

Network Working Group
Request for Comments: 2486
Category: Standards Track

B. Aboba
Microsoft
M. Beadles
WorldCom Advanced Networks
January 1999

The Network Access Identifier

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (1999). All Rights Reserved.

1. Abstract

In order to enhance the interoperability of roaming and tunneling services, it is desirable to have a standardized method for identifying users. This document proposes syntax for the Network Access Identifier (NAI), the userID submitted by the client during PPP authentication. It is expected that this will be of interest for support of roaming as well as tunneling. "Roaming capability" may be loosely defined as the ability to use any one of multiple Internet service providers (ISPs), while maintaining a formal, customer-vendor relationship with only one. Examples of where roaming capabilities might be required include ISP "confederations" and ISP-provided corporate network access support.

2. Introduction

Considerable interest has arisen recently in a set of features that fit within the general category of "roaming capability" for dialup Internet users. Interested parties have included:

Regional Internet Service Providers (ISPs) operating within a particular state or province, looking to combine their efforts with those of other regional providers to offer dialup service over a wider area.

National ISPs wishing to combine their operations with those of one or more ISPs in another nation to offer more comprehensive dialup service in a group of countries or on a continent.

Businesses desiring to offer their employees a comprehensive package of dialup services on a global basis. Those services may include Internet access as well as secure access to corporate intranets via a Virtual Private Network (VPN), enabled by tunneling protocols such as PPTP, L2F, L2TP, and IPSEC tunnel mode.

In order to enhance the interoperability of roaming and tunneling services, it is desirable to have a standardized method for identifying users. This document proposes syntax for the Network Access Identifier (NAI). Examples of implementations that use the NAI, and descriptions of its semantics, can be found in [1].

2.1. Terminology

This document frequently uses the following terms:

Network Access Identifier

The Network Access Identifier (NAI) is the userID submitted by the client during PPP authentication. In roaming, the purpose of the NAI is to identify the user as well as to assist in the routing of the authentication request. Please note that the NAI may not necessarily be the same as the user's e-mail address or the userID submitted in an application layer authentication.

Network Access Server

The Network Access Server (NAS) is the device that clients dial in order to get access to the network. In PPTP terminology this is referred to as the PPTP Access Concentrator (PAC), and in L2TP terminology, it is referred to as the L2TP Access Concentrator (LAC).

Roaming Capability

Roaming capability can be loosely defined as the ability to use any one of multiple Internet service providers (ISPs), while maintaining a formal, customer-vendor relationship with only one. Examples of cases where roaming capability might be required include ISP "confederations" and ISP-provided corporate network access support.

Tunneling Service

A tunneling service is any network service enabled by tunneling protocols such as PPTP, L2F, L2TP, and IPSEC tunnel mode. One example of a tunneling service is secure access to corporate intranets via a Virtual Private Network (VPN).

2.2. Requirements language

In this document, the key words "MAY", "MUST", "MUST NOT", "optional", "recommended", "SHOULD", and "SHOULD NOT", are to be interpreted as described in [9].

2.3. Purpose

As described in [1], there are now a number of services implementing dialup roaming, and the number of Internet Service Providers involved in roaming consortia is increasing rapidly.

In order to be able to offer roaming capability, one of the requirements is to be able to identify the user's home authentication server. For use in roaming, this function is accomplished via the Network Access Identifier (NAI) submitted by the user to the NAS in the initial PPP authentication. It is also expected that NASes will use the NAI as part of the process of opening a new tunnel, in order to determine the tunnel endpoint.

2.4. Notes for Implementors

As proposed in this document, the Network Access Identifier is of the form user@realm. Please note that while the user portion of the NAI conforms to the BNF described in [5], the BNF of the realm portion allows the realm to begin with a digit, which is not permitted by the BNF described in [4]. This change was made to reflect current practice; although not permitted by the BNF described in [4], FQDNs such as 3com.com are commonly used, and accepted by current software.

Please note that NAS vendors may need to modify their devices so as to support the NAI as described in this document. Devices handling NAIs MUST support an NAI length of at least 72 octets.

3. Formal definition of the NAI

The grammar for the NAI is given below, described in ABNF as documented in [7]. The grammar for the username is taken from [5], and the grammar for the realm is an updated version of [4].

```
nai           = username / ( username "@" realm )
```

username = dot-string

realm = realm "." label

label = let-dig * (ldh-str)

ldh-str = *(Alpha / Digit / "-") let-dig

dot-string = string / (dot-string "." string)

string = char / (string char)

char = c / ("\" x)

let-dig = Alpha / Digit

Alpha = %x41-5A / %x61-7A ; A-Z / a-z

Digit = %x30-39 ; 0-9

c = < any one of the 128 ASCII characters, but
not any special or SP >

x = %x00-7F
; all 127 ASCII characters, no exception

SP = %x20 ; Space character

special = "<" / ">" / "(" / ")" / "[" / "]" / "\" / "."
/ "," / ";" / ":" / "@" / %x22 / Ctl

Ctl = %x00-1F / %x7F
; the control characters (ASCII codes 0 through 31
; inclusive and 127)

Examples of valid Network Access Identifiers include:

fred@3com.com
fred@foo-9.com
fred_smith@big-co.com
fred=?#\$%*+~^smith@bigco.com
fred@bigco.com
nancy@eng.bigu.edu
eng!nancy@bigu.edu
eng%nancy@bigu.edu

Examples of invalid Network Access Identifiers include:

```
fred@foo
fred@foo_9.com
@howard.edu
fred@bigco.com@smallco.com
eng:nancy@bigu.edu
eng;nancy@bigu.edu
<nancy>@bigu.edu
```

4. References

- [1] Aboba, B., Lu J., Alsop J., Ding J. and W. Wang, "Review of Roaming Implementations", RFC 2194, September 1997.
- [2] Rigney C., Rubens A., Simpson W. and S. Willens, "Remote Authentication Dial In User Service (RADIUS)", RFC 2138, April 1997.
- [3] Rigney C., "RADIUS Accounting", RFC 2139, April 1997.
- [4] Mockapetris, P., "Domain Names - Implementation and Specification", STD 13, RFC 1035, November 1987.
- [5] Postel, J., "Simple Mail Transfer Protocol", STD 10, RFC 821, August 1982.
- [6] Gulbrandsen A. and P. Vixie, "A DNS RR for specifying the location of services (DNS SRV)", RFC 2052, October 1996.
- [7] Crocker, D. and P. Overrell, "Augmented BNF for Syntax Specifications: ABNF", RFC 2234, November 1997.
- [8] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", RFC 2401, November 1998.
- [9] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

5. Security Considerations

Since an NAI reveals the home affiliation of a user, it may assist an attacker in further probing the username space. Typically this problem is of most concern in protocols which transmit the user name in clear-text across the Internet, such as in RADIUS, described in [2] and [3]. In order to prevent snooping of the user name, protocols may use confidentiality services provided by IPSEC, described in [8].

6. IANA Considerations

This document defines a new namespace that will need to be administered, namely the NAI realm namespace. In order to avoid creating any new administrative procedures, administration of the NAI realm namespace will piggyback on the administration of the DNS namespace.

NAI realm names are required to be unique and the rights to use a given NAI realm for roaming purposes are obtained coincident with acquiring the rights to use a particular fully qualified domain name (FQDN). Those wishing to use an NAI realm name should first acquire the rights to use the corresponding FQDN. Using an NAI realm without ownership of the corresponding FQDN creates the possibility of conflict and therefore is to be discouraged.

Note that the use of an FQDN as the realm name does not imply use of the DNS for location of the authentication server or for authentication routing. Since to date roaming has been implemented on a relatively small scale, existing implementations typically handle location of authentication servers within a domain and perform authentication routing based on local knowledge expressed in proxy configuration files. The implementations described in [1] have not found a need for use of DNS for location of the authentication server within a domain, although this can be accomplished via use of the DNS SRV record, described in [6]. Similarly, existing implementations have not found a need for dynamic routing protocols, or propagation of global routing information. Note also that there is no requirement that the NAI represent a valid email address.

7. Acknowledgements

Thanks to Glen Zorn of Microsoft for many useful discussions of this problem space.

8. Authors' Addresses

Bernard Aboba
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052

Phone: 425-936-6605
EMail: bernarda@microsoft.com

Mark A. Beadles
WorldCom Advanced Networks
5000 Britton Rd.
Hilliard, OH 43026

Phone: 614-723-1941
EMail: mbeadles@wcom.net

9. Full Copyright Statement

Copyright (C) The Internet Society (1999). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS 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.

