

Timer Enhancements to Reduce Failover Times for the
Virtual Router Redundancy Protocol for IPv4

<[draft-ietf-vrrp-ipv4-timers-02.txt](#)>

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

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>.

This Internet-Draft will expire on September 6, 2006.

Abstract

The router survivability capability provided by the Virtual Router Redundancy Protocol for IPv4 (VRRPv4) satisfies the requirements for many LAN environments. There are, however, LAN environments that have sub-second failover requirements and thus a need for finer granularity of the VRRP timers. This draft proposes extensions to VRRPv4 [[RFC 3768](#)] for specifying sub-second Advertisement Intervals. A new message type is introduced which permits the timer granularity for the Advertisement Interval to be specified. In addition, a new field is introduced permitting the specification of the number of missed ADVERTISEMENTS before a Virtual Router Master is declared down.

Table of Contents

1.	Introduction.....	3
1.1.	Requirements notation.....	3
1.2.	Contributors.....	3
1.3.	Scope.....	4
2.	Update to the VRRPv4 Protocol.....	4
2.1.	Updates to the VRRPv4 Packet Format.....	4
2.2.	Updates to the VRRPv4 Field Descriptions.....	5
3.	Update to the VRRPv4 Protocol State Machine.....	6
3.1.	Updates to the Parenters per Virtual Router.....	6
3.2.	Update to the Timers.....	8
3.3.	Updates to the State Descriptions.....	8
4.	Updates for Sending and Receiving VRRPv4 Packets.....	15
4.1.	Receiving VRRPv4 Packets.....	15
4.2.	Transmitting VRRPv4 Packets.....	16
5.	Operational Issues.....	17
5.1.	Sub-second Timers.....	17
5.2.	Interoperability / Backward Compatibility.....	17
6.	Security Considerations.....	17
7.	Intellectual Property.....	17
8.	Acknowledgments.....	18
9.	IANA Considerations.....	18
10.	Normative References.....	18
11.	Informative References.....	18
12.	Authors' Address.....	19
13.	Disclaimer of Validity.....	19
14.	Copyright Statement.....	19

1. Introduction

VRRPv4 [[RFC 3768](#)] specifies an election protocol that dynamically assigns responsibility for a virtual router to one of the VRRP routers on a LAN. This election process provides dynamic fail over in the forwarding responsibility should the Master become unavailable. While this capability may meet the survivability requirements for many LAN environments, there are environments in which sub-second recovery from outages is required.

To achieve sub-second failovers for VRRPv4, the granularity of the timers within VRRPv4 must be reduced from one second intervals to sub-second intervals. This document proposes an optional message type which permits the specification of the timer granularity. By specifying finer granularity timers, the Advertisement Interval can be specified, in increments, based upon the timer granularity.

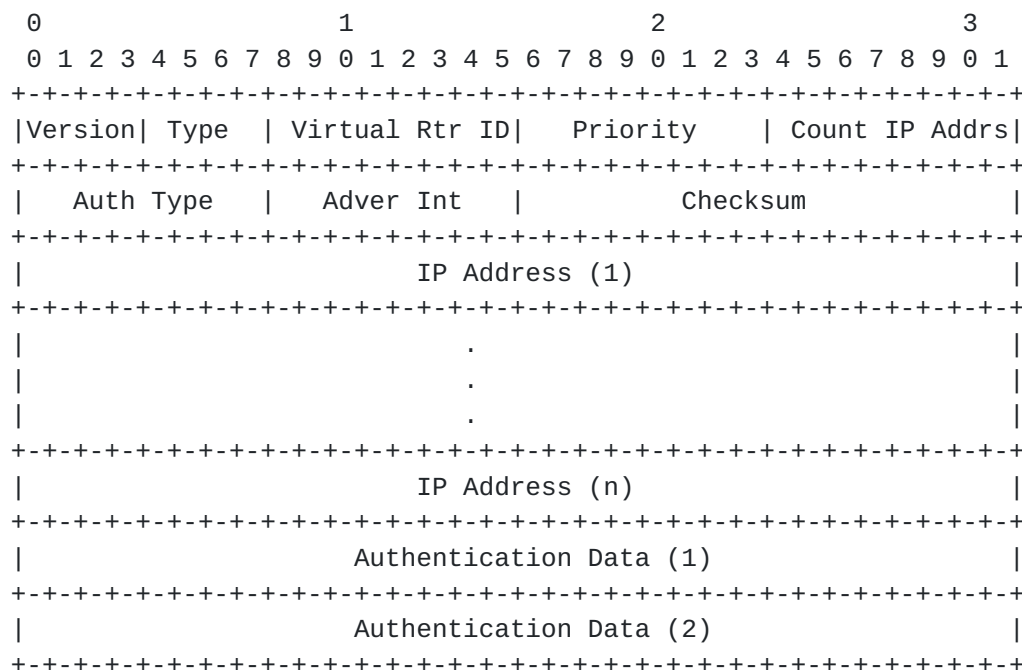
In addition to specifying the timer granularity and the Advertisement Interval in these new time increments, this document also proposes the ability to specify the number of ADVERTISEMENTS that must be missed prior to declaring a MASTER inactive. In the current specification for VRRPv4 [[RFC 3768](#)], the number of ADVERTISEMENTS missed prior to declaring a MASTER inactive is three, based upon the calculation of the Master_Down_Interval. In permitting the sub-second Advertisement Interval, the potential for VRRP instability is increased. Instability could occur due to processing requirements within the router preventing the processing of ADVERTISEMENTS or loading conditions on the network preventing reception of these ADVERTISEMENTS. Specifying the number of ADVERTISEMENTS that can be missed offers a mechanism to address stability issues with VRRPv4, regardless of the timer granularity specified.

1.1. Requirements notation

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](#)].

1.2. Contributors

The following people contributed to the text in this document: K. O'Donoghue, R. Hinden, S. Bates, S. Mathur, and M. Gupta. In addition, the authors of VRRPv4 [[RFC 3768](#)] are recognized as the basis for the text and concepts within this document.



This field is only present in the FAST ADVERTISEMENT, Type 2, message. The Advertisement Count field specifies the number of FAST ADVERTISEMENT messages that can be missed before a BACKUP will declare a MASTER down. Configurable item in the range 1-15. Default is 3.

2.2.3. Advertisement Interval Granularity (AIG)

This field is only present in the FAST ADVERTISEMENT, Type 2, message. The Advertisement Interval Granularity field specifies the timer granularity. The currently recognized values are

- | | |
|---|-------------------|
| 0 | seconds (default) |
| 1 | centiseconds |
| 2 | milliseconds |

A packet with an unknown Advertisement Interval Granularity MUST be discarded.

2.2.4. Advertisement Interval (Adver Int)

The Advertisement Interval field is present in both the ADVERTISEMENT and the FAST ADVERTISEMENT messages. The field indicates the number of time intervals between the respective advertisements. For the ADVERTISEMENT, Type 1, message the interval is measured in seconds and is an 8 bit field. For the FAST ADVERTISEMENT, Type 2, message the interval is measured in the units specified in the AIG field and is a 10 bit field. The default for the Advertisement Interval is 1 second.

3. Update to the VRRPv4 Protocol State Machine

This section outlines the changes to [Section 6](#) (Protocol State Machine) of the VRRPv4 [[RFC 3768](#)] to accommodate the optional FAST ADVERTISEMENT, Type 2, message.

With the introduction of the optional FAST ADVERTISEMENT message, time values MUST reflect the granularity supported in the FAST ADVERTISEMENT message. Thus all time values for both the ADVERTISEMENT and FAST ADVERTISEMENT message are specified according to the message format but are reflected internally in milliseconds.

The state machines have also been updated to deal with interoperability issues due to addition of the new message type.

3.1. Updates to the Parameters per Virtual Router

VR_Type	The message type for this VRRP router Configurable item with values 1 for ADVERTISEMENT and 2 for FAST ADVERTISEMENT.
---------	---

VR_Mode	The mode of operation for this VRRP router environment. Values are 0 for a homogeneous VRRP router environment and 1 for a heterogeneous VRRP router environment (i.e., both message types are used). The default is 0.
VR_AIG	The Advertisement Interval Granularity for this VRRP router. Configurable item with values 0 for seconds (default), 1 for centiseconds, and 2 for milliseconds.
VR_Adver_Interval	Time interval between ADVERTISEMENTS or FAST ADVERTISEMENTS (milliseconds). Configurable item. Default is 1,000 milliseconds (1 second) for both message types.
Advertisement_Interval	Time interval between ADVERTISEMENTS or FAST ADVERTISEMENTS (milliseconds).
Type1_Adver_Interval	Time interval between ADVERTISEMENTS (seconds).
Skew_Time	Time to skew Master_Down_Interval in milliseconds. Calculated as: $\left(\left(256 - \text{Priority} \right) * \text{Advertisement_Interval} \right) / 256$
VR_Adver_Count	The number of ADVERTISEMENTS or FAST ADVERTISEMENTS that can be missed before a BACKUP will declare a MASTER down. Configurable item. When ADVERTISEMENTS are used this is set to 3. When FAST ADVERTISEMENTS are used, the range is 1-15. The default is 3.
Advertisement_Count	The number of ADVERTISEMENTS or FAST ADVERTISEMENTS that can be missed before a BACKUP will declare a MASTER down.
Master_Down_Interval	Time interval for Backup to declare Master down (milliseconds). Calculated as: $\left(\text{Advertisement_Count} * \text{Advertisement_Interval} \right) + \text{Skew_time}$

3.2. Updates to the Timers

With the introduction of the optional FAST ADVERTISEMENT message, time values MUST reflect the granularity supported in the FAST ADVERTISEMENT message. Thus all timers MUST reflect the granularity for FAST ADVERTISEMENT messages (milliseconds).

Master_Down_Timer	Timer that fires when ADVERTISEMENT or FAST ADVERTISEMENT has not been heard for Master_Down_Interval.
Adver_Timer	Timer that fires to trigger sending of ADVERTISEMENT based on Advertisement_Interval.
Adver_2_Timer	Timer that fires to trigger sending of FAST ADVERTISEMENT based on Advertisement_Interval.

3.3. Updates to the State Descriptions

The State Transition Diagram does not change with the addition of the FAST ADVERTISEMENT message type. The following updates the descriptions for the three states.

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.

3.3.1. Initialize

The purpose of this state is to wait for a Startup event. If a Startup event is received, then:

- o Initialize local Virtual Router settings (VR_Mode, VR_Type, VR_AIG, VR_Adver_Interval, and VR_Adver_Count)
- If the Priority = 255, then: (i.e., the router owns the IP address(es) associated with the virtual router)

If the VR_Type is 1, then:

- o Send an ADVERTISEMENT
- o Set the Adver_Timer to Advertisement_Interval

else:

- o Send a FAST ADVERTISEMENT
- o Set the Adver_2_Timer to Advertisement_Interval

endif

- o Broadcast a gratuitous ARP request containing the virtual router MAC address for each IP address associated with the virtual router.

- o Transition to the {Master} state

else:

- o Set the Master_Down_Timer to Master_Down_Interval
- o Transition to the {Backup} state

endif

3.3.2. 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} state

endif

- If the Master_Down_Timer fires, then:

If the VR_Mode is 0, then:

If the VR_Type is 1, then:


```
    o Send an ADVERTISEMENT
    o Set the Adver_Timer to Advertisement_Interval

else:

    o Send a FAST ADVERTISEMENT
    o Set the Adver_2_Timer to Advertisement_Interval

endif

else:

    o Send an ADVERTISEMENT
    o Send a FAST ADVERTISEMENT
    o Set the Adver_Timer to Maximum of (1 second,
      Advertisement_Interval)
    o Set the Adver_2_Timer to Advertisement_Interval

endif

    o Broadcast a gratuitous ARP request containing the virtual
      router MAC address for each IP address associated with the
      virtual router
    o Transition to the {Master} state

endif

- If an ADVERTISEMENT is received, then:

  If VR_Type is 2, then:

    o Set the VR_Mode to 1 (mixed message type VRRP
      environment)

  endif

  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

- If a FAST ADVERTISEMENT is received, then:

    If VR_Type is 1, then:

        o Set the VR_Mode to 1 (mixed message type VRRP
          environment)

    endif

    If the Priority in the FAST 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
        FAST 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 FAST ADVERTISEMENT

        endif
    endif
endif
```

3.3.3. 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 Cancel the Adver_2_Timer

If the VR_Mode is 0, then

If the VR_Type is 1, then:

- o Send an ADVERTISEMENT with Priority = 0

else

- o Send a FAST ADVERTISEMENT with Priority = 0

endif

else

- o Send an ADVERTISEMENT with Priority = 0
- o Send a FAST ADVERTISEMENT with Priority = 0

endif

- If the Adver_Timer fires, then:

If the VR_Mode is 1

or

if the VR_Type is 1, then:

- o Send an ADVERTISEMENT
- o Reset the Adver_Timer to Maximum of (1 second, VR_Adver_Interval)

endif

endif

- If the Adver_2_Timer fires, then:

If the VR_Mode is 1

or

if the VR_Type is 2, then:


```
    o Send a FAST ADVERTISEMENT
    o Reset the Adver_2_Timer to VR_Adver_Interval

endif
endif

- If an ADVERTISEMENT is received, then:

  If the VR_Mode is 0
  and
  If the VR_Type is 2, then

    o Set VR_Mode to 1 (heterogeneous VRRP environment)
    o Set the Adver_Timer to Maximum of 1 second and
      VR_Adver_Interval

  endif

  If the Priority in the ADVERTISEMENT is Zero, then:

    If the VR_Mode is 0, then

      If the VR_Type is 1, then:

        o Send an ADVERTISEMENT
        o Reset the Adver_Timer to Maximum of 1 second and
          VR_Adver_Interval

      else

        o Send a FAST ADVERTISEMENT
        o Reset the Adver_2_Timer to VR_Adver_Interval

      endif

    else

      o Send an ADVERTISEMENT
      o Send a FAST ADVERTISEMENT
      o Reset the Adver_Timer to Maximum of 1 second and
        VR_Adver_Interval
      o Reset the Adver_2_Timer to VR_Adver_Interval

    endif

  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  Cancel Adver_2_Timer
        o  Set Master_Down_Timer to Master_Down_Interval
        o  Transition to the {Backup} state

    else:

        o  Discard ADVERTISEMENT

    endif
endif
endif

- If a FAST ADVERTISEMENT is received, then:

    If the VR_Mode is 0
    and
    If the VR_Type is 1, then

        o  Set VR_Mode to 1 (heterogeneous VRRP environment)
        o  Set the Adver_2_Timer to VR_Adver_Interval

    endif

    If the Priority in the FAST ADVERTISEMENT is Zero, then:

        If the VR_Mode is 0, then

            o  Send a FAST ADVERTISEMENT
            o  Reset the Adver_2_Timer to VR_Adver_Interval

        else

            o  Send an ADVERTISEMENT
            o  Send a FAST ADVERTISEMENT
            o  Reset the Adver_Timer to Maximum of 1 second and
              VR_Adver_Interval
            o  Reset the Adver_2_Timer to VR_Adver_Interval

        endif

    else:
```



```
If the Priority in the FAST ADVERTISEMENT is greater
than the local Priority,
or
If the Priority in the FAST 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 Cancel Adver_2_Timer
    o Set Master_Down_Timer to Master_Down_Interval
    o Transition to the {Backup} state

else:

    o Discard FAST ADVERTISEMENT

endif
endif
endif
```

4. Updates for Sending and Receiving VRRPv4 Packets

This section outlines the changes to [Section 7](#) (Sending and Receiving VRRP Packets) of the VRRPv4 [[RFC 3768](#)] to accommodate the optional FAST ADVERTISEMENT, Type 2, message.

4.1. Receiving VRRPv4 Packets

Perform 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 and variable fields) for either Type 1 or Type 2 messages.
- 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)).
- For Type 1 (ADVERTISEMENT) messages, 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 the message Type matches the locally

configured VRRP Advertisement Type for the virtual router

[draft-ietf-vrrp-ipv4-timers-02.txt](#)

[Page 15]

(either Type 1 for ADVERTISEMENT or Type 2 for FAST ADVERTISEMENT).

If the above check fails, the receiver SHOULD log the event and MAY indicate via network management that a misconfiguration was detected.

- 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.

- For Type 2 (FAST ADVERTISEMENT) messages, MUST verify that the Advertisement Count is the same as locally configured for this virtual router.
- For Type 2 (FAST ADVERTISEMENT) messages, MUST verify that the Advertisement Interval Granularity is the same as locally configured for this virtual router.
- MUST verify that the Advertisement Interval in the packet is the same as locally configured for this virtual router.

If any of the above checks fail, the receiver SHOULD log the event and MAY indicate via network management that a misconfiguration was detected.

4.2. Transmitting VRRPv4 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 (based on the message Type)
- 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.

5. Operational Issues

5.1 Sub-second Timers

The changes proposed to VRRP for IPv4 are intended to provide additional capabilities to address specific operational requirements, specifically, sub-second fail over from the Master. The use of sub-second timers is not recommended for general purpose networking environments. Missed ADVERTISEMENTS can lead to fail overs. Reducing the time period for fail over, the MASTER_DOWN_TIMER, increases the potential for missed ADVERTISEMENTS, due to router processing requirements, network congestion, or even denial of service attacks.

The new message type provides extensions to VRRPv4 allowing the specification of sub-second timers. It also provides the ability to specify the number of advertisement messages that can be missed by a BACKUP before declaring a MASTER down.

5.2. Interoperability / Backward Compatibility

The addition of the new message type introduces the potential for routers that do not support the new message type configured on the same network with routers that use the new message type. The state machines have been updated to interoperate with routers only supporting Type 1 Advertisements. When routers configured to send Type 2 Fast Advertisements discover routers sending Type 1 Advertisements, it sends both types of advertisements. In the Type 1 Advertisements, the Advertisement Interval is set to the larger of the interval value from the MASTER or one second (the minimum setting for Type 1 Advertisements). Type 2 messages will not support authentication.

6. Security Considerations

This draft does not address the security issues relating to VRRP that have been identified in [RFC 3768](#).

7. 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.

8. Acknowledgments

The work presented in this document is based upon the VRRP specification in [RFC3768](#) and the current work in progress for VRRP for IPv6. The authors and contributors of these works are R. Hinden, P. Higginson, R. Hinden, P. Hunt, S. Knight, A. Lindem, D. Mitzel, M. Shand, D. Weaver, and D. Whipple.

The author of this document would also like to thank Karen O'Donoghue, Leon Sangroniz, Changming Liu, Sonum Mathur, Mukesh Gupta, Steve Bates, and Bob Hinden for their guidance and helpful suggestions.

9. IANA Considerations

This document has no actions for IANA.

10. Normative References

[RFC3768] Hinden, R., Ed., "Virtual Router Redundancy Protocol (VRRP)", [RFC 3768](#), April 2004.

11. Informative References

[RFC2338] Knight, S., et. al., "Virtual Router Redundancy Protocol", [RFC2338](#), April 1998.

[VRRP-IPv6]
Hinden, R., "Virtual Router Redundancy Protocol for IPv6", [draft-ietf-vrrp-ipv6-spec-07](#) (work in progress), September 2004.

12. Author's Address

Robert W. Hott
Naval Surface Warfare Center Dahlgren Division
Code B35
17320 Dahlgren Road
Dahlgren, VA 22448-5100
USA

Phone: +1 540 653-1497
EMail: robert.hott@navy.mil

13. Disclaimer of Validity

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.

14. Copyright Statement

Copyright (C) The Internet Society (2006). 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.

