

Versions: 01

PIM WG
Internet-Draft
Intended status: Informational
Expires: August 16, 2013

J. Asghar
IJ. Wijnands
S. Krishnaswamy
A. Karan
Cisco Systems
V. Arya
Directv, Inc.

February 19, 2013

Explicit RPF Vector
draft-asghar-pim-explicit-rpf-vector-01

Abstract

This document describes a use of the Reverse Path Forwarding (RPF) Vector TLV as defined in [RPC 5496] to build multicast trees via an explicitly configured path sent in the PIM join.

Status of this Memo

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

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/lid-abstracts.html>

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

This Internet-Draft will expire April 15, 2013.

Copyright Notice

Copyright (c) 2012 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. Specification of Requirements	3
3. Soluton Requirements	3
4. Use of the Explicit RPF Vector	4
5. Explicit RPF Vector Attribute	4
6. Conflicting RPF Vectors	4
7. Explicit RPF Vector Attribute TLV Format	4
8. IANA Considerations	5
9. Security Considerations	5
10. Acknowledgments	6
11. Normative References	6
Authors' Addresses	7

1. Introduction

In some applications it might be useful to have a way to specify the explicit path along which the PIM join is propagated.

This document defines a new TLV in the PIM Join Attribute message [RFC5384] for specifying the explicit path.

The procedures in [RFC5496] define how a RPF vector can be used to influence the path selection in the absence of a route to the Source. However, the same procedures can be used to override a route to the Source when it exists. It is possible to include multiple RPF vectors in the stack where each router along the path will perform a unicast route lookup on the first vector in the attribute list. Once the router owning the address of the RPF vector is reached, following the procedures in [RFC5496], the RPF vector will be removed from the attribute list. This will result in a 'loosely' routed path based on the unicast reachability of the RPF vector(s). We call this loosely because we still depend on unicast routing reachability to the RPF Vector.

In some scenarios we don't want to rely on the unicast reachability to the RPF vector address and we want to build a path strictly based on the RPF vectors. In that case the RPF vector(s) represent a list of directly connected PIM neighbors along the path. For these vectors we MUST NOT do a unicast route lookup. We call these 'explicit' RPF vector addresses. If a router receiving an explicit RPF vector does not have a PIM neighbor matching the explicit RPF vector address it MUST NOT fall back to loosely routing the JOIN. Since the behavior of the explicit RPF vector differs from the loose RPF vector as defined [RFC5496], we're defining a new attribute called the explicit RPF Vector.

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 [RFC2119].

3. Solution Requirements

Some broadcast video transport networks use a multicast PIM live-live resiliency model for video delivery based on PIM SSM or PIM ASM. Live-Live implies using 2 active-active spatially diverse multicast trees to transport video flows from root to leaf multicast routers. The leaf multicast router receives 2 copies from the PIM multicast core and will replicate 1 copy towards the receivers [draft-mofrr-karan].

One of the main requirements of PIM live-live resiliency model is to ensure path-diversity of the active-active PIM trees in the core such that they do not intersect to avoid a single point of failure. IGP routed RPF paths of active-active PIM trees could be routed over the same transit router and create a single point of failure. It might be useful to have a way to specify the explicit path along which the PIM join is propagated.

How the explicit RPF vector stack is determined is outside the scope of this document. It may either be manually configured by the network operator or procedures may be implemented on the egress router to dynamically calculate the vector stack based on a link state database protocol, like OSPF or ISIS.

Due to the fact that the leaf router receives two copies of the multicast stream via two diverse paths, there is no need for PIM to repair the broken path immediately. It is up to the egress router to either wait for the broken path to be repaired or build a new explicit path using a new RPF vector stack. Which method is applied depends very much on how the vector stack was determined initially. Double failures are not considered and outside the scope of this document

4. Use of the PIM Explicit RPF Vector

Figure 1 provides an example multicast join path R4->R3->R6->R5->R2->R1, where the multicast JOIN is explicitly routed to the source hop-by-hop using the explicit RPF vector list.




Figure 1


5. Explicit RPF Vector Attribute

This draft uses vector attribute 4 for specifying an explicit RPF vector.

6. Conflicting RPF Vectors

It is possible that a PIM router has multiple downstream neighbors. If for the same multicast route there is inconsistency between the Explicit RPF Vector stacks provided by the downstream PIM neighbor, the procedures as documented in RFC5384 section 3.3.3 apply.

7. Explicit RPF Vector Attribute TLV Format



Authors' Addresses

Javed Asghar
Cisco Systems, Inc.
725, Alder Drive
Milpitas, CA 95035

Email: jasghar@cisco.com

IJsbrand Wijnands
Cisco Systems, Inc.
De kleetlaan 6a
Diegem 1831
Belgium

EMail: ice@cisco.com

Sowmya Krishnaswamy
Cisco Systems, Inc.
3750 Cisco Way
San Jose, CA 95134

EMail: sowkrish@cisco.com

Apoorva Karan
Cisco Systems, Inc.
3750 Cisco Way
San Jose, CA 95134

EMail: apoorva@cisco.com

Vishal Arya
DIRECTV Inc.
2230 E Imperial Hwy
El Segundo, CA 90245

Email: varya@directv.com