

Internet Area Working Group
Internet-Draft
Intended status: Standards Track
Expires: 8 February 2026

B. Fenner
R. Thomas
Arista Networks
7 August 2025

Adding Extensions to ICMP Errors for Originating Node Identification
draft-ietf-intarea-extended-icmp-nodeid-03

Abstract

RFC5837 describes a mechanism for Extending ICMP for Interface and Next-Hop Identification, which allows providing additional information in an ICMP error that helps identify interfaces participating in the path. This is especially useful in environments where a given interface may not have a unique IP address to respond to, e.g., a traceroute.

This document introduces a similar ICMP extension for Node Identification. It allows providing a unique IP address and/or a textual name for the node, in the case where each node may not have a unique IP address (e.g., a deployment in which all interfaces have IPv6 addresses and all nexthops are IPv6 nexthops, even for IPv4 routes).

About This Document

This note is to be removed before publishing as an RFC.

The latest revision of this draft can be found at <https://fenner.github.io/icmp-node-id/draft-ietf-intarea-extended-icmp-nodeid.html>. Status information for this document may be found at <https://datatracker.ietf.org/doc/draft-ietf-intarea-extended-icmp-nodeid/>.

Discussion of this document takes place on the Internet Area Working Group Working Group mailing list (<mailto:int-area@ietf.org>), which is archived at <https://mailarchive.ietf.org/arch/browse/int-area/>. Subscribe at <https://www.ietf.org/mailman/listinfo/int-area/>.

Source for this draft and an issue tracker can be found at <https://github.com/fenner/icmp-node-id>.

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 8 February 2026.

Copyright Notice

Copyright (c) 2025 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
2. Conventions and Definitions	3
3. Node Identification Object	4
3.1. C-Type Meaning in a Node Identification Object	5
3.1.1. Behavior when additional bits are reserved	6
3.2. IP Address Sub-Object	6
3.3. Name Sub-Object	7
4. Addition of Node Identification Object by Intermediate Nodes	8
5. Security Considerations	8
6. IANA Considerations	8
7. References	9
7.1. Normative References	9
7.2. Informative References	10
Appendix A. Change history	11
A.1. Changes since draft-fenner-intarea-extended-icmp-hostid-00	11
A.2. Changes since draft-fenner-intarea-extended-icmp-hostid-01	11

A.3. Changes since	
draft-fenner-intarea-extended-icmp-hostid-02	11
A.4. Changes since	
draft-ietf-intarea-extended-icmp-nodeid-00	11
A.5. Changes since	
draft-ietf-intarea-extended-icmp-nodeid-01	11
A.6. Changes since	
draft-ietf-intarea-extended-icmp-nodeid-02	12
Acknowledgments	12
Authors' Addresses	12

1. Introduction

In addition to adding incoming interface information to a traceroute using the mechanisms described in [RFC5837], a network operator may be interested in adding information to unambiguously identify nodes themselves. For example, [I-D.chroboczek-intarea-v4-via-v6] describes a scenario in which individual nodes do not have unique IPv4 addresses to use to reply to an IPv4 traceroute, so additional information is needed. Another scenario is described in [I-D.equinox-v6ops-icmpext-xlat-v6only-source]: when an IPv6-only node runs the customer-side translator (CLAT, [RFC6877]), traceroute to an IPv4 destination can not represent intermediate IPv6-only routers.

The goal of this specification is to have a mechanism to provide additional useful information about the identification of a node sending an ICMP error, which depends on the actual context and scope of the message being delivered. To this end, it is RECOMMENDED to use a combination of IP Address and Name sub-objects (including combinations where one of the sub-objects is not used) that is unique and meaningful in the actual context and scope. It is explicitly permitted to use an IP address that may have only local meaning (e.g., ULA [RFC4193]), since that information can then be provided to the operator of the domain who can then determine the local meaning.

This document defines an ICMP extension that fills that void.

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.

ICMPv4 is used to refer to Internet Control Message Protocol (ICMP) specified in [RFC0792].

ICMPv6 is used to refer to Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) specified in [RFC4443].

ICMP is used to refer to both ICMPv4 and ICMPv6.

3. Node Identification Object

This section defines the Node Identification Object, an ICMP Extension Object with a Class-Num (Object Class Value) of 5 (see Section 6).

Similar to Section 4 of [RFC5837], this object can be appended to the following messages:

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

For reasons described in [RFC4884], this extension cannot be appended to any of the currently defined ICMPv4 or ICMPv6 messages other than those listed above.

The extension defined herein MAY be appended to any of the above listed messages and SHOULD be appended whenever required to identify the node and when local policy or security considerations do not supersede this requirement. See Section 5 for suggested configuration regarding including these messages.

Similarly to the Interface Identification Object defined in [RFC5837], there are two different pieces of information that can appear in a Node Identification Object:

1. An IP Address Sub-Object MAY be included, containing an address of sufficient scope to identify the node within the domain. The IP Address Sub-Object is defined in Section 3.2 of this memo.
2. A Name Sub-Object MAY be included, as specified in Section 3.3, containing up to 63 octets of the YANG sys:hostname ([RFC7317]) or another appropriate name uniquely identifying the node.

3.1. C-Type Meaning in a Node Identification Object

The C-Type contains a bitmask describing what information is included in this Node Identification Object (Figure 1). The fields in this bitmask are chosen so that the IPAddr and name bits overlap with the same bits as defined in [RFC5837], so that an implementation that supports exactly these bits can reuse packet generation and parsing code.




Figure 1: C-Type for the Node Identification Object

The following are bit-field definitions for C-Type:

- * Unassigned (bits 0-4): These bits are reserved for future use and MUST be set to 0 on transmit and ignored on receipt.
- * IP Addr (bit 5) : When set, an IP Address Sub-Object is present. When clear, an IP Address Sub-Object is not present. The IP Address Sub-Object is described in Section 3.2 of this memo.
- * Name (bit 6): When set, a Name Sub-Object is included. When clear, it is not included. The Name Sub-Object is described in Section 3.3 of this memo.
- * Un2 (bit 7): This bit is reserved for future use and MUST be set to 0 on transmit and ignored on receipt.

The information included does not self-identify, so this specification defines a specific ordering for sending the information that must be followed.

If bit 5 (IP Address) is set, an IP Address Sub-Object MUST be sent first. If bit 6 (Name) is set, a Name Sub-Object MUST be sent next. The information order is thus: IP Address Sub-Object, Name Sub-Object. Any or all pieces of information may be present or absent, as indicated by the C-Type. Any data that follows these optional pieces of information MUST be ignored.

It is valid (though pointless until additional bits are assigned by IANA) to receive a Node Identification Object where bits 5 and 6 are both 0; this MUST NOT generate a warning or error. A packet with such a Node Identification Object SHOULD be treated as though it contains no Node Identification Object.

3.1.1. Behavior when additional bits are reserved

Bit values SHOULD be assigned from left to right in the diagram above, i.e., starting at zero. The sub-objects associated with each new bit MUST be placed in the packet after the sub-objects defined in this memo. For example, if bit 0 is assigned to the Fooblewomp, a packet with bits 0 and 5 set MUST contain the IP Address Sub-Object, followed by the Fooblewomp sub-object.

If a bit is set that a receiver does not support, followed by a bit that the receiver does support, the receiver MUST ignore all of the additional data, since the length of the unsupported data is unknown.

3.2. IP Address Sub-Object

If the Node Identification Object identifies the node by address, the Object Payload contains an address sufficient to identify the node within the appropriate scope - global or as otherwise configured - as depicted in Figure 2.

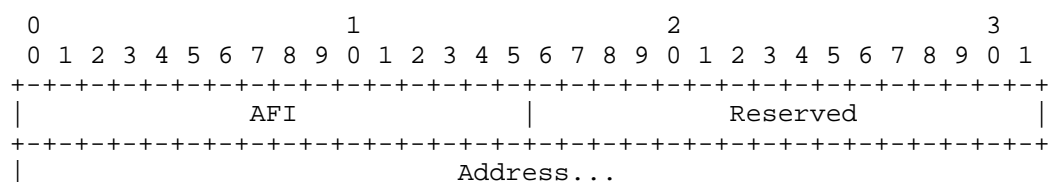


Figure 2: IP Address Sub-Object

Payload fields are defined as follows:

- * Address Family Identifier (AFI): This 16-bit field identifies the type of address represented by the Address field. Values for this field represent a subset of values found in the IANA registry of Address Family Numbers (available from [IANA.address-family-numbers]). Valid values are as follows:
 - 1: 32-bit IPv4 address
 - 2: 128-bit IPv6 address.
- * Reserved: This field MUST be set to 0 and ignored upon receipt.
- * Address: This variable-length field represents an address of appropriate scope (global, if none other defined) that can be used to identify the node. The length of this field is derived from the AFI (i.e., 32 bits if the AFI field is set to 1, and 128 bits if the AFI is set to 2).

3.3. Name Sub-Object

Figure 3 depicts the Name Sub-Object:

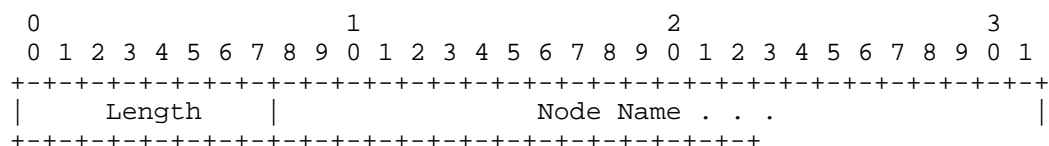
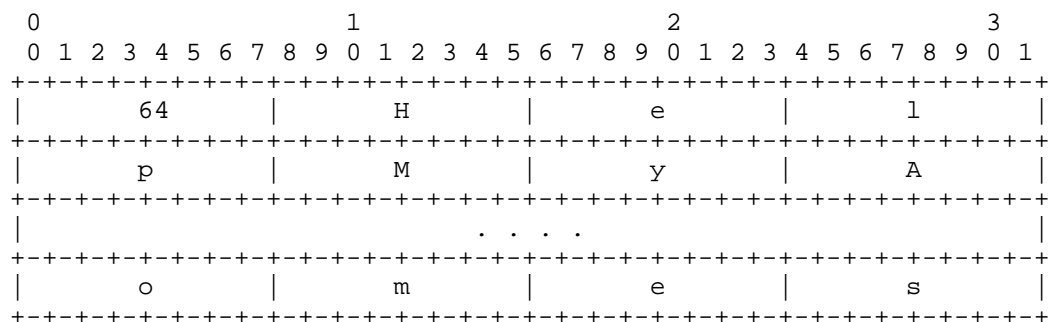


Figure 3: Name Sub-Object

The Name Sub-Object **MUST** have a length that is a multiple of 4 octets and **MUST NOT** exceed 64 octets. If the length field exceeds 64 octets, the receiver **MUST** ignore the sub-object.

The Length field represents the length of the Name Sub-Object, including the length, the node name, and any padding, in octets. The maximum valid length is 64 octets. The length is constrained to ensure there is space for the start of the original packet and additional information.

The second field contains the human-readable node name. The node name **SHOULD** be the YANG `sys:hostname` [RFC7317], if less than 64 octets, or the first 63 octets of the `sys:hostname`, if the `sys:hostname` is longer. The node name **MAY** be some other human-meaningful name of the node. The node name **MUST** be padded with ASCII NUL characters if the object would not otherwise terminate on a 4-octet boundary. An example of truncation of a 66-octet node name, beginning "HelpMyAscii" and ending "Homestar" would be encoded as follows:



The node name **MUST** be represented in the UTF-8 charset [RFC3629] using the Default Language [RFC2277].

4. Addition of Node Identification Object by Intermediate Nodes

An IP/ICMP translator MAY use this extension when translating an ICMP message listed above to include the pre-translation source address of a packet. When doing so, it MUST include the IP Address Sub-Object. If an ICMP Extension Structure is already present in the packet being translated, this Extension Object is appended to the existing ICMP Extension Structure and the checksum is updated. If an ICMP Extension Structure is not present in the packet being translated, one is added using the rules of [RFC4884]. Further details of this mode of operation are outside the scope of this memo.

5. Security Considerations

A node name may reveal sensitive information, especially when it encodes semantic information. It may not be desirable to allow this information to be sent to an arbitrary receiver. The addition of this information to the ICMP responses listed in Section 3 is configurable, and SHOULD be disabled by default, with the exception of IP/ICMP translators [RFC7915]. Those translators SHOULD add the Node Identification Extension Object with the IP Address Sub-Object, as described in [I-D.equinox-v6ops-icmpext-xlat-v6only-source]. An implementation may determine what objects may be appended to a given message based on the destination IP address of the ICMP message that will contain the objects. Access control lists (ACLs) may be used to filter the destinations to which this information may be communicated.

This document does not specify an authentication mechanism for the extension that it defines. Application developers should be aware that ICMP messages and their contents are easily spoofed.

6. IANA Considerations

This IANA has allocated the ICMP Extension Object Class value 5 to the extension described above.

Class Value	Class Name	Reference
5	Node Identification Object	[This document]

Table 1

The corresponding Class Sub-types Registry is as follows:

C-Type (Value)	Description	Reference
0-4	Unassigned - allocatable with Standards Action	[This document]
5	IP Address Sub-object included	[This document]
6	Name Sub-object included	[This document]
7	Unassigned - allocatable with Standards Action	[This document]

Table 2

As mentioned in Section 3.1.1, IANA is requested to assign additional bits starting at zero.

7. References

7.1. Normative References

- [IANA.address-family-numbers] IANA, "Address Family Numbers", <<https://www.iana.org/assignments/address-family-numbers>>.
- [RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, DOI 10.17487/RFC0792, September 1981, <<https://www.rfc-editor.org/rfc/rfc792>>.
- [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>>.
- [RFC2277] Alvestrand, H., "IETF Policy on Character Sets and Languages", BCP 18, RFC 2277, DOI 10.17487/RFC2277, January 1998, <<https://www.rfc-editor.org/rfc/rfc2277>>.
- [RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, DOI 10.17487/RFC3629, November 2003, <<https://www.rfc-editor.org/rfc/rfc3629>>.

- [RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", STD 89, RFC 4443, DOI 10.17487/RFC4443, March 2006, <<https://www.rfc-editor.org/rfc/rfc4443>>.
- [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>>.
- [RFC7317] Bierman, A. and M. Bjorklund, "A YANG Data Model for System Management", RFC 7317, DOI 10.17487/RFC7317, August 2014, <<https://www.rfc-editor.org/rfc/rfc7317>>.
- [RFC7915] Bao, C., Li, X., Baker, F., Anderson, T., and F. Gont, "IP/ICMP Translation Algorithm", RFC 7915, DOI 10.17487/RFC7915, June 2016, <<https://www.rfc-editor.org/rfc/rfc7915>>.
- [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>>.

7.2. Informative References

- [I-D.chroboczek-intarea-v4-via-v6]
Chroboczek, J., Kumari, W., and T. Håland-Jørgensen, "IPv4 routes with an IPv6 next hop", Work in Progress, Internet-Draft, draft-chroboczek-intarea-v4-via-v6-03, 20 January 2025, <<https://datatracker.ietf.org/doc/html/draft-chroboczek-intarea-v4-via-v6-03>>.
- [I-D.equinox-v6ops-icmpext-xlat-v6only-source]
Lamparter, D. and J. Linkova, "Using Dummy IPv4 Address and Node Identification Extensions for IP/ICMP translators (XLATs)", Work in Progress, Internet-Draft, draft-equinox-v6ops-icmpext-xlat-v6only-source-02, 21 March 2025, <<https://datatracker.ietf.org/doc/html/draft-equinox-v6ops-icmpext-xlat-v6only-source-02>>.

- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, DOI 10.17487/RFC4193, October 2005, <<https://www.rfc-editor.org/rfc/rfc4193>>.
- [RFC6877] Mawatari, M., Kawashima, M., and C. Byrne, "464XLAT: Combination of Stateful and Stateless Translation", RFC 6877, DOI 10.17487/RFC6877, April 2013, <<https://www.rfc-editor.org/rfc/rfc6877>>.

Appendix A. Change history

This section is to be removed before publishing as an RFC.

A.1. Changes since draft-fenner-intarea-extended-icmp-hostid-00

- * Instead of having two different messages with the same Class Value and different CType values, we copy the bitmap implementation from [RFC5837]. The re-use of bit positions means that packet parsing and generation code can be largely reused from existing [RFC5837] code.

A.2. Changes since draft-fenner-intarea-extended-icmp-hostid-01

- * Fixed several copy-pasta errors that still referred to interface names instead of node name.

A.3. Changes since draft-fenner-intarea-extended-icmp-hostid-02

- * Renamed to draft-ietf-intarea-extended-icmp-nodeid-00 to reflect adoption by WG

A.4. Changes since draft-ietf-intarea-extended-icmp-nodeid-00

- * Several edits suggested by Med Boucadair
- * Added Section 4 to address the needs of XLAT
- * Changed title to "Adding Extensions to ICMP Errors for Originating Node Identification"

A.5. Changes since draft-ietf-intarea-extended-icmp-nodeid-01

- * Added the request to assign bits starting at 0 to the IANA Considerations
- * Updated IANA Considerations wording based on RFC8126

- * Shortened sub-object names to "IP Address" and "Name", eliminating "Node".

A.6. Changes since draft-ietf-intarea-extended-icmp-nodeid-02

- * Added table to reflect Object Class assignment to IANA Considerations
- * Use SVG for packet figures
- * Add example of an encoded, truncated node name
- * Be explicit on treatment of a packet with no bits set
- * Clarified "defaults to off" in security considerations
- * Clarified use of IP address and names

Acknowledgments

This document derives text heavily from [RFC5837], since the underlying mechanism is identical, and only the semantics of the message differs. Thanks are therefore due to that document's authors: Alia K. Atlas, Ronald P. Bonica, Carlos Pignataro, Naiming Shen, and JR. Rivers.

Further thanks are due to the following who have provided valuable contributions to this document: Med Boucadair, Jen Linkova, David Lamparter, and Luigi Iannone.

Authors' Addresses

Bill Fenner
Arista Networks
5453 Great America Parkway
Santa Clara, California 95054
United States of America
Email: fenner@fenron.com

Reji Thomas
Arista Networks
Global Tech Park
Bangalore 560103
Karnataka
India
Email: reji.thomas@arista.com