PANA WG
Internet Draft
Category: Informational
Document: draft-yacine-pana-paa-ep-reqs-00.txt
Expires: December 2003

Yacine El Mghazli
Alcatel

June 2003

PANA
PAA-EP Protocol Requirements

Status of this Memo

Abstract

This document specifies the requirements that the PAA-EP protocol must satisfy in order to meet the needs of PANA when the PAA is separated from EP(s).

Conventions used in this document

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].

Table of Contents

1. Glossary

PANA Protocol for Carying Authentication for Network Access.

PaC (PANA Client):

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

DI (Device Identifier):

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

PAA (PANA Authentication Agent):

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

EP (Enforcement Point):

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

2. Introduction

Client access authentication should be followed by access control to make sure only authenticated and authorized clients can send and receive IP packets via access network. Access control can involve setting access control lists on the enforcement points. Identification of clients which are authorized to access the network is done by the PANA protocol exchange.

PANA does not assume that the PAA is always co-located with the EP(s). Network access enforcement can be provided by one or more nodes on the same IP subnet as the client (e.g., multiple routers), or on another subnet in the access domain (e.g., gateway to the Internet, depending on the network architecture). When the PAA and the EP(s) are separated, there needs to be another transport for client provisioning. This transport is needed to create access

control lists to allow authenticated and authorized clients' traffic through the EPs. PANA Working Group will preferably identify an existing protocol solution that allows the PAA to deliver the authorization information to one or more EPs when the PAA is separated from EPs.

2.1 Scope

The following section 3 discusses the PANA framework assumptions that are being made within the PANA working group. It deals with crucial issues around the authentication process, when the PAA is separated from EP(s).

From this standpoint, section 4 details the requirements that the PAA-EP protocol must satisfy in order to meet the needs of such a framework.

3. PANA framework Assumptions/Issues

3.1 Multiple PAAs

Multiple PAAs may be used for redundancy, load sharing, distributed authentication, or other purposes:

a) Redundancy is the case where one or more PAAs are prepared to take over if an active PAA fails.

b) Load sharing is the case where two or more PAAs are concurrently active and any PaC that can be authenticated by one of the PAAs can also be authenticated by any of the other PAAs.

For both redundancy and load sharing, the PAAs involved are equivalently capable. The only difference between these two cases a) and b) is in terms of how many active PAAs there are.

c) Distributed authentication is the case where two or more PAAs are concurrently active but certain PANA requests using PANA can only be serviced by certain PAAs. The logical separation can be based on:

. Topology: One given PAA is in charge of authenticating a pool of PaCs belonging to the same topological area.

. The ISP: One given PAA is in charge of authenticating the PaCs clients to a given ISP. Then it forwards the PANA requests based on the NAI or other identifier.

. Etc.

Clearly stating the motivation for having multiples PAAs authenticating PaCs and provisoning EPs in an access network has direct consequences on both PaC-PAA and EP-PAA relations.

3.1.1 PAA-PaC relation assumption

According to [PANA] (section "Discovery and Initial Handshake Phase"), "There can be multiple PAAs on the link. The result does not depend on which PAA PaC chooses. By default PaC chooses the PAA that sent the first response."
Then, it is straighforward that the assumption that is being made here is that two or more PAAs are concurrently active and any PaC that can be authenticated by one of the PAAs can also be authenticated by any of the other PAAs. We are clearly in the case where the PaCs load is shared between the multiple PAAs (b).

Do note that discovery issues are raised with allowing muliples PAAs to authenticate the various PaCs. [PANA] solves the problem simply stating that the chosen PAA corresponds to the first response. It is consistent with case b).

3.1.2 PAA-EP relation issue

In a similar manner, it is crucial for identifying the various PAA-EP protocol requirements to clearly identify the context for having multiples PAAs with respect to the EPs provisoning.

One PAA have to communicate with several EPs once a PaC is authenticated is a requirement for the PAA-EP protocol (see section 4.3). In the case where there is a single PAA, the assumption being made is that the PAA will provision all the EPs. However, it remains an issue in case we have multiple PAAs.

When multiple PAAs authenticate the PaC, a given PAA can either:

a) Redundancy:
provisions all the EPs of the underlying access network and each EP has a single active PAA. A back-up PAA is ready to take over if the first one fails.

b) Load sharing:
provisions all the EPs of the underlying access network and each EP can be controlled by another active PAA.

c) Distributed control:
provisions a pool of EPs within a given area. The pool can be identified based on topological criteria for instance.

The choice between these options is motivated by PANA-specific considerations. Typically, these can be:

. Scalability:
How many EPs are managed by the PAA(s) ?

. Symetry:
Does all the EPs need to be configured with the same rules ?

. Etc.

3.2 Inter-PAAs communication

When multiple PAAs are employed, their internal organization is considered an implementation issue that is beyond the scope of PANA. PAAs are wholly responsible for coordinating amongst themselves to provide consistency and synchronization. However, PANA does not define the implementation or protocols used between PAAs, nor does it define how to distribute functionality among PAAs. Nevertheless, PANA will support mechanisms for PAA redundancy or fail over, and it is expected that vendors will provide redundancy or fail over solutions within the PANA framework.

4. PAA-EP Protocol Requirements

4.1 Push model

PAA must be able to "push" the provisioning information down to EPs, without any of the EPs "pulling" it. Since PANA exchange takes place between PaC and PAA, EPs are unlikely to be aware of it.

EP provisioning takes place once the PaC is authenticated and authorized, hence the event triggering the EP configuration takes place at the PAA. Then it's straighforward to initiate the exchange at the PAA.

4.2 Pull model

The PUSH model (PAA-initiated configuration) should be used for the communication between PAA and EP.

On the other hand, the PULL model (EP-initiated configuration) might be supported also for the following purposes:

1. Initial EP registration/Recovery:

   When a EP is newly connected to the network, it needs to register itself to the PAA.

   In a similar manner, when an EP crashes and comes up again, it needs to re-connect its PAA. In general, when a failure is detected, the EP must try to reconnect to the remote PAA or attempt to connect to PAA.

2. (Optional) traffic-driven config (a.k.a. event-notification):

   As stated in [PANA], PaC may also choose to start sending packets before getting authenticated. In that case, the network should detect this and send an unsolicited PANA_start message to PaC. EP is the node that can detect such activity. If EP and PAA are co-located, then an internal mechanism (e.g. API) between the EP module and the PAA module on the same host can prompt PAA to start PANA. In case they are separate, there needs an explicit message to prompt PAA.

   Upon detecting the need to authenticate a client, EP can send a trigger message to the PAA on behalf of the PaC. This can be one of the messages provided by the PAA-EP protocol, or, in the absence of such a facility, PAC_discovery can be used as well. This message MUST carry the device identifier of the PaC. So that,

PAA can send the unsolicited PANA_start message directly to the PaC.

4.3 One-to-many PAA-EP relation

One PAA have to communicate with several Eps once a PaC is authenticated. The PAA-EP protocol must be able to handle this 1:n communication.

4.4 Inactive peer detection

The protocol used between PAA and EP should be able to detect inactive peer within an appropriate time period.

This can be achieved by having both the EP and remote PAA constantly verify their connection to each other via keep-alive messages: a heartbeat in fact.

4.5 Stateful approach

The protocol must allow to maintain some states in the PAA in order for an EP that went down and came back up, or an EP that is being introduced in the network to (re-)synchronize with the PAA.

In general terms, the PAA-EP protocol needs to support the stateful model between the PAA and the EP(s) and some other mechanism for the EP to learn the policies currently in effect on that access network.

4.6 Recovery

TBD.

If there is awareness of a connection state in the application level and we have a keepalive and/or use a reliable transport protocol (similar to routing protocols), all detection and synchronizing of state will come naturally.

4.7 General Security Requirements

The PAA-EP protocol must provide for message authentication, confidentiality, and integrity.

The PAA-EP protocol must define mechanisms to mitigate replay attacks on the control messages.

4.8 Accounting/Feedback from the EPs

The PAA must have an efficient way to to get the accounting information of PaC from EP since the PAA may be a client of the AAA backend infrastructure.

4.9 Re-use of an existing protocol

This work hopefully will not involve any new protocol design, it may involve definition of new AVPs for existing protocols. The PANA working group should try to re-use one of the many protocols around to do this. The following protocols were mentioned for consideration:

- SNMP [SNMP]
- COPS-PR [COPS-PR]
- DIAMETER [DIAMETER]
- RADIUS [RADIUS]
- ForCES [ForCES]

Security Considerations

See section 4.7

References

[STD] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.

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

[PANAREQ] R. Penno, et al., "Protocol for Carrying Authentication for Network Access (PANA) Requirements and Terminology" (draft-ietf-pana-requirements-07.txt).

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

[COPS-PR] K. Chan, D. Durham, S. Gai, S. Herzog, K. McCloghrie, F. Reichmeyer, J. Seligson, A. Smith, R. Yavatkar, "COPS Usage for Policy Provisioning,", RFC 3084, March 2001

[PANA] D. Forsberg, Y. Ohba, B. Pati, H. Tschofenig, A. Yegin, "Protocol for Carrying Authentication for Network Access (PANA)"(draft-ietf-pana-pana-00.txt).

[DIAMETER] P. Calhoun, J. Arkko, E. Guttman, G. Zorn, J. Loughney, "Diameter Base Protocol" (draft-ietf-aaa-diameter-15.txt).

Author's Addresses

Yacine El Mghazli
Alcatel
Route de Nozay
91460 Marcoussis cedex
Phone: +33 (0)1 69 63 41 87
Email: yacine.el_mghazli@alcatel.fr

Full Copyright Statement