

Network Working Group
Internet-Draft
Updates: [4187](#) (if approved)
Intended status: Informational
Expires: May 22, 2009

J. Arkko
V. Lehtovirta
Ericsson
P. Eronen
Nokia
November 18, 2008

Improved Extensible Authentication Protocol Method for 3rd Generation
Authentication and Key Agreement (EAP-AKA')
draft-arkko-eap-aka-kdf-10

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

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 obsoleted 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 can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

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

This Internet-Draft will expire on May 22, 2009.

Abstract

This specification defines a new EAP method, EAP-AKA', a small revision of the EAP-AKA method. The change is a new key derivation function that binds the keys derived within the method to the name of the access network. The new key derivation mechanism has been defined in the 3rd Generation Partnership Project (3GPP). This specification allows its use in EAP in an interoperable manner. In addition, EAP-AKA' employs SHA-256 instead of SHA-1.

This specification also updates [RFC 4187](#) EAP-AKA to prevent bidding

Internet-Draft

EAP-AKA'

November 2008

down attacks from EAP-AKA'.

Table of Contents

1.	Introduction	3
2.	Requirements language	4
3.	EAP-AKA'	4
3.1.	AT_KDF_INPUT	6
3.2.	AT_KDF	8
3.3.	Key Generation	10
3.4.	Hash Functions	12
3.4.1.	PRF'	12
3.4.2.	AT_MAC	12
3.4.3.	AT_CHECKCODE	12
4.	Bidding Down Prevention for EAP-AKA	13
5.	Security Considerations	14
5.1.	Security Properties of Binding Network Names	17
6.	IANA Considerations	18
6.1.	Type Value	18
6.2.	Attribute Type Values	18
6.3.	Key Derivation Function Namespace	19
7.	Acknowledgments	19
8.	References	19
8.1.	Normative References	19
8.2.	Informative References	20
Appendix A.	Changes from RFC 4187	21
Appendix B.	Importance of Explicit Negotiation	21
	Authors' Addresses	22
	Intellectual Property and Copyright Statements	23

Internet-Draft

EAP-AKA'

November 2008

1. Introduction

This specification defines a new Extensible Authentication Protocol (EAP)[[RFC3748](#)] method, EAP-AKA', a small revision of the EAP-AKA method originally defined in [[RFC4187](#)]. What is new in EAP-AKA' is that it has a new key derivation function specified in [[3GPP.33.402](#)]. This function binds the keys derived within the method to the name of the access network. This limits the effects of compromised access network nodes and keys. This specification defines the EAP encapsulation for AKA when the new key derivation mechanism is in use.

3GPP has defined a number of applications for the revised AKA mechanism, some based on native encapsulation of AKA over 3GPP radio access networks and others based on the use of EAP.

For making the new key derivation mechanisms usable in EAP-AKA additional protocol mechanisms are necessary. Given that [RFC 4187](#) calls for the use of CK (the encryption key) and IK (the integrity key) from AKA, existing implementations continue to use these. Any change of the key derivation must be unambiguous to both sides in the protocol. That is, it must not be possible to accidentally connect old equipment to new equipment and get the key derivation wrong or attempt to use wrong keys without getting a proper error message. The change must also be secure against bidding down attacks that attempt to force the participants to use the least secure mechanism.

This specification therefore introduces a variant of the EAP-AKA method, called EAP-AKA'. This method can employ the derived keys CK' and IK' from the 3GPP specification and updates the used hash function to SHA-256 [[FIPS.180-2.2002](#)]. But it is otherwise equivalent to [RFC 4187](#). Given that a different EAP method Type value is used for EAP-AKA and EAP-AKA', a mutually supported method may be negotiated using the standard mechanisms in EAP [[RFC3748](#)].

Note: [Appendix B](#) explains why it is important to be explicit about

the change of semantics for the keys, and why other approaches would lead to severe interoperability problems.

The rest of this specification is structured as follows. [Section 3](#) defines the EAP-AKA' method. [Section 4](#) adds support to EAP-AKA to prevent bidding down attacks from EAP-AKA'. [Section 5](#) explains the security differences between EAP-AKA and EAP-AKA'. [Section 6](#) describes the IANA considerations and [Appendix A](#) explains what updates to [RFC 4187](#) EAP-AKA have been made in this specification. Finally, [Appendix B](#) explains some of the design rationale for creating EAP-AKA'.

Editor's Note: The publication of this RFC depends on its normative references [[3GPP.24.302](#)], [[3GPP.33.102](#)], and [[3GPP.33.402](#)] from 3GPP reaching their final Release 8 status at 3GPP. This is expected to happen shortly. The RFC Editor should check with the 3GPP liaisons that this has happened. RFC Editor: Please delete this note upon publication of this specification as an RFC.

[2.](#) Requirements language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

[3.](#) EAP-AKA'

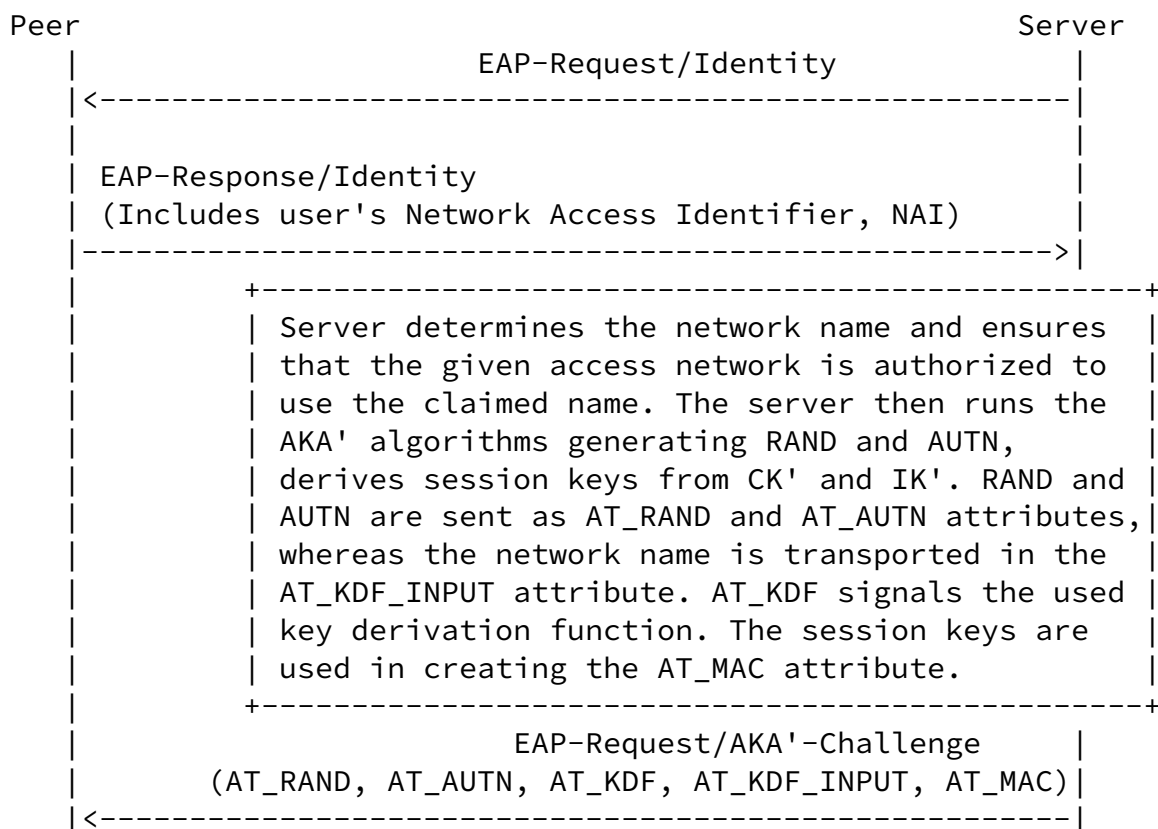
EAP-AKA' is a new EAP method that follows the EAP-AKA specification [[RFC4187](#)] in all respects except the following:

- o It uses the Type code TBA1 BY IANA, not 23 which is used by EAP-AKA.
- o It carries the AT_KDF_INPUT attribute, as defined in [Section 3.1](#) to ensure that both the peer and server know the name of the access network.
- o It supports key derivation function negotiation via the AT_KDF

attribute ([Section 3.2](#)), to allow for future extensions.

- o It calculates keys as defined in [Section 3.3](#), not as defined in EAP-AKA.
- o It employs SHA-256 [[FIPS.180-2.2002](#)], not SHA-1 [[FIPS.180-1.1995](#)] ([Section 3.4](#)).

Figure 1 shows an example of the authentication process. Each message AKA'-Challenge and so on represents the corresponding message from EAP-AKA, but with EAP-AKA' Type code. The definition of these messages, along with the definition of attributes AT_RANDOM, AT_AUTN, AT_MAC, and AT_RES can be found in [[RFC4187](#)].



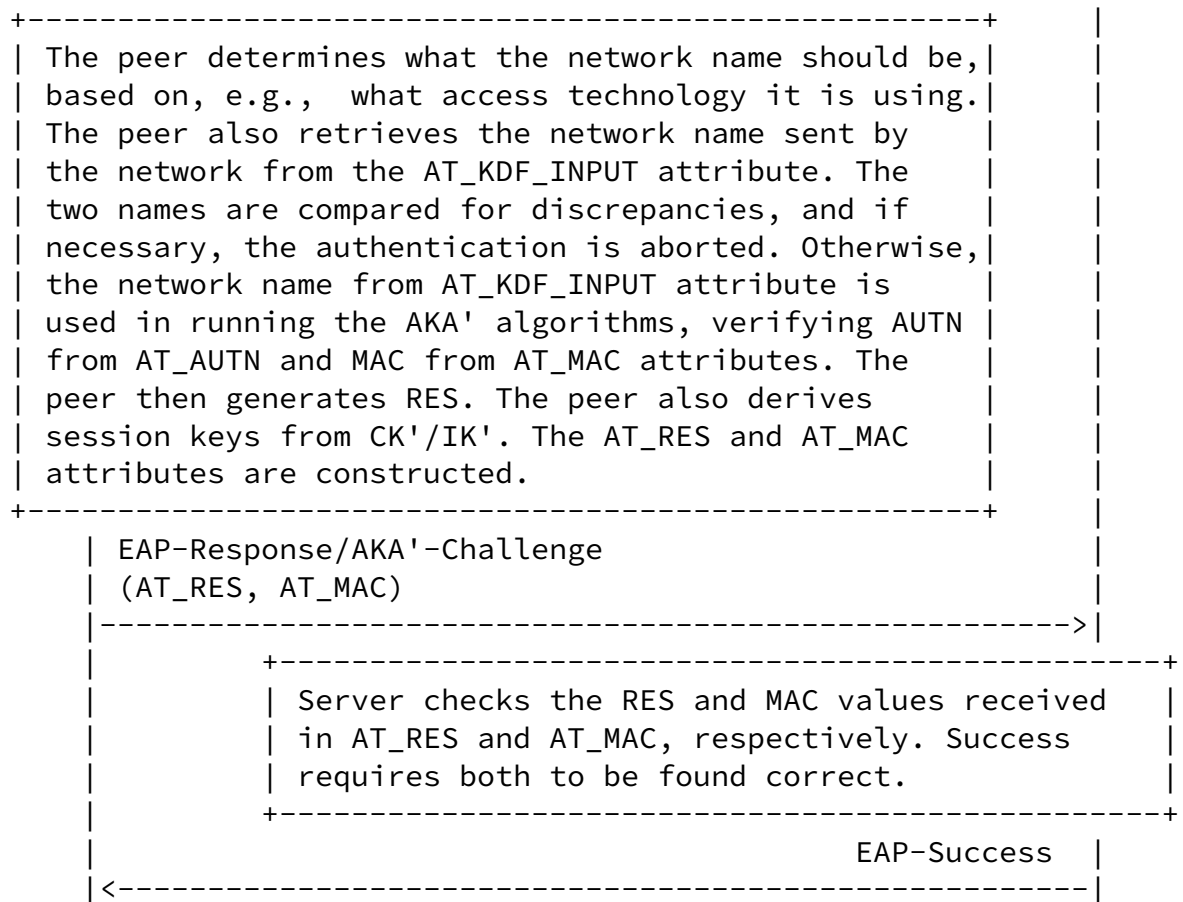
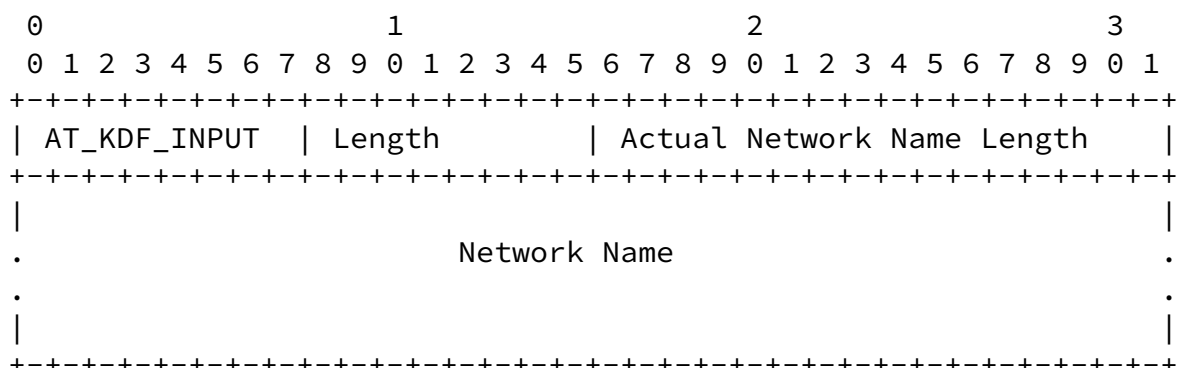


Figure 1: EAP-AKA' Authentication Process

[3.1.](#) AT_KDF_INPUT

The format of the AT_KDF_INPUT attribute is shown below.



The fields are as follows:

AT_KDF_INPUT

This is set to TBA2 BY IANA.

Length

The length of the attribute, calculated as defined in [\[RFC4187\]](#) [Section 8.1](#).

Actual Network Name Length

This a 2-byte actual length field, needed due to the requirement that the previous field is expressed in multiples of 4 bytes per the usual EAP-AKA rules. The Actual Network Name Length field provides the length of the Network Name in bytes.

Network Name

This field contains the network name of the access network for which the authentication is being performed. The name does not include any terminating null characters. Because the length of the entire attribute must be a multiple of 4 bytes, the sender pads the name with one, two, or three bytes of all zero bits when necessary.

Only the server sends the AT_KDF_INPUT attribute. Per [\[3GPP.33.402\]](#), the server always verifies the authorization of a given access network to use a particular name before sending it to the peer over EAP-AKA'. The value of the AT_KDF_INPUT attribute from the server MUST be non-empty. If it is empty, the peer behaves as if AUTN had

been incorrect and authentication fails. See [Section 3](#) and Figure 3 of [\[RFC4187\]](#) for an overview of how authentication failures are handled.

In addition, the peer MAY check the received value against its own understanding of the network name. Upon detecting a discrepancy, the peer either warns the user and continues, or fails the authentication process. More specifically, the peer SHOULD have a configurable

policy which it can follow under these circumstances. If the policy indicates that it can continue, the peer SHOULD log a warning message or display it to the user. If the peer chooses to proceed, it MUST use the network name as received in the AT_KDF_INPUT attribute. If the policy indicates that the authentication should fail, the peer behaves as if AUTN had been incorrect and authentication fails.

The Network Name field contains an UTF-8 string. This string MUST be constructed as specified in [3GPP.24.302] for "Access Network Identity". The string is structured as fields separated by colons (:). The algorithms and mechanisms to construct the identity string depend on the used access technology.

On the network side, the network name construction is a configuration issue in an access network and an authorization check in the authentication server. On the peer, the network name is constructed based on the local observations. For instance, the peer knows which access technology it is using on the link, it can see information in a link layer beacon, and so on. The construction rules specify how this information maps to an access network name. Typically, the network name consists of the name of the access technology, or name of the access technology followed by some operator identifier that was advertised in a link layer beacon. In all cases [3GPP.24.302] is the normative specification for the construction in both the network and peer side. If the peer policy allows running EAP-AKA' over an access technology for which that specification does not provide network name construction rules, the peer SHOULD rely only on the information from the AT_KDF_INPUT attribute and not perform a comparison.

If a comparison of the locally determined network name and the one received over EAP-AKA' is performed on the peer, it MUST be done as follows. First, each name is broken down to the fields separated by colon. If one of the names has more colons and fields than the other one, the additional fields are ignored. The remaining sequences of fields are compared, and they match only if they are equal character by character. This algorithm allows a prefix match where the peer would be able to match "", "FOO", and "FOO:BAR" against the value "FOO:BAR" received from the server. This capability is important in order to allow possible updates to the specifications that dictate

how the network names are constructed. For instance, if a peer knows

that it is running on access technology "FOO" it can use the string "FOO" even if the server uses an additional, more accurate description "FOO:BAR" that contains more information.


The allocation procedures in [3GPP.24.302] ensure that conflicts potentially arising from using the same name in different types of networks are avoided. The specification also has detailed rules about how a client can determine these based on information available to the client, such as the type of protocol used to attach to the network, beacons sent out by the network, and so on. Information that the client cannot directly observe (such as the type or version of the home network) is not used by this algorithm.

The AT_KDF_INPUT attribute MUST be sent and processed as explained above when AT_KDF attribute has the value 1. Future definitions of new AT_KDF values MUST define how this attribute is sent and processed.

3.2. AT_KDF

AT_KDF is an attribute that the server uses to reference a specific key derivation function. It offers a negotiation capability that can be useful for future evolution of the key derivation functions.

The format of the AT_KDF attribute is shown below.



The fields are as follows:

AT_KDF

This is set to TBA3 BY IANA.

Length

The length of the attribute, MUST be set to 1.

Key Derivation Function

An enumerated value representing the key derivation function that the server (or peer) wishes to use. Value 1 represents the

default key derivation function for EAP-AKA', i.e., employing CK' and IK' as defined in [Section 3.3](#).

Servers MUST send one or more AT_KDF attributes in the EAP-Request/AKA'-Challenge message. These attributes represent the desired functions ordered by preference, the most preferred function being the first attribute.

Upon receiving a set of these attributes, if the peer supports and is willing to use the key derivation function indicated by the first attribute, the function is taken into use without any further negotiation. However, if the peer does not support this function or is unwilling to use it, it responds with the EAP-Response/AKA'-Challenge message that contains only one attribute, AT_KDF with the value set to the selected alternative. If there is no suitable alternative, the peer behaves as if AUTN had been incorrect and authentication fails (see Figure 3 of [\[RFC4187\]](#)). The peer fails the authentication also if there are any duplicate values within the list of AT_KDF attributes (except where the duplication is due to a request to change the key derivation function; see below for further information).

Upon receiving an EAP-Response/AKA'-Challenge with AT_KDF from the peer, the server checks that the suggested AT_KDF value was one of the alternatives in its offer. The first AT_KDF value in the message from the server is not a valid alternative. If the peer has replied with the first AT_KDF value, the server behaves as if AT_MAC of the response had been incorrect and fails the authentication. For an overview of the failed authentication process in the server side, see [Section 3](#) and Figure 2 in [\[RFC4187\]](#). Otherwise, the server re-sends the EAP-Response/AKA'-Challenge message, but adds the selected alternative to the beginning of the list of AT_KDF attributes, and retains the entire list following it. Note that this means that the selected alternative appears twice in the set of AT_KDF values. Responding to the peer's request to change the key derivation function is the only legal situation where such duplication may occur.

When the peer receives the new EAP-Request/AKA'-Challenge message, it MUST check that the requested change, and only the requested change occurred in the list of AT_KDF attributes. If yes, it continues. If not, it behaves as if AT_MAC had been incorrect and fails the authentication. If the peer receives multiple EAP-Request/AKA'-Challenge messages with differing AT_KDF attributes without having requested negotiation, the peer MUST behave as if AT_MAC had been incorrect and fail the authentication.

Internet-Draft

EAP-AKA'

November 2008

[3.3.](#) Key Generation

Both the peer and server MUST derive the keys as follows.

AT_KDF set to 1

In this case MK is derived and used as follows:

```
MK = PRF'(IK'|CK',"EAP-AKA'|Identity)
K_encr = MK[0..127]
K_aut  = MK[128..383]
K_re   = MK[384..639]
MSK    = MK[640..1151]
EMSK   = MK[1152..1663]
```

Here [n..m] denotes the substring from bit n to m. PRF' is a new pseudo random function specified in [Section 3.4](#). The 1664 first bits from its output are used for K_encr (encryption key, 128 bits), K_aut (authentication key, 256 bits), K_re (re-authentication key, 256 bits), MSK (Master Session Key, 512 bits) and EMSK (Extended Master Session Key, 512 bits). These keys are used by the subsequent EAP-AKA' process. K_encr is used by the AT_ENCR_DATA attribute, and K_aut by the AT_MAC attribute. K_re is used later in this section. MSK and EMSK are outputs from a successful EAP method run [[RFC3748](#)].

IK' and CK' are derived as specified in [[3GPP.33.402](#)]. The functions that derive IK' and CK' take the following parameters: CK and IK produced by the AKA algorithm, and value of the Network Name field (without length or padding) from AT_KDF_INPUT.

The value "EAP-AKA'" is an eight characters long ASCII string. It is used as is, without any trailing NUL characters.

Identity is the peer identity as specified in [Section 7 of \[RFC4187\]](#).

When the server creates an AKA challenge and corresponding AUTN,

CK, CK', IK, and IK' values, it MUST set the AMF separation bit to 1 in the AKA algorithm [[3GPP.33.102](#)]. Similarly, the peer MUST check that the AMF separation bit set is to 1. If the bit is not set to 1, the peer behaves as if the AUTN had been incorrect and fails the authentication.

On fast re-authentication, the following keys are calculated:

```
MK = PRF'(K_re,"EAP-AKA' re-auth"|Identity|counter|NONCE_S)
MSK  = MK[0..511]
EMSK = MK[512..1023]
```

MSK and EMSK are the resulting 512 bit keys, taking the first 1024 bits from the result of PRF'. Note that K_encr and K_aut are not re-derived on fast re-authentication. K_re is the re-authentication key from the preceding full authentication and stays unchanged over any fast re-authentication(s) that may happen based on it. The value "EAP-AKA' re-auth" is a sixteen characters long ASCII string, again represented without any trailing NUL characters. Identity is the fast re-authentication identity, counter is the value from the AT_COUNTER attribute, NONCE_S is the nonce value from the AT_NONCE_S attribute, all as specified in [Section 7 of \[RFC4187\]](#). To prevent the use of compromised keys on other places, it is forbidden to change the network name when going from the full to the fast re-authentication process. The peer SHOULD NOT attempt fast re-authentication when it knows that the network name in the current access network is different from the one in the initial, full authentication. Upon seeing a re-authentication request with a changed network name, the server SHOULD behave as if the re-authentication identifier had been unrecognized and fall back to full authentication. The server observes the change in the name by comparing where the fast re-authentication and full authentication EAP transactions were received from at the Authentication, Authorization, and Accounting (AAA) protocol level.

AT_KDF has any other value

Future variations of key derivation functions may be defined, and

they will be represented by new values of AT_KDF. If the peer does not recognize the value it cannot calculate the keys and behaves as explained in [Section 3.2](#).

AT_KDF is missing

The peer behaves as if the AUTN had been incorrect and MUST fail the authentication.

If the peer supports a given key derivation function but is unwilling to perform it for policy reasons, it refuses to calculate the keys and behaves as explained in [Section 3.2](#).

[3.4](#). Hash Functions

EAP-AKA' uses SHA-256 [[FIPS.180-2.2002](#)], not SHA-1 [[FIPS.180-1.1995](#)] as in EAP-AKA. This requires a change to the pseudo random function (PRF) as well as the AT_MAC and AT_CHECKCODE attributes.

[3.4.1](#). PRF'

The PRF' construction is the same one as IKEv2 uses (see [Section 2.13 in \[RFC4306\]](#)). The function takes two arguments. K is a 256 bit value and S is an octetstring of arbitrary length. PRF' is defined as follows:

$$\text{PRF}'(K, S) = T1 \mid T2 \mid T3 \mid T4 \mid \dots$$

where:

T1 = HMAC-SHA-256 (K, S | 0x01)

T2 = HMAC-SHA-256 (K, T1 | S | 0x02)

T3 = HMAC-SHA-256 (K, T2 | S | 0x03)

T4 = HMAC-SHA-256 (K, T3 | S | 0x04)

...

PRF' produces as many bits of output as is needed. HMAC-SHA-256 is the application of HMAC [[RFC2104](#)] to SHA-256.

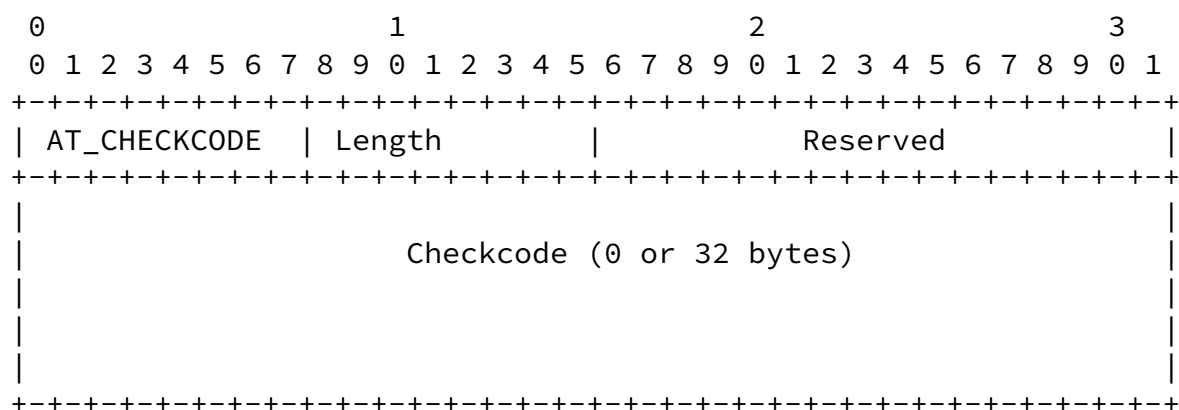
3.4.2. AT_MAC

When used within EAP-AKA', the AT_MAC attribute is changed as follows. The MAC algorithm is HMAC-SHA-256-128, a keyed hash value. The HMAC-SHA-256-128 value is obtained from the 32-byte HMAC-SHA-256 value by truncating the output to the first 16 bytes. Hence, the length of the MAC is 16 bytes.

Otherwise the use of AT_MAC in EAP-AKA' follows [Section 10.15 of \[RFC4187\]](#).

3.4.3. AT_CHECKCODE

When used within EAP-AKA', the AT_CHECKCODE attribute is changed as follows. First, a 32 byte value is needed to accommodate a 256 bit hash output:



Second, the checkcode is a hash value, calculated with SHA-256 [\[FIPS.180-2.2002\]](#), over the data specified in [Section 10.13 of \[RFC4187\]](#).

4. Bidding Down Prevention for EAP-AKA

willing to use it, and prefers it over EAP-AKA. Otherwise it should be set to 0.

Reserved

This field MUST be set to zero when sent and ignored on receipt.

The server sends this attribute in the EAP-Request/AKA-Challenge message. If the peer supports EAP-AKA', it compares the received value to its own capabilities. If it turns out that both the server and peer would have been able to use EAP-AKA' and preferred it over EAP-AKA, the peer behaves as if AUTN had been incorrect, and fails the authentication (see Figure 3 of [[RFC4187](#)]). A peer not supporting EAP-AKA' will simply ignore this attribute. In all cases, the attribute is protected by the integrity mechanisms of EAP-AKA, so it cannot be removed by a man-in-the-middle attacker.

Note that we assume ([Section 5](#)) that EAP-AKA' is always stronger than EAP-AKA. As a result, there is no need to prevent bidding "down" attacks in the other direction, i.e., attackers forcing the endpoints to use EAP-AKA'.

[5](#). Security Considerations

A summary of the security properties of EAP-AKA' follows. These properties are very similar to those in EAP-AKA. We assume that SHA-256 is at least as secure as SHA-1. This is called the SHA-256 assumption in the remainder of this section. Under this assumption EAP-AKA' is at least as secure as EAP-AKA.

If the AT_KDF attribute has value 1, then the security properties of EAP-AKA' are as follows:

Protected ciphersuite negotiation

EAP-AKA' has no ciphersuite negotiation mechanisms. It does have a negotiation mechanism for selecting the key derivation

functions. This mechanism is secure against bidding down attacks. The negotiation mechanism allows changing the offered key derivation function, but the change is visible in the final EAP-Request/AKA'-Challenge message that the server sends to the peer. This message is authenticated via the AT_MAC attribute, and carries both the chosen alternative and the initially offered list. The peer refuses to accept a change it did not initiate. As a result, both parties are aware that a change is being made and what the original offer was.

Mutual authentication

Under the SHA-256 assumption, the properties of EAP-AKA' are at least as good as those of EAP-AKA in this respect. Refer to [\[RFC4187\] Section 12](#) for further details.

Integrity protection

Under the SHA-256 assumption, the properties of EAP-AKA' are at least as good (most likely better) as those of EAP-AKA in this respect. Refer to [\[RFC4187\] Section 12](#) for further details. The only difference is that a stronger hash algorithm, SHA-256 is used instead of SHA-1.

Replay protection

Under the SHA-256 assumption, the properties of EAP-AKA' are at least as good as those of EAP-AKA in this respect. Refer to [\[RFC4187\] Section 12](#) for further details.

Confidentiality

The properties of EAP-AKA' are exactly the same as those of EAP-AKA in this respect. Refer to [\[RFC4187\] Section 12](#) for further details.

Key derivation

EAP-AKA' supports key derivation with an effective key strength against brute force attacks equal to the minimum of the length of the derived keys and the length of the AKA base key, i.e. 128-bits or more. The key hierarchy is specified in [Section 3.3](#).

The Transient EAP Keys used to protect EAP-AKA packets (K_encr,

K_{aut}, K_{re}), the MSK, and the EMSK are cryptographically separate. If we make the assumption that SHA-256 behaves as a pseudo-random function, an attacker is incapable of deriving any non-trivial information about any of these keys based on the other keys. An attacker also cannot calculate the pre-shared secret from IK, CK, IK', CK', K_{encr}, K_{aut}, K_{re}, MSK, or the EMSK by any practically feasible means.

EAP-AKA' adds an additional layer of key derivation functions within itself to protect against the use of compromised keys. This is discussed further in [Section 5.1](#).

EAP-AKA' uses a pseudo random function modeled after the one used in IKEv2 [[RFC4306](#)] together with SHA-256.

Key strength

See above.

Dictionary attack resistance

Under the SHA-256 assumption, the properties of EAP-AKA' are at least as good as those of EAP-AKA in this respect. Refer to [[RFC4187](#)] [Section 12](#) for further details.

Fast reconnect

Under the SHA-256 assumption, the properties of EAP-AKA' are at least as good as those of EAP-AKA in this respect. Refer to [[RFC4187](#)] [Section 12](#) for further details. Note that implementations MUST prevent performing a fast reconnect across method types.

Cryptographic binding

Note that this term refers to a very specific form of binding, something that is performed between two layers of authentication. It is not the same as the binding to a particular network name. The properties of EAP-AKA' are exactly as those of EAP-AKA in this respect, i.e., as it is not a tunnel method this property is not applicable to it. Refer to [[RFC4187](#)] [Section 12](#) for further details.

Session independence

The properties of EAP-AKA' are exactly the same as those of EAP-AKA in this respect. Refer to [[RFC4187](#)] [Section 12](#) for further

details.

Fragmentation

The properties of EAP-AKA' are exactly the same as those of EAP-AKA in this respect. Refer to [\[RFC4187\] Section 12](#) for further details.

Channel binding

EAP-AKA', like EAP-AKA, does not provide channel bindings as they're defined in [\[RFC3748\]](#) and [\[RFC5247\]](#). New skippable attributes can be used to add channel binding support in the future, if required.

However, including the network name field in the AKA' algorithms (which are also used for other purposes than EAP-AKA') does provide a form of cryptographic separation between different network names, which resembles channel bindings. However, the network name does not typically identify the EAP (pass-through) authenticator. See the following section for more discussion.

[5.1.](#) Security Properties of Binding Network Names

The ability of EAP-AKA' to bind the network name into the used keys provides some additional protection against key leakage to inappropriate parties. The keys used in the protocol are specific to a particular network name. If key leakage occurs due to an accident, access node compromise, or another attack, the leaked keys are only useful when providing access with that name. For instance, a malicious access point cannot claim to be network Y if has stolen keys from network X. Obviously, if an access point is compromised, the malicious node can still represent the compromised node. As a result, neither EAP-AKA' or any other extension can prevent such attacks, but the binding to a particular name limits the attacker's choices, allows better tracking of attacks, makes it possible to identify compromised networks, and applies good cryptographic hygiene.

The server receives the EAP transaction from a given access network, and verifies that the claim from the access network corresponds to the name that this access network should be using. It becomes

impossible for an access network to claim over AAA that it is another access network. In addition, if the peer checks that the information it has received locally over the network access link layer matches with the information the server has given it via EAP-AKA', it becomes impossible for the access network to tell one story to the AAA network and another one to the peer. These checks prevent some "lying NAS" (Network Access Server) attacks. For instance, a roaming partner, R, might claim that it is the home network H in an effort to

lure peers to connect to itself. Such an attack would be beneficial for the roaming partner if it can attract more users, and damaging for the users if their access costs in R are higher than those in other alternative networks, such as H.

Any attacker who gets hold of the keys CK, IK produced by the AKA algorithm can compute the keys CK', IK' and hence the master key MK according to the rules in [Section 3.3](#). The attacker could then act as a lying NAS. In 3GPP systems in general, the keys CK and IK have been distributed to, for instance, nodes in a visited access network where they may be vulnerable. In order to reduce this risk the AKA algorithm MUST be computed with the AMF separation bit set to 1, and the peer MUST check that this is indeed the case whenever it runs EAP-AKA'. Furthermore, [\[3GPP.33.402\]](#) requires that no CK, IK keys computed in this way ever leave the home subscriber system.

The additional security benefits obtained from the binding depend obviously on the way names are assigned to different access networks. This is specified in [\[3GPP.24.302\]](#). See also [\[3GPP.23.003\]](#). Ideally, the names allow separating each different access technology, each different access network, and each different NAS within a domain. If this is not possible, the full benefits may not be achieved. For instance, if the names identify just an access technology, use of compromised keys in a different technology can be prevented, but it is not possible to prevent their use by other domains or devices using the same technology.

[6.](#) IANA Considerations

[6.1.](#) Type Value

EAP-AKA' has the EAP Type value TBA1 BY IANA in the Extensible

Authentication Protocol (EAP) Registry under Method Types. Per [\[RFC3748\] Section 6.2](#), this allocation can be made with Designated Expert and Specification Required.

[6.2.](#) Attribute Type Values

EAP-AKA' shares its attribute space and subtypes with EAP-SIM [\[RFC4186\]](#) and EAP-AKA [\[RFC4186\]](#). No new registries are needed.

However, a new Attribute Type value (TBA2) in the non-skippable range needs to be assigned for AT_KDF_INPUT ([Section 3.1](#)) in the EAP-AKA and EAP-SIM Parameters registry under Attribute Types.

Also, a new Attribute Type value (TBA3) in the non-skippable range needs to be assigned for AT_KDF ([Section 3.2](#)).

Finally, a new Attribute Type value (TBA4) in the skippable range needs to be assigned for AT_BIDDING ([Section 4](#)).

[6.3.](#) Key Derivation Function Namespace

IANA also needs to create a new namespace for EAP-AKA' AT_KDF Key Derivation Function values. This namespace can exist under the EAP-AKA and EAP-SIM Parameters registry. The initial contents of this namespace are given below; new values can be created through Specification Required policy [\[RFC5226\]](#).

Value	Description	Reference
-----	-----	-----
0	Reserved	
1	EAP-AKA' with CK'/IK'	[this document]
2-65535	Unassigned	

[7.](#) Acknowledgments

The authors would like to thank Guenther Horn, Joe Salowey, Mats Naslund, Adrian Escott, Brian Rosenberg, Lakshminath Dondeti, Ahmad Muhanna, Stefan Rommer, Miguel Garcia, Jan Kall, Ankur Agarwal, Jouni Malinen, Brian Weis, Russ Housley, and Alfred Hoenes for their in-depth reviews and interesting discussions in this problem space.

[8.](#) References

[8.1.](#) Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2104] Krawczyk, H., Bellare, M., and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", [RFC 2104](#), February 1997.
- [RFC3748] Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H. Levkowetz, "Extensible Authentication Protocol (EAP)", [RFC 3748](#), June 2004.
- [RFC4187] Arkko, J. and H. Haverinen, "Extensible Authentication Protocol Method for 3rd Generation Authentication and Key Agreement (EAP-AKA)", [RFC 4187](#), January 2006.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 5226](#),

Arkko, et al.

Expires May 22, 2009

[Page 19]

Internet-Draft

EAP-AKA'

November 2008

May 2008.

[3GPP.24.302]

3GPP, "3rd Generation Partnership Project; Technical Specification Group Core Network and Terminals; Access to the 3GPP Evolved Packet Core (EPC) via non-3GPP access networks; Stage 3; (Release 8)", 3GPP Draft Technical Specification 24.302 v 1.0.0, September 2008.

[3GPP.33.102]

3GPP, "3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; 3G Security; Security architecture (Release 8)", 3GPP Draft Technical Specification 33.102 v 8.0.0, June 2008.

[3GPP.33.402]

3GPP, "3GPP System Architecture Evolution (SAE); Security aspects of non-3GPP accesses; Release 8", 3GPP Draft Technical Specification 33.402 v 8.0.0, June 2008.

[FIPS.180-2.2002]

National Institute of Standards and Technology, "Secure Hash Standard", FIPS PUB 180-2, August 2002, <<http://csrc.nist.gov/publications/fips/fips180-2/fips180-2.pdf>>.

8.2. Informative References

- [RFC4186] Haverinen, H. and J. Salowey, "Extensible Authentication Protocol Method for Global System for Mobile Communications (GSM) Subscriber Identity Modules (EAP-SIM)", [RFC 4186](#), January 2006.
- [RFC4284] Adrangi, F., Lortz, V., Bari, F., and P. Eronen, "Identity Selection Hints for the Extensible Authentication Protocol (EAP)", [RFC 4284](#), January 2006.
- [RFC4306] Kaufman, C., "Internet Key Exchange (IKEv2) Protocol", [RFC 4306](#), December 2005.
- [RFC5113] Arkko, J., Aboba, B., Korhonen, J., and F. Bari, "Network Discovery and Selection Problem", [RFC 5113](#), January 2008.
- [RFC5247] Aboba, B., Simon, D., and P. Eronen, "Extensible Authentication Protocol (EAP) Key Management Framework", [RFC 5247](#), August 2008.
- [3GPP.23.003] 3GPP, "3rd Generation Partnership Project; Technical

Arkko, et al.

Expires May 22, 2009

[Page 20]

Internet-Draft

EAP-AKA'

November 2008

Specification Group Core Network and Terminals; Numbering, addressing and identification (Release 8)", 3GPP Draft Technical Specification 23.003 v 8.0.0, June 2008.

[FIPS.180-1.1995]

National Institute of Standards and Technology, "Secure Hash Standard", FIPS PUB 180-1, April 1995, <<http://www.itl.nist.gov/fipspubs/fip180-1.htm>>.

Appendix A. Changes from [RFC 4187](#)

The changes to [RFC 4187](#) relate only to the bidding down prevention support defined in [Section 4](#). In particular, this document does not change how the Master Key (MK) is calculated in [RFC 4187](#) (it uses CK and IK, not CK' and IK'); neither is any processing of the AMF bit added to [RFC 4187](#).

[Appendix B](#). Importance of Explicit Negotiation

Choosing between the traditional and revised AKA key derivation functions is easy when their use is unambiguously tied to a particular radio access network, e.g Long Term Evolution (LTE) as defined by 3GPP or evolved High Rate Packet Data (eHRPD) as defined by 3GPP2. There is no possibility for interoperability problems if this radio access network is always used in conjunction with new protocols that cannot be mixed with the old ones; clients will always know whether they are connecting to the old or new system.

However, using the new key derivation functions over EAP introduces several degrees of separation, making the choice of the correct key derivation functions much harder. Many different types of networks employ EAP. Most of these networks have no means to carry any information about what is expected from the authentication process. EAP itself is severely limited in carrying any additional information, as noted in [\[RFC4284\]](#) and [\[RFC5113\]](#). Even if these networks or EAP were extended to carry additional information, it would not affect millions of deployed access networks and clients attaching to them.

Simply changing the key derivation functions that EAP-AKA [\[RFC4187\]](#) uses would cause interoperability problems with all of the existing implementations. Perhaps it would be possible to employ strict separation into domain names that should be used by the new clients and networks. Only these new devices would then employ the new key derivation mechanism. While this can be made to work for specific cases, it would be an extremely brittle mechanism, ripe to result in

problems whenever client configuration, routing of authentication requests, or server configuration does not match expectations. It also does not help to assume that the EAP client and server are running a particular release of 3GPP network specifications. Network vendors often provide features from the future releases early or do

not provide all features of the current release. And obviously, there are many EAP and even some EAP-AKA implementations that are not bundled with the 3GPP network offerings. In general, these approaches are expected to lead to hard-to-diagnose problems and increased support calls.

Authors' Addresses

Jari Arkko
Ericsson
Jorvas 02420
Finland

Email: jari.arkko@piuha.net

Vesa Lehtovirta
Ericsson
Jorvas 02420
Finland

Email: vesa.lehtovirta@ericsson.com

Pasi Eronen
Nokia Research Center
P.O. Box 407
FIN-00045 Nokia Group
Finland

Email: pasi.eronen@nokia.com

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

