

Network Working Group
Internet-Draft
Intended status: Informational
Expires: October 18, 2014

S. Bellovin
Columbia University
R. Bush
Internet Initiative Japan
D. Ward
Cisco Systems
April 16, 2014

Security Requirements for BGP Path Validation
draft-ietf-sidr-bgpsec-reqs-10

Abstract

This document describes requirements for a BGP security protocol design to provide cryptographic assurance that the origin AS had the right to announce the prefix and to provide assurance of the AS Path of the announcement.

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 [RFC 2119](#) [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 <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 October 18, 2014.

Copyright Notice

Copyright (c) 2014 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.

Table of Contents

1.	Introduction	2
2.	Recommended Reading	3
3.	General Requirements	3
4.	BGP UPDATE Security Requirements	5
5.	IANA Considerations	6
6.	Security Considerations	6
7.	Acknowledgments	6
8.	References	7
8.1.	Normative References	7
8.2.	Informative References	7
	Authors' Addresses	8

[1.](#) Introduction

RPKI-based Origin Validation, [[RFC6811](#)], provides a measure of resilience to accidental mis-origination of prefixes. But it provides neither cryptographic assurance (announcements are not signed), nor assurance of the AS Path of the announcement.

This document describes requirements to be placed on a BGP security protocol, herein termed BGPsec, intended to rectify these gaps.

The threat model assumed here is documented in [[RFC4593](#)] and [[I-D.ietf-sidr-bgpsec-threats](#)].

As noted in the threat model, [[I-D.ietf-sidr-bgpsec-threats](#)], this work is limited to threats to the BGP protocol. Issues of business relationship conformance, while quite important to operators, are not security issues per se, and are outside the scope of this document. It is hoped that these issues will be better understood in the future.

2. Recommended Reading

This document assumes knowledge of the RPKI see [[RFC6480](#)], the RPKI Repository Structure, see [[RFC6481](#)].

This document assumes ongoing incremental deployment of ROAs, see [[RFC6482](#)], the RPKI to Router Protocol, see [[RFC6810](#)], and RPKI-based Prefix Validation, see [[RFC6811](#)].

And, of course, a knowledge of BGP [[RFC4271](#)] is required.

3. General Requirements

The following are general requirements for a BGPsec protocol:

- 3.1 A BGPsec design must allow the receiver of a BGP announcement to determine, to a strong level of certainty, that the originating AS in the received PATH attribute possessed the authority to announce the prefix in the announcement.
- 3.2 A BGPsec design must allow the receiver of a BGP announcement to determine, to a strong level of certainty, that the received PATH attribute accurately represents the sequence of eBGP exchanges that propagated the prefix from the origin AS to the receiver, particularly if an AS has added or deleted any AS number other than its own in the path attribute. This includes modification to the number of AS prepends.
- 3.3 BGP attributes other than the AS_PATH are used only locally, or have meaning only between immediate neighbors, may be modified by intermediate systems, and figure less prominently in the decision process. Consequently, it is not appropriate to try to protect such attributes in a BGPsec design.
- 3.4 A BGPsec design MUST be amenable to incremental deployment. This implies that incompatible protocol capabilities MUST be negotiated.
- 3.5 A BGPsec design MUST provide analysis of the operational considerations for deployment and particularly of incremental deployment, e.g, contiguous islands, non-contiguous islands, universal deployment, etc.
- 3.6 As proofs of possession and authentication may require cryptographic payloads and/or storage and computation, likely increasing processing and memory requirements on routers, a BGPsec design MAY require use of new hardware. I.e.

compatibility with current hardware abilities is not a requirement that this document imposes on a solution.

- 3.7 A BGPsec design need not prevent attacks on data plane traffic. It need not provide assurance that the data plane even follows the control plane.
- 3.8 A BGPsec design MUST resist attacks by an enemy who has access to the inter-router link layer, per [Section 3.1.1.2 of \[RFC4593\]](#). In particular, such a design must provide mechanisms for authentication of all data, including protecting against message insertion, deletion, modification, or replay. Mechanisms that suffice include TCP sessions authenticated with TCP-AO [\[RFC5925\]](#), IPsec [\[RFC4301\]](#), or TLS [\[RFC5246\]](#).
- 3.9 It is assumed that a BGPsec design will require information about holdings of address space and ASNs, and assertions about binding of address space to ASNs. A BGPsec design MAY make use of a security infrastructure (e.g., a PKI) to distribute such authenticated data.
- 3.10 It is entirely OPTIONAL to secure AS SETs and prefix aggregation. The long range solution to this is the deprecation of AS-SETs, see [\[RFC6472\]](#).
- 3.11 If a BGPsec design uses signed prefixes, given the difficulty of splitting a signed message while preserving the signature, it need NOT handle multiple prefixes in a single UPDATE PDU.
- 3.12 A BGPsec design MUST enable each BGPsec speaker to configure use of the security mechanism on a per-peer basis.
- 3.13 A BGPsec design MUST provide backward compatibility in the message formatting, transmission, and processing of routing information carried through a mixed security environment. Message formatting in a fully secured environment MAY be handled in a non-backward compatible manner.
- 3.14 While the trust level of a route should be determined by the BGPsec protocol, local routing preference and policy MUST then be applied to best path and other routing decisions.
- 3.15 A BGPsec design MUST support 'transparent' route servers, meaning that the AS of the route server is not counted in downstream BGP AS-path-length tie-breaking decisions.
- 3.16 A BGPsec design MUST support AS aliasing. This technique is not well-defined or universally implemented, but is being

documented in [[I-D.ga-idr-as-migration](#)]. A BGPsec design SHOULD accommodate AS 'migration' techniques such as common proprietary and non-standard methods which allow a router to have two AS identities, without lengthening the effective AS Path.

- 3.17 If a BGPsec design makes use of a security infrastructure, that infrastructure SHOULD enable each network operator to select the entities it will trust when authenticating data in the security infrastructure. See, for example, [[I-D.ietf-sidr-lta-use-cases](#)].
- 3.18 A BGPsec design MUST NOT require operators to reveal more than is currently revealed in the operational inter-domain routing environment, other than the inclusion of necessary security credentials to allow others to ascertain for themselves the necessary degree of assurance regarding the validity of NLRI received via BGPsec. This includes peering, customer, and provider relationships, an ISP's internal infrastructure, etc. It is understood that some data are revealed to the savvy seeker by BGP, traceroute, etc. today.
- 3.19 A BGPsec design MUST signal (logging, SNMP, ...) security exceptions which are significant to the operator. The specific data to be signaled are an implementation matter.
- 3.20 Any routing information database MUST be re-authenticated periodically or in an event-driven manner, especially in response to events such as, for example, PKI updates.
- 3.21 Any inter-AS use of cryptographic hashes or signatures, MUST provide mechanisms for algorithm agility.
- 3.22 A BGPsec design SHOULD NOT presume to know the intent of the originator of a NLRI, nor that of any AS on the AS Path, other than that they intended to pass it to the next AS in the Path.
- 3.23 A BGPsec listener SHOULD NOT trust non-BGPsec markings, such as communities, across trust boundaries.

4. BGP UPDATE Security Requirements

The following requirements MUST be met in the processing of BGP UPDATE messages:

- 4.1 A BGPsec design MUST enable each recipient of an UPDATE to formally validate that the origin AS in the message is

authorized to originate a route to the prefix(es) in the message.

- 4.2 A BGPsec design MUST enable the recipient of an UPDATE to formally determine that the NLRI has traversed the AS path indicated in the UPDATE. Note that this is more stringent than showing that the path is merely not impossible.
- 4.3 Replay of BGP UPDATE messages need not be completely prevented, but a BGPsec design SHOULD provide a mechanism to control the window of exposure to replay attacks.
- 4.4 A BGPsec design SHOULD provide some level of assurance that the origin of a prefix is still 'alive', i.e. that a monkey in the middle has not withheld a WITHDRAW message or the effects thereof.
- 4.5 The AS Path of an UPDATE message SHOULD be able to be authenticated as the message is processed.
- 4.6 Normal sanity checks of received announcements MUST be done, e.g. verification that the first element of the AS_PATH list corresponds to the locally configured AS of the peer from which the UPDATE was received.
- 4.7 The output of a router applying BGPsec validation to a received UPDATE MUST be unequivocal and conform to a fully specified state in the design.

5. IANA Considerations

This document asks nothing of the IANA.

6. Security Considerations

The data plane might not follow the control plane.

Security for subscriber traffic is outside the scope of this document, and of BGP security in general. IETF standards for payload data security should be employed. While adoption of BGP security measures may ameliorate some classes of attacks on traffic, these measures are not a substitute for use of subscriber-based security.

7. Acknowledgments

The author wishes to thank the authors of [[I-D.ietf-rpsec-bgpsec](#)] from whom we liberally stole, Russ Housley, Geoff Huston, Steve Kent,

Sandy Murphy, Eric Osterweil, John Scudder, Kotikalapudi Sriram, Sam Weiler, and a number of others.

8. References

8.1. Normative References

- [I-D.ietf-sidr-bgpsec-threats]
Kent, S. and A. Chi, "Threat Model for BGP Path Security", [draft-ietf-sidr-bgpsec-threats-05](#) (work in progress), March 2013.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4593] Barbir, A., Murphy, S., and Y. Yang, "Generic Threats to Routing Protocols", [RFC 4593](#), October 2006.
- [RFC5925] Touch, J., Mankin, A., and R. Bonica, "The TCP Authentication Option", [RFC 5925](#), June 2010.

8.2. Informative References

- [I-D.ga-idr-as-migration]
George, W. and S. Amante, "Autonomous System (AS) Migration Features and Their Effects on the BGP AS_PATH Attribute", [draft-ga-idr-as-migration-01](#) (work in progress), February 2013.
- [I-D.ietf-rpsec-bgpsecrec]
Christian, B. and T. Tauber, "BGP Security Requirements", [draft-ietf-rpsec-bgpsecrec-10](#) (work in progress), November 2008.
- [I-D.ietf-sidr-lta-use-cases]
Bush, R., "RPKI Local Trust Anchor Use Cases", [draft-ietf-sidr-lta-use-cases-00](#) (work in progress), February 2014.
- [RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), January 2006.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", [RFC 4301](#), December 2005.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", [RFC 5246](#), August 2008.

- [RFC6472] Kumari, W. and K. Sriram, "Recommendation for Not Using AS_SET and AS_CONFED_SET in BGP", [BCP 172](#), [RFC 6472](#), December 2011.
- [RFC6480] Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", [RFC 6480](#), February 2012.
- [RFC6481] Huston, G., Loomans, R., and G. Michaelson, "A Profile for Resource Certificate Repository Structure", [RFC 6481](#), February 2012.
- [RFC6482] Lepinski, M., Kent, S., and D. Kong, "A Profile for Route Origin Authorizations (ROAs)", [RFC 6482](#), February 2012.
- [RFC6810] Bush, R. and R. Austein, "The Resource Public Key Infrastructure (RPKI) to Router Protocol", [RFC 6810](#), January 2013.
- [RFC6811] Mohapatra, P., Scudder, J., Ward, D., Bush, R., and R. Austein, "BGP Prefix Origin Validation", [RFC 6811](#), January 2013.

Authors' Addresses

Steven M. Bellovin
Columbia University
1214 Amsterdam Avenue, MC 0401
New York, New York 10027
US

Phone: +1 212 939 7149
Email: bellovin@acm.org

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

Email: randy@psg.com

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

Email: dward@cisco.com