

Network Working Group
Request for Comments: 4272
Category: Informational

S. Murphy
Sparta, Inc.
January 2006

BGP Security Vulnerabilities Analysis

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

Border Gateway Protocol 4 (BGP-4), along with a host of other infrastructure protocols designed before the Internet environment became perilous, was originally designed with little consideration for protection of the information it carries. There are no mechanisms internal to BGP that protect against attacks that modify, delete, forge, or replay data, any of which has the potential to disrupt overall network routing behavior.

This document discusses some of the security issues with BGP routing data dissemination. This document does not discuss security issues with forwarding of packets.

Table of Contents

1. Introduction	3
1.1. Specification of Requirements	5
2. Attacks	6
3. Vulnerabilities and Risks	7
3.1. Vulnerabilities in BGP Messages	8
3.1.1. Message Header	9
3.1.2. OPEN	9
3.1.3. KEEPALIVE	11
3.1.4. NOTIFICATION	11
3.1.5. UPDATE	11
3.1.5.1. Unfeasible Routes Length, Total Path Attribute Length	12
3.1.5.2. Withdrawn Routes	13
3.1.5.3. Path Attributes	13
3.1.5.4. NLRI	16
3.2. Vulnerabilities through Other Protocols	16
3.2.1. TCP Messages	16
3.2.1.1. TCP SYN	16
3.2.1.2. TCP SYN ACK	17
3.2.1.3. TCP ACK	17
3.2.1.4. TCP RST/FIN/FIN-ACK	17
3.2.1.5. DoS and DDos	18
3.2.2. Other Supporting Protocols	18
3.2.2.1. Manual Stop	18
3.2.2.2. Open Collision Dump	18
3.2.2.3. Timer Events	18
4. Security Considerations	19
4.1. Residual Risk	19
4.2. Operational Protections	19
5. References	21
5.1. Normative References	21
5.2. Informative References	21

1. Introduction

The inter-domain routing protocol BGP was created when the Internet environment had not yet reached the present, contentious state. Consequently, the BGP design did not include protections against deliberate or accidental errors that could cause disruptions of routing behavior.

This document discusses the vulnerabilities of BGP, based on the BGP specification [RFC4271]. Readers are expected to be familiar with the BGP RFC and the behavior of BGP.

It is clear that the Internet is vulnerable to attack through its routing protocols and BGP is no exception. Faulty, misconfigured, or deliberately malicious sources can disrupt overall Internet behavior by injecting bogus routing information into the BGP-distributed routing database (by modifying, forging, or replaying BGP packets). The same methods can also be used to disrupt local and overall network behavior by breaking the distributed communication of information between BGP peers. The sources of bogus information can be either outsiders or true BGP peers.

Cryptographic authentication of peer-peer communication is not an integral part of BGP. As a TCP/IP protocol, BGP is subject to all TCP/IP attacks, e.g., IP spoofing, session stealing, etc. Any outsider can inject believable BGP messages into the communication between BGP peers, and thereby inject bogus routing information or break the peer-peer connection. Any break in the peer-peer communication has a ripple effect on routing that can be widespread. Furthermore, outsider sources can also disrupt communications between BGP peers by breaking their TCP connection with spoofed packets. Outsider sources of bogus BGP information can reside anywhere in the world.

Consequently, the current BGP specification requires that a BGP implementation must support the authentication mechanism specified in [TCPMD5]. However, the requirement for support of that authentication mechanism cannot ensure that the mechanism is configured for use. The mechanism of [TCPMD5] is based on a pre-installed, shared secret; it does not have the capability of IPsec [IPsec] to agree on a shared secret dynamically. Consequently, the use of [TCPMD5] must be a deliberate decision, not an automatic feature or a default.

The current BGP specification also allows for implementations that would accept connections from "unconfigured peers" ([RFC4271] Section 8). However, the specification is not clear as to what an unconfigured peer might be, or how the protections of [TCPMD5] would

apply in such a case. Therefore, it is not possible to include an analysis of the security issues of this feature. When a specification that describes this feature more fully is released, a security analysis should be part of that specification.

BGP speakers themselves can inject bogus routing information, either by masquerading as any other legitimate BGP speaker, or by distributing unauthorized routing information as themselves. Historically, misconfigured and faulty routers have been responsible for widespread disruptions in the Internet. The legitimate BGP peers have the context and information to produce believable, yet bogus, routing information, and therefore have the opportunity to cause great damage. The cryptographic protections of [TCPMD5] and operational protections cannot exclude the bogus information arising from a legitimate peer. The risk of disruptions caused by legitimate BGP speakers is real and cannot be ignored.

Bogus routing information can have many different effects on routing behavior. If the bogus information removes routing information for a particular network, that network can become unreachable for the portion of the Internet that accepts the bogus information. If the bogus information changes the route to a network, then packets destined for that network may be forwarded by a sub-optimal path, or by a path that does not follow the expected policy, or by a path that will not forward the traffic. Consequently, traffic to that network could be delayed by a path that is longer than necessary. The network could become unreachable from areas where the bogus information is accepted. Traffic might also be forwarded along a path that permits some adversary to view or modify the data. If the bogus information makes it appear that an autonomous system originates a network when it does not, then packets for that network may not be deliverable for the portion of the Internet that accepts the bogus information. A false announcement that an autonomous systems originates a network may also fragment aggregated address blocks in other parts of the Internet and cause routing problems for other networks.

The damages that might result from these attacks include:

starvation: Data traffic destined for a node is forwarded to a part of the network that cannot deliver it.

network congestion: More data traffic is forwarded through some portion of the network than would otherwise need to carry the traffic.

blackhole: Large amounts of traffic are directed to be forwarded through one router that cannot handle the increased level of traffic and drops many/most/all packets.

delay: Data traffic destined for a node is forwarded along a path that is in some way inferior to the path it would otherwise take.

looping: Data traffic is forwarded along a path that loops, so that the data is never delivered.

eavesdrop: Data traffic is forwarded through some router or network that would otherwise not see the traffic, affording an opportunity to see the data.

partition: Some portion of the network believes that it is partitioned from the rest of the network, when, in fact, it is not.

cut: Some portion of the network believes that it has no route to some network to which it is, in fact, connected.

churn: The forwarding in the network changes at a rapid pace, resulting in large variations in the data delivery patterns (and adversely affecting congestion control techniques).

instability: BGP becomes unstable in such a way that convergence on a global forwarding state is not achieved.

overload: The BGP messages themselves become a significant portion of the traffic the network carries.

resource exhaustion: The BGP messages themselves cause exhaustion of critical router resources, such as table space.

address-spoofing: Data traffic is forwarded through some router or network that is spoofing the legitimate address, thus enabling an active attack by affording the opportunity to modify the data.

These consequences can fall exclusively on one end-system prefix or may effect the operation of the network as a whole.

1.1. Specification of Requirements

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 RFC2119 [RFC2119].

2. Attacks

BGP, in and of itself, is subject to the following attacks. (The list is taken from the IAB RFC that provides guidelines for the "Security Considerations" section of RFCs [SecCons].)

confidentiality violations: The routing data carried in BGP is carried in cleartext, so eavesdropping is a possible attack against routing data confidentiality. (Routing data confidentiality is not a common requirement.)

replay: BGP does not provide for replay protection of its messages.

message insertion: BGP does not provide protection against insertion of messages. However, because BGP uses TCP, when the connection is fully established, message insertion by an outsider would require accurate sequence number prediction (not entirely out of the question, but more difficult with mature TCP implementations) or session-stealing attacks.

message deletion: BGP does not provide protection against deletion of messages. Again, this attack is more difficult against a mature TCP implementation, but is not entirely out of the question.

message modification: BGP does not provide protection against modification of messages. A modification that was syntactically correct and did not change the length of the TCP payload would in general not be detectable.

man-in-the-middle: BGP does not provide protection against man-in-the-middle attacks. As BGP does not perform peer entity authentication, a man-in-the-middle attack is child's play.

denial of service: While bogus routing data can present a denial of service attack on the end systems that are trying to transmit data through the network and on the network infrastructure itself, certain bogus information can represent a denial of service on the BGP routing protocol. For example, advertising large numbers of more specific routes (i.e., longer prefixes) can cause BGP traffic and router table size to increase, even explode.

The mandatory-to-support mechanism of [TCPMD5] will counter message insertion, deletion, and modification, man-in-the-middle and denial of service attacks from outsiders. The use of [TCPMD5] does not protect against eavesdropping attacks, but routing data confidentiality is not a goal of BGP. The mechanism of [TCPMD5] does

not protect against replay attacks, so the only protection against replay is provided by the TCP sequence number processing. Therefore, a replay attack could be mounted against a BGP connection protected with [TCPMD5] but only in very carefully timed circumstances. The mechanism of [TCPMD5] cannot protect against bogus routing information that originates from an insider.

3. Vulnerabilities and Risks

The risks in BGP arise from three fundamental vulnerabilities:

- (1) BGP has no internal mechanism that provides strong protection of the integrity, freshness, and peer entity authenticity of the messages in peer-peer BGP communications.
- (2) no mechanism has been specified within BGP to validate the authority of an AS to announce NLRI information.
- (3) no mechanism has been specified within BGP to ensure the authenticity of the path attributes announced by an AS.

The first fundamental vulnerability motivated the mandated support of [TCPMD5] in the BGP specification. When the support of [TCPMD5] is employed, message integrity and peer entity authentication are provided. The mechanism of [TCPMD5] assumes that the MD5 algorithm is secure and that the shared secret is protected and chosen to be difficult to guess.

In the discussion that follows, the vulnerabilities are described in terms of the BGP Finite State Machine events. The events are defined and discussed in section 8 of [RFC4271]. The events mentioned here are:

[Administrative Events]

Event 2: ManualStop

Event 8: AutomaticStop

[Timer Events]

Event 9: ConnectRetryTimer_Expires

Event 10: HoldTimer_Expires

Event 11: KeepaliveTimer_Expires

Event 12: DelayOpenTimer_Expires

Event 13: IdleHoldTimer_Expires

[TCP Connection based Events]

Event 14: TcpConnection_Valid

Event 16: Tcp_CR_Acked

Event 17: TcpConnectionConfirmed

Event 18: TcpConnectionFails

[BGP Messages based Events]

Event 19: BGPOpen

Event 20: BGPOpen with DelayOpenTimer running

Event 21: BGPHeaderErr

Event 22: BGPOpenMsgErr

Event 23: OpenCollisionDump

Event 24: NotifMsgVerErr

Event 25: NotifMsg

Event 26: KeepAliveMsg

Event 27: UpdateMsg

Event 28: UpdateMsgErr

3.1. Vulnerabilities in BGP Messages

There are four different BGP message types - OPEN, KEEPALIVE, NOTIFICATION, and UPDATE. This section contains a discussion of the vulnerabilities arising from each message and the ability of outsiders or BGP peers to exploit the vulnerabilities. To summarize, outsiders can use bogus OPEN, KEEPALIVE, NOTIFICATION, or UPDATE messages to disrupt the BGP peer-peer connections. They can use bogus UPDATE messages to disrupt routing without breaking the peer-peer connection. Outsiders can also disrupt BGP peer-peer connections by inserting bogus TCP packets that disrupt the TCP connection processing. In general, the ability of outsiders to use bogus BGP and TCP messages is limited, but not eliminated, by the TCP sequence number processing. The use of [TCPMD5] can counter these

outsider attacks. BGP peers themselves are permitted to break peer-peer connections, at any time, using NOTIFICATION messages. Thus, there is no additional risk of broken connections through their use of OPEN, KEEPALIVE, or UPDATE messages. However, BGP peers can disrupt routing (in impermissible ways) by issuing UPDATE messages that contain bogus routing information. In particular, bogus ATOMIC_AGGREGATE, NEXT_HOP and AS_PATH attributes and bogus NLRI in UPDATE messages can disrupt routing. The use of [TCPMD5] will not counter these attacks from BGP peers.

Each message introduces certain vulnerabilities and risks, which are discussed in the following sections.

3.1.1. Message Header

Event 21: Each BGP message starts with a standard header. In all cases, syntactic errors in the message header will cause the BGP speaker to close the connection, release all associated BGP resources, delete all routes learned through that connection, run its decision process to decide on new routes, and cause the state to return to Idle. Also, optionally, an implementation-specific peer oscillation damping may be performed. The peer oscillation damping process can affect how soon the connection can be restarted. An outsider who could spoof messages with message header errors could cause disruptions in routing over a wide area.

3.1.2. OPEN

Event 19: Receipt of an OPEN message in states Connect or Active will cause the BGP speaker to bring down the connection, release all associated BGP resources, delete all associated routes, run its decision process, and cause the state to return to Idle. The deletion of routes can cause a cascading effect in which routing changes propagate through other peers. Also, optionally, an implementation-specific peer oscillation damping may be performed. The peer oscillation damping process can affect how soon the connection can be restarted.

In state OpenConfirm or Established, the arrival of an OPEN may indicate a connection collision has occurred. If this connection is to be dropped, then Event 23 will be issued. (Event 23, discussed below, results in the same set of disruptive actions as mentioned above for states Connect or Active.)

In state OpenSent, the arrival of an OPEN message will cause the BGP speaker to transition to the OpenConfirm state. If an outsider was able to spoof an OPEN message (requiring very careful timing), then the later arrival of the legitimate peer's OPEN message might lead

the BGP speaker to declare a connection collision. The collision detection procedure may cause the legitimate connection to be dropped.

Consequently, the ability of an outsider to spoof this message can lead to a severe disruption of routing over a wide area.

Event 20: If an OPEN message arrives when the DelayOpen timer is running when the connection is in state OpenSent, OpenConfirm or Established, the BGP speaker will bring down the connection, release all associated BGP resources, delete all associated routes, run its decision process, and cause the state to return to Idle. The deletion of routes can cause a cascading effect in which routing changes propagate through other peers. Also, optionally, an implementation-specific peer oscillation damping may be performed. The peer oscillation damping process can affect how soon the connection can be restarted. However, because the OpenDelay timer should never be running in these states, this effect could only be caused by an error in the implementation (a NOTIFICATION is sent with the error code "Finite State Machine Error"). It would be difficult, if not impossible, for an outsider to induce this Finite State Machine error.

In states Connect and Active, this event will cause a transition to the OpenConfirm state. As in Event 19, if an outsider were able to spoof an OPEN, which arrived while the DelayOpen timer was running, then a later arriving OPEN (from the legitimate peer) might be considered a connection collision and the legitimate connection could be dropped.

Consequently, the ability of an outsider to spoof this message can lead to a severe disruption of routing over a wide area.

Event 22: Errors in the OPEN message (e.g., unacceptable Hold state, malformed Optional Parameter, unsupported version, etc.) will cause the BGP speaker to bring down the connection, release all associated BGP resources, delete all associated routes, run its decision process, and cause the state to return to Idle. The deletion of routes can cause a cascading effect in which routing changes propagate through other peers. Also, optionally, an implementation-specific peer oscillation damping may be performed. The peer oscillation damping process can affect how soon the connection can be restarted. Consequently, the ability of an outsider to spoof this message can lead to a severe disruption of routing over a wide area.

3.1.3. KEEPALIVE

Event 26: Receipt of a KEEPALIVE message, when the peering connection is in the Connect, Active, and OpenSent states, would cause the BGP speaker to transition to the Idle state and fail to establish a connection. Also, optionally, an implementation-specific peer oscillation damping may be performed. The peer oscillation damping process can affect how soon the connection can be restarted. The ability of an outsider to spoof this message can lead to a disruption of routing. To exploit this vulnerability deliberately, the KEEPALIVE must be carefully timed in the sequence of messages exchanged between the peers; otherwise, it causes no damage.

3.1.4. NOTIFICATION

Event 25: Receipt of a NOTIFICATION message in any state will cause the BGP speaker to bring down the connection, release all associated BGP resources, delete all associated routes, run its decision process, and cause the state to return to Idle. The deletion of routes can cause a cascading effect in which routing changes propagate through other peers. Also, optionally, in any state but Established, an implementation-specific peer oscillation damping may be performed. The peer oscillation damping process can affect how soon the connection can be restarted. Consequently, the ability of an outsider to spoof this message can lead to a severe disruption of routing over a wide area.

Event 24: A NOTIFICATION message carrying an error code of "Version Error" behaves the same as in Event 25, with the exception that the optional peer oscillation damping is not performed in states OpenSent or OpenConfirm, or in states Connect or Active if the DelayOpen timer is running. Therefore, the damage caused is one small bit less, because restarting the connection is not affected.

3.1.5. UPDATE

Event 8: A BGP speaker may optionally choose to automatically disconnect a BGP connection if the total number of prefixes exceeds a configured maximum. In such a case, an UPDATE may carry a number of prefixes that would result in that maximum being exceeded. The BGP speaker would disconnect the connection, release all associated BGP resources, delete all associated routes, run its decision process, and cause the state to return to Idle. The deletion of routes can cause a cascading effect in which routing changes propagate through other peers. Also, optionally, an implementation-specific peer oscillation damping may be performed. The peer oscillation damping

process can affect how soon the connection can be restarted. Consequently, the ability of an outsider to spoof this message can lead to a severe disruption of routing over a wide area.

Event 28: If the UPDATE message is malformed, then the BGP speaker will bring down the connection, release all associated BGP resources, delete all associated routes, run its decision process, and cause the state to return to Idle. (Here, "malformed" refers to improper Withdrawn Routes Length, Total Attribute Length, or Attribute Length, missing mandatory well-known attributes, Attribute Flags that conflict with the Attribute Type Codes, syntactic errors in the ORIGIN, NEXT_HOP or AS_PATH, etc.) The deletion of routes can cause a cascading effect in which routing changes propagate through other peers. Also, optionally, an implementation-specific peer oscillation damping may be performed. The peer oscillation damping process can affect how soon the connection can be restarted. Consequently, the ability of an outsider to spoof this message could cause widespread disruption of routing. As a BGP speaker has the authority to close a connection whenever it wants, this message gives BGP speakers no additional opportunity to cause damage.

Event 27: An Update message that arrives in any state except Established will cause the BGP speaker to bring down the connection, release all associated BGP resources, delete all associated routes, run its decision process, and cause the state to return to Idle. The deletion of routes can cause a cascading effect in which routing changes propagate through other peers. Also, optionally, an implementation-specific peer oscillation damping may be performed. The peer oscillation damping process can affect how soon the connection can be restarted. Consequently, the ability of an outsider to spoof this message can lead to a severe disruption of routing over a wide area.

In the Established state, the Update message carries the routing information. The ability to spoof any part of this message can lead to a disruption of routing, whether the source of the message is an outsider or a legitimate BGP speaker.

3.1.5.1. Unfeasible Routes Length, Total Path Attribute Length

There is a vulnerability arising from the ability to modify these fields. If a length is modified, the message is not likely to parse properly, resulting in an error, the transmission of a NOTIFICATION message and the close of the connection (see Event 28, above). As a true BGP speaker is able to close a connection at any time, this vulnerability represents an additional risk only when the source is not the configured BGP peer, i.e., it presents no additional risk from BGP speakers.

3.1.5.2. Withdrawn Routes

An outsider could cause the elimination of existing legitimate routes by forging or modifying this field. An outsider could also cause the elimination of reestablished routes by replaying this withdrawal information from earlier packets.

A BGP speaker could "falsely" withdraw feasible routes using this field. However, as the BGP speaker is authoritative for the routes it will announce, it is allowed to withdraw any previously announced routes that it wants. As the receiving BGP speaker will only withdraw routes associated with the sending BGP speaker, there is no opportunity for a BGP speaker to withdraw another BGP speaker's routes. Therefore, there is no additional risk from BGP peers via this field.

3.1.5.3. Path Attributes

The path attributes present many different vulnerabilities and risks.

- o Attribute Flags, Attribute Type Codes, Attribute Length

A BGP peer or an outsider could modify the attribute length or attribute type (flags and type codes) not to reflect the attribute values that followed. If the flags were modified, the flags and type code could become incompatible (i.e., a mandatory attribute marked as partial), or an optional attribute could be interpreted as a mandatory attribute or vice versa. If the type code were modified, the attribute value could be interpreted as if it were the data type and value of a different attribute.

The most likely result from modifying the attribute length, flags, or type code would be a parse error of the UPDATE message. A parse error would cause the transmission of a NOTIFICATION message and the close of the connection (see Event 28, above). As a true BGP speaker is able to close a connection at any time, this vulnerability represents an additional risk only when the source is an outsider, i.e., it presents no additional risk from a BGP peer.

- o ORIGIN

This field indicates whether the information was learned from IGP or EGP information. This field is used in making routing decisions, so there is some small vulnerability of being able to affect the receiving BGP speaker's routing decision by modifying this field.

- o AS_PATH

A BGP peer or outsider could announce an AS_PATH that was not accurate for the associated NLRI.

Because a BGP peer might not verify that a received AS_PATH begins with the AS number of its peer, a malicious BGP peer could announce a path that begins with the AS of any BGP speaker, with little impact on itself. This could affect the receiving BGP speaker's decision procedure and choice of installed route. The malicious peer could considerably shorten the AS_PATH, which will increase that route's chances of being chosen, possibly giving the malicious peer access to traffic it would otherwise not receive. The shortened AS_PATH also could result in routing loops, as it does not contain the information needed to prevent loops.

It is possible for a BGP speaker to be configured to accept routes with its own AS number in the AS path. Such operational considerations are defined to be "outside the scope" of the BGP specification. But because AS_PATHS can legitimately have loops, implementations cannot automatically reject routes with loops. Each BGP speaker verifies only that its own AS number does not appear in the AS_PATH.

Coupled with the ability to use any value for the NEXT_HOP, this provides a malicious BGP speaker considerable control over the path traffic will take.

- o Originating Routes

A special case of announcing a false AS_PATH occurs when the AS_PATH advertises a direct connection to a specific network address. A BGP peer or outsider could disrupt routing to the network(s) listed in the NLRI field by falsely advertising a direct connection to the network. The NLRI would become unreachable to the portion of the network that accepted this false route, unless the ultimate AS on the AS_PATH undertook to tunnel the packets it was forwarded for this NLRI toward their true destination AS by a valid path. But even when the packets are tunneled to the correct destination AS, the route followed may not be optimal, or may not follow the intended policy. Additionally, routing for other networks in the Internet could be affected if the false advertisement fragmented an aggregated address block, forcing the routers to handle (issue UPDATES, store, manage) the multiple fragments rather than the single aggregate. False originations for multiple addresses can result in routers and transit networks along the announced route to become flooded with misdirected traffic.

- o NEXT_HOP

The NEXT_HOP attribute defines the IP address of the border router that should be used as the next hop when forwarding the NLRI listed in the UPDATE message. If the recipient is an external peer, then the recipient and the NEXT_HOP address must share a subnet. It is clear that an outsider who modified this field could disrupt the forwarding of traffic between the two ASes.

If the recipient of the message is an external peer of an AS and the route was learned from another peer AS (this is one of two forms of "third party" NEXT_HOP), then the BGP speaker advertising the route has the opportunity to direct the recipient to forward traffic to a BGP speaker at the NEXT_HOP address. This affords the opportunity to direct traffic at a router that may not be able to continue forwarding the traffic. A malicious BGP speaker can also use this technique to force another AS to carry traffic it would otherwise not have to carry. In some cases, this could be to the malicious BGP speaker's benefit, as it could cause traffic to be carried long-haul by the victim AS to some other peering point it shared with the victim.

- o MULTI_EXIT_DISC

The MULTI_EXIT_DISC attribute is used in UPDATE messages transmitted between inter-AS BGP peers. While the MULTI_EXIT_DISC received from an inter-AS peer may be propagated within an AS, it may not be propagated to other ASes. Consequently, this field is only used in making routing decisions internal to one AS. Modifying this field, whether by an outsider or a BGP peer, could influence routing within an AS to be sub-optimal, but the effect should be limited in scope.

- o LOCAL_PREF

The LOCAL_PREF attribute must be included in all messages with internal peers, and excluded from messages with external peers. Consequently, modification of the LOCAL_PREF could effect the routing process within the AS only. Note that there is no requirement in the BGP RFC that the LOCAL_PREF be consistent among the internal BGP speakers of an AS. Because BGP peers are free to choose the LOCAL_PREF, modification of this field is a vulnerability with respect to outsiders only.

- o ATOMIC_AGGREGATE

The ATOMIC_AGGREGATE field indicates that an AS somewhere along the way has aggregated several routes and advertised the aggregate NLRI without the AS_SET being formed as usual from the ASes in the aggregated routes' AS_PATHs. BGP speakers receiving a route with ATOMIC_AGGREGATE are restricted from making the NLRI any more specific. Removing the ATOMIC_AGGREGATE attribute would remove the restriction, possibly causing traffic intended for the more specific NLRI to be routed incorrectly. Adding the ATOMIC_AGGREGATE attribute, when no aggregation was done, would have little effect beyond restricting the un-aggregated NLRI from being made more specific. This vulnerability exists whether the source is a BGP peer or an outsider.

- o AGGREGATOR

This field may be included by a BGP speaker who has computed the routes represented in the UPDATE message by aggregating other routes. The field contains the AS number and IP address of the last aggregator of the route. It is not used in making any routing decisions, so it does not represent a vulnerability.

3.1.5.4. NLRI

By modifying or forging this field, either an outsider or BGP peer source could cause disruption of routing to the announced network, overwhelm a router along the announced route, cause data loss when the announced route will not forward traffic to the announced network, route traffic by a sub-optimal route, etc.

3.2. Vulnerabilities through Other Protocols

3.2.1. TCP Messages

BGP runs over TCP, listening on port 179. Therefore, BGP is subject to attack through attacks on TCP.

3.2.1.1. TCP SYN

SYN flooding: Like other protocols, BGP is subject to the effects on the TCP implementation of SYN flooding attacks, and must rely on the implementation's protections against these attacks.

Event 14: If an outsider were able to send a SYN to the BGP speaker at the appropriate time during connection establishment, then the legitimate peer's SYN would appear to be a second connection. If the outsider were able to continue with a sequence of packets resulting

in a BGP connection (guessing the BGP speaker's choice for sequence number on the SYN ACK, for example), then the outsider's connection and the legitimate peer's connection would appear to be a connection collision. Depending on the outcome of the collision detection (i.e., if the outsider chooses a BGP identifier so as to win the race), the legitimate peer's true connection could be destroyed. The use of [TCPMD5] can counter this attack.

3.2.1.2. TCP SYN ACK

Event 16: If an outsider were able to respond to a BGP speaker's SYN before the legitimate peer, then the legitimate peer's SYN-ACK would receive an empty ACK reply, causing the legitimate peer to issue a RST that would break the connection. The BGP speaker would bring down the connection, release all associated BGP resources, delete all associated routes, and run its decision process. This attack requires that the outsider be able to predict the sequence number used in the SYN. The use of [TCPMD5] can counter this attack.

3.2.1.3. TCP ACK

Event 17: If an outsider were able to spoof an ACK at the appropriate time during connection establishment, then the BGP speaker would consider the connection complete, send an OPEN (Event 17), and transition to the OpenSent state. The arrival of the legitimate peer's ACK would not be delivered to the BGP process, as it would look like a duplicate packet. Thus, this message does not present a vulnerability to BGP during connection establishment. Spoofing an ACK after connection establishment requires knowledge of the sequence numbers in use, and is, in general, a very difficult task. The use of [TCPMD5] can counter this attack.

3.2.1.4. TCP RST/FIN/FIN-ACK

Event 18: If an outsider were able to spoof a RST, the BGP speaker would bring down the connection, release all associated BGP resources, delete all associated routes, and run its decision process. If an outsider were able to spoof a FIN, then data could still be transmitted, but any attempt to receive it would trigger a notification that the connection is closing. In most cases, this results in the connection being placed in an Idle state. But if the connection is in the Connect state or the OpenSent state at the time, the connection will return to an Active state.

Spoofing a RST in this situation requires an outsider to guess a sequence number that need only be within the receive window [Watson04]. This is generally an easier task than guessing the exact

sequence number required to spoof a FIN. The use of [TCPMD5] can counter this attack.

3.2.1.5. DoS and DDos

Because the packets directed to TCP port 179 are passed to the BGP process, which potentially resides on a slower processor in the router, flooding a router with TCP port 179 packets is an avenue for DoS attacks against the router. No BGP mechanism can defeat such attacks; other mechanisms must be employed.

3.2.2. Other Supporting Protocols

3.2.2.1. Manual Stop

Event 2: A manual stop event causes the BGP speaker to bring down the connection, release all associated BGP resources, delete all associated routes, and run its decision process. If the mechanism by which a BGP speaker was informed of a manual stop is not carefully protected, the BGP connection could be destroyed by an outsider. Consequently, BGP security is secondarily dependent on the security of the management and configuration protocols that are used to signal this event.

3.2.2.2. Open Collision Dump

Event 23: The OpenCollisionDump event may be generated administratively when a connection collision event is detected and the connection has been selected to be disconnected. When this event occurs in any state, the BGP connection is dropped, the BGP resources are released, the associated routes are deleted, etc. Consequently, BGP security is secondarily dependent on the security of the management and configuration protocols that are used to signal this event.

3.2.2.3. Timer Events

Events 9-13: BGP employs five timers (ConnectRetry, Hold, Keepalive, MinASOrigination-Interval, and MinRouteAdvertisementInterval) and two optional timers (DelayOpen and IdleHold). These timers are critical to BGP operation. For example, if the Hold timer value were changed, the remote peer might consider the connection unresponsive and bring the connection down, thus releasing resources, deleting associated routes, etc. Consequently, BGP security is secondarily dependent on the security of the operation, management, and configuration protocols that are used to modify the timers.

4. Security Considerations

This entire memo is about security, describing an analysis of the vulnerabilities that exist in BGP.

Use of the mandatory-to-support mechanisms of [TCPMD5] counters the message insertion, deletion, and modification attacks, as well as man-in-the-middle attacks by outsiders. If routing data confidentiality is desired (there is some controversy as to whether it is a desirable security service), the use of IPsec ESP could provide that service.

4.1. Residual Risk

As cryptographic-based mechanisms, both [TCPMD5] and IPsec [IPsec] assume that the cryptographic algorithms are secure, that secrets used are protected from exposure and are chosen well so as not to be guessable, that the platforms are securely managed and operated to prevent break-ins, etc.

These mechanisms do not prevent attacks that arise from a router's legitimate BGP peers. There are several possible solutions to prevent a BGP speaker from inserting bogus information in its advertisements to its peers (i.e., from mounting an attack on a network's origination or AS-PATH):

- (1) Origination Protection: sign the originating AS.
- (2) Origination and Adjacency Protection: sign the originating AS and predecessor information ([Smith96])
- (3) Origination and Route Protection: sign the originating AS, and nest signatures of AS_PATHs to the number of consecutive bad routers you want to prevent from causing damage. ([SBGP00])
- (4) Filtering: rely on a registry to verify the AS_PATH and NLRI originating AS ([RPSL]).

Filtering is in use near some customer attachment points, but is not effective near the Internet center. The other mechanisms are still controversial and are not yet in common use.

4.2. Operational Protections

BGP is primarily used as a means to provide reachability information to Autonomous Systems (AS) and to distribute external reachability internally within an AS. BGP is the routing protocol used to

distribute global routing information in the Internet. Therefore, BGP is used by all major Internet Service Providers (ISP), as well as many smaller providers and other organizations.

BGP's role in the Internet puts BGP implementations in unique conditions, and places unique security requirements on BGP. BGP is operated over interprovider interfaces in which traffic levels push the state of the art in specialized packet forwarding hardware and exceed the performance capabilities of hardware implementation of decryption by many orders of magnitude. The capability of an attacker using a single workstation with high speed interface to generate false traffic for denial of service (DoS) far exceeds the capability of software-based decryption or appropriately-priced cryptographic hardware to detect the false traffic. Under such conditions, one means to protect the network elements from DoS attacks is to use packet-based filtering techniques based on relatively simple inspections of packets. As a result, for an ISP carrying large volumes of traffic, the ability to packet filter on the basis of port numbers is an important protection against DoS attacks, and a necessary adjunct to cryptographic strength in encapsulation.

Current practice in ISP operation is to use certain common filtering techniques to reduce the exposure to attacks from outside the ISP. To protect Internal BGP (IBGP) sessions, filters are applied at all borders to an ISP network. This removes all traffic destined for network elements' internal addresses (typically contained within a single prefix) and the BGP port number (179). If the BGP port number is found, packets from within an ISP are not forwarded from an internal interface to the BGP speaker's address (on which External BGP (EBGP) sessions are supported), or to a peer's EBGP address. Appropriate router design can limit the risk of compromise when a BGP peer fails to provide adequate filtering. The risk can be limited to the peering session on which filtering is not performed by the peer, or to the interface or line card on which the peering is supported. There is substantial motivation, and little effort is required, for ISPs to maintain such filters.

These operational practices can considerably raise the difficulty for an outsider to launch a DoS attack against an ISP. Prevented from injecting sufficient traffic from outside a network to effect a DoS attack, the attacker would have to undertake more difficult tasks, such as compromising the ISP network elements or undetected tapping into physical media.

5. References

5.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119, BCP 14, March 1997.
- [TCPMD5] Heffernan, A., "Protection of BGP Sessions via the TCP MD5 Signature Option", RFC 2385, August 1998.
- [RFC4271] Rekhter, Y., Li, T., and S. Hares, Eds., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.

5.2. Informative References

- [IPsec] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", RFC 2401, November 1998.
- [SBGP00] Kent, S., Lynn, C. and Seo, K., "Secure Border Gateway Protocol (Secure-BGP)", IEEE Journal on Selected Areas in Communications, Vol. 18, No. 4, April 2000, pp. 582-592.
- [SecCons] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, July 2003.
- [Smith96] Smith, B. and Garcia-Luna-Aceves, J.J., "Securing the Border Gateway Routing Protocol", Proc. Global Internet '96, London, UK, 20-21 November 1996.
- [RPSL] Villamizar, C., Alaettinoglu, C., Meyer, D., and S. Murphy, "Routing Policy System Security", RFC 2725, December 1999.
- [Watson04] Watson, P., "Slipping In The Window: TCP Reset Attacks", CanSecWest 2004, April 2004.

Author's Address

Sandra Murphy
Sparta, Inc.
7075 Samuel Morse Drive
Columbia, MD 21046

EMail: Sandy@tislabs.com

Full Copyright Statement

Copyright (C) The Internet Society (2006).

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 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.

Acknowledgement

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

