

A Forensic Case Study on AS Hijacking: The Attacker’s Perspective

Johann Schlamp
TU München
Dept. of Computer Science
schlamp@in.tum.de

Georg Carle
TU München
Dept. of Computer Science
carle@in.tum.de

Ernst W. Biersack
Eurecom
Sophia Antipolis
erbi@eurecom.fr

ABSTRACT

The Border Gateway Protocol (BGP) was designed without security in mind. Until today, this fact makes the Internet vulnerable to hijacking attacks that intercept or blackhole Internet traffic. So far, significant effort has been put into the detection of IP prefix hijacking, while AS hijacking has received little attention. AS hijacking is more sophisticated than IP prefix hijacking, and is aimed at a long-term benefit such as over a duration of months.

In this paper, we study a malicious case of AS hijacking, carried out in order to send spam from the victim’s network. We thoroughly investigate this AS hijacking incident using live data from both the control and the data plane. Our analysis yields insights into how an attacker proceeded in order to covertly hijack a whole autonomous system, how he misled an upstream provider, and how he used an unallocated address space. We further show that state of the art techniques to prevent hijacking are not fully capable of dealing with this kind of attack. We also derive guidelines on how to conduct future forensic studies of AS hijacking. Our findings show that there is a need for preventive measures that would allow to anticipate AS hijacking and we outline the design of an early warning system.

Categories and Subject Descriptors

C.2.0 [Computer-Communication Networks]:
General—*Security and protection*; C.2.3 [Computer-
Communication Networks]: Network Operations—
Network monitoring

General Terms

Security, Measurement, Experimentation

Keywords

AS Hijacking, Prefix Hijacking, BGP, Monitoring, Case Study

1. INTRODUCTION

The Internet is a loose federation of heterogenous and independent autonomous systems (AS). Data packets between end hosts traverse multiple ASes, whose border routers are connected by the Border Gateway Protocol (BGP). BGP was initially designed for a small group of participants, but not for sustaining a world-wide network of networks with billions of hosts.

Until today, routing in the Internet is still vulnerable to a variety of attacks. This paper focuses on specific weaknesses

based on the lack of resource ownership validation. Network resources like IP address blocks or AS numbers are managed by one of the five regional Internet registrars (RIR). If an institution applies for resources and is able to prove its need, the responsible RIR will allocate the requested resources. Some RIR policies dictate at least two upstream providers for an AS in order to ensure reliable operation. A formless *letter of authorization* is often accepted by Internet service providers (ISP) as a legitimation to advertise a customer’s AS and its resources. To check authenticity, RIR-operated databases can be queried by *whois clients*. These databases hold information about Internet resources and corresponding resource holders. Such resource objects cannot be modified without valid access credentials and are thus used for ownership validation in practice.

RIR registration systems do not fully prevent attackers from claiming ownership of a victim’s prefixes and injecting such claims into the global routing system. Inattentive upstream providers might accept a customer’s prefix announcements without proper origin validation. Besides, resource objects in RIR databases are maintained by the resource holders themselves and are thus prone to be incomplete. It further may be possible for an attacker to gain control over such resource objects. With the Resource Public Key Infrastructure (RPKI) [10] deployed globally in the future, reliable origin validation will be enabled.

In this paper, we focus on elaborate *AS hijacking* attacks. This type of attack allows an attacker to claim ownership of a whole autonomous system and its prefixes despite origin validation. So far, no corresponding incident has been studied in detail, in contrast to well-known IP prefix hijacking incidents. We overcome this lack of evidential data by thoroughly analyzing a real case of AS hijacking. Our analysis yields insights into the attacker’s proceeding, which gives us the possibility to infer tangible goals and relevant preconditions that enabled the attack. We further discuss limitations of state of the art prevention techniques and prove the need for additional preventive measures. Our understanding of the intention behind the studied attack leads to the design of an early warning system with preventive capabilities.

We organize our paper as follows: Section 2 gives a technical overview of AS hijacking and the difference to common IP prefix hijacking. In Section 3, we evaluate the applicability of current prevention techniques and review the RPKI’s potential to prevent AS hijacking. A detailed case study of a real AS hijacking attack follows in Section 4. We outline the design of an early warning system to prevent AS hijacking in Section 5, and conclude in Section 6.

2. AS HIJACKING

AS hijacking attacks aim at the impersonation of a victim’s organization. The motivation behind this type of attack is a malicious use: activities conducted with the hijacked networks are masked and appear to be carried out by the victim itself. Such attacks are characterized by an attacker announcing the victim’s prefixes originating at the victim’s AS, which is called a last hop attack in literature [9]. Technically, this can be achieved by pretending to own the victim’s AS, or by pretending to peer with it (i.e. by announcing a false link, which implies access to another AS). In contrast, ordinary IP prefix hijacking incidents often arise from misconfiguration or are motivated by a use without the need for deception, e.g. to blackhole networks or to intercept traffic. In general, illicit announcements propagate via an upstream provider, which implies that an attacker is able to pass or to avoid its ownership validation mechanism.

Prefix hijacking attacks lead to noticeable changes in the Internet topology. Hijacked prefixes originate from both the victim’s AS and the attacker’s AS, which is called a *multi-origin AS* (MOAS). AS hijacking attacks by contrast only add another upstream link to the victim’s AS. Figure 1 shows the topological differences between prefix hijacking and AS hijacking. MOAS are generally considered suspicious, although valid causes exist. Multiple upstream links for an AS however do not create suspicion in general. In some RIR regions, policies even enforce a newly established AS to be connected to at least two upstream providers.

Pretending to own a victim’s AS and establishing a fraudulent business relationship with an upstream provider is surprisingly easy. From a technical point of view, a leased server running a software BGP router is sufficient to operate a hijacked AS. Payment can be arranged anonymously and even the technical setup does not depend on face-to-face interaction. In order to enable an upstream provider to announce prefixes originating at a customer’s AS, a formless *letter of authorization* from the legitimate holder is sufficient in general, which can be easily forged by an attacker. To assure authenticity, the attacker might use means of social engineering or try to get hold of the resources’ RIR database objects. Control over those objects is generally considered a proof of prefix ownership. It can be gained by convincing a RIR of recent changes in responsibility for the resources in question (e.g. with forged papers of a company acquisition), by exploiting flaws in the RIR database software, or by taking over the victim’s DNS domain.

According to [3] for instance, an attacker fraudulently authorized an upstream provider to announce a victim’s prefix. In 2003, a large US defence company failed to renew the registration of a DNS domain associated with one of its prefixes. The attacker re-registered that domain to gain control over the mail address that was provided in ARIN’s RIR database. After using it to prove prefix ownership to a U.S. upstream provider, the attacker massively sent spam from the hijacked prefix. It took the victim more than two months to notice and to counter the attack. Note that the burden of proof for such attacks lies with the victim.

Hijacking whole ASes in addition to single prefixes masks technical evidence due to untampered prefix origins. This can greatly extend the lifetime of an attack. In the following, we discuss state of the art prevention techniques and show that prevention of AS hijacking is not fully covered yet.

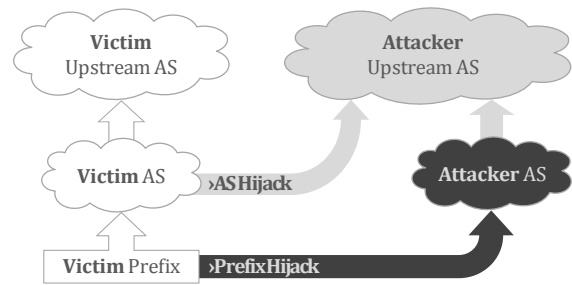


Figure 1: Differences between AS and prefix hijacking

3. RELATED WORK

3.1 Detection techniques

Related work mostly focuses on the detection of hijacking attacks and the application of counter-measures after successful detection. The Prefix Hijack Alert System (PHAS) [9] monitors global route changes by aggregating multiple BGP feeds to effectively detect multi-origin ASes. Several techniques have been proposed to decrease its rate of false positives (e.g. [14]). A complementary approach is taken by the iSPY detection system [18], which is locally deployed at an ISP’s network. It is based on the observation that an ongoing hijacking attempt will lead to significant loss of connectivity for the victim’s hosts. There is also a distributed technique to detect hijacking incidents without a victim’s cooperation [19]. It is based on the evaluation of hop count metrics as a result of traceroute measurements. In [5], the combination of BGP monitoring to identify MOAS target lists and active network fingerprinting carried out from multiple vantage points was shown to be beneficial. Further efforts, e.g. to locate the attacker [15], can be applied as a second step after the detection of an attack.

3.2 Prevention techniques

To prevent hijacking attacks, much effort has been put into the improvement of BGP security. Concepts like S-BGP [8], soBGP [17] or psBGP [13] proposed cryptographic means to attest a valid route origin with so-called *route origin authorizations* (ROA), but were not deployed in practice.

BGPsec [11] is the latest approach to secure BGP, developed by the secure inter-domain routing (SIDR) working group. It specifies a ROA infrastructure for route origin validation (RPKI) [10] and additional components for AS path validation. While the RPKI is well advanced and already deployed by RIRs in prototypical environments, path validation is still at an early stage [6].

Limitations of RPKI: The current design of RPKI is based on route origin attestations and incapable to prevent AS hijacking due to untampered origins: ROAs generally validate both the victim’s and the attacker’s route origin. Nevertheless, following best practices can make an AS unattractive for covertly acting attackers. An IETF draft document recommends to create ROA for unused prefixes bound to the AS number 0, which would effectively prevent an attacker from hijacking an AS’ unannounced prefixes (see [12], Section 3.7). This however does not prevent an attacker to hijack an AS and its already announced prefixes only. The draft further states that *adjacency validation* is beyond scope (see [12], Section 4). We suggest to develop additional prevention techniques, and, if RPKI is extended in the future, to provide **mutual** adjacency attestation objects [7] until BGPsec router certificates are in place.

3.3 Forward-confirmed reverse DNS

Forward-confirmed reverse DNS (FCrDNS) checks are often used for blacklisting or whitelisting IP addresses. Control over reverse DNS can thus be beneficial for an attacker to abuse hijacked networks. A variety of service implementations, e.g. for mail, SSH, FTP and IRC, are capable to perform FCrDNS checks. Failures are often reported by default¹, or might even be used to block connections. Due to valid reasons for FCrDNS checks to fail, it is discouraged to block requests solely based on this criterion, although some operators decide to do nevertheless. However, using FCrDNS checks for whitelisting purposes is wide-spread: validated hosts are generally considered trustworthy.

FCrDNS checks match ownership of both a domain name and the corresponding authoritative reverse DNS server. Such checks are carried out by doing a reverse DNS lookup on a client's IP address to obtain a list of PTR records (i.e. domain names) and a forward lookup on each of the resulting names. If the client's IP address can be mapped to one of those results, the check is passed (RFC5451).

Reverse DNS lookups are delegated by RIRs to name servers of prefix holders. Redelegation requires access to the RIR database (which we do not assume for an attacker). If however the delegation points to a server name within an expired domain re-registered by an attacker, or if its IP address is located within a hijacked prefix, control over reverse DNS and the ability to pass FCrDNS checks is gained as a side-effect of the attack. Note that DNSSEC (RFC4033) effectively protects against this threat.

4. CASE STUDY: "LINKTEL INCIDENT"

4.1 Introduction

We study a real case of long-term AS hijacking that was carried out in order to send spam. Our analysis thereby complements a recent study on spammers, which connects spam to short-lived prefix hijacking events [16].

On August 20, 2011 a representative of the Russian ISP Link Telecom (AS31733) sent a distress mail to the North American Network Operators' Group (NANOG)². The subject of this message was a suspected prefix hijacking observed at the ISP Internap (AS12812) located in the USA. The author explained that Link Telecom had struggled with the financial crisis and got almost bankrupt, but was now on the verge of recovery due to a new investor. While trying to get their prefixes back online, Link Telecom's operators recognized massive blocking of the ISP's traffic and learned that all prefixes were listed on Spamhaus [1] spam blacklists. We refer to this event as the "LinkTel incident".

According to the author, Internap received a forged letter of authorization from Link Telecom on June 9, 2011 and started to advertise routes to AS31733 and its prefixes. By end of July 2011, Link Telecom closely investigated the issue, found its prefixes announced via Internap, and complained to this ISP. Internap consequently referred to a valid letter of authorization and refused to take actions. In addition, Link Telecom apparently was contacted by a person claiming ownership of the prefixes in question:

"(...) someone named Michael Lindsay contacted us and said it is his network!"³

¹ e.g. "sshd: reverse mapping checking failed - POSSIBLE BREAK-IN ATTEMPT!"

² <http://mailman.nanog.org/pipermail/nanog/2011-August/039379.html>

³ <http://mailman.nanog.org/pipermail/nanog/2011-August/039568.html>

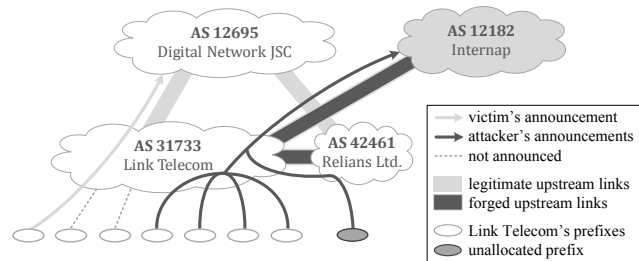


Figure 2: Resulting AS topology after the hijacking phase

On August 25, Link Telecom informed Internap's upstream providers (which at that time were Telia, Cogent, NTT, Global Crossing and Tinet), and started announcing more specific prefixes while redelegating reverse DNS. After Tinet (AS3257) and NTT (AS2914) started to filter out illicit announcements on August 29, the attack ended on August 30.

While this story already gives an overview of the attacker's proceeding, we are able to disclose the full course of action by studying archived control plane data for the corresponding period of time. We further analyze manually obtained meta data to complement our findings. The evaluation of data plane information finally allows us to understand the attacker's intention in great detail.

4.2 Control plane analysis

We analyzed publicly available data feeds of RouteViews Oregon's BGP router [2] from April 1 till September 30, 2011⁴. An animated visualization of all BGP path changes in chronological order can be found on our website⁵.

We learned that the attack started on April 15, 2011 with the announcement of 188.164.0.0/16. Within eight weeks, the attacker gradually took over further parts of Link Telecom's networks (namely the prefixes 94.250.128.0/18, 83.223.224.0/19 and 46.96.0.0/16). Note that the final announcement happened on June 9, which is the same day when Internap supposedly received the forged letter of authorization. This indicates that Internap announced the attacker's routes for two months without formal authorization.

The attacker spared Link Telecom's remaining prefixes 86.59.224.0/17, 79.174.128.0/18 and 94.250.192.0/18. The first of these prefixes was partially announced by Link Telecom before and during the attack. All three prefixes were never touched by the attacker. Figure 2 shows the resulting AS topology including all events described so far.

On May 12, the attacker further attempted to announce an unallocated address space (89.145.168.0/21). For that, a second AS was hijacked and used to originate routes via the already hijacked AS: Relians Ltd. (AS42461). At that time Relians Ltd. was connected to the same Russian upstream provider as the one used by the victim, Digital Network JSC (AS12695), see Figure 2 for details. We assume that this AS was used as a decoy to test announcing unallocated space without risking the whole attack. On July 11, this prefix was globally withdrawn. This indicates that the attacker – or someone else along the topological way – suddenly decided to stop routing that prefix. Our finding confirms that invalid routes can still leak into the global routing system [4].

Link Telecom started to recover on August 24 by announcing its prefixes at full size and also more specific prefixes on

⁴ We parsed 17,573 BGP dump files and obtained 12,751 update messages related to the LinkTel incident.

⁵ <http://hijacking.net.in.tum.de/>

HIJACKING PHASE	
Mar 12	The attacker re-registers the recently expired DNS domain link-telecom.biz to hijack AS31733 (Link Telecom)
Apr 15	The attacker announces prefix 188.164.0.0/16 originating at AS31733 via AS12812 (Internap)
May 06	The attacker announces prefix 94.250.128.0/18 (2x /19)
May 12	The attacker hijacks AS42461 (Relians Ltd.) and uses it to announce the unallocated prefix 89.145.168.0/21
May 28	The attacker announces prefix 83.223.224.0/19
Jun 09	The attacker announces prefix 46.96.0.0/16
Jun 09	<i>Internap receives a faked letter of authorization</i>
PRODUCTIVE PHASE	
Jul 11	Global withdraw of unalloc. prefix 89.145.168.0/21
Jul 28	Spamhaus blacklists all remaining hijacked prefixes
RECOVERY PHASE	
Aug 11	<i>Link Telecom sends complaints to Internap</i>
Aug 24	Spamhaus starts to close spam cases
Aug 24	Link Telecom announces all prefixes at full size
Aug 25	<i>Link Telecom sends complaints to upstream ISPs and redelegates reverse DNS</i>
Aug 28	Link Telecom announces more specific prefixes
Aug 29	<i>Link Telecom receives responses to complaints</i>
Aug 30	NTT and Tinet withdraw routes to hijacked prefixes
Sep 03	BGP converges, no further related events

Table 1: Full disclosure of the LinkTel incident (2011). *Italic entries* could not be confirmed by our analyses.

August 29. As a consequence, some of the routes to the hijacked networks were withdrawn. On August 30, major upstream providers (including NTT and Tinet) withdrew further routes to these networks. The remaining routes finally vanished 141 days after the first illicit announcement. The attacker never tried to announce more specific prefixes and did not fight back in any other observable way.

4.3 Analysis of meta data

There are few archives of additional historical data, which makes it difficult to carry out forensic analyses. Note that the owner of Link Telecom’s AS and prefixes as well as the unallocated address space and the decoy AS’s operating company have changed since the attack. We were able to obtain an accurate view on the attack’s time frame by searching the Internet for evidential data. We found a RIPE database dump from June 3, 2011 on the UK mirror service⁶ providing the necessary data. To allow for future analyses, we suggest to archive RIR database dumps where available.

We found evidence that the attacker deceived Internap of being authorized to advertise the hijacked resources. Link Telecom’s DNS domain link-telecom.biz was taken over by the attacker as reported on NANOG’s mailing list. The aforementioned RIPE database dump revealed corresponding mail addresses associated to Link Telecom’s resources (i.e. implicitly given in the resources’ *changed* attributes⁷). A DNS whois lookup still shows that the domain expired on March 11, 2011, and was re-registered 6 hours later by a proxy registrar protecting the buyer’s identity. Note again that for incidents lying further in the past, DNS registration data is inaccessible due to the lack of historical archives.

The database dump also revealed that reverse DNS was delegated to name servers within the re-registered domain. By assigning `{ns1|ns2}.link-telecom.biz` to a malicious

host, all reverse DNS queries for the hijacked networks could have been intercepted by the attacker. This implies that the attacker had the capability to pass FCrDNS checks.

Finally, we discovered the exact point in time at which the prefixes first appeared on Spamhaus blacklists. Note that this information is not available anymore due to limited long-term archives. We found a discussion on the RIPE abuse working group’s mailing list⁸. Its subject is beyond the scope of this paper – however, the initial message held a copy of the latest Spamhaus listings. This excerpt shows that all hijacked prefixes were blacklisted on July 28, 2011. According to the Spamhaus Register Of Known Spam Operations (ROKSO)⁹, all spam cases were closed between August 24 and September 8. This indicates that the attacker did not send spam after August 24, which correlates with our control plane results and the victim’s efforts to counteract.

Up to this point, we learned the attacker’s moves to hijack Link Telecom’s AS. Table 1 summarizes all events and identifies specific attack phases. In the next section, we study the attacker’s objectives in more detail.

4.4 Data plane analysis

In order to better understand the purpose of the attack and to confirm an abusive use of the hijacked AS, it is helpful to analyze traffic related to the attack. We have access to archived netflow data of the Münchner Wissenschaftsnetz (MWN) – Munich’s scientific network – which connects more than 80,000 end hosts. It is used by researchers, students, and administrative personnel, who generate monthly upstream and downstream traffic volumes of more than 300 and 600 Terabyte, respectively.

From April 19 till September 2, 2011, we extracted a total of 603 flows related to the LinkTel incident. No further flows were observed for at least one month before and after this time frame, i.e. random traffic from spoofed IP addresses is unlikely. In Figure 3, we see a correlation between the attack phases identified in Table 1 and flows originating from the hijacked prefixes. Note that outgoing traffic that is directed towards unannounced prefixes before the attack (e.g. by end of April) has been analyzed and traced back to research activities on one of MWN’s planetlab nodes.

Not a single flow left the MWN in response to incoming traffic from the unallocated prefix (marked by (ua) in Figure 3), which indicates a blocking of traffic to illicit destinations at our site. As described above, massive BGP withdraw messages were received for the unallocated space on July 11. On that day, we also observed the last incoming flow from the corresponding prefix.

We already know that the attacker sent spam from the hijacked networks because all prefixes were blacklisted by Spamhaus on July 28. However, the attacker’s first announcement was received on April 15, and we observed the first related flow on April 18. It might be the case that Spamhaus’ blacklisting techniques take a significant amount of time to detect spam operations, but the attacker could also have carried out other activities without attracting attention within the three months between the initial hijacking and the blacklisting. To reconstruct such actions, we take a closer look on the ports observed in the attacker’s traffic.

Figure 4 shows the frequency of ports observed in all extracted flows. These ports were equally observed in traffic

⁶<http://www.mirrormirror.org/sites/ftp.ripe.net/ripe/dbase/split/>

⁷RIPE database dumps do not include further individual-related attributes

⁸<http://lists.ripe.net/pipermail/anti-abuse-wg/2011-July/000838.html>

⁹http://www.spamhaus.org/rokso/sbl_archived/SPM792/zombies




Figure 3: Observed flows within Munich’s scientific network (MWN) from and to hijacked prefixes. Vertical lines show the specific date of illicit BGP announcements. The prefix marked by (ua) was unallocated at that time.

for all hijacked prefixes. A significant fraction of flows represents bidirectional traffic to port 80 (HTTP). Note that incoming flows to ports 6667 (IRC) and 445 (NetBIOS) were hardly answered. IRC traffic might indicate connections to chat channels or even activities related to botnet command and control. Traffic to port 445 is often used to exploit vulnerabilities of Windows for remote code execution.

Traffic with source port 53 (DNS) implies services hosted in the hijacked prefixes, which is backed up by outgoing HTTP traffic. We further observed incoming flows to port 443 (HTTPS). One of these HTTPS connections was established to a webserver under our control. By analyzing its log files, we were able to track the content of the connection and learned, that the attacker created a ticket in our project management system with the following content:

currency trading
<http://theforexsoftwaretrader.com> currency trading

We queried the Internet archive¹⁰ for that website and found a snapshot from July 23, 2011. At that time, this website advertised software and tutorials for forex trading beginners, among others with tips on auto pilot software to automatically make funds and stock trading strategies based on the utilization of “Fibonacci retracement” algorithms.

¹⁰ <http://www.archive.org/>




Figure 4: Observed ports within MWN’s flows from and to hijacked prefixes.

4.5 Lessons learned

Our analysis indicates that, beyond sending spam, the attacker hosted services in the hijacked prefixes, scanned for client vulnerabilities, and placed adverts for questionable products on websites and possibly in IRC channels. We assume that the attacker’s ability to pass FCrDNS checks supported the abusive use of the hijacked networks.

During our analysis, we repeatedly faced the fact that archives for certain historical data were not available. For forensic analyses, access to RIR databases, spam blacklists, and DNS registration data stemming from the time of an attack is indispensable. We suggest to periodically monitor and archive these data sources, and to add further sources, e.g. by deploying spam traps. In addition, we encourage RIRs to provide database snapshots including details about resource holders. A comprehensive data archive allows to study hijacking incidents from the past, and more importantly, it can form the base for an early warning system.

5. AN EARLY WARNING SYSTEM

Our analysis of the LinkTel incident revealed a great deal about the attacker’s proceeding. We observed a spamming attack that was enabled by the following preconditions:

1. The victim’s DNS domain was registered in the RIPE database, and it was going to expire.
2. Sending spam from the victim’s networks was possible due to appropriate reverse DNS delegations.
3. Most of the victim’s prefixes were unannounced without recent changes in BGP activity.

We conclude that the victim has been carefully selected. This was unlikely a manual operation: various data sources had to be combined to assess the victim’s eligibility, which suggests that the attacker had access to automated tools for spotting vulnerable targets. We propose to take this insight into account for the design of an early warning system to prevent similar attacks in the future. Its purpose is to inform vulnerable ASes to deploy counter-measures in time.

Our warning system is based on the evaluation of multiple data sources. First, DNS domains associated to ASes are extracted from RIR database dumps (where available). The expiry dates for those domains are queried by a whois client. ASes with an expiring domain in the near future are




Figure 5: Threat levels for AS hijacking vulnerability

considered at risk: an attacker might be able to re-register that domain and claim ownership of the AS and its resources by sending mails from the operator’s mail address.

The threat level escalates for ASes with vulnerable reverse DNS delegations, i.e. DNSSEC is not deployed and the name servers can be captured. This is the case for servers that are assigned a host name of the expiring domain or for servers with an IP address lying within a prefix of the targeted AS. Both cases can be easily checked by obtaining the reverse DNS delegation for the AS’ prefixes from the RIR database.

The attacker’s proceeding in the LinkTel incident confirmed that hijacking unannounced prefixes of ASes with little BGP activity has great potential for covert attacks. We consider this fact in our warning system by raising the threat level for such ASes. To evaluate BGP activity, we integrate publicly available BGP feeds into our system, and we extract an AS’ unannounced prefixes from the RIR database.

Threat levels will change over time, and have to be repeatedly assessed on a periodical base. Operators of ASes that meet all preconditions can be readily informed about the imminent threat: corresponding mail addresses are already obtained in the first step.

In addition to contacting the victim, ASes with high threat values should be consequently monitored regarding DNS re-registration, suspicious BGP activity and appearance on spam blacklists until re-evaluation yields lowered threat levels. We further suggest the recurrent application of state of the art detection techniques for those ASes.

Similar case studies on future hijacking incidents might identify additional criteria for selecting a victim as well as other purposes of hijacking an AS. These insights could be used to improve our proposed early warning system.

6. CONCLUSION

We studied AS hijacking attacks, which aim at a long-term benefit, and outlined that current prevention techniques are not fully capable to deal with this kind of attack. We further provided forensic evidence for such attacks by thoroughly analyzing a malicious case of AS hijacking. With the evaluation of BGP control plane data and additional meta data, we were able to disclose the attacker’s activities and to reconstruct the full sequence of steps during the attack. We saw that the attack was carried out in a professional manner in order to send spam from the hijacked networks. Detailed studies of data plane information revealed further objectives.

By reflecting on the technical insights we gained from the attack, we profiled the attacker and understood that he must have had access to automated tools in order to detect potential victims. Given this fact, we outlined the design of an early warning system for AS hijacking to prevent such incidents in the future.

7. ACKNOWLEDGEMENTS

We thank Pierre-Antoine Vervier for introducing us to the LinkTel incident, and Lothar Braun for providing us access to archived MWN netflow data.

This work is partly supported by the German BMBF within the project Peeroskop (<http://peeroskop.realmv6.org>).

The research of Ernst Biersack is supported by the European Commission’s Seventh Framework Programme (FP7 2007-2013) under grant agreement no. 257495 (VIS-SENSE).

8. REFERENCES

- [1] The Spamhaus project. <http://www.spamhaus.org/>.
- [2] University of Oregon RouteViews project. <http://www.routeviews.org/>.
- [3] L. Benkis. Practical BGP security: architecture, techniques and tools, 2008. http://www.renesys.com/tech/notes/WP_BGP_rev6.pdf.
- [4] N. Feamster, J. Jung, and H. Balakrishnan. An empirical study of “bogon” route advertisements. *ACM SIGCOMM CCR '05*, pages 63–70, 2005.
- [5] X. Hu and Z. M. Mao. Accurate real-time identification of IP prefix hijacking. In *IEEE SP '07*, pages 3–17, 2007.
- [6] G. Huston and R. Bush. Securing BGP and SIDR. *IETF Journal*, pages 1815–1828, 2011.
- [7] G. Huston and G. Michaelson. A profile for AS adjacency attestation objects. *IETF*, 2009.
- [8] S. Kent, C. Lynn, J. Mikkelsen, and K. Seo. Secure border gateway protocol (S-BGP). *IEEE J-SAC '00*, pages 103–116, 2000.
- [9] M. Lad, D. Massey, D. Pei, Y. Wu, B. Zhang, and L. Zhang. PHAS: A prefix hijack alert system. In *USENIX-Security '06*, 2006.
- [10] M. Lepinski and S. Kent. An infrastructure to support secure internet routing. *IETF*, 2012. RFC6480.
- [11] M. Lepinski and S. Turner. An overview of BGPSEC. *IETF*, 2012.
- [12] T. Manderson, K. Sriram, and R. White. Use cases and interpretation of RPKI objects for issuers and relying parties. *IETF*, 2012.
- [13] P. v. Oorschot, T. Wan, and E. Kranakis. On interdomain routing security and pretty secure BGP (psBGP). *ACM TISSEC '07*, 2007.
- [14] J. Qiu and L. Gao. Detecting bogus BGP route information: going beyond prefix hijacking. In *SecureComm '07*, 2007.
- [15] T. Qiu, L. Ji, D. Pei, J. Wang, J. Xu, and H. Ballani. Locating prefix hijackers using LOCK. In *USENIX-Security '09*, pages 135–150, 2009.
- [16] A. Ramachandran and N. Feamster. Understanding the network-level behavior of spammers. In *SIGCOMM '10*, pages 291–302, 2006.
- [17] R. White. Securing BGP through secure origin BGP. *The Internet Protocol Journal*, 2003.
- [18] Z. Zhang, Y. Zhang, Y. C. Hu, Z. M. Mao, and R. Bush. iSPY: Detecting IP prefix hijacking on my own. *IEEE/ACM ToN*, pages 1815–1828, 2010.
- [19] C. Zheng, L. Ji, D. Pei, J. Wang, and P. Francis. A light-weight distributed scheme for detecting IP prefix hijacks in real-time. In *SIGCOMM '07*, pages 277–288, 2007.