



Vera C. Rubin Observatory
Rubin Observatory Operations


Rubin Data and Information Security Plan

William O'Mullane and Russ Allbery and Richard Dubois and Kian-Tat
Lim

RTN-030

Latest Revision: 2025-03-07

DRAFT



Abstract

This document covers data security and recovery, threats and mechanisms for dealing with them. Where possible it will refer to other documents for details. This replaces the construction era LDM-324 for operations.

Draft

Change Record

Version	Date	Description	Owner name
0.1	2022-09-09	draft	William O'Mullane
0.2	2023-11-28	Chile disaster recovery, DP is now DM, incident response	William O'Mullane
0.3	2024-06-28	SLAC disaster recovery, UK and French security policies, trace to NIST now in RTN-082, tidy and issue for review	William O'Mullane
1.0	2024-07-09	Issue one remove draft, final action complete. CCB approved CCB-4	William O'Mullane
1.1	2024-07-14	Clarify Alert Vetting - change approved by director under CCB-4	William O'Mullane
1.2	2025-03-07	Comments from SLAC DM-46258	William O'Mullane, KT Lim

Document source location: <https://github.com/lst/rtn-030>

Contents

1 Introduction	1
2 Information types in Rubin Observatory	2
2.1 Presence of controlled information	2
2.2 Controlled Catalog	2
2.3 Embargoed Data	2
2.4 Unembargoed Data	3
2.5 Unreleased Data	3
3 Cyber incident response	3
4 System description	4
4.1 Data verification and quality assurance	4
4.2 Data processing	4
4.2.1 Prompt Processing	4
4.2.2 Data Release Processing	5
4.3 Data access	5
4.4 Rubin Director’s Office	5
5 Data Management system architecture	6
5.1 US Data Facility	6
5.1.1 Prompt Zone	6
5.1.2 Satellite Catalog Sub-Zone	8
5.1.3 Offline Production Zone	9
5.1.4 Archive Zone	10
5.1.5 Development and Integration Zone	11
5.2 United States Data Facility (USDF) Data Access Center (DAC) Zone	11
5.2.1 Threats and Security infrastructure	12
5.2.2 Disaster recovery	12

5.3 Chile DAC Zone	12
5.3.1 Threats and Security infrastructure	13
5.3.2 Disaster recovery	13
5.4 FrDF Processing Zone	13
5.4.1 Threats and Security infrastructure	13
5.4.2 Disaster recovery	14
5.5 UKDF Processing Zone	14
5.5.1 Threats and Security infrastructure	14
5.5.2 Disaster recovery	14
5.6 External entities	14
5.6.1 Threats and Security infrastructure	14
5.6.2 Disaster recovery	14
6 Telescope and Site System architecture	15
6.1 Summit Systems	15
6.1.1 Threats and Security infrastructure	16
6.1.2 Disaster recovery	16
6.2 Base Test Stand	17
6.2.1 Threats and Security infrastructure	17
6.2.2 Disaster recovery	17
6.3 Tucson Test Stand	17
6.3.1 Threats and Security infrastructure	17
6.3.2 Disaster recovery	17
7 Rubin Directors Office	18
7.1 Threats and Security infrastructure	18
7.2 Disaster recovery	18
A References	18

Rubin Data and Information Security Plan

1 Introduction

This is the Rubin information and data security plan for operations. This renders LDM-324 obsolete.

This security plan conforms with the construction era security documentation:

- “LSST Information Classification Policy” [LPM-122]
- “LSST Master Information Security Policy” [LPM-121]

Note has been taken of Adams et al. (2021).

During the ramp up to operations further security requirements were given to the project by the agencies. The response to these is in DM Technical Note (DMTN)-199.

The Rubin Data Management (RDM) department is concerned with the operations and maintenance of Rubin constructed software, hardware and networks. Development remains open-source in nature.

Data from the telescope on Cerro Pachón is transmitted over Rubin controlled networks to SLAC National Accelerator Laboratory (SLAC) in Menlo Park USA. After an agreed embargo period all raw data will be transmitted to the French Data Facility at the computing center for Institut National de Physique Nucléaire et de Physique des Particules (IN2P3) in Lyon, France, and at least some raw data is transmitted to the United Kingdom Data Facility operating on the IRIS infrastructure and in particular utilizing facilities at Rutherford Appleton Laboratory (United Kingdom (UK)) (RAL) and Lancaster University. Processing occurs in all locations using a coordinated release of the Rubin Science Pipelines.

Access to data will be provided to authorized users via a US DAC. A data access center similar to the United States (US) DAC will also be provided in Chile.

For data taken initially we have been requested to comply with NIST.SP.800-171r3; the com-

pliance matrix is provided in RTN-082. SLAC is subject to NIST.800-53

This document does not do a detailed threat analysis though this should be and is done for individual systems and Zones. Adams et al. (2021) is a good guide for software threats.

2 Information types in Rubin Observatory

Vera C. Rubin Observatory is an open source project with no confidentiality requirements on the software. The software project's integrity requirements are met through the combination of processes and controls which provide verified user access and protected credentials. The majority of software testing is conducted using open simulated and observational data sets.

All data, after the embargo period, is classified as Internal for two years before it becomes public. Internal here means available to data rights holders as per RDO-013. Data rights holders are instructed not to disseminate data outside of the collaboration. This is the customary protection for this data in the field of optical astronomy.

2.1 Presence of controlled information

Rubin observatory data is not considered Controlled Unclassified Information (CUI), however we have been asked to embargo pixel data for 80 hours during Operations and for 30 days during Commissioning, with the ability to extend the embargo for designated images.

2.2 Controlled Catalog

To avoid alerting on certain man-made objects Rubin has access to a satellite catalog which is designated CUI NIST.800-53 *Moderate*. This catalog is held only in a "Satellite Catalog Sub-Zone" inside the "Embargo Rack" within SLAC Shared Scientific Data Facility (S3DF). All access to it is via queries into that system for specific times and positions.

2.3 Embargoed Data

While pixel data is embargoed it is kept on encrypted storage within the embargo rack in S3DF. This is in a locked rack with limited access and observed by cameras. Access is only available

to Rubin and Commissioning staff.

2.4 Unembargoed Data

Once the raw pixel data reaches its embargo age it moves to the Archive Zone within S3DF where it is available to data rights holders. It is still not public but is no longer of a sensitive nature. Prompt pixel image data products are released from embargo at the same time.

2.5 Unreleased Data

Data products generated within the Offline Production Zone are not embargoed, but they are restricted to staff until they are released, at which point they move to the Archive Zone and are made available to data rights holders.

3 Cyber incident response

In this document many threat vectors are identified. Rubin and SLAC IT, Chile DevOps Team and many other Data Management (DM) staff are highly aware of these threat vectors. If an incident is observed we:

- Notify the Information Security Officer (ISO) [LPM-121]
- ISO coordinates the security team analysis to determine the scope of the incident and establish it is a real threat
- ISO notifies Project Management Office (PMO) (project manager and deputies)
- The ISO decides which information and how widely to broadcast to all staff.
- Investigation and remediation are initiated - which may include disaster recovery.
- The combined Rubin Information Technology (IT) team (SLAC, and AURA) take action to isolate, remove or switch off affected systems.
- Staff are informed of the plan to return systems to service.
- ISO organizes a Post Incident Review

4 System description

Vera C. Rubin Observatory produces around 20TB of astronomical data per night for the 10 year Legacy Survey of Space and Time (formerly Large Synoptic Survey Telescope) (LSST) The control of the the observatory is part of Telescope and site and is covered in section 6. Data are processed in SLAC, IN2P3 and Royal Observatory Edinburgh (ROE). A publicly available alert stream emanates from the USDF at SLAC. Processing is the responsibility's of Rubin Data Production (Obsolete use RDM) (RDP) department while Quality Assurance is carried out under Rubin system PerFormance (RPF) department.

The Rubin Operations Plan RDO-018 gives more details. The USDF specification is in DMTN-189 but some architecture details are provided in section 5.

4.1 Data verification and quality assurance

Members of the RPF department assess the data quality at short and long timescales. They require access to tools installed at the summit as well as in the data facilities. In particular they will need access to images as they are processed to generate alerts.

These team members will have to have SLAC accounts to perform this work. SLAC accounts are governed by SLAC policies¹ and SLAC IT specific policies².

In addition they will need summit accounts as governed by ITTN-045 and ITTN-010.

4.2 Data processing

4.2.1 Prompt Processing

Prompt processing is performed on embargoed data with in SLAC's prompt processing zone subsection 5.1.1. This data is subject to (Marshall, ACP).

¹<https://policies.slac.stanford.edu>

²<https://policies.slac.stanford.edu/categories/information-technology>

4.2.2 Data Release Processing

DRP is performed only on unembargoed data which is fully available to Rubin data rights holders. Processing is carried out at three sites. Each site has its own security policies:

- SLAC lists cyber policies on <https://policies.slac.stanford.edu/>.
- UKDF falls under the e-Infrastructure for Research and Innovation for UK Science and Technology Facilities Council (STFC) (IRIS) security policies listed on <https://www.iris.ac.uk/security/>.
- FrDF lists cyber policies and compliance on <https://doc.lsst.eu/cybersecurity/cybersecurity.html>.

4.3 Data access

After a short embargo period image and derived data are made available to all data rights holders as defined in LDO-13.

Front end data access via the Rubin Science Platform (RSP) will be hosted on a cloud provider such as Google. Thus not requiring community science users to have SLAC accounts.

There will be a DAC in Chile to support Chilean users. The UK intends to host a DAC for UK users.

In addition there will be a set of Independent Data Access Center (IDAC)s which will usually serve a portion of the data e.g. perhaps only catalogs or only the object catalog.

4.4 Rubin Director's Office

Rubin headquarters are in Tucson Arizona where a few services are also deployed such as Jira and Confluence. These are discussed more in section 7

5 Data Management system architecture

The overall system architecture is available in LDM-148. Details on the USDF specifications are given in DMTN-189.

section 4 gives a high level overview of the system. Architecturally we look at this as a set of zones, each of which is allocated compute, storage, and networking. As images are processed in the Prompt and Offline Production zones, their resulting data products are stored in the Archive Zone and made available to the DAC Zone where data rights holders can access and analyze them. In addition, Rubin Observatory staff will use the Development/Integration Zone to maintain the Observatory's software tools and systems and to develop new versions of them.

These zones are further described here and for each a series of subsections explore :

1. Threats and Security infrastructure
2. Disaster recovery

5.1 US Data Facility

A number of zones are within the USDF. Responsibility for the design, operations, maintenance, and security of these systems is held by the USDF Lead at SLAC, who may delegate certain functions to the USDF Deputy Lead, USDF Technical Lead, other USDF staff, S3DF staff, and SLAC IT staff.

In particular, at present the USDF Technical Lead serves as the Rubin Information Systems Security Manager for all USDF systems.

5.1.1 Prompt Zone

The Prompt Zone holds and processes pixel data that is subject to embargo, meaning that it is to be accessed only by Rubin and Commissioning staff, not by the Rubin data rights community.

It receives images from the Observatory facilities in Chile via a Long Haul Network connection. It stores these and processes them into Prompt data products of three main types: alerts for things that have moved or changed, measurement catalogs, and processed images. The alerts can be further subdivided into *streak* alerts for objects that have moved a long distance and *non-streak* alerts for all other objects. Measurements in the catalogs follow the same subdivision. Images may be Commissioning images used for testing and characterizing the Observatory systems, normal science images without significant *streaks*, or delayed science images that do contain significant *streaks*.

Streak alerts corresponding to satellites in the Official Use Only (OUO) catalog at SLAC will not be released. Uncatalogued *Streak* alerts agreed with the Jet Propulsion Laboratory (dark energy (DE) ephemerides) (JPL) Near-Earth Object (NEO) group are to be published at least to Minor Planet Center (MPC). *Non-streak* alerts are to be published to the world at large within 60 seconds of the original raw image being taken. Normal science images are made available to data rights holders in the DAC after an 80 hour embargo period. Delayed science images, as identified by the agencies, are released after their specific embargo period. Commissioning images are made available to data rights holders after a 30 day embargo period.

All Prompt data products are checked for quality by automated systems but also by human operators from the Rubin Observatory staff, who have access to all images and data products in order to perform spot checks or follow ups.

5.1.1.1 Threats and Security infrastructure The obvious threat surfaces here are :

1. Transmission of Data from Chile. IPSec built into the routers will be used on the long haul network (LHN). DMTN-108 discusses threats in this realm a little more.
2. Transmission to JPL. This transmission will only include measurements of potential NEOs, not pixel data, and will be over internet using Transport Layer Security (TLS).
3. Pipeline access for Prompt Processing. The pipelines will run on resources dedicated to the Prompt Zone and without direct external Internet connections for the duration of the processing, mitigating external, intra-node, and side channel attacks.
4. Staff access for Quality Assurance (QA). All the usual user threats such as phishing apply

- these users are however governed by SLAC security policies³.
- 5. QA tools. The web accessible QA tools should have a threat analysis performed by SLAC or our Security consultants although they will be behind SLAC authentication (including 2FA), Rubin authorization (via Gafaelawr DMTN-234) and the S3DF HTTPS Ingress. These tools are also only accessible by staff and probably pose a low risk.
- 6. Physical access to storage. For the initial 80 hours (30 days in commissioning) after acquisition the data is maintained on encrypted storage within the physically locked Embargo rack.

5.1.1.2 Disaster recovery All embargoed raw data is also stored on a secure server in Chile as part of the Observatory Operations Data Service (OODS), hence it can be retransmitted as needed. For a limited time greater than 80 hours, embargoed raw data is also available from the LSSTCam Camera Control System, giving a third data location during the Operations embargo period. Embargoed data products can be regenerated from the raw data. In the case of a total wipe out of the Prompt Zone systems, use of Chef, Kubernetes, etc. allow rapid redeployment. See also the USDF Disaster Recovery Plan RTN-078.

5.1.2 Satellite Catalog Sub-Zone

This Zone holds the only OUO data on the project, a catalog of satellite orbit information. This catalog will be held on two dedicated servers (for availability) within the Embargo Rack. Access to these servers is limited to a highly vetted list of administrators and operators, including the USDF Lead, USDF Deputy Lead, and USDF Technical Lead, who are a (very) small subset of Rubin staff.

5.1.2.1 Threats and Security infrastructure

1. Catalog retrieval. The catalog is downloaded from an external source via HTTPS. Connections originate in the Sub-Zone from a dedicated server and are authenticated and authorized via a service account provided by the external source. This minimizes access

³<https://policies.slac.stanford.edu/> see also <https://it.slac.stanford.edu/cybersecurity/compliance>

to the catalog in transit and avoids intra-node and side channel attacks on the catalog service.

2. Catalog lookup. The catalog is held in the memory of a service that provides match/no match results for each streak measurement in a list. The service is only accessible by S3DF systems, not the Internet. This minimizes attacks via the service Application Programming Interface (API) and prevents attacks via storage.

5.1.2.2 Disaster recovery Since the catalog is downloaded at startup and periodically, it will be refreshed automatically after any disaster.

5.1.3 Offline Production Zone

The Offline Production Zone holds data that is not embargoed but is not yet released to the data rights community. Access to this data is limited to Rubin staff.

Each year (or more frequently), the Offline Production Zone takes the raw images accumulated to date in the Archive and reprocesses them to generate highly accurate, consistent images and measurement catalogs, known as a Data Release. These data products are stored in the Archive and made available to data rights holders in the DAC after they have been checked by automated systems and after Rubin Observatory staff has vetted, characterized, and documented them. Offline Production is split between the USDF and the French Data Facility (FrDF) and UKDF. Each Data Facility performs part of the computations and exchanges its results with the others, so all have a complete set of data products at release time.

5.1.3.1 Threats and Security infrastructure

- Offline production data is no longer embargoed ergo not considered under threat. Early release of data to the data rights community or even release of small amounts of data to the world at large is more of a “science” problem, less a “security” problem.
- Although data exchange among the facilities uses encryption (secure HTTP over TLS), if data were to be intercepted in transfer between sites, this could only occur after the embargo period hence the security risk is low.

- Malicious users could disrupt data or processing. We are using standard tooling from High Energy Physics (HEP) which has been in use for many years and gives a level confidence of their suitability in this scientific endeavor. Still internal users remain a major risk - we maintain an inclusive project and try to avoid disgruntled team members.

5.1.3.2 Disaster recovery The Offline Production systems run in batch and Kubernetes and can be reconstituted after a disaster. The data being processed can be regenerated from the Archive Zone below.

5.1.4 Archive Zone

The raw images, released data products, and other records of the survey such as commands, events, and telemetry from Observatory systems are all stored in the Archive. As the permanent scientific record of the survey, no more than 1% of the raw images or telemetry may be lost or corrupted according to Rubin requirements.

To help ensure this, the French Data Facility maintains a disaster recovery copy of all raw images and selected data products. Additional copies of some raw images and data products will be stored in Observatory systems in Chile. Raw images and key data products are also stored on tape backup at SLAC.

5.1.4.1 Threats and Security infrastructure

- Archive data is no longer embargoed ergo not considered under threat. All Archive data becomes public after the two year proprietary period, so any disclosure of small portions to the world at large is a premature release and not really a security issue.

5.1.4.2 Disaster recovery Post embargo FrDF keeps a full copy of the raw data.

5.1.5 Development and Integration Zone

Rubin Observatory and USDF staff will use this zone to build and test new versions of software and services to be deployed in the other zone.

5.1.5.1 Threats and Security infrastructure

- Developers have a higher level of access than data rights holders. This is a necessary and accepted risk.
- All developers must have SLAC accounts and therefore adhere to SLAC access rules e.g. Foreign Access Central Tracking System (FACTS) checking etc.

5.1.5.2 Disaster recovery SLAC keep tape backups.

All code is deployed using Kubernetes or Chef and hence fairly easily recoverable in case of catastrophic failure.

5.2 USDF DAC Zone

The USDF DAC is hosted on Google and is the responsibility of SLAC. All deployments on the USDF DAC are made by Science Quality and Reliability Engineering (SQuaRE) using Phalanx.

Data rights holders will use the services and systems in this zone to work with the survey data products. It is therefore a general-purpose scientific computing facility. Generally users will interact with the Rubin Science Platform (RSP), which is composed of a web-based Portal Aspect providing a guided user interface for accessing and analyzing the data, a Notebook Aspect providing an interactive, flexible, programming-oriented interface, and an API Aspect providing an programmable access service. Users of the DAC may connect from anywhere in the world over the Internet; all such users will be authenticated and authorized before accessing any RSP service. The RSP is hosted on a cloud service, currently Google Cloud Platform.

The DAC retrieves the released data products from the Archive Zone via protocols and services authenticated at a service account level only. While end-user identities may be included for

audit and accounting purposes, fundamentally the DAC exists to provide access to all Archive contents.

5.2.1 Threats and Security infrastructure

The RSP is an attractive generic target due to its computing resources. There is some user generated data which is mildly sensitive. Hosting it on a cloud provider reduces risk considerably for the Archive zone, and also leverages the security products and services made available by the hosting provider. SQR-041 provides a risk assessment for the RSP. DMTN-193 provides a more in depth web risk analysis.

- We will have a lot of users which could be problematic. Keeping the data rights holders on the cloud allows a clean separation of concerns between SLAC for processing and archive and the more public facing RSP.
- Backend archive services could provide another attack surface. These are governed by SLAC security.

5.2.2 Disaster recovery

For the software and deployed systems, all information needed to reconstitute the US DAC is stored in public repositories of container images and configuration files.

For user spaces we rely on cloud provider redundancy/backup/recovery.

The data in the cloud is merely a cache; a full copy is always held at the USDF hence any Rubin data at the DAC is expendable.

Further considerations are covered in RTN-059.

5.3 Chile DAC Zone

This proposed DAC in Chile is covered in O'Mullane (LDM-572). We will start work on this in 2025 nearer the start of operations. Chile DevOps team are responsible for the Chile DAC.

Most applications deployed on the Chile DAC will be deployed by DM's SQuaRE team. The applications will be same as deployed on the USDF DAC.

5.3.1 Threats and Security infrastructure

- The Chile DAC is within the Recinto data center and covered by Association of Universities for Research in Astronomy (AURA)/COS security measures.
- All Rubin traffic is run through a security appliance (currently Zeek).
- Selected Chilean users have access to the DAC. We will keep the DAC and the users confined with least privileges. We will use a caching mechanism analogous to the Cloud DAC system to restrict access to the object store for the external users.
- All access will be via RSP pods and hence containerized - escalation potential from inside the container will be carefully monitored.

5.3.2 Disaster recovery

The Chile disaster recovery plan will cover the Chile DAC ITTN-055.

5.4 FrDF Processing Zone

40% of Data Release Production (DRP) will be done at IN2P3. A full back up of the raw data will also be held there. The IN2P3 computing infrastructure is described in <https://doc.lsst.eu/>. The Centre national de la recherche scientifique (CNRS) staff at IN2P3 are fully responsible for the FrDF.

5.4.1 Threats and Security infrastructure

The considerations for this Zone are a combination of those for the Offline Production Zone and Archive Zone. Access will be granted only to IN2P3 staff and Rubin staff. IN2P3 have their own cyber security procedures which will be adhered to.

5.4.2 Disaster recovery

All Raw data is also at SLAC and can be resent over a period of time.

5.5 UKDF Processing Zone

25% of processing will be done on IRIS. ROE staff are responsible for the United Kingdom Data Facility (UKDF) noting that IRIS is a shared computing facility beyond their control.

5.5.1 Threats and Security infrastructure

The considerations for this Zone are a combination of those for the Offline Production Zone and Archive Zone. Access will be granted only to UKDF staff and Rubin staff. UKDF have their own cyber security procedures which will be adhered to.

5.5.2 Disaster recovery

All Raw data is at IN2P3 and can be resent over a period of time.

5.6 External entities

There are a number of IDACs which will have and serve catalogs and or images. These are within our realm of security to some extent but not entirely - we rely on trust at some level.

5.6.1 Threats and Security infrastructure

The obvious threat here is unauthorized access to the data rights accessible data. Any IDAC must adhere to our user access protocols so this should not happen. If unauthorized access occurs the impact is low in terms of system integrity - it may reflect badly on Rubin Observatory and erode the brand and the entire notion of restricted access to the data.

5.6.2 Disaster recovery

We are not concerned with disasters at IDACs. We can resend the appropriate data to them.

6 Telescope and Site System architecture

We concern our selves here mainly with the software architecture of telescope and site, this includes the control system but also the controlled devices and various test stands.

The control system architecture is given in LSE-150. Broadly this is a message bus architecture with various controllable components such as the Camera, Environmental Control, etc. attached to it. The components can receive control messages and telemetry from the bus by listening to various queues. The script queue component allows for orchestrated commanding of various components.

This set may be seen for the Main Telescope as well as the Auxiliary Telescope. In addition there is a test stand in the Base and one in Tucson which have physical DAQ hardware to emulate the camera and can simulate many other physical components for testing the control system.

We consider these systems under the same headings used in section 5.

6.1 Summit Systems

The summit is the crown jewel of Rubin; the network and control system touch all the hardware on the summit. The control system, which touches all hardware, is described in LSE-150. This is a message bus system allowing command of all commandable devices from the Telescope Mount Assembly (TMA) to the Heating, Ventilation, and Air Conditioning (HVAC). Most of the computing hardware lives in the summit computer room on the second floor which is card accessible and has cameras in place. Racks in the computing room are locked with individual codes known to Rubin IT. Combinations are kept in a password vault. Access to services on the summit is more restricted than to the rest of Rubin, see the on boarding procedure ITTN-045.

Access to most controls is through the control room which is a key card accessible room on the second floor of the observatory

Underlying some of this is the virtualization system as described ITTN-036.

6.1.1 Threats and Security infrastructure

The summit has several security features coming from the Chilean infrastructure:

- The summit has firewalls and 2FA enabled VPN access.
- Accessibility to the summit is via the access road which has a physical security check.
- The Control and Computer rooms as well as the Dome can only be accessed by authorized personnel with key cards.

There are also many threats:

- DMTN-108 discusses some issues such as fiber taps to access data.
- Assuming access was gained to the computer room physical disks could be removed, our infrastructure as code approach allows us to quickly rebuild servers, the data is also available at SLAC. Disks on the summit are encrypted meaning it would be quite difficult for anyone to retrieve data from any physical disks removed from the computer room.
- As always our network may be vulnerable to attack, we follow NIST advice and will have a contract with a cyber security firm to assist in this area.

6.1.2 Disaster recovery

- We can deploy most systems from scratch using the Rancher, puppet, kubernetes Infrastructure as Code (IaC) approach. This is relatively quick (hours) provided machines are available. This also means machines are interchangeable and we keep at least one spare on the summit.
- Other systems such as the coating chamber control computer have spares since we can not rebuild them easily.
- Should we have an all out attack on the system via the LHN we have an out of bounds link which still provides access and monitoring (allowing shutdown if needed).
- Though the software *could* command systems out of limits all the physical devices have engineering safety stops build in.

6.2 Base Test Stand

In the computer room on the base facility in Las Serena we have the Base (La Serena) Test Stand (BTS). This is a full Data Acquisition System (DAQ) identical to that attached to the camera on the summit, as well as a set of supporting machines which allow deployment of both control components and simulators. This allows full scale testing of the summit control systems and especially the Camera readout.

6.2.1 Threats and Security infrastructure

This system is behind the La Serena firewall. It is accessible by VPN. Access is restricted to the computer room and cameras are in place.

6.2.2 Disaster recovery

Apart from the DAQ itself the machines here are standard and the system can be rebuilt using our IaC approach. There is no irreplaceable data on the system.

6.3 Tucson Test Stand

The Tucson Test Stand (TTS) is located in the computer room on Cherry Ave in Tucson. It is similar to the BTS subsection 6.2 but has a smaller DAQ more ComCam sized. This is still very useful for testing.

6.3.1 Threats and Security infrastructure

This system is behind the Tucson firewall. It is accessible by VPN. Access is restricted to the computer room with only a few AURA employees allowed to access it.

6.3.2 Disaster recovery

Apart from the DAQ itself the machines here are standard and the system can be rebuilt using our IaC approach.

7 Rubin Directors Office

The directors office is in Tucson Arizona and hosts several observatory functions.

These include :

- Active directory and Data Services.
- Websites such as Drupal.
- Databases such as Docushare, Contacts Database, Primavera.
- Terminal servers for access to some windows based services such as as primavera.
- Outlook mail server.
- Support of Jira/Confluence Cloud.

7.1 Threats and Security infrastructure

The “LSST Tucson Site IT Cybersecurity Policy” [LPM-125] is the policy for directors office.

The main threats are against our web interfaces such as Drupal and Outlook mail server. IT keep these servers up to date with security patches and we look out for any threat warnings.

7.2 Disaster recovery

The construction era “LSST Tucson Site Disaster Recovery Plan” [LPM-101] covers disaster recovery for the directors office.

A References

Adams, A., Avila, K., Heymann, E., et al., 2021, Guide to securing scientific software, URL <https://zenodo.org/record/5777646#.YfSEvmB1C3o>

- [SQR-041]**, Allbery, R., 2022, Science Platform security risk assessment, URL <https://sqr-041.lsst.io/>,
Vera C. Rubin Observatory SQuaRE Technical Note SQR-041
- [DMTN-193]**, Allbery, R., 2025, Web security for the Science Platform, URL <https://dmtn-193.lsst.io/>,
Vera C. Rubin Observatory Data Management Technical Note DMTN-193
- [DMTN-234]**, Allbery, R., 2025, RSP identity management design, URL <https://dmtn-234.lsst.io/>,
Vera C. Rubin Observatory Data Management Technical Note DMTN-234
- [RDO-018]**, Blum, R., 2021, PLAN for the OPERATIONS of the VERA C. RUBIN OBSERVATORY, URL <https://docushare.lsstcorp.org/docushare/dsweb/Get/RDO-18>,
Vera C. Rubin Observatory RDO-018
- [RDO-013]**, Blum, R., the Rubin Operations Team, 2020, Vera C. Rubin Observatory Data Policy, URL <https://ls.st/RDO-013>,
Vera C. Rubin Observatory RDO-013
- [LDO-13]**, Blum, R., et al., 2019, LSST Data Policy, URL <https://ls.st/LDO-13>,
Vera C. Rubin Observatory LDO-13
- [RTN-078]**, Dubois, R., 2024, USDF Disaster Recovery Plan, URL <https://rtn-078.lsst.io/>,
Vera C. Rubin Observatory Technical Note RTN-078
- [NIST.800-53]**, FORCE, J.T., 2020, COMPUTER SECURITY, URL <https://doi.org/10.6028/NIST.SP.800-53r5>
- [LPM-101]**, Goodenow, I., McKercher, R., 2013, Tucson Site Disaster Recovery Plan, URL <https://ls.st/LPM-101>,
Vera C. Rubin Observatory LPM-101
- [LDM-324]**, Kantor, J., 2016, Data Management Information Security Plan, URL <https://ls.st/LDM-324>,
Vera C. Rubin Observatory LDM-324
- [LPM-125]**, Krabendam, V., Goodenow, I., 2016, Project Management Office Information Security Plan, URL <https://ls.st/LPM-125>,
Vera C. Rubin Observatory LPM-125

- [DMTN-189]**, Lim, K.T., 2021, Data Facility Specifications, URL <https://dmtn-189.lsst.io/>,
Vera C. Rubin Observatory Data Management Technical Note DMTN-189
- [LDM-148]**, Lim, K.T., Bosch, J., Dubois-Felsmann, G., et al., 2020, Data Management System Design, URL <https://ldm-148.lsst.io/>,
Vera C. Rubin Observatory Data Management Controlled Document LDM-148
- [ACP]**, Marshall, P., 2024, Access Control Plan for the Vera C. Rubin Observatory U.S. Data Facility Embargo Rack , URL <https://ls.st/ACP>,
Internal document
- [DMTN-108]**, O'Mullane, W., 2021, Security of Rubin Observatory data, URL <https://dmtn-108.lsst.io/>,
Vera C. Rubin Observatory Data Management Technical Note DMTN-108
- [LDM-572]**, O'Mullane, W., 2021, Chilean Data Access Center, URL <https://ldm-572.lsst.io/>,
Vera C. Rubin Observatory Data Management Controlled Document LDM-572
- [RTN-082]**, O'Mullane, W., 2024, Pixel Zone system security plan, URL <https://rtn-082.lsst.io/>,
Vera C. Rubin Observatory Technical Note RTN-082
- [LPM-122]**, Petravick, D., 2015, LSST Information Classification Policy, URL <https://ls.st/LPM-122>,
Vera C. Rubin Observatory LPM-122
- [LPM-121]**, Petravick, D.L., Withers, A., 2016, LSST Master Information Security Policy, URL <https://ls.st/LPM-121>,
Vera C. Rubin Observatory LPM-121
- [ITTN-036]**, Reinking, H., 2021, Virtualization Cluster Topology and Design, URL <https://ittn-036.lsst.io/>,
Vera C. Rubin Observatory ITTN-036
- [LSE-150]**, Ribeiro, T., O'Mullane, W., Axelrod, T., Mills, D., 2020, Control Software Architecture, URL <https://lse-150.lsst.io/>,
Vera C. Rubin Observatory LSE-150
- [NIST.SP.800-171r3]**, Ross, R., Pillitteri, V., 2024, Special publication 800-171, protecting controlled unclassified information in nonfederal systems and organizations, URL <https://doi.org/10.6028/NIST.SP.800-171r3>

[ITTN-055], Silva, C., 2023, Disaster Recovery, URL <https://ittn-055.lsst.io/>,
Vera C. Rubin Observatory ITTN-055

[ITTN-045], Tapia, D., Silva, C., 2024, Summit Onboarding Procedure, URL <https://ittn-045.lsst.io/>,
Vera C. Rubin Observatory ITTN-045

[ITTN-010], Thebo, A., Hoblitt, J., 2023, User Identification and Authorization, URL <https://ittn-010.lsst.io/>,
Vera C. Rubin Observatory ITTN-010

[RTN-059], White, B., 2023, Rubin Data Retention Implementation Strategy, URL <https://rtn-059.lsst.io/>,
Vera C. Rubin Observatory Technical Note RTN-059

B Glossary

API Application Programming Interface.

Archive The repository for documents required by the NSF to be kept. These include documents related to design and development, construction, integration, test, and operations of the LSST observatory system. The archive is maintained using the enterprise content management system DocuShare, which is accessible through a link on the project website www.project.lsst.org.

Association of Universities for Research in Astronomy consortium of US institutions and international affiliates that operates world-class astronomical observatories, AURA is the legal entity responsible for managing what it calls independent operating Centers, including LSST, under respective cooperative agreements with the National Science Foundation. AURA assumes fiducial responsibility for the funds provided through those cooperative agreements. AURA also is the legal owner of the AURA Observatory properties in Chile.

AURA Association of Universities for Research in Astronomy.

BTS Base (La Serena) Test Stand.

Camera The LSST subsystem responsible for the 3.2-gigapixel LSST camera, which will take more than 800 panoramic images of the sky every night. SLAC leads a consortium of Department of Energy laboratories to design and build the camera sensors, optics, electronics, cryostat, filters and filter exchange mechanism, and camera control system.

Center An entity managed by AURA that is responsible for execution of a federally funded project.

cloud A visible mass of condensed water vapor floating in the atmosphere, typically high above the ground or in interstellar space acting as the birthplace for stars. Also a way of computing (on other peoples computers leveraging their services and availability)..

CNRS Centre national de la recherche scientifique.

Commissioning A two-year phase at the end of the Construction project during which a technical team a) integrates the various technical components of the three subsystems; b) shows their compliance with ICDs and system-level requirements as detailed in the LSST Observatory System Specifications document (OSS, LSE-30); and c) performs science verification to show compliance with the survey performance specifications as detailed in the LSST Science Requirements Document (SRD, LPM-17).

CUI Controlled Unclassified Information.

DAC Data Access Center.

DAQ Data Acquisition System.

Data Access Center Part of the LSST Data Management System, the US and Chilean DACs will provide authorized access to the released LSST data products, software such as the Science Platform, and computational resources for data analysis. The US DAC also includes a service for distributing bulk data on daily and annual (Data Release) timescales to partner institutions, collaborations, and LSST Education and Public Outreach (EPO)..

Data Management The LSST Subsystem responsible for the Data Management System (DMS), which will capture, store, catalog, and serve the LSST dataset to the scientific community and public. The DM team is responsible for the DMS architecture, applications, middleware, infrastructure, algorithms, and Observatory Network Design. DM is a distributed team working at LSST and partner institutions, with the DM Subsystem Manager located at LSST headquarters in Tucson.

Data Release Production An episode of (re)processing all of the accumulated LSST images, during which all output DR data products are generated. These episodes are planned to occur annually during the LSST survey, and the processing will be executed at the Archive Center. This includes Difference Imaging Analysis, generating deep Coadd Images, Source detection and association, creating Object and Solar System Object catalogs, and related metadata.

DE dark energy.

Director The person responsible for the overall conduct of the project; the LSST director is charged with ensuring that both the scientific goals and management constraints on

the project are met. S/he is the principal public spokesperson for the project in all matters and represents the project to the scientific community, AURA, the member institutions of LSSTC, and the funding agencies.

DM Data Management.

DMTN DM Technical Note.

DRP Data Release Production.

FACTS Foreign Access Central Tracking System.

FrDF French Data Facility.

HEP High Energy Physics.

HVAC Heating, Ventilation, and Air Conditioning.

IaC Infrastructure as Code.

IDAC Independent Data Access Center.

IN2P3 Institut National de Physique Nucléaire et de Physique des Particules.

Independent Data Access Center Externally supported and administered versions of the DAC to serve the full, or a limited subset of, the LSST data products and/or software to authorized users..

IRIS e-Infrastructure for Research and Innovation for STFC.

ISO Information Security Officer.

IT Information Technology.

JPL Jet Propulsion Laboratory (DE ephemerides).

Kubernetes A system for automating application deployment and management using software containers (e.g. Docker); <https://kubernetes.io>.

LHN long haul network.

LSST Legacy Survey of Space and Time (formerly Large Synoptic Survey Telescope).

MPC Minor Planet Center.

NEO Near-Earth Object.

Object In LSST nomenclature this refers to an astronomical object, such as a star, galaxy, or other physical entity. E.g., comets, asteroids are also Objects but typically called a Moving Object or a Solar System Object (SSObject). One of the DRP data products is a table of Objects detected by LSST which can be static, or change brightness or position with time.

Operations The 10-year period following construction and commissioning during which the LSST Observatory conducts its survey.

OUO Official Use Only.

PMO Project Management Office.

Project Management Office the work element responsible for achieving the project's objec-

tives.

QA Quality Assurance.

Quality Assurance All activities, deliverables, services, documents, procedures or artifacts which are designed to ensure the quality of DM deliverables. This may include QC systems, in so far as they are covered in the charge described in LDM-622. Note that contrasts with the LDM-522 definition of “QA” as “Quality Analysis”, a manual process which occurs only during commissioning and operations. See also: Quality Control.

RAL Rutherford Appleton Laboratory (UK).

RDM Rubin Data Management.

RDP Rubin Data Production (Obsolete use RDM).

ROE Royal Observatory Edinburgh.

RPF Rubin system PerFormance.

RSP Rubin Science Platform.

S3DF SLAC Shared Scientific Data Facility.

Science Pipelines The library of software components and the algorithms and processing pipelines assembled from them that are being developed by DM to generate science-ready data products from LSST images. The Pipelines may be executed at scale as part of LSST Prompt or Data Release processing, or pieces of them may be used in a standalone mode or executed through the Rubin Science Platform. The Science Pipelines are one component of the LSST Software Stack.

Science Platform A set of integrated web applications and services deployed at the LSST Data Access Centers (DACs) through which the scientific community will access, visualize, and perform next-to-the-data analysis of the LSST data products.

SLAC SLAC National Accelerator Laboratory.

SLAC National Accelerator Laboratory A national laboratory funded by the US Department of Energy (DOE); SLAC leads a consortium of DOE laboratories that has assumed responsibility for providing the LSST camera. Although the Camera project manages its own schedule and budget, including contingency, the Camera team’s schedule and requirements are integrated with the larger Project. The camera effort is accountable to the LSSTPO..

software The programs and other operating information used by a computer..

SQuARE Science Quality and Reliability Engineering.

STFC UK Science and Technology Facilities Council.

TLS Transport Layer Security.

TMA Telescope Mount Assembly.

TTS Tucson Test Stand.

UK United Kingdom.

UKDF United Kingdom Data Facility.

US United States.

USDF United States Data Facility.

Zone Individually defined portions of the computational resources at the Summit, Base, USDF, and Satellite Facilities, such as the Pixel Zone, Prompt Zone, Archive Zone, etc..

Draft