

Rubin Observatory

Vera C. Rubin Observatory
Systems Engineering


Data Management LSST Science Platform Requirements

Gregory Dubois-Felsmann, David Ciardi, Fritz Mueller,
Frossie Economou

LDM-554

Latest Revision: 2019-04-01

Draft Revision NOT YET Approved - This Rubin Observatory document has been approved as a Content-Controlled Document by the Rubin Observatory DM Change Control Board. If this document is changed or superseded, the new document will retain the Handle designation shown above. The control is on the most recent digital document with this Handle in the Rubin Observatory digital archive and not printed versions. Additional information may be found in the corresponding DM RFC. - **Draft Revision NOT YET Approved**



Change Record

Version	Date	Description	Owner name
	2017-07-03	Initial draft.	G. Dubois-Felsmann
	2018-05-24	Add notebook and API requirements	G. Dubois-Felsmann, T. Jenness
1.0	2018-06-18	Approved in RFC-485.	J. Swinbank
1.1	2019-04-01	Approved in RFC-588. Issued for LSP review, April 2019	G. Dubois-Felsmann, L. P. Guy

Document source location: MagicDraw SysML

Version from source repository: 223

Contents

1 Science Platform	1
1.1 Overview	1
1.1.1 Access to All Released or Authorized Data Products	1
1.1.2 Portal Aspect	1
1.1.3 Notebook Aspect	1
1.1.4 API (Data Access) Aspect	2
1.1.5 Linkage of Aspects	2
1.1.6 Use of VO Standards	2
1.1.7 Abide by the Data Access Policies	3
1.2 Discover Data	3
1.2.1 Semantic Linkage	3
1.2.2 Semantic Linkage: Uncertainties	4
1.2.3 Transfer of Portal Data References to Notebook	4
1.3 User Storage	4
1.3.1 User File Workspace	4
1.3.2 User Database Workspace	5
1.3.3 User Workspace Access Controls	5
1.4 Transfer Data	5
1.4.1 Download Data	5
1.4.2 Upload Data	6
1.4.3 Transfer Data to Workspace	6
1.4.4 Tabular Data Download File Formats	6
1.4.5 Image Data Download File Format	7
1.5 Performance	7
1.5.1 Peak Volume for Moderate-Sized Queries	7
1.5.2 Peak Volume for Queries on all Objects	7
1.5.3 Peak Volume of In-process Queries	8
1.5.4 Query Result Download Bandwidth	8

1.6	Control and Management	9
1.6.1	Alert Subscription	9
1.7	Documentation	9
1.7.1	Documentation	9
1.8	General	9
1.8.1	Identity and Security	9
1.8.2	Implementation	12
2	Portal Aspect	13
2.1	Overview	13
2.1.1	Portal is a Web Application	13
2.2	Discover Data	13
2.2.1	Portal Discovery of all Data Products	13
2.2.2	Portal Access to Workspace	14
2.2.3	Semantic Linkage: Portal Workflows	14
2.2.4	Access to Calibration Products	15
2.2.5	Coadded Image to Single-Epoch Image Associations	15
2.2.6	Access to External Archives	15
2.2.7	API for Access to Portal Session State	16
2.3	Query Data	16
2.3.1	Support Synchronous and Asynchronous Queries	16
2.3.2	Long Query Backgrounding	16
2.3.3	Query Status and Termination Notification	17
2.3.4	Query Results Size Limitation	17
2.3.5	Query History Inspection	18
2.3.6	Query Saving - Portal	18
2.3.7	General Tabular Queries	18
2.3.8	Catalog Queries	25
2.3.9	Image Metadata Queries	27
2.3.10	Image Cutout Queries	28

2.4	Display Data	29
2.4.1	Visualization of Tabular and Image Data	29
2.4.2	Visualization of Ancillary Information	30
2.4.3	Linking Visualization of Image Data to Tabular Data	30
2.4.4	Visualization of Uploaded Tabular and Image Data	30
2.4.5	Visualization of Workspace Data	31
2.4.6	Display Tabular Data	31
2.4.7	Display Image Data	36
2.4.8	All-Sky Visualization	41
2.4.9	Coordinate Display Tools	43
2.5	Exploration and Analysis	45
2.5.1	Table Manipulation	46
2.5.2	Image Manipulation	48
2.5.3	Special Workflows	48
2.6	User Storage	52
2.6.1	Access to User Databases	52
2.7	Transfer Data	52
2.7.1	Tabular Data Download	52
2.7.2	Image Data Download	53
2.7.3	Selected Image Download	53
2.7.4	Download Volume Estimation	53
2.7.5	Long Download Completion Notification	53
2.8	Performance	54
2.9	Control and Management	54
2.9.1	APIs for Visualization Components	54
2.9.2	Storage Quotas User Interface	54
2.9.3	Computational Quotas User Interface	55
2.9.4	Portal Display Preferences	55
2.9.5	Alert Subscription	55

2.10	Documentation	56
2.10.1	Access to Observatory Documentation	57
2.10.2	Portal User Documentation	57
2.10.3	Portal API Documentation	57
2.11	General	57
2.11.1	Identity and Security	57
2.11.2	Implementation	58
3	Notebook Aspect	58
3.1	Environment	58
3.1.1	Interactive Python Environment	58
3.1.2	Unix Shell Access	59
3.1.3	Pre-installed Containerized Software Releases	59
3.1.4	Release Deployment Latency	59
3.1.5	Data Access Middleware Availability	59
3.1.6	Common Astronomy Package Availability	60
3.1.7	User Package Installation	60
3.1.8	User Development Environment	60
3.1.9	Persistent User Home File Space	60
3.1.10	Documentation	61
3.1.11	New-User Onboarding	61
3.2	Resources	61
3.2.1	Shared File Space	61
3.2.2	Access to the API and Portal Aspects	62
3.2.3	User File Workspace Access	62
3.2.4	VOSpace Access	62
3.2.5	User Database Workspace Access	62
3.2.6	Batch System Access	63
3.2.7	Compute and Storage Quotas	63
3.2.8	Access to All Data Products	63

3.3	Architecture	63
3.3.1	Ease of Deployment	63
3.3.2	Deployment Workload in Kubernetes	64
3.3.3	System Health Monitoring	64
3.4	Display Data	64
3.4.1	Image Visualization	64
3.4.2	Scientific Plotting	65
3.4.3	Visualization Linkage	65
3.4.4	Visualization Interactivity	65
3.4.5	Visualization Scaling	66
3.5	Interaction with Portal	66
3.5.1	Access to Portal-Initiated Queries	66
3.5.2	Access to Portal Visualization API	66
3.5.3	Notebook-Launching Interface	66
3.6	General	67
3.6.1	Identity and Security	67
3.6.2	Implementation	68
4	API (Data Access) Aspect	68
4.1	Discover Data	68
4.1.1	Use of CAOM2	68
4.1.2	Access to Image and Visit Metadata	68
4.1.3	Access to Catalog Data Products	68
4.1.4	Access to Observatory Metadata	69
4.1.5	Enforcement of Information Classification	69
4.1.6	Access to Reference Catalogs	69
4.1.7	Access to Virtual Data Products	69
4.1.8	Access to Image Data in FITS Format	70
4.1.9	Multiple Data Releases	70
4.1.10	Catalog Metadata Service	70

4.2	Query Data	70
4.2.1	TAP Service - Tabular Queries	70
4.2.2	SIA Service - Image Availability	72
4.2.3	SODA Service - Image Data Retrieval	72
4.2.4	Query Persistence	73
4.2.5	Python Interface	74
4.3	User Storage	74
4.3.1	VOSpace Service	74
4.3.2	WebDAV Service	75
4.4	Transfer Data	75
4.4.1	Tabular Data Output	75
4.4.2	Workspace Interaction	77
4.5	Performance	78
4.6	Control and Management	78
4.7	Documentation	78
4.8	General	78
4.8.1	Identity and Security	78
4.8.2	Implementation	79
5	Proposed Requirements	80
5.1	Large Result Set Support	80
5.2	Moderate Result Set Support	80

Data Management LSST Science Platform Requirements

These requirements set out the major components of the LSST Science Platform (LSP).

1 Science Platform

1.1 Overview

1.1.1 Access to All Released or Authorized Data Products

ID: DMS-LSP-REQ-0001

Specification: The LSP shall provide the capability to access all the Project's released data products, including, but not limited to, the data products enumerated in the DPDD (LSE-163), as well as all user data products to which a user has access.

1.1.2 Portal Aspect

ID: DMS-LSP-REQ-0002

Specification: The LSP shall provide a Web-based "Portal" means of access to all the LSST data products, and to user storage resources.

Discussion: The Portal is defined by further requirements.

1.1.3 Notebook Aspect

ID: DMS-LSP-REQ-0003

Specification: The LSP shall provide an interactive Python computing environment, accessible through a Web browser, with access to all the LSST data products and to user computing and storage resources.

Discussion: The term "Notebook aspect" (and, *a fortiori*, "Jupyter[Lab] aspect", which has also been used in DM discussions) refers to the reference implementation of this aspect of the LSP in terms of Jupyter notebooks and the forthcoming JupyterLab successor to that technology, with an IPython kernel back end. For the purposes of this requirements document, however, it seems preferable to avoid using an external brand name in the normative text.

1.1.4 API (Data Access) Aspect

ID: DMS-LSP-REQ-0004

Specification: The LSP shall provide a Web API for access to all the LSST data products and the user storage resources.

Discussion: The Web API will use VO standards as much as is practical and supported by community expectations; see DMS-LSP-REQ-0006 below. It may also include additional APIs to provide services unique to LSST.

1.1.5 Linkage of Aspects

ID: DMS-LSP-REQ-0005

Specification: The LSP shall facilitate access to the same LSST and user data through multiple aspects.

Discussion: It should be possible to identify or create data in one aspect and retrieve it in another. This requirement is satisfied principally by building the various aspects over the same underlying data services.

1.1.6 Use of VO Standards

ID: DMS-LSP-REQ-0006

Specification: The LSP shall utilize stable and accepted Virtual Observatory standards for publically offered APIs wherever feasible.

Discussion: The intent is to establish a "VO First" posture, and require explicit exception be requested for any APIs needing to deviate from this requirement. This requirement applies most importantly to the API Aspect's externally-facing data access APIs, but also in other areas. E.g., from this requirement we may derive that the Portal must support access to externally-source reference catalog via VO-compliant queries, or that a Portal service for pushing data to user applications should support SAMP.

1.1.7 Abide by the Data Access Policies

ID: DMS-LSP-REQ-0007

Specification: The LSP shall comply with the public data access policy and access restrictions defined by the LSST Project and operations organization.

Discussion: This includes both the enforcement of Project-level data rights (e.g., to the Level 1 and Level 2 data) and user-provided access controls to user-created data (e.g., Level 3 data products).

1.2 Discover Data

1.2.1 Semantic Linkage

ID: DMS-LSP-REQ-0008

Specification: The LSP shall support the identification of linkages between data items that reflect their provenance and data dependencies.

Discussion: For instance, from a calibrated image it should be possible to identify the raw image from which it was generated, and the calibration data used in its processing; from a catalog entry it should be possible to identify the image(s) on which the measurement was made. In practice this has DM system-level implications. For the LSP it relies on upstream components recording the appropriate provenance metadata. At the Data Access Services level it mainly just requires exposing the tables and columns containing this metadata, and identifying them as such (e.g., by reporting foreign-key relationships) in the Reflection APIs. At the Notebook level this is largely satisfied by providing Butler-level access to this information.

1.2.2 Semantic Linkage: Uncertainties

ID: DMS-LSP-REQ-0009

Specification: The LSP shall support the identification of relationships between data items (notably database columns) that represent a quantity and its uncertainty(ies).

Discussion: This should extend to relationships between groups of quantities and other quantities that represent their covariances, but this is not an open-ended requirement to support all such cases. These relationships should be expressed as appropriate to the LSP aspect, and require "upstream" support at the point of definition or creation of the data.

1.2.3 Transfer of Portal Data References to Notebook

ID: DMS-LSP-REQ-0010

Specification: The LSP shall facilitate the transfer to the Notebook aspect of references allowing retrieval in a notebook of the data explored in the Portal session.

Discussion: This allows a user to locate and preview data in the Portal environment and then readily transfer their work to the Notebook aspect for detailed analysis.

1.3 User Storage

1.3.1 User File Workspace

ID: DMS-LSP-REQ-0011

Specification: The LSP shall provide a "User File Workspace": resources for the storage of file oriented User Generated data, which shall be accessible from all three aspects.

Discussion: All users with data rights will have a "birthright" quota of User File Workspace storage. A proposal-based process will be available for requests for additional storage. The Workspace will not, in general, place constraints on the format of the data stored. It may be used for image data as well as for tabular data in file-oriented storage forms. User databases

are a separate mechanism. The Workspace is expected to be made available through an implementation of the IVOA VOSpace standard.

1.3.2 User Database Workspace

ID: DMS-LSP-REQ-0012

Specification: The LSP shall provide for the creation, use, and management of user databases (User Generated tabular data products), and shall enable interaction with user databases with the same facilities as for Project-created database to the extent feasible.

Discussion: Some database-related capabilities of the LSP rely on the availability of detailed metadata on the Project-created databases that goes beyond the normal content of a database schema (e.g., IVOA UCDs for table columns). Users will be enabled, but not required, to supply such metadata for their own databases (and they may do so incorrectly), so LSP functionality that depends on it may not be available for user databases.

1.3.3 User Workspace Access Controls

ID: DMS-LSP-REQ-0013

Specification: The LSP shall permit users to place access restrictions on data in the User File and Database Workspaces, based on both user and user-group identities, and shall enforce these restrictions in all its aspects.

1.4 Transfer Data

1.4.1 Download Data

ID: DMS-LSP-REQ-0014

Specification: The LSP shall provide means for downloading data resulting from queries or other operations, or from the Workspace, to the user's system.

Discussion: In general the API aspect can be used to retrieve data to remote sites. Additional

mechanisms will be provided as appropriate to the various aspects, and are described in lower-level requirements. For the DAC instances of the LSP, the "user's system" will generally be on the public Internet; for other instances it may be on project-internal systems.

1.4.2 Upload Data

ID: DMS-LSP-REQ-0015

Specification: The LSP shall provide means for uploading data from the user's system for use in the LSP aspects, including for storage in the Workspace.

1.4.3 Transfer Data to Workspace

ID: DMS-LSP-REQ-0016

Specification: Transfer of data to and from the Workspace shall be usable as an alternative in all features of the LSP where download or upload, respectively, are available.

Discussion: The user can "download" data, e.g., a query result, to the Workspace, and can "upload" data from the Workspace into other functions of the LSP, e.g., to use as a target list in a multi-object search.

1.4.4 Tabular Data Download File Formats

ID: DMS-LSP-REQ-0017

Specification: The LSP shall allow tabular search results, including but not limited to data from the source and object tables and the image metadata tables, to be downloaded or saved to the workspace in at least the following formats: FITS table, VOTable, and ASCII delimiter-separated table (e.g., CSV).

Discussion: A final set of formats needs to be discussed and approved by the Project. This is primarily a requirement on the API aspect and Data Access Services.

1.4.5 Image Data Download File Format

ID: DMS-LSP-REQ-0018

Specification: The LSP shall allow LSST image data products to be downloaded or saved to the workspace as FITS files including the appropriate metadata.

Discussion: This needs some discussion as to what other formats we may want and must be readable by future releases.

1.5 Performance

1.5.1 Peak Volume for Moderate-Sized Queries

ID: DMS-LSP-REQ-0028

Specification: The LSP shall handle at peak usage 50 simultaneous queries without degradation, with the following properties: input selection of up to 1E7 objects in the catalog, result data set of up to 0.1GB, and a response time of 10 seconds.

Discussion: This requirement flows down from several requirements in the DMSR (LSE-61) which constrain both the performance of the database systems (via LDM-555) and the Science Platform. In the Science Platform context, this requirement implies that the LSP will not degrade the performance supported by the underlying database systems. This applies to queries originating from any of the Aspects, and covers all of the stages of the query: query generation, query running, results generation, display of results, and downloading and saving of results.

1.5.2 Peak Volume for Queries on all Objects

ID: DMS-LSP-REQ-0029

Specification: The LSP shall handle at peak usage 20 simultaneous queries without degradation, with the following properties: input selection of up the entire object database, result data set of up to 6 GB, and a response time of one hour.

Discussion: This requirement flows down from several requirements in the DMSR (LSE-61) which constrain both the performance of the database systems (via LDM-555) and the Science Platform. In the Science Platform context, this requirement implies that the LSP will not degrade the performance supported by the underlying database systems. This applies to queries originating from any of the Aspects, and covers all of the stages of the query: query generation, query running, results generation, display of results, and downloading and saving of results.

1.5.3 Peak Volume of In-process Queries

ID: DMS-LSP-REQ-0030

Specification: The LSP shall simultaneously handle at peak usage $20 * 6 \text{ GB} = 120 \text{ GB}$ downloads

Discussion: This requirement flows down from several requirements in the DMSR (LSE-61) which constrain both the performance of the database systems (via LDM-555) and the Science Platform. In the Science Platform context, this requirement implies that the LSP will not degrade the performance supported by the underlying database systems. This applies to queries originating from any of the Aspects, and covers all of the stages of the query: query generation, query running, results generation, display of results, and downloading and saving of results.

1.5.4 Query Result Download Bandwidth

ID: DMS-LSP-REQ-0031

Specification: The LSP shall support a download rate of 6 Gbps for query results including results tables and images.

Discussion: This requirement flows down from several requirements in the DMSR (LSE-61) which constrain both the performance of the database systems (via LDM-555) and the Science Platform. In the Science Platform context, this requirement implies that the LSP will not degrade the performance supported by the underlying database systems. This applies to queries originating from any of the Aspects, and covers all of the stages of the query: query

generation, query running, results generation, display of results, and downloading and saving of results.

1.6 Control and Management

1.6.1 Alert Subscription

1.7 Documentation

1.7.1 Documentation

ID: DMS-LSP-REQ-0019

Specification: The LSP shall provide user and reference documentation for all its aspects.

1.8 General

1.8.1 Identity and Security

1.8.1.1 Authenticated User Access

ID: DMS-LSP-REQ-0020

Specification: The functions and services of the LSP, including all three aspects, shall be available only to authenticated users, except where other requirements or other change-controlled specifications authorize or mandate otherwise.

1.8.1.2 New-user Support

ID: DMS-LSP-REQ-0021

Specification: The Portal and Notebook aspects shall provide guidance to unauthenticated users as to how to establish an identity as usable for authentication to the LSP.

Discussion: This could be as simple as a link to a "register for LSST access" page from the login screens of the Portal and Notebook.

1.8.1.3 Common Identity

ID: DMS-LSP-REQ-0022

Specification: A user shall be able to use the same credentials to authenticate to all aspects of the LSP, and to receive access to any personal data or other state that is available cross-aspects.

Discussion: This does not explicitly mandate "strong single-sign-on" in the sense that someone who has logged in to the Portal aspect can then proceed to the Notebook aspect without a separate login. (TBR: This behavior would be highly desirable and we may wish to adopt this as a requirement after all.) It *does* require "weak single-sign-on" - the same credentials work everywhere.

1.8.1.4 Use of External Identity Providers

ID: DMS-LSP-REQ-0023

Specification: The LSP shall permit users to authenticate to the system using external credentials, from identity providers determined to be trusted by the LSST project or its operations organization.

Discussion: This means that a user should be able to authenticate to an instance of the LSP using, for example, Github credentials, or credentials from a home institution. The policies for how LSST determines that an external user, with external credentials, has data rights and may establish an identity in LSST systems are set forth in other documents.

1.8.1.5 Use of Multiple Sets of Credentials

ID: DMS-LSP-REQ-0024

Specification: The LSP shall permit users to associate multiple sets of credentials, from different providers, with the same identity within the LSP.

1.8.1.6 Acceptable Use Policy

ID: DMS-LSP-REQ-0025

Specification: Non-project-staff users of the LSP shall be required to agree to and abide by an Acceptable Use Policy, to be determined by the LSST project or its operations organization, as a condition of access to any Project instance of the LSP.

Discussion: The policies for project staff are set in other documents.

1.8.1.7 Using secure protocols

ID: DMS-LSP-REQ-0026

Specification: All external connections to the LSP shall be encrypted using protocols and cipher suites compliant with LSST cybersecurity policy.

Discussion: All connections from remote clients to LSST servers (including any web sockets to enable AJAX-like functionality) will be encrypted. It is expected that this will mean general use of "https:" protocols. Connections back out to external archives that do not support secure communications may be left unencrypted.

1.8.1.8 Privacy of User Activities

ID: DMS-LSP-REQ-0027

Specification: The LSP shall ensure that a user's activities on the LSP are not visible to other users without the originating user's explicit authorization.

Discussion: Users will expect, for instance, that the queries they perform are not revealed to

other users without their consent.

1.8.2 Implementation

1.8.2.1 Multiple installations

ID: DMS-LSP-REQ-0032

Specification: The LSP design shall facilitate the installation and maintenance of multiple instances of the LSP and shall support both instances that are accessible from the public Internet and instances that are accessible only within the LSST Project.

Discussion: In addition to the public instances at the Project-provided Data Access Centers, it is expected that additional instances will be used to support internal processes such as commissioning. It is also desirable that the LSP be designed and implemented so that its components can readily be installed outside the Project-provided facilities (e.g., at IN2P3), but (TBR) this is not a formal requirement.

1.8.2.2 Internet-Accessible (IPv4)

ID: DMS-LSP-REQ-0033

Specification: The LSP shall support access from the public Internet using IPv4 protocols.

Discussion: This covers all three aspects of the LSP and thus both browser and other Web-API access.

1.8.2.3 Internet-Accessible (IPv6)

ID: DMS-LSP-REQ-0034

Specification: The LSP shall support access from the public Internet using IPv6 protocols.

Discussion: This covers all three aspects of the LSP and thus both browser and other Web-API

access.

1.8.2.4 System-Availability Indication

ID: DMS-LSP-REQ-0035

Specification: The LSP aspects shall provide means to inform users when their services are unavailable, including for reasons of maintenance or excessive load.

Discussion: This is a DM requirement, not just a Science Platform requirement, as it interacts with lower-level system-availability issues. Even at LSP level this is a complicated problem, as availability of query services through the Portal, for instance, involves time budgeting and resource management for the all of the stages of the query: query generation, query running, results generation, display of results, and downloading and saving of results.

2 Portal Aspect

2.1 Overview

2.1.1 Portal is a Web Application

ID: DMS-PRTL-REQ-0001

Specification: The Portal aspect shall be accessible through commonly used desktop web browsers without requiring users to download and install local software packages.

Discussion: The supported browsers will be documented by the Project.

2.2 Discover Data

2.2.1 Portal Discovery of all Data Products

ID: DMS-PRTL-REQ-0002

Specification: The Portal aspect shall provide the capability to discover and access all the Project's released data products, including, but not limited to, the data products enumerated in the DPDD (LSE-163), the calibration database, and the Reformatted EFD, as well as all user data products to which a user has access.

Discussion: The Portal's workflows should allow a user to learn what data exist: what data releases are available, what image and catalog data they contain, the names of all databases, tables, and columns, etc. For all tabular data products the Generic Query requirements below cover the basic level of access provided.

2.2.2 Portal Access to Workspace

ID: DMS-PRTL-REQ-0003

Specification: The Portal aspect shall have the capability to discover all data in the user's Workspace.

Discussion: This will allow for data and images to be retrieved from the environment for use within the portal.

2.2.3 Semantic Linkage: Portal Workflows

ID: DMS-PRTL-REQ-0004

Specification: The Portal aspect shall provide for the identification and retrieval of semantically linked data.

Discussion: The Portal should provide straightforward UI workflows for starting from a selected data item (image or catalog entry) and identifying related data, including both direct data-dependency and provenance linkages and more scientifically oriented linkages such as the ability to navigate from an Object to its associated ForcedSources. Some of these linkages will be supported by the automatic exploitation of metadata exposed through the underlying data access APIs; others will be specifically designed workflows reflecting scientific understanding of the relationships among LSST data products and their processing.

2.2.4 Access to Calibration Products

ID: DMS-PRTL-REQ-0005

Specification: The Portal aspect shall enable access to Project calibration data products, both directly and via linkages from science data products generated using them.

Discussion: This is a sub-requirement of "Semantic Linkage: Portal" that carries forward a specific requirement originating from early SUIT specifications.

2.2.5 Coadded Image to Single-Epoch Image Associations

ID: DMS-PRTL-REQ-0006

Specification: The Portal aspect shall facilitate users following the associations between coadded images and the single epoch images that were used to generate them.

Discussion: This is a sub-requirement of "Semantic Linkage: Portal" that carries forward a specific requirement originating from early SUIT specifications.

2.2.6 Access to External Archives

ID: DMS-PRTL-REQ-0007

Specification: The Portal aspect shall provide an interface to outside catalog and image data services that are available via standard astronomical VO interfaces to enable a user to determine what external astronomical data are associated with a given location on the sky and return that data for use within the Portal.

Discussion: The purpose of this requirement is to help users connect LSST data with other data already available in community archives (e.g., IRSA, NED, MAST, HEASARC), and should support any VO-compliant archive. If access to a non-VO compliant archive is critical for the needs of the LSST project that could be considered here.

2.2.7 API for Access to Portal Session State

ID: DMS-PRTL-REQ-0008

Specification: The Portal aspect shall provide a network API that allows authenticated remote access by a user to aspects of their session state in the Portal. The minimal requirement is for access to the list of queries performed in that session.

Discussion: Access to the list of queries allows user code in the Notebook aspect to use those query references to retrieve the data which was being explored in the Portal.

2.3 Query Data

2.3.1 Support Synchronous and Asynchronous Queries

ID: DMS-PRTL-REQ-0009

Specification: The Portal aspect shall provide UI models for both synchronous and asynchronous queries, based on user preference, loading, and resource capabilities.

Discussion: This Portal capability should include an interface to initiate, monitor, and control the execution of both sync and async queries, as well as browse their results. Long running queries may be forced to be asynchronous. There is not a 1:1 relationship between this Portal capability and the analogous capability in the API aspect; for instance, the Portal may provide UI support for placing a long-running, technically synchronous (from the Data Access Services perspective), query into the background from the user's perspective.

2.3.2 Long Query Backgrounding

ID: DMS-PRTL-REQ-0010

Specification: The Portal aspect shall notify the user if a query is estimated to take longer than 60 seconds and will allow the user to put the query in background if desired.

Discussion: This requirement arose from user panel feedback and the original SUIT require-

ments review. It requires support from Data Access Services/database for query estimation and that is still TBD. Having the query "in the background" refers to the user's perspective of being able to proceed with other work while waiting for it to complete, and may or may not involve a truly asynchronous Data Access Services query.

2.3.3 Query Status and Termination Notification

ID: DMS-PRTL-REQ-0011

Specification: The Portal aspect shall notify the user of the status of user-initiated database queries, including whether the query has been terminated for any reason.

Discussion: There needs to be an operations discussion about when and why database queries are terminated, and what sorts of notifications will be generated by Data Access Services in such cases.

How the Portal Aspect will notify the user is not specified by this requirement; it might be by email to a registered address, or it might be through a notification mechanism in the Portal UI, or something else.

2.3.4 Query Results Size Limitation

ID: DMS-PRTL-REQ-0012

Specification: The Portal aspect shall have a mechanism to notify a user that a query result is predicted to exceed, or has exceeded, threshold(s) for the maximum results size allowed and that the query has been disallowed or terminated as a result.

Discussion: If the size limitation is applied at run time as query results accumulate, this is an important concrete sub-case of the DMS-PRTL-REQ-0011 requirement; if it is applied earlier, it is a separate mechanism.

The underlying Data Access Services capabilities required to support this remain to be specified in detail.

2.3.5 Query History Inspection

ID: DMS-PRTL-REQ-0013

Specification: The Portal aspect shall provide a user interface for browsing the history of LSST project and user database queries performed by the user, for re-executing a selected query on demand, and, for recent queries, re-retrieving their results.

Discussion: The implementation of this capability is assumed to rest on the API aspect, and the underlying database systems, supporting this functionality. The Portal interface to this is a thin UI. This should be imagined as a CASJobs-like query history. Note that because the API aspect is shared between the Portal and Notebook, this capability can be used to retrieve in the Portal the results of queries recently executed in the Notebook aspect; this allows complex ADQL queries to be formulated programmatically in the Notebook while still permitting their results to be inspected in the Portal. Because of the reliance on the API aspect, the Portal is not required to provide this capability for external archives.

2.3.6 Query Saving - Portal

ID: DMS-PRTL-REQ-0014

Specification: The Portal aspect shall provide a UI for the saving of a specification artifact for a user-performed query, either for downloading or for saving to the Workspace, and a UI for re-executing a saved query found in the Workspace or uploaded remotely.

2.3.7 General Tabular Queries

2.3.7.1 Generic Query

ID: DMS-PRTL-REQ-0015

Specification: The Portal aspect shall enable the generation of queries against any tabular data exposed in the API aspect.

Discussion: Provision of this relies on the Discovery and Reflection API requirement. The

idea is that some level of query will automatically be provided for every table even for tables for which no deliberately-designed custom search screens are available. Note that image metadata is tabular data, so a basic level of ability to query for lists of images is available through "generic queries", though more image-specific workflows will be provided.

2.3.7.2 Generic Query - Form-based

ID: DMS-PRTL-REQ-0016

Specification: The Portal aspect shall provide a search-builder form-based interface for generic table queries. This facility may have reduced functionality for user tables for which the user has not provided full, or accurate, metadata.

Discussion: The Discovery and Reflection API will be used to construct a form allowing query against any attribute of a table, taking the column metadata provided by the Reflection API into account to display units and other information that will help the user construct a meaningful query. We wish to allow and facilitate, but not require, users to provide detailed metadata for the tables they create; when they do, the system will take advantage of it.

2.3.7.3 Generic Query - ADQL-based

ID: DMS-PRTL-REQ-0017

Specification: The Portal aspect shall provide a means for entering a query against any table directly in ADQL. This facility shall be available for every table, including user-supplied tables.

Discussion: This is essentially a pass-through to the underlying API aspect TAP service. UI support for selecting from available identifiers (e.g. table or column names) for use in the ADQL, based on the Discovery and Reflection API, would be highly desirable. UI support for the specification of coordinate values for use in the ADQL, based on the Spatial Query Parameters requirements below, would also be desirable.

2.3.7.4 Query Result Size

ID: DMS-PRTL-REQ-0018

Specification: The Portal aspect shall provide UI support for any mechanism provided by the API Aspect for determining or estimating the number of rows matching the query criteria without generating a full return set.

Discussion: What level of support for this will be provided by the Data Access Services remains to be determined, so the requirement has been written to accommodate that uncertainty.

A proper implementation would rely on row-counting in the ADQL submitted to the Data Access Services.

2.3.7.5 Specific Query Types

This section covers two key types of queries, temporal and spatial, for which it is often necessary to perform transformations from user specifications to internal forms, and for which there are specific community expectations. Queries of these types may apply to any temporally or spatially organized tables, respectively, including both catalogs and image metadata tables. The spirit of all the requirements in this section is that they should be satisfied by providing appropriate UI elements, not merely by the provision of the raw ADQL query capability to, e.g., code a box search. It is assumed that the underlying query capabilities of the API aspect / Data Access Services / database, and in particular the support for tabular queries with ADQL, will enable the implementation. The availability of these queries is driven by the table and column metadata.

2.3.7.5.1 Identifier Queries

In general the Portal is required to support retrieval of data on any LSST data product entity (e.g., catalog row) by its unique ID. The requirements in this section elaborate on that.

QUERY BY IDENTIFIER

ID: DMS-PRTL-REQ-0028

Specification: The Portal aspect shall provide UI support for performing queries for data on any LSST data product entity with a unique ID by that ID.

Discussion: Given, e.g., an Object ID or a Visit ID, it should be possible to perform a query by that ID and retrieve all, or a selection of, the data for the corresponding Object or Visit. This requirement is limited to the direct content of the corresponding table row; the following requirements add semantic workflows returning additional data. It is desirable for this capability to be available wherever such a unique ID is displayed in the UI, though the mandatory requirement is only for a dedicated form for such a query.

QUERY BY LSST OBJECT AND SOURCE IDENTIFIERS: SPECIFIC MATCH TO IDENTIFIER

ID: DMS-PRTL-REQ-0029

Specification: The Portal aspect shall provide UI support to query and return data associated with a specific LSST (DIA)Object, (DIA)Source, or ForcedSource identifier, including catalog data associated with the entity as well as the image data and metadata directly associated with the measurement.

Discussion: For example, for an Object ID query, the Portal will, upon request, return information on the coadded images on which the Object was observed. For a ForcedSource ID, the Portal should return information on the associated Object and on the single-epoch image on which the ForcedSource measurement was made.

QUERY BY SOLAR SYSTEM OBJECTS: SPECIFIC MATCH TO IDENTIFIER

ID: DMS-PRTL-REQ-0030

Specification: The Portal aspect shall provide UI support to query and return data specifically associated with a Solar System Object.

Discussion: The Portal will need to be able to return the data (catalog and images) associated with a specific solar system identifier and that cover the time range wanted. This requires that the SSO database will have identified known solar objects that have been observed.

2.3.7.5.2 Spatial Queries

Much of this section covers both catalog and image metadata queries. In general the interpretation of a spatial query on image data is more complex than for catalog data. The requirements at present do not attempt to fully define the specific variants of spatial image searches supported (e.g., "contains a specified point", "contains a specified region", "overlaps with a region", "is fully contained within a region", etc.). We expect to support most or all common variants, but details of the underlying database systems may determine which variants are more or less performant.

SPATIAL QUERY PARAMETERS

This section covers the means of specification of coordinates for spatial queries in the Portal UI.

Positional Query: Position on the Sky

ID: DMS-PRTL-REQ-0020

Specification: The Portal aspect shall support queries based on an astrophysical position (i.e., coordinates) on the sky.

Discussion: The intent of this requirement is to enable users to search for data at or near locations in equatorial, ecliptic or galactic coordinates and not a specific astrophysical object. The Portal is responsible for any conversion required to put the coordinate query parameters into the forms accepted by the underlying API aspect / Data Access Services.

Positional Query: Multiple Positions/Objects

ID: DMS-PRTL-REQ-0021

Specification: The Portal aspect shall support list-based positional queries, with the coordinates used specified by any of the means of specifying positions required elsewhere herein.

Discussion: Lists of positions may be uploaded or taken from the Workspace. Efficient implementation of list-based queries requires a corresponding API aspect / Data Access Web

API service, to avoid the submission of large numbers of separate queries. The system is not required to, but may, accept lists in which different elements are specified in more than one of the supported means of specifying positions; e.g., it is not required to support a list that is partially in equatorial coordinates and partially in LSST object IDs, because of the parsing and interpretation complexities involved. The implementation should not preclude adding the ability to mix specification types later on, guided by feedback from the user community as to what is most useful. Specific coordinates required elsewhere in this document include "equatorial, ecliptic or galactic coordinates" in DMS-PRTL-REQ-0020 and DMS-PRTL-REQ-0022, "astrophysical source name" (DMS-PRTL-REQ-0023), "LSST catalog entry identifiers" (DMS-PRTL-REQ-0024), and "external Solar System Object identifiers" (DMS-PRTL-REQ-0025).

Positional Query: Astrophysical Coordinate Systems

ID: DMS-PRTL-REQ-0022

Specification: The Portal aspect shall support positional queries based on the following astrophysical coordinate systems: equatorial, ecliptic, and galactic.

Discussion: The UI should include information on (and possibly selection of) the particular definition of each available coordinate system.

Positional Query: Astrophysical Source Name Lookup

ID: DMS-PRTL-REQ-0023

Specification: The Portal aspect shall support the specification of coordinates for use within all positional queries by the use of source names in common community-established astrophysical source name lookup services.

Discussion: Services include, but are not limited to, NED, SIMBAD, and Horizons

Positional Query: LSST Object and Source Identifiers

ID: DMS-PRTL-REQ-0024

Specification: The Portal aspect shall support the specification of coordinates for use within

all positional queries by the use of specific LSST catalog entry identifiers, including those for the Object, DIAObject, Source, and DIASource tables. The default choice of coordinate columns within these tables to use for the ID-to-coordinate translation shall be documented and shall be able to be determined from the UI.

Discussion: LSST object and source identifiers are not necessarily associated with previously known astrophysical sources. The Portal will need to be able to interpret source and object names and return positions that can be used in positional-based searches.

Positional Query: Solar System Object Names

ID: DMS-PRTL-REQ-0025

Specification: The Portal aspect shall support positional queries based on external Solar System Object identifiers, including names from, but not limited to, NASA's Navigation and Ancillary Information Facility (NAIF), the Minor Planet Center, and JPL's Horizons, coupled with a date/time range specification.

Discussion: The intent here is to enable a user to enter a solar system object name (e.g., 25155 van Belle) and, for a given time range, get back a list of observations that LSST may have made of that object because the survey has overlapped the position of that object at the appropriate times. This capability is expected to be available for tables that are both spatially and temporally organized, e.g., Visit or DIASource.

SPATIAL QUERY TYPES

Positional Query by Region: Cone-Search

ID: DMS-PRTL-REQ-0026

Specification: The Portal aspect shall support position-based queries based on a cone-shaped radial search.

Positional Query by Region: Box-Search

ID: DMS-PRTL-REQ-0027

Specification: The Portal aspect shall support position-based queries based on a coordinate-system box search.

2.3.7.5.3 Temporal Queries

QUERY BY DATE AND TIME: TIME RANGE OF OBSERVATION

ID: DMS-PRTL-REQ-0019

Specification: The Portal aspect shall support queries based on times and ranges of date/time values in both UT and (barycentric) Julian date.

Discussion: The intent of this requirement is to enable users to search for an event within a range of times/dates and is akin to a positional box search of images that overlap that time range or are fully enveloped in that time range. This type of query should be available for all time-point based tables in the LSST data products (e.g., Visit, Source, ForcedSource, DIASource, and EFD), and more generally for all tables with columns indicated as times by their metadata (e.g., by UCDs). The Portal is responsible for conversion of user-supplied times to the appropriate form for use in the underlying API aspect and database queries.

2.3.8 Catalog Queries

This section contains requirements on queries related to specific LSST catalogs.

2.3.8.1 Tabular Data Query Specifications

ID: DMS-PRTL-REQ-0031

Specification: The Portal aspect shall provide a user interface to execute queries of the (DIA)Object and (DIA)Source tables, driven by the data dictionary associated with the tables.

Discussion: This should be satisfied almost completely by the "Generic Query - Form-Based" requirement above, but with some additional work on the UI to produce a more friendly workflow.

2.3.8.2 Query Tabular Data based upon Image MetaData

ID: DMS-PRTL-REQ-0032

Specification: The Portal aspect shall be able to support queries of catalog data that include constraints on the properties of the images on which the catalog measurements were made.

Discussion: This allows, for instance, limiting the return of Source catalog entries to measurements made on images taken with constraints on airmass, moon angle, etc.

2.3.8.3 Queries on the Alerts Database

ID: DMS-PRTL-REQ-0033

Specification: The Portal aspect shall provide a query interface to the Alert Database, allowing searches based on parameters which shall include, but may not be limited to, Alert ID, time of alert, position on the sky, filter, and alert characteristics.

Discussion: This capability, as all others in this section, is limited to data rights holders. Non-data-rights holders have access to alerts only through the alerts stream(s) sent to public brokers.

2.3.8.4 Access to Original Alert State

ID: DMS-PRTL-REQ-0034

Specification: The Portal aspect shall provide access to the alerts as they were originally raised.

Discussion: This is in the context of features elsewhere in the system that may allow for a chain of annotations of alerts.

2.3.9 Image Metadata Queries

Image metadata is a special case of tabular data. Some dedicated search capabilities are provided for it beyond the generic tabular searches.

2.3.9.1 Query for Single Epoch Visit Images

ID: DMS-PRTL-REQ-0035

Specification: The Portal aspect shall enable a user to proceed from a visit-selection query or a list of visits and return a list of all single-epoch images of a specified type corresponding to those visits.

Discussion: The common image types will be raw, PVI (processed, i.e., calibrated, visit image), and difference image.

2.3.9.2 Query for Single Epoch Raft Images

ID: DMS-PRTL-REQ-0036

Specification: The Portal aspect shall enable a user to limit the list of images selected by a single-epoch visit image query to those from a specified raft.

Discussion: The intent is to enable the return of every image from RAFT X as a function of time or filter to enable viewing just that raft.

2.3.9.3 Query for Single Epoch CCD Image

ID: DMS-PRTL-REQ-0037

Specification: The Portal aspect shall enable a user to limit the list of images selected by a single-epoch visit image query to those from a specified CCD.

Discussion: The intent is to enable the return of every image from CCD X as a function of time or filter to enable viewing just that CCD.

2.3.9.4 Single-Epoch Image Query Specifications

ID: DMS-PRTL-REQ-0038

Specification: The Portal aspect shall provide UI support for queries for visits and their single-epoch images of specified type, based on image metadata parameters including pointing, time and date, and filter selection, as well as on parameters from the Reformatted EFD.

Discussion: The parameters specifically named are expected to be highlighted in the UI, rather than requiring the user to scroll through a long generic-table-query form to find the appropriate fields. The UI will provide support for generating a join query including tables from the R-EFD, and for selecting the R-EFD tables and columns to use.

2.3.9.5 Coadded Image Query Specifications

ID: DMS-PRTL-REQ-0039

Specification: The Portal aspect shall provide UI support for queries for coadded images based on the image metadata that describe the provenance of the images (e.g., filters, position on the sky, date, number of single-epoch images, coverage, survey depth).

2.3.10 Image Cutout Queries

The queries which follow return the actual image data.

2.3.10.1 Query for Coadded Image Cutouts

ID: DMS-PRTL-REQ-0041

Specification: The Portal aspect shall enable a user to perform a coadded image query, as

above, and additionally return a list of sub-images (i.e., cutouts) from the all-sky co-added images based upon user-specified center position and image size, including the appropriate metadata for describing the image cut-outs.

Discussion: This is a front end to a cutout capability in the API aspect.

2.3.10.2 Query for Single Epoch Image Cutouts

ID: DMS-PRTL-REQ-0040

Specification: The Portal aspect shall enable a user to perform a single-epoch image query, as above, and additionally return a list of sub-images (i.e., cutouts) from them based upon a specified center position, time range, and image size, including the appropriate metadata for describing the image cut-outs.

Discussion: This is a front end to a cutout capability in the API aspect.

2.4 Display Data

The Display Data requirements are unique to the Portal. The Notebook and API aspects of the LSP do not provide visualization services as such, though of course the Notebook permits the use of a wide variety of visualization libraries.

2.4.1 Visualization of Tabular and Image Data

ID: DMS-PRTL-REQ-0042

Specification: The Portal aspect shall provide the capability to visualize all tabular and image data products in the DPDD, as well as user data products.

Discussion: The products in the DPDD are the primary data products for use by the LSST users. The "tabular and image" qualification indicates that the Portal is not required to provide a dedicated visualization for all data products that do not naturally fall into one of those categories. For user data products, the amount of detail and labeling, and the amount of UI

support, will be less if they lack the full level of metadata that comes with the Project's own data products.

2.4.2 Visualization of Ancillary Information

ID: DMS-PRTL-REQ-0043

Specification: The Portal aspect shall include the ability to visualize selected ancillary information produced by the LSST pipeline including, but not limited to, image regions, image bit-planes, survey footprints, focal-plane footprints and PSF representations.

Discussion: The intent here is to call attention to the fact there is more than just the survey images and coadds that are have a "2-dimension" form that need to be visualized and presented to the user in the interface. The specific ancillary data products to visualize will be determined during construction, based in part on feedback received during PDAC operation and the use of the Portal tools by developers. It is desirable that custom visualizations be available for important and frequently used ones such as Footprints (which can readily be displayed as pixel overlays). Where dedicated Portal visualizations are not available, however, users should be able to use either LSST-provided or community libraries in the Notebook aspect to create custom visualizations.

2.4.3 Linking Visualization of Image Data to Tabular Data

ID: DMS-PRTL-REQ-0044

Specification: The Portal aspect shall provide the capability for the user to navigate between visualized tabular data and visualized image data.

Discussion: For instance, there should be very simple UI support to display an image based on a row in an image metadata table, or to navigate from a selected source overplotted on an image to further information about that source.

2.4.4 Visualization of Uploaded Tabular and Image Data

ID: DMS-PRTL-REQ-0045

Specification: The Portal aspect shall support a convenient workflow for the visualization of uploaded tabular and image data products.

Discussion: The idea is to provide something close to a one-button "show me this" workflow.

2.4.5 Visualization of Workspace Data

ID: DMS-PRTL-REQ-0046

Specification: The Portal aspect shall support a convenient workflow for the visualization of data selected in a workspace browser.

Discussion: This should appear as a standard "select and open" workflow, with the Portal determining a reasonable action to take based on its determination of the type of Workspace data selected.

2.4.6 Display Tabular Data

2.4.6.1 Property Sheets

2.4.6.1.1 Table Row Property Sheet

ID: DMS-PRTL-REQ-0047

Specification: The Portal aspect shall permit the inspection of all the data in a single row of a tabular data query result as a "property sheet" for that row, taking advantage of available metadata to supply units and other semantic information for each column value.

Discussion: Resources permitting, the property sheet may be elaborated to provide additional functionality (typically, further queries) associated with particular data items displayed.

Property sheets should, where enabled by metadata, appropriately exhibit relationships between columns, such by displaying a value and its uncertainty together.

The system *must* provide a generic property sheet functionality for any table for which full metadata is available. It *may* also provide custom property sheets for commonly-queried tables such as Object, ForcedSource, Visit, etc. that provide a more scientifically useful layout, and additional available workflows, than possible just from the metadata.

2.4.6.1.2 Alert Visualization

ID: DMS-PRTL-REQ-0048

Specification: The Portal aspect shall provide for the users a "property sheet" for the contents of an alert packet including, but not necessarily limited to, the alert postage stamp image, the postage stamp time series, the photometric time series, the source and object information (e.g., position, brightness).

Discussion: This display is based on a query performed on the alert database. It may also be desirable to provide this functionality based on an actual published alert packet, e.g., by permitting the upload of such a packet for display. The alert property sheet should facilitate further exploration based on, e.g., the associated Object IDs in the alert.

2.4.6.2 Table Viewers

2.4.6.2.1 Display of Tabular Data

ID: DMS-PRTL-REQ-0049

Specification: The Portal aspect provide the capability to display tabular data in an interactive environment which displays the tables by columns and rows.

Discussion: The intent of this requirement is to capture that the database query returns are displayed.

2.4.6.2.2 Column Selection of Tabular Data

ID: DMS-PRTL-REQ-0050

Specification: The Portal aspect shall provide the capability to select, for display and downloading, specific columns within the tabular data viewer.

Discussion: The intent of this requirement is to enable users to decide which columns are desired for display and download.

2.4.6.2.3 Display Order of Columns of Tabular Data

ID: DMS-PRTL-REQ-0051

Specification: The Portal aspect shall provide the capability to change the display order of the columns for tabular data.

Discussion: The intent of this requirement is to enable users to decide which order to view the columns.

2.4.6.2.4 Copying of Tabular Data

ID: DMS-PRTL-REQ-0052

Specification: The Portal aspect shall provide the capability of interactively selecting and copying data within a displayed data table.

Discussion: The intent of this requirement is to enable users to use the mouse to select fields within a displayed table and utilize standard copy mechanisms.

2.4.6.2.5 Row Selection of Tabular Data

ID: DMS-PRTL-REQ-0053

Specification: The Portal aspect shall provide the capability to select, for display and down-

loading, specific rows within the tabular data.

Discussion: The intent of this requirement is to enable users to decide which rows are desired for display and download and these may be different from the filtered rows (see filtering requirement below).

2.4.6.2.6 Paging of Tabular Data

ID: DMS-PRTL-REQ-0054

Specification: The Portal aspect shall provide the capability to display tabular data in a paged format.

Discussion: The intent of this requirement is to capture that the database query returns may be too large to display all at once to the user and a form of paging will be necessary.

2.4.6.3 Tabular Data Graphing

2.4.6.3.1 XY Scatter Plots

ID: DMS-PRTL-REQ-0055

Specification: The Portal aspect shall enable the creation and display of 2-dimensional xy-plots from tabular data.

2.4.6.3.2 Histograms

ID: DMS-PRTL-REQ-0056

Specification: The Portal aspect shall enable the creation and display of 1-dimensional and 2-dimensional histograms of tabular data.

Discussion: This requirement is about producing traditional histograms in 1 and 2 dimen-

sions where the number of items within a bin are reported

2.4.6.3.3 Symbol Size, Shape, and Color Coding in XY(Z) Scatter Plots

ID: DMS-PRTL-REQ-0057

Specification: The Portal aspect shall enable the use of symbol size, shape, and color as indicators of additional tabular data associated with the XY(Z)-data plotted.

Discussion: This is a generalized requirement to enable the visualization of multi-dimensional data in a 2-d scatter plot environment.

2.4.6.3.4 Plot Quantitative Uncertainties

ID: DMS-PRTL-REQ-0058

Specification: The Portal aspect shall be able represent uncertainties in the plotting of data.

Discussion: This flows down from higher-level requirements above, and has implications for the creation of the necessary metadata to support this in the first place.

2.4.6.3.5 Plot Asymmetric Quantitative Uncertainties

ID: DMS-PRTL-REQ-0059

Specification: The Portal aspect shall be able represent uncertainties in the plotting of data that are unequal in value for the positive and negative directions.

Discussion: Uncertainties often have different limits in the positive and negative directions and as a result representation of the uncertainties will be different.

2.4.6.3.6 Plot Upper and Lower Quantitative Limits

ID: DMS-PRTL-REQ-0060

Specification: The Portal aspect shall be able represent upper and lower limits in the plotting of tabular data.

Discussion: Often the values are non-detections and a limit is estimated on that value. The limit can be an upper limit (e.g., flux) or a lower limit (e.g., magnitude). Typically, the limit is represented by a flat line at the position of the value and an up or down arrow from that position. For tables with full metadata including UCDs this can be handled generically: the "stat.min" and "stat.max" UCDs can be used to trigger the display of the distinctive symbols.

2.4.6.3.7 Multiple XY-Plots on the Same Display

ID: DMS-PRTL-REQ-0061

Specification: The Portal aspect shall be able to overlay multiple plots on the same display, differentiated by plotting colors, symbols, line styles, and shading.

Discussion: The intent of this requirement is enable the plotting of multiple graphs on the same plotting canvas. An example would be RA vs Time and Dec vs Time or u-band mag vs i-z color and r-band mag vs i-z color.

2.4.7 Display Image Data

NB: These requirements are solely about *displaying* images in the Portal, not querying for them. It is highly likely that they will all be satisfied with the same SUIT component; they are enumerated separately if for no other reason than to facilitate verification that the common data products of LSST can actually be displayed in useful ways.

2.4.7.1 Display Native Single-Visit Image Data Products

ID: DMS-PRTL-REQ-0062

Specification: The Portal aspect shall have the ability to display the native single-visit image

data products, including raw images, Processed Visit Images (PVIs), and difference images, as well as the standard single-exposure calibration images used as inputs for flats, bias frames, etc.

Discussion: The native raw data will contain amplifier-level data with full pre-scan, serial overscan, and parallel overscan. The other native single-visit data products will be at CCD level.

2.4.7.2 Display Raft- and Focal-Plane-Level Single-Visit Image Data

ID: DMS-PRTL-REQ-0063

Specification: The Portal aspect shall have the capability to generate a synthetic display of image data at raft level and at full focal plane (FPA) level.

Discussion: No such data products will exist *per se*; the requirement is for the Portal to be able to show, e.g., the coverage of raft-level or FPA-level visit imaging on the sky.

2.4.7.3 Display Single Visit Image Cut-Out

ID: DMS-PRTL-REQ-0064

Specification: The Portal aspect shall have the capability to display a cutout from a single visit image.

Discussion: Corresponds to the ability to request the generation of cutouts and cross-CCD mosaics.

2.4.7.4 Display Native Coadded Image Data Products

ID: DMS-PRTL-REQ-0065

Specification: The Portal aspect shall have the capability to display the native coadded image

data products, i.e., the patch-level images.

2.4.7.5 Display Coadded Image Cutouts / Mosaics

ID: DMS-PRTL-REQ-0066

Specification: The Portal aspect shall have the capability to display cutouts and mosaics from coadded image data products, as delivered from the API aspect.

Discussion: Corresponds to the ability to request the generation of cutouts and mosaics. These may cover areas both smaller than and larger than the native "patch" scale.

2.4.7.6 Display Calibration Image Data Products

ID: DMS-PRTL-REQ-0067

Specification: The Portal aspect shall have the capability to display the calibration image data products such as synthetic flats, bias frames, and the like.

Discussion: Note that these images may not have WCS information.

2.4.7.7 Display User-provided Images

ID: DMS-PRTL-REQ-0068

Specification: The Portal aspect shall have the capability to display user-provided images in widely-used astronomical community formats, including FITS, and shall properly interpret a variety of commonly-used WCS specifications in the image headers.

Discussion: This would be used, for instance, to allow a user to view LSST catalog data superposed on a user-provided image. FITS is the only currently supported image file format; others will be considered as community usage develops or as part of supporting all-sky visualization.

2.4.7.8 Image Property Sheet

ID: DMS-PRTL-REQ-0069

Specification: The Portal shall have the ability to display a property sheet for an image data product or user-provided image, displaying image format and other header data.

Discussion: This is information on the image artifact itself, e.g., on the contents of FITS headers. It should not be confused with the property sheet associated with an image metadata table entry, though there may be a strong overlap in content, and for UX purposes they may be displayed in combined screens.

2.4.7.9 Image Display Tools

2.4.7.9.1 Provide Coordinate Display Tools for Images

ID: DMS-PRTL-REQ-0070

Specification: The Portal shall provide all the capabilities in the Coordinate Display Tools section herein for image displays. Specific capabilities will depend on the availability of WCS information for an image.

2.4.7.9.2 Image Pixel Content Display

ID: DMS-PRTL-REQ-0071

Specification: The Portal aspect shall have the capability to inspect the pixel content of an image at the position of the mouse cursor. This capability shall be integrated with the Point Coordinate Display capability for the image.

2.4.7.9.3 Image Spatial Manipulation

ID: DMS-PRTL-REQ-0072

Specification: The Portal aspect shall have the capability to resize, rescale, re-project, zoom, and crop the image display and save or download the current view.

2.4.7.9.4 Multi-Image Scaling and Aligning

ID: DMS-PRTL-REQ-0073

Specification: The Portal aspect shall have the capability to display multiple images on a common astrophysical coordinate scale and aligned on the screen in a common astrophysical orientation.

Discussion: The point behind this requirement is to enable viewing the same part of the sky in different filters aligned to same orientation and scaled to the same screen resolution for both single-frame and coadded images.

2.4.7.9.5 Image Appearance Manipulation

ID: DMS-PRTL-REQ-0074

Specification: The Portal aspect shall enable the user to change the view of a displayed image including, but not necessarily limited to, the color table, the stretch function, and the displayed data range.

2.4.7.10 Image Overlays

These capabilities concern the overlay of additional information on an image by matching to one of its coordinate systems.

2.4.7.10.1 Image Mask and Variance Overlays

ID: DMS-PRTL-REQ-0075

Specification: The Portal aspect shall enable the overlaying of additional pixel-oriented data

on an image, including image masks (bit planes) and variance data.

Discussion: This also enables, but does not require, the overplotting of two-dimensional density plots such as depth maps.

2.4.7.10.2 Image Plot Overlays

ID: DMS-PRTL-REQ-0076

Specification: The Portal aspect shall enable the overlaying of tabular data on an image, either based on pixel coordinates or astrophysical coordinates, as supported by the availability of coordinate system information.

Discussion: More generally, this should work for any two-dimensional plot data that shares a coordinate system with an image. The catalog data need not be semantically linked with the image by anything other than the coordinate system used.

2.4.7.10.3 Image Overlays: Adjustment of Colors and Positions

ID: DMS-PRTL-REQ-0077

Specification: The Portal aspect shall have the capability for a user to configure the annotations, colors, transparency, and positions (where applicable) of any image overlays, including those resulting from the use of the Coordinate Tools.

Discussion: Often the default color and position of the overlay needs to be changed for clarification.

2.4.8 All-Sky Visualization

2.4.8.1 Display All-Sky HEALPix Image

ID: DMS-PRTL-REQ-0078

Specification: The Portal aspect shall be able to display an all-sky image in the HEALPix format.

Discussion: In addition to flux images such as all-sky coadds, the LSST pipelines and/or quality assessment processes may generate a variety of all-sky metrics, diagnostics, and other artifacts in this format.

2.4.8.2 Zoom In and Out on a HEALPix Image

ID: DMS-PRTL-REQ-0079

Specification: The Portal aspect shall enable a user to zoom in and out on a HEALPix all-sky image.

Discussion: The zooming in and out will enable a user to change effective spatial resolution of the image on the screen, transferring across levels of the image hierarchy.

2.4.8.3 Pan Around on a HEALPix Image

ID: DMS-PRTL-REQ-0080

Specification: The Portal aspect shall enable a user to move around within a HEALPix all-sky image when the full image is not displayed on the screen.

Discussion: The panning is intended to enable a user to move around on the sky with a fixed zoom level. Panning does not apply if the full all-sky image is visible on the screen.

2.4.8.4 HEALPix Pixel Selection

ID: DMS-PRTL-REQ-0081

Specification: The Portal aspect shall enable a user to select individual HEALPix pixels or groups of pixels and obtain references to them which can be used in other LSP aspects.

Discussion: UI selection of pixels, at selectable scales, can be done in the Portal, with the selections then used in other aspects.

2.4.8.5 Retrieve HEALPix-Associated Data

ID: DMS-PRTL-REQ-0082

Specification: The Portal aspect shall enable a user to retrieve metadata and data associated with selected HEALPixes and display that information as tabular or image data as appropriate.

Discussion: The HEALPix pixels will be associated with metadata (e.g., which objects are associated with that position on the sky) or data (e.g., what is the FWHM of all of the sources within that pixel on the sky). That metadata/data will be retrievable via a selection of the HEALPix pixels.

2.4.9 Coordinate Display Tools

These tools are intended to be useful in any two-dimensional display of data in a well-defined spatial coordinate system: images as well as x-y plots where the axes are spatial.

2.4.9.1 Coordinate Display Applicability

ID: DMS-PRTL-REQ-0083

Specification: The Portal aspect shall have the capability to provide the set of coordinate system display and measurement tools in this section for any two-dimensional data display where both coordinates have a spatial interpretation, except as further specified below.

Discussion: The knowledge that coordinates in an x-y plot are spatial will in general depend on the availability of suitable metadata to define the coordinates.

2.4.9.2 Point Coordinate Display

ID: DMS-PRTL-REQ-0084

Specification: The Portal aspect shall have the capability to inspect the coordinates in a two-dimensional display that correspond to the position of the mouse cursor. When coordinate system conversion information is available, display of all available coordinates shall be supported. In particular, when available (e.g., for data associated with single-epoch and calibration images) the coordinate display shall include focal plane array electronic as well as spatial coordinates. When available (i.e., for an image, when a WCS is present) the coordinate display shall include astrophysical coordinates.

Discussion: When applied to images and at high zoom levels, the UI shall make clear whether pixel center coordinates or continuous cursor-location coordinates are being displayed.

2.4.9.3 Distance Measurement Tool

ID: DMS-PRTL-REQ-0085

Specification: The Portal aspect shall have the capability to determine the distance between two positions within an image or 2-dimensional plot in both image/plot coordinates (electronic or spatial X and Y) and in astrophysical coordinates (if applicable). Calculations shall be performed in spherical geometry where appropriate.

Discussion: The point behind this requirement is to enable distance determinations of equatorial, galactic, and ecliptic coordinates and make sure spherical geometry is used. When applied to general two-dimensional plots, distance measurement should only be supported when a metric and a geometry, or at least the relative scale between the two displayed coordinates, is known. This means that this capability is particularly dependent on metadata support.

2.4.9.4 Coordinate Grid Overlays

ID: DMS-PRTL-REQ-0086

Specification: The Portal aspect shall have the capability to display one or more coordinate

grids on top of images or two-dimensional plots with a known astrophysical coordinate system.

Discussion: The point behind this requirement is to enable viewing of equatorial, galactic, and ecliptic coordinates overlays at the same time.

2.4.9.5 Astrophysical Compass Overlay

ID: DMS-PRTL-REQ-0087

Specification: The Portal aspect shall have the capability to display a North-East Compass on an image or two-dimensional plot with a known astrophysical coordinate system.

Discussion: So the user knows that the (particularly for the generally randomly rotated single-epoch data) which directions are North and East.

2.4.9.6 Geometric Figure Overlays

ID: DMS-PRTL-REQ-0088

Specification: The Portal aspect shall enable the drawing, display, and selection of a closed 2-dimensional polygon on any two-dimensional image.

Discussion: This is a general requirement that enables the overlay of a polygon on an image or a plot. A polygon could be a circle, an ellipse, or an N-vertices polygon. The purpose of this is enable area selection on the images or plots.

2.5 Exploration and Analysis

Requirements in this section are unique to the Portal aspect. The equivalent functionality in the Notebook aspect simply arises from its nature as an open programming environment.

2.5.1 Table Manipulation

2.5.1.1 Sorting of Tabular Data by Column

ID: DMS-PRTL-REQ-0089

Specification: The Portal aspect shall provide the capability to sort tabular data by a single column within the table and redisplay the sorted table.

Discussion: This is the traditional sorting by one column.

2.5.1.2 Simple Filtering of Tabular Data

ID: DMS-PRTL-REQ-0090

Specification: The Portal aspect shall provide the capability to filter tabular data by a single column, including but not limited to less than (<), less than or equal (<=), greater than (>), greater than or equal (>=), equal (=), not equal (!=) and not null (!=null).

Discussion: The intent of this requirement is enable simple one-dimensional filtering on a single column or a series of columns and-ed together.

2.5.1.3 Calculated Filtering of Tabular Data

ID: DMS-PRTL-REQ-0091

Specification: The Portal aspect shall provide the capability to filter a table by single column where the filter has simple arithmetic calculations applied to the column values, including but not limited to sqrt, log, log10, exponentials and trigonometric functions.

Discussion: The intent of this requirement is enable simple arithmetic functions on the values in the columns prior and then do the filtering on that value (e.g., $X*0.33 < 10$).

2.5.1.4 Filtering of Tabular Data by Multiple Columns

ID: DMS-PRTL-REQ-0092

Specification: The Portal aspect shall provide the capability to filter tabular data by multiple columns within the table and redisplay the filtered table.

Discussion: Apply filters from various columns.

2.5.1.5 Calculated Quantities on Tabular Data

ID: DMS-PRTL-REQ-0093

Specification: The Portal aspect shall provide the capability to add an additional column to the displayed table based upon an arithmetic operations on columns within the displayed table and display the new column.

Discussion: This intent of this requirement is enable the user to combine columns in the table to form a new column that can be included in the table and used in the same manner as the intrinsic columns.

2.5.1.6 Statistical Measurements on Tabular Data

ID: DMS-PRTL-REQ-0094

Specification: The Portal aspect shall enable the capability to perform a set of statistical measurements (e.g., mean, median, RMS, skew, kurtosis) on tabular data selected by the user.

2.5.1.7 Saving Displayed Tabular Data

ID: DMS-PRTL-REQ-0095

Specification: The Portal aspect shall provide the capability to save and or download tabular data as it is displayed in the interface maintaining the content, filtering, and sorting.

Discussion: This intent of this requirement is enable the user to manipulate the table (e.g., sorting, filtering, calculated quantities and save that table to either workspace or in an offline download).

2.5.2 Image Manipulation

2.5.2.1 False-color Images Creation and Display

ID: DMS-PRTL-REQ-0096

Specification: The Portal aspect shall have the capability to create and display false-color images composed from any user-selectable set of filters from multiple filter views of the same region.

2.5.2.2 Statistical Measurements on Image Data

ID: DMS-PRTL-REQ-0097

Specification: The Portal aspect shall enable the capability to perform a set of statistical measurements (e.g., mean, median, RMS, skew, kurtosis) on user-selected regions in images.

2.5.3 Special Workflows

The following requirements describe specific workflows or screens that combine various Portal capabilities to deliver scientifically useful functions.

2.5.3.1 Overlay Catalog of Sources and Objects on Images

ID: DMS-PRTL-REQ-0098

Specification: The Portal aspect shall be able to overlay the positions of catalog sources

and objects on a displayed image based upon astrophysically-based or observatory-based coordinates.

2.5.3.2 Overlay LSST-Derived Orbits

ID: DMS-PRTL-REQ-0099

Specification: The Portal aspect shall have the capability to overlay predicted positions from the orbits of solar system objects in the LSST catalog on to images.

Discussion: This is envisioned as the ability to display a specific prediction for a position along an orbit on a single-epoch image, as well as a set of predictions for an orbit on a coadded image or all-sky map.

It would also be useful to support overlay of predicted positions from user-supplied orbits in community-standard forms. The capabilities to be provided in this area will be determined during construction.

It might further be useful to be able to overlay intermediate data products such as tracks and tracklets; whether it is desirable and feasible to provide this will be determined during construction.

2.5.3.3 Overlay User-provided Catalogs on Images

ID: DMS-PRTL-REQ-0100

Specification: The Portal aspect shall be able to overlay user-provided source lists or catalogs on images.

2.5.3.4 Overlay User-provided Region Files on Images

ID: DMS-PRTL-REQ-0101

Specification: The Portal aspect shall be able to overlay user-provided region files (e.g., DS9 region files, focal plane outlines) on images.

2.5.3.5 Display of Camera Artifacts as Overlays

ID: DMS-PRTL-REQ-0102

Specification: The Portal aspect shall have the capability to display camera artifacts including but not limited to image crosstalk matrices, ghost image identifications, saturation, and column bleeding.

Discussion: The intent of this requirement is to enable the users to be able to see where artifacts may be affecting the data. These artifacts may not be stored in image format and may need to be reconstructed algorithmically.

2.5.3.6 Single-Object Time-Domain Image View

ID: DMS-PRTL-REQ-0103

Specification: The Portal aspect shall have the capability to view an image time series that maintains the same physical scale, photometric scale, and image size display of a cutout area centered on an LSST object

Discussion: If the object moves, then the images stay centered on the object.

2.5.3.7 Position-based Time-Domain Image View

ID: DMS-PRTL-REQ-0104

Specification: The Portal aspect shall have the capability to view an image time series that maintains the same physical scale, photometric scale, and image size display of a specified position on the sky.

Discussion: If the object moves, then the images stay centered on the sky and the object appears to move.

2.5.3.8 Brightness Light Curves

ID: DMS-PRTL-REQ-0105

Specification: The Portal aspect shall have the capability to display graphically the brightness/magnitude of an LSST Object or Source or Forced Source as a function of time.

Discussion: This is a specific implementation of the xy-plot capabilities

2.5.3.9 Linked Tables, Plots, and Images

ID: DMS-PRTL-REQ-0106

Specification: The Portal aspect shall have the capability to have tabular data, plots, and images with overlays connected via brushing and linking.

Discussion: Updates to the data in any one visualization tool (e.g., plot, image, table) creates an update in other visualization tools. For example, selection of a set of photometry points in a color-color plot causes the corresponding rows in an Object table and symbols on an image to be highlighted; or, application of a selection predicate to a table causes the corresponding points in a plot to be highlighted.

2.5.3.10 Data Selection from a Plot or Image

ID: DMS-PRTL-REQ-0107

Specification: The Portal aspect shall enable the selection of data contained inside or outside a closed 2-dimensional polygon on an xy-plot, 2-dimensional data structure (e.g., density plot), and a 2-dimensional image.

Discussion: This is a general requirement that enables the selection of data from inside or outside a polygon.

2.5.3.11 Saving Data Selection from a Plot or Image

ID: DMS-PRTL-REQ-0108

Specification: The Portal aspect shall enable the saving of data selected via a polygon selection across the linked images, tables, and plots.

Discussion: An example here is to have an image up; draw a polygon on the image to select the area on the sky. All the tabular data associated with sources and objects in that part of the sky would be selected and saved.

2.6 User Storage

2.6.1 Access to User Databases

ID: DMS-PRTL-REQ-0109

Specification: The Portal aspect shall provide read/write access to user databases (Level 3 tabular data products) and shall implement any access restrictions placed on such data.

2.7 Transfer Data

2.7.1 Tabular Data Download

ID: DMS-PRTL-REQ-0110

Specification: The Portal aspect shall include a mechanism for a user to download to a remote site, Workspace, or to an existing or new user database the tabular results from a database query, including for catalog or image metadata.

Discussion: This may be implemented as a pass-through to the API aspect when applied to

the results of a query against an LSST database.

2.7.2 Image Data Download

ID: DMS-PRTL-REQ-0111

Specification: The Portal aspect shall include mechanisms for a user to download image data to a remote site or to the Workspace, from both screens displaying images and screens displaying lists of image metadata.

Discussion: Again, this should be implemented as a pass-through to the API aspect.

2.7.3 Selected Image Download

ID: DMS-PRTL-REQ-0112

Specification: The Portal aspect shall support user selection for download of a subset of the images in an image metadata table or image cutout table.

2.7.4 Download Volume Estimation

ID: DMS-PRTL-REQ-0113

Specification: The Portal aspect shall provide an estimate of the data download volume prior to a user confirming the download option.

Discussion: The intent of this requirement is to enable the users to understand how large a query result or data set may be prior to the full retrieval and downloading of that data to either the workspace environment or their own local environment.

2.7.5 Long Download Completion Notification

ID: DMS-PRTL-REQ-0114

Specification: The Portal aspect shall notify the user with an estimate of how long a down-

load is expected to take. The user can continue to monitor the download; an option shall be provided to notify the user when the download has completed.

Discussion: Data Access Services requirement for async.

2.8 Performance

Substantially more work needs to be done on codifying performance requirement for the components of the LSP.

2.9 Control and Management

2.9.1 APIs for Visualization Components

ID: DMS-PRTL-REQ-0115

Specification: The Portal aspect shall provide a documented application program interface that allows users and services at any location to access and manipulate the Portal's visualization services

Discussion: This is intended to be enable API control of the visualization components and tool-level visualization services to be called and controlled through an API. There will be a Web API as well as a Python wrapper for it.

2.9.2 Storage Quotas User Interface

ID: DMS-PRTL-REQ-0116

Specification: The Portal aspect shall provide the user with an understanding of the current status of their storage allocations

Discussion: This requirement is about the SUIT implementing the quotas defined by the DM system. Those requirements need to defined and work through by the Project and the DM.

2.9.3 Computational Quotas User Interface

ID: DMS-PRTL-REQ-0117

Specification: The Portal aspect shall provide the user with an understanding of the current status of their allocations.

Discussion: This requirement is about the SUIT implementing the quotas defined by the DM system. Those requirements need to be defined and worked through by the Project and the DM.

2.9.4 Portal Display Preferences

ID: DMS-PRTL-REQ-0118

Specification: The Portal aspect shall enable a user to establish and save viewing preferences, including, but not limited to, which tabular data columns to view, how tables should be sorted by default, which calculated quantities appear within a table, what image stretch and color tables, what types of plots are generated, how data are overlaid on images.

Discussion: The intent behind this requirement is to enable user to set up a working environment within the default portal or the workspace environment and be able to save the state of the workflow.

2.9.5 Alert Subscription

2.9.5.1 Alert Subscription Service

ID: DMS-PRTL-REQ-0119

Specification: The Portal aspect shall provide an interface to the alert subscription service that allows authenticated users with LSST data rights to subscribe to a stream of alert events.

Discussion: This is just a UI for the underlying capability developed under 02C.03.03. Note that users without data rights will have to subscribe to alerts through brokers external to the LSST project.

2.9.5.2 Pre-defined Alert Filters

ID: DMS-PRTL-REQ-0120

Specification: The Portal aspect shall provide an interface to permit alert subscriptions to be configured with Project-provided alert filters.

2.9.5.3 User-defined Alert Filters

ID: DMS-PRTL-REQ-0121

Specification: The Portal aspect shall provide an interface to permit alert subscriptions to be configured with user-defined alert filters.

2.9.5.4 Alert Subscription Monitoring

ID: DMS-PRTL-REQ-0127

Specification: The Portal aspect shall report feedback about the status and performance of a user's filters in the alert subscription service.

Discussion: This is a front end to information exposed by an API provided by the Simple Filtering Service, and is expected to encompass information such as filter status (enabled/disabled, error), statistics on the number of alerts seen and the number transmitted by the filter, and optional debugging log information.

2.10 Documentation

"Documentation" refers, essentially, to curated human-written information that is largely static. E.g., a description of the survey strategy is documentation; a Web page displaying the current progress of the survey is not.

2.10.1 Access to Observatory Documentation

ID: DMS-PRTL-REQ-0122

Specification: The Portal aspect shall provide access to Project-provided documentation on the design, construction, and operation of the LSST.

Discussion: Much of this will be just a link or links to documentation whose creation is out of the scope of DM. The DM documentation will include documentation on the code, on the pipeline processing, and on the delivered data quality. Note that documentation on the Science Platform itself will also be included.

2.10.2 Portal User Documentation

ID: DMS-PRTL-REQ-0123

Specification: The Portal aspect shall provide user-guide-style documentation on the use of the Portal.

2.10.3 Portal API Documentation

ID: DMS-PRTL-REQ-0124

Specification: The Portal aspect shall provide reference-manual-style documentation on its public network and programmatic APIs.

Discussion: This may include network APIs that allow interaction with a Portal session's state, Python APIs including wrappers for the network APIs, and JavaScript APIs for the components of the Portal application. This requirement is somewhat redundant with the general DM requirement that all code be supplied with reference documentation, but emphasizes a coherent presentation of the APIs relevant to the Portal.

2.11 General

2.11.1 Identity and Security

2.11.2 Implementation

2.11.2.1 Tolerance of Production Database Changes

ID: DMS-PRTL-REQ-0125

Specification: The Portal aspect shall be designed to facilitate accommodation of database expansion and changes and metadata extension and changes associated with the evolution of the Level 1 data, Level 2 data releases, and other planned data sources.

Discussion: The LSP needs to accommodate the database changes associated with Level 1 and Level 2 updates; this also has implications for Data Access Services - see the "Discovery and Reflection APIs" requirement - and for the availability of the data as they are released with the processing (nightly and annually). This requires stability for the APIs (web and client), and usability of the Portal across releases (and for multiple releases).

2.11.2.2 System-Busy Indication

ID: DMS-PRTL-REQ-0126

Specification: The Portal aspect shall provide a means to inform users when elements of the system are unavailable.

Discussion: This might be due to maintenance or to excessive load.

3 Notebook Aspect

3.1 Environment

3.1.1 Interactive Python Environment

ID: DMS-NB-REQ-0005

Specification: The Notebook Aspect shall provide an interactive Python environment through

both a notebook interface and via a Python interactive interpreter.

3.1.2 Unix Shell Access

ID: DMS-NB-REQ-0006

Specification: The Notebook Aspect shall provide command line access to a Unix shell with the same environment as DMS-NB-REQ-0005.

3.1.3 Pre-installed Containerized Software Releases

ID: DMS-NB-REQ-0007

Specification: Users of the Notebook Aspect shall be able to choose from a curated list of pre-built containers (including version of LSST stack) for their notebooks (and any other provided interactive environment) to execute in.

Discussion: The use of the Docker container technology is anticipated.

3.1.4 Release Deployment Latency

ID: DMS-NB-REQ-0008

Specification: It shall be possible to add a new environment (with a new version of the LSST stack) to the curated list of available execution environments in less than four hours.

3.1.5 Data Access Middleware Availability

ID: DMS-NB-REQ-0009

Specification: Users of the Notebook Aspect shall be able to make use of the LSST Python I/O middleware layer to perform data discovery, data access and any other supported functions (e.g., provenance information).

Discussion: Notably, the Data Butler is available in the Notebook Python environment, with

full access to all authorized data products available on that instance of the Science Platform.

3.1.6 Common Astronomy Package Availability

ID: DMS-NB-REQ-0010

Specification: The Notebook Aspect shall provide select standard astronomy and data analysis packages in the interactive environments.

Discussion: These may include, for example, astropy, pandas, scipy, scikit-learn, matplotlib, bokeh, and seaborn.

3.1.7 User Package Installation

ID: DMS-NB-REQ-0011

Specification: The Notebook Aspect shall have a process that allows users to add new packages to their environment

Discussion: It is intended that operations like "pip install" will be usable.

3.1.8 User Development Environment

ID: DMS-NB-REQ-0012

Specification: The Notebook Aspect environment shall permit a user to edit and build their own version of any LSST science pipeline package in their container.

Discussion: This implies the availability of both a C++ and a Python development environment.

3.1.9 Persistent User Home File Space

ID: DMS-NB-REQ-0013

Specification: The Notebook Aspect shall provide a persistent home space such that per user configuration survives shutdown and restart of the environment.

Discussion: This space appears as a home directory from Python and in the Unix shell environment. This includes things like `.bashrc`, `.pythonrc`, and user installed python libs.

3.1.10 Documentation

ID: DMS-NB-REQ-0014

Specification: The Notebook Aspect shall provide documentation of each of the constituent features as well as tutorial notebooks demonstrating the use of the Aspect.

3.1.11 New-User Onboarding

ID: DMS-NB-REQ-0015

Specification: The Notebook Aspect shall provide clear documentation on how to obtain credentials for accessing the Notebook Aspect.

3.2 Resources

3.2.1 Shared File Space

ID: DMS-NB-REQ-0016

Specification: The Notebook Aspect shall provide access to a shared read/write filesystem visible to all users of an instance of the Science Platform.

Discussion: In the LDF this is intended to be implemented as `.`. DMS-REQ-0340 means that permissions will be controlled on a variety of granularity including user and group.

3.2.2 Access to the API and Portal Aspects

ID: DMS-NB-REQ-0017

Specification: The Notebook Aspect shall be able to utilise the data access services provided by other Aspects.

Discussion: In particular, a Notebook user can use standard VO services to access LSST Data Releases.

3.2.3 User File Workspace Access

ID: DMS-NB-REQ-0018

Specification: The Notebook Aspect shall be able to access the User File Workspace available as a POSIX filesystem from within the Python kernels and shell-prompt sessions it supports.

3.2.4 VOSpace Access

ID: DMS-NB-REQ-0019

Specification: The Notebook Aspect shall be able to interact with VOSpace services available through project or external services.

Discussion: Users will be able to directly use VOSpace APIs within a Notebook. It is not yet decided whether there will be support for user-mode mounting of non-LSP VOSpace (or WebDAV) services as virtual POSIX filesystems.

3.2.5 User Database Workspace Access

ID: DMS-NB-REQ-0020

Specification: Users will be able to interact with their User Database through the Notebook Aspect to insert, delete, and control access to their tables.

Discussion: This will be possible via TAP, at least, and possibly through lower-level access.

3.2.6 Batch System Access

ID: DMS-NB-REQ-0021

Specification: The Notebook Aspect shall provide access to a batch processing system via shell access.

Discussion: This is a cluster of computers scheduled through a standard scheduler like slurm, condor, or pbs.

3.2.7 Compute and Storage Quotas

ID: DMS-NB-REQ-0022

Specification: The Notebook Aspect shall have a quota system for compute and storage authorized access via an authentication system.

3.2.8 Access to All Data Products

ID: DMS-NB-REQ-0023

Specification: An authorized user of the Notebook Aspect shall be able to access the Transformed Engineering and Facilities Database (EFD) and all other LSST released data products.

3.3 Architecture

3.3.1 Ease of Deployment

ID: DMS-NB-REQ-0024

Specification: The Notebook Aspect shall be deployable to multiple instances and contexts,

both private and public.

Discussion: Such as the Commissioning Cluster and the LDF, but also collaborator clusters, subject to the underlying resources available in the specific instance. (What level of effort? 2 days/week/month; one click deployable on a common standard platform: e.g., Kubernetes.)

3.3.2 Deployment Workload in Kubernetes

ID: DMS-NB-REQ-0025

Specification: Given a Kubernetes cluster with a configuration meeting a documented standard set of specifications, it shall take an engineer with admin rights no more than 2 days to deploy the Notebook Aspect in that context.

Discussion: The specification is expected to constrain factors such as software versions for Kubernetes and related packages, available storage, a shared file system, and an authentication system.

3.3.3 System Health Monitoring

ID: DMS-NB-REQ-0026

Specification: The Notebook Aspect shall provide a service health microservice and a dynamic web page hostable on separate resources that provides a view of the health status.

3.4 Display Data

3.4.1 Image Visualization

ID: DMS-NB-REQ-0032

Specification: The Notebook Aspect shall provide a tool for displaying image like datasets produced by LSST stack tools.

Discussion: This requirement could be satisfied simply by inclusion of `afw.display` in the pre-

installed stack. However, it is anticipated that we will also provide a way to use Portal Aspect Javascript components in JupyterLab.

3.4.2 Scientific Plotting

ID: DMS-NB-REQ-0033

Specification: The Notebook Aspect shall provide common plotting methods: scatter plots, raster images, histograms, 2D histograms, contours, line traces, polygons, compositions of these (contours on scatter plots), density images

Discussion: This requirement could be satisfied simply by ensuring that matplotlib is usable within JupyterLab.

3.4.3 Visualization Linkage

ID: DMS-NB-REQ-0034

Specification: The Notebook Aspect shall provide "drill down" functionality in plots: brushing and linking between plots, interactively discover metadata about particular points, drill down to imaging from measurements

Discussion: Metadata can be visit properties for a measurement, git commits, etc. (gpdf is concerned that this is too vaguely defined to be verifiable. Should this be in a design document instead?)

3.4.4 Visualization Interactivity

ID: DMS-NB-REQ-0035

Specification: The Notebook Aspect shall provide interactive plots for certain visualizations: Linked axes on multiple plots, zoom, pan, data point selection

Discussion: (gpdf is concerned that this is too vaguely defined to be verifiable. Should this be in a design document instead?)

3.4.5 Visualization Scaling

ID: DMS-NB-REQ-0036

Specification: The Notebook Aspect shall provide interactive plots that scale to include at least 1E6 datapoints.

Discussion: This may be done through an adaptive refinement scheme like datashader.

3.5 Interaction with Portal

3.5.1 Access to Portal-Initiated Queries

ID: DMS-NB-REQ-0029

Specification: A user of the Notebook Aspect shall have access to search queries they performed in the Portal Aspect.

Discussion: This depends on underlying services from the API Aspect.

3.5.2 Access to Portal Visualization API

ID: DMS-NB-REQ-0030

Specification: The Notebook Aspect shall provide a mechanism for "pushing" specific types of data to the Portal API.

Discussion: For instance, this allows a user to plot a catalog of coordinates over an image display using the Portal's Firefly components. This is supported by DMS-PRTL-REQ-0115 on the Portal side.

3.5.3 Notebook-Launching Interface

ID: DMS-NB-REQ-0031

Specification: The Notebook Aspect shall provide a means to trigger the opening of a notebook with access to the results of a query performed in the Portal.

Discussion: This is intended to permit a Portal user to perform a query and then quickly obtain a Notebook session with that data available for further analysis. The UI element for this might be in either the Portal or Notebook system Uis, depending on implementation issues.

3.6 General

3.6.1 Identity and Security

3.6.1.1 Secure Protocol

ID: DMS-NB-REQ-0001

Specification: The Notebook Aspect shall be accessible through an HTTPS endpoint.

3.6.1.2 Authentication and Authorization

ID: DMS-NB-REQ-0002

Specification: The Notebook Aspect shall authenticate users for the purpose of establishing authorized use and only permit access to authenticated users using the LSST Data Facility authentication and authorisation service.

3.6.1.3 Secure Implementation

ID: DMS-NB-REQ-0003

Specification: The Notebook Aspect shall prevent users from circumventing authorisation controls.

Discussion: The Notebook Aspect relies on other services, such as authentication, file system

permissions etc to prevent access to unauthorized data. It should not be possible for a user to spoof another user in a way that permits access to unauthorized data

3.6.2 Implementation

3.6.2.1 IPV6 Access

ID: DMS-NB-REQ-0004

Specification: Access to the Notebook Aspect shall support access using IPV6 protocols.

4 API (Data Access) Aspect

4.1 Discover Data

4.1.1 Use of CAOM2

ID: DMS-API-REQ-0021

Specification: The API Aspect Web APIs shall present image and visit metadata organized in accordance with the CAOM2 data model.

4.1.2 Access to Image and Visit Metadata

ID: DMS-API-REQ-0022

Specification: The API Aspect shall provide for retrieval of image and visit metadata via TAP ADQL queries.

4.1.3 Access to Catalog Data Products

ID: DMS-API-REQ-0023

Specification: The API Aspect shall provide for retrieval of all Prompt and Data Release cata-

log data (per LSE-163) via TAP ADQL queries.

4.1.4 Access to Observatory Metadata

ID: DMS-API-REQ-0024

Specification: The API Aspect shall provide for retrieval of observatory metadata (including the Transformed EFD) via TAP ADQL queries.

4.1.5 Enforcement of Information Classification

ID: DMS-API-REQ-0025

Specification: The API Aspect shall **NOT** allow access to Sensitive or Highly Sensitive (per LPM-122) observatory metadata.

Discussion: Information classified as "Internal" should only be provided to project staff.

4.1.6 Access to Reference Catalogs

ID: DMS-API-REQ-0026

Specification: The API Aspect shall provide for retrieval of all reference catalog data via TAP ADQL queries. For the purposes of this requirement a "reference catalog" is an externally sourced catalog used during data production activities.

Discussion: This requirement is not to imply *linkage* of reference catalogs/images, but should be interpreted to mean that the reference catalogs that we use will also be queryable.

4.1.7 Access to Virtual Data Products

ID: DMS-API-REQ-0027

Specification: The API Aspect shall provide services to initiate regeneration of, and facilitate

retrieval of, virtual data products on demand.

Discussion: For image data products, this would likely be provided via the SODA endpoint.

4.1.8 Access to Image Data in FITS Format

ID: DMS-API-REQ-0028

Specification: The API Aspect shall deliver image data in FITS format, and **MAY** deliver images in additional formats.

4.1.9 Multiple Data Releases

ID: DMS-API-REQ-0029

Specification: The API Aspect Web APIs shall provide unambiguous access to data products and metadata from more than one Data Release simultaneously

Discussion: The requirement is explicitly silent on the question of whether data from multiple releases will be available from a single endpoint.

4.1.10 Catalog Metadata Service

ID: DMS-API-REQ-0030

Specification: The API Aspect shall provide complete metadata for all tables within each data release, including per-column a description, IVOA UCD when appropriate, unit when appropriate, and any relationship with other columns

4.2 Query Data

4.2.1 TAP Service - Tabular Queries

The API Aspect includes a TAP service as the primary means of access to tabular data.

4.2.1.1 TAP Service for Tabular Queries

ID: DMS-API-REQ-0006

Specification: The API Aspect Web APIs shall include an endpoint conforming to IVOA TAP 1.1 for the purpose of accessing tabularly structured data.

4.2.1.2 Synchronous TAP Support

ID: DMS-API-REQ-0007

Specification: The API Aspect TAP endpoint shall support synchronous queries as described by the IVOA TAP 1.1 specification.

Discussion: Synchronous queries are primarily expected to be used for small results. It is TBD whether this mode will be suitable for use with shared-scan queries.

4.2.1.3 Asynchronous TAP Support

ID: DMS-API-REQ-0008

Specification: The API Aspect TAP endpoint shall support asynchronous queries as described by the IVOA TAP 1.1 specification.

Discussion: Asynchronous queries are expected to be the primary means for carrying out user-driven queries from the Portal aspect (so that the query is assigned an ID that enables multiple retrievals of its results, e.g., in both the Portal and Notebook Aspects). (The Portal Aspect implementation will still use synchronous queries for internal functions such as retrieving metadata needed for page configuration, etc.) Asynchronous queries will also be used across all aspects for queries where large result sets and/or long run times are expected.

4.2.1.4 ADQL Support

ID: DMS-API-REQ-0009

Specification: The API Aspect TAP endpoint shall support IVOA ADQL 2.1 as a query language, **BUT** supported query syntax for database targets **MAY** be limited by practical considerations of individual underlying database technologies

4.2.2 SIA Service - Image Availability

4.2.2.1 SIA Service for Image Availability

ID: DMS-API-REQ-0016

Specification: The API Aspect Web APIs shall include an endpoint conforming to IVOA SIA V2 for the purpose of locating available images

4.2.3 SODA Service - Image Data Retrieval

The retrieval of actual image pixel data is provided by an IVOA SODA service. This service provides both for the retrieval of as-archived images and for the retrieval of post-processed image data such as cutouts.

It is expected that some form of mosaicing service will be provided, but a specific requirement is still to be written.

4.2.3.1 SODA Service for Image Data

ID: DMS-API-REQ-0017

Specification: The API Aspect Web APIs shall include an endpoint conforming to IVOA SODA 1.0 for the purpose of retrieving image data.

4.2.3.2 Cutout Service

ID: DMS-API-REQ-0018

Specification: The API Aspect SODA endpoint shall support performing cutouts on all released image data types, **BUT** supported filter predicates **MAY** exclude POLYGON

4.2.4 Query Persistence

Requirements in this section apply to the persistence of query texts and the limited persistence of query results.

4.2.4.1 Query History Retrieval

ID: DMS-API-REQ-0038

Specification: The API aspect shall provide interfaces for retrieving the history of queries for a user.

Discussion: This capability is essential for cross-Aspect linking, where a query is created in one Aspect and accessed, or re-executed, in another.

4.2.4.2 Cached Query Result Retrieval

ID: DMS-API-REQ-0039

Specification: The API Aspect shall provide for the caching of results of queries for a limited time, and their retrieval based on information from the query history or on query identifiers previously returned from asynchronous query services.

Discussion: Caching is subject to resource constraints. The system may use a combination of a central buffer and quota from the user's Workspace to implement caching; the details are still under design.

4.2.4.3 Query Specification Retrieval

ID: DMS-API-REQ-0040

Specification: The API Aspect shall provide interfaces that return an artifact containing a complete specification for a query, and that permit that artifact to be used at a later time to re-execute the same query.

Discussion: The results of re-executing the same query depend on the database(s) being queried. For Level 2 (Data Releases) the content of a data release is nominally frozen upon release and the same query repeated later should always return the same result. (We assume that if a correctable error is found in a data release after its release it will only be repaired, if at all, by adding additional database(s) or table(s) with the corrected data, or just by providing correction recipes that could be applied to the results of a query.) For the continuously updated Level 1 database(s) update times will be recorded that will permit queries to be repeated precisely, or, optionally, with new data taken into account. For user databases no guarantees can be made.

4.2.5 Python Interface

4.2.5.1 Butler Interface to Data Products

ID: DMS-API-REQ-0034

Specification: The API Aspect shall provide connection between the Data Butler (Generation 3) instances within notebooks hosted in a LDF instance and backend file system, database, and object data stores within that same LDF instance, for the purpose of allowing notebook aspect users to access data release data products and user generated data products as Python objects

Discussion: See LDM-556.

4.3 User Storage

Requirements for the User Database Workspace are still being developed.

4.3.1 VOSpace Service

ID: DMS-API-REQ-0019

Specification: The API Aspect Web APIs shall include an endpoint conforming to IVOA VOSpace 2.0 for the purpose of persistence and retrieval of user-generated file-oriented data products in the User Workspace defined in DMS-LSP-REQ-0011.

4.3.2 WebDAV Service

ID: DMS-API-REQ-0020

Specification: The API Aspect Web APIs shall include an endpoint conforming to WebDAV for the purpose of persistence and retrieval of user-generated file-oriented data products in the User Workspace defined in DMS-LSP-REQ-0011.

4.4 Transfer Data

4.4.1 Tabular Data Output

Requirements in this section apply to the return of tabular data from API Aspect services.

4.4.1.1 VOTable Output for TAP

ID: DMS-API-REQ-0010

Specification: The API Aspect TAP endpoint shall support IVOA VOTable 1.3 as an available output format

4.4.1.2 VOTable TABLEDATA Payload

ID: DMS-API-REQ-0011

Specification: API Aspect services that support returning results in VOTable format shall support the return of a VOTable data payload in the XML-based TABLEDATA serialization.

Discussion: This payload form is not intended for large tables, but is provided for compat-

ibility with community tools. The API Aspect TAP service may place an upper bound on the size of a table that can be returned in this format.

4.4.1.3 VOTable BINARY2 Payload

ID: DMS-API-REQ-0012

Specification: API Aspect services that support returning results in VOTable format shall support the return of a VOTable data payload in the BINARY2 serialization.

Discussion: This payload form supports larger tabular results but is most likely still not optimal for the largest queries. The API Aspect team is still investigating additional options, including the use of the FITS serialization, as well as the use of non-VOTable formats, for efficient and metadata-rich tabular results.

4.4.1.4 JSON Output for TAP

ID: DMS-API-REQ-0013

Specification: The API Aspect TAP endpoint shall support JSON as an alternative available output format

4.4.1.5 CSV Output for TAP

ID: DMS-API-REQ-0014

Specification: The API Aspect TAP endpoint shall support CSV as and alternative available output format. This output format is not required to meet requirements otherwise in force on the return of table and column metadata.

Discussion: The CSV format inherently is unsuitable for returning rich metadata with a table.

4.4.1.6 SQLite Output for TAP

ID: DMS-API-REQ-0015

Specification: The API Aspect TAP endpoint SHOULD support SQLite as an alternative available output format

Discussion: The mandatory status of this requirement should be settled one way or the other as soon as possible. It is a candidate for a efficient and readily consumed format for large results with rich metadata.

4.4.2 Workspace Interaction

4.4.2.1 Tabular Result Download to Workspace

ID: DMS-API-REQ-0031

Specification: The API Aspect shall provide a capability for users to save their query results as VOTables in their allocated VOSpace, subject to limitations of a resource quota system

Discussion: Or any of the other supported formats?

4.4.2.2 Tabular Upload to Workspace

ID: DMS-API-REQ-0032

Specification: The API Aspect shall provide a capability for users to upload catalog data products (formatted as VOTables) residing within their allocated VOSpace, such that the catalog products after upload may be joined in queries against data release catalog products, subject to limitations of a resource quota system

Discussion: Or any of the other supported formats?

4.4.2.3 Deletion from Workspace

ID: DMS-API-REQ-0033

Specification: The API Aspect shall provide a capability for users to drop previously uploaded user catalog data products

4.5 Performance

Performance requirements are still being finalized.

4.6 Control and Management

4.7 Documentation

4.8 General

4.8.1 Identity and Security

4.8.1.1 Secure Protocols

ID: DMS-API-REQ-0001

Specification: The API Aspect Web APIs shall be accessible through HTTPS endpoints.

4.8.1.2 Authentication

ID: DMS-API-REQ-0003

Specification: The API Aspect Web APIs shall accept authenticated requests for the purpose of establishing user identity.

4.8.1.3 Authorization

ID: DMS-API-REQ-0004

Specification: The API Aspect Web APIs shall interact with project authorization infrastruc-

ture for the purpose of establishing authorized use.

4.8.1.4 Secure Implementation

ID: DMS-API-REQ-0005

Specification: The API Aspect Web APIs shall prevent users from circumventing authorization controls.

4.8.2 Implementation

4.8.2.1 Containerized Deployment

ID: DMS-API-REQ-0035

Specification: The API Aspect services shall be delivered as containerized applications.

4.8.2.2 Result Compression

ID: DMS-API-REQ-0002

Specification: The API Aspect Web APIs shall support gzip HTTP content-encoding for the purpose of returning compressed data.

4.8.2.3 Upgradability

ID: DMS-API-REQ-0036

Specification: The API Aspect service software shall be upgradable in place with minimal user downtime.

4.8.2.4 Logging and Monitoring

ID: DMS-API-REQ-0037

Specification: The API Aspect services shall provide logging and monitoring capabilities for the purpose of supporting service operators

5 Proposed Requirements

This section includes requirements that are still under review and have not been baselined.

5.1 Large Result Set Support

ID: DMS-NB-REQ-0027

Specification: The Notebook Aspect, when deployed in an LSST Data Access Center, shall be sized to support 20 simultaneous users and be able to provide visualization for 10 billion objects and result sets of up to 6GB.

Discussion: This requirement does not apply to all instances, as it is dependent on resource allocations.

5.2 Moderate Result Set Support

ID: DMS-NB-REQ-0028

Specification: The Notebook Aspect, when deployed in an LSST Data Access Center, shall be sized to support 50 simultaneous users and be able to provide visualization for 10 million objects and result sets of up to 0.1GB.

Discussion: This requirement does not apply to all instances, as it is dependent on resource allocations.

References

- [1] **[LDM-555]**, Becla, J., 2017, *Data Management Database Requirements*, LDM-555, URL <https://ls.st/LDM-555>
- [2] **[LSE-61]**, Dubois-Felsmann, G., Jenness, T., 2018, *LSST Data Management Subsystem Requirements*, LSE-61, URL <https://ls.st/LSE-61>
- [3] **[LDM-556]**, Dubois-Felsmann, G., Jenness, T., Bosch, J., et al., 2018, *Data Management Middleware Requirements*, LDM-556, URL <https://ls.st/LDM-556>
- [4] **[LSE-163]**, Jurić, M., et al., 2017, *LSST Data Products Definition Document*, LSE-163, URL <https://ls.st/LSE-163>
- [5] **[LPM-122]**, Petravick, D., 2015, *LSST Information Classification Policy*, LPM-122, URL <https://ls.st/LPM-122>