Network Working Group
Internet-Draft
Expires: August 25, 2008

Y. Ohba (Editor)
Toshiba
February 22, 2008

EAP Pre-authentication Problem Statement
draft-ietf-hokey-preauth-ps-02

Status of this Memo

Copyright Notice

Abstract

EAP pre-authentication is defined as the utilization of EAP to pre-establish EAP keying material on an authenticator prior to arrival of the peer at the access network managed by that authenticator. This draft discusses EAP pre-authentication problems in details.

Table of Contents

1. Contributors

The following people contributed to this document.

Yoshihiro Ohba (yohba@tari.toshiba.com)

Ashutosh Dutta (adutta@research.telcordia.com)

Srinivas Sreemanthula (srinivas.sreemanthula@nokia.com)

Alper E. Yegin (alper.yegin@yegin.org)

Madjid Nakhjiri (madjid.nakhjiri@motorola.com)

Mahalingam Mani (mmani@avaya.com)

2. Introduction

When a mobile during an active communication session moves from one access network to another access network and changes its point of attachment it is subjected to disruption in the continuity of service because of the associated handover operation. During the handover process, when the mobile changes its point-of-attachment in the network, it may change its subnet or administrative domain it is connected to. We provide in Appendix A some performance requirement that are needed to support an interactive real-time communication such as VoIP and thus can serve as the guidelines for handover optimization.

Handover often requires authentication and authorization for acquisition or modification of resources assigned to a mobile and the authorization needs interaction with a central authority in a domain. In many cases an authorization procedure during a handover procedure follows an authentication procedure that also requires interaction with a central authority in a domain. The delay introduced due to such an authentication and authorization procedure adds to the handover latency and consequently affects the ongoing multimedia sessions [MQ7]. The authentication and authorization procedure may include EAP authentication [RFC3748] where an AAA server may be involved in EAP messaging during the handover. Depending upon the type of architecture, in some cases the AAA signals traverse all the way to the AAA server in the home domain of the mobile as well before the network service is granted to the mobile in the new network.

Real-time communication and interactive traffic such as VoIP is very sensitive to the delay. Thus it is desirable that interactions between the mobile and AAA servers must be avoided or be reduced

during the handover.

This draft discusses EAP pre-authentication problems in details where EAP pre-authentication is defined as the utilization of EAP to pre-establish EAP keying material on an authenticator prior to arrival of the peer at the access network served by that authenticator.

2.1. Specification of Requirements

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

Basic mechanism of handover is a three-step procedure involving i) discovery of potential points of attachment and their authenticators, ii) network selection procedure to determine the appropriate candidate network point of attachment and iii) handover or setting up of L2 and L3 connectivity to the target network point of attachment. Currently, security mechanisms for authentication and authorization are performed as part of the third step directly with the target network. For example, in basic IEEE 802.11 based wireless networks, the security mechanism involves performing a new IEEE 802.1X message exchange with the authenticator in the target AP (Access Point) to initiate an EAP exchange to the authentication server [WPA]. Following a successful authentication, a secure association protocol named four-way handshake with the wireless station derives a new set of the session keys for use in data communications. Unless PMK (Pairwise Master Key) is not cached in the target AP, this mechanism is same as the initial setup to the AP with no particular optimizations for the handover scenario. The handover latency introduced by this security mechanism has proven to be larger than what is acceptable for some handover scenarios [MQ7]. Hence, improvement in the handover latency performance due to security procedures is a necessary objective for such scenarios.

For example, if a mobile only needs 250 ms for "fast reconnect" then if it is moving at 60 mph (87 feet/second), then the mobile will have moved roughly 22 feet during the EAP authentication process. This is larger than the average coverage overlap of a wireless LAN (WLAN).

There is relevant work undertaken by various standards organizations. But these efforts are scoped to a specific access technology. IEEE 802.11f has defined context transfer between APs. IEEE 802.11i

defines a pre-authentication mechanism for use in 802.11 variant wireless networks. This mechanism allows mobile devices to pre-authenticate using EAP to one or more candidate authenticators over the wired medium, by way of the serving authenticator. IEEE 802.11r [802.11r] defines Fast BSS transition mechanisms involving a definition of key management hierarchy and setup of session keys before the re-association to the target AP. These mechanisms, as indicated before, are defined for IEEE 802.11 technologies and are only applicable within a certain access domain and fall short when it comes to inter-access technology handovers. They also require L2 (e.g., Ethernet) connectivity for transfer of encapsulated signaling to a candidate or the target AP. Especially, a solution is needed to enable EAP pre-authentication in IEEE 802.11 to work even if the station and AP are not members of the same VLAN.

As various flavors of wireless technologies are increasingly available, there is a growing demand for seamless inter-access technology mobility and handovers. This is particularly beneficial in the presence of high bandwidth wireless technologies (e.g., IEEE 802.11a/b/g) with only hotspot like coverages while the overlay licensed wireless/cellular coverages are expensive and relatively lower bandwidth. There is a strong motivation to allow seamless inter-technology handovers for all kinds of data communications. Hence, the security optimization mechanisms for better handover performance must be looked at from the IP level so as to make it a common method over different access technologies.

Solutions for inter-authenticator mobility security optimizations can be largely seen as security context transfer, handover keying or EAP pre-authentication. Security context transfer involves transfer of reusable key context in the new point of attachment. However, the recent AAA key management requirement [RFC4962] does not recommend horizontal context transfer of reusable key context because of domino effect in which a compromise of an authenticator will lead to a compromise of another authenticator. Handover keying and re-authentication [I-D.ietf-hokey-reauth-ps] uses an existing EAP-generated key for deriving a re-authentication key to be distributed to a HOKEY server in a visited domain in order to reduce the handover delay, which eliminates the need for running a full EAP authentication with the EAP server in the home domain for handovers within the visited domain. On the other hand, there are certain cases where an EAP-generated key does not exist or is not usable for handover keying at the time of handover and an EAP run is not avoidable to generate a key for the candidate authenticator. One case is an inter-domain handover without any trust relationship between domains. Another case is an intra-domain handover where the access networks and/or the AAA infrastructure in the visited domain do not support handover keying and low-latency re-authentication.

EAP pre-authentication discussed in this document is mainly to deal with an environment where the mobile device and candidate authenticators are not in the same subnet or of the same link-layer technology. Such use of EAP pre-authentication would enable the mobile device to authenticate and setup keys prior to connecting to one of the candidate authenticators.

This framework has general applicability to various deployment scenarios in which proactive signaling can take effect. In other words, applicability of EAP pre-authentication is limited to the scenarios where candidate authenticators can be easily discovered, an accurate prediction of movement can be easily made. Also the effectiveness of EAP pre-authentication may be less significant for particular inter-technology handover scenarios where simultaneous use of multiple technologies is not a major concern or where there is sufficient radio-coverage overlap among different technologies.

Note that EAP pre-authentication problem for intra-technology intra-subnet handover could be solved by each link-layer and is thus out of the scope of this document while a general solution developed at IETF can be used for intra-technology and intra-subnet scenarios as well.

In EAP pre-authentication, AAA authentication and authorization for a candidate authenticator is performed while ongoing data communications are in progress via the serving network. The goal of EAP pre-authentication is to avoid AAA signaling for EAP when or soon after the device moves. There are several AAA issues related to EAP pre-authentication. The pre-authentication AAA issues are described in Section 6.

Figure 1 shows the functional elements that are related to EAP pre-authentication.

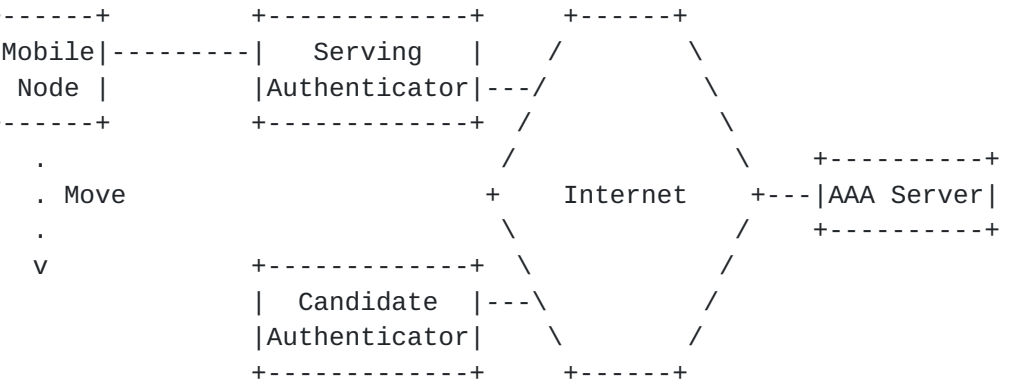

Figure 1: EAP Pre-authentication Functional Elements

A mobile node is attached to the serving access network. Before the

mobile node performs handover from the serving access network to a candidate access network, it performs EAP pre-authentication with a candidate authenticator, an authenticator in the candidate access network, via the serving access network. The mobile node may perform EAP pre-authentication with one or more candidate authenticators. It is assumed that each authenticator has an IP address when authenticators are on different IP links. 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 has the functionality of EAP authenticator which is either standalone EAP authenticator or pass-through EAP authenticator. When an authenticator acts as a standalone EAP authenticator, it also has the functionality of EAP server. On the other hand, when an authenticator acts as a pass-through EAP authenticator, it communicates with EAP server typically implemented on a AAA server using a AAA protocol such as RADIUS and Diameter.

If the candidate authenticator is of an existing link-layer technology that uses an MSK (Master Session Key) [I-D.ietf-eap-keying] for generating lower-layer ciphering keys, EAP pre-authentication is used for proactively generating the MSK for the candidate authenticator.

4. Usage Scenarios

There are two scenarios on how EAP pre-authentication signaling can happen among a mobile node, a serving authenticator, a candidate authenticator and a AAA server, depending on how the serving authenticator is involved in the EAP pre-authentication signaling. No security association between the serving authenticator and the candidate authenticator is required for both pre-authentication scenarios (see Section 7 for more detailed discussion).

4.1. Direct Pre-authentication

Direct pre-authentication signaling is shown in Figure 2.

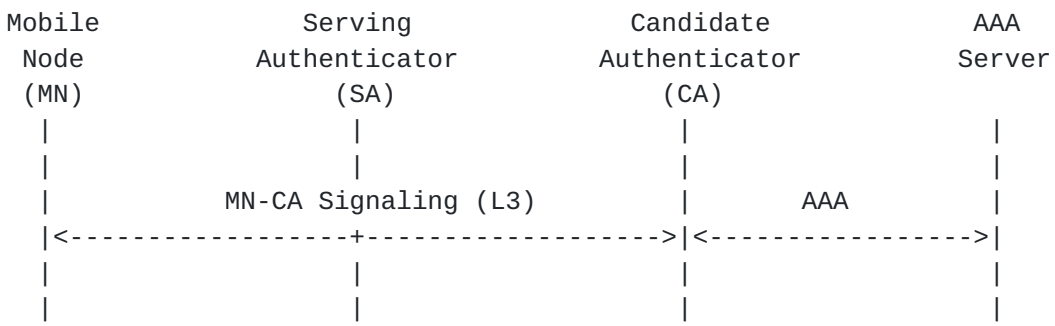

Figure 2: Direct Pre-authentication

In this type of pre-authentication, the serving authenticator 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.

[I-D.ietf-pana-preauth] is identified as a protocol to realize direct pre-authentication.

4.2. Indirect Pre-authentication

Indirect pre-authentication signaling is shown in Figure 3.

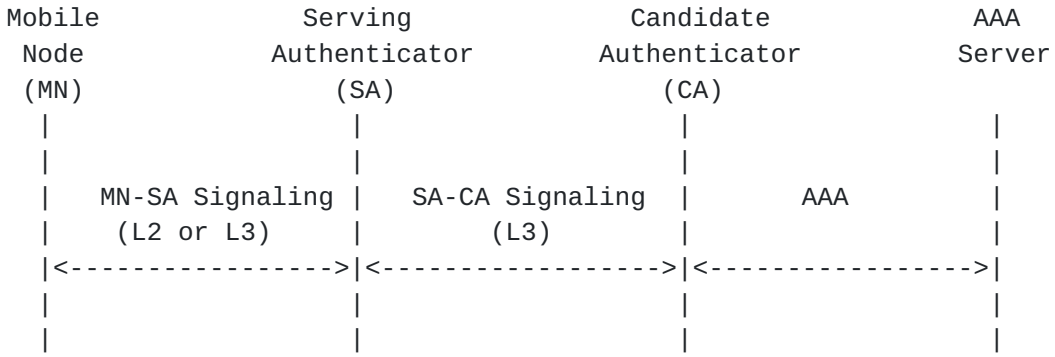

Figure 3: Indirect Pre-authentication

With indirect pre-authentication, the serving authenticator is involved in EAP pre-authentication signaling. Indirect pre-authentication is needed if the MN cannot discover the CA's IP address or if IP communication is not allowed between the candidate authenticator and unauthorized nodes for security reasons.

Indirect pre-authentication signaling is spliced into mobile node to serving authenticator signaling (MN-SA signaling) and serving authenticator to candidate authenticator signaling (SA-CA signaling).

SA-CA signaling is performed over L3.

MN-SA signaling is performed over L2 or L3.

The role of the serving authenticator in indirect pre-authentication is to forward EAP pre-authentication signaling between the mobile node and the candidate authenticator and not to act as an EAP authenticator, while it acts as an EAP authenticator for normal authentication signaling. This is illustrated in Figure 4.

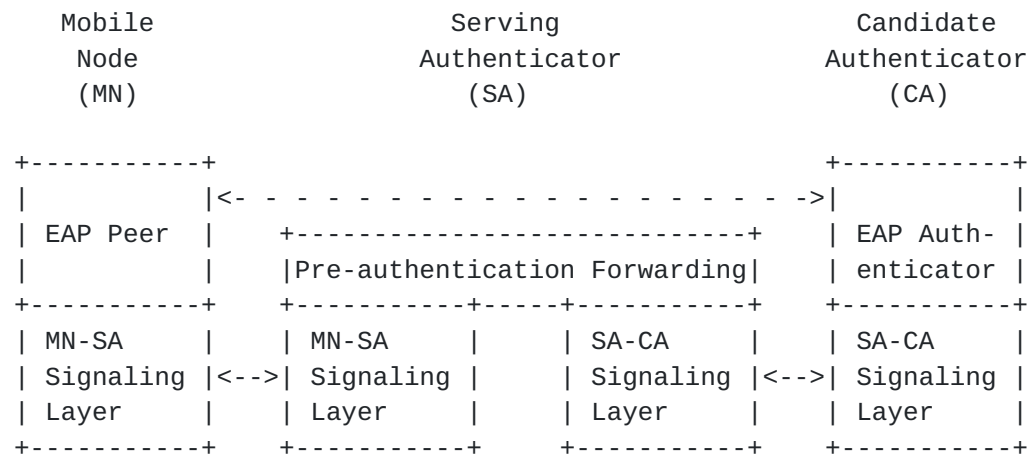

Figure 4: Indirect Pre-authentication Layering Model

5. Architectural Considerations

There are two architectural issues relating to pre-authentication, i.e., authenticator discovery and context binding.

5.1. Authenticator Discovery

In general, pre-authentication requires an address of a candidate authenticator to be discovered either by a mobile node or by a serving authenticator prior to handover. An authenticator discovery protocol is typically defined as a separated protocol from a pre-authentication protocol. When pre-authentication is used for inter-technology or inter-subnet handover, a candidate authenticator needs to have a global IP address and a mechanism for discovering the candidate authenticators IP address is needed. For example, IEEE 802.21 Information Service (IS) [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.

An authenticator discovery mechanism requires a database on the

neighboring network information. 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 pre-authentication, a mechanisms is needed to bind link-layer independent context carried over pre-authentication signaling to the link-layer specific context of the link to be established between the mobile node 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 mobile node and the candidate authenticator.

There are two possible approaches to address the context binding issue. The first approach is based on communicating the lower-layer context as opaque data via pre-authentication signaling and perform the link-layer specific secure association procedure after handover. The second approach is based on running EAP over the link-layer of the candidate authenticator after handover using short-term credentials generated via pre-authentication, followed by the link-layer specific secure association procedure. In this case, the short-term credentials are shared between the mobile node and the candidate authenticator, and hence the EAP server for the post-handover EAP resides in the candidate authenticator. In both approaches, the binding needs to be securely made between the peer and the candidate authenticator using a security association established via pre-authentication.

6. AAA Issues

Most of the AAA documentations today do not distinguish between a full authentication and a pre-authentication and this creates a set of open issues:

Pre-authentication authorization: Many users may not be allowed to have more than one logon session at the time. This means, when such users actively engage in an active session (as a result of a previously valid authentication), they will not be able to perform pre-authentication. The AAA server currently has no way of distinguishing between a full authentication request and a pre-authentication request.

Pre-authentication life time: Currently AAA protocols define attributes (AVPs) carrying life time information for a full authentication session. Even when a user profile and the AAA server support pre-authentication function, after the pre-authentication of a peer is complete, since the pre-authentication may be accompanied with a pre-authorization, the pre-authentication is typically valid only for a short amount of time. It is currently not possible for a AAA server to indicate to the AAA client or a peer what the life time of the pre- authenticated session is. In other words, it is not clear to the peer or the NAS, when the pre-authentication will expire.

Pre-authentication retries: It is typically expected that shortly following the pre-authentication process, the mobile entity moves to the new point of attachment and converts the pre-authentication state to a full authentication state (the procedure for which is not the topic of this particular subsection). However, if the peer has yet not moved to the new location and realizes that the pre-authentication is expiring, it may perform another pre-authentication. In order to avoid unlimited number of pre-authentication tries, it is quite possible that the network policy sets a limit on the maximum number of pre-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 pre-authenticated to fully attached and authorized. There may need to be a mechanism within the AAA protocol to provide this indication to the AAA server.

Session Resumption: In case the peer ping pongs between a network N1 with which it has a full authentication state to another network N2 and then back to N1, it should be possible to simply convert the full authentication state to a pre-authenticated state. The problems around handling session life time and keying material caching needs to be dealt with.

Multiple candidate authenticators: There may be situations where the mobile node may need to make a selection between a number of candidate attachment points. In such cases, it is desirable for the mobile to perform pre-authentication with multiple authenticators. In such cases the AAA server may need to be aware of the situation.

Roaming support: In case the pre-authentication is being performed through a serving network that is foreign to the MN's home AAA server, the AAA server needs to obtain the information about the serving network in addition to the information about the candidate network, so that the AAA server can make authorization decisions

accordingly, e.g., depending on the authorization policy, the home AAA server may not allow pre-authentication via a particular serving network.

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

Network controlled handovers: It is becoming quite common for the network operators to maintain the control over the handovers for various reasons including load balancing and performance. Hence the network may need to direct the MN to perform pre-authentication to a set of candidate authenticators in an anticipation for a handover. The AAA protocol extensions for carrying out such procedures need to be provided.

7. Security Considerations

Since pre-authentication described in this document needs to work across multiple authenticators, any solution for this problem needs considerations on the following security threats.

First, a possible resource consumption denial of service attack where an attacker that is not on the same IP link as the mobile node or the candidate authenticator may send unprotected pre-authentication messages to the mobile node or the candidate authenticator to let the legitimate mobile node and candidate authenticator spend their computational and bandwidth resources. This attack is possible for both direct and indirect pre-authentication scenarios. To mitigate this attack, the candidate network or authenticator should apply non-cryptograhpic packet filtering so that pre-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 pre-authentication and not allow EAP pre-authentication initiation from authenticator side.

Second, consideration for the Channel Binding problem described in [I-D.ietf-eap-keying] 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, when

normal authentication uses link-layer EAP transport, it would be easier to launch such an impersonation attack for pre-authentication than normal authentication because an attacker does not need to be physically on the same link as the legitimate peer to send a pre-authentication trigger to the peer.

8. IANA Considerations

This document has no actions for IANA.

9. Acknowledgments

The authors would like to thank Bernard Aboba, Jari Arkko, Ajay Rajkumar and Maryna Komarova for their valuable input.

10. References

10.1. 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. 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.

[I-D.ietf-eap-keying] Aboba, B., Simon, D., and P. Eronen, "Extensible Authentication Protocol (EAP) Key Management Framework", draft-ietf-eap-keying-22 (work in progress), November 2007.

10.2. Informative References

[I-D.ietf-hokey-reauth-ps] Clancy, C., Nakhjiri, M., Narayanan, V., and L. Dondeti, "Handover Key Management and Re-authentication Problem Statement", draft-ietf-hokey-reauth-ps-08 (work in progress), February 2008.

[I-D.aboba-radext-wlan]

Malinen, J. and B. Aboba, "RADIUS Attributes for IEEE 802 Networks", draft-aboba-radext-wlan-06 (work in progress), July 2007.

[I-D.ietf-pana-preauth]
Ohba, Y., "Pre-authentication Support for PANA", draft-ietf-pana-preauth-02 (work in progress), November 2007.

[802.21] IEEE, "Draft Standard for Local and Metropolitan Area Networks: Media Independent Handover Services", LAN MAN Standards Committee of the IEEE Computer Society 802.21 D9.0 2008.

[802.11r] IEEE, "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", LAN MAN Standards Committee of the IEEE Computer Society 802.11r D9.0 2008.

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

[ETSI] ETSI, "Telecommunications and Internet Protocol Harmonization Over Networks (TIPHON) Release 3: End-to-end Quality of Service in TIPHON systems; Part 1: General aspects of Quality of Service.", ETSI TR 101 329-6 V2.1.1.

[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", ACM Mobiquitous 2007.

Appendix A. Performance Requirements

In order to provide the desirable quality of service for interactive VoIP and streaming traffic during handoff, one needs to limit the value of end-to-end delay, jitter and packet loss to a certain threshold level. ITU-T and ITU-R standards define the acceptable values for these parameters. For example for one-way delay, ITU-T G.114 [ITU] recommends 150 ms as the upper limit for most of the

applications, and 400 ms as generally unacceptable delay. One way delay tolerance for video conferencing is in the range of 200 to 300 ms. Also if an out-of-order packet is received after a certain threshold, it is considered lost. The performance requirement will vary based on the type of application and its characteristics such as delay tolerance and loss tolerance limit. Interactive traffic such as VoIP and streaming traffic will have different tolerance for delay and packet loss. For example, according to ETSI TR 101 [ETSI] a normal voice conversation can tolerate up to 2% packet loss. Similarly there are other factors such as Transmission Rating Factor (R) standardized within ITU-T G.107, End to End delay (one way mouth-to-ear) and call blocking ratio that determine the QoS metrics. An R value of 50 is considered to be poor and a value of 90 can be considered as the best that provides most user satisfaction. As an example, a class B QoS which is equivalent to cellular telephony has a R factor that is greater than 70, E2E delay of less than 150 ms and call blocking ratio which is less than or equal to 0.15. Class A QoS that is the highest and is equivalent to fixed phone quality has an R value that is more than 80 and an end-to-end delay that is less than 100 ms. Similarly, 3GPP TS23.107 defines 4 application classes: conversational, streaming, interactive and background each with different set of end-to-end delay and QoS requirement. The streaming class has the tolerable packet (SDU) error rates ranging from 0.1 to 0.00001 and the transfer delay of less than 300ms. In short, the delay and packet loss tolerance value will depend upon the type of application and different standard bodies and vendors provide different specification for each type of application and thus any optimized handoff mechanism will need to take these values into consideration.

It is desirable to support a heterogeneous handover that is agnostic to link-layer technologies in an optimized and secure fashion without incurring unreasonable complexity while providing seamless handover experience to the user. As a mobile goes through a handover process, it is subjected to handover delay because of the rebinding of properties at several layers of the protocol stack, such as layer 2, layer 3 and application layer. There are several common properties that contribute to the re-establishment or modification of these layers during handover. These properties can mostly be attributed to things such as access characteristics (e.g., bandwidth, channel characteristics, channel scan, access point association), physical-layer access methods (e.g., CDMA, TDMA), MAC layer protocols (e.g., CSMA/CA), configuration of layer 3 parameters such as IP address acquisition, re-authentication, re-authorization, rebinding of security association at all layers, binding update etc. Although each of the components during the handover process that contributes to the handover delay needs to be optimized, we focus our discussion on optimizing the delay due to authentication and authorization.

Author's Address

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

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

Full Copyright Statement

Intellectual Property

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).