Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: December 31, 2016

T. King
C. Dietzel
DE-CIX Management GmbH
J. Snijders
NTT
G. Doering
SpaceNet AG
G. Hankins
Nokia

June 29, 2016

BLACKHOLE BGP Community for Blackholing
draft-ietf-grow-blackholing-01

Abstract

This document describes the use of a well-known Border Gateway Protocol (BGP) community for blackholing in IP networks. This well-known advisory transitive BGP community, namely BLACKHOLE, allows an origin AS to specify that a neighboring network should blackhole a specific IP prefix.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in [RFC2119] only when they appear in all upper case. They may also appear in lower or mixed case as English words, without normative meaning.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

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."

This Internet-Draft will expire on December 31, 2016.

Table of Contents

1. Introduction

Network infrastructures have been increasingly hampered by DDoS attacks. In order to dampen the effects of these DDoS attacks, IP networks have offered BGP blackholing to neighboring networks via various mechanisms such as described in [RFC3882] and [RFC5635].

DDoS attacks targeting a certain IP address may cause congestion of links used to connect to other networks. In order to limit the impact of such a scenario on legitimate traffic, networks adopted a mechanism called BGP blackholing. A network that wants to trigger blackholing needs to understand the triggering mechanism adopted by its neighboring networks. Different networks provide different mechanisms to trigger blackholing, including but not limited to pre-defined blackhole next-hop IP addresses, specific BGP communities or via an out-of-band BGP session with a special BGP speaker.

Having several different mechanisms to trigger blackholing in different networks makes it an unnecessarily complex, error-prone and cumbersome task for network operators. Therefore a well-known BGP community [RFC1997] is defined for operational ease.

Having such a well-known BGP community for blackholing also supports networks because:

- implementing and monitoring blackholing becomes easier when implementation and operational guides do not cover many options that trigger blackholing.
- the number of support requests from customers about how to trigger blackholing in a particular neighboring network will be reduced as the codepoint for common blackholing mechanisms is unified.

Making it considerably easier for network operators to utilize blackholing makes operations easier.

2. BLACKHOLE Attribute

This document defines the use of a new well-known BGP transitive community, BLACKHOLE.

The semantics of this attribute allow a network to interpret the presence of this community as an advisory qualification to drop any traffic being sent towards this prefix.

3. Operational Recommendations

3.1. IP Prefix Announcements with BLACKHOLE Community Attached

When a network is under DDoS duress, it MAY announce an IP prefix covering the victim's IP address(es) for the purpose of signaling to neighboring networks that any traffic destined for these IP address(es) should be discarded. In such a scenario, the network operator SHOULD attach BLACKHOLE BGP community.

3.2. Local Scope of Blackholes

A BGP speaker receiving a BGP announcement tagged with the BLACKHOLE BGP community SHOULD add a NO_ADVERTISE, NO_EXPORT or similar community to prevent propagation of this route outside the local AS.

Unintentional leaking of more specific IP prefixes to neighboring networks can have adverse effects. Extreme caution should be used when purposefully propagating IP prefixes tagged with the BLACKHOLE BGP community outside the local routing domain.

3.3. Accepting Blackholed IP Prefixes

It has been observed that announcements of IP prefixes larger than /24 for IPv4 and /48 for IPv6 are usually not accepted on the Internet (see section 6.1.3 [RFC7454]). However, blackhole routes should be as small as possible in order to limit the impact of discarding traffic for adjacent IP space that is not under DDoS duress. Typically, the blackhole route's prefix length is as specific as /32 for IPv4 and /128 for IPv6.

BGP speakers SHOULD only accept and honor BGP announcements carrying the BLACKHOLE community if the announced prefix is covered by a shorter prefix for which the neighboring network is authorized to advertise.

4. IANA Considerations

The IANA is requested to register BLACKHOLE as a well-known BGP community with global significance:

   BLACKHOLE (= 0xFFFF029A)

The low-order two octets in decimal are 666, amongst network operators a value commonly associated with BGP blackholing.

5. Security Considerations

BGP contains no specific mechanism to prevent the unauthorized modification of information by the forwarding agent. This allows routing information to be modified, removed, or false information to be added by forwarding agents. Recipients of routing information are not able to detect this modification. Also, RPKI [RFC6810] and BGPSec [I-D.ietf-sidr-bgpsec-overview] do not fully resolve this situation. For instance, BGP communities can still be added or altered by a forwarding agent even if RPKI and BGPSec are in place.

The BLACKHOLE BGP community does not alter this situation.

A new additional attack vector is introduced into BGP by using the BLACKHOLE BGP community: denial of service attacks for IP prefixes.

The unauthorized addition of the BLACKHOLE BGP community to an IP prefix by a forwarding agent may cause a denial of service attack based on denial of reachability. The denial of service will happen if a network offering blackholing is traversed. However, denial of service attack vectors to BGP are not new as the injection of false routing information is already possible.

In order to further limit the impact of unauthorized BGP announcements carrying the BLACKHOLE BGP community, the receiving BGP speaker SHOULD verify by applying strict filtering (see section 6.2.1.1.2. [RFC7454]) that the peer announcing the prefix is authorized to do so. If not, the BGP announcement should be filtered out.

The presence of this BLACKHOLE BGP community may introduce a resource exhaustion attack to BGP speakers. If a BGP speaker receives many IP prefixes containing the BLACKHOLE BGP community, its internal resources such as CPU power, memory or FIB capacity might exhaust, especially if usual prefix sanity checks (e.g. such as IP prefix length or number of prefixes) are disabled (see Section 3.3).

6. References

6.1. Normative References

[RFC1997] Chandra, R., Traina, P., and T. Li, "BGP Communities Attribute", RFC 1997, DOI 10.17487/RFC1997, August 1996, <http://www.rfc-editor.org/info/rfc1997>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

6.2. Informative References

[I-D.ietf-sidr-bgpsec-overview] Lepinski, M. and S. Turner, "An Overview of BGPsec", draft-ietf-sidr-bgpsec-overview-08 (work in progress), June 2016.

[RFC3882] Turk, D., "Configuring BGP to Block Denial-of-Service Attacks", RFC 3882, DOI 10.17487/RFC3882, September 2004, <http://www.rfc-editor.org/info/rfc3882>.

[RFC5635] Kumari, W. and D. McPherson, "Remote Triggered Black Hole Filtering with Unicast Reverse Path Forwarding (uRPF)", RFC 5635, DOI 10.17487/RFC5635, August 2009, <http://www.rfc-editor.org/info/rfc5635>.

[RFC6810] Bush, R. and R. Austein, "The Resource Public Key Infrastructure (RPKI) to Router Protocol", RFC 6810, DOI 10.17487/RFC6810, January 2013, <http://www.rfc-editor.org/info/rfc6810>.

[RFC7454] Durand, J., Pepelnjak, I., and G. Doering, "BGP Operations and Security", BCP 194, RFC 7454, DOI 10.17487/RFC7454, February 2015, <http://www.rfc-editor.org/info/rfc7454>.

Appendix A. Acknowledgements

The authors would like to gratefully acknowledge many people who have contributed discussions and ideas to the making of this proposal. They include Petr Jiran, Yordan Kritski, Christian Seitz, Nick Hilliard, Joel Jaeggli, Christopher Morrow, Thomas Mangin, Will Hargrave and Niels Bakker.

Authors' Addresses

Thomas King
DE-CIX Management GmbH
Lichtstrasse 43i
Cologne 50825
Germany

Email: thomas.king@de-cix.net

Christoph Dietzel
DE-CIX Management GmbH
Lichtstrasse 43i
Cologne 50825
Germany

Email: christoph.dietzel@de-cix.net

Job Snijders
NTT Communications, Inc.
Theodorus Majofskistraat 100
Amsterdam 1065 SZ
NL

Email: job@ntt.net

Gert Doering
SpaceNet AG
Joseph-Dollinger-Bogen 14
Munich 80807
Germany

Email: gert@space.net

Greg Hankins
Nokia
777 E. Middlefield Road
Mountain View, CA 94043
USA

Email: greg.hankins@nokia.com