

EMU Working Group
Internet-Draft
Intended status: Standards Track
Expires: November 25, 2012

S. Hartman, Ed.
Painless Security
T. Clancy
Department of Electrical
Engineering and Computer Science
K. Hooper
Motorola Solutions, Inc.
May 24, 2012

Channel Binding Support for EAP Methods
draft-ietf-emu-chbind-16.txt

Abstract

This document defines how to implement channel bindings for Extensible Authentication Protocol (EAP) methods to address the lying Network Access Service (NAS) as well as the lying provider problem.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on November 25, 2012.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1. Introduction	5
2. Terminology	6
3. Problem Statement	6
4. Channel Bindings	8
4.1. Types of EAP Channel Bindings	8
4.2. Channel Bindings in the Secure Association Protocol	10
4.3. Channel Bindings Scope	11
5. Channel Binding Process	13
5.1. Protocol Operation	13
5.2. Channel Binding Consistency Check	15
5.3. EAP Protocol	16
5.3.1. Channel Binding Codes	18
5.3.2. Namespace Identifiers	18
5.3.3. RADIUS Namespace	18
6. System Requirements	19
6.1. General Transport Protocol Requirements	19
6.2. EAP Method Requirements	20
7. Channel Binding TLV	20
7.1. Requirements for Lower-Layer Bindings	20
7.2. EAP Lower Layer Attribute	21
8. AAA-Layer Bindings	21
9. Security Considerations	22
9.1. Trust Model	22
9.2. Consequences of Trust Violation	24
9.3. Bid-Down Attacks	25
9.4. Privacy Violations	25
10. Operations and Management Considerations	26
11. IANA Considerations	26
11.1. EAP Lower Layers Registry	27
11.2. RADIUS Registration	27
12. Acknowledgments	27
13. References	28
13.1. Normative References	28
13.2. Informative References	28

Appendix A. Attacks Prevented by Channel Bindings	29
A.1. Enterprise Subnetwork Masquerading	30
A.2. Forced Roaming	30
A.3. Downgrading attacks	30
A.4. Bogus Beacons in IEEE 802.11r	31
A.5. Forcing false authorization in IEEE 802.11i	31
Appendix B. Change History	31
B.1. Changes Since 09	31
B.2. Changes since Version 06	32
B.3. Changes since version 04	32
Authors' Addresses	32

1. Introduction

The so-called "lying NAS" problem is a well-documented problem with the current Extensible Authentication Protocol (EAP) architecture [RFC3748] when used in pass-through authenticator mode. Here, a Network Access Server (NAS), or pass-through authenticator, may represent one set of information (e.g. network identity, capabilities, configuration, etc) to the backend Authentication, Authorization, and Accounting (AAA) infrastructure, while representing contrary information to EAP peers. Another possibility is that the same false information could be provided to both the EAP peer and EAP server by the NAS. A "lying" entity can also be located anywhere on the AAA path between the NAS and the EAP server.

This problem results when the same credentials are used to access multiple services that differ in some interesting property. The EAP server learns which client credentials are in use. The client knows which EAP credentials are used, but cannot distinguish between servers that use those credentials. For methods that distinguish between client and server credentials, either using different server credentials for access to the different services or having client credentials with access to a disjoint set of services can potentially defend against the attack.

As a concrete example, consider an organization with two different IEEE 802.11 wireless networks. One is a relatively low-security network for accessing the web while the other has access to valuable confidential information. An access point on the web network could act as a lying NAS, sending the SSID of the confidential network in its beacons. This access point could gain an advantage by doing so if it tricks clients intending to connect to the confidential network to connect to it and disclose confidential information.

A similar problem can be observed in the context of roaming. Here, the lying entity is located in a visited service provider network, e.g. attempting to lure peers to connect to the network based on false advertized roaming rates. This is referred to as "the lying provider" problem in the remainder of this document. The lying entity's motivation often is financial; the entity may be paid whenever peers roam to its service. However a lying entity in a provider network can also gain access to traffic that it might not otherwise see.

This document defines and implements EAP channel bindings to solve the lying NAS and the lying provider problems, using a process in which the EAP peer provides information about the characteristics of the service provided by the authenticator to the AAA server protected within the EAP method. This allows the server to verify the

authenticator is providing information to the peer that is consistent with the information received from this authenticator as well as the information stored about this authenticator. "AAA Payloads" defined in [I-D.clancy-emu-aaapay] served as the starting point for the mechanism proposed in this specification to carry this information.

2. Terminology

In this document, several words are used to signify the requirements of the specification. These words are often capitalized. 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. Problem Statement

In a [RFC4017] compliant EAP authentication, the EAP peer and EAP server mutually authenticate each other, and derive keying material. However, when operating in pass-through mode, the EAP server can be far removed from the authenticator both in terms of network distance and number of entities who need to be trusted in order to establish trusted communication. A malicious or compromised authenticator may represent incorrect information about the network to the peer in an effort to affect its operation in some way. Additionally, while an authenticator may not be compromised, other compromised elements in the network (such as proxies) could provide false information to the authenticator that it could simply be relaying to EAP peers. Hence, the goal must be to ensure that the authenticator is providing correct information to the EAP peer during the initial network discovery, selection, and authentication.

There are two different types of networks to consider: enterprise networks and service provider networks. In enterprise networks, assuming a single administrative domain, it is feasible for an EAP server to have information about all the authenticators in the network. In service provider networks, global knowledge is infeasible due to indirection via roaming. When a peer is outside its home administrative domain, the goal is to ensure that the level of service received by the peer is consistent with the contractual agreement between the two service providers. The same EAP server may need to support both types of networks. For example an enterprise may have a roaming agreement permitting its users to use the networks of third-party service providers. In these situations, the EAP server may authenticate for an enterprise and provider network.

The following are example attacks possible by presenting false

network information to peers.

- o Enterprise Network: A corporate network may have multiple virtual LANs (VLANs) available throughout their campus network, and have IEEE 802.11 access points connected to each VLAN. Assume one VLAN connects users to the firewalled corporate network, while the other connects users to a public guest network. The corporate network is assumed to be free of adversarial elements, while the guest network is assumed to possibly have malicious elements. Access Points on both VLANs are serviced by the same EAP server, but broadcast different SSIDs to differentiate. A compromised access point connected to the guest network but not the corporate network could advertise the SSID of the corporate network in an effort to lure peers to connect to a network with a false sense of security regarding their traffic. Conditions and further details of this attack can be found in the Appendix.
- o Enterprise network: The EAP GSS-API mechanism [I-D.ietf-abfab-gss-eap] mechanism provides a way to use EAP to authenticate to mail servers, instant messaging servers and other non-network services. Without EAP channel binding, an attacker could trick the user into connecting to a relatively untrusted service instead of a relatively trusted service. For example the instant messaging service could impersonate the mail server.
- o Service Provider Network: An EAP-enabled mobile phone provider could advertise very competitive flat rates but send per minute rates to the home server, thus, luring peers to connect to their network and overcharging them. In more elaborate attacks, peers can be tricked into roaming without their knowledge. For example, a mobile phone provider operating along a geo-political boundary could boost their cell towers' transmission power and advertise the network identity of the neighboring country's indigenous provider. This would cause unknowing handsets to associate with an unintended operator, and consequently be subject to high roaming fees without realizing they had roamed off their home provider's network. These types of scenarios can be considered as "lying provider" problem, because here the provider configures its NAS to broadcast false information. For the purpose of channel bindings as defined in this draft, it does not matter which local entity (or entities) is "lying" in a service provider network (local NAS, local authentication server and/or local proxies), because the only information received from the visited network that is verified by channel bindings is the information the home authentication server received from the last hop in the communication chain. In other words, channel bindings enable the detection of inconsistencies in the information from a visited network, but cannot determine which entity is lying. Naturally,

channel bindings for EAP methods can only verify the endpoints and, if desirable, intermediate hops need to be protected by the employed AAA protocol.

- o Enterprise and provider networks: In a situation where an enterprise has roaming agreements with providers, a compromised access point in a provider network could masquerade as the enterprise network in an attempt to gain confidential information. Today this could potentially be solved by using different credentials for internal and external access. Depending on the type of credential this may introduce usability or man-in-the-middle security issues.

To address these problems, a mechanism is required to validate unauthenticated information advertised by EAP authenticators.

4. Channel Bindings

EAP channel bindings seek to authenticate previously unauthenticated information provided by the authenticator to the EAP peer, by allowing the peer and server to compare their perception of network properties in a secure channel.

It should be noted that the definition of EAP channel bindings differs somewhat from channel bindings documented in [RFC5056], which seek to securely bind together the end points of a multi-layer protocol, allowing lower layers to protect data from higher layers. Unlike [RFC5056], EAP channel bindings do not ensure the binding of different layers of a session but rather the information advertised to EAP peer by an authenticator acting as pass-through device during an EAP execution. The term channel bindings was independently adopted by these two related concepts; by the time the conflict was discovered, a wide body of literature existed for each usage. EAP channel bindings could be used to provide RFC 5056 channel bindings. In particular, an inner EAP method could be bound to an outer method by including the RFC 5056 channel binding data for the outer channel in the inner EAP method's channel bindings. Doing so would provide a facility similar to EAP cryptographic binding, except that a man-in-the-middle could not extract the inner method from the tunnel. This specification does not weigh the advantages of doing so nor specify how to do so; the example is provided only to illustrate how EAP channel binding and RFC 5056 channel binding overlap.

4.1. Types of EAP Channel Bindings

There are two categories of approach to EAP channel bindings:

- o After keys have been derived during an EAP execution, the peer and server can, in an integrity-protected channel, exchange plaintext information about the network with each other, and verify consistency and correctness.
- o The peer and server can both uniquely encode their respective view of the network information without exchanging it, resulting into an opaque blob that can be included directly into the derivation of EAP session keys.

Both approaches are only applicable to key deriving EAP methods and both have advantages and disadvantages. Various hybrid approaches are also possible. Advantages of exchanging plaintext information include:

- o It allows for policy-based comparisons of network properties, rather than requiring precise matches for every field, which achieves a policy-defined consistency, rather than bitwise equality. This allows network operators to define which properties are important and even verifiable in their network.
- o EAP methods that support extensible, integrity-protected channels can easily include support for exchanging this network information. In contrast, direct inclusion into the key derivation would require more extensive revisions to existing EAP methods or a wrapper EAP method.
- o Given it doesn't affect the key derivation, this approach facilitates debugging, incremental deployment, backward compatibility and a logging mode in which verification results are recorded but do not have an effect on the remainder of the EAP execution. The exact use of the verification results can be subject to the network policy. Additionally, consistent information canonicalization and formatting for the key derivation approach would likely cause significant deployment problems.

The following are advantages of directly including channel binding information in the key derivation:

- o EAP methods not supporting extensible, integrity-protected channels could still be supported, either by revising their key derivation, revising EAP, or wrapping them in a universal method that supports channel binding.
- o It can guarantee proper channel information, since subsequent communication would be impossible if differences in channel information yielded different session keys on the EAP peer and server.

4.2. Channel Bindings in the Secure Association Protocol

This document describes channel bindings performed by transporting channel binding information as part of an integrity-protected exchange within an EAP method. Alternatively, some future document could specify a mechanism for transporting channel bindings within the lower layer's secure association protocol. Such a specification would need to describe how channel bindings are exchanged over the lower layer protocol between the peer and authenticator. In addition, since the EAP exchange concludes before the secure association protocol begins, a mechanism for transporting the channel bindings from the authenticator to the EAP server needs to be specified. A mechanism for transporting a protected result from the EAP server, through the authenticator, back to the peer needs to be specified.

The channel bindings **MUST** be transported with integrity protection based on a key known only to the peer and EAP server. The channel bindings **SHOULD** be confidentiality protected using a key known only to the peer and EAP server. For the system to function, the EAP server or AAA server needs access to the channel binding information from the peer as well as the AAA attributes and a local database described later in this document.

The primary advantage of sending channel bindings as part of the secure association protocol is that EAP methods need not be changed. The disadvantage is that a new AAA exchange is required, and secure association protocols need to be changed. As the result of the secure association protocol change, every NAS needs to be upgraded to support channel bindings within the secure association protocol.

For many deployments, changing all the NASes is expensive and adding channel binding support to enough EAP methods to meet the goals of the deployment will be cheaper. However for deployment of new equipment, or especially deployment of a new lower layer technology, changing the NASes may be cheaper than changing EAP methods. Especially if such a deployment needed to support a large number of EAP methods, sending channel bindings in the secure association protocol might make sense. Sending channel bindings in the secure association protocol can work even with ERP [RFC5296] in which previously established EAP key material is used for the secure association protocol without carrying out any EAP method during re-authentication.

If channel bindings using a secure association protocol is specified, semantics as well as the set of information that peers exchange can be shared with the mechanism described in this document.

4.3. Channel Bindings Scope

The scope of EAP channel bindings differs somewhat depending on the type of deployment in which they are being used. In enterprise networks, they can be used to authenticate very specific properties of the authenticator (e.g. MAC address, supported link types and data rates, etc), while in service provider networks they can generally only authenticate broader information about a roaming partner's network (e.g. network name, roaming information, link security requirements, etc). The reason for the difference has to do with the amount of information about the authenticator and/or network to which the peer is connected the home EAP server is expected to have access to. In roaming cases, the home server is likely to only have access to information contained in their roaming agreements.

With any multi-hop AAA infrastructure, many of the NAS-specific AAA attributes are obscured by the AAA proxy that's decrypting, reframing, and retransmitting the underlying AAA messages. Especially service provider networks are affected by this and the AAA information received from the last hop may not contain much verifiable information any longer. For example, information carried in AAA attributes such as the NAS IP address may have been lost in transition and are thus not known to the EAP server. Even worse, information may still be available but be useless, for example representing the identity of a device on a private network or a middlebox. This affects the ability of the EAP server to verify specific NAS properties. However, often verification of the MAC or IP address of the NAS is not useful for improving the overall security posture of a network. More often the best approach is to make policy decisions about services being offered to peers. For example, in an IEEE 802.11 network, the EAP server may wish to ensure that peers connecting to the corporate intranet are using secure link-layer encryption, while link-layer security requirements for peers connecting to the guest network could be less stringent. These types of policy decisions can be made without knowing or being able to verify the IP address of the NAS through which the peer is connecting.

The properties of the network that the peer wishes to validate depend on the specific deployment. In a mobile phone network, peers generally don't care what the name of the network is, as long as they can make their phone call and are charged the expected amount for the call. However, in an enterprise network the administrators of a peer may be more concerned with specifics of where their network traffic is being routed and what VLAN is in use. To establish policies surrounding these requirements administrators would capture some attribute such as SSID to describe the properties of the network they care about. Channel bindings could validate the SSID. The

administrator would need to make sure that the network guarantees that when an authenticator trusted by the AAA infrastructure to offer a particular SSID to clients does offer this SSID, that network has the intended properties. Generally it is not possible for channel bindings to detect lying NAS behavior when the NAS is authorized to claim a particular service. That is, if the same physical authenticator is permitted to advertise two networks, the AAA infrastructure is unlikely to be able to determine when this authenticator lies.

As discussed in the next section, some of the most important information to verify cannot come from AAA attributes but instead comes from local configuration. For example in the mobile phone case, the expected roaming rate cannot come from the roaming provider without being verified against the contract between the two providers. Similarly, in an enterprise, the SSID a particular access point is expected to advertise is a matter of configuration rather than something that can be trusted because it is included in an AAA exchange.

The peer and authenticator do not initially have a basis for trust. The peer has a credential with the EAP server that forms a basis for trust. The EAP server and authenticator have a potentially indirect trust path using the AAA infrastructure. Channel binding leverages the trust between the peer and EAP server to build trust in certain attributes between the peer and authenticator.

Channel bindings can be important for forming areas of trust, especially when provider networks are involved, and exact information is not available to the EAP server. Without channel bindings, all entities in the system need to be held to the standards of the most trusted entity that could be accessed using the EAP credential. Otherwise, a less trusted entity can impersonate a more trusted entity. However when channel bindings are used, the EAP server can use information supplied by the peer, AAA protocols and local database to distinguish less trusted entities from more trusted entities. One possible deployment involves being able to verify a number of characteristics about relatively trusted entities while for other entities simply verifying that they are less trusted.

Any deployment of channel bindings should take into consideration both what information the EAP server is likely to know or have access to, and also what type of network information the peer would want and need authenticated.

5. Channel Binding Process

This section defines the process for verifying channel binding information during an EAP authentication. The protocol uses the approach where plaintext data is exchanged, since it allows channel bindings to be used more flexibly in varied deployment models (see Section 4.1). In the first subsection, the general communication infrastructure is outlined, the messages used for channel binding verifications are specified, and the protocol flows are defined. The second subsection explores the difficulties of checking the different pieces of information that are exchanged during the channel binding protocol for consistency. The third subsection describes the information carried in the EAP exchange.

5.1. Protocol Operation

Channel bindings are always provided between two communication endpoints, here the EAP peer and the EAP server, who communicate through an authenticator typically in pass-through mode. Specifications treat the AAA server and EAP server as distinct entities. However there is no standardized protocol for the AAA server and EAP server to communicate with each other. For the channel binding protocol presented in this draft to work, the EAP server needs to be able to access information from the AAA server that is utilized during the EAP session (i2 below) and a local database. For example, the EAP server and the local database can be co-located with the AAA server, as illustrated in Figure 1. An alternate architecture would be to provide a mechanism for the EAP server to inform the AAA server what channel binding attributes were supplied and the AAA server to inform the EAP server about what channel binding attributes it considered when making its decision.




Figure 1: Overview of Channel Binding Protocol

During network advertisement, selection, and authentication, the authenticator presents unauthenticated information, labeled i1, about the network to the peer. Message i1 could include an authenticator identifier and the identity of the network it represents, in addition to advertised network information such as offered services and roaming information. Information may be communicated implicitly in i1, such as the type of media in use. As there is no established trust relationship between the peer and authenticator, there is no way for the peer to validate this information.

Additionally, during the transaction the authenticator presents a number of information properties in form of AAA attributes about itself and the current request to the AAA infrastructure which may or may not be valid. This information is labeled i2. Message i2 is the information the AAA server receives from the last hop in the AAA proxy chain which is not necessarily the authenticator.

AAA hops between the authenticator and AAA server can validate some of i2. Whether the AAA server will be able to depend on this depends significantly on the business relationship executed with these proxies and on the structure of the AAA network.

The local database is perhaps the most important part of this system. In order for the EAP server or AAA server to know whether i1 and i2 are correct, they need access to trustworthy information, since an authenticator could include false information in both i1 and i2. Additional reasons why such a database is necessary for channel bindings to work are discussed in the next subsection. The information contained within the database could involve wildcards. For example, this could be used to check whether WiFi access points on a particular IP subnet all use a specific SSID. The exact IP address is immaterial, provided it is on the correct subnet.

During an EAP method execution with channel bindings, the peer sends i1 to the EAP server using the mechanism described in Section 5.3. the EAP server verifies the consistency of i1 provided by the peer, i2 provided by the authenticator, and the information in the local database. Upon the check, the EAP server sends a message to the peer indicating whether the channel binding validation check succeeded or failed and includes the attributes that were used in the check. The message flow is illustrated in Figure 1.

Above, the EAP server is described as performing the channel binding validation. In most deployments, this will be a necessary implementation constraint. The EAP exchange needs to include an indication of channel binding success or failure. Most existing implementations do not have a way to have an exchange between the EAP server and another AAA entity during the EAP server's processing of a

single EAP message. However another AAA entity can provide information to the EAP server to make its decision.

If the compliance of i1 or i2 information with the authoritative policy source is mandatory and a consistency check failed, then after sending a protected indication of failed consistency, the EAP server MUST send an EAP-Failure message to terminate the session. If the EAP server is otherwise configured, it MUST allow the EAP session to complete normally, and leave the decision about network access up to the peer's policy. If i1 or i2 does not comply with policy, the EAP server MUST NOT list information that failed to comply in the set of information used to perform channel binding. In this case the EAP server SHOULD indicate channel binding failure; this requirement may be upgraded to a MUST in the future.

5.2. Channel Binding Consistency Check

The validation check that is the core of the channel binding protocol described in the previous subsection, consists of two parts in which the server checks whether:

1. the authenticator is lying to the peer, i.e. i1 contains false information,
2. the authenticator or any entity on the AAA path to the AAA server provides false information in form of AAA attributes, i.e. i2 contains false information,

These checks enable the EAP server to detect lying NAS/authenticator in enterprise networks and lying providers in service provider networks.

Checking the consistency of i1 and i2 is nontrivial, as has been pointed out already in [HC07]. First, i1 can contain any type of information propagated by the authenticator, whereas i2 is restricted to information that can be carried in AAA attributes. Second, because the authenticator typically communicates over different link layers with the peer and the AAA infrastructure, different type of identifiers and addresses may have been presented to both communication endpoints. Whether these different identifiers and addresses belong to the same device cannot be directly checked by the EAP server or AAA server without additional information. Finally, i2 may be different from the original information sent by the authenticator because of en route processing or malicious modifications. As a result, in the service provider model, typically the i1 information available to the EAP server can only be verified against the last-hop portion of i2, or values propagated by proxy servers. In addition, checking the consistency of i1 and i2 alone is

insufficient because an authenticator could lie to both, the peer and the EAP server, i.e. i1 and i2 may be consistent but both contain false information.

A local database is required to leverage the above mentioned shortcomings and support the consistency and validation checks. In particular, information stored for each NAS/authenticator (enterprise scenario) or each roaming partner (service provider scenario) enables a comparison of any information received in i1 with AAA attributes in i2 as well as additionally stored AAA attributes that might have been lost in transition. Furthermore, only such a database enables the EAP server and AAA server to check the received information against trusted information about the network including roaming agreements.

Section 7 describes lower-layer specific properties that can be exchanged as a part of i1. Section 8 describes specific AAA attributes that can be included and evaluated in i2. The EAP server reports back the results from the channel binding validation check that compares the consistency of all the values with those in the local database. The challenges of setting up such a local database are discussed in Section 10.

5.3. EAP Protocol

EAP methods supporting channel binding consistent with this specification provide a mechanism for carrying channel binding data from the peer to the EAP server and a channel binding response from the EAP server to the peer. The specifics of this mechanism are dependent on the method, although the content of the channel binding data and channel binding response are defined by this section.

Typically the lower layer will communicate a set of attributes to the EAP implementation on the peer that should be part of channel binding. The EAP implementation may need to indicate to the lower layer that channel binding information cannot be sent. Reasons for failing to send channel binding information include an EAP method that does not support channel binding is selected, or channel binding data is too big for the EAP method selected. Peers SHOULD provide appropriate policy controls to select channel binding or mandate its success.

The EAP server receives the channel binding data and performs the validation. The EAP method provides a way to return a response; the channel binding response uses the same basic format as the channel binding data.

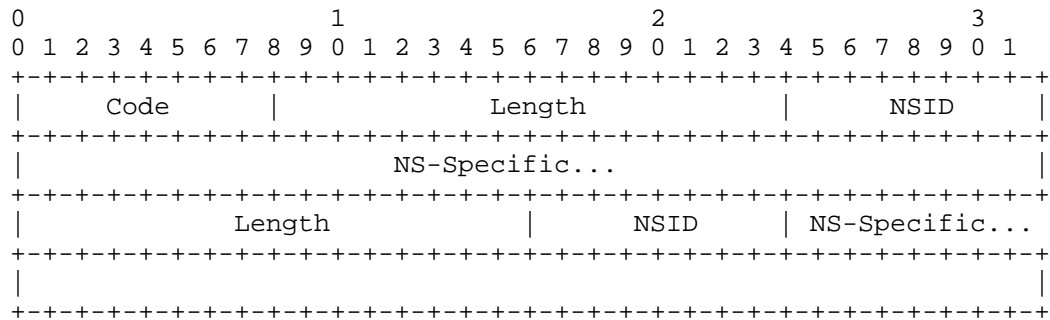


Figure 2: Channel Binding Encoding

Both the channel binding data and response use the format illustrated in Figure 2. The protocol starts with a one byte code; see Section 5.3.1. Then for each type of attributes contained in the channel binding data, the following information is encoded:

Length: Two octets of length in network byte order, indicating the length of the NS-SPECIFIC data. The NSID and length octets are not included.

NSID: One octet describing the namespace from which the attributes are drawn. See Section 5.3.3 for a description of how to encode RADIUS attributes in channel binding data and responses. RADIUS uses a namespace identifier of 1.

NS-SPECIFIC: The encoding of the attributes in a manner specific to the type of attribute.

A given NSID MUST NOT appear more than once in a channel binding data or channel binding response. Instead, all NS-SPECIFIC data for a particular NSID must occur inside one (NSID, Length, NS-Specific) field. This set of fields may be repeated if multiple namespaces are included.

In channel binding data, the code is set to 1 (channel binding data) and the full attributes and values that the peer wishes the EAP server to validate are included.

In a channel binding response, the server selects the code; see Section 5.3.1. For successful channel binding, the server returns code 2. The set of attributes that the EAP server returns depend on the code. For success, the server returns the attributes that were considered by the server in making the determination that channel bindings are successfully validated; attributes that the server is

unable to check or that failed to validate against what is sent by the peer MUST NOT be returned in a success response. Generally, servers will not return a success response if any attributes were checked and failed to validate those specified by the peer. Special circumstances such as a new attribute being phased in at a server MAY require servers to return success when such an attribute fails to validate. The server returns the value supplied by the peer when returning an attribute in channel binding responses.

For channel binding failure (code 3), the server SHOULD include any attributes that were successfully validated. This code means that server policy indicates that the attributes sent by the client do not accurately describe the authenticator. Servers MAY include no attributes in this response, for example if all the attributes supplied by the peer that the server can check failed to be consistent.

Peers MUST treat unknown codes as Channel binding Failure. Peers MUST ignore differences between attribute values sent in the channel binding data and those sent in the response. Peers and servers MUST ignore any attributes contained in a field with an unknown NSID. Peers MUST ignore any attributes in a response not present in the channel binding data.

5.3.1. Channel Binding Codes

Code	Meaning
1	Channel Binding data from client
2	Channel binding response: success
3	Channel binding response: failure

5.3.2. Namespace Identifiers

ID	Namespace	Reference
1	RADIUS	Section 5.3.3
255	Private Use	

5.3.3. RADIUS Namespace

RADIUS AVPs are encoded with a one-octet attribute type followed by a one-octet length followed by the value of the RADIUS attribute being encoded. The length includes the type and length octets; the minimum

legal length is 3. Attributes are concatenated to form the namespace specific portion of the packet.




Figure 3: RADIUS AVP Encoding

The full value of an attribute is included in the channel binding data and response.

6. System Requirements

This section defines requirements on components used to implement the channel bindings protocol.

The channel binding protocol defined in this document must be transported after keying material has been derived between the EAP peer and server, and before the peer would suffer adverse affects from joining an adversarial network. This document describes a protocol for performing channel binding within EAP methods. As discussed in Section 4.2, an alternative approach for meeting this requirement is to perform channel bindings during the secure association protocol of the lower layer.

6.1. General Transport Protocol Requirements

The transport protocol for carrying channel binding information MUST support end-to-end (i.e. between the EAP peer and server) message integrity protection to prevent the adversarial NAS or AAA device from manipulating the transported data. The transport protocol SHOULD provide confidentiality. The motivation for this is that the channel bindings could contain private information, including peer identities, which SHOULD be protected. If confidentiality cannot be provided, private information MUST NOT be sent as part of the channel binding information.

Any transport needs to be careful not to exceed the MTU for its lower-layer medium. In particular, if channel binding information is exchanged within protected EAP method channels, these methods may or may not support fragmentation. In order to work with all methods, the channel binding messages must fit within the available payload. For example, if the EAP MTU is 1020 octets, and EAP-GPSK is used as

the authentication method, and maximal-length identities are used, a maximum of 384 octets are available for conveying channel binding information. Other methods, such as EAP-TTLS, support fragmentation and could carry significantly longer payloads.

6.2. EAP Method Requirements

If transporting data directly within an EAP method, it MUST be able to carry integrity protected data from the EAP peer to server. EAP methods SHOULD provide a mechanism to carry protected data from server to peer. EAP methods MUST exchange channel binding data with the AAA subsystem hosting the EAP server. EAP methods MUST be able to import channel binding data from the lower layer on the EAP peer.

7. Channel Binding TLV

This section defines some channel binding TLVs. While message `il` is not limited to AAA attributes, for the sake of tangible attributes that are already in place, this section discusses AAA AVPs that are appropriate for carrying channel bindings (i.e. data from `il` in Section 5).

For any lower-layer protocol, network information of interest to the peer and server can be encapsulated in AVPs or other defined payload containers. The appropriate AVPs depend on the lower layer protocol as well as on the network type (i.e. enterprise network or service provider network) and its application.

7.1. Requirements for Lower-Layer Bindings

Lower-layer protocols MUST support EAP in order to support EAP channel bindings. These lower layers MUST support EAP methods that derive keying material, as otherwise no integrity-protected channel would be available to execute the channel bindings protocol. Lower-layer protocols need not support traffic encryption, since this is independent of the authentication phase.


The data conveyed within the AVP type MUST NOT conflict with the externally-defined usage of the AVP. Additional TLV types MAY be defined for values that are not communicated within AAA attributes.

In general, lower layers will need to specify what information should be included in `il`. Existing lower layers will probably require new documents to specify this information. Lower layer specifications need to include sufficient information in `il` to uniquely identify which lower layer is involved. The preferred way to do this is to include the `eap-lower-layer` attribute defined in the next section.

This MUST be included in i1 unless an attribute specific to a particular lower layer is included in i1.

7.2. EAP Lower Layer Attribute

A new RADIUS attribute is defined to carry information on which EAP lower layer is used for this EAP authentication. This Attribute provides information relating to the lower layer over which EAP is transported. This Attribute MAY be sent by the NAS to the RADIUS server in an Access-Request or an Accounting-Request packet. A summary of the EAP-Lower-Layer Attribute format is shown below. The fields are transmitted from left to right.



The code is TBD, the length is 6 and the value is a 32-bit unsigned integer in network byte order. The value specifies the EAP lower layer in use. Values are taken from the IANA registry established in Section 11.1.

8. AAA-Layer Bindings

This section discusses which AAA attributes in a AAA Access-Request message can and should be validated by a EAP server (i.e. data from i2 in Section 5). As noted before, this data can be manipulated by AAA proxies either to enable functionality (e.g. removing realm information after messages have been proxied) or maliciously (e.g. in the case of a lying provider). As such, this data cannot always be easily validated. However as thorough of a validation as possible should be conducted in an effort to detect possible attacks.

NAS-IP-Address: This value is typically the IP address of the authenticator, but in a proxied connection it likely will not match the source IP address of an Access-Request. A consistency check MAY verify the subnet of the IP address was correct based on the last-hop proxy.

NAS-IPv6-Address: This value is typically the IPv6 address of the authenticator, but in a proxied connection it likely will not match the source IPv6 address of an Access-Request. A consistency check MAY verify the subnet of the IPv6 address was correct based on the last-hop proxy.

NAS-Identifier: This is an identifier populated by the NAS to identify the NAS to the AAA server; it SHOULD be validated against the local database.

NAS-Port-Type: This specifies the underlying link technology. It SHOULD be validated against the value received from the peer in the information exchange, and against a database of authorized link-layer technologies.

9. Security Considerations

This section discusses security considerations surrounding the use of EAP channel bindings.

9.1. Trust Model

In the considered trust model, EAP peer and authentication server are honest while the authenticator is maliciously sending false information to peer and/or server. In the model, the peer and server trust each other, which is not an unreasonable assumption, considering they already have a trust relationship. The following are the trust relationships:

- o The server trusts that the channel binding information received from the peer is the information that the peer received from the authenticator.
- o The peer trusts the channel binding result received from the server.
- o The server trusts the information contained within its local database.

In order to establish the first two trust relationships during an EAP execution, an EAP method MUST provide the following:

- o mutual authentication between peer and server
- o derivation of keying material including a key for integrity protection of channel binding messages known to the peer and EAP server but not the authenticator
- o sending channel binding request from peer to server over an integrity-protected channel

- o sending the channel binding result from server to peer over an integrity-protected channel

This trust model is a significant departure from the standard EAP model. In many EAP deployments today attacks where one authenticator can impersonate another are not a significant concern because all authenticators provide the same service. A authenticator does not gain significant advantage by impersonating another authenticator. The use of EAP in situations where different authenticators provide different services may give an attacker who can impersonate a authenticator greater advantage. The system as a whole needs to be analyzed to evaluate cases where one authenticator may impersonate another and to evaluate the impact of this impersonation.

One attractive implementation strategy for channel binding is to add channel binding support to a tunnel method which can tunnel an inner EAP authentication. This way, channel binding can be achieved with any method that can act as an inner method even if that inner method does not have native channel binding support. The requirement for mutual authentication and key derivation is at the layer of EAP that actually performs the channel binding. Tunnel methods sometimes use cryptographic binding, a process where a peer proves that the peer for the outer method is the same as the peer for an inner method to tie authentication at one layer together with an inner layer. Cryptographic binding does not always provide mutual authentication; its definition does not require the server to prove that the inner server and outer server are the same. Even when cryptographic binding does attempt to confirm that the inner and outer server are the same, the Master Session Key (MSK) from the inner method is typically used to protect the binding. An attacker such as an authenticator that wishes to subvert channel binding could establish an outer tunnel terminating at the authenticator. If the outer method tunnel terminates on the authenticator, the MSK is disclosed to the authenticator, which can typically attack cryptographic binding. If the authenticator controls cryptographic binding then it typically controls the channel binding parameters and results. If the channel binding process is used to differentiate one authenticator from another then the authenticator can claim to support services that it was not authorized to. This attack was not in scope for existing threat models for cryptographic binding because differentiated authenticators was not a consideration. Thus, existing cryptographic binding does not typically provide mutual authentication of the inner method server required for channel binding. Other methods besides cryptographic binding are available to provide mutual authentication required by channel binding. As an example, if server certificates are validated and names checked, mutual authentication can be provided directly by the tunnel.

9.2. Consequences of Trust Violation

If any of the trust relationships listed in Section 9.1 are violated, channel binding cannot be provided. In other words, if mutual authentication with key establishment as part of the EAP method as well as protected database access are not provided, then achieving channel binding is not feasible.

Dishonest peers can only manipulate the first message *il* of the channel binding protocol. In this scenario, a peer sends *il'* to the server. If *il'* is invalid, the channel binding validation will fail. On the other hand if *il'* passes the validation, either the original *il* was wrong and *il'* corrected the problem or both *il* and *il'* constitute valid information. A peer could potentially gain an advantage in auditing or charging if both are valid and information from *il'* is used for auditing or charging. Such peers can be detected by including the information in *i2* and checking *il* against *i2*.

If information from *il* does not validate, an EAP server cannot generally determine whether the authenticator advertised incorrect information or whether the peer is dishonest. This should be considered before using channel binding validation failures to determine the reputation either of the peer or authenticator.

Dishonest servers can send EAP-Failure messages and abort the EAP authentication even if the received *il* is valid. However, servers can always abort any EAP session independent of whether channel binding is offered or not. On the other hand, dishonest servers can claim a successful validation even if *il* contains invalid information. This can be seen as collaboration of authenticator and server. Channel binding can neither prevent nor detect such attacks. In general such attacks cannot be prevented by cryptographic means and should be addressed using policies making servers liable for their provided information and services.

Additional network entities (such as proxies) might be on the communication path between peer and server and may attempt to manipulate the channel binding protocol. If these entities do not possess the keying material used for integrity protection of the channel binding messages, the same threat analysis applies as for the dishonest authenticators. Hence, such entities can neither manipulate single channel binding messages nor the outcome. On the other hand, entities with access to the keying material must be treated like a server in a threat analysis. Hence such entities are able to manipulate the channel binding protocol without being detected. However, the required knowledge of keying material is unlikely since channel binding is executed before the EAP method is

completed, and thus before keying material is typically transported to other entities.

9.3. Bid-Down Attacks

EAP methods that add channel binding will typically negotiate its use. Even for entirely new EAP methods designed with channel binding from the first version, some deployments may not use it. It is desirable to protect against attacks on the negotiation of channel bindings. An attacker including the NAS SHOULD NOT be able to prevent a peer and server who support channel bindings from using it.

Unfortunately existing EAP methods may make it difficult or impossible to protect against attacks on negotiation. For example, many EAP state machines will accept a success message at any point after key derivation to terminate authentication. EAP success methods are not integrity protected; an attacker who could insert a message can generate one. The NAS is always in a position to generate a success message. Common EAP servers take advantage of state machines accepting success messages even in cases where an EAP method might support a protected indication of success. It may be challenging to define channel binding support for existing EAP methods in a manner that permits peers to distinguish an old EAP server that sends a success indication and does not support channel binding from an attacker injecting a success indication.

9.4. Privacy Violations

While the channel binding information exchanged between EAP peer and EAP server (i.e. i1 and the result message) must always be integrity-protected it may not be encrypted. In the case that these messages contain identifiers of peer and/or network entities, the privacy property of the executed EAP method may be violated. Hence, in order to maintain the privacy of an EAP method, the exchanged channel binding information must be encrypted. If encryption is not available, private information is not sent as part of the channel binding information, as described in Section 6.1.

Privacy implications of attributes selected for channel binding need to be considered. Consider channel binding the username attribute. A peer sends a privacy protecting anonymous identifier in its EAP identity message, but sends the full username in the protected i1 message. However the authenticator would like to learn the full username. It makes a guess and sends that in i2 rather than the anonymous identifier. If the EAP server validates this attribute and fails when the username from the peer mismatches i2, then the EAP server confirms the authenticator's guess. Similar privacy exposures may result whenever one party is in a position to guess channel

binding information provided by another party.

10. Operations and Management Considerations

As with any extension to existing protocols, there will be an impact on existing systems. Typically the goal is to develop an extension that minimizes the impact on both development and deployment of the new system, subject to the system requirements. This section discusses the impact on existing devices that currently utilize EAP, assuming the channel binding information is transported within the EAP method execution.

The EAP peer will need an API between the EAP lower layer and the EAP method that exposes the necessary information from the NAS to be validated to the EAP peer, which can then feed that information into the EAP methods for transport. For example, an IEEE 802.11 system would need to make available the various information elements that require validation to the EAP peer which would properly format them and pass them to the EAP method. Additionally, the EAP peer will require updated EAP methods that support transporting channel binding information. While most method documents are written modularly to allow incorporating arbitrary protected information, implementations of those methods would need to be revised to support these extensions. Driver updates are also required so methods can access the required information.

No changes to the pass-through authenticator would be required.

The EAP server would need an API between the database storing NAS information and the individual EAP server. The database may already exist on the AAA server in which case the EAP server passes the parameters to the AAA server for validation. The EAP methods need to be able to export received channel binding information to the EAP server so it can be validated.

11. IANA Considerations

A new top level registry is created for "EAP Channel Binding Parameters." This registry consists of several sub registries.

The "Channel Binding Codes" sub-registry defines values for the code field in the channel binding data and channel binding response packet. See the table in Section 5.3.1 for initial registrations. This registry requires standards action [RFC5226] for new registrations. Early allocation [RFC4020] is allowed. An additional reference column should be added to the table for the registry,

pointing all codes in the initial registration to this specification. Valid values in this sub-registry range from 0-255; 0 is reserved.

The "Channel Binding Namespaces" sub-registry contains registrations for the NSID field in the channel binding data and channel binding response. Initial registrations are found in the table in Section 5.3.2. Registrations in this registry require IETF review. Valid values range from 0-255; 0 is reserved. As with the Channel Binding Codes sub-registry a reference column should be included and should point to this document for initial registrations.

11.1. EAP Lower Layers Registry

A new sub registry in the EAP Numbers registry at <http://www.iana.org/assignments/eap-numbers> is created for EAP Lower Layers. Registration requires expert review; the primary role of the expert is to prevent multiple registrations for the same lower layer.

The following table gives the initial registrations for this registry.

Value	Lower Layer
1	Wired IEEE 802.1X
2	IEEE 802.11 (no-pre-auth)
3	IEEE 802.11 (pre-authentication)
4	IEEE 802.16e
5	IKEv2
6	PPP
7	PANA (no pre-authentication) [RFC5191]
8	GSS-API [I-D.ietf-abfab-gss-eap]
9	PANA (pre-authentication) [RFC5873]

11.2. RADIUS Registration

A new RADIUS attribute is registered with the name EAP-Lower-Layer; TBD should be replaced with the number corresponding to this attribute. The RADIUS attributes are in the registry at <http://www.iana.org/assignments/radius-types/radius-types.xml>

12. Acknowledgments

The authors and editor would like to thank Bernard Aboba, Glen Zorn, Joe Salowey, Stephen Hanna, and Klaas Wierenga for their valuable inputs that helped to improve and shape this document over the time.

Sam Hartman's work on this specification is funded by JANET(UK).

The EAP-Lower-Layer attribute was taken from draft-aboba-radext-wlan [I-D.aboba-radext-wlan].

13. References

13.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3748] Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H. Levkowitz, "Extensible Authentication Protocol (EAP)", RFC 3748, June 2004.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC4020] Kompella, K. and A. Zinin, "Early IANA Allocation of Standards Track Code Points", BCP 100, RFC 4020, February 2005.

13.2. Informative References

- [I-D.aboba-radext-wlan] Aboba, B., Malinen, J., Congdon, P., and J. Salowey, "RADIUS Attributes for IEEE 802 Networks", draft-aboba-radext-wlan-15 (work in progress), October 2011.
- [I-D.clancy-emu-aaapay] Clancy, T., Lior, A., and G. Zorn, Ed., "EAP Method Support for Transporting AAA Payloads", Internet Draft draft-clancy-emu-aaapay-02, May 2009.
- [RFC4017] Stanley, D., Walker, J., and B. Aboba, "Extensible Authentication Protocol (EAP) Method Requirements for Wireless LANs", RFC 4017, March 2005.
- [RFC5056] Williams, N., "On the Use of Channel Bindings to Secure Channels", RFC 5056, November 2007.
- [HC07] Hoeper, K. and L. Chen, "Where EAP Security Claims Fail", ICST QShine, August 2007.

[80211U-D4.01]

"Information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications - Amendment 7: Interworking with External Networks", IEEE Draft Standard 802.11u, November 2008.

[I-D.ietf-abfab-gss-eap]

Hartman, S. and J. Howlett, "A GSS-API Mechanism for the Extensible Authentication Protocol", draft-ietf-abfab-gss-eap-06 (work in progress), April 2012.

[RFC5873] Ohba, Y. and A. Yegin, "Pre-Authentication Support for the Protocol for Carrying Authentication for Network Access (PANA)", RFC 5873, May 2010.

[RFC5296] Narayanan, V. and L. Dondeti, "EAP Extensions for EAP Re-authentication Protocol (ERP)", RFC 5296, August 2008.

[RFC5191] Forsberg, D., Ohba, Y., Patil, B., Tschofenig, H., and A. Yegin, "Protocol for Carrying Authentication for Network Access (PANA)", RFC 5191, May 2008.

Appendix A. Attacks Prevented by Channel Bindings

In the following it is demonstrated how the presented channel bindings can prevent attacks by malicious authenticators (representing the lying NAS problem) as well as malicious visited networks (representing the lying provider problem). This document only provides part of the solution necessary to realize a defense against these attacks. In addition, lower-layer protocols need to describe what attributes should be included in channel binding requests. EAP methods need to be updated in order to describe how the channel binding request and response are carried. In addition, deployments may need to decide what information is populated in the local database. The following sections describe types of attacks that can be prevented by this framework with appropriate lower-layer attributes carried in channel bindings, EAP methods with channel binding support and appropriate local database information at the EAP server.

A.1. Enterprise Subnetwork Masquerading

As outlined in Section 3, an enterprise network may have multiple VLANs providing different levels of security. In an attack, a malicious NAS connecting to a guest network with lesser security protection could broadcast the SSID of a subnetwork with higher protection. This could lead peers to believe that they are accessing the network over secure connections, and, e.g., transmit confidential information that they normally would not send over a weakly protected connection. This attack works under the conditions that peers use the same set of credentials to authenticate to the different kinds of VLANs and that the VLANs support at least one common EAP method. If these conditions are not met, the EAP server would not authorize the peers to connect to the guest network, because the peers used credentials and/or an EAP method that is associated with the corporate network.

A.2. Forced Roaming

Mobile phone providers boosting their cell tower's transmission power to get more users to use their networks have occurred in the past. The increased transmission range combined with a NAS sending a false network identity lures users to connect to the network without being aware of that they are roaming.

Channel bindings would detect the bogus network identifier because the network identifier sent to the authentication server in `il` will neither match information `i2` nor the stored data. The verification fails because the info in `il` claims to come from the peer's home network while the home authentication server knows that the connection is through a visited network outside the home domain. In the same context, channel bindings can be utilized to provide a "home zone" feature that notifies users every time they are about to connect to a NAS outside their home domain.

A.3. Downgrading attacks

A malicious authenticator could modify the set of offered EAP methods in its Beacon to force the peer to choose from only the weakest EAP method(s) accepted by the authentication server. For instance, instead of having a choice between EAP-MD5-CHAP, EAP-FAST and some other methods, the authenticator reduces the choice for the peer to the weaker EAP-MD5-CHAP method. Assuming that weak EAP methods are supported by the authentication server, such a downgrading attack can enable the authenticator to attack the integrity and confidentiality of the remaining EAP execution and/or break the authentication and key exchange. The presented channel bindings prevent such downgrading attacks, because peers submit the offered EAP method

selection that they have received in the beacon as part of il to the authentication server. As a result, the authentication server recognizes the modification when comparing the information to the respective information in its policy database. This presumes that all acceptable EAP methods support channel binding and that an attacker cannot break the EAP method in real-time.

A.4. Bogus Beacons in IEEE 802.11r

In IEEE 802.11r, the SSID is bound to the TSK calculations, so that the TSK needs to be consistent with the SSID advertised in an authenticator's Beacon. While this prevents outsiders from spoofing a Beacon it does not stop a "lying NAS" from sending a bogus Beacon and calculating the TSK accordingly.

By implementing channel bindings, as described in this draft, in IEEE 802.11r, the verification by the authentication server would detect the inconsistencies between the information the authenticator has sent to the peer and the information the server received from the authenticator and stores in the policy database.

A.5. Forcing false authorization in IEEE 802.11i

In IEEE 802.11i a malicious NAS can modify the beacon to make the peer believe it is connected to a network different from the one the peer is actually connected to.

In addition, a malicious NAS can force an authentication server into authorizing access by sending an incorrect Called-Station-ID that belongs to an authorized NAS in the network. This could cause the authentication server to believe it had granted access to a different network or even provider than the one the peer got access to.

Both attacks can be prevented by implementing channel bindings, because the server can compare the information that was sent to the peer, with information it received from the authenticator during the AAA communication as well as the information stored in the policy database.

Appendix B. Change History

RFC editor, remove this section prior to publication.

B.1. Changes Since 09

Based on WG discussion, all assigned numbers start at 1, including NSIDs and codes.

Based on WG discussion we include the value of attributes in the RADIUS namespace in channel binding responses.

B.2. Changes since Version 06

The purpose of this revision is to provide a specific candidate protocol for channel binding data and channel binding responses.

B.3. Changes since version 04

- o Clarify examples in introduction.
- o In problem statement note that one EAP server may deal with both enterprise and provider networks.
- o Update discussion of the architecture. Talk about channel bindings as a mechanism to introduce levels of trust.
- o Indicate that this document is focusing on EAP channel bindings within methods while trying to do a better job of describing the SAP approach in more detail.
- o Claim that we're using the encoding from draft-clancy-emu-aaapay. The WG almost certainly doesn't have consensus on this, but in the interest of actually describing what the protocol might be like, it is a good straw-man proposal.
- o Update protocol description.

Authors' Addresses

Sam Hartman (editor)
Painless Security
356 Abbott ST
North Andover, MA 01845
USA

Email: hartmans-ietf@mit.edu

T. Charles Clancy
Department of Electrical Engineering and Computer Science
Virginia Tech
Arlington, VA 22203
USA

Email: tcc@vt.edu

Katrin Hoeper
Motorola, Inc.
1301 E. Algonquin Road
Schaumburg, IL 60196
USA

Email: khoeper@motorolasolutions.com

Network Working Group
Internet-Draft
Intended status: Informational
Expires: December 30, 2012

S. Hartman
M. Wasserman
Painless Security
D. Zhang
Huawei
June 28, 2012

EAP Mutual Cryptographic Binding
draft-ietf-emu-crypto-bind-00.txt

Abstract

As the Extensible Authentication Protocol (EAP) evolves, EAP peers rely increasingly on information received from the EAP server. EAP extensions such as channel binding or network posture information are often carried in tunnel methods; peers are likely to rely on this information. [RFC 3748] is a facility that protects tunnel methods against man-in-the-middle attacks. However, cryptographic binding focuses on protecting the server rather than the peer. This memo explores attacks possible when the peer is not protected from man-in-the-middle attacks and recommends mutual cryptographic binding, a new form of cryptographic binding that protects both peer and server along with other mitigations.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on December 30, 2012.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal

Provisions Relating to IETF Documents
(<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents


1. Introduction	3
2. An Example Problem	5
3. The Server insertion Attack	7
3.1. Conditions for the Attack	7
3.2. Mitigation Strategies	8
3.2.1. Server Authentication	8
3.2.2. Server Policy	9
3.2.3. Existing Cryptographic Binding	12
3.2.4. Introducing EMSK-based Cryptographic Binding	13
3.2.5. Mix Key into Long-Term Credentials	14
3.3. Intended Intermediates	14
4. Recommendations	16
4.1. Mutual Cryptographic Binding	16
4.2. State Tracking	16
4.3. Certificate Naming	16
4.4. Inner Mixing	17
5. Survey of Tunnel Methods	18
6. Survey of Inner Methods	19
7. Security Considerations	20
8. Acknowledgements	21
9. References	22
9.1. Normative References	22
9.2. Informative References	22
Authors' Addresses	24

1. Introduction

The Extensible Authentication Protocol [RFC3748] provides authentication between a peer (a party accessing some service) and a authentication server. Traditionally, peers have not relied significantly on information received from EAP servers. However facilities such as EAP Channel Binding [I-D.ietf-emu-chbind] provide the peer with confirmation of information about the resource it is accessing. Other facilities such as EAP Posture Transport [I-D.ietf-nea-pt-eap] permit a peer and EAP server to discuss the security properties of accessed networks. Both of these facilities provide peers with information they need to rely on and provide attackers who are able to impersonate an EAP server to a peer with new opportunities for attack.

Instead of adding these new facilities to all EAP methods, work has focused on adding support to tunnel methods [I-D.ietf-emu-eaptunnel-req]. There are numerous tunnel methods including [RFC4851], [RFC5281], and work on building a standards track tunnel method [I-D.ietf-emu-eap-tunnel-method]. These tunnel methods are extensible. By adding an extension to support a facility such as channel binding to a tunnel method, it can be used with any inner method carried in the tunnel.

Tunnel methods need to be careful about man-in-the-middle attacks. See section 3.2 and 4.6.3 in [I-D.ietf-emu-eaptunnel-req] and [TUNNEL-MITM] for a detailed description of these attacks. An example of the attack can happen when a peer is willing to perform authentication inside and outside a tunnel. An attacker can impersonate the EAP server and offer the inner method to the peer. However, on the other side, the attacker acts as a man-in-the-middle and opens a tunnel to the real EAP server. Figure 1 illustrates this attack. At the end of the attack, the EAP server believes it is talking to the peer. At the inner method level, this is true. At the outer method level, however, the server is talking to the attacker.



A classic tunnel attack where the attacker inserts an extra tunnel between the attacker and EAP server.

Figure 1: Classic Tunnel Attack


There are several mitigation strategies for this classic attack. First, security policy can be set up so that the same method is not offered by a server both inside and outside a tunnel. A technical solution is available if the inner method is sufficiently strong: cryptographic binding is a security property of a tunnel method under which the EAP server confirms that the inner and outer parties are the same. One common way to do this is to ask the outer party (the other end of the tunnel) to prove knowledge of the Master Session Key (MSK) of the inner method. As defined in RFC 3748, cryptographic binding may prove to the peer that the inner and outer exchanges are with the same party, but it typically does not make this proof; instead it is typically limited to proving to the server that the inner and outer peer are the same.

2. An Example Problem

The GSS-EAP mechanism [I-D.ietf-abfab-gss-eap] provides application authentication using EAP. A peer could reasonably trust some applications significantly more than others. If the peer sends confidential information to some applications, an attacker may gain significant value from convincing the peer that the attacker is the trusted application. Channel bindings are used to tell the peer which application service is being connected to. Prior to channel bindings, peers could not distinguish one Network Access Service (NAS) from another, so attacks where one NAS impersonated another were out-of-scope. However channel bindings add this capability and thus expands the threat model of EAP. The GSS-EAP mechanism requires distinguishing one service from another.

A relatively untrusted service, say a print server, has been compromised. A user is attempting to connect to a trusted service such as a financial application. Both the print server and the financial application use an Authentication, Authorization and Accounting protocol (AAA) to transport EAP authentication back to the user's EAP server. The print server mounts a man-in-the-middle attack on the user's connection to the financial application and claims to be the application.

The print server offers a tunnel method towards the peer. The print server extracts the inner method from the tunnel and sends it on towards the AAA server. Channel binding happens at the tunnel method though. So, the print server is happy to confirm that it is the financial application. After the inner method completes, the EAP server sends the MSK to the print server over the AAA protocol. If only the MSK is needed for cryptographic binding then the print server can successfully perform cryptographic binding and may be able to impersonate the financial application to the peer.



A modified tunnel attack when an extra server rather than extra client is inserted.

Figure 2: Channel Binding Requires More than Crypto Binding

This attack is not specific to GSS-EAP. The channel bindings specification [I-D.ietf-emu-chbind] describes a number of situations where channel bindings are important for network access. In these situations one NAS could impersonate another by using a similar attack.

3. The Server insertion Attack

The previous section described an example of the server insertion attack. In this attack, one party adds a layer of tunneling such that from the perspective of the EAP peer, there are more methods than from the perspective of the EAP server. This attack is most beneficial when the party inserting the extra tunnel is a legitimate NAS, so mitigations need to be able to prevent a legitimate NAS from inappropriately adding a layer of tunneling. Some deployments utilize an intentional intermediary that adds an extra level of EAP tunneling between the peer and the EAP server; see Section 3.3 for a discussion.

3.1. Conditions for the Attack

For an inserted server attack to have value, the attacker needs to gain an advantage from its attack. An advantage to the attacker could come from:

- o The attacker can send information to a peer that the peer would trust from the EAP server but not the attacker. Examples of this include channel binding responses.
- o The peer sending information to the attacker that was intended for the EAP server. For example, the inner user identity may disclose privacy-sensitive information. The channel binding request may disclose what service the peer wishes to connect to.
- o The attacker may influence session parameters. For example, if the attacker can influence the MSK, then the attacker may be able to read or influence session traffic and mount an attack on the confidentiality or integrity of the resulting session.
- o An attacker may impact availability of the session. In practice though, an attacker that can mount a server insertion attack is likely to be able to impact availability in other ways.

For this attack to be possible, the following conditions need to hold:

1. The attacker needs to be able to establish a tunnel method with the peer over which the peer will authenticate.
2. The attacker needs to be able to respond to any inner authentication. For example an attacker who is a legitimate NAS can forward the inner authentication over AAA towards the EAP server. Note that the inner authentication may not be EAP.

3. Typically, the attacker needs to be able to complete the tunnel method after inner authentication. This may not be necessary if the attacker is gaining advantage from information sent by the peer over the tunnel.
4. In some cases the attacker may need to complete a Secure Association Protocol (SAP) or otherwise demonstrate knowledge of the MSK after the tunnel method successfully completes.

Attackers who are legitimate NASes are the primary focus of this memo. Previous work has provided mitigation against attackers who are not a NAS; these mitigations are briefly discussed.

3.2. Mitigation Strategies

3.2.1. Server Authentication

If the peer confirms the identity of the party that the tunnel method is established with, the peer prevents the first condition (attacker establishing a tunnel method). Many tunnel methods rely on TLS [RFC5281] [I-D.ietf-emu-eap-tunnel-method]. The specifications for these methods tend to encourage or mandate certificate checking. If the TLS certificate is validated back to a trust anchor and the identity of the tunnel method server confirmed, then the first attack condition cannot be met.

Many challenges make server authentication difficult. There is not an obvious name by which to identify a tunnel method server. It is not obvious where in the tunnel server certificate the name should be found. One particularly problematic practice is to use a certificate that names the host on which the tunnel server runs. Given such a name it is very difficult for a peer to understand whether that server is intended to be a tunnel method server for the realm.

It's not clear what trust anchors to use for tunnel servers. Using commercial Certificate Authorities (CAs) is probably undesirable because tunnel servers often operate in a closed community and are often provisioned with certificates issued by that community. Using commercial CAs can be particularly problematic with peers that support hostnames in certificates. Then anyone who can obtain a certificate for any host in the domain being contacted can impersonate a tunnel server.

These difficulties lead to poor deployment of good certificate validation. Many peers make it easy to disable certificate validation. Other peers validate back to trust anchors but do not check names of certificates. What name types are supported and what configuration is easy to perform depends significantly on the peer in

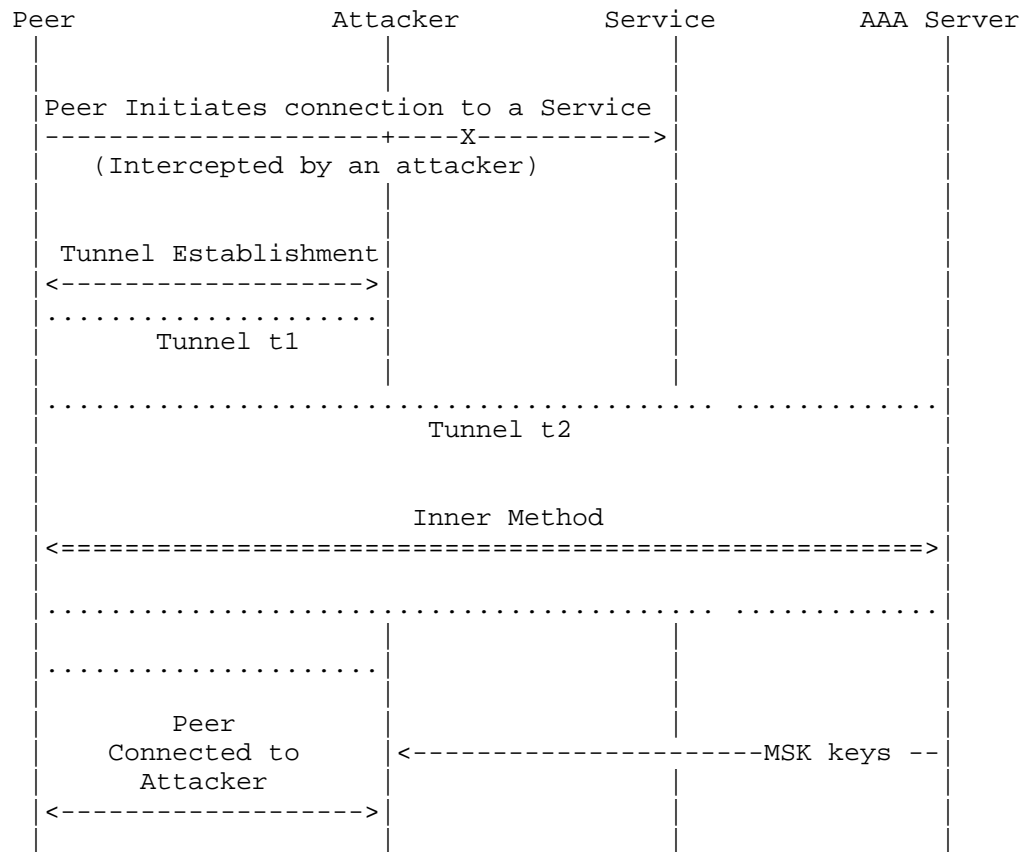
question.

Specifications also make the problem worse. For example [RFC5281] indicates that the only impact of failing to perform certificate validation is that the inner method can be attacked. Administrators and implementors believing this claim may believe that protection from passive attacks is sufficient.

In addition, some deployments such as provisioning or strong inner methods are designed to work without certificate validation. Section 3.9 of the tunnel requirements [I-D.ietf-emu-eaptunnel-req] discusses this requirement.

3.2.2. Server Policy

Server policy can potentially prevent the second condition (attacker being able to respond to inner authentication) from being possible. If the server only performs a particular inner authentication within a tunnel, then the attacker cannot gain a response to the inner authentication without their being such a tunnel. The attacker may be able to add a second layer of tunnels; see Figure 3. The inner tunnel may limit the attacker's capabilities; for example if channel binding is performed over tunnel t2 in the figure, then an attacker cannot observe or influence it.



A tunnel t1 from the peer to the attacker contains a tunnel t2 from the peer to the home EAP server. Inside t2 is an inner authentication.

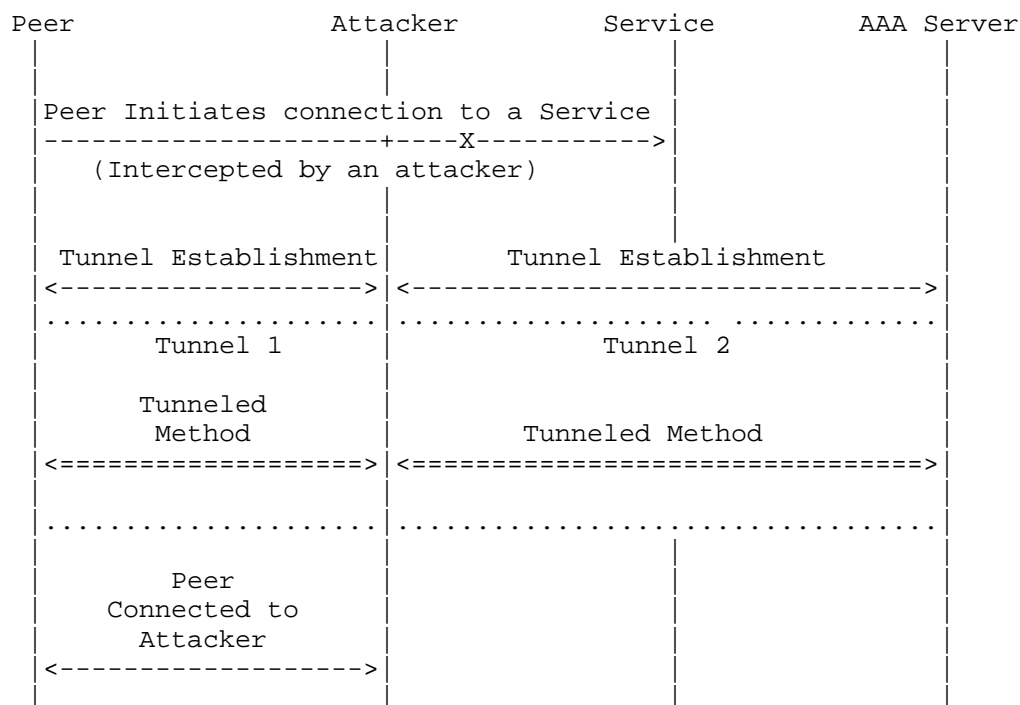
Figure 3: Multiple Layered Tunnels

Peer policy can be combined with this server policy to help prevent conditions 1 (attacker can establish a tunnel the peer will use) and 2 (attacker can respond to inner authentication). If the peer requires exactly one tunnel of a particular type and the EAP server only performs inner authentication over a tunnel of this type, then the attacker cannot establish tunnel t1 in the figure above.

An attacker may be able to mount a more traditional man-in-the-middle attack in this instance; see Figure 4. This policy on the peer and EAP server combined with a tunnel method that supports cryptographic binding will allow the EAP server to detect the attacker. This means

the attacker cannot act as a legitimate NAS and in particular does not obtain the MSK. So, if the tunnel between the attacker and peer also requires cryptographic binding and if the cryptographic binding requires both the EAP server and peer to prove knowledge of the inner MSK, then the authentication will fail. If cryptographic binding is not performed, then this attack may succeed.

Please view in a fixed-width font such as Courier.



A tunnel t1 extends from the peer to the attacker. a tunnel t2 extends from the attacker to the home EAP server. An inner EAP authentication is forwarded unmodified by the attacker from t1 to t2. The attacker can observe this inner authentication.

Figure 4: A Traditional Man-in-the-Middle Attack

Cryptographic binding is only a valuable component of a defense if the inner authentication is a key-deriving EAP method. Most tunnel methods also support non-EAP inner authentication such as Microsoft Chap version 2 [RFC2759]. This may undermine cryptographic binding in a number of ways. An attacker may be able to convert an EAP method into a compatible non-EAP form of the same credential to suppress cryptographic binding. In addition, an inner authentication

may be available through an entirely different means. For example, a Lightweight Directory Access Protocol [RFC4510] or other directory server may provide an attacker a way to get challenges and provide responses for an authentication mechanism entirely outside of the AAA/EAP context. An attacker with this capability may be able to get around server policy requiring an inner authentication be used only in a given type of tunnel.

An attacker can convert an inner authentication using an EAP method to a inner authentication that does not use EAP in some cases. This may avoid cryptographic binding.

Converting EAP Inner Authentication

An attacker may contact another authentication resource to gain a challenge useful for an inner authentication.

Non-EAP Sources of Inner Authentication

To Recap, the following policy conditions appear sufficient to prevent a server insertion attack:

1. Peer and EAP server require a particular inner EAP method used within a particular tunnel method
2. The inner EAP method's authentication is only available within the tunnel and through no other means including non-EAP means
3. The inner EAP method produces a key
4. The tunnel method uses cryptographic binding and the peer requires the other end of the tunnel to prove knowledge of the inner MSK.

3.2.3. Existing Cryptographic Binding

The most advanced examples of cryptographic binding today work at two levels. First, the server and peer prove to each other knowledge of the inner MSK. Then, the inner MSK is combined into some outer key material to form the tunnel's keys. This is sufficient to detect an inserted server or peer provided that the attacker does not learn the inner MSK. This seems sufficient to defend against attackers who cannot act as a legitimate NAS.

The definition of cryptographic binding in RFC 3748 does not require these steps. To meet that definition it would be sufficient for a peer to prove knowledge of the inner key to the EAP server. This would open some additional attacks. For example by indicating success an attacker might be able to mask a cryptographic binding

failure. Especially if only the tunnel key material is used for the final keys, the peer is unlikely to be able to detect the failure.

As discussed in the previous section, cryptographic binding is only effective when the inner method is EAP.

3.2.4. Introducing EMSK-based Cryptographic Binding

Cryptographic binding can be strengthened when the inner EAP method supports an Extended Master Session Key (EMSK). The EMSK is never disclosed to any party other than the EAP server or peer, so even a legitimate NAS cannot learn the EMSK. So, if the same techniques currently applied to the inner MSK are applied to the inner EMSK, then condition 3 (completing tunnel authentication) will not hold because the attacker cannot complete this new form of cryptographic binding. This does not prevent the attacker from learning confidential information such as a channel binding request sent over the tunnel prior to cryptographic binding.

Obviously as with all forms of cryptographic binding, cryptographic binding only works for key-deriving inner EAP methods. Also, some deployments (see Section 3.3) insert intermediates between the peer and the EAP server. EMSK-based cryptographic binding is incompatible with these deployments because the intermediate cannot learn the EMSK.

Formally, EMSK-based cryptographic binding is a security claim for EAP tunnel methods that holds when:

1. The peer proves to the server that the peer participating in any inner method is the same as the peer for the tunnel method.
2. The server proves to the peer that the server for any inner method is the same as the server for the tunnel method.
3. The MSK and EMSK for the tunnel depend on the MSK and EMSK of inner methods.
4. The peer MUST be able to force the authentication to fail if the peer is unable to confirm the identity of the server.
5. Proofs offered need to be secure even against attackers who know the inner method MSK.

If EMSK-based cryptographic binding is not an optional facility it provides a strong defense against server insertion attacks and other tunnel MITM attacks for inner methods that provide an EMSK. The strength of the defense is dependent on the strength of the inner

method. EMSK-Based cryptographic binding MAY be provided as an optional facility. The value of EMSK-based cryptographic binding is reduced somewhat if it is an optional feature. It permits configurations where a peer uses other means to authenticate the server if the peer has sufficient information configured to validate the certificate and identity of an EAP server while using EMSK-based cryptographic binding for deployments where that is possible.

If EMSK-based cryptographic binding is an optional facility, the negotiation of whether to use it MUST be protected by the inner MSK or EMSK. Typically the MSK will be used as the primary advantage of making EMSK-based cryptographic binding an optional facility is to permit intermediates who know only the MSK to decline to use EMSK-based cryptographic binding. The peer MUST have an opportunity to fail the authentication after the server declines to use EMSK-based cryptographic binding.

3.2.5. Mix Key into Long-Term Credentials

Another defense against tunnel MITM attacks potentially including server insertion attacks is to use a different credential for tunneled methods from other authentications. This may prevent the second condition (attacker being able to respond to inner authentication) from taking place. For example, if key material from the tunnel is mixed into a shared secret or password that is the basis of the inner authentication, then the second condition will not hold unless the attacker already knows this shared secret. The advantage of this approach is that it seems to be the only way to strengthen non-EAP inner authentications within a tunnel.

There are several disadvantages. Choosing a function to mix the tunnel key material into the inner authentication will be very dependent on the inner authentication. In addition, this appears to involve a layering violation. However, exploring the possibility of providing a solution like this seems important because it can function for inner authentications where no other approach will work.

3.3. Intended Intermediates

Some deployments introduce a tunnel server separate from the EAP server; see [RFC5281] for an example of this style of deployment. The only difference between such an intermediate and an attacker is that the intermediate provides some function valuable to the peer or EAP server and that the intermediate is trusted by the peer. If peers are configured with the necessary information to validate certificates of these intermediates and to confirm their identity, then tunnel MITM and inserted server attacks can be defended against. The intermediates need to be trusted with regard to channel binding

and other services that the peer depends on.

Support for trusted intermediates is not a requirement according to the tunnel method requirements.

It seems reasonable to treat trusted intermediates as a special case if they are supported and to focus on the security of the case where there are not intermediates in the tunnel as the common case.

4. Recommendations

4.1. Mutual Cryptographic Binding

The EAP Tunnel Method [I-D.ietf-emu-eap-tunnel-method] should gain support for EMSK-based cryptographic binding.

As channel binding support is added to existing EAP methods, EMSK-based cryptographic binding or some other form of cryptographic binding that protects against server insertion should also be added to these methods. Mutual cryptographic binding may also be valuable when other services are added to EAP methods that may require a peer trust an EAP server.

4.2. State Tracking

Today, mutual authentication in EAP is thought of as a security claim about a method. However, in practice it's an attribute of a particular exchange. Mutual authentication can be obtained via checking certificates, through mutual cryptographic binding, or in very controlled cases through carefully crafted peer and server policy combined with existing cryptographic binding. Using services like channel binding that involve the peer trusting the EAP server should require mutual authentication be present in the session.

to accomplish this, implementations including channel binding or other peer services MUST track whether mutual authentication has happened. They SHOULD default to not permitting these peer services unless mutual authentication has happened. They SHOULD support a configuration where the peer fails to authenticate unless mutual authentication takes place. Discussion of whether this configuration should be recommended as a default is required.

The EAP Tunnel Method should permit peers to force authentication failure if they are unable to perform mutual authentication. The protocol should permit this to be deferred until after mutual cryptographic binding is considered.

Services such as channel binding should be deferred until after cryptographic binding/mutual cryptographic binding.

4.3. Certificate Naming

Work is required to promote interoperable deployment of server certificate validation by peers. A standard way to name EAP servers is required. Recommendations for what name forms peers should implement is required.

4.4. Inner Mixing

More consideration of the proposal to mix some key material into inner authentications is desired. As stated today, the proposal is under-defined and fairly invasive. Are there versions of this proposal that would be valuable? Is there a way to view it as something more abstract so that it does not involve tunnel and inner method specific combinatorial explosion?

5. Survey of Tunnel Methods

6. Survey of Inner Methods

7. Security Considerations

8. Acknowledgements

The authors would like to thank Alan DeKok for helping to explore these attacks. Alan focused the discussion on the importance of inner authentications that are not EAP and proposed mixing in key material as a way to resolve these authentications.

Jari Arkko provided a review of the attack and valuable context on past efforts in developing cryptographic binding.

Sam Hartman's and margaret Wasserman's work on this memo is funded by Huawei.

9. References

9.1. Normative References

- [RFC3748] Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H. Levkowetz, "Extensible Authentication Protocol (EAP)", RFC 3748, June 2004.

9.2. Informative References

- [I-D.ietf-abfab-gss-eap]
Hartman, S. and J. Howlett, "A GSS-API Mechanism for the Extensible Authentication Protocol", draft-ietf-abfab-gss-eap-08 (work in progress), June 2012.
- [I-D.ietf-emu-chbind]
Hartman, S., Clancy, T., and K. Hoeper, "Channel Binding Support for EAP Methods", draft-ietf-emu-chbind-16 (work in progress), May 2012.
- [I-D.ietf-emu-eap-tunnel-method]
Zhou, H., Cam-Winget, N., Salowey, J., and S. Hanna, "Tunnel EAP Method (TEAP) Version 1", draft-ietf-emu-eap-tunnel-method-03 (work in progress), June 2012.
- [I-D.ietf-emu-eaptunnel-req]
Zhou, H., Salowey, J., Hoeper, K., and S. Hanna, "Requirements for a Tunnel Based EAP Method", draft-ietf-emu-eaptunnel-req-09 (work in progress), December 2010.
- [I-D.ietf-nea-pt-eap]
Cam-Winget, N. and P. Sangster, "PT-EAP: Posture Transport (PT) Protocol For EAP Tunnel Methods", draft-ietf-nea-pt-eap-02 (work in progress), May 2012.
- [RFC2759] Zorn, G., "Microsoft PPP CHAP Extensions, Version 2", RFC 2759, January 2000.
- [RFC4510] Zeilenga, K., "Lightweight Directory Access Protocol (LDAP): Technical Specification Road Map", RFC 4510, June 2006.
- [RFC4851] Cam-Winget, N., McGrew, D., Salowey, J., and H. Zhou, "The Flexible Authentication via Secure Tunneling Extensible Authentication Protocol Method (EAP-FAST)", RFC 4851, May 2007.

[RFC5281] Funk, P. and S. Blake-Wilson, "Extensible Authentication Protocol Tunneled Transport Layer Security Authenticated Protocol Version 0 (EAP-TTLSv0)", RFC 5281, August 2008.

[TUNNEL-MITM]
 "".

Authors' Addresses

Sam Hartman
Painless Security

Email: hartmans-ietf@mit.edu

Margaret Wasserman
Painless Security

Email: mrw@painless-security.com
URI: <http://www.painless-security.com/>

Dacheng Zhang
Huawei

Email: zhangdacheng@huawei.com

EMU Working Group
Internet-Draft
Intended status: Standards Track
Expires: December 23, 2012

H. Zhou
N. Cam-Winget
J. Salowey
Cisco Systems
S. Hanna
Juniper Networks
June 21, 2012

Tunnel EAP Method (TEAP) Version 1
draft-ietf-emu-eap-tunnel-method-03.txt

Abstract

This document defines the Tunnel Extensible Authentication Protocol (TEAP) version 1. TEAP is a tunnel based EAP method that enables secure communication between a peer and a server by using the Transport Layer Security (TLS) to establish a mutually authenticated tunnel. Within the tunnel, Type-Length-Value (TLV) objects are used to convey authentication related data between the EAP peer and the EAP server.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on December 23, 2012.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents

carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	5
1.1. Specification Requirements	5
1.2. Design Goals	5
1.3. Terminology	7
2. Protocol Overview	8
2.1. Architectural Model	8
2.2. Protocol Layering Model	9
3. TEAP Protocol	10
3.1. Version Negotiation	10
3.2. TEAP Authentication Phase 1: Tunnel Establishment	11
3.2.1. TLS Session Resume Using Server State	12
3.2.2. TLS Session Resume Using a PAC	13
3.2.3. Transition between Abbreviated and Full TLS Handshake	14
3.3. TEAP Authentication Phase 2: Tunneled Authentication	15
3.3.1. EAP Sequences	15
3.3.2. Optional Password Authentication	16
3.3.3. Protected Termination and Acknowledged Result Indication	16
3.4. Determining Peer-Id and Server-Id	17
3.5. TEAP Session Identifier	18
3.6. Error Handling	18
3.6.1. TLS Layer Errors	19
3.6.2. Phase 2 Errors	19
3.7. Fragmentation	20
3.8. PAC Provisioning	21
3.9. Certificate Provisioning Within the Tunnel	21
3.10. Server Unauthenticated Provisioning Mode	22
4. Message Formats	22
4.1. TEAP Message Format	22
4.2. TEAP TLV Format and Support	25
4.2.1. General TLV Format	26
4.2.2. Authority-ID TLV	28
4.2.3. Identity-Type TLV	28
4.2.4. Result TLV	30
4.2.5. NAK TLV	31
4.2.6. Error TLV	32
4.2.7. Channel-Binding TLV	34
4.2.8. Vendor-Specific TLV	35
4.2.9. Request-Action TLV	36

4.2.10. EAP-Payload TLV	37
4.2.11. Intermediate-Result TLV	39
4.2.12. PAC TLV Format	40
4.2.12.1. Formats for PAC Attributes	41
4.2.12.2. PAC-Key	42
4.2.12.3. PAC-Opaque	42
4.2.12.4. PAC-Info	43
4.2.12.5. PAC-Acknowledgement TLV	45
4.2.12.6. PAC-Type TLV	46
4.2.13. Crypto-Binding TLV	47
4.2.14. Basic-Password-Auth-Req TLV	49
4.2.15. Basic-Password-Auth-Resp TLV	50
4.2.16. PKCS#7 TLV	52
4.2.17. PKCS#10 TLV	53
4.2.18. Trusted-Server-Root TLV	54
4.3. Table of TLVs	55
5. Cryptographic Calculations	56
5.1. TEAP Authentication Phase 1: Key Derivations	56
5.2. Intermediate Compound Key Derivations	57
5.3. Computing the Compound MAC	58
5.4. EAP Master Session Key Generation	59
6. IANA Considerations	59
7. Security Considerations	62
7.1. Mutual Authentication and Integrity Protection	63
7.2. Method Negotiation	63
7.3. Separation of Phase 1 and Phase 2 Servers	63
7.4. Mitigation of Known Vulnerabilities and Protocol Deficiencies	64
7.4.1. User Identity Protection and Verification	65
7.4.2. Dictionary Attack Resistance	66
7.4.3. Protection against Man-in-the-Middle Attacks	66
7.4.4. PAC Binding to User Identity	67
7.5. Protecting against Forged Clear Text EAP Packets	67
7.6. Server Certificate Validation	67
7.7. Tunnel PAC Considerations	68
7.8. Security Claims	68
8. Acknowledgements	70
9. References	70
9.1. Normative References	70
9.2. Informative References	72
Appendix A. Evaluation Against Tunnel Based EAP Method Requirements	75
A.1. Requirement 4.1.1 RFC Compliance	75
A.2. Requirement 4.2.1 TLS Requirements	75
A.3. Requirement 4.2.1.1.1 Cipher Suite Negotiation	75
A.4. Requirement 4.2.1.1.2 Tunnel Data Protection Algorithms	75
A.5. Requirement 4.2.1.1.3 Tunnel Authentication and Key Establishment	76

A.6.	Requirement 4.2.1.2 Tunnel Replay Protection	76
A.7.	Requirement 4.2.1.3 TLS Extensions	76
A.8.	Requirement 4.2.1.4 Peer Identity Privacy	76
A.9.	Requirement 4.2.1.5 Session Resumption	76
A.10.	Requirement 4.2.2 Fragmentation	76
A.11.	Requirement 4.2.3 Protection of Data External to Tunnel .	76
A.12.	Requirement 4.3.1 Extensible Attribute Types	77
A.13.	Requirement 4.3.2 Request/Challenge Response Operation .	77
A.14.	Requirement 4.3.3 Indicating Criticality of Attributes .	77
A.15.	Requirement 4.3.4 Vendor Specific Support	77
A.16.	Requirement 4.3.5 Result Indication	77
A.17.	Requirement 4.3.6 Internationalization of Display Strings	77
A.18.	Requirement 4.4 EAP Channel Binding Requirements	77
A.19.	Requirement 4.5.1.1 Confidentiality and Integrity	77
A.20.	Requirement 4.5.1.2 Authentication of Server	78
A.21.	Requirement 4.5.1.3 Server Certificate Revocation Checking	78
A.22.	Requirement 4.5.2 Internationalization	78
A.23.	Requirement 4.5.3 Meta-data	78
A.24.	Requirement 4.5.4 Password Change	78
A.25.	Requirement 4.6.1 Method Negotiation	78
A.26.	Requirement 4.6.2 Chained Methods	78
A.27.	Requirement 4.6.3 Cryptographic Binding with the TLS Tunnel	78
A.28.	Requirement 4.6.4 Peer Initiated	79
A.29.	Requirement 4.6.5 Method Meta-data	79
Appendix B.	Major Differences from EAP-FAST	79
Appendix C.	Examples	79
C.1.	Successful Authentication	79
C.2.	Failed Authentication	81
C.3.	Full TLS Handshake using Certificate-based Cipher Suite .	83
C.4.	Client authentication during Phase 1 with identity privacy	84
C.5.	Fragmentation and Reassembly	86
C.6.	Sequence of EAP Methods	88
C.7.	Failed Crypto-binding	90
C.8.	Sequence of EAP Method with Vendor-Specific TLV Exchange	91
C.9.	Peer Requests Inner Method After Server Sends Result TLV	93
C.10.	Channel Binding	95
Appendix D.	Major Differences from Previous Revisions	96
D.1.	Changes from -02	96
D.2.	Changes from -01	97
D.3.	Changes from -00	98

1. Introduction

An Extensible Authentication Protocol (EAP) tunnel method is an EAP method that establishes a secure tunnel and executes other EAP methods under the protection of that secure tunnel. An EAP tunnel method can be used in any lower layer protocol that supports EAP authentication. There are several existing EAP tunnel methods that use Transport Layer Security (TLS) [RFC5246] to establish the secure tunnel. EAP methods supporting this include Protected EAP (PEAP) [PEAP], Tunneled Transport Layer Security EAP (TTLS) [RFC5281] and EAP Flexible Authentication via Secure Tunneling (EAP-FAST) [RFC4851]. However, they all are either vendor specific or informational and industry calls for a standard-track tunnel EAP method. [I-D.ietf-emu-eaptunnel-req] outlines the list of requirements for a standard tunnel based EAP method.

Since the introduction of EAP-FAST [RFC4851] a few years ago, it has been widely adopted in variety of devices and platforms due to its strong security, flexibility and ease of deployment. It has been adopted by EMU working group as the basis for the standard tunnel based EAP method. This document describes Tunnel Extensible Authentication Protocol (TEAP) version 1, based on EAP-FAST [RFC4851] with some minor changes, to meet the requirements outlined in [I-D.ietf-emu-eaptunnel-req] for a standard tunnel based EAP method.

1.1. Specification Requirements

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] .

1.2. Design Goals

Network access solutions requiring user friendly and easily deployable secure authentication mechanisms highlight the need for strong mutual authentication protocols that enable the use of weaker user credentials. This document defines an Extensible Authentication Protocol (EAP) which consists of establishing a Transport Layer Security (TLS) tunnel using TLS version 1.2 [RFC5246] or a successor version supported by both parties. Once the tunnel is established, the protocol further exchanges data in the form of Type-Length-Value (TLV) objects to perform further authentication. TEAP supports the TLS extension defined in [RFC5077] to support fast re-establishment of the secure tunnel without having to maintain per-session state on the server.

TEAP's design motivations included:

- o Mutual authentication: an EAP server must be able to verify the identity and authenticity of the peer, and the peer must be able to verify the authenticity of the EAP server.
- o Immunity to passive dictionary attacks: many authentication protocols require a password to be explicitly provided (either as cleartext or hashed) by the peer to the EAP server; at minimum, the communication of the weak credential (e.g., password) must be immune from eavesdropping.
- o Immunity to man-in-the-middle (MitM) attacks: in establishing a mutually authenticated protected tunnel, the protocol must prevent adversaries from successfully interjecting information into the conversation between the peer and the EAP server.
- o Flexibility to enable support for most password authentication interfaces: as many different password interfaces (e.g., Microsoft Challenge Handshake Authentication Protocol (MS-CHAP), Lightweight Directory Access Protocol (LDAP), One-Time Password (OTP), etc.) exist to authenticate a peer, the protocol must provide this support for legacy password authentication seamlessly.
- o Cryptographic algorithm agility: a cryptographic algorithm's strength is not perpetual, as weaknesses in an algorithm are discovered or increased processing power overtakes an algorithm over time. Hence, the protocol must not be tied to any single cryptographic algorithm. Instead, it MUST support run-time negotiation to select among an extensible set of cryptographic algorithms and also allow users to choose the algorithm that best meets their needs.
- o Sequence of chained EAP methods: Several circumstances are best addressed by using chained EAP methods. For example, it may be desirable to authenticate the user and also authenticate the device being used. The protocol must support chained EAP methods while including protection against attacks on method chaining.

With these motivational goals defined, further secondary design criteria are imposed:

- o Flexibility to extend the communications inside the tunnel: with the growing complexity in network infrastructures, the need to gain authentication, authorization, and accounting is also evolving. For instance, there may be instances in which multiple existing authentication protocols are required to achieve mutual authentication. Similarly, different protected conversations may be required to achieve the proper authorization once a peer has successfully authenticated.

- o Minimize the authentication server's per user authentication state requirements: with large deployments, it is typical to have servers authenticating many peers. With many different authentication servers deployed, a peer's session state may need to be replicated to allow for high availability or mobility scenarios. To facilitate scalable authentication server deployments and more efficient per user state management, it is desirable for a peer to cache its session state that has been securely encapsulated by the authentication server infrastructure.
- o Efficiency: specifically when using wireless media, peers will be limited in computational and power resources. The protocol must enable the network access communication to be computationally lightweight.
- o Channel bindings: EAP channel bindings seek to authenticate previously unauthenticated information provided by the authenticator to the EAP peer, by allowing the peer and server to compare their perception of network properties in a secure channel. It is used to solve the lying NAS and the lying provider problems. The protocol should provide support for EAP channel bindings as defined in [I-D.ietf-emu-chbind].

1.3. Terminology

Much of the terminology in this document comes from [RFC3748]. Additional terms are defined below:

Protected Access Credential (PAC)

Credentials distributed to a peer for future optimized network authentication. The PAC consists of a minimum of two components: a shared secret and an opaque element. The shared secret component contains the pre-shared key between the peer and the authentication server. The opaque part is provided to the peer and is presented to the authentication server when the peer wishes to obtain access to network resources. The opaque element and shared secret are used with TLS stateless session resumption defined in RFC 5077 [RFC5077] to establish a protected TLS session. The secret key and opaque part may distributed using RFC 5077 messages or using TLVs within the TEAP tunnel. Finally, a PAC may optionally include other information that may be useful to the peer.

Type-Length-Value (TLV)

The TEAP protocol utilizes objects in Type Length Value (TLV) format. The TLV format is defined in Section 4.2.

2. Protocol Overview


TEAP authentication occurs in two phases. In the first phase, TEAP employs the TLS [RFC5246] handshake to provide an authenticated key exchange and to establish a protected tunnel. Once the tunnel is established, the second phase begins with the peer and server engaging in further conversations to establish the required authentication and authorization policies. TEAP makes use of Type-Length-Value objects (TLVs) to carry out the inner authentication, results and other information, such as channel binding information.

TEAP makes use of the TLS enhancements in Ticket Extension [RFC5077] to enable an optimized TLS tunnel session resume while minimizing server state. The ticket is referred to as the Protected Access Credential opaque data (or PAC-Opaque). The PAC-Opaque may be distributed through the use of the NewSessionTicket message or through a mechanism that uses TLVs within phase 2 of TEAP. The secret key used to resume the session in TEAP is referred to as the Protected Access Credential key (or PAC-Key). When the NewSessionTicket message is being used to distribute the PAC-Opaque, the PAC-Key is the Master Secret for the session. If TEAP phase 2 is used to distribute the PAC-Opaque, then the PAC-Key is distributed along with the PAC-Opaque. TEAP implementations MUST support the RFC 5077 mechanism for distributing a PAC-Opaque and it is RECOMMENDED that implementations support the capability to distribute the ticket and secret key within the TEAP tunnel.

The TEAP conversation is used to establish or resume an existing session to typically establish network connectivity between a peer and the network. Upon successful execution of TEAP, both EAP peer and EAP server derive strong session key material that can then be communicated to the network access server (NAS) for use in establishing a link layer security association.

2.1. Architectural Model

The network architectural model for TEAP usage is shown below:




TEAP Architectural Model

The entities depicted above are logical entities and may or may not correspond to separate network components. For example, the TEAP server and inner method server might be a single entity; or the authenticator and TEAP server might be a single entity; or the functions of the authenticator, TEAP server, and inner method server might be combined into a single physical device. For example, typical IEEE 802.11 deployments place the Authenticator in an access point (AP) while a Radius server may provide the TEAP and inner method server components. The above diagram illustrates the division of labor among entities in a general manner and shows how a distributed system might be constructed; however, actual systems might be realized more simply. The security considerations Section 7.3 provides an additional discussion of the implications of separating the TEAP server from the inner method server.

2.2. Protocol Layering Model

TEAP packets are encapsulated within EAP; EAP in turn requires a carrier protocol for transport. TEAP packets encapsulate TLS, which is then used to encapsulate user authentication information. Thus, TEAP messaging can be described using a layered model, where each layer encapsulates the layer above it. The following diagram clarifies the relationship between protocols:



Protocol Layering Model

The TLV layer is a payload with Type-Length-Value (TLV) Objects defined in Section 4.2. The TLV objects are used to carry arbitrary parameters between an EAP peer and an EAP server. All conversations in the TEAP protected tunnel must be encapsulated in a TLV layer.

TEAP packets may include TLVs both inside and outside the TLS tunnel.

The term "Outer TLVs" is used to refer to optional TLVs outside the TLS tunnel, which are only allowed in the first two messages in the TEAP protocol. That is the first EAP server to peer message and first peer to EAP server message. If the message is fragmented, the whole set of messages is counted as one message. The term "Inner TLVs" is used to refer to TLVs sent within the TLS tunnel. In TEAP Phase 1, Outer TLVs are used to help establishing the TLS tunnel, but no Inner TLVs are used. In Phase 2 of the TEAP conversation, TLS records may encapsulate zero or more Inner TLVs, but no Outer TLVs.

Methods for encapsulating EAP within carrier protocols are already defined. For example, IEEE 802.1X [IEEE.802-1X.2004] may be used to transport EAP between the peer and the authenticator; RADIUS [RFC3579] or Diameter [RFC4072] may be used to transport EAP between the authenticator and the EAP server.

3. TEAP Protocol

TEAP authentication occurs in two phases. In the first phase, TEAP employs the TLS handshake to provide an authenticated key exchange and to establish a protected tunnel. Once the tunnel is established the second phase begins with the peer and server engaging in further conversations to establish the required authentication and authorization policies. The operation of the protocol, including Phase 1 and Phase 2, is the topic of this section. The format of TEAP messages is given in Section 4 and the cryptographic calculations are given in Section 5.

3.1. Version Negotiation

TEAP packets contain a 3-bit version field, following the TLS Flags field, which enables future TEAP implementations to be backward compatible with previous versions of the protocol. This specification documents the TEAP version 1 protocol; implementations of this specification MUST use a version field set to 1.

Version negotiation proceeds as follows:

In the first EAP-Request sent with EAP type=TEAP, the EAP server must set the version field to the highest supported version number.

If the EAP peer supports this version of the protocol, it MUST respond with an EAP-Response of EAP type=TEAP, and the version number proposed by the TEAP server.

If the TEAP peer does not support this version but supports the version that is lower than the version proposed by the TEAP server, it responds with an EAP-Response of EAP type=TEAP and the highest supported version number. If the TEAP peer only supports the version that is higher than the version proposed by the TEAP server, then use of TEAP will not be possible. In this case, the TEAP peer should send back an EAP-Nak with other proposed EAP method if available.

If the TEAP server does not support the version number proposed by the TEAP peer, it MAY terminate the conversation with EAP-Failure or negotiate for another EAP type. Otherwise the TEAP conversation continues.

The version negotiation procedure guarantees that the TEAP peer and server will agree to the latest version supported by both parties. If version negotiation fails, then use of TEAP will not be possible, and another mutually acceptable EAP method will need to be negotiated if authentication is to proceed.

The TEAP version is not protected by TLS; and hence can be modified in transit. In order to detect a modification of the TEAP version, the peers MUST exchange the TEAP version number received during version negotiation using the Crypto-Binding TLV described in Section 4.2.13. The receiver of the Crypto-Binding TLV MUST verify that the version received in the Crypto-Binding TLV matches the version sent by the receiver in the TEAP version negotiation.

3.2. TEAP Authentication Phase 1: Tunnel Establishment

TEAP is based on the TLS handshake [RFC5246] to establish an authenticated and protected tunnel. The TLS version offered by the peer and server MUST be TLS version 1.2 [RFC5246] or later. This version of the TEAP implementation MUST support the following TLS ciphersuites:

TLS_RSA_WITH_AES_128_CBC_SHA [RFC5246]

TLS_DHE_RSA_WITH_AES_128_CBC_SHA [RFC5246]

Other ciphersuites MAY be supported. It is REQUIRED that anonymous ciphersuites such as TLS_DH_anon_WITH_AES_128_CBC_SHA [RFC5246] only be used in the case when the inner authentication method provides mutual authentication, key generation, and resistance to man-in-the-middle and dictionary attack. During the TEAP Phase 1 conversation, the TEAP endpoints MAY negotiate TLS compression. During TLS tunnel establishment, TLS extensions MAY be used. For instance, Certificate

Status Request extension [RFC6066] can be used to leverage a certificate-status protocol such as OCSP [RFC2560] to check the validity of server certificates. TLS renegotiation indications defined in RFC 5746 [RFC5746] MUST be supported.

The EAP server initiates the TEAP conversation with an EAP request containing a TEAP/Start packet. This packet includes a set Start (S) bit, the TEAP version as specified in Section 3.1, and an authority identity TLV. The TLS payload in the initial packet is empty. The authority identity TLV (Authority-ID TLV) is used to provide the peer a hint of the server's identity that may be useful in helping the peer select the appropriate credential to use. Assuming that the peer supports TEAP, the conversation continues with the peer sending an EAP-Response packet with EAP type of TEAP with the Start (S) bit clear and the version as specified in Section 3.1. This message encapsulates one or more TLS records containing the TLS handshake messages. If the TEAP version negotiation is successful then the TEAP conversation continues until the EAP server and EAP peer are ready to enter Phase 2. When the full TLS handshake is performed, then the first payload of TEAP Phase 2 MAY be sent along with server-finished handshake message to reduce the number of round trips.

TEAP implementations MUST support client authentication during tunnel establishment using the TLS ciphersuites specified in Section 3.2. The EAP peer does not need to authenticate as part of the TLS exchange, but can alternatively be authenticated through additional exchanges carried out in Phase 2.

The TEAP tunnel protects peer identity information exchanged during phase 2 from disclosure outside the tunnel. Implementations that wish to provide identity privacy for the peer identity must carefully consider what information is disclosed outside the tunnel prior to phase 2. TEAP implementations SHOULD support the immediate renegotiation of a TLS session to initiate a new handshake message exchange under the protection of the current cipher suite. This allows support for protection of the peer's identity when using TLS client authentication.

The following sections describe resuming a TLS session based on server-side or client-side state.

3.2.1. TLS Session Resume Using Server State

TEAP session resumption is achieved in the same manner TLS achieves session resume. To support session resumption, the server and peer must minimally cache the Session ID, master secret, and ciphersuite. The peer attempts to resume a session by including a valid Session ID from a previous handshake in its ClientHello message. If the server

finds a match for the Session ID and is willing to establish a new connection using the specified session state, the server will respond with the same Session ID and proceed with the TEAP Phase 1 tunnel establishment based on a TLS abbreviated handshake. After a successful conclusion of the TEAP Phase 1 conversation, the conversation then continues on to Phase 2.

3.2.2. TLS Session Resume Using a PAC

TEAP supports the resumption of sessions based on server state being stored on the client side using the TLS SessionTicket extension techniques described in [RFC5077]. This version of TEAP supports the provisioning of a ticket called a Protected Access Credential (PAC) through the use of the NewSessionTicket handshake described in [RFC5077], as well as provisioning of a PAC inside the protected tunnel. Implementations may provide additional ways to provision the PAC, such as manual configuration. Since the PAC mentioned here is used for establishing the TLS Tunnel, it is more specifically referred to as the Tunnel PAC. The Tunnel PAC is a security credential provided by the EAP server to a peer and comprised of:

1. PAC-Key: this is the key used by the peer as the TLS master secret to establish the TEAP Phase 1 tunnel. The PAC-Key is a strong high-entropy at minimum 48-octet key and is typically the master secret from a previous TLS session. The PAC-Key is a secret and MUST be treated accordingly. In the case that a PAC-Key is provisioned to the client through another means it must have its confidentiality and integrity protected by a mechanism, such as the TEAP phase 2 tunnel. The PAC-Key must be stored securely by the peer.
2. PAC-Opaque: this is a variable length field containing the ticket that is sent to the EAP server during the TEAP Phase 1 tunnel establishment based on RFC 5077. The PAC-Opaque can only be interpreted by the EAP server to recover the required information for the server to validate the peer's identity and authentication. The PAC-Opaque includes the PAC-Key and other TLS session parameters. It may contain the PAC's peer identity. The PAC-Opaque format and contents are specific to the PAC issuing server. The PAC-Opaque may be presented in the clear, so an attacker MUST NOT be able to gain useful information from the PAC-Opaque itself. The server issuing the PAC-Opaque must ensure it is protected with strong cryptographic keys and algorithms. The PAC-Opaque may be distributed using the NewSessionTicket message defined in RFC 5077 or it may be distributed through another mechanism such as the phase 2 TLVs defined in this document.

3. PAC-Info: this is an optional variable length field used to provide, at a minimum, the authority identity of the PAC issuer. Other useful but not mandatory information, such as the PAC-Key lifetime, may also be conveyed by the PAC issuing server to the peer during PAC provisioning or refreshment. PAC-Info is not included if the NewSessionTicket message is used to provision the PAC.

The use of the PAC is based on the SessionTicket extension defined in [RFC5077]. The EAP server initiates the TEAP conversation as normal. Upon receiving the Authority-ID TLV from the server, the peer checks to see if it has an existing valid PAC-Key and PAC-Opaque for the server. If it does, then it obtains the PAC-Opaque and puts it in the SessionTicket extension in the ClientHello. It is RECOMMENDED in TEAP that the peer include an empty Session ID in a ClientHello containing a PAC-Opaque. This version of TEAP supports the NewSessionTicket Handshake message as described in [RFC5077] for distribution of a new PAC, as well as the provisioning of PAC inside the protected tunnel. If the PAC-Opaque included in the SessionTicket extension is valid and the EAP server permits the abbreviated TLS handshake, it will select the cipher suite from information within the PAC-Opaque and finish with the abbreviated TLS handshake. If the server receives a Session ID and a PAC-Opaque in the SessionTicket extension in a ClientHello, it should place the same Session ID in the ServerHello if it is resuming a session based on the PAC-Opaque. The conversation then proceeds as described in [RFC5077] until the handshake completes or a fatal error occurs. After the abbreviated handshake completes, the peer and the server are ready to commence Phase 2.

3.2.3. Transition between Abbreviated and Full TLS Handshake

If session resumption based on server-side or client-side state fails, the server can gracefully fall back to a full TLS handshake. If the ServerHello received by the peer contains an empty Session ID or a Session ID that is different than in the ClientHello, the server may fall back to a full handshake. The peer can distinguish the server's intent of negotiating full or abbreviated TLS handshake by checking the next TLS handshake messages in the server response to the ClientHello. If ChangeCipherSpec follows the ServerHello in response to the ClientHello, then the server has accepted the session resumption and intends to negotiate the abbreviated handshake. Otherwise, the server intends to negotiate the full TLS handshake. A peer can request for a new PAC to be provisioned after the full TLS handshake and mutual authentication of the peer and the server. In order to facilitate the fallback to a full handshake the peer SHOULD include cipher suites that allow for a full handshake and possibly PAC provisioning so the server can select one of these in case

session resumption fails. An example of the transition is shown in Appendix C.

3.3. TEAP Authentication Phase 2: Tunneled Authentication

The second portion of the TEAP Authentication occurs immediately after successful completion of Phase 1. Phase 2 occurs even if both peer and authenticator are authenticated in the Phase 1 TLS negotiation. Phase 2 **MUST NOT** occur if the Phase 1 TLS handshake fails. Phase 2 consists of a series of requests and responses encapsulated in TLV objects defined in Section 4.2. Phase 2 **MUST** always end with a crypto-binding TLV exchange described in Section 4.2.13 and a protected termination exchange described in Section 3.3.3. The TLV exchange may include the execution of zero or more EAP methods within the protected tunnel as described in Section 3.3.1. A server **MAY** proceed directly to the protected termination exchange if it does not wish to request further authentication from the peer. However, the peer and server must not assume that either will skip inner EAP methods or other TLV exchanges. The peer may have roamed to a network that requires conformance with a different authentication policy, or the peer may request the server take additional action (e.g., channel binding) through the use of the Request-Action TLV as defined in Section 4.2.9.

3.3.1. EAP Sequences

EAP [RFC3748] prohibits use of multiple authentication methods within a single EAP conversation in order to limit vulnerabilities to man-in-the-middle attacks. TEAP addresses man-in-the-middle attacks through support for cryptographic protection of the inner EAP exchange and cryptographic binding of the inner authentication method(s) to the protected tunnel. EAP methods are executed serially in a sequence. This version of TEAP does not support initiating multiple EAP methods simultaneously in parallel. The methods need not be distinct. For example, EAP-TLS could be run twice as an inner method, first using machine credentials followed by a second instance using user credentials.

EAP method messages are carried within EAP-Payload TLVs defined in Section 4.2.10. If more than one method is going to be executed in the tunnel, then upon method completion, the server **MUST** send an Intermediate-Result TLV indicating the result. The peer **MUST** respond to the Intermediate-Result TLV indicating its result. If the result indicates success, the Intermediate-Result TLV **MUST** be accompanied by a Crypto-Binding TLV. The Crypto-Binding TLV is further discussed in Section 4.2.13 and Section 5.3. The Intermediate-Result TLVs can be included with other TLVs such as EAP-Payload TLVs starting a new EAP

conversation or with the Result TLV used in the protected termination exchange.

If both peer and server indicate success, then the method is considered complete. If either indicates failure, then the method is considered failed. The result of failure of an EAP method does not always imply a failure of the overall authentication. If one authentication method fails, the server may attempt to authenticate the peer with a different method.

3.3.2. Optional Password Authentication

The use of EAP-FAST-GTC as defined in RFC 5421 [RFC5421] is not recommended with TEAPv1. Implementations should instead make use of the password authentication TLVs defined in this specification. The authentication server initiates password authentication by sending a Basic-Password-Auth-Req TLV defined in Section 4.2.14. If the peer wishes to participate in password authentication then it responds with a Basic-Password-Auth-Resp TLV as defined in Section 4.2.15 that contains the username and password. If it does not wish to perform password authentication then it responds with a NAK TLV indicating the rejection of the Basic-Password-Auth-Req TLV. Upon receiving the response, the server indicates the success or failure of the exchange using an Intermediate-Result TLV. Multiple roundtrips of password authentication requests and responses MAY be used to support some "housecleaning" functions such as password change, change pin, etc. before a user is authenticated.

3.3.3. Protected Termination and Acknowledged Result Indication

A successful TEAP Phase 2 conversation MUST always end in a successful Crypto-Binding TLV and Result TLV exchange. A TEAP server may initiate the Crypto-Binding TLV and Result TLV exchange without initiating any EAP conversation in TEAP Phase 2. After the final Result TLV exchange, the TLS tunnel is terminated and a clear text EAP-Success or EAP-Failure is sent by the server. Peers implementing TEAP MUST NOT accept a clear-text EAP success or failure packet prior to the peer and server reaching synchronized protected result indication.

The Crypto-Binding TLV exchange is used to prove that both the peer and server participated in the tunnel establishment and sequence of authentications. It also provides verification of the TEAP type, version negotiated, outer TLVs exchanged before the TLS tunnel establishment. The Crypto-Binding TLV MUST be exchanged and verified before the final Result TLV exchange, regardless whether there is an inner EAP method authentication or not. It MUST be included with the Intermediate-Result TLV to perform Cryptographic Binding after each

successful EAP method in a sequence of EAP methods, before proceeding with another inner EAP method. The server may send the final Result TLV along with an Intermediate-Result TLV and a Crypto-Binding TLV to indicate its intention to end the conversation. If the peer requires nothing more from the server, it will respond with a Result TLV indicating success accompanied by a Crypto-Binding TLV and Intermediate-Result TLV if necessary. The server then tears down the tunnel and sends a clear text EAP-Success or EAP-Failure.

If the peer receives a Result TLV indicating success from the server, but its authentication policies are not satisfied (for example it requires a particular authentication mechanism be run or it wants to request a PAC), it may request further action from the server using the Request-Action TLV. The Request-Action TLV is sent with a Status field indicating what EAP Success/Failure result the peer would expect if the requested action is not granted. The value of the Action field indicates what the peer would like to do next. The format and values for the Request-Action TLV are defined in Section 4.2.9.

Upon receiving the Request-Action TLV the server may process the request or ignore it, based on its policy. If the server ignores the request, it proceeds with termination of the tunnel and send the clear text EAP Success or Failure message based on the value of the peer's result TLV. If the server honors and processes the request, it continues with the requested action. The conversation completes with a Result TLV exchange. The Result TLV may be included with the TLV that completes the requested action.

Error handling for Phase 2 is discussed in Section 3.6.2.

3.4. Determining Peer-Id and Server-Id

The Peer-Id and Server-Id [RFC5247] may be determined based on the types of credentials used during either the TEAP tunnel creation or authentication. In the case of multiple peer authentications, all authenticated peer identities need to be exported.

When X.509 certificates are used for peer authentication, the Peer-Id is determined by the subject or subjectAltName fields in the peer certificate. As noted in [RFC5280]:

The subject field identifies the entity associated with the public key stored in the subject public key field. The subject name MAY be carried in the subject field and/or the subjectAltName extension.... If subject naming information is present only in the subjectAltName extension (e.g., a key bound only to an email address or URI), then the subject name MUST be an empty sequence

and the subjectAltName extension MUST be critical.

Where it is non-empty, the subject field MUST contain an X.500 distinguished name (DN).

If an inner EAP method is run, then the Peer-Id is obtained from the inner method.

When the server uses an X.509 certificate to establish the TLS tunnel, the Server-Id is determined in a similar fashion as stated above for the Peer-Id; e.g., the subject or subjectAltName field in the server certificate defines the Server-Id.

3.5. TEAP Session Identifier

The EAP session identifier [RFC5247] is constructed using the `tls_unique` from the TLS tunnel establishment as defined by [RFC5929]. The Session-Id is defined as follows:

Session-Id = `teap_type || tls_unique`

where `teap_type` is the EAP method type assigned to TEAP.

`tls_unique` = `tls_unique` for the phase 1 outer tunnel as defined by [RFC5929].

3.6. Error Handling

TEAP uses the following error handling rules summarized below:

1. Errors in the TLS layer are communicated via TLS alert messages in all phases of TEAP.
2. The Intermediate-Result TLVs carry success or failure indications of the individual EAP methods in TEAP Phase 2. Errors within the EAP conversation in Phase 2 are expected to be handled by individual EAP methods.
3. Violations of the TLV rules are handled using Result TLVs together with Error TLVs.
4. Tunnel compromised errors (errors caused by Crypto-Binding failed or missing) are handled using Result TLVs and Error TLVs.

3.6.1. TLS Layer Errors

If the TEAP server detects an error at any point in the TLS Handshake or the TLS layer, the server SHOULD send a TEAP request encapsulating a TLS record containing the appropriate TLS alert message rather than immediately terminating the conversation so as to allow the peer to inform the user of the cause of the failure and possibly allow for a restart of the conversation. The peer MUST send a TEAP response to an alert message. The EAP-Response packet sent by the peer may encapsulate a TLS ClientHello handshake message, in which case the TEAP server MAY allow the TEAP conversation to be restarted, or it MAY contain a TEAP response with a zero-length message, in which case the server MUST terminate the conversation with an EAP-Failure packet. It is up to the TEAP server whether to allow restarts, and if so, how many times the conversation can be restarted. A TEAP server implementing restart capability SHOULD impose a limit on the number of restarts, so as to protect against denial-of-service attacks.

If the TEAP peer detects an error at any point in the TLS layer, the TEAP peer should send a TEAP response encapsulating a TLS record containing the appropriate TLS alert message. The server may restart the conversation by sending an TEAP request packet encapsulating the TLS HelloRequest handshake message. The peer may allow the TEAP conversation to be restarted or it may terminate the conversation by sending an TEAP response with an zero-length message.

3.6.2. Phase 2 Errors

Any time the peer or the server finds a fatal error outside of the TLS layer during Phase 2 TLV processing, it MUST send a Result TLV of failure and an Error TLV with the appropriate error code. For errors involving the processing of the sequence of exchanges, such as a violation of TLV rules (e.g., multiple EAP-Payload TLVs), the error code is `Unexpected_TLVs_Exchanged`. For errors involving a tunnel compromise, the error-code is `Tunnel_Compromise_Error`. Upon sending a Result TLV with a fatal Error TLV the sender terminates the TLS tunnel. Note that a server will still wait for a message from the peer after it sends a failure, however the server does not need to process the contents of the response message.

If a server receives a Result TLV of failure with a fatal Error TLV, it SHOULD send a clear text EAP-Failure. If a peer receives a Result TLV of failure, it MUST respond with a Result TLV indicating failure. If the server has sent a Result TLV of failure, it ignores the peer response, and it SHOULD send a clear text EAP-Failure.

3.7. Fragmentation

A single TLS record may be up to 16384 octets in length, but a TLS message may span multiple TLS records, and a TLS certificate message may in principle be as long as 16 MB. This is larger than the maximum size for a message on most media types, therefore it is desirable to support fragmentation. Note that in order to protect against reassembly lockup and denial-of-service attacks, it may be desirable for an implementation to set a maximum size for one such group of TLS messages. Since a typical certificate chain is rarely longer than a few thousand octets, and no other field is likely to be anywhere near as long, a reasonable choice of maximum acceptable message length might be 64 KB. This is still a fairly large message packet size so an TEAP implementation MUST provide its own support for fragmentation and reassembly.

Since EAP is a lock-step protocol, fragmentation support can be added in a simple manner. In EAP, fragments that are lost or damaged in transit will be retransmitted, and since sequencing information is provided by the Identifier field in EAP, there is no need for a fragment offset field.

TEAP fragmentation support is provided through the addition of flag bits within the EAP-Response and EAP-Request packets, as well as a TLS Message Length field of four octets. Flags include the Length included (L), More fragments (M), and TEAP Start (S) bits. The L flag is set to indicate the presence of the four-octet TLS Message Length field, and MUST be set for the first fragment of a fragmented TLS message or set of messages. The M flag is set on all but the last fragment. The S flag is set only within the TEAP start message sent from the EAP server to the peer. The TLS Message Length field is four octets, and provides the total length of the TLS message or set of messages that is being fragmented; this simplifies buffer allocation.

When a TEAP peer receives an EAP-Request packet with the M bit set, it MUST respond with an EAP-Response with EAP-Type of TEAP and no data. This serves as a fragment ACK. The EAP server must wait until it receives the EAP-Response before sending another fragment. In order to prevent errors in processing of fragments, the EAP server MUST increment the Identifier field for each fragment contained within an EAP-Request, and the peer must include this Identifier value in the fragment ACK contained within the EAP-Response. Retransmitted fragments will contain the same Identifier value.

Similarly, when the TEAP server receives an EAP-Response with the M bit set, it must respond with an EAP-Request with EAP-Type of TEAP and no data. This serves as a fragment ACK. The EAP peer MUST wait

until it receives the EAP-Request before sending another fragment. In order to prevent errors in the processing of fragments, the EAP server MUST increment the Identifier value for each fragment ACK contained within an EAP-Request, and the peer MUST include this Identifier value in the subsequent fragment contained within an EAP-Response.

3.8. PAC Provisioning

To request provisioning of a PAC, a peer sends a PAC TLV as defined in Section 4.2.12 containing a PAC Attribute as defined in Section 4.2.12.1 of PAC Type set to the appropriate value. The request MAY be issued after the peer has determined that it has successfully authenticated the EAP server and validated the Crypto-Binding TLV as defined in Section 4.2.13 to ensure that the TLS tunnel's integrity is intact. The peer MUST send separate PAC TLVs for each type of PAC it wants to be provisioned. Multiple PAC TLVs can be sent in the same packet or different packets. The EAP server will send the PACs after its internal policy has been satisfied, or it MAY ignore the request or request additional authentications if its policy dictates. If a peer receives a PAC with an unknown type, it MUST ignore it.

A PAC-TLV containing PAC-Acknowledge attribute MUST be sent by the peer to acknowledge the receipt of the Tunnel PAC. A PAC-TLV containing PAC-Acknowledge attribute MUST NOT be used by the peer to acknowledge the receipt of other types of PACs.

3.9. Certificate Provisioning Within the Tunnel

Provisioning of a peer's certificate is supported in TEAP by performing the Simple PKI Request/Response from [RFC5272] using PKCS#10 and PKCS#7 TLVs, respectively. A peer sends the Simple PKI Request using a PKCS#10 CertificateRequest [RFC2986] encoded into the body of a PKCS#10 TLV (see section Section 4.2.17). The TEAP Server issues a Simple PKI Response using a PKCS#7 [RFC2315] degenerate "certs-only" message encoded into the body of a PKCS#7 TLV (see section Section 4.2.16).

In order to provide linking identity and proof-of-possession by including information specific to the current authenticated TLS session within the signed certification request, the client generating the request SHOULD obtain the tls-unique value as defined in Channel Bindings for TLS [RFC5929] from the TLS subsystem, encode it using base64 encoding, and place the resulting string in the certification request challenge password field. The tls-unique value used MUST be from the first TLS handshake. TEAP client and server must use their tls-unique implementation specific synchronization

methods to obtain this first tls-unique value. The server SHOULD verify the tls-unique information. This ensures that the authenticated TEAP client is in possession of the private key used to sign the certification request.

The Simple PKI Request/Response generation and processing rules of [RFC5272] SHALL apply to TEAP, with the exception of error conditions. In the event of an error, the TEAP Server SHOULD respond with an Error TLV using the most descriptive error code possible; it MAY ignore the PKCS#10 request which generated the error.

3.10. Server Unauthenticated Provisioning Mode

In Server Unauthenticated Provisioning Mode, an unauthenticated tunnel is established in phase 1 and the peer and server negotiate an EAP method in phase 2 that supports mutual authentication and key derivation that is resistant to attacks such as Man-in-the-middle and dictionary attacks. This provisioning mode enables the bootstrapping of peers when the peer lacks a strong credential usable for mutual authentication with the server during phase 1.

Upon successful completion of the EAP method in phase 2, the peer and server exchange a Crypto-Binding TLV to bind the inner method with the outer tunnel and ensure that a man-in-the-middle attack has not been attempted.


Support for the Server Unauthenticated Provisioning Mode is optional. The cipher suite TLS_DH_anon_WITH_AES_128_CBC_SHA is RECOMMENDED when using server unauthenticated mode, but other anonymous ciphersuites MAY be supported as long as the TLS pre-master secret is generated from contribution from both peers. Phase 2 EAP methods used in Server Unauthenticated Provisioning Mode MUST provide mutual authentication, key generation, and be resistant to dictionary attack. Example inner methods include EAP-pwd [RFC5931] and EAP-EKE [RFC6124].

4. Message Formats

The following sections describe the message formats used in TEAP. The fields are transmitted from left to right in network byte order.

4.1. TEAP Message Format

A summary of the TEAP Request/Response packet format is shown below.



Code

The code field is one octet in length defined as follows:

- 1 Request
- 2 Response

Identifier

The Identifier field is one octet and aids in matching responses with requests. The Identifier field MUST be changed on each Request packet. The Identifier field in the Response packet MUST match the Identifier field from the corresponding request.

Length

The Length field is two octets and indicates the length of the EAP packet including the Code, Identifier, Length, Type, Flags, Ver, Message Length, TLS Data, and Outer TLVs fields. Octets outside the range of the Length field should be treated as Data Link Layer padding and should be ignored on reception.

Type

TBD for TEAP

Flags

```
  0 1 2 3 4
+---+---+---+
|L M S R R|
+---+---+---+
```

L Length included; set to indicate the presence of the four octet Message Length field

M More fragments; set on all but the last fragment

S TEAP start; set in a TEAP Start message

R Reserved (must be zero)

Ver

This field contains the version of the protocol. This document describes version 1 (001 in binary) of TEAP.

Message Length

The Message Length field is four octets, and is present only if the L bit is set. This field provides the total length of the message that may be fragmented over the data fields of multiple packets.

TLS Data

When the Data field is present, it consists of an encapsulated TLS packet in TLS record format. A TEAP packet with Flags and Version fields, but with zero length TLS data field, is used to indicate TEAP acknowledgement for either a fragmented message, a TLS Alert message or a TLS Finished message.

Outer TLVs

The Outer-TLVs consist of the optional data used to help establishing the TLS tunnel in TLV format. They are only allowed in the first two messages in the TEAP protocol. That is the first EAP server to peer message and first peer to EAP server message. The start of the Outer-TLV can be derived from the EAP Length field and Message Length field.

4.2. TEAP TLV Format and Support

The TLVs defined here are standard Type-Length-Value (TLV) objects. The TLV objects could be used to carry arbitrary parameters between EAP peer and EAP server within the protected TLS tunnel.

The EAP peer may not necessarily implement all the TLVs supported by the EAP server. To allow for interoperability, TLVs are designed to allow an EAP server to discover if a TLV is supported by the EAP peer, using the NAK TLV. The mandatory bit in a TLV indicates whether support of the TLV is required. If the peer or server does not support a TLV marked mandatory, then it MUST send a NAK TLV in the response, and all the other TLVs in the message MUST be ignored. If an EAP peer or server finds an unsupported TLV that is marked as optional, it can ignore the unsupported TLV. It MUST NOT send an NAK TLV for a TLV that is not marked mandatory. If all TLVs in a message are marked optional and none are understood by the peer, then an EMPTY TEAP Phase 2 message must still be sent to the other side in order to continue the conversation.

Note that a peer or server may support a TLV with the mandatory bit set, but may not understand the contents. The appropriate response to a supported TLV with content that is not understood is defined by the individual TLV specification.

EAP implementations compliant with this specification MUST support TLV exchanges, as well as the processing of mandatory/optional settings on the TLV. Implementations conforming to this specification MUST support the following TLVs:

Authority-ID TLV

Identity-Type TLV

Result TLV

NAK TLV

Error TLV

Request-Action TLV

EAP-Payload TLV

Intermediate-Result TLV


Crypto-Binding TLV

Basic-Password-Auth-Req TLV

Basic-Password-Auth-Resp TLV

4.2.1. General TLV Format

TLVs are defined as described below. The fields are transmitted from left to right.



M

- 0 Optional TLV
- 1 Mandatory TLV

R

Reserved, set to zero (0)

TLV Type

A 14-bit field, denoting the TLV type. Allocated Types include:

- 0 Unassigned
- 1 Authority-ID TLV (Section 4.2.2)

- 2 Identity-Type TLV (Section 4.2.3)
- 3 Result TLV (Section 4.2.4)
- 4 NAK TLV (Section 4.2.5)
- 5 Error TLV (Section 4.2.6)
- 6 Channel-Binding TLV (Section 4.2.7)
- 7 Vendor-Specific TLV (Section 4.2.8)
- 8 Request-Action TLV (Section 4.2.9)
- 9 EAP-Payload TLV (Section 4.2.10)
- 10 Intermediate-Result TLV (Section 4.2.11)
- 11 PAC TLV (Section 4.2.12)
- 12 Crypto-Binding TLV (Section 4.2.13)
- 13 Basic-Password-Auth-Req TLV (Section 4.2.14)
- 14 Basic-Password-Auth-Resp TLV (Section 4.2.15)
- 15 PKCS#7 TLV (Section 4.2.16)
- 16 PKCS#10 TLV (Section 4.2.17)
- 17 Server-Trusted-Root TLV (Section 4.2.18)


Length

The length of the Value field in octets.

Value

The value of the TLV.

4.2.2. Authority-ID TLV



M

Mandatory, set to (0)

R

Reserved, set to zero (0)

Type

The Type field is two octets. It is set to 1 for Authority ID

Length

The Length field is two octets, which contains the length of the ID field in octets.

ID


Hint of the identity of the server, to help the peer to match the credentials available for the server. It should be unique across the deployment.

4.2.3. Identity-Type TLV

The Identity-Type TLV allows an EAP server to send a hint to help the EAP peer select the right type of identity; for example; user or machine. TEAPv1 implementations MUST support this TLV. If the Identity-Type field does not contain one of the known values or if the EAP peer does not have an identity corresponding to the identity-type, then the peer SHOULD respond with an Identity-Type TLV with the one of available identity types. If the server receives an identity

type in the response that does not match the requested type, then the peer does not possess the requested credential type and the server SHOULD proceed with authentication for the credential type proposed by the peer or proceed with requesting another credential type, or simply apply the network policy based on the configured policy, e.g., sending Result TLV with Failure.

The Identity-Type TLV is defined as follows:




- M
 - 0 (Optional)
- R
 - Reserved, set to zero (0)
- TLV Type
 - 2 for Identity-Type TLV
- Length
 - 2
- Identity-Type
 - The Identity-Type field is two octets. Values include:
 - 1 User

2 Machine

4.2.4. Result TLV

The Result TLV provides support for acknowledged success and failure messages for protected termination within TEAP. If the Status field does not contain one of the known values, then the peer or EAP server MUST treat this as a fatal error of Unexpected_TLVs_Exchanged. The behavior of the Result TLV is further discussed in Section 3.3.3 and Section 3.6.2. A Result TLV indicating failure MUST NOT be accompanied by the following TLVs: NAK, EAP-Payload TLV, or Crypto-Binding TLV. The Result TLV is defined as follows:



M
Mandatory, set to one (1)

R
Reserved, set to zero (0)

TLV Type
3 for Result TLV

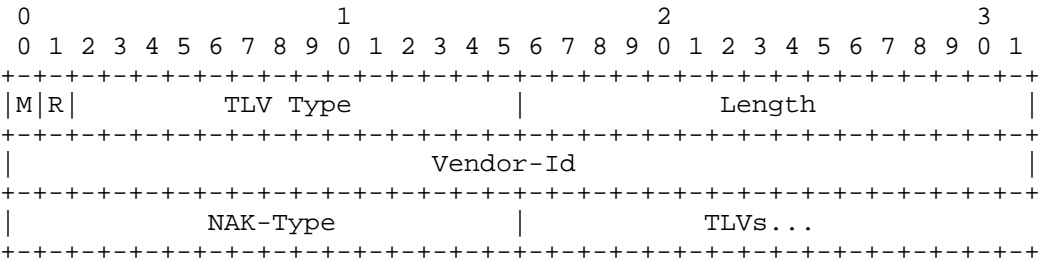
Length
2

Status
The Status field is two octets. Values include:

- 1 Success
- 2 Failure

4.2.5. NAK TLV

The NAK TLV allows a peer to detect TLVs that are not supported by the other peer. A TEAP packet can contain 0 or more NAK TLVs. A NAK TLV should not be accompanied by other TLVs. A NAK TLV MUST NOT be sent in response to a message containing a Result TLV, instead a Result TLV of failure should be sent indicating failure and an Error TLV of Unexpected_TLVs_Exchanged. The NAK TLV is defined as follows:



M
Mandatory, set to one (1)

R
Reserved, set to zero (0)

TLV Type
4 for NAK TLV

Length
>=6

Vendor-Id

The Vendor-Id field is four octets, and contains the Vendor-Id of the TLV that was not supported. The high-order octet is 0 and the low-order three octets are the Structure of Management Information (SMI) Network Management Private Enterprise Code of the Vendor in network byte order. The Vendor-Id field MUST be zero for TLVs that are not Vendor-Specific TLVs.

NAK-Type


The NAK-Type field is two octets. The field contains the Type of the TLV that was not supported. A TLV of this Type MUST have been included in the previous packet.

TLVs

This field contains a list of zero or more TLVs, each of which MUST NOT have the mandatory bit set. These optional TLVs are for future extensibility to communicate why the offending TLV was determined to be unsupported.

4.2.6. Error TLV

The Error TLV allows an EAP peer or server to indicate errors to the other party. A TEAP packet can contain 0 or more Error TLVs. The Error-Code field describes the type of error. Error Codes 1-999 represent successful outcomes (informative messages), 1000-1999 represent warnings, and codes 2000-2999 represent fatal errors. A fatal Error TLV MUST be accompanied by a Result TLV indicating failure and the conversation must be terminated as described in Section 3.6.2. The Error TLV is defined as follows:



M

Mandatory, set to one (1)

R

Reserved, set to zero (0)

TLV Type

5 for Error TLV

Length

4

Error-Code

The Error-Code field is four octets. Currently defined values for Error-Code include:

2001 Tunnel_Compromise_Error

2002 Unexpected_TLVs_Exchanged

2003 Unsupported_Algorithm_In_CertificateSigning_Request

2004 Unsupported_Extension_In_CertificateSigning_Request

2005 Bad_Identity_In_CertificateSigning_Request


2006 Bad_CertificateSigning_Request

2007 Internal_CA_Error

2008 General_PKI_Error

4.2.7. Channel-Binding TLV

The Channel-Binding TLV provides a mechanism for carrying channel binding data from the peer to the EAP server and a channel binding response from the EAP server to the peer as described in [I-D.ietf-emu-chbind]. TEAPv1 implementations MAY support this TLV, which cannot be responded to with a NAK TLV. If the Channel-Binding data field does not contain one of the known values or if the EAP server does not support this TLV, then the server MUST ignore the value. The Channel-Binding TLV is defined as follows:



M

0 (Optional)

R

Reserved, set to zero (0)

TLV Type

6 for Channel-Binding TLV

Length

variable

Data

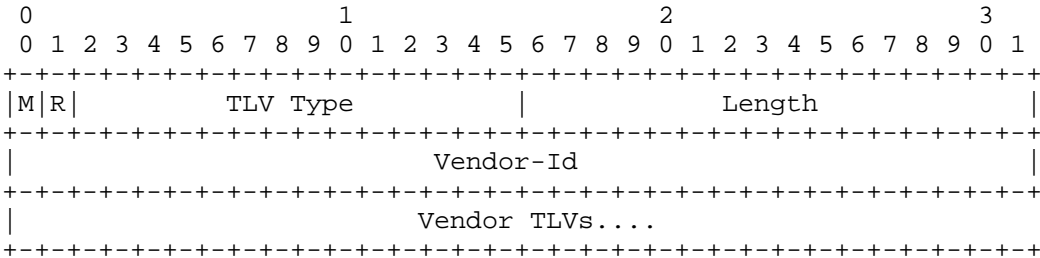
The data field contains channel binding data defined in [I-D.ietf-emu-chbind].

4.2.8. Vendor-Specific TLV

The Vendor-Specific TLV is available to allow vendors to support their own extended attributes not suitable for general usage. A Vendor-Specific TLV attribute can contain one or more TLVs, referred to as Vendor TLVs. The TLV-type of a Vendor-TLV is defined by the vendor. All the Vendor TLVs inside a single Vendor-Specific TLV belong to the same vendor. There can be multiple Vendor-Specific TLVs from different vendors in the same message.

Vendor TLVs may be optional or mandatory. Vendor TLVs sent with Result TLVs MUST be marked as optional.

The Vendor-Specific TLV is defined as follows:



- M
 - 0 or 1
- R
 - Reserved, set to zero (0)
- TLV Type
 - 7 for Vendor Specific TLV
- Length
 - 4 + cumulative length of all included Vendor TLVs

Vendor-Id

The Vendor-Id field is four octets, and contains the Vendor-Id of the TLV. The high-order octet is 0 and the low-order 3 octets are the SMI Network Management Private Enterprise Code of the Vendor in network byte order.

Vendor TLVs

This field is of indefinite length. It contains vendor-specific TLVs, in a format defined by the vendor.

4.2.9. Request-Action TLV

The Request-Action TLV MAY be sent by the peer in response to a server's successful Result TLV. It allows the peer to request the EAP server to negotiate additional EAP methods or process TLVs specified in the response packet. The server MAY ignore this TLV.

The peer MAY send multiple Request-Action TLVs to the server. Two Request TLVs MUST NOT occur in the same response packet if they have the same Status value. The order of processing multiple Request TLVs is implementation dependent. If the server process the optional (non-fatal) items first, it is possible that the fatal items will disappear at a later time. If the server process the fatal items first, the communication time will be shorter.

The client MAY return a new set of Request-Action TLVs after one or more of the requested items has been processed and the server has signaled it wants to end the EAP conversation.

The Request-Action TLV is defined as follows:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							
M R										TLV Type										Length																			
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							
										Status										Action										TLVs....									
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+																																							

M

Mandatory set to one (1)

R

Reserved, set to zero (0)

TLV Type

8 for Request-Action TLV

Length

2 + cumulative length of all included TLVs

Status

The Status field is one octet. This indicates the result if the server does not process the action requested by the peer. Values include:

1 Success

2 Failure

Action

The Action field is one octet. Values include:

1 Process-TLV

2 Negotiate-EAP


TLVs

This field is of indefinite length. It contains TLVs that the peer wants the server to process.

4.2.10. EAP-Payload TLV

To allow piggybacking an EAP request or response with other TLVs, the EAP-Payload TLV is defined, which includes an encapsulated EAP packet and a list of optional TLVs. The optional TLVs are provided for future extensibility to provide hints about the current EAP

authentication. Only one EAP-Payload TLV is allowed in a message. The EAP-Payload TLV is defined as follows:




- M
 - Mandatory, set to (1)
- R
 - Reserved, set to zero (0)
- TLV Type
 - 9 for EAP-Payload TLV
- Length
 - length of embedded EAP packet + cumulative length of additional TLVs
- EAP packet
 - This field contains a complete EAP packet, including the EAP header (Code, Identifier, Length, Type) fields. The length of this field is determined by the Length field of the encapsulated EAP packet.
- TLVs
 - This (optional) field contains a list of TLVs associated with the EAP packet field. The TLVs MUST NOT have the mandatory bit

set. The total length of this field is equal to the Length field of the EAP-Payload TLV, minus the Length field in the EAP header of the EAP packet field.

4.2.11. Intermediate-Result TLV

The Intermediate-Result TLV provides support for acknowledged intermediate Success and Failure messages between multiple inner EAP methods within EAP. An Intermediate-Result TLV indicating success MUST be accompanied by a Crypto-Binding TLV. The optional TLVs associated with this TLV are provided for future extensibility to provide hints about the current result. The Intermediate-Result TLV is defined as follows:



- M
- Mandatory, set to (1)
- R
- Reserved, set to zero (0)

TLV Type

10 for Intermediate-Result TLV

Length

2 + cumulative length of the embedded associated TLVs

Status

The Status field is two octets. Values include:

- 1 Success
- 2 Failure

TLVs

This field is of indeterminate length, and contains zero or more of the TLVs associated with the Intermediate Result TLV. The TLVs in this field MUST NOT have the mandatory bit set.

4.2.12. PAC TLV Format

The PAC TLV provides support for provisioning the Protected Access Credential (PAC) defined within [RFC4851]. The PAC TLV carries the PAC and related information within PAC attribute fields. Additionally, the PAC TLV MAY be used by the peer to request provisioning of a PAC of the type specified in the PAC Type PAC attribute. The PAC TLV MUST only be used in a protected tunnel providing encryption and integrity protection. A general PAC TLV format is defined as follows:

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|M|R|          TLV Type          |          Length          |
+-----+-----+-----+-----+-----+-----+-----+-----+
|          PAC Attributes...          |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

M

- 0 - Non-mandatory TLV
- 1 - Mandatory TLV

R

Reserved, set to zero (0)

TLV Type

- 11 - PAC TLV

Length


Two octets containing the length of the PAC attributes field in octets.

PAC Attributes

A list of PAC attributes in the TLV format.

4.2.12.1. Formats for PAC Attributes

Each PAC attribute in a PAC TLV is formatted as a TLV defined as follows:



Type

The Type field is two octets, denoting the attribute type. Allocated Types include:

- 1 - PAC-Key
- 2 - PAC-Opaque
- 3 - PAC-Lifetime
- 4 - A-ID
- 5 - I-ID
- 6 - Reserved
- 7 - A-ID-Info
- 8 - PAC-Acknowledgement
- 9 - PAC-Info
- 10 - PAC-Type

Length


Two octets containing the length of the Value field in octets.

Value

The value of the PAC attribute.

4.2.12.2. PAC-Key

The PAC-Key is a secret key distributed in a PAC attribute of type PAC-Key. The PAC-Key attribute is included within the PAC TLV whenever the server wishes to issue or renew a PAC that is bound to a key such as a Tunnel PAC. The key is a randomly generated octet string, which is 48 octets in length. The generator of this key is the issuer of the credential, which is identified by the Authority Identifier (A-ID).



Type

1 - PAC-Key

Length

2-octet length indicating the length of the key

Key

The value of the PAC-Key.


4.2.12.3. PAC-Opaque

The PAC-Opaque attribute is included within the PAC TLV whenever the server wishes to issue or renew a PAC.

The PAC-Opaque is opaque to the peer and thus the peer MUST NOT attempt to interpret it. A peer that has been issued a PAC-Opaque by a server stores that data and presents it back to the server according to its PAC Type. The Tunnel PAC is used in the ClientHello SessionTicket extension field defined in [RFC5077]. If a peer has

opaque data issued to it by multiple servers, then it stores the data issued by each server separately according to the A-ID. This requirement allows the peer to maintain and use each opaque datum as an independent PAC pairing, with a PAC-Key mapping to a PAC-Opaque identified by the A-ID. As there is a one-to-one correspondence between the PAC-Key and PAC-Opaque, the peer determines the PAC-Key and corresponding PAC-Opaque based on the A-ID provided in the TEAP/Start message and the A-ID provided in the PAC-Info when it was provisioned with a PAC-Opaque.

The PAC-Opaque attribute format is summarized as follows:



Type

2 - PAC-Opaque

Length


The Length field is two octets, which contains the length of the Value field in octets.

Value

The Value field contains the actual data for the PAC-Opaque. It is specific to the server implementation.

4.2.12.4. PAC-Info

The PAC-Info is comprised of a set of PAC attributes as defined in Section 4.2.12.1. The PAC-Info attribute MUST contain the A-ID, A-ID-Info, and PAC-Type attributes. Other attributes MAY be included in the PAC-Info to provide more information to the peer. The PAC-Info attribute MUST NOT contain the PAC-Key, PAC-Acknowledgement, PAC-Info, or PAC-Opaque attributes. The PAC-Info attribute is included within the PAC TLV whenever the server wishes to issue or renew a PAC.



Type

- 9 - PAC-Info

Length

2-octet Length field containing the length of the attributes field in octets.

Attributes

The attributes field contains a list of PAC attributes. Each mandatory and optional field type is defined as follows:

- 3 - PAC-LIFETIME

This is a 4-octet quantity representing the expiration time of the credential expressed as the number of seconds, excluding leap seconds, after midnight UTC, January 1, 1970. This attribute MAY be provided to the peer as part of the PAC-Info.

- 4 - A-ID

The A-ID is the identity of the authority that issued the PAC. The A-ID is intended to be unique across all issuing servers to avoid namespace collisions. The A-ID is used by the peer to determine which PAC to employ. The A-ID is treated as an opaque octet string. This attribute MUST be included in the PAC-Info attribute. The A-ID MUST match the Authority-ID the server used to establish the tunnel. One method for generating the A-ID is to use a high-quality random number generator to generate a random number. An alternate method would be to take the hash of the public key or public key certificate belonging a server represented by the A-ID.

5 - I-ID

Initiator identifier (I-ID) is the peer identity associated with the credential. This identity is derived from the inner EAP exchange or from the client-side authentication during tunnel establishment if inner EAP method authentication is not used. The server employs the I-ID in the TEAP phase 2 conversation to validate that the same peer identity used to execute TEAP phase 1 is also used in at minimum one inner EAP method in TEAP phase 2. If the server is enforcing the I-ID validation on the inner EAP method, then the I-ID MUST be included in the PAC-Info, to enable the peer to also enforce a unique PAC for each unique user. If the I-ID is missing from the PAC-Info, it is assumed that the Tunnel PAC can be used for multiple users and the peer will not enforce the unique-Tunnel-PAC-per-user policy.

7 - A-ID-Info


Authority Identifier Information is intended to provide a user-friendly name for the A-ID. It may contain the enterprise name and server name in a human-readable format. This TLV serves as an aid to the peer to better inform the end-user about the A-ID. The name is encoded in UTF-8 [RFC3629] format. This attribute MUST be included in the PAC-Info.

10 - PAC-type

The PAC-Type is intended to provide the type of PAC. This attribute SHOULD be included in the PAC-Info. If the PAC-Type is not present, then it defaults to a Tunnel PAC (Type 1).

4.2.12.5. PAC-Acknowledgement TLV

The PAC-Acknowledgement is used to acknowledge the receipt of the Tunnel PAC by the peer. The peer includes the PAC-Acknowledgement TLV in a PAC-TLV sent to the server to indicate the result of the processing and storing of a newly provisioned Tunnel PAC. This TLV is only used when Tunnel PAC is provisioned.



+-----+

Type

8 - PAC-Acknowledgement

Length

The length of this field is two octets containing a value of 2.


Result

The resulting value MUST be one of the following:

- 1 - Success
- 2 - Failure

4.2.12.6. PAC-Type TLV

The PAC-Type TLV is a TLV intended to specify the PAC type. It is included in a PAC-TLV sent by the peer to request PAC provisioning from the server. Its format is described below:



Type

10 - PAC-Type

Length

2-octet Length field with a value of 2

PAC Type

This 2-octet field defines the type of PAC being requested or provisioned. The following values are defined:

1 - Tunnel PAC

4.2.13. Crypto-Binding TLV

The Crypto-Binding TLV is used to prove that both the peer and server participated in the tunnel establishment and sequence of authentications. It also provides verification of the TEAP type, version negotiated, outer TLVs exchanged before the TLS tunnel establishment.

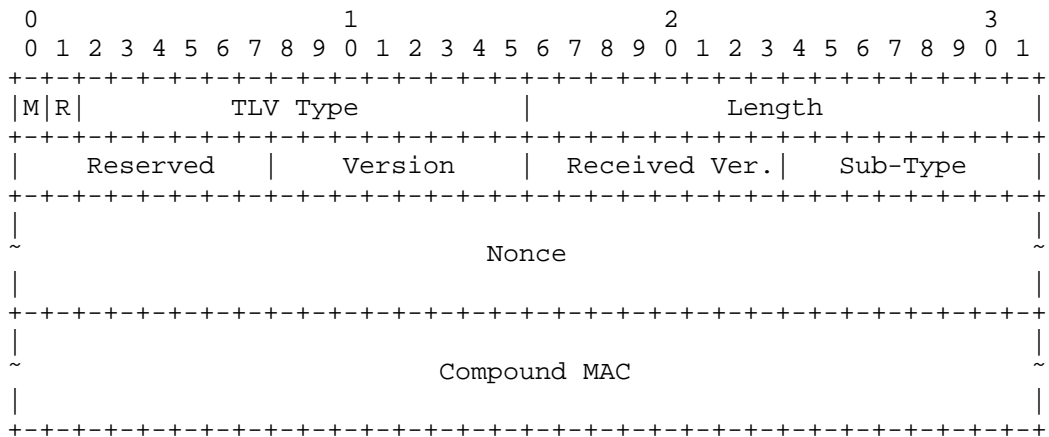
The Crypto-Binding TLV MUST be exchanged and verified before the final Result TLV exchange, regardless whether there is an inner EAP method authentication or not. It MUST be included with the Intermediate-Result TLV to perform Cryptographic Binding after each successful EAP method in a sequence of EAP methods, before proceeding with another inner EAP method.

The Crypto-Binding TLV is valid only if the following checks pass:

- o The Crypto-Binding TLV version is supported
- o The MAC verifies correctly
- o The received version in the Crypto-Binding TLV matches the version sent by the receiver during the EAP version negotiation
- o The subtype is set to the correct value

If any of the above checks fails, then the TLV is invalid. An invalid Crypto-Binding TLV is a fatal error and is handled as described in Section 3.6.2

The Crypto-Binding TLV is defined as follows:



M

Mandatory, set to (1)

R

Reserved, set to zero (0)

TLV Type

12 for Crypto-Binding TLV

Length

56

Reserved

Reserved, set to zero (0)

Version

The Version field is a single octet, which is set to the version of Crypto-Binding TLV the EAP method is using. For an implementation compliant with this version of TEAP, the version

number MUST be set to 1.

Received Version

The Received Version field is a single octet and MUST be set to the EAP version number received during version negotiation. Note that this field only provides protection against downgrade attacks, where a version of EAP requiring support for this TLV is required on both sides.

Sub-Type

The Sub-Type field is one octet. Defined values include

- 0 Binding Request
- 1 Binding Response

Nonce

The Nonce field is 32 octets. It contains a 256-bit nonce that is temporally unique, used for compound MAC key derivation at each end. The nonce in a request MUST have its least significant bit set to 0 and the nonce in a response MUST have the same value as the request nonce except the least significant bit MUST be set to 1.


Compound MAC

The Compound MAC field is 20 octets. This can be the Server MAC (B1_MAC) or the Client MAC (B2_MAC). The computation of the MAC is described in Section 5.3.

4.2.14. Basic-Password-Auth-Req TLV

The Basic-Password-Auth-Req TLV is used by the authentication server to request a username and password from the peer. It contains an optional user prompt message for the request. The peer is expected to obtain the username and password and send them in a Basic-Password-Auth-Resp TLV.

The Basic-Password-Auth-Req TLV is defined as follows:



M

0 (Optional)

R

Reserved, set to zero (0)

TLV Type

13 for Basic-Password-Auth-Req TLV

Length

variable


Prompt

optional user prompt message in UTF-8 format

4.2.15. Basic-Password-Auth-Resp TLV

The Basic-Password-Auth-Resp TLV is used by the peer to respond to a Basic-Password-Auth-Req TLV with a username and password. The TLV contains a username and password. The username and password are in UTF-8 format.

The Basic-Password-Auth-Resp TLV is defined as follows:



M

0 (Optional)

R

Reserved, set to zero (0)

TLV Type

14 for Basic-Password-Auth-Resp TLV

Length

variable

Userlen

Length of Username field in octets

Username

Username in UTF-8 format

Passlen

Length of Password field in octets

Password

Password in UTF-8 format

4.2.16. PKCS#7 TLV

The PKCS#7 TLV is used by the EAP server to deliver (a) certificate(s) to the peer. The format consists of a certificate or certificate chain in a degenerate certificates-only PKCS#7 SignedData Content as defined in [RFC2311]. When used in response to a Trusted-Server-Root TLV request from the peer, the EAP server MUST send the PKCS#7 TLV inside a Trusted-Server-Root TLV. When used in response to a PKCS#10 certificate enrollment request from the peer, the EAP server MUST send the PKCS#7 TLV without a Trusted-Server-Root TLV. The PKCS#7 TLV is always marked as optional, which cannot be responded to with a NAK TLV. TEAP implementations that support the Trusted-Server-Root TLV or the PKCS#10 TLV MUST support this TLV. Peers MUST NOT assume that the certificates in a PKCS#7 TLV are in any order. TEAP Servers SHOULD include all intermediate certificates needed to form complete certificate paths to one or more trust anchors, and not just return the newly issued certificate(s). TEAP Servers MAY return CRLs in the CRL bag. TEAP Servers MAY return self-signed certificates. Peers that handle self-signed certificates or trust anchors MUST NOT implicitly trust these certificates merely due to their presence in the certificate bag. Note: Peer's are advised to take great care in deciding whether to use a received certificate as a trust anchor. The authenticated nature of the tunnel in which a PKCS#7 bag is received can provide a level of authenticity to the certificates contained therein. Peers are advised to take into account the implied authority of the EAP server and to constrain the trust it can achieve through the trust anchor received in a PKCS#7 TLV.

The PKCS#7 TLV is defined as follows:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|M|R|          TLV Type          |          Length          |
+-----+-----+-----+-----+-----+-----+-----+-----+
|          PKCS #7 Data...
+-----+-----+-----+-----+-----+-----+-----+-----+

```

M

0 - Optional TLV

 \mathbb{R}

Reserved, set to zero (0)

TLV Type

15 - PKCS#7 TLV

Length

The length of the PKCS #7 Data field.

PKCS #7 Data

This field contains the X.509 certificate or certificate chain in a Certificates-Only PKCS#7 SignedData message.

4.2.17. PKCS#10 TLV

The PKCS#10 TLV is used by the peer to initiate the "simple PKI" Request/Response from [RFC5272]. The format of the request is as specified in Section 6.4 of [RFC4945]. The PKCS#10 TLV is always marked as optional, which cannot be responded to with a NAK TLV.

The PKCS#10 TLV is defined as follows:

[illegible]

M

0 - Optional TLV

R

Reserved, set to zero (0)

TLV Type

16 - PKCS#10 TLV

Length

The length of the PKCS #10 Data field.

PKCS #10 Data

This field contains the PKCS#10 certificate request.

4.2.18. Trusted-Server-Root TLV

Trusted-Server-Root TLV facilitates the request and delivery of a trusted server root certificate. The Trusted-Server-Root TLV can be exchanged in regular TEAP authentication mode or provisioning mode. The Trusted-Server-Root TLV is always marked as optional, and cannot be responded to with a Negative Acknowledgement (NAK) TLV. The Trusted-Server-Root TLV MUST only be sent as an inner TLV (inside the protection of the tunnel).


After the peer has determined that it has successfully authenticated the EAP server and validated the Crypto-Binding TLV, it MAY send one or more Trusted-Server-Root TLVs (marked as optional) to request the trusted server root certificates from the EAP server. The EAP server MAY send one or more root certificates with a Public Key Cryptographic System #7 (PKCS#7) TLV inside Server-Trusted-Root TLV. The EAP server MAY also choose not to honor the request.

The Trusted-Server-Root TLV allows the peer to send a request to the EAP server for a list of trusted roots. The server may respond with one or more root certificates in PKCS#7 [RFC2315] format.

If the EAP server sets the credential format to PKCS#7-Server-Certificate-Root, then the Trusted-Server-Root TLV should contain the root of the certificate chain of the certificate issued to the EAP server packaged in a PKCS#7 TLV. If the Server certificate is a self-signed certificate, then the root is the self-signed certificate.

If the Trusted-Server-Root TLV credential format contains a value unknown to the peer, then the EAP peer should ignore the TLV.

The Trusted-Server-Root TLV is defined as follows:



M

0 - Non-mandatory TLV

R

Reserved, set to zero (0)

TLV Type

17 - Trusted-Server-Root TLV [RFC4851]

Length

>=2 octets

Credential-Format

The Credential-Format field is two octets. Values include:

1 - PKCS#7-Server-Certificate-Root

Cred TLVs

This field is of indefinite length. It contains TLVs associated with the credential format. The peer may leave this field empty when using this TLV to request server trust roots.

4.3. Table of TLVs

The following table provides a guide to which TLVs may be found in which kinds of messages, and in what quantity. The messages are as follows: Request is a TEAP Request, Response is a TEAP Response, Success is a message containing a successful Result TLV, and Failure is a message containing a failed Result TLV.

Request	Response	Success	Failure	TLVs
0-1	0-1	0	0	Authority-ID
0-1	0-1	0	0	Identity-Type
0-1	0-1	1	1	Result
0+	0+	0	0	NAK
0+	0+	0+	0+	Error
0-1	0-1	0	0	Channel-Binding
0+	0+	0+	0+	Vendor-Specific [NOTE1]
0	0-1	0-1	0-1	Request-Action
0-1	0-1	0	0	EAP-Payload
0-1	0-1	0-1	0-1	Intermediate-Result
0+	0+	0+	0	PAC-TLV
0-1	0-1	0-1	0-1	Crypto-Binding
0-1	0	0	0	Basic-Password-Auth-Req
0	0-1	0	0	Basic-Password-Auth-Resp
0-1	0	0-1	0	PKCS#7
0	0-1	0	0	PKCS#10
0-1	0-1	0-1	0	Server-Trusted-Root

[NOTE1] Vendor TLVs (included in Vendor-Specific TLVs) sent with a Result TLV MUST be marked as optional.

The following table defines the meaning of the table entries in the sections below:

0 This TLV MUST NOT be present in the message.

0+ Zero or more instances of this TLV MAY be present in the message.

0-1 Zero or one instance of this TLV MAY be present in the message.

1 Exactly one instance of this TLV MUST be present in the message.

5. Cryptographic Calculations

5.1. TEAP Authentication Phase 1: Key Derivations

With TEAPv1, the TLS master secret is generated as specified in TLS. If a PAC is used then the master secret is obtained as described in [RFC5077].

TEAPv1 makes use of the TLS Keying Material Exporters defined in [RFC5705] to derive the `session_key_seed`. The Label used in the derivation is "teap session key seed". The length of the session key seed material is 40 octets. No context data is used in the export process.

The `session_key_seed` is used by the TEAP Authentication Phase 2

conversation to both cryptographically bind the inner method(s) to the tunnel as well as generate the resulting TEAP session keys. The other quantities are used as they are defined in [RFC5246].

5.2. Intermediate Compound Key Derivations

The `session_key_seed` derived as part of TEAP Phase 2 is used in TEAP Phase 2 to generate an Intermediate Compound Key (IMCK) used to verify the integrity of the TLS tunnel after each successful inner authentication and in the generation of Master Session Key (MSK) and Extended Master Session Key (EMSK) defined in [RFC3748]. Note that the IMCK must be recalculated after each successful inner EAP method.

The first step in these calculations is the generation of the base compound key, `IMCK[n]` from the `session_key_seed` and any session keys derived from the successful execution of `n`th inner EAP methods. The inner EAP method(s) may provide Inner Method Session Keys (IMSK), `IMSK1..IMSKn`, corresponding to inner method 1 through `n`.

If an inner method supports export of an Extended Master Session Key (EMSK), then the IMSK is derived from the EMSK as defined in [RFC5295]. The usage label used is "TEAPbindkey@ietf.org" and the length is 64 octets. Optional data parameter is not used in the derivation.

```
IMSK = First 32 octets of KDF(EMSK, "TEAPbindkey@ietf.org" | "\0"
| 64)
```

where the KDF is defined in [RFC5295].

If an inner method does not support export of an Extended Master Session Key (EMSK), then IMSK is the MSK of the inner method. The MSK is truncated at 32 octets if it is longer than 32 octets or padded to a length of 32 octets with zeros if it is less than 32 octets.

If the `i`th inner method does not generate an EMSK or MSK, then `IMSKi` is set to zero (e.g., `MSKi` = 32 octets of 0x00s). If an inner method fails, then it is not included in this calculation. The derivations of S-IMCK is as follows:

```
S-IMCK[0] = session_key_seed
For j = 1 to n-1 do
    IMCK[j] = TLS-PRF(S-IMCK[j-1], "Inner Methods Compound Keys",
        IMSK[j], 60)
    S-IMCK[j] = first 40 octets of IMCK[j]
    CMK[j] = last 20 octets of IMCK[j]
```

where TLS-PRF is the PRF negotiated as part of TLS handshake [RFC5246].

5.3. Computing the Compound MAC

For authentication methods that generate keying material, further protection against man-in-the-middle attacks is provided through cryptographically binding keying material established by both TEAP Phase 1 and TEAP Phase 2 conversations. After each successful inner EAP authentication, EAP MSKs are cryptographically combined with key material from TEAP Phase 1 to generate a compound session key, CMK. The CMK is used to calculate the Compound MAC as part of the Crypto-Binding TLV described in Section 4.2.13, which helps provide assurance that the same entities are involved in all communications in TEAP. During the calculation of the Compound-MAC the MAC field is filled with zeros.

The Compound MAC computation is as follows:

```
CMK = CMK[j]
Compound-MAC = HMAC-HASH( CMK, BUFFER )
```

where j is the number of the last successfully executed inner EAP method, HASH is the default hash function or the alternative hash function negotiated in TLS 1.2 [RFC5246], and BUFFER is created after concatenating these fields in the following order:

- 1 The entire Crypto-Binding TLV attribute with the MAC field zeroed out.
- 2 The EAP Type sent by the other party in the first TEAP message.
- 3 All the Outer-TLVs from the first TEAP message sent by EAP server to peer. If a single TEAP message is fragmented into multiple TEAP packets; then the Outer-TLVs in all the fragments of that message MUST be included.
- 4 All the Outer-TLVs from the first TEAP message sent by the peer to the EAP server. If a single TEAP message is fragmented into multiple TEAP packets, then the Outer-TLVs in all the fragments of that message MUST be included.

5.4. EAP Master Session Key Generation

TEAP Authentication assures the master session key (MSK) and Extended Master Session Key (EMSK) output from the EAP method are the result of all authentication conversations by generating an Intermediate Compound Key (IMCK). The IMCK is mutually derived by the peer and the server as described in Section 5.2 by combining the MSKs from inner EAP methods with key material from TEAP Phase 1. The resulting MSK and EMSK are generated as part of the IMCKn key hierarchy as follows:

```
MSK  = TLS-PRF(S-IMCK[j], "Session Key Generating Function", 64)
EMSK = TLS-PRF(S-IMCK[j],
               "Extended Session Key Generating Function", 64)
```

where j is the number of the last successfully executed inner EAP method.

The EMSK is typically only known to the TEAP peer and server and is not provided to a third party. The derivation of additional keys and transportation of these keys to a third party is outside the scope of this document.

If no EAP methods have been negotiated inside the tunnel or no EAP methods have been successfully completed inside the tunnel, the MSK and EMSK will be generated directly from the `session_key_seed` meaning `S-IMCK = session_key_seed`.

6. IANA Considerations

This section provides guidance to the Internet Assigned Numbers Authority (IANA) regarding registration of values related to the TEAP protocol, in accordance with BCP 26, [RFC5226].

The EAP Method Type number for TEAP needs to be assigned.

The document defines a registry for TEAP TLV types, which may be assigned by Specification Required as defined in [RFC5226]. Section 4.2 defines the TLV types that initially populate the registry. A summary of the TEAP TLV types is given below:

0 Unassigned

- 1 Authority-ID TLV
- 2 Identity-Type TLV
- 3 Result TLV
- 4 NAK TLV
- 5 Error TLV
- 6 Channel-Binding TLV
- 7 Vendor-Specific TLV
- 8 Request-Action TLV
- 9 EAP-Payload TLV
- 10 Intermediate-Result TLV
- 11 PAC TLV
- 12 Crypto-Binding TLV
- 13 Basic-Password-Auth-Req TLV
- 14 Basic-Password-Auth-Resp TLV
- 15 PKCS#7 TLV
- 16 PKCS#10 TLV
- 17 Trusted-Server-Root TLV

The Identity-Type defined in Section 4.2.3 contains an Identity Type code which is assigned on a Specification Required basis as defined in [RFC5226]. The initial types defined are:

- 1 User
- 2 Machine

The Result TLV defined in Section 4.2.4, Request-Action TLV defined in Section 4.2.9, and Intermediate-Result TLV defined in Section 4.2.11 contain a Status code which is assigned on a Specification Required basis as defined in [RFC5226]. The initial types defined are:

- 1 Success

- 2 Failure

The Error-TLV defined in Section 4.2.6 requires an error-code. TEAP Error-TLV error-codes are assigned based on Specification Required as defined in [RFC5226]. The initial list of error codes is as follows:

- 2001 Tunnel_Compromise_Error

- 2002 Unexpected_TLVs_Exchanged

- 2003 Unsupported_Algorithm_In_CertificateSigning_Request

- 2004 Unsupported_Extension_In_CertificateSigning_Request

- 2005 Bad_Identity_In_CertificateSigning_Request

- 2006 Bad_CertificateSigning_Request

- 2007 Internal_CA_Error

- 2008 General_PKI_Error

The Request-Action TLV defined in Section 4.2.9 contains an action code which is assigned on a Specification Required basis as defined in [RFC5226]. The initial actions defined are:

- 1 Process-TLV

- 2 Negotiate-EAP

The PAC Attribute defined in Section 4.2.12.1 contains a Type code which is assigned on a Specification Required basis as defined in [RFC5226]. The initial types defined are:

- 1 PAC-key

- 2 PAC-Opaque

- 3 PAC-Lifetime

- 4 A-ID

- 5 I-ID
- 6 Reserved
- 7 A-ID-Info
- 8 PAC-Acknowledgement
- 9 PAC-Info
- 10 PAC-Type

The PAC-Type defined in Section 4.2.12.6 contains a Type code which is assigned on a Specification Required basis as defined in [RFC5226]. The initial types defined are:

- 1 Tunnel PAC

The Trusted-Server-Root TLV defined in Section 4.2.18 contains a Credential-Format code which is assigned on a Specification Required basis as defined in [RFC5226]. The initial types defined are:

- 1 PKCS#7-Server-Certificate-Root

The various values under Vendor-Specific TLV are assigned by Private Use and do not need to be assigned by IANA.

TEAP makes use of the TLS Keying Material Exporters defined in [RFC5705]. The Label used in the derivation as defined in Section 5.1 is "teap session key seed".

TEAP registers a TEAP binding usage label from the "USRK Key Labels" name space defined in [RFC5295] with a value "TEAPbindkey@ietf.org".

7. Security Considerations

TEAP is designed with a focus on wireless media, where the medium itself is inherent to eavesdropping. Whereas in wired media, an attacker would have to gain physical access to the wired medium; wireless media enables anyone to capture information as it is transmitted over the air, enabling passive attacks. Thus, physical security can not be assumed and security vulnerabilities are far greater. The threat model used for the security evaluation of TEAP is defined in the EAP [RFC3748].

7.1. Mutual Authentication and Integrity Protection

TEAP as a whole, provides message and integrity protection by establishing a secure tunnel for protecting the authentication method(s). The confidentiality and integrity protection is defined by TLS and provides the same security strengths afforded by TLS employing a strong entropy shared master secret. The integrity of the key generating authentication methods executed within the TEAP tunnel is verified through the calculation of the Crypto-Binding TLV. This ensures that the tunnel endpoints are the same as the inner method endpoints.

The Result TLV is protected and conveys the true Success or Failure of TEAP, and should be used as the indicator of its success or failure respectively. However, as EAP must terminate with a clear text EAP Success or Failure, a peer will also receive a clear text EAP Success or Failure. The received clear text EAP Success or Failure must match that received in the Result TLV; the peer SHOULD silently discard those clear text EAP success or failure messages that do not coincide with the status sent in the protected Result TLV.

7.2. Method Negotiation

As is true for any negotiated EAP protocol, NAK packets used to suggest an alternate authentication method are sent unprotected and as such, are subject to spoofing. During unprotected EAP method negotiation, NAK packets may be interjected as active attacks to negotiate down to a weaker form of authentication, such as EAP-MD5 (which only provides one-way authentication and does not derive a key). Both the peer and server should have a method selection policy that prevents them from negotiating down to weaker methods. Inner method negotiation resists attacks because it is protected by the mutually authenticated TLS tunnel established. Selection of TEAP as an authentication method does not limit the potential inner authentication methods, so TEAP should be selected when available.

An attacker cannot readily determine the inner EAP method used, except perhaps by traffic analysis. It is also important that peer implementations limit the use of credentials with an unauthenticated or unauthorized server.

7.3. Separation of Phase 1 and Phase 2 Servers

Separation of the TEAP Phase 1 from the Phase 2 conversation is NOT RECOMMENDED. Allowing the Phase 1 conversation to be terminated at a different server than the Phase 2 conversation can introduce vulnerabilities if there is not a proper trust relationship and

protection for the protocol between the two servers. Some vulnerabilities include:

- o Loss of identity protection
- o Offline dictionary attacks
- o Lack of policy enforcement
- o Man-in-the-middle attacks (as described in [I-D.hartman-emu-mutual-crypto-bind])

There may be cases where a trust relationship exists between the Phase 1 and Phase 2 servers, such as on a campus or between two offices within the same company, where there is no danger in revealing the inner identity and credentials of the peer to entities between the two servers. In these cases, using a proxy solution without end-to-end protection of TEAP MAY be used. The TEAP encrypting/decrypting gateway SHOULD, at a minimum, provide support for IPsec or similar protection in order to provide confidentiality for the portion of the conversation between the gateway and the EAP server. In addition, separation of the inner and outer method servers allows for crypto-binding based on the inner method MSK to be thwarted as described in [I-D.hartman-emu-mutual-crypto-bind]. Implentor and deployment SHOULD adopt various mitigation strategies described in [I-D.hartman-emu-mutual-crypto-bind]. If the inner method is deriving EMSK, then this threat is mitigated as TEAP utilizes the mutual crypto-binding based on EMSK as described in [I-D.hartman-emu-mutual-crypto-bind].

7.4. Mitigation of Known Vulnerabilities and Protocol Deficiencies

TEAP addresses the known deficiencies and weaknesses in the EAP method. By employing a shared secret between the peer and server to establish a secured tunnel, TEAP enables:

- o Per packet confidentiality and integrity protection
- o User identity protection
- o Better support for notification messages
- o Protected EAP inner method negotiation
- o Sequencing of EAP methods
- o Strong mutually derived master session keys

- o Acknowledged success/failure indication
- o Faster re-authentications through session resumption
- o Mitigation of dictionary attacks
- o Mitigation of man-in-the-middle attacks
- o Mitigation of some denial-of-service attacks

It should be noted that TEAP, as in many other authentication protocols, a denial-of-service attack can be mounted by adversaries sending erroneous traffic to disrupt the protocol. This is a problem in many authentication or key agreement protocols and is therefore noted for TEAP as well.

TEAP was designed with a focus on protected authentication methods that typically rely on weak credentials, such as password-based secrets. To that extent, the TEAP Authentication mitigates several vulnerabilities, such as dictionary attacks, by protecting the weak credential-based authentication method. The protection is based on strong cryptographic algorithms in TLS to provide message confidentiality and integrity. The keys derived for the protection relies on strong random challenges provided by both peer and server as well as an established key with strong entropy. Implementations should follow the recommendation in [RFC4086] when generating random numbers.

7.4.1. User Identity Protection and Verification

The initial identity request response exchange is sent in cleartext outside the protection of TEAP. Typically the Network Access Identifier (NAI) [RFC4282] in the identity response is useful only for the realm information that is used to route the authentication requests to the right EAP server. This means that the identity response may contain an anonymous identity and just contain realm information. In other cases, the identity exchange may be eliminated altogether if there are other means for establishing the destination realm of the request. In no case should an intermediary place any trust in the identity information in the identity response since it is unauthenticated and may not have any relevance to the authenticated identity. TEAP implementations should not attempt to compare any identity disclosed in the initial cleartext EAP Identity response packet with those Identities authenticated in Phase 2.

Identity request-response exchanges sent after the TEAP tunnel is established are protected from modification and eavesdropping by attackers.

Note that since TLS client certificates are sent in the clear, if identity protection is required, then it is possible for the TLS authentication to be re-negotiated after the first server authentication. To accomplish this, the server will typically not request a certificate in the `server_hello`, then after the `server_finished` message is sent, and before TEAP Phase 2, the server MAY send a TLS `hello_request`. This allows the client to perform client authentication by sending a `client_hello` if it wants to, or send a `no_renegotiation` alert to the server indicating that it wants to continue with TEAP Phase 2 instead. Assuming that the client permits renegotiation by sending a `client_hello`, then the server will respond with `server_hello`, a certificate and `certificate_request` messages. The client replies with `certificate`, `client_key_exchange` and `certificate_verify` messages. Since this re-negotiation occurs within the encrypted TLS channel, it does not reveal client certificate details. It is possible to perform certificate authentication using an EAP method (for example: EAP-TLS) within the TLS session in TEAP Phase 2 instead of using TLS handshake renegotiation.

7.4.2. Dictionary Attack Resistance

TEAP was designed with a focus on protected authentication methods that typically rely on weak credentials, such as password-based secrets. TEAP mitigates dictionary attacks by allowing the establishment of a mutually authenticated encrypted TLS tunnel providing confidentiality and integrity to protect the weak credential based authentication method.

7.4.3. Protection against Man-in-the-Middle Attacks

Allowing methods to be executed both with and without the protection of a secure tunnel opens up a possibility of a man-in-the-middle attack. To avoid man-in-the-middle attacks it is recommended to always deploy authentication methods with protection of TEAP. TEAP provides protection from man-in-the-middle attacks even if a deployment chooses to execute inner EAP methods both with and without TEAP protection, TEAP prevents this attack in two ways:

1. By using the PAC-Key to mutually authenticate the peer and server during TEAP Authentication Phase 1 establishment of a secure tunnel.
2. By using the keys generated by the inner authentication method (if the inner methods are key generating) in the crypto-binding exchange and in the generation of the key material exported by the EAP method described in Section 5.

7.4.4. PAC Binding to User Identity

A PAC may be bound to a user identity. A compliant implementation of TEAP MUST validate that an identity obtained in the PAC-Opaque field matches at minimum one of the identities provided in the TEAP Phase 2 authentication method. This validation provides another binding to ensure that the intended peer (based on identity) has successfully completed the TEAP Phase 1 and proved identity in the Phase 2 conversations.

7.5. Protecting against Forged Clear Text EAP Packets

EAP Success and EAP Failure packets are, in general, sent in clear text and may be forged by an attacker without detection. Forged EAP Failure packets can be used to attempt to convince an EAP peer to disconnect. Forged EAP Success packets may be used to attempt to convince a peer that authentication has succeeded, even though the authenticator has not authenticated itself to the peer.

By providing message confidentiality and integrity, TEAP provides protection against these attacks. Once the peer and AS initiate the TEAP Authentication Phase 2, compliant TEAP implementations must silently discard all clear text EAP messages, unless both the TEAP peer and server have indicated success or failure using a protected mechanism. Protected mechanisms include TLS alert mechanism and the protected termination mechanism described in Section 3.3.3.

The success/failure decisions within the TEAP tunnel indicate the final decision of the TEAP authentication conversation. After a success/failure result has been indicated by a protected mechanism, the TEAP peer can process unprotected EAP Success and EAP Failure messages; however the peer MUST ignore any unprotected EAP success or failure messages where the result does not match the result of the protected mechanism.

To abide by [RFC3748], the server must send a clear text EAP Success or EAP Failure packet to terminate the EAP conversation. However, since EAP Success and EAP Failure packets are not retransmitted, the final packet may be lost. While a TEAP protected EAP Success or EAP Failure packet should not be a final packet in a TEAP conversation, it may occur based on the conditions stated above, so an EAP peer should not rely upon the unprotected EAP success and failure messages.

7.6. Server Certificate Validation

As part of the TLS negotiation, the server presents a certificate to the peer. The peer MUST verify the validity of the EAP server

certificate, and SHOULD also examine the EAP server name presented in the certificate, in order to determine whether the EAP server can be trusted. When performing server certificate validation implementations MUST provide support rules in [RFC5280] for validating certificates against a known trust anchor. In addition, implementations SHOULD support matching the realm portion of the client's NAI against a SubjectAltName of type dNSName within the server certificate. Please note that in the case where the EAP authentication is remoted, the EAP server will not reside on the same machine as the authenticator, and therefore the name in the EAP server's certificate cannot be expected to match that of the intended destination. In this case, a more appropriate test might be whether the EAP server's certificate is signed by a CA controlling the intended domain and whether the authenticator can be authorized by a server in that domain.

7.7. Tunnel PAC Considerations

Since the Tunnel PAC is stored by the peer, special care should be given to the overall security of the peer. The Tunnel PAC must be securely stored by the peer to prevent theft or forgery of any of the Tunnel PAC components. In particular, the peer must securely store the PAC-Key and protect it from disclosure or modification. Disclosure of the PAC-Key enables an attacker to establish the TEAP tunnel; however, disclosure of the PAC-Key does not reveal the peer or server identity or compromise any other peer's PAC credentials. Modification of the PAC-Key or PAC-Opaque components of the Tunnel PAC may also lead to denial of service as the tunnel establishment will fail. The PAC-Opaque component is the effective TLS ticket extension used to establish the tunnel using the techniques of [RFC5077]. Thus, the security considerations defined by [RFC5077] also apply to the PAC- Opaque. The PAC-Info may contain information about the Tunnel PAC such as the identity of the PAC issuer and the Tunnel PAC lifetime for use in the management of the Tunnel PAC. The PAC-Info should be securely stored by the peer to protect it from disclosure and modification.

7.8. Security Claims

This section provides the needed security claim requirement for EAP [RFC3748].

Auth. mechanism:	Certificate based, shared secret based and various tunneled authentication mechanisms.
------------------	--

Ciphersuite negotiation: Yes

Mutual authentication: Yes

Integrity protection: Yes, Any method executed within the TEAP tunnel is integrity protected. The cleartext EAP headers outside the tunnel are not integrity protected.

Replay protection: Yes

Confidentiality: Yes

Key derivation: Yes

Key strength: See Note 1 below.

Dictionary attack prot.: Yes

Fast reconnect: Yes

Cryptographic binding: Yes

Session independence: Yes

Fragmentation: Yes

Key Hierarchy: Yes

Channel binding: Yes

Notes

1. BCP 86 [RFC3766] offers advice on appropriate key sizes. The National Institute for Standards and Technology (NIST) also offers advice on appropriate key sizes in [NIST-SP-800-57]. [RFC3766] Section 5 advises use of the following required RSA or DH module and DSA subgroup size in bits, for a given level of attack resistance in bits. Based on the table below, a 2048-bit RSA key is required to provide 128-bit equivalent key strength:

Attack Resistance (bits)	RSA or DH Modulus size (bits)	DSA subgroup size (bits)
-----	-----	-----
70	947	129
80	1228	148
90	1553	167
100	1926	186
150	4575	284
200	8719	383
250	14596	482

8. Acknowledgements

The TEAP v1 design and protocol specification is based on EAP-FAST [RFC4851], which included the ideas and hard efforts of Nancy Cam-Winget, David McGrew, Joe Salowey, Hao Zhou, Pad Jakkahalli, Mark Krischer, Doug Smith, and Glen Zorn of Cisco Systems, Inc.

The TLV processing was inspired from work on the Protected Extensible Authentication Protocol version 2 (PEAPv2) with Ashwin Palekar, Dan Smith, Sean Turner and Simon Josefsson.

Helpful review comments were provided by Russ Housley, Jari Arkko, Ilan Frenkel, Jeremy Steiglitz, Dan Harkins, Sam Hartman, and Jim Schaad.

9. References

9.1. Normative References

- | | |
|-----------------------|---|
| [I-D.ietf-emu-chbind] | Hartman, S., Clancy, T., and K. Hoeper, "Channel Binding Support for EAP Methods", draft-ietf-emu-chbind-15 (work in progress), May 2012. |
| [RFC2119] | Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997. |
| [RFC3748] | Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H. Levkowetz, "Extensible Authentication Protocol (EAP)", RFC 3748, June 2004. |
| [RFC4851] | Cam-Winget, N., McGrew, D., |

- Salowey, J., and H. Zhou, "The Flexible Authentication via Secure Tunneling Extensible Authentication Protocol Method (EAP-FAST)", RFC 4851, May 2007.
- [RFC5077] Salowey, J., Zhou, H., Eronen, P., and H. Tschofenig, "Transport Layer Security (TLS) Session Resumption without Server-Side State", RFC 5077, January 2008.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.
- [RFC5295] Salowey, J., Dondeti, L., Narayanan, V., and M. Nakhjiri, "Specification for the Derivation of Root Keys from an Extended Master Session Key (EMSK)", RFC 5295, August 2008.
- [RFC5705] Rescorla, E., "Keying Material Exporters for Transport Layer Security (TLS)", RFC 5705, March 2010.
- [RFC5746] Rescorla, E., Ray, M., Dispensa, S., and N. Oskov, "Transport Layer Security (TLS) Renegotiation Indication Extension", RFC 5746, February 2010.
- [RFC5929] Altman, J., Williams, N., and L. Zhu, "Channel Bindings for TLS", RFC 5929, July 2010.

9.2. Informative References

- [I-D.hartman-emu-mutual-crypto-bind] Hartman, S., Wasserman, M., and D. Zhang, "EAP Mutual Cryptographic Binding", draft-hartman-emu-mutual-crypto-bind-00 (work in progress), March 2012.
- [I-D.ietf-emu-eaptunnel-req] Zhou, H., Salowey, J., Hoepfer, K., and S. Hanna, "Requirements for a Tunnel Based EAP Method", draft-ietf-emu-eaptunnel-req-09 (work in progress), December 2010.
- [IEEE.802-1X.2004] "Local and Metropolitan Area Networks: Port-Based Network Access Control", IEEE Standard 802.1X, December 2004.
- [NIST-SP-800-57] National Institute of Standards and Technology, "Recommendation for Key Management", NIST Special Publication 800-57, May 2006.
- [PEAP] Microsoft Corporation, "[MS-PEAP]: Protected Extensible Authentication Protocol (PEAP) Specification", August 2009.
- [RFC2311] Dusse, S., Hoffman, P., Ramsdell, B., Lundblade, L., and L. Repka, "S/MIME Version 2 Message Specification", RFC 2311, March 1998.
- [RFC2315] Kaliski, B., "PKCS #7: Cryptographic Message Syntax Version 1.5", RFC 2315, March 1998.
- [RFC2560] Myers, M., Ankney, R., Malpani, A., Galperin, S., and C. Adams, "X.509 Internet Public Key Infrastructure Online Certificate Status Protocol -

- OCSP", RFC 2560, June 1999.
- [RFC2986] Nystrom, M. and B. Kaliski,
"PKCS #10: Certification
Request Syntax Specification
Version 1.7", RFC 2986,
November 2000.
- [RFC3579] Aboba, B. and P. Calhoun,
"RADIUS (Remote Authentication
Dial In User Service) Support
For Extensible Authentication
Protocol (EAP)", RFC 3579,
September 2003.
- [RFC3629] Yergeau, F., "UTF-8, a
transformation format of ISO
10646", STD 63, RFC 3629,
November 2003.
- [RFC3766] Orman, H. and P. Hoffman,
"Determining Strengths For
Public Keys Used For Exchanging
Symmetric Keys", BCP 86,
RFC 3766, April 2004.
- [RFC4072] Eronen, P., Hiller, T., and G.
Zorn, "Diameter Extensible
Authentication Protocol (EAP)
Application", RFC 4072,
August 2005.
- [RFC4086] Eastlake, D., Schiller, J., and
S. Crocker, "Randomness
Requirements for Security",
BCP 106, RFC 4086, June 2005.
- [RFC4282] Aboba, B., Beadles, M., Arkko,
J., and P. Eronen, "The Network
Access Identifier", RFC 4282,
December 2005.
- [RFC4945] Korver, B., "The Internet IP
Security PKI Profile of IKEv1/
ISAKMP, IKEv2, and PKIX",
RFC 4945, August 2007.
- [RFC5247] Aboba, B., Simon, D., and P.

- Eronen, "Extensible Authentication Protocol (EAP) Key Management Framework", RFC 5247, August 2008.
- [RFC5272] Schaad, J. and M. Myers, "Certificate Management over CMS (CMC)", RFC 5272, June 2008.
- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, May 2008.
- [RFC5281] Funk, P. and S. Blake-Wilson, "Extensible Authentication Protocol Tunneled Transport Layer Security Authenticated Protocol Version 0 (EAP-TLSv0)", RFC 5281, August 2008.
- [RFC5421] Cam-Winget, N. and H. Zhou, "Basic Password Exchange within the Flexible Authentication via Secure Tunneling Extensible Authentication Protocol (EAP-FAST)", RFC 5421, March 2009.
- [RFC5931] Harkins, D. and G. Zorn, "Extensible Authentication Protocol (EAP) Authentication Using Only a Password", RFC 5931, August 2010.
- [RFC6066] Eastlake, D., "Transport Layer Security (TLS) Extensions: Extension Definitions", RFC 6066, January 2011.
- [RFC6124] Sheffer, Y., Zorn, G., Tschofenig, H., and S. Fluhrer, "An EAP Authentication Method

Based on the Encrypted Key
Exchange (EKE) Protocol",
RFC 6124, February 2011.

Appendix A. Evaluation Against Tunnel Based EAP Method Requirements

This section evaluates all tunnel based EAP method requirements described in [I-D.ietf-emu-eaptunnel-req] against TEAP version 1.

A.1. Requirement 4.1.1 RFC Compliance

TEAP v1 meets this requirement by being compliant to RFC 3748, RFC 4017, RFC 5247, and RFC 4962. It is also compliant with the "cryptographic algorithm agility" requirement by leveraging TLS 1.2 for all cryptographic algorithm negotiation.

A.2. Requirement 4.2.1 TLS Requirements

Requirement 4.2.1 states:

The tunnel based method MUST support TLS version 1.2 [RFC5246] and may support earlier versions greater than SSL 2.0 to enable the possibility of backwards compatibility.

TEAP v1 meets this requirement by mandating TLS version 1.2 support as defined in Section 3.2.

A.3. Requirement 4.2.1.1.1 Cipher Suite Negotiation

Requirement 4.2.1.1.1 states:

Hence, the tunnel method MUST provide integrity protected cipher suite negotiation with secure integrity algorithms and integrity keys.

TEAP v1 meets this requirement by using TLS to provide protected cipher suite negotiation.

A.4. Requirement 4.2.1.1.2 Tunnel Data Protection Algorithms

Requirement 4.2.1.1.2 states:

The tunnel method MUST provide at least one mandatory to implement cipher suite that provides the equivalent security of 128-bit AES for encryption and message authentication.

TEAP v1 meets this requirement by mandating TLS_RSA_WITH_AES_128_CBC_SHA as a mandatory to implement cipher suite

as defined in Section 3.2.

A.5. Requirement 4.2.1.1.3 Tunnel Authentication and Key Establishment

TEAP v1 meets this requirement by mandating TLS_RSA_WITH_AES_128_CBC_SHA as a mandatory to implement cipher suite which provides certificate-based authentication of the server and is approved by NIST. The mandatory to implement cipher suites only include cipher suites that use strong cryptographic algorithms. They do not include cipher suites providing mutually anonymous authentication or static Diffie-Hellman cipher suites as defined in Section 3.2.

A.6. Requirement 4.2.1.2 Tunnel Replay Protection

TEAP v1 meets this requirement by using TLS to provide sufficient replay protection.

A.7. Requirement 4.2.1.3 TLS Extensions

TEAP v1 meets this requirement by allowing TLS extensions, such as TLS Certificate Status Request extension [RFC6066] and SessionTicket extension [RFC5077] to be used during TLS tunnel establishment.

A.8. Requirement 4.2.1.4 Peer Identity Privacy

TEAP v1 meets this requirement by establishment of the TLS tunnel and protection of inner method specific identities. In addition, the peer certificate can be sent confidentially (i.e. encrypted).

A.9. Requirement 4.2.1.5 Session Resumption

TEAP v1 meets this requirement by mandating support of TLS session resumption as defined in Section 3.2.1 and TLS Session Resume Using a PAC as defined in Section 3.2.2 .

A.10. Requirement 4.2.2 Fragmentation

TEAP v1 meets this requirement by leveraging fragmentation support provided by TLS as defined in Section 3.7.

A.11. Requirement 4.2.3 Protection of Data External to Tunnel

TEAP v1 meets this requirement by including TEAP version number received in the computation of crypto-binding TLV as defined in Section 4.2.13.

A.12. Requirement 4.3.1 Extensible Attribute Types

TEAP v1 meets this requirement by using an extensible TLV data layer inside the tunnel as defined in Section 4.2.

A.13. Requirement 4.3.2 Request/Challenge Response Operation

TEAP v1 meets this requirement by allowing multiple TLVs to be sent in a single EAP request or response packet, while maintaining the half-duplex operation typical of EAP.

A.14. Requirement 4.3.3 Indicating Criticality of Attributes

TEAP v1 meets this requirement by having a mandatory bit in TLV to indicate whether it is mandatory to support or not as defined in Section 4.2.

A.15. Requirement 4.3.4 Vendor Specific Support

TEAP v1 meets this requirement by having a Vendor-Specific TLV to allow vendors to define their own attributes as defined in Section 4.2.8.

A.16. Requirement 4.3.5 Result Indication

TEAP v1 meets this requirement by having a Result TLV to exchange the final result of the EAP authentication so both the peer and server have a synchronized state as defined in Section 4.2.4.

A.17. Requirement 4.3.6 Internationalization of Display Strings

TEAP v1 meets this requirement by supporting UTF-8 format in Basic-Password-Auth-Req TLV as defined in Section 4.2.14 and Basic-Password-Auth-Resp TLV as defined in Section 4.2.15.

A.18. Requirement 4.4 EAP Channel Binding Requirements

TEAP v1 meets this requirement by having a Channel-Binding TLV to exchange the EAP channel binding data as defined in Section 4.2.7.

A.19. Requirement 4.5.1.1 Confidentiality and Integrity

TEAP v1 meets this requirement by running the password authentication inside a protected TLS tunnel.

A.20. Requirement 4.5.1.2 Authentication of Server

TEAP v1 meets this requirement by mandating authentication of the server before establishment of the protected TLS and then running inner password authentication as defined in Section 3.2.

A.21. Requirement 4.5.1.3 Server Certificate Revocation Checking

TEAP v1 meets this requirement by supporting TLS Certificate Status Request extension [RFC6066] during tunnel establishment.

A.22. Requirement 4.5.2 Internationalization

TEAP v1 meets this requirement by supporting UTF-8 format in Basic-Password-Auth-Req TLV as defined in Section 4.2.14 and Basic-Password-Auth-Resp TLV as defined in Section 4.2.15.

A.23. Requirement 4.5.3 Meta-data

TEAP v1 meets this requirement by supporting Identity-Type TLV as defined in Section 4.2.3 to indicate whether the authentication is for a user or a machine.

A.24. Requirement 4.5.4 Password Change

TEAP v1 meets this requirement by supporting multiple Basic-Password-Auth-Req TLV and Basic-Password-Auth-Resp TLV exchanges within a single EAP authentication, which allows "housekeeping" functions such as password change.

A.25. Requirement 4.6.1 Method Negotiation

TEAP v1 meets this requirement by supporting inner EAP method negotiation within the protected TLS tunnel.

A.26. Requirement 4.6.2 Chained Methods

TEAP v1 meets this requirement by supporting inner EAP method chaining within protected TLS tunnel as defined in Section 3.3.1.

A.27. Requirement 4.6.3 Cryptographic Binding with the TLS Tunnel

TEAP v1 meets this requirement by supporting cryptographic binding of the inner EAP method keys with the keys derived from the TLS tunnel as defined in Section 4.2.13.

A.28. Requirement 4.6.4 Peer Initiated

TEAP v1 meets this requirement by supporting Request-Action TLV as defined in Section 4.2.9 to allow peer to initiate another inner EAP method.

A.29. Requirement 4.6.5 Method Meta-data

TEAP v1 meets this requirement by supporting Identity-Type TLV as defined in Section 4.2.3 to indicate whether the authentication is for a user or a machine.

Appendix B. Major Differences from EAP-FAST

This document is a new standard tunnel EAP method based on revision of the EAP-FAST version 1 [RFC4851] which contains improved flexibility, particularly for negotiation of cryptographic algorithms. The major changes are:

1. The EAP method name have been changed from EAP-FAST to TEAP, hence it would require a new EAP method type to be assigned.
2. This version of TEAP MUST support TLS 1.2 [RFC5246].
3. The key derivation now makes use of TLS keying material exporters [RFC5705] and the PRF and hash function negotiated in TLS. This is to simplify implementation and better support cryptographic algorithm agility.
4. TEAP is in full conformance with TLS Ticket extension [RFC5077] as described in Section 3.2.2.
5. Support of passing optional outer TLVs in the first two message exchanges, in addition to the Authority-ID TLV data in EAP-FAST.
6. Basic password authentication on the TLV level has been added in addition to the existing inner EAP method.
7. Additional TLV types have been defined to support EAP channel binding and meta-data. They are Identity-Type TLV and Channel-Binding TLVs, defined in Section 4.2.

Appendix C. Examples

C.1. Successful Authentication

The following exchanges show a successful TEAP authentication with basic password authentication and optional PAC refreshment, the

conversation will appear as follows:

```
Authenticating Peer      Authenticator
-----
EAP-Response/
Identity (MyID1) ->

                                <- EAP-Request/
                                Identity

                                <- EAP-Request/
                                EAP-Type=TEAP, V=1
                                (TEAP Start, S bit set, Authority-ID)

EAP-Response/
EAP-Type=TEAP, V=1
(TLS client_hello with
 PAC-Opaque in SessionTicket extension)->

                                <- EAP-Request/
                                EAP-Type=TEAP, V=1
                                (TLS server_hello,
                                (TLS change_cipher_spec,
                                TLS finished)

EAP-Response/
EAP-Type=TEAP, V=1 ->
(TLS change_cipher_spec,
 TLS finished)

TLS channel established
(messages sent within the TLS channel)

                                <- Basic-Password-Auth-Req TLV, Challenge

Basic-Password-Auth-Resp TLV, Response with both
user name and password) ->

optional additional exchanges (new pin mode,
password change etc.) ...

                                <- Crypto-Binding TLV (Request),
                                Result TLV (Success),
                                (Optional PAC TLV)
```

```
Crypto-Binding TLV(Response),  
Result TLV (Success),  
(PAC TLV Acknowledgment) ->
```

```
TLS channel torn down  
(messages sent in clear text)
```

```
<- EAP-Success
```

C.2. Failed Authentication

The following exchanges show a failed TEAP authentication due to wrong user credentials, the conversation will appear as follows:

```
Authenticating Peer      Authenticator
-----
<- EAP-Request/
Identity

EAP-Response/
Identity (MyID1) ->

<- EAP-Request/
EAP-Type=TEAP, V=1
(TEAP Start, S bit set, Authority-ID)

EAP-Response/
EAP-Type=TEAP, V=1
(TLS client_hello with
PAC-Opaque in SessionTicket extension)->

<- EAP-Request/
EAP-Type=TEAP, V=1
(TLS server_hello,
(TLS change_cipher_spec,
TLS finished)

EAP-Response/
EAP-Type=TEAP, V=1 ->
(TLS change_cipher_spec,
TLS finished)

TLS channel established
(messages sent within the TLS channel)

<- Basic-Password-Auth-Req TLV, Challenge

Basic-Password-Auth-Resp TLV, Response with both
user name and password) ->

<- Result TLV (Failure)

Result TLV (Failure) ->

TLS channel torn down
(messages sent in clear text)

<- EAP-Failure
```

C.3. Full TLS Handshake using Certificate-based Cipher Suite

In the case where an abbreviated TLS handshake is tried and failed and falls back to certificate based full TLS handshake occurs within TEAP Phase 1, the conversation will appear as follows:

Authenticating Peer -----	Authenticator -----
	<- EAP-Request/Identity
EAP-Response/ Identity (MyID1) ->	
<pre>// Identity sent in the clear. May be a hint to help route the authentication request to EAP server, instead of the full user identity.</pre>	
	<pre><- EAP-Request/ EAP-Type=TEAP, V=1 (TEAP Start, S bit set, Authority-ID)</pre>
EAP-Response/ EAP-Type=TEAP, V=1 (TLS client_hello [PAC-Opaque extension])->	
<pre>// Peer sends PAC-Opaque of Tunnel PAC along with a list of ciphersuites supported. If the server rejects the PAC- Opaque, it falls through to the full TLS handshake</pre>	
	<pre><- EAP-Request/ EAP-Type=TEAP, V=1 (TLS server_hello, TLS certificate, [TLS server_key_exchange,] [TLS certificate_request,] TLS server_hello_done)</pre>
EAP-Response/ EAP-Type=TEAP, V=1 ([TLS certificate,] TLS client_key_exchange, [TLS certificate_verify,] TLS change_cipher_spec, TLS finished) ->	
	<pre><- EAP-Request/ EAP-Type=TEAP, V=1 (TLS change_cipher_spec, TLS finished, EAP-Payload-TLV[EAP-Request/ Identity])</pre>

```

// TLS channel established
  (messages sent within the TLS channel)

// First EAP Payload TLV is piggybacked to the TLS Finished as
  Application Data and protected by the TLS tunnel

EAP-Payload-TLV
[EAP-Response/Identity (MyID2)]->

// identity protected by TLS.

                                <- EAP-Payload-TLV
                                [EAP-Request/EAP-Type=X]

EAP-Payload-TLV
[EAP-Response/EAP-Type=X] ->

// Method X exchanges followed by Protected Termination

                                <- Intermediate-Result-TLV (Success),
                                Crypto-Binding TLV (Request),
                                Result TLV (Success)

Intermediate-Result-TLV (Success),
Crypto-Binding TLV (Response),
Result-TLV (Success) ->

// TLS channel torn down
(messages sent in clear text)

                                <- EAP-Success

```

C.4. Client authentication during Phase 1 with identity privacy

In the case where a certificate based TLS handshake occurs within TEAP Phase 1, and client certificate authentication and identity privacy is desired, therefore TLS renegotiation is being used to transmit the peer credentials in the protected TLS tunnel, the conversation will appear as follows:

Authenticating Peer	Authenticator
-----	-----
	<- EAP-Request/Identity
EAP-Response/ Identity (MyID1) ->	
// Identity sent in the clear. May be a hint to help route the authentication request to EAP server, instead of the	

full user identity.

```

                                <- EAP-Request/
                                EAP-Type=TEAP, V=1
                                (TEAP Start, S bit set, Authority-ID)
EAP-Response/
EAP-Type=TEAP, V=1
(TLS client_hello)->
                                <- EAP-Request/
                                EAP-Type=TEAP, V=1
                                (TLS server_hello,
                                 TLS certificate,
                                 [TLS server_key_exchange,]
                                 [TLS certificate_request,]
                                 TLS server_hello_done)
EAP-Response/
EAP-Type=TEAP, V=1
(TLS client_key_exchange,
 TLS change_cipher_spec,
 TLS finished) ->
                                <- EAP-Request/
                                EAP-Type=TEAP, V=1
                                (TLS change_cipher_spec,
                                 TLS finished,
                                 EAP-Payload-TLV[EAP-Request/
                                 Identity])

// TLS channel established
// (EAP Payload messages sent within the TLS channel)

// peer sends TLS client_hello to request TLS renegotiation

TLS client_hello ->
                                <- TLS server_hello,
                                TLS certificate,
                                [TLS server_key_exchange,]
                                [TLS certificate_request,]
                                TLS server_hello_done
[TLS certificate,]
TLS client_key_exchange,
[TLS certificate_verify,]
TLS change_cipher_spec,
TLS finished ->
                                <- TLS change_cipher_spec,
                                TLS finished,

```



```
Crypto-Binding TLV (Request),
Result TLV (Success)
```

```
Crypto-Binding TLV (Response),
Result-TLV (Success)) ->
```

```
//TLS channel torn down
(messages sent in clear text)
```

```
<- EAP-Success
```

C.5. Fragmentation and Reassembly

In the case where TEAP fragmentation is required, the conversation will appear as follows:

Authenticating Peer -----	Authenticator -----
	<- EAP-Request/ Identity
EAP-Response/ Identity (MyID) ->	
	<- EAP-Request/ EAP-Type=TEAP, V=1 (TEAP Start, S bit set, Authority-ID)
EAP-Response/ EAP-Type=TEAP, V=1 (TLS client_hello)->	
	<- EAP-Request/ EAP-Type=TEAP, V=1 (TLS server_hello, TLS certificate, [TLS server_key_exchange,] [TLS certificate_request,] TLS server_hello_done) (Fragment 1: L, M bits set)
EAP-Response/ EAP-Type=TEAP, V=1 ->	
	<- EAP-Request/ EAP-Type=TEAP, V=1 (Fragment 2: M bit set)
EAP-Response/ EAP-Type=TEAP, V=1 ->	
	<- EAP-Request/

```

                                EAP-Type=TEAP, V=1
                                (Fragment 3)

EAP-Response/
EAP-Type=TEAP, V=1
([TLS certificate,]
 TLS client_key_exchange,
[TLS certificate_verify,]
 TLS change_cipher_spec,
 TLS finished)
(Fragment 1: L, M bits set)->

                                <- EAP-Request/
                                EAP-Type=TEAP, V=1

EAP-Response/
EAP-Type=TEAP, V=1
(Fragment 2)->

                                <- EAP-Request/
                                EAP-Type=TEAP, V=1
                                (TLS change_cipher_spec,
                                 TLS finished,
                                 [EAP-Payload-TLV[
                                 EAP-Request/Identity]])

// TLS channel established
// (messages sent within the TLS channel)

// First EAP Payload TLV is piggybacked to the TLS Finished as
// Application Data and protected by the TLS tunnel

EAP-Payload-TLV
[EAP-Response/Identity (MyID2)]->

// identity protected by TLS.

                                <- EAP-Payload-TLV
                                [EAP-Request/EAP-Type=X]

EAP-Payload-TLV
[EAP-Response/EAP-Type=X] ->

// Method X exchanges followed by Protected Termination

                                <- Intermediate-Result-TLV (Success),
                                Crypto-Binding TLV (Request),
                                Result TLV (Success)

Intermediate-Result-TLV (Success),
Crypto-Binding TLV (Response),

```

```

Result-TLV (Success) ->

// TLS channel torn down
(messages sent in clear text)

<- EAP-Success

```

C.6. Sequence of EAP Methods

When TEAP is negotiated, with a sequence of EAP method X followed by method Y, the conversation will occur as follows:

Authenticating Peer -----	Authenticator -----
	<- EAP-Request/ Identity
EAP-Response/ Identity (MyID1) ->	
	<- EAP-Request/ EAP-Type=TEAP, V=1 (TEAP Start, S bit set, Authority-ID)
EAP-Response/ EAP-Type=TEAP, V=1 (TLS client_hello)->	
	<- EAP-Request/ EAP-Type=TEAP, V=1 (TLS server_hello, TLS certificate, [TLS server_key_exchange,] [TLS certificate_request,] TLS server_hello_done)
EAP-Response/ EAP-Type=TEAP, V=1 ([TLS certificate,] TLS client_key_exchange, [TLS certificate_verify,] TLS change_cipher_spec, TLS finished) ->	
	<- EAP-Request/ EAP-Type=TEAP, V=1 (TLS change_cipher_spec, TLS finished, Identity-Type TLV, EAP-Payload-TLV[EAP-Request/Identity])

```
// TLS channel established
  (messages sent within the TLS channel)

// First EAP Payload TLV is piggybacked to the TLS Finished as
  Application Data and protected by the TLS tunnel

Identity_Type TLV
EAP-Payload-TLV
[EAP-Response/Identity] ->

      <- EAP-Payload-TLV
      [EAP-Request/EAP-Type=X]

EAP-Payload-TLV
[EAP-Response/EAP-Type=X] ->

    // Optional additional X Method exchanges...

      <- EAP-Payload-TLV
      [EAP-Request/EAP-Type=X]

EAP-Payload-TLV
[EAP-Response/EAP-Type=X]->

      <- Intermediate Result TLV (Success),
      Crypto-Binding TLV (Request),
      Identity-Type TLV,
      EAP Payload TLV [EAP-Type=Y],

// Next EAP conversation started after successful completion
  of previous method X. The Intermediate-Result and Crypto-
  Binding TLVs are sent in next packet to minimize round-
  trips. In this example, identity request is not sent
  before negotiating EAP-Type=Y.

// Compound MAC calculated using Keys generated from
  EAP methods X and the TLS tunnel.

Intermediate Result TLV (Success),
Crypto-Binding TLV (Response),
EAP-Payload-TLV [EAP-Type=Y] ->

    // Optional additional Y Method exchanges...

      <- EAP Payload TLV [
      EAP-Type=Y]

EAP Payload TLV
```

```

[EAP-Type=Y] ->

                                <- Intermediate-Result-TLV (Success),
                                Crypto-Binding TLV (Request),
                                Result TLV (Success)

Intermediate-Result-TLV (Success),
Crypto-Binding TLV (Response),
Result-TLV (Success) ->

// Compound MAC calculated using Keys generated from EAP
// methods X and Y and the TLS tunnel. Compound Keys
// generated using Keys generated from EAP methods X and Y;
// and the TLS tunnel.

// TLS channel torn down (messages sent in clear text)

                                <- EAP-Success

```

C.7. Failed Crypto-binding

The following exchanges show a failed crypto-binding validation. The conversation will appear as follows:

Authenticating Peer	Authenticator
-----	-----
	<- EAP-Request/ Identity
EAP-Response/ Identity (MyID1) ->	
	<- EAP-Request/ EAP-Type=TEAP, V=1 (TEAP Start, S bit set, Authority-ID)
EAP-Response/ EAP-Type=TEAP, V=1 (TLS client_hello without PAC-Opaque extension)->	
	<- EAP-Request/ EAP-Type=TEAP, V=1 (TLS Server Key Exchange TLS Server Hello Done)
EAP-Response/ EAP-Type=TEAP, V=1 -> (TLS Client Key Exchange TLS change_cipher_spec, TLS finished)	

```

        <- EAP-Request/
        EAP-Type=TEAP, V=1
        (TLS change_cipher_spec
         TLS finished)
        EAP-Payload-TLV[
        EAP-Request/Identity])

// TLS channel established
// (messages sent within the TLS channel)

// First EAP Payload TLV is piggybacked to the TLS Finished as
// Application Data and protected by the TLS tunnel

EAP-Payload TLV/
EAP Identity Response ->

        <- EAP Payload TLV, EAP-Request,
        (EAP-MSCHAPV2, Challenge)

EAP Payload TLV, EAP-Response,
(EAP-MSCHAPV2, Response) ->

        <- EAP Payload TLV, EAP-Request,
        (EAP-MSCHAPV2, Success Request)

EAP Payload TLV, EAP-Response,
(EAP-MSCHAPV2, Success Response) ->

        <- Intermediate-Result-TLV (Success),
        Crypto-Binding TLV (Request),
        Result TLV (Success)

Intermediate-Result-TLV (Success),
Result TLV (Failure)
Error TLV with
(Error Code = 2001) ->

// TLS channel torn down
// (messages sent in clear text)

        <- EAP-Failure

```

C.8. Sequence of EAP Method with Vendor-Specific TLV Exchange

When TEAP is negotiated, with a sequence of EAP method followed by Vendor-Specific TLV exchange, the conversation will occur as follows:

Authenticating Peer	Authenticator
---------------------	---------------

```

-----
EAP-Response/
Identity (MyID1) ->
    <- EAP-Request/
    Identity

EAP-Response/
EAP-Type=TEAP, V=1
(TLS client_hello)->
    <- EAP-Request/
    EAP-Type=TEAP, V=1
    (TEAP Start, S bit set, Authority-ID)

EAP-Response/
EAP-Type=TEAP, V=1
(TLS client_hello)->
    <- EAP-Request/
    EAP-Type=TEAP, V=1
    (TLS server_hello,
     TLS certificate,
     [TLS server_key_exchange,]
     [TLS certificate_request,]
     TLS server_hello_done)

EAP-Response/
EAP-Type=TEAP, V=1
([TLS certificate,]
 TLS client_key_exchange,
 [TLS certificate_verify,]
 TLS change_cipher_spec,
 TLS finished) ->
    <- EAP-Request/
    EAP-Type=TEAP, V=1
    (TLS change_cipher_spec,
     TLS finished,
     EAP-Payload-TLV[
     EAP-Request/Identity])

// TLS channel established
// (messages sent within the TLS channel)

// First EAP Payload TLV is piggybacked to the TLS Finished as
// Application Data and protected by the TLS tunnel

EAP-Payload-TLV
[EAP-Response/Identity] ->
    <- EAP-Payload-TLV
    [EAP-Request/EAP-Type=X]

EAP-Payload-TLV
[EAP-Response/EAP-Type=X] ->

```

```

                                <- EAP-Payload-TLV
                                [EAP-Request/EAP-Type=X]

EAP-Payload-TLV
[EAP-Response/EAP-Type=X]->

                                <- Intermediate Result TLV (Success),
                                Crypto-Binding TLV (Request),
                                Vendor-Specific TLV,

// Vendor Specific TLV exchange started after successful
// completion of previous method X. The Intermediate-Result
// and Crypto-Binding TLVs are sent with Vendor Specific TLV
// in next packet to minimize round-trips.

// Compound MAC calculated using Keys generated from
// EAP methods X and the TLS tunnel.

Intermediate Result TLV (Success),
Crypto-Binding TLV (Response),
Vendor-Specific TLV ->

    // Optional additional Vendor-Specific TLV exchanges...

                                <- Vendor-Specific TLV

Vendor Specific TLV ->
                                <- Result TLV (Success)

Result-TLV (Success) ->

// TLS channel torn down (messages sent in clear text)

                                <- EAP-Success

```

C.9. Peer Requests Inner Method After Server Sends Result TLV

In the case where the peer is authenticated during Phase 1 and server sends back result TLV, but the peer wants to request another inner method, the conversation will appear as follows:

Authenticating Peer	Authenticator
-----	-----
	<- EAP-Request/Identity
EAP-Response/ Identity (MyID1) ->	
// Identity sent in the clear. May be a hint to help route	

the authentication request to EAP server, instead of the full user identity.

```

                                <- EAP-Request/
                                EAP-Type=TEAP, V=1
                                (TEAP Start, S bit set, Authority-ID)
EAP-Response/
EAP-Type=TEAP, V=1
(TLS client_hello)->
                                <- EAP-Request/
                                EAP-Type=TEAP, V=1
                                (TLS server_hello,
                                 TLS certificate,
                                 [TLS server_key_exchange,]
                                 [TLS certificate_request,]
                                 TLS server_hello_done)
EAP-Response/
EAP-Type=TEAP, V=1
[TLS certificate,]
  TLS client_key_exchange,
[TLS certificate_verify,]
  TLS change_cipher_spec,
  TLS finished ->
                                <- EAP-Request/
                                EAP-Type=TEAP, V=1
                                (TLS change_cipher_spec,
                                 TLS finished,
                                 Crypto-Binding TLV (Request),
                                 Result TLV (Success))

// TLS channel established
  (TLV Payload messages sent within the TLS channel)

Crypto-Binding TLV(Response),
Request-Action TLV
(Status=Failure, Action=Negotiate-EAP)->
                                <- EAP-Payload-TLV
                                [EAP-Request/Identity]

EAP-Payload-TLV
[EAP-Response/Identity] ->
                                <- EAP-Payload-TLV
                                [EAP-Request/EAP-Type=X]

EAP-Payload-TLV
[EAP-Response/EAP-Type=X] ->
```

```

                                <- EAP-Payload-TLV
                                [EAP-Request/EAP-Type=X]

EAP-Payload-TLV
[EAP-Response/EAP-Type=X]->

                                <- Intermediate Result TLV (Success),
                                Crypto-Binding TLV (Request),
                                Result TLV (Success)

Intermediate Result TLV (Success),
Crypto-Binding TLV (Response),
Result-TLV (Success)) ->

//TLS channel torn down
(messages sent in clear text)

                                <- EAP-Success

```

C.10. Channel Binding

The following exchanges show a successful TEAP authentication with basic password authentication and channel binding using Request-Action TLV, the conversation will appear as follows:

Authenticating Peer	Authenticator
-----	-----
	<- EAP-Request/ Identity
EAP-Response/ Identity (MyID1) ->	
	<- EAP-Request/ EAP-Type=TEAP, V=1 (TEAP Start, S bit set, Authority-ID)
EAP-Response/ EAP-Type=TEAP, V=1 (TLS client_hello with PAC-Opaque in SessionTicket extension)->	
	<- EAP-Request/ EAP-Type=TEAP, V=1 (TLS server_hello, (TLS change_cipher_spec, TLS finished)
EAP-Response/	

```
EAP-Type=TEAP, V=1 ->
(TLS change_cipher_spec,
 TLS finished)

TLS channel established
(messages sent within the TLS channel)

        <- Basic-Password-Auth-Req TLV, Challenge

Basic-Password-Auth-Resp TLV, Response with both
user name and password) ->

optional additional exchanges (new pin mode,
password change etc.) ...

        <- Crypto-Binding TLV (Request),
        Result TLV (Success),

Crypto-Binding TLV(Response),
Request-Action TLV
(Status=Failure, Action=Process-TLV,
TLV=Channel-Binding TLV)->

        <- Channel-Binding TLV (Response),
        Result TLV (Success),

Result-TLV (Success) ->

TLS channel torn down
(messages sent in clear text)

        <- EAP-Success
```

Appendix D. Major Differences from Previous Revisions

D.1. Changes from -02

- 1 Section 3.3.3, clarified protected termination and use of crypto-binding TLV.
- 2 Section 3.5, changed Session ID to use tls-unique and added reference to RFC5247.
- 3 Section 3.9, added the use of tls-unique to the certificate enrollment request.

- 4 Section 4.2.9, modified Request-Action TLV to include Status code and optional TLVs.
- 5 Section 3.4, clarified that all authenticated Peer-Ids need to be exported.
- 6 Section 5.1, changed TLS Keying Material Exporter label to "teap session key seed".
- 7 Section 5.2, changed Intermediate Compound Key Derivation from MSK to EMSK generated by inner method.
- 8 Section 6, added missing IANA considerations.
- 9 Section 7.3, added more security considerations for separation of Phase 1 and Phase 2 servers.
- 10 Appendix C, updated examples with Request-Action TLV, channel binding, and sending certificate after TLS renegotiation.

D.2. Changes from -01

- 1 In Version Negotiation section, clarified what the peer needs to do if the supported version is higher than what the server proposed.
- 2 Section 3.2, clarified the requirement for using anonymous cipher suites.
- 3 Clarified that Crypto-binding TLV is always exchanged and validated, even without inner methods.
- 4 Section 3.4, clarified that all authenticated Peer-Ids need to be exported.
- 5 Clarified that channel-binding TLV can be used to transmit data bidirectionally.
- 6 Updated obsolete RFC references
- 7 Renumbered TLVs to eliminate gaps
- 8 Updated examples with basic password authentication TLVs.
- 9 Added Certificate Provisioning Within the Tunnel.

10 Added Server Unauthenticated Provisioning Mode.

D.3. Changes from -00

- 1 Changed protocol name to TEAP: Tunnel EAP Method
- 2 Changed version of protocol to version 1
- 3 Revised introduction
- 4 Moved differences section to appendix
- 5 Revised design goals section
- 6 Revised PAC definition
- 7 Revised protocol description to be in line with RFC 5077 PAC distribution
- 8 Revised EAP Sequences Section
- 9 Added section on PAC provisioning within tunnel
- 10 Added outer TLVs to the message format
- 11 Renumbered TLVs
- 12 Included PAC TLVs
- 13 Added Authority ID TLV
- 14 Added PKCS#7 and server trust root TLV definitions
- 15 Added PKCS#10 TLV
- 16 PKCS#10 TLV
- 17 Added EAP-Type and outer TLVs to crypto binding compound MAC

Authors' Addresses

Hao Zhou
Cisco Systems
4125 Highlander Parkway
Richfield, OH 44286
US

EMail: hzhou@cisco.com

Nancy Cam-Winget
Cisco Systems
3625 Cisco Way
San Jose, CA 95134
US

EMail: ncamwing@cisco.com

Joseph Salowey
Cisco Systems
2901 3rd Ave
Seattle, WA 98121
US

EMail: jsalowey@cisco.com

Stephen Hanna
Juniper Networks
79 Parsons Street
Brighton, MA 02135
US

EMail: shanna@juniper.net

EMU Working Group
Internet-Draft
Intended status: Informational
Expires: June 19, 2011

K. Hoeper
Motorola, Inc.
S. Hanna
Juniper Networks
H. Zhou
J. Salowey, Ed.
Cisco Systems, Inc.
December 16, 2010

Requirements for a Tunnel Based EAP Method
draft-ietf-emu-eaptunnel-req-09.txt

Abstract

This memo defines the requirements for a tunnel-based Extensible Authentication Protocol (EAP) Method. This tunnel method will use Transport Layer Security (TLS) to establish a secure tunnel. The tunnel will provide support for password authentication, EAP authentication and the transport of additional data for other purposes.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on June 19, 2011.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents

carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1. Introduction	5
2. Conventions Used In This Document	5
3. Use Cases	6
3.1. Password Authentication	6
3.2. Protection of Weak EAP Methods	6
3.3. Chained EAP Methods	7
3.4. Identity Protection	7
3.5. Anonymous Service Access	7
3.6. Network Endpoint Assessment	8
3.7. Client Authentication During Tunnel Establishment	8
3.8. Extensibility	8
3.9. Certificate-less Authentication and Generic EAP Method Extension	9
4. Requirements	9
4.1. General Requirements	9
4.1.1. RFC Compliance	9
4.2. Tunnel Requirements	10
4.2.1. TLS Requirements	10
4.2.1.1. Cipher Suites	10
4.2.1.1.1. Cipher Suite Negotiation	10
4.2.1.1.2. Tunnel Data Protection Algorithms	11
4.2.1.1.3. Tunnel Authentication and Key Establishment	11
4.2.1.2. Tunnel Replay Protection	12
4.2.1.3. TLS Extensions	12
4.2.1.4. Peer Identity Privacy	12
4.2.1.5. Session Resumption	12
4.2.2. Fragmentation	12
4.2.3. Protection of Data External to Tunnel	13
4.3. Tunnel Payload Requirements	13
4.3.1. Extensible Attribute Types	13
4.3.2. Request/Challenge Response Operation	13
4.3.3. Indicating Criticality of Attributes	13
4.3.4. Vendor Specific Support	13
4.3.5. Result Indication	14
4.3.6. Internationalization of Display Strings	14
4.4. EAP Channel Binding Requirements	14
4.5. Requirements Associated with Carrying Username and Passwords	14
4.5.1. Security	15
4.5.1.1. Confidentiality and Integrity	15
4.5.1.2. Authentication of Server	15
4.5.1.3. Server Certificate Revocation Checking	15
4.5.2. Internationalization	15

4.5.3. Meta-data	16
4.5.4. Password Change	16
4.6. Requirements Associated with Carrying EAP Methods	16
4.6.1. Method Negotiation	16
4.6.2. Chained Methods	16
4.6.3. Cryptographic Binding with the TLS Tunnel	16
4.6.4. Peer Initiated	18
4.6.5. Method Meta-data	18
5. IANA Considerations	18
6. Security Considerations	18
6.1. Cipher Suite Selection	18
6.2. Tunneled Authentication	19
6.3. Data External to Tunnel	20
6.4. Separation of TLS Tunnel and Inner Authentication Termination	20
7. References	20
7.1. Normative References	20
7.2. Informative References	21
Appendix A. Changes from -01	23
Appendix B. Changes from -02	23
Appendix C. changes from -03	23
Authors' Addresses	24

1. Introduction

An Extensible Authentication Protocol (EAP) tunnel method is an EAP method that establishes a secure tunnel and executes other EAP methods under the protection of that secure tunnel. An EAP tunnel method can be used in any lower layer protocol that supports EAP authentication. There are several existing EAP tunnel methods that use Transport Layer Security (TLS) to establish the secure tunnel. EAP methods supporting this include Protected EAP (PEAP) [PEAP], Tunneled Transport Layer Security EAP (TTLS) [RFC5281] and EAP Flexible Authentication via Secure Tunneling (EAP-FAST) [RFC4851]. In general this has worked well so there is consensus to continue to use TLS as the basis for a tunnel method. There have been various reasons for employing a protected tunnel for EAP processes. They include protecting weak authentication exchanges, such as username and password. In addition a protected tunnel can provide means to provide peer identity protection and EAP method chaining. Finally, systems have found it useful to transport additional types of data within the protected tunnel.

This document describes the requirements for a EAP tunnel method as well as for a password protocol supporting legacy password verification within the tunnel method.

2. Conventions Used In This Document

Use of each capitalized word within a sentence or phrase carries the following meaning during the EAP Method Update (EMU) WG's method selection process:

MUST - indicates an absolute requirement

MUST NOT - indicates something absolutely prohibited

SHOULD - indicates a strong recommendation of a desired result

SHOULD NOT - indicates a strong recommendation against a result

MAY - indicates a willingness to allow an optional outcome

Lower case uses of "MUST", "MUST NOT", "SHOULD", "SHOULD NOT" and "MAY" carry their normal meaning and are not subject to these definitions.

3. Use Cases

To motivate and explain the requirements in this document, a representative set of use cases for the EAP tunnel method are supplied here. It is mandatory for a candidate tunnel method to support all of the use cases that are marked below as "MUST".

3.1. Password Authentication

Many legacy systems only support user authentication with passwords. Some of these systems require transport of the actual username and password to the authentication server. This is true for systems where the authentication server does not have access to the cleartext password or a consistent transform of the cleartext password. Example of such systems are some one time password (OTP) systems and other systems where the username and password are submitted to an external party for validation. The tunnel method MUST support transporting cleartext username and password to the EAP server. It MUST NOT reveal information about the username and password to parties in the communication path between the peer and the EAP Server. The advantage any attacker gains against the tunnel method when employing a username and password for authentication MUST be through interaction and not computation. The tunnel MUST support protection from man-in-the-middle attacks. The combination of the tunnel authentication and password authentication MUST enable mutual authentication.

Since EAP authentication occurs before network access is granted the tunnel method SHOULD enable an inner exchange to provide support for minimal password management tasks including password change, "new PIN mode", and "next token mode" required by some systems.

3.2. Protection of Weak EAP Methods

Some existing EAP methods have vulnerabilities that could be eliminated or reduced by running them inside a protected tunnel. For example, a EAP-MD5 does not provide mutual authentication or protection from dictionary attacks. Without extra protection, EAP tunnel methods are vulnerable to a special type of tunnel man-in-the-middle attack [TUNNEL-MITM]. This attack is referred to as "tunnel MitM attack" in the remainder of this document. The additional protection needed to thwart tunnel MitM attacks depends on the inner method executed within the tunnel. When weak methods are used, these attacks can be mitigated via security policies that require the method to be used only within a tunnel. On the other hand, a technical solution (so-called cryptographic bindings) can be used whenever the inner method derives key material and is not susceptible to attacks outside a tunnel. Only the latter mitigation technique

can be made an actual requirement for EAP tunnel methods (see Section 4.6.3), while security policies are outside the scope of this requirement draft. Please refer to the NIST Recommendation for EAP Methods Used in Wireless Network Access Authentication [NIST SP 800-120] and [LCN 2010] for a discussion on security policies and complete solutions for thwarting tunnel MitM attacks.

The tunnel method MUST support protection of weak EAP methods. Cryptographic protection from tunnel MitM attacks MUST be provided for all key generating methods. In combination with an appropriate security policy this will thwart MitM attacks against inner methods.

3.3. Chained EAP Methods

Several circumstances are best addressed by using chained EAP methods. For example, it may be desirable to authenticate the user and also authenticate the device being used. However, chained EAP methods from different conversations can be re-directed into the same conversation by an attacker giving the authenticator the impression that both conversations terminate at the same end-point. Cryptographic binding can be used to bind the results of chained key generating methods together or to an encompassing tunnel.

The tunnel method MUST support chained EAP methods while including protection against attacks on method chaining.

3.4. Identity Protection

When performing an EAP authentication, the peer may want to protect its identity and only disclose it to a trusted EAP server. This helps to maintain peer privacy.

The tunnel method MUST support identity protection, therefore the identity of the peer used for authentication purposes MUST NOT be obtainable by any entity other than the EAP server terminating the tunnel method. Peer identity protection provided by the tunnel method applies to tunnel method and inner method specific identities. Note that the peer may need to expose the realm portion of the EAP outer identity in the Network Access Identifier (NAI) [RFC4282] in a roaming scenario in order to reach the appropriate authentication server.

3.5. Anonymous Service Access

When network service is provided, it is sometimes desirable for a user to gain network access in order to access limited services for emergency communication or troubleshooting information. To avoid eavesdropping, it's best to negotiate link layer security as with any

other authentication.

Therefore, the tunnel method SHOULD allow anonymous peers or server-only authentication, while still deriving keys that can be used for link layer security. The tunnel method MAY also allow for the bypass of server authentication processing on the client.

Forgoing user or server authentication increases the chance of man-in-the-middle and other types of attacks that can compromise the derived keys used for link layer security. Therefore, passwords and other sensitive information MUST NOT be disclosed to an unauthenticated server, or to a server that is not authorized to authenticate the user.

3.6. Network Endpoint Assessment

The Network Endpoint Assessment (NEA) protocols and reference model described in [RFC5209] provide a standard way to check the health ("posture") of a device at or after the time it connects to a network. If the device does not comply with the network's requirements, it can be denied access to the network or granted limited access to remediate itself. EAP is a convenient place for conducting an NEA exchange.

The tunnel method SHOULD support carrying NEA protocols such as PB-TNC [RFC5793]. Depending on the specifics of the tunnel method, these protocols may be required to be carried in an EAP method.

3.7. Client Authentication During Tunnel Establishment

In some cases the peer will have credentials that allow it to authenticate during tunnel establishment. These credentials may only partially authenticate the identity of the peer and additional authentication may be required inside the tunnel. For example, a communication device may be authenticated during tunnel establishment, in addition user authentication may be required to satisfy authentication policy. The tunnel method MUST be capable of providing client side authentication during tunnel establishment.

3.8. Extensibility

The tunnel method MUST provide extensibility so that additional data related to authentication, authorization and network access can be carried inside the tunnel in the future. This removes the need to develop new tunneling methods for specific purposes.

An application for extensibility is credential provisioning. When a peer has authenticated with EAP, this is a convenient time to

distribute credentials to that peer that may be used for later authentication exchanges. For example, the authentication server can provide a private key or shared key to the peer that can be used by the peer to perform rapid re-authentication or roaming. In addition there have been proposals to perform enrollment within EAP, such as [I-D.mahy-eap-enrollment]. Another use for extensibility is support for alternate authentication frameworks within the tunnel.

3.9. Certificate-less Authentication and Generic EAP Method Extension

In some cases the peer will not have a way to verify a server certificate and in some cases a server might not have a certificate to verify. Therefore, it is desirable to support certificate-less authentication. An application for this is credential provisioning where the peer and server authenticate each other with a shared password and credentials for subsequent authentication (e.g. a key pair and certificate or a shared key) can be passed inside the tunnel. Another application is to extend existing EAP methods with new features such as EAP channel bindings.

Great care must be taken when using tunnels with no server authentication for the protection of an inner method. For example, the client may lack the appropriate trust roots to fully authenticate the server, but may still establish the tunnel to execute an inner EAP method to perform mutual authentication and key derivation. In these cases, the inner EAP method **MUST** provide resistance to dictionary attack and a cryptographic binding between the inner method and the tunnel method **MUST** be established. Furthermore, the cipher suite used to establish the tunnel **MUST** derive the master key using contribution from both client and server, as in ephemeral Diffie-Hellman cipher suites.

The tunnel method **MAY** allow for certificate-less authentication.

4. Requirements

4.1. General Requirements

4.1.1. RFC Compliance

The tunnel method **MUST** include a Security Claims section with all security claims specified in Section 7.2 in RFC 3748 [RFC3748]. In addition, it **MUST** meet the requirement in Sections 2.1 and 7.4 of RFC 3748 that tunnel methods **MUST** support protection against man-in-the-middle attacks. Furthermore, the tunnel method **MUST** support identity protection as specified in Section 7.3 in RFC 3748.

The tunnel method MUST be unconditionally compliant with RFC 4017 [RFC4017] (using the definition of "unconditionally compliant" contained in section 1.1 of RFC 4017). This means that the method MUST satisfy all the MUST, MUST NOT, SHOULD, and SHOULD NOT requirements in RFC 4017.

The tunnel method MUST meet all the MUST and SHOULD requirements relevant to EAP methods contained in the EAP Key Management Framework [RFC5247] or any successor. This includes the generation of the MSK, EMSK, Peer-Id, Server-Id and Session-Id. These requirements will enable the tunnel method to properly fit into the EAP key management framework, maintaining all of the security properties and guarantees of that framework.

The tunnel method MUST NOT be tied to any single cryptographic algorithm. Instead, it MUST support run-time negotiation to select among an extensible set of cryptographic algorithms, such as algorithms used with certificates presented during tunnel establishment. This "cryptographic algorithm agility" provides several advantages. Most important, when a weakness in an algorithm is discovered or increased processing power overtakes an algorithm, users can easily transition to a new algorithm. Also, users can choose the algorithm that best meets their needs.

The tunnel method MUST meet the SHOULD and MUST requirements pertinent to EAP method contained in Section 3 of RFC 4962 [RFC4962]. This includes: cryptographic algorithm independence; strong, fresh session keys; replay detection; keying material confidentiality and integrity; and confirmation of cipher suite selection.

4.2. Tunnel Requirements

The following section discusses requirements for TLS Tunnel Establishment.

4.2.1. TLS Requirements

The tunnel based method MUST support TLS version 1.2 [RFC5246] and may support earlier versions greater than SSL 2.0 to enable the possibility of backwards compatibility.

4.2.1.1. Cipher Suites

4.2.1.1.1. Cipher Suite Negotiation

Cipher suite negotiations always suffer from downgrading attacks when they are not secured by any kind of integrity protection. A common practice is a post integrity check in which, as soon as available,

the established keys (here the tunnel key) are used to derive integrity keys. These integrity keys are then used by peer and authentication server to verify whether the cipher suite negotiation has been maliciously altered by another party.

Integrity checks prevent downgrading attacks only if the derived integrity keys and the employed integrity algorithms cannot be broken in real-time. See Section 6.1 or [KHL07] for more information on this. Hence, the tunnel method MUST provide integrity protected cipher suite negotiation with secure integrity algorithms and integrity keys.

TLS provides protected cipher suite negotiation as long as all the cipher suites supported provide authentication, key establishment and data integrity protection as discussed in Section 6.1.

4.2.1.1.2. Tunnel Data Protection Algorithms

In order to prevent attacks on the cryptographic algorithms employed by inner authentication methods, a tunnel protocol's protection needs to provide a basic level of algorithm strength. The tunnel method MUST provide at least one mandatory to implement cipher suite that provides the equivalent security of 128-bit AES for encryption and message authentication. See Part 1 of the NIST Recommendation for Key Management [NIST SP 800-57] for a discussion of the relative strengths of common algorithms.

4.2.1.1.3. Tunnel Authentication and Key Establishment

A tunnel method MUST provide unidirectional authentication from authentication server to EAP peer and mutual authentication between authentication server and EAP peer. The tunnel method MUST provide at least one mandatory to implement cipher suite that provides certificate-based authentication of the server and provides optional certificate-based authentication of the client. Other types of authentication MAY be supported.

At least one mandatory to implement cipher suite MUST be approved by NIST Draft Recommendation for Key Management, Part 3 [NIST SP 800-57p3], i.e., the ciphersuite MUST be listed in Table 4-1, 4-2 or 4-3 in that document.

The mandatory to implement cipher suites MUST only include cipher suites that use strong cryptographic algorithms. They MUST NOT include cipher suites providing mutually anonymous authentication or static Diffie-Hellman cipher suites.

Other ciphersuites MAY be selected following the security

requirements for tunnel protocols in NIST DRAFT Recommendation for EAP Methods Used in Wireless Network Access Authentication [NIST SP 800-120].

4.2.1.2. Tunnel Replay Protection

In order to prevent replay attacks on a tunnel protocol, the message authentication MUST be generated using a time-variant input such as timestamps, sequence numbers, nonces, or a combination of these so that any re-use of the authentication data can be detected as invalid. TLS provides sufficient replay protection to meet this requirements as long as weak cipher suites discussed in Section 6.1 are avoided.

4.2.1.3. TLS Extensions

In order to meet the requirements in this document TLS extensions MAY be used. For example, TLS extensions may be useful in providing certificate revocation information via the TLS Online Certificate Status Protocol (OCSP) extension [I-D.ietf-tls-rfc4366-bis] (thus meeting the requirement in Section 4.5.1.3).

4.2.1.4. Peer Identity Privacy

A tunnel protocol MUST support peer privacy. This requires that the username and other attributes associated with the peer are not transmitted in the clear or to an unauthenticated, unauthorized party. Peer identity protection provided by the tunnel method applies to establishment of the tunnel and protection of inner method specific identities. If applicable, the peer certificate is sent confidentially (i.e. encrypted).

4.2.1.5. Session Resumption

The tunnel method MUST support TLS session resumption as defined in [RFC5246]. The tunnel method MAY support other methods of session resumption such as those defined in [RFC5077].

4.2.2. Fragmentation

Tunnel establishment sometimes requires the exchange of information that exceeds what can be carried in a single EAP message. In addition information carried within the tunnel may also exceed this limit. Therefore a tunnel method MUST support fragmentation and reassembly.

4.2.3. Protection of Data External to Tunnel

A man-in-the-middle attacker can modify clear text values such as protocol version and type code information communicated outside the TLS tunnel. The tunnel method MUST provide implicit or explicit protection of the protocol version and type code. If modification of other information external to the tunnel can cause exploitable vulnerabilities, the tunnel method MUST provide protection against modification of this additional data.

4.3. Tunnel Payload Requirements

This section describes the payload requirements inside the tunnel. These requirements frequently express features that a candidate protocol must be capable of offering so that a deployer can decide whether to make use of that feature. This section does not state requirements about what features of each protocol must be used during a deployment.

4.3.1. Extensible Attribute Types

The payload MUST be extensible. Some standard payload attribute types will be defined to meet known requirements listed below, such as password authentication, inner EAP method, vendor specific attributes, and result indication. Additional payload attributes MAY be defined in the future to support additional features and data types.

4.3.2. Request/Challenge Response Operation

The payload MUST support request and response type of half-duplex operation typical of EAP. Multiple attributes may be sent in a single payload. The payload MAY support carrying on multiple authentications in a single payload packet.

4.3.3. Indicating Criticality of Attributes

It is expected that new attributes will be defined to be carried within the tunnel method. In some cases it is necessary for the sender to know if the receiver did not understand the attribute. To support this, there MUST be a way for the sender to mark attributes such that the receiver will indicate if an attribute is not understood.

4.3.4. Vendor Specific Support

The payload MUST support communication of an extensible set of vendor-specific attributes. These attributes will be segmented into

uniquely identified vendor specific name spaces. They can be used for experiments or vendor specific features.

4.3.5. Result Indication

The payload MUST support result indication and its acknowledgement, so both the EAP peer and server will end up with a synchronized state. The result indication is needed after each chained inner authentication method and at the end of the authentication, so separate result indication for intermediate and final result MUST be supported.

4.3.6. Internationalization of Display Strings

The payload MAY provide a standard attribute format that supports international strings. This attribute format MUST support encoding strings in UTF-8 [RFC3629] format. Any strings sent by the server intended for display to the user MUST be sent in UTF-8 format and SHOULD be able to be marked with language information and adapted to the user's language preference as indicated by RFC 5646 [RFC5646]. Note that in some cases, such as when transmitting error codes, it is acceptable to exchange numeric codes that can be translated by the client to support the particular local language. These numeric codes are not subject internationalization during transmission.

4.4. EAP Channel Binding Requirements

The tunnel method MUST be capable of meeting EAP channel binding requirements described in [I-D.ietf-emu-chbind]. As discussed in [RFC5056], EAP Channel bindings differ from channel bindings discussed in other contexts. Cryptographic binding between the TLS tunnel and the inner method discussed in Section 4.6.3 relates directly to the non-EAP channel binding concepts discussed in RFC 5056.

4.5. Requirements Associated with Carrying Username and Passwords

This section describes the requirements associated with tunneled password authentication. The password authentication mentioned here refers to user or machine authentication using a legacy password database or verifier, such as LDAP [RFC4511], OTP, etc. These typically require the password in its original text form in order to authenticate the peer, hence they require the peer to send the clear text user name and password to the EAP server.

4.5.1. Security

Many internal EAP methods have the peer send its password in the clear to the EAP server. Other methods (e.g. challenge-response methods) are vulnerable to attacks if an eavesdropper can intercept the traffic. For any such methods, the security measures in the following sections MUST be met.

4.5.1.1. Confidentiality and Integrity

The clear text password exchange MUST be integrity and confidentiality protected. As long as the password exchange occurs inside an authenticated and encrypted tunnel, this requirement is met.

4.5.1.2. Authentication of Server

The EAP server MUST be authenticated before the peer sends the clear text password to the server.

4.5.1.3. Server Certificate Revocation Checking

When certificate authentication is used during tunnel establishment the EAP peer may need to present its password to the server before it has network access to check the revocation status of the server's credentials. Therefore, the tunnel method MUST support mechanisms to check the revocation status of a credential. The tunnel method SHOULD make use of Online Certificate Status Protocol (OCSP) [RFC2560] or Server-based Certificate Validation Protocol (SCVP) [RFC5055] to obtain the revocation status of the EAP server certificate.

4.5.2. Internationalization

The password authentication exchange MUST support user names and passwords in international languages. It MUST support encoding of user name and password strings in UTF-8 [RFC3629] format. The method MUST specify how username and password normalizations and/or comparisons is performed in reference to SASLPrep [RFC4013], Net-UTF-8 [RFC5198] or their replacement.

Any strings sent by the server intended for display to the user MUST be sent in UTF-8 format and SHOULD be able to be marked with language information and adapted to the user's language preference as indicated by RFC 5646 [RFC5646]. Note that in some cases, such as when transmitting error codes, it is acceptable to exchange numeric codes that can be translated by the client to support the particular local language. These numeric codes are not subject to

internationalization during transmission.

4.5.3. Meta-data

The password authentication exchange SHOULD support additional associated meta-data which can be used to indicate whether the authentication is for a user or a machine. This allows the EAP server and peer to request and negotiate authentication specifically for a user or machine. This is useful in the case of multiple inner authentications where the user and machine both need to be authenticated.

4.5.4. Password Change

The password authentication exchange MUST support password change. The exchange SHOULD be extensible to support other "housekeeping" functions, such as the management of PINs or other data, required by some systems.

4.6. Requirements Associated with Carrying EAP Methods

The tunnel method MUST be able to carry inner EAP methods without modifying them. EAP methods MUST NOT be redefined inside the tunnel.

4.6.1. Method Negotiation

The tunnel method MUST support the protected negotiation of the inner EAP method. It MUST NOT allow the inner EAP method negotiation to be manipulated by intermediaries.

4.6.2. Chained Methods

The tunnel method SHOULD support the chaining of multiple EAP methods. The tunnel method MUST allow for the communication of intermediate result and verification of compound binding between executed inner methods when chained methods are employed.

4.6.3. Cryptographic Binding with the TLS Tunnel

The tunnel method MUST provide a mechanism to bind the tunnel protocol and the inner EAP method. This property is referred to as cryptographic binding. Such bindings are an important tool for mitigating the tunnel MitM attacks on tunnel methods [TUNNEL-MITM]. Cryptographic bindings enable the complete prevention of tunnel MitM attacks without the need of additional security policies as long as the inner method derives keys and is not vulnerable to attacks outside a protected tunnel [LCN 2010]. Even though weak or non-key deriving inner methods may be permitted, and thus security policies

preventing tunnel MitM attacks are still necessary, the tunnel method MUST provide cryptographic bindings, because only this allows migrating to more secure, policy-independent implementations.

Cryptographic bindings are typically achieved by securely mixing the established keying material (say tunnel key TK) from the tunnel protocol with the established keying material (say method key MK) from the inner authentication method(s) in order to derive fresh keying material. If chained EAP methods are executed in the tunnel, all derived inner keys are combined with the tunnel key to create a new compound tunnel key (CTK). In particular, CTK is used to derive the EAP MSK, EMSK and other transient keys (TEK), such as transient encryption keys and integrity protection keys. The key hierarchy for tunnel methods executions that derive compound keys for the purpose of cryptographic binding is depicted in Figure 1.

In the case of the sequential executions of n inner methods, a chained compound key CTK_i MUST be computed upon the completion of each inner method i such that it contains the compound key of all previous inner methods, i.e. $CTK_i = f(CTK_{i-1}, MK_i)$ with $0 < i \leq n$ and $CTK_0 = TK$, where $f()$ is a key derivation function, such as one that complies with NIST Recommendation for Key Derivation Using Pseudorandom Functions [NIST SP 800-108]. CTK_n SHOULD serve as the key to derive further keys. Figure 1 depicts the key hierarchy in the case of a single inner method. Transient keys derived from the compound key CTK are used in a cryptographic protocol to verify the integrity of the tunnel and the inner authentication method.




Figure 1: Compound Keys

Furthermore, all compound keys CTK_i and all keys derived from it

SHOULD follow the recommendations for key derivations and key hierarchies as specified in [NIST SP 800-108]. In particular, all derived keys MUST have a lifetime assigned that does not exceed the lifetime of any key higher in the key hierarchy. The derivation MUST prevent a compromise in one part of the system from leading to compromises in other parts of the system that relay on keys at the same or higher level in the hierarchy.

4.6.4. Peer Initiated

The tunnel method SHOULD allow for the peer to initiate an inner EAP authentication in order to meet its policy requirements for authenticating the server.

4.6.5. Method Meta-data

The tunnel method SHOULD allow for the communication of additional data associated with an EAP method. This can be used to indicate whether the authentication is for a user or a machine. This allows the EAP server and peer to request and negotiate authentication specifically for a user or machine. This is useful in the case of multiple inner EAP authentications where the user and machine both need to be authenticated.

5. IANA Considerations

This document has no IANA considerations.

6. Security Considerations

A tunnel method is often deployed to provide mutual authentication between EAP Peer and EAP Server and to generate key material for use in protecting lower layer protocols. In addition the tunnel is used to protect the communication of additional data, including peer identity between the EAP Peer and EAP Server from disclosure to or modification by an attacker. These sections cover considerations that affect the ability for a method to achieve these goals.

6.1. Cipher Suite Selection

TLS supports a wide variety of cipher suites providing a variety of security properties. The selection of cipher suites is critical to the security of the tunnel method. Selection of a cipher suite with weak or no authentication, such as an anonymous Diffie-Hellman based cipher suite will greatly increase the risk of system compromise. Since a tunnel method uses the TLS tunnel to transport data, the

selection of a ciphersuite with weak data encryption and integrity algorithms will also increase the vulnerability of the method to attacks.

A tunnel protocol is prone to downgrading attacks if the tunnel protocol supports any key establishment algorithm that can be broken on-line. In a successful downgrading attack, an adversary breaks the selected "weak" key establishment algorithm and optionally the "weak" authentication algorithm without being detected. Here, "weak" refers to a key establishment algorithm that can be broken in real-time, and an authentication scheme that can be broken off-line, respectively. See [KHLCO7] for more details. The requirements in this document disapprove the use of key establishment algorithms that can be broken on-line.

Mutually anonymous tunnel protocols are prone to man-in-the-middle attacks described in [KHLCO7]. During such an attack, an adversary establishes a tunnel with each the peer and the authentication server, while peer and server believe that they established a tunnel with each other. Once both tunnels have been established, the adversary can eavesdrop on all communications within the tunnels, i.e. the execution of the inner authentication method(s). Consequently, the adversary can eavesdrop on the identifiers that are exchanged as part of the EAP method and thus, the privacy of peer and/or authentication server is compromised along with any other data transmitted within the tunnels. This document requires server authentication to avoid the risks associated with anonymous cipher suites.

6.2. Tunneled Authentication

In many cases a tunnel method provides mutual authentication by authenticating the server during tunnel establishment and authenticating the peer within the tunnel using an EAP method. As described in [TUNNEL-MITM], this mode of operation can allow tunnel man-in-the-middle attackers to authenticate to the server as the peer by tunneling the inner EAP protocol messages to and from a peer executing the method outside a tunnel or with an untrustworthy server. Cryptographic binding between the established keying material from the inner authentication method(s) and the tunnel protocol verifies that the endpoints of the tunnel and the inner authentication method(s) are the same. This can thwart the attack if the inner method derived keys of sufficient strength that they cannot be broken in real-time.

In cases where the inner authentication method does not generate any or only weak key material, security policies MUST be enforced such that the peer cannot execute the inner method with the same

credentials outside a protective tunnel or with an untrustworthy server.

6.3. Data External to Tunnel

The tunnel method will use data that is outside the TLS tunnel such as the EAP type code or version numbers. If an attacker can compromise the protocol by modifying these values the tunnel method MUST protect this data from modification. In some cases external data may not need additional protection because it is implicitly verified during the protocol operation.

6.4. Separation of TLS Tunnel and Inner Authentication Termination

Terminating the inner method at a different location than the outer tunnel needs careful consideration. The inner method data may be vulnerable to modification and eavesdropping between the server that terminates the tunnel and the server that terminates the inner method. For example if a clear text password is used then it may be sent to the inner method server in a RADIUS password attribute which uses weak encryption that may not be suitable protection for many environments.

In some cases terminating the tunnel at a different location may make it difficult for a peer to authenticate the server and trust it for further communication. For example, if the TLS tunnel is terminated by a different organization the peer needs to be able to authenticate and authorize the tunnel server to handle secret credentials that it shares with the home server that terminates the inner method. This may not meet the security policy of many environments.

7. References

7.1. Normative References

- [I-D.ietf-emu-chbind]
Hartman, S., Clancy, C., and K. Hoeper, "Channel Binding Support for EAP Methods", draft-ietf-emu-chbind-06 (work in progress), October 2010.
- [RFC2560] Myers, M., Ankney, R., Malpani, A., Galperin, S., and C. Adams, "X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP", RFC 2560, June 1999.
- [RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, November 2003.

- [RFC3748] Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H. Levkowitz, "Extensible Authentication Protocol (EAP)", RFC 3748, June 2004.
- [RFC4017] Stanley, D., Walker, J., and B. Aboba, "Extensible Authentication Protocol (EAP) Method Requirements for Wireless LANs", RFC 4017, March 2005.
- [RFC4962] Housley, R. and B. Aboba, "Guidance for Authentication, Authorization, and Accounting (AAA) Key Management", BCP 132, RFC 4962, July 2007.
- [RFC5055] Freeman, T., Housley, R., Malpani, A., Cooper, D., and W. Polk, "Server-Based Certificate Validation Protocol (SCVP)", RFC 5055, December 2007.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.
- [RFC5247] Aboba, B., Simon, D., and P. Eronen, "Extensible Authentication Protocol (EAP) Key Management Framework", RFC 5247, August 2008.

7.2. Informative References

- [I-D.ietf-tls-rfc4366-bis]
3rd, D., "Transport Layer Security (TLS) Extensions: Extension Definitions", draft-ietf-tls-rfc4366-bis-12 (work in progress), September 2010.
- [I-D.mahy-eap-enrollment]
Mahy, R., "An Extensible Authentication Protocol (EAP) Enrollment Method", draft-mahy-eap-enrollment-01 (work in progress), March 2006.
- [KHLCO7] Hoeper, K. and L. Chen, "Where EAP Security Claims Fail", ICST QShine , August 2007.
- [LCN 2010]
Hoeper, K. and L. Chen, "An Inconvenient Truth about Tunneled Authentications", Proceedings of 35th Annual IEEE Conference on Local Computer Networks (LCN 2010) , September 2009.
- [NIST SP 800-108]
Chen, L., "Recommendation for Key Derivation Using Pseudorandom Functions", Draft NIST Special Publication 800-108, April 2008.

- [NIST SP 800-120]
Hoeper, K. and L. Chen, "Recommendation for EAP Methods
Used in Wireless Network Access Authentication", NIST
Special Publication 800-120, September 2009.
- [NIST SP 800-57]
Barker, E., Barker, W., Burr, W., Polk, W., and M. Smid,
"Recommendation for Key Management - Part 1: General
(Revised)", NIST Special Publication 800-57, part 1,
March 2007.
- [NIST SP 800-57p3]
Barker, E., Burr, W., Jones, A., Polk, W., , S., and M.
Smid, "Recommendation for Key Management, Part 3
Application-Specific Key Management Guidance", Draft NIST
Special Publication 800-57, part 3, October 2008.
- [PEAP] Microsoft Corporation, "[MS-PEAP]: Protected Extensible
Authentication Protocol (PEAP) Specification",
August 2009.
- [RFC4013] Zeilenga, K., "SASLprep: Stringprep Profile for User Names
and Passwords", RFC 4013, February 2005.
- [RFC4282] Aboba, B., Beadles, M., Arkko, J., and P. Eronen, "The
Network Access Identifier", RFC 4282, December 2005.
- [RFC4511] Sermersheim, J., "Lightweight Directory Access Protocol
(LDAP): The Protocol", RFC 4511, June 2006.
- [RFC4851] Cam-Winget, N., McGrew, D., Salowey, J., and H. Zhou, "The
Flexible Authentication via Secure Tunneling Extensible
Authentication Protocol Method (EAP-FAST)", RFC 4851,
May 2007.
- [RFC5056] Williams, N., "On the Use of Channel Bindings to Secure
Channels", RFC 5056, November 2007.
- [RFC5077] Salowey, J., Zhou, H., Eronen, P., and H. Tschofenig,
"Transport Layer Security (TLS) Session Resumption without
Server-Side State", RFC 5077, January 2008.
- [RFC5198] Klensin, J. and M. Padlipsky, "Unicode Format for Network
Interchange", RFC 5198, March 2008.
- [RFC5209] Sangster, P., Khosravi, H., Mani, M., Narayan, K., and J.
Tardo, "Network Endpoint Assessment (NEA): Overview and
Requirements", RFC 5209, June 2008.

- [RFC5281] Funk, P. and S. Blake-Wilson, "Extensible Authentication Protocol Tunneled Transport Layer Security Authenticated Protocol Version 0 (EAP-TTLSv0)", RFC 5281, August 2008.
- [RFC5646] Phillips, A. and M. Davis, "Tags for Identifying Languages", BCP 47, RFC 5646, September 2009.
- [RFC5793] Sahita, R., Hanna, S., Hurst, R., and K. Narayan, "PB-TNC: A Posture Broker (PB) Protocol Compatible with Trusted Network Connect (TNC)", RFC 5793, March 2010.
- [TUNNEL-MITM] Asokan, N., Niemi, V., and K. Nyberg, "Man-in-the-Middle in Tunnelled Authentication Protocols", Cryptology ePrint Archive: Report 2002/163, November 2002.

Appendix A. Changes from -01

- o Added combined mutual authentication in section 3.1
- o Changed reference from SP 800-52 to SP 800-57, part 3
- o In section 6.2 changed terminology to tunnel MitM and security policy enforcement
- o Reworded text in section 3.2 to clarify MITM protection
- o Added more specific text about derivation of the CTK
- o Removed resource constrained section
- o Added support for Non EAP authentication as a use for extensibility
- o Added text to emergency services section to emphasize that sensitive information should not be sent if the server is unauthenticated.
- o Reworded TLS requirements
- o Reworded external data protection requirements
- o Added text to section 4.6 that states method must not be re-defined inside the tunnel.
- o Editorial fixes

Appendix B. Changes from -02

- o Editorial Fixes
- o Clarified client authentication during tunnel establishment
- o Changed text so that the tunnel method MUST meet all MUST and SHOULD requirements relevant to EAP methods in RFCs 4962 and 5247

Appendix C. changes from -03

- o Resolution of open issues:
<http://trac.tools.ietf.org/wg/emu/trac/report/9>

- o Revised section 2 to match other similar RFC(Issue 6)
- o Cleaned up section 3.2 (issue 8)
- o Clarified identity protection scope in section 3.4 and 4.2.1.4(issue 9)
- o Changed Emergency Services to anonymous authentication(section 3.5)(issue 10)
- o Clarified section 4.1.1 (issue 15)
- o Cleaned up TLS requirements in section 4.2.1(issue 11)
- o Replaced text in 4.2.1.1.3 with suitable reference
- o Improved wording in 4.2.3 and 6.3 (issue 13)
- o Update internationalization requirements in 4.3.6 and 4.5.2 (Issues 25,18)
- o Updated text in 4.5.1 (issue 16)
- o Changed meta-data to SHOULD in 4.5.3 and 4.6.5(Issue 20)
- o Changed chained methods to SHOULD in 4.6.2(issue 19)
- o Added security consideration for inner method termination(issue 24)
- o Updated references
- o Editorial changes(issues 7,22,17)

Authors' Addresses

Katrin Hoeper
Motorola, Inc.
1301 E Algonquin Rd
Schaumburg, IL 60196
USA

Email: khoeper@motorola.com

Stephen Hanna
Juniper Networks
3 Beverly Road
Bedford, MA 01730
USA

Email: shanna@juniper.net

Hao Zhou
Cisco Systems, Inc.
4125 Highlander Parkway
Richfield, OH 44286
USA

Email: hzhou@cisco.com

Joseph Salowey (editor)
Cisco Systems, Inc.
2901 3rd. Ave
Seattle, WA 98121
USA

Email: jsalowey@cisco.com

