

RTGWG
Internet-Draft
Updates: 9568 (if approved)
Intended status: Standards Track
Expires: 25 January 2026

G. Mirsky
Ericsson
J. Tantsura
NVIDIA
G. Mishra
Verizon Inc.
24 July 2025

Applicability of Bidirectional Forwarding Detection (BFD) for Multi-
point Networks in Virtual Router Redundancy Protocol (VRRP)
draft-ietf-rtgwg-vrrp-p2mp-bfd-13

Abstract

This document explores the applicability of Bidirectional Forwarding Detection (BFD) in multipoint networks to enable sub-second convergence in the Virtual Router Redundancy Protocol (VRRP) for determining the Active Router. Additionally, it defines extensions to bootstrap point-to-multipoint BFD sessions using an IPv4/IPv6 VRRP Advertisement message, and, thus, updates RFC 9568.

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 25 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
1.1. Conventions used in this document	3
1.1.1. Terminology	3
1.1.2. Requirements Language	3
2. Applicability of p2mp BFD	3
2.1. Multipoint BFD Encapsulation	4
3. IANA Considerations	4
4. Security Considerations	4
5. Acknowledgements	5
6. Normative References	5
Authors' Addresses	5

1. Introduction

The [RFC9568] is the current Virtual Router Redundancy Protocol (VRRP) specification for IPv4 and IPv6 networks. VRRPv3 allows for a faster switchover to a Backup Router. A router may be part of several Virtual Router Redundancy groups, such as Active in some and Backup in others. Supporting sub-second mode for VRRPv3 [RFC9568] for all these roles without specialized support in the data plane may prove challenging because of the increased load on the control plane. However, it may still be possible to deploy VRRP and provide sub-second detection of Active Router failure by Backup Routers.

Bidirectional Forwarding Detection (BFD) [RFC5880] had been originally defined to detect failure of point-to-point (p2p) paths: single-hop [RFC5881], multihop [RFC5883]. Single-hop BFD may enable a Backup router to detect a failure of the Active router within sub-seconds.

[RFC8562] extends [RFC5880] for multipoint and multicast networks, which matches the deployment scenarios for VRRP over the LAN segment. This document demonstrates how point-to-multipoint (p2mp) BFD can enable faster detection of the Active Router failure and thus minimize service disruption in a VRRP domain. It also defines the extension to VRRP [RFC9568] to bootstrap a VRRP Backup Router to join a p2mp BFD session.

1.1. Conventions used in this document

1.1.1. Terminology

BFD: Bidirectional Forwarding Detection

p2mp: Pont-to-Multipoint

VRRP: Virtual Router Redundancy Protocol

1.1.2. Requirements Language

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.

2. Applicability of p2mp BFD

[RFC8562] may provide an efficient and scalable solution for a fast-converging environment that uses the default route rather than dynamic routing. Each redundancy group presents itself as a p2mp BFD session, with its Active Router being the head and Backup Routers being the tails of the p2mp BFD session.

The Active Router, configured to use p2mp BFD to support faster convergence of VRRP, starts transmitting BFD control packets with IPvX address associated with the Virtual Router [RFC9568] as a source IP address and the value of the My Discriminator field ([RFC5880]) locally selected according to the following rules:

For IPv4, the Active Router uses 32 bits of the IPv4 source address as described in Section 5.1.1.1 of [RFC9568].

For IPv6, the Active Router uses the 32 least-significant bits of IPv6 source address, as described in Section 5.1.2.1 of [RFC9568].

A Backup Router demultiplexes p2mp BFD test sessions based on IPvX address associated with the Virtual Router that it has been configured with and the non-zero My Discriminator value, it deduces from the received VRRP Advertisement packet according to the rules listed above. When a Backup router detects the failure of the Active Router, according to Section 5.11 [RFC8562], it re-evaluates its role in the Virtual Router. As a result, the Backup Router may become the Active router of the given Virtual Router or continue as a Backup Router.

If the former is the case, then the new Active router will start transmitting p2mp BFD control packets using the Active Router IP address as the source IP address for p2mp BFD control packets and thus bootstraps a new p2mp BFD session on a Backup Router.

If the latter is the case, the Backup Router MUST close and remove the p2mp BFD session associated with the failed Active Router. The VRRP Advertisement packet from the new VRRP Active Router will bootstrap the new p2mp BFD session.

2.1. Multipoint BFD Encapsulation

The MultipointHead of p2mp BFD session when transmitting BFD control packet:

Set the source MAC address according to rules in Section 7.3 of [RFC9568];

MUST set TTL or Hop Limit value to 255 (Section 5 [RFC5881]). Similarly, all received BFD Control packets that are demultiplexed to the session MUST be discarded if the received TTL or Hop Limit is not equal to 255;

SHOULD use group address VRRP ('224.0.0.18' for IPv4 and 'FF02:0:0:0:0:0:0:12' for IPv6) as destination IP address;

MAY use network broadcast address for IPv4 or link-local all nodes multicast group for IPv6 as destination IP address;

MUST set destination UDP port value to 3784 when transmitting BFD control packets, as defined in [RFC8562];

Source UDP port value selection follows the rules defined in Section 4 of [RFC5881];

MUST use the Active Router IP address as the source IP address.

3. IANA Considerations

This document makes no requests for IANA allocations. This section may be deleted by RFC Editor.

4. Security Considerations

This document defines an alternative way, to the one defined in [RFC9568], to accelerate detecting a failure that affects VRRP functionality using p2mp BFD. The operation of either protocol is not changed.

Security considerations discussed in [RFC9568], [RFC5880], [RFC5881], and [RFC8562], apply to this document.

5. Acknowledgements

The authors appreciate comments and suggestions shared by Acee Lindem and Alexander "Sasha" Vainshtein that helped simplify the solution.

6. 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>>.
- [RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, DOI 10.17487/RFC5880, June 2010, <<https://www.rfc-editor.org/info/rfc5880>>.
- [RFC5881] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop)", RFC 5881, DOI 10.17487/RFC5881, June 2010, <<https://www.rfc-editor.org/info/rfc5881>>.
- [RFC5883] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for Multihop Paths", RFC 5883, DOI 10.17487/RFC5883, June 2010, <<https://www.rfc-editor.org/info/rfc5883>>.
- [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>>.
- [RFC8562] Katz, D., Ward, D., Pallagatti, S., Ed., and G. Mirsky, Ed., "Bidirectional Forwarding Detection (BFD) for Multipoint Networks", RFC 8562, DOI 10.17487/RFC8562, April 2019, <<https://www.rfc-editor.org/info/rfc8562>>.
- [RFC9568] Lindem, A. and A. Dogra, "Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6", RFC 9568, DOI 10.17487/RFC9568, April 2024, <<https://www.rfc-editor.org/info/rfc9568>>.

Authors' Addresses

Greg Mirsky
Ericsson
Email: gregimirsky@gmail.com

Jeff Tantsura
NVIDIA
Email: jefftant.ietf@gmail.com

Gyan Mishra
Verizon Inc.
Email: gyan.s.mishra@verizon.com