How to Prevent AS Hijacking Attacks

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 aspects in mind. This fact makes the Internet vulnerable to attacks: complete networks can be hijacked to blackhole or intercept traffic. In this work, we extend the set of known hijacking attacks with a real case study on *AS hijacking*, carried out in order to send spam from a victim's network. This type of attack is more sophisticated than common IP prefix hijacking, and is aimed at a long-term benefit, with effective use for several months. On our poster, we thoroughly investigate the aforementioned incident based on live data from both the control and the data plane. Our analysis yields insights into the attacker's proceeding to covertly hijack a whole autonomous system, mislead an upstream provider and abuse an unallocated address space. We further discuss the potential for prevention and reveal shortcomings of state of the art BGP security extensions like RPKI. Based on these findings, we outline the concept of an early warning system for AS hijacking with pre-emptive capabilities.

Categories and Subject Descriptors

C.2.0 [**Computer-Communication Networks**]: General—*Security and protection*

Keywords

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

1. INTRODUCTION

The Internet is a federation of autonomous systems (AS). Data packets between end hosts traverse multiple AS connected by the Border Gateway Protocol (BGP). BGP was initially designed for a small group of participants, but not for sustaining a world-wide Internet with billions of hosts.

Until today, routing in the Internet is 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 their 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 customers' AS and its resources. To check authenticity, RIR-operated databases can be queried by *whois clients*. These databases hold information about resource holders that cannot be modified without valid access credentials and are thus the base for ownership validation in practice. However, RIR registration systems do not fully prevent attackers from claiming ownership of a victim's prefixes. With the Resource Public Key Infrastructure (RPKI) [4] deployed globally in the future, reliable origin validation will be enabled.

In our work, we study 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. Based on the presentation of a real case study on our poster, we outline an early-warning system and discuss best practices for RPKI to prevent such incidents.

2. AS HIJACKING

We distinguish between *prefix hijacking* and *AS hijacking*. Prefix hijacking attacks are characterized by an attacker's AS announcing a victim's prefixes. In contrast, AS hijacking is carried out by announcing prefixes on behalf of the victim's AS, which are routed to the attacker's network. Both cases imply that the attacker is able to pass or avoid an upstream provider's ownership validation, e.g. by impersonating the victim's organization.

Prefix hijacking attacks lead to noticable changes in the topology: the prefix originates from two different 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.

MOAS conflicts are generally considered suspicious, although valid causes exist. Multiple upstream links 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.

Establishing a fraudulent business relationship with an upstream provider on behalf of a victim's organization is surprisingly easy. 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

CoNEXT Student'12, December 10, 2012, Nice, France.

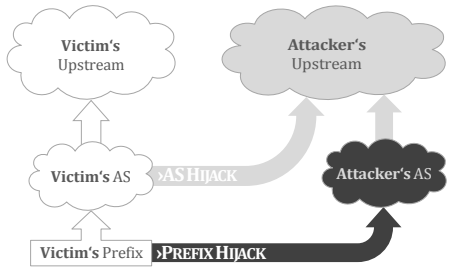

Figure 1: Differences between AS and prefix hijacking

to announce an AS and corresponding prefixes, a formless *letter of authorization* from the legitimate holder is sufficient, 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 ownership, and can be gained by convincing a RIR of recent changes in responsibility for the resources in question (i.e. with forged papers of a company acquisition), exploiting flaws in the RIR database software, or getting hold of the victim's DNS domain.

Hijacking whole AS instead of single prefixes masks technical evidence due to untampered prefix origins. This avoids detection and can greatly extend the lifetime of an attack. It is also much harder to prevent.

3. HIJACKING PREVENTION TODAY

Previous work mostly focused on the analysis and detection of IP prefix hijacking attacks [1, 3, 8, 9]. In contrast, we study techniques to prevent hijacking attacks, in particular for the aforementioned AS hijacking attack.

Much effort has been put into the improvement of security in BGP. In the past, cryptographic concepts were proposed to attest a valid route origin with so-called *route origin authorizations* (ROA), but were not deployed in practice. BGPSec [5] is the latest approach to secure BGP, developed by the SIDR working group. It specifies a ROA infrastructure for route origin validation (RPKI) [4] and additional components for AS path validation, which would effectively prevent both prefix and AS hijacking attacks. While the RPKI is well advanced and partially deployed by RIR, path validation is still at an early stage and not to be deployed soon. RPKI however does not provide adequate means to prevent AS hijacking: ROA validate both the victim's and the attacker's route origin.

4. OUR CONTRIBUTION

On our poster, we present a real case of long-term AS hijacking. This incident has been reported to a NANOG mailing list[1], we refer to it as the *"LinkTel incident"*. We analyzed publicly available BGP data feeds, historical meta data and live traffic within Munich's scientific network for the attack's time period. We learned that the attacker reregistered the victim's expiring DNS domain to prove ownership of its resources. Our analysis further indicates that the attacker massively sent spam, hosted services in the hijacked prefixes, scanned for client vulnerabilities, and placed adverts for questionable products. We also complement a study that connects spam to short-lived hijacking [7].

[1] http://mailman.nanog.org/pipermail/nanog/2011-August/039379.html

4.1 An early-warning system

The LinkTel incident represents a long-term hijacking attack carried out in order to send spam. We thoroughly analyzed preconditions that enabled the attack, and 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. Our poster outlines the design of an early warning system based on the lessons learned from the attack. Its main purpose is to identify AS that are vulnerable to hijacking attacks intended for spamming. This implies to find AS meeting our identified preconditions. We propose a threat escalation model that can be deployed to monitor arbitrary AS, and to readily inform operators about an immediate threat in order to take preventive measures.

4.2 Best practices for RPKI

The current design of RPKI is based on route origin attestations, and is thus incapable to prevent AS hijacking due to untampered origins. Following best practices however 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's unannounced prefixes (see [6], Section 3.7). The draft further states that *adjacency validation* is beyond scope (see [6], Section 4). If RPKI is extended in the future, we suggest to provide **mutual** adjacency attestation objects [2] until router certificates as specified in BGPSec [5] are in place.

5. CONCLUSION

We studied AS hijacking attacks, which aim at a long-term benefit, and outlined that prevention within RPKI is limited. On our poster, we provide forensic evidence for such attacks by thoroughly analyzing a real case of AS hijacking. We profiled the attacker and understood that he must have had access to automated tools for identifying vulnerable targets. Given this fact, we outlined the design of an early warning system for AS hijacking. We further discussed best practices for RPKI to eliminate this threat in the future.

6. REFERENCES

[1] S. Goldberg, M. Schapira, P. Hummon, and J. Rexford. How secure are secure interdomain routing protocols. In *SIGCOMM '10*, pages 87–98, 2010.

[2] G. Huston and G. Michaelson. *A Profile for AS Adjacency Attestation Objects*. IETF, 2009.

[3] M. Lad, D. Massey, D. Pei, Y. Wu, B. Zhang, and L. Zhang. PHAS: A prefix hijack alert system. In *USENIX-Security '06*, 2006.

[4] M. Lepinski and S. Kent. *An Infrastructure to Support Secure Internet Routing*. IETF, 2012. RFC6480.

[5] M. Lepinski and S. Turner. *An Overview of BGPSEC*. IETF, 2012.

[6] T. Manderson, K. Sriram, and R. White. *Use Cases and Interpretation of RPKI Objects for Issuers and Relying*. IETF, 2012.

[7] A. Ramachandran and N. Feamster. Understanding the network-level behavior of spammers. In *SIGCOMM '10*, pages 291–302, 2006.

[8] Z. Zhang, Y. Zhang, Y. C. Hu, Z. M. Mao, and R. Bush. iSPY: Detecting IP prefix hijacking on my own. *IEEE/ACM Transactions on Networking*, pages 1815–1828, 2010.

[9] 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.