

Qserv: Questions from in2p3 to SLAC and related answers

Emmitters Emmanuel Medernach, computer scientist, developer and EGI TIER2 manager
Fabrice Jammes, computer scientist, developer and head of LPC/in2p3 IT department

Approval Emmanuel Gangler, scientist, head of LSST team at LPC/in2p3

Diffusion

For action SLAC (Douglas Smith, Daniel Wang, Jacek Becla)

For information LSST-France

Change history

Date	Author	Description of the update
6Feb2013	LPC	Creation
13Feb2013	LPC	Answer proposal to questions 1 to 8 (2,3,5,6 7 closed), thanks to SLAC help.
28Feb2013	LPC	Questions 4, 8, 9, 10, 11 still opened or added. (#1 closed)

N°	Theme	Priority	Remark	Status
1	<p data-bbox="293 244 398 276">Install</p> <p data-bbox="203 323 443 355">Ticket/ #2511</p>	<p data-bbox="551 244 631 276">High</p>	<p data-bbox="689 244 1406 276"><u>IN2P3 Question (submitted on 1Feb2013) :</u></p> <p data-bbox="689 292 1899 395">Is the code review for ticket #2511 (scons automated install) done ? If not, what is the timeline ? This is a priority item as further development like automated tests, are dependent on this ticket.</p> <p data-bbox="689 451 1348 483"><u>SLAC Answer (submitted on 9Feb2013) :</u></p> <p data-bbox="689 499 1906 754">Daniel send this email to Douglas : Okay, please bring the branch up to date with the master, confirm that things still work, and then put it on master. See steps 4 and 5 in the example: https://dev.lsstcorp.org/trac/wiki/GitDemoAndTutorial#LSSTgitworkflowandbranchmanagementpolicy Thanks! -Daniel</p> <p data-bbox="689 818 1429 850"><u>IN2P3 Question (submitted on 13Feb2013) :</u></p> <p data-bbox="689 866 1182 898">Could you do the merge please ?</p> <p data-bbox="689 962 1406 994"><u>IN2P3 Answer (submitted on 28Feb2013) :</u></p> <p data-bbox="689 1010 1025 1042">We've done the merge.</p>	<p data-bbox="1946 244 2051 308">Closed 2</p>
2	<p data-bbox="237 1046 454 1078">Data-loading</p>	<p data-bbox="551 1046 631 1078">Med</p>	<p data-bbox="689 1046 1406 1078"><u>IN2P3 Question (submitted on 1Feb2013) :</u></p> <p data-bbox="689 1094 1899 1158">Could we have some information about PT1.1 'Source' loading, as it would allow to complete the data loading and to execute queries for Qserv validation.</p> <p data-bbox="689 1174 1899 1398">Is the next command ok : <code>/opt/qserv-dev/bin/python /opt/qserv-dev/qserv/master/examples/partition.py --output-dir /data/lsst/pt11_partition --chunk-prefix Source --theta-column 33 --phi-column 34 /data/lsst/pt11/Source.txt --num-stripes 10 --num-sub-stripes 2</code> are these parameters correct : --theta-column 33 --phi-column 34</p>	<p data-bbox="1946 1046 2051 1078">Closed</p>

			<p><u>SLAC Answer (meeting of 6Feb2013) :</u> In tests/case01/data/Source.schema zero index based of raObject and declObject are 33 and 34.</p> <p>'Source' table fields ' raObject' and 'declObject' are strictly equals with fields 'ra' and 'decl' of 'Object' table. Object moves aren't modeled in these parameters.</p>	
3	Automated tests	Low	<p><u>IN2P3 Question (submitted on 6Feb2013) :</u> Is the development on tests/runTest.py inside the perimeter of ticket #2014 or will a new ticket be created ? At the moment development is done in branch "u/fjammes/automatedTest"</p> <p><u>SLAC Answer (submitted on 6Feb2013) :</u> DST : Well, yes, that can be part of the ticket. I think the important point there is that the ticket should be for only running runTest.py from the installer. Other dev. of the use of runTest.py should be another ticket.</p> <p><u>Meeting (7Feb2013) :</u> DST : remove test data after test execution not to conflict with other data (like PT1.1 for example). Indeed, 'LSST__Object' and 'LSST__Source' tables in 'qservMeta' database are the same and so conflict for each dataset.</p>	Closed
4	Automated tests / Data-loading	Low	<p><u>IN2P3 Question (submitted on 6Feb2013) :</u> The automated tests with qserv mode needs to load partitioned data. In your tests/runTest.py comments, you propose this : <i>#use the loadPartitionedObjectTables.py script to generate loadO</i> but loadPartitionedObjectTables.py doesn't exist. Should we use instead <i>master/examples/loader.py</i>, or re-implement in Python the load() function of <i>admin/custom/bin/qserv-admin.pl</i>?</p> <p><u>SLAC Answer (submitted on 6Feb2013) :</u></p>	Open

			<p>DST : Hmm... well good question. Sounds like we need to create a re-usable python loader class, maybe need a new ticket there ?</p> <p><u>IN2P3 Question (submitted on 13Feb2013) :</u> What do you think of next script (see mail sent to qserv-l@slac.stanford.edu on 8Feb) :</p> <pre>python master/examples/loader.py --help</pre> <p>It seems to do the job, it may have been written by Jacek.</p> <p><u>IN2P3 Question (submitted on 28Feb2013) :</u> We now use <code>master/examples/loader.py</code> in <code>runTests.py</code> with success, and we plan tu use it to load PT1.1 Sources. Do you think it is a good solution ?</p>	
5	Data-loading	Med	<p><u>IN2P3 Question (submitted on 6Feb2013) :</u> Furthermore, in order to allow Qserv users to save expensive storage cost., we propose to develop a dynamic loader, which load data in DB directly at partitionning, without having to store partitioned data in intermediate files ? Do you think it is interesting ?</p> <p><u>SLAC Answer (submitted on 6Feb2013) :</u> DST : Yeah, well, we've thought of this, and there was some talk with Serge about having a type of partitioner which is more of a streaming executable. Taking data from one input, and loading it into the database as output. There is a concern, since on large data sets, these processes take days, and are then quite fragile for errors, and you don't want the loading left in some unknown state on failure, such that you need to re-start everything from the beginning. It would be good to be able to batch the partitioning and loading into smaller when contained parts, and be able to re-start on the part left over on failure. But not sure that is the best thing to work on just yet, I think we might need more experience on handle large datasets with the partitioning and loading still.</p>	Closed

			<p>Meeting (7Feb2013) : DST : complicated and cost-effective : error on failure should be managed.</p>	
6	Minutes	Med	<p><u>IN2P3 Question : Wiki page for Phone meeting minutes (submitted 6Feb2013)</u> We would like to have a centralized place to put phone meeting minutes, could it be possible to use your wiki pages for that ?</p> <p><u>SLAC Answer (submitted on 6Feb2013) :</u> DST : Um, maybe, does trac scale enough to have a new page added for each meeting?</p> <p>Meeting (7Feb2013) : LPC team has now created a page on the wiki : https://dev.lsstcorp.org/trac/wiki/db/Qserv/IN2P3/Meetings</p>	Closed
7	Benchmarking	Med	<p><u>IN2P3 Question : Standard performance tests for QServ (submitted 6Feb2013)</u> Thanks for giving us next link : https://dev.lsstcorp.org/trac/wiki/dbQserv/perTableKeyCache . Does the queries defined in https://dev.lsstcorp.org/trac/wiki/dbQueriesForPerfTest also fit for Qserv master and node performance and benchmarking tests ? Beyond individual queries, is there a document describing a typical queries workload to be used for benchmarking ? Could we define standard performance tests for QServ master and nodes ? Could you provide the ones you used before ?</p> <p><u>SLAC Answer (submitted on 6Feb2013) :</u> JBA : Regarding #7, for sample queries related to PT1.1, I'd suggest to use: https://dev.lsstcorp.org/trac/wiki/db/Qserv/250NodeTestPlan</p> <p>However, PT1.1 is not good enough for testing this: https://dev.lsstcorp.org/trac/wiki/db/Qserv/PerTableKeyCache</p>	Closed

8	Data-loading	<p>Med</p> <p><u>IN2P3 Question : Standard performance tests for Qserv (submitted 7Feb2013)</u></p> <p>How could we generated Qserv meta LSST_Object and LSST_Source tables with good (chunk, subchunk) data.</p> <p>Meeting (7Feb2013) :</p> <p>DST : watch qserv-admin.pl and try to use it.</p> <p><u>IN2P3 Question (submitted on 13Feb2013) :</u></p> <p>We've changed some partitionning parameters (see tests/runTests.py in ticket #2014) so that chunk and subchunk data are added in Source_ChunkIds and</p>	Open

because this test exposes the fact that we have one huge table (RunDeepForcedSource), and some small tables, such as the Science_Ccd_Exposure, or RefObject.

It'd make more sense to use the Winter13 data set for this test. That data set is on the large side... a little over 2 TB. I believe we will be moving more towards relying on that data set in the future, so it makes sense to get it. We are in the process of partitioning it right now on slac machines. We are happy to make it available to you, if you have space to take it, either as csv files, or partitioned data.

Meeting (7Feb2013) :

DST : test with PT1.1 data with MySQL tuning parameters

[IN2P3 Answer \(submitted on 13Feb2013\) :](#)

This point must be delayed as in2p3 priority is to run Qserv on 250 nodes soon. So in2p3 team must focus now on installation/configuration/administration and data-loading procedures.

[IN2P3 Question : Standard performance tests for Qserv \(submitted 7Feb2013\)](#)

How could we generated Qserv meta LSST_Object and LSST_Source tables with good (chunk, subchunk) data.

Meeting (7Feb2013) :

DST : watch qserv-admin.pl and try to use it.

[IN2P3 Question \(submitted on 13Feb2013\) :](#)

We've changed some partitionning parameters (see tests/runTests.py in ticket #2014) so that chunk and subchunk data are added in Source_ChunkIds and

			<p>Object_ChunkIds table during data-loading. Then meta are added in LSST_Object and LSST_Source via of bunk of mysql queries of this type :</p> <pre># mysql -u<u> -p<p> qservTest_case01_q qservMeta -e "insert into LSST__Object SELECT objectId, chunkId, subChunkId from qservTest_case01_q.Object_100"</pre> <p>(inspired from initial Jacek comments in tests/runTests.py). It is quite simple and efficient enough for the tiny test dataset. Does it seems correct ?</p> <p><u>IN2P3 Question (submitted on 28Feb2013) :</u> We now use this technique in runTests.py with success, and we plan tu use it to load PT1.1 Sources. Do you think it is a good solution ?</p>	
9	Data-loading	Med	<p><u>IN2P3 Question : W13 data transfert at CC-IN2P3 (submitted on 28Feb2013)</u> Could you please transfer us W13 data at CC-IN2P3 in /sps/lst/Qserv/data/Winter2013/ ?</p>	Open
10	Data-loading	Med	<p><u>IN2P3 Question : W13 data transfert at CC-IN2P3 (submitted on 28Feb2013)</u> How the chunk number (7200) is computed depending on stripes and substripes, it seems to be computed as $2*60*60$, what is 60 ?</p>	Open
11	Data-loading	Med	<p><u>IN2P3 Question : Design of runTests.py (submitted on 28Feb2013)</u> Would you be you interested in generalizing our design used in tests/runtTests.py, so that we continue to improve it. It mainly consists in :</p> <ul style="list-style-type: none"> - read and analyze input data config in a separate class (type of files, extensions, use of a config file with meta data) - Use of QservDataLoader which could be used for larget tests, like for PT1.1 or W13 for example. This class could be generalized for loading data on worker. 	Open