

Network Working Group
Request for Comments: 5191
Category: Standards Track

D. Forsberg
Nokia
Y. Ohba, Ed.
Toshiba
B. Patil
H. Tschafenig
Nokia Siemens Networks
A. Yegin
Samsung
May 2008

Protocol for Carrying Authentication for Network Access (PANA)

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

This document defines the Protocol for Carrying Authentication for Network Access (PANA), a network-layer transport for Extensible Authentication Protocol (EAP) to enable network access authentication between clients and access networks. In EAP terms, PANA is a UDP-based EAP lower layer that runs between the EAP peer and the EAP authenticator.

Table of Contents

1. Introduction	3
1.1. Specification of Requirements	4
2. Terminology	4
3. Protocol Overview	6
4. Protocol Details	7
4.1. Authentication and Authorization Phase	7
4.2. Access Phase	11
4.3. Re-Authentication Phase	11
4.4. Termination Phase	13
5. Processing Rules	13
5.1. Fragmentation	13
5.2. Sequence Number and Retransmission	14
5.3. PANA Security Association	15
5.4. Message Authentication	17
5.5. Message Validity Check	17
5.6. PaC Updating Its IP Address	19
5.7. Session Lifetime	19
6. Message Format	20
6.1. IP and UDP Headers	20
6.2. PANA Message Header	20
6.3. AVP Format	22
7. PANA Messages	24
7.1. PANA-Client-Initiation (PCI)	27
7.2. PANA-Auth-Request (PAR)	28
7.3. PANA-Auth-Answer (PAN)	28
7.4. PANA-Termination-Request (PTR)	28
7.5. PANA-Termination-Answer (PTA)	29
7.6. PANA-Notification-Request (PNR)	29
7.7. PANA-Notification-Answer (PNA)	29
8. AVPs in PANA	29
8.1. AUTH AVP	30
8.2. EAP-Payload AVP	30
8.3. Integrity-Algorithm AVP	31
8.4. Key-Id AVP	31
8.5. Nonce AVP	31
8.6. PRF-Algorithm AVP	32
8.7. Result-Code AVP	32
8.8. Session-Lifetime AVP	32
8.9. Termination-Cause AVP	33
9. Retransmission Timers	33
9.1. Transmission and Retransmission Parameters	35
10. IANA Considerations	35
10.1. PANA UDP Port Number	36
10.2. PANA Message Header	36
10.2.1. Message Type	36
10.2.2. Flags	36

10.3.	AVP Header	36
10.3.1.	AVP Code	37
10.3.2.	Flags	37
10.4.	AVP Values	37
10.4.1.	Result-Code AVP Values	37
10.4.2.	Termination-Cause AVP Values	38
11.	Security Considerations	38
11.1.	General Security Measures	38
11.2.	Initial Exchange	40
11.3.	EAP Methods	40
11.4.	Cryptographic Keys	40
11.5.	Per-Packet Ciphering	41
11.6.	PAA-to-EP Communication	41
11.7.	Liveness Test	41
11.8.	Early Termination of a Session	42
12.	Acknowledgments	42
13.	References	42
13.1.	Normative References	42
13.2.	Informative References	43

1. Introduction

Providing secure network access service requires access control based on the authentication and authorization of the clients and the access networks. 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 methods between the client and the access network.

The scope of this work is identified as designing a network-layer transport for network access authentication methods. The Extensible Authentication Protocol (EAP) [RFC3748] provides such authentication methods. In other words, PANA carries EAP, which can carry various authentication methods. By the virtue of enabling the transport of EAP above IP, any authentication method that can be carried as an EAP method is made available to PANA and hence to any link-layer technology. There is a clear division of labor between PANA (an EAP lower layer), EAP, and EAP methods as described in [RFC3748].

Various environments and usage models for PANA are identified in Appendix A of [RFC4058]. Potential security threats for network-layer access authentication protocol are discussed in [RFC4016]. These have been essential in defining the requirements [RFC4058] of the PANA protocol. Note that some of these requirements are imposed by the chosen payload, EAP [RFC3748].

There are components that are part of a complete secure network access solution but are outside of the PANA protocol specification, including PANA Authentication Agent (PAA) discovery, authentication method choice, PANA Authentication Agent-Enforcement Point (PAA-EP) protocol, access control filter creation, and data traffic protection. These components are described in separate documents (see [RFC5193] and [RFC5192]).

1.1. 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].

2. Terminology

PANA Client (PaC):

The client side of the protocol that resides in the access device (e.g., laptop, PDA, etc.). It is responsible for providing the credentials in order to prove its identity (authentication) for network access authorization. The PaC and the EAP peer are colocated in the same access device.

PANA Authentication Agent (PAA):

The protocol entity in the access network whose responsibility it is to verify the credentials provided by a PANA client (PaC) and authorize network access to the access device. The PAA and the EAP authenticator (and optionally the EAP server) are colocated in the same node. Note the authentication and authorization procedure can, according to the EAP model, also be offloaded to the back end Authentication, Authorization, and Accounting (AAA) infrastructure.

PANA Session:

A PANA session is established between the PANA Client (PaC) and the PANA Authentication Agent (PAA), and it terminates as a result of an authentication and authorization or liveness test failure, a message delivery failure after retransmissions reach maximum values, session lifetime expiration, an explicit termination message or any event that causes discontinuation of the access service. A fixed session identifier is maintained throughout a session. A session cannot be shared across multiple network interfaces.

Session Lifetime:

A duration that is associated with a PANA session. For an established PANA session, the session lifetime is bound to the lifetime of the current authorization given to the PaC. The session lifetime can be extended by a new round of EAP authentication before it expires. Until a PANA session is established, the lifetime SHOULD be set to a value that allows the PaC to detect a failed session in a reasonable amount of time.

Session Identifier:

This identifier is used to uniquely identify a PANA session on the PaC and the PAA. It is included in PANA messages to bind the message to a specific PANA session. This bidirectional identifier is allocated by the PAA in the initial request message and freed when the session terminates. The session identifier is assigned by the PAA and is unique within the PAA.

PANA Security Association (PANA SA):

A PANA security association is formed between the PaC and the PAA by sharing cryptographic keying material and associated context. The formed duplex security association is used to protect the bidirectional PANA signaling traffic between the PaC and PAA.

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 access devices. The EP and the PAA may be colocated. EPs should prevent data traffic from and to any unauthorized client, unless that data traffic is either PANA or one of the other allowed traffic types (e.g., Address Resolution Protocol (ARP), IPv6 neighbor discovery, DHCP, etc.).

Master Session Key (MSK):

A key derived by the EAP peer and the EAP server and transported to the EAP authenticator [RFC3748].

For additional terminology definitions, see the PANA framework document [RFC5193].

3. Protocol Overview

The PANA protocol is run between a client (PaC) and a server (PAA) in order to perform authentication and authorization for the network access service.

The protocol messaging consists of a series of requests and answers, some of which may be initiated by either end. Each message can carry zero or more AVPs (Attribute-Value Pairs) within the payload. The main payload of PANA is EAP, which performs authentication. PANA helps the PaC and PAA establish an EAP session.

PANA is a UDP-based protocol. It has its own retransmission mechanism to reliably deliver messages.

PANA messages are sent between the PaC and PAA as part of a PANA session. A PANA session consists of distinct phases:

- o Authentication and authorization phase: This is the phase that initiates a new PANA session and executes EAP between the PAA and PaC. The PANA session can be initiated by both the PaC and the PAA. The EAP payload (which carries an EAP method inside) is what is used for authentication. The PAA conveys the result of authentication and authorization to the PaC at the end of this phase.
- o Access phase: After successful authentication and authorization, the access device gains access to the network and can send and receive IP traffic through the EP(s). At any time during this phase, the PaC and PAA may optionally send PANA notification messages to test liveness of the PANA session on the peer.

- o Re-authentication phase: During the access phase, the PAA may, and the PaC should, initiate re-authentication if they want to update the PANA session lifetime before the PANA session lifetime expires. EAP is carried by PANA to perform re-authentication. This phase may be optionally triggered by both the PaC and the PAA without any respect to the session lifetime. The re-authentication phase is a sub-phase of the access phase. The session moves to this sub-phase from the access phase when re-authentication starts, and returns back there upon successful re-authentication.
- o Termination phase: The PaC or PAA may choose to discontinue the access service at any time. An explicit disconnect message can be sent by either end. If either the PaC or the PAA disconnects without engaging in termination messaging, it is expected that either the expiration of a finite session lifetime or failed liveness tests would clean up the session at the other end.

Cryptographic protection of messages between the PaC and PAA is possible as soon as EAP in conjunction with the EAP method exports a shared key. That shared key is used to create a PANA SA. The PANA SA helps generate per-message authentication codes that provide integrity protection and authentication.

4. Protocol Details

The following sections explain in detail the various phases of a PANA session.

4.1. Authentication and Authorization Phase

The main task of the authentication and authorization phase is to establish a PANA session and carry EAP messages between the PaC and the PAA. The PANA session can be initiated by either the PaC or the PAA.

PaC-initiated Session:

When the PaC initiates a PANA session, it sends a PANA-Client-Initiation message to the PAA. When the PaC is not configured with an IP address of the PAA before initiating the PANA session, DHCP [RFC5192] is used as the default method for dynamically configuring the IP address of the PAA. Alternative methods for dynamically discovering the IP address of the PAA may be used for PaC-initiated sessions, but they are outside the scope of this specification. The PAA that receives the PANA-Client-Initiation message MUST respond to the PaC with a PANA-Auth-Request message.

PAA-initiated Session:

When the PAA knows the IP address of the PaC, it MAY send an unsolicited PANA-Auth-Request to the PaC. The details of how PAA can learn the IP address of the PaC are outside the scope of this specification.

A session identifier for the session is assigned by the PAA and carried in the initial PANA-Auth-Request message. The same session identifier MUST be carried in the subsequent messages exchanged between the PAA and PaC throughout the session.

When the PaC receives the initial PANA-Auth-Request message from a PAA, it responds with a PANA-Auth-Answer message, if it wishes to continue the PANA session. Otherwise, it silently discards the PANA-Auth-Request message.

The initial PANA-Auth-Request and PANA-Auth-Answer messages MUST have the 'S' (Start) bit set, regardless of whether the session is initiated by the PaC or the PAA. Non-initial PANA-Auth-Request and PANA-Auth-Answer messages as well as any other messages MUST NOT have the 'S' (Start) bit set.

It is recommended that the PAA limit the rate at which it processes incoming PANA-Client-Initiation messages to provide robustness against denial of service (DoS) attacks. The details of rate limiting are outside the scope of this specification.

If a PANA SA needs to be established with use of a key-generating EAP method, the Pseudo-Random Function (PRF) and integrity algorithms to be used for PANA_AUTH_KEY derivation (see Section 5.3) and AUTH AVP calculation (see Section 5.4) are negotiated as follows: the PAA sends the initial PANA-Auth-Request carrying one or more PRF-Algorithm AVPs and one or more Integrity-Algorithm AVPs for the PRF and integrity algorithms supported by it, respectively. The PaC then selects one PRF algorithm and one integrity algorithm from these AVPs carried in the initial PANA-Auth-Request, and it responds with the initial PANA-Auth-Answer carrying one PRF-Algorithm AVP and one Integrity-Algorithm AVP for the selected algorithms. The negotiation is protected after the MSK is available, as described in Section 5.3.

If the PAA wants to stay stateless in response to a PANA-Client-Initiation message, it doesn't include an EAP-Payload AVP in the initial PANA-Auth-Request message, and it should not retransmit the message on a timer. For this reason, the PaC MUST retransmit the PANA-Client-Initiation message until it receives the second PANA-Auth-Request message (not a retransmission of the initial one) from the PAA.

It is possible that both the PAA and the PaC initiate the PANA session at the same time, i.e., the PAA sends the initial PANA-Auth-Request message without solicitation while the PaC sends a PANA-Client-Initiation message. To resolve the race condition, the PAA MUST silently discard the PANA-Client-Initiation message received from the PaC after it has sent the initial PANA-Auth-Request message. The PAA uses the source IP address and the source port number of the PANA-Client-Initiation message to identify the PaC among multiple PANA-Client-Initiation messages sent from different PaCs.

EAP messages are carried in PANA-Auth-Request messages. PANA-Auth-Answer messages are simply used to acknowledge receipt of the requests. As an optimization, a PANA-Auth-Answer message sent from the PaC MAY include the EAP message. This optimization SHOULD NOT be used when it takes time to generate the EAP message (due to, e.g., intervention of human input), in which case returning an PANA-Auth-Answer message without piggybacking an EAP message can avoid unnecessary retransmission of the PANA-Auth-Request message.

A Nonce AVP MUST be included in the first PANA-Auth-Request and PANA-Auth-Answer messages following the initial PANA-Auth-Request and PANA-Auth-Answer messages (i.e., with the 'S' (Start) bit set), and MUST NOT be included in any other message, except during re-authentication procedures (see Section 4.3).

The result of PANA authentication is carried in the last PANA-Auth-Request message sent from the PAA to the PaC. This message carries the EAP authentication result and the result of PANA authentication. The last PANA-Auth-Request message MUST be acknowledged with a PANA-Auth-Answer message. The last PANA-Auth-Request and PANA-Auth-Answer messages MUST have the 'C' (Complete) bit set, and any other message MUST NOT have the 'C' (Complete) bit set. Figure 1 shows an example sequence in the authentication and authorization phase for a PaC-initiated session.

```

PaC      PAA  Message(sequence number)[AVPs]
-----
----->    PANA-Client-Initiation(0)
<-----    PANA-Auth-Request(x)[PRF-Algorithm,Integrity-Algorithm]
              // The 'S' (Start) bit set
----->    PANA-Auth-Answer(x)[PRF-Algorithm, Integrity-Algorithm]
              // The 'S' (Start) bit set
<-----    PANA-Auth-Request(x+1)[Nonce, EAP-Payload]
----->    PANA-Auth-Answer(x+1)[Nonce] // No piggybacking EAP
----->    PANA-Auth-Request(y)[EAP-Payload]
<-----    PANA-Auth-Answer(y)
<-----    PANA-Auth-Request(x+2)[EAP-Payload]
----->    PANA-Auth-Answer(x+2)[EAP-Payload]
              // Piggybacking EAP
<-----    PANA-Auth-Request(x+3)[Result-Code, EAP-Payload,
              Key-Id, Session-Lifetime, AUTH]
              // The 'C' (Complete) bit set
----->    PANA-Auth-Answer(x+3)[Key-Id, AUTH]
              // The 'C' (Complete) bit set

```

Figure 1: Example sequence for the authentication and authorization phase for a PaC-initiated session ("Piggybacking EAP" is the case in which an EAP-Payload AVP is carried in PAN)

If a PANA SA needs to be established with use of a key-generating EAP method and an MSK is successfully generated, the last PANA-Auth-Request message with the 'C' (Complete) bit set MUST contain a Key-Id AVP and an AUTH AVP for the first derivation of keys in the session, and any subsequent message MUST contain an AUTH AVP.

EAP authentication can fail at a pass-through authenticator without sending an EAP Failure message [RFC4137]. When this occurs, the PAA SHOULD silently terminate the session, expecting that a session timeout on the PaC will clean up the state on the PaC.

There is a case where EAP authentication succeeds with producing an EAP Success message, but network access authorization fails due to, e.g., authorization rejected by a AAA server or authorization locally rejected by the PAA. When this occurs, the PAA MUST send the last PANA-Auth-Request with a result code PANA_AUTHORIZATION_REJECTED. If an MSK is available, the last PANA-Auth-Request and PANA-Auth-Answer messages with the 'C' (Complete) bit set MUST be protected with an AUTH AVP and carry a Key-Id AVP. The PANA session MUST be terminated immediately after the last PANA-Auth message exchange.

For reasons described in Section 3 of [RFC5193], the PaC may need to reconfigure the IP address after a successful authentication and authorization phase to obtain an IP address that is usable for

exchanging data traffic through EP. In this case, the PAA sets the 'I' (IP Reconfiguration) bit of PANA-Auth-Request messages in the authentication and authorization phase to indicate to the PaC the need for IP address reconfiguration. How IP address reconfiguration is performed is outside the scope of this document.

4.2. Access Phase

Once the authentication and authorization phase successfully completes, the PaC gains access to the network and can send and receive IP data traffic through the EP(s), and the PANA session enters the access phase. In this phase, PANA-Notification-Request and PANA-Notification-Answer messages with the 'P' (Ping) bit set (ping request and ping answer messages, respectively) can be used for testing the liveness of the PANA session on the PANA peer. Both the PaC and the PAA are allowed to send a ping request to the communicating peer whenever they need to ensure the availability of the session on the peer, and they expect the peer to return a ping answer message. The ping request and answer messages MUST be protected with an AUTH AVP when a PANA SA is available. A ping request MUST NOT be sent in the authentication and authorization phase, re-authentication phase, and termination phase.

Implementations MUST limit the rate of performing this test. The PaC and the PAA can handle rate limitation on their own, they do not have to perform any coordination with each other. There is no negotiation of timers for this purpose. Additionally, an implementation MAY rate limit processing the incoming ping requests. It should be noted that if a PAA or PaC that considers its connectivity lost after a relatively small number of unresponsive pings is coupled with a peer that is aggressively rate limiting the ping request and answer messages, then false-positives could result. Therefore, a PAA or PaC should not rely on frequent ping operation to quickly determine loss of connectivity.

4.3. Re-Authentication Phase

The PANA session in the access phase can enter the re-authentication phase to extend the current session lifetime by re-executing EAP. Once the re-authentication phase successfully completes, the session re-enters the access phase. Otherwise, the session is terminated.

When the PaC initiates re-authentication, it sends a PANA-Notification-Request message with the 'A' (re-Authentication) bit set (a re-authentication request message) to the PAA. This message MUST contain the session identifier assigned to the session being re-authenticated. If the PAA already has an established PANA session for the PaC with the matching session identifier, it MUST

first respond with a PANA-Notification-Answer message with the 'A' (re-Authentication) bit set (a re-authentication answer message), followed by a PANA-Auth-Request message that starts a new EAP authentication. If the PAA cannot identify the session, it MUST silently discard the message. The first PANA-Auth-Request and PANA-Auth-Answer messages in the re-authentication phase MUST have the 'S' (Start) bit cleared and carry a Nonce AVP.

The PaC may receive a PANA-Auth-Request before receiving the answer to its outstanding re-authentication request message. This condition can arise due to packet re-ordering or a race condition between the PaC and PAA when they both attempt to engage in re-authentication. The PaC MUST keep discarding the received PANA-Auth-Requests until it receives the answer to its request.

When the PAA initiates re-authentication, it sends a PANA-Auth-Request message containing the session identifier for the PaC. The PAA MUST initiate EAP re-authentication before the current session lifetime expires.

Re-authentication of an ongoing PANA session MUST NOT reset the sequence numbers.

For any re-authentication, if there is an established PANA SA, re-authentication request and answer messages and subsequent PANA-Auth-Request and PANA-Auth-Answer messages MUST be protected with an AUTH AVP. The final PANA-Auth-Request and PANA-Auth-Answer messages and any subsequent PANA message MUST be protected by using the key generated from the latest EAP authentication.

PaC	PAA	Message(sequence number)[AVPs]

----->		PANA-Notification-Request(q)[AUTH] // The 'A' (re-Authentication) bit set
<-----		PANA-Notification-Answer(q)[AUTH] // The 'A' (re-Authentication) bit set
<-----		PANA-Auth-Request(p)[EAP-Payload, Nonce, AUTH]
----->		PANA-Auth-Answer(p)[AUTH, Nonce]
----->		PANA-Auth-Request(q+1)[EAP-Payload, AUTH]
<-----		PANA-Auth-Answer(q+1)[AUTH]
<-----		PANA-Auth-Request(p+1)[EAP-Payload, AUTH]
----->		PANA-Auth-Answer(p+1)[EAP-Payload, AUTH]
<-----		PANA-Auth-Request(p+2)[Result-Code, EAP-Payload, Key-Id, Session-Lifetime, AUTH] // The 'C' (Complete) bit set
----->		PANA-Auth-Answer(p+2)[Key-Id, AUTH] // The 'C' (Complete) bit set

Figure 2: Example sequence for the re-authentication phase initiated by PaC

4.4. Termination Phase

A procedure for explicitly terminating a PANA session can be initiated either from the PaC (i.e., disconnect indication) or from the PAA (i.e., session revocation). The PANA-Termination-Request and PANA-Termination-Answer message exchanges are used for disconnect-indication and session-revocation procedures.

The reason for termination is indicated in the Termination-Cause AVP. When there is an established PANA SA between the PaC and the PAA, all messages exchanged during the termination phase MUST be protected with an AUTH AVP. When the sender of the PANA-Termination-Request message receives a valid acknowledgment, all states maintained for the PANA session MUST be terminated immediately.

5. Processing Rules

5.1. Fragmentation

PANA does not provide fragmentation of PANA messages. Instead, it relies on fragmentation provided by EAP methods and IP layer when needed.

5.2. Sequence Number and Retransmission

PANA uses sequence numbers to provide ordered and reliable delivery of messages.

The PaC and PAA maintain two sequence numbers: one is for setting the sequence number of the next outgoing request; the other is for matching the sequence number of the next incoming request. These sequence numbers are 32-bit unsigned numbers. They are monotonically incremented by 1 as new requests are generated and received, and wrapped to zero on the next message after $2^{32}-1$. Answers always contain the same sequence number as the corresponding request. Retransmissions reuse the sequence number contained in the original packet.

The initial sequence numbers (ISN) are randomly picked by the PaC and PAA as they send their very first request messages. PANA-Client-Initiation message carries sequence number 0.

When a request message is received, it is considered valid in terms of sequence numbers if and only if its sequence number matches the expected value. This check does not apply to the PANA-Client-Initiation message and the initial PANA-Auth-Request message.

When an answer message is received, it is considered valid in terms of sequence numbers if and only if its sequence number matches that of the currently outstanding request. A peer can only have one outstanding request at a time.

PANA request messages are retransmitted based on a timer until an answer is received (in which case the retransmission timer is stopped) or the number of retransmission reaches the maximum value (in which case the PANA session MUST be terminated immediately).

The retransmission timers SHOULD be calculated as described in Section 9, unless a given deployment chooses to use its own retransmission timers optimized for the underlying link-layer characteristics.

Unless dropped due to rate limiting, the PaC and PAA MUST respond to all duplicate request messages received. The last transmitted answer MAY be cached in case it is not received by the peer, which generates a retransmission of the last request. When available, the cached answer can be used instead of fully processing the retransmitted request and forming a new answer from scratch.

5.3. PANA Security Association

A PANA SA is created as an attribute of a PANA session when EAP authentication succeeds with a creation of an MSK. A PANA SA is not created when the PANA authentication fails or no MSK is produced by the EAP authentication method. When a new MSK is derived in the PANA re-authentication phase, any key derived from the old MSK MUST be updated to a new one that is derived from the new MSK. In order to distinguish the new MSK from old ones, one Key-Id AVP MUST be carried in the last PANA-Auth-Request and PANA-Auth-Answer messages with the 'C' (Complete) bit set at the end of the EAP authentication, which resulted in deriving a new MSK. The Key-Id AVP is of type Unsigned32 and MUST contain a value that uniquely identifies the MSK within the PANA session. The last PANA-Auth-Answer message with the 'C' (Complete) bit set in response to the last PANA-Auth-Request message with the 'C' (Complete) bit set MUST contain a Key-Id AVP with the same MSK identifier carried in the request. The last PANA-Auth-Request and PANA-Auth-Answer messages with a Key-Id AVP MUST also carry an AUTH AVP whose value is computed by using the new PANA_AUTH_KEY derived from the new MSK. Although the specification does not mandate a particular method for calculation of the Key-Id AVP value, a simple method is to use monotonically increasing numbers.

The PANA session lifetime is bounded by the authorization lifetime granted by the authentication server (same as the MSK lifetime). The lifetime of the PANA SA (hence the PANA_AUTH_KEY) is the same as the lifetime of the PANA session. The created PANA SA is deleted when the corresponding PANA session is terminated.

PANA SA attributes as well as PANA session attributes are listed below:

PANA Session attributes:

- * Session Identifier
- * IP address and UDP port number of the PaC
- * IP address and UDP port number of the PAA
- * Sequence number for the next outgoing request
- * Sequence number for the next incoming request
- * Last transmitted message payload
- * Retransmission interval

- * Session lifetime
- * PANA SA attributes

PANA SA attributes:

- * Nonce generated by PaC (PaC_nonce)
- * Nonce generated by PAA (PAA_nonce)
- * MSK
- * MSK Identifier
- * PANA_AUTH_KEY
- * Pseudo-random function
- * Integrity algorithm

The PANA_AUTH_KEY is derived from the available MSK, and it is used to integrity protect PANA messages. The PANA_AUTH_KEY is computed in the following way:

$$\text{PANA_AUTH_KEY} = \text{prf+}(\text{MSK}, \text{"IETF PANA"} | \text{I_PAR} | \text{I_PAN} | \text{PaC_nonce} | \text{PAA_nonce} | \text{Key_ID})$$

where:

- The prf+ function is defined in IKEv2 [RFC4306]. The pseudo-random function to be used for the prf+ function is negotiated using PRF-Algorithm AVP in the initial PANA-Auth-Request and PANA-Auth-Answer exchange with 'S' (Start) bit set.
- MSK is the master session key generated by the EAP method.
- "IETF PANA" is the ASCII code representation of the non-NULL terminated string (excluding the double quotes around it).
- I_PAR and I_PAN are the initial PANA-Auth-Request and PANA-Auth-Answer messages (the PANA header and the following PANA AVPs) with 'S' (Start) bit set, respectively.
- PaC_nonce and PAA_nonce are values of the Nonce AVP carried in the first non-initial PANA-Auth-Answer and PANA-Auth-Request messages in the authentication and authorization phase or the first PANA-Auth-Answer and PANA-Auth-Request messages in the re-authentication phase, respectively.

- Key_ID is the value of the Key-Id AVP.

The length of PANA_AUTH_KEY depends on the integrity algorithm in use. See Section 5.4 for the detailed usage of the PANA_AUTH_KEY.

5.4. Message Authentication

A PANA message can contain an AUTH AVP for cryptographically protecting the message.

When an AUTH AVP is included in a PANA message, the Value field of the AUTH AVP is calculated by using the PANA_AUTH_KEY in the following way:

AUTH AVP value = PANA_AUTH_HASH(PANA_AUTH_KEY, PANA_PDU)

where PANA_PDU is the PANA message including the PANA header, with the AUTH AVP Value field first initialized to 0. PANA_AUTH_HASH represents the integrity algorithm negotiated using Integrity-Algorithm AVP in the initial PANA-Auth-Request and PANA-Auth-Answer exchange with 'S' (Start) bit set. The PaC and PAA MUST use the same integrity algorithm to calculate an AUTH AVP they originate and receive.

5.5. Message Validity Check

When a PANA message is received, the message is considered to be invalid, at least when one of the following conditions are not met:

- o Each field in the message header contains a valid value including sequence number, message length, message type, flags, session identifier, etc.
- o The message type is one of the expected types in the current state. Specifically, the following messages are unexpected and invalid:
 - * In the authentication and authorization phase:
 - + PANA-Client-Initiation after completion of the initial PANA-Auth-Request and PANA-Auth-Answer exchange with 'S' (Start) bit set.
 - + Re-authentication request.
 - + Ping request.

- + The last PANA-Auth-Request with 'C' (Complete) bit set before completion of the initial PANA-Auth-Request and PANA-Auth-Answer exchange with 'S' (Start) bit set.
- + The initial PANA-Auth-Request with 'S' (Start) bit set after a PaC receives a valid non-initial PANA-Auth-Request with 'S' (Start) bit cleared.
- + PANA-Termination-Request.
- * In the re-authentication phase:
 - + PANA-Client-Initiation.
 - + The initial PANA-Auth-Request.
- * In the access phase:
 - + PANA-Auth-Request.
 - + PANA-Client-Initiation.
- * In the termination phase:
 - + PANA-Client-Initiation.
 - + All requests but PANA-Termination-Request and ping request.
- o The message payload contains a valid set of AVPs allowed for the message type. There is no missing AVP that needs to be included in the payload, and no AVP, which needs to be at a fixed position, is included in a position different from this fixed position.
- o Each AVP is recognized and decoded correctly.
- o Once the PANA authentication succeeds in using a key-generating EAP method, the PANA-Auth-Request message that carries the EAP Success and any subsequent message in that session contains an AUTH AVP. The AVP value matches the hash value computed against the received message.

Invalid messages MUST be discarded in order to provide robustness against DoS attacks.

5.6. PaC Updating Its IP Address

A PaC's IP address used for PANA can change in certain situations, e.g., when IP address reconfiguration is needed for the PaC to obtain an IP address after successful PANA authentication (see Section 3 of [RFC5193]) or when the PaC moves from one IP link to another within the same PAA's realm. In order to maintain the PANA session, the PAA needs to be notified about the change of PaC address.

After the PaC has changed its IP address used for PANA, it MUST send any valid PANA message. If the message that carries the new PaC IP address in the Source Address field of the IP header is valid, the PAA MUST update the PANA session with the new PaC address. If there is an established PANA SA, the message MUST be protected with an AUTH AVP.

5.7. Session Lifetime

The authentication and authorization phase determines the PANA session lifetime, and the lifetime is indicated to the PaC when the network access authorization succeeds. For this purpose, when the last PANA-Auth-Request message (i.e., with the 'C' (Complete) bit set) in authentication and authorization phase or re-authentication phase carries a Result-Code AVP with a value of PANA_SUCCESS, a Session-Lifetime AVP MUST also be carried in the message. A Session-Lifetime AVP MUST be ignored when included in other PANA messages.

The lifetime is a non-negotiable parameter that can be used by the PaC to manage PANA-related state. The PaC MUST initiate the re-authentication phase before the current session lifetime expires, if it wants to extend the session.

The PaC and the PAA MAY use information obtained outside PANA (e.g., lower-layer indications) to expedite the detection of a disconnected peer. Availability and reliability of such indications MAY depend on a specific link-layer or network topology and are therefore only hints. A PANA peer SHOULD use the ping request and answer exchange to verify that a peer is, in fact, no longer alive, unless information obtained outside PANA is being used to expedite the detection of a disconnected peer.

The session lifetime parameter is not related to the transmission of ping request messages. These messages can be used for asynchronously verifying the liveness of the peer. The decision to send a ping request message is made locally and does not require coordination between the peers.

6. Message Format

This section defines message formats for PANA protocol.

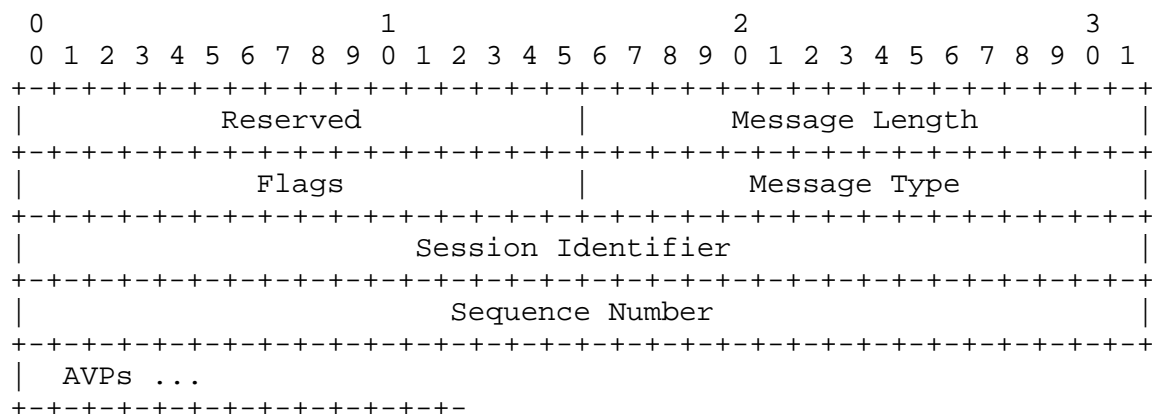
6.1. IP and UDP Headers

Any PANA message is unicast between the PaC and the PAA.

For any PANA message sent from the peer that has initiated the PANA session, the UDP source port is set to any number on which the peer can receive incoming PANA messages, and the destination port is set to the assigned PANA port number (716). For any PANA message sent from the other peer, the source port is set to the assigned PANA port number (716), and the destination port is copied from the source port of the last received message. In case both the PaC and PAA initiate the session (i.e., PANA-Client-Initiation and unsolicited PANA-Auth-Request messages cross each other), then the PaC is identified as the initiator. All PANA peers MUST listen on the assigned PANA port number (716).

6.2. PANA Message Header

A summary of the PANA message header format is shown below. The fields are transmitted in network byte order.



Reserved

This 16-bit field is reserved for future use. It MUST be set to zero and ignored by the receiver.

Message Length

The Message Length field is two octets and indicates the length of the PANA message including the header fields.

Flags

The Flags field is two octets. The following bits are assigned:

```

      0                               1
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+
|R S C A P I r r r r r r r r r r|
+---+---+---+---+---+---+---+---+

```

R (Request)

If set, the message is a request. If cleared, the message is an answer.

S (Start)

If set, the message is the first PANA-Auth-Request or PANA-Auth-Answer in authentication and authorization phase. For other messages, this bit MUST be cleared.

C (Complete)

If set, the message is the last PANA-Auth-Request or PANA-Auth-Answer in authentication and authorization phase. For other messages, this bit MUST be cleared.

A (re-Authentication)

If set, the message is a PANA-Notification-Request or PANA-Notification-Answer to initiate re-authentication. For other messages, this bit MUST be cleared.

P (Ping)

If set, the message is a PANA-Notification-Request or PANA-Notification-Answer for liveness test. For other messages, this bit MUST be cleared.

I (IP Reconfiguration)

If set, it indicates that the PaC is required to perform IP address reconfiguration after successful authentication and authorization phase to configure an IP address that is usable for exchanging data traffic across EP. This bit is set by the PAA only for PANA-Auth-Request messages in the authentication and authorization phase. For other messages, this bit MUST be cleared.

r (reserved)

These flag bits are reserved for future use. They MUST be set to zero and ignored by the receiver.

Message Type

The Message Type field is two octets, and it is used in order to communicate the message type with the message. Message Type allocation is managed by IANA [IANAWEB].

Session Identifier

This field contains a 32-bit session identifier.

Sequence Number

This field contains a 32-bit sequence number.

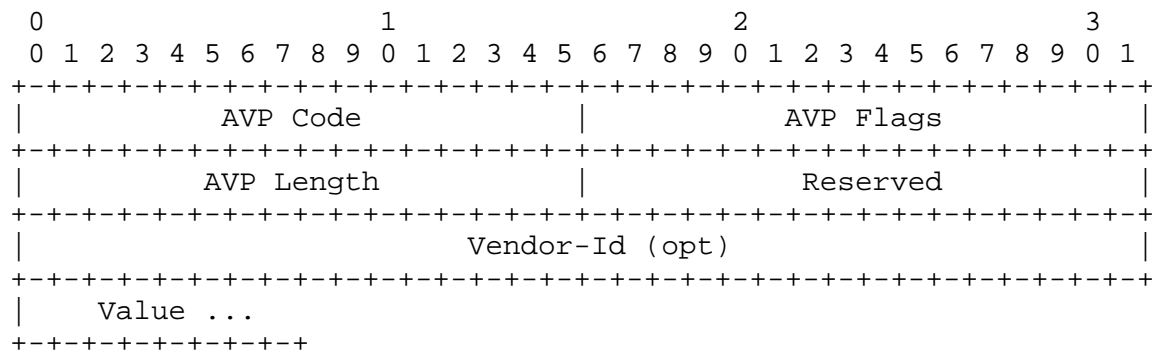
AVPs

AVPs are a method of encapsulating information relevant to the PANA message. See Section 6.3 for more information on AVPs.

6.3. AVP Format

Each AVP of type OctetString MUST be padded to align on a 32-bit boundary, while other AVP types align naturally. A number of zero-valued bytes are added to the end of the AVP Value field until a word boundary is reached. The length of the padding is not reflected in the AVP Length field [RFC3588].

The fields in the AVP are sent in network byte order. The AVP format is:




AVP Code

The AVP Code, together with the optional Vendor-Id field, identifies an attribute that follows. If the V-bit is not set, then the Vendor-Id is not present and the AVP Code refers to an IETF attribute.

AVP Flags

The AVP Flags field is two octets. The following bits are assigned:



V (Vendor)

The 'V' (Vendor) bit indicates whether the optional Vendor-Id field is present in the AVP header. When set, the AVP Code belongs to the specific vendor code address space. All AVPs defined in this document MUST have the 'V' (Vendor) bit cleared.

r (reserved)

These flag bits are reserved for future use. They MUST be set to zero and ignored by the receiver.

AVP Length

The AVP Length field is two octets, and indicates the number of octets in the Value field. The length of the AVP Code, AVP Length, AVP Flags, Reserved and Vendor-Id fields are not counted in the AVP Length value.

Reserved

This two-octet field is reserved for future use. It MUST be set to zero and ignored by the receiver.

Vendor-Id

The Vendor-Id field is present if the 'V' (Vendor) bit is set in the AVP Flags field. The optional four-octet Vendor-Id field contains the IANA assigned "SMI Network Management Private Enterprise Codes" [IANAWEB] value, encoded in network byte order. Any vendor wishing to implement a vendor-specific PANA AVP MUST use their own Vendor-Id along with their privately managed AVP address space, guaranteeing that they will not collide with any other vendor's vendor-specific AVP(s) nor with future IETF applications.

Value

The Value field is zero or more octets and contains information specific to the Attribute. The format of the Value field is determined by the AVP Code and Vendor-Id fields. The length of the Value field is determined by the AVP Length field.

7. PANA Messages

Each Request/Answer message pair is assigned a sequence number, and the sub-type (i.e., request or answer) is identified via the 'R' (Request) bit in the Message Flags field of the PANA message header.

Every PANA message MUST contain a message type in its header's Message Type field, which is used to determine the action that is to be taken for a particular message. Figure 3 lists all PANA messages defined in this document:

Message Name	Abbrev.	Message Type	PaC<->PAA	Ref.
PANA-Client-Initiation	PCI	1	----->	7.1
PANA-Auth-Request	PAR	2	<----->	7.2
PANA-Auth-Answer	PAN	2	<----->	7.3
PANA-Termination-Request	PTR	3	<----->	7.4
PANA-Termination-Answer	PTA	3	<----->	7.5
PANA-Notification-Request	PNR	4	<----->	7.6
PANA-Notification-Answer	PNA	4	<----->	7.7

Figure 3: Table of PANA Messages

The language used for PANA message definitions (i.e., AVPs valid for that PANA message type), in Section 7.1 through Section 7.7, is defined using ABNF [RFC5234] as follows:

```

message-def      = Message-Name LWSP "::~=" LWSP PANA-message

Message-Name     = PANA-name

PANA-name        = ALPHA *(ALPHA / DIGIT / "-")

PANA-message     = header LWSP *fixed LWSP *required
                  LWSP *optional LWSP *fixed

header           = "<" LWSP "PANA-Header:" LWSP Message-Type
                  [r-bit] [s-bit] [c-bit] [a-bit] [p-bit] [i-bit]
                  LWSP ">"

Message-Type     = 1*DIGIT
                  ; The Message Type assigned to the message

r-bit            = ",REQ"
                  ; If present, the 'R' (Request) bit in the Message
                  ; Flags is set, indicating that the message
                  ; is a request, as opposed to an answer.

```

s-bit = ",STA"
; If present, the 'S' (Start) bit in the Message
; Flags is set, indicating that the message
; is the initial PAR or PAN in authentication
; and authorization phase.

c-bit = ",COM"
; If present, the 'C' bit in the Message
; Flags is set, indicating that the message
; is the final PAR and PAN in authentication
; and authorization phase or re-authentication
; phase.

a-bit = ",REA"
; If present, the 'A' (re-Authentication) bit
; in the Message Flags is set, indicating that
; the message is a re-authentication request or
; answer.

p-bit = ",PIN"
; If present, the 'P' (Ping) bit in the Message
; Flags is set, indicating that the message
; is a ping request or answer.

i-bit = ",IPR"
; If present, the 'I' (IP Reconfiguration) bit
; in the Message Flags is set, indicating that
; the PaC requires IP address reconfiguration
; after successful authentication and
; authorization phase.

fixed = [qual] "<" LWSP avp-spec LWSP ">"
; Defines the fixed position of an AVP.

required = [qual] "{" LWSP avp-spec LWSP "}"
; The AVP MUST be present and can appear
; anywhere in the message.

optional = [qual] "[" LWSP avp-name LWSP "]"
; The avp-name in the 'optional' rule cannot
; evaluate any AVP Name that is included
; in a fixed or required rule. The AVP can
; appear anywhere in the message.

```

qual          = [min] "*" [max]
                ; See ABNF conventions, RFC 5234 Section 3.6.
                ; The absence of any qualifiers depends on whether
                ; it precedes a fixed, required, or optional
                ; rule. If a fixed or required rule has no
                ; qualifier, then exactly one such AVP MUST
                ; be present. If an optional rule has no
                ; qualifier, then 0 or 1 such AVP may be
                ; present.
                ;
                ; NOTE: "[" and "]" have a different meaning
                ; than in ABNF (see the optional rule, above).
                ; These braces cannot be used to express
                ; optional fixed rules (such as an optional
                ; AUTH at the end). To do this, the convention
                ; is '0*1fixed'.

min           = 1*DIGIT
                ; The minimum number of times the element may
                ; be present. The default value is zero.

max          = 1*DIGIT
                ; The maximum number of times the element may
                ; be present. The default value is infinity. A
                ; value of zero implies the AVP MUST NOT be
                ; present.

avp-spec      = PANA-name
                ; The avp-spec has to be an AVP Name, defined
                ; in the base or extended PANA protocol
                ; specifications.

avp-name      = avp-spec / "AVP"
                ; The string "AVP" stands for *any* arbitrary
                ; AVP Name, which does not conflict with the
                ; required or fixed position AVPs defined in
                ; the message definition.

```

7.1. PANA-Client-Initiation (PCI)

The PANA-Client-Initiation (PCI) message is used for PaC-initiated session. The Sequence Number and Session Identifier fields in this message MUST be set to zero (0).

```

PANA-Client-Initiation ::= < PANA-Header: 1 >
                        *[ AVP ]

```

7.2. PANA-Auth-Request (PAR)

The PANA-Auth-Request (PAR) message is either sent by the PAA or the PaC.

The message MUST NOT have both the 'S' (Start) and 'C' (Complete) bits set.

```
PANA-Auth-Request ::= < PANA-Header: 2,REQ[,STA][,COM][,IPR] >
    [ EAP-Payload ]
    [ Nonce ]
    *[ PRF-Algorithm ]
    *[ Integrity-Algorithm ]
    [ Result-Code ]
    [ Session-Lifetime ]
    [ Key-Id ]
    *[ AVP ]
    0*1< AUTH >
```

7.3. PANA-Auth-Answer (PAN)

The PANA-Auth-Answer (PAN) message is sent by either the PaC or the PAA in response to a PANA-Auth-Request message.

The message MUST NOT have both the 'S' (Start) and 'C' (Complete) bits set.

```
PANA-Auth-Answer ::= < PANA-Header: 2[,STA][,COM] >
    [ Nonce ]
    [ PRF-Algorithm ]
    [ Integrity-Algorithm ]
    [ EAP-Payload ]
    [ Key-Id ]
    *[ AVP ]
    0*1< AUTH >
```

7.4. PANA-Termination-Request (PTR)

The PANA-Termination-Request (PTR) message is sent either by the PaC or the PAA to terminate a PANA session.

```
PANA-Termination-Request ::= < PANA-Header: 3,REQ >
    < Termination-Cause >
    *[ AVP ]
    0*1< AUTH >
```

7.5. PANA-Termination-Answer (PTA)

The PANA-Termination-Answer (PTA) message is sent either by the PaC or the PAA in response to PANA-Termination-Request.

```
PANA-Termination-Answer ::= < PANA-Header: 3 >
                             * [ AVP ]
                             0 * 1 < AUTH >
```

7.6. PANA-Notification-Request (PNR)

The PANA-Notification-Request (PNR) message is used for signaling re-authentication and performing liveness test. See Section 4.3 and Section 4.2 for details on re-authentication and liveness test, respectively.

The message MUST have one of the 'A' (re-Authentication) and 'P' (Ping) bits exclusively set.

```
PANA-Notification-Request ::= < PANA-Header: 4, REQ[, REA][, PIN] >
                             * [ AVP ]
                             0 * 1 < AUTH >
```

7.7. PANA-Notification-Answer (PNA)

The PANA-Notification-Answer (PNA) message is sent by the PAA (PaC) to the PaC (PAA) in response to a PANA-Notification-Request from the PaC (PAA).

The message MUST have one of the 'A' (re-Authentication) and 'P' (Ping) bits exclusively set.

```
PANA-Notification-Answer ::= < PANA-Header: 4[, REA][, PIN] >
                             * [ AVP ]
                             0 * 1 < AUTH >
```

8. AVPs in PANA

This document uses AVP Value Format such as 'OctetString' and 'Unsigned32' as defined in Section 4.2 of [RFC3588]. The definitions of these data formats are not repeated in this document.

The following table lists the AVPs used in this document, and specifies in which PANA messages they MAY or MAY NOT be present.

The table uses the following symbols:

- 0 The AVP MUST NOT be present in the message.
- 0-1 Zero or one instance of the AVP MAY be present in the message.
It is considered an error if there is more than one instance of the AVP.
- 1 One instance of the AVP MUST be present in the message.
- 0+ Zero or more instances of the AVP MAY be present in the message.

Attribute Name	Message Type						
	PCI	PAR	PAN	PTR	PTA	PNR	PNA
AUTH	0	0-1	0-1	0-1	0-1	0-1	0-1
EAP-Payload	0	0-1	0-1	0	0	0	0
Integrity-Algorithm	0	0+	0-1	0	0	0	0
Key-Id	0	0-1	0-1	0	0	0	0
Nonce	0	0-1	0-1	0	0	0	0
PRF-Algorithm	0	0+	0-1	0	0	0	0
Result-Code	0	0-1	0	0	0	0	0
Session-Lifetime	0	0-1	0	0	0	0	0
Termination-Cause	0	0	0	1	0	0	0

Figure 4: AVP Occurrence Table

8.1. AUTH AVP

The AUTH AVP (AVP Code 1) is used to integrity protect PANA messages. The AVP data payload contains the Message Authentication Code encoded in network byte order. The AVP length varies depending on the integrity algorithm used. The AVP data is of type OctetString.

8.2. EAP-Payload AVP

The EAP-Payload AVP (AVP Code 2) is used for encapsulating the actual EAP message that is being exchanged between the EAP peer and the EAP authenticator. The AVP data is of type OctetString.

8.3. Integrity-Algorithm AVP

The Integrity-Algorithm AVP (AVP Code 3) is used for conveying the integrity algorithm to compute an AUTH AVP. The AVP data is of type Unsigned32. The AVP data contains an Internet Key Exchange Protocol version 2 (IKEv2) Transform ID of Transform Type 3 [RFC4306] for the integrity algorithm. All PANA implementations MUST support AUTH_HMAC_SHA1_160 (7) [RFC4595].

8.4. Key-Id AVP

The Key-Id AVP (AVP Code 4) is of type Integer32 and contains an MSK identifier. The MSK identifier is assigned by PAA and MUST be unique within the PANA session.

8.5. Nonce AVP

The Nonce AVP (AVP Code 5) carries a randomly chosen value that is used in cryptographic key computations. The recommendations in [RFC4086] apply with regard to generation of random values. The AVP data is of type OctetString, and it contains a randomly generated value in opaque format. The data length MUST be between 8 and 256 octets, inclusive.

The length of the nonces are determined based on the available pseudo-random functions (PRFs) and the degree of trust placed into the PaC and the PAA to compute random values. The length of the random value for the nonce is determined in one of two ways, depending on whether:

1. The PaC and the PAA each are likely to be able to compute a random nonce (according to [RFC4086]). The length of the nonce has to be 1/2 the length of the PRF key (e.g., 10 octets in the case of HMAC-SHA1).
2. The PaC and the PAA each are not trusted with regard to the computation of a random nonce (according to [RFC4086]). The length of the nonce has to have the full length of the PRF key (e.g., 20 octets in the case of HMAC-SHA1).

Furthermore, the strongest available PRF for PANA has to be considered in this computation. Currently, only a single PRF (namely HMAC-SHA1) is available and therefore the maximum output length is 20 octets. Therefore, the maximum length of the nonce value SHOULD be 20 octets.

8.6. PRF-Algorithm AVP

The PRF-Algorithm AVP (AVP Code 6) is used for conveying the pseudo-random function to derive PANA_AUTH_KEY. The AVP data is of type Unsigned32. The AVP data contains an IKEv2 Transform ID of Transform Type 2 [RFC4306]. All PANA implementations MUST support PRF_HMAC_SHA1 (2) [RFC2104].

8.7. Result-Code AVP

The Result-Code AVP (AVP Code 7) is of type Unsigned32 and indicates whether an EAP authentication was completed successfully. Result-Code AVP values are described below.

PANA_SUCCESS 0

Both authentication and authorization processes are successful.

PANA_AUTHENTICATION_REJECTED 1

Authentication has failed. When authentication fails, authorization is also considered to have failed.

PANA_AUTHORIZATION_REJECTED 2

The authorization process has failed. This error could occur when authorization is rejected by a AAA server or rejected locally by a PAA, even if the authentication procedure has succeeded.

8.8. Session-Lifetime AVP

The Session-Lifetime AVP (AVP Code 8) contains the number of seconds remaining before the current session is considered expired. The AVP data is of type Unsigned32.

8.9. Termination-Cause AVP

The Termination-Cause AVP (AVP Code 9) is used for indicating the reason why a session is terminated by the requester. The AVP data is of type Enumerated. The following Termination-Cause data values are used with PANA.

LOGOUT 1 (PaC -> PAA)

The client initiated a disconnect.

ADMINISTRATIVE 4 (PAA -> PaC)

The client was not granted access or was disconnected due to administrative reasons.

```
SESSION_TIMEOUT      8    (PAA -> PaC)
```

The session has timed out, and service has been terminated.

9. Retransmission Timers

The PANA protocol provides retransmissions for the PANA-Client-Initiation message and all request messages.

PANA retransmission timers are based on the model used in DHCPv6 [RFC3315]. Variables used here are also borrowed from this specification. PANA is a request/response-based protocol. The message exchange terminates when the requester successfully receives the answer, or the message exchange is considered to have failed according to the retransmission mechanism described below.

The retransmission behavior is controlled and described by the following variables:

RT	Retransmission timeout from the previous (re)transmission
----	---

IRT	Base value for RT for the initial retransmission
-----	--

MRC Maximum retransmission count

MRT Maximum retransmission time

MRD Maximum retransmission duration

RAND Randomization factor

With each message transmission or retransmission, the sender sets RT according to the rules given below. If RT expires before the message exchange terminates, the sender recomputes RT and retransmits the message.

Each of the computations of a new RT include a randomization factor (RAND), which is a random number chosen with a uniform distribution between -0.1 and +0.1. The randomization factor is included to minimize the synchronization of messages.

The algorithm for choosing a random number does not need to be cryptographically sound. The algorithm SHOULD produce a different sequence of random numbers from each invocation.

RT for the first message retransmission is based on IRT:

$$RT = IRT + RAND * IRT$$

RT for each subsequent message retransmission is based on the previous value of RT:

$$RT = 2 * RT_{prev} + RAND * RT_{prev}$$

MRT specifies an upper bound on the value of RT (disregarding the randomization added by the use of RAND). If MRT has a value of 0, there is no upper limit on the value of RT. Otherwise:

$$\begin{aligned} &\text{if } (RT > MRT) \\ &\quad RT = MRT + RAND * MRT \end{aligned}$$

MRC specifies an upper bound on the number of times a sender may retransmit a message. Unless MRC is zero, the message exchange fails once the sender has transmitted the message MRC times.

MRD specifies an upper bound on the length of time a sender may retransmit a message. Unless MRD is zero, the message exchange fails once MRD seconds have elapsed since the client first transmitted the message.

If both MRC and MRD are non-zero, the message exchange fails whenever either of the conditions specified in the previous two paragraphs are met.

If both MRC and MRD are zero, the client continues to transmit the message until it receives a response.

9.1. Transmission and Retransmission Parameters

This section presents a table of values used to describe the message retransmission behavior of PANA requests (REQ_*) and PANA-Client-Initiation message (PCI_*). The table shows default values.

Parameter	Default	Description
PCI_IRT	1 sec	Initial PCI timeout.
PCI_MRT	120 secs	Max PCI timeout value.
PCI_MRC	0	Max PCI retransmission attempts.
PCI_MRD	0	Max PCI retransmission duration.
REQ_IRT	1 sec	Initial Request timeout.
REQ_MRT	30 secs	Max Request timeout value.
REQ_MRC	10	Max Request retransmission attempts.
REQ_MRD	0	Max Request retransmission duration.

So, for example, the first RT for the PANA-Auth-Request (PAR) message is calculated using REQ_IRT as the IRT:

$$RT = REQ_IRT + RAND * REQ_IRT$$

10. IANA Considerations

This section provides guidance to the Internet Assigned Numbers Authority (IANA) regarding the registration of values related to the PANA protocol, in accordance with BCP 26 [IANA]. The following policies are used here with the meanings defined in BCP 26: "Private Use", "First Come First Served", "Expert Review", "Specification Required", "IETF Consensus", and "Standards Action".

This section explains the criteria to be used by the IANA for assignment of numbers within namespaces defined within this document.

For registration requests where a Designated Expert should be consulted, the responsible IESG Area Director should appoint the Designated Expert. For Designated Expert with Specification Required, the request is posted to the PANA WG mailing list (or, if it has been disbanded, a successor designated by the Area Director) for comment and review, and MUST include a pointer to a public specification. Before a period of 30 days has passed, the Designated Expert will either approve or deny the registration request and

publish a notice of the decision to the PANA WG mailing list or its successor. A denial notice must be justified by an explanation and, in the cases where it is possible, concrete suggestions on how the request can be modified so as to become acceptable.

IANA has created a registry for PANA.

10.1. PANA UDP Port Number

PANA uses one well-known UDP port number (see Section 6.1), which has been assigned by the IANA (716).

10.2. PANA Message Header

As defined in Section 6.2, the PANA message header contains two fields that require IANA namespace management; the Message Type and Flags fields.

10.2.1. Message Type

The Message Type namespace is used to identify PANA messages. Message Type 0 is not used and is not assigned by IANA. The range of values 1 - 65,519 are for permanent, standard message types, allocated by IETF Consensus [IANA]. This document defines the range of values 1 - 4. The same Message Type is used for both the request and the answer messages, except for type 1. The Request bit distinguishes requests from answers. See Section 7 for the assignment of the namespace in this specification.

The range of values 65,520 - 65,535 (hexadecimal values 0xffff0 - 0xfffff) are reserved for experimental messages. As these codes are only for experimental and testing purposes, no guarantee is made for interoperability between the communicating PaC and PAA using experimental commands, as outlined in [IANA-EXP].

10.2.2. Flags

There are 16 bits in the Flags field of the PANA message header. This document assigns bit 0 ('R'), 1 ('S'), 2 ('C'), 3 ('A'), 4 ('P'), and 5 ('I') in Section 6.2. The remaining bits MUST only be assigned via a Standards Action [IANA].

10.3. AVP Header

As defined in Section 6.3, the AVP header contains three fields that require IANA namespace management; the AVP Code, AVP Flags, and Vendor-Id fields, where only the AVP Code and AVP Flags created new namespaces.

10.3.1. AVP Code

The 16-bit AVP code namespace is used to identify attributes. There are multiple namespaces. Vendors can have their own AVP codes namespace, which will be identified by their Vendor-Id (also known as Enterprise-Number), and they control the assignments of their vendor-specific AVP codes within their own namespace. The absence of a Vendor-Id identifies the IETF IANA controlled AVP codes namespace. The AVP codes, and sometimes also possible values in an AVP, are controlled and maintained by IANA.

AVP Code 0 is not used and is not assigned by IANA. This document defines the AVP Codes 1-9. See Section 8.1 through Section 8.9 for the assignment of the namespace in this specification.

AVPs may be allocated following Designated Expert Review with Specification Required [IANA] or Standards Action.

Note that PANA defines a mechanism for Vendor-Specific AVPs, where the Vendor-Id field in the AVP header is set to a non-zero value. Vendor-Specific AVP codes are for Private Use and should be encouraged instead of allocation of global attribute types, for functions specific only to one vendor's implementation of PANA, where no interoperability is deemed useful. Where a Vendor-Specific AVP is implemented by more than one vendor, allocation of global AVPs should be encouraged instead.

10.3.2. Flags

There are 16 bits in the AVP Flags field of the AVP header, defined in Section 6.3. This document assigns bit 0 ('V'). The remaining bits should only be assigned via a Standards Action .

10.4. AVP Values

Certain AVPs in PANA define a list of values with various meanings. For attributes other than those specified in this section, adding additional values to the list can be done on a First Come, First Served basis by IANA [IANA].

10.4.1. Result-Code AVP Values

As defined in Section 8.7, the Result-Code AVP (AVP Code 7) defines the values 0-2.

All remaining values are available for assignment via IETF Consensus [IANA].

10.4.2. Termination-Cause AVP Values

As defined in Section 8.9, the Termination-Cause AVP (AVP Code 9) defines the values 1, 4, and 8.

All remaining values are available for assignment via IETF Consensus [IANA].

11. Security Considerations

The PANA protocol defines a UDP-based EAP encapsulation that runs between two IP-enabled nodes. Various security threats that are relevant to a protocol of this nature are outlined in [RFC4016]. Security considerations stemming from the use of EAP and EAP methods are discussed in [RFC3748] [EAP-KEYING]. This section provides a discussion on the security-related issues that are related to PANA framework and protocol design.

An important element in assessing the security of PANA design and deployment in a network is the presence of lower-layer security. In the context of this document, lower layers are said to be secure if the environment provides adequate protection against spoofing and confidentiality based on its operational needs. For example, DSL and cdma2000 networks' lower-layer security is enabled even before running the first PANA-based authentication. In the absence of such a preestablished secure channel prior to running PANA, one can be created after the successful PANA authentication using a link-layer or network-layer cryptographic mechanism (e.g., IPsec).

11.1. General Security Measures

PANA provides multiple mechanisms to secure a PANA session.

PANA messages carry sequence numbers, which are monotonically incremented by 1 with every new request message. These numbers are randomly initialized at the beginning of the session, and they are verified against expected numbers upon receipt. A message whose sequence number is different than the expected one is silently discarded. In addition to accomplishing orderly delivery of EAP

messages and duplicate elimination, this scheme also helps prevent an adversary from spoofing messages to disturb ongoing PANA and EAP sessions unless it can also eavesdrop to synchronize with the expected sequence number. Furthermore, impact of replay attacks is reduced as any stale message (i.e., a request or answer with an unexpected sequence number and/or a session identifier for a non-existing session) and any duplicate answer are immediately discarded, and a duplicate request can trigger transmission of the cached answer (i.e., no need to process the request and generate a new answer).

The PANA framework defines EP, which is ideally located on a network device that can filter traffic from the PaCs before the traffic enters the Internet/intranet. A set of filters can be used to discard unauthorized packets, such as the initial PANA-Auth-Request message that is received from the segment of the access network, where only the PaCs are supposed to be connected (i.e., preventing PAA impersonation).

The protocol also provides authentication and integrity protection to PANA messages when the used EAP method can generate cryptographic session keys. A PANA SA is generated based on the MSK exported by the EAP method. This SA is used for generating an AUTH AVP to protect the PANA message header and payload (including the complete EAP message).

The cryptographic protection prevents an adversary from acting as a man-in-the-middle, injecting messages, replaying messages and modifying the content of the exchanged messages. Any packet that fails to pass the AUTH verification is silently discarded. The earliest this protection can be enabled is when the PANA-Auth-Request message that signals a successful authentication (EAP Success) is generated. Starting with these messages, any subsequent PANA message can be cryptographically protected until the session gets torn down.

The lifetime of the PANA SA is set to the PANA session lifetime, which is bounded by the authorization lifetime granted by the authentication server. An implementation MAY add a grace period to that value. Unless the PANA session is extended by executing another EAP authentication, the PANA SA is removed when the current session expires.

The ability to use cryptographic protection within PANA is determined by the used EAP method, which is generally dictated by the deployment environment. Insecure lower layers necessitate the use of key-generating EAP methods. In networks where lower layers are already secured, cryptographic protection of PANA messages is not necessary.

11.2. Initial Exchange

The initial PANA-Auth-Request and PANA-Auth-Answer exchange is vulnerable to spoofing attacks as these messages are not authenticated and integrity protected. In order to prevent very basic DoS attacks, an adversary should not be able to cause state creation by sending PANA-Client-Initiation messages to the PAA. This protection is achieved by allowing the responder (PAA) to create as little state as possible in the initial message exchange. However, it is difficult to prevent all spoofing attacks in the initial message exchange entirely.

11.3. EAP Methods

Eavesdropping EAP messages might cause problems when the EAP method is weak and enables dictionary or replay attacks or even allows an adversary to learn the long-term password directly. Furthermore, if the optional EAP Response/Identity payload is used, then it allows the adversary to learn the identity of the PaC. In such a case, a privacy problem is prevalent.

To prevent these threats, [RFC5193] suggests using proper EAP methods for particular environments. Depending on the deployment environment, an EAP authentication method that supports user-identity confidentiality, protection against dictionary attacks, and session-key establishment must be used. It is therefore the responsibility of the network operators and users to choose a proper EAP method.

11.4. Cryptographic Keys

When the EAP method exports an MSK, this key is used to produce a PANA SA with PANA_AUTH_KEY with a distinct key ID. The PANA_AUTH_KEY is unique to the PANA session, and it takes PANA-based nonce values into computation to cryptographically separate itself from the MSK.

The PANA_AUTH_KEY is solely used for the authentication and integrity protection of the PANA messages within the designated session.

The PANA SA lifetime is bounded by the MSK lifetime. Another execution of the EAP method yields a new MSK, and it updates the PANA SA, PANA_AUTH_KEY, and key ID.

11.5. Per-Packet Ciphering

Networks that are not secured at the lower layers prior to running PANA can rely on enabling per-packet data-traffic ciphering upon successful PANA SA establishment. The PANA framework allows generation of cryptographic keys from the PANA SA and uses the keys with a secure association protocol to enable per-packet cryptographic protection, such as link-layer or IPsec-based ciphering [PANA-IPSEC]. These mechanisms ultimately establish a cryptographic binding between the data traffic generated by and for a client and the authenticated identity of the client. Data traffic can be data origin authenticated, replay and integrity protected, and optionally encrypted using the cryptographic keys. How these keys are generated from the PANA SA and used with a secure association protocol is outside the scope of this document.

11.6. PAA-to-EP Communication

The PANA framework allows separation of PAA from EP. The protocol exchange between the PAA and EP for provisioning authorized PaC information on the EP must be protected for authentication, integrity, and replay protection.

11.7. Liveness Test

A PANA session is associated with a session lifetime. The session is terminated unless it is refreshed by a new round of EAP authentication before it expires. Therefore, the latest a disconnected client can be detected is when its session expires. A disconnect may also be detected earlier by using PANA ping messages.

A request message can be generated by either PaC or PAA at any time in access phase with the expectation that the peer responds with an answer message. A successful round-trip of this exchange is a simple verification that the peer is alive.

This test can be engaged when there is a possibility that the peer might have disconnected (e.g., after the discontinuation of data traffic for an extended period of time). Periodic use of this exchange as a keep-alive requires additional care, as it might result in congestion and hence false alarms.

This exchange is cryptographically protected when a PANA SA is available in order to prevent threats associated with the abuse of this functionality.

Any valid PANA answer message received in response to a recently sent request message can be taken as an indication of a peer's liveness. The PaC or PAA MAY forgo sending an explicit ping request message if a recent exchange has already confirmed that the peer is alive.

11.8. Early Termination of a Session

The PANA protocol supports the ability for both the PaC and the PAA to transmit a tear-down message before the session lifetime expires. This message causes state removal, a stop of the accounting procedure and removes the installed per-PaC state on the EP(s). This message is cryptographically protected when PANA SA is present.

12. Acknowledgments

We would like to thank Mark Townsley, Jari Arkko, Mohan Parthasarathy, Julien Bournelle, Rafael Marin Lopez, Pasi Eronen, Randy Turner, Erik Nordmark, Lionel Morand, Avi Lior, Susan Thomson, Giaretta Gerardo, Joseph Salowey, Sasikanth Bharadwaj, Spencer Dawkins, Tom Yu, Bernard Aboba, Subir Das, John Vollbrecht, Prakash Jayaraman, and all members of the PANA working group for their valuable comments on this document.

13. References

13.1. Normative References

- [RFC2104] Krawczyk, H., Bellare, M., and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", RFC 2104, February 1997.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

- [RFC3588] Calhoun, P., Loughney, J., Guttman, E., Zorn, G., and J. Arkko, "Diameter Base Protocol", RFC 3588, September 2003.
- [RFC3748] Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H. Levkowetz, Ed., "Extensible Authentication Protocol (EAP)", RFC 3748, June 2004.
- [RFC4086] Eastlake, D., 3rd, Schiller, J., and S. Crocker, "Randomness Requirements for Security", BCP 106, RFC 4086, June 2005.
- [RFC5234] Crocker, D., Ed., and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, January 2008.
- [RFC5192] Morand, L., Yegin A., Kumar S., and S. Madanapalli, "DHCP Options for Protocol for Carrying Authentication for Network Access (PANA) Authentication Agents", RFC 5192, May 2008.
- [IANA] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.

13.2. Informative References

- [RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.
- [RFC4016] Parthasarathy, M., "Protocol for Carrying Authentication and Network Access (PANA) Threat Analysis and Security Requirements", RFC 4016, March 2005.
- [RFC4058] Yegin, A., Ed., Ohba, Y., Penno, R., Tsirtsis, G., and C. Wang, "Protocol for Carrying Authentication for Network Access (PANA) Requirements", RFC 4058, May 2005.
- [RFC4137] Vollbrecht, J., Eronen, P., Petroni, N., and Y. Ohba, "State Machines for Extensible Authentication Protocol (EAP) Peer and Authenticator", RFC 4137, August 2005.
- [RFC4306] Kaufman, C., Ed., "Internet Key Exchange (IKEv2) Protocol", RFC 4306, December 2005.

- [RFC4595] Maino, F. and D. Black, "Use of IKEv2 in the Fibre Channel Security Association Management Protocol", RFC 4595, July 2006.
- [RFC5193] Jayaraman, P., Lopez R., Ohba Y., Ed., Parthasarathy, M., and A. Yegin, "Protocol for Carrying Authentication for Network Access (PANA) Framework", RFC 5193, May 2008.
- [EAP-KEYING] Aboba, B., Simon D., and P. Eronen, "Extensible Authentication Protocol (EAP) Key Management Framework", Work in Progress, November 2007.
- [PANA-IPSEC] Parthasarathy, M., "PANA Enabling IPsec based Access Control", Work in progress, July 2005.
- [IANAWEB] IANA, "Number assignment", <http://www.iana.org>.
- [IANA-EXP] Narten, T., "Assigning Experimental and Testing Numbers Considered Useful", BCP 82, RFC 3692, January 2004.

Authors' Addresses

Dan Forsberg
Nokia Research Center
P.O. Box 407
FIN-00045 NOKIA GROUP
Finland

Phone: +358 50 4839470
EMail: dan.forsberg@nokia.com

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

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

Basavaraj Patil
Nokia Siemens Networks
6000 Connection Drive
Irving, TX 75039
USA

EMail: basavaraj.patil@nsn.com

Hannes Tschofenig
Nokia Siemens Networks
Linnoitustie 6 Espoo 02600
Finland

Phone: +358 (50) 4871445
EMail: Hannes.Tschofenig@nsn.com
URI: <http://www.tschofenig.priv.at>

Alper E. Yegin
Samsung
Istanbul, Turkey

EMail: a.yegin@partner.samsung.com

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

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

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

Intellectual Property

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

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

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

