INTERNET-DRAFT
<draft-ietf-schema-rqmts-list-01.txt>
Expires: October 21, 1998

C. Apple
AT&T Labs
21 April 1998

Requirements for the Initial Release of a Directory Schema Listing Service
<draft-ietf-schema-rqmts-list-01.txt>

Status of this Memo

Abstract

This memo documents requirements for listing directory services schema in a centrally operated, administered, and maintained repository. This repository will be available as a resource to directory protocol and service implementors to facilitate schema discovery.

Table of Contents

1.0 Introduction

The fastest route to interoperable directory services is through standard object classes and attribute types. There is a growing number of places where schema for Internet Directory Services and Internet Operations are being defined, with varying degrees of documentation. This plethora of schema is unavoidable in the light of the needs of different service communities, but it makes it difficult for directory service builders to find and make use of an existing schema that will serve their needs and increase interoperability with other systems. A listing service providing a single point of discovery for directory service schema will promote schema reuse, reduce duplication of effort, and thus promote directory service interoperability.

The intent is to offer a schema listing service with public read access and restricted, moderated write access. Many hard-coded choices and constraints have been reflected in this requirements document for the purpose of expediting deployment. Future releases of the service may require an update of this document.

Initially, such a listing service will be centrally operated, administered, and maintained. The schema listing repository database may also be mirrored to ensure some level of redundancy for read access in the event of service interruption. Eventually, the operations, administration, and maintenance of such a listing service may evolve to use a more distributed deployment scenario.

The schema listing service is also intended to be largely automated, with minimal human involvement. Human involvement is likely to be limited to the following types of activities:

+ handling repository access problems
+ trouble resolution for computing and communications facilities
+ dealing with reasonable requests that fall outside of the scope of normal schema listing repository operations
+ reviewing schema listing requests on a mailing list prior to publishing in the listing repository

Future releases of the service may automate some of these tasks.

1.1 Scope

Requirements for the initial release of a directory schema listing service are inside the scope of this document.

Specifications for syntaxes and grammars to be used in the initial release of the directory schema listing service are outside the scope

of this document.

Documentation of schema listing procedures is outside the scope of this document.

1.2 Terms and Definitions

Information Object - a descriptive abstraction of some real-world object

Object Attribute - a descriptive property of an information object; typically, object attributes are defined in terms of semantic and syntactic definitions

Schema - a collection of definitions for related information objects

Schema Unit - a related or grouped set of object attributes that form a discrete unit within the context of a schema for a particular protocol; examples include an LDAP object class or a WHOIS++ template

Schema Pak - a related or grouped set of schema units that collectively specify a schema associated with a particular protocol; an example of a schema pak is the set of LDAP object classes specified in [RFC2256]

Metadata - characteristics that differentiate one schema unit or schema pak from another; used to catalog listing service content; structured using a profile of [MIMEDIR]; also contains references to files stored within and outside of a listing repository

Schema Unit Content - a formal specification of a schema unit using a profile of [MIMEDIR]

Schema Unit Listing - the combination of a single schema unit content file intended for use within the context of a particular protocol and a file containing metadata describing the schema unit specified within that schema unit content file

Schema Pak Listing - a single metadata file containing information describing and referring to a set of related or grouped schema unit content files

Repository - a database in which listings are stored

Listing Request - a proposed schema unit listing or schema pak listing formatted using [MIME] constructs that is submitted for consideration as a listing to be published in a repository

Operator - an organization that administers and maintains a repository

Primary Repository - the repository that masters the schema listings database

Shadow Repository - a repository that mirrors the primary repository

Contact Person - the name of the individual who holds the authority to update a listing and who should be contacted if questions or concerns arise related to a listing or listing request

Listing Authority Contact - the name of the individual who holds authority to replace a contact person; can be either the contact person for a listing or an alternate contact within the organization to which the contact person belongs (this allows one person organizations to list schema)

The terms for specifying requirement level defined in [RFC2119] are used in this document.

1.3 Usage Scenarios

1.3.1 Location/Retrieval of the vCard Schema for LDAPv3

A user of the schema listing service wants to locate a copy of the vCard schema for LDAPv3 [RFC2251] so that they can use it in a prototyping project. First, they point their web browser at a schema listing repository web site and download the list of available schema. Next, they use the search command on their browser to locate occurances of the string "vCard". The browser automatically scrolls down to the appropriate place in the list of available schema and the user clicks on a link to view the listing metadata to verify that this is indeed the vCard schema for use with version 3 of the LDAP protocol. Included in the web-based representation of the listing metadata are ftp URLs pointing to available profiles containing listing content for this schema. The user clicks on the link for the profile that they can use.

1.3.2 Submission of a New Schema Listing via SMTP

A schema writer wishes to list a schema they have created and prepares the listing metadata and listing content according to one or more appropriate [MIMEDIR] profiles. The schema writer will obtain a permanent, unique schema listing name for the request.

The schema writer sends an SMTP message including the listing meta data and all available listing content in multipart-related [MIME]

format to a listing request review mailing list. After a short review period, the listing repository operator validates the request, and if properly formed, publishes the listing according to the listing procedures. An announcement of the newly published schema listing is sent to a mailing list reserved for this purpose.

2.0 Listing Service Requirements

2.1 Functional Requirements

Listed schema MAY be published as an RFC.

A list of available listings MUST be maintained.

Listings MUST be named according to the namespace requirements defined in section 3.

The listing service SHALL maintain information about schema units, beyond their definition. This information is referred to as metadata and will consist of information used for cataloging listings in the repositories. The particular set of metadata elements used during the initial deployment of the listing service will be defined in other documents.

Listing metadata and listing content MUST be parsable.

2.2 Operational Requirements

The process of listing schema MUST be centralized for the initial deployment.

All versions of all listings MUST be retained. A simple method for getting the most recent version of a particular listing MUST be provided.

The contact person for a listing MAY give an earlier listing a higher version number, or MAY request that the listing get a new name.

The listing repository MUST be centrally administered.

The listing repository MAY be mirrored.

The primary repository operator MUST obtain an OID subtree for which they hold sub-allocation authority for use in the schema listing service.

Listing requests MAY be signed using PGP/MIME as described in [RFC2015]. The primary listing repository operator MUST be able to

accept schema listing requests in PGP/MIME messages, although they are NOT REQUIRED to validate the signatures. The method for validating and determining trust of signatures is outside the scope of this document and is determined by the parties in the exchange. The method for determining and validating trust in an unsigned request is outside the scope of this document, as is the method for determining trust in schema listing repository or its content.

A mailing list MUST be created for the purpose of submitting listing requests for review prior to publishing in the schema listing repository. The schema listing repository publication process MUST be moderated via this mailing list. Listing requests MUST be subjected to community review on this mailing list for a period of at least 2 weeks. If no comments are received, properly formed schema listing requests SHALL be published as listings; otherwise, the request MAY either denied or the listing MAY published subject to incorporation of comments.

A mailing list MUST be created for announcing new and updated listings.

A mailbox MUST be created for the purpose of receiving service trouble requests from users.

Listing repository operators (of primary and shadow sites) MUST provide a free means of accessing the listing service consistent with the functionality documented in paragraph 2.3.

2.3 Repository Access Functionality

The following schema listing repository access protocols MUST be supported: FTP [RFC959], HTTP 1.1 [RFC2068], and SMTP [RFC821].

The following access functions are REQUIRED:

a) browse and retrieve schema unit content, metadata, and a list of available listings:

+ via HTTP requests

+ via FTP clients

+ via requests through an SMTP server

b) search a list of available listings:

+ via HTTP, retrieving either HTML or text listings that can then be searched by the requestor

+ via HTTP by accessing repository-based searching facilities such as keyword searching; this can return listing names, schema unit listings, schema pak listings, metadata, or other useful information

c) add and update listings by submitting a formatted request to a mailing list for community review:

+ via SMTP using appropriate MIME constructs as described in section 4.0

Other access functions, including the following, MAY be supported, but will be defined in other documents in the future:

a) search schema unit content

b) search metadata

3.0 Listing Service Namespace Requirements

The listing service namespace MUST be protocol-independent.

The listing service namespace SHALL be based on OIDs.

Listing names:

+ MUST be permanent

+ MUST be globally unique

+ MUST be publicly available

+ MUST NOT be recycled or re-used

+ MUST be created within the OID subtree reserved for use in the schema listing service and administered by the primary listing repository operator

4.0 Listing Requirements

Schema unit content SHALL be limited to the information actually required to specify and encode the schema for storage and transfer.

Metadata SHALL be composed of information used to catalog listings.

Metadata element syntax SHALL be defined based on the concept of tagged attribute type-value pairs.

Language tags as specified in [RFC1766] MUST be used in all listings.

Metadata element values MUST be encoded using the UCS Transformation Format - 8 bit form [RFC2044].

For the purposes of submitting a listing request, schema unit content and metadata SHOULD be structured according to appropriate profiles of [MIMEDIR] defined in other documents.

Content associated with a listing, but not stored in the schema listing repository (such as large copyright notices and vendor logo images) MAY be included by reference in the metadata. If such external references are included in a particular schema listing, a fingerprint of the external content generated prior to schema listing request creation MUST be included along with these references in the request. Details associated with the creation of these external content references, including the algorithm to be used for generation of a content fingerprint and the syntax of the reference, will be defined in the [MIMEDIR] profile used to format and encode listing metadata for storage and transfer.

5.0 Listing Storage Requirements

Listing repository file names MUST be permanent, globally unique, and publicly available.

Listing repository file names SHOULD be constructed in a manner that allows human and machine users to determine the nature of file content by inspecting the file names.

Schema unit content and metadata MUST be stored in separate files.

6.0 Security Considerations

6.1 Compromisable Assets

One or more of the following assets could be compromised if the service is attacked:

+ Metadata
+ Schema unit content
+ Repository Hardware & Software
+ Networking Facilities Connecting Repository to the Internet
+ Repository Mirror Sites

6.2 Attack Scenarios

Allowable methods for submitting listing requests are:

a) sending an e-mail message to a mail box

b) submitting requests using web-based forms

Based on these request submission methods, there are a number of known repository attack scenarios that must be considered during the implementation of schema listing procedures and the software and operational processes required to support them.

6.2.1 Denial-of-Service Attack Scenarios

Scenario A: someone could send in a large number of improperly formed requests

Scenario B: someone could send in a large number of properly formed, but frivolous, useless, or trivial requests

6.2.2 Confuse-the-User Attack Scenarios

Scenario A: someone could send in a large number of valid, but frivolous, useless, or trivial requests and some or all of these requests actually become listings in the repository

Scenario B: someone could maliciously submit one or more slightly modified versions of existing listings which are popular or widely used

6.3 Security Requirements on Schema Listing Procedures

The following contextual definitions apply to the requirements listed in the remainder of this paragraph:

Verification - a process of determining authenticity of facts implied or explicitly specified by a contact person during the process of submitting a schema listing request; the methods used to implement such a process MAY or MAY NOT be based on an IETF-sanctioned security protocol; specification of the methods used to implement such a process as well as the trust relationships relevant to the process are outside the scope of this document.

High-Quality Directory Schema - a directory schema that serves some useful purpose (e.g., a related set of attribute and object class definitions for holding information about people in a LDAP directory); a schema that is _not_ merely trivial or frivolous (e.g., a trivial schema might consist of a related set of attribute and object class definitions for holding information about the two possible binary bit values in a directory).

The schema listing procedures SHOULD be designed to enable:

a) verification that all properly formed schema listing requests are submitted by the contact person claiming to originate them

b) methods of ensuring that only properly-formed, high-quality directory schema are published in the schema listing repository

c) verifcation that requests to change the identity of the contact person for a listing originate from the listing authority contact or the contact person

d) coping with the situation in which the contact person and/or listing authority contact for a schema is no longer available or is unable to submit updates to the listing for which they hold update authority

For the initial release of the service, there is NO REQUIREMENT on any participant, user, or application to retain signature information as it applies to an entire schema listing request.

Fingerprints included with external content reference metadata elements MUST be retained and included in published listing request. Users of the schema listing service SHOULD verify that fingerprints of referenced content match corresponding fingerprints included with external references as a part of the published schema listing. If purported (included in the listing) and actual (computed by the user) fingerprints are different, users of the service SHOULD consider the possibility that the referenced content has changed since publication of the schema listing and that such a change could affect the way in which associated content can be used.

Referenced content is outside of the control of the schema listing service. A caveat explaining this concept MUST be included in the metadata of all published listings if external references are included in corresponding listing requests.

7.0 Acknowledgements

Leslie Daigle of Bunyip Information Systems and Chris Weider of Microsoft provided valuable comments on multiple versions of this document.

The engineering team for listing service requirements:

Chris Apple - AT&T Labs

Sanjay Jain - Oracle
Michael Mealling - NSI
John Strassner - Cisco
Sam Sun - CNRI
Mark Wahl - Critical Angle
Chris Weider - Microsoft

Paul Hoffman, for review and comment during his effort to implement the primary directory schema listing service platform.

The members of the ietf-schema-reg@imc.org mailing list.

8.0 References

[CHARSET] Internet Assigned Numbers Authority, "CHARACTER SETS", <URL:ftp://ftp.isi.edu/in-notes/iana/assignments/character-sets>.

[MIME] [RFC2045], [RFC2046], and [RFC2047].

[MIMEDIR] T. Howes, M. Smith, "A MIME Content-Type for Directory Information", INTERNET-DRAFT <draft-ietf-asid-mime-direct-04.txt>, July 1997.

[RFC821] J. Postel, "Simple Mail Transfer Protocol", RFC 821, August 1982.

[RFC959] J. Postel, J.K. Reynolds, "File Transfer Protocol", RFC 959, October 1985.

[RFC1630] T. Berners-Lee, "Universal Resource Identifiers in WWW", RFC 1630, June 1994.

[RFC1766] H. Alvestrand, "Tags for the Identification of Languages", RFC 1766, March 1995.

[RFC2015] M. Elkins, "MIME Security with Pretty Good Privacy (PGP)", RFC 2015, October 1996.

[RFC2044] F. Yergeau, "UTF-8, a transformation format of Unicode and ISO 10646", RFC 2044, October 1996.

[RFC2045] N. Freed, N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies", RFC 2045, November 1996.

[RFC2046] N. Freed & N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part Two: Media Types", RFC 2046, November 1996.

[RFC2047] K. Moore, "MIME (Multipurpose Internet Mail Extensions) Part Three: Message Header Extensions for Non-ASCII Text", RFC 2047, November 1996.

[RFC2068] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2068, January 1997.

[RFC2119] S. Bradner, "Key words for use in RFCs to Indicate Requirement Level", RFC 2119, March 1997.

[RFC2251] M. Wahl, T. Howes, S. Kille, "Lightweight Directory Access Protocol (Version 3)", RFC 2251, December 1997.

9.0 Author's Address

Chris Apple
AT&T Labs
600 - 700 Mountain Ave., Room 2F-165
Murray Hill, NJ 07974-0636
USA

E-Mail: capple@att.com
Phone: +1 908 582 2409
FAX: +1 908 582 3296

This INTERNET-DRAFT expires on October 21, 1998.