

Routing Protocol Security
Requirements
Internet-Draft
Intended status: Informational
Expires: August 17, 2007

B. Christian, Ed.
KMC Telecom Solutions
T. Tauber, Ed.
Comcast
February 13, 2007

BGP Security Requirements
draft-ietf-rpsec-bgpsecrec-07

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on August 17, 2007.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

The security of BGP, the Border Gateway Protocol, is critical to the proper operation of large-scale internetworks, both public and private. While securing the information transmitted between two BGP speakers is a relatively easy technical matter, securing BGP, as a routing system, is more complex. This document describes a set of requirements for securing BGP, including communications between BGP speakers, and the routing information carried within BGP.

Table of Contents

1.	Introduction	3
1.1.	System Description	3
1.2.	Threats	3
1.3.	Areas to secure	4
2.	Underlying Assumptions regarding BGP	5
3.	Operational Requirements	7
3.1.	Convergence speed	8
3.2.	Incremental deployment	8
3.3.	Conditions for initialization	9
3.4.	Local controls for secure UPDATE acceptance	9
3.5.	Processing on Routers	10
3.6.	Configuration on Routers	11
4.	Infrastructure Requirements	11
5.	The Trust Model	11
6.	The AS_PATH Attribute and NLRI Authentication	12
7.	Address Allocation and Advertisement	14
8.	Logging	14
9.	NLRI and Path Attribute Tracking	15
10.	Transport Layer Protection	15
11.	Key Management	15
12.	IANA Considerations	16
13.	Security Considerations	16
14.	References	16
14.1.	Normative References	16
14.2.	Informative References	16
Appendix 1.	Acknowledgements	17
Authors' Addresses	17
Intellectual Property and Copyright Statements	18

1. Introduction

1.1. System Description

BGP is described in [RFC1771](#) [4], and, more recently, in an updated specification, [RFC4271](#) [1], as a path-vector routing protocol. BGP speakers typically exchange information about reachable destinations (expressed as address prefixes) in an internetwork through pair-wise peering sessions. Once this information has been exchanged, each BGP speaker locally determines a loop free path to each reachable destination, based on local policy or policy indicators such as community values and LOCAL_PREF which may be carried in the UPDATE, and the AS_PATH data carried in the BGP UPDATE messages.

Each BGP speaker represents an Autonomous System (AS). All of the BGP speakers within an AS operate under a common administrative policy.

1.2. Threats

Violations of security for network and information systems generally fall under one of the three categories as defined in [RFC 2196](#) [2]:

- o Unauthorized access to resources and/or information
- o Unintended and/or unauthorized disclosure of information
- o Denial of service

A number of attacks can be realized which, if exploited, can lead to one of the above mentioned security violations. Attacks against communications are typically classified as passive or active wiretapping attacks. Passive attacks are ones where an attacker simply observes information traversing the network, violating confidentiality or identifying a means of engaging in further attacks. Active attacks are ones where the attacker modifies data in transit. Such attacks include replay attacks, message insertion, message deletion, and message modification attacks. Some attacks may be effected by sending data from any where in the Internet. Other active attacks require a "man-in-the-middle" capability, i.e., the attacker must be in a position where traffic passes through an attacker-controlled device. Attacks against BGP may be used by an attacker to facilitate a wide variety of active or passive wiretapping attacks against subscriber traffic.

Attacks that do not involve direct manipulation of BGP, and the information contained within BGP, are outside the scope of this document.

Because ASes are autonomous in their operation, it is not possible to mandate secure operation by all ASes, nor would it be advisable to assume such operation. Thus the primary goal of BGP security measures is to provide data to AS operators to enable BGP speakers to reject advertisements (UPDATE messages) that are not valid. For example, UPDATE messages that represent erroneous binding of prefixes to an origin AS, or that advertise invalid paths (as defined later in this document) should be rejected. Because BGP peering sessions take place in the context of TCP, the authentication and integrity guarantees usually associated with TCP need to be provided in the face of possible active wiretapping attacks. Using the terminology established in [RFC 3552](#) [3], these peering sessions should be afforded data origin and peer entity authentication and connection-oriented integrity.

Security for subscriber traffic is outside the scope of this document, and of BGP security in general. IETF standards for subscriber data security, e.g., IPsec, TLS, and S/MIME should be employed for such purposes. While adoption of BGP security measures may preclude certain classes of attacks on subscriber traffic, these measures are not a substitute for use of subscriber-based security mechanisms of the sort noted above.

1.3. Areas to secure

There are two primary points where BGP may be secured; the data payload of the protocol and the data semantics of the protocol.

The session between two BGP speakers can be secured such that the BGP data received by the BGP speakers can be cryptographically verified to have been transmitted by the peer BGP speaker and not a replay of previously transmitted legitimate data. There are several existing IETF standards to choose from to ensure that this system functions with greater effectiveness than the current system. An example might be IPsec. Some in the Operator community have expressed concerns that a requiring cryptographic validation could open another vector for a denial-of-service attack by flooding the processor with bogus packets which must be cryptographically invalidated before being discarded.

There are also several questions we can ask about the information contained within a received UPDATE.

- o Is the originating Autonomous System authorized to propagate the prefix we have received?
- o Does the AS_PATH, received via an UPDATE, represent a valid path through the network?

The determination of AS_PATH validity falls into two distinct categories. These categories are ordered from least to most rigorous.

- o Does the AS_PATH specified actually exist as a path in the network topology and, based on the AS_PATH, is it possible to traverse that path to reach a given prefix? This AS_PATH Feasibility Check will be referred to later in this document.
- o Has the UPDATE actually traveled the path?

2. Underlying Assumptions regarding BGP

In order to properly identify security requirements it is important to articulate the fundamental aspects of BGP as related to security requirements. The following list presents the basic parameters and application concepts of BGP that are assumed by this document.

- o Peer Communication: BGP traffic travels over TCP between peers, so BGP speakers assume the data delivery guarantees of TCP in a benign environment. This includes ordered, error-free delivery of application traffic from a peer identified by an IP address, plus integrity of the control aspects of TCP. From a security perspective, these guarantees need to be enforced in the context of possible active wiretapping.
- o Routing and Reachability: BGP is a protocol used to convey routing and reachability information both internal and external to an Autonomous System. Typically, interior BGP (iBGP) is used to distribute prefix reachability information in conjunction with an Interior Gateway Protocol (IGP) and is used by a distinct network administrative entity to convey internal routing policy regarding external and internal information. Exterior BGP (eBGP) is typically used to distribute route/prefix reachability information between two distinct routing entities and is used to signal eBGP preferences and policy decisions.
- o Inter-AS UPDATE Message assumptions: When an AS distributes reachability information to a peer it is done with the intent of affecting routing decisions by the peer. For example, with regard to a block of addresses represented by a prefix, AS-A may send peer AS-B an advertisement which is less specific (shorter in length of mask) and peer AS-C a more specific advertisement (longer mask). This prefix distribution decision may have been made to provide a means for failure resolution between AS-A and AS-C, i.e., to provide a backup path for the addresses in question. However, it should be noted that while AS-A tries to

influence the routing decisions of AS-B and downstream ASes, AS-A is only providing inputs to a local decision by AS-B, a decision that is ultimately controlled by AS-B's local policy over which AS-A has no control. Update messages are sent between AS peers with the tacit authorization for those messages to be forwarded to others. A notable exception to this assumption is the use of policy-based mechanisms between peers such as the NO-EXPORT community. It is important to note that an UPDATE message itself generally is not re-transmitted. Instead, the specific UPDATE message is regenerated continually as it passes from BGP speaker to BGP speaker. Furthermore, UPDATE messages have no mechanism to indicate freshness (e.g., timestamps or sequence numbers). This implies that messages may appear valid at any point in the life of a BGP peering session. While the AS_PATH information is typically transitive it is, currently, not clearly mandated and many times is modified for various utilitarian reasons.

- o It is important to note that while preferences regarding routing can be explicitly managed with direct peers it is markedly more difficult to influence routing decisions by ASes that are not directly adjacent.
- o Inter-AS withdrawal message assumptions: The processing model of BGP [RFC4271](#) [1] indicates that only the peer advertising NLRI information may withdraw it. There are several instances where a withdrawal may occur. Typical reasons for withdrawal include the determination of a better path, peer session failure, or local policy change. There is no specified mechanism for indicating to a peer the reason for a route withdrawal. Each withdrawal received via a valid peering session must be taken at face value. There is no existing method to ensure that an AS will properly respond to a withdrawal message, e.g., withdraw the route and send such announcement to its neighbors. Nor do mechanisms exist to ensure that old UPDATES are not re-propagated after a route was withdrawn before it is legitimately re-advertised.
- o AS_PATH assumptions: Aside from the use of AS_SET, the AS_PATH is defined as an ordered list of the Autonomous Systems that an UPDATE has traversed. The rightmost AS in the list is understood to be the originator of the BGP announcement. Specifications state that the AS routing graph MUST be loop free. This indicates that UPDATES received from an external peer which contain the local AS will be rejected. Prepending one or more instances of an AS number on inbound advertisements (where the external peer's AS number is prepended) and outbound advertisements (where the local AS number is prepended) is a commonly used method to bias routing. Prepending a peer AS number on inbound UPDATES is employed for biasing internal routing and forwarding management while

prepending one's own AS number on outbound advertisement is typically used to bias forwarding and routing changes in external networks. The latter practice is explicitly permitted by [RFC4271 \[1\]](#), but the former is not. Some operators, insert a remote AS number in an UPDATE, in order to cause the UPDATE to be dropped by that AS so that traffic will not traverse a given path. Though this practice appears to run counter to the design of BGP, anecdotal evidence is that its use is not totally insignificant. While such a practice can be beneficial to legitimate operators, it presents a strong potential for misuse. A proposed security system SHOULD address how to either address this concern or give specific information on this topic for consideration by the Operational community.

- o Route Origination: BGP speakers may originate routes based either configured internal data or via data received from peers via UPDATES. An Autonomous System SHOULD only originate a prefix to its external peers if that prefix has been allocated to the administrators of that system, or if authorized by the prefix holder.
- o Originating a route without the ability to forward the traffic associated with that route is, in most cases, in conflict with the intent of the BGP specification, notable exceptions include:
 - * Deployments that make use of route servers which are separate from forwarding devices
 - * Deployments that use the temporary propagation of prefixes in order to effectively block high bandwidth attacks (e.g., DDoS) against specific IP addresses (and the associated oversubscription of resources)
- o Aggregation and de-aggregation: According to [RFC4271 \[1\]](#), if a BGP speaker chooses to aggregate a set of more specific prefixes into a less specific prefix then the ATOMIC_AGGREGATE attribute SHOULD be set. This creates a significant challenge for solutions to secure BGP because some origination information is removed (i.e. the more-specific information which triggered the generation of the aggregate). Proposed solutions MUST indicate how aggregation will be accommodated.

3. Operational Requirements

We have determined, through discussion with several large internetwork operators and equipment vendors, that the following attributes are important to the ongoing performance of interdomain

routing systems such as BGP.

3.1. Convergence speed

Convergence speed is a major concern to many operators of large scale internetworking systems. Networks, and internetworks, are carrying ever increasing amounts of information that is time and delay sensitive; increasing convergence times can adversely affect the usability of the network, and the ability of an internetwork to grow. BGP's convergence speed, with a security system in operation, SHOULD strive towards equivalence to BGP running without the security system in operation. This includes the preservation of optimizations currently used to produce acceptable convergence speeds on current hardware, including UPDATE packing, peer groups, etc. Two types of verification MAY be offered for the NLRI and the AS_PATH in order to allow for a selection of optimizations:

- o Contents of the UPDATE message SHOULD be authenticated in real-time as the UPDATE message is processed.
- o The route information base MAY be authenticated periodically or in an event-driven manner by scanning the route-table data and verifying the originating AS and the validity of the AS_PATH list.

All BGP implementations that implement security MUST utilize at least one of the above methods for validating routing information. Real time verification is preferred in order to prevent transitive failures based on periodic or event-driven scan intervals. See the section on "Local controls ..." below for more discussion.

It is recognized that achieving all of these goals might prove very difficult or even impossible.

3.2. Incremental deployment

It will not be feasible to deploy a newly secured BGP protocol throughout the public Internet instantaneously. It also may not be possible to deploy such a protocol to all routers in a large AS at one time. Any proposed solution MUST support an incremental deployment which will provide some benefit for those who participate. Because of this, there are several requirements that any proposed mechanism to secure BGP must consider.

- o A BGP security mechanism MUST enable each BGP speaker to configure use of the security mechanism on a per-peer basis.
- o A BGP security mechanism 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 fashion though care must be taken to ensure UPDATES can traverse intermediate routers which don't support the new format.

- o In an environment where both secured and non-secured systems are interoperating a mechanism MUST exist for secured systems to identify whether an originator intended the information to be secured.
- o Proposed solutions MUST provide comment and analysis of what the security services the solution will provide in the case of incremental deployment scenarios (e.g, contiguous islands, discontinuous islands, universal deployment).
- o In an environment where secured service is in the process of being deployed a mechanism MUST exist to support a transition free of service interruption.

3.3. Conditions for initialization

A key factor in the robust nature of the existing internal and external relationships maintained in today's Internet is the ability to maintain and return to a significantly converged state without the need to rely on systems external to the routing system (the equipment that is performing the forwarding). In order to ensure the rapid initialization and/or return to service of failed nodes it is important to reduce reliance on these external systems to the greatest extent possible. Therefore, proposed systems SHOULD NOT require connections to external systems, beyond those directly involved in peering relationships, in order to return to full service. Proposed systems MAY require post initialization synchronization with external systems in order to synchronize security information.

3.4. Local controls for secure UPDATE acceptance

Each secured environment (e.g., public Internet vs. private internetwork) may have different metrics of what is acceptable or unacceptable with regard to routing security. In environments that require strict security it may not be acceptable to temporarily route to a destination while waiting for path validation to be performed. However, in many environments the rapidity of route installation may be of paramount importance, e.g., in order to facilitate the common occurrence of route withdrawal due to network failure. Based on the two divergent requirements, the following criteria apply:

- o The security system **MUST** support a range of possible outputs for local determination of the trust level for a specific route so that routing preference and policy can be applied to its inclusion in the RIB. Any given route should be trustable to a locally configured degree, based on the completeness of security information with a received UPDATE and other factors. However, experience in the security community suggests that trying to assign trust ratings to inputs to a decision process usually adds considerable complexity to the management of the process. This complexity, in turn, may undermine the security offered by the process.
- o The security system **SHOULD** allow the operator to determine whether speed of convergence is more important than security, or whether security is more important than the speed of convergence. This facilitates the incremental deployment of security on systems not designed to support increased processing requirements imposed by the security system.

3.5. Processing on Routers

The introduction of mechanisms to improve routing security will generally increase the processing performed by a router. The increased processing typically will result from additional checks performed to determine the validity of UPDATES, especially if these checks entail cryptographic operations. Since currently deployed routers generally do not have hardware to accelerate cryptographic operations, these operations could impose a significant processing burden under some circumstances. Thus proposed solutions should be evaluated carefully with regard to the processing burden they may impose, since deployment may be impeded if network operators perceive that a solution will impose a processing burden which either:

- o provokes substantial capital expense, or
- o threatens to destabilize routers.

Given the pervasive number of BGP-speaking routers in a typical ISP deployment, solutions can increase their appeal by minimizing the burden imposed on all BGP routers in favor of confining significant work loads to a relatively small number of devices.

Optional features or increased assurance which provokes more pervasive processing load **MAY** be made available for deployments where the additional resources are economically justifiable.

Some statement as to the expected performance measures and scaling as a function of prefixes, peers, NLRI, etc. **MUST** be included with any

proposed solution.

3.6. Configuration on Routers

It is undesirable to have long or very detailed configuration on the routers, especially it needs to be synchronized on all of them. Long configuration makes operating the device more difficult, and having to do very detailed configuration may hinder the adoption of the security solution; it should be possible to "just start using it" if possible.

As above, a statement as to the expected configuration burden as a function of routers, peers, NLRI, ASNs, etc. MUST be included with any proposed solution. Additionally, some consideration SHOULD be given and statement made as to frequency of changes in the case of dynamic data.

4. Infrastructure Requirements

BGP security mechanisms MAY make use of a security infrastructure to distribute authenticated data that is an input to routing decisions. Such data may be needed to verify whether a given AS is authorized to originate an advertisement for a specified prefix, whether an given organization is the recognized holder of a block of address space or of an AS number, etc. Any infrastructure used to distribute data in support of BGP security is subject to the following criteria:

- o It MUST be resilient to attacks on the integrity of the data it contains.
- o It MUST enable network operators to verify the entity which originated the data.
- o It MUST be sufficiently available so as to not degrade the existing pace of network operations.
- o It SHOULD not introduce new organizational entities that have to be trusted in order to establish the authenticity of the data.

5. The Trust Model

In discussion with the operations community, concerns have emerged regarding the viability of a security system that requires agreement on a trust model dependent on a single root. Current operational practice has many providers engaging in bilateral agreements and preserving the primacy of local policy choices. The viability of a

solution may well rest on the business imperatives of the provider community who may be unwilling to surrender their perceived autonomy or unable to come to communal agreement on this topic.

In other environments, deployments may require an authority which has been selected by law or other institutional mandate. Moreover, these two deployment types (single-rooted hierarchy or arbitrary association) may "touch" (i.e. be part of the same co-extensive BGP topology).

Solutions **MUST** account for these differing types of deployments.

If two internetworks using differing trust models are interconnected they **MUST** be able to interoperate using locally determined levels of assurance to compensate for differences in these trust models. Some acknowledgement is made that this requirement might render it difficult to discern an attack from a difference in trust model or implementation. Any proposed solution **MUST** mitigate this risk.

6. The AS_PATH Attribute and NLRI Authentication

BGP distributes routing information across the Internet (between BGP speakers) using BGP UPDATE messages. The UPDATE message contains withdrawn routes, path attributes and NLRI (Network Layer Reachability Information, synonymous with advertised prefix(es)). For the remainder of this section, we will focus on the AS_PATH Attribute and the NLRI. Attributes such as MED are not transitive and, as such, are protected by BGP session security.

The AS_PATH for specific prefixes may be protected in any proposed security system in four ways, outlined below. Special Note: On the first two categories below, the community has reached consensus; on the latter two (AS_PATH Feasibility Check and Update Transit Check), the community has not reached consensus.

- o Authorization of Originating AS: For the purposes of authorization of the originating AS, authorization means that it **MUST** be possible to verify that the origin AS has been authorized to originate the route by the prefix holder(s).
- o Announcing AS Check: For all BGP peers, a BGP Implementation **MUST** ensure that the first element of the AS_PATH list corresponds to the locally configured AS of the peer from which the UPDATE was received.
- o AS_PATH Feasibility Check: The AS_PATH list **MUST** correspond to a valid list of autonomous systems according to the first

verification category listed in the "Areas to Secure" Section above.

- o Update Transit Check: Routing information carried through BGP SHOULD include information that can be used to verify the re-advertisement or modification by each autonomous system through which the UPDATE has passed. This check is more rigorous than the "valid list of autonomous systems" above.

The results of all of these checks SHOULD be made available to network operators. Each network operator will decide, on a local basis, which of these checks to enable.

There are many ways in which any difference between the speed of prefix/AS path attribute propagation and the availability of the information needed to validate the prefix/AS_PATH attribute information can be exploited to attack the routing system on a transient basis. These types of attacks primarily exploit the time it takes to follow the withdrawal of a route via an UPDATE. As a result of this potential for temporary disruption, BGP security solutions MUST be capable of distributing security information at the same rate as the BGP announcements and withdrawals propagate.

All data needed by BGP routers to evaluate the validity of an advertisement MUST be made available to the routers in a timeframe consistent with the rate at which advertisement characteristics change. Two examples are:

- o the distribution of information about the AS(es) authorized to advertise a given block of IP addresses,
- o the distribution of information about connectivity between autonomous systems and about autonomous system policies

Note that in today's operational Internet, the first two pieces of information, or their analogues, are not a part of the BGP routing system per se (e.g., information in Routing or Address registries.) They are consulted by most operators on an irregular basis and are not consulted in real time by the routing system. Policy information that is explicitly carried in the routing system is inconsistently expressed and consulted in Routing registries by operators. For instance, most providers are reticent to define their interconnection arrangements as transit or non-transit in Routing registries; some may do so, most do not. However, the ability to change inter-AS traffic flows in real time is an important feature of the current Internet.

7. Address Allocation and Advertisement

As part of the regular operation of the Internet, addresses allocated to one organization may be, and are quite commonly, advertised by ASes belonging to other organizations. Common reasons for this practice include multi-homing and route reduction for the purposes of resource conservation (e.g., aggregation). There are two modes of delegation:

- o A BGP speaker and listener have chosen to restrict the number of received prefixes for the listener. The listener has chosen to honor route announcements sent in a summary fashion by the speaker.
- o Address space that is being delegated is part of a larger allocation that is held by an autonomous system. The holder then delegates the smaller block to another AS for purposes of advertisement. This mode is commonly observed in multi-homing.

These two modes lead to a single common requirement: Any BGP Security solution **MUST** support the ability of an address block holder to declare (in a secure fashion) the AS(es) that the holder authorizes to originate routes to its address block(s) or any portion thereof regardless of the relationship of the entities.

An associated delegation criteria is the requirement to allow for non-BGP stub networks. As a result, all secured BGP implementations **MUST** allow for the contemporaneous origination of a route for a prefix by more than one AS.

8. Logging

In order to facilitate auditing and troubleshooting, a logging capability **MUST** be implemented that will indicate both negative and positive event behaviors. This data **SHALL** be for consumption of the AS operating the device that is producing the logs. Further, the information **MAY** be combined with data from other ASes or devices with different implementations within the same AS for purposes of event correlation and tracking. Here follow some considerations in this regard:

- o The data generated by logging may be very large depending on the number of peers, the number of prefixes received, the authentication model used, and routing policies. As such, efficient data structures and storage mechanisms **MUST** be developed to allow for an effective means of reproducing incidents and outages

- o Path and NLRI attributes MUST be logged using a standard format. The format MUST be scalable with the amount of data logged and the frequency of log generation. The frequency of log generation should be controllable by the operator. The logging mechanisms for the tracked information MUST be standardized across all platforms. Logging ability both on and off line is considered highly desirable.

9. NLRI and Path Attribute Tracking

The ability for a receiver to know the identity of each AS that originates and/or forwards a routing UPDATE is a desirable trait. In order to rapidly identify attack points and parties at fault for route table disruption, it is important to be able to track and log prefix origination information along with associated security information.

This capability can be afforded by implementation of the aforementioned directive that any security system SHOULD provide a method to allow the receiver of an UPDATE to verify that the originator is actually authorized to originate the update, and that the AS's listed in the AS_PATH actually forwarded the update.

10. Transport Layer Protection

Transport protection is an important aspect of BGP routing protocol security. The potential to create a linked transport/NLRI/AS_PATH authentication mechanism should not be overlooked and may provide for the accelerated deployment of a BGP security system. Current security mechanisms for BGP transport (e.g., TCP-MD5 [5] and GTSM [7]) are inadequate and require significant operator interaction to maintain a respectable level of security.

Transport protection systems MAY function as a component of the BGP routing protocol security mechanism. This includes the use of the same key generation/management systems as the rest of the security system.

Any proposed security mechanism MUST include provisions for securing both internal BGP and external BGP peering sessions.

11. Key Management

Current implementations and deployments of TCP-MD5 [5] exhibit serious shortcomings with regard of key management as described in

[RFC 3562](#) [6].

Key management can be especially onerous for operators. The number of keys required and the maintenance of keys (issue/revoke/renew) has had an additive effect as a barrier to deployment. Thus automated means of managing keys, to reduce operational burdens, **MUST** be available in proposed BGP security systems. These security systems **MUST** be resistant to compromise of session-level or device-level keys, i.e., the security implications of such compromises **MUST** be limited.

12. IANA Considerations

This document asks nothing of IANA.

13. Security Considerations

This document describes requirements for securing BGP as envisioned by the community. Its completeness is likely not exhaustive but represents the broadest consensus. As the understanding of the issues and possible residual vulnerabilities are refined, so these requirements may be revised in successor documents.

14. References

14.1. Normative References

- [1] Rekhter, Li, and Hares, "[RFC 4271](#) - A Border Gateway Protocol 4 (BGP-4)", October 2005.

14.2. Informative References

- [2] Fraser, "[RFC 2196](#) - Site Security Handbook", September 1997.
- [3] Rescorla, Korver, and Internet Architecture Board, "[RFC 3552](#) - Guidelines for Writing RFC Text on Security Considerations", July 2003.
- [4] Rekhter and Li, "[RFC 1771](#) - A Border Gateway Protocol 4 (BGP-4)", March 1995.
- [5] Heffernan, "[RFC 2385](#) - Protection of BGP Sessions via the TCP MD5 Signature Option", August 1998.
- [6] Leech, "[RFC 3562](#) - Key Management Considerations for the TCP MD5

Signature Option", July 2003.

- [7] Gill, Heasley, and Meyer, "[RFC 3682](#) - The Generalized TTL Security Mechanism (GTSM)", February 2004.

1. Acknowledgements

The following individuals contributed to the development and review of this draft. Steve Kent, Russ White, Sandy Murphy, Jeff Haas, Bora Akyol, Susan Hares, Mike Tibodeau, Thomas Renzy, Kaarthik Sivakumar, Tao Wan, Radia Perlman, Pekka Savola and Merike Kaeo.

This draft was developed based on conversations with various network operators including Chris Morrow, Jared Mauch, Tim Battles, and Ryan McDowell.

Authors' Addresses

Blaine Christian (editor)
KMC Telecom Solutions
1545 U.S. Highway 206
Bedminster, NJ 07921
US

Tony Tauber (editor)
Comcast
27 Industrial Avenue
Chelmsford, MA 01824
US

Email: ttauber@1-4-5.net

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

