

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: November 23, 2012

P. Mohapatra
Cisco Systems
J. Scudder
Juniper Networks
D. Ward
Cisco Systems
R. Bush
Internet Initiative Japan, Inc.
R. Austein
Dragon Research Labs
May 22, 2012

BGP Prefix Origin Validation
draft-ietf-sidr-pfx-validate-06

Abstract

To help reduce well-known threats against BGP including prefix mis-announcing and monkey-in-the-middle attacks, one of the security requirements is the ability to validate the origination AS of BGP routes. More specifically, one needs to validate that the AS number claiming to originate an address prefix (as derived from the AS_PATH attribute of the BGP route) is in fact authorized by the prefix holder to do so. This document describes a simple validation mechanism to partially satisfy this requirement.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

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

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

This Internet-Draft will expire on November 23, 2012.

Copyright Notice

Copyright (c) 2012 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 (<http://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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction	4
1.1.	Requirements Language	5
2.	Prefix-to-AS Mapping Database	5
2.1.	Pseudo-Code	7
3.	Policy Control	8
4.	Interaction with Local Cache	8
5.	Deployment Considerations	8
6.	Acknowledgments	9
7.	IANA Considerations	9
8.	Security Considerations	9
9.	References	10
9.1.	Normative References	10
9.2.	Informational References	10
	Authors' Addresses	10

1. Introduction

A BGP route associates an address prefix with a set of autonomous systems (AS) that identify the interdomain path the prefix has traversed in the form of BGP announcements. This set is represented as the AS_PATH attribute in BGP [[RFC4271](#)] and starts with the AS that originated the prefix. To help reduce well-known threats against BGP including prefix mis-announcing and monkey-in-the-middle attacks, one of the security requirements is the ability to validate the origination AS of BGP routes. More specifically, one needs to validate that the AS number claiming to originate an address prefix (as derived from the AS_PATH attribute of the BGP route) is in fact authorized by the prefix holder to do so. This document describes a simple validation mechanism to partially satisfy this requirement.

The Resource Public Key Infrastructure (RPKI) describes an approach to build a formally verifiable database of IP addresses and AS numbers as resources. The overall architecture of RPKI as defined in [[RFC6480](#)] consists of three main components:

- o A public key infrastructure (PKI) with the necessary certificate objects,
- o Digitally signed routing objects,
- o A distributed repository system to hold the objects that would also support periodic retrieval.

The RPKI system is based on resource certificates that define extensions to X.509 to represent IP addresses and AS identifiers [[RFC3779](#)], thus the name RPKI. Route Origin Authorizations (ROA) [[RFC6482](#)] are separate digitally signed objects that define associations between ASes and IP address blocks. Finally the repository system is operated in a distributed fashion through the IANA, RIR hierarchy, and ISPs.

In order to benefit from the RPKI system, it is envisioned that relying parties either at AS or organization level obtain a local copy of the signed object collection, verify the signatures, and process them. The cache must also be refreshed periodically. The exact access mechanism used to retrieve the local cache is beyond the scope of this document.

Individual BGP speakers can utilize the processed data contained in the local cache to validate BGP announcements. The protocol details to retrieve the processed data from the local cache to the BGP speakers is beyond the scope of this document (refer to [[I-D.ietf-sidr-rpki-rtr](#)] for such a mechanism). This document

proposes a means by which a BGP speaker can make use of the processed data in order to assign a "validation state" to each prefix in a received BGP UPDATE message.

Note that the complete path attestation against the AS_PATH attribute of a route is outside the scope of this document.

Although RPKI provides the context for this draft, it is equally possible to use any other database which is able to map prefixes to their authorized origin ASes. Each distinct database will have its own particular operational and security characteristics; such characteristics are beyond the scope of this document.

1.1. Requirements Language

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 [RFC 2119](#) [[RFC2119](#)].

2. Prefix-to-AS Mapping Database

The BGP speaker loads validated objects from the cache into local storage. The objects loaded have the content (IP address, prefix length, maximum length, origin AS number). We refer to such a locally stored object as a "Validated ROA Payload" or "VRP".

We define several terms in addition to "VRP". Where these terms are used, they are capitalized:

- o Prefix: (IP address, prefix length), interpreted as is customary (see [[RFC4632](#)]).
- o Route: Data derived from a received BGP UPDATE, as defined in [[RFC4271](#)], [Section 1.1](#). The Route includes one Prefix and an AS_PATH; it may include other attributes to characterize the prefix.
- o VRP Prefix: The Prefix from a VRP.
- o VRP ASN: The origin AS number from a VRP.
- o Route Prefix: The Prefix derived from a route.
- o Route Origin ASN: The origin AS number derived from a Route as follows:

- * the rightmost AS in the final segment of the AS_PATH attribute in the Route if that segment is of type AS_SEQUENCE, or
 - * the BGP speaker's own AS number if that segment is of type AS_CONFED_SEQUENCE or AS_CONFED_SET or if the AS_PATH is empty, or
 - * the distinguished value "NONE" if the final segment of the AS_PATH attribute is of any other type.
- o Covered: A Route Prefix is said to be Covered by a VRP when the VRP prefix length is less than or equal to the Route prefix length, and the VRP prefix address and the Route prefix address are identical for all bits specified by the VRP prefix length. (I.e. the Route prefix is either identical to the VRP prefix or a more specific of the VRP prefix.)
 - o Matched: A Route Prefix is said to be Matched by a VRP when the Route Prefix is Covered by that VRP and in addition, the Route prefix length is less than or equal to the VRP maximum length and the Route Origin ASN is equal to the VRP ASN.

Given these definitions, any given BGP Route will be found to have one of the following "validation states":

- o NotFound: No VRP Covers the Route Prefix.
- o Valid: At least one VRP Matches the Route Prefix.
- o Invalid: At least one VRP Covers the Route Prefix, but no VRP Matches it.

We observe that no VRP can have the value "NONE" as its VRP ASN. Thus a Route whose Origin ASN is "NONE" cannot be Matched by any VRP. Similarly, no valid Route can have an Origin ASN of zero [[I-D.ietf-idr-as0](#)]. Thus no Route can be Matched by a VRP whose ASN is zero.

When a BGP speaker receives an UPDATE from a neighbor, it SHOULD perform a lookup as described above for each of the Routes in the UPDATE message. The lookup SHOULD also be applied to routes which are redistributed into BGP from another source, such as another protocol or a locally defined static route. An implementation MAY provide configuration options to control which routes the lookup is applied to. The "validation state" of the Route MUST be set to reflect the result of the lookup. The implementation should consider the "validation state" as described in the document as a local property or attribute of the Route. If validation is not performed

on a Route, the implementation SHOULD initialize the "validation state" of such a route to "NotFound".

Use of the validation state is discussed in [Section 3](#) and [Section 5](#). An implementation MUST NOT exclude a route from the Adj-RIB-In or from consideration in the decision process as a side-effect of its validation state, unless explicitly configured to do so.

We observe that a Route can be Matched or Covered by more than one VRP. This procedure does not mandate an order in which VRPs must be visited; however, the "validation state" output is fully determined.

[2.1.](#) Pseudo-Code

The following pseudo-code illustrates the procedure above. In case of ambiguity, the procedure above, rather than the pseudo-code, should be taken as authoritative.

```
result = BGP_PFXV_STATE_NOT_FOUND;

//Iterate through all the Covering entries in the local VRP
//database, pfx_validate_table.
entry = next_lookup_result(pfx_validate_table, route_prefix);

while (entry != NULL) {
    prefix_exists = TRUE;

    if (route_prefix_length <= entry->max_length) {
        if (route_origin_as != NONE
            && entry->origin_as != 0
            && route_origin_as == entry->origin_as) {
            result = BGP_PFXV_STATE_VALID;
            return (result);
        }
    }
    entry = next_lookup_result(pfx_validate_table, input.prefix);
}

//If one or more VRP entries Covered the route prefix, but
//no one Matched, return "Invalid" validation state.
if (prefix_exists == TRUE) {
    result = BGP_PFXV_STATE_INVALID;
}

return (result);
```


3. Policy Control

An implementation MUST provide the ability to match and set the validation state of routes as part of its route policy filtering function. Use of validation state in route policy is elaborated in [Section 5](#). For more details on operational policy considerations, see [\[I-D.ietf-sidr-origin-ops\]](#).

An implementation MUST support Four-Octet AS Numbers, [\[RFC4893\]](#).

4. Interaction with Local Cache

Each BGP speaker supporting prefix validation as described in this document is expected to communicate with one or more RPKI caches, each of which stores a local copy of the global RPKI database. The protocol mechanisms used to gather and validate these data and present them to BGP speakers are described in [\[I-D.ietf-sidr-rpki-rtr\]](#).

The prefix-to-AS mappings used by the BGP speaker are expected to be updated over time. When a mapping is added or deleted, the implementation MUST re-validate any affected prefixes. An "affected prefix" is any prefix that was matched by a deleted or updated mapping, or could be matched by an added mapping.

5. Deployment Considerations

Once a Route is selected for validation, it is categorized according the procedure given in [Section 2](#). Subsequently, routing policy as discussed in [Section 3](#) can be used to take action based on the validation state.

Policies which could be implemented include filtering routes based on validation state (for example, rejecting all "invalid" routes) or adjusting a route's degree of preference in the selection algorithm based on its validation state. The latter could be accomplished by adjusting the value of such attributes as LOCAL_PREF. Considering invalid routes for BGP decision process is a pure local policy matter and should be done with utmost care.

In some cases (particularly when the selection algorithm is influenced by the adjustment of a route property that is not propagated into IBGP) it could be necessary for routing correctness to propagate the validation state to the IBGP peer. This can be accomplished on the sending side by setting a community or extended community based on the validation state, and on the receiving side by

matching the (extended) community and setting the validation state.

6. Acknowledgments

The authors wish to thank Rex Fernando, Hannes Gredler, Mouhcine Guennoun, Russ Housley, Junaid Israr, Miya Kohno, Shin Miyakawa, Taka Mizuguchi, Hussein Mouftah, Keyur Patel, Tomoya Yoshida, Kannan Varadhan, Wes George, Jay Borkenhagen, and Sandra Murphy. The authors are grateful for the feedback from the members of the SIDR working group.

Junaid Israr's contribution to this specification was part of his PhD research work and thesis at University of Ottawa.

7. IANA Considerations

8. Security Considerations

Although this specification discusses one portion of a system to validate BGP routes, it should be noted that it relies on a database (RPKI or other) to provide validation information. As such, the security properties of that database must be considered in order to determine the security provided by the overall solution. If "invalid" routes are blocked as this specification suggests, the overall system provides a possible denial-of-service vector, for example if an attacker is able to inject or remove one or more records in the validation database, it could lead an otherwise valid route to be marked as invalid.

In addition, this system is only able to provide limited protection against a determined attacker -- the attacker need only prepend the "valid" source AS to a forged BGP route announcement in order to defeat the protection provided by this system.

This mechanism does not protect against "AS in the middle attacks" or provide any path validation. It only attempts to verify the origin. In general, this system should be thought of more as a protection against misconfiguration than as true "security" in the strong sense.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3779] Lynn, C., Kent, S., and K. Seo, "X.509 Extensions for IP Addresses and AS Identifiers", [RFC 3779](#), June 2004.
- [RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), January 2006.
- [RFC4632] Fuller, V. and T. Li, "Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan", [BCP 122](#), [RFC 4632](#), August 2006.
- [RFC4893] Vohra, Q. and E. Chen, "BGP Support for Four-octet AS Number Space", [RFC 4893](#), May 2007.
- [RFC6482] Lepinski, M., Kent, S., and D. Kong, "A Profile for Route Origin Authorizations (ROAs)", [RFC 6482](#), February 2012.

9.2. Informational References

- [I-D.ietf-idr-as0]
Kumari, W., Bush, R., Schiller, H., and K. Patel,
"Codification of AS 0 processing.", [draft-ietf-idr-as0-04](#)
(work in progress), May 2012.
- [I-D.ietf-sidr-origin-ops]
Bush, R., "RPKI-Based Origin Validation Operation",
[draft-ietf-sidr-origin-ops-15](#) (work in progress),
March 2012.
- [I-D.ietf-sidr-rpki-rtr]
Bush, R. and R. Austein, "The RPKI/Router Protocol",
[draft-ietf-sidr-rpki-rtr-26](#) (work in progress),
February 2012.
- [RFC6480] Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", [RFC 6480](#), February 2012.

Authors' Addresses

Pradosh Mohapatra
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95134
USA

Email: pmohapat@cisco.com

John Scudder
Juniper Networks
1194 N. Mathilda Ave
Sunnyvale, CA 94089
USA

Email: jgs@juniper.net

David Ward
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95134
USA

Email: dward@cisco.com

Randy Bush
Internet Initiative Japan, Inc.
5147 Crystal Springs
Bainbridge Island, Washington 98110
USA

Email: randy@psg.com

Rob Austein
Dragon Research Labs

Email: sra@hactrn.net

