INTERNET-DRAFT
July 28, 1997

S. Knight
D. Weaver
Ascend Communications, Inc.
D. Whipple
Microsoft, Inc.
R. Hinden
D. Mitzel
Ipsilon Networks, Inc.

Virtual Router Redundancy Protocol

<draft-ietf-vrrp-spec-01.txt>

Status of this Memo

Abstract

This memo defines the Virtual Router Redundancy Protocol (VRRP). VRRP specifies an election protocol that dynamically assigns responsibility for a virtual IP address to a single router among a collection of VRRP routers. The VRRP router controlling the virtual IP address is called the Master router, and forwards packets sent to the virtual IP address. The election process provides dynamic fail over in the forwarding responsibility should the Master become unavailable. The virtual IP address can then be used as the default

first hop router by end-hosts. The advantage gained from using the VRRP virtual IP address is a higher availability default path without requiring configuration of dynamic routing or router discovery protocols on every end-host.

This memo describes the features and theory of operation of VRRP. The protocol processing and state machine that guarantee convergence to a single Master router is presented. Also issues related to MAC address mapping, handling ARP requests, generating ICMP redirects, and security issues are addressed.

Table of Contents

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, which 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 IP address to a single router among a collection of VRRP routers. The VRRP router controlling the virtual IP address is called the Master router, and forwards packets sent to the virtual IP address. The election process provides dynamic fail-over in the forwarding responsibility should the Master become unavailable. The virtual IP address can then be used as the default first hop router by end-hosts. The advantage gained from using the VRRP virtual IP address 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 a Cisco Systems, Inc. proprietary protocol named Hot Standby Router Protocol (HSRP) [HSRP].

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this

document are to be interpreted as described in [RFC 2119].

1.1 Scope

The remainder of this document describes the features, design goals, and theory of operation of VRRP. The message formats, protocol processing rules and state machine that guarantee convergence to a single Master router 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.

1.2 Definitions

Cluster	The set of routers participating in VRRP to emulate a virtual router.
Master Router	The VRRP router controlling the virtual IP address and assuming the responsibility of forwarding packets sent to the virtual router.
Backup Router	The set of routers in the quiescent state with regard to the virtual router operation. This set includes all active VRRP routers within a cluster that are not the Master router.

2.0 Required Features

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

2.1 Virtual IP Management

Management of the virtual IP address is the primary function of the virtual router protocol. While providing election of a Master router 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 or in 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 more preferential than the current Master. It may be useful to support an override of the immediate convergence to the preferred path.

2.4 Extensible Security

The virtual router functionality is applicable to a wide range of internetworking environments that may employ different security policies. The protocol should require minimal configuration and overhead in the insecure operation, provide for strong authentication when increased security is required, and allow integration of new security mechanisms without breaking backwards compatible operation.

2.5 Efficient Operation over Extended LANs

Sending IP packets on a multiaccess LAN requires mapping from the virtual 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.

3.0 VRRP Overview

VRRP assumes that each router has a consistent set of routes. The mechanism used to learn or configure this routing state and ensure its consistency is beyond the scope of this specification.

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 messages sent by the Master router to enable bridge learning in an extended LAN.

A virtual router is identified by its virtual IP address, and associated with a VRRP cluster. The virtual IP address must not match the real IP address of any host or the virtual IP address of any other VRRP cluster on the LAN. Each VRRP router assigned to the cluster must be configured with the same virtual IP address and must have a real IP address with a prefix matching the virtual router address. In addition, each VRRP router is assigned a priority to indicate the preference for Master election. Multiple virtual routers can be elected on a network by associating them with different VRRP clusters, and a single router can participate in multiple VRRP clusters by maintaining independent state machines for each cluster.

To minimize network traffic, only the Master router sends periodic Advertisement messages. A Backup router will not attempt to pre-empt

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 pre-emption attempts. 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.

VRRP defines three types of authentication providing simple deployment in insecure environments, added protection against misconfiguration, and strong sender authentication in security conscious environments. Analysis of the protection provided and vulnerability of each mechanism is deferred to Section 11.0 Security Considerations. In addition new authentication types and data can be defined in the future without affecting the format of the fixed portion of the protocol packet, thus preserving backward compatible operation.

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 in a VRRP cluster (i.e., a Master and one Backup), 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 (<< 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 VRRP network.

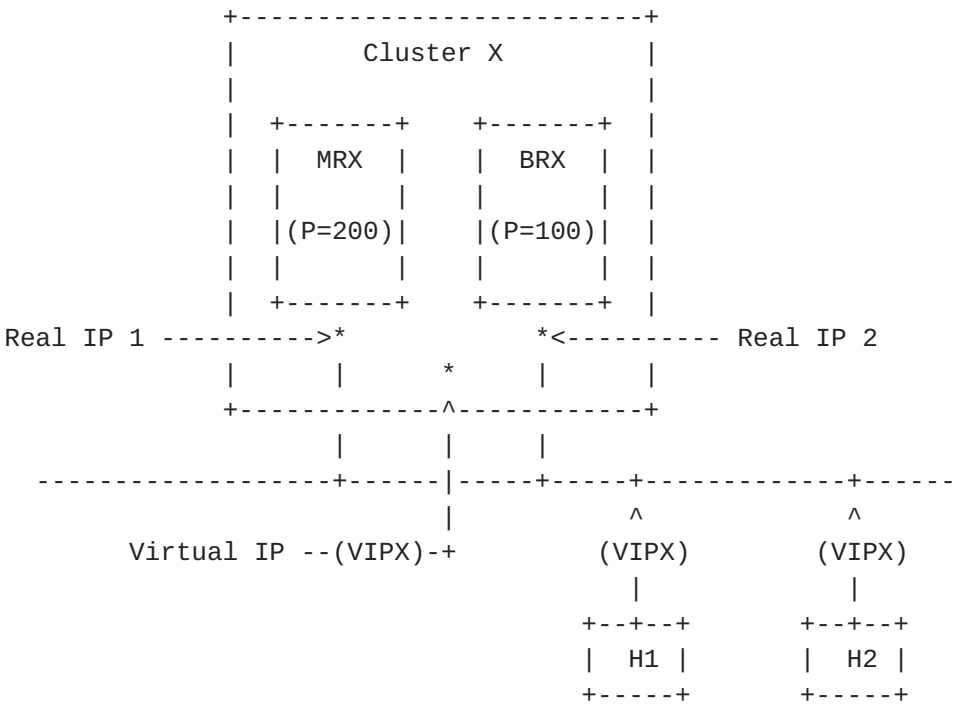

Legend:

```
---+---+---+--  =  802 network, Ethernet or FDDI
             H  =  Host computer
            MR  =  Master Router (Priority=200)
            BR  =  Backup Router (Priority=100)
             *  =  IP Address
           VIP  =  default router for hosts (Virtual IP)
```

The above configuration shows a typical VRRP scenario. In this configuration, the end-hosts install a default route to the virtual IP address (VIPX), and the routers run VRRP to elect the Master router. The router on the left (MRX) becomes the Master router because it has the highest priority and the router on the right (BRX) becomes the backup router.

4.2 Sample Configuration 2

The following figure shows a configuration with two clusters.

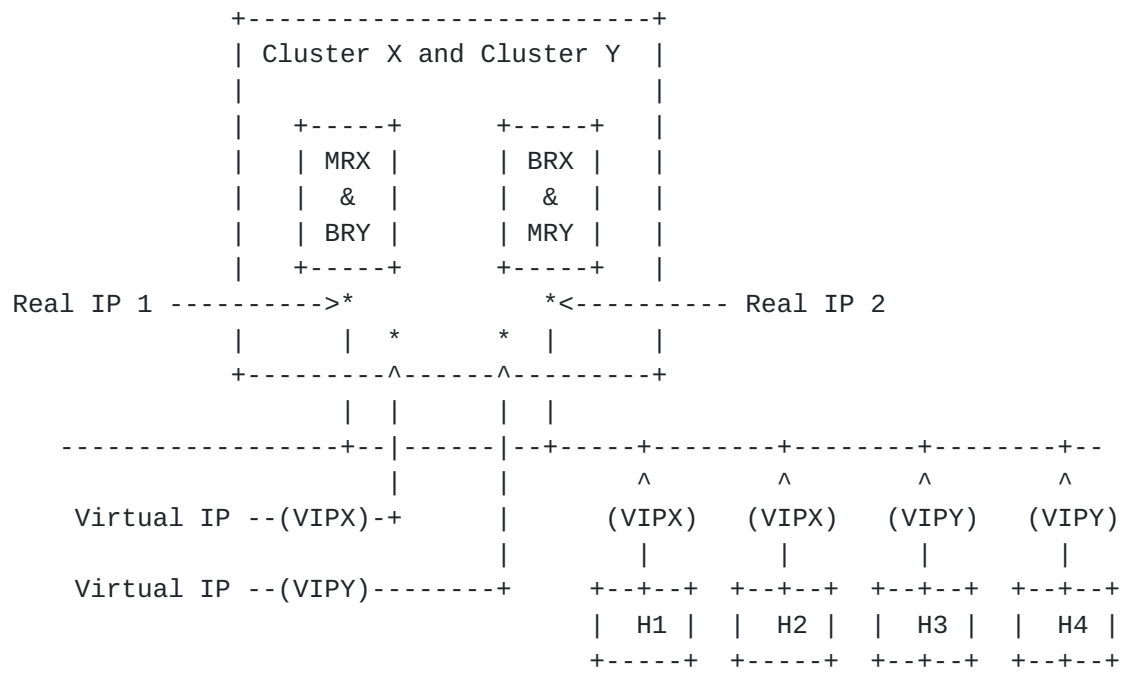

Legend:

---+---+---+--	=	802 network, Ethernet or FDDI
H	=	Host computer
MR	=	Master Router
BR	=	Backup Router
*	=	IP Address
VIP	=	default router for hosts (Virtual IP)

In the above configuration, half of the hosts install a default route to cluster X's virtual IP address (VIPX), and the other half of the hosts install a default route to cluster Y's virtual IP address (VIPY). This has the effect of load balancing the outgoing traffic, while also providing full redundancy.

5.0 Protocol

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

VRRP packets are sent encapsulated in IP packets. They are sent to an 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.

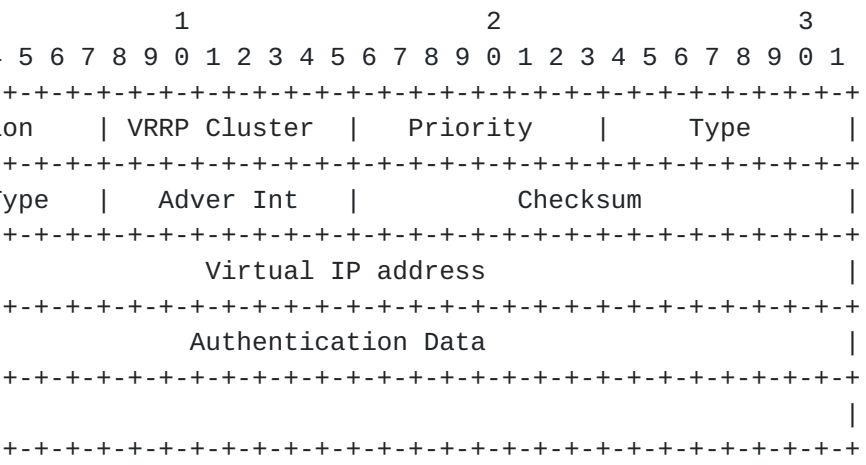

5.2 IP Field Descriptions

5.2.1 Source Address

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

5.2.2 Destination Address

The VRRP IP multicast address assigned by the IANA. It is defined to be:

224.0.0.(TBD IANA assignment)

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 VRRP IP protocol number assigned by the IANA. It is defined to be (TBD).

5.3 VRRP Field Descriptions

5.3.1 Version

The version field specifies the VRRP protocol version of this packet. This document defines version 1.

5.3.2 VRRP Cluster

The VRRP Cluster field specifies the cluster this packet applies to. Note: The interface may participate in more than one VRRP cluster simultaneously, perhaps serving as Master in one cluster, while simultaneously serving as backup in other clusters.

5.3.3 Priority

The priority field specifies the router's priority for the Virtual IP address and cluster. Higher values equal higher priority. This field is an 8 bit unsigned field, giving 1 as the minimum priority, and 255 as the maximum priority. The default priority is 100 (decimal).

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

In the event that two or more routers within a cluster have equal priority, and that priority is the highest priority for the cluster, initially the router with the higher real interface IP address (interpreted as a 32 bit unsigned integer) will become Master. Any router joining the cluster with the same priority will not become Master even if it has a higher IP address unless the current Master goes down.

5.3.4 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.

5.3.5 Authentication Type

The authentication type field identifies the authentication method being utilized. The authentication type field is an 8 bit number. A packet with unknown authentication type or that does not match the locally configured authentication method MUST be discarded.

The authentication methods currently defined are:

0 - No Authentication
1 - Simple Text Password
2 - IP Authentication Header

5.3.5.1 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.

5.3.5.2 Simple Text Password

The use of this authentication type means that VRRP protocol exchanges are authenticated by a clear text password. The contents of the Authentication Data field should be set to the locally configured password on transmission. There is no default password. The receiver MUST check that the Authentication Data in the packet matches its configured authentication string. Packets that do not match MUST be discarded.

5.3.5.3 IP Authentication Header

The use of this authentication type means the VRRP protocol exchanges are authenticated using the mechanisms defined by the IP Authentication Header [AUTH] using HMAC: Keyed-Hashing for Message Authentication [HMAC]. Keys may be either configured manually or via a key distribution protocol.

If a packet is received that does not pass the authentication check due to a missing authentication header or incorrect message digest, then the packet MUST be discarded. The contents of the Authentication Data field should be set to zero on transmission and ignored on reception.

5.3.6 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.

5.3.7 Checksum

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

The checksum is the 16-bit one's complement of the one's complement sum of the entire VRRP message starting with the version field. For computing the checksum, the checksum field is set to zero.

5.3.8 Virtual IP address

The virtual IP address field specifies the Virtual IP (VIP) address associated with the particular cluster. This field is used for troubleshooting misconfigured routers.

The VIP MUST be an IP address assigned from the subnet that the interface is attached and does not match any hosts real IP or cluster VIP address.

5.3.9 Authentication Data

The authentication string is currently only utilized for simple text authentication, similar to the simple text authentication found in the Open Shortest Path First routing protocol [OSPF]. It is up to 8 characters of plain text. If the configured authentication string is shorter than 8 bytes, the remaining space MUST be zero-filled. Any VRRP packet with an authentication string that does not match its configured authentication string SHOULD be discarded. The authentication string is unique on a per interface basis.

There is no default value for this field.

6. Protocol State Machine

6.1 Parameters

Cluster_ID	Cluster identifier. Configured item. There is no default.
Priority	Priority value for this cluster. Configured item. Range is between 1-255. Default is 100 (decimal).
Virtual_IP	Virtual IP Address for this cluster. Configured item.
Advertisement_Interval	Time interval between ADVERTISEMENTS in seconds. Default is 1 second.
Skew_Time	Calculated time to skew Master_Down_Interval in seconds. Defined to be: ((256 - Priority) / 256)
Master_Down_Interval	Time interval for Backup to declare Master down in seconds. Defined to be: (3 * Advertisement_Interval) + Skew_time
Preempt_Mode	Configuration switch controlling whether a higher priority VRRP router preempts a lower priority VRRP Master. Values are True to preempt and False to not preempt. Default is True.

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

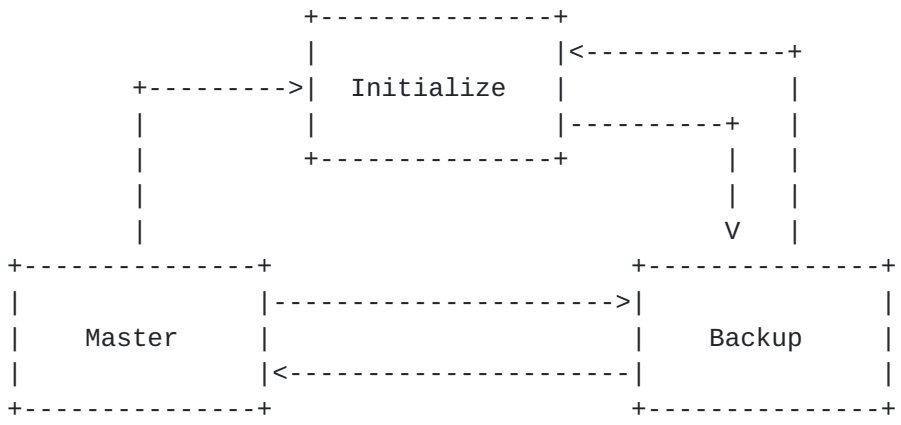

6.4 State Descriptions

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

6.4.1 Initialize

{Initialize} is the state a virtual router takes when VRRP is inactive. The purpose of this state is to wait for a Startup event. If a Startup event is received, then:

- Set the Master_Down_Timer to Master_Down_Interval

- Transition to the {Backup} state

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, an virtual router MUST do the following:

- MUST NOT respond to ARP requests for the virtual router IP address

- MUST discard packets with a destination link layer MAC address equal to the virtual router MAC address

- MUST not accept packets addressed to the Virtual IP address

```
    - If a Shutdown event is received, then:

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

      endif

    - If the Master_Down_Timer fires, then:

       o Send an ADVERTISEMENT
       o Set the Adver_Timer to Advertisement_Interval
       o Transition to the {Master} state

      endif

    - If an ADVERTISEMENT is received, then:

         If the Priority in the ADVERTISEMENT is Zero, then:

          o Set the Master_Down_Timer to Skew_Time

         else:

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

             o Reset the Master_Down_Timer to Master_Down_Interval

            else:

             o Discard the ADVERTISEMENT

            endif
         endif
      endif
```

6.4.3 Master

While in the {Master} state the router functions as the physical router for the Virtual IP address.

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

- MUST respond to ARP requests for the VIP address with the virtual router MAC address

```
    - Must accept and forward packets with a destination link layer MAC
      address equal to the virtual router MAC address

    - Must accept packets addressed to the VIP address

    - 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 IP Address of the sender is greater than
            the local 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

The following actions MUST be performed when a VRRP packet is received:

- Verify that the IP TTL is 255.
- Verify that the received packet length is greater than or equal to the VRRP header length
- Verify the VRRP checksum
- Verify the VRRP version
- Perform authentication specified by Auth Type

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.

- Verify that the Cluster identifier and the VIP are valid on the receiving interface
- Verify that the VIP in packet is same as the configured VIP for this cluster

If any one of the above checks fails, the receiver MUST discard the packet.

- 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 an error occurred.

7.2 Transmitting Packets

The following operations MUST be performed prior to 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
- Send the VRRP packet to the VRRP IP multicast group

Note: VRRP packets are transmitted with the virtual 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-XX-XX-{cluster id} (in hex in internet standard bit-order)

The first three octets are derived from the IANA's OUI. The next two octets (to be assigned by the IANA) indicate the address block assigned to the VRRP protocol. {cluster id} is the VRRP cluster identifier. This mapping provides for up to 255 VRRP clusters on a network.

8. Host Operation

8.1 Host ARP Requests

When a host sends an ARP request for the virtual IP address, the Master router MUST respond to the ARP request with the virtual MAC address for the virtual router. This allows the client to always use the same MAC address regardless of the current Master router. The request MUST be handled as a standard ARP reply.

9. Operational Issues

9.1 ICMP Redirects

VRRP operation relies on hosts only using the Virtual IP address. It is important that client hosts do not learn the real IP address of any VRRP router on the LAN segment. Consequently VRRP routers MUST NOT send ICMP Redirects on any interface they are running VRRP on.

9.2 Proxy ARP

If Proxy ARP is to be used on a router running VRRP, 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 IP address of the VRRP routers.

9.3 Network Management

It is important that network management tools (e.g., SNMP, Telnet, etc.) always use the real IP addresses of a VRRP router. This ensures that network management is aware of the status of the real

routers (e.g., to detect that a router has failed so that it can be repaired).

10. Operation over FDDI and Token Ring

10.1 Operation over FDDI

FDDI interfaces strip 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 stripped off the ring during the Master convergence, and convergence will fail.

To avoid this an implementations 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 stripping any ADVERTISEMENTS it did not originate.

10.2 Operation over Token Ring

Token Ring has several characteristics which make running VRRP problematic. This includes:

- No general multicast mechanism. Required use of "functional addresses" as a substitute, which may collide with other usage of the same "functional addresses".
- Token Ring interfaces may have a limited ability to receive on multiple MAC addresses.
- In order to switch to a new master located on a different physical ring from the previous master when using source route bridges, a mechanism is required to update cached source route information.

Due the these issues and the limited knowledge about the detailed operation of Token Ring by the authors, this version of VRRP does not work over Token Ring networks. This may be remedied in new version of this document, or in a separate document.

11. Security Considerations

VRRP is designed for a range of internetworking environments that may employ different security policies. The protocol includes several authentication methods ranging from no authentication, simple clear text passwords, and strong authentication using IP Authentication with HMAC. The details on each approach including possible attacks and recommended environments follows.

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.

11.1 No Authentication

The use of this authentication type means that VRRP protocol exchanges are not authenticated. This type of authentication SHOULD only be used in environments were there is minimal security risk and little chance for configuration errors (e.g., two VRRP routers in a single cluster on a link).

11.2 Simple Text Password

The use of this authentication type means that VRRP protocol exchanges are authenticated by a simple clear text password.

This type of authentication is useful to protect against accidental misconfiguration of routers on a link. It protects against routers inadvertently becoming a member of a VRRP cluster. A new router must first be configured with the correct password before it can become a member of the VRRP cluster. This type of authentication does not protect against hostile attacks where the password can be learned by a node snooping VRRP packets on the link. The Simple Text Authentication combined with the TTL check makes it difficult for a VRRP packet to be sent from another link to disrupt VRRP operation.

This type of authentication is RECOMMENDED when there is minimal risk of nodes on the link actively disrupting VRRP operation.

11.3 IP Authentication Header

The use of this authentication type means the VRRP protocol exchanges are authenticated using the mechanisms defined by the IP Authentication Header [AUTH] using HMAC: Keyed-Hashing for Message

Authentication [HMAC]. This provides strong protection against configuration errors, replay attacks, and packet corruption/modification.

This type of authentication is RECOMMENDED when there is limited control over the administration of nodes on the link. While this type of authentication does protect the operation of VRRP, there are other types of attacks that may be employed on shared media links (e.g., generation of bogus ARP replies) which are independent from VRRP and are not protected.

12. References

[AUTH] Atkinson, R., "IP Authentication Header", RFC-1826, August 1995.

[DISC] Deering, S., "ICMP Router Discovery Messages", RFC-1256, September 1991.

[DHCP] Droms, R., "Dynamic Host Configuration Protocol", RFC-1541, October 1993.

[HMAC] Krawczyk, H., M. Bellare, R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", RFC-2104, February 1997.

[HSRP] Li, T., B. Cole, P. Morton, D. Li, "Hot Standby Router Protocol (HSRP)", Internet Draft, <draft-li-hsrp-00.txt>, June 1997.

[OSPF] Moy, J., "OSPF version 2", RFC-1583, July 1997.

[RIP] Hedrick, C., "Routing Information Protocol" , RFC-1058, June 1988.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", RFC-2119, BCP14, March 1997.

13. Author's Addresses

Steven Knight
Ascend Communications
High Performance Network Division
10250 Valley View Road, Suite 113
Eden Prairie, MN USA 55344
USA

Phone: +1 612 943-8990
EMail: Steven.Knight@ascend.com

Douglas Weaver
Ascend Communications
High Performance Network Division
10250 Valley View Road, Suite 113
Eden Prairie, MN USA 55344
USA

Phone: +1 612 943-8990
EMail: Doug.Weaver@ascend.com

David Whipple
Microsoft Corporation
One Microsoft Way
Redmond, WA USA 98052-6399
USA
Phone: +1 206 703-3876
EMail: dwhipple@microsoft.com

Robert Hinden
Ipsilon Networks, Inc.
232 Java Drive
Sunnyvale, CA 94089
USA
Phone: +1 408 990-2004
EMail: hinden@ipsilon.com

Danny Mitzel
Ipsilon Networks, Inc.
232 Java Drive
Sunnyvale, CA 94089
USA
Phone: +1 408 990-2037
EMail: mitzel@ipsilon.com

14. Acknowledgments

The authors would like to thank Glen Zorn, and Michael Lane, Clark Bremer, Hal Peterson, Peter Hunt, Tony Li, Barbara Denny, and Steve Bellovin for their comments and suggestions.

15. Changes from Previous Drafts

Changes from <draft-ietf-vrrp-spec-00.txt>

- Added Preempt_Mode to allow user control over preemption independent of configured priorities.
- Rewrote authentication section and expanded security considerations.
- Expanded State Description section and removed State Table which become redundant and impossible to edit.
- Changed authentication to be on a per interface basis (not per cluster).
- Clarified text on disabling ICMP Redirects.
- Added text on FDDI and Token Ring issues.
- Added HSRP acknowledgment.
- Rewrote Introduction, Required Features, and VRRP Overview sections.
- Many small text clarifications.

Changes from <draft-hinden-vrrp-00.txt>

- Changed default behavior to stay with current master when priorities are equal. This behavior can be changed by configuring explicit priorities.
- Changed Master state behavior to not send Advertisements when receiving Advertisement with lower priority. Change reduces worst case election message overhead to "n", where "n" is number of configured equal priority VRRP routers.
- Added Skew_Time parameter and changed receiving advertisement with zero priority behavior to cause resulting advertisement sent to be skewed by priority.
- Changed sending behavior to send VRRP packets with VMAC as source MAC and added text describing why this is important for bridged environments.
- Changed definition of VMAC to be in IANA assigned unicast MAC block.
- Added Advertisement Interval to VRRP header.
- Added text regarding ICMP Redirects, Proxy ARP, and network management issues.
- Various small text clarifications.