Internet Engineering Task Force (IETF)
Request for Comments: 9568
Obsoletes: 5798
Category: Standards Track
ISSN: 2070-1721

A. Lindem
LabN Consulting, L.L.C.
A. Dogra
Cisco Systems

April 2024

Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6

Abstract

This document defines version 3 of the Virtual Router Redundancy Protocol (VRRP) for IPv4 and IPv6. It obsoletes RFC 5798, which previously specified VRRP (version 3). RFC 5798 obsoleted RFC 3768, which specified VRRP (version 2) for IPv4. VRRP specifies an election protocol that dynamically assigns responsibility for a Virtual Router to one of the VRRP Routers on a LAN. The VRRP Router controlling the IPv4 or IPv6 address(es) associated with a Virtual Router is called the Active Router, and it forwards packets routed to these IPv4 or IPv6 addresses. Active Routers are configured with virtual IPv4 or IPv6 addresses, and Backup Routers infer the address family of the virtual addresses being advertised based on the IP protocol version. Within a VRRP Router, the Virtual Routers in each of the IPv4 and IPv6 address families are independent of one another and always treated as separate Virtual Router instances. The election process provides dynamic failover in the forwarding responsibility should the Active Router become unavailable. For IPv4, the advantage gained from using VRRP is a higher-availability default path without requiring configuration of dynamic routing or router discovery protocols on every end-host. For IPv6, the advantage gained from using VRRP for IPv6 is a quicker switchover to Backup Routers than can be obtained with standard IPv6 Neighbor Discovery mechanisms.

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 https://www.rfc-editor.org/info/rfc9568.

Table of Contents

1. Introduction

This document defines version 3 of the Virtual Router Redundancy Protocol (VRRP) for IPv4 and IPv6. It obsoletes [RFC5798], which previously specified VRRP (version 3). [RFC5798] obsoleted [RFC3768], which specified VRRP (version 2) for IPv4. VRRP specifies an election protocol that dynamically assigns responsibility for a Virtual Router (refer to Section 1.7) to one of the VRRP Routers on a LAN. The VRRP Router controlling the IPv4 or IPv6 address(es) associated with a Virtual Router is called the Active Router, and it forwards packets routed to these IPv4 or IPv6 addresses (except for packets addressed to these addresses as described in Section 8.3.1). VRRP Active Routers are configured with virtual IPv4 or IPv6 addresses, and Backup Routers infer the address family of the virtual addresses being advertised based on the IP protocol version. Within a VRRP Router, the Virtual Routers in each of the IPv4 and IPv6 address families are independent of one another and always treated as separate Virtual Router instances. The election process provides dynamic failover in the forwarding responsibility should the Active Router become unavailable.

VRRP provides a function similar to the proprietary protocols Hot Standby Router Protocol (HSRP) [RFC2281] and IP Standby Protocol [IPSTB].

1.1. Differences from RFC 5798

The following changes have been made from [RFC5798]:

1. The VRRP terminology has been updated to conform to inclusive language guidelines for IETF technologies. The IETF has designated the National Institute of Standards and Technology (NIST) document "Guidance for NIST Staff on Using Inclusive Language in Documentary Standards" [NISTIR8366] for its inclusive language guidelines.

2. The term for the VRRP Router assuming forwarding responsibility has been changed to "Active Router" to be consistent with IETF inclusive terminology. Additionally, inconsistencies in the terminology of [RFC5798] for both "Active Router" and "Backup Router" were corrected. Additionally, the undesirable term for attracting and dropping unreachable packets has been changed.

3. Errata pertaining to the state machines in Section 6 were corrected.

4. The checksum calculation in Section 5.2.8 has been clarified to specify precisely what is included and that it does not include the pseudo-header for IPv4.

5. When a VRRP advertisement is received from a lower priority VRRP

Router, the Active VRRP Router will immediately send a VRRP advertisement to assure learning bridges will bridge the packets to the correct Ethernet segment (refer to Section 6.4.3).

6. Appendices describing operation over legacy technologies (Fiber Distributed Data Interface (FDDI), Token Ring, and ATM LAN Emulation) were removed.

7. A recommendation was added indicating that IPv6 Unsolicited Neighbor Advertisements SHOULD be accepted by the Active and Backup Routers (Section 8.2.4).

8. Checking that the Maximum Advertisement Intervals match is recommended, although this will not result in the VRRP packet being dropped (Section 7.1).

9. Miscellaneous editorial changes were made for readability.

10. The IANA Considerations section was augmented to include all the IPv4/IPv6 multicast address allocations and Ethernet Media Access Control (MAC) address allocations.

1.2. A Note on Terminology

This document discusses both IPv4 and IPv6 operations, and with respect to the VRRP protocol, many of the descriptions and procedures are common. In this document, it would be less verbose to be able to refer to "IP" to mean either "IPv4 or IPv6". However, historically, the term "IP" often refers to IPv4. For this reason, in this specification, the term "IPvX" (where X is 4 or 6) is introduced to mean either "IPv4" or "IPv6". In this text, where the IP version matters, the appropriate term is used, and the use of the term "IP" is avoided.

1.3. IPv4

There are a number of methods that an IPv4 end-host can use to determine its first-hop router for a particular IPv4 destination. These include running (or snooping) a dynamic routing protocol such as Routing Information Protocol (RIP) [RFC2453] or OSPF version 2 [RFC2328], running an ICMP router discovery client [RFC1256], running DHCPv4 [RFC2131], or using a statically configured default route.

Running a dynamic routing protocol on every end-host may not be feasible for a number of reasons, including administrative overhead, processing overhead, security issues, or the lack of an implementation for a particular platform. Neighbor or router discovery protocols may require active participation by all hosts on a network, requiring large timer values to reduce protocol overhead associated with the protocol packet processing for each host. This can result in a significant delay in the detection of an unreachable router, and such a delay may introduce unacceptably long periods of unreachability for the default route.

The use of a manually configured default route (either via a static route or DHCPv4) is quite popular since it minimizes configuration and processing overhead on the end-host and is supported by virtually every IPv4 implementation. However, this creates a single point of failure. Loss of the default router results in a catastrophic event, isolating all end-hosts that are unable to detect an available alternate path.

The Virtual Router Redundancy Protocol (VRRP) is designed to eliminate the single point of failure inherent in a network utilizing default routing. VRRP specifies an election protocol that dynamically assigns responsibility for a Virtual Router to one of the

VRRP Routers on a LAN. The VRRP Router controlling the IPv4 address(es) associated with a Virtual Router is called the Active Router and forwards packets sent to these IPv4 addresses. The election process provides dynamic failover of the forwarding responsibility should the Active Router become unavailable. Any of the Virtual Router's IPv4 addresses on a LAN can then be used as the default first-hop router by end-hosts. The advantage gained from using VRRP is a higher availability default path without requiring configuration of dynamic routing or a router discovery protocol on every end-host.

1.4. IPv6

IPv6 hosts on a LAN will usually learn about one or more default routers by receiving Router Advertisements sent using the IPv6 Neighbor Discovery (ND) protocol [RFC4861]. The Router Advertisements are periodically multicast at a rate such that the hosts can take more than 10 seconds to learn the default routers on a LAN. They are not sent frequently enough to rely on the absence of the Router Advertisement to detect router failures.

The ND protocol includes a mechanism called Neighbor Unreachability Detection to detect the failure of a neighbor node (router or host) or the forwarding path to a neighbor. This is done by sending unicast ND Neighbor Solicitation messages to the neighbor node. To reduce the overhead of sending Neighbor Solicitations, they are only sent to neighbors to which the node is actively sending traffic and only after there has been no positive indication that the router is up for a period of time. Using the default parameters in ND, it can take a host more than 10 seconds to learn that a router is unreachable before it will switch to another default router. This delay would be very noticeable to users and cause some transport protocol implementations to time out.

While the Neighbor Unreachability Detection could be made quicker by configuring the timer intervals to be more aggressive (note that the current lower limit for this is 5 seconds), this would have the downside of significantly increasing the overhead of ND traffic, especially when there are many hosts all trying to determine the reachability of one or more routers.

The Virtual Router Redundancy Protocol for IPv6 provides a much faster switchover to an alternate default router than can be obtained using standard ND procedures. Using VRRP, a Backup Router can take over for a failed default router in around three seconds (using VRRP default parameters). This is done without any interaction with the hosts and a minimum amount of VRRP traffic.

1.5. 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.

1.6. Scope

The remainder of this document describes the features, design goals, and theory of operation of VRRP. The message formats, protocol processing rules, and state machine that guarantee convergence to a single Active Router are presented. Finally, operational issues related to MAC address mapping, handling of ARP messages, generation of ICMP redirect messages, and security issues are addressed.

1.7. Definitions

VRRP Router	A router running the Virtual Router Redundancy Protocol. It may participate as one or more Virtual Routers.
Virtual Router	An abstract object managed by VRRP that acts as a default router for hosts on a shared LAN. It consists of a Virtual Router Identifier and either a set of associated IPv4 addresses or a set of associated IPv6 addresses across a common LAN. A VRRP Router can serve as a Backup Router for one or more Virtual Routers.
Virtual Router Identifier	An integer value (1-255) identifying an instance of a Virtual Router on a LAN. Also referred by its acronym, VRID.
Virtual Router MAC Address	The multicast Ethernet MAC address used for VRRP advertisements for a VRID. Refer to Section 7.3.
IP Address Owner	The VRRP Router that has the Virtual Router's IPvX address(es) as real interface address(es). This is the router that, when up, will respond to packets addressed to one of these IPvX addresses for ICMP pings, TCP connection requests, etc.
Primary IP Address	In IPv4, an IPv4 address selected from the set of real interface addresses. One possible selection algorithm is to always select the first address. In IPv4, VRRP advertisements are always sent using the primary IPv4 address as the source of the IPv4 packet. In IPv6, the link-local address of the interface over which the packet is transmitted is used.
Forwarding Responsibility	The responsibility for forwarding packets sent to the IPvX address(es) associated with the Virtual Router. This includes receiving packets sent to the Virtual Router MAC address, forwarding these packets based on the local Routing Information Base (RIB) / Forwarding Information Base (FIB), answering ARP requests for the IPv4 address(es), and answering ND requests for the IPv6 address(es).
Active Router	The VRRP Router that is assuming the responsibility of forwarding packets sent to the IPvX address(es) associated with the Virtual Router, answering ARP requests for the IPv4 address(es), and answering ND requests for the IPv6 address(es). Note that if the IPvX address owner is available, then it will always be the Active Router.
Backup Router(s)	The set of VRRP Routers available to assume forwarding responsibility for a Virtual Router should the current Active Router fail.
Drop Route	A route installed in the Routing Information Base (RIB) that will result in traffic with a destination address that matches the route to

be dropped.

2. Required Features

This section describes the set of features that were considered mandatory and that guided the design of VRRP.

2.1. IPvX Address Backup

Backup of an IPvX address or addresses is the primary function of VRRP. When providing election of an Active Router and the additional functionality described below, the protocol should strive to:

* minimize the duration of unreachability,

* minimize the steady-state bandwidth overhead and processing complexity,

* function over a wide variety of multiaccess LAN technologies capable of supporting IPvX traffic,

* allow multiple Virtual Routers on a network for load-balancing, and

* support multiple logical IPvX subnets on a single LAN segment.

2.2. Preferred Path Indication

A simple model of Active Router election among a set of redundant routers is to treat each router with equal preference and claim victory after converging to any router as the Active Router. However, there are likely to be many environments where there is a distinct preference (or range of preferences) among the set of redundant routers. For example, this preference may be based upon access link cost or speed, router performance or reliability, or other policy considerations. The protocol should allow the expression of this relative path preference in an intuitive manner and guarantee Active Router convergence to the most preferred Virtual Router currently available.

2.3. Minimization of Unnecessary Service Disruptions

Once Active Router election has been performed, any unnecessary transition between Active and Backup Routers can result in a disruption of service. The protocol should ensure that, after Active Router election, no state transition is triggered by any Backup Router of equal or lower preference as long as the Active Router continues to function properly.

Some environments may find it beneficial to avoid the state transition triggered when a router that is preferred over the current Active Router becomes available. It may be useful to support an override of the immediate restoration to the preferred path.

2.4. Efficient Operation over Extended LANs

Sending IPvX packets, i.e., sending either IPv4 or IPv6, on a multiaccess LAN requires mapping from an IPvX address to a MAC address. The use of the Virtual Router MAC address in an extended LAN employing learning bridges can have a significant effect on the bandwidth overhead of packets sent to the Virtual Router. If the Virtual Router MAC address is never used as the source address in a link-level frame, then the MAC address location is never learned, resulting in flooding of all packets sent to the Virtual Router. To improve the efficiency in this environment, the protocol should do the following:

1. Use the Virtual Router MAC address as the source in a packet sent by the Active Router to trigger MAC learning.

2. Trigger a message immediately after transitioning to the Active Router to update MAC learning.

3. Trigger periodic messages from the Active Router to maintain the MAC address cache.

2.5. Sub-second Operation for IPv4 and IPv6

Sub-second detection of Active Router failure is needed in both IPv4 and IPv6 environments. Earlier work proposed that sub-second operation was for IPv6, and this specification leverages that earlier approach for both IPv4 and IPv6.

One possible problematic scenario that may occur when using a small Advertisement_Interval (refer to Section 6.1) is when a VRRP Router is generating more packets than it can transmit, and a queue builds up on the VRRP Router. When this occurs, it is possible that packets being transmitted onto the VRRP-protected LAN could see a larger queueing delay than the smallest Advertisement_Interval. In this case, the Active_Down_Interval (refer to Section 6.1) may be small enough that normal queuing delays might cause a Backup Router to conclude that the Active Router is down and, hence, promote itself to Active Router. Very shortly afterwards, the delayed VRRP packets from the original Active Router cause the VRRP Router to switch back to Backup Router. Furthermore, this process can repeat many times per second, causing a significant disruption of traffic. To mitigate this problem, giving VRRP packets priority on egress interface queues should be considered. If the Active Router observes that this is occurring, it SHOULD log the problem (subject to rate-limiting).

3. VRRP Overview

VRRP specifies an election protocol to provide the Virtual Router function described earlier. All protocol messaging is performed using either IPv4 or IPv6 multicast datagrams. Thus, the protocol can operate over a variety of multiaccess LAN technologies supporting IPvX multicast. Each link of a VRRP Virtual Router has a single well-known MAC address allocated to it. This document currently only details the mapping to networks using an IEEE 802 48-bit MAC address. The Virtual Router MAC address is used as the source in all periodic VRRP messages sent by the Active Router to enable MAC learning by Layer 2 (L2) bridges on an extended LAN.

A Virtual Router is defined by its Virtual Router Identifier (VRID) and a set of either IPv4 or IPv6 address(es). A VRRP Router may associate a Virtual Router with its real address on an interface. The scope of each Virtual Router is restricted to a single LAN. A VRRP Router may be configured with additional Virtual Router mappings and priority for Virtual Routers it is willing to back up. The mapping between the VRID and its IPvX address(es) must be coordinated among all VRRP Routers on a LAN.

There is no restriction against reusing a VRID with a different address mapping on different LANs, nor is there a restriction against using the same VRID number for a set of IPv4 addresses and a set of IPv6 addresses. However, these are two different Virtual Routers.

To minimize network traffic, only the Active Router for each Virtual Router sends periodic VRRP Advertisement messages. A Backup Router will not attempt to preempt the Active Router unless the Backup Router has a higher priority. This eliminates service disruption unless a more preferred path becomes available. It’s also possible

to administratively prohibit Active Router preemption attempts. The only exception is that a VRRP Router will always become the Active Router for any Virtual Router associated with address(es) it owns. If the Active Router becomes unavailable, then the highest-priority Backup Router will transition to the Active Router after a short delay, providing a controlled transition of Virtual Router responsibility with minimal service interruption.

The VRRP protocol design provides rapid transition from the Backup Router to the Active Router to minimize service interruption and incorporates optimizations that reduce protocol complexity while guaranteeing controlled Active Router transition for typical operational scenarios. These optimizations result in an election protocol with minimal runtime state requirements, minimal active protocol states, and a single message type and sender. The typical operational scenarios are defined to be two redundant routers and/or distinct path preferences for each router. A side effect when these assumptions are violated, i.e., more than two redundant paths with equal preference, is that duplicate packets may be forwarded for a brief period during Active Router election. However, the typical scenario assumptions are likely to cover the vast majority of deployments, loss of the Active Router is infrequent, and the expected duration for Active Router election convergence is quite small (< 4 seconds when using the default Advertisement_Interval and configurable to < 1/25 second). Thus, the VRRP optimizations represent significant simplifications in the protocol design while incurring an insignificant probability of brief network disruption.

4. Sample VRRP Networks

4.1. Sample VRRP Network 1

The following figure shows a simple network with two VRRP Routers implementing one Virtual Router.

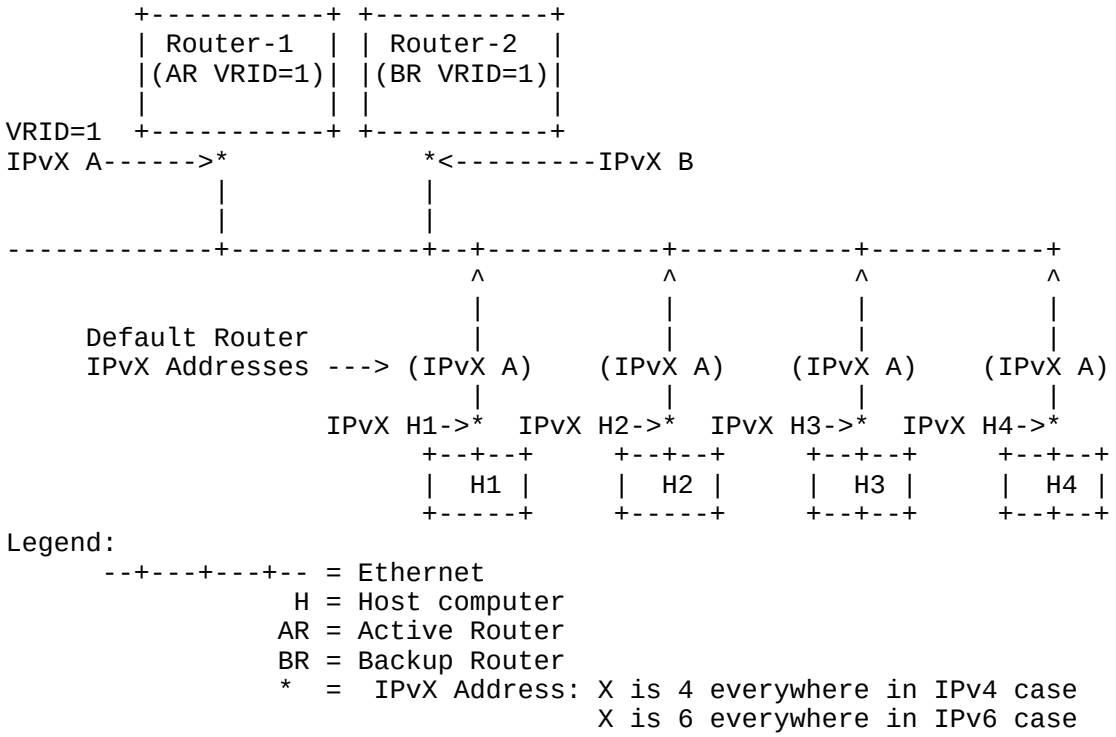

Figure 1: Sample VRRP Network 1

In the IPv4 case, i.e., IPvX is IPv4 everywhere in the figure, each router is permanently assigned an IPv4 address on the LAN interface

(Router-1 is assigned IPv4 A and Router-2 is assigned IPv4 B), and each host installs a default route (learned through DHCPv4 or via a configured static route) through one of the routers (in this example, they all use Router-1's IPv4 A).

In the IPv6 case, i.e., IPvX is IPv6 everywhere in the figure, each router has its own link-local IPv6 address on the LAN interface and a link-local IPv6 address per VRID that is shared with the other routers that serve the same VRID. Each host learns a default route from Router Advertisements through one of the routers (in this example, they all use Router-1's IPv6 Link-Local A).

In an IPv4 VRRP environment, each router supports reception and transmission for the exact same IPv4 address. Router-1 is said to be the IPv4 address owner of IPv4 A, and Router-2 is the IPv4 address owner of IPv4 B. A Virtual Router is then defined by associating a unique identifier (the VRID) with the address owned by Router-1.

In an IPv6 VRRP environment, each router will support transmission and reception for the IPv6 addresses associated with the VRID. Router-1 is said to be the IPv6 address owner of IPv6 A, and Router-2 is the IPv6 address owner of IPv6 B. A Virtual Router is then defined by associating a unique identifier (the VRID) with the address owned by Router-1.

Finally, in both the IPv4 and IPv6 cases, the VRRP protocol manages Virtual Router failover to a Backup Router.

The IPvX example above shows a Virtual Router configured to cover the IPvX address owned by Router-1 (VRID=1, IPvX_Address=A). When VRRP is enabled on Router-1 for VRID=1, it will assert itself as the Active Router, with priority = 255, since it is the IPvX address owner for the Virtual Router IPvX address. When VRRP is enabled on Router-2 for VRID=1, it will transition to the Backup Router, with priority = 100 (the default priority is 100), since it is not the IPvX address owner. If Router-1 should fail, then the VRRP protocol will transition Router-2 to the Active Router, temporarily taking over forwarding responsibility for IPvX A to provide uninterrupted service to the hosts.

Note that in both cases in this example, IPvX B is not backed up and it is only used by Router-2 as its interface address. In order to back up IPvX B, a second Virtual Router must be configured. This is shown in the next section.

4.2. Sample VRRP Network 2

The following figure shows a configuration with two Virtual Routers with the hosts splitting their traffic between them.

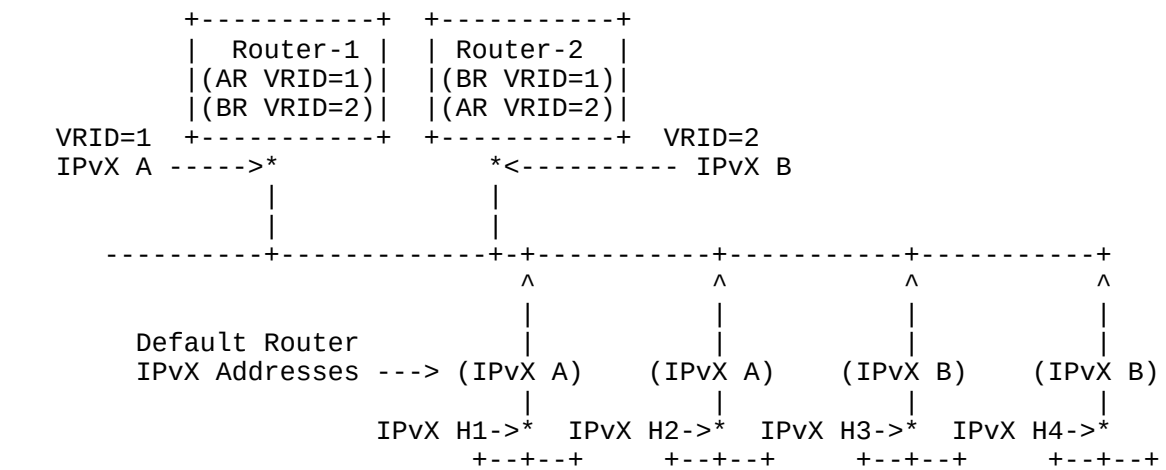

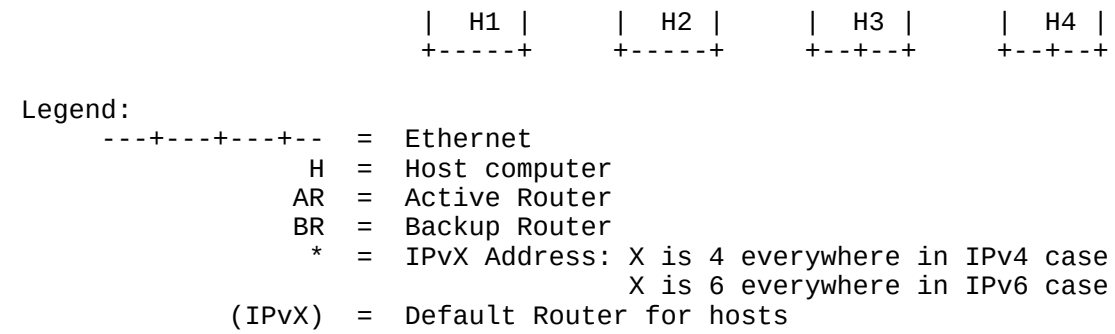

Figure 2: Sample VRRP Network 2

In the IPv4 example above, i.e., IPvX is IPv4 everywhere in the figure, half of the hosts have configured a static default route through Router-1's IPv4 A, and half are using Router-2's IPv4 B. The configuration of Virtual Router VRID=1 is exactly the same as in the first example (see Section 4.1), and a second Virtual Router has been added to cover the IPv4 address owned by Router-2 (VRID=2, IPv4_Address=B). In this case, Router-2 will assert itself as the Active Router for VRID=2, while Router-1 will act as a Backup Router. This scenario demonstrates a deployment providing load-splitting when both routers are available, while providing full redundancy for robustness.

In the IPv6 example above, i.e., IPvX is IPv6 everywhere in the figure, half of the hosts are using a default route through Router-1's IPv6 A, and half are using Router-2's IPv6 B. The configuration of Virtual Router VRID=1 is exactly the same as in the first example (see Section 4.1), and a second Virtual Router has been added to cover the IPv6 address owned by Router-2 (VRID=2, IPv6_Address=B). In this case, Router-2 will assert itself as the Active Router for VRID=2, while Router-1 will act as a Backup Router. This scenario demonstrates a deployment providing load-splitting when both routers are available while providing full redundancy for robustness.

Note that the details of load-balancing are out of scope of this document. However, in a case where the servers need different weights, it may not make sense to rely on Router Advertisements alone to balance the host traffic between the routers [RFC4311].

5. Protocol

The purpose of the VRRP Advertisement is to communicate to all VRRP Routers the priority, Maximum Advertisement Interval, and IPvX addresses of the Active Router associated with the VRID.

When VRRP is protecting an IPv4 address, VRRP packets are sent encapsulated in IPv4 packets. They are sent to the IPv4 multicast address assigned to VRRP.

When VRRP is protecting an IPv6 address, VRRP packets are sent encapsulated in IPv6 packets. They are sent to the IPv6 multicast address assigned to VRRP.

5.1. VRRP Packet Format

This section defines the format of the VRRP packet and the relevant fields in the IPvX header (in conjunction with the address family).

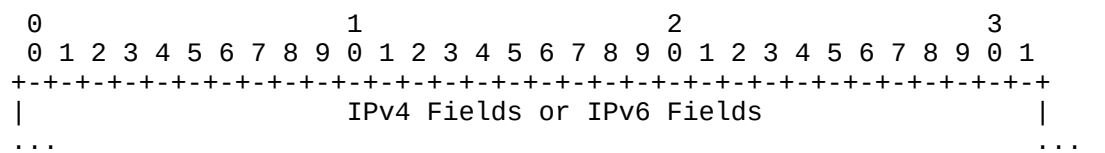

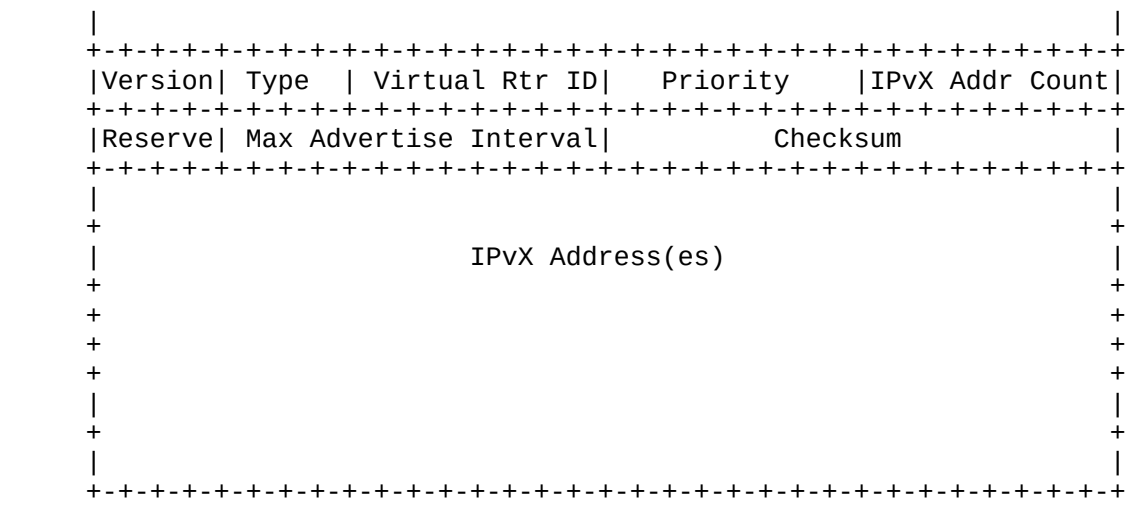

Figure 3: IPv4/IPv6 VRRP Advertisement Packet Format

5.1.1. IPv4 Field Descriptions

5.1.1.1. Source Address

This is the primary IPv4 address of the interface from which the packet is being sent.

5.1.1.2. Destination Address

The IPv4 multicast address as assigned by the IANA for VRRP is:

    224.0.0.18

This is a link-local scope multicast address. Routers MUST NOT forward a datagram with this destination address, regardless of its TTL.

5.1.1.3. TTL

The TTL MUST be set to 255. A VRRP Router receiving a packet with the TTL not equal to 255 MUST discard the packet [RFC5082].

5.1.1.4. Protocol

The IPv4 protocol number assigned by the IANA for VRRP is 112 (decimal).

5.1.2. IPv6 Field Descriptions

5.1.2.1. Source Address

This is the IPv6 link-local address of the interface from which the packet is being sent.

5.1.2.2. Destination Address

The IPv6 multicast address assigned by the IANA for VRRP is:

    ff02:0:0:0:0:0:0:12

This is a link-local scope multicast address. Routers MUST NOT forward a datagram with this destination address, regardless of its Hop Limit.

5.1.2.3. Hop Limit

The Hop Limit MUST be set to 255. A VRRP Router receiving a packet with the Hop Limit not equal to 255 MUST discard the packet [RFC5082].

5.1.2.4. Next Header

The IPv6 Next Header protocol assigned by the IANA for VRRP is 112 (decimal).

5.2. VRRP Field Descriptions

5.2.1. Version

The Version field specifies the VRRP protocol version of this packet. This document defines version 3.

5.2.2. Type

The Type field specifies the type of this VRRP packet. The only packet type defined in this version of the protocol is:

1 - ADVERTISEMENT

A packet with unknown type MUST be discarded.

5.2.3. Virtual Rtr ID (VRID)

The Virtual Rtr ID field identifies the Virtual Router for which this packet is reporting status.

5.2.4. Priority

The Priority field specifies sending the VRRP Router’s priority for the Virtual Router. Higher values indicate higher priority. This field is an 8-bit unsigned integer field.

The priority value for the VRRP Router that owns the IPvX address associated with the Virtual Router MUST be 255 (decimal).

VRRP Routers backing up a Virtual Router MUST use priority values between 1-254 (decimal). The default priority value for VRRP Routers backing up a Virtual Router is 100 (decimal). Refer to Section 8.3.2 for recommendations on setting the priority.

The priority value zero (0) has special meaning, indicating that the current Active Router has stopped participating in VRRP. This is used to trigger Backup Routers to quickly transition to the Active Router without having to wait for the current Active_Down_Interval (refer to Section 6.1).

5.2.5. IPvX Addr Count

The IPvX Addr Count field is the number of either IPv4 addresses or IPv6 addresses contained in this VRRP advertisement. The minimum value is 1. If the received count is 0, the VRRP advertisement MUST be ignored.

5.2.6. Reserve

The Reserve field MUST be set to zero on transmission and ignored on reception.

5.2.7. Maximum Advertisement Interval (Max Advertise Interval)

The Max Advertise Interval is a 12-bit field that indicates the time interval (in centiseconds) between advertisements. The default is

100 centiseconds (1 second).

Note that higher-priority Active Routers with slower transmission rates than their Backup Routers are unstable. This is because lower-priority Backup Routers configured to faster rates could join the LAN and decide they should be Active Routers before they have heard anything from the higher-priority Active Router with a slower rate. When this happens, it is temporary, i.e., once the lower-priority node does hear from the higher-priority Active Router, it will relinquish Active Router status.

5.2.8. Checksum

The Checksum field is used to detect data corruption in the VRRP message.

For both the IPv4 and IPv6 address families, the checksum is the 16-bit one's complement of the one's complement sum of the VRRP message. For computing the checksum, the Checksum field is set to zero. See [RFC1071] for more details.

For the IPv4 address family, the checksum calculation only includes the VRRP message starting with the Version field and ending after the last IPv4 address (refer to Section 5.2).

For the IPv6 address family, the checksum calculation also includes a prepended "pseudo-header", as defined in Section 8.1 of [RFC8200]. The Next Header field in the "pseudo-header" should be set to 112 (decimal) for VRRP.

5.2.9. IPvX Address(es)

This refers to one or more IPvX addresses associated with the Virtual Router. The number of addresses included is specified in the IPvX Addr Count field. These fields are used for troubleshooting misconfigured routers. If more than one address is sent, it is recommended that all routers be configured to send these addresses in the same order to simplify comparisons.

For IPv4 addresses, this refers to one or more IPv4 addresses that are backed up by the Virtual Router.

For IPv6, the first address MUST be the IPv6 link-local address associated with the Virtual Router.

This field contains either one or more IPv4 addresses or one or more IPv6 addresses. The address family of the addresses, IPv4 or IPv6 but not both, MUST be the same as the VRRP packet's IPvX header address family.

6. Protocol State Machine

6.1. Parameters per Virtual Router

VRID	Virtual Router Identifier. Configurable value in the range 1-255 (decimal). There is no default.
Priority	Priority value to be used by this VRRP Router in Active Router election for this Virtual Router. The value of 255 (decimal) is reserved for the router that owns the IPvX address associated with the Virtual Router. The value of 0 (zero) is reserved for the Active Router to indicate it is relinquishing

	responsibility for the Virtual Router. The range 1-254 (decimal) is available for VRRP Routers backing up the Virtual Router. Higher values indicate higher priorities. The default value is 100 (decimal).
IPv4_Addresses	One or more IPv4 addresses associated with this Virtual Router. Configured list of addresses with no default.
IPv6_Addresses	One or more IPv6 addresses associated with this Virtual Router. Configured list of addresses with no default. The first address MUST be the Link-Local address associated with the Virtual Router.
IPvX_Addresses	Refer to either the IPv4 or IPv6 address associated with this Virtual Router (see IPv4_Addresses and IPv6_Addresses above).
Advertisement_Interval	Time interval between VRRP Advertisements (centiseconds) sent by this Virtual Router. Default is 100 centiseconds (1 second).
Active_Adver_Interval	Advertisement interval contained in VRRP Advertisements received from the Active Router (in centiseconds). This value is saved by Virtual Routers in the Backup state and used to compute Skew_Time (as specified in Section 8.3.2) and Active_Down_Interval. The initial value is the same as Advertisement_Interval.
Skew_Time	Time to skew Active_Down_Interval in centiseconds. Calculated as: (((256 - Priority) * Active_Adver_Interval) / 256)
Active_Down_Interval	Time interval for the Backup Router to declare the Active Router down (centiseconds). Calculated as: (3 * Active_Adver_Interval) + Skew_Time
Preempt_Mode	Controls whether a (starting or restarting) higher-priority Backup Router preempts a lower-priority Active Router. Values are True to allow preemption and False to prohibit preemption. Default is True. Note: The exception is that the router that owns the IPvX address associated with the Virtual Router always preempts, independent of the setting of this flag.
Accept_Mode	Controls whether a Virtual Router in Active state will accept packets addressed to the address owner’s IPvX address as its own even if it is not the IPvX address owner. The default is

	False. Deployments that rely on, for example, pinging the address owner's IPvX address may wish to configure Accept_Mode to True. Note: IPv6 Neighbor Solicitations and Neighbor Advertisements MUST NOT be dropped when Accept_Mode is False.
Virtual_Router_MAC_Address	The MAC address used for the source MAC address in VRRP advertisements and advertised in ARP/ND messages as the MAC address to use for IPvX_Addresses.

6.2. Timers

Active_Down_Timer	Timer that fires when a VRRP Advertisement has not been received for Active_Down_Interval (Backup Routers only).
Adver_Timer	Timer that fires to trigger transmission of a VRRP Advertisement based on the Advertisement_Interval (Active Routers only).

6.3. State Transition Diagram

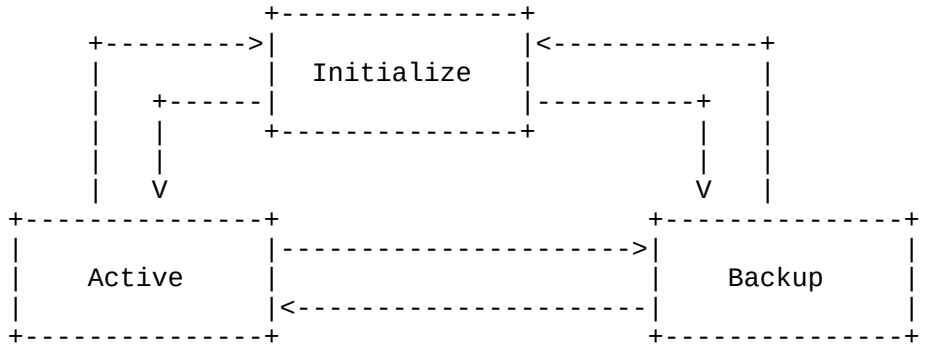

Figure 4: State Transition Diagram

6.4. State Descriptions

In the state descriptions below, the state names are identified by {state-name}, and the packets are identified by all-uppercase characters.

A VRRP Router implements an instance of the state machine for each Virtual Router in which it is participating.

6.4.1. Initialize

The purpose of this state is to wait for a Startup event, that is, an implementation-defined mechanism that initiates the protocol once it has been configured. The configuration mechanism is out of scope for this specification.

If a Startup event is received, then:

* If the Priority = 255, i.e., the router owns the IPvX address(es) associated with the Virtual Router, then:

  - Send an ADVERTISEMENT

  - If the protected IPvX address is an IPv4 address, then:

      o For each IPv4 address associated with the Virtual Router, broadcast a gratuitous ARP message containing the Virtual Router MAC address and with the target link-layer address set to the Virtual Router MAC address.

   - else // IPv6

      o For each IPv6 address associated with the Virtual Router, send an unsolicited ND Neighbor Advertisement with the Router Flag (R) set, the Solicited Flag (S) clear, the Override flag (O) set, the target address set to the IPv6 address of the Virtual Router, and the target link-layer address set to the Virtual Router MAC address.

   - endif // was protected address IPv4?

   - Set the Adver_Timer to Advertisement_Interval

   - Transition to the {Active} state

* else // Router is not the address owner

   - Set the Active_Adver_Interval to Advertisement_Interval

   - Set the Active_Down_Timer to Active_Down_Interval

   - Transition to the {Backup} state

* endif // was priority 255?

endif // Startup event was received

6.4.2. Backup

The purpose of the {Backup} state is to monitor the availability and state of the Active Router. The Solicited-Node multicast address [RFC4291] is referenced in the pseudocode below.

While in the {Backup} state, a VRRP Router MUST do the following:

* If the protected IPvX address is an IPv4 address, then:

   - It MUST NOT respond to ARP requests for the IPv4 address(es) associated with the Virtual Router.

* else // protected address is IPv6

   - It MUST NOT respond to ND Neighbor Solicitation messages for the IPv6 address(es) associated with the Virtual Router.

   - It MUST NOT send ND Router Advertisement messages for the Virtual Router.

* endif // was protected address IPv4?

* It MUST discard packets with a destination link-layer MAC address equal to the Virtual Router MAC address.

* It MUST NOT accept packets addressed to the IPvX address(es) associated with the Virtual Router.

* If a Shutdown event is received, then:

   - Cancel the Active_Down_Timer

- Transition to the {Initialize} state

* endif // Shutdown event received

* If the Active_Down_Timer fires, then:

  - Send an ADVERTISEMENT

  - If the protected IPvX address is an IPv4 address, then:

    o For each IPv4 address associated with the Virtual Router, broadcast a gratuitous ARP message containing the Virtual Router MAC address and with the target link-layer address set to the Virtual Router MAC address.

  - else // IPv6

    o Compute and join the Solicited-Node multicast address [RFC4291] for the IPv6 address(es) associated with the Virtual Router.

    o For each IPv6 address associated with the Virtual Router, send an unsolicited ND Neighbor Advertisement with the Router Flag (R) set, the Solicited Flag (S) clear, the Override flag (O) set, the target address set to the IPv6 address of the Virtual Router, and the target link-layer address set to the Virtual Router MAC address.

  - endif // was protected address IPv4?

  - Set the Adver_Timer to Advertisement_Interval

  - Transition to the {Active} state

* endif // Active_Down_Timer fired

* If an ADVERTISEMENT is received, then:

  - If the Priority in the ADVERTISEMENT is 0, then:

    o Set the Active_Down_Timer to Skew_Time

  - else // priority non-zero

    o If Preempt_Mode is False, or if the Priority in the ADVERTISEMENT is greater than or equal to the local Priority, then:

      + Set the Active_Adver_Interval to the Max Advertise Interval contained in the ADVERTISEMENT

      + Recompute the Skew_Time

      + Recompute the Active_Down_Interval

      + Set the Active_Down_Timer to Active_Down_Interval

    o else // preempt was true and priority was less than the local priority

      + Discard the ADVERTISEMENT

    o endif // preempt test

  - endif // was priority 0?

* endif // was advertisement received?

endwhile // {Backup} state

6.4.3. Active

While in the {Active} state, the router functions as the forwarding router for the IPvX address(es) associated with the Virtual Router.

Note that in the {Active} state, the Preempt_Mode Flag is not considered.

While in the {Active} state, a VRRP Router MUST do the following:

* If the protected IPvX address is an IPv4 address, then:
  - It MUST respond to ARP requests for the IPv4 address(es) associated with the Virtual Router.
* else // IPv6
  - It MUST be a member of the Solicited-Node multicast address for the IPv6 address(es) associated with the Virtual Router.
  - It MUST respond to ND Neighbor Solicitation messages (with the Router Flag (R) set) for the IPv6 address(es) associated with the Virtual Router.
  - It MUST send ND Router Advertisements for the Virtual Router.
  - If Accept_Mode is False:
    o It MUST NOT drop IPv6 Neighbor Solicitations and Neighbor Advertisements.
* endif // IPv4?
* It MUST forward packets with a destination link-layer MAC address equal to the Virtual Router MAC address.
* It MUST accept packets addressed to the IPvX address(es) associated with the Virtual Router if it is the IPvX address owner or if Accept_Mode is True. Otherwise, it MUST NOT accept these packets.
* If a Shutdown event is received, then:
  - Cancel the Adver_Timer
  - Send an ADVERTISEMENT with Priority = 0
  - Transition to the {Initialize} state
* endif // shutdown received
* If the Adver_Timer fires, then:
  - Send an ADVERTISEMENT
  - Reset the Adver_Timer to Advertisement_Interval
* endif // advertisement timer fired
* If an ADVERTISEMENT is received, then:
  - If the Priority in the ADVERTISEMENT is 0, then:

- o Send an ADVERTISEMENT

- o Reset the Adver_Timer to Advertisement_Interval

- - else // priority was non-zero

- o If the Priority in the ADVERTISEMENT is greater than the local Priority or the Priority in the ADVERTISEMENT is equal to the local Priority and the primary IPvX address of the sender is greater than the local primary IPvX address (based on an unsigned integer comparison of the IPvX addresses in network byte order), then:

- + Cancel Adver_Timer

- + Set the Active_Adver_Interval to the Max Advertise Interval contained in the ADVERTISEMENT

- + Recompute the Skew_Time

- + Recompute the Active_Down_Interval

- + Set the Active_Down_Timer to Active_Down_Interval

- + Transition to the {Backup} state

- o else // new Active Router logic

- + Discard the ADVERTISEMENT

- + Send an ADVERTISEMENT immediately to assert the {Active} state to the sending VRRP Router and to update any learning bridges with the correct Active VRRP Router path.

- o endif // new Active Router detected

- - endif // was priority zero?

* endif // advert received

endwhile // in {Active} state

Note: VRRP packets are transmitted with the Virtual Router MAC address as the source MAC address to ensure that learning bridges correctly determine the LAN segment to which the Virtual Router is attached.

7. Sending and Receiving VRRP Packets

7.1. Receiving VRRP Packets

The following functions must be performed when a VRRP packet is received:

* If the received packet is an IPv4 packet, then:

- It MUST verify that the IPv4 TTL is 255.

* else // IPv6 VRRP packet received

- It MUST verify that the IPv6 Hop Limit is 255.

* endif

* It MUST verify that the VRRP version is 3.

* It MUST verify that the VRRP packet type is 1 (ADVERTISEMENT).

* It MUST verify that the received packet contains the complete VRRP packet (including fixed fields and the IPvX address).

* It MUST verify the VRRP checksum.

* It MUST verify that the VRID is configured on the receiving interface and the local router is not the IPvX address owner (Priority = 255 (decimal)).

If any one of the above checks fails, the receiver MUST discard the packet, SHOULD log the event (subject to rate-limiting), and MAY indicate via network management that an error occurred.

A receiver SHOULD also verify that the Max Advertise Interval in the received VRRP packet matches the Advertisement_Interval configured for the VRID. Instability can occur with differing intervals (refer to Section 5.2.7). If this check fails, the receiver SHOULD log the event (subject to rate-limiting) and MAY indicate via network management that a misconfiguration was detected.

A receiver MAY also verify that "IPvX Addr Count" and the list of IPvX address(es) match the IPvX address(es) configured for the VRID. If this check fails, the receiver SHOULD log (subject to rate-limiting) the event and MAY indicate via network management that a misconfiguration was detected.

7.2. Transmitting VRRP Packets

The following operations MUST be performed when transmitting a VRRP packet:

* Fill in the VRRP packet fields with the appropriate Virtual Router configuration state

* Compute the VRRP checksum

* Set the source MAC address to the Virtual Router MAC address

* If the protected address is an IPv4 address, then:

  - Set the source IPv4 address to the interface’s primary IPv4 address

* else // IPv6

  - Set the source IPv6 address to the interface’s link-local IPv6 address

* endif

* Set the IPvX protocol to VRRP

* Send the VRRP packet to the VRRP IPvX multicast group

Note: VRRP packets are transmitted with the Virtual Router MAC address as the source MAC address to ensure that learning bridges correctly determine the LAN segment to which the Virtual Router is attached.

7.3. Virtual Router MAC Address

The Virtual Router MAC address associated with a Virtual Router is an

IEEE 802 MAC address [RFC9542] in the following format:

IPv4 case: 00-00-5E-00-01-{VRID} (in hex, in network byte order)

The first three octets are derived from the IANA's Organizationally Unique Identifier (OUI). The next two octets (00-01) indicate the address block assigned to the VRRP protocol for the IPv4 protocol. {VRID} is the Virtual Router Identifier. This mapping provides for up to 255 IPv4 VRRP Routers on a LAN.

IPv6 case: 00-00-5E-00-02-{VRID} (in hex, in network byte order)

The first three octets are derived from the IANA's OUI. The next two octets (00-02) indicate the address block assigned to the VRRP protocol for the IPv6 protocol. {VRID} is the Virtual Router Identifier. This mapping provides for up to 255 IPv6 VRRP Routers on a LAN.

7.4. IPv6 Interface Identifiers

[RFC8064] specifies that [RFC7217] be used as the default scheme for generating a stable address in IPv6 Stateless Address Autoconfiguration (SLAAC) [RFC4862]. The Virtual Router MAC MUST NOT be used for the Net_Iface parameter used in the Interface Identifier (IID) derivation algorithms in [RFC7217] and [RFC8981].

This VRRP specification describes how to advertise and resolve the VRRP Router's IPv6 link-local address and other associated IPv6 addresses into the Virtual Router MAC address.

8. Operational Issues

8.1. IPv4

8.1.1. ICMP Redirects

ICMP redirects can be used normally when VRRP is running among a group of routers. This allows VRRP to be used in environments where the topology is not symmetric.

The IPv4 source address of an ICMP redirect should be the address that the end-host used when making its next-hop routing decision. If a VRRP Router is acting as the Active Router for Virtual Router(s) containing address(es) it does not own, then it must determine to which Virtual Router the packet was sent when selecting the redirect source address. One method to deduce the Virtual Router used is to examine the destination MAC address in the packet that triggered the redirect.

It may be useful to disable redirects for specific cases where VRRP is being used to load-share traffic among a number of routers in a symmetric topology.

8.1.2. Host ARP Requests

When a host sends an ARP request for one of the Virtual Router IPv4 addresses, the Active Router MUST respond to the ARP request with an ARP response that indicates the Virtual Router MAC address for the Virtual Router. Note that the source address of the Ethernet frame of this ARP response is the physical MAC address of the physical router. The Active Router MUST NOT respond with its physical MAC address in the ARP response. This allows the host to always use the same MAC address, regardless of the current Active Router.

When a VRRP Router restarts or boots, it SHOULD NOT send any ARP messages using its physical MAC address for an IPv4 address for which

it is the IPv4 address owner (as defined in Section 1.7), and it should only send ARP messages that include Virtual Router MAC addresses.

This entails the following:

* When configuring an interface, Active Routers SHOULD broadcast a gratuitous ARP message containing the Virtual Router MAC address for each IPv4 address on that interface.

* At system boot, when initializing interfaces for VRRP operation, gratuitous ARP messages MUST be delayed until both the IPv4 address and the Virtual Router MAC address are configured.

* When, for example, Secure Shell (SSH) access to a particular VRRP Router is required, an IPv4 address known to belong to that router SHOULD be used.

8.1.3. Proxy ARP

If Proxy ARP is to be used on a VRRP Router, then the VRRP Router MUST advertise the Virtual Router MAC address in the Proxy ARP message. Doing otherwise could cause hosts to learn the real MAC address of the VRRP Router.

8.2. IPv6

8.2.1. ICMPv6 Redirects

ICMPv6 redirects can be used normally when VRRP is running among a group of routers [RFC4443]. This allows VRRP to be used in environments where the topology is not symmetric, e.g., the VRRP Routers do not connect to the same destinations.

The IPv6 source address of an ICMPv6 redirect SHOULD be the address that the end-host used when making its next-hop routing decision. If a VRRP Router is acting as the Active Router for Virtual Router(s) containing address(es) it does not own, then it has to determine to which Virtual Router the packet was sent when selecting the redirect source address. A method to deduce the Virtual Router used is to examine the destination MAC address in the packet that triggered the redirect.

8.2.2. ND Neighbor Solicitation

When a host sends an ND Neighbor Solicitation message for a Virtual Router IPv6 address, the Active Router MUST respond to the ND Neighbor Solicitation message with the Virtual Router MAC address for the Virtual Router. The Active Router MUST NOT respond with its physical MAC address. This allows the host to always use the same MAC address, regardless of the current Active Router.

When an Active Router sends an ND Neighbor Solicitation message for a host's IPv6 address, the Active Router MUST include the Virtual Router MAC address for the Virtual Router if it sends a source link-layer address option in the Neighbor Solicitation message. It MUST NOT use its physical MAC address in the source link-layer address option.

When a VRRP Router restarts or boots, it SHOULD NOT send any ND messages with its physical MAC address for the IPv6 address it owns and it should only send ND messages that include Virtual Router MAC addresses.

This entails the following:

* When configuring an interface, Active Routers SHOULD send an unsolicited ND Neighbor Advertisement message containing the Virtual Router MAC address for the IPv6 address on that interface.

* At system boot, when initializing interfaces for VRRP operation, all ND Router Advertisements, ND Neighbor Advertisements, and ND Neighbor Solicitation messages MUST be delayed until both the IPv6 address and the Virtual Router MAC address are configured.

Note that on a restarting Active Router where the VRRP protected address is an interface address, i.e., the address owner, Duplicate Address Detection may fail, as the Backup Router MAY answer that it owns the address. One solution is to not run Duplicate Address Detection in this case.

8.2.3. Router Advertisements

When a Backup VRRP Router has become the Active Router for a Virtual Router, it is responsible for sending Router Advertisements for the Virtual Router, as specified in Section 6.4.3. The Backup Routers MUST be configured to send the same Router Advertisement options as the address owner.

Router Advertisement options that advertise special services, e.g., Home Agent Information Option, that are present in the address owner SHOULD NOT be sent by the address owner unless the Backup Routers are prepared to assume these services in full and have a complete and synchronized database for this service.

8.2.4. Unsolicited Neighbor Advertisements

A VRRP Router acting as either an IPv6 Active Router or Backup Router SHOULD accept Unsolicited Neighbor Advertisements and update the corresponding neighbor cache [RFC4861]. Since these are sent to the IPv6 all-nodes multicast address (ff02::1) [RFC4861] or the IPv6 all-routers multicast address (ff02::2), they will be received. Unsolicited Neighbor Advertisements are sent both in the case where the link-level addresses change [RFC4861] and for gratuitous neighbor discovery by first-hop routers [RFC9131]. Additional configuration may be required in order for Unsolicited Neighbor Advertisements to update the corresponding neighbor cache.

8.3. IPvX

8.3.1. Potential Forwarding Loop

If it is not the address owner, a VRRP Router SHOULD NOT forward packets addressed to the IPvX address for which it becomes the Active Router. Forwarding these packets would result in unnecessary traffic. Also, in the case of LANs that receive packets they transmit, this can result in a forwarding loop that is only terminated when the IPvX TTL expires.

One mechanism for VRRP Routers to avoid these forwarding loops is to add/delete a host Drop Route for each non-owned IPvX address when transitioning to/from the Active state.

8.3.2. Recommendations Regarding Setting Priority Values

A priority value of 255 designates a particular router as the "IPvX address owner" for the VRID. VRRP Routers with priority 255 will, as soon as they start up, preempt all lower-priority routers. For a VRID, only a single VRRP Router on the link SHOULD be configured with priority 255. If multiple VRRP Routers advertising priority 255 are detected, the condition SHOULD be logged (subject to rate-limiting). If no VRRP Router has this priority, and preemption is disabled, then

no preemption will occur.

In order to avoid two or more Backup Routers simultaneously becoming Active Routers after the previous Active Router fails or is shut down, all Virtual Routers SHOULD be configured with different priorities and with sufficient differences in the priorities so that lower priority Backup Routers do not transition to the Active state before receiving an advertisement from the highest priority Backup Router when it transitions to the Active Router. If multiple VRRP Routers advertising the same priority are detected, this condition MAY be logged as a warning (subject to rate-limiting).

Since the Skew_Time is reduced as the priority is increased, faster convergence can be obtained by using a higher priority for the preferred Backup Router. However, with multiple Backup Routers, the priorities should have sufficient differences, as previously recommended.

8.4. VRRPv3 and VRRPv2 Interoperation

8.4.1. Assumptions

1. VRRPv2 and VRRPv3 interoperation is optional.

2. Mixing VRRPv2 and VRRPv3 should only be done when transitioning from VRRPv2 to VRRPv3. Mixing the two versions should not be considered a permanent solution.

8.4.2. VRRPv3 Support of VRRPv2 Interoperation

As mentioned above, this support is intended for upgrade scenarios and is NOT RECOMMENDED for permanent deployments.

An implementation MAY implement a configuration flag that tells it to listen for and send both VRRPv2 and VRRPv3 advertisements.

When a Virtual Router is configured this way and is the Active Router, it MUST send both types at the configured rate, even if it is sub-second.

When a Virtual Router is configured this way and is the Backup Router, it MUST time out based on the rate advertised by the Active Router. In the case of a VRRPv2 Active Router, this means it MUST translate the timeout value it receives (in seconds) into centiseconds. Also, a Backup Router SHOULD ignore VRRPv2 advertisements from the current Active Router if it is also receiving VRRPv3 packets from it. It MAY report when a VRRPv3 Active Router is not sending VRRPv2 packets, as this suggests they don’t agree on whether they’re supporting VRRPv2 interoperation.

8.4.2.1. Interoperation Considerations

8.4.2.1.1. Slow, High-Priority Active Routers

See also Section 5.2.7, "Maximum Advertisement Interval (Max Advertise Interval)".

The VRRPv2 Active Router interacting with a sub-second VRRPv3 Backup Router is the most important example of this.

A VRRPv2 implementation SHOULD NOT be given a higher priority than a VRRPv2 or VRRPv3 implementation with which it is interoperating if the VRRPv2 or VRRPv3 router’s advertisement rate is sub-second.

8.4.2.1.2. Overwhelming VRRPv2 Backups

It seems possible that a VRRPv3 Active Router sending at centisecond rates could potentially overwhelm a VRRPv2 Backup Router with potentially non-deterministic results.

In this upgrade case, a deployment should initially run the VRRPv3 Active Routers with lower frequencies, e.g., 100 centiseconds, until the VRRPv2 routers are upgraded. Then, once the deployment has verified that VRRPv3 is working properly, the VRRPv2 support may be disabled and the desired sub-second rates may be configured.

9. Security Considerations

VRRP for IPvX does not currently include any type of authentication. Earlier versions of the VRRP specification included several types of authentication, ranging from no authentication to strong authentication. Operational experience and further analysis determined that these did not provide sufficient security to overcome the vulnerability of misconfigured secrets, causing multiple Active Routers to be elected. Due to the nature of the VRRP protocol, even if VRRP messages are cryptographically protected, it does not prevent hostile nodes from behaving as if they are an Active Router, creating multiple Active Routers. Authentication of VRRP messages could have prevented a hostile node from causing all properly functioning routers from going into the Backup state. However, having multiple Active Routers can cause as much disruption as no routers, which authentication cannot prevent. Also, even if a hostile node could not disrupt VRRP, it can disrupt ARP/ND and create the same effect as having all routers go into the Backup state.

Some L2 switches provide the capability to filter out, for example, ARP and/or ND messages from end-hosts on a switch-port basis. This mechanism could also filter VRRP messages from switch ports associated with end-hosts and can be considered for deployments with untrusted hosts.

It should be noted that these attacks are not worse and are a subset of the attacks that any node attached to a LAN can do independently of VRRP. The kind of attacks a malicious node on a LAN can perform include:

* promiscuously receiving packets for any router’s MAC address,

* sending packets with the router’s MAC address as the source MAC address in the L2 header to tell the L2 switches to send packets addressed to the router to the malicious node instead of the router,

* sending redirects to tell hosts to send their traffic somewhere else,

* sending unsolicited ND replies,

* answering ND requests, etc.

All of these can be done independently of implementing VRRP. VRRP does not add to these vulnerabilities, and most of these vulnerabilities are addressed independently, e.g., SEcure Neighbor Discovery (SEND) [RFC3971].

VRRP includes a mechanism (setting IPv4 TTL or IPv6 Hop Limit to 255 and checking the value on receipt) that protects against VRRP packets being injected from another remote network [RFC5082]. This limits most vulnerabilities to attacks on the local network.

VRRP does not provide any confidentiality. Confidentiality is not necessary for the correct operation of VRRP, and there is no

information in the VRRP messages that must be kept secret from other nodes on the LAN.

In the context of IPv6 operation, if SEND is deployed, VRRP is compatible with the "trust anchor" and "trust anchor or CGA" modes of SEND [RFC3971]. The SEND configuration needs to give the Active and Backup Routers the same prefix delegation in the certificates so that Active and Backup Routers advertise the same set of subnet prefixes. However, the Active and Backup Routers should have their own key pairs to avoid private key sharing.

Also in the context of IPv6 operation, it is RECOMMENDED that the link-level security guidelines in Section 2.3 of [RFC9099] be followed.

10. IANA Considerations

IANA has updated all IANA registry references to [RFC5798] to references to RFC 9568, i.e., this document. The individual IANA references are listed below.

The value 112 is assigned to VRRP in the "Assigned Internet Protocol Numbers" registry.

In the "Local Network Control Block (224.0.0.0 - 224.0.0.255 (224.0.0/24))" registry of the "IPv4 Multicast Address Space Registry" [RFC5771], IANA has assigned the IPv4 multicast address 224.0.0.18 for VRRP.

In the "Link-Local Scope Multicast Addresses" registry of the "IPv6 Multicast Address Space Registry" [RFC3307], IANA has assigned the IPv6 link-local scope multicast address ff02:0:0:0:0:0:0:12 for VRRP for IPv6.

In the "IANA MAC ADDRESS BLOCK" registry [RFC9542], IANA has assigned blocks of Ethernet unicast addresses as follows (in hexadecimal):

Addresses	Usage	Reference
00-01-00 to 00-01-FF	VRRP (Virtual Router Redundancy Protocol)	RFC 9568
00-02-00 to 00-02-FF	VRRP IPv6 (Virtual Router Redundancy Protocol IPv6)	RFC 9568

Table 1

11. References

11.1. 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>.

[RFC3307] Haberman, B., "Allocation Guidelines for IPv6 Multicast Addresses", RFC 3307, DOI 10.17487/RFC3307, August 2002, <https://www.rfc-editor.org/info/rfc3307>.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <https://www.rfc-editor.org/info/rfc4291>.

[RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", STD 89, RFC 4443, DOI 10.17487/RFC4443, March 2006, <https://www.rfc-editor.org/info/rfc4443>.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <https://www.rfc-editor.org/info/rfc4861>.

[RFC5082] Gill, V., Heasley, J., Meyer, D., Savola, P., Ed., and C. Pignataro, "The Generalized TTL Security Mechanism (GTSM)", RFC 5082, DOI 10.17487/RFC5082, October 2007, <https://www.rfc-editor.org/info/rfc5082>.

[RFC5771] Cotton, M., Vegoda, L., and D. Meyer, "IANA Guidelines for IPv4 Multicast Address Assignments", BCP 51, RFC 5771, DOI 10.17487/RFC5771, March 2010, <https://www.rfc-editor.org/info/rfc5771>.

[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>.

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <https://www.rfc-editor.org/info/rfc8200>.

[RFC9542] Eastlake 3rd, D., Abley, J., and Y. Li, "IANA Considerations and IETF Protocol and Documentation Usage for IEEE 802 Parameters", BCP 141, RFC 9542, DOI 10.17487/RFC9542, April 2024, <https://www.rfc-editor.org/info/rfc9542>.

11.2. Informative References

[IPSTB] Higginson, P. and M. Shand, "Development of Router Clusters to Provide Fast Failover in IP Networks", Digital Technical Journal, Volume 9, Number 3, 1997.

[NISTIR8366] National Institute of Standards and Technology (NIST), "Guidance for NIST Staff on Using Inclusive Language in Documentary Standards,", NISTIR 8366, DOI 10.6028/NIST.IR.8366, April 2021, <https://doi.org/10.6028/NIST.IR.8366>.

[RFC1071] Braden, R., Borman, D., and C. Partridge, "Computing the Internet checksum", RFC 1071, DOI 10.17487/RFC1071, September 1988, <https://www.rfc-editor.org/info/rfc1071>.

[RFC1256] Deering, S., Ed., "ICMP Router Discovery Messages", RFC 1256, DOI 10.17487/RFC1256, September 1991, <https://www.rfc-editor.org/info/rfc1256>.

[RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, DOI 10.17487/RFC2131, March 1997, <https://www.rfc-editor.org/info/rfc2131>.

[RFC2281] Li, T., Cole, B., Morton, P., and D. Li, "Cisco Hot Standby Router Protocol (HSRP)", RFC 2281, DOI 10.17487/RFC2281, March 1998, <https://www.rfc-editor.org/info/rfc2281>.

[RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328,
DOI 10.17487/RFC2328, April 1998,
<https://www.rfc-editor.org/info/rfc2328>.

[RFC2338] Knight, S., Weaver, D., Whipple, D., Hinden, R., Mitzel,
D., Hunt, P., Higginson, P., Shand, M., and A. Lindem,
"Virtual Router Redundancy Protocol", RFC 2338,
DOI 10.17487/RFC2338, April 1998,
<https://www.rfc-editor.org/info/rfc2338>.

[RFC2453] Malkin, G., "RIP Version 2", STD 56, RFC 2453,
DOI 10.17487/RFC2453, November 1998,
<https://www.rfc-editor.org/info/rfc2453>.

[RFC3768] Hinden, R., Ed., "Virtual Router Redundancy Protocol
(VRRP)", RFC 3768, DOI 10.17487/RFC3768, April 2004,
<https://www.rfc-editor.org/info/rfc3768>.

[RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander,
"SEcure Neighbor Discovery (SEND)", RFC 3971,
DOI 10.17487/RFC3971, March 2005,
<https://www.rfc-editor.org/info/rfc3971>.

[RFC4311] Hinden, R. and D. Thaler, "IPv6 Host-to-Router Load
Sharing", RFC 4311, DOI 10.17487/RFC4311, November 2005,
<https://www.rfc-editor.org/info/rfc4311>.

[RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless
Address Autoconfiguration", RFC 4862,
DOI 10.17487/RFC4862, September 2007,
<https://www.rfc-editor.org/info/rfc4862>.

[RFC5798] Nadas, S., Ed., "Virtual Router Redundancy Protocol (VRRP)
Version 3 for IPv4 and IPv6", RFC 5798,
DOI 10.17487/RFC5798, March 2010,
<https://www.rfc-editor.org/info/rfc5798>.

[RFC7217] Gont, F., "A Method for Generating Semantically Opaque
Interface Identifiers with IPv6 Stateless Address
Autoconfiguration (SLAAC)", RFC 7217,
DOI 10.17487/RFC7217, April 2014,
<https://www.rfc-editor.org/info/rfc7217>.

[RFC8064] Gont, F., Cooper, A., Thaler, D., and W. Liu,
"Recommendation on Stable IPv6 Interface Identifiers",
RFC 8064, DOI 10.17487/RFC8064, February 2017,
<https://www.rfc-editor.org/info/rfc8064>.

[RFC8981] Gont, F., Krishnan, S., Narten, T., and R. Draves,
"Temporary Address Extensions for Stateless Address
Autoconfiguration in IPv6", RFC 8981,
DOI 10.17487/RFC8981, February 2021,
<https://www.rfc-editor.org/info/rfc8981>.

[RFC9099] Vyncke, ., Chittimaneni, K., Kaeo, M., and E. Rey,
"Operational Security Considerations for IPv6 Networks",
RFC 9099, DOI 10.17487/RFC9099, August 2021,
<https://www.rfc-editor.org/info/rfc9099>.

[RFC9131] Linkova, J., "Gratuitous Neighbor Discovery: Creating
Neighbor Cache Entries on First-Hop Routers", RFC 9131,
DOI 10.17487/RFC9131, October 2021,
<https://www.rfc-editor.org/info/rfc9131>.

[VRRP-IPv6]
Hinden, R. and J. Cruz, "Virtual Router Redundancy

Protocol for IPv6", Work in Progress, Internet-Draft, draft-ietf-vrrp-ipv6-spec-08, 5 March 2007, <https://datatracker.ietf.org/doc/html/draft-ietf-vrrp-ipv6-spec-08>.

Acknowledgments

The IPv6 text in this specification is based on [RFC2338]. The authors of [RFC2338] are S. Knight, D. Weaver, D. Whipple, R. Hinden, D. Mitzel, P. Hunt, P. Higginson, M. Shand, and A. Lindem.

The authors of [VRRP-IPv6] would also like to thank Erik Nordmark, Thomas Narten, Steve Deering, Radia Perlman, Danny Mitzel, Mukesh Gupta, Don Provan, Mark Hollinger, John Cruz, and Melissa Johnson for their helpful suggestions.

The IPv4 text in this specification is based on [RFC3768]. The authors of that specification would like to thank Glen Zorn, Michael Lane, Clark Bremer, Hal Peterson, Tony Li, Barbara Denny, Joel Halpern, Steve M. Bellovin, Thomas Narten, Rob Montgomery, Rob Coltun, Radia Perlman, Russ Housley, Harald Alvestrand, Ned Freed, Ted Hardie, Bert Wijnen, Bill Fenner, and Alex Zinin for their comments and suggestions.

Thanks to Steve Nadas for his work merging/editing [RFC3768] and [VRRP-IPv6] into the document that eventually became [RFC5798].

Thanks to Stewart Bryant, Sasha Vainshtein, Pascal Thubert, Alexander Okonnikov, Ben Niven-Jenkins, Tim Chown, Malia Vuini, Russ White, Donald Eastlake, Dave Thaler, Eric Kline, and Vijay Gurbani for comments on the current document (RFC 9568). Thanks to Gyan Mishra, Paul Congdon, and Jon Rosen for discussions related to the removal of legacy technology appendices. Thanks to Dhruv Dhody and Donald Eastlake for comments and suggestions for improving the IANA section. Thanks to Sasha Vainshtein for recommending "Maximum Advertisement Interval" validation. Thanks to Tim Chown and Fernando Gont for discussions and updates related to IPv6 SLAAC.

Special thanks to Quentin Armitage for a detailed review and extensive comments on the current document (RFC 9568).

Authors' Addresses

Acee Lindem
LabN Consulting, L.L.C.
301 Midenhall Way
Cary, NC 27513
United States of America
Email: acee.ietf@gmail.com

Aditya Dogra
Cisco Systems
Sarjapur Outer Ring Road
Bangalore 560103
Karnataka
India
Email: addogra@cisco.com