Secure and efficient BGP routing

The Border Gateway Protocol (BGP) is a fundamental feature of the internet, helping to stitch together the plethora of networks that make up the web and carry the vast quantities of traffic that cross the globe each day.

As a crucial piece of the internet, it is imperative BGP functions smoothly and avoids vulnerabilities, disruptions and security issues.

The internet relies on network operators doing the right thing to ensure the correct information goes to the correct parties. However, the provision of robust, foolproof BGP route filtering presents a challenge for many organizations.

The BGP protocol works very well. However, it has been around since long before today's specific security issues and works in coordination with networks scattered across the globe, leaving it potentially open to hijacks and leaks. The Internet Society estimated that there were more than 5,000 route leaks and hijacks in 2017.

The information we have information gathered on this phenomenon has shown some patterns in the US. Most of these leaks seem to happen in the middle of the week, around Tuesday, but there is also a peak on Fridays, just before the weekend starts. There is a need to be extra-vigilant around these times.

BGP route leaks involve the accidental misconfiguration or illegitimate advertisement of prefixes, or blocks of IP addresses, that propagate across networks and result in suboptimal routing or hijacking of traffic.

It's well worth implementing filters to counteract the issues this can cause because these types of leaks have continued year in, year out over the past decade.

There are several such mechanisms that can help protect against such leaks, and we are well-prepared to help you.

Contents

Peerlock "Lite"

One way of filtering uses a method we refer to as "Peerlock Lite" to reject prefixes which have passed through Tier 1 networks received from customers or peers.

For example, we are only reachable via settlement-free peering. Therefore, any routes to our network you receive from a customer or peer are leaks. By using protection mechanisms deployed at private peering or internet exchanges to deny these Tier 1 networks, it is possible to block out many potential problems before they even occur.

The drawback to this method is that it requires a static list of autonomous system numbers (ASNs), which identify networks. Their static nature means if the function of one of these networks is modified or it stops being transit-free, the list needs updating. This makes it crucial to ensure that a half-yearly or yearly review is carried out to keep the list accurate. It also requires an inferred or explicit knowledge of the locked ASNs' transit relationships. Nonetheless, this method is one of the most effective ways to snuff out route leaks.

BGP communities

Another method for blocking leaks is to ensure that prefixes received from settlement-free peering partners are never announced to other such peering partners. One way of doing this is by tagging routes with BGP communities, or labels for routes that share a common property. Filters can then be set so that prefixes without the appropriate communities are rejected upon egress from a border router.

Furthermore, if there is no community associated with particular routes, it would be advisable for these not to be announced to another party. This ensures that, if those prefixes somehow get into your network, it will never propagate them. In this way, the use of BGP communities can be a key tool in preventing route leaks.

It is possible to **block out many potential problems** before they even occur by using protection mechanisms deployed at private peering or internet exchanges.

Among the best-known BGP communities are "no-export" and "no-advertise." The first of these is associated with routes not to be advertised beyond the company's own ASN and the second with those not to be advertised beyond the receiving router. It is important to understand the behavior of these communities before using them, thus helping ensure that the required level of availability is maintained on routes.

Communities can be specified for categories including where routes were learned from, such as a transit customer or peering partner, or for locations such as Europe or a city. However, the range of possibilities for BGP communities is wide. This flexibility provides significant scope for harnessing BGP communities.

As an example of the use of such communities, we have one that can be applied for suppressing announcements to the carrier's peering partners. A customer can use this if, for example, it wants to steer traffic away from a peer that has congestion in their network or is suffering an outage, and towards another one. An alternative option allows traffic to be steered away, but route announcements to the peer to be left as a backup of last resort in case of hitches with connections to other peers too.

These options can be applied to all peers or just selected ones, with the aim of giving maximum flexibility for customers to determine how route announcements are handled in a way that best fits their business needs.

Meanwhile, we offer broader communities for regionally based choices and community-triggered blackholing as well. This includes selective and regional blackholing, which provide tools for customers with even more granular options.

Whitelists

Another approach is to apply a so-called "whitelist" of prefixes that a customer can announce to every customer-facing external BGP (eBGP) session, making operations more secure.

This is a method that we employ for all such sessions by using data from internet routing registries (IRRs). Indeed, the company uses a unique whitelist for each customer, dramatically reducing the chances and extent of damage, and giving it tight control over the routes the company accepts from customers.

As well as our mechanisms to deal with this, there are a number of opensource tools that can come in helpful for applying prefix filters and can be converted into a format suitable for the particular router platform, such as BGPQ3.

We use a **unique** whitelist for **each customer**, dramatically reducing the chances and extent of damage.

Maximum prefix limits

Yet another method for preventing route leaks is for maximum prefix limits to be applied. For example, a limit of 1,000 routes can be applied for an eBGP session so that the session is automatically closed if that number is exceeded.

These prefix limits provide a key safety measure for helping the network respond in a way that causes minimal harm to the global routing system and guards against leaks, providing protection to routers and networks. They can act as a highly effective way of protecting the network if there is in fact a route leak, because it prevents this from being propagated.

Maximum prefix limits can be applied either pre- or post-policy, though the maximum effect can be gained from doing this pre-policy to help avert any significant issues before they happen rather than risking some leaked prefixes being allowed through. Nevertheless, prefix filtering policies vary by routing platform, with some only enabling this to be done post-policy.

Our peerlocking mechanism can significantly **reduce the impact** and proliferation of route leaks.

Peerlocking

We have successfully deployed a more comprehensive form of peerlocking. Using this can massively reduce the risk on a global scale of prefixes being accepted via unauthorized routes.

The essence of this approach is to putting human insight into the network. In basic terms, it relies on peering partners telling us which networks, if any, are authorized transit providers, with the partners that provide this information known as “protected ASNs.” Routes can then be “locked out” if they come from unauthorized transit providers.

We recommend you inform networks you are trying to protect so they can agree to these filters being deployed and avoid unexpected surprises down the line. Partners always need to be aware of what is happening with the network, and engagement is central to this.

It is also essential for such peerlock filters to be applied to every eBGP session, whether customer-facing or peering ones, to ensure that full use is made of this key protection mechanism.

In a nutshell, our peerlocking offers a highly efficient method for halting route leaks, and we have seen significant improvements for networks that agreed to become protected ASNs.

Flexibility

We offer regional expectations, giving global settlement-free peers flexibility if they run operations differently in different continents. We also have a manual that we generate for each peer that it enables locking for. This is useful in setting out the documentation on the ins and outs of the technology and how it functions, as well as for a company to retain this knowledge as employees come and go.

All in all, the our peerlocking mechanism can therefore significantly reduce the impact and proliferation of route leaks, helping through active monitoring of the default-free zone.

One key to our success in the deployment of these technologies is our industry-leading GIN Unified Management System (GUMS) SDN controller. The use of GUMS allows us to deploy changes to whitelists, communities, peerlocking, and BGP policies in general in a programmatic fashion. This leads to consistently deployed configurations and much lower error rates in configuration.

Operators make changes in the GUMS Web UI and deploy their changes from the GUMS server rather than logging into routers and making changes manually. The ability to do this makes the process more efficient, improving the effectiveness of the system as a whole.


NTT