

INTERNET-DRAFT
[draft-gill-btsh-01.txt](#)

Category
Expires: November 2003

Vijay Gill
John Heasley
David Meyer
Informational
May 2003

The BGP TTL Security Hack (BTSH)
<[draft-gill-btsh-02.txt](#)>

Status of this Document

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 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/ietf/1id-abstracts.txt>

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

This document is a product of an individual. Comments are solicited and should be addressed to the author(s).

Copyright Notice

Copyright (C) The Internet Society (2003). All Rights Reserved.

Abstract

The BGP TTL Security Hack (BTSH) is designed to protect the BGP [[RFC1771](#)] infrastructure from CPU-utilization based attacks. While BTSH is most effective in protecting directly connected BGP peers, it can also provide a lower level of protection to multi-hop sessions.

Table of Contents

- [1. Introduction](#) [2](#)
- [2. Assumptions Underlying BTSH.](#) [3](#)
 - [2.1. Assumptions on Attack Sophistication.](#) [3](#)
- [3. BTSH Procedure](#) [4](#)
 - [3.1. Multi-hop Scenarios](#) [5](#)
 - [3.1.1. iBGP Handling.](#) [5](#)
- [4. Intellectual Property.](#) [5](#)
- [5. Acknowledgments.](#) [6](#)
- [6. Security Considerations.](#) [6](#)
- [7. IANA Considerations.](#) [6](#)
- [8. References](#) [6](#)
 - [8.1. Normative References.](#) [7](#)
 - [8.2. Informative References.](#) [7](#)
- [9. Author's Addresses](#) [8](#)
- [10. Full Copyright Statement.](#) [8](#)

1. Introduction

The BGP TTL Security Hack (BTSH) is designed to protect the BGP [[RFC1771](#)] infrastructure from CPU-utilization based attacks. In particular, while cryptographic techniques can protect the routed infrastructure from a wide variety of attacks, many attacks based on CPU-overload can be prevented by the simple mechanism described in this document.

BTSH is based on the fact that the vast majority of ISP eBGP peerings are established between routers that are adjacent [[PEERING](#)]. Thus most eBGP peerings are either directly between connected interfaces or at the worst case, are between loopback and loopback, with static routes to loopbacks. Since TTL spoofing [[BALDWIN2001](#)] is considered nearly impossible, a mechanism based on an expected TTL value can provide a simple and reasonably robust defense from infrastructure attacks based on forged BGP packets.

The keywords MUST, MUST NOT, MAY, OPTIONAL, REQUIRED, RECOMMENDED, SHALL, SHALL NOT, SHOULD, SHOULD NOT are to be interpreted as defined in [RFC 2119](#) [[RFC2119](#)].

2. Assumptions Underlying BTSH

BTSH is predicated upon the following assumptions:

- (i). The vast majority of eBGP peerings are between adjacent routers [[PEERING](#)].
- (ii). It is common practice for many service providers to ingress filter (deny) packets that have the provider's loopback addresses as the source IP address.
- (iii). Use of BTSH is OPTIONAL, and can be configured on a per-peer (group) basis.
- (iv). The router supports a method of classifying traffic destined for the route processor into interesting/control and not-control queues.
- (iv). The peer routers both implement BTSH.

2.1. Assumptions on Attack Sophistication

Throughout this document, we assume that attackers have evolved in both sophistication and access to the point that they can send control traffic to a BGP session, and that this traffic appears to be valid control traffic (i.e., has the source/destination of configured peer routers).

We also assume that each router in the path between the attacker and the victim BGP speaker decrements TTL properly (clearly, if the either the path or the adjacent peer is compromised, then there are worse problems we have to worry about).

Since the vast majority of our peerings are between adjacent routers, we can set the TTL on the BGP packets to 255 (the maximum possible for IP) and then reject any BGP packets that come in from configured peers which do NOT have a TTL in the range 255-254. That is, the receive TTL is expected to be within a small range of 1 or 2 (254-255). The actual value depends upon the architecture, but is it is expected that the receiver will verify the range.

BTSH can be disabled for applications such as route-servers and other large diameter multi-hop peerings. In the event that an the attack comes in from a compromised multi-hop peering, that peering can be

shut down (a method to reduce exposure to multi-hop attacks is outlined below).

3. BTSH Procedure

BTSH SHOULD not be enabled by default. The following process described the per-peer behavior:

(i). If BTSH is enabled, do the following:

(a). For directly connected routers,

- o Set the TCP TTL for the BGP connection a value in the range 255-254.

- o For each configured eBGP peer:

Update the receive path ACL/firewall to only allow BGP packets to pass onto the Route Processor (RP) that have the correct <source,destination,TTL> tuple. The TTL must either be in the range 255-254 (directly connected peer), or 255-(configured-range-of-hops) for a multi-hop peer. We specify a range here to achieve some robustness to changes in topology. The connected check should be disabled for such non-direct peerings.

It is assumed that a receive path ACL is an ACL that is designed to control which packets are allowed to go to the RP. This procedure will only allow BGP packets from adjacent router to pass onto the RP.

- (c). If the TTL is not in the range 255-254 (or 255-(configured-range-of-hops) for multi-hop peers), punt into low priority queue, log, or silently discard.

(ii). If BTSH is not enabled for a particular peering, normal [RFC 1772](#) [[RFC1772](#)] protocol behavior is followed.

3.1. Multi-hop Scenarios

When a multi-hop BGP session is required, we set the expected TTL value to be 255-(configured-range-of-acceptable-of-hops). While this approach provides a qualitatively lower degree of security for BGP (i.e., an DoS attack could be theoretically be launched by compromising some box in the path). However, BTSH will still catch the vast majority of observed DDoS attacks against eBGP.

3.1.1. iBGP Handling

BTSH is not used for iBGP peer groups. Current best practice is to protect peers (both eBGP and iBGP) with an MD5 signature [[RFC2385](#)]. Such sessions can be further protected by filtering (deny) at the network edge for any packet that has a source address of one of the loopbacks addresses used for iBGP peering.

4. Intellectual Property

The IETF takes no position regarding the validity or scope of any intellectual property 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; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in [BCP-11](#). Copies of claims of rights made available for publication 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 implementors or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

5. Acknowledgments

The BTSH concept originated with many different people, including Paul Traina and Jon Stewart. Ryan McDowell also suggested a similar idea. Steve Bellovin, Jay Borkenhagen and Randy Bush also provided useful feedback on early versions of this document.

6. Security Considerations

BTSH is a simple procedure that protects single hop BGP sessions, except in those cases where the directly connected peer has been compromised. While the method is less effective for multi-hop BGP sessions, it still closes the window on several forms of attack.

Protection of the BGP infrastructure beyond this method will likely require cryptographic machinery such as is envisioned by Secure BGP (S-BGP) [[SBGP1](#), [SBGP2](#)], and/or other extensions. For example, consider the class of attacks based on forged SYN packets directed to port 179/tcp on a large core infrastructure routers. In this case, the routers respond with SYN/ACKs (or ICMP messages) towards the victim, resulting in flooding of the victim's link being flooded with SYN/ACK or ICMP traffic. Preventing such attacks will likely require that BGP speakers send SYN/ACKs only to configured neighbors, and they never send ICMP messages related to these events.

Finally, note that in the multi-hop case described above, we specify a range of acceptable TTLs in order to achieve some robustness to topology changes. This robustness to topological change comes at the cost of the loss some robustness to different forms of attack.

7. IANA Considerations

This document creates a no new requirements on IANA namespaces [[RFC2434](#)].

8. References

8.1. Normative References

- [RFC1771] "A Border Gateway Protocol (BGP-4)", Y. Rekhter, T. Li, Editors, [RFC 1771](#), March, 1995
- [RFC1772] "Application of the Border Gateway Protocol in the Internet", Y. Rekhter, P. Gross, [RFC 1772](#), March, 1995
- [RFC2385] "Protection of BGP Sessions via the TCP MD5 Signature Option", A. Heffernan, [RFC 2385](#), August, 1998.
- [SBGP1] "Secure Border Gateway Protocol (Secure-BGP)", Stephen Kent and Charles Lynn and Karen Seo, IEEE Journal on Selected Areas in Communications, volume 18, number 4, April, 2000.
- [SBGP2] "Secure Border Gateway Protocol (S-BGP) -- Real World Performance and Deployment Issues", Stephen Kent and Charles Lynn and Joanne Mikkelsen and Karen Seo, Proceedings of the IEEE Network and Distributed System Security Symposium, February, 2000.

8.2. Informative References

- [BALDWIN2001] http://www.sekure.net/docs/detecting_spoof.txt
- [PEERING] Empirical data gathered from the Sprint and AOL backbones, October, 2002.
- [RFC2119] "Key words for use in RFCs to Indicate Requirement Levels", S. Bradner, [RFC 2119](#), March, 1997.
- [[RFC2434](#)] Narten, T., and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [RFC 2434](#)/BCP 0026, October, 1998.

9. Author's Addresses

Vijay Gill
AOL
Email: vijay@umbc.edu

John Heasley
Verio
Email: heas@shrubbery.net

David Meyer
Sprint
Email: dmm@1-4-5.net

10. Full Copyright Statement

Copyright (C) The Internet Society (2003). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS 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.

