Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: July 18, 2008

S. Decugis
University of Tokyo

January 15, 2008

Issues of the Key Management Mobility Capability (K) flag in Mobile IPv6.
draft-sdecugis-mext-kbit-issues-00

Status of this Memo

Copyright Notice

Abstract

Mobile IPv6 specification (RFC 3775) introduces a flag called "Key Management Mobility Capability (K)" that is used during home registration procedure. This document describes the behavior of implementations when the capability is supported or not. The purpose of this description is to highlight the expected behavior of

implementations at the end of home registration process with regards to the exchanged (K) flag value.

Table of Contents

1. Introduction

This introduction briefly presents the use of dynamic keying with Mobile IPv6, and the exact meaning of the 'K' flag in BU/BA messages. Previous knowledge of the Mobile IPv6 [RFC3775], IPsec [RFC4301], and IKEv2 [RFC4306] [RFC4877] RFCs is highly recommended.

This document mainly deals with IKEv2 protocol, but it applies to IKE also with very few differences.

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. Mobile IPv6

Mobile IPv6 provides the ability for a node (called "Mobile Node", MN) to be reachable at a permanent IPv6 address (its "Home Address", HoA) while its point of attachment to the network is changing (the temporary IPv6 address is called the "Care-of Address", CoA). One of the routers serving the prefix of the HoA has a special role in the mobility (this router is called the "Home Agent", HA.) Basically, the MN registers its CoA to its HA, then the traffic directed to the HoA is tunneled to the MN at its CoA. The correspondents can always use the HoA to reach the MN.

1.3. Dynamic keying

Mobile IPv6 also makes mandatory the use of IPsec to protect the messages (Binding Update, Binding Acknowledgement -- BU / BA) involved during the MN's registration with its HA. Dynamic keying represents the situation where a Key Management Protocol (such as IKE or IKEv2) is used to negotiate the Security Association ("SA") pair that will protect the BU/BA exchange and optionally other messages.

We can distinguish two steps in the KMP exchanges. First, the peers negotiate a secured channel ("IKE_SA"), based on some trust mechanism (shared key, certificates, ...). Then, using this secured channel, the peers can negotiate the SA parameters for protecting the BU/BA and optionally other SA parameters for other kinds of traffic. Since the BU/BA exchange has not yet been completed when the SA are negotiated, the HoA cannot be used, and therefore the CoA is used for the negotiation of the IKE_SA on the MN side. The child SA (the IPsec SA that protects the BU/BA) is established between the HoA and the HA.

1.4. Movement

When the MN changes its CoA, the IPsec rules are updated to reflect this change. For the transport-mode SA, there is no update needed since the SA are established between the HoA and the HA. The tunnel-mode SA and SP entries MUST be updated, as required by RFC 3775. The IKE_SA (used only to protect IKE messages) endpoints may or may not be updated. The "K flag" as discussed in this document is meant to represent the ability of the peer to update the IKE_SA endpoints.

2. Implementation dependent behavior

In order to describe properly the implementation behavior, we have to make some assumptions on the reaction to some events, when this reaction is not clearly defined by a RFC.

2.1. Local address is not usable

When an IKE session is established, the IKE module stores both the local and remote IPv6 addresses of this session. In Mobile IPv6, it may happen that the local address is not available anymore. In this case, if a new IKE message is to be sent using this session, the implementation may have two different behaviors, depending if a fallback mechanism is available or not.

- It may just fail to send the message, and consider the IKE session as dead. It means that the IKE session is removed, as well as dependent SA. Next time an ACQUIRE message is received, a new IKE

session will be established. How the source address will be determined is outside the scope of this document. This behavior will be later referenced as IMPL_1A.

- It may otherwise try to use another available address as the source of the message. Note that the HoA MUST NOT be used in this case. This behavior will be later referenced as IMPL_1B. Note that in the case where the IKE module is able to determine the correct CoA to use, then we can assume that the module has the Key Management Mobility Capability.

2.2. Unexpected source address

Once an IKE session is established, the IKE module may receive an IKE message with a source address different from the address that was stored in the IKE session. Once again we may have different implementation behaviors.

- The IKE module may choose to simply discard the message. Even though there is no assumption in the RFC about the source address of the message, nothing forbids this behavior either. It can be used as a protection against some DoS attacks, for example. This behavior will be referenced as IMPL_2A.

- More likely, the IKE module will process the message. In this case, according to IKEv2 specification, the reply MUST be sent to the source address of the message. This new address additionally may or may not be stored in the IKE session, to be used next time a message has to be sent to the remote peer. This behavior will be referenced as IMPL_2B. Note that since the outer IP source address is not protected, it is not recommended to store this address.

3. Description of movement and consequences

For the remaining of this document, we are considering the following initial situation:

MN is on a Foreign Link 1 (with CoA1) and registered to its HA.

The following SA are established:

* IKE_SA established between CoA1 and HA

* SA1 (transport) between HoA and HA.

* SA2 (tunnel) between CoA1 and HA.

Then MN moves to Foreign Link 2 with CoA2. The MIPv6 module on MN detects the movement and sends a new BU to its HA, protected by SA1. HA replies a BA. SA2 is updated on MN and on HA, and now established between CoA2 and HA.

Now we will consider the following situations. "k=0" means that the change of IP address of the MN is not notified to the IKE session in IKE module. "k=1" means that the IKE session is updated with the new CoA.

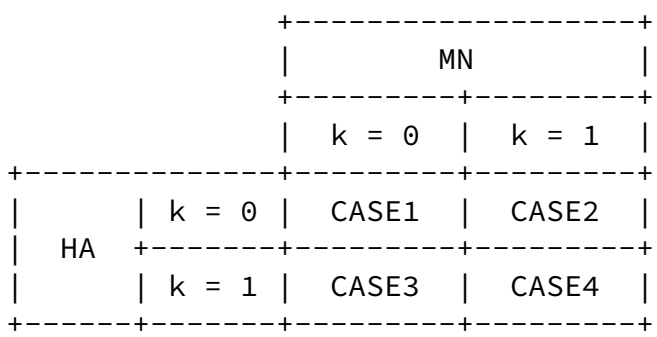

Matrix of situations.

Figure 1

3.1. CASE1: No support for movement.

In this case, neither MN nor HA update their IKE session after movement.

3.1.1. MN sends next IKE message

We are now considering the case where MN needs to send the next IKE message. This is more likely to happen, since MN is the initiator of IKE session.

- If MN behaves as described in IMPL_1A, then the IKE session is destroyed on MN side, with consequence that the SA (SA1 and SA2) are also destroyed (immediately in IKEv2, after expiry in IKE). In this case, the upper-layer traffic that was protected by SA2 is interrupted. HA is not notified, and will not be able to send anything to the MN until a new IKE session is established, and SA1 and SA2 renegotiated. This will happen on next ACQUIRE message on MN (provided that the IKE module is able to pickup the proper IP address.)

- If MN behaves as described in IMPL_1B, then the situation is the same as in CASE2, as long as the packet reaches the HA. Otherwise, for example if the chosen IP address belongs to a private network, then the IKE module will have to timeout before cleaning the IKE session and dependent SA.

3.1.2. HA sends next IKE message

It may also happen that HA needs to send an IKE message to the MN (to check for its liveness for example.) In that case, since the message is sent to the old CoA, no answer will be received and the IKE session will be deleted after a timeout (and the SA in the case of IKEv2). Since MN is not notified of the change, it may continue to

send encrypted packets which will be lost. This situation will last until the MN initiates a new IKE exchange -- which could take a long time, during which the upper-layer flow is interrupted.

3.2. CASE2: Only MN's IKE session is updated.

3.2.1. MN sends next IKE message

In this situation, the MN is able to send the correct message to HA. If HA behaves as described in IMPL_2A, the IKE message is ignored, resulting in the IKE session being destroyed after a timeout on MN. The HA will not be able to send anything to the MN until MN initiates a new IKE session. It's the same as in CASE1.

If the HA behaves as described in IMPL_2B, the IKE message is received correctly and the IKE session survives. Depending on whether the new MN CoA is saved in the IKE session on HA's side, the situation described here bellow when HA needs to send a message can occur later or otherwise we are in the same situation as in CASE4 (no problem).

3.2.2. HA sends next IKE message

Since the HA IKE session does not have the new CoA of the MN, the IKE message sent by HA receives no reply. This will result in HA detecting the session is broken after a timeout, and the session being removed. It's the same situation as in CASE1, the interruption in upper-layer flow will last until the MN has to send a new IKE message, detect the problem, and establish a new IKE session and dependent SA, which can last very long time.

3.3. CASE3: Only HA's IKE session is updated.

3.3.1. MN sends next IKE message

As in CASE1, we have two possibilities for MN. In case it behaves as described in IMPL_1A, we have the same consequences, the IKE session is destroyed on MN, and child SAs also. The HA is not notified. In case the MN behaves as described in IMPL_1B, the IKE message is sent using another IP source address (but not the HoA). In this case, if the CoA is selected, all goes well as described in CASE4. If another address is selected, the packet may or may not reach the HA, or be rejected by the HA (IMPL_2A or IMPL_2B.) If the HA accepts the packet, then we are in the same situation as CASE4 (all OK, no interruption in upper-layer flow.) Otherwise, the MN will let the session timeout and discard the SA. The upper-layer flow is interrupted until a new IKE session is negotiated.

3.3.2. HA sends the next IKE message

Here the behavior depends on if the MN's IKE module behaves as described in IMPL_2A or IMPL_2B. In the first case, then the HA will not receive a reply, and delete the IKE session after a timeout. The upper-layer flow is broken until MN has to send a new IKE message -- very long time. In the other case (MN behaves as described in IMPL_2B), HA receives a reply and the IKE session is not broken. In that case, the upper-layer flow is not interrupted.

3.4. CASE4: Both peers update the IKE session

In this case, the movement is transparent to IKE modules. Note that there may be issues in the timing of events, in the case where the movement occurs during an IKE exchange, but it is outside the scope of this document.

4. Conclusion

As we have seen previously, only the situation where both MN and HA update their IKE session with the new Care-of Address will result in smooth operation. If the HA does not update the IKE session endpoints, there is a possibility to be faced with a situation where the upper-layer traffic is interrupted and it will be restored only when MN needs to send a new IKE message, which means very long interruption. If the HA updates its IKE session, but not the MN, the result is better but depending on MN's IKE implementation.

The conclusion to this analysis are as follow:

- To ensure a smooth operation, we need to have both peers support the migration of their IKE session endpoints.

- In other cases, detected thanks to the (K) flag in BU / BA exchanges, then we can avoid very long interruption in upper-layer flow by forcing the deletion of the IKE session and all child SAs after each movement. This avoids the case where HA got rid of the SA and waits for the MN to detect something is going wrong.

In addition, the (K) flag is a feature negotiated by the MIPv6 module, but representing a capability of the IKE module. It should therefore be negotiated locally between the IKE module and the MIPv6 module.

As a final note, the recommendation is that the support for migrating the IKE session should be made mandatory, in which case the (K) flag should be deprecated. In case it is not possible, then at least it

should be requested that the IKE session is destroyed immediately after the movement is completed, in the case where at least one host does not support IKE session movement.

5. Contributors

Thanks to Arnaud Ebalard, Shinta Sugimoto and Francis Dupont for their reviews and comments.

6. IANA Considerations

This memo includes no request to IANA.

7. Security Considerations

This is not relevant for this draft.

8. References

8.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

8.2. Informative References

[I-D.narten-iana-considerations-rfc2434bis] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", draft-narten-iana-considerations-rfc2434bis-08 (work in progress), October 2007.

[RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", RFC 2629, June 1999.

[RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, July 2003.

[RFC3775] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.

[RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.

[RFC4306] Kaufman, C., "Internet Key Exchange (IKEv2) Protocol", RFC 4306, December 2005.

[RFC4877] Devarapalli, V. and F. Dupont, "Mobile IPv6 Operation with IKEv2 and the Revised IPsec Architecture", RFC 4877, April 2007.

Author's Address

Sebastien Decugis
University of Tokyo
Keio University Murai Lab., 144-8 Ogura, Saiwai-ku
Kawasaki, Kanagawa 212-0054
JP

Phone: +81 44 580 1600
Email: sdecugis@hongo.wide.ad.jp

Full Copyright Statement

Intellectual Property

specification can be obtained from the IETF on-line IPR repository at http://www.ietf.org/ipr.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).