

## Communication Networks

Prof. Laurent Vanbever

### Exercise 6 – Border Gateway Protocol (BGP)


#### 6.1 BGP Sessions

Before BGP can receive and also advertise any routes, BGP sessions between neighboring ASes need to be set up.

- Explain how eBGP sessions can still be established using IP addresses even though without BGP there is no IP routing between ASes yet.
- Why is it important for BGP sessions to have a keepalive message? Think about what happens if a BGP router freezes and stops sending BGP session messages.
- BGP session packets are sent over the same links as the rest of the traffic. Can you think of possible ways this can be used to disrupt BGP sessions and can you think of ways to mitigate them?
- The BGP process will sometimes wait to send route updates until either a certain number of updates have happened or a certain time threshold was reached. Can you think of a reason why the BGP process wouldn't just send the update as fast as possible?

#### 6.2 Route Propagation

Each AS might receive multiple advertisements for the same prefix. In this task you will investigate how they propagate in the network and which ASes will see which advertisements. Assume that business relationships are followed, where the arrows point from provider to customer. Peers are shown with an arrows pointing in both directions.




- Assume that every link has the same delay of 1ms and that computation is negligible. AS 1 sends out an advertisement for its prefix 1.0.0.0/8 out at time 0 to all of its neighbors concurrently. Write down all the advertisements that AS 3 receives by writing down their AS path.
- Some advertisements are not propagated to AS 3 since only the best known route will be further propagated. Which additional AS paths can be seen at AS3, if AS 1 only advertises the prefix on the link 1-5.

### 6.3 BGP Decision Process

BGP elects only one single route for each prefix using a specific priority of metrics for the decision. Please sort the routes (shown in the following table) by their precedence.

MED	AS-PATH length	egress IP	eBGP?	IGP metric	Local Pref	Order
50	5	1.5.3.2	no	8	100	
20	5	1.8.4.2	no	10	200	
10	1	1.9.1.5	yes	-	50	
50	5	3.8.4.1	no	4	100	
50	8	2.6.7.8	no	1	200	
50	3	1.2.3.8	no	4	50	
80	2	2.3.8.4	yes	-	100	
10	1	2.6.7.8	no	2	200	
50	8	1.4.4.3	no	1	200	
10	5	6.5.1.9	no	5	100	
10	1	1.5.2.2	yes	-	200	

### 6.4 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.