

Table of Contents

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

Site

Question	<i>Response</i>
Site and Endpoints	
What is the site name?	T1-JINR
Which endpoint URLs do your archival systems expose?	se-hd02-mss.jinr-t1.ru
How is tape storage selected for a write (choice of endpoint, specification of a spacetoken, namespace prefix).	The whole storage is dedicated to single client CMS
Queue	
What limits should clients respect?	Only physical tape limits (all tapes)
---> Max number of outstanding requests in number of files or data volume	
---> 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?	
---> How is this signalled to client?	
---> For which operations?	
Is it advantageous to group requests by a particular criterion (e.g. tape family, date)?	
---> What	

critterion?	
Prioritisation	
Can you handle priority requests?	No.
---> How is this requested?	
Protocol support	
Are there any unsupported or partially supported operations (e.g. pinning) ?	All supported by dCache.
Timeouts	
What timeouts do you recommend?	No timeouts. Tasks beyond reasonable time are handled manually.
Do you have hardcoded or default timeouts?	No timeouts
Operations and metrics	
Can you provide total sum of data stored by VO in the archive to 100TB accuracy?	6300 TB
Can you provide space occupied on tapes by VO (includes deleted data, but not yet reclaimed space) to 100TB accuracy?	7190 TB
How do you allocate free tape space to VOs?	The whole space to CMS VO.
What is the frequency with which you run repack operations to reclaim space on tapes after data deletion?	After massive deletion.
Recommendations for clients	
Recommendation 1	
---> Information required by users to follow advice	
Recommendation 2	
Buffer Management	

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?	
Additional questions	
Should any other questions appear in subsequent iterations of this survey?	

-- OliverKeeble - 2018-01-30

This topic: HEPTape > JINR_Survey
 Topic revision: r3 - 2018-04-19 - 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