

Internet Engineering Task Force (IETF)
Request for Comments: 9279
Category: Standards Track
ISSN: 2070-1721

M. Sivakumar
Juniper Networks
S. Venaas
Cisco Systems, Inc.
Z. Zhang
ZTE Corporation
H. Asaeda
NICT
August 2022

Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Version 2 (MLDv2) Message Extension

Abstract

This document specifies a generic mechanism to extend IGMPv3 and Multicast Listener Discovery Version 2 (MLDv2) by using a list of TLVs (Type, Length, and Value).

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9279>.

Copyright Notice

Copyright (c) 2022 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
2. Conventions Used in This Document
3. Extension Format
 - 3.1. Multicast Listener Query Extension
 - 3.2. Version 2 Multicast Listener Report Extension
 - 3.3. IGMP Membership Query Extension
 - 3.4. IGMP Version 3 Membership Report Extension
4. No-op TLV
5. Processing the Extension
6. Applicability and Backwards Compatibility
7. Security Considerations

9. References

9.2. Informative References

Authors' Addresses

This document defines a generic method to extend IGMPv3 [RFC3376] and MLDv2 [RFC3810] messages to accommodate information other than what is contained in the current message formats. This is done by allowing a list of TLVs to be used in the Additional Data section of IGMPv3 and MLDv2 messages. This document defines a registry for such TLVs. Other documents will define their specific types, and their values and semantics. The extension would only be used when at least one TLV is to be added to the message. This extension also applies to the lightweight versions of IGMPv3 and MLDv2 as defined in [RFC5790].

Additional Data is defined for Query messages in IGMPv3 (Section 4.1.10 of [RFC3376]) and MLDv2 (Section 5.1.12 of [RFC3810]), and for Report messages in IGMPv3 (Section 4.2.11 of [RFC3376]) and MLDv2 (Section 5.2.11 of [RFC3810]).

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.

For each of the IGMPv3 and MLDv2 headers, a previously reserved bit is used to indicate the presence of this extension. When this extension is used, the Additional Data of IGMPv3 and MLDv2 messages is formatted as follows. Note that this format contains a variable number of TLVs. It MUST contain at least one TLV.

[illegible]

Figure 1: Extension Format

Extension Type: 2 octets. This identifies a particular Extension Type as defined in the "IGMP/MLD Extension Types" registry. If this is not the first TLV, it will follow immediately after the end of the previous one. There is no alignment or padding.

Extension Length: 2 octets. This specifies the length in octets of the following Extension Value field. The length may be zero if no value is needed.

Extension Value: This field contains the value. The specification defining the Extension Type describes the length and contents of this field.

IGMPv3 and MLDv2 messages are defined so they can fit within the network MTU in order to avoid fragmentation. An IGMPv3/MLDv2 Report message contains a number of records. The records are called Group Records for IGMPv3 and Address Records for MLDv2. When this extension mechanism is used, the number of records in each Report message SHOULD be kept small enough so that the entire message, including any extension TLVs, can fit within the network MTU.

3.1. Multicast Listener Query Extension

The MLDv2 Query message format [RFC3810] with extension is shown below. The E-bit MUST be set to 1 to indicate that the extension is present. Otherwise, it MUST be 0.



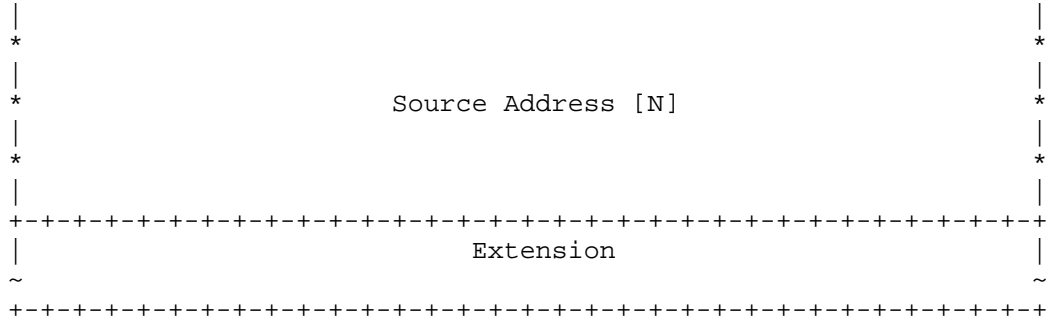


Figure 2: MLD Query Extension

3.2. Version 2 Multicast Listener Report Extension

The MLDv2 Report message format [RFC3810] with extension is shown below. The E-bit MUST be set to 1 to indicate that the extension is present. Otherwise, it MUST be 0.

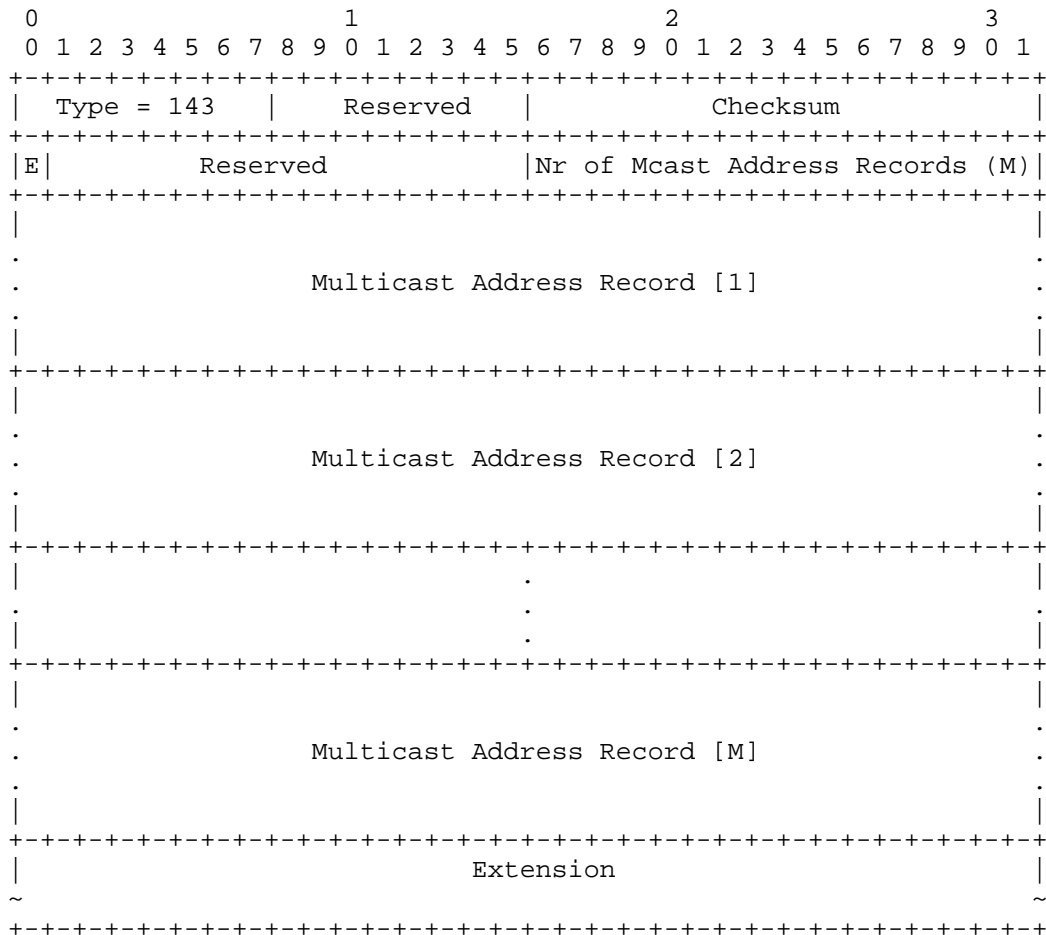



Figure 3: MLD Report Extension

3.3. IGMP Membership Query Extension

The IGMPv3 Query message format [RFC3376] with the extension is shown below. The E-bit MUST be set to 1 to indicate that the extension is present. Otherwise, it MUST be 0.



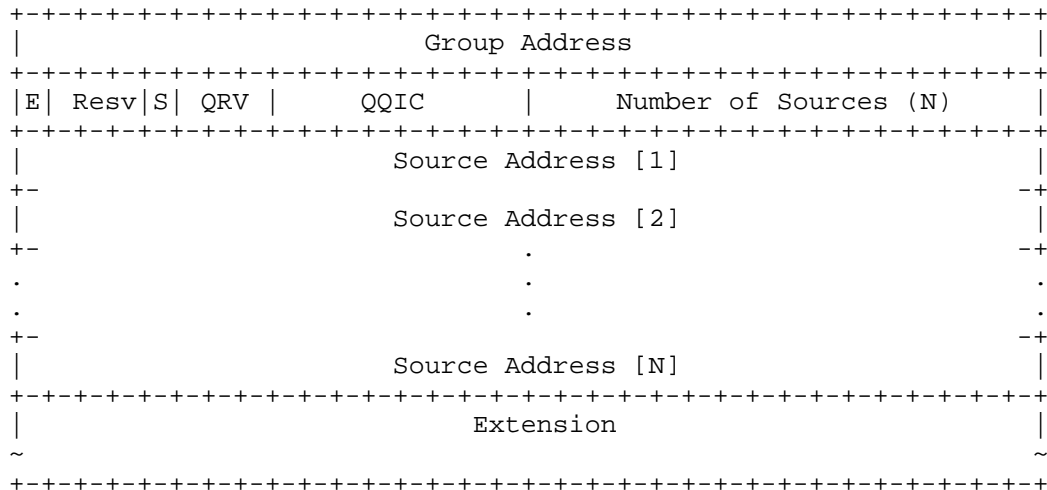


Figure 4: IGMP Query Extension

3.4. IGMP Version 3 Membership Report Extension

The IGMPv3 Report message format [RFC3376] with the extension is shown below. The E-bit MUST be set to 1 to indicate that the extension is present. Otherwise, it MUST be 0.




Figure 5: IGMP Report Extension

4. No-op TLV

The No-op TLV is a No-Operation TLV that MUST be ignored during processing. This TLV may be used to verify that the extension mechanism has been implemented correctly. Note that there is no

alignment requirement, so there is no need to use this Extension Type to provide alignment.




Figure 6: No-op TLV Format

No-op Type: 2 octets. The type of the No-op TLV extension is 0.

Extension Length: 2 octets. This specifies the length in octets of the following Value field. The length may be zero if no value is needed.

Value: This field contains the value. As this Extension Type is always ignored, the value can be arbitrary data. The number of octets used MUST match the specified length.

5. Processing the Extension

The procedure specified in this document only applies when the E-bit is set.

If the validation of the TLVs fails, the entire Additional Data field MUST be ignored as specified in IGMPv3 [RFC3376] and MLDv2 [RFC3810]. The following checks must pass for the validation of the TLVs not to fail:

- * At least one TLV MUST be present.
- * There MUST NOT be any data in the IP payload after the last TLV. To check this, the parser needs to walk through each of the TLVs until there are less than four octets left in the IP payload. If there are any octets left, validation fails.
- * The total length of the Extension MUST NOT exceed the remainder of the IP payload length. For this validation, only the content of the Extension Length fields is examined.

Future documents defining a new Extension Type MUST specify any additional processing and validation. These rules, if any, will be examined only after the general validation succeeds.

TLVs with unsupported Extension Types MUST be ignored.

6. Applicability and Backwards Compatibility

IGMP and MLD implementations, particularly implementations on hosts, rarely change. The adoption process of this extension mechanism is expected to be slow. As new extension TLVs are defined, it may take a long time for them to be supported. Due to this, defining new extension TLVs should not be taken lightly, and it is crucial to consider backwards compatibility.

Implementations that do not support this extension mechanism will ignore it, as specified in [RFC3376] and [RFC3810]. As mentioned in the previous section, unsupported extension TLVs are ignored.

It is possible that a new extension TLV will only apply to queries or

only to reports, or that there may be other specific conditions for when it is to be used. A document defining a new Extension Type MUST specify the conditions under which the new Extension Type should be used, including which message types. It MUST also be specified what the behavior should be if a message is not used in the defined manner, e.g., if it is present in a Query message, when it was only expected to be used in reports.

When defining new Extension Types, the effect of partial support for the new TLV, by either the hosts or routers, on the same link should be carefully considered. Further, whether there are any dependencies or restrictions on combinations between the new Extension Types and any preexisting Extension Types must be considered.

This document defines an extension mechanism only for IGMPv3 and MLDv2. Hence, this mechanism does not apply if hosts or routers send older version messages.

7. Security Considerations

The Security Considerations of [RFC3376] and [RFC3810] also apply here.

This document extends the IGMP and MLD message formats, allowing for a variable number of TLVs. Implementations must take care not to exceed the packet boundary when parsing the TLVs, because an attacker could intentionally specify a TLV with a length exceeding the boundary.

An implementation could add a large number of minimal TLVs in a message to increase the cost of processing the message. This would magnify a denial-of-service attack.

8. IANA Considerations

IANA has created a new registry called "IGMP/MLD Extension Types" in the "Internet Group Management Protocol (IGMP) Type Numbers" section and lists this document as the reference. The registration procedure is "IETF Review" [RFC8126]. The registry is common for IGMP and MLD.

Two Extension Types (65534 and 65535) are provided for "Experimental Use" [RFC8126]. Any experiments should be confined to closed environments where it is unlikely that they may conflict with other experiments; see [RFC3692].

IANA has initially populated the registry as shown in Table 1

Extension Type	Length	Name	Reference
0	variable	No-op	RFC 9279
1-65533		Unassigned	
65534-65535	variable	Reserved for Experimental Use	

Table 1: IGMP/MLD Extension Types

9. References

9.1. 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>>.

[RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, DOI 10.17487/RFC3376, October 2002, <<https://www.rfc-editor.org/info/rfc3376>>.

[RFC3810] Vida, R., Ed. and L. Costa, Ed., "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, DOI 10.17487/RFC3810, June 2004, <<https://www.rfc-editor.org/info/rfc3810>>.

[RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.

[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>>.

9.2. Informative References

[RFC3692] Narten, T., "Assigning Experimental and Testing Numbers Considered Useful", BCP 82, RFC 3692, DOI 10.17487/RFC3692, January 2004, <<https://www.rfc-editor.org/info/rfc3692>>.

[RFC5790] Liu, H., Cao, W., and H. Asaeda, "Lightweight Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Version 2 (MLDv2) Protocols", RFC 5790, DOI 10.17487/RFC5790, February 2010, <<https://www.rfc-editor.org/info/rfc5790>>.

Acknowledgements

The authors thank Ron Bonica, Ian Duncan, Wesley Eddy, Leonard Giuliano, Jake Holland, Tommy Pauly, Pete Resnick, Alvaro Retana, and Zhaohui Zhang for reviewing the document and providing valuable feedback.

Authors' Addresses

Mahesh Sivakumar
Juniper Networks
64 Butler St
Milpitas, CA 95035
United States of America
Email: sivakumar.mahesh@gmail.com

Stig Venaas
Cisco Systems, Inc.
Tasman Drive
San Jose, CA 95134
United States of America
Email: stig@cisco.com

Zheng(Sandy) Zhang
ZTE Corporation
No. 50 Software Ave, Yuhuatai District
Nanjing
210000
China

Email: zhang.zheng@zte.com.cn

Hitoshi Asaeda
National Institute of Information and Communications Technology
4-2-1 Nukui-Kitamachi, Koganei, Tokyo
184-8795
Japan
Email: asaeda@nict.go.jp