SIDR
Internet-Draft
Intended status: Informational
Expires: May 11, 2013

D. McPherson
Verisign, Inc.
S. Amante
Level 3 Communications, Inc.
E. Osterweil
Verisign, Inc.

November 7, 2012

Route Leaks & MITM Attacks Against BGPSEC
draft-foo-sidr-simple-leak-attack-bgpsec-no-help-02

Abstract

This document describes a very simple attack vector that illustrates how RPKI-enabled BGPSEC machinery as currently defined can be easily circumvented in order to launch a Man In The Middle (MITM) attack via BGP. It is meant to serve as input to the IETF's Secure Inter-Domain Routing working group during routing security requirements discussions and subsequent specification.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on May 11, 2013.

Table of Contents

1. Introduction

This document describes a very simple attack vector that illustrates how RPKI-enabled BGPSEC [I-D.ietf-sidr-bgpsec-protocol] machinery, as currently defined, can be easily circumvented in order to launch a Man In The Middle (MITM) attack via BGP [RFC4271]. It is meant to serve as input to the IETF's SIDR Working Group during routing security requirements discussions and subsequent specification.

This draft shows evidence that the attack vector described herein is extremely common, with over 9.6 million candidate instances being recorded since 2007. As a result of this evidence (and additional contextual knowledge), the authors believe the capability to prevent leaks and MITM leak-attacks should be a first-order engineering objective in any secure routing architecture.

While the formal definition of a route leak has proven elusive in the literature, their rampant occurrence and persistent operational threats have proven to be anything but elusive. This document is intended to serve as an existence proof for this threat vector, and any supplementary formal models are left for future work.

2. Discussion

In order to understand how a MITM attack can be launched with this attack vector, assume a multi-homed Autonomous System (AS), AS1, connects to two ISPs (ISP1 & ISP2), and wishes to insert themselves in the data-path between a target network (prefix P) connected to ISP2 and systems in ISP1's network in order to launch a Man In The Middle (MITM) attack. Further, assume that an RPKI-enabled BGPSEC [I-D.ietf-sidr-bgpsec-protocol] as currently defined is fully deployed by all parties in this scenario and functioning as designed.

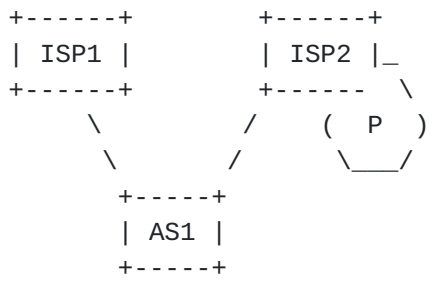

This figure depicts a multi-homed AS1, who is connected to two upstream ISPs (ISP1 and ISP2).

Network operators on the Internet today typically prefer customer routes over routes learned from bi-lateral or settlement free peers. Network operators commonly accomplish this via application of one or more BGP [RFC4271] Path Attributes, most commonly, LOCAL_PREF as illustrated in [RFC1998], that are evaluated earlier in the BGP Path Selection process than AS_PATH length.

As currently defined, [I-D.ietf-sidr-bgpsec-protocol] only provides two functions:

1. Is an Autonomous System authorized to originate an IP prefix?
2. Is the AS_PATH (or any similarly derived attribute such as BGPSEC_Path) in the route the same as the list of ASes through which the NLRI traveled?

In order for an attacker (AS1) to divert traffic from ISP1 for prefix P through their AS they simply fail to scope the propagation of the target prefix P (received from ISP2) by announcing a (syntactically correct) BGPSEC update for prefix P to ISP1. This vulnerability is what the authors refer to as a 'route leak' or a 'leak-attack' (when intent aligns with actions). It is important to note that the default behavior in BGP [RFC4271] is to announce all best paths to external BGP peers, unless explicitly scoped by a BGP speaker through configuration. Because ISP1 prefers prefixes learned from customers (AS1) over prefixes learned from peers (ISP2), they begin forwarding traffic for prefix P destinations through the attacker's AS (AS1). Voila!

It is important to note that the route leaks described herein are NOT 'misorginiations.' Rather, these are cases in which routes are propagated with their authentic origins, but are misdirected through unexpected intermediaries.

It should be understood that any multi-homed AS can potentially launch such an attack, even if through simple misconfiguration, as is a common occurrence today on the Internet. As a matter of fact, advertising these prefixes is the default behavior is many BGP implementations, and explicit action must be taken to not advertise all prefixes learned in BGP. Such occurrences have been historically archived [ROUTE_LEAK_DETECTION_TOOL] and presented to the operational community [NANOG_LEAK_TALK] since 2007. To date, over 9.6 million such events have been recorded and are queriable [ROUTE_LEAK_DETECTION_TOOL]. This corpus serves as a low pass filter, and likely contains some degree of false positives. Thus, while some may debate how many of the occurrences were malicious, or how many were actually real leaks, the corpus itself (and its sheer size) serves as evidence of the large magnitude of this problem. Determination of benign versus malicious intent in these situations

is usually imperceptible, and as such, preventative controls are requisite.

To illustrate the above point, consider the events that transpired on November 6th, 2012 [LEAK_ATTACK_ON_GOOGLE]. On that day a large Internet property (who services hundreds of billions of public end user transactions every day) became unreachable for roughly 27 minutes. At a transaction volume like that, an outage of 27 minutes is a very visible (and likely financially measurable) event. In this case, services became unreachable because a peered AS improperly propagated the impacted party's AS' prefix(s). In leaks such as this, there exists public acknowledgment by the impacted party that [RFC6480] and [I-D.ietf-sidr-bgpsec-protocol] would be unable to detect or remediate this attack.

In an environment where [I-D.ietf-sidr-bgpsec-protocol] is fully deployed, it is expected that there would be high assurances that guard the syntactic integrity of the AS_PATH (or BGPSEC_Path) attribute. As such, one would expect that such an attribute would, indeed, accurately reflect the attacker's AS number in the appropriate location of the AS_PATH; however, it would not prevent an attacker from inserting his AS in the first place. That is, nothing in [I-D.ietf-sidr-bgpsec-protocol] would stop an attacker from launching this type of leak-attack.

Discussion of out of band methods to mitigate this attack are beyond the scope of this document, as its objective is to inform routing protocol design choices currently being considered within the IETF's SIDR Working Group.

3. Acknowledgements

4. IANA Considerations

5. Security Considerations

This document describes an attack on an RPKI-enabled BGPSEC and is meant to inform the IETF Secure Inter-Domain Routing working group on the vulnerability that exists as a result of "leaks" and attacks that conform to this type of behavior.

The authors believe the capability to prevent leaks and leak-attacks should be a first-order engineering objective in any secure routing architecture.

6. Informative References

[I-D.ietf-sidr-bgpsec-protocol]
Lepinski, M., "BGPSEC Protocol Specification", draft-ietf-sidr-bgpsec-protocol-06 (work in progress), October 2012.

[LEAK_ATTACK_ON_GOOGLE]
CloudFlare, CF., "Why Google Went Offline Today and a Bit about How the Internet Works", November 2012, <http://blog.cloudflare.com/why-google-went-offline-today-and-a-bit-about>.

[NANOG_LEAK_TALK]
Mauch, J., "Detecting Routing Leaks by Counting", October 2007, <http://www.nanog.org/meetings/nanog41/presentations/mauch-lightning.pdf>.

[RFC1998] Chen, E. and T. Bates, "An Application of the BGP Community Attribute in Multi-home Routing", RFC 1998, August 1996.

[RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.

[RFC6480] Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", RFC 6480, February 2012.

[ROUTE_LEAK_DETECTION_TOOL]
Mauch, J., "BGP Routing Leak Detection System Routing Leak Detection System", September 2007, <http://puck.nether.net/bgp/leakinfo.cgi>.

Authors' Addresses

Danny McPherson
Verisign, Inc.
12061 Bluemont Way
Reston, VA 20190
USA

Phone: +1 703.948.3200
Email: dmcpherson@verisign.com

Shane Amante
Level 3 Communications, Inc.
1025 Eldorado Boulevard
Broomfield, CO 80021
US

Phone: +1 720.888.1000
Email: shane@level3.net

Eric Osterweil
Verisign, Inc.
12061 Bluemont Way
Reston, VA 20190
USA

Email: eosterweil@verisign.com