show how AS *K* could still deliver the traffic to the real destination (AS *G*) by poisoning the AS path while attracting as much traffic as possible. In addition, list the ASes from which it can attract the traffic.

Poisoning the AS path means that AS *K* would put specific AS numbers in the AS path of the hijacked prefix in order to abuse the BGP loop prevention mechanism. This way, specific ASes will drop the hijacked prefix instead of forwarding it.

d) Can you think of a different way for AS *K* to achieve similar results as in *c*) without poisoning the AS path? Explain.