


## Communication Networks

Prof. Laurent Vanbever

### Exercise 7 – Border Gateway Protocol (BGP)

## BGP


### 7.1 Circularity



Chicken or the egg?

Consider two directly-connected routers A and B that have an eBGP session (running over TCP) between them. Explain how BGP, which is used to propagate *IP-based* routing information, can rely on establishing a TCP connection to exchange routes, which itself relies on IP. How is this circularity resolved? How do the two TCP endpoints manage to reach each other?

### 7.2 Not-so-reliable Internet



Which messages are exchanged?

Consider the BGP network composed of 5 ASes shown on the left which uses the normal 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:

- the link between AS 0 and AS 2 fails (event 1)
- the link between AS 1 and AS 4 fails (event 2)
- the link between AS 1 and AS 2 fails (event 3)

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


### 7.3 Visibility (Exam Question 2016)

Consider now the network depicted on the left. Single-headed plain arrows point from providers to their customers (AS A is the provider of AS D), while double-headed dashed arrows connect peers (AS D and AS E are peers). Each AS in the network originates a unique prefix that it advertises to all its BGP neighbors. Each AS also applies the default selection and exportation BGP policies based on their customers, peers and providers.



A simple BGP network

- What path (sequence of ASes) is followed when AS G sends packets destined to the prefix originated by AS E?
- What path (sequence of ASes) is followed when AS F sends packets destined to the prefix originated by AS G?
- Suppose AS A and AS C give you a “dump” of all the BGP routes they *learn* for every destination. You then extract all links from the AS paths seen in those “dumps” and use them to construct a view of the AS-level topology. Draw the resulting AS-level topology in the figure below.




Draw the inferred AS-level topology.

- Give the minimum set of ASes that must provide a “dump” of each route they learn s.t. all the edges (the ones in the figure on the left) are visible? Justify your answer.

### 7.4 Left? Right? Both? (Exam Question 2017)

Consider the BGP network composed of 4 routers depicted in Figure below. Two of these routers, R1 and R4 are egress routers and maintain eBGP sessions with external neighbors. R1 is configured to associate a local-preference of 100 to externally-learned routes, while R4 is configured to associate a local-preference of 200 to externally-learned routes. R2 and R3 are internal routers. All four routers are connected in an iBGP full-mesh. OSPF is used as intra-domain routing protocol. The link weights are indicated in the figure, e.g. the (R1, R2) link is configured with a weight of 20. The Figure also indicates the propagation delay for each link (e.g., it takes 5ms for a packet to propagate between R1 and R2).



A simple BGP network learning external routes via eBGP on R1 and R4.

- a) Considering the above configuration, indicate the next-hop used by each router in the steady state, *i.e.*, once the network has fully converged. Use the keyword “external” to indicate that an edge router is forwarding outside of the domain. Note that we are not looking for the *BGP* next-hop but rather the next-hop a packet would take when being forwarded.
- b) One of the network operator decides to lower the local-preference associated by R4 to externally-learned routes to 50 (instead of the original 200). Indicate the sequence of BGP messages sent which is triggered following that change along with the timestamps at which they are generated. You can consider that the BGP process on each router is infinitely fast meaning only propagation delay matters. Only indicate when messages are sent, not when messages are received.

Use this template to answer (replace the content within the square brackets):

Timestamp [YY ms] [RX] sends the message [msg\_content] to [RA, RB, and RC]

- c) Was a forwarding loop induced due to the configuration change? Briefly explain why or why not. If a loop was created, also indicate its duration (in ms).
- d) It turns out that the network operator changed her mind. This time, she configures R4 to associate a local-preference of 100 to externally-learned routes (*i.e.* the same local-preference value as on R1). Indicate the next-hop used by each router in the steady state (once the network has fully converged). Again use the keyword “external” to indicate that an egress router is forwarding outside of the domain.
- e) Soon after the network has fully converged due to the configuration change of R4, a failure happens disconnecting R4 from all its external neighbors. The connection between R4 and R3 is still working fine though. Indicate the sequence of BGP messages sent following that failure along with the timestamps at which they are generated. Only indicate when messages are sent, not when messages are received.

Use this template to answer (replace the content within the square brackets):

Timestamp [YY ms] [RX] sends the message [msg\_content] to [RA, RB, and RC]

- f) Was a forwarding loop induced due to the failure? Briefly explain why or why not. If a loop was created, also indicate its duration (in ms).