Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: June 6, 2013

P. Saint-Andre
Cisco Systems, Inc.

December 3, 2012

The 'acct' URI Scheme
draft-ietf-appsawg-acct-uri-02

Abstract

This document defines the 'acct' URI scheme as a way to identify a user's account at a service provider, irrespective of the particular protocols that can be used to interact with the account.

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 June 6, 2013.

Copyright Notice

Table of Contents

1. Introduction

Existing URI schemes that enable interaction with, or that identify resources associated with, a user's account at a service provider are tied to particular services or application protocols. Two examples are the 'mailto' scheme (which enables interaction with a user's email account) and the 'http' scheme (which enables retrieval of web files controlled by a user or interaction with interfaces providing information about a user). However, there exists no URI scheme that generically identifies a user's account at a service provider without specifying a particular protocol to use when interacting with the account. This specification fills that gap.

2. Rationale

During formalization of the WebFinger protocol [I-D.ietf-appsawg-webfinger], much discussion occurred regarding the appropriate URI scheme to include when specifying a user's account as a web link [RFC5988]. Although both the 'mailto' [RFC6068] and 'http' [RFC2616] schemes were proposed, not all service providers offer email services or web interfaces on behalf of user accounts (e.g., a microblogging or instant messaging provider might not offer email services, or an enterprise might not offer HTTP interfaces to information about its employees). Therefore, the discussants recognized that it would be helpful to define a URI scheme that could be used to generically identify a user's account at a service provider, irrespective of the particular application protocols used to interact with the account. The result was the 'acct' URI scheme defined in this document.

3. Definition

The syntax of the 'acct' URI scheme is defined under Section 4 of this document. Although 'acct' URIs take the form "user@host", the

scheme is designed for the purpose of identification instead of interaction (regarding this distinction, see Section 1.2.2 of [RFC3986]). The "Internet resource" identified by an 'acct' URI is a user's account hosted at a service provider, where the service provider is typically associated with a DNS domain name. Thus a particular 'acct' URI is formed by setting the "user" portion to the user's account name at the service provider and by setting the "host" portion to the DNS domain name of the service provider.

Consider the case of a user with an account name of "foobar" on a microblogging service "status.example.net". It is taken as convention that the string "foobar@status.example.net" designates

that account. This is expressed as a URI using the 'acct' scheme as "acct:foobar@status.example.net".

It is not assumed that an entity will necessarily be able to interact with a user's account using any particular application protocol, such as email; to enable such interaction, an entity would need to use the appropriate URI scheme for such a protocol, such as the 'mailto' scheme. While it might be true that the 'acct' URI minus the scheme name (e.g., "user@example.com" derived from "acct:user@example.com") can be reached via email or some other application protocol, that fact would be purely contingent and dependent upon the deployment practices of the provider.

Because an 'acct' URI enables identification only and not interaction, it cannot be deferenced on its own and is not employed directly in a protocol, e.g., as the value of the 'href' attribute of an HTML anchor element. For example, an 'acct' URI would not be used as follows:

```
<a href='acct:bob@example.com'>find out more</a>
```

Instead, an 'acct' URI is employed indirectly and typically is passed around as a parameter in the background within a protocol flow so that an entity can interact with the resource identified by the 'acct' URI in a particular way or for a particular purpose. For example, in the WebFinger protocol [I-D.ietf-appsawg-webfinger] an 'acct' URI is used to identify the resource about which an entity would like to discover metadata expressed as "web links" [RFC5988]; the relevant HTTP request passes an 'acct' URI (or some other URI) as

the value of a "resource" parameter, as shown in the following example:

```
GET /.well-known/webfinger?resource=acct%3Abob%40example.com HTTP/1.1
```

Therefore, any protocol that uses 'acct' URIs, such as the WebFinger protocol [I-D.ietf-appsawg-webfinger] or the Simple Web Discovery protocol [I-D.jones-simple-web-discovery], is responsible for specifying how an 'acct' URI is employed in the context of that protocol (in particular, how it is dereferenced or resolved; see [RFC3986]). As a concrete example, in the WebFinger protocol an 'acct' URI is passed as a parameter in an HTTP request for metadata (i.e., web links) about the resource; the service retrieves the metadata associated with the account identified by that URI and then provides that metadata to the requesting entity in an HTTP response (see [I-D.ietf-appsawg-webfinger] for details). Similar functionality is envisioned for other uses of 'acct' URIs.

4. IANA Considerations

In accordance with the guidelines and registration procedures for new URI schemes [RFC4395], this section provides the information needed to register the 'acct' URI scheme.

4.1. URI Scheme Name

acct

4.2. Status

permanent

4.3. URI Scheme Syntax

The 'acct' URI syntax is defined here in Augmented Backus-Naur Form (ABNF) [RFC5234], borrowing the 'host', 'pct-encoded', 'sub-delims', 'unreserved' rules from [RFC3986]:

```
acctURI      =  "acct" ":" userpart "@" host
userpart     =  1*( unreserved / pct-encoded / sub-delims )
```

4.4. URI Scheme Semantics

The 'acct' URI scheme identifies accounts hosted at service providers. It is used only for identification, not interaction. A protocol that employs the 'acct' URI scheme is responsible for specifying how an 'acct' URI is dereferenced in the context of that protocol. There is no media type associated with the 'acct' URI scheme.

4.5. Encoding Considerations

The 'acct' URI scheme allows any character from the Unicode repertoire [UNICODE] encoded as UTF-8 [RFC3629] and then percent-encoded into valid ASCII [RFC20] as specified in [RFC3986]. Note that domain labels need to be encoded as A-labels (see [RFC5890]) in order to support internationalized domain names (IDNs).

4.6. Applications/Protocols That Use This URI Scheme Name

At the time of this writing, only the WebFinger protocol uses the 'acct' URI scheme. However, use is not restricted to the WebFinger protocol, and the scheme might be considered for use in other protocols, such as Simple Web Discovery.

4.7. Interoperability Considerations

There are no known interoperability concerns related to use of the 'acct' URI scheme.

4.8. Security Considerations

See Section 5 of RFCXXXX. [Note to RFC Editor: please replace XXXX with the number issued to this document.]

4.9. Contact

Peter Saint-Andre, psaintan@cisco.com

4.10. Author/Change Controller

This scheme is registered under the IETF tree. As such, the IETF maintains change control.

4.11. References

None.

5. Security Considerations

Because the 'acct' URI scheme does not directly enable interaction with a user's account at a service provider, possible security concerns are minimized.

Protocols that make use of 'acct' URIs are responsible for defining security considerations related to such usage, e.g., the risks involved in dereferencing an 'acct' URI and the authentication and authorization methods that could be used to control access to personally identifying information associated with a user's account at a service.

6. References

6.1. Normative References

[RFC20] Cerf, V., "ASCII format for network interchange", RFC 20, October 1969.

[RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, November 2003.

[RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.

[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/>.

6.2. Informative References

[I-D.ietf-appsawg-webfinger]
Jones, P., Salgueiro, G., and J. Smarr, "WebFinger", draft-ietf-appsawg-webfinger-07 (work in progress), December 2012.

[I-D.jones-simple-web-discovery]
Jones, M. and Y. Goland, "Simple Web Discovery (SWD)", draft-jones-simple-web-discovery-04 (work in progress), November 2012.

[RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999.

[RFC4395] Hansen, T., Hardie, T., and L. Masinter, "Guidelines and Registration Procedures for New URI Schemes", BCP 35, RFC 4395, February 2006.

[RFC5890] Klensin, J., "Internationalized Domain Names for Applications (IDNA): Definitions and Document Framework", RFC 5890, August 2010.

[RFC5988] Nottingham, M., "Web Linking", RFC 5988, October 2010.

[RFC6068] Duerst, M., Masinter, L., and J. Zawinski, "The 'mailto' URI Scheme", RFC 6068, October 2010.

Appendix A. Acknowledgements

The 'acct' URI scheme was originally proposed during work on the WebFinger protocol; special thanks are due to Blaine Cook, Brad Fitzpatrick, and Eran Hammer-Lahav for their early work on the concept (which in turn was partially inspired by work on Extensible Resource Indentifiers at OASIS). The scheme was first formally

specified in [I-D.ietf-appsawg-webfinger]; the authors of that

specification (Paul Jones, Gonzalo Salgueiro, and Joseph Smarr) are gratefully acknowledged. Thanks are also due to Martin Duerst, Graham Klyne, Barry Leiba, and many other participants in the IETF APPSAWG for their feedback.

Author's Address

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

Email: psaintan@cisco.com