

Network Working Group
Request for Comments: 4058
Category: Informational

A. Yegin, Ed.
Samsung AIT
Y. Ohba
Toshiba
R. Penno
Juniper Networks
G. Tsirtsis
Flarion
C. Wang
ARO/NCSU
May 2005

Protocol for Carrying Authentication for Network Access (PANA)
Requirements

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2005).

Abstract

It is expected that future IP devices will have a variety of access technologies to gain network connectivity. Currently there are access-specific mechanisms for providing client information to the network for authentication and authorization purposes. In addition to being limited to specific access media (e.g., 802.1X for IEEE 802 links), some of these protocols are limited to specific network topologies (e.g., PPP for point-to-point links). The goal of this document is to identify the requirements for a link-layer agnostic protocol that allows a host and a network to authenticate each other for network access. This protocol will run between a client's device and an agent in the network where the agent might be a client of the AAA infrastructure.

Table of Contents

1. Introduction	3
2. Requirements Notation	3
3. Terminology	4
4. Requirements	4
4.1. Authentication	4
4.1.1. Authentication of Client	4
4.1.2. Authorization, Accounting, and Access Control	6
4.1.3. Authentication Backend	7
4.1.4. Identifiers	7
4.2. IP Address Assignment	7
4.3. EAP Lower Layer Requirements	7
4.4. PAA-to-EP Protocol	8
4.5. Network	8
4.5.1. Multi-access	8
4.5.2. Disconnect Indication	8
4.5.3. Location of PAA	9
4.5.4. Secure Channel	9
4.6. Interaction with Other Protocols	10
4.7. Performance	10
4.8. Congestion Control	10
4.9. IP Version Independence	10
4.10. Denial of Service Attacks	10
4.11. Client Identity Privacy	10
5. Security Considerations	11
6. Acknowledgements	11
A. Problem Statement	12
B. Usage Scenarios	13
References	16
Normative References	16
Informative References	16

1. Introduction

Secure network access service requires access control based on the authentication and authorization of the clients and the access networks. Initial and subsequent client-to-network authentication provides parameters that are needed to police the traffic flow through the enforcement points. A protocol is needed to carry authentication parameters between the client and the access network. See Appendix A for the associated problem statement.

The protocol design will be limited to defining a messaging protocol (i.e., a carrier) that will allow authentication payload to be carried between the host/client and an agent/server in the access network for authentication and authorization purposes regardless of the AAA infrastructure that may (or may not) reside on the network. As a network-layer protocol, it will be independent of the underlying access technologies and applicable to any network topology.

The intent is not to invent new security protocols and mechanisms but to reuse existing mechanisms such as EAP [RFC3748]. In particular, the requirements do not mandate the need to define new authentication protocols (e.g., EAP-TLS [RFC2716]), key distribution or key agreement protocols, or key derivation methods. The desired protocol can be viewed as the front-end of the AAA protocol or any other protocol/mechanisms the network is running at the background to authenticate its clients. It will act as a carrier for an already defined security protocol or mechanism.

An example of a protocol being extended for use in authenticating a host for network access is Mobile IPv4. A Mobile IPv4 registration request message is used as a carrier for authentication extensions (MN-FA [RFC3344] or MN-AAA [RFC3012]) that allows a foreign agent to authenticate mobile nodes before providing forwarding service. The goal of PANA is similar in that it aims to define a network-layer transport for authentication information. However, PANA will be decoupled from mobility management and will rely on other specifications for the definition of authentication payloads.

This document defines common terminology and identifies requirements of a protocol for PANA that will be used to define and limit the scope of the work to be done in this group.

2. Requirements Notation

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. Terminology

PANA Client (PaC)

The client side of the protocol that resides in the host device which is responsible for providing the credentials to prove its identity for network access authorization.

PANA Client Identifier (PaCI)

The identifier that is presented by the PaC to the PAA for network access authentication. A simple username and NAI [RFC2794] are examples of PANA client identifiers.

Device Identifier (DI)

The identifier used by the network as a handle to control and police the network access of a client. Depending on the access technology, this identifier might contain an IP address, a link-layer address, or a switch port number, etc. of a connected device.

PANA Authentication Agent (PAA)

The access network side entity of the protocol whose responsibility is to verify the credentials provided by a PANA client and grant network access service to the device associated with the client and identified by a DI.

Enforcement Point (EP)

A node on the access network where per-packet enforcement policies (i.e., filters) are applied on the inbound and outbound traffic of client devices. Information such as DI and (optionally) cryptographic keys are provided by PAA per client for constructing filters on the EP.

4. Requirements

4.1. Authentication

4.1.1. Authentication of Client

PANA MUST enable authentication of PaCs for network access. A PaC's identity can be authenticated by verifying the credentials (e.g., identifier, authenticator) supplied by one of the users of the device or the device itself. PANA MUST only grant network access service to the device identified by the DI, rather than separate access to

multiple simultaneous users of the device. Once network access is granted to the device, methods used by the device on arbitrating which user can access the network is outside the scope of PANA.

PANA MUST NOT define new security protocols or mechanisms. Instead, it MUST be defined as a "carrier" for such protocols. PANA MUST identify which specific security protocol(s) or mechanism(s) it can carry (the "payload"). EAP is a candidate protocol that satisfies many requirements for authentication. PANA would be a carrier protocol for EAP. If the PANA Working Group decides that extensions to EAP are needed, it will define requirements for the EAP WG instead of designing such extensions.

Providing authentication, integrity and replay protection for data traffic after a successful PANA exchange is outside the scope of this protocol. In networks where physical layer security is not present, link-layer or network-layer ciphering (e.g., IPsec) can be used to provide such security. These mechanisms require the presence of cryptographic keying material at PaC and EP. Although PANA does not deal with key derivation or distribution, it enables this by carrying EAP and allowing appropriate EAP method selection. Various EAP methods are capable of generating basic keying material that cannot be directly used with IPsec because it lacks the properties of an IPsec SA (security association) including secure cipher suite negotiation, mutual proof of possession of keying material, freshness of transient session keys, key naming, etc. These basic (initial) EAP keys can be used with an IPsec key management protocol, like IKE, to generate the required security associations. A separate protocol, called secure association protocol, is required to generate IPsec SAs based on the basic EAP keys. This protocol MUST be capable of enabling IPsec-based access control on the EPs. IPsec SAs MUST enable authentication, integrity and replay protection of data packets as they are sent between the EP and PaC.

Providing a complete secure network access solution by securing router discovery [RFC1256], neighbor discovery [RFC2461], and address resolution protocols [RFC826] is outside the scope as well.

Some access networks might require or allow their clients to get authenticated and authorized by the network access provider (NAP) and ISP before the clients gain network access. NAP is the owner of the access network who provides physical and link-layer connectivity to the clients. PANA MUST be capable of enabling two independent authentication operations (i.e., execution of two separate EAP methods) for the same client. Determining the authorization parameters as a result of two separate authentications is an operational issue and therefore outside the scope of PANA.

Both the PaC and the PAA MUST be able to perform mutual authentication for network access. Providing only the capability of a PAA authenticating the PaC is not sufficient. Mutual authentication capability is required in some environments but not in all of them. For example, clients might not need to authenticate the access network when physical security is available (e.g., dial-up networks).

PANA MUST be capable of carrying out both periodic and on-demand re-authentication. Both the PaC and the PAA MUST be able to initiate both the initial authentication and the re-authentication process.

Certain types of service theft are possible when the DI is not protected during or after the PANA exchange [RFC4016]. PANA MUST have the capability to exchange DI securely between the PaC and PAA where the network is vulnerable to man-in-the-middle attacks. While PANA MUST provide such a capability, its utility relies on the use of an authentication method that can generate keys for cryptographic computations on PaC and PAA.

4.1.2. Authorization, Accounting, and Access Control

After a device is authenticated by using PANA, it MUST be authorized for "network access." That is, the core requirement of PANA is to verify the authorization of a PaC so that PaC's device may send and receive any IP packets. It may also be possible to provide finer granularity authorization, such as authorization for QoS or individual services (e.g., http vs. ssh). However, while a backend authorization infrastructure (e.g., RADIUS or Diameter based AAA infra) might provide such indications to the PAA, explicit support for them is outside the scope of PANA. For instance, PANA is not required to carry any indication of the services authorized for the authenticated device.

Providing access control functionality in the network is outside the scope of PANA. Client access authentication SHOULD be followed by access control to make sure only authenticated and authorized clients can send and receive IP packets via the access network. Access control can involve setting access control lists on the EPs. PANA protocol exchange identifies clients that are authorized to access the network. If IPsec-based access control is deployed in an access network, PaC and EPs should have the required IPsec SA in place. Generating the IPsec SAs based on EAP keys is outside the scope of PANA protocol. This transformation MUST be handled by a separate secure association protocol (see section 4.1.1).

Carrying accounting data is outside the scope of PANA.

4.1.3. Authentication Backend

PANA protocol MUST NOT make any assumptions on the backend authentication protocol or mechanisms. A PAA MAY interact with backend AAA infrastructures such as RADIUS or Diameter, but it is not a requirement. When the access network does not rely on an IETF-defined AAA protocol (e.g., RADIUS, Diameter), it can still use a proprietary backend system, or rely on the information locally stored on the authentication agents.

The interaction between the PAA and the backend authentication entities is outside the scope of PANA.

4.1.4. Identifiers

PANA SHOULD allow various types of identifiers to be used as the PaCI (e.g., username, Network Access Identifier (NAI), Fully Qualified Domain Name (FQDN), etc.). This requirement generally relies on the client identifiers supported by various EAP methods.

PANA SHOULD allow various types of identifiers to be used as the DI (e.g., IP address, link-layer address, port number of a switch, etc.).

A PAA MUST be able to create a binding between the PaCI and the associated DI upon successful PANA exchange. This can be achieved by PANA communicating the PaCI and DI to the PAA during the protocol exchange. The DI can be carried either explicitly as part of the PANA payload, or implicitly as the source of the PANA message, or both. Multi-access networks also require use of a cryptographic protection along with DI filtering to prevent unauthorized access [RFC4016]. The keying material required by the cryptographic methods needs to be indexed by the DI. As described in section 4.1.2, the binding between DI and PaCI is used for access control and accounting in the network.

4.2. IP Address Assignment

Assigning an IP address to the client is outside the scope of PANA. PaC MUST configure an IP address before running PANA.

4.3. EAP Lower Layer Requirements

The EAP protocol imposes various requirements on its transport protocols that are based on the nature of the EAP protocol, and need to be satisfied for correct operation. Please see [RFC3748] for the generic transport requirements that MUST be satisfied by PANA.

4.4. PAA-to-EP Protocol

PANA does not assume that the PAA is always co-located with the EP(s). Network access enforcement can be provided by one or more nodes on the same IP subnet as the client (e.g., multiple routers), or on another subnet in the access domain (e.g., gateway to the Internet, depending on the network architecture). When the PAA and the EP(s) are separated, another transport for client provisioning is necessary. This transport is needed to create access control lists in order to allow authenticated and authorized clients' traffic through the EPs. PANA Working Group will preferably identify an existing protocol solution that allows the PAA to deliver the authorization information to one or more EPs when the PAA is separated from EPs. Possible candidates include, but are not limited to COPS, SNMP, Diameter, etc.

The communication between PAA and EP(s) MUST be secure. The objective of using a PAA-to-EP protocol is to provide filtering rules to EP(s) for allowing network access of a recently authenticated and authorized PaC. The chosen protocol MUST be capable of carrying DI and cryptographic keys for a given PaC from PAA to EP. Depending on the PANA protocol design, support for either of the pull model (i.e., EP initiating the PAA-to-EP protocol exchange per PaC) or the push model (i.e., PAA initiating the PAA-to-EP protocol exchange per PaC), or both may be required. For example, if the design is such that the EP allows the PANA traffic to pass through even for unauthenticated PaCs, the EP should also allow and expect the PAA to send the filtering information at the end of a successful PANA exchange without the EP ever sending a request.

4.5. Network

4.5.1. Multi-access

PANA MUST support PaCs with multiple interfaces, and networks with multiple routers on multi-access links. In other words, PANA MUST NOT assume that the PaC has only one network interface, that the access network has only one first hop router, or that the PaC is using a point-to-point link.

4.5.2. Disconnect Indication

PANA MUST NOT assume that the link is connection-oriented. Links may or may not provide disconnect indication. Such notification is desirable in order for the PAA to clean up resources when a client moves away from the network (e.g., inform the enforcement points that the client is no longer connected). PANA SHOULD have a mechanism to

provide disconnect indication. PANA MUST be capable of securing disconnect messages in order to prevent malicious nodes from leveraging this extension for DoS attacks.

This mechanism MUST allow the PAA to be notified about the departure of a PaC from the network. This mechanism MUST also allow a PaC to be notified about the discontinuation of the network access service. Access discontinuation can occur due to various reasons such as network systems going down or a change in the access policy.

In case the clients cannot send explicit disconnect messages to the PAA, it can still detect their departure by relying on periodic authentication requests.

4.5.3. Location of PAA

The PAA and PaC MUST be exactly one IP hop away from each other. That is, there must be no IP routers between the two. Note that this does not mean they are on the same physical link. Bridging and tunneling (e.g., IP-in-IP, GRE, L2TP, etc.) techniques can place two nodes just exactly one IP hop away from each other although they might be connected to separate physical links. A PAA can be on the network access server (NAS) or WLAN access point or first hop router. The use of PANA when the PAA is multiple IP hops away from the PaC is outside the scope of PANA.

A PaC may or may not be pre-configured with the IP address of PAA. Therefore the PANA protocol MUST define a dynamic discovery method. Given that the PAA is one hop away from the PaC, there are a number of discovery techniques that could be used (e.g., multicast or anycast) by the PaC to find out the address of the PAA.

4.5.4. Secure Channel

PANA MUST NOT assume the presence of a secure channel between the PaC and the PAA. PANA MUST be able to provide authentication especially in networks which are not protected against eavesdropping and spoofing. PANA MUST enable protection against replay attacks on both PaCs and PAAs.

This requirement partially relies on the EAP protocol and the EAP methods carried over PANA. Use of EAP methods that provide mutual authentication and key derivation/distribution is essential for satisfying this requirement. EAP does not make a secure channel assumption, and supports various authentication methods that can be used in such environments. Additionally, PANA MUST ensure that its design does not contain vulnerabilities that can be exploited when it is used over insecure channels. PANA MAY provide a secure channel by

deploying a two-phase authentication. The first phase can be used for creation of the secure channel, and the second phase for client and network authentication.

4.6. Interaction with Other Protocols

Mobility management is outside the scope of PANA. However, PANA MUST be able to co-exist and MUST NOT unintentionally interfere with various mobility management protocols, such as Mobile IPv4 [RFC3344], Mobile IPv6 [RFC3775], fast handover protocols [FMIPv6] [FMIPv4], and other standard protocols like IPv6 stateless address auto-configuration [RFC2461] (including privacy extensions [RFC3041]), and DHCP [RFC2131] [RFC3315]. PANA MUST NOT make any assumptions on the protocols or mechanisms used for IP address configuration of the PaC.

4.7. Performance

PANA design SHOULD efficiently handle the authentication process in order to gain network access with minimum latency. For example, it may minimize the protocol signaling by creating local security associations.

4.8. Congestion Control

PANA MUST provide congestion control for the protocol messaging. Under certain conditions PaCs might unintentionally get synchronized when sending their requests to the PAA (e.g., upon recovering from a power outage on the access network). The network congestion generated from such events can be avoided by using techniques like delayed initialization and exponential back off.

4.9. IP Version Independence

PANA MUST work with both IPv4 and IPv6.

4.10. Denial of Service Attacks

PANA MUST be robust against a class of DoS attacks such as blind masquerade attacks through IP spoofing. These attacks would swamp the PAA, causing it to spend resources and prevent network access by legitimate clients.

4.11. Client Identity Privacy

Some clients might prefer hiding their identity from visited access networks for privacy reasons. Providing identity protection for clients is outside the scope of PANA. Note that some authentication

methods may already have this capability. Where necessary, identity protection can be achieved by letting PANA carry such authentication methods.

5. Security Considerations

This document identifies requirements for the PANA protocol design. Due to the nature of this protocol, most of the requirements are security related. The actual protocol design is not specified in this document. A thorough discussion on PANA security threats can be found in PANA Threat Analysis and Security Requirements [RFC4016]. Security threats identified in that document are already included in this general PANA requirements document.

6. Acknowledgements

Authors would like to thank Bernard Aboba, Derek Atkins, Steven Bellovin, Julien Bournelle, Subir Das, Francis Dupont, Dan Forsberg, Pete McCann, Lionel Morand, Thomas Narten, Mohan Parthasarathy, Basavaraj Patil, Hesham Soliman, and the PANA Working Group members for their valuable contributions to the discussions and preparation of this document.

Appendix A. Problem Statement

Access networks in most cases require some form of authentication in order to prevent unauthorized usage. In the absence of physical security (and sometimes in addition to it) a higher layer (L2+) access authentication mechanism is needed. Depending on the deployment scenarios, a number of features are expected from the authentication mechanism. For example, support for various authentication methods (e.g., MD5, TLS, SIM, etc.), network roaming, network service provider discovery and selection, separate authentication for access (L1+L2) service provider and ISP (L3), etc. In the absence of a link-layer authentication mechanism that can satisfy these needs, operators are forced to either use non-standard ad-hoc solutions at layers above the link, insert additional shim layers for authentication, or misuse some of the existing protocols in ways that were not intended by design. PANA will be developed to fill this gap by defining a standard network-layer access authentication protocol. As a network-layer access authentication protocol, PANA can be used over any link-layer that supports IP.

DSL networks are a specific example where PANA has the potential for addressing some of the deployment scenarios. Some DSL deployments do not use PPP [RFC1661] as the access link-layer (IP is carried over ATM and the subscriber device is either statically or DHCP-configured). The operators of these networks are left either using an application-layer web-based login (captive portal) scheme for subscriber authentication, or providing a best-effort service only as they cannot perform subscriber authentication required for the differentiated services. The captive portal scheme is a non-standard solution that has various limitations and security flaws.

PPP-based authentication can provide some of the required functionality. But using PPP only for authentication is not a good choice, as it incurs additional messaging during the connection setup and extra per-packet processing. It also forces the network topology to a point-to-point model. Aside from resistance to incorporating PPP into an architecture unless it is absolutely necessary, there is even interest in the community in removing PPP from some of the existing architectures and deployments (e.g., 3GPP2, DSL).

Using Mobile IPv4 authentication with a foreign agent instead of proper network access authentication is an example of protocol misuse. The Registration Required flag allows a foreign agent to force authentication even when the agent is not involved in any Mobile IPv4 signalling (co-located care-of address case). This enables the use of a mobility-specific protocol for an unrelated functionality.

PANA will carry EAP above IP in order to enable any authentication method on any link-layer. EAP can already be carried by [IEEE-802.1X] and PPP. IEEE 802.1X can only be used on unbridged IEEE 802 links, hence it only applies to limited link types. Inserting PPP between IP and a link-layer can be perceived as a way to enable EAP over that particular link-layer, but using PPP for this reason has the aforementioned drawbacks and is not a good choice. While IEEE 802.1X and PPP can continue to be used in their own domains, they do not take away the need to have a protocol like PANA.

Appendix B. Usage Scenarios

PANA will be applicable to various types of networks. Based on the presence of lower-layer security prior to running PANA, the following types cover all possibilities:

- a) Physically secured networks (e.g., DSL networks). Although data traffic is always carried over a physically secured link, the client might need to be authenticated and authorized when accessing the IP services.
- b) Networks where L1-L2 is already cryptographically secured before enabling IP (e.g., cdma2000 networks). Although the client is authenticated on the radio link before enabling ciphering, it additionally needs to get authenticated and authorized for accessing the IP services.
- c) No lower-layer security present before enabling IP. PANA is run in an insecure network. PANA-based access authentication is used to bootstrap cryptographic per-packet authentication and integrity protection.

PANA is applicable to not only large-scale operator deployments with full AAA infrastructure, but also to small disconnected deployments like home networks and personal area networks.

Since PANA enables decoupling AAA from the link-layer procedures, network access authentication does not have to take place during the link establishment. This allows deferring client authentication until the client attempts to access differentiated services (e.g., high bandwidth, unlimited access, etc.) in some deployments. Additionally, multiple simultaneous network access sessions over the same link-layer connection can occur as well.

The following five scenarios capture the PANA usage model in different network architectures with reference to its placement of logical elements such as the PANA Client (PaC) and the PANA Authentication Agent (PAA) with respect to the Enforcement Point (EP) and the Access Router (AR). Note that PAA may or may not use AAA infrastructure to verify the credentials of PaC in order to authorize network access.

Scenario 1: PAA co-located with EP but separated from AR

In this scenario (Figure 1), PAA is co-located with the enforcement point on which access control is performed. This might be the case where PAA is co-located with the L2 access device (e.g., an IP-capable switch).

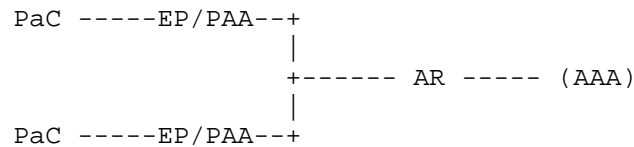


Figure 1: PAA co-located with EP but separated from AR.

Scenario 2: PAA co-located with AR but separated from EP

In this scenario, PAA is not co-located with EPs but is placed on the AR. Although we have shown only one AR here, there could be multiple ARs, one of which is co-located with the PAA. Access control parameters have to be distributed to the respective enforcement points so that the corresponding device on which PaC is authenticated can access the network. A separate protocol is needed between PAA and EP to carry access control parameters.

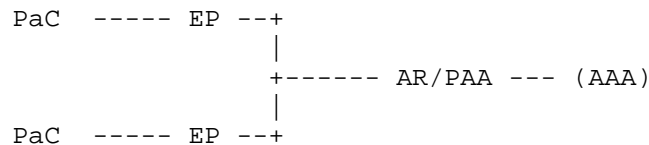


Figure 2: PAA co-located with AR but separated from EP

Scenario 3: PAA co-located with EP and AR

In this scenario (Figure 3), PAA is co-located with the EP and AR on which access control and routing are performed.

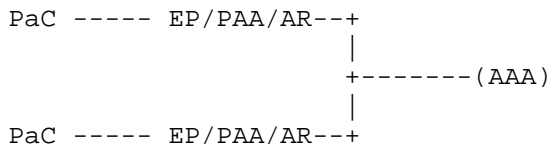


Figure 3: PAA co-located with EP and AR.

Scenario 4: Separated PAA, EP, and AR

In this scenario, PAA is neither co-located with EPs nor with ARs. It still resides on the same IP link as ARs. After successful authentication, access control parameters will be distributed to respective enforcement points via a separate protocol and PANA does not play any explicit role in this.

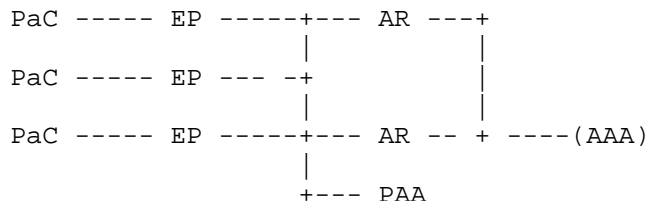


Figure 4: PAA, EP and AR separated.

Scenario 5: PAA separated from co-located EP and AR

In this scenario, EP and AR are co-located with each other but separated from PAA. PAA still resides on the same IP link as ARs. After successful authentication, access control parameters will be distributed to respective enforcement points via a separate protocol and PANA does not play any explicit role in this.

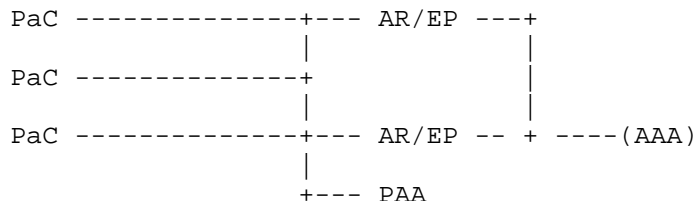


Figure 5: PAA separated from EP and AR.

References

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.
- [RFC4016] Parthasarathy, M., "Protocol for Carrying Authentication and Network Access (PANA) Threat Analysis and Security Requirements", RFC 4016, March 2005.

Informative References

- [FMIPv4] Malki, K., "Low Latency Handoffs in Mobile IPv4", Work in Progress, June 2004.
- [IEEE-802.1X] Institute of Electrical and Electronics Engineers, "Local and Metropolitan Area Networks: Port-Based Network Access Control", IEEE Standard 802.1X, September 2001.
- [RFC826] Plummer, D., "Ethernet Address Resolution Protocol: Or converting network protocol addresses to 48.bit Ethernet address for transmission on Ethernet hardware", STD 37, RFC 826, November 1982.
- [RFC1256] Deering, S., "ICMP Router Discovery Messages", RFC 1256, September 1991.
- [RFC1661] Simpson, W., "The Point-to-Point Protocol (PPP)", STD 51, RFC 1661, July 1994.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.

- [RFC2461] Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.
- [RFC2716] Aboba, B. and D. Simon, "PPP EAP TLS Authentication Protocol", RFC 2716, October 1999.
- [RFC2794] Calhoun, P. and C. Perkins, "Mobile IP Network Access Identifier Extension for IPv4", RFC 2794, March 2000.
- [RFC3012] Perkins, C. and P. Calhoun, "Mobile IPv4 Challenge/Response Extensions", RFC 3012, November 2000.
- [RFC3041] Narten, T. and R. Draves, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 3041, January 2001.
- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.
- [RFC3344] Perkins, C., "IP Mobility Support for IPv4", RFC 3344, August 2002.
- [RFC3775] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.
- [FMIPv6] Koodli, R., Ed., "Fast Handovers for Mobile IPv6", Work in Progress.

Authors' Addresses

Alper E. Yegin (editor)
Samsung Advanced Institute of Technology
75 West Plumeria Drive
San Jose, CA 95134
USA

Phone: +1 408 544 5656
EMail: alper.yegin@samsung.com

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

Phone: +1 732 699 5305
EMail: yohba@tari.toshiba.com

Reinaldo Penno
Juniper Networks
10 Technology Park Drive
Westford, MA 01886-3146
USA

EMail: rpenno@juniper.net

George Tsirtsis
Flarion
Bedminster One
135 Route 202/206 South
Bedminster, NJ 07921
USA

Phone: +44 20 88260073
EMail: G.Tsirtsis@Flarion.com

Cliff Wang
ARO/NCSU
316 Riggsbee Farm
Morrisville, NC 27560
USA

Phone: +1 919 548 4207
EMail: cliffwangmail@yahoo.com

Full Copyright Statement

Copyright (C) The Internet Society (2005).

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

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

Intellectual Property

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

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

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

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.