

Network Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: September 23, 2021

V. Kamath  
VMware  
R. Chokkanathapuram Sundaram  
Cisco Systems, Inc.  
R. Banthia  
Apstra  
A. Gopal  
Cisco Systems, Inc.  
March 22, 2021

**PIM Null-Register packing**  
**draft-ietf-pim-null-register-packing-08**

Abstract

In PIM-SM networks PIM 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 Null-Register message, in a packed format. This document also discusses the interoperability between the PIM routers which do not understand the packed message format with multiple multicast source and group details.

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 September 23, 2021.

## Copyright Notice

Copyright (c) 2021 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 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

<a href="#">1.</a>	Introduction . . . . .	<a href="#">2</a>
<a href="#">1.1.</a>	Conventions used in this document . . . . .	<a href="#">3</a>
<a href="#">1.2.</a>	Terminology . . . . .	<a href="#">3</a>
<a href="#">2.</a>	Packed Null-Register Capability . . . . .	<a href="#">3</a>
<a href="#">3.</a>	PIM Packed Null-Register message . . . . .	<a href="#">4</a>
<a href="#">4.</a>	PIM Packed Register-Stop message format . . . . .	<a href="#">5</a>
<a href="#">5.</a>	Protocol operation . . . . .	<a href="#">6</a>
<a href="#">6.</a>	PIM Anycast RP considerations . . . . .	<a href="#">7</a>
<a href="#">7.</a>	PIM RP router version downgrade . . . . .	<a href="#">7</a>
<a href="#">8.</a>	Fragmentation consideration . . . . .	<a href="#">7</a>
<a href="#">9.</a>	Security Considerations . . . . .	<a href="#">7</a>
<a href="#">10.</a>	IANA Considerations . . . . .	<a href="#">8</a>
<a href="#">11.</a>	Acknowledgments . . . . .	<a href="#">8</a>
<a href="#">12.</a>	Normative References . . . . .	<a href="#">8</a>
	Authors' Addresses . . . . .	<a href="#">9</a>

## [1.](#) Introduction

PIM Null-Registers are sent by the DR periodically for Multicast streams to keep the states active on the RP, as long as the multicast source is alive. 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.

The control plane policing (COPP), monitors the packets that are processed by the control plane. The high rate at which Null-Registers are received at the RP can lead to COPP drops of Multicast PIM Null-Register messages. This draft proposes a method to efficiently pack multiple PIM Null-Registers [[RFC7761](#)] ([Section 4.4](#)) and Register-Stops [[RFC7761](#)] ([Section 3.2](#)) into a



single message as these packets anyway do not contain encapsulated data.

The draft also discusses interoperability with PIM routers that do not understand the new packet format.

### **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](#)] 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 Capability**

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

Conversely, a DR that does not support the new format can continue generating the PIM Null-Register the current way. To exchange the capability information in the Register-Stop message, the "reserved" field can be used to indicate this capability in those Register-Stop messages. One bit of the reserved field is used to indicate the "packing" capability (P bit). The rest of the bits in the "Reserved" field will be retained for future use.



Same as [\[RFC7761\]](#) ([Section 4.9.3](#))



Type, SubType:

The new packed Null-Register Type and SubType values TBD.  
[[RFC8736](#)]

Count:

The number of packed Null-Register records. A record consists of a Group Address and Source Address pair.

Group Address, Source Address:

Same as [[RFC7761](#)] ([Section 4.9.4](#))

#### 4. PIM Packed Register-Stop message format

The PIM Packed Register-Stop message includes a count to indicate the number of records that are present in the message.

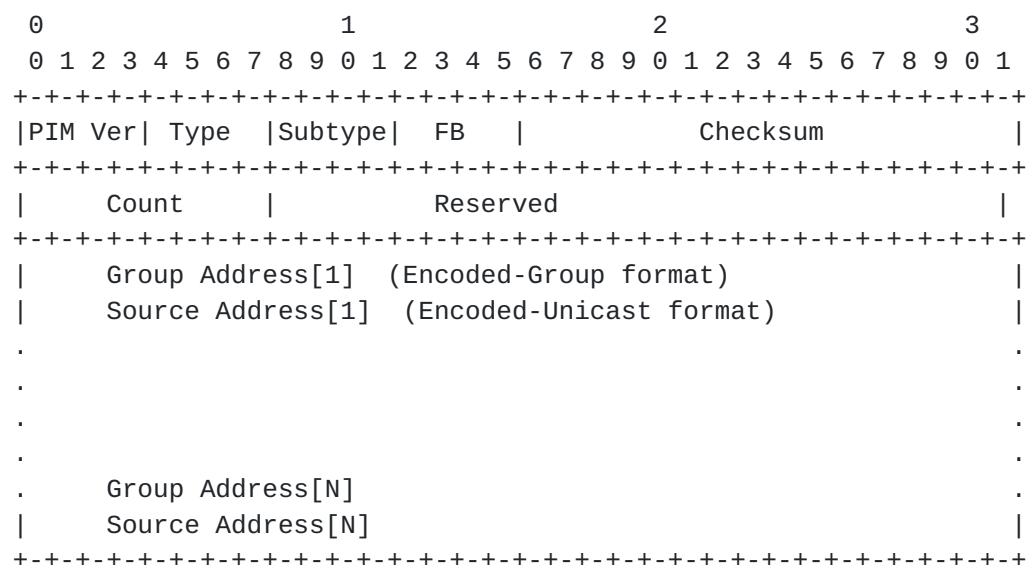


Figure 3: PIM Packed Register-Stop message format

PIM Version, Reserved, Checksum:

Same as [[RFC7761](#)] ([Section 4.9.4](#))

Type:

The new Register Stop Type and SubType values TBD

Count:



The number of packed Register-Stop records. A record consists of a Group Address and Source Address pair.

Group Address, Source Address:

Same as [\[RFC7761\]](#) ([Section 4.9.4](#))

## **5. Protocol operation**

The following combinations exist -

### **1. DR and RP both support the PIM Packed Null-Register format**

- \* As specified in [\[\[RFC7761\]\]](#), the DR sends PIM Register messages towards the RP when a new source is detected.
- \* An RP supporting this specification MUST set the P-bit in the corresponding Register-Stop messages.
- \* When a Register-Stop message with the P-bit set is received, the DR MAY send Packed Null-Register messages ([Section 3](#)) to the RP instead of multiple Register messages with the N-bit set ([\[\[RFC7761\]\]](#)).
- \* The RP, after receiving a Packed Null-Register message MAY start sending Packed Register-Stop messages ([Section 4](#)) to the corresponding DR instead of individual Register-Stop messages.

### **2. DR supports but RP does not support PIM Packed Null-Register format**

- \* As specified in [\[\[RFC7761\]\]](#), DR sends PIM Register towards the RP.
- \* RP sends a normal Register-Stop without any capability information.
- \* DR then sends Null-Registers in the old format. [\[\[RFC7761\]\]](#)

### **3. RP supports but DR doesn't support the new PIM Packed Null-Register format**

- \* As specified in [\[\[RFC7761\]\]](#), DR sends the PIM Register towards the RP.
- \* P sends a PIM Packed Register-Stop towards the DR that includes capability information.



- \* Since DR does not support the new format, it sends Null-Registers in the old format. [[[RFC7761](#)]]

## **6. PIM Anycast RP considerations**

The PIM Packed Null-Register format should be enabled only if it is supported by all PIM Anycast RP [[[RFC4610](#)]] members in the RP set for the RP address.

## **7. PIM RP router version downgrade**

Consider a PIM RP router that supports PIM Register Packing and then downgrades to a software version which does not support PIM Register Packing. The DR that sends the PIM Packed Null-Register message will not get a PIM Register-Stop message back. In such scenarios the DR can send an unpacked PIM Null-Register and check the PIM Register-Stop to see if the capability bit (P-bit) for PIM Packed Null-Register is set or not. If it is not set then the DR will continue sending unpacked PIM Null-Register messages.

## **8. Fragmentation consideration**

When building a PIM Packed Null-Register message or PIM Packed Register-Stop message, a router should include as many records as possible based on the path MTU towards RP, if path MTU discovery is done. Otherwise, the number of records should be limited by the MTU of the outgoing interface.

## **9. Security Considerations**

General Register messages security considerations from [RFC7761](#) apply. As mentioned in [RFC7761](#), PIM Null-Register messages and Register-Stop messages are forwarded by intermediate routers to their destination using normal IP forwarding. Without data origin authentication, an attacker who is located anywhere in the network may be able to forge a Null-Register or Register-Stop message. We next consider the effect of a forgery of each of these messages. By forging a Register message, an attacker can cause the RP to inject forged traffic onto the shared multicast tree.

By forging a Register-Stop message, an attacker can prevent a legitimate DR from registering packets to the RP. This can prevent local hosts on that LAN from sending multicast packets. The above two PIM messages are not changed by intermediate routers and need only be examined by the intended receiver. Thus, these messages can be authenticated end-to-end. Attacks on Register and Register-Stop messages do not apply to a PIM-SSM-only implementation, as these messages are not used in PIM-SSM.



There is another case where a spoofed Register-Stop can be sent to make it appear that is from the RP, and that the RP supports this new packed capability when it does not. This can cause Null-Registers to be sent to an RP that doesn't support this packed format. But standard methods to prevent spoofing should take care of this case. For example, uRPF can be used to filter out packets coming from the outside from addresses that belong to routers inside.

## **10. IANA Considerations**

This document requires the assignment of Capability bit (P-bit), flag bit 7 in the PIM Register-Stop message.

This document requires the assignment of 2 new PIM message types for the PIM Packed Null-Register and PIM Packed Register-Stop.

## **11. Acknowledgments**

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

## **12. 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>>.
- [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>>.
- [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>>.
- [RFC8736] Venaas, S. and A. Retana, "PIM Message Type Space Extension and Reserved Bits", [RFC 8736](#), DOI 10.17487/RFC8736, February 2020, <<https://www.rfc-editor.org/info/rfc8736>>.



Authors' Addresses

Vikas Ramesh Kamath  
VMware  
3401 Hillview Ave  
Palo Alto CA 94304  
USA

Email: vkamath@vmware.com

Ramakrishnan Chokkanathapuram Sundaram  
Cisco Systems, Inc.  
Tasman Drive  
San Jose CA 95134  
USA

Email: ramaksun@cisco.com

Raunak Banthia  
Apstra  
333 Middlefield Rd STE 200  
Menlo Park CA 94025  
USA

Email: rbanthia@apstra.com

Ananya Gopal  
Cisco Systems, Inc.  
Tasman Drive  
San Jose CA 95134  
USA

Email: ananygop@cisco.com

