

Virtual Router Redundancy Protocol for IPv6

<[draft-ietf-vrrp-ipv6-spec-06.txt](#)>

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of \[RFC2026\]](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

To view the list Internet-Draft Shadow Directories, see <http://www.ietf.org/shadow.html>.

This internet draft expires on August 18, 2004.

Abstract

This memo defines the Virtual Router Redundancy Protocol (VRRP) for IPv6. It is version three (3) of the protocol. It is based on the original version of VRRP (version 2) for IPv4 that is defined in [RFC2338](#).

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 IP address(es) associated with a virtual router is called the Master, and forwards packets sent to these IP addresses. The election process provides dynamic fail over in the forwarding responsibility should the Master become unavailable. The advantage gained from using VRRP for IPv6 is a quicker switch over to back up routers than can be obtained with standard IPv6 Neighbor Discovery [[ND](#)] mechanisms.

Table of Contents

1.	Introduction.....	3
2.	Required Features.....	5
3.	VRRP Overview.....	6
4.	Sample Configurations.....	8
5.	Protocol.....	10
5.1	VRRP Packet Format.....	10
5.2	IP Field Descriptions.....	11
5.3	VRRP Field Descriptions.....	11
6.	Protocol State Machine.....	13
6.1	Parameters per Virtual Router.....	13
6.2	Timers.....	13
6.3	State Transition Diagram.....	15
6.4	State Descriptions.....	15
7.	Sending and Receiving VRRP Packets.....	19
7.1	Receiving VRRP Packets.....	19
7.2	Transmitting Packets.....	19
7.3	Virtual MAC Address.....	20
7.4	IPv6 Interface Identifiers.....	20
8.	Operational Issues.....	21
8.1	ICMPv6 Redirects.....	21
8.2	ND Neighbor Solicitation.....	21
8.3	Potential Forwarding Loop.....	22
8.4	Recommendations regarding setting priority values.....	22
9.	Operation over FDDI, Token Ring, and ATM LANE.....	22
9.1	Operation over FDDI.....	22
9.2	Operation over Token Ring.....	22
9.3	Operation over ATM LANE.....	25
10.	Security Considerations.....	25
11.	Intellectual Property.....	26
12.	Acknowledgments.....	26
13.	IANA Considerations.....	27
14.	Normative References.....	27
15.	Informative References.....	28
16.	Authors' Address.....	28
17.	Changes from RFC2338.....	29

1. Introduction

IPv6 hosts on a LAN will usually learn about one or more default routers by receiving Router Advertisements sent using the IPv6 Neighbor Discovery protocol [[ND](#)]. The Router Advertisements are multicast periodically at a rate that the hosts will learn about the default routers in a few minutes. They are not sent frequently enough to rely on the absence of the router advertisement to detect router failures.

Neighbor Discovery (ND) 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 will take a host about 38 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 timeout.

While the ND unreachability detection could be speeded up by changing the parameters 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 a one of more routers.

The Virtual Router Redundancy Protocol for IPv6 provides a much faster switch over 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 with out any interaction with the hosts and a minimum amount of VRRP traffic.

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 IP address(es) associated with a virtual router is called the Master, and forwards packets sent to these IP addresses. The election process provides dynamic fail over in the forwarding responsibility should the Master become unavailable.

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

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC 2119](#)].

[1.1](#) Scope

The remainder of this document describes the features, design goals, and theory of operation of VRRP for IPv6. The message formats, protocol processing rules and state machine that guarantee convergence to a single Virtual Router Master are presented. Finally, operational issues related to MAC address mapping, handling of Neighbor Discovery requests, generation of ICMPv6 redirect messages, and security issues are addressed.

This protocol is intended for use with IPv6 routers only. VRRP for IPv4 is defined in [[VRRP-V4](#)].

[1.2](#) Definitions

VRRP Router	A router running the Virtual Router Redundancy Protocol. It may participate in 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 an a set of associated IPv6 address(es) across a common LAN. A VRRP Router may backup one or more virtual routers.
IPv6 Address Owner	The VRRP router that has the virtual router's IPv6 address(es) as real interface address. This is the router that, when up, will respond to packets addressed to the IPv6 address(es) for ICMPv6 pings, TCP connections, etc.
Virtual Router Master	The VRRP router that is assuming the responsibility of forwarding packets sent to the IPv6 address(es) associated with the virtual router, and answering ND requests for these IPv6 address(es). Note that if the IPv6 address owner is available, then it will always become the Master.
Virtual Router Backup	The set of VRRP routers available to assume forwarding responsibility for a virtual router

should the current Master fail.

2.0 Required Features

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

2.1 IPv6 Address Backup

Backup of an IPv6 address(es) is the primary function of the Virtual Router Redundancy Protocol. While providing election of a Virtual Router Master and the additional functionality described below, the protocol should strive to:

- Minimize the duration of black holes.
- Minimize the steady state bandwidth overhead and processing complexity.
- Function over a wide variety of multiaccess LAN technologies capable of supporting IPv6 traffic.
- Provide for election of multiple virtual routers on a network for load balancing
- Support of multiple logical IPv6 subnets on a single LAN segment.

2.2 Preferred Path Indication

A simple model of Master election among a set of redundant routers is to treat each router with equal preference and claim victory after converging to any router as Master. 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 Master convergence to the most preferential router currently available.

2.3 Minimization of Unnecessary Service Disruptions

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

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

2.4 Efficient Operation over Extended LANs

Sending IPv6 packets on a multiaccess LAN requires mapping from an IPv6 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 station 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: 1) use the virtual router MAC as the source in a packet sent by the Master to trigger station learning; 2) trigger a message immediately after transitioning to Master to update the station learning; and 3) trigger periodic messages from the Master to maintain the station learning cache.

3.0 VRRP Overview

VRRP specifies an election protocol to provide the virtual router function described earlier. All protocol messaging is performed using IPv6 multicast datagrams, thus the protocol can operate over a variety of multiaccess LAN technologies supporting IPv6 multicast. Each VRRP virtual router has a single well-known MAC address allocated to it. This document currently only details the mapping to networks using the IEEE 802 48-bit MAC address. The virtual router MAC address is used as the source in all periodic VRRP messages sent by the Master router to enable bridge learning in an extended LAN.

A virtual router is defined by its virtual router identifier (VRID) and a set of IPv6 address(es). A VRRP router may associate a virtual router with its real address on an interface, and may also be configured with additional virtual router mappings and priority for virtual routers it is willing to backup. The mapping between VRID and its IPv6 address(es) must be coordinated among all VRRP routers on a LAN. However, there is no restriction against reusing a VRID with a different address mapping on different LANs. The scope of each virtual router is restricted to a single LAN.

To minimize network traffic, only the Master for each virtual router sends periodic VRRP Advertisement messages. A Backup router will not attempt to preempt the Master unless it has higher priority. This

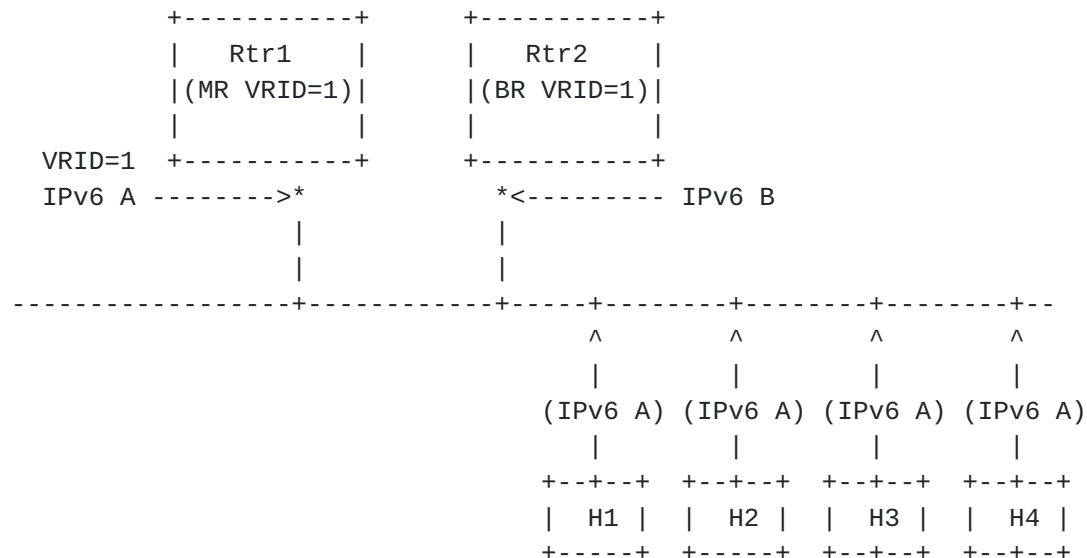
eliminates service disruption unless a more preferred path becomes available. It's also possible to administratively prohibit all preemption attempts. The only exception is that a VRRP router will always become Master of any virtual router associated with address it owns. If the Master becomes unavailable then the highest priority Backup will transition to Master after a short delay, providing a controlled transition of the virtual router responsibility with minimal service interruption.

The VRRP protocol design provides rapid transition from Backup to Master to minimize service interruption, and incorporates optimizations that reduce protocol complexity while guaranteeing controlled Master transition for typical operational scenarios. The 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 among each router. A side effect when these assumptions are violated (i.e., more than two redundant paths all with equal preference) is that duplicate packets may be forwarded for a brief period during Master election. However, the typical scenario assumptions are likely to cover the vast majority of deployments, loss of the Master router is infrequent, and the expected duration in Master election convergence is quite small (\ll 1 second). Thus the VRRP optimizations represent significant simplifications in the protocol design while incurring an insignificant probability of brief network degradation.

4. Sample Configurations

4.1 Sample Configuration 1

The following figure shows a simple network with two VRRP routers implementing one virtual router. Note that this example is provided to help understand the protocol, but is not expected to occur in actual practice.



Legend:

```

---+---+---+-- = Ethernet, Token Ring, or FDDI
      H = Host computer
      MR = Master Router
      BR = Backup Router
      * = IPv6 Address
      (IPv6) = default router for hosts

```

Eliminating all mention of VRRP (VRID=1) from the figure above leaves it as a typical IPv6 deployment. Each router has a link-local IPv6 address on the LAN interface (Rtr1 is assigned IPv6 Link-Local A and Rtr2 is assigned IPv6 Link-Local B), and each host learns a default route from Router Advertisements through one of the routers (in this example they all use Rtr1's IPv6 Link-Local A).

Moving to the VRRP environment, each router has the exact same Link-Local IPv6 address. Rtr1 is said to be the IPv6 address owner of IPv6 A, and Rtr2 is the IPv6 address owner of IPv6 B. A virtual router is then defined by associating a unique identifier (the virtual router ID) with the address owned by a router. Finally, the VRRP protocol manages virtual router fail over to a backup router.


The example above shows a virtual router configured to cover the IPv6

address owned by Rtr1 (VRID=1,IPv6_Address=A). When VRRP is enabled on Rtr1 for VRID=1 it will assert itself as Master, with priority=255, since it is the IPv6 address owner for the virtual router IPv6 address. When VRRP is enabled on Rtr2 for VRID=1 it will transition to Backup, with priority=100, since it is not the IPv6 address owner. If Rtr1 should fail then the VRRP protocol will transition Rtr2 to Master, temporarily taking over forwarding responsibility for IPv6 A to provide uninterrupted service to the hosts.

Note that in this example IPv6 B is not backed up, it is only used by Rtr2 as its interface address. In order to backup IPv6 B, a second virtual router must be configured. This is shown in the next section.

4.2 Sample Configuration 2

The following figure shows a configuration with two virtual routers with the hosts splitting their traffic between them. This example is expected to be common in actual practice.



Legend:

```

---+---+---+--- = Ethernet, Token Ring, or FDDI
H = Host computer
MR = Master Router
BR = Backup Router
* = IPv6 Address
(IPv6) = default router for hosts

```

In the example above, half of the hosts have learned a default route through Rtr1's IPv6 A and half are using Rtr2'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 Rtr2 (VRID=2, IPv6_Address=B). In this case Rtr2 will assert itself as Master for VRID=2 while Rtr1 will act as a backup. This scenario demonstrates a deployment providing load splitting when both routers are available while providing full redundancy for robustness.

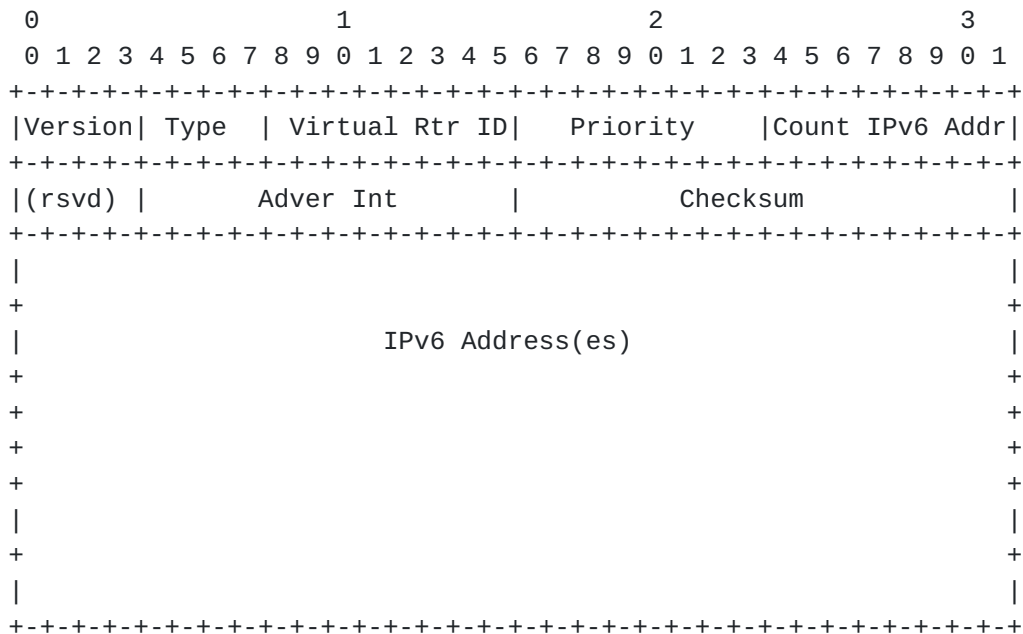
5.0 Protocol

The purpose of the VRRP packet is to communicate to all VRRP routers the priority and the state of the Master router associated with the Virtual Router ID.

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 IPv6 header.



[5.2](#) IPv6 Field Descriptions

[5.2.1](#) Source Address

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

[5.2.2](#) Destination Address

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

FF02:0:0:0:0:0:XXXX:XXXX

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

[5.2.3](#) 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.

[5.2.4](#) Next Header

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

[5.3](#) VRRP Field Descriptions

[5.3.1](#) Version

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

[5.3.2](#) 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.

[5.3.3](#) Virtual Rtr ID (VRID)

The Virtual Router Identifier (VRID) field identifies the virtual router this packet is reporting status for.

5.3.4 Priority

The priority field specifies the sending VRRP router's priority for the virtual router. Higher values equal higher priority. This field is an 8 bit unsigned integer field.

The priority value for the VRRP router that owns the IPv6 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).

The priority value zero (0) has special meaning indicating that the current Master has stopped participating in VRRP. This is used to trigger Backup routers to quickly transition to Master without having to wait for the current Master to timeout.

5.3.5 Count IPv6 Addr

The number of IPv6 addresses contained in this VRRP advertisement. The minimum value is 1.

5.3.5 Rsvd

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

5.3.6 Advertisement Interval (Adver Int)

The Advertisement interval is a 12-bit field that indicates the time interval (in centiseconds) between ADVERTISEMENTS. The default is 100 centiseconds (1 second). This field is used for troubleshooting misconfigured routers.

5.3.7 Checksum

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

The checksum is the 16-bit one's complement of the one's complement sum of the entire VRRP message starting with the version field and a "pseudo-header" as defined in [section 8.1 of RFC2460 \[IPv6\]](#). The next header field in the "pseudo-header" should be set to 112 (decimal) for VRRP. For computing the checksum, the checksum field is set to zero. See [RFC1071](#) for more detail [[CKSM](#)].

5.3.8 IPv6 Address(es)

One or more IPv6 addresses associated associated with the virtual router. The number of addresses included is specified in the "Count IP Addr" field. The first address must be the IPv6 link-local address associated with the virtual router. These fields are used for troubleshooting misconfigured routers.

6. Protocol State Machine

6.1 Parameters per Virtual Router


VRID	Virtual Router Identifier. Configurable item in the range 1-255 (decimal). There is no default.
Priority	Priority value to be used by this VRRP router in Master election for this virtual router. The value of 255 (decimal) is reserved for the router that owns the IPv6 address associated with the virtual router. The value of 0 (zero) is reserved for Master router to indicate it is releasing responsibility for the virtual router. The range 1-254 (decimal) is available for VRRP routers backing up the virtual router. The default value is 100 (decimal).
IPv6_Addresses	One or more IPv6 addresses associated with this virtual router. Configured item. No default. The first address must be the Link-Local address associated with the virtual router.
Advertisement_Interval	Time interval between ADVERTISEMENTS (centiseconds). Default is 100 centiseconds (1 second).
Skew_Time	Time to skew Master_Down_Interval in centiseconds. Calculated as: $(((256 - \text{Priority}) / 256) * \text{Advertisement_Interval})$

Master_Down_Interval	Time interval for Backup to declare Master down (centiseconds). Calculated as: $(3 * \text{Advertisement_Interval}) + \text{Skew_time}$
Preempt_Mode	Controls whether a higher priority Backup router preempts a lower priority Master. Values are True to allow preemption and False to prohibit preemption. Default is True. Note: Exception is that the router that owns the IPv6 address associated with the virtual router always preempts independent of the setting of this flag.
Accept_Mode	Controls whether a virtual router in Master state will accept packets addressed to the address owner's IPv6 address as its own if it is not the IPv6 address owner. Default is False.

[6.2](#) Timers

Master_Down_Timer	Timer that fires when ADVERTISEMENT has not been heard for Master_Down_Interval.
Adver_Timer	Timer that fires to trigger sending of ADVERTISEMENT based on Advertisement_Interval.

6.3 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 upper case characters.

A VRRP router implements an instance of the state machine for each virtual router election it is participating in.

6.4.1 Initialize

The purpose of this state is to wait for a Startup event. If a Startup event is received, then:

- If the Priority = 255 (i.e., the router owns the IPv6 address associated with the virtual router)
 - o Send an ADVERTISEMENT
 - o Send an unsolicited ND Neighbor Advertisement with the Router Flag (R) set, the Solicited Flag (S) unset, the Override flag (O) set, the Target Address set to the IPv6 link-local address of the Virtual Router, and the Target Link Layer address set to the virtual router MAC address.
 - o Set the Adver_Timer to Advertisement_Interval
 - o Transition to the {Master} state

else

- o Set the Master_Down_Timer to Master_Down_Interval
- o Transition to the {Backup} state

endif

6.4.2 Backup

The purpose of the {Backup} state is to monitor the availability and state of the Master Router.

While in this state, a VRRP router MUST do the following:

- MUST NOT respond to ND Neighbor Solicitation messages for the IPv6 address(es) associated with the virtual router.
- MUST NOT send ND Router Advertisement messages for the virtual router.
- MUST discard packets with a destination link layer MAC address equal to the virtual router MAC address.
- MUST NOT accept packets addressed to the IPv6 address(es) associated with the virtual router.

- If a Shutdown event is received, then:

- o Cancel the Master_Down_Timer
- o Transition to the {Initialize} state

endif

- If the Master_Down_Timer fires, then:

- o Send an ADVERTISEMENT
- o Compute and join the Solicited-Node multicast address [[ADD-ARH](#)] for the IPv6 address(es) addresses associated with the the Virtual Router.
- o Send an unsolicited ND Neighbor Advertisement with the Router Flag (R) set, the Solicited Flag (S) unset, the Override flag (O) set, the Target Address set to the IPv6 link-local address of the Virtual Router, and the Target Link Layer address set to the virtual router MAC address.
- o Set the Adver_Timer to Advertisement_Interval
- o Transition to the {Master} state

endif

- If an ADVERTISEMENT is received, then:

If the Priority in the ADVERTISEMENT is Zero, then:

- o Set the Master_Down_Timer to Skew_Time


```
else:

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

        o Reset the Master_Down_Timer to Master_Down_Interval

    else:

        o Discard the ADVERTISEMENT

    endif
endif
endif
```

6.4.3 Master

While in the {Master} state the router functions as the forwarding router for the IPv6 address associated with the virtual router.

While in this state, a VRRP router MUST do the following:

- MUST be a member of the Solicited-Node multicast address for the IPv6 link-local address associated with the virtual router.
- MUST respond to ND Neighbor Solicitation message for the IPv6 address(es) associated with the virtual router.
- MUST send ND Router Advertisements for the virtual router.
- MUST respond to ND Router Solicitation message for the virtual router.
- MUST forward packets with a destination link layer MAC address equal to the virtual router MAC address.
- MUST accept packets addressed to the IPv6 address(es) associated with the virtual router if it is the IPv6 address owner or if Accept_Mode is True. Otherwise, MUST NOT accept these packets.
- If a Shutdown event is received, then:
 - o Cancel the Adver_Timer
 - o Send an ADVERTISEMENT with Priority = 0
 - o Transition to the {Initialize} state

```
endif
```



```
- If the Adver_Timer fires, then:

    o Send an ADVERTISEMENT
    o Reset the Adver_Timer to Advertisement_Interval

endif

- If an ADVERTISEMENT is received, then:

    If the Priority in the ADVERTISEMENT is Zero, then:

        o Send an ADVERTISEMENT
        o Reset the Adver_Timer to Advertisement_Interval

    else:

        If the Priority in the ADVERTISEMENT is greater than the
        local Priority,
        or
        If the Priority in the ADVERTISEMENT is equal to the local
        Priority and the IPv6 Address of the sender is greater than
        the local IPv6 Address, then:

            o Cancel Adver_Timer
            o Set Master_Down_Timer to Master_Down_Interval
            o Transition to the {Backup} state

        else:

            o Discard ADVERTISEMENT

        endif
    endif
endif
```


7. Sending and Receiving VRRP Packets

7.1 Receiving VRRP Packets

Performed the following functions when a VRRP packet is received:

- MUST verify that the IPv6 Hop Limit is 255.
- MUST verify the VRRP version is 3
- MUST verify that the received packet contains the complete VRRP packet (including fixed fields, and IPv6 Address.
- MUST verify the VRRP checksum
- MUST verify that the VRID is configured on the receiving interface and the local router is not the IPv6 Address owner (Priority equals 255 (decimal)).

If any one of the above checks fails, the receiver MUST discard the packet, SHOULD log the event and SHOULD indicate via network management that an error occurred.

- MAY verify that the IPv6 Address matches the IPv6_Address configured for the VRID.

If the above check fails, the receiver SHOULD log the event and SHOULD indicate via network management that a misconfiguration was detected. If the packet was not generated by the address owner (Priority does not equal 255 (decimal)), the receiver MUST drop the packet, otherwise continue processing.

- MUST verify that the Adver Interval in the packet is the same as the locally configured for this virtual router

If the above check fails, the receiver SHOULD log the event and SHOULD 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 Virtual Router MAC Address
- Set the source IPv6 address to interface link-local IPv6 address
- Set the IPv6 protocol to VRRP
- Send the VRRP packet to the VRRP IP 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 the virtual router is attached to.

7.3 Virtual Router MAC Address

The virtual router MAC address associated with a virtual router is an IEEE 802 MAC Address in the following format:

00-00-5E-00-02-{VRID} (in hex in internet standard bit-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 for IPv6 protocol. {VRID} is the VRRP Virtual Router Identifier. This mapping provides for up to 255 VRRP routers on a network.

7.4 IPv6 Interface Identifiers

IPv6 Routers running VRRP MUST create their Interface Identifiers in the normal manner (e.g., [RFC2464](#) "Transmission of IPv6 Packets over Ethernet"). They MUST NOT use the Virtual Router MAC address to create the Modified EUI-64 identifiers.

This VRRP specification describes how to advertise and resolve the VRRP routers IPv6 link local address into the Virtual Router MAC address.

8. Operational Issues

8.1 ICMPv6 Redirects

ICMPv6 Redirects may be used normally when VRRP is running between a group of routers [[ICMPv6](#)]. 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 the end host used when making its next hop routing decision. If a VRRP router is acting as Master for virtual router(s) containing addresses it does not own, then it must determine which virtual router the packet was sent to 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.

8.2 ND Neighbor Solicitation

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

When a Master virtual router sends an ND Neighbor Solicitation message for a host's IPv6 address, the Master virtual router **MUST** include the virtual 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, it should only send ND messages that include Virtual MAC addresses. This may entail:

- When configuring an interface, VRRP 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; delay all ND Router and Neighbor Advertisements and Solicitation messages until both the IPv6 address and the virtual router MAC address are configured.

8.3 Potential Forwarding Loop

A VRRP router SHOULD not forward packets addressed to the IPv6 Address it becomes Master for if it is not the owner. Forwarding these packets would result in unnecessary traffic. Also in the case of LANs that receive packets they transmit (e.g., token ring) this can result in a forwarding loop that is only terminated when the IPv6 TTL expires.

One such mechanism for VRRP routers is to add/delete a reject host route for each adopted IPv6 address when transitioning to/from MASTER state.

8.4 Recommendations regarding setting priority values

A priority value of 255 designates a particular router as the "IPv6 address owner". Care must be taken not to configure more than one router on the link in this way for a single VRID.

Routers with priority 255 will, as soon as they start up, preempt all lower priority routers. Configure no more than one router on the link with priority 255, especially if preemption is set. If no router has this priority, and preemption is disabled, then no preemption will occur.

When there are multiple Backup routers, their priority values should be uniformly distributed. For example, if one Backup routers has the default priority of 100 and another BR is added, a priority of 50 would be a better choice for it than 99 or 100 to facilitate faster convergence.

9. Operation over FDDI, Token Ring, and ATM LANE

9.1 Operation over FDDI

FDDI interfaces remove from the FDDI ring frames that have a source MAC address matching the device's hardware address. Under some conditions, such as router isolations, ring failures, protocol transitions, etc., VRRP may cause there to be more than one Master router. If a Master router installs the virtual router MAC address as the hardware address on a FDDI device, then other Masters' ADVERTISEMENTS will be removed from the ring during the Master convergence, and convergence will fail.

To avoid this an implementation SHOULD configure the virtual router

MAC address by adding a unicast MAC filter in the FDDI device, rather than changing its hardware MAC address. This will prevent a Master router from removing any ADVERTISEMENTS it did not originate.

9.2 Operation over Token Ring

Token ring has several characteristics that make running VRRP difficult. These include:

- In order to switch to a new master located on a different bridge token ring segment from the previous master when using source route bridges, a mechanism is required to update cached source route information.
- No general multicast mechanism supported across old and new token ring adapter implementations. While many newer token ring adapters support group addresses, token ring functional address support is the only generally available multicast mechanism. Due to the limited number of token ring functional addresses these may collide with other usage of the same token ring functional addresses.

Due to these difficulties, the preferred mode of operation over token ring will be to use a token ring functional address for the VRID virtual MAC address. Token ring functional addresses have the two high order bits in the first MAC address octet set to B'1'. They range from 03-00-00-00-00-80 to 03-00-02-00-00-00 (canonical format). However, unlike multicast addresses, there is only one unique functional address per bit position. The functional addresses 03-00-00-10-00-00 through 03-00-02-00-00-00 are reserved by the Token Ring Architecture [[TKARCH](#)] for user-defined applications. However, since there are only 12 user-defined token ring functional addresses, there may be other non-IP protocols using the same functional address. Since the Novell IPX [[IPX](#)] protocol uses the 03-00-00-10-00-00 functional address, operation of VRRP over token ring will avoid use of this functional address. In general, token ring VRRP users will be responsible for resolution of other user-defined token ring functional address conflicts.

VRIDs are mapped directly to token ring functional addresses. In order to decrease the likelihood of functional address conflicts, allocation will begin with the largest functional address. Most non-IP protocols use the first or first couple user-defined functional addresses and it is expected that VRRP users will choose VRIDs sequentially starting with 1.

VRID	Token Ring Functional Address
----	-----
1	03-00-02-00-00-00
2	03-00-04-00-00-00
3	03-00-08-00-00-00
4	03-00-10-00-00-00
5	03-00-20-00-00-00
6	03-00-40-00-00-00
7	03-00-80-00-00-00
8	03-00-00-01-00-00
9	03-00-00-02-00-00
10	03-00-00-04-00-00
11	03-00-00-08-00-00

Or more succinctly, octets 3 and 4 of the functional address are equal to $(0x4000 \gg (VRID - 1))$ in non-canonical format.

Since a functional address cannot be used as a MAC level source address, the real MAC address is used as the MAC source address in VRRP advertisements. This is not a problem for bridges since packets addressed to functional addresses will be sent on the spanning-tree explorer path [[802.1D](#)].

The functional address mode of operation MUST be implemented by routers supporting VRRP on token ring.

Additionally, routers MAY support unicast mode of operation to take advantage of newer token ring adapter implementations that support non-promiscuous reception for multiple unicast MAC addresses and to avoid both the multicast traffic and usage conflicts associated with the use of token ring functional addresses. Unicast mode uses the same mapping of VRIDs to virtual MAC addresses as Ethernet. However, one important difference exists. ND request/reply packets contain the virtual MAC address as the source MAC address. The reason for this is that some token ring driver implementations keep a cache of MAC address/source routing information independent of the ND cache. Hence, these implementations need have to receive a packet with the virtual MAC address as the source address in order to transmit to that MAC address in a source-route bridged network.

Unicast mode on token ring has one limitation that should be considered. If there are VRID routers on different source-route bridge segments and there are host implementations that keep their source-route information in the ND cache and do not listen to gratuitous NDs, these hosts will not update their ND source-route information correctly when a switch-over occurs. The only possible solution is to put all routers with the same VRID on the same source-bridge segment and use techniques to prevent that bridge segment from

being a single point of failure. These techniques are beyond the scope of this document.

For both the multicast and unicast mode of operation, VRRP advertisements sent to 224.0.0.18 should be encapsulated as described in [\[RFC1469\]](#).

9.3 Operation over ATM LANE

Operation of VRRP over ATM LANE on routers with ATM LANE interfaces and/or routers behind proxy LEC's are beyond the scope of this document.

10. Security Considerations

VRRP for IPv6 does not currently include any type of authentication. Earlier versions of the VRRP (for IPv4) specification included several types of authentication ranging from none to strong. Operational experience and further analysis determined that these did not provide any real measure of security. Due to the nature of the VRRP protocol, even if VRRP messages are cryptographically protected, it does not prevent hostile routers from behaving as if they are a VRRP master, creating multiple masters. Authentication of VRRP messages could have prevented a hostile router from causing all properly functioning routers from going into backup state. However, having multiple masters can cause as much disruption as no routers, which authentication cannot prevent. Also, even if a hostile router could not disrupt VRRP, it can disrupt ARP and create the same effect as having all routers go into backup.

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 do include promiscuously receiving packets for any routers MAC address, sending packets with the routers MAC address as the source MAC addresses in the L2 header to tell the L2 switches to send packets addressed to the router to the malicious node instead of the router, send redirects to tell the hosts to send their traffic somewhere else, send unsolicited ND replies, answer ND requests, etc., etc. All of this can be done independently of implementing VRRP. VRRP does not add to these vulnerabilities.

Independent of any authentication type VRRP includes a mechanism (setting TTL=255, checking on receipt) that protects against VRRP packets being injected from another remote network. This limits most vulnerabilities to local attacks.

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.

11. Intellectual Property

The IETF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in [BCP-11](#). Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementors or users of this specification can be obtained from the IETF Secretariat. See the IETF IPR web page at <http://www.ietf.org/ipr.html> for additional information.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

The IETF has been notified of intellectual property rights claimed in regard to some or all of the specification contained in this document. For more information consult the online list of claimed rights.

12. Acknowledgments

This specification is based on [RFC2238](#). The authors of [RFC2238](#) are S. Knight, D. Weaver, D. Whipple, R. Hinden, D. Mitzel, P. Hunt, P. Higginson, M. Shand, and A. Lindem.

The author of this document would also like to thank Erik Nordmark, Thomas Narten, Steve Deering, Radia Perlman, Danny Mitzel, Mukesh Gupta, Don Provan, and Mark Hollinger for their helpful suggestions.

13. IANA Considerations

VRRP for IPv6 needs an IPv6 link-local scope multicast address assigned by the IANA for this specification. The IPv6 multicast address should be of the following form:

FF02:0:0:0:0:0:XXXX:XXXX

The values assigned address should be entered into [section 5.2.2](#).

A convenient assignment of this link-local scope multicast would be:

FF02:0:0:0:0:0:0:12

as this would be consistent with the IPv4 assignment for VRRP.

The IANA should also reserve a block of IANA Ethernet unicast addresses from:

00-00-5E-00-02-00 to 00-00-5E-00-02-FF in hex

for VRRP for IPv6. Similar assignments are documented in:

<http://www.iana.org/assignments/ethernet-numbers>

14. Normative References

- [802.1D] International Standard ISO/IEC 10038: 1993, ANSI/IEEE Std 802.1D, 1993 edition.
- [ADD-ARH] Hinden, R., S. Deering, "IP Version 6 Addressing Architecture", [RFC3513](#), April 2003.
- [CKSM] Braden, R., D. Borman, C. Partridge, "Computing the Internet Checksum", [RFC1071](#), September 1988.
- [ICMPv6] Conta, A., S. Deering, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6)", [RFC2463](#), December 1998.
- [IPv6] Deering, S., R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC2460](#), December 1998.
- [IPX] Novell Incorporated., "IPX Router Specification", Version 1.10, October 1992.
- [ND] Narten, T., E. Nordmark, W. Simpson, "Neighbor Discovery

for IP Version 6 (IPv6)", [RFC2461](#), December 1998.

[RFC1469] Pusateri, T., "IP Multicast over Token Ring Local Area Networks", [RFC1469](#), June 1993.

[RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", [RFC2026](#), [BCP00009](#), October 1996.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [RFC2119](#), [BCP0014](#), March 1997.

[TKARCH] IBM Token-Ring Network, Architecture Reference, Publication SC30-3374-02, Third Edition, (September, 1989).

[VRRP-V4] Knight, S., et. al., "Virtual Router Redundancy Protocol", [RFC2338](#), April 1998.

15. Informative References

[HSRP] Li, T., B. Cole, P. Morton, D. Li, "Cisco Hot Standby Router Protocol (HSRP)", [RFC2281](#), March 1998.

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

[OSPF] Moy, J., "OSPF version 2", [RFC2328](#), STD0054, April 1998.

[RIP] Malkin, G., "RIP Version 2", [RFC2453](#), STD0056, November 1998.

16. Author's Address

Robert Hinden
Nokia
313 Fairchild Drive
Mountain View, CA 94043
USA

Phone: +1 650 625-2004
EMail: bob.hinden@nokia.com

17. Changes from [RFC2338](#)

- Added new subsection (8.4) with recommendations about setting priority values and it's relationship to the preempt flag.
- Changed rules for receiving VRRP packets to not drop the packet if the Adver Interval is not consistent with the local configuration for the virtual router. Only log and notify network management.
- Reduced granularity of the Advertisement_Interval to centiseconds (i.e., 1/100 of a second). Changes include:
 - o Made Adver Int field in the header 12-bits to allow range from 1 to 4096 centiseconds.
 - o Change Skew_Timer calculation to skew over one Advertisement_Interval.
- Added switch (Accept_Mode) to control whether a virtual router in Master state will accept packets addresses to the address owner's IPv6 address as its own if it is not the IPv6 address owner.
- Changed VMAC assignments to a separate block of IANA Ethernet addresses and added this to the IANA considerations section.
- Removed different authentication methods, header fields, and updated the security considerations section to explain the reasons for doing this.
- General rewrite to change protocol to provide virtual router functionality from IPv4 to IPv6. Specific changes include:
 - o Increment VRRP version to 3.
 - o Change packet format to support an 128-bit IPv6 address.
 - o Rewrote text to specify IPv6 Neighbor Discovery mechanisms instead of ARP.
 - o Changed state machine actions to use Neighbor Discovery mechanisms. This includes sending unsolicited Neighbor Advertisements, Receiving Neighbor Solicitations, joining the appropriate solicited node multicast group, sending Router Advertisements, and receiving Router Solicitations.
- Revised the [section 4](#) examples text with a clearer description of mapping of IPv6 address owner, priorities, etc.
- Clarify the [section 7.1](#) text describing address list validation.
- Corrected text in Preempt_Mode definition.
- Changed authentication to be per Virtual Router instead of per Interface.
- Added new subsection (9.3) stating that VRRP over ATM LANE is beyond the scope of this document.
- Clarified text describing received packet length check.
- Clarified text describing received authentication check.
- Clarified text describing VRID verification check.
- Added new subsection (8.3) describing need to not forward packets for adopted IPv6 addresses.
- Added clarification to the security considerations section.
- Added reference for computing the internet checksum.
- Updated references and author information.

