

Internet Engineering Task Force (IETF)
Request for Comments: 5798
Obsoletes: 3768
Category: Standards Track
ISSN: 2070-1721

S. Nadas, Ed.
Ericsson
March 2010

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

Abstract

This memo defines the Virtual Router Redundancy Protocol (VRRP) for IPv4 and IPv6. It is version three (3) of the protocol, and it is based on VRRP (version 2) for IPv4 that is defined in RFC 3768 and in "Virtual Router Redundancy Protocol for IPv6". 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 Master, and it forwards packets sent to these IPv4 or IPv6 addresses. VRRP Master routers are configured with virtual IPv4 or IPv6 addresses, and VRRP Backup routers infer the address family of the virtual addresses being carried based on the transport protocol. Within a VRRP router, the virtual routers in each of the IPv4 and IPv6 address families are a domain unto themselves and do not overlap. The election process provides dynamic failover in the forwarding responsibility should the Master 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 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc5798>.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1. Introduction	4
1.1. A Note on Terminology	4
1.2. IPv4	5
1.3. IPv6	6
1.4. Requirements Language	6
1.5. Scope	7
1.6. Definitions	7
2. Required Features	8
2.1. IPvX Address Backup	8
2.2. Preferred Path Indication	8
2.3. Minimization of Unnecessary Service Disruptions	9
2.4. Efficient Operation over Extended LANs	9
2.5. Sub-Second Operation for IPv4 and IPv6	9
3. VRRP Overview	10
4. Sample Configurations	11
4.1. Sample Configuration 1	11
4.2. Sample Configuration 2	13

5. Protocol	14
5.1. VRRP Packet Format	15
5.1.1. IPv4 Field Descriptions	15
5.1.1.1. Source Address	15
5.1.1.2. Destination Address	15
5.1.1.3. TTL	16
5.1.1.4. Protocol	16
5.1.2. IPv6 Field Descriptions	16
5.1.2.1. Source Address	16
5.1.2.2. Destination Address	16
5.1.2.3. Hop Limit	16
5.1.2.4. Next Header	16
5.2. VRRP Field Descriptions	16
5.2.1. Version	16
5.2.2. Type	17
5.2.3. Virtual Rtr ID (VRID)	17
5.2.4. Priority	17
5.2.5. Count IPvX Addr	17
5.2.6. Rsvd	17
5.2.7. Maximum Advertisement Interval (Max Adver Int)	17
5.2.8. Checksum	18
5.2.9. IPvX Address(es)	18
6. Protocol State Machine	18
6.1. Parameters Per Virtual Router	18
6.2. Timers	20
6.3. State Transition Diagram	21
6.4. State Descriptions	21
6.4.1. Initialize	21
6.4.2. Backup	22
6.4.3. Master	24
7. Sending and Receiving VRRP Packets	26
7.1. Receiving VRRP Packets	26
7.2. Transmitting VRRP Packets	27
7.3. Virtual Router MAC Address	28
7.4. IPv6 Interface Identifiers	28
8. Operational Issues	29
8.1. IPv4	29
8.1.1. ICMP Redirects	29
8.1.2. Host ARP Requests	29
8.1.3. Proxy ARP	30
8.2. IPv6	30
8.2.1. ICMPv6 Redirects	30
8.2.2. ND Neighbor Solicitation	30
8.2.3. Router Advertisements	31
8.3. IPvX	31
8.3.1. Potential Forwarding Loop	31
8.3.2. Recommendations Regarding Setting Priority Values	32

8.4. VRRPv3 and VRRPv2 Interoperation	32
8.4.1. Assumptions	32
8.4.2. VRRPv3 Support of VRRPv2	32
8.4.3. VRRPv3 Support of VRRPv2 Considerations	33
8.4.3.1. Slow, High-Priority Masters	33
8.4.3.2. Overwhelming VRRPv2 Backups	33
9. Security Considerations	33
10. Contributors and Acknowledgments	34
11. IANA Considerations	35
12. References	35
12.1. Normative References	35
12.2. Informative References	36
Appendix A. Operation over FDDI, Token Ring, and ATM LANE	38
A.1. Operation over FDDI	38
A.2. Operation over Token Ring	38
A.3. Operation over ATM LANE	40

1. Introduction

This memo defines the Virtual Router Redundancy Protocol (VRRP) for IPv4 and IPv6. It is version three (3) of the protocol. It is based on VRRP (version 2) for IPv4 that is defined in [RFC3768] and in [VRRP-IPv6]. 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 Master, and it forwards packets sent to these IPv4 or IPv6 addresses. VRRP Master routers are configured with virtual IPv4 or IPv6 addresses, and VRRP Backup routers infer the address family of the virtual addresses being carried based on the transport protocol. Within a VRRP router, the virtual routers in each of the IPv4 and IPv6 address families are a domain unto themselves and do not overlap. The election process provides dynamic failover in the forwarding responsibility should the Master become unavailable.

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

1.1. A Note on Terminology

This document discusses both IPv4 and IPv6 operation, 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" usually 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.

1.2. IPv4

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

Running a dynamic routing protocol on every end-host may be infeasible for a number of reasons, including administrative overhead, processing overhead, security issues, or lack of a protocol implementation for some platforms. Neighbor or router discovery protocols may require active participation by all hosts on a network, leading to large timer values to reduce protocol overhead in the face of large numbers of hosts. This can result in a significant delay in the detection of a lost (i.e., dead) neighbor; such a delay may introduce unacceptably long "black hole" periods.

The use of a statically configured default route is quite popular; it minimizes configuration and processing overhead on the end-host and is supported by virtually every IPv4 implementation. This mode of operation is likely to persist as dynamic host configuration protocols [RFC2131] are deployed, which typically provide configuration for an end-host IPv4 address and default gateway. 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 any alternate path that may be available.

The Virtual Router Redundancy Protocol (VRRP) is designed to eliminate the single point of failure inherent in the static default-routed environment. 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 Master and forwards packets sent to these IPv4 addresses. The election process provides dynamic failover in the forwarding responsibility should the Master 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 router discovery protocols on every end-host.

1.3. 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 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 time out.

While the ND unreachability detection could be made quicker 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 one of 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.

1.4. Requirements Language

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 [RFC2119].

1.5. 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 Virtual Router Master are presented. Finally, operational issues related to MAC address mapping, handling of ARP requests, generation of ICMP redirect messages, and security issues are addressed.

1.6. 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 may back up one or more virtual routers.
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 connections, 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 mode, 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.
Virtual Router Master	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 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. Required Features

This section outlines 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. While providing election of a Virtual Router Master and the additional functionality described below, the protocol should strive to:

- o Minimize the duration of black holes.
- o Minimize the steady-state bandwidth overhead and processing complexity.
- o Function over a wide variety of multiaccess LAN technologies capable of supporting IPvX traffic.
- o Allow multiple virtual routers on a network for load balancing.
- o Support multiple logical IPvX 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, 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 that is preferred over the current Master becomes available. It may be useful to support an override of the immediate convergence to the preferred path.

2.4. Efficient Operation over Extended LANs

Sending IPvX packets (that is, 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 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 address as the source in a packet sent by the Master to trigger station learning;
- 2) trigger a message immediately after transitioning to the Master to update the station learning;
- and 3) trigger periodic messages from the Master to maintain the station learning cache.

2.5. Sub-Second Operation for IPv4 and IPv6

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

One possible problematic scenario when using small VRRP_Advertisement_Intervals may occur when a router is delivering more packets onto the LAN than can be accommodated, and so a queue builds up in the router. It is possible that packets being transmitted onto the VRRP-protected LAN could see larger queueing delay than the smallest VRRP Advertisement_Interval. In this case, the Master_Down_Interval will be small enough so that normal queueing delays might cause a VRRP Backup to conclude that the Master is down, and therefore promote itself to Master. Very shortly afterwards, the delayed VRRP packets from the Master cause a switch back to Backup status. Furthermore, this process can repeat many times per second,

causing significant disruption to traffic. To mitigate this problem, priority forwarding of VRRP packets should be considered. It should be possible for a VRRP Master to observe that this situation is occurring frequently and at least log the problem.

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 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 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 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 addresses 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.

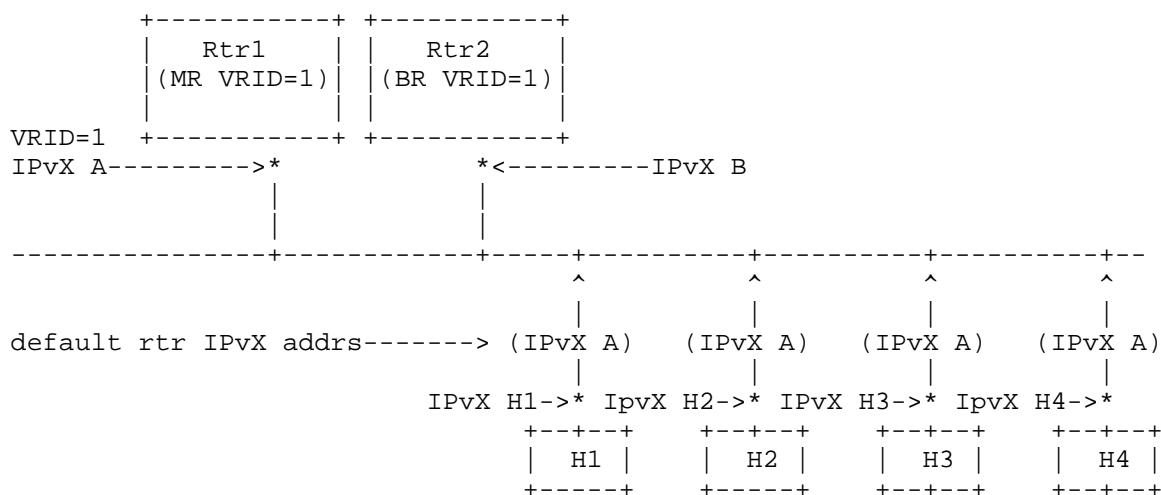
The VRRP protocol design provides rapid transition from Backup to Master to minimize service interruption and incorporates optimizations that reduce protocol complexity while guaranteeing

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.



Legend:

--+---+---+-- = Ethernet, Token Ring, or FDDI

H = Host computer

MR = Master Router

BR = Backup Router

* = IPvX Address; X is 4 everywhere in IPv4 case

X is 6 everywhere in IPv6 case

```
(IPvX) = default router for hosts
```

Eliminating all mention of VRRP (VRID=1) from the figure above leaves it as a typical deployment.

In the IPv4 case (that is, IPvX is IPv4 everywhere in the figure), each router is permanently assigned an IPv4 address on the LAN interface (Rtr1 is assigned IPv4 A and Rtr2 is assigned IPv4 B), and each host installs a static default route through one of the routers (in this example they all use Rtr1's IPv4 A).

In the IPv6 case (that is, IPvX is IPv6 everywhere in the figure), 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 an IPv4 VRRP environment, each router has the exact same permanently assigned IPv4 address. Rtr1 is said to be the IPv4 address owner of IPv4 A, and Rtr2 is the IPv4 address owner of IPv4 B. A virtual router is then defined by associating a unique identifier (the virtual router ID) with the address owned by a router.

Moving to an IPv6 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, in both the IPv4 and IPv6 cases, the VRRP protocol manages virtual router failover to a Backup router.

The IPv4 example above shows a virtual router configured to cover the IPv4 address owned by Rtr1 (VRID=1, IPv4_Address=A). When VRRP is enabled on Rtr1 for VRID=1, it will assert itself as Master, with priority = 255, since it is the IP address owner for the virtual router IP address. When VRRP is enabled on Rtr2 for VRID=1, it will transition to Backup, with priority = 100 (the default priority is 100), since it is not the IPv4 address owner. If Rtr1 should fail, then the VRRP protocol will transition Rtr2 to Master, temporarily taking over forwarding responsibility for IPv4 A to provide uninterrupted service to the hosts. When Rtr1 returns to service, it will re-assert itself as Master.

The IPv6 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

cover the IPv4 address owned by Rtr2 (VRID=2, IPv4_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.

In the IPv6 example above (that is, IPvX is IPv6 everywhere in the figure), 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.

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 load between the routers.

5. 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 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.

5.1.1.3. 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.

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 the packet is being sent from.

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.

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 this packet is reporting status for.

5.2.4. 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 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).

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 time out.

5.2.5. Count IPvX Addr

This is the number of either IPv4 addresses or IPv6 addresses contained in this VRRP advertisement. The minimum value is 1.

5.2.6. Rsvd

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

5.2.7. Maximum Advertisement Interval (Max Adver Int)

The Maximum Advertisement 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 Master routers with slower transmission rates than their Backup routers are unstable. This is because low-priority nodes configured to faster rates could come online and decide they should be Masters before they have heard anything from the higher-priority Master with a slower rate. When this happens, it is temporary: once the lower-priority node does hear from the higher-priority Master, it will relinquish mastership.

5.2.8. 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]. 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 [RFC1071].

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 "Count IP Addr" 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 make it easier to do this comparison.

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, that is, IPv4 and IPv6 MUST NOT both be carried in one IPvX Address field.

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 IPvX address associated with the virtual router. The value of 0 (zero) is reserved for the 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. Higher values indicate higher priorities. The default value is 100 (decimal).
IPv4_Addresses	One or more IPv4 addresses associated with this virtual router. Configured item with no default.
IPv6_Addresses	One or more IPv6 addresses associated with this virtual router. Configured item with 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).
Master_Adver_Interval	Advertisement interval contained in ADVERTISEMENTS received from the Master (centiseconds). This value is saved by virtual routers in the Backup state and used to compute Skew_Time and Master_Down_Interval. The initial value is the same as Advertisement_Interval.
Skew_Time	Time to skew Master_Down_Interval in centiseconds. Calculated as $(((256 - \text{priority}) * \text{Master_Adver_Interval}) / 256)$
Master_Down_Interval	Time interval for Backup to declare Master down (centiseconds). Calculated as $(3 * \text{Master_Adver_Interval}) + \text{Skew_time}$

Preempt_Mode Controls whether a (starting or restarting) higher-priority Backup router preempts a lower-priority Master 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 Master state will accept packets addressed to the address owner's IPvX address as its own 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 responses as the MAC address to use for IP_Addresses.

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-uppercase 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, that is, an implementation-defined mechanism that initiates the protocol once it has been configured. The configuration mechanism is out of scope of this specification.

(100) If a Startup event is received, then:

(105) - If the Priority = 255 (i.e., the router owns the IPvX address associated with the virtual router), then:

(110) + Send an ADVERTISEMENT

(115) + If the protected IPvX address is an IPv4 address, then:

(120) * Broadcast a gratuitous ARP request containing the virtual router MAC address for each IP address associated with the virtual router.

(125) + else // IPv6

(130) * 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) unset, 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.

```
(135) +endif // was protected addr IPv4?

(140) + Set the Adver_Timer to Advertisement_Interval

(145) + Transition to the {Master} state

(150) - else // rtr does not own virt addr

(155) + Set Master_Adver_Interval to Advertisement_Interval

(160) + Set the Master_Down_Timer to Master_Down_Interval

(165) + Transition to the {Backup} state

(170) -endif // priority was not 255

(175) endif // startup event was recv
```

6.4.2. Backup

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

```
(300) While in this state, a VRRP router MUST do the following:

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

(310) + MUST NOT respond to ARP requests for the IPv4
address(es) associated with the virtual router.

(315) - else // protected addr is IPv6

(320) + MUST NOT respond to ND Neighbor Solicitation messages
for the IPv6 address(es) associated with the virtual router.

(325) + MUST NOT send ND Router Advertisement messages for the
virtual router.

(330) -endif // was protected addr IPv4?

(335) - MUST discard packets with a destination link-layer MAC
address equal to the virtual router MAC address.
```

```
(340) - MUST NOT accept packets addressed to the IPvX address(es)
associated with the virtual router.

(345) - If a Shutdown event is received, then:

    (350) + Cancel the Master_Down_Timer

    (355) + Transition to the {Initialize} state

(360) -endif // shutdown recv

(365) - If the Master_Down_Timer fires, then:

    (370) + Send an ADVERTISEMENT

    (375) + If the protected IPvX address is an IPv4 address, then:

        (380) * Broadcast a gratuitous ARP request on that interface
        containing the virtual router MAC address for each IPv4
        address associated with the virtual router.

    (385) + else // ipv6

        (390) * Compute and join the Solicited-Node multicast
        address [RFC4291] for the IPv6 address(es) associated with
        the virtual router.

        (395) * 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) unset, 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.

    (400) +endif // was protected addr ipv4?

    (405) + Set the Adver_Timer to Advertisement_Interval

    (410) + Transition to the {Master} state

(415) -endif // Master_Down_Timer fired

(420) - If an ADVERTISEMENT is received, then:

    (425) + If the Priority in the ADVERTISEMENT is zero, then:

        (430) * Set the Master_Down_Timer to Skew_Time
```

```
(440) + else // priority non-zero

(445) * If Preempt_Mode is False, or if the Priority in the
ADVERTISEMENT is greater than or equal to the local
Priority, then:

(450) @ Set Master_Adver_Interval to Adver Interval
contained in the ADVERTISEMENT

(455) @ Recompute the Master_Down_Interval

(460) @ Reset the Master_Down_Timer to
Master_Down_Interval

(465) * else // preempt was true or priority was less

(470) @ Discard the ADVERTISEMENT

(475) *endif // preempt test

(480) +endif // was priority zero?

(485) -endif // was advertisement recv?

(490) endwhile // Backup state
```

6.4.3. Master

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

Note that in the Master state, the Preempt_Mode Flag is not considered.

```
(600) While in this state, a VRRP router MUST do the following:

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

(610) + MUST respond to ARP requests for the IPv4 address(es)
associated with the virtual router.

(615) - else // ipv6

(620) + MUST be a member of the Solicited-Node multicast
address for the IPv6 address(es) associated with the virtual
router.
```

```
(625) + MUST respond to ND Neighbor Solicitation message for
the IPv6 address(es) associated with the virtual router.

(630) ++ MUST send ND Router Advertisements for the virtual
router.

(635) ++ If Accept_Mode is False: MUST NOT drop IPv6 Neighbor
Solicitations and Neighbor Advertisements.

(640) +-endif // ipv4?

(645) - MUST forward packets with a destination link-layer MAC
address equal to the virtual router MAC address.

(650) - 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, MUST NOT accept these
packets.

(655) - If a Shutdown event is received, then:

    (660) + Cancel the Adver_Timer

    (665) + Send an ADVERTISEMENT with Priority = 0

    (670) + Transition to the {Initialize} state

(675) -endif // shutdown recv

(680) - If the Adver_Timer fires, then:

    (685) + Send an ADVERTISEMENT

    (690) + Reset the Adver_Timer to Advertisement_Interval

(695) -endif // advertisement timer fired

(700) - If an ADVERTISEMENT is received, then:

    (705) -- If the Priority in the ADVERTISEMENT is zero, then:

        (710) -* Send an ADVERTISEMENT

        (715) -* Reset the Adver_Timer to Advertisement_Interval

    (720) -- else // priority was non-zero
```

```
(725) -* If the Priority in the ADVERTISEMENT is greater
than the local Priority,

(730) -* or

(735) -* If 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, then:

    (740) -@ Cancel Adver_Timer

    (745) -@ Set Master_Adver_Interval to Adver Interval
    contained in the ADVERTISEMENT

    (750) -@ Recompute the Skew_Time

    (755) @ Recompute the Master_Down_Interval

    (760) @ Set Master_Down_Timer to Master_Down_Interval

    (765) @ Transition to the {Backup} state

(770) * else // new Master logic

    (775) @ Discard ADVERTISEMENT

(780) *endif // new Master detected

(785) +endif // was priority zero?

(790) -endif // advert recv

(795) endwhile // in Master
```

7. Sending and Receiving VRRP Packets

7.1. Receiving VRRP Packets

The following functions are performed when a VRRP packet is received:

- If the received packet is an IPv4 packet, then:
 - + MUST verify that the IPv4 TTL is 255.
- else // ipv6 recv
 - + MUST verify that the IPv6 Hop Limit is 255.

-endif

- MUST verify that the VRRP version is 3.
- MUST verify that the received packet contains the complete VRRP packet (including fixed fields, and IPvX address).
- MUST verify the VRRP checksum.
- 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, and MAY indicate via network management that an error occurred.

- MAY verify that "Count IPvX Addrs" and the list of IPvX address(es) match the IPvX Address(es) configured for the VRID.

If the above check fails, the receiver SHOULD log 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
 - If the protected address is an IPv4 address, then:
 - + Set the source MAC address to virtual router MAC Address
 - + Set the source IPv4 address to interface primary IPv4 address
 - else // ipv6
 - + Set the source MAC address to virtual router MAC Address
 - + Set the source IPv6 address to interface 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 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:

IPv4 case: 00-00-5E-00-01-{VRID} (in hex, in Internet-standard bit-order)

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

IPv6 case: 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 IPv6 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., "Transmission of IPv6 Packets over Ethernet Networks" [RFC2464]). They MUST NOT use the virtual router MAC address to create the Modified Extended Unique Identifier (EUI)-64 identifiers.

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 may be used normally when VRRP is running between 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 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.

It may be useful to disable redirects for specific cases where VRRP is being used to load-share traffic between 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 Virtual Router Master MUST respond to the ARP request with an ARP response that indicates the virtual 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 Virtual Router Master MUST NOT respond with its physical MAC address in the ARP response. This allows the client to always use the same MAC address regardless of the current Master router.

When a VRRP router restarts or boots, it SHOULD NOT send any ARP messages using its physical MAC address for the IPv4 address it owns; it should only send ARP messages that include virtual MAC addresses.

This may entail the following:

- o When configuring an interface, Virtual Router Master routers should broadcast a gratuitous ARP request containing the virtual router MAC address for each IPv4 address on that interface.
- o At system boot, when initializing interfaces for VRRP operation, delay gratuitous ARP requests and ARP responses until both the IPv4 address and the virtual router MAC address are configured.

- o When, for example, ssh access to a particular VRRP router is required, an IP address known to belong to that router must 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 may be used normally when VRRP is running between 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 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. 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 the virtual router IPv6 address, the Virtual Router Master MUST respond to the ND Neighbor Solicitation message with the virtual MAC address for the virtual router. The Virtual Router Master 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 Virtual Router Master sends an ND Neighbor Solicitation message for a host's IPv6 address, the Virtual Router Master 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 the following:

- o When configuring an interface, Virtual Router Master routers should send an unsolicited ND Neighbor Advertisement message containing the virtual router MAC address for the IPv6 address on that interface.
- o At system boot, when initializing interfaces for VRRP operation, all ND Router and Neighbor Advertisements and Solicitation messages must be delayed until both the IPv6 address and the virtual router MAC address are configured.

Note that on a restarting Master router where the VRRP protected address is the interface address, (that is, priority 255) duplicate address detection (DAD) may fail, as the Backup router may answer that it owns the address. One solution is to not run DAD in this case.

8.2.3. Router Advertisements

When a Backup VRRP router has become Master 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.

8.3. 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 Master. 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 IPvX TTL expires.

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

8.3.2. Recommendations Regarding Setting Priority Values

A priority value of 255 designates a particular router as the "IPvX 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. No more than one router on the link is to be configured 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 router has the default priority of 100 and another Backup Router is added, a priority of 50 would be a better choice for it than 99 or 100, in order to facilitate faster convergence.

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

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 Master, it MUST send both types at the configured rate, even if sub-second.

When a virtual router is configured this way and is the Backup, it should time out based on the rate advertised by the Master; in the case of a VRRPv2 Master, this means it must translate the timeout value it receives (in seconds) into centiseconds. Also, a Backup should ignore VRRPv2 advertisements from the current Master if it is also receiving VRRPv3 packets from it. It MAY report when a VRRPv3 Master is **not** sending VRRPv2 packets: that suggests they don't agree on whether they're supporting VRRPv2 routers.

8.4.3. VRRPv3 Support of VRRPv2 Considerations

8.4.3.1. Slow, High-Priority Masters

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

The VRRPv2 Master 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/VRRPv3 implementation it is interacting with if the VRRPv2/VRRPv3 rate is sub-second.

8.4.3.2. Overwhelming VRRPv2 Backups

It seems possible that a VRRPv3 Master router sending at centisecond rates could potentially overwhelm a VRRPv2 Backup router with potentially unclear results.

In this upgrade case, a deployment should initially run the VRRPv3 Master routers with lower frequencies (e.g., 100 centiseconds) until the VRRPv2 routers are upgraded. Then, once the deployment has convinced itself that VRRPv3 is working properly, the VRRPv2 support may be unconfigured and then the desired sub-second rates configured.

9. Security Considerations

VRRP for IPvX 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 sufficient security to overcome the vulnerability of misconfigured secrets, causing multiple Masters 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 a VRRP Master, creating multiple Masters. Authentication of VRRP messages could have prevented a hostile node 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 node could not disrupt VRRP, it can disrupt ARP and create the same effect as having all routers go into Backup.

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 do 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; send redirects to tell the hosts to send their traffic somewhere else; send unsolicited ND replies; answer ND requests; 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.

In the context of IPv6 operation, if SEcure Neighbor Discovery (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 Master and Backup routers the same prefix delegation in the certificates so that Master and Backup routers advertise the same set of subnet prefixes. However, the Master and Backup routers should have their own key pairs to avoid private key sharing.

10. Contributors and Acknowledgments

The editor would like to thank V. Ullanatt for his review of an early version. This document consists of very little new material (there is some new text in Appendix A) and was created by merging and "xml-izing" [VRRP-IPv6] and [RFC3768], and then adding in the changes discussed recently on the Virtual Router Redundancy Protocol working group's mailing list. R. Hinden is the author and J. Cruz the editor of the former. The contributors for the latter appear below.

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 author 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 Bellovin, Thomas Narten, Rob Montgomery, Rob Coltun, Radia Perlman, Russ Housley, Harald Alvestrand, Steve Bellovin, Ned Freed, Ted Hardie, Russ Housley, Bert Wijnen, Bill Fenner, and Alex Zinin for their comments and suggestions.

11. IANA Considerations

IANA has assigned an IPv6 link-local scope multicast address for VRRP for IPv6. The IPv6 multicast address is as follows:

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

The values assigned address should be entered into Section 5.1.2.2.

The IANA has reserved a block of IANA Ethernet unicast addresses for VRRP for IPv6 in the range

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

Similar assignments are documented at:

<http://www.iana.org>

12. References

12.1. Normative References

- [ISO.10038.1993] International Organization for Standardization, "Information technology - Telecommunications and information exchange between systems - Local area networks - Media access control (MAC) bridges", ISO Standard 10038, 1993.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- [RFC3768] Hinden, R., "Virtual Router Redundancy Protocol (VRRP)", RFC 3768, April 2004.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, March 2006.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.

12.2. Informative References

- [VRRP-IPv6] Hinden, R. and J. Cruz, "Virtual Router Redundancy Protocol for IPv6", Work in Progress, March 2007.
- [IPSTB] Higginson, P. and M. Shand, "Development of Router Clusters to Provide Fast Failover in IP Networks", Digital Technical Journal, Volume 9 Number 3, Winter 1997.
- [IPX] Novell Incorporated, "IPX Router Specification Version 1.10", October 1992.
- [RFC1071] Braden, R., Borman, D., Partridge, C., and W. Plummer, "Computing the Internet checksum", RFC 1071, September 1988.
- [RFC1256] Deering, S., Ed., "ICMP Router Discovery Messages", RFC 1256, September 1991.
- [RFC1469] Pusateri, T., "IP Multicast over Token-Ring Local Area Networks", RFC 1469, June 1993.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.
- [RFC2281] Li, T., Cole, B., Morton, P., and D. Li, "Cisco Hot Standby Router Protocol (HSRP)", RFC 2281, March 1998.

- [RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.
- [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, April 1998.
- [RFC2453] Malkin, G., "RIP Version 2", STD 56, RFC 2453, November 1998.
- [RFC2464] Crawford, M., "Transmission of IPv6 Packets over Ethernet Networks", RFC 2464, December 1998.
- [RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, March 2005.
- [TKARCH] IBM Incorporated, "IBM Token-Ring Network, Architecture Specification, Publication SC30-3374-02, Third Edition", September 1989.

Appendix A. Operation over FDDI, Token Ring, and ATM LANE

A.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.

A.2. Operation over Token Ring

Token Ring has several characteristics that make running VRRP difficult. These include the following:

- o 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.
- o No general multicast mechanism is 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-IPvX 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-IPvX 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 >> (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 [ISO.10038.1993].

The functional-address mode of operation MUST be implemented by routers supporting VRRP on Token Ring.

Additionally, routers MAY support the 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 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 switchover occurs. The only possible solution is to put all routers with the same VRID on the same source-route-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].

A.3. Operation over ATM LANE

Operation of VRRP over ATM LANE on routers with ATM LANE interfaces and/or routers behind proxy LAN Emulation Clients (LECs) are beyond the scope of this document.

Author's Address

Stephen Nadas (editor)
Ericsson
900 Chelmsford St., T3 4th Floor
Lowell, MA 01851
USA

Phone: +1 978 275 7448
EMail: stephen.nadas@ericsson.com

