



# CDMI Extension: Partial Upload

## **Version 1.1a**

"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
08-28-2014	1.1a	Marie McMinn, Cloud TWG Editor	New version of Partial Load extension created for the CDMI 1.1.0 revision of the standard. Derived from the 1.0e draft.

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](mailto: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 © 2014 Storage Networking Industry Association.

# Partial Upload Extension

## Overview

CDMI 1.0.2 provides limited support for partial uploads. The following extension to the X-CDMI-Partial header is proposed to support managing conflicting uploads and specifying completion conditions for parallel uploads:

```
X-CDMI-Partial: [ true | false ] | [ upload-id=<upload-id> [ ; [
count=<integer> | range=<byte-range> ] ] [ ; replace= { true | false } ] ]
```

Partial uploads by multiple concurrent clients are accommodated by specifying a unique "upload-id" for each set of partial uploads.

A condition is associated with an upload id to indicate when the partial upload will be considered complete. The condition may be associated with an upload id at any time. Once a condition is associated with an upload id, specifying a different condition is considered an error.

A count condition handles the scenario when a partial upload is to be completed when a specific number of partial uploads are received by the server.

A range condition handles the scenario when a partial upload is to be completed when the server receives a specific byte range.

Replace indicates if the partial uploads replace the entire object or just the specified byte ranges.

When partial uploads associated with an upload id have not completed within a given time, the upload id will time out, and the partial uploads associated with that upload id will be discarded by the server.

## Changes to CDMI 1.1.0:

### 1) Add a new section 5.17 CDMI Partial Header

CDMI defines a custom "X-CDMI-Partial" header that indicates when a partial upload is being performed. The value of this header is formatted according to the following BNF:

```
[ true | false ] | [ upload-id=<upload-id> [ ; [ count=<integer> |
range=<byte-range> ] ] [ ; replace= { true | false } ] ]
```

- "true" – Indicates that the newly created object is part of a series of uploads and the value has not yet been fully populated. These uploads are considered to have a null upload-id.
- "false" (or header not present) – Indicates that the set of uploads associated with a null upload-id shall be considered complete, and the object shall be updated.
- "upload-id" – Indicates that requests with the same upload-id are part of the set of partial uploads.
- "count" – An integer value greater than zero that indicates that when exactly this many partial uploads with the same upload ID are received, the set of uploads is considered complete, and the object shall be updated. If the number of received uploads is larger than the specified count, an HTTP status code of 400 *Bad Request* shall be returned.

- "range" – A byte range as specified in section 14.35.1 of RFC 2616 that indicates that when a given byte range of partial uploads with the same upload-id are received, the set of uploads is considered complete, and the object shall be updated.
- "replace" – When multiple ranges are sent as part of a partial upload, if this flag has the value "true", this indicates that the entire object shall be replaced by the set of ranged uploads (with any range gaps zero-filled). If it has the value "false" or not present, the set of uploads shall be used to update the object, not replace it.

For a given upload-id, if a condition or replace flag is received that is different from a previously received condition or replace flag, an HTTP status code of 400 *Bad Request* shall be returned.

For a given upload-id, if a partial upload is received with an content-range exactly the same as a previously received content-range, the value associated with that range shall be replaced with the newer value. This allows partial uploads to be retried without error. In this case, the count shall not change.

For a given upload-id, if a partial upload is received with a content-range that overlaps a previously received content-range and is not exactly the same as a previously received content-range, an HTTP status code of 400 *Bad Request* shall be returned.

If a set of uploads for a particular upload-id is not complete and no messages for that upload-id are received before the timeout value specified in the `cdmi_partial_timeout` expires, then the server shall terminate the set and discard all previously received messages for that upload-id.

If a new object is being created using the X-CDMI-Partial header, the `completionStatus` field in the response body shall be set to "Processing" and the value of the object shall not be returned to clients until the partial upload is considered complete.

If an existing object is being updated or replaced using the X-CDMI-Partial header, the object shall not be updated until the partial upload is considered complete.

## 2) Insert at end of table Table 100 - System-Wide Capabilities

Capability Name	Type	Description
<code>cdmi_partial</code>	JSON String	If present and "true", this capability indicates that the cloud storage system shall support the X-CDMI-Partial header "true" and "false" values.
<code>cdmi_partial_uploadid</code>	JSON String	If present and "true", this capability indicates that the cloud storage system shall support the X-CDMI-Partial header upload-id values.
<code>cdmi_partial_count</code>	JSON String	If present and "true", this capability indicates that the cloud storage system shall support the X-CDMI-Partial header count completion condition.
<code>cdmi_partial_range</code>	JSON String	If present and "true", this capability indicates that the cloud storage system shall support the X-CDMI-Partial header range completion condition.
<code>cdmi_partial_replace</code>	JSON String	If present and "true", this capability indicates that the cloud storage system shall support the X-CDMI-Partial header replace flag.

Capability Name	Type	Description
cdmi_partial_timeout	JSON String	If present, this capability indicates the upload-id timeout duration in seconds.

**3) Replace the "X-CDMI-Partial" Description in 8.2.4, Table 20 - Request Headers for Creating a CDMI Data Object using CDMI**

**4) Replace the "X-CDMI-Partial" Description in 6.2.3 Table 6 - Request Headers - Create a CDMI Data Object using HTTP**

**5) Replace the "X-CDMI-Partial" Description in 8.4.3, Table 29 - Request Headers - Update a CDMI Data Object using CDMI**

**6) Replace the "X-CDMI-Partial" Description in 6.4.3, Table 11 - Request Headers - Update a CDMI Data Object using HTTP**

Indicates that a partial upload is being performed. See 5.16 Object References.