



Vera C. Rubin Observatory
Systems Engineering

Project-wide Documentation Proposal for Rubin Observatory Operations

(Report from the Rubin Observatory Documentation Working Group)

Chuck Claver (chair), David Cabrera (co-Chair), Rob McKercher (co-chair), John Andrew, Diane Hascall, Patrick Ingraham, Tony Johnson, Kristen Metzger, Austin Roberts, Jonathan Sick, Matthew Rumore.

SITCOMTN-014

Latest Revision: 2022-05-03

DRAFT



Abstract

This technical note is a report to recommend a future state for Rubin Observatory Operations documentation by the Project-wide Documentation Working Group after an initial documentation survey (<https://sitcomtn-012.lsst.io/>) and further analysis and discussion. As a final deliverable for the charge described in LSE-489.lsst.io, this proposal includes requirements, strategy, methodology and organization for the technical documentation package needed to establish construction completeness and operational readiness from the Construction Project and a method of sustained maintenance throughout Operations. The presented implementation plan for transitioning includes an estimated level of effort, resources, scheduling and risks if not implemented.

Change Record

Version	Date	Description	Owner name
0.0	2021-06-26	First draft of preliminary content.	Chuck Claver
0.1	2021-08-30	Input to Documentation Portal.	Jonathan Sick
0.2	2021-09-14	Preliminary content, Portal architecture and effort.	Chuck Claver, Jonathan Sick
0.3	XXXX-XX-XX	Incorporate draft from Confluence pages and shared files.	David Cabrera, Matthew Rumore

Document source location: <https://github.com/lsst-sitcom/sitcomtn-014>

Contents

1 Introduction	1
2 Documentation Views	1
2.1 Product View	3
2.2 Access View	3
2.3 Storage View	3
3 Implementation Plan	3
3.1 Primary Documentation Sources	4
3.2 The Documentation Portal	4
3.2.1 Design requirements	4
3.2.2 Technical architecture	5
3.2.3 Support for multiple views	8
3.2.4 Information security	9
3.2.5 Effort estimate	10
3.3 The Path Forward with DocuShare	10
3.4 The Path Forward with PDWM Works	12
3.5 Required Resources & Subsystem Responsibilities	12
3.6 schedule	13
4 Transition Plan and Workflow	13
5 Risk Assessment	13
A References	14
B Acronyms	14

Project-wide Documentation Proposal for Rubin Observatory Operations

1 Introduction

Put basic description of approach and report structure here.

This report include details on the following topics:

- DocumentationViews
 - ProductView
 - AccessView
 - StorageView
- Implementation Plan
 - PrimarySources
 - DocumentationPortal
 - DocuSharePath
 - PDMWorksPath
 - Resources & Responsibilities
 - Schedule
 - TransitionPlan
- RiskAssessment

2 Documentation Views

The section describe three different view of the documentation content:

We have developed 3 views to help organized the future (e.g. delivered) state of the Rubin Construction Project documentation. These are:

- The Product View;
- The Access View; and
- The Storage View.

Within each view, particularly the Access view, we have subcategorized the view into access for Users/operator, Engineers/Technical and Scientists. The Product and Access view derive their content from the Storage view. The Storage view is a means to determine where the source of "truth" lies.

The "view" are represented by "trees" provide structure for different areas of the project involved in the Rubin Observatory Documentation Process, which involves the Information Retrieval, distribution, and archiving to the future operation phase of the Vera C Rubin Observatory, to provide a robust and orderly structure to future users, to accomplish their expectations on searching and reference of documentation and at the same time, provide a structure to include the new documentation, to register in detail, the information that is being produced during the Operational Phase of the telescope Subsystems and corresponding facilities.

The different trees shall embrace the different steps that are involved in a General and wide Documentation Process Requirement, in which the different Managing Areas of the project will provide and develop their corresponding and common Operation Documentation with an approach on the users' expectation to a reliable searching process, through all the Rubin Observatory associated groups.

It is intended that the implementation of these views will constitute a deliverable from the construction project to the operations team.

It is also important to note that no content current in storage structure (see the Documentation Inventory assessment in SITCOMTN-012) will be lost. This content will be maintained as is, but not necessary fully incorporated into the construct of the final delivered technical documentation packed as described below.

2.1 Product View

Describe here the properties of the product View and provide diagrams showing the proposed structure for this view.

2.2 Access View

Describe here the properties of the Access View and provide diagrams showing the proposed structure for this view.

2.3 Storage View

Describe here the properties of the Storage View and provide diagrams showing the proposed structure for this view.

3 Implementation Plan

Add here an overview of the proposed implementation plan including the following specific areas of implementation:

- Primary Sources
- Documentation Portal
- DocuShare
- PDM Works
- Resources and Responsibilities
- Schedule
- Transition Workflow

3.1 Primary Documentation Sources

Describe here the proposed location of the primary documentation sources. These are a subset of the identified source listed in SITCOMTN-012.

3.2 The Documentation Portal

The working group recommends the creation of a *documentation portal* web application as a means of making documentation content discoverable and accessible to Rubin Observatory staff. The documentation portal provides interfaces for both searching (based on content and metadata) and browsing (based on hierarchical categorization) of documentation resources. Documentation does not reside within the portal itself. Rather, the portal's objective is to efficiently link the user to the document where it is stored in any of the observatory's adopted storage platforms (Section 3.1). The portal discussed here is based upon the www.lsst.io website, which provides a search and browsing interface for the Rubin Observatory's public-facing technical documentation. The new portal will be accessible only those with Rubin Observatory staff credentials, and will be purpose-built for observatory and survey operations. This section describes the design principles, technical architecture, security model, and cost estimate of the Rubin Observatory documentation portal.

3.2.1 Design requirements

The design requirements of the documentation portal reflect the recommendations made by the Documentation Working Group elsewhere in this report:

1. The role of the documentation portal is to link to documentation resources. The portal itself does not host the content itself, or provide user interfaces for creating and maintaining new versions of documentation content. This requirement reflects the working group's recommendation that documentation content should be hosted on a select set of platforms that are idiomatic for the content and the teams that work with that content (Section 3.1).
2. The documentation portal must provide equal support for content stored in any of the working group's recommended storage platforms.
3. The documentation portal must be capable of supporting several hierarchical browsing

schemes for accessing content based on different organizational views of the documentation (Section 2).

4. The documentation portal should automatically update and sort documentation content, to the greatest extent possible. In other words, curators should only need to maintain documentation in the document's storage platform, without any administrative action through the documentation portal's UI. A consequence of this requirement is that the documentation portal should not persist information about a document that is not available from the document's own storage platform.
5. The documentation portal should be secured so that it is only accessible to users with Rubin Observatory staff credentials.
6. The documentation portal should not maintain fine-grained access control for specific documents or categories of documents. For secured documents, the portal relies upon the security mechanisms of the document's own storage platform. The portal should also reduce its metadata storage of confidential documents to ensure that content cannot be inferred from a search, for example.

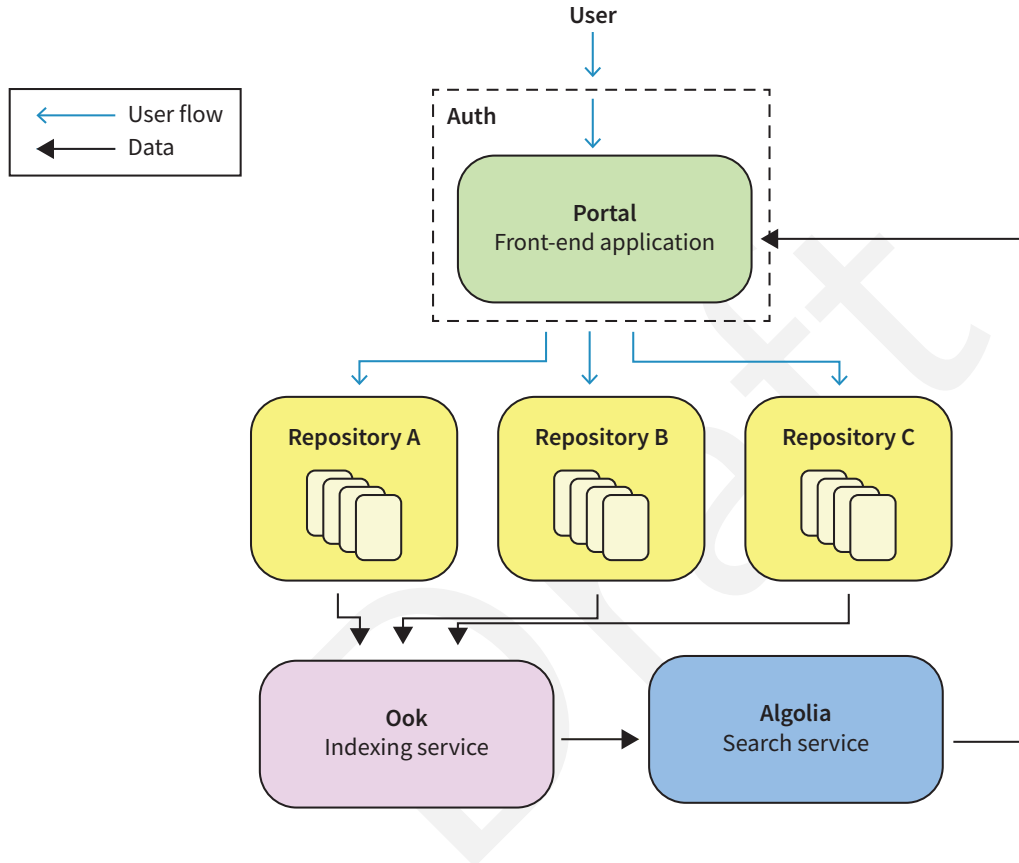
3.2.2 Technical architecture

The architecture described here is based upon that which is already put into production with the www.lsst.io portal for public-facing documentation. Starting from this working archetype relieves a great deal of technical risk and development from the new portal's implementation. Both portals share the use of Algolia as a search backend and Ook as a content indexing service. The documentation portal discussed here will have an independent front-end to support the specific access views recommended by the working group. The documentation portal will also use a separate instance of the Algolia database to eliminate any risks associated with leaking internal documentation to the public-facing www.lsst.io portal.

Figure 1 depicts the components of the documentation portal.

3.2.2.1 Algolia The core function of the portal is to enable access to documentation through browsing and search. To implement this, the portal needs a back-end service that contains metadata about Rubin's documentation holdings and provides interfaces to access and query that metadata from the front-end (website). This search database and interface could be

FIGURE 1: Architecture of the documentation portal. Users find documents on the portal application, which in turn provides links into the original documentation repositories. Data in the portal application is supplied by the Algolia search service, which in turn gets its metadata from the original documents in their repositories via the Ook indexing service.



made for “free” with entirely open-source components such as Elasticsearch and in-house web service. However, a search database and service are sufficiently generic that we cannot add value by making it in-house, and in fact developing, tuning, and operating this service, would costly in terms of labor. For the www.lsst.io public documentation portal, we opted to use Algolia and recommend that we make the same choice for the observatory’s documentation portal.

Although www.lsst.io is currently operating on a free open-source license of Algolia, the observatory’s internal documentation portal would be an entirely paid license. Algolia prices based on record counts and request rates. In operating www.lsst.io, we found record count to be the limiting factor. At the moment, 1000 records costs \$1 per month. 1,000,000 records would cost \$850 per month with volume discounts. For reference, the www.lsst.io service

currently uses 110,000 records to host all technical notes and change-controlled documents with drafts hosted on lsst.io.

Note that a single document is composed of potentially many records in Algolia. To optimize full-text search, we break a document into smaller records, generally across section boundaries. Our current algorithms generally produce small record, so it is possible to reduce costs by tuning how we segment content into Algolia records. Furthermore, each sorting option requires a separate pre-sort index. Sorting documents by date, and document number, in addition to relevance, consumes three times as many records as only sorting by relevance.

3.2.2.2 Ook The Ook service is responsible for continuously indexing content into Algolia. Whereas we chose Algolia to provide a turn-key search database service, for www.lsst.io we chose to build the indexing service in-house to have complete control over how documents are indexed, and what metadata is associated with each document. For example, LaTeX-based documents are indexed based on metadata exported from the Lander PDF landing page generator (also developed in-house), which in turn parses LaTeX syntax in the document source to access metadata such as titles, authors, and so on. The configurability of Ook is beneficial to indexing other types of highly specialized documentation.

The key design principle of Ook is that metadata is extracted from the document as it appears in its repository, rather than requiring direct human interaction with Ook or Algolia to curate the data. This allows Ook to scale well across an organization as large and varied as Rubin because individual teams manage documents as they already do in the repositories they are already familiar with.

Ook is built such that new content types can be added by writing additional Python-based workflows for each content type. Ook itself provides utilities for queuing ingests, converting content and formatting data for Algolia, and working with the Algolia service itself.

Ook indexing operations can be triggered several different ways. For example, the lsst.io service publishes messages to a Kafka cluster whenever documentation is published on that platform; Ook subscribes to those messages and queues indexing workflows. Ook indexing operations can also be scheduled through an HTTP API. Generally, the goal is to trigger indexing operations automatically whenever the source material changes.

The existing Ook indexing workflows work by downloading content from websites and web services (HTTP APIs). If content is not easily accessible, it would be possible to develop an alternative workflow, such as submitting copies of the document directly to Ook for indexing. Some document repositories may offer online access but have hard-to-use APIs (such as DocuShare). These difficulties can be worked around, for example by emulating a web browser to download content and metadata, but at the cost of more fragile indexing workflows.

Ook is currently operated as a Kubernetes application in the Google Cloud. This arrangement is ideal for minimizing operations cost, and providing convenient scaling.

3.2.2.3 Front-end application The front-end web application is how users (Rubin Observatory staff) find documentation. The web application does not provide the document itself—instead, the web application provides a search result card that the user can click on to access the document in its original repository. The Algolia services provides all browsing and search functionality; the front-end application provides the user interface over top of Algolia.

In addition to providing a link to the original document, the portal can also provide an immediate view of a document's metadata. Although the front-end application can show a basic view of a document's metadata based on metadata common to all records, the website can be developed to show additional metadata for different types of documents.

For www.lsst.io, we built the site as a React JavaScript application. This allowed us to use and customize the pre-made widgets provided by Algolia for building the user interface.

The front-end application will be accessible only to users who log in. The simplest way to approach this is by putting the application behind a VPN so that the application is completely separate from security concerns. Another approach would be to place the application behind an OAuth proxy to provide a slightly better user experience.

3.2.3 Support for multiple views

In Section 2, the documentation working group outlines several views for hierarchically arranging documents trees. These views correspond to navigational structures in the front-end application. Although the front-end application code is generally “aware” of the different trees, individual documents are placed in the tree on the basis of metadata in their Algolia records,

so that the Algolia service can pre-sort and filter documents into the trees. Since Ook supplies metadata to Algolia, and Ook in turn leverages metadata native to the document and the document's repository, the responsibility for curating documents into different views is the responsibility of individuals managing documents in each repository.

3.2.4 Information security

Documentation has multiple types of security concerns, such as control over who and how documents are updated, control over who can access documents, and ensuring the long term integrity and preservation of information. Since the documentation portal is not the canonical repository for any documentation, the portal is not involved in controlling document updates and preservation. The portal's key security concern is access control.

The portal is designed to only provide authentication-based access control. Any Rubin staff member with credentials can access any metadata records contained within the documentation portal. Once a user selects a document to view, they are forwarded to that documentation repository and must authenticate with that repository and be subject to its access control rules.

However, the metadata contained in the documentation portal can be potentially rich, even including the full-text content of a document to enable search functionality. It is conceivable that some of this metadata may not be appropriate for observatory-wide access (such as information with export controls, for example). In these cases, the most realistic approach to preserving strict access controls in these situations is to limit what metadata is available. In order of strictness, the following approaches can be used:

1. Omit full-text content of a document from Algolia records.
2. Limit or obfuscate other metadata (such as titles) in the Algolia records.
3. Omit the individual documents altogether from Algolia and instead link to a documentation landing page hosted by the secure document repository itself.

In all cases, controlling how documents are indexed is done by configuring the indexing service, Ook.

3.2.5 Effort estimate

The key tasks for implementing this plan are:

1. Building additional content indexing workflows into Ook and updating the metadata schema to accommodate all use cases.
2. Build the documentation portal application.
3. Curate documents so that they are uniformly present in their document repositories and have any metadata that is expected by the Ook indexing workflows.

The last task category can be distributed to teams working in each document repository.

To provide a sense of the scale of the first two tasks, the implementation of the www.lsst.io portal is a useful reference. Building the first version of Ook, which indexes LaTeX documents and Sphinx-based technical notes in lsst.io required 4 weeks of work. Building the front-end website required 6 weeks.

Since that original implementation, we have gained more experience building React web applications and working with Algolia, so we can expect less than 4 weeks of work to build out the internal web application. Although Ook can be expanded as is, the key uncertainty is in the number and complexity of the additional indexing pipelines that need to be built for additional document types and repositories. Generally, though, an internal documentation portal can likely be stood up within 1 to 3 months, with a potential long tail of effort to continue to add support for additional content types.

3.3 The Path Forward with DocuShare

The working group recommends that Rubin continue to use DocuShare in Operations; further, we recommend that Operations use the DocuShare instance currently in use by Construction with the Archive Server add-on enabled. The working group believes the continuity afforded by doing so is important and useful. DocuShare's version control, version history, permissions and co-location capabilities are valuable tools for effective document management and should be maintained. Additionally, continuing DocuShare use will ease the transition from

Construction to Operations by providing a single shared platform, thereby avoiding compatibility issues, eliminating the need to migrate documents, preventing loss of documentation and avoiding confusion by retaining legacy handles and references. Enabling the Archive Server will allow Operations to structure and manage DocuShare according to its needs while retaining access to all documentation accumulated during Construction. The working group believes the recommendation's benefits outweigh any required new and/or continuing costs and labor resources.

The recommendation is consistent with NOIRLab document management, which has selected DocuShare as the lab's document repository. By also using DocuShare, Rubin Operations' repository will be compatible with the knowledge and skills held by the business unit responsible for NOIRLab document management. While Rubin Operations has the option to use NOIRLab's DocuShare, the working group believes that to be less advantageous than retaining Rubin Construction's DocuShare instance. Switching to NOIRLab's DocuShare would result in loss of legacy document handles and would require labor-intensive and time-consuming document migration. Continuing to use Rubin DocuShare will retain documents' version histories and allow persistence of handles for foundation documents, many of which are known within the project and community as much by their handles as they are by their titles. Lastly, each instance of DocuShare has a 2,000,000 document limit. If the Construction project DocuShare instance is retained, the entire document limit is available to Rubin. If NOIRLab's DocuShare instance is used, Rubin would have to share the document limit with other lab units and programs.

However, the working group recognizes that simply continuing with the Rubin DocuShare as-is fails to address long-standing issues of clutter, inadequate metadata, and poor classification; therefore, we recommend enabling the Archive Server add-on. The add-on provides a second server into which documents whose lifecycles have ended can be sequestered, freeing the active server to be structured and managed according to best practices and Operations' needs. While archived documents no longer reside in the active server and are not represented in its directory structure, they are viewable and searchable and can be restored to the active server if necessary. The two servers are connected and can be interacted with using a single user interface. Under this model, Construction project content would be moved to the archive server, and the active server would contain only Operations and Operations-relevant Construction content following an agreed-upon directory structure.

Further, to mitigate the risk of replicating the clutter and less-than-optimal classification of

Construction-era DocuShare, the working group recommends Rubin adopt a more formal workflow for how and by whom DocuShare content is created, formatted, and updated. This workflow would involve something similar to NOIRLab's steward concept. NOIRLab's document management plan expects each business unit or program to designate a person or persons to lead the group's documentation activities in DocuShare. These stewards are to be highly trained so they can perform DocuShare functions, provide guidance to others, enforce standards, and correct errors. The idea is to have fewer but better-trained users with full DocuShare permissions and functionality. While this will result in a less distributed workload, consistency and quality will improve, leading to greater tool utility. The Operations departments should follow a similar approach, and as much as possible, the workflow should occur within DocuShare.

Retaining the current Rubin DocuShare instance and enabling the Archive Server add-on requires both new and continuing costs and labor. The calendar year 2020 renewal cost for Rubin's current DocuShare instance was \$22,000, and Rubin's IT resources are expended to maintain the service, which is hosted on a physical server in Tucson. The archive server add-on carries a one-time cost of \$7,500 plus an additional \$1,400 per year for support, according to a December 2020 quote from Xerox. Additionally, IT resources will be needed to configure the add-on, and some training likely will be required. Despite these considerations, the working group believes the benefits outweigh the costs, and the required resources are less burdensome than those that would be needed to adopt another repository.

A more detailed analysis of pro and cons is available in DocuShare Options Trade Study for the Documentation Working Group (Document-36788).

3.4 The Path Forward with PDWM Works

Describe the maintenance plan of the existing CAD models in PDWM works.

3.5 Required Resources & Subsystem Responsibilities

Description of the estimated resources need to carry out this implementation plan and what the roles and responsibilities are for each subsystem in this context.

3.6 schedule

Outline the schedule needed for implementation to meet the objective delivering a coherent technical document packed at the end of the Project Construction effort.

4 Transition Plan and Workflow

Describe the process for transition from the current state to the future state.

5 Risk Assessment

Describe the risk exposure if various part of the whole of this implementation is not conducted.

A References

B Acronyms

Acronym	Description
API	Application Programming Interface
HTTP	HyperText Transfer Protocol
LaTeX	(Leslie) Lamport TeX (document markup language and document preparation system)
PDF	Portable Document Format
IT	Information Technology
SE	System Engineering
UI	User Interface
VPN	virtual private network