

Protocol for Carrying Authentication and Network Access (PANA)
Threat Analysis and Security Requirements

Status of This Memo

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

Copyright Notice

Copyright (C) The Internet Society (2005).

Abstract

This document discusses the threats to protocols used to carry authentication for network access. The security requirements arising from these threats will be used as additional input to the Protocol for Carrying Authentication for Network Access (PANA) Working Group for designing the IP based network access authentication protocol.

Table of Contents

1. Introduction	2
2. Keywords	2
3. Terminology and Definitions.	2
4. Usage Scenarios.	3
5. Trust Relationships.	4
6. Threat Scenarios	5
6.1. PAA Discovery.	6
6.2. Authentication	6
6.3. PaC Leaving the Network.	9
6.4. Service Theft.	10
6.5. PAA-EP Communication.	11
6.6. Miscellaneous Attacks.	12
7. Summary of Requirements.	13
8. Security Considerations.	13
9. Normative References	14
10. Informative References	14
11. Acknowledgements	14
Author's Address	14
Full Copyright Statement	15

1. Introduction

The Protocol for Carrying Authentication for Network Access (PANA) Working Group is developing methods for authenticating clients to the access network using IP based protocols. This document discusses the threats to such IP based protocols.

A client wishing to get access to the network must carry on multiple steps. First, it needs to discover the IP address of the PANA authentication agent (PAA) and then execute an authentication protocol to authenticate itself to the network. Once the client is authenticated, there might be other messages exchanged during the lifetime of the network access. This document discusses the threats in these steps without discussing any solutions. The requirements arising out of these threats will be used as input to the PANA Working Group. The use of word co-located in this document means that the referred entities are present on the same node.

2. Keywords

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 [KEYWORDS].

3. Terminology and Definitions

Client Access Device

A network element (e.g., notebook computer, PDA) that requires access to a provider's network.

Network Access Server (NAS)

Network device that provides access to the network.

PANA Client (PaC)

An entity in the edge subnet that seeks to obtain network access from a PANA authentication agent within a network. A PANA client is associated with a device and a set of credentials to prove its identity within the scope of PANA.

PANA Authentication Agent (PAA)

An entity whose responsibility is to authenticate the PANA client and to grant network access service to the client's device.

Authentication Server (AS)

An entity that authenticates the PANA client. It may be co-located with the PANA authentication agent or part of the back-end infrastructure.

Device Identifier (DI)

The identifier used by the network to control and police the network access of a client. Depending on the access technology, the identifier might contain the IP address, link-layer address, switch port number, etc., of a device. The PANA authentication agent keeps a table for binding device identifiers to the PANA clients. At most one PANA client should be associated with a DI on a PANA authentication agent.

Enforcement Point (EP)

A node capable of filtering packets sent by the PANA client by using the DI information authorized by PANA authentication agent.

Compound methods

Authentication protocol in which methods are used in a sequence one after another or in which methods are tunneled inside another independently established tunnel between the client and server [TUN-EAP].

4. Usage Scenarios

PANA is intended to be used in an environment where there is no a priori trust relationship or security association between the PaC and other nodes, such as the PAA and EP. In these environments, one may observe the following:

- o The link between PaC and PAA may be a shared medium (e.g., Ethernet) or may not be a shared medium (e.g., DSL network).
- o All the PaCs may be authenticated to the access network at layer 2 (e.g., 3GPP2 CDMA network) and share a security association with a layer 2 authentication agent (e.g., BTS). The PaCs still don't trust each other; any PaC can pretend to be a PAA, spoof IP addresses, and launch various other attacks.

The scenarios mentioned above affect the threat model of PANA. This document discusses the various threats in the context of the above network access scenarios for a better understanding of the threats. In the following discussion, any reference to a link that is not

shared (or non-shared) is assumed to be physically secure. If such an assumption cannot be made about the link, then the case becomes the same as that for a link being shared by more than one node.

5. Trust Relationships

PANA authentication involves a client (PaC), a PANA agent (PAA), an Authentication server (AS), and an Enforcement point (EP). The AS here refers to the AAA server that resides in the home realm of the PaC.

The entities that have a priori trust relationships before PANA begins are as follows:

- 1) PAA and AS: The PaC belonging to the same administrative domain that the AS does often has to use resources provided by a PAA that belongs to another administrative domain. A PAA authenticates the PaC before providing local network access. The credentials provided by the PaC for authentication may or may not be understood by the PAA. If the PAA does not understand the credentials, it needs to communicate with the AS in a different domain to verify the credentials. The threats in the communication path between the PAA and AS are already covered in [RAD-EAP]. To counter these threats, the communication between the PAA and AS is secured by using a static or dynamic security association.
- 2) PAA and EP: The PAA and EP belong to the same administrative domain. Hence, the network operator can set up a security association to protect the traffic exchanged between them. This document discusses the threats in this path.
- 3) PaC and AS: The PaC and AS belong to the same administrative domain and share a trust relationship. When the PaC uses a different domain than its home for network access, it provides its credentials to the PAA in the visited network for authentication. The information provided by the PaC traverses the PaC-PAA and PAA-AS paths. The threats in the PAA-AS path are already discussed in [RAD-EAP]. This document discusses the threats in the PaC-PAA path.

It is possible that some of the entities such as the PAA, AS, and EP are co-located. In those cases, it can be safely assumed that there are no significant external threats in their communication.

The entities that do not have any trust relationship before PANA begins are as follows:

- 1) PaC and PAA: The PaC and PAA normally belong to two different administrative domains. They do not necessarily share a trust relationship initially. They establish a security association in the process of authentication. All messages exchanged between the PaC and PAA are subject to various threats, which are discussed in this document.
- 2) PaC and EP: The EP belongs to the same administrative domain as the PAA. Hence, the PaC and EP do not necessarily share a trust relationship initially. When the PaC is successfully authenticated, it may result in key establishment between the PaC and PAA, which can be further used to secure the link between the PaC and EP. For example, the EAP keying framework, [EAP-KEY], defines a three party EAP exchange in which the clients derive the transient sessions keys to secure the link between the peer and NAS in their final step. Similarly, PANA will provide the ability to establish keys between the PaC and EP that can be used to secure the link further. This is discussed further in Section 6.4 below.

6. Threat Scenarios

First, the PaC needs to discover the PAA. This involves either sending solicitations or waiting for advertisements. Once it has discovered the PAA, the two will enter authentication exchange. Once the access is granted, the PaC will most likely exchange data with other nodes in the Internet. These steps are vulnerable to man-in-the-middle (MITM), denial of service (DoS), and service theft attacks, which are discussed below.

The threats are grouped by the various stages the client goes through to gain access to the network. Section 6.1 discusses the threats related to PAA discovery. Section 6.2 discusses the threats related to authentication itself. Section 6.3 discusses the threats involved when leaving the network. Section 6.4 discusses service theft. Section 6.5 discusses the threats in the PAA-EP path. Section 6.6 discusses the miscellaneous threats.

Some of the threats discussed in the following sections may be specific to shared links. The threat may be absent on non-shared links. Hence, it is only required to prevent the threat on shared links. Instead of specifying a separate set of requirements for shared links and non-shared links, this document specifies one set of requirements with the following wording: "PANA MUST be able to prevent threat X". This means that the PANA protocol should be capable of preventing threat X. The feature that prevents threat X may or may not be used depending on the deployment.

6.1. PAA Discovery

The PAA is discovered by sending solicitations or receiving advertisements. The following are possible threats.

T6.1.1: A malicious node can pretend to be a PAA by sending a spoofed advertisement.

In existing dial-up networks, the clients authenticate to the network but generally do not verify the authenticity of the messages coming from Network Access Server (NAS). This mostly works because the link between the device and the NAS is not shared with other nodes (assuming that nobody tampers with the physical link), and clients trust the NAS and the phone network to provide the service. Spoofing attacks are not present in this environment, as the PaC may assume that the other end of the link is the PAA.

In environments where the link is shared, this threat is present, as any node can pretend to be a PAA. Even if the nodes are authenticated at layer 2, the threat remains present. It is difficult to protect the discovery process, as there is no a priori trust relationship between the PAA and PaC. In deployments where EP can police the packets that are sent among the PaCs, it is possible to filter out the unauthorized PANA packets (e.g., PAA advertisements sent by PaC) to prevent this threat.

The advertisement may be used to include information (such as supported authentication methods) other than the discovery of the PAA itself. This can lead to a bidding down attack, as a malicious node can send a spoofed advertisement with capabilities that indicate authentication methods less secure than those that the real PAA supports, thereby fooling the PaC into negotiating an authentication method less secure than would otherwise be available.

Requirement 1

PANA MUST not assume that the discovery process is protected.

6.2. Authentication

This section discusses the threats specific to the authentication protocol. Section 6.2.1 discusses the possible threat associated with success/failure indications that are transmitted to PaC at the end of the authentication. Section 6.2.2 discusses the man-in-the-middle attack when compound methods are used. Section 6.2.3 discusses the replay attack, and Section 6.2.4 discusses the device identifier attack.

6.2.1. Success or Failure Indications

Some authentication protocols (e.g., EAP) have a special message to indicate success or failure. An attacker can send a false authentication success or failure message to the PaC. By sending a false failure message, the attacker can prevent the client from accessing the network. By sending a false success message, the attacker can prematurely end the authentication exchange, effectively denying service for the PaC.

If the link is not shared, then this threat is absent, as ingress filtering can prevent the attacker from impersonating the PAA.

If the link is shared, it is easy to spoof these packets. If layer 2 provides per-packet encryption with pair-wise keys, it might make it hard for the attacker to guess the success or failure packet that the client would accept. Even if the node is already authenticated at layer 2, it can still pretend to be a PAA and spoof the success or failure.

This attack is possible if the success or failure indication is not protected by using a security association between the PaC and the PAA. In order to avoid this attack, the PaC and PAA should mutually authenticate each other. In this process, they should be able to establish keys to protect the success or failure indications. It may not always be possible to protect the indication, as the keys may not be established prior to transmitting the success or failure packet. If the client is re-authenticating to the network, it can use the previously established security association to protect the success or failure indications. Similarly, all PANA messages exchanged during the authentication prior to key establishment may not be protected.

Requirement 2

PANA MUST be able to mutually authenticate the PaC and PAA. PANA MUST be able to establish keys between the PaC and PAA to protect the PANA messages.

6.2.2. MITM Attack

A malicious node can claim to be the PAA to the real PaC and claim to be the PaC to the real PAA. This is a man-in-the-middle (MITM) attack, whereby the PaC is fooled to think that it is communicating with the real PAA and the PAA is fooled to think that it is communicating with the real PaC.

If the link is not shared, this threat is absent, as ingress filtering can prevent the attacker from acting as a man-in-the-middle.

If the link is shared, this threat is present. Even if the layer 2 provides per-packet protection, the attacker can act as a man-in-the-middle and launch this attack. An instance of MITM attack, in which compound authentication methods are used is described in [TUN-EAP]. In these attacks, the server first authenticates to the client. As the client has not proven its identity yet, the server acts as the man-in-the-middle, tunneling the identity of the legitimate client to gain access to the network. The attack is possible because there is no verification that the same entities participated among the compound methods. It is not possible to do such verification if compound methods are used without being able to create a cryptographic binding among them. This implies that PANA will be vulnerable to such attacks if compound methods are used without being able to cryptographically bind them. Note that the attack does not exist if the keys derived during the tunnel establishment are not used to authenticate the client (e.g., tunnel keys are used for just protecting the identity of the client).

Requirement 3

When compound authentication methods are used in PANA, the methods MUST be cryptographically bound.

6.2.3. Replay Attack

A malicious node can replay the messages that caused authentication failure or success at a later time to create false failures or success. The attacker can also potentially replay other messages of the PANA protocol to deny service to the PaC.

If the link is not shared, this threat is absent, as ingress filtering can prevent the attacker from impersonating the PAA to replay the packets.

If the link is shared, this threat is present. If the packets are encrypted at layer 2 by using pair-wise keys, it will make it hard for the attacker to learn the unencrypted (i.e., original) packet that needs to be replayed. Even if layer 2 provides replay protection, the attacker can still replay the PANA messages (layer 3) for denying service to the client.

Requirement 4

PANA MUST be able to protect itself against replay attacks.

6.2.4. Device Identifier Attack

When the client is successfully authenticated, the PAA sends access control information to the EP for granting access to the network. The access control information typically contains the device identifier of the PaC, which is either obtained from the IP headers and MAC headers of the packets exchanged during the authentication process or carried explicitly in the PANA protocol field. The attacker can gain unauthorized access into the network by taking the following steps.

- o An attacker pretends to be a PAA and sends advertisements. The PaC is fooled and starts exchanging packets with the attacker.
- o The attacker modifies the IP source address on the packet, adjusts the UDP/TCP checksum, and forwards the packet to the real PAA. It also does the same on return packets.
- o When the real PaC is successfully authenticated, the attacker gains access to the network, as the packets contained the IP address (and potentially the MAC address also) of the attacker.

If the link is not shared, this threat is absent, as the attacker cannot impersonate the PAA and intercept the packets from the PaC.

If the link is shared, this threat is present. If the layer 2 provides per-packet protection, it is not possible to change the MAC address, and hence this threat may be absent in such cases if EP filters on both the IP and MAC address.

Requirement 5

PANA MUST be able to protect the device identifier against spoofing when it is exchanged between the PaC and PAA.

6.3. PaC Leaving the Network

When the PaC leaves the network, it can inform the PAA before disconnecting from the network so that the resources used by PaC can be accounted properly. The PAA may also choose to revoke the access anytime it deems necessary. The following are possible threats:

T6.3.1: A malicious node can pretend to be a PAA and revoke the access to PaC.

T6.3.2: A malicious node can pretend to be a real PaC and transmit a disconnect message.

T6.3.3: The PaC can leave the network without notifying the PAA or EP (e.g., the Ethernet cable is unplugged, system crash). An attacker can pretend to be the PaC and start using the network.

If the link is not shared, threats T6.3.1 and T6.3.2 are absent. Threat T6.3.3 may still be present. If there is no layer 2 indication, or if the layer 2 indication cannot be relied upon, then threat T6.3.3 is still present on non-shared links.

If the link is shared, all of the above threats are present, as any node on the link can spoof the disconnect message. Even if layer 2 has per-packet authentication, the attacker can pretend to be a PaC (e.g., by spoofing the IP address) and disconnect from the network. Similarly, any node can pretend to be a PAA and revoke the access to the PaC. Therefore, T6.3.1 and T6.3.2 are possible even on links where layer 2 is secured. Threat T6.3.3 can be prevented if layer 2 provides per-packet authentication. The attacker cannot subsume the PaC that left the network without knowing the keys that protect the packet at layer 2.

Requirement 6

PANA MUST be able to protect disconnect and revocation messages.
PANA MUST NOT depend on the PaC sending a disconnect message.

6.4. Service Theft

An attacker can gain unauthorized access into the network by stealing the service from another client. Once the real PaC is successfully authenticated, the EP will have filters in place to prevent unauthorized access into the network. The filters will be based on something that will be carried on every packet. For example, the filter could be based on the IP and MAC addresses, where the packets will be dropped unless the packets coming with certain IP addresses also match the MAC addresses. The following are possible threats:

T6.4.1: An attacker can spoof both the IP and MAC addresses of an authorized client to gain unauthorized access. The attacker can launch this attack easily by just sniffing the wire for IP and MAC addresses. This lets the attacker use the network without any authorization, getting a free service.

T6.4.2: The PaC can leave the network without notifying the PAA or EP (e.g., the Ethernet cable is unplugged, system crash). An attacker can pretend to be the PaC and start using the network.

Service theft allows the possibility of exploiting the weakness in other authentication protocols that use IP address for authentication. It also allows the interception of traffic destined for other nodes by spoofing the IP address.

If the link is not shared, T6.4.1 is absent, as there is only one client on the link, and ingress filtering can prevent the use of the authorized IP and MAC addresses by the attacker on another link. Threat T6.4.2 exists, as the attacker can use the IP or MAC address of the real PaC to gain access to the network.

If the link is shared, both the threats are present. If layer 2 provides per-packet protection using pair-wise keys, both the threats can be prevented.

Requirement 7

PANA MUST securely bind the authenticated session to the device identifier of the client, to prevent service theft. PANA MUST be able to bootstrap a shared secret between the PaC and PAA that can be further used to set up a security association between the PaC and EP to provide cryptographic protection against service theft.

6.5. PAA-EP Communication

After a successful authentication, the PAA needs to communicate the access control information of the PaC to the EP so that the PaC will be allowed to access the network. The information communicated would contain at least the device identifier of the PaC. If strong security is needed, the PAA will communicate a shared secret known only to the PaC and PAA, for setting up a security association between the PaC and EP. The following are possible threats:

T6.5.1: An attacker can eavesdrop to learn the information communicated between the PAA and EP. The attacker can further use this information to spoof the real PaC and also to set up security association for gaining access to the network. This threat is absent if the attacker cannot eavesdrop on the link; e.g., the PAA and EP communicate on a link separate from that of visiting PaCs.

T6.5.2: An attacker can pretend to be a PAA and send false information to an EP to gain access to the network. In the case of stronger security, the attacker has to send its own device identifier and also a shared secret, so that the EP will let the attacker access the network.

If the communication between the PAA and EP is protected, these threats are absent.

Requirement 8

The communication between the PAA and EP MUST be protected against eavesdropping and spoofing attacks.

6.6. Miscellaneous Attacks

T6.6.1: There are various forms of DoS attacks that can be launched on the PAA or AS. A few are mentioned below. As it is hard to defend against some of the DoS attacks, the protocol should be designed carefully to mitigate or prevent such attacks.

- o An attacker can bombard the PAA with lots of authentication requests. If the PAA and AS are not co-located, the PAA may have to allocate resources to store some state about the PaC locally before it receives the response from the back-end AS. This can deplete memory resources on the PAA.
- o With minimal effort, an attacker can force the PAA or AS to make computationally intensive operations with minimal effort, that can deplete the CPU resources of the PAA or AS.

T6.6.2: PaC acquires an IP address by using stateful or stateless mechanisms before PANA authentication begins [PANAREQ]. When the IP addresses are assigned before the client authentication, it opens up the possibility of DoS attacks in which unauthenticated malicious nodes can deplete the IP address space by acquiring multiple IP addresses or deny allocation to others by responding to every duplicate address detection (DAD) query.

Depleting a /64 IPv6 link-local address space or a /8 RFC1918 private address space requires a brute-force attack. Such an attack is part of a DoS class that can equally target the link capacity or the CPU cycles on the target system by bombarding arbitrary packets. Therefore, solely handling the IP address depletion attack is not going to improve the security, as a more general solution is needed to tackle the whole class of brute-force attacks.

The DAD attack can be prevented by deploying secure address resolution that does not depend on the client authentication,

such as [SEND]. The attack may also be prevented if the EP is placed between the PaCs to monitor the ND/ARP activity and to detect DAD attacks (excessive NA/ARP replies). If none of these solutions are applicable to a deployment, the PaCs can send arbitrary packets to each other without going through the EP, which enables a class of attacks that are based on interfering with the PANA messaging (See T6.1.1). Since there will always be a threat in this class (e.g., insecure discovery), it is not going to improve the overall security by addressing DAD.

7. Summary of Requirements

1. PANA MUST not assume that the discovery process is protected.
2. PANA MUST be able to mutually authenticate the PaC and PAA. PANA MUST be able to establish keys between the PaC and PAA to protect the PANA messages.
3. When compound authentication methods are used in PANA, the methods MUST be cryptographically bound.
4. PANA MUST be able to protect itself against replay attacks.
5. PANA MUST be able to protect the device identifier against spoofing when it is exchanged between the PaC and PAA.
6. PANA MUST be able to protect disconnect and revocation messages. PANA MUST NOT depend on whether the PaC sends a disconnect message.
7. PANA MUST securely bind the authenticated session to the device identifier of the client, to prevent service theft. PANA MUST be able to bootstrap a shared secret between the PaC and PAA that can be further used to set up a security association between the PaC and EP to provide cryptographic protection against service theft.
8. The communication between the PAA and EP MUST be protected against eavesdropping and spoofing attacks.

8. Security Considerations

This document discusses various threats with IP based network access authentication protocol. Though this document discusses the threats for shared and unshared links separately, it may be difficult to make such a distinction in practice (e.g., a dial-up link may be a point-to-point IP tunnel). Hence, the link should be assumed to be a shared link for most of the threats in this document.

9. Normative References

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

10. Informative References

[PANAREQ] Yegin, A., Ed., Ohba, Y., Penno, R., Tsirtsis, G., and C. Wang, "Protocol for Carrying Authentication for Network Access (PANA) Requirements and Terminology", Work in Progress, August 2004.

[EAP-KEY] Aboba, B., et al., "EAP keying framework", Work in Progress.

[RAD-EAP] Aboba, B. and P. Calhoun, "RADIUS (Remote Authentication Dial In User Service) Support For Extensible Authentication Protocol (EAP)", RFC 3579, September 2003.

[TUN-EAP] Puthenkulam, J., et al., "The compound authentication binding problem", Work in Progress.

[SEND] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, March 2005.

11. Acknowledgements

The author would like to thank the following people (in no specific order) for providing valuable comments: Alper Yegin, Basavaraj Patil, Pekka Nikander, Bernard Aboba, Francis Dupont, Michael Thomas, Yoshihiro Ohba, Gabriel Montenegro, Tschofenig Hannes, Bill Sommerfeld, N. Asokan, Pete McCan, Derek Atkins, and Thomas Narten.

Author's Address

Mohan Parthasarathy
Nokia
313 Fairchild Drive
Mountain View, CA-94303

EMail: mohanp@sbcglobal.net

Full Copyright Statement

Copyright (C) The Internet Society (2005).

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

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

Intellectual Property

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

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

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

Acknowledgement

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

