Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: July 10, 2019

A. Azimov
E. Uskov
Qrator Labs
R. Bush
Internet Initiative Japan
K. Patel
Arrcus
J. Snijders
NTT
R. Housley
Vigil Security

January 6, 2019

A Profile for Autonomous System Provider Authorization
draft-azimov-sidrops-aspa-profile-01

Abstract

This document defines a standard profile for Autonomous System Provider Authorization in the Resource Public Key Infrastructure. An Autonomous System Provider Authorization is a digitally signed object that provides a means of verifying that a Customer Autonomous System holder has authorized a Provider Autonomous System to be its upstream provider and for the Provider to send prefixes received from the Customer Autonomous System in all directions including providers and peers.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

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 July 10, 2019.

Table of Contents

1. Introduction

The primary purpose of the Resource Public Key Infrastructure (RPKI) is to improve routing security. (See [RFC6480] for more information.) As part of this infrastructure, a mechanism is needed to verify that a Provider AS (PAS) has permission from a Customer AS (CAS) holder to send routes in all directions. The digitally signed Autonomous System Provider Authorization (ASPA) object provides this verification mechanism.

The ASPA uses the template for RPKI digitally signed objects [RFC6488], which defines a Cryptographic Message Syntax (CMS) [RFC5652] wrapper for the ASPA content as well as a generic validation procedure for RPKI signed objects. As ASPAs need to be validated with RPKI certificates issued by the current infrastructure, we assume the mandatory-to-implement algorithms in [RFC6485], or its successor.

To complete the specification of the ASPA (see Section 4 of [RFC6488]), this document defines:

1. The object identifier (OID) that identifies the ASPA signed object. This OID appears in the eContentType field of the encapContentInfo object as well as the content-type signed attribute within the signerInfo structure).

2. The ASN.1 syntax for the ASPA content, which is the payload signed by the CAS. The ASPA content is encoded using the ASN.1 [X680] Distinguished Encoding Rules (DER) [X690].

3. The steps required to validate an ASPA beyond the validation steps specified in [RFC6488]).

2. The ASPA Content Type

The content-type for an ASPA is defined as id-cct-ASPA, which has the numerical value of 1.2.840.113549.1.9.16.1.TBD. This OID MUST appear both within the eContentType in the encapContentInfo structure as well as the content-type signed attribute within the signerInfo structure (see [RFC6488]).

3. The ASPA eContent

The content of an ASPA identifies the Customer AS (CAS) as well as the Provider AS (PAS) that is authorized to further propagate announcements received from the customer. If customer has multiple providers, it issues multiple ASPAs, one for each provider AS. An ASPA is formally defined as:

```
ct-ASPA CONTENT-TYPE ::=
    { ASProviderAttestation IDENTIFIED BY id-ct-ASPA }

id-ct-ASPA OBJECT IDENTIFIER ::= { id-ct TBD }

ASProviderAttestation ::= SEQUENCE {
    version [0] ASPAVersion DEFAULT v0,
    AFI  AddressFamilyIdentifier,
    customerASID  ASID,
    providerASID  ASID }

ASPAVersion ::= INTEGER  { v0(0) }

AddressFamilyIdentifier ::= INTEGER

ASID ::= INTEGER
```

Note that this content appears as the eContent within the encapContentInfo as specified in [RFC6488].

3.1. version

The version number of the ASProviderAttestation MUST be v0.

3.2. AFI

The AFI field contains Address Family Identifier for which the relation between customer and provider ASes is authorized. Presently defined values for the Address Family Identifier field are specified in the IANA's Address Family Numbers registry [IANA-AF].

3.3. customerASID

The customerASID field contains the AS number of the Autonomous System that authorizes an upstream provider (listed in the providerASId) to propagate prefixes in the specified address family other ASes.

3.4. providerASID

The providerASID contains the AS number that is authorized to further propagate announcements in the specified address family received from the customer.

4. ASPA Validation

Before a relying party can use an ASPA to validate a routing announcement, the relying party MUST first validate the ASPA object itself. To validate an ASPA, the relying party MUST perform all the validation checks specified in [RFC6488] as well as the following additional ASPA-specific validation step.

- The autonomous system identifier delegation extension [RFC3779] is present in the end-entity (EE) certificate (contained within the ASPA), and the customer AS number in the ASPA is contained within the set of AS numbers specified by the EE certificate's autonomous system identifier delegation extension.

5. ASN.1 Module for the ASPA Content Type

```
RPKI-ASPA-2018
    { iso(1) member-body(2) us(840) rsadsi(113549) pkcs(1)
       pkcs-9(9) smime(16) modules(0) id-mod-rpki-aspa-2018(TBD2) }
DEFINITIONS IMPLICIT TAGS ::=
BEGIN
IMPORTS

CONTENT-TYPE
FROM CryptographicMessageSyntax-2010  -- RFC 6268
    { iso(1) member-body(2) us(840) rsadsi(113549) pkcs(1)
       pkcs-9(9) smime(16) modules(0) id-mod-cms-2009(58) } ;

ContentSet CONTENT-TYPE ::= { ct-ASPA, ... }

--
-- ASPA Content Type
--

id-smime OBJECT IDENTIFIER ::= { iso(1) member-body(2)
     us(840) rsadsi(113549) pkcs(1) pkcs-9(9) 16 }

id-ct OBJECT IDENTIFIER ::= { id-smime 1 }

id-ct-ASPA OBJECT IDENTIFIER ::= { id-ct TBD }

ct-ASPA CONTENT-TYPE ::=
    { TYPE ASProviderAttestation IDENTIFIED BY id-ct-ASPA }

ASProviderAttestation ::= SEQUENCE {
    version [0] ASPAVersion DEFAULT v0,
    AFI  AddressFamilyIdentifier,
    customerASID  ASID,
    providerASID  ASID }

ASPAVersion ::= INTEGER  { v0(0) }

AddressFamilyIdentifier ::= INTEGER

ASID ::= INTEGER

END
```

6. IANA Considerations

Please add the id-mod-rpki-aspa-2018 to the SMI Security for S/MIME Module Identifier (1.2.840.113549.1.9.16.0) registry (https://www.iana.org/assignments/smi-numbers/smi-numbers.xml#security-smime-0) as follows:

Decimal	Description	Specification
TBD2	id-mod-rpki-aspa-2018	[ThisRFC]

Please add the ASPA to the SMI Security for S/MIME CMS Content Type (1.2.840.113549.1.9.16.1) registry (https://www.iana.org/assignments/smi-numbers/smi-numbers.xml#security-smime-1) as follows:

Decimal	Description	Specification
TBD	id-ct-ASPA	[ThisRFC]

Please add the ASPA to the RPKI Signed Object registry (https://www.iana.org/assignments/rpki/rpki.xhtml#signed-objects) as follows:

Name	OID	Specification
ASPA	1.2.840.113549.1.9.16.1.TBD	[ThisRFC]

7. Security Considerations

8. Acknowledgments

9. References

9.1. Normative References

[IANA-AF] IANA, "Address Family Numbers", <http://www.iana.org/numbers.html>.

[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>.

[RFC3779] Lynn, C., Kent, S., and K. Seo, "X.509 Extensions for IP Addresses and AS Identifiers", RFC 3779, DOI 10.17487/RFC3779, June 2004, <https://www.rfc-editor.org/info/rfc3779>.

[RFC5652] Housley, R., "Cryptographic Message Syntax (CMS)", STD 70, RFC 5652, DOI 10.17487/RFC5652, September 2009, <https://www.rfc-editor.org/info/rfc5652>.

[RFC6485] Huston, G., "The Profile for Algorithms and Key Sizes for Use in the Resource Public Key Infrastructure (RPKI)", RFC 6485, DOI 10.17487/RFC6485, February 2012, <https://www.rfc-editor.org/info/rfc6485>.

[RFC6488] Lepinski, M., Chi, A., and S. Kent, "Signed Object Template for the Resource Public Key Infrastructure (RPKI)", RFC 6488, DOI 10.17487/RFC6488, February 2012, <https://www.rfc-editor.org/info/rfc6488>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[X680] ITU-T, "Information technology -- Abstract Syntax Notation One (ASN.1): Specification of basic notation", ITU-T Recommendation X.680, 2015.

[X690] ITU-T, "Information Technology -- ASN.1 encoding rules: Specification of Basic Encoding Rules (BER), Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER)", ITU-T Recommendation X.690, 2015.

9.2. Informative References

[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>.

Authors' Addresses

Alexander Azimov
Qrator Labs

Email: a.e.azimov@gmail.com

Eugene Uskov
Qrator Labs

Email: eu@qrator.net

Randy Bush
Internet Initiative Japan

Email: randy@psg.com

Keyur Patel
Arrcus, Inc.

Email: keyur@arrcus.com

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

Email: job@ntt.net

Russ Housley
Vigil Security, LLC
918 Spring Knoll Drive
Herndon, VA 20170
USA

Email: housley@vigilsec.com

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 25, 2019

A. Azimov
E. Bogomazov
Qrator Labs
R. Bush
Internet Initiative Japan
K. Patel
Arrcus, Inc.
J. Snijders
NTT

October 22, 2018

Verification of AS_PATH Using the Resource Certificate Public Key Infrastructure and Autonomous System Provider Authorization
draft-azimov-sidrops-aspa-verification-01

Abstract

This document defines the semantics of an Autonomous System Provider Authorization object in the Resource Public Key Infrastructure to verify the AS_PATH attribute of routes advertised in the Border Gateway Protocol.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

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 April 25, 2019.

Table of Contents

1. Introduction

The Border Gateway Protocol (BGP) was designed with no mechanisms to validate BGP attributes. Two consequences are BGP Hijacks and BGP Route Leaks [RFC7908]. BGP extensions are able to partially solve these problems. For example, ROA-based Origin Validation [RFC6483] can be used to detect and filter accidental mis-originations, and [I-D.ymbk-idr-bgp-eotr-policy] can be used to detect accidental route leaks. While these upgrades to BGP are quite useful, they still rely on transitive BGP attributes, i.e. AS_PATH, that can be manipulated by attackers.

BGPSec [RFC8205] was designed to solve the problem of AS_PATH validation. Unfortunately, strict cryptographic validation brought unaffordable computational overhead for BGP routers. BGPSec also proved to be vulnerable to downgrade attacks that can nullify all the work of AS_PATH signing. As a result, to abuse the AS_PATH or any

other signed transit attribute, an attacker merely needs to downgrade to 'old' BGP-4.

An alternative approach was introduced with soBGP [I-D.white-sobgp-architecture]. Instead of strong cryptographic AS_PATH validation, it was suggested to create an AS_PATH security function based on a shared database of ASN adjacencies. While such an approach has reasonable computational cost, the two side adjacencies don't provide a way to automate anomaly detection without high adoption rate - an attacker can easily up a one-way adjacency. SO-BGP suggested sharing data about adjacencies using additional BGP messages, which is recursively complex thus significantly increasing adoption complexity. In addition, the general goal to verify all AS_PATHs was not achievable given the indirect adjacencies at internet exchange points.

Instead of the general goal of checking AS_PATH correctness, this document focuses on solving real-world operational problems - automatic detection of malicious hijacks and route leaks. To achieve this goal a new AS_PATH verification procedure is defined which is able to automatically detect invalid (malformed) AS_PATHs in announcements that are received from customers and peers. This procedure uses a shared signed database of customer-to-provider relationships that is built using a new RPKI object - Autonomous System Provider Authorization (ASPA). This technique provides benefits for the participants even in a state of early adoption.

2. Anomaly Propagation

Both route leaks and hijacks have similar effects on ISP operations - they redirect traffic, resulting in increased latency, packet loss, or possible MiTM attacks. But the level of risk depends significantly on the propagation of these BGP anomalies. For example, a hijack that is propagated only to customers may concentrate traffic in a particular ISP's customer cone; while if the anomaly is propagated through peers, upstreams, or reaches Tier-1 networks, thus distributing globally, traffic may be redirected at the level of entire countries and/or global providers.

The ability to constrain propagation of BGP anomalies to upstreams and peers, without requiring support from the source of the anomaly (which is critical if source has malicious intent), should significantly improve the security of inter-domain routing and solve the majority of problems.

3. Autonomous System Provider Authorization

As described in [RFC6480], the RPKI is based on a hierarchy of resource certificates that are aligned to the Internet Number Resource allocation structure. Resource certificates are X.509 certificates that conform to the PKIX profile [RFC5280], and to the extensions for IP addresses and AS identifiers [RFC3779]. A resource certificate is a binding by an issuer of IP address blocks and Autonomous System (AS) numbers to the subject of a certificate, identified by the unique association of the subject's private key with the public key contained in the resource certificate. The RPKI is structured so that each current resource certificate matches a current resource allocation or assignment.

ASPAs are digitally signed objects that bind a selected AFI Provider AS number to a Customer AS number (in terms of BGP announcements not business), and are signed by the holder of the Customer AS. An ASPA attests that a Customer AS holder (CAS) has authorized a particular Provider AS (PAS) to propagate the Customer's IPv4/IPv6 announcements onward, e.g. to the Provider's upstream providers or peers. The ASPA record profile is described in [I-D.azimov-sidrops-aspa-profile].

4. Customer-Provider Verification Procedure

This section describes an abstract procedure that checks that pair of ASNs (AS1, AS2) is included in the set of signed ASPAs. The semantics of its usa are defined in next section. The procedure takes (AS1, AS2, ROUTE_AFI) as input parameters and returns three types of results: "valid", "invalid" and "unknown".

A relying party (RP) must have access to a local cache of the complete set of cryptographically valid ASPAs when performing customer-provider verification procedure.

1. Retrieve all cryptographically valid ASPAs in a selected AFI with a customer value of AS1. This selection forms the set of "candidate ASPAs."

2. If the set of candidate ASPAs is empty, then the procedure exits with an outcome of "unknown."

3. If there is at least one candidate ASPA where the provider field is AS2, then the procedure exits with an outcome of "valid."

4. Otherwise, the procedure exits with an outcome of "invalid."

Since an AS1 may have different set providers in different AFI, it should also have different set of corresponding ASPAs. In this case,

the output of this procedure with input (AS1, AS2, ROUTE_AFI) may have different output for different ROUTE_AFI values.

5. AS_PATH Verification

The AS_PATH attribute identifies the autonomous systems through which an UPDATE message has passed. AS_PATH may contain two types of components: ordered AS_SEQes and unordered AS_SETs, as defined in [RFC4271].

The value of each AS_SEQ component can be described as set of pairs {(AS(I), prepend(I)), (AS(I-1), prepend(I-1))...}. In this case, the sequence {AS(I), AS(I-1),...} represents different ASNs, that packet should pass towards the destination. When a route is received from a customer or a literal peer, each pair (AS(I-1), AS(I)) MUST belong to customer-provider or sibling relationship. If there are other types of relationships, it means that the route was leaked or the AS_PATH attribute was malformed. The goal of the above procedure is to check the correctness of this statement.

For 32-bit AS number compatible BGP speakers, if a route from ROUTE_AFI address family is received from a customer or peer, its AS_PATH MUST be verified as follows:

1. If the closest AS in the AS_PATH is not the receiver's neighbor ASN then procedure halts with the outcome "invalid";

2. If in one of AS_SEQ segments there is a pair (AS(I-1), AS(I)), and customer-provider verification procedure (Section 4) with parameters (AS(I-1), AS(I), ROUTE_AFI) returns "invalid" then the procedure also halts with the outcome "invalid";

3. If the AS_PATH has at least one AS_SET segment then procedure halts with the outcome "unverifiable";

4. Otherwise, the procedure halts with an outcome of "valid".

For BGP speakers that are not 32-bit AS compatible, the above procedure is slightly different. In point 2 if at least one AS(I-1), AS(I) is equal to AS_TRANS(23456), the corresponding pair must be passed without check using the customer-provider verification procedure.

If the output of the AS_PATH verification procedure is "invalid" the LOCAL_PREF SHOULD be set to 0 or the route MAY be dropped. If an "invalid" route has no alternative route(s) and it is propagated to other ASes despite the above, it MUST be marked with the GRACEFUL_SHUTDOWN community to avoid possible stable oscillations,

when an unchecked route received from a provider becomes preferred over an invalid route received from a customer. This also allows customers to detect malformed routes received from upstream providers.

If the output of the AS_PATH verification procedure is 'unverifiable' it means that AS_PATH can't be fully verified. Such routes should be treated with caution and SHOULD be processed the same way as "invalid" routes. This policy goes with full correspondence to [I-D.kumari-deprecate-as-set-confed-set].

The above AS_PATH verification procedure is able to check routes received from customers and peers. The ASPA mechanism combined with BGP Roles [I-D.ietf-idr-bgp-open-policy] and ROA-based Origin Validation [RFC6483] provide a fully automated solution to detect and filter hijacks and route leaks, including malicious ones.

6. Disavowal of Provider Authorizaion

An ASPA is a positive attestation that an AS holder has authorized its provider to redistribute received routes to the provider's providers and peers. This does not preclude the provider AS from redistribution to its other customers. By creating an ASPA where the provider AS is 0, the customer indicates that no provider should further announce its routes. Specifically, AS 0 is reserved to identify provider-free networks, Internet exchange meshes, etc.

An ASPA with a provider AS of 0 is a statement by the customer AS that the its routes should not be received by any relying party AS from any of its customers or peers.

By convention, an ASPA with a provider AS of 0 should be the only ASPA issued by a given AS holder; although this is not a strict requirement. A provider 0 ASPA may coexist with ASPAs that have different provider AS values; though in such cases, the presence or absence of the provider AS 0 ASPA does not alter the AS_PATH verification procedure.

7. Siblings (Complex Relations)

There are peering relationships which can not be described as strictly simple peer-peer or customer-provider; e.g. when both parties are intentionally sending prefixes received from each other to their peers and/or upstreams.

In this case, two symmetric ASPAs records {(AS1, AS2), (AS2, AS1)} must be created by AS1 and AS2 respectively.

8. Security Considerations

ASPA issuers should be aware of the verification implication in issuing an ASPA - an ASPA implicitly invalidates all routes passed to upstream providers other than the provider ASs listed in the collection of ASPAs. It is the Customer AS's duty to maintain a correct set of ASPAs.

While the ASPA provides a check of an AS_PATH for routes received from customers and peers, it doesn't provide full support for routes that are received from upstream providers. So, this mechanism guarantees detection of both malicious and accidental route leaks and provides partial support for detection of malicious hijacks: upstream transit ISPs will still be able to send hijacked prefixes with malformed AS_PATHs to their customers.

9. Acknowledgments

The authors wish to thank authors of [RFC6483] since its text was used as an example while writing this document.

10. References

10.1. 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>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

10.2. Informative References

[I-D.azimov-sidrops-aspa-profile] Azimov, A., Uskov, E., Bush, R., Patel, K., Snijders, J., and R. Housley, "A Profile for Autonomous System Provider Authorization", draft-azimov-sidrops-aspa-profile-00 (work in progress), June 2018.

[I-D.ietf-idr-bgp-open-policy] Azimov, A., Bogomazov, E., Bush, R., Patel, K., and K. Sriram, "Route Leak Prevention using Roles in Update and Open messages", draft-ietf-idr-bgp-open-policy-02 (work in progress), January 2018.

[I-D.kumari-deprecate-as-set-confed-set]
Kumari, W. and K. Sriram, "Deprecation of AS_SET and AS_CONFED_SET in BGP", draft-kumari-deprecate-as-set-confed-set-12 (work in progress), July 2018.

[I-D.white-sobgp-architecture]
White, R., "Architecture and Deployment Considerations for Secure Origin BGP (soBGP)", draft-white-sobgp-architecture-02 (work in progress), June 2006.

[I-D.ymbk-idr-bgp-eotr-policy]
Azimov, A., Bogomazov, E., Bush, R., and K. Patel, "Route Leak Detection and Filtering using Roles in Update and Open messages", draft-ymbk-idr-bgp-eotr-policy-02 (work in progress), March 2018.

[RFC3779] Lynn, C., Kent, S., and K. Seo, "X.509 Extensions for IP Addresses and AS Identifiers", RFC 3779, DOI 10.17487/RFC3779, June 2004, <https://www.rfc-editor.org/info/rfc3779>.

[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>.

[RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <https://www.rfc-editor.org/info/rfc5280>.

[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>.

[RFC6483] Huston, G. and G. Michaelson, "Validation of Route Origination Using the Resource Certificate Public Key Infrastructure (PKI) and Route Origin Authorizations (ROAs)", RFC 6483, DOI 10.17487/RFC6483, February 2012, <https://www.rfc-editor.org/info/rfc6483>.

[RFC7908] Sriram, K., Montgomery, D., McPherson, D., Osterweil, E., and B. Dickson, "Problem Definition and Classification of BGP Route Leaks", RFC 7908, DOI 10.17487/RFC7908, June 2016, <https://www.rfc-editor.org/info/rfc7908>.

[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

Alexander Azimov
Qrator Labs

Email: aa@qrator.net

Eugene Bogomazov
Qrator Labs

Email: eb@qrator.net

Randy Bush
Internet Initiative Japan

Email: randy@psg.com

Keyur Patel
Arrcus, Inc.

Email: keyur@arrcus.com

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

Email: job@ntt.net

Network Working Group
Internet-Draft
Intended status: Best Current Practice
Expires: April 25, 2019

Y. Gilad
S. Goldberg
Boston University
K. Sriram
USA NIST
J. Snijders
NTT
B. Maddison
Workonline Communications
October 22, 2018

The Use of Maxlength in the RPKI
draft-ietf-sidrops-rpkimaxlen-01

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].

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

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 19, 2019

T. Bruijnzeels
NLnet Labs
C. Martinez
LACNIC
R. Austein
Dragon Research Labs

October 16, 2018

RPKI Signed Object for Trust Anchor Keys
draft-ietf-sidrops-signed-tal-02

Abstract

Trust Anchor Locators (TALs) [I-D.ietf-sidrops-https-tal] are used by Relying Parties in the RPKI to locate and validate Trust Anchor certificates used in RPKI validation. This document defines an RPKI signed object for Trust Anchor Keys (TAK), that can be used by Trust Anchors to signal their set of current keys and the location(s) of the accompanying CA certiifcates to Relying Parties, as well as changes to this set in the form of revoked keys and new keys, in order to support both planned and unplanned key rolls without impacting RPKI validation.

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 April 19, 2019.

Table of Contents

1. Requirements notation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Overview

Trust Anchor Locators (TALs) [I-D.ietf-sidrops-https-tal] are used by Relying Parties in the RPKI to locate and validate Trust Anchor (TA) certificates used in RPKI validation. However, until now there has been no formal way of notifying Relying Parties (RP) of updates to a TAL. Such updates may be needed in particular in case a Trust Anchor needs to perform a planned, or unplanned, key roll.

This document defines a new RPKI signed object that can be used to document the current set of keys and the loctaion(s) of the accompanying CA certificates, as well as any changes to this set. This allows RPs to be notified automatically of such changes, and enables Trust Anchors to pre-stage a number of operational keys so that planned and unplanned key rolls can be performed without risking the invalidation of the RPKI tree under the TA. We call this object the Trust Anchor Keys (TAK) object.

When Relying Parties (RPs) are first bootstrapped, they use any current TAL to discover a key and location(s) of the TA certificate(s) for a TA. The RP can then retreive and validating the TA certificate, and subsequently validate the manifest [RFC6486] and CRL [section 5 of @!RFC6487]. However, before processing any other objects it will then first validate the TAK object, if present. All enumarated new keys (and locations) are then added to a new list of current TA keys for this TA. The RP will then recursively fetch and validate the TA certificates, manifest, CRL and TAK objects for each of these keys. As a part of this process the RP will also compile a list of revoked keys enumarated by any of the validly signed TAK objects. As the final step the RP will then filter out any revoked TA keys from its new set. This new set now replaces the previous set.

If the key used to start this process is still considered current, then validation continues. But if the key was revoked, then validation is restarted using one of the remaining keys in the set.

This process allows Trust Anchors to operate a set of N current keys, where any key can effectively revoke any or all of the other keys to perform either a planned, or an unplanned, key roll. This also

allows Trust Anchors to produce long lived TAK objects as forward pointers to RPs, and retire its old key when doing a key roll.

While the generic process is quite involved, the amount of work needed to support an envisioned normal key roll is fairly limited. Under normal circumstances a TA will typically have two current keys, so that is can perform an emergency roll over in case one of the keys is lost. This means that the RP will need to validate two TAK objects. However, typically these files will agree that both keys are current and validation continues.

When a key roll is executed a TA will remove one old key, and introduce one new (back-up) key. The RP will remove the old key from its set, and it will not be queried again, and it will add the new key and its TA certifcate location(s).

Only in a situation where an RP is very outdated can it be expected that the RP will have to discover several chained TAK object. But, since it will remove the outdated TALs in this process, this presents a one time cost only.

Note that in theory a TA can revoke all of its keys and make itself obsolete. In practice however, a well operated TA will have measures in place to prevent this. Furthermore they can protect themselves against key loss to adversaries through the use of such as the use of a Hardware Security Module (HSM) to protect keys. Protecting against this mis-operation would incur complexity and guesswork on the RPs. Therefore it is believed that it is best to keep the process straightforward, and offer a solution for the more likely issues of loss of a key, e.g. because an HSM or card set is broken, and planned key rolls.

3. TAK Object definition

The TAK object makes use of the template for RPKI digitally signed objects [RFC6488], which defines a Crytopgraphic Message Syntax (CMS) [RFC5652] wrapper for the Signed TALs content as well as a generic validation procedure for RPKI signed objects. Therefore, to complete the specification of the TAK object (see Section 4 of [RFC6488]), this document defines:

- The OID defined in Section 3.1 that identifies the signed object as being a TAK. (This OID appears within the eContentType in the encapContentInfo object as well as the content-type signed attribute in the signerInfo object).

- The ASN.1 syntax for the TAK eContent defined in Section 3.2.

- Additional steps to the validation steps specified in [RFC6488] required to validate the TAK, defined in Section 3.3.

3.1. The TAK Object Content Type

This document requests an OID for TAK objects as follows:

```
signed-Tal OBJECT IDENTIFIER ::= { iso(1) member-body(2) us(840)
           rsadsi(113549) pkcs(1) pkcs9(9) 16 id-smime (1) TBD }
```

This OID MUST appear both within the eContentType in the encapContentInfo object as well as the content-type signed attribute in the signerInfo object (see [RFC6488])

3.2. The TAK Object eContent

The content of a TAK object is ASN.1 encoded using the Distinguished Encoding Rules (DER) [X.690], and is defined as follows:

```
TAK ::= SEQUENCE {
   version   INTEGER DEFAULT 0,
   current   ::= SEQUENCE SIZE (1..MAX) OF CurrentKey,
   revoked   ::= SEQUENCE OF SubjectPublicKeyInfo
}

CurrentKey ::= SEQUENCE {
   certificateURIs       SEQUENCE SIZE (1..MAX) OF CertificateURI,
   subjectPublicKeyInfo  SubjectPublicKeyInfo
}

CertificateURI ::= IA5String

SubjectPublicKeyInfo ::= SEQUENCE {
     algorithm            AlgorithmIdentifier,
     subjectPublicKey     BIT STRING
}
```

3.2.1. version

The version number of the TAK object MUST be 0.

3.2.2. current

This field defines the set of current keys (CurrentKey) according to the signer of this Signed TALs object.

3.2.2.1. CurrentKey

This field defines a current TA Key, equivalent to [I-D.ietf-sidrops-https-tal]. This structure contains a sequence of one or more URIs and a SubjectPublicKeyInfo.

3.2.2.1.1. certificateURIs

This field is equivalent to the URI section in section 2.1 of [I-D.ietf-sidrops-https-tal]. It MUST contain at least one CertificateURI element. Each CertificateURI element contains the IA5String representation of either an rsync URI [RFC5781], or an HTTPS URI [RFC7230].

3.2.2.1.2. subjectPublicKeyInfo

This field contains a SubjectPublicKeyInfo [section 4.1.2.7 or @!RFC5280] in DER format [X.690].

3.2.3. revoked

This field contains the list of keys, identified by SubjectPublicKeyInfo, that are no longer to be used according to the signer of this document.

3.3. TAK Object Validation

To determine whether a TAK object is valid, the RP MUST perform the following steps in addition to those specified in [RFC6488]:

- The eContentType OID matches the OID described in Section 3.1
- The TAK object appears as the product of a Trust Anchor CA certificate.
- This Trust Anchor CA has published only one TAK object in its repository for this key, and this object appears on the Manifest as the only entry using the ".tak" extension (see [RFC6481]). In case more than one TAK object is found, all such objects MUST be considered invalid.
- The EE certificate of this TAK object describes its Internet Number Resources (INRs) using the "inherit" attribute
- The decoded TAK content conforms to the format defined in Section 3.2.

If the above procedure indicates that the manifest is invalid, then the TAK object MUST be discarded and treated as though no TAK object were present.

4. Maintaining multiple TA keys

As described in Section 6 a TA will most likely choose to operate two keys at any one time in order to be prepared for an emergency key roll. When a TA operates multiple keys, each key MUST use its own CA repository publication point as described in [RFC6481]. The CRL and Manifest [RFC6486] for each of these keys will be unique to each key, but the TA MUST ensure that equivalent CA certificates and RPKI signed objects are issued under each key. Note that this is similar to how such certificates and RPKI signed objects are re-issued as part of a lower level CA key roll, described in section 4 of [RFC6489].

4.1. Prepare a new TA key

The Trust Anchors MUST generate a new key pair and generate a new TA Certificate. For the Subject Information Access (see section 4.8.8.1 of [RFC6487]) this MUST use URIs that will be used by the new key to publish objects. These URIs MUST be uniqe for use by this new key only. The Internet Number Resources on this new certificate MUST be equivalent to those found on the current certificate.

The new TA certificate MUST be published under one or more new Certificate URIs for use by this new key only.

As decribed above, the TA MUST issue and publish equivalent CA certificates and RPKI signed objects under this new key.

It is RECOMMENDED that the TA now generates a new TAL [I-D.ietf-sidrops-https-tal] and verifies that the new Trust Anchor certificate can be retrieved from all locations, and that it generates the same results when it is used for top-down validation instead of (any of) the current TA key(s).

Note that the TA MAY choose to make this TAL available to Relying Parties, in particular to those that do not support TAK objects, and for inclusion in the distribution of RP software in order to minimise the overhead in bootstrapping fresh installations.

4.2. Publishing for Multiple TA Keys

If a TA uses a single remote publication server for its keys using the RPKI publication protocol [RFC8181], then it MUST include all <publish/> and <withdraw/> PDUs for the products of each of its keys

in a single query in order to ensure that they will reflect the same content at all times.

If a TA uses multiple publication servers then it is by definition inevitable that the content of different keys will be out of sync at times. In such cases the TA SHOULD ensure that the duration of these moments are limited to the shortest possible time. Furthermore the following should be observed:

- It is strongly RECOMMENDED that TAs do not issue any RPKI Signed Objects, such as ROAs [RFC6482], but limit their operations to maintaining a CRL, Manifest and CA certificates only. If an organisation maintaining a TA has an operational need for such objects then it is strongly RECOMMENDED that they operate a separate non-TA CA as a child of their TA for these operations. If this approach is used the remaining issues regarding temporary inconsistencies between multiple TA key repository publication points is greatly reduced.

- In cases where a CA certificate is revoked completely, or replaced by a certifcate with a reduced set of resources, these changes will not take effect fully until all the TA keys repository publication points have been updated. Given that TA key operations are normally performed infrequently we don't expect that this is a problem. I.e. if the revocation or shrinking of an issued CA certificate is staged for days, or weeks anyway, then experiencing a delay of several minutes for the repository publication points to all be updated is fairly insignificant.

- In cases where a CA certificate is replaced by a certifcate with an extend set of resources the TA MUST inform the receiving CA only after all its repository publication points have been updated. This ensures that the receiveing CA will not issue any products that could be invalid if an RP uses a TA key just before the CA certificate was due to be updated.

5. TAK Object Generation and Publication

A TA MAY choose to use TAK objects to communicate its set of current, and revoked keys. If a TA chooses to use TAK objects, then it SHOULD generate and publish TAK objects under each of its current keys. An exception to this rule exists when a TA has lost permanent access to one of its keys or the accompanying repository publication point. In such cases however, the key in question MUST be revoked as described below in Section 6.

A non-normative guideline for naming this object is that the filename chosen for the Signed TAL Object in the publication repository be a

value derived from the public key part of the entity's key pair, using the algorithm described for CRLs in section 2.2 of [RFC6481] for generation of filenames. The filename extension of ".tak" MUST be used to denote the object as a TAK. Note that this is in-line with filename extensions defined in section 7.2 of [RFC6481]

In order to generate the TAK Objects, the TA MUST perform the following actions:

- The TA MUST generate a key pair for a "one-time-use" EE certificate to use for the TAK
- The TA MUST generate a one-time-use EE certificate for the TAK
- This EE certificate MUST have an SIA extension access description field with an accessMethod OID value of id-ad-signedobject, where the associated accessLocation references the publication point of the TAK as an object URL.
- As described in [RFC6487], an [RFC3779] extension is required in the EE certificate used for this object. However, because the resource set is irrelevant to this object type, this certificate MUST describe its Internet Number Resources (INRs) using the "inherit" attribute, rather than explicit description of a resource set.
- This EE certificate MUST have a "notBefore" time that matches, or predates the moment that the TAK will be published.
- This EE certificate MUST have a "notAfter" time that reflects the intended duration for which this TAK will be published. If the EE certificate for a Signed TAL is expired, it MUST no longer be published, but it MAY be replaced by a newly generated TAK object with equivalent content and an updated "notAfter" time.
- The same set of current keys (see Section 3.2.2) MUST be included on each TAK object for each current key.
- The TAK object MUST include all revoked keys (see Section 3.2.3) that became revoked while the key signing the TAK in question was current.

6. Performing TA Key Rolls

6.1. Opting in to Key Rolls

6.1.1. Trust Anchor

For simplicitly let's start with a situation where a TA has only one key. The TA wants to start using TAK objects to perform key rolls in future, so it introduces a TAK object under its single key 'A'. The repository structure looks as follows (irrelevant details omitted):

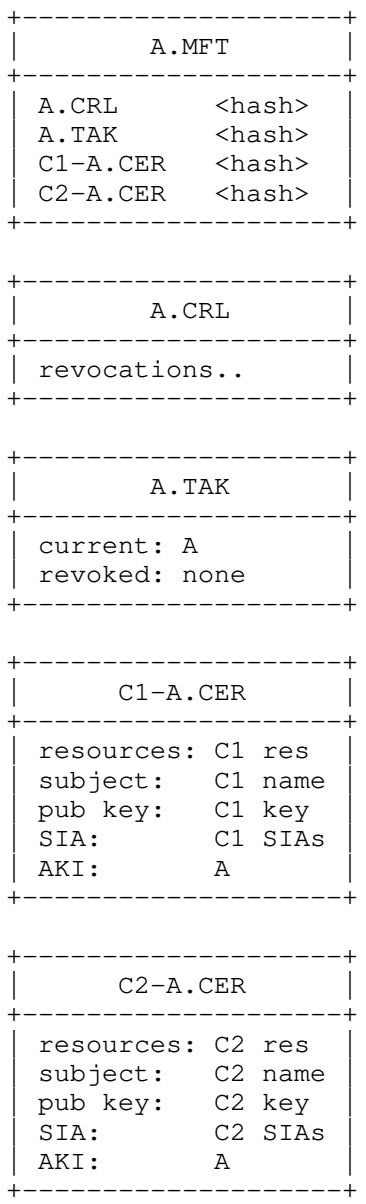

So, the TA publishes a CRL and MFT under its key A, listing a TAK object and in this case two certificates issued to children 'C1' and 'C2' signed using key A. The TAK object lists key 'A' as the only current key, and has no revoked keys.

6.1.2. Relying Parties

Relying Parties who have a TAL for key 'A' configured will discover the TAK object. If the RP does not support this object, it will reject this object but continue to validate the remaining RPKI tree as usual. If the RP does support TAK objects it will conclude that key 'A' is the one and only current key, and will proceed to validate the remaining RPKI tree as usual.

6.2. Pre-stage a New Key

6.2.1. Trust Anchor

Now the TA prestages a new key 'B' and produces equivalent CA certificates for children 'C1' and 'C2', i.e. the resources, subject name, public key and SIA etc are all equivalent, but these certificates are signed under key 'B'. (See Section 4 for a more thorough description of this). The TAK object for key 'B' recognises both keys 'A' and 'B' as current.

The repostory structure and TAK object for key B are then as follows:

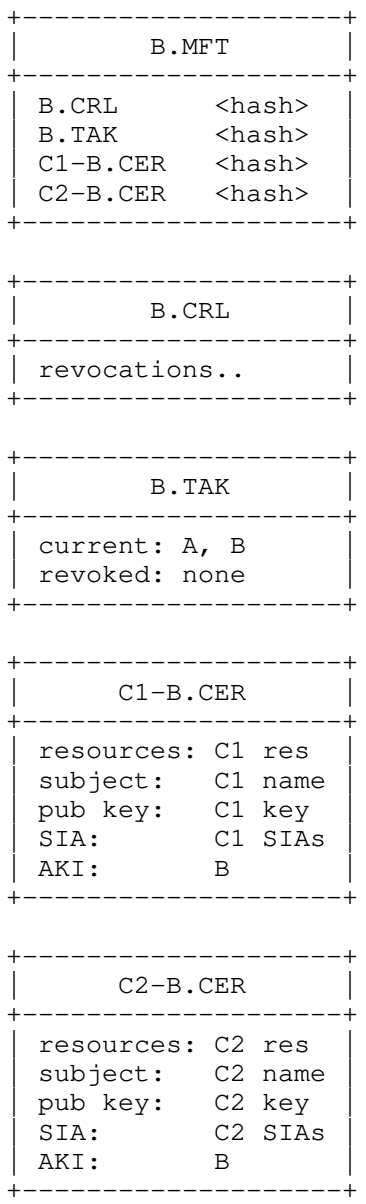

When the TA is certain that the content for key 'B' is correct, it can also update the TAK object for key 'A' to include 'B':

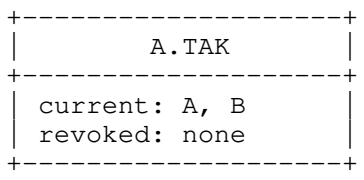

One way to do this is by generating a TAL [I-D.ietf-sidrops-https-tal] for key B and verifying that validation using this yields the same results as validation using the TAL for key A would. However, note, that it is preferred that this is done as part of an automated process that is sufficiently well tested, and that the contents of the repositories for keys 'A' and 'B' are updated as a single delta if the publication protocol [RFC8181] is used (see also: Section 5).

6.2.2. Relying Parties

Relying Parties who have a TAL for key 'A' configured will discover the TAK object. If the RP does not support this object, it will reject this object but continue to validate the remaining RPKI tree as usual. If the RP does support TAK objects it will conclude that there are now two keys 'A' and 'B', and no revoked keys that it should be aware of. Since key 'A' is still current, the RP will continue to validate the RPKI tree structure using the repository for key 'A', ignoring the non-TAK objects in the repository for key 'B'.

The result will be the same for Relying Parties who have a TAL for key 'B' configured, because both keys are equivalent at this time.

6.3. Planned Key Revocation

6.3.1. Trust Anchor

The TA has now decided that key 'A' must be revoked. It still has access to this key and the repository, so it can perform a planned key roll. In addition to revoking key 'A', the TA will also generate new key 'C' to ensure that it has at least two current keys at all times for redundancy.

Keys 'B' and 'C' will become current keys on the TAK objects for all keys: 'A', 'B' and 'C'. Key 'A' will become part of the revoked keys on the TAK objects for keys 'A' and 'B'. Note that it is not needed to list key 'A' as revoked on the TAK file for key 'C', because RPs will only learn about key 'C' at the same time as learning about the revocation of key 'A' (see also below).

The TA will publish a long-lived TAK file and MFT and CRL only for key 'A' and publish these objects as waypointers for RPs that have a TAL pointing at key 'A' before destroying key 'A'.

The resulting structure for key 'A' will be as follows:

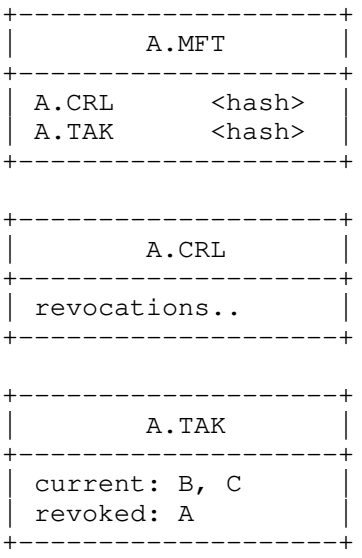

The resulting structures for keys 'B' and 'C' will be as follows:

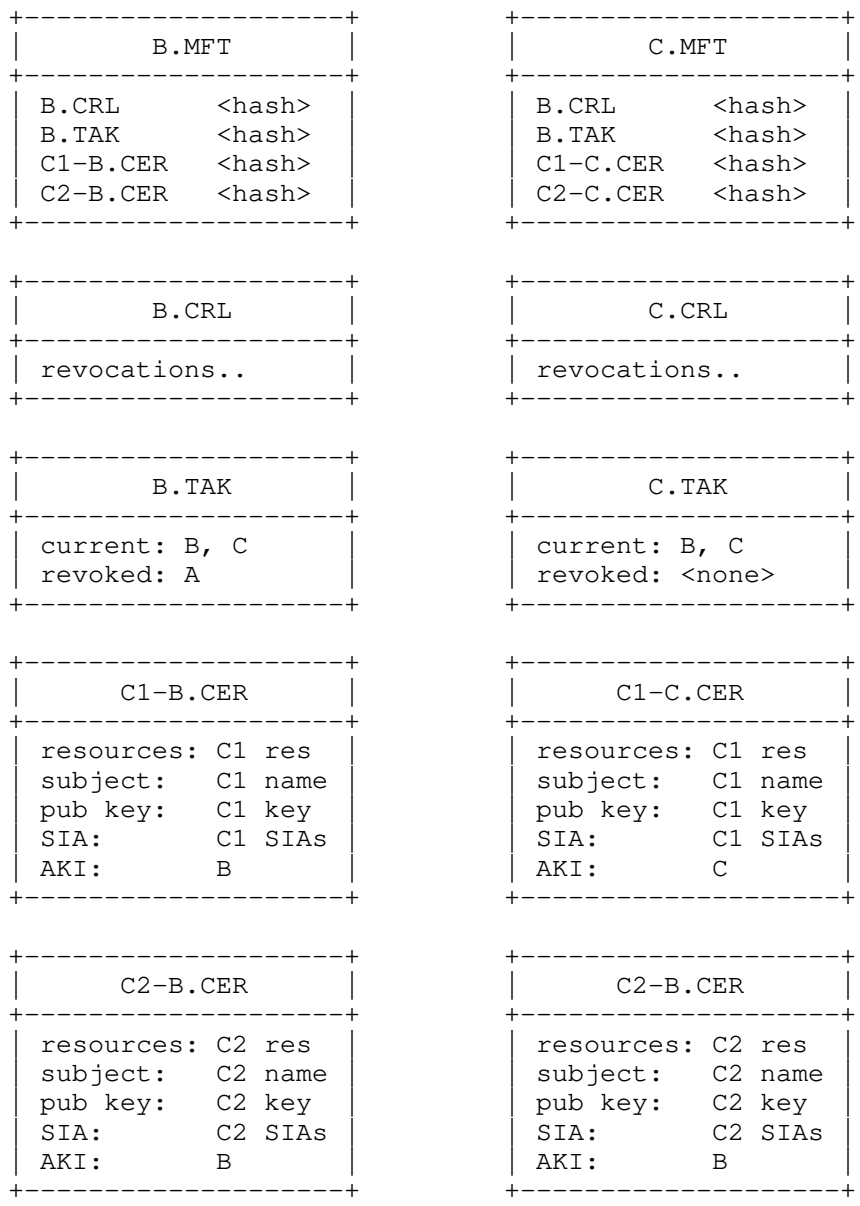

In addition to this the TA SHOULD reach out to RP vendors so that they can update the TAL included in the RP software distribution to use key 'B'.

6.3.2. Relying Parties

Relying Parties who have a TAL for key 'A' configured will discover the TAK object. If the RP does not support this object, it will reject this object but continue to validate the remaining RPKI tree as usual. In this case that means that validation will stop, because there are no more objects under key 'A'. Therefore it is important that RPs that do not support TAK files are updated to use the TAL for key 'B' through some other process.

If the RP uses a TAL for key 'A' and it supports TAK objects, it will discover that the TAL for key 'A' has keys 'B' and 'C' as current, and revokes itself. It will then proceed to process keys 'B' and 'C' and find TALs which list the same current keys. So, it will now replace its notion of the current key set for this TA based on its TAL (key 'A') with what it learned. To keep things simple the RP will now conclude that it should re-start validation using a remaining current key, in this case key either 'B' or 'C' may be used.

If the RP already had a TAL for key 'B' and it supports TAK objects, or it simply started with key 'B' because it added it to its set of current keys when this key was pre-staged (see Section 6.2), it will learn that key 'A' is revoked and therefore will not attempt to verify the TAK file for key 'A'. It will also learn about key 'C' and inspect this key's TAL, and discover that only keys 'B' and 'C' are considered current. Since it started the validation process with a key that is still current, it can proceed to validate the RPKI tree using the repository under key 'B'.

6.4. Unplanned revocation

6.4.1. Trust Anchor

Now keys 'B' and 'C' are current. The TA may have intended to revoke key 'B', essentially rolling over to key 'C' and a new key 'D', but let us suppose that the TA lost access to key 'C'. In this case the TA will simply revoke key 'C' instead, and still introduce a new key 'D'.

The major difference with the process described above for planned rolls, is that now the TA will not be able to update the TAK object, MFT or CRL for key 'C'. However, because all TAL objects for current keys are evaluated before tree validation is performed, it is safe to leave these objects in a repository. Keys 'B' and 'D' will simply mark key 'C' as being revoked.

If an RP still has a TAL pointing at key 'C' it will discover that key 'D' is added, and that key 'B' has been revoked through the TAK object published for keys 'B' and 'D'. At least, as long as the the MFT and TAK EE certificates have not expired, and the CRL and MFT are not stale.

If the TA is absolutely sure that the TAL for key 'C' never shipped with any RP distribution, then it would also be safe to delete the repository key 'C' altogether. RPs will learn that 'C' is revoked, and therefore will not even attempt to download the TAK object. However, it is hard to be certain of this and there this is NOT RECOMMENDED.

7. Deployment Considerations

Including Signed TAL objects while RP tools do not support this standard will result in these RPs rejecting these objects. It is not expected that this will result in the invalidation of any other object under a Trust Anchor.

That said, the flagging mechanism introduced here can only be relied on once a majority of RPs support it. Defining when that moment arrives is by definition something that cannot be established at the time of writing this document. Until such time, TAs SHOULD continue to generate unsigned TAL files [I-D.ietf-sidrops-https-tal], and indicate which should be considered their current TAL, and communicate them to RPs through other means.

However, once a majority of RPs support this mechanism it would be RECOMMENDED that Trust Anchor operators perform key rolls regularly. The most assured way to know that such key rolls will work is by making them a part of normal operations. Determining when this moment arrives is by definition out of scope for this document, as it should be based on operational experience.

8. IANA Considerations

8.1. OID

IANA is to add the following to the "RPKI Signed Objects" registry:

Decimal	Description	References
TBD	Trust Anchor Keys	[section 3.1]

8.2. File Extension

IANA is to add an item for the Signed TAL file extension to the "RPKI Repository Name Scheme" created by [RFC6481] as follows:

Extension	RPKI Object	References
.tak	Trust Anchor Keys	[this document]

9. Security Considerations

TBD

10. Acknowledgements

The authors wish to thank Martin Hoffmann for a thourough review of this document.

11. References

11.1. Normative References

[I-D.ietf-sidrops-https-tal]
Huston, G., Weiler, S., Michaelson, G., Kent, S., and T. Bruijnzeels, "Resource Public Key Infrastructure (RPKI) Trust Anchor Locator", draft-ietf-sidrops-https-tal-05 (work in progress), October 2018.

[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>.

[RFC3779] Lynn, C., Kent, S., and K. Seo, "X.509 Extensions for IP Addresses and AS Identifiers", RFC 3779, DOI 10.17487/RFC3779, June 2004, <https://www.rfc-editor.org/info/rfc3779>.

[RFC5781] Weiler, S., Ward, D., and R. Housley, "The rsync URI Scheme", RFC 5781, DOI 10.17487/RFC5781, February 2010, <https://www.rfc-editor.org/info/rfc5781>.

[RFC6481] Huston, G., Loomans, R., and G. Michaelson, "A Profile for Resource Certificate Repository Structure", RFC 6481, DOI 10.17487/RFC6481, February 2012, <https://www.rfc-editor.org/info/rfc6481>.

[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>.

[RFC6486] Austein, R., Huston, G., Kent, S., and M. Lepinski, "Manifests for the Resource Public Key Infrastructure (RPKI)", RFC 6486, DOI 10.17487/RFC6486, February 2012, <https://www.rfc-editor.org/info/rfc6486>.

[RFC6487] Huston, G., Michaelson, G., and R. Loomans, "A Profile for X.509 PKIX Resource Certificates", RFC 6487, DOI 10.17487/RFC6487, February 2012, <https://www.rfc-editor.org/info/rfc6487>.

[RFC6488] Lepinski, M., Chi, A., and S. Kent, "Signed Object Template for the Resource Public Key Infrastructure (RPKI)", RFC 6488, DOI 10.17487/RFC6488, February 2012, <https://www.rfc-editor.org/info/rfc6488>.

[RFC6489] Huston, G., Michaelson, G., and S. Kent, "Certification Authority (CA) Key Rollover in the Resource Public Key Infrastructure (RPKI)", BCP 174, RFC 6489, DOI 10.17487/RFC6489, February 2012, <https://www.rfc-editor.org/info/rfc6489>.

[RFC7230] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, DOI 10.17487/RFC7230, June 2014, <https://www.rfc-editor.org/info/rfc7230>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8181] Weiler, S., Sonalker, A., and R. Austein, "A Publication Protocol for the Resource Public Key Infrastructure (RPKI)", RFC 8181, DOI 10.17487/RFC8181, July 2017, <https://www.rfc-editor.org/info/rfc8181>.

[X.690] ITU-T Recommendation X.690 (2002) | ISO/IEC 8825-1:2002, "Information technology - ASN.1 encoding rules: Specification of Basic Encoding Rules (BER), Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER)", 2002.

11.2. Informative References

[RFC5652] Housley, R., "Cryptographic Message Syntax (CMS)", STD 70, RFC 5652, DOI 10.17487/RFC5652, September 2009, <https://www.rfc-editor.org/info/rfc5652>.

Authors' Addresses

Tim Bruijnzeels
NLnet Labs

Email: tim@nlnetlabs.nl
URI: https://www.nlnetlabs.nl/

Carlos Martinez
LACNIC

Email: carlos@lacnic.net
URI: https://www.lacnic.net/

Rob Austein
Dragon Research Labs

Email: sra@hactrn.net

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: July 2, 2019

R. Bush
Internet Initiative Japan
K. Patel
Arrcus
December 29, 2018

Origin Validation Signaling
draft-ymbk-sidrops-ov-signal-02

Abstract

Within a trust boundary, e.g. an operator's PoP, it may be useful to have only a few central devices do full Origin Validation using the Resource Public Key Infrastructure, and be able to signal to an internal sender that a received route fails Origin Validation. E.g. route reflectors could perform Origin Validation for a cluster and signal back to a sending client that it sent an invalid route. Routers capable of sending and receiving this signal can use the extended community described in [RFC8097]

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in [RFC2119] only when they appear in all upper case. They may also appear in lower or mixed case as English words, without normative meaning.

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 July 2, 2019.

1. Introduction

Within a routing trust boundary, e.g. an operator's Point of Presence (PoP), it may not be desirable or necessary for all routers to perform Origin Validation using the Resource Public Key Infrastructure (RPKI) per [RFC6811]. A good example is route reflectors (see [RFC4456]).

An RPKI-enabled device, an Evaluator, SHOULD signal receipt of an Invalid route back to the sender by announcing that route back to the sender marked with the BGP Prefix Origin Validation State Extended Community as defined in [RFC8097] with a last octet having the value 2, meaning "Invalid." In the rest of this document we take the liberty of calling it the "community."

We use the term "Sender" to refer to the router announcing routes to the device evaluating the Origin Validation of the announcements. Beware that the Sender receives signaling back from the Evaluator, which can be somewhat confusing.

We use the term "Evaluator" to describe the device receiving routing announcements from senders, applying RPKI-based Origin Validation, and possibly signaling route Invalidity back to the sender(s).

2. Suggested Reading

It is assumed that the reader understands BGP, [RFC4271], the RPKI, [RFC6480], RPKI-based Prefix Validation, [RFC6811], and the BGP Prefix Origin Validation State Extended Community as described in [RFC8097].

3. Trust Boundary

As a general rule, we discourage 'outsourcing trust,' i.e. letting others make security decisions for us. But there are operational environments with a somewhat wide trust boundary, a single operator's PoP for example.

This is not outsourcing trust; this is remote decision making. It is not letting a third party make the decision; it is simply doing it on a different computer. It's trust in a distributed system, where what is (sometimes) called the Policy Decision Point is not the same as the Policy Enforcement Point.

As described in [RFC7115], a PoP might have a single RPKI Cache, hence all trust is vested in it. So it is reasonable that routers in that PoP could share Origin Validation results instead of each doing full validation.

An [RFC4456] Route Reflector Cluster is an obvious candidate for this approach. The route reflector(s) would perform Origin Validation and signal an Invalid route back to the sending client.

[RFC8097] provides the obvious signaling mechanism, the BGP Prefix Origin Validation State Extended Community. The device performing OV SHOULD signal back to the sender by announcing the offending prefix marked with the extended community with the last octet having the value 2, indicating an Invalid route.

4. The OV Signaling Capability

Unfortunately, the router sending the Invalid announcement is not normally expecting to receive it back. Therefore, both parties MUST agree on this feature by using a BGP Capability [RFC5492].

To advertise the OV Signaling Capability to a peer, a BGP speaker uses BGP Capabilities Advertisement [RFC5492]. By advertising the OV Signaling Capability to a peer, a BGP speaker conveys that it is able to send, receive, and properly handle OV Signaling using the community.

A peer which does not advertise this capability MUST NOT send OV Signaling, and BGP OV Signaling MUST NOT be sent to it.

The OV Signaling Capability is a new BGP Capability defined with Capability code [TBD] and Capability length 0.

5. Recommended Action

This section assumes that the OV Signaling Capability has been negotiated by the sending and receiving routers.

An Evaluating device which performs Origin Validation on a route received from a capable sender and finds a prefix with a particular origin AS to be Invalid (in the [RFC6811] sense), MUST announce that prefix back to the sending router from which it was received with the Invalid origin AS and the addition of the community with the last octet being 2.

A sender receiving the returned prefix announcement so marked MUST treat it the way it would treat an Invalid origin that it itself detected. It should withdraw all routes it had announced to that prefix with the Invalid origin AS. This includes withdrawing any instances of additional paths with that origin AS advertised under [RFC7911].

For a sender to properly evaluate the community returned by the evaluator, the sender MUST recognize the community before loop detection. This is a change to the Phase 2 Route Selection process of [RFC4271] Section 9.1.2.

If a sender originally received the Invalid route from an evaluator within its trust boundary with which it has negotiated the OV Signaling Capability, it MAY also propagate that signal to the original sender.

6. Security Considerations

As with all communities which cause semantic change, this use of the community may be abused as an attack vector. Therefore the operator MUST configure their incoming external border to strip the community.

As the BGP sessions are already established using whatever channel security the operator chooses or not, this change specifies no additional channel or object security. Of course, the BGP transport should be protected for integrity and authentication. TCP-MD5 [RFC2385] is available on almost all platforms. If more modern methods are available, they should be used.

Outsourcing security is usually considered bad policy.
Section Section 3 above discusses why that is not really the case here.

Otherwise, this document does not create security considerations beyond those of [RFC6811].

7. IANA Considerations

This document requests the IANA assign the "OV Signaling Capability" to the BGP Capabilities described in Section 2.1 in the "Capability Codes" registry's "IETF Review" range [RFC8126].. This document is the reference for the new capability.

8. Acknowledgments

Thanks to Steve Bellovin for a serious security review, and Rob Austein for a useful security snark.

9. References

9.1. 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, <http://www.rfc-editor.org/info/rfc2119>.

[RFC2385] Heffernan, A., "Protection of BGP Sessions via the TCP MD5 Signature Option", RFC 2385, DOI 10.17487/RFC2385, August 1998, <http://www.rfc-editor.org/info/rfc2385>.

[RFC4456] Bates, T., Chen, E., and R. Chandra, "BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)", RFC 4456, DOI 10.17487/RFC4456, April 2006, <http://www.rfc-editor.org/info/rfc4456>.

[RFC5492] Scudder, J. and R. Chandra, "Capabilities Advertisement with BGP-4", RFC 5492, DOI 10.17487/RFC5492, February 2009, <http://www.rfc-editor.org/info/rfc5492>.

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

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

[RFC7911] Walton, D., Retana, A., Chen, E., and J. Scudder, "Advertisement of Multiple Paths in BGP", RFC 7911, DOI 10.17487/RFC7911, July 2016, <http://www.rfc-editor.org/info/rfc7911>.

[RFC8097] Mohapatra, P., Patel, K., Scudder, J., Ward, D., and R. Bush, "BGP Prefix Origin Validation State Extended Community", RFC 8097, DOI 10.17487/RFC8097, March 2017, <http://www.rfc-editor.org/info/rfc8097>.

[RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <http://www.rfc-editor.org/info/rfc8126>.

9.2. Informative References

[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, <http://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, <http://www.rfc-editor.org/info/rfc6480>.

Authors' Addresses

Randy Bush
Internet Initiative Japan
5147 Crystal Springs
Bainbridge Island, Washington 98110
US

Email: randy@psg.com

Keyur Patel
Arrcus
2077 Gateway Place, Suite #250
San Jose, CA 95119
United States of America

Email: keyur@arrcus.com