

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

Tdoc S3z020057

Network Working Group
INTERNET-DRAFT
<draft-arkko-sip-sec-agree-02.txt>
April 2002

Jari Arkko
Vesa Torvinen
Ericsson
Tao Haukka
Nokia
Sanjoy Sen
Lee Valerius
Nortel Networks

Security Mechanism Agreement for SIP Sessions

1. Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC2026. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or made obsolete by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as work in progress.

The list of current Internet-Drafts may be found at
<http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories may be found at
<http://www.ietf.org/shadow.html>.

The distribution of this memo is unlimited. It is filed as <draft-arkko-sip-sec-agree-02.txt>, and expires October, 2002. Please send comments to the author or to SIPPING or SIP working group.

2. Abstract

SIP has a number of security mechanisms for hop-by-hop and end-to-end protection. Some of the security mechanisms have been built in to the SIP protocol, such as HTTP authentication or secure attachments. In these mechanisms there are even alternative algorithms and parameters. Currently, HTTP authentication is known to be vulnerable to so called Bidding-Down attacks where a Man-In-The-Middle attacker simply modifies messages in a way that leads parties to believe the other side only supports weaker algorithms than they actually do. Also, currently it isn't possible to select which security mechanisms to use over a connection. In particular, even if some mechanisms such as OPTIONS were used to make this selection, the selection would be again vulnerable against the Bidding-Down attack. On small networks configuration and software update methods are sufficient to deal with this type of attacks, but on large networks that evolve over time, the security implications are serious: either you deny connections from large amounts of older equipment, or risk losing all value of new algorithms through attacks that are trivial to the attackers. This document

defines new headers and procedures for negotiating the security mechanisms within SIP. A SIP entity applying this mechanism must always require some minimum security (i.e. integrity protection) from all communicating parties in order to secure the negotiation.

3. Contents

1.	Status of this Memo.....	1
2.	Abstract.....	1

3.	Contents.....	2
4.	Introduction.....	2
5.	The Problem.....	3
6.	Requirements.....	4
6.1.	Solution.....	5
6.2.	Header descriptions.....	7
7.	Examples.....	8
7.1.	Selecting Between New and Old Mechanisms.....	8
7.2.	Selections Along the Path.....	9
8.	Security Considerations.....	11
9.	Conclusions.....	11
10.	Modifications.....	11
11.	Acknowledgments.....	12
12.	References.....	12
13.	Author's Address.....	12

4. Introduction

Traditionally, security protocols have included facilities to agree on the used mechanisms, algorithms, and other security parameters. The reason for this is that experience has shown algorithm development uncovers problems in old algorithms and produces new ones. Furthermore, different algorithms are suitable for different situations. Typically, protocols also select other parameters beyond algorithms at the same time.

The purpose of this paper is to study whether similar functionality is necessary in SIP [1]. SIP has some security functionality built-in such as HTTP authentication [4], secure attachments such as S/MIME, and can also use underlying security protocols such as IPSec/IKE [2], TLS [3]. Some of the built-in security functionality has also alternative algorithms and other parameters. While some work within the SIP Working Group has been looking towards reducing the number of recommended security solutions (e.g. recommend just one lower layer security protocol), we can not expect to cut down the number of items in the whole list to one. There will still be multiple security solutions in SIP. Furthermore, given that security work around SIP is in its early stages, it is likely that new methods will appear in the future, to complete the methods that exist today.

Chapter 5 shows that without a secure method to choose between security mechanisms and/or their parameters, SIP is vulnerable to certain attacks. As the HTTP authentication RFC [4] points out, authentication and integrity protection using multiple alternative methods and

algorithms is vulnerable to Man-in-the-Middle (MITM) attacks. More seriously, it is hard to know if a SIP peer entity truly can't perform e.g. auth-int QOP in Digest, TLS, or S/MIME, or if a MITM attack is in progress. In small workstation networks these issues are not very relevant, but the deployment of hundreds of millions of small devices with little or no possibilities for coordinated security policies, let alone software upgrades makes these issues much worse. This conclusion is supported by the requirements from 3GPP [5].

Chapter 6 documents the proposed solution, and chapter 7 gives some demonstrative examples.

5. The Problem

SIP has alternative security mechanisms such as HTTP authentication / integrity protection, lower layer security protocol(s), S/MIME. It is likely that their use will continue in the future. SIP security is developing, and is likely to see also new solutions in the future, for example along the introduction of SIP for new network access technologies. Future services may also bring with themselves different security requirements and methods.

Deployment of large number of SIP-based consumer devices such as 3GPP

terminals requires all network devices to be able to accommodate both current and future mechanisms; there is no possibility for instantaneous change since the new solutions are coming gradually in as new standards and product releases occur. It isn't even possible to upgrade some of the devices without getting completely new hardware.

So, the basic security problem that such a large SIP-based network must consider, how do security mechanisms get selected? It would be desirable to take advantage of new mechanisms as they become available in products.

Firstly, we need to know somehow what security should be applied, and preferably find this out without too many additional roundtrips.

Secondly, selection of security mechanisms MUST be secure. Traditionally, all security protocols use a secure form of negotiation. For instance, after establishing mutual keys through Diffie-Hellman, IKE sends hashes of the previously sent data -- including the offered crypto mechanisms. This allows the peers to detect if the initial, unprotected offers were tampered with.

The security implications of this are subtle, but do have a fundamental importance in building large networks that change over time. Given that the hashes are produced also using algorithms agreed in the first unprotected messages, one could ask what the difference in security really is. Assuming integrity protection is mandatory and only secure algorithms are used, we still need to prevent MITM attackers from modifying other parameters, such as whether encryption is provided or not. Let us first assume two peers capable of using both strong and weak security. If the initial offers are not protected in any way,

any attacker can easily "downgrade" the offers by removing the strong options. This would force the two peers to use weak security between them. But if the offers are protected in some way -- such as by hashing, or repeating them later when the selected security is really on -- the situation is different. It would not be sufficient for the attacker to modify a single message. Instead, the attacker would have to modify both the offer message, as well as the message that contains the hash/repetition. More importantly, the attacker would have to forge the weak security that is present in the second message, and would have to do so in real time between the sent offers and the later messages. Otherwise, the peers would notice that the hash is incorrect. If the attacker is able to break the weak security, the security method and/or the algorithm should not be used.

In conclusion, the security difference is making a trivial attack possible versus demanding the attacker to break algorithms. An example of where this has a serious consequence is when a network is first deployed with integrity protection (such as HTTP Digest [4]), and then later new devices are added that support also encryption (such as S/MIME [1]). In this situation, an insecure negotiation procedure allows attackers to trivially force even new devices to use only integrity protection.

It can be asked why the devices would be allowing both weak and strong security in the first place. The answer lies in understanding how networks are deployed, and in the logistical and economical problems in upgrading global networks instantaneously. These issues are of particularly high relevance for networks with a large number of devices, such as the third generation mobile networks. Once millions or even hundreds of millions of devices have been sold to customers, it becomes impossible to replace them with new devices. Therefore, network equipment such as SIP proxies must continue to accept even the older equipment that are less capable in terms of security. Similarly, clients wishing to stay in contact regardless of who they call or where they are, have a need to allow both weaker and stronger mechanisms. Naturally, broken security mechanisms must not be used even when communicating with older devices and network equipment.

6. Solution

The solution to the SIP security negotiation problem should have the following properties:

(a) It allows the selection of security mechanisms, such as lower layer security protocols or secure attachments. It also allows the selection of individual algorithms and parameters where the security functions are integrated in SIP (such as in the case of HTTP authentication or secure attachments).

(b) It allows both end-to-end and hop-by-hop negotiation.

(c) It is secure, i.e. prevents bidding down attacks.

(d) It is capable of running without additional roundtrips. This is important in the cellular environment, where an additional roundtrip could cost 1000 to 1500 ms for the call set up delay.

(e) It does not introduce any additional state to servers and proxies.

Currently, SIP does not have any mechanism which fulfills the requirements above. The basic SIP features such as OPTIONS and Require, Supported headers are capable of informing peers about various capabilities including security mechanisms. However, the straightforward use of these features does not guarantee a secured agreement. HTTP Digest algorithm lists [4] are not secure for picking among the digest

integrity algorithms, as is described in the RFC itself. More seriously, they have no provisions for allowing encryption to be negotiated. Hence, it would be hard to turn on possible future encryption schemes in a secure manner.

6.1. Procedure

In this solution, the security features are represented as regular option tags in SIP. If there will ever be any features that require parameters such as key lengths, the option tags can be associated with an optional value field.

- The clients MUST announce a list of supported option tags in their first request. The servers MUST use this information in preparing their response, such as including a challenge if the first commonly supported mechanism is HTTP Digest. It isn't necessary, however, for the server to remember the clients preferences beyond the response. If the list is not present in the request when arrived to the server, the server MAY respond with an error message. The error message MAY include a list of preferred security mechanisms as well as HTTP authentication challenge if appropriate.

- The servers MUST announce a list of supported option tags in their first response. This list MUST NOT depend on the contents of the list sent by the client in the first message because this list may have been modified by an attacker. For the same reason, the server SHOULD continue the process even if there were no common security mechanisms in the client's and server's lists. Typically, the server's list of supported option tags is static. However, a server MAY maintain several lists, e.g. one for each different access channel.

- The client makes the selection of the used security mechanism based on its own preferences and the server's list. The client MUST start to use the selected security mechanism from the second request message. Note that non-adjacent SIP entities can not use hop-by-hop security mechanisms such as TLS or IPsec. If a client receives a list of hop-by-hop security mechanisms from a server several hops away, it MUST NOT try to use these mechanisms with the first hop proxy. The client MAY try to contact the server several hops away directly leaving the other proxies in between away.

INTERNET-DRAFT

SIP Sec Agreement

4 April 2002

- In the client's second request, the client MUST return the server's list.

The security of the agreement comes from the client's repetition of the server's list of option tags in the second request message. The server can then proceed to verify that the list has not been modified. If a modification is detected, the server MUST return an error or disconnect. The server MUST send a positive answer if and only if the list was not modified. The server does not need to memorize the lists it has sent in earlier responses, provided that the set of security mechanisms supported by the server is constant, which seems like a reasonable assumption. A proxy or server that implements this mechanism must have a policy on whether or not it accepts such requests that does not include security agreement headers. Also these requests may have been modified by an attacker.

Note that one SIP request MAY include several independent security agreement lists. However, only one header SHOULD be used between two SIP entities.

Note also that once the security has been negotiated between two SIP entities, the same SIP entities MAY use the same security when communicating with each other in different SIP roles. For example, if a UAC in a end-user equipment and a UAS in a proxy negotiate some security, they may try to use the same security for terminating requests.

Attackers could try to modify the repeated list in the second request from the client. However, if the selected security mechanism uses encryption this may not be possible, and if it uses integrity protection any modifications will be detected by the server. In order to ensure this, all clients that implement this specification MUST select HTTP Digest, S/MIME, TLS, IPsec, or any stronger method for the protection of the second request. If HTTP Digest is used alone, the security agreement headers MUST be protected. This can be done with HTTP Digest if combined with MIME/SIP tunneling, for example.

Attackers could also try to modify the client's list of security mechanisms in the first message. This would either be revealed to the participants, because of unexpected challenges in the server's first response, or would have no effect because the client picks its own security method only based on its local information and the server's static list.

The client's first protected request can be a real request such as INVITE, as the server MUST check the correctness of the lists before it proceeds to execute the requested operation.

This approach explicitly lists the recipients of the security method agreement. This is intended to allow a negotiation of the first-hop security mechanism while at the same time running e.g. a REGISTER with Digest authentication to a server some hops further away.

This approach could also be trivially extended to support security agreement over a full path. However, since the sips: URI scheme

already solves the most pressing issue in that area we have chosen to not support this.

6.2. Header descriptions

The Security-Method header indicates who wants security towards whom, and what kind of security. The following ABNF describes the syntax of this header and extends section 25.1 in [1]:

```
"Security-Method" HCOLON to-uri COMMA from-uri COMMA mechlist
```

Where

```
to-uri = addr-spec  
from-uri = addr-spec  
mechlist = mechopts *( COMMA mechopts )  
mechopts = mechtags *( SEMI mechtags )  
mechtags = option-tag [EQUAL token]
```

The meaning of these fields is as follows:

- The "to-uri" indicates the desired receiver of the information. The value of this field should be a SIP URI. When sent by a client, the value would typically (but not necessarily) contain just the host and port number parts.
- The "from-uri" indicates the sender of the security agreement information. The value of this is also a SIP URI. When sent by a client, the value would typically (but not necessarily) include a username part.
- The "mechlist" represents a list of alternative security mechanisms. Inside one "mechlist" entry we can have multiple alternative mechanisms and algorithms. The order of the mechanisms in the list represents the server preferences; first 'mechopts' in the list is most favourable mechanism. For instance, the list "org.iana.sip.digest; org.iana.sip.tls, org.iana.sip.ike" would represent the requirement that one must run HTTP Digest and TLS simultaneously, but if that is not possible, also IPsec/IKE is acceptable.

- The "mechopts" represents a list of security mechanisms, all of which must be supported simultaneously on the same connection (such as both Digest and TLS).

The "mechtag" represent one individual mechanism. The "option-tag" syntax is used for these in order to facilitate the easy addition of new mechanisms. All option tags starting with "org.iana.sip." MUST be documented in Internet Drafts or RFCs. The initial list of standardized option-tags is presented below:

```
org.iana.sip.cdigest: client authentication with HTTP Digest
org.iana.sip.digest: mutual authentication with HTTP Digest
org.iana.sip.stls: server authentication with TLS
org.iana.sip.tls: mutual authentication with TLS
```

```
org.iana.sip.smime: S/MIME
org.iana.sip.ike: IPsec/IKE
```

Note that HTTP Digest alone does not fulfill the minimum security requirements (i.e. integrity protection). In order to use HTTP Digest alone, some variant of MIME tunneling should be used. HTTP Digest can also be combined with TLS, for example.

The optional "token" parameter associated with an "option-tag" can be used to assign a parameter value to certain options. This may be useful to select algorithms, key lengths, or other similar parameters in mechanisms integrated to SIP. No such parameters are defined for the four above mechanisms, however.

Multiple instances of the same header field can appear in SIP messages. Typically, the client inserts its own Security-Method header when it sends a request, and the server/proxy adds its own response. The parameters are in all cases set in an appropriate manner to indicate in the "to-uri" parameter the party who inserted the header. Or rather -- since the client is copying some of the server's responses -- whose security capabilities the header applies to.

7. Examples

7.1. Selecting Between New and Old Mechanisms

In this example we demonstrate the use of the framework for securing the first hop using some security mechanism, without knowing beforehand which methods the server supports. We assume that the client is not willing to reveal any information on what it intends to do, so it uses OPTIONS in the first message that is sent in the clear. The example starts by a client sending a message to the server, indicating that it is of the new variant that supports TLS in Step 1. In Step 2, the server responds that with its own list of security mechanisms -- S/MIME or TLS in this case -- and the peers start only common security service i.e. TLS at Step 3. In Step 4, the client resends the server's Security-Method header, which the server verifies, and responds with 200 OK.

1. Client -> Server:

```
OPTIONS server SIP/2.0
Security-Method: sip:client, sip:server, org.iana.sip.tls
```

2. Server -> Client:

```
200 OK
Security-Method: sip:server, sip:client, org.iana.sip.smime,
                org.iana.sip.tls
```

3. Security handshake at a lower layer i.e. TLS

4. Client -> Server:

INTERNET-DRAFT

SIP Sec Agreement

4 April 2002

```
INVITE server SIP/2.0
Security-Method: sip:server, sip:client, org.iana.sip.smime,
                org.iana.sip.tls
```

5. Server -> Client:

```
200 OK
```

In the example we have omitted the returned values of Security-Method in replies for clarity. Typically in SIP the servers do not remove header fields as they answer, they only add new headers.

If this example was run without Security-Method in Step 2, the client would not know what kind of security the other one supports, and would be forced to error-prone trials.

More seriously, if the Security-Method was omitted in Step 4, the whole process would be prone for MITM attacks. An attacker could spoof "ICMP Port Unreachable" message on the trials, or remove the stronger security option from the header in Step 1, therefore substantially reducing the security.

7.2. Selections Along the Path

This example attempts to show how selections can be made e.g. between a client and the first-hop proxy while the actual SIP messages are still destined to a server further on in the network. This example also demonstrates how we can fulfill the 3GPP requirements on being

able to securely agree on the security mechanism between the client and its first hop proxy, without adding roundtrips.

In 3GPP networks, the clients make REGISTER operation in their first message, in order to inform the home network that they are at a particular location. Due to the properties of 3GPP radio interfaces, it is necessary to optimize the number of roundtrips needed in the whole process. Therefore, we try to parallelize the tasks. It should be noted that the same functionality could be achieved using additional OPTIONS messages.

This example does not assume anything on the security methods used in 3GPP. Instead, it demonstrates a general case in which the first hop proxy and the client terminal may support one or more of the three alternative security solutions: a) TLS with HTTP Digest, b) IPsec/IKE, and c) plain HTTP Digest (which may use some variant of MIME/SIP tunneling in order to provide full integrity protection).

The example starts by a client coming to a new area and learning the address of the local proxy. The proxy is of a new version, so it supports all security mechanisms. The client supports alternatives b) and c). The client also knows its home server address. We assume that some trust has already been established between the client and the home, and between the client and the proxy. Perhaps this trust is in the form of the nodes belonging under the same PKI, or having distributed shared secrets beforehand.

In Step 1 the client contacts the proxy using a REGISTER message. We omit the details of the communications with the home server in this

discussion, but the proxy forwards the messages onwards in Step 2. In Step 3, the proxy responds indicating that it is of the new variant that supports TLS with HTTP Digest, IPsec/IKE, and plain HTTP Digest. In Step 4, the client selects the first method it supports (IPsec/IKE in this case), the protection is turned on. In Step 5, the client sends the next round of REGISTER messages to the server. This includes the repetition of the original security capabilities of the server. The server verifies this list, and in Step 7 it responds with a 200 OK.

1. Client -> Proxy:

```
REGISTER server SIP/2.0
Security-Method: sip:client, sip:proxy, org.iana.sip.ike,
                org.iana.sip.digest
```

2. Proxy communicates with the Server.

3. Proxy -> Client:

```
401 Authentication Required
(HTTP Digest challenge from the proxy to the client)
Security-Method: sip:proxy, sip:client, org.iana.sip.stls;
                org.iana.sip.cdigest, org.iana.sip.ike,
                org.iana.sip.digest
```

4. Security handshake at a lower layer i.e. IPsec/IKE

5. Client -> Proxy:

```
REGISTER server SIP/2.0
Security-Method: sip:proxy, sip:client, org.iana.sip.stls;
                org.iana.sip.cdigest, org.iana.sip.ike,
                org.iana.sip.digest
```

6. Proxy communicates with the Server.

7. Proxy -> Client:

200 OK

As in the previous example, if this was run without Security-Method in Step 3, the client would not know what kind of algorithms the server supports. In this example we demonstrate also the need for the client to send its own mechanism list in Step 1. If this wasn't known to the proxy when it responds in Step 3, it could not have provided a suitable HTTP Digest challenge because at that point the proxy would not have known if the client supports that.

J. Arkko et al

Expires October 2002

[Page 10]

INTERNET-DRAFT

SIP Sec Agreement

4 April 2002

As in the previous example, removing the repetition of the Security-Method header in Step 5 would open the system to MITM attacks.

8. Security Considerations

This draft is about making it possible to select between various SIP security mechanisms in a secure manner. In particular, the method presented here allow current networks using e.g. Digest later securely upgrade to e.g. S/MIME without requiring a simultaneous modification in all equipment.

The method presented in this draft is secure only if the weakest proposed mechanism offers at least integrity protection. Therefore, at least HTTP Digest authentication in combination with some variant of MIME tunneling SHOULD be used in conjunction with our approach.

9. Conclusions

The presented method appear to secure the selection between different security mechanisms. The authors encourage security analysis of the proposal.

10. Modifications

The -02 version of this draft introduced the following modifications:

- Clarified that the solution requires always some base level of security (i.e. integrity) in order to work. Even 'the weak security' must not be broken.
- References to Enhanced HTTP Digest removed.
- Option tags has been modified. This version proposes option tags also for cases in which a security method is used for one direction only (e.g. if TLS is used for server authentication only).
- Motivation section has been shortened since this is now a WG item.
- Text related to alternative solutions shortened and moved to a new place.
- New rules for possible error and special cases has been added, e.g. for the case in which an non-adjacent SIP entities try to negotiate hop-by-hop security mechanisms.
- The semantics of the 'mechlist' and 'mechopts' has been changed in order to cover more potential use cases. The comma is now 'OR' and the semicolon is now 'AND'.

The -01 version of this draft introduced the following modifications:

- Reversed approach to make servers stateless

INTERNET-DRAFT

SIP Sec Agreement

4 April 2002

- Removed discussion of the use of this for Digest algorithm selection, since Enhanced Digest already has bidding-down protection
- Renamed org.iana.sip.digest to org.iana.sip.edigest and removed the parameters, as we can rely on Enhanced Digest to perform the algorithm selection.
- Removed agreements for full paths.
- Simplified syntax

11. Acknowledgments

The authors wish to thank Rolf Blom, James Undery, Jonathan Rosenberg, Hugh Shieh, Gunther Horn, Krister Boman, David Castellanos-Zamora, Aki Niemi, Miguel Garcia, Valtteri Niemi, Martin Euchner, Eric Rescorla, and members of the 3GPP SA3 group for interesting discussions in this problem space.

12. References

- [1] Handley, M., Schulzrinne, H, Schooler, E. and Rosenberg, J., "SIP: Session Initiation Protocol", Work In Progress, draft-ietf-sip-rfc2543bis-08.txt, IETF, February 2002.
- [2] S. Kent, R. Atkinson, "Security Architecture for the Internet Protocol", RFC 2401, November 1998.
- [3] T. Dierks, C. Allen, "The TLS Protocol Version 1.0", RFC 2246, January 1999.

[4] Franks, J. et al, "HTTP Authentication: Basic and Digest Access Authentication", RFC 2617, June 1999.

[5] M. Garcia, D. Mills, G. Bajko, G. Mayer, F. Derome, H. Shieh, A. Allen, S. Chotai, K. Drage, J. Bharatia, "3GPP requirements on SIP", draft-garcia-sipping-3gpp-reqs-00.txt. Work In Progress, IETF, October 2001.

13. Author's Address

Jari Arkko, Vesa Torvinen
Ericsson
02420 Jorvas
Finland
EMail: Jari.Arkko@ericsson.com, Vesa.Torvinen@ericsson.fi

Tao Haukka
Nokia
Finland
EMail: Tao.Haukka@nokia.com

Sanjoy Sen

J. Arkko et al

Expires October 2002

[Page 12]

INTERNET-DRAFT

SIP Sec Agreement

4 April 2002

Nortel Networks
2735-B Glenville Drive
Richardson, TX 75082, USA
EMail: sanjoy@nortelnetworks.com

Lee Valerius

Nortel Networks

2201 Lakeside Blvd

Richards, TX 75082, USA

EMail: valerius@nortelnetworks.com

