**3GPP TSG SA WG3 Security — S3#22** **S3-020082**

**25 - 28 February 2002**

**Bristol, UK**

**Source:** **Nortel Networks**

**Title:** **Updates to SIP-Level Solution for IMS Integrity Protection**

**Document for:** **Discussion/Decision**

**Agenda Item:** ***TBD*, IMS Access Security: SIP Integrity Protection**

**Abstract**

*This contribution shows updates to the description of the "SIP-level security solution" for IMS message integrity in draft TS 33.203. The updates reflect the work done since IETF 52 to enhance HTTP Digest such that it is a viable solution mechanism for SIP message integrity in the IMS. The updates also address inadequacies in the present 33.203 material in relation to the anti-replay protection feature of the SIP-level security solution. A previously identified problem with loss of calls due to counter synchronisation failures is addressed. Procedures for management of the values of the Digest 'nonce' and 'nonce-count' directives are adjusted.*

## 1. Updates from IETF

Discussion took place at IETF 52 regarding enhancing HTTP Digest [1] to provide more complete integrity protection in one-hop situations and to enable last-hop (proxy to User Agent Server) authentication to work properly. The SIP Working Group agreed that Digest authentication as is currently specified is inadequate and that enhanced Digest mechanisms may be required in SIP. A principal aspect of the work to propose such enhancements is represented by the Internet draft 'draft-undery-sip-auth-00.txt' [4], which is slated for discussion at IETF 53.

The Internet draft proposes to define a new SIP response, 492 *Proxies Unauthorized*, to permit origin servers to challenge proxies from which they receive SIP requests, and new SIP headers to permit proxies to authenticate themselves to origin servers. New values of the Digest "qop-options" directive direct the client that receives a challenge to apply integrity protection with an extended scope when computing the response.

Section 3 of this contribution proposes specific text adjustments to 33.203 section C.2 to reflect the work done since IETF 52 that relates to the "SIP-level security solution" for IMS message integrity protection.

## 2. Anti-Replay Protection: Counter Synchronisation and Reflection Attacks

It is possible to have anti-replay protection as a feature of the integrity protection of SIP messages (both requests and responses) travelling in either direction between the UE and Proxy CSCF. The combination of directives 'nonce' and 'nonce-count' from HTTP Digest [1] comprise the basic mechanism. However, the solution for IMS must accommodate the client-server model upon which HTTP Digest is based.

Document S3-010664 [2] described a problem with the anti-replay protection component of the SIP-level message integrity solution that is described in current draft TS 33.203. The problem is that counter synchronisation failures can occur in certain scenarios when a single nonce counter is incremented by both SIP entities that are participating in a Security Association; the result is a false detection of message replay. Investigation of the problem points out that it is not sufficient to have only a single instance of the nonce counter in an implementation of one-hop bi-directional message integrity protection. The description of the Digest 'nonce-count' directive dictates roles and responsibilities for both the client and server functions – the server issues the nonce and detects replays, and the client increments 'nonce-count' for each SIP request that is integrity protected using the nonce. When the

UA (UE) sends an INVITE or other request towards the Proxy (P-CSCF), the UA is the client and the Proxy is the server. When the Proxy sends (or re-submits) an INVITE towards the UA, the Proxy acts as the client and the UA acts as the server. To accommodate the Digest client-server model, then, a separate nonce counter must be maintained for protecting SIP requests that are sent in each direction.

It is also possible for a Man-in-the-Middle (MITM) to mount what is called a "reflection attack"; this is an attack in which a MITM intercepts a genuine message and then reflects it back to its sender. This type of attack was described in [3]. The Digest framework thwarts this attack by ensuring that a message will not pass the integrity check unless a server-issued nonce is used in the computation of the message 'digest' (integrity credential). It is worth mentioning the attack and its mitigation in the 33.203 description of the "SIP-level Security Solution".

The next section proposes specific text adjustments to 33.203 section C.2 to correct the treatment of anti-replay protection for the "SIP-level security solution". It is proposed to adjust the description of the procedures for anti-replay protection that operate in the UE and the P-CSCF, and to correct the example information flow that illustrates the usage of 'nonce' and 'nonce-count' in a one-hop bi-directional context.

## 3. Proposed Text Adjustments for 33.203 Section C.2

*[Editors note: There seems to be an unexpected shortcoming in the way SIP provides integrity protection on messages between UE and Proxies. In current SIP, HTTP Digest can be used to partially integrity protect the messages originated by an UE. However, SIP fails to provide integrity for Proxy to UE communication, i.e. for terminating INVITEs, for example. Proxies are not able to add Authorization headers on these messages, thus leaving the messages unprotected.*

*For the reason above, the headers and field names used in this section may not be final. However, the found inconsistency will probably make it easier for 3GPP to discuss about new SIP level integrity protection schemes with IETF.]*

HTTP Digest shall provide integrity protection of SIP signalling between the UE and the P-CSCF, protecting all SIP signalling messages at the SIP level.

### C.2.1 6.3.1 Security Association Setup

The SA that is required for Digest integrity protection shall use the 128-bit integrity key IK generated through IMS AKA, as specified in section 6.1. The integrity algorithm and key are identical for integrity protection applied to messages travelling in either direction. Negotiation of the integrity algorithm to use occurs in the following way: The UE communicates the set of integrity algorithms that it supports to the P-CSCF through the Security-setup header field of the REGISTER message, as described in section 7.2. The P-CSCF selects an algorithm to use from the set of algorithm capabilities common to both the UE and the P-CSCF. The P-CSCF indicates the algorithm to use in the "algorithm" directive of the Digest challenge that is subsequently issued to the UE.

### C.2.2 6.3.2 Scope of Integrity Protection

**Digest supports integrity protection of the SIP message body (not the headers) when the "qop-options" directive within the Digest challenge is set to the value "auth-int". Digest supports integrity protection of the SIP message body plus a named list of headers when the "qop-options" directive is set to the value "auth-hdr-int". Digest supports integrity protection of the entire SIP message when the "qop-options" directive within the Digest challenge is set to the value "~~extended~~auth-extd-int".** ~~(Use of either of these values of "qop-options" assumes that a context of client authentication has been previously established.)~~ To provide for protection of the entire SIP message, the P-CSCF shall issue a Digest challenge to the UE specifying the value "~~extended~~auth-extd-int" for the "qop-options" directive.

### C.2.3 6.3.3 Computation of Integrity Protection Credential

The message 'digest', or message authentication code, is conveyed in the "response" directive of the Digest response. The rules for computing "response" are as described in [1] with the following consideration: if the UE receives a Digest challenge with the "qop-options" directive set to either "int" or "~~extended-int~~auth-extd-int", and the associated authentication challenge was an IMS AKA challenge, then the UE substitutes IK for the "password" component of A1 when computing "response=" in the

Digest response. The UE sets the "username" component of A1 to a fixed value (e.g., "ims-user"). When sending messages to the UE that are to be integrity protected, the P-CSCF applies the same rules when computing "response". In this manner, the whole SIP message is always protected.

## C.2.4 6.3.4 Anti-Replay Protection

The Digest framework specifies that a server-initiated nonce is to be used by the client as a random number input to the production of the message digest. This nonce, along with a counter ('nonce-count') that is incremented by ~~either endpoint~~the client when sending ~~a message~~each SIP request that is to be protected, facilitate anti-replay protection. The anti-replay protection feature of the integrity protection mechanism is as described in RFC 2617 with the following considerations. Per RFC 2617, the role of the server is to issue the nonce and to detect replays (through validation of 'nonce-count'), and the client must increment 'nonce-count' when computing the digest for each new SIP request that is to be integrity protected. In the one-hop environment that exists for the UE and the P-CSCF in the IMS, both the UE and the P-CSCF may fill either the client or server role in particular operational situations. When the UE sends an INVITE or other request towards the P-CSCF, the UE is the client and the P-CSCF is the server. When the P-CSCF sends (or re-submits) an INVITE towards the UE, the P-CSCF acts as the client and the UE acts as the server. The implications of supporting the Digest client-server model, then, are that both the UE and the P-CSCF must: 1) be able to issue Digest challenges, which includes issuing nonces; and 2) maintain its own counter for the 'nonce-count' directive for use when operating in the client role.

New nonce values are communicated by the server to the client in two ways: 1) through the 'nonce' directive that is an obligatory part of the Digest challenge; and 2) through the 'nextnonce' directive that is an obligatory part of the Digest authentication of SIP responses (e.g., Authentication-Info header). Nonce values themselves are selected entirely by the server implementation – counter-based, clock- or other random number-based, and hybrid implementations are all possible. It is also a matter of server implementation how frequently new nonces are to be issued. To minimize the number of "stale" authentication attempts (generation of credentials by the client using an older nonce), the server should maintain a list of reasonable size of previously issued nonce values.

Expected behaviour of the UE and P-CSCF in relation to anti-replay protection is illustrated in the example information flow that follows in this section.

## C.2.5 6.3.5 Mitigation of 'Reflection Attacks'

When either the UE or P-CSCF receives a SIP request (i.e. is acting as Digest server), it expects the sending entity (acting as client) to use in the computation of the message digest a nonce that it (the server) has previously issued. If an unrecognized nonce appears in the Digest response, the receiving entity will deem the message to have failed the integrity check. In this way the Digest framework mitigates "reflection attacks" (attacks in which a Man-in-the-Middle reflects a genuine message from an entity back to its sender). It is possible that in the course of generating random nonces the UE and P-CSCF, while operating in the server role, happen to issue identical nonces for use; by making the nonces of sufficient length, the chance of such an occurrence is minimized.

## C.2.6 6.3.6 Digest Operation and Syntax in SIP

In the 3GPP IMS, then, normal operation of the Digest challenge-response mechanism for integrity protection is as follows:

~~Per RFC 2617, t~~The Digest challenge-related directives are carried in ~~either~~ the WWW-Authenticate, ~~or~~ Proxy-Authenticate or UAS-Authenticate header fields. The P-CSCF adds a Proxy-Authenticate header field to the 4xx Auth_Challenge that is sent by the S-CSCF (SIP registrar) toward the UE; the Proxy-Authenticate contains the Digest challenge that has been constructed by the P-CSCF.

~~Per RFC 2617, T~~the Digest response-related directives are carried in ~~either~~ the Authorization, ~~or~~ Proxy-Authorization or UAS-Authorization header fields, depending upon which header field carried the corresponding Digest challenge. These directives contain the credentials for the message integrity check. In the IMS context, the UE responds to the initial Digest challenge by adding a Proxy-Authorization header field to the REGISTER toward the S-CSCF (registrar). The UE pre-emptively adds a Proxy-Authorization header field to all subsequent UE-initiated SIP requests. The ~~UE and the~~ P-CSCF adds the Proxy-Authentication-Info header to all SIP responses. ~~**Finally, t**~~**The P-CSCF adds**

**an ~~Integrity~~ UAS-Authorization header field to all SIP requests sent toward the UE. Finally, the UE adds the UAS-Authentication-Info header to all SIP responses.**

## C.2.7 6.3.7 Example Information Flow

The simplified message flow shown below illustrates the relevant header fields and contents for the SIP-level integrity protection mechanism. Please note that the message flow contains three cases: a registration (1-3), and two SIP sessions: one UE initiated (4-5) and one UE terminated (6-~~7~~9).

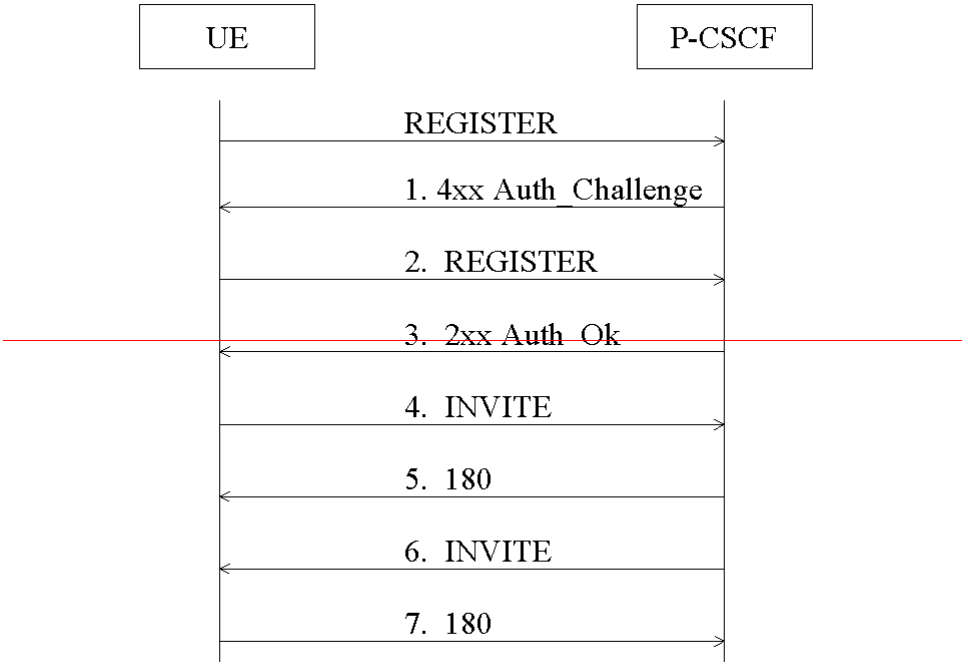

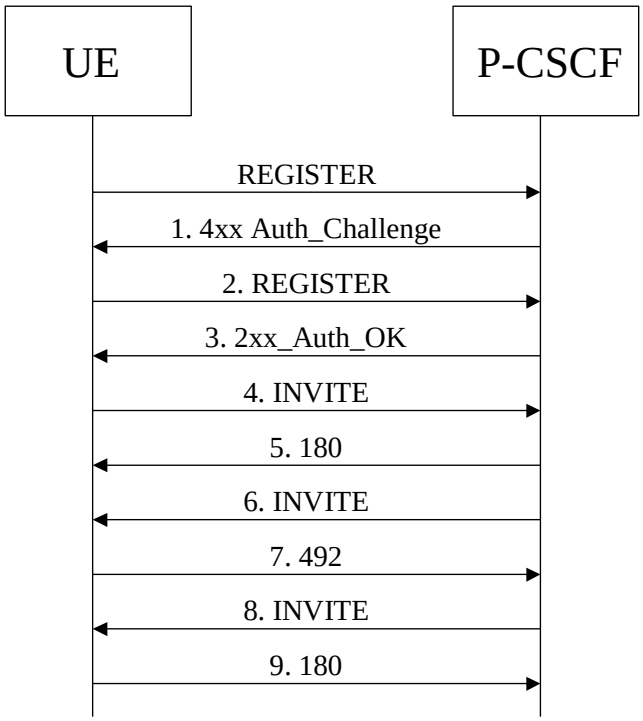


1. **4xx response – this carries both the IMS AKA challenge (from the registrar) and the Digest challenge for integrity protection (from the P-CSCF):**

   SIP/2.0 4xx Auth_Challenge

   WWW-Authenticate: EAP <RAND AUTN>

   Proxy-Authenticate: Digest realm=3GPP-IMS nonce=<~~random-number~~P-nonce1> algorithm=MD5 qop=~~extended~~auth-extd-int

2. **Integrity protection is turned on with the next REGISTER – the integrity credentials are placed in the Digest response:**

   REGISTER sip: ... SIP/2.0

   Authorization: EAP <RES>

   Proxy-Authorization: Digest username=ims-user, realm=3GPP-IMS, nonce=<~~echo-random-number~~P-nonce1>, uri=<SIP-URI>, response=<message-digest>, cnonce=<value>, nc=1, qop=~~extended-int~~auth-extd-int

3. **The 2xx response is also integrity protected – the P-CSCF adds the Proxy-Authentication-Info header to carry the message digest:**

   SIP/2.0 2xx Auth_Ok

   Proxy-Authentication-Info: nextnonce=<P-nonce2>, qop=~~extended-int~~auth-extd-int, rspauth=<message-digest>, nc=~~2~~1, cnonce=<value>

4. **A subsequent INVITE request must also be integrity protected – the UE pre-emptively adds the Proxy-Authorization header:**

   INVITE sip: … SIP/2.0

Proxy-Authorization: Digest username=ims-user, realm=3GPP-IMS, nonce=<~~echo-random-number~~P-nonce2>, uri=<SIP-URI>, response=<message-digest>, cnonce=<value>, nc=~~3~~1, qop=~~extended-int~~auth-extd-int

*Note*: The client (UE) may re-use the previously issued nonce (i.e. set "nonce" to <P-nonce1> and "nc" to 2), but the Digest specification recommends against this. If the 2xx message containing 'nextnonce' were lost and not received by the UE, the UE would then use <P-nonce1> in the computation of the credential.

5. **The 180 is integrity protected in the same fashion was the 2xx response (message #3):**

   SIP/2.0 180 Ringing

   Proxy-Authentication-Info: nextnonce=<P-nonce3>, qop=~~extended-int~~auth-extd-int, rspauth=<message-digest>, nc=~~4~~1, cnonce=<value>

6. **An incoming INVITE must also be integrity protected – the first terminating SIP request, however, must be sent without the integrity credential (this permits the UE to issue a Digest challenge containing its own server-provided nonce).**

7. **The UE issues a 492 response containing a Digest challenge:**

   SIP/2.0 492 Proxies Unauthorized

   UAS-Authenticate: Digest realm=3GPP-IMS, nonce=<UE-nonce1>, algorithm=MD5, qop=auth-extd-int, target=<address>

8. **The P-CSCF adds the UAS-Authorization header, which has similar syntax to Proxy-Authorization:**

   INVITE sip: … SIP/2.0

   ~~Integrity~~UAS-Authorization: Digest username=ims-user, realm=3GPP-IMS, nonce=<~~echo-random-number~~UE-nonce1>, uri=<SIP-URI>, response=<message-digest>, cnonce=<value>, nc=~~5~~1, qop=~~extended-int~~auth-extd-int, responder=<address>

~~7.~~**9. The UE protects the 180 response by adding UAS-Authentication-Info:**

   SIP/2.0 180 Ringing

   UAS-Authentication-Info: realm=3GPP-IMS, nextnonce=<UE-nonce2>, qop=~~extended-int~~auth-extd-int, rspauth=<message-digest>, nc=~~6~~1, cnonce=<value>

~~*[Editors Note: Further details will be provided on how replay protection is accomplished. It has been identified that the scheme above needs to be enhanced since otherwise unnecessary loss of calls can occur. The reason for that is that both originating and terminating calls can occur and the counters in the P-CSCF and in the UE are not independent.]*~~

*[Editors Note: A description of the security mode setup headers shall be included in this Annex. Furthermore the message flows need to be enhanced.]*

## 4. Recommendation

It is recommended that SA3 adopt the text adjustments in section 3 of this contribution for inclusion in draft TS 33.203 to: 1) reflect updates based on the present IETF work in the area of SIP message integrity protection; and 2) correct the treatment of anti-replay protection for the "SIP-level security solution".

## REFERENCES

[1] "HTTP Authentication: Basic and Digest Access Authentication", IETF RFC 2617

[2] 3GPP TSG SA3 document S3-010664 [Siemens]: “Problems with the replay protection scheme in the SIP level integrity solution”

[3] 3GPP TSG SA3 document S3z020008 [Vodafone]: “Reflection Attacks in IMS”

[4] J. Undery et al. “SIP Digest Authentication: Extensions to HTTP Digest Authentication”, IETF Internet-draft, Work in progress, January 2002.