Internet Area Working Group
Internet-Draft
Intended status: Standards Track
Expires: 19 February 2026

J. Rajamanickam, Ed.
D. Dukes
M. Sankaranarayanan
Cisco Systems, Inc.

18 August 2025

ICMP extension to include underlay information
draft-jags-intarea-icmp-ext-underlay-info-02

Abstract

Network operators operating overlay networks require the ability to identify hops in an underlay network when traceroute in the overlay. This document defines an ICMP Error extension message to carry the underlay error information to the overlay network endpoint.

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 19 February 2026.

Table of Contents

1. Introduction

The mechanism allowing ICMP messages to carry additional information is [RFC4884]. ICMP message extensions, describing the mechanisms for extending ICMP messages to carry additional information about the system where the error occurred are defined in [RFC5837], [RFC8335] and [RFC8883]. These messages are transmitted to the source node to provide deeper insight into the error in relation to the node where it occurred.

Network operators who administer and overlay and underlay network, such as those with VPN segmentation within their network and an SRv6 core connecting them, find it particularly useful to have ICMP underlay information transmitted to source nodes for the purpose of using traceroute. This ICMP underlay information provides details about errors and failures in the underlay network.

The underlay error information described in this document satisfy the need of these network operators.

2. Conventions and Definitions

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.

3. Underlay Information Object

This section defines a new ICMP extension object called Underlay Information Object (UIO) that is encoded as part of ICMP extension message. A new Class-Num value TBA (To Be Assigned) is assigned to identify the UIO. As per [RFC4884], this object MAY be appended to one of the following ICMP messages:

- ICMPv4 Time Exceeded
- ICMPv4 Destination Unreachable
- ICMPv4 Parameter Problem
- ICMPv6 Time Exceeded
- ICMPv6 Destination Unreachable

3.1. UIO Object Format

This section described the UIO object format.

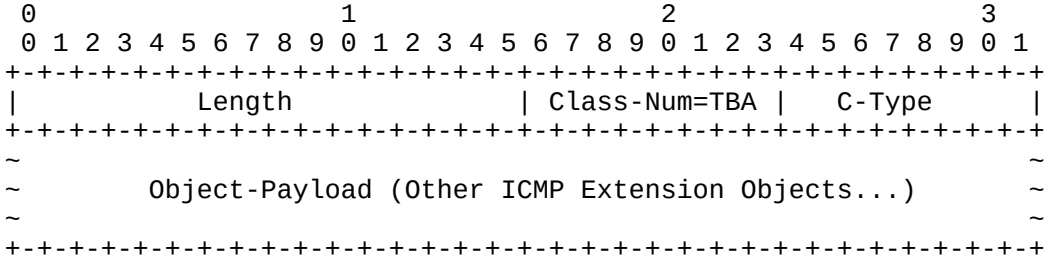

Figure 1: Underlay Information Object Format

UIO ICMP extension object contains the following fields:

Length: 16 bits

Length of the object, measured in octets, including UIO ICMP extension object header and object payload

Class-Num: 8 bits

Identifies object class. In the case of UIO, a new value TBA will be assigned from IANA "ICMP Extension Object Classes and Class Sub-types" registry. This object could co-exist with the other ICMP extension objects

C-Type: 8 bits

Identifies object sub-type. This value MUST be transmitted as zeros and ignored upon receipt.

Object-Payload: Variable Length

UIO Object payload contains one or more other ICMP Extension Objects that are related to the underlay nodes.

This ICMP extension object acts as an envelope to carry other ICMP extension objects related to the underlay. Primarily, the UIO ICMP extension object is encoded in the ICMP extension message by the underlay head-end when it receives an ICMP error message from one of its intermediate nodes.

This UIO ICMP extension object can encapsulate one or more relevant ICMP extension objects that are related to the underlay node. When the underlay head-end encodes its ICMP extension object, the first object MUST contain the ICMP extension object that carries IP address or the hostname of the node where the initial ICMP error was generated. The ICMP extension objects encoded within the UIO ICMP extension objects can belong to any address family, irrespective of the address family of the source node that decapsulates the UIO ICMP extension objects, as opposed to what is stated in [RFC5837] Section 4.2.

When the UIO is encoded, the total length of the ICMP message, including extensions, MUST NOT exceed the minimum reassembly buffer size.

If the node decoding the ICMP extension header does not recognize the UIO ICMP extension message, it SHOULD ignore this message and continue processing the other messages.

3.2. Underlay Information Object Encoding Process

An underlay head-end node generates an ICMP error directed to a host in the overlay network. The specifics of the ICMP error content are beyond the scope of this document. The error is triggered by the receipt of an underlay error message from an IPv4/IPv6 interface. The encoding and processing of the underlay error are also outside the scope of this document.

The underlay head-end appends a UIO ICMP extension object to the ICMP error it generates to the overlay host.

4. Security Considerations

In most of the cases the overlay and underlay networks are owned by different operators, so it may not be advisable to permit the transmission of the underlay information to an arbitrary recipient. The inclusion of this information should be configurable and must default to being disabled. An implementation should decide which objects can be appended as part of the UIO ICMP extension message.

The extensions defined in this document are intended for use in administrative debugging and troubleshooting. They provide additional information in ICMP responses. These mechanisms are not designed for use in non-debugging applications.

5. IANA Considerations

This document requests that IANA allocate a value (TBA - To Be Assigned) for the UIO ICMP extension object class value from the "ICMP Extension Object Classes and Class Sub-types" registry to indicate the presence of the UIO ICMP extension class, referred to as TBA above.

6. Appendix

6.1. UIO ICMP Extension Message Examples

This section lists the UIO example encoding format

6.1.1. UIO carrying IPv6 information to the IPv4 source

In this example, a host receives an IPv4 ICMPv4 Time Exceeded error message in response to an ICMP Echo Request as part of the traceroute application. It also contains an UIO ICMP extension object with IPv6 interface address information as follows.

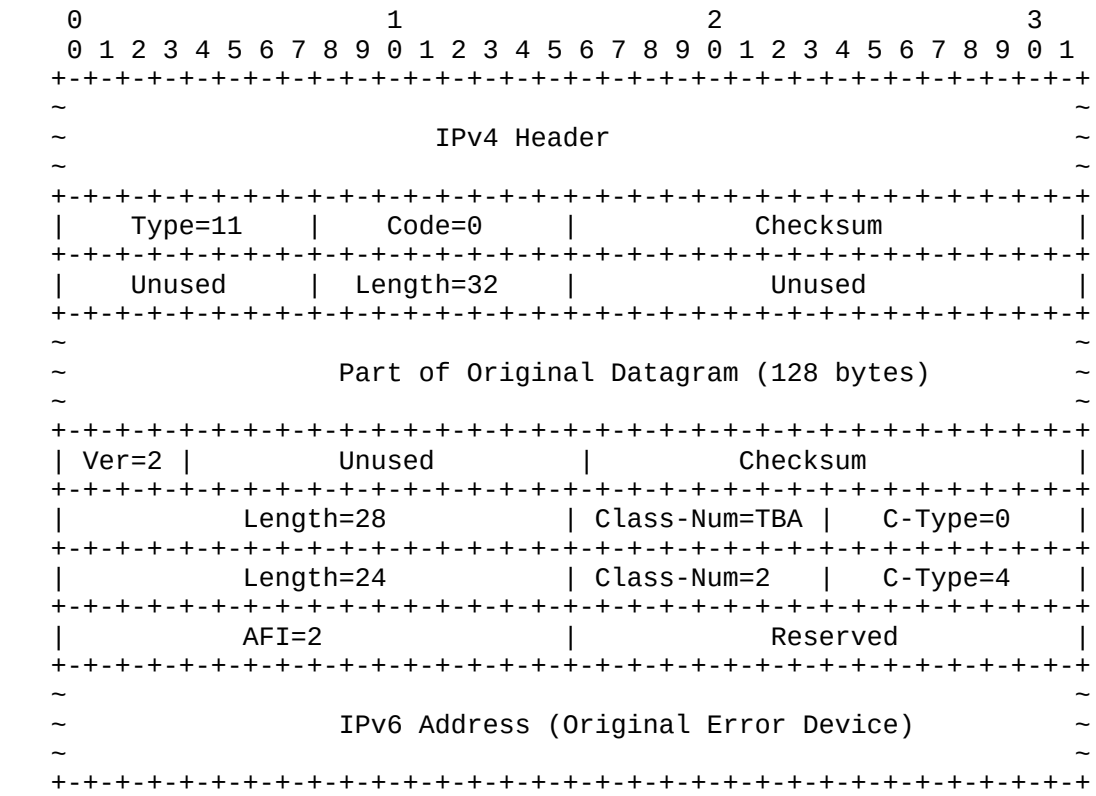

Figure 2: ICMPv4 packet carrying UIO ICMP extension

The traceroute application displays the IPv6 Address in the UIO to allow an administrator to trace the underlay path of the route being traced.

6.1.2. UIO carrying IPv4 information to the IPv6 source

In this example, a host receives an IPv6 ICMPv6 Time Exceeded error message in response to an ICMP Echo Request as part of the traceroute application. It contains a UIO ICMP extension object with IPv4 interface address information as follows.

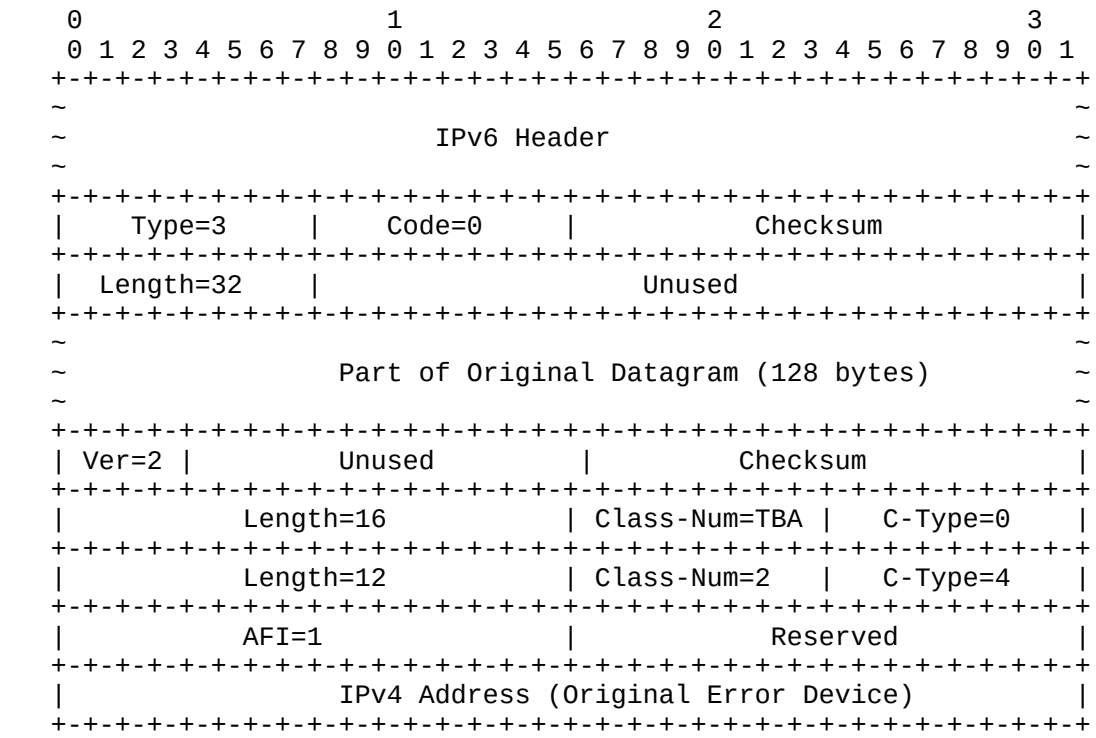

Figure 3: UIO carrying IPv4 information to the IPv6 source

The traceroute application displays the IPv4 Address in the UIO to allow an administrator to trace the underlay path of the route being traced.

7. Normative References

[IANA.address-family-numbers]
IANA, "Address Family Numbers",
<http://www.iana.org/assignments/address-family-numbers>.

[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/rfc/rfc2119>.

[RFC4884] Bonica, R., Gan, D., Tappan, D., and C. Pignataro, "Extended ICMP to Support Multi-Part Messages", RFC 4884, DOI 10.17487/RFC4884, April 2007, <https://www.rfc-editor.org/rfc/rfc4884>.

[RFC5837] Atlas, A., Ed., Bonica, R., Ed., Pignataro, C., Ed., Shen, N., and JR. Rivers, "Extending ICMP for Interface and Next-Hop Identification", RFC 5837, DOI 10.17487/RFC5837, April 2010, <https://www.rfc-editor.org/rfc/rfc5837>.

[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/rfc/rfc8174>.

[RFC8335] Bonica, R., Thomas, R., Linkova, J., Lenart, C., and M. Boucadair, "PROBE: A Utility for Probing Interfaces", RFC 8335, DOI 10.17487/RFC8335, February 2018, <https://www.rfc-editor.org/rfc/rfc8335>.

[RFC8883] Herbert, T., "ICMPv6 Errors for Discarding Packets Due to Processing Limits", RFC 8883, DOI 10.17487/RFC8883, September 2020, <https://www.rfc-editor.org/rfc/rfc8883>.

Acknowledgments

This document derives text heavily from [RFC4884] and [RFC5837].

Contributors

The following people have substantially contributed to this document:

Tamilselvan Murugan
Cisco Systems, Inc.
Email: tammurug@cisco.com

Dhilip Sekar
Cisco Systems, Inc.
Email: dhsekar@cisco.com

Figure 4

Authors' Addresses

Jaganbabu Rajamanickam (editor)
Cisco Systems, Inc.
Canada
Email: jrajaman@cisco.com

Darren Dukes
Cisco Systems, Inc.
Canada
Email: ddukes@cisco.com

Madhan Sankaranarayanan
Cisco Systems, Inc.
India
Email: madsanka@cisco.com