PANA Working Group
Internet-Draft
Expires: January 14, 2005

P. Jayaraman
Net.Com
R. Lopez
Univ. of Murcia
Y. Ohba (Ed.)
Toshiba
M. Parthasarathy
Nokia
A. Yegin
Samsung
July 16, 2004

PANA Framework
draft-ietf-pana-framework-01

Status of this Memo

Abstract

PANA design provides support for various types of deployments. Access networks can differ based on the availability of lower-layer

security, placement of PANA entities, choice of client IP configuration and authentication methods, etc. This I-D defines a general framework for describing how these various deployment choices are handled by PANA and the access network architectures. Additionally, two possible deployments are described in detail: using PANA over DSL networks and WLAN networks.

Table of Contents

1. Introduction

PANA is a link-layer agnostic network access authentication protocol that runs between a node that wants to gain access to the network and a server on the network side. PANA defines a new EAP [RFC3748] lower layer that uses IP between the protocol end points.

The motivation to define such a protocol and the requirements are described in [I-D.ietf-pana-requirements]. Protocol details are documented in [I-D.ietf-pana-pana]. [I-D.ietf-pana-ipsec] describes use of IPsec for access control following PANA-based authentication. IPsec can be used for per-packet access control, but nevertheless it is not the only way to achieve this functionality. Alternatives include reliance on physical security and link-layer ciphering. Separation of PANA server from the entity enforcing the access

control is envisaged as an optional deployment choice. SNMP [I-D.ietf-pana-snmp] is chosen as the protocol to carry associated information between the separate nodes.

PANA design provides support for various types of deployments. Access networks can differ based on the availability of lower-layer security, placement of PANA entities, choice of client IP configuration and authentication methods, etc.

PANA can be used in any access network regardless of the underlying security. For example, the network might be physically secured, or secured by means of cryptographic mechanisms after the successful client-network authentication.

The PANA client, PANA authentication agent, authentication server, and enforcement point are the relevant functional entities in this design. PANA authentication agent and enforcement point(s) can be placed on various elements in the access network (e.g., access point, access router, dedicated host).

IP address configuration mechanisms vary as well. Static configuration, DHCP, stateless address autoconfiguration are possible mechanisms to choose from. If the client configures an IPsec tunnel for enabling per-packet security, configuring IP addresses inside the tunnel becomes relevant, for which there are additional choices such as IKE.

This I-D defines a general framework for describing how these various deployment choices are handled by PANA and the access network architectures. Additionally, two possible deployments are described in detail: PANA over DSL networks and WLAN networks.

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

Pre-PANA address (PRPA)

This is an IP address configured on a PANA client before starting the PANA protocol exchange.

Post-PANA address (POPA)

This is an IP address (optionally) configured on a PANA client after a successful authentication.

IPsec Tunnel Inner Address (IPsec-TIA)

This is an IP address configured on a PANA client as the inner address of an IPsec tunnel mode SA.

IPsec Tunnel Outer Address (IPsec-TOA)

This is the address configured on a PANA client as the outer address of an IPsec tunnel mode SA.

Secure Association Protocol

A protocol that provides a cryptographic binding between the initial entity authentication (and authorization) exchange to the subsequent exchange of data packets. Examples of secure association protocols include the 4-way handshake in IEEE 802.11i [802.11i], and IKE [RFC2409] in IP-based access control.

3. General PANA Framework

PANA protocol is designed to facilitate authentication and authorization of clients in access networks. PANA is an EAP [RFC3748] lower-layer that carries EAP authentication methods encapsulated inside EAP between a client host and an agent in the access network. While PANA enables the authentication process between the two entities, it is only a part of an overall AAA and access control framework. A AAA and access control framework using PANA is comprised of four functional entities.

PANA Client (PaC):

The PaC is the client implementation of the PANA protocol. This entity resides on the end host that is requesting network access. The end hosts are, for example, laptops, PDAs, cell phones, desktop pcs that are connected to a network via a wired or wireless interface. A PaC is responsible for requesting network access and engaging in the authentication process using the PANA protocol

PANA Authentication Agent (PAA):

The PAA is the server implementation of the PANA protocol. A PAA is in charge of interfacing with the PaCs for authenticating and authorizing them for the network access service.
The PAA consults an authentication server in order to verify the credentials and rights of a PaC. If the authentication server resides on the same host as the PAA, an API is sufficient for this interaction. When they are separated (a much more common case in public access networks), a protocol needs to run between the two. LDAP [RFC3377] and AAA protocols like RADIUS [RFC2865] and Diameter [RFC3588] are commonly used for this purpose.
The PAA is also responsible for updating the access control state (i.e., filters) depending on the creation and deletion of the authorization state. The PAA communicates the updated state to the enforcement points in the network. If the PAA and EP are residing on the same host, an API is sufficient for this communication. Otherwise, a protocol is required to carry the authorized client attributes from the PAA to the EP. While not prohibiting other protocols, currently SNMP [I-D.ietf-pana-snmp] is suggested for this task.
The PAA resides on a node that is typically called a NAS (network access server) in the local area network. PAA can be hosted on any IP-enabled node on the same IP subnet as the PaC. For example on a BAS (broadband access server) in DSL networks, or PDSN in 3GPP2 networks.

Authentication Server (AS):

The server implementation that is in charge of verifying the credentials of a PaC that is requesting the network access service. The AS receives requests from the PAA on behalf of the PaCs, and responds with the result of verification together with the authorization parameters (e.g., allowed bandwidth, IP configuration, etc). The AS might be hosted on the same host as

the PAA, on a dedicated host on the access network, or on a central server somewhere on the Internet.

Enforcement Point (EP):

The access control implementation that is in charge of allowing access to authorized clients while preventing access by others. An EP learns the attributes of the authorized clients from the PAA.
The EP uses non-cryptographic or cryptographic filters to selectively allow and discard data packets. These filters may be applied at the link-layer or the IP-layer. When cryptographic access control is used, a secure association protocol needs to run between the PaC and EP. Link or network layer protection (for example TKIP, IPsec ESP) is used after the secure association protocol established the necessary security association to enable integrity protection, data origin authentication, replay protection and optionally confidentiality protection.
An EP must be located strategically in a local area network to minimize the access of unauthorized clients to the network. For example, the EP can be hosted on the switch that is directly connected to the clients in a wired network. That way the EP can drop unauthorized packets before they reach any other client host or beyond the local area network.

Figure 1 illustrates these functional entities and the interfaces (protocols, APIs) among them.

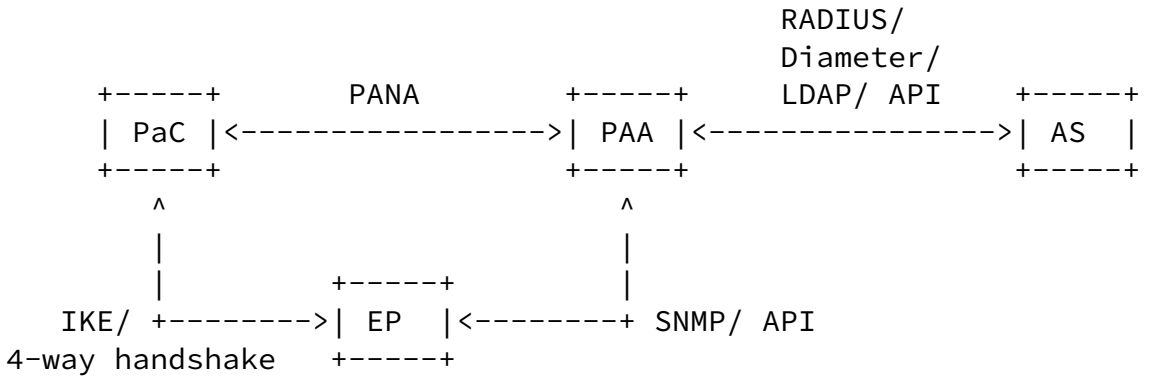

Figure 1: PANA Functional Model

Some of the entities may be co-located depending on the deployment scenario. For example, the PAA and EP would be on the same node (BAS) in DSL networks. In that case a simple API is sufficient between the PAA and EP. In small enterprise deployments the PAA and AS may be hosted on the same node (access router) that eliminates the need for a protocol run between the two. The decision to co-locate these entities or otherwise, and their precise location in the network topology are deployment decisions.

Use of IKE or 4-way handshake protocols for secure association is only required in the absence of any lower-layer security prior to running PANA. Physically secured networks (e.g., DSL) or the networks that are already cryptographically secured on the link-layer prior to PANA run (e.g., cdma2000) do not require additional secure association and per-packet ciphering. These networks can bind the PANA authentication and authorization to the lower-layer secure channel that is already available.

Figure 2 illustrates the signaling flow for authorizing a client for network access.

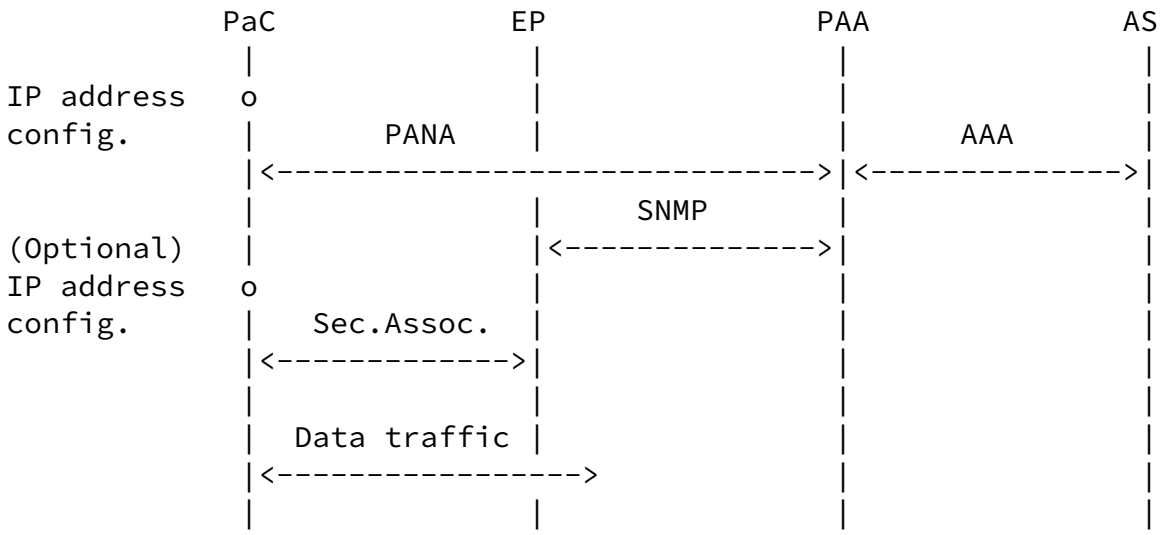

Figure 2: PANA Signaling Flow

The EP on the access network allows general data traffic from any authorized PaC, whereas it allows only limited type of traffic (e.g., PANA, DHCP, router discovery) for the unauthorized PaCs. This ensures that the newly attached clients have the minimum access service to engage in PANA and get authorized for the unlimited service.

The PaC MUST configure an IP address prior to running PANA. After the successful PANA authentication, depending on the deployment scenario the PaC MAY need to re-configure its IP address or configure additional IP address(es). The additional address configuration MAY be executed as part of the secure association protocol run.

An initially unauthorized PaC starts the PANA authentication by discovering the PAA on the access network, followed by the EAP exchange over PANA. The PAA interacts with the AS during this process. Upon receiving the authentication and authorization result from the AS, the PAA informs the PaC about the result of its network access request.

If the PaC is authorized to gain the access to the network, the PAA also sends the PaC-specific attributes (e.g., IP address, cryptographic keys, etc.) to the EP by using SNMP. The EP uses this information to alter its filters for allowing data traffic from and to the PaC to pass through.

In case cryptographic access control needs to be enabled after the PANA authentication, a secure association protocol runs between the PaC and the EP. The PaC should already have the input parameters to this process as a result of the successful PANA exchange. Similarly, the EP should have obtained them from the PAA via SNMP. Secure

association exchange produces the required security associations between the PaC and the EP to enable cryptographic data traffic protection. Per-packet cryptographic data traffic protection introduces additional per-packet overhead but the overhead exists only between the PaC and EP and will not affect communications beyond the EP. In this sense it is important to place the EP as close to the edge of the network as possible.

Finally data traffic can start flowing from and to the newly authorized PaC.

4. Environments

The PANA protocol can be used on any network environment whether there is a lower-layer secure channel prior to PANA, or one has to be enabled upon successful PANA authentication.

With regard to network access authentication two types of networks need to be considered:

a. Networks where a secure channel is already available prior to running PANA

This type of network is characterized by the existence of protection against spoofing and eavesdropping. Nevertheless, user authentication and authorization is required for network connectivity.
One example is a DSL network where the lower-layer security is provided by physical means (a.1). Physical protection of the network wiring ensures that practically there is only one client that can send and receive IP packets on the link. Another example is a cdma2000 network where the lower-layer security is provided by means of cryptographic protection (a.2). By the time the client requests access to the network-layer services, it is already authenticated and authorized for accessing the radio channel, and link-layer ciphering is enabled.
The presence of a secure channel before PANA exchange eliminates the need for executing a secure association protocol after PANA. The PANA session can be bound to the communication channel it was carried over. Also, the choice of EAP authentication method depends on the presence of this security during PANA run. Use of some authentication methods outside a secure channel is not recommended (e.g., EAP-MD5).

b. Networks where a secure channel is created after running PANA

These are the networks where there is no lower-layer protection prior to running PANA. A successful PANA authentication enables generation of cryptographic keys that are used with a secure association protocol to enable per-packet cryptographic protection.
PANA authentication is run on an insecure channel that is vulnerable to eavesdropping and spoofing. The choice of EAP method must be resilient to the possible attacks associated with such an environment. Furthermore, the EAP method must be able to create cryptographic keys that will later be used by the secure association protocol.

Whether to use a link-layer per-frame security (b.1) or a network layer security (b.2) is a deployment decision. This decision also dictates the choice of the secure association protocol. If link-layer protection is used, the protocol would be link-layer specific. If IP-layer protection is used, the secure association protocol would be IKE and the per-packet security would be provided by IPsec AH/ESP regardless of the underlying link-layer technology.

5. IP Address Configuration

The PaC configures an IP address before the PANA protocol exchange begins. This address is called a pre-PANA address (PRPA). After a successful authentication, the client may have to configure a post-PANA address (POPA) for communication with other nodes, if PRPA is a local-use (e.g., link-local or private address) or a temporarily allocated IP address. An operator might choose allocating a POPA only after successful PANA authorization either to prevent waste of premium IP resources until the client is authorized, or to enable client identity based address assignment.

In case the PaC is a IPv4/IPv6 dual-stacked host, it may configure more than one PRPA. After a successful PANA authentication the PaC may configure multiple POPAs.

There are different methods by which a PRPA can be configured.

1. In some deployments (e.g., DSL networks) the PaC may be statically configured with an IP address. This address SHOULD be used as PRPA.

2. In IPv4, most clients attempt to configure an address dynamically using DHCP [RFC2131]. If they are unable to configure an address using DHCP, they can configure a link-local address using [I-D.ietf-zeroconf-ipv4-linklocal].

   When the network access provider is able to run a DHCP server on the access link, the client would configure the PRPA using DHCP. This address may be from a private address pool [RFC1918]. Also,

the lease time on the address may vary. For example, a PRPA configured solely for running PANA can have a short lease time. PRPA may be used for local-use only (i.e., only for on-link communication, such as for PANA and IPsec tunneling with EP), or also for ultimate data communication.

In case there is no running DHCP server on the link, the client would fall back to configuring a PRPA via zeroconfiguration technique [I-D.ietf-zeroconf-ipv4-linklocal]. This yields a long-term address that can only be used for on-link communication.

3. In IPv6, clients configure a link-local address [RFC2462] when they initialize an interface. This address SHOULD be used as a PRPA.

When a PRPA is configured, the client starts the PANA protocol exchange. By that time, a dual-stacked client might have configured both an IPv4 address, and one or more IPv6 addresses as PRPAs. When

the client successfully authenticates to the network, it may be required to configure POPAs for its subsequent data communication with the other nodes.

If the client is already configured with a non-temporary address that can be used with data communication, it is not required to configure a POPA. Otherwise, the PANA-Bind-Request message allows the PAA to indicate the available configuration methods to the PaC. The PaC can choose one of the methods and act accordingly.

1. If the network relies on physical or link-layer security, the PaC can configure a POPA using DHCP [RFC2131][RFC3315] or using IPv6 stateless auto-configuration [RFC2461]. If IPv4 is being used, a PRPA is likely to be a link-local or private address, or an address with a short DHCP lease. An IPv4 PRPA SHOULD be unconfigured when the POPA is configured to prevent IPv4 address

selection problem [I-D.ietf-zeroconf-ipv4-linklocal]. If IPv6 is used, the link-local PRPA SHOULD NOT be unconfigured [RFC3484].

If the PaC is a dual-stacked host, it can configure both IPv4 and IPv6 type POPAs.

The PaC with a PRPA and the PAA with IP address X can perform on-link communication as required by PANA. In IPv4, the PRPA and IP address X have the same on-link prefix. In IPv6, the two addresses are link-local addresses. When the PaC replaces its IPv4 PRPA with an IPv4 POPA, the PaC and PAA SHOULD create host routes to each other, as they do not share the same on-link prefix any more. This is needed for the PaC with the IPv4 POPA and the PAA with IP address X to continue on-link communication. In this case, the PaC SHOULD create a host route to IP address X, and the PAA SHOULD create a host route to the IPv4 POPA. PANA defines a mechanism for the PaC to report the POPA to the PAA.

2. If the network uses IPsec for protecting the traffic on the link subsequent to PANA authentication [I-D.ietf-pana-ipsec], the PaC would use the PRPA as the outer address of IPsec tunnel mode SA (IPsec-TOA). The PaC also needs to configure an inner address (IPsec-TIA). There are different ways to configure an IPsec-TIA which are indicated in a PANA-Bind-Request message.

   When an IPv4 PRPA is configured, the same address may be used as both IPsec-TOA and IPsec-TIA. In this case, a POPA is not configured. Alternatively, an IPsec-TIA can be obtained via the configuration method available within [RFC3456] for IPv4, and [I-D.ietf-ipsec-ikev2] for both IPv4 and IPv6. This newly configured address constitutes a POPA. Please refer to [I-D.ietf-pana-ipsec] for more details.

IKEv2 can enable configuration of one IPv4 IPsec-TIA and one IPv6 IPsec-TIA for the same IPsec tunnel mode SA. Therefore, IKEv2 is recommended for handling dual-stacked PaCs where single execution of PANA and IKE is desired.

Although there are potentially a number of different ways to

configure a PRPA, and POPA when necessary, it should be noted that the ultimate decision to use one or more of these in a deployment depends on the operator. The decision is dictated by the operator's choice of per-packet protection capability (physical and link-layer vs network-layer), PRPA type (local and temporary vs global and long-term), and POPA configuration mechanisms available in the network.

6. Data Traffic Protection

Protecting data traffic of authenticated and authorized client from others is another component of providing a complete secure network access solution. Authentication, integrity and replay protection of data packets is needed to prevent spoofing when the underlying network is not physically secured. Encryption is needed when eavesdropping is a concern in the network.

When the network is physically secured, or the link-layer ciphering is already enabled prior to PANA, data traffic protection is already in place. In other cases, enabling link-layer ciphering or network-layer ciphering might rely on PANA authentication. The user and network have to make sure that an appropriate EAP method which generates keying materials is used. Once the keying material is available, it needs to be provided to the EP(s) for use with data traffic protection.

Network-layer protection, i.e., IPsec, can be used when data traffic protection is required but link-layer protection is not available. Note that the keying material generated by an EAP method is not readily usable by IPsec AH/ESP or most link layer mechanisms. A fresh and unique session key derived from the EAP method is still insufficient to produce an IPsec SA since both traffic selectors and other IPsec SA parameters are missing. The shared secret can be used in conjunction with a key management protocol like IKE [RFC2409] to turn a session key into the required IPsec SA. The details of such a mechanism is outside the scope of PANA protocol and is specified in [I-D.ietf-pana-ipsec]. PANA provides bootstrapping functionality for such a mechanism by carrying EAP methods that can generate initial keying material.

Using network-layer ciphers should be regarded as a substitute for link-layer ciphers when the latter is not available. Network-layer ciphering can also be used in addition to link-layer ciphering if the added benefits outweigh its cost to the user and the network. In this case, PANA bootstraps only the network-layer ciphering and link-layer is protected using any of the existing link-layer specific methods.

7. PAA-EP Protocol

The PANA protocol provides client authentication and authorization functionality for securing network access. The other component of a complete solution is the access control which ensures that only authenticated and authorized clients can gain access to the network. PANA enables access control by distinguishing authenticated and authorized clients from others and generating filtering information for access control mechanisms.

Access control can be achieved by placing EPs (Enforcement Points) in the network for policing the traffic flow. EPs should prevent data traffic from and to any unauthorized client unless it's either PANA or one of the other allowed traffic types (e.g., ARP, IPv6 neighbor discovery, DHCP, etc.).

When a PaC is authenticated and authorized, the PAA should notify EP(s) and ask for installing filtering rules to allow the PaC to sent and receive data traffic. SNMP is used between PAA and EP(s) for this purpose when these entities are not co-located [I-D.ietf-pana-snmp].

This section describes the possible models on the location of PAA and EP, as well as the basic authorization information that needs to be exchanged between PAA and EP. When PAA and EP are not co-located in a single device, there are other issues such as dead or rebooted peer detection and consideration for specific authorization and accounting models. However, these issues are closely related to the PAA-EP protocol solution and thus not discussed in this document. See

[I-D.ietf-pana-snmp] for further discussion.

7.1 PAA and EP Locations

EPs' location in the network topology should be appropriate for performing access control functionality. The closest IP-capable access device to the client devices is the logical choice. PAA and EPs on an access network should be aware of each other as this is necessary for access control. Generally this can be achieved by manual configuration. Dynamic discovery is another possibility, but this is left to implementations and outside the scope of this document.

Since PANA allows the separation of EP and PAA, there are several models depending on the number of EPs and PAAs and their locations. This section describes all possible models on the placement of PAA(s) and EP(s).

7.1.1 Single PAA, Single EP, Co-located

This model corresponds to the legacy NAS model. Since the PAA and the EP are co-located, the PAA-EP communication can be implemented locally by using, e.g., IPC (Inter-Process Communication). The only difference from the legacy NAS model is the case where there are multiple co-located PAA/EP devices on the same IP link and the PAA/EP devices are L2 switches or access points. In this case, for a PaC that attaches to a given PAA/EP device, other PAA/EP devices should not be discovered by the PaC even if those devices are on the same IP link. Otherwise, the PaC may result in finding a PAA that is not the closest one to it during the PANA discovery and initial handshake phase and performing PANA with the PAA, which does not correspond to the legacy NAS model. To prevent this, each PAA/EP device on an L2 switch or access point should not forward multicast PANA discovery message sent by PaCs attached to it to other devices.

7.1.2 Separate PAA and EP

When PAA is separated from EP, two cases are possible with regard to whether PAA and EP are located in parallel or serial when viewed from PaC, for each of models described in this section.

In the first case, PAA is located behind EP. The EP should be configured to always pass through PANA messages and address configuration protocol messages used for configuring an IP address used for initial PANA messaging. This case can typically happen when the EP is an L2 switch or an access point (the EP also has an IP stack to communicate with PAA via a PAA-EP protocol).

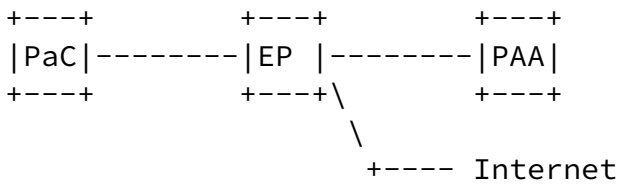

Figure 3: PAA and EP in Serial

In the other case, PAA is located in parallel to EP. Since the EP is not on the communication path between PaC and PAA, the EP does not have to configure to pass through PANA messages or address configuration protocol messages in this case. This case can typically happen when the EP is a router and the PAA is an authentication gateway without IP routing functionality.

+---+

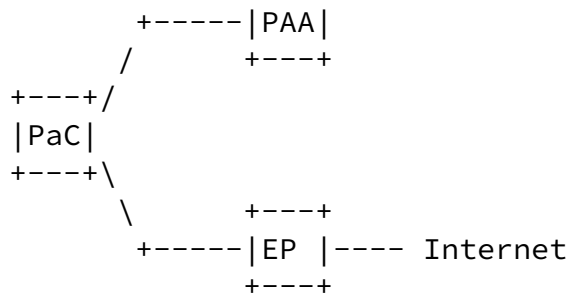

Figure 4: PAA and EP in Parallel

In the remaining of this section, PaC is not shown in figures and the figures cover both the serial and parallel models.

7.1.2.1 Single PAA, Single EP, Separated

The model benefits from separation of data traffic handling and AAA. The EP does not need to have a AAA protocol implementation which might be updated relatively more frequently than the per-packet access enforcement implementation.

7.1.2.2 Single PAA, Multiple EPs

In this model, a single PAA controls multiple EPs. The PAA may be separated from any EP or co-located with a particular EP. This model might be useful where it is preferable to run a AAA protocol at a single, manageable point. This model is particularly useful in an access network that consists of a large number of access points on which per-packet access enforcement is made. When a PaC is authenticated to the PAA, the PAA should install access control filters to each of the EPs under control of the PAA if the PAA cannot tell which EP the PaC is attached to. Even if the PAA can tell which EP the PaC is attached to, the PAA may install access control filters to those EPs if the PaC is a mobile device that can roam among the EPs. Such pre-installation of filters can reduce handoff latency. If different access authorization policies are applied to different EPs, different filter rules for a PaC may be installed on different EPs.

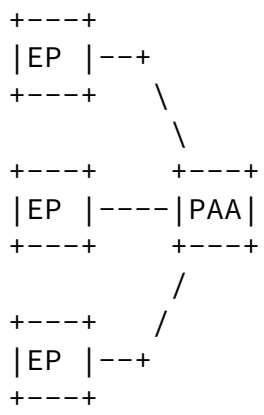

Figure 5: An example model for a single PAA and multiple EPs

7.1.2.3 Multiple PAAs

The PANA protocol allows multiple PAAs to exist on the same IP link and to be visible to PaCs on the link. PAAs may or may not be co-located with EPs as long as authorization results do not depend on whichever PAAs are chosen by a PaC [I-D.ietf-pana-pana].

7.2 Notification of PaC Presence

When PAA and EP are separated and PAA is configured to be the initiator of the discovery and initial handshake phase of PANA, EP has the responsibility to detect presence of a new PaC and notifies the PAA(s) of the presence [I-D.ietf-pana-requirements]. Such a presence notification is carried in a PAA-EP protocol message [I-D.ietf-pana-snmp].

7.3 Filter Rule Installation

Filtering rules to be installed on EP generally include a device identifier of PaC, and also cryptographic keying material (e.g., IKE pre-shared key [RFC2409]) when cryptographic data traffic protection is needed (See Section 6). Each keying material is uniquely identified with a keying material name (e.g., ID_KEY_ID in IKE [RFC2409]) and has a lifetime for key management, accounting, access control and security reasons in general. In addition to the device identifier and keying material, other filter rules, such as the IP filter rules specified in NAS-Filter-Rule AVPs carried in Diameter EAP application [I-D.ietf-aaa-eap] may be installed on EP.

8. Network Selection

The network selection problem statement is made in [I-D.ietf-eap-netsel-problem]. The PANA protocol [I-D.ietf-pana-pana] provides a way for networks to advertise which ISPs are available and for PaC to choose one ISP from the advertised information. When a PaC chooses an ISP in the PANA protocol exchange, the ultimate destination of the AAA exchange is determined based on the identity of the chosen service provider. It is also possible that the PaC does not choose a specific ISP in the PANA protocol exchange. In this case, both the ISP choice and the AAA destination are determined based on the PaC's identity, where the identity may be an NAI [RFC2486] or the physical port number or L2 address of the subscriber.

As described in [I-D.ietf-eap-netsel-problem], network selection is not only related to AAA routing but also related to payload routing. Once an ISP is chosen and the PaC is successfully authenticated and authorized, PaC is assigned an address by the ISP whose IP prefix may

be different from that of the AR. This affects the routing of the subsequent data traffic between AR and PaC. A suggested solution is to add host route from AR to PaC's POPA address and host route from PaC to AR.

Consider a typical DSL network where the AR, EP, and PAA are co-located on a BAS (Broadband Access Server) in the access network operated by a NAP (Network Access Provider). Figure 6 shows a typical model for ISP selection.

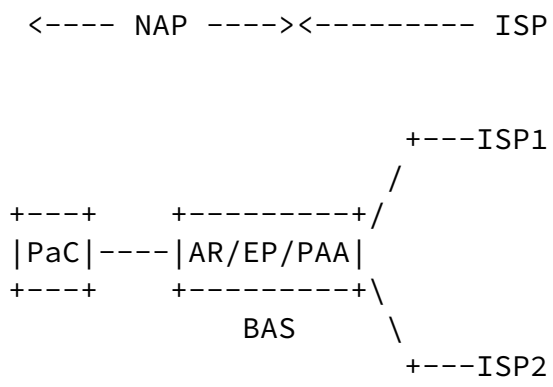

Figure 6: A Network Selection Model

When network selection is made at L3 with the use of the PANA protocol instead of L2-specific authentication mechanisms, the IP link between PaC and PAA needs to exist prior to doing PANA (and prior to network selection). In this model, the PRPA is either given by NAP or a link-local address is auto-configured. After the successful authentication with the ISP, PaC may acquire an address (POPA) from the ISP. It also learns the address of the AR, e.g.,

through DHCP, to be used as its default router. The address of the AR may or may not be in the same IP subnet as that of the PaC's POPA. If they don't share the same prefix, they SHOULD use host routes to reach each other. Note that the physically secured DSL networks do not require IPsec-based access control. Therefore the PaCs use one

IP address at a time where POPA replaces PRPA upon configuration.

9. Authentication Method Choice

Authentication methods' capabilities and therefore applicability to various environments differ among them. Not all methods provide support for mutual authentication, key derivation or distribution, and DoS attack resiliency that are necessary for operating in insecure networks. Such networks might be susceptible to eavesdropping and spoofing, therefore a stronger authentication method needs to be used to prevent attacks on the client and the network.

The authentication method choice is a function of the underlying security of the network (e.g., physically secured, shared link, etc.). It is the responsibility of the user and the network operator to pick the right method for authentication. PANA carries EAP regardless of the EAP method used. It is outside the scope of PANA to mandate, recommend, or limit use of any authentication methods. PANA cannot increase the strength of a weak authentication method to make it suitable for an insecure environment. There are some EAP-based approaches to achieve this goal (see [I-D.josefsson-pppext-eap-tls-eap],[I-D.ietf-pppext-eap-ttls] ,[I-D.tschofenig-eap-ikev2]). PANA can carry these EAP encapsulating methods but it does not concern itself with how they achieve protection for the weak methods (i.e., their EAP method payloads).

10. Example Cases

10.1 DSL Access Network

In a DSL access network, PANA is seen applicable in the following scenarios.

A typical DSL access consists of a NAS device at the DSL-access provider and a DSL-modem (CPE) at the customer premises. The CPE devices support multiple modes of operation and PANA is applicable in each of these modes.

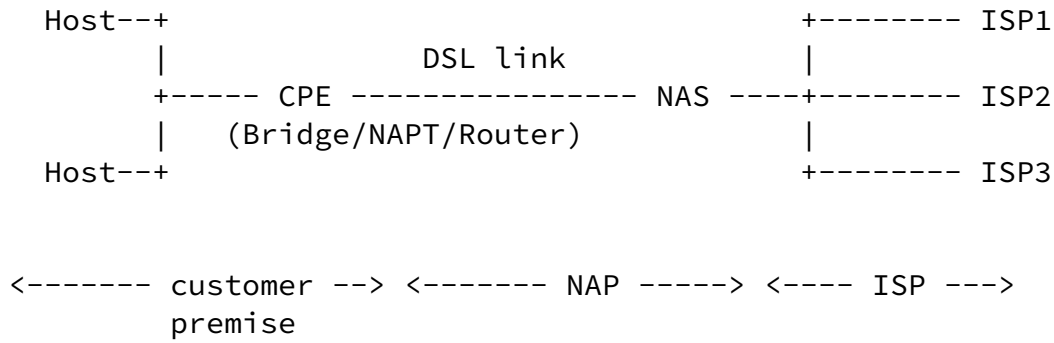

Figure 7: DSL Model

The devices at the customer premises have been shown as "hosts" in the above network.

DSL networks are protected by physical means. Eavesdropping and

spoofing attacks are prevented by keeping the unintended users physically away from the network media. Therefore, generally cryptographic protection of data traffic is not necessary. Nevertheless, if enhanced security is deemed necessary for any reason, IPsec-based access control can be enabled on DSL networks as well by using the method described in [I-D.ietf-pana-ipsec].

10.1.1 Bridging Mode

In the bridging mode, the CPE acts as a simple layer-2 bridge. The hosts at the customer premises will function as clients to obtain addresses from the NAS device by using DHCP or PPPoE.

If PPPoE is used, authentication is typically performed using CHAP or MS-CHAP.

PANA will be applicable when the hosts use DHCP to obtain IP address. DHCP does not support authentication of the devices on either side of the DSL access line. In the simplest method of address assignment, the NAS will allocate the IP address to a host with a lease time reasonably sufficient to complete a full PANA based authentication

which will be triggered immediately after the address assignment. The hosts will perform the PaC function and the NAS will perform the PAA, EP and AR functions.

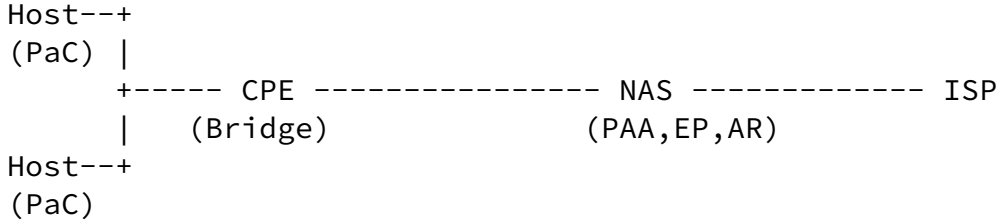

Figure 8: Bridge Mode

The DSL service provider's trunk network should not be accessible to any host that has not successfully completed the PANA authentication phase.

10.1.2 Router Mode

In this mode, the CPE acts as a router for the customer premises network. The CPE itself may obtain the IP address using DHCP or be configured with a static IP address. Once the CPE is authenticated using PANA and is provided access to the service provider's network, the NAS should begin exchanging routing updates with the CPE. All devices at the customer premises will then have access to the service provider's network.

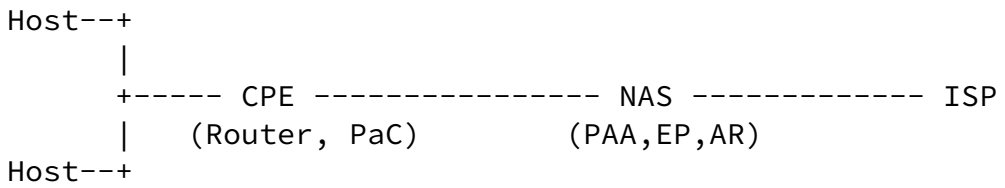

Figure 9: Router Mode

It is possible that both ends of the DSL link are configured with static IP addresses. PANA-based mutual authentication of CPE and NAS is desirable before data-traffic is exchanged between the customer premises network and the service provider network. The CPE router may also use NAPT (Network Address Port Tranlation).

10.1.3 PANA and Dynamic Internet Service Provider Selection

In some installations, a NAS device is shared by multiple service providers. Each service provider configures the NAS with a certain IP address space.

The devices at the customer premises network indicate their choice of

service provider and the NAS chooses the IP address from the appropriate service provider's pool. In many cases, the address is assigned not by the NAS but by the AAA server that is managed fully by the service provider.

This simplifies the management of the DSL access network as it is not always necessary to configure each DSL access line with the service provider's identity. The service provider is chosen dynamically by the CPE device. This is typically known as "dynamic Internet Service Provider selection". The AAA function is usually overloaded to perform dynamic ISP selection.

If the CPE device uses a PPP based protocol (PPP or PPPoE), the ISP is chosen by mapping the username field of a CHAP response to a provider.

If the CPE uses DHCP, the 'client-id' field of the DHCP-discover or DHCP-request packet is mapped to the provider.

10.1.3.1 Selection as Part of the DHCP protocol or an Attribute of DSL Access Line

The ISP selection, therefore the IP address pool, can be conveyed based on the DHCP protocol exchange (as explained earlier), or by associating the DSL access line to the service provider before the PANA authentication begins. When any of these schemes is used, the IP address used during PANA authentication (PRPA) is the ultimate IP address and it does not have to be changed upon successful authorization.

10.1.3.2 Selection as Part of the PANA Authentication

The ISP selection of the client can be explicitly conveyed during the PANA authentication. In that case, the client can be assigned a temporary IP address (PRPA) prior to PANA authentication. This IP address might be obtained via DHCP with a lease reasonably long to complete PANA authentication, or via the zeroconf technique [I-D.ietf-zeroconf-ipv4-linklocal]. In either case, successful PANA authentication signalling prompts the client to obtain a new (long

term) IP address via DHCP. This new IP address (POPA) replaces the previously allocated temporary IP address.

10.2 Wireless LAN Example

This section describes how PANA can be used on WLAN networks. In most common WLAN deployments the IP addresses are dynamically configured. Therefore this section does not cover the scenarios where the IP address is statically configured. There are two models

depending on which layer security is bootstrapped from PANA authentication, L2 or L3. When PANA authentication is used for bootstrapping L3 security, L2 security is not necessarily to exist even after PANA authentication. Instead, IPsec-based data traffic protection is bootstrapped from PANA. The PAA can indicate the PaC as to whether L2 or L3 protection is needed, by including a Protection-Capability AVP in PANA-Bind-Request message. In both cases, the most common deployment would be illustrated in Figure 10, where EP is typically co-located with AP (access point) when PANA is used for bootstrapping L2 security or with AR when PANA is used for bootstrapping L3 security.

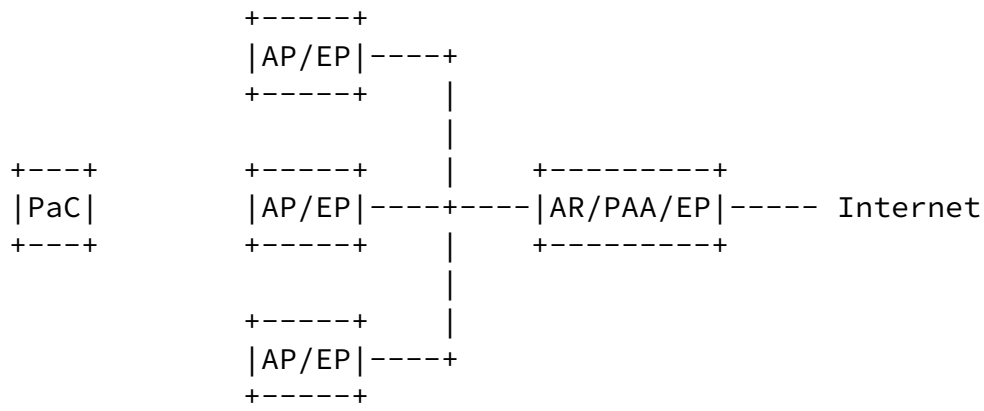

Figure 10: PANA Wireless LAN Model

10.2.1 PANA with Bootstrapping IPsec

In this model, data traffic is protected by using IPsec tunnel mode SA and an IP address is used as the device identifier of PaC (see Section 5 for details). Some or all of AP, DHCPv4 Server (including PRPA DHCPv4 Server and IPsec-TIA DHCPv4 Server), DHCPv6 Server, PAA and EP may be co-located in a single device. EP is always co-located with AR and may be co-located with PAA. When EP and PAA are not co-located, PAA-EP protocol is used for communication between PAA and EP.

Note that for all of the cases described in this section, PBR (PANA-Bind-Request) and PBA (PANA-Bind-Answer) exchange in PANA [I-D.ietf-pana-pana] should occur after installing the authorization parameter to AR, so that IKE can be performed immediately after PANA is successfully completed.

10.2.1.1 IPv4

Case A: IPsec-TIA obtained by using DHCPv4

In this case, the IPsec-TIA and IPsec-TOA are the same as the PRPA, and all configuration information including the IP address is obtained by using DHCPv4 [RFC2131]. No POPA is configured. Case A is the simplest compared to other ones and might be used in a network where IP address depletion attack on DHCP is not a significant concern. The PRPA needs to be a routable address unless NAT is performed on AR.

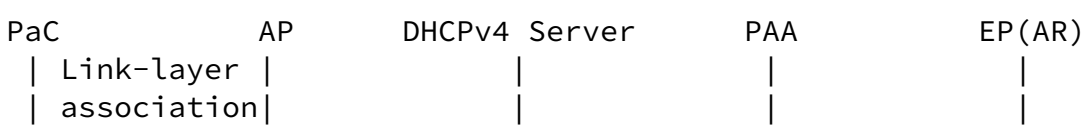

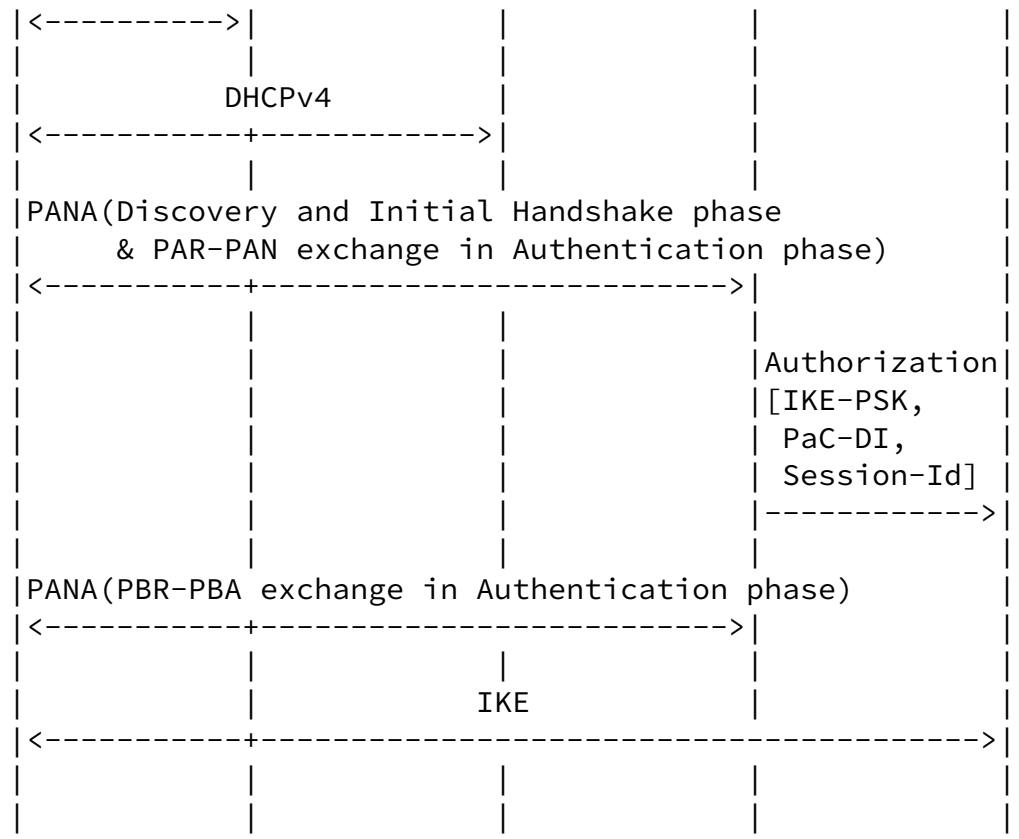

Figure 11: An example case for configuring IPsec-TIA by using DHCPv4

Case B: IPsec-TIA obtained by using IKE

In this case, the PRPA is obtained by using DHCPv4 and used as IPsec-TOA. The POPA is obtained by using IKE (via a Configuration Payload exchange or equivalent) and used as IPsec-TIA.

Case C: IPsec-TIA obtained by using RFC3456

Like Case B, the PRPA is obtained by using DHCPv4. The difference is that the POPA (eventually used as IPsec-TIA) and other configuration parameter are configured by running DHCPv4 over a special IPsec tunnel mode SA [RFC3456]. Note that the PRPA DHCPv4 Server and IPsec-TIA DHCPv4 Server may be co-located on the same node.
Note: this case may be used only when IKEv1 is used as the IPsec key management protocol (IKEv2 does not seem to support RFC3456 equivalent case).

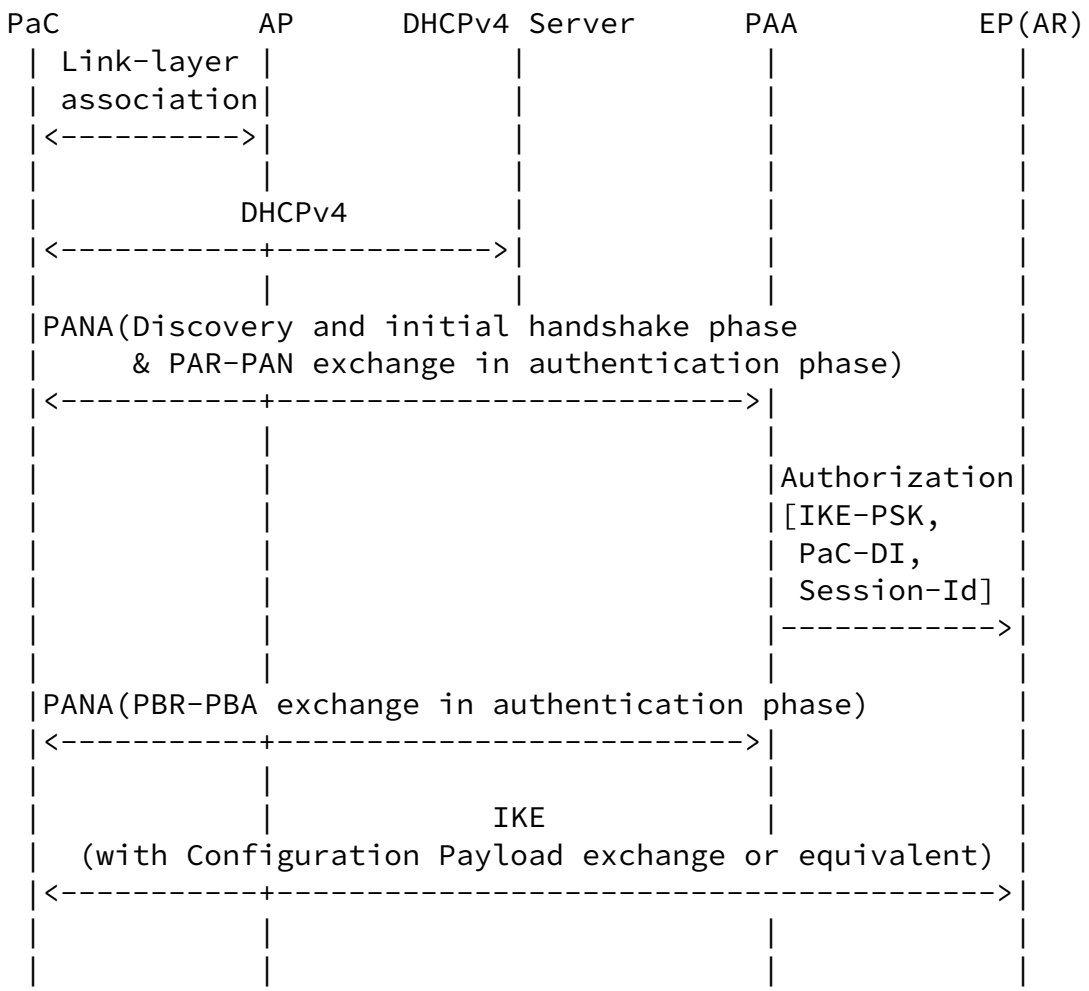

Figure 12: An example case for IPsec-TIA obtained by using IKE

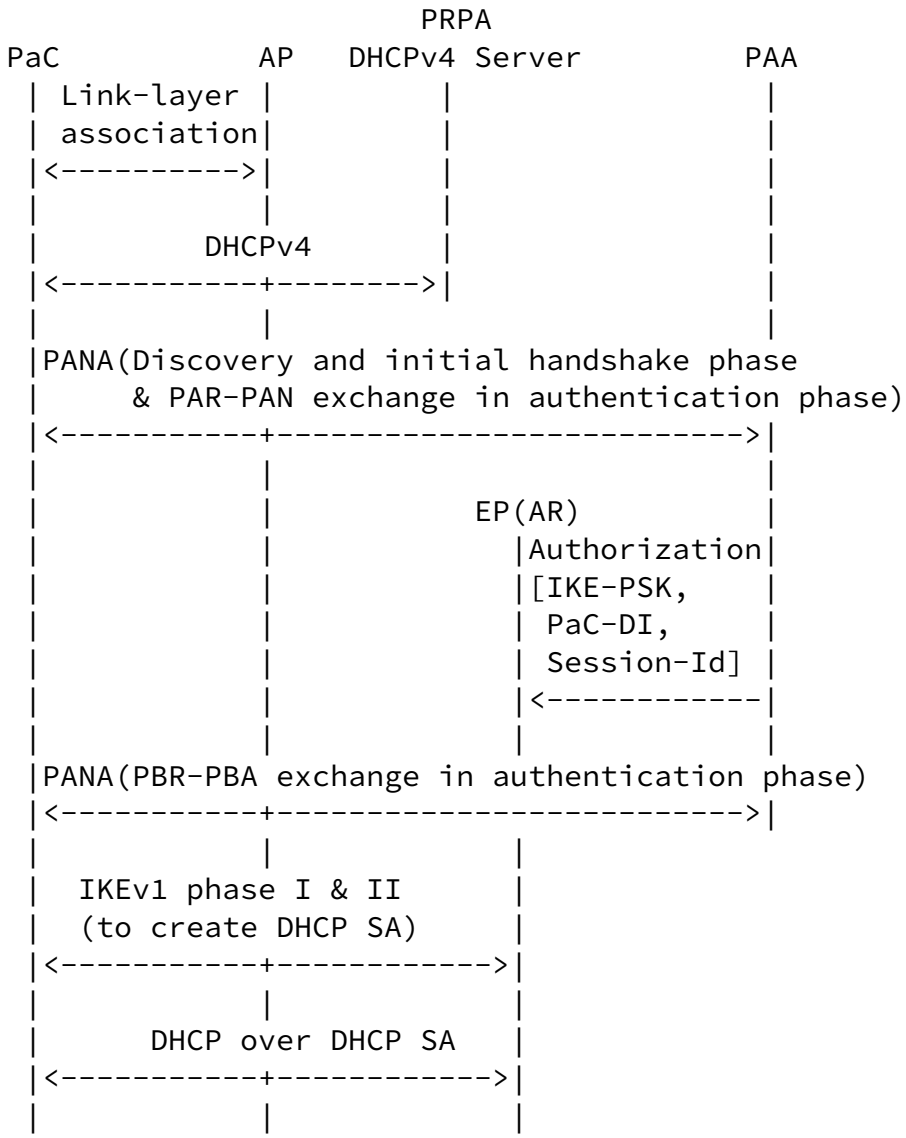
PRPA
PaC
AP
DHCPv4 Server
PAA
Link-layer association
DHCPv4
PANA(Discovery and initial handshake phase
& PAR-PAN exchange in authentication phase)
EP(AR)
Authorization
[IKE-PSK,
PaC-DI,
Session-Id]
PANA(PBR-PBA exchange in authentication phase)
IKEv1 phase I & II
(to create DHCP SA)
DHCP over DHCP SA

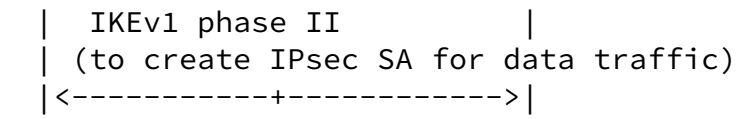

Figure 13: An example case for configuring IPsec-TIA by using RFC3456

10.2.1.2 IPv6

In the case of IPv6, the IPsec-TOA (PRPA) is the IPv6 link-local address. IPsec-TIA (POPA) is obtained by using Configuration Payload exchange of IKE version 2 (Note that there is no standard method for configuring IPsec-TIA in IKEv1). Other configuration information may be obtained in the same Configuration Payload exchange or may be obtained by running an additional DHCPv6.

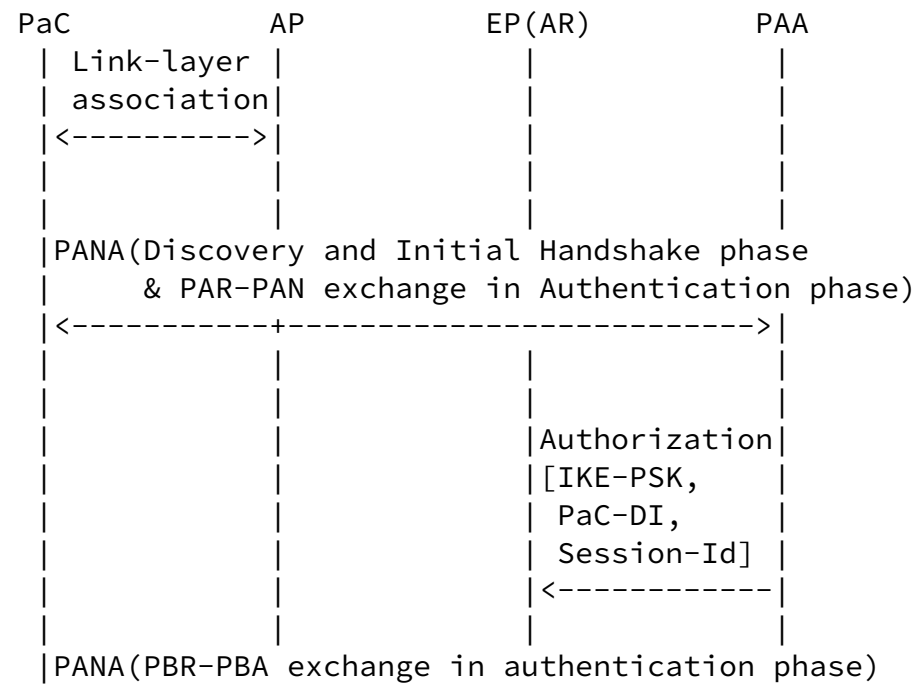

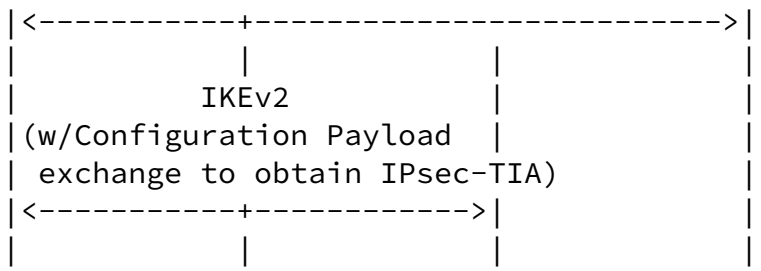

Figure 14: An example sequence for configuring IPsec-TIA in IPv6

10.2.2 PANA with Bootstrapping WPA/IEEE 802.11i

In this model, PANA is used for authentication and authorization, and L2 ciphering is used for access control, the latter is enabled by the former. The L2 ciphering is based on using PSK (Pre-Shared Key) mode of WPA (Wi-Fi Protected Access) [WPA] or IEEE 802.11i [802.11i], which is derived from the EAP MSK as a result of successful PANA authentication. In this document, the pre-shared key shared between station and AP is referred to as PMK (Pair-wise Master Key). In this model, MAC address is used as the device identifier in PANA.

This model allows the separation of PAA from APs (EPs). A typical purpose of using this model is to reduce AP management cost by allowing physical separation of RADIUS/Diameter client from access points, where AP management can be a significant issue when deploying a large number of access points.

By bootstrapping PSK mode of WPA and IEEE 802.11i from PANA it is also possible to improve wireless LAN security by providing protected

disconnection procedure at L3.

This model does not require any change in the current WPA and IEEE

802.11i specifications. This also means that PANA doesn't provide any L2 security features beyond those already provided for in WPA and IEEE 802.11i.

The IEEE 802.11 specification [802.11] allows Class 1 data frames to be received in any state. Also, the latest version of IEEE 802.11i [802.11i] optionally allows higher-layer data traffic to be received and processed on their IEEE 802.1X Uncontrolled Ports. This feature allows processing IP-based traffic (such as ARP, IPv6 neighbor discovery, DHCP, and PANA) on IEEE 802.1X Uncontrolled Port prior to client authentication. (Note: WPA and its corresponding version of IEEE 802.11i draft do not explicitly define this operation, so it may be safer not to use this in WPA).

Until the PaC is successfully authenticated, only a selected type of IP traffic is allowed over the IEEE 802.1X Uncontrolled Port. Any other IP traffic is dropped on the AP without being forwarded to the DS (Distribution System). Upon successful PANA authentication, the traffic switches to the controlled port. Host configuration, including obtaining an (potentially new) IP address, takes place on this port. Usual DHCP-based, and also in the case of IPv6 stateless autoconfiguration, mechanism is available to the PaC. After this point, the rest of the IP traffic, including PANA exchanges, are processed on the controlled port.

When a PaC does not have a PMK for the AP, the following procedure is taken:

1. The PaC associates with the AP.

2. The PaC configures a PRPA by using DHCP (in the case of IPv4) or configures a link-local address (in the case of IPv6), and then runs PANA by using the address.

3. Upon successful authentication, the PaC obtains a PMK for each AP controlled by the PAA.

4. The AP initiates IEEE 802.11i 4-way handshake to establish a PTK (Pair-wise Transient Key) with the PaC, by using the PMK.

5. The PaC obtains a POPA by using any method that the client normally uses.

10.2.3 Capability Discovery

When a PaC is a mobile, there may be multiple APs available in its vicinity. Each AP are connected to one of the following types of access networks.

a) Free access network

There is no IEEE 802.1X or PANA authentication in this access network.

b) PANA-secured network

There is PANA authentication in this access network.

c) IEEE 802.1X-secured network

There is IEEE 802.1X authentication in this access network.

Type (c) is distinguished from others by checking the capability information advertised in IEEE 802.11 Beacon frames (IEEE 802.11i defines RSN Information Element for this purpose). Types (a) and (b) are not distinguishable until the PaC associates with the AP, get an IP address, and engage in PANA discovery. The default PaC behavior would be to act as if this is a free network and attempt DHCP. This would be detected by the access network and trigger unsolicited PANA discovery. A type (b) network would send a PANA-Start-Request to the client and block general purpose data traffic. This helps the client discover whether the network is type (a) or type (b). Or if the PaC

is pre-provisioned with the information that this is a PANA enabled network, it can attempt PAA discovery immediately. The PaC behavior after connecting to an AP of type (b) network is described in Section 10.2, Section 10.2.1 and Section 10.2.2.

11. Open Issue

Certain combination of network environments and timing cases may require further considerations.

If PaC changes access point right after a successful PANA authorization but before POPA configuration, it might be confused whether the new IP address configured via the new access point is the POPA of the ongoing session or a PRPA of a new session. When the PRPA and POPA types or configuration mechanisms are different this is not a problem. One possible combination where such clues are not available is when PaC moves from a Case A of Section 10.2.1.1 (IPsec-TIA obtained by using DHCPv4) network to a Case D of Appendix A.1 (IPsec-TIA obtained by using RFC3118) network.

In another example, when PaC switches from one wired Ethernet hub to another (without the use of bootstrapping IPsec) before configuring

the POPA. In this case, PaC may not able to know whether an obtained address is the POPA in the same subnet or a PRPA in a new subnet.

For these cases, a mechanism to remove the ambiguity (PRPA vs. POPA) may need to be defined.

12. Acknowledgments

We would like to thank Bernard Aboba, Yacine El Mghazli, Randy Turner and Hannes Tschofenig for their valuable comments.

13. References

13.1 Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

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

[I-D.ietf-zeroconf-ipv4-linklocal] Aboba, B., "Dynamic Configuration of Link-Local IPv4 Addresses", draft-ietf-zeroconf-ipv4-linklocal-17 (work in progress), July 2004.

[RFC2462] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.

[RFC2461] Narten, T., Nordmark, E. and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.

[RFC3484] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", RFC 3484, February 2003.

[RFC2409] Harkins, D. and D. Carrel, "The Internet Key Exchange (IKE)", RFC 2409, November 1998.

[I-D.ietf-ipsec-ikev2] Kaufman, C., "Internet Key Exchange (IKEv2) Protocol", draft-ietf-ipsec-ikev2-14 (work in progress), June 2004.

[I-D.ietf-pana-snmp] Mghazli, Y., Ohba, Y. and J. Bournelle, "SNMP usage for PAA-2-EP interface", draft-ietf-pana-snmp-00 (work in progress), April 2004.

[I-D.ietf-pana-pana]
Forsberg, D., Ohba, Y., Patil, B., Tschofenig, H. and A. Yegin, "Protocol for Carrying Authentication for Network Access (PANA)", draft-ietf-pana-pana-04 (work in progress), May 2004.

[I-D.ietf-pana-ipsec]
Parthasarathy, M., "PANA enabling IPsec based Access Control", draft-ietf-pana-ipsec-03 (work in progress), May 2004.

[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.

[RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.

[RFC3456] Patel, B., Aboba, B., Kelly, S. and V. Gupta, "Dynamic Host Configuration Protocol (DHCPv4) Configuration of IPsec Tunnel Mode", RFC 3456, January 2003.

13.2 Informative References

[RFC3377] Hodges, J. and R. Morgan, "Lightweight Directory Access Protocol (v3): Technical Specification", RFC 3377, September 2002.

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

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

[RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G. and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

[I-D.ietf-eap-netsel-problem] Arkko, J. and B. Aboba, "Network Discovery and Selection Problem", draft-ietf-eap-netsel-problem-00 (work in progress), January 2004.

[RFC2486] Aboba, B. and M. Beadles, "The Network Access Identifier", RFC 2486, January 1999.

[I-D.ietf-pana-requirements] Yegin, A. and Y. Ohba, "Protocol for Carrying Authentication for Network Access (PANA)Requirements", draft-ietf-pana-requirements-08 (work in progress), June 2004.

[I-D.ietf-aaa-eap] Eronen, P., Hiller, T. and G. Zorn, "Diameter Extensible Authentication Protocol (EAP) Application", draft-ietf-aaa-eap-08 (work in progress), June 2004.

[I-D.yegin-eap-boot-rfc3118] Yegin, A., Tschofenig, H. and D. Forsberg, "Bootstrapping RFC3118 Delayed DHCP Authentication Using EAP-based Network Access Authentication", draft-yegin-eap-boot-rfc3118-00 (work in progress), February 2004.

[RFC3118] Droms, R. and W. Arbaugh, "Authentication for DHCP Messages", RFC 3118, June 2001.

[I-D.josefsson-pppext-eap-tls-eap] Josefsson, S., Palekar, A., Simon, D. and G. Zorn, "Protected EAP Protocol (PEAP)", draft-josefsson-pppext-eap-tls-eap-07 (work in progress), October 2003.

[I-D.ietf-pppext-eap-ttls] Funk, P. and S. Blake-Wilson, "EAP Tunneled TLS Authentication Protocol (EAP-TTLS)", draft-ietf-pppext-eap-ttls-04 (work in progress), April 2004.

[I-D.tschofenig-eap-ikev2] Tschofenig, H. and D. Kroeselberg, "EAP IKEv2 Method (EAP-IKEv2)", draft-tschofenig-eap-ikev2-03 (work in progress), February 2004.

[DSL] DSL Forum Architecture and Transport Working Group, "DSL Forum TR-058 Multi-Service Architecture and Framework Requirements", September 2003.

[802.11i] Institute of Electrical and Electronics Engineers, "Draft supplement to standard for telecommunications and information exchange between systems - lan/man specific requirements - part 11: Wireless medium access control (mac) and physical layer (phy) specifications: Specification for enhanced security", IEEE 802.11i/D10.0, 2004.

[802.11] Institute of Electrical and Electronics Engineers, "Information technology - telecommunications and information exchange between systems - local and metropolitan area networks - specific requirements part 11: Wireless lan medium access control (mac) and physical layer (phy) specifications", IEEE Standard 802.11, 1999(R2003).

[WPA] The Wi-Fi Alliance, "WPA (Wi-Fi Protected Access)", Wi-Fi WPA v2.0, 2003.

Authors' Addresses

Prakash Jayaraman
Network Equipment Technologies, Inc.
6900 Paseo Padre Parkway
Fremont, CA 94555
USA

Phone: +1 510 574 2305
EMail: prakash_jayaraman@net.com

Rafa Marin Lopez
University of Murcia
30071 Murcia
Spain

EMail: rafa@dif.um.es

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

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

Mohan Parthasarathy
Nokia
313 Fairchild Drive

Mountain View, CA 94043
USA

Phone: +1 408 734 8820
EMail: mohanp@sbcglobal.net

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

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

Appendix A. Other Possible Cases for PANA with Bootstrapping IPsec in Wireless LAN

This section describes other possible cases for PANA with Bootstrapping IPsec in wireless LAN environments.

Appendix A.1 IPv4

Case D: IPsec-TIA obtained by using RFC3118

In this case, the POPA is configured and used as both IPsec-TIA and IPsec-TOA. The IPsec-TIA is assigned by using RFC3118 (authenticated DHCP) [RFC3118] before running IKE. The DHCP-PSK needed for authenticated DHCP is distributed from the PAA to the POPA DHCPv4 server by using the method specified in [I-D.yegin-eap-boot-rfc3118]. The PRPA is assigned by using DHCPv4 and may be assigned with a short lease period in order to

provide some level of robustness against IP address depletion attack. The IPsec-TIA is bound to an IPsec SA by using specifying the IPsec-TIA as the SA Identification in IKEv1 phase II or IKEv2 CREATE_CHILD_SA exchange as specified in [I-D.ietf-pana-ipsec]. Once the IPsec-TIA is obtained, the PANA re-authentication based on PUR (PANA-Update-Rquest) and PUA (PANA-Update-Answer) exchange is performed with using the obtained IPsec-TIA in order to inform PAA of the update of PaC-DI. The IKE procedure should occur after the PUR-PUA exchange procedure. The PaC unconfigures the PRPA immediately after the IPsec-TIA is obtained.

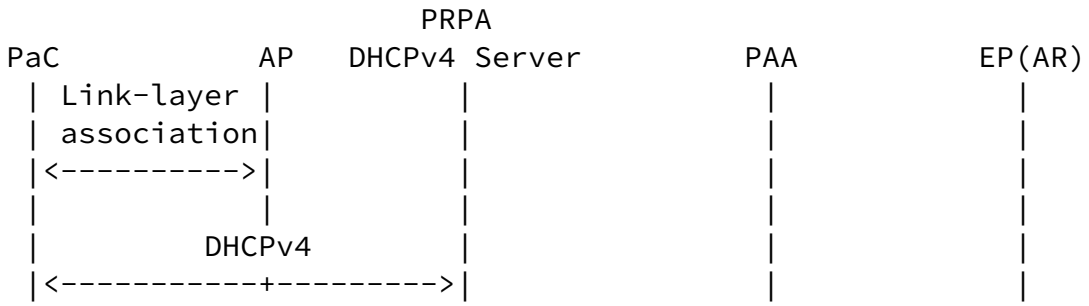

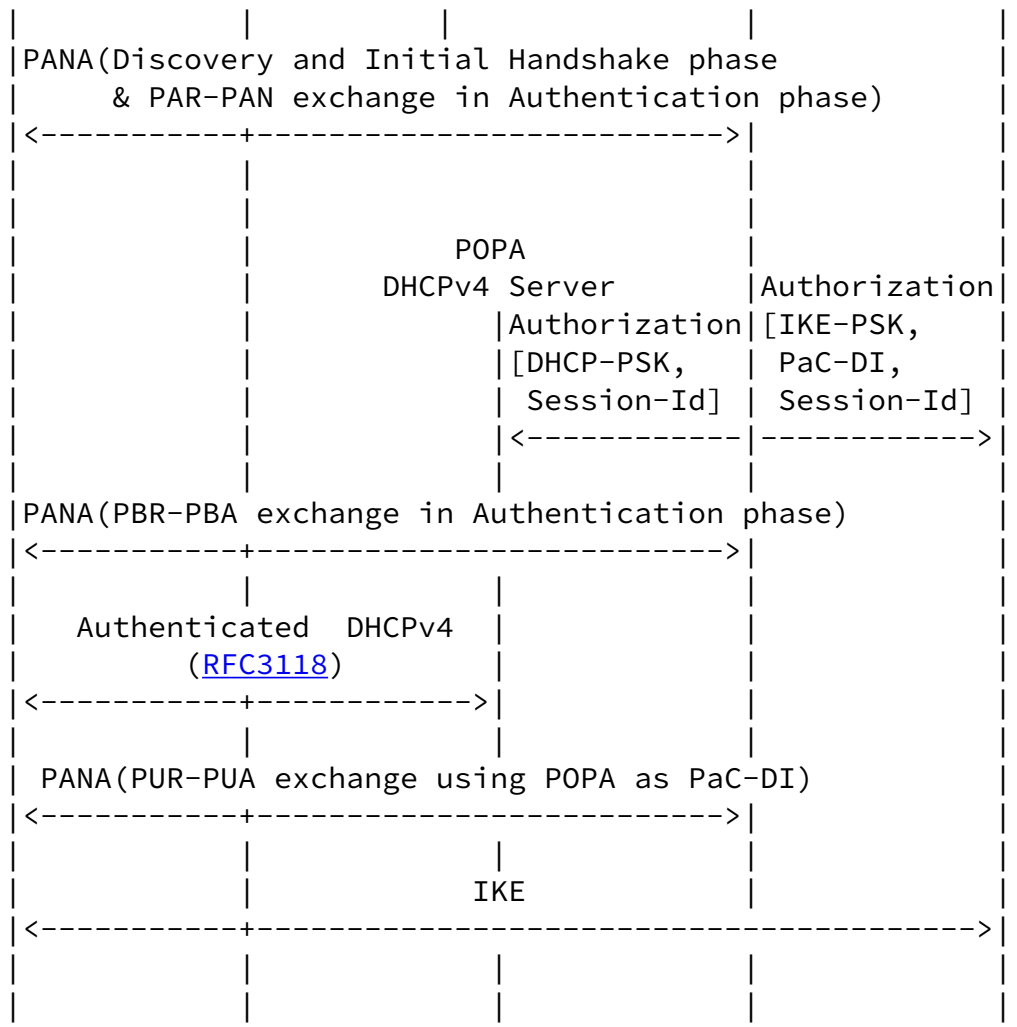

Figure 15: An example case for IPsec-TIA obtained by using RFC3118

Appendix A.2 IPv6

Case A: IPsec-TIA obtained by using DHCPv6

This case is similar to Case A in IPv4, except that a link-local address is used as the PRPA and IPsec-TOA, and that the DHCPv6 procedure can occur at any time after link-layer association and before IKE.

This case is not recommended since there is an ambiguity on whether IPv6 Neighbor Discovery for the POPA should run on the physical interface or inside the IPsec tunnel or both.

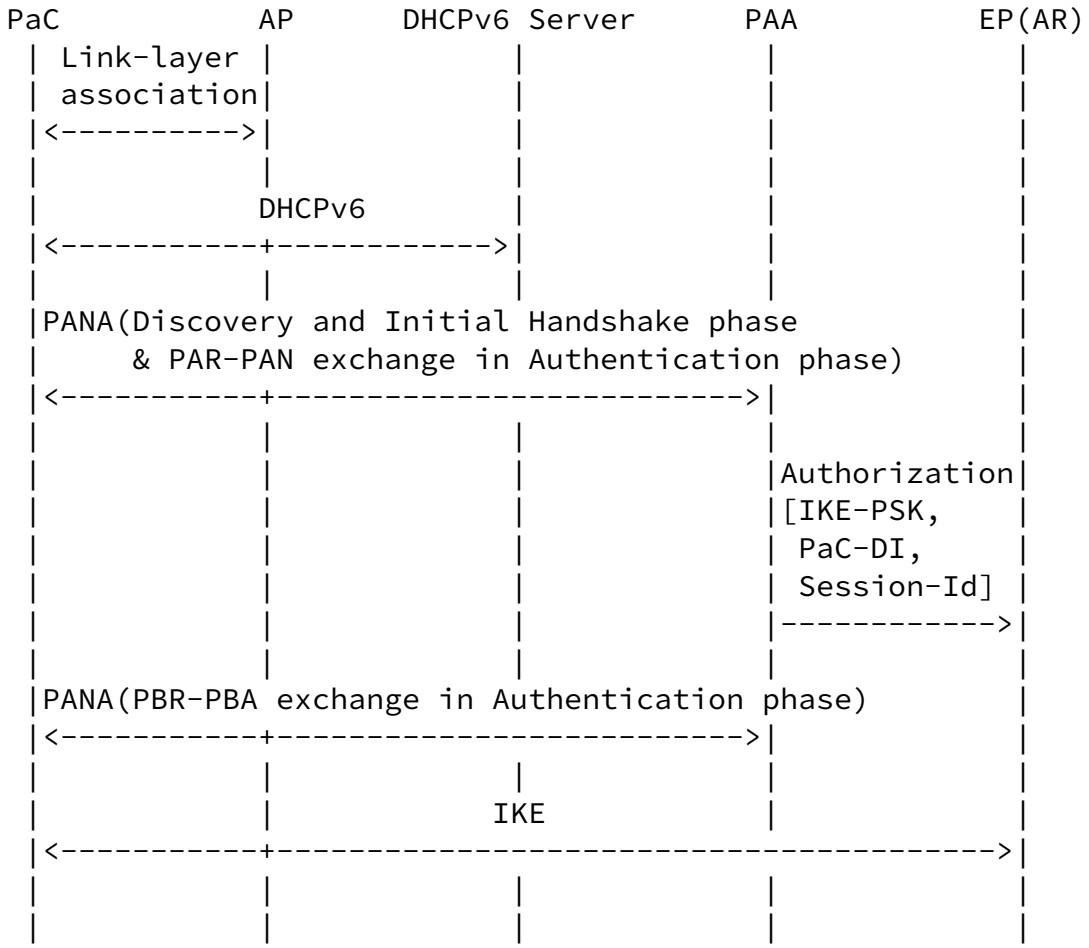

Figure 16: An example case for IPsec-TIA obtained by using DHCPv6

Case B: IPsec-TIA obtained by using IPv6 stateless address autoconfiguration

In this case, the IPsec-TOA (link-local address) and IPsec-TIA (global address) are configured through IPv6 stateless address autoconfiguration before running IKE. Other configuration information can be obtained by using several methods including authenticated DHCPv6, Configuration Payload exchange and DHCPv6 over IPsec SA.

This case is not recommended for the same reason as Case A of IPv6.

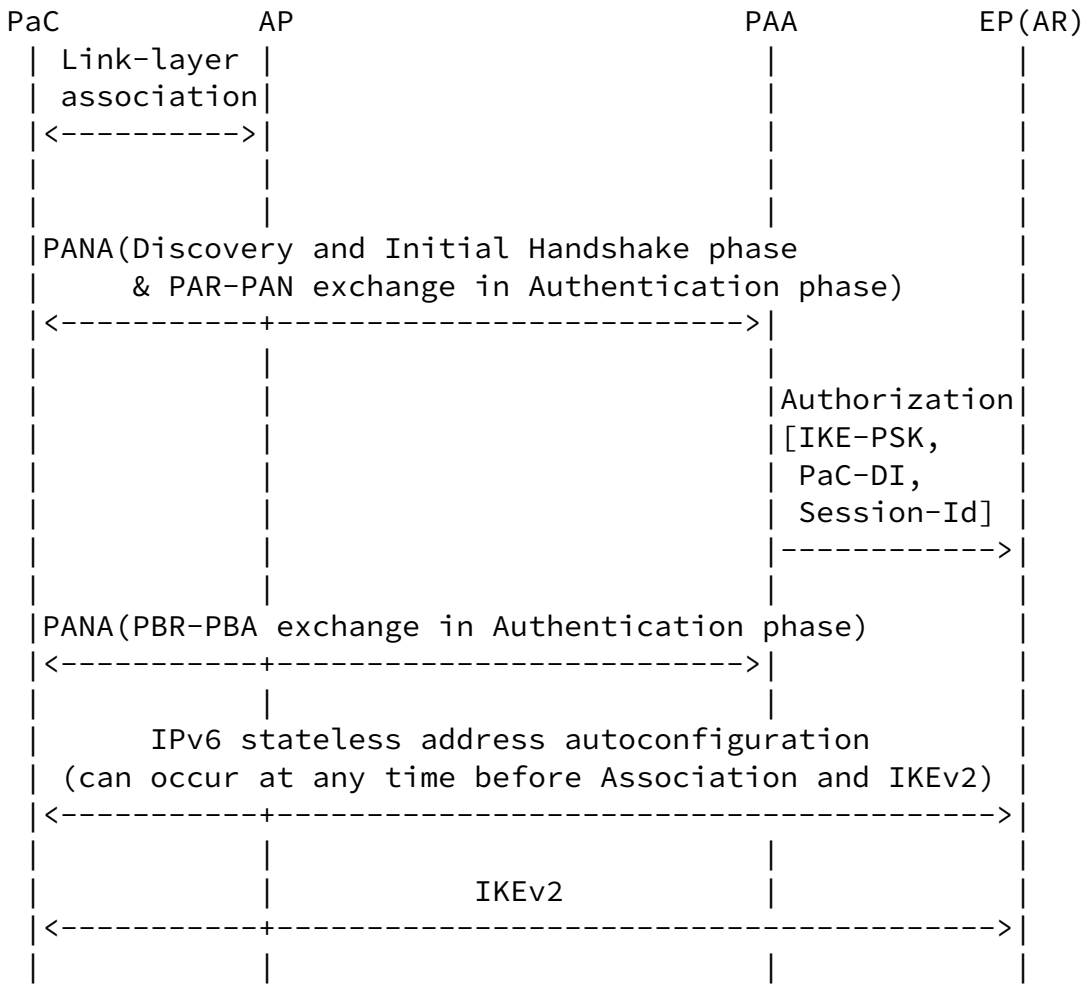

Figure 17: An example sequence for IPsec-TIA obtained by using IPv6 stateless address autoconfiguration

Case C: IPsec-TIA obtained by using authenticated DHCPv6

This case is similar to Case C of IPv4, except that a link-local address is used as the PRPA, and that there is no need for additional PUR-PUA exchange to update the PaC-DI.

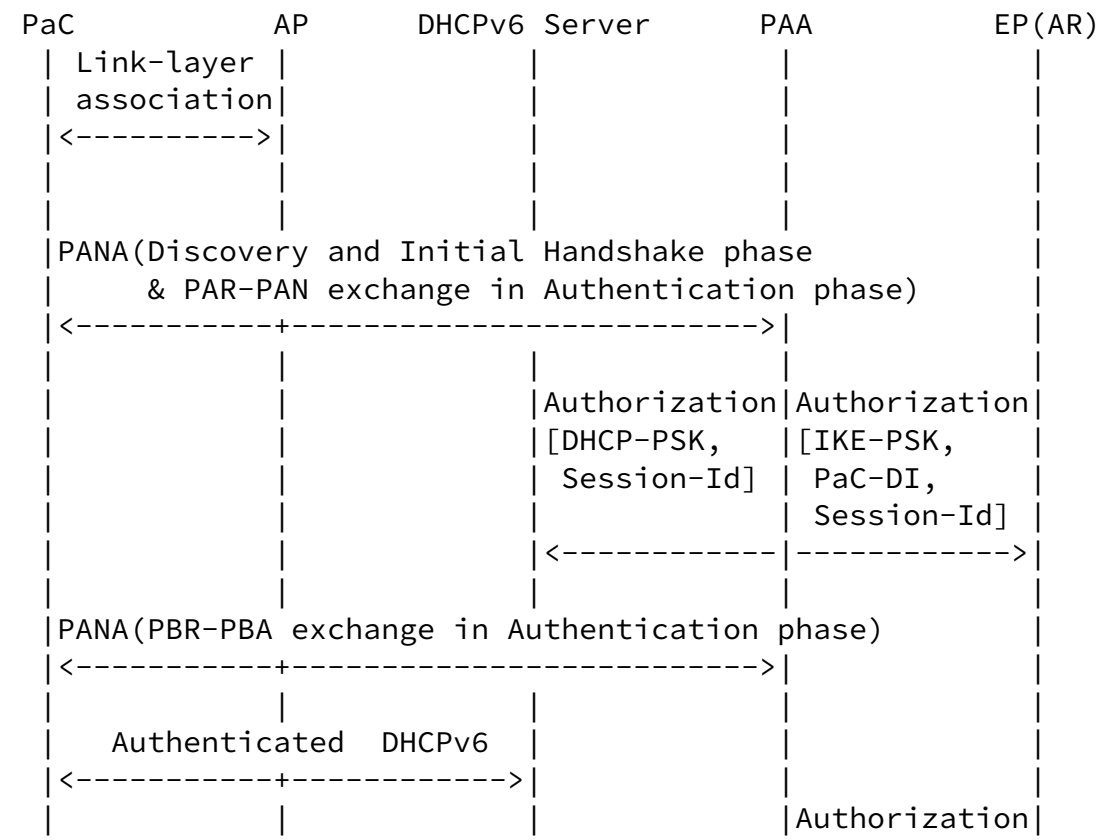

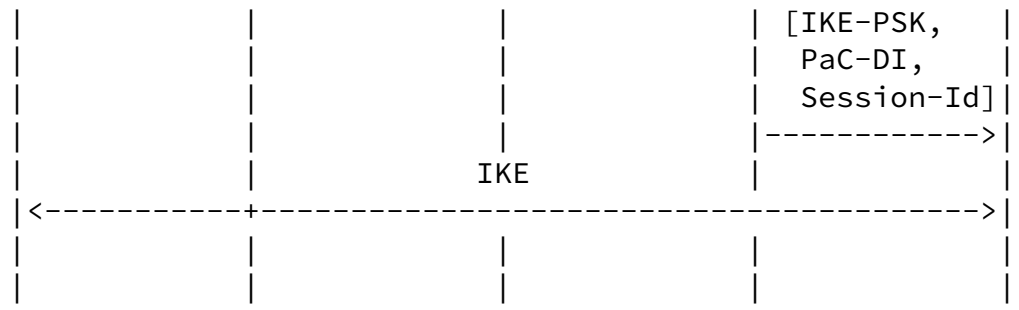

Figure 18: An example case for configuring IPsec-TIA by using authenticated DHCPv6

Intellectual Property Statement

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.

Disclaimer of Validity

Copyright Statement

Acknowledgment

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