

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 14 September 2023

V. Kamath
VMware
R. Chokkanathapuram Sundaram
Cisco Systems, Inc.
R. Banthia
Apstra
A. Gopal
Cisco Systems, Inc.
13 March 2023

PIM Null-Register packing
draft-ietf-pim-null-register-packing-16

Abstract

In PIM-SM networks, PIM Null-Register messages are sent by the Designated Router (DR) to the Rendezvous Point (RP) to signal the presence of Multicast sources in the network. There are periodic PIM Null-Registers sent from the DR to the RP to keep the state alive at the RP as long as the source is active. The PIM Null-Register message carries information about a single Multicast source and group.

This document defines a standard to send multiple Multicast source and group information in a single PIM message. This document refers to the new messages as the PIM Packed Null-Register message and PIM Packed Register-Stop 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 <https://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 14 September 2023.

Copyright Notice

Copyright (c) 2023 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
1.1. Conventions used in this document	3
1.2. Terminology	3
2. Packed Null-Register Packing Capability	3
3. PIM Packed Null-Register message format	3
4. PIM Packed Register-Stop message format	4
5. Protocol operation	5
6. Operational Considerations	6
6.1. PIM Anycast RP Considerations	6
6.2. Interoperability between different versions	6
6.3. Disabling PIM Packed Message Support at RP and/or DR . .	6
7. Fragmentation Considerations	7
8. Security Considerations	7
9. IANA Considerations	7
10. Acknowledgments	8
11. Normative References	8
Authors' Addresses	9

1. Introduction

The DR periodically sends PIM Null-Registers to keep the state of existing multicast sources active on the RP. As the number of multicast sources increases, the number of PIM Null-Register messages that are sent also increases. This results in more PIM packet processing at the RP and the DR.

This document specifies a method to efficiently pack the content of multiple PIM Null-Register and Register-Stop messages [RFC7761] into a single message.

The document also discusses interoperability between PIM routers that support PIM Packed Null-Registers and PIM Packed Register-Stops and PIM routers that do not.

1.1. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

1.2. Terminology

RP: Rendezvous Point

DR: Designated Router

2. Packed Null-Register Packing Capability

The RP indicates its ability to receive PIM Packed Null-Register messages (Section 3) and send PIM Packed Register-Stop messages (Section 4) with a Packing Capability bit (P-bit) in the PIM Register-Stop message. The P-bit is allocated in Section 9.

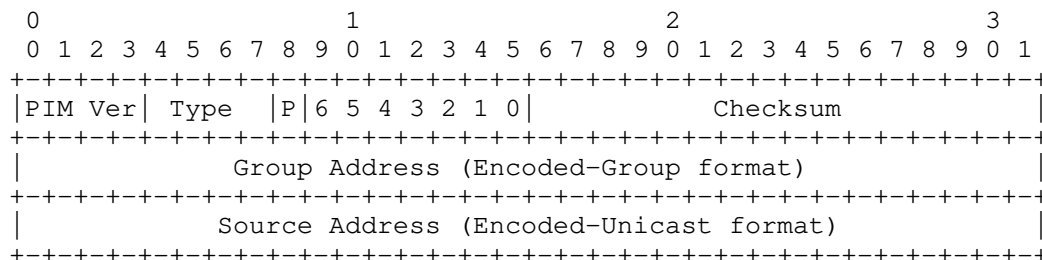


Figure 1: PIM Register-Stop message with Packing Capability option

The fields in the PIM Register-Stop message are defined in Section 4.9.4 of [RFC7761], and the common header in [I-D.venaas-pim-rfc8736bis].

Packing Capability bit (P-bit / Flag Bit TBD1): When set, it indicates the ability of the RP to receive PIM Packed Null-Register messages, and send PIM Packed Register-Stop messages.

3. PIM Packed Null-Register message format

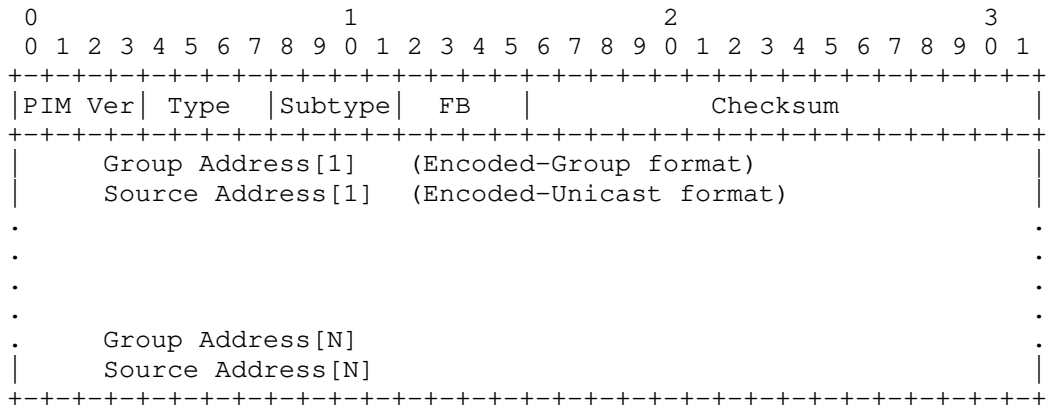


Figure 2: PIM Packed Null-Register message format

The fields in the PIM Packed Null-Register message are defined in Section 4.9.4 of [RFC7761], and the common header in [I-D.venaas-pim-rfc8736bis]

Type, Subtype: The PIM Packed Null-Register Type value TBD2. [I-D.venaas-pim-rfc8736bis]

N: The total number of records; A record consists of a Group Address and Source Address pair.

After parsing the PIM common header, individual records are then parsed one by one until the length of the PIM Packed Null-Register message. This length is inferred from the IP layer.

Sending or receiving a PIM Packed Null-Register message is the equivalent, for all purposes, of sending or receiving an individual Null-Register message for each record represented in the PIM Packed Null-Register message.

4. PIM Packed Register-Stop message format

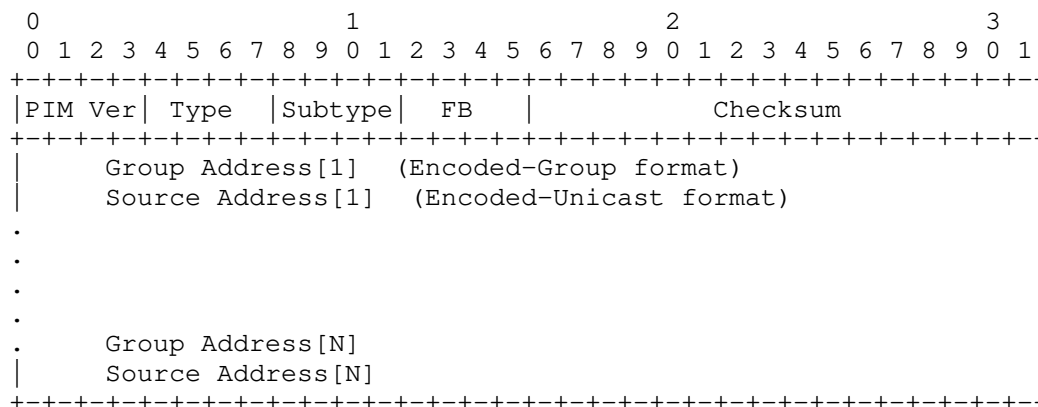


Figure 3: PIM Packed Register-Stop message format

The fields in the PIM Packed Register-Stop message are defined in Section 4.9.4 of [RFC7761], and the common header in [I-D.venaas-pim-rfc8736bis]

Type, Subtype: The PIM Packed Register-Stop Type TBD3

N: The total number of records; A record consists of a Group Address and Source Address pair.

After parsing the PIM common header, individual records are then parsed one by one until the length of the PIM Packed Register-Stop message. This length is inferred from the IP layer.

Sending or receiving a PIM Packed Register-Stop message is the equivalent, for all purposes, of sending or receiving an individual Null-Register message for each record represented in the PIM Packed Register-Stop.

5. Protocol operation

As specified in [RFC7761], the DR sends PIM Register messages towards the RP when a new source is detected.

When this feature is enabled/configured, an RP supporting this specification MUST set the P-bit (Flag bit TBD1) in all Register-Stop messages.

When a Register-Stop message with the P-bit set is received, the DR SHOULD send PIM Packed Null-Register messages (Section 3) to the RP instead of multiple Register messages with the N-bit set [RFC7761]. The DR MAY use a mixture of PIM Packed Null-Register messages and

Register messages. The decision is up to the implementation and out of the scope of this document. However, it is RECOMMENDED to stick to the PIM Packed Null-Register and PIM Packed Register-Stop formats as long as the RP and DR have the feature enabled.

The RP, after receiving a PIM Packed Null-Register message, SHOULD start sending PIM Packed Register-Stop messages (Section 4) to the corresponding DR instead of individual Register-Stop messages. The RP MAY use a mixture of PIM Packed Register-Stop messages and individual Register-Stop messages. The decision is up to the implementation and out of the scope of this document. However, it is RECOMMENDED to stick to the PIM Packed Null-Register and PIM Packed Register-Stop formats as long as the RP and DR have the feature enabled.

6. Operational Considerations

6.1. PIM Anycast RP Considerations

The PIM Packed Null-Register packet format should be enabled only if it is supported by all the routers in the Anycast-RP set [RFC4610]. This consideration applies to PIM Anycast RP with MSDP [RFC3446] as well.

6.2. Interoperability between different versions

A router (DR) can decide to use the PIM Packed Null-Register message format based on the Packing Capability received from the RP as part of the PIM Register-Stop. This ensures compatibility with routers that do not support processing of the new packet format. The Packing Capability information MUST be indicated by the RP via the PIM Register-Stop message sent to the DR. Thus, a DR will switch to the new packet format only when it learns that the RP is capable of handling the PIM Packed Null-Register messages.

Conversely, a DR that does not support the packed format can continue generating the PIM Null-Register as defined in [RFC7761] (Section 4.4).

6.3. Disabling PIM Packed Message Support at RP and/or DR

Consider a PIM RP router that supports PIM Packed Null-Registers and PIM Packed Register-Stops. In scenarios where this router now no longer supports this feature, for example, in case of a software downgrade, it will not send a PIM Register-Stop message to the DR in response to a PIM Packed Null-Register message.

When the DR switches to Data Registers from Null-Registers, it MUST start a `Packed_Register_Probe_Time` timer. If no PIM Packed Register-Stop or Register-Stop with the P-bit set is received within `Packed_Register_Probe_Time` seconds, the DR can decide that the RP no longer supports PIM Packed Null-Registers. The `Packed_Register_Probe_Time` timer is configurable; its default value is 60 seconds.

When `Packed_Register_Probe_Time` expires, The DR MAY also send an unpacked PIM Null-Register and check the PIM Register-Stop to see if the P-bit is set or not. If it is not set then the DR will continue sending unpacked PIM Null-Register messages.

In case the network manager disables the Packing Capability at the RP, or in other words, disables the feature from the RP, the router MUST NOT advertise the Packing Capability. However, an implementation MAY choose to still parse any packed registers if they are received. This may be particularly useful in the transitional period after the network manager disables it.

7. Fragmentation Considerations

As explained in Section 4.4.1 of [RFC7761], the DR may perform Path MTU Discovery to the RP before sending PIM Packed Null-Register messages. Similarly, the RP may perform Path MTU Discovery to the DR before sending PIM Packed Register-Stop messages. In both cases, the number of records in a message should be limited such that it can fit within the Path MTU.

8. Security Considerations

The Security Considerations from [RFC7761] apply to this document. In particular, the effect of forging a PIM Packed Null-Register or Register-Stop message would be amplified to all the records included instead of just one.

By forging a PIM Register-Stop message and setting the P-bit, an attacker can trigger the use of PIM Packed Null-Register messages by a DR thus creating unnecessary churn in the network.

9. IANA Considerations

When this document is published, IANA is asked to assign a Packing Capability bit (TBD1) in the PIM Register-Stop Common Header from the PIM Message Types registry.

When this document is published, IANA is asked to assign a PIM message type (TBD2) for the PIM Packed Null-Register from the PIM Message Types registry. The Flag Bits (0-3) for PIM message type (TBD2) are requested to be "Unassigned".

When this document is published, IANA is asked to assign a PIM message type (TBD3) for the PIM Packed Register-Stop from the PIM Message Types registry. The Flag Bits (0-3) for PIM message type (TBD3) are requested to be "Unassigned".

10. Acknowledgments

The authors would like to thank Stig Venaas, Alvaro Retana, Anish Peter, Zheng Zhang and Umesh Dudani for their helpful comments on the document.

11. 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, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [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, <<https://www.rfc-editor.org/info/rfc7761>>.
- [RFC4610] Farinacci, D. and Y. Cai, "Anycast-RP Using Protocol Independent Multicast (PIM)", RFC 4610, DOI 10.17487/RFC4610, August 2006, <<https://www.rfc-editor.org/info/rfc4610>>.
- [I-D.venaas-pim-rfc8736bis] Venaas, S. and A. Retana, "PIM Message Type Space Extension and Reserved Bits", Work in Progress, Internet-Draft, draft-venaas-pim-rfc8736bis-00, 1 March 2023, <<https://datatracker.ietf.org/doc/html/draft-venaas-pim-rfc8736bis-00>>.

[RFC3446] Kim, D., Meyer, D., Kilmer, H., and D. Farinacci, "Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)", RFC 3446, DOI 10.17487/RFC3446, January 2003, <<https://www.rfc-editor.org/info/rfc3446>>.

Authors' Addresses

Vikas Ramesh Kamath
VMware
3401 Hillview Ave
Palo Alto, CA 94304
United States of America
Email: vkamath@vmware.com

Ramakrishnan Chokkanathapuram Sundaram
Cisco Systems, Inc.
Tasman Drive
San Jose, CA 95134
United States of America
Email: ramaksun@cisco.com

Raunak Banthia
Apstra
333 Middlefield Rd STE 200
Menlo Park, CA 94025
United States of America
Email: rbanthia@apstra.com

Ananya Gopal
Cisco Systems, Inc.
Tasman Drive
San Jose, CA 95134
United States of America
Email: ananygop@cisco.com