Large Synoptic Survey Telescope (LSST)

Data Management System Design

K.-T. Lim, J. Bosch, G. Dubois-Felsmann, T. Jenness, J. Kantor, W. O'Mullane, D. Petravick, and the DM Leadership Team

LDM-148

Latest Revision: 2017-07-04

Abstract

The LSST Data Management System (DMS) is a set of services employing a variety of software components running on computational and networking infrastructure that combine to deliver science data products to the observatory's users and support observatory operations. This document describes the components, their service instances, and their deployment environments as well as the interfaces among them, the rest of the LSST system, and the outside world.

Change Record

Version	Date	Description	Owner name
2	2011-08-09	Copied from MREFC Proposal into LDM-148 handle, reformatted	Robert McKercher
3	2011-08-15	Updated for Preliminary Design Review	Tim Axelrod, K-T Lim, Mike Freemon, Jeffrey Kantor
4	2013-10-09	Updated for Final Design Review	Mario Juric, K-T Lim, Jeffrey Kantor
5.0	2017-07-04	Rewritten for Construction and Operations. Approved in RFC-358.	K-T Lim

Document curator: Kian-Tat Lim

Document source location: `https://github.com/lsst/LDM-148`

Contents

The contents of this document are subject to configuration control by the LSST DM Technical Control Team.

Data Management System Design

1   Introduction

The purpose of the LSST Data Management System (DMS) is to deliver science data products to the observatory's users and to support observatory operations. The DMS is a set of services employing a variety of software components running on computational and networking infrastructure. The DMS is constructed by the DM subsystem in the NSF MREFC project; in the Operations era, it is operated by a combination of the LSST Data Facility, Science Operations, and Observatory Operations departments.

The data products to be delivered are defined and described in the Data Products Definition Document (LSE-163). These are divided into three major categories.

One category of data products is generated on a nightly or daily cadence and comprises raw, processed/calibrated, and difference images as well as alerts of transient, moving, and variable objects detected from the images, published within 60 seconds, and recorded in searchable catalogs. These data products can be considered "online", as they are driven primarily by the observing cadence of the observatory. This category has historically been referred to as "Level 1". These products are intended to enable detection and follow-up of time-sensitive time-domain events.

A second category of data products is generated on an annual cadence and represents a complete reprocessing of the set of images taken to date to generate astronomical catalogs containing measurements and characterization of tens of billions of stars and galaxies with high and uniform astrometric and photometric accuracy. As part of this reprocessing, all of the first category of data products is regenerated, often using more accurate algorithms. This category also includes other data products such as calibration products and templates that are generated in an "offline" mode, not directly tied to the observing cadence. This category has historically been referred to as "Level 2", including the regenerated data products from the first category.

The third category of data products is not generated by the LSST DMS but is instead generated, created, or imported by science users for their own science goals. These products derive value from their close association with or derivation from other LSST data products.

The DMS is responsible for providing facilities, services, and software for their generation and storage. This category has historically been referred to as "Level 3".

Data products are delivered to science users through Data Access Centers (DACs). In addition, streams of near-realtime alerts and telescope pointing predictions are provided. Each LSST data product has associated metadata providing provenance and quality metrics and tracing it to relevant calibration information in the archive. The DACs are composed of modest but significant computational, storage, networking, and other resources intended for use as a flexible, multi-tenant environment for professional astronomers with LSST data rights to retrieve, manipulate, and annotate LSST data products in order to perform scientific discovery and inquiry.

The first section of this document describes how the DMS components work together to generate and distribute the data products. The next section describes how the size of the DMS computing environments was estimated. Subsequent sections describe the individual components of the DMS in more detail, including their interfaces with each other, with other LSST subsystems, and with the outside world.

2   Summary Concept of Operations

The principal functions of the DMS are to:

- Process the incoming stream of images generated by the camera system during observing by archiving raw images, generating transient alerts, and updating difference source and object catalogs.

- Periodically (at least annually) process the accumulated survey data to provide a uniform photometric and astrometric calibration, measure the properties of fainter objects, and characterize the time-dependent behavior of objects. The results of such a processing run form a data release (DR), which is a static, self-consistent data set for use in performing scientific analysis of LSST data and publication of the results. All data releases are archived for the entire operational life of the LSST archive.

- Periodically create new calibration data products, such as bias frames and flat fields, to be used by the other processing functions.

- Make all LSST data available through an interface that utilizes, to the maximum possible extent, community-based standards such as those being developed by the Virtual Observatory (VO) in collaboration with the International Virtual Observatory Alliance (IVOA). Provide enough processing, storage, and network bandwidth to enable user analysis of the data without petabyte-scale data transfers.

The latency requirements for alerts determine several aspects of the DMS design and overall cost. An alert is triggered by an unexpected excursion in brightness of a known object or the appearance of a previously undetected object such as a supernova or a GRB. The astrophysical time scale of some of these events may warrant follow-up by other telescopes on short time scales. These excursions in brightness must be recognized by the pipeline, and the resulting alert data product sent on its way, within 60 seconds. This drives the DMS design in the decision to acquire high-bandwidth/high-reliability long-haul networking from the Summit at Cerro Pachon to the Base in La Serena and from Chile to the U.S. These networks allow the significant computational resources necessary for promptly processing incoming images to be located in cost-effective locations: the Base has far fewer limitations on power, cooling, and rack space capacity than the Summit, and placing the scientific processing at NCSA allows for far greater flexibility in the allocation of resources to ensure that deadlines are met. Performing cross-talk correction on the data in the data acquisition system and parallelizing the alert processing at the amplifier and CCD levels, where possible, also help to minimize the latency to alert delivery.

The Data Release processing requires extensive computation, combining information from all images of an object in order to measure it as accurately as possible. A sophisticated workload and workflow management system and Task Framework are used to divide the processing into manageable units of work that can be assigned to available resources, including the two dedicated processing clusters at NCSA and CC-IN2P3.

Calibration data products must be created and updated at cadences in between the Alert and Data Release periods. The stability of the system is expected to require daily, monthly, and annual calibration productions. The daily production must be synchronized with the observatory schedule, occurring after raw calibration frames have been taken but well before science observing is planned. This requirement necessitates the inclusion of a service that allows the Observatory Control System to trigger remote calibration processing at NCSA.

The DACs are a key component of the DMS, giving the community resources and an interface to interact with and utilize the LSST data products to perform science. An instance of the LSST Science Platform (LSP) is deployed in each DAC to support science users with its Portal, JupyterLab (notebook), and Web API aspects. Substantial compute, storage, and storage bandwidth is devoted to ensuring that the LSP is responsive and allows for exploration of the vast LSST data products.

Underlying all of the above is a Data Backbone that provides storage, tracking, and replication for all LSST data products. The Data Backbone links all of the computational environments and the Data Access Centers, acting as the spine that supports them all.

3   Sizing

A fundamental question is how large the LSST Data Management System must be. To this end, a complex analytical model has been developed driven by input from the requirements specifications. Specifications from the science requirements and other subsystem designs, and the observing strategy, translate directly into numbers of detected sources and astronomical objects, and ultimately into required network bandwidths and the size of storage systems. Specific science requirements of the survey determine the data quality that must be maintained in the DMS products, which in turn determine the algorithmic requirements and the computer power necessary to execute them. The relationship of the elements of this model and their flow-down from systems and DMS requirements is shown in Figure 1. Detailed sizing computations and associated explanations appear in LSST Documents listed on the Figure.

Key input parameters include camera characteristics, the expected cadence of observations, the number of observed stars and galaxies expected per band, the processing operations per data element, the data transfer rates between and within processing locations, the ingest and query rates of input and output data, the alert generation rates, and latency and throughput requirements for all data products.

Processing requirements were extrapolated from the functional model of operations, prototype pipelines and algorithms, and existing precursor pipelines adjusted to LSST scale. As a part of every Data Release, all data previously processed are reprocessed with the latest algorithms, calibration products, and parameters. This causes the processing requirements to

DM System Design
LDM-148
DM Applications Design
LDM-151
DM Middleware Design
LDM-152
DM Database Design
LDM-135

LSST Science and Project Sizing Inputs
LSE-81
Explanation:
LSE-82

LSST Science Requirements
LPM-17
LSST System Requirements
LSE-29
Observatory System Specs
LSE-30
DM Subsystem Requirements
LSE-61
Data Products Definition Document
LSE-163
Community Access White Paper
Document-5373

DM Compute Sizing
LDM-138
Explanation:
LDM-140

DM Storage & I/O Sizing
LDM-141
Explanation:
LDM-139

Network Sizing
LDM-142
Explanation:
LSE-78

DM Infrastructure Design
LDM-129

DM Site Specific Infrastructure Architecture/Design/Costs
**LDM-144**
Explanation:
LDM-143

FIGURE 1: DMS Infrastructure Sizing and Estimation.

TABLE 1: DMS Compute Infrastructure Sizing; growth from Survey Year 1 to Year 10; hwm = high-water mark

		Archive Site	Base Site
Compute	TeraFLOPS (sustained)	197 → 970	30 → 62
	Nodes	436 → 305 (455 hwm)	56 → 17 (59 hwm)
	Cores	18K → 62K	3K → 4K
Database	TeraFLOPS (sustained)	40 → 310	55 → 306
	Nodes	94 → 113 (148 hwm)	108 → 98 (133 hwm)
Facilities	Floor Space	826 → 744 ft^2 (834 hwm)	278 → 195 ft^2 (435 hwm)
	Power	274 → 273 kW (309 hwm)	158 → 248 kW (248 hwm)
	Cooling	0.9 → 0.9 mmbtu (1.1 hwm)	0.5 → 0.8 mmbtu (0.8 hwm)

TABLE 2: DMS Storage Infrastructure Sizing; growth from Survey Year 1 to Year 10

		Archive Site	Base Site
File Storage	Capacity	24 → 81 PB	
	Drives	1602 → 862	597 → 249
	Bandwidth	493 → 714 GB/s (752 hwm)	223 → 231 GB/s (236 hwm)
Database	Capacity	29 → 99 PB	16 → 72 PB
	Drives	3921 → 2288	2190 → 1642
	Bandwidth	1484 → 2040 GB/s (2163 hwm)	829 → 1169 GB/s (1615 hwm)
Tape Storage	Capacity	31 → 242 PB	
	Tapes	2413 → 3691 (4117 hwm)	
	Tape Bandwidth	36 → 65 GB/s	

increase with time. Advances in hardware performance, however, are expected to reduce the number of nodes needed and the power and cooling devoted to them. This causes some of the performance figures in Table 1 to reach a high-water mark during the survey.

Storage and input/output requirements were extrapolated from the data model of LSST data products, the DMS and precursor database schemas, and existing database management system overhead factors in precursor surveys and experiments adjusted to LSST scale. A summary of key numbers is in Table 2.

Communications requirements were developed and modeled for the data transfers and user query/response load, extrapolated from existing surveys and adjusted to LSST scale. These requirements are illustrated in Figure 2 for per-visit transfers. Peak bandwidths assume a 3 second budget for Summit to Base image transfer and a 5 second budget for international image transfer.

The Summit to Base and Base to NCSA network links have been significantly over-engineered

for four main reasons: first, because the incremental costs of higher bandwidth once the link has been provisioned at all have been small; second, to allow key functions, such as interfacing with the Camera Data System or performing image analysis and measurement to generate alerts, to be performed in appropriate locations; third, to increase reliability of the system; and fourth, to simplify certain components such as the Forwarders and Archivers that interface with or depend on the networks. A high-speed Base to NCSA link also enables Data Releases to be transferred south to the Chilean DAC over the network rather than through physical media, as originally planned, decoupling science from maintenance and upgrade activities to a greater extent.

FIGURE 2: Near Real-Time Data Flows

In all of the above, industry-provided technology trends (LDM-143) were used to extrapolate to the LSST construction and operations phases in which the technology will be acquired, configured, deployed, operated, and maintained. A just-in-time acquisition strategy is employed to leverage favorable cost/performance trends.

The resulting performance and sizing requirements show the DMS to be a supercomputing-
-class system with correspondingly large data input/output and network bandwidth rates.
Despite this size, technology trends show this to be well within the anticipated performance
of commodity-based systems during the construction and operations time frame.

4 Component Overview

The services that make up the DMS are in turn made up of software and underlying service
components, instantiated in a particular configuration in a particular computing environment
to perform a particular function. Some software components are specific to a service; oth-
ers are general-purpose and reused across multiple services. Many services have only one
instance in the production system; others have several, and all have additional instances in
the development and integration environments for testing purposes.

The DMS services can be considered to consist of four tiers of software components. The
top tier is the LSST Science Platform, which is deployed in the DACs and other computational
environments to provide a user interface and analysis environment for science users and
LSST staff. The detailed design of this tier is given in "LSST Science Platform," LDM-542. The
next tier is composed of science "applications" software that generates data products. This
software is used to build "payloads", sequences of pipelines, that perform particular data
analysis and product generation tasks. It is also used by science users and staff to analyze
the data products. The detailed design of the components in this tier is given in Data Manage-
ment Science Pipelines Design, LDM-151. A lower tier is "middleware" software components
and services that execute the science application payloads and isolate them from their en-
vironment, including changes to underlying technologies. These components also provide
data access for science users and staff. The detailed design of the components in this tier is
given in Data Management Middleware Design, LDM-152. The bottom tier is "infrastructure":
hardware, networking, and low-level software and services that provide a computing environ-
ment. The detailed design of components in this tier is given in Data Management Services &
Infrastructure, LDM-129, and Network Design, LSE-78.

The DMS computing environments reside in four main physical locations: the Summit Site
including the main Observatory and Auxiliary Telescope buildings on Cerro Pachon, Chile;
the Base Facility data center located at the Base Site in La Serena, Chile; the Archive Facility
data center at the National Center for Supercomputing Applications (NCSA) in Urbana, Illi-

nois, USA; and the Satellite Computing Facility at CC-IN2P3 in Lyon, France. These are linked by high-speed networks to allow rapid data movement. The Base Facility includes the Base Center production computational environments and also the Chilean Data Access Center. The Archive Facility also hosts the US Data Access Center. In addition, a Commissioning Cluster computational environment also resides at the Base Facility.

The DMS service instances at NCSA can be broken down into three main functional domains: a near-realtime online domain (L1) closely linked to the rest of the Observatory; an offline Level 2 domain (L2) organized primarily around the annual Data Release Production; and an analysis and developer support domain (ADS) encompassing environments that operations staff use for science validation, software development, system integration, and system testing. In addition, an underlying infrastructure domain (Infra) hosts services supporting all of the other domains, including a common Data Backbone that provides data transport and archiving and that is the primary connection between all of the domains. These domains are distinguished by having different users, operations timescales, interfaces, and often components.

The service instances that make up the DMS include (with the computational environment or domains they are in noted):

- Image and EFD Archiving (Base)

- Prompt Processing Ingest (Base)

- Observatory Control System (OCS) Driven Batch Control (Base)

- Telemetry Gateway (Base)

- Prompt Processing (L1)

- OCS Driven Batch Processing (L1)

- Offline Processing (L1)

- Alert Distribution (L1)

- Alert Filtering (L1)

- Level 1 Quality Control (QC) (L1)

- Template and Calibration Products Production Execution (L2)

The contents of this document are subject to configuration control by the LSST DM Technical Control Team.

- Data Release Production Execution (L2)

- Data Release Production Satellite Processing (Satellite Computing)

- Level 2 QC (L2)

- LSST Science Platform Commissioning Cluster instance (Commissioning Cluster)

- LSST Science Platform Data Access Center instances (DACs)

- Bulk Data Distribution (DAC)

- LSST Science Platform Science Validation instance (ADS)

- Developer Services (ADS)

- Integration and Test (ADS)

- Data Backbone (Infra)

- Management/Monitoring (Infra)

- Provisioning/Deployment (Infra)

- Workload/Workflow (Infra)

- HTCondor Batch Processing (Infra)

- Identity Management (Infra)

The relationships between these services, their deployment environments, functional domains, and science application "payloads" can be visualized in Figure 3.

The common infrastructure services are illustrated in Figure 4.

The science application software for the Alert Production, daytime processing, Data Release Production, and calibration processing is built out of a set of frameworks that accept plugins. In turn, those frameworks build on middleware that provides portability and scalability. The relationships between the packages implementing these se frameworks and plugins and the underlying middleware packages are shown in Figure 5.

Key applications software components include:

FIGURE 3: Data Management System Deployment
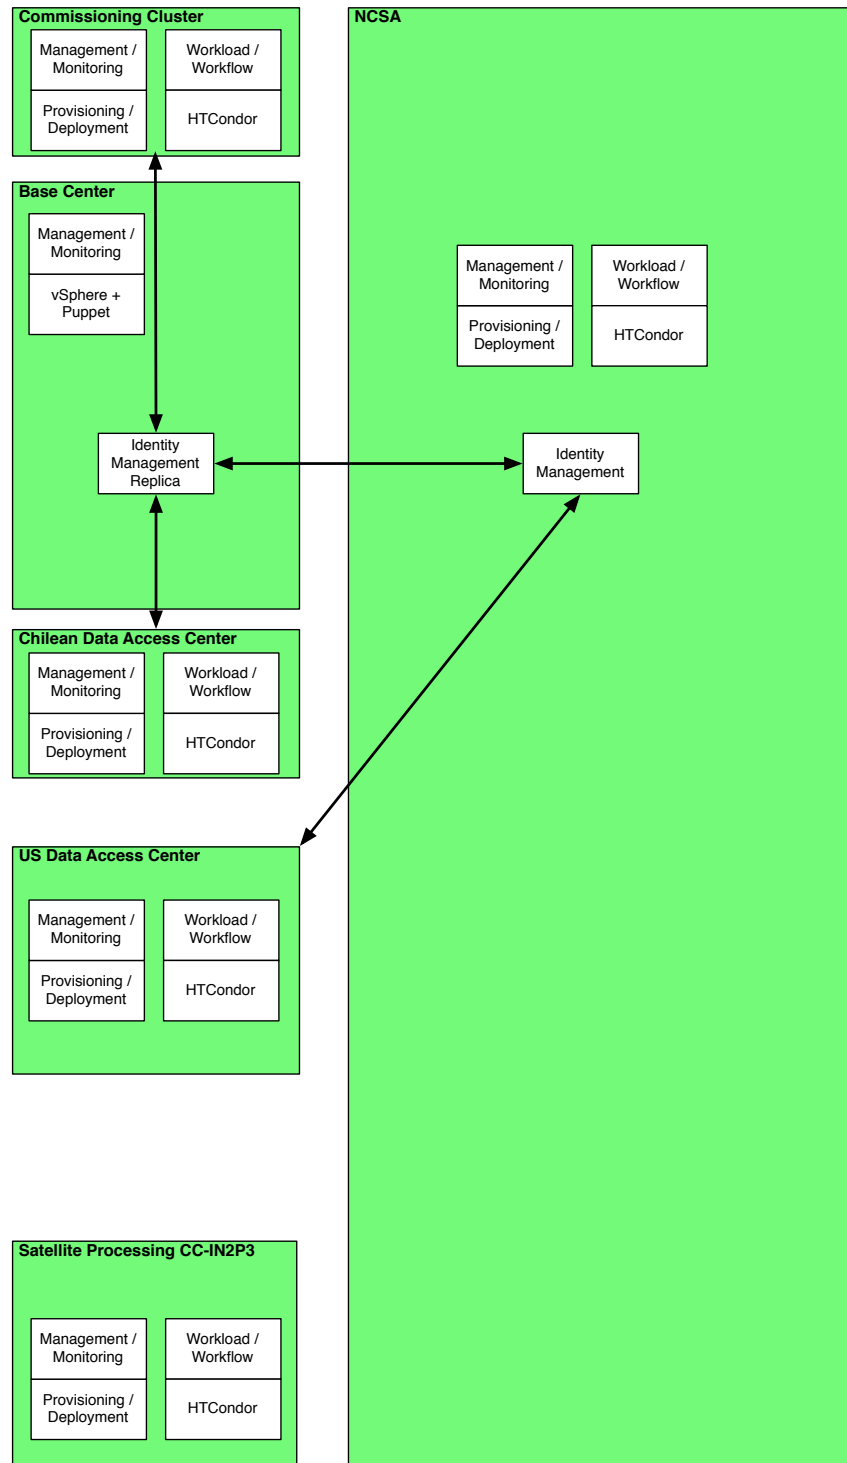
FIGURE 4: Data Management System Common Infrastructure Services

- Low-level astronomical software primitives and data structures (`afw`)

- Image processing and measurement framework with core algorithms (`ip_*`, `meas_*`)

- Additional image processing and measurement algorithms (`meas_extensions_*`)

- High-level algorithms and driver scripts that define pipelines (`pipe_tasks`, `pipe_drivers`)

- Camera-specific customizations (`obs_*`)
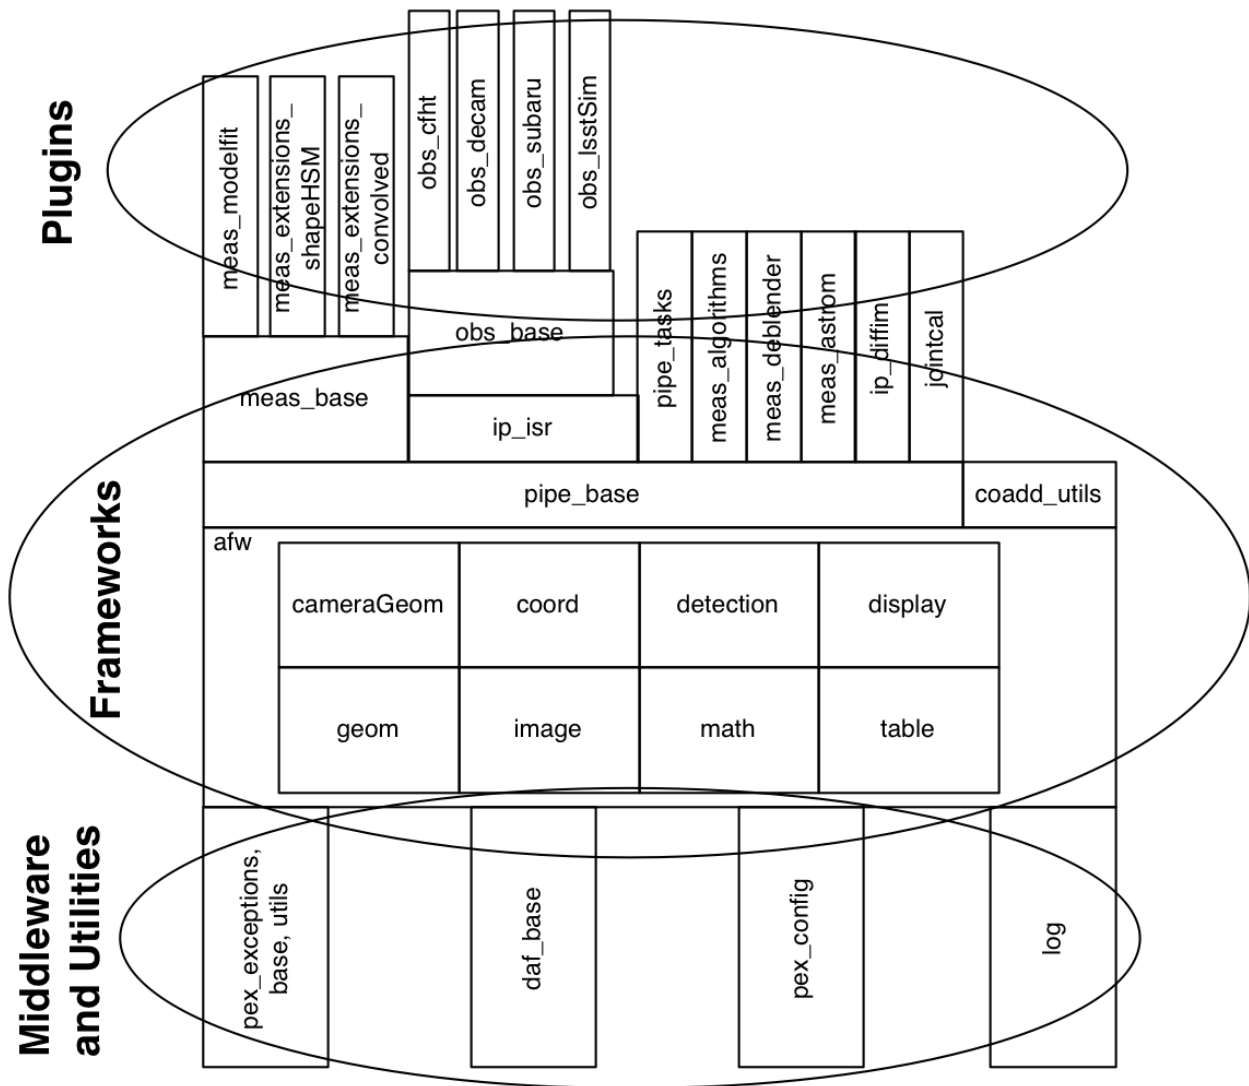
FIGURE 5: Data Management Science Pipelines Software "Stack"

Key middleware components include:

- Data access client (Data Butler) (`daf_persistence`)

- Parallel distributed database (`qserv`)

- Task framework (`pex_*`, `log`, `pipe_base`, `ctrl_pool`)

- Workflow and orchestration for production control (`ctrl_*`)

Infrastructure components include:

- Other databases (typically relational)

- Filesystems

- Authentication and authorization (identity management)

- Provisioning and resource management

- Monitoring

The relationships between the middleware and infrastructure components are illustrated in Figure 6.

5   Base Center

Services located in this environment are located at the Base solely because they must interact with the OCS or the Camera Data System (also known as the Camera DAQ) or both. In several cases, services located here interact closely with corresponding services in NCSA's Level 1 Domain, to the point where the Base service cannot function if the NCSA service is not operational. This reliance has been taken into account in the fault tolerance strategies used.

The primary goals of the services in this environment are to transfer data to appropriate locations, either to NCSA, from NCSA, or to the Data Backbone.

The services in this environment and their partners in the NCSA Level 1 Domain need to run rapidly and reliably. They run at times (outside office hours) and with latencies that are not amenable to a human-in-the-loop design. Instead, they are designed to execute
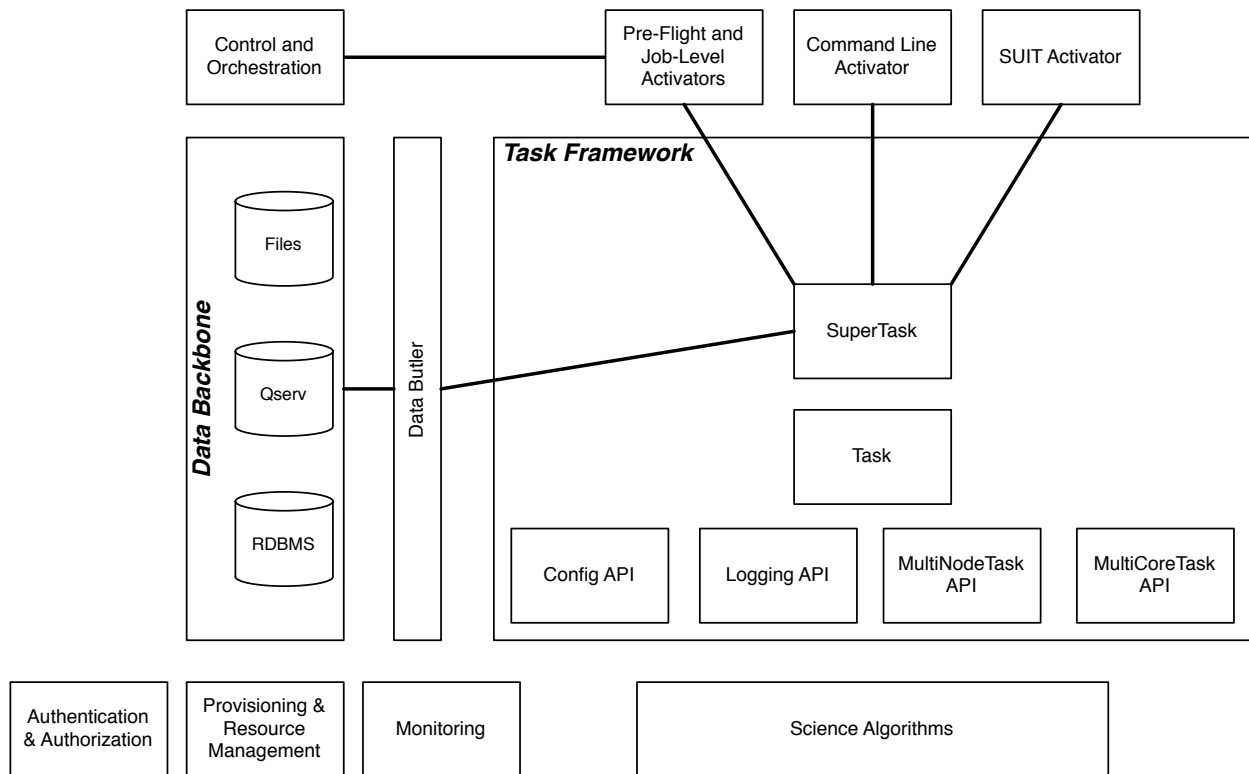
FIGURE 6: Data Management Middleware and Infrastructure

autonomously, often under the control of the OCS, with human oversight, monitoring, and control only at the highest level.

5.1 Service Descriptions

Detailed concepts of operations for each service can be found in "Concept of Operations for the LSST Production Services" (LDM-230).

5.1.1 Image and EFD Archiving

This component is composed of several Image Archiving service and Catch-Up Image Archiving instances: one pair each for the LSSTCam, the ComCam, and the Auxiliary Telescope Spectrograph, all of which may be operated simultaneously. These capture raw images taken by each camera, including the wavefront sensors and the guide sensors of the LSSTCam or ComCam when so configured, retrieving them from their respective Camera Data System instances. They also capture specific sets of metadata associated with the images, including telemetry values and event timings, from the OCS publish/subscribe middleware and/or from the EFD. The image pixels and metadata are then permanently archived in the Data Backbone via the Observatory Operations Data Service (OODS). The catch-up versions archive into the OODS and Data Backbone any raw images and metadata that were missed by the primary archiving services due to network or other outage, retrieving them from the flash storage in the Camera Data System instances and the EFD.

This component also includes an EFD Transformation service that extracts all information (including telemetry, events, configurations, and commands) from the EFD and its large file annex, transforms it into a form more suitable for querying by image timestamp, and loads it into the permanently archived "Transformed EFD" database in the Data Backbone.

5.1.2 Observatory Operations Data Service

This component provides low-latency access to images and metadata for use by Observatory systems and the Commissioning Cluster LSP instance. It maintains a higher level of service availability than the Data Backbone and so provides a staging area for to-be-archived data prior to ingestion into the Data Backbone. After images and metadata are ingested, they remain available through the OODS for a policy-configured amount of time.

5.1.3   Prompt Processing Ingest

This component is composed of two instances that capture crosstalk-corrected images from the LSSTCam and ComCam Camera Data Systems along with selected metadata from the OCS and/or EFD and transfer them to the Prompt Processing service in the NCSA Level 1 Domain.

There is no Prompt Processing Ingest instance for the auxiliary telescope spectrograph.

5.1.4   OCS Driven Batch Control

This service receives commands from the OCS and invokes the OCS Driven Batch Processing service in the NCSA Level 1 Domain to execute corresponding science payloads. It is used for modest-latency analysis of images during Commissioning and for processing daily calibration images in normal observing operations. A summary status for the processing performed is returned to the OCS for each command, following the normal OCS commanding protocol.

5.1.5   Telemetry Gateway

This service obtains information from the NCSA Level 1 Domain, particularly status and quality metrics from Prompt Processing of images and the Level 1 Quality Control service, and transmits it to the OCS as specified in the Data Management-OCS Software Communication Interface (LSE-72). Note that more detailed information on the status and performance of DMS services will also be available to Observatory operators through remote displays originated from the Management/Monitoring infrastructure services in all DMS computational environments.

5.2   Interfaces

OCS to all Base Center services: these interface through the SAL library provided by the OCS subsystem.

Archiver and Catch-Up Archiver to Observatory Operations Data Service: image files with associated metadata are written to OODS storage.

Observatory Operations Data Service to Data Backbone: files are copied to Data Backbone

The contents of this document are subject to configuration control by the LSST DM Technical Control Team.

17

storage via a file transfer mechanism, and their information and metadata are registered with Data Backbone management dataabases. The Data Butler is not used for this low-level, non-science-payload interface.

Observatory Operations Data Service to Commissioning Cluster LSP and other users: the OODS provides a mountable POSIX filesystem interface. A web interface (e.g. WebDAV) may also be provided, but this will be as simple as possible to enable maintaining a very high level of service availability and reliability.

EFD to EFD Transformer: this interface is via connection to the databases that make up the EFD as well as file transfer from the EFD's Large File Annex.

EFD Transformer to Data Backbone: Transformed EFD entries are inserted into the "Transformed EFD" database resident within the Data Backbone.

Camera Data System to Archiver, Catch-Up Archiver, Prompt Processing Ingest: these interface through the custom library provided by the Camera Data System.

Prompt Processing Ingest to Prompt Processing: BBFTP is used to transfer files over the international network from the ingest service to the processing service.

OCS Driven Batch Control to OCS Driven Batch Processing: HTCondor is used to transfer execution instructions over the international network from the control service to the processing service and return status and result information.

Telemetry Gateway from NCSA Level 1 Domain services: RabbitMQ is used to transfer status and quality metrics to the gateway over the international network.

6 NCSA Level 1 Domain

This domain is responsible for the compute-intensive processing for all near-realtime operations and other operations closely tied with the Observatory. Its primary goals are to process images and metadata from the Observatory into "online" science data products and publish them to the DACs, alert subscribers, and back to the OCS.

The Prompt Processing, OCS Driven Batch Processing, and Offline Processing services support execution of science payloads in three different ways. Prompt Processing is tightly integrated with the observing cadence and is intended to function in near-realtime with strict result deadlines. OCS Driven Batch Processing is invoked by the OCS but has more modest latency requirements. Offline Processing is not invoked by the OCS but operates under DMS control, typically during the daytime.

The Alert Distribution and Alert Filtering services receive batches of alerts resulting from Prompt Processing of each scienc visit; they then provide alert streams to community alert brokers and LSST data rights holders, respectively.

The Level 1 Quality Control service monitors the "online" science data products, including alerts, notifying operators if any anomalies are found.

Like the services in the Base Center, these services need to run rapidly and reliably and so are designed to execute autonomously.

6.1   Service Descriptions

Detailed concepts of operations for each service can be found in "Concept of Operations for the LSST Production Services" (LDM-230).

6.1.1   Prompt Processing

This service receives crosstalk-corrected images and metadata from the Prompt Processing Ingest service at the Base and executes the Alert Production science payload on them, generating "online" data products that are stored in the Data Backbone. The Alert Production payload then sends alerts to the Alert Distribution service.

The Prompt Processing service has calibration (including Collimated Beam Projector images), science, and deep drilling modes. In calibration mode, it executes a Raw Calibration Validation payload that provides rapid feedback of raw calibration image quality. In normal science mode, two consecutive exposures are grouped and processed as a single visit. Definitions of exposure groupings to be processed as visits in deep drilling and other modes are TBD. The service is required to deliver Alerts within 60 seconds of the final camera readout of a

standard science visit with 98% reliability.

There is no Prompt Processing service instance for the Auxiliary Telescope Spectrograph.

6.1.2   OCS Driven Batch Processing

This service executes science payloads in response to commands from the OCS Driven Batch Control service at the Base and thus indirectly from the Observatory Control System. It is used for modest-latency analysis of images during Commissioning and for processing daily calibration images in normal observing operations. Images and metadata are taken from the Data Backbone, and results are provided back to the Data Backbone; there is no direct connection from this service to the Camera Data System. This obviously bounds the minimum latency from image acquisition to processing start by the latency of the Archiving service and Data Backbone transfer. A summary status for the processing performed is sent to the OCS Driven Batch Control service to be returned to the OCS.

6.1.3   Offline Processing

This service executes science payloads to ensure that all Level 1 data products are generated within 24 hours. In particular, this service executes the daytime Moving Object Processing System payload. It also may execute a variant of the Alert Production payload if the Prompt Processing service encounters difficulties. Images and metadata are taken from the Data Backbone, and results are provided back to the Data Backbone.

6.1.4   Level 1 Quality Control

This service collects information on Level 1 science and calibration payload execution, post-processes the science data products from the Data Backbone to generate additional measurements, and monitors the measurement values against defined thresholds, providing an automated quality control capability for potentially detecting issues with the environment, telescope, camera, data acquisition, or data processing. Alarms stemming from threshold crossings are delivered to Observatory operators and to LSST Data Facility Production Scientists for verification, analysis, and resolution.

6.1.5   Alert Distribution

This service obtains alerts generated by the Alert Production science payload and distributes them to community alert brokers and to the Alert Filtering service.

6.1.6   Alert Filtering

This service obtains an alert feed from the Alert Broker Feed service and allows individual LSST data rights holders to execute limited filters against it, producing filtered feeds that are then distributed to the individuals.

6.2   Interfaces

Prompt Processing to Alert Distribution and Alert Filtering: these interface through a reliable transport system.

Prompt Processing to Offline Processing: in the event that Prompt Processing runs over its allotted time window, processing can be cancelled and the failure recorded, after which Offline Processing will redo the processing at a later time. Note that it may be possible, if sufficient computational resources have been provisioned, for the Prompt Processing to be allowed to continue to run, with spare capacity used to maintain latency for future visits. In that case, there would effectively be an infinite time window.

Science Payloads to Data Backbone: payloads use the Data Butler as a client to access files and catalog databases within the Data Backbone.

7   NCSA Level 2 Domain

This domain is responsible for all longer-period data processing operations, including the largest and most complex payloads supported by the DMS: the annual Data Release Production (DRP) and periodic Calibration Products Productions (CPPs). Note that CPPs will execute even while the annual DRP is executing, hence the need for a separate service. The Level 2 Quality Control Service monitors the science data products, notifying operators if any anomalies are found.

The services in this domain need to run efficiently and reliably over long periods of time, spanning weeks or months. They need to execute millions or billions of tasks when their input data becomes available while tracking the status of each and preserving its output. They are designed to execute autonomously with human oversight, monitoring, and control primarily at the highest level, although provisions are made for manual intervention if absolutely necessary.

This domain does not have direct users (besides the operators of its services); the services within it obtain inputs from the Data Backbone and place their outputs into the Data Backbone.

7.1 Service Descriptions

7.1.1 Template and Calibration Products Production Execution

This service executes various CPP science payloads at intervals to generate Master Calibration Images and populate the Calibration Database with information derived from analysis of raw calibration images from the Data Backbone and information in the Transformed EFD. This includes the computation of crosstalk correction matrices. Although not a calibration product, the templates used by Alert Production are also generated by this service, based on raw science images from the Data Backbone. Additional information such as external catalogs are also taken from the Data Backbone. The intervals at which this service executes will depend on the stability of Observatory systems but are expected to include at least monthly and annual executions. The annual execution is a prerequisite for the subsequent execution of the Data Release Production. The service involves human scientist/operator input to determine initial configurations of the payload, to monitor and analyze the results, and possibly to provide additional configuration information during execution.

7.1.2 Data Release Production Execution

This service executes the DRP science payload annually to generate all Level 2 data products after the annual CPP is executed. A small-scale (about 10% of the sky) mini-production is executed first to ensure readiness, followed by the full production. Raw science images are taken from the Data Backbone along with Master Calibration Images and information from the Transformed EFD. Additional information such as external catalogs may also be taken

---

from the Data Backbone. Computing is performed in conjunction with the DRP Satellite Processing service at CC-IN2P3, which will have capacity for half of the DRP processing. Output data products from both the mini-production and the main production are loaded into the Data Backbone, including both images and catalogs. From there, they are analyzed by LSST staff scientists and selected external scientists using the Science Validation instance of the LSST Science Platform to ensure quality and readiness for release. The to-be-released data products are loaded into the Data Access Center services, and access is then enabled on the release date. The service involves human scientist/operator/programmer input to determine initial configurations of the payload, to monitor and analyze results, and, when absolutely necessary, to make "hot fixes" during execution that maintain adequate consistency of the resulting data products.

7.1.3 Level 2 Quality Control

This collects information on Level 2 science payload execution, post-processes the science data products from the Data Backbone to generate additional measurements, and monitors the measurement values against defined thresholds, providing an automated quality control capability for potentially detecting issues with the data processing but also the environment, telescope, camera, or data acquisition. Alarms stemming from threshold crossings are delivered to LSST Data Facility Production Scientists for verification, analysis, and resolution.

7.2 Interfaces

Calibration Products Production Execution and Data Release Production Execution to Data Backbone: for large-scale productions, a workflow system is expected to stage files and selected database entries from the Data Backbone to local storage for access by the science payloads via the Data Butler. Similarly, the staging system will ingest output images and catalogs into the Data Backbone.

8 Satellite Computing Center

8.1 Service Description

8.1.1   Data Release Production Satellite Processing

This service controls the processing of jobs on the CC-IN2P3 satellite computing facilities under the overall workload and workflow management of the Data Release Production Execution service at NCSA.

8.2   Interfaces

DRP Satellite Processing to DRP Execution: DRP Execution will stage data to the Satellite Processing facility. All raw data, including images, metadata, and the Transformed EFD, will be transferred. Intermediate data products will be transferred back to the NCSA Level 2 Domain.

9   Data Access Centers

There are two Data Access Centers, one in the US at NCSA and one in Chile at the Base. These DACs are responsible for all science-user-facing services, primarily instances of the LSST Science Platform (LSP). The LSP is the preferred analytic interface to LSST data products in the DAC. It provides computation and data access on both interactive and asynchronous timescales. The US DAC also includes a service for distributing bulk data on daily and annual (Data Release) timescales to partner institutions, collaborations, and LSST Education and Public Outreach (EPO).

The services in this domain must support multiple users simultaneously and securely. The LSP must be responsive to science user needs; updates are likely to occur at a different cadence from the other domains as a result. The LSP must operate reliably enough that scientific work is not impeded.

9.1   Service Descriptions

9.1.1   Bulk Data Distribution

This service is used to transmit Level 1 and Level 2 data products to partners such as LSST Education and Public Outreach, the UK LSST project, and the Dark Energy Science Collaboration. It extracts data products from the Data Backbone and transmits them over high bandwidth

---

connections to designated, pre-subscribed partners.

9.1.2   LSST Science Platform DAC instances

This service provides an exploratory analysis environment for science users. It can be further broken down into three "Aspects" that it presents to end users, along with underlying "backend services" that users can take advantage of, as illustrated in Figure 7.
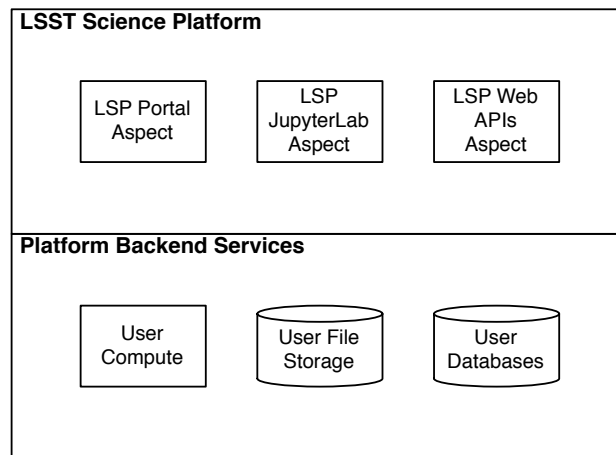
FIGURE 7: LSST Science Platform

The "Portal" Aspect provides a pre-specified yet flexible discovery, query, and viewing tool. The "JupyterLab" Aspect provides a fully flexible ("notebook") environment incorporating rendering of images, catalogs, and plots and providing for execution of LSST-provided and custom algorithms. The "Web API" Aspect provides a language-independent, VO-compliant Web Services data access API with extensions for LSST capabilities and volumes. Access is provided via all three Aspects to all data products, including images, catalogs, and metadata. The Web API Aspect regenerates "virtual" data products on demand when required.

The backend services provide general-purpose user computation, including batch job submission; user file storage accessible to all three Aspects; and user database storage for relational tables. Data may be shared with individual users, with groups, or with all DAC users (data rights holders). Resource management of the backend services is based on a small "birthright" quota with additional resources allocated by a committee.

LSST Science Platform instances access read-only data products, including files and databases, from the Data Backbone. In addition, files may be cached within the LSP instance for speed,

and large-scale catalogs are typically loaded into an instance of the Qserv database management system for efficient query and analysis.

All usage of any LSST Science Platform instance requires authentication to ensure availability only to LSST data rights holders or LSST operations staff.

9.2   Interfaces

Bulk Data Distribution and LSST Science Platform to Data Backbone: Both DAC-resident services retrieve their data, including raw images, nightly and annual image and catalog data products, metadata, and provenance, from the Data Backbone. The LSP Portal Aspect uses the LSP Web APIs to retrieve data. The LSP JupyterLab Aspect can use the LSP Web APIs and also can use the Data Butler client library to access the Data Backbone.

Bulk Data Distribution to partners: The exact delivery mechanism for large-scale data distribution is TBD.

10   NCSA Analysis and Developer Support Domain

This domain encompasses environments for analysts, developers, and integration and test. Its users are the Observatory staff as they analyze raw data and processed data products to characterize them, develop new algorithms and systems, and test new versions of components and services before deployment.

10.1   Service Descriptions

10.1.1   LSST Science Platform Science Validation instance

This instance of the LSST Science Platform is customized to allow access to unreleased and intermediate data products from the Alert, Calibration Products, and Data Release Productions. It is optimized for usage by scientists within the LSST Operations team, although selected external scientists can be granted access to assist with Science Validation. Part of the optimization is to size and configure the three Aspects of the LSP appropriately; in particular, more JupyterLab usage and less portal usage is expected.

10.1.2   Developer Services

Software version control service, packaging, build and unit test service, software release management, ticket tracking service, documentation services, etc.

10.1.3   Integration and Testing

Integration environments representing various deployment environments, deployment services, test datasets, test execution services, metric measurement and tracking services, etc. This service includes the DM instance of the Camera DAQ Test Stand, used in an integration environment simulating the Base Center, as well as the Prototype Data Access Center (PDAC), which is an integration environment simulating a Data Access Center.

10.2   Interfaces

LSP Science Validation instance and Integration and Testing to Data Backbone: All three services in this domain interface with the Data Backbone. The LSP Science Validation instance is used to inspect, analyze, and validate the data products of the Data Release Production prior to their release and so has access to those products in the Data Backbone; since it may be used to annotate the data products, it can also write to the Data Backbone. The Integration and Testing environment may use raw data, intermediate data products, and final data products to perform tests.

Developer Services do not have direct interfaces with the rest of the operational system; they communicate via the distributed source version control system, the package management system, and the configuration system. There is a separate Developer Services instance of the Data Backbone to manage and track datasets used during development.

11   Commissioning Cluster

11.1   Service Description

11.1.1 LSST Science Platform Commissioning instance

This instance of the LSST Science Platform for Science Validation runs on the Commissioning Cluster at the Base Facility (but also has access to computational resources at the Archive) and accesses a Base endpoint for the Data Backbone. This location at the Base lowers the latency of both access to Data Backbone-resident data (which does not have to wait for transfer over the international network) and, perhaps more importantly, for user interface operations for staff in Chile, which are served locally. Note that the Commissioning Cluster does not have direct access to the Camera Data System; it relies on the Archiver service to obtain data. The Commissioning Cluster will have direct access to the OCS's Base replica of the EFD (before transformation).

11.2   Interfaces

Commissioning Cluster to Data Backbone: The Commissioning Cluster relies on the Data Backbone for its data, like the other instances of the LSST Science Platform.

Commissioning Cluster to EFD: The Commissioning Cluster has direct read-only client access to the Base replica of the EFD (before transformation).

12   Infrastructure Domain

This domain encompasses the underlying services and systems that form the computing environments in which the other domains are deployed and operate. It interfaces with the other domains but has no direct users.

12.1   Service Descriptions

12.1.1   Data Backbone

The Data Backbone is a key component that provides for data storage, transport, and replication, allowing data products to move between computational environments. This service provides policy-based replication of files (in the Science Image Archive) and databases (in the Science Catalog Archive) across multiple physical locations, including the Base, Commission-

ing Cluster, NCSA, and DACs. It manages caches of files at each endpoint as well as persistence to long-term archival storage (e.g. tape). It provides a registration mechanism for new datasets and database entries and a retrieval mechanism compatible with the Data Butler.

The relationships between the Data Backbone components are illustrated in Figure 8.
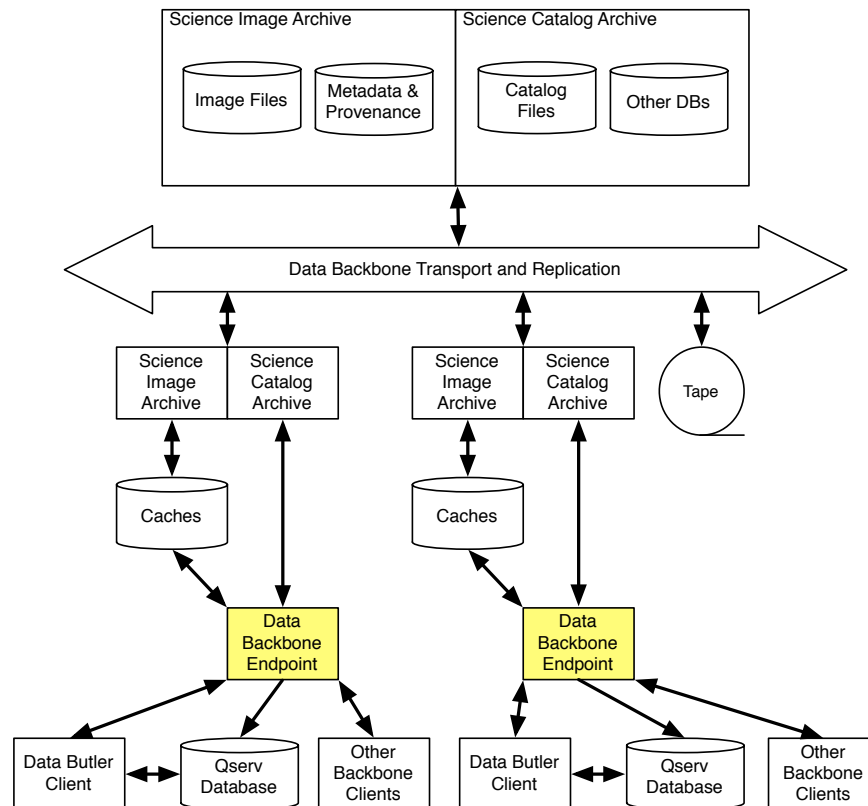
FIGURE 8: Data Backbone

12.1.2   Management/Monitoring

These services provide management and monitoring at service and infrastructure systems levels for each computing environment.

12.1.3   Provisioning/Deployment

These services provide compute, local-to-node storage, and local-to-LAN storage resources for all processing, including Prompt Processing, Batch Processing, and the Science Platforms. They allow allocation of compute and storage resources as well as reproducible, controlled

deployment of services onto those resources.

Some compute resources are reserved for particular uses, but others can be flexibly provisioned, up to a certain maximum quota, if needed to deal with surges in processing.

The priority order for processing is:

- Prompt processing

- Offline processing

- OCS-controlled batch processing

- LSP Commissioning Cluster processing

- LSP Science Validation processing

- LSP Data Access Center processing

- Template and Calibration Products Production

- Data Release Production

The Base Center's services are not highly dynamic or flexible, as they primarily provide interfacing to the OCS and Camera Data System. The baseline provisioning for them is using vSphere; they will be deployed using Puppet.

12.1.4 Workload/Workflow

These services provide management of the execution of science payloads ranging from a single pipeline to a series of "campaigns", each consisting of multiple pipelines. They are able to handle massively distributed computing, executing jobs when their inputs become available and tracking their status and outputs. They ensure that the data needed for a job is accessible to it and that outputs (including log files, if any) are preserved. They can allocate work across multiple computing environments, in particular between NCSA and the Satellite Computing Facility at CC-IN2P3.

12.1.5   Batch Processing

This service provides execution of batch jobs with a variety of priorities from a variety of users in a variety of environments (e.g. OS and software configurations) on the underlying provisioned compute resources. It will use containerization to handle heterogeneity of environments. HTCondor is the baseline technology choice for this service.

12.1.6   Identity Management

This service provides authentication and authorization for all users of any DMS component, especially the LSST Science Platform instances.

12.2   Interfaces

The infrastructure services generally interact with all other deployed services. In the Base Center, there is no workload/workflow management or batch processing; provisioning and deployment are also simplified.

Identity management instances are present in the Base Center and at NCSA. (Another replica will be maintained at the Summit.) These are used to support authentication and authorization for the other physically co-located environments: the Commissioning Cluster and the two Data Access Centers.

13   Software Components

13.1   Science Payloads

These payloads are described in more detail in the DM Applications Design Document (LDM-151). Payloads are built from application software components.

13.1.1   Alert Production Payload

Executes under control of the Prompt Processing service. Generates all Level 1 science data products including Alerts (with the exception of Solar System object orbits) and loads them

---

into the Data Backbone and Level 1 Database. Transmits Alerts to Alert Distribution service. Generates image quality feedback to the OCS and observers via the Telemetry Gateway. Uses crosstalk-corrected science images and associated metadata delivered by the Prompt Processing service; uses Master Calibration Images, Template Images, Level 1 Database, and Calibration Database information from the Data Backbone.

13.1.2　MOPS Payload

Executes under control of the Offline Processing service after a night's observations are complete. Generates entries in the MOPS Database and the Level 1 Database, including Solar System Object records, measurements, and orbits. Performs precovery forced photometry of transients. Uses Level 1 Database entries and images from the Data Backbone.

13.1.3　Raw Calibration Validation Payload

Executes under control of the Prompt Processing service. Generates raw calibration image quality feedback to the OCS and observers via the Telemetry Gateway. Uses crosstalk-corrected science images and associated metadata delivered by the Prompt Processing service, Master Calibration Images, and Calibration Database information from the Data Backbone.

13.1.4　Daily Calibration Products Update Payload

Executes under control of the OCS-controlled batch processing service so that its execution can be synchronized with the observing schedule. Uses raw calibration images and information from the Transformed EFD to generate a subset of Master Calibration Images and Calibration Database entries in the Data Backbone.

13.1.5　Periodic Calibration Products Production Payload

Executes under control of the Template and CPP Execution service at nominally monthly intervals but perhaps as frequently as weekly or as infrequently as quarterly, depending on the stability of Observatory systems and their calibrations. Uses raw calibration images and information from the Transformed EFD to generate a subset of Master Calibration Images and Calibration Database entries in the Data Backbone.

The contents of this document are subject to configuration control by the LSST DM Technical Control Team.

32

13.1.6   Template Generation Payload

Executes under control of the Template and CPP Execution service if necessary to generate templates for Alert Production in between annual Data Release Productions. Uses raw science images to generate the templates, placing them in the Data Backbone.

13.1.7   Annual Calibration Products Production Payload

Executes under control of the Template and CPP Execution service at annual intervals prior to the start of the Data Release Production. Uses raw calibration images, information from the Transformed EFD, information from the Auxiliary Telescope Spectrograph, and external catalogs to generate Master Calibration Images and Calibration Database entries in the Data Backbone.

13.1.8   Data Release Production Payload

Executes under control of the DRP Execution service at annual intervals, first running a "mini-DRP" over a small portion of the sky, followed by the full DRP over the entire sky. Produces science data products in the Data Backbone.

13.2   SUIT

The Science User Interface and Tools provide visualization, plotting, catalog rendering, browsing, and searching elements that can be assembled into predetermined "portals" but can also be used flexibly within dynamic "notebook" environments.

13.3   Middleware

13.3.1   Data Butler Access Client

The Data Butler provides an access abstraction for all science payloads that enables their underlying data sources and destinations to be configured at runtime with a variety of back-ends ranging from local disk to network locations and a variety of serializations ranging from YAML and FITS files (extensible to HDF5 or ASDF) to database tables. The Butler client is also

available within the LSST Science Platform JupyterLab environment.

13.3.2   Parallel Distributed Database (Qserv)

Underlying the catalog data access web service is a parallel distributed database required to handle the petabyte-scale, tens-of-trillions-of-rows catalogs produced by LSST.

13.3.3   Task Framework

The Task Framework is a Python class library that provides a structure (standardized class entry points and conventions) to organize low-level algorithms into potentially-reusable algorithmic components (Tasks; e.g. dark frame subtraction, object detection, object measurement), and to organize tasks into basic pipelines (SuperTasks; e.g., process a single visit, build a coadd, difference a visit). The algorithmic code is written into (Super)Tasks by overriding classes and providing implementation for standard entry points. The Task Framework allows the pipelines to be constructed and run at the level of a single node or a group of tightly-synchronized nodes. It allows for sub-node parallelization: trivial parallelization of Task execution, as well as providing (in the future) parallelization primitives for development of multi-core Tasks and synchronized multi-node Tasks.

The Task Framework serves as an interface layer between orchestration and the algorithmic code. It exposes a standard interface to "activators" (command-line runners as well as the orchestration layer and QA systems), which use it to execute the code wrapped in tasks. The Task Framework does not concern itself with fault-tolerant massively parallel execution of the pipelines over multiple (thousands) of nodes nor any staging of data that might be required; this is the concern of the orchestration middleware.

The Task Framework exposes to the orchestration system needs and capabilities of the underlying algorithmic code (i.e., the number of cores needed, expected memory-per-core, expected need for data). It may also receive from the orchestration layer the information on how to optimally run the particular task (i.e., which level of intra-node parallelization is be desired).

It also includes a configuration API and a logging API.

14 Design and Implementation Standards

Standards have been adopted by the DM Change Control Board (CCB) that apply to all component designs within the LSST DM System. Coding standards and the like that are not pertinent to design may be found in the LSST DM Developer Guide ([1]).

14.1 Python Version

All LSST DM code will run under Python 3 by 2018-01-31. Package authors may choose to have their code backward-compatible with Python 2.7, but this is only required for Science Pipelines code that is a dependency of `lsst_apps`, and then only until external dependencies become Python 3-only.

Python 2.7 will be deprecated by the time LSST begins operations, and many external dependencies are in the process of moving to being Python 3-only.

14.2 HTTPS Protocol

In the absence of a specific technical justification and acceptance by the LSST Information Security Officer and DM Change Control Board, all Web-enabled user interfaces and Web services exposed to users and the public Internet will use the HTTPS protocol and not the HTTP protocol. To reiterate: this is only a default, and exceptions can be made when justified.

The covered interfaces include those of the three LSP Aspects (Portal, JupyterLab, and Web APIs).

The requirement to implement data access policies limiting data access to identified rights holders will require all, or nearly all, data access to be authenticated provides a strong technical justification. In addition, it appears to be appropriate "technical best practice" in the current Internet environment, in the absence of good reasons to do otherwise.

15   Appendix: Traceability

15.1   Requirement to Component Traceability

Note that DMS-REQ-0006 has no components; this requirement is in the process of being dropped.

Requirement	Components
DMS-REQ-0059 Bad Pixel Map	OCS Driven Batch, Calibration Products Production Execution, Annual Calibration, Daily Calibration Update, Periodic Calibration
DMS-REQ-0060 Bias Residual Image	OCS Driven Batch, Calibration Products Production Execution, Annual Calibration, Daily Calibration Update, Periodic Calibration
DMS-REQ-0061 Crosstalk Correction Matrix	OCS Driven Batch, Calibration Products Production Execution, Annual Calibration, Daily Calibration Update, Periodic Calibration
DMS-REQ-0062 Illumination Correction Frame	OCS Driven Batch, Calibration Products Production Execution, Annual Calibration, Daily Calibration Update, Periodic Calibration
DMS-REQ-0063 Monochromatic Flatfield Data Cube	OCS Driven Batch, Calibration Products Production Execution, Annual Calibration, Daily Calibration Update, Periodic Calibration
DMS-REQ-0130 Calibration Data Products	OCS Driven Batch, Calibration Products Production Execution, Annual Calibration, Daily Calibration Update, Periodic Calibration
DMS-REQ-0132 Calibration Image Provenance	OCS Driven Batch, Calibration Products Production Execution, Annual Calibration, Daily Calibration Update, Periodic Calibration
DMS-REQ-0282 Dark Current Correction Frame	OCS Driven Batch, Calibration Products Production Execution, Annual Calibration, Daily Calibration Update, Periodic Calibration
DMS-REQ-0283 Fringe Correction Frame	OCS Driven Batch, Calibration Products Production Execution, Annual Calibration, Daily Calibration Update, Periodic Calibration
DMS-REQ-0018 Raw Science Image Data Acquisition	Image and EFD Archiving

Requirement	Components
DMS-REQ-0020 Wavefront Sensor Data Acquisition	Image and EFD Archiving
DMS-REQ-0022 Crosstalk Corrected Science Image Data Acquisition	Prompt Processing
DMS-REQ-0024 Raw Image Assembly	Image and EFD Archiving
DMS-REQ-0068 Raw Science Image Metadata	Image and EFD Archiving
DMS-REQ-0265 Guider Calibration Data Acquisition	Image and EFD Archiving, OCS Driven Batch, Calibration Products Production Execution, Annual Calibration, Daily Calibration Update, Periodic Calibration
DMS-REQ-0326 Storing Approximations of Per-pixel Metadata	Data Release Production
DMS-REQ-0331 Computing Derived Quantities	LSP Web APIs, Data Release Production
DMS-REQ-0332 Denormalizing Database Tables	LSP Web APIs
DMS-REQ-0333 Maximum Likelihood Values and Covariances	Data Access Client, Alert Production, Data Release Production
DMS-REQ-0346 Data Availability	Image and EFD Archiving, OCS Driven Batch, Offline Processing, Prompt Processing, Calibration Products Production Execution, Data Release Production Execution
DMS-REQ-0347 Measurements in catalogs	Data Access Client, LSP Web APIs, Alert Production, Data Release Production
DMS-REQ-0004 Nightly Data Accessible Within 24 hrs	Alert Distribution, Alert Filtering, Offline Processing, Prompt Processing, Alert Production, MOPS and Forced Photometry
DMS-REQ-0010 Difference Exposures	Alert Production
DMS-REQ-0074 Difference Exposure Attributes	Alert Production
DMS-REQ-0069 Processed Visit Images	Alert Production
DMS-REQ-0029 Generate Photometric Zeropoint for Visit Image	Alert Production
DMS-REQ-0030 Generate WCS for Visit Images	Alert Production
DMS-REQ-0070 Generate PSF for Visit Images	Alert Production
DMS-REQ-0072 Processed Visit Image Content	Alert Production
DMS-REQ-0327 Background Model Calculation	Alert Production
DMS-REQ-0328 Documenting Image Characterization	Alert Production

Requirement	Components
DMS-REQ-0097 Level 1 Data Quality Report Definition	QC System, Alert Production
DMS-REQ-0099 Level 1 Performance Report Definition	Prompt Processing, QC System
DMS-REQ-0101 Level 1 Calibration Report Definition	OCS Driven Batch, Daily Calibration Update, Raw Calibration
DMS-REQ-0266 Exposure Catalog	Alert Production
DMS-REQ-0269 DIASource Catalog	Alert Production
DMS-REQ-0270 Faint DIASource Measurements	Alert Production
DMS-REQ-0271 DIAObject Catalog	Alert Production
DMS-REQ-0272 DIAObject Attributes	Alert Production
DMS-REQ-0273 SSObject Catalog	MOPS and Forced Photometry
DMS-REQ-0274 Alert Content	Alert Production
DMS-REQ-0317 DIAForcedSource Catalog	Alert Production, MOPS and Forced Photometry
DMS-REQ-0319 Characterizing Variability	Alert Production, MOPS and Forced Photometry
DMS-REQ-0323 Calculating SSObject Parameters	LSP Web APIs
DMS-REQ-0324 Matching DIASources to Objects	LSP Web APIs, Alert Production
DMS-REQ-0325 Regenerating L1 Data Products During Data Release Processing	Data Release Production
DMS-REQ-0034 Associate Sources to Objects	Data Release Production
DMS-REQ-0047 Provide PSF for Coadded Images	Data Release Production
DMS-REQ-0103 Produce Images for EPO	Data Release Production
DMS-REQ-0106 Coadded Image Provenance	Data Release Production
DMS-REQ-0267 Source Catalog	Data Release Production
DMS-REQ-0268 Forced-Source Catalog	Data Release Production
DMS-REQ-0275 Object Catalog	Data Release Production
DMS-REQ-0046 Provide Photometric Redshifts of Galaxies	Data Release Production
DMS-REQ-0276 Object Characterization	Data Release Production
DMS-REQ-0277 Coadd Source Catalog	Data Release Production
DMS-REQ-0349 Detecting extended low surface brightness objects	Data Release Production
DMS-REQ-0278 Coadd Image Method Constraints	Data Release Production

Requirement	Components
DMS-REQ-0279 Deep Detection Coadds	Data Release Production
DMS-REQ-0280 Template Coadds	Data Release Production, Template Generation
DMS-REQ-0281 Multi-band Coadds	Data Release Production
DMS-REQ-0329 All-Sky Visualization of Data Releases	Data Release Production
DMS-REQ-0330 Best Seeing Coadds	Data Release Production
DMS-REQ-0334 Persisting Data Products	Data Backbone, LSP Web APIs
DMS-REQ-0335 PSF-Matched Coadds	Data Release Production
DMS-REQ-0336 Regenerating Data Products from Previous Data Releases	Data Backbone, LSP Web APIs
DMS-REQ-0337 Detecting faint variable objects	Data Release Production
DMS-REQ-0338 Targeted Coadds	Data Backbone, LSP Web APIs
DMS-REQ-0339 Tracking Characterization Changes Between Data Releases	Data Backbone, LSP Web APIs
DMS-REQ-0320 Processing of Data From Special Programs	Task Execution Framework, Data Release Production
DMS-REQ-0321 Level 1 Processing of Special Programs Data	Offline Processing, Prompt Processing, Alert Production, MOPS and Forced Photometry
DMS-REQ-0322 Special Programs Database	Data Backbone, LSP Web APIs
DMS-REQ-0344 Constraints on Level 1 Special Program Products Generation	Offline Processing, Prompt Processing, Data Backbone, LSP Web APIs, Alert Production, MOPS and Forced Photometry
DMS-REQ-0185 Archive Center	Bulk Distribution, Data Backbone, ITC Environments
DMS-REQ-0186 Archive Center Disaster Recovery	Bulk Distribution, Data Backbone, ITC Environments
DMS-REQ-0187 Archive Center Co-Location with Existing Facility	ITC Environments
DMS-REQ-0188 Archive to Data Access Center Network	Networks
DMS-REQ-0189 Archive to Data Access Center Network Availability	Networks
DMS-REQ-0190 Archive to Data Access Center Network Reliability	Networks
DMS-REQ-0191 Archive to Data Access Center Network Secondary Link	Networks
DMS-REQ-0176 Base Facility Infrastructure	Data Backbone, ITC Environments
DMS-REQ-0177 Base Facility Temporary Storage	Data Backbone

Requirement	Components
DMS-REQ-0178 Base Facility Co-Location with Existing Facility	ITC Environments
DMS-REQ-0316 Commissioning Cluster	ITC Environments
DMS-REQ-0180 Base to Archive Network	Networks
DMS-REQ-0181 Base to Archive Network Availability	Networks
DMS-REQ-0182 Base to Archive Network Reliability	Networks
DMS-REQ-0183 Base to Archive Network Secondary Link	Networks
DMS-REQ-0008 Pipeline Availability	Alert Distribution, Alert Filtering, Image and EFD Archiving, OCS Driven Batch, Offline Processing, Prompt Processing, Telemetry Gateway, Calibration Products Production Execution, Data Release Production Execution, Data Backbone, ITC Environments, Networks
DMS-REQ-0161 Optimization of Cost, Reliability and Availability in Order	Alert Distribution, Alert Filtering, Data Backbone, ITC Environments, Networks, LSP Web APIs, LSP JupyterLab, LSP Portal
DMS-REQ-0162 Pipeline Throughput	Alert Distribution, Alert Filtering, Image and EFD Archiving, OCS Driven Batch, Offline Processing, Prompt Processing, Data Backbone, ITC Environments, Networks
DMS-REQ-0163 Re-processing Capacity	Calibration Products Production Execution, Data Release Production Execution, Data Backbone, ITC Environments, Networks
DMS-REQ-0164 Temporary Storage for Communications Links	Data Backbone
DMS-REQ-0165 Infrastructure Sizing for "catching up"	Image and EFD Archiving, Offline Processing, Data Backbone, ITC Environments, Networks
DMS-REQ-0166 Incorporate Fault-Tolerance	Data Backbone
DMS-REQ-0167 Incorporate Autonomics	Alert Distribution, Alert Filtering, Image and EFD Archiving, OCS Driven Batch, Offline Processing, Prompt Processing, Calibration Products Production Execution, Data Release Production Execution, Data Backbone, ITC Environments, Networks
DMS-REQ-0314 Compute Platform Heterogeneity	Alert Distribution, Alert Filtering, Image and EFD Archiving, OCS Driven Batch, Offline Processing, Prompt Processing, Telemetry Gateway, Calibration Products Production Execution, Data Release Production Execution, Bulk Distribution, Proposal Manager, Developer Services, Integration and Test, Data Backbone, Identity Manager, ITC Environments, Networks, Data Access Client, Task Execution Framework, LSP Web APIs, LSP JupyterLab, LSP Portal, QC System

Requirement	Components
DMS-REQ-0318 Data Management Unscheduled Downtime	Alert Distribution, Alert Filtering, Image and EFD Archiving, OCS Driven Batch, Offline Processing, Prompt Processing, Telemetry Gateway, Calibration Products Production Execution, Data Release Production Execution, Bulk Distribution, Proposal Manager, Data Backbone, Identity Manager, ITC Environments, Networks, LSP Web APIs, LSP JupyterLab, LSP Portal, QC System
DMS-REQ-0122 Access to catalogs for external Level 3 processing	Bulk Distribution, Data Backbone
DMS-REQ-0123 Access to input catalogs for DAC-based Level 3 processing	Data Backbone, LSP Web APIs, LSP JupyterLab, LSP Portal
DMS-REQ-0124 Federation with external catalogs	Data Backbone, LSP Web APIs, LSP JupyterLab, LSP Portal
DMS-REQ-0126 Access to images for external Level 3 processing	Bulk Distribution, Data Backbone
DMS-REQ-0127 Access to input images for DAC-based Level 3 processing	Data Backbone, LSP Web APIs, LSP JupyterLab, LSP Portal
DMS-REQ-0193 Data Access Centers	Bulk Distribution, Data Backbone, ITC Environments, Networks, LSP Web APIs, LSP JupyterLab, LSP Portal
DMS-REQ-0194 Data Access Center Simultaneous Connections	Networks, LSP JupyterLab, LSP Portal
DMS-REQ-0196 Data Access Center Geographical Distribution	ITC Environments
DMS-REQ-0197 No Limit on Data Access Centers	Data Backbone, Identity Manager, ITC Environments, Networks, LSP Web APIs, LSP JupyterLab, LSP Portal
DMS-REQ-0075 Catalog Queries	LSP Web APIs
DMS-REQ-0077 Maintain Archive Publicly Accessible	Data Backbone
DMS-REQ-0078 Catalog Export Formats	LSP Web APIs
DMS-REQ-0094 Keep Historical Alert Archive	Data Backbone
DMS-REQ-0102 Provide Engineering and Facility Database Archive	Image and EFD Archiving, Data Backbone
DMS-REQ-0309 Raw Data Archiving Reliability	Image and EFD Archiving, Data Backbone
DMS-REQ-0310 Un-Archived Data Product Cache	Data Backbone
DMS-REQ-0311 Regenerate Un-archived Data Products	Data Backbone, LSP Web APIs
DMS-REQ-0312 Level 1 Data Product Access	Prompt Processing, Data Backbone, LSP Web APIs, Alert Production

Requirement	Components
DMS-REQ-0313 Level 1 and 2 Catalog Access	Data Backbone, LSP Web APIs
DMS-REQ-0341 Providing a Precovery Service	Offline Processing, LSP Portal, MOPS and Forced Photometry
DMS-REQ-0345 Logging of catalog queries	LSP Web APIs
DMS-REQ-0168 Summit Facility Data Communications	Networks
DMS-REQ-0170 Prefer Computing and Storage Down	ITC Environments
DMS-REQ-0315 DMS Communication with OCS	Image and EFD Archiving, OCS Driven Batch, Prompt Processing, Telemetry Gateway
DMS-REQ-0171 Summit to Base Network	Networks
DMS-REQ-0172 Summit to Base Network Availability	Networks
DMS-REQ-0173 Summit to Base Network Reliability	Networks
DMS-REQ-0174 Summit to Base Network Secondary Link	Networks
DMS-REQ-0175 Summit to Base Network Ownership and Operation	Networks
DMS-REQ-0002 Transient Alert Distribution	Alert Distribution, Alert Filtering, Prompt Processing, Alert Production
DMS-REQ-0089 Solar System Objects Available Within Specified Time	Offline Processing, Data Backbone, LSP Web APIs, MOPS and Forced Photometry
DMS-REQ-0096 Generate Data Quality Report Within Specified Time	QC System
DMS-REQ-0098 Generate DMS Performance Report Within Specified Time	QC System
DMS-REQ-0100 Generate Calibration Report Within Specified Time	QC System, Daily Calibration Update
DMS-REQ-0131 Calibration Images Available Within Specified Time	OCS Driven Batch, Prompt Processing, Daily Calibration Update, Raw Calibration
DMS-REQ-0284 Level-1 Production Completeness	Image and EFD Archiving, Offline Processing, Prompt Processing
DMS-REQ-0285 Level 1 Source Association	Alert Production
DMS-REQ-0286 SSObject Precovery	MOPS and Forced Photometry
DMS-REQ-0287 DIASource Precovery	Data Backbone, LSP Web APIs, MOPS and Forced Photometry
DMS-REQ-0288 Use of External Orbit Catalogs	Alert Production, MOPS and Forced Photometry
DMS-REQ-0342 Alert Filtering Service	Alert Filtering
DMS-REQ-0343 Performance Requirements for LSST Alert Filtering Service	Alert Distribution, Alert Filtering

Requirement	Components
DMS-REQ-0348 Pre-defined alert filters	Alert Filtering
DMS-REQ-0289 Calibration Production Processing	OCS Driven Batch, Calibration Products Production Execution, Annual Calibration, Daily Calibration Update, Periodic Calibration
DMS-REQ-0006 Timely Publication of Level 2 Data Releases	
DMS-REQ-0350 Associating Objects across data releases	Data Backbone, Data Release Production
DMS-REQ-0291 Query Repeatibility	Data Backbone, LSP Web APIs
DMS-REQ-0292 Uniqueness of IDs Across Data Releases	Data Backbone, LSP Web APIs
DMS-REQ-0293 Selection of Datasets	Data Backbone, LSP Web APIs
DMS-REQ-0294 Processing of Datasets	OCS Driven Batch, Offline Processing, Prompt Processing, Calibration Products Production Execution, Data Release Production Execution, Data Backbone
DMS-REQ-0295 Transparent Data Access	LSP Web APIs
DMS-REQ-0119 DAC resource allocation for Level 3 processing	Bulk Distribution, Proposal Manager, LSP Web APIs, LSP JupyterLab, LSP Portal
DMS-REQ-0120 Level 3 Data Product Self Consistency	Data Backbone, Workload+Workflow, LSP Web APIs
DMS-REQ-0121 Provenance for Level 3 processing at DACs	Data Backbone, Workload+Workflow, Data Access Client, Task Execution Framework, LSP Web APIs
DMS-REQ-0125 Software framework for Level 3 catalog processing	Data Backbone, Workload+Workflow, Data Access Client, Task Execution Framework, LSP Web APIs
DMS-REQ-0128 Software framework for Level 3 image processing	Data Backbone, Workload+Workflow, Data Access Client, Task Execution Framework, LSP Web APIs
DMS-REQ-0290 Level 3 Data Import	Data Backbone, LSP Web APIs
DMS-REQ-0340 Access Controls of Level 3 Data Products	Data Backbone, Identity Manager
DMS-REQ-0009 Simulated Data	Alert Production, Annual Calibration, Daily Calibration Update, Data Release Production, MOPS and Forced Photometry, Periodic Calibration, Raw Calibration, Template Generation
DMS-REQ-0032 Image Differencing	Science Algorithms, Alert Production, Data Release Production
DMS-REQ-0033 Provide Source Detection Software	Science Algorithms, Alert Production, Data Release Production
DMS-REQ-0042 Provide Astrometric Model	Science Algorithms, Alert Production, Data Release Production, Science Primitives
DMS-REQ-0043 Provide Calibrated Photometry	Science Algorithms, Alert Production, Data Release Production, Science Primitives

Requirement	Components
DMS-REQ-0052 Enable a Range of Shape Measurement Approaches	Science Algorithms, Alert Production, Data Release Production, Science Primitives
DMS-REQ-0160 Provide User Interface Services	LSP Portal
DMS-REQ-0296 Pre-cursor, and Real Data	Data Backbone, Data Access Client, Science Algorithms
DMS-REQ-0351 Provide Beam Projector Coordinate Calculation Software	Science Primitives
DMS-REQ-0308 Software Architecture to Enable Community Re-Use	Science Algorithms, Alert Production, Annual Calibration, Daily Calibration Update, Data Release Production, MOPS and Forced Photometry, Periodic Calibration, Raw Calibration, Template Generation
DMS-REQ-0065 Provide Image Access Services	LSP Web APIs
DMS-REQ-0155 Provide Data Access Services	LSP Web APIs
DMS-REQ-0298 Data Product and Raw Data Access	Data Backbone, LSP Web APIs
DMS-REQ-0299 Data Product Ingest	Data Backbone, LSP Web APIs
DMS-REQ-0300 Bulk Download Service	Bulk Distribution, Data Backbone
DMS-REQ-0156 Provide Pipeline Execution Services	Workload+Workflow
DMS-REQ-0302 Production Orchestration	Workload+Workflow
DMS-REQ-0303 Production Monitoring	Workload+Workflow
DMS-REQ-0304 Production Fault Tolerance	Workload+Workflow, Task Execution Framework
DMS-REQ-0158 Provide Pipeline Construction Services	Task Execution Framework
DMS-REQ-0305 Task Specification	Task Execution Framework
DMS-REQ-0306 Task Configuration	Task Execution Framework
DMS-REQ-0297 DMS Initialization Component	ITC Environments
DMS-REQ-0301 Control of Level-1 Production	Prompt Processing
DMS-REQ-0307 Unique Processing Coverage	Workload+Workflow

15.2   Component to Requirement Traceability

Note that only "leaf" components are traced to requirements.

Component	Requirements
Data Management	

Component	Requirements
Level 1 System	
Alert Distribution	• DMS-REQ-0004 Nightly Data Accessible Within 24 hrs • DMS-REQ-0008 Pipeline Availability • DMS-REQ-0161 Optimization of Cost, Reliability and Availability in Order • DMS-REQ-0162 Pipeline Throughput • DMS-REQ-0167 Incorporate Autonomics • DMS-REQ-0314 Compute Platform Heterogeneity • DMS-REQ-0318 Data Management Unscheduled Downtime • DMS-REQ-0002 Transient Alert Distribution • DMS-REQ-0343 Performance Requirements for LSST Alert Filtering Service
Alert Filtering	• DMS-REQ-0004 Nightly Data Accessible Within 24 hrs • DMS-REQ-0008 Pipeline Availability • DMS-REQ-0161 Optimization of Cost, Reliability and Availability in Order • DMS-REQ-0162 Pipeline Throughput • DMS-REQ-0167 Incorporate Autonomics • DMS-REQ-0314 Compute Platform Heterogeneity • DMS-REQ-0318 Data Management Unscheduled Downtime • DMS-REQ-0002 Transient Alert Distribution • DMS-REQ-0342 Alert Filtering Service • DMS-REQ-0343 Performance Requirements for LSST Alert Filtering Service • DMS-REQ-0348 Pre-defined alert filters
Image and EFD Archiving	• DMS-REQ-0018 Raw Science Image Data Acquisition • DMS-REQ-0020 Wavefront Sensor Data Acquisition • DMS-REQ-0024 Raw Image Assembly • DMS-REQ-0068 Raw Science Image Metadata • DMS-REQ-0265 Guider Calibration Data Acquisition • DMS-REQ-0346 Data Availability • DMS-REQ-0008 Pipeline Availability • DMS-REQ-0162 Pipeline Throughput • DMS-REQ-0165 Infrastructure Sizing for "catching up" • DMS-REQ-0167 Incorporate Autonomics • DMS-REQ-0314 Compute Platform Heterogeneity • DMS-REQ-0318 Data Management Unscheduled Downtime • DMS-REQ-0102 Provide Engineering and Facility Database Archive • DMS-REQ-0309 Raw Data Archiving Reliability • DMS-REQ-0315 DMS Communication with OCS • DMS-REQ-0284 Level-1 Production Completeness

Component	Requirements
OCS Driven Batch	• DMS-REQ-0059 Bad Pixel Map
	• DMS-REQ-0060 Bias Residual Image
	• DMS-REQ-0061 Crosstalk Correction Matrix
	• DMS-REQ-0062 Illumination Correction Frame
	• DMS-REQ-0063 Monochromatic Flatfield Data Cube
	• DMS-REQ-0130 Calibration Data Products
	• DMS-REQ-0132 Calibration Image Provenance
	• DMS-REQ-0282 Dark Current Correction Frame
	• DMS-REQ-0283 Fringe Correction Frame
	• DMS-REQ-0265 Guider Calibration Data Acquisition
	• DMS-REQ-0346 Data Availability
	• DMS-REQ-0101 Level 1 Calibration Report Definition
	• DMS-REQ-0008 Pipeline Availability
	• DMS-REQ-0162 Pipeline Throughput
	• DMS-REQ-0167 Incorporate Autonomics
	• DMS-REQ-0314 Compute Platform Heterogeneity
	• DMS-REQ-0318 Data Management Unscheduled Downtime
	• DMS-REQ-0315 DMS Communication with OCS
	• DMS-REQ-0131 Calibration Images Available Within Specified Time
	• DMS-REQ-0289 Calibration Production Processing
	• DMS-REQ-0294 Processing of Datasets

Component	Requirements
Offline Processing	• DMS-REQ-0346 Data Availability • DMS-REQ-0004 Nightly Data Accessible Within 24 hrs • DMS-REQ-0321 Level 1 Processing of Special Programs Data • DMS-REQ-0344 Constraints on Level 1 Special Program Products Generation • DMS-REQ-0008 Pipeline Availability • DMS-REQ-0162 Pipeline Throughput • DMS-REQ-0165 Infrastructure Sizing for "catching up" • DMS-REQ-0167 Incorporate Autonomics • DMS-REQ-0314 Compute Platform Heterogeneity • DMS-REQ-0318 Data Management Unscheduled Downtime • DMS-REQ-0341 Providing a Precovery Service • DMS-REQ-0089 Solar System Objects Available Within Specified Time • DMS-REQ-0284 Level-1 Production Completeness • DMS-REQ-0294 Processing of Datasets
Prompt Processing	• DMS-REQ-0022 Crosstalk Corrected Science Image Data Acquisition • DMS-REQ-0346 Data Availability • DMS-REQ-0004 Nightly Data Accessible Within 24 hrs • DMS-REQ-0099 Level 1 Performance Report Definition • DMS-REQ-0321 Level 1 Processing of Special Programs Data • DMS-REQ-0344 Constraints on Level 1 Special Program Products Generation • DMS-REQ-0008 Pipeline Availability • DMS-REQ-0162 Pipeline Throughput • DMS-REQ-0167 Incorporate Autonomics • DMS-REQ-0314 Compute Platform Heterogeneity • DMS-REQ-0318 Data Management Unscheduled Downtime • DMS-REQ-0312 Level 1 Data Product Access • DMS-REQ-0315 DMS Communication with OCS • DMS-REQ-0002 Transient Alert Distribution • DMS-REQ-0131 Calibration Images Available Within Specified Time • DMS-REQ-0284 Level-1 Production Completeness • DMS-REQ-0294 Processing of Datasets • DMS-REQ-0301 Control of Level-1 Production

Component	Requirements
Telemetry Gateway	• DMS-REQ-0008 Pipeline Availability • DMS-REQ-0314 Compute Platform Heterogeneity • DMS-REQ-0318 Data Management Unscheduled Downtime • DMS-REQ-0315 DMS Communication with OCS
Level 2 System	
Calibration Products Production Execution	• DMS-REQ-0059 Bad Pixel Map • DMS-REQ-0060 Bias Residual Image • DMS-REQ-0061 Crosstalk Correction Matrix • DMS-REQ-0062 Illumination Correction Frame • DMS-REQ-0063 Monochromatic Flatfield Data Cube • DMS-REQ-0130 Calibration Data Products • DMS-REQ-0132 Calibration Image Provenance • DMS-REQ-0282 Dark Current Correction Frame • DMS-REQ-0283 Fringe Correction Frame • DMS-REQ-0265 Guider Calibration Data Acquisition • DMS-REQ-0346 Data Availability • DMS-REQ-0008 Pipeline Availability • DMS-REQ-0163 Re-processing Capacity • DMS-REQ-0167 Incorporate Autonomics • DMS-REQ-0314 Compute Platform Heterogeneity • DMS-REQ-0318 Data Management Unscheduled Downtime • DMS-REQ-0289 Calibration Production Processing • DMS-REQ-0294 Processing of Datasets

The contents of this document are subject to configuration control by the LSST DM Technical Control Team.

48

Component	Requirements
Data Release Production Execution	• DMS-REQ-0346 Data Availability • DMS-REQ-0008 Pipeline Availability • DMS-REQ-0163 Re-processing Capacity • DMS-REQ-0167 Incorporate Autonomics • DMS-REQ-0314 Compute Platform Heterogeneity • DMS-REQ-0318 Data Management Unscheduled Downtime • DMS-REQ-0294 Processing of Datasets
Data Access Center	
Bulk Distribution	• DMS-REQ-0185 Archive Center • DMS-REQ-0186 Archive Center Disaster Recovery • DMS-REQ-0314 Compute Platform Heterogeneity • DMS-REQ-0318 Data Management Unscheduled Downtime • DMS-REQ-0122 Access to catalogs for external Level 3 processing • DMS-REQ-0126 Access to images for external Level 3 processing • DMS-REQ-0193 Data Access Centers • DMS-REQ-0119 DAC resource allocation for Level 3 processing • DMS-REQ-0300 Bulk Download Service
Proposal Manager	• DMS-REQ-0314 Compute Platform Heterogeneity • DMS-REQ-0318 Data Management Unscheduled Downtime • DMS-REQ-0119 DAC resource allocation for Level 3 processing
Analysis and Developer Services	
Developer Services	• DMS-REQ-0314 Compute Platform Heterogeneity
Integration and Test	• DMS-REQ-0314 Compute Platform Heterogeneity
Infrastructure	
Data Backbone	• DMS-REQ-0334 Persisting Data Products • DMS-REQ-0336 Regenerating Data Products from Previous Data Releases • DMS-REQ-0338 Targeted Coadds • DMS-REQ-0339 Tracking Characterization Changes Between Data Releases • DMS-REQ-0322 Special Programs Database • DMS-REQ-0344 Constraints on Level 1 Special Program Products Generation • DMS-REQ-0185 Archive Center • DMS-REQ-0186 Archive Center Disaster Recovery • DMS-REQ-0176 Base Facility Infrastructure • DMS-REQ-0177 Base Facility Temporary Storage

The contents of this document are subject to configuration control by the LSST DM Technical Control Team.

Component	Requirements
Data Backbone (continued)	• DMS-REQ-0008 Pipeline Availability
	• DMS-REQ-0161 Optimization of Cost, Reliability and Availability in Order
	• DMS-REQ-0162 Pipeline Throughput
	• DMS-REQ-0163 Re-processing Capacity
	• DMS-REQ-0164 Temporary Storage for Communications Links
	• DMS-REQ-0165 Infrastructure Sizing for "catching up"
	• DMS-REQ-0166 Incorporate Fault-Tolerance
	• DMS-REQ-0167 Incorporate Autonomics
	• DMS-REQ-0314 Compute Platform Heterogeneity
	• DMS-REQ-0318 Data Management Unscheduled Downtime
	• DMS-REQ-0122 Access to catalogs for external Level 3 processing
	• DMS-REQ-0123 Access to input catalogs for DAC-based Level 3 processing
	• DMS-REQ-0124 Federation with external catalogs
	• DMS-REQ-0126 Access to images for external Level 3 processing
	• DMS-REQ-0127 Access to input images for DAC-based Level 3 processing
	• DMS-REQ-0193 Data Access Centers
	• DMS-REQ-0197 No Limit on Data Access Centers
	• DMS-REQ-0077 Maintain Archive Publicly Accessible
	• DMS-REQ-0094 Keep Historical Alert Archive
	• DMS-REQ-0102 Provide Engineering and Facility Database Archive
	• DMS-REQ-0309 Raw Data Archiving Reliability
	• DMS-REQ-0310 Un-Archived Data Product Cache
	• DMS-REQ-0311 Regenerate Un-archived Data Products
	• DMS-REQ-0312 Level 1 Data Product Access
	• DMS-REQ-0313 Level 1 and 2 Catalog Access
	• DMS-REQ-0089 Solar System Objects Available Within Specified Time
	• DMS-REQ-0287 DIASource Precovery
	• DMS-REQ-0350 Associating Objects across data releases
	• DMS-REQ-0291 Query Repeatibility
	• DMS-REQ-0292 Uniqueness of IDs Across Data Releases
	• DMS-REQ-0293 Selection of Datasets
	• DMS-REQ-0294 Processing of Datasets
	• DMS-REQ-0120 Level 3 Data Product Self Consistency
	• DMS-REQ-0121 Provenance for Level 3 processing at DACs
	• DMS-REQ-0125 Software framework for Level 3 catalog processing
	• DMS-REQ-0128 Software framework for Level 3 image processing
	• DMS-REQ-0290 Level 3 Data Import
	• DMS-REQ-0340 Access Controls of Level 3 Data Products
	• DMS-REQ-0296 Pre-cursor, and Real Data
	• DMS-REQ-0298 Data Product and Raw Data Access

The contents of this document are subject to configuration control by the LSST DM Technical Control Team.

Component	Requirements
Data Backbone (continued)	• DMS-REQ-0299 Data Product Ingest • DMS-REQ-0300 Bulk Download Service
Identity Manager	• DMS-REQ-0314 Compute Platform Heterogeneity • DMS-REQ-0318 Data Management Unscheduled Downtime • DMS-REQ-0197 No Limit on Data Access Centers • DMS-REQ-0340 Access Controls of Level 3 Data Products
ITC Environments	• DMS-REQ-0185 Archive Center • DMS-REQ-0186 Archive Center Disaster Recovery • DMS-REQ-0187 Archive Center Co-Location with Existing Facility • DMS-REQ-0176 Base Facility Infrastructure • DMS-REQ-0178 Base Facility Co-Location with Existing Facility • DMS-REQ-0316 Commissioning Cluster • DMS-REQ-0008 Pipeline Availability • DMS-REQ-0161 Optimization of Cost, Reliability and Availability in Order • DMS-REQ-0162 Pipeline Throughput • DMS-REQ-0163 Re-processing Capacity • DMS-REQ-0165 Infrastructure Sizing for "catching up" • DMS-REQ-0167 Incorporate Autonomics • DMS-REQ-0314 Compute Platform Heterogeneity • DMS-REQ-0318 Data Management Unscheduled Downtime • DMS-REQ-0193 Data Access Centers • DMS-REQ-0196 Data Access Center Geographical Distribution • DMS-REQ-0197 No Limit on Data Access Centers • DMS-REQ-0170 Prefer Computing and Storage Down • DMS-REQ-0297 DMS Initialization Component

Component	Requirements
Networks	• DMS-REQ-0188 Archive to Data Access Center Network
	• DMS-REQ-0189 Archive to Data Access Center Network Availability
	• DMS-REQ-0190 Archive to Data Access Center Network Reliability
	• DMS-REQ-0191 Archive to Data Access Center Network Secondary Link
	• DMS-REQ-0180 Base to Archive Network
	• DMS-REQ-0181 Base to Archive Network Availability
	• DMS-REQ-0182 Base to Archive Network Reliability
	• DMS-REQ-0183 Base to Archive Network Secondary Link
	• DMS-REQ-0008 Pipeline Availability
	• DMS-REQ-0161 Optimization of Cost, Reliability and Availability in Order
	• DMS-REQ-0162 Pipeline Throughput
	• DMS-REQ-0163 Re-processing Capacity
	• DMS-REQ-0165 Infrastructure Sizing for "catching up"
	• DMS-REQ-0167 Incorporate Autonomics
	• DMS-REQ-0314 Compute Platform Heterogeneity
	• DMS-REQ-0318 Data Management Unscheduled Downtime
	• DMS-REQ-0193 Data Access Centers
	• DMS-REQ-0194 Data Access Center Simultaneous Connections
	• DMS-REQ-0197 No Limit on Data Access Centers
	• DMS-REQ-0168 Summit Facility Data Communications
	• DMS-REQ-0171 Summit to Base Network
	• DMS-REQ-0172 Summit to Base Network Availability
	• DMS-REQ-0173 Summit to Base Network Reliability
	• DMS-REQ-0174 Summit to Base Network Secondary Link
	• DMS-REQ-0175 Summit to Base Network Ownership and Operation

Component	Requirements
Workload+Workflow	• DMS-REQ-0120 Level 3 Data Product Self Consistency • DMS-REQ-0121 Provenance for Level 3 processing at DACs • DMS-REQ-0125 Software framework for Level 3 catalog processing • DMS-REQ-0128 Software framework for Level 3 image processing • DMS-REQ-0156 Provide Pipeline Execution Services • DMS-REQ-0302 Production Orchestration • DMS-REQ-0303 Production Monitoring • DMS-REQ-0304 Production Fault Tolerance
Data Access Client	• DMS-REQ-0333 Maximum Likelihood Values and Covariances • DMS-REQ-0347 Measurements in catalogs • DMS-REQ-0314 Compute Platform Heterogeneity • DMS-REQ-0121 Provenance for Level 3 processing at DACs • DMS-REQ-0125 Software framework for Level 3 catalog processing • DMS-REQ-0128 Software framework for Level 3 image processing • DMS-REQ-0296 Pre-cursor, and Real Data
Task Execution Framework	• DMS-REQ-0320 Processing of Data From Special Programs • DMS-REQ-0314 Compute Platform Heterogeneity • DMS-REQ-0121 Provenance for Level 3 processing at DACs • DMS-REQ-0125 Software framework for Level 3 catalog processing • DMS-REQ-0128 Software framework for Level 3 image processing • DMS-REQ-0304 Production Fault Tolerance • DMS-REQ-0158 Provide Pipeline Construction Services • DMS-REQ-0305 Task Specification • DMS-REQ-0306 Task Configuration
Science Platform	
LSP Web APIs	• DMS-REQ-0331 Computing Derived Quantities • DMS-REQ-0332 Denormalizing Database Tables • DMS-REQ-0347 Measurements in catalogs • DMS-REQ-0323 Calculating SSObject Parameters • DMS-REQ-0324 Matching DIASources to Objects • DMS-REQ-0334 Persisting Data Products • DMS-REQ-0336 Regenerating Data Products from Previous Data Releases • DMS-REQ-0338 Targeted Coadds • DMS-REQ-0339 Tracking Characterization Changes Between Data Releases • DMS-REQ-0322 Special Programs Database • DMS-REQ-0344 Constraints on Level 1 Special Program Products Generation • DMS-REQ-0161 Optimization of Cost, Reliability and Availability in Order

Component	Requirements
LSP Web APIs (continued)	• DMS-REQ-0314 Compute Platform Heterogeneity
	• DMS-REQ-0318 Data Management Unscheduled Downtime
	• DMS-REQ-0123 Access to input catalogs for DAC-based Level 3 processing
	• DMS-REQ-0124 Federation with external catalogs
	• DMS-REQ-0127 Access to input images for DAC-based Level 3 processing
	• DMS-REQ-0193 Data Access Centers
	• DMS-REQ-0197 No Limit on Data Access Centers
	• DMS-REQ-0075 Catalog Queries
	• DMS-REQ-0078 Catalog Export Formats
	• DMS-REQ-0311 Regenerate Un-archived Data Products
	• DMS-REQ-0312 Level 1 Data Product Access
	• DMS-REQ-0313 Level 1 and 2 Catalog Access
	• DMS-REQ-0345 Logging of catalog queries
	• DMS-REQ-0089 Solar System Objects Available Within Specified Time
	• DMS-REQ-0287 DIASource Precovery
	• DMS-REQ-0291 Query Repeatibility
	• DMS-REQ-0292 Uniqueness of IDs Across Data Releases
	• DMS-REQ-0293 Selection of Datasets
	• DMS-REQ-0295 Transparent Data Access
	• DMS-REQ-0119 DAC resource allocation for Level 3 processing
	• DMS-REQ-0120 Level 3 Data Product Self Consistency
	• DMS-REQ-0121 Provenance for Level 3 processing at DACs
	• DMS-REQ-0125 Software framework for Level 3 catalog processing
	• DMS-REQ-0128 Software framework for Level 3 image processing
	• DMS-REQ-0290 Level 3 Data Import
	• DMS-REQ-0065 Provide Image Access Services
	• DMS-REQ-0155 Provide Data Access Services
	• DMS-REQ-0298 Data Product and Raw Data Access
	• DMS-REQ-0299 Data Product Ingest

Component	Requirements
LSP JupyterLab	• DMS-REQ-0161 Optimization of Cost, Reliability and Availability in Order • DMS-REQ-0314 Compute Platform Heterogeneity • DMS-REQ-0318 Data Management Unscheduled Downtime • DMS-REQ-0123 Access to input catalogs for DAC-based Level 3 processing • DMS-REQ-0124 Federation with external catalogs • DMS-REQ-0127 Access to input images for DAC-based Level 3 processing • DMS-REQ-0193 Data Access Centers • DMS-REQ-0194 Data Access Center Simultaneous Connections • DMS-REQ-0197 No Limit on Data Access Centers • DMS-REQ-0119 DAC resource allocation for Level 3 processing
LSP Portal	• DMS-REQ-0161 Optimization of Cost, Reliability and Availability in Order • DMS-REQ-0314 Compute Platform Heterogeneity • DMS-REQ-0318 Data Management Unscheduled Downtime • DMS-REQ-0123 Access to input catalogs for DAC-based Level 3 processing • DMS-REQ-0124 Federation with external catalogs • DMS-REQ-0127 Access to input images for DAC-based Level 3 processing • DMS-REQ-0193 Data Access Centers • DMS-REQ-0194 Data Access Center Simultaneous Connections • DMS-REQ-0197 No Limit on Data Access Centers • DMS-REQ-0341 Providing a Precovery Service • DMS-REQ-0119 DAC resource allocation for Level 3 processing • DMS-REQ-0160 Provide User Interface Services

Component	Requirements
QC System	• DMS-REQ-0097 Level 1 Data Quality Report Definition • DMS-REQ-0099 Level 1 Performance Report Definition • DMS-REQ-0314 Compute Platform Heterogeneity • DMS-REQ-0318 Data Management Unscheduled Downtime • DMS-REQ-0096 Generate Data Quality Report Within Specified Time • DMS-REQ-0098 Generate DMS Performance Report Within Specified Time • DMS-REQ-0100 Generate Calibration Report Within Specified Time
Science Algorithms	• DMS-REQ-0032 Image Differencing • DMS-REQ-0033 Provide Source Detection Software • DMS-REQ-0042 Provide Astrometric Model • DMS-REQ-0043 Provide Calibrated Photometry • DMS-REQ-0052 Enable a Range of Shape Measurement Approaches • DMS-REQ-0296 Pre-cursor, and Real Data • DMS-REQ-0308 Software Architecture to Enable Community Re-Use
Science Payloads	
Alert Production	• DMS-REQ-0333 Maximum Likelihood Values and Covariances • DMS-REQ-0347 Measurements in catalogs • DMS-REQ-0004 Nightly Data Accessible Within 24 hrs • DMS-REQ-0010 Difference Exposures • DMS-REQ-0074 Difference Exposure Attributes • DMS-REQ-0069 Processed Visit Images • DMS-REQ-0029 Generate Photometric Zeropoint for Visit Image • DMS-REQ-0030 Generate WCS for Visit Images • DMS-REQ-0070 Generate PSF for Visit Images • DMS-REQ-0072 Processed Visit Image Content • DMS-REQ-0327 Background Model Calculation • DMS-REQ-0328 Documenting Image Characterization • DMS-REQ-0097 Level 1 Data Quality Report Definition • DMS-REQ-0266 Exposure Catalog • DMS-REQ-0269 DIASource Catalog • DMS-REQ-0270 Faint DIASource Measurements • DMS-REQ-0271 DIAObject Catalog • DMS-REQ-0272 DIAObject Attributes • DMS-REQ-0274 Alert Content • DMS-REQ-0317 DIAForcedSource Catalog • DMS-REQ-0319 Characterizing Variability • DMS-REQ-0324 Matching DIASources to Objects • DMS-REQ-0321 Level 1 Processing of Special Programs Data

Component	Requirements
Alert Production (continued)	• DMS-REQ-0344 Constraints on Level 1 Special Program Products Generation • DMS-REQ-0312 Level 1 Data Product Access • DMS-REQ-0002 Transient Alert Distribution • DMS-REQ-0285 Level 1 Source Association • DMS-REQ-0288 Use of External Orbit Catalogs • DMS-REQ-0009 Simulated Data • DMS-REQ-0032 Image Differencing • DMS-REQ-0033 Provide Source Detection Software • DMS-REQ-0042 Provide Astrometric Model • DMS-REQ-0043 Provide Calibrated Photometry • DMS-REQ-0052 Enable a Range of Shape Measurement Approaches • DMS-REQ-0308 Software Architecture to Enable Community Re-Use
Annual Calibration	• DMS-REQ-0059 Bad Pixel Map • DMS-REQ-0060 Bias Residual Image • DMS-REQ-0061 Crosstalk Correction Matrix • DMS-REQ-0062 Illumination Correction Frame • DMS-REQ-0063 Monochromatic Flatfield Data Cube • DMS-REQ-0130 Calibration Data Products • DMS-REQ-0132 Calibration Image Provenance • DMS-REQ-0282 Dark Current Correction Frame • DMS-REQ-0283 Fringe Correction Frame • DMS-REQ-0265 Guider Calibration Data Acquisition • DMS-REQ-0289 Calibration Production Processing • DMS-REQ-0009 Simulated Data • DMS-REQ-0308 Software Architecture to Enable Community Re-Use

Component	Requirements
Daily Calibration Update	• DMS-REQ-0059 Bad Pixel Map • DMS-REQ-0060 Bias Residual Image • DMS-REQ-0061 Crosstalk Correction Matrix • DMS-REQ-0062 Illumination Correction Frame • DMS-REQ-0063 Monochromatic Flatfield Data Cube • DMS-REQ-0130 Calibration Data Products • DMS-REQ-0132 Calibration Image Provenance • DMS-REQ-0282 Dark Current Correction Frame • DMS-REQ-0283 Fringe Correction Frame • DMS-REQ-0265 Guider Calibration Data Acquisition • DMS-REQ-0101 Level 1 Calibration Report Definition • DMS-REQ-0100 Generate Calibration Report Within Specified Time • DMS-REQ-0131 Calibration Images Available Within Specified Time • DMS-REQ-0289 Calibration Production Processing • DMS-REQ-0009 Simulated Data • DMS-REQ-0308 Software Architecture to Enable Community Re-Use
Data Release Production	• DMS-REQ-0326 Storing Approximations of Per-pixel Metadata • DMS-REQ-0331 Computing Derived Quantities • DMS-REQ-0333 Maximum Likelihood Values and Covariances • DMS-REQ-0347 Measurements in catalogs • DMS-REQ-0325 Regenerating L1 Data Products During Data Release Processing • DMS-REQ-0034 Associate Sources to Objects • DMS-REQ-0047 Provide PSF for Coadded Images • DMS-REQ-0103 Produce Images for EPO • DMS-REQ-0106 Coadded Image Provenance • DMS-REQ-0267 Source Catalog • DMS-REQ-0268 Forced-Source Catalog • DMS-REQ-0275 Object Catalog • DMS-REQ-0046 Provide Photometric Redshifts of Galaxies • DMS-REQ-0276 Object Characterization • DMS-REQ-0277 Coadd Source Catalog • DMS-REQ-0349 Detecting extended low surface brightness objects • DMS-REQ-0278 Coadd Image Method Constraints • DMS-REQ-0279 Deep Detection Coadds • DMS-REQ-0280 Template Coadds • DMS-REQ-0281 Multi-band Coadds • DMS-REQ-0329 All-Sky Visualization of Data Releases • DMS-REQ-0330 Best Seeing Coadds • DMS-REQ-0335 PSF-Matched Coadds • DMS-REQ-0337 Detecting faint variable objects

Component	Requirements
Data Release Production (continued)	• DMS-REQ-0320 Processing of Data From Special Programs • DMS-REQ-0350 Associating Objects across data releases • DMS-REQ-0009 Simulated Data • DMS-REQ-0032 Image Differencing • DMS-REQ-0033 Provide Source Detection Software • DMS-REQ-0042 Provide Astrometric Model • DMS-REQ-0043 Provide Calibrated Photometry • DMS-REQ-0052 Enable a Range of Shape Measurement Approaches • DMS-REQ-0308 Software Architecture to Enable Community Re-Use
MOPS and Forced Photometry	• DMS-REQ-0004 Nightly Data Accessible Within 24 hrs • DMS-REQ-0273 SSObject Catalog • DMS-REQ-0317 DIAForcedSource Catalog • DMS-REQ-0319 Characterizing Variability • DMS-REQ-0321 Level 1 Processing of Special Programs Data • DMS-REQ-0344 Constraints on Level 1 Special Program Products Generation • DMS-REQ-0341 Providing a Precovery Service • DMS-REQ-0089 Solar System Objects Available Within Specified Time • DMS-REQ-0286 SSObject Precovery • DMS-REQ-0287 DIASource Precovery • DMS-REQ-0288 Use of External Orbit Catalogs • DMS-REQ-0009 Simulated Data • DMS-REQ-0308 Software Architecture to Enable Community Re-Use

Component	Requirements
Periodic Calibration	• DMS-REQ-0059 Bad Pixel Map • DMS-REQ-0060 Bias Residual Image • DMS-REQ-0061 Crosstalk Correction Matrix • DMS-REQ-0062 Illumination Correction Frame • DMS-REQ-0063 Monochromatic Flatfield Data Cube • DMS-REQ-0130 Calibration Data Products • DMS-REQ-0132 Calibration Image Provenance • DMS-REQ-0282 Dark Current Correction Frame • DMS-REQ-0283 Fringe Correction Frame • DMS-REQ-0265 Guider Calibration Data Acquisition • DMS-REQ-0289 Calibration Production Processing • DMS-REQ-0009 Simulated Data • DMS-REQ-0308 Software Architecture to Enable Community Re-Use
Raw Calibration	• DMS-REQ-0101 Level 1 Calibration Report Definition • DMS-REQ-0131 Calibration Images Available Within Specified Time • DMS-REQ-0009 Simulated Data • DMS-REQ-0308 Software Architecture to Enable Community Re-Use
Template Generation	• DMS-REQ-0280 Template Coadds • DMS-REQ-0009 Simulated Data • DMS-REQ-0308 Software Architecture to Enable Community Re-Use
Science Primitives	• DMS-REQ-0042 Provide Astrometric Model • DMS-REQ-0043 Provide Calibrated Photometry • DMS-REQ-0052 Enable a Range of Shape Measurement Approaches • DMS-REQ-0351 Provide Beam Projector Coordinate Calculation Software

16   References

[1] LSST DM Developer Guide, URL `https://developer.lsst.io/`

[2] **[LDM-139]**, Becla, J., Lim, K.T., 2013, *Data Management Storage Sizing and I/O Model Explanation*, LDM-139, URL `https://ls.st/LDM-139`

[3] **[LDM-141]**, Becla, J., Lim, K.T., 2013, *Data Management Storage Sizing and I/O Model*, LDM-141, URL `https://ls.st/LDM-141`

The contents of this document are subject to configuration control by the LSST DM Technical Control Team.

[4] **[LDM-135]**, Becla, J., Wang, D., Monkewitz, S., et al., 2013, *Database Design*, LDM-135, URL `https://ls.st/LDM-135`

[5] **[LSE-29]**, Claver, C.F., The LSST Systems Engineering Integrated Project Team, 2016, *LSST System Requirements*, LSE-29, URL `https://ls.st/LSE-29`

[6] **[LSE-30]**, Claver, C.F., The LSST Systems Engineering Integrated Project Team, 2016, *LSST System Requirements*, LSE-30, URL `https://ls.st/LSE-30`

[7] **[LSE-72]**, Dubious-Felsmann, G., Schumacher, G., Selvy, B., 2014, *OCS Command Dictionary for Data Management*, LSE-72, URL `https://ls.st/LSE-72`

[8] **[LSE-81]**, Dubois-Felsmann, G., 2013, *LSST Science and Project Sizing Inputs*, LSE-81, URL `https://ls.st/LSE-81`

[9] **[LSE-61]**, Dubois-Felsmann, G., 2016, *LSST Data Management Subsystem Requirements*, LSE-61, URL `https://ls.st/LSE-61`

[10] **[LSE-82]**, Dubois-Felsmann, G., Lim, K.T., 2013, *Science and Project Sizing Inputs Explanation*, LSE-82, URL `https://ls.st/LSE-82`

[11] **[LDM-542]**, Dubois-Felsmann, G., Lim, K.T., Wu, X., et al., 2017, *LSST Science Platform Design*, LDM-542, URL `https://ls.st/LDM-542`

[12] **[LDM-129]**, Freemon, M., Kantor, J., 2013, *Data Management Infrastructure Design*, LDM-129, URL `https://ls.st/LDM-129`

[13] **[LDM-143]**, Freemon, M., Pietrowicz, S., 2013, *Site Specific Infrastructure Estimation Explanation*, LDM-143, URL `https://ls.st/LDM-143`

[14] **[LDM-144]**, Freemon, M., Pietrowicz, S., Alt, J., 2016, *Site Specific Infrastructure Estimation Model*, LDM-144, URL `https://ls.st/LDM-144`

[15] **[LPM-17]**, Ivezić, Ž., The LSST Science Collaboration, 2011, *LSST Science Requirements Document*, LPM-17, URL `https://ls.st/LPM-17`

[16] **[LSE-163]**, Jurić, M., et al., 2017, *LSST Data Products Definition Document*, LSE-163, URL `https://ls.st/LSE-163`

[17] **[LDM-142]**, Kantor, J., 2017, *Network Sizing Model*, LDM-142, URL `https://ls.st/LDM-142`

[18] **[LDM-138]**, Kantor, J., Axelrod, T., Lim, K.T., 2013, *Data Management Compute Sizing Model*, LDM-138, URL `https://ls.st/LDM-138`

[19] **[LSE-78]**, Lambert, R., Kantor, J., Huffer, M., et al., 2017, *LSST Observatory Network Design*, LSE-78, URL `https://ls.st/LSE-78`

[20] **[LDM-230]**, Lim, K.T., 2013, *Automated Operation of the LSST Data Management System*, LDM-230, URL `https://ls.st/LDM-230`

[21] **[LDM-140]**, Lim, K.T., Smith, C., Axelrod, T., Dubois-Felsmann, G., Freemon, M., 2013, *Data Management Compute Sizing Explanation*, LDM-140, URL `https://ls.st/LDM-140`

The contents of this document are subject to configuration control by the LSST DM Technical Control Team.

[22] **[LDM-148]**, Lim, K.T., Bosch, J., Dubois-Felsmann, G., et al., 2017, *Data Management System Design*, LDM-148, URL `https://ls.st/LDM-148`

[23] **[LDM-152]**, Lim, K.T., Dubois-Felsmann, G., Johnson, M., Jurić, M., Petravick, D., 2017, *Data Management Middleware Design*, LDM-152, URL `https://ls.st/LDM-152`

[24] **[Document-5373]**, Pinto, P., Kantor, J., Strauss, M., Sweeney, D., 2008, *Data Access White Paper*, Document-5373, URL `https://ls.st/Document-5373`

[25] **[LDM-151]**, Swinbank, J.D., et al., 2017, *Data Management Science Pipelines Design*, LDM-151, URL `https://ls.st/LDM-151`