

Virtual Router Redundancy Protocol (VRRP)

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2004). All Rights Reserved.

Abstract

This memo defines the Virtual Router Redundancy Protocol (VRRP). 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. This allows any of the virtual router IP addresses on the LAN to 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.

Table of Contents

1.	Introduction.	2
1.1.	Contributors.	3
1.2.	Scope	4
1.3.	Definitions	4
2.	Required Features	5
2.1.	IP Address Backup	5
2.2.	Preferred Path Indication	5
2.3.	Minimization of Unnecessary Service Disruptions	5
2.4.	Efficient Operation over Extended LANs.	6
3.	VRRP Overview	6
4.	Sample Configurations	7

4.1.	Sample Configuration 1.	7
4.2.	Sample Configuration 2.	9
5.	Protocol.	10
5.1.	VRRP Packet Format.	10
5.2.	IP Field Descriptions	10
5.3.	VRRP Field Descriptions	11
6.	Protocol State Machine.	13
6.1.	Parameters per Virtual Router	13
6.2.	Timers.	14
6.3.	State Transition Diagram.	15
6.4.	State Descriptions.	15
7.	Sending and Receiving VRRP Packets.	18
7.1.	Receiving VRRP Packets.	18
7.2.	Transmitting Packets.	19
7.3.	Virtual MAC Address	19
8.	Operational Issues.	20
8.1.	ICMP Redirects.	20
8.2.	Host ARP Requests	20
8.3.	Proxy ARP	20
8.4.	Potential Forwarding Loop	21
9.	Operation over FDDI, Token Ring, and ATM LANE	21
9.1.	Operation over FDDI	21
9.2.	Operation over Token Ring	21
9.3.	Operation over ATM LANE	23
10.	Security Considerations	23
11.	Acknowledgements.	24
12.	References.	24
12.1.	Normative References.	24
12.2.	Informative References.	25
13.	Changes from RFC2338.	25
14.	Editor's Address.	26
15.	Full Copyright Statement.	27

1. Introduction

There are a number of methods that an end-host can use to determine its first hop router towards a particular IP destination. These include running (or snooping) a dynamic routing protocol such as Routing Information Protocol [RIP] or OSPF version 2 [OSPF], running an ICMP router discovery client [DISC] 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, that 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 IP implementation. This mode of operation is likely to persist as dynamic host configuration protocols [DHCP] are deployed, which typically provide configuration for an end-host IP 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 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. Any of the virtual router's IP 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.

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 [RFC2119].

1.1. Contributors

The following people, who are the authors of the RFC 2338 that this document is based on and replaces, contributed to the text in this document. They are P. Higginson, R. Hinden, P. Hunt, S. Knight, A. Lindem, D. Mitzel, M. Shand, D. Weaver, and D. Whipple. They are not listed as authors of the document due to current RFC-Editor policies.

1.2. 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.

This protocol is intended for use with IPv4 routers only. A separate specification will be produced if it is decided that similar functionality is desirable in an IPv6 environment.

1.3. 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 a set of associated IP address(es) across a common LAN. A VRRP Router may backup one or more virtual routers.
IP Address Owner	The VRRP router that has the virtual router's IP address(es) as real interface address(es). This is the router that, when up, will respond to packets addressed to one of these IP addresses for ICMP pings, TCP connections, etc.
Primary IP Address	An IP address selected from the set of real interface addresses. One possible selection algorithm is to always select the first address. VRRP advertisements are always sent using the primary IP address as the source of the IP packet.
Virtual Router Master	The VRRP router that is assuming the responsibility of forwarding packets sent to the IP address(es) associated with the virtual router, and answering ARP requests for these IP addresses. Note that if the IP 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. IP Address Backup

Backup of IP addresses 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 IP traffic.
- Provide for election of multiple virtual routers on a network for load balancing.
- Support of multiple logical IP 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 IP packets on a multiaccess LAN requires mapping from an IP 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. VRRP Overview

VRRP specifies an election protocol to provide the virtual router function described earlier. All protocol messaging is performed using IP multicast datagrams, thus the protocol can operate over a variety of multiaccess LAN technologies supporting IP 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 IP addresses. A VRRP router may associate a virtual router with its real addresses 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 addresses 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

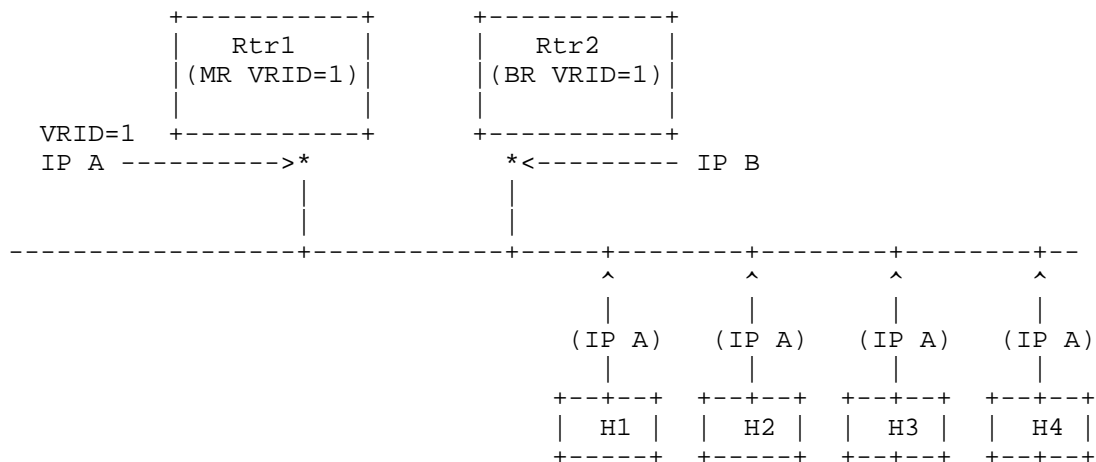
preemption attempts. The only exception is that a VRRP router will always become Master of any virtual router associated with addresses it owns. If the Master becomes unavailable then the highest priority Backup will transition to Master after a short delay, providing a controlled transition of the virtual 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. 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
* = IP Address
(IP) = default router for hosts

Eliminating all mention of VRRP (VRID=1) from the figure above leaves it as a typical IP deployment. Each router is permanently assigned an IP address on the LAN interface (Rtr1 is assigned IP A and Rtr2 is assigned IP B), and each host installs a static default route through one of the routers (in this example they all use Rtr1's IP A).

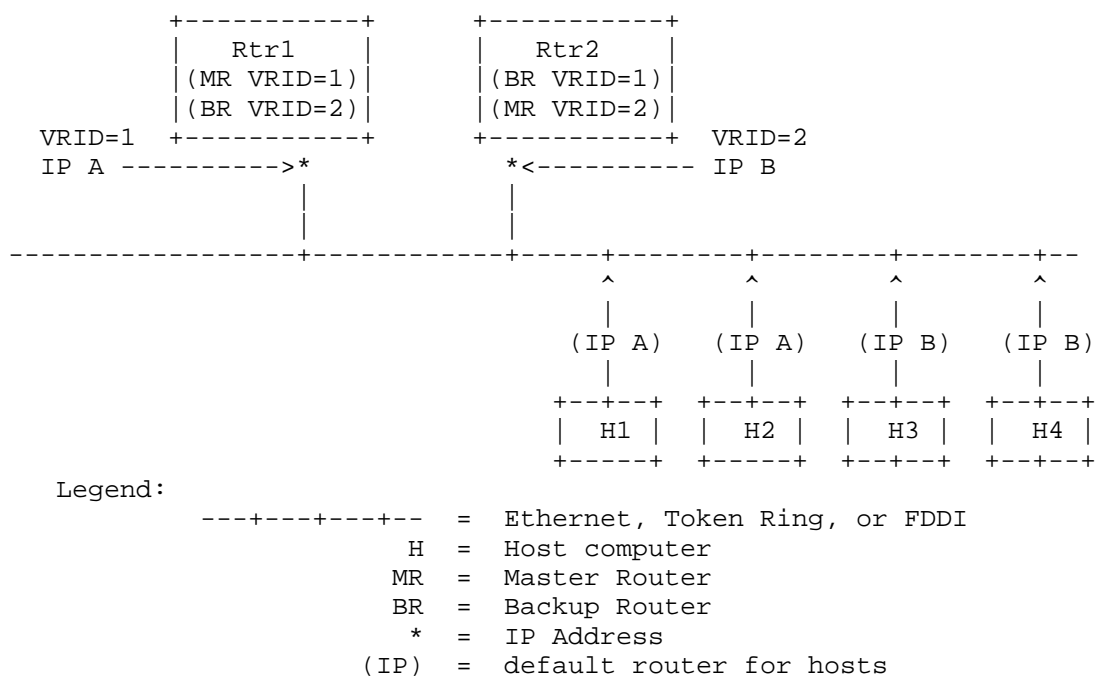
Moving to the VRRP environment, each router has the exact same permanently assigned IP address. Rtr1 is said to be the IP address owner of IP A, and Rtr2 is the IP address owner of IP B. A virtual router is then defined by associating a unique identifier (the virtual router ID) with the address owned 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 IP address owned by Rtr1 (VRID=1, IP_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, since it is not the IP address owner. If Rtr1 should fail then the VRRP protocol will transition Rtr2 to Master, temporarily taking over forwarding responsibility for IP A to provide uninterrupted service to the hosts.

Note that in this example IP B is not backed up, it is only used by Rtr2 as its interface address. In order to backup IP 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 spitting their traffic between them. This example is expected to be very common in actual practice.



In the example above, half of the hosts have configured a static route through Rtr1's IP A and half are using Rtr2's IP 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 IP address owned by Rtr2 (VRID=2, IP_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.

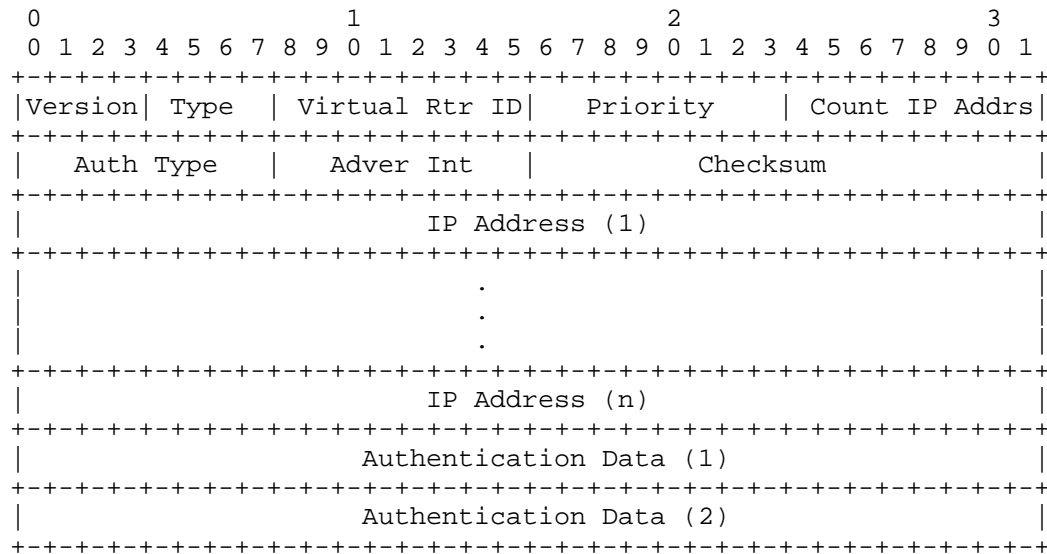
5. Protocol

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 IP packets. They are sent to the IPv4 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 IP header.



5.2. IP Field Descriptions

5.2.1. Source Address

The primary IP address of the interface the packet is being sent from.

5.2.2. Destination Address

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

224.0.0.18

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

5.2.3. 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.

5.2.4. Protocol

The IP protocol number 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 2.

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. Configurable item in the range 1-255 (decimal). There is no default.

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 IP address(es) 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 IP Addr

The number of IP addresses contained in this VRRP advertisement.

5.3.6. Authentication Type

The authentication type field identifies the authentication method being utilized. Authentication type is unique on a Virtual Router basis. The authentication type field is an 8 bit unsigned integer. A packet with unknown authentication type or that does not match the locally configured authentication method MUST be discarded.

Note: Earlier version of the VRRP specification had several defined authentication types [RFC2338]. These were removed in this specification because operational experience showed that they did not provide any real security and would only cause multiple masters to be created.

The authentication methods currently defined are:

- 0 - No Authentication
- 1 - Reserved
- 2 - Reserved

5.3.6.1. Authentication Type 0 - No Authentication

The use of this authentication type means that VRRP protocol exchanges are not authenticated. The contents of the Authentication Data field should be set to zero on transmission and ignored on reception.

5.3.6.2. Authentication Type 1 - Reserved

This authentication type is reserved to maintain backwards compatibility with RFC 2338.

5.3.6.3. Authentication Type 2 - Reserved

This authentication type is reserved to maintain backwards compatibility with RFC 2338.

5.3.7. Advertisement Interval (Adver Int)

The Advertisement interval indicates the time interval (in seconds) between ADVERTISEMENTS. The default is 1 second. This field is used for troubleshooting misconfigured routers.

5.3.8. 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. For computing the checksum, the checksum field is set to zero. See RFC 1071 for more detail [CKSM].

5.3.9. IP Address(es)

One or more IP addresses that are associated with the virtual router. The number of addresses included is specified in the "Count IP Addr's" field. These fields are used for troubleshooting misconfigured routers.

5.3.10. Authentication Data

The authentication string is currently only used to maintain backwards compatibility with RFC 2338. It SHOULD be set to zero on transmission and ignored on reception.

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 IP addresses 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).

IP_Addresses One or more IP addresses associated with this virtual router. Configured item. No default.

Advertisement_Interval Time interval between ADVERTISEMENTS (seconds). Default is 1 second.

Skew_Time Time to skew Master_Down_Interval in seconds. Calculated as:

$$((256 - \text{Priority}) / 256)$$

Master_Down_Interval Time interval for Backup to declare Master down (seconds). 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 IP address(es) associated with the virtual router always preempts independent of the setting of this flag.

Authentication_Type Type of authentication being used. Values are defined in section 5.3.6.


Authentication_Data Authentication data specific to the Authentication_Type being used.

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 IP address(es) associated with the virtual router)
 - o Send an ADVERTISEMENT
 - o Broadcast a gratuitous ARP request containing the virtual router MAC address for each IP address associated with the virtual router.
 - 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 ARP requests for the IP address(s) associated with 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 IP 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} stateendif
- If the Master_Down_Timer fires, then:
 - o Send an ADVERTISEMENT
 - o Broadcast a gratuitous ARP request containing the virtual router MAC address for each IP address associated with the virtual router
 - o Set the Adver_Timer to Advertisement_Interval
 - o Transition to the {Master} stateendif
- If an ADVERTISEMENT is received, then:

If the Priority in the ADVERTISEMENT is Zero, then:

 - o Set the Master_Down_Timer to Skew_Timeelse:

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 IP address(es) associated with the virtual router.

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

- MUST respond to ARP requests for the IP address(es) associated with the virtual router.
 - MUST forward packets with a destination link layer MAC address equal to the virtual router MAC address.
 - MUST NOT accept packets addressed to the IP address(es) associated with the virtual router if it is not the IP address owner.
 - MUST accept packets addressed to the IP address(es) associated with the virtual router if it is the IP address owner.
 - 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 primary IP Address of the sender is greater
    than the local primary IP 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 IP TTL is 255.
- MUST verify the VRRP version is 2.
- MUST verify that the received packet contains the complete VRRP packet (including fixed fields, IP Address(es), and Authentication Data).
- MUST verify the VRRP checksum.
- MUST verify that the VRID is configured on the receiving interface and the local router is not the IP Address owner (Priority equals 255 (decimal)).
- MUST verify that the Auth Type matches the locally configured authentication method for the virtual router and perform that authentication method.

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 IP Addrs" and the list of IP Address matches the IP_Addresses 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. 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 MUST discard the packet, 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
- Set the source MAC address to Virtual Router MAC Address
- Set the source IP address to interface primary IP address
- Set the IP 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-00-01- $\{VRID\}$ (in hex in internet standard bit-order)

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

8. Operational Issues

8.1. 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 IP source address of an ICMP 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.

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.

8.2. Host ARP Requests

When a host sends an ARP request for one of the virtual router IP addresses, the Master virtual router MUST respond to the ARP request 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 VRRP router restarts or boots, it SHOULD not send any ARP messages with its physical MAC address for the IP address it owns, it should only send ARP messages that include Virtual MAC addresses. This may entail:

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

8.3. 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.

8.4. Potential Forwarding Loop

A VRRP router SHOULD not forward packets addressed to the IP Address(es) 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 IP TTL expires.

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

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 >> (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. ARP 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 ARP cache. Hence, these implementations need 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 ARP cache and do not listen to gratuitous ARPs, these hosts will not update their ARP 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 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 does not currently include any type of authentication. Earlier versions of the VRRP 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 ARP replies, answer ARP 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. Acknowledgements

The authors would like to thank Glen Zorn, and 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.

12. References

12.1. Normative References

- [802.1D] International Standard ISO/IEC 10038: 1993, ANSI/IEEE Std 802.1D, 1993 edition.
- [CKSM] Braden, R., Borman, D. and C. Partridge, "Computing the Internet checksum", RFC 1071, September 1988.

- [HSRP] Li, T., Cole, B., Morton, P. and D. Li, "Cisco Hot Standby Router Protocol (HSRP)", RFC 2281, March 1998.
- [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.
- [RFC1469] Pusateri, T., "IP Multicast over Token Ring Local Area Networks", RFC 1469, June 1993.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [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.
- [TKARCH] IBM Token-Ring Network, Architecture Reference, Publication SC30-3374-02, Third Edition, (September, 1989).

12.2. Informative References

- [DISC] Deering, S., Ed., "ICMP Router Discovery Messages", RFC 1256, September 1991.
- [DHCP] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.
- [OSPF] Moy, J., "OSPF version 2", STD 54, RFC 2328, April 1998.
- [RIP] Malkin, G., "RIP Version 2", STD 56, RFC 2453, November 1998.

13. Changes from RFC 2338

- Moved authors of RFC 2338 to new Contributors section to comply with RFC editor policy and listed R. Hinden as Editor.
- Removed authentication methods from VRRP. Changes included:
 - o Removed the values for password and IPSEC based authentication. The fields and values are retained to keep backwards compatibility with RFC 2338.
 - o Removed section on extensible security
 - o Updated security consideration section to remove discussion of different authentication methods and added new text explaining motivation for change and describe vulnerabilities.

- Revised the section 4 examples text with a clearer description of mapping of IP 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 (8.4) describing need to not forward packets for adopted IP addresses.
- Added clarification to the security considerations section.
- Added reference for computing the internet checksum.
- Updated references and author information.
- Various small editorial changes.

14. Editor's Address

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

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

15. Full Copyright Statement

Copyright (C) The Internet Society (2004). This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights 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; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat 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 implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

