

# Table of Contents

Site.....	1
-----------	---

# Site

Question	<i>Response</i>
<b>Site and Endpoints</b>	
What is the site name?	FZK_LCG2
Which endpoint URLs do your archival systems expose?	srm: {atlassrm-fzk,cmssrm-kit,lhcbsrm-kit}.gridka.de
How is tape storage selected for a write (choice of endpoint, specification of a spacetoken, namespace prefix).	dCache and xrootd both archive into the same tape storage backend.
<b>Queue</b>	
What limits should clients respect?	
Max number of outstanding requests in number of files or data volume	dCache assigns flush and stage tasks to pools, which all have an upper limit for concurrent active tasks, usually 2k. Requests beyond that are queued. For xrootd the limit is a total of 3200 concurrent flushing and staging tasks.
Max submission rate for recalls or queries	
Min/Max bulk request size (srmBringOnline or equivalent) in files or data volume	
Should clients back off under certain circumstances?	SRM feature with dCache: A limit can be set for every request type, including srm-bring-online (10k by default). Once more requests are accumulated, SRM will block and return "overloaded" error. For xrootd, there is no such feature that would recognise an overload situation. If a file cannot be staged from tape, xrootd will fail on each subsequent request immediately.
How is this signalled to client?	srm-bring-online and accessing a file will fail.
For which operations?	srm-bring-online / open
Is it advantageous to group requests by a particular criterion (e.g. tape family, date)?	In theory, yes, that would be advantageous. But we cannot guarantee that it will stay in that order or grouping. There is only a loose chronological order.

What criterion?	Timestamps would be most useful, since tape families don't necessarily match what the VOs use as classification.
<b>Prioritisation</b>	
Can you handle priority requests?	No.
How is this requested?	
<b>Protocol support</b>	
Are there any unsupported or partially supported operations (e.g. pinning) ?	All features that dCache and xrootd support natively should work for GridKa, too.
<b>Timeouts</b>	
What timeouts do you recommend?	
Do you have hardcoded or default timeouts?	Yes, we have default timeouts with dCache for flushing and staging of at least 24 hours (may be larger on request). No timeouts are enforced with xrootd.
<b>Operations and metrics</b>	
Can you provide total sum of data stored by VO in the archive to 100TB accuracy?	Yes
Can you provide space occupied on tapes by VO (includes deleted data, but not yet reclaimed space) to 100TB accuracy?	Yes
How do you allocate free tape space to VOs?	We do not allocate space on tape for any VO.
What is the frequency with which you run repack operations to reclaim space on tapes after data deletion?	We have defined a threshold for "tape occupancy", which will trigger reclamation per tape.
<b>Recommendations for clients</b>	
Recommendation 1	
---> Information required by users to follow advice	
Recommendation 2	
<b>Buffer Management</b>	

Should a client stop submitting recalls if the available buffer space reaches a threshold?	
---> How can a client determine the buffer used and free space?	
---> What is the threshold (high water mark)?	
---> When should the client restart submission (low water mark)?	
If the client does not have to back off on a full buffer, and you support pinning, how is the buffer managed?	
---> Is data moved from buffer to another local disk, either by the HSM or by an external agent?	
<b>Additional questions</b>	
Should any other questions appear in subsequent iterations of this survey?	

History:

OliverKeeble - 2018-01-30

Created article.

XavierMol - 2018-04-16

Inserted answers for KIT/GridKa on behalve of Doris Ressmann.

XavierMol - 2018-04-16

Changed format of the tabel

This topic: HEPTape > KIT-GridKa\_Survey

Topic revision: r4 - 2018-05-04 - OliverKeeble



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback