# VRRP White paper

## Virtual Router Redundancy Protocol (VRRP)

The Virtual Router Redundancy Protocol (VRRP) is designed to eliminate the single point of failure inherent in the static default routed environment by automatically providing alternate router paths. As specified by RFC 2338, VRRP uses 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 failover of the forwarding responsibility should the Master become unavailable. The Virtual Router associated with a given alternate path supported by VRRP uses the same IP address and MAC address as the routers for other paths. As a result, the host's gateway information does not change, no matter what path is used. Because of this, VRRP-based redundancy significantly reduces administrative overhead when compared to redundancy schemes that require hosts to be configured with multiple default gateways.

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, RFC 2338 states that 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 multi-access LAN technologies capable of supporting IP traffic.
- Provide for election of multiple virtual routers on a network for load balancing
- Support multiple logical IP subnets on a single LAN segment.

## VRRP Terminology

Each physical router running VRRP is known as a VRRP Router. Two or more VRRP Routers can be configured to form a Virtual Router. Each VRRP Router may participate in one or more Virtual Routers. Additionally, each VRRP-capable device is autonomous.
There is no requirement that Virtual Routers be identically configured. Different router models with different numbers of ports and different enabled services may be used in a Virtual Router.

A Virtual Router acts as a default or next hop gateway for hosts on a LAN. The Virtual Router is managed by each of the routers running VRRP. Each Virtual Router consists of a user-configured Virtual Router Identifier (VRID) and an IP address or set of IP addresses on the shared LAN.

The VRID is used to build the Virtual Router MAC Address. The five highest order octets of the Virtual Router MAC Address are the standard MAC prefix (00-00-5E-00-01) defined in RFC 2338. The VRID is used to form the lowest order octet (the last two digits in the MAC address). One, but not more than one, of the VRRP Routers in a Virtual Router may be configured as the IP Address Owner. This router has the Virtual Router's IP address as its real interface address. This router, when up, responds to packets addressed to the Virtual Router's IP address for ICMP pings, TCP connections, etc.

The Virtual Router Master forwards packets sent to the Virtual Router. It also responds to ARP requests the Virtual Router's IP address. Finally, it sends out periodic advertisements to let other VRRP Routers know it is alive and its priority (explained below). Within a Virtual Router, the VRRP routers not selected to be the Master are known as Virtual Router Backups. Should the Virtual Router Master fail, one of the Virtual Router Backups becomes the Master and assumes its responsibilities.

## A summary of these definitions is contained below in Table 1 below

| Component | Definition |
|---|---|
| Virtual Router: | The IP address or set of addresses shared by the VRRP Routers and addressed by a unique identifier. |
| VRRP Router: | A router running Virtual Router Redundancy Protocol. |
| IP Address Owner: | The VRRP router that has the Virtual Router's IP address configured as its real interface address. There cannot be more than one IP Address Owner in a Virtual Router. An optional designation, the IP Address Owner will always win the election process to become Master in a Virtual Router. |
| Renter: | All VRRP Routers that do not have the Virtual Router's IP address configured as their real interface address. This designation is not defined by RFC 2338. |
| Master: | Determined by an election process, the Master is the router that assumes responsibility for forwarding packets sent to the Virtual Router and answering ARP requests for the Virtual Router. |
| Backup: | A VRRP Router that is available to assume Master duties should the current Master fail. |

**Figure 1 Simple two-node virtual router using VRRP**

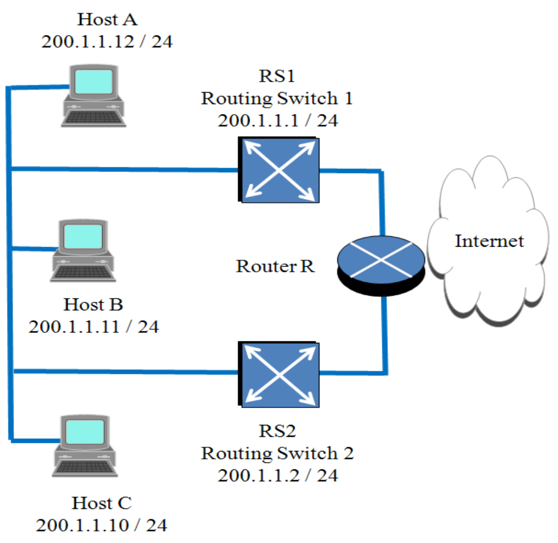

Figure 1 illustrates a simple, two-node Virtual Router using VRRP. The routers in the diagram have been configured as VRRP Routers. They form a Virtual Router, Router/Switch 1 has its real interface configured with the IP address of the Virtual Router, 200.1.1.1/24, and is therefore the IP Address Owner.

As long as this VRRP Router is active and available, it will be the Virtual Router Master. Router/switch 2 below is a Virtual Router Backup. Its real interface is configured with an IP address that is on the same subnet as that of the Virtual Router but that is not the IP address of the Virtual Router. As a result, it is a Renter and is also a Virtual Router Backup.

In the example above, the Virtual Router will also have been assigned a Virtual Router ID (VRID). If the VRID is 1, both of the VRRP Routers will use the MAC address of 00- 00-5E-00-01-01.

## VRRP: How It Works

The hosts shown in Figure 1 are configured with the Virtual Router's IP address as its default gateway. As mentioned above, the Master forwards packets destined to remote subnets and responds to ARP requests. Since, in this example, the Master is also the Virtual Interface Router's IP Address Owner, it also responds to ICMP ping requests and IP datagram’s destined for the Virtual Interface Router's IP address. The Backup does not forward any traffic on behalf of the Virtual Interface Router, nor does it respond to ARP requests.

If the Master (also the Owner in this case) is not available, the Backup becomes the Master and takes over responsibility for packet forwarding and responding to ARP requests. However, since this new Master router is not the IP Address Owner, it does not respond to ICMP ping requests and IP datagram’s destined to that address. Each VRRP Router that is a Renter is configured with a priority between 1 and 254. According to the VRRP standard, an Owner has a priority of 255. In Figure 1 above, once RS1 is configured for VRRP, it looks at the IP address of the virtual router and compares it with the IP addresses of its own interface that is configured for VRRP on that VRID. Since the RS1 owns the Virtual Router's IP address, it declares itself the Master and sends out an advertisement to all of the other VRRP Routers. The IP Address Owner is always the Master as long as it is available. It is not necessary for the Virtual Router IP address to be owned by one of the VRRP Routers connecting the LAN to the Internet in Figure 1. In this case, however, the bidding process to determine the Master is different. The process involves comparing two criteria. First, the VRRP Router with the highest priority becomes the Master. Second, if the priorities are the same, the higher IP address wins and becomes the Master. Understanding the VRRP advertisement packet is key to understanding the bidding process.

## VRRP Packets

VRRP packets are sent encapsulated in IP packets. Figure 2 shows the layout of the packet's IP header and the packet itself.

<table>
<tr><td colspan="7"><b>Figure 2 – Contents of a VRRP Packet</b></td></tr>
<tr><td>0</td><td>4</td><td colspan="2">8</td><td>16 18</td><td></td><td>24 31</td></tr>
<tr><td>VERS</td><td colspan="2">HLEN</td><td>SERVICE TYPE</td><td colspan="3">TOTAL LENGTH</td></tr>
<tr><td colspan="4">IDENTIFICATION</td><td>FLAGS</td><td colspan="2">FRAGMENT OFFSET</td></tr>
<tr><td colspan="3">TIME TO LIVE</td><td>PROTOCOL</td><td colspan="3">HEADER CHECKSUM</td></tr>
<tr><td colspan="7">SOURCE IP ADDRESS</td></tr>
<tr><td colspan="7">DESTINATION IP ADDRESS</td></tr>
<tr><td colspan="6">IP OPTIONS</td><td>PADDING</td></tr>
<tr><td>VERS</td><td>TYPE</td><td colspan="2">VRID</td><td colspan="2">PRIORITY</td><td>COUNT IP ADDRESS</td></tr>
<tr><td colspan="2">AUTH TYPE</td><td colspan="2">ADVERTISEMENT Inter</td><td colspan="3">CHECKSUM</td></tr>
<tr><td colspan="7">IP ADDRESS (1)</td></tr>
<tr><td colspan="7">:<br>:</td></tr>
<tr><td colspan="7">IP ADDRESS (n)</td></tr>
<tr><td colspan="7">AUTHENTICATION DATA (1)</td></tr>
<tr><td colspan="7">AUTHENTICATION DATA (2)</td></tr>
</table>

The important fields in the IP header (in terms of VRRP) are explained below.

**Source IP Address:** This is a 32-bit field. The source address is the primary IP address of the interface from which the packet is being sent. This is the IP address of the master router's interface connected to the LAN.

**Destination IP Address:** This is a 32-bit field. It is the IP multicast address assigned by the IANA for VRRP. This multicast IP address is 224.0.0.18. All the routers running VRRP receive this multicast

**Time To Live (TTL):** This is an 8-bit field; the value in this field must be equal to 255. Any VRRP packet received with TTL not equal to 255 is discarded. The router does not forward a datagram with VRRP multicast destination address, regardless of its TTL.

**Protocol:** This is an 8-bit field that specifies the protocol being used. The IP protocol number assigned by IANA for VRRP is 112.
The following fields are in the VRRP packet:

**Version (VERS):** This is a 4-bit field that specifies the VRRP version. The version that is available is 2.

**Type:** This is a 4-bit field that specifies the type of VRRP packet. The only type is ADVERTISEMENT.

**Virtual Router Identifier (VRID):** Identifies the virtual router for which this packet is reporting status.

**Priority:** This 8-bit field specifies the sending VRRP router's priority for the virtual router. A higher value means a higher priority. The priority value of the VRRP router that owns the IP address associated with the virtual router must be 255. The default priority value is 100, but you can assign any value between 1 and 254. A priority of 0 is a special value that specifies the master has stopped working, and the backup router needs to transition to master state.

**Count IP Addresses:** This 8-bit field specifies the number of IP addresses contained in this VRRP advertisement.

**Authentication Type:** This 8-bit field specifies the authentication type being used. A packet with an unknown authentication type or one that does not match the locally configured authentication type is discarded. The authentication methods defined by the RFC are:

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

No authentication 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. Simple Text Password authentication indicates 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 receiving VRRP Router must check that the Authentication Data in the packet matches its configured authentication string. Packets that do not match are discarded automatically. The use of the IP Authentication Header type means the VRRP protocol exchanges are authenticated using the mechanisms defined by the IP Authentication Header. Although described briefly in the RFC, this authentication type has not been implemented.

**Advertisement Interval:** This 8-bit field specifies the time interval between advertisements sent from the master, to let the backup router know that it is alive. It is important that all routers with the same VRID should have the same advertisement interval.

**Checksum:** This 16-bit field is used to detect data corruption in the VRRP message.

**IP Address(es):** This is a 32-bit field. The IP address is the Virtual Router's IP address that the Master is backing up. Each address associated with the Virtual Router is included in a separate 32-bit field within the announcement. Not all VRRP implementations fully support this part of RFC 2338. Some, like Nortel's Passport, only support sending a single IP address with each announcement.

**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, and must match the locally configured string or be discarded.
Returning to Figure 1 above, the Master sends out an advertisement at a specified interval to the VRRP IP multicast address (224.0.0.18, defined by IANA) declaring itself as the Master. As long as the Backups receive these advertisements, they remain in the backup state.

If a Backup does not receive an advertisement for three advertisement intervals, it starts a bidding process to determine which VRRP Router has the highest priority. The VRRP Router with the highest priority takes over as Master.

It is important that all VRRP routers have a physical interface configured with an IP address in the same subnet as the Virtual Router. The VRRP protocol sends only IP addresses and not subnet information. Without the corresponding subnet information, the VRRP Router will add the Virtual Router address as a single IP address with a host (/32 or 255.255.255.255) netmask. This will prevent routing from working properly, as the Virtual Router will not listen to broadcasts from the local network.

If, at any time, a Backup determines that it has higher priority than the current Master does, it can preempt the Master, unless it is configured not to do so. In pre-emption, the Backup begins to send its own advertisements. The current Master will see that the Backup has higher priority and stop functioning as the Master. The Backup will then see that the Master has stopped sending advertisements and assume the role of Master. While pre-emption can ensure that a primary router will return to Master status once it returns to service, pre-emption also causes a brief outage while the election process takes place. Disabling pre-emption will ensure maximum uptime on the network, but will not always result in the primary, or highest priority, router acting as Master.

## VRRP: Perspective of the Host

All the decisions regarding who is going to be the Master for a particular LAN are made on the VRRP Routers. The host is oblivious to the whole process. When a host must send a message to another host on a different network connected by the VRRP routers, it sends an ARP request for the MAC address of the default gateway.

Normally, when a host “ARPs for” (resolves) the MAC address, the router replies with its own physical address. But when VRRP is deployed, the Master replies with a virtual MAC address instead of its actual MAC address. The benefit of this virtual MAC address is that when the Master goes down and a Backup router becomes the Master, it does not make any difference to the host because it uses the same MAC address.

The virtual MAC address belongs to the virtual IP address, which belongs to the Master for that VRID.

For instance, return to Figure 1 above. Host A wants to send data to a host on the Internet. In this case, RS1 is acting as the Master, and RS2 is the backup. Host A will ARP for the MAC address to the default gateway whose IP address is 200.1.1.1/24. In return, RS1 replies with the virtual router's MAC address (which is 00-00-5E-00-01- <VRID>).

Then, the host sends the packets to this MAC address. The datagrams are then routed out of the LAN and to the Internet. If RS1 goes down, and RS2 takes over as the virtual master, all forwarding and ARP tasks are performed by RS2.

Therefore, when Host A sends an ARP for the MAC address to the default gateway, RS2 replies to that with the virtual router's MAC address (again, 00-00-5E-00-01-<VRID>). Another scenario is that the host already had an ARP table and knows that if it needs to send any information to the 200.1.1.1/24 IP address (which is its default gateway), it will send it to the 00-00-5E-00-01-<VRID> MAC address. So, it sends it to the virtual router's MAC address, and the information flows via RS2 instead of RS1. For the host, it is all the same.

## Applications: Load Sharing

Referring to Figure 1 in our initial example, the Master is forwarding all of the traffic. The other router is an idle Backup. To utilize the bandwidth efficiently, we can create two different VRIDs, sending some of the traffic through RS1 and other traffic through RS2. To do this, we configure RS1 to be the default gateway for a certain number of hosts, and RS2 for the rest of them.

**Figure 3 – Load sharing example**

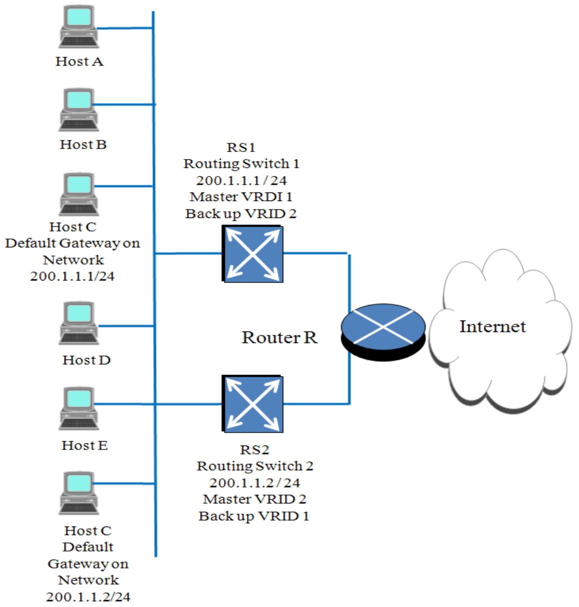

In Figure 3, RS1 is the default gateway for the three hosts at the top, and RS2 is the default gateway for the three hosts at the bottom. There are two VRIDs: 1 and 2. RS1 (with VRID 1) is the Master for host A, B and C, and Backup for the hosts D, E and F. On the other hand, RS2 (with VRID 2) is the Master for the hosts D, E and F, and the Backup for A, B and C. This way, the traffic going out of the LAN 200.1.1.0/24 is shared between the two routers, thus efficiently utilizing the routers and bandwidth.

## Applications: Maintenance Windows without Downtime

VRRP can also be used to eliminate downtime due to maintenance, the leading cause of router downtime on today's networks. An average router requires software upgrades 2 to 4 times a year, requiring a reboot. Depending on the complexity of the configuration, the Mean Time To Restoration (MTTR) can range from 5 to 40 minutes, during which the router is "down" (i.e., not passing traffic).

In Figure 1 above, if RS1 must be upgraded, RS2 can take over routing for temporarily and the hosts see the Virtual Router as “up” (i.e. passing traffic). The following calculations in Table 2 show, for a 500 router network typical of a medium-sized nationwide service provider, an average expected savings 99% of router downtime, or 123.75 hours, in our calculations.

## TABLE 2 Router Downtime Calculations

### Router Downtime Calculations

Assume a Router will encounter, on average, a total of 3 software upgrades or control module failures per year. (Actual range: 2-6) Assume an average router reboot takes 5 minutes (Actual range: 4- 40min.):

**Without VRRP**
3 reboots X 5 min = 15 min/year downtime*

**With VRRP**
3 VRRP Master elections x 3 seconds = 9 sec/year downtime* Assume a Service Provider with 500 routers:

**Without VRRP**
500 routers x 15 min/year/router= 125 hours

**With VRRP**
500 routers x 9 sec/year/router= 1.25 hours

**Savings: 123.75 hours router downtime/year**
* From failures or software upgrades. There may be other sources of downtime not included here.