

Table of Contents

HTTP TF Recommendations.....	1
Document Status.....	1
Basic Methods.....	1
Authentication and Authorisation.....	1
Space Reporting.....	1
Checksums.....	1
3rd party copy.....	2
ACL Management.....	2
Traffic Monitoring.....	2
BDII Integration.....	2

HTTP TF Recommendations

Document Status

First rough draft based on notes from the 3rd Jun meeting.

Basic Methods

Storage should support PUT, GET, DELETE and HEAD (HTTP) and PROPFIND and MOVE (WebDAV). Reference clients can be considered to be davix and curl, storage is NOT expected to support all clients for all operations. In particular, gateways are not obligatory and clients/applications can be expected to follow the necessary redirections.

We note that substantial performance testing remains to be done and the results of this process may throw up functional requirements (for example to allow throttling).

Default encodings are considered sufficient.

Q: Should the TF define a list of DAV properties required?

Authentication and Authorisation

The support of RFC proxies for authentication is required (legacy proxies are not). Support for VOMS attribute certificates is required. The underlying authorisation model is not prescribed and is assumed to be whatever the storage system already uses.

Space Reporting

The support of RFC 4331 on DAV collections is desirable and is required in the single-protocol case.

Q: Can support for `DAV:quota-available-bytes` be dropped from requirements? Could bytes available simply be taken from pledges?

Where DAV collections are functioning as directories in the storage namespace, storage providers may optimise by implementing only the top two levels of the experiment's hierarchy with this support.

Data freshness "of order hours" is acceptable.

Checksums

The only required checksum operation is `get checksum of type X for resource Y`. The storage can retrieve this from metadata or calculate it on demand. This should be implemented according to RFC 3230 (want-digest in request header) and <http://www.iana.org/assignments/http-dig-alg/http-dig-alg.xhtml>. This is currently implemented in dCache and DPM.

Q: Should we define a DAV property for the checksum? The advantage is that a full response to a HEAD request does not have to be assembled by the server.

Q: The only checksum type mentioned in the TF twiki is MD5 which types are actually required? GridFTP typically creates Adler32. LHCb requires at least Adler32.

An upload with checksum , where the checksum is supplied along with the upload, in order to allow the response to reflect the status of checksum validation, is desirable but not required.

Q: Should we define/find a standard interface for upload with checksum ? One possible interface is provided by ownCloud (is this standardised anywhere?)

3rd party copy

Status dCache and DPM have agreed on an interface for delegating credentials to storage (gridsite delegation protocol) and requesting that the storage COPY a file to a remote server which need not be modified to support the operation. The latest versions of those systems implement this interface.

The following were considered to be necessary in the single protocol scenario:

3rd party copy between a pair of compliant storage systems where either endpoint can be requested to initiate the transfer (ie either the source or the destination can be addressed). Internal implementation details (push/pull) are left to the providers.

The following was seen as desirable:

The HTTP storage should be able to retrieve files from or put files onto unmodified remote HTTP/WebDAV or S3 servers, presenting the same interface to initiate the operation as described above.

Interaction with other protocols such as GridFTP was discussed but not considered important.

ACL Management

Making ACL information visible over HTTP is desirable, and ACLs must be honoured for HTTP access, but management features can be exposed only to the storage administrator and are not required over HTTP.

Traffic Monitoring

To be discussed - see the minutes of the meeting on 15th July - <https://indico.cern.ch/event/390085/>.

BDII Integration

Desirable, but we have to agree details here.

This topic: LCG > HTTPStorageRecommendations

Topic revision: r4 - 2015-09-09 - OliverKeeble



Copyright &© 2008-2019 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback