

Network Working Group
Internet-Draft
Intended status: Best Current Practice
Expires: August 26, 2021

Y. Gilad
Hebrew University of Jerusalem
S. Goldberg
Boston University
K. Sriram
USA NIST
J. Snijders
Fastly
B. Maddison
Workonline Communications
February 22, 2021

**The Use of Maxlength in the RPKI
draft-ietf-sidrops-rpkimaxlen-06**

Abstract

This document recommends ways to reduce forged-origin hijack attack surface by prudently limiting the set of IP prefixes that are included in a Route Origin Authorization (ROA). One recommendation is to avoid using the maxLength attribute in ROAs except in some specific cases. The recommendations complement and extend those in [RFC 7115](#). The document also discusses creation of ROAs for facilitating the use of Distributed Denial of Service (DDoS) mitigation services. Considerations related to ROAs and origin validation in the context of destination-based Remote Triggered Black Hole (RTBH) filtering are also highlighted.

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 August 26, 2021.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
1.1.	Requirements	3
1.2.	Documentation Prefixes	4
2.	Suggested Reading	4
3.	Forged-Origin Subprefix Hijack	4
4.	Measurements of Today's RPKI	6
5.	Recommendations about Minimal ROAs and maxLength	7
5.1.	Facilitating Ad-hoc Routing Changes and DDoS Mitigation	7
6.	ROAs and Origin Validation for RTBH Filtering Scenario	9
7.	Acknowledgments	10
8.	References	10
8.1.	Normative References	10
8.2.	Informative References	11
	Authors' Addresses	12

[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 that prefix. Each ROA contains a set of IP prefixes, and an AS number of an AS authorized to 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, measurements have shown that 84% of the prefixes specified in ROAs that use the maxLength attribute, are vulnerable to a forged-origin subprefix hijack [[HARMFUL](#)]. The forged-origin prefix or subprefix hijack involves inserting the legitimate AS as specified in the ROA as the origin AS in the AS_PATH, and can be launched against any IP prefix/subprefix that has a ROA. Consider a prefix/subprefix that has a ROA but is unused, i.e., not announced in BGP by a legitimate AS. A forged origin hijack involving such a prefix/subprefix can propagate widely throughout the Internet. On the other hand, if the prefix/subprefix were announced by the legitimate AS, then the propagation of the forged-origin hijack is somewhat limited because of its increased AS_PATH length relative to the legitimate announcement. Of course, forged-origin hijacks are harmful in both cases but the extent of harm is greater for unannounced prefixes.

For this reason, this document recommends that, whenever possible, operators SHOULD use "minimal ROAs" that authorize only those IP prefixes that are actually originated in BGP, and no other prefixes. Further, it recommends ways to reduce forged-origin attack surface by prudently limiting the address space that is included in Route Origin Authorizations (ROAs). One recommendation is to avoid using the maxLength attribute in ROAs except in some specific cases. The recommendations complement and extend those in [[RFC7115](#)]. The document also discusses creation of ROAs for facilitating the use of Distributed Denial of Service (DDoS) mitigation services. Considerations related to ROAs and origin validation in the context of destination-based Remote Triggered Black Hole (RTBH) filtering are also highlighted.

One ideal place to implement the ROA related 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.

Best current practices described in this document require 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](#)].

1.2. Documentation Prefixes

The documentation prefixes recommended in [RFC5737] are insufficient for use as example prefixes in this document. Therefore, this document uses [RFC1918] address space for constructing example prefixes.

2. Suggested Reading

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

3. Forged-Origin Subprefix Hijack

A detailed description and discussion of forged-origin subprefix hijacks are presented here, especially considering the case when the subprefix is not announced in BGP. 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.

Note that this set of assumptions accurately describes a substantial, and growing, number of large Internet networks at the time writing.

The forged-origin subprefix hijack [RFC7115] [GCHSS] is described here using a running example.

Consider the IP prefix 192.168.0.0/16 which is allocated to an organization that also operates AS 64496. In BGP, AS 64496 originates the IP prefix 192.168.0.0/16 as well as its subprefix 192.168.225.0/24. Therefore, the RPKI should contain a ROA authorizing AS 64496 to originate these two IP prefixes.

Suppose, however, the organization issues and publishes a ROA including a maxLength value of 24:

```
ROA:(192.168.0.0/16-24, AS 64496)
```

We refer to the above as a "loose ROA" since it authorizes AS 64496 to originate any subprefix of 192.168.0.0/16 up to and including

length /24, rather than only those prefixes that are intended to be announced in BGP.

Because AS 64496 only originates two prefixes in BGP: 192.168.0.0/16 and 192.168.225.0/24, all other prefixes authorized by the "loose ROA" (for instance, 192.168.0.0/24), are vulnerable to the following forged-origin subprefix hijack [[RFC7115](#)] [[GCHSS](#)]:

The hijacker AS 64511 sends a BGP announcement "192.168.0.0/24: 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 192.168.0.0/24. In fact, the IP prefix 192.168.0.0/24 is not originated by AS 64496.

The hijacker's BGP announcement is valid according to the RPKI, since the ROA (192.168.0.0/16-24, AS 64496) authorizes AS 64496 to originate BGP routes for 192.168.0.0/24.

Because AS 64496 does not actually originate a route for 192.168.0.0/24, the hijacker's route is the **only** route to the 192.168.0.0/24. Longest-prefix-match routing ensures that the hijacker's route to the subprefix 192.168.0.0/24 is always preferred over the legitimate route to 192.168.0.0/16 originated by AS 64496.

Thus, the hijacker's route propagates through the Internet, the traffic destined for IP addresses in 192.168.0.0/24 will be delivered to the hijacker.

The forged-origin **subprefix** hijack would have failed if a "minimal ROA" described below was used instead of the "loose ROA". In this example, a "minimal ROA" would be:

ROA:(192.168.0.0/16, 192.168.225.0/24, 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](#)].

The "minimal ROA" renders AS 64511's BGP announcement invalid, because:

(1) this ROA "covers" the attacker's announcement (since 192.168.0.0/24 is a subprefix of 192.168.0.0/16), and

(2) there is no ROA "matching" the attacker's announcement (there is no ROA for AS 64511 and IP prefix 192.168.0.0/24) [[RFC6811](#)].

If routers ignore invalid BGP announcements, the minimal ROA above ensures that the subprefix hijack will fail.

Thus, if a "minimal ROA" had been used, 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 "192.168.0.0/16: 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:

AS 64496 legitimately originates 192.168.0.0/16 in BGP, so the hijacker AS 64511 is not presenting the **only** route to 192.168.0.0/16.

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 summary, a forged-origin subprefix hijack has the same impact as a regular subprefix hijack, despite the increased AS_PATH length of the illegitimate route. A forged-origin **subprefix** hijack is also more damaging than forged-origin **prefix** hijack.

4. Measurements of Today's RPKI

Network measurements have shown that 12% of the IP prefixes authorized in ROAs have a maxLength longer than their prefix length. Of these, the vast majority (84%) are non-minimal, as they include subprefixes that are not announced in BGP by the legitimate AS, and are thus vulnerable to forged origin subprefix hijacks. See [[GSG17](#)] for details.

These measurements suggest that operators commonly misconfigure the maxLength attribute, and unwittingly open themselves up to forged-origin subprefix hijacks. That is, they are exposing a much larger attack surface for forged-origin hijacks than necessary.

5. Recommendations about Minimal ROAs and maxLength

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.)

In general, except in some special cases, operators SHOULD avoid using the maxLength attribute in their ROAs, since its inclusion will usually make the ROA non-minimal.

One such exception may be when all more specific prefixes permitted by the maxLength are actually announced by the AS in the ROA. Another exception is where: (a) the maxLength is substantially larger compared to the specified prefix length in the ROA, and (b) a large number of more specific prefixes in that range are announced by the AS in the ROA. This case should occur rarely in practice (if at all). Operator discretion is necessary in this case.

This practice requires no changes to the RPKI specification and need 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. Facilitating Ad-hoc Routing Changes and DDoS Mitigation

Operational requirements may require that a route for an IP prefix be originated on an ad-hoc basis, with little or no prior warning. An example of such a situation arises where an operator wishes to make use of DDoS mitigation services that use BGP to redirect traffic via a "scrubbing center".

In order to ensure that such ad-hoc routing changes are effective, there should exist a ROA validating the new route. However a difficulty arises due to the fact that newly created objects in the RPKI are made visible to relying parties considerably more slowly than routing updates in BGP.

Ideally, it would not be necessary to pre-create the ROA which validates the ad-hoc route, and instead create it "on-the-fly" as required. However, this is practical only if the latency imposed by the propagation of RPKI data is guaranteed to be within acceptable limits in the circumstances. For time-critical interventions such as responding to a DDoS attack, this is unlikely to be the case.

Thus, the ROA in question will usually need to be created well in advance of the routing intervention, but such a ROA will be non-

minimal, since it 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 unless doing so has no impact on the set of included prefixes.

The running example is now extended to illustrate one situation where it is not possible to issue a minimal ROA.

Consider the following scenario prior to deployment of RPKI. Suppose AS 64496 announced 192.168.0.0/16 and has a contract with a Distributed Denial of Service (DDoS) mitigation service provider that holds AS 64500. Further, assume that the DDoS mitigation service contract applies to all IP addresses covered by 192.168.0.0/22. When a DDoS attack is detected and reported by AS 64496, AS 64500 immediately originates 192.168.0.0/22, 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 AS 64496's /16 prefix, and so all the traffic (destined to addresses in 192.168.0.0/22) that normally goes to AS 64496 goes to AS 64500 instead. In some deployments, the origination of the /22 route is performed by AS 64496 and announced only to AS 64500, which then announces transit for that prefix. This variation does not change the properties considered here.

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 DDoS mitigation (and traffic scrubbing) scheme would not work. That is, if AS 64500 originates the /22 prefix in BGP during DDoS attacks, the announcement would be invalid [[RFC6811](#)].

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

ROA:(192.168.0.0/16, 192.168.225.0/24, AS 64496)

ROA:(192.168.0.0/22, 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 192.168.0.0/22 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 "192.168.0.0/22: AS 64511, AS 64500", falsely claiming that AS 64511 is a neighbor of AS 64500 and falsely claiming that AS 64500 originates 192.168.0.0/22.)

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 192.168.0.0/24. 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:(192.168.0.0/16, 192.168.225.0/24, AS 64496)

ROA:(192.168.0.0/22-24, AS 64500)

The second ROA uses the maxLength attribute because it is designed to explicitly enable AS 64500 to originate *any* /24 subprefix of 192.168.0.0/22.

As before, the second ROA is not "minimal" because it contains prefixes that are not originated by anyone in BGP during normal operations. As before, all IP addresses in 192.168.0.0/22 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 192.168.0.0/24 during a DDoS attack in that space, it also makes the *other* /24 prefixes covered by the /22 prefix (i.e., 192.168.1.0/24, 192.168.2.0/24, 192.168.3.0/24) vulnerable to a forged-origin subprefix attacks.

6. ROAs and Origin Validation for RTBH Filtering Scenario

Considerations related to ROAs and origin validation [[RFC6811](#)] for the case of destination-based Remote Triggered Black Hole (RTBH) filtering are addressed here. In RTBH filtering, highly specific

prefixes (greater than /24 in IPv4 and greater than /48 in IPv6; possibly even /32 (IPv4) and /128 (IPv6)) are announced in BGP. These announcements are tagged with a BLACKHOLE Community [RFC7999]. It is obviously not desirable to use large maxLength or include any such highly specific prefixes in the ROAs to accommodate destination-based RTBH filtering, for the reasons set out above.

As a result, RPKI based route origin validation [RFC6811] is a poor fit for the validation of RTBH routes. Specification of new procedures to address this use case through the use of the RPKI is outside the scope of this document.

Therefore:

- o Operators SHOULD NOT create non-minimal ROAs (either by creating additional ROAs, or through the use of maxLength) for the purpose of advertising RTBH routes; and
- o Operators providing a means for operators of neighboring autonomous systems to advertise RTBH routes via BGP MUST NOT make the creation of non-minimal ROAs a pre-requisite for its use.

7. 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). Thanks are also due to Matthias Waehlich (Free University of Berlin) for comments and suggestions.

8. References

8.1. Normative References

- [RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <<https://www.rfc-editor.org/info/rfc1918>>.
- [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>>.

8.2. 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/>>.
- [NIST-800-189] Sriram, K. and D. Montgomery, "Resilient Interdomain Traffic Exchange: BGP Security and DDoS Mitigation", NIST Special Publication, NIST SP 800-189, December 2019, <<https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-189.pdf>>.
- [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>>.
- [RFC7999] King, T., Dietzel, C., Snijders, J., Doering, G., and G. Hankins, "BLACKHOLE Community", [RFC 7999](#), DOI 10.17487/RFC7999, October 2016, <<https://www.rfc-editor.org/info/rfc7999>>.
- [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
Hebrew University of Jerusalem
Rothburg Family Buildings, Edmond J. Safra Campus
Jerusalem 9190416
Israel

EMail: yossigi@cs.huji.ac.il

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
Fastly
Amsterdam
Netherlands

EMail: job@fastly.com

Ben Maddison
Workonline Communications
114 West St
Johannesburg 2196
South Africa

EMail: benm@workonline.africa

