



BGP in our network: How does the internet GPS work

APRIL 3, 2018

When you are viewing your favorite internet website, obviously, you are not directly connected to the server of the website. Technically speaking, packets have to travel through a bunch of routers to reach their destination. But how do the packets

know which route to take when the **global routing table** consists of approximately 660 thousand routes to date? And how do they know which is the best path if there is more than just one exit?

The magic behind each packet being sent and delivered on the internet is nothing else but routing on the layer L3 – the so called **network layer**. It is the routing protocols which enable the communication between different networks on the TCP/IP (internet) to work. Each network has its own autonomous system (AS) number assigned. Routing protocols are designed to exchange routing information within your AS as well as between other AS'es.

The most popular routing protocol is **Border Gateway Protocol (BGP)**, which is common in the internet service provider world. Without BGP, internet communication would be impossible or at least less flexible. Let's take a closer look at BGP and find out if it is a good fit for your network needs.


What is BGP

BGP is designed to route packets both, throughout the internet (eBGP) as well as within your internal network (iBGP) – within a single AS. It determines the **optimal routing path** to send/receive packets based on a set of rules and attributes. BGP configuration is represented by mechanisms determining a degree of preference of routes – supporting **Classless Interdomain Routing (CIDR)**. A BGP router can receive more copies of routes (multiple advertisements) from multiple sources. The BGP **best path selection algorithm** gives the router an instruction about which path to prefer based on your BGP settings.

Do I need BGP?


There are a few network scenarios when the implementation of BGP is highly recommended. But before we go on with this, you have to first understand the difference between default route and full routing table.

Let's assume you run your own network AS1 with a few end users and you use a service provider – ISP. Because you only have one ISP – your network is single-homed, there is only one exit path, therefore you advertise the default route using BGP. With **default route**, you are not able to see what networks are passed on the way from source to destination in the route table (because you do not see path). Your ISP will send default route toward your AS1 network.



Full routing table comes into play in a scenario when AS1 is multi-homed – uses two different ISPs – ISP1 and ISP2 to [ensure redundancy and network optimization](#).

BGP network – full routing table



AS1 announces its prefixes to ISP1 and ISP2 via BGP – AS1 is a multi-homed BGP network. AS1 now receives full BGP routing table and sees all AS'es on the path. This way networks don't have to be connected to each other in order to learn their prefixes. But how does BGP know if a packet should take the path through ISP1 or ISP2? BGP determines the best path based on a set of attributes and filters – the so called BGP best path selection.


BGP best path selection

As I mentioned in the previous section, a BGP router receives **multiple copies of routes from multiple providers**. Therefore there is an algorithm for comparing the BGP routing tables and selecting the most efficient path which will go into the IP routing table on the local BGP speaker. Choosing one path over another has direct influence on network operational costs and the overall quality of service (delay, packet loss, etc.).


We don't want you to feel overwhelmed by describing the functionality and role of each and every BGP attribute. Let's just take a look at three examples of BGP use.

Example #1 – AS path prepending:

AS path is probably the most commonly used BGP attribute. Imagine a network (AS1) be a multi-homed network which uses two internet exchange points to exchange its traffic with other AS'es. AS1 announces its prefixes to both exchanges – IXP1 and IXP2. Behind both IXP1 and IXP2 are AS2, AS3 and AS4. AS paths from AS2, AS3 and AS4 toward AS1 are of the same length. AS2 and AS3 send their traffic through IXP1 and AS4 through IXP2.




Let's say that AS1 changes its peering policy and prefers to receive prefixes from AS2 and AS3 also via IXP2 instead of IXP1. It prepends its AS path towards IXP1 so that it looks longer.




All the traffic from networks connected to both internet exchanges now goes through IXP2 if no other filters and policies have been set.

Example #2 – Local Preference:

AS1 is a multi-homed customer of a service provider AS2 connected with one 10GE link and one 100GE link.



AS1 wants to use the lower capacity 10GE link as a backup and send all traffic through the 100GE link. AS2 allows AS1 to use Local Preference to influence their routing so that AS1 can send all its announced routes through the higher bandwidth link by setting a higher local preference of 200 on the interface toward R2.



Example #3 – AS path access lists (another Local Preference case):

Let's assume we would like to optimize routing to Swisscom and prefer routes via operator 1 over operator 2. AS path access lists are the best tool to do so. An AS path access list is a list based on AS path information and rules learned by a router. Every new route update is filtered through an access list and then based on filters either permitted or denied. Routes can then be assigned a BGP attribute, e.g. higher/lower Local Preference, Weight, etc.

In our example, we use access lists to accept everything with last hop being Swisscom AS and set the route via operator 1 a higher Local Preference.