

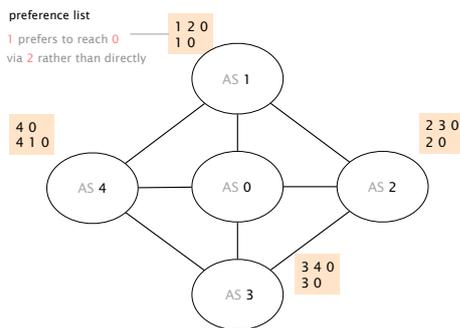
Communication Networks

Prof. Laurent Vanbever

Solution: Exercises week 9 – Routing Policies and Security

Convergence

Consider this BGP network composed of 5 ASes. Each AS is assigned with a list of paths which indicates its preferences to reach AS 0.



Does this network ever converge?

Considering that only AS 0 originates prefixes, does that BGP network have a unique, stable solution?

- If yes, indicate the path that each AS selects in the stable solution.
- If not, describe an example of oscillation. For instance, by describing a sequence of messages that repeats itself.

Solution: This BGP network does have a unique, stable solution in which:

- AS 1 selects $[1, 2, 0]$ (preferred path);
- AS 2 selects $[2, 0]$;
- AS 3 selects $[3, 4, 0]$ (preferred path);
- AS 4 selects $[4, 0]$ (preferred path).

Not-so-reliable Internet

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].

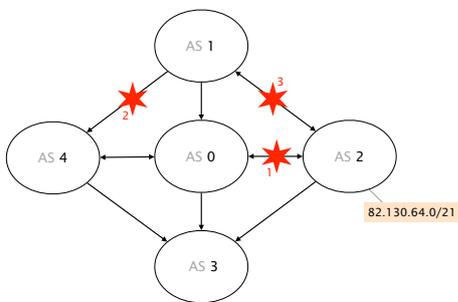
- b) the link between AS 1 and AS 4 fails (event 2):

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;



Which messages are exchanged?

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.

Peering war

Consider two ASes, say ATT and Google, which are peering with each other in multiple locations (using a peer-peer relationship). Oftentimes, peering agreements include a clause that mandate peers to announce BGP routes with the same AS-PATH length at *every* peering location. Explain why this clause makes sense.

Solution: This clause prevents ATT and Google from tweaking their advertisements so as to use the other network to transit their traffic across long distances on their behalf.

As an illustration, consider that ATT and Google peer in New York (NYC) and San Francisco (SFO) and that one of Google's prefixes, say 8.8.8.0/24, is connected directly to one of its NYC routers (draw a picture). Without the clause, Google could advertise 8.8.8.0/24 to ATT with a short AS-PATH in NYC and a relatively longer one in SFO. As a consequence, any traffic for 8.8.8.0/24 entering ATT via its SFO router would be transited *internally* to ATT's NYC router instead of being delivered directly to the Google router in SFO.