Workgroup: Network Working Group
Internet-Draft: draft-jiang-asymmetric-ipv6-04
Published: 15 November 2020
Intended Status: Informational
Expires: 19 May 2021
Authors: S. Jiang
Huawei Technologies Co., Ltd
G. Li
Huawei Technologies
B. E. Carpenter
Univ. of Auckland

Asymmetric IPv6 for Resource-constrained IoT Networks

Abstract

This document describes a new approach to IPv6 header compression for use in scenarios where minimizing packet size is crucial but routing performance must be maximised.

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 May 2021.

Copyright Notice

Table of Contents

1. Introduction

The large address space of IPv6 is essential for the massive expansion of the network edge that will be caused by "Internet of Things" (IoT) technology over low-power or 5G links. However, the size of a raw IPv6 packet header causes difficulty due to the small maximum transmission units (MTU) allowed by typical low-power, low-cost link layers. For 5G, the importance of header overhead in small packets is discussed in [NGMN-5G]. Thus header compression, including address compression, is an important issue. This decreases the size of raw packets, but compressed IP addresses are not routeable except by decompressing them completely in every forwarding node. There are two issues here. The first is the extra computation resource needed for compressing or decompressing in constrained IoT nodes. The second is that full-length IPv6 routing will consume more memory to store routing tables and packet queues (assuming that routing is not bypassed by tunnelling). Such resource consumption is very undesirable in constrained nodes with limited storage, CPU power, and battery capacity.

To mitigate these issues, here we propose a solution enabling the shortening of IPv6 addresses inside packets, and the routing of packets according to short addresses, without needing the overhead of a decompression step prior to route lookup. Considering that the scale and size of edge networks may vary widely, different lengths of short address can be used in different domains.

As an illustrative example, consider an edge network which is known to never require more than a few hundred nodes, which in most cases will communicate either with each other, or with application layer gateways to the rest of the Internet. Rather than needing 128-bit addresses, such a network could very well operate with 16-bit addresses. Also, it could very likely operate without needing

enhancements such as differentiated services, ECN or flow labels. If only IPv6 is supported, the version number field is pointless. There is no reason for IPv6 packets within such a network to contain 40-byte headers as specified in [RFC8200]. Therefore, the useful information could be carried in 8 bytes (see Figure 1). Furthermore, routers within the edge network can route packets directly on 16-bit addresses, reducing RIB and FIB sizes and the lookup time.

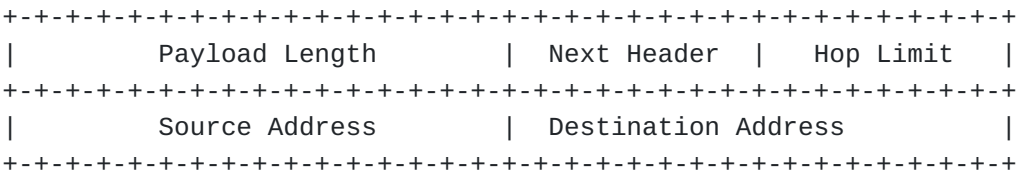

Figure 1

This work is distinct from previous work on address compression [RFC6282] [RFC7400]. Although those solutions tackle the problem of small MTU size, they do not address the problem of decompression overhead.

This work is also distinct from the work on static context header compression [RFC8724], as discussed in more detail below.

Finally, this work is distinct from the 6LoWPAN Routing Header [RFC8138], which can support truncated addresses in a different way.

2. Proposed Solution

The use of IPv6 naturally implies 128-bit addresses for both source and destination. However, this address size is huge by the standards of IoT edge networks. We propose the use of a context parameter to indicate the effective length of the IP address for every node in a local domain. If the effective length is N bits, then all addresses in the domain are assumed to be preceded by a common prefix of 128-N bits, when a full size IPv6 address is needed. Any node in the domain that needs the full address, such as a gateway node to the Internet, can therefore easily synthesize it. If a client communicates with a server that is in the local domain, short addresses will be used end-to-end.

The address length parameter may be needed by every node in the domain. It can be spread by various techniques:

*Configure the address length in every node.

*Obtain the address length from a gateway (next hop router) node.

*Negotiate the address length between neighbors.

The solution operates by shortening IP address fields to save overhead. To enhance this, we propose a new field named Flexible Header Encoding (FHE). It consists of 8 bits, each indicating whether the corresponding IPv6 header field [RFC8200] exists.

- *Bit 0 indicates the Modified Version field
- *Bit 1 indicates the Traffic Class field
- *Bit 2 indicates the Flow Label field.
- *Bit 3 indicates the Payload Length field.
- *Bit 4 indicates the Next Header field. (Zero implies "No Next Header", value 59)
- *Bit 5 indicates the Hop Limit field.
- *Bit 6 indicates the Source Address field.
- *Bit 7 indicates the Destination Address field.

The "Version" field is a special case. In the context of FHE, all packets are presumed to be IPv6 so the normal version field has no purpose. The Modified Version field, if present, has the following encoded meanings:

- *0b0000: The source address (if exist) has pre-determined length inside the domain and the destination address (if exist) uses standard 128-bit IPv6 address. (Outward traffic)
- *0b0001: The source address (if exist) uses standard 128-bit IPv6 address and the destination address (if exist) has pre-determined length inside the domain. (Inward traffic)
- *0b0010: The source address and destination address have the same length inside the domain. The address length will be pre-determined.
- *0b0110: Reserved for IPv6 compatible case.
- *0b0100: Reserved for IPv4 compatible case.
- *0b0011~0b1111(except 0b0110, 0b0100): Reserved.

All fields, including the Modified Version field, follow the FHE in the same order as in [RFC8200], with no padding. There are no alignment requirements, but when a packet is decompressed to a normal IPv6 format, padding options as defined in RFC8200 must be inserted.

Compared to the illustrative example in Figure 1, the actual packet size would therefore be 10 bytes, a considerable improvement on the standard 40 bytes.

One implication of the above is that the source and destination addresses may be elided completely if they are implicit. Sourceless packets were originally suggested in [Crowcroft].

Figure 2 illustrates an example of the FHE format. In this example the traffic class, flow label and source address are elided, and the destination address is truncated to 16 bits. The modified version field could be 0b0001 or 0b0010.

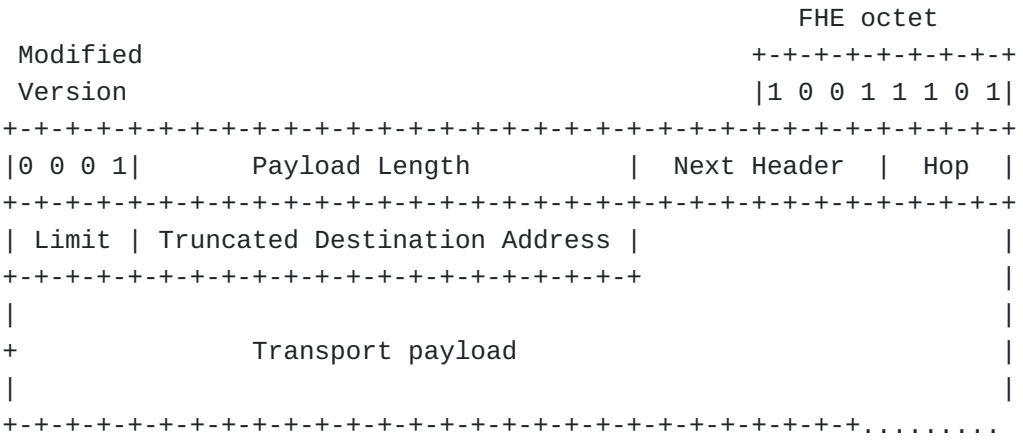

Figure 2

Note that Asymmetric IPv6 does not contain any special handling for IPv6 fragmentation, which will operate exactly as described in [RFC8200], with Asymmetric IPv6 applied to each fragment packet. However, we assume that in IoT deployment scenarios, packets whose length exceeds the IPv6 minimum link MTU before applying Asymmetric IPv6 will be rare. If the underlying link layer cannot carry complete packets even after applying Asymmetric IPv6 compression, an adaptation layer will be necessary exactly as for normal IPv6.

3. Address Transformation at the Gateway

Truncated intra-domain addresses will be used to identify nodes inside the domain. When a packet is sent from an IoT node to an external IPv6 host , the node's intra-domain address, which is unique in the domain, will be carried in the source address field. When the packet is forwarded outside the domain by a gateway, the intra-domain address will be transformed to a complete IPv6 address. To achieve this, the gateway should will maintain a globally routeable prefix for all the nodes in the domain. When a packet with an intra-domain source address is received, the gateway extracts this address and concatenates it to the prefix to form a standard,

globally unique IPv6 address. Vice versa, when IPv6 packets are received from the Internet, the prefix will be removed to recover the intra-domain short address.

There are two options for handling the addresses of external hosts within the domain. One is to use their full IPv6 addresses via Modified Version codes 0b0000 and 0b0001. The other is effectively a specialized form of Network Address Translation. Here, the gateway will maintain a dynamic mapping table between synthetic intra-domain addresses and IPv6 addresses. As packets are received, the gateway performs the appropriate mapping. The transformation must be checksum-neutral for the transport layer, so the methods designed for NAT46 should be adapted [RFC6145].

It is an engineering choice whether this method is preferable to carrying full 128-bit addresses on the IOT side. Which type of resource is more expensive should be seriously considered to choose the appropriate ways, e.g. computing, memory, or transmiting in various resource-constrained IoT networks.

4. Routing without Decompression

Routing mechanisms may readily be adapted to truncated address sizes. If there is routing with an HFE domain, we assume that the truncated address size will be split into a prefix and an interface identifier, but this will not be at the traditional /64 boundary. If routing between different length addresses is required, a suitably modified Forwarding Information Base (FIB) structure is needed, as for any variable length addressing scheme. A truncated address needs to be virtually expanded to 128 bits at the router's inbound interface, although this may not be the physical implementation.

A possible routing choice for IOT edge networks is RPL [RFC6550], although a more complete survey can be found in [Talwar].

5. Address Configuration

The simplest approach to address configuration is simply to run normal IPv6 procedures (SLAAC or DHCPv6), on the argument that this is a rare process and the overhead does not matter. If the truncated address size is less than 64 bits, it will be necessary to use shorter interface identifiers than normal, but this is not a major change. Once a node has acquired an IPv6 address and has learned the local address length parameter as outlined in Section 2, it can continue in FHE mode.

6. Compatibility with Existing Protocols

Although HFE nodes can only talk directly to each other, they are essentially a special form of IPv6 node and they can communicate

with the whole IPv6 Internet via gateways. The complexity is not greater than 6LoWPAN. If appropriate, the 6LoWPAN adaptation layer [RFC4944] could be used, with a specific dispatch type.

7. Relationship to Static Context Header Compression

Static Context Header Compression (SCHC) [RFC8724] is a powerful mechanism for reducing IPv6 packet size in an IoT application environment. In particular it includes a profile for UDP over IPv6, and a somewhat modified version of this profile could achieve much of what Asymmetric IPv6 proposes. In addition, SCHC provides support for fragmentation in the case of very small link MTUs. However, SCHC is by design static, and once a context is established the fields to be compressed do not change. Asymmetric IPv6 transmits the FHE and Modified Version bytes with every packet, so it provides dynamic choice as to which header elements are compressed or elided.

In a context where the desirable compression is fixed, e.g. every address is the same length, the flow label is never used, etc., SCHC can used to the same effect as Asymmetric IPv6. However, if the behavior needs to be dynamic, the signaling power of the FHE and Modified Version bytes in Asymmetric IPv6 is needed.

Further study is needed whether the advantages of the two mechanisms can be combined.

8. Security Considerations

HFE is essentially only a non-cryptographic compression technique so it neither adds to nor reduces the intrinsic security of an IPv6 packet. The address length parameter is not a secret, since all nodes in the domain must know it. The mechanism for distributing this parameter must be no less secure than any other configuration mechanism in us.

Address-based privacy issues must be considered in deciding on the address length. If the number of bits available for the interface identifier is significantly less than the 64 currently in use, address traceability and guessability will be affected. However, if the traffic with short addresses is confined to within the edge network, the privacy issue will be minimized. [RFC7721] and [RFC7217] should be consulted prior to deciding the address length.

9. IANA Considerations

This document makes no request of the IANA.

NOTE IN DRAFT: If the solution of a 6LoWPAN dispatch type is adopted, a suitable assignment request will be added.

10. Acknowledgements

Useful comments were received from Uma Chunduri, Cheng Li, Pascal Thubert, Laurent Toutain and others.

11. References

[Crowcroft] Crowcroft, J. and M. Bagnulo, "SNA: Sourceless Network Architecture", University of Cambridge Computer Laboratory Technical Report UCAM-CL-TR-849, 2014.

[NGMN-5G] Thibault, I., "5G Extreme Requirements: Operators' views on fundamental trade-offs", NGMN Alliance , 2017.

[RFC4944] Montenegro, G., Kushalnagar, N., Hui, J., and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks", RFC 4944, DOI 10.17487/RFC4944, September 2007, <https://www.rfc-editor.org/info/rfc4944>.

[RFC6145] Li, X., Bao, C., and F. Baker, "IP/ICMP Translation Algorithm", RFC 6145, DOI 10.17487/RFC6145, April 2011, <https://www.rfc-editor.org/info/rfc6145>.

[RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", RFC 6282, DOI 10.17487/RFC6282, September 2011, <https://www.rfc-editor.org/info/rfc6282>.

[RFC6550] Winter, T., Ed., Thubert, P., Ed., Brandt, A., Hui, J., Kelsey, R., Levis, P., Pister, K., Struik, R., Vasseur, JP., and R. Alexander, "RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks", RFC 6550, DOI 10.17487/RFC6550, March 2012, <https://www.rfc-editor.org/info/rfc6550>.

[RFC7217] Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", RFC 7217, DOI 10.17487/RFC7217, April 2014, <https://www.rfc-editor.org/info/rfc7217>.

[RFC7400] Bormann, C., "6LoWPAN-GHC: Generic Header Compression for IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs)", RFC 7400, DOI 10.17487/RFC7400, November 2014, <https://www.rfc-editor.org/info/rfc7400>.

[RFC7721] Cooper, A., Gont, F., and D. Thaler, "Security and Privacy Considerations for IPv6 Address Generation

Mechanisms", RFC 7721, DOI 10.17487/RFC7721, March 2016, <https://www.rfc-editor.org/info/rfc7721>.

[RFC8138] Thubert, P., Ed., Bormann, C., Toutain, L., and R. Cragie, "IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN) Routing Header", RFC 8138, DOI 10.17487/RFC8138, April 2017, <https://www.rfc-editor.org/info/rfc8138>.

[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>.

[RFC8724] Minaburo, A., Toutain, L., Gomez, C., Barthel, D., and JC. Zúñiga, "SCHC: Generic Framework for Static Context Header Compression and Fragmentation", RFC 8724, DOI 10.17487/RFC8724, April 2020, <https://www.rfc-editor.org/info/rfc8724>.

[Talwar] Talwar, M., "Routing Techniques and Protocols for Internet of Things: a Survey", Indian J.Sci.Res. 12(1): 417-423, 2015.

Appendix A. Change log [RFC Editor: Please remove]

*draft-jiang-asymmetric-ipv6-00, 2019-06-03:

-Initial version

*draft-jiang-asymmetric-ipv6-01, 2019-06-21:

-Fixed reference error

*draft-jiang-asymmetric-ipv6-02, 2019-10-29:

-Added illustrative example

-Discussed fragmentation

-Discussed relationship to SCHC

-Fixed bit pattern errors

*draft-jiang-asymmetric-ipv6-03, 2020-05-15:

-Minor technical and editorial fixes

-Converted to xml2rfc v3

*draft-jiang-asymmetric-ipv6-04, 2020-11-15:

-Explicitly limit the scope to resource-constrained domain

-Add engineering choice considerations accordingly

Authors' Addresses

Sheng Jiang
Huawei Technologies Co., Ltd
Q14, Huawei Campus, No.156 Beiqing Road
Hai-Dian District, Beijing, 100095
P.R. China

Email: jiangsheng@huawei.com

Guangpeng Li
Huawei Technologies
Q14, Huawei Campus
No.156 Beiqing Road
Hai-Dian District, Beijing
100095
P.R. China

Email: liguangpeng@huawei.com

Brian Carpenter
The University of Auckland
School of Computer Science
University of Auckland
PB 92019
Auckland 1142
New Zealand

Email: brian.e.carpenter@gmail.com