



Vera C. Rubin Observatory
Rubin Observatory Operations

Rubin Data and Information Security Plan

William O'Mullane and Russ Allbery and Richard Dubois and KT. Lim

RTN-030

Latest Revision: 2024-06-18

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 | 2024-06-25 | SLAC disaster recovery, UK security policies, tidy and issue for review | William O'Mullane |

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

Draft

Contents

| | |
|---|----------|
| 1 Introduction | 1 |
| 2 Presence of controlled information | 2 |
| 3 Cyber incident response | 2 |
| 4 System description | 3 |
| 4.1 Data verification and quality assurance | 3 |
| 4.2 Data processing | 3 |
| 4.2.1 Data Release Processing | 3 |
| 4.3 Data access | 4 |
| 4.4 Rubin Director’s Office | 4 |
| 5 Data Management system architecture | 4 |
| 5.1 Prompt Enclave | 5 |
| 5.1.1 Threats and Security infrastructure | 5 |
| 5.1.2 Disaster recovery | 6 |
| 5.2 Offline Production Enclave | 6 |
| 5.2.1 Threats and Security infrastructure | 7 |
| 5.2.2 Disaster recovery | 7 |
| 5.3 Archive Enclave | 7 |
| 5.3.1 Threats and Security infrastructure | 7 |
| 5.3.2 Disaster recovery | 8 |
| 5.4 USDF Data Access Center (DAC) Enclave | 8 |
| 5.4.1 Threats and Security infrastructure | 8 |
| 5.4.2 Disaster recovery | 9 |
| 5.5 Chile DAC Enclave | 9 |
| 5.5.1 Threats and Security infrastructure | 9 |
| 5.5.2 Disaster recovery | 9 |

| | |
|---|-----------|
| 5.6 Development and Integration Enclave | 9 |
| 5.6.1 Threats and Security infrastructure | 10 |
| 5.6.2 Disaster recovery | 10 |
| 5.7 FrDF Processing Enclave | 10 |
| 5.7.1 Threats and Security infrastructure | 10 |
| 5.7.2 Disaster recovery | 10 |
| 5.8 UKDF Processing Enclave | 10 |
| 5.8.1 Threats and Security infrastructure | 10 |
| 5.8.2 Disaster recovery | 11 |
| 5.9 External entities | 11 |
| 5.9.1 Threats and Security infrastructure | 11 |
| 5.9.2 Disaster recovery | 11 |
| 6 Telescope and Site System architecture | 11 |
| 6.1 Summit Systems | 12 |
| 6.1.1 Threats and Security infrastructure | 12 |
| 6.1.2 Disaster recovery | 13 |
| 6.2 Base Test Stand | 13 |
| 6.2.1 Threats and Security infrastructure | 13 |
| 6.2.2 Disaster recovery | 14 |
| 6.3 Tucson Test Stand | 14 |
| 6.3.1 Threats and Security infrastructure | 14 |
| 6.3.2 Disaster recovery | 14 |
| 7 Rubin Directors Office | 14 |
| 7.1 Threats and Security infrastructure | 15 |
| 7.2 Disaster recovery | 15 |
| A References | 15 |
| B Compliance with NIST800-171 | 17 |

C Actions (temporary)

21

Draft

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)-199¹.

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 data will be transmitted to the Institut National de Physique Nucléaire et de Physique des Particules (IN2P3) in Lyon, France, and at least some data is transmitted to the Royal Observatory Edinburgh. Processing occurs in all three locations using a coordinate 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.800-171r2 the compliance matrix is provided in Appendix B

¹This is a limited availability document

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

2 Presence of controlled information

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.

Rubin observatory data is not considered Controlled Unclassified Information (CUI), however we have been asked to embargo data for eighty hours and up to ten days for some images. 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.

3 Cyber incident response

In this document many threat vectors are identified all IT, DevOps and many other DM staff are highly aware of these. If an incident is observed we:

- notify the Information Security Officer (ISO) [LPM-121]
- ISO notifies Project Management Office (PMO) (project manager and deputies)
- the ISO decides which information and how widely to broadcast to all staff
- IT team take action to isolate, remove or switch off affected systems.
- investigation and remediation are initiated - which may be disaster recovery.
- staff are informed of the plan to return systems to service.

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 United States Data Facility (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 procedures ? and IT (2024). In addition they will need summit accounts as governed by ITTN-045 and ITTN-010.

4.2 Data processing

4.2.1 Data Release Processing

DRP is performed only on unencumbered 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://it.slac.stanford.edu/cybersecurity/compliance>.
- 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 period of encumbrance 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 United Kingdom (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 enclaves. As images are processed in the Prompt and Offline Production enclaves, their resulting data products are stored in the Archive enclave and made available to the DAC enclave where data rights holders can access and analyze them. In addition, Rubin Observatory staff will use the Development/Integration enclave to maintain the Observatory's software tools and systems and to develop new versions of them.

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

1. Threats and Security infrastructure
2. Disaster recovery

5.1 Prompt Enclave

The Prompt enclave 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 are transmitted to an Alert Vetting System (AVS) located at a Trusted Broker Facility at Lawrence Livermore National Lab. *Non-streak* alerts and *streak* alerts approved by the Alert Vetting System (AVS) 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 AVS, and any unapproved *streak* measurements taken from them are made available after a 10 day 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 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 Lawrence Livermore National Laboratory (LLNL). This will be over internet using TLS.
3. Staff access for Quality Assurance (QA). All the usual user threats such as phishing apply - these users are however governed by SLAC security policies².
4. 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 virtual private network (VPN) and 2FA. These tools are also only accessible by staff and probably pose a low risk.
5. The for the initial 10 days (30 in commissioning) after acquisition the data is maintain on encrypted disks with in the restricted Other Government Agencies (OGA) rack.

5.1.2 Disaster recovery

All embargoed data is also stored on a secure server in Chile hence it can be retransmitted as needed. In the case of a total wipe out of the OGA systems use of Chef, docker etc allow redeployment rapidly. See also the SLAC Rubin disaster recovery plan RTN-078. The back up for the embargoed data is the Chile OGA rack.

5.2 Offline Production Enclave

Each year (or more frequently), the Offline Production enclave 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.

²<https://it.slac.stanford.edu/cybersecurity/compliance>

5.2.1 Threats and Security infrastructure

- Offline production data is no longer embargoed ergo not considered under threat.
- 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.

5.2.2 Disaster recovery

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

5.3 Archive Enclave

The raw images, 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.

We intend to have a hybrid access model for the archive where the RSP users and user processing will be on Google cloud while the data resides at SLAC. A cache of images will be held on Google with a dedicated client pulling needed images from SLAC to google.

5.3.1 Threats and Security infrastructure

- Archive data is no longer embargoed ergo not considered under threat.

- 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.

5.3.2 Disaster recovery

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

5.4 USDF DAC Enclave

Data rights holders will use the services and systems in this enclave 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 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 Enclave 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.

5.4.1 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 enclave, 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.

Backend archive services could provide another attack surface. These are governed by SLAC security.

AI S-1

5.4.2 Disaster recovery

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

Our data is cached 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.

5.5 Chile DAC Enclave

This proposed DAC in Chile is covered in O'Mullane (LDM-572). We will not start work on this until near the start of operations.

5.5.1 Threats and Security infrastructure

- The Chile DAC is within the Recinto data center and covered by 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.

5.5.2 Disaster recovery

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

5.6 Development and Integration Enclave

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

5.6.1 Threats and Security infrastructure

- Developers have a higher level of access than most users. This is a necessary and accepted risk.
- All developers must adhere to SLAC access rules e.g. FACTS checking etc.

5.6.2 Disaster recovery

SLAC keep tape backups.

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

5.7 FrDF Processing Enclave

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/>.

5.7.1 Threats and Security infrastructure

IN2P3 have their own cyber security procedures which will be adhered to.

5.7.2 Disaster recovery

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

5.8 UKDF Processing Enclave

25% of processing will be done on IRIS.

5.8.1 Threats and Security infrastructure

UKDF have their own cyber security procedures which will be adhered to.

5.8.2 Disaster recovery

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

5.9 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.

5.9.1 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.

5.9.2 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 iOS 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 as set of supporting machines which allow deployment of both control components and simulators. This allow 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 commuter 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, Jira, Confluence .
- Databases such as Docushare, Contacts Database, Primavera.
- Terminal servers for access to some windows based services such as as primavera.

Jira, Confluence, Docushare as well as engineering oriented services such as remain in place for the directors office.

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 Confluence and Drupal. 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>

[DMTN-193], Allbery, R., 2022, Web security for the Science Platform, URL <https://dmtn-193.lsst.io/>,
Vera C. Rubin Observatory Data Management Technical Note DMTN-193

[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

[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

[LPM-101], Goodenow, I., McKercher, R., 2013, Tucson Site Disaster Recovery Plan, URL <https://ls.st/LPM-101>,
Vera C. Rubin Observatory LPM-101

IT, S., 2024, SLAC New Hire IT Orientation, URL <https://it.slac.stanford.edu/sites/default/files/2024-05/SLAC%20New%20Hire%20IT%20Orientation.pdf>

[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

[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

[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.800-171r2]**, ROSS, R., VISCUSO, P., GUISSANIE, G., DEMPSEY, K., RIDDLE, M., 2020, Special publication 800-171, protecting controlled unclassified information in nonfederal systems and organizations, URL <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-171r2.pdf>
- [ITTN-055]**, Silva, C., 2023, Disaster Recovery, URL <https://ittn-055.lsst.io/>,
Vera C. Rubin Observatory ITTN-055
- [NIST.800-114]**, Souppaya, M., Scarfone, K., 2016, COMPUTER SECURITY, URL <https://doi.org/10.6028/NIST.SP.800-114r1>
- [NIST.800-46]**, Souppaya, M., Scarfone, K., 2016, COMPUTER SECURITY, URL <https://doi.org/10.6028/NIST.SP.800-46r2>
- [ITTN-045]**, Tapia, D., Silva, C., 2023, 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 Compliance with NIST800-171

Table 1: This table provides an overview of the NIST.SP.800-171r2 and Rubin compliance with it.

| NIST 800-171 | 2024 Status | Intended Compliance | Note |
|---|-------------|---------------------|---|
| 3.1 ACCESS CONTROL | | | |
| 3.1.1 Limit system access to authorized users, processes acting on behalf of authorized users, and devices (including other systems). | Y | Y | |
| 3.1.2 Limit system access to the types of transactions and functions that authorized users are permitted to execute. | Y | Y | IPA groups are in place on the summit restricting users abilities. Legacy systems use the active directory groups for this. |
| 3.1.3 Control the flow of CUI in accordance with approved authorizations. | Y | Y | |
| 3.1.4 Separate the duties of individuals to reduce the risk of malevolent activity without collusion. | Y | Y | Principle of least privilege is applied. Many users have access to hosts that is unneeded. |
| 3.1.5 Employ the principle of least privilege, including for specific security functions and privileged accounts. | N | Y | Targeted sudo rules are needed for common operations. IPA controls sudo centrally |
| 3.1.6 Use non-privileged accounts or roles when accessing nonsecurity functions. | Y | Y | |
| 3.1.7 Prevent non-privileged users from executing privileged functions and capture the execution of such functions in audit logs. | | Y | This is probably sudo attempts audits. Full commands can be logged in at the cost of extra load for the servers. |
| 3.1.8 Limit unsuccessful login attempts. | Y | Y | Web Services such as love, foreman, ipa console, nublado, etc. may need rate limiting. We don't use passwords in ssh hosts, only ssh keys (which many consider more secure). We are not aware of a retry limit for ssh-key access; an appropriate extra level of security would be to not use the default port 22. However, we do limit attempts to 6 with a block of 600 minutes, which will effectively block failed SUDO logins. |
| 3.1.9 Provide privacy and security notices consistent with applicable CUI rules. | N | Y | Check login notices etc. A login banner can be displayed upon login |
| 3.1.10 Use session lock with pattern-hiding displays to prevent access and viewing of data after a period of inactivity. | Y | Y | This is our policy. |
| 3.1.11 Terminate (automatically) a user session after a defined condition. | Y | Y | ssh sessions are generally not limited on hosts but VPN will timeout daily; some network equipment has timeouts set; |
| 3.1.12 Monitor and control remote access sessions. | N | Y | We currently check who and from where is connecting. |
| 3.1.13 Employ cryptographic mechanisms to protect the confidentiality of remote access sessions. | Y | Y | VPN is in use |
| 3.1.14 Route remote access via managed access control points. | Y | Y | Bastion nodes – IPSec routers now in place also. |
| 3.1.15 Authorize remote execution of privileged commands and remote access to security-relevant information. | Y | Y | |
| 3.1.16 Authorize wireless access prior to allowing such connections. | Y | Y | All devices attaching in Chile need to be registered by Mac address. |
| 3.1.17 Protect wireless access using authentication and encryption. | Y | Y | |
| 3.1.18 Control connection of mobile devices. | Y | Y | In the sense there is no open wifi, and on the summit devices must be registered. |
| 3.1.19 Encrypt CUI on mobile devices and mobile computing platforms. ²³ | Y | Y | Data will not exist on mobile devices - in the case where an image may exist on say commissioning team laptop we will have disk encryption enabled. |
| 3.1.20 Verify and control/limit connections to and use of external systems. | Y | Y | This implies vetting of devices that connect to the control network - we use mac address for laptops and personal mobile phones can not connect to the control network. We also have a separation with the LHN Service Set Identifier (SSID) and Virtual Local Area Network (VLAN)s. |
| 3.1.21 Limit use of portable storage devices on external systems. | N | Y | Can be rolled out with puppet but there are some servers that need usb. |
| 3.1.22 Control CUI posted or processed on publicly accessible systems. | Y | Y | We do not intend to post images on publicly accessible systems. (DMTN-286) |
| 3.2 AWARENESS AND TRAINING | | | |
| 3.2.1 Ensure that managers, systems administrators, and users of organizational systems are made aware of the security risks associated with their activities and of the applicable policies, standards, and procedures related to the security of those systems. | Y | Y | |
| 3.2.2 Ensure that personnel are trained to carry out their assigned information security-related duties and responsibilities. | P | Y | OUO trainign at SLAC |
| 3.2.3 Provide security awareness training on recognizing and reporting potential indicators of insider threat. | Y | Y | We would like to do more here like capture flag exercises for developers or blue/red teams events |
| 3.3 AUDIT AND ACCOUNTABILITY | | | |
| 3.3.1 Create and retain system audit logs and records to the extent needed to enable the monitoring, analysis, investigation, and reporting of unlawful or unauthorized system activity. | Y | Y | |
| 3.3.2 Ensure that the actions of individual system users can be uniquely traced to those users, so they can be held accountable for their actions. | Y | Y | |
| 3.3.3 Review and update logged events. | P | Y | We may look for a third party contract for this. |

| | | | |
|--|---|---|---|
| 3.3.4 Alert in the event of an audit logging process failure. | N | Y | |
| 3.3.5 Correlate audit record review, analysis, and reporting processes for investigation and response to indications of unlawful, unauthorized, suspicious, or unusual activity. | N | Y | Again shall look for third party contract for this |
| 3.3.6 Provide audit record reduction and report generation to support on-demand analysis and reporting. | Y | Y | Observability system |
| 3.3.7 Provide a system capability that compares and synchronizes internal system clocks with an authoritative source to generate timestamps for audit records. | Y | Y | |
| 3.3.8 Protect audit information and audit logging tools from unauthorized access, modification, and deletion. | Y | Y | |
| 3.3.9 Limit management of audit logging functionality to a subset of privileged users. | Y | Y | |
| 3.4 CONFIGURATION MANAGEMENT | | | |
| 3.4.1 Establish and maintain baseline configurations and inventories of organizational systems (including hardware, software, firmware, and documentation) throughout the respective system development life cycles. | Y | Y | We use mainly infrastructure as code approaches so the software is well tracked. IT inventory all the hardware. |
| 3.4.2 Establish and enforce security configuration settings for information technology products employed in organizational systems. | Y | Y | |
| 3.4.3 Track, review, approve or disapprove, and log changes to organizational systems. | Y | Y | We have CCBs and code change process in place which also cover the infrastructure as code. |
| 3.4.4 Analyze the security impact of changes prior to implementation. | Y | Y | |
| 3.4.5 Define, document, approve, and enforce physical and logical access restrictions associated with changes to organizational systems. | Y | Y | |
| 3.4.6 Employ the principle of least functionality by configuring organizational systems to provide only essential capabilities. | Y | Y | |
| 3.4.7 Restrict, disable, or prevent the use of nonessential programs, functions, ports, protocols, and services. | Y | Y | We get a lot of this by mainly containerizing the applications and having users work within deployed containers. |
| 3.4.8 Apply deny-by-exception (blacklisting) policy to prevent the use of unauthorized software or deny-all, permit-by-exception (whitelisting) policy to allow the execution of authorized software. | Y | Y | SUDO lists to restrict access so users can not install applicaitons on the summit nor in SLAC (outside a container). |
| 3.4.9 Control and monitor user-installed software. | Y | Y | |
| 3.5 IDENTIFICATION AND AUTHENTICATION | | | |
| 3.5.1 Identify system users, processes acting on behalf of users, and devices. | Y | Y | |
| 3.5.2 Authenticate (or verify) the identities of users, processes, or devices, as a prerequisite to allowing access to organizational systems. | Y | Y | |
| 3.5.3 Use multifactor authentication for local and network access to privileged accounts and for network access to non-privileged accounts. | Y | Y | |
| 3.5.4 Employ replay-resistant authentication mechanisms for network access to privileged and non-privileged accounts. | Y | Y | |
| 3.5.5 Prevent reuse of identifiers for a defined period. | N | Y | |
| 3.5.6 Disable identifiers after a defined period of inactivity. | Y | Y | |
| 3.5.7 Enforce a minimum password complexity and change of characters when new passwords are created. | Y | Y | |
| 3.5.8 Prohibit password reuse for a specified number of generations. | Y | Y | |
| 3.5.9 Allow temporary password use for system logons with an immediate change to a permanent password. | Y | Y | |
| 3.5.10 Store and transmit only cryptographically-protected passwords. | Y | Y | |
| 3.5.11 Obscure feedback of authentication information. | Y | Y | |
| 3.6 INCIDENT RESPONSE | | | |
| 3.6.1 Establish an operational incident-handling capability for organizational systems that includes preparation, detection, analysis, containment, recovery, and user response activities. | Y | Y | AURA have insurance which covers this. But we really should have a contract to look over logs etc. to note when we are hit. |
| 3.6.2 Track, document, and report incidents to designated officials and/or authorities both internal and external to the organization. | Y | Y | |
| 3.6.3 Test the organizational incident response capability. | Y | Y | This was done at least with the PEN testing - which we shall repeat. |
| 3.7 MAINTENANCE | | | |
| 3.7.1 Perform maintenance on organizational systems. | Y | Y | |
| 3.7.2 Provide controls on the tools, techniques, mechanisms, and personnel used to conduct system maintenance. | Y | Y | |
| 3.7.3 Ensure equipment removed for off-site maintenance is sanitized of any CUI. | Y | Y | |
| 3.7.4 Check media containing diagnostic and test programs for malicious code before the media are used in organizational systems. | Y | Y | |
| 3.7.5 Require multifactor authentication to establish nonlocal maintenance sessions via external network connections and terminate such connections when nonlocal maintenance is complete. | Y | Y | |
| 3.7.6 Supervise the maintenance activities of maintenance personnel without required access authorization. | Y | Y | |
| 3.8 MEDIA PROTECTION | | | |

| | | | |
|---|---|---|--|
| 3.8.1 Protect (i.e., physically control and securely store) system media containing CUI, both paper and digital. | Y | Y | Pixel Zone and Embago Rack |
| 3.8.2 Limit access to CUI on system media to authorized users. | Y | Y | Pixel Zone and Embago Rack |
| 3.8.3 Sanitize or destroy system media containing CUI before disposal or release for reuse. | Y | Y | |
| 3.8.4 Mark media with necessary CUI markings and distribution limitations. | N | Y | We decline to label rooms and machines according to https://www.archives.gov/files/cui/20161206-cui-marking-handbook-v1-1.pdf |
| 3.8.5 Control access to media containing CUI and maintain accountability for media during transport outside of controlled areas. | Y | Y | |
| 3.8.6 Implement cryptographic mechanisms to protect the confidentiality of CUI stored on digital media during transport unless otherwise protected by alternative physical safeguards. | Y | Y | |
| 3.8.7 Control the use of removable media on system components. | N | Y | Can be rolled out with puppet but there are some servers that need usb. |
| 3.8.8 Prohibit the use of portable storage devices when such devices have no identifiable owner. | Y | Y | |
| 3.8.9 Protect the confidentiality of backup CUI at storage locations. | Y | Y | |
| 3.9 PERSONNEL SECURITY | | | |
| 3.9.1 Screen individuals prior to authorizing access to organizational systems containing CUI. | Y | Y | Only project team members will have access to early images - all are known individuals. This doesn't suggest background security screening and it was also explicitly not required by the agencies in section 2 of the requirements document. |
| 3.9.2 Ensure that organizational systems containing CUI are protected during and after personnel actions such as terminations and transfers. | Y | Y | |
| 3.10 PHYSICAL PROTECTION | | | |
| 3.10.1 Limit physical access to organizational systems, equipment, and the respective operating environments to authorized individuals. | Y | Y | This physical access includes locks on server cabinets and key card access in base. (Contracted for summit) |
| 3.10.2 Protect and monitor the physical facility and support infrastructure for organizational systems. | Y | Y | Security is in place on Cerro Pachon and at the entrance to the mountain - though not only for Rubin so not permanently at the observatory. |
| 3.10.3 Escort visitors and monitor visitor activity. | Y | Y | Actual visitors are escorted on the summit - contractors are considered more like staff. |
| 3.10.4 Maintain audit logs of physical access. | Y | Y | NOIRLab can currently store 80 gigs of data for audit logs of physical access, which will last at least three years - all the equipment being installed is HID and complies with section 889 of the John S. McCain National Defense Authorization Act (NDAA) |
| 3.10.5 Control and manage physical access devices. | Y | Y | |
| 3.10.6 Enforce safeguarding measures for CUI at alternate work sites. | Y | Y | This brings in NIST.800-46 and NIST.800-114. Threat analysis suggested. NAT considered bad. |
| 3.11 RISK ASSESSMENT | | | |
| 3.11.1 Periodically assess the risk to organizational operations (including mission, functions, image, or reputation), organizational assets, and individuals, resulting from the operation of organizational systems and the associated processing, storage, or transmission of CUI. | Y | Y | |
| 3.11.2 Scan for vulnerabilities in organizational systems and applications periodically and when new vulnerabilities affecting those systems and applications are identified. | Y | Y | Third party contract PEN test |
| 3.12 SECURITY ASSESSMENT | | | |
| 3.12.1 Periodically assess the security controls in organizational systems to determine if the controls are effective in their application. | Y | Y | |
| 3.12.2 Develop and implement plans of action designed to correct deficiencies and reduce or eliminate vulnerabilities in organizational systems. | Y | Y | |
| 3.12.3 Monitor security controls on an ongoing basis to ensure the continued effectiveness of the controls. | Y | Y | |
| 3.12.4 Develop, document, and periodically update system security plans that describe system boundaries, system environments of operation, how security requirements are implemented, and the relationships with or connections to other systems. | Y | Y | |
| 3.13 SYSTEM AND COMMUNICATIONS PROTECTION | | | |
| 3.13.1 Monitor, control, and protect communications (i.e., information transmitted or received by organizational systems) at the external boundaries and key internal boundaries of organizational systems. | Y | Y | |
| 3.13.2 Employ architectural designs, software development techniques, and systems engineering principles that promote effective information security within organizational systems. | Y | Y | We can do more here. |
| 3.13.3 Separate user functionality from system management functionality. | Y | Y | Commissioners may be fairly privileged users. |
| 3.13.4 Prevent unauthorized and unintended information transfer via shared system resources. | Y | Y | DMTN-286 and SITCOMTN-076 cover ground rules on this |

| | | | |
|--|---|------------|--|
| 3.13.5 Implement subnetworks for publicly accessible system components that are physically or logically separated from internal networks. | Y | Y | |
| 3.13.6 Deny network communications traffic by default and allow network communications traffic by exception (i.e., deny all, permit by exception). | Y | Y | We may need to bring up iptables on each host |
| 3.13.7 Prevent remote devices from simultaneously establishing non-remote connections with organizational systems and communicating via some other connection to resources in external networks (i.e., split tunneling). | Y | Y | |
| 3.13.8 Implement cryptographic mechanisms to prevent unauthorized disclosure of CUI during transmission unless otherwise protected by alternative physical safeguards. | Y | Y | IPSec and encryption at rest |
| 3.13.9 Terminate network connections associated with communications sessions at the end of the sessions or after a defined period of inactivity. | Y | Y | |
| 3.13.10 Establish and manage cryptographic keys for cryptography employed in organizational systems. | Y | Y | |
| 3.13.11 Employ FIPS-validated cryptography when used to protect the confidentiality of CUI. | Y | Y | |
| 3.13.12 Prohibit remote activation of collaborative computing devices and provide indication of devices in use to users present at the device. | Y | Y | We should take care with the new roaming camera. |
| 3.13.13 Control and monitor the use of mobile code. | Y | Y | Currently we have no mobile code |
| 3.13.14 Control and monitor the use of Voice over Internet Protocol (VoIP) technologies. | N | N | |
| 3.13.15 Protect the authenticity of communications sessions. | Y | Y | |
| 3.13.16 Protect the confidentiality of CUI at rest. | Y | Y | |
| 3.14 SYSTEM AND INFORMATION INTEGRITY | | | |
| 3.14.1 Identify, report, and correct system flaws in a timely manner. | Y | Y | |
| 3.14.2 Provide protection from malicious code at designated locations within organizational systems. | Y | Y | |
| 3.14.3 Monitor system security alerts and advisories and take action in response. | Y | Y | |
| 3.14.4 Update malicious code protection mechanisms when new releases are available. | Y | Y | |
| 3.14.5 Perform periodic scans of organizational systems and real-time scans of files from external sources as files are downloaded, opened, or executed. | Y | Y | |
| 3.14.6 Monitor organizational systems, including inbound and outbound communications traffic, to detect attacks and indicators of potential attacks. | Y | Y | |
| Total requirements | | 108 | |
| Total Rubin Intends to comply with | | 107 | |
| Total Rubin Complies with in 2024 | | 95 | |

C Glossary

Alert A packet of information for each source detected with signal-to-noise ratio > 5 in a difference image by Alert Production, containing measurement and characterization parameters based on the past 12 months of LSST observations plus small cutouts of the single-visit, template, and difference images, distributed via the internet.

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.

AVS Alert Vetting System.

BTS Base (La Serena) Test Stand.

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)..

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.

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.

DMTN DM Technical Note.

DRP Data Release Production.

Enclave Individually defined portions of the computational resources at the Summit, Base,

NCSA, and Satellite Facilities, such as the Prompt Enclave, the Archive Enclave, etc..

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.

LHN long haul network.

LLNL Lawrence Livermore National Laboratory.

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

OGA Other Government Agencies.

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

PMO Project Management Office.

Project Management Office the work element responsible for achieving the project's objectives.

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.

RDM Rubin Data Management.

RDP Rubin Data Production (Obsolete use RDM).

ROE Royal Observatory Edinburgh.

RPF Rubin system PerFormance.

RSP Rubin Science Platform.

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..

SSID Service Set Identifier.

STFC UK Science and Technology Facilities Council.

TMA Telescope Mount Assembly.

TTS Tucson Test Stand.

UK United Kingdom.

US United States.

USDF United States Data Facility.

VLAN Virtual Local Area Network.

VPN virtual private network.

D Actions (temporary)

This section should disappear when these are all done.

| Id | Actionee | Due Date | Action |
|-----|----------|-----------|--|
| S-1 | JC | June 2024 | Reference for security of archive services |