SDC
STORAGE DEVELOPER CONFERENCE
SNIA ■ SANTA CLARA, 2014

Introducing CDMI 1.1

David Slik
NetApp

Introducing CDMI 1.1
Table of Contents

- A Quick Review of CDMI
- CDMI Adoption
- CDMI 1.1 - Document Organization
- CDMI 1.1 - Clarifications
- CDMI 1.1 - New Functionality
- CDMI 1.1 - Extensions
- How to transition to CDMI 1.1

A Quick Review of CDMI
History of the standard

- Technical Working Group founded in 2009
  - Published TWG Charter and Use Cases
- CDMI Timeline:
  - 2010 - CDMI 1.0 Technical Architecture
  - 2011 - CDMI 1.0.1 Errata
  - 2012 - CDMI 1.0.2 Errata
  - 2013 - Adopted as ISO/IEC 17826
  - 2014 - CDMI 1.1.0 Technical Architecture

A Quick Review of CDMI
What is CDMI?

- CDMI standardizes the following:
  - How is data stored in the cloud
    - Data objects, queues, metadata
  - How is data stored in the cloud organized
    - Containers, object IDs, query, snapshots
  - How is data stored in the cloud transferred
    - Client to cloud, cloud-to-cloud, exports, serialization, notifications
  - How is data stored in the cloud secured

A Quick Review of CDMI
What is CDMI?

- CDMI works with existing protocols:
  - File: NFS, CIFS, LTFS, etc.
  - Block: iSCSI, VMDKs, etc.
  - Object: S3, Swift, etc.
- CDMI fills many gaps in the above protocols
  - Unified storage management
  - Global and hierarchical namespaces
  - Query, notification and workflow

A Quick Review of CDMI
CDMI Adoption

- 22 publically announced CDMI servers
  - Major vendors (NetApp, DDN,etc)
  - Startup companies
  - Open source projects
- Widespread adoption in government
  - USA DoD, UK, Italy, etc.
- Supported in OpenStack Swift

CDMI 1.1 Changes
Document Organization

- To improve readability, the CDMI specification has been split into five parts:

Preamble	pp 1 – 25	25
Basic Cloud Storage	pp 27 – 41	15
CDMI Core	pp 42 – 108	67
CDMI Advanced	pp 109 – 235	127
CDMI Annexes	pp 235 – 256	22

CDMI 1.1 Changes
Document Organization

- Section 1 - **Preamble**
  - References and Terms
  - Provides an overview of cloud storage
  - Provides an overview of the CDMI standard
  - Defines the CDMI model for cloud storage and metadata
  - Introduces general CDMI concepts:
    - Object Types, Object IDs, Time, use of HTTP, Security, Backwards Compatibility

CDMI 1.1 Changes
Document Organization

- Section 2 – **Basic Cloud Storage**
  - Formerly “Non-CDMI” operations
  - Defines basic RESTful operations for data objects and containers
  - Compatible subset of CDMI, S3, Swift, etc.
    - Provides guidance for multi-protocol support
    - Also see the Header-based Metadata Extension
  - Minimal baseline for cloud storage

CDMI 1.1 Changes
Document Organization

- Section 3 – **CDMI Core**
  - Defines operations for CDMI Data Objects
  - Defines operations for CDMI Containers
  - Minimal baseline for CDMI-based systems
    - Containers optional
    - By ID only objects optional
    - Etc.

CDMI 1.1 Changes
Document Organization

- Section 4 – **CDMI Advanced**
  - Defines operations for CDMI Domain Objects
  - Defines operations for CDMI Queue Objects
  - Defines operations for CDMI Capabilities
  - Advanced Features of CDMI
    - Exports, Snapshots, Serialization, Metadata, Retention and Hold, Logging, Notifications, Query

CDMI 1.1 Changes
Document Organization

- Section 5 – **CDMI Annexes**
  - Extensions to the CDMI standard implemented by at least one vendor
  - Includes:
    - Summary Metadata for Bandwidth
    - Expiring ACLs
    - Group Storage System Metadata
    - Versioning

CDMI 1.1 Changes – Co-existence
Clarification - #904, #907, #918, #919, #931

- Clause 6 and 7 reworked
  - Clarifies that Non-CDMI operations represent basic RESTful HTTP operations that are consistent with most object storage protocols
- cdmi_authentication_methods
  - Text added explaining how S3, Keystone, etc work with CDMI

CDMI 1.1 Changes – Copy/Move
Clarification - #440, #504, #815, #847

- Copy and Move
  - Copying data to an existing or new object has been clarified.
    - Behaviours are documented when fields in the source URI are omitted or specified
  - Copying between and from queues have been clarified
    - Added cdmi_copy_dataobject_from_queue
  - Domain move capability missing

CDMI 1.1 Changes – Container Fields Clarification - #476

- childrenrange/children now optional on a container create
  - Eliminates an edge case where copying or deserializing a container could result in a large listing of children being returned

- **Servers**: No change required
- **Clients**: No longer depend on these fields being returned

CDMI 1.1 Changes – Container Fields
Clarification - #651

- Clarified contents of parentURI and parentID for root containers
  - Multiple vendors had chosen different approaches
  - Selected approach was “best compromise”

- **Servers**: Changes may be required
- **Clients**: Changes if depend on these fields

CDMI 1.1 Changes – Metadata
Clarification - #517, #566, #833

- Metadata updates, additions and deletions have been clarified in a new section: 16.6
  - Examples are already present in 1.0.2
  - Additional examples added
- Mutability of storage and data system metadata
- Default values of storage system metadata

CDMI 1.1 Changes – HTTP Headers
Clarification - #536

- The Location header must be an absolute URI
  - A close reading of RFC 2616 should have already confirmed this for implementers.

CDMI 1.1 Changes – ACLs
Clarification - #812, #817, #890

- Clarified field results when ACL deny access to specific parts of objects
  - CDMI 1.0.2 approach was viewed as standard and intuitive, but needed to be specified normatively.
- Now indicates which status code to return
- Clarified that hex and string forms are allowed
- **Servers**: Changes may be required
- **Clients**: Changes if depend on these fields

CDMI 1.1 Changes – Retention and Hold Clarification - #894

- Additional examples added

CDMI 1.1 Changes – Scopes
New Functionality - #483, #508, #902

- CDMI Scopes have been enhanced to handle JSON arrays.
  - Required for querying against ACLs
- AND statements now use JSON arrays to avoid the use of duplicate keys
- Numeric query broken out

- **Servers**: Add new functionality if supported
- **Clients**: Changes required for numeric matching

CDMI 1.1 Changes – Queues
New Functionality - #515

- CDMI Queues now allow deletion by range
  - Allows idempotent deletes

- **Servers**: Add new functionality if supported
- **Clients**: No changes required

CDMI 1.1 Changes – Data Object Updates
New Functionality - #881

- Update range spec was unnecessarily restrictive
  - Now allows appends and sparse updates

- **Servers**: Add new functionality if supported
- **Clients**: No changes required

CDMI 1.1 Changes – Extensions
Multi-part MIME

- Allows the data object value to be sent as a separate MIME part, without requiring encoding
  - Improves efficiency of binary transfers

CDMI 1.1 Changes – Extensions
Domain Authentication Methods

- Allows a client to discover which authentication methods a server supports
- Allows an administrator to restrict which authentication methods are supported for a given domain

CDMI 1.1 Changes – Extensions
Group Storage System Metadata

- Allows objects to have a specified “owner” that ACLs can refer to
- Allows broader compatibility with NFS and CIFS ACLs

CDMI 1.1 Changes – Extensions
Domain Summary Metadata

- Allows domain summaries to include information about bandwidth consumed
  - Network bytes
  - Reads & Writes
  - Public and Private (Internal) I/O

CDMI 1.1 Changes – Extensions
Expiring ACEs

- Allows ACEs to have an expiration time, when the ACE will no longer be evaluated as part of the ACL
  - Allows time-limited access
  - Allows content to become public after a period of time

CDMI 1.1 Changes – Extensions
Versioning

- Allows data objects to retain historical versions as changes are made
  - Historical versions are accessed by ID
  - Historical versions are enumerated as a tree
  - Enables multi-writer conflict resolution
  - Simplifies federation, distributed storage and disconnected operation

Guidance for CDMI 1.1 Adoption

- Review the CDMI 1.1 spec and errata, which is posted on the SNIA public review site

  http://www.snia.org/tech_activities/publicreview
- Most vendors will be able to support 1.0.2 and 1.1 clients without changes.
  - Review error handling, data object update and container field behaviours
  - Consider adding support for Multi-part MIME

Guidance for CDMI 1.1 Adoption

Come visit us at the

Cloud Plugfest

Held here at SDC 2014

Introducing CDMI 1.1

Questions and Answers

Contact Info:
dslik@netapp.com