3GPP TSG SA 3 Meeting #12
Stockholm, Sweden, 11-14 April 2000

***Document* S3-000253**
e.g. for 3GPP use the format TP-99xxx
or for SMG, use the format P-99-xxx

CHANGE REQUEST

Please see embedded help file at the bottom of this page for instructions on how to fill in this form correctly.

33.102 CR Current Version: 3.3.1

GSM (AA.BB) or 3G (AA.BBB) specification number ↑ *↑ CR number as allocated by MCC support team*

For submission to: TSG SA #7
list expected approval meeting # here
↑

for approval	☒	strategic	□	*(for SMG*
for information	□	non-strategic	□	*use only)*

Form: CR cover sheet, version 2 for 3GPP and SMG *The latest version of this form is available from: ftp://ftp.3gpp.org/Information/CR-Form-v2.doc*

Proposed change affects: (U)SIM □ ME □ UTRAN / Radio □ Core Network □
(at least one should be marked with an X)

Source: Siemens Atea **Date:** 2000-04-04

Subject: Authentication and key agreement

Work item: Security

Category:
(only one category shall be marked with an X)

F	Correction	□	**Release:**	Phase 2	□
A	Corresponds to a correction in an earlier release	□		Release 96	□
B	Addition of feature	□		Release 97	□
C	Functional modification of feature	□		Release 98	□
D	Editorial modification	☒		Release 99	☒
				Release 00	☒

Reason for change: Better presentation of the mechanism for authentication and key agreement. The mechanism for authentication and the mechanism for re-synchronisation are discussed separately. The procedures are discussed separately from the mechanisms, and the functions implemented in the USIM and the HLR/AuC are discussed separately from the procedures.

Clauses affected: 6.3

Other specs affected:

Other 3G core specifications	□	→ List of CRs:
Other GSM core specifications	□	→ List of CRs:
MS test specifications	□	→ List of CRs:
BSS test specifications	□	→ List of CRs:
O&M specifications	□	→ List of CRs:

Other comments: The existing 6.3 should be replaced by the attached text. (The traditional lay-out with change bars is not used as it does not improve readability in this case).

help.doc

<--------- double-click here for help and instructions on how to create a CR.

6.3 Authentication and key agreement

6.3.1 General

The mechanism described here achieves mutual entity authentication and the establishment of a shared secret cipher key and integrity key between the MS at the user side and the VLR or SGSN on behalf of the user's HLR/AuC at the network side.

The mechanism uses symmetric key techniques using a secret subscriber authentication key K that is shared between and available only to the USIM and the AuC in the user’s HE. In addition, the AuC keeps track of a counter SQN_{HE} and the USIM keeps track of a counter SQN_{MS} and stores additional data to support network authentication and to provide the user with assurance of key freshness.

The method was chosen in such a way as to achieve maximum compatibility with the current GSM security architecture and facilitate migration from GSM to UMTS. The method is composed of a challenge/response protocol identical to the GSM subscriber authentication and key establishment protocol combined with a sequence number-based one-pass protocol for network authentication derived from the ISO standard ISO/IEC 9798-4 (section 5.1.1).

The HE, that manages both the HLR/AuC and the USIM, has some flexibility in the management of sequence numbers, but some requirements need to be fulfilled:

a) The mechanism shall support re-synchronisation of the counter SQN_{HE} in the AuC to the value of the counter SQN_{MS} in the USIM, as described in section 6.3.2.2;

b) The mechanism shall protect against wrap around of the counter SQN_{MS} in the USIM. A mechanism to achieve this is provided in C.2.

c) The mechanism should not compromise user identity and location confidentiality. If consecutive sequence numbers for the same user are highly correlated, sending them in the clear should be considered as a compromise of user identity and location confidentiality, and the use of concealment of the sequence number, described as an option throughout 6.4.3, is recommended. In case however, the sequence numbers SQN are partly derived from time, such correlation is minimised, and the concealment may not be required.

d) The mechanisms for verifying the freshness of sequence numbers in the USIM shall to some extent allow the out-of-order use of sequence numbers. This is to ensure that the authentication failure rate due to synchronisation failures is sufficiently low. This requires the capability of the USIM to store information on past successful authentication events (e.g. sequence numbers or relevant parts thereof). The mechanism shall ensure that a sequence number can still be accepted if it is among the last 50 sequence numbers generated.

Note 1: This shall not preclude that a sequence number is rejected for other reasons such as a limit on the age for time-based sequence numbers.

Note 2: The same minimum number (50) needs to be used across the systems to guarantee that the synchronisation failure rate is sufficiently low under various usage scenarios, in particular simultaneous registration in the CS- and the PS-service domains, user movement between VLRs and/or SGSNs that do not exchange authentication data and super-charged networks.

Annex C contains a detailed description of a sequence number management scheme that satisfies the above conditions.

6.3.2 Mechanisms

6.3.2.1 Authentication and key agreement

An overview of the mechanism for authentication and key agreement is shown in Figure 6.3.1.

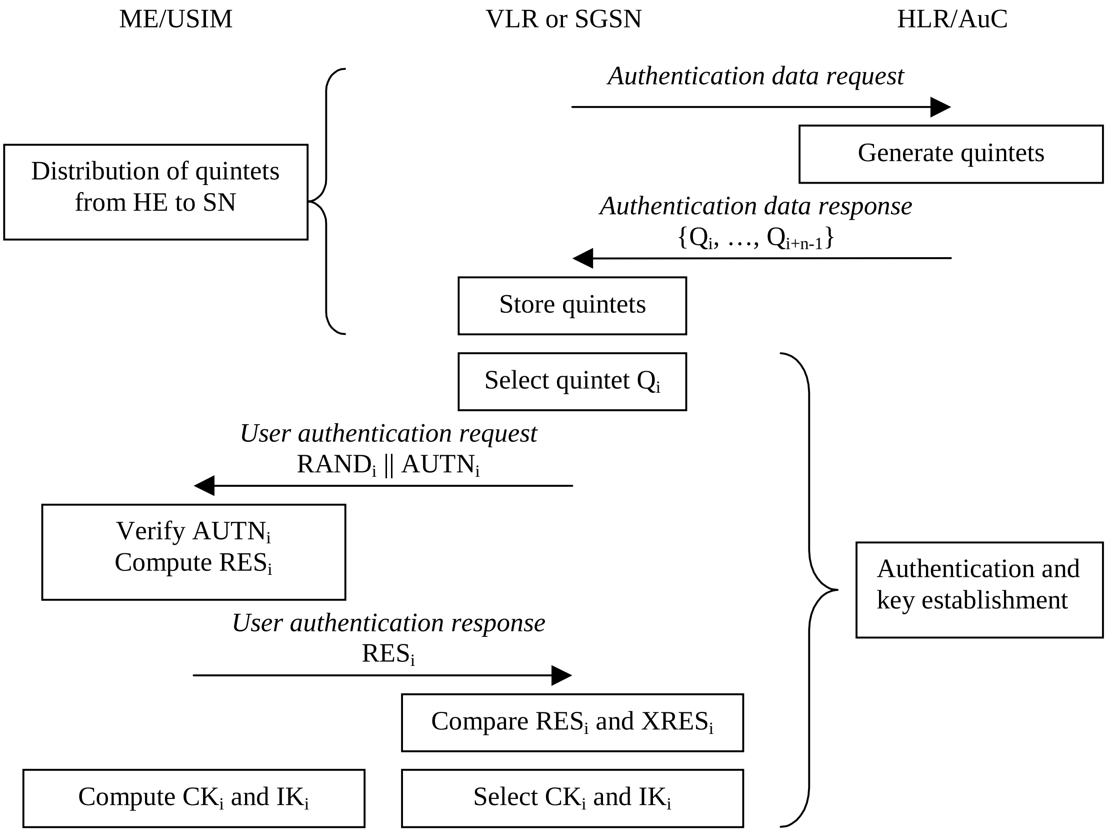

Figure 6.3.1: Authentication and key agreement mechanism

The procedure for distribution of authentication data from the HE to a service domain in the SN (described in 6.3.3.1) starts with the VLR or SGSN sending a request to the user's HLR/AuC. Upon receipt of that request the HLR/AuC sends a quintet (the equivalent of a GSM "triplet") or an ordered array of n quintets to the VLR or SGSN. Each quintet consists of the following components: a challenge RAND, an expected response XRES, a cipher key CK, an integrity key IK and an authentication token AUTN. Each quintet is good for one authentication and key agreement between the VLR or SGSN and the MS.

When the VLR or SGSN initiates the over-the-air authentication and key agreement procedure (described in 6.3.3.2), it selects the next quintet from the array and sends the parameters RAND and AUTN to the user. The USIM checks whether AUTN can be accepted and, if so, produces a response RES which is sent back to the VLR or SGSN. The USIM also computes CK and IK. The VLR or SGSN compares the received RES with XRES. If they match the VLR or SGSN considers the authentication and key agreement exchange to be successfully completed and selects the corresponding CK and IK from the quintet. The established keys CK and IK are transferred by the USIM to the ME and by the VLR or SGSN to RNC; the entities that perform ciphering and integrity protection. The USIM stores the established cipher/integrity keys until the next successful authentication and key agreement.

If the USIM also supports cipher key agreement for the GSM radio interface, the USIM in addition derives a GSM cipher key Kc that is passed along to the ME (see 6.8.1).

The over-the-air authentication and key agreement procedure can fail for three reasons:

a) The USIM may successfully verify the integrity of the (RAND, AUTN) pair, but may be unable to verify the freshness of the (RAND, AUTN) pair. In this case the USIM shall trigger the re-synchronisation mechanism (see 6.3.3.2).

b) The USIM may find that the integrity of the (RAND, AUTN) pair could not be verified. In that case, the user informs the VLR or SGSN of the failure and of its nature, but no parameters are sent. The VLR or SGSN shall inform the HLR/AuC (see 6.3.3.4) about the failure and may request for new quintets (see 6.3.3.1). The VLR or SGSN may also decide to initiate a new identification and authentication procedure towards the user.

c) The VLR or SGSN may find that the user response RES and the expected response XRES do not match. In that case the VLR or SGSN sends *user authentication reject* to the MS. The VLR or SGSN shall inform the HLR/AuC (see 6.3.3.4) about the failure and may request for new quintets (see 6.3.3.1).

The VLR and SGSN shall use or attempt to use a quintet only once. Hence, quintets cannot be re-used. A VLR or SGSN can serve a user securely even when links to the user's HLR/AuC are unavailable by means of re-use of previously derived cipher and integrity keys. In this manner a secure connection can be set up without the need for an authentication and key agreement and a fresh quintet. Authentication is in that case based on a shared integrity key, by means of data integrity protection of signalling messages (see 6.4).

6.3.2.2 Re-synchronisation

An overview of the mechanism for resynchronisation is shown in Figure 6.3.2:

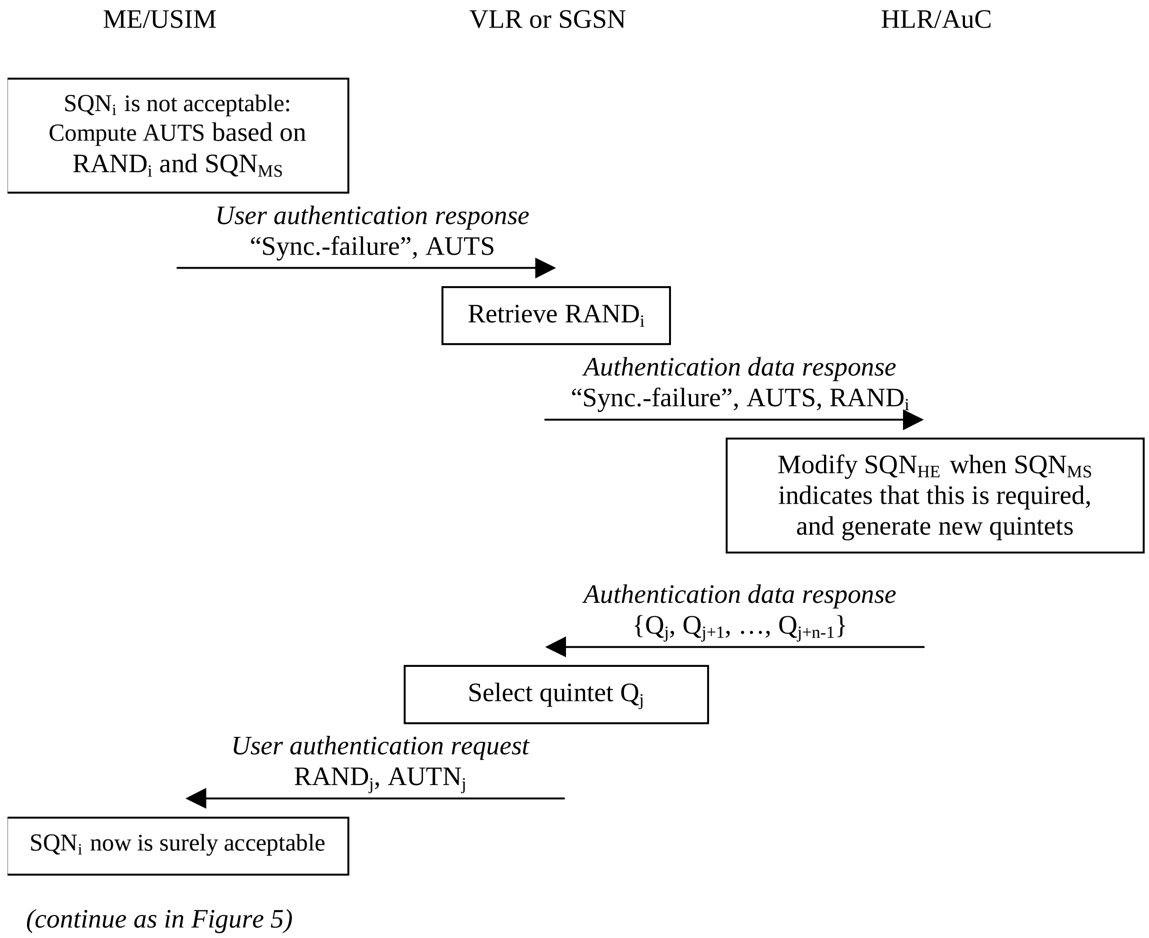
ME/USIM
VLR or SGSN
HLR/AuC
SQN$_i$ is not acceptable: Compute AUTS based on RAND$_i$ and SQN$_{MS}$
User authentication response
"Sync.-failure", AUTS
Retrieve RAND$_i$
Authentication data response
"Sync.-failure", AUTS, RAND$_i$
Modify SQN$_{HE}$ when SQN$_{MS}$ indicates that this is required, and generate new quintets
Authentication data response
$\{Q_j, Q_{j+1}, \ldots, Q_{j+n-1}\}$
Select quintet Q$_j$
User authentication request
RAND$_j$, AUTN$_j$
SQN$_i$ now is surely acceptable
(continue as in Figure 5)

Figure 6.3.2: Re-synchronisation mechanism

The mechanism for re-synchronisation is triggered by the unsuccessful verification by the USIM of the freshness of SQN$_i$ that is included in AUTN$_i$ (see 6.3.2.1). The USIM then sends a *user authentication response* to the VLR or SGSN including an indication of synchronisation failure and a re-synchronisation token AUTS, that includes the current value of the counter SQN$_{MS}$. The VLR or SGSN appends the challenge RAND$_i$ and sends an *authentication data request* to the HLR/AuC with indication of synchronisation failure and including (RAND$_i$, AUTS). Upon receipt of such a request, the HLR/AuC verifies whether the value of SQN$_{MS}$ mandates that the SQN$_{HE}$ needs to be modified. If necessary, the HLR/AuC shall set SQN$_{HE}$ equal to SQN$_{MS}$. Consecutively, the HLR/AuC sends the VLR quintets generated from the current SQN$_{HE}$, which are forwarded to the user.

The new quintet will now surely be acceptable to the user. For a formal proof see TR 33.902.

6.3.3 Procedures

6.3.3.1 Distribution of authentication data from HE to SN

The purpose of this procedure is to provide the VLR or SGSN with an array of fresh quintets from the user's HE to perform a number of authentication and key agreement exchanges.

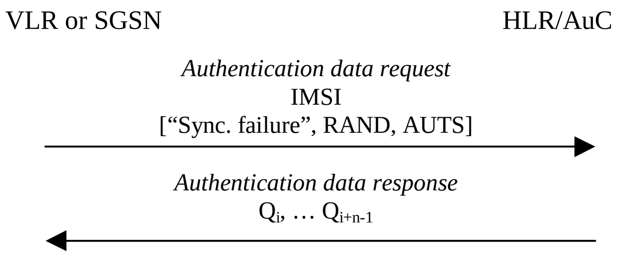
VLR or SGSN
HLR/AuC
Authentication data request
IMSI
[“Sync. failure”, RAND, AUTS]
Authentication data response
Q_i, … Q_{i+n-1}

Figure 6.3.3: Distribution of authentication data from HE to SN

The VLR or SGSN invokes the procedures by requesting quintets to the HLR/AuC.

The protocol steps are as follows:

a) The VLR or SGSN sends an *authentication data request* to the HLR/AuC; this message shall contain the IMSI and may contain an indication of synchronisation failure and shall in that case also contain a re-synchronisation token AUTS and a challenge RAND.

b) Upon receipt of an *authentication data request* with an indication of synchronisation failure the HLR/AuC acts as described in 6.3.4.4. It shall verify whether the counter SQN_{HE} needs to be modified and contingent on the outcome set SQN_{HE} to SQN_{MS}.

c) The HLR/AuC then sends an authentication data response back to the VLR or SGSN that includes a quintet Q or an ordered array of quintets {Q_i, …, Q_{i+n-1}} that have/has been generated as described in 6.3.4.1. The quintet(s) may have been generated in advance or on demand. In case a synchronisation failure caused the counter SQN_{HE} to be reset the quintet(s) are/is generated on demand.

d) Upon receipt the VLR or SGSN stores the user identity and/or quintets it receives, maintaining the ordering.

6.3.2.2 Authentication and key agreement

The purpose of this procedure is to authenticate the user and establish a new pair of cipher and integrity keys between the VLR or SGSN and the ME/USIM. During the authentication, the user verifies data origin, the integrity and the freshness of the quintet that is used. The procedure is shown in Figure 6.3.4.

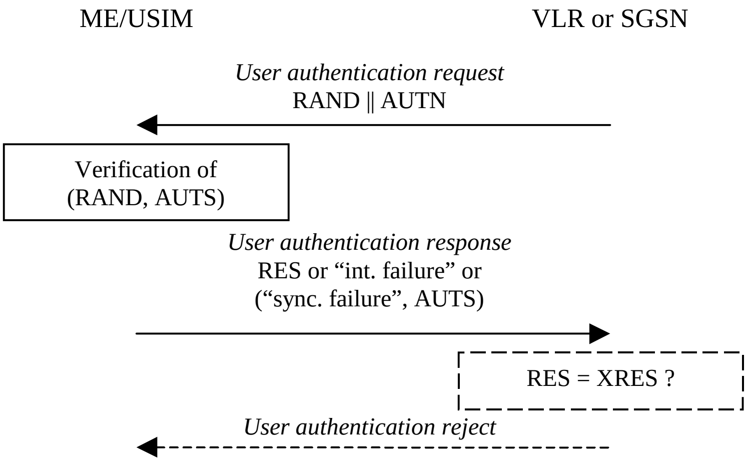

Figure 6.3.4: Over-the-air authentication and key agreement procedure

The VLR or SGSN invokes the procedure by selecting the next unused quintet from the ordered array of quintets in the VLR or SGSN database.

The protocol steps are the following:

a) The VLR or SGSN sends to the user a *user authentication request*, including the network challenge RAND and the authentication token AUTN from the selected quintet.

b) The USIM then verifies the (RAND, AUTN) pair as described in 6.3.4.2, and contingent on the outcome acts as follows:

   i) In case the data origin and integrity of (RAND, AUTN) is successfully verified, and the sequence number is acceptable, the ME sends a *user authentication response* back with an indication of success and including the user response RES;

   ii) In case the data origin and integrity of (RAND, AUTN) is not successfully verified, the ME sends a *user authentication response* back with and indication of integrity failure (without any parameter);

   iii) In case the data origin and integrity of (RAND, AUTN) is successfully verified, but the sequence number is not acceptable, the ME sends a *user authentication response* back with and indication of synchronisation failure and including the re-synchronisation token AUTS.

c) Upon receipt of the *user authentication response,* the VLR or SGSN acts as follows:

   i) In case of success, the VLR or SGSN compares the received response RES with the expected response XRES. In case there is a match, the VLR or SGSN selects the CK and IK and authentication ends successfully. On the other hand, in case there is a mismatch, the VLR or SGSN sends *user authentication reject* to the user and authentication ends unsuccessfully. The VLR or SGSN should in that case report the failure to the HE, as described in 6.3.2.4.

   ii) In case of integrity failure, the VLR or SGSN may report the failure to the HE, as described in 6.3.2.4 or may request for new quintets using the procedure described in 6.3.2.1.

   iii) In case of synchronisation failure, the VLR or SGSN may report the failure to the HE, as described in 6.3.2.4 but should request for new quintets using the procedure described in 6.3.2.1, include an indication of synchronisation failure, the parameter AUTS and the parameter RAND. The VLR or SGSN deletes the user's quintets from the database.

The VLR or SGSN shall discard any unsolicited user authentication response, in particular, it shall discard an unsolicited user authentication response with indication of synchronisation failure.

The VLR or SGSN shall not send a user new *user authentication requests* before it has received a response of the user or the before a certain time period has elapsed.

6.3.2.3 Distribution of IMSI and temporary authentication data within one serving network domain

The purpose of this procedure is to provide a newly visited VLR or SGSN with temporary authentication data from a previously visited VLR or SGSN within the same serving network domain.

The procedure is shown in Figure 6.3.5.

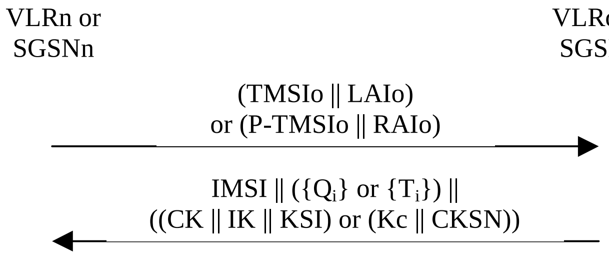
VLRn or SGSNn
VLRo or SGSNo
(TMSIo || LAIo)
or (P-TMSIo || RAIo)
IMSI || ({Q_i} or {T_i}) ||
((CK || IK || KSI) or (Kc || CKSN))

Figure 6.3.5: Distribution of IMSI and temporary authentication data within one serving network domain

The procedure shall be invoked by the newly visited VLRn (resp. SGSNn) after the receipt of a location update request (resp. routing area update request) from the user wherein the user is identified by means of a temporary user identity TMSIo (resp. P-TMSIo) and the location area identity LAIo (resp. routing area identity RAIo) under the jurisdiction of a previously visited VLRo or SGSNo that belongs to the same serving network domain as the newly visited VLRn or SGSNn.

The protocol steps are as follows:

a) The VLRn (resp. SGSNn) sends a *user identity request* to the VLRo (or SGSNo), this message contains TMSIo and LAIo (resp. P-TMSIo and RAIo).

b) The VLRo (resp. SGSNo) searches the user data in the database.

If the user is found, the VLRo (resp. SGSNo) shall send a *user identity response* back that

i) shall include the IMSI,

ii) may include a number of unused authentication vectors (quintets or triplets) and

iii) may include the current security context data: CK, IK and KSI (UMTS) or Kc and CKSN (GSM).

The VLRo or SGSNo subsequently deletes the authentication vectors which have been sent and the data elements on the current security context.

If the user cannot be identified the VLRo or SGSNo shall send a *user identity response* indicating that the user identity cannot be retrieved.

c) If the VLRn or SGSNn receives a *user identity response* with an IMSI, it creates an entry and stores any authentication vectors and any data on the current security context that may be included.

If the VLRn or SGSNn receives a *user identity response* indicating that the user could not be identified, it shall initiate the user identification procedure described in 6.2.

6.3.2.4 Reporting authentication failures from SN to HE

The purpose of this procedure is to provide a mechanism for reporting authentication failures from the serving environment back to the home environment.

The procedure is shown in Figure 6.3.6.

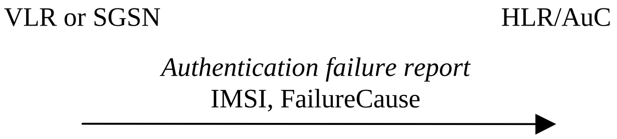

Figure 6.3.6: Reporting authentication failures from SN to HE

The procedure is invoked by the serving network VLR or SGSN when the authentication procedure fails. The *authentication failure report* shall contain the subscriber identity and a failure cause code. The possible failure causes are either that the network signature was wrong or that the user response was wrong.

The HE may decide to cancel the location of the user after receiving an *authentication failure report*.

6.3.4 Functions

6.3.4.1 Generation of quintets in the AuC

For each user the HLR/AuC keeps track of a counter: SQN_{HE}.

The AuC generates quintets as follows:

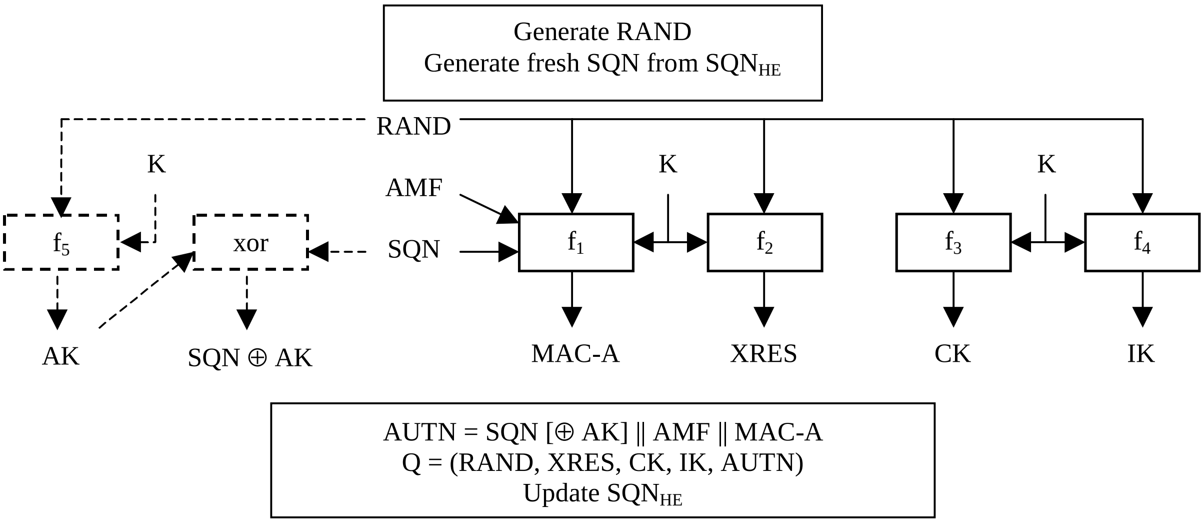
Generate RAND
Generate fresh SQN from SQN_HE
RAND
K
AMF
K
K
SQN
f5
xor
f1
f2
f3
f4
AK
SQN ⊕ AK
MAC-A
XRES
CK
IK
AUTN = SQN [⊕ AK] || AMF || MAC-A
Q = (RAND, XRES, CK, IK, AUTN)
Update SQN_HE

Figure 6.3.7: Generation of quintets in the AuC

a) The HLR/AuC generates a fresh sequence number SQN from the counter SQN_{HE}. The HE has some flexibility in the management of sequence numbers, but the requirements listed in 6.3.1 need to be fulfilled, in particular, the generation mechanism needs to support the re-synchronisation mechanism described in 6.3.2.2. Annex C.1 contains a detailed description of a mechanism to generate sequence numbers that satisfies all conditions.

b) The HLR/AuC generates an unpredictable challenge RAND.

c) The HLR/AuC then computes

i) a message authentication code for authentication MAC-A = $f1_K$(SQN || RAND || AMF) where f1 is a message authentication function;

ii) an expected response XRES = $f2_K$ (RAND) where f2 is a (possibly truncated) message authentication function;

iii) a cipher key CK = $f3_K$ (RAND) where f3 is a key generating function;

iv) an integrity key IK = $f4_K$ (RAND) where f4 is a key generating function;

d) If SQN is to be concealed, in addition the HLR/AuC computes an anonymity key AK = $f5_K$ (RAND) where f5 is a key generating function and computes the concealed sequence number SQN ⊕ AK = SQN xor AK.

e) Finally, the HLR/AuC assembles the authentication token AUTN = SQN [⊕ AK] || AMF || MAC-A and the quintet Q = (RAND, XRES, CK, IK, AUTN) and updates the counter SQN_{HE}.

An authentication and key management field AMF is included in the authentication token of each quintet. Example uses of this field are included in Annex F.

The concealment of the sequence number is optional. Concealment is recommended when sequence numbers are derived from counters whereby strong correlation exists between consecutive sequence numbers that are sent to the same user. In that the concealment is required to provide location and identity confidentiality. However, when time-based counters are used to derive sequence numbers from, this lowers the correlation considerably, and the concealment can safely be omitted.

6.3.4.2 Authentication and key derivation in the USIM

Upon receipt of a (RAND, AUTN) pair the USIM acts as follows:

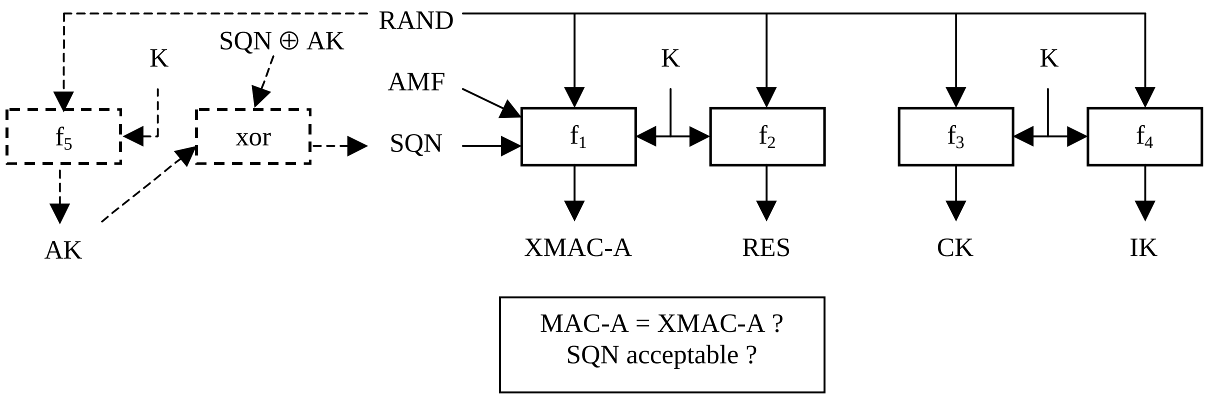

Figure 6.3.8: Authentication and key derivation in the USIM

a) If the sequence number is concealed, the USIM computes the anonymity key AK = $f5_K$(RAND) and retrieves the unconcealed sequence number SQN = (SQN ⊕ AK) xor AK.

b) The USIM then computes XMAC-A = $f1_K$ (SQN || RAND || AMF) and compares XMAC-A with MAC-A included in AUTN.

c) If they are different, the USIM triggers the ME to send back a *user authentication response* with indication of integrity failure to the VLR or SGSN and abandons the procedure. The remainder of this paragraph applies thus for the case where XMAC-A and MAC-A are equal.

d) Next the USIM verifies that the received sequence number SQN is acceptable. The HE has some flexibility in the management of sequence numbers, but the requirements listed in 6.3.1 need to be fulfilled, in particular, the verification mechanism needs to protect against wrap around and allow to a certain extent the out-of-order use of quintets. Annex C.2 contains a detailed description of a mechanism to generate sequence numbers that satisfies all conditions.

e) If the sequence number SQN is not acceptable, the USIM computes the re-synchronisation token AUTS as described in 6.4.3.3 and triggers the ME to send back a *user authentication response* back to the VLR or SGSN, with an indication of synchronisation failure, including the re-synchronisation token AUTS and abandons the procedure. The remainder of this paragraph applies thus for the case where SQN is acceptable.

f) The USIM then computes the response RES = $f2_K$(RAND) and triggers the ME to send back a user authentication response back to the VLR or SGSN, with an indication of successful receipt of the signed challenge and including the response RES.

g) Finally the user computes the cipher key CK = $f3_K$ (RAND) and the integrity key IK = $f4_K$ (RAND).

Note: If this is more efficient, RES, CK and IK could also be computed earlier at any time after receiving RAND.

6.3.4.3 Generation of re-synchronisation token in the USIM

Upon the assertion of a synchronisation failure, the USIM generates a re-synchronisation token as follows:

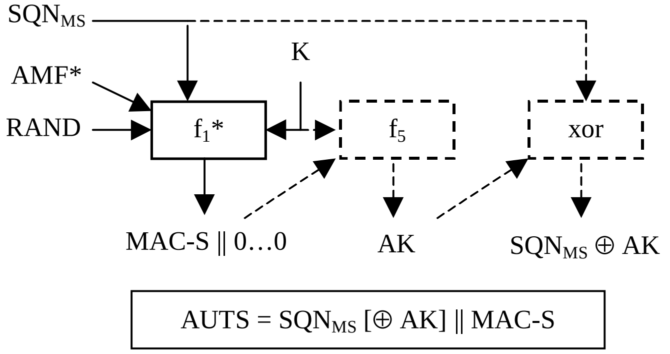

Figure 6.3.9: Generation of re-synchronisation token in the USIM

a) The USIM computes MAC-S = $f1^*_K(SQN_{MS} \| RAND \| AMF^*)$, whereby f1* is a message authentication function and whereby AMF* is a default value for AMF used in re-synchronisation.

b) If SQN_{MS} is to be concealed with an anonymity key AK, the USIM computes AK = $f5_K(MAC\text{-}S \| 0\ldots0)$ and the concealed counter value is then computed as $SQN_{MS} \oplus AK$.

c) The re-synchronisation token is constructed as AUTS = SQN_{MS} [$\oplus$ AK] || MAC-S.

6.3.4.4 Re-synchronisation in the HLR/AuC

Upon receipt of an indication of synchronisation failure and a (AUTS, RAND) pair, the HLR/AuC acts as follows:

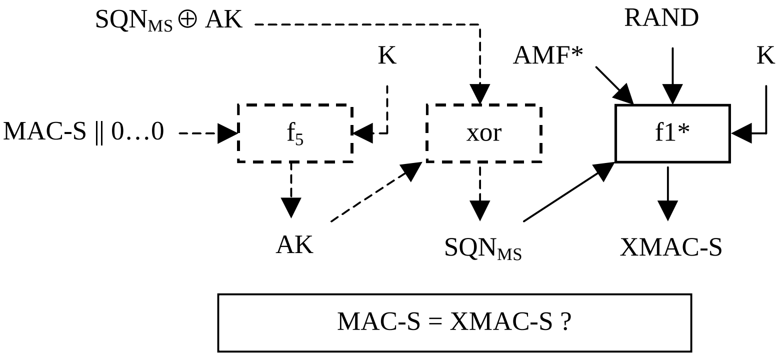

Figure 6.3.10: Re-synchronisation in the HLR/AuC

a) If SQN_{MS} is concealed with an anonymity key AK, the HLR/AuC computes AK = $f5_K(MAC\text{-}S \| 0\ldots0)$ and retrieves the unconcealed counter value as $SQN_{MS} = (SQN_{MS} \oplus AK)$ xor AK.

b) The HLR/AuC now verifies whether SQN generated from SQN_{HE} would be acceptable for a USIM that has SQN_{MS}. This test is identical to the test performed by the USIM described in 6.3.4.2. If SQN generated from SQN_{HE} would be acceptable, then the value of SQN_{HE} need not be modified and the function is aborted.

c) If SQN generated from SQN_{HE} would not be acceptable, then the HLR/AuC computes XMAC-S = $f1^*_K(SQN_{MS} \| RAND \| AMF^*)$, whereby AMF* is a default value for AMF used in re-synchronisation and the HLR/AuC then compares MAC-S and XMAC-S. If there is a match, the need to modify SQN_{HE} is recognised, otherwise again, it is decided that SQN_{HE} should not be modified.

Note: When a synchronisation failure is caused by an out-of-order use of a quintet, SQN_{HE} will be such that SQN generated from SQN_{HE} would be acceptable for a USIM that has SQN_{MS}. Therefore SQN_{HE} will not have to be modified and XMAC-S need not be computed. If SQN_{MS} is not concealed no cryptographic computation is required in this case.