

PDAC Meeting 2016-11-03

Date

03 Nov 2016

Attendees




- [Gregory Dubois-Felsmann](#)
- [Kian-Tat Lim](#)
- [Unknown User \(xiuqin\)](#)
- [Brian Van Klaveren](#)
- [Unknown User \(jalt\)](#)
- [Greg Daues](#)
- [Trey Roby](#)
- [John Gates](#)
- [Fritz Mueller](#)

Goals

- Review progress toward PDAC completion at the end of the month

Discussion items

Time	Item	Who	Notes
	NCSA / infrastructure	Unknown User (jalt)	<ul style="list-style-type: none">• Going to announce on Monday that <code>lsst-dev01</code> will be in full operation in two weeks• Looking at the sizing model this month, looking for ways to improve it<ul style="list-style-type: none">◦ Fritz Mueller: Andy Hanushevsky has been looking at it too and may have some suggestions• Looking at DRP/Princeton needs for QA• Discussion of handling 100 sq. deg. of HSC public-data<ul style="list-style-type: none">◦ Comparison of this to the Data Access group's need for further scaling tests◦ QA system appears to need the interactive creation of user databases◦ Discussion of whether the QA system can be hosted on the same cluster as the S82/WISE PDAC• All hosts in the PDAC cluster are functioning

SLAC / Data Access - Qserv	Igor Gaponenko	<ul style="list-style-type: none"> Stripe 82 catalog databases have been loaded into Qserv; simple tests seem to be working <ul style="list-style-type: none"> Have not tried performance / large-scale tests yet Lots of changes to the loading procedure were required in order to make this efficient & correct Procedure is now highly parallel Next step: produce a writeup on loading, especially for Unknown User (speckins)'s use <ul style="list-style-type: none"> Should be possible to experiment with loading additional databases without affecting S82 data Working on better file organization for the image data <ul style="list-style-type: none"> Looking at following Unknown User (jalt)'s proposal for data immutability (see <div data-bbox="467 363 1243 531" style="border: 1px solid orange; padding: 10px; margin: 10px 0;">  RFC-249 - Jira project doesn't exist or you don't have permission to view it. </div> <ul style="list-style-type: none"> Ran into some problems with respect to time stamps / time zones in creating Docker containers (see <div data-bbox="423 558 1200 726" style="border: 1px solid orange; padding: 10px; margin: 10px 0;">  DM-8119 - Jira project doesn't exist or you don't have permission to view it. </div> <ul style="list-style-type: none"> Interpreting logs when containers have a different time zone than the host OS - KTL says that <div data-bbox="467 747 1243 915" style="border: 1px solid orange; padding: 10px; margin: 10px 0;">  DM-1203 - Jira project doesn't exist or you don't have permission to view it. </div> <p>- Set single time for all</p> <p>project software & common machines Done is meant to be the current policy on log timezones and formatting; if NCSA can now support UTC-only systems, that would be preferable.</p> <ul style="list-style-type: none"> Interpretation of time stamps in the database data - KTL says that database columns are meant to be explicit in the schema (usually in the column name itself) as to which time scale they are in (TAI or UTC) and should never be in local time. But while this appears to have been followed in the baseline schema, it doesn't seem to have been written down. MySQL (MariaDB) appears to store TIMESTAMP columns as integers containing YYYYMMDDHHMMSS interpreted as UTC but they appear to ignore leap seconds and do not accept 60 as a seconds value. Unknown User (jalt): NCSA is very happy to go to uniform UTC time zone settings in production services
SLAC / Data Access - DAX	Brian Van Klaveren	<ul style="list-style-type: none"> Ready to redirect dbserve at the Qserv cluster Qserv is now up, so we should try to do this today Still working on getting metaserv up today and tomorrow John Gates made progress on imgserv and we will make some images available today John Gates: raw images and cutouts are now working <ul style="list-style-type: none"> Need Jim Bosch's changes to afw deployed in order to get codas working, may take a couple of days IPAC input: please prioritize getting Qserv going behind {{dbserve}} Everybody is OK with giving first priority to dbserve-Qserv and imgserv-raw/calexp Some discussion with Unknown User (ymei) of whether DAX is giving accurate reports of float vs. double column types <ul style="list-style-type: none"> Conclusion: prioritize getting metaserv up to produce accurate answers in preference to fixing the immediate issue
IPAC / SUIT	Trey Roby	<ul style="list-style-type: none"> Working on getting multiple types of searches attached to dbserve from the UI. Have demonstrated full round trips from the UI to the current temporary dbserve. Really need the actual Qserv service in order to continue to make progress. <ul style="list-style-type: none"> These search processors are already running on the PDAC Firefly deployment Load-balancing server is up and running <ul style="list-style-type: none"> Planning to put the DAX services behind the load-balancer as well Planning to do direct DAX queries from an IPython notebook as a proof-of-concept, with Firefly visualization in the notebook Have started the creation of the PDAC Firefly application page <ul style="list-style-type: none"> Will bring in the IRSA light curve viewer once the images are available

Action items

- ✓ Igor Gaponenko: change the Docker container production/configuration procedures in such a way that internal timezone setting of the containers were inherited from their host environment. The implementation of this will be conducted in a context of



DM-8119 - Jira project doesn't exist or you don't have permission to view

it.

. Implement this before the next meeting.