3GPP TSG SA WG3 #22bis aSIP ad hoc
Fort Lauderdale, Florida USA
8th April 2002

Tdoc S3z020054

Source: **Hutchison 3G UK**

Title: **SA handling**

Agenda item: **7.1**

Document for: **Discussion**

1 Introduction

The aim of this paper is to discuss a problem in the current SA handling procedures and propose some new text for TS 33.203 to overcome the problem.

2 Problem with current SA handling procedures

The identified problem with the current SA handling procedures is the fact that it is possible for the P-CSCF to end up with three security association pairs for a UE and nothing in the procedures defines how the P-CSCF deals with this situation. Figure 1 gives the security association set-up flows from section 7.2 of TS 33.203 for reference.

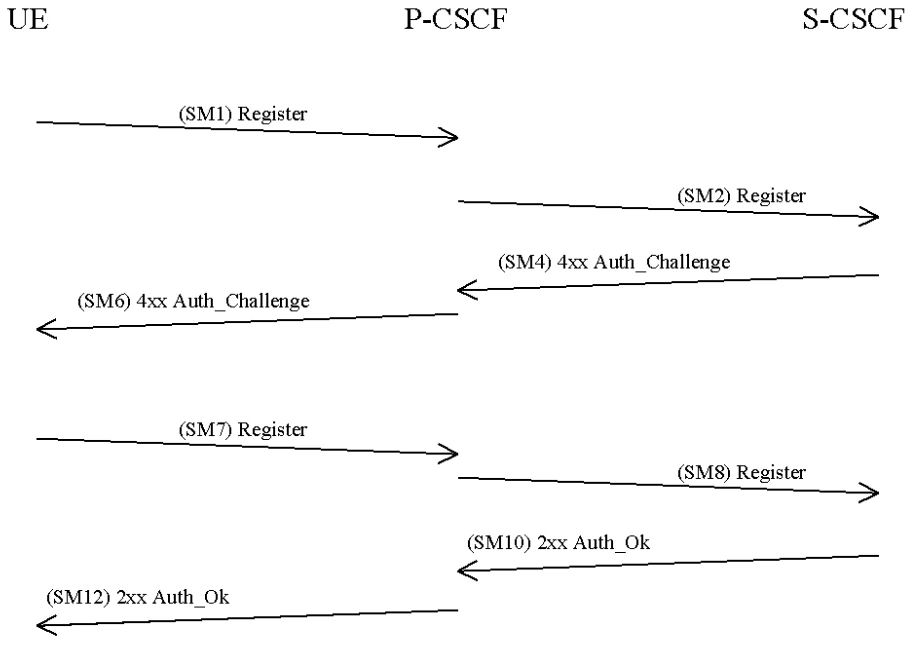

Figure 1: Security Association set-up flows from section 7.2 of TS 33.203

The situation occurs under the following circumstances. Assume that the UE has performed a successful registration with user authentication. This creates a pair of SAs at the P-CSCF, SA1_d (the downlink SA) and SA1_u (the uplink SA) to be used between the UE and P-CSCF. Assume then that at a later time there is another registration with user authentication that creates a new pair of SAs, SA2_d and SA2_u at the P-CSCF. These will be created after the P-CSCF receives SM4 and kept if the P-CSCF receives SM10 in the above flows. The P-CSCF does not throw the old SAs away at this stage, in case the UE does not receive SM12 (see section 7.3.3 of TS 33.203). This means the P-CSCF now has two sets of SAs.

It is then possible for an unprotected REGISTER to arrive at the P-CSCF. This will be forwarded to the S-CSCF, which will respond with a challenge. On receipt of this challenge (SM4), the P-CSCF will create another set of SAs, SA3_d and SA3_u. Nothing in the current procedures states how to deal with this situation.

The unprotected REGISTER could have been sent under the following circumstances (among others):

- A UE that wants to re-register some IMPUs but does not have a current SA e.g. due to being powered off.

- A malicious attacker sending the REGISTER, as none of the data needed in the REGISTER is strictly secret.

The need for the P-CSCF to be able to cope with a third set of SAs is due to the fact that the SA based on parameters received in SM4 is used to protect message SM7, before authentication is complete (i.e. the successful receipt of message SM8 by S-CSCF).

This means that either the P-CSCF must be able to either store three pairs of SAs or delete one or more of the three that it has. The following paragraphs detail the issues involved in either keeping all the SAs or deleting one particular set.

Clearly the P-CSCF must keep the SA3_d and SA3_u, in case it is a genuine attempt by a UE to register. Hence deleting SA3_d and SA3_u is not a viable possibility.

Suppose the P-CSCF deletes SA2_d and SA2_u. If the unprotected REGISTER came from attacker, the P-CSCF will have deleted the new SAs that the UE will use (assuming it successfully received SM12 of the 2nd registration). Hence deleting SA2_u and SA2_d is not a viable possibility.

Suppose the P-CSCF deletes SA1_d and SA1_u. If the third unprotected REGISTER came from an attacker, the UE and P-CSCF will not share a pair of SAs if the UE has not received SM12 of the 2nd registration, i.e. the UE will discard SA2_u and SA2_d if the 2nd registration does not complete successfully. Furthermore any message that the UE has sent protected with SA1_d will be discarded as the P-CSCF has discarded SA1 pair. It is normal behaviour for the UE to use SA1_d until it receives SM12. Hence deleting SA1_u and SA1_d will cause some traffic to be discarded under normal behaviour of the system and further problems if the SM12 of the 2nd registration is lost.

An alternative is to assume that the P-CSCF keeps all three sets of SAs. If the 3rd registration procedure completes successfully, then the P-CSCF can delete SA1_d, SA1_u, SA2_d and SA2_u, as the UE has none of these because it did not protect the initial REGISTER in the 3rd registration (this assumes that the UE is mandated to protect a REGISTER, if it has SAs that it believes are valid). If the registration procedure fails, it can delete SA3_d and SA3_u. Hence it will not need to keep all three pairs of SAs for long.

According to this analysis, the only possibility that does not introduce problems is to keep all three pairs of SAs.

3 Multiple Simultaneous REGISTERs

Another problem with the SA handling procedure is that it does not describe what to do in case there are several consecutive REGISTERs (either protected or unprotected) in a row. The S-CSCF might send a challenge in response to several of these REGISTERs, which would mean the P-CSCF would need to set up a SA for each one. This behaviour is currently not covered in the specifications.

From a UE perspective, there seems to be no value in performing several simultaneous registrations, as these might all need to be user authenticated. If they are done serially, only one should require user authentication (once one registration been authenticated, registration immediately after this one should not require a user authentication – as it should be protected using the SA from the previous registration). Hence it seems to be a benefit to limit a UE to be involved in one registration at a time (this seems to be inline with the latest SIP draft, see page 41 of draft-ietf-rfc2543bis-09).

Of course this does not stop an attacker sending a series of unprotected REGISTERs to cause the P-CSCF to have lots of unnecessary SAs and the S-CSCF to send out several AVs. It seems difficult to prevent a DoS attack to stop a subscriber performing an initial registration. It should be possible to perform a registration and/or authentication of an already registered user, as the P-CSCF and UE share an SA. Section 5 contains some proposed text to cover the behaviour of the P-CSCF and S-CSCF that receive multiple simultaneous registration requests.

4 Further Discussion on Modifications to the Current Text

This section discusses some further issues relating to proposed new text to TS 33.203.

Section 2 showed there is a possibility that the P-CSCF has to deal with three different pairs of SA at simultaneously. For this reason, the proposed text distinguishes between three types of SA; **current** SAs that are used to protect all outbound messages (that are protected) with the exception of SM12, **valid** SAs that have been created by a completed successful registration process but have yet to be used by the UE other than for the SM7 message and **registration** SAs that are created during a registration procedure, have been used to protect SM7 and SM12 are turned into valid SAs on successful completion of the registration procedure i.e. receipt of SM10.

A typical evolution of SAs at the P-CSCF will be as follows:

1. The UE starts a registration procedure and the P-CSCF subsequently receives a message containing a challenge. The P-CSCF creates the registration SA pair. It has only the registration SA pair.

2. The registration procedure is successful (SM10 received correctly). The registration SA pair becomes the current SA pair, as there is no current pair. The UE has only the current SA pair.

3. The UE starts a registration procedure (protected with an SA) and the P-CSCF subsequently receives a message containing a challenge. The P-CSCF creates the registration SA pair. It now has both registration and current SA pairs.

4. The registration procedure is successful (SM10 received correctly). The registration SA pair becomes the valid SA pair. The P-CSCF now has current and valid SA pairs.

5. When the P-CSCF receives a message protected with the inbound valid SA, it makes the valid SA pair the current SA pair. The P-CSCF has only the current SA pair.

Similarly at the UE, the current solution suggests having two types of SAs, old and new. This is not enough, as there are two states of the UE when it has the old and new SAs. Firstly during the registration procedure that produces the new SA, the old SA should be used for non-registration outbound traffic. Secondly when the registration procedure that produces the new SA is successfully completed, the new SA should be used to protect outbound non-registration traffic. There is a good technical reason for this behaviour at the UE. If the UE protects a non-registration message with the new SA before the end of the registration procedure that creates the new SA, the P-CSCF will delete the old SAs and on a registration failure the UE will delete the new SAs. This will leave the UE and P-CSCF without a common SA.

For this reason is it proposed to have three types of SA at the UE; **current** SAs that are used to protect all outbound traffic except SM7, **registration** SAs that are created during a registration procedure, used to protect SM7and SM12, and are turned into the current SAs on successful completion of the registration procedure and **old** inbound SA that is the previous current inbound SA and is kept until the P-CSCF protects some traffic with the UE's current inbound SA. Although there are three sorts of SA, the UE will never need to keep both old and registration SAs simultaneously.

A typical evolution of SAs at the UE will be as follows:

1. UE starts a registration procedure and subsequently receives a correct challenge. The UE creates the registration SA pair. It has only the registration SA pair.

2. The registration procedure is successful. The registration SA pair becomes the current SA pair. The UE has only the current SA pair.

3. UE starts a registration procedure (protected with the current outbound SA) and subsequently receives a correct challenge. The UE creates the registration SA pair. It now has both registration and current SA pairs.

4. The registration procedure is successful, the current inbound SA becomes the old inbound SA and the registration SA pair becomes the current SA pair. The UE now has the current SA pair and an inbound old SA.

5. When the UE receives a message protected with the inbound current SA, it deletes the old inbound SA. The UE has only the current SA pair.

5 Proposed Modifications to the Current Text

This section contains two proposed sets of new text for TS 33.203. Firstly there is some text proposed to explain how the P-CSCF and S-CSCF deal with simultaneous registrations. Some of the text may be more suitable for the Stage 3 document. The text is proposed to initiate discussion on how the network deals with multiple simultaneous registration requests.

Secondly, it is proposed to replace section 7.3.3.1 and 7.3.3.2 with a new section 7.4 to clarify the SA handling procedures. Section 7.4 is written assuming that all three sets of SAs are kept and the UE is only involved in one simultaneous registration at a time. SA3 are asked to accept the proposed text as a basis for a CR to TS 33.203.

************** PROPOSED CHANGE FOR SIMULTANEOUS REGISTRATION ****************

6.1.5 Simultaneous registration requests

The UE should only be involved in one simultaneous registration attempt. It is possible for the network to be involved in more than one simultaneous registration attempt for a particular IMPI because of an attacker or the UE terminating a registration attempt and starting a new before the network has done terminated the original one.

6.1.5.1 Behaviour of the P-CSCF

The P-CSCF shall forward all protected REGISTER to the S-CSCF. The P-CSCF can choose not to forward an unprotected REGISTER if it is already involved in a registration (protected or unprotected).

On receipt of a response carrying a challenge from the S-CSCF that relates to a protected REGISTER request, the P-CSCF shall forward the message to the UE, store the relevant registration information, i.e.

SA details etc. and delete all other registration attempt information relating to that IMPI (but **not** any existing SAs).

On receipt of a response carrying a challenge from the S-CSCF that relates to a unprotected REGISTER request, the P-CSCF shall forward the message to the UE, store the relevant registration information, i.e. SA details etc. if it is not already involved in a protected registration with that IMPI. If it is already involved in a protected registration, it can either ignore the message (e.g. the P-CSCF assumes that there is an attacker) or forward the message to the UE, store the relevant registration information, i.e. SA details etc. In this case it must not overwrite the details of a protected registration.

6.1.5.2 Behaviour of the S-CSCF

The S-CSCF must allow at least one simultaneous registration attempt for an IMPI. It is operator choice whether the S-CSCF allows more than one. The S-CSCF can choose not to respond to an unprotected REGISTER request for an IMPI that already has at least one registration in process. The S-CSCF shall respond to a protected REGISTER request. After responding to the protected REGISTER request, it shall delete all other registration attempt information relating to that IMPI.

*********************** PROPOSED NEW TEXT FOR SA HANDLING ***********************

7.4 Management and Use of Security Associations

Every successful registration or re-registration procedure that includes a user authentication produces a new pair of security associations (SAs). These new SAs shall then replace the previous SAs. This section describes how the UE and P-CSCF shall handle this replacement and which SA to apply to which message.

7.4.1 Management of security associations

Security associations may be unidirectional or bi-directional. This section assumes that security associations are unidirectional, as this is the general case. Under normal behaviour for release 5, at any one time, there shall be at most two security associations stored at a UE, and at most three security associations per UE stored at a P-CSCF, for each direction.

Whenever a user is registered there is a **current SA** for each direction. In addition there may be either a **registration SA** for each direction or an inbound **old SA** at the UE and either a **registration SA** or a **valid SA** for each direction or rarely both at the P-CSCF. Their use is explained in the following. They are denoted as follows:

SA_in_cur	current inbound SA
SA_out_cur	current outbound SA
SA_in_reg	registration inbound SA
SA_out_reg	registration outbound SA
SA_in_old	old inbound SA (in UE only)
SA_in_val	valid inbound SA (in P-CSCF only)
SA_out_val	valid outbound SA (in P-CSCF only)

This notation has local significance only. That means that SA_in_cur at the UE is not always the same as SA_out_cur at the P-CSCF, and similarly for the other SAs.

The SAs shall be distinguished by different SA_IDs.

7.4.1.1 Management of security associations in the UE

The UE shall be involved in only one authenticated registration procedure at a time. Upon starting a new registration procedure, any existing registration SAs are deleted.

The UE shall delete any SA whose expiry time is exceeded.

When SA_out_cur exists, the UE shall it to protect all outbound traffic, except at the specified cases in the registration procedure, when the UE shall use SA_reg_out.

The UE shall ensure that all inbound traffic outside registration procedures is protected with either SA_in_cur or SA_in_old. Traffic inside registration procedures shall be protected with SA_in_cur, if the initial REGISTER was protected and not protected otherwise, except where stated in the registration flows when they shall be protected with SA_reg_in. If the wrong SA is used, the message should be discarded.

When a SIP message protected with SA_in_cur is successfully received from the P-CSCF, the UE shall delete SA_in_old if it exists.

A successful registration with authentication proceeds in the following steps:

- The UE sends the initial message to register with the IMS. It should be integrity-protected using SA_out_cur if available.
- The UE receives an authentication challenge in a message from the P-CSCF. This message shall be integrity-protected using SA_in_cur if the UE's initial message was integrity-protected.
- If this message can be successfully processed by the UE, the UE creates the new SAs, SA_in_reg and SA_out_reg, which are derived according to section 7.2. The UE then sends its response to the P-CSCF, which shall be protected with SA_out_reg.
- The UE receives a registration successful message from the P-CSCF, which shall be protected using SA_in_reg.
- After the successful processing of the registration successful message by the UE, the registration is complete. SA_in_cur becomes the new SA_in_old, SA_out_reg becomes the new SA_out_cur and SA_in_reg becomes the new SA_in_cur.

A failure in the registration means the UE should delete SA_in_reg and SA_out_reg.

7.4.1.2 Management of security associations in the P-CSCF

The P-CSCF shall delete any SA whose expiry time is exceeded. If the current SAs are deleted and there exist valid SAs, then the P-CSCF makes the SA_out_val the new SA_out_cur and SA_in_val the new SA_in_cur, and removes the valid SAs.

To protect outbound traffic, the P-CSCF shall always use SA_out_cur if it exists, except in the registration procedure where it shall use SA_out_reg where specified and no protection if the initial message in the registration procedure was sent unprotected by the UE.

The P-CSCF shall ensure that inbound traffic is protected with either SA_in_cur or SA_in_val, except in registrations procedure where it shall ensure SA_in_reg was used where specified and other messages shall be unprotected if the initial message in the registration procedure was sent unprotected by the UE.

When the P-CSCF successfully receives a SIP message protected with SA_in_val from the UE, then SA_in_val and SA_out_val becomes the new SA_in_cur and SA_out_cur respectively, and there are no more valid SAs.

The P-CSCF associates the IMPI given in the registration procedure and all the successfully registered IMPUs related to that IMPI with the SAs.

There are two possible registration cases, one when the first message is integrity protected and one when it is not.

A successful registration with the first message protected proceeds in the following steps:

- The P-CSCF receives the initial register message. It should be integrity-protect using SA_in_cur or SA_in_val.

Editor's note: there can exist three SAs here. This can be avoided the addition of "If SA_in_cur is used to protect the message, then the P-CSCF deletes SA_in_val and SA_out_val."

- The P-CSCF forwards the message containing the challenge to the UE. This shall be integrity-protected using SA_out_cur.
- The P-CSCF then creates the new SAs, SA_in_reg and SA_out_reg, which are derived according to section 7.2.
- The P-CSCF receives the message carrying the response from the UE. It shall be protected using SA_in_reg.
- The P-CSCF sends registration to the UE, which shall be protected using SA_out_reg. This completes the registration procedure for the P-CSCF. SA_out_reg becomes SA_out_val and SA_in_reg becomes SA_in_val (overwriting any previous valid SAs).

A successful registration with the first message unprotected proceeds in the following steps:

- The P-CSCF receives the initial register message.
- The P-CSCF sends the message containing the challenge to the UE.
- The P-CSCF then creates the new SAs, SA_in_reg and SA_out_reg, which are derived according to section 7.2.

Editor's note: there can exist 3 sets of SAs at this point.

- The P-CSCF receives the message carrying the response from the UE. It shall be protected using SA_in_reg.
- The P-CSCF sends registration to the UE, which shall be protected using SA_out_reg. This completes the registration procedure for the P-CSCF. SA_out_reg becomes SA_out_cur and SA_in_reg becomes SA_in_cur, and there all valid and registration SAs are deleted.

A failure in the registration means the P-CSCF should delete SA_in_reg and SA_out_reg.

Editor's note: What happens to the expiry times of SAs at the P-CSCF.