

Network Working Group
Internet-Draft
Updates: 5384 (if approved)
Intended status: Standards Track
Expires: August 10, 2013

S. Venaas
I. Kouvelas
J. Arango
Cisco Systems
February 6, 2013

Hierarchical Join/Prune Attributes
draft-venaas-pim-hierarchicaljoinattr-00.txt

Abstract

This document defines a hierarchical method of encoding Join attributes, providing a more efficient encoding when the same attribute values need to be specified for multiple sources in a PIM Join/Prune message.

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 10, 2013.

Copyright Notice

Copyright (c) 2013 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	3
2. Requirements Notation	4
3. Hierarchical Join/Prune Attribute Definition	5
4. PIM Address Encoding Types	8
5. Hierarchical Join/Prune Attribute Hello Option	9
6. Security Considerations	10
7. IANA Considerations	11
8. Acknowledgments	12
9. Normative References	13
Authors' Addresses	14

1. Introduction

PIM Join attributes as defined in [RFC5384] allow for specifying a set of attributes for each of the joined or pruned sources in a PIM Join/Prune message. Attributes must be separately specified for each individual source in the message. However, in some cases the same attributes and values need to be specified for some, or even all, the sources in the message. The attributes and their values then need to be repeated for each of the sources where they apply.

This document provides a hierarchical way of encoding attributes and their values in a Join/Prune message, so that if the same attribute and value is to apply for all the sources, it needs only be specified once in the message. Similarly, if all the sources in a specific group set share a specific attribute and value, it needs only be specified once for the entire group set.

This document updates [RFC5384] which defines an encoding to be used for Encoded-Source Addresses. This document extends this by specifying the same encoding type also for Encoded-Unicast and Encoded-Group formats. This document defines a new IANA registry for PIM encoding types which is to be used for all the fields in PIM messages where encoding types are used, replacing the old registry that is specific to Encoded-Source Addresses. The encoding type used for Join attributes is however still limited to be used in Join/Prune messages. Note that Join attributes, as they are referred to in [RFC5384], also apply to pruned sources in a Join/Prune message. Thus the more correct name Join/Prune attributes will be used throughout the rest of this document.


This document allows Join/Prune attributes to be specified in the Upstream Neighbor Address field, and also in the Multicast Group Address field, of a Join/Prune message. It defines how this is used to specify the same Join/Prune attribute and value for multiple sources. This document also introduces a new Hello Option to indicate support for the hierarchical encoding specified.

2. Requirements Notation

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. Hierarchical Join/Prune Attribute Definition

The format of a PIM Join/Prune message is defined in [RFC4601] as follows:



The message contains a single Upstream Neighbor Address, and one or more group sets. Each group set contains a Group Address and two source lists, the Joined Sources and the Pruned Sources. The Upstream Neighbor Address, the group addresses and the source addresses are all encoded in Encoded-Unicast format, Encoded-Group format and Encoded-Source format, respectively. In this document we make use of this to allow Join/Prune attributes in each of these addresses, using the encoding in Section 4.

For a Join/Prune message we define a hierarchy of Join/Prune attributes. At the highest level, that is the least specific, we have attributes that apply to every source in the message. These are encoded in the Upstream Neighbor Address. At the next more specific level we have attributes that apply to every source in a group set. They are encoded in a Group Address. And finally at the most specific level, we have attributes that just apply to a single source, encoded in the source address as defined in [RFC5384].

The complete set of attributes that apply to a given source is obtained by combining the message wide attributes, the attributes of the group set that the source belongs to, and the source specific attributes. However, if the same attribute is specified at multiple levels, then the one at the most specific level overrides the other instances of the attribute.

Note that Join/Prune attributes are still applied to sources as specified in [RFC5384]. This document does not change the meaning of any attributes, it is simply a more compact way of encoding an attribute when the same attribute and value applies to multiple sources.


4. PIM Address Encoding Types

Addresses in PIM messages are specified together with an address family and an encoding type. This applies to Encoded-Unicast, Encoded-Group and Encoded-Source addresses. The encoding types allow the address to be encoded according to different schemes. While it is possible to have the same encoding type value indicate different encodings depending on whether it is a Unicast, Group or Source address, it is simpler to have the same encoding type value indicate the same encoding independent of where it is used. This means that as currently defined, 0 means a native encoding, and 1 means there are Join/Prune attributes, encoded according to [RFC5384]. Even if the encoding type space is shared between the different address types (Encoded-Unicast, Encoded-Group and Encoded-Source), one could have a specific encoding apply to a specific address type if needed.

The current IANA PIM Encoded-Source Address Encoding Type Field registry should be changed into a PIM Address Encoding Type registry.

5. Hierarchical Join/Prune Attribute Hello Option

A PIM router indicates that it supports the mechanism specified in this document by including the Hierarchical Join/Prune Attribute Hello Option in its PIM Hello message. Note that it also needs to include the Join-Attribute Hello option as specified in [RFC5384]. The format of the Hierarchical Join/Prune Attribute Hello Option is defined to be:



OptionType = TBD, OptionLength = 0. Note that there is no option value included.

A PIM router MUST NOT send a Join/Prune message with Join/Prune attributes encoded in the Upstream Neighbor Address or any of the group addresses out any interface on which there is a PIM neighbor that has not included this option in its Hellos. Even a router that is not the upstream neighbor must be able to parse the message in order to do Join suppression or Prune overriding.

6. Security Considerations

This document specifies a more compact encoding of Join/Prune attributes. Use of the encoding has no impact on security.

7. IANA Considerations

The current PIM Encoded-Source Address Encoding Type Field registry should be changed into a PIM Address Encoding Type registry. The only required change is the name of the registry. The contents remain the same.

A new PIM Hello Option type needs to be assigned. The string TBD needs to be replaced with the permanently assigned value.

8. Acknowledgments

Acknowledgments to be added.

9. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4601] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.
- [RFC5384] Boers, A., Wijnands, I., and E. Rosen, "The Protocol Independent Multicast (PIM) Join Attribute Format", RFC 5384, November 2008.

Authors' Addresses

Stig Venaas
Cisco Systems
Tasman Drive
San Jose, CA 95134
USA

Email: stig@cisco.com

Isidor Kouvelas
Cisco Systems
Tasman Drive
San Jose, CA 95134
USA

Email: kouvelas@cisco.com

Jesus Arango
Cisco Systems
Tasman Drive
San Jose, CA 95134
USA

Email: jearango@cisco.com

