

Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 5, 2012

Xiangyang Zhang
Tina Tsou
Futurewei Technologies, Inc
October 6, 2011

IPsec anti-replay algorithm without bit-shifting
draft-zhang-ipsecme-anti-replay-06

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), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Distribution of this memo is unlimited.

This Internet-Draft will expire on April 5, 2012.

Copyright Notice

Copyright (c) 2011 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.

Abstract

This document presents an alternate method to do the anti-replay checks and updates for IP Authentication Header (AH) and Encapsulating Security Protocol (ESP). The method defined in this document obviates the need for bit-shifting and it reduces the number of times anti-replay window is adjusted.

Table of Contents

- 1. Introduction 3
- 2. Description of new anti-replay algorithm 3
- 3. Example of new anti-replay algorithm 6
- 4. Acknowledgements 10
- 5. Security considerations 10
- 6. IANA Considerations 10
- 7. Normative References 10
- Author's Address 10

1. Introduction

IP Authentication Header (AH) [RFC4302] and IP Encapsulating Security Payload (ESP) [RFC4303] define an anti-replay service that employs a sliding window mechanism. The mechanism, when enabled by a receiver, uses an anti-replay window of size W . This window limits how far out of order a packet can be, relative to the packet with the highest sequence number that has been authenticated so far. The window can be represented by a range $[WB, WT]$, where $WB=WT-W+1$. The whole anti-replay window can be thought of as a string of bits. The value of each bit indicates whether or not a packet with that sequence number has been received and authenticated, so that replay packet can be detected and rejected. If the packet is received, the receiver gets the sequence number S in the packet. If S is inside window ($S \leq WT$ and $S \geq WB$), then checks the corresponding bit (location is $S-WB$) in the window to see if this S has already been seen. If $S < WB$, the packet is dropped. If $S > WT$ and is validated, the window is advanced by $(S-WT)$ bits. The new window becomes $[WB+S-WT, S]$. The new bits in this new window are set to indicate that no packets with those sequence numbers have been received. The typical implementation (for example, [RFC4302] algorithm) is done by shifting $(S-WT)$ bits. In normal cases, the packets arrive in order, which results in constant update and bit shifting operation.

[RFC4302][RFC4303] defined minimum window sizes of 32 and 64. But no requirement is established for minimum or recommended window sizes beyond 64-packet. The window size needs to be based on reasonable expectations for packet re-ordering. For a high-end multi-core network processor with multiple crypto cores, a window size bigger than 64 or 128 is needed due to the varied IPsec processing latency caused by different cores. In such a case, the window sliding is tremendous costly even with hardware acceleration to do the bit shifting. This draft describes an alternate method to avoid bit-shifting. It only discusses the anti-replay processing at the receiving side. The processing is always safe and has no interoperability effects. Even with the window size beyond the usual 32 or 64 bit window, it does not cause any interoperability issue.

2 Description of new anti-replay algorithm

Here we present an easy way to only update the window index and also reduce the times of updating the window. The basic idea is illustrated in the following figures. Suppose that we configure the window size W , which consists of $M-1$ blocks, where M is power of two (2). Each block

contains N bits, where N is also power of two (2). It can be a byte (8 bit) or word (32bit), or multiple words. The supported sliding window size is (M-1)*N. However, it covers up M blocks (four blocks as shown in Figure 1). All these M blocks are circulated and become a ring of blocks, each with N bits. In this way, the supported sliding window (M-1 blocks) is always a subset window of the actual window when window slides.

Initially the actual window is defined by low and high end index [WB, WT], as illustrated in Figure 1.




Figure 1: the sliding window [WB, WT], in which WT is last validated sequence number and the supported window size W is WT-WB+1. (x=don't care bit, c=check bit)

If we receive a packet with the sequence number (S) greater than WT, we slide the window. But we only change the window index by adding the difference (S-WT) to both WT (WB is automatically changed as window size is fixed). So S becomes the largest sequence number of the received packets. Figure 2 shows the case that the packet with sequence number S=WT+1 is received.




Figure 2: the sliding window [WB, WT] after S=WT+1

If S is in the different block from where WT is, we have to initialize all bit values in the blocks to 0 without bit shifting. If S passes several blocks, we have to initialize several blocks instead of only

one block. Figure 3 shows that the sequence number already pass the block boundary. Immediately after update, all the check bits should be 0 in the block where WT resides.




Figure 3: the sliding window [WB, WT] after S pass the boundary

After update, the new window still covers the configured window. This means the configured sub-window also slides, conforming to the sliding window protocol. The actual effect is somewhat like shifting the block. In this way, the bit-shifting is deemed unnecessary.

It is also easier and much faster to check the window with the sequence number because the sequence number check does not depend on the lowest index WB. Instead, it only depends on the sequence number of the received packet. If we receive a sequence number S, the bit location is the lowest several bits of the sequence number, which only depends on the block size (N). The block index is several bits before the location bits, which only depends on the window size (M).

We do not specify how many redundancy bits needed except that it should be power of two (2) for computation efficiency. If microprocessor is 32 bit, 32 might be a better choice while 64 might be better for 64 bit microprocessor. For microprocessor with cache support, one cache line is also a good choice. It also depends on how big the sliding window size is. If we have N redundancy bits (for example, 32 bit in the above description), we only need 1/N times update of blocks, comparing to the bit-shifting algorithm in [RFC4302].

The cost of this method is extra byte or bytes used as redundant window. The cost will be minimal if the window size is big enough. Actually the extra redundant bits are not completely wasted. We could reuse the unused bits in the block where index WB resides, i.e. the supported window size could be (M-1)*N, plus the unused bits in the last block.

3 Example of new anti-replay algorithm

Here is the example code to implement the algorithm of anti-replay check and update, which is described in the previous sections.

<CODE BEGINS>

```
/**
 * Copyright (c) 2011 IETF Trust, Xiangyang Zhang and Tina Tsou.
 * All rights reserved.
 *
 * Redistribution and use in source and binary forms, with or without
 * modification, is permitted pursuant to, and subject to the license
 * terms contained in, the Simplified BSD License set forth in Section
 * 4.c of the IETF Trust's Legal Provisions Relating to IETF Documents
 * (http://trustee.ietf.org/license-info).
 *
 */

/**
 * In this algorithm, the hidden window size must be a power of two,
 * for example, 1024 bits. The redundant bits must also be a power of
 * two, for example 32 bits. Thus, the supported anti-replay window
 * size is the hidden window size minus the redundant bits. It is 992
 * in this example. The size of integer depends on microprocessor
 * architecture. In this example, we assume that the software runs on
 * 32 bit microprocessor. So the size of integer is 32. In order to
 * convert the bitmap into an array of integer, the total number of
 * integers is the hidden window size divided by size of integer.
 *
 * struct ipsec_sa contains the window and window related parameters,
 * such as the window size, the last acknowledged sequence number.
 *
 * all the value of macro can be changed, but must follow the rule
 * defined in the algorithm.
 */
```

```
#define SIZE_OF_INTEGER      32 /** 32 bit microprocessor */
#define BITMAP_LEN          (1024/ SIZE_OF_INTEGER)
                          /** in terms of 32 bit integer */
#define BITMAP_INDEX_MASK  (IPSEC_BITMAP_LEN-1)
#define REDUNDANT_BIT_SHIFTS  5
#define REDUNDANT_BITS      (1<<REDUNDANT_BIT_SHIFTS)
#define BITMAP_LOC_MASK    (IPSEC_REDUNDANT_BITS-1)

int
ipsec_check_replay_window (struct ipsec_sa *ipsa,
                          uint32_t sequence_number)
{
    int bit_location;
    int index;

    /**
     * replay shut off
     */
    if (ipsa->replaywin_size == 0) {
        return 1;
    }

    /**
     * first == 0 or wrapped
     */
    if (sequence_number == 0) {
        return 0;
    }

    /**
     * first check if the sequence number is in the range
     */
    if (sequence_number>ipsa->replaywin_lastseq) {
        return 1; /** larger is always good */
    }

    /**
     * The packet is too old and out of the window
     */
    if ((sequence_number + ipsa->replaywin_size) <
        ipsa->replaywin_lastseq) {
        return 0;
    }
}
```

```
/**
 * The sequence is inside the sliding window
 * now check the bit in the bitmap
 * bit location only depends on sequence number
 */
bit_location = sequence_number&BITMAP_LOC_MASK;
index = (sequence_number>>REDUNDANT_BIT_SHIFTS)&BITMAP_INDEX_MASK;

/*
 * this packet already seen
 */
if (ippsa->replaywin_bitmap[index]&(1<<bit_location)) {
    return 0;
}

return 1;
}

int
ipsec_update_replay_window (struct ipsec_sa *ippsa,
                           uint32_t sequence_number)
{
    int bit_location;
    int index, index_cur, id;
    int diff;

    if (ippsa->replaywin_size == 0) { /** replay shut off */
        return 1;
    }

    if (sequence_number == 0) {
        return 0; /** first == 0 or wrapped */
    }

    /**
     * the packet is too old, no need to update
     */
    if ((ippsa->replaywin_size + sequence_number) <
        ippsa->replaywin_lastseq) {
        return 0;
    }
}
```



```
/**
 * now update the bit
 */
index = (sequence_number>>REDUNDANT_BIT_SHIFTS);

/**
 * first check if the sequence number is in the range
 */
if (sequence_number>ipsa->replaywin_lastseq) {
    index_cur = ipsa->replaywin_lastseq>>REDUNDANT_BIT_SHIFTS;
    diff = index - index_cur;
    if (diff > BITMAP_LEN) { /* something unusual in this case */
        diff = BITMAP_LEN;
    }

    for (id = 0; id < diff; ++id) {
        ipsa->replaywin_bitmap[(id+index_cur+1)&BITMAP_INDEX_MASK]
            = 0;
    }

    ipsa->replaywin_lastseq = sequence_number;
}

index &= BITMAP_INDEX_MASK;
bit_location = sequence_number&BITMAP_LOC_MASK;

/* this packet already seen */
if (ipsa->replaywin_bitmap[index]&(1<<bit_location)) {
    return 0;
}

ipsa->replaywin_bitmap[index] |= (1<<bit_location);

return 1;
}

<CODE ENDS>
```

4. Acknowledgements

The idea in this document came from the software design on one high-performance multi-core network processor. The new network processor core integrates a dozen of crypto core in distributed way, which makes hardware anti-replay service impossible.

5. Security considerations

This document does not change [RFC4302] or [RFC4303]. It provides an alternate method for anti-replay.

6. IANA Considerations

None.

7. Normative References

- [RFC4302] "IP Authentication Header", RFC 4302.
- [RFC4303] "IP Encapsulating Security Payload (ESP)", RFC 4303.

Author's address

Xiangyang Zhang
Futurewei Technologies
2330 Central Expressway
Santa Clara, California 95051
USA

Phone: +1-408-330-4545
Email: xiangyang.zhang@huawei.com

Tina TSOU (Ting ZOU)
Futurewei Technologies
2330 Central Expressway
Santa Clara, California 95051
USA

Phone: +1-408-859-4996
Email: tena@huawei.com