

lpwan Working Group
Internet-Draft
Intended status: Informational
Expires: May 1, 2018

D. Barthel
Orange SA
L. Toutain
Institut MINES TELECOM; IMT Atlantique
A. Kandasamy
Acklio
October 28, 2017

LPWAN Static Context Header Compression (SCHC) for ICMPv6
draft-barthel-icmpv6-schc-00

Abstract

ICMPv6 is a companion protocol to IPv6. It defines messages that inform the source of IPv6 packets of errors during packet delivery. It also defines the Echo Request/Reply messages that are used for basic network troubleshooting (ping command). ICMPv6 messages are transported on IPv6.

This document describes how to adapt ICMPv6 to Low Power Wide Area Networks (LPWANs) by compressing ICMPv6/IPv6 headers and by protecting the LPWAN network and the Device from undesirable ICMPv6 traffic.

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 <http://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 May 1, 2018.

Copyright Notice

Copyright (c) 2017 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 (<http://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

| | |
|---|----|
| 1. Introduction | 2 |
| 2. Terminology | 3 |
| 3. Use cases | 3 |
| 4. Detailed behavior | 4 |
| 4.1. Device is the source of an ICMPv6 error message | 4 |
| 4.2. Device is the destination of an ICMPv6 error message | 4 |
| 4.2.1. ICMPv6 error message compression. | 5 |
| 4.3. Device does a ping | 6 |
| 4.4. Device is ping'ed | 8 |
| 5. Traceroute | 9 |
| 6. Security considerations | 10 |
| 7. IANA Considerations | 10 |
| 8. References | 11 |
| 8.1. Normative References | 11 |
| 8.2. Informative References | 11 |
| Authors' Addresses | 12 |

1. Introduction

ICMPv6 [RFC4443] is a companion protocol to IPv6 [RFC8200].

[RFC4443] defines a generic message format. This format is used for messages to be sent back to the source of an IPv6 packet to inform it about errors during packet delivery.

More specifically, [RFC4443] defines 4 error messages: Destination Unreachable, Packet Too Big, Time Exceeded and Parameter Problem.

[RFC4443] also defines the Echo Request and Echo Reply messages, which provide support for the ping application.

Other ICMPv6 messages are defined in other RFCs, such as an extended format of the same messages [RFC4884] and other messages used by the Neighbor Discovery Protocol [RFC4861].

This document focuses on using Static Context Header Compression (SCHC) to compress [RFC4443] messages that need to be transmitted over the LPWAN network, and on having the LPWAN gateway proxying the Device to save it the unwanted traffic.

LPWANs' salient characteristics are described in [I-D.ietf-lpwan-overview]

2. Terminology

This draft re-uses the Terminology defined in [I-D.ietf-lpwan-ipv6-static-context-hc].

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. Use cases

In the LPWAN architecture, we can distinguish the following cases:

- o the Device is the (purported) source of an ICMP error message, mainly in response to an incorrect incoming IPv6 message, or in response to a ping request. In this case, as much as possible, the core SCHC C/D should act as a proxy and originate the ICMP message, so that the Device and the LPWAN network are protected from this unwanted traffic.
- o the Device is the destination of the ICMP message, mainly in response to a packet sent by the Device to the network that generates an error. In this case, we want the ICMP message to reach the Device, and this document describes in section Section 4.2.1 what SCHC compression should be applied.
- o the Device is the originator of an Echo Request message, and therefore the destination of the Echo Reply message.
- o the Device is the destination of an Echo Request message, and therefore the purported source of an Echo Reply message.

These cases are further described in Section 4.

4. Detailed behavior

4.1. Device is the source of an ICMPv6 error message

As stated in [RFC4443], a node should generate an ICMPv6 message in response to an IPv6 packet that is malformed or which cannot be processed due to some incorrect field value.

The general intent of this document is to spare both the Device and the LPWAN network this un-necessary traffic. The incorrect packets should be caught at the core SCHC C/D and the ICMPv6 notification should be sent back from there.

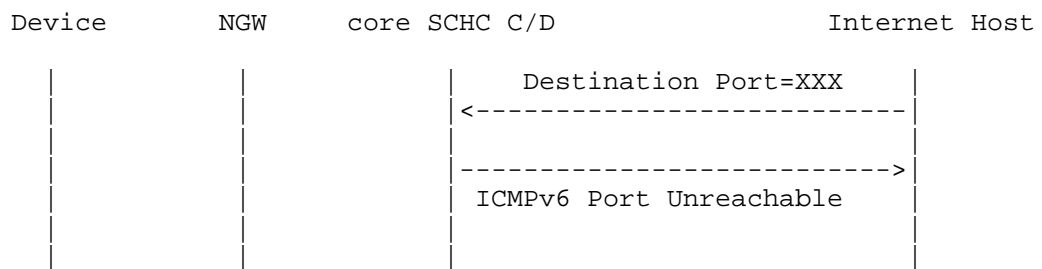


Figure 1: Example of ICMPv6 error message sent back to the Internet

Figure 1 shows an example of an IPv6 packet trying to reach a Device. Let's assume that the port number used as destination port is not "known" (needs better definition) from the core SCHC C/D. Instead of sending the packet over the LPWAN and having this packet rejected by the Device, the core SCHC C/D issues an ICMPv6 error message "Destination Unreachable" (Type 1) with Code 1 ("Port Unreachable") on behalf of the Device.

TODO: This assumes that all ports that the Device listens to will be matched by a SCHC rule. Is this the basic assumption of SCHC that all packets that do not match a rule are rejected? If yes, why do have fragmentation also for uncompressed packets?

TODO: discuss the various Type/Code that are expected to be generated in response to various errors.

4.2. Device is the destination of an ICMPv6 error message

In this situation, we assume that a Device has been configured to send information to a server on the Internet. If this server becomes no longer accessible, an ICMPv6 message will be generated back

towards the Device by an intermediate router. This information can be useful to the Device, for example for reducing the reporting rate in case of periodic reporting of data. Therefore, we compress the ICMPv6 message using SCHC and forward it to the Device over the LPWAN.




Figure 2: Example of ICMPv6 error message sent back to the Device

Figure 2 illustrates this behavior. The ICMPv6 error message is compressed as described in Section 4.2.1 and forwarded over the LPWAN to the Device.

4.2.1. ICMPv6 error message compression.

The ICMPv6 error messages defined in [RFC4443] contain the fields shown in Figure 3.

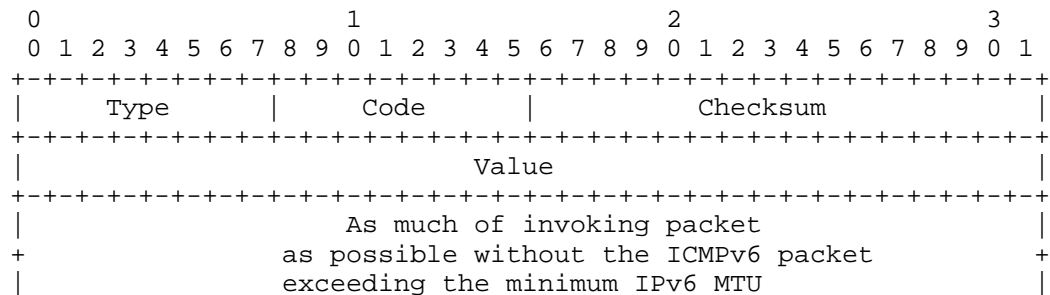


Figure 3: ICMPv6 Error Message format

[RFC4443] states that Type can take the values 1 to 4, and Code can be set to values between 0 and 6. Value is unused for the Destination Unreachable and Time Exceeded messages. It contains the MTU for the Packet Too Big message and a pointer to the byte causing

the error for the Parameter Error message. Therefore, Value is never expected to be greater than 1280 in LPWAN networks.

The following generic rule can therefore be used to compress all ICMPv6 error messages as defined today. More specific rules can also be defined to achieve better compression of some error messages.

The Type field can be associated to a matching list [1, 2, 3, 4] and is therefore compressed down to 2 bits. Code can be reduced to 3 bits using the LSB CDA. Value can be sent on 11 bits using the LSB CDA, but if the Device is known to send smaller packets, then the size of this field can be further reduced.

By [RFC4443], the rest of the ICMPv6 message must contain as much as possible of the IPv6 offending (invoking) packet that triggered this ICMPv6 error message. This information is used to try and identify the SCHC rule that was used to decompress the offending IPv6 packet. If the rule can be found then the Rule Id is added at the end of the compressed ICMPv6 message. Otherwise the compressed packet ends with the compressed Value field.

[RFC4443] states that the "ICMPv6 error message MUST include as much of the IPv6 offending (invoking) packet ... as possible". In order to comply with this requirement, if there is enough information in the incoming ICMPv6 message for the core SCHC C/D to identify the rule that has been used to decompress the erroneous IPv6 packet, this Rule Id must be sent in the compressed ICMPv6 message to the Device. TODO: the erroneous IPv6 packet header (not just the Rule Id) should be sent back. This includes the Rule Id and the compression residue. This means the SCHC C/D uses the context backwards (in the reverse direction). How does the Device know it must also use the context backwards?

TODO: how does one know that the "payload" of a compressed-header packet is in fact another compressed header?

4.3. Device does a ping

If a ping request is generated by a Device, then SCHC compression applies.

The format of an ICMPv6 Echo Request message is described in Figure 4, with Type=128 and Code=0.




Figure 4: ICMPv6 Echo Request message format

If we assume that one rule will be devoted to compressing Echo Request messages, then Type and Code are known in the rule to be 128 and 0 and can therefore be elided with the not-sent CDA.

Checksum can be reconstructed with the compute-checksum CDA and therefore is not transmitted.

[RFC4443] states that Identifier and Sequence Number are meant to "aid in matching Echo Replies to this Echo Request" and that they "may be zero". Data is "zero or more bytes of arbitrary data".

We recommend that Identifier be zero, Sequence Number be a counter on 3 bits, and Data be zero bytes (absent). Therefore, Identifier is elided with the not-sent CDA, Sequence Number is transmitted on 3 bits with the LSB CDA and no Data is transmitted.

The transmission cost of the Echo Request message is therefore the size of the Rule Id + 3 bits.

When the destination receives the Echo Request message, it will respond back with a Echo Reply message. This message bears the same format as the Echo Request message but with Type = 129 (see Figure 4).

[RFC4443] states that the Identifier, Sequence Number and Data fields of the Echo Reply message shall contain the same values as the invoking Echo Request message. Therefore, a rule shall be used similar to that used for compressing the Echo Request message.

TODO: how about a shared rule for Echo Request and Echo Reply with an LSB(1) CDA on the Type field? Or exploiting the Up/Down direction field in the rule?

4.4. Device is ping'ed

If the Device is ping'ed (i.e., is the destination of an Echo Request message), the default behavior is to avoid propagating the Echo Request message over the LPWAN.

This is the recommended behavior with the Code 0 (default value) of the Echo Request message. In addition, this document defines two other Code values to achieve two other behaviors.

The resulting three behaviors are shown on Figure 5 and described below:

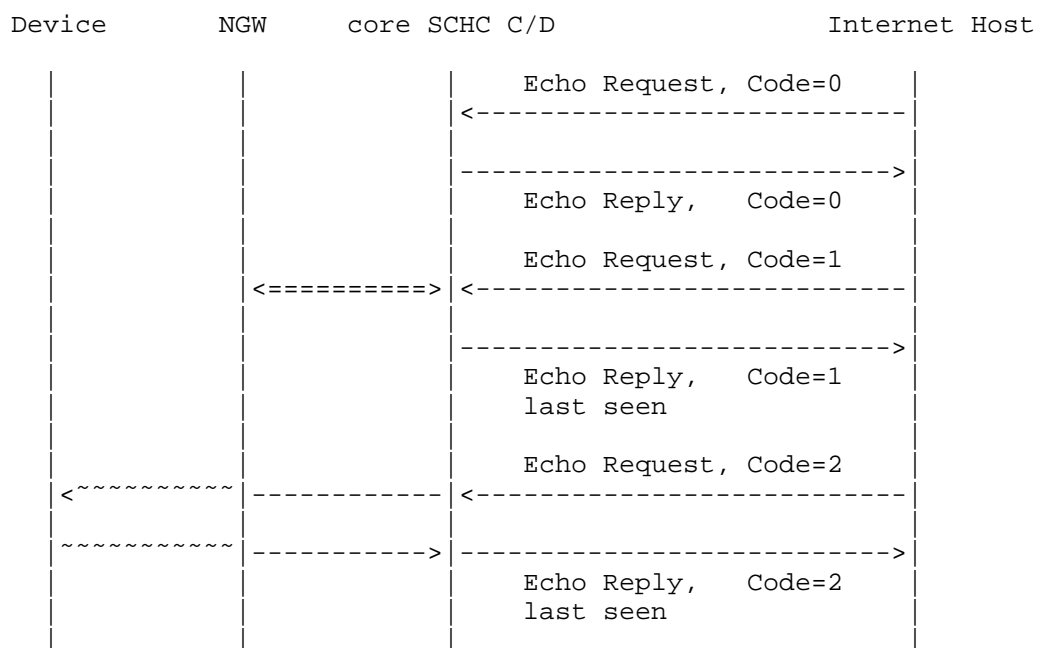


Figure 5: Examples of ICMPv6 Echo Request/Reply

- o Code = 0: The Echo Request message is not propagated on the LPWAN to the Device. If the SCHC C/D finds a rule in the context with the IPv6 address of the Device, it responds with an Echo Reply on behalf of the Device. If no rule is found with that IPv6 address, the SCHC C/D does not respond.

TODO: again, we are assuming that no compression rule is equivalent to the device not providing the service.

- o Code = 1: the SCHC C/D queries the NGW (or maintains a local database) and answers with the number of seconds since the Device last transmission.

TODO: what does it mean to answer "with the number of seconds ..."? There is no such field in an Echo Reply message. Do we overwrite one of the fields (Identifier, Sequencer Number, Data)? They are all supposed to be copied from the Echo Request. Do we change their definition with Code==1 or Code==2?

- o Code = 2: the SCHC C/D compresses the ICMPv6 message and forwards it to the Device. The Echo Reply message sent by the Device is also compressed. Since the Echo Request message comes from the Internet, the values of the Identifier, Sequence Number and Data fields cannot be known in advance, and therefore must be transmitted. However, it is likely that the Echo Request with Code 2 will be firewalled from the Internet and restricted to authorized users. Therefore, the Echo Request message can be assumed to have the same content as recommended in Section 4.3, and the same compression rules apply.

5. Traceroute

The `traceroute6` program sends successive probe packets destined to a chosen target but with the Hop Limit value successively incremented from the initial value 1.

It expects to receive a "Time Exceeded" (Type = 3) "Hop Limit" (Code = 0) ICMPv6 error message back from the successive routers along the path to the destination.

The probe packet is usually a UDP datagram, but can also be a TCP datagram or even an ICMPv6 message. The destination port is chosen in the unassigned range in hope that the destination, when eventually reached, will respond with a "Destination Unreachable" (Type = 1) "Port Unreachable" (Code = 4) ICMPv6 error message.

It is not anticipated that a Device will want to traceroute a destination on the Internet.

By contrast, a host on the Internet may attempt to traceroute an IPv6 address that is assigned to an LPWAN device. This is described in Figure 6.




Figure 6: Example of traceroute to the LPWAN Device

When the probe packet first reaches the core SCHC C/D, its remaining Hop Limit is 1. The core SCHC C/D will respond back with a "Time Exceeded" (Type = 3) "Hop Limit" (Code = 0) ICMPv6 error message. Later on, when the probe packet reaches the core SCHC C/D with a Hop Limit value of 2, the core SCHC C/D will, as explained in Section 4.1, answer back with a "Destination Unreachable" (Type = 1) "Port Unreachable" (Code = 4) ICMPv6 error message. This is what the traceroute6 command expects. Therefore, the traceroute6 command will work with LPWAN IPv6 destinations, except for the time displayed for the destination, which is actually the time to its proxy.

However, if the probe packet happens to hit a port that matches a SCHC rule for that Device, the packet will be compressed with this rule and sent over the LPWAN, which is unfortunate. Forwarding of packets to the Device over the LPWAN should only be done from authenticated/trusted sources anyway. Rate-limitation on top of authentication will mitigate this nuisance.

6. Security considerations

TODO

7. IANA Considerations

TODO

8. References

8.1. Normative References

- [I-D.ietf-lpwan-ipv6-static-context-hc]
Minaburo, A., Toutain, L., and C. Gomez, "LPWAN Static Context Header Compression (SCHC) and fragmentation for IPv6 and UDP", draft-ietf-lpwan-ipv6-static-context-hc-07 (work in progress), October 2017.
- [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>>.
- [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>>.
- [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>>.
- [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>>.

8.2. Informative References

- [I-D.ietf-lpwan-overview]
Farrell, S., "LPWAN Overview", draft-ietf-lpwan-overview-07 (work in progress), October 2017.

Authors' Addresses

Dominique Barthel
Orange SA
28 chemin du Vieux Chene
BP 98
38243 Meylan Cedex
France

Email: dominique.barthel@orange.com

Laurent Toutain
Institut MINES TELECOM; IMT Atlantique
2 rue de la Chataigneraie
CS 17607
35576 Cesson-Sevigne Cedex
France

Email: laurent.toutain@imt-atlantique.fr

Arunprabhu Kandasamy
Acklio
2bis rue de la Chataigneraie
35510 Cesson-Sevigne Cedex
France

Email: arun@ackl.io