Title: SRDP Operations Plan	Owner: Lacy	Date: 3/31/2019
Document No. 530-SRDP-035-OPER		Version: 1.1

Science Ready Data Products

Operations Plan
Project 530
Draft

PREPARED BY	**ORGANIZATION**
Mark Lacy	Organization, Title
VERSION RELEASE BY	**ORGANIZATION**
Mark Lacy, Operations Manager	NRAO

APPROVALS	**ORGANIZATION**	**SIGNATURE**
Jeff Kern	NRAO SRDP, Project Director	see SharePoint workflow for approvals
Lewis Ball, Sponsor	NRAO SSR AD	

Change Record

VERSION	DATE	RESPONSIBLE	REASON
1.0	3/4/19	M. Lacy	First version for distribution to stakeholders
1.1	3/11/19	M. Lacy	Revised per stakeholder comments

Title: SRDP Operations Plan	Owner: Lacy	Date: 3/31/2019
Document No. 530-SRDP-035-OPER		Version: 1.1

Table of Contents

Introduction

1 Background

High Level goals, objectives, and a description of the SRDP project can be found in the Project Charter [RD01]. A detailed description of the programmatic project elements is addressed in the Project Management Plan (PMP) [RD02], while the detailed technical elements are addressed in the System Engineering Management Plan (SEMP) [RD03].

The SRDP project (in collaboration with existing observatory management structures) will develop and implement the software tools, scientific heuristics, and operations structures to deliver science quality data products to the NRAO user community, both as it exists now, and the foreseen expansion of the community to include astronomers who do not have a detailed knowledge of radio astronomy.

1.1 Scope of this Document

In this document, we describe the Operations of SRDP, including the plan for Quality Assurance of the products and estimate required personnel and computing resources.

1.2 SRDP in the context of the NRAO/ALMA data management infrastructure

SRDP activities need to be folded into the other NRAO data management activities. These include ALMA data processing operations (pipeline operations are scheduled to move to JAO eventually, but manual processing will continue to be done at the NA ARC). ALMA pipeline testing, CASA testing and visitor support (including remote interactive use) and, currently, processing of VLA Sky Survey (VLASS) data. How these are prioritized is outside the scope of this document.

1.3 Reference Documents

[RD01] SRDP Project Charter 530-SRDP-001-MGMT

[RD02] SRDP Project Management Plan 530-SRDP-003-MGMT

[RD03] SRDP System Engineering Management Plan 530-SRDP-010-MGMT

[RD04] SRDP System Concept Document 530-SRDP-014-MGMT-SRDP_System_Concept

Lacy et al. 2019 (in preparation): The VLA Sky Survey

Project documents may be found on Sharepoint:
https://sharepoint.nrao.edu/pmd/projects/530%20Science%20Ready%20Data%20Products

1.4 Abbreviations and Acronyms

ALMA – the Atacama Large Mm/submm Array
API – Application Program Interface
ARC – ALMA Regional Centre
AWS – Amazon Web Services
DMS – NRAO's Data Management and Software department
JAO – Joint ALMA Observatory
MS – Measurement Set
NA – North America
ngPST – next generation proposal submission tool

QA – Quality Assurance
SB – Scheduling Block
SRDP – Science-Ready Data Products
VLA – the Karl G. Jansky Very Large Array
VLASS – the VLA Sky Survey
WF - workflow

1.5 Assumptions

1.5.1 Scope of the SRDP project

Initially, SRDP will be restricted to making products for the VLA and ALMA. We assume that the pipeline has already been run on ALMA data to produce science ready calibration tables and standard image products as part of the ALMA Observatory deliverables. VLASS data processing operations will be assumed by SRDP at some point, but not before summer 2019. At that time, we will use the VLASS Operations Plan to add the needs of VLASS to the resource estimates in this document. In addition, although VLBA processing may be supported in the future, it is premature to estimate the operational input, particularly as no CASA pipeline currently exists for these data. Both VLBA and GBT raw data will be stored in, and made available from, the NRAO archive.

1.5.2 Use Cases/workflows for Support

The SRDP System Concept Document [RD04] identifies several use cases, that we have mapped into workflows (WFs) for support:

- WF1: Standard Calibration (VLA only): running the calibration pipeline until satisfactory (see Section 8) calibration tables are produced, then archiving.
- WF2: Standard Imaging (VLA only): restoring calibration tables (or using cached calibrated measurement sets on disk), flagging target data as needed, then running the imaging pipeline until satisfactory images are produced, then archiving.
- WF3: Optimized imaging (VLA and ALMA): restoring calibration tables, running the imaging pipeline with user-specified parameters, archiving products that pass QA.
- WF4: Archive (VLA and ALMA): data discovery and download, visualization.
- WF5: Restoration of Calibration (VLA and ALMA): including small modifications such as smoothing in time or frequency, generation of MS etc)
- WF6: Recalibration (VLA and ALMA): running the calibration pipeline with user-specified parameters (including line emission information so as to avoid flagging spectral lines as RFI, until these can be specified in the ngPST), then archiving products that pass QA.
- WF7: Combined imaging (VLA and ALMA): restoring calibration tables, running the imaging combination pipeline with user-specified parameters, archiving products that pass QA.
- WF8: Time Critical observations (VLA only): for pre-identified SBs, running the calibration and imaging pipelines (for a restricted set of images) on individual execution blocks as soon as they are observed, without waiting for the full set of SB executions to be completed, and making the results available to the PI.
- WF9: Large & Commensal projects (VLA and ALMA): ingest of user-supplied products into the archive following QA, eventually special processing

- WF10: Curation and reproducibility (relevant for all the above workflows, but not a workflow in itself). (These topics will be the responsibility of the SRDP project, and will be achieved via curation of the appropriate metadata in the archive.)
- WF11: Commissioning and validation of SRDP workflows and pipelines (VLA and ALMA).

2 Roles and responsibilities

2.1 Overview

The SRDP Operations Manager bears the primary responsibility for successful SRDP operations. The Operations Manager, in conjunction with the SRDP Project Director, negotiates for the necessary compute, storage, and data analyst resources at each site. The managers of the data analysts at each site (the Data Analyst Manager in Charlottesville and the head of the Science User Support Group in Socorro) are then responsible for allocating effort for SRDP.

2.2 Feedback to telescope operations

An important part of data processing is the ability to feed back information on telescope issues to the telescope operations teams. For ALMA, this is formalized through the PRTSPR JIRA Department. Although SRDP products will only be reprocessings of the ALMA data, PRTSPR tickets should be filed by data analysts (with the SRDP Operations Manager added as a watcher) if problems are found with the data or the ALMA archive content. For the VLA, the feedback process is via a weekly meeting on Thursdays that discusses the results of the weekly telescope stress tests and other operational issues identified by the data analysts.

2.3 Feedback to pipeline development

Individual data analysts may provide any feedback on pipeline operations and heuristics to the Operations Manager through the operations meetings, Google spreadsheets, or an SRDP Operations JIRA ticket. The Operations manager, in conjunction with the Project Scientist, will then synthesize input into appropriately themed and prioritized SRDP Project tickets, and assign to the SRDP Project Scientist. The Project Scientist will take these requests to the Pipeline Stakeholders meeting.

User choices for the imaging and recalibration workflows, pipeline run durations, and pipeline memory use will be stored in the workspace database. This will allow for analysis with the aim of optimizing the standard pipelines to reflect the most common needs of users.

3 Communications

3.1 Meetings

SRDP operations will be coordinated by means of regular tag-ups with the operations teams in Socorro and Charlottesville. In Charlottesville these will be the regular 9:45am twice-weekly tag-ups (where other stakeholders will also be present such as NA ALMA Ops). In Socorro, we will hold one short (30min) meeting per week with those data analysts involved in SRDP operations, science staff with responsibilities for SRDP operations (including the Operations Manager and Head of Science User Support), and technical experts (as needed).

3.2 Communication of resource needs

The resource model described in the Appendix will be used to predict resource needs (computing, storage and Data Analyst effort) on a quarterly basis. Every quarter the parameters of the model for each active SRDP workflow will be reviewed and adjusted in order to make predictions for the following quarter. For small changes, we expect that iteration with the line managers will be sufficient for us to obtain the resources we need. Large changes will need to be escalated to the appropriate Observatory management (DMS, SSR or ALMA), and any temporary rescoping of SRDP efforts and timeline will be negotiated with them.

3.3 JIRA and Helpdesk use

SRDP operations will make extensive use of NRAO's JIRA and Helpdesk systems. We envisage the following (in addition to the existing CASA and SSA software projects/departments for bug reporting, and the SRDP JIRA department):

- A JIRA SRDP Operations project. In this project, data analysts may file tickets for difficult or marginal QAs that require the attention of a Scientist, any operational issues they come across, and suggestions for pipeline heuristics improvements. (As an alternative to the JIRA workflow, spreadsheets may also be maintained to collate this information for the Operations Manager.) Tickets will be assigned by default to the SRDP Operations Manager, who will reassign them as needed.
- An SRDP Helpdesk Department in the NRAO helpdesk. This will be the mechanism by which users report problems with SRDP products, and ask general questions about SRDP. These tickets will be assigned to the SRDP Operations Manager by default.

3.4 Reporting

A subset of operations metrics (see Section 7.4), including the numbers of datasets processed through each workflow, will be part of the quarterly SRDP report sent to NRAO management.

4 Risk Management

A risk register for SRDP is maintained on sharepoint, under the Project Management Plan. Risks that are tracked that are relevant for SRDP Operations include: SRDP-9 (uncertainty in definition of product QA), SRDP-14 (operations staff oversubscription if user-requested processing proves very popular), SRDP-15 (computing resources oversubscription), SRDP-25 (lack of workspace for pilot). We will also track the risk associated with bulk reprocessing of SRDP data products if a problem is found in the pipeline.

5 Operations Roadmap

SRDP will be commissioned in Waves through the 5-year duration of the project, beginning with a pilot project in 2019. An operations roadmap for SRDP is maintained on Confluence at https://open-confluence.nrao.edu/display/SRDP/Operations+Roadmap. Here we describe the pilot and each of the SRDP waves based on the operations roadmap at the time of writing. Note that VLASS will be added to the SRDP effort at some time to be decided, and that we will follow the timeline in the VLASS Operations Plan for VLASS processing.

5.1 Pilot (end FY19)

A pilot study will run June-September 2019. This will consist of the VLA pipeline calibration workflow (WF1),

the ALMA optimized imaging workflow (WF3) and restoration of calibration for both VLA and ALMA (WF5). In addition, ingestion of products from the realfast commensal project into the archive will be in operation (WF9, but run by the realfast project in practice). A separate Pilot Program Plan (in prep) discusses the Pilot in detail, an operations summary is given below:

5.1.1 VLA Calibration in the pilot

The VLA calibration workflow (WF1) is expected to be very similar to that presently carried out for VLASS. Pipelines will be rerun in cases where manual flagging or parameter changes are needed and data products ingested into the archive. We will initially restrict the bands where calibration is attempted to high-frequency bands (Ku-band and above), as these are likely to be the most straightforward in terms of RFI rejection and calibration. Only SBs written using standard calibration best practices (as determined by the DA team) will be calibrated. Calibration needs to keep up with the rate of data taking of the VLA, in practice this means SRDP calibration products need to be delivered within ~30d (which is also a high-level SRDP requirement). Some flexibility in this timescale may be need to be made shortly after telescope moves are complete as a backlog can build up while the final antenna positions are being updated. As usual, the raw visibilities will be made available as soon as they are archived for those who do not wish to wait for SRDP products.

As the SRDP workflow manager has been descoped from the pilot, spreadsheets will be used to track the processing. After the first six weeks of pilot operations the process will be reviewed and we may decide to change the restrictions on the data that qualifies for pilot processing. It is estimated that there will be 316 hr of VLA data in Ku-band and above to calibrate, corresponding to approximately 158 execution blocks (one month of the pilot will be spent observing VLASS in the BnA configuration, which is not included in the SRDP estimates).

5.1.2 ALMA optimized imaging in the pilot

During the pilot period, user access will be restricted to a small group (for example, individuals suggested by the NRAO Users' Committee, ALMA Ambassadors, REU summer students etc). As the workflow manager will not be available for the pilot, use will be made of externally-hosted (i.e. Google) spreadsheets to track workflows and user interactions, and to communicate within the SRDP project. Pilot operations will be restricted to no more than two concurrent job requests from a given user, and overall to ten job requests per week for the first six weeks of the pilot, after which time these limits will be reviewed. Any jobs that are found to consume excessive amounts of computing resources (in the view of the Operations Manager) will be terminated. Information on such jobs will be collected for review following the close out of the pilot, and before the start of Wave 1 so that better ways of identifying problematic jobs can be found. The June-Sep 2019 time frame of the pilot will result in resource contention with ALMA pipeline and end-to-end testing. We have considered this in planning the scope of this pilot.

5.1.3 Restoration of calibrations

By the start of the pilot, restoration of calibrated measurement sets is expected to be mature for both ALMA and the VLA. The pilot will allow for more extensive testing and user feedback on this new archive capability.

5.1.4 Realfast support

Realfast will be operating by the start of the pilot, and ingest of realfast products will continue throughout

under WF9 (large projects).

5.2 Wave 1 (FY20)

Wave-1 capabilities for external users will be the same as the pilot, with a slow expansion of the parameter space covered (for example, inclusion of low frequency VLA calibration). Workflow management software will be available to replace the spreadsheets used for the pilot. Other large/commensal programs besides realfast will be ingested.

5.3 Wave 2 (FY21)

Wave 2 will expand user services to include a pilot of the VLA optimized imaging workflow. An image cutout service will also be made available.

5.4 Wave 3 (FY22)

In Wave 3, several new workflows will be introduced. VLA standard imaging (WF2) will begin, along with Time Critical observations (WF8), ALMA multi-configuration imaging (WF7), and ALMA and VLA recalibration (WF6), and large program processing (WF9).

5.5 Wave 4 (FY23)

Wave 4 will add VLA multi-configuration imaging as a pilot.

5.6 Wave 5 (FY24)

In Wave 5, all capabilities will be in production.

6 Resource requirements

The resource requirements for SRDP Operations, in terms of compute nodes, storage and personnel are estimated using the resource model described in the Appendix. The intent of the model is not necessarily to provide a high fidelity estimate of SRDP resource requirements, but instead is intended to be used as a tool to provide approximate guidance, and also allow the easy investigation of the effects of any changes in capabilities and resources on the total SRDP effort. Some of the assumptions for that model are described below, followed by the predictions.

6.1 Archive storage

6.1.1 Socorro

The VLA data rate is limited to ~25 MB/s (90 GB/hr) for regular observations (though may be exceeded in certain circumstances). If such a data rate were maintained for 8000 hr/year this would result in a data rate of 720 TB/yr. In practice, the actual rate (once low data rate projects and maintenance time are allowed for) is about 400 TB/yr (S. Witz, personal communication). As the VLA is now a fairly mature system, we do not expect a significant increase in this data rate until the ngVLA starts to ramp up, earliest 2025.

Archival storage for SRDP products in New Mexico will be small (a few percent of the raw data rate) until the imaging workflows begin in Wave-2. At this time, the sizes of the generated image products become a significant fraction of that of the raw data from which they are derived. Although the SRDP project will not produce a full set of continuum images and cubes for every observation, we can probably assume (based on

current ALMA values, where a full set of images is only attempted in the smaller configurations due to processing limitations) that the image products will be about 50% of the raw data. The image volume is very dependent on configuration, as the image sizes scale with the square of the configuration linear size (although this variation will be somewhat mitigated by the fact that most spectral line projects that will generate cubes need the greater brightness sensitivity of the closer configurations).

6.1.2 Charlottesville

The ALMA raw data volume in Cycle 6 (Oct 2018-Sep 2019) is predicted to be 128 TB (218 TB including products). As ALMA operations are now fairly mature, we expect this to only rise by a few percent per year over the next few years. A planned correlator upgrade in 2022, however, means the raw data rate could rise to ~260 TB/yr on a 4 to 5 year timescale. Image products currently constitute a data volume equivalent to 70% of the raw data, this fraction is expected to rise slowly as the pipeline performance is improved. SRDP products for ALMA are expected to be only a small fraction (~10%) of the total ALMA data volume, as only Optimized Imaging and multi-configuration imaging products (with perhaps high-level products from 1-2 large projects per Cycle) will be ingested into the SRDP archive.

6.2 Data Analysts

Data analysts (DAs) will be responsible for communication with the users, overseeing the workflows and reviewing products and weblogs for quality assessment (QA). Both Charlottesville and Socorro have experienced DAs who will need little additional training to take on these tasks. The additional effort for SRDP will, however, likely require additional data analyst hires. Three DAs have been added in NM to help support VLASS, and a further DA is being added in NM to augment the staff during the pilot.

6.3 Compute nodes and scratch storage

Both the Charlottesville and Socorro clusters, and their lustre scratch storage are heavily utilized. Currently, the Charlottesville cluster consists of 64 nodes with 256 GB of memory each, connected to a 1.1 PB Lustre filesystem. New Mexico has 60 nodes (plus 30 currently reserved for VLASS) with up to 512 GB of memory, connected to a 1.7 PB Lustre system. There is also the possibility of obtaining further resources from XSEDE, AWS or through the RADIAL program, by partnering with other institutions.

In the early waves of SRDP, the additional processing load from SRDP workflows is expected to be low. There will be additional imaging pipeline runs in Charlottesville to accommodate the Optimized imaging workflow. In New Mexico, there will be additional calibration pipeline runs (the ones following flagging). Later waves, which use the imaging pipeline as part of the standard SRDP processing in Socorro will have much more impact there, and, based on the performance of currently available software, the SRDP project will either need external processing resources to manage the load, or will need to restrict the scope of the imaging performed.

6.4 Resource model predictions

The results from the model can be compared against the available resources (Table 1). Note that we are ignoring other uses of the processing clusters. Planning is on a quarterly basis to align with software development and POP milestones.

Table 1: available resources

Resource	NM	CV
Number of nodes	60	64
Cores per node	16	16
Memory per node (GB)	512	256
Lustre space (TB)	1700	1100
Node hr/quarter	129600	138240
Core hr/quarter	2073600	2211840

6.4.1 Pilot and Wave 1

The pilot (2019 Jun – 2019 Sep) and Wave-1 (2019 Nov – 2020 Sep) operational requirements are very similar. The pilot operations will be more restricted, but also using incomplete software, so the DA effort per SB/MOUS will be larger, and more pipeline reruns may be required.).

6.4.1.1 Charlottesville operations

Table 2 (in the Appendix) shows the estimated resources per quarter for the pilot and Table 3 that for Wave 1. These assume 10 user-requested pipeline jobs per week. Actual demand may be higher (particularly early on for restores – the current operations model assumes ~80 jobs per week, but that could be higher if we have users who need a large fraction of the archival data calibrated) and we anticipate more user interaction while bugs are ironed out, so the request to the NAASC for SRDP operations will be for 50 TB of Lustre space and 4 weeks (160 hr) of Data Analyst time per quarter, which would be reviewed after the pilot is complete. Assuming the pilot runs for one quarter and Wave 1 for three quarters, about 22 TB of NGAS storage will be needed.

6.4.1.2 Socorro Operations

Socorro operations for the pilot and Wave 1 are given in Tables 4 and 5 in the Appendix. The effort and resource needs will be dominated by the calibration workflow (WF1). Approximately one FTE data analyst will be required on the pilot to operate the calibration pipeline and perform QA (with further 1.5 FTE needed for Wave 1). The estimated 7000 node hr for Wave 1 calibrations will be well within the available number of node hr for the Socorro cluster (140000). We estimate 250TB of Lustre space will be needed for Wave 1 operations, and a total of 2.5 FTE of data analyst effort.

6.4.2 Wave 2 Operations

Wave 2 adds pilot VLA Optimized Images. An additional 20000 node hours and 60hr of DA time will be needed for VLA Optimized Images pilot.

6.4.3 Wave 3+ ("Steady State") Operations

In Wave 3, VLA standard imaging will be rolled out. This will have a major impact on SRDP operations in Socorro. Several minor new capabilities (e.g. multi-configuration imaging, recalibration workflows) also will be rolled out in Wave 3+. These are likely to have a small impact on overall resource requirements (except for DA time), but are included for completeness. Tables 6 and 7 indicate likely long-term resource requirements for SRDP in Charlottesville and Socorro, respectively. *Note that the current cluster resources will likely be oversubscribed in Socorro, even without the additional load from VLASS.*

7 Operations requirements

7.1 Workflow process monitoring

Workflows are initiated either by a workflow manager, in the case of pipeline processing (Workflows 1, 2, 8), or by a user/DA in Workflows 3-7 and 9. Once a job begins, it will be tracked through the stages of being queued, running and completion (with or without an error flag). Jobs which crash during execution also need to be caught even if the software does not exit cleanly.

Once software is available to track these states through the workflow manager this will be the DA's primary source of information. During pilot operations states may need to be tracked by hand, which will add significantly to the human resources required.

7.2 Pipeline characterization and profiling

Pipeline development and testing is typically focused on heuristics, rather than operations. There is a need for testing of the pipeline in operations to profile memory, i/o and CPU usage so as to optimize the use of resources for pipeline runs. For example, more memory results in small improvements to some pipeline results, so memory allocations tend to be conservative. Similarly, running the pipeline highly parallel reduces the time required for any given dataset, but increases the i/o load on the lustre system as a whole. This profiling will be carried out by the group in Socorro led by James Robnett.

7.3 Remote processing

It is anticipated that NRAO computing resources will not be sufficient to support all processing. Therefore, remote processing workflows need to be supported. We can assume that the workflows would be triggered by NRAO staff, who will be able to make the appropriate changes in the state system and trigger the running of the process. Process monitoring becomes more complicated in this instance as performing sufficient messaging from the remote site may not be straightforward, and this will need to be borne in mind in the design phase. Once the process is complete, the weblog could either be transferred back to NRAO along with the contents of the products directory, or remain on the remote site and be viewed either remotely or by staff located at the processing center. This latter model might be especially appropriate for the RADIAL program, where training of participants at the remote processing centers in QA processes may be incorporated. Once the products are deemed satisfactory, they will be transferred back to NRAO and archived. There are many parallels here to the ALMA ARC processing workflow, and we will investigate what concepts and components can be reused from that effort.

7.4 Metrics

Accurate and complete metrics will be needed for SRDP for both system management and reporting. Metric definitions are discussed on confluence (https://open-confluence.nrao.edu/display/SRDP/SRDP+Metrics+Definition). There are several categories, listed below:

7.4.1 Archive Content Metrics

Describe the number and volume of SRDP products stored in the archive (including raw data, number of visibilities per EB, number of EBs per SB/MOUS, and volume of any reprocessings).

7.4.2 Archive Usage Metrics

Number of downloads by product, method and user type (PI/co-I, archival user).

7.4.3 Capability Usage Metrics

Number of runs of each type of SRDP workflow, customization choices made by users, location of pipeline run (e.g. CV, NM, AWS).

7.4.4 Workflow Metrics

Metrics describing pipeline runs (time to fetch data from the archive, total pipeline run duration, number of repeats (with reasons), resource usage etc). This will also include the total elapsed time between submission and delivery for each job request.

7.4.5 Product Quality Metrics

Quantify the number of times users (or the Observatory) request reprocessing of datasets, and whether those requests benefit from reprocessing.

7.4.6 Deployment Metrics

Describe the usage of the deployed system in terms of simultaneous users, number of processings requested, and system reliability.

8 QA plan

A uniform set of QA standards will be developed, which will include criteria on dynamic range, RMS compared to that theoretically achievable, fraction of data flagged (criteria will depend on telescope/band), uniformity of mosaic noise and relative flux accuracy over a mosaic, beam size and shape achieved in the image (compared to the nominal for the relevant telescope configuration), level and severity of any image artifacts, continuum subtraction, and calibration quality (gain, flux, phase transfer and polarization). These will be associated with scores. A data product with a low QA score on one or more criteria will not usually be archived (but may be given to a user with appropriate warnings).

The Data Analysts will provide the bulk of the QA assessments. In difficult or marginal cases, the Operations Manager will be alerted via an SRDP Operations JIRA ticket, and will assign QA to themselves, or to an a local SRDP scientist filling an "Astronomer on Duty" or similar role.

8.1 Standard use cases

The QA procedure will depend on the workflow:

- **WF1:** Standard Calibration: a DA will assess calibration weblog. If needed, they will add flagging or make other parameter changes as necessary and rerun (they can also request input from a scientist if they feel it is needed to resolve problems with the data). This procedure may be repeated until the calibration is considered to be science quality (comparable to the result that could be obtained by an experienced scientist processing the data by hand, and thereby satisfying the SRDP QA requirements), at which point archive ingest is performed and the dataset calibration is marked as completed in the workflow manager.
- **WF2:** Standard Imaging: a DA will assess the imaging weblog and image product. They will change pipeline parameters in the PPR, add target flags and rerun if needed to obtain a scientifically-valid result. Input from a scientist may be requested by the DA. This procedure may be repeated until the image reaches science quality, at which point archive ingest is performed and the dataset imaging is marked as completed in the workflow manager. If it is discovered that the calibration needs reprocessing as part of the imaging, new calibration products will be added to the archive (during the pilot period, if a problem is found with an ALMA observation, an ALMA PRTSPR ticket, and possibly a QA3 report will be generated to fix the calibration in the ALMA archive).
- **WF3:** Optimized imaging: a DA will assess the imaging weblog and image product. They will suggest changes to the parameters to the users if it is clear that doing so will improve the results. If the final run is judged by the DA to be of sufficient quality, it will be ingested into the archive (otherwise, or during the pilot phase where none of these products will be ingested, it will remain in the user's workspace for them to pick up). In the pilot stage of the project, before software is developed to check user inputs, the DAs will also perform a check on the pipeline parameters suggested by the user for each one to ensure that they are sensible and will not result in an overuse of resources. If this is not thought to be the case, the DA will contact the user requesting changes before the pipeline can be run.
- **WF4:** Archive download: no QA intervention should be required.
- **WF5:** Restoration of calibration: no QA intervention should be required, unless it repeatedly fails for a given dataset, in which case the Operations Manager should be alerted.
- **WF6:** Recalibration: a DA will assess the calibration weblog. They will suggest changes to the parameters to the users if it is clear that doing so will improve the results. If the final run is judged by the DA to be of sufficient quality, it will be ingested into the archive (otherwise, or during the pilot phase where none of these products will be ingested, it will remain in the users' workspace for them to pick up). In the pilot stage of the project, before software is developed to check user inputs, the DAs will also perform a check on the pipeline parameters suggested by the user for each one to ensure that they are sensible. If this is not thought to be the case, the DA will contact the user requesting changes before the pipeline is run.
- **WF7:** Combined imaging: QA for this will follow WF3, but more intervention from a DA may be required to pick suitable parameters and ensure that any necessary reprocessing is done to obtain consistent weights in the input uv data. Again, if the final run is judged by the DA to be of sufficient quality, it will be ingested into the archive.
- **WF8:** Time critical observations: no QA will be performed on reductions of time critical observations, however, some work will be needed by DAs to start pipeline runs and communicate the results to the user. Processing will be repeated in the regular WF1 and WF2 if needed before QA and ingestion into the archive.
- **WF9:** Large and commensal projects: the project team will work with the data analysts to ensure that QA is performed on any products that are ingested into the NRAO archive.

8.2 Erroneous SRDP products in the archive ("QA3")

If a user (either as a PI/member of the proposing team or an archival user) finds a problem with a dataset generated under the auspices of SRDP, this will be reported via a helpdesk ticket. The ticket will be assigned to the SRDP operations manager who will in turn assign a data analyst to look into the issue. Based on the report of the data analyst the following actions will be taken:

1) The QA3 report is determined to be based on a misconception or error on the part of the user. The user will be notified through a response on the helpdesk ticket and no further action will be taken.
2) The QA3 report is determined to be based on a problem with the observations. In the case of an ALMA observation the ticket will be forwarded to the NAASC as an ALMA QA3 request. In the case of another telescope, the user will be notified through a response on the helpdesk ticket. If it is determined that the product can be improved (e.g. by flagging of bad data), then it will be rerun through the relevant pipeline, a new product will be added to the pipeline, and the QA status of the prior product in the archive changed to fail.
3) The QA3 report is determined to be based on a problem with the data processing. If possible, the data will be rerun with corrected parameters and/or a new version of the pipeline to resolve the problem and re-ingested into the archive. If this is not possible (because the problem is due to a newly found or unfixed bug in the pipeline), then a CASA ticket will be opened and the user informed. The data product (and ones affected by the same issue) may be blocked from download from the archive if the issue is deemed to affect the scientific validity of the data by the SRDP operations manager, in consultation with other SRDP and CASA scientists.

QA3 actions also sometimes arise as a result of Observatory staff finding a bug in the pipeline or related software. In these cases, the SRDP project will initiate reprocessing of data after conducting an assessment in conjunction with the pipeline team. The decision to reprocess affected datasets or not, and how to communicate problems to the users will be taken by the SRDP project management team.

9 Curation and Reproducibility

9.1 Data Curation

The NRAO archive, provides access to data from all NRAO telescopes (ALMA, VLA, and VLBA) as well as from the Green Bank Observatory. The architecture of the Archive is outside the scope of this document. However, for clarity we note that meta-data for data from all telescopes is co-located in a database located in Socorro. The data itself resides on NGAS nodes located in both Socorro and Charlottesville (backup versions of all NGAS nodes are maintained). Products produced by the ALMA project also reside on NGAS nodes in Charlottesville, these are maintained as a distinct set to allow replication with other ALMA executives.

9.2 Data provenance

With each SRDP product, sufficient metadata will be stored that would allow a hypothetical user to reproduce the product should they desire. This information will be stored in the archive database at the same time the product is ingested. These metadata shall contain:

1. VO Obscore 1.1 values or their equivalents such that an Obscore 1.1 view of the archive can be built from them. These include basic parameters of the observation such as footprint on the sky, frequency range and resolution, and date/time of observation, but can be extended to include details of the products (for example, axes of a cube, including frequency and polarization).
2. The version of CASA that was run to produce the product
3. The operating system and python version under which the pipeline was run.

4. The hardware configuration under which the code was run (memory and number of cores), and the site at which it was run (CVille, Socorro or e.g. AWS, XCEDE).

9.3 DOIs

Digital Object Identifiers (DOIs) will be associated with each set of archived SRDP image products to ensure the existence of a persistent link to the data even if, for example, the archive API changes. This will allow users to cite SRDP data in publications.

10 Authorization and Authentication

In order to run any user-driven SRDP workflow other than an archive download or restore (WF4 and WF5), a user must be logged into the system (with their my.nrao account or equivalent). This ensures traceability of the job and allows linkage to any helpdesk tickets.

SRDP operations will follow the policies of the different observatories, allowing any user to enact relevant workflows while data are out of their proprietary period. For non-public data, only the PIs or PI-delegated users may perform SRDP workflows on ALMA data, and only PIs or CoIs can run SRDP workflows for VLA data. Checks for proprietary periods need to bear in mind that in the case of ALMA, public data can have its release rescinded if a problem is found.

11 User Policies

User policies for acceptable use of the user-driven workflows for restoration, imaging and recalibration will be developed. These will apply to both storage and compute resources. Initially, we expect a period where these policies are still under development, and the Operations Manager will have discretion to approve or disapprove requests for large amounts of processing. Large requests may also be reduced in priority on the queues so as not to prevent the timely execution of smaller jobs. User created data will remain on disk and available for download for a fixed time period (if it is not directly imported into the archive). This time period remains to be determined, but is likely to be around 30 days (parameter Duration in the operations model).

12 Test plan, including end-to-end and regression testing

12.1 Overview

A separate test environment will be used to test new SRDP functionality. SRDP will have rolling wave development, so testing will be ongoing in parallel with operations. Testing of SRDP workflows and software is estimated to require 10% of the total resources. Note that this does not include testing of the pipeline heuristics or CASA, which are managed separately, although coordination of the two types of tests may increase efficiency.

12.2 New features

As major new functionality ("Epics") are introduced, there will be associated science validation testing following the handoff of new software by DMS. These tests will be conducted under the auspices of SRDP, and testing will be led by the Project Scientist or their designee.

12.3 Regression testing

A regression test suite will be identified at each site and run prior to the implementation of each new Wave in production after handoff from DMS. This will consist of a set of workflows from the prior wave of

implementation. These tests will also be led by the Project Scientist.

12.4 Operations testing

Before the start of every Wave, the SRDP operations staff will test an example of each active workflow to verify that the system works as expected, that any data products are archived along with the appropriate metadata, and that appropriate notifications are sent to the Data Analysts, Operations Manager and User. In addition, if major new workflow functionality is introduced during a wave, operations staff will test it in a test environment (along with any regression tests the Operations Manager deems necessary) before deployment.

Appendix: Operations model

The SRDP operational model has ten operational use cases, the demand for some of which is currently unknown, thus quantifying the necessary resources is difficult. Nevertheless, we have come up with a simple model to allow us to scale the likely resource requests, the key parameters of which are:

1. MaxImgSize: Maximum size of Image products (10GPix, 40GB). It is suggested that this be communicated to user as the maximum they can ask for, but how it is distributed (e.g. 40 1GB images, 4 10GB images or even a single 40GB cube) can be left up to the user.
2. PlRunsPerCalProd: Number of pipeline runs per calibration product (1.5). Based on Drew's assessment for VLASS.
3. PlRunsPerImgProd. Number of pipeline runs per image product (1.2)
4. WeblogTimeCal: time for weblog review - calibration, including flagging if needed) (1hr)
5. WeblogTimeImg: for weblog review - imaging (0.25hr)
6. UserInteraction3: Time for user interaction per job, optimized imaging (0.25hr)
7. JobsPerWeek3: Number of jobs/week, use case 3 (10)
8. JobDurationImg10Gpix: Duration of a 10Gpix imaging job: (7d). This is currently uncertain and depends on the effectiveness of parallelization etc. More investigation required.
9. JobsPerNodeCal: Number of jobs to stack per node, calibration (8)
10. JobsPerNodeImg: Number of jobs to stack per node, imaging (1)
11. TestingFrac: Fraction of time/resources for testing: 10% (floor or 20hr/Q for DA time)
12. TimeMultCal: ratio of time to process data to time to take data – calibration (20)
13. TimeMultRestore: ratio of time to process data to time to take data – restore (2)
14. SpaceMultCal: ratio of raw to processed data volume – calibration (3)
15. SpaceMultImg: ratio of primary+auxillary image products to primary products (5)
16. Duration: length of time products stay on Lustre (30d)
17. VLARawDataRate: VLA raw data rate 400TB/yr
18. ALMARawDataRate: ALMA raw data rate 120TB/yr
19. VLASBDuration: average VLA SB duration: hr
20. ALMASBDuration: average ALMA SB duration: 2hr
21. ImgSuppLargeProj: how much extra support will be needed for imaging large projects (as a fraction of all image processing): 0.2 (not currently used; NumImgLrg used instead).
22. UserInteraction9: time per large project to interact with PI team: 40hr
23. NumLargeProj: number of large project ingests per quarter per facility: 1
24. JobsPerWeek6: Number of recalibration jobs per week: 10
25. JobsPerWeek7: Number of multiconfiguration jobs per week: 5
26. NumTimeCrit: Number of time critical SBs per quarter: 1

27. NumImgLrg: Average number of "special" images to be made per large project: 20
28. CalProdSizeFrac: size of archived calibration products as a fraction of the raw data rate (based on Myers' estimate for VLASS): 0.013
29. DaysPerQ: number of days per Quarter (or other reporting period): 90
30. NumCoresImg: number of cores actually used in an imaging pipeline run: 9
31. FracALMAMan: fraction of ALMA datasets that are manually calibrated (for which SRDP will store the calibrated MS)
32. TimePerMSCopy: DA time to initiate the copying of a manually calibrated ALMA MS to the archive staging area and trigger ingest: 0.1hr

There are also two constants: NumHourPerQ – 1500 hr/quarter observed; NumCoresPerNode=16
Resource requirements per quarter are estimated as follows:

WF1 (VLA only):

- Number of core hr = NumHourPerQ* TimeMultCal* PlRunsPerCalProd
- Number of node hr = NumHourPerQ* TimeMultCal* PlRunsPerCalProd/JobsPerNodeCal
- Lustre space (TB) = VLARawDataRate/(365/DaysPerQ)*SpaceMultCal*Duration/DaysPerQ
- DA hrs = NumHourPerQ*VLASBDuration*PlRunsPerCalProd*WeblogTimeCal
- Archive storage = VLARawDataRate/(365/DaysPerQ)*CalProdSizeFrac

WF2 (VLA only):

- Number of core hr = NumHourPerQ/VLASBDuration*PlRunsPerImgProd* MaxImgSize/40GB* JobDurationImg10Gpix*24hr*NumCoresImg
- Number of node hr = NumHourPerQ/VLASBDuration*PlRunsPerImgProd* MaxImgSize/40GB* JobDurationImg10Gpix*24hr/JobsPerNodeImg
- Lustre space (TB) = NumHourPerQ/VLASBDuration* MaxImgSize/40GB*0.04TB* SpaceMultImg*Duration/DaysPerQ
- DA hrs = NumHourPerQ/VLASBDuration* PlRunsPerImgProd*WeblogTimeImg
- Archive storage (TB) = NumHourPerQ/VLASBDuration*MaxImgSize/1000

WF3 (ALMA & VLA):

- Number of core hr = JobsPerWeek3*DaysPerQ/7* MaxImgSize/40GB *JobDurationImg10Gpix *24hr *NumCoresImg
- Number of node hr = JobsPerWeek3*DaysPerQ/7*MaxImgSize/40GB*JobDurationImg10Gpix *24hr/JobsPerNodeImg
- Lustre space (TB) = JobsPerWeek3*DaysPerQ/7* Duration/120d *SpaceMultImg *MaxImgSize/1000
- DA hrs = JobsPerWeek3*DaysPerQ/7*(WeblogTimeImg+UserInteraction3)
- Archive Storage (TB) = JobsPerWeek3*DaysPerQ/7*MaxImgSize/1000

WF4 (ALMA & VLA):

- Number of core hr = NumHourPerQ (assumes staging time=observation time)
- Number of node hr = NumHourPerQ/JobsPerNodeCal
- Lustre space (TB) = [ALMA/VLA]RawDataRate/(365/DaysPerQ)*Duration/DaysPerQ
- DA hr = 0 (unless there are issues)
- Archive Storage = 0

WF5 (ALMA & VLA):

- Number of core hr = NumHourPerQ*TimeMultRestore
- Number of node hr = NumHourPerQ*TimeMultRestore / JobsPerNodeCal
- Lustre space (TB) = [ALMA/VLA]RawDataRate/(365/DaysPerQ)*Duration/DaysPerQ *SpaceMultCal
- DA hr = 0 (unless there are issues)
- Archive Storage = 0

WF6 (ALMA & VLA):

- Number of core hr = JobsPerWeek6*DaysPerQ/7*[ALMA/VLASBDuration]*TimeMultCal *NumCoresPerNode /JobsPerNodeCal
- Number of node hr = JobsPerWeek6*DaysPerQ/7*[ALMA/VLASBDuration]*TimeMultCal /JobsPerNodeCal
- Lustre space (TB) = JobsPerWeek6*DaysPerQ/7*[ALMA/VLASBDuration]* ALMA/VLA]RawDataRate/(365/DaysPerQ)/NumHourPerQ* SpaceMultCal *Duration/DaysPerQ
- DA hr = JobsPerWeek6*DaysPerQ/7*(WeblogTimeCal+UserInteracton3)
- Archive Storage (TB) = JobsPerWeek6*DaysPerQ/7*[ALMA/VLA]RawDataRate/4*[ALMA/VLA]SBDuration/NumHourPerQ

WF7 (VLA & ALMA):

- Number of core hr = JobsPerWeek7*DaysPerQ/7* MaxImgSize/40GB *JobDurationImg10Gpix *24hr *NumCoresPerNode/JobsPerNodeImg
- Number of node hr = JobsPerWeek7*DaysPerQ/7*MaxImgSize/40GB*JobDurationImg10Gpix *24hr/JobsPerNodeImg
- Lustre space (TB) = JobsPerWeek7*DaysPerQ/7*Duration/DaysPerQ *SpaceMultImg *MaxImgSize/1000
- DA hrs = JobsPerWeek7*DaysPerQ/7*(WeblogTimeImg+UserInteraction3)
- Archive Storage (TB) = JobsPerWeek7*DaysPerQ/7*MaxImgSize/1000

WF8 (VLA only):

- Number of core hr = NumTimeCrit*(VLASBDuration*(TimeMultCal)+ JobDurationImg10Gpix*24)* NumCoresPerNode/JobsPerNodeImg

- Number of node hr = NumTimeCrit*(VLASBDuration*(TimeMultCal)+ JobDurationImg10Gpix*24)* NumCoresPerNode/JobsPerNodeImg
- Lustre space (TB) = NumTimeCrit*(VLASBDuration*VLARawDataRate/4/NumHourPerQ* SpaceMultCal+MaxImgSize/1000 *SpaceMultImg)*Duration/DaysPerQ
- DA hrs = NumTimeCrit*(WeblogTimeCal+WeblogTimeImg+UserInteraction3)
- Archive Storage = 0

WF9 (ALMA & VLA large projects [excluding VLASS]; assumes calibration is absorbed into WF1 or ALMA equivalent):

- Number of core hr = NumLargeProj*NumImgLrg* JobDurationImg10Gpix*24hr* NumCoresPerNode/JobsPerNodeImg
- Number of node hr = NumLargeProj*NumImgLrg* JobDurationImg10Gpix *24hr/JobsPerNodeImg
- Lustre space (TB) = NumLargeProj*NumImgLrg*MaxImgSize/1000*SpaceMultImg *Duration/DaysPerQ
- DA hrs = NumLargeProj* NumImgLrg * (WeblogTimeImg+UserInteraction3)
- Archive Storage (TB) = NumLargeProj*NumImgLrg*MaxImgSize/1000

WF10: N/A

WF11 (testing/validation):

10% of total of the above

Pilot	Core hr	Node hr	Lustre TB (average)	DA hrs	Archive storage (TB)
WF1 (CAL)	0	0	0		
WF2 (IMG)	0	0	0		
WF3 (User Opt Img)	194400	21600	10	64	5
WF4 (Archive tasks)	1500	188	10		
WF5 (Restore etc)	3000	188	15		
WF6 (User recalibration)	0	0	0		
WF7 (Combined Img)	0	0	0		
WF8 (VLA Time critical)	0	0	0		
WF9 (Large Projects)	0	0	1	10	1
WF11 (Testing)	19890	2198	4	20	
Total	218790	24173	39	94	6
Fraction of available	0.10	0.17	0.00		

Table 2: Charlottesville Pilot

Wave1	Core hr	Node hr	Lustre TB (average)	DA hrs	Archive storage (TB)
WF1 (CAL)	0	0	0		
WF2 (IMG)	0	0	0		
WF3 (User Opt Img)	194400	21600	9	64	5
WF4 (Archive tasks)	1500	188	10		
WF5 (Restore etc)	3000	188	15		
WF6 (User recalibration)	0	0	0		
WF7 (Combined Img)	0	0			
WF8 (VLA Time critical)	0	0	0		
WF9 (Large Projects)	0	0	1	10	1
WF11 (Testing)	19890	2198	3	20	
Total	218790	24173	38	94	6
Fraction of available	0.10	0.17	0.03		

Table 3: Charlottesville Wave 1

Pilot	Core hr	Node hr	Lustre TB (average)	DA hrs	Archive storage (TB)
WF1 (CAL)	7110	889	99	237	0
WF2 (IMG)	0	0	0	0	
WF3 (User Opt Img)	0	0	0	0	
WF4 (Archive tasks)	1500	188	33	0	
WF5 (Restore etc)	3000	375	100	0	
WF6 (User recalibration)	0	0	0	0	
WF7 (Combined Img)	0	0		0	
WF8 (Time critical)	0	0	0	0	
WF9 (Large Projects)	0	0	0	0	
WF11 (Testing)	1161	145	23	30	
Total	12771	1596	255	267	0
Fraction of available	0.01	<0.01	<0.01		

Table4: Socorro Pilot

Wave1	Core hr	Node hr	Lustre TB (average)	DA hrs	Archive storage (TB)
WF1 (CAL)	33750	4219	99	1125	0
WF2 (IMG)	0	0	0	0	
WF3 (User Opt Img)	0	0	0	0	
WF4 (Archive tasks)	1500	188	33	0	
WF5 (Restore etc)	3000	375	100	0	
WF6 (User recalibration)	0	0	0	0	
WF7 (Combined Img)	0	0		0	
WF8 (Time critical)	0	0	0	0	
WF9 (Large Projects)	0	0	0	40	
WF11 (Testing)	3825	478	23	30	
Total	42075	5259	255	1195	0
Fraction of available	0.02	0.04	0.15		

Table5: Socorro Wave 1

Steady State	Core hr	Node hr	Lustre TB (average)	DA hrs	Archive storage (TB)
WF1 (CAL)	0	0	0		
WF2 (IMG)	0	0	0		
WF3 (User Opt Img)	194400	20160	9	60	5
WF4 (Archive tasks)	1500	188	10		
WF5 (Restore etc)	3000	375	30		
WF6 (User recalibration)	3857	482	17	60	0
WF7 (Combined Img)	97200	10800	3	30	
WF8 (Time critical)	0	0	0		
WF9 (Large Projects)	53760	3360	1	10	1
WF11 (Testing)	35372	3536	7	20	
Total	389089	38901	77	180	6
Fraction of available	0.18	0.28	0.07		

Table6: Steady State Charlottesville

Steady State	Core hr	Node hr	Lustre TB (average)	DA hrs	Archive storage (TB)
WF1 (CAL)	33750	4219	99	1125	0
WF2 (IMG)	1360800	151200	50	225	30
WF3 (User Opt Img)	194400	20160	9	60	5
WF4 (Archive tasks)	1500	188	25	0	
WF5 (Restore etc)	3000	375	101	0	
WF6 (User recalibration)	3857	482	17	150	0
WF7 (Combined Img)	97200	10800	4	30	3
WF8 (Time critical)	3168	198	67	2	
WF9 (Large Projects)	53760	3360	1	10	1
WF11 (Testing)	175144	19098	37	30	
Total	1926579	210080	410	1632	39
Fraction of available	0.93	1.62	0.24		

Table7: Steady State Socorro