Network Working Group   Y. Ohba, Ed.
Internet-Draft   Toshiba
Intended status: Informational   Q. Wu, Ed.
Expires: November 16, 2009   Huawei
                                                           G. Zorn, Ed.
                                                          Network Zen
                                                          May 15, 2009

                  EAP Early Authentication Problem Statement
                        draft-ietf-hokey-preauth-ps-07

Status of this Memo

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

   Internet-Drafts are working documents of the Internet Engineering
   Task Force (IETF), its areas, and its working groups. Note that
   other groups may also distribute working documents as Internet-
   Drafts.

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

   The list of current Internet-Drafts can be accessed at
   http://www.ietf.org/ietf/1id-abstracts.txt.

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

   This Internet-Draft will expire on November 16, 2009.

Copyright Notice

   Copyright (c) 2009 IETF Trust and the persons identified as the

Abstract

   EAP (Extensible Authentication Protocol) early authentication may be
   defined as the use of EAP to establish authenticated keying material
   on a target authenticator prior to arrival of the peer at the access
   network managed by that authenticator. This draft discusses the EAP
   early authentication problem in detail.

Table of Contents

1. Introduction

When a mobile device, during an active communication session, moves
from one access network to another and changes its point of
attachment, the session may be subjected to disruption in the
continuity of service due to the delay associated with the handover
operation. The performance requirements of a real-time application
will vary based on the type of application and its characteristics
such as delay and packet loss tolerance. For VoIP applications,
ITU-T G.114 [ITU] recommends a steady-state end-to-end delay of 150
ms as the upper limit and rates 400 ms as generally unacceptable
delay. Similarly, a streaming application has a tolerable packet
(SDU) error rates ranging from 0.1 to 0.00001 with a transfer delay
of less than 300 ms. Any help that an optimized handoff mechanism
can provide toward meeting these objectives is useful. The ultimate
objective is to achieve seamless handover with low latency, even when
handover is between different link technologies or between different
AAA domains.

As a mobile device goes through a handover process, it is subjected
to delay because of the rebinding of its association at or across
several layers of the protocol stack and because of the additional
round trips needed for a new EAP exchange. Delays incurred within
each protocol layer affect the ongoing multimedia application and
data traffic within the client [WCM].

The handover process often requires authentication and authorization
for acquisition or modification of resources assigned to the mobile
device. In most cases, this authentication and authorization needs
interaction with a central authority in a domain. In some cases the
central authority may be placed far away from the mobile device. The
delay introduced due to such an authentication and authorization
procedure adds to the handover latency and consequently affects
ongoing application sessions[MQ7]. The discussion in this document
is focused on mitigating delay due to network access authentication
and authorization.

2. Terminology

AAA Authentication, Authorization, and Accounting. AAA protocols
   RADIUS [RFC2865] and Diameter [RFC3588].

AAA domain
   The set of access networks within the scope of a specific AAA
   server. Thus, if a peer changes from one point of attachment to
   another within the same AAA domain, it continues to be served by
   the same AAA server.

Access Point (AP)
   A network point of attachment in a IEEE 802.11 wireless LAN
   [IEEE.802-11.2007].

Authenticator
   See [RFC3748].

Basic Service Set (BSS)
   The basic building block of an IEEE 802.11 wireless LAN
   [IEEE.802-11.2007]. A BSS consists of a group of any number of
   802.11 stations.

Candidate Access Network
   An access network that can potentially become the target access
   network for a peer. There can be multiple candidate access
   networks for the peer.

Candidate Authenticator (CA)
   An authenticator that can potentially become the target
   authenticator for a peer. There can be multiple candidate
   authenticators for the peer.

EAP Server
   See [RFC3748].

EAP Early Authentication (EEA)
   The utilization of EAP to pre-establish EAP keying material on an
   EAP authenticator prior to arrival on a link served by that
   authenticator of the mobile device that acts as an EAP peer.

Extended Service Set (ESS)
   A set of infrastructure BSSs in IEEE 802.11 wireless LAN
   [IEEE.802-11.2007], where the access points communicate amongst
   themselves to forward traffic from one BSS to another to
   facilitate movement of stations between BSSs.

Inter-AAA-Domain Handover (Inter-Domain Handover)
   A handover across multiple AAA domains.

Inter-Authenticator Handover
   A handover across multiple authenticators. An inter-access-domain
   handover, an inter-ESS handover, an inter-AAA-domain handover, an
   inter-technology handover can be view as examples of inter-
   authenticator handover.

   Inter-ESS Handover
      An 802.11 handover across multiple ESSs.

   Inter-Technology Handover
      A handover across different link layer technologies.

   Intra-AAA-Domain Handover (Intra-Domain Handover)
      A handover within the same AAA domain. Intra-AAA-domain handover
      include a handover across different authenticators within the same
      AAA domain.

   Intra-Technology Handover
      A handover within the same link layer technology.

   Master Session Key (MSK)
      See [RFC3748].

   Peer
      The entity that responds to the authenticator (below); for
      details, see [RFC3748].

   Serving Access Network
      An access network that is currently serving the peer.

   Serving Authenticator (SA)
      An authenticator that is currently serving the peer.

   Target Access Network
      An access network that has been chosen to be the new serving
      access network for a peer.

   Target Authenticator (TA)
      An authenticator that has been chosen to be the new serving
      authenticator for a peer.

3. Problem Statement

   The basic mechanism of handover is a two-step procedure involving

   o   handover preparation iand

   o   handover execution

   Handover preparation includes the discovery of candidate network
   points of attachment and selection of an appropriate target
   attachment point from the candidate set. Handover execution consists
   of setting up L2 and L3 connectivity with the target. Currently, as

part of the second step, network access authentication and
authorization is performed directly with the target network.
Following a successful EAP authentication, a secure association
procedure is performed between the peer and the target authenticator
to derive a new set of link-layer ciphering keys from EAP keying
material such as the MSK. The second step may require full EAP
authentication in the absence of any particular optimization for
handover key management. The handover latency introduced by full EAP
authentication has proven to be larger than what is acceptable for
real-time application scenarios as described in [MQ7]. Hence,
improvement in the handover latency performance due to EAP is a
necessary objective for such scenarios.

As an example of the second step, in IEEE 802.11 wireless LANs
[IEEE.802-11.2007]the network access authentication and authorization
involves performing a new IEEE 802.1X message exchange with the
authenticator in the target AP to initiate an EAP exchange to the
authentication server[WPA].

As another example, in 3GPP Technical Specification TS 33.402
[TS33.402]], network access authentication and authorization happens
after L2 connection is established between the mobile device and a
non-3GPP target access network, and involves EAP exchange between the
mobile device and 3GPP AAA server through the non-3GPP target access
network.

There has been relevant optimization work undertaken by various
standards organizations, but these efforts have generally been scoped
to specific link layer technologies. The work done in the IEEE
802.11f ([IEEE.802-11F.2003] and 802.11r [IEEE.802-11R.2008]) Task
Groups applies only to transfers within one 802.11 ESS or AAA domain.
[TS33.402] defines the authentication and key management procedures
performed during interworking between non-3GPP access networks and
the Evolved Packet System (EPS). These procedures are not really
independent of link technology, since they assume either that the
authenticator lies in the EPS network or that separate
authentications are performed in the access network and then in the
EPS network. Therefore, a solution is still needed to enable EAP
early authentication for inter-AAA-domain handovers and inter-
technology handovers. A search for solutions at the IP level may
offer the necessary technology independence.

Optimized solutions for secure inter-authenticator handovers can be
seen either as security context transfer (e.g., using the EAP
Extensions for EAP Re-authentication Protocol (ERP)) [RFC5296], or as
EAP early authentication. Security context transfer involves
transfer of reusable key context to the new point of attachment.
Horizontal context transfer of reusable key context is not

recommended [RFC4962] because of the possibility that the compromise
of one authenticator might lead to the compromise of another
authenticator. ERP uses existing EAP keying material obtained from
the AAA server in the home realm to derive a cryptologically
independent re-authentication key to be distributed to an ERP server
in a visited domain. This reduces handover delay by eliminating the
need to run full EAP authentication with the EAP server in the home
domain for intra-domain handovers.

However, there are certain cases where ERP is not applicable or an
additional optimization mechanism is needed to establish a key for
the candidate authenticator:

o One case is an inter-domain handover. A trust relationship is
   required between the home and visited AAA domains. Given that
   trust relationship and assuming the visited AAA domain supports
   ERP, full EAP authentication with the EAP server in the home AAA
   domain is still needed to distribute the existing keying materials
   to the ERP server when the mobile device first enters the visited
   AAA domain.

o Another case is an inter-technology handover where the candidate
   and serving authenticator are different entities belonging to two
   different visited AAA domains and the AAA is same in the home AAA
   domain.

Applicability of EAP early authentication is limited to the scenarios
where candidate authenticators can be discovered and an accurate
prediction of movement can be easily made; also, the effectiveness of
EAP early authentication may be less significant for particular
inter-technology handover scenarios where simultaneous use of
multiple technologies is not a major concern.

In EAP early authentication, AAA-based authentication and
authorization for a candidate authenticator is performed while
ongoing data communication is in progress via the serving network.
The goal of EAP early authentication is to complete AAA signaling for
EAP before the peer moves. There are several AAA issues related to
EAP early authentication. These issues are described in Section 6.

Figure 1 shows the functional elements that are related to EAP early
authentication. These functional elements include a peer, a serving
authenticator, a candidate authenticator and an AAA/EAP server (or
AAA/EAP servers, if this is an inter-AAA-domain handover). When the
serving and candidate authenticators belong to different AAA domains,
the candidate authenticator may use a different AAA server and user
credentials than those were used by the serving authenticator to
authenticate the peer. Alternatively, the candidate authenticator

and the serving authenticator may rely on the same AAA server, which
is located in the home domain of the mobile device.

```
+------+      +-------------+     +------+
| Peer |------|   Serving   |    /        \
|      |      |Authenticator|---/          \
+------+      +-------------+  /            \
   .                         /              \    +-----------------+
   . Move                   +   IP Network   +---|AAA/EAP Server(s)|
   .                         \              /    +-----------------+
   v            +-------------+  \          /
                |  Candidate  |---\        /
                |Authenticator|    \      /
                +-------------+     +------+
```
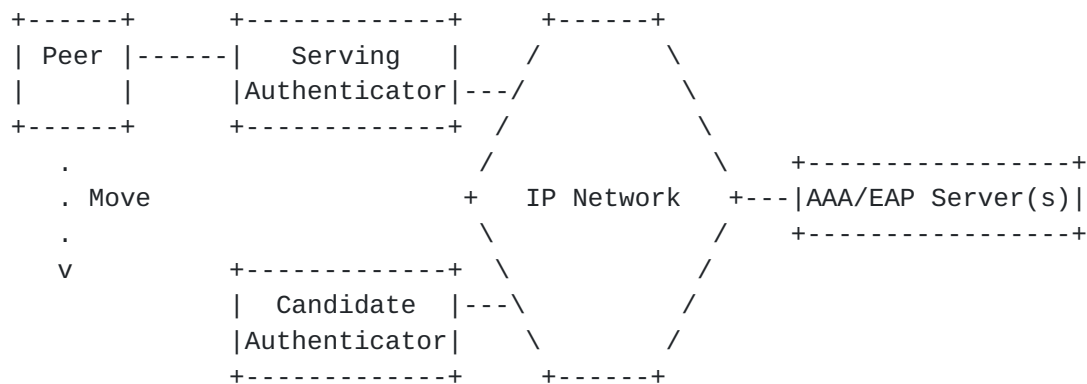
Figure 1: EAP Pre-authentication Functional Elements

A peer is attached to the serving access network. Before the peer
performs handover from the serving access network to a candidate
access network, it performs EAP early authentication with a candidate
authenticator via the serving access network. The peer may perform
EAP early authentication with one or more candidate authenticators.
It is assumed that each authenticator has an IP address. It is
assumed that there is at least one candidate authenticator in each
candidate access network while the serving access network may or may
not have a serving authenticator. The serving and candidate access
networks may use different link layer technologies.

Each authenticator is either a standalone authenticator or pass-
through authenticator [RFC3748]. When an authenticator acts as a
standalone authenticator, it also has the functionality of an EAP
server. When an authenticator acts as a pass-through authenticator,
it communicates with the EAP server typically implemented on a AAA
server using a AAA transport protocol such as RADIUS [RFC2865] and
Diameter [RFC3588].

If the candidate authenticator uses an MSK [RFC5247] for generating
lower-layer ciphering keys, EAP early authentication is used for
proactively generating an MSK for the candidate authenticator.

3.1. Topological Classification of Handover Scenarios

The complexity of the authentication and authorization portion of
handover depends on whether the handover involves a change of
authenticator, and whether it involves a change in EAP Server.
Consider first the case where the authenticators operate in pass-
through mode, so that the EAP Server is a AAA server. Then there is

a strict hierarchy of complexity, as follows:

1. intra-authenticator handover: the candidate and serving
    authenticator are identical. The authenticator can continue to
    use the same keying material. The early authentication problem
    is simply how to recognize this situation.

2. inter-authenticator handover with common AAA server: the
    candidate and serving authenticator are different entities, but
    the AAA server is the same. There are two sub-cases here:

    (a) the AAA server is common because both authenticators lie
         within the same network, or

    (b) the AAA server is common because AAA entities in the serving
         and candidate networks proxy to a AAA server in the home
         domain.

3. inter-AAA-domain handover: the candidate and serving
    authenticator are different entities, and the respective AAA
    servers also differ. As a result, authentication in the
    candidate network requires a second set of user credentials.

A fourth case is where one or both authenticators is collocated with
an EAP Server. This has some of the characteristics of an inter-AAA-
domain handover, but offers less flexibility for resolution of the
early authentication problem.

Orthogonally to this classification, one can distinguish intra-
technology handover from inter-technology handover, thinking of the
link technologies involved. In the inter-technology case, it is
highly probable that the authenticators will differ. The most likely
cases are 2(b) or 3 in the above list.

[4](#). **Early Authentication Usage Models**

As noted in [Section 3](#), there are cases where early authentication is
applicable while ERP does not work. This section concentrates on
providing some usage models around which we can build our analysis of
the EAP early authentication problem. Different usage models can be
defined depending on whether

o the serving authenticator is not involved in early authentication
   (direct pre-authentication usage model),

o the serving authenticator interacts only with the candidate
   authenticator (indirect pre-authentication usage model), or

o the serving authenticator interacts with the AAA server (the
   authenticated anticipatory keying usage model).

It is assumed that the serving and candidate authenticators are
different entities (case 1 of Section 3.1 excluded). It is further
assumed in describing these models that there is no direct L2
connectivity between the peer and a candidate authenticator.

4.1. EAP Pre-authentication Usage Models

In the EAP Pre-authentication usage model, the serving authenticator
does not interact with the AAA server directly. Depending on how the
serving authenticator is involved in the pre-authentication
signaling, the EAP pre-authentication usage model can be further
categorized into the following two submodels.

4.1.1. The Direct Pre-authentication Model

In this model, the serving authenticator is not involved in the EAP
exchange and only forwards the EAP pre-authentication traffic as it
would any other data traffic, or there may be no serving
authenticator at all in the serving access network. This model is
applicable to any of the cases described in Section 3.1 except case
1.

The direct pre-authentication signaling for the usage model is shown
in Figure 3.

```
         Peer                Candidate              AAA Server
                             Authenticator
                             (CA)


         +---------+   +---------------------+   +---------+
         |         |   |      EAP Auth-      |   | EAP     |
         |EAP Peer |   |       enticator     |   | Server  |
         |         |   |                     |   |         |
         +---------+   +---------------------+   +---------+
         |Peer-SA  |   |Peer-CA  | |CA-AAA    |   |CA-AAA   |
         |Signaling<-->|Signaling| |Signaling|<--->Signaling|
         |Layer    |   |Layer    | |Layer    |   |Layer    |
         +---------+   +---------+ +---------+   +---------+
```

           Figure 2: Direct Pre-authentication Usage Model

```
        Peer                Serving              Candidate               AAA/EAP
                          Authenticator        Authenticator             Server
                             (SA)                  (CA)
          |                   |                     |                       |
          |                   |                     |                       |
          |        Peer-CA Signaling (L3)           |        AAA            |
          |<-----------------+------------------->|<----------------->|
          |                   |                     |                       |
          |                   |                     |                       |
```
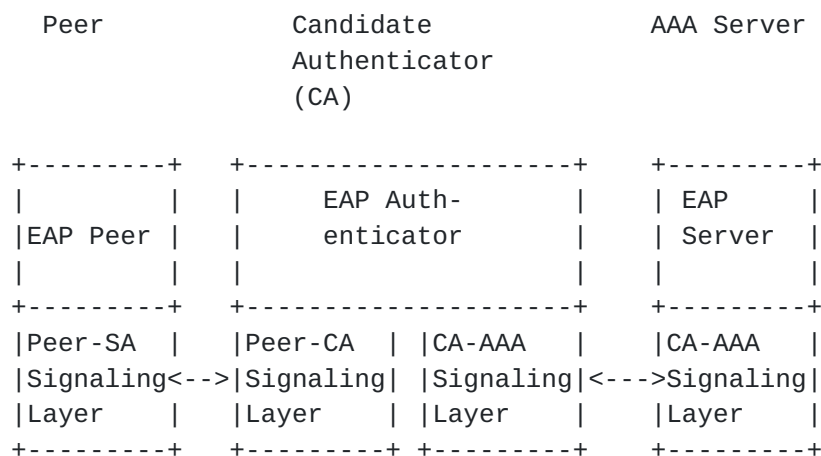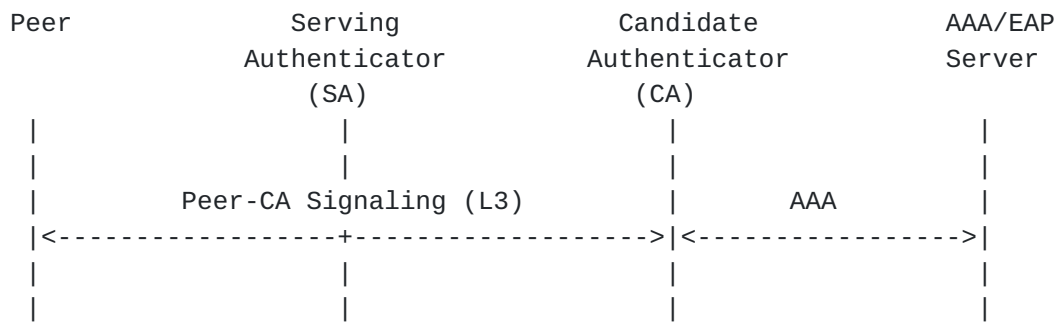
               Figure 3: Direct Pre-authentication Signaling for the Usage Model

4.1.2. The Indirect Pre-authentication Usage Model

   The indirect pre-authentication usage model is illustrated in
   Figure 4

```
        Peer                Serving              Candidate           AAA Server
                          Authenticator        Authenticator
                             (SA)                  (CA)

    +----------+                        +-------------------+  +------+
    |          <- - - - - - - - - - - - ->|                 |  <->|      |
    | EAP Peer |   +-------------------|   |   EAP Auth-     |  |EAP   |
    |          |   |Pre-authentication |   |    enticator    |  |Server|
    |          |   |Forwarding         |   |                 |  |      |
    +----------+   +---------++---------|   +-------------------+  +------+
    | Peer-SA  |   |Peer-SA  ||SA-CA    |   |SA-CA    ||CA-AAA   |  |CA-AAA|
    | Signaling<-->|Signaling||Signaling<-->|Signaling||Signaling<-->Signa-|
    | Layer    |   |Layer    ||Layer    |   |Layer    ||Layer    |  |ling  |
    +----------+   +---------++---------+   +---------++---------+  |Layer |
                                                                   +------+
```
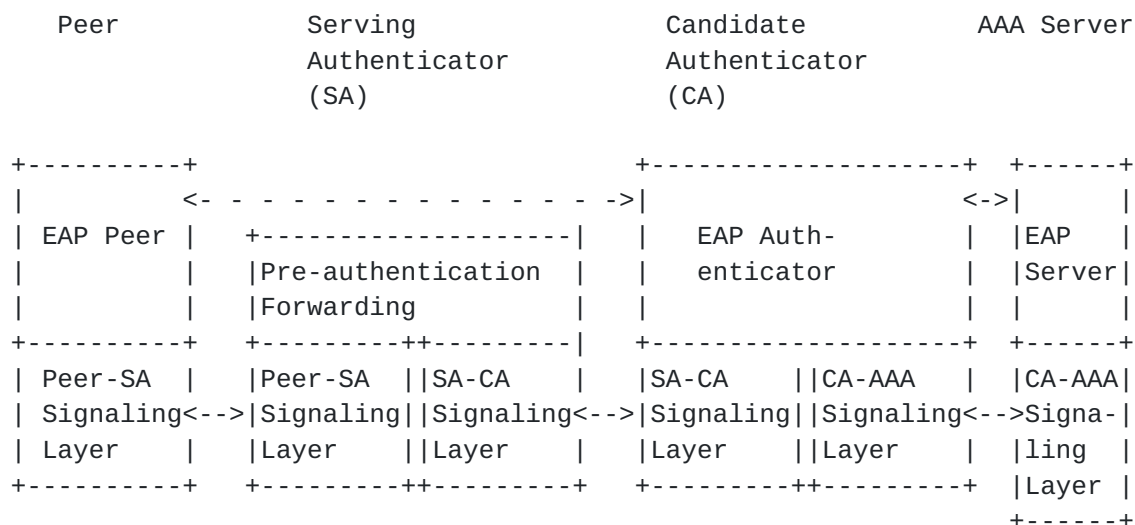
                   Figure 4: Indirect Pre-authentication Usage Model

   In this indirect pre-authentication model, it is assumed that a trust
   relationship exists between the serving network (or serving AAA
   domain) and candidate network (or candidate AAA domain). The serving
   authenticator is involved in EAP pre-authentication signaling. This
   pre-authentication model is needed if the peer cannot discover the
   candidate authenticator's Identity or if IP communication is not
   available due to security or network topology reasons.

   The role of the serving authenticator in this pre-authentication
   model is to forward EAP pre-authentication signaling between the peer
   and candidate authenticator and not to act as an authenticator for
   the candidate point of access. It continues to act as an
   authenticator for the serving point of access. The role of the

candidate authenticator is to forward EAP pre-authentication
signaling between the peer (via the serving authenticator) and EAP
server and receive the transported keying materials from the EAP
server as an authenticator.

The pre-authentication signaling for this model is shown in Figure 5.

```
  Peer                 Serving               Candidate               AAA/EAP
                     Authenticator          Authenticator            Server
                        (SA)                   (CA)
    |                    |                      |                      |
    |                    |                      |                      |
    | Peer-SA Signaling  |   SA-CA Signaling    |        AAA           |
    |    (L2 or L3)      |        (L3)          |                      |
    |<----------------->|<------------------->|<------------------->|
    |                    |                      |                      |
    |                    |                      |                      |
```
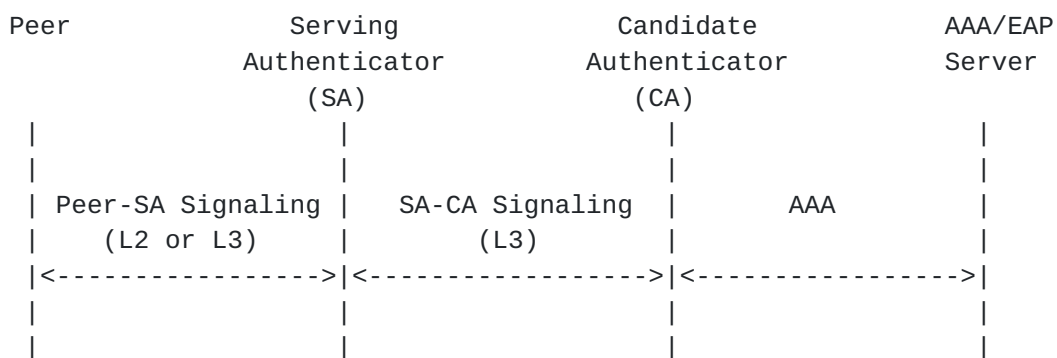
 Figure 5: Indirect Pre-authentication Signaling for the Usage Model

In this model, the pre-authentication signaling path between a peer
and a candidate authenticator consists of two segments: peer to
serving authenticator signaling (Peer-SA signaling) and serving
authenticator to candidate authenticator signaling (SA-CA signaling).

Peer-SA signaling is performed over L2 or L3.

SA-CA signaling is performed over L3.

4.2. The Authenticated Anticipatory Keying Usage Model

In the anticipated authentication keying usage model, the serving
authenticator is required to interact with the AAA server directly.
The authenticated anticipatory keying usage model is illustrated in
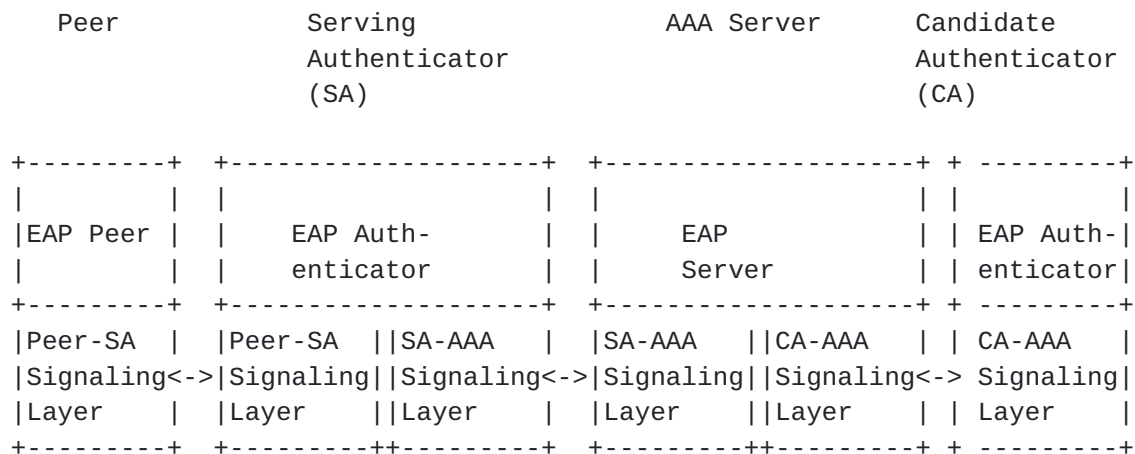Figure 6.

```
      Peer                Serving              AAA Server      Candidate
                       Authenticator                        Authenticator
                           (SA)                                 (CA)


   +---------+  +------------------+  +------------------+ + ---------+
   |         |  |                  |  |                  | |          |
   |EAP Peer |  |    EAP Auth-     |  |      EAP         | | EAP Auth-|
   |         |  |    enticator     |  |     Server       | | enticator|
   +---------+  +------------------+  +------------------+ + ---------+
   |Peer-SA  |  |Peer-SA  ||SA-AAA   |  |SA-AAA   ||CA-AAA   | | CA-AAA   |
   |Signaling<->|Signaling||Signaling<->|Signaling||Signaling<-> Signaling|
   |Layer    |  |Layer    ||Layer    |  |Layer    ||Layer    | | Layer    |
   +---------+  +---------++---------+  +---------++---------+ + ---------+
```

                Figure 6: Authenticated Anticipatory Keying Usage Model

   In this usage model, it is assumed that there is no trust
   relationship between the serving authenticator and the candidate
   authenticator. The serving authenticator is involved in EAP
   authenticated anticipatory keying signaling.

   The role of the serving authenticator in this usage model is to
   communicate with the peer on one side and exchange authenticated
   anticipatory keying signaling with the EAP server on the other side.
   This is not the simple mediation function of an authenticator,
   because the SA-AAA signaling in this case must identify the candidate
   authenticator to which keying material must be pushed. The role of
   the candidate authenticator is to receive the transported keying
   materials from the EAP server and to act as an authenticator after
   handover occurs. The Peer-SA signaling is performed over L2 or L3.
   The SA-AAA and AAA-CA segments operate over L3.

5. **Architectural Considerations**

   There are two architectural issues relating to early authentication:
   authenticator discovery and context binding.

5.1. **Authenticator Discovery**

   In general, early authentication requires the identity of a candidate
   authenticator to be discovered by a peer, by a serving authenticator,
   or by some other entity prior to handover. An authenticator
   discovery protocol is typically defined as a separate protocol from
   an early authentication protocol. For example, the IEEE 802.21
   Information Service (IS) [IEEE.802-21] provides a link-layer-
   independent mechanism for obtaining neighboring network information
   by defining a set of Information Elements (IEs), where one of the IEs

is defined to contain an IP address of a point of attachment. IEEE
802.21 IS queries for such an IE may be used as a method for
authenticator discovery.

If IEEE 802.21 IS or a similar mechanism is used, authenticator
discovery requires a database of information regarding the target
network; the provisioning of a server with such a database is another
issue.

5.2. Context Binding

When a candidate authenticator uses different EAP transport protocols
for normal authentication and early authentication, a mechanism is
needed to bind link-layer-independent context carried over early
authentication signaling to the link-layer-specific context of the
link to be established between the peer and the candidate
authenticator. The link-layer-independent context includes the
identities of the peer and authenticator as well as the MSK. The
link-layer-specific context includes link layer addresses of the peer
and the candidate authenticator. Such context binding can happen
before or after the peer changes its point of attachment.

There are at least two possible approaches to address the context
binding issue. The first approach is based on communicating the link
layer context as opaque data via early authentication signaling. The
second approach is based on running EAP over the link layer of the
candidate authenticator after the peer arrives at the authenticator,
using short-term credentials generated via early authentication. In
this case, the short-term credentials are shared between the peer and
the candidate authenticator. In both approaches, context binding
needs to be securely made between the peer and the candidate
authenticator. Also, the peer is not fully authorized by the
candidate authenticator until the peer completes the link-layer-
specific secure association procedure with the authenticator using
link layer signaling.

6. AAA Issues

Most of the AAA documents today do not distinguish between a normal
authentication and a early authentication and this creates a set of
open issues:

Early authentication authorization
   Users may not be allowed to have more than one logon session at
   the time. This means that while such users actively engage in a
   session (as a result of a previously valid authentication), they
   will not be able to perform early authentication. The AAA server

currently has no way of distinguishing between a normal
authentication request and an early authentication request.

Early authentication lifetime
Currently, AAA protocols define attributes carrying lifetime
information for a normal authentication session. Even when a user
profile and the AAA server support early authentication, the
lifetime for a early authentication session is typically valid
only for a short amount of time because the peer has not completed
its authentication at the target link layer. It is currently not
possible for a AAA server to indicate to the AAA client or a peer
the lifetime of the early authenticated session unless AAA
protocols are extended to carry early authentication session
lifetime information. In other words, it is not clear to the peer
or the authenticator when the early authentication session will
expire.

Early authentication retries
It is typically expected that shortly following the early
authentication process, the peer moves to the new point of
attachment and converts the early authentication state to a normal
authentication state (the procedure for which is not the topic of
this particular subsection). However, if the peer has not yet
moved to the new location and realizes that the early
authentication is expiring, it may perform another early
authentication. Some limiting mechanism is needed to avoid an
unlimited number of early-authentication attempts.

Completion of network attachment
Once the peer has successfully attached to the new point of
attachment, it needs to convert its authentication state from
early authenticated to fully attached and authorized. If the AAA
server needs to differentiate between early authentication and
normal authentication, there may need to be a mechanism within the
AAA protocol to provide this indication to the AAA server. This
may be important from a billing perspective if the billing policy
does not charge for a early authenticated peer until the peer is
fully attached to the target authenticator.

Session resumption
In the case where the peer cycles between a network N1 with which
it has a normal authentication state to another network N2 and
then back to N1, it should be possible to simply convert the full
authentication state to an early authenticated state. The
problems around handling session lifetime and keying material
caching need to be dealt with.

   Multiple candidate authenticators
      There may be situations where the peer needs from among from among
      number of candidate authenticators. In such cases, it is
      desirable for the peer to perform early authentication with
      multiple candidate authenticators. This amplifies the
      difficulties noted under the point "Early authentication
      authorization"

   Inter-domain handover support
      There may be situations where the peer moves out of the home
      domain or across different visited domains, in such cases, the
      early authentication should be performed through the visited AAA
      domain with the AAA server in the home AAA domain. It also
      requires the peer or the authenticator in the visited domain to
      acquire the identity information of the visited domain or the home
      domain for routing the EAP early authentication traffic.
      Knowledge of domain identities is required by both the peer and
      the authenticator to generate the early authentication key for
      mutual authentication between the peer and the visited AAA server.

   Inter-technology support
      Current specifications on early authentication mostly deal with
      homogeneous 802.11 networks. AAA attributes such as Calling-
      Station-ID [I-D.aboba-radext-wlan] may need to be expanded to
      cover other access technologies. Furthermore, inter-technology
      handovers may require a change of the peer identifier as part of
      the handover. Investigation on the best type of identifiers for
      peers that support multiple access technologies is required.

7. Security Considerations

   This section specifically covers threats introduced to the EAP model
   by early authentication. Security issues on general EAP and handover
   are described in other documents such as [RFC3748], [RFC4962],
   [RFC5169] and [RFC5247].

   Since early authentication described in this document needs to work
   across multiple authenticators, any solution needs to consider the
   following security threats.

   First, a resource consumption denial of service attack is possible,
   where an attacker that is not on the same IP link as the legitimate
   peer or the candidate authenticator may send unprotected early
   authentication messages to the legitimate peer or the candidate
   authenticator. As a result, the latter may spend computational and
   bandwidth resources on processing early authentication messages sent
   by the attacker. This attack is possible for both direct and

indirect pre-authentication scenarios. To mitigate this attack, the
candidate network or authenticator may apply non-cryptographic packet
filtering so that early authentication messages received from only a
specific set of serving networks or authenticators are processed. In
addition, a simple solution for the peer side would be to let the
peer always initiate EAP early authentication and not allow EAP early
authentication initiation from an authenticator.

Second, consideration for the channel binding problem described in
[RFC5247] is needed as lack of channel binding may enable an
authenticator to impersonate another authenticator or communicate
incorrect information via out-of-band mechanisms (such as via a AAA
or lower layer protocol) [RFC3748]. It should be noted that it is
relatively easier to launch such an impersonation attack for early
authentication than normal authentication because an attacker does
not need to be physically on the same link as the legitimate peer to
send a early authentication trigger to the peer.

8. IANA Considerations

This document makes no requests for IANA action.

9. Acknowledgments

The authors would like to thank Bernard Aboba, Jari Arkko, Ajay
Rajkumar, Maryna Komarova, Charles Clancy, Subir Das, Shubhranshu
Singh, Preetida Vinayakray and Rafa Marin Lopez for their valuable
input.

10. Contributors

The following people contributed to this document: Ashutosh Dutta,
Srinivas Sreemanthula, Alper E. Yegin, Madjid Nakhjiri, Mahalingam
Mani and Tom Taylor.

11. References

11.1. Normative References

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

[RFC4962]   Housley, R. and B. Aboba, "Guidance for Authentication,

Authorization, and Accounting (AAA) Key Management",
BCP 132, RFC 4962, July 2007.

[RFC5247]   Aboba, B., Simon, D., and P. Eronen, "Extensible
Authentication Protocol (EAP) Key Management Framework",
RFC 5247, August 2008.

11.2. Informative References

[RFC2865]   Rigney, C., Willens, S., Rubens, A., and W. Simpson,
"Remote Authentication Dial In User Service (RADIUS)",
RFC 2865, June 2000.

[RFC3588]   Calhoun, P., Loughney, J., Guttman, E., Zorn, G., and J.
Arkko, "Diameter Base Protocol", RFC 3588, September 2003.

[RFC5169]   Clancy, T., Nakhjiri, M., Narayanan, V., and L. Dondeti,
"Handover Key Management and Re-Authentication Problem
Statement", RFC 5169, March 2008.

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

[I-D.aboba-radext-wlan]
Aboba, B., Malinen, J., Congdon, P., and J. Salowey,
"RADIUS Attributes for IEEE 802 Networks",
draft-aboba-radext-wlan-11 (work in progress), April 2009.

[IEEE.802-21]
"Draft Standard for Local and Metropolitan Area Networks:
Media Independent Handover Services", IEEE , 2008.

[IEEE.802-11.2007]
"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", IEEE Standard 802.11, 2007, <
http://standards.ieee.org/getieee802/download/
802.11-2007.pdf>.

[IEEE.802-11R.2008]
"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 2: Fast BSS
Transition", IEEE Standard 802.11R, 2008, <http://

                standards.ieee.org/getieee802/download/802.11r-2008.pdf>.

    [IEEE.802-11F.2003]
                "IEEE Trial-Use Recommended Practice for Multi-Vendor
                Access Point Interoperability via an Inter-Access Point
                Protocol Across Distribution Systems Supporting IEEE
                802.11 Operation", IEEE Recommendation 802.11F, 2003, <htt
                p://standards.ieee.org/getieee802/download/
                802.11F-2003.pdf>.

    [TS33.402]
                3GPP, "System Architecture Evolution (SAE):Security
                aspects of non-3GPP accesses (Release 8)", 3GPP TS33.402,
                V8.3.1 , 2009.

    [ITU]   ITU-T, "General Characteristics of International Telephone
                Connections and International Telephone Circuits: One-Way
                Transmission Time", ITU-T Recommendation G.114 , 1998.

    [WPA]   The Wi-Fi Alliance, "WPA (Wi-Fi Protected Access)", Wi-
                Fi WPA v3.1, 2004.

    [MQ7]   Lopez, R., Dutta, A., Ohba, Y., Schulzrinne, H., and A.
                Skarmeta, "Network-layer Assisted Mechanism to Optimize
                Authentication Delay During Handoff in 802.11 Networks",
                The 4th Annual International Conference on Mobile and
                Ubiquitous Systems: Computing, Networking and Services
                (MOBIQUITOUS 2007) , 2007.

    [WCM]   Dutta, A., Famorali, D., Das, S., Ohba, Y., and R. Lopez,
                "Media-independent pre-authentication supporting secure
                interdomain handover optimization", IEEE Wireless
                Communications Volume 15, Issue 2, April 2008.

Authors' Addresses

    Yoshihiro Ohba (editor)
    Toshiba America Research, Inc.
    1 Telcordia Drive
    Piscataway, NJ 08854
    USA

    Phone: +1 (732) 699-5365
    Email: yohba@tari.toshiba.com

   Qin Wu (editor)
   Huawei Technologies Co.,Ltd
   SiteB, Floor 12F,Huihong Mansion, No.91.,Baixia Rd.
   Nanjing, JiangSu 210001
   PRC

   Phone: +86 2584565892
   Email: sunseawq@huawei.com

   Glen Zorn (editor)
   Network Zen
   1310 East Thomas Street
   Seattle, Washington 98102
   US

   Phone: +1 (206) 377-9035
   Email: gwz@net-zen.net