

Data Management Monthly Report

March 2017

High-level Summary

There was good attendance at the Joint Technical Meeting in Pasadena March 6-8th. This included the first presentation by William O'Mullane (WOM) to the DM assembly on the direction of DM.

At JTM a discussion on the first Camera raft data was held and post JTM a module for the camera has been created to load this data in the LSST stack.

Also at JTM the Prototype Data Access Centre (PDAC) hosted at NCSA was shown off to much acclaim. Work continues to further refine the PDAC especially in the area of the Science Platform. This will allow some scientific validation of the platform.

Work continues on dealing with the heterogeneous focal plane. The "Photometric Redshifts with the LSST: Evaluating Survey Observing Strategies with a New Nearest-Neighbors Estimator" (Graham et al.) manuscript is ready and awaiting the publication board process.

Project Status

Wil O'Mullane stayed on after JTM with Krabendam to visit IPAC and discuss detail planning and scope, this was a good session clarifying roles and guidelines for IPAC helping to get costs under control.

The initial work on the camera module starts work on the code needed for the commissioning camera and should help define this interface for DM.

Selected technical accomplishments:

- The Architecture Team contributed, through activities at the Joint Technical Meeting and subsequent meetings and writings, to defining interfaces, system components, and verification activities. In addition, the infrastructure for writing and publishing LaTeX-based DM documents was substantially enhanced.
- The AP Team
 - Implemented a significantly improved matching algorithm for single chip to reference catalog

- Significant progress in the area of Wcs, XYTransform and coordinate transforms in general
- Plugged jointcal into the SQuaRE validation system
- The DRP Team
 - Completed transitioning the technology used to interface C++ code with Python in the DM Stack from [Swig](#) to [pybind11](#). This new system is simpler, more flexible, and easier to maintain: it represents a major step forward in the usability of our codebase.
 - Investigated and resolved a major regression in weak lensing science measurements, representing a major success for our QA efforts.
 - Completed the first fully functional prototype of the “donut” fitting pipeline. This enables us to measure the wavefront, and hence estimate the PSF, across the focal plane.
- The SUIT Team
 - Deployed a new version of PDAC at NCSA with updated time series function.
 - Prototyped Jupyter notebook and Firefly server connection to JupyterHub using PDAC, running LSST stack to do forced photometry calculation.
- The DAX Team
 - Continued support of integration activities at the PDAC via problem investigations and bug fixes.
 - Completed a significant refactor of the Data Butler camera mapper so that it no longer accesses repository data by making filesystem calls directly.
- The Net&BS Team
 - La Serena - Santiago Fiber: REUNA delivered the report that documents the completion of the long haul fiber with the measurements of length and attenuation of the fiber in each segment. These measures were sent to Coriant, for them to verify the equipment design, and the result was positive. It is relevant to note that as the fiber is brand new the attenuation parameters are of a very good quality, less than the theoretical parameters.
 - DWDM Equipment: The AURA requisition went out to Coriant for the DWDM equipment. Expect arrival of the equipment at the end of May, when we will have 30 days to inspect the equipment before payment. REUNA sent the formal announcement to Coriant, stating they are the awarded solution.
 - FIU CIARA hosted the first monthly meeting of the LSST Network Engineering Team (NET), which is the cross-institution team of network management and technical staff responsible for designing, engineering, deployment, and supporting operations of the LSST long-haul networks.
- The NCSA Team

- Thirteen NCSA project team members attended the 2017 Joint Technical Meeting in Glendale, California.
- Members of the Process Middleware team met with developers from USC to discuss use of the Pegasus workflow management system for LSST Batch Production Services.
- We completed the Operations Plan (LPM-181) for the LSST Data Facility elements.
- The SQuaRE Team
 - Made a major SQuaSH release into production
 - Set up a binary distribution server and is beta-testing it
 - Completed the 13.0 release activities

Risk Management

The DM Risk Register was reviewed in the monthly process. No new risks were added and no significant changes to existing risk exposure were made.

Detailed Project Progress and Status

LSST Program Office

DM Project Management and Control

All milestones from the DM plan were integrated into a master plan in Tucson. Old milestones were merged or superseded accordingly. On first inspection this seems to tally well with guidelines for the plan set out initially by Juric.

Current accomplishments:

The DM Subsystem Lead and the DM Project Manager:
New DM project manager will start April 3.

Planned activities:

The DM Subsystem Lead and the DM Project Manager will prepare for the AMCL meeting in Washington. Bring together current ETC for DM and present it there. The other main areas of

work are to update the DM project management plan while in parallel defining the Verification and Validation Scheme for DM. The latter will shape the document pack for the Review in July.

DM Science

Current accomplishments:

DM Subsystem Science staff has completed the work on the study of impacts of the heterogeneous focal plane to DM image differencing development (projecting no additional impact). We continue to work on two ongoing studies (photometric redshifts, and star/galaxy separation), science validation of the Prototype DAC, and organizational activities.

The "Photometric Redshifts with the LSST: Evaluating Survey Observing Strategies with a New Nearest-Neighbors Estimator" (Graham et al.) manuscript is ready and awaiting the publication board process. The write-up of the star-galaxy separation paper (Slater et al.) continues. Work has begun on the review of DM requirements and data products for special programs (e.g., the Deep Drilling fields).

We are performing the science validation of the Prototype DAC (PDAC) environment (Suberlak et al.). A draft report on the PDAC has been shared with the PDAC team, with a technical note forthcoming.

Initial draft of the Science Platform document and the Science Validation document have been put into place (Juric et al.). These now await input by other members of the DM team.

The work has finished on the impact of hybrid (multiple CCD vendors) focal plane on image differencing false positive rates. A report (Slater, Bellm, et al.) has been finalized is being reviewed by the Pipeline Scientist. The key conclusion is we do not expect additional impact from the selection of the hybrid focal plane.

DM System Science Team establishment and organization has been completed. The first staff meeting was held on March 17th, with a meeting of the broader team via telecon on March 24th. The Confluence page for the team can be found at <https://confluence.lsstcorp.org/x/cgBFAw>.

Planned activities:

Continue work on requirements for Special Programs. (Graham)
Complete and submit the photo-Z paper. (Graham)
Complete the first phase of PDAC science validation. (Suberlak)

Deliver the first version of the Science Platform document. (Juric et al.)

Support the LSST:UK meeting. (Juric et al.)

DM System Engineering

Current accomplishments:

DM System Architect and Deputy System Architect:

- Attended, led and moderated sessions at, and reached agreements at the Joint Technical Meeting.
- Contributed to Operations Planning documents.
- Began a series of interviews with stakeholders to understand production compute environment provisioning and deployment requirements.
- Conducted a “seminar” on the history and structure of the current DM sizing/cost model in preparation for updating it.
- Prepared a change request covering the updated LSE-61 and matching changes to LSE-163. The change request has not yet been submitted pending further review.
- Significantly enhanced the new LSST LaTeX document styles to match the look and feel of standard LSST Word documents. This included work on references and handling documents subject to change control.
- Made some cleanups to the LDM-151 draft and created DMTN-038.

DM System Interfaces Scientist:

- Attended, led and moderated sessions at, and reached agreements at the Joint Technical Meeting.
- Led the gathering of input for and initial decisions of the SuperTask/Butler Working Group.
- Helped define goals and tests for Early Integration (OCS-CCS-TCS-DM) activity to be held in April at NCSA.

Planned activities:

DM System Architect and Deputy System Architect will:

- Complete the provisioning and deployment interviews.
- Iterate further on the LDM-148 DM design document.
- Help with submissions to the LSST CCB relating to processed visit images, previous data releases and storing of heavy footprints.
- Discuss missing flowdown of requirements from LSR with Systems Engineering.

- Document the LSST LaTeX system.
- Attend MagicDraw SysML training.

DM System Interfaces Scientist will:

- Attend and assist the Early Integration (OCS-CCS-TCS-DM) activity at NCSA.
- Attend MagicDraw SysML training.

DM Science Quality and Reliability Engineering (SQuaRE)

Current accomplishments:

Last CI job on April 17, 2017, 12:34 a.m.	
# of metrics being monitored:	9
# of LSST packages being monitored:	83
# of CI jobs:	754
# of measurements:	3308
Supported datasets:	cfht, decam, hsc

- A major SQuaSH release was deployed into production. This includes months of development including the “drill-down” views into selected metrics and rolls in the extensive `validate_drp` refactoring. It even includes a metrics display on the metrics harness! (see above) Note: we have major noted issues with very new versions of Safari dropping the bokeh websockets. We have spent some time investigating them but unless we can successfully isolate them to the satisfaction of Webkit developers we may have to drop Safari from the supported list.
- The documentation for the 13.0 release was completed and the release was announced to the community. Many thanks to all who contributed.
- As part of the binary release improvements we now have a binary distribution server that receives eups binary tarballs from the CI system in an automated fashion. A more complex directory structure that the source-only server had to be developed. In general this works well; it has not been released due to some issues with OSX binaries and final user-level documentation updates. The new site is in testing as `eups.lsst.codes`

- In the SQuaSH prototype, we had been passing parameters via the Referer HTTP header. This is blocked by some browser advanced privacy setting so we now pass our own custom headers around.
- A micro-service that indexes technotes, an important part towards DocHub, was deployed. These are now pushed in quasi-raw format to www.lsst.io - look-and-feel work still needs to be done but is de-prioritised for now
- We wrote a python logging class to ship logs to SQuaRE's logging service (logging.lsst.codes) and converts some of our microservices to using
- LSST branding and "edit on Github" to facilitate contributions was added to the developer guide (developer.lsst.io) and the stack release documentation (pipelines.lsst.io)
- Lsst_ci now runs obs_cfht and obs_decam examples via scones (see [DM-9740](#) for details)
- Microservices were ported to uwsgi (from standalone flask) to improve handling of concurrent requests
- We now sync Travis CI with Github more often to minimise the elapsed time that it takes for new Travis CI jobs to be set up (eg. when new technotes are created)
- Participated in the Joint Technical meeting and supported the hack activities there (such as splitting off lsst_sims into its own CI job to "protect" it from the pybind11 breaking merge)

Planned activities:

- Continue on the metrics framework development to support in-task metric collections (using jointcal as the testcase)
- Evaluate Jupyterlab via prototype activities with an eye to Science Platform activities
- Continue productizing automated binary distribution

Recruitment update:

None planned

University of Washington

Current accomplishments:

02C.03.00 - Alert Production Management Engineering and Integration

All team members prepared for and attended the JTM in Glendale CA.

Rawls worked on the AP prototype validation.

Hiring

No activity.

02C.03.01 - Single Frame Processing Pipelines

Findeisen worked to translate the current 2-D to 2-D generic coordinate transform framework to use the new astshim backed system (DM-9595, DM-9597, DM-9826, DM-9598).

Owen also worked on the transforms system (DM-8439 and DM-9846).

Morrison implemented a prototype matcher that is much improved in many ways over the current matcher (DM-8113).

02C.03.02 - Association Pipelines

No advancements in this area.

02C.03.03 - Alert Generation Pipelines

Patterson drafted the alert distribution testing and scaling plan for the distribution system (DM-8821).

02C.03.04 - Image Differencing Pipeline

Reiss implemented the classes necessary for doing spatially varying decorrelation with applications to the image differencing algorithm (DM-9629; DM-7611).

Sullivan finished writing and testing a prototype version of the DCR corrected template generation algorithm (DM-8801).

02C.03.05 - Application Framework for Exposures

Rawls supported her AP validation work by improving the way the DECam instrument handles calibration products (the current prototype uses data from DECam; DM -5466).

Parejko ported jointcal to use pybind11 for wrappers (DM-9187). Parajko also added an integration interface to the classes that will be used for the spatially varying photometric calibration (DM-9882).

Findiesen and Owen did significant bug handling and general improvements. Highlights are: fixing build problems, improved interfaces, documentation improvements and sanitization of API interfaces (DM-7913, DM-9966, DM-7892, DM-9808, DM-9899, DM-9946, DM-9895, DM-9758, DM-9925, DM-9863).

Patterson updated the docker images for the alert stream system to take advantage of advancements in underlying technology (DM-9802).

02C.03.06 - Moving Object Pipelines

No advancements in this area.

02C.03.07 - Photometric Calibration Pipeline

No advancements in this area.

02C.03.08 - Astrometric Calibration Pipeline

Parejko added metrics for jointcal to the SQuaRE validation framework and set them to run automatically on lsst continuous integration (DM-9542, DM-9752).

Planned activities:

02C.03.00 - Alert Production Management Engineering and Integration

Morrison will attend the DESC hack week at Fermilab.

Rawls will continue working on the prototype AP system.

02C.03.01 - Single Frame Processing Pipelines

Owen and Findeisen will continue work on the new transform framework. Specifically, they will be replacing old versions of the system with new ones.

02C.03.02 - Association Pipelines

No planned work.

02C.03.03 - Alert Generation Pipelines

Patterson will be looking into stream processing technologies that work the current ecosystem of tools.

02C.03.04 - Image Differencing Pipeline

Reiss will finish implementing spatially varying ZOGY to test against the current stack baseline algorithm.

Sullivan will finish writing up the proposed DCR correction algorithm as a DMTN

02C.03.05 - Application Framework for Exposures

Bug fixes and technical debt.

02C.03.06 - Moving Object Pipelines

No planned work.

02C.03.07 - Photometric Calibration Pipeline

Parejko will work on improving the spatially variable photometric model in jointcal.

02C.03.08 - Astrometric Calibration Pipeline

Morrison will execute large scale tests to determine if the improved matcher will handle the dynamic range we need for precursor datasets, out of the box.

Recruitment update:

N/A

Princeton University and University of California, Davis

Current accomplishments:

02C.04.00 - Data Release Production Management Engineering and Integration

- Travel & meetings:
 - The DRP team joined the rest of the LSST staff at the 2017 [Joint Technical Meeting](#) in Glendale, CA. This meeting provided useful opportunities for coordination and collaboration with the wider project. Particularly useful was the opportunity to introduce members of the DRP team to the science raft data, discussions on the future of the “Jointcal” calibration tool, and co-working in the Hack Sessions.
- Data Management replan [[DM-8494](#)]:
 - The revised Calibration Products Production section of [LDM-151](#), previously reported on in January, has now been reviewed and merged into the working draft of the document. [[DM-8066](#), [DM-9131](#), [DM-9132](#), [DM-9136](#)]
- Quality assurance [[DM-8299](#)]:
 - Members of the DRP team have been cooperating with colleagues at NCSA to plan the reprocessing of the first Hyper Suprime-Cam (HSC) public data release on the LSST Verification Cluster. We expect this work to be undertaken during the remainder of the S17 cycle. [[DM-9886](#)]

- A regression in the correction being applied for the Brighter-Fatter effect was observed during QA analysis of using the DM stack to process HSC data. This regression caused discrepancies of up to 1% in the trace radii being measured for sources, which has a significant impact on weak lensing science applications.
- The problem was traced to a new algorithm being used to compute second derivatives. The old algorithm was restored to resolve the regression. Further work on the Brighter-Fatter effect is scheduled for later in construction, and this will be investigated in detail then. [[DM-8533](#), [DM-9535](#), [DM-9553](#)]
- An error was identified in the way that the meas_mosaic calibration package is handling the WCS for rotated CCDs. This affects about half the HSC focal plane, rendering it much less useful as a QA tool. Although the issue has been isolated, a fix has not yet been verified at time of writing. [[DM-9383](#), [DM-9862](#)]
- Task documentation prototype [[DM-8559](#)]:
 - This activity aimed to study the topic of how best to document Tasks, which form the building blocks of the core Science Pipelines logic in the DM stack.
 - A “topic template” and accompanying notes describing how to document a Task has now been produced, together with a number of example documents. This marks the successful completion of this work. [[DM-8717](#), [DM-8723](#)]
 - Some technical work remains to be tackled by the SQuaRE team before these can be deployed in practice.
 - The member of staff who was working on this activity has now left the project.

02C.04.01 - Application Framework for Catalogs

- Middleware evolution [[DM-1109](#)]:
 - The DRP team has continued to participate in the the SuperTask Working Group, taking part in discussions and contributing [a document](#) describing the changes in data units and/or parallelization axis which will be encountered while producing a data release. [[DM-9604](#)]
- Implementation of new Footprint system [[DM-3559](#)]:
 - Work has continued through this month to convert all existing Science Pipelines code to use the new Footprint system. [[DM-8108](#)]
 - This work is expected to merge to master in early April.
- Conversion from [SWIG](#) to [pybind11](#) [[DM-7717](#)]:
 - During the “hack session” at the Joint Technical Meeting the conversion of all of the released parts of the DM stack to pybind11 was completed.
 - This work has now been merged to master; all stack releases from w_2017_11 onwards depend on pybind11.

- This was accompanied by documentation describing the new system being added to the [Developer Guide](#). This includes both a tutorial introduction to wrapping code with pybind11 and a set of guidelines to ensure that all DM use of pybind11 will be consistent.
- This completes the scheduled pybind11 work package.
- Some minor ongoing work is expected to make our usage of pybind11 idiomatic and optimized for both performance and size across the whole stack. This will be included in our regular technical debt reduction activities.
- Emergent work and reduction of technical debt [[DM-8136](#)]:
 - A series of minor improvements were made across the codebase during this period. Highlights include:
 - Updated the software versions available in the “shared” software stack used for development. [[DM-9682](#), [DM-9639](#)]
 - Resolved build warnings in afw. [[DM-9531](#)]
 - Updated coding guidelines. [[DM-9087](#)]

02C.04.02 - Calibration Products Pipeline

- Auxiliary telescope development [[DM-8151](#)]:
 - Work has focused on producing a camera (or “obs”) package, which enable us to ingest the test data taken on the 0.9m telescope at CTIO into the LSST software stack. [[DM-9363](#), [DM-9991](#), [DM-9992](#)]
 - This activity has dovetailed with the parallel implementation of a camera package to provide access to camera test stand data (see below).
- Access to camera test stand data [[DM-8292](#)]:
 - Following a successful session at the Joint Technical Meeting, work has focused on enabling easy access to camera test stand data from the DM stack. This has involved making the data available on the LSST owned compute hardware at NCSA and constructing a camera package which can load it into the stack. This camera package will ultimately scale to act as our interface to ComCam (the commissioning camera). [[RFC-301](#), [RFC-304](#), [DM-9400](#), [DM-9872](#), [DM-10006](#)]
- Photometric Calibration [[DM-8276](#)]:
 - A wavelength-dependent global-photometric fit will be necessary for testing and commissioning pipelines before the full calibration system becomes available. This will be based upon the Forward Global Calibration Model, which is being upgraded for use with LSST. Modernization of this codebase is underway with an expected completion date of April. [[DM-9558](#)]

02C.04.03 - PSF Estimation

- Wavefront measurement and PSF reconstruction [[DM-1909](#)]:
 - In order to fit “donut” images, it’s necessary to understand the details of the projection of the spider onto the CCD, and hence the state of the telescope rotator. This is now being extracted and stored from HSC images when they are processed. The machinery exists to do the same for other instrumentation, but that has not yet been implemented (and is not required to completed this work package). [[DM-9111](#)]
 - A detailed characterization of the HSC instrument has been completed and incorporated into the “donut” fitting pipeline. This includes focal length, pupil diameter and vignetting. [[DM-8589](#)]
 - Based on the above, a complete prototype of the “donut” pipeline has been completed. It includes the ability to process donuts taken at any rotator angle, and both intra- and extra-focal donuts. Upcoming work will translate this prototype pipeline into the regular DM stack framework. [[DM-8565](#)]

02C.04.04 - Image Coaddition Pipeline

- Research properties of covariance under coaddition [[DM-6176](#)]:
 - Due to the departure of the member of staff who was working on this project, no further work is expected in this work package. The results achieved to date represent a useful step toward an upgraded coaddition system which will be implemented in future cycles.
- PSF homogenization [[DM-8289](#)]:
 - High-quality PSF matched coadds can now be produced by the DM stack. An example is shown in the figure. Minor code cleanliness and integration work remains before this work package is completed.



An example PSF-matched coadd produced by the LSST stack. This is a subsection of the COSMOS field as observed by Hyper Suprime-Cam in the g, r and i bands. This image approximates the depth of the LSST survey. Data from the Hyper Suprime-Cam Strategic Survey Program.

02C.04.05 - Object Detection and Deblending

- Deblender development [[DM-8140](#)]:
 - Development and testing of the Non-negative Matrix Factorization (NMF) deblender has continued through this period.
 - Building on the excellent results previously reported, efforts have focused on testing and validation of the work. This has included developing a suite of validation metrics and producing a sets of both real and simulated data which can be used for testing. [[DM-9172](#), [DM-9640](#), [DM-9169](#)]
 - Algorithmic improvements include a 2 dimensional PSF convolution operator and improvements to the symmetry and monotonicity operators used by the deblender. These have been incorporated into a working deblender based on an alternating least squares procedure. [[DM-9170](#), [DM-9171](#), [DM-9561](#), [DM-8675](#)].

02C.04.06 - Object Characterization Pipeline

- Enhancements and upgrades to CModel flux measurement [[DM-1111](#)]:
 - An error in the weighting of CModel priors was corrected, making the priors much more informative. In addition, a configuration option was added so that the relative weighting of the prior may be tuned to meet particular user requirements. [[DM-9795](#)]
- Emergent work and reduction of technical debt [[DM-8306](#)]:
 - A series of minor improvements were made across the pipelines during this period. Highlights include:
 - Corrected propagation of celestial coordinates on coadds. [[DM-9556](#)]

- Made it possible to use rectangular (rather than square) bin sizes when estimating backgrounds. [[DM-9828](#)]
- Upgraded the FlagHandler system, which is used for recording errors encountered when running measurement algorithms, to make it more reliable and robust against programmer error. [[DM-9249](#)]

Planned activities:

02C.04.00 - Data Release Production Management Engineering and Integration

- Work with the new DM Project Manager to complete the DM replan.
- Continue large scale QA work, in particular:
 - Resolve issues surrounding rotated CCDs in meas_mosaic;
 - Work with the NCSA team to bootstrap processing of HSC data on the LSST Verification Cluster.

02C.04.01 - Application Framework for Catalogs

- Complete the merge of the new Footprint system to master.
- Continued participation in the SuperTask Working Group.

02C.04.02 - Calibration Products Pipeline

- Deploy camera packages capable of handling both CTIO 0.9m and camera test stand data.
- Begin conversion of the prototype CTIO 0.9m spectroscopic pipeline to DM stack conventions.

02C.04.03 - PSF Estimation

- Convert the prototype “donut” fitting pipeline into a fully integrated part of stack processing, using our regular code primitives and conventions.

02C.04.04 - Image Coaddition Pipeline

- Complete final merge of all PSF-matched coaddition work.
- Start work on warped image comparison; that is, an upgraded routine for rejecting artefacts during coaddition.

02C.04.05 - Object Detection and Deblending

- Continued testing and optimizing of the NMF deblender.

02C.04.06 - Object Characterization Pipeline

- Start work on tools for visualizing the results of CModel measurement.

Staffing update:

- Vishal Kasliwal and Mandeep S. S. Gill have both left the project.

IPAC / California Institute of Technology

Current accomplishments:

02C.05.00

- Continued to work with IPAC IRSA group on collaboration in Firefly development, plan and schedule coordination. IRSA started the time series tool testing and gave good feedback.
- Continued to work on SUIT requirements and resource estimates to work packages, send a revision of work packages to Kevin Long.
- Attended JTM at Glendale, CA
- Science Platform and SUIT technical, priorities, and budget discussions with Wil, Mario, and Victor at IPAC after JTM.
- Hosted the tiger team meeting led by Jeff Kantor in IPAC after JTM.

02C.05.01 Basic Archive Access Tools

- Published the documentation guide for JavaScript using JSDoc
<https://developer.lsst.io/docs/jsdoc.html>
- Build Firefly Python API doc using Sphinx
- PDAC:
 - Time series data display issues found in test.
 - Launch the time series data view from the tri-view
 - Fixed a bug in table download when user chooses columns to save.
 - Deployed updated version. Please see access guide here:
<https://confluence.lsstcorp.org/display/DM/Guide+to+PDAC+version+1>
 - Prototype Jupyter notebook and Firefly server connection to JupyterHub using PDAC, running LSST stack to do forced photometry calculation

- Java code refactoring and adding unit tests.
- Preliminary research of using embedded database in Firefly for table data handling, DM-9888. Report is here:
<https://confluence.lsstcorp.org/display/DM/Research+database+options+to+support+table+display+functions>.

02C.05.02 Data Analysis and Visualization Tools

- Improvement and bug fixes in scattered plot.
- Improvement in Firefly Python API firefly_client to handle multiple base URLs.
- Researched Plotly.js library to replace Highcharts, and started working on it.
- New features for 2D charts: add asymmetric error for scatter plot, more options for histogram.

02C.05.03 Alert/Notification Toolkit

- No new work done.

02C.05.05 User workspace

- Block concept design diagram for workspace as part of science platform with JupyterHub. DM-8164

Planned activities: (for Apr., 2017)

02C.05.00

- Continue to work with IPAC IRSA group on collaboration in Firefly development, plan and schedule coordination.
- Support NCSA application interview.

02C.05.01 Basic Archive Access Tools

- Continue working on SUIT design and design document as part of science platform design.
- Junit tests for Java classes.
- PDAC development:
 - Multi-object catalog search
 - Improvement in time series data display UI with user feedback
 - Search and display WISE data in PDAC

02C.05.02 Data Analysis and Visualization Tools

- Improvement to UI due to feedback from users
- New functions: Better control of chart data and charts type
- Replace Highcharts library with Plotly.js
- Support Camera visualization work by attend the weekly meeting

02C.05.03 Alert/Notification Toolkit

- No work planned

02C.05.05 User workspace

- No work planned

Recruitment update:

- No activities

SLAC / Stanford University

Current accomplishments:

02C.06.00 Science Data Archive and Application Services Management Engineering and Integration

- Entire team attended 2017 JTM in Glendale, CA
- Team continued improving robustness of DAX services and otherwise supported SUIT integration efforts in the PDAC as necessary
- Thukral continued gathering 20% DR1 KPMs on the IN2P3 Qserv cluster. Mueller and Thukral investigated and addressed a column statistics issue with the new 20% DR1 dataset which had been negatively impacting certain KPMs
- Mueller and Hanushevsky had vacations during this month

02C.06.01.01 Catalogs, Alerts and Metadata

- Salnikov continued experiments with L1/AP db prototype, now working with up to 1 month's worth of simulated data. Performance continues to drop linearly per expectation. [DM-7182]

02C.06.01.02 Image and File Archive

- Lo implemented coadd cutouts in the imgserv service for use by SUIT in the PDAC.[DM-9631]

02C.06.02.01 Data Access Client Framework

- Pease completed a Data Butler refactor allowing child repositories to perform registry lookups in parent repositories using information passed in initialization parameters instead of direct parent access. [DM-8686]
- Pease completed a refactor of the Data Butler camera mapper so repository data is accessed via a storage abstraction instead of direct POSIX filesystem calls. [DM-8125]
- Pease addressed some concurrency issues in the reworked Data Butler camera mapper that were exposed by obs_subaru unit tests in CI. [DM-9848]
- Pease addressed warning generated incorrectly in some circumstances during Data Butler initialization. [DM-9669]
- Salnikov and Pease continued their participation in the SuperTask Working Group.

02C.06.02.02 Web Services

- No work was carried out in this WBS element this month

02C.06.02.03 Query Services

- Gates utilized some recently added XRootD features to implement threading and throttling improvements in Qserv which help keep returned large results from choking outbound interactive queries. [DM-9321]
- Gates implemented some query templating improvements which increase query dispatch performance. [DM-9670]
- Gates addressed a performance issue related to dead job tracking that was brought to light by the new WISE dataset in the PDAC, which has a much higher chunk count than had previously been used/tested. [DM-9672]
- Gates addressed an issue with secondary-index lookups on non-integer director keys that was brought to light by the WISE dataset in the PDAC. [DM-9706]
- Jammes continued investigations/prototyping of Kubernetes for Qserv container deployment and management on the IN2P3 clusters. [DM-9019, DM-9776]

02C.06.02.04 Image Services

- No work was carried out in this WBS element this month.

02C.06.02.05 Catalog Services

- No work was carried out in this WBS element this month.

Planned activities:

02C.06.00 Science Data Archive and Application Services Management Engineering and Integration

- Team to continue improving robustness of DAX services and otherwise support SUIT. integration efforts in the PDAC as necessary.
- Thukral to write up 20% DR1 KPM results.
- Gates has a vacation scheduled this month.

02C.06.01.01 Catalogs, Alerts and Metadata

- Salnikov to continue L1/AP db prototype investigations.

02C.06.01.02 Image and File Archive

- Lo and Van Klaveren to work with SUIT team to revise/refine DAX service APIs.

02C.06.02.01 Data Access Client Framework

- Pease to begin work on db storage support for Data Butler.
- Salnikov and Pease to continue their participation in the SuperTask Working Group.

02C.06.02.02 Web Services

- Van Klaveren to reload and deploy metaserv in the PDAC.

02C.06.02.03 Query Services

- Gates to address various Qserv bugs uncovered by ongoing PDAC and KPM efforts.

02C.06.02.04 Image Services

- No work is scheduled for this WBS element this month.

02C.06.02.05 Catalog Services

- No work is scheduled for this WBS element this month.

Recruitment update:

No activity.

NCSA / University of Illinois

Current accomplishments:

02C.07.00 Processing Control and Site Infrastructure Management, Engineering and Integration

Planning and Management

These activities involve the creation of project plans for the construction phase for the DM-wide replan, and planning for the commissioning and operations phases. NCSA leads planning for operations of the Data Products Production (DPP) department.

We refined the WBS ITC and foundational services section, working with representatives from each major cyberinfrastructure group at NCSA (storage, security, systems, data and databases, networking) to assess completeness of WBS and effort estimates. We did preliminary scheduling of the WBS, considering together the construction, commissioning and pre-operations staffing profile as services are constructed, come online, and are operated. The schedule follows project milestones where they are defined (e.g., the arrival of instrument test stands and phases of instrument commissioning), and relies on experience and engineering judgement where no milestones have been defined by the project (e.g., operational dress rehearsals, data challenges, verification tests, etc.).

We continued refinement of the operations plan document (LPM-181), including labor and non-labor costs, for the LSST Data Facility at NCSA. The operations plan document was complete on March 31st.

We participated in Joint Coordination Committee meetings with CC-IN2P3 to discuss technical aspects of inter-site operations and CC-IN2P3's potential role in disaster recovery. We prepared for and attended the Joint Technical Meeting, and prepared for the AMCL meeting in collaboration with the incoming DM project manager.

We performed preliminary investigations of elastic compute use cases and alternative disaster recovery implementations, as well as investigation into US network peering costs.

DM-9656 Replanning, Commissioning, Operations Planning (S17b)

Service Management and Monitoring

The Service Manager develops and maintains the service catalog, including service-level targets, service-level agreements, and principal technical and managerial liaisons. This includes monitoring and managing availability, capacity, and IT continuity of services, and other operational matters, as well as forming a definitive opinion about the satisfaction of each customer for the services provided. The service manager supports deployment of services to operation and service transitions.

We continued refining the prototype Service Catalog and associated processes, completing review of initial SLA templates, beginning to develop ITSM processes for incident response and change management, and continuing to map high-level services to internal foundational services on which they. We also trained on using MagicDraw, which is replacing Enterprise Architect (EA) as the LSST Systems Engineering's tool for MBSE project-level architecture models and requirements documents. We ported the Data Facility service models from EA to MagicDraw.

DM-8515 Service Management Definition and Refinement

DM-9657 Service Management Definition & Refinement (S17b)

DM-9658 Enterprise Monitoring Design - Phase 2

Hardware acquisition and provisioning planning

These activities involve high-level planning and design of physical systems related to computation, storage, networking, administration, and IT security on which DPP services are running or will be deployed. This includes consideration of near-term construction phase needs, as well as preparation for commissioning and operational needs.

We continued work to design and implement the data backbone hardware, proposing a proof-of-concept set of technologies and hardware. The first procurement activity, based on the FY2017 capability acquisition strategy, was sent to AURA for final agreement before purchasing.

Planning for system deployments in Chile this summer, which will support network-based security services and authentication and authorization services, began with investigating and documenting international shipping requirements and procedures.

DM-8512 Planning Deployment in Chile

DM-8511 Degenerate Backbone Design and Implementation

DM-8506 Costing Model Redesign

Hiring:

Scott Schreiner was hired as an IT Systems Engineer, and will help construct and operate compute systems that support development and production services at NCSA.

02C.07.01 Processing Control

Batch Production Services

These activities support the construction of services that enable the production of data products in a batch environment (e.g., Level 2 data products). The batch production services will execute processing campaigns on computing resources to produce the desired LSST data products, where campaigns are defined as sets of pipelines (ordered ensembles of computational steps), inputs they are being run against, and methods handling their outputs.

We continued refining the Batch Service codebase by tracking changes to the CameraMapper and Butler APIs. We began genericizing the design to enable specification of workflows using a standard graphical language, which can be interpreted by the Pegasus workflow management system.

We continued preparations for processing HSC data by documenting the proposed processing and submitting the plan for approval, preparing the workflow tasks, and staging datasets.

We also participated in the SuperTask/Butler Working Group, providing input on technical and operational requirements of the production processing services.

DM-8332 Run selected science pipeline using Pegasus

DM-8333 Run selected existing science codes with initial Batch Production Service façade

DM-9662 Workflow creation improvements and expansion

DM-9664 SuperTask Working Group Activities

Level 1 Services

These activities support the construction of services that will support Observatory Operations and enable the production, transport, and archiving of Level 1 data, including the nightly stream of images and events collected and processed in near real time.

We continued integrating the existing prompt processing system orchestration code with the newly developed front end elements for Level 1 services, including enhancements for OCS bridge components. We continued working with Telescope & Site developers to identify sources and contents of the EFD and OCS metadata available to build headers, and began integrating this information into test code. We also began preparing for hosting the Pathfinder Integration Activity #3, which will take place at NCSA in April.

DM-8317 Enhancements for OCS Bridge components

DM-9659 Header construction and distribution architecture

DM-9661 Reintegration of archiving and prompt processing

Common Middleware and Other Tasks

These activities consist of developing and maintaining general processing control software components, including common software used in production processing and by DM developers, as well as emergent unplanned support tasks.

Work in March involved improvements and fixes to `ctrl_*` orchestration packages, as well as migrating orchestration documentation to Data Management user guide format.

DM-9663 Emergent Middleware Work (S17b)

02C.07.02 Infrastructure Services

Data Backbone

These activities support the construction of services which archive, catalog, and distribute data to compute resources and data access endpoints across all sites—ingestion, file management, database hosting, etc.

We continued to work with SLAC and the EFD architect to detail the operational requirements of hosting identified databases across the NCSA and Chilean sites, as well as provenance and metadata catalogs that are required for implementing file ingestion, distribution, and retention services within the Data Backbone.

DM-7645 Prototype bookkeeping and interfaces for the data backbone

02C.07.03 Environment and Tools

QA Environment

These activities provide environments and tools to support general DM developer activities.

In March, we resumed work to provide access to datasets through a Jupyter hub instance on the NCSA Nebula cluster to support upcoming QA activities. We created a Docker image for the Nebula system, deployed a JupyterHub instance, added authentication using LSST LDAP, and began working toward persistent access using shared file systems.

DM-8330 Instantiate VM with Jupyter Hub and GPFS access

02C.07.04 Site Infrastructure

Hardware Acquisition, Deployment and Provisioning

These activities consist of the acquisition, provisioning, configuration, maintenance, and decommissioning of physical capabilities and associated systems which support LSST services running at NCSA.

The move from NFS to GPFS file systems continued with a resolution to user issues with quotas and file migration, making a final backup of the /home partition, and beginning to identify remaining services and non-curated data. Disaster Recovery design and implementation continued, with monitoring of data transfer. We also began prototyping a framework for ITC-level and service-level monitoring.

Hardware acquisition for the L1 Test Stand system continued with the procurement of a 10G network card which will be needed for installing and supporting the Camera DAQ system from SLAC.

DM-8503 Migrate Users from NFS

DM-8504 Migrate Services from NFS

DM-8502 Disaster Recovery for Science Datasets

Service Management for LSST development resources, Emergent and Miscellaneous Work

These activities consist of supporting and communicating with users of the current services provided by NCSA to LSST DM developers.

Work for March included addressing user- and developer-driven requests and issues related to the GPFS file system transition, as well as preparing the Nebula systems for hosting Integration Activity #3.

DM-9666 Service Management & Emergent Work (S17b)

Planned activities:

02C.07.00 Processing Control and Site Infrastructure Management, Engineering and Integration

Planning and Management

For the DM-wide replanning effort, we will continue investigation of value engineering options, and integration with commissioning and early operations plans. We will also break down the WBS into planning packages in preparation for the DM review. We will continue to refine the commissioning plan within the commissioning planning team as needed.

We will construct a capacity model for the pre-operations phase of the project based on use cases of archiving raw data from test stands and from instruments as they are commissioned, as well as processing use cases for continued development support, service integration, and production. Requirements from DM, the commissioning team, pre-operations plan, and EPO will be gathered and synthesized.

DM-9656 Replanning, Commissioning, Operations Planning (S17b)

DM-7632 Develop planning packages from DM replanning deliverables

Service Management and Monitoring

Development of the DPP services model within the LSST Enterprise Architect framework will continue, with refinement of the mapping of high-level services to the internal service architecture and development of incident and request response processes. Work on service monitoring will continue with configuring and testing performance thresholds, and investigating/integrating prototype real-time metrics and key performance indicators.

DM-8515 Service Management Definition and Refinement

DM-8516 Enterprise Monitoring Integration – Phase 1

DM-9658 Enterprise Monitoring Design – Phase 2

DM-9657 Service Management Definition & Refinement (S17b)

Hardware acquisition and provisioning planning

We will continue planning for hardware deployment in Chile that will occur later this year, verifying needs against planned procurements, coordinating with Chilean staff and beginning the procurement process. This will lead to the deployment of project-wide network-based security services, bringing together elements of computing hardware, networking, storage, security, and system administration.

DM-8512 Planning Deployment in Chile

02C.07.01 Processing Control

Batch Production Services

In April, we will continue to enhance this service façade with the ability to orchestrate and run more complex pipelines in the planned workload management framework using Pegasus and HTCondor, adding orchestration elements as they become available. This includes participating in ongoing work to detail the interface between the production system and SuperTask. Work also will continue in preparing a production-scale batch processing of HSC data in the as-is software

framework. We will also continue our involvement in the SuperTask/Butler Working Group, including contributions to requirements and design documentation.

DM-8333 Run selected existing science codes with initial Batch Production Service façade

DM-9662 Workflow creation improvements & expansion

DM-9664 SuperTask Working Group Activities

Level 1 Services

We will continue integrating the existing prompt processing system orchestration code with the newly developed front end elements for Level 1 services, and continue work on the OCS system, prototyping logging/monitoring as a first step in performance evaluation. We will also continue working with Telescope & Site developers to identify sources and contents of the EFD and OCS metadata available to build headers.

We will participate in Pathfinder Integration Activity #3, which will be hosted for two days at NCSA.

We have proposed work on a design for acquiring header information from OCS in “catch-up” mode, in which archiving of camera data occurs asynchronously from data acquisition by the camera; if accepted, work will begin on this task.

DM-8317 Enhancements for OCS Bridge components

DM-8575 Deploy Mock for OCS with monitoring info

DM-9659 Header construction and distribution architecture

DM-9661 Reintegration of archiving and prompt processing

Common Middleware and Other Tasks

Work for April is anticipated to include changes to common library functions necessary to support Batch Production and Level 1 Services, and addressing further user- and developer-driven issues as they arise.

DM-9663 Emergent Middleware Work (S17b)

02C.07.02 Infrastructure Services

Data Backbone

We will finish work with SLAC and the DM architect to gather requirements about the operational model of hosting the 60+ identified databases, including estimated query loads, capacity, and other administrative concerns. We will also begin prototyping a data backbone

façade to determine the feasibility of an interface sufficient to meet the needs of near-term DM activities and milestones. We will investigate methods for data movement over a WAN supporting the spectrograph test stand in January 2018, as well as long-term solutions for data distribution in the Data Backbone.

DM-7645 Prototype bookkeeping and interfaces for the data backbone

DM-9665 Prototype Data Backbone Façade

02C.07.03 Environment and Tools

QA Environment

Work to provide a Jupyter hub instance and access to datasets to support upcoming QA activities will continue with working toward persistent access using shared file systems, and containerization using kubernetes.

DM-8330 Instantiate VM with Jupyter Hub and GPFS access.

02C.07.04 Site Infrastructure

Hardware Acquisition, Deployment and Provisioning

Planned work includes finalizing the migration from NFS to GPFS file systems, continuing migration of services, and beginning the decommissioning of the NFS file system. Work on implementing disaster recovery procedures for datasets will be completed. We will continue planning and procurement activities for the systems to be used in raft-scale testing of the L1 Services and finalizing plans for deployments in Chile.

DM-8503 Migrate Users from NFS

DM-8504 Migrate Services from NFS

DM-8505 Decommissioning NFS

DM-8502 Implementation of disaster recovery for /datasets (Phase 2)

DM-9685 Procurement Activities for L1 Test Systems

Service Management for LSST development resources, Emergent and Miscellaneous Work

Work for April is anticipated to include addressing further user- and developer-driven requests and issues with migration from and retirement of the NFS file system, and related system needs. We will continue to investigate and prototype additional means to monitor the services provided on the development cluster. We will also address technical debt related to system administration.

DM-9666 Service Management & Emergent Work (S17b)

DM-9660 System Administration Technical Debt (S17)
DM-9667 Cluster Service Monitoring Development

Recruitment update:

Scott Schreiner was hired as an IT Systems Engineer, and will help construct and operate compute systems that support development and production services at NCSA.

AURA

This section covers accomplishments and planned activities in WBS 02C.08 International Communications and Base Site.

Current accomplishments:

02C.08.00 International Communications and Base Site Management Engineering and Integration

02C.08.01 Base Center

- Summit Base ITC Tiger Team: Continued work on the Summit Base ITC Design Document (SBIDD) that will be proposed to the LSST CCB to become a baselined LSE document.
<https://confluence.lsstcorp.org/display/SYSENG/LSST+Summit+-+Base+ITC+Design>. Created draft of the document for discussion at the Joint Technical Meeting in Glendale, March 5- 9, 2017. Only missing sections now are on Visualization and Common Hardware standards. Created a draft Summit Computer Room Requirements Document to parallel LSE-239 Base Data Center Requirements. While construction is already in progress on the Summit, and basic construction parameters are fixed, there are still specifications needed for the build-out of the Computer Room.
It is understood that the Tiger Team will end with the baselined SBIDD NLT the August LSST 2017 meeting. In reality we need to be done before then, by end of May latest to ensure adequate implementation time to support the Auxilliary Telescope and Early Operations.. At that time, the SBIDD will be baselined and complete for Summit, Summit - Base, and Base. Summit sections will be more mature than Base sections when submitted to CCB for the first time. This is unavoidable due to more advanced schedule for Summit Facility and Network. For example, we will not have complete Control Room designs by August 2017, given an 8-month INRIA design contract, but we will have enough to define the commissioning cluster, network, servers, and DM software

contribution. When the Tiger Team disbands SBIDD responsibility will migrate to "regular" System Engineering team.

- Summit and Base Networking and Computing: Planning of racks and network equipment continued this month. The date of reception for the Summit LAN RFP vendor proposals was postponed to 30 March at their request. Work continued on refining the costing of the Summit LAN. Provided support for planning and overseeing for the arrival of EIE dome team and installation of phones and internet in their containers on Pachon.

02C.08.02 Chilean Data Access Center

- Provided input for analysis underway at NCSA for a “thin DAC” approach.

02C.08.03 Long-Haul Networks

02C.08.03.01 Chile National WAN

- Summits - AURA Gatehouse Network: No change from last month. We are still waiting for the final measurements for the fibers from Gatehouse to Summits. We will perform these measurements at the same time as measuring from La Serena to the AURA Gatehouse, to avoid having to split the already fused fibers.
- DWDM Equipment: The AURA requisition went out to Coriant for the DWDM equipment. Expect arrival of the equipment at the end of May, when we will have 30 days to inspect the equipment before payment. REUNA sent the formal announcement to Coriant, stating they are the awarded solution. REUNA sent Annex IV Contract 1, which defines the activation of the 1st optical channel from AURA/LSST to Santiago, to AURA. AURA revised and agreed on the terms, only the final signature remains. With the amendment signed REUNA will sign the contract with Coriant/Raylex. REUNA also worked with Telefonica and Raylex to start planning the equipment deployment phase.
- Santiago-La Serena: REUNA delivered the report that documents the completion of the long haul fiber with the measurements of length and attenuation of the fiber in each segment. These measures were sent to Coriant, for them to verify the equipment design, and the result was positive. It is relevant to note that as the fiber is brand new the attenuation parameters are of a very good quality, less than the theoretical parameters.
- La Serena – AURA Gatehouse: The new ducts into the AURA compound are now ready and waiting for Telefonica to reinstall the existing fiber cable and the new fibers coming from the LSST/AURA contracts. We are still waiting for the final measurements for the fibers from La Serena to the Gatehouse. We will perform these measurements at the same time as measuring from the Summits to the AURA Gatehouse, to avoid having to split the already fused fibers.

02C.08.03.02 International Chile - US WAN

- 100 Gbps Managed Ring: No planned activity in this period. The ring is performing as expected.
- Management and Coordination Contract:
Florida International University (FIU) hosted the Global Environment for Network Innovations (GENI) Engineering and GENI GEC25 Conferences. GENI provides a virtual laboratory for networking and distributed systems research and education relevant to the LSST long-haul networks.
FIU CIARA hosted the first monthly meeting of the LSST Network Engineering Team (NET), which is the cross-institution team of network management and technical staff responsible for designing, engineering, deployment, and supporting operations of the LSST long-haul networks. The NET members are:
 - Technical:
 - Ron Lambert (LSST)
 - Jeronimo Bezerra (CIARA FIU)
 - Matt Kollross (NSCA)
 - Albert Astudillo (REUNA)
 - Humberto Galiza (AmLight RNP)
 - Shahram Belldeh
 - Management:
 - Jeff Kantor (LSST)
 - Julio Ibarra (CIARA FIU)
 - Sandra Jaque (REUNA)

The early focus of the NET is on the Network Verification Plan and Matrix. This plan covers the final verification tests to be performed on the network segments to confirm that they are meeting the LSST requirements.

- Spectrum Contract: The subaward from FIU to Angola Cables (AC) is drafted and pending the commitment form from AC to FIU to fully execute the agreement. AmLight Co-PI Chip Cox met with AC on March 10th to review the subaward.
- US National WAN: AmLight Co-PI Ibarra and Jeff Kantor planned a follow up meeting with ESNNet to review and modify the initial proposal for the US National WAN. That meeting has not been scheduled yet. A DOE “advocate” for LSST networking needs to be supported by ESNNet was identified at FNAL. A strategy meeting between AURA, FIU, and NCSA is planned prior to engaging the FNAL advocate.

Planned activities:

02C.08.01 Base Center

- Summit Base ITC Tiger Team: Continue work to complete draft of Summit - Base ITC Design Document for submission to CCB. Complete Summit Computer Room Requirements Document and submit to CCB.
- Summit and Base Networking and Computing: Planning of racks and network equipment will continue to be completed by May. Receive the Summit LAN RFP proposals and score with the selection committee.

02C.08.03 Long-Haul Networks

02C.08.03.01 Chile National WAN

- Summits - AURA Gatehouse Network: Complete measurements of the fibers between the Summits and Gatehouse. Expect to have the Telefonica fiber cable installed in the existing NOAO computer room.
- DWDM Equipment: The Purchase Order will be issued to Coriant for AURA/LSST DWDM equipment. REUNA will acquire from CORIANT/RAYLEX the long haul DWDM equipment plus the installation services of the whole. To do this, they will execute the following steps:
 - A letter from AURA (draft was prepared and received) to formally inform the evaluation committee of the tender process that it was accepted to award to CORIANT.
 - REUNA will sign a Contract Amendment with AURA to purchase the equipment. A draft was developed and has been agreed.
 - With the amendment REUNA will sign the contract with CORIANT/Raylex.
 - REUNA will prepare for the deployment phase with a site survey to confirm that all the physical conditions including power, are in place and will be ready when the equipment arrives. Deployment will start in the Santiago sites, with five sites involved. After that the plan is to deploy in the other sites from Santiago to La Serena.
- Santiago-La Serena: REUNA to prepare the report related to the reception of the fiber, Delivery #4 of Contract 1 Annex II and send it to AURA.
- La Serena - AURA Gatehouse: Expect to have the Telefonica fiber cable installed in the existing NOAO computer room.

02C.08.03.02 International Chile - US WAN

- 100 Gbps Managed Ring: Next planned activities are scheduled for May/June 2017, when the Pacific 100G circuit will be opened in Chile.
- Management and Coordination Contract: Conduct meeting of the LSST Network Engineering Team (NET) to resolve issues/comments on Network Verification Plan and Matrix.
- Spectrum Contract: FIU and AC complete the operations and maintenance agreement and payment timing.
- US National WAN: Negotiations to refine and improve the ESNet service and costs will continue during April.

Recruitment update:

- The hiring committee scored and selected six candidates for the Network Engineer. Telecons were carried out with those six and two picked for F2F interviews in La Serena sometime in April.
- IT Technicians were selected and four who live in the La Serena area will be advised to visit AURA during a week in April for F2F interviews. If no success with any of these candidates then a search further afield in Chile will take place.

