

Internet Engineering Task Force (IETF)
Request for Comments: 7887
Updates: 5384
Category: Standards Track
ISSN: 2070-1721

S. Venaas
J. Arango
Cisco Systems
I. Kouvelas
Arista Networks
June 2016

Hierarchical Join/Prune Attributes

Abstract

This document defines a hierarchical method of encoding Join/Prune attributes that provides a more efficient encoding when the same attribute values need to be specified for multiple sources in a PIM Join/Prune message. This document updates RFC 5384 by renaming the encoding type registry specified there.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7887>.

Copyright Notice

Copyright (c) 2016 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	3
3. Hierarchical Join/Prune Attribute Definition	3
4. PIM Address Encoding Types	6
5. Hierarchical Join/Prune Attribute Hello Option	6
6. Security Considerations	7
7. IANA Considerations	7
8. Normative References	7
Authors' Addresses	8

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 only needs to be specified once in the message. Similarly, if all the sources in a specific group set share a specific attribute and value, it only needs to be specified once for the entire group set.

This document extends [RFC5384] by specifying that the encoding type defined there also applies to Encoded-Unicast and Encoded-Group formats. This document also updates [RFC5384] by renaming the "PIM Encoded-Source Address Encoding Type Field" registry to "PIM Address Encoding Types". The content of the registry remains the same. The encoding type used for Join attributes is, however, still limited to use 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 defines 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 [RFC7761] 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 encoded in Encoded-Unicast format, Encoded-Group format, and Encoded-Source format, respectively. This document extends the use of the source address encoding defined in [RFC5384] to also apply to the Upstream Neighbor Address and the Group Address fields (see Section 4).

For a Join/Prune message, a hierarchy of Join/Prune attributes is defined. Attributes at the highest level, which is the least specific, apply to every source in the message. These are encoded in the Upstream Neighbor Address. Attributes at the next, more-specific level apply to every source in a group set. They are encoded in a Group Address. And finally, there are attributes that apply to a single source and are 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 the set of attributes and their values is formed before processing the attributes. Hence, a value that is invalid for a given type might override a valid value at a higher level.

As an example, say that for a given source, we have attributes T_1 with value V_1, T_2 with value V_2, and T_3 with value V_3. Also assume that in the Group Address of the source's group set, we have attributes T_1 with value V_6 and T_4 with value V_4. And assume that we in the Upstream Neighbor Address have encoded the attributes T_1 with value V_7, T_4 with value V_8, and T_5 with value V_5. The attributes applied to the given source will be T_1 with value V_1, T_2 with value V_2, T_3 with value V_3, T_4 with value V_4, and T_5 with value V_5. Here we have T_1 with different values at each level, so we use the value specified at the source level. Also, we have T_4 with different values at the group and message levels, so we use the value at the group level. Here it could be that V_1 is not a valid value for T_1, but it still overrides the values at the higher levels as we do not process the attributes until after forming the set.

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

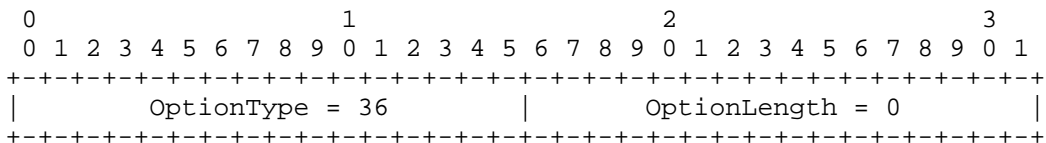
attribute when the same attribute and value applies to multiple sources, e.g., with the example above, we would have the exact same meaning if we instead had encoded all the attributes T1, ..., T5 with the respective values V1, ..., V5 in the source address.

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. An encoding type indicates how an address is encoded irrespective of address type, Encoded-Unicast, Encoded-Group, or Encoded-Source. It is possible that there will be future encoding types that do not apply to all address types though. This means that as currently defined, 0 is native encoding [RFC7761], and 1 is Join/Prune attributes encoding [RFC5384]. Note that as specified in [RFC5384], a type 1 Encoded Address MUST contain at least one Join/Prune attribute.

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. When this new Hello Option is included, it MUST also 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 = 36, 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 of 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 perform Join suppression and Prune override.

6. Security Considerations

This document specifies a more compact encoding of Join/Prune attributes. Use of the encoding has no impact on security aside from using the encoding in [RFC5384]. For instance, an attack with a forged message with certain attribute values is equally difficult independent of which encoding is used. If an attribute that applies to the entire message is wrong, then that may cause an issue for all the sources in the message. But without this encoding, one would instead include that attribute for every single source, and that would also cause an issue for all the sources in the message.

7. IANA Considerations

IANA has renamed the "PIM Encoded-Source Address Encoding Type Field" registry to "PIM Address Encoding Types".

The Hierarchical Join/Prune Attribute (36) has been added to the "PIM-Hello Options" registry.

8. Normative References

- [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>>.
- [RFC5384] Boers, A., Wijnands, I., and E. Rosen, "The Protocol Independent Multicast (PIM) Join Attribute Format", RFC 5384, DOI 10.17487/RFC5384, November 2008, <<http://www.rfc-editor.org/info/rfc5384>>.
- [RFC7761] Fenner, B., Handley, M., Holbrook, H., Kouvelas, I., Parekh, R., Zhang, Z., and L. Zheng, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", STD 83, RFC 7761, DOI 10.17487/RFC7761, March 2016, <<http://www.rfc-editor.org/info/rfc7761>>.

Authors' Addresses

Stig Venaas
Cisco Systems
Tasman Drive
San Jose, CA 95134
United States

Email: stig@cisco.com

Jesus Arango
Cisco Systems
Tasman Drive
San Jose, CA 95134
United States

Email: jeearango@cisco.com

Isidor Kouvelas
Arista Networks
5453 Great America Parkway
Santa Clara, CA 95054
United States

Email: kouvelas@arista.com