

Large Synoptic Survey Telescope Data Management Monthly Report

October 2017

High-level Summary

During this month the named DM “points of contact” were rolled-out to the various science collaborations. We also unveiled a “Data QA” section on the community.lsst.org site. This provides a forum in which members of the DM Subsystem Science Team will respond to questions from scientists about the details of LSST data products. Within the Subsystem Science Team, work started on a proposal to simplify LSST’s solar system data products (while delivering equivalent capabilities to the end user).

A priority for DM as we move forward will be to undertake a series of internal reviews of important components of the system. This effort got underway with a successful design review of the Header Service.

The Data Butler Working Group continued this month, with substantial design and prototyping work being undertaken on the new, “third generation” I/O abstraction layer. This work will be presented to the DM Leadership Team during the early November face-to-face meeting.

Version 14.0 of the LSST Science Pipelines was publicly released and is now available from the pipelines.lsst.io site. This new version provides a range of new and upgraded capabilities, which are described in the extensive notes that accompany the release.

Moving beyond the 14.0 release, the Pipelines group delivered a detailed description of the algorithm which is being proposed to deliver template images which are corrected for differential chromatic refraction and an early implementation of the source association system within the end-to-end alert production test system. The new deblending algorithm was presented to various science collaborations and external workshops, collecting valuable feedback.

HyperSQL table data support has been added to the Firefly server making it possible to use standard SQL for table data manipulation.

We supported the Pathfinder Early integration Activity “Next Visit” milestone and workshop (hosted by NCSA).

Work on Summit Network progressed. Power cable trays and cables were installed in the Computer Room, and we updated the design with seismic bracing of supports for the

raised floor.

Delays in the provisioning of the upgraded PDAC Qserv head node, and additional evolving hardware and GPFS/NFS instabilities within the PDAC are creating schedule pressure that could put some upcoming end-of-cycle deliverables at risk.

Several team members attended the ADASS meeting in Santiago, Chile and some stayed for the subsequent IVOA Interoperability Meeting. This included a very productive “birds of a feather” session on science platforms, which was an extension of the earlier discussions involving NOAO, STScI and LSST. We conducted discussions on VOTable data representation, hierarchical image representation, authenticated service listings in VO registries, and the applicability of other software and VO standards to LSST at both ADASS the IVOA meeting.

In staffing news, Mario Jurić has formally stepped down as DM Subsystem Scientist and transitioned to his new role as DM System Science Team Coordinator. We have carried out telephone interviews with several candidates to replace Mario as the Subsystem Scientist, and will move to face to face interviews next month for three of them. We have also interviewed candidates for the DM Release Manager/Engineer and expect to make an offer next month.

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.

Milestone Summary

Delayed Milestones

ID	Description	Due Date
DM System Architecture		
DLP-538, DLP-539, DLP-541, DLP-544	Assorted ICDs updated to Phase 3 Work is in progress to remove all TBDs and TBRs from these ICDs and achieve “substantial completion” by the end of the calendar year.	2017-05-31
Science Data Archive & Application Services		
DLP-802	AP/L1 Database Design Functional prototype has been delivered to AP team. Further scale/performance experiments pending deployment of Oracle consolidated db systems at NCSA (currently included in the S18 NC	2016-11-30

	<i>procurement.)</i>	
DLP-472	Qserv Data Distribution <i>Delayed due to personnel reassignments. Gaponenko assigned. Scheduled for completion in F17; however, PDAC head node installation delays and evolving GPFS/NFS instability at NCSA are prolonging Gaponenko's scheduled database load activities which could ripple to further delays in this milestone.</i>	2017-05-31
LSST Data Facility		
DLP-366	Mid-scale processing of eligible camera data <i>This was postponed due to SuperTask framework status. Upcoming milestone LDM-503-2 will demonstrate mid-scale processing of eligible camera data using an alternate production-capable framework. Expected completion date is November 30, 2017.</i>	2016-11-30
International Communications & Base Site		
DMTC-7400-242 0	REUNA La Serena - Santiago DWDM equipment <i>Delayed due to vendor delay. Expected completion November 10, 2017</i>	2017-07-31
DMTC-7400-209 0	Report on functional fiber connections, including AURA equipment <i>Delayed due to REUNA DWDM equipment delay. Expected completion November 15, 2017</i>	2017-08-18
DMTC-7400-248 0	AURA DWDM equipment in operation <i>Delayed due to delay in getting access to Cerro Pachon Caseta. Expected completion November 15, 2017</i>	2017-10-30
System Integration & Test		
DLP-579	Usability and developer efficiency <i>This is currently being held up by Orchestration and Supertask framework.</i>	2016-08-31

Detailed Project Progress

1.02C.01: System Management

Current accomplishments

We held a cross project meeting with NOAO and STScI Oct 9 and 10 in Tucson. One major topic was the science platform and how we might best collaborate to make that a success across all institutes. This was also a precursor to the Birds of a Feather (BoF) session in

ADASS. The BoF was highly attended and a write up will be in the ADASS proceedings (draft <https://github.com/lstt-dm/adass27-womullan/blob/master/B5/B5.pdf>) There are some areas which we will investigate further. LSST probably has the most advanced thinking in this area.

We also prepared for and attended the Directors' Operations Proposal review and interviewed two candidates for DM Release Manager/Engineer position; we expect to make an offer imminently.

Planned activities

We will interview candidates for DM Subsystem Scientist and make an offer for the Release Manager/Engineer. Preparations need to be made for the Chilean community workshop.

1.02C.02.01: Data Management Science

Current accomplishments

We have rolled out the DM points of contact for the LSST Science Collaborations, and opened the "Science > Data Q&A" section on community.lsst.org. These set up a well-defined mechanism for DM to communicate with the LSST DM stakeholders. Melissa Graham will serve as the editor of the "Science > Data Q&A" section. We have also delivered a proposal for the change of data product terminology from "Level 1/2/3" to "Prompt", "Data Release", and "Added-Value" products, delivering it to the Directorate.

We have continued work on assessment of DM preparedness for execution of special survey programs (e.g., deep drilling fields and others), led by Melissa Graham. The final report and actionable issues will be issued in November. Colin Slater and Chris Suberlak continue to work on the development of the test metrics and data sets for crowded fields.

Mario Juric met with the Director of the Minor Planet Center to discuss LSST-MPC interfaces. We have drafted a proposal on changes to LSST's Solar System data products, with the goal to speed up submission of LSST data to the MPC, and use the MPC catalog as LSST's Level 1 orbit catalog. This proposal, as well as the status and plans for Solar System data products, were presented to the Solar System Science Collaboration at the annual AAS Division of Planetary Sciences Meeting in mid-October. The feedback was overwhelmingly positive.

Mario Juric has stepped down as the DM Subsystem Scientist and assumed the role of the DM System Science Team Coordinator. His primary responsibilities through the end of FY'18 will be the completion of DM SST studies currently in progress, coordination of SST members' work, and help in on-boarding of the new DM Subsystem Scientist, once hired.

Planned activities

- Finalizing the Special Programs DM preparedness report.
- Initial report and test dataset for LSST crowded fields.

1.02C.02.02: DM System Architecture

Current accomplishments

Architecture Team design activities included conducting a design review of the Header Service (and providing a list of telemetry items it should capture in initial tests); completing a first draft of the Butler Use Cases document (LDM-592); clarifying current designs and capabilities with scientists as input to a technical note describing Special Programs processing (DMTN-065); and conducting discussions on VOTable data representation, hierarchical image representation, authenticated service listings in VO registries, and the applicability of other software and VO standards to LSST at ADASS XXVII, the subsequent IVOA meeting, and an NOAO/STScI/LSST synergy meeting.

The team also drove an update of the early integration activity plan and interviewed candidates for the DM Release Manager position, as well as reviewing the DM risk register.

Planned activities

The team will file a change request to incorporate the updated early integration schedule into the project baseline; attend the DMLT face to face meeting and the Mauna Kea Python Workshop; submit the products of the Butler Working Group for review; and complete the administrative tasks needed to bring a new software systems engineer into the team.

1.02C.03: Alert Production

Current accomplishments

02C.03.00 – Alert Production Management Engineering and Integration

- Travel & meetings:
 - Eric Bellm spent the first three weeks of this month in Pasadena working on commissioning the [ZTF](#) system. This work is not funded by LSST, but the experience gained with ZTF will directly feed into the LSST construction effort.
- Construction of an “end-to-end” Alert Production test system [[DM-9676](#), [DM-10770](#), [DM-10771](#), [DM-10773](#), [DM-10775](#)]:
 - A new task which can be used to select only those images which reach some required level of seeing has been added to the stack. This is used to select

the images which will be used to build the coadds which serve as templates for image differencing in the end-to-end system. [[DM-10977](#)]

- The simple source association system provided by AssociationTask has now been incorporated into the prototype end-to-end pipeline. [[DM-12256](#), [DM-12257](#)]
- A detailed consideration of options for collecting verification metrics from running pipeline tasks was completed by Krzysztof Findeisen. This substantial piece of work has been published as [DMTN-057](#) and serves as important input to the design not only of the prototype AP system, but to the wider metric collection and analysis systems being developed across DM. [[DM-11772](#)]

02C.03.01 – Single Frame Processing Pipelines

- Replacement of old WCS classes [[DM-9679](#)]:
 - Work continues to replace the old Coord class hierarchy and [WCSSLIB](#)-backed WCS system with versions based on SpherePoint and [AST](#). [[DM-10765](#), [DM-11162](#)]

02C.03.02 – Association Pipelines

- Prototype source association system [[DM-10768](#)]:
 - The AssociationTask, encapsulating the simple source matching algorithm used to assemble DIAObjects, has now been made available in the DM Stack. [[DM-11747](#), [DM-11921](#)]

02C.03.03 – Alert Generation Pipelines

- Testing prototype alert distribution system [[DM-7355](#)]:
 - Our existing tests on the prototype alert distribution system have been scaled up to include a larger total number of reports and data volume. The results are reported in [DMTN-027](#). [[DM-11804](#)]

02C.03.04 – Image Differencing Pipeline

- Research and implementation of techniques for mitigating Differential Chromatic Refraction (DCR) [[DM-6245](#), [DM-9613](#)]:
 - [DMTN-037](#), an extensive description of the approach being proposed to mitigate DCR in the DM Stack, was published this month. This is the culmination of several months of effort by Ian Sullivan (UW), and marks an important step forward in the design of the DM system. The algorithm being proposed will undergo extensive testing and trial implementation over the next several months. [[DM-7446](#)]
 - The basic algorithm being proposed for generating DCR-corrected templates makes the assumption that there is negligible DCR across each DCR model

plane's bandwidth. This assumption is questionable when observing with high airmass. An extension to the algorithm to correct for this effect has been developed, and is described in [DMTN-037](#). [[DM-12380](#)]

02C.03.05 – Application Framework for Exposures

- Emergent work [[DM-10068](#), [DM-11798](#)]:
 - A number of minor bugs and issues were addressed across the codebase. Highlights include:
 - Resolved an issue in which AST-backed WCS objects could not be properly round-tripped to persistence storage. [[DM-12270](#)]
 - Resolved an issue in which photometric models dependent on Chebyshev polynomials of greater than the seventh degree could not be saved to persistent storage. [[DM-11957](#)]
 - Documentation for using the popular Sublime Text editor in compliance with the DM coding standards was made available to the project. [[DM-9464](#)]
 - The definition of the PhotoCalib object, used to represent photometric calibration within the DM Stack, has been revised so that it is multiplicative rather than based on division. [[DM-11905](#)]

02C.03.06 – Moving Objects Pipeline

- No work was undertaken in this WBS during this month.

02C.03.07 – Photometric Calibration Pipeline

- The photometric model used by the "Jointcal" system for simultaneously fitting for photometry and astrometry across multiple visits has been upgraded. The new model is described in [DMTN-036](#). [[DM-9195](#), [DM-11456](#), [DM-11905](#)]
- Comparison of Jointcal with meas_mosaic [[DM-11783](#)]:
 - Work is ongoing to compare Jointcal to meas_mosaic, the older, HSC-specific system which it is destined to replace. Our immediate goal is to demonstrate that Jointcal performs at least as well as meas_mosaic under all (plausible) circumstances, so that the latter can be dropped from the DM Stack.
 - A dataset has been identified which will be used to perform this comparison. [[DM-11784](#)]
 - The Jointcal reference object loader has been updated to support the meas_mosaic comparison project. [[DM-12254](#)]

02C.03.08 – Astrometric Calibration Pipeline

- No work was undertaken in this WBS during this month.

Planned activities

02C.03.00 – Alert Production Management Engineering and Integration

- Eric Bellm and John Swinbank will attend the DMLT face-to-face meeting in Tucson.
- The team will focus on planning for the Spring 2018 development cycle, which starts in December.
- Work will continue on the end-to-end system, with the immediate goal of meeting the LDM-503-3 (“Alert Generation Distribution”) milestone by the end of November.

02C.03.01 – Single Frame Processing Pipelines

- The replacement of the WCS classes with AST-backed alternatives will reach its conclusion.

02C.03.02 – Association Pipelines

- Work will continue on the association system in support of the development of the end-to-end system.

02C.03.03 – Alert Generation Pipelines

- Benchmarking and monitoring of the alert generation system will continue.

02C.03.04 – Image Differencing Pipeline

- Work will focus on implementing the DCR correction algorithm described in [DMTN-037](#) as a Task within the DM Stack framework.

02C.03.05 – Application Framework for Exposures

- We will continue to service technical dept and emergent requirements.

02C.03.06 – Moving Objects Pipeline

- No work is planned in this WBS during this month.

02C.03.07 – Photometric Calibration Pipeline

- Ongoing comparison of Jointcal and meas_mosaic with the ultimate aim of retiring the latter.

02C.03.08 – Astrometric Calibration Pipeline

- No work is planned in this WBS during this month.

Recruitment update

- John Swinbank joined the team on 16 October. He will serve as T/CAM for both the Alert Production and Data Release Production groups and as Deputy Data Management Project Manager.
- [Currently advertising](#) for a Research Scientist with a specialism in image differencing techniques to join the team at the University of Washington.

1.02C.04: Data Release Production

Current accomplishments

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

- Travel & meetings:
 - Moolekamp visited the [Space Telescope Science Institute](#) during the week of 2 October to participate in a workshop on PSF photometry.
- Management & planning [[DM-10816](#)]:
 - Documentation and release notes were developed in support of the release of version 14.0 of the DM Stack. [[DM-12044](#)]
 - The way that teams and labels are used on the DM JIRA project were revised and the new system properly documented. This makes it easier to identify the responsible manager and any interested groups for each ticket in JIRA. [[DM-11943](#), [DM-11945](#)]
- Forward Global Calibration Method (FGCM; [Burke et al., 2018](#)) [[DM-10584](#)]:
 - Work through this month continued to focus on integrating the FGCM codebase with DM by exposing it using the standard “task” functionality. [[DM-11313](#)]
 - FGCM was updated to support multiple filters which can be cross-calibrated into a single “band”. [[DM-11960](#)]
- Emergent work and pipeline support [[DM-10382](#)]:
 - A number of minor issues and bugs have been addressed across the codebase. [[DM-11902](#), [DM-11948](#)]
- Pipeline QA [[DM-10571](#)]:
 - The pipe_analysis package, which provides plotting and other tools to assist in pipeline QA work, has been updated to work with Python 3 and modernized in terms of the DM Stack APIs being used. [[DM-11976](#), [DM-12173](#)]
 - The pipe_analysis scripts now dump data to [Parquet](#) format files which are appropriate for loading into our [Bokeh](#)-based interactive QA environment.
 - A series of updates have been made to our Jupyter-notebook based QA system. These include generating “linked” plots, where the user can plot

multiple views of the same dataset and have selections or updates made in one plot automatically reflected in the others, and retrieving data IDs and other statistics based on data interactively selected by the user from a figure. [[DM-11455](#), [DM-11682](#)]

02C.04.01 – Application Framework for Catalogs

- Middleware and framework development [[DM-10586](#)]:
 - The DM team continued to play an active role in the [Data Butler Working Group](#) throughout this month. In addition to contributing to use case collection and requirements development, the DM team has been leading design and development effort on a proof-of-concept of the vision for the Butler that the working group is proposing. This will be presented to the DM Leadership Team in early November. [[DM-11749](#), [DM-11750](#), [DM-11752](#), [DM-11892](#), [DM-12139](#), [DM-12141](#)]
- Emergent work and reduction of technical debt [[DM-10383](#)]:
 - A number of improvements have been implemented across the codebase. [[DM-12417](#), [DM-12264](#)]

02C.04.02 – Calibration Products Pipeline

- Auxiliary telescope development [[DM-10581](#)]:
 - Work has continued to test and integrated the obs_ctio0m9 camera support package. [[DM-10964](#)]
- Processing camera test stand data [[DM-10897](#)]:
 - The obs_comcam and obs_ctio0m9 camera support packages were updated to work with Python 3 and have been added to the standard DM software distribution metapackage (lsst_distrib). [[DM-11653](#), [DM-12070](#), [DM-12080](#)]
 - Work has continued throughout this month on implementing some functionality from [eotest](#), the Camera Team's electro-optical acceptance testing tool, within the DM stack framework. [[DN-11479](#)]

02C.04.03 – PSF Estimation

- Wavefront measurement and PSF reconstruction [[DM-10355](#)]:
 - The capability has been developed to use intra- and extra-focal images to predict the optical PSF of in-focus images. [[DM-11315](#)]
 - Effort is now focused on developing a summary document describing the work carried out during this cycle and suggesting future directions for development. This document was not complete as of the end of the month being reported upon. [[DM-9988](#), [DM-12472](#)]

02C.04.04 – Image Coaddition Pipeline

- Warped image comparison [[DM-8290](#)]:

- Throughout this month work has continued to evaluate the scientific impacts of the new “robust” coaddition algorithm and to determine the most reliable default configuration for deploying it in production. This work will continue into next month. [[DM-11445](#), [DM-12445](#)]

02C.04.05 – Object Detection and Deblending

- Deblender development [[DM-10353](#)]:
 - The deblending algorithm currently under development was presented to both the relevant [Dark Energy Science Collaboration](#) working group and the STScI PSF Photometry Workshop. These meetings provided valuable input to guide further development of the algorithm. [[DM-12068](#), [DM-12232](#)]
 - A new set of simulated data for deblender testing was created. These images have a better understood distribution of objects to allow more reliable and easier-to-interpret tests to be undertaken. [[DM-11642](#)]
 - Most work this month was invested in the encapsulation of the new deblender logic in the Task framework used by the DM Stack. As of the end of the month, this work was nearing completion, but had not yet been reviewed and merged. [[DM-11329](#)]

02C.04.06 – Object Characterization Pipeline

- Experiments in shear measurement on coadds [[DM-10579](#)]:
 - Work has continued through this month to investigate the methodologies and trade-offs which can be employed to measure high-quality galaxy shears directly from coadded images. We expect this work to culminate next month in an internal report which will then be developed into a publication.
 - Through October, efforts have focused on increasing the complexity of the coadds being used to test the shear measurement techniques. [[DM-11311](#)]
- Improved galaxy flux measurement algorithms [[DM-10580](#)]:
 - The bulk of the work in this area has focused on completing the integration of [Synpipe](#) with the DM Stack. This has involved a number of minor bug fixes and some significant book-keeping effort, but was almost complete by the end of the month. [[DM-11318](#), [DM-11319](#), [DM-11897](#), [DM-11898](#), [DM-12108](#)]
 - Preparatory work has now begun to develop a galaxy shape fitting code. This will be the focus of work throughout November and early in the Spring 2018 development cycle. [[DM-9999](#), [DM-10847](#)]

Planned activities

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

- Jim Bosch and John Swinbank will attend the DM Leadership Team Face-to-Face meeting in Tucson during the week of 30 October.
- Jim Bosch will attend the [Mauna Kea Python Workshop](#) during the week of 6

November.

- The plan for the Spring 2018 development cycle will be developed and baselined.
- Focus on delivering the LDM-503-2 (“HSC Reprocessing”) milestone.

02C.04.01 – Application Framework for Catalogs

- Continue working with the Data Butler Working Group to refine Butler requirements and develop the Butler design sketch. This will culminate in the acceptance of the new Butler design by the DM Leadership Team.

02C.04.02 – Calibration Products Pipeline

- Demonstrate successful reduction of CTIO 0.9m data based on the DM Stack.
- Demonstrate successful derivation of detector gains based on test stand measurements processed with the DM Stack.

02C.04.03 – PSF Estimation

- Complete a technical note (“DMTN”) describing the work which has been performed to date on PSF estimation.

02C.04.04 – Image Coaddition Pipeline

- Complete validation, test and characterization of the new robust coaddition system.

02C.04.05 – Object Detection and Deblending

- First at-scale testing of the NMF deblender.

02C.04.06 – Object Characterization Pipeline

- Produce basic prototype of new shape measurement code.
- Complete an initial report describing the results of experiments to measure galaxy shears from coadded images.

Recruitment update

- John Swinbank left Princeton to move to the University of Washington this month. He will continue as T/CAM for the DRP team remotely, while also serving as T/CAM for Alert Production and as Deputy Data Management Project Manager.
- Recruitment is underway for an 0.5 FTE project management position located in Princeton to act as a deputy to Swinbank.

1.02C.05: Science User Interface & Tools

Current accomplishments

02C.05.00

- Prep work for ISST Science Platform workshop to be held in December 4-8, 2017.
- Prep work for DM and SE allhands meeting in March 2018, hosted by IPAC.
- Gregory and Xiuqin attended ADASS in Chile. Gregory gave a presentation at the "Science Platforms" BoF session.
- Gregory attended IVOA interoperability meeting in Chile.

02C.05.01 Basic Archive Access Tools

- Bug fixes. (DM-12112, DM-12208)

02C.05.02 Data Analysis and Visualization Tools

- New functions
 - UI support to show the active trace in the multi-trace chart (DM-11604).
 - Create a framework to incorporate the multi-trace charts in application, specifically tr-view (DM-10833).
- Code refactoring and improvement:
 - using HyperSQL as table data support, making it possible to use standard SQL for table data manipulation (DM-11814).
 - Upgrade the nom.tam.fits Java package used by Firefly (DM-11883).
 - Performance improvement for Plotly chart over 1 million points (DM-12089).
 - Other miscellaneous ones (DM-12112, DM-12189).

Planned activities

02C.05.00

- Prep work for ISST Science Platform workshop to be held in IPAC.
- Prep work for DM allhands meeting in March 2018, possibly hosted by IPAC.
- S18 planning.

02C.05.01 Basic Archive Access Tools

- Bug fixes and improvements.
- UI elements for access user workspace (DM-10855).
- Work with DAX on ImageServ and MetaServ API v1.

02C.05.02 Data Analysis and Visualization Tools

- New functions:
 - Upgrade the nom.tam.fits java package in Firefly (DM-11883).
 - HiPS map visualization in Firefly (DM_12305, DM-12551).
- Firefly_client update to be in sync with all the new JS features (DM-11817).
- Finish Firefly_client Python documents and examples, publish the doc.

Recruitment update

- None.

1.02C.06: Science Data Archive & Application Services

Current accomplishments

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

- Team continued improving robustness of DAX services and otherwise supported SUIT integration efforts in the PDAC as necessary.
- Mueller, Hanushevsky, and Pease attended XLDB 2017 in Clermont-Ferrand, France.
- Van Klaveren and Lo attended ADASS XXVII and IVOA Interop meetings in Santiago, Chile.
- Salnikov took vacation during this month.

02C.06.01.01 Catalogs, Alerts and Metadata

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

02C.06.01.02 Image and File Archive

- Lo completed initial implementation of image service supporting image cutouts [DM-9929].
- Lo implemented code to keep intermediate files in memory rather than persisting to disk in imgserv [DM-12069].

02C.06.02.01 Data Access Client Framework

- Van Klaveren continued participation in the Butler 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 worked on synthesis and load a 30% DR1 test dataset on one of the Qserv

development clusters at CC-IN2P3.

- Hanushevsky completed converting czar code to use the new XRootD SSI v2 API for content-addressed messaging between czar and workers [DM-9737].
- Gaponenko continued work on Qserv data distribution/replication framework.
- Salnikov extended unit tests to cover the newly introduced asynchronous/disconnected query capabilities [DM-12360].
- Salnikov implemented cleanup of stale entries in Qmeta database following czar crash/restart [DM-12204].
- Salnikov implemented minor fixes to build scripts [DM-12361, DM-12362].
- Mueller, Pease, Hanushevsky, and Jammes met for two days following XLDB 2017 and planned future development path for Qserv Kubernetes deployment tooling.
- Jammes added debug tools to containers generated by CI [DM-11795].
- Mueller fixed Qserv dev container and Travis CI builds [DM-12364].
- Mueller updated XRootD from upstream and addressed several small gcc 7.2 incompatibilities and build issues [DM-11677, DM-12308, DM-12363, DM-12382].

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.
- Mueller to attend DMLT meeting in Tucson.

02C.06.01.01 Catalogs, Alerts and Metadata

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

02C.06.01.02 Image and File Archive

- Lo to add mosaic and image stitching features to imgserv.

02C.06.02.01 Data Access Client Framework

- Further development on the current Data Butler implementation is now frozen, pending the outcome of the Butler Working Group. Pease to work as needed, however, to address any issues in the current Butler implementation that may be blocking other teams' work.

02C.06.02.02 Web Services

- Van Klaveren to add support for asynchronous queries at the web service layer.

02C.06.02.03 Query Services

- Gates to complete synthesis and load of 30% DR1 test dataset on the IN2P3 cluster, and Thukral to take KPMs using that dataset.
- Gaponenko to continue work on Qserv data distribution/replication framework.
- Jammes to continue adapting Qserv for deployment within Kubernetes.
- Pease to relocate subchunk query expansion from czar to workers for greater parallelization and to remove a related bottleneck in the current czar implementation.

02C.06.02.04 Image Services

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

02C.06.02.05 Catalog Services

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

Recruitment update

- No recruitment activity this month.

1.02C.07: LSST Data Facility

Current accomplishments

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

We continued service management work on refining service change management processes.

02C.07.02 Infrastructure Services

We completed our commitment to attend Butler Working Group meetings and provide input and feedback, pending feedback from DM leads and overall progress on deliverables.

We completed current activities toward researching containerization options, publishing a tech note on the use of OpenShift for container management and security controls.

We completed planned work toward file transfer capabilities, testing various file transfer clients against a WebDAV service. This work will support archiving of data from the Spectrograph test stand in early 2018, as well as long-term solutions for data distribution in the Data Backbone.

02C.07.04 Site Infrastructure

We continued work toward creating policy for file system reorganization for data sets, including guidelines for use, retention policies, and best practices.

We continued acquisition and provisioning of final FY17 hardware procurements, including a capacity increase for the lsst-db system, expanded L1 complete test stand integration platform, configuration of the prototype Kubernetes cluster, completion of the NCSA 3003 lab upgrade, and initial setup for the Base Authentication and Authorization hardware.

We continued work in consolidation of ITC management systems, investigating the use of packages such as Foreman, Pakrat, and Katello to provide maintainable processes and procedures for managing multiple enclaves at NCSA, at the Base, and the network-based security systems at the Summit.

02C.07.05 LSST Data Facility Management, Service Architecture, and Project Controls

We continued working on articulating requirements documents and defining verification tests consistent with the Data Facility and overall DM roadmap. We continued articulating mid-level architecture of LSST Data Facility in the Archi modeling tool, including creation of WBS layered views of logical service constructions and cross-cutting views to show service dependencies.

We finished breaking down planning packages into work packages for the next planning half-cycle.

We participated in the Data Access Working Group. We continued leading the Lossy Compression Working Group. Work in October included developing plans and identifying datasets and pipelines for testing lossy compression.

We continued pursuing incremental improvements to Service Management processes, as well as design of service level monitoring capabilities. This included preparing for 24/7 incident response as needed to support systems in Chile AA Base, documenting incident response processes, and solidify monitoring design with technical staff.

02C.07.06.01 LDF-offered Services

We continued periodic reprocessing of RC datasets in support of biweekly stack release testing and ongoing pipeline development.

02C.07.06.02 Reusable Production Services

We continued supporting the in-place prototype Kubernetes installation to support JupyterLab development. We also continued work towards developing a stable containerized management architecture, beginning to document installation procedures for provisioning new cluster hardware.

02C.07.07 Data and Compute Services

We continued work putting in place a file transfer mechanism toward support of the

Spectrograph test stand, including definition of minimum necessary components.

02C.07.08 LDF Service Software

We continued work on L1 Service software, focusing on support for the DAQ, implementing the first version of Spectrograph controller functionality, unit tests for L1 components, and responding to outcomes of integration activities.

We held a design review for Header Client software design, and began enhancements to the software based on recommendations.

We began Python 3 conversion for DESDM framework batch production code. We made changes as necessary to support ongoing dataset reprocessing to support ongoing periodic reprocessing as well as reprocessing for milestone LDM-503-2.

We prepared software and systems to host and participate in the Pathfinder Integration Activity "Next Visit" milestone and workshop.

02C.07.09 ITC and Facilities

We began acquisitions the first phase of implementing the FY18 acquisition plan, including planning for hardware provisioning. The initial focus was on general enclave infrastructure expansion, Kubernetes cluster planning, Level 1 Camera Control System (CCS) hardware planning, and Level 1 Integration system expansion provisioning. We began improvements to the existing disaster recovery system, as well as making general improvements and additions to existing system monitoring.

Planned activities

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

We will work toward completion of outstanding service management issues, including refinement of the service management change control process.

02C.07.04 Site Infrastructure

Work toward consolidation of ITC management systems will continue with evaluation and testing of third-party packages for managing configuration and management.

Provisioning of remaining hardware systems from the FY17 acquisition plan will continue through November. These include the L1 Test System, the Isst-db database system, and the Base Authentication and Authorization systems.

We will finish an annual audit of NCSA-hosted LSST system user groups.

02C.07.05 LSST Data Facility Management, Service Architecture, and Project Controls

We will continue breaking down replan planning packages into work packages for the next full planning cycle. We will continue articulating mid-level architecture of LSST Data Facility

in the Archi modeling tool.

We will continue work planning for provisioning of FY18 hardware acquisitions, focusing on the L1 system expansion, the infrastructure expansion necessary to provision subsequent systems, and emerging project needs for Kubernetes capability.

We will continue activities leading the Lossy Compression Working Group.

We will continue pursuing incremental improvements to Service Management processes, as well as design of service level monitoring capabilities. This includes preparing for 24/7 incident response as needed to support systems in Chile AA Base as well as general incident and change request response.

02C.07.06.01 LDF-offered Services

We will continue periodic reprocessing of datasets in support of stack testing and pipeline development.

02C.07.06.02 Reusable Production Services

We will continue supporting the in-place prototype Kubernetes installation to support JupyterLab development. We will also continue work towards developing a stable containerized management architecture.

02C.07.07 Data and Compute Services

We will continue follow-on work investigating file transfer mechanism toward support of the Spectrograph test stand.

02C.07.08 LDF Service Software

We will continue work on L1 Service software, focusing on support for the DAQ, implementing the first version of Spectrograph controller functionality, and responding to needs to support upcoming integration activities.

We will continue work on implementing Header Client software, focusing on supporting the Spectrograph test installation in early 2018.

We will continue Python 3 conversion for DESDM framework batch production code.

In support of the LDM-503-2 milestone, we will continue dataset reprocessing and software modification, and provide input into documenting milestone verification,

02C.07.09 ITC and Facilities

We will continue planning for provisioning of the first phase of implementing the FY18 acquisition plan. The initial focus will be planning general enclave infrastructure expansion, Kubernetes cluster installation, Level 1 Camera Control System (CCS), and Level 1 Integration system expansion. We will continue improvements to the existing disaster recovery system. We will begin installing and testing system-level monitoring services for the AA system prior to delivery to Chile, as well as making general improvements and additions to existing system monitoring.

Recruitment update

A requisition for a full-time database administrator is in progress.

1.02C.08: International Communications & Base Site

Current accomplishments

02C.08.01 – Base Center

- Summit Base ITC Tiger Team: No activity. Regular meetings suspended pending completion of Summit and Summit - Base Network installations. Will resume in CY18 with focus on the Base Computer Center.
- Summit and Base Networking and Computing: Based on state of the existing data center and need for a more detailed power and cooling design, we postponed the NCSA visit to Chile to install the security system until February or March, 2018. The Cisco summit network equipment is in storage until occupancy is granted to the Summit Computer Room. Contract additions were made with Elypsa, the electrical contractor at the Summit, to move the Besalco-installed power cable trays from above the racks, to below the raised floor. The trays were moved, and a distribution box and cables for one phase of power were installed. Analysis is underway to revisit the LSE-299 requirement for a second phase of power to the racks, which if implemented, will be cabled before the racks are installed. Wiring the additional phase from the first floor to the Computer Room will occur at a later point in time. On examination of the Besalco-provided raised floor supports, it was determined that additional bracing under the racks is advisable for seismic protection. Ordering commercial supports instead of the delivered supports will take too much time, so we are exploring adding bracing to the supports ourselves. In any case, this will occur in the next week, in order to allow Besalco to complete the as-contracted raised floor and to secure site occupancy from Besalco as soon as possible.

02C.08.02 – Chilean Data Access Center

- No activity this month.

02C.08.03 – Long-Haul Networks

- Jeff Kantor arrived in Santiago, Chile on October 17, attended the South American Astronomy Coordination Committee (SAACC) meeting and the ADASS conference, then continued on to La Serena on October 27, to coordinate the Summit and Summit - Base installations and configuration and test. He will return to Tucson on approximately January 9, 2018.

02C.08.03.01 – Chile National WAN

- Summits–AURA Gatehouse Network: No activity. This segment is accepted.
- DWDM Equipment: Over 80% of the REUNA DWDM equipment is already installed (in Santiago, just need to install La Serena node).
- Santiago–La Serena: No activity. This segment is accepted. The primary link to Level3 is operational.
- La Serena–AURA Gatehouse: No activity. This segment is accepted.

02C.08.03.02 – International Chile–US WAN

- **100 Gbps Managed Ring:** The ring is performing as expected. As a preparation for the 100G installation in Chile, Jeronimo Bezerra performed maintenance at the Level3 Data Center. In total, three loaner modules were installed in the AndesLight1 switch: one 8x10G module (to connect to the 10G nodes for future tests), one 2x100G plus 2x transceivers (to connect to the 100G links), and one High Speed Fabric (to enable the 100G module to achieve maximum throughput). Both switches' firmware were upgraded to the latest 5.8 version (5.8g) and new connections were made to CLARA to improve the out-of-band IP connectivity. Also, a power supply was replaced. A plan for a future management upgrade is being designed to increase network access security and management.
- **Management and Coordination Contract:** The LSST Network Engineering Team (NET) meeting was held on October 12. A new logical network diagram with the most current design of the USA and International Links was presented, showing all the components of the end-to-end network, including dedicated primary, secondary, and back-up paths with internat2 and ESnet. Four of the aggregation sites are located in USA: Miami, Boca Raton, Jacksonville, and Atlanta. AmLight ExPress will be introducing a spectrum connection between the North and South America in 2018. Future protection for natural disasters, like hurricanes, was addressed as well. Multiple scenarios for single point of failure were discussed. Members of the NET attended conferences and presented papers relevant to LSST Networks:
 - South American Astronomy Coordination Committee October 19, 2017 Santiago.
 - ADASS XXVII October 22 - 26, 2017 Santiago.
 - Internet2 Technology Exchange October 15-18 San Francisco.
- **Spectrum Contract:** The Angola Cable contract is signed. Activity centered on finishing Operations and Management contract and payment schedule.
- **US National WAN:** ESNet collaboration negotiations with DOE and with FNAL on an operations role continued, and approval was granted to proceed to technical implementation discussions. FIU continued working with FLR and other fiber providers to implement the most cost effective and greatest throughput for transit

from Florida to Atlanta.

Planned activities

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

- Jeff Kantor will continue in La Serena.

02C.08.01 – Base Center

- Summit to Base ITC Tiger Team: Tiger Team completed its last regularly scheduled meeting in CY2017. Will resume with the focus on the Base Facility after Jeff Kantor returns from Chile in January, 2018.
- Summit and Base Networking and Computing: We are installing an initial configuration of a single Spine Switch, 2 Leaf Switches, and 2 Controllers in La Serena, to start working with the ACI architecture in a sandbox.

02C.08.03 – Long-Haul Networks

02C.08.03.01 – Chile National WAN

- Summits - AURA Gatehouse Network: No activity. This segment is accepted.
- DWDM Equipment: We expect to do the installation and test of the REUNA DWDM and AURA DWDM equipment in La Serena in early November. The LSST DWDM equipment will not be installed until March 2018 or later.
- Santiago-La Serena: No activity. This segment is accepted.
- La Serena - AURA Gatehouse: No activity. This segment is accepted.

02C.08.03.02 – International Chile - US WAN

- 100 Gbps Managed Ring: We are still waiting for CLARA's cross-connects to be installed in Santiago. LAN and AmLight are ready and will be able to proceed as soon as we receive these cross-connects. FIU is waiting for AURA/LSST's confirmation regarding the final selection among 3 options to be installed in Chile.
- Management and Coordination Contract: Coordinate next NET meeting.
- Spectrum Contract: Continue work on the Operations and Maintenance Agreement.
- US National WAN: Discussions with ESNet for LSST within DOE to define the ESNet service and costs will continue. FIU will continue working with FLR and other fiber providers to determine the most cost effective and greatest throughput for transit from AMPATH North.

Recruitment update

- No activity, all open positions have been filled.

1.02C.10: Science Quality and Reliability Engineering

Current accomplishments

Highlights of work completed this month include:

- The [14.0](#) version of the stack and its documentation suite including the current [characterisation report](#) was [publicly released](#). This is the first release offering native binary installation.
- The transition of SQuaSH to a Kubernetes-based deployment is complete and deployed to [squash.lsst.codes](#). This completed epic [DM-10742](#).
- The technote creation capability of sqrobot (the SQuaRE chatbot) was re-implemented as a Celery-based microservice to allow for asynchronous interaction with the user command. This resolves the Travis CI timeouts users were experiencing during technote creation.
- The `validate_drp` package now successfully runs y-band HSC data after a fix to the filter maps.
- Notable external interactions:
 - SK attended the [workshop](#) Searching for Dwarf Companions of the Milky Way and Beyond in the Era of LSST.
 - FE attended the Github Universe conference and visited SLAC.
 - FE participated in a conceptual design review of the Target and Observation Manager project led by Las Cumbres Observatory.

Planned activities

- Port `validate_drp` to the new verification framework.
- Prototype bokeh interactions in JupyterLab.
- Cycle-end activities.

Recruitment update

- None planned for SQuaRE. We participated in Architecture's recruitment of the Release Engineer.