

Communication Networks

Prof. Laurent Vanbever

Solution: 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.

- a) Explain how eBGP sessions can still be established using IP addresses even though without BGP there is no IP routing between ASes yet.

Solution: IP routing is not necessary in L2 networks to reach other IPs that are within the same subnet (using ARP). If the session is not established within a L2 network, the other option would be to use static routes to reach the neighboring BGP router.

- b) 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.

Solution: If a router freezes and stops sending updates, both reachability and convergence could be impacted, since important route updates are not distributed properly anymore. At the same time all the received routes are still being used. In the case the router crashed all packets sent to that router will be lost. This is why it is important to detect such problems as fast as possible using keepalive messages.

- c) 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?

Solution:

- (i) A malicious actor could try to spoof the IP address of the neighboring BGP router and try to inject malicious updates.

Mitigation: Installing firewall rules to block packets using any of the IP addresses used in the BGP sessions from all other network interfaces.

- (ii) Malicious actors could overwhelm the link by sending a lot of traffic bringing the session to a halt if keepalives are dropped.


Mitigation: Prioritize BGP session traffic over normal traffic so that control messages are still delivered even if the link is congested.

- d) 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?

Solution: The main reason to delay updates is to reduce the computational load on the devices. By batching multiple updates this load can be reduced and computations can be optimized. Something like route flapping (fast ongoing changes between a route being valid and invalid) for example could put enormous strain on the control plane by recomputing the same state over and over.

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.



- a) 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.

Solution: The following AS paths arrive at AS3:

1,
 1->5,
 1->2->4,
 1->5->6,
 1->2->4->8

- b) 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.

Solution: The following additional AS path will show up at AS3: 1->5->6->8

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	9
20	5	1.8.4.2	no	10	200	3
10	1	1.9.1.5	yes	-	50	10
50	5	3.8.4.1	no	4	100	8
50	8	2.6.7.8	no	1	200	5
50	3	1.2.3.8	no	4	50	11
80	2	2.3.8.4	yes	-	100	6
10	1	2.6.7.8	no	2	200	2
50	8	1.4.4.3	no	1	200	4
10	5	6.5.1.9	no	5	100	7
10	1	1.5.2.2	yes	-	200	1

6.4 Not-so-reliable Internet

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:

- 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 PATH 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;

- d) 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.



Which messages are exchanged?