

Secure Inter-Domain Routing
Internet-Draft
Intended status: Informational
Expires: September 27, 2014

K. Sriram
D. Montgomery
US NIST
March 26, 2014

Design Discussion and Comparison of Replay-Attack Protection Mechanisms
for BGPSEC

[draft-sriram-replay-protection-design-discussion-03](#)

Abstract

The BGPSEC protocol requires a method for protection from replay attacks, at least to control the window of exposure. In the context of BGPSEC, a replay attack occurs when an adversary suppresses a prefix withdrawal (implicit or explicit) or replays a previously received BGPSEC announcement for a prefix that has since been withdrawn. This informational document provides design discussion and comparison of multiple alternative replay-attack protection mechanisms weighing their pros and cons. It is meant to be a companion document to the standards track I-D.-ietf-sidr-bgpsec-rollover that will specify a method to be used with BGPSEC for replay-attack protection.

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 September 27, 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.	Definition of Replay Attack	3
3.	Classification of Solutions	4
4.	Expire Time Method	4
5.	Key Rollover Method	5
5.1.	Periodic Key Rollover Method	6
5.2.	Event-driven Key Rollover Method	8
5.2.1.	EKR-A: EKR where Update Expiry is Enforced by CRL . .	9
5.2.2.	EKR-B: EKR where Update Expiry is Enforced by NotValidAfter Time	10
5.2.3.	EKR with Separate Key for Each Incoming-Outgoing Peering-Pair	11
6.	Summary of Pros and Cons	12
7.	Summary and Conclusions	14
8.	Acknowledgements	15
9.	IANA Considerations	15
10.	Security Considerations	15
11.	Informative References	15
	Authors' Addresses	16

[1.](#) Introduction

The BGPSEC protocol [[bgpsec-protocol](#)] requires a method for protection from replay attacks, at least to control the window of exposure [[bgpsec-reqs](#)]. In the context of BGPSEC, a replay attack occurs when an adversary suppresses a prefix withdrawal or replays a previously received BGPSEC announcement for a prefix that has since been withdrawn.

In this informational document, we provide design discussion and comparison of various replay-attack protection mechanisms that may be used in conjunction with the BGPSEC protocol. It is meant to be a companion document to the standards track document [[bgpsec-rollover](#)] that will specify a method to be used with BGPSEC for replay-attack protection. Here we consider four alternative mechanisms - one based on the explicit Expire Time approach and three different variants based on the Key Rollover approach. We provide a detailed comparison between these mechanisms weighing their pros and cons. This document

is meant to help inform the decision process leading to an exact description for the mechanism to be finalized and formally specified in [[bgpsec-rollover](#)].

2. Definition of Replay Attack

In the context of BGPSEC, a replay attack occurs when an adversary suppresses a prefix withdrawal (implicit or explicit). A replay attack occurs also when the adversary replays a previously received BGPSEC announcement for a prefix that has since been withdrawn. In the rest of this document, we will refer to either of these two situations as replay attack. The following are examples of replay attacks:

Example 1: AS1 has AS2 and AS3 as eBGPSEC peers. At time x , AS1 had announced a prefix P to AS2 and AS3. At a later time $x+d$, AS1 sends a Withdraw for prefix P to AS2. AS2 suppresses the Withdraw (does not send to its peers any explicit or implicit Withdraw). AS2 continues to attract some of the data for prefix P towards itself by pretending to still have a signed and valid route for P . In effect, AS2 can conduct a DOS attack on a server located at AS1 at prefix P . (See slide #15 in [[replay-discussion](#)] for an illustration.)

Example 2: AS1 has AS2 and AS3 as eBGPSEC peers. AS2 and AS3 are also eBGPSEC peers. At time x , AS1 had announced a prefix P to AS2 and AS3. AS3 also propagates to AS2 its route (via AS1) for prefix P . At a later time $x+d$, AS1 discontinues its peering with AS2. AS2 should propagate an alternate longer path via AS3 for prefix P and thus send an implicit Withdraw. However, AS2 suppresses it. AS2 can thus make a significant part of traffic destined for prefix P to flow via itself and eavesdrop on the data but not cause a DOS attack. (See slide #16 in [[replay-discussion](#)] for an illustration.)

Example 3: AS1 has AS2 and AS3 as eBGPSEC peers. AS2 and AS3 are also eBGPSEC peers. At time x , AS1 had announced a prefix P to AS2 without prepending (Update: AS1{pCount=1} P) but announced the same prefix to AS3 with prepending (Update: AS1{pCount=2} P). Thus AS1 had preferred its ingress data traffic for prefix P to come in via AS2. At a later time $x+d$, AS1 switches ingress data path preference to AS3 over AS2 - announces prefix P without prepending (Update: AS1{pCount=1} P) to AS3 and with prepending (Update: AS1{pCount=2} P) to AS2. AS2 suppresses the new prepended path announcement (does not send to its peers any new update about P). Thus AS2 carries more of AS1's ingress data traffic and generates more revenue for itself at the expense of AS1. (See slide #17 in [[replay-discussion](#)] for an illustration.)

Thus the scenarios and motivations for replay attacks may differ as illustrated by the examples above.

A requirement for replay-attack protection can be stated as follows. The update that AS1 sent to AS2 at time x should expire at time $x+w$. That means, AS2 can suppress the Withdraw or possibly replay the update from AS1 for prefix P until at most $x+w$. This limits the replay vulnerability window. (Note: If no peering or policy change affecting prefix P occurs during the vulnerability window, then a typical solution would include a method for extending the validity period of the route(s) beyond $x+w$.)

3. Classification of Solutions

Mechanisms for replay-attack protection can be classified into two broad categories as follows:

- o Expire Time (ET) Method: This method uses an explicit expire time field in the BGPSEC update.
- o Key Rollover (KR) Method: In this method, the update expiry is enforced by a key rollover. Router rolls over to a new signing cert with a new pair of keys, and the previous router cert either expires or is revoked.

The Key Rollover method can be further characterized into the following sub categories:

- o Periodic Key Rollover (PKR): Key rollovers happen at periodic intervals.
- o Event-driven Key Rollover (EKR): Key rollovers happen only when peering or policy change events occur.
 - * EKR-A: EKR where expiry of previous update is enforced by CRL.
 - * EKR-B: EKR where expiry of previous update is controlled by NotValidAfter time.

In [Section 4](#), [Section 5](#), and [Section 6](#) we describe the various methods listed above, and discuss their pros and cons.

4. Expire Time Method

The details of the Expire Time (ET) method are as follow:

- o Explicit Expire Time is used for origin's signature.

- o Expire Time field is required in the BGPSEC update.
- o Periodic re-origination (beaconing) of prefixes is performed by origin ASes. The value in the ET field in the update is extended at beaconing time, and thereby the update is refreshed. Every prefix in the Internet is re-originated and propagates through the Internet once every 'beacon' interval.
- o These beacons are distributed actions by prefix owners and jittered in time by design to reduce burstiness. The beacon interval can be different at different originating ASes.
- o Beacon interval granularity: TBD but preferably in fairly granular units (days).

Discussion of Pros and Cons:

Pro: This method is easy on transit routers. In the event of peering or policy change, BGPSEC with the ET method behaves the same way as BGP-4 in terms of which prefix routes are propagated. That is, the router re-evaluates best paths factoring in peering or policy changes, and propagates only those prefix routes that have a change in best path. In other words, there is no necessity for the BGPSEC router to re-propagate and refresh prefixes on all peering links. This is because prefix updates are refreshed anyway once every beacon interval by all prefix originators. There is low steady-state traffic associated with beaconing (see Figure on slide #8 in [\[replay-discussion\]](#)), but there are no huge bursts or spikes in workload due to peering or policy change events at transit routers.

Con: Equipment vendor can potentially facilitate unnecessary frequent beaconing if ISP urges and pays (dollar attack!). This possibility is mitigated by having a well thought-out granularity for ET, for example, if the unit of ET is one day (rather than one minute).

Con: A change in on-the-wire BGPSEC protocol would be needed in case the unit of the ET field (granularity) needs to be changed.

5. Key Rollover Method

Key Rollover (KR) method has three variations as outlined in [Section 3](#). Those will be discussed later in this section. The following features are common to all variants of the KR method:

- o In the KR method, it is best if the BGPSEC router has two pairs of certs as follows: A pair of origination certs (current and next) for signing prefixes being originated by the AS of the router, and

a pair of transit certs (current and next) for signing transit prefixes.

- o Note: If a BGPSEC router only originates prefixes (i.e., has no transit prefixes), then it needs to maintain only a pair of origination certs and need not maintain the extra pair of transit certs.
- o The three KR methods differ in how the rollover of certs (or keys) is done:
 - * Cert rollovers are Periodic vs. Event-driven.
 - * In the Event-driven method, the expiry of old update is (A) Enforced by CRL vs. (B) Controlled by NotValidAfter time.
 - * In (A), cert's NotValidAfter field is set to a very large value and CRL is issued to revoke the cert when necessary. In (B), NotValidAfter field set to a permissible vulnerability window time and CRL to revoke cert is not required.

Discussion of Pros and Cons (common to all Key Rollover methods):

Pro: The KR method functions by manipulating the RPKI objects (certs, keys, NotValidAfter field in cert, etc.) to refresh updates or to cause expiry of previously propagated updates. Unlike the ET method, it does not rely on any explicit field in the update. Hence, an advantage of the KR method over the ET method is that in case any parameters need to change or if the method itself is modified, then there is no impact on the BGPSEC protocol on the wire.

Con: The KR method introduces additional churn in the global RPKI system.

Con: There is also added update churn. The amount of update churn varies depending on the type of KR method used (see [Section 5.1](#) and [Section 5.2](#)).

We will now describe and discuss in detail the variants of the KR method.

[5.1](#). Periodic Key Rollover Method

The details of the Periodic Key Rollover (PKR) method are as follow.

- o Router's origination cert's NotValidAfter time is used as the implicit expire time for origin's signature.

- o Each origination router re-originate (i.e., beacons) before NotValidAfter time of the current cert. Beacons are periodic re-origination of prefixes by origin ASes.
- o At beaconing time, next cert becomes the new current cert, and update is signed with the private key of this new current cert and re-originate.
- o A new 'next' cert is created and propagated at beaconing time. This can also be done with a good lead time. In practice, multiple 'next' certs can be kept in the pipeline. They must have contiguous or slightly overlapping validity periods.
- o Every prefix in the Internet is re-originate and propagates through the Internet once every 'beacon' interval.
- o The re-originations or beacons are distributed actions by prefix owners and jittered in time by design to reduce burstiness. The beacon interval can be different at different originating ASes.
- o Beacon (or re-origination) interval granularity: TBD but preferably in fairly granular units (days).
- o Transit certs can have very large NotValidAfter time (say ~years).
- o When a peering or policy change event occurs at a transit router, the router (i.e. BGPSEC router with PKR) does not perform any key rollover. The router re-evaluates best paths factoring in peering or policy changes, and propagates only those prefix routes that have a change in best path (similar to BGP-4). There is no necessity for the BGPSEC router to re-propagate and refresh prefixes on all peering links. This is because prefix updates are refreshed anyway once every re-origination (i.e. beaconing) interval by all prefix originators.

Discussion of Pros and Cons:

Several of the same pros/cons of the Expire Time method also apply here for the PKR method.

Pro: The main pro for the PKR method is the same as that for the Expire Time (ET) method. That is, being easy on transit routers as discussed in [Section 4](#). Just as in the ET method, there is low steady-state traffic associated with periodic re-originations (i.e. beaconing) (see Figure on slide #8 in [[replay-discussion](#)]), but there are no huge bursts or spikes in workload due to peering or policy change events at transit routers. (See comparisons with the EKR methods in [Section 5.2](#).)

Pro: The pro discussed above for the KR method regarding parameter changes (e.g., beacon interval units) not requiring change of protocol on the wire is naturally applicable here.

Con: Churn in the RPKI is of concern. Every BGPSEC router rolls two origination certs (current and next) once in every beacon (i.e., re-origination) interval.

5.2. Event-driven Key Rollover Method

The common details of the Event-driven Key Rollover (EKR) methods are as follow.

- o Key rollover is reactive to events (not periodic).
- o If a peering or policy change event involves only prefixes being originated at the AS of the router, then the router rolls only the origination key.
- o If a peering change event involves transit prefixes at the AS of the router, then the router rolls the transit key as well as the origination key.
- o If a key rollover takes place, then a corresponding (origination or transit) new 'next' cert is propagated in RPKI.

Discussion of Pros and Cons:

Pro: As long as no triggering events occur, there is no added update churn in BGPSEC.

Con: Whenever the transit key is rolled, there is a storm of BGPSEC updates at routers in transit ASes. For example, consider BGPSEC capable transit AS5 that is connected to four BGPSEC non-stub customers (AS1, AS2, AS3, AS4). Assume each AS has a single BGPSEC router in it. AS1 through AS4 each receives almost full table (400K signed prefix updates) from AS5. Assume also that AS1 and its customers together originate 100 prefixes in total; likewise for AS2, AS3 and AS4. Now consider that an event occurs whereby the peering between AS1 and AS5 is discontinued. As a result of this event, in the EKR method, the AS5 router signs and re-propagates approximately $3 \times 400K = 1.2$ Million signed prefix updates to AS2, AS3 and AS4 combined. In addition, it also sends $4 \times 100 = 400$ Withdraws, which are negligible. In comparison, in the PKR method, following the same event, the router at AS5 sends only $4 \times 100 = 400$ Withdraws and signs/re-propagates ZERO prefix updates. (An illustration can be found in slide #9 in [[replay-discussion](#)]). Also, additional peering change

scenarios and quantitative comparisons can be found in slides #10 and #11 in [[replay-discussion](#)].)

It remains to be seen through measurement and modeling how the impact of such large bursts of workload in the ETR method at the time of event occurrence can be managed in route processors, e.g., by jittering and throttling the workload.

[5.2.1](#). EKR-A: EKR where Update Expiry is Enforced by CRL

EKR-A builds on the common principles as described for EKR above in [Section 5.2](#). The additional details of EKR-A operation are as follow:

- o NotValidAfter time of origination and transit certs is set to a large value (~year).
- o Whenever key rollover (for origination or transit) occurs, then CRL is propagated for the old cert. So the old update expires (due to invalid state) only when the CRL propagates and reaches the relying router.
- o This method relies on end-to-end CRL propagation through the RPKI system to enforce expiry of a previous update whenever the need arises.
- o The cert CRL either propagates all the way to the relying router, or the RPKI cache server of the router receives the CRL and then sends a withdrawal of the {AS, SKI, Pub Key} tuple to the router. Either way, the CRL must in effect propagate all the way to the relying router.
- o Thus the attack vulnerability window with the EKR-A method is governed by the end-to-end CRL propagation time.

Discussion of Pros and Cons:

The following pro and con for the EKR-A method are in addition to the common pros and cons listed above for the KR and EKR methods ([Section 5](#) and [Section 5.2](#)).

Pro: EKR-A has much less RPKI churn than PKR or EKR-B (see [Section 5.2.2](#)).

Con: Router needs to receive a CRL or a withdraw of {AS, SKI, Pub Key} tuple in order to know an update has expired. Hence, the replay-attack vulnerability window is determined by the CRL propagation time which can vary widely from one relying router to

another router that may be in different regions. It is anticipated that this would be no worse than 24 hours, but needs to be confirmed by measurements in an operational or emulated RPKI systems [[rpki-delay](#)].

5.2.2. EKR-B: EKR where Update Expiry is Enforced by NotValidAfter Time

EKR-B builds on the common principles as described for EKR above in [Section 5.2](#). The additional details of EKR-B operation are as follow:

- o NotValidAfter time of current origination and transit certs is set to a value determined by the desired vulnerability window (~day).
- o Update expiry is controlled by NotValidAfter time and CRL is not sent for the old cert when key rollover happens.
- o If no triggering event occurs to cause origination key rollover within a pre-set time (NotValidAfter), then new origination (current and next) certs are issued only to extend the NotValidAfter time but the corresponding key pairs and SKIs remain unchanged.
- o A previous update automatically becomes invalid at the earliest NotValidAfter time of the certs used in the signatures unless each of those certs' NotValidAfter time has been extended.
- o Likewise for the transit (current and next) certs and keys.
- o Changes in certs to extend their NotValidAfter time need not propagate end-to-end (all the way to the relying routers); they may propagate only up to the RPKI cache server of the relying router. RPKI cache server would send a withdraw for an {AS, SKI, Pub Key} tuple to a relying router if the NotValidAfter time of the cert has passed.
- o The changes in certs to advance NotValidAfter time can be scheduled and propagated in RPKI well in advance.

Discussion of Pros and Cons:

The following pro and con for EKR-B are in addition to the common pros and cons listed above for the KR and EKR methods ([Section 5](#) and [Section 5.2](#)).

Pro: Update expiry is automatic in case the NotValidAfter time of any of the certs used to sign the update has not been extended. So the

replay-attack vulnerability window is predictable and not influenced by the RPKI end-to-end propagation time.

Pro: Routers do not get any RPKI updates from the RPKI cache server when cert changes but the key pair and SKI remain unchanged. Routers do not receive NotValidAfter time from their RPKI cache server. There is no need for it. Instead, the RPKI cache server keeps track of NotValidAfter time, and provides to routers only valid {AS, SKI, Pub Key} tuples. This saves some RPKI state maintenance workload at the routers.

Con: EKR-B has much more RPKI churn than EKR-A because both origination and transit certs need to be reissued periodically to extend their validity time (in the absence of any events).

5.2.3. EKR with Separate Key for Each Incoming-Outgoing Peering-Pair

This is a place holder section where we mention another variant of the EKR method. This idea has not been considered or whetted by the SIDR WG yet. So we only mention it here briefly.

As noted earlier, the EKR methods considered so far generate a huge spike in workload whenever the transit key rollover takes place at a router. One way to reduce that workload is to have a separate signing key for each incoming-outgoing peering pair. For example, consider a BGPSEC router in AS4 that has peers in AS1, AS2, and AS3. The router will hold six signing keys, one each corresponding to (AS1, AS2), (AS2, AS1), (AS1, AS3), (AS3, AS1), (AS2, AS3), and (AS3, AS2) peering-pairs. Note that the directionality of peering is included here and is necessary. The key corresponding to (AS-i, AS-j) would only be used to sign updates received from AS-i and being forwarded to AS-j. In the general case, when the BGPSEC router has n peers, the number of transit keys will be $n(n-1)$. Since there would be a Current and a Next key (for rollover), the number of transit keys held in the router for signing will be actually $2n(n-1)$. When a peering or policy change occurs, the router would rollover only those specific keys that correspond to the peering-pairs over which the prefix updates are affected. In the above example, suppose a policy change between AS4 and AS1 causes AS4 to prepend prefixes sent to AS1 (pCount changed from 1 to 2). Then AS4 would do key rollover only for (AS2, AS1) and (AS3, AS1) peering-pairs, and not for any of the others. This would substantially reduce the quantity of prefix updates that are signed and re-propagated. In general, when peering or policy changes occur, this method will reduce the number of prefix updates to be re-propagated to exactly the same as that with normal BGP. That means that this method would also be on par with the ET and PKR methods in terms of update churn when a peering or policy

change takes place. The downside of this method is that the router needs to maintain $2n(n-1)$ key pairs if it has n BGPSEC peers.

Detailed discussion and comparison of this method with other methods can be provided in a later version of this document if the idea picks up interest in the WG.

6. Summary of Pros and Cons

Table 1 below summarizes the pros and cons for the various replay-attack protection methods. This summary follows from the discussion above in [Section 4](#) and [Section 5](#).

Method	Pros	Cons
Expire Time (ET)	1. The background load due to beaconing is low and not bursty. --- 2. Transit AS does NOT have a huge spike in workload even when a peering or policy change happens at that AS. Beaconing facilitates this. --- 3. Does not add to RPKI churn.	1. Prefix owner can abuse by beaconing too frequently. --- 2. Any change to the units (granularity) of ET field entails a change to on-the-wire BGPSEC protocol. ---
Periodic Key Rollover (PKR)	1. The background load due to beaconing is low and not bursty. --- 2. Transit AS does NOT have a huge spike in workload even when a peering change happens at that AS. Beaconing (i.e. periodic re-origination) facilitates this. --- 3. If the periodic re-origination (i.e.,	1. Prefix owner can abuse by beaconing (i.e. re-originating) too frequently. --- 2. Adds to RPKI churn. A pair of certs (current and next) for each origination router are rolled once every beacon (i.e. re-origination) interval. Significantly more RPKI churn than that with EKR-A or EKR-B methods. ---

	beaconing) interval units change, BGPSEC protocol on the wire remains unaffected.	
	---	---
	4. Changes in the method (while still based on Key Rollover) can be accommodated without requiring any change to on-the-wire BGPSEC protocol.	
-----	-----	-----
Event driven Key Rollover Type A (EKR-A)	1. No update churn for long periods when no peering or policy changes occur.	1. Whenever the transit key is rolled (in response to a peering or policy change event), there is a storm of BGPSEC updates, especially at routers in large transit ASes.
	---	---
	2. The added churn in RPKI is much lower than that in the EKR-B method.	2. The replay-attack vulnerability window is dependent on end-to-end CRL propagation. It may vary significantly from one relying router to another that may be in different regions.
	---	---
	3. Same as Pro #4 for the PKR method.	
-----	-----	-----
Event driven Key Rollover Type B (EKR-B)	1. Same as Pro #1 for the EKR-A method.	1. Same as Con #1 for the EKR-A method.
	---	---
	2. The replay-attack vulnerability window is enforced by NotValidAfter time in certs and is therefore predictable.	2. The added churn in RPKI is much higher than that in the EKR-A method.
	---	---
	3. Same as Pro #4 for the PKR method.	

1. The Expire Time (ET) method is best (on par with the PKR method) in terms of preventing huge update workloads during peering and policy change events at transit routers with several peers. It has no added RPKI churn. But the ET method has the disadvantage of requiring on-the-wire protocol change if some parameters (e.g., the units of beacon interval) change.
2. The Periodic Key Rollover (PKR) method operates the same way as the ET method for preventing huge update workloads during peering and policy change events at transit routers with several peers. It does not have the disadvantage of requiring on-the-wire protocol change if some parameters (e.g., the units of beaconing/re-origination periodicity) change. But it has the downside of added RPKI churn.
3. The Event-driven Key Roll (EKR-A and EKR-B) methods have significantly less RPKI churn than the PKR method. They also have no BGPSEC update churn during long quiet periods when no peering or policy change events occur. But they suffer the drawback of creating huge update workloads during peering and policy change events at transit routers with several peers. Can this workload be jittered or flow controlled to spread it over time without convergence delay concerns? May be - needs further study.
4. The EKR-A method relies on end-to-end CRL propagation through the RPKI system to enforce expiry of a previous update when needed. By contrast, in the EKR-B method the update expiry is controlled by NotValidAfter time of the certs used in update signatures. In EKR-B, previous update automatically becomes invalid at the earliest NotValidAfter time of the certs used in the signatures unless each of those certs' NotValidAfter time has been extended.

In the latter method, changes in certs to extend their NotValidAfter time need not propagate end-to-end (all the way to the relying routers); they may propagate only up to the RPKI cache server of the relying router (see [Section 5.2.2](#)). The changes in certs to advance NotValidAfter time can be scheduled and propagated in RPKI well in advance.

5. Besides being out-of-band relative to the BGPSEC protocol on the wire, the other good thing about the Key Rollover method is that once the basics of the mechanism are implemented, there may be flexibility to implement PKR, EKR-A or EKR-B on top of it. It may also be possible to switch from one method to another (within this class) if necessary based on operational experience; this transition would not require any change to on-the-wire BGPSEC protocol.

8. Acknowledgements

The authors would like to thank Roque Gagliano, Brian Weis and Steve Kent for helpful discussions. Further, we are thankful to fellow NIST BGP team members for comments and suggestions.

9. IANA Considerations

This memo includes no request to IANA.

10. Security Considerations

This memo requires no security considerations of its own since it is targeted to be an informational RFC in support of [[bgpsec-rollover](#)] and [[bgpsec-protocol](#)]. The reader is therefore directed to the security considerations provided in those documents.

11. Informative References

[[bgpsec-protocol](#)]

Lepinski (Ed.), M., "BGPSEC Protocol Specification", Work in Progress, February 2013, <<http://datatracker.ietf.org/doc/draft-ietf-sidr-bgpsec-protocol/>>.

[[bgpsec-reqs](#)]

Belloven, S., Bush, R., and D. Ward, "Security Requirements for BGP Path Validation", Work in Progress, April 2013, <<http://datatracker.ietf.org/doc/draft-ietf-sidr-bgpsec-reqs/>>.

[bgpsec-rollover]

Gagliano, R., Patel, K., and B. Weis, "BGPSEC router key rollover as an alternative to beaconing", Work in Progress, April 2013, <<http://datatracker.ietf.org/doc/draft-ietf-sidr-bgpsec-rollover/>>.

[replay-discussion]

Sriram, K. and D. Montgomery, "Discussion of Key Rollover Mechanisms for Replay-Attack Protection", Presented at IETF-85 SIDR WG Meeting, November 2012, <<http://www.ietf.org/proceedings/85/slides/slides-85-sidr-4.pdf>>.

[rpki-delay]

Kent, S. and K. Sriram, "RPKI rsync Download Delay Modeling", Presented at IETF-86 SIDR WG Meeting, March 2013, <<http://www.ietf.org/proceedings/86/slides/slides-86-sidr-1.pdf>>.

Authors' Addresses

Kotikalapudi Sriram
US NIST

Email: ksriram@nist.gov

Doug Montgomery
US NIST

Email: doug@nist.gov

