Network Working Group
INTERNET DRAFT
draft-murphy-bgp-secr-04.txt

Sandra Murphy
NAI Labs
November 2001

BGP Security Analysis

Status of this Memo

This document is an Internet-Draft and is subject to all provisions of Section 10 of RFC2026.

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/1id-abstracts.html

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

Abstract

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

This internet draft discusses some of the security issues with BGP routing data dissemination, and possible security solutions and the costs of those solutions. This internet draft does not discuss security issues with forwarding of packets.

Table of Contents

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 protocol was not designed with protection against deliberate or accidental errors causing disruptions of routing behavior.

We here discuss the vulnerabilities of BGP, based on the BGP RFC [1] and on [2]. Readers are expected to be familiar with the BGP RFC and the behavior of BGP. We propose several different security solutions to protect these vulnerabilities and discuss the benefits derived from each solution and its cost.

It is clear that the Internet is vulnerable to attack through its routing protocols. 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.

Under the present BGP design, cryptographic authentication of the peer-peer communication is not mandated. As a TCP/IP protocol, BGP is subject to all the TCP/IP attacks, like 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 to peer connection. Any break in the peer to peer communication has a ripple effect on routing that can be wide spread. Furthermore, outsider sources can also disrupt communications between BGP peers by breaking their TCP connection with spoofed RST packets. Outsider sources of bogus BGP information can reside anywhere in the world.

BGP speakers themselves can inject bogus routing information, either by masquerading as information from any other legitimate BGP speaker, or by distributing unauthentic routing information as themselves. Historically, misconfigured and faulty routers have been responsible for widespread disruptions in the Internet.

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 a path that does not follow the expected policy, or a path that will not forward the traffic. As a consequence, traffic to that network could be delayed by a longer than necessary path. 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 a view of 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 damage that might result from these attacks are:

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 it is not,

cut: some portion of the network believes that it has no route to some network that 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 become unstable so that convergence on a global forwarding state is not achieved, and

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

2. Vulnerabilities and Risks

The risks in BGP arise from three fundamental vulnerabilities:

no mechanism has been mandated that provides strong protection of the integrity, freshness and source authenticity of the messages in peer-peer BGP communications.

no mechanism has been specified to validate the authority of an AS to announce NLRI information.

no mechanism has been specified to ensure the authenticity of the AS_PATH announced by an AS.

There are four different BGP message types - OPEN, KEEPALIVE, NOTIFICATION, and UPDATE. A discussion of the vulnerabilities arising from each message and the ability of outsiders or BGP peers to exploit the vulnerabilities is contained in Appendix A. To summarize, outsiders can use bogus OPEN, KEEPALIVE, or NOTIFICATION messages to disrupt the BGP peer-peer connections and can use bogus UPDATE messages to disrupt routing. Outsiders can also disrupt BGP peer-peer connections by inserting bogus TCP RST packets. BGP peers themselves are permitted to break peer-peer connections at any time, and so they cannot be said to be issuing "bogus" OPEN, KEEPALIVE or NOTIFICATION messages. However, BGP peers can disrupt routing by issuing bogus UPDATE messages. In particular, bogus ATOMIC_AGGREGATE and AS_PATH attributes and bogus NLRI in UPDATE messages can disrupt routing.

3. Possible Protections

3.1. Threats from outsiders

Outsiders can be prevented from disrupting routing by providing cryptographic protection of the BGP peer-peer connection. The cryptography chosen should protect the source authenticity and integrity

of the message and should also protect against replay. As the protection is only of the peer-peer communications, asymmetric cryptography is not needed. Two different protection against spoofing in the peer-peer connection have been suggested:

IP level protection as defined by IPSEC [7]

TCP level protection as defined by [8]

Prevention of IP spoofing completely removes any risk associated with bogus OPEN, KEEPALIVE, or NOTIFICATION messages, as the only vulnerabilities from these messages come from outsiders. It also protects against all threats from bogus UPDATE messages and from bogus TCP messages arising from outsiders.

3.1.1. IP level protection

Protection as specified for the IPSEC [7] can be used to provide connectionless integrity, data origin authentication, and a anti-replay service.

3.1.2. TCP level protection

It is possible to protect the peer-peer connection by applying cryptographic protection at the TCP level to provide connectionless integrity and data origin authentication. This has been in use with some vendors for some time as specified in [8]. Note, however, that the protections specified in [8] were put in place some time ago. The IPSEC protections have advanced since that time. In particular, the protections of [8] use MD5 directly, where the IPSEC protections mandate the use of HMAC-MD5 because of recently publicized concerns of collisions in MD5. Also, [8] does not have the provisions for anti-reply found in IPSEC. The TCP sequence numbers do provide some protection against replay. But as some packets, notably a RST packet, need only be within the receive window to be accepted, the TCP sequence number protection is not complete. Finally, [8] has no provisions for multiple keys to be used in rekeying. As these are pairwise keys used for long-lived sessions, the inability to specify multiple keys may not cause operational difficulties.

Although the TCP level protection specified in [8] has deficiencies when compared with the protection of IPSEC [7], it is vastly preferable to a unprotected connection. If IPSEC is not available, then the TCP level protection of [8] should definitely be used. When IPSEC is available, IPSEC is preferable. One or the other must be used.

3.2. Threats from BGP peers

Protection of the communication between BGP peers does nothing to protect against errors introduced by the BGP speakers themselves. BGP speakers can introduce bogus routing information, e.g., invalid AS_PATHs, false origination announcements of NLRI, etc., at any time. Furthermore, detecting the BGP source of bogus information (if and when the bogus-ness is detected) can be difficult if not impossible.

There are several possible solutions to prevent a BGP speaker from inserting bogus information in its advertisements to its peers.

(1) Origination Protection: sign the originating AS.

(2) Origination and Adjacency Protection: sign the originating AS and predecessor information.

(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.

(4) Filtering: rely on a registry to verify the AS_PATH and NLRI originating AS.

The first three solutions are implemented within the protocol exchanges. The last solution is an operational protection and leaves the protocol messages unchanged.

3.3. Origination Protection

It would be beneficial to authenticate the AS that first advertises a route to a NLRI. This could be done by including a new path attribute which would contain the originating AS number and the originating AS's digital signature [4] of the AS number plus the NLRI advertised. A digital signature (as opposed to a message integrity check using a shared secret) is preferable because the number and identities of all eventual recipients are not known and because non-repudiation would be desired. If routing was disturbed by the presence of this advertisement, then the culprit could be determined.

If an infra-structure exists representing ownership of network addresses, then the owner as well as the originator could sign. This signature would indicate that the AS was authorized by the owner to originate the network. A BGP speaker receiving a protected origination could verify taht the origination was legitimate.

If routes are aggregated by a BGP speaker and the aggregated route advertised, then the idea of "originator" and "owner" become less useful. There might be several "originators" and "owners" represented among the aggregated routes. The AGGREGATOR field should be mandatory and an aggregating BGP speaker should append its signature of the AGGREGATOR field and the aggregated NLRI. Unfortunately, aggregation prevents identification of the specific culprit should it be discovered that a network is being originated incorrectly.

This protection does not ensure that the AS_PATH is authentic or current.

3.4. Origination and Adjacency Protection

Reference [3] suggests several different types of cryptographic protection of BGP. The suggested protection against possibly faulty BGP speakers introduces some link state topology information (as in OSPF [5]) that can be used to verify AS_PATHs.

To obviate the need to trust BGP speakers regarding NLRI information not specific to their own AS, [3] suggests adding the following information to the UPDATE message in new path attributes:

- a sequence number for the UPDATE

- the AS originating the information (either the aggregator or the originator of a direct route) and

- the "predecessor" of the originating AS (i.e., the neighbor to which the NLRI is first advertised).

This information is digitally signed by the BGP originator along with the NLRI, the ATOMIC_AGGREGATE, the ORIGIN, and the AGGREGATOR fields of the UPDATE message. The signature and the predecessor information must be included as the route in the UPDATE message propagates across the network, i.e., it is transitive.

This information distributes a bit of link state topology information, concerning just the last hop before the destination network's AS, into the usual BGP distance vector (some say "path vector") protocol. Each BGP speaker will propogate this "predecessor" information. Any BGP speaker could use the signed predecessor information received to verify that each of the adjacencies represented in an AS_PATH is legitimate. That is, when a segment of an AS_PATH is a sequence, each adjacent pair in the sequence could be verified to correspond to a received signed

predecessor tuple.

The protection provided by the signed predecessor information becomes more difficult to use past an aggregation point where a BGP speaker advertises a less specific route which includes the originated NLRI. In particular, the rules for verifying an AS_PATH containing a segment that is a set would be either very lenient or very complex.

While this predecessor information assures that each adjacency in a sequence of an AS_PATH is valid, it does not ensure that the AS_PATH as a whole is valid. Each AS's decision regarding routes it will advertise and traffic it will transit is individual and totally unconstrained. The fact that a valid path of ASs exists to a destination does not ensure that the corresponding AS_PATH is valid. This mechanism also does not assure that any information which comes from one router alone (LOCAL_PREF, NEXT_HOP, AGGREGATOR, etc.) is accurate. A router, then, can still falsely announce that its neighbor should be forwarded the traffic for an NLRI.

3.5. Sign Originating AS and AS_PATH

A protection against possibly faulty BGP speakers that does provide some assurance that the AS_PATH is valid is described in [9]. The protection requires digital signatures of nested prefixes of the AS_PATH, carried in a transitive path attribute.

Each BGP speaker would receive signed route information (including the AS_PATH, the ATOMIC_AGGREGATE attribute and NLRI) from its peer. The signature represents permission from the peer for the speaker to advertise the route. After making its routing decision, the BGP speaker would augment the chosen AS_PATH with its own AS and sign the resulting route (NLRI + AS_PATH) and ATOMIC_AGGREGATE. The BGP speaker would pass to its peers the augmented AS_PATH, its signature, and the signature it received from its peer which covers the received AS_PATH. The neighbor receiving this information can verify that the received AS_PATH was indeed constructed from an authorized path, by verifying the signature of the BGP speaker's peer over the tail of the received AS_PATH. The BGP speaker's signature will be passed on by the neighbor to provide similar assurance that it constructed its advertised AS_PATH legitimately.

A BGP speaker could snip out a suffix of the data it received as well as the associated signatures and pass those on as the proof that its AS_PATH was based on reality. To prevent this, the signatures generated must cover not only the route information but the intended receiver as

well.

This procedure as described protects against one consecutive faulty router in the path. If it is desired to protect against more than one possibly faulty routers in the path, then the procedure can be nested arbitrarily. To protect against K consecutive faulty routers, each router would receive signed AS_PATH information from its neighbor along with K signatures of the preceding K BGP speakers in the path, each successive signature covering a shorter suffix of the AS_PATH. It would pass on a newly constructed AS_PATH along with its own signature, its neighbor's signature and K-1 of the included nested signatures.

For example, consider a case where it was decided to protect against two faulty BGP speakers. Suppose AS1 receives an AS_PATH from AS2 of 'AS2 AS3 AS4 ... ASk'. (In this discussion as well as the next, ATOMIC_AGGREGATE would be included in the signature, but is omitted for brevity.) Then AS1 should also receive AS2's signature of "AS1, 'AS2 AS3 ... ASk', NLRI" as well as AS3's signature of "AS2, 'AS3 ... ASk', NLRI" and AS4's signature of "AS3, 'AS4 ... ASk', NLRI". AS1 would verify the authenticity of AS2's route by verifying the authorizing signatures from AS3 and AS4.

AS1 uses AS2's signature in authorizing its own announcements. AS1 would compute a path as 'AS1 AS2 AS3 ... ASk', and pass to its neighbor AS0 this path along with its signature of "AS0, 'AS1 AS2 AS3 ... ASk', NLRI" and the authorizing signatures of AS2 and AS3.

Because the intended recipient is included in the signature, each BGP speaker with N peers generates N signatures for each route announced, one for each peer.

This discussion is predicated on the use of asymmetric cryptography. Each autonomous system would be required to possess an asymmetric key pair. The public key would have to be accessible to all recipients of the digital signature. The private key would have to be available to all BGP speakers in the autonomous system, but protected from exposure.

3.5.1. Remaining Issues

This scheme becomes more complicated when one of the BGP speakers performs aggregation of a set of routes. To assure recipients of the validity of the aggregated route, it would be necessary to pass on the text and signatures of each of the aggregated component routes. This means an enormous increase in transmission bandwidth at each aggregation point and a similar increase in verification time at each verification

point's peers. This cost would not have to be passed on to further neighbors further than K (the nesting level of signatures) hops away, but it does violate the spirit of aggregation. Alternatively, an aggregation point could be treated as another type of origination point, and signatures and verification would stop at that point. Unfortunately, that provides a mechanism for malicious BGP peers to announce bogus routes, simply by claiming to have aggregated the information. Aggregation also prevents identification of the specific culprit should it be discovered that a network is being originated in error.

Neither this scheme nor the Origination and Adjacency Protection assures that the received route is the best route the peer could have computed. In particular, it does not guard against a peer that does not announce the best result of its decision process - for example, a peer that replays previous updates instead of forwarding a withdrawal, or does not change its announcements when it receives withdrawn routes, or replays withdrawal information after a route is reestablished, or replays an update after having forwarded a withdrawal. These are a matter for the internal correct operation of the router and cannot be precluded by authentication or authorization protection. This mechanism also does not assure that any information which comes from one router alone (LOCAL_PREF, NEXT_HOP, AGGREGATOR, etc.) is accurate. With these fields, a router can still falsely announce that its neighbor should be forwarded the traffic for an NLRI.

The verification process chooses the key to use based on the AS's mentioned in the AS_PATH. If implementations do not verify that the peer BGP speaker prepended its own AS to the AS_PATH, i.e. if a BGP speaker might pass on an AS_PATH in which it's own AS is not the first in the PATH, then the AS's mentioned in the AS_PATH are not necessarily the AS's that produced the signature. In that case, this verification process could be using the wrong key. This signature scheme would have to be complicated by requiring either

- that the sender's AS be included in the UPDATE message and the signature (so that the verification process could find the appropriate key for the signature)

- or that each sender know the private key of the first AS in the AS_PATH (so that the signature and verification processes would be using the same key).

Sharing keys between AS's makes those AS's indistinguishable to the cryptography; the second alternative design should only be chosen with

that caution clearly understood.

3.6. Filtering

The Internet registries can provide data that will help to assure that AS_PATHS and NLRI origination data are authorized. If the data registered in the Internet registries is available, then an NLRI origination can be verified against the registry. Also, peering information in the registry can be used to verify that the AS_PATH as a whole was valid.

The assurance provided by this protection would rely on the registry data being:

complete: AS's that do not register their peering relationships or assigned networks would hamper the ability to protect the BGP data.

accurate: AS's must register peering relationships that exactly match their policies. If the peering relationships are described too broadly, bogus routes will pass the filter. If the peering relationships are described too narrowly, legitimate routes will fail.

timely: Additions and changes to the registry data must be communicated to the AS's in a timely manner.

secure: The registries must be known to ensure the authenticity, integrity and authorization of provided information while in storage and in transmission to an AS.

The assurance provided by this protection relies on the openness of the data recorded in the registries. To be truly useful, each autonomous system's policy would have to be recorded in the registry, in order that the AS_PATH as a whole can be verified to be valid. Information about policy, however, can be sensitive to an autonomous system and not openly available to every other autonomous system. Any restrictions on the availability of information stored in the registry will restrict its applicability as a protection mechanism.

Filtering based on the registries can also not ensure the authenticity of the received routes. The registries currently contain permitted routes, not necessarily the current forwarding routes. If, for example, an AS will accept the same network from multiple peer, with one route to serve as a backup, filtering based on the registry would not be able to distinguish which is the route currently in use.

4. Security Costs

Choosing the protection to apply in any situation depends on the perception of the risk of attack, of the damage that can result, of the benefits derived from providing the protection, and of the cost of providing the protection. This section discusses the cost of each of the protection options mentioned above.

4.1. Protecting the Peer-Peer Link

The cost of this protection is the processing required for the cryptography and the need to establish and manage the cryptographic keys. The cryptography need not be computationally expensive; HMAC or similar algorithms can be used. Shared secret keys are adequate for this protection, as the protection applies only to the communication between peers, so the key management cost is low. A separate key must be used for each BGP peering. Using one key for multiple peerings even at a peering point reduces the level of protection that is provided.

This level of protection is low cost and protects against the vast number of outsiders who pose a threat.

4.2. Origination Protection

The cost of signing the originating AS of each route is the cost of the processing required for the cryptography -- generating and verifying digital signatures -- as well as the need to establish and manage the cryptographic keys. For digital signatures, establishing and managing the cryptographic keys means providing a public key infrastructure to generate an asymmetric key pair for each autonomous system and to distribute and maintain certifications of the public keys associated with each autonomous system.

While the key distribution and signature generation and verification provides for authentication and integrity of the messages, there is also a need to ensure the authorization of the origination. This requires another infrastructure that would provide certifications of the authority of an AS to originate a route to a network.

The infrastructures required for authentication of AS messages and for authorization of origination information might be colocated. The establishment and secure operation of the security infrastructure as well as the cost of secure communication with the infrastructure are additional costs of this technique.

4.3. Origination and Adjacency Protection

This scheme requires the same public key infrastructure and origination infrastructure as is needed for Origination Protection It also requires that each adjacency for a BGP speaker be signed (the "predecessor" information) and transmitted along with an AS_PATH. Essentially, each BGP speaker must announce its peers, something that does not currently occur in the BGP protocol.

Each BGP speaker must compute one signature for each NLRI in each UPDATE message transmitted and must verify one signature for each NLRI in each UPDATE message received. Each UPDATE message must be separately signed because the mechanism described in [3] includes a sequence number, the Withdrawn Routes and the Unfeasible Route Length fields, so the information to be signed changes with each message. (These fields are protected from outsiders by the peer-peer communication protection and so do not need to be digitally signed. If only the NLRI, ATOMIC_AGGREGATE and predecessor information were signed each time, then the signature might not have to be computed with each new UPDATE message, i.e., AS_PATH changes would not induce new signature computations.)

The signed predecessor information in each route must be stored by the recipient indefinitely. Each route received must be verified by comparison to the store of predecessor information previously received in UPDATE messages from all AS sources.

4.4. Origination and AS_PATH Protection

The operational cost of this scheme is described in detail in [10].

This scheme requires the same public key infrastructure and origination infrastructure as is needed for Origination Protection.

For each UPDATE message received, K signatures (where K is the level of protection desired from consecutive faulty routers) must be verified.

For each UPDATE message transmitted, one signature must be computed for each NLRI per recipient. As discussed before, prevention of cut and paste attacks requires that the signature include the recipient, so computing one signature per AS_PATH and NLRI announced is insufficient.

The signatures contained in an UPDATE must be retained for all received routes in case the route is ever chosen and announced to the peers. This increases the size of the routing tables.

For efficiency in situations where route flapping is occuring, withdrawn AS_PATHs and their signatures could be retained. This would mean that new announcements of flapped routes still retained would not require new signature verifications.

4.5. Rely on Registries

The cost of relying on registries would vary considerably depending on the protection provided to the information in storage and in transit. Any latency, above that caused by the use of cryptography, would depend on the mechanism used to transmit the registry information (e.g., anything from frequent complete download to real-time query and response).

5. Security Considerations

This entire memo is about security considerations.

6. References

[1] Y. Rekhter and T. Li, "A Border Gateway Protocol 4 (BGP-4)", RFC1771, March 1995.

[2] Y. Rekhter and T. Li, "A Border Gateway Protocol 4 (BGP-4)", work in progress, November 2001. available as <<draft-ietf-idr-bgp4-15.txt>> at Internet-Draft shadow sites.

[2] B. Smith and J.J. Garcia-Luna-Aceves, "Securing the Border Gateway Routing Protocol", Proc. Global Internet'96, London, UK, 20-21 November 1996.

[3] Bruce Schneier. Applied Cryptography Protocols, Algorithms, and Source Code in C, John Wiley & Sons, Inc 1994.

[4] John Moy, "OSPF Version 2", RFC 1583, March 1994.

[5] C. Villamizar, C. Alaettinoglu, D. Meyer, S. Murphy and C. Orange, "Routing Policy System Security", RFC 2725, December, 1999.

[7] S. Kent and R.Atkinson, "Security Architecture for the Internet Protocol", RFC2401, November 1998.

[8] A. Heffernan, "Protection of BGP Sessions via the TCP MD5 Signature Option", RFC2385, August 1998.

[10] S.Kent, C.Lynn, and K. Seo, "Secure Border Gateway Protocol (Secure-BGP)", IEEE Journal on Selected Areas in Communications, Vol. 18, No. 4, April 2000, pp. 582-592.

[10] S. Kent, C. Lynn, J. Mikkelson, and K. Seo, "Secure Border Gateway Protocol (S-BGP) -- Real World Performance and Deployment Issues", Proceedings of the IEEE Network and Distributed System Security Symposium (NDSS 2000), February 2000.

7. Author's Address

Sandra Murphy
Network Associates, Inc.
NAI Labs
3060 Washington Road
Glenwood, MD 21738
EMail: Sandy@tislabs.com

A. Appendix A - Vulnerabilities and Risks

Each message introduces certain different vulnerabilities and risks:

A.1. OPEN

Because receipt of a new OPEN message in the Established state will cause the close of the BGP peering session and thereby induce the release of all resources and deletion of all associated routes, the ability to spoof this message can lead to a severe disruption of routing.

A.2. KEEPALIVE

Receipt of a KEEPALIVE message when the peering connection is in the OpenSent state can lead to a failure to establish a connection. The ability to spoof this message is a vulnerability. 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.

A.3. NOTIFICATION

Receipt of a NOTIFICATION message will cause the close of the BGP peering session and thereby induce the release of all resources and deletion of all associated routes. Therefore, the ability to spoof this message can lead to a severe disruption of routing.

A.4. UPDATE

The Update message carries the routing information. The ability to spoof any part of this message can lead to a disruption of routing.

A.4.1. 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. As a true BGP speaker is always able to close a connection at any time, this vulnerability represents an additional risk only when the source is a non-BGP speaker, i.e., it presents no additional risk from BGP sources.

A.4.2. 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.

A.4.3. Path Attributes

The path attributes present many different vulnerabilities and risks.

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) so they did not 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 a 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. As a true BGP speaker is always 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.

ORIGIN

This field indicates whether the information was learned from IGP or EGP information. If the route is used in inter-AS multicast routing, a value of INCOMPLETE may be used. This field is not used in making routing decisions, so there are no vulnerabilities arising from this field, either from BGP peers or outsiders.

AS_PATH

A BGP peer or outsider could announce an AS_PATH that was not accurate for the associated NLRI. Forwarding for the NLRI associated with the AS_PATH could potentially be induced to follow a sub-optimal path, a path that did not follow some intended policy, or even a path that would not forward the traffic. If the path would not forward the traffic, the NLRI would become unreachable for that portion of the Internet that accepted the false path. If much traffic is mis-directed, some routers and transit networks along the announced route could become flooded with the mis-directed traffic.

It is not clear how far an inaccurate AS_PATH could deviate from the true AS_PATH. It may be that the first AS in the AS_PATH, at least, must be a legal hop. The RFC states that a BGP speaker prepends its own AS to an AS_PATH before announcing it to a neighbor. If the BGP speaker must prepend its own AS, then a BGP speaker that produced a bogus AS_PATH could end up receiving the traffic for the associated NLRI. This could be desirable if the error was deliberate and the intent was to receive traffic that would not otherwise be received. Receiving the mis-routed traffic could be undesirable for the faulty BGP speaker if it were not prepared to handle the extra (mis-routed) traffic. So, requiring a BGP peer to prepend its own AS to the AS_PATH, might encourage or discourage it from inventing an arbitrary AS_PATH, depending on its resources and intent.

If BGP peers need not prepend its own AS (or implementations do not ensure that this is done), then a malicious BGP peer could announce a path that begins with the AS of any BGP speaker with little impact on itself.

Whether such an arbitrary AS_PATH is a vulnerability would depend on whether BGP implementations check the AS_PATH (to see if the first AS is the neighbor) and would catch the error. If there are legitimate situations in which a BGP speaker could pass an AS_PATH to a neighbor without putting its own AS at the head of the AS_PATH, then there is no way for implementations to detect totally bogus AS_PATHs.

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. An 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 on 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 mis-directed traffic.

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 modifying this field could disrupt the forwarding of traffic between the two AS's.

In the case that 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 addresss. This affords the opportunity to direct traffic at a router that may not be able to continue forwarding the traffic. It is unclear whether this would also require the advertising BGP speaker to construct an AS_PATH mentioning the NEXT_HOP external peer's AS.

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 AS's. Consequently, this field is only used in making routing decisions internal to one AS. Modifying this field, whether by an outsider or an BGP peer, could influence routing within an AS to be sub-optimal, but the effect should be limited in scope.

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. As BGP peers are free to choose the LOCAL_PREF as they wish, modification of this field is a vulnerability with respect to outsiders only.

ATOMIC_AGGREGATE

The ATOMIC_AGGREGATE field indicates that an AS somewhere along the way has received a more specific and a less specific route to the NLRI and installed the aggregated route. This route cannot be de-aggregated because it is not certain that the route to more specific prefixes will follow the listed AS_PATH. Consequently, 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.

AGGREGATOR

This field may be included by a BGP speaker who has computed the routes represented in the UPDATE message from aggregation of 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.

A.4.4. 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.