

Considerations in Validating the Path in BGP

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.

IESG Note

After consultation with the RPSEC WG, the IESG thinks that this work is related to IETF work done in WG RPSEC, but this does not prevent publishing.

This RFC is not a candidate for any level of Internet Standard. The IETF disclaims any knowledge of the fitness of this RFC for any purpose and in particular notes that the decision to publish is not based on IETF review for such things as security, congestion control, or inappropriate interaction with deployed protocols. The RFC Editor has chosen to publish this document at its discretion. Readers of this document should exercise caution in evaluating its value for implementation and deployment. See RFC 3932 for more information.

Abstract

This document examines the implications of hop-by-hop forwarding, route aggregation, and route filtering on the concept of validation within a BGP Autonomous System (AS) Path.

1. Background

A good deal of thought has gone into, and is currently being given to, validating the path to a destination advertised by BGP. The purpose of this work is to explain the issues in validating a BGP AS Path, in the expectation that it will help in the evaluation of schemes seeking to improve path validation. The first section defines at least some of the types of questions a BGP speaker receiving an update from a peer not in the local autonomous system (AS) could ask about the information within the routing update. The following sections examine the answers to these questions in consideration of specific deployments of BGP.

The examples given in this document are intended to distill deployments down to their most critical components, making the examples easier to understand and consider. In many situations, the specific path taken in the example may not be relevant, but that does not nullify the principles considered in each example. It has been suggested that these examples are "red herrings", because they do not illustrate actual problems with specific policies. On the contrary, these examples are powerful because they are simple. Any topology in which one of these example topologies is a subtopology will exhibit the characteristics explained in this document. Rather than focusing on a specific topology, then dismissing that single topology as a "corner case", this document shows the basic issues with assertions about the AS Path attribute within BGP. These generalized issues can then be applied to more specific cases.

With the heightened interest in network security, the security of the information carried within routing systems running BGP, as described in [RFC4271], is being looked at with great interest. While there are techniques available for securing the relationship between two devices exchanging routing protocol information, such as [BGP-MD5], these techniques do not ensure various aspects of the information carried within routing protocols are valid or authorized.

The following small internetwork is used to examine the concepts of validity and authorization within this document, providing definitions used through the remainder of the document.

10.1.1.0/24--(AS65000)---(AS65001)--(AS65002)

Assume a BGP speaker in AS65002 has received an advertisement for 10.1.1.0/24 from a BGP speaker in AS65001, with an AS Path of {65000, 65001}.

1.1. Is the Originating AS Authorized to Advertise Reachability to the Destination?

The most obvious question the receiving BGP speaker can ask about this advertisement is whether or not the originating AS, in this case AS65000, is authorized to advertise the prefix contained within the advertisement, in this case 10.1.1.0/24. Whether or not a BGP speaker receiving a route to 10.1.1.0/24 originating in AS65000 can verify that AS65000 is, indeed, authorized to advertise 10.1.1.0/24 is outside the scope of this document.

1.2. Is the Path Contained in the Advertised Routing Information Valid?

If a BGP speaker receives an advertisement from a peer outside the local autonomous system (AS), the peer sending the update has a path to the destination prefix in the update. Specifically, are the autonomous systems within the internetwork connected in such a way that the receiver, following the AS Path listed in the BGP update itself, can reach the originating AS listed in the received AS Path? Within this document, this is called path validation.

Path validation, in the context of this small internetwork, asserts that when a BGP speaker in AS65002 receives an advertisement from a BGP speaker in AS65001 with the AS Path {65000, 65001}, the speaker can assume that AS65001 is attached to the local AS, and that AS65001 is also attached to AS65000.

1.3. Is the Advertisement Authorized?

There are at least three senses in which the readvertisement of a received advertisement can be authorized in BGP:

- o The transmitter is authorized to advertise the specific routing information contained in the route. This treats the routing information as a single, atomic unit, regardless of the information the route actually contains. A route to 10.1.1.0/24 and another route to 10.1.0.0/16 are considered completely different advertisements of routing information, so an AS may be authorized to advertise 10.1.0.0/16 without regard to its authorization to advertise 10.1.1.0/24, since these are two separate routes. This is called route authorization throughout this document.
- o The transmitter is authorized to advertise the specific reachable destination(s) contained in the route. This treats the routing information as a set of destinations. 10.1.1.0/24 is contained within 10.1.0.0/16, and authorization to advertise the latter implies authorization to advertise the former. This is called reachability authorization throughout this document.
- o The transmitter is authorized to transit traffic to the destinations contained within the route. This ties the concepts of the route to what the route is used for. If a BGP speaker is advertising reachability to 10.1.1.0/24, it is authorized to transit traffic to all reachable destinations within 10.1.1.0/24 along the path advertised. This is called transit authorization throughout this document.

There is considerable tension between these three definitions of authorization; much of this document is built around exploring the relationships between these different types of authorization, and how they may, or may not, work in various internetworks. One of the conclusions reached by this document is that route authorization, reachability authorization, and transit authorization are often at odds with each other. Showing one type of authorization to be true does not show any other types of authorization to be true, and route authorization is of questionable value because of the intertwined nature of these three types of authorization in a routing system.

1.4. Will Traffic Forwarded to an Advertising Speaker Follow the Described AS Path?

If a BGP speaker receives an advertisement from a peer not in the local AS, will traffic forwarded to the peer advertising the update follow the path described in the AS Path? In this document, this is called forwarding consistency.

In terms of the small example internetwork, if a BGP speaker in AS65002 receives an advertisement from a peer in AS65001 for the destination 10.1.1.0/24, with an AS Path {65000, 65001}, will traffic forwarded to the BGP speaker in AS65001 actually be forwarded through routers within AS65001, then AS65000, to reach its destination?

1.5. Is a Withdrawing Speaker Authorized to Withdraw the Routing Information?

If an advertisement is withdrawn, the withdrawing BGP peer was originally advertising the prefix (or was authorized to advertise the prefix). This assertion is out of the scope of this document.

2. Analysis

To begin, we review some of the concepts of routing, since we need to keep these concepts fixed firmly in place while we examine these questions. After this, four examples will be undertaken with BGP to show the tension between the various types of authorization in a path vector routing system.

2.1. A Short Analysis of Routing

Routing protocols are designed, in short, to discover a set of loop-free paths to each reachable destination within a network (or internetwork). The loop-free path chosen to reach a specific destination may not be the shortest path, and it may not always be

the "best" path (depending on the definition of "best"), but it should always be a loop-free path, otherwise the routing protocol has failed.

This sheds some light on the purpose of the path included in a path vector protocol's routing update: the path is there to prove the path is loop free, rather than to provide any other information. While Dijkstra's Sender Policy Framework (SPF) and the Diffusing Update Algorithm (DUAL) both base their loop-free path calculations on the cost of a path, path vector protocols, such as BGP, prove a path is loop free by carrying a list of nodes the advertisement itself has traversed. BGP specifically uses an AS Path-based mechanism to prove loop freeness for any given update so each AS along the path may implement local policy without risking a loop in the routing system caused by competing metrics.

We need to keep this principle in mind when considering the use of the path carried in a path-vector protocol, and its use by a receiving BGP speaker for setting policy or gauging the route's security level.

2.2. First Example: Manual Intervention in the Path Choice

In the small network:

```

          +---(AS65002)---+
(AS65000)--(AS65001)      (AS65004)--10.1.1.0/24
          +---(AS65003)---+

```

A BGP speaker in AS65000 may receive an advertisement from a peer that 10.1.1.0/24 is reachable along the path {65004, 65002, 65001}. Based on this information, the BGP speaker may forward packets to its peer in AS65001, expecting them to take the path described. However, within AS65001, the network administrator may have configured a static route making the next hop to 10.1.1.0/24 the edge router with AS65003.

It's useful to note that while [RFC4271] states: "...we assume that a BGP speaker advertises to its peers only those routes that it itself uses...", the definition of the term "use" is rather loose in all known widely deployed BGP implementations. Rather than meaning: "A BGP speaker will only advertise destinations the BGP process on the speaker has installed in the routing table", it generally means: "A BGP speaker will only advertise destinations that the local routing table has a matching route installed for, no matter what process on the local router installed the route". A naive reaction may be to simply change the BGP specifications and all existing implementations so a BGP speaker will only advertise a route to a

peer if the BGP process on the router actually installed the route in the local routing table. This, however, ignores the complex interactions between interior and exterior gateway protocols, which most often run on the same device, and the complexities of route origination.

Although this is an "extreme" example, since we can hardly claim the information within the routing protocol is actually insufficient, we still find this example instructive in light of the questions outlined in Section 1:

- o Is the AS Path valid? The AS Path the receiving BGP speaker in AS65000 receives from its peer in AS65001, {65004, 65002, 65001}, does exist, and is valid.
- o Is the advertisement authorized? Since we have no knowledge of any autonomous system level policy within this network, we have no way of answering this question. We can assume that AS65001 has both route and reachability authorization, but it is difficult to see how transit authorization can be accomplished in this situation. Even if the BGP speaker in AS65000 could verify AS65001 is authorized to transit AS65002 to reach 10.1.1.0/24, this implies nothing about the authorization to transit traffic through the path traffic will actually take, which is through AS65003.
- o Is the AS Path consistent with the forwarding path (does forwarding consistency exist)? No, the advertised AS Path is {65004, 65002, 65001}, while the actual path is {65004, 65003, 65001}.

From this example, we can see forwarding consistency is not possible to validate in a deployed network; just because a BGP speaker advertises a specific path to reach a given destination, there is no reason to assume traffic forwarded to the speaker will actually follow the path advertised. In fact, we can reason this from the nature of packet-forwarding networks; each router along a path makes a completely separate decision about where to forward received traffic. Any router along the path could actually change the path due to network conditions without notifying, in any way, upstream routers. Therefore, at any given time, an upstream router may be forwarding traffic along a path that no longer exists, or is no longer optimal, and downstream routers could be correcting the forwarding decision by placing the traffic on another available path unknown to the upstream router.

As a corollary, we can see that authorizing transit through a specific AS Path is not possible, either. If the source of a

specific flow cannot know what path the traffic within that flow will take to reach the destination, then there is no meaningful sense in which downstream routers can authorize the source to use available paths for transiting traffic.

2.3. Second Example: An Unintended Reachable Destination

In this internetwork, we assume a single policy: the system administrator of AS65000 would not like the destination 10.1.1.0/24 to be reachable from any autonomous system beyond AS65001. In other words, 10.1.1.0/24 should be reachable within AS65001, but not to autonomous systems connected to AS65001, such as AS65002.

```
10.1.1.0/24---(AS65000)---(AS65001)---(AS65002)
```

The system administrator can implement this policy by causing BGP speakers within AS65000 to advertise 10.1.1.0/24 to peers within AS65001 with a signal that the BGP speakers in AS65001 should not readvertise the reachability of this routing information. For example, BGP speakers in AS65000 could advertise the route to 10.1.1.0/24 with the NO_ADVERTISE community attached, as described in [RFC4271]. If the BGP speakers in AS65001 are configured to respond to this community (and we assume they are for the purposes of this document) correctly, they should accept this advertisement, but not readvertise reachability to 10.1.1.0/24 into AS65002.

However, unknown to the system administrator of AS65000, AS65001 is actually advertising a default route to AS65002 with an AS Path of {65001}, and not a full routing table. If some host within AS65002, then, originates a packet destined to 10.1.1.1, what will happen? The packet will be routed according to the default route from AS65002 into AS65001. Routers within AS65001 will forward the packet along the 10.1.1.0/24 route, eventually forwarding the traffic into AS65000.

- o Is the AS Path valid? This is a difficult question to answer, since there are actually two different advertisements in the example. From the perspective of the BGP speaker in AS65002 receiving a default route in an advertisement from a peer in AS65001, the AS Path to the default route is valid. However, there is no route to 10.1.1.0/24 for the BGP speaker in AS65002 to examine for validity, so the question appears to be out of scope for this example.
- o Is the AS Path consistent with the forwarding path (is there forwarding consistency)? From the perspective of a BGP speaker in AS65002, traffic forwarded to AS65001 towards a destination within 10.1.1.0/24 is going to actually terminate within AS65001, since

that is the entire AS Path for the default route. However, this traffic actually transits AS65001 towards AS65000. Therefore, forwarding consistency does not exist in this example, in which we are dealing with a case of aggregation, and as Section 9.1.4 of [RFC4271], in reference to aggregated routing information, states: "Forwarding along such a route does not guarantee that IP packets will actually traverse only ASes listed in the AS_PATH attribute of the route".

2.3.1. Advertisement Authorization

Is the advertisement authorized? This example highlights the tension between the three different types of authorization. The three following sections discuss issues with different advertisements AS65001 may send to AS65002.

2.3.1.1. Valid Unauthorized Aggregates

The first issue that comes up in this example is the case where there is no expectation of authorization for aggregation. The most common example of this is the advertising and accepting of the default route (0/0). This is a common practice typically done by agreement between the two parties. Obviously, there is not an authorization process for such an advertisement. This advertisement may extend reachability beyond the originator's intention (along the lines of the previous example). It may cause packets to take paths not known by the sender (since the path on 0/0 is simply the advertising AS). It may violate other security constraints. This places limits on the power and applicability of efforts to secure the AS path and AS policies. It does not vitiate the underlying value in such efforts.

2.3.1.2. Owner Aggregation

In the current instantiation of IP address allocation, an AS may receive authorization to advertise 10.1.0.0/16, for instance, and may authorize some other party to use (or own) 10.1.1.0/24, a subblock of the space authorized. This is called a suballocation.

For instance, in this example, if AS65001 were authorized to originate 10.1.0.0/16, it could advertise 10.1.0.0/16 towards AS65002, rather than a default route. Assume there is some form of authorization mechanism AS65002 can consult to verify AS65001 is authorized to advertise 10.1.0.0/16. In this case, AS65002 could examine the authorization of AS65001 to originate 10.1.0.0/16, and assume that if AS65002 is authorized to advertise 10.1.0.0/16, it is also authorized to transit traffic towards every possible subblock of (every destination within) 10.1.0.0/16. To put this in more distinct terms:

- o AS65002 verifies route authorization by examining the authorization of AS65001 to advertise 10.1.0.0/16.
- o AS65002 assumes destination authorization, that AS65001 has the authorization to advertise every possible subblock of 10.1.0.0/16, because AS65001 is authorized to advertise 10.1.0.0/16.
- o AS65002 assumes transit authorization, that AS65001 has the authorization to transit traffic to every possible subblock of 10.1.0.0/16, because AS65001 is authorized to advertise 10.1.0.0/16.

From the example given, however, it is obvious route authorization does not equal destination or transit authorization. While AS65001 does have route authorization to advertise 10.1.0.0/16, it does not have destination authorization to advertise 10.1.1.0/24, nor transit authorization for destinations with 10.1.1.0/24.

The naive reply to this would be to simply state that destination and transit authorization should not be assumed from route authorization. However, the problem is not that simple to resolve. The assumption of destination authorization and transit authorization are not decisions AS65002 actually makes; they are embedded in the way the routing system works. The route itself, within the design of routing, carries these implications.

Why does routing intertwine these three types of authorization? Most simply, because AS65002 does not have any information about subblocks that are part of 10.1.0.0/16. There is no way for AS65002 to check for destination and transit authorization because this information is removed from the system altogether. In order to show destination and transit authorization, this information must be reinjected into the routing system in some way.

For instance, considering destination authorization alone, it is possible to envision a system where AS65001, when suballocating part of 10.1.0.0/16 to the originator, must also obtain permission to continue advertising the original address block as an aggregate, to attempt to resolve this problem. However, this raises some other issues:

- o If the originator did not agree to AS65001 advertising an aggregate containing 10.1.1.0/24, then AS65001 would be forced to advertise some collection of advertisements not containing the suballocated block.
- o If AS65001 actually does obtain permission to advertise the aggregate, we must find some way to provide AS65002 with

information about this authorization for each possible subblock of 10.1.0.0/16.

In other words, either AS65002 must receive the actual routes for each suballocation of 10.1.0.0/16, or it must receive some form of authorization allowing AS65001 to advertise each suballocation of 10.1.0.0/16. This appears to defeat the purpose of aggregation, rendering routing systems much less scalable than current design allows. Further, this does not resolve the issue of how AS65002 would actually know what all the suballocations of 10.1.0.0/16 actually are. Some possible solutions could be:

- o The suballocator must advertise all suballocations. This could potentially expose business relationships and patterns that many large commercial providers do not want to expose, and degrades the hierarchical nature of suballocation altogether.
- o The IP address space must be reconstructed so everyone using IP address space will know, based on the construction of the IP address space, what subblocks exist. For instance, the longest allowed subblock could be set at a /24, and authorization must be available for every possible /24 in the address space, either for origination, or as part of an aggregate. This sort of solution would be an extreme burden on the routing system.

2.3.1.3. Proxy Aggregation

It is also possible for AS65001 to have some form of agreement with AS65002 to aggregate blocks of address space it does not own towards AS65002. This might be done to reduce the burden on BGP speakers within AS65002. This is called proxy aggregation. While proxy aggregation is rare, it is useful to examine the result of agreed upon proxy aggregation in this situation.

Assume AS65001 has an advertisement for 10.1.0.0/24 from some unknown source, and decides to advertise both 10.1.0.0/24 and 10.1.1.0/24 as 10.1.0.0/23 to AS65002. If there exists an agreement for AS65001 to advertise proxy aggregates to AS65002, the latter will accept the advertisement regardless of any attached authorization to advertise. This may represent a security risk for AS65002, but it might be seen as an acceptable risk considering the trade-offs, from AS65002's perspective.

The problem is, however, this also impacts the policies of AS65000, which is originating one of the two routes being aggregated by AS65001. There is no way for AS65002 to know about this policy, nor to react to it, and there is actually no incentive for AS65002 to react to a security threat posed to AS65000, which it has no direct

relationship with. There doesn't appear to be any immediately available solution to this problem, other than to disallow proxy aggregation, even between two cooperating autonomous systems.

2.3.2. Implications

The basic problem is that AS65002 assumes when AS65001 advertises an authorized route containing 10.1.1.0/24, either through a valid unauthorized aggregate, an owner aggregated route, or a proxy aggregation, AS65001 also has destination authorization for the subblock 10.1.1.0/24, and transit authorization for destinations within 10.1.1.0/24. These are, in fact, invalid assumptions, but they are tied to the way routing actually works. This shows the value of route authorization is questionable, unless there is some way to untie destination and transit authorization from route advertisements, which are deeply intertwined today.

2.4. Third Example: Following a Specific Path

This example is slightly more complex than the last two. Given the following small network, assume that A and D have a mutually agreed upon policy of not allowing traffic to transit B to reach destinations within 10.1.1.0/25.

```

10.1.1.0/25--A---B---C---D
           |           |   |
           E-----F---G

```

Assume the following:

- o A advertises 10.1.1.0/25 to B, and 10.1.1.0/24 to E.
- o B advertises 10.1.1.0/25 {B,A} to C.
- o E advertises 10.1.1.0/24 {E,A} to F, filtering 10.1.1.0/25 {E,A} based on some local policy.
- o F advertises 10.1.1.0/24 {F,E,A} to C and to G.
- o C advertises 10.1.1.0/24 {C,F,E,A} to D, filtering 10.1.1.0/25 {B,A} based on some local policy.
- o G advertises 10.1.1.0/24 {G,F,E,A} to D.
- o D chooses 10.1.1.0/24 {C,F,E,A} over 10.1.1.0/24 {G,F,E,A}.

What path will traffic forwarded to C destined to hosts within 10.1.1.0/25 actually follow?

- o D forwards this traffic to C, since its best path is through 10.1.1.0/24 {C,F,E,A}.
- o C forwards this traffic to B, since its best path is through 10.1.1.0/25 {B,A}.
- o B forwards this traffic to A, since its best path is through 10.1.1.0/25 {A}.

Considering this result:


- o Is the AS Path valid? Both {G, F, E, A} and {C, F, E, A} are valid AS Paths, so both AS Paths in this example are valid.
- o Is the advertisement authorized? Assuming A is authorized to advertise 10.1.1.0/24, and all the autonomous systems in the example are authorized to readvertise 10.1.1.0/24, the route is authorized. However, C does not have destination nor transit authorization to 10.1.1.0/25, since B is the best path from C to 10.1.1.0/25, and D and A have explicit policies not to transit this path. In effect, then C does not have destination or transit authorization for 10.1.1.0/24, because it contains 10.1.1.0/25.
- o Is the AS Path consistent with the forwarding path (is there forwarding consistency)? While C is advertising the AS Path {C, F, E, A} to D to reach destinations within 10.1.1.0/24, it is actually forwarding along a different path to some destinations within this advertisement. Forwarding consistency does not exist within this internetwork.

In this example, A assumes that D will receive both the advertisement for 10.1.1.0/24 and the advertisement for 10.1.1.0/25, and will be able to use the included AS Path to impose their mutually agreed on policy not to transit B. Information about 10.1.1.0/25, however, is removed from the routing system by policies outside the knowledge or control of A or D. The information remaining in the routing system implies D may correctly route to destinations within 10.1.1.0/25 through C, since 10.1.1.0/25 is contained within 10.1.1.0/24, which C is legally advertising.

The tension between route authorization, destination authorization, and transit authorization can be seen clearly in this slightly more complex example. Route authorization implies destination and transit authorization in routing, but route authorization does not include destination or prefix authorization in this example.

2.5. Fourth Example: Interior and Exterior Paths Mismatch

This is the most complex example we will cover in this document. It can be argued that the configuration described in this example is a misconfiguration, but we have chosen this example for its simplicity, as an illustration of the complexity of the interaction between interior and exterior gateway protocols within an autonomous system. BGP route reflectors, particularly when configured in a hierarchy, provide many examples of forwarding inconsistency, but they are much more complex than this small example.



In this diagram, each router is labeled, with the AS in which it is contained, in parenthesis following the router label. As its best path to 10.1.1.0/24:

- o Router E is using its local (intra-AS) path.
- o Router C is using the path through AS3.
- o Router D is using the path through Router E.
- o Router B is using the path through Router E.

Examining the case of Router B more closely, however, we discover that while Router B prefers the path it has learned from Router E, that path has been advertised with a next hop of Router E itself. However, Router B's best path to this next hop (i.e., Router E), as determined by the interior routing protocol, is actually through Router C. Thus, Router B advertises the path {2, 5} to Router A, but traffic actually follows the path {2, 3, 5} when Router B receives it.

The system administrator of AS1 has determined there is an attacker in AS3, and has set the policy on router A to avoid any route with AS3 in the AS Path. So, beginning with this rule, it discards the path learned from AS9. It now examines the two remaining paths,

learned from AS2 (B) and AS6, and determines the best path is {2, 5}, through AS2 (B). However, unknown to A, AS2 (B) is also connected to AS3, and is transiting traffic to AS5 via the path {2, 3, 5}.

Returning to our questions:

- o Is the AS Path valid? The AS Path {2, 3, 5} is a valid AS Path.
- o Is the route authorized? Assuming each AS along the path is authorized to originate, or readvertise, 10.1.1.0/24, the route is authorized. Destination authorization is also clear in this situation, since 10.1.1.0/24 is the single destination throughout the example. Transit authorization is a little more difficult to determine, since the traffic doesn't actually flow along the AS Path contained in the routing advertisement. It's impossible to claim the AS Path {2,3,5} is a valid path from either the route originator or the traffic originator, since that AS Path is not the AS Path advertised. Essentially, Router E assumes transit authorization from route authorization, when there is no way to determine that AS3 is actually authorized to transit traffic to 10.1.1.0/24.
- o Is the AS Path consistent with the forwarding path (is there forwarding consistency)? The advertised AS Path is {2, 5}, while the traffic forwarded to the destination actually transits {2, 3, 5}. Forwarding is not consistent in this example.

3. Summary

The examples given in this document are not the only possible examples of forwarding that is inconsistent with the advertised AS Path; [ROUTINGLOGIC] also provides some examples, as well. [ASTRACEROUTE] provides some interesting background on the practical impact of forwarding that is inconsistent with the advertised AS Path, in the context of attempting to trace the actual path of packets through a large internetwork, running BGP as an exterior gateway protocol.

Routing strongly intertwines the concepts of route authorization, destination authorization, and transit authorization. If a BGP speaker is authorized to advertise a specific route, routing assumes that it is also authorized to advertise every possible subblock within the destination prefix, and the advertiser is authorized to transit packets to every destination within the route. By surveying these examples, we see that route authorization does not, in fact, equal destination authorization or transit authorization, calling into question the value of route authorization.

There are no easy or obviously scalable solutions to this problem.

4. Acknowledgements

We would like to thank Steve Kent for his contributions and comments on this document. We would also like to thank Joel Halpern for his work in clarifying many sections of the document, including additional text in critical sections.

5. Security Considerations

This document does not propose any new extensions or additions to existing or proposed protocols, and so does not impact the security of any protocol.

6. Informative References

- [ASTRACEROUTE] Feamster, N. and H. Balakrishnan, "Towards an Accurate ASLevel Traceroute Tool", SIGCOMM ACM SIGCOMM, 2003.
- [BGP-MD5] Heffernan, A., "Protection of BGP Sessions via the TCP MD5 Signature Option", RFC 2385, August 1998.
- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.
- [ROUTINGLOGIC] Feamster, N. and H. Balakrishnan, "Towards a Logic for Wide Area Routing", SIGCOMM ACM SIGCOMM Workshop on Future Directions in Network Architecture, Germany, August 2003.
- [SOBGP] White, R., "Architecture and Deployment Considerations for Secure Origin BGP (soBGP)", Work in Progress.

Authors' Addresses

Russ White
Cisco Systems

EMail: riw@cisco.com

Bora Akyol
Cisco Systems

EMail: bora@cisco.com

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in BCP 78 and at www.rfc-editor.org/copyright.html, and except as set forth therein, the authors retain all their rights.

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

Intellectual Property

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

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

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

