Network Working Group
Internet-Draft
Intended status: Best Current Practice
Expires: October 31, 2018

Y. Gilad
S. Goldberg
Boston University
K. Sriram
USA NIST
J. Snijders
NTT
B. Maddison
Workonline Communications

April 29, 2018

The Use of Maxlength in the RPKI
draft-sidrops-rpkimaxlen-00

Abstract

This document recommends that operators avoid using the maxLength attribute when issuing Route Origin Authorizations (ROAs) in the Resource Public Key Infrastructure (RPKI). These recommendations complement those in [RFC7115].

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 https://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 October 31, 2018.

Table of Contents

1. Introduction

The RPKI [RFC6480] uses Route Origin Authorizations (ROAs) to create a cryptographically verifiable mapping from an IP prefix to a set of autonomous systems (ASes) that are authorized to originate this prefix. Each ROA contains a set of IP prefixes, and an AS number of an AS authorized originate all the IP prefixes in the set [RFC6482]. The ROA is cryptographically signed by the party that holds a certificate for the set of IP prefixes.

The ROA format also supports a maxLength attribute. According to [RFC6482], "When present, the maxLength specifies the maximum length of the IP address prefix that the AS is authorized to advertise." Thus, rather than requiring the ROA to list each prefix the AS is authorized to originate, the maxLength attribute provides a shorthand that authorizes an AS to originate a set of IP prefixes.

However, measurements of current RPKI deployments have found that use of the maxLength in ROAs tends to lead to security problems. Specifically, as of June 2017, 84% of the prefixes specified in ROAs that use the maxLength attribute, are vulnerable to a forged-origin subprefix hijack [HARMFUL]. The forged-origin subprefix hijack, as described below, can be launched against any IP prefix that is authorized in ROA but is not originated in BGP. The impact of such an attack is the same as that of a subprefix hijack in the absence of ROA-based protection.

For this reason, this document recommends that, whenever possible, operators SHOULD use "minimal ROAs" that include only those IP prefixes that are actually originated in BGP, and no other prefixes. Operators SHOULD also avoid using the maxLength attribute in their ROAS whenever possible. One ideal place to implement these recommendations is in the user interfaces for configuring ROAs: thus this document further recommends that designers and/or providers of such user interfaces SHOULD provide warnings to draw the user's attention to the risks of using the maxLength attribute.

The recommendations in this document clarify and extend the following recommendation from [RFC7115]:

> One advantage of minimal ROA length is that the forged origin attack does not work for sub-prefixes that are not covered by overly long max length. For example, if, instead of 10.0.0.0/16-24, one issues 10.0.0.0/16 and 10.0.42.0/24, a forged origin attack cannot succeed against 10.0.666.0/24. They must attack the whole /16, which is more likely to be noticed because of its size.

This best current practice requires no changes to the RPKI specification and will not increase the number of signed ROAs in the RPKI, because ROAs already support lists of IP prefixes [RFC6482].

1.1. Requirements

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2. Suggested Reading

It is assumed that the reader understands BGP [RFC4271], the RPKI [RFC6480] Route Origin Authorizations (ROAs) [RFC6482], RPKI-based Prefix Validation [RFC6811], and BGPSEC [RFC8205].

3. Forged Origin Subprefix Hijack

The forged-origin subprefix hijack is relevant to a scenario in which (1) the RPKI [RFC6480] is deployed, and (2) routers use RPKI origin validation to drop invalid routes [RFC6811], but (3) BGPSEC [RFC8205] (or any similar method to validate the truthfulness of the BGP AS_PATH attribute) is not deployed.

We describe the forged-origin subprefix hijack [RFC7115] [GCHSS] using a running example.

Consider the IP prefix 168.122.0.0/16 which is allocated to an organization that also operates AS 64496. In BGP, AS 64496 originates the IP prefix 168.122.0.0/16 as well as its subprefix 168.122.225.0/24. Therefore, the RPKI should contain a ROA authorizing AS 64496 to originate these two IP prefixes. That is, the ROA should be

ROA:(168.122.0.0/16,168.122.225.0/24, AS 64496)

This ROA is "minimal" because it includes only those IP prefixes that AS 64496 originates in BGP, but no other IP prefixes [RFC6907].

Now suppose an attacking AS 64511 originates a BGP announcement for a subprefix 168.122.0.0/24. This is a standard "subprefix hijack".

In the absence of the minimal ROA above, AS 64511 could intercept traffic for the addresses in 168.122.0.0/24. This is because routers perform a longest-prefix match when deciding where to forward IP packets, and 168.122.0.0/24 originated by AS 64511 is a longer prefix than 168.122.0.0/16 originated by AS 64496.

However, the minimal ROA renders AS 64511's BGP announcement invalid, because (1) this ROA "covers" the attacker's announcement (since 168.122.0.0/24 is a subprefix of 168.122.0.0/16), and (2) there is no ROA "matching" the attacker's announcement (there is no ROA for AS 64511 and IP prefix 168.122.0.0/24) [RFC6811]. If routers ignore invalid BGP announcements, the minimal ROA above ensures that the subprefix hijack will fail.

Now suppose that the "minimal ROA" was replaced with a "loose ROA" that used maxLength as a shorthand for set of IP prefixes that AS 64496 is authorized to originate. The "loose ROA" would be:

ROA:(168.122.0.0/16-24, AS 64496)

This "loose ROA" authorizes AS 64496 to originate any subprefix of 168.122.0.0/16, up to length /24. That is, AS 64496 could originate 168.122.225.0/24 as well as all of 168.122.0.0/17, 168.122.128.0/17, ..., 168.122.255.0/24 but not 168.122.0.0/25.

However, AS 64496 only originates two prefixes in BGP: 168.122.0.0/16 and 168.122.255.0/24. This means that all other prefixes authorized by the "loose ROA" (for instance, 168.122.0.0/24), are vulnerable to the following forged-origin subprefix hijack [RFC7115], [GCHSS]:

The hijacker AS 64511 sends a BGP announcement "168.122.0.0/24: AS 64511, AS 64496", falsely claiming that AS 64511 is a neighbor of AS 64496 and falsely claiming that AS 64496 originates the IP

prefix 168.122.0.0/24. In fact, the IP prefix 168.122.0.0/24 is not originated by AS 64496.

The hijacker's BGP announcement is valid according to the RPKI, since the ROA (168.122.0.0/16-24, AS 64496) authorizes AS 64496 to originate BGP routes for 168.122.0.0/24. Because AS 64496 does not actually originate a route for 168.122.0.0/24, the hijacker's route is the *only* route to the 168.122.0.0/24. Longest-prefix-match routing ensures that the hijacker's route to the subprefix 168.122.0.0/24 is always preferred over the legitimate route to 168.122.0.0/16 originated by AS 64496. Thus, the hijacker's route propagates through the Internet, the traffic destined for IP addresses in 168.122.0.0/24 will be delivered to the hijacker.

The forged origin *subprefix* hijack would have failed if the "minimal ROA" described above was used instead of the "loose ROA". If the "minimal ROA" had been used instead, the attacker would be forced to launch a forged origin *prefix* hijack in order to attract traffic, as follows:

The hijacker AS 64511 sends a BGP announcement "168.122.0.0/16: AS 64511, AS 64496", falsely claiming that AS 64511 is a neighbor of AS 64496.

This forged-origin *prefix* hijack is significantly less damaging than the forged-origin *subprefix* hijack. With a forged-origin *prefix* hijack, AS 64496 legitimately originates 168.122.0.0/16 in BGP, so the hijacker AS 64511 is not presenting the *only* route to 168.122.0.0/16. Moreover, the path originated by AS 64511 is one hop longer than the path originated by the legitimate origin AS 64496. As discussed in [LSG16], this means that the hijacker will attract less traffic than he would have in the forged origin *subprefix* hijack, where the hijacker presents the *only* route to the hijacked subprefix.

In sum, a forged-origin subprefix hijack has the same impact as a regular subprefix hijack. A forged-origin *subprefix* hijack is also more damaging than forged-origin *prefix* hijack.

4. Measurements of Today's RPKI

Network measurements from June 1, 2017 show that 12% of the IP prefixes authorized in ROAs have a maxLength longer than their prefix length. The vast majority of these (84%) of these are vulnerable to forged-origin subprefix hijacks. Even large providers are vulnerable to these attacks. See [GSG17] for details.

These measurements suggest that operators commonly misconfigure the maxLength attribute, and unwittingly open themselves up to forged-origin subprefix hijacks.

5. Use Minimal ROAs without Maxlength

Operators SHOULD avoid using the maxLength attribute in their ROAs.

Operators SHOULD use "minimal ROAs" whenever possible. A minimal ROA contains only those IP prefixes that are actually originated by an AS in BGP, and no other IP prefixes. (See Section 3 for an example.)

This practice requires no changes to the RPKI specification and will not increase the number of signed ROAs in the RPKI, because ROAs already support lists of IP prefixes [RFC6482]. See also [GSG17] for further discussion of why this practice will have minimal impact on the performance of the RPKI ecosystem.

5.1. When a Minimal ROA Cannot Be Used?

Sometimes, it is not possible to use a "minimal ROA", because an operator wants to issue a ROA that includes an IP prefix that is sometimes (but not always) originated in BGP.

In this case, the ROA SHOULD include (1) the set of IP prefixes that are always originated in BGP, and (2) the set IP prefixes that are sometimes, but not always, originated in BGP. The ROA SHOULD NOT include any IP prefixes that the operator knows will not be originated in BGP. Whenever possible, the ROA SHOULD also avoid the use of the maxlength attribute.

We now extend our running example to illustrate one situation where where it is not possible to issue a minimal ROA.

Consider the following scenario prior to deployment of RPKI. Suppose AS 64496 announced 168.122.0.0/16 and has a contract with a DDoS mitigation service provider that holds AS 64500. Further, assume that the DDoS mitigation service contract applies to all IP addresses covered by 168.122.0.0/22. When a DDoS attack is detected and reported by AS 64496, AS 64500 immediately originates 168.122.0.0/22, thus attracting all the DDoS traffic to itself. The traffic is scrubbed at AS 64500 and then sent back to AS 64496 over a backhaul data link. Notice that, during a DDoS attack, the DDoS mitigation service provider AS 64500 originates a /22 prefix that is longer than than AS 64496's /16 prefix, and so all the traffic (destined to addresses in 168.122.0.0/22) that normally goes to AS 64496 goes to AS 64500 instead.

First, suppose the RPKI only had the minimal ROA for AS 64496, as described in Section 3. But, if there is no ROA authorizing AS 64500 to announce the /22 prefix, then the traffic-scrubbing scheme would not work. That is, if AS 64500 originates the /22 prefix in BGP during a DDoS attack, the announcement would be invalid [RFC6811].

Therefore, the RPKI should have two ROAs: one for AS 64496 and one for AS 64500.

ROA:(168.122.0.0/16,168.122.225.0/24, AS 64496)

ROA:(168.122.0.0/22, AS 64500)

Neither ROA uses the maxLength attribute. But, the second ROA is not "minimal" because it contains a /22 prefix that is not originated by anyone in BGP during normal operations. The /22 prefix is only originated by AS 64500 as part of its DDoS mitigation service during a DDoS attack.

Notice, however, that this scheme does not come without risks. Namely, all IP addresses in 168.122.0.0/22 are vulnerable to a forged-origin subprefix hijack during normal operations, when the /22 prefix is not originated. (The hijacker AS 64511 would send the BGP announcement "168.122.0.0/22: AS 64511, AS 64500", falsely claiming that AS 64511 is a neighbor of AS 64500 and falsely claiming that AS 64500 originates 168.122.0.0/22.)

In some situations, the DDoS mitigation service at AS 64500 might want to limit the amount of DDoS traffic that it attracts and scrubs. Suppose that a DDoS attack only targets IP addresses in 168.122.0.0/24. Then, the DDoS mitigation service at AS 64500 only wants to attract the traffic designated for the /24 prefix that is under attack, but not the entire /22 prefix. To allow for this, the RPKI should have two ROAs: one for AS 64496 and one for AS 64500.

ROA:(168.122.0.0/16,168.122.225.0/24, AS 64496)

ROA:(168.122.0.0/22-24, AS 64500)

The second ROA uses the maxLength attribute because it is designed to explicitly enable AS 64500 to originate *any* /24 subprefix of 168.122.0.0/22.

As before, the second ROA is also not "minimal" because it contains prefixes that are not originated by anyone in BGP during normal operations. As before, all IP addresses in 168.122.0.0/22 are vulnerable to a forged-origin subprefix hijack during normal operations, when the /22 prefix is not originated.

The use of maxLength in this second ROA also comes with an additional risk. While it permits the DDoS mitigation service at AS 64500 to originate prefix 168.122.0.0/24 during a DDoS attack in that space, it also makes the *other* /24 prefixes covered by the /22 prefix (i.e., 168.122.1.0/24, 168.122.2.0/24, 168.122.3.0/24) vulnerable to a forged-origin subprefix attacks.

6. Acknowledgments

The authors would like to thank the following people for their review and contributions to this document: Omar Sagga (Boston University) and Aris Lambrianidis (AMS-IX).

7. References

7.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <https://www.rfc-editor.org/info/rfc4271>.

[RFC6480] Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", RFC 6480, DOI 10.17487/RFC6480, February 2012, <https://www.rfc-editor.org/info/rfc6480>.

[RFC6482] Lepinski, M., Kent, S., and D. Kong, "A Profile for Route Origin Authorizations (ROAs)", RFC 6482, DOI 10.17487/RFC6482, February 2012, <https://www.rfc-editor.org/info/rfc6482>.

[RFC6811] Mohapatra, P., Scudder, J., Ward, D., Bush, R., and R. Austein, "BGP Prefix Origin Validation", RFC 6811, DOI 10.17487/RFC6811, January 2013, <https://www.rfc-editor.org/info/rfc6811>.

7.2. Informative References

[GCHSS] Gilad, Y., Cohen, A., Herzberg, A., Schapira, M., and H. Shulman, "Are We There Yet? On RPKI's Deployment and Security", in NDSS 2017, February 2017, <https://eprint.iacr.org/2016/1010.pdf>.

[GSG17] Gilad, Y., Sagga, O., and S. Goldberg, "Maxlength Considered Harmful to the RPKI", in ACM CoNEXT 2017, December 2017, <https://eprint.iacr.org/2016/1015.pdf>.

[HARMFUL] Gilad, Y., Sagga, O., and S. Goldberg, "MaxLength Considered Harmful to the RPKI", 2017, <https://eprint.iacr.org/2016/1015.pdf>.

[LSG16] Lychev, R., Shapira, M., and S. Goldberg, "Rethinking Security for Internet Routing", in Communications of the ACM, October 2016, <http://cacm.acm.org/magazines/2016/10/207763-rethinking-security-for-internet-routing/>.

[RFC6907] Manderson, T., Sriram, K., and R. White, "Use Cases and Interpretations of Resource Public Key Infrastructure (RPKI) Objects for Issuers and Relying Parties", RFC 6907, DOI 10.17487/RFC6907, March 2013, <https://www.rfc-editor.org/info/rfc6907>.

[RFC7115] Bush, R., "Origin Validation Operation Based on the Resource Public Key Infrastructure (RPKI)", BCP 185, RFC 7115, DOI 10.17487/RFC7115, January 2014, <https://www.rfc-editor.org/info/rfc7115>.

[RFC8205] Lepinski, M., Ed. and K. Sriram, Ed., "BGPsec Protocol Specification", RFC 8205, DOI 10.17487/RFC8205, September 2017, <https://www.rfc-editor.org/info/rfc8205>.

Authors' Addresses

Yossi Gilad
Boston University
111 Cummington St, MCS135
Boston, MA 02215
USA

EMail: yossigi@bu.edu

Sharon Goldberg
Boston University
111 Cummington St, MCS135
Boston, MA 02215
USA

EMail: goldbe@cs.bu.edu

Kotikalapudi Sriram
USA National Institute of Standards and Technology
100 Bureau Drive
Gaithersburg, MD 20899
USA

EMail: kotikalapudi.sriram@nist.gov

Job Snijders
NTT Communications
Theodorus Majofskistraat 100
Amsterdam 1065 SZ
The Netherlands

EMail: job@ntt.net

Ben Maddison
Workonline Communications
30 Waterkant St
Cape Town 8001
South Africa

EMail: benm@workonline.co.za