



UC Curation Center / CDL

CAN: A Simple File System-Based Object Store

Rev. 0.10 – 2010-01-14

1 Introduction

A Content Access Node (CAN) is a simple file system convention for storing digital objects. It imposes minimal architectural and policy constraints while reserving a small set of file system names (directories and files) that place certain salient object store features, if available, in well-known locations within a single directory hierarchy that comprises the object store.

While CAN can be deployed usefully on its own, it was designed to interoperate cleanly with other independent, but related, specifications such as Pairtree [Pairtree], and Dflat [Dflat]. All three of these specifications start from the assumption that a file system is an appropriate storage abstraction for the effective management of digital objects. Assuming that these objects are arranged in a tree-like directory hierarchy, CAN specifies the organization and global properties of the tree; Pairtree specifies the form of the branches of the tree; and Dflat specifies the local structure of the leaves of the tree.

2 Notation

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 [RFC2119].

Note that some care **MUST** be taken in reading this specification so as not to misinterpret uses of the acronym “CAN” for an imperative indicating optionality. The term “MAY” will always be used to indicate such optionality.

Angle brackets and italics are used to indicate an arbitrary, as opposed to prescribed, file or directory name; for example, the arbitrarily-named file “<*file*>”. When referring to directory names and symbolic links in examples, the names are followed by a solidus (“/”), for example “*directory*/”; this suffix is indicative of type and is not part of the name. All directory and file names are case sensitive.

In examples of file system directory hierarchies, non-required directories or files are enclosed by square brackets (“[“ and “]”), a number sign (“#”) introduces an informative comment, and an ellipsis (“...”) indicates arbitrary repetition of the previous element.

Augmented Backus-Naur Form (ABNF) [RFC5234] is used to define the syntax of specific files required or recommended by this specification. Syntax rules names left undefined in this specification (for example, *ALPHA*) SHALL be interpreted as core ABNF names.

This specification is intended to be applicable, and implementable, in both Unix/Linux and Windows/DOS environments. Consequently, all uses of the term “directory” are interchangeable with “folder” without loss of meaning. Similarly, all uses of the solidus (“/”) as a directory path separator are interchangeable with a reverse solidus (“\”).

The complete set of CAN conventions is provided in Appendix A.



3 Content Access Node

A Content Access Node, or CAN, is a file system hierarchy that comprises a store for digital objects. The directory that is the structural root of a CAN hierarchy is known as the *CAN home directory*. The CAN specification reserves a few key file system names within the home directory and its descendent sub-directories, but it does not dictate how the home directory itself is named, as that name is not visible from inside the CAN itself.

A CAN looks like this:

```

<can_home>/
    [ 0=can_0.10 ]      # CAN home directory
    [ admin/ ]         # CAN Namaste signature
    [ can-info.txt ]   # administrative declarations
    [ lock.txt ]       # CAN properties file
    [ log/ ]           # write lock
    [ store/ ]         # log directory
                        # object store

```

A CAN home directory SHOULD contain a file named “0=can_0.10” that is its Namaste [Namaste] signature. A Namaste signature plays the same role for a directory that a magic number plays for a file.

The home directory MAY contain a sub-directory named “admin” that holds administrative declarations about the CAN itself, as opposed to the objects managed in it.

The CAN home directory SHOULD contain a file named “can-info.txt” that defines the global properties of the CAN itself.

The home directory MAY contain a file named “lock.txt” that indicates a write operation is in-process and that the CAN may be in a temporarily unknown, inconsistent, or incomplete state. If present, the lock file MUST conform to the syntax, semantics, and normative obligations defined by the LockIt specification [LockIt].

The home directory MAY contain a sub-directory named “log” holding log information about the use of the CAN and the objects managed in it.

The home directory MUST contain a directory named “store”, which is the root of the hierarchical object store.

3.1 Namaste signature file (0=can_<version>)

The RECOMMENDED Namaste signature file “0=can_0.10” self-identifies a directory as a CAN home directory. The “<version>” portion of the file name asserts the version of the CAN specification to which the directory conforms. If present, the contents of the file MUST contain the specification name and version, separated by a forward slash, followed by a CR, CRLF, or LF end-of-line (EOL) marker.

```
CAN/0.10
```



Note that this value is duplicative of the “nodeScheme” property of the global properties file “can-info.txt”. If both are present, the global properties file is considered authoritative.

3.2 Administrative directory (admin/)

The OPTIONAL “admin” directory holds administrative declarations associated with the CAN itself, as opposed to the objects managed in it.

```
admin/
```

3.3 Properties file (can-info.txt)

The RECOMMENDED properties file “can-info.txt” defines the global properties of that CAN itself, as opposed to the objects managed in it. These properties are expressed in terms of ANVL [ANVL] name/value pairs, for example:

```
name: Primary
identifier: 12
description: Primary UC3 storage node
nodeScheme: CAN/0.9
branchScheme: Pairtree/0.1
leafScheme: Dflat/0.18
classScheme: CLOP/0.3
mediaType: magnetic-disk
accessMode: on-line
verifyOnRead: true
verifyOnWrite: true
baseURI: http://can01.cdlib.org/
supportURI: mailto:merritt-support@ucop.edu
```

The “name” property indicates the name of the CAN, which SHOULD be assigned to be unique within the administrative regime in which the CAN operates.

The “identifier” property is an identifier for the CAN that MUST be assigned to be unique within the administrative regime in which the CAN operates.

The “description” property provides a short textual description of the CAN.

The “nodeScheme” property indicates the version of the CAN specification to which the CAN conforms. Note that this value is duplicative to the information provided by the OPTIONAL Namaste tag. If both are present, the properties file is considered authoritative.

The “branchScheme” property indicates the specification name and version of the convention used for the branches of the object store hierarchy.

The “leafScheme” property indicates the specification name and version of the convention used for the leaves of the object store hierarchy.

The “classScheme” property indicates the specification name and version of the convention used for managing the administrative properties of managed objects.



The “mediaType” property indicates the storage technology underlying the CAN: magnetic disk, magnetic tape, optical disk, or solid-state.

The “accessMode” property indicates level of availability of stored data: on-line, near-line, or off-line.

The “verifyOnRead” property indicates whether a message digest verification SHOULD be performed prior to responding to a request to retrieve an object, version, or file (but *not* their states) from the CAN. The “verifyOnWrite” property indicates whether a message digest verification SHOULD be performed following the addition of a new object version to the CAN. A directive to verify on read or write MAY be ignored if the storage media underlying the CAN is not amenable to such verification, as may be case, for example, with tape storage.

The “baseURI” property indicates the base URI used for CAN method invocations.

The “supportURI” property indicates the URI for user support.

3.4 Log directory (log/)

The OPTIONAL “log” directory log holds log information about the use of the CAN and the objects managed in it.

```
log/
  [ last-activity.txt ]
  [ lock.txt ]
  [ log-<year><month><day>.txt ]
  [ summary-stats.txt ]
```

The directory SHOULD contain a log file named “last-activity.txt” containing the date/timestamp, in fully-qualified W3C form [DateTime], of the last instance of various activities performed on the objects in the CAN and a unique identifier of the activity’s process (such as a process number or thread identifier), for example:

```
lastAddVersion: 2008-11-23T08:17:53-08:00
lastObjectDelete: 2008-05-03T13:04:12-08:00
lastFixity: 2008-12-22T23:00:10-08:00
```

The sub-directory MAY contain a file named “lock.txt” that conforms to the syntax, semantics, and normative obligations defined in by the LockIt specification.

The sub-directory MAY contain a file named in the form “log-<year><month><day>.txt”, that holds log information for the CAN for the specified month. Log files for previous months MAY be compressed using GZIP [RFC1952] or ZIP [ZIP], resulting in file names of the form, “log-<year><month><day>.txt.gz” and “log-<year><month><day>.txt.zip”, respectively.

The administrative directory MAY contain a file named “summary-stats.txt” holding summary statistics about the CAN expressed as ANVL name/value pairs, for example:



```

numObjects: 18302
numVersions: 27551
numFiles: 405833
totalSize: 730415172

```

The “numObjects”, “numVersions”, and “numFiles” properties indicate the number of objects, versions, and files, respectively, managed in the CAN.

The “totalSize” property indicates the total size of the CAN, in bytes.

3.5 Object store directory (store/)

The OPTIONAL object store directory “store” is the structural root of the file system hierarchy in which digital objects are managed. The local structure of the “store” directory and its descendent sub-directories is subject to other specifications, such as Pairtree [Pairtee] and Dflat [Dflat].

```
store/
```

The particular schemes used to define the branches and leaves of the CAN’s object store MUST be declared in the signature file.

4 Implementation

The CAN specification was developed with the intention that it could be implemented on top of any file system that supports hierarchical directory structures and arbitrary directory and file names, such as POSIX [POSIX].

5 Security Considerations

CAN poses no direct risk to computers or networks. As a file system convention, CAN is capable of holding files that might contain malicious executable content, but it is no more vulnerable in this regard than any file system.

Appendix A: Complete CAN Conventions

The overall file system structure of a CAN is:

```

<can_home>/
  [ 0=can <version> ]
  [ admin/ ]
  [ can-info.txt ]
  [ lock.txt ]
  [ log/
    [ last-activity.txt ]
    [ lock.txt ]
    [ log-<year><month><day>.txt
    [ log-<year><month><day>.txt.gz ]
    [ log-<year><month><day>.txt.zip ]]
  [ summary-stats.txt ] ]

```



store/

The production rule for the CAN Namaste signature file “0=can_<version>” is:

```

<cansig>      = “CAN/” <version> EOL
<version>    = NONNEG 0* (“.” 1*DIGIT)
NONNEG       = “0” / (1*POSDIG 0*DIGIT)
POSDIG       = “1” / “2” / “3” / “4” / “5” / “6” / “7” / “8” / “9”
EOL          = CR / CRLF / LF

```

The generic production rules for all ANVL-based files are:

```

<file>       = 1*<line>
<line>       = <name> “:” 1*WSP <value> EOL
<name>       = 1*VCHAR
<value>      = 1*VCHAR

```

Matching of all ANVL property names **MUST** be performed on a case-insensitive basis. More specific rules **MAY** be defined by each such ANVL-based file. It is **RECOMMENDED** that only a single white space character (WSP) is used to separate the name/value pairs in these files.

The production rule for the last activity log file “last-activity.txt” is:

```

<activity>   = 1*(<activities> 1*WSP <date-time> EOL)
<activities> = <add> / <object> / <version> / <fixity>
<add>        = “lastAddVersion:”
<object>     = “lastDeleteObject:”
<version>    = “lastDeleteVersion:”
<fixity>     = “lastFixity:” 1*WSP <date-time>
<date-time>  = <date> “T”<time>
<date>       = <year> “-“ <month> “-“ <day>
<year>       = 4DIGIT
<month>      = 2DIGIT ; normal constraints apply, 01-12
<day>        = 2DIGIT ; normal constraints apply, 01-31
<time>       = <hour> “:” <minute> “:” <second> <zzzz>
<hour>       = 2DIGIT ; normal constraints apply, 00-23
<minute>     = 2DIGIT ; normal constraints apply, 00-59
<second>     = 2DIGIT ; normal constraints apply, 00-59
<zzzz>      = “Z” / (“+” | “-” ) <hour> “:” <minute>
<process>   = 1*VCHAR

```

It is **RECOMMENDED** that only a single white space character (WSP) is used to demarcate the date/timestamp and process identifier.

The production rules for the summary statistics file “summary-stats.txt” are:

```

<stats>      = 1*(<stat> EOL)
<stat>       = <objects> / <versions> / <files> / <size>
<objects>    = “numObjects:” 1*WSP NONNEG
<versions>   = “numVersions:” 1*WSP NONNEG

```



```

<files>      = "numFiles:" 1*WSP NONNEG
<size>       = "totalSize:" 1*WSP NONNEG

```

The production rules for the CAN global properties file “can-info.txt” are:

```

<properties> = 1*( <property> EOL )
<property>   = <canname> / <identifier> / <description> / <node> /
              <branch> / <leaf> / <class> / <media> / <mode> /
              <read> / <write> / <base> / <support>

<canname>    = "name:" 1*WSP 1*VCHAR
<identifier> = "identifier:" 1*WSP 1*VCHAR
<description> = "description:" 1*WSP 1*VCHAR
<node>       = "nodeScheme:" 1*WSP <scheme>
<branch>     = "branchScheme:" 1*WSP <scheme>
<leaf>       = "leafScheme:" 1*WSP <scheme>
<class>      = "classScheme:" 1*WSP <scheme>
<scheme>     = <name> "/" <version>
<name>       = 1*VCHAR
<media>      = "mediaType:" 1*WSP ("magnetic-disk" /
                                  "magnetic-tape" /
                                  "optical-disk" / "solid-state")
<mode>       = "accessMode" 1*WSP ("on-line" / "near-line" /
                                  "off-line")
<read>       = "verifyOnRead:" 1*WSP ("true" / "false")
<write>      = "verifyOnWrite:" 1*WSP ("true" / "false")
<base>       = "baseURI:" 1*WSP <rfc-3986-compliant-uri>
<support>    = "supportURI:" 1*WSP <rfc-3986-compliant-uri>

```

References

- [ANVL] J. Kunze, B. Kahle, J. Masanes, and G. Mohr, *A Name-Value Language (ANVL)*, Internet draft, February 14, 2005 <<http://www.ietf.org/internet-draft/draft-kunze-anvl-01.txt>>.
- [CLOP] *CLOP: A Class-based System for Managing Object Properties*, rev. 0.3, April 3, 2009.
- [DateTime] Misha Wolf and Charles Wicksteed, *Date and Time Formats*, September 15, 1997 <<http://www.w3.org/TR/NOTE-datetime>>.
- [Dflat] *DFlat: A Simple File System Convention for Object Storage*, rev. 0.10, April 4, 2009.
- [LockIt] UC3, *LockIt: A Simple File-based Convention for Resource Locking*, 2009.
- [Pairtree] J. Kunze, M. Haye, E. Hetzner, M. Reyes, and C. Snively, *Pairtrees for Object Storage (V0.1)*, Internet draft, November 25, 2008 <<http://www.ietf.org/internet-drafts/draft-kunze-pairtree-01.txt>>.
- [POSIX] ISO/IEC 9945:2003/IEEE Std 1003.1, *Information technology – Portable operating system interface (POSIX) – Part 2: System interface*.
- [RFC1952] P. Deutsch, *GZIP File Format Specification 4.3*, RFC 1952, May 1996 <<http://www.ietf.org/rfc/rfc1952.txt>>.
- [RFC2119] S. Bradner, *Key Words for Use in RFCs to Indicate Requirement Levels*, BCP 14, RFC 2119, March 1997 <<http://www.ietf.org/rfc/rfc2119.txt>>.
- [RFC5234] D. Crocker (ed.) and P. Overell, *Augmented BNF for Syntax Specifications: ABNF*, STD 68, RFC

UC3



5234, January 2008 <<http://www.ietf.org/rfc/rfc5234.txt>>.

[ZIP]

PKWARE, Inc., *.ZIP File Format Specification*, Version 6.3.2, September 28, 2007
<<http://www.pkware.com/documents/casestudies/APPNOTE.TXT>>.