

intarea  
Internet-Draft  
Intended status: Standards Track  
Expires: April 4, 2019

R. Patterson  
Sky UK  
M. Abrahamsson  
T-Systems  
October 01, 2018

**IP over Ethernet (IPoE) Session Health Checking**  
**draft-patterson-intarea-ipoe-health-05**

Abstract

PPP over Ethernet clients have the functionality to detect path unavailability by using PPP Keepalives. IP over Ethernet does not have this functionality, and it is not specified in the IETF when an IP over Ethernet client should consider its WAN connectivity down, unless there is a physical layer link down event.

This document describes a mechanism for IP over Ethernet clients to achieve connectivity validation, similar to that of PPP over Ethernet, by using BFD Echo, or an alternative health check mechanism.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on April 4, 2019.

Copyright Notice

Copyright (c) 2018 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of

publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

<a href="#">1.</a>	<a href="#">Introduction</a>	<a href="#">2</a>
<a href="#">1.1.</a>	<a href="#">Requirements Language</a>	<a href="#">3</a>
<a href="#">1.2.</a>	<a href="#">Terminology</a>	<a href="#">3</a>
<a href="#">2.</a>	<a href="#">Alternative Mitigations</a>	<a href="#">4</a>
<a href="#">3.</a>	<a href="#">IPoE Health Checks</a>	<a href="#">4</a>
<a href="#">3.1.</a>	<a href="#">Parameters</a>	<a href="#">4</a>
<a href="#">3.2.</a>	<a href="#">Startup</a>	<a href="#">5</a>
<a href="#">3.3.</a>	<a href="#">BFD Echo</a>	<a href="#">5</a>
<a href="#">3.4.</a>	<a href="#">IPoE Health Check Probe</a>	<a href="#">5</a>
<a href="#">4.</a>	<a href="#">IPoE Health Check DHCP Options</a>	<a href="#">6</a>
<a href="#">4.1.</a>	<a href="#">DHCPv6</a>	<a href="#">6</a>
<a href="#">4.2.</a>	<a href="#">DHCPv4</a>	<a href="#">7</a>
<a href="#">5.</a>	<a href="#">Recovery Behaviour</a>	<a href="#">7</a>
<a href="#">5.1.</a>	<a href="#">LAN Considerations</a>	<a href="#">9</a>
<a href="#">6.</a>	<a href="#">Multihomed Clients</a>	<a href="#">9</a>
<a href="#">7.</a>	<a href="#">Security Considerations</a>	<a href="#">9</a>
<a href="#">8.</a>	<a href="#">IANA Considerations</a>	<a href="#">9</a>
<a href="#">9.</a>	<a href="#">Acknowledgements</a>	<a href="#">10</a>
<a href="#">10.</a>	<a href="#">Appendix A. Changes from -00</a>	<a href="#">10</a>
<a href="#">11.</a>	<a href="#">Appendix B. Changes from -01</a>	<a href="#">10</a>
<a href="#">12.</a>	<a href="#">Appendix C. Changes from -02</a>	<a href="#">10</a>
<a href="#">13.</a>	<a href="#">Appendix D. Changes from -03</a>	<a href="#">11</a>
<a href="#">14.</a>	<a href="#">Appendix E. Changes from -04</a>	<a href="#">11</a>
<a href="#">15.</a>	<a href="#">References</a>	<a href="#">11</a>
<a href="#">15.1.</a>	<a href="#">Normative References</a>	<a href="#">11</a>
<a href="#">15.2.</a>	<a href="#">Informative References</a>	<a href="#">12</a>
<a href="#">15.3.</a>	<a href="#">URIs</a>	<a href="#">13</a>
	<a href="#">Authors' Addresses</a>	<a href="#">13</a>

## [1. Introduction](#)

PPP [[RFC1661](#)] makes use of regular LCP echos and replies to continually test the data link layer, if the peer fails to respond to a predetermined number of LCP echos, the PPP session is terminated and will return to the Link Dead phase, ready for reestablishing. IPoE currently lacks this functionality.



Physical link state change on an IPoE client can trigger the renewing of a DHCP lease, however any indirect upstream network changes are not visible to the IPoE client.

An outage or planned maintenance work on, for example, a Broadband Network Gateways (BNG) or intermediate DHCP Relay, can leave an IPoE client with a stale DHCP lease for up to the Valid Lifetime.

IPoE Session Health Check allows for an IPoE client to proactively or passively monitor the state of upstream connectivity, and defines several actions that may be taken to help the client recover.

[[TR-146](#)], Section 6.2 describes this problem, while [[TR-124](#)] identifies some requirements to solve the problem.

Several vendors have implemented subscriber connectivity checking on their BNG, using ARP and Neighbor Discovery as per Sections [6.2.4](#) and [6.2.5](#) of [[TR-146](#)]. This allows the BNG to detect loss of connectivity and to update local session state and DHCP lease bindings. Without reciprocal checking, this puts the CE at further risk of being in a stale state.

### **[1.1.](#) Requirements Language**

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [BCP 14](#) [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

### **[1.2.](#) Terminology**

This document makes use of the following terms:

- o BNG: Broadband Network Gateway. Often also running a DHCP server or relay.
- o CE: Customer Equipment. aka. Customer Premise Equipment (CPE), Residential Gateway (RG).
- o IPoE: IP over Ethernet.
- o IPoE Client: A network device, often a CE, running a DHCPv4 and/or DHCPv6 client.
- o IPoE Health Check: The name of the process described in this document.



## **2. Alternative Mitigations**

- o Short DHCP lease times reduce the time a client may be left in a stale state, but scale poorly, putting extra load on the DHCP server.
- o Broadband Forum's [\[TR-146\]](#) and [\[TR-124\]](#) discuss this problem and recommend the use of BFD echo [\[RFC5880\]](#). This document acknowledges TR-146 and recommends the use of BFD echo for health checks, but like the Broadband Forum, acknowledges that it is not widely available within consumer CEs.  
The renew action specified in TR-146 is often insufficient for a network to authenticate an IPoE Client, leaving the client in a stale state for up to the lease expiry, and no better off than without the check.  
This document introduces an alternative action to help expedite client recovery.
- o For planned maintenance, network engineers could include DHCPv4 Force Renew [\[RFC3203\]](#) or DHCPv6 Reconfigure [\[RFC3315\]](#)-bis in their maintenance plans, however neither of these have been widely adopted by CE or BNG vendors due to authentication complexity.

## **3. IPoE Health Checks**

### **3.1. Parameters**

IPoE Health Check uses the following parameters:

- o Interval (Integer): The frequency in seconds, which health checks are sent by the IPoE client. Default value: 120 seconds.
- o Retry Interval (Integer): The frequency in seconds, which health checks are sent by the IPoE client, after a failure. Default value: 10 seconds.
- o Limit (Integer): The number of failed consecutive checks before an action is taken. Default value: 3.
- o Release (Boolean): Instructs the client to send a DHCP RELEASE and to not attempt a renewal of the current lease. Default value: 0.

An IPoE client MAY be statically configured for IPoE health checks. Non-default static parameters SHOULD override any signalled via a dynamic means (e.g, DHCP or TR-69).

An IPoE client MAY use default parameters in lieu of manually configured, or dynamically signalled parameters (DHCP for example).



Statically configured or dynamically signalled parameters SHOULD override any default parameters.

### **3.2. Startup**

An IPoE client supporting IPoE Health Check MUST begin sending health checks at the Retry Interval ([Section 3.1](#)) specified, upon successful binding of a lease that contains a valid IPoE Health Check DHCP Option ([Section 4](#)), or for which it has been statically configured.

After Limit IPoE health checks succeeds consecutively, the IPoE client MUST begin sending health checks at the regular Interval.

If Limit IPoE health checks fail consecutively, IPoE Health Check should be considered unusable and the IPoE client MUST cease the sending of health checks, ignoring the recovery behaviour ([Section 5](#)).

### **3.3. BFD Echo**

If the IPoE Client supports BFD [[RFC5880](#)], it SHOULD use BFD Echo as the health check mechanism.

If the IPoE Client does not support BFD, or if it is unable to establish a BFD session with the upstream router, it MUST use IPoE Health Check Probe [Section 3.4](#) as the health check mechanism.

### **3.4. IPoE Health Check Probe**

Similar in format to a BFD Echo packet, the IPoE Health Check probe is a simple UDP/IP packet with a destination IP address from the lease being checked, and a destination MAC address of the upstream router. However, unlike BFD Echo, an IPoE Health Check probe does not require a BFD session to be established between peers; allowing for less state and configuration on a BNG and a simpler CPE implementation.

The destination IP MUST be an address locally bound on the IPoE client and MUST be from the lease triggering the IPoE Health Check.

The source IP SHOULD be the same as the destination address, but MAY be another address bound to the same interface the health check probe is being sent out. E.g. The link-local or a DHCPv6 NA assigned address.

The source MAC address MUST belong to the IPoE Client interface of the lease being checked.

The destination MAC address MUST be address of the upstream router.





Using an IP packet as the health check probe not only validates the layer 2 forwarding path, but also validates the DHCP lease state.

#### 4. IPoE Health Check DHCP Options

This document defines a new option for both DHCPv4 and DHCPv6 servers to signal suggested health check parameters to clients. IPoE clients SHOULD use these values when no statically configured parameters have been defined.

The option data fields are common between DHCPv6 and DHCPv4.

##### 4.1. DHCPv6

For DHCPv6, this option (Figure 1) MUST be within a specific Identity Association as an IPoE client MAY have multiple IAs with different health check parameters.




Figure 1: DHCPv6 IPoE Check Option Format

The description of the fields is as follows:

option-code: OPTION\_IPOE\_HEALTH (TBA1).

option-len: 12.

limit: Consecutive failed checks, before an action is taken.

R: Release flag. Instructs the client to send a DHCP RELEASE when a failure is detected, and to skip the renew step.

interval: Indicates how often a health check should be sent when no failure is encountered. Expressed in units of seconds.

retry-interval: Indicates how often a health check should be sent after a previous failure. Expressed in units of seconds.



## 4.2. DHCPv4

The DHCPv4 client can retrieve IPoE Health Check information by including `OPTION_IPOE_HEALTH` in a Parameter Request List option [[RFC2132](#)].

Figure 2 shows the DHCPv4 IPoE Health Check option.




Figure 2: DHCPv4 IPoE Health Check Option Format

The description of the fields is as follows:

option-code: `OPTION_IPOE_HEALTH` (TBA2).

option-len: 10.

limit: Consecutive failed checks, before an action is taken.

R: Release flag. Instructs the client to send a DHCP RELEASE when a failure is detected, and to skip the renew step.

interval: Indicates how often a health check should be sent when no failure is encountered. Expressed in units of seconds.

retry-interval: Indicates how often a health check should be sent after a previous failure. Expressed in units of seconds.

## 5. Recovery Behaviour

IPoE Health Check defines the behaviour that an IPoE Client should take once the timeout threshold has been reached. This behaviour makes use of existing procedures outlined in [[RFC2131](#)], [Section 4.4.5](#) for DHCPv4, and [[RFC3315](#)]-bis, [Sections 18.2.4](#), [18.2.5](#) for DHCPv6.

When triggering a DISCOVER or SOLICIT, an IPoE client may also choose to use Rapid Commit [[RFC4039](#)], [[RFC3315](#)], [Section 22.14](#) to help expedite the recovery process.

After Limit ([Section 3.1](#)) consecutive check failures, T1 and T2 MUST both be set to zero.



If the Release Flag ([Section 3.1](#)) is unset, the IPoE client SHOULD immediately attempt to renew the current lease from the original server. If connectivity to the original DHCP server has recovered, and the server can satisfy the request, the lease may be renewed and timers updated.

If the renew is unsuccessful, the IPoE client MUST move to the discovery or solicit phase.

If the Release Flag is unset, the IPoE client MUST keep the address or prefix in the preferred state until the preferred lifetime expires, and MUST keep the address or prefix until the valid lifetime expires.

If the Release Flag is set the IPoE Client MUST NOT send a DHCP RENEW, it MUST send a RELEASE as per [\[RFC2131\]](#), [Section 3.1](#) for DHCPv4 and [\[RFC3315\]](#)-bis, [Section 18.2.7](#).

If the IPoE client is already in the renew or rebind state when this behaviour is triggered, the client MUST cease renew or rebind attempts and wait for any outstanding messages to time out before sending a RELEASE.

If an outstanding renew or rebind attempt is successful, the IPoE client MUST update T1, T2 and lease lifetimes appropriately, and MUST NOT continue with this behaviour.

Once the RELEASE process has completed, the IPoE Client should move to the discovery or solicit phase.

The IPoE client MUST include the current address or prefix in the IA Address or IA\_PD options within the DHCPv6 SOLICIT, or in the Requested IP Address Option of a DHCPv4 DISCOVER [\[RFC2131\]](#), [Section 4.4.2](#).

The DHCPv6 SOLICIT DUID and IAID values MUST be the same as used in the current lease.

If the DHCP server assigns the current address or prefix, the IPoE client MUST update the current lease timers, and any differing parameters.

If the DHCP server assigns an alternate address or prefix, the IPoE client MUST deprecate the current lease and follow the actions outlined in requirement L-13 [\[RFC7084\]](#), [Section 4.3](#)



### 5.1. LAN Considerations

If all DHCPv6 leases have expired, either naturally or proactively with IPoE health checks, an IPoE client acting as a router, SHOULD withdraw itself as a default router on the LAN, following requirement G-5 of [\[RFC7084\]](#), [Section 4.1](#).

## 6. Multihomed Clients

An IPoE client may have multiple leases from the same, or different DHCP servers. These leases may have different IPoE health check parameters, and health checks MUST be treated distinctly, tracking the particular lease that they belong to.

Local network administrators may choose to override DHCP-signalled parameters in order to facilitate appropriate IPoE Health Check operation in a multihomed environment.

## 7. Security Considerations

IPoE Health Check frequency would typically be controlled by the network using DHCP options, but overly aggressive, statically configured IPoE Health Checks, could have an adverse impact. For example, this may induce an overload on the IP access nodes. However, BFD Echo and the IPoE Health Check probe defined in this document, both use an IP packet destined for the IPoE client, the remote peer forwards the packet back to the IPoE client without any local processing.

The inclusion of the current lease address or prefix when sending a DISCOVER or SOLICIT, introduces a privacy risk, possibly leaking lease information if the IPoE client has been moved to a different network, e.g., from one fixed line provider to another.

## 8. IANA Considerations

IANA is requested to assign a new DHCPv6 Option Code in the registry maintained in <http://www.iana.org/assignments/dhcpv6-parameters> [1]:

Option Name	Value
-----	-----
OPTION_IPOE_HEALTH	TBA1

Also, IANA is requested to assign the following new DHCPv4 Option Code in the registry maintained in <http://www.iana.org/assignments/bootp-dhcp-parameters> [2]:





Option Name	Tag	Data Length	Meaning
-----			
OPTION_IPOE_HEALTH	TBA2	14	Provides a set of IPoE check configuration information.

## **9. Acknowledgements**

The authors would like to thank Ian Farrer, Dusan Mudric, Mohamed Boucadair, Jean-Yves Cloarec, Bernie Volz, Barbara Stark, Dave Freedman, and Job Snijders for their review and comments on this and previous versions of this document.

## **10. [Appendix A](#). Changes from -00**

This section should be removed by the RFC Editor.

- o Added reference to TR-146.
- o Added BFD Echo section, and wording to prefer it as the health check mechanism over ARP/ND, if available.

## **11. [Appendix B](#). Changes from -01**

This section should be removed by the RFC Editor.

- o Emphasised preference for use of BFD echo as the health check mechanism.
- o Removed lifetime expiration from Behaviour 2 and clarified usage.
- o Updated Behaviour 3 with instructions for whilst mid-renew/rebind.
- o Reworded multihoming section.
- o Added Acknowledgements.

## **12. [Appendix C](#). Changes from -02**

This section should be removed by the RFC Editor.

- o Added DHCP option flag to force ARP/ND for health checks.
- o Populated IANA Considerations.
- o Added Retry Interval distinct timer for between failed checks.
- o Added default parameter values.



### **13.** **Appendix D.** Changes from -03

This section should be removed by the RFC Editor.

- o Reduced default Limit value.
- o Formatting and minor cosmetic changes.

### **14.** **Appendix E.** Changes from -04

This section should be removed by the RFC Editor.

This revision is a major rewrite. Changes include:

- o Consolidated the multiple behaviours down to one.
- o Added a Release flag instead of a distinct behaviour.
- o Added custom IPoE Health Check Probe definition.
- o Removed ARP/ND and Passive checks.
- o Removed alternative target address parameter.
- o Removed IANA registry request for Behaviour values.

## **15.** **References**

### **15.1.** **Normative References**

- [RFC0826] Plummer, D., "An Ethernet Address Resolution Protocol: Or Converting Network Protocol Addresses to 48.bit Ethernet Address for Transmission on Ethernet Hardware", STD 37, [RFC 826](https://www.rfc-editor.org/info/rfc826), DOI 10.17487/RFC0826, November 1982, <<https://www.rfc-editor.org/info/rfc826>>.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", [RFC 2131](https://www.rfc-editor.org/info/rfc2131), DOI 10.17487/RFC2131, March 1997, <<https://www.rfc-editor.org/info/rfc2131>>.
- [RFC2132] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", [RFC 2132](https://www.rfc-editor.org/info/rfc2132), DOI 10.17487/RFC2132, March 1997, <<https://www.rfc-editor.org/info/rfc2132>>.
- [RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 3315](https://www.rfc-editor.org/info/rfc3315), DOI 10.17487/RFC3315, July 2003, <<https://www.rfc-editor.org/info/rfc3315>>.



- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", [RFC 4861](#), DOI 10.17487/RFC4861, September 2007, <<https://www.rfc-editor.org/info/rfc4861>>.
- [RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", [RFC 5880](#), DOI 10.17487/RFC5880, June 2010, <<https://www.rfc-editor.org/info/rfc5880>>.

## **15.2. Informative References**

- [RFC1661] Simpson, W., Ed., "The Point-to-Point Protocol (PPP)", STD 51, [RFC 1661](#), DOI 10.17487/RFC1661, July 1994, <<https://www.rfc-editor.org/info/rfc1661>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC3203] T'Joens, Y., Hublet, C., and P. De Schrijver, "DHCP reconfigure extension", [RFC 3203](#), DOI 10.17487/RFC3203, December 2001, <<https://www.rfc-editor.org/info/rfc3203>>.
- [RFC4039] Park, S., Kim, P., and B. Volz, "Rapid Commit Option for the Dynamic Host Configuration Protocol version 4 (DHCPv4)", [RFC 4039](#), DOI 10.17487/RFC4039, March 2005, <<https://www.rfc-editor.org/info/rfc4039>>.
- [RFC6192] Dugal, D., Pignataro, C., and R. Dunn, "Protecting the Router Control Plane", [RFC 6192](#), DOI 10.17487/RFC6192, March 2011, <<https://www.rfc-editor.org/info/rfc6192>>.
- [RFC7084] Singh, H., Beebe, W., Donley, C., and B. Stark, "Basic Requirements for IPv6 Customer Edge Routers", [RFC 7084](#), DOI 10.17487/RFC7084, November 2013, <<https://www.rfc-editor.org/info/rfc7084>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [TR-124] "Functional Requirements for Broadband Residential Gateway Devices", 2006, <<https://www.broadband-forum.org/technical/download/TR-124.pdf>>.
- [TR-146] "Subscriber Sessions", 2013, <<https://www.broadband-forum.org/technical/download/TR-146.pdf>>.



### **15.3. URIs**

- [1] <http://www.iana.org/assignments/dhcpv6-parameters>
- [2] <http://www.iana.org/assignments/bootp-dhcp-parameters>

#### Authors' Addresses

Richard Patterson  
Sky UK

Email: richard.patterson@sky.uk

Mikael Abrahamsson  
T-Systems

Email: mikael.abrahamsson@t-systems.se



