sysmocom

sysmocom - s.f.m.c. GmbH

Specification for IMSI Pseudonymization on the Radio Interface for 2G/3G/4G

by Oliver Smith, Neels Hofmeyr, and Harald Welte

HISTORY

NUMBER	DATE	DESCRIPTION	NAME
1	2020-04-03	Initial version.	OS

Contents

1 Introduction

1.1 Protecting the IMSI on the Radio Interface is Desirable

A long-standing issue in the 3GPP specifications for cellular mobile communications starting from 2G (GSM) is, that mobile phones and other mobile equipment (ME) have to send the International Mobile Subscriber Identity (IMSI) unencrypted over the air. Each IMSI is a unique identifier for the subscriber. Therefore, most people can be uniquely identified by recording the IMSI that their ME is sending.

The 3GPP specifications provide means for implementations to send the IMSI less often by using the Temporary Mobile Subscriber Identity (TMSI) where possible. However, the decision on when to use IMSI or TMSI is entirely on the networks side, without any control by the ME or even the subscriber.

This leads to a variety of attacks on subscriber location privacy, including the use of passive air-interface sniffing as well as false base station attacks, where an attacker impersonates a base station which subsequently inquires every ME about its IMSI.

Some related devices have been termed *IMSI catchers* or *Stingray* in both scientific literature as well as mainstream media. IMSI catchers have become small and affordable during the last decade; criminals actors and in some cases even tabloid journalists without much budget have reportedly used them to track anybody with a mobile phone.

5G addresses this problem with the Subscriber Concealed Identifier (SUCI), which uses public-key cryptography to ensure that the permanent subscriber identity (IMSI) is not transmitted over the air interface anymore. Rather, a concealed version of it is transmitted (3GPP TS 33.501, Section 6.12.2). The 5G SUCI mechanism can not be adapted easily for previous generations of cellular networks as it relies on introducing an entirely new mobile identity type of larger size (SUCI) than any of the existing ones (e.g. IMSI), causing significant implications on protocol stacks and implementations all across the protocol stack of all network elements, including the ME.

No mechanism for increasing subscriber identity and location privacy on the radio interface has been specified for the previous cellular technologies 2G (GSM), 3G (UMTS) and 4G (LTE). Meanwhile, pure 5G networks are and will remain rare for many years to come, as operators have to support billions of deployed legacy pre-5G ME. Operating combined 5G + previous technology networks enables the so-called "downgrade attacks" where the attacker blocks access to 5G e.g. by means of jamming/interference, and hence triggers the ME to use a previous generation which is still susceptible to the attacks.

This specification proposes a different approach to conceal the IMSI for legacy 2G, 3G and 4G networks.

1.2 Summary of Proposed Solution

The solution presented in this document is to periodically change the IMSI of the ME to a new pseudonymous IMSI allocated by the Home Location Register (HLR) or Home Subscriber Service (HSS). The next pseudonymous IMSI is sent to the SIM/USIM via Short Message Service (SMS), then a SIM applet overwrites the IMSI of the SIM/USIM with the new value. The only components in the network that need to be changed in order to support this mechanism are the SIM/USIM and the HLR/HSS. All other elements (like BTS, NodeB, eNodeB, BSC, RNC, MME, MSC/VLR, SGSN, GGSN, S-GW, P-GW, ...) remain as-is, without any changes to their specification or implementation.

Constraining the required changes to only two elements in the network enables quick adoption potential for the proposed mechanism. Furthermore, as SIM/USIM and HLR/HSS are the only two elements under control of a Mobile Virtual Network Operator (MVNO), this mechanism can be deployed by a MVNO without any changes to the operators of the physical infrastructure (MNO).

1.3 Summary of Existing Location Updating Procedures in RAN and CN

Every subscriber's SIM/USIM is provisioned with the IMSI and secret cryptographic keys (Ki or K+OP[c]). The same IMSI and key data is also provisioned into the HLR/HSS, the central subscriber database of a cellular network.

In a number of different situations, the IMSI is sent over the air interface and back-haul towards the Core Network (CN), where it is validated by the HLR/HSS. The involved components vary by the generation of the network and whether the SIM/USIM is attempting a Circuit Switched (CS) or Packet Switched (PS) connection, but the principle is the same. This document uses 2G CS Location Updating for reference, as in Figure 1.

The IMSI is transmitted in the Location Updating Request from ME. The VLR needs an authentication challenge specific to the secret keys on the SIM/USIM to authenticate the SIM/USIM, and looks the authentication challenges up by the IMSI. If the VLR does not have any more authentication challenges for the IMSI (as it happens when the VLR sees the IMSI for the first time), the VLR requests new authentication challenges from the HLR/HSS. Then the HLR/HSS verifies that the IMSI is known and, if it is unknown, sends back an error that will terminate the Location Updating procedure.

After the VLR found the authentication challenge, it authenticates the SIM/USIM, and performs a Classmark Enquiry and Physical Channel Reconfiguration. Then the VLR has the required information to finish the Location Updating, and continues with Process Update_Location_HLR (3GPP TS 29.002). Afterwards, the VLR assigns a new TMSI with the Location Updating Accept, which is acknowledged by the TMSI Reallocation Complete. In following Location Updates with the same MSC, the ME sends the TMSI instead of the IMSI in the Location Updating Request.

However, the allocation of the TMSI is optional (the network may choose to not perform it), and particularly at mobility changes across the MSC/VLR boundary, or even across the PLMN boundary, the TMSI allocated by the previouis network element may not be known, and an IMSI based Location Updating procedure is used.

Furthermore, at any given point in time, a legitimate network or a rogue base station can inquire the IMSI from the ME using the "MM IDENTITY REQUEST (IMSI)" command.

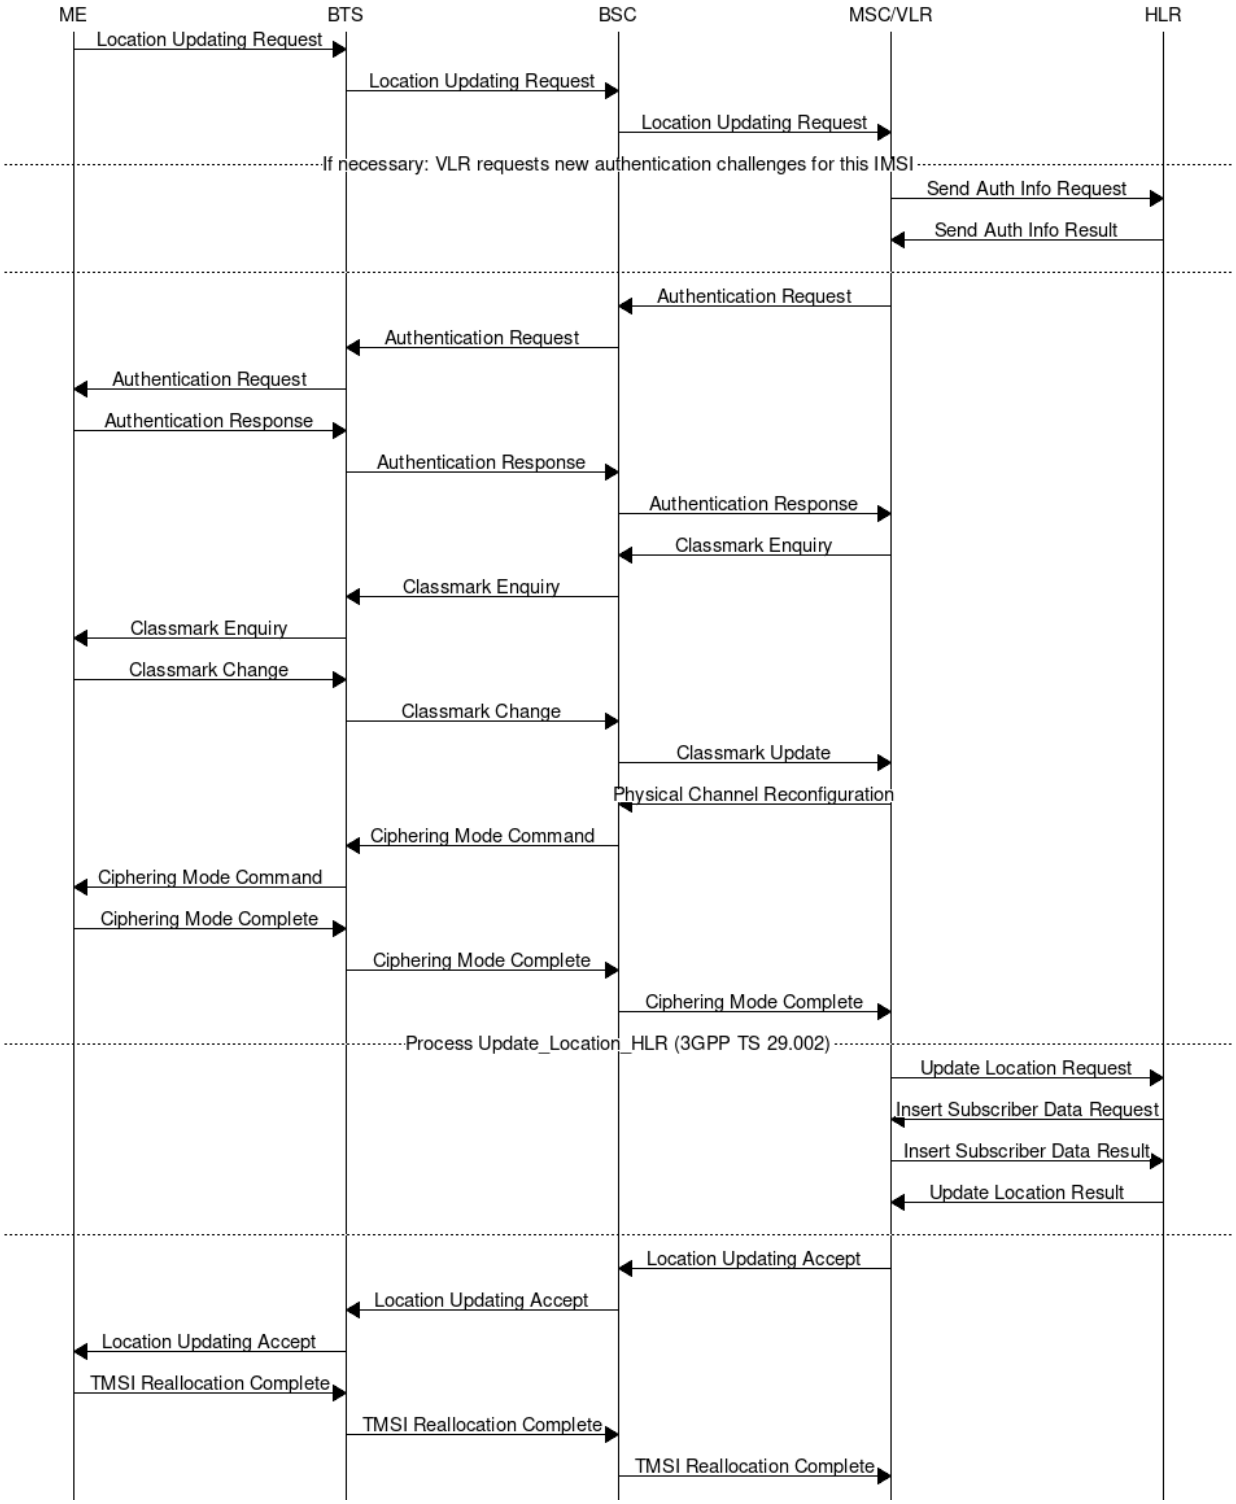

Figure 1: Location Updating in 2G CS with IMSI

2 Required Changes

This section covers the changes / enhancements required compared to the existing 3GPP specifications.

2.1 Pseudonymous IMSI Storage in the HLR/HSS

The HLR/HSS must store up to two pseudonymous IMSIs (`imsi_pseudo`) and their related counters (`imsi_pseudo_i`) per subscriber. Each subscriber initially has one pseudonymous IMSI allocated. A subscriber has two valid pseudonymous IMSIs only during the transition phase from the old pseudonymous IMSI to the new one.

Subsequently, the amount of available IMSIs must be higher than the amount of subscribers registered with the HLR/HSS. If the amount of available IMSIs is too small, the HLR/HSS could delay assigning new pseudonymous IMSIs until new IMSIs are available again.

Table 1: Examples for additional subscriber data in HLR

Subscriber ID	**imsi_pseudo**	**imsi_pseudo_i**
123	310150123456789	1
234	502130123456789	1
234	460001357924680	2

2.1.1 imsi_pseudo

The value for `imsi_pseudo` is a random choice from the pool of available IMSIs that the HLR/HSS controls. The pseudonymous IMSI must not be used by any subscriber as pseudonymous IMSI yet, but may be the real IMSI of a subscriber.

2.1.2 imsi_pseudo_i

The counter `imsi_pseudo_i` indicates how often a subscribers pseudonymous IMSI was changed. The value is 1 for the first allocated pseudonymous IMSI of a subscriber. When allocating a new pseudonymous IMSI for the same subscriber, the new `imsi_pseudo_i` value is increased by 1. The counter is used by the SIM/USIM applet to detect and ignore outdated requests related to changing the pseudonymous IMSI.

2.2 SIM/USIM Provisioning

IMSI pseudonymization as specified by this document works with traditional SIM (used in 2G), as well as with USIM (used from 3G onwards).

The initial IMSI provisioned in the SIM/USIM is provisioned as the initial pseudonymous IMSI in the HLR/HSS.

2.2.1 SIM applet

SIM/USIM have long supported the installation and operation of additional applets on the card itself. The programming language and runtime environment for such applets is an implementation detail. However, the industry has converged around JavaCards with related additional APIs specific to SIM, UICC and USIM. Depending on the card profile / provisioning, it is possible for such applets to access the card file system and modify files on the card, such as the file storing the IMSI.

A SIM/USIM compatible with this specification is provisioned with a SIM applet, which is able to change the IMSI once the next pseudonymous IMSI arrives from the HLR/HSS. A reference implementation is provided in Section 5.

2.2.1.1 Counter Storage

The following counter variables are stored in the SIM applet.

Name	Initial value	Description
imsi_pseudo_i	1	See Section 2.1.2.
imsi_pseudo_lu	0	Amount of Location Updating procedures done with the same pseudonymous IMSI.
imsi_pseudo_lu_max	(decided by operator)	Maximum amount of Location Updating procedures done with the same pseudonymous IMSI, before the SIM applet shows a warning to the subscriber.

2.2.1.2 Switch to Next Pseudonymous IMSI

The SIM applet registers to a suitable SMS trigger (3GPP TS 43.019, Section 6.2). When an SMS from the HLR/HSS in the structure of Section 2.3.4 arrives, the applet must verify that the SMS is not outdated by comparing `imsi_pseudo_i` from the SMS with the last `imsi_pseudo_i` that was used when changing the IMSI (initially 1 as in Section 2.1.2). The new value must be higher, otherwise the SMS should not be processed further.

The SIM applet registers a timer with `min_sleep_time` from the SMS. When the timer triggers, EF_{IMSI} of the SIM/USIM is overwritten with the new pseudonymous IMSI. The TMSI and related data (EF_{LOCI}, EF_{PSLOCI}) and ciphering keys (EF_{Kc}, EF_{KcGPRS}, EF_{Keys}, EF_{KeysPS}) are invalidated (see 3GPP TS 31.102). The current `imsi_pseudo_i` from the SMS is stored in the SIM applet to compare it with the next SMS. `imsi_pseudo_lu` is reset to 0. Afterwards, the EF_{IMSI} changing procedure in 3GPP TS 11.14, Section 6.4.7.1 is executed to apply the new IMSI.

2.2.1.3 Warning the Subscriber If the Pseudonymous IMSI Does Not Change

An attacker could potentially block the next pseudonymous IMSI SMS on purpose. Because the SIM applet cannot decide the next pseudonymous IMSI, it would have the same pseudonymous IMSI for a long time. Then it could become feasible for an attacker to track the subscriber by their pseudonymous IMSI. Therefore the SIM applet should warn the subscriber if the pseudonymous IMSI does not change.

The SIM applet registers to EVENT_EVENT_DOWNLOAD_LOCATION_STATUS (3GPP TS 03.19, Section 6.2) and increases `imsi_pseudo_lu` by 1 when the event is triggered. If `imsi_pseudo_lu` reaches `imsi_pseudo_lu_max`, the SIM applet displays a warning to the subscriber.

2.3 Process Update_Location_HLR

All IMSI Pseudonymization related changes to Process Update_Location_HLR (3GPP TS 29.002) are optional. Deviations from the existing specification that are outlined in this section are expected to be enabled or disabled entirely where IMSI pseudonymization is implemented.

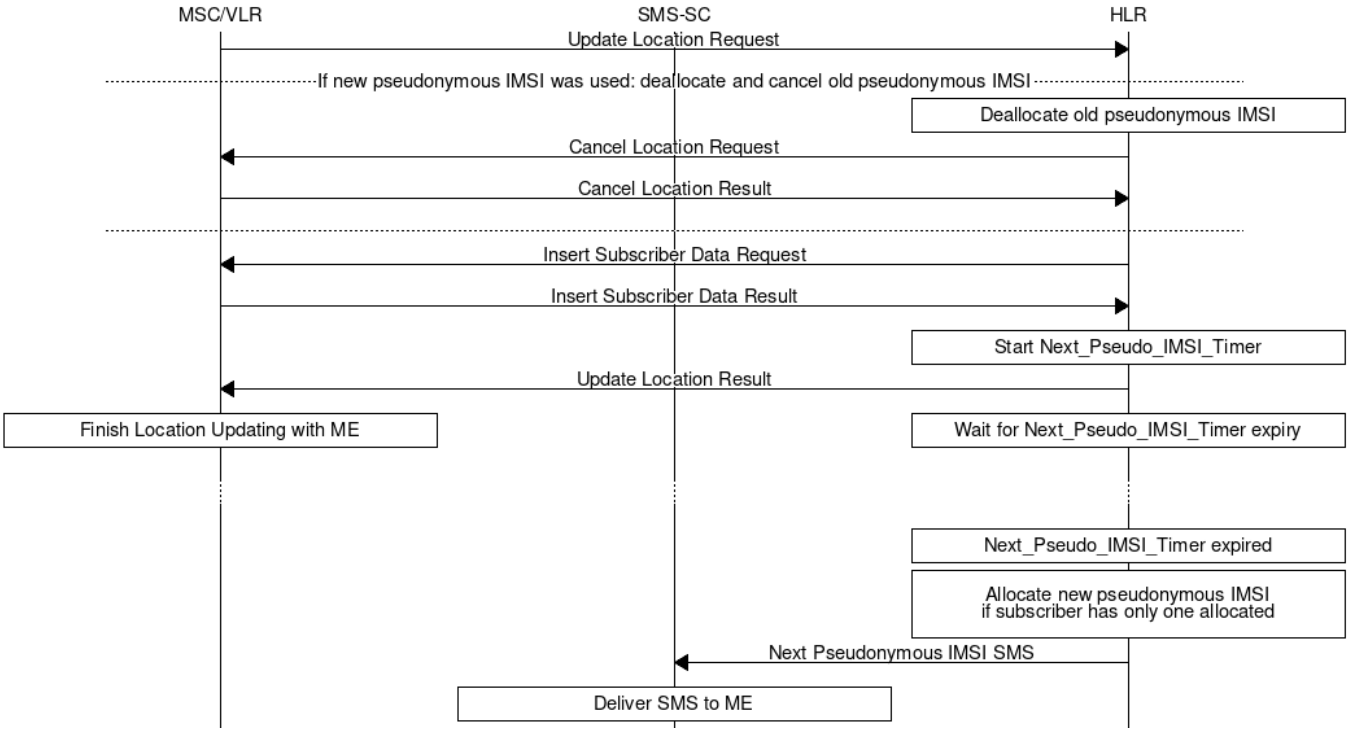

Figure 2: Process Update_Location_HLR with IMSI pseudonymization changes

2.3.1 Update Location Request

When Update Location Request arrives, the HLR/HSS does not look up the subscriber by the IMSI, but by the pseudonymous IMSI instead. Unless the subscriber has two pseudonymous IMSI allocated and used the new pseudonymous IMSI in the Update Location Request, this is followed by the existing logic to continue with Insert Subscriber Data Request.

2.3.1.1 Update Location Request With New Pseudonymous IMSI

If the subscriber has two pseudonymous IMSIs allocated, and the newer entry was used (higher `imsi_pseudo_i`, see Section 2.1.2), this section applies. The older pseudonymous IMSI is deallocated in the HLR/HSS. This is done as early as possible, so the timeframe where two pseudonymous IMSI are allocated for one subscriber is short.

A Cancel Location Request with the old pseudonymous IMSI is sent to the VLR, so the conflicting subscriber entry with the old pseudonymous IMSI is deleted from the VLR. Receiving a Cancel Location Result is followed by the existing logic to continue with Insert Subscriber Data Request.

2.3.1.2 Update Location Request With Old Pseudonymous IMSI

If the subscriber has two pseudonymous IMSIs allocated, and the older entry was used (lower `imsi_pseudo_i`, see Section 2.1.2), the newer entry is *not* deallocated. This could lock out the subscriber from the network if the SMS with the new pseudonymous IMSI arrives with a delay.

2.3.2 Insert Subscriber Data Result

When Insert Subscriber Data Result arrives, a subscriber specific Next_Pseudo_IMSI_Timer starts.

2.3.3 Next_Pseudo_IMSI_Timer Expires

If the subscriber has only one pseudonymous IMSI allocated, and the amount of available IMSIs in the HLR/HSS is high enough, a second pseudonymous IMSI and related `imsi_pseudo_i` gets allocated for the subscriber (as described in Section 2.1).

If the subscriber still has only one pseudonymous IMSI, because not enough IMSIs were available in the HLR/HSS, the process is aborted here and no SMS with a next pseudonymous IMSI is sent to the subscriber. The subscriber will get a new pseudonymous IMSI during the next Location Updating Procedure, if the HLR/HSS has enough IMSIs available at that point.

An SMS is sent to the SMS - Service Centre (SMS-SC) with the newer pseudonymous IMSI (higher `imsi_pseudo_i`, see Section 2.1.2) and related `imsi_pseudo_i` value.

2.3.4 Next Pseudonymous IMSI SMS Structure

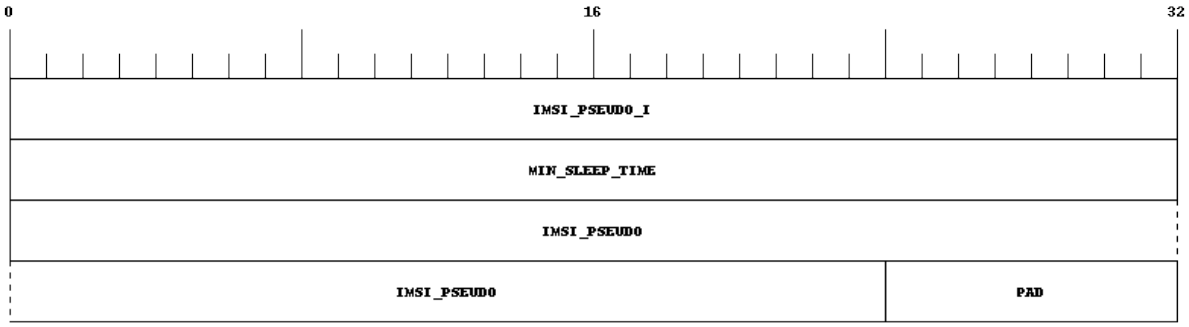

Figure 3: Next pseudonymous IMSI SMS structure

Important
This is a draft. The structure is likely to change after the reference implementation phase.

IMSI_PSEUDO_I: 32 bits
See Section 2.1.2.

MIN_SLEEP_TIME: 32 bits
Amount of seconds, which the SIM applet should wait before changing to the new pseudonymous IMSI. Since it is unclear when the SMS will arrive (ME might be turned off), this is a minimum amount.

IMSI_PSEUDO: 60 bits
Telephony Binary Coded Decimal (TBCD, 3GPP TS 29.002) version of the next pseudonymous IMSI.

PAD: 8 bits
Padding at the end, should be filled with 1111 as in the TBCD specification.

3 Error Scenarios

3.1 Next Pseudonymous IMSI SMS is Lost

If the SMS with the next pseudonymous IMSI does not arrive, the SIM/USIM will start the next Location Updating Procedure with the old pseudonymous IMSI. Because the HLR/HSS has both the old and the new pseudonymous IMSI allocated at this point, the subscriber is not locked out of the network.

3.2 Next Pseudonymous IMSI SMS Arrives Out of Order

The next pseudonymous IMSI SMS may arrive out of order. Either, because the network is not able to deliver them in order, or even because an attacker would perform a replay attack.

If the SMS arrives out of order, the `imsi_pseudo_i` counter will not be higher than the value the SIM applet (Section 2.2.1) has stored. Therefore, the applet will discard the message and the subscriber is not locked out of the network.

4 Recommendations for Real-World Implementations

4.1 BCCH SI3: ATT = 0

When changing from one pseudonymous IMSI to the next, it is important that the ME does not detach from the network. Otherwise it would be trivial for an attacker to correlate the detach with the attach of the same ME with the next pseudonymous IMSI.

This is controlled with the ATT flag in the SYSTEM INFORMATION TYPE 3 (SI3) message on the Broadcast Control Channel (BCCH), see 3GPP TS 44.018 Section 10.5.2.11. It must be set to 0.

4.2 End to End Encryption of SMS

When deploying the IMSI pseudonymization, the operator should make sure that the next pseudonymous IMSI SMS (Section 2.3.4) cannot be read or modified by third parties. Otherwise, the next pseudonymous IMSI is leaked, and if the pseudonymous IMSI in the SMS was changed, the SIM/USIM would be locked out of the network.

The safest way to protect the next pseudonymous IMSI SMS is a layer of end to end encryption from the HLR/HSS to the SIM/USIM. The existing means for OTA SMS security (3GPP TS 23.048) provide mechanisms for integrity protection, confidentiality as well as replay protection and must be implemented when using IMSI pseudonymization.

4.3 User-configurable Minimum Duration Between IMSI Changes

It may be desirable to let subscribers configure their minimum duration between IMSI changes. This allows subscribers with a high privacy requirement to switch their pseudonymous IMSI more often, and it allows the pseudonymous IMSI change to happen less frequently if it is distracting to the subscriber.

How distracting the pseudonymous IMSI change is, depends on the ME. The following examples were observed:

- A Samsung GT-I9100 Galaxy SII smartphone with Android 4.0.3 displays a message at the bottom of the screen for about 5 seconds, but the user interface remains usable.
- A Samsung GT-E1200 feature phone displays a waiting screen for 16 to 17 seconds and is unusable during that time.

5 Reference Implementation with Source Code

A reference implementation for the SIM applet (Section 2.2.1) is available in source code under the Apache-2.0 license at:

https://osmocom.org/projects/imsi-pseudo

The HLR/HSS modifications described in Section 2.1 and Section 2.3 were implemented for reference in OsmoHLR from the Osmocom project, licensed under AGPL-3.0. Information about the source code and related branches for IMSI pseudonymization can be found at the above URL as well.

