IP Security Maintenance and Extensions (IPsecME) WG

IETF 114, Monday, July 25th, 2022

Chairs: Tero Kivinen
Yoav Nir

Responsible AD: Roman Danyliw

Note Well

This is a reminder of IETF policies in effect on various topics such as patents or code of conduct. It is only meant to point you in the right direction. Exceptions may apply. The IETF's patent policy and the definition of an IETF "contribution" and "participation" are set forth in BCP 79; please read it carefully.

As a reminder:

• By participating in the IETF, you agree to follow IETF processes and policies.

• If you are aware that any IETF contribution is covered by patents or patent applications that are owned or controlled by you or your sponsor, you must disclose that fact, or not participate in the discussion.

• As a participant in or attendee to any IETF activity you acknowledge that written, audio, video, and photographic records of meetings may be made public.

• Personal information that you provide to IETF will be handled in accordance with the IETF Privacy Statement.

• As a participant or attendee, you agree to work respectfully with other participants; please contact the ombudsteam (https://www.ietf.org/contact/ombudsteam/) if you have questions or concerns about this.

Definitive information is in the documents listed below and other IETF BCPs. For advice, please talk to WG chairs or ADs:
•BCP 9 (Internet Standards Process)
•BCP 25 (Working Group processes)
•BCP 25 (Anti-Harassment Procedures)
•BCP 54 (Code of Conduct)
•BCP 78 (Copyright)
•BCP 79 (Patents, Participation)
•https://www.ietf.org/privacy-policy/ (Privacy Policy)

Administrative Tasks

We need volunteers to be:

- Two note takers

Jabber: xmpp:ipsecme@jabber.ietf.org?join

MeetEcho: https://meetings.conf.meetecho.com/ietf114/?group=ipsecme&short=&item=1

Notes: https://notes.ietf.org/notes-ietf-114-ipsecme

Agenda

- Note Well, technical difficulties and agenda bashing – Chairs (5 min) (15:00-15:05)
- Document Status – Chairs (15 min) (15:05-15:20)
- Work items
  - IKEv2 Downstream Fragmentation Notification Extension – Daniel Migault (20 min) (15:20-15:40)
  - IKEv2 Count Based SA Extension – Daniel Migault (20 min) (15:40-16:00)
  - IPsec Multi SA – Paul Wouters (10 min) (16:00-16:10)
  - Charter Update and Status – Chairs (10 min) (16:10-16:20)
- AOB + Open Mic (60 min) (16:20-17:00)

WG Status Report

- Published as RFC
  - draft-ietf-ipsecme-ikev2-intermediate published as RFC9242
- Publication requested:
  - draft-ietf-ipsecme-iptfs IESG Eval
  - draft-ietf-ipsecme-rfc8229bis IESG Eval
  - draft-ietf-ipsecme-yang-iptfs IETF Last Call
  - draft-ietf-ipsecme-ikev1-algo-to-historic AD eval
  - draft-ietf-ipsecme-mib-iptfs AD eval
  - draft-ietf-ipsecme-ikev2-multiple-ke Pub req
- Waiting for write-up / Chair review:
  - draft-ietf-ipsecme-labeled-ipsec
- Work in progress:
  - draft-ietf-ipsecme-auth-announce
  - draft-ietf-ipsecme-g-ikev2
  - draft-ietf-ipsecme-add-ike

Presentations

- IKEv2 Downstream Fragmentation Notification Extension –
  Daniel Migault
- IKEv2 Count Based SA Extension –
  Daniel Migault
- IPsec Multi SA –
  Paul Wouters

Presentations

- **IKEv2 Downstream Fragmentation Notification Extension – Daniel Migault**
- IKEv2 Count Based SA Extension – Daniel Migault
- IPsec Multi SA – Paul Wouters

IKEv2 IPv4 Downstream Fragmentation Notification Extension

Liu, Zhang. Migault

Problem Statement (Fragmentation)

Reassembling packets by a downstream security gateway:

- requires additional resources which under heavy load results in service degradations.
- the 16-bit IPv4 identification field that is not large enough to prevent duplication making fragmentation not sufficiently robust at high data rates.

The problem is only for IPv4

When the DF bit is set to 1, the security gateway is unlikely to receive the unencrypted ICMPv4 message Packet Too Big (PTB).

- this results in black holing

As a result DF is set to 0 and the security gateway does the reassembly operation.

We propose that the downstream security gateway advertises the other security gateway that fragmentation is ongoing with an indication of the acceptable MTU.

- similar to PTB and RFC8900 recommends each layer handles fragmentation at their layer and to reduce the reliance on IP fragmentation to the greatest degree possible

```
Initiator                              Responder
-------------------------------------------------------------------
HDR, SA, KEi, Ni -->
                            <-- HDR, SA, KEr, Nr
HDR, SK {IDi, AUTH,
     SA, TSi, TSr,
     N(IP4_DOWNSTREAM_FRAGMENTATION_SUPPORTED)} -->
                            <-- HDR, SK {IDr, AUTH,
                                SA, TSi, TSr,
                                N(IP4_DOWNSTREAM_FRAGMENTATION_SUPPORTED)}

Receiving Security Gateway                    Sending Security Gateway
-------------------------------------------------------------------
HDR SK { N(IP4_DOWNSTREAM_FRAGMENTATION)} -->
```

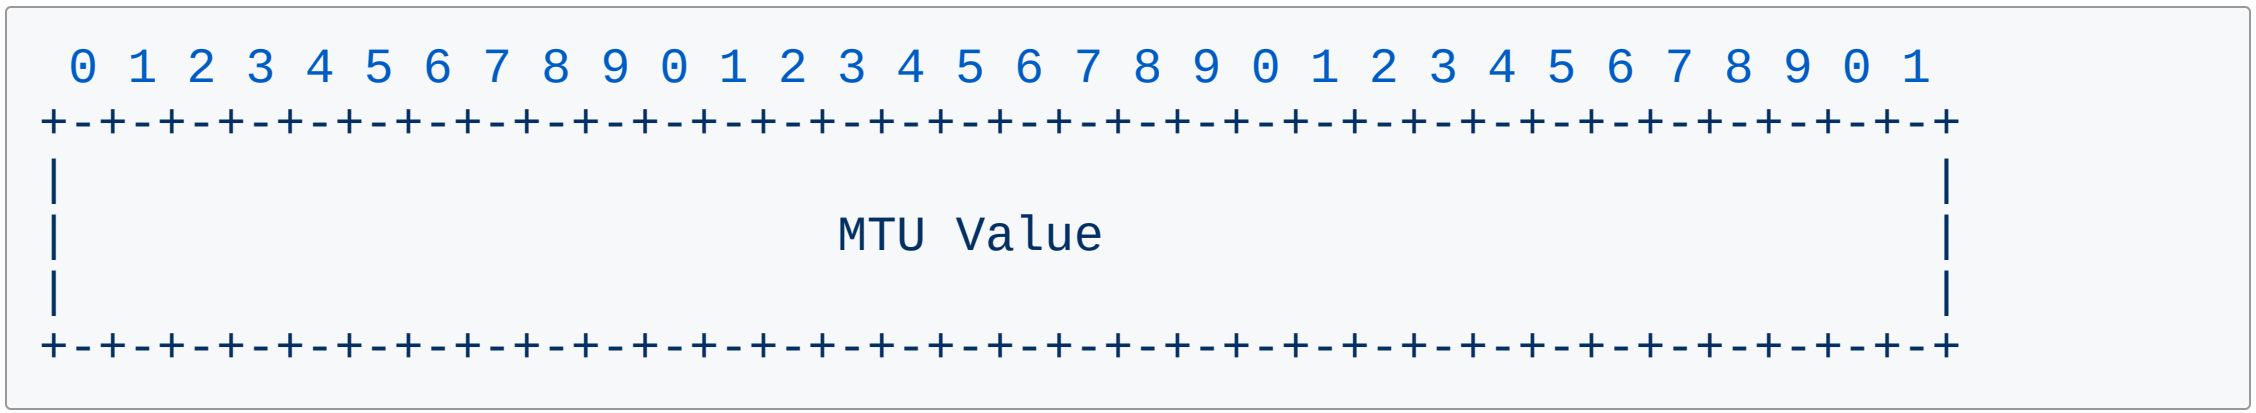
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
MTU Value

Upon receiving a IP4_DOWNSTREAM_FRAGMENTATION:

1. The sending security gateway SHOULD advertise sending nodes of the MTU expected for the inner packet and discard those that are larger.
2. If inner packet has its DF bit set to 0, the security gateway MAY perform inner fragmentation
3. The security gateway MAY perform outer fragmentation.

- DF is then set to 1 and communication may be exposed to blackholing

Thanks!

Presentations

- IKEv2 Downstream Fragmentation Notification Extension –
  Daniel Migault
- **IKEv2 Count Based SA Extension –**
  Daniel Migault
- IPsec Multi SA –
  Paul Wouters

IKEv2 Count Based SA Extension

Migault, Liu, Zhang

Problem Statement

Hardware accelerated IPsec:

- Are designed for a fix number of SAs
- SAs that cannot be created result in traffic being rejected.

Simultaneous IKEv2 rekeys result in the creation of redundant SAs

- underutilisation of the hardware component

SAs life time can be expressed using a time or a byte count limit.

Time limit:

- makes limit predictable (over time)
  - it is easy to anticipate the expiration time
- Uniformly randomizing time limit distributes the IKEv2 rekey uniformly
  - and works pretty well

byte count limit:

- are hard to predict - as it depends on the traffic.
  - our implementation checks every 2s which SAs needs to be rekeys
- the randomization does not compensate the traffic bursts

As a result, count bases SA results in multiple redundant rekey

Why we want to use byte count limit ?

- a direct expression of the lifetime of the cryptographic key
- binding the SA life time to traffic is appropriated for device that can be in long sleeping mode.

The document describes an IKEv2 extension the prevents redundant rekey:

- no magic: peers agree who is expected to start the rekey

```
Initiator                                Responder
-----------------------------------------------
HDR, SAi1, KEi, Ni  -->
                          <--  HDR, SAr1, KEr, Nr, [CERTREQ]
HDR, SK {IDi, [CERT,] [CERTREQ,]
    [IDr,] AUTH, SAi2,
    TSi, TSr, N(COUNT_BASED_SA_PROPOSED)}  -->
                          <--  HDR, SK {IDr, [CERT,] AUTH,
                             SAr2, TSi, TSr, N(COUNT_BASED_SA_SELECTED)}
```

COUNT_BASED_SA_PROPOSED Notification Data contains for each Transform ID:

- Acceptable count base life time
- Rekey Value (random)

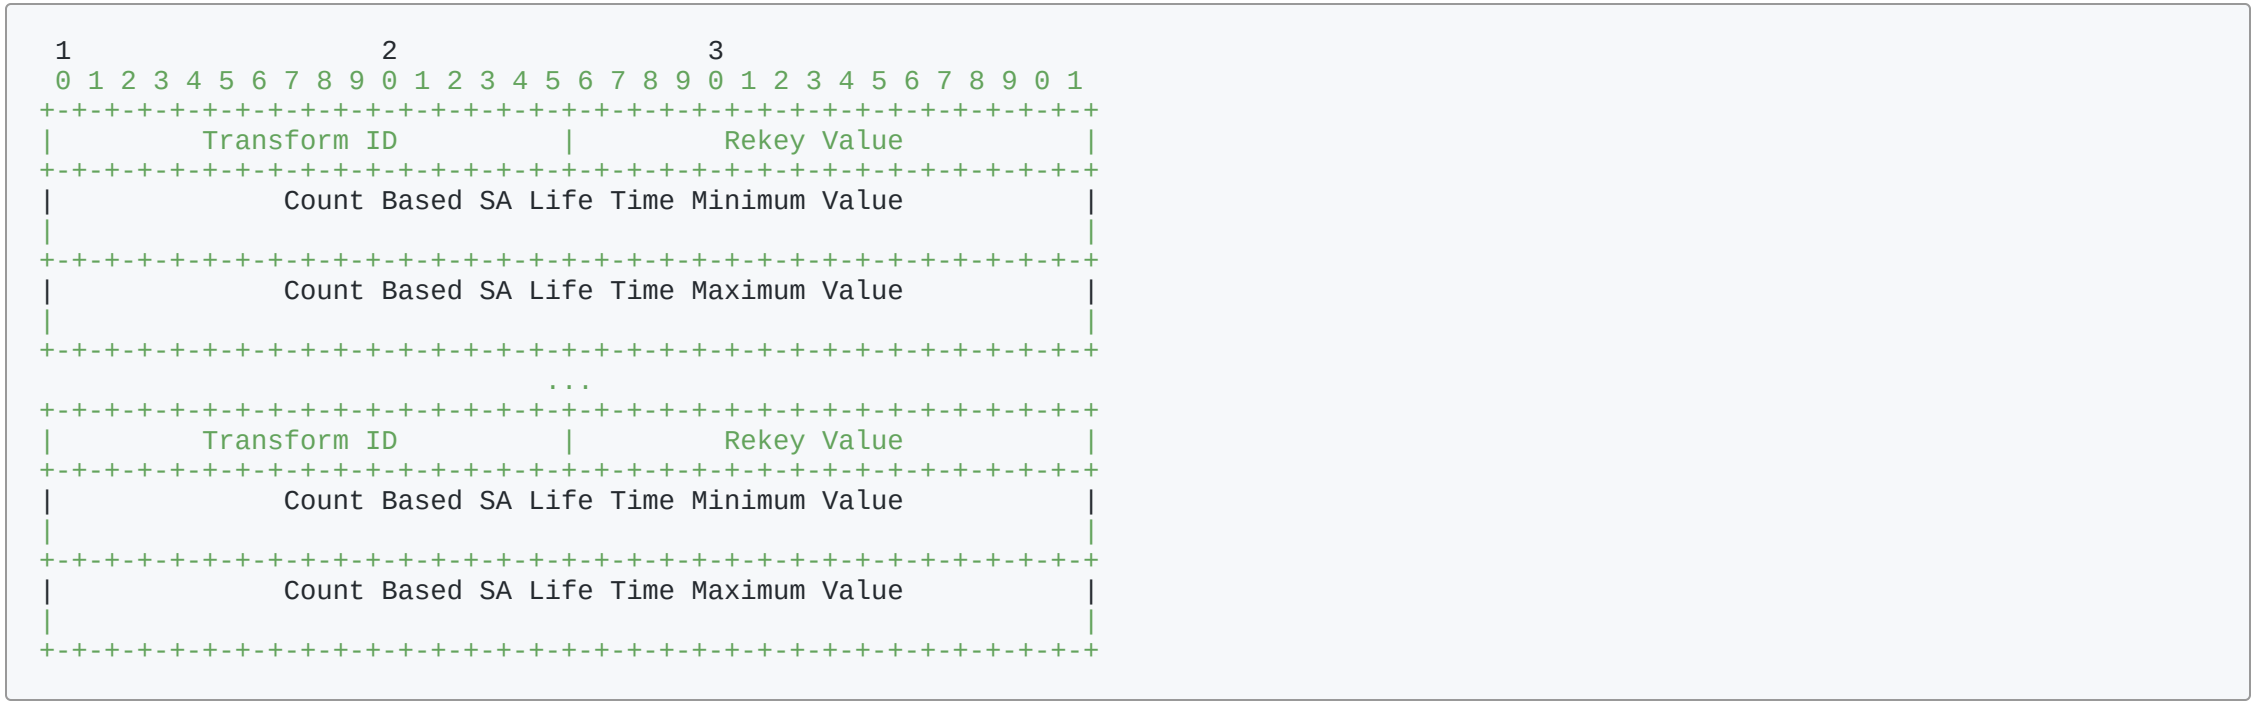

COUNT_BASED_SA_SELECTED Notification Data

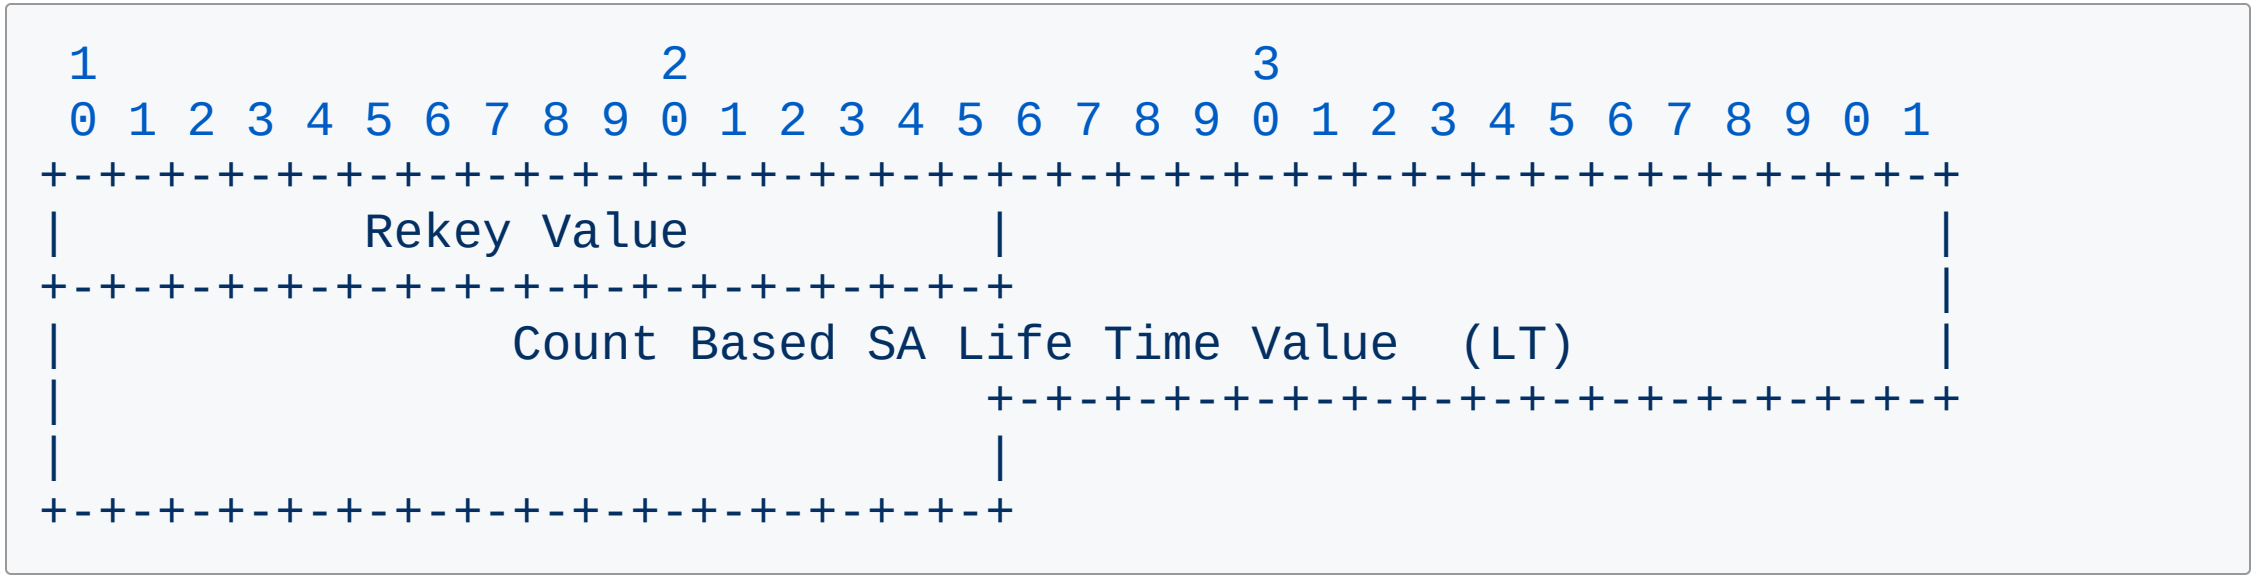

1. Selection of the initiator for the next rekey:

- The peer with the greatest Rekey Value is designated to initiate the next rekey.
- In case of equality, the current initiator remains the initiator.

2. setting the Hard (H) and Soft (S) count base SA lifetime:

- initiator:
  - S = X_i * LT + rand(0, 5% LT) with X_i <= 80%
  - H = LT
- responder :
  - S = X_r LT + rand(0, 5% LT) with X_r >= 95%
  - H = LT

Thanks!

Presentations

- IKEv2 Downstream Fragmentation Notification Extension –
  Daniel Migault
- IKEv2 Count Based SA Extension –
  Daniel Migault
- **IPsec Multi SA –**
  Paul Wouters

draft-pwouters-ipsecme-multi-sa-performance

Goal of the draft

- Increase performance of IPsec
  - Setup regular Child SA, signal support in IKE via NOTIFY
  - Additionally, trigger ACQUIREs per CPU
    - Negotiate additional Child SAs with same Traffic Selectors with CPU id in Notify for “pinning”
- Multiple identical Child SAs already possible, but implementations often see these as replacement SAs.

Status

- 00 draft included support for multi CPUs and QoS
- 01 draft removed QoS support, removed SA count negotiation
  - No further changes in 1.5 years – stable code

- Implemented IPsec code for Linux kernel
- Implemented IKEv2 code in strongSwan and Libreswan

Performance numbers

- Using bare metal with Intel i40e cards
- In forwarding mode (gateway to gateway, not host to host)
- Linux 5.x using XFRM

- Using traditional IPsec: 3.5 gbps
- Using 8 per-CPU Child SAs: 27 gbps

Need standardization for per-CPU capability

- Please (re)read draft-pwouters-multi-sa-performance
- If you strongly object to assigning two NOTIFY payloads, please say so on the ipsec list.
- If you like performance, ask chairs for WG adoption call
- Linux kernel code is waiting on this to merge in per-cpu code.
- strongSwan and Libreswan waiting on kernel code before merging

Charter Update and Status

- Items finished (given to IESG):
  - PPK – RFC8784
  - Implicit IV – RFC8750
  - Post Quantum – RFC 9242 + Multiple ke
  - RFC8229bis
  - TFC + Yang + Mib for it
- Items still in charter:
  - G-doi
  - Constrained ESP
  - Signature algorithm negotiation
  - Security Labels

Charter Other Items

- Work items not in charter, but agreed as WG items:
  - ADD – IKEv2 config for Encrypted DNS
  - IKEv1 to historic
- Items not in charter, but being worked on:
  - draft-pwouters-ipsecme-multi-sa-performance
  - draft-liu-ipsecme-mtu-detect
  - draft-liu-ipsecme-ikev2-rekey-redundant-sas
- Old items, nothing happening lately:
  - draft-kampati-ipsecme-ikev2-sa-ts-payloads-opt
  - draft-mglt-ipsecme-diet-esp
  - draft-smyslov-ipsecme-ikev2-cookie-revised
  - draft-tjhai-ikev2-beyond-64k-limit

Open Discussion

- Other points of interest?