Network Working Group
Internet-Draft
Intended status: BCP
Expires: May 12, 2011

R. Bush
IIJ

November 8, 2010

RPKI-Based Origin Validation Operations
draft-ymbk-rpki-origin-ops-00

Abstract

Deployment of the RPKI-based BGP origin validation has many operational considerations. This document attempts to collect and present them. It is expected to evolve as RPKI-based origin validation is deployed and the dynamics are better understood.

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].

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79. This document may not be modified, and derivative works of it may not be created, and it may not be published except as an Internet-Draft.

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 May 12, 2011.

Table of Contents

1. Introduction

RPKI-based origin validation relies on widespread propagation of the Resource Public Key Infrastructure (RPKI) [I-D.ietf-sidr-arch]. How the RPKI is distributed and maintained globally is a serious concern from many aspects.

The global RPKI has yet to be deployed, only a testbed exists, and some beta testing is being done by the IANA and some RIRs. It is expected to be deployed incrementally over a number of years. It is thought that origin validation based on the RPKI will deploy over the next year to five years.

Origin validation only need be done by an AS's border routers and is designed so that it can be used to protect announcements which are originated by large providers, upstreams and downstreams, and by small stub/entetprise/edge routers.

Origin validation has been designed to be deployed on current routers without hardware upgrade. It should be used by everyone from large backbones to small stub/entetprise/edge routers.

RPKI-based origin validation has been designed so that, with prudent local routing policies, there is no liability that normal Internet routing is threatened by unprudent deployment of the global RPKI, see Section 5.

2. Suggested Reading

It is assumed that the reader understands BGP, [RFC4271], the RPKI, see [I-D.ietf-sidr-arch], the RPKI Repository Structure, see [I-D.ietf-sidr-repos-struct], ROAs, see [I-D.ietf-sidr-roa-format], the RPKI to Router Protocol, see [I-D.ymbk-rpki-rtr-protocol], and RPKI-based Prefix Validation, see [I-D.pmohapat-sidr-pfx-validate].

3. RPKI Distribution and Maintenance

The RPKI is a distributed database containing certificates, CRLs, manifests, ROAs, and Ghostbuster Records as described in [I-D.ietf-sidr-repos-struct]. Policies and considerations for RPKI object generation and maintenance are discussed elsewhere.

A local valid cache containing all RPKI data may be gathered from the global distributed database using the rsync protocol and a validation tool such as rcynic.

Validated caches may also be created and maintained from other validated caches. An operator should take maximum advantage of this feature to minimize load on the global distributed RPKI database.

As RPKI-based origin validation relies on the availability of RPKI data, operators will likely want border routers to have one or more nearby caches.

For redundancy, a router may peer with more than one cache at the same time. Peering with two or more, one local and others remote, is recommended.

If an operator or site trusts upstreams to carry their traffic, they might as well trust the RPKI data those upstreams cache and feed off of those caches. Note that this places an obligation on those upstreams to maintain fresh and reliable caches.

A transit provider or a network with peers will want to validate origins in announcements made by downstreams and peers. They still may choose to trust the caches provided by their upstreams.

4. Within a Network

Origin validation need only be done by edge routers in a network, those which border other networks/ASs.

A validating router will use the result of origin validation to influence local policy within its network, see Section 5. In deployment this policy should fit into the AS's existing policy, preferences, etc. This allows a network to incrementally deploy validation capable border routers.

eBGP speakers which face more critical peers or up/downstreams would be candidates for the earliest deployment. Validating more critical received announcements should be considered in partial deployment.

5. Routing Policy

Origin validation based on the RPKI merely marks a received announcement as having an origin which is Validated, Unknown, or Invalid. How this is used in routing is up to the router operator's local policy. See [I-D.pmohapat-sidr-pfx-validate].

Reasonable application of local policy should be designed eliminate the threat of unroutability of prefixes due to ill-advised or incorrect certification policies.

As origin validation will be rolled out over years coverage will be spotty for a long time. Hence a normal operator's policy should not be overly strict, perhaps preferring valid announcements and giving very low preference, but still using, invalid announcements.

Some may choose to use the large Local-Preference hammer. Others might choose to let AS-Path rule and set their internal metric, which comes after AS-Path in the BGP decision process.

Certainly, routing on unknown validity state will be prevalent for a long time.

Until the community feels comfortable relying on RPKI data, routing on invalid origin validity, though at a low preference, may be prevalent for a long time.

Announcements with valid origins SHOULD be preferred over those with unknown or invalid origins.

Announcements with unvalidatable origins SHOULD be preferred over those with invalid origins.

Announcements with invalid origins MAY be used, but SHOULD be less preferred than those with valid or unknown.

6. Notes

Like the DNS, the global RPKI presents only a loosely consistent view, depending on timing, updating, fetching, etc. Thus, one cache or router may have different data about a particular prefix than another cache or router. There is no 'fix' for this, it is the nature of distributed data with distributed caches.

There is some uncertainty about the origin AS of aggregates and what, if any, ROA can be used. The long range solution to this is the deprecation of AS-SETs, see [I-D.wkumari-deprecate-as-sets].

7. Security Considerations

As the BGP origin is not signed, origin validation is open to malicious spoofing. It is only designed to deal with inadvertent mis-advertisement.

Origin validation does nothing about AS-Path validation and therefore is open to monkey in the middle path attacks.

The data plane may not follow the control plane.

8. IANA Considerations

This document has no IANA Considerations.

9. Acknowledgments

The author wishes to thank Rob Austein, Steve Bellovin, Pradosh Mohapatra, Chris Morrow, Keyur Patel, Heather and Jason Schiller, John Scudder, and Dave Ward.

10. References

10.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[I-D.ietf-sidr-arch] Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", draft-ietf-sidr-arch-11 (work in progress), September 2010.

[I-D.ietf-sidr-repos-struct] Huston, G., Loomans, R., and G. Michaelson, "A Profile for Resource Certificate Repository Structure", draft-ietf-sidr-repos-struct-05 (work in progress), October 2010.

[I-D.ietf-sidr-roa-format] Lepinski, M., Kent, S., and D. Kong, "A Profile for Route Origin Authorizations (ROAs)", draft-ietf-sidr-roa-format-07 (work in progress), July 2010.

[I-D.ymbk-rpki-rtr-protocol] Bush, R. and R. Austein, "The RPKI/Router Protocol", draft-ymbk-rpki-rtr-protocol-06 (work in progress), July 2010.

[I-D.pmohapat-sidr-pfx-validate] Mohapatra, P., Scudder, J., Ward, D., Bush, R., and R. Austein, "BGP Prefix Origin Validation", draft-pmohapat-sidr-pfx-validate-07 (work in progress),

April 2010.

[I-D.wkumari-deprecate-as-sets]
Kumari, W., "Deprecation of BGP AS_SET, AS_CONFED_SET.", draft-wkumari-deprecate-as-sets-01 (work in progress), September 2010.

10.2. Informative References

[RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.

Author's Address

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

Phone: +1 206 780 0431 x1
Email: randy@psg.com