Internet-Draft
Expires: July, 2002

Yoshihiro Ohba
Subir Das
Henry Haverinen
Basavaraj Patil
Phil Roberts
Hesham Soliman
Barani Subbiah

January 28, 2002

PANA Usage Scenarios

<draft-ietf-pana-usage-scenarios-00.txt>

Status of This Memo

Abstract

Many commercial networks today require users to provide their authentication information before being allowed access to network resources. Resources could include basic access to the network or could be more specific services in the network or a certain grade of service, etc. Currently, the authentication process depends upon the type of network that a user is attaching to and in most cases it is specific to an access technology. Therefore, a common protocol for performing user authentication (which is independent of access technologies) at the network layer (IP) or above is deemed necessary. This document attempts to capture various such scenarios where a higher layer protocol, such as, PANA (Protocol for carrying

Authentication for Network Access) would be appropriate. The purpose of this draft is to help in understanding the problem space clearly, issues involved for different usage scenarios, and finally to facilitate the discussion for PANA requirements and framework.

Table of Contents

1 Introduction

Many commercial networks today require users to provide their authentication information before being allowed access to network resources. Resources could include basic access to the network or could be more specific services in the network or a certain grade of service (e.g., free vs. paid services), etc. Although authentication process varies from network to network, current best practices are to perform the authentication at the link layer. Since existing solutions are specific to access technologies, network providers need either a new transport mechanism or an extension to existing mechanism to authenticate users each time a new access technology is being standardized. A common protocol designed at the network layer (IP) or above could solve this problem.

This document attempts to capture various scenarios where a higher layer protocol, such as, PANA (Protocol for carrying Authentication for Network Access) would be appropriate. The purpose of this draft is to help in understanding the problem space clearly, issues

involved for different usage scenarios, and finally to facilitate the discussion for PANA requirements and framework.

1.1. Requirements Language

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].

1.2. Acronyms

AAA: Authentication, Authorization and Accounting

AP: Access Point

AR: Access Router

DSL: Digital Subscriber Line

GPRS: General Packet Radio Service

ISP: Internet Service Provider

NAS: Network Access Server

PPP: Point-to-Point Protocol

RADIUS: Remote Authentication Dial In User Service

SA: Security Association

WISP: Wireless ISP

WLAN: Wireless Local Area Network

1.3. Terminology

Following terminologies are defined for this document.

Device

A user's equipment (namely notebook computers, PDAs, etc.) that requires access to a provider's network.

Local Network

The immediate network(s) that is available to the device/user for Connecting. The network that the device is connecting to may be operated by anyone (not necessarily the users home network operator).

PAA (PANA Authentication Agent)

The functional element in the access network that a user device communicates with and provides it with the credentials used for authentication. PAA MAY also have an interface to AAA backend

infrastructures for authenticating the device/user or may use any other authentication mechanism.

PANA (Protocol for carrying Authentication for Network Access)

A protocol which is used for carrying authentication information

between a device (acting on behalf of a user) and a PAA in order for the device to be allowed to access the network.

Credentials

Information that is transferred or presented to establish either a claimed identity or the authorizations of a system entity [RFC2828]. Credentials include things such as, private keys, trusted roots, tickets, or the private part of a Personal Security Environment (PSE) [RFC3157,RFC2510].

Initial Authentication

Authentication performed by PAA when a device needs to be authenticated and authorized for network access.

Re-authentication

Authentication performed after successful initial authentication when a device needs to be authenticated again in order to extend the authorization lifetime for the network access.

LSA (Local Security Association)

A temporary security association between a device and a PAA, which is derived from a credential that is provided by the device on behalf of a user during initial authentication.

2. Current Mechanisms

Today's technologies mostly rely on access specific mechanisms. For example, dial-up networks have relied on PPP's [RFC1661] capability to do user authentication in conjunction with AAA (RADIUS) [RFC2865] as a means to initial authentication. However, PPP may not be appropriate for most new type of networks. Moreover, node configuration through PPP is not always necessary. Today's wireless networks, especially cellular networks, also have relied on specific layer 2 identifiers and layer 3 messaging (NOT IP) to accomplish user authentication. Access control in most current technologies is performed on layer 2 based on device identification combined with layer 2 security. As new type of networks (including wireless LANs) are deployed in hotspot areas the legacy mechanisms may not be appropriate in such scenarios. Different kinds of service models are

also emerging in today's Internet. For example, while basic connectivity may be user-agnostic, enhanced services will be available only to authorized users (since charging is involved for such services).

The following scenarios are described in the next section:

- NAS
- WLAN
- GPRS
- Intra-domain, inter-technology handoff

3. Scenarios Where PANA is Applicable

3.1. NAS

Basic NAS functionality includes authentication of users with an authentication server. It has also hooks for access control. Although PPP offers these functionalities, it assumes few network characteristics:

i) PPP always assumes a link layer disconnect indication.

-- This feature is not available in networks such as Wireless LANs and others. So there has to be some mechanism to detect when the client disconnects. This might mean that there is also a need to re-authenticate client X to appear just after client A left, spoof the MAC and IP addresses of the client A, and basically continue to use client A's authenticated access.

ii) Since PPP works in scenarios where dedicated links exist, it normally assumes a single access router in that link.

-- However, this is not true in multi-access links, such as WLANs. Multiple ARs are required for efficient control and robustness. The multi-AR scenario is discussed in section 3.3.

iii) PPP does address configuration to the client.

-- Address configuration should be decoupled from AAA functionalities since in many cases devices may use other address configuration mechanisms such as DHCP, stateless address autoconfiguration, etc.

In view of the above, we need a flexible NAS type of functionality. There are definitely architectural tradeoffs relating to the relationship between the PAA and the access control enforcement point in terms of their relative location (same box, or different box) as well as how close or far away from the client. Also there may be a need for defining some rules or policies on how this will interwork with L2 authentication and access control mechanisms, such as 802.1X

and cellular systems.

3.1.1. Layer Two Specific Network Access Mechanisms

WISPs are becoming prevalent and they are also starting Internet access services in public areas such as airports and hotel lobbies. The access control technologies they are currently considering would

be (i) web-based access control or (ii) L2 access control such as 802.1X. The web-based access control is performed at a higher layer after obtaining an IP address, whereas the L2 access control is performed before obtaining an IP address. There are pros and cons as to which layer access control is performed. However, a higher layer access control would be suitable for realizing user-agnostic network access services in which a user can access local information such as local-area map and flight information for free of charge, while access to specific external web-sites is subject to charge.

The web-based access control could provide such a service, but there is no standard protocol or method which could help WISPs to support roaming users. While web-based access control can support devices with web browsers, there are number of other applications such as FTP, VoIP require a standard client application or protocol at the client. The client will use this protocol to initiate the network access for any non browser based applications. In addition to charge a user for external web access, a local WISP can also provide charge-based information delivery services such as music download, VoIP, etc. These services are provided by the local network and external connectivity is supported by gateways. It is also envisioned that NAS can be triggered either by the device or PAA, based on provider business model. PANA could provide a standardized way of L2-independent user-agnostic network access with roaming enabled.

3.1.2. Multiple Access Routers

In multi-access environments such as Ethernet and 802.11a/b, it is possible for multiple ARs to exist on the same subnet. This is a fundamental difference from PPP in which a single AR is always associated with a PPP tunnel and the association never changes throughout the lifetime of the PPP connection.

It would be desirable to use multiple ARs in such a way that traffic coming from and going to a specific device always goes through a single AR for ease of access control, but traffic among different devices diverges for the purpose of load balancing and redundancy.

Although many NAS-based networks support multiple ARs for inbound traffic, however, for outbound traffic (originated from user devices connected via PPP tunnels to a single NAS) a single AR is always used. In an 802.11a/b network, it is possible to support multiple ARs if all first hop routers reside behind the layer-2 access points. On the other hand, if WLAN access points act as first hop access routers then it boils down to the same NAS type of model as described earlier.

Thus in order to achieve load balancing and redundancy in a more flexible way we may need a different type of solution and in which case PANA seems to be appropriate. Consider the following scenario (see Fig.1 for IPv4 example):

o Unlike the existing PPP-based NAS model, the device randomly (or by using some other criteria) selects one of the ARs as the

default router and always uses the selected AR as the next hop for the outgoing traffic. The network is also configured in such a way that ICMP Redirect would not occur in the edge subnet to avoid divergence of traffic coming from the user terminal.

- Like the existing PPP-based NAS model, each AR has at least two interfaces, one (Ie) is attached to the edge subnet and the other (Ic) is attached to the core network. The subnet prefix assigned to Interface Ic covers that is assigned to Interface Ie. When an AR receives an address resolution request (ARP REQUEST or Neighbor Solicitation) from R (who is searching for the device on Interface Ic) the AR that is selected by the device replies to the address resolution request on behalf of the device and thereby guarantees that all incoming traffic going to the device passes through the AR. This part is the same as the existing PPP-based NAS scenario.

For example, assume that devices D1 and D3 perform initial authentication with AR1 and that devices D2 and D4 perform initial authentication with AR2. If initial authentication is successful, AR1 and AR2 will open firewall and create proxy ARP (or ND) entries for D1/D3 and D2/D4, respectively. Packets originated from D1 or D3 are forwarded to AR1. Packets originated from D2 or D4 are forwarded to AR2. On the other hand, in terms of the reverse direction, when core router R receives a packet destined for either D1 or D2, it is delivered to the final destination through AR1 in the following way. If R does not yet have an ARP/Neighbor cache of the destination, it performs address resolution since it thinks that the destination is on-link according to the prefix assignment. Only AR1 makes a response to the address resolution request with specifying the MAC address of Ic1. Then, R forwards the packet to AR1 which delivers it to the final destination device. In the same way, when R receives a packet destined for either D3 or D4, it is delivered to the final destination through AR2.

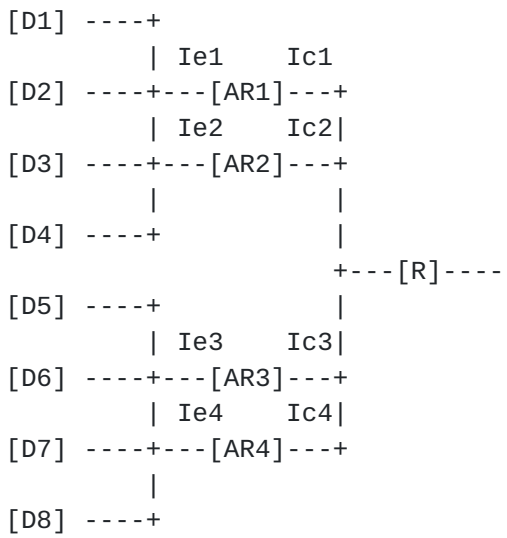

Fig.1: A possible multi-AR environment

In the above case, since both the incoming and outgoing traffic with regard to a specific device are controlled to pass though a single AR, it would be easier to perform authentication via the same AR. However, this would change if the topology is different or if traffic coming from and going to a specific device diverges among different ARs as a result of ICMP Redirect or the use of "Default Router Preferences and More-Specific Routes" [Draves] in IPv6. Detailed architecture tradeoffs will be discussed in framework document.

3.1.3. Multiple Interfaces on a Single Device

PANA would be useful when a host has multiple interfaces of homogeneous or heterogeneous technologies. A typical example is a device with a Bluetooth interface and and an 802.11 interface or with multiple Bluetooth interfaces. There are three possible scenarios: interface switching, multi-homing, and interface sharing.

3.1.3.1. Interface Switching

If multiple interfaces are connectable to the same local network,

only one of those interfaces may be activated at a time for the purpose of battery saving, and perform interface switching when other interface is to be activated, with or without changing IP address. In such an environment, the device may have to perform initial authentication each time interface switching occurs, as well as periodical re-authentication for the activated interface. This would not be desired if initial authentication or re-authentication involves communication with the AAA infrastructure which would

increase AAA signaling traffic in the core network unless the user's credential is stored in the L2 network access point.

The LSA established by using PANA as a result of successful initial authentication with an interface can be used for local re-authentication which would be performed periodically or at every interface switching event. See section "Using LSA for details.

Note that this kind of optimization by PANA would not be possible if (i) multiple interfaces are connected to different local networks each managed by an independent PAA and (ii) re-authentication is not required.

3.1.3.2. Multi-homing

If multiple interfaces are connectable to the same local network, those interfaces may be activated at the same time for the purpose of bandwidth increase and/or load balancing.

In such an environment, the device may have to perform initial authentication multiple times, one for each interface, and periodical re-authentication for each interface. This would not be desired for the same reason described in section "Interface Switching".

The LSA established by using PANA as a result of successful initial authentication with an interface can be can be shared among all the interfaces and used for local re-authentication which would be performed periodically or when an additional interface comes up. See section "Using LSA for Re-authentication" for details.

Note that this kind of optimization by PANA would not be possible if (i) multiple interfaces are connected to different local networks each managed by an independent PAA and (ii) re-authentication is not required.

3.1.3.3. Interface Sharing

There may be cases in which a device (i.e., a gateway device) has one provider interface and multiple private interfaces, where only the provider interface is connected to the ISP's local network and other interfaces are connected to other devices (child devices) under administration of the user (i.e., both the gateway device and child devices are owned by the same user). Traffic coming from the private interfaces would be bridged or routed to the provider interface and vice versa. Such a scenario would be more realistic in IPv6 where a /64 prefix could be allocated to the user, enabling a distinct IP

address within the allocated prefix to be assigned to each of the devices of the same user. In terms of both access control overhead and AAA signaling overhead, it would be desirable to perform access control per prefix rather than per child device. PANA would be very suitable for such a prefix-based access control, especially in DSL or cable modem environment (see Fig.2).

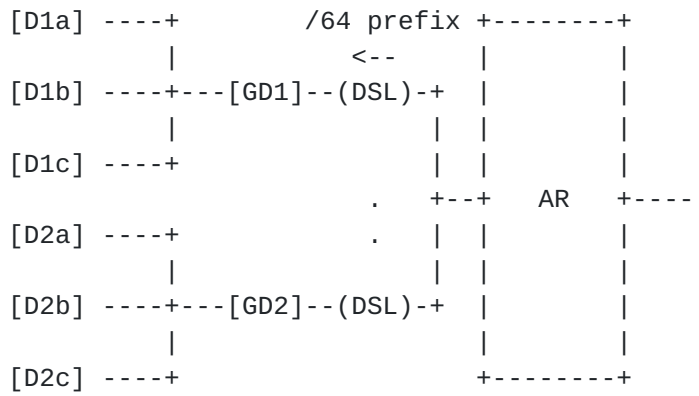

Fig.2: PANA usage in DSL

In this usage scenario, it would be required to prevent other users from using the advertised prefix especially when the AR uses a single Ethernet interface for multiple gateway devices (i.e., multiple gateway devices of different users are on the same shared media, as shown in Fig.2). Details will be discussed in framework document.

Note that if the gateway device and child devices are owned by different users, access control would be performed per-device instead of per-prefix. Then it boils down to one of the models as described earlier in which each device performs PANA with PAA.

3.2. WLAN

WLAN is a one of the many scenarios where PANA could be useful. While 802.1X specification is being deployed by WISPs today, it would not meet the future needs such as inter-technology hand-offs, differentiation between free and paid access and others.

While interworking with 802.1X needs to be taken into consideration, a solution at the higher layer is clearly another option, especially where only legacy 802.11a/b interface cards and APs are available. An alternative is that L2 authentication could be used to get free local access, and PANA could then later be used to access specific paid services or access to global Internet.

PANA and 802.1X-based authentication together may not be required for all situations rather it would be preferable to use them in different places - 802.1X authentication is preferred in public operator networks where there are no free local services (or if all services are charged at a flat rate), and PANA is preferred in hotels and other places with free intranets. However, both mechanisms can coexist in a service provider's network since they belong to two different layers.

3.3. GPRS

In GPRS there are only two methods of authenticating/acquiring an address from networks other than the GPRS access network (corporate or ISP). In one mode PPP is terminated by the GGSN which then authenticates the user (using the RADIUS client) to the "other" network, being the ISP the user connects to or the "home" corporate network depending on the requested APN. In the second mode, PPP is terminated in the MT, ONLY the Challenge Response is piggybacked on PDP context requests. The RADIUS client in the GGSN then relays this info to the RADIUS server wherever that maybe (Based on the APN).

As yet another example, there is no mechanism for exchanging a second level of authentication and addressing information, only the interaction with the cellular systems' authentication and address assignment mechanism is possible (although Mobile IP could be used).

A protocol that decouples the authentication and addressing from PPP (and Mobile IP) could allow an operator to provide remote authentication and authorization of home network services and address assignment within such a domain using one of the available IP layer tunneling mechanisms to deliver packets.

3.4. Intra-domain, Inter-technology Handoff

Although the PANA WG will not directly work on solutions relating to mobility of the device. However, it is noted in the PANA WG charter that the ability to re-authenticate locally with the PAA, can be an important element in allowing efficient handling of mobile devices. This section describe possible situations where an LSA established by using PANA would be useful for mobile nodes.

Different radio technologies will have different characteristics in terms of BW, cost, resilience and speed. Different applications may prefer to use access technologies that are most suited to their needs. It is foreseen that network operators will overlay different access technologies on top of their existing IP backbones to satisfy users' needs, hot spot coverages, etc. For an operator to have control over the access to their networks some mechanism for user authentication and access control is required. Currently these mechanisms are access dependent (e.g.,. HLR in GPRS and 802.1X specific mechanisms). Such access dependence requires an access dependent identity for the user. Hence while connected to a Cellular network, a user is identified by an IMSI, on a WLAN or Bluetooth network a user will have a different identity.

For users to be able to roam seamlessly within an operator's domain between different access technologies they need to avoid re-

authentication all the way back to their home AAA servers each time the move. Context Transfer (CT) may help, however if the identity of a user is different in each access technology, CT will not help and re-authentication will become necessary. Hence an access independent mechanism that uses a standard access independent identity is need to identify the user regardless of which access technology he/she is connected to. The LSA established by using PANA can be used for minimize the re-authentication overhead so that re-authentication is

performed only locally in a similar way that is described in section "Interface Switching". CT and MIP can solve the rest of the problem for seamless mobility.

Another possible scenario in Mobile IPv6 is described in [Soliman].

4. LSA Usages

Establishing an LSA between a device and a PAA is equivalent to having a shared secret between them. The shared secret for the LSA can be derived from credentials of a user. The credentials could be a symmetric cryptographic key (i.e., shared key) or an asymmetric cryptographic key (i.e., public key). A typical example of symmetric cryptographic key is a password stored in the database of the user's home ISP. A typical example of asymmetric cryptographic key is a PKI digital certificate issued by a trusted 3rd party.

4.1. Using LSA for Re-authentication

Once initial authentication is successful, re-authentication would be necessary in the following situations:

- When there is a change in attributes of either the device or PAA. For example, re-authentication would be necessary when the device's IP address changes due to e.g., interface switching or handoff.

- If the underlying access network does not have a capability to detect physical disconnection of devices, periodical re-authentication is necessary for (i) preventing connection hijacking from malicious users which would possibly occur when the device is shutdown or the user leaves the local network with the device without performing explicit log-off from the local network or (ii) improve the accuracy of accounting.

On the other hand, it is desired that periodical re-authentication is performed locally between the device and PAA in order to minimize the signaling traffic.

Once the LSA is established, re-authentication could be performed locally based on the shared secret for the established LSA.

4.2. Using LSA for IKE

The shared key for the LSA or other shared key derived from the LSA

can also be used as an IKE credential with which an IPsec SA between the device and PAA is established via IKE[RFC2409]. This would be useful especially in roaming environment where it is difficult to share the IKE credential a priori between the device and an IPsec entity in the local network.

The established IPsec SA can be used for protecting PANA message exchange, which would be useful for quick re-authentication and/or

secure explicit log-off.

It is also possible to utilize the LSA to establish an IPsec tunnel between a device and an IPsec gateway. The established IPsec tunnel can be used for protecting user data packets in the access network at the IP layer. This would provide access network security without relying on L2 security mechanisms, which is important considering the recent exposure of WEP vulnerabilities.

5. Conclusion

The scenarios described above show a clear need for a common edge protocol that would provide network or service providers a vehicle for user authentication irrespective of what access technology they are using. It is anticipated that a higher layer solution like PANA can support future flexible service models. However, there are several issues that need to be resolved before we deploy this solution. We hope this draft will help the WG to understand different usage scenarios either existing or upcoming and the rationale behind this approach.

6. Acknowledgments

The authors would like to thank James Kempf, David Spence, and the rest of the PANA Working Group for the ideas and support they have given to this document.

7. References

[DIAMETER] P. Calhoun, et al., "Diameter Base Protocol", Work in progress, November 2001.

[Draves] R. Draves, "Default Router Preferences and More-Specific Routes", Work in Progress, May 2001.

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

[RFC1661] W. Simpson, "The Point-to-Point Protocol (PPP)", RFC 1661 (STD 51), July 1994.

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

[RFC2510] C. Adams, et al., "Internet X.509 Public Key Infrastructure Certificate Management Protocols", RFC 2510, March 1999.

[RFC2828] R. Shirey, "Internet Security Glossary", RFC 2828, May 2000.

[RFC2865] C. Rigney, "Remote Authentication Dial In User Service (RADIUS)", RFC 2865, June 2000.

[RFC3157] A. Arsenault, et. al, "Securely Available Credentials - Requirements", RFC 3157, August 2001.

[Soliman] H. Soliman, et al., "Security Association establishment for Mobile IPv6 Route Optimisation using AAA", Internet-Draft, Work in progress, July 2001.

8. Authors' Information

Yoshihiro Ohba
Toshiba America Research, Inc.
P.O. Box 136
Convent Station, NJ 07961-0136
USA
Phone: +1 973 829 5174
Fax: +1 973 829 5601
Email: yohba@tari.toshiba.com

Subir Das
MCC 1D210R, Telcordia Technologies
445 South Street, Morristown, NJ 07960
Phone: +1 973 829 4959
email: subir@research.telcordia.com

Henry Haverinen
Nokia Mobile Phones
P.O. Box 88
FIN-33721 Tampere
Finland
E-mail: henry.haverinen@nokia.com
Phone: +358 50 594 4899
Fax: +358 3 318 3690

Basavaraj Patil
Nokia
6000 Connection Dr.
Irving, TX. 75039
USA
Phone: +1 972-894-6709
Email: Basavaraj.Patil@nokia.com

Phil Roberts
Megisto Corp.
Suite 120
20251 Century Blvd
Germantown MD 20874
USA

Phone: +1 847-202-9314
Email: PRoberts@MEGISTO.com

Hesham Soliman
Ericsson Radio Systems AB
Torshamnsgatan 29,
Kista, Stockholm 16480
Sweden

Phone: +46 8 7578162
Fax: +46 8 4043630
E-mail: Hesham.Soliman@era.ericsson.se

Barani Subbiah
3Com Corporation
5400 Bayfront Plaza
Santa Clara CA 95052
Email: barani_subbiah@3com.com