



Immediate Query Extension

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
2012-11-01	1.0a	David Slik, NetApp Inc.	Initial creation based on document produced at the Portland Face to Face meeting
2012-11-02	1.0b	David Slik, NetApp Inc.	Edits and refinements based on internal review feedback
2013-01-28	1.0c	Dauid Slik, NetApp Inc.	Fixes to backwards compatibility to address ballot comments.

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.

Immediate Query Extension

Overview

CDMI provides a query mechanism based around the concept of persistence. A query queue is created, metadata is specified that defines the query operation, the query is performed asynchronously, results are populated in the queue, and then read by the client as separate operations.

This architecture, while providing significant value, is complex for clients that do not need to persist the results of a query. Specifically, a client must: a) asynchronously poll the query queue to determine when results are present and when the query has completed. b) delete the queue when results are no longer needed.

In order to provide a simpler interface for simple queries where a small number of results are expected and persistence is not required, the TWG has proposed the following approach to allow query queues to optionally not be persistent, with the results being returned immediately as the response to the initial query queue creation.

In addition, functionality permitting the immediate returning of results has been added to asynchronous query queue creation.

Modifications to the CDMI 1.0.2 spec:

1) Modify existing "cdmi_query" entry in table "Table 101 - System-Wide Capabilities"

Capability Name	Type	Description
cdmi_query	JSON String	If present and "true", this capability indicates that the cloud storage system supports persistent query queues.

2) Insert below "cdmi_query" in table "Table 101 - System-Wide Capabilities"

Capability Name	Type	Description
cdmi_query_immediate	JSON String	If present and "true", this capability indicates that the cloud storage system supports immediate query support.

3) Replace the first paragraph of Clause 22.1 with:

"A cloud storage system may optionally implement metadata and/or full-text query functionality. The implementation of query is indicated by the presence of the cloud storage system-wide capabilities for query, and requires support for CDMI™ queues when persisting query results."

4) Replace the third paragraph of Clause 22.1 with:

"When a client wishes to perform queries, it shall first check if the system is capable of providing query functionality by checking for the presence of the cdmi_query, or cdmi_query_immediate capabilities in the root container capabilities. If these capabilities are not present, and queues

are supported, creating a query queue shall be successful, but no query results shall be enqueued into the query queue."

5) Modify existing "cdmi_queue_type" in table "124 - Required Metadata for a Query Queue"

Metadata Name	Type	Description	Requirement
cdmi_queue_type	JSON String	The queue type indicates how the cloud storage system shall manage the queue object. Defined values are: <ul style="list-style-type: none"> • cdmi_query_queue – Perform an asynchronous query, which may return none, some, or all results in the request response body. A new queue object shall be created. • cdmi_query_immediate – Perform a synchronous query, returning all matching results in the request response body. The query queue object may not be accessible and shall be automatically deleted when the query completes. 	Optional

6) Insert new clause "22.2 Immediate Queries"

If "cdmi_query_immediate" is specified in cdmi_queue_type, all query results shall be immediately returned in the response body as shown in the below example.

EXAMPLE 3:

```
PUT /MyContainer/myQuery HTTP/1.1
Host: cloud.example.com
Accept: application/cdmi-queue
Content-Type: application/cdmi-queue
X-CDMI-Specification-Version: 1.0.2
```

```
{
  "metadata" : {
    "cdmi_queue_type" : "cdmi_query_immediate",
    "cdmi_scope_specification" : [
      {
        "domainURI" : "=/cdmi_domains/MyDomain/",
        "parentURI" : "starts /sandbox",
        "metadata" : {
          "cdmi_size" : "#> 100000"
        }
      }
    ],
    "cdmi_results_specification" : {
      "objectID" : "",
      "metadata" : {
        "cdmi_size" : ""
      }
    }
  }
}
```

```
201 Created
Content-Type: application/cdmi-queue
X-CDMI-Specification-Version: 1.0.2
Location: http://cloud.example.com/MyContainer/myQuery
```

```

{
  "objectType" : "application/cdmf-queue",
  "objectID" : "00007E7F00104BE66AB53A9572F9F51E",
  "objectName" : "myQuery",
  "parentURI" : "/MyContainer/",
  "parentID" : "0000706D0010B84FAD185C425D8B537E",
  "domainURI" : "/cdmf_domains/MyDomain/",
  "capabilitiesURI" : "/cdmf_capabilities/queue/",
  "completionStatus" : "Complete",
  "metadata" : {
    "cdmf_queue_type" : "cdmf_query_immediate",
    "cdmf_scope_specification" : [
      {
        "domainURI" : "=/cdmf_domains/MyDomain/",
        "parentURI" : "starts /sandbox",
        "metadata" : {
          "cdmf_size" : "#> 100000"
        }
      }
    ],
    "cdmf_results_specification" : {
      "objectID" : "",
      "metadata" : {
        "cdmf_size" : ""
      }
    }
  },
  "queueValues" : "0-0",
  "mimetype" : [
    "application/json"
  ],
  "valuerange" : [
    "0-111"
  ],
  "valuetransferencoding" : [
    "base64"
  ],
  "value":
  "ew0KCQkJIm9iamVjdE1E1iA6IC1wMDAwN0U3RjAwMTBFQjkwOTJCMj1GNkNENkFENjgyNCIsDQoJJCQkibWV0YWRhdGEiIDogew0KCQkJCSJjZG1pX3NpemUiIDogIjEwODI2MyINCgkJCX0NCgkJfQ0K"
}

```

Where the value of the above base64 encoded value is:

```

{
  "objectID" : "00007E7F0010EB9092B29F6CD6AD6824",
  "metadata" : {
    "cdmf_size" : "108263"
  }
}

```