ETH zürich

Networked Systems
ETH Zürich — seit 2015

Communication Networks

Prof. Laurent Vanbever

Solution: Exercise 8 – Border Gateway Protocol (BGP)

8.1 Traffic Engineering

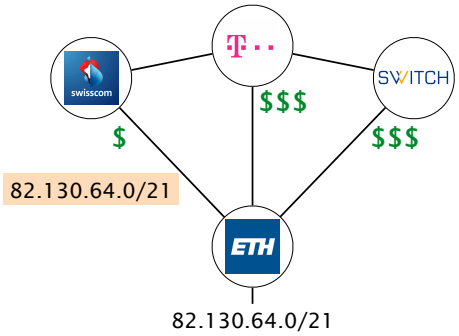

ETH is connected to three providers with different costs.

Assume that ETH has only one prefix: 82.130.64.0/21. As depicted on the left, the ETH network is connected to three providers (Swisscom, Deutsche Telekom and Switch) and the providers are interconnected with each other. The contract with Swisscom is the cheapest one (indicated by the dollar symbols). For this reason, ETH wants to receive all the incoming traffic over the Swisscom link and therefore announces its prefix only to Swisscom.

a) Do you think that is a good configuration? What happens if the link between ETH and Swisscom fails?

Solution:
Not a good solution. If the link fails, ETH will no longer receive any traffic. ETH is no longer reachable from other networks.

b) To improve the connectivity in case of a link failure between ETH and Swisscom, ETH wants to optimize its announcements. Write down the prefixes which ETH announces to Swisscom, Deutsche Telekom and Switch. During normal operation (no link failure) ETH should still receive all incoming traffic over the Swisscom link.

Solution:
To Swisscom: 82.130.64.0/22 and 82.130.68.0/22 (other splits are also possible)
To Deutsche Telekom: 82.130.64.0/21
To Switch: 82.130.64.0/21

c) After further investigations, ETH decides that only traffic towards 82.130.68.0/23 has to be received over the Swisscom link. All the other traffic can enter over any of the providers. Which prefixes do you have to announce to achieve this traffic distribution?

Solution:
To Swisscom: 82.130.68.0/23 and 82.130.64.0/21
To Deutsche Telekom: 82.130.64.0/21
To Switch: 82.130.64.0/21

8.2 Not-so-reliable Internet

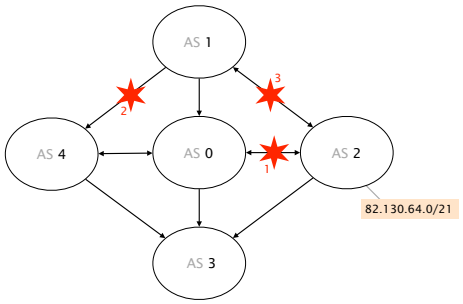

Which messages are exchanged?

Consider now the same BGP network composed of 5 ASes but assuming customer-provider and peer-to-peer policies. Providers are connected to their customers with a single-headed arrow pointing to their customers (AS 1 is the provider of AS 4), while peers are connected with double-headed arrows (AS 1 and AS 2 are peers).

Assume that AS 2 is the only one to advertise an IPv4 prefix: 82.130.64.0/21 (to *all* its neighbors) and that the Internet has converged. Which BGP messages are exchanged after the following events happen, one after the other:

a) the link between AS 0 and AS 2 fails (event 1)

Solution:

(i) AS 0 sends a WITHDRAW for 82.130.64.0/21 to AS 3 (optional);

(ii) AS 0 sends an UPDATE for 82.130.64.0/21 to AS 3 with AS-PATH [0,1,2].

b) the link between AS 1 and AS 4 fails (event 2)

Solution:

AS 4 sends a WITHDRAW for 82.130.64.0/21 to AS 3.

c) the link between AS 1 and AS 2 fails (event 3)

Solution:

(i) AS 1 sends a WITHDRAW for 82.130.64.0/21 to AS 0;

(ii) AS 0 sends a WITHDRAW for 82.130.64.0/21 to AS 3;

Is the network still connected at the end? If not, list the ASes that cannot reach the prefix anymore.

Solution:

No. The BGP network is not connected anymore. Only AS 3 is able to reach 82.130.64.0/21 via its direct link with AS 2. Observe that the physical graph *is* still connected yet as BGP policies prevent paths to be used, blackholes appear nonetheless.

8.3 BGP and IGP: Very creative! (Exam Question 2020)

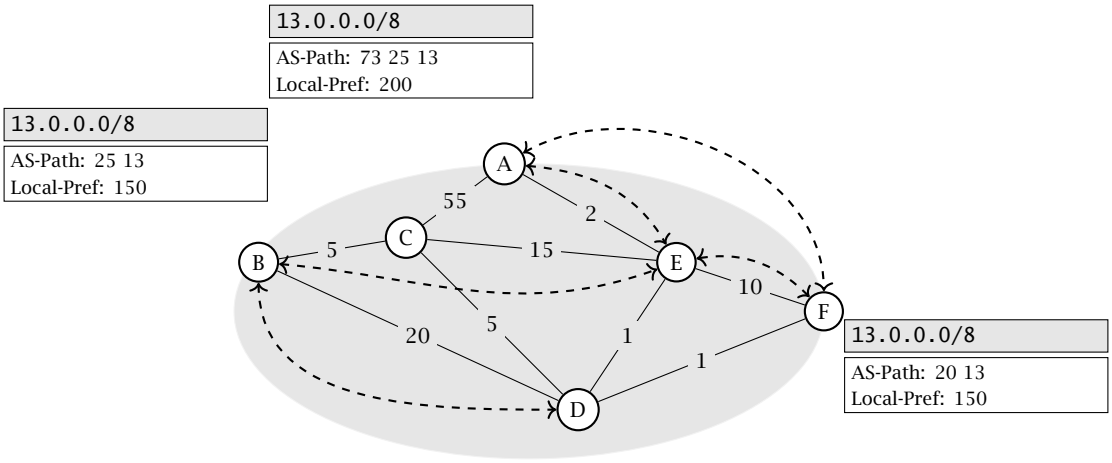

A simple BGP network **not** forming an iBGP full-mesh.

Consider the AS above with three border routers (A, B, F) and three internal routers (C, D, E). All three border routers receive a route announcement for the prefix `13.0.0.0/8` from their eBGP neighbors (not depicted), which they distribute internally. The iBGP sessions are depicted by double-headed dashed arrows (e.g., router A and F maintain an iBGP session). All routers follow the standard BGP decision process. The three border routers have `next-hop-self` configured on all iBGP sessions.

a) For every router, list *(i)* the BGP next-hop, *(ii)* the path taken by the traffic and *(iii)* indicate whether the router's traffic can actually reach the destination. If the next-hop is external, put EXT. If there is no next-hop, put NO.

Solution:

Router	BGP next-hop	Path taken by the traffic	Reachable?
A	EXT	A → EXT	Yes
B	EXT	B → EXT	Yes
C	NO	C → ∅	No
D	B	D → C → ∅	No
E	A	E → A → EXT	Yes
F	A	F → D → C → ∅	No

b) Assume the eBGP session of router A fails and consequently, **the external route of A is not available anymore**. List for every router *(i)* the BGP next-hop, *(ii)* the path taken by the traffic and *(iii)* indicate whether the router's traffic can reach the destination. If the next-hop is external, put EXT. If there is no next-hop, put NO.

Solution:

Router	BGP next-hop	Path taken by the traffic	Reachable?
A	F	A → E → D → C → ∅	No
B	EXT	B → EXT	Yes
C	NO	C → ∅	No
D	B	D → C → ∅	No
E	F	E → D → C → ∅	No
F	EXT	F → EXT	Yes

c) The network operator reacted and **added a new iBGP session between routers B and C**. The failure still persists, i.e., the external route of A is not available. List for every router *(i)* the BGP next-hop, *(ii)* the path taken by the traffic and *(iii)* indicate whether the router's traffic can reach the destination. If the next-hop is external, put EXT. If there is no next-hop, put NO.

Solution:

Router	BGP next-hop	Path taken by the traffic	Reachable?
A	F	A → E → D → C → B → EXT	Yes
B	EXT	B → EXT	Yes
C	B	C → B → EXT	Yes
D	B	D → C → B → EXT	Yes
E	F	E → D → C → B → EXT	Yes
F	EXT	F → EXT	Yes