Routing Protocol Security
Requirements
Internet-Draft
Expires: June 13, 2005

B. Christian
KMC Telecom Solutions
B. Akyol
R. White
Cisco Systems
J. Haas
Next Hop Technologies
S. Murphy
Trusted Information Systems
December 13, 2004

BGP Security Requirements
draft-ietf-rpsec-bgpsecrec-00

Status of this Memo

Copyright Notice

Abstract

The security of BGP 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 routed system, is more complex. This document describes a set of requirements for securing BGP, including securing peering relationships between BGP speakers, and authenticating the routing information carried within BGP.

1. Introduction

Threats to networking protocols generally fall under one of the three categories as defined in RFC 2196 [1]:

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

A number of attacks can be realized which, if exploited, can lead to one of the above mentioned threats. These are typically classified as passive attacks and active attacks. Passive attacks are ones where an attacker simply reads information off the network and obtains confidential and/or private information. Active attacks are ones where the attacker writes data to the network and can include replay attacks, message insertion, message deletion, message modification and man-in-the-middle attacks. These attacks are often combined.

Attacks that do not involve direct manipulation of BGP, and the information contained within it, are outside the scope of this document. When possible, the requirements will attempt to minimize the extent of the damage that occurs when end systems come under attack.

The intent of this requirements document is to prevent attacks that originate false data or create invalid routing paths and therefore addresses issues relating to data integrity and peer entity authentication. As described in RFC 3552 [2], data integrity protection ensures that data is not modified in transit and peer entity authentication ensures that there is a reasonable guarantee that the sender and recipient of the data are the intended parties.

Guaranteed packet delivery is not part of the BGP protocol security model. Just because a packet is addressed to a specific destination does not mean it will be received, even with a "secure" route. For example: an attacker could have compromised an intermediate router and installed a static route for target address A.B.C.D pointing to an inappropriate direction or an attacker might splice into a circuit between two secure routers and install a device that diverts A.B.C.D traffic without requiring the compromise of control plane devices.

2. Deployment Requirements

We have determined, through discussion with 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:

- 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 be equivalent 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, and others. Current timers, including hold timers, keepalive timers, and the peering process, SHOULD NOT be impacted by the security system. Two types of verification MAY be offered for the NLRI and the AS_PATH in order to allow for a selection of optimizations:
  - Contents of the UPDATE message SHOULD be authenticated in real-time as the UPDATE message is processed.
  - The route information base MAY be authenticated periodically or in an event driven manner by scanning the data and verifying the originating AS and the verifiability 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.
- Incremental Deployment: We will not be able to deploy a newly secured BGP protocol instantaneously and will be unable to dictate a partitioning of large internetworks by the operators. BGP MUST support both secured and unsecured routes with the security system in place. The security system MUST allow the forming of peering relationships between secure and non-secure BGP speakers, and MUST be backward compatible 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.
- Trust level Variability: Each secured environment may have different levels of requirements in terms of what is acceptable or unacceptable. In environments that require strict security it may not be acceptable to temporarily route to a destination while waiting for security verification to be performed. However, in many commercial environments the rapidity of route installation may be of paramount importance; in order to facilitate the more common occurence of route withdrawl due to network failure. Based

on the two divergent requirements, the security system MUST support a range of possible outputs for local determination of the trust level for a specific route. Any given route should be trustable to a locally configured degree, based on the completeness of security information for the update and other factors. The security system SHOULD allow the operator to determine whether the speed of convergence is more important than security operations, or security operations are 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. The Trust Model

In examining the various environments in which BGP is deployed, and through discussions with various operators working with the context of the public Internet, and other internetworks, it is apparent that trust models are largely environment specific. For instance, in the public Internet, a distributed trust model, following the current transitive trust pattern of contractual and peering arrangements, would fit the the business models of the participants. In other environments a hierarchical trust model would work better. Thus, any trust system specified in a security mechanism designed for BGP must be flexible, and support both a true distributed trust model and a fully hierarchical trust model.

Since hierarchical trust models are a subset (or a special case of) a distributed trust model, any security system designed for BGP MUST support a distributed trust model, and MUST also support a hierarchical trust model, if desired.

If two internetworks using differing trust models are interconnected they MUST be able to interoperate using locally determined levels of trust to compensate for differences in their trust models.

4. 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 one or more NLRIs (Network Layer Reachability Information is synonymous with advertised prefix). For the remainder of this section, we will focus on the AS-Path Attribute and the NLRI. Attributes such as local pref are locally specific and, as such, are protected by BGP session security.

The AS_PATH for specific prefixes must be protected in any proposed security system in three ways:

- Authorization of Originating AS: For all prefixes announced in BGP, the originating AS MUST be verifiable through the trust model as the authorized announcer of the prefix. The verification mechanism must account for existing BGP mechanisms such as summarization. For the purpose of this document the term verifiable is defined as the resultant of a secured routing systems as described in this document. The term specifically indicates the ability to validate the originator of a specific prefix (or block of IP addresses) and the ability to validate the session through which the prefix was received
- The AS_PATH list MUST correspond to a verifiable list of autonomous systems based on the peering topology of the network.
- 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 that peer.

There are many ways in which a differential between the speed of prefix/AS path attribute propagation and the information validating the the prefix/AS path attribute information can be exploited to attack the routing system on a temporary basis. These types of attacks are dominantly exploitative of the moment in time it takes to follow the withdraw of a NLRI with an update. As a result of this potential for temporary disruption, BGP security solutions MUST propagate security information at the same rate as the BGP updates and withdrawls. The following items are required to propagate at the same rate:

- The distribution of key information used by individual actors within the system, including the keys used by individual autonomous systems to sign certificates and other objects
- The distribution of information about the AS authorized to advertise a given block of IP addresses (or an address space)
- The distribution of information about connectivity between autonomous systems and autonomous system polices, if such information is to be distributed within the security system.

5. Address Allocation and advertisement

As part of the regular operation of the Internet, addresses that are allocated to an organization may be, and are quite commonly, advertised by a different organizations. Common reasons for this practice include multi-homing and route reduction for the purposes of resource conservation. There are two modes of delegation:

- A BGP speaker and listener have chosen to restrict the amount of received prefixes for the listener. The listener has chosen to honor route announcements sent in a summary fashion by the speaker.
- Address space that is being delegated is part of a larger allocation that is owned by an autonomous system. The owner 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 delegation of an address block of any size regardless of its relationship to other address blocks to another entity via verifiable means.

An associated delegation criteria is the requirement to allow for non-BGP IP end user implementations. As a result, all secured BGP implementations MUST allow for the propagation of a prefix by more than one originator AS within normal network convergence times.

6. NLRI and Path Attribute Tracking

Non-repudiation of routing updates, the ability for a receiver to know exactly who originated and forwarded a routing update, is a desirable trait. In order to rapidly identify agressors and parties at fault for route table disruption it is important to track and log prefix origination information along with associated security information.

Any security system SHOULD provide a method to allow the receiver of an update to verify that the originator actually originated the update, and that the AS's listed in the AS_PATH actually forwarded the update.

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

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.

7. Transport 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 are inadequate and require significant operator interaction to maintain a respectable level of security.

Any proposed security mechanism MUST include provisions for securing both internal BGP and external BGP peering sessions. Key maintenance can be especially onerous to the operators. The number of keys required and the maintenance of keys (update/withdraw/renew) may have an additive affect to a barrier to deployment. A highly securable BGP routing system SHOULD require no more than three keys and each key should be updateable within similar timeframes as prefix propagation. The preferred number of keys is ONE per AS.

Transport protection systems SHOULD 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.

8 References

[1] Fraser, "RFC 2196 - Site Security Handbook", September 1997.

[2] Rescorla, Korver and Internet Architecture Board, "RFC 3552 - Guidelines for Writing RFC Text on Security Considerations", July 2003.

Authors' Addresses

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

Bora Akyol
Cisco Systems
170 Tasman Drive
San Jose, CA 95134
US

Russ White
Cisco Systems
7025 Kit Creek Road
Research Triangle Park, NC 95134
US

Jeffrey Haas
Next Hop Technologies
825 Victors Way Suite 100
Ann Arbor, MI 48108
US

Sandy Murphy
Trusted Information Systems
3060 Washington Road
Glenwood, MD 21378
US

Appendix A. Acknowledgements

The following individuals contributed to the development and review of this draft. Steve Kent, Mike Tibodeau, Thomas Renzy, Kaarthik Sivakumar, Tao Wan, Radia Perlman, and Merike Kaeo.

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

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.