

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 3 January 2026

M. Mishra
A. Budhiraja
Cisco Systems
H. Bidgoli
Nokia
2 July 2025

Sticky PIM DR
draft-mankamana-pim-pim-sticky-dr-00

Abstract

This draft proposes a mechanism to enhance the stability of Protocol Independent Multicast (PIM) Designated Router (DR) election by introducing the concept of a Sticky DR. In traditional PIM operations, the DR role on a LAN may change dynamically in response to router availability or priority changes, leading to potential multicast service disruptions or state reinitialization. The Sticky PIM DR approach aims to retain the previous DR role whenever feasible. By minimizing DR churn, this mechanism improves multicast session continuity, reduces unnecessary Join/Prune message propagation, and enhances operational stability in dense multicast deployments. The draft outlines the procedural extensions to existing PIM DR election, compatibility considerations, and configuration guidelines to support backward interoperability.

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 3 January 2026.

Copyright Notice

Copyright (c) 2025 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
2. Conventions Used in This Document	3
3. Control plane operations	3
3.1. Initial PIM DR election	3
3.2. Sticky PIM DR Procedures	4
3.3. Overriding PIM DR with Sticky DR	4
3.4. Overriding PIM DR with Sticky DR	5
4. Interoperability with Legacy PIM Routers	5
5. Security Considerations	6
6. IANA Considerations	6
7. Acknowledgments	6
8. Normative References	6
Authors' Addresses	7

1. Introduction

In Protocol Independent Multicast (PIM) [RFC7761], the Designated Router (DR) on a multi-access network plays a critical role in initiating PIM Join/Prune messages for multicast group membership and acting as a liaison for local receivers. The DR is elected based on the highest IP address or configured priority among routers on the subnet. However, this election is inherently dynamic and can result in frequent DR changes due to config change or new higher priority router joining the shared LAN.

Several operational deployments, including those in enterprise and service provider networks, have observed that frequent DR changes can lead to undesirable multicast behavior, including temporary multicast traffic loss, source registration issues (in PIM-SM), and unnecessary Join/Prune state transitions. These disruptions are particularly impactful in high-availability environments, such as IPTV, financial trading platforms, or real-time sensor networks.

This document proposes a s procedure for Sticky PIM DR behavior to introduces minimal extensions to the existing PIM DR election logic, is backward-compatible with existing routers, and preserves PIM' s core operational principles while improving robustness in the face of control-plane volatility.

2. 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.

3. Control plane operations


3.1. Initial PIM DR election

Let us consider a scenario with three PIM routers operating on a shared LAN segment. The routers are identified as Router 1, Router 2, and Router 3, with the following IP addresses and configured PIM priorities:

Router 1: IP address = 10.1.1.1, PIM Priority = 100

Router 2: IP address = 20.1.1.1, PIM Priority = 200

Router 3: IP address = 30.1.1.1, PIM Priority = 300



These routers participate in the standard PIM DR election process as defined in [RFC7761], where the router with the highest PIM priority is elected as the Designated Router (DR). In the event of a tie in priority, the router with the highest IP address is selected. With the procedure defined in [RFC7761] Router 3 becomes PIM DR.

3.2. Sticky PIM DR Procedures

Once the initial PIM DR election has been completed and the network is provisioned to support Sticky PIM DR behavior, the elected DR is permitted to advertise an updated, elevated DR priority value to its neighbors in order to retain its role across reboots or transient link failures.

Specifically, the DR will begin advertising a priority value of $(\text{MAX_PIM_DR} - 1)$, where MAX_PIM_DR is defined as $0xFFFFFFFF$. This mechanism ensures that the current DR maintains its role, even in the presence of routers that may later join the LAN with a higher configured priority.

For example, assume Router 3 becomes the DR through the standard election process. Upon activation of Sticky PIM DR behavior, Router 3 MUST begin advertising its PIM DR priority as $(0xFFFFFFFF - 1)$. As a result, even if a new router (e.g., with configured priority 400) subsequently joins the LAN, no DR re-election will be triggered, since Router 3 now advertises a higher effective priority.

This behavior allows the DR role to be "sticky" to the original elected router, providing stability and avoiding unnecessary multicast state churn during topology changes or restarts.

3.3. Overriding PIM DR with Sticky DR

In networks that support the concept of a Backup PIM Designated Router (Backup DR), this specification allows for enhanced coordination by assigning a specific priority value to the Backup DR role. In such cases, the router elected as the Backup DR MUST advertise a PIM DR Priority of $\text{MAX_PIM_DR} - 2$ (i.e., one value lower than the Sticky DR Priority of $\text{MAX_PIM_DR} - 1$).

The election of the Backup DR follows the standard PIM DR election rules as defined in [RFC7761], and all other operational procedures remain unchanged. The use of this distinct priority value allows other routers on the LAN to recognize the Backup DR without affecting the active DR selection process.

This approach ensures a predictable and deterministic Backup DR assignment, while maintaining compatibility with both traditional PIM routers and the sticky DR mechanism introduced in this document.

3.4. Overriding PIM DR with Sticky DR

There may be operational scenarios where the network operator requires the ability to override the Sticky PIM DR behavior and trigger a DR re-election. This need can arise due to a variety of reasons, such as the existing DR becoming non-operational, requiring a planned maintenance window, or necessitating network upgrades that involve role reassignment.

To support such scenarios, this document defines a configurable mechanism that allows the operator to force a DR change even when Sticky PIM DR is enabled. A new router can be configured with a user-defined PIM DR priority, along with an optional setting to trigger DR re-election. When this option is enabled, the router will temporarily advertise a PIM DR priority equal to MAX_PIM_DR (0xFFFFFFFF) to initiate the re-election process. Upon receiving this advertisement, the currently active Sticky DR MUST relinquish its DR role and release its acquired Sticky DR priority of (MAX_PIM_DR - 1).

Once the existing DR steps down, the initiating router will revert to advertising its originally configured PIM DR priority. The standard PIM DR election procedure is then executed among all routers on the LAN. After election completion, the newly elected DR regardless of which router wins will resume the Sticky DR behavior by advertising a priority of (MAX_PIM_DR - 1).

4. Interoperability with Legacy PIM Routers

This specification has been designed to ensure backward compatibility with existing PIM implementations that do not support Sticky DR behavior. In scenarios where some routers on the LAN are traditional PIM routers compliant with [RFC7761] and do not implement the extensions defined in this document, the behavior of the network remains consistent and predictable.

If a new router that supports this specification and is eligible to become the DR (based on configured priority) joins the LAN, it will begin advertising the elevated DR priority value as defined by this draft (i.e., MAX_PIM_DR - 1), following the Sticky DR procedure. Traditional routers, which honor DR priority per [RFC7761], will accept the advertised value and defer DR election accordingly. Thus, no disruption occurs in the election process, and interoperability is preserved.

Conversely, if a traditional PIM router that does not support this specification becomes the DR, it will operate based on standard [RFC7761] behavior and will not advertise any Sticky DR priority. In

this case, the sticky mechanism is effectively bypassed, but the DR election and forwarding functionality continue to work as expected, without violating protocol behavior.

Additionally, if a legacy router is manually configured with a reserved DR priority value (e.g., MAX_PIM_DR or (MAX_PIM_DR - 1)), the election process will naturally fall back to the default PIM DR election rules, using priority and IP address as defined in [RFC7761]. The Sticky DR behavior will not take effect unless explicitly supported by the router.

To assist with operational visibility, routers that implement this specification SHOULD log a warning if a reserved DR priority value is detected in the network from a device which never originated any other DR priority. This provides the operator with the opportunity to detect and correct any misconfiguration or unsupported behavior proactively.

5. Security Considerations

When it comes to general PIM message security, see [RFC7761].

This document does not introduce any new security consideration on top of what has been already covered by [RFC7761].

6. IANA Considerations

This document requires the assignment of a DR priority value as below.

MAX_PIM_DR: 0xFFFFFFFF

Sticky PIM DR Priority: 0xFFFFFFFF - 1

Sticky PIM Backup DR Priority: 0xFFFFFFFF - 2

7. Acknowledgments

The authors would like to thank Joya Neema for her valuable discussions and insightful comments that contributed to the development of this document.

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, <<https://www.rfc-editor.org/info/rfc2119>>.

- [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>>.
- [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>>.

Authors' Addresses

Mankamana Mishra
Cisco Systems
Tasman Drive
San Jose, CA 95134
United States of America
Email: mankamis@cisco.com

Anuj Budhiraja
Cisco Systems
Tasman Drive
San Jose, CA 95134
United States of America
Email: abudhira@cisco.com

Hooman Bidgoli
Nokia
March Road
Ottawa Ontario K2K 2T6
Canada
Email: hooman.bidgoli@nokia.com