

Anonymous Simple Authentication and Security Layer (SASL) Mechanism

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 (2006).

Abstract

On the Internet, it is common practice to permit anonymous access to various services. Traditionally, this has been done with a plain-text password mechanism using "anonymous" as the user name and using optional trace information, such as an email address, as the password. As plain-text login commands are not permitted in new IETF protocols, a new way to provide anonymous login is needed within the context of the Simple Authentication and Security Layer (SASL) framework.

1. Introduction

This document defines an anonymous mechanism for the Simple Authentication and Security Layer ([SASL]) framework. The name associated with this mechanism is "ANONYMOUS".

Unlike many other SASL mechanisms, whose purpose is to authenticate and identify the user to a server, the purpose of this SASL mechanism is to allow the user to gain access to services or resources without requiring the user to establish or otherwise disclose their identity to the server. That is, this mechanism provides an anonymous login method.

This mechanism does not provide a security layer.

This document replaces RFC 2245. Changes since RFC 2245 are detailed in Appendix A.

2. The Anonymous Mechanism

The mechanism consists of a single message from the client to the server. The client may include in this message trace information in the form of a string of [UTF-8]-encoded [Unicode] characters prepared in accordance with [StringPrep] and the "trace" stringprep profile defined in Section 3 of this document. The trace information, which has no semantical value, should take one of two forms: an Internet email address, or an opaque string that does not contain the '@' (U+0040) character and that can be interpreted by the system administrator of the client's domain. For privacy reasons, an Internet email address or other information identifying the user should only be used with permission from the user.

A server that permits anonymous access will announce support for the ANONYMOUS mechanism and allow anyone to log in using that mechanism, usually with restricted access.

A formal grammar for the client message using Augmented BNF [ABNF] is provided below as a tool for understanding this technical specification.

```
message      = [ email / token ]  
              ;; to be prepared in accordance with Section 3  
  
UTF1         = %x00-3F / %x41-7F ;; less '@' (U+0040)  
UTF2         = %xC2-DF UTF0  
UTF3         = %xE0 %xA0-BF UTF0 / %xE1-EC 2(UTF0) /  
              %xED %x80-9F UTF0 / %xEE-EF 2(UTF0)  
UTF4         = %xF0 %x90-BF 2(UTF0) / %xF1-F3 3(UTF0) /  
              %xF4 %x80-8F 2(UTF0)  
UTF0         = %x80-BF  
  
TCHAR       = UTF1 / UTF2 / UTF3 / UTF4  
              ;; any UTF-8 encoded Unicode character  
              ;; except '@' (U+0040)  
  
email        = addr-spec  
              ;; as defined in [IMAIL]  
  
token        = 1*255TCHAR
```

Note to implementors:

The <token> production is restricted to 255 UTF-8-encoded Unicode characters. As the encoding of a characters uses a sequence of 1 to 4 octets, a token may be as long as 1020 octets.

3. The "trace" Profile of "Stringprep"

This section defines the "trace" profile of [StringPrep]. This profile is designed for use with the SASL ANONYMOUS Mechanism. Specifically, the client is to prepare the <message> production in accordance with this profile.

The character repertoire of this profile is Unicode 3.2 [Unicode].

No mapping is required by this profile.

No Unicode normalization is required by this profile.

The list of unassigned code points for this profile is that provided in Appendix A of [StringPrep]. Unassigned code points are not prohibited.

Characters from the following tables of [StringPrep] are prohibited:

- C.2.1 (ASCII control characters)
- C.2.2 (Non-ASCII control characters)
- C.3 (Private use characters)
- C.4 (Non-character code points)
- C.5 (Surrogate codes)
- C.6 (Inappropriate for plain text)
- C.8 (Change display properties are deprecated)
- C.9 (Tagging characters)

No additional characters are prohibited.

This profile requires bidirectional character checking per Section 6 of [StringPrep].

4. Example

Here is a sample ANONYMOUS login between an IMAP client and server. In this example, "C:" and "S:" indicate lines sent by the client and server, respectively. If such lines are wrapped without a new "C:" or "S:" label, then the wrapping is for editorial clarity and is not part of the command.

Note that this example uses the IMAP profile [IMAP4] of SASL. The base64 encoding of challenges and responses as well as the "+ " preceding the responses are part of the IMAP4 profile, not part of SASL itself. Additionally, protocols with SASL profiles permitting an initial client response will be able to avoid the extra round trip below (the server response with an empty "+ ").

In this example, the trace information is "sirhc".

```
S: * OK IMAP4 server ready
C: A001 CAPABILITY
S: * CAPABILITY IMAP4 IMAP4rev1 AUTH=DIGEST-MD5 AUTH=ANONYMOUS
S: A001 OK done
C: A002 AUTHENTICATE ANONYMOUS
S: +
C: c2lyaGM=
S: A003 OK Welcome, trace information has been logged.
```

5. Security Considerations

The ANONYMOUS mechanism grants access to services and/or resources by anyone. For this reason, it should be disabled by default so that the administrator can make an explicit decision to enable it.

If the anonymous user has any write privileges, a denial-of-service attack is possible by filling up all available space. This can be prevented by disabling all write access by anonymous users.

If anonymous users have read and write access to the same area, the server can be used as a communication mechanism to exchange information anonymously. Servers that accept anonymous submissions should implement the common "drop box" model, which forbids anonymous read access to the area where anonymous submissions are accepted.

If the anonymous user can run many expensive operations (e.g., an IMAP SEARCH BODY command), this could enable a denial-of-service attack. Servers are encouraged to reduce the priority of anonymous users or limit their resource usage.

While servers may impose a limit on the number of anonymous users, note that such limits enable denial-of-service attacks and should be used with caution.

The trace information is not authenticated, so it can be falsified. This can be used as an attempt to get someone else in trouble for access to questionable information. Administrators investigating abuse need to realize that this trace information may be falsified.

A client that uses the user's correct email address as trace information without explicit permission may violate that user's privacy. Anyone who accesses an anonymous archive on a sensitive subject (e.g., sexual abuse) likely has strong privacy needs. Clients should not send the email address without the explicit permission of the user and should offer the option of supplying no trace information, thus only exposing the source IP address and time.

Anonymous proxy servers could enhance this privacy but would have to consider the resulting potential denial-of-service attacks.

Anonymous connections are susceptible to man-in-the-middle attacks that view or alter the data transferred. Clients and servers are encouraged to support external data security services.

Protocols that fail to require an explicit anonymous login are more susceptible to break-ins given certain common implementation techniques. Specifically, Unix servers that offer user login may initially start up as root and switch to the appropriate user id after an explicit login command. Normally, such servers refuse all data access commands prior to explicit login and may enter a restricted security environment (e.g., the Unix `chroot(2)` function) for anonymous users. If anonymous access is not explicitly requested, the entire data access machinery is exposed to external security attacks without the chance for explicit protective measures. Protocols that offer restricted data access should not allow anonymous data access without an explicit login step.

General [SASL] security considerations apply to this mechanism.

[StringPrep] security considerations and [Unicode] security considerations discussed in [StringPrep] apply to this mechanism. [UTF-8] security considerations also apply.

6. IANA Considerations

The SASL Mechanism registry [IANA-SASL] entry for the ANONYMOUS mechanism has been updated by the IANA to reflect that this document now provides its technical specification.

To: iana@iana.org

Subject: Updated Registration of SASL mechanism ANONYMOUS

SASL mechanism name: ANONYMOUS

Security considerations: See RFC 4505.

Published specification (optional, recommended): RFC 4505

Person & email address to contact for further information:

Kurt Zeilenga <Kurt@OpenLDAP.org>

Chris Newman <Chris.Newman@sun.com>

Intended usage: COMMON

Author/Change controller: IESG <iesg@ietf.org>

Note: Updates existing entry for ANONYMOUS

The [StringPrep] profile "trace", first defined in this RFC, has been registered:

To: iana@iana.org
Subject: Initial Registration of Stringprep "trace" profile

Stringprep profile: trace
Published specification: RFC 4505
Person & email address to contact for further information:
Kurt Zeilenga <kurt@openldap.org>

7. Acknowledgement

This document is a revision of RFC 2245 by Chris Newman. Portions of the grammar defined in Section 1 were borrowed from RFC 3629 by Francois Yergeau.

This document is a product of the IETF SASL WG.

8. Normative References

- [ABNF] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", RFC 4234, October 2005.
- [IMAIL] Resnick, P., "Internet Message Format", RFC 2822, April 2001.
- [SASL] Melnikov, A., Ed. and K. Zeilenga, Ed., "Simple Authentication and Security Layer (SASL)", RFC 4422, June 2006.
- [StringPrep] Hoffman, P. and M. Blanchet, "Preparation of Internationalized Strings ('stringprep')", RFC 3454, December 2002.
- [Unicode] The Unicode Consortium, "The Unicode Standard, Version 3.2.0" is defined by "The Unicode Standard, Version 3.0" (Reading, MA, Addison-Wesley, 2000. ISBN 0-201-61633-5), as amended by the "Unicode Standard Annex #27: Unicode 3.1" (<http://www.unicode.org/reports/tr27/>) and by the "Unicode Standard Annex #28: Unicode 3.2" (<http://www.unicode.org/reports/tr28/>).
- [UTF-8] Yergeau, F., "UTF-8, a transformation format of ISO 10646", RFC 3629 (also STD 63), November 2003.

9. Informative References

- [IMAP4] Crispin, M., "INTERNET MESSAGE ACCESS PROTOCOL - VERSION 4rev1", RFC 3501, March 2003.
- [IANA-SASL] IANA, "SIMPLE AUTHENTICATION AND SECURITY LAYER (SASL) MECHANISMS", <<http://www.iana.org/assignments/sasl-mechanisms>>.

Appendix A. Changes since RFC 2245

This appendix is non-normative.

RFC 2245 allows the client to include optional trace information in the form of a human readable string. RFC 2245 restricted this string to US-ASCII. As the Internet is international, this document uses a string restricted to UTF-8 encoded Unicode characters. A "stringprep" profile is defined to precisely define which Unicode characters are allowed in this string. While the string remains restricted to 255 characters, the encoded length of each character may now range from 1 to 4 octets.

Additionally, a number of editorial changes were made.

Editor's Address

Kurt D. Zeilenga
OpenLDAP Foundation

EMail: Kurt@OpenLDAP.org

Full Copyright Statement

Copyright (C) The Internet Society (2006).

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 provided by the IETF Administrative Support Activity (IASA).

