

DM JTM 2016-02 Meeting Notes

Error rendering macro 'toc'

```
[com.ctc.wstx.exc.WstxLazyException] com.ctc.wstx.exc.  
WstxUnexpectedCharException: Unexpected character "'" (code 34) expected '=' at  
[row,col {unknown-source}]: [1,25184]
```

Link to meeting agenda: [DM Agenda for Joint Technical Meeting 2016-02-22 to 24](#)

Science Pipelines and SQuaRE

Eliminate boilerplate in source files

Jonathan is working on this. Will be done after the HSC fork is merged. Will be an automatic transition managed by SQuaRE.

Switching to NumPy docstring format

Waiting for the new documentation build system, "LSST the Docs". Should be in the next week or two. Will be discussed in the "decision making" session tomorrow.

Still work to be done in determining the best way to document tasks and how tutorials and examples should be published

Continuous Integration

What does it mean to have "enough" CI to perform major changes? Suggestion that we need all the supported cameras processed through to measurement on coadds every week. This should include regression tests as well as tests for absolute values and how they relate to the science requirements.

What is a supported camera? SQuaRE emphasize that we should not guarantee to support in perpetuity every `obs_` package that is submitted. Agreed that they should have designated "sponsors" who will maintain them. However, if a change to the stack causes a test to break for a specific camera, our first assumption should be that it's an algorithmic error rather than a bug in `obs_`.

Plans for nightly QA runs that will run large scale (~hours) tests. Regular CI will run shorter (minutes) tests.

Limited developer support in Spring 16 and `lsst-dev` shared stack

SQuaRE will provide only limited developer support during the short Summer 16 issue. Science Pipelines identify one issue deserving of immediate support as the shared stack on `lsst-dev`. JDS will act as point of contact for SQuaRE on this.

Process Control and Data Access

Issues to discuss in details

- Abstracting access to OCS through butler
- Fetching data from staging area to disk (e.g., templates, exposures for forced dia photometry)
- provisioning hardware for L1: when, what
- batch processing related needs for L2/L3
- data backbone for L2/L3
- cross-site failure recovery
- cross-site upgrades (are all sites required to be on the same release etc)
- replicating L3 across different sites / L3 user mydb synchronization across multiple sites
- distributing DRP products - via network or physical shipping?
- storage technology for large files (object store)
- capturing provenance (Cooper starting to think about it)
 - gray area: who captures provenance about OS/hardware
- how much containerization should we be doing?
 - does it simplify provenance capture?
- NCSA is writing docs about L2
 - Margaret will expose to Data Access when the docs are in reasonable shape
- need to bring IN2P3 to Data Access <-> NCSA into discussions (via JCC)

SUIT and Architecture

- User workspace discussion. SUIT presented a diagram of the preliminary thoughts on workspace. The workspace could be like user home. Users can save their work, install their software, run a LSST task from the workspace. Users should be able to access the extended storage space (at LSST or not) and access the computation facility (at LSST or not). iPlant was suggested as a possible option for workspace. We will have to get the relevant parties together to discuss more and plan a workshop at some stage.



- SuperTask's role in workspace.

Additional notes from Paul Wefel

- The network made remote NCSA participation (4 people) challenging but eventually managed to get Jason Alt connected.
- There is a little bit of a chicken / egg problem here. SUIT would like to know the DM environment / constraints and DM is probably waiting on SUIT for their environment requirements
- High level overview of SUIT process
 - Users will come in through a web page / portal
 - Authenticate through the web page
 - Now interacting with SUIT, SUIT wants to act as the user for all processes run
 - Launch python processes behind the SUIT (using extension architecture that has been proposed)
 - example - looking at an image in detail
 - how the background process is launched is TBD
 - would like the background process to run as the UID/GID

Another mechanism being explored for SUIT is through real interactive environments where a VM or container running on a user workstation/laptop is pre-configured with SUIT software stack.

One question that came up: what is the resource limit for a Science user running jobs.? (Asked to Jason. Couldn't hear the answer)

VM vs. Batch processing, SUIT understands Don's point from yesterday on using a batch job

**SUTI would like to have a two / three day meeting with NCSA to work out system interaction details (action item that isn't owned)

Jason and Gregory to talk

Science Pipelines and Architecture

Previous coding standard decisions

Python 3, idiomatic Python, etc.

Almost all of this work has been blocking on the HSC merge. However, the onus is then on Science Pipelines to decide if and when to schedule the work: Architecture will not make this decision. Should assess:

- Is work done as a big bang or incrementally?
- If work is incremental, how does it impact ongoing tickets?

(KTL adds: Prioritization needs to happen through the usual mechanisms with *input* from Science Pipelines and Architecture. Science Pipelines does not get to make this decision unilaterally either.)

Rough estimate that the effort to make our Python code "idiomatic" would be two weeks of work for three people. Important to have enough test coverage in place to make sure this doesn't break anything; running an HSC data release candidate should be sufficient.

The plan is to carry out the Python 3 work in a "big bang" hack session at the August All Hands.

AstroPy

There will be a meeting with the AstroPy developers at UW in ~1 month. Where we draw the line in terms of stack integration has a huge impact on the amount of effort required. Aim to have a rough idea of this before the meeting, but expect to refine it based on input from the AstroPy folks. An ideal outcome might be to produce AstroPy affiliated packages which provide C++ APIs that we can use in the stack.

Many-core architectures

It is likely that future architectures, as we might reasonably expect to run on in operations, will skew towards large numbers of cores per system, with relatively small memory and storage per core. This suggests we should move towards threading in our high-performance C++ code rather than relying on multiprocessing at a higher level. The likely, but not definitive, technology choice is OpenMP. We do not regard Science Pipelines as responsible for "blue skies" research in this area, but it is reasonable to expect that they will provide examples and requirements. We do not think this decision can be driven purely "bottom-up" but will ultimately require input from the Architecture team.

Process Control and SQuaRE

SUIT and Data Access

We all would like to exercise the SUIT to Data Access system in a basic way regularly, deploy nightly out of NCSA. The idea is to get this going with a very small repository. Then, when we receive the PanSTARRS data, SUIT can start to use that data to develop new features for users to try and write a robust set of regression tests.

Jacek would like to get the small set of available around July-August 2016. Before then SUIT will deliver example queries so we can verify we're prepared to run the tests. SUIT would like to have the web portal to access PanSTARRS data ready by the end of November, which means that data should be ready for access by end of September the latest.

Senior management should prioritize how important it is to deliver Pan-STARRS data through Qserv/SUIT to our users, and define timeline. Two months ago the top two priorities were: end-to-end system, and serving Pan-STARRS data

Once the Pan-STARRS data arrives it should take Data Access about a month to get the pan stars data ready. Note, we don't know what format it's coming in; we will need to think about how to load it. After loaded, keep it available "for ever". OK to limit access for users during scheduled stress DB tests.

Remote Butler - SUIT really wants a Java client to the remote Butler, for the Firefly service. We need to discuss it with the architecture group, but perhaps we should consider "not" doing the python butler client for now, and doing the Java one instead. Java client useful for SUIT server side (performance reasons, does not want to fork python process). Python access still needed for SUIT user access.

Would be nice to have simple prototype in Fall 2016 that demonstrates credentials acquiring/passing

Science Pipelines and Process Control

Computing resources for developers

There's a modest cluster hanging off the back of `lsst-dev` which is available for developers to use. Use the `ctrl_orca` middleware to access it. Documentation is available on Confluence.

This cluster is not addressable by the `ctrl_pool` middleware recently ported from HSC. Adding this functionality to `ctrl_pool` should be a relatively modest task for Science Pipelines developers, although it probably falls strictly outside their remit per WBS. Getting this working will be a requirement in the intermediate future.

Future middleware development

Some discussion of the SuperTask concept and how it might relate to Science Pipelines. Nobody in the room is confident in describing either when it will be available or what impact it might have on existing tasks. We hope this will be clarified soon.

Hsin-Fang has volunteered to be a contact at NCSA for fielding short term maintenance issues with the existing middleware (`pipe_base`, `pex_config` etc). The Science Pipelines group will be proactive about putting larger requests to the Middleware Group well in advance of future cycles so that they can be included in the plan.

The group at NCSA which will be working on middleware is gradually ramping up in terms of staffing and experience.

Many-core architectures

Agreement that there will be effort available at NCSA during Summer 2016 to support an investigation into future stack parallelization strategies.

Feedback on L1 presentation

Simon is concerned that the authors/owners of the various boxes in Don's presentation from this morning should be clearly stated, and that the interfaces where alert data passes from Science Pipelines to NCSA, or vice versa, should be documented.

Architecture and Data Access

Public Interfaces to DAX

- VO
 - Architecture team owns decision which subset. The plan is to pick protocols that are in wide spread use
 - Currently on NCSA WBS. Arch team and Mario wants it to be in Data Access WBS. This would increase scope of Data Access (need to estimate, will need ~1FTE for some time).
 - This work also includes discussing with VO community, pushing back when protocols are insufficient for us.
 - Global DM-wide VO expert was considered, we decided it is not a good idea
- RESTful webservices

Internal interfaces

- butler interfaces
- SQL / mysql. Note, that there is strong desire to NOT expose SQL / mysql interface to public, because once we do, we will be stuck with supporting mysql interface for ever

Metadata / butler

- metadata depends on butler. It can use remote access to get data via butler
- butler depends on metadata. Usecase: metadata can manage repository configurations (versions), the versions are used by butler

Saving queries for published papers

- will save query text, that is easy
- reproducibility is harder to guarantee
- we could push results to cloud (especially if requestor pays)
- side note: need to think about sharing Qserv across multiple data releases and how we would upgrade UDFs without sacrificing reproducibility

Global catalog

- need more definitions of interfaces to global catalog. This needs to be discussed with NCSA. Not sure when we will know
- Since we need it soon, prototyping is useful
- but keep an eye on existing systems, like Fermi DataCat

SUIT and SQuaRE

- VO protocol support. VO table binary format was suggested for large data amount for data transfer. Using VO protocol as internal interface was suggested.
- Calibration and EFD information are important to SQuaRE and to SUIT since QA and users would like to know. More discussion and understanding are needed.
- Documentation. SUIT will invite Jonathan Sick to IPAC for a day or two to work together on documentation summer.

Science Pipelines and SUIT

Large scale visualization

It was agreed that the large scale visualization of HSC data as presented by Robert was a compelling use case for SUIT. There was some discussion about what it would take to achieve this technically. It was agreed that it would be technically feasible to generate PNG (or equivalent) images which would be used in such a visualization as part of the data release processing (rather than post-processing the data release by the SUIT group). The key is to pre-generate the large images in different resolutions.

Interfacing Firefly with Science Pipelines

Following some work in 2015 to integrate Firefly with the `afw.display` system, this effort has been languishing. The SUIT group are keen to see Science Pipelines developers using, and providing feedback on, Firefly; Science Pipelines developers are keen to have access to better visualization and debugging tools. However, the barriers to entry for getting Firefly running on individual laptops are -- arguably! -- offputting.

It was agreed that we will aim to establish a Firefly service on `lsst-dev` which will enable developers to visualize data stored on the filesystem. We hope this will help build a critical mass of Firefly users.

A compelling user case for future development would be to enable visualization of `afw.tables` through an interface similar to `afw.display`.

SQuaRE and Data Access

SQuaRE wants

- TAP interface
- write access via TAP (or RESTful), but not direct SQL
- does not care about authorization in the short term
- remote access to butler (python client), read and write
- S3 plugin backend for butler
- [Unknown User \(npease\)](#) to visit them for ~ a week

Data Access wants

- integrating Qserv integration tests with CI

Architecture and Process Control

Science Pipelines and Data Access

Integrating calibration data with the Butler

Important to be able to select calibration data according to complex criteria, e.g. "give me the calibration data that would have been used were I reducing *this* data on *that* day". A variety of approaches were discussed to addressing this problem, including educating the Butler about all the "calibration roots" (corresponding to different versions of calibration data). Closely related to capturing provenance (see below): need to be able to reproduce results by capturing not only the version of the code but the specific calibration data used to generate them.

Capturing provenance

- Need to understand how to capture provenance about software, eg versions of calibrations, need to capture calibration root (location where the contents is)
- Create child rerun for each case when command / output changes

Better Documentation

- lots of features going into butler, driven by experienced Science Pipelines developers with complex use cases (e.g. Robert and Jim), translated by K-T into concrete requirements; mapping back from there to practical use in pipeline code is not obvious, particularly for less experienced Science Pipelines devs.
- write technical note(s), target to specific audience.
- possible solution is butler dev to do demo with science stakeholder, and science stakeholder writes technical note for science users. This would serve as acceptance testing for butler feature.

Feature requests and requirements

- Discussed the difference between feature requests (e.g. data being indexed in a particular way) and requirements (e.g. data of type X must be accessible with latency no more than Y under conditions Z; should be traceable to baseline documentation).
- Science Pipelines developers may have ideas about the implementation necessary to meet specific requirements (e.g. caching), but ultimately that's a call for the Data Access Team: not appropriate for us to simply dump them into JIRA.
- However, we do need to figure out a good way to capture this sort of material.

Fetching pixel data

- Science pipe know which pixels they want
- choosing order and getting from disk is middleware job
- Need to understand interfaces

Reference catalog

- Want to do joining with reference catalog through butler. [John Swinbank](#) to ensure that somebody from Science Pipelines writes a specification for the required behaviour.

- How do we get reference catalog (for a given circle, or box). Through butler?
- eg might use usno for astrometry, and something else for photometry
- currently stores catalogs per shard, shards based on htm

Interactions w/database for L1, for L2

Need to go through the whole flow, in details

Table-valued functions

Table-valued functions are a major feature of the SciServer infrastructure. Not currently available in MariaDB or Qserv (although less-than-optimal workarounds may be possible). Members of the Data Access will be meeting with Monty Widenius next week and may discuss it with him.

SQuaRE and Architecture

SUIT and Process Control

Margret and Don had other session to attend. Jason from NCSA called in for the discussion. We touched on workspace, resource management, and authorization/authentication. Xiuqin would like to start a regular discussion on work space which leads to a 2-3 day face to face design meeting of all parties.

[workspace diagram](#)

DM-wide Decision Making

Virtual Observatory

- Data access protocols should be handled by Data Access team.
- We will choose a set of current commonly-used community standards and write it into the baseline. That can be changed later if needed.
- We should push back on things that need change in the standards.
- "Author" components that are VOEvent Transport Protocol compliant endpoints are the responsibility of the AP team.

Unicode

- Focus on adding Unicode test cases for strings that would come from outside.
- Make sure that C++ interfaces handle Unicode properly.
- Move to Python 3 to avoid having to add "u" prefixes to all constant strings in Python code.

Data and file formats

- Need to answer:
 - Outside export requirements (could be more than FITS)
 - Internal format needs to be efficient for access and export
 - How self-describing and what does that really mean?
 - How do we handle multi-part objects?
- [Unknown User \(ciardi\)](#) will take the lead in defining the file format(s) which LSST will ultimately provide to end users.
- The Architecture team will take the lead in defining the file format(s) which are used within the stack. Efficient conversion to end-user-facing format(s), where applicable, will be an important (but not the only) consideration.
- No timeline was established for these definitions being made available.

Naming the "Science Pipelines"

- A consensus was not quite established that establishing a new "brand" for the Science Pipelines was a good idea: [Mario Juric](#) in particular had reservations about whether this was necessary or desirable.
 - There was some concern that leaving e.g. the middleware unbranded could foster a perceived divide between the science and engineering parts of the project.
 - It was agreed that if a new name were to be chosen, the appropriate top-level product to which it would apply would be `lsst_apps`. It would exclude the `obs_` camera packages.
 - The next step is to write a complete description of what this change would involve so that it is possible to estimate the total cost of making it. This action was assigned to [Tim Jenness](#).
-

Actions

- ✓ Jacek Becla coordinate Data Access / NCSA discussions (didn't really happen because key people were unavailable / away for ~ two months, now coordination is done through DPS-WG)
- ✓ Unknown User (npease) visit Princeton ~1 week
- ✓ Kian-Tat Lim follow up with senior mgmt wrt prioritization of making PanSTARRS data available through Qserv/SUIT
- ✓ Frossie Economou look into Russell's OSX build issue

- ✓ Frossie Economou schedule Josh to work with Fabrice on qserv Docker-based test deployment
- ✓ Frossie Economou advise pipelines (John Swinbank) how to self-maintain lsst-dev stack
- ✓ Unknown User (xiuqin) start a telecon to kick start the workspace discussion, leading to a 2-3 day workshop. Workspace has been expanded into LSST Science Platform. I think Kian-Tat Lim should be leading this effort since many teams are involved for requirement gathering, design, and implementation. The page Science Platform captures the initial definition from K-T.
- ✓ Unknown User (xiuqin) study iPlant to see if we can adopt its workspace ([DM-6663](#) - Getting issue details... STATUS is created for this action)
- ✓ Unknown User (xiuqin), Frossie Economou, Gregory Dubois-Felsmann, Fritz Mueller requirement for remote butler
- ✓ Unknown User (xiuqin), Jacek Becla, Fritz Mueller butler in Java?
- ✓ [Unknown User \(ciardi\)](#) What does user community want in the external file format?
- ✓ Simon Krughoff setup a meeting to talk about AP communication in the context of association (Data Access, Process Control, Alert Production) (KSK: Similar to the situation below. We need to sort all this out at a very fine level, but given the replan, I think we can put it off until after the AHM. Please disagree Fritz Mueller, Donald Petravick, or Andrew Connolly)
- ✓ Simon Krughoff verify division of labor in terms of defining standards, authoring packets and publishing streams in the context of the Event Production Pipeline (VOEvent and VOEvent Transport Protocol)
- ✓ Simon Krughoff add request for butler to think about how repeated access to the same data (calibration products) could be made performant
- ✓ Simon Krughoff look into adding the ability to get reference catalogs from the butler. Look into how multiple catalogs and external repositories for reference catalogs can be handled. (DM-6658)
- ✓ Tim Jenness prepare a complete description of work involved in rebranding `lsst_apps`.
- ✓ John Swinbank, Jacek Becla, Simon Krughoff Arrange a full walkthrough of the AP and DRP interaction with Data Access services including relevant members of all three teams. (KSK: I don't think this is relevant right at this moment. I think we need to finish the replan of the pipelines before this can be maximally useful. I'm checking it off, but we need to do this at some point. Please feel free to disagree Fritz Mueller and John Swinbank)
- ✓ ~~John Swinbank arrange for the Data Access team to be provided with a specification for in-Butler catalogue joins. (This is~~

[DM-6662](#) - Getting issue details... STATUS

~~; removing from the list here.)~~
- ✓ ~~Unknown User (xiuqin), John Swinbank Arrange for Firefly-as-a-service on lsst-dev. (JDS: I believe this work is being coordinated by SUIT; I have nothing to do here until they ask for support. Unknown User (xiuqin), do you agree?) (Epic~~

[DM-5591](#) - Getting issue details... STATUS

~~was created to address this issue)~~
- ✓ John Swinbank, Jim Bosch Compile short-term middleware feature requests and supply to Margaret Gelman well in advance of the next cycle.
- ✓ John Swinbank, Margaret Gelman Decide on an appropriate division of labour for updating `ctrl_pool` middleware to address Condor. (JDS: I believe this is obsolete due to the ongoing replanning exercise and discussions at the May DMLT)