INTERNET PLUMBING FOR SECURITY PROFESSIONALS: THE STATE OF BGP SECURITY

Wim Remes – Rapid7
Blackhat USA 2015 – Las Vegas

Introduction

The Border Gateway Protocol, commonly known as BGP, is a fundamental protocol in the functioning of the internet as we know it today. In essence, BGP allows Autonomous Systems (ASNs) to inform eachother of the networks they know a route to. This allows those ASNs to build a resilient network that does not have a single point of failure for most IP prefixes, given that enough peering relationships between ASNs are established.

Together with technologies such as DNS, SSL/TLS, and HTTP, BGP forms the foundational structure of the internet and, as such, supports many critical services globally. Several events over the past few years, that will be illustrated later in this paper, have shown how it is fairly easy to impact large parts of the internet. While several initiatives have been taken to improve security of the BGP infrastructure, they are plagued by low adoption rates and a general misunderstanding of their impact. The ultimate goal of this whitepaper and its accompanying presentation is to make the BGP problem space understood by a wider audience and to drive adoption of the available counter-measures in order to improve the security of the foundational internet infrastructure.

A few key things are necessary to be understood how BGP works and where its weaknesses are.

BGP, for the purpose of this paper, is looked at in its Exterior Border Gateway Protocol (eBGP) implementation. That is where it is configured to communicate with BGP peers outside its own **autonomous system**.

An **autonomous system (ASN)** is a collection of connected IP routing prefixes under the control of one or more network operators on behalf of a single administrative entity. An ASN is represented by its Autonomous System Number, as assigned by IANA or its authorized RIRs (Regional Internet Registries).

A **BGP Peer** is a BGP router that uses BGP to communicate routing information about its own IP prefixes or IP prefixes it knows a route to, to its peers.

Peers, in general, are BGP routers that have been configured by their owners to exchange routing information with each other. The owners of the BGP routers, or their corresponding ASNs, establish a trust relationship with eachother.

Lastly, BGP has always been regarded as a subject for networking experts. While there are definitely things that can go horribly wrong if the BGP configuration of one or more major ASNs is impacted, BGP in itself is a fairly simple protocol. As an example, unlike other routing protocols, BGP does not take bandwidth or other link parameters (apart from path length) into account in the messages that are exchanged between peers.

The goal of this paper and its accompanying presentation is to provide network professionals with information on how they can leverage BGP data for security intelligence, what tools are available to secure the BGP backbone of the internet, and how they can help to ultimately make the internet more secure.

The BGP Threat Model

BGP, being one of the older protocols still in use in the core infrastructure of the internet, could be considered to be intentionally brittle. Its lack of default security features provides its users with often much-needed flexibility. BGP routers provide routing updates to their BGP peers which, in turn, implement routing policies that are intended to control the propagation of the received routes to other peers. The lack of strict routing policies are often the key reason why BGP hijacking, intentional or not, is possible.

Specifically for BGP, we include the following potential attacks in our threat model.

One or more routes to prefixes that the ASN is not authorized for is provided through a routing update to its upstream peers. This results in the traffic for those prefixes to be routed to the announcing ASN if one of these conditions is met:

- The ASN announces the shortest path for the prefix
- The announced prefix is more specific than the prefix announced by the ASN authorized for the prefix

> Note that pre-pending is an often-used technique to influence BGP routing decisions. Pre-pending is the practice of adding the ASN multiple times to the routing path. As BGP, in absence of other parameters, prefers the shortest routing path, the erroneous route often has a shorter path than the real route and will be preferred.

An adversary gains access to a BGP router, directly, through the compromise of a computer with access to the routing infrastructure, or through collusion with an ASN owner, allowing them to:

- Influence the path length of routes of interest.
- Announce prefixes the compromised ASN is not authorized for.
- Reconfigure routing policies to facilitate BGP hijacking from upstream ASNs.

There are several scenarios where these threats could manifest themselves:

- An adversary wants to obtain access to the data sent from and to specific prefixes.
- An adversary wants to economically impact the owner of specific prefixes.
- An adversary wants to gain a competitive advantage on the owner of specific prefixes.
- An adversary has a strategic need to manipulate routing to and from specific prefixes

While the attack surface for BGP is admittedly small, the impact of BGP-related events can be very high and, as illustrated below, the probability of BGP-related events is relatively high. We have come to a point where not acting on our need to guarantee the integrity of the core internet infrastructure is no longer an option. In order to maintain the already tried and tested trust in the internet, we have an obligation to do better.

Real world BGP attack/impact examples

Given that it is very difficult to initially discern a BGP attack from an administrative mistake on the BGP infrastructure, I hereby provide a brief analysis of major BGP-related events, malicious or not, over the past few years that illustrate the impact of such events.

The AS 7007 Incident

In 1997, autonomous system 7007 leaked a big amount of routes to /24 prefixes to its peers. Because the leaked prefixes were more specific than the original prefixes this resulted in several networks becoming unreachable through what is commonly known as a routing blackhole. This event is commonly attributed to an extremely unlikely combination of events and configurations, leading to one of the most notorious BGP incidents on the, then nascent, internet.

Youtube Hijacking
On February 24th 2008, it is assumed that Pakistan attempted to block Youtube access within its country through BGP table manipulation. One of the IP prefixes owned by Youtube (208.65.153.0/24) leaked through Pakistan Telecom and its downstream providers to the wider internet. This meant that traffic destined for the Youtube prefix was then routed to the Pakistani internet infrastructure. While the intention was not to advertise the new route to the wider internet, the result was exactly that. The duration of the event was approximately 80 minutes.

Chinese Hijack
In 2010, a Chinese ISP suddenly announced approximately 37,000 routes to prefixes they were not authorized for. While only 10% of the announced routes were further propagated, this impacted major websites such as cnn.com, Amazon.de, and rapidshare.com. It is assumed the BGP announcement was not malicious but again the impact was significant.

Malaysian route leak
In June 2015, a Malaysian ISP leaked almost half of the global routing table to their downstream provider (Level3), severely impacting routing across the internet for a significant amount of time. Again, this was not a malicious attack.

Intential BGP hijacking
In June 2014, an Italian company lost access to – what was for them – a critical server because the hoster they worked with went out of business. The company collaborated with a local ISP to announce the prefixes previously used by the hoster in order to recover clients that would connect to the orphaned server. This type of intentional hijacking is definitely to be categorized as malicious as it severely undermines the trust between ASNs.

The list of BGP incidents goes on as there isn't a day that goes by where an ASN somewhere does not erroneously announces a prefix it is not authorized for. While the impact, in most cases is limited, it illustrates the criticality of the protocol and the potential effects on the reliability and trustworthiness of the internet, and the security of its users.

BGP Security, the options

First and foremost, ASNs should maintain high standards for their core routing infrastructure and its supporting systems. This includes the hardening of BGP routers, hardening of computer systems used to configure and monitor the BGP infrastructure, strict access control to the management infrastructure, and monitoring of their core routing infrastructure.

While prevention is obviously the first goal, ASN owners should be aware that an incident can always happen. It is therefore also recommended that an incident response plan specifically for BGP incidents is put in place. This allows the ASN owner to swiftly react in case of an incident.

Additionally, there is currently very limited communication around BGP-related incidents. P-mortem analysis of such events is often provided by third party monitoring providers, CERTs, or otherwise interested parties such as BGPMON (now owned by OpenDNS), Dyn Research, or Team Cymru. Inter-ASN, and public, communications about BGP-related events and especially the lessons learned should not only be encouraged but demanded.

Just like with DNS, several efforts to implement security in BGP have been considered and tested. While no proposed technology has provided a solution to all potential threats, the technology that has proven the most suitable is Resource Public Key Infrastructure (RPKI).

With RPKI the regional registries (Afrinic, APNIC, RIPE, etc.) act as trust anchors in the infrastructure. As they are the authorities that assign Autonomous System Numbers and IP

prefixes for their regions, they are the default point of trust for these resources. Owners of an ASN can then request a Route Origin Authorization (ROA) which is a signed statement, based on the X.509 certificate standard, that identifies the ASN, the prefixes it is authorized to announce, the minimum and maximum length of those prefixes, and the validity of routes. In essence looks like this:

```
Origin ASN        : 1234
Not valid before  : 2015-08-05 00:00:00
Not valid after   : 2016-08-04 23:59:59
Prefixes          : 1.2.3.0/24 (max length /28)
                    2.3.4.0/18 (max length /32)
```

With these ROAs in place, routers can now rely on an RPKI validator to validate the routes announced by their peers. When an RPKI validator is used, there are 3 possible outcomes from the validation process:

- validation succeeded
- validation failed (ROA found but no validation success)
- validation unknown (ROA not found)

A BGP router that validates ROAs can then implement routing policies based on the validation results rather than using arbitrary configurations as is the case now.

A long way ahead for ROA validation

Unfortunately, there is very little momentum for the adoption of RPKI in the BGP infrastructure. According to rpki.surfnet.nl, a resource that tracks the adoption of RPKI, validators can only validate 6.5% of the currently announced prefixes, valid or invalid. Of the Alexa top 500, only 16 websites lie within IP prefixes that can be validated. The best-known examples are facebook.com, Mozilla.org, web.de, etc. but other immensely popular websites do not have a valid ROA.

When we look on a per country level, we see that Azerbeidjan has the highest percentage of prefixes with an ROA (based on assigned prefixes). If we look at raw numbers, we see that France has the highest total number of prefixes with an ROA. While I don't know why Azerbeidjan has such a high RPKI adoption rate, it is no accident that France is in the lead when it comes to total number of validated prefixes. France has done a considerable effort to create awareness around the BGP problem space. This includes guidance for ASN owners and recommendations for secure BGP router configurations.

In the past 6 months I have worked with a national CERT in EMEA to achieve the same in their country. We have started by identifying the most pressing BGP security risks through table top exercises, which have the added benefit that they contribute to the knowledge level of the participants. Based on the results from these exercises, we have started to create BGP-related tooling and guidance.

As an example, the CERT is currently hosting an RPKI validator that local ASN owners can leverage immediately. The CERT will also publish a BGP dashboard and open source the data collection and analysis software behind it. The initiative is by no means ground breaking but instead of leaving it to each individual ASN owner to figure everything out, it significantly lowers the bar (and necessary effort) to participate. I would encourage each organization that has a direct incentive to participate in a more secure BGP infrastructure to aim for the same :

- Focus on collaboration and joined learning
- Focus on information sharing
- Find common ground in the goal of the project and contribute according to your own means

At this moment it is estimated that RPKI adoption will be at 50% of the announced prefixes by 2020. This is, in my honest opinion, not enough. With a collaborative effort, I am confident that we can move much quicker without negative consequences.

Each ASN owner can, through their regional RIR, request ROAs for their ip prefixes right now. There is no direct negative impact for this. Additionally, each ASN owner can set up an RPKI validator using the open source validator tool offered through RIPE and configure their routers to validate routes presented to them by their peers. Again, this has no immediate negative impact, it just needs to happen.

The reason why the actions described above have no immediate negative consequences is that routing decisions are not immediately made based on ROA validation status. Individual ASN owners will have to configure routing policies that can leverage ROA validation status in their decision tree.

Increasing the number of prefixes with an ROA should not take another 5 years. Knowing what we know right now, ASN owners should seriously consider to start using ROAs right now. Once there is enough adoption across the BGP infrastructure, ASN owners can gradually start implementing routing policies that take ROA validation status into consideration. As we will evolve through a learning phase while ROA adoption is increasing, we are able to make the routing policies stricter as we go.

BGP Data Analysis and tools

The need for BGP monitoring is not immediately obvious for organizations that are not ASN owners. However, this can be very valuable as illustrated in this section of the paper.

First and foremost, all traffic from and to your network is influenced by BGP. This means that any traffic that has business value to you, whatever that may be, can be directly impacted through BGP manipulation.

To ensure a secure path for all traffic, it could be very useful to know which ASNs know the path to the ASN that owns the prefixes your assets use IP addresses in and whenever changes to that path happen. There are both commercial and open source solutions that can be leveraged for this purpose. I believe it is important to understand data and its usefullness before deciding to pay for solution and as it happens to be, there are several useful data sets and tools available that can be leveraged at minimal cost.

Colorado State University

This university maintains peering relationships with several ASN owners for the sole purpose of gathering BGP-related data. While their own BGPMON tool can be very useful for BGP monitoring, they also provide a continuous XML data stream that contains all BGP announcements they receive through their peers in real time. This includes route announcements and route withdrawals. The team at Colorado State University is always on the lookout for new peering relationships, especially in currently underserved regions, as it allows them to increase the quality of their data set. I would encourage anybody that owns an ASN to get into a conversation with the team and start to contribute to a very useful effort. The streaming data is very easy to consume using a variety of programming languages and provides real time insight into BGP events.

bgp.he.net

Hurricane Electric (HE) is probably best known with people that started looking into IPv6 at a very early stage. They were one of the first to offer 4-to-6 tunneling back in the day. Their BGP dashboard and dataset is very useful for anybody looking into BGP issues. HE documents ASN to prefix mapping, administrative details, peering relationships, etc. etc. in a very easy to use web interface. While an API could be useful for this dataset, it doesn't exist yet. That, however, does not diminish the value of the dataset.

CIRCL BGP Ranking

CIRCL is a Luxembourg CERT that has a long tradition of making their tools and datasets publicly available. Their BGP Ranking project correlates BGP data (ASN and prefixes) to data gathered from malware intelligence, ip address blacklists (e.g. SMTP, SSH, etc. etc.), SANS DSHIELD information, etc. This allows them to score each ASN for the amount of recorded malicious behavior. This data and scoring is then visualized on a world map. Additionally, CIRCL provides an easy search interface to retrieve information related to individual ASNs. Again this information is extremely useful for BGP researchers.

Team Cymru BGP Monitoring

Team Cymru is a collective of individual technologists with a common interest in making the internet more secure. They operate a website where they provide a significant number of reports and data points around BGP that are, again, useful in combination with other data sources.

Team Cymru also operates a so-called looking glass server that you can remotely connect to to obtain up to date BGP information.

Agence national de la sécurité des systèmes d'information -ANSSI (France)

ANSII is one of the very few government organizations that provides a best practices guide for their constituent ASN owners. While there are other guides (like those from NIST), the ANSSI guide is by far one of the most comprehensive guides available, including configuration recommendations for all BGP routers from major vendors. It is a document that can be built upon to create recommendation guides for other locales.

Routing Resilience Manifesto – MANRS

The manifesto shares a common objective with this paper : encouraging network operators to adhere to a minimum set of standards that allow for a more resilient internet backbone.

Zebra Dump Parser

This tool, written and released by Marco d'Itri is specifically helpful to parse the BGP information published by RIPE, in addition to data gathered using the open source routing software Quagga/Zebra.

BGPMON

BGPMON is a commercial service developed by Andree Toonk, and recently acquired by OpenDNS, that provides a real time alerting service for BGP-related events. There are several other BGP monitoring services and products, but I believe BGPMON is the most comprehensive and useful among them. Not in the least thanks to the efforts of Mr. Toonk before it turned into a commercial service.

Conclusion

BGP incidents happen on a daily basis. They impact everybody's online resources and the trust in the internet. With all the knowledge we have about the BGP attack surface and the tools we have available to protect against it, now is the time to actually do it.

- ASN owners have a responsibility to everybody that uses online technology to provide the most secure environment possible.
- RPKI technology can help to add security to the routing decision workflow.
- We need to do better than a projected 50% RPKI adoption rate by 2020.
- ASN owners need to join in a collaborative effort, similar to what happened with DNSSEC.
- There is no such thing as 100% secure or a 100% effective solution but there is something much better than doing nothing.
- Organizations should consider monitoring for BGP events related to the online resources they value.
- We need an internet that has ownership and accountability built into it.

References

http://rpki.surfnet.nl/index.html
http://www.bgpmon.net
http://www.washingtonpost.com/sf/business/2015/05/31/net-of-insecurity-part-2/
http://www.ssi.gouv.fr/uploads/IMG/pdf/guide_configuration_BGP.pdf
http://bgp.he.net
http://www.potaroo.net/bgp/iana/asn-ctl.txt
https://en.wikipedia.org/wiki/IP_hijacking
http://circl.lu/projects/bgpranking/
http://www.bgpmon.io/
http://www.team-cymru.org/BGP-monitoring.html