PDAC Meeting 2016-09-09

Date

09 Sep 2016

Attendees

- Gregory Dubois-Felsmann
- Unknown User (xiuqin)
- Kian-Tat Lim
- Fritz Mueller
- Brian Van Klaveren
- Igor Gaponenko
- Unknown User (jalt)

Goals

- Review readiness to begin assembling the system

Discussion items

Time	Item	Who	Notes
	SLAC status	Fritz Mueller	- One month behind schedule in getting cluster access - Igor has found all the catalog data, is generating .tsv files and should be able to start loading a test slice of the data by next week 16 Sep 2016, assuming that qserv has been able to be deployed across the cluster - Discussion of whether we have Object tables for all five bands (as opposed to just *i* band). We *do* have ForcedSource in all bands. Difficult to resolve this question because the processing configuration may not have been preserved. The provenance would have been in the output repository that contained the original output afw.table FITS files. Igor has found these files - the catalog data that has been found was in the form of loaded databases. - Igor has made progress in tracking down the calexps, starting to transfer them from IN2P3 iRODS to NCSA. Looking into how to structure the data for use in the contemporary Butler. Most likely the original Butler registry was preserved and may not need to be recreated. Igor and Brian will be responsible for ensuring that the image repository is Butlerized. The goal is to complete this work within the next 30 days. - Still believe that ~20% of the calexps were lost. - Planning to bring up slices of both catalog and image data. Will try to make sure that they are overlapping so that catalog-image mapping can be tested. - Will ask John to document the capabilities of the cutout service - Containerizations of DAX services are in progress, should finish next week (Brian Van Klavern) - Working on getting tagging, integration with eups right - Trying to start with the six-monthly release just made by Frossie - Only Object (~200M rows) and ForcedSource (~20G rows) will be partitioned; all other tables will be replicated and therefore all queries will be through the qserv master - Kian-Tat Lim: there is also an SDSS-provided table of data quality information, at per-CCD per-pointing granularity; this could be added to the set of replicated tables
	NCSA status	Unknown User (jalt)	- Systems are up (all but one), everyone should have two-factor access - Will send access instructions - Two things missing: - Filesystem mounts are not there - Admins are still implementing the access policy - Work is likely complete on Monday - SLAC group requirements: - Docker and ssh access - Will probably use shmux for initial configuration of a qserv cluster - Docker is not currently pre-installed, but NCSA will be happy to assist in the creation of a recovery image that contains it. SLAC will advise about this once initial testing is done.
	IPAC status	Unknown User (xiuqin)	- Starting to work on understanding the table schemas (see case03) and how the semantic connections across tables can be made - Would like to do wget-level testing of the DAX services as soon as they are up - Ready to start trying to running basic test servers on the new machines as soon as we know their names and can log in - Will self-install a recent Java 1.8 from Oracle (have not tested the very latest point release), then will ask to have it added to the recovery image. - Currently developing on Tomcat 7, may use either 7 or 8, will self-install

Action items

- ☑ Fritz Mueller Apply "pdac" label to PDAC-related stories and epics in JIRA for Data Access 16 Sep 2016 (**note that JIRA labels are case-sensitive**)

- ☑ Unknown User (xiuqin) Apply "pdac" label to PDAC-related stories and epics in JIRA for SUIT 16 Sep 2016
- ☑ Unknown User (jalt) Apply "pdac" label to PDAC-related stories and epics in JIRA for infrastructure and service deployment at NCSA 16 Sep 2016
- ☑ Gregory Dubois-Felsmann Create a dedicated HipChat room for PDAC discussions 14 Sep 2016