Network Working Group
Internet Draft
Expiration Date: October 2004
File Name: draft-white-pathconsiderations-02.txt

Russ White
Bora Akyol
Cisco Systems
Nick Feamster
MIT
April 2004

Considerations in Validating the Path in Routing Protocols
draft-white-pathconsiderations-02.txt

Status of this Memo

This document is an Internet-Draft and is in full conformance with 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. Internet Drafts may be updated, replaced, or obsoleted by other documents at any time. It is not appropriate to use Internet Drafts as reference material or to cite them other than as a "working draft" or "work in progress".

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

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

Abstract

A good deal of consideration has gone into, and is currently being given to, validating the path to a destination advertised by an adjacent router or peer, such as [S-BGP], [SOBGP-DEPLOY], and [IRV]. Since much of this effort has been focused on BGP, this draft discusses some issues with this work in terms of BGP.

One of the primary assumptions in much of this work is that the authentication of a given advertisement received by a specific BGP speaker is the same as authorization to use the path advertised. In other words, it is generally assumed that if a BGP speaker receives an advertisement for which the AS Path can somehow be verified, the speaker is authorized to transit traffic along the path specified contained in the update, and the traffic forwarded to the destination contained in the update will actually follow the path advertised.

This draft shows these two assumptions cannot be held to be true in a path vector routing system.

1. Background

With the heightened interest in network security, the security of the information carried within the routing system 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. One issue that cannot be addressed through peer authentication is the validity of the path represented by a BGP speaker when advertising reachability to a specific prefix.

To place this in more direct terms, consider this small network.

10.1.1.0/24--A---B--C

Assume C has received an advertisement for 10.1.1.0/24 from B, with an AS Path of {A, B}. We can ask three questions about this update:

- Does a path actually exist from the advertising router to the destination advertised?
- Is C authorized, though receiving this advertisement, to transit traffic along this path to each reachable destination within the prefix advertised?
- Will traffic forwarded to some destination within 10.1.1.0/24 actually follow the path described in the update advertised by B?

The primary question we would like to examine in this draft is which of these three questions can actually be answered witin a path vector protocol, such as BGP. This draft contends that the second and third meanings of path validity cannot be verified in a distance or path vector protocol.

We will first examine some fundamental concepts of routing and path selection in general, then we will proceed through some examples, and re-examine each question above in light of each of those examples. In each example, we will discuss policy, in terms of an acceptable path for the receiver of the traffic (the originator of the advertisement) or the transmitter of the traffic.

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 why the second of two questions cannot be answered in a path vector routing system. Finally, a short section on transitive authorization in a path vector protocol is provided, which considers the reasons behind the results we find in the examples.

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

This sheds some light on the purpose of the path included in an 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 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.

We need to keep this principle in mind when considering the use of the path carried in a path vector protocol for setting policy.

2.2. First Example: Manual Intervention in the Path Choice

In the small network:

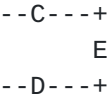

A may receive an advertisement from B that E is reachable along the path {B, C, E}. Based on this information, A may forward packets to B, expecting them to take the path described. However, at B's edge router receiving this traffic, the network administrator may have configured a static route making the next hop to E the edge router with D.

Although this is an "extreme" example, since we can hardly claim the

information within the routing protocol is actually insufficient, we will still find it instructive to examine this example in light of the original three questions:

- Does a path actually exist from the advertising router to the destination advertised? Yes, it clearly does.

- Is C authorized, though receiving this advertisement, to transit traffic along this path to each reachable destination within the prefix advertised? This question isn't addressed in this example, since we have no idea what A or E's policies are.

- Will traffic forwarded to some destination within 10.1.1.0/24 actually follow the path described in the update advertised by B? No, traffic forwarded by A torwards B will not follow the path described in B's update.

  There is no way to account for the overriding of a routing protocol's information through static configuration or through other routing protocols running on the same devices, since routing is a hop by hop endeavor.

2.3. Second Example: An Unintended Reachable Destination

Here, we return the small network outlined earlier in this draft.

10.1.1.0/24---A---B---C

We will assume, for argument's sake, that A and C are competitors, and A would like to prevent hosts within C's network from reaching anything within its network. A has implemented this policy by advertising 10.1.1.0/24 to B with some restriction (we can use the NO_ADVERTISE community described in [BGP] for this purpose) so B cannot readvertise the destination to C.

However, unknown to A, B is actually advertising a default route only to C, and not a full routing table. If some host within C, then, originates a packet destined to 10.1.1.1, what will happen? The packet will be routed according to the default route advertised by B. When the edge router between B and C receives the packet, it will forward the packet along the 10.1.1.0/24 route learned from A, forwarding the traffic into A's network.

Returning to our questions:

- Does a path actually exist from the advertising router to the destination advertised? Yes, it does. If B doesn't know of some specific host connected to the Internetwork, we can assume that host doesn't exist, thus the default route is a valid route for B to advertise in this case.

- Is C authorized, though receiving this advertisement, to transit traffic along this path to each reachable destination within the prefix advertised? No. In fact, A has explicity attempted to prevent C from using this path to reach any hosts within its network.

- Will traffic forwarded to some destination within 10.1.1.0/24 actually follow the path described in the update advertised by B? No. The path advertised with the default route ends in B, while the traffic transits beyond B, into A, which is hidden in the AS Path B advertises.

  The basic problem here is that A is assuming that because B doesn't receive an advertisement for 10.1.1.0/24, it cannot reach 10.1.1.0/24. We see, however, that lack of routing information does not imply lack of authorization, because aggregates cover many possible destinations, and the default is just the shortest prefix aggregate available.

2.4. Third Example: Following a Specific Path

This example is slightly more complex than the last two. Given the following small network:

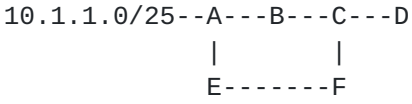

Assume the following:

- A advertises 10.1.1.0/25 to B and E.

- B advertises 10.1.1.0/24 to C.

- E advertises the aggregate 10.1.1.0/24 to F.

- F advertises the aggregate 10.1.1.0/24 to C.

- C advertises the aggregate 10.1.1.0/24 to D, but not the more specific 10.1.1.0/25.

There are a number of reasons C might advertise the aggregate, and not the more specific, to D, including (but not limited to):

- B and C both accept prefixes with a length of /25, while D does not, so D filters the 10.1.1.0/25 inbound from C.
- A has a policy that nothing originating in D may traverse B, so it advertises the update in such a way to prevent C from readvertising 10.1.1.0/25 to D.
- D has a policy that anything destined to A cannot traverse B, so it blocks 10.1.1.0/25 at its border with C (because it finds B in the AS Path).
- B has a policy that traffic originating in D will not transit B's network to reach A.
- C notes it has two advertisements covering the same address space, and advertises only one of them to D.

So, there are several possible reasons information about 10.1.1.0/25 is removed from the routing system at this point. What is the practical result of removing this information? Suppose some host in D originates a packet destined to 10.1.1.1. The packet will be forwarded based on the route to 10.1.1.0/24 in D, to C. The edge router in C finds it has a route to that destination, 10.1.1.0/25, and forwards the traffic to B, for final transmission to A.

Let's return to our questions:

- Does a path actually exist from the advertising router to the destination advertised? Yes, C really does have a route to 10.1.1.1.
- Is C authorized, though receiving this advertisement, to transit traffic along this path to each reachable destination within the prefix advertised? If the reason C doesn't readvertise the 10.1.1.0/25 route to D is because of A's, B's, or D's policy, then no, D is not authorized to transit the path the traffic actually takes to reach this destination.
- Will traffic forwarded to some destination within 10.1.1.0/24 actually follow the path described in the update advertised by B? No. The path described in C's advertisement to D is {C, F, E}, while the path the traffic actually takes is {C, B, A}.

2.5. Fourth Example: A Mismatch Between the Interior and Exterior Paths

This is the most complex example we will cover in this draft. Many people will note the configuration described is a misconfiguration, but there are many such possible situations in the interaction between BGP and interior gateway protocols. Note this example doesn't involve the removal of information from the routing system, and it is specific only to BGP, not to path vector protocols in general.

Assume we have the following small internetwork:

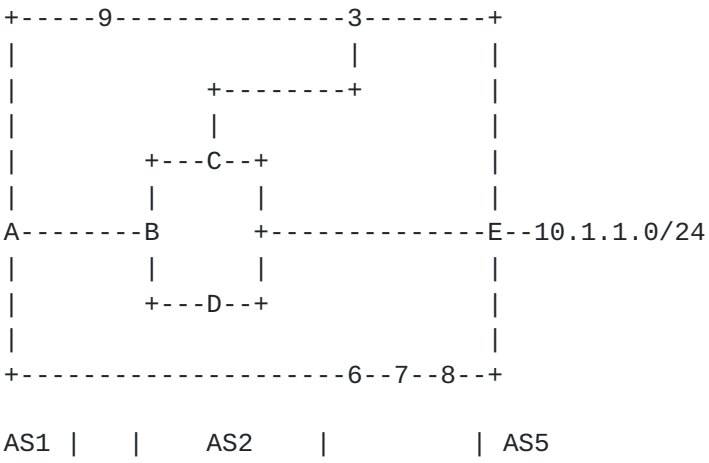

In this diagram, routers are represented by letters, and autonomous systems by numbers. So:

- Router A is in AS1
- Routers B, C, and D are in AS2
- Router E is in AS5

Each router is using, as its best path to 10.1.1.0/24:

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

- Does a path actually exist from the advertising router to the destination advertised? Yes.
- Is A authorized, though receiving this advertisement, to transit traffic along this path to each reachable destination within the prefix advertised? There is no mention of policy within this example, so we can't answer this question.
- Will traffic forwarded to some destination within 10.1.1.0/24 actually follow the path described in the update advertised by B? No. Router B advertises the path {2, 5} to Router A, but traffic actually follows the path {2, 3, 5}.

This is only one given example of the interaction between BGP and interior gateway protocols resulting in one path being advertised, and another path being taken. It's actually common in the case of route reflectors, for instance.

2.6. Trasitive Authorization in Path Vector Protocols

A route is carried as a prefix and its associated attributes, one of which is the AS Path, [BGP]. It is possible to verify that the prefix is originated by its authorized owner AS by multiple means including an encrypted certificate or the use of a route or address registry and checking that the first AS in the path is the AS that owns the prefix. However, this authorization can not be carried over to infer that the path associated with this prefix is in fact authorized by the originating AS. As we have shown in the examples above, BGP does not transmit routing information intact across autonomous systems.

In fact, routing information is frequently summarized or filtered,

with more specific prefixes hidden and sometimes completely ignored via application of routing policy. Due to application of routing policy as well as the hop by hop nature of IP routing, the only facts that can be inferred from a prefix and its path attribute received by BGP are:

- Originator AS is the authorized advertiser of the prefix: This can be achieved by means of use of a routing registry, or via security extensions to BGP [soBGP, SBGP].

- The path being carried is _plausible_; that is, the path is a path that is likely to carry the packets destined for that prefix to the originating AS. This can be achieved by either knowledge of peering information (as can be obtained by means of a routing registry) or via security extensions to BGP.

Therefore, from a security perspective, a prefix and its path can be classified in two dimensions: Originating AS :={Authorized, Unauthorized}, Path := {Plausible, Implausible}.

3. Summary

While it is tempting to set policy, or to infer policy, from the existence or non-existence of information within a routing system, it isn't possible to do so, since routing systems remove information on a regular basis. Further, it appears logical that policy could be set based on the path advertised in a path vector protocol, however, since routing information is regularly removed from the routing system, it isn't possible to do so.

[ROUTINGLOGIC] also provides some instances in which information is removed from the routing system, through other means (such as route reflectors), which could result in situations similar to the ones cited above. [ASTRACEROUTE] also provides some interesting background on the problems involved in attempting to map a packet's path to an AS Path advertised in BGP.

Informative References

[BGP] Rekhter, Y., and T. Li, "A Border Gateway Protocol 4 (BGP-4)", RFC 1771, March 1995.

[S-BGP]
Lynn, C, et al., "Secure BGP (S-BGP)", draft-clynn-s-bgp-protocol-01.txt, June 2003

[SOBGP-DEPLOY]
White, R. (editor), "Architecture and Deployment Considerations for Secure Origin BGP (soBGP) Deployment", draft-white-sobgp-deployment-02, April 2004

[IRV] Goodell, G., et al., Working Around BGP: An Incremental Approach to Improving Security and Accuracy of Interdomain Routing, http://www.isoc.org/isoc/conferences/ndss/03/proceedings/papers/5.pdf

[BGP-MD5]
Heffernon, A., Protection of BGP Sessions via the TCP MD5 Signature Option, RFC 2385, August 1998

[ROUTINGLOGIC]
Nick Feamster and Hari Balakrishnan, Towards a Logic for Wide-Area Internet Routing, ACM SIGCOMM Worshop on Future Directions in Network Architecture, Germany, August 2003

[ASTRACEROUTE]
Zhuoqing Morley Mao, et al., Towards an Accurate ASLevel Traceroute Tool, SIGCOMM 2003

4. Author's Addresses

Russ White
Cisco Systems
7025 Kit Creek Road
Research Triangle Park, NC 27709
riw@cisco.com

Nick Feamster
200 Technology Square, NE43-504
Cambridge, MA 02139-3578
617-253-7341
feamster@lcs.mit.edu

Bora Akyol
Cisco Systems
170 West Tasman Drive
San Jose, CA 95134
bora@cisco.com