INTERNET-DRAFT   Eric A. Hall
Document: draft-ietf-crisp-firs-contact-03.txt   August 2003
Expires: March, 2004
Category: Standards-Track

Defining and Locating Contact Information in the Federated Internet Registry Service

Status of this Memo

Copyright Notice

Abstract

This document defines LDAP schema and searching rules for contact persons, in support of the Federated Internet Registry Service (FIRS) described in [FIRS-ARCH] and [FIRS-CORE].

Table of Contents

1. Introduction

This specification defines the naming syntax, object classes, attributes, matching filters, and query processing rules for storing and locating contact persons in the FIRS service. Refer to [FIRS-ARCH] for information on the FIRS architecture and [FIRS-CORE] for the schema definitions and rules which govern the FIRS service as a whole.

The definitions in this specification are intended to be used with FIRS. Their usage outside of FIRS is not prohibited, but any such usage is beyond this specification's scope of authority.

2. Prerequisites and Terminology

The complete set of specifications in the FIRS collection cumulative define a structured and distributed information service using LDAPv3 for the data-formatting and transport functions. This specification should be read in the context of that set, which currently includes [FIRS-ARCH], [FIRS-CORE], [FIRS-DNS], [FIRS-DNSRR], [FIRS-ASN], [FIRS-IPV4] and [FIRS-IPV6].

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 RFC 2119.

3.   Naming Syntax

The naming syntax for contact entries in FIRS MUST follow the form of "cn=<inetContactSyntax>,cn=inetResources,<partition>", where <inetContactSyntax> is an email address representing a contact resource, and where <partition> is a sequence of domainComponent relative distinguished names which identifies the scope of authority for the selected directory partition.

The inetContactSyntax is unstructured, in that it uses standardized procedures to produce heavily-normalized email addresses rather than using structured syntax rules. The principle reason for this is due to conflicting syntax rules in different canonical email addressing rules, with these rules preventing the use of a common syntax.

The normalization procedure produces UTF-8 [RFC2279] email addresses as output, with these domain names being suitable for direct comparisons, substring searches, and other lightweight comparisons. Servers tend to be more heavily-loaded than clients, and requiring the data to be normalized before it is used for comparison operations ensures that a broader range of comparison operations can be performed with minimal impact on those servers.

This normalization procedure is as follows:

a. Email addresses MUST contain three elements, which are a localpart element, an "at" sign ("@") separator character, and a domain element.

b. The localpart element is currently unspecified, pending ongoing effort to internationalize this element. Subsequent versions of this specification may define specific handling rules for this element.

c. The domain element MUST be normalized according to the inetDnsDomainSyntax procedure defined in [FIRS-DNS].

Once all of these steps have successfully completed, the email address can be stored in the directory or used as an assertion value. Any fatal error conditions encountered during these conversions MUST result in a local failure; FIRS-aware applications MUST NOT store or transmit non-normalized email addresses for any purposes.

The inetContactSyntax syntax is as follows:

```
     inetContactSyntax
     ( 1.3.6.1.4.1.7161.1.4.0
       NAME 'inetContactSyntax'
       DESC 'A fully-qualified email address.' )
```

Note that the use of the "at" separator character is illegal as data in URLs, and these characters will be escaped before they are stored in a URL as data.

Also note that UTF-8 characters use character codes which are frequently illegal as data in URLs, and many of those octet values will probably be escaped before they are stored in a URL as data.

4. Object Classes and Attributes

Contact entries in FIRS MUST use the inetOrgPerson object class as defined in RFC 2798 [RFC2798], in addition to the mandatory object classes defined in [FIRS-CORE]. Contact entries MUST be treated as containers capable of holding subordinate entries.

If an entry exists as a referral source, the entry MUST be defined with the referral object class, in addition to the other object classes defined above. Referral sources MUST NOT contain subordinate entries. Refer to section 3.5 of [FIRS-CORE] for more information on referral entries in FIRS.

The inetOrgPerson object class is a structural object class. The inetOrgPerson object class has three mandatory attributes (cn, sn, and objectClass), and has several optional attributes. Contact entries also inherit the attributes defined in the inetResources object class when they are used with FIRS.

Refer to [RFC2798] for the inetOrgPerson schema definitions.

Note that the "mail" attribute defined for use with the inetOrgPerson object class is restricted to seven-bit character codes and is typically interpreted as [US-ASCII], and is therefore not compatible with the inetContactSyntax rules defined in section 3. As such, if the mail domain uses an internationalized domain name, the domain element of the mail attribute MUST be reduced to its ASCII-compatible form using the ToASCII process defined in [RFC3490], and MUST NOT use the UTF-8 encoding.

Note that International postal regulations generally require that the recipient address on an envelope be provided in a language and charset which is native to the recipient's country, with the exception of the destination country name which should be provided in a language and charset that is native to the sender's country. This model ensures that the sender's post office will be able to route the mail to the recipient's country, while also ensuring that the destination country's post office will be able to perform local delivery. In order to facilitate this usage, the country attribute value MAY (encouraged) be localized to the local user's nomenclature for a country, but other postal address information SHOULD NOT be localized.

Notwithstanding the above, it is ENCOURAGED that contact names be provided in English forms in order to facilitate inter-party communications, using the mechanisms offered by [RFC2596]. For example, the default contact entry for a person in Japan SHOULD be provided in the native form for that person, but an English form is also ENCOURAGED in order to allow non-Japanese users to properly address that person in subsequent communications. As stated in the preceding paragraph however, any postal communications for that person SHOULD use the native-language representation (at least on the envelope) in order to facilitate the delivery of postal mail.

An example of the inetOrgPerson object class in use is shown in Figure 1 below. The example includes attributes from the inetOrgPerson, inetResources, and inetAssociatedResources object classes.

```
cn=admins@example.com,cn=inetResources,dc=example,dc=com
[top object class]
[inetResources object class]
[inetOrgPerson object class]
[inetAssociatedResources object class]
|
+-attribute: mail
| value: "admins@example.com"
|
+-attribute: inetAssociatedIpv4Network
  value: "192.0.2.0/24"
```

Figure 1: The entry for the admins@example.com contact in the dc=netsol,dc=com partition.

5. Query Processing Rules

Queries for contact entries have several special requirements, as discussed in the following sections.

Refer to [FIRS-CORE] for general information about FIRS queries.

5.1. Query Pre-Processing

Clients MUST ensure that the query input is normalized according to the rules specified in section 3 before the input is used as the assertion value to the resulting LDAP query.

The authoritative partition for a contact entry is determined by mapping the domain element of a normalized email address to a sequence of domainComponent labels.

Since the domainComponent attribute is restricted to seven-bit characters, the domain element MUST be converted to its IDNA form using the "ToASCII" conversion operation specified in [RFC3490], with the "UseSTD3ASCIIRules" flag disabled (FIRS applications MAY reuse the output from the conversion performed in step 3.c if the entire conversion process is known to have completed successfully). The resulting sequence of ASCII labels are used to form the domainComponent sequence which represents the authoritative partition for the email address.

As a simple example, "admins@example.com" would be mapped to the "dc=example,dc=com" authoritative partition, with this partition being used to seed the query process.

5.2. Query Bootstrapping

FIRS clients MUST use the bottom-up bootstrap model by default for contact queries. As such, the search base for default queries would be set to the complete sequence of domainComponent relative distinguished names of the authoritative partition.

FIRS clients MAY use the targeted or top-down bootstrap models for queries if necessary or desirable. However, it is not likely that entries will be found for all possible contacts using these models (the "dc=com" partition is not likely to have entries for all of the possible contacts with mailboxes in the "com" hierarchy, for example). As such, the bottom-up bootstrap model will be the most useful in most cases, and MUST be used by default.

Note that registration bodies can allocate email addresses within their own managed portion of the DNS namespace if predictability is at a premium. For example, a registrar could assign "user@registrar.com" email addresses to the contact entries that it creates, thereby ensuring that the contact entries are always locatable and managed.

5.3. LDAP Matching

If the server advertises the inetOrgPerson object class and the inetContactMatch matching filter in the inetResourcesControl server control, FIRS clients MUST use the inetDnsDomainMatch matching filter in LDAP searches for contact entries.

The inetContactMatch filter provides an identifier and search string format which collectively inform a queried server that a specific contact identifier should be searched for, and that any inetOrgPerson object class entries which match the assertion value should be returned.

The inetContactMatch filter is defined as follows:

```
     inetContactMatch
     ( 1.3.6.1.4.1.7161.1.4.0.1
       NAME 'inetDnsDomainMatch'
       SYNTAX 1.3.6.1.4.1.7161.1.4.0 )
```

Clients MUST ensure that the query input is normalized according to the rules specified in section 3 before the input is used as the assertion value to the resulting LDAP query.

A FIRS server MUST compare the assertion value against the distinguished name of all entries within and beneath the container specified by the search base of the query. Any entry in that hierarchy with an object class of inetOrgPerson and a distinguished name component that is equal to the assertion value MUST be returned to the client (this specifically includes any child entries, such as referral stubs). Entries which do not have an object class of inetOrgPerson MUST NOT be returned.

The matching filters defined in this specification MUST be supported by FIRS clients and servers. FIRS servers MAY support additional matching filters, although FIRS clients MUST NOT expect any additional filters to be available.

If the server does not advertise support for the inetContactMatch matching filter in the inetResourcesControl server control, the client MAY choose to emulate the matching filter through the use of locally-constructed equalityMatch filters. However, this process can result in incomplete answers in some cases, so if the server advertises support for the inetContactMatch matching filter in the inetResourcesControl control, the client MUST use it.

5.4. Example Query

The following example assumes that the user has specified "admins@example.com" as the query value:

a. Normalize the input, which is "admins@example.com" in this case.

b. Determine the canonical authoritative partition, which is "dc=example,dc=com" in this case. By default, queries for contacts use the bottom-up model, meaning that the fully-qualified distinguished name of "dc=example,dc=com" will be used.

c. Determine the search base for the query, which will be "cn=inetResources,dc=example,dc=com" if the defaults are used.

d. Initiate a DNS lookup for the SRV resource records associated with "_ldap._tcp.example.com." For the purpose of this example, assume that this lookup succeeds, with the DNS response message indicating that "firs.example.com" is the preferred LDAP server.

e. Submit an LDAPv3 query to the specified server, using "(1.3.6.1.4.1.7161.1.4.0.1:=admins@example.com)" as the matching filter, "cn=inetResources,dc=example, dc=com" as the search base, and the global query defaults defined in [FIRS-CORE].

f. Assume that no referrals are received. Display the answer data which has been received and exit the query.

6. Security Considerations

Security considerations are discussed in [FIRS-ARCH].

7. IANA Considerations

IANA considerations are discussed in [FIRS-ARCH].

8. Normative References

[FIRS-ARCH] Hall, E. "The Federated Internet Registry Service: Architecture and Implementation Guide", draft-ietf-crisp-firs-arch-03, August 2003.

[FIRS-ASN] Hall, E. "Defining and Locating Autonomous System Numbers in the Federated Internet Registry Service", draft-ietf-crisp-firs-asn-03, August 2003.

[FIRS-CORE] Hall, E. "The Federated Internet Registry Service: Core Elements", draft-ietf-crisp-firs-core-03, August 2003.

[FIRS-DNS] Hall, E. "Defining and Locating DNS Domains in the Federated Internet Registry Service", draft-ietf-crisp-firs-dns-03, August 2003.

[FIRS-DNSRR] Hall, E. "Defining and Locating DNS Resource Records in the Federated Internet Registry Service", draft-ietf-crisp-firs-dnsrr-03, July 2003.

[FIRS-IPV4] Hall, E. "Defining and Locating IPv4 Address Blocks in the Federated Internet Registry Service", draft-ietf-crisp-firs-ipv4-03, August 2003.

[FIRS-IPV6] Hall, E. "Defining and Locating IPv6 Address Blocks in the Federated Internet Registry Service", draft-ietf-crisp-firs-ipv6-03, August 2003.

[RFC2247] Kille, S., Wahl, M., Grimstad, A., Huber, R., and Sataluri, S. "Using Domains in LDAP/X.500 DNs", RFC 2247, January 1998.

[RFC2251] Wahl, M., Howes, T., and Kille, S. "Lightweight Directory Access Protocol (v3)", RFC 2251, December 1997.

[RFC2252] Wahl, M., Coulbeck, A., Howes, T., and Kille, S. "Lightweight Directory Access Protocol (v3): Attribute Syntax Definitions", RFC 2252, December 1997.

[RFC2254]   Howes, T. "The String Representation of LDAP Search Filters", RFC 2254, December 1997.

[RFC2279]   Yergeau, F. "UTF-8, a transformation format of ISO 10646", RFC 2279, January 1998.

[RFC2596]   Wahl, M., and Howes, T. "Use of Language Codes in LDAP", RFC 2596, May 1999.

[RFC2798]   Smith, M. "Definition of the inetOrgPerson LDAP Object Class", RFC 2798, April 2000.

[RFC3490]   Faltstrom, P., Hoffman, P., and Costello, A. "Internationalizing Domain Names in Applications (IDNA)", RFC 3490, March 2003.

[US-ASCII]   Cerf, V. "ASCII format for Network Interchange", RFC 20, October 1969.

9. Changes from Previous Versions

draft-ietf-crisp-firs-contact-03:

- Several clarifications and corrections have been made.
- The inetContactMatch matching filter was defined. The use of equalityMatch and extensibleMatch has been deprecated.

draft-ietf-crisp-firs-contact-02:

- Several clarifications and corrections have been made.

draft-ietf-crisp-firs-contact-01:

- Several clarifications and corrections have been made.
- Several attributes had their OIDs changed. NOTE THAT THIS IS AN INTERNET DRAFT, AND THAT THE OIDS ARE SUBJECT TO ADDITIONAL CHANGES AS THIS DOCUMENT IS EDITED.

draft-ietf-crisp-firs-contact-00:

- Restructured the document set.
- "Attribute references" have been eliminated from the specification. All referential attributes now provide actual data instead of URL pointers to data. Clients that

wish to retrieve these values will need to start new queries using the data values instead of URLs.

draft-ietf-crisp-lw-user-01:

- Removed references to LDAPS (LDAP-over-SSL), which is not a standards-track protocol.
- Added a discussion on localization considerations.
- Moved attribute-specific security requirements to the Security section.

10. Author's Addresses

Eric A. Hall
ehall@ehsco.com

11. Acknowledgments

Funding for the RFC editor function is currently provided by the Internet Society.

Portions of this document were funded by VeriSign Labs.

The first version of this specification was co-authored by Andrew Newton of Verisign Labs, and subsequent versions continue to be developed with his active participation.

12. Full Copyright Statement