PANA Working Group
Internet Draft
Document: draft-ietf-pana-threats-eval-07.txt
Expires: January 2005

M. Parthasarathy
Nokia

August 2004

Protocol for Carrying Authentication and Network Access Threat Analysis and Security Requirements

Status of this Memo

Copyright Notice

Abstract

This document discusses the threats to protocols that are used to carry authentication for network access. The security requirements arising out of these threats will be used as additional input to the PANA WG for designing the IP based network access authentication protocol.

Table of Contents

1.0 Introduction

The PANA (Protocol for Carrying Authentication for Network Access) 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.

2.0 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].

3.0 Terminology and Definitions

Client Access Device

A network element (e.g., notebook computer, PDA, etc.) 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, who is wishing 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 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 PANA authentication agent or part of the back-end infrastructure.

Device Identifier (DI)

The identifier used by the network as a handle to control and police the network access of a client. Depending on the access technology, identifier might contain any of IP address, link-layer address, switch port number, etc. of a device. 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 that is capable of filtering packets sent by the PANA client using the DI information authorized by PANA authentication agent.

Compound methods

Authentication protocol where, sequence of methods are used one after an other or where methods are tunneled inside an another independently established tunnel between the client and server [TUN-EAP].

4.0 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 like PAA and EP. In these environments, one may observe the following.

- 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).

- All the PaCs may be authenticated to the access network at layer 2 (e.g., 3GPP2 CDMA network) and share a security association with layer 2 authentication agent (e.g., BTS). The PaCs still don't trust each other i.e., 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 it becomes the same as the link that is being shared by more than one node.

5.0 Trust Relationships

PANA authentication involves a client (PaC), PANA agent (PAA), 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 as the AS, often needs to use resources provided by PAA that belongs to another administrative domain. PAA authenticates the PaC before providing local network access. The credentials

provided by PaC for authentication may or may not be understood by PAA. If 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 PAA and AS are already covered in [RAD-EAP]. To counter these threats, the communication between PAA and AS are secured 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 setup 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 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 PaC traverses the PaC-PAA path and PAA-AS path. The threats in PAA-AS path are already discussed in [RAD-EAP]. This document discusses the threats in PaC-PAA path.

It is possible that some of the entities like 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 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 PAA and hence PaC and EP do not necessarily share a trust relationship initially. When PaC is successfully authenticated, it may result in key establishment between PaC and PAA, which can be further used to secure the link between PaC and EP. For example, EAP keying framework [EAP-KEY], defines a three party EAP exchange where 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 PaC and EP that can be used to secure the link further. This is further discussed in section 6.4 below.

6.0 Threat Scenarios

The PANA authentication client (PaC) needs to discover the PAA first. This involves either sending solicitations or waiting for advertisements. Once it has discovered the PAA, it will lead to authentication exchange with PAA. Once the access is granted, 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 while leaving the network. Section 6.4 discusses service theft. Section 6.5 discusses the threats in 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 just 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. Following are the 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 because 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, this threat is present. It is difficult to protect the

discovery process, as there is no a priori trust relationship between 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) and prevent this threat.

The advertisement may be used to include other information like supported authentication methods etc., besides 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 less secure authentication methods than what the real PAA supports, thereby fooling the PaC into negotiating a less secure authentication method than what 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 about 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 false authentication success or failure message to the PaC. By sending false failure message, the attacker can prevent the client from accessing the network. By sending 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 as 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 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 the process of mutually authenticating each other, they should be able to establish keys to protect the success or failure indications. It may not be possible to protect the success or failure indication always 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 that are 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 PAA to the real PaC and claim to be PaC to the real PAA. This is a man in the middle (MITM) attack where the PaC is fooled to think that it is communicating with real PAA and the real PAA is fooled to think that it is communicating with real PaC.

If the link is not shared, this threat is absent as ingress filtering can prevent the attacker from acting as 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 man in the middle and launch this attack. An instance of MITM attack, when 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 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 for authenticating 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 as PAA and replay the packets.

If the link is shared, this threat is present. If the packets are encrypted at layer 2 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 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 using the following steps.

- An attacker pretends to be a PAA and sends advertisements. PaC gets fooled and starts exchanging packets with the attacker.
- The attacker modifies the IP source address on the packet, adjusts the UDP/TCP checksum and forwards the packet to the real PAA. It does the same on return packets also.
- 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 as PAA and intercept the packets from 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 both on 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 any time if it deems necessary. Following are the 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, the threats T6.3.1 and T6.3.2 are absent. The threat T6.3.3 may still be present. If there is no layer 2 indication or the layer 2 indication cannot be relied up on, then the 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 the 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. Hence, T6.3.1 and T6.3.2 are possible even on links where layer 2 is secured. The 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, 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 IP and MAC address where the packets will be dropped unless the packets coming with certain IP address match the MAC address also. Following are the possible threats.

T6.4.1: Attacker can spoof both the IP and MAC address of an authorized client to gain unauthorized access. Attacker can launch this attack easily by just sniffing the wire for IP and MAC address. 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 for intercepting 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 authorized IP and MAC address by the attacker on another link. The threat T6.4.2 exists as the attacker can use the IP address 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 which can be further used to setup a security association between 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 EP so that 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, PAA will communicate a shared secret known only to PaC and PAA, for setting up a security association between PaC and EP. Following are the possible threats.

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

T6.5.2: Attacker can pretend to be PAA and send false information to EP for gaining access to the network. The attacker has to send its own device identifier and also a shared secret in the case of stronger security so that EP will let the attacker access the network.

If the communication between 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.

- Attacker can bombard the PAA with lots of authentication requests. If PAA and AS are not collocated, PAA may have to allocate resources to store some state about PaC locally before it receives the response from the backend AS. This can deplete memory resources on PAA.

- The 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 where unauthenticated malicious nodes can deplete the IP address space by acquiring multiple IP addresses, or denying 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 in between the PaCs to monitor the ND/ARP activity and 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.

7.0 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 the PaC sending 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 which can be further used to setup a security association between 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.

8.0 Security Considerations

This document discusses various threats with IP based network access authentication protocol. Though this document discusses the threats separately for shared and unshared links, it may be difficult to make such 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.

9.0 IANA Considerations

This document has no actions for IANA.

10.0 Normative References

[KEYWORDS] S. Bradner, "Key words for use in RFCS to indicate requirement levels", RFC 2119, March 1997.

11.0 Informative References

[PANAREQ] A. Yegin et al., "Protocol for Carrying Authentication for Network Access (PANA) Requirements and Terminology", draft-ietf-pana-requirements-08.txt.

[RADIUS] C. Rigney et. al, "Remote Authentication Dial In User Service", RFC2865, June 2000.

[EAP-KEY] B. Aboba et. al, "EAP keying framework", draft-ietf-eap-keying-00.txt.

[ADDRCONF] S. Thomson et. al, "IPv6 Stateless Address Autoconfiguration", RFC2462, December 1998.

[RAD-EAP] B. Aboba, et. al, "Radius support for Extensible authentication protocol", RFC3579, September 2003.

[TUN-EAP] J. Puthenkulam et. al, "The compound authentication binding problem", draft-puthenkulam-eap-binding-04.txt.

[IPSEC] S. Kent et. al, "Security architecture for the Internet Protocol", RFC 2401, November 1998.

[SEND] J. Arkko et. al, "Secure Neighbor Discovery (SEND)", draft-ietf-send-ndopt-05.txt.

[IEEE-802.11i] Institute of Electrical and Electronics Engineers "Unapproved Draft Supplement to Standard for Telecommunications and Information Exchange Between systems - LAN/MAN Specific Requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Specification for Enhanced Security", IEEE Draft 802.11i (work in progress), 2003.

[IEEE-802.11] Institute of Electrical and Electronics Engineers, "Information Technology - Telecommunications and Information Exchange between Systems - Local and Metropolitan Area Network - Specific Requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications", IEEE Standard 802.11, 1999.

12.0 Acknowledgments

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.

13.0 Revision Log

Changes between 06 and 07

-Updates after IESG review

Changes between 05 and 06

-IANA considerations section added.

Changes between 04 and 05

-Updates after AD review.

Changes between 03 and 04

-Added a new requirement for the disconnect notification.
-Trust relationship section was rewritten.
-Device identifier attack requirements was rewritten.
-Service theft requirement was rewritten.
-Added a new section for PAA-EP threats.

Changes between revision 02 and 03

-Changed Requirement 1 to include text about weak authentication suites.
-Rearranged the order of definitions in terminology section.
-Removed some confusing text with respect to IPsec from the Service theft section.

Changes between revision 01 and 02

-Renamed the section "Assumptions" to "Trust relationships" and added more text to clarify the relationship between PaC and EP.
-Added more text for threats in the path between PAA and AS.
-Merged the "Type of Attacks" section into "Threat Scenarios"
-Removed the requirement on DoS attack.
-Reworded most of the requirements.

Changes between revision 00 and 01

-Removed unused terms from section 3.0.
-Removed identity protection as a threat after feedback from Atlanta IETF55 meeting.
-Renamed the section "Attacks on Normal Data communication" to "Service theft". Removed confidentiality as a requirement from that section.
-Added a new threat "Device Identifier attack".

14.0 Author's Address

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

Email: mohanp@sbcglobal.net

Intellectual Property Statement

Disclaimer of Validity

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.

Copyright Statement

Acknowledgment

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