6MAN Working Group
Internet-Draft
Updates: 4620, 4884 (if approved)
Intended status: Standards Track
Expires: 12 December 2025

X. Min
ZTE Corp.
G. Mirsky
Ericsson

10 June 2025

IPv6 Query for Enabled In-situ OAM Capabilities
draft-ietf-6man-icmpv6-ioam-conf-state-08

Abstract

This document describes the application of the mechanism of discovering In-situ OAM (IOAM) capabilities, described in RFC 9359 "Echo Request/Reply for Enabled In Situ OAM (IOAM) Capabilities", in IPv6 networks. IPv6 Node IOAM Request uses the IPv6 Node Information messages, allowing the IOAM encapsulating node to discover the enabled IOAM capabilities of each IOAM transit and IOAM decapsulating node.

This document updates RFCs 4620 and 4884.

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 12 December 2025.

Copyright Notice

Table of Contents

1. Introduction

IPv6 encapsulation for In-situ OAM (IOAM) data is defined in [RFC9486], which uses the IPv6 hop-by-hop and destination options to carry IOAM data fields ([RFC9197], [RFC9326]).

As specified in [RFC9359], the echo request/reply can be used by the IOAM encapsulating node to discover the enabled IOAM capabilities at the IOAM transit and decapsulating nodes.

As specified in [RFC4443], the Internet Control Message Protocol for IPv6 (ICMPv6) is an integral part of IPv6. ICMPv6 messages include error messages and informational messages. [RFC4884] defines ICMPv6 Extension Structure by which multi-part ICMPv6 error messages are supported. [RFC8335] updates [RFC4884] by adding two ICMPv6 informational messages, ICMPv6 Extended Echo Request message and ICMPv6 Extended Echo Reply message, to the supporting list of ICMPv6 Extension Structure. Both [RFC4884] and [RFC8335] provide sound principles and examples on extending ICMPv6 messages.

As specified in [RFC4620], two types of IPv6 Node Information messages, the Node Information Query (or NI Query) and the Node Information Reply (or NI Reply), also known as two ICMPv6 informational messages, are used for a Querier node to query information of a Responder node.

This document describes the IPv6 Node IOAM Query functionality, which uses the IPv6 Node Information messages, allowing the IOAM encapsulating node to discover the enabled IOAM capabilities of each IOAM transit and IOAM decapsulating node.

The IOAM encapsulating node sends an NI Query to each IOAM transit and decapsulating node. Upon receiving the query, each node executes access control procedures. If access is granted, the node returns an NI Reply indicating its enabled IOAM capabilities. The NI Reply contains an ICMPv6 Extension Structure customized to this message, and the ICMPv6 Extension Structure contains one or more IOAM Capabilities Objects.

Before the IOAM encapsulating node sends the NI Query, it must know the IPv6 address of each node along the transport path of the data packet to which IOAM data will be added. This can be achieved by executing an ICMPv6/UDP traceroute or by provisioning an explicit path at the IOAM encapsulating node. In an Equal-Cost Multipath (ECMP) scenario, the same values in any ECMP-affecting fields (e.g., the 3-tuple of the Flow Label, Source Address, and Destination Address fields as per [RFC6437]) of the IOAM data packets MUST be populated in the NI Query, ensuring fate sharing between the NI Query and the IOAM data packets.

2. Conventions Used in This Document

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. Node IOAM Request

The Node IOAM Request message is encapsulated in an IPv6 header [RFC8200], like any ICMPv6 message.

The Node IOAM Request message has the following format:

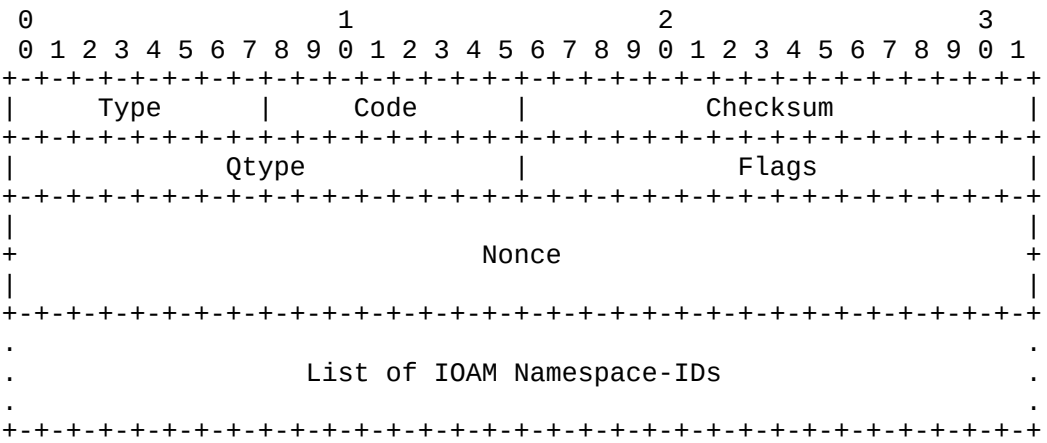

Figure 1: Node IOAM Request Message

IPv6 Header fields:

* Source Address: The Source Address identifies the IOAM encapsulating node. It MUST be a valid IPv6 unicast address.

* Destination Address: The Destination Address identifies the IOAM transit or decapsulating node. It MUST be a valid IPv6 unicast address.

ICMPv6 fields:

* Type: NI Query. The value is 139 as allocated for [RFC4620].

* Code: The value is (TBD1). The Data field contains a list of IOAM Namespace-IDs, which are the subject of this query.

* Checksum: The ICMPv6 checksum.

* Qtype: The value is (TBD2). Which indicates the NI Query is a node IOAM capabilities query.

* Flags: The same as defined in [RFC4620]. Flags are Qtype-specific, the NI Query Qtype used in this document has no defined flags.

* Nonce: The same as defined in [RFC4620].

* Data: Following the NI Query header, the Data field is a List of IOAM Namespace-IDs, which is also called IOAM Capabilities Query Container payload in Section 3.1 of [RFC9359].

3.1. Examples of the Node IOAM Request

The format of a Node IOAM Request can vary from deployment to deployment.

In a deployment where only the default Namespace-ID is used, the Node IOAM Request is depicted as the following:

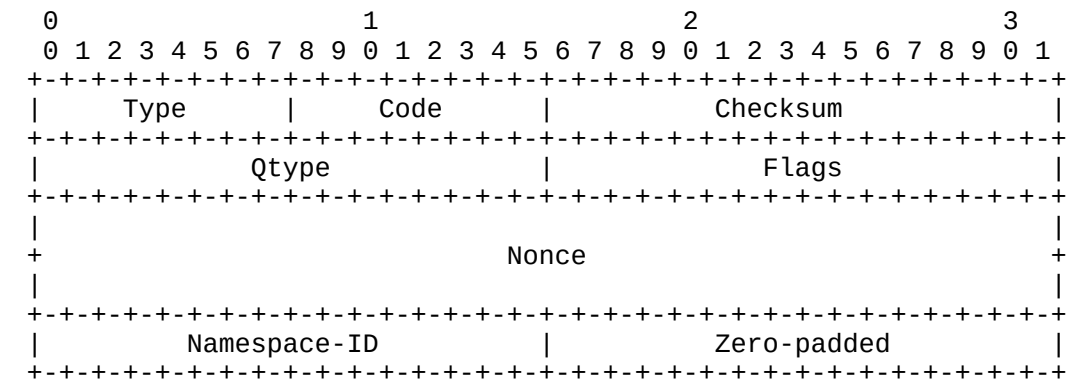

Figure 2: Node IOAM Request of the Default IOAM Namespace

In a deployment where two Namespace-IDs (Namespace-ID1 and Namespace-ID2) are used, the Node IOAM Request is depicted as the following:

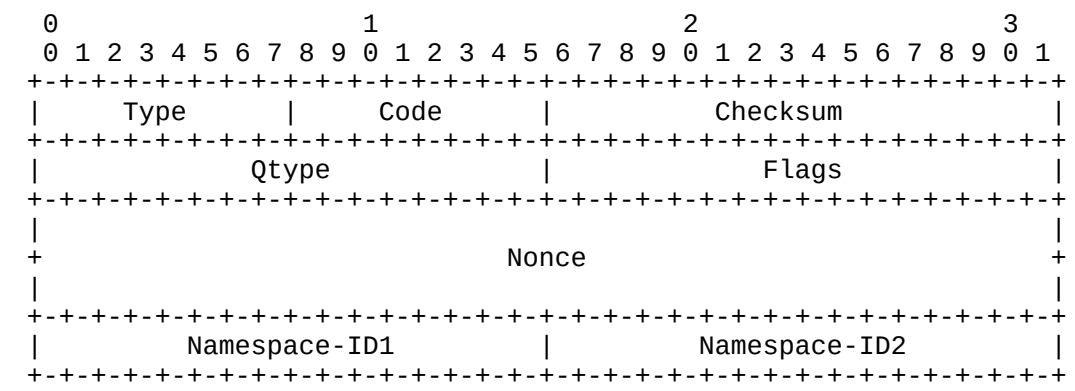

Figure 3: Node IOAM Request of the Two IOAM Namespaces

When a Node IOAM Request message is received, the length of the message is determined by the Payload Length field in the IPv6 Header, as specified in [RFC8200].

4. Node IOAM Reply

The Node IOAM Reply message is encapsulated in an IPv6 header [RFC8200], like any ICMPv6 message.

The Node IOAM Reply message has the following format:

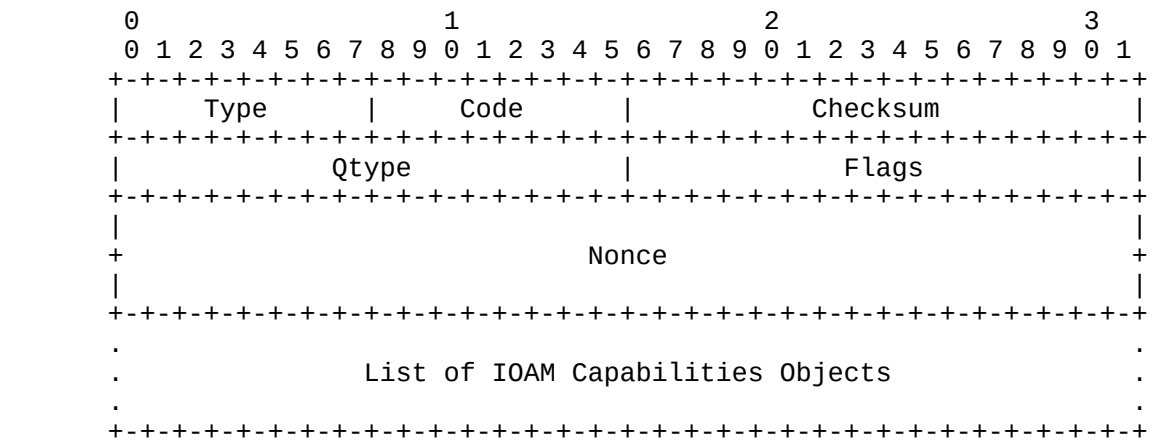

Figure 4: Node IOAM Reply Message

IPv6 Header fields:

* Source Address: Copied from the Destination Address field of the invoking Node IOAM Request packet.

* Destination Address: Copied from the Source Address field of the invoking Node IOAM Request packet.

ICMPv6 fields:

* Type: NI Reply. The value is 140 as allocated for [RFC4620].

* Code: The values are (TBD3) No Matched Namespace-ID, and (TBD4) Exceed the minimum IPv6 MTU. See Section 5 for details.

* Checksum: The ICMPv6 checksum.

* Qtype: Copied from the Qtype field of the invoking Node IOAM Request.

* Flags: The same as defined in [RFC4620]. Flags are Qtype-specific, the NI Reply Qtype used in this document has no defined flags.

* Nonce: Copied from the Nonce field of the invoking Node IOAM Request.

* Data: Following the NI Reply header, the Data field is a List of IOAM Capabilities Objects, which is also called IOAM Capabilities Response Container payload in Section 3.2 of [RFC9359]. Section 7 of [RFC4884] defines the ICMP Extension Structure. As per RFC 4884, the Extension Structure contains exactly one Extension Header followed by one or more objects. When applied to the Node IOAM Reply message, the ICMP Extension Structure MUST contain one or more IOAM Capabilities Objects. Also note that [I-D.ietf-intarea-icmp-exten-hdr-len] updates [RFC4884] by adding a Length field to the ICMP Extension Header.

4.1. IOAM Capabilities Objects

All ICMPv6 IOAM Capabilities Objects are encapsulated in a Node IOAM Reply message.

Each ICMPv6 IOAM Capabilities Object has the following format:

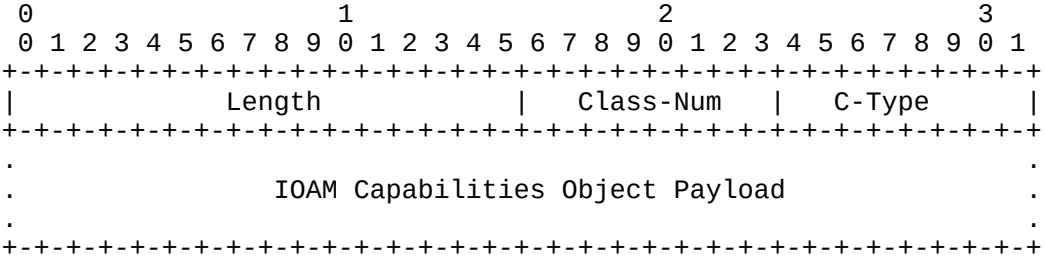

Figure 5: IOAM Capabilities Object

Object fields:

* Class-Num: IOAM Capabilities Objects. The values are listed as the following:

Value	Object Name
TBD5	IOAM Tracing Capabilities Object
TBD6	IOAM Proof of Transit Capabilities Object
TBD7	IOAM Edge-to-Edge Capabilities Object
TBD8	IOAM DEX Capabilities Object
TBD9	IOAM End-of-Domain Object

* C-Type: Values are listed as the following:

Class-Num	C-Type	C-Type Name
TBD5	0	Reserved
	1	Pre-allocated Tracing
TBD6	0	Reserved
TBD7	0	Reserved
TBD8	0	Reserved
TBD9	0	Reserved

* Length: Length of the object, measured in octets, including the Object Header and payload.

* Object payload: Following the IOAM Capabilities Object Header is the IOAM Capabilities Object payload, which is defined in Sections 3.2.1, 3.2.3, 3.2.4, 3.2.5, and 3.2.6 of [RFC9359].

4.2. Examples of the Node IOAM Reply

The format of a Node IOAM Reply can vary from deployment to deployment.

In a deployment where only the default Namespace-ID is used, the IOAM Pre-allocated Tracing Capabilities and the IOAM Proof of Transit Capabilities are enabled at the IOAM transit node that received a Node IOAM Request, the Node IOAM Reply is depicted as the following:

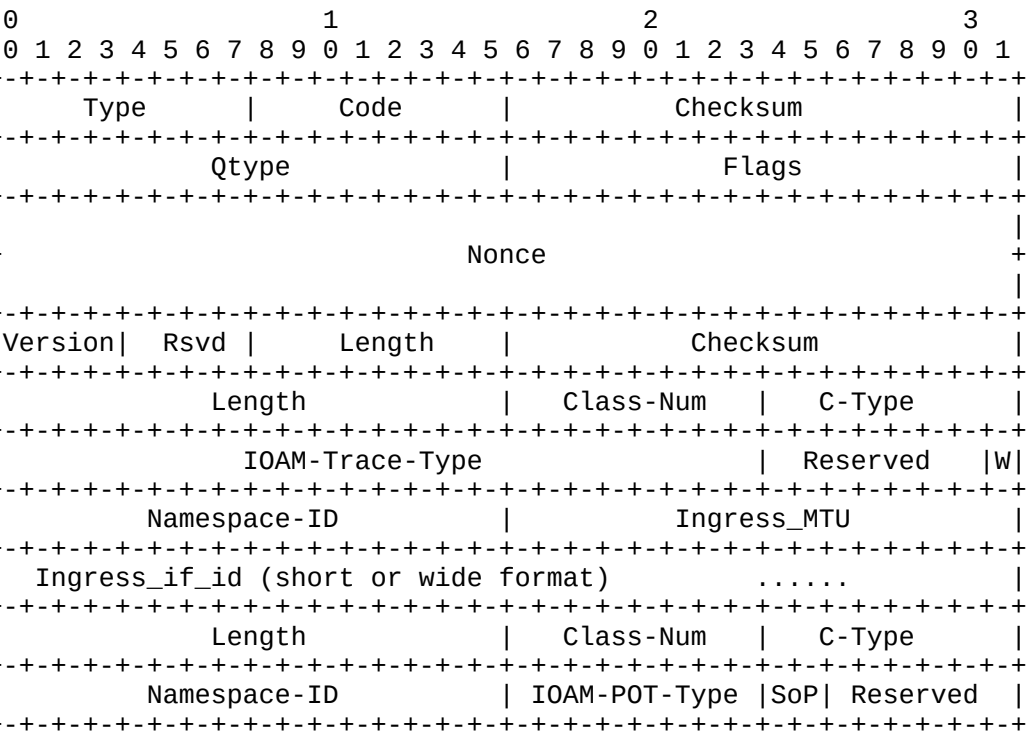

Figure 6: Example 1 of the Node IOAM Reply

In a deployment where two Namespace-IDs (Namespace-ID1 and Namespace-ID2) are used, for both Namespace-ID1 and Namespace-ID2 the IOAM Pre-allocated Tracing Capabilities and the IOAM Proof of Transit Capabilities are enabled at the IOAM transit node that received a Node IOAM Request, the Node IOAM Reply is depicted as the following:

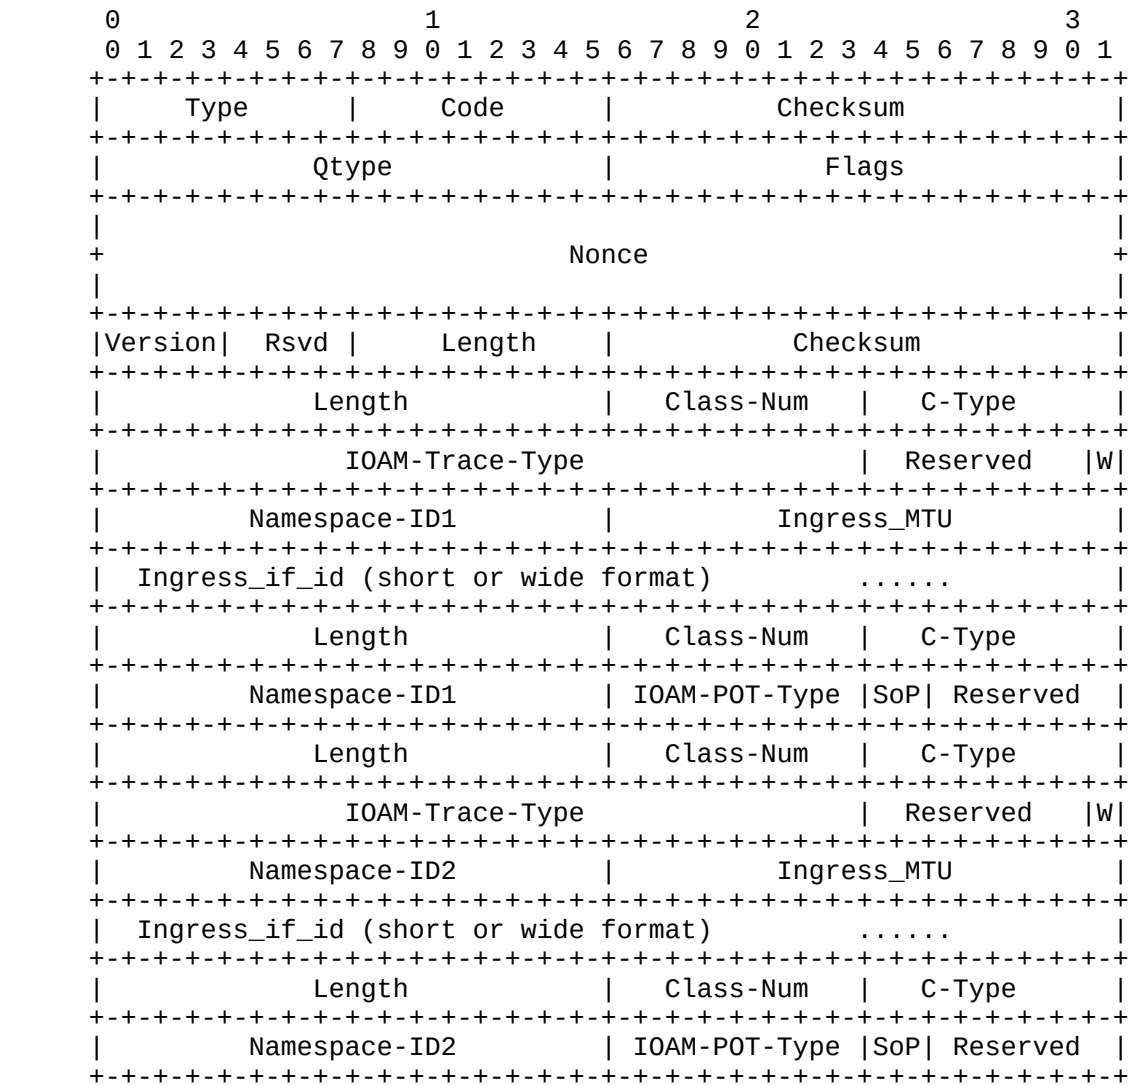

Figure 7: Example 2 of the Node IOAM Reply

In a deployment where only the default Namespace-ID is used, the IOAM Pre-allocated Tracing Capabilities, the IOAM Proof of Transit Capabilities, and the IOAM Edge-to-Edge Capabilities are enabled at the IOAM decapsulating node that received a Node IOAM Request, the Node IOAM Reply is depicted as the following:

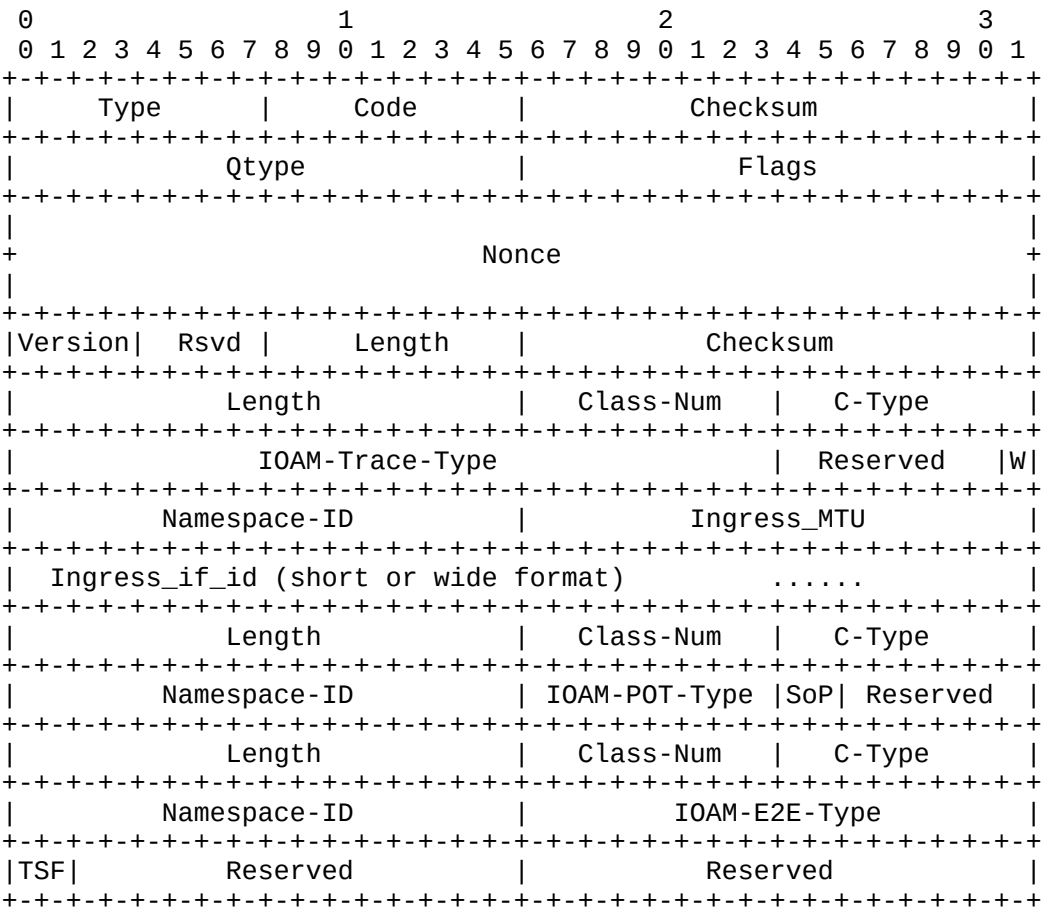

Figure 8: Example 3 of the Node IOAM Reply

When a Node IOAM Reply message is received, the length of the message is determined by the Payload Length field in the IPv6 Header, as specified in [RFC8200].

5. Code Field Processing

The Code field in the Node IOAM Reply MUST be set to (TBD3) No Matched Namespace-ID if any of the following conditions apply:

* The Node IOAM Request does not include any Namespace-ID.

* None of the contained list of IOAM Namespace-IDs is recognized.

* None of the contained list of IOAM Namespace-IDs is enabled.

The Code field in the Node IOAM Reply MUST be set to (TBD4) Exceed the minimum IPv6 MTU if the formatted NI Reply packet exceeds the minimum IPv6 MTU (i.e., 1280 octets). In this case, all objects MUST be stripped before forwarding the Node IOAM Reply to its destination.

6. Updates to RFC 4884

Section 4.6 of [RFC4884] provides a list of extensible ICMP messages (i.e., messages that can carry the ICMP Extension Structure). This document adds the IPv6 Node Information Reply message to that list.

7. IANA Considerations

This document requests the following IANA actions:

* Add the following Code to the "Type 139 - ICMP Node Information Query" sub-registry:
  - (TBD1) The Data field contains a List of IOAM Namespace-IDs which is the Subject of this Query
* Add the following to the "Qtypes" registry:
  - TBD2 Node IOAM Capabilities
* Add the following Codes to the "Type 140 - ICMP Node Information Response" sub-registry:
  - (TBD3) No Matched Namespace-ID
  - (TBD4) Exceed the minimum IPv6 MTU
* Add the following to the "ICMP Extension Object Classes and Class Sub-types" registry:
  - (TBD5) IOAM Tracing Capabilities Object
* Add the following C-types to the "Sub-types - Class TBD5 - IOAM Tracing Capabilities Object" sub-registry:
  - (0) Reserved
  - (1) Pre-allocated Tracing
* Add the following to the "ICMP Extension Object Classes and Class Sub-types" registry:
  - (TBD6) IOAM Proof of Transit Capabilities Object

* Add the following C-types to the "Sub-types - Class TBD6 - IOAM Proof of Transit Capabilities Object" sub-registry:

  - (0) Reserved

* Add the following to the "ICMP Extension Object Classes and Class Sub-types" registry:

  - (TBD7) IOAM Edge-to-Edge Capabilities Object

* Add the following C-types to the "Sub-types - Class TBD7 - IOAM Edge-to-Edge Capabilities Object" sub-registry:

  - (0) Reserved

* Add the following to the "ICMP Extension Object Classes and Class Sub-types" registry:

  - (TBD8) IOAM DEX Capabilities Object

* Add the following C-types to the "Sub-types - Class TBD8 - IOAM DEX Capabilities Object" sub-registry:

  - (0) Reserved

* Add the following to the "ICMP Extension Object Classes and Class Sub-types" registry:

  - (TBD9) IOAM End-of-Domain Object

* Add the following C-types to the "Sub-types - Class TBD9 - IOAM End-of-Domain Object" sub-registry:

  - (0) Reserved

All codes mentioned above are assigned on a First Come First Serve (FCFS) basis with a range of 0-255.

8. Security Considerations

Security issues discussed in [RFC4620] and [RFC9359] apply to this document.

This document recommends using IP Authentication Header [RFC4302] or IP Encapsulating Security Payload Header [RFC4303] to provide integrity protection for IOAM capabilities information.

This document recommends using IP Encapsulating Security Payload Header [RFC4303] to provide privacy protection for IOAM capabilities information.

This document recommends that the network operators establish policies that restrict access to IPv6 Node IOAM Query functionality. In order to enforce these policies, nodes that support IPv6 Node IOAM Query functionality MUST support the following configuration options:

* Enable/disable IPv6 Node IOAM Query functionality. By default, IPv6 Node IOAM Query functionality is disabled.

* Define enabled Namespace-IDs. By default, all Namespace-IDs except the default one (i.e., Namespace-ID 0x0000) are disabled.

* For each enabled Namespace-ID, define the prefixes from which Node IOAM Request messages are permitted.

In order to protect local resources, implementations SHOULD rate-limit incoming Node IOAM Request messages.

Considering the packet size of the Node IOAM Reply could be much larger than that of the Node IOAM Request, to mitigate the potential amplification attack by using the Node IOAM Request with a spoofed source address, which is similar to the amplification attack by sending an ICMPv6 ECHO_REQUEST to ff02::1 with a spoofed source address (refer to Section 2.3.5 of [RFC9099]), an implementation that supports this specification MUST support an option of padding a Node IOAM Request packet to the Path MTU or the minimum IPv6 MTU [RFC8200], which can ensure that the Node IOAM Reply packet would not be larger than the invoking Node IOAM Request packet. The network operators can choose to enforce the padding option or not in their networks.

9. Acknowledgements

The authors would like to acknowledge Eric Vyncke and Erik Kline for their valuable suggestions on using IPv6 Node Information Queries as the basis.

The authors would like to acknowledge Bob Hinden for his valuable suggestions on the ICMPv6 message format.

The authors would like to acknowledge Chongfeng Xie, Zhenqiang Li, David Lamparter, and Daniel King for their review and helpful comments.

10. References

10.1. Normative References

[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>.

[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/info/rfc4443>.

[RFC4620] Crawford, M. and B. Haberman, Ed., "IPv6 Node Information Queries", RFC 4620, DOI 10.17487/RFC4620, August 2006, <https://www.rfc-editor.org/info/rfc4620>.

[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/info/rfc4884>.

[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>.

[RFC9359] Min, X., Mirsky, G., and L. Bo, "Echo Request/Reply for Enabled In Situ OAM (IOAM) Capabilities", RFC 9359, DOI 10.17487/RFC9359, April 2023, <https://www.rfc-editor.org/info/rfc9359>.

[RFC9486] Bhandari, S., Ed. and F. Brockners, Ed., "IPv6 Options for In Situ Operations, Administration, and Maintenance (IOAM)", RFC 9486, DOI 10.17487/RFC9486, September 2023, <https://www.rfc-editor.org/info/rfc9486>.

10.2. Informative References

[I-D.ietf-intarea-icmp-exten-hdr-len] Bonica, R., hexiaoming, X., Min, X., and T. Mizrahi, "ICMP Extension Header Length Field", Work in Progress, Internet-Draft, draft-ietf-intarea-icmp-exten-hdr-len-00, 15 March 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-intarea-icmp-exten-hdr-len-00>.

[RFC4302] Kent, S., "IP Authentication Header", RFC 4302, DOI 10.17487/RFC4302, December 2005, <https://www.rfc-editor.org/info/rfc4302>.

[RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, DOI 10.17487/RFC4303, December 2005, <https://www.rfc-editor.org/info/rfc4303>.

[RFC6437] Amante, S., Carpenter, B., Jiang, S., and J. Rajahalme, "IPv6 Flow Label Specification", RFC 6437, DOI 10.17487/RFC6437, November 2011, <https://www.rfc-editor.org/info/rfc6437>.

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <https://www.rfc-editor.org/info/rfc8200>.

[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/info/rfc8335>.

[RFC9099] Vyncke, ., Chittimaneni, K., Kaeo, M., and E. Rey, "Operational Security Considerations for IPv6 Networks", RFC 9099, DOI 10.17487/RFC9099, August 2021, <https://www.rfc-editor.org/info/rfc9099>.

[RFC9197] Brockners, F., Ed., Bhandari, S., Ed., and T. Mizrahi, Ed., "Data Fields for In Situ Operations, Administration, and Maintenance (IOAM)", RFC 9197, DOI 10.17487/RFC9197, May 2022, <https://www.rfc-editor.org/info/rfc9197>.

[RFC9326] Song, H., Gafni, B., Brockners, F., Bhandari, S., and T. Mizrahi, "In Situ Operations, Administration, and Maintenance (IOAM) Direct Exporting", RFC 9326, DOI 10.17487/RFC9326, November 2022, <https://www.rfc-editor.org/info/rfc9326>.

Authors' Addresses

Xiao Min
ZTE Corp.
Nanjing
China
Phone: +86 18061680168
Email: xiao.min2@zte.com.cn

Greg Mirsky
Ericsson
United States of America

Email: gregimirsky@gmail.com