Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: May 4, 2016

J. Mauch
J. Snijders
NTT

November 1, 2015

By default reject propagation when no policy is associated with a BGP peering session.
draft-mauch-bgp-reject-01.txt

Abstract

This document defines the default behaviour of a BGP speaker when no explicit policy is associated with a BGP peering session.

Foreword

A placeholder to list general observations about this document.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [1].

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 May 4, 2016.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (http://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction 2
2. Definitions and Acronyms 3
3. Solution Requirements 3
4. Acknowledgements 3
5. Security Considerations 3
6. References 3
6.1. Normative References 3
6.2. Informative References 4
Authors' Addresses 4

1. Introduction

BGP speakers have many default settings which need to be revisited as part of improving the routing ecosystem. There is a need to provide guidace to BGP implementors for the default behaviors of a well functioning internet ecosystem. Routing leaks [3] are part of the problem, but software defects and operator misconfigurations are just a few of the attacks on internet stability we aim to address.

Usually BGP speakers accept all routes from a configured peer or neighbor. This practice dates back to the early days of internet protocols in being very permissive in offering routing information to allow all networks to reach each other. With the core of the internet becoming more densely interconnected the risk of a misbehaving edge device or BGP speaking customer poses signficiant risks to the reachability of critical services.

This proposal intends to solve this situation with the requiring the explicity configuration of BGP policy for any non-iBGP speaking session such as customers, peers or confederation boundaries. When this solution is implemented, devices will no longer pass routes without explicit policy.

2. Definitions and Acronyms

- BGP: Border Gateway Protocol [2]

3. Solution Requirements

The following requirements apply to the solution described in this document:

- Software MUST mark any routes from an eBGP peer as 'invalid' in the Adj-RIB-In, if no explicit policy was configured.
- Software MUST NOT advertise any routes to an eBGP peer without an operator configuring a policy
- Software MUST NOT require a configuration directive to operate in this mode.
- Software MUST provide protection from internal failures preventing the advertisement and acceptance of routes
- Software MAY provide a configuration option to disable this security capability.

4. Acknowledgements

The authors would like to thank the following people for their comments and support: Shane Amante, Christopher Morrow, Robert Raszuk.

5. Security Considerations

This document addresses the basic security posture of a BGP speaking device within a network. Operators have a need for implementors to address the problem through a behavior change to mitigate against possible attacks from a permissive security posture. Attacks and inadvertent advertisements cause business impact necessitating this default behavior.

6. References

6.1. Normative References

[1] 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>.

[2] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <http://www.rfc-editor.org/info/rfc4271>.

6.2. Informative References

[3] "Methods for Detection and Mitigation of BGP Route Leaks", <https://tools.ietf.org/html/draft-sriram-idr-route-leak-detection-mitigation>.

Authors' Addresses

Jared Mauch
NTT Communications, Inc.
8285 Reese Lane
Ann Arbor Michigan 48103
US

Email: jmauch@us.ntt.net

Job Snijders
NTT Communications, Inc.
Amsterdam
NL

Email: job@ntt.net

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 30, 2016

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

July 29, 2015

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

Abstract

This document describes the use of a well-known Border Gateway Protocol (BGP) community for blackholing at IP networks and Internet Exchange Points (IXP). This well-known advisory transitive BGP community, namely BLACKHOLE, allows an origin AS to specify that a neighboring IP network or IXP 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 January 30, 2016.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (http://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction 2
2. BLACKHOLE Attribute 3
3. Operational Recommendations 3
3.1. IP Prefix Announcements with BLACKHOLE Community Attached 3
3.2. Local Scope of Blackholes 3
3.3. Accepting Blackholed IP Prefixes 4
3.4. IXPs: Peering at Route Servers 4
4. IANA Considerations 4
5. Security Considerations 5
6. References 5
6.1. Normative References 5
6.2. Informative References 6
6.3. URIs 6
Appendix A. Acknowledgements 6
Authors' Addresses 7

1. Introduction

The network infrastructure has been getting hammered by DDoS attacks for years. In order to block DDoS attacks, IP networks have offered BGP blackholing to neighboring networks (iBGP scenarios [RFC3882] and RTBH filtering [RFC5635]), much like some IXPs have recently started to do.

DDoS attacks targeting a certain IP network may cause congestion of links used to connect to other networks. In order to limit the impact of such a scenario on legitimate traffic, IP networks and IXPs adopted a mechanism called BGP blackholing. A network that wants to trigger blackholing needs to understand the triggering mechanism adopted by its neighboring IP networks and IXPs. Different IP networks and IXPs provide different BGP mechanism to trigger

blackholing including pre-defined blackhole next- hop IP addresses and pre-defined BGP communities.

Having several different mechanisms to trigger blackholing at different IP networks and IXPs 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 IP networks and IXPs as

- o implementing and monitoring blackholing gets easier if implementation and operational guides do not cover many options to trigger blackholing
- o the amount of support requests from customers about how to trigger blackholing at a particular IP network or IXP will be reduced as the mechanism is unified

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

2. BLACKHOLE Attribute

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

The semantics of this attribute is to 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 an IP 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 IP networks or IXPs that any traffic destinated 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 communities 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 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.

3.4. IXPs: Peering at Route Servers

Many IXPs provide the so-called policy control feature as part of their route servers [I-D.ietf-idr-ix-bgp-route-server] (see e.g. the LINX website [1]). Policy control allows members to specify by using BGP communities which ASNs connected to the route server receive a particular BGP announcement.

Combined usage of the BGP communities for blackholing and policy control allows a fine-grained control of a blackhole.

In some implementations of blackholing at IXPs, the route server after receiving a BGP announcement tagged with the BLACKHOLE BGP community rewrites the next-hop IP address to the pre-defined blackholing IP address before redistributing the announcement.

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 IP 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.

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 an IP network or IXP 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 and/or memory might get consumed, 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-idr-ix-bgp-route-server] Jasinska, E., Hilliard, N., Raszuk, R., and N. Bakker, "Internet Exchange BGP Route Server", draft-ietf-idr-ix-bgp-route-server-07 (work in progress), June 2015.

[I-D.ietf-sidr-bgpsec-overview] Lepinski, M., "An Overview of BGPsec", draft-ietf-sidr-bgpsec-overview-07 (work in progress), June 2015.

[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>.

6.3. URIs

[1] https://www.linx.net/members/support/route-servers.html

Appendix A. Acknowledgements

The authors gratefully acknowledges the contributions of:

- Petr Jiran, NIX.CZ, Milesovska 1136/5, Praha 130 00, Czech Republic, Email: pj@nix.cz
- Yordan Kritski, NetIX Ltd., 3 Grigorii Gorbatenko Str., Sofia 1784, Bulgaria, Email: ykritski@netix.net
- Christian Seitz, STRATO AG, Pascalstr. 10, Berlin 10587, Germany, Email: seitz@strato.de

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
Alcatel-Lucent
777 E. Middlefield Road
Mountain View, CA 94043
USA

Email: greg.hankins@alcatel-lucent.com