Directory Schema Listing File Names
<draft-ietf-schema-file-list-01.txt>

Status of this Memo

   This document is an Internet-Draft. Internet-Drafts are working
   documents of the Internet Engineering Task Force (IETF), its areas,
   and its working groups. Note that other groups may also distribute
   working documents as Internet-Drafts.

   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.''

   To learn the current status of any Internet-Draft, please check the
   ``1id-abstracts.txt'' listing contained in the Internet-Drafts Shadow
   Directories on ftp.is.co.za (Africa), nic.nordu.net (Europe),
   munnari.oz.au (Pacific Rim), or ftp.isi.edu (US West Coast).

Abstract

   This memo specifies a file name syntax for use by the primary listing
   repository operator of the directory schema listing service.

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 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. Schema listings will be stored in multiple files
   based on the different types of information associated with a
   listing: meta data and one or more syntax specifications.

1.1 Scope

A file name syntax specification intended for use during the initial
release of a directory schema listing service is inside 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.

2.0 File Name Syntax

All file names for listing meta data and listing content MUST comply with the following BNF [RFC822] grammar:

```
        file-name = sequence "." listversion "." type


        sequence = ("0" / "current") / NZDIGIT 0*<DIGIT>
                ; initialized to one (1) for first schema listing
                ; increments by one (1) for each successive schema
                ; listing name


        type = "meta-unit" /     ; these <type> values are defined
                "ldap" /         ; for the initial release of the
                "pak-ldap" /     ; schema listing service
                "whois++" /
                "pak-whois++" /  ; other <type> values may be defined
                "rwhois" /       ; according to community needs in
                "pak-rwhois" /   ; the future
                "whois" /
                "pak-whois"      ; this document will be updated or
                                 ; obsoleted when additional <type>
                                 ; values are defined


        listversion = 1*<DIGIT>
```

```
        NZDIGIT = <any DIGIT except "0" (0x30)>


        DIGIT = <any ASCII decimal digit (0x30 - 0x39)>
```


   Other possible values of the type component of a file name MAY be
   defined in the future to accomodate schema listings specified using
   [MIMEDIR] profiles other than those defined for containing LDAP
   [RFC2251], WHOIS++ [RFC1835], and RWHOIS [RFC1714] schema listing
   content.

3.0 Intended Use of File Names

   Schema writers, implementors, and users of the schema listing service
   SHOULD make use of the form of file names which includes descriptive
   alphabetic tokens as the value for the <type> part of a file name.

   Filenames MAY be specified as an OID by prepending the OID value used
   as a root for the service filename and swapping alphabetic tokens for
   their numeric equivalent according to the following table:

```
        Token            Number
        -----------      ------
        current            0
        meta-unit          0
        ldap               1
        pak-ldap           2
        whoispp            3
        pak-whoispp        4
        rwhois             5
        pak-rwhois         6
        whois              7
        pak-whois          8
```


   For the initial release of the service the behaviors documented in
   Section 4.0 for file retrieval based on file name will be supported.
   Schema writers, implementors, and users of the schema listing service
   SHOULD NOT rely on future support of such file retrieval behavior for
   the file name examples that are missing alphabetic tokens.

   The behavior of file retrieval based on file names containing
   alphabetic tokens MUST be preserved permanently by the schema listing
   repository operators.

4.0 Example File Names

Generally, file names will be of the following form:

    "sequence.listversion.type"

The 'sequence' part of a file name consists of a serial number
generated by the primary listing repository operator and is unique
within the context of the schema listing service.

When referring to a listing, a 'listversion' of "0" always represents
the most current version (the highest current listversion number)
published in the repository. Alternately, the token "current" may be
used to request the most current version of a listing file.
Otherwise, the listversion part of a file name represents the version
number of a listing within the context of the schema listing service.

The 'type' part of a file name consists of a token or number
representing a file type. This token is unique within the context of
the schema listing service and reflects the nature of file content.

If an OID is used to retrieve a file, the base OID used by the
primary listing repository operator MUST be prepended to the numeric
representation of the filename.

Retrieval of files will exhibit the following behavior for the
initial release of the service (NOTE: a value of 1 is used as the
base OID in these examples, the real base OID will be different):

1.12.4.0: returns schema unit metadata for version 4 of listing 12.

12.4.meta-unit: returns schema unit metadata for version 4 of listing
12

1.12.0.0: returns schema unit metadata for latest version of listing
12

12.current.meta-unit: returns schema unit metadata for latest version
of listing 12

1.12.4.1: returns ldap schema unit content for version 4 of listing
12

12.4.ldap: returns ldap schema unit content for version 4 of listing
12

1.12.0.1: returns ldap schema unit content for latest version of
listing 12

12.current.ldap: returns ldap schema unit content for latest version
of listing 12

1.13.2.2: returns metadata for version 4 of listing 12

13.2.pak-ldap: returns ldap schema pak metadata for version 2 of
listing 13

1.13.0.2: returns ldap schema pak metadata for latest version of
listing 13

13.current.pak-ldap: returns ldap schema pak metadata for latest
version of listing 13

5.0 Security Considerations

There are no known security concerns associated with the file name
syntax specified in this document.

6.0 Acknowledgements

Leslie Daigle of Bunyip Information Systems reviewed and provided
valuable comments on the syntax specification content in this
document.

The schema listing service engineering team:

    Chris Apple - AT&T Labs
    Sanjay Sain - Oracle
    Michael Mealling - NSI
    John Strassner - Cisco
    Sam Sun - CNRI
    Mark Wahl - Critical Angle
    Chris Weider - Microsoft

Paul Hoffman for review and comment resulting from his effort to
develop a platform for the initial release of the listing service.

7.0 References

[MIMEDIR] T. Howes, M. Smith, "A MIME Content-Type for Directory
Information", INTERNET-DRAFT <draft-ietf-asid-mime-direct-05.txt>,
November 1997.

[RFC822] D. Crocker, "Standard of the Format of ARPA-Internet Text
Messages", STD 11, RFC 822, August 1982.

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

[RFC1835] P. Deutsch, R. Schoultz, P. Faltstrom, C. Weider,
"Architecture of the WHOIS++ Service", RFC 1835, August, 1995.

[RFC1714] S. Williamson, M. Kosters,"Referral Whois Protocol
(RWhois)", RFC 1714, November 1994

[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.

8.0 **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.