

Transport Area Working Group
Internet-Draft
Intended status: Standards Track
Expires: July 31, 2017

J. Saldana
University of Zaragoza
January 27, 2017

Simplemux. A generic multiplexing protocol
draft-saldana-tsvwg-simplemux-06

Abstract

The high amount of small packets present in nowadays networks results in a low efficiency, as the size of the headers and the payload of these packets can be in the same order of magnitude. In some situations, multiplexing a number of small packets into a bigger one is desirable in order to improve the efficiency. For example, a number of small packets can be sent together between a pair of machines if they share a common network path. Thus, the traffic profile can be shifted from small to larger packets, reducing the network overhead and the number of packets per second to be managed by intermediate routers.

This document describes Simplemux, a protocol able to encapsulate a number of packets belonging to different protocols into a single packet. Small headers (separators) are added at the beginning of each multiplexed packet, including some flags, the packet length and a "Protocol" field. This allows the inclusion of a number of packets belonging to different protocols (the "multiplexed packets") on a packet of another protocol (the "tunneling protocol").

In order to reduce the overhead, the size of the multiplexing headers is kept very low (it may be a single byte when multiplexing small packets).

Status of This Memo

This Internet-Draft is submitted to IETF 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 July 31, 2017.

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.

Table of Contents

1. Introduction	2
1.1. Requirements Language	3
1.2. Existing multiplexing protocols	3
1.3. Benefits of multiplexing	5
2. Description of the scenario	6
3. Protocol description	6
4. Acknowledgements	13
5. IANA Considerations	13
6. Security Considerations	13
7. References	13
7.1. Normative References	13
7.2. Informative References	14
Author's Address	14

1. Introduction

The high amount of small packets present in nowadays networks results in a low efficiency, when the size of the headers and the payload are in the same order of magnitude. In some situations, multiplexing a number of small packets into a bigger one is desirable in order to improve the efficiency. For example, a number of small packets can be sent together between a pair of machines if they share a common network path. Thus, the traffic profile can be shifted from small to larger packets, thus reducing the network overhead and the number of packets per second to be managed by intermediate routers.

This document describes Simplemux, a protocol able to encapsulate a number of packets belonging to different protocols into a single packet. This can be useful e.g. for grouping small packets and thus reducing the number of packets per second in a network.

Simplemux is a generic multiplexing protocol, i.e. it can be used to aggregate a number of packets belonging to a protocol, on a single packet belonging to other (or the same) protocol. Thus, in this document we will talk about the "multiplexed" protocol, and the "tunneling" protocol, being Simplemux the "multiplexing" protocol. The "external header" will be the one of the "tunneling" protocol (see the figure (Figure 1))




Figure 1

As an example, if a number of small IPv6 packets have to travel over an IPv4 network, they can be multiplexed and put into a single IPv4 packet. In this case, IPv4 is the "tunneling" protocol and IPv6 is the "multiplexed" protocol. The IPv4 header is called in this case the "tunneling" or the "external" header. The simplified scheme of this packet would be:

```
|IPv4 hdr||Simplemux hdr|IPv6 packet||Simplemux hdr|IPv6 packet|...|
```

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

1.2. Existing multiplexing protocols

Different multiplexing protocols have been approved by the IETF in the past:

- o TMux [RFC1692]

TMux is able to combine multiple short transport segments, independent of application type, and send them between a server and host pair. As stated in the reference, "The TMux protocol is intended to optimize the transmission of large numbers of small data packets. In particular, communication load is not measured only in bits per seconds but also in packets per seconds, and in many

situation the latter is the true performance limit, not the former. The proposed multiplexing is aimed at alleviating this situation."

A TMux message appears as:

```
|IP hdr||TMux hdr|Transport segment||TMux hdr|Transport segment||...|
```

Therefore, the Transport Segment is not an entire IP packet, since it does not include the IP header.

TMux works "between a server and host pair," so it multiplexes a number of segments between the same pair of machines. However, there are scenarios where a number of low-efficiency flows share a common path, but they do not travel between the same pair of machines.

o PPPMux [RFC3153]

PPPMux "sends multiple PPP encapsulated packets in a single PPP frame. As a result, the PPP overhead per packet is reduced." Thus, it is able to multiplex complete IP packets, using separators.

However, the use of PPPMux requires the use of PPP and L2TP in order to multiplex a number of packets together, as done in TCRTP [RFC4170]. Thus, it introduces more overhead and complexity.

An IP packet including a number of them using PPPMux appears as:

```
|IP hdr|L2TP hdr|PPP hdr||PPPMux hdr|packet||PPPMux hdr|packet||...|
```

The scheme proposed by PPPMux is similar to the Compound-Frames of PPP LCP Extensions [RFC1570]. The key differences are that PPPMux is more efficient and that it allows concatenation of variable sized frames.

The definition of a protocol able to multiplex complete packets, avoiding the need of other protocols as PPP is seen as convenient. The multiplexed packets can be of any kind, since a "Protocol Number" field can be added to each of them. Not all the packets multiplexed together must belong to the same protocol. The general scheme of Simplemux is:

```
|tunnel hdr||Simplemux hdr|packet||Simplemux hdr|packet||...|
```

The Simplemux header includes the "Protocol Number" field, so it permits the multiplexing of different kinds of packets in the same bundle.

We will also refer to the Simplemux header with the terms "separator," "Simplemux separator" or "mux separator". In the figures we will also use the abbreviation "Smux".

When applied to IP packets, the scheme of a multiplexed packet becomes:

```
|tunnel hdr||Simplemux hdr|IP packet||Simplemux hdr|IP packet||...|
```

1.3. Benefits of multiplexing

The benefits of multiplexing are:

- Tunneling a number of packets together. If a number of packets have to be tunneled through a network segment, they can be multiplexed and then sent together using a single external header. This will avoid the need for adding a tunneling header to each of the packets, thus reducing the overhead.
- Reduction of the amount of packets per second in the network. It is desirable for two main reasons: first, network equipment has a limitation in terms of the number of packets per second it can manage, i.e. many devices are not able to send small packets back to back due to processing delay.
- Bandwidth reduction. The presence of high rates of tiny packets translates into an inefficient usage of network resources, so there is a need for mechanisms able to reduce the overhead introduced by low-efficiency flows. When combined with header compression, as done in TCRTTP [RFC4170] multiplexing may produce significant bandwidth savings, which are interesting for network operators, since they may alleviate the traffic load in their networks.
- Energy savings: a lower amount of packets per second will reduce energy consumption in network equipment since, according to [Bolla], internal packet processing engines and switching fabric require 60% and 18% of the power consumption of high-end routers respectively. Thus, reducing the number of packets to be managed and switched will reduce the overall energy consumption. The measurements deployed in [Chabarek] on commercial routers corroborate this. A study using different packet sizes was presented, and the tests with big packets showed that energy consumption gets reduced, since a non-negligible amount of energy is associated to header processing tasks, and not only to the sending of the packet itself.

2. Description of the scenario

Simplemux works between a pair of machines. It creates a tunnel between an "ingress" and an "egress". They MAY be the endpoints of the communication, but they MAY also be middleboxes able to multiplex packets belonging to different flows. Different mechanisms MAY be used in order to classify flows according to some criteria (sharing a common path, kind of service, etc.) and to select the flows to be multiplexed and sent to the egress (see Figure 2).




Figure 2

3. Protocol description

A Simplemux packet consists of:

- An external header that is used as the tunneling header for the whole packet.
- A series of pairs "Simplemux header" + "packet" of the multiplexed protocol.

This is the scheme of a Simplemux packet:

```
|tun hdr||Simplemux hdr|packet||Simplemux hdr|packet||...|
```

The Simplemux header has two different forms: one for the "First Simplemux header," and another one for the rest of the Simplemux headers (called "Non-first Simplemux headers"):

- o First Simplemux header (after the tunneling header, and before the first multiplexed packet):

In order to allow the multiplexing of packets of any length, the number of bytes expressing the length is variable, and a field called "Length Extension" (LXT, one bit) is used to flag if the current byte is the last one including length information. This is the structure of a First Simplemux header:

```
|SPB(1 bit)|LXT(1 bit)|length (6 bits)||LXT(1 bit)|length (7
bits)||...||Protocol (8 bits)|
```

- Single Protocol Bit (SPB, one bit) only appears in the first Simplemux header. It is set to 1 if all the multiplexed packets belong to the same protocol (in this case, the "Protocol" field will only appear in the first Simplemux header). It is set to 0 when each packet MAY belong to a different protocol.

- Length Extension (LXT, one bit) is 0 if the current byte is the last byte where the length of the first packet is included, and 1 in other case.

- Length (LEN, 6, 13, 20, etc. bits): This is the length of the multiplexed packet (in bytes), not including the length field. If the length of the multiplexed packet is less than 64 bytes (less than or equal to 63 bytes), the first LXT is set to 0 and the 6 bits of the length field are the length of the multiplexed packet. If the length of the multiplexed packet is equal or greater than 64 bytes, additional bytes are added. The first bit of each of the added bytes is the LXT. If LXT is set to 1, it means that there is an additional byte for expressing the length. This allows to multiplex packets of any length (see the next figures).

- Protocol (8 bits) is the Protocol field of the multiplexed packet, according to IANA "Assigned Internet Protocol Numbers."

As an example, a First Simplemux header before a packet smaller than 64 (2^6) bytes would be 2 bytes long:




Figure 3

A First Simplemux header before a packet with a length greater or equal to 64 bytes, and smaller than 8192 bytes (2^{13}) will be 3 bytes long:




Figure 4

In this case, the length of the packet will be the number expressed by the concatenation of the bits of Length 1 - Length 2 (total 13 bits). Length 1 includes the 6 most significant bits and Length 2 the 7 less significant bits.

A First Simplemux header before a packet with a length greater of equal to 8192 bytes, and smaller than 1048576 bytes (2^{20}) would be 4 bytes long:




Figure 5

In this case, the length of the packet will be the number expressed by the concatenation of the bits of Length 1 - Length 2 - Length 3 (total 20 bits). Length 1 includes the 6 most significant bits and Length 3 the less 7 significant bits.

More bytes can be added to the length if required, using the same scheme: 1 LXT byte plus 7 bits for expressing the length.

- o Subsequent (Non-first) Simplemux headers (before the other packets):

The Non-first Simplemux headers also employ a format allowing the multiplexing of packets of any length, so the number of bytes expressing the length is variable, and the field Length Extension

(LXT, one bit) is used to flag if the current byte is the last one including length information. This is the structure of a Non-first Simplemux header:

```
|LXT(1 bit)|length (7 bits)||LXT(1 bit)|length (7
bits)||...||Protocol (8 bits, optional)|
```

- Length Extension (LXT, one bit) is 0 if the current byte is the last byte where the length of the packet is included, and 1 in other case.
- Length (LEN, 7, 14, 21, etc. bits): This is the length of the multiplexed packet (in bytes), not including the length field. If the length of the multiplexed packet is less than 128 bytes (less than or equal to 127 bytes), LXT is set to 0 and the 7 bits of the length field represent the length of the multiplexed packet. If the length of the multiplexed packet is greater than 127 bytes, additional bytes are added. The first bit of each of the added bytes is the LXT. If LXT is set to 1, it means that there is an additional byte for expressing the length. This allows to multiplex packets of any length (see the next figures).
- Protocol (8 bits) is the Protocol field of the multiplexed packet, according to IANA "Assigned Internet Protocol Numbers". It only appears in Non-first headers if the Single Protocol Bit (SPB) of the First Simplemux header is set to 1.

As an example, a Non-first Simplemux header before a packet smaller than 128 bytes, when the protocol bit has been set to 0 in the first header, would be 1 byte long:


```
0
0 1 2 3 4 5 6 7
+---+---+---+---+---+---+
|L|           |
|X|   Length   |
|T| (7 bits)   |
+---+---+---+---+---+
^
0
```

SPB = 0 in the first header

Figure 6

A Non-first Simplemux header before a packet with a length greater or equal to 128 bytes, and smaller than 16384 (2^{14}), when the


protocol bit has been set to 0 in the first header, will be 2 bytes long:



SPB = 0 in the first header

Figure 7

A Non-first Simplemux header before a packet with a length greater or equal to 16384 bytes, and smaller than 2097152 bytes (2^{21}), when the protocol bit has been set to 0 in the first header, will be 3 bytes long:




SPB = 0 in the first header

Figure 8

In this case, the length of the packet will be the number expressed by the concatenation of the bits of Length 1 - Length 2 - Length 3 (total 21 bits). Length 1 includes the 7 most significant bits and Length 3 the 7 less significant bits.

More bytes can be added to the length if required, using the same scheme: 1 LXT byte plus 7 bits for expressing the length.


A Non-first Simplemux header before a packet smaller than 128 bytes, when the protocol bit has been set to 1 in the first header, will be 2 bytes long:



SPB = 1 in the first header

Figure 9


A Non-first Simplemux header before a packet with a length greater or equal to 128 bytes, and smaller than 16384 (2^14), when the protocol bit has been set to 1 in the first header, will be 3 bytes long:



SPB = 1 in the first header

Figure 10

A Non-first Simplemux header before a packet with a length greater of equal to 16384 bytes, and smaller than 2097152 bytes (2^21), when the protocol bit has been set to 1 in the first header, will be 4 bytes long:



SPB = 1 in the first header

Figure 11

In this case, the length of the packet will be the number expressed by the concatenation of the bits of Length 1 - Length 2 - Length 3 (total 21 bits). Length 1 includes the 7 most significant bits and Length 3 the 7 less significant bits.

More bytes can be added to the length if required, using the same scheme: 1 LXT byte plus 7 bits for expressing the length.

These would be some examples of the whole bundles:

Case 1: All the packets belong to the same protocol: The first Simplemux header would be 2 or 3 bytes (for usual packet sizes), and the other Simplemux headers would be 1 or 2 bytes. For small packets (< 128 bytes), the Simplemux header would only require one byte.

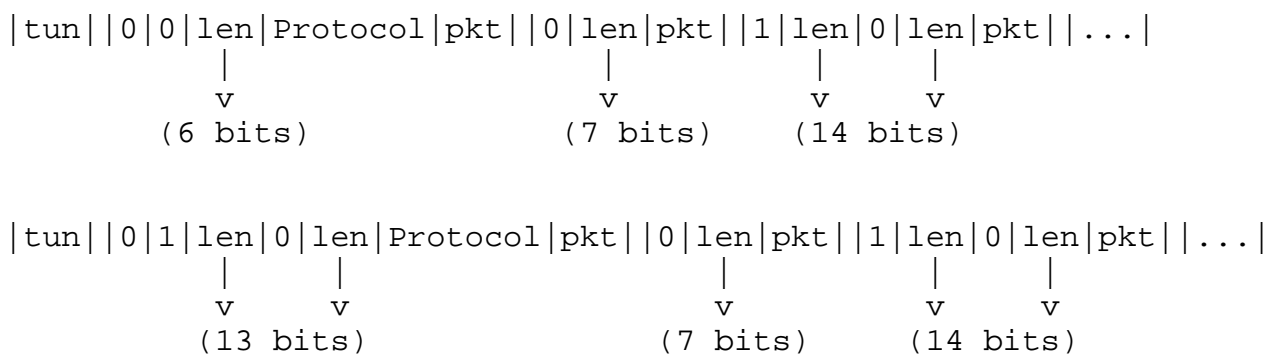


Figure 12

Case 2: Each packet may belong to a different protocol: All the Simplemux headers would be 2 or 3 bytes (for usual packet sizes).




Figure 13

4. Acknowledgements

Jose Saldana was funded by the EU H2020 Wi-5 project (Grant Agreement no: 644262).

5. IANA Considerations

A protocol number for Simplemux should be requested to IANA.

As a provisional solution for IP networks, the ingress and the egress optimizers may agree on a UDP port, and use IP/UDP as the multiplexing protocol.

6. Security Considerations

7. References

7.1. Normative References

- [RFC1570] Simpson, W., Ed., "PPP LCP Extensions", RFC 1570, DOI 10.17487/RFC1570, January 1994, <<http://www.rfc-editor.org/info/rfc1570>>.
- [RFC1692] Cameron, P., Crocker, D., Cohen, D., and J. Postel, "Transport Multiplexing Protocol (TMux)", RFC 1692, DOI 10.17487/RFC1692, August 1994, <<http://www.rfc-editor.org/info/rfc1692>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.

- [RFC3153] Pazhyannur, R., Ali, I., and C. Fox, "PPP Multiplexing", RFC 3153, DOI 10.17487/RFC3153, August 2001, <<http://www.rfc-editor.org/info/rfc3153>>.
- [RFC4170] Thompson, B., Koren, T., and D. Wing, "Tunneling Multiplexed Compressed RTP (TCRTP)", BCP 110, RFC 4170, DOI 10.17487/RFC4170, November 2005, <<http://www.rfc-editor.org/info/rfc4170>>.

7.2. Informative References

- [Bolla] Bolla, R., Bruschi, R., Davoli, F., and F. Cucchietti, "Energy Efficiency in the Future Internet: A Survey of Existing Approaches and Trends in Energy-Aware Fixed Network Infrastructures", IEEE Communications Surveys and Tutorials vol.13, no.2, pp.223,244, 2011.
- [Chabarek] Chabarek, J., Sommers, J., Barford, P., Estan, C., Tsiang, D., and S. Wright, "Power Awareness in Network Design and Routing", INFOCOM 2008. The 27th Conference on Computer Communications. IEEE pp.457,465, 2008.

Author's Address

Jose Saldana
University of Zaragoza
Dpt. IEC Ada Byron Building
Zaragoza 50018
Spain

Phone: +34 976 762 698
Email: jsaldana@unizar.es