Data Facilities Transition Plan

Work Package 2.1, Deliverable D2.1.5

Project Acronym	LUSC-B
Project Title	UK Involvement in the Legacy Survey of Space and Time
Document Number	LUSC-B-28

Submission date	25/MAR/22
Version	1.2
Status	Final
Author(s) inc. institutional affiliation	George Beckett (Edinburgh); Pete Clarke (Edinburgh)
Reviewer(s)	Aprajita Verma (Oxford); Aaron Watkins (Hertfordshire)

Dissemination level	
Public	

Version History

Version	Date	Comments, Changes, Status	Authors, contributors, reviewers
0.1	27/OCT/21	Initial skeleton and content	PC
0.2	28/OCT/21	Minor changes	MGB
0.3	5/NOV/21	Extended introduction and motivation	MGB
0.4	12/NOV/21	Further revisions	MGB, PC
0.5	18/NOV/21	Added forward look	MGB
0.6	17/JAN/22	Further enhancements for Exec Group	MGB
1.0	25/MAR/22	Updated and finalised	MGB
1.1	2/JUN/22	Updated in light of reviewer feedback	
1.2	14/SEP/22	Dissemination level updated to public following confirmation from MGB	TMS

Table of Contents

1 Introduction

This document describes the work that needs to be done within LSST:UK, to ensure that the UK Consortium can deliver effectively on the in-kind commitment to complete 25% of the Rubin Observatory Data Release Processing (DRP) workload. It also documents the required roles and resource levels, during Rubin Operations, based on guidance provided by the Rubin Data Management team, US and French Data Facilities Coordinators, and contacts within the IRIS programme.

We begin by describing the context of this work (Section 2) and summarising initial activities that have been possible with in-kind effort from GridPP, working with peers in the USA and France, during the latter part of 2021 (Section 3).

We then propose a high-level work plan for DRP commissioning contributions that are needed, from LSST:UK, in the remaining 12 months or so of Phase B (Section 4). The work plan is aligned with a Rubin Observatory in-preparation document [RTN-021] and assumes telescope operations from April 2024.

We then describe the anticipated resource requirements (skills and effort levels) for during Rubin Operations (Phase C and D), plus the immediate resource requirements to complete the proposed Phase B work plan (Section 5).

1.1 Glossary of Acronyms

DR	Data Release
DRP	Data Release Processing
IDF	Interim Data Facility
LSST	Legacy Survey of Space and Time
LSST-Comp-Team	The LSST global compute team led from the USA. Comprising those responsible to construct the DRP processing chain and those recently appointed to run it at SLAC
SE	Storage Element (dCache, or DPM or xRootD)
RSE	Rucio Storage Element (means an RE register with the Rucio server)
CE	Compute Element (either ARC-CE or HTCondor)
CVMFS	Cern-VM File System
FTS	File Transfer Service
NREN	National Research and Education Network
VOMS	Virtual Organisation Membership Service

2 Overview

The Vera C. Rubin Observatory, which is a US-led facility, under construction in Northern Chile, plans to undertake a 10-year survey of the southern sky, referred to as the Legacy Survey of Space and Time (LSST). It will conduct the most ambitious optical sky survey yet planned, imaging the whole visible portion of the southern sky twice each week for 10 years from early 2024 (start date to be confirmed).

It will visit each field more than 800 times during that decade, with exposures taken through six filters (ugrizy) spanning the visible portion of the spectrum, from 320nm to 1,050nm, and is expected to observe and catalogue around 35 billion astronomical sources (galaxies and stars, and solar-system objects). The integrated depth of the final stacked images (r~27.5 mag, 5 point-source) will be five magnitudes deeper than the Sloan Digital Sky Survey, and the survey dataset will constitute an unprecedented multi-colour movie of the Universe, which will be used to study astrophysical phenomena all the way from the closest asteroids to the nature of the dark energy.

The LSST survey will be delivered to the science community as a collection of products, including a (typically) annual update to the survey, called a Data Release (DR). Each Data Release is a substantial, multi-Petabyte resource consisting of various science-ready datasets:

- Processed versions of telescope images, corrected for instrumental signatures, etc.
- Deep Coadds which are created by stacking multiple images of the same region of the sky to create an image of increased sensitivity and detail.
- Various science catalogues (relational-database tables) that document detected astronomical objects along with standard measurements on those objects (location, colour, flux, etc.).
- Various ancillary and intermediate products, to support different science use cases with different requirements, and to support downstream processing and the generation of derived datasets.

DRP is a substantial piece of work involving significant computational and storage infrastructure. It is a multi-stage process in which images progress through various processes. There are science-relevant decisions needed for many of these processes, meaning there may be a need to vary the configuration (branches in the pipeline) used at specific points, to support different science use cases. Thus, in addition to the end products noted above, it is expected that some (at least) of the intermediate products will be preserved (or should be readily reproducible).

DRP is undertaken using a large and bespoke suite of software commonly referred to as the LSST Science Pipeline or LSST Stack. This is an Open-Source project, which is rapidly developing at the time of writing, consisting of a collection of Python libraries and command-line tools backed by C++ kernels for computationally intensive elements.

DRP progress is recorded in a database repository called the Data Butler. This is intended to track the locations and contents of datasets as they are processed. It is heavily used by the LSST Stack and, upon the completion of a DR, becomes the primary metadata enabling science users to interact with the survey images and coadds.

The overall responsibility for DRP is with the Rubin Observatory. They are responsible for developing the pipeline software, selecting appropriate middleware on which to deploy the processing workflow, and for undertaking quality assurance on the output data products. The UK (and French) Data Facilities are responsible for completing their agreed share of the workload, as a follows:

- The Rubin Observatory will complete 25% of the processing at the United States Data Facility (USDF) which will be hosted at the SLAC National Accelerator Laboratory in California.
- The Institut national de physique nucléaire et de physique des particules (IN2P3) will complete 50% of the processing at their facilities in Lyon, France.
- The LSST:UK Consortium will complete 25% of the processing on the UK IRIS shared infrastructure.

At a high-level, DRP involves the following workflow:

- Raw images, captured at the telescope are transferred to the USDF over a dedicated network link, with very low latency (target is 7 seconds from time of shutter close until it arrives at SLAC).
- The portion of the raw images to be processed in France and the UK are then transmitted (on a timescale to be determined) on to those facilities, over the public Internet, along with calibration images and other ancillary products.
- Once each facility has the raw image data it is to process, it may proceed with processing. Processing can effectively be completed independently at each facility, though there will be a requirement for the USDF to have an overview of fine-grain progress at each facility, plus there should be a capability to reassign processing work on the granularity of days/ weeks, in response to processing problems.
  It is likely that some processing can be undertaken as data arrives: that is, before the complete set of raw images has been received.
- Once processed, output data is transferred back to the USDF (again, over the public Internet) to be assembled into a Data Release.
- Processed data needs to be validated (by Rubin Observatory) before it is confirmed to be ready for publication. It is unclear whether validation work will be undertaken when data is returned to the USDF or can be done earlier, at each data facility. This validation will be done by the Rubin Verification and Validation Team.
- It is likely that some later (hopefully, not computationally intensive) steps aggregate data from across the whole DR. This may or may not be possible with DR products distributed across the three sites, so may have to be done at the USDF.
- It is also intended that a full copy of all DR-related output data will be held at the French DF. This may serve as an online replica, for disaster recovery, though that may/ may not be the primary motivation for doing this.
- The UK DF requires a full copy of the *output* data from a DRP campaign, in order that it can serve science-ready data to a subset of the Rubin science community from a UK-based Data Access Centre. It is not anticipated, however, that the UK DF will need or want all the raw images: just those images that are to be processed in the UK. Details are to be finalised, but the working assumption of DM staff is that the UK DF will process one quarter of the observed sky.

It is intended that each DRP campaign will reprocess all images to date (that is, going back to the beginning of the survey). This is required to ensure consistent processing of all images (the processing stack will evolve throughout the survey), though does mean that the volume of DRP-related computing and storage will grow year on year.

The preparation of a Data Release – that is, a DRP campaign – should be completed within six months of the end of observing period. This includes assembly of complete datasets in both the US DF and French DF. It may also include the distribution of products to Independent Data Access Centres, though that should also be done in a timely manner, so as not to delay the release of a DR to the science community.

The decision to distribute DRP across three facilities increases the complexity of the task significantly. The advantage is that it does significantly reduce the financial and infrastructure burden for the Rubin Observatory as well as providing resilience (i.e., no single point of failure at SLAC). Compared to a single-DF approach, the following additional challenges are envisaged:

- **Software synchronisation** – to ensure that each site undertakes precisely the same processing, involving not just the same versions of the LSST stack, but also the same (or guaranteed compatible) versions of third-party supporting libraries, tools, and supporting services.
- **Distributed campaign management** – is required to ensure that the many tasks involved to turn a collection of raw images into a full and complete Data Release is completed in an efficient, reliable, and timely manner.
- **Data movement and staging** – significant additional data movement is required to stage input files to the relevant Data Facility and receive back outputs into the *master archives* held at the USDF and in France. This movement needs to happen in a timely manner, even though it is taking place on public Internet capacity with consequent variations in performance and capacity.
- **Campaign monitoring** – to ensure that each facility progresses as expected through its portion of the DRP campaign, but also to provide contingencies if one or more of the facilities falls behind with processing (maybe because of infrastructure issues or because of unforeseen complexities in image data or DRP).
- **Data aggregation** – the output data produced at each DF needs to be assembled into a consistent Data Release (along with required intermediate products). A Data Butler instance needs to be created that captures the provenance of the processing campaign, in a way that is indistinguishable from a DRP undertaken at a single facility.
- **Data publication** – possibly an advantage of distributed DRP, rather than a challenge but, at the end of each DRP campaign, the Data Release (in part or whole) needs to be distributed to a number (around 10) Data Access Centres, internationally, from where it is made available to the Rubin Science Community.
- **Authentication and authorisation** – despite the heterogeneity of the infrastructure, across the three facilities, it is necessary for DRP staff from across the facilities to contribute to the DRP campaign – e.g., for USDF staff to have seamless access to data products being created in the UK and even to intervene in processing, should the need arise.
- **Accounting** – the three-site configuration is underpinned by an in-kind agreement which translates contributions to the DRP (in France and the UK) into data rights for the relevant France-based and UK-based science communities. It is therefore likely that the DFs will need to be able to record and present evidence that they have contributed resources in line with the intent of DRP.

DRP has been observed to share several analogues with large-scale data processing undertaken for LHC experiments (e.g., ATLAS). That observation has motivated consideration of tools, technologies, and processes from the LHC community.

Several specific tools and technologies have been highlighted, though more may be required:

- **Rucio:** A data management system from the ATLAS experiment at the LHC, used widely in particle physics (and possibly SKA). This is a policy-driven system for making (multiple) copies of data at any number sites globally, whilst maintain a coherent global catalogue. Sites are registered as a Rucio Storage element (RSEs). Typically, data is identified, moved, and accessed by a URL-like string.
- **PanDA**: large-scale, distributed workflow orchestration is likely to fulfil many of the requirements of processing campaign management.

- **FTS:** The File Transfer Service that takes care of reliable (used in the sense of TCP reliable packet delivery) end-to-end file transfer, including third party file transfer. FTS guarantees a file will be delivered, regardless of any short-term failures in any individual copy.
- **VOMS** (a software service for managing credentials for Virtual Organisations, http://italiangrid.github.io/voms/) is likely to be able to provide a common authentication/ authorisation platform
- **CVMFS** (the Cern-VM File System, https://cernvm.cern.ch/fs/) is likely to be able to automate the distribution and curation of software suites for DRP.
- Various other underlying components to implement what is known as a "Storage Element" (SE). Typical SEs are dCache or xRootD , but may be other things.
- Various other components to implement what is known as a Compute Element (CE). Typical are ARC-CE, HTCondor-CE.

A likely advantage of these components is that they are known to work, as they have been in full-scale production for decades at the LHC. The question is whether they are appropriate for the specific LSST data context and anticipated ways of working.

Complementing this, several pre-existing Rubin Observatory technologies are considered fundamental to DRP and hence need to be incorporated into the three-site DRP, including:

- **LSST Stack** – as noted above
- **Data Butler** – as noted above
- **Qserv** – a bespoke, distributed-memory, relational database into which the science catalogues are ingested for presentation to science users. It is understood that Qserv is not integral to DRP, though the outputs of DRP need to be ingested into Qserv, which is itself a challenging process.
- **Rubin Science Platform (RSP)** – a mix of client software and services which is used to interrogate raw, intermediate, and processed data. Again, RSP is not integral to DRP. However, it is understood that Science Validation teams may wish to use RSP to interrogate candidate DR products and even in-progress DRP campaign data. The primary intermediaries between the RSP and a DR are Qserv and the Data Butler.
- **Kafka** – the stream processing platform that is used to serve a complement to DRP, called Alert Processing, which will serve pseudo-real-time information on transient astronomical activity to the science community. The Kafka platform is independent of DRP and likely only to involve the USDF. Therefore, any overlap between AP and DRP should be limited to the USDF.

Where these technologies have been developed in-house, by Rubin staff, it is likely they were not envisaged to be employed for a multi-site campaign.

3 Topical Activities

3.1 Coordination

To support and inform early experiments, a series of Data Facility Workshops has been scheduled, involving representatives from the Rubin Construction Project, Rubin Operations, the three Data Facilities, and technology experts from the high-energy physics community.

The first two workshop were held in June 2021 and January 2022, respectively. The material from these workshops is available on the Rubin Confluence site (login required). At the time of writing, a third workshop is in preparation, and will run during 6th—8th April 2022. Further workshops are expected in due course.

The outputs from these workshops (including the activities that are presented) are driving the development of two Project technical notes:

- RTN-021: Data Facilities Transition Plan
- DMTN-213: Multi-Site Data Release Processing Using PanDA and Rucio

—which represent the topical understanding of how the data facilities will complete each DRP campaign.

In addition to this, there are several, regular meetings to deal with day-to-day management and activity planning:

- The three Data Facilities Liaisons meet every two weeks
- A technical team, working on early Rucio and PanDA experiments also meets on a fortnightly basis.

3.2 Rubin Observatory Planning

RTN-021 outlines a series of escalating DRP experiments, with the following provisional timeline:

- **DP0.2 (July 2021—September 2022)** – a DRP experiment to reprocess data from DESC DC2 using the LSST Pipeline software. This DRP experiment only involves a single data facility (the Interim Data Facility on Google Cloud) but is a vital source of information for the three-site DRP capability, as it involves all the fundamental elements of DRP.
- **DP0.3 (to be confirmed)** – a re-run of the DRP workflow used to create DP07.2, though running across the three Data Facilities rather than on the IDF. It will be worthwhile to compare the properties of the DP0.2 and DP0.3 campaigns, considering efficiency, resilience, and effectiveness, for example. This experiment requires agreement from all three Data Facilities before it can be confirmed.
- **DP1 (April 2023—October 2023)** – data from the Commissioning Camera (ComCam), installed at the Observatory, will be processed twice: first, at the US DF; and then, across the three Data Facilities.
- **DP2 (April 2024—October 2024)** – Commissioning Data from the full LSST Camera, installed at the Observatory, will be processed across the three Data Facilities, aiming to mimic the conditions and timeline of a production DRP campaign as closely as possible.

RTN-021 assumes that Rubin Observatory operations will begin in April 2024, with the raw images that will constitute the first, in-production, Data Release (DR1) being ready to process in October 2024.

Considering this timeline within the LSST:UK roadmap (that is, the phases of funding), we anticipate:

- DP0.3 should be possible within Phase B.
- DP1 and DP2 will be completed within Phase C.

Although the existence and timing for DP0.3 has not yet been confirmed, some experiments have begun involving candidate technologies and infrastructure, with in-kind, UK contributions from GridPP, as described below.

3.3 DRP Experimentation

At the time of writing, a team is developing the way-of-working to use the candidate distributed-data management system, Rucio. The team involves:

- Rubin, DOE and NSF staff in the USA, coordinated by Michelle Butler (NCSA)
- Staff from the IN2P3 laboratory in Lyon, France.
- GridPP staff in the UK (funded initially by GridPP).

Initial Rucio work, underway at the time of writing, addresses:

- Setup of a multi-site Rucio infrastructure at SLAC, RAL, Lancaster, and IN2P3
- Undertaking a series of increasingly sophisticated data-transfer experiments:
    - Initial file registration
    - One-directional transfers
    - Bidirectional transfers between pairs of Data Facility sites
- Working through the relationship and interaction of Rucio with the Data Butler

Alongside Rucio, work has also begun to create a PanDA testbed, with execution sites in the US, France, and UK, collocated with the Rucio testbed noted above. At the time of writing, experimentation is at the *hello-world* scale, which yields useful information about authentication and authorisation matters, technology interfaces, and fine-grain job management.

3.4 Wide-area Networking

The international NREN (National Research and Education Network) infrastructure is already excellent.

- In the USA, it is taken care of by ESNET (DOE Energy Sciences Network) which connects all DOE labs (including SLAC, FNAL and BNL) at very high capacity.
- ESNET has over 300 Gbit/s connectivity to Europe to support LHC and other science. Connections from the USA *peer* in Europe, in London, Amsterdam and Geneva.
- The European network, Geant, also has multi-100 Gbit/s links to ESNET.
- From a UK point of view ESNET, Geant and JANET all peer in London.
- RAL is connected to JANET at 2x 200Gbits/s.

An active dialogue with the NREN stakeholders has begun and we therefore believe that the NREN infrastructure is more than sufficient for the needs of Rubin.

Several GridPP members are in regular contact with JISC (operators of JANET) and they attend as necessary the LSST NET group monthly meeting.

4 Phase B (and Early Phase C) Work Plan

Given the substantial task to prepare for DRP in Operations, UK preparations need to begin before the end of Phase B, building on the early GridPP experiments and discussions noted above, and ensuring UK interests and requirements are fully represented in the development of the DRP operational plan, so that the UK can fulfil its commitment to DRP in an effective and high-quality manner.

4.1 Aims

Prior to the beginning of Rubin Operations, the three involved Data Facilities need to define, validate, and implement a distributed infrastructure (software and hardware) on which DRP can be completed within the constraints laid out by the Rubin Project, including:

- Processing each observing campaign within six months of the end of the campaign.
- Each Data Facility completing their agreed share of the computation load of a DRP.
- Delivering Data Release Products to Rubin and Independent Data Access Centres.

By around the end of LSST:UK Phase B, a distributed infrastructure should be both documented and implemented, including at least the following details:

- A **design report** detailing a proposal for how DRP should operation, for the consideration of the Rubin Observatory Directorate, in sufficient detail to be reproducible based solely on said design. The proposal should include details of the Rubin and third-party infrastructure that is required, a timeline of how a DRP campaign would be completed (highlighting key milestones and progress points) and listing risks to the timely and accurate completion of DRP along with reasonable contingencies.
- **Resourcing estimates**, for staff, etc., to operate the infrastructure (in the run-up to) and during telescope operations (that is from the time of writing through until October 2035).
- **A record of supporting technologies and services** (those that are outside of the control of the Rubin Observatory) should be detailed, along with information on how the Data Facilities can ensure that the supporting technologies and services will continue to provide the required functionality for DRP.
- **The DP0.3 Demonstrator** should be completed and documented, adhering as closely as possible to the conditions of telescope operations (with any deviations documented and justified).

These outputs should allow the three Data Facilities to run full-scale validation runs on DP1 and to confidently take the lead role in preparing DP2.

4.2 Work Package Breakdown

The basic approach is to partake in the series of increasingly realistic tests, culminating in a full DRP demonstration, based on Rubin Observatory Data Preview 1 (data from the Commissioning Camera). The precise form of this demonstration is still to be defined (it will be documented in the WP4.1 design report), though would include a series of measures and success criteria, agreed in advance, to ensure the demonstration fulfilled the aims of the project.

The number and details of these tests is still to be confirmed though, in line with the current draft of RTN-021, it is anticipated to include at least a hypothetical run, in which key stakeholders would talk through the process from beginning to end, identifying key steps, potential bottlenecks, technology limitations and omissions, etc., followed by two actual processing campaigns for DP0.3 and DP1. It is possible that DP0.3 will be undertaken multiple times, exploiting different infrastructure configurations and technology choices.

The full DRP demonstration (reprocessing of DP1) would need to be completed at least six months before the beginning of operations (that is, by October 2023), in time for a full, *dress rehearsal* of DRP using DP2.

More details of the implementation are described below.

4.2.1 On-boarding of Staff

Staff in the UK, who will work on this activity, should be on-boarded into the LSST:UK Science Centre team, following the induction process laid out in the LSST:UK Project Manual. On-boarding is a relatively lightweight process, which can be completed in around one week.

In addition, it is likely that some or all WP4 staff will need to be onboarded to the Rubin project. This will provide access to project resources, such as internal-documentation, communications channels, and project-management tools. Again, there is a process for doing this, which can be completed in around one week.

4.2.2 Hypothetical Run-through of Demonstrator

We will schedule a workshop (estimated to be 2—3 hours) involving key stakeholders—from the three Data Facilities, the Rubin DM team and key contributors from DOE Labs (to PanDA and Rucio)—to talk through the steps involved in a DRP campaign, aiming to identify the sequence of events involved, pre-requisites, success criteria, etc. It is likely that early experiences from DP0.2 would be crucial inputs into this workshop.

For example, the workshop would be expected to confirm:

- what raw data would be the subject of each Demonstrator
- where the raw data would be sourced from (and how similar would that source be to the delivery from the telescope)
- how would the work assigned to individual DFs be decided
- how would necessary input data be transferred to each DF
- how long would it take
- how would one know that data had been transferred
- what technology would be used to complete the transfer
- who would action the transfer and what level of manual involvement would be required
- what ingestion work would be required at each DF to prepare for processing
- how would pipeline software be deployed to and validated at each DF

—and so on.

The proceedings of the workshop would be documented, and an action plan for the subsequent demonstrators would be agreed.

4.2.3 Intermediate Experiments

Based on the action plan, a set of experiments would be defined and completed, towards preparation for a final demonstrator. The number and form of these experiments would be agreed during the workshop, with options for:

- running partial demonstrators (e.g., subsets of the DRP campaign);
- technology evaluations, to demonstrate one can reach target performance levels, or to compare different technology options

- running lean versions of the final demonstrator, using subsets of the data or proxy datasets.

The initial UK contributions, already underway and described earlier, which have taken advantage of opportunities to engage with Construction Project experiments and use in-kind effort from GridPP staff should be factored into this work package as intermediate experiments.

4.2.4 First Full Demonstrator (DP0.3)

A key milestone will be a first full demonstration, in which the agreed input data would be processed, across the three data facilities, and then reassembled and ingested into the RSP (to be confirmed) for validation.

4.2.5 DRP Rehearsal #1 (DP1)

Having completed the preparatory experiments and DP0.3 demonstrator, the team would contribute towards preparations for a larger-scale rehearsal, engaging operational staff, to (re)process the raw data from ComCam for DP1, in pseudo-real-time as one continuous campaign, with all the various infrastructure technologies and processes operating in unison, as they are envisaged for DPR in operations.

5 UK DRP Staffing and Resource Requirements

5.1 DRP Operations (Phase C)

To fulfil the UK's in-kind commitment to DRP, during Rubin Operations and for the later processing rehearsals (that is, DP1 and DP2), will require a UK-based team (the UK Data Release Processing Team, with experts in DRP (compute-oriented astronomers embedded in Rubin) as well as experts in distributed compute and data management (that is, staff at IRIS sites maintaining the compute fabric and software infrastructure). These two groups of experts must work in synergy.

The Observatory has confirmed they anticipate each DF will require at least 5 FTE of technical expertise to undertake DRP and, based on the candidate technologies and guidance from the Rubin Data Facilities Coordinator and from the UK GridPP leadership, we have defined the activities that will be required below.

5.1.1 DRP Data Management

Each DRP campaign involves a complicated sequence of workflow steps, each consuming and producing significant volumes of data. To ensure a DRP campaign can proceed at a sufficient pace, in a multi-site setting, it is crucial that the workflow sequence is optimised to reduce the requirement for data movement and that required data movement is undertaken in a timely manner to avoid interruption. Technology can partly address this requirement: the preferred data-management platform, Rucio, uses a policy-based approach to organise and distribute data to where it is needed. However, the complexity of the campaign, plus the impact of external factors, such as variations to the observing schedule, means data management cannot be address by technology alone. The DRP Data Management team will be responsible for overseeing the UK element of each campaign, working closely with peers at SLAC and IN2P3, and responding to deviations from the intended campaign flow. The team will also ensure that input data (most notably, raw observations and calibration data) which is needed for multiple DRP campaigns is archived in the UK ready to be recalled when needed.

5.1.2 DRP Pipeline Operations

As noted above, PanDA will be used as the workload manager for DRP. Each DR will consist of thousands of individual jobs, which will be distributed between the three Data Facilities by a campaign manager running in the USDF. In the UK, jobs will be completed by IRIS sites running GridPP Compute Element (CE) services, alongside LHC workloads.

A Processing Scientist will be responsible for ensuring UK (IRIS) compute resources are available and correctly configured to fulfil the UK contribution to each DRP campaign. They will work with the Rubin Infrastructure Group to harmonise compute-resource provision across the three facilities, to complete pre-campaign testing, and to track progress of the UK share of the workload. The Processing Scientist will be responsible for day-to-day operations to ensure problems are detected and resolved quickly and to ensure sufficient compute resources are available to keep to the campaign schedule. They will also maintain regular and frequent communications with the Rubin Workflow Management Engineer and Processing Scientists at the US and French facilities, about campaign progress. The Processing Scientist will also liaise with the Production Scientist (see WP4.4) and other members of V&V, to arrange appropriate access to logging and outputs from in-progress and recently completed computations. Pipeline Operations will also respond to occasional issues at UK processing sites - for example, site outages or infrastructure-configuration problems - and will work with the Rubin to remediate any consequent impacts on the campaign.

5.1.3 Scientific Validation

The UK DRP team will contribute a Production Scientist as part of the Rubin Verification and Validation team, to ensure the DRP outputs are scientifically correct and complete, to address and remediate any scientific issues with a processing campaign, in a timely manner, and to help troubleshoot algorithm/pipeline-level problems. The Production Scientist will require a thorough understanding of the Science Pipeline and have extensive experience of optical, survey astronomy. They will work closely with the Rubin Algorithms and Pipelines team, include peer roles in the French and US Data Facilities, and the Calibration Support team at the Observatory site.

As noted in WP2.3, the RSP is a convenient interface for the V&V team to interrogate the in-progress calculation and, to this end, the Production Scientist will work with the WP4.3 Pipeline Operations team and the WP2.3 Science Platform team to make in-progress image and catalogue products available via an RSP instance promptly. The V&V team will also coordinate any remedial action that is required, during a campaign - for example, updating configuration parameters, coordinating the re-processing of any problematic datasets and, in rare cases, overseeing the roll-out of live updates to the Science Pipeline – and the Production Scientist will have a significant role in such action.

5.1.4 Infrastructure Support and Troubleshooting

Experience from other large-scale distributing computing campaigns (for example, the LHC ATLAS experiment) has highlighted how critical it is to have experienced, on-site staff at facilities, with experience of both DRP application requirements and of the facilities' services and technologies. While LHC technologies deliver significant levels of automation and implicitly ensure homogeneity across the services provided at different sites, the complex and tightly coupled nature of key services such as Rucio and PanDA benefit from on-site staff who can mediate between the application requirements and the service capabilities as well as to represent the needs of the application about routine facility operations, such as maintenance and infrastructure upgrades.

The on-site staff have a significant role not only to ensure smooth day-to-day operation, but also contributing to the medium-term roadmap for the facility considering DRP forecast requirements. The on-site team will work also closely with the Data Facility Coordinator (described in WP1.5), advising them of local issues that have implications for other data facilities and to help secure access to the scale of resources documented in the UK Data Facility Infrastructure Sizing plan.

5.2 DRP Commissioning Team (Phase B and Year 1 of Phase C)

Prior to full operations, a small DRP Commissioning Team will be sufficient to undertake the work proposed in Section 4. Where possible, this team would be a continuation of the GridPP-funded effort that is engaging in early experiments, described in Section 3, as follows:

- 0.25 FTE for Edinburgh, for the UK Data Facilities Coordinator, to facilitate the commissioning experiments outlined above, liaise with peers at the US and French DFs and in Rubin management structure, represent UK interests in the development of the DRP operational plan.
- 0.5 FTE for RAL, for storage management and support (RAL will provide and support 9 PB of storage for LSST:UK by 2024) to include responsibility for Rucio and Butler commissioning, and responsibility to run all data infrastructure testing as requested by LSST-Comp-Team.
- 1.0 FTE for Lancaster University, supporting DRP testing and commissioning, data challenges etc, supporting use of PanDA, any AAI work, and responsibility to run all UK testing as requested by LSST-Comp-Team. This will also include an LSST-GridPP liaison responsibility. Initial thoughts are that this would naturally be at

Lancaster which has both GridPP and LSST interests and has the PanDA experts already in place.
- 0.5 FTE for Cambridge University, for a DRP software expert embedded in an astronomy group.

Other liaison roles, including networking, have been confirmed to be undertaken by GridPP on behalf of LSST for the near-term future.

6 References

[1] R. Dubois, W. O'Mullane, *RTN-021: Data Facilities Transition Plan*, Rubin Technical Note (in preparation, last updated November 2021)

[2] R. Dubois, W. O'Mullane, *RTN-021: Data Facilities Transition Plan*, Rubin Technical Note (in preparation, last updated November 2021)

[3] K-T Lim, *DMTN-189: Data Facility Specifications*, Rubin Technical Note (April 2021)

[4] K-T Lim, *DMTN-213: Multi-Site Data Release Processing Using PanDA and Rucio*, Rubin Technical Note (in preparation, last updated January 2022)

[5] W. O'Mullane, R. Dubois, M. Butler, K-T Lim, *DMTN-135: DM sizing model and cost plan for construction and operations*, Rubin Technical Note (November 2021)