



LARGE SYNOPTIC SURVEY TELESCOPE

Large Synoptic Survey Telescope (LSST)  
Data Management

# Conda Environment Proposal for Science Pipelines

Gabriele Comoretto

DMTN-110

Latest Revision: 2019-06-19

**DRAFT**

## Abstract

This proposal is an improved way to deal with Conda environments when building the Science Pipelines software. The benefits of implementing it can be seen in both the development and release processes. The effort required for its implementation is limited since it reuses all tooling already in place.

## Change Record

Version	Date	Description	Owner name
1.0	2019-06-19	Document Issue 1	Gabriele Comoretto
	2019-03-01	First draft	Gabriele Comoretto

*Document source location:* <https://github.com/lst-dm/DMTN-110>

*Version from source repository:* 1214904

Draft

# Contents

<b>1 Introduction</b>	<b>1</b>
1.1 Applicable Documents . . . . .	1
1.2 Definitions . . . . .	1
<b>2 Current Implementation</b>	<b>3</b>
2.1 Development Workspace Preparation . . . . .	3
2.2 Science Pipelines Distribution . . . . .	3
2.3 Conda Environment Management . . . . .	4
<b>3 Problems</b>	<b>5</b>
3.1 Static Environment . . . . .	5
3.2 Unknown Environment Status . . . . .	5
3.3 High Time to Rebuild . . . . .	5
3.4 Missing Usecases . . . . .	6
<b>4 Proposed Implementation</b>	<b>7</b>
4.1 Define Environment References . . . . .	7
4.2 Use Different Environments in lsstsw Git Package . . . . .	7
4.2.1 Update deploy Command . . . . .	7
4.2.2 Update rebuild Command . . . . .	8
4.2.3 New envconfig Command . . . . .	8
4.3 Changes to lsst/lsst Git Package . . . . .	9
4.4 Changes to Jenkins builds . . . . .	9
<b>5 Following activities</b>	<b>10</b>
5.1 Conda Packaging 3rd Party Libraries . . . . .	10
5.1.1 Conda Packages Managed by DM . . . . .	10
5.1.2 PyPi Packages . . . . .	11
5.2 Change Control . . . . .	11
5.3 More Packages to Conda . . . . .	11



<b>A Packages to be Moved</b>	<b>13</b>
<b>A References</b>	<b>16</b>
<b>B Acronyms used in this document</b>	<b>16</b>

Draft

# Conda Environment Proposal for Science Pipelines

## 1 Introduction

The proposal made in this document is oriented to have a more functional approach in the environment management for Science Pipelines builds.

The benefits that can be obtained by implementing this proposal are mainly for developers, who will have more control over the environments they are using. This will permit a more dynamic evolution of the used environment, allowing more frequent upgrades to newer versions of packages in the environment.

The release process will benefit as the global build time for the Science Pipelines will be reduced. A reduced build time allows more frequent releases and faster development turnaround both locally and using continuous integration. This will be very important during commissioning and the initial operations phase, when multiple releases per day can be needed.

At the moment there are few EUPS packages *stubs* that are not always working as expected. Implementing this proposal will remove these packages, and the overall build process will be more robust.

Implementing this procedure will have very low impact on the DM development team. The immediate benefit will be a better management of the conda environment. Activities documented in section 5 can be scheduled afterward, depending on resource availability and avoiding impacts on project milestones.

### 1.1 Applicable Documents

### 1.2 Definitions

Following terms are used in this document which are often overloaded.

Please refer to DMTN-106, section 2, where their definition is under consolidation.

- **Environment** DMTN-106, subsection 2.1
- **Software Product** DMTN-106, subsection 2.2

- **Software Release** DMTN-106, subsection 2.4

Draft

## 2 Current Implementation

### 2.1 Development Workspace Preparation

As of March 2019, the following steps are used to configure a build workspace for the Science Pipelines:

- clone the GitHub package <https://github.com/lsst/lsstsw>
- deploy the environment using `./bin/deploy`
- activate the environment by sourcing `bin/setup.sh`

This will make available locally all tooling required to build the science pipelines source code using the command `rebuild <-r ref> lsst_distrib`.

The environment deployed and activate using the above steps corresponds to a fixed revision of the *scipipe\_conda\_env* Git package. This revision is hardcoded in the *lsstsw* build tool. It is also hardcoded in the continuous integration scripts versioned in the *jenkins-dm-jobs* GitHub repository.

### 2.2 Science Pipelines Distribution

Following the instructions provided in [pipelines.lsst.io](https://pipelines.lsst.io), it is possible to configure a build workspace using the *newinstall.sh* script. This script is available in the *lsst/lsst* Git repository and it is tagged weekly together with the Science Pipelines weekly build.

The instantiation of the work space is described in the [pipelines.lsst.io](https://pipelines.lsst.io) documentation page, and can be summarized as follows:

- download from GitHub the *newinstall.sh* script in an empty folder.
- execute the script: it will deploy the tools and the environment needed for the build. It will also create an additional shell script, *loadLSST.bash*, to be execute to activate the Conda environment. The script is generated different shell flavors also: `csh`, `ksh` and `zsh`.
- source the *loadLSST.bash* script in order to activate the environment

The workspace will be ready to build the Science Pipelines using `eups distrib install ...` builds. If `newinstall.sh` has been executed with the options `-ct`, binary packages will be downloaded from EUPS instead of building the source code, when building.

Note that the Conda environment reference is hardcoded in the `newinstall.sh` script.

The script may have changed from one week to another. This implies that each EUPS weekly build may work only with its corresponding `newinstall.sh`.

## 2.3 Conda Environment Management

The repository `scipipe_conda_env` defines a list of Conda packages to be installed in the Conda environment. These packages are resolved by Conda from the default channels.

The `scipipe_conda_env` repository includes two different types of environment definition:

- **bleed**: the list of the Conda packages that are required by the Science Pipelines. No version is specified.
- **pinned**: the list of all Conda packages resolved in the environment. This includes the bleed packages and all required dependencies resolved automatically by Conda. The precise resolved version is given for each package listed in the file.

The *pinned* version of the environment definition is the one used to instantiate the environment for the Science Pipelines build as specified in sections 2.1 and 2.2.

The *bleed* version is used to update the *pinned* definition to the latest packages version that have been made available in the public Conda channels since the previous update. The *pinned* file is updated every 6 months approximately.

For the *linux* platform, an extra package is required, `nomkl`<sup>1</sup>. For this reason two different bleed files and two different pinned files are provided.

---

<sup>1</sup>See JiraDM-8146



## 3 Problems

### 3.1 Static Environment

As described in section 2, the environment deployed and activated using the actual version of the tooling is hardcoded in a few places.

This implies that, each time the environment is updated, the hardcoded value need to be updated, the tools redeployed, and restart from scratch any build or development activities.

The tools used should be able to deploy and activate a different environment, without the need to have it hardcoded.

### 3.2 Unknown Environment Status

In some cases, the science pipelines build is failing due to some packages misalignment. It is very common to see support requests on *slack* channels, that end up been a problem of Conda environment.

Since developers are not constrained to use the recommended environment, there should be an easy way to compare the active environment with a reference one, locally deployed or a reference in *scipipe\_conda\_env* Git package.

### 3.3 High Time to Rebuild

The 17.0 release of the Science Pipelines Stack comprises 128 packages.

All packages are rebuilt all the times a new *lsstsw* development workspace is configured. Each time a release or a release candidate are created, they are rebuilt also, with the effect to slow down considerably the release process.

Many of these packages are 3rd party libraries that are not available in the public Conda channels, or that require some changes in order to be used in the Science Pipelines Stack.

These packages are not DM software, they are not formally part of the release, and they should be considered as part of the Conda environment. Moving them to the conda environment will decrease the build time and improve the release as is.

In order to move these packages to the Conda environment, there are two problems:

- create and maintain Conda packages suitable for DM usage
- dynamically manage the Science Pipelines Conda environment, in order to be able to include and use new versions of the required packages. This is the critical point.

See appendix A for a list of packages proposed to be moved to the conda environment.

The expected benefits are:

- more efficient environment management;
- reduced build time and therefore easier to release the Science Pipelines;
- have a more stable build system, since reducing the number of package to build every time will reduce the possibility of random failures.

### 3.4 Missing Usecases

As a result of the problems exposed above, following use cases should be taken in account:

- be able to deploy a new Conda environment without the need to update and have a new clone of *lsstsw*.
- be able to activate different Conda environments using a *lsstsw* command.
- be able to check the differences between Conda environments.
- record which Conda environment has been used for a specific build, *lsstsw* rebuild, and be able to activate it again.
- be able to activate the correct environment before running and EUPS build using `eups distrib install`.
- be able to run a CI build using a specific environment reference.
- be able to manage Conda packages for the third party libraries that requires DM rebuild or customization.

## 4 Proposed Implementation

Here follows the proposed changes required to satisfy the usecases from the previous section.

### 4.1 Define Environment References

The package *scipipe\_conda\_env* to be tagged in GitHub for each weekly and each time an EUPS build is created. This will guarantee that each EUPS weekly build will have a corresponding well known Conda environment, that can be used for rebuilt them. The Conda environment tag shall be in line with the EUPS tag. The SHA1 can be identified also, but should not be strictly needed.

All changes implemented in the *scipipe\_conda\_env* repository shall follow the usual development process. The tooling improvements proposed bellow, will permit the validation of those changes and include them in the development workflow.

### 4.2 Use Different Environments in lsstsw Git Package

A new command, or command line option, shall be added to lsstsw to deploy and activate a specific Conda environment reference from the *scipipe\_conda\_env* GitHub package.

The possible options should be:

- deploy/activate a specific *scipipe\_conda\_env* hash
- deploy/activate a specific *scipipe\_conda\_env* branch
- deploy/activate a specific *scipipe\_conda\_env* tag
- deploy/activate the tip of *scipipe\_conda\_env* master

This can be done with the following changes in lsstsw package.

#### 4.2.1 Update deploy Command

The **deploy** lsstsw command should accept the environment reference **envref** as an additional command line option.

The default behavior should be the same as now. If no `envref` is provided, the usual fixed hash shall be deployed.

In case the `envref` is *master*, it shall deploy the latest master, if not already deployed. In this way, this option can be used to check if there is a newer master version of the conda environment and deploy it.

The environment shall have a unique name, not the fixed one that is currently used.

#### 4.2.2 Update rebuild Command

The **rebuild** command should record the environment reference used for each build.

This will permit in the future to retrieve the information, in order to manually activate the same conda environment before running a build.

The **rebuild** command could also raise a warning, in case the active environment is different from the one used previously on the same package. However this is not a problem related to the environment management but a potential improvement in the build tooling.

#### 4.2.3 New `envconfig` Command

A new command shall be added, that permits to choose an available environment to activate. The environment need to be already available locally, deployed for example with *deploy* command describe above in `refsec:deploy`.

It shall accept one of following options:

- a *scipipe\_conda\_env* environment reference to activate
- an existing build number, in order to activate the environment that has been used for that build.
- list the available environments, ready to be activated
- compare the active environment with an other one, a *scipipe\_conda\_env* reference or a deployed environment

If no option is provided, the default reference shall be activated, in order to keep the current

behavior.

### 4.3 Changes to lsst/lst Git Package

The distribution script *newinstall.sh* should be updated in order to be able to deploy and activate the environment corresponding to the EUPS build that is going to be installed.

This imply that the generate scripts *loadLSST.\** shall accept an environment parameter as input, and deploy the environment reference before activating it, if not already deployed.

### 4.4 Changes to Jenkins builds

The Jenkins jobs, for example *run-rebuild*, shall be updated in order to accept an environment reference as an optional parameter. In this way it will be possible to verify that the science pipelines builds with a new environment.

## 5 Following activities

The changes proposed in the previous section 4 can be implemented with minimal impact to the day to day activities.

The new functionalities introduced can already facilitate the management of the conda environment and development activities in general.

However, the resolution of the problems exposed in 3, requires the reduction of the number of packages included in the shared stack, moving third party libraries in the conda environment.

### 5.1 Conda Packaging 3rd Party Libraries

Some of the packages listed in appendix A are already available in Conda public channels or *PyPi* (see following 5.1.2). They are now included in the Science Pipelines Stack as EUPS packages since when they were initially needed, they were not available in any Conda channels.

Some other packages require to be updated more often, therefore they have been included as EUPS packages, since the Conda environment is not sufficiently dynamic.

Packages that require customizations with substantial code changes, need to remain part of the stack, at least in a first phase. Packages that require minor patches can be handled by a new recipe.

A small team shall be defined in DM to maintain the environment and the needed Conda recipes.

#### 5.1.1 Conda Packages Managed by DM

DM shall use the packages from the Conda public channels, if available.

DM shall maintain a Conda recipes for all those packages that are not available in Conda / *PyPi*, or that requires to be rebuild with a specific flag. This is the case of *boost*. These packages shall have a specific naming convention, so it will be clear they have been *re-packaged* by LSST DM.

A DM maintained Conda package shall be dismissed, not maintained anymore, as soon as

the corresponding package available in the public Conda channel can be used in the Science Pipelines build.

### 5.1.2 PyPi Packages

The conda environment definition files in *scipipe\_conda\_env* can accept also pip package.

This implies that packages available in *PyPi* package repository and that can be installed in the environment. However, if available, the preference shall be given to the Conda package.

In case a pip package will bring into the environment too many PyPi dependencies, DM can decide to maintain a Conda recipe for it.

## 5.2 Change Control

Since maintaining conda recipes may require a considerable amount of effort, the DMCCB shall monitor when new packages are added to the environment or are used in the codebase.

Changes to the environment shall be proposed using RFC Jira issues.

Third party packages shall also be documented and included the DM product tree. Each package shall have an internal reference person.

## 5.3 More Packages to Conda

Considering the inclusion of the 3rd party libraries listed in appendix A as a first step, follows a list of activities that can improve the overall DM codebase management and facilitate the development activities.

- create conda packages for the main DM tools used for the build process, such as EUPS, sconsUtils, lsst\_build.
- include in the Conda environment those third party libraries that are customized by DM with specific code changes, such as *starlink\_ast*.
- define environments for different uses, for example development and software debug, operations, science data analysis, etc.
- create conda packages for all DM software.

Implementing the first bullet should not require much effort, but will simplify the tooling we are using for instantiate a new development workspace, such as *lsstsw/deploy* or *newinstall.sh*.

The second bulled is a long term activity, that will permit to apply semantic versioning to each Git package, and therefore have proper dependency management.

Draft



## A Packages to be Moved

The following (EUPS) packages are proposed to be moved from science pipelines to the conda environment:

- apr (\*)
- apr\_util (\*)
- astrometry\_net
- astropy (\*)
- boost (\*)
- cfitsio (\*)
- coord
- doxygen (\*)
- eigen (\*)
- esutil
- fftw (\*)
- flake8
- galsim
- gsl (\*)
- healpy
- libyaml
- lmfit
- log4cxx (\*)
- matplotlib (\*)
- minuit2 (\*)
- mpi

- mpi4py
- mpich
- numpy (\*)
- ndarray (\*)
- pep8\_naming (\*)
- pybind11 (\*)
- pycodestyle
- pyflakes
- pykg\_config
- python\_coverage
- python\_execnet
- python\_future
- python\_psutil
- python\_py
- pytest (\*)
- pytest\_cov (\*)
- pytest\_forked
- pytest\_flake8 (\*)
- pytest\_session2file (\*)
- pytest\_xdist (\*)
- python\_mccabe
- pyyaml (\*)
- requests
- scipy

- `scons (*)`
- `sqlalchemy`
- `starlink_ast`
- `treecorr`
- `wcslib`
- `ws4py`

In the DM-15495 exercise, we manage to build *afw* moving the 3rd party libraries marked with **(\*)** to the conda environment definition.

Other packages, like for example *starlink\_ast*, may require additional work in order to have them available as a conda package. This can be done in a second time.

## A References

### References

[DMTN-106], Comoretto, G., 2019, DM Release Process, URL <https://dmtn-106.lsst.io/>

## B Acronyms used in this document

Acronym	Description
CI	Continuous Integration
DM	Data Management
DMCCB	DM Change Control Board
DMTN	DM Technical Note
EUPS	Extended Unix Product System
LSST	Large Synoptic Survey Telescope
RFC	Request For Comment