



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Hacker Tools, Techniques, and Incident Handling (Security 504)"
at <http://www.giac.org/registration/gcih>

BGP Hijinks and Hijacks - Incident Response When Your Backbone Is Your Enemy

Author: Tim Collyer, tccollier@gmail.com

Advisor: Rob Vandenbrink

Accepted: 11-19-2016

Abstract

The Border Gateway Protocol (BGP) is used to route packets across the Internet, usually at the level of the Internet backbone where Internet Service Providers (ISPs) pass traffic amongst themselves. Unfortunately, BGP was not designed with security in mind, like many of the protocols used in modern networks such as the Internet. Lack of security within BGP means that traffic is susceptible to misdirection and manipulation through either misconfiguration or malicious intent. Among the traffic manipulation possible within BGP routing is Autonomous System (AS) path injection, in which a new router can insert itself into the routing path of traffic. This can create a man-in-the-middle condition if the path injection is malicious in nature. Differentiation between a malicious incident and mere misconfiguration can be extremely challenging. Even more difficult for an affected company is to conduct incident response during a BGP-related incident. This paper explores the incident response options currently available to security teams to prevent, detect, and where possible, respond should a BGP incident arise.

1. Introduction

When tasked with the challenge of securing the information used and created by an organization, Information Security professionals have plenty of areas to consider. A good security program tackles securing applications, the endpoints, the network perimeter, the network interior, user credentials and identities, and more. The security industry has numerous aphorisms to help its practitioners keep their focus in the right places. “The network perimeter is more permeable than ever,” helps to remind teams that it is no longer enough to try to protect a corporate network for example because devices and data routinely need to leave the protected network. “The user is the new perimeter,” is another favorite. This saying is a reminder that attackers are finding it easier to compromise machines and access data by going through the user - through phishing or other social engineering attacks. All this is in contrast to attacks which attempt to exploit servers and network services directly.

These reminders can be very helpful to summarize the prevailing wisdom. But what about attack vectors which fall outside the prevailing wisdom or outside our basic assumptions? If one were a fish living in the ocean, it would be natural to worry about sharks and other predators and to look for secure hiding spots in your coral reef. But do fish worry about pollution in the water in which they swim? Similarly, though concerns around endpoints, users, and the trusted network are all valid and worthwhile, there are vulnerabilities in foundational technologies that are equivalent to problems in the water where fish swim. These kinds of problems usually arise in technologies which were created, for various reasons, without considering security. While busily hardening our endpoints, patching our applications, and securing our passwords, it may be easy, or more convenient, to forget that the underpinnings of the network itself may be subverted to bypass the other controls.

BGP is an example of such a technology. BGP is a routing protocol in use across the Internet to not only help deliver traffic to its destination but also provide the resiliency and reactivity which supports a fast and stable Internet. The uses of modern

Tim Collyer, tccollier@gmail.com

networks and attacks against them are vastly different than they were even in the mid-90s, which wasn't so very long ago, whereas the protocols which allow networks to function have not significantly changed in the same period of time.

BGP-4 was first described in RFC 1771 (Rekhter & Li, 1995) in March of 1995. The authors of BGP-4 were likely focused on the needs of the time - functionality being paramount. BGP as a protocol was designed without security in mind and there is no better proof of this than in RFC 1771 itself. There is, in fact, a section titled "Security Considerations," but it is almost a footnote at the end of the document. All it says is, "Security issues are not discussed in this document."

While this may be understandable, the ramifications of this lack of attention to security are still being felt today. RFC 4271 (Rekhter, Li, & Hares, 2006), published in 2006, provided an update to BGP-4 and included several recommendations about how to better secure BGP, including a reference to RFC 4272, "BGP Security Vulnerabilities Analysis." RFC 4271, including all of its revisions, are in part an attempt to bolt security on top of something which never intended to support it. The fact that BGP hijacks are still relevant today is an indication that this issue has not yet been solved despite these ongoing efforts.

A well-positioned attacker can use BGP to subvert, monitor, and tamper with traffic on the Internet. Security professionals may work hard to train users to avoid falling victim to phishing emails and clicking on malicious links, but that may not matter if an attacker uses BGP to redirect legitimate traffic to illegitimate locations. Since BGP attacks take place outside of an organization's network boundary, it can be extremely challenging to detect, defend, and respond to them. Not only can users be directed to bad destinations without their knowledge and without any flawed choices on their part, but an attacker could use BGP to eavesdrop on almost all network traffic going to an intended victim. Or rather than eavesdrop, an attacker might just drop that traffic, effectively severing all connections to the Internet without ever touching equipment owned by a victim. In summary, attacks against BGP could be leveraged to compromise the confidentiality, integrity, and/or the availability of a company's network traffic.

The lack of BGP security is not newly discovered. Peter Zatko, a.k.a. "Mudge", from the hacking group L0pht was testifying before Congress about it as early as 1998

Tim Collyer, tccollier@gmail.com

(Zetter, 2008). While much has been said and written about the security of BGP and how to improve it, this paper will examine BGP attacks through the lens of incident response. It will explore how to prevent, detect, and respond to BGP attacks (where possible) and map those actions back to the standard steps in the incident response process.

2. Attacking BGP

Before exploring how best to respond to a BGP incident, it is important to have a fundamental understanding of what BGP is, how it works, and what kinds of attacks are possible against it.

The Border Gateway Protocol is used to provide routing and resiliency between destinations on the Internet. Those ‘destinations’ are usually described as Autonomous Systems (AS). RFC 1930 defines an AS: “An AS is a connected group of one or more IP prefixes run by one or more network operators which has a SINGLE and CLEARLY DEFINED routing policy,” (Hawkinson & Bates, 1996). In this description, an “IP prefix” can be considered to be a group of IP addresses, usually defined by some sort of CIDR notation, e.g. 123.123.0.0/22. Fundamentally, an AS is a group of IPs which are managed by a single entity. Within an AS, routing is handled with a coherent internal plan usually using Interior Gateway Protocols (IGPs) for routing. IGPs can include protocols such as RIP, OSPF, IS-IS, and others (Hummel, 2013). How those work and even what the constellation of initials mean is beyond the scope of this paper, but what’s important is that an AS is defined in part by who is responsible for what happens within - including how things are routed. An Exterior Gateway Protocol (EGP), most commonly BGP, is used to route traffic *between* one AS and another.

By analogy, an AS is an apartment building owned by a single company. Residents may exchange messages in a variety of ways like hallway gossip or slipping notes under each other’s doors (e.g. the various Interior Gateway Protocols). However, when a resident needs to send something to another building or city, they use the postal service (e.g. an Exterior Gateway Protocol, likely BGP).

BGP routing can be complex and the intricacies of the protocol will not be explored here. In broad strokes, BGP peers advertise the ASs which they represent and

use these advertisements to build a route to a given destination. This route is called an “AS path” - to get to AS ‘F’ from AS ‘A’, a packet follows an AS path of $A \rightarrow B \rightarrow C \rightarrow D \rightarrow E \rightarrow F$. An important aspect of BGP design is to ensure that an AS path is not a static description of how packets should flow. If in the simple example above, router C fails, BGP will adapt and send packets along a different path, choosing the shortest AS path available.

The inherent vulnerability in BGP comes from the AS advertisements - there is no verification of authenticity or integrity within them. Without verification, there is an assumption of trust that all BGP AS advertisements are valid, which is an opportunity for abuse. If a malicious individual were able to inject false AS advertisements into the BGP network they would be accepted as valid because there is no way differentiate good advertisements from bad. The attacker would, therefore, be able direct the flow of traffic across the Internet. This has happened numerous times, a prominent example of which was in Pakistan in 2008 when, in an attempt to censor access to YouTube, the country inadvertently prevented the entire Internet from accessing the website by manipulating BGP (RIPE Network Coordination Centre, 2008). Even in misconfiguration, BGP attacks can be powerful.

There are a variety of objectives an attacker might achieve by manipulating BGP traffic. They range from a denial of service (DoS) attack by disrupting traffic routing ("Protecting Border Gateway Protocol for the Enterprise - Cisco," n.d.), stealing unallocated IP addresses for sending spam (Vervier, Thonnard, & Dacier, 2015), or through man-in-the-middle (MitM) attacks. Of these various possibilities, this paper will focus on the MitM attack specifically when it comes to incident response as a MitM attack has the most opportunity for ongoing and potentially long-term malicious actions.

2.1 Example Attack

Recently, security journalist Brian Krebs reported on a BGP hijack which was used to take down a popular distributed denial of service (DDoS) provider (Krebs, 2016). A brief examination of the attack can illustrate what a BGP hijack looks like. To start with, this is what network traffic looked like prior to the hijack:

Tim Collyer, tccollier@gmail.com

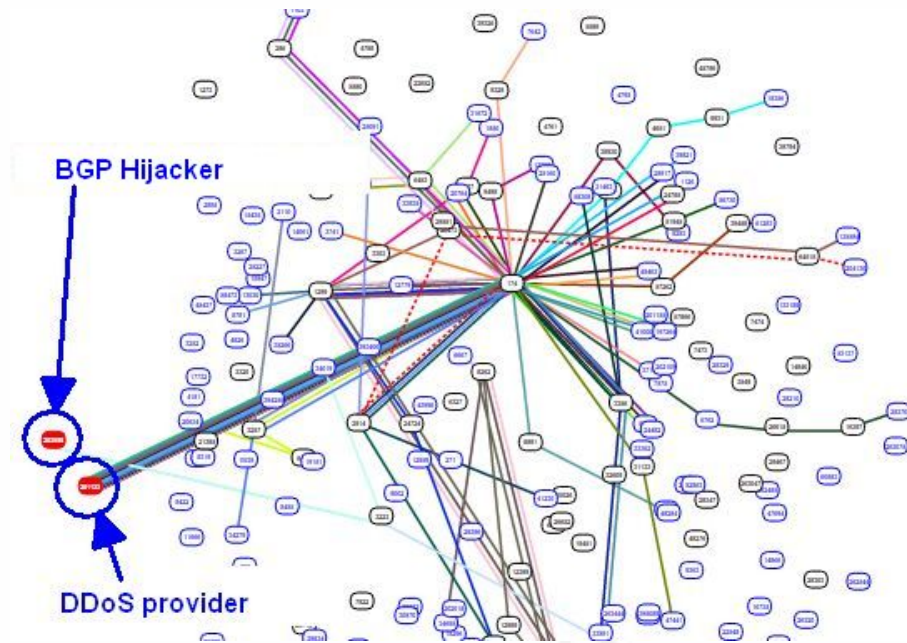


Figure 1 - Normal traffic before BGP hijack ("BGPStream event #54711," 2016)

This diagram can appear confusing, but the details of it are less important to understand than the overall picture for the purposes of this explanation. Each circle on the diagram represents a router on the internet, and the lines connecting them represent traffic flowing as they relate to the source AS, 201133 (labeled DDoS provider). The takeaway here is that most traffic flows between AS 201133 and the center of the spider's web of traffic lines (AS 174, an ISP).

When the BGP hijack began, a new route was announced which directed traffic destined for the IP addresses in AS 201133 instead through AS 3223 and then to AS 203959 (which is labeled "BGP Hijacker" in the diagram). Described with words, the BGP traffic looked like this:

- The new route 13124 **3223 203959** has been announced
- The route 33891 3356 174 201133 is changed to 33891 **3223 203959**
- The route 6667 8262 48452 48452 1299 174 201133 is changed to 6667 **3223 203959**
- etc.

There followed numerous similar BGP route updates, describing similar changes until the flow of traffic looked like this:

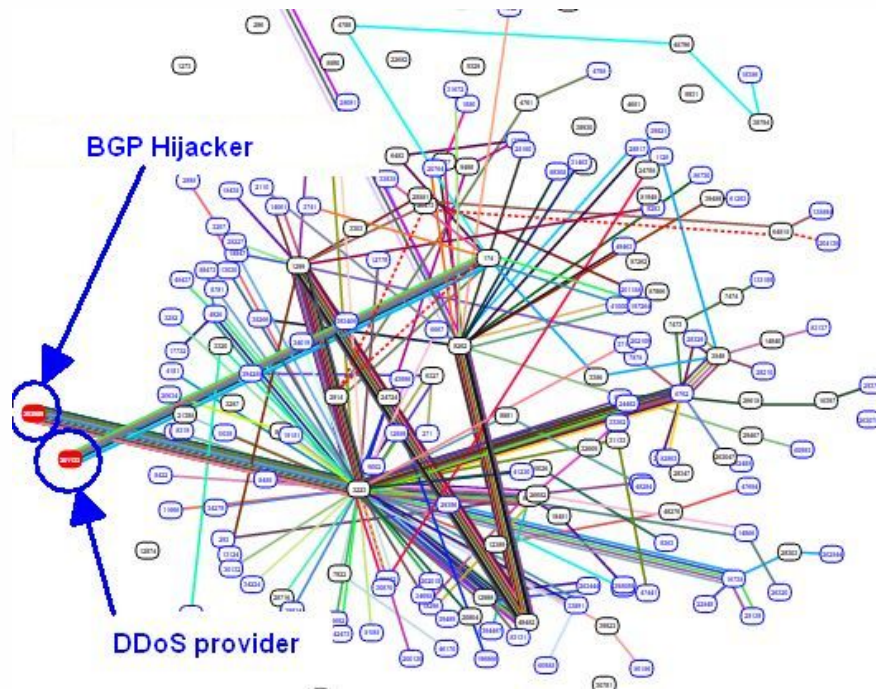


Figure 2 - Traffic flow during the BGP hijack ("BGPStream event #54711," 2016)

The flow of traffic was clearly significantly altered and most traffic in this diagram which was destined for the IP addresses in AS 201133 was now headed to the wrong place. This change in traffic took approximately 2 minutes to achieve. It is also important to note that not *all* traffic was changed. The valid BGP routes were still present and some routers still sent traffic through the correct path.




Figure 3. - BGP route hijack penetration. Red represents the hijacked route, purple the valid route (Madory, 2016).

Also, it is important to note that if this had been a MitM attack, the new destination (AS 203959) would have needed a ‘clean path’ down which to send traffic (Cowie, 2013). That is, the traffic needs to get to the true destination eventually or the attack is merely a DoS, not a MitM, so the false destination needs a way to send the traffic back out to the proper recipient. This is done by maintaining an AS path which is not polluted with the hijack information and which has the correct destination.

In point of fact, the BGP hijacker in this case (BackConnect Inc.) described the purpose behind the attack as, “to collect intelligence on the actors behind the bot net as well as identify the attack servers used by the booter service,” (Townsend, 2016). The comment from BackConnect highlights the power of just monitoring the traffic via BGP hijack, not even going to the lengths to establish a full MitM.

3. Incident Response

The incident response process is frequently defined as having 6 steps: Preparation, Identification, Containment, Eradication, Recovery, and Lessons Learned (Computer incident response guidebook, 1996, p. 7). These steps provide an excellent framework or thought model from which to construct an incident response process. The steps may seem linear but what actually happens during a real incident is likely to have blurry boundaries between one activity and another or require returning to an earlier step as the understanding of the scope becomes clearer. When considering how to approach incident response to a BGP hijack, this framework provides a convenient way to examine the options but true incident response will need to incorporate an iterative process, returning to previous steps as needed.

There are two aspects to a BGP incident that are worth considering. The first part is the attack on BGP itself - manipulating the protocol and the flow of traffic as a means to an end. The second part is the end itself – some malicious purpose. The nature of a BGP hijack, particularly when it leads to a MitM condition, affords the attacker the opportunity to achieve standard goals (i.e. installing malware on a system) in non-standard ways (i.e. no user clicking on a link required). Those different mechanisms are worth considering from an incident response perspective where the response required deviates from more typical incident response.

3.1 Preparation

Because of the fact that BGP attacks happen outside of the purview of an organization's network, many of the steps in the incident response process are constrained in what can be done. Successful incident response is generally the result of an effective Preparation phase. This is true for other types of incidents of course but is even more so for BGP attacks.

Some view the Preparation phase for incident response as considering what to do when preventative controls fail. However, this paper will stretch the strict definition of the Preparation phase to include prevention as well. After all, it might be said that the best type of preparation for an incident is the kind which prevents it from happening. Some consideration of prevention is included here because of the extremely limited

Tim Collyer, tccollier@gmail.com

options for effectively responding to a BGP attack. Though an incident response team can pick up the pieces and track down what went wrong as a result of a BGP hijack, the only real way to get ahead of damage to an organization is to prevent the attack. Once a BGP attack has begun, some damage is likely to occur, though the extent of the damage depends upon the nature of the attack and the response.

There are a few preventative measures available for BGP attacks. The simplest measure, in some ways, is to implement some amount of route filtering and specifying of an AS Path. The challenge with this is that the more one specifies exactly what path traffic should take, the less resilient the routing system becomes when faced with system failures. Route filtering can be an effective security control for BGP configurations that are at the edge of the internet backbone which should in general not be involved with routing decisions made for distant Autonomous Systems. However, for BGP routers at the core of the internet, the dynamic shifting of routes is a required feature of BGP. In such cases, BGP supports neighbor authentication by using MD5 hashes with a pre-shared key as a salt.

Another promising security measure is through Resource Public Key Infrastructure (RPKI) which helps to perform Route Origin Authorization (ROA) ("Protecting Border Gateway Protocol for the Enterprise - Cisco," n.d.). In essence, this just uses a Public Key Infrastructure like that used to support certificates for authenticating websites (i.e. the certificates used in HTTPS), to authorize a specific AS to advertise for IP prefixes (Huston & Michaelson, 2012). This attempts to address the underlying problem with BGP - anyone can send out a route advertisement for anything and it will be trusted. With ROA, only approved sources can advertise for a specific range of IP addresses.

There are several problems with this approach for incident response teams. First, implementations of BGP security features such as neighbor authentication or RPKI require changes to router software and/or configurations which limits the rate of adoption of such features. For example, NIST monitors the adoption of RPKI and although adoption is growing, as of this writing it is still under 6% of unique IPv4 prefix/origin pairs.

Tim Collyer, tccollier@gmail.com




Figure 4 - IPv4 prefix/origin pairs validated by RPKI; note the scale on the Y axis is logarithmic (NIST, n.d.)

Until a critical mass of routers support it however, it is difficult to rely on RPKI since many legitimate routes will not incorporate it and to be truly effective it needs to be implemented across each hop. As with many difficult-to-implement security measures however, an ISP may not choose to undertake rolling out RPKI without some indication from customers that there is a demand for that level of security, so it is worth inquiring of an ISP about their involvement with that type of security measure to drive adoption forward.

The bigger issue is the same one that permeates much of the discussion around BGP attacks in that by definition most BGP events happen outside of an organization's network. So even if RPKI were the silver bullet to solve all of BGP's security problems, an organization may have little say in the externally facing router configuration without an appropriate agreement with an ISP. The real preparation step here is one of 'security through contract negotiation.'

Tim Collyer, tccollier@gmail.com

Organizations should not simply trust that everything will be OK once traffic is outside their networks. It is extremely important that security elements of the relationship with an ISP be agreed upon during the contract negotiation period. At the least, they should specify visibility into router configurations and ideally some ability to dictate changes in that configuration. Also, if preventative measures like RPKI are not available, other options for BGP incident response should be discussed up front. The worst time for an incident response team to explore what options there are with an ISP is during or after a BGP attack. At that point, the security team is likely to speak with the ISP sales team, not engineers who might actually be able to help.

It is unlikely that BGP security measures will actually be incorporated into a contract, but contract negotiation is when service providers are most pliable and likely to provide appropriate contact information and other requested access. For example, visibility into the BGP traffic arriving at the last hop can provide pertinent insight, in addition to third-party BGP monitoring which will be discussed elsewhere. Routers may also be set to log any BGP update events to help give visibility into changes. However, unless this is enumerated ahead of time, most easily done at the time of contract negotiation, an organization will not have access to that BGP traffic during or after an incident to help explain the timeline of events. One other element to keep in mind, is that normally BGP and other networking agreements may happen under the purview of the networking team and without the security team present to make requests. A good security team cannot be successful while operating in a vacuum. Establishing communication and relationships with other teams, networking in this particular case, is another key step in the Preparation phase. Naturally, this advice applies to more than just BGP incident response.

Provided that the security team is at the table when an organization sets the working relationship with an ISP, another important preparation step is perhaps even more basic. Because a good portion of the incident response for BGP (as opposed to BGP attack after-effects) is in the hands of other entities, an incident response team will need to lean on the relationship with and contacts at the ISP. Unless the ISP is aware that a company is interested in responding to BGP incidents, it is unlikely that the appropriate resources will be available at the right time. The first step therefore is to open the lines of

Tim Collyer, tccollier@gmail.com

communication by bringing up the subject of BGP hijack incident response with the ISP. The resulting conversation can explore what options and resources the ISP can provide in the heat of the moment. This is especially important because the ISP may need to contact other ISPs to eliminate the false BGP information and reconstruct the origin of the hijack. They may not be prepared to conduct such operations without prior warning so that they can set a procedure to do so.

3.1.1 Secondary Attacks

If preventative measures have not been effective to avoid a BGP hijack of a malicious nature, an incident response team needs to consider what the purpose behind the BGP hijack might be. The most damaging form of BGP hijack, depending upon a victim's business, is one in which no network disruption occurs and an attacker is sitting in as a MitM. From that vantage, an attacker can execute any number of secondary attacks, installing malware, capturing credentials, providing disinformation, etc. When an attacker can manipulate traffic in that way, the entire Internet becomes potentially hostile (or rather, even more so). Some normal methods to detect malicious attacks - monitoring email attachments and links for example - will not be effective to identify that a user may have been attacked. Part of the incident response process for a BGP hijack incident is therefore to identify what the secondary objective is for attackers that have successfully hijacked BGP. What happened after the BGP attack began?

In order to identify some of the more subtle possibilities from a MitM type of attack, it is most helpful to be able to examine the traffic arriving from the network to look for anomalies. Once again, preparation is the only way to be successful in such a situation. The first step is to ensure that the security team is capturing all traffic crossing the network perimeter. This means having legal permission, appropriate policies and expectation of privacy notifications (i.e. no expectation of privacy) for users, etc. Then there are the hardware requirements - network taps and devices capable of processing and recording the traffic. There are many options for full packet capture, including commercial offerings like RSA's NetWitness technologies as well as open source software designed to run on available hardware - Security Onion being a leading

example. Both of these product suites handle capturing the raw traffic (from a software perspective) as well as some amount of analysis of the traffic.

Chances are that if and when full packet capture is running across an organization's perimeter, the retention period of all of that raw packet data will be limited. Even with relatively cheap storage available, the quantity of data becomes unmanageable over any significant period of time. In comparison, the traffic analysis and metadata is only a percentage of the raw traffic and as such it can have a much longer retention period before it needs to be purged to make space for newer data. This relates to BGP hijack incident response in that even if containment and other response efforts are successful to keep the period of the hijack to relatively short, there is still likely to be a fairly large volume of packets to comb through above and beyond whatever default analysis of network traffic may be happening.

Detection of secondary attacks resulting from a BGP hijack likely requires an analysis environment to more intensively search through network traffic collected during the window of the hijack event. Again, this is a capability which needs to be considered ahead of time. The analysis environment needs enough storage to be able to contain the full packet capture from the hijacked Autonomous System traffic. It also needs the processing power and memory to run through the saved packets in an effective and timely manner. And just as important, it needs personnel familiar with trying to find needles in that particular type of haystack. It takes time to develop scripts (or applications) designed to iteratively process multiple pcap files and look for evidence of SSL certificate tampering, DNS manipulation, or just plain injection of malware. That time should be spent *before* an incident, so that the time for incident response does not also include time for development of response capabilities.

To help set a sense of scale, I was able to obtain information from a mid-size enterprise which was conducting incident response on the traffic from a BGP hijack. Though the hijack window was only 50 minutes, the volume of traffic collected from the perimeter topped out at around 3 Tb. This included approximately 2.75 billion sessions which involved 3.7 million destination IP addresses. Just peeling the DNS queries observed within that volume of traffic yielded approximately 4.5 million DNS requests, of which ~70,000 were unique. Not all environments are that large (though this was by

Tim Collyer, tccollier@gmail.com

no means the far end of the spectrum of network size), nevertheless, these numbers indicate that an analysis infrastructure is required to process that much traffic. This cannot be handled in an ad hoc and manual way by one or two analysts. Eventually, the automated analysis can help to pare down the results to something manageable and interesting enough to warrant manual review by an analyst; but it is important that such automation be configured ahead of time for effective incident response to the after effects of a BGP hijack.

Another related item in the preparation category addresses the specifics of looking for anomalies in DNS or SSL certificate information. If an attacker is able to manipulate DNS traffic as a result of a BGP hijack, incident responders need to be able to compare valid DNS responses, for example, with those received during the hijack window in order to check for signs of tampering. Again, this is a capability to be set up before an incident. Even with historical data near the time of an incident, it may be helpful to have a third party perspective/database to reference. Incorporating the ability to query such databases is best included within the Preparation phase, though the mechanics of using such capabilities are discussed more during the Identification section.

3.2 Identification

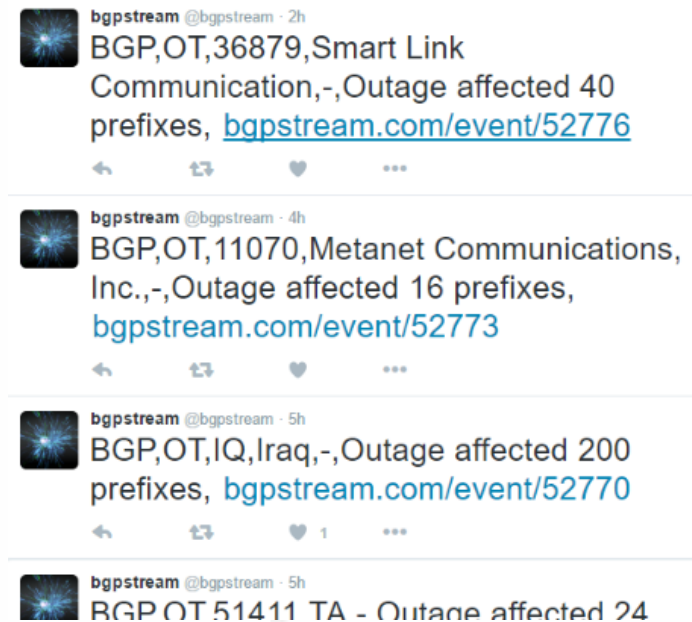
The first reactive step in any incident response process is to recognize that an incident is indeed happening. Ideally, incident identification happens as early as possible so that response can help to avoid damage, not just recover from it. But sometimes ‘identification’ happens with something like the legendary call from an FBI agent, “You have a problem.” Or worse, a company identifies that there was an incident because someone else comes to market first with what was supposed to be proprietary technology indicating that there must have been a leak or loss of intellectual property.

With respect to BGP hijacks, part of the challenge is that the incident starts outside of the trusted network. This complicates identification efforts because monitoring traffic within the trusted network will not reveal any indication of something unusual with BGP. The solution to this is fairly simple - an organization needs to monitor BGP routes which relate to their AS. Typically this is done through third parties, of which there are many. BGPMon (<http://www.bgpmon.net/>) is one such example. One can sign

Tim Collyer, tccollier@gmail.com

up for free to monitor a limited number of prefixes or pay a subscription fee for more full-fledged monitoring of a larger group of prefixes.

BGPMon also maintains BGPStream which is an automated service that provides some alerting. The easiest way to view BGPStream alerts is through its Twitter feed:



*Figure 5 - Sample from the bgpstream Twitter feed of BGP alerts
([https://twitter.com/bgpstream](\"https://twitter.com/bgpstream\"))*

The links in each BGPStream tweet provide a more detailed view of the BGP event and allow for a replay to watch as the BGP routes change:

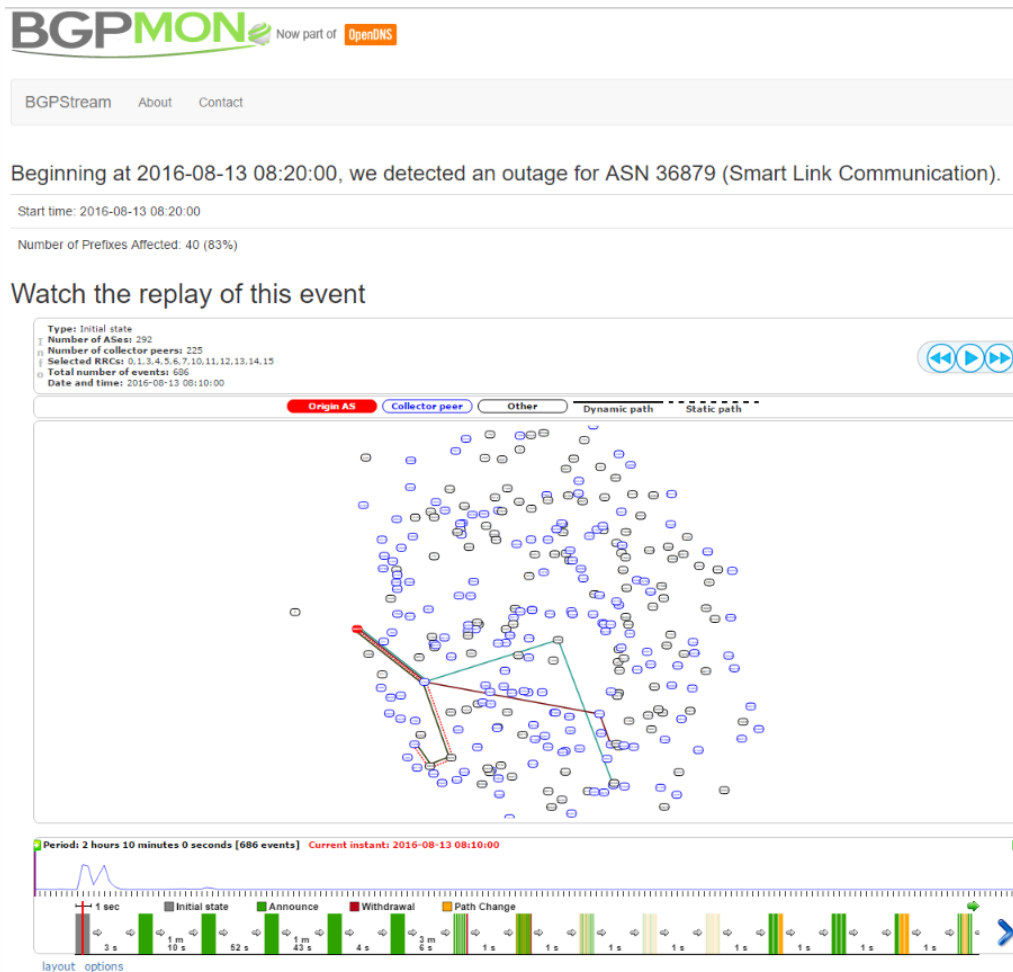


Figure 6 - Sample event reply from BGPMon/BGPStream ("BGPStream event #52776," 2016)

Through the use of a monitoring service like BGPMon or “Internet Intelligence - Transit” from Dyn.com (<http://dyn.com/ip-transit-intelligence/>), an organization can receive an alert that something unusual is occurring within the BGP routing tables that relate to their AS. This is the most effective way of identifying a BGP attack and initiating the rest of the incident response process.

3.2.1 Secondary Attacks

When it comes to follow-up attacks after a successful BGP hijack (specifically with the aim of obtaining a MitM position), the possibilities are numerous. From an incident response perspective however, the identification process boils down to looking at the traffic collected during the hijack window and answering the question, “What has changed in this traffic that would not have been there without a MitM?” Perhaps simple

Tim Collyer, tccollier@gmail.com

to ask, but not very simple to answer, especially at scale. It should also be noted that there is also the passive portion of exposure to a MitM. This can be summarized with the question, “What has been exposed to the attacker?”

When an attacker wants to move from passively recording traffic, to taking actions against a target, they must inevitably alter some aspect of the network traffic. Not all alterations will be detectable as changes by the incident response team. For example, simply injecting malware into the return stream of traffic, with no other redirect etc., would not be obvious as an injection (assuming full control over a packet’s contents, including checksums) unless the malware itself triggered alerting. Similarly, disinformation - intentionally seeding traffic with false information - would be an alteration which would be difficult to detect but which could have extremely detrimental effects.

However, many of the actions which an attacker might take could leave more traces in the network traffic, even if slight. For example, rather than have a user click on a link which leads to a malicious website, an attacker in a MitM position can instead provide a malicious response for a legitimate request. This can be done in a variety of ways, but poisoning the DNS response to a legitimate DNS request might be an easy start.

By comparing previously captured DNS traffic with traffic from the period of a BGP hijack, a response team may be able to catch unusual IP addresses which have been substituted for the proper destination. Historical DNS logging from internal DNS servers can help answer these questions and the most cost effective and efficient means for analyzing DNS traffic for signs of tampering. However, an organization must be capturing and storing DNS requests/responses ahead of time. If that capability has not been configured at the time of an incident, it is possible to use a third-party service for such data, for a price.

Services such as DomainTools or PassiveTotal track data for Internet domains, including DNS and WhoIs information and they support a good API to be able to enrich large data sets, as might be found in the course of responding to a BGP hijack. For example, the DomainTools Hosting History (<https://www.domaintools.com/resources/api-documentation/hosting-history>) provides

Tim Collyer, tccollier@gmail.com

this historical information needed to be able to compare what a DNS response ought to have been with that which was observed during the course of an incident.

The DomainTools dataset is powerful, but is not free. Running through every unique DNS request which left the internal network over the course of a hijack may be prohibitive. Any reduction in the dataset by using local resources, or locally observed DNS traffic is helpful. In addition, it is certainly possible to use a quick script to run iteratively through a command-line utility like ‘dig’ or ‘nslookup’ to compare live DNS responses with those captured during an incident. The further that the investigation gets from the actual incident, the greater the potential drift from normal churn as websites and IP addresses change. This can lead to false positives - variance from what was observed during the attack but which are benign in nature. A historical dataset like that maintained by PassiveTotal or DomainTools is helpful to reduce or eliminate such problems.

Another potential avenue of attack would be to send users to false websites with ‘valid’ but stolen certificates. Such an attack might build upon the DNS poisoning. For example, if a user attempted to visit <https://www.example.com> and received an attacker-altered DNS entry, the user would be sent to the attacker IP address instead of the legitimate IP address which hosts example.com. However, in order for the attacker to maintain the flow of traffic without alerting the user, they would need to provide a valid certificate for example.com so that SSL encryption could be negotiated. It might be possible to steal a copy of the actual example.com certificate, or to compromise a Certificate Authority and generate a valid but illegitimate certificate for example.com. Easiest of all might be to convince personnel at a valid Certificate Authority to issue false certificates, as Symantec did for Google in 2015 (Goodin, 2015). SSL certificate misuse is exactly the type of thing that Certificate Transparency (<https://www.certificate-transparency.org/what-is-ct>) has been designed to prevent. The Certificate Transparency project aims to provide audit capabilities of SSL certificates which have been issued and maintains tamper-proof logs to support that auditing. As of this writing, the project has accumulated logs for approximately 83 million certificates, though Chrome only requires Certificate Transparency to be enabled on Extended Validation Certificates. Google is moving toward requiring Certificate Transparency for all X.509 certificates, but has set no enforcement date at this time.

Tim Collyer, tccollier@gmail.com

In order to extract the DNS or SSL information from packet captures, an incident response team can use a variety of tools to parse through the traffic from open source to commercial solutions. Bro Network Security Monitoring (<https://www.bro.org/>), is a great open source tool which can carve almost anything out of a packet capture in one pass. Bro can be configured to work with a cluster of servers or to run in standalone mode. One aspect to note is that Bro is designed to parse traffic as it flows by in real-time. Bro definitely supports reading from a saved pcap file, but that is not the primary use-case.

Even if an organization has Bro configured in production, it may or may not be combing through traffic as finely as would be needed to detect subtle secondary attacks from a BGP hijack. This is why it is useful to have an environment ready to export the network traffic observed during an attack of this sort. One challenge is that Bro does not inherently support reading from multiple pcaps iteratively, as might be an expected requirement when going back through a large collection of network traffic from a BGP hijack. A little bash for loop work can take care of this. Below is an example of one way to solve this:

```
for file in $(ls <path to pcaps>/<location>/*.pcap); do t=$(echo $file |  
cut -d '/' -f 5 | cut -d '.' -f 1); mkdir /var/bro/$t; bro -C -r $file  
/usr/local/bro/share/bro/policy/frameworks/files/extract-all-files.bro; mv  
*.log /var/bro/$t; done
```

This command assumes that there may be multiple locations, e.g. egress points, where network traffic would be captured during a BGP hijack. The command then creates appropriately named directories, based upon the name of the pcap files, in which to put the Bro logs. The command also specifies the extract-all-files.bro configuration file (the path to that file may vary depending upon the system and how Bro was installed) which will carve all files possible out of the traffic. This is a brute force approach but can be useful if something suspicious is discovered during analysis. The “-C” option tells Bro to ignore checksums for a cleaner reading of a saved file. This command is not intended to represent the most optimized way to parse multiple pcaps, but merely serve as an example of one way to do this.

Tim Collyer, tccollier@gmail.com

Another important element to note when attempting to use Bro in this fashion (reading large volumes of pcap files) and with such an aggressive approach to carving files out is that this will push the available resources of a server, including the file table. This reveals a bug in the ext4 file system which uses a feature called `dir_index` to help with performance. `Dir_index` stores file information as a hash in a hashtable, but unfortunately does not grow its hashtable when it starts to fill up (Wagner, 2013) as it will when reading large volumes of traffic and carving out all files. As a result, the operating system reports that it has “no space left on device.” It is possible to disable `dir_index` in ext4, but the performance of reading pcaps suffers significantly. It is better to use a different file system altogether, like XFS, to avoid this problem.

Bro will extract any kind of file it encounters in a packet capture, including SSL certificates. SSL certificates extracted by Bro will be DER-encoded (in binary form, not ASCII), but these can be processed iteratively with a simple bash for loop and openssl, e.g.:

```
for f in <directory with extracted certs>*; do openssl x509 -in $f -inform der -fingerprint -subject -serial -noout && echo ','; done> ssl.csv
```

This command tells openssl to look for X509 certificates which are DER-encoded and to extract the fingerprint, subject, and serial for each one and output the results to a csv file. By querying a repository such as that maintained by the Certificate Transparency project for the same certificate subject, an incident response team can look for SSL certificates which have not been properly issued and may indicate malicious activity. As with hunting through DNS responses, examining SSL certificates will require some automation through scripting.

DNS and SSL manipulation are not the only options available to attackers (and therefore to incident responders in the Identification phase), but they provide an example of the kind of activity which might be visible as well as the challenges of finding it. Traditional malware detection in network traffic and on hosts will continue to be helpful during the Identification phase of a BGP hijack.

3.3 Containment

Often, the Containment phase of the incident response process considers a discrete source of compromise - malware in the network being the easiest example. When a malware worm is spreading through a network, the concept of containment is easily understood: prevent the spread of the worm to more systems! But what does containment mean with respect to a BGP hijack?

Because BGP relies upon the concept of ‘advertisements’ for new routes, the injection of a malicious BGP route does not have an instantaneous effect upon the flow of traffic across the Internet. Rather, that route needs to spread - to propagate to other routers. Preventing further spread may seem to be an opportunity for containment efforts, but it is unlikely to be an effective point. Circumstances will vary, but the time that it takes for a BGP route to propagate is measured in minutes, not hours or days. Containing the propagation is unlikely to be possible within the short window available to respond.

However, it is possible to strive for containment by limiting the exposure time to BGP traffic manipulation. This requires removing the false route advertisements by either stopping the source if that is even possible, or convincing other ISPs to drop or filter the false advertisements. Whatever the response, it needs to be done by the organization that owns the routers - the ISPs. This means that the victim needs to have effective contacts with the various ISPs (an element of the Preparation step).

3.3.1 Secondary Attacks

The final possibility for containment relates to the possible purpose behind a BGP hijack, specifically an AS path injection. Containment of the next phases of an attack depends upon the nature of those next steps. If the purpose of the BGP hijack is to sniff traffic for later analysis, the best containment option is to limit the amount of time (and therefore volume) that traffic is exposed. Attackers can use their position to actively alter unencrypted conversations by redirecting users to malicious websites for example. This would allow an attacker to silently achieve the same ends as sending a phishing email but without requiring a user to click on a link. In the end though, the purpose of the malicious website is the same, often to deliver malware of some sort to an end user. At that point

one can fall back to more traditional incident response containment efforts - find the malware and limit its ability to function and spread.

3.4 Eradication

For a more standard incident response situation (e.g. malware on the network), the Eradication phase is where the response team cleans up the malware. It is very important that containment happen first and be thorough and complete before beginning eradication efforts or reinfection will undermine the process. However, eradication within the context of a BGP hijack is blurred with containment. Eradication of the BGP hijack requires that the false route advertisements stop. As already discussed, this means administrative intervention from the source ISP and therefore the ability to contact that ISP and convince them to take such action.

3.4.1 Secondary Attacks

Eradication of the effects of secondary attacks following a BGP hijack does not significantly differ from standard incident response steps in the Eradication phase except insofar as it may be difficult to define the full scope of exposure. Strong execution of the Identification phase will help to make Eradication phase more effective.

3.5 Recovery

The Recovery step of the 6 steps of the incident response process revolves around restoring business processes and ideally fixing the vulnerability. When it comes to BGP hijacks, especially AS path manipulations which do not lead to denial of service, it can be difficult for business to perceive this malicious activity as there is no obvious disruption to the network. Restoration of service is therefore not an issue, and since BGP configuration and BGP security happens largely outside an organization's network, there is little to be done as part of the incident response process for the Recovery phase.

3.5.1 Secondary Attacks

Secondary impacts from the BGP hijack such as delivery of malware may have recovery needs of their own of course, but they do not differ from the standard incident response efforts at this phase.

Tim Collyer, tccollier@gmail.com

3.6 Lessons Learned

Every incident is an opportunity. Having identified that an incident has occurred, the incident response team has an opportunity to improve the security posture of an organization by addressing previous flaws now discovered in the course of an incident (especially in the Recovery phase). However, just because the Recovery phase in a BGP hijack incident has such limited scope, does not also mean that this is true of the Lessons Learned phase.

Successfully perpetrating a BGP hijack which does not noticeably disrupt network traffic requires a high level of understanding as well as access to a well-positioned BGP router, probably maintained by an ISP. The NSA is not the only organization capable of this kind of attack (Weaver, 2014), but there *is* a degree of sophistication and required resources which places such attacks out of reach of most of the ne'er-do-wells on the Internet. Therefore one of the lessons a victim can take away from such an attack is just the fact that they qualify to be a victim. Knowing that you are a target of an adversary capable of such attacks helps to underscore the need to have defenses capable of meeting attacks from an adversary of such skill.

3.6.1 Secondary Attacks

As has been previously stated, the BGP hijack itself is only one step along the way toward the actual goal of the attack. Detection of a BGP attack and careful monitoring of subsequent actions can yield valuable information about the purpose behind the attack. Now the victim knows not only the caliber of an adversary that may be targeting and attacking them, but what the attacker's goal may be. And that's the true value - once an organization knows what attackers find valuable, it becomes much easier to know where to apply additional security resources or to confirm that existing security priorities are correct.

4. Conclusion

Examining BGP and its vulnerabilities from the perspective of incident response paints a grim picture. The protocol has fundamental security flaws in that it makes assumptions about running in a trusted environment. While there are some preventative controls, they are only somewhat effective or else require widespread adoption (currently

Tim Collyer, tccollier@gmail.com

incomplete) across the Internet in order to be effective. Much of the response process, as it relates to BGP directly, involves communication and actions on the part of third parties, outside the control of incident response teams.

After a BGP hijack, the next steps in an attack which might leverage a man-in-the-middle position have a steep set of requirements which need to be in place *before* an attack in order for an incident response team to be effective. These include policy and legal work, significant hardware investment, personnel training, software purchase or script development, etc. All of this represents a not insignificant financial investment which may or may not be available, especially for smaller organizations.

Any organizations with a connection to the Internet endures many types of malicious attacks every day. Phishing campaigns are a good example of an attack which offers a constant barrage of events, alerts, and incidents for security teams to handle. BGP hijacks, however, seem to be less prevalent or at least less monitored as evidenced by the fact that common yearly threat summaries such as the Verizon DBIR Report (www.verizonenterprise.com/resources/reports/rp_DBIR_2016_Report_en_xg.pdf) and McAfee Labs Threat Report (<http://www.mcafee.com/us/resources/reports/rp-quarterly-threats-mar-2016.pdf>) never mention BGP even once. While a BGP attack that ends in a denial of service may catch an organization's attention, the more sophisticated BGP hijacks can avoid significant network disruption and therefore attention. What security teams are left with is a low frequency attack with limited preventative measures available. The attack may not be connected to obvious impact to an organization but it costs a fair amount of time and money to put the required elements in place in order to fully respond to the secondary objectives of the attack. If (detected) impact drives budget, defenses for BGP MitM attacks may be difficult to fund.

If costly and difficult defenses are the rock, the hard place is that if an organization chooses to ignore the possibility of BGP hijacks because of the presumed comparatively low frequency and the high cost of risk reduction, they leave themselves open to the attacks. Without proper monitoring, an organization may never know if an attack happened, and without appropriate incident response capabilities, the company would not know what the intent was behind the attack or if it were successful. Due to the powerful position of a MitM attacker, the impact to an organization could be extremely

Tim Collyer, tccollier@gmail.com

significant. The far end of the spectrum for impacts could definitely lead to the loss of the company or worse. While some organizations might be more likely targets than others, in general, most companies will find themselves in a high risk, high cost, low probability state, which is just about the worst section of a risk matrix to be in when it comes to deciding how to apply limited funding and other resources.

The good news is that the other efforts of security teams are helpful when dealing with secondary attacks resulting from a BGP hijack. Malware detection, behavioral monitoring, lateral movement detection, and other measures will continue to be effective in detecting some actions which may be taken by attackers during or after a successful BGP hijack. Moreover, BGP hijack response is hardly the only reason to implement elements like full packet capture at the network perimeter. There are numerous other situations in which a security team could benefit from having a record of the packets that flowed in and out of a network which makes the justification for implementing such capabilities much easier. Similarly, there are other benefits from including the security team in contract negotiations with an ISP (or other service providers) besides just the BGP use-case.

There are several key takeaways from the consideration of BGP attacks and associated incident response. By far the best situation is if such attacks can be prevented outright, whether through efforts such as RPKI or other mechanisms. Once an attack on BGP is successful, a victim is much more at the mercy of the attacker and response efforts are difficult and messy. However, by far the most important phase of BGP-related incident response is the Preparation phase. If an organization has not taken steps to consider how to handle such a situation beforehand and has not taken steps to have the technical capabilities in place ahead of time, there is little that can be done except to shrug and hope that it doesn't happen to you. And with that kind of approach, you wouldn't know it was happening until it was far too late anyway.

Acknowledgements

I would like to thank the following people for their instrumental help in the research for this paper: Girithar Anthay Suthakaran, Alex Pulver, Donny Hubener, and Awake Networks.

References

- 2016 Data breach investigations report*. (2016). Retrieved from Verizon website:
www.verizonenterprise.com/resources/reports/rp_DBIR_2016_Report_en_xg.pdf
- BGPStream event #52776. (2016, August). Retrieved from
<https://bgpstream.com/event/52776>
- BGPStream event #54711. (2016, September). Retrieved from
<https://bgpstream.com/event/54711>
- Computer incident response guidebook* (NAVSO P-5239-19). (1996). Retrieved from Department of the Navy website:
<http://trygstad.rice.iit.edu:8000/Government%20Documents/Navy-Marine/NAVSO%20P5239-19%20CIRT%20Guide.pdf>
- Cowie, J. (2013, November 19). The new threat: Targeted internet traffic misdirection. Retrieved from <http://research.dyn.com/2013/11/mitm-internet-hijacking/>
- Gavrichenkov, A. (2015, August). Breaking HTTPS with BGP hijacking. Retrieved from <https://www.blackhat.com/docs/us-15/materials/us-15-Gavrichenkov-Breaking-HTTPS-With-BGP-Hijacking-wp.pdf>
- Grand, J. (1998, May 19). Hackers Testifying at the United States Senate, May 19, 1998 (L0pht Heavy Industries) [Video file]. Retrieved from https://www.youtube.com/watch?v=VVJldn_MmMY
- Goodin, D. (2015, September 21). Symantec employees fired for issuing rogue HTTPS certificate for Google. Retrieved from <http://arstechnica.com/security/2015/09/symantec-employees-fired-for-issuing-rogue-https-certificate-for-google/>

Tim Collyer, tccollier@gmail.com

- Hawkinson, J., & Bates, T. (1996, March). RFC 1930 - Guidelines for creation, selection, and registration of an Autonomous System (AS). Retrieved from <https://tools.ietf.org/html/rfc1930>
- Hepner, C., & Zmijewski, E. (2009, February). Defending against BGP man-in-the-middle attacks. Retrieved from <https://www.blackhat.com/presentations/bh-dc-09/Zmijewski/BlackHat-DC-09-Zmijewski-Defend-BGP-MITM.pdf>
- Hummel, S. (2013, February 27). Routing protocol selection guide - IGRP, EIGRP, OSPF, IS-IS, BGP. Retrieved from <https://supportforums.cisco.com/document/127851/routing-protocol-selection-guide-igrp-eigrp-ospf-bgp>
- Huston, G., & Michaelson, G. (2012, February). RFC 6483 - Validation of route origination using the resource certificate PKI and ROAs. Retrieved from <https://tools.ietf.org/html/rfc6483>
- Litke, P., & Stewart, J. (2014, August 7). BGP Hijacking for Cryptocurrency Profit. Retrieved from <http://www.secureworks.com/cyber-threat-intelligence/threats/bgp-hijacking-for-cryptocurrency-profit/>
- Krebs, B. (2016, September). Alleged vDoS proprietors arrested in Israel. Retrieved from <http://krebsonsecurity.com/2016/09/alleged-vdos-proprietors-arrested-in-israel/>
- Madory, D. (2016, September 20). BackConnect's suspicious BGP hijacks. Retrieved from <http://research.dyn.com/2016/09/backconnects-suspicious-bgp-hijacks>
- NIST. (n.d.). RPKI deployment monitor. Retrieved September 17, 2019, from <https://rpki-monitor.antd.nist.gov/>
- Protecting Border Gateway Protocol for the Enterprise - Cisco. (n.d.). Retrieved from <http://www.cisco.com/c/en/us/about/security-center/protecting-border-gateway-protocol.html>
- Rekhter, Y., & Li, T. (1995, March). RFC 1771 - A Border Gateway Protocol 4 (BGP-4). Retrieved from <https://tools.ietf.org/html/rfc1771>
- Rekhter, Y., Li, T., & Hares, S. (2006, January). RFC 4271 - A Border Gateway Protocol 4 (BGP-4). Retrieved from <https://tools.ietf.org/html/rfc4271>
- Remes, W. (2015, August). Internet plumbing for security professionals: the state of BGP security. Retrieved from <https://www.blackhat.com/docs/us-15/materials/us-15->

Tim Collyer, tccollier@gmail.com

Remes-Internet-Plumbing-For-Security-Professionals-The-State-Of-BGP-Security-wp.pdf

- RIPE Network Coordination Centre. (2008). YouTube hijacking: a RIPE NCC RIS case study. Retrieved from <https://www.ripe.net/publications/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study>
- Townsend, B. (2016, September 13). nanog: Re: "Defensive" BGP hijacking? [Electronic mailing list message]. Retrieved from <http://seclists.org/nanog/2016/Sep/122>
- US Navy. (1996, August). Computer incident response guidebook. Retrieved from <http://trygstad.rice.iit.edu:8000/Government%20Documents/Navy-Marine/NAVSO%20P5239-19%20CIRT%20Guide.pdf>
- Vervier, P., Thonnard, O., & Dacier, M. (2015). Mind Your Blocks: On the Stealthiness of Malicious BGP Hijacks. Retrieved from http://www.internetsociety.org/sites/default/files/NDSS2015_Mind_Your_Blocks_Stealthiness_Malicious_BGP_Attacks.pdf
- Wagner, A. (2013, September 20). ext4: Mysterious “No space left on device”-errors. Retrieved from <http://blog.merovius.de/2013/10/20/ext4-mysterious-no-space-left-on.html>
- Weaver, N. (2014, February 13). A close look at the NSA’s most powerful Internet attack tool. Retrieved from <http://www.wired.com/2014/03/quantum>
- Zetter, K. (2008, July 26). Revealed: the Internet’s biggest security hole. Retrieved from <https://www.wired.com/2008/08/revealed-the-in>

Tim Collyer, tccollier@gmail.com