3GPP TSG SA WG3 Security — S3#29 **S3-030415**

15 – 18 July 2003

San Francisco, USA

Source: **Siemens**
Title: **Discussion paper on solutions regarding the behaviour of SIP over TCP and SA handling in re-authentications**
Document for: **Discussion**
Agenda Item: **7.1 IMS**

Abstract

This contribution evolved from the discussion paper sent by Siemens to the SA3 list. It is meant to explain the proposed changes implemented in two accompanying change requests.

---

1. Introduction

Two issues with IMS security, as currently specified in TS 33.203 v550, have been brought up in 3GPP SA3 recently. The problems are stated in the second section of the paper, and solutions are proposed in the third section.

The two issues are:

- While the current specification, which assumes one IPsec security association between UE and P-CSCF in each direction, seems to work fine for SIP over UDP, the same may not hold for SIP over TCP due to the different communication behaviour in the two cases, as specified in RFC 3261. Two IPsec security associations between UE and P-CSCF in each direction may be needed in the TCP case. A secondary issue related to SIP over TCP behaviour is the number of ports at the UE.
- In an authenticated re-registration, the UE port changes according to the current specification. It is stated in S3-030258 (Lucent) that potentially many other SIP entities may need to be informed about this port change. This would clearly be undesirable.

2. Problem Statement

2.1 Different behaviour of SIP over TCP and UDP

SIP can be used with TCP and UDP. In some respects, the behaviour of SIP depends on the transport protocol. When used with UDP, the typical behaviour is as follows:
the P-CSCF listens on a port well-known to all the UEs registered with this P-CSCF. In fact, in IMS there are two such ports, an unprotected port (e.g. 5060) to receive unprotected REGISTER messages, and a protected port to receive protected messages. The protected port was established during the registration procedure. When the P-CSCF sends a (protected) SIP message to the UE it typically uses a port different from the protected port (which we call server port here) on which the P-CSCF listens. Nevertheless, the UE sends the response back to the server port at the P-CSCF which was specified in the via header, and not to the port at the P-CSCF where the request came from. In this way, all protected messages from a particular UE to the P-CSCF can be received at a fixed port at the P-CSCF, and all protected messages from the P-CSCF to a particular UE can be sent from another port (which may be assumed fixed for that UE, as negotiated in the registration procedure). So, one uni-directional security association (SA) is needed to protect messages from the P-CSCF to the UE and another uni-directional SA is needed to protect messages in the opposite direction.

This does not work with TCP. It is specified in RFC3261, section 18.2.2, that a SIP response has to use the same TCP connection as the corresponding SIP request. If the P-CSCF sends a request to UEs over TCP from a port different from the server port (which we assume it typically does for implementation reasons), then the corresponding response from the UE is also received on that port. On the other hand, a request over TCP is sent from the UE to the server port at the P-CSCF, and then also the response has to be sent from the server port at the P-CSCF. Consequently, protected messages may be going in both directions over two different ports at the P-CSCF. As, unfortunately, the current IPsec specification [1](in RFC2401) does not allow a list of ports as selector associated with an IPsec SA, and RFC2401 does not seem entirely clear about whether it must be possible to map two different entries of the SPD (security policy database) onto one entry of the SAD (security association database), two different IPsec SAs would be needed in each direction to protect the SIP messages between the P-CSCF and a particular UE.
But TS 33.203 v550 currently specifies only one IPsec SA per direction. Therefore, it appears that the behaviour of SIP over TCP as described above is in conflict with the current TS 33.203.

Another issue related to the behaviour of SIP over TCP concerns the number of ports at the UE. TS 33.203 v550 specifies that the same port is used for sending and receiving, for both TCP and UDP. For TCP, however, the typical behaviour [2] is that requests are sent from a port, used by the User Agent Client, different from the listening port, used by the User Agent Server.

2.2 Change of ports in authenticated re-registrations

In an authenticated re-registration, the UE port changes according to the current specification. It is stated in S3-030258 (Lucent), which was submitted to SA3#28, that potentially many other SIP entities may need to be informed about this port change. This would clearly be undesirable.
The discussion on the SA3 mailing list concluded that there really is a problem.

3. Proposed approaches to solutions

In the following, we always distinguish between a client port and a server port. In general, client and server port are different. These ports are used with both, UDP and TCP. The use of these ports may differ for TCP and UDP, as follows:
UDP case: the UE receives requests and responses protected with ESP on the port *port_us* (the"protected server port"). The UE sends requests and responses protected with ESP on the port *port_uc* (the "protected client port"). For every protected request towards the P-CSCF, the UE shall insert the protected server port *port_us* into the Via header. The protected responses are then sent to port_us.
TCP case: the UE receives requests and sends responses protected with ESP on the port *port_us*. The UE sends requests and receives responses protected with ESP on the port *port_uc* (the "protected client port").

An analogous description applies to the P-CSCF side.

We always assume that the same port numbers are used for UDP and TCP (in accordance with TS 33.203 v560).

Notation: *port_uc* = client port at the UE
port_us = server port at the UE
port_pc = client port at the P-CSCF
port_ps = server port at the P-CSCF

[1] A draft RFC2401bis has been submitted to the IETF. It allows lists of ports as selectors. However, it does not seem appropriate to base TS 33.203 on RFC2401bis as it may take a long while before we see products supporting RFC2401bis.
[2] An interesting document in this context is draft-ietf-sipping-connect-reuse-reqs-00.txt which was recently submitted to the IETF. It addresses the use of ports and TCP connections with SIP.

3.1 How to address the different behaviour of SIP over TCP and UDP

The problem described in section 2.1 can be solved by establishing two pairs of uni-directional[3] IPsec security associations (SAs) between UE and P-CSCF, instead of only one pair as in the current TS 33.203. One pair of SAs would be used to protect the communication between the client port at the UE and the server port at the P-CSCF, and the other pair would be used to protect the communication between the server client port at the UE and the client port at the P-CSCF.

The solution is illustrated in Figure 1 below with a set of example message exchanges protected by the respective IPsec SAs.

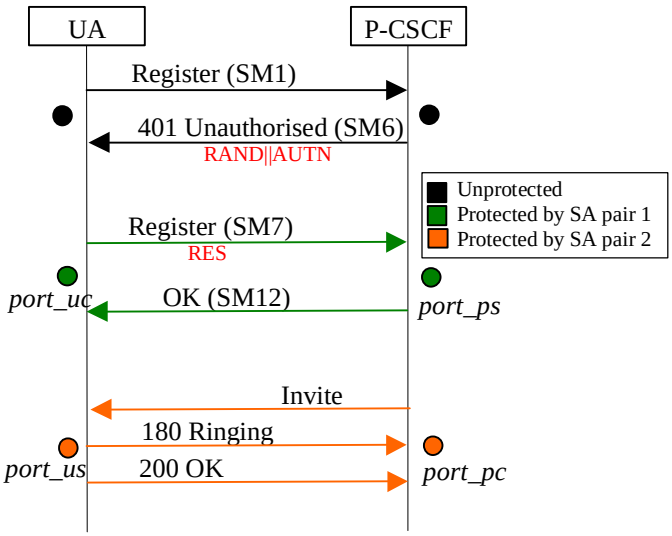

Figure 1

The following IPsec port selectors have to be used to configure the IPsec SAs at the UE and the P-CSCF respectively:

UAC:
SA1 outbound (src_port *port_uc*, dst_port *port_ps*)
SA1 inbound (src_port *port_ps*, dst_port *port_uc*)

UAS:
SA2 outbound (src_port *port_us*, dst_port *port_pc*)
SA2 inbound (src_port *port_pc*, dst_port *port_us*)

P-CSCF (server):
SA1 outbound (src_port *port_ps*, dst_port *port_uc*)
SA1 inbound (src_port *port_uc*, dst_port *port_ps*)

P-CSCF (client):
SA2 outbound (src_port *port_pc*, dst_port *port_us*)
SA2 inbound (src_port *port_us*, dst_port *port_pc*)

3.2 Mitigation of the man-in-the-middle attack against sip-sec-agree

A security problem, which is not present in the current specification, needs to be addressed in case the change proposed in section 3.1 is made: if there are two different ports at the UE, a client and a server port, then also two different ports need to be negotiated in the registration procedure (through sip-sec-agree[4]). Similarly, two different SPIs need to be negotiated. But in sip-sec-agree (RFC 3329),

[3] Note that IPsec security associations are unidirectional by definition.
[4] RFC 3329 “Security Mechanism Agreement for the Session Initiation Protocol (SIP)”.

the UE parameters are sent only in the unprotected REGISTER message (SM1 in TS 33.203, section 7.2) and are not explicitly confirmed in the protected REGISTER message (SM7 in TS 33.203, section 7.2). The idea of sip-sec-agree is that the UE parameters are implicitly confirmed by using them in the security association which protects SM7. But SM7 would be sent from the client port. The UE server port, if different, would not be implicitly protected as it would not be not used. Consequently, a man-in-the-middle could tamper with the UE server port. A similar problem exists with the SPIs.
For the solution proposed in section 3.1, the pairs of ports or SPIs have to be communicated to the other side in the registration procedure before the security associations can be set up. We propose the following: the pairs (*port_uc, port_us*) and *(SPI_uc, SPI_us)* are sent from the UE to the P-CSCF in SM1 (see Figure 1). The pairs (*port_ps, port_pc*) and *(SPI_ps, SPI_pc)* are sent from the P-CSCF to the UE in SM6. All ports and SPIs are sent again in SM7. SM7 is protected, so the ports and SPIs are confirmed to the P-CSCF. If they had been tampered with in SM1 or SM6 by an attacker, the P-CSCF would notice and would not send a 200 OK in SM12.

3.3 Change of ports in authenticated re-registrations

When changing security associations in an authenticated re-registration, there is a need[5] for the UE and P-CSCF to be able to distinguish whether the old or the new SA (or none) was used to protect a certain message (cf. TS. 33203, section 7.4). The only information available to the SIP application, which characterises an SA is the set of selectors, i.e. source and destination IP addresses, source and destination ports and transport protocol. The SPI is network layer information, and is not available to the application. As IP addresses and transport protocol remain the same in a re-registration, only the ports can be used to distinguish between the SAs.

In TS 33.203 it was decided to use the UE port (there is only one in TS 33.203 v560) to distinguish old and new SAs. Consequently, TS 33.203 v560 specifies that the UE port is changed in an authenticated re-registration. This may, however, lead to the problems described in S3-030258 (cf. section 2.2 above).

We propose the following solution:

Keep the protected server ports (port_us and port_ps) always fixed when changing SAs in a re-authentication, while varying the protected client ports (port_uc and port_pc) to distinguish between old and new SAs. As the end-points of an SA always include a client port (see section 3.1 and figure 1), a change of client ports always implies a change of the SA selectors. Hence the SA selectors can be used to distinguish between old and new SAs. With this solution, the remote SIP entity would not be affected any more, as it destines its SIP messages to the UE server port, which remains fixed. So, the problem raised in S3-030258 does not occur any more.

There is a potential problem with this solution, however. In case, the SA and the TCP connection have changed after the UE or the P-CSCF received a request and before it sends a response, the UE or the P-CSCF can no longer use the old TCP connection to send the response. There are two solutions to this problem:

1) keep old TCP connections and the associated SAs open until the transaction is completed. This, however, would complicate the SA handling. It would also mean that several TCP connections may have to be kept open in case of overlapping transactions. But please note that old and new TCP connections need to co-exist for some time anyhow, if ports are changed in re-authentications.

2) send the response back to the client port relating to the new SA and TCP connection. This would avoid the need to keep several TCP connections open in parallel longer than absolutely necessary. This would conflict, however, with text in RFC3261, section 18.2.2, which says how a SIP entity SHOULD behave in case the TCP connection on which the request was received was no longer available. But, according to the IETF rules, "SHOULD" means that a different behaviour is allowed if there are valid reasons. We believe that such valid reasons exist in this case.

In the accompanying CR, we implemented solution 1). But we believe that solution 2) is also acceptable.

[5] This is still under discussion on the mailing list at the time of writing

Conclusion

Two problems, associated with the use of SIP over TCP in IMS, and with the change of UE ports in re-authentications, are not addressed in the current version of TS 33.203. This contribution showed how the problems can be solved. Accompanying CRs implement the proposed solutions.