

# Table of Contents

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

# Site

Question	<i>Response</i>
<b>Site and Endpoints</b>	
What is the site name?	TRIUMF
Which endpoint URLs do your archival systems expose?	srm://triumf.ca/atlas/tape/
How is tape storage selected for a write (choice of endpoint, specification of a spacetoken, namespace prefix).	endpoint same as above, ATLAS only
<b>Queue</b>	
What limits should clients respect?	We do accept any kind of recalls, but prefer bulk requests. We purposely delay requests to get processed in order to get bulk requests
---> Max number of outstanding requests in number of files or data volume	No exactly number, there was one peak number more than hundreds of k during ATLAS test, no problem for us
---> Max submission rate for recalls or queries	
---> Min/Max bulk request size (srmBringOnline or equivalent) in files or data volume	5k-30k (per session) is good, even 1k is ok, but few at a time is not welcomed few TB - 200TB
Should clients back off under certain circumstances?	Ideally no, our HSM is able to handle hundreds of k requests without load problem, however there is a hard limit from disk buffer size, we don't use any extra disk buffer for tape operations, the disk buffer that tape use is also used for ATLAS((dcache hsm pools), space is limited to that, so realistically speaking, few TB-200TB a day is good enough, though can be reached to 500TB
---> How is this signalled to client?	through SRM
---> For which operations?	Depends on ATLAS how launch and check requests
Is it advantageous to group requests by a particular criterion (e.g. tape family, date)?	tape family grouped by datatype, dataset, and date

Triumf\_Survey < HEPTape < TWiki

---> What criterion?	data will wait for at least 45 hours within a dataset if the dataset size not exceed a tape capacity, or will be processed when the dataset size > a single tape capacity, different datasets will be packed together by project, data type, for example: data17_900GeV, mc15_5TeV, further grouped by datatype, datatape, mctape etc..
<b>Prioritisation</b>	
Can you handle priority requests?	quite often tape is quiet, no need yet, also no priority flat in ATLAS operations, can be implemented if particular circumstance is identified
---> How is this requested?	
<b>Protocol support</b>	
Are there any unsupported or partially supported operations (e.g. pinning) ?	we can always pin or restore a file through dcache HSM interface, bulk or individual
<b>Timeouts</b>	
What timeouts do you recommend?	All tape requests will be served . 24 hours or even longer if large amounts data requested, few hours if requests are small number
Do you have hardcoded or default timeouts?	No
<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,all deleted data on tape is logical delete
How do you allocate free tape space to VOs?	Through info publish
What is the frequency with which you run repack operations to reclaim space on tapes after data deletion?	we do repack when space is needed or media upgrade, just done LTO5-> lto7 migration this early year, 2600 LTO5 tapes, 2.5PB(on tape since 2012) migrated to LTO7 tapes at new site, no data lost
<b>Recommendations for clients</b>	
Recommendation 1	Bulk requests by dataset, you will get your data quicker, our tape system pick the tapes has most requests first
---> Information required by users to follow advice	

Recommendation 2	Let us know how you use tape data, then we can tweak our tape data packing policy, you will get fast readback and data safe protection, tape has mounting limits
<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?	

-- OliverKeeble - 2018-01-30

This topic: HEPTape > Triumf\_Survey

Topic revision: r4 - 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