

PRECIS
Internet-Draft
Obsoletes: [4013](#) (if approved)
Intended status: Standards Track
Expires: February 5, 2014

P. Saint-Andre
Cisco Systems, Inc.
A. Melnikov
Isode Ltd
August 4, 2013

**Preparation and Comparison of Internationalized Strings Representing
Usernames and Passwords
draft-ietf-precis-saslprepbis-04**

Abstract

This document describes how to handle Unicode strings representing usernames and passwords. This profile is intended to be used by protocols that exchange or otherwise make use of usernames and passwords. This document obsoletes [RFC 4013](#).

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on February 5, 2014.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as

described in the Simplified BSD License.

Table of Contents

- [1. Introduction](#) [3](#)
- [2. What the Username and Password Profiles Provide](#) [3](#)
- [3. Terminology](#) [4](#)
- [4. Usernames](#) [4](#)
 - [4.1. Definition](#) [4](#)
 - [4.2. Preparation](#) [5](#)
- [5. Passwords](#) [6](#)
 - [5.1. Definition](#) [7](#)
 - [5.2. Preparation](#) [7](#)
- [6. Migration](#) [8](#)
 - [6.1. Usernames](#) [8](#)
 - [6.2. Passwords](#) [9](#)
- [7. Security Considerations](#) [10](#)
 - [7.1. Password/Passphrase Strength](#) [10](#)
 - [7.2. Identifier Comparison](#) [10](#)
 - [7.3. Reuse of PRECIS](#) [10](#)
 - [7.4. Reuse of Unicode](#) [11](#)
- [8. IANA Considerations](#) [11](#)
 - [8.1. Use of IdentifierClass](#) [11](#)
 - [8.2. Use of FreeformClass](#) [11](#)
- [9. References](#) [12](#)
 - [9.1. Normative References](#) [12](#)
 - [9.2. Informative References](#) [12](#)
- [Appendix A. Differences from \[RFC 4013\]\(#\)](#) [13](#)
- [Appendix B. Acknowledgements](#) [14](#)
- [Authors' Addresses](#) [14](#)

1. Introduction

Usernames and passwords are widely used for authentication and authorization on the Internet, either directly when provided in plaintext (as in the SASL PLAIN mechanism [[RFC4616](#)] or the HTTP Basic scheme [[RFC2617](#)]) or indirectly when provided as the input to a cryptographic algorithm such as a hash function (as in the SASL SCRAM mechanism [[RFC5802](#)] or the HTTP Digest scheme [[RFC2617](#)]). To increase the likelihood that the input and comparison of usernames and passwords will work in ways that make sense for typical users throughout the world, this document defines rules for preparing and comparing internationalized strings that represent usernames and passwords.

The methods specified in this document define a PRECIS profile as explained in the PRECIS framework specification [[I-D.ietf-precis-framework](#)]. This document assumes that all strings are comprised of characters from the Unicode character set [[UNICODE](#)], with special attention to characters outside the ASCII range [[RFC20](#)]. The methods defined here might be applicable wherever usernames or passwords are used. However, the methods are not intended for use in preparing strings that are not usernames (e.g., email addresses and LDAP distinguished names), nor in cases where identifiers or secrets are not strings (e.g., keys and certificates) or require specialized handling.

This document obsoletes [RFC 4013](#) (the "SASLprep" profile of stringprep [[RFC3454](#)]) but can be used by technologies other than the Simple Authentication and Security Layer (SASL) [[RFC4422](#)], such as HTTP authentication [[RFC2617](#)].

2. What the Username and Password Profiles Provide

Profiles of the PRECIS framework enable software to handle Unicode characters outside the ASCII range in an automated way, so that such characters are treated carefully and consistently in application protocols. In large measure, these profiles are designed to protect application developers from the potentially negative consequences of supporting the full range of Unicode characters. For instance, in almost all application protocols it would be dangerous to treat the Unicode character SUPERSCRIPT ONE (U+0089) as equivalent to DIGIT ONE (U+0031), since that would result in false positives during comparison, authentication, and authorization (e.g., an attacker could easily spoof an account "user1@example.com").

Whereas a naive use of Unicode would make such attacks trivially easy, the Username PRECIS profile defined in this document generally

protects applications from inadvertently causing such problems. (Similar considerations apply to passwords, although here it is desirable to support a wider range of characters so as to maximize entropy during authentication.)

3. Terminology

Many important terms used in this document are defined in [\[I-D.ietf-precis-framework\]](#), [\[RFC5890\]](#), [\[RFC6365\]](#), and [\[UNICODE\]](#). The term "non-ASCII space" refers to any Unicode code point having a general category of "Zs", with the exception of U+0020 (here called "ASCII space").

As used here, the term "password" is not literally limited to a word; i.e., a password could be a passphrase consisting of more than one word, perhaps separated by spaces or other such characters.

Some SASL mechanisms (e.g., CRAM-MD5, DIGEST-MD5, and SCRAM) specify that the authentication identity used in the context of such mechanisms is a "simple user name" (see [Section 2 of \[RFC4422\]](#) as well as [\[RFC4013\]](#)). Various application technologies also assume that the identity of a user or account takes the form of a username (e.g., authentication for the HyperText Transfer Protocol [\[RFC2617\]](#)), whether or not they use SASL. Note well that the exact form of a username in any particular SASL mechanism or application technology is a matter for implementation and deployment, and that a username does not necessarily map to any particular application identifier (such as the localpart of an email address).

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [\[RFC2119\]](#).

4. Usernames

4.1. Definition

This document specifies that a username is a string of Unicode code points [\[UNICODE\]](#), encoded using UTF-8 [\[RFC3629\]](#), and structured either as an ordered sequence of "userparts" (where the complete username can consist of a single userpart or a space-separated sequence of userparts) or as a userpart@domainpart (where the domainpart is an IP literal, an IPv4 address, or a fully-qualified domain name).

The syntax for a username is defined as follows using the Augmented Backus-Naur Form (ABNF) [[RFC5234](#)].

```
username = userpart [1*(1*SP userpart)]
          / userpart '@' domainpart
userpart  = 1*(idpoint)
          ;
          ; an "idpoint" is a UTF-8 encoded Unicode code point
          ; that conforms to the PRECIS "IdentifierClass"
          ;
domainpart = IP-literal / IPv4address / ifqdn
          ;
          ; the "IPv4address" and "IP-literal" rules are
          ; defined in RFC 3986, and the first-match-wins
          ; (a.k.a. "greedy") algorithm described in RFC 3986
          ; applies
          ;
          ; reuse of the IP-literal rule from RFC 3986 implies
          ; that IPv6 addresses are enclosed in square brackets
          ; (i.e., beginning with '[' and ending with ']')
          ;
ifqdn     = 1*1023(domainpoint)
          ;
          ; a "domainpoint" is a UTF-8 encoded Unicode code
          ; point that conforms to RFC 5890
          ;
```

All code points and blocks not explicitly allowed in the PRECIS IdentifierClass are disallowed; this includes private use characters, surrogate code points, and the other code points and blocks that were defined as "Prohibited Output" in [[RFC4013](#)]. In addition, common constructions such as "user@example.com" are allowed as usernames under this specification, as they were under [[RFC4013](#)].

[4.2.](#) Preparation

A username MUST NOT be zero bytes in length. This rule is to be enforced after any normalization and mapping of code points.

Each userpart of a username MUST conform to the definition of the PRECIS IdentifierClass provided in [[I-D.ietf-precis-framework](#)], where the width mapping, additional mapping, case mapping, normalization, and directionality rules are as follows.

1. Fullwidth and halfwidth characters MUST be mapped to their decomposition equivalents.

2. So-called additional mappings MAY be applied, such as those defined in [[I-D.ietf-precis-mappings](#)].
3. Uppercase and titlecase characters SHOULD be mapped to their lowercase equivalents (not doing so can lead to false positives during authentication and authorization, as described in [[RFC6943](#)]).
4. Unicode Normalization Form C (NFC) MUST be applied to all characters.

With regard to directionality, the "Bidi Rule" provided in [[RFC5893](#)] applies.

SASL mechanisms that directly re-use this profile MUST specify whether and when case mapping is to be applied to authentication identifiers. SASL mechanisms SHOULD delay any case mapping to the last possible moment, such as when doing a lookup by username, username comparisons, or generating a cryptographic salt from a username. In keeping with [RFC4422](#), SASL mechanisms are not to apply this or any other profile to authorization identifiers.

Application protocols that use SASL (such as IMAP [[RFC4616](#)] and XMPP [[RFC6120](#)]) and that directly re-use this profile MUST specify whether case mapping is to be applied to authorization identifiers. Such "SASL application protocols" SHOULD delay any case mapping of authorization identifiers to the last possible moment, which happens to necessarily be on the server side. In keeping with [RFC4422](#), SASL application protocols are not to apply this or any other profile to authentication identifiers.

Application protocols that do not use SASL (such as HTTP authentication with the Basic and Digest schemes [[RFC2617](#)]) MUST specify whether and when case mapping is to be applied to authentication identifiers and authorization identifiers. Such application protocols SHOULD delay any case mapping to the last possible moment, such as when doing a lookup by username, username comparisons, or generating a cryptographic salt from a username.

In protocols that provide usernames as input to a cryptographic algorithm such as a hash function, the client will need to perform proper preparation of the username before applying the algorithm, since the username is not available to the server in plaintext form.

5. Passwords

5.1. Definition

This document specifies that a password is a string of Unicode code points [[UNICODE](#)], encoded using UTF-8 [[RFC3629](#)], and conformant to the PRECIS FreeformClass.

The syntax for a password is defined as follows using the Augmented Backus-Naur Form (ABNF) [[RFC5234](#)].

```
password      = 1*(freepoint)
               ;
               ; a "freepoint" is a UTF-8 encoded
               ; Unicode code point that conforms to
               ; the PRECIS "FreeformClass"
               ;
```

All code points and blocks not explicitly allowed in the PRECIS FreeformClass are disallowed; this includes private use characters, surrogate code points, and the other code points and blocks defined as "Prohibited Output" in [Section 2.3 of RFC 4013](#).

5.2. Preparation

A password MUST NOT be zero bytes in length. This rule is to be enforced after any normalization and mapping of code points.

A password MUST conform to the definition of the PRECIS FreeformClass provided in [[I-D.ietf-precis-framework](#)], where the width mapping, additional mapping, case mapping, normalization, and directionality rules are as described below.

1. Fullwidth and halfwidth characters MUST NOT be mapped to their decomposition equivalents.
2. Any instances of non-ASCII space MUST be mapped to ASCII space (U+0020).
3. So-called additional mappings MAY be applied, such as those defined in [[I-D.ietf-precis-mappings](#)].
4. Uppercase and titlecase characters MUST NOT be mapped to their lowercase equivalents.
5. Unicode Normalization Form C (NFC) MUST be applied to all characters.

With regard to directionality, the "Bidi Rule" (defined in [[RFC5893](#)]) and similar rules are unnecessary and inapplicable to passwords, since they can reduce the range of characters that are allowed in a string and therefore reduce the amount of entropy that is possible in a password. Furthermore, such rules are intended to minimize the possibility that the same string will be displayed differently on a

system set for right-to-left display and a system set for left-to-right display; however, passwords are typically not displayed at all and are rarely meant to be interoperable across different systems in the way that non-secret strings like domain names and usernames are.

In protocols that provide passwords as input to a cryptographic algorithm such as a hash function, the client will need to perform proper preparation of the password before applying the algorithm, since the password is not available to the server in plaintext form.

6. Migration

The rules defined in this specification differ slightly from those defined by the SASLprep specification [[RFC4013](#)]. The following sections describe these differences, along with their implications for migration, in more detail.

6.1. Usernames

Deployments that currently use SASLprep for handling usernames might need to scrub existing data when migrating to use of the rules defined in this specification. In particular:

- o SASLprep specified the use of Unicode Normalization Form KC (NFKC), whereas this usage of the PRECIS IdentifierClass employs Unicode Normalization Form C (NFC). In practice this change is unlikely to cause significant problems, because NFKC provides methods for mapping Unicode code points with compatibility equivalents to those equivalents, whereas the PRECIS IdentifierClass entirely disallows Unicode code points with compatibility equivalents (i.e., during comparison NFKC is more "aggressive" about finding matches than is NFC). A few examples might suffice to indicate the nature of the problem: (1) U+017F LATIN SMALL LETTER LONG S is compatibility equivalent to U+0073 LATIN SMALL LETTER S (2) U+2163 ROMAN NUMERAL FOUR is compatibility equivalent to U+0049 LATIN CAPITAL LETTER I and U+0056 LATIN CAPITAL LETTER V (3) U+FB01 LATIN SMALL LIGATURE FI is compatibility equivalent to U+0066 LATIN SMALL LETTER F and U+0069 LATIN SMALL LETTER I. Under SASLprep, the use of NFKC also handled the mapping of fullwidth and halfwidth code points to their decomposition equivalents (see [[I-D.ietf-precis-mappings](#)]). Although it is expected that code points with compatibility equivalents are rare in existing usernames, for migration purposes deployments might want to search their database of usernames for Unicode code points with compatibility equivalents and map those code points to their compatibility equivalents.

- o SASLprep mapped non-ASCII spaces to ASCII space (U+0020), whereas the PRECIS IdentifierClass entirely disallows non-ASCII spaces. The non-ASCII space characters are U+00A0 NO-BREAK SPACE, U+1680 OGHAM SPACE MARK, U+180E MONGOLIAN VOWEL SEPARATOR, U+2000 EN QUAD through U+200A HAIR SPACE, U+202F NARROW NO-BREAK SPACE, U+205F MEDIUM MATHEMATICAL SPACE, and U+3000 IDEOGRAPHIC SPACE. For migration purposes, deployments might want to convert non-ASCII space characters to ASCII space in usernames.
- o SASLprep mapped the "characters commonly mapped to nothing" from [Appendix B.1 of \[RFC3454\]](#) to nothing, whereas the PRECIS IdentifierClass entirely disallows most of these characters, which correspond to the code points from the "M" category defined under Section 6.13 of [\[I-D.ietf-precis-framework\]](#) (with the exception of U+1806 MONGOLIAN TODO SOFT HYPHEN, which was "commonly mapped to nothing" in Unicode 3.2 but at the time of this writing does not have a derived property of Default_Ignorable_Code_Point in Unicode 6.2). For migration purposes, deployments might want to remove code points contained in the PRECIS "M" category from usernames.
- o SASLprep allowed uppercase and titlecase characters, whereas this usage of the PRECIS IdentifierClass maps uppercase and titlecase characters to their lowercase equivalents. For migration purposes, deployments can either convert uppercase and titlecase characters to their lowercase equivalents in usernames (thus losing the case information) or preserve uppercase and titlecase characters and ignore the case difference when comparing usernames.

[6.2.](#) Passwords

Depending on local service policy, migration from [RFC 4013](#) to this specification might not involve any scrubbing of data (since passwords might not be stored in the clear anyway); however, service providers need to be aware of possible issues that might arise during migration. In particular:

- o SASLprep specified the use of Unicode Normalization Form KC (NFKC), whereas this usage of the PRECIS FreeformClass employs Unicode Normalization Form C (NFC). Because NFKC is more aggressive about finding matches than NFC, in practice this change is unlikely to cause significant problems and indeed has the security benefit of probably resulting in fewer false positives when comparing passwords. A few examples might suffice to indicate the nature of the problem: (1) U+017F LATIN SMALL LETTER LONG S is compatibility equivalent to U+0073 LATIN SMALL LETTER S (2) U+2163 ROMAN NUMERAL FOUR is compatibility equivalent to U+0049 LATIN CAPITAL LETTER I and U+0056 LATIN CAPITAL LETTER V

(3) U+FB01 LATIN SMALL LIGATURE FI is compatibility equivalent to U+0066 LATIN SMALL LETTER F and U+0069 LATIN SMALL LETTER I. Under SASLprep, the use of NFKC also handled the mapping of fullwidth and halfwidth code points to their decomposition equivalents (see [[I-D.ietf-precis-mappings](#)]). Although it is expected that code points with compatibility equivalents are rare in existing passwords, some passwords that matched when SASLprep was used might no longer work when the rules in this specification are applied.

- o SASLprep mapped the "characters commonly mapped to nothing" from [Appendix B.1 of \[RFC3454\]](#) to nothing, whereas the PRECIS FreeformClass entirely disallows such characters, which correspond to the code points from the "M" category defined under [Section 6.13 of \[I-D.ietf-precis-framework\]](#) (with the exception of U+1806 MONGOLIAN TODO SOFT HYPHEN, which was commonly mapped to nothing in Unicode 3.2 but at the time of this writing is allowed by Unicode 6.2). In practice, this change will probably have no effect on comparison, but user-oriented software might reject such code points instead of ignoring them during password preparation.

[7.](#) Security Considerations

[7.1.](#) Password/Passphrase Strength

The ability to include a wide range of characters in passwords and passphrases can increase the potential for creating a strong password with high entropy. However, in practice, the ability to include such characters ought to be weighed against the possible need to reproduce them on various devices using various input methods.

[7.2.](#) Identifier Comparison

The process of comparing identifiers (such as SASL simple user names, authentication identifiers, and authorization identifiers) can lead to either false negatives or false positives, both of which have security implications. A more detailed discussion can be found in [[RFC6943](#)].

[7.3.](#) Reuse of PRECIS

The security considerations described in [[I-D.ietf-precis-framework](#)] apply to the "IdentifierClass" and "FreeformClass" base string classes used in this document for usernames and passwords, respectively.

7.4. Reuse of Unicode

The security considerations described in [[UTR39](#)] apply to the use of Unicode characters in usernames and passwords.

8. IANA Considerations

[Note to RFC Editor: please change XXXX to the number issued for this specification.]

8.1. Use of IdentifierClass

The IANA shall add an entry to the PRECIS Usage Registry for reuse of the PRECIS IdentifierClass, as follows:

Applicability: Usernames in security and application protocols.
Base Class: IdentifierClass.
Subclass: No.
Replaces: The SASLprep profile of Stringprep.
Width Mapping: Map fullwidth and halfwidth characters to their decomposition equivalents.
Additional Mappings: None.
Case Mapping: Not recommended, but to be defined by application protocols that use this profile.
Normalization: NFC.
Directionality: The "Bidi Rule" defined in [RFC 5893](#) applies.
Specification: RFC XXXX.

8.2. Use of FreeformClass

The IANA shall add an entry to the PRECIS Usage Registry for reuse of the PRECIS FreeformClass, as follows:

Applicability: Passwords in security and application protocols.
Base Class: FreeformClass
Subclass: No.
Replaces: The SASLprep profile of Stringprep.
Width Mapping: None.
Additional Mappings: Map non-ASCII space characters to ASCII space.
Case Mapping: None.
Normalization: NFC.
Directionality: None.
Specification: RFC XXXX.

9. References

9.1. Normative References

- [I-D.ietf-precis-framework]
Saint-Andre, P. and M. Blanchet, "Precis Framework: Handling Internationalized Strings in Protocols", [draft-ietf-precis-framework-09](#) (work in progress), July 2013.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, [RFC 3629](#), November 2003.
- [RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, [RFC 5234](#), January 2008.
- [UNICODE] The Unicode Consortium, "The Unicode Standard, Version 6.1", 2012,
<<http://www.unicode.org/versions/Unicode6.1.0/>>.

9.2. Informative References

- [I-D.ietf-precis-mappings]
YONEYA, Y. and T. NEMOTO, "Mapping characters for PRECIS classes", [draft-ietf-precis-mappings-02](#) (work in progress), May 2013.
- [RFC20] Cerf, V., "ASCII format for network interchange", [RFC 20](#), October 1969.
- [RFC2617] Franks, J., Hallam-Baker, P., Hostetler, J., Lawrence, S., Leach, P., Luotonen, A., and L. Stewart, "HTTP Authentication: Basic and Digest Access Authentication", [RFC 2617](#), June 1999.
- [RFC3454] Hoffman, P. and M. Blanchet, "Preparation of Internationalized Strings ("stringprep")", [RFC 3454](#), December 2002.
- [RFC4013] Zeilenga, K., "SASLprep: Stringprep Profile for User Names and Passwords", [RFC 4013](#), February 2005.
- [RFC4422] Melnikov, A., Ed. and K. Zeilenga, Ed., "Simple Authentication and Security Layer (SASL)", [RFC 4422](#), June 2006.
- [RFC4616] Zeilenga, K., "The PLAIN Simple Authentication and

- Security Layer (SASL) Mechanism", [RFC 4616](#), August 2006.
- [RFC5802] Newman, C., Menon-Sen, A., Melnikov, A., and N. Williams, "Salted Challenge Response Authentication Mechanism (SCRAM) SASL and GSS-API Mechanisms", [RFC 5802](#), July 2010.
- [RFC5890] Klensin, J., "Internationalized Domain Names for Applications (IDNA): Definitions and Document Framework", [RFC 5890](#), August 2010.
- [RFC5891] Klensin, J., "Internationalized Domain Names in Applications (IDNA): Protocol", [RFC 5891](#), August 2010.
- [RFC5893] Alvestrand, H. and C. Karp, "Right-to-Left Scripts for Internationalized Domain Names for Applications (IDNA)", [RFC 5893](#), August 2010.
- [RFC5894] Klensin, J., "Internationalized Domain Names for Applications (IDNA): Background, Explanation, and Rationale", [RFC 5894](#), August 2010.
- [RFC6120] Saint-Andre, P., "Extensible Messaging and Presence Protocol (XMPP): Core", [RFC 6120](#), March 2011.
- [RFC6365] Hoffman, P. and J. Klensin, "Terminology Used in Internationalization in the IETF", [BCP 166](#), [RFC 6365](#), September 2011.
- [RFC6943] Thaler, D., "Issues in Identifier Comparison for Security Purposes", [RFC 6943](#), May 2013.
- [UTR39] The Unicode Consortium, "Unicode Technical Report #39: Unicode Security Mechanisms", August 2010, <<http://unicode.org/reports/tr39/>>.

[Appendix A](#). Differences from [RFC 4013](#)

This document builds upon the PRECIS framework defined in [[I-D.ietf-precis-framework](#)], which differs fundamentally from the stringprep technology [[RFC3454](#)] used in SASLprep [[RFC4013](#)]. The primary difference is that stringprep profiles allowed all characters except those which were explicitly disallowed, whereas PRECIS profiles disallow all characters except those which are explicitly allowed (this "inclusion model" was originally used for internationalized domain names in [[RFC5891](#)]; see [[RFC5894](#)] for further discussion). It is important to keep this distinction in mind when comparing the technology defined in this document to

SASLprep [[RFC4013](#)].

The following substantive modifications were made from [RFC 4013](#).

- o A single SASLprep algorithm was replaced by two separate algorithms: one for usernames and another for passwords.
- o The new preparation algorithms use PRECIS instead of a stringprep profile. The new algorithms work independently of Unicode versions.
- o As recommended in the PRECIS framework, changed the Unicode normalization form to NFC (from NFKC).
- o Some Unicode code points that were mapped to nothing in [RFC 4013](#) are simply disallowed by PRECIS.

[Appendix B](#). Acknowledgements

The following individuals provided helpful feedback on this document: Marc Blanchet, Alan DeKok, Joe Hildebrand, Jeffrey Hutzelman, Simon Josefsson, Jonathan Lennox, Matt Miller, Chris Newman, Yutaka OIWA, Pete Resnick, Andrew Sullivan, and Nico Williams (Nico in particular provided text that was used in [Section 2.2](#)). Thanks also to Yoshiro YONEYA and Takahiro NEMOTO for implementation feedback.

This document borrows some text from [[RFC4013](#)] and [[RFC6120](#)].

Authors' Addresses

Peter Saint-Andre
Cisco Systems, Inc.
1899 Wynkoop Street, Suite 600
Denver, CO 80202
USA

Phone: +1-303-308-3282
Email: psaintan@cisco.com

Alexey Melnikov
Isode Ltd
5 Castle Business Village
36 Station Road
Hampton, Middlesex TW12 2BX
UK

Email: Alexey.Melnikov@isode.com

