

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: August 8, 2014

H. Jeng
AT&T
J. Haas
Y. Rekhter
J. Zhang
Juniper Networks
February 4, 2014

Multicast Distribution Reachability Signaling
draft-ietf-idr-mdrs-00

Abstract

This document describes a mechanism whereby a subscriber's Internet service provider may signal in BGP the ability of the subscriber network to receive content using multicast connectivity. This mechanism is called Multicast Distribution Reachability Signaling (MDRS).

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 August 8, 2014.

Copyright Notice

Copyright (c) 2014 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

Internet-Draft

MDRS

February 2014

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	3
2.	Specification of Requirements	3
2.1.	Multicast Distribution Reachability Signaling	4
3.	IANA Considerations	5
4.	Security Considerations	5
5.	Acknowledgements	5
6.	Normative References	5
	Authors' Addresses	5

Internet-Draft

MDRS

February 2014

1. Introduction

Consider a content provider that wants to deliver a particular content to a set of customers/subscribers, where the provider and the subscribers are connected by an IP service provider. This content provider can deliver its content via unicast connectivity or, if supported by the subscriber network, multicast connectivity. A mechanism is required to determine if the subscriber network supports delivery of content to subscribers via multicast connectivity.

This document describes a mechanism whereby the subscriber's Internet service provider may signal in BGP the ability of the subscriber network to receive the content using multicast connectivity. This mechanism is called Multicast Distribution Reachability Signaling (MDRS).

For the purpose of this document we assume that a content provider consists of one or more Content Servers, and one or more Content Distribution Controllers. While this document assumes communication between Content Servers and Content Distribution Controllers, the procedures for implementing such communication is outside the scope of this document.

Content Servers are connected to one or more IP service providers (ISPs) that can offer both multicast and unicast connectivity service to the subscribers of the content provider. Content providers use these ISPs to deliver content to their subscribers.

Subscribers are connected to the Egress Routers (ERs) of the ISP. Note that the multicast connectivity service provided by the ISP extends all the way to the ERs. Such service could be provided by either deploying IP multicast natively, or with some tunneling mechanism like AMT, or by a combination of both within the ISP. However, between the ERs and the subscribers there may, or may not be multicast connectivity.

In the case where a particular subscriber of a given content provider does not have multicast connectivity to its ER, the content provider would use IP unicast service provided by the ISP to transmit the particular content to that subscriber.

[2.](#) Specification of Requirements

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](#) [[RFC2119](#)].

[2.1.](#) Multicast Distribution Reachability Signaling

Multicast distribution reachability signaling is responsible for giving a content provider, and specifically Content Distribution Controller(s) of the content provider the information of whether a particular subscriber of that content provider has multicast connectivity to an ER of an ISP that the content provider uses for distributing its content.

This document assumes that each ER can determine the multicast reachability status for each of the subscriber connected to that ER. Procedures by which an ER accomplishes this are outside the scope of this document.

To indicate whether a given ER has multicast reachability to a subscriber (be that either a native multicast or AMT) this document uses BGP as follows. An ER originates into IBGP routes for the subscribers connected to that ER for which the ER has multicast reachability. These routes are carried using BGP multi-protocol capabilities [[RFC4760](#)] with AFI 1 or 2, and the MCAST-REACH SAFI. The NLRI field in the MP_REACH_NLRI/MP_UNREACH_NLRI attribute of these routes contains subscribers' IP addresses encoded as IP address prefixes. The value of the AFI field in the MP_REACH_NLRI/MP_UNREACH_NLRI attribute of these routes determines whether subscribers' addresses are IPv4 or IPv6 (AFI 1 indicates IPv4 addresses, AFI 2 indicates IPv6 addresses).

A Content Distribution Controller, when it receives such routes, uses them to determine whether the content could be delivered to the

subscribers via the ISP who owns the ERs using the multicast connectivity service provided by the ISP.

To constrain the flow of BGP routes that carry multicast distribution reachability information such routes carry a particular Route Target (RT) Extended Community [[RFC4360](#)], and Content Distribution Controller(s) are provisioned to import routes with such a RT.

RTs carried by routes with AFI 1 and MCAST-REACH SAFI SHOULD NOT be re-used by routes with any other AFI and/or SAFI. Likewise, RTs carried by routes with AFI 2 and MCAST-REACH SAFI SHOULD NOT be re-used by routes with any other AFI and/or SAFI.

To facilitate such constrained distribution of multicast distribution reachability information one MAY use Constrained Route Distribution [[RFC4684](#)].

[3.](#) IANA Considerations

This document defines a new BGP Subsequent Address Family Identifier (SAFI) value, MCAST-REACH. The authors request assignment of a value from the First Come, First Served portion of this registry.

[4.](#) Security Considerations

TBD

[5.](#) Acknowledgements

The authors would like to thank Han Nguyen for his contributions to this document.

[6.](#) Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

- [RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", [RFC 4360](#), February 2006.
- [RFC4684] Marques, P., Bonica, R., Fang, L., Martini, L., Raszuk, R., Patel, K., and J. Guichard, "Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)", [RFC 4684](#), November 2006.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", [RFC 4760](#), January 2007.

Authors' Addresses

Huajin Jeng
AT&T

Phone:
Email: hj2387@att.com

Jeng, et al.

Expires August 8, 2014

[Page 5]

Internet-Draft

MDRS

February 2014

Jeffrey Haas
Juniper Networks
1194 N. Mathida Ave.
Sunnyvale, CA 94089
US

Email: jhaas@juniper.net

Yakov Rekhter
Juniper Networks
1194 N. Mathida Ave.
Sunnyvale, CA 94089
US

Email: yakov@juniper.net

Jeffrey (Zhaohui) Zhang
Juniper Networks
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
US

Email: zzhang@juniper.net