

**VRRP Extention to Support Dynamic Topology Changes  
of Virtual Router topology  
draft-celer-vrrp-ext-00.txt**

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#) except that the right to produce derivative works is not granted.

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

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Abstract

This draft defines an extention to Virtual Router Redundancy Protocol (VRRP) [1] which limits the amount of provisioning required to support the VRRP operations. It specifically addresses the scenario when the primary owner of the interfaces leaves the Virtual Router group.

Table of Contents

<a href="#">1</a>	Introduction .....	<a href="#">2</a>
<a href="#">2</a>	VRRP State Machine .....	<a href="#">2</a>
<a href="#">3</a>	Sending and Receiving VRRP Packets .....	<a href="#">5</a>
<a href="#">4</a>	Security Considerations .....	<a href="#">6</a>
<a href="#">5</a>	Message Formats .....	<a href="#">6</a>
<a href="#">6</a>	Backward Compatibility .....	<a href="#">8</a>
<a href="#">7</a>	References .....	<a href="#">8</a>
<a href="#">8</a>	Author's Address .....	<a href="#">8</a>

## **1. Introduction**

The Virtual Router Redundancy Protocol (VRRP) specifies an election protocol that dynamically assigns responsibility for a virtual router to one of the VRRP routers on the LAN. The VRRP router controlling the IP address(es) associated with the virtual router is called the Master, and forwards packets sent to these IP addresses. In case of the Master's failure or deliberate decision to leave the group, the back-up virtual router will assume the forwarding responsibility for these IP addresses.

Current protocol definition does not allow for the dynamic removal of the IP address(es) if the router decides to leave group of routers it belongs to. Currently, the router sends the VRRP advertisement message, specifying priority zero for the addresses. Another router in this virtual router group will assume ownership over these addresses. If the IP address(es) are removed from under Virtual Router control, all routers which participate in this virtual router have to be de-provisioned.

The proposed extension to the protocol will allow dynamic removal of IP address(es) associated with router which leaves the virtual router. This draft introduces the new VRRP message type: LEAVE\_GROUP which can be used to remove the addresses from the virtual router topology.

## **2. VRRP State Machine**

Introduction of the new VRRP message types does not change the Virtual Router definition of states in VRRP state machine. New message type processing is included in appropriate sections.

### **2.1 Parameters**

New attribute is added to describe the 'persistence' of the IP address. It can have one of two values: Dynamic and Static.

- Dynamic - defines the IP address which should be removed from the Virtual Router in case that shutdown event is received
- Static - defines the IP address which should remain part of the Virtual Router in case the shutdown event is received.

### **2.2 Timers**

The set of timers does not change.

### **2.3 Initialize State**

The Initialize state does not change.

## **2.4 Backup State**

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} state

endif

- 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} 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

Celer, A

Internet Draft

[Page 3]

```
endif
endif
endif
```

- If a LEAVE\_GROUP is received, then:

```
    if priority in LEAVE_GROUP message is 255, then:
        o Remove the IP address(es) in the message from the database
    endif
endif
```

## [2.5](#) Master State

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 If none of the routers in the Virtual Router group runs VRRPv.2
      - o send an LEAVE\_GROUP message with the list of the interfaces to be removed from the Virtual Router; priority is set to 255
- ```
endif
```
- 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
- If LEAVE\_GROUP is received, then:
  - If Priority in the LEAVE\_GROUP message is 255, then
    - o remove included IP addresses from the database
  - endif
- endif

### **3. Sending and Receiving VRRP Packets**

The rules for sending and receiving VRRP packets do not differ from those defined in [1] .

#### 4. Security Considerations

The rules do not differ from those defined in [1] .

#### 5. Message Formats

This section defines the format of the VRRP packet. Relevant fields in the IP header are the same as described in [1].

```

      0             1             2             3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|Version| Type  | Virtual Rtr ID|   Priority   | Count IP Adrs|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  Auth Type  |  Adver Int  |           Checksum           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|
|                               IP Address (1)
|                               |
|                               .
|                               .
|                               .
|                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|
|                               IP Address (n)
|                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|
|                               Authentication Data (1)
|                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|
|                               Authentication Data (2)
|                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

##### 5.1 Version

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

##### 5.2 Type

The type field specifies the type of this VRRP packet. There are two packet types defined:

- 1        ADVERTISEMENT
- 2        LEAVE\_GROUP

A packet with unknown type MUST be discarded.

##### 5.3 Virtual Rtr ID (VRID)

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

## **5.2 Type**

The type field specifies the type of this VRRP packet. There are two packet types defined:

- |   |               |
|---|---------------|
| 1 | ADVERTISEMENT |
| 2 | LEAVE_GROUP   |

A packet with unknown type MUST be discarded.

## **5.3 Virtual Rtr ID (VRID)**

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

## **5.4 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.

In case of LEAVE\_GROUP message priority is set to 255 .

## **5.5 Count IP Addrs**

The number of IP addresses contained in this VRRP advertisement.

## **5.6 Authentication Type**

This draft does not change definition of the Authentication types as defined in [1].

## **5.7 Advertisement Interval (Adver Int)**

This draft does not change definition of the advertisement interval as defined in [1].



### **5.8 Checksum**

This draft does not change definition of the checksum as defined in [1].

### **5.9 IP Address(es)**

This draft does not change definition of the IP address(es) as defined in [1].

### **5.10 Authentication Data**

This draft does not change definition of the authentication data as defined in [1].

## **6. Backward Compatibility**

This extension is backward compatible.

The router which implements version 3 of VRRP protocol will recognize the presence of version 2 compatible routers in the group, and if it receives the {shutdown} event, it will send the ADVERTISEMENT message with priority zero (0) only. LEAVE\_GROUP message, if removal of its IP address(es) is required will be sent only when all routers which are members of the Virtual Router run VRRP v.3 .

In case that at least one router in Virtual Router group does not recognize LEAVE\_GROUP message, runs VRRP v.2, it will drop the message, and take ownership of the IP address(es) which belonged to router leaving the Virtual Router group. To remove the IP address(es) from the list, all routers in Virtual Router group will have to go through {shutdown} event.

## **7. References**

- [1] [RFC2338](#) Virtual Router Redundancy Protocol

## **8. Author's Address**

Alicja Celer  
Nortel (Northern Telecom), Ltd.  
PO Box 3511 Station C  
Ottawa ON K1Y 4H7  
Canada

EMail: [aceler@nortelnetworks.com](mailto:aceler@nortelnetworks.com)