3GPP TSG SA WG3 Security — S3#21b **S3z020026**

31st January - 1st February, 2002

Antwerp, Belgium

Source: **Hutchison 3G UK**

Title: **Use and Handling of SAs**

Document for: **Discussion/Decision**

Agenda Item: **IMS-8**

1 Introduction

This contribution suggests new text for the handling of SAs and a section stating which security association should be applied to which message.

The current text in TS 33.203 describing the handling of SAs was written assuming that SAs could only be updated by authenticated re-registrations of an IMPU, which took place one at a time.

Now there is going to be only one SA between the P-CSCF and UE. This SA can be updated by either authenticated registrations of previously unregistered IMPU or authenticated re-registrations of registered IMPUs. This means that two authentications, which would both change the SAs could be happening simultaneously. This would imply that there could be two new SAs created one after the other. Nothing in TS 33.203 currently describes how to deal with this.

Furthermore currently in TS 33.203, there is no clear statement of which SA should be used to protect traffic between the UE and P-CSCF. Most of the information prensented in this contribution is already contained in main body of TS 33.203, the appendices of TS 33.203 or suggested in an email to SA3 list. The aim of the contribution is to gather the information in one place, in order to make it easy to understand and keep it independent as possible of the integrity and confidentiality methods used.

2 Assumptions on SA handling and use

The following are the assumptions that were used in designing a SA use and renewal strategy for IMS

1. As many messages as possible should be integrity protected.
2. The UE can be involved in at most a (specified?) small number of attempted (re-) registrations at one time.
3. Until RES has been checked, the SAs should not be used to protect messages other than those in that particular registration flow (i.e. message carrying RES from UE to P-CSCF and message carrying registration success from P-CSCF to UE). It must be used to protect those messages.
4. There must be a pair of SAs that expires later than the registration of all IMPUs.
5. A message from a UE that is outside a registration flow (this include de-register messages) must be protected with an SA that was set up no earlier than the included IMPU was last authenticated or implicitly registered (this assumes every register/re-register ia authenticed).
6. A REGISTER message that results in an IMPU being registered must be integrity protected.
7. If the UE protects the first message in a registration procedure, all messages in that registration procedure must be protected.

Assumption 5 is very important as a successful authentication should update the keys used to protect the traffic from then onwards. The use of new security associations needs to be enforced by the network to avoid the attack given in the next section.

3 Problem of not using new SAs

Suppose we have a subscriber with private identity, IMPI and two public IMPU1 and IMPU2. Furthermore let us suppose that IMPU1 belongs to a service profile that has only “low value” services, whereas IMPU2 belongs to a service profile with “high value” services. Accordingly the S-CSCF registers IMPU1 for eight hours until the next authentication, while IMPU2 is only registered for one hour between authentications.

Now suppose the following sequence of events happens. The subscriber successfully registers (and authenticates) IMPU1. As a result of this the P-CSCF and UE share security association SA1 (strictly there is a pair, but for the purposes of this discussion it is enough to think of just one).

Two hours later the subscriber register (and authenticates) IMPU2 and hence sets up security association SA2 (as an aside, is SA2 replaces SA1 it must not be set to expire before SA1 was set to expire or the UE and P-CSCF could be left with no valid security association). From now the UE is expected to uses SA2 to secure traffic between it and the P-CSCF. There is no mechanism to force the UE to change security associations, as the P-CSCF will only change once it receives a message protected with SA2 from the UE.

The result of this is that keys that are over 2 hours old (in SA1) could be used to protect traffic (for IMPU2) that is supposed to be authenticated every hour.

The above assume a UE that is not functioning correctly as it should change to the new security associations after the second authentication. This is not the point, as if possible the network should not rely on the UE functioning as expected to securely deliver services.

The situation is also much worse if we assume that the keys in SA1 are somehow compromised (unknown to the network and UE). This would allow the attacker to sit between the UE and P-CSCF and access the services available to IMPU1 and IMPU2 as well as any other IMPU the subscriber registers before SA1 times out.

This is serious as a compromise of keys allows the attacker to carry on attacking the network even after a new authentication. If the attacker can block the traffic between the UE and P-CSCF, it will be difficult to stop them using the services available to IMPU1 but they should not get access the services available to IMPU2.

4 Which node should enforce update of SAs

There are two nodes that could enforce the update of SAs. Strictly it is not possible to force the UE use the new SAs, but it is possible to reject traffic that is not protected by the correct SA. In the above example the UE should be allowed to gain access to IMPU1 services by protecting the traffic with either SA1 or SA2, but IMPU2 services should only be accessed by traffic protected by SA2.

The P-CSCF and the S-CSCF both could decide whether a suitable security association was used to protect the IMS signalling.

If the S-CSCF trusts the P-CSCF to check the integrity, it should be able to trust it to ensure the correct security association was used to protect the traffic. To decide if a correct security association (or otherwise) was used to protect the traffic, the P-CSCF needs to know which IMPUs were registered, when each security association was established (CN1 has agreed on a method to transmit implicitly registered IMPUs from S-CSCF to P-CSCF, see see S3-020030). Then any security association established after the registration of an IMPU (note: could there be a problem, if during multiple registrations traffic arrives out of sequence) is a valid one to protect traffic with this IMPU (of course registration and re-registration messages may need to be treated differently way, although de-register message must be treated the same). This method seems to be directly applicable post-Release 5 when IMPUs from one subscriber could be registered at more than one S-CSCF. This is because the P-CSCF is still sent all the information relating to the security associations and registrations.

If the S-CSCF is to check a valid security association is used to protect the traffic. It needs the P-CSCF to pass to it information about the security association used to protect the traffic. This could have advantages in this would allow an IMPU to be registered using only one SIP REGISTER message or even allow INVITE with an unregistered IMPU (see section 8 for further explanation of this possible functionality). Post-Release-5 with more than S-CSCF, the information passed to the S-CSCF will need to consist of the length of time the SA has been used valid and possibly more), as the S-CSCF cannot know the parameter of all the SAs.

It seems that the P-CSCF is the most natural place to ensure that the new SAs are used. This is because it requires fewer amendments to the current flows given acceptance of the transfer of implicitly registered IMPUs from the S-CSCF to the P-CSCF.

The assumption that there should only be one SA between the UE and P-CSCF even when there is more than one S-CSCF is an interesting one. The situation means that one S-CSCF is relying on the security association that was

generated using an authentication via another S-CSCF. This means that a false S-CSCF could force the P-CSCF to use keys that it knows. This attack would not even require valid authentication vectors, if there were a colluding false UE to respond to the challenge. This contribution does not discuss the feasibility of mounting this attack, as it is a post-Release-5 issue.

5 SA use and renewal at the P-CSCF and UE

This section contains a method of ensuring that the correct SAs are used to protect traffic. It assumes that the CN1 suggestion of transferring the implicitly registered IMPUs for the S-CSCF to the P-CSCF is used and the P-CSCF is responsible for ensuring the correct SA is used.

When a UE receives the SM4 message (see TS 33.203 v1.0.0, page 20), it should have enough information to set a pair of SAs. These SAs will not be used for general traffic, until the registration flows are successfully completed, i.e. the UE has successfully received SM8. It is proposed that these are called registration SAs as there only use at this time is for registration. The expiry time of registrations SAs should be set to some short time (expiry timer only arrives in SM8 and the SAs can only be used for the registration of the associated IMPU, which should happen quickly). Registrations SAs need to be stored by UE, along with the associated IMPU. The number of pairs of registration SAs that a UE needs to store is no more than the number of registrations a UE can deal with simultaneously.

When a UE receive the SM8 message, it takes this SA as the current SA and stores the previous one to recieve inbound traffic as described in TS 33.203 v1.0.0 section 7.3.3.1.

When a P-CSCF receives the SM3 message, it should have enough information to set a pair of SAs. These SAs will not be used for general traffic, until the registration flows are successfully completed, i.e. the UE has successfully received SM8. It is proposed that these are called registration SAs as there only use at this time is for registration. The expiry time of registrations SAs should be set to some short time (expiry timer only arrives in SM8 and the SAs can only be used for the registration of the associated IMPU, which should happen quickly). Registrations SAs need to be stored by P-CSCF, along with the associated IMPU. The number of pairs of registration SAs that a P-CSCF needs to store is no more than the number of registrations a UE can deal with simultaneously.

Once the P-CSCF has received the message SM7 (registration complete) from the S-CSCF, it considers the currently negotiated pair of registation SAs to be valid for use between the UE and P-CSCF. Hence it is proposed to call these pairs of SAs valid. Once an SA is considered valid, it is no longer considered a registration SA.

Valid SA pairs are stored in the P-CSCF. The information held about each pair of SA is the following:

- SA_ID_U and corresponding SA information.
- SA_ID_P and corresponding SA information.
- IMPI.
- List of IMPUs that can be protected using the inbound SA.
- Time at which the pair of SAs became valid.

The list of IMPUs contains the IMPU used during the registration flows that created the pair of SAs, any IMPUs that were implicitly registered in the same flows and all the IMPUs that are in the list of older SAs with the same IMPI. When a new valid pair of SAs is created, the IMPU used during the registration flows that created it and all implicitly registered IMPU should be removed from the list of all older valid SAs. For an IMPI, one pair of valid SAs is the considered the current pair. If the P-CSCF runs out of space to store valid pairs of SAs, it should overwrite the oldest one.

The P-CSCF should use the current valid SA to protect all traffic towards the UE that does not require the registration SA. The P-CSCF should only accept an INVITE from and IMPU that was integrity protected with an SA for which that IMPU is valid. The number of non-current valid pairs of SAs that the P-CSCF needs to store is no more than the number of registrations a UE can deal with simultaneously. This would avoid any problems of the P-CSCF over-writing the pair of SAs a UE wants to use, as long as the UE always integrity protected the first message in a registration or re-registration.

Once the P-CSCF receives a SIP message that is integrity protected using a valid SA (note: this message cannot be the response in a message flow, as a pair of SAs do not become valid until after the SM7 message has been received). That SA becomes the current SA and the P-CSCF should delete all valid SA pairs that are older.

6 Proposed changes to TS 33.203

The proposed changes for TS 33.203 is in the attached document.

Firstly sections 7.3.3.1 and 7.3.3.2 are rewritten as a more general method to update SAs is needed because there can be simultaneous (re-) registrations affecting the update of SAs. The suggested method builds on the one already given. One main difference is the requirement for the P-CSCF to store a several pairs of SAs is several (re-) registrations happen simultaneously. This can not be avoided as there is not way to guarantee the last message in a (re-) registration flow will reach the UE. The second differnce is a clear separation between SAs that have been created by completed successful (re-) registrations and ones that are created from incomplete (re-) registrations. This avoid overwriting a usable pair of SAs with one that would get deleted if the (re-) registration attempt fails. As the re-written text applies equalling to authenticated registrations as authenticted re-registartion, the text can no longer be in section 7.3.3. It is proposed to put it into a new section 7.4 that will deal with the management and use of security associations.

Secondly another new section is proposed. This section describes exactly which SA should be used to protect the various kinds of traffic between the UE and P-CSCF. As previously stated most of this information contained in this section has already been given, but the aim here is to gather it together clearly in one place. It is also proposed to include this text in section 7.4.

7 Outstanding Issue with SA handling

There are a couple of potential problems with the suggested SA update method. The first problem is with the handling old SA in the UE. If the P-CSCF never recieved a message protected by a current otbound SA in UE, then the UE would make this current SA pair the old pair when a new current SA was created while the P-CSCF would not make the corresponding SA pair the current one. This means that P-CSCF would use an SA to protect traffic that the UE does not have. This problem exists with the current SA handling system.

The problem can be solved by making the UE keep a sequence of old pairs.

8 Further considerations

This section contains some discussions about further considerations. First it discusses the possibility of reducing the number of authentications by allowing registration of an IMPU, if the REGISTER message was integrity protected. Further discussion considers if an INVITE from a non-registered IMPU is received integrity protected whether it could be allowed.

Currently there is no method for accepting a subscriber registration without an authentication. This means that a subscriber must register and authenticate at least one IMPU from each service profile to fully register for its services. This means there is more than one authentication of the same IMPI in very quick succession. This is a waste of both network resources and authentication vectors. A potentially worse situation is if two IMPUs from the same service profile are registered in quick succession (not all IMPUs from a service profile have to be implicitly registered). The “authentication” of an IMPU could be piggy-backed onto the integrity protection at the P-CSCF, if there was a mechanism for the P-CSCF to inform the S-CSCF of details of the security association used to protect that REGISTER message. This requires putting no more trust into the P-CSCF than there is already. This is being discussed in CN1, see S3-020029.

The same method as suggested above could be extended to allowing INVITEs from unregistered IMPUs (assuming of course that at least one IMPU belong to the same IMPI is registered). The P-CSCF could pass the integrity protection details onto the S-CSCF as above. If the S_CSCF would have been happy to register that IMPU, it should be OK to accept an INVITE from that IMPU as well. This leads onto to issues like can an INVITE kick-off an authentication etc. and also may break some of the service assumptions, e.g. needing to be registered before using INVITE. A colleague tells me that this would be a “more natural” way to use SIP.

8 Conclusions

This contribution proposes some changes to TS 33.203. These changes aim to generalise the SA handling process to allow for simultaneous registrations and re-registration that might update the SAs used between the UE and P-CSCF. They also aim to clearly state which SA should be used to protect what traffic between the UE and P_CSCF. The proposed changes also discuss the of ensuring the UE uses the correct SA to protect traffic and the expiry times of SAs at both P-CSCF and UE.

SA3 is asked to approve the proposed changes to TS 33.203.

********** FIRST MODIFIED SECTION **********

6.1 Authentication and key agreement

The scheme for authentication and key agreement in the IMS is called IMS AKA. The IMS AKA achieves mutual authentication between the ISIM and the HN, cf. Figure 3. The identity used for authenticating a subscriber is the private identity, IMPI, which has the form of a NAI, cf. [3]. The HSS and the ISIM share a long-term key associated with the IMPI.

For the IMS the ISIM and the HSS keeps track of the counters SQN_{ISIM} and SQN_{HSS}~~,~~ respectively. The handling of the SQN can be as in [1]. The HN shall choose the IMS AKA scheme for authenticating an IM subscriber accessing through UMTS. The security parameters e.g. keys generated by the IMS AKA scheme are transported by SIP and embedded in EAP, cf. [7]-[9].

[Editors Note: Shall the HN choose EAP AKA for 3GPP-access or is it to be an option for the HN to choose either EAP AKA or perhaps any other mechanism e.g. HTTP digest depending on policy?]

The generation of the authentication vector AV that includes RAND, XRES, CK, IK and AUTN shall be done in the same way as specified in [1]. For each user it is the HSS that keeps track of the counter SQN_{HSS}. The requirements on the SQN handling both in the Home Network i.e. the HSS and the ISIM are specified in [1]. The AMF field can be used in the same way as in [1].

Furthermore a security association is established between the UE and the P-CSCF. The subscriber may have several IMPUs associated with one IMPI. These can ~~and~~ belong to the same or different service profiles. Only one SA shall be active between the UE and the P-CSCF. This single SA shall be updated when a new successful authenticated registration or re-registration has occurred, cf. section 7.~~3.3~~4.1. It is the policy of the HN that decides if an authentication shall take place for the registration of different IMPUs e.g. belonging to same or different service profiles.

********** NEXT MODIFIED SECTION **********

~~#### 7.3.3.1 Handling of security associations in authenticated re-registrations (successful case)~~

~~Before re-registration begins the following SAs exist:~~

~~- SA1 from UE to P-CSCF~~

~~- SA2 from P-CSCF to UE~~

~~The re-registration then is as follows:~~

~~1) The UE sends SM1 to re-register with the IMS, using the existing SA1 to the P-CSCF. As in the case of a new registration, a list of parameters to be negotiated in a security association set-up is included.~~

~~*[Editors Note: It is FFS if the SA1 shall be used for SM1 or not]*~~

~~2) The P-CSCF waits for the response SM3 from the S-CSCF and then sends SM4 to the UE, using SA2. As in the case of a new registration, the parameters selected for the new security associations are included. The P-CSCF then creates two new security associations, in parallel to the existing ones, in its database:~~

~~- SA11 from UE to P-CSCF~~

~~- SA12 from P-CSCF to UE~~

~~3) If SM4 could be successfully processed by the UE, the UE also creates the new SAs SA11 and SA12 in its database. The UE then sends SM5 to the P-CSCF. As in the case of a new registration, the authentication response and the list of parameters repeated from message 1 are included. SM5 is protected with the new SA11.~~

~~4) The P-CSCF waits for the response SM7 from the S-CSCF and then sends SM8 to the UE, using the new SA 12.~~

~~5) After the reception of SM8 by the UE, the re-registration is complete.~~

~~The UE now uses the new SAs for all subsequent messages. The old (outbound) SA1 is deleted. The old (inbound) SA2 must be kept until a further SIP message protected with the new inbound SA12 is successfully received from the P-CSCF.~~

~~The P-CSCF keeps all four SAs with the UE active until a further SIP message protected with the new inbound SA11 is successfully received from the UE. In the meantime, the P-CSCF continues to use the old SA2 for outbound traffic to the UE.~~

~~7.3.3.2 Error cases related to authenticated re-registration~~

~~Whenever an expected message is not received after a time-out the receiving entity considers the registration to have failed. The receiving entity then deletes any new security associations it may have established and continues to use the old ones if they have not yet expired.~~

~~If the registration protocol goes well up to the last message SM8, and SM8 is sent by the P-CSCF, but not received by the UE , then the UE has only the olds SAs available (after the time-out), but the P-CSCF cannot know this. Therefore, the P-CSCF continues to use the old SA2 for outbound traffic to the UE, but keeps both, old and new SAs. The new SAs are deleted when a message is received from the UE which is protected with the old SA, or if a REGISTER message is received on the port where the P-CSCF accepts specific unprotected messages.~~

********** NEXT MODIFIED SECTION **********

7.4 Management and Use of Security Associations

Every authenticated (re-) registration procedure potentially produces a new pair of security associations. This new pair of SAs should only be used to protect some message in its associated registration procedure until the authenticated (re-) registration is complete. If the (re-) registration procedure was unsuccessful, the SAs should be deleted. If it was successful, the pair of SAs can then be used to protect general signalling and replace any previous SAs. The following sections describe how the UE and P-CSCF use and manage SAs.

7.4.1 Management of security associations

7.4.1.1 Management of security associations in the UE

In the UE, the SAs are considered to be in the different states defined below:

- a **Registration** pair of SAs is a pair that has been created by a (re-) registration procedure for which the authentication challenge has been successfully received but the (re-) registration procedure is incomplete.

- the **Current** pair of SAs is the pair that was that was created by the most recent successful authenticated (re-) registration based on the CSeq number of the last message in the (re-) registration procedure.

- the **Old** pair of SAs is the pair of SAs is the pair, whose outbound SA was last used to protect a message outside the (re-) registration procedure that created that pair of SAs. The Old pair of SAs is different from the current pair (note strictly with the Old pair of SAs, only the inbound SA needs to be kept).

Registrations SAs are used only to protect traffic during the (re-) registration procedure that created them. They should be given a short lifetime and deleted if there is some failure in the (re-) registration procedure. The UE needs to store one registration pair of SAs for each registration and re-registration the UE is currently involved in.

Upon the successful completion of an authenticated (re-) registration, the pair of SAs created during that (re-) registration procedure becomes the current SA pair at the UE, if it is the most recent pair of SAs based on the CSeq number of the last message in the (re-) registration procedure.

The expiry time of the Current SA must be kept later than the expiry time of all registered IMPUs,

If the Current SA pair is changed, the original Current pair become the Old pair, if the outbound SA was ever used to protect a message outside the (re-) registration procedure that created it or there is no Old pair.

If the UE ever receives a message protected by the Current inbound SA, it can delete the Old SA pair.

7.4.1.2 Management of security associations in the P-CSCF

In the UE, the SAs are considered to be in the different states defined below:

- a **Registration** pair of SAs is a pair that has been created by a (re-) registration procedure for which the authentication challenge message has been successfully received by the P_CSCF but the (re-) registration procedure is incomplete.

- a **Valid** pair of SAs is the pair that was that was created by a successful authenticated (re-) registration.

- the **Current** pair of SAs is the pair whose inbound SA was last used by the UE to protect a message outside the registration flow that created that pair of SAs (note the Current pair of SAs is considered Valid).

Registrations SAs are used only to protect signalling during the (re-) registration procedure that created them. They should be given a short lifetime and deleted if there is some failure in the (re-) registration procedure. The P-CSCF should be capable of storing as least as many pairs of registration SAs as the total number of registration and re-registration the UE is capable of being simultaneously involved in.

Once the P-CSCF has received an authenticated registration successful response, the P-CSCF considers the pair of SAs created by the registration procedure to be Valid. If there is no Current pair of SAs, this pair becomes the Current pair.

The P-CSCF stores the Valid SA pairs from a UE in order based on the CSeq number. The expiry timer of the new pair of SAs should be set to the maximum of the expiry time given in the successful registration message and the expiry time of all stored older Valid SAs. This must be enforced even if messages arrive out of order.

Each Valid pair of SAs contains a list of IMPUs whose traffic can be protected with that SA. The list of IMPUs contains the IMPU used during the registration procedure that created the pair of SAs, any IMPUs that were implicitly registered at the same and all the IMPUs that are in the list of older SAs with the same

IMPI. When a new Valid pair of SAs is created, the IMPU used during the registration procedure that created it and all implicitly registered IMPU must be removed from the list of all older Valid SAs

If it receives a message protected with a the inbound SA of a Valid pair, the P-CSCF make that Valid pair of SAs, the Current pair of SAs and deletes all the older Valid pairs of SAs. If the P-CSCF runs out of space to hold Valid pairs of SAs, it should delete the oldest pair, which is not Current, in order to store the new one. To be robust the P-CSCF needs to be able to store as many Valid pairs of SAs as the total number of registration and re-registration the UE is capable of being simultaneously involved plus one (to allow for the Current one).

7.4.2 Correct Use of the Security Associations to Protect Signalling

7.4.2.1 Signalling from UE to P-CSCF

The UE must protect SIP signalling towards the P-CSCF according to the following rules:

- Requests or Responses, that are not REGISTER requests must be protected with the Current SA.
- A REGISTER request that has its expiry time set to 0 (i.e. a de-register message) must be protected with the Current SA.
- REGISTER request carrying a RES must be protected by the Registration SA created during that (re-) registration procedure.
- Other REGISTER requests can be protected with either the Current SA or not protected. In general the UE should try to protect these requests. Furthermore if the first message in a (re-) registration procedure is protected, all further messages must be protected.

On receiving a Request or Response, which is not a REGISTER request, the P-CSCF must ensure that the SA used to protect was a Valid SA with the given IMPU in its list. If the wrong SA was used to protect the message or there was no SA applied, the P-CSCF must discard the message.

On receiving a REGISTER request, the P-CSCF does the following:

- If it is a de-register message, i.e. expiry timer set to 0, then the P-CSCF must ensure that it was protected with a Valid SA with the given IMPU in its list.
- If it carries a RES, then the P-CSCF must ensure that it was protected with the Registration SA created during that (re-) registration procedure.
- Otherwise, the P-CSCF must ensure that it was protected with a Valid SA if the first message in the registration procedure was protected or if not and it was protected, it was done with a Valid SA for that UE.

If the wrong SA was used to protect the message or no SA was used when one was required, the P-CSCF must discard the message.

7.4.2.2 Signalling from P-CSCF to UE

The UE must protect SIP signalling towards the P-CSCF according to the following rules:

- Messages outside (re-) registration flows must be protected with either the Current SA or if there is not a Current SA, the most recent Valid SA.
- A successful registration acknowledgement must be protected with the Registration SA created during that (re-) registration procedure.

- Other messages in (re-) registration procedure must be protected with the Current SA if it exists and the UE initiated the (re-) registration with a protected message.

On receiving a message outside a (re-) registration procedure, the UE must ensure it was protected with either the Current SA or an Old SA. If the wrong SA or no SA was used to protect the message, the UE must discard the message.

On receiving a successful registration acknowledgement, the UE must ensure it was protected with the Registration SA created during that (re-) registration procedure. If the wrong SA or no SA was used to protect the message, the UE must discard the message.

On receiving any other message from a (re-) registration procedure, the UE must ensure it was protected with either the Current SA or an Old SA if it initiated the (re-) registration procedure with a protected message or no protection otherwise. If the wrong SA or no SA was used when it should have been, the UE must discard the message.