



Cloud Data Management Interface Extension: Server Side Partial-Value Copy

Version 1.0c

"Publication of this Working Draft for review and comment has been approved by the Cloud Storage Technical Working Group. This draft represents a "best effort" attempt by the Cloud Storage Technical Working Group to reach preliminary consensus, and it may be updated, replaced, or made obsolete at any time. This document should not be used as reference material or cited as other than a 'work in progress.' Suggestion for revision should be directed to <http://snia.org/feedback>."

Working Draft

Revision History

Date	Version	By	Comments
2011-12-05	1.0a	David Slik, NetApp, Inc.	Original draft
2012-01-16	1.0b	Marie McMinn	Updates to include standard SNIA front matter and technical edit
2012-01-26	1.0c	David Slik, NetApp, Inc.	Updated to fix typos found by Kevin, added updates for clause 9.9 and queues.

The SNIA hereby grants permission for individuals to use this document for personal use only, and for corporations and other business entities to use this document for internal use only (including internal copying, distribution, and display) provided that:

- Any text, diagram, chart, table, or definition reproduced shall be reproduced in its entirety with no alteration, and,
- Any document, printed or electronic, in which material from this document (or any portion hereof) is reproduced shall acknowledge the SNIA copyright on that material, and shall credit the SNIA for granting permission for its reuse.

Other than as explicitly provided above, you may not make any commercial use of this document, sell any excerpt or this entire document, or distribute this document to third parties. All rights not explicitly granted are expressly reserved to SNIA.

Permission to use this document for purposes other than those enumerated above may be requested by e-mailing tcmd@snia.org. Please include the identity of the requesting individual and/or company and a brief description of the purpose, nature, and scope of the requested use.

Copyright © 2012 Storage Networking Industry Association.

CDMI Extension: Server-Side Partial Value Copy

Overview

CDMI supports the ability to create a new data object that contains a copy of the value of an existing data object and the ability to update the value of an existing data object using the value of a second existing data object. However, the 1.0.1 spec does not document the behavior when either the source URI or the destination URI is a range of a data object value. This extension defines the behaviors that allow clients to perform scatter-gather operations, where segments of a data object can be extracted (scattered) to the server into new data objects and multiple data objects can be combined (gathered) into a single data object without having to retrieve the data from the client and then send it back to the server.

Server-side partial value copy is based on the existing server-side value copy functionality in the 1.0.1 spec and leverages the existing fields defined for copy.

Modifications to the CDMI 1.0.1 spec:

1) Insert into Clause 8.2.3 - "Capabilities", add:

- If the new data object is a copy of a range of the value of an existing data object, support for the ability to copy the range is indicated by the presence of the "cdmi_copy_dataobject_range" capability in the parent container.

2) Insert into Clause 8.2.5 - "Request Message Body", Table 6 - "Request Message Body - Create a Data Object using CDMI Content Type", "copy" row:

Replace:

"URI of a CDMI data object or queue that shall be copied into the new data object"

with:

"URI of a CDMI data object, data object value range, or queue that shall be copied into the new data object"

3) Insert into Clause 8.2.9 - "Examples", add:

EXAMPLE 3 PUT to the container URI to copy an existing data object:

```
PUT /MyContainer/DestinationDataObject.txt HTTP/1.1
Host: cloud.example.com
Accept: application/cdmi-object
Content-Type: application/cdmi-object
X-CDMI-Specification-Version: 1.0.1

{
  "copy" : "/MyContainer/SourceDataObject.txt"
}
```

The following shows the response.

```
HTTP/1.1 201 Created
Content-Type: application/cdmi-object
X-CDMI-Specification-Version: 1.0.1
```

```
{
  "objectType": "application/cdmi-object",
  "objectID": "0000706D0010374085EF1A5C7018D774",
  "objectName": "DestinationDataObject.txt",
  "parentURI": "/MyContainer/",
  "parentID": "00007E7F00102E230ED82694DAA975D2",
  "domainURI": "/cdmi_domains/MyDomain/",
  "capabilitiesURI": "/cdmi_capabilities/dataobject/",
  "completionStatus": "Complete",
  "mimetype": "text/plain",
  "metadata": {
    "cdmi_size": "37"
  }
}
```

EXAMPLE 4 PUT to the container URI to copy a range of an existing data object:

```
PUT /MyContainer/DestinationDataObject.txt HTTP/1.1
Host: cloud.example.com
Accept: application/cdmi-object
Content-Type: application/cdmi-object
X-CDMI-Specification-Version: 1.0.1

{
  "copy" : "/MyContainer/SourceDataObject.txt?value:10-19"
}
```

The following shows the response.

```
HTTP/1.1 201 Created
Content-Type: application/cdmi-object
X-CDMI-Specification-Version: 1.0.1

{
  "objectType": "application/cdmi-object",
  "objectID": "0000706D0010374085EF1A5C7018D774",
  "objectName": "DestinationDataObject.txt",
  "parentURI": "/MyContainer/",
  "parentID": "00007E7F00102E230ED82694DAA975D2",
  "domainURI": "/cdmi_domains/MyDomain/",
  "capabilitiesURI": "/cdmi_capabilities/dataobject/",
  "completionStatus": "Complete",
  "mimetype": "text/plain",
  "metadata": {
    "cdmi_size": "10"
  }
}
```

4) Insert into Clause 8.6.2 - "Capabilities", add:

- Support for the ability to modify the value of an existing data object by copying a specified byte range from another data object is indicated by the presence of the "cdmi_copy_dataobject_value_range" capability in the specified object.

5) Insert into Clause 8.6.4 - "Request Message Body", Table 20 - "Request Message Body - Update a Data Object using CDMI Content Type", "copy" row:

Replace:

"URI of a CDMI data object or queue that shall be copied into the existing data object."

with:

"URI of a CDMI data object, data object value range, or queue that shall be copied into the existing data object.

When updating a value range, if the source data object value, data object value range, or queue value is not equal in length to the value specified for the existing object, the server shall return a HTTP status code of 400 Bad Request.

When updating a value range, if the update value range extends beyond the end of the existing value range, data shall be appended to the existing data object."

6) Insert into Clause 8.6.8 - "Examples", add:

EXAMPLE 7 PUT to the data object URI to copy an existing data object over a second existing data object:

```
PUT /MyContainer/DestinationDataObject.txt HTTP/1.1
Host: cloud.example.com
Content-Type: application/cdm-object
X-CDMI-Specification-Version: 1.0.1

{
  "copy" : "/MyContainer/SourceDataObject.txt"
}
```

The following shows the response.

```
HTTP/1.1 204 No Content
```

EXAMPLE 8 PUT to the data object URI to copy a range of a first existing data object over a range in a second existing data object:

```
PUT /MyContainer/DestinationDataObject.txt?value:0-9 HTTP/1.1
Host: cloud.example.com
Content-Type: application/cdm-object
X-CDMI-Specification-Version: 1.0.1

{
  "copy" : "/MyContainer/SourceDataObject.txt?value:10-19"
}
```

The following shows the response.

```
HTTP/1.1 204 No Content
```

7) Insert into Clause 9.9.3 - "Capabilities", add below bullet for "cdmi_copy_dataobject_by_ID":

- If the new data object being created in "/cdmi_objectid/" is a range copy of an existing data object, support for the ability to copy is indicated by the presence of the "cdmi_copy_dataobject_range_to_ID" system capability.

8) Insert into Clause 9.9.3 - "Capabilities", add below bullet for "cdmi_copy_dataobject":

- If the new data object is a copy of a range of the value of an existing data object, support for the ability to copy the range is indicated by the presence of the "cdmi_copy_dataobject_range" capability in the parent container object.

9) Insert into Clause 9.9.5 - "Request Message Body", Table 50 - "Request Message Body - Create a New Data Object using CDMI Content Type", "copy" row:

Replace:

"URI of a CDMI data object or queue that shall be copied into the new data object"

with:

"URI of a CDMI data object, data object value range, or queue that shall be copied into the new data object"

10) Insert into Clause 9.9.9 - "Examples", add:

EXAMPLE 3 POST to the object ID URI to copy an existing data object:

```
POST /MyContainer/ HTTP/1.1
Host: cloud.example.com
Accept: application/cdmi-object
Content-Type: application/cdmi-object
X-CDMI-Specification-Version: 1.0.1

{
  "copy" : "/MyContainer/SourceDataObject.txt"
}
```

The following shows the response.

```
HTTP/1.1 201 Created
Content-Type: application/cdmi-object
X-CDMI-Specification-Version: 1.0.1
Location: http://cloud.example.com/MyContainer/00007ED90010AB253635A213DD5EEE6B

{
  "objectType": "application/cdmi-object",
  "objectID": "00007ED90010AB253635A213DD5EEE6B",
  "objectName": "00007ED90010AB253635A213DD5EEE6B",
  "parentURI": "/MyContainer/",
  "parentID": "00007E7F00102E230ED82694DAA975D2",
  "domainURI": "/cdmi_domains/MyDomain/",
  "capabilitiesURI": "/cdmi_capabilities/dataobject/",
  "completionStatus": "Complete",
  "mimetype": "text/plain",
  "metadata": {
    "cdmi_size": "37"
  }
}
```

EXAMPLE 4 POST to the object ID URI to copy a range of an existing data object:

```
POST /MyContainer/ HTTP/1.1
Host: cloud.example.com
Accept: application/cdmi-object
Content-Type: application/cdmi-object
X-CDMI-Specification-Version: 1.0.1

{
  "copy" : "/MyContainer/SourceDataObject.txt?value:10-19"
}
```

The following shows the response.

```

HTTP/1.1 201 Created
Content-Type: application/cdm-object
X-CDMI-Specification-Version: 1.0.1
Location: http://cloud.example.com/MyContainer/00007ED90010E92F4313580A56A6739C

```

```

{
  "objectType": "application/cdm-object",
  "objectID": "00007ED90010E92F4313580A56A6739C",
  "objectName": "00007ED90010E92F4313580A56A6739C",
  "parentURI": "/MyContainer/",
  "parentID": "00007E7F00102E230ED82694DAA975D2",
  "domainURI": "/cdmi_domains/MyDomain/",
  "capabilitiesURI": "/cdmi_capabilities/dataobject/",
  "completionStatus": "Complete",
  "mimetype": "text/plain",
  "metadata": {
    "cdmi_size": "10"
  }
}

```

11) Insert into Clause 11.6.2 - "Capabilities", add at end:

- If the enqueued value is a copy of a range of the value of an existing data object, support for the ability to copy the range is indicated by the presence of the "cdmi_copy_dataobject_range" capability in the specified queue object.

12) Insert into Clause 11.6.4 - "Request Message Body", Table 98 - "Request Message Body - Enqueue a New Queue Value using CDMI Content Type", "copy" row:

Replace:

"If a URI to a data object is provided, the value, mimetype, and valuetransferencoding from the data object is used to enqueue the new item into the queue object, and the data object is atomically deleted."

with:

"If a URI to a data object is provided, the value or value range, mimetype, and valuetransferencoding from the data object is used to enqueue the new item into the queue object"

13) Insert into Clause 12.1.1 - "Cloud Storage System-Wide Capabilities", Table 105 - "System-Wide Capabilities":

Capability Name	Type	Description
cdmi_copy_dataobject_range_to_ID	JSON String	If present and "true", this capability indicates that the system allows the copying a range of an existing data object when creating a new data object by ID via POST to "/cdmi_objectid/".

14) Insert into Clause 12.1.4 - "Data Object Capabilities", Table 105 - "Capabilities for Data Objects":

Capability Name	Type	Description
cdmi_copy_dataobject_value_range	JSON String	If present and "true", this capability indicates that the object's value may be modified with a byte range copied from another data object.

15) Insert into Clause 12.1.5 - "Container Capabilities", Table 106 - "Capabilities for Containers":

Capability Name	Type	Description
cdmi_copy_dataobject_value_range	JSON String	If present and "true", this capability indicates that the object's value may be created with a byte range copied from another data object.