Reduction of MeerKAT interferometric data in PUNCH4NFDI

Nicola Malavasi, Kristof Rozgonyi, Joe Mohr

Collaborators: Manuel Gieffels, Benoît Roland, Matthias Hoeft, Alexander Drabent, Harry Henke

AG Annual Meeting - 12/09/2023

LMU

LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

PUNCH 4NFDI

Outline

- **PUNCH4NFDI**: a new consortium for research infrastructure
- **Reduction**: description of the pipeline workflow
- **MeerKAT interferometric data**: peculiar aspects of the data and scientific interest
- Challenges and achievements

What is PUNCH4NFDI?

The **Particles, Universe, NuClei and Hadrons** (PUNCH) is a consortium funded by the National Research Data Infrastructure (NFDI).

GOAL: set up a new science data platform, offering infrastructures and solutions to advance scientific research.

- Federated (distributed)
- Adhering to EU's FAIR principles
- Providing access and use of data

PUNCH4NFDI resources: Storage and Computing

What is Compute4PUNCH?

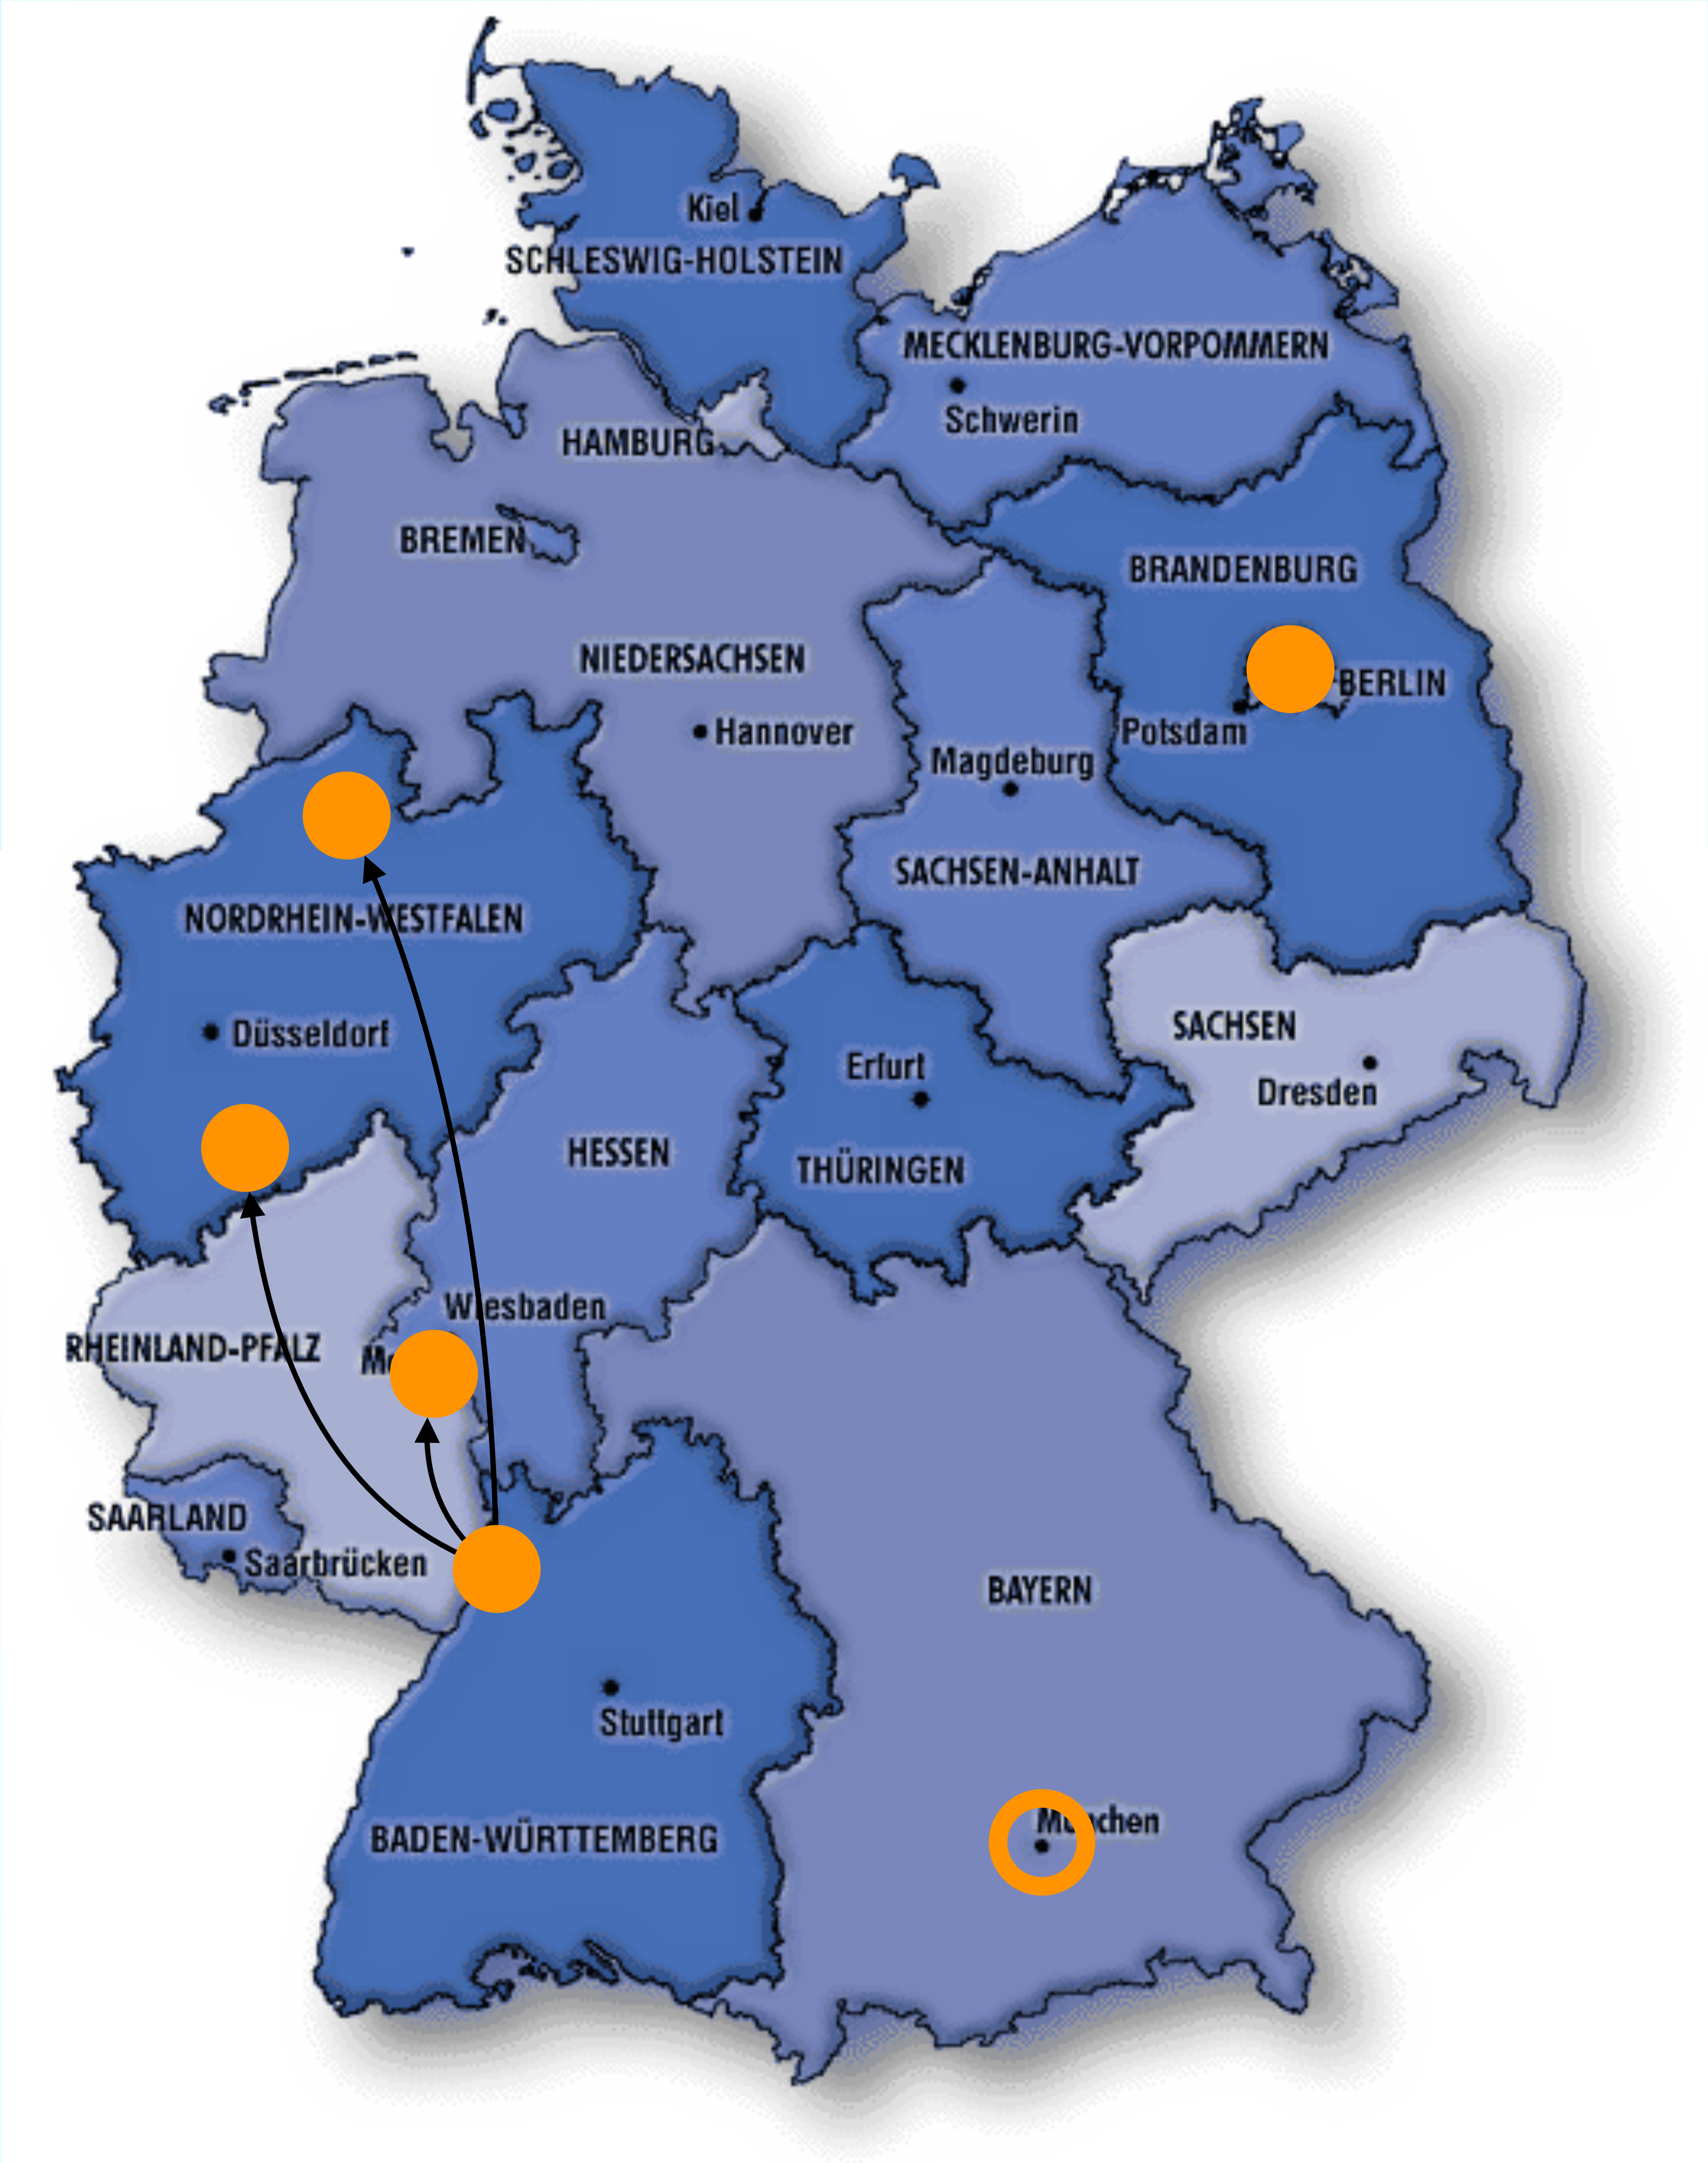

Computing

Compute4PUNCH is the PUNCH4NFDI branch responsible to provide computing infrastructure.

Federated infrastructure, distributed all over Germany, single access point!

What is Compute4PUNCH?

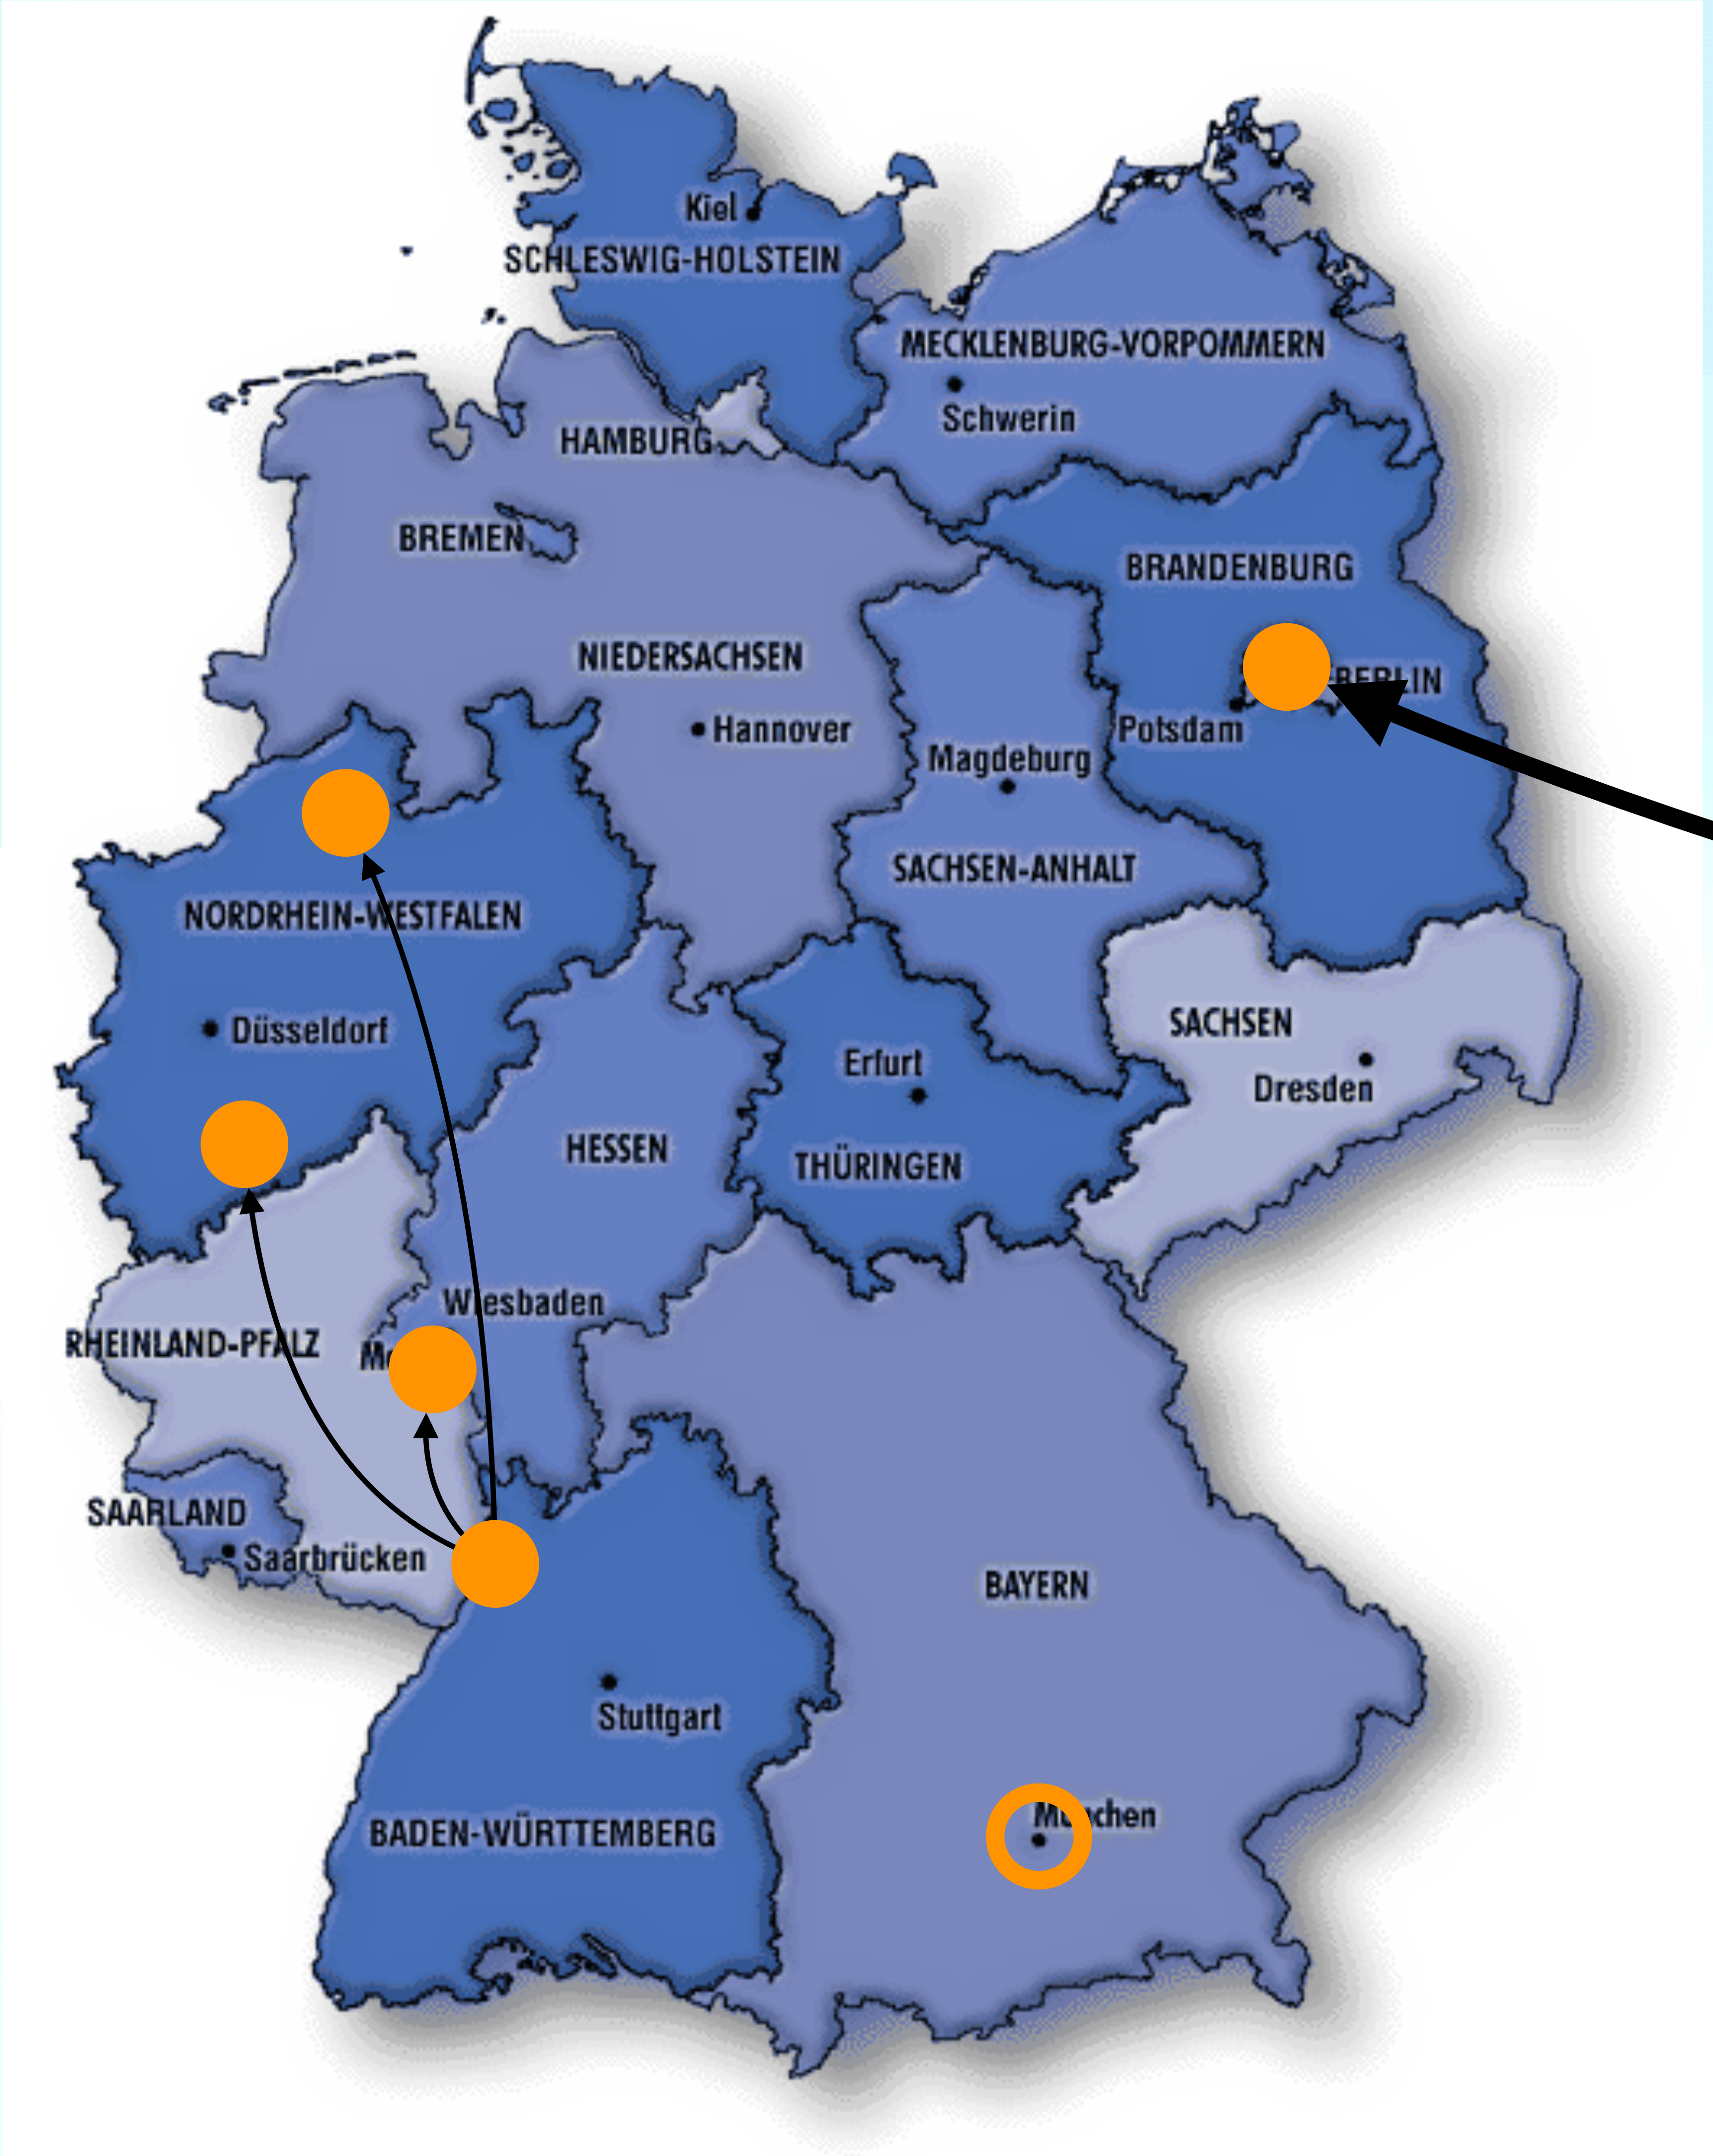

Computing

Compute4PUNCH is the PUNCH4NFDI branch responsible to provide computing infrastructure.

Federated infrastructure, distributed all over Germany, single access point!

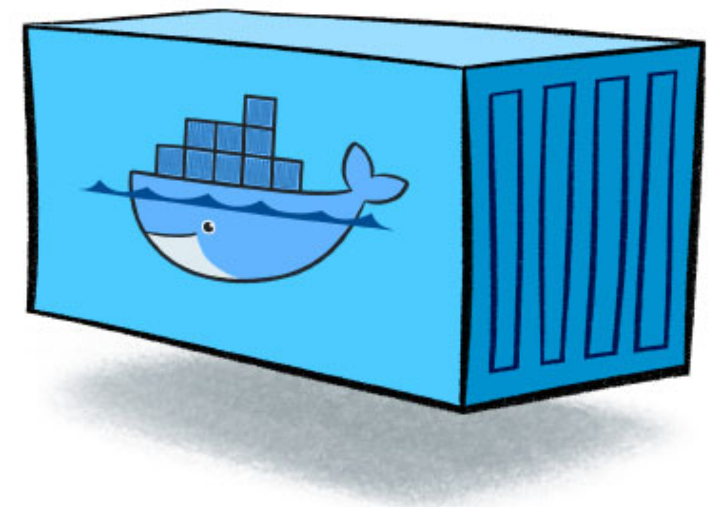

Container

What is Compute4PUNCH?

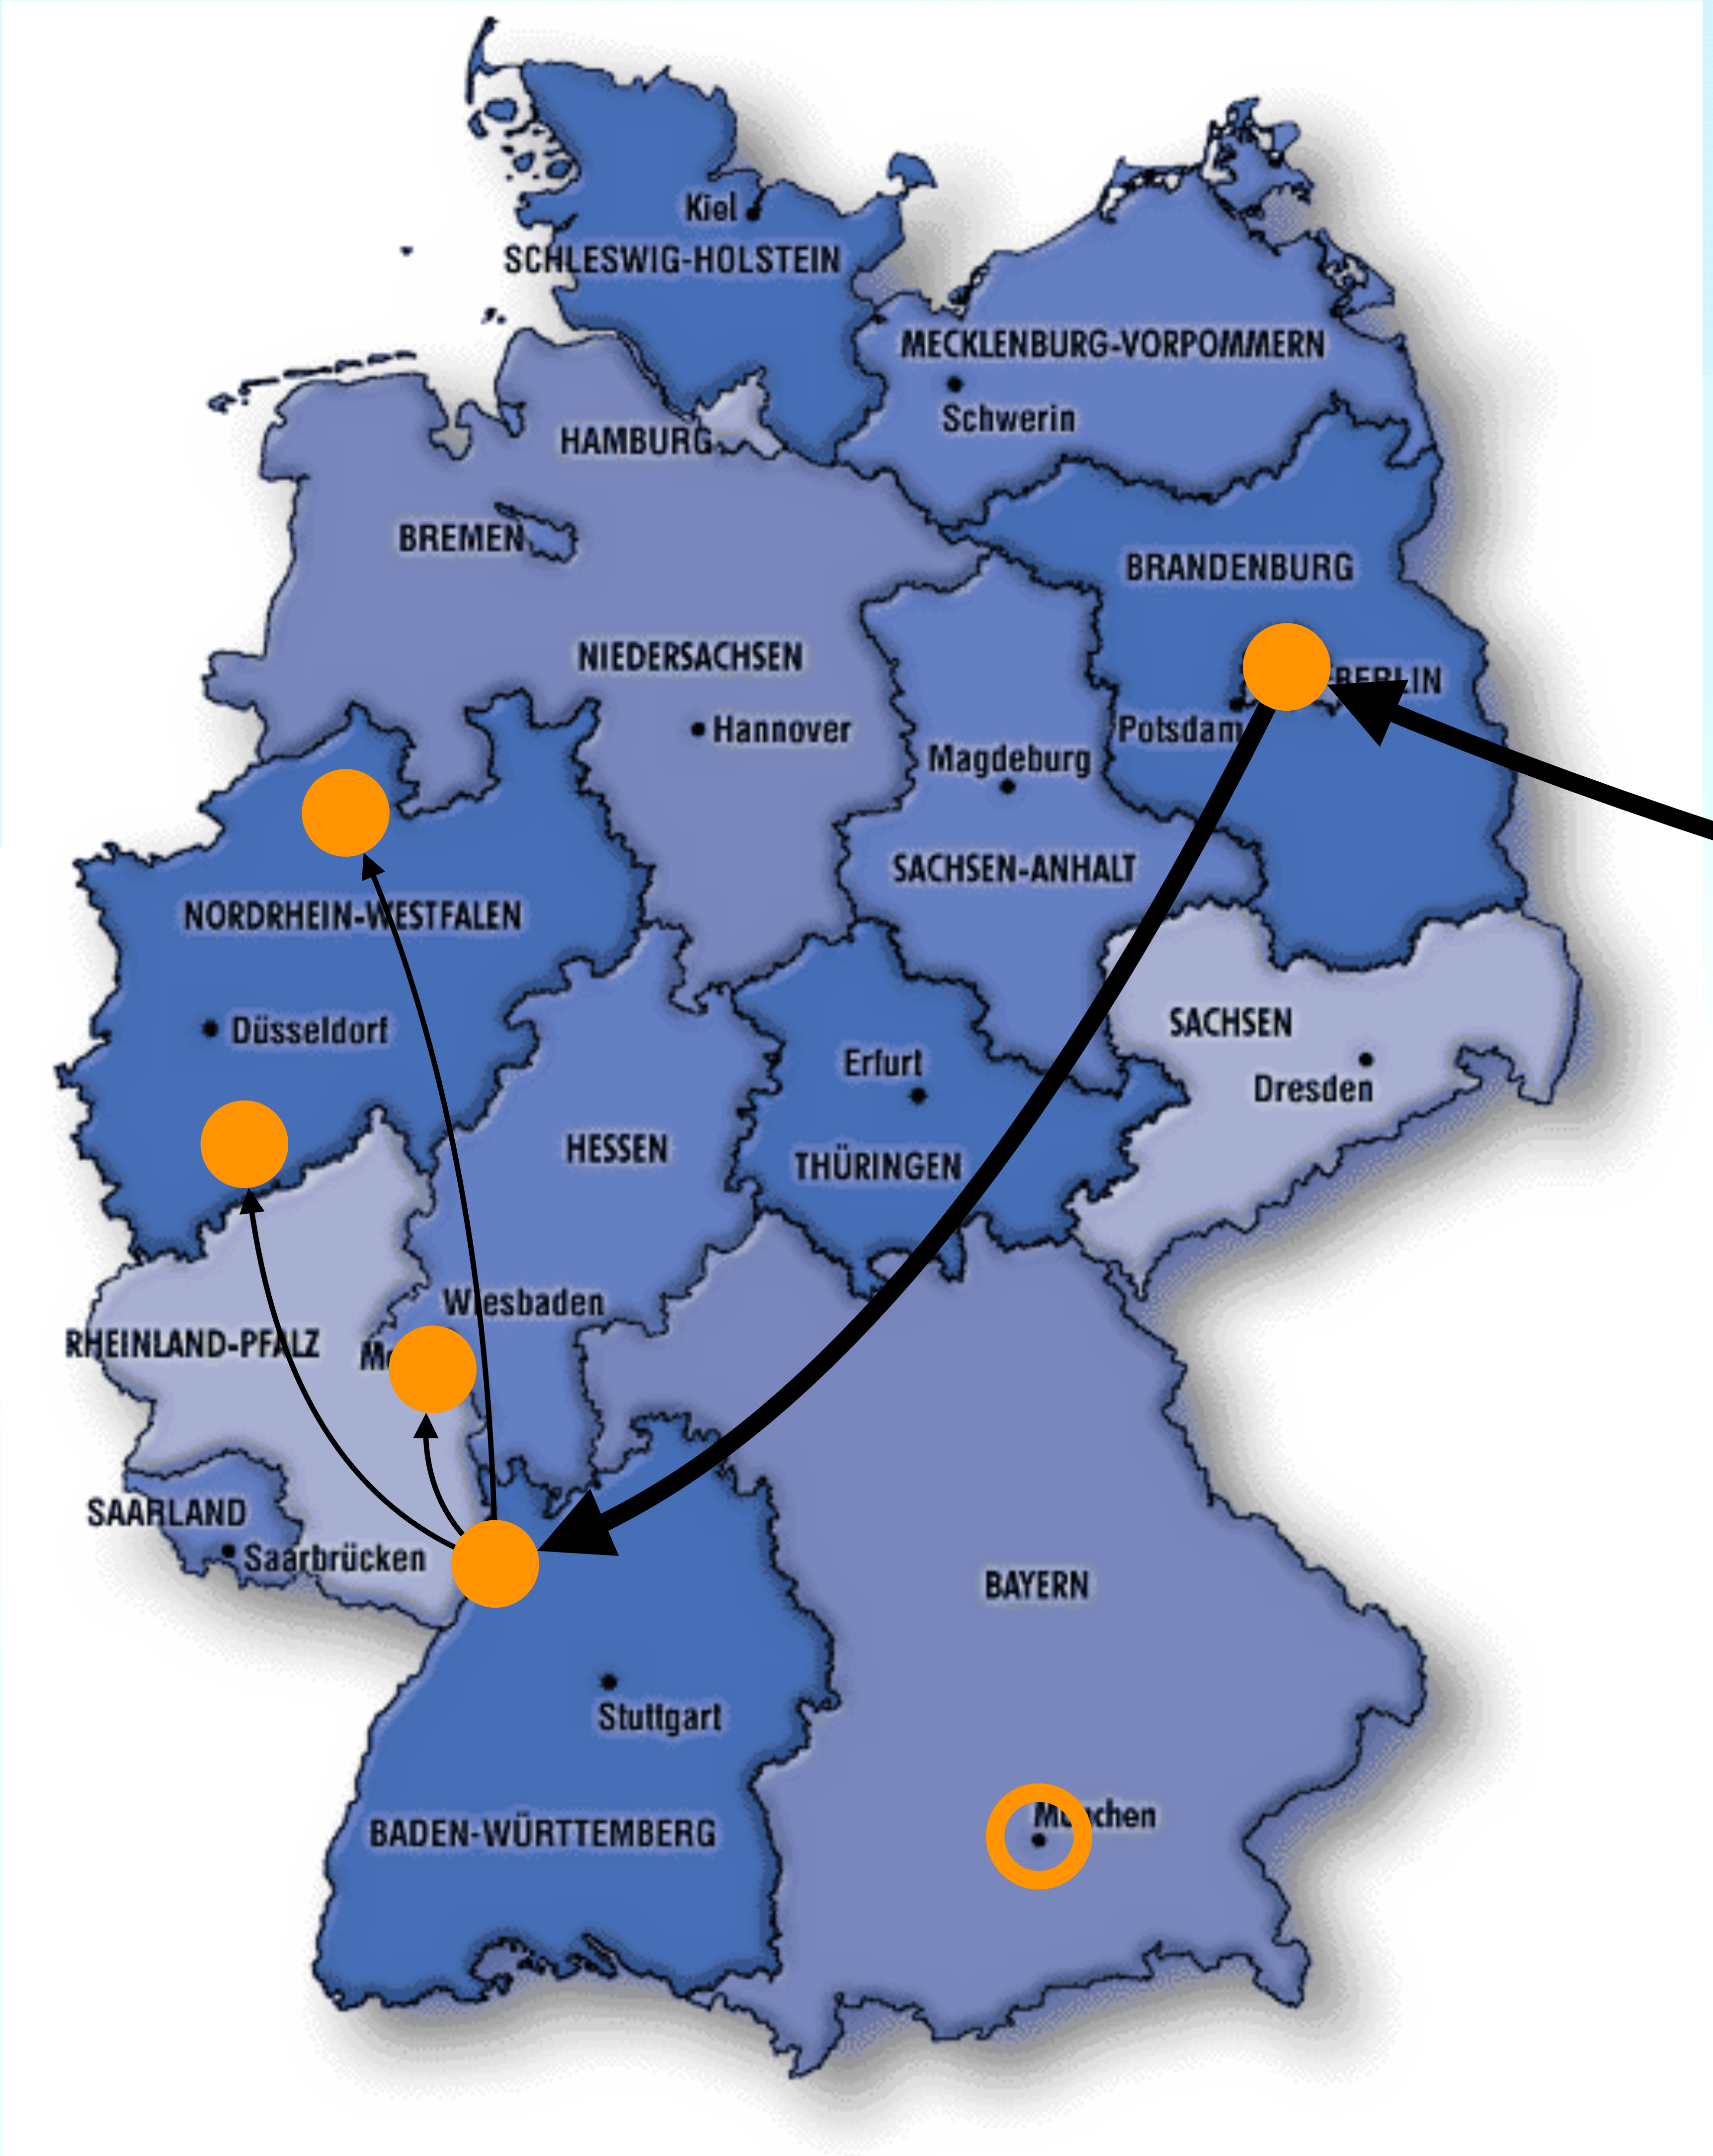

Computing

Compute4PUNCH is the PUNCH4NFDI branch responsible to provide computing infrastructure.

Federated infrastructure, distributed all over Germany, single access point!

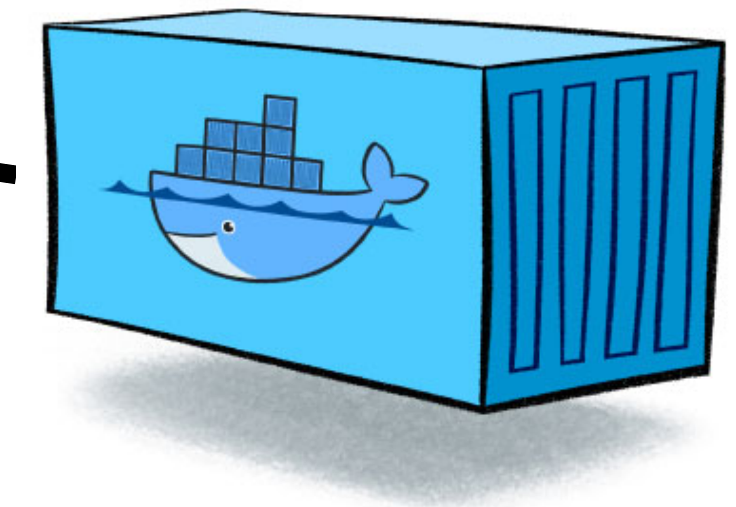

Container

What is Compute4PUNCH?

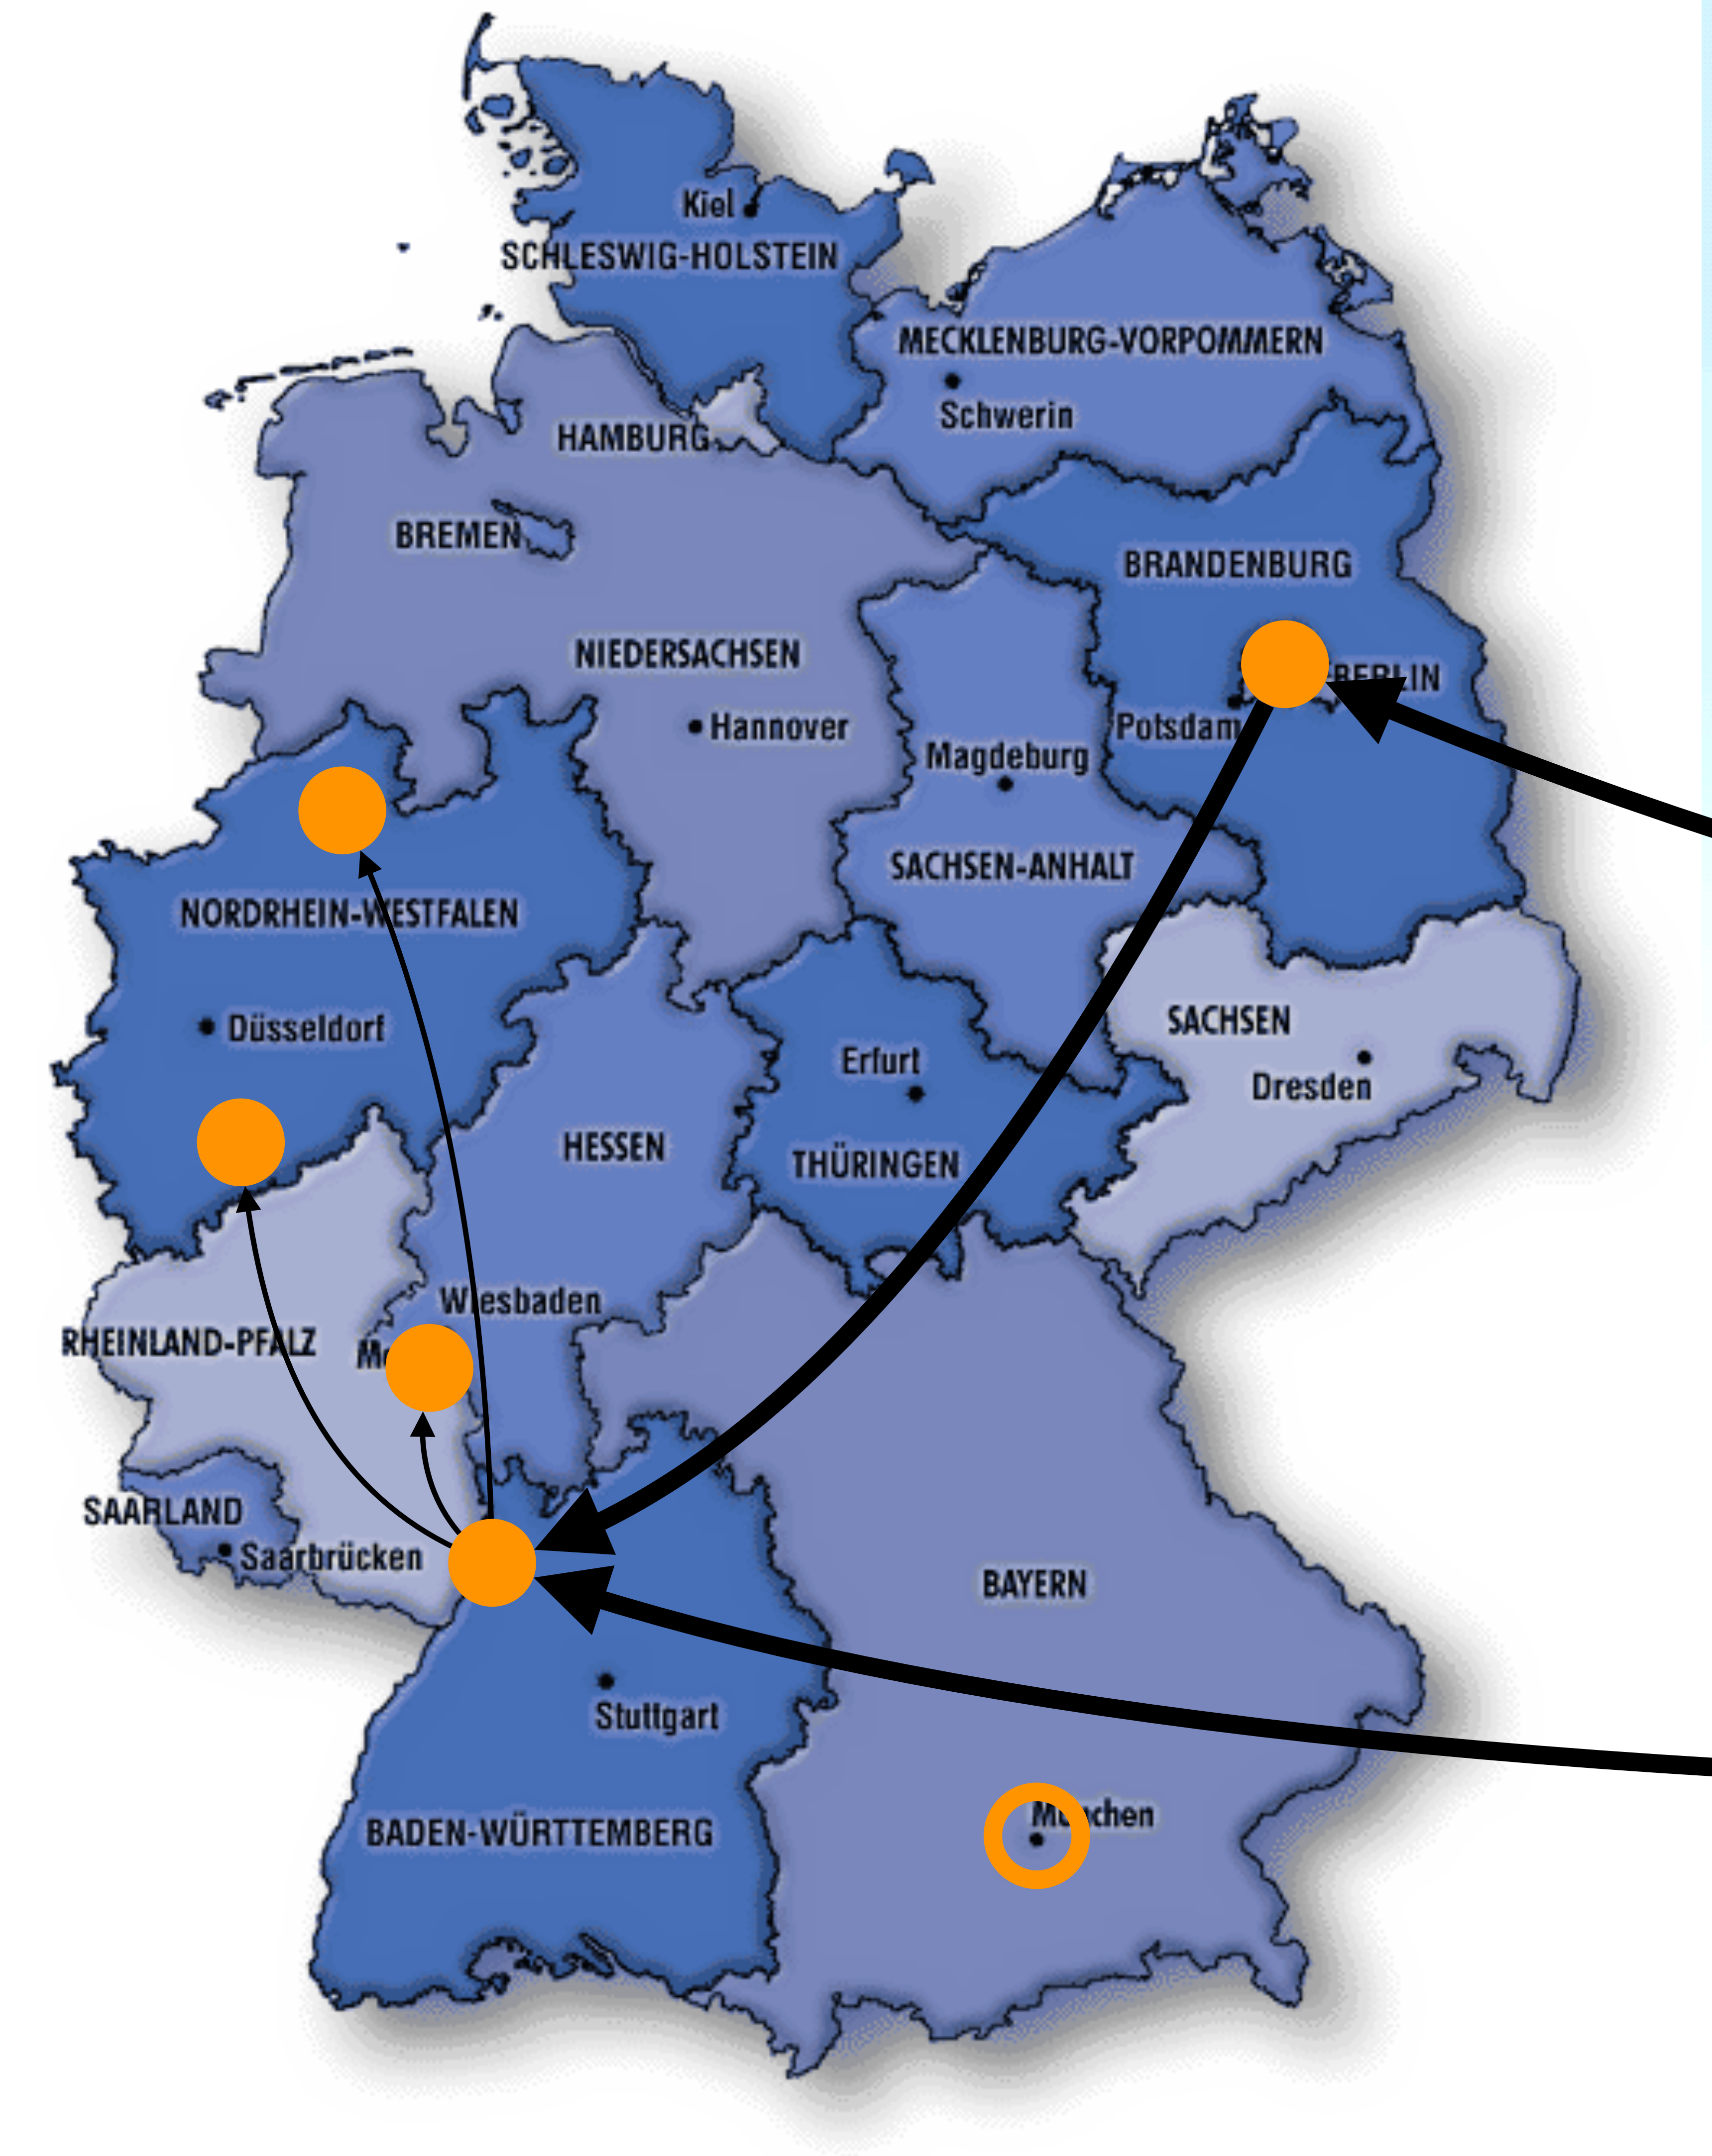

Computing

Compute4PUNCH is the PUNCH4NFDI branch responsible to provide computing infrastructure.

Federated infrastructure, distributed all over Germany, single access point!

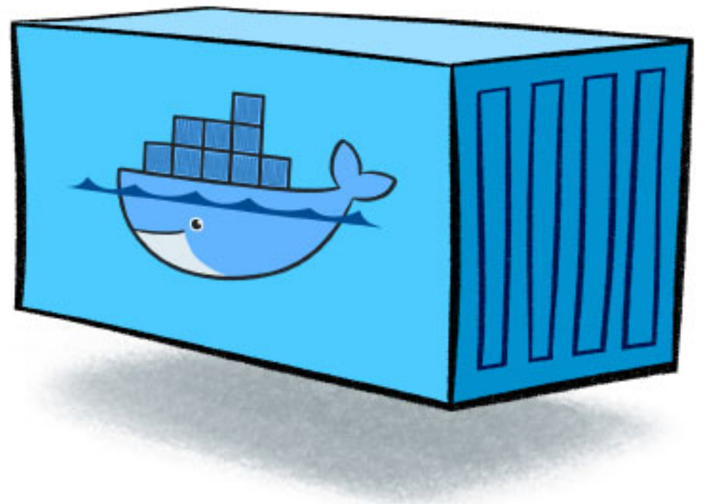

Container

Execution script

HTCondor scheduler

What is Storage4PUNCH?

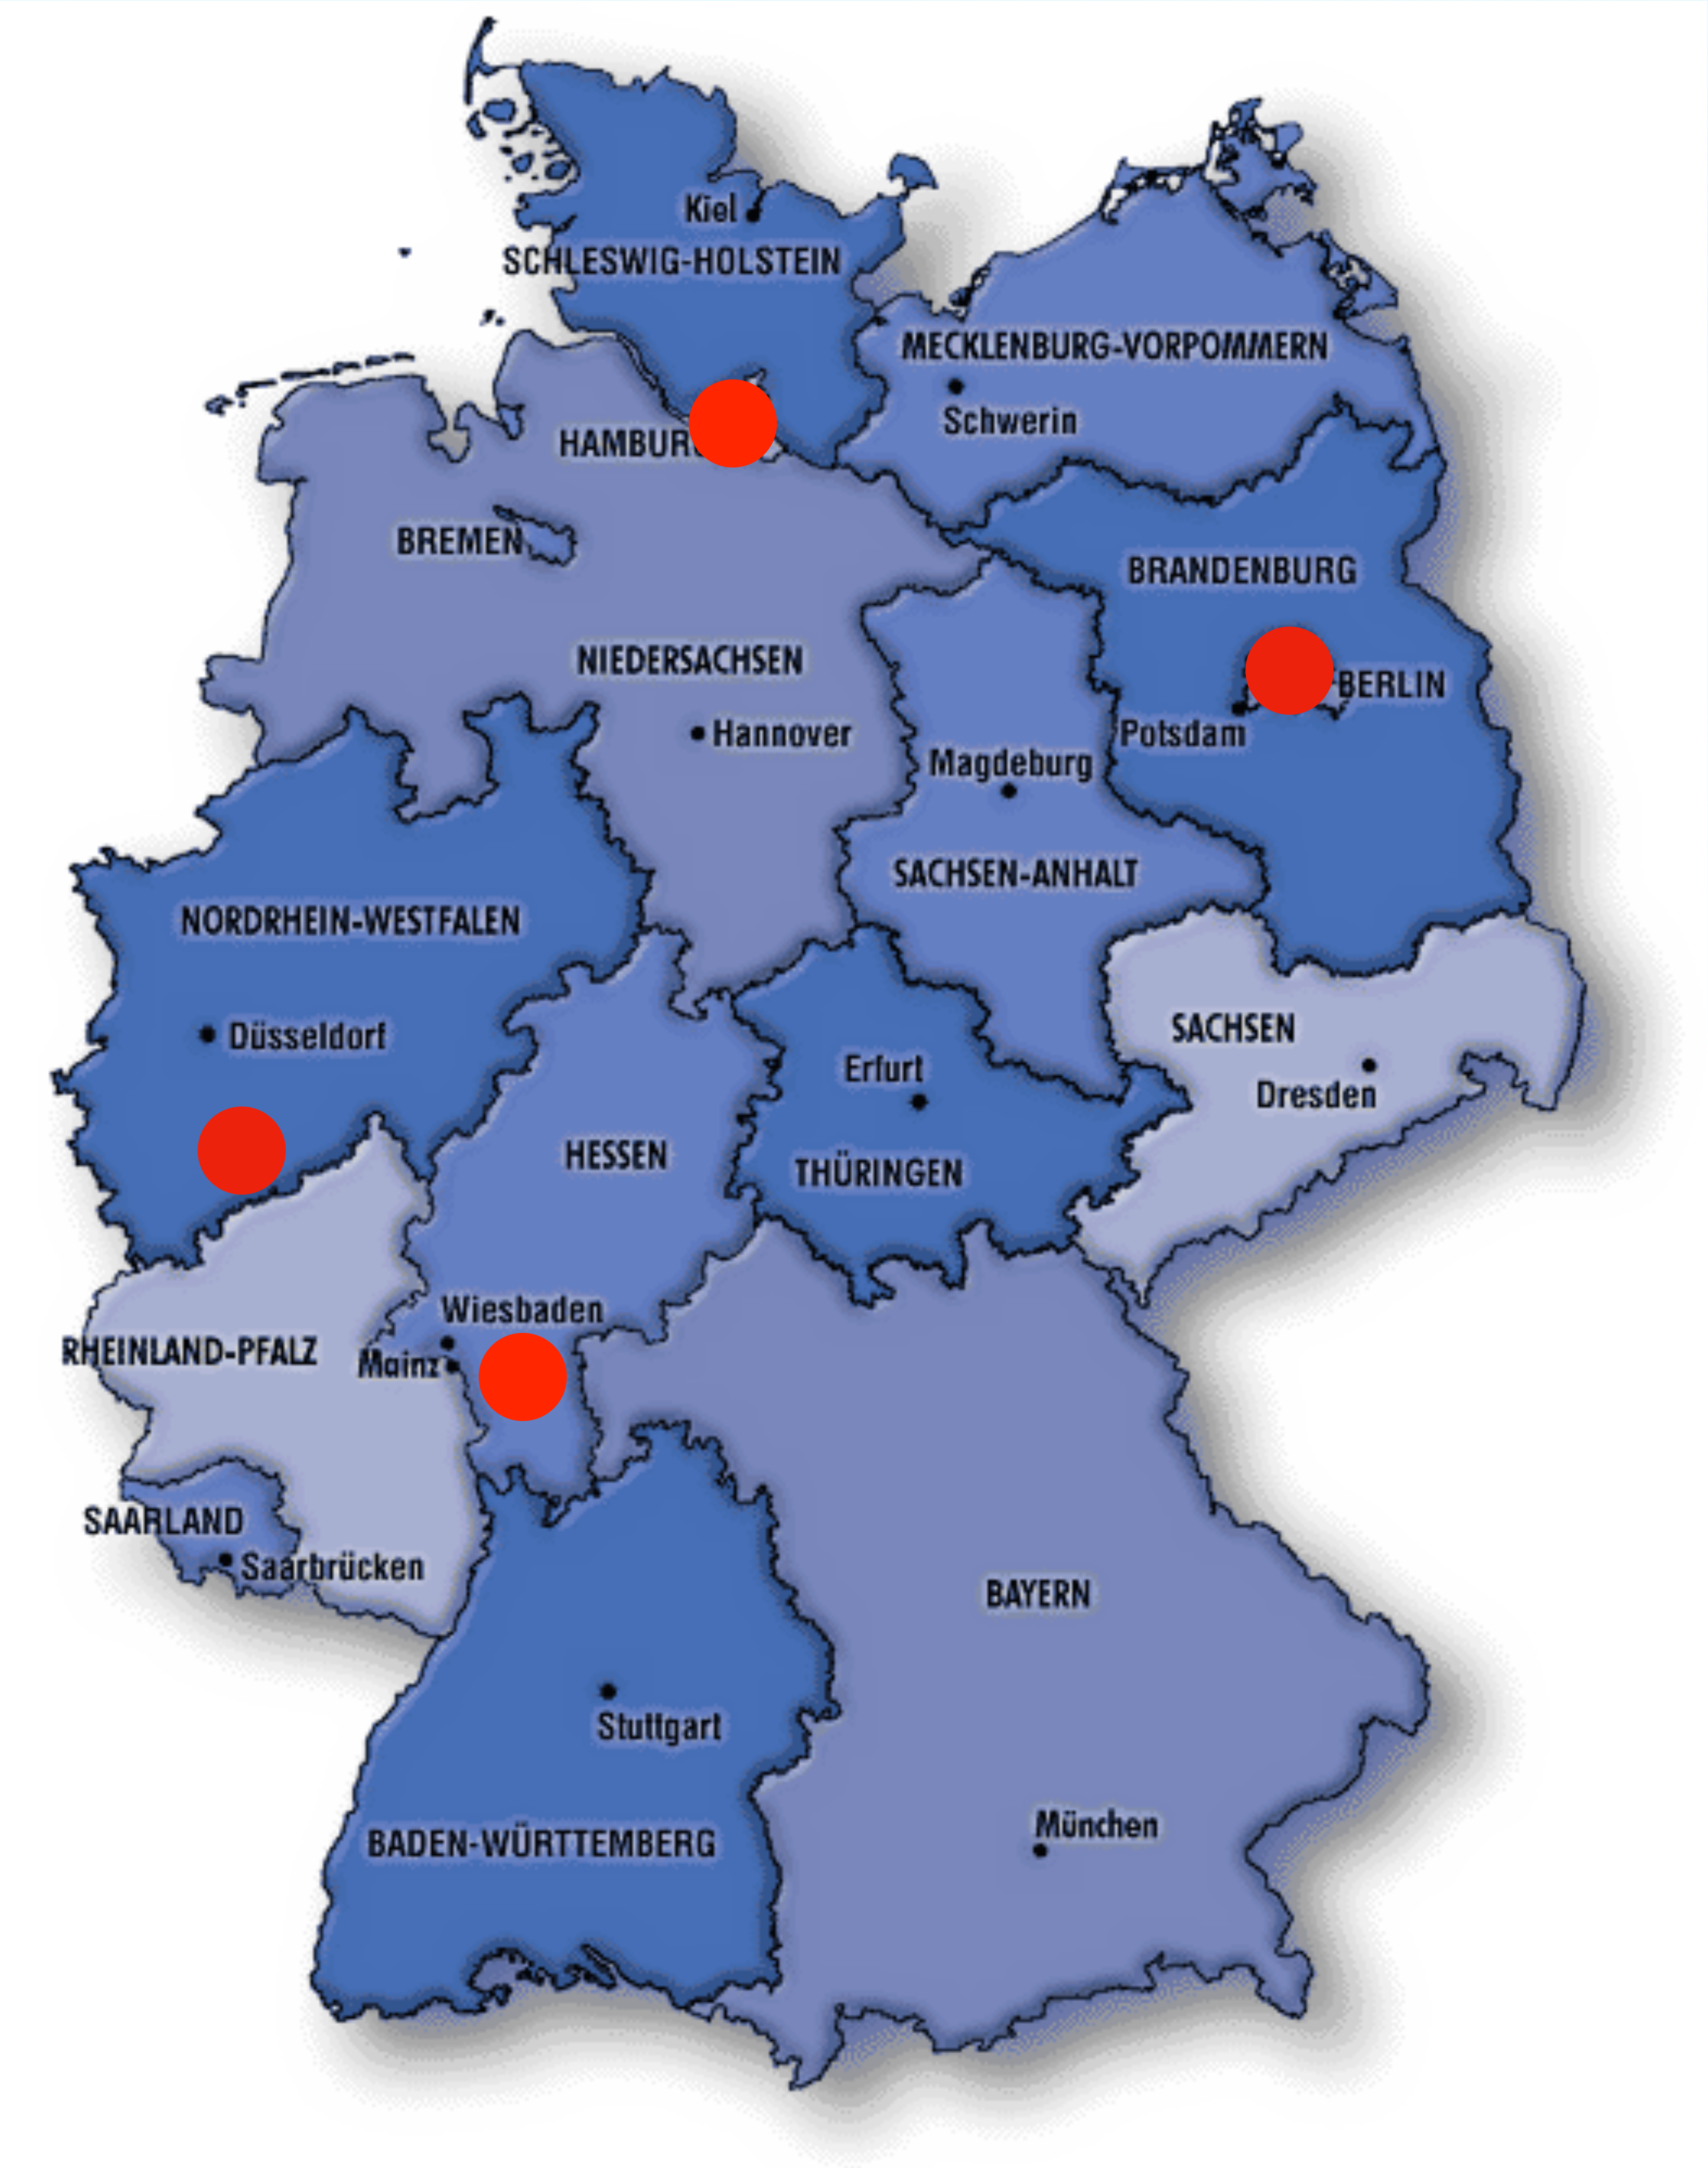

Storage

Storage4PUNCH is the PUNCH4NFDI branch responsible to provide storage infrastructure.

Federated infrastructure, distributed all over Germany, single access point!

Example: dCache instance

Streaming of data, mounting of devices

Future: streaming directly into computing nodes

Beyond computing and storage

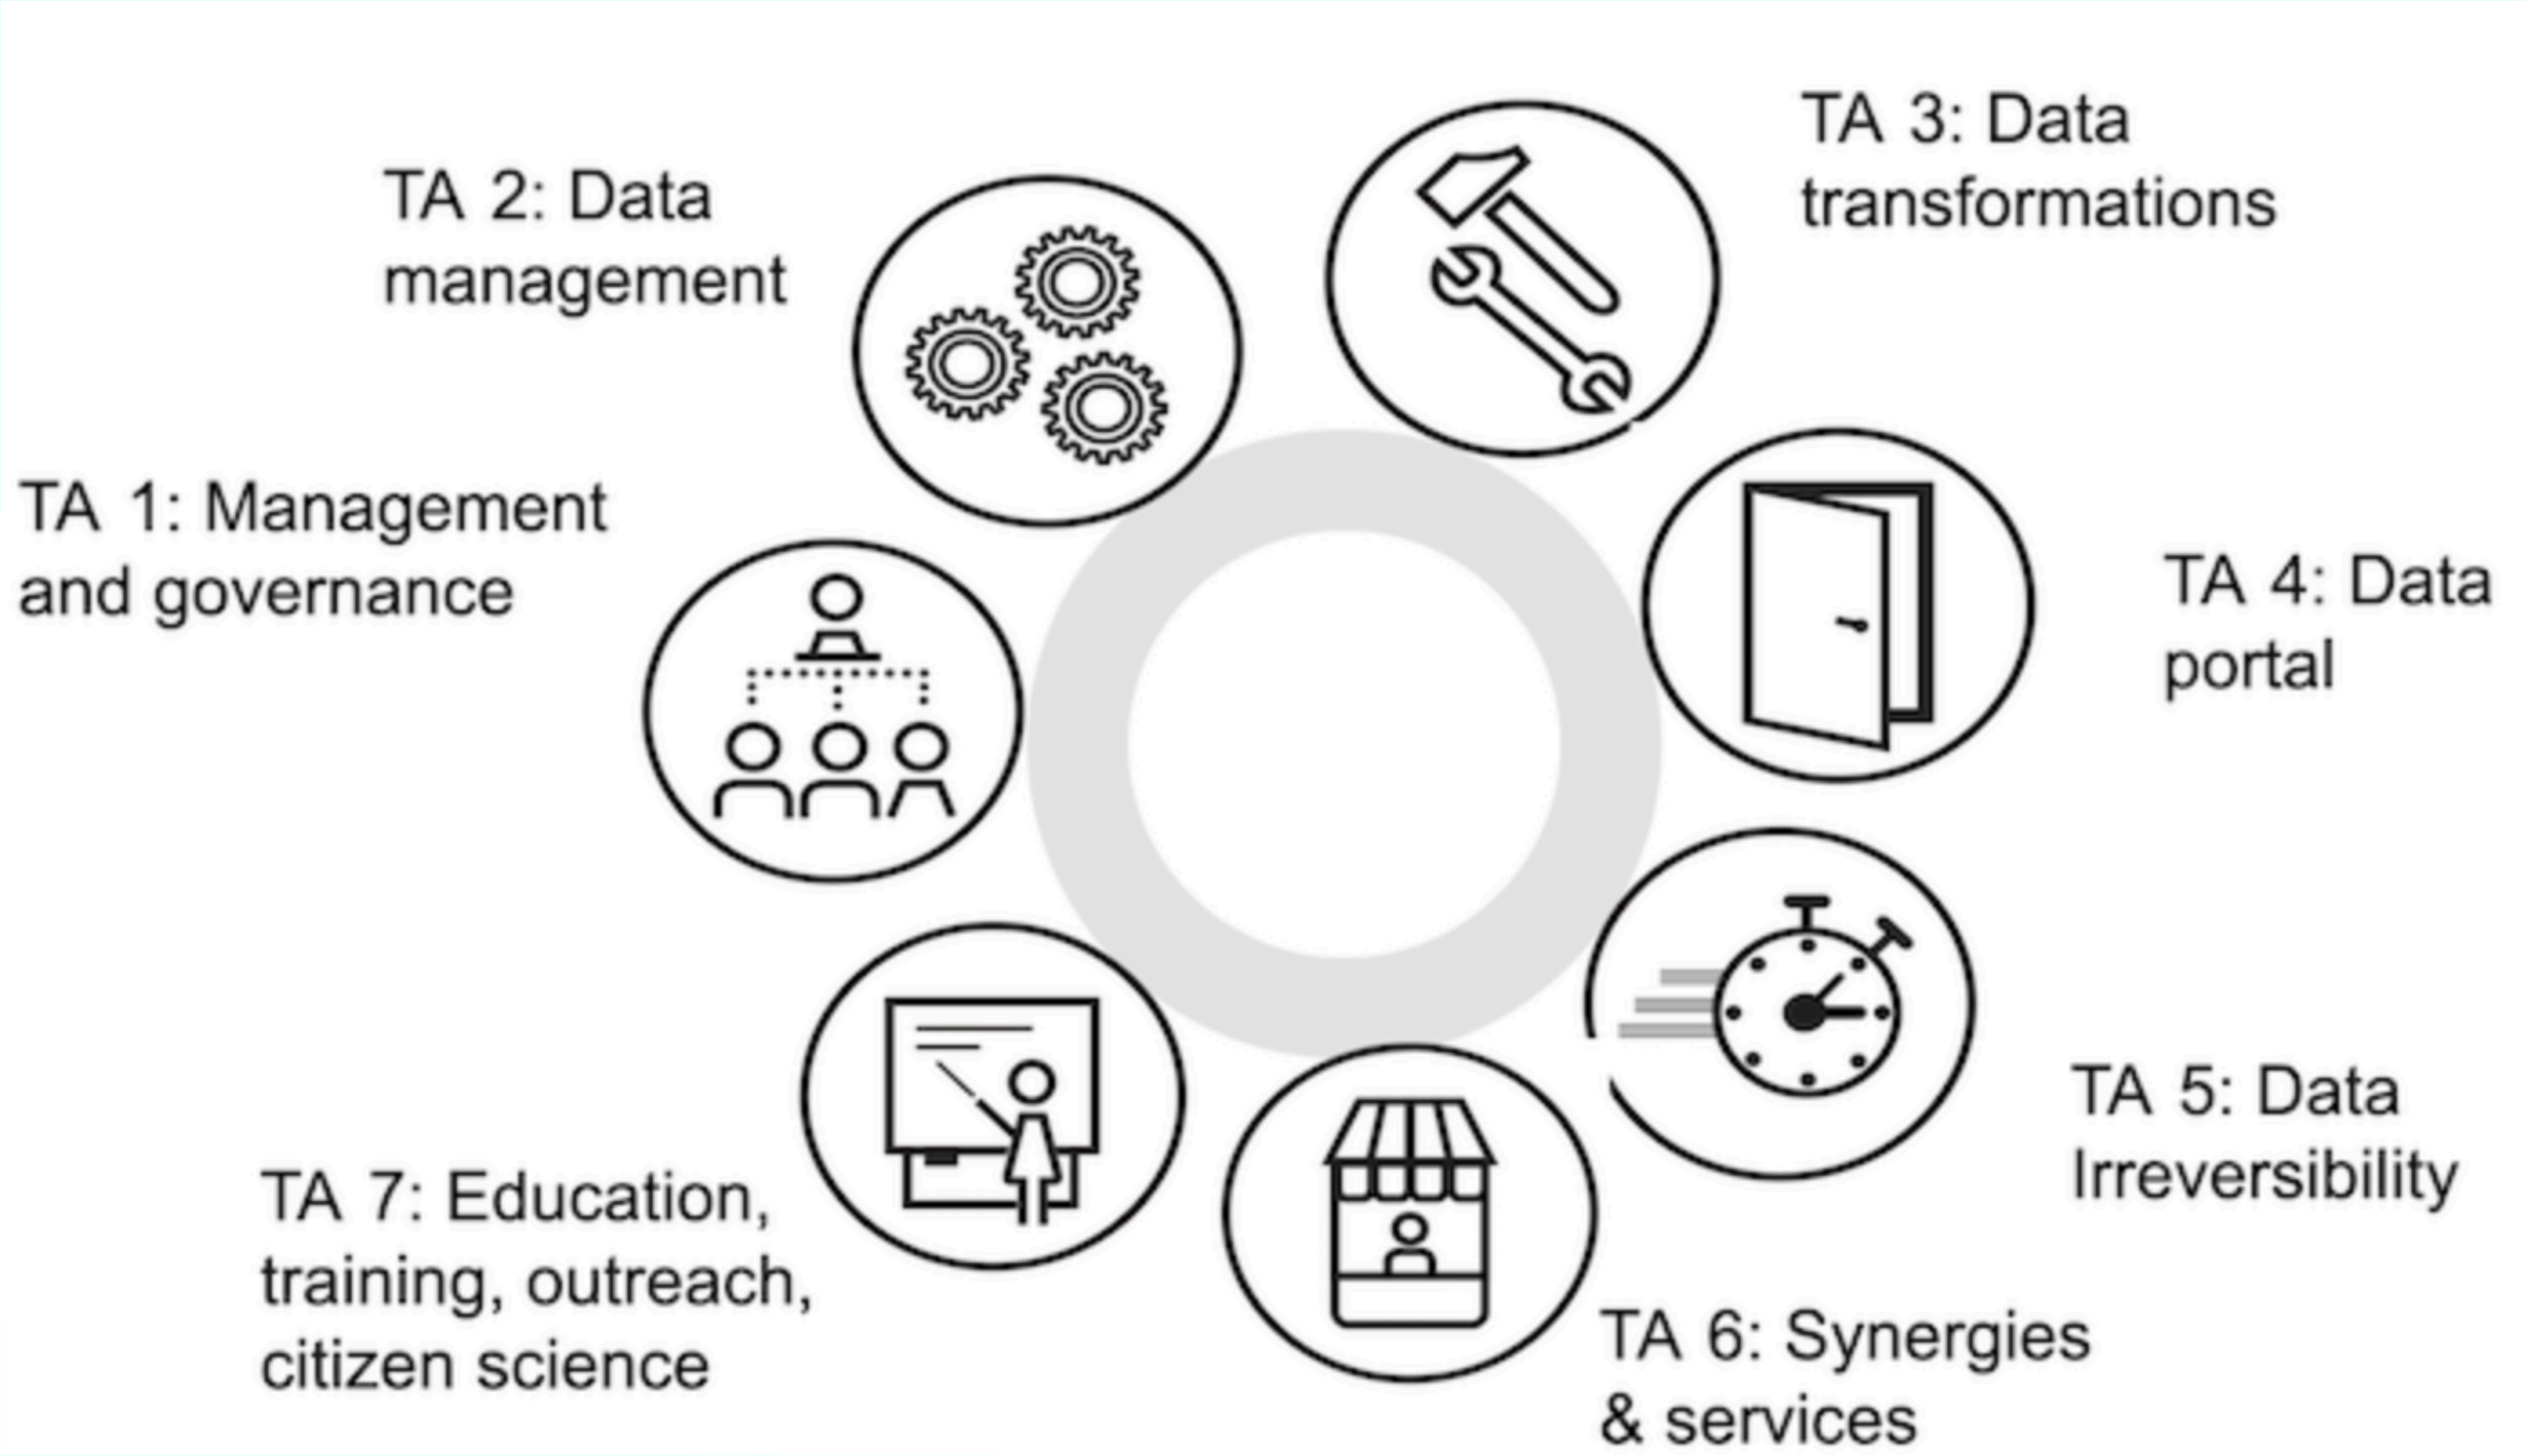

PUNCH is much more than just federated computing and storage resources.

- **Research product availability through portal**
- **Metadata of stored products**
- **Workflow management**
- **Science reproducibility**

A use case for this infrastructure

A scientific analysis that tests, provides feedback to the development of this infrastructure, and exploits synergy among its various parts.

Requirements:

- Large data volume/need for intensive data access $\longrightarrow$ to test storage
- Complex workflow $\longrightarrow$ to test workflow management
- Computationally intensive/need for computation resources $\longrightarrow$ to test computing
- Delivering results/making results available through PUNCH $\longrightarrow$ important for metadata section, science reproducibility (see Yori's talk)

Identified use case: radio data reduction

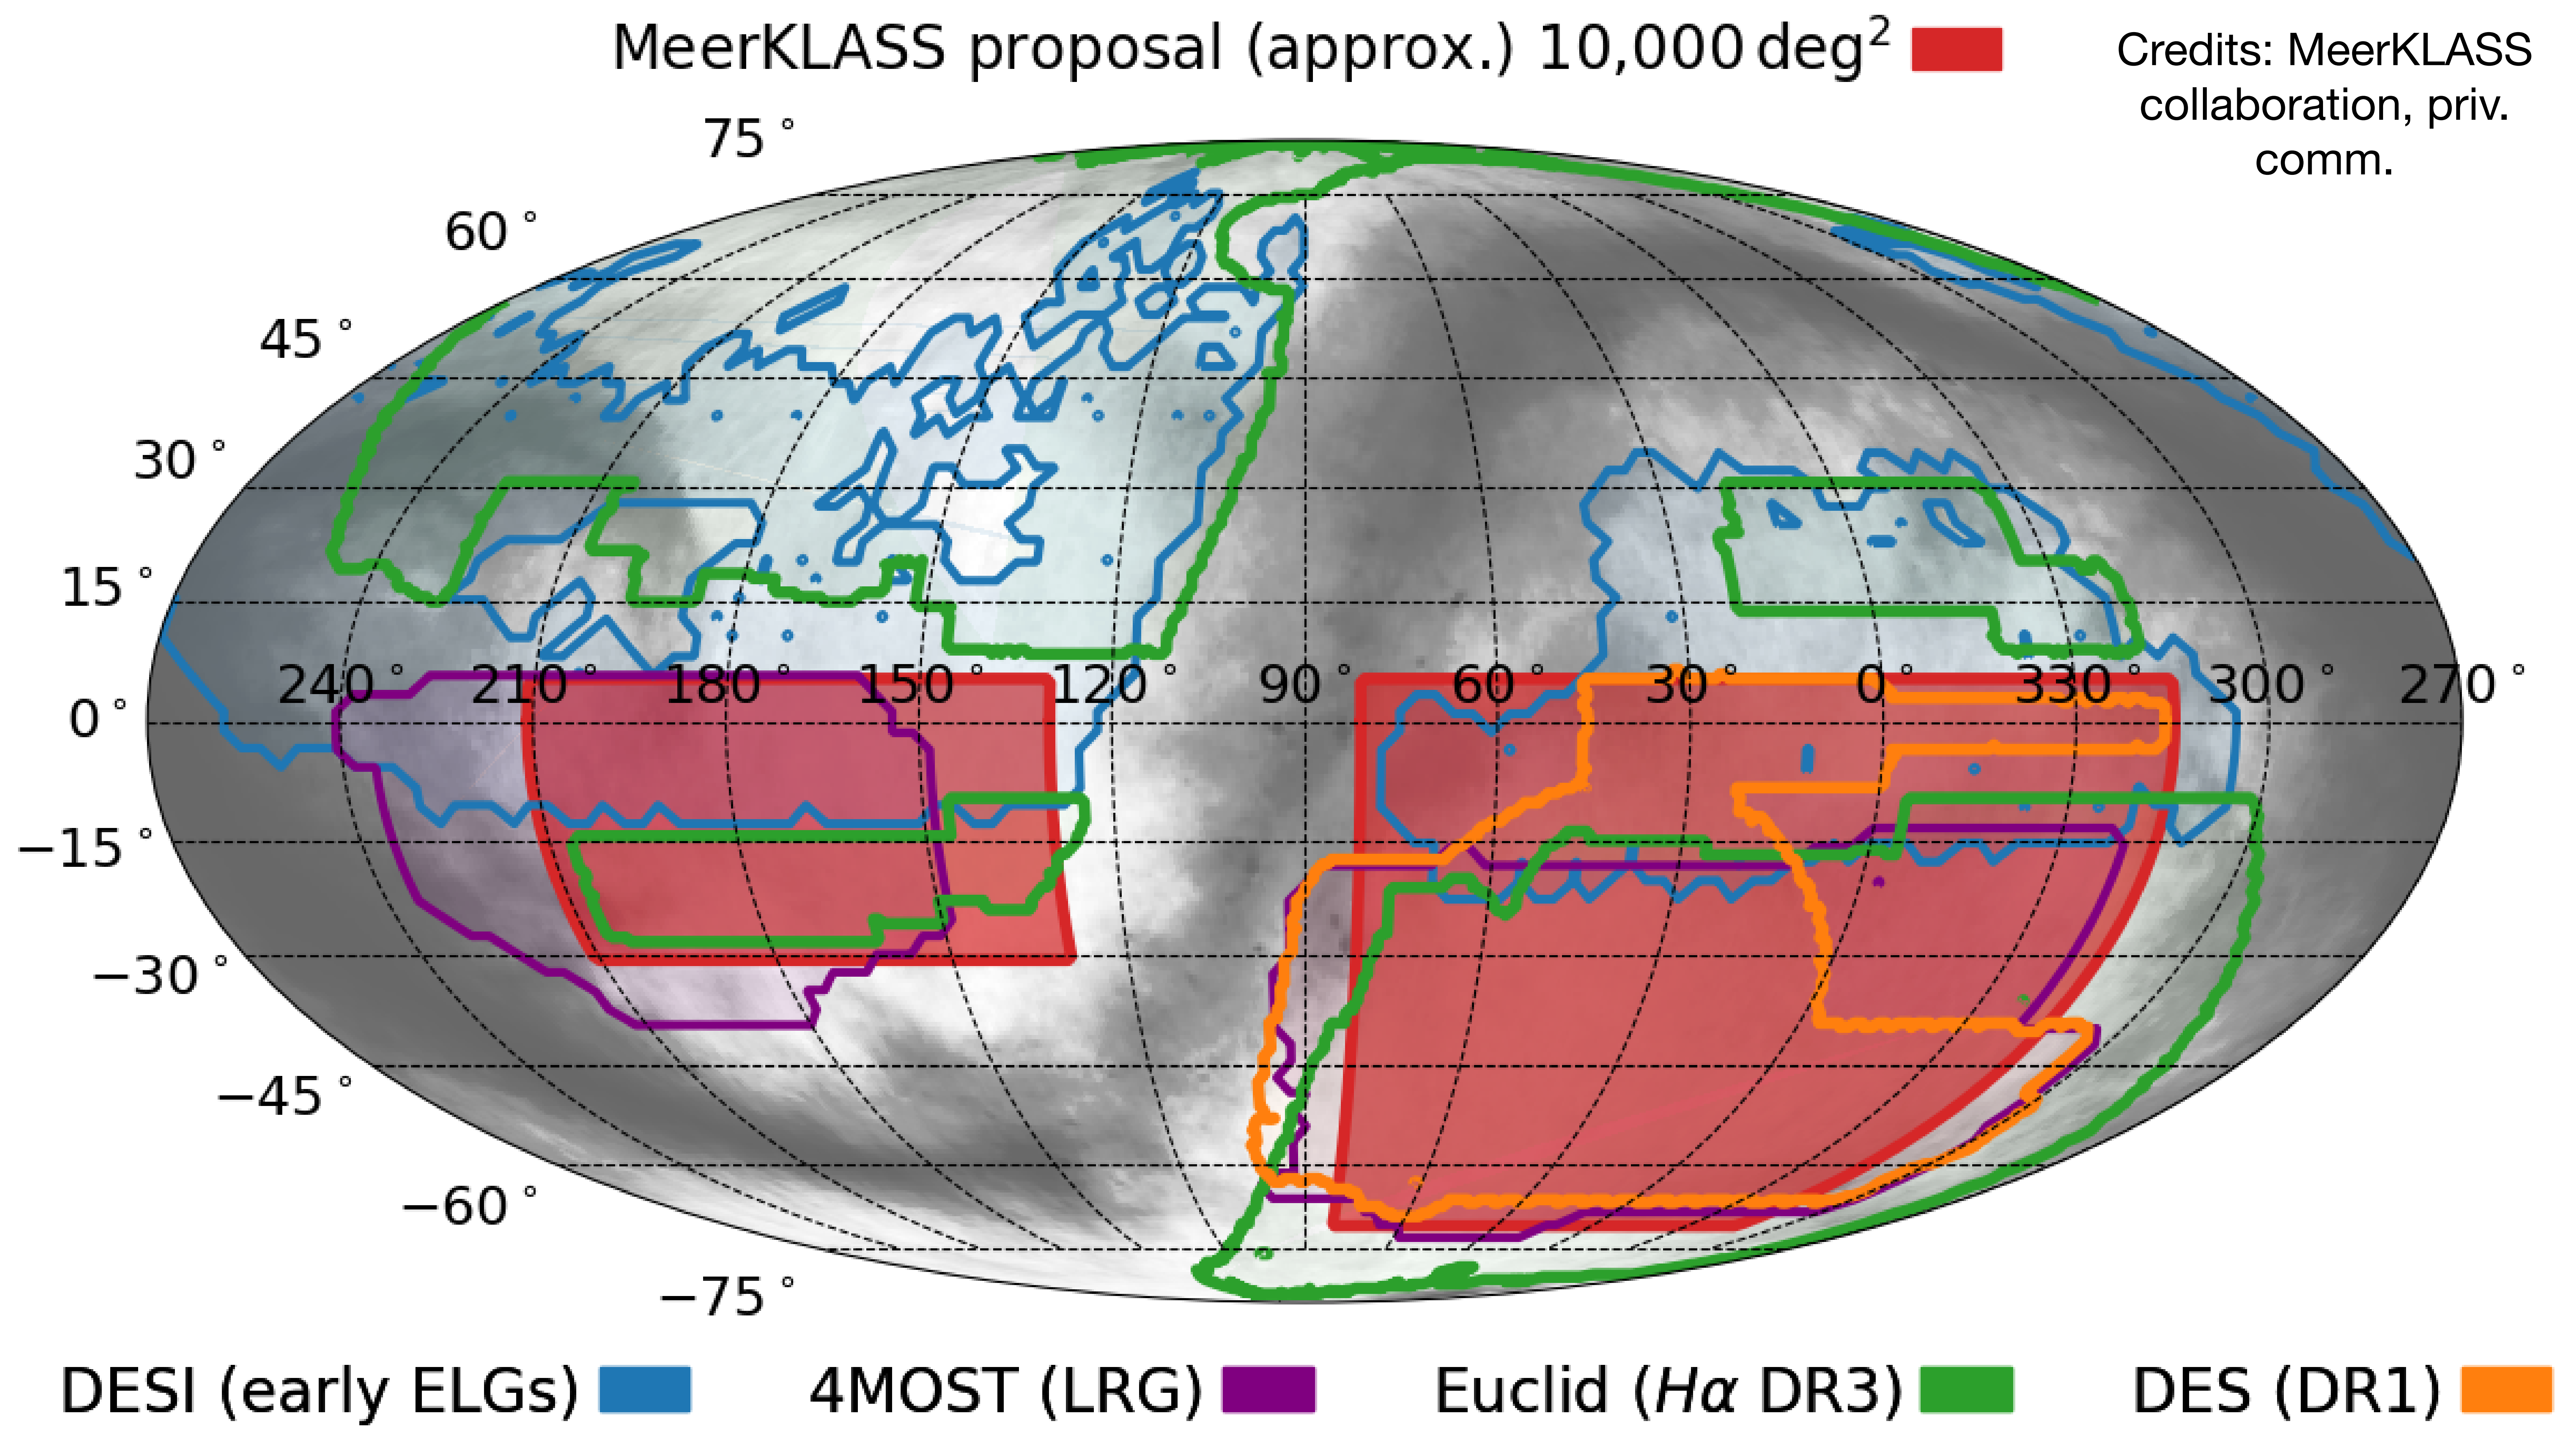

Fig. 2: The proposed MeerKLASS survey footprint (shaded red).

The MeerKLASS survey

- 10'000 sqdeg.
- ~2500 h of observations.
- 300 sqdeg already available (MeerKLASS pilot observations).
- New observing mode for MeerKAT: On The Fly interferometric mode.

Radio-interferometric data

- Radio interferometry works by combining the signal of different radio telescopes observing the same object.
- This allows us to increase sensitivity and angular resolution at the expense of the field of view.
- Usually radio interferometers are used in fixed pointing mode: they observe their target and don’t move.

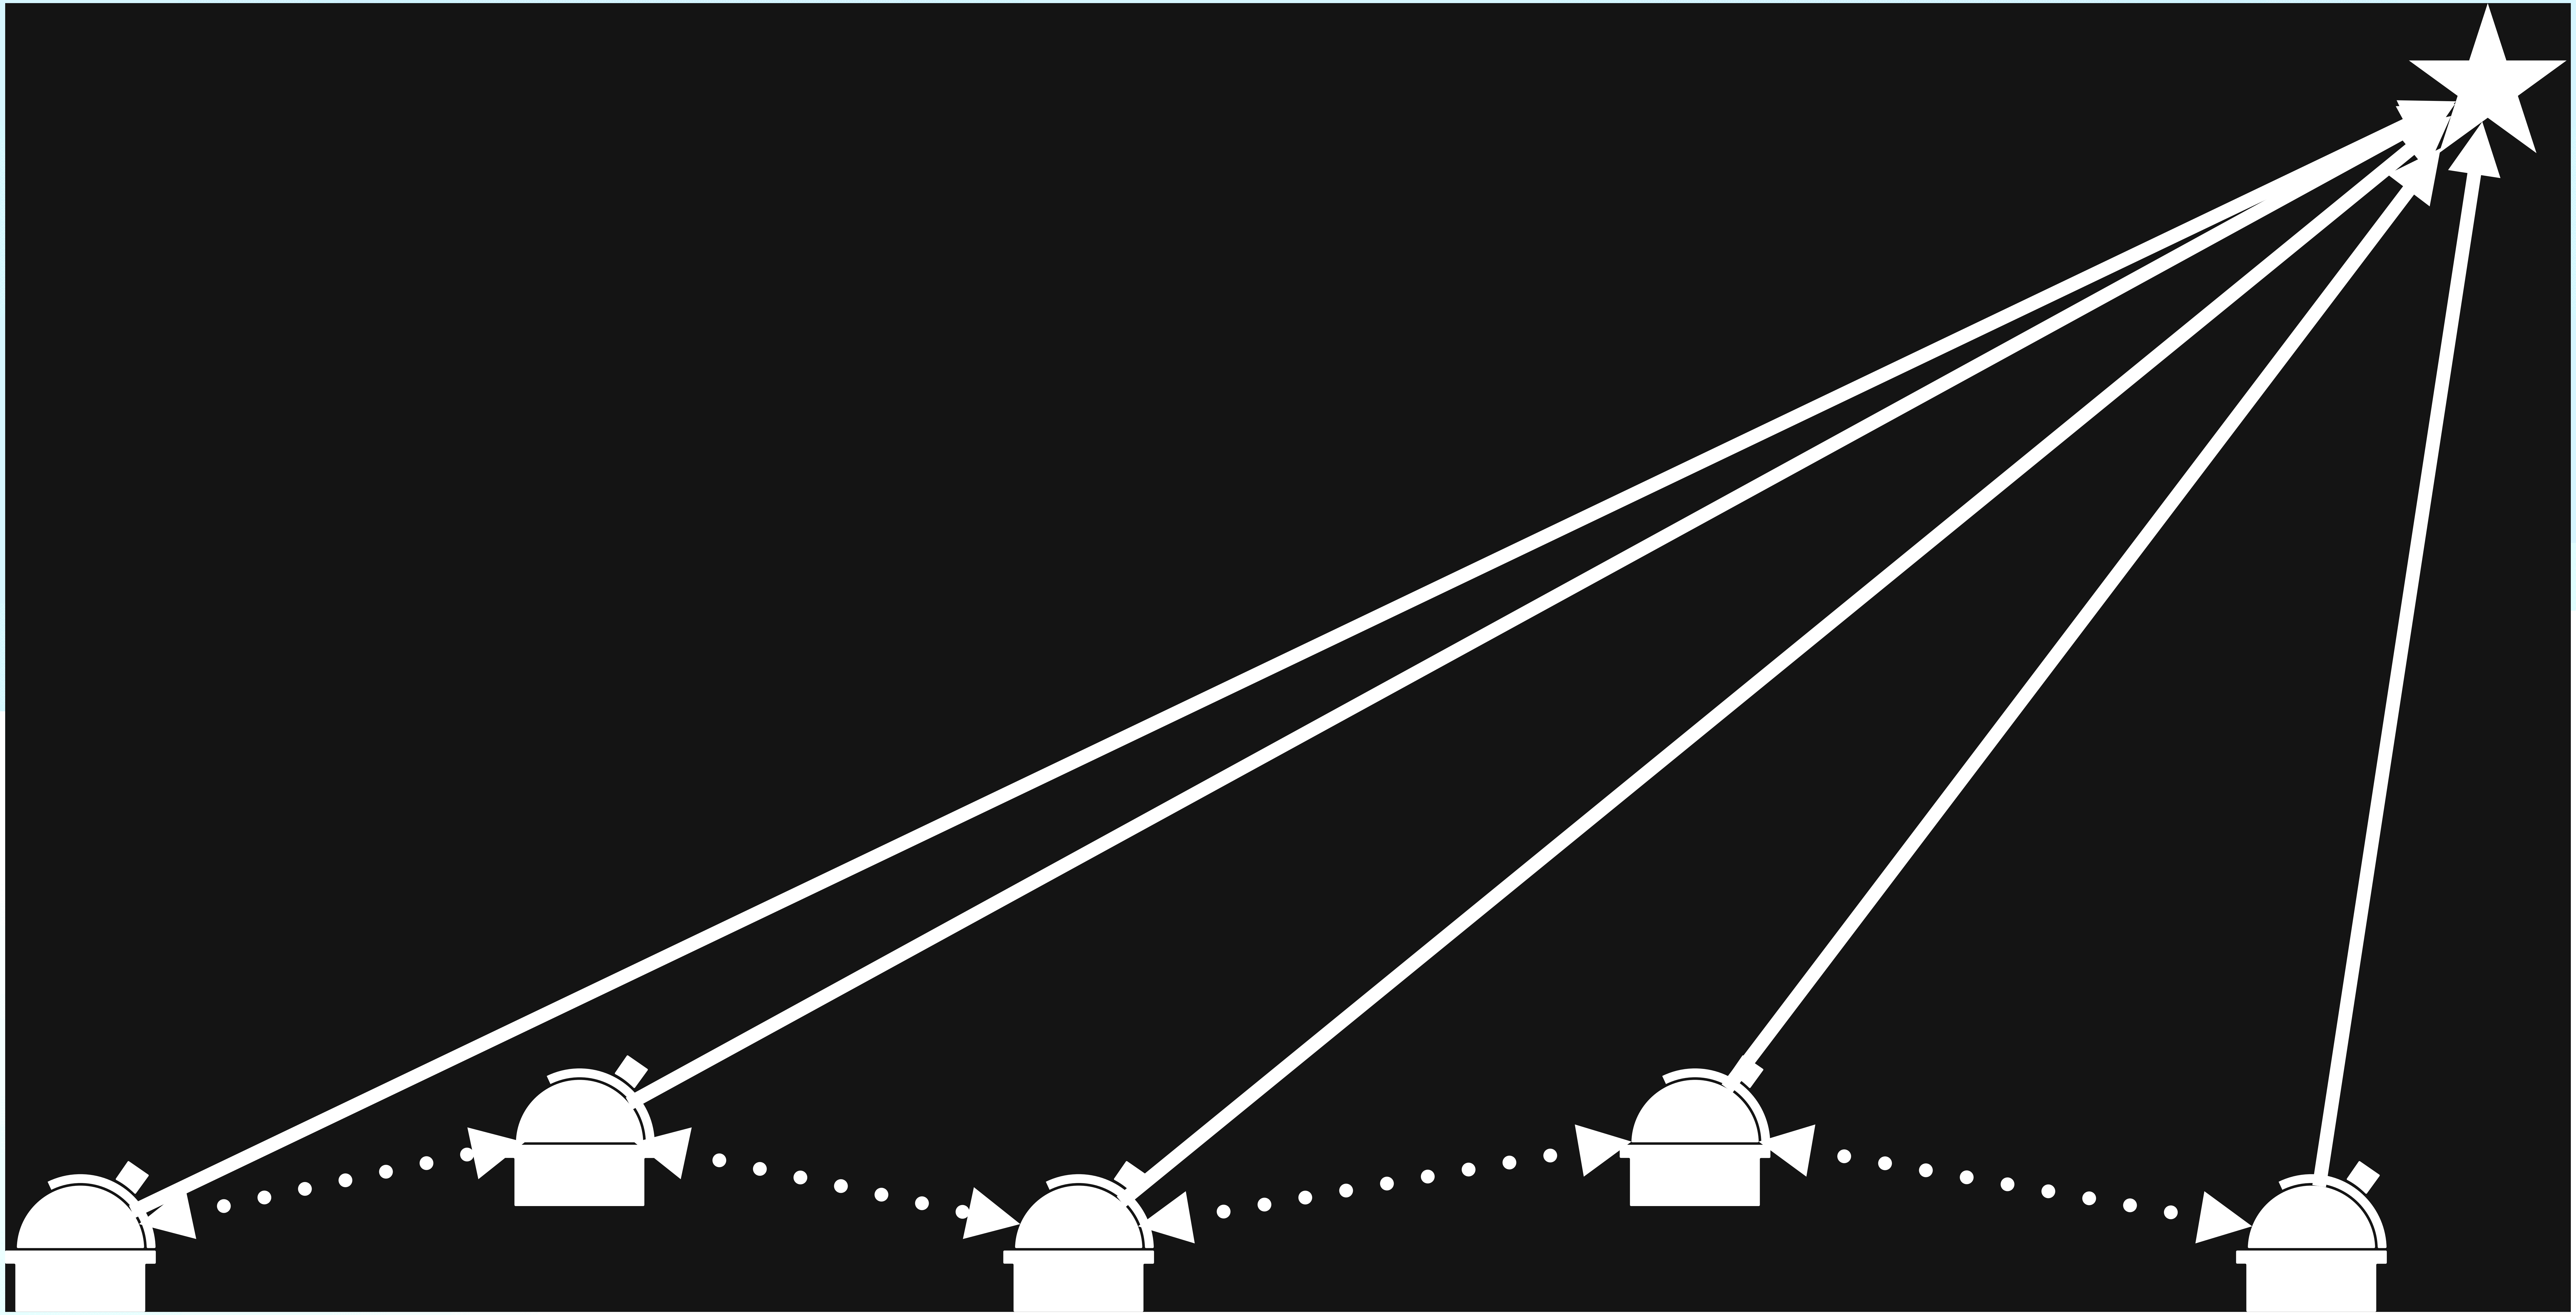

New method: on the fly interferometry

- D-MeerKAT collaboration (Kristof Rozgonyi) has developed a new way to operate a radio-interferometer: scanning mode.
- The radio-interferometer is not observing a fixed target but scanning the sky.

New method: on the fly interferometry

- D-Mee… y to operat…
- The ra… the sky.

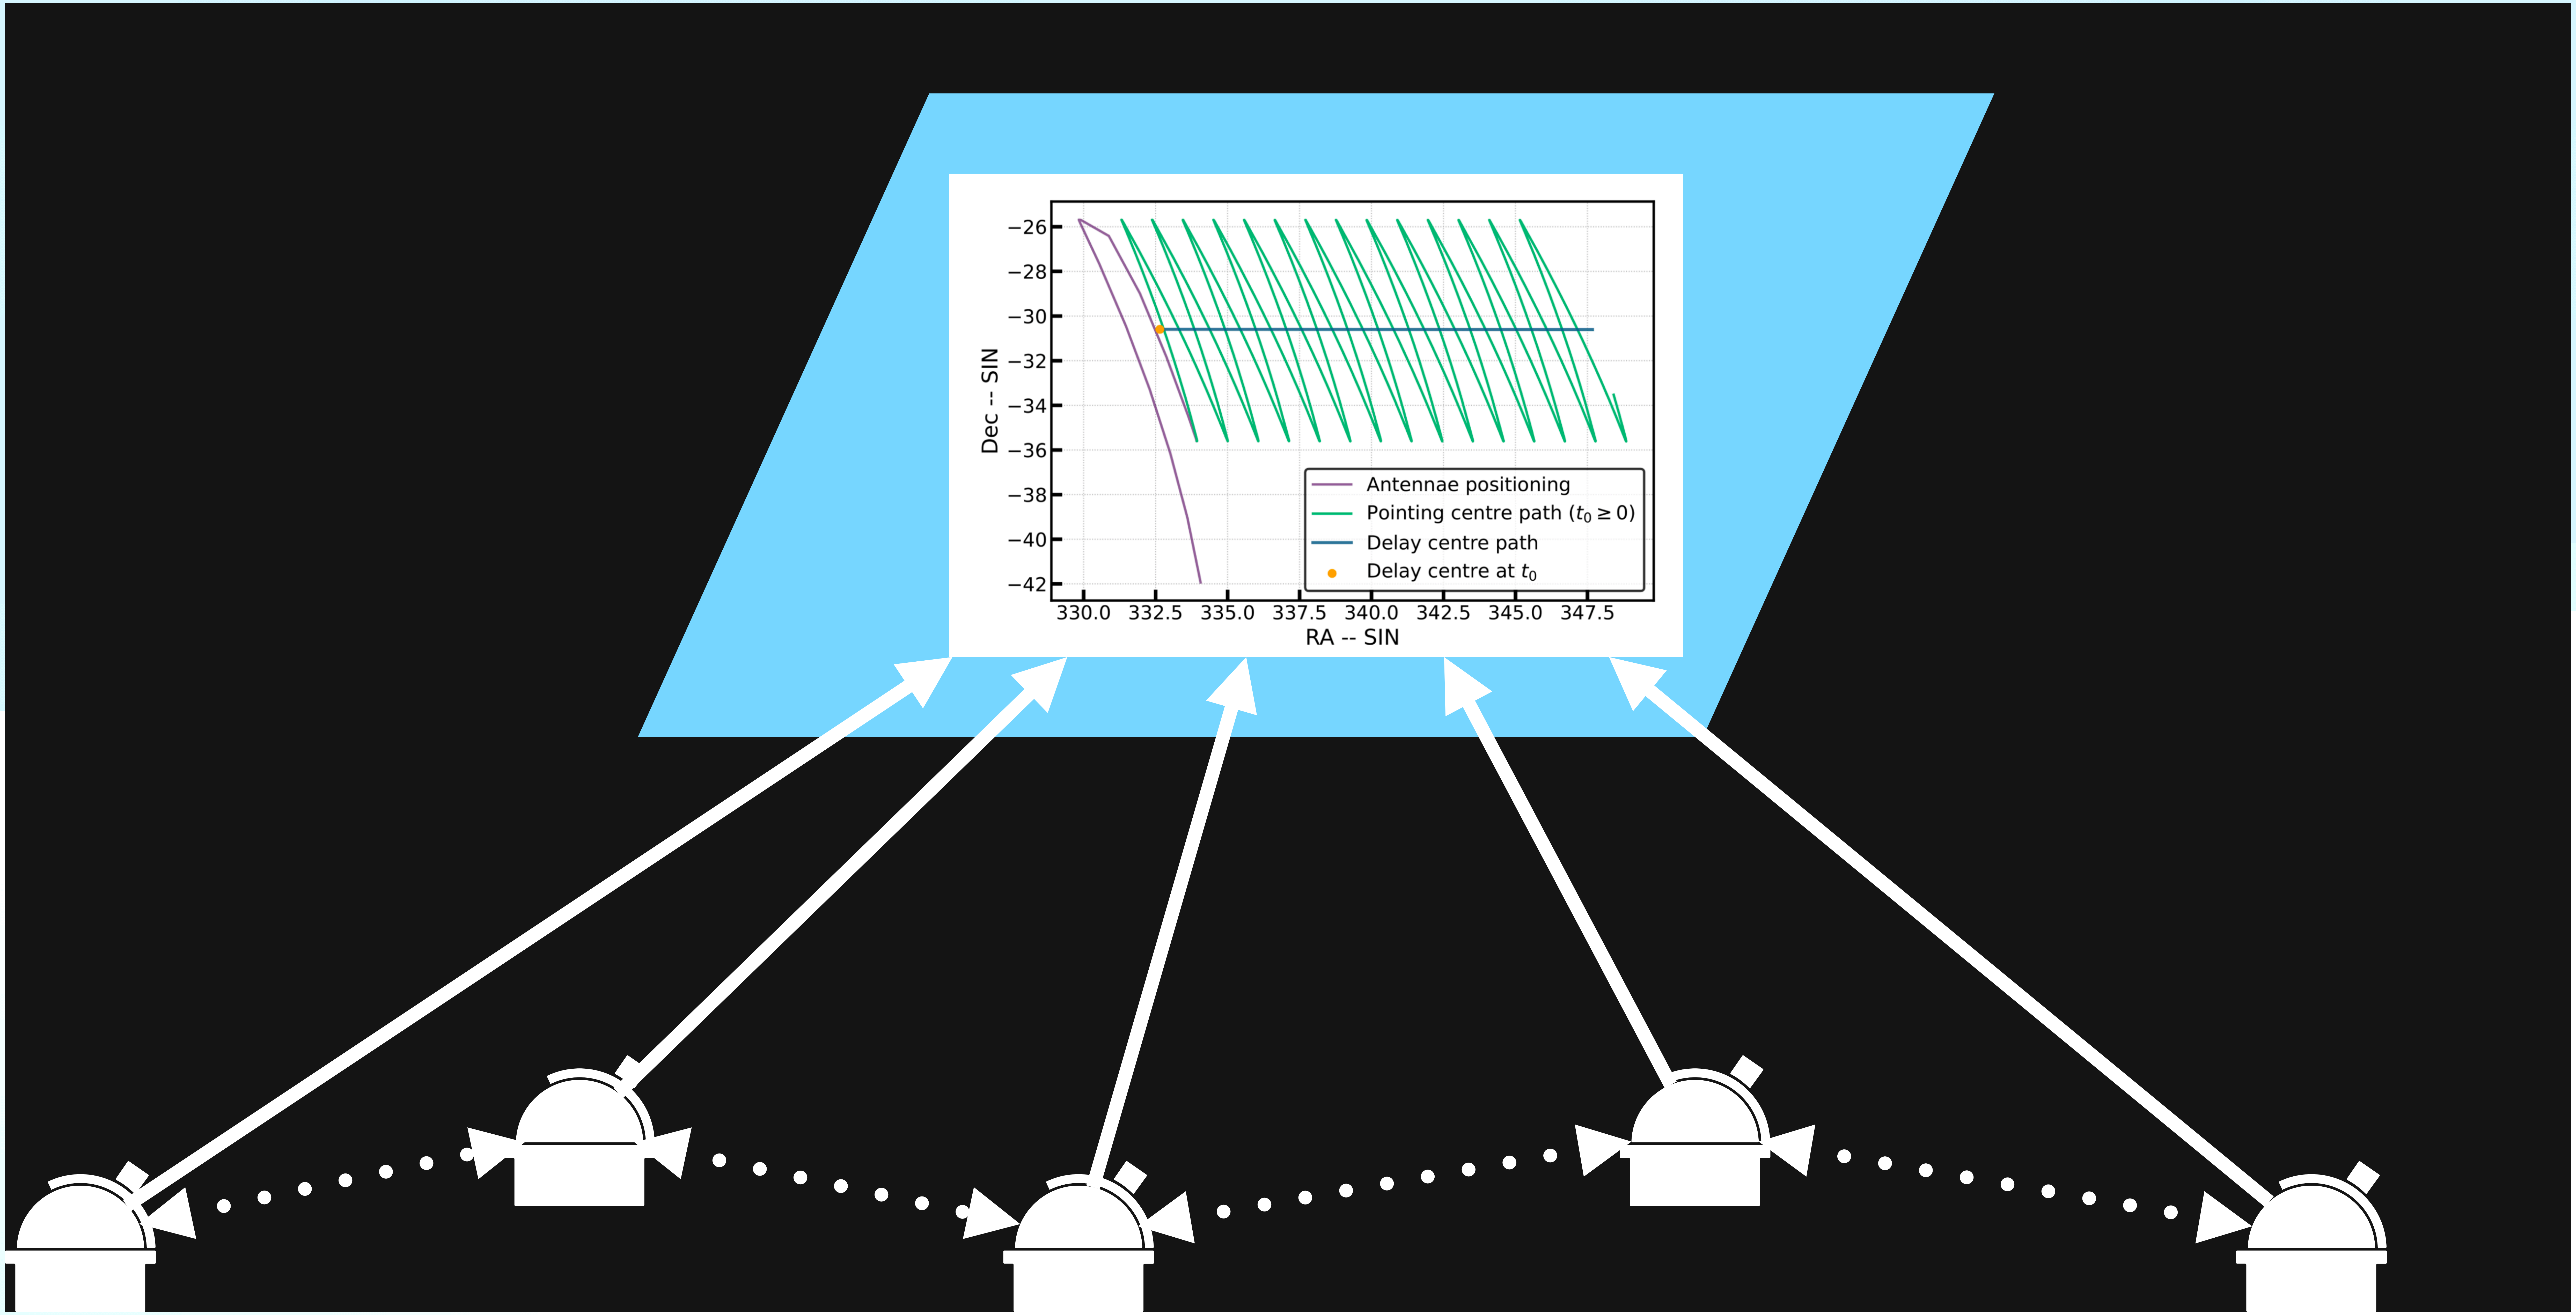

−26
−28
−30
−32
−34
−36
−38
−40
−42
Dec -- SIN
330.0
332.5
335.0
337.5
340.0
342.5
345.0
347.5
RA -- SIN
Antennae positioning
Pointing centre path ($t_0 \geq 0$)
Delay centre path
Delay centre at t_0

OTF correction for pointing centers

- In order to correlate the signal from the telescopes, they are assumed to be pointing at the same direction in the sky.
- But the telescopes scan a region of the sky: their pointing direction changes with time.
- Need a correction (rotation) of the assumed pointing center into the actual pointing center.

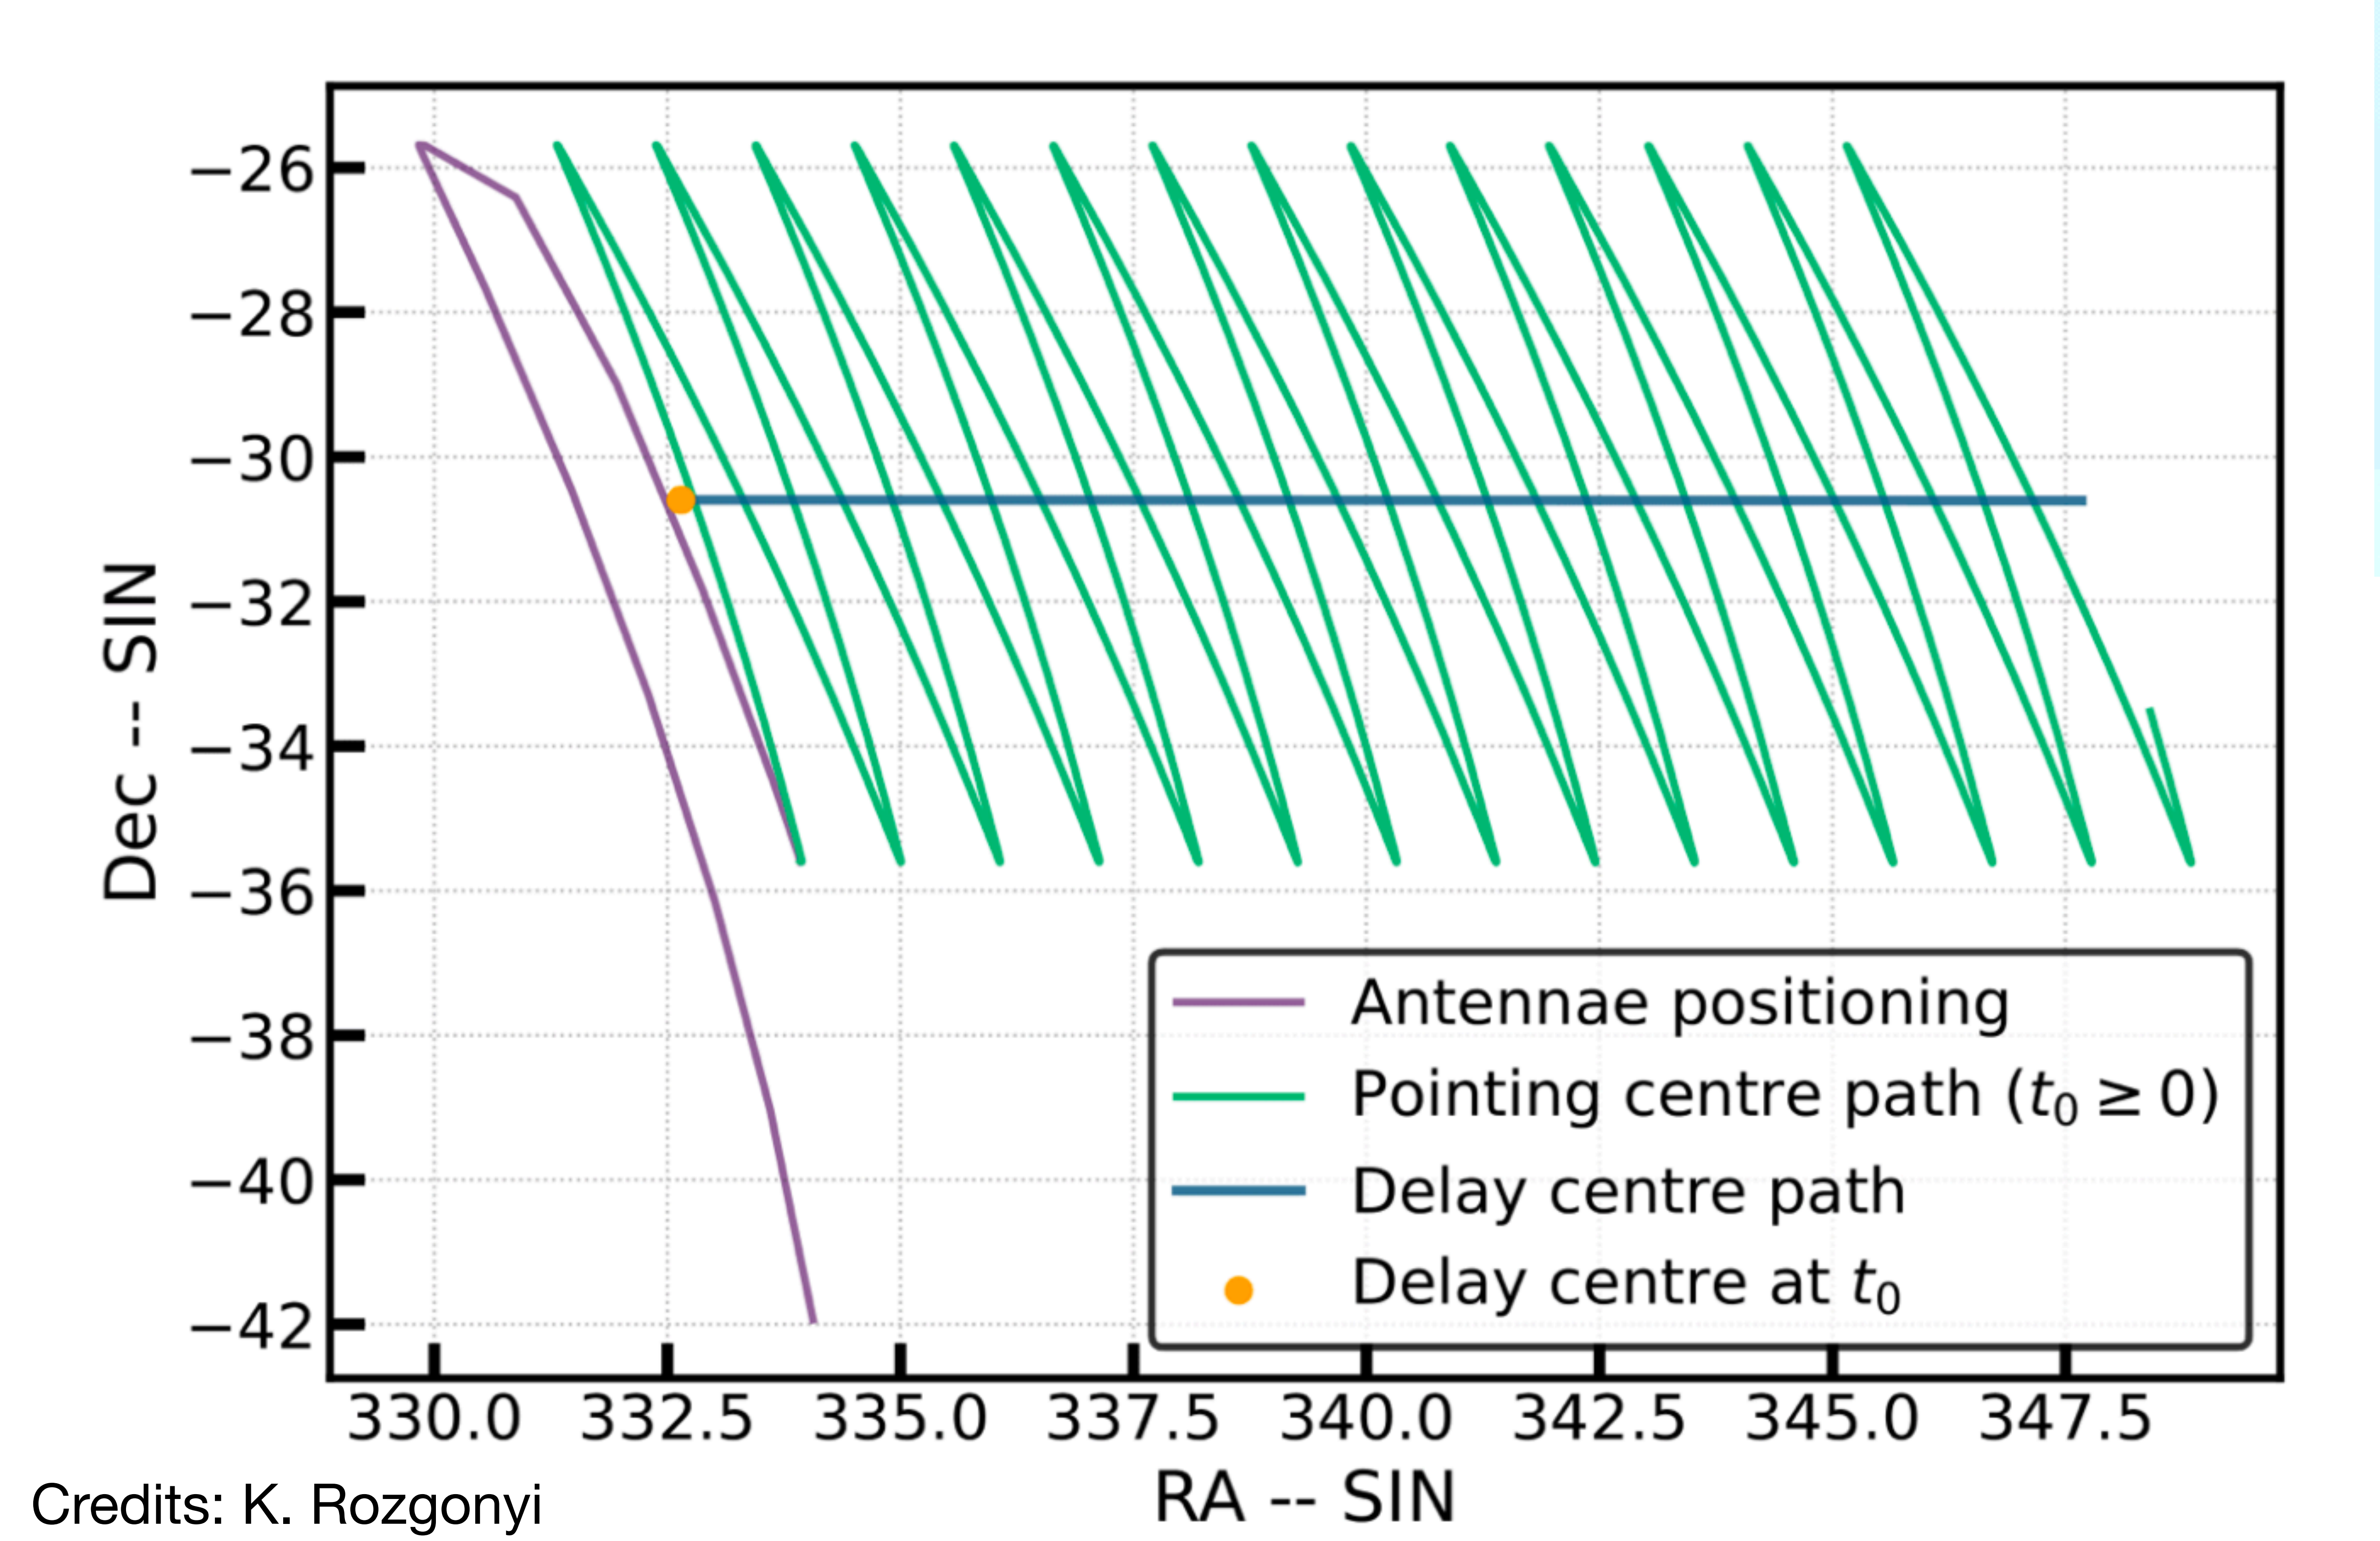

Credits: K. Rozgonyi

Novelty aspects of the pipeline

- Standard radioastronomical data reduction pipeline cannot be fully used (OTF part is not implemented, massive parallelization to deal with large data volumes is needed).
- Need imaging, co-adding, and source extraction (this is a survey).
- New ad-hoc parts need to be developed.
- Will use existing software and software appositely written.

Ad-hoc reduction pipeline

Credits: K. Rozgonyi

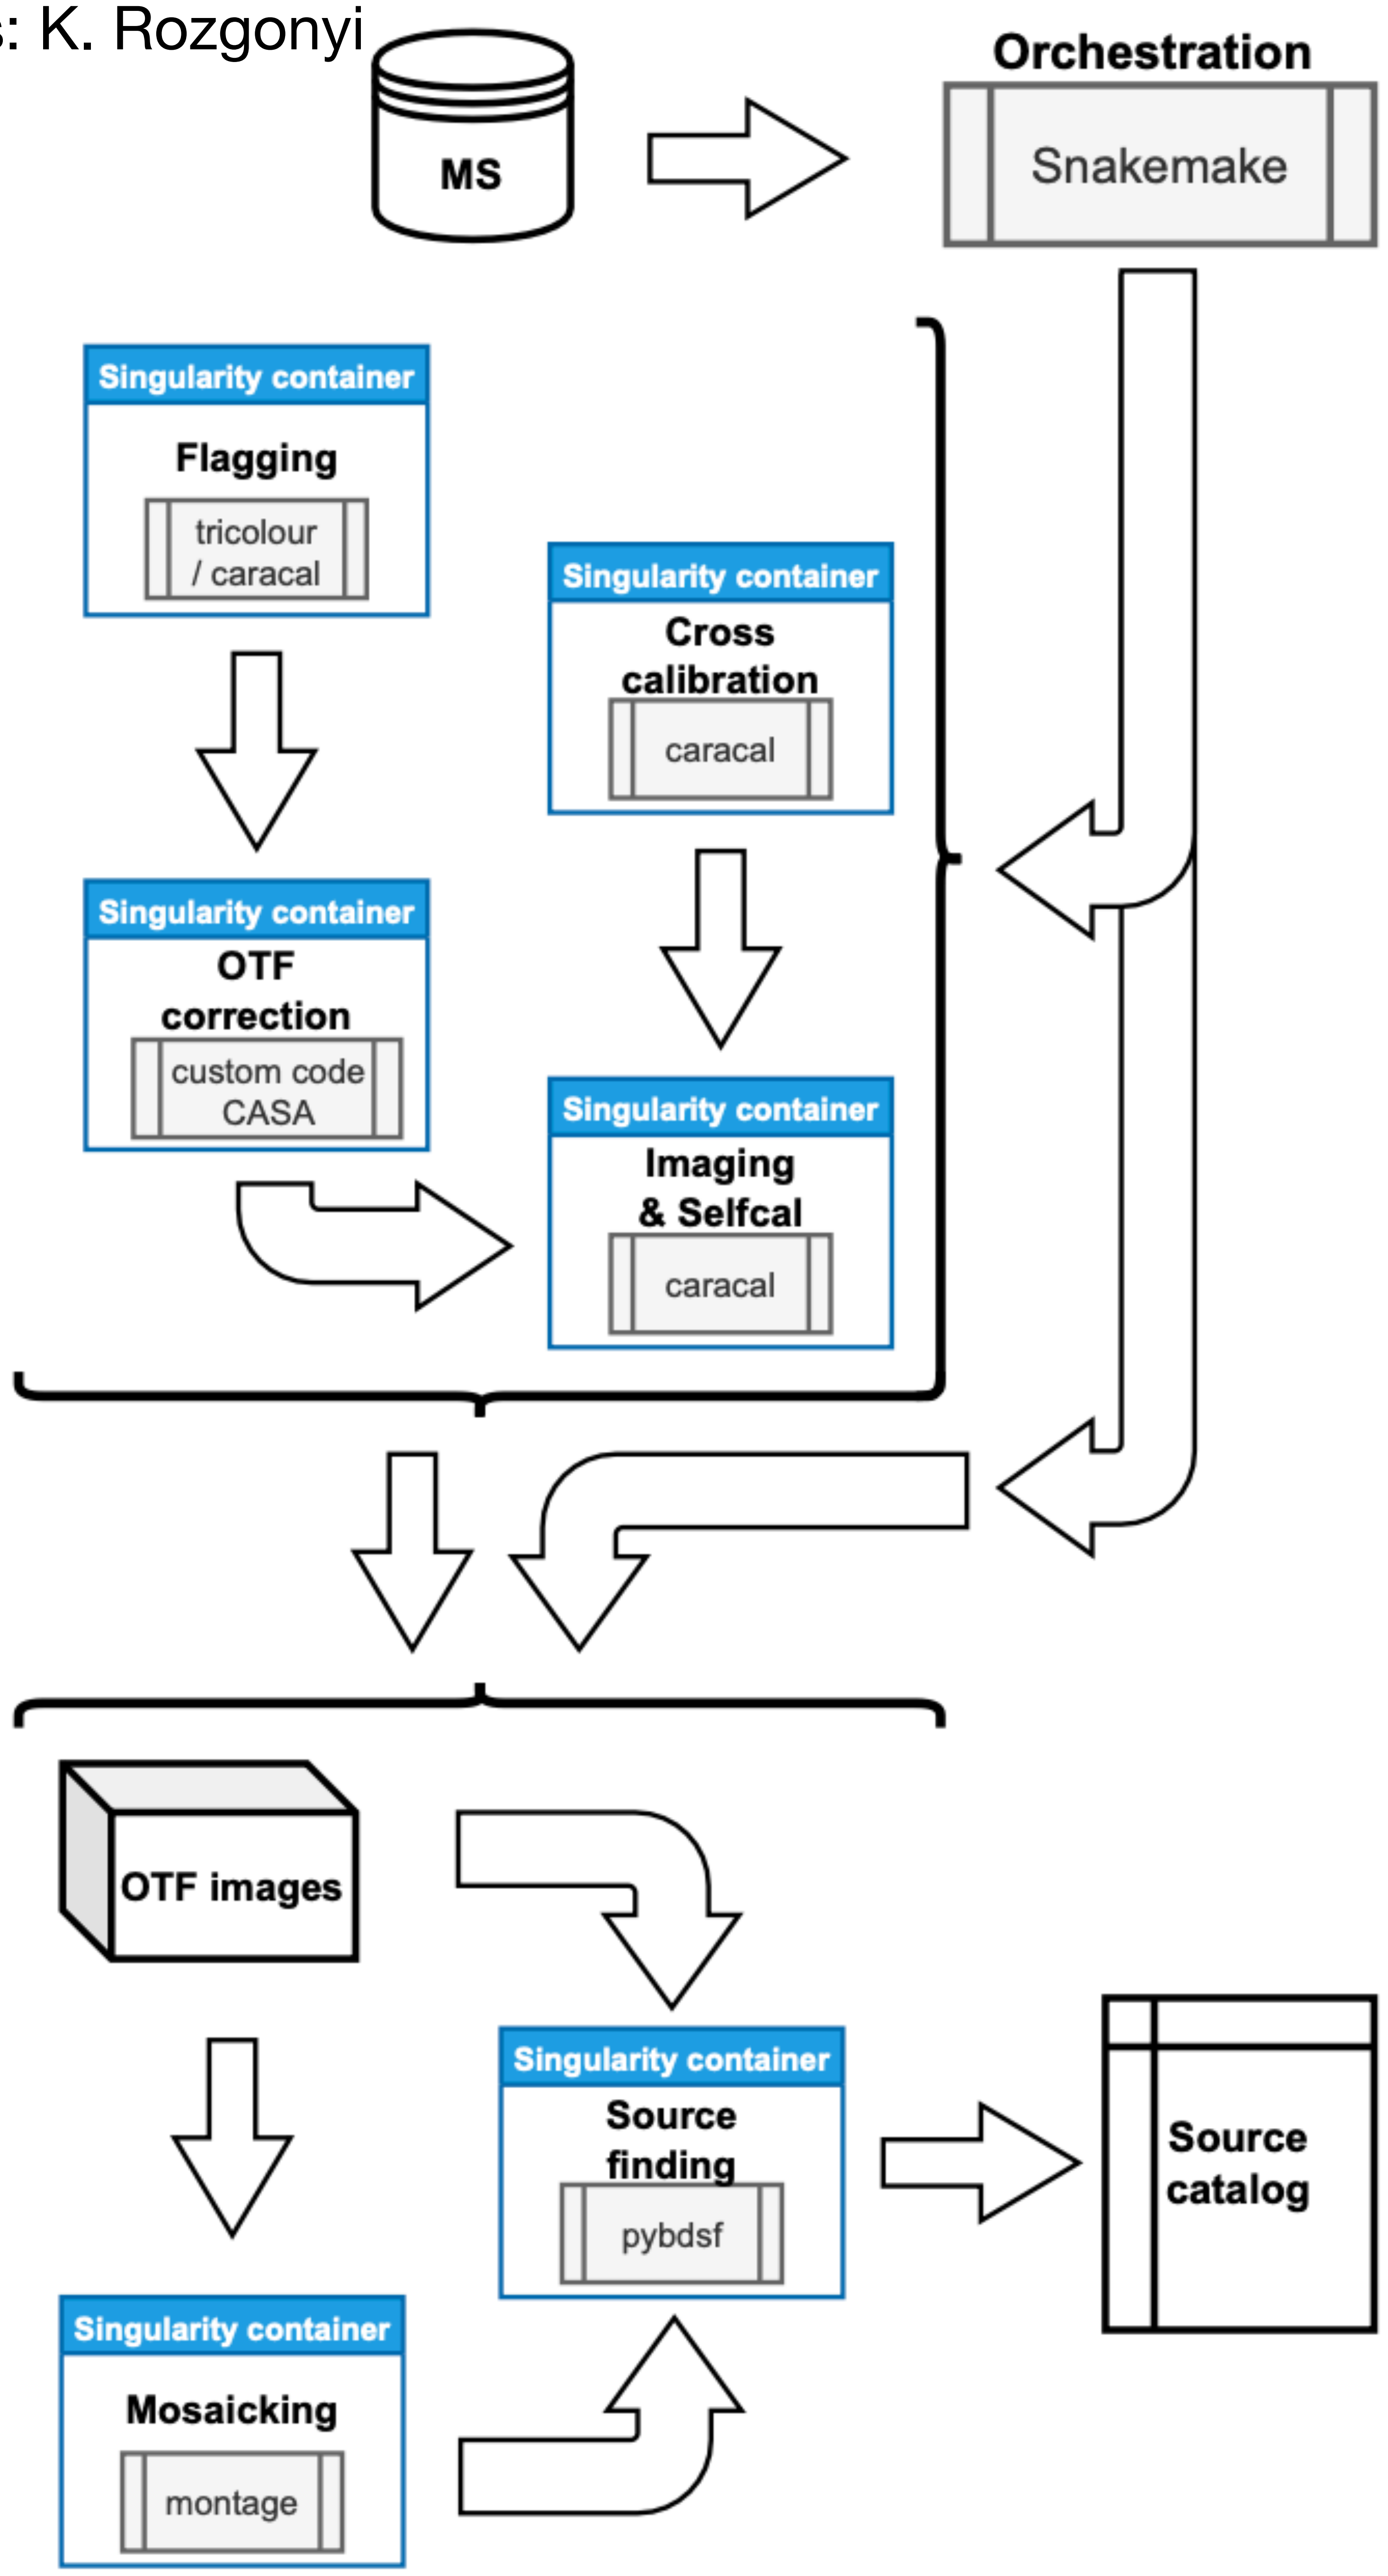

Ad-hoc reduction pipeline

- Flagging, cross- and self-calibration prepare the data (rather standard).

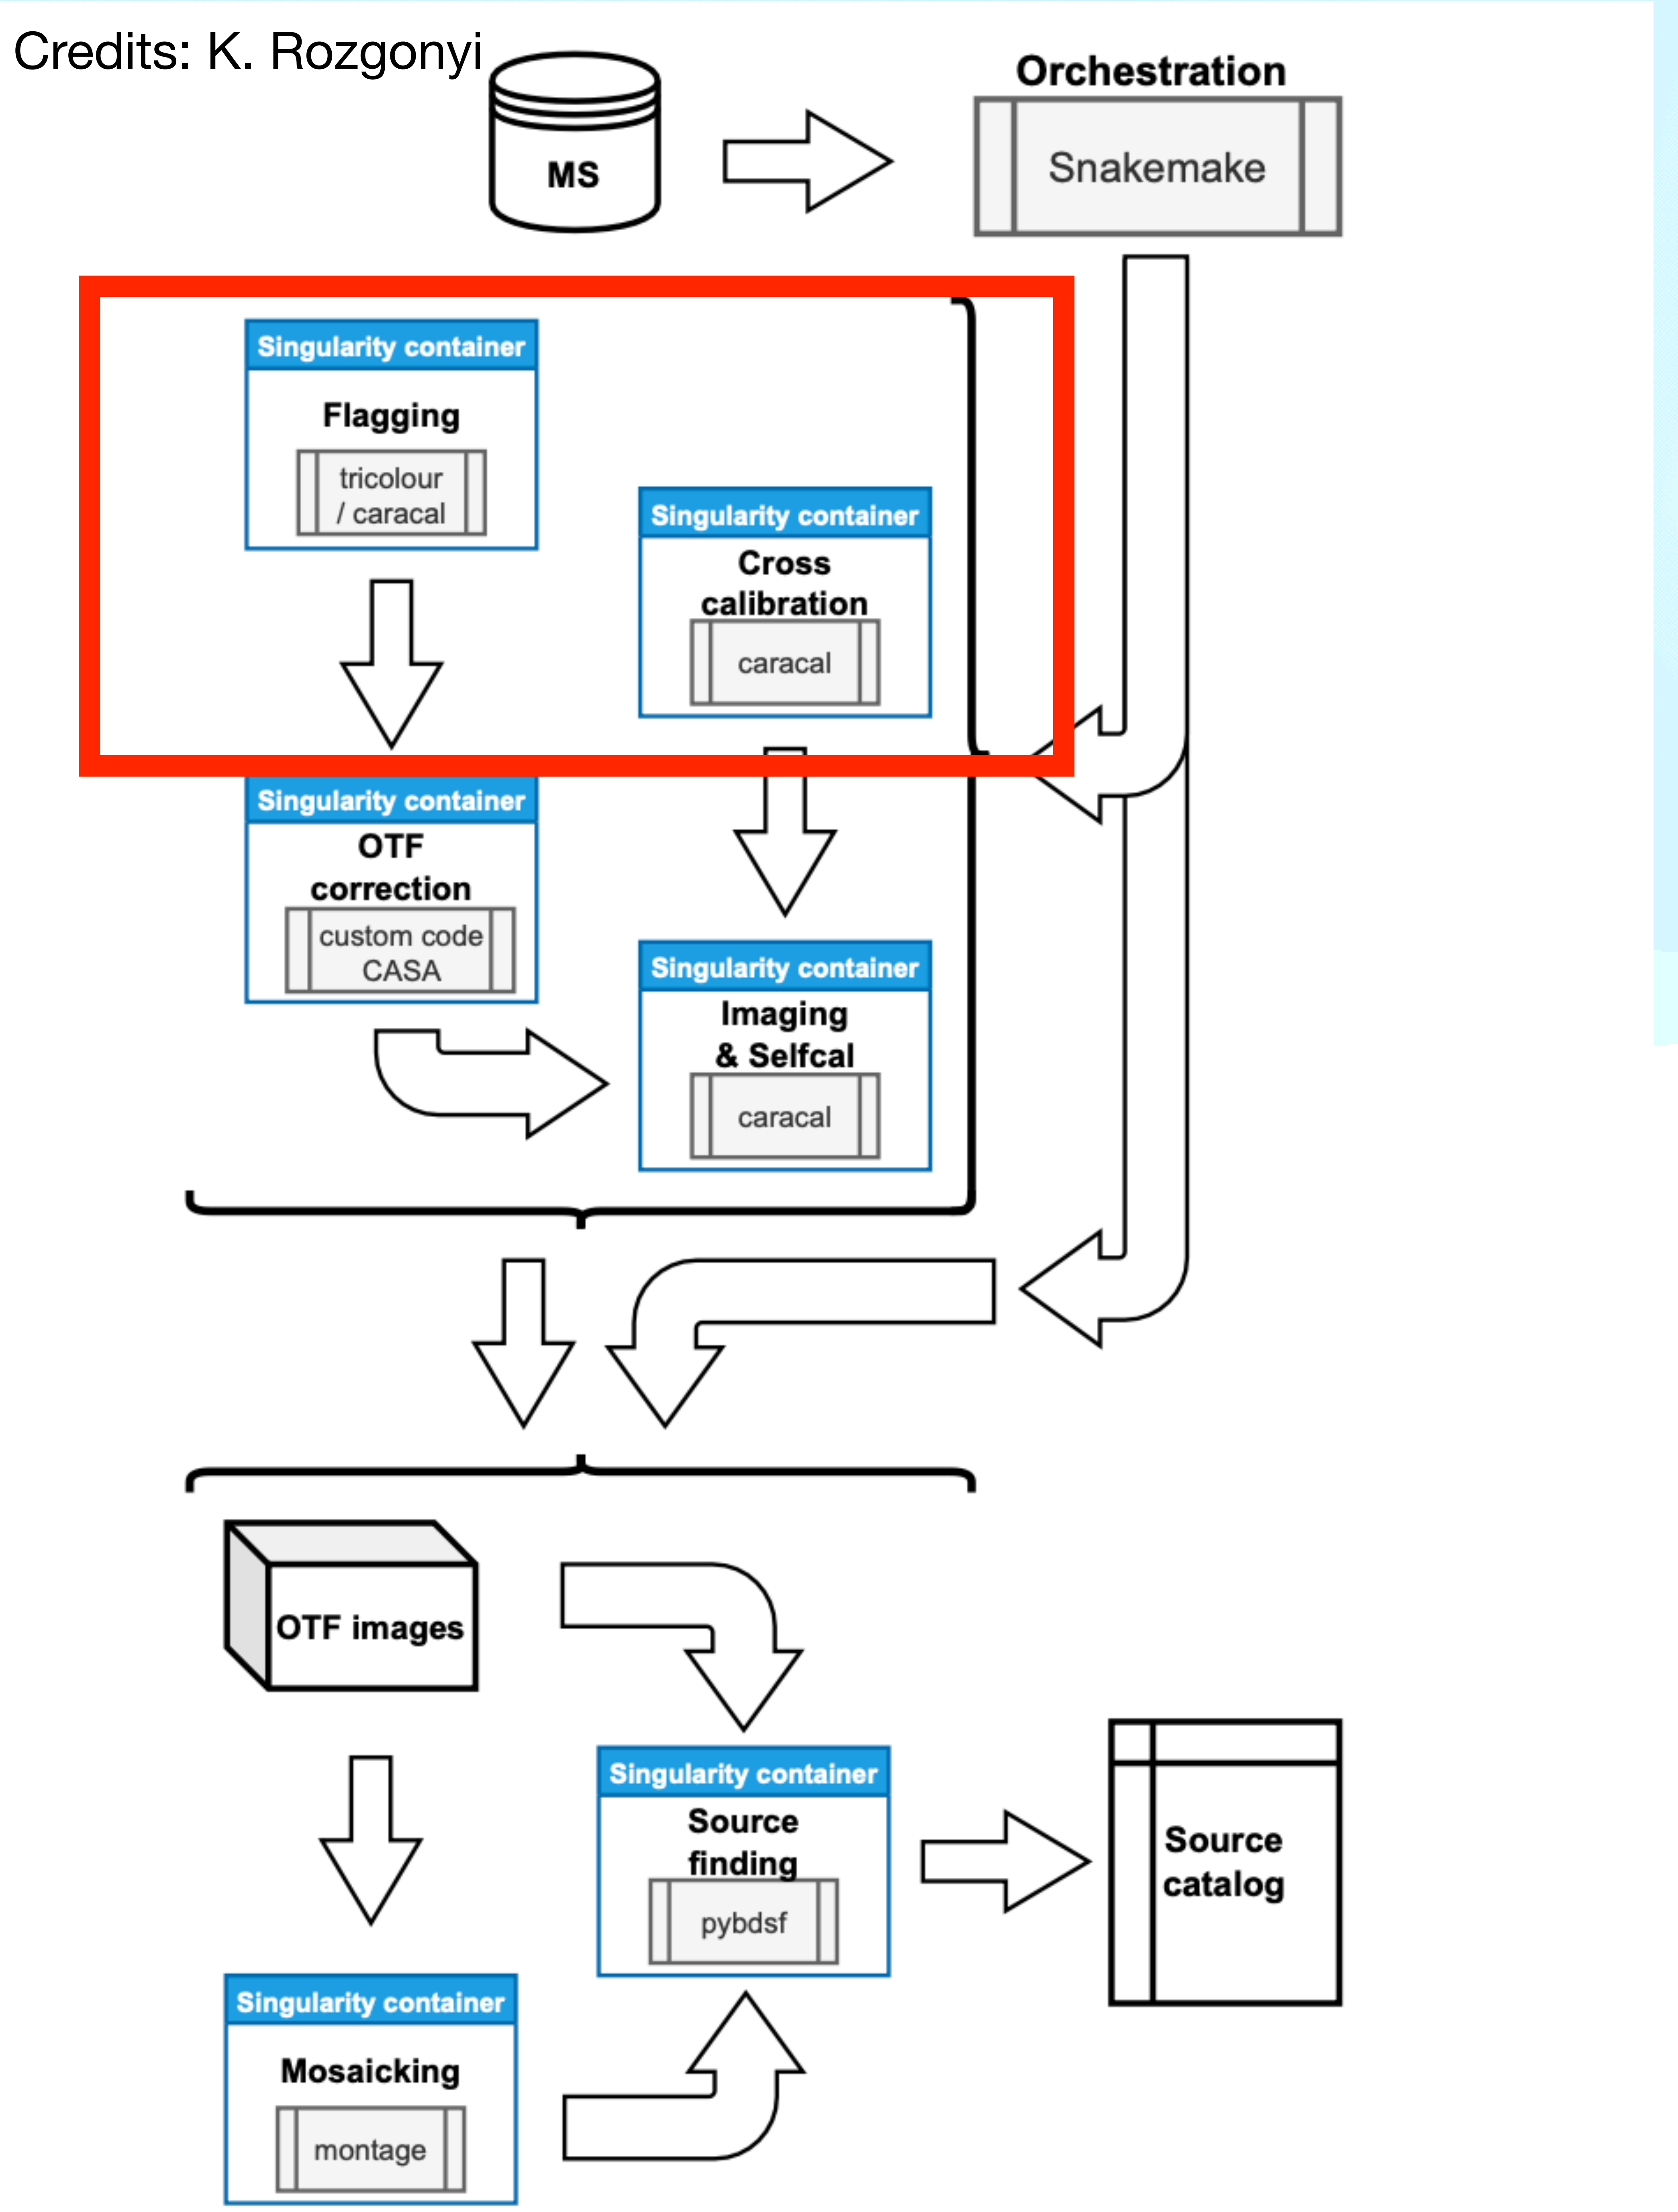

Ad-hoc reduction pipeline

- Flagging, cross- and self-calibration prepare the data (rather standard).
- OTF correction (NEW).

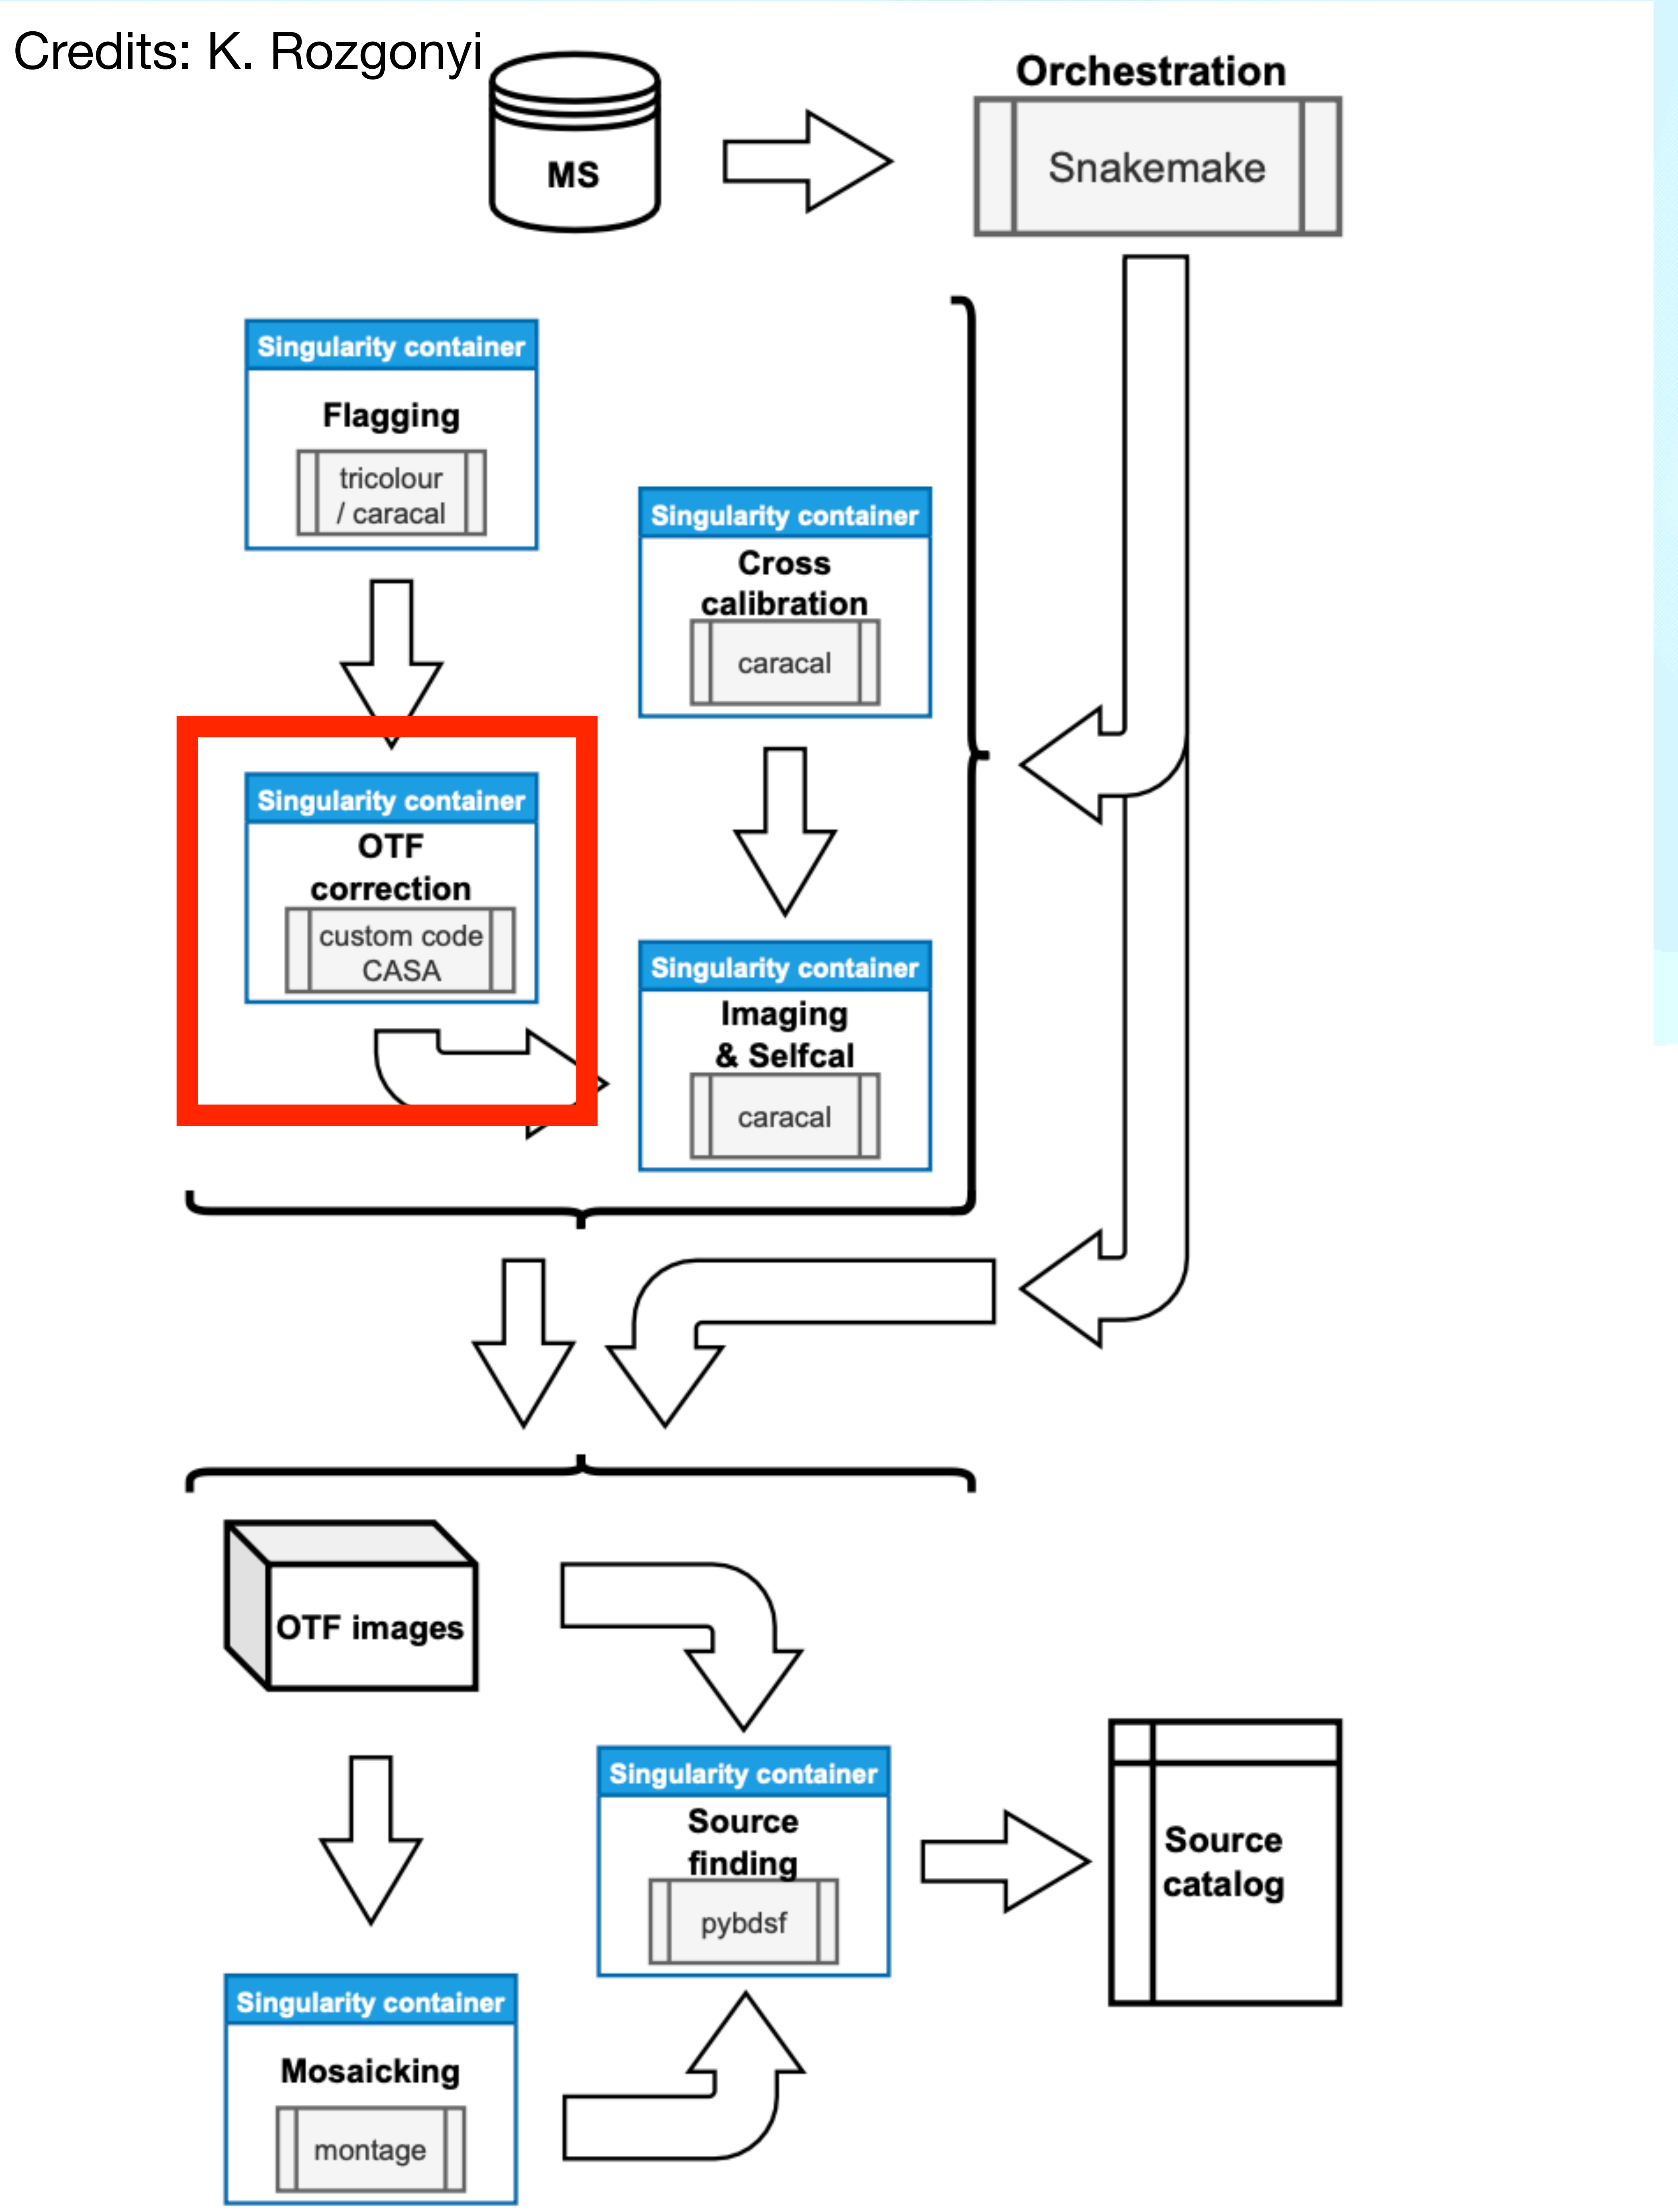

Ad-hoc reduction pipeline

- Flagging, cross- and self-calibration prepare the data (rather standard).
- OTF correction (NEW).
- Imaging (1GB input data —> ~5GB output data; 3h using ~30 cores).

Credits: K. Rozgonyi

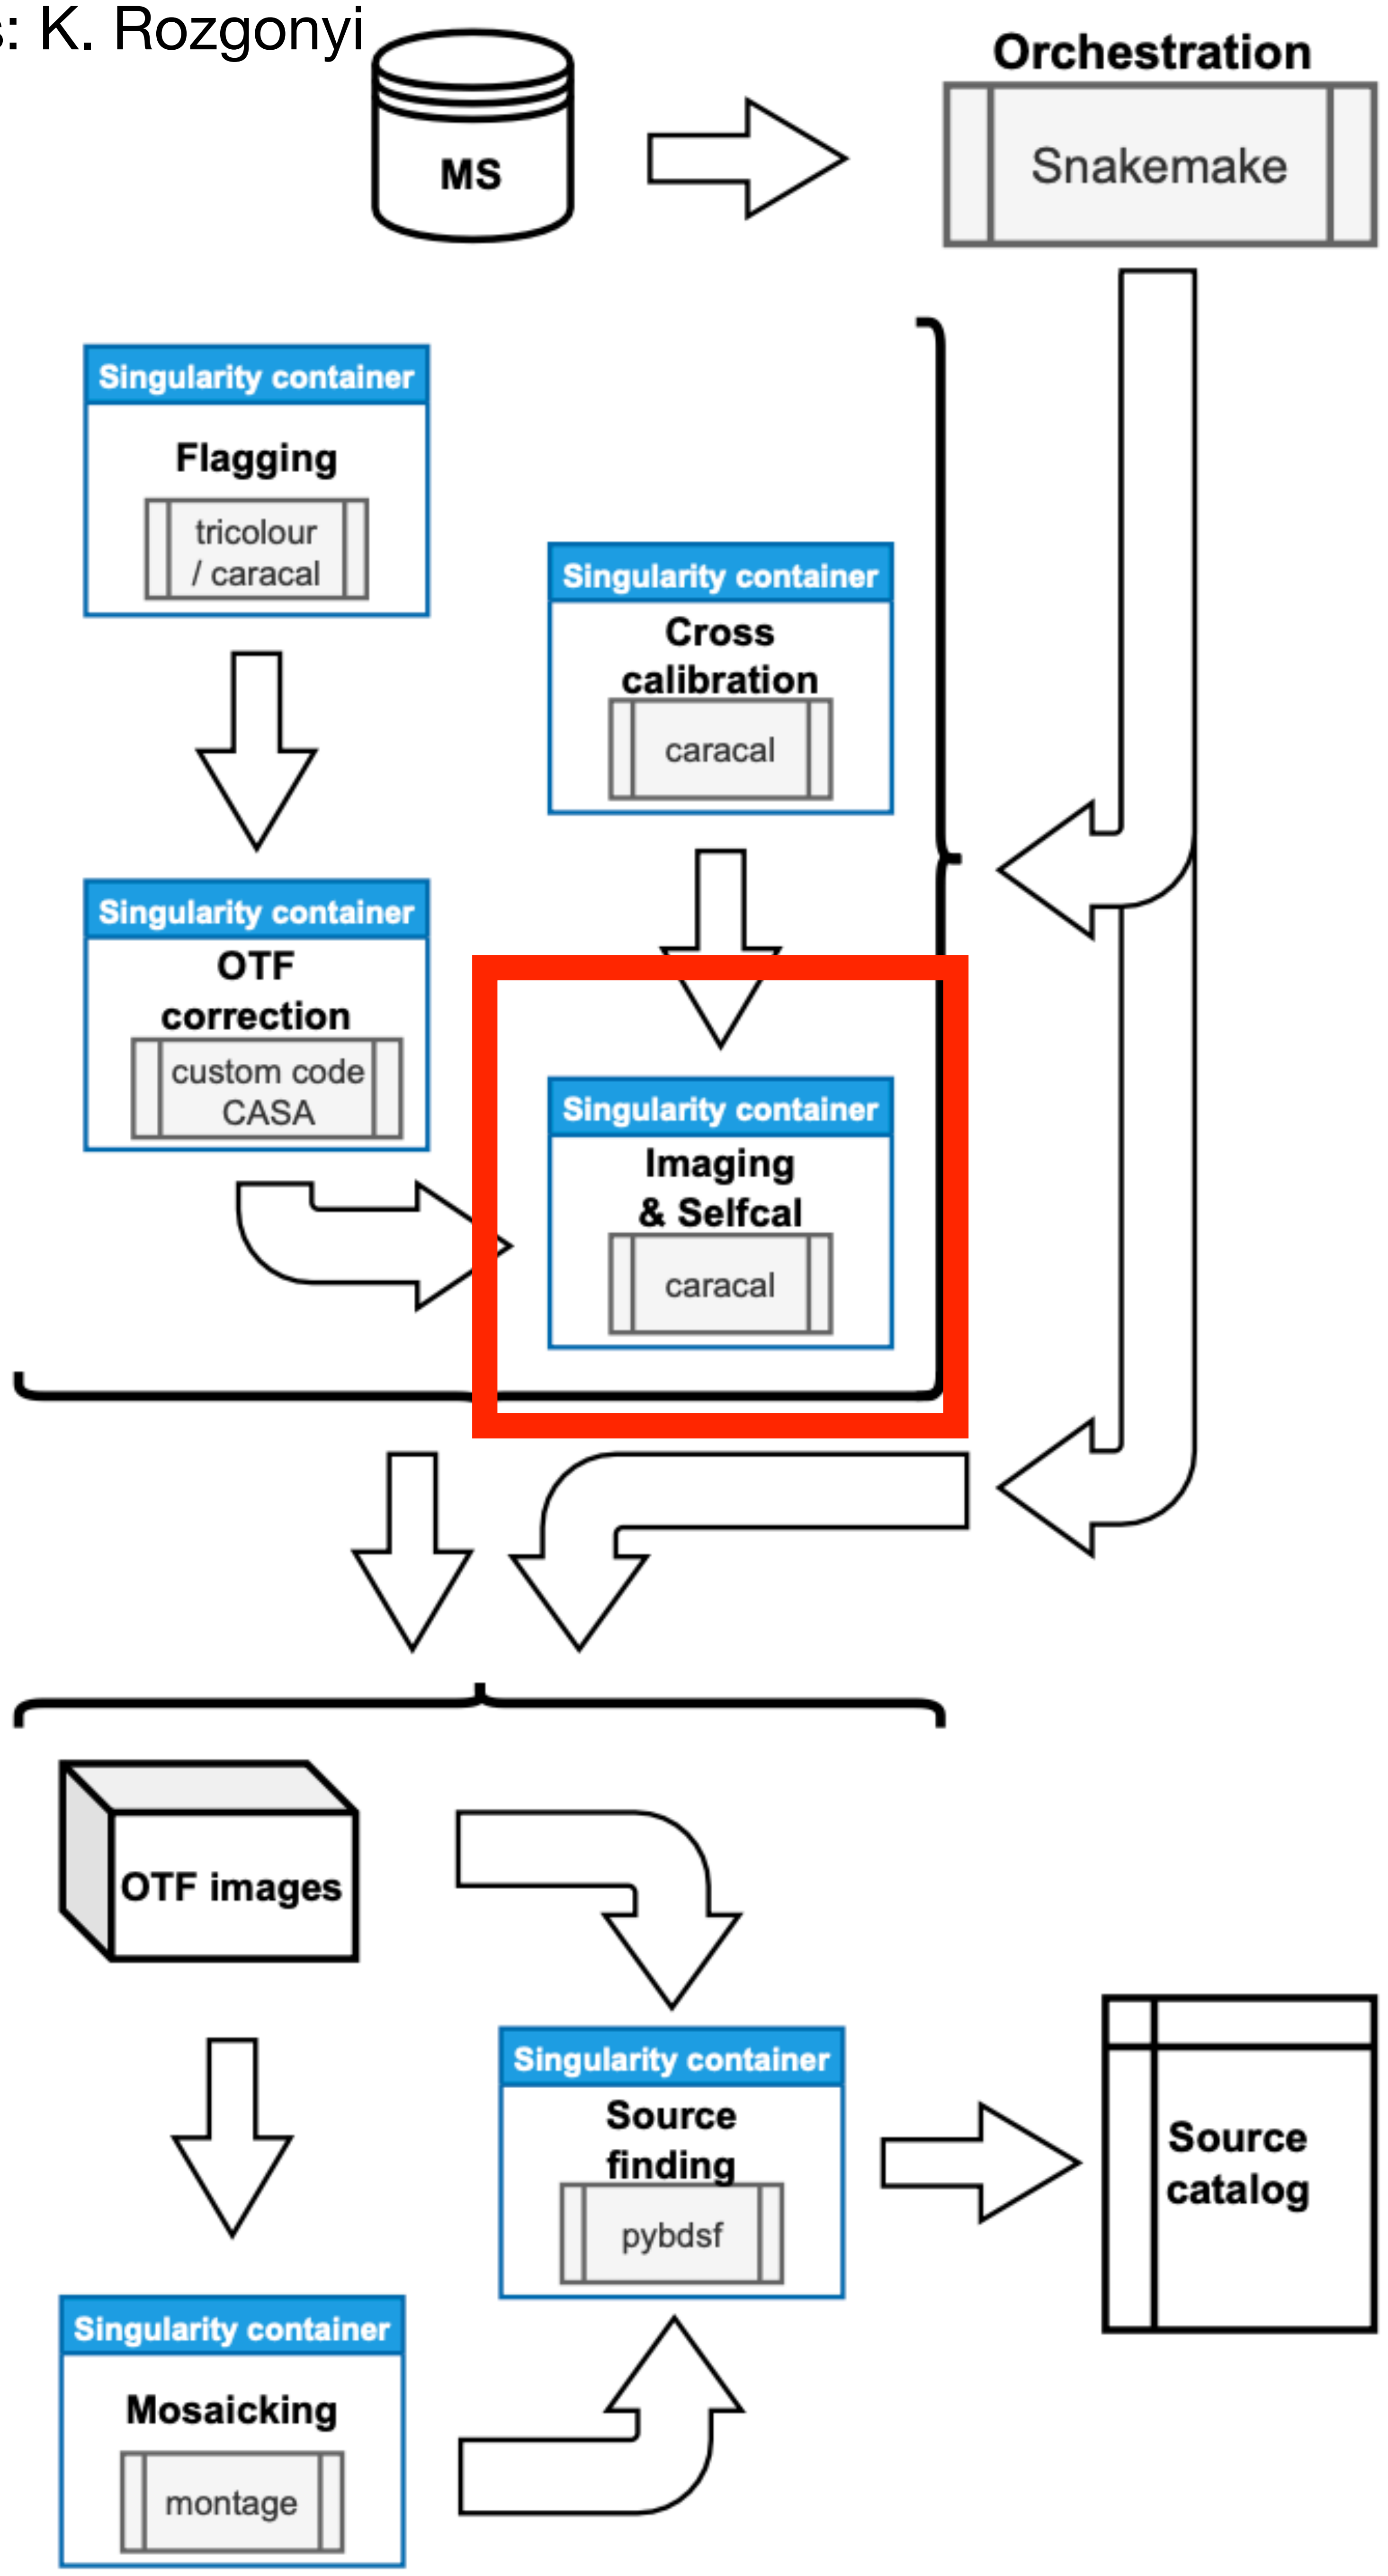

Ad-hoc reduction pipeline

- Flagging, cross- and self-calibration prepare the data (rather standard).
- OTF correction (NEW).
- Imaging (1GB input data —> ~5GB output data; 3h using ~30 cores).
- Mosaicking: finding all the tiles (1 deg region of the sky) that cover a given area of scientific interest on the sky.
- Source extraction, catalogue creation (metadata).

Credits: K. Rozgonyi

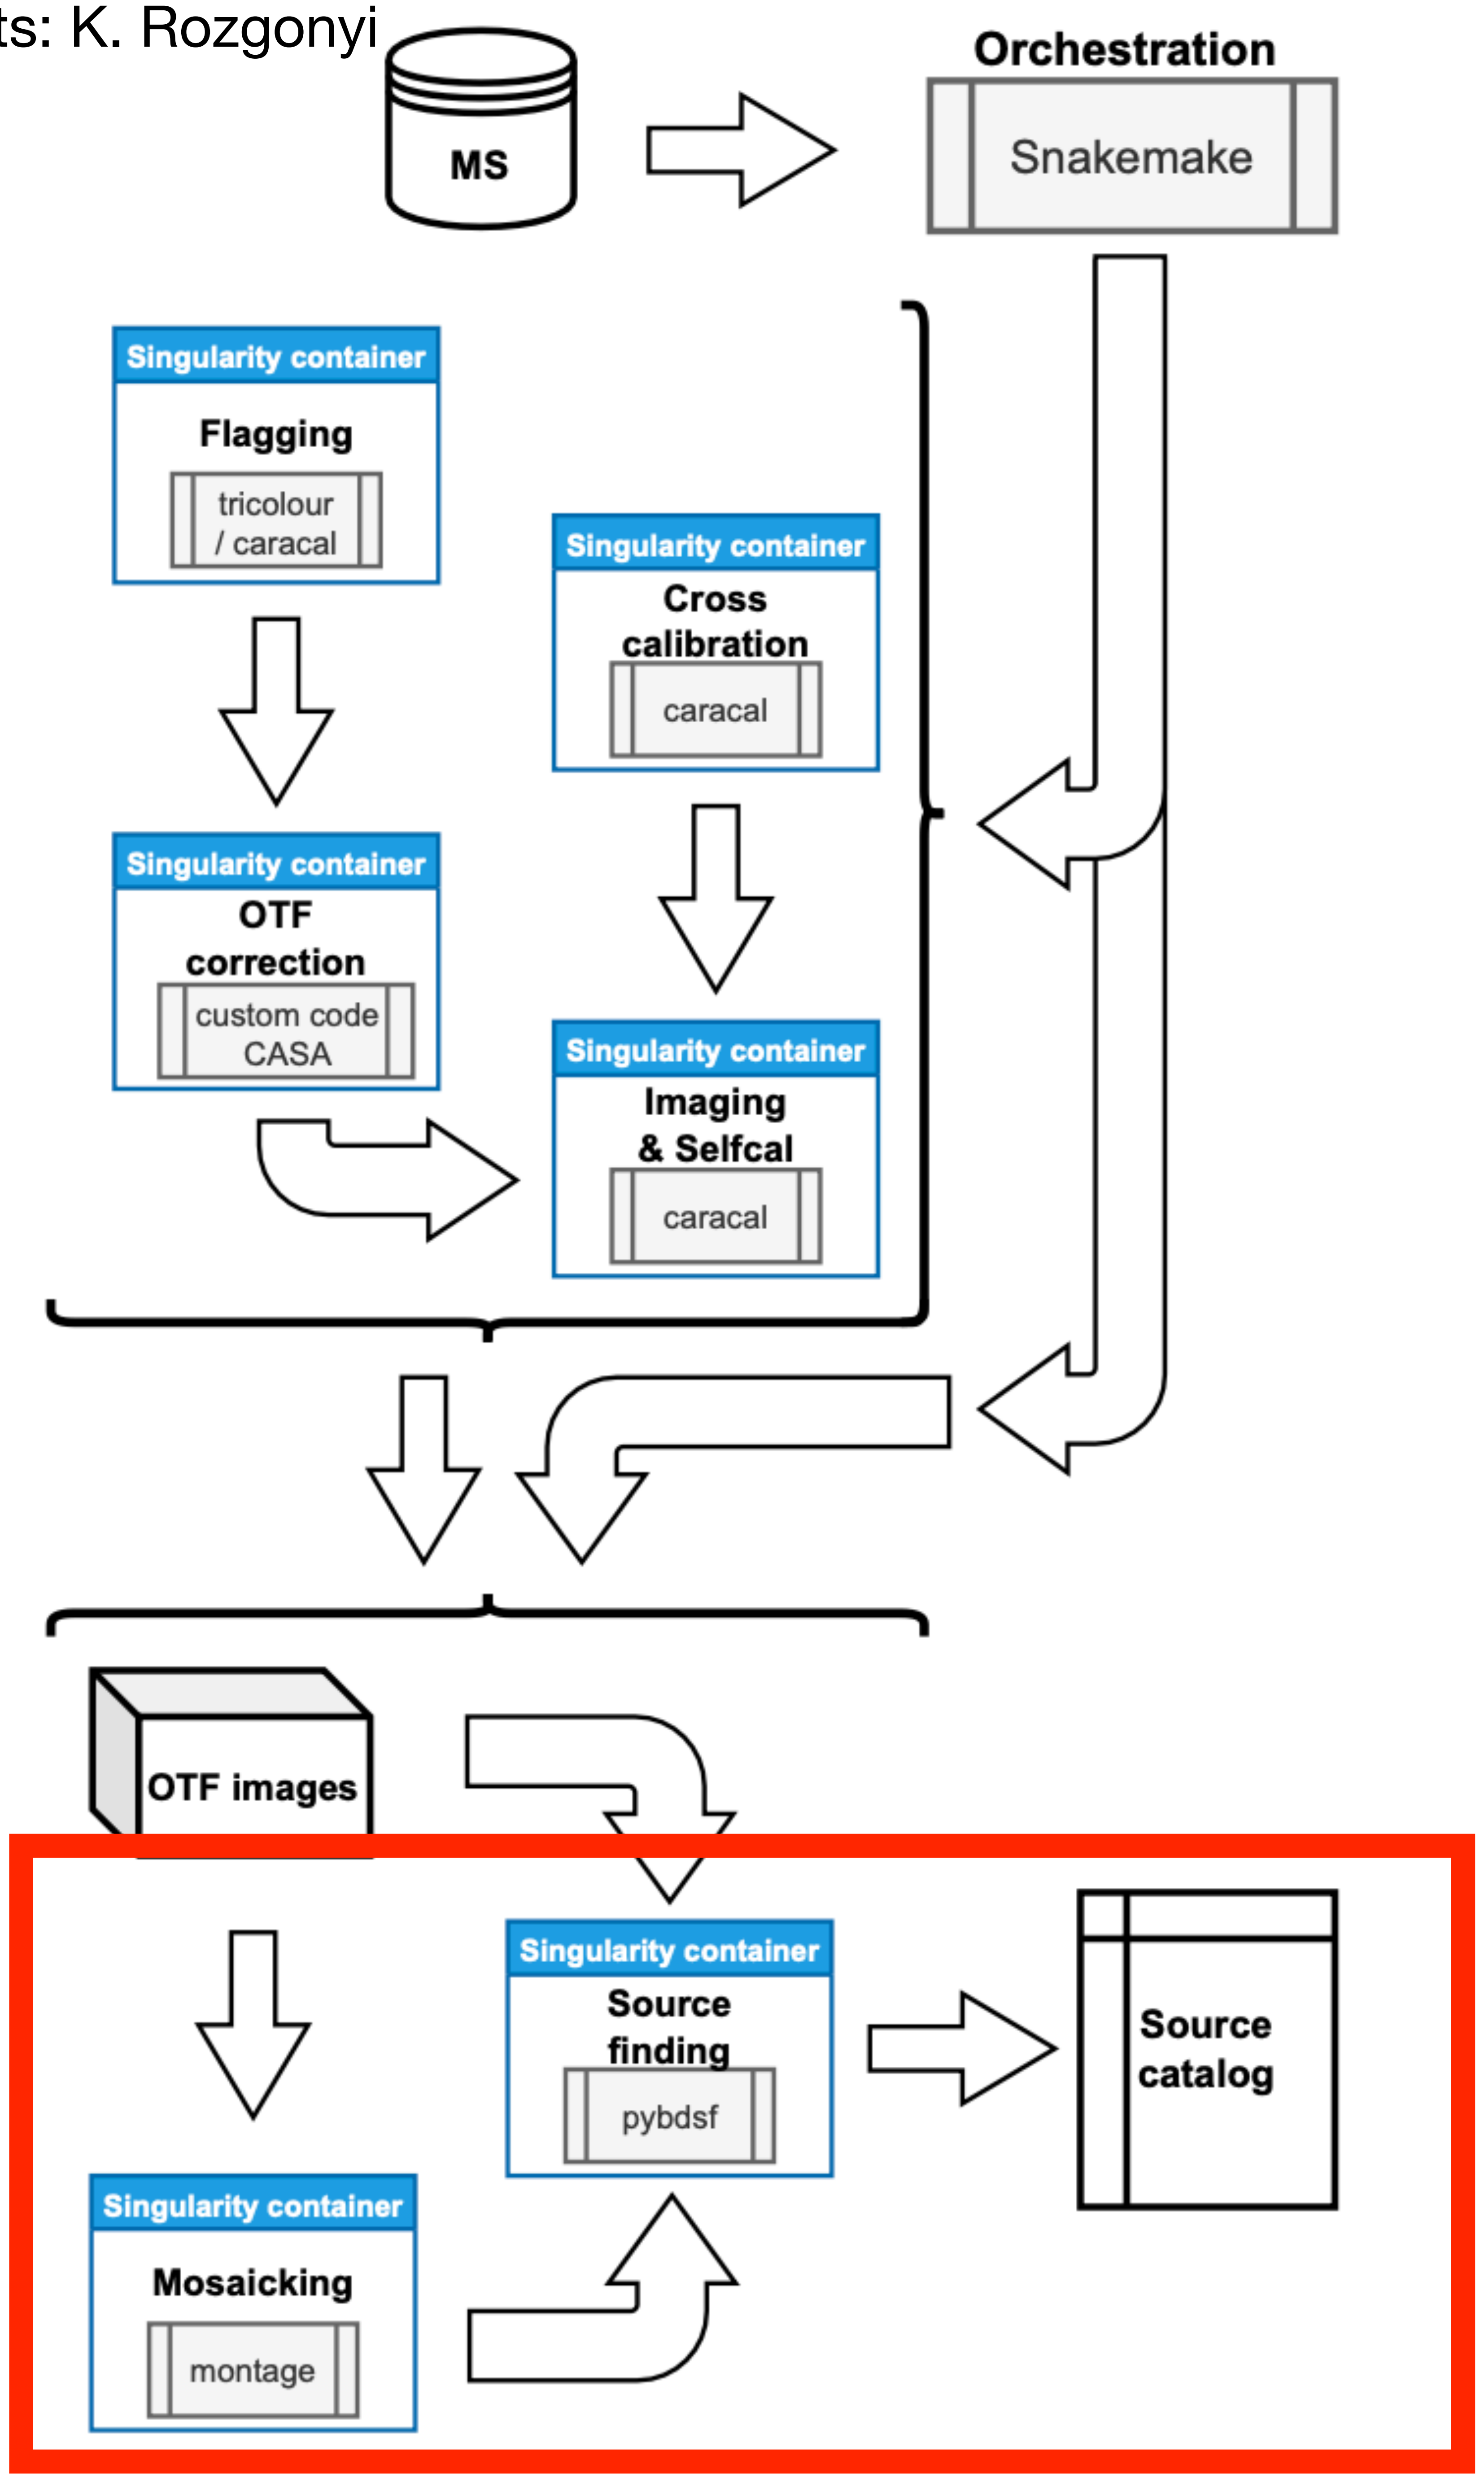

Ad-hoc reduction pipeline

- Flagging, cross- and self-calibration prepare the data (rather standard).
- OTF correction (NEW).
- Imaging (1GB input data —> ~5GB output data; 3h using ~30 cores).
- Mosaicking: finding all the tiles (1 deg region of the sky) that cover a given area of scientific interest on the sky.
- Source extraction, catalogue creation (metadata).
- Everything managed by orchestration/workflow manager.

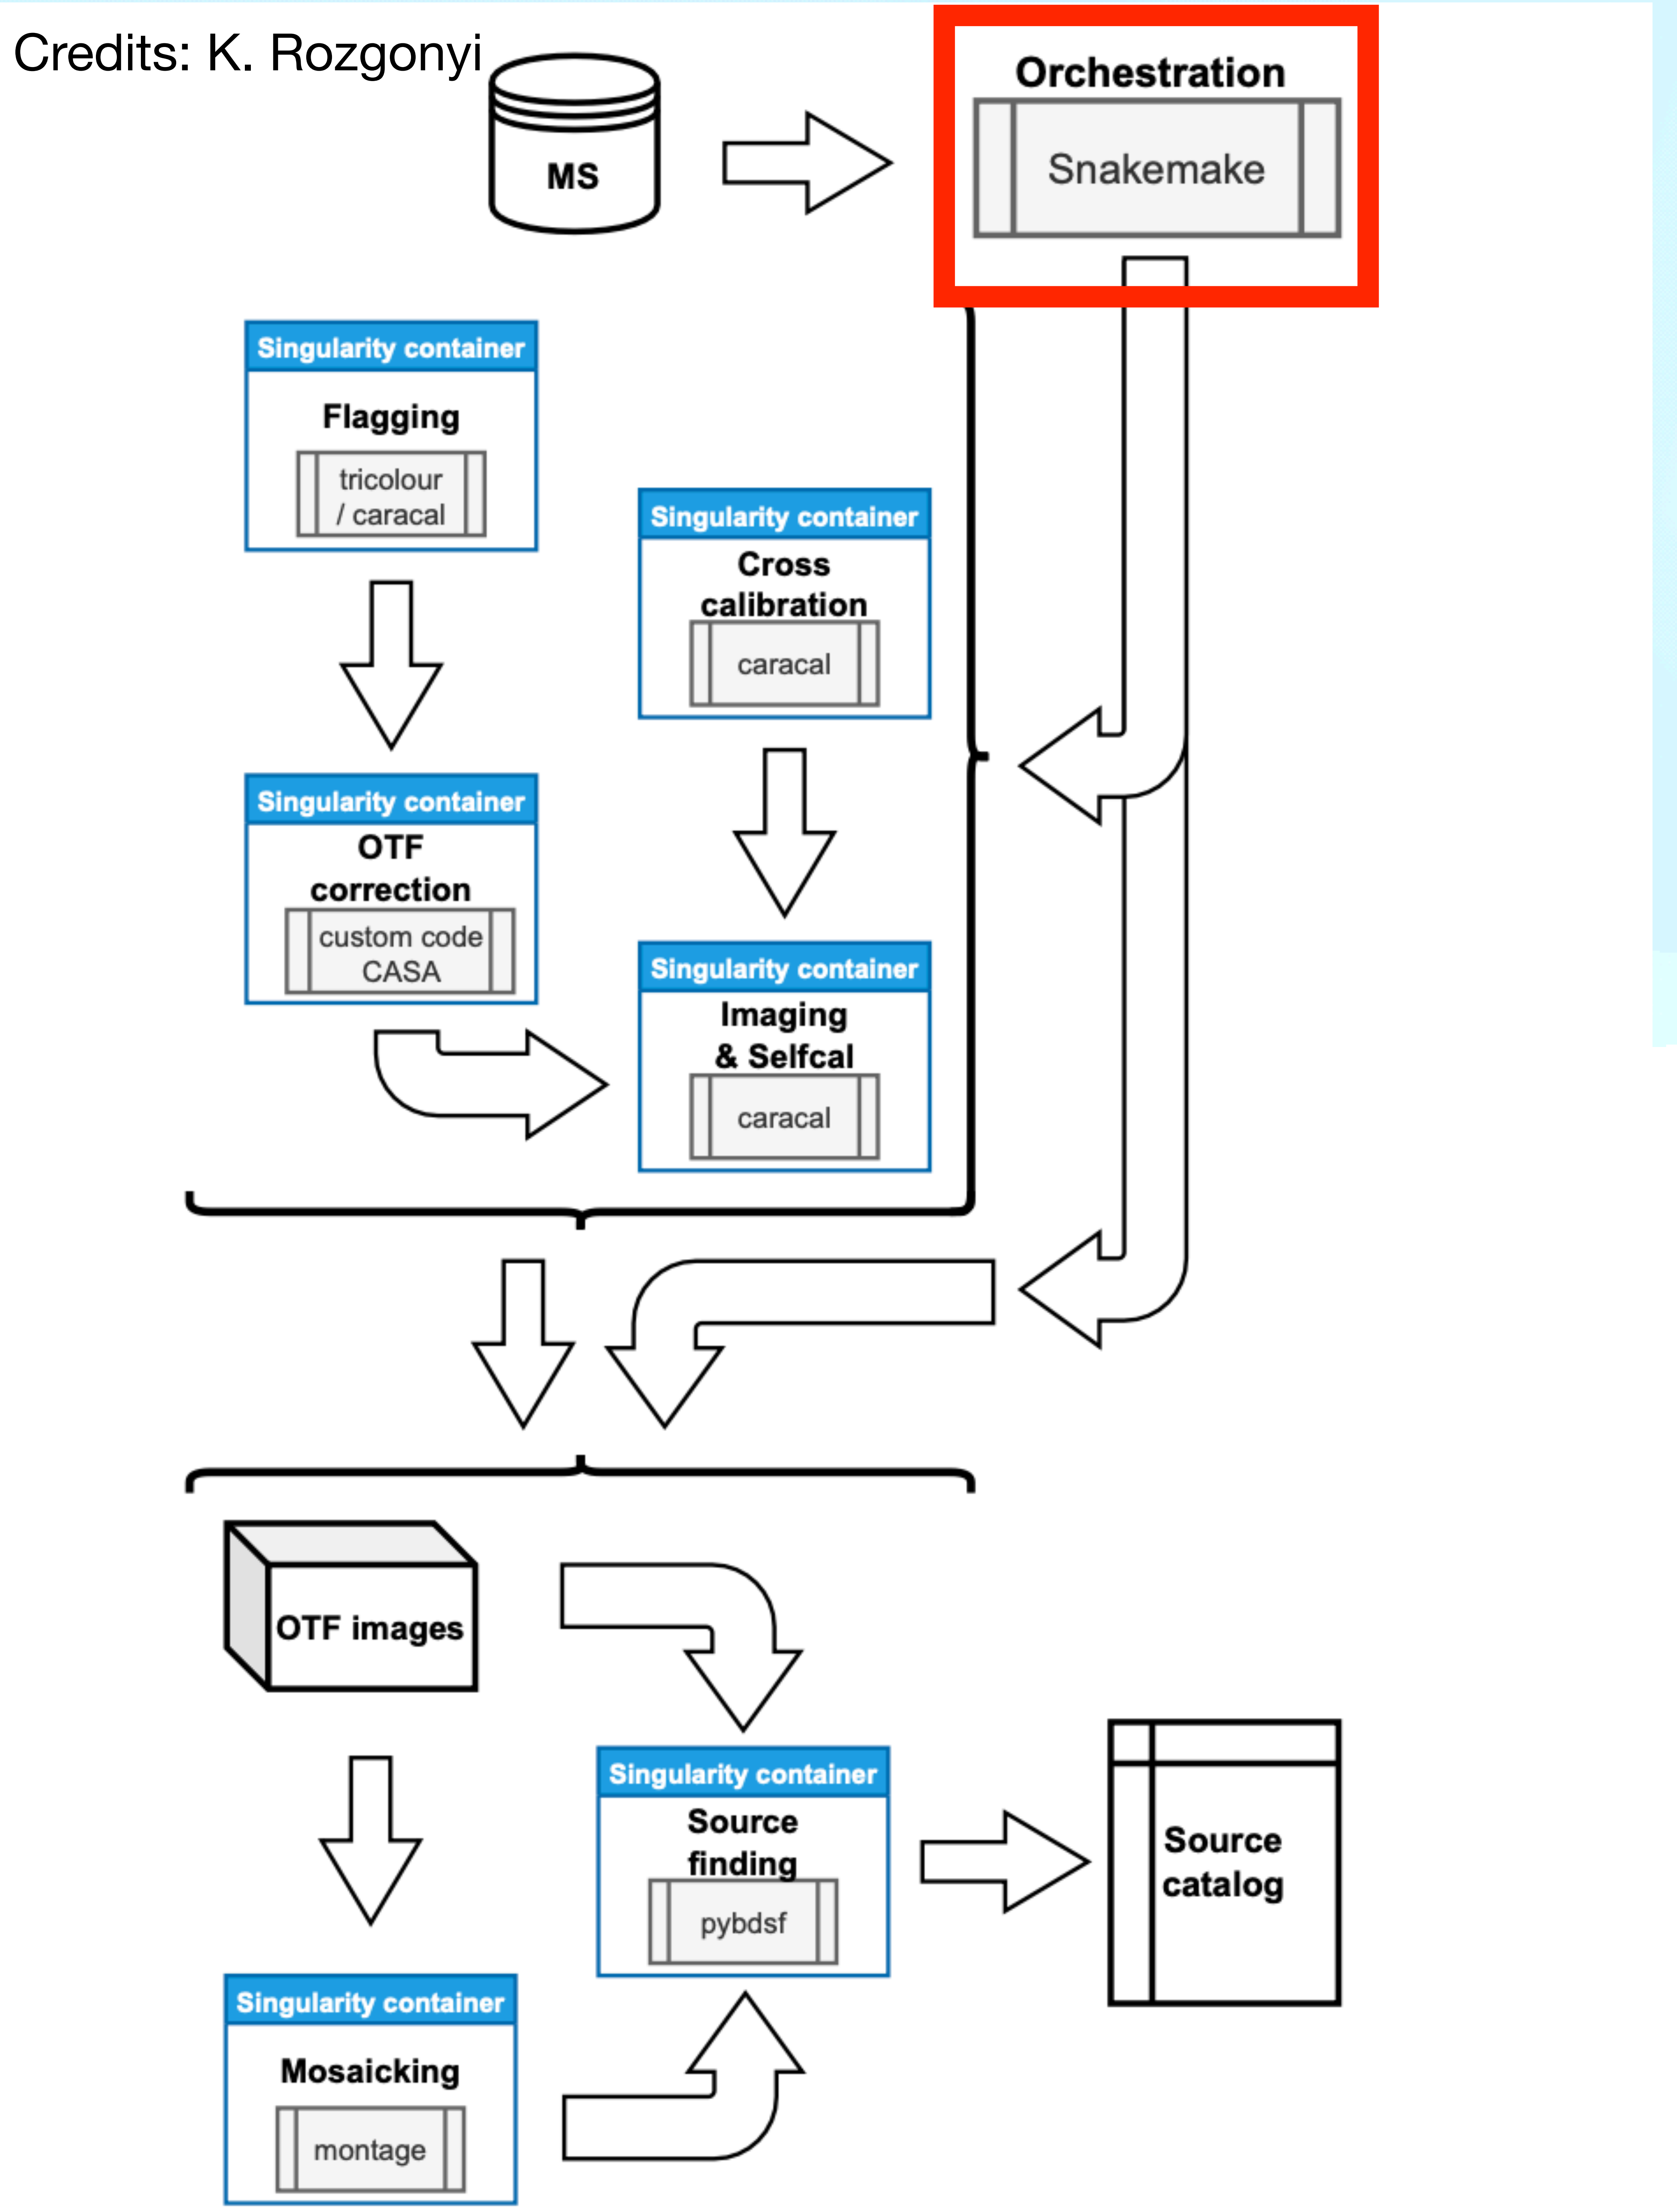

Challenges

Challenges are related to the functioning of the PUNCH4NFDI infrastructure and code requirements/complexity of the pipeline.

- Containers need to be created for the various parts for it to run on Compute4PUNCH
- Workflow managers need to interact with the computing facility
- As storage and computing are not connected yet, data movement must be managed

Achievements - OTF correction

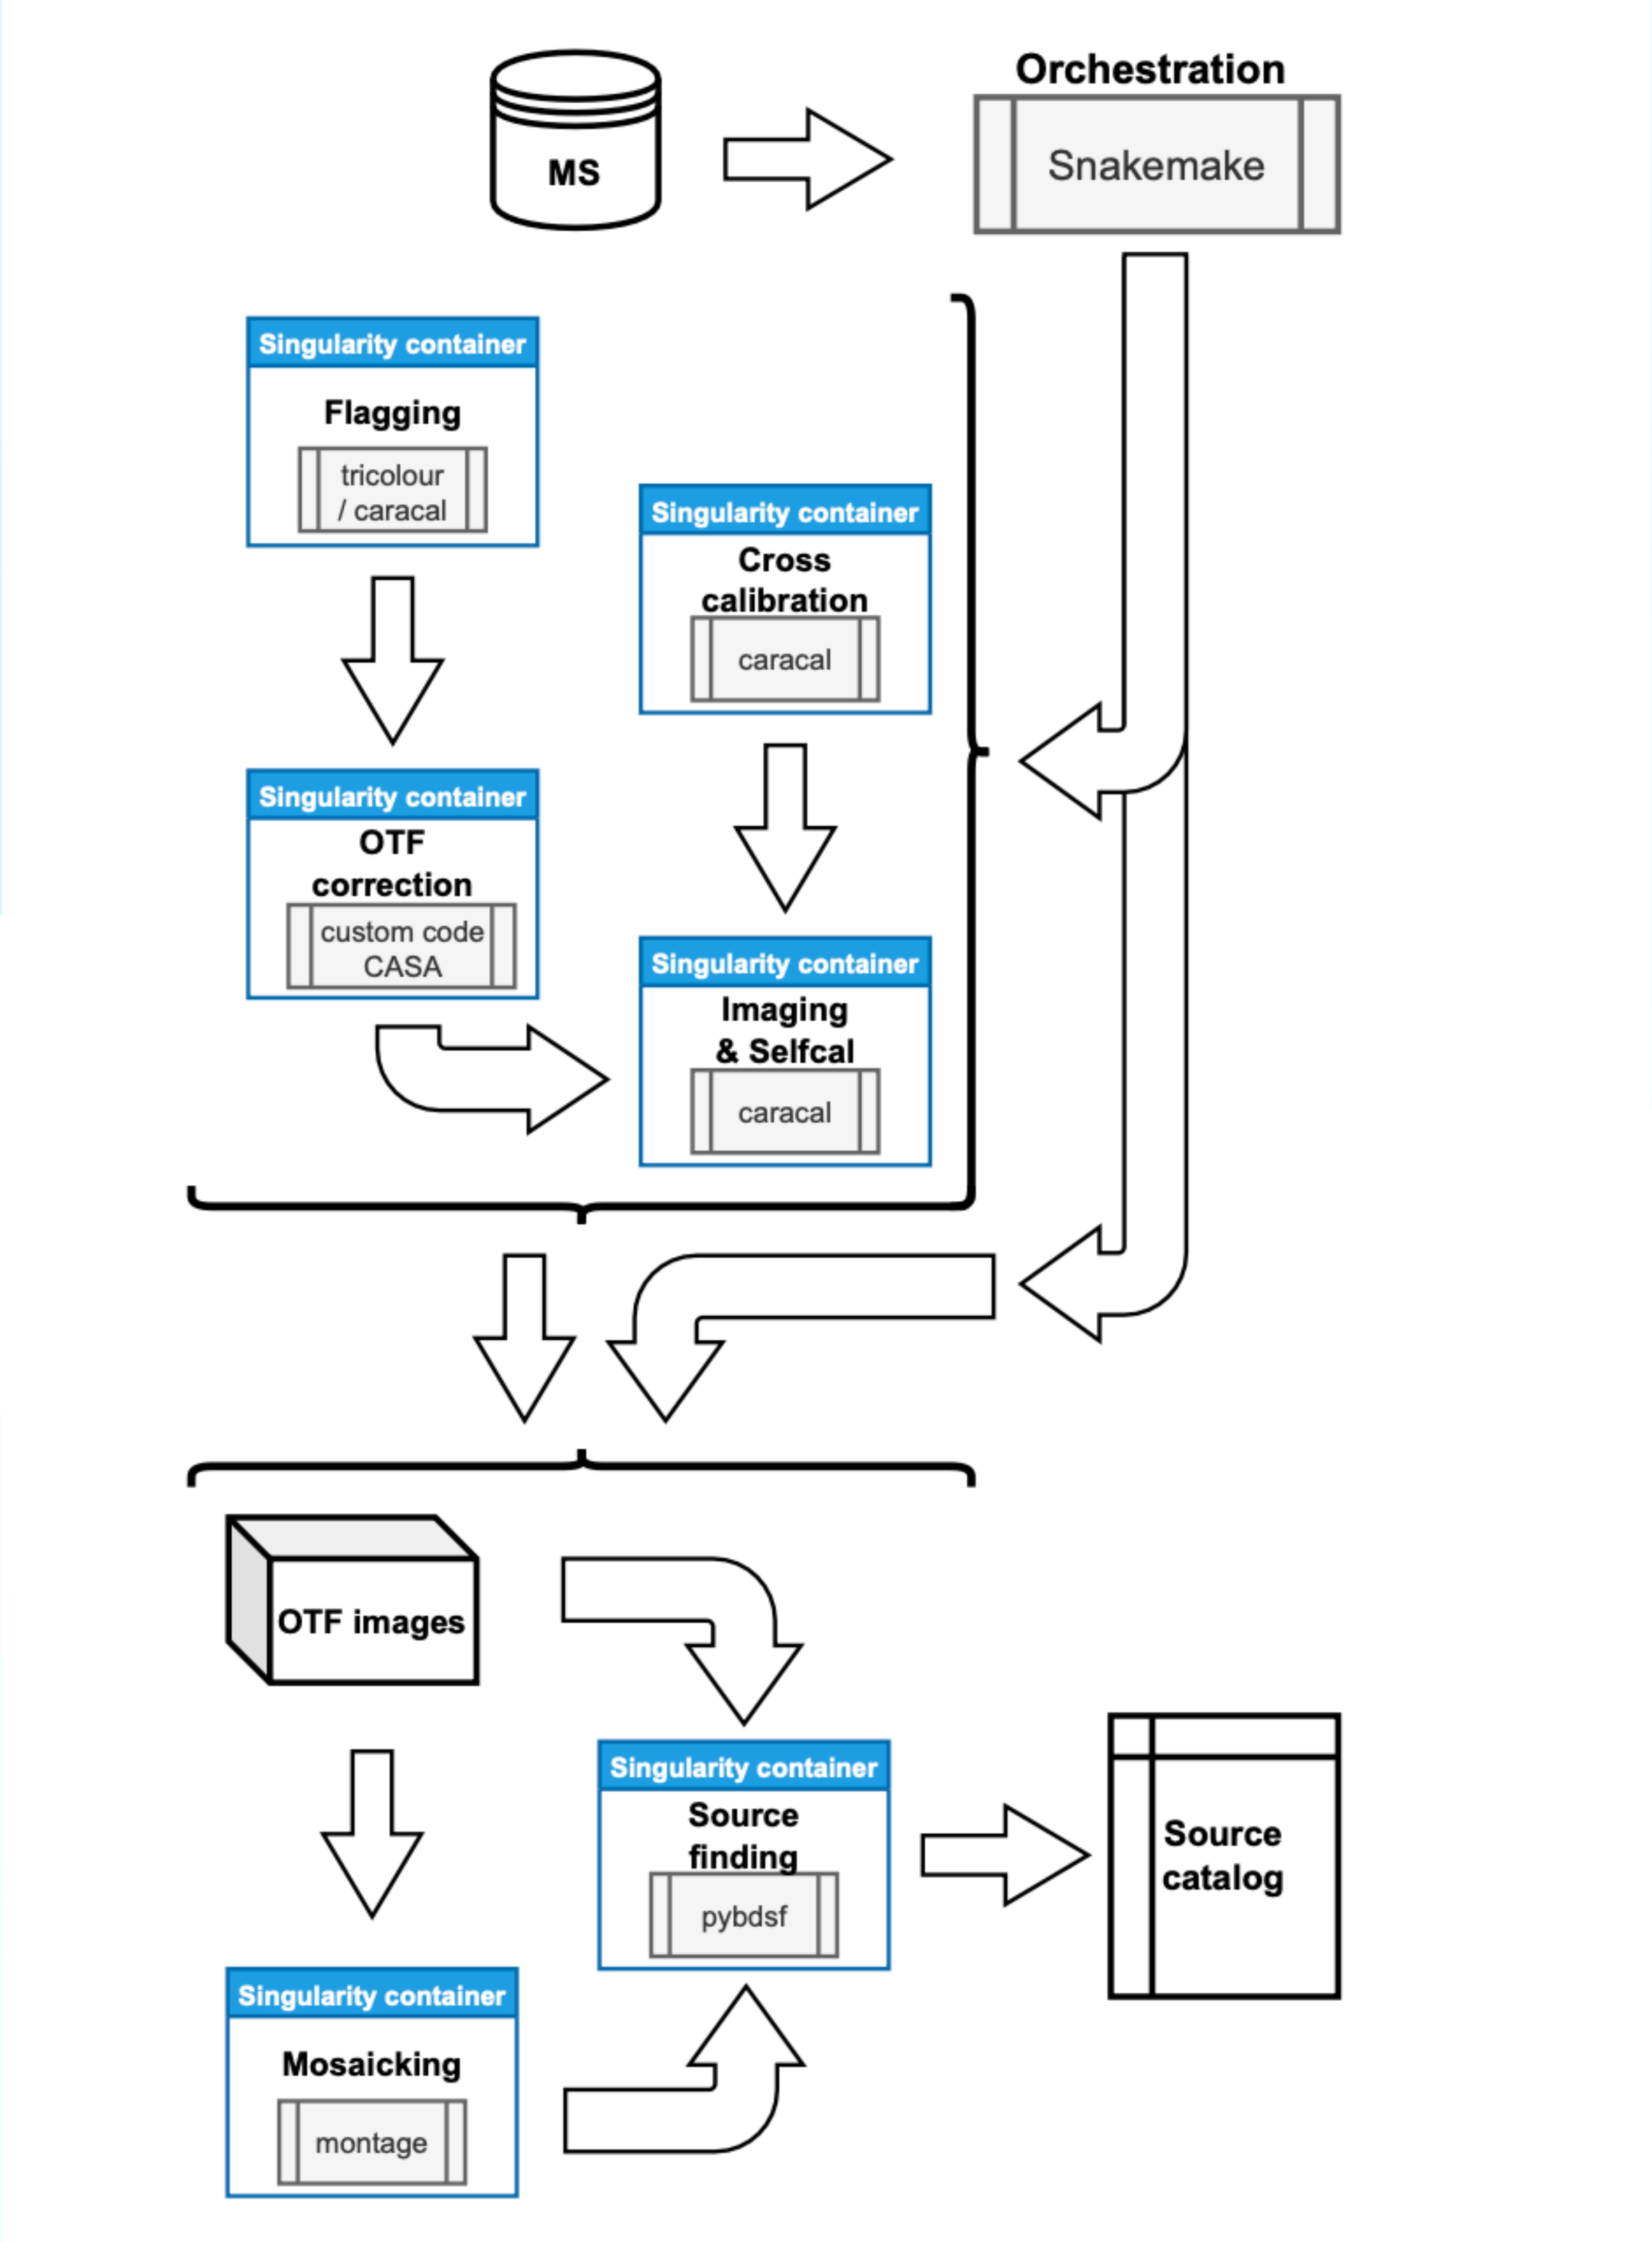

First building block of the pipeline to be implemented in Compute4PUNCH.

It tested container creation for custom software and the use of a workflow manager.

Achievements - OTF correction

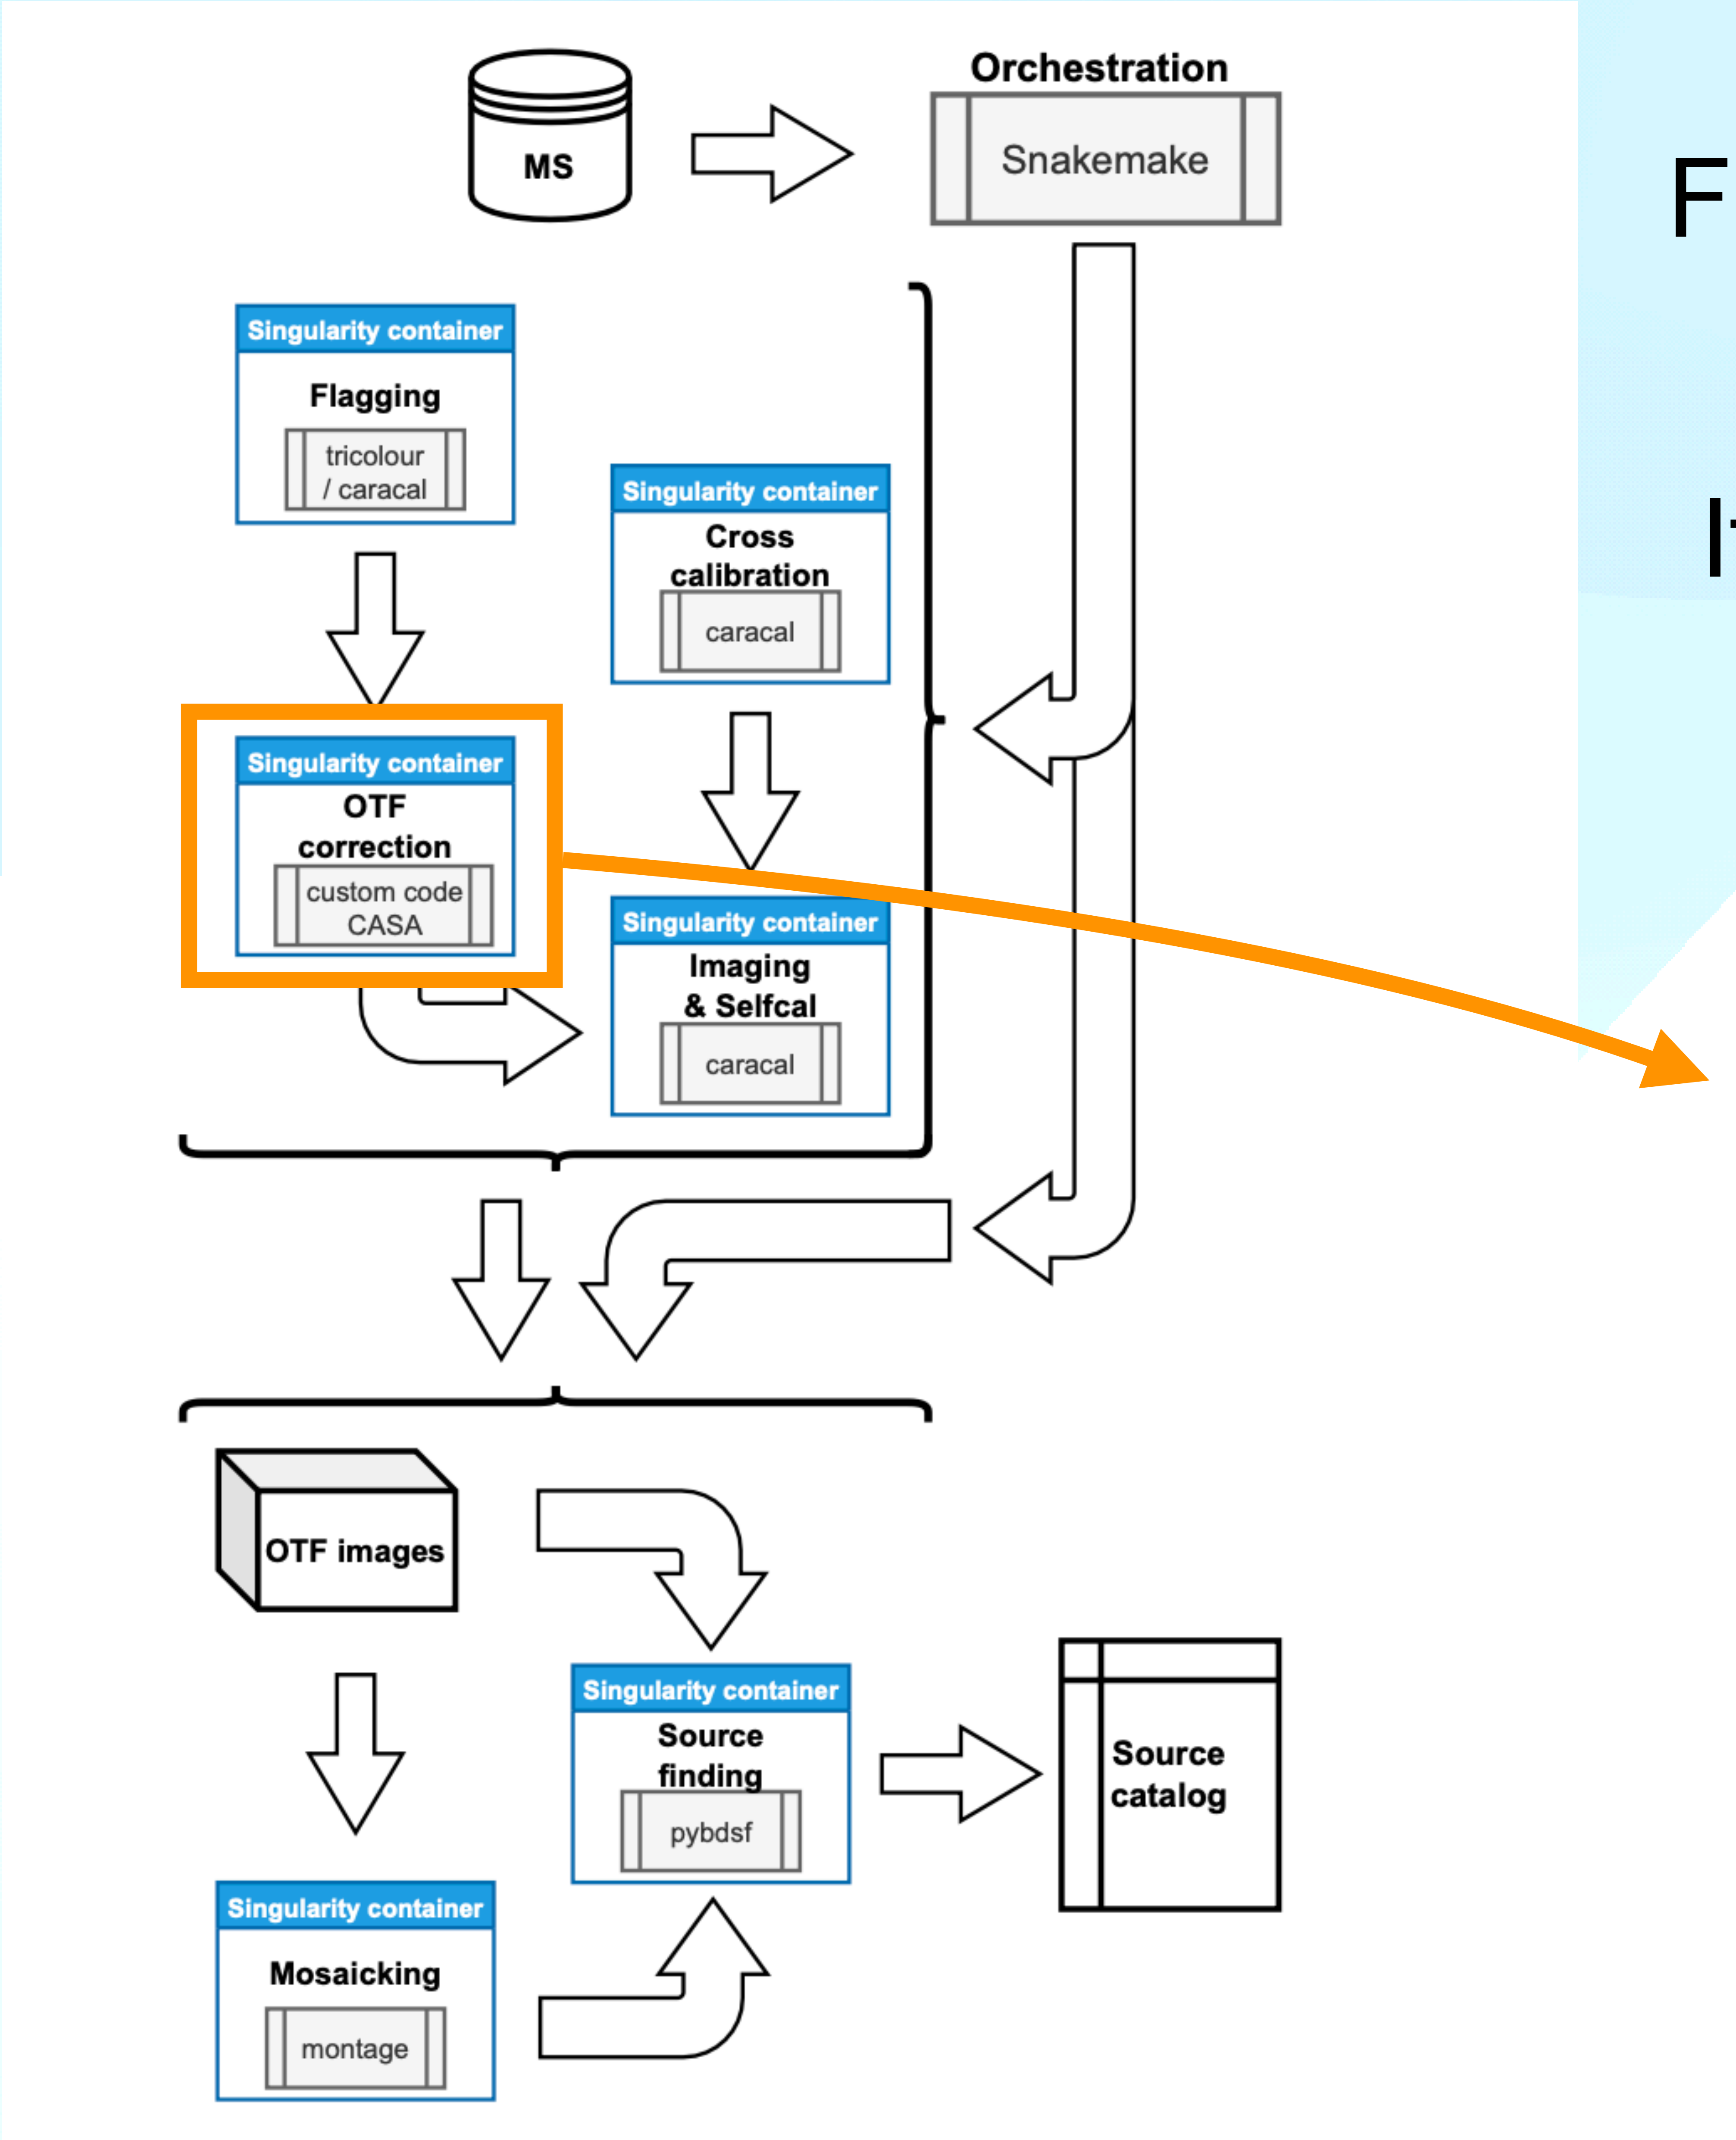

First building block of the pipeline to be implemented in Compute4PUNCH.

It tested container creation for custom software and the use of a workflow manager.

Achievements - OTF correction

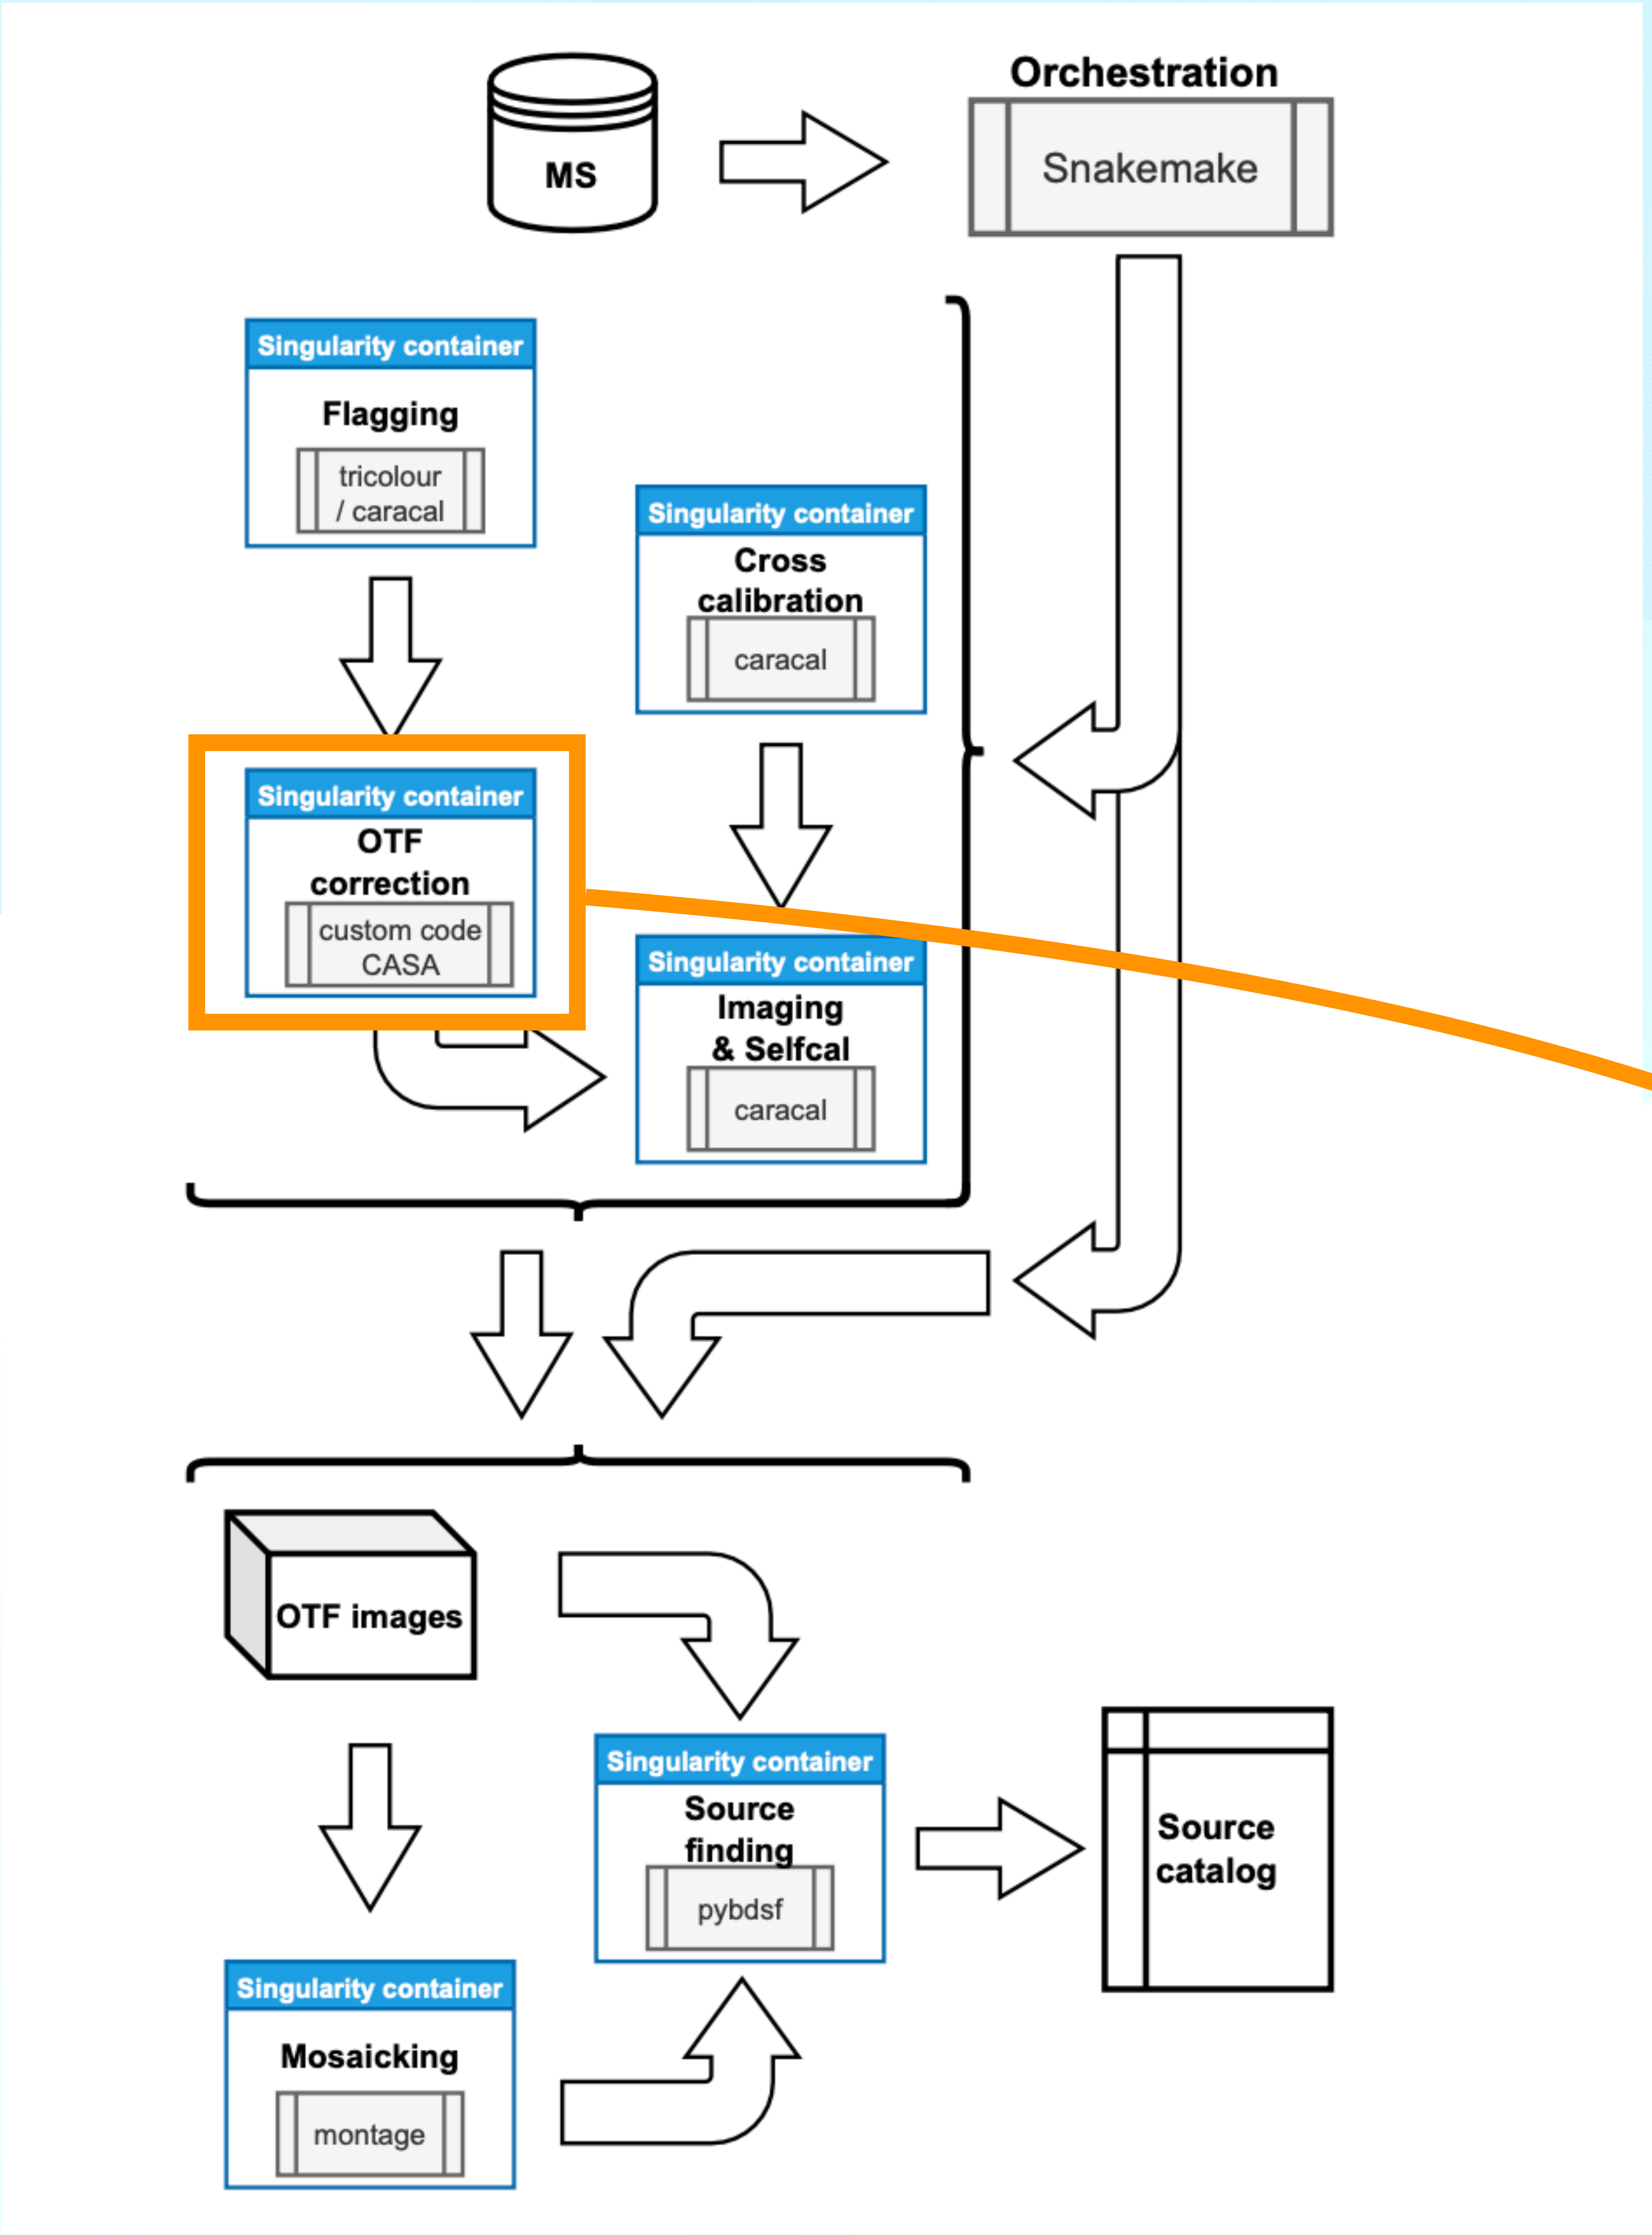

First building block of the pipeline to be implemented in Compute4PUNCH.

It tested container creation for custom software and the use of a workflow manager.

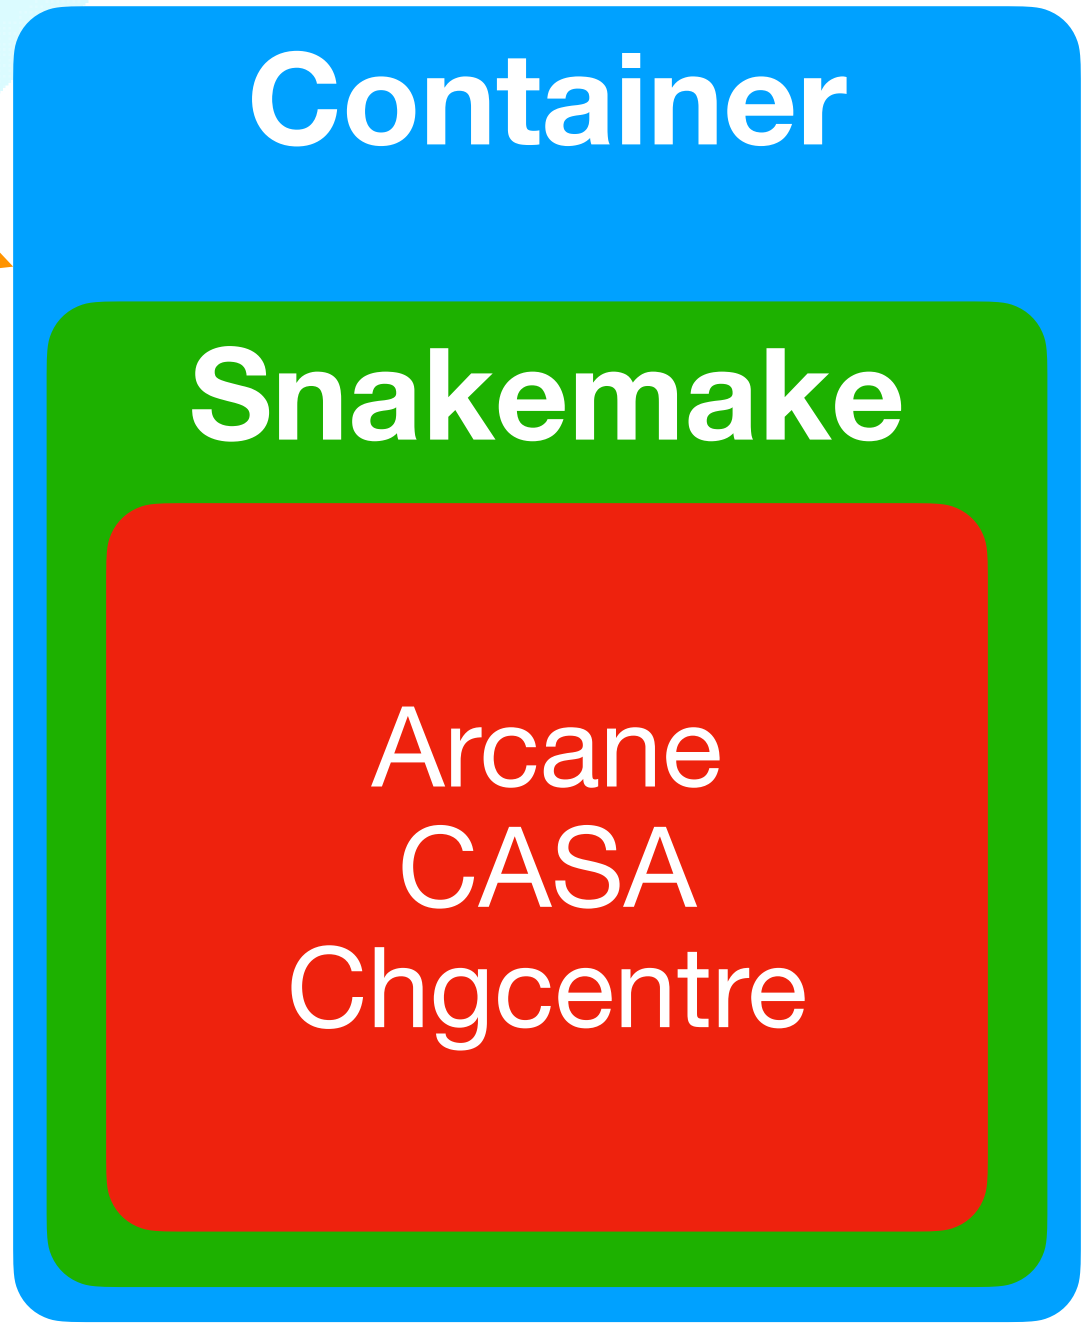

Achievements - OTF correction

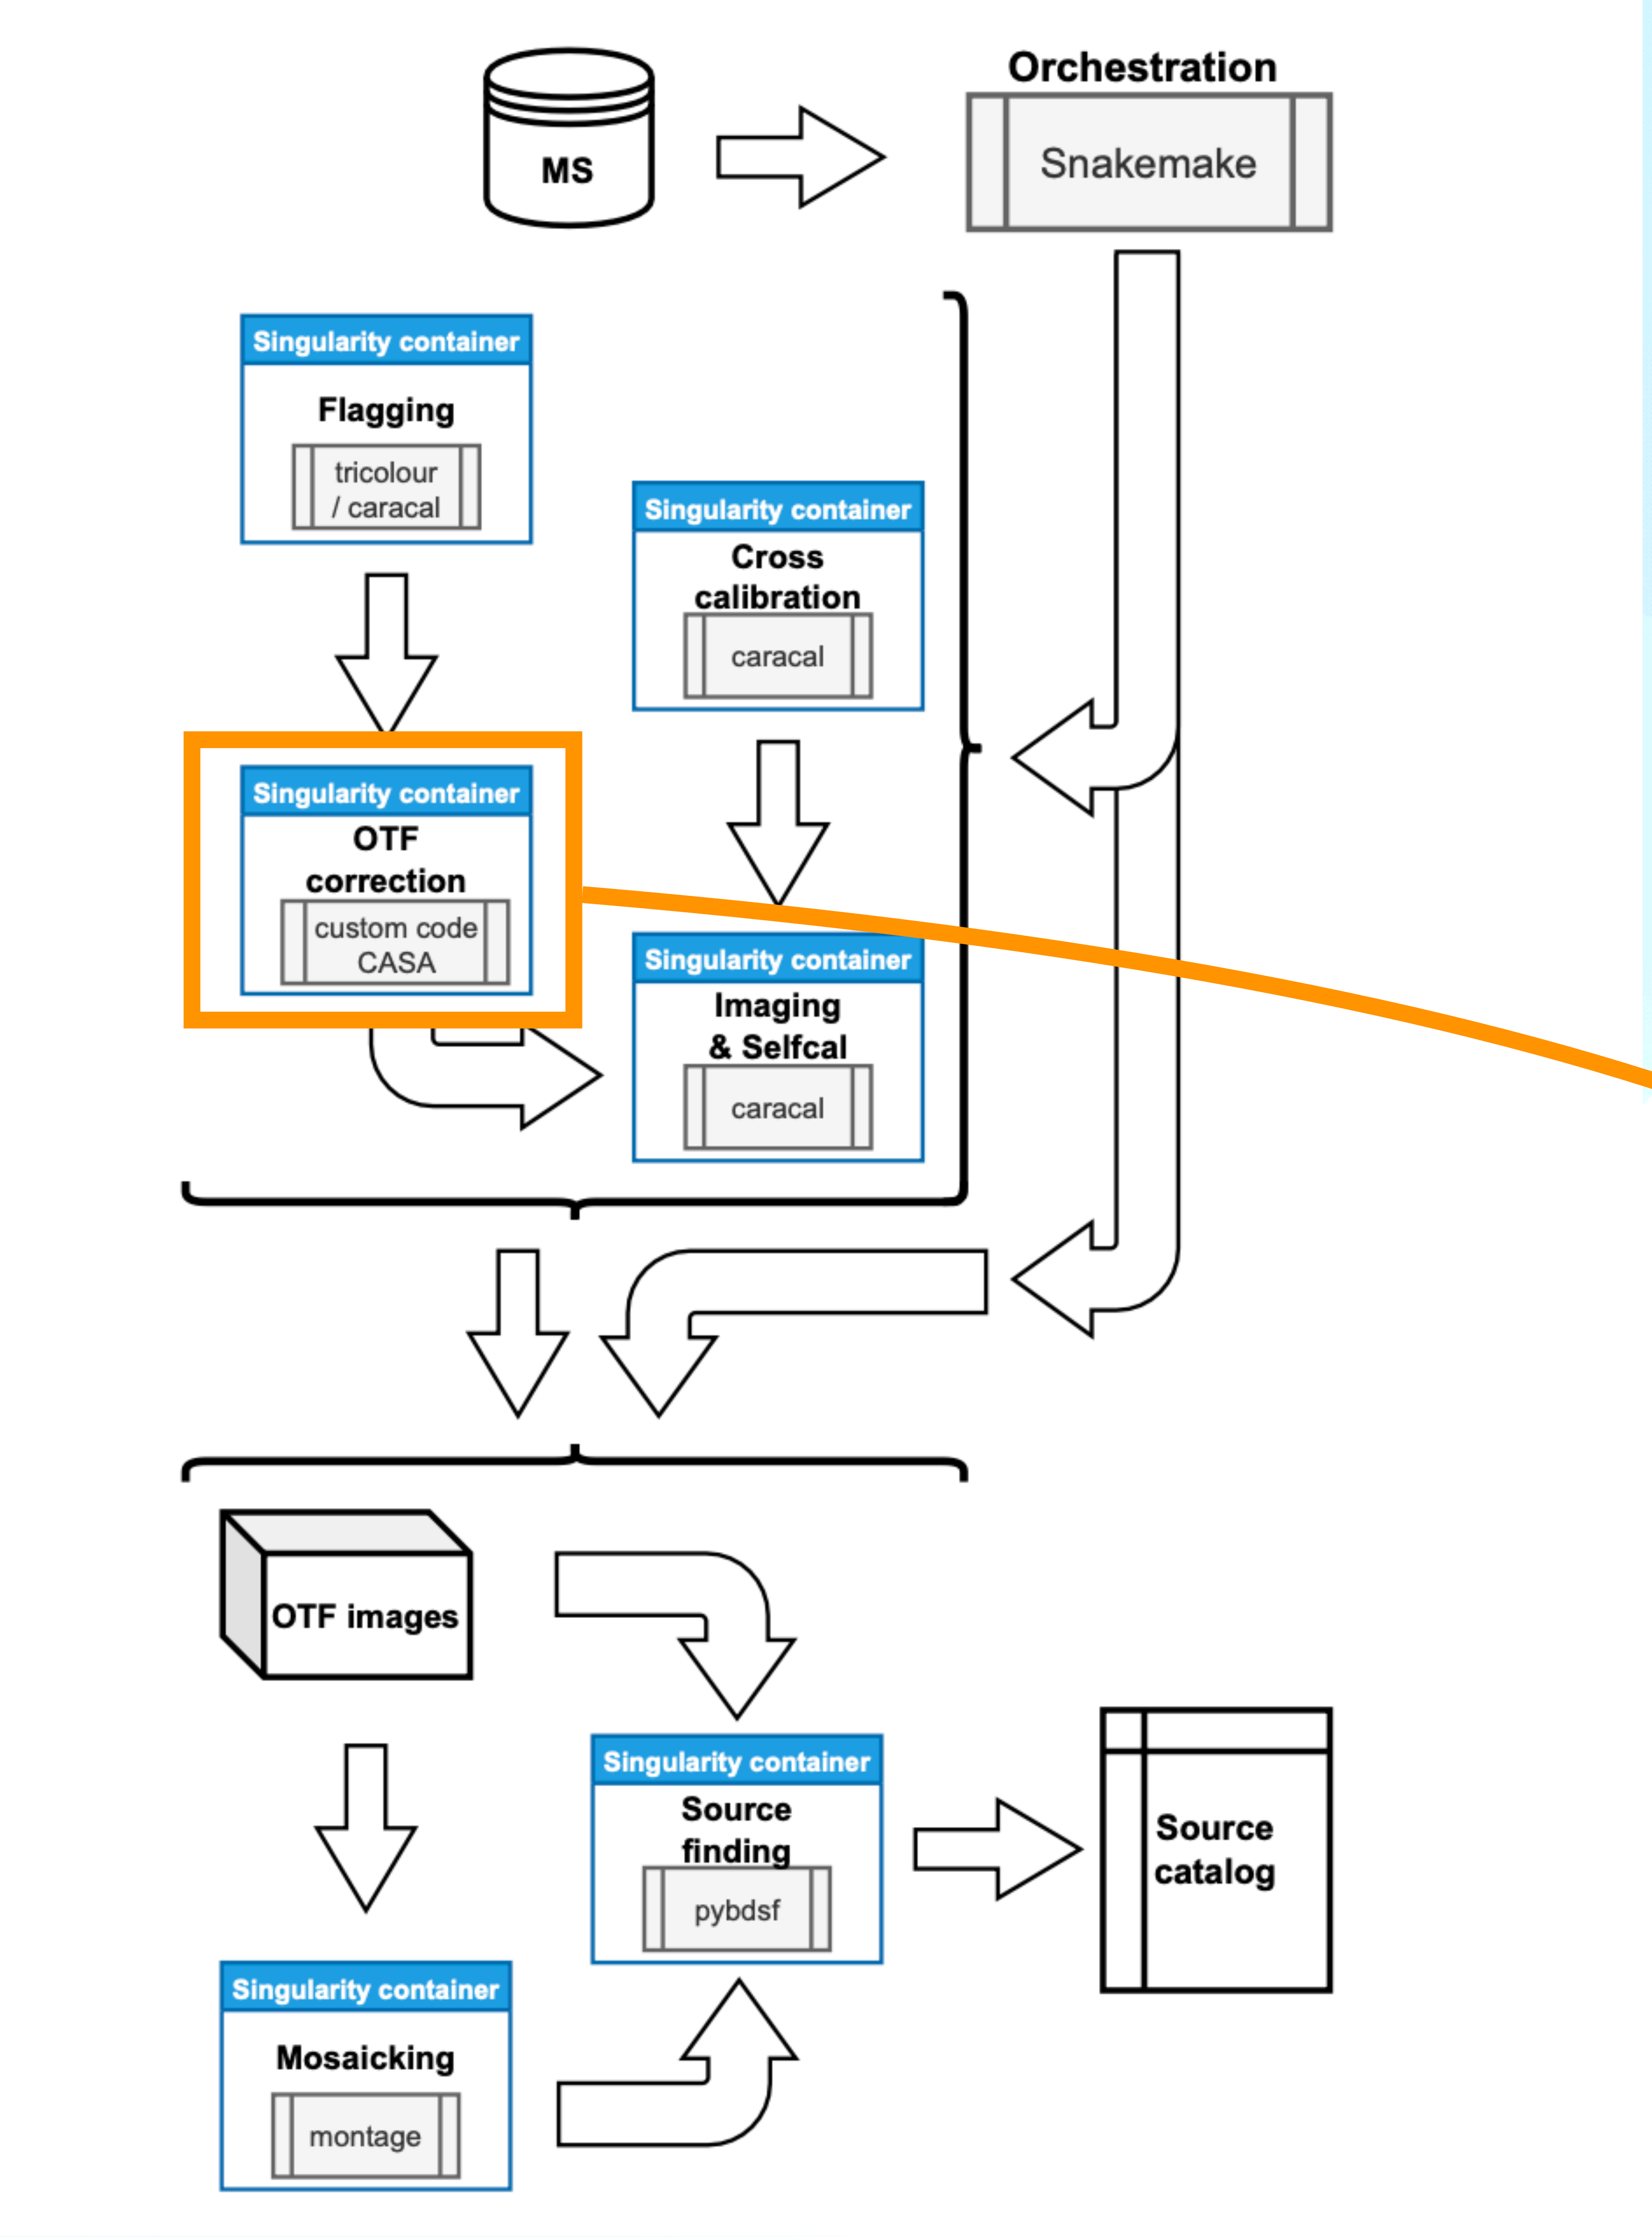

First building block of the pipeline to be implemented in Compute4PUNCH.

It tested container creation for custom software and the use of a workflow manager.

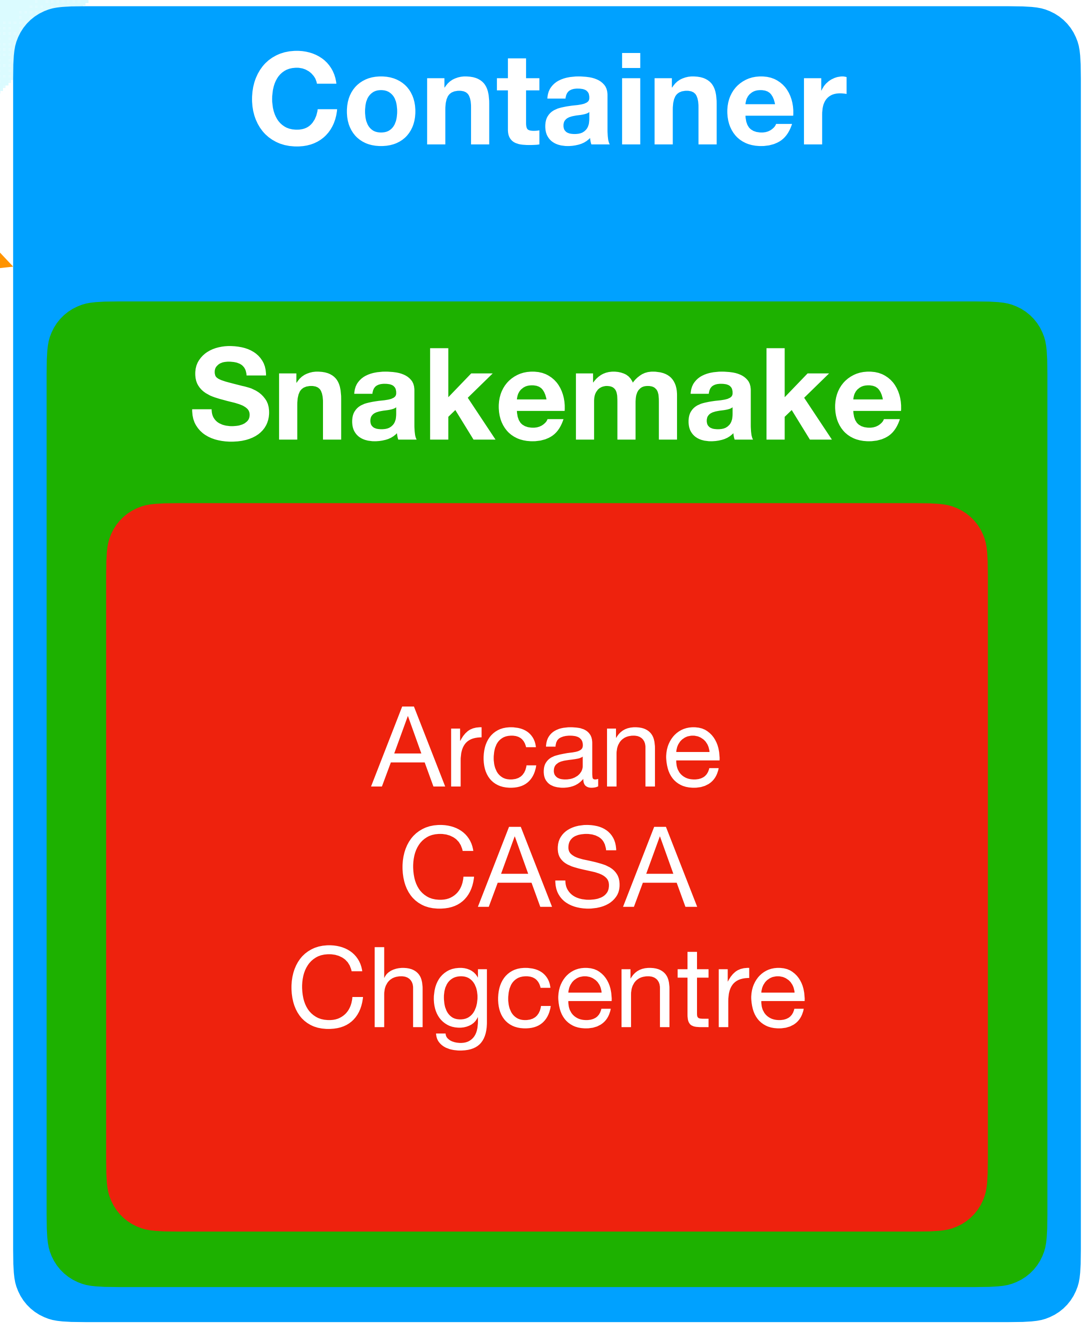

- Container succesfully submitted to C4P container stack.
- Test job ran requesting container from list.
- Container available to use on other clusters, **solution open to be scaled up**.

Achievements - orchestration

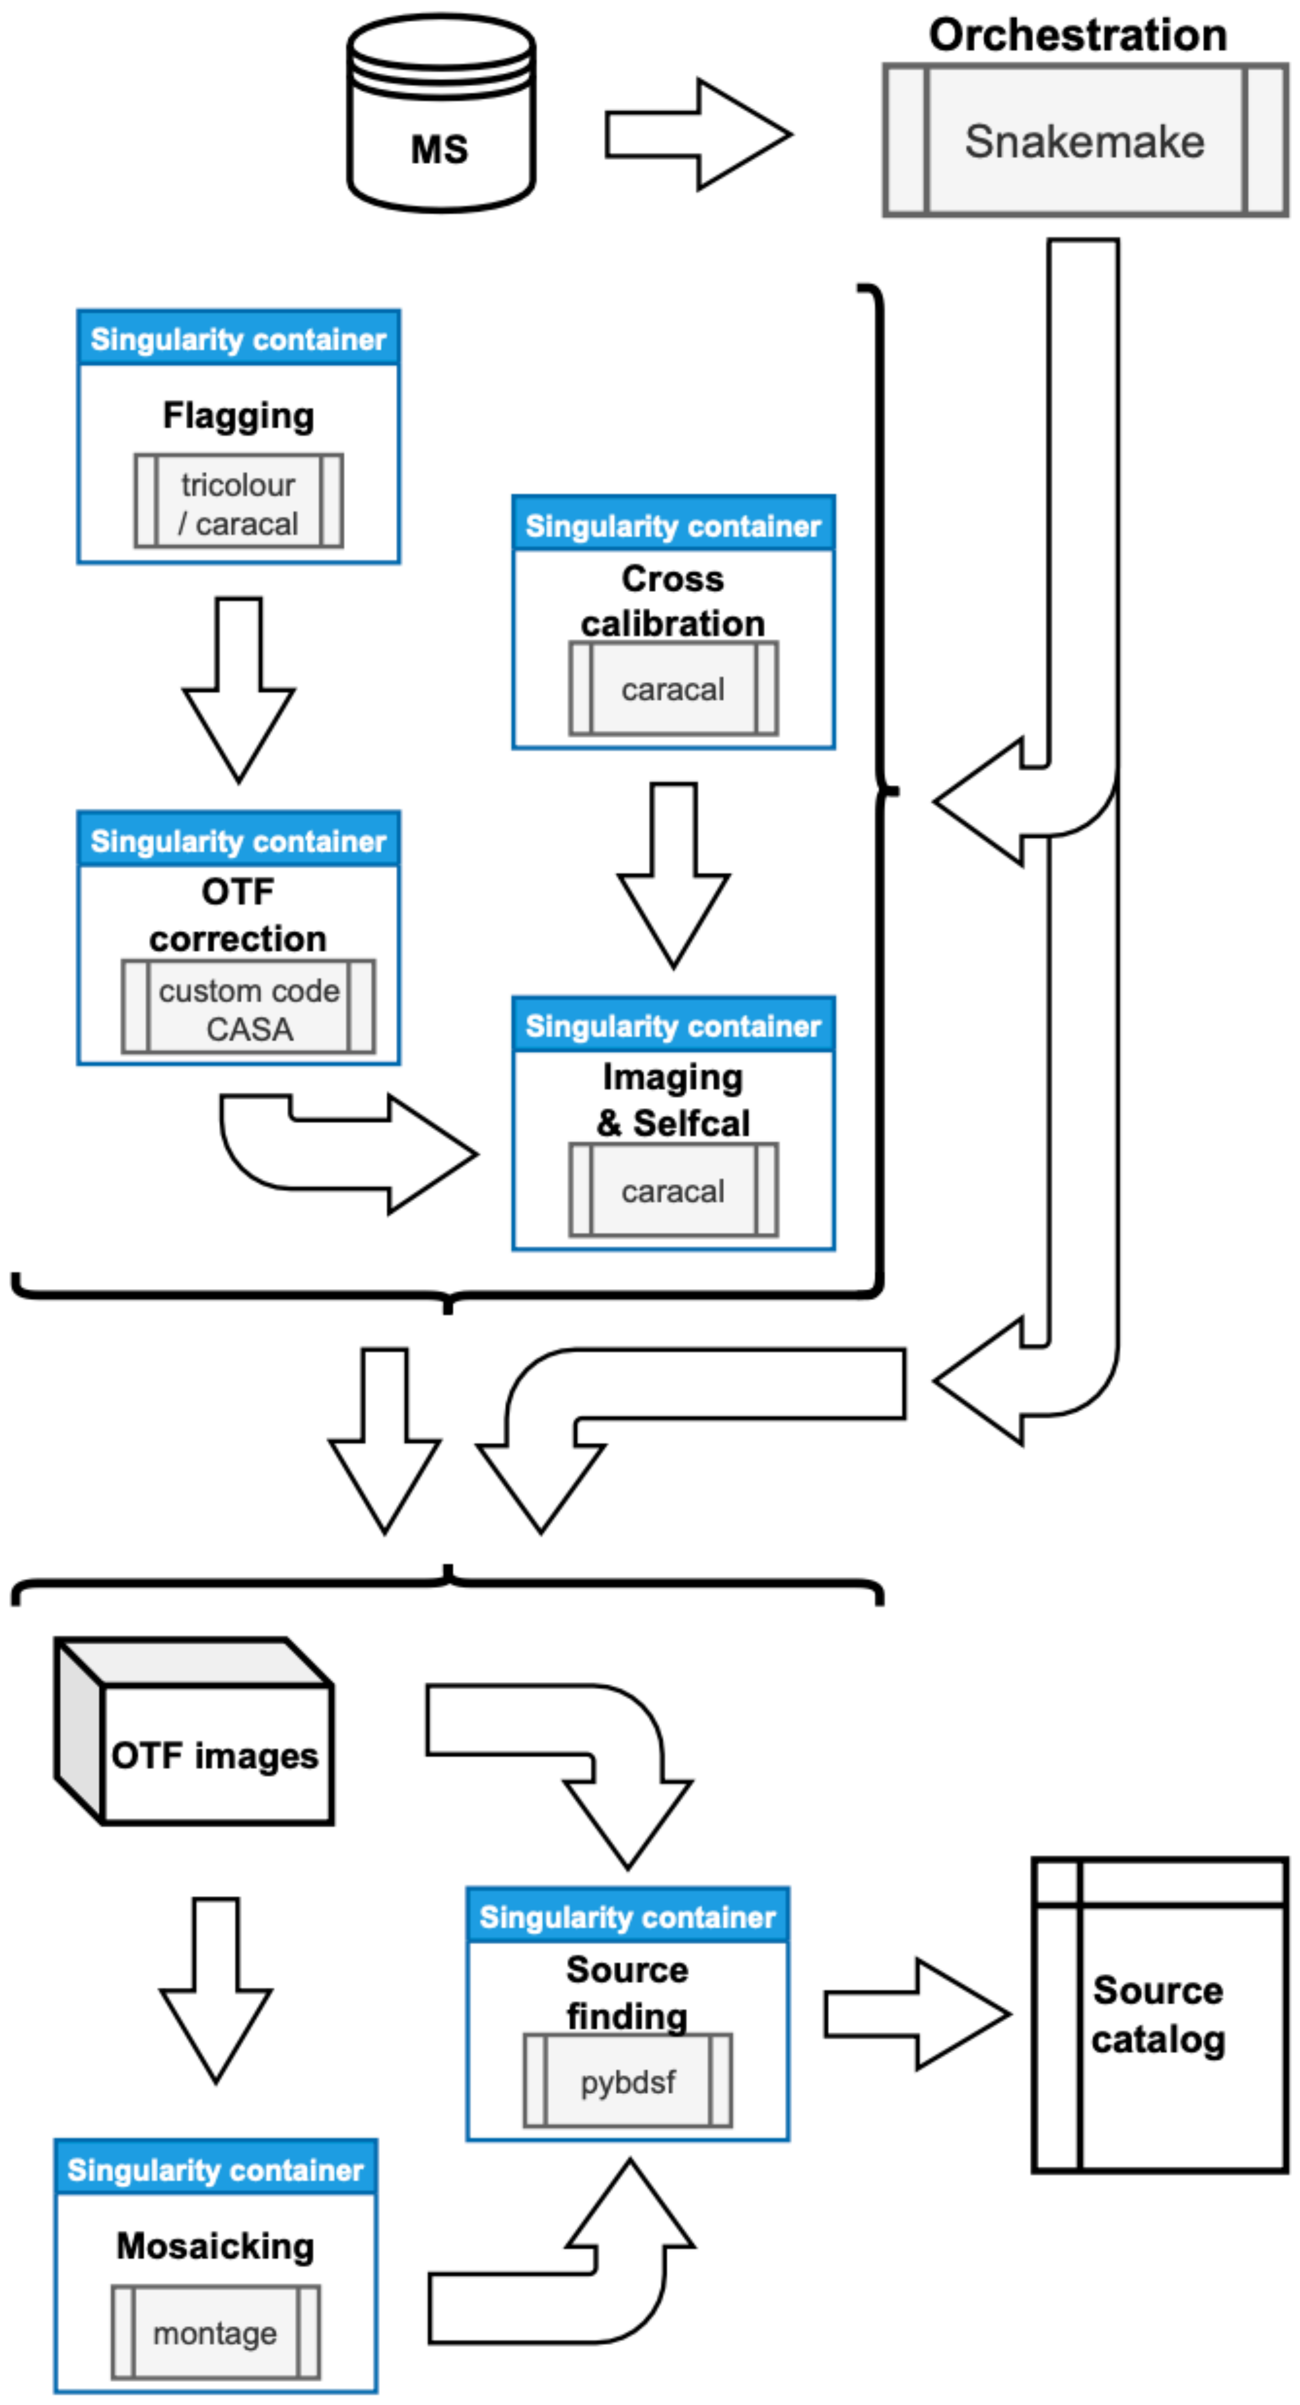

Achievements - orchestration

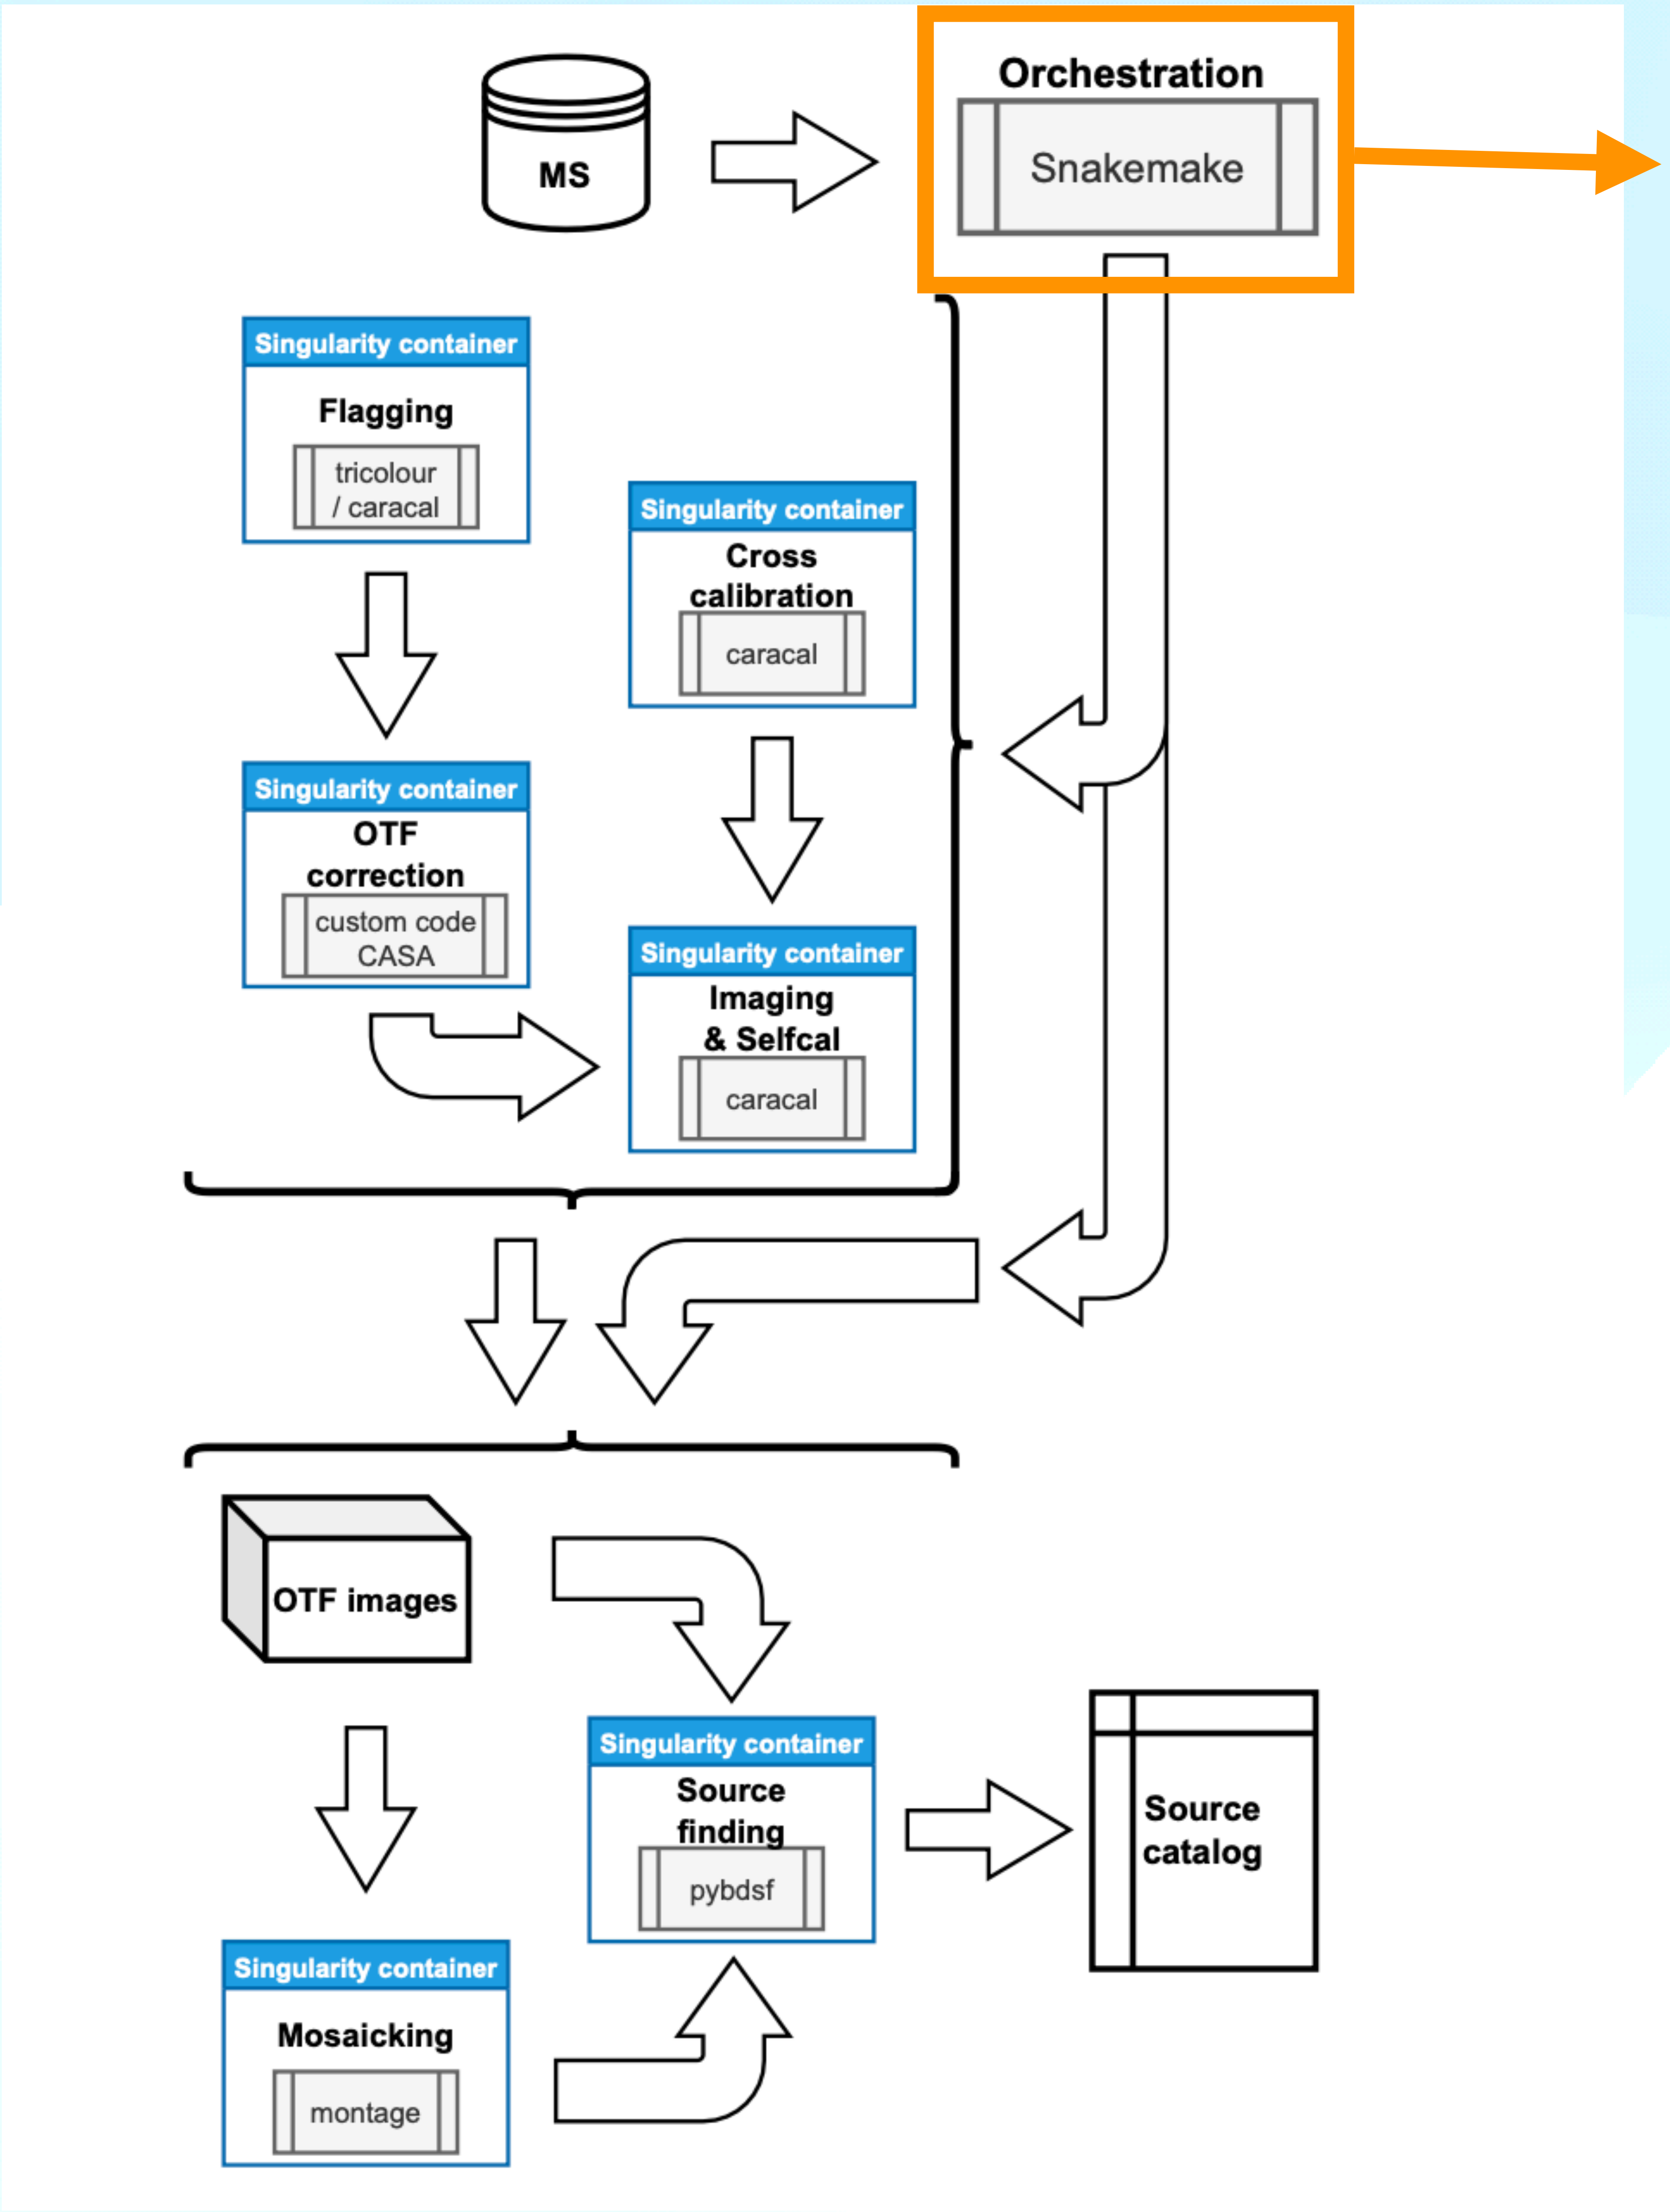

Achievements - orchestration

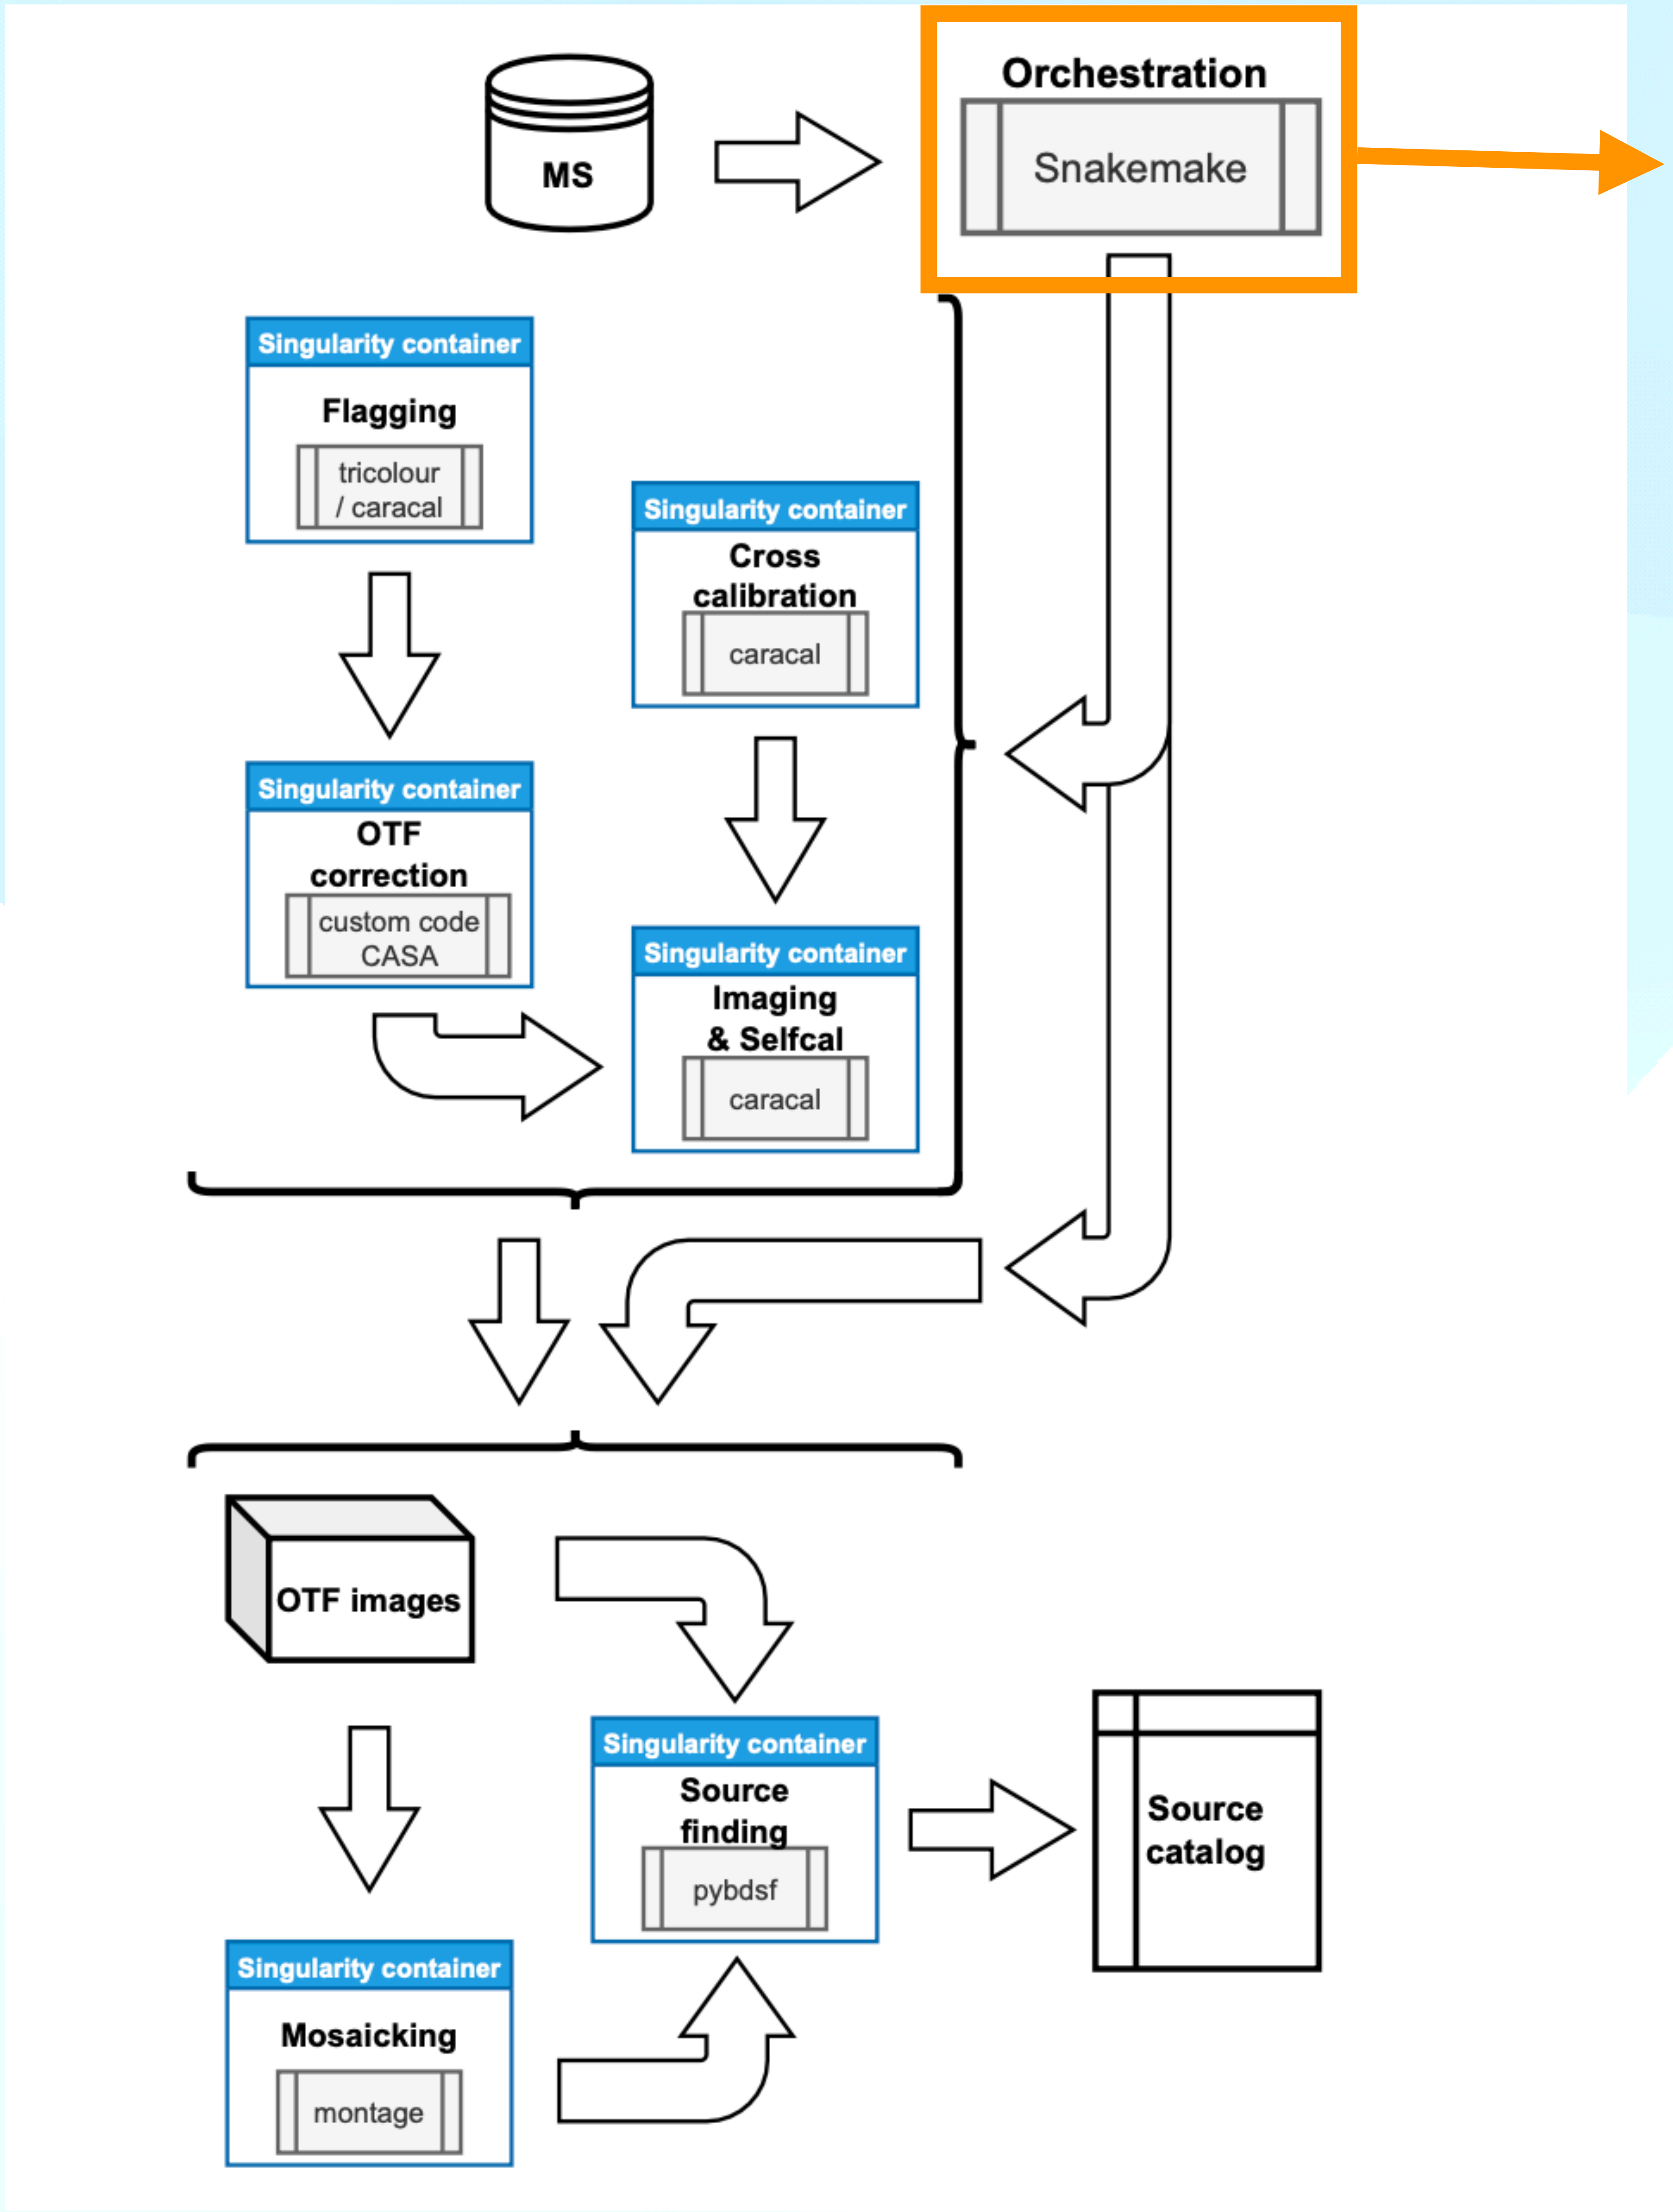

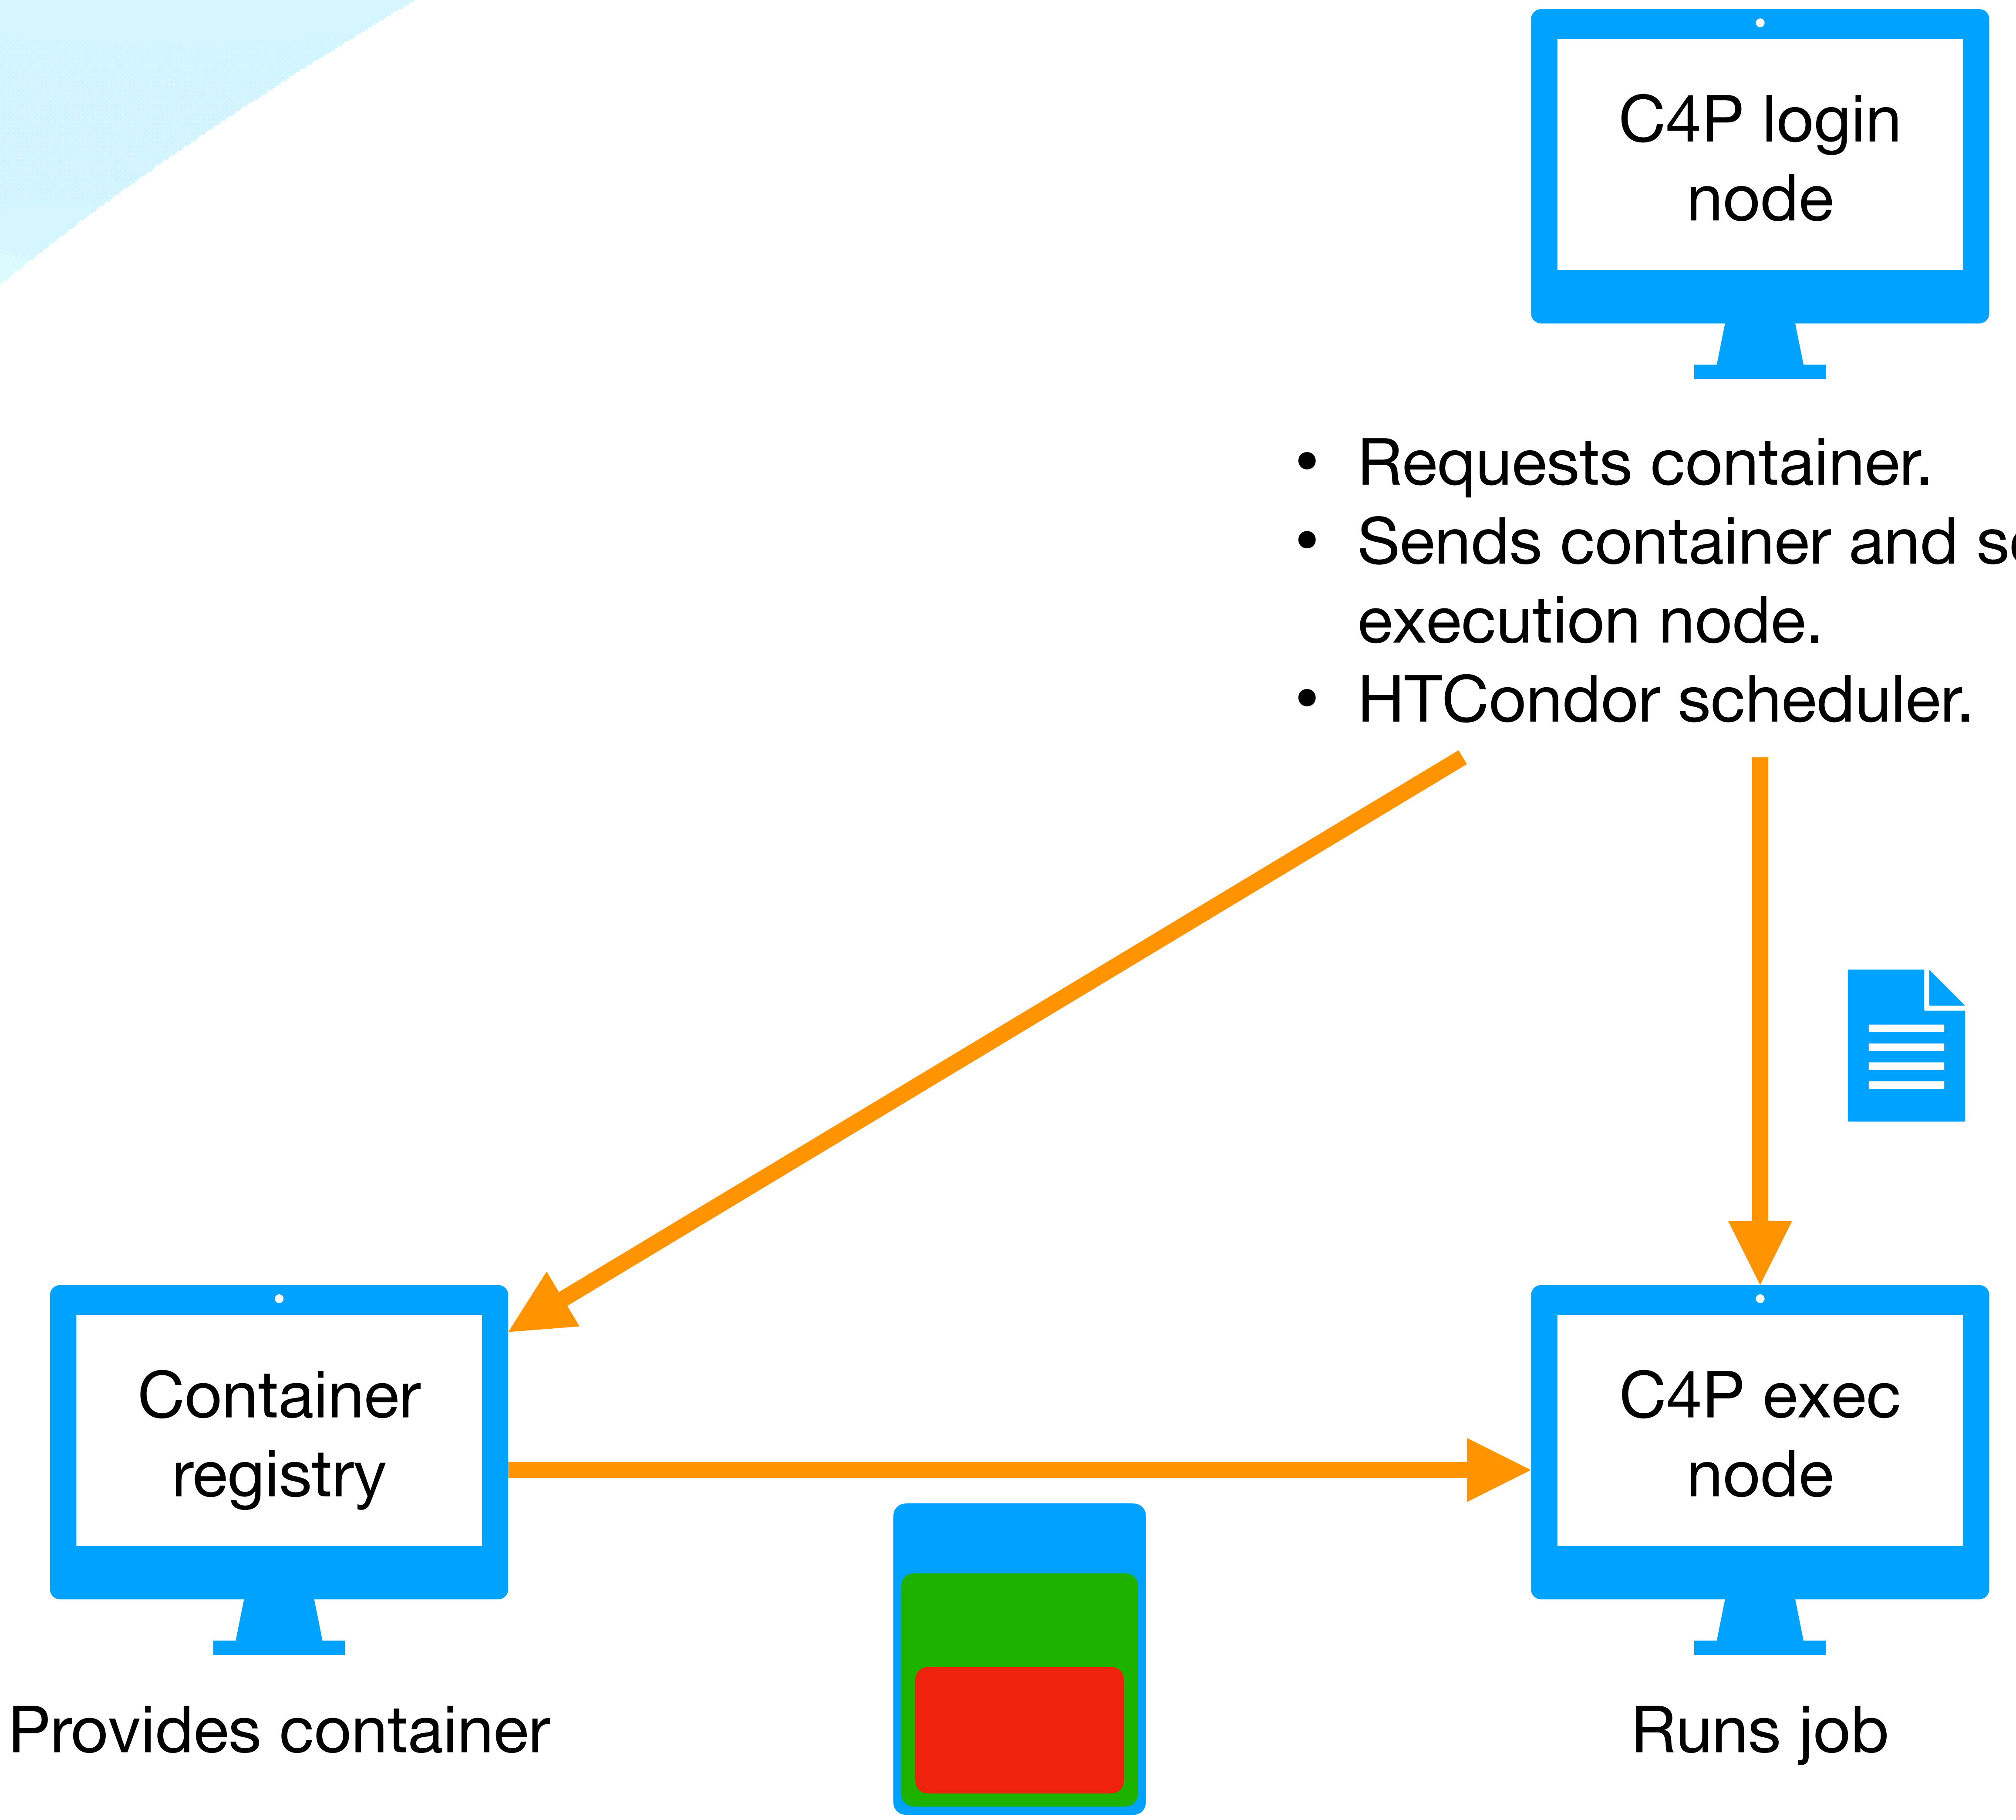

- Requests container.
- Sends container and script to execution node.
- HTCondor scheduler.

Achievements - orchestration

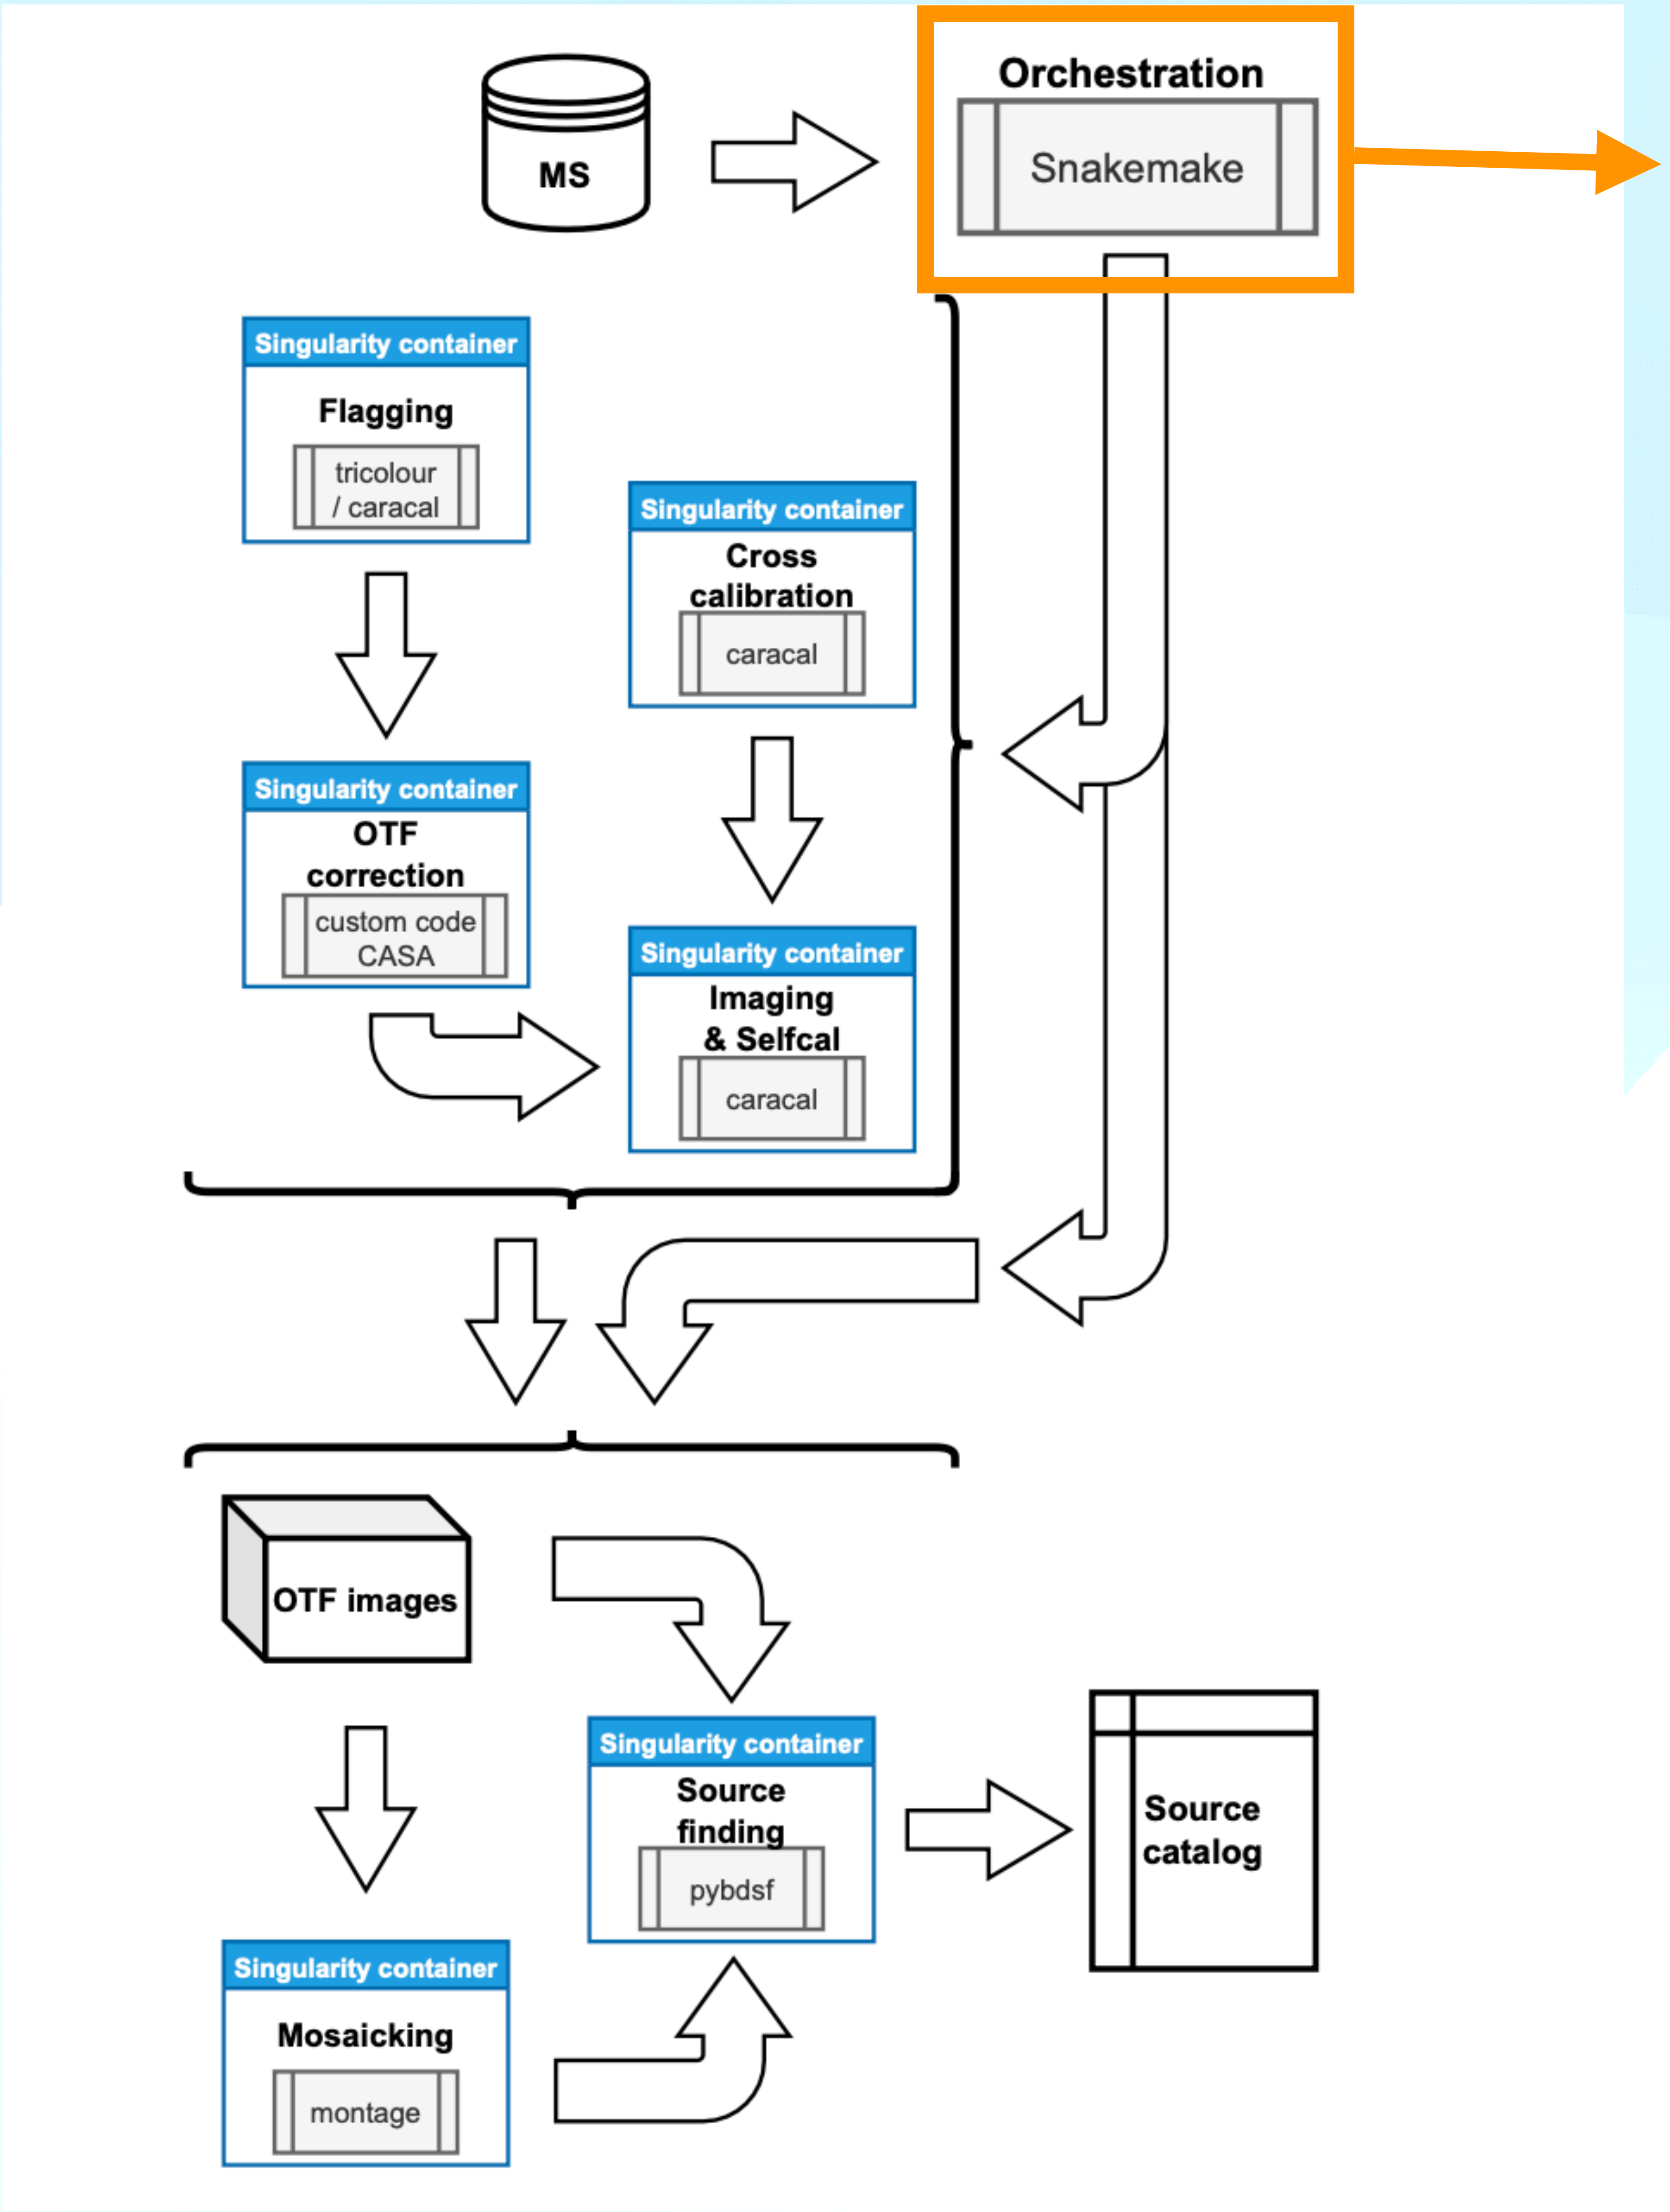

Job submission to Compute4PUNCH done from external cluster via snakemake+ssh command to HTCondor.

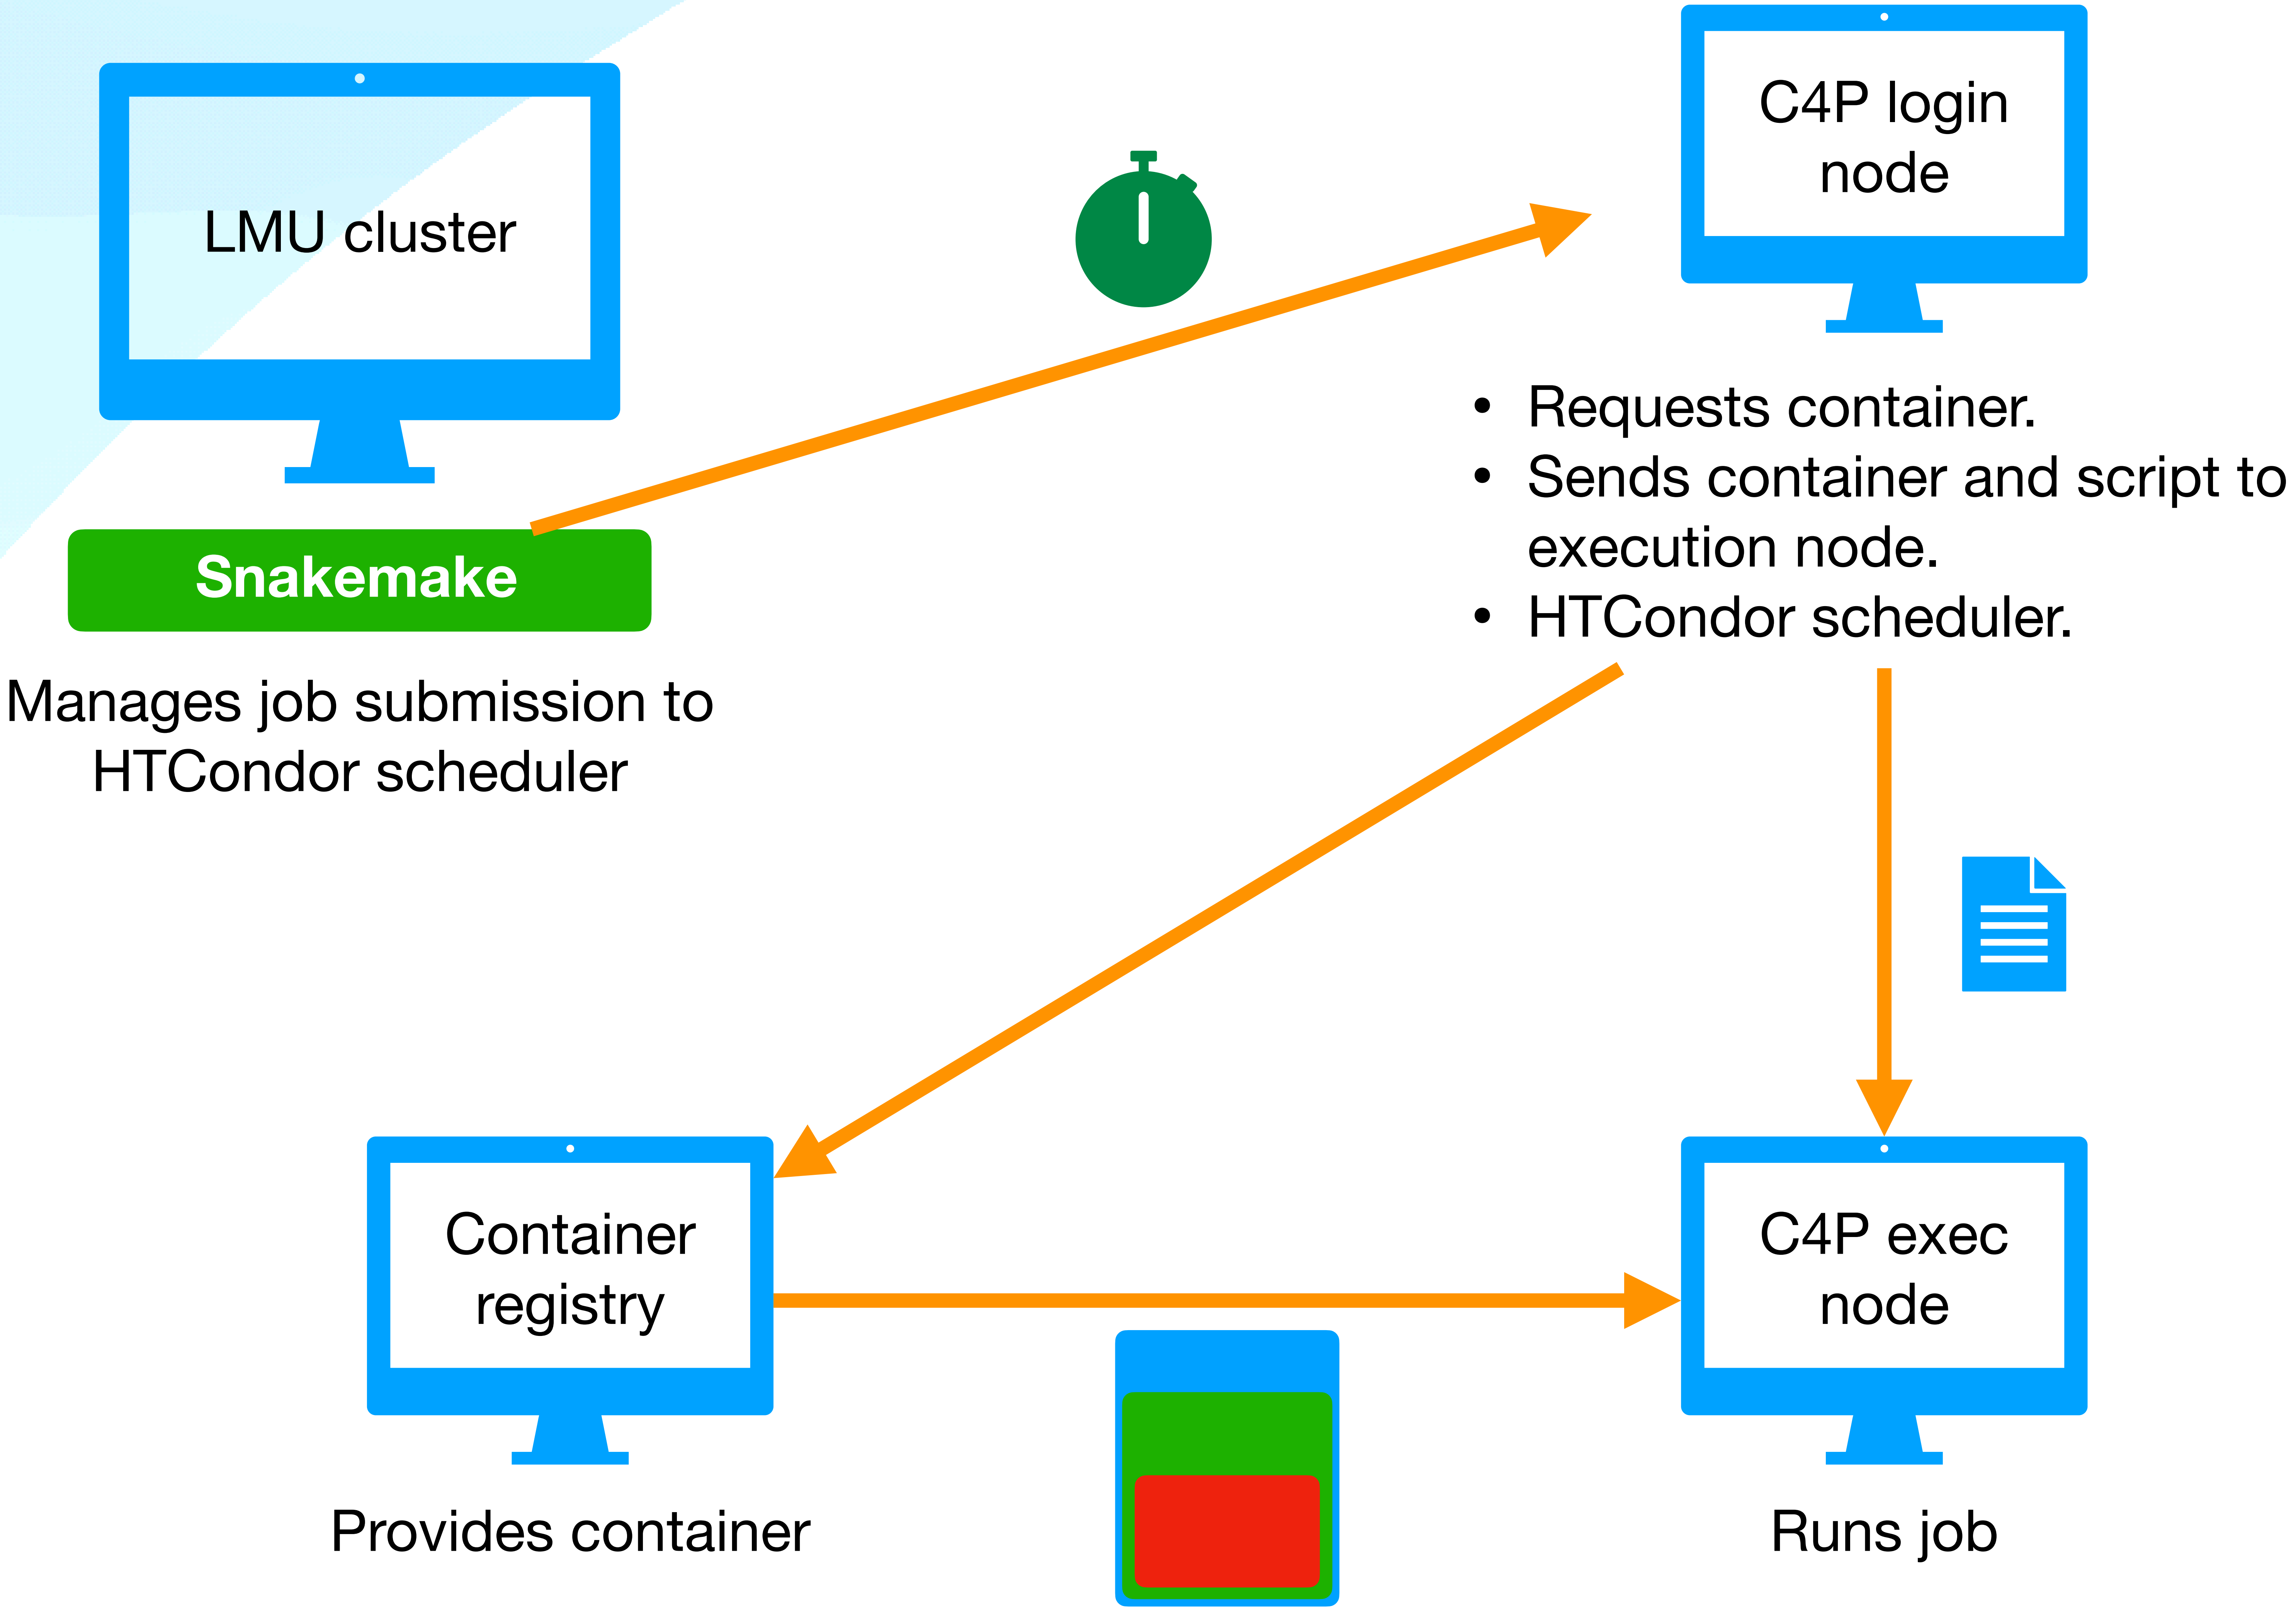

Achievements - orchestration

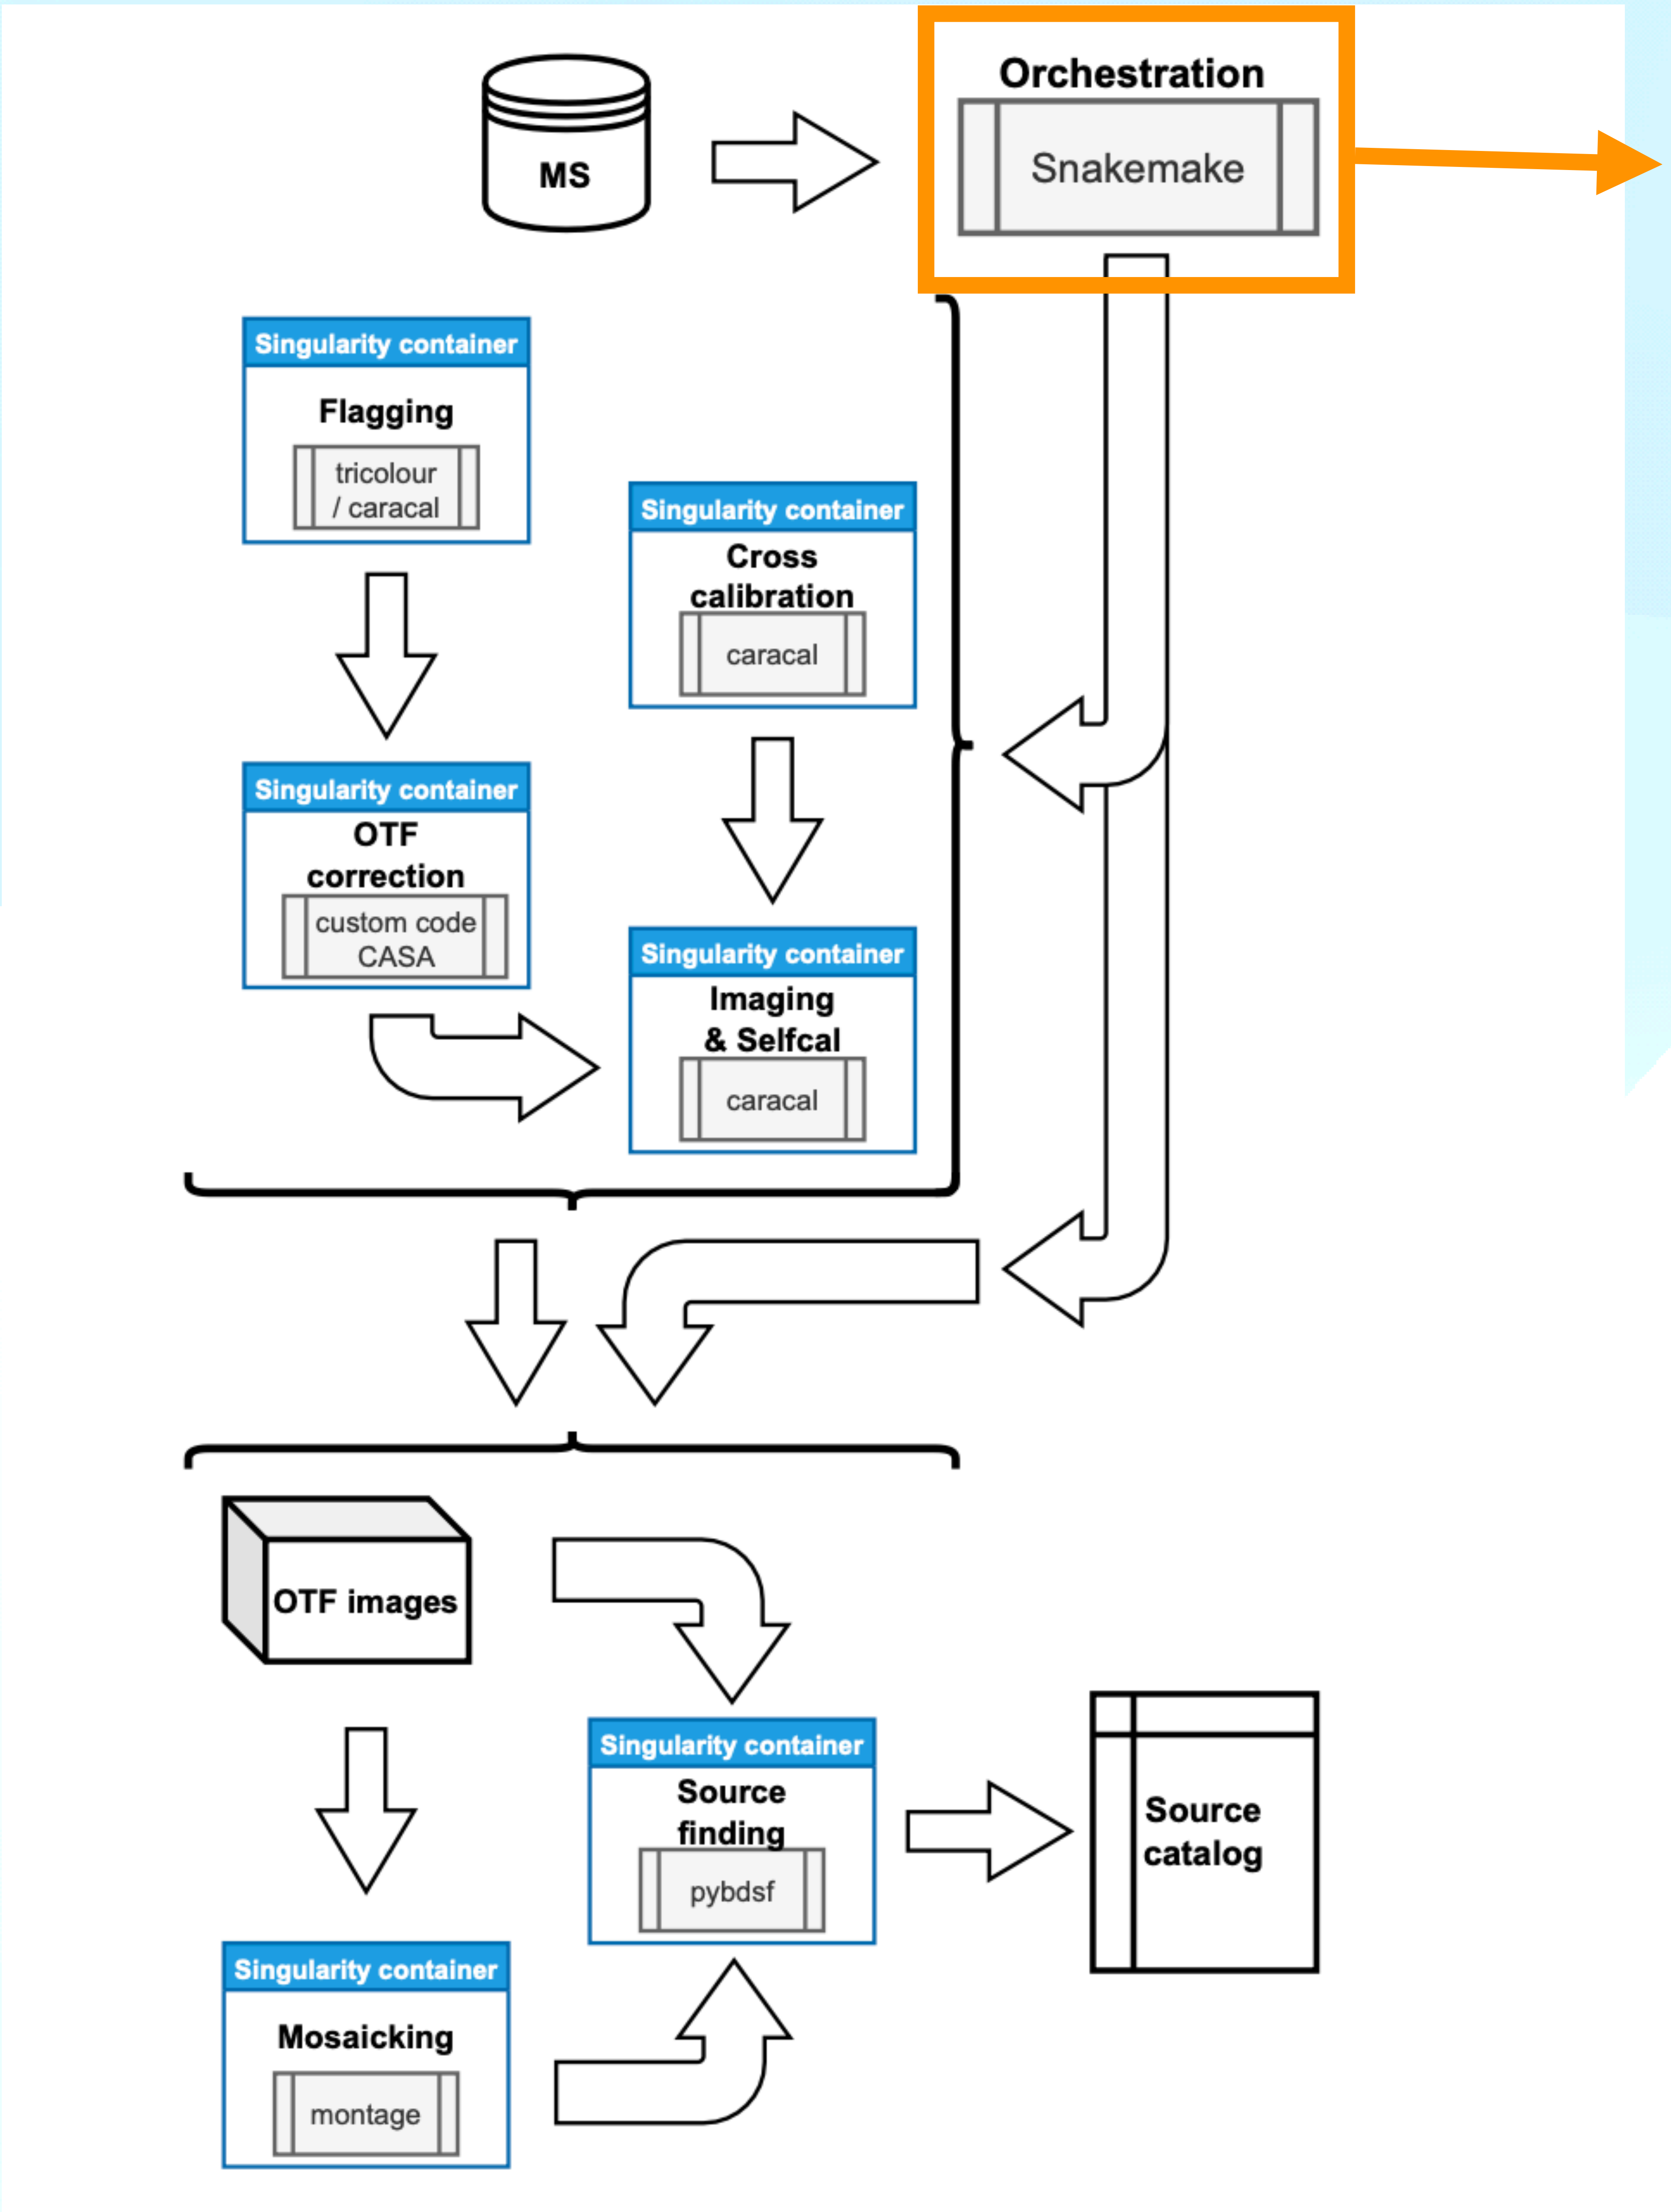

Future implementation will use REANA

Achievements - Imaging

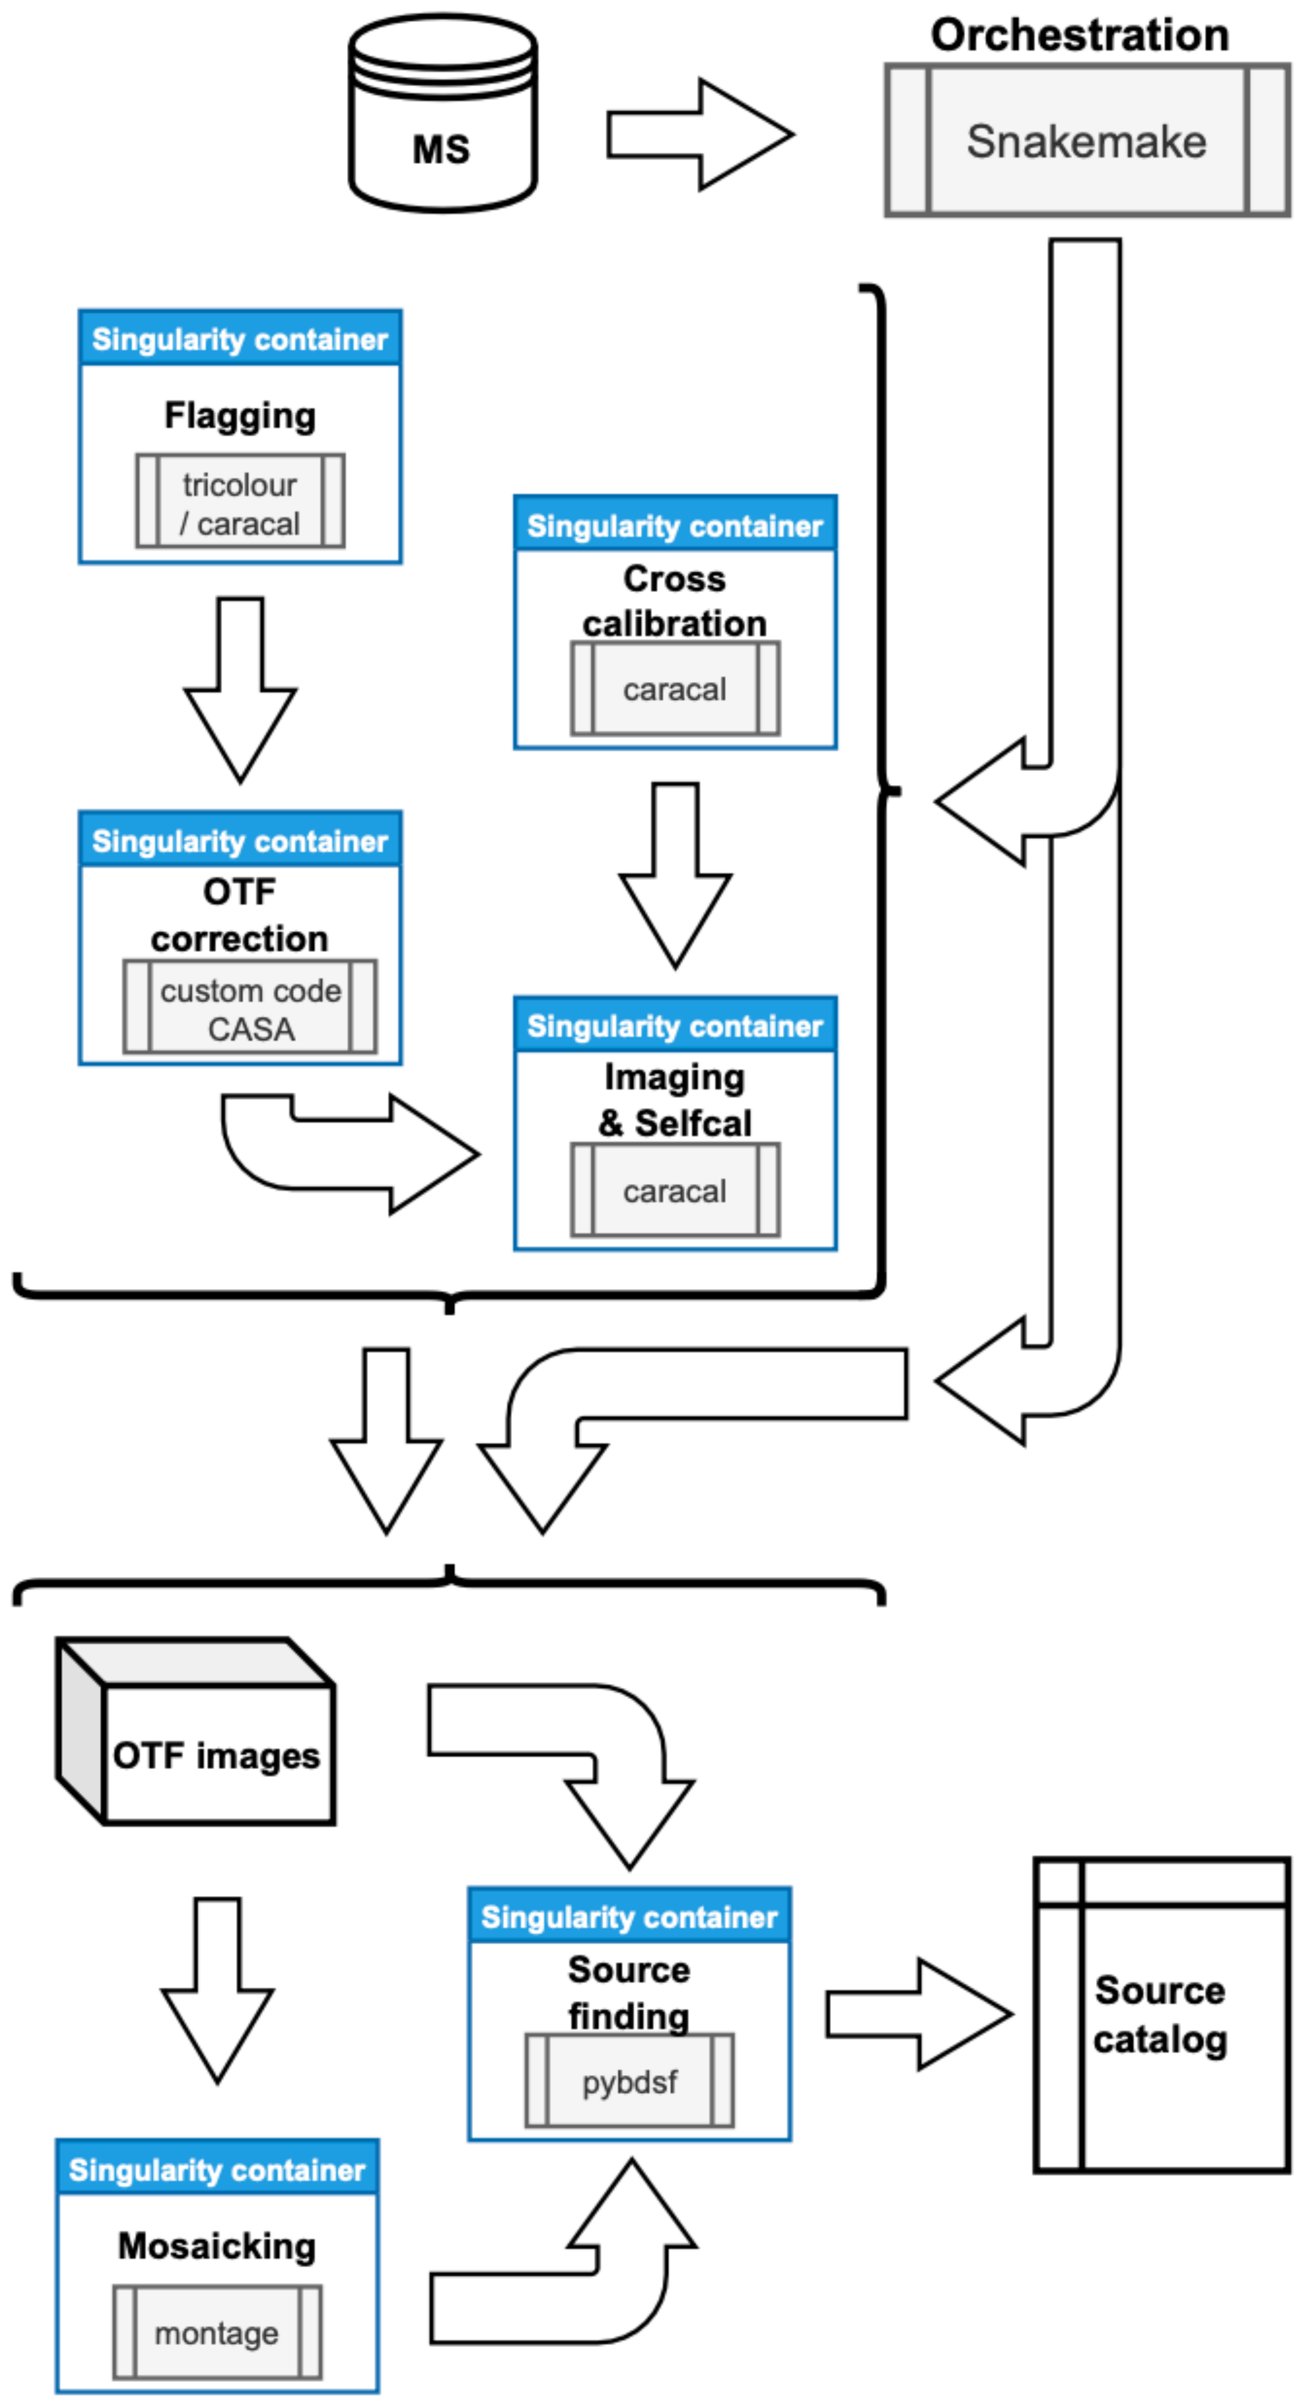

Biggest challenge faced so far: **how to use pre-containerized software**?

Imaging portion needs the software caracal.

Achievements - Imaging

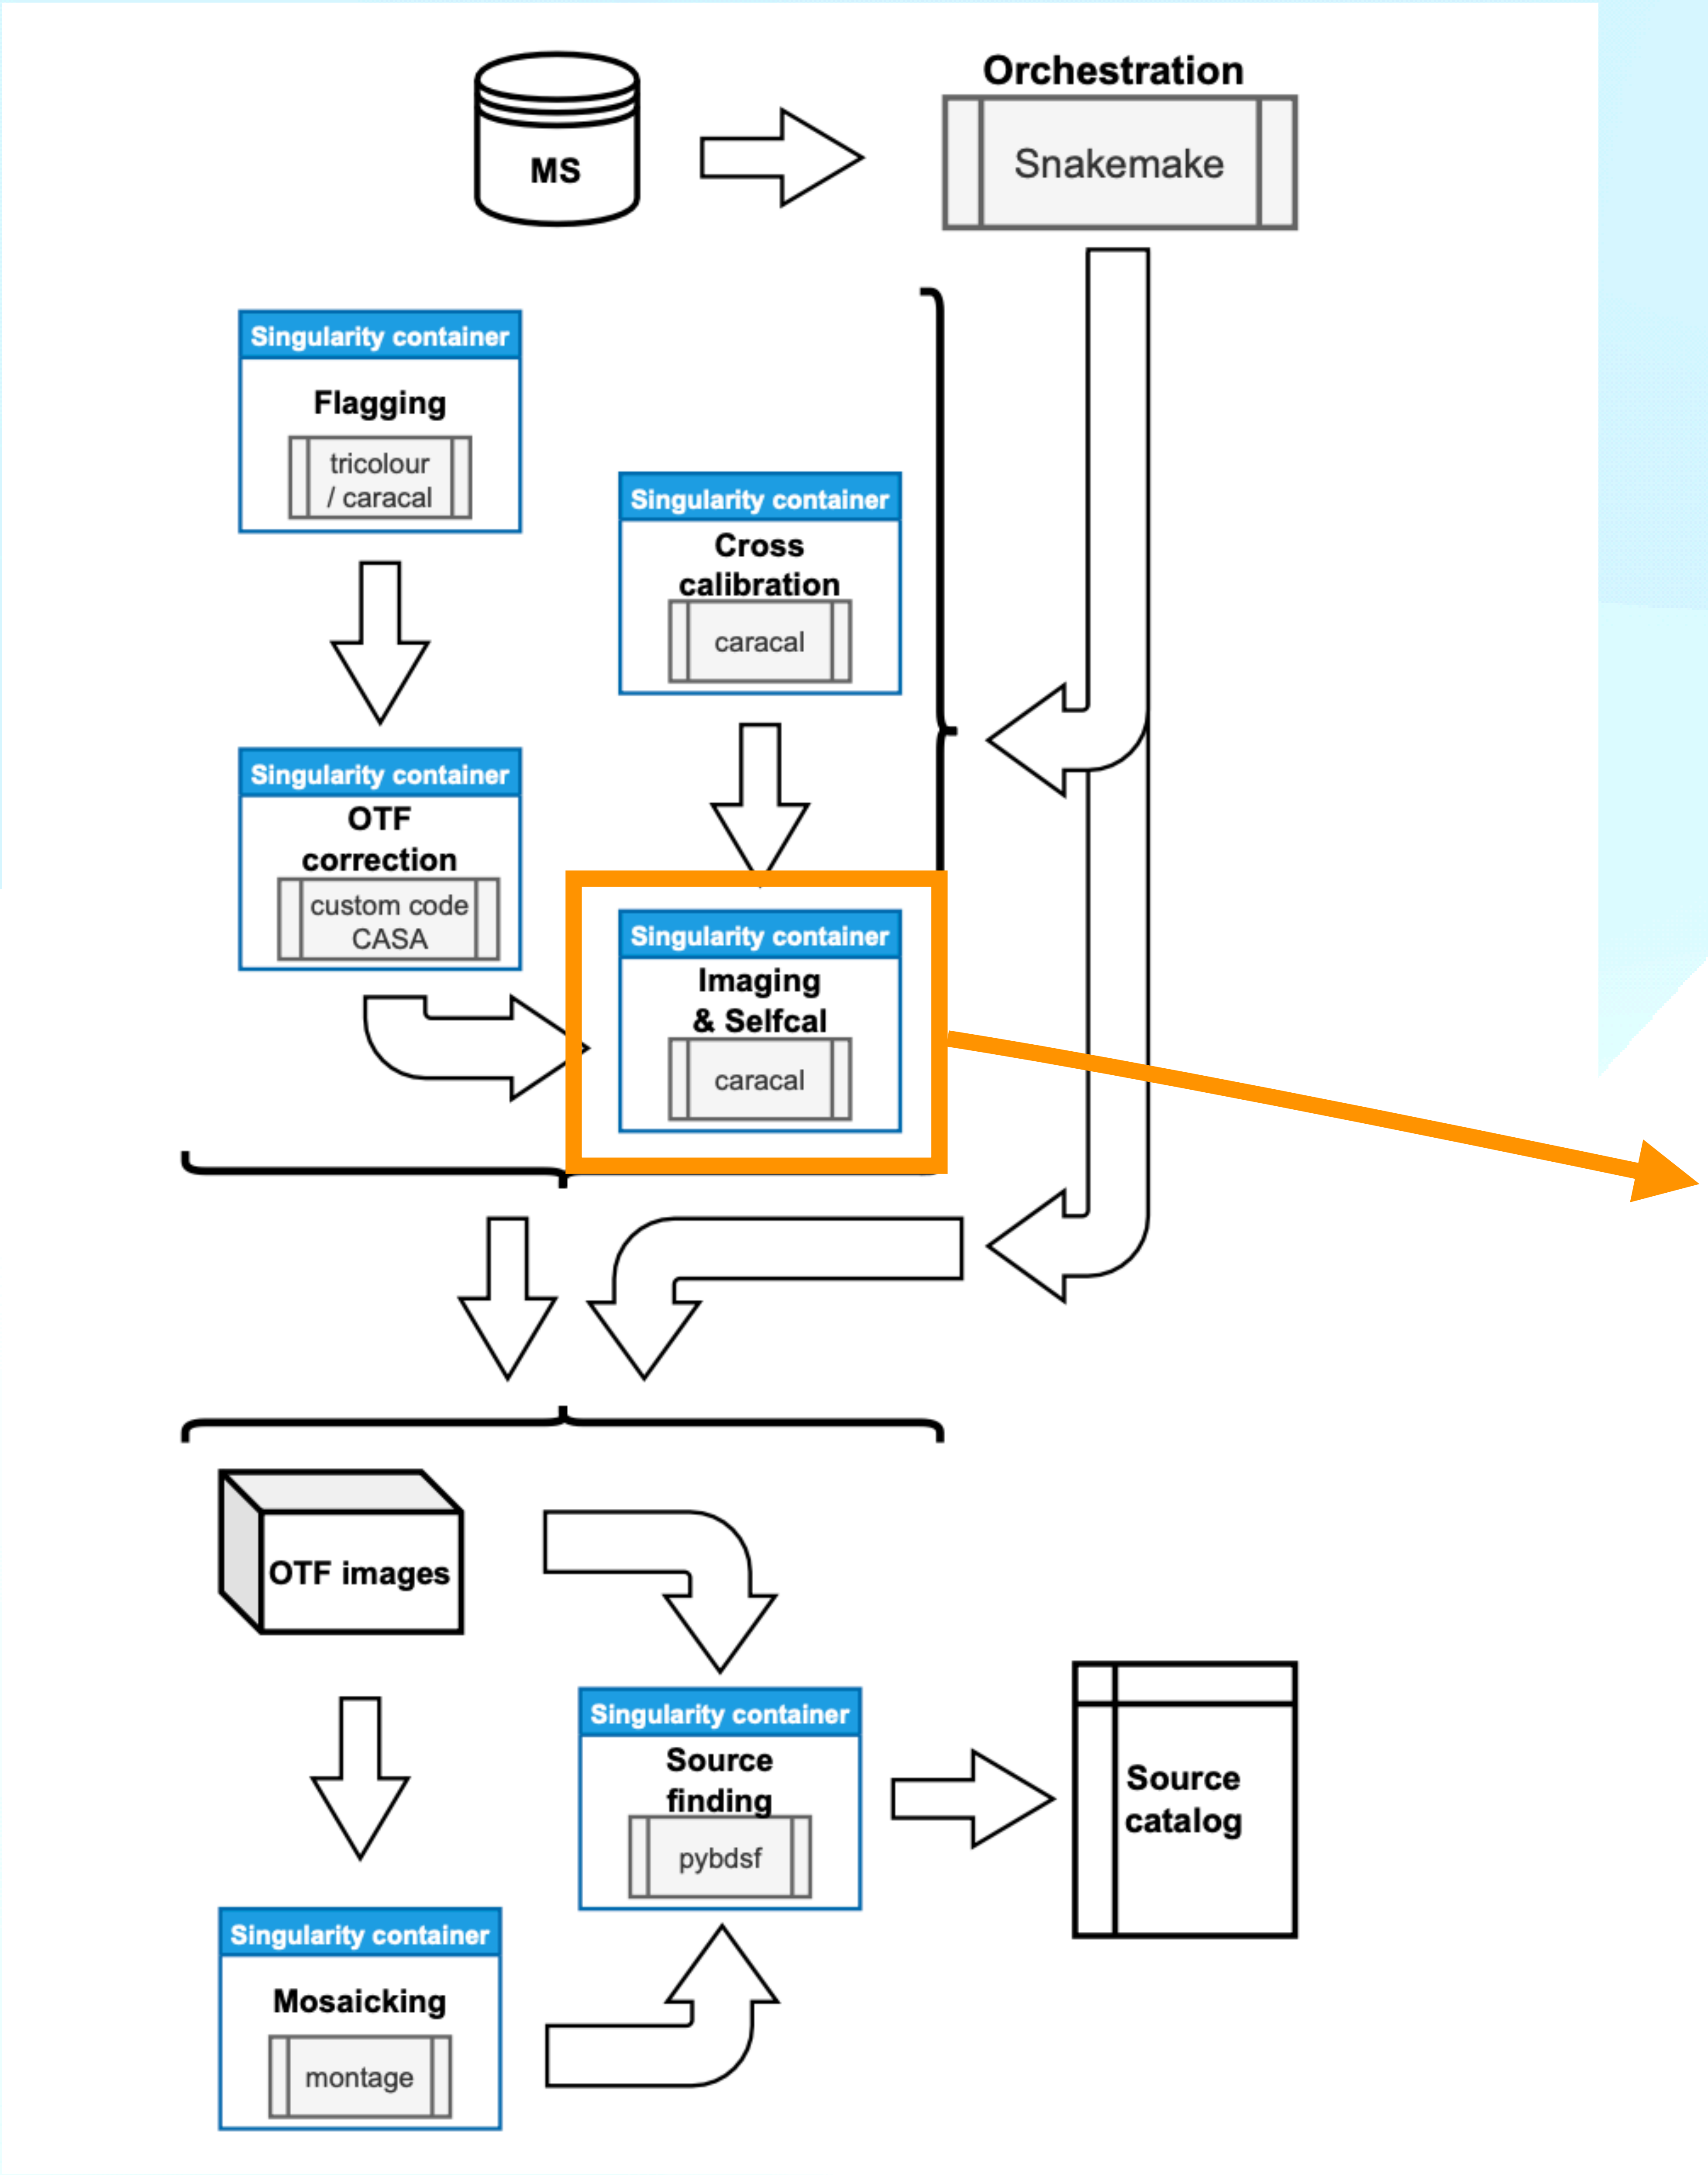

Biggest challenge faced so far: **how to use pre-containerized software**?

Imaging portion needs the software caracal.

Achievements - Imaging

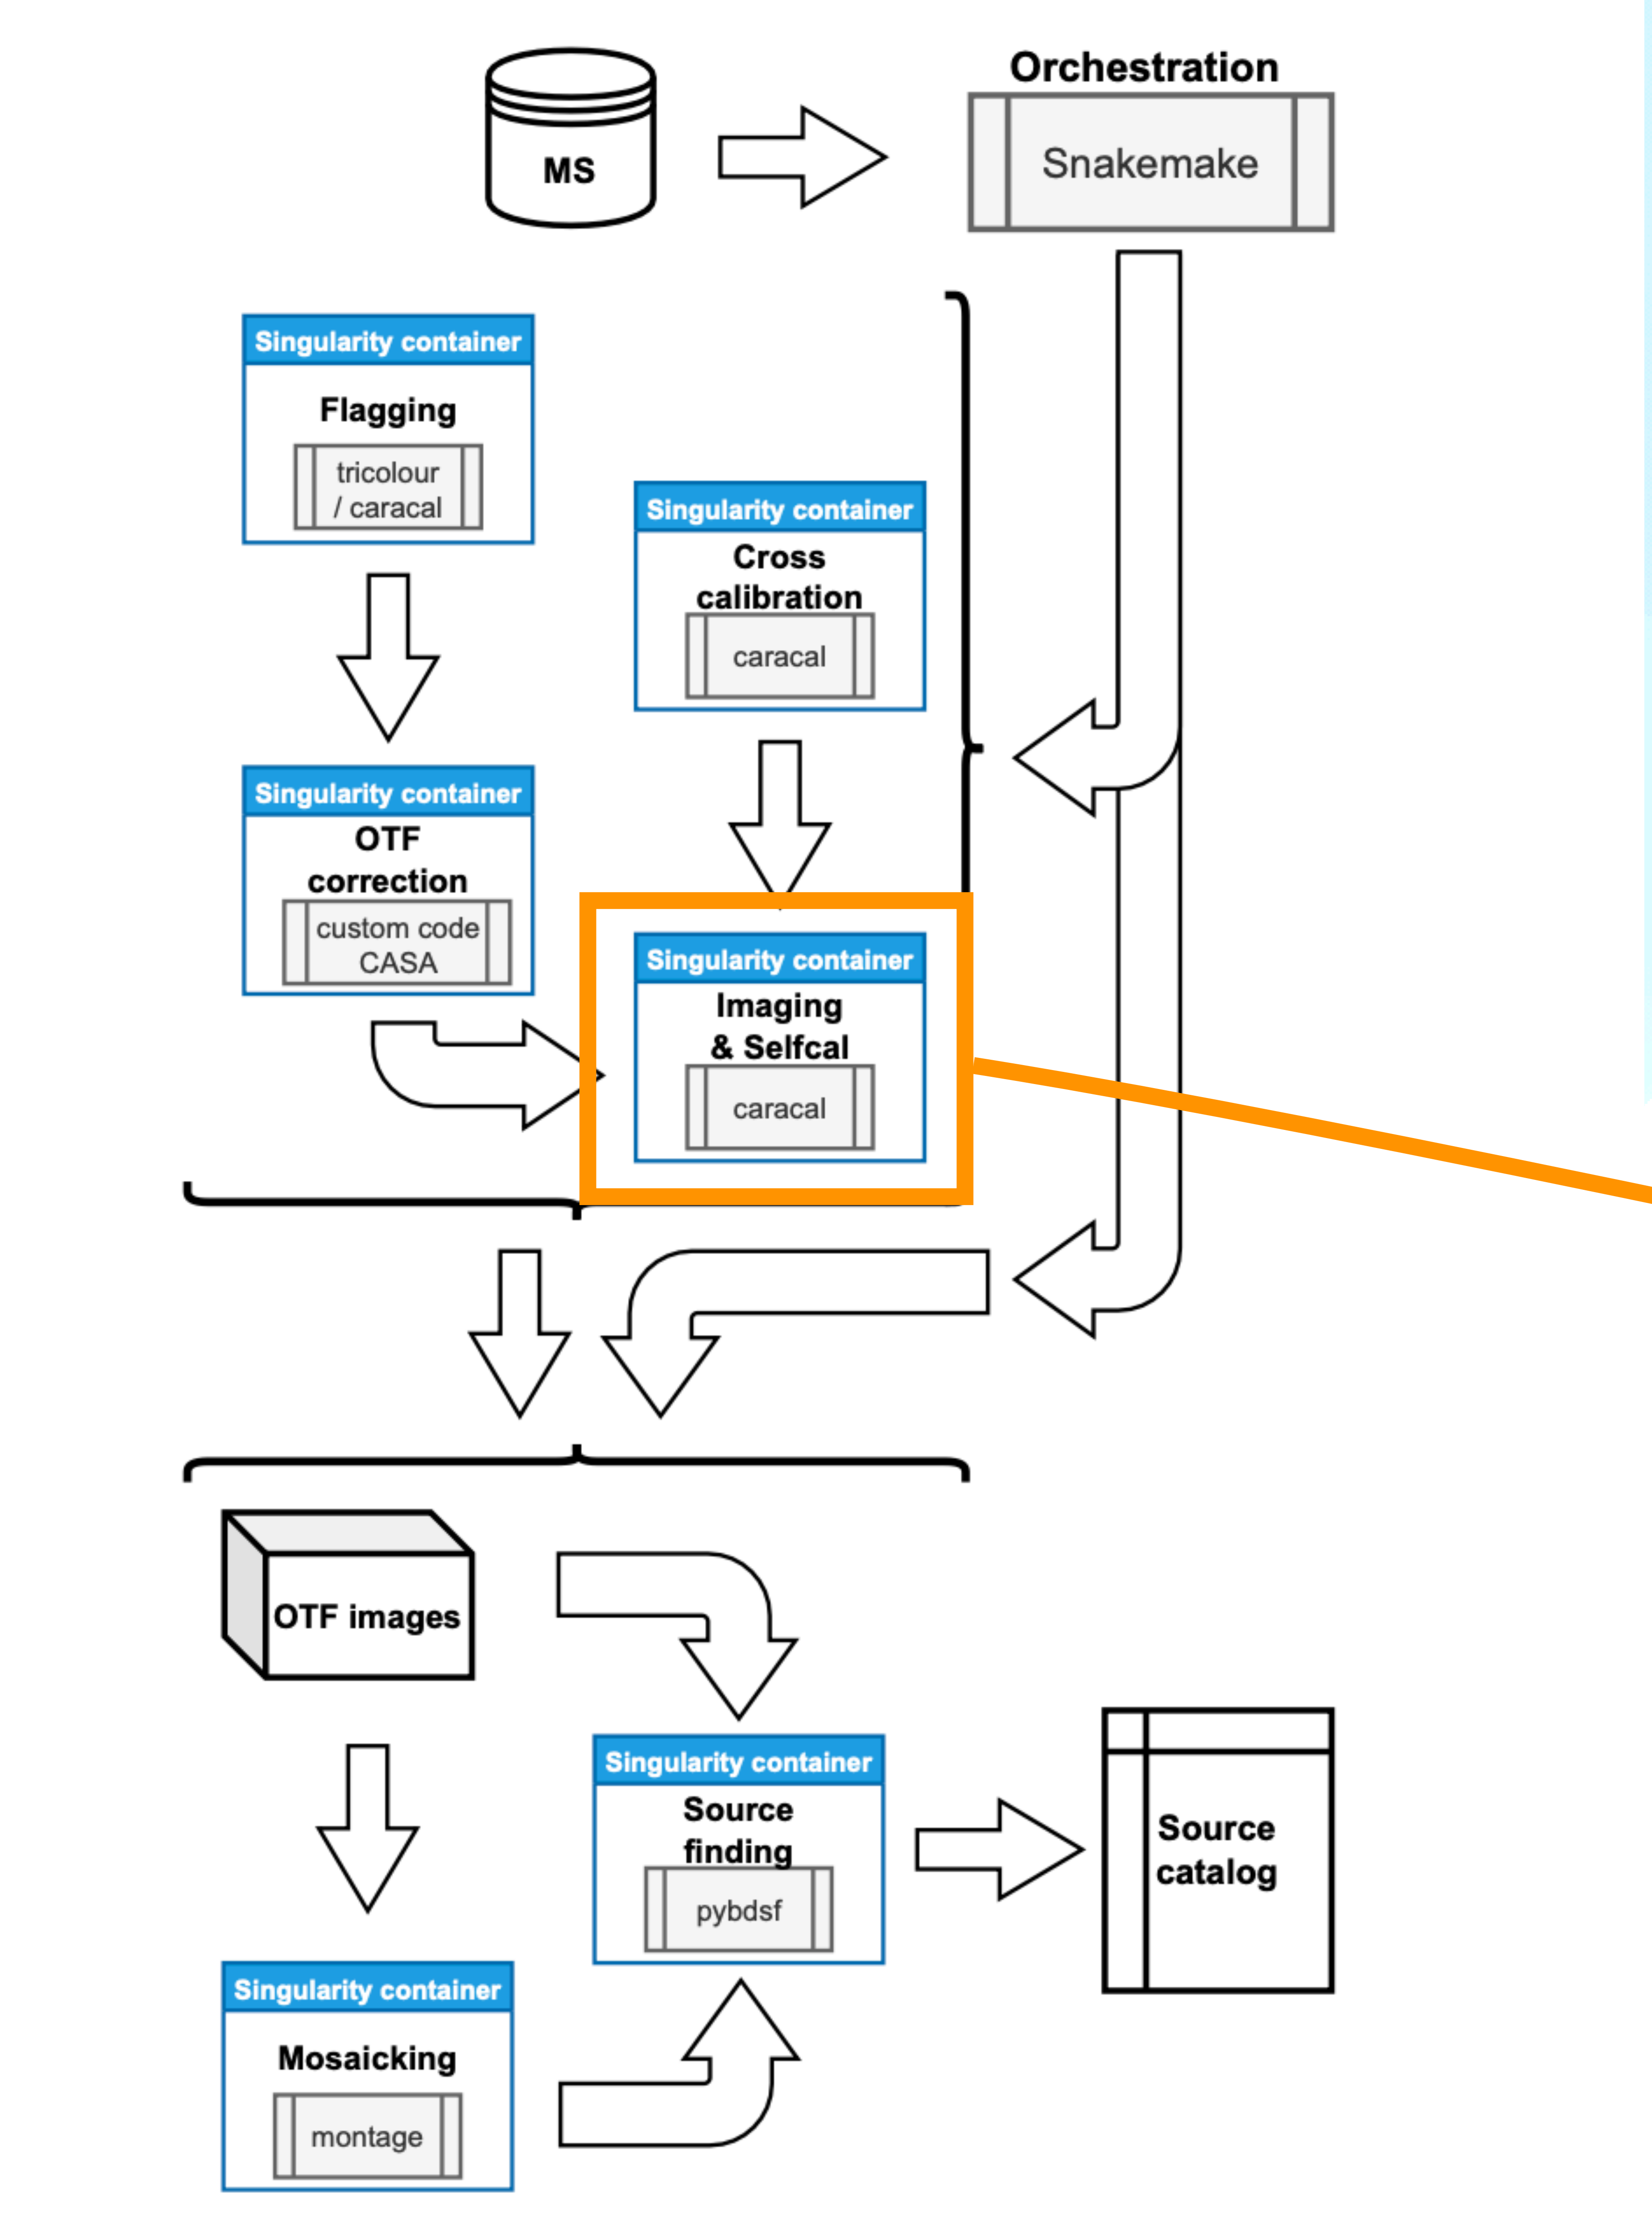

Biggest challenge faced so far: **how to use pre-containerized software**?

Imaging portion needs the software caracal.

- To be run on Compute4PUNCH caracal needs to be containerized.

Achievements - Imaging

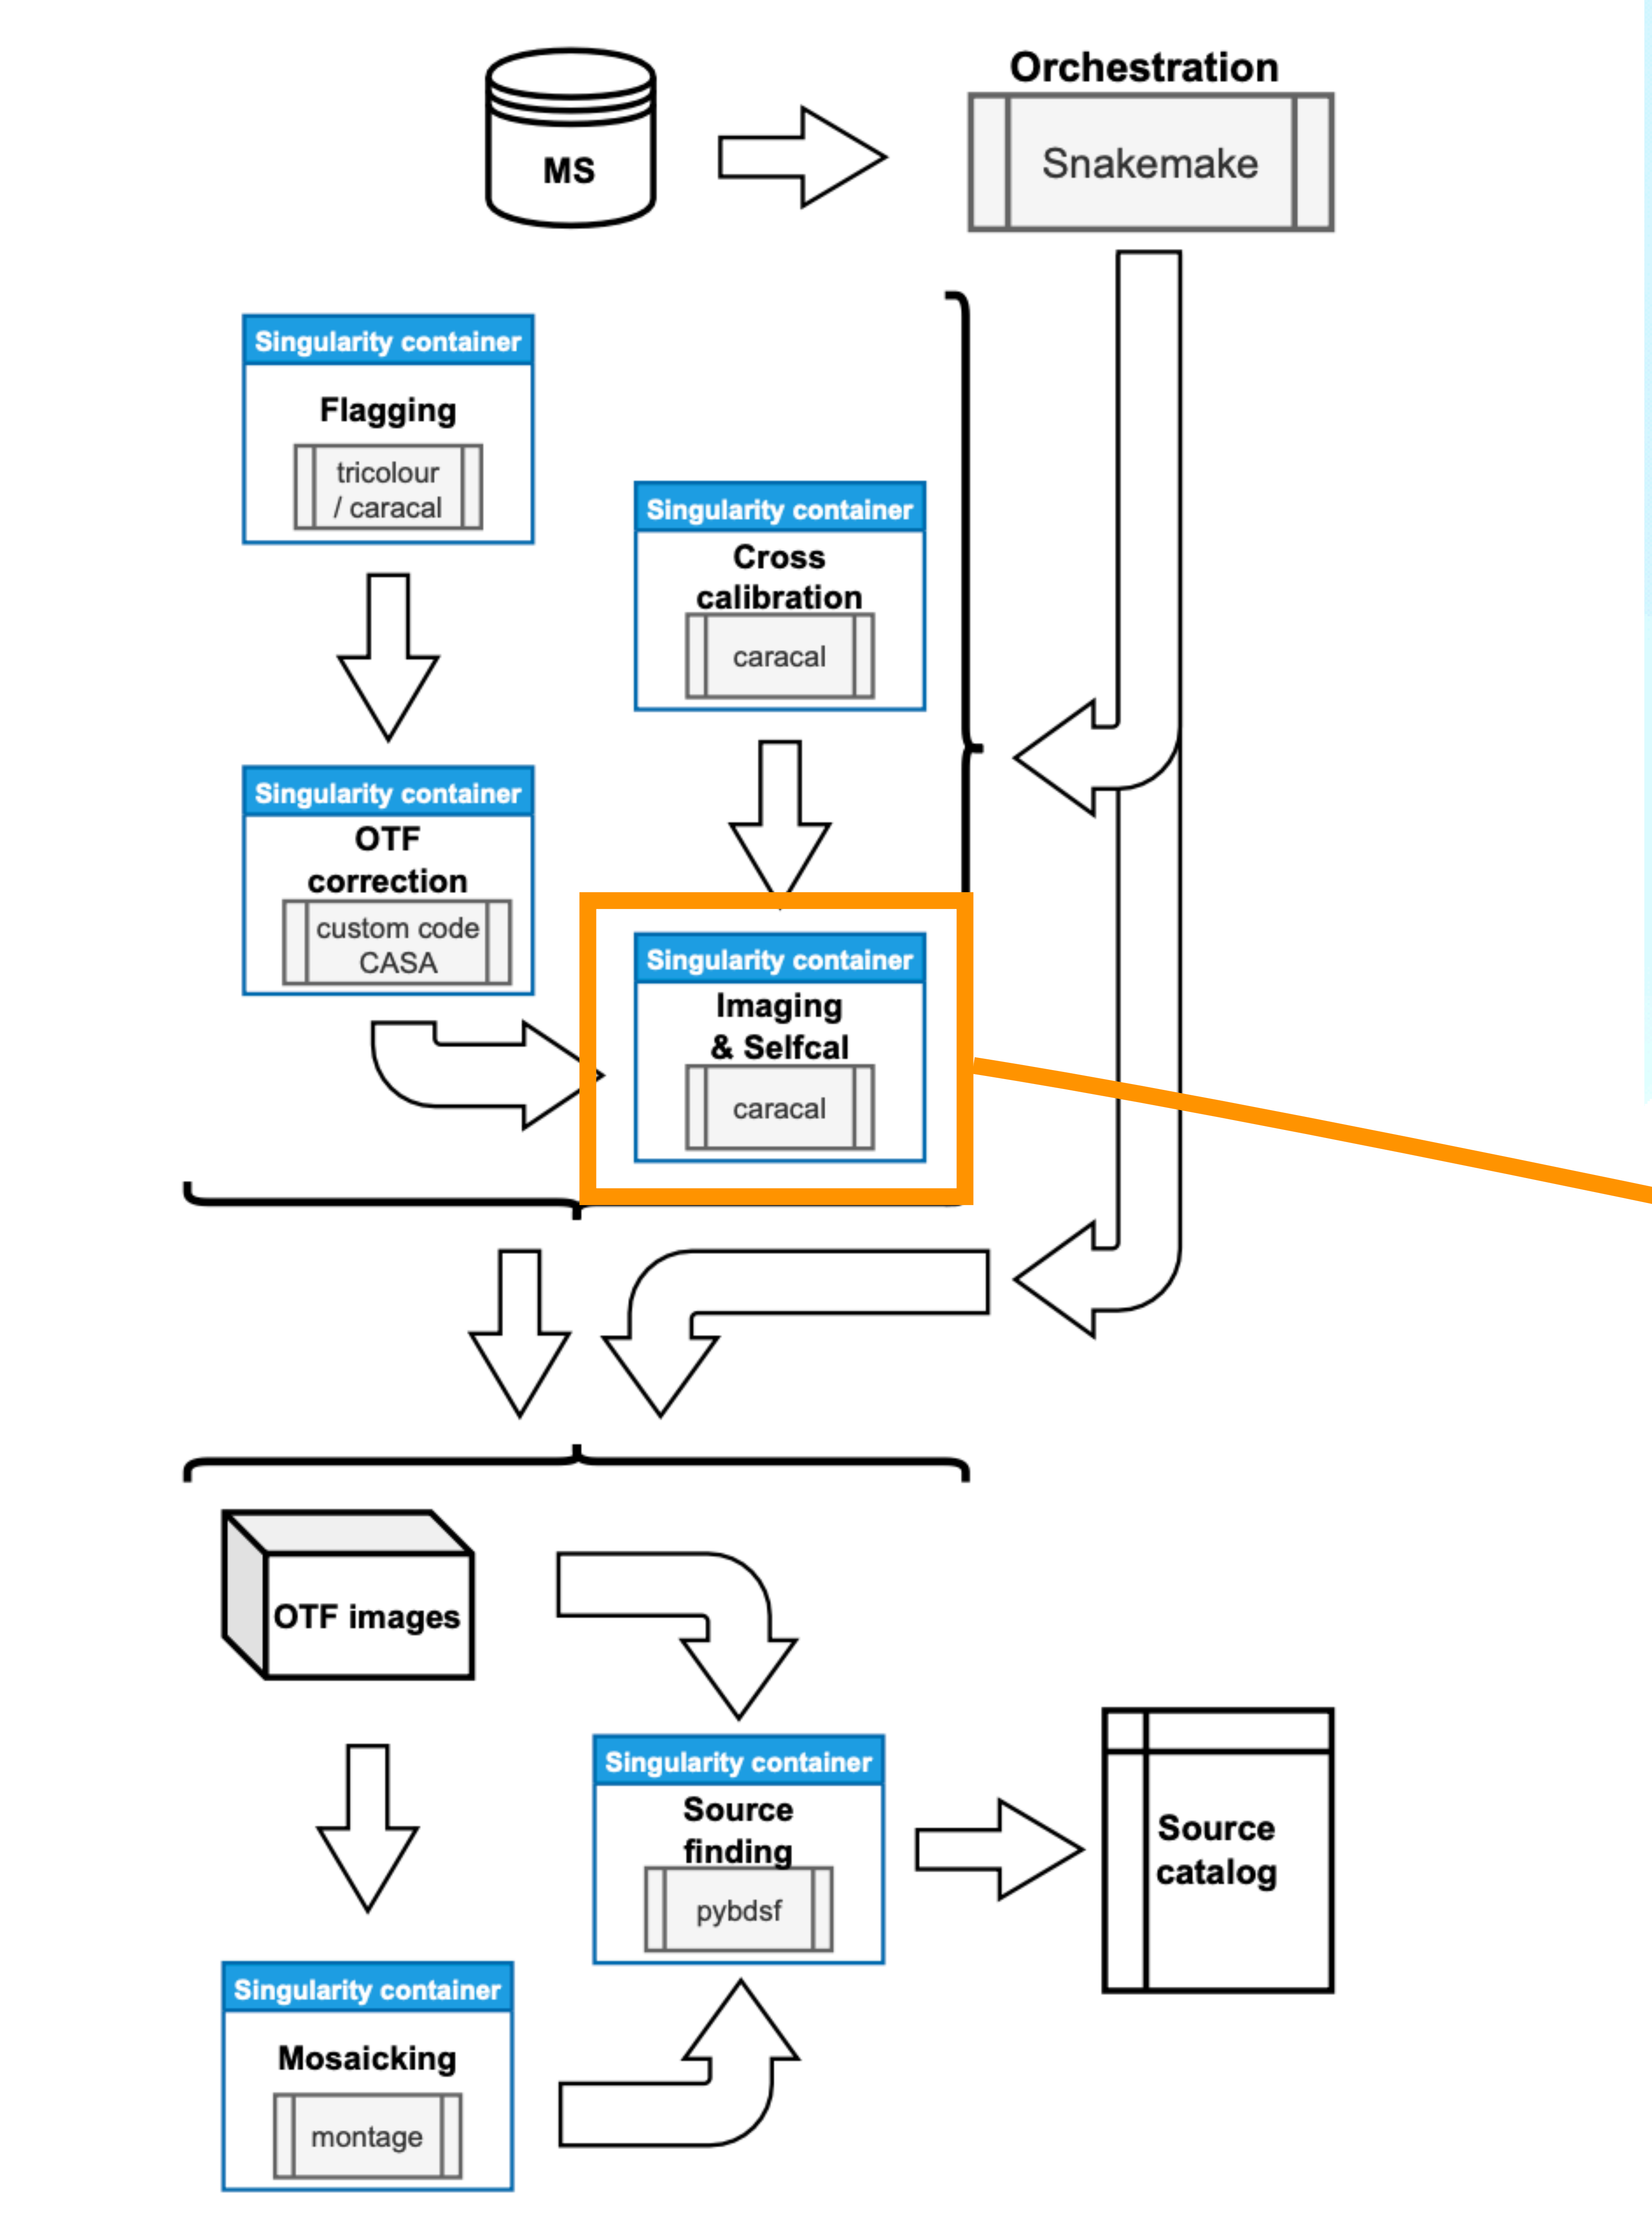

Biggest challenge faced so far: **how to use pre-containerized software**?

Imaging portion needs the software caracal.

- To be run on Compute4PUNCH caracal needs to be containerized.
- PROBLEM: caracal uses stimela which provides containerized software for specific tasks (called “cabs”).

Achievements - Imaging

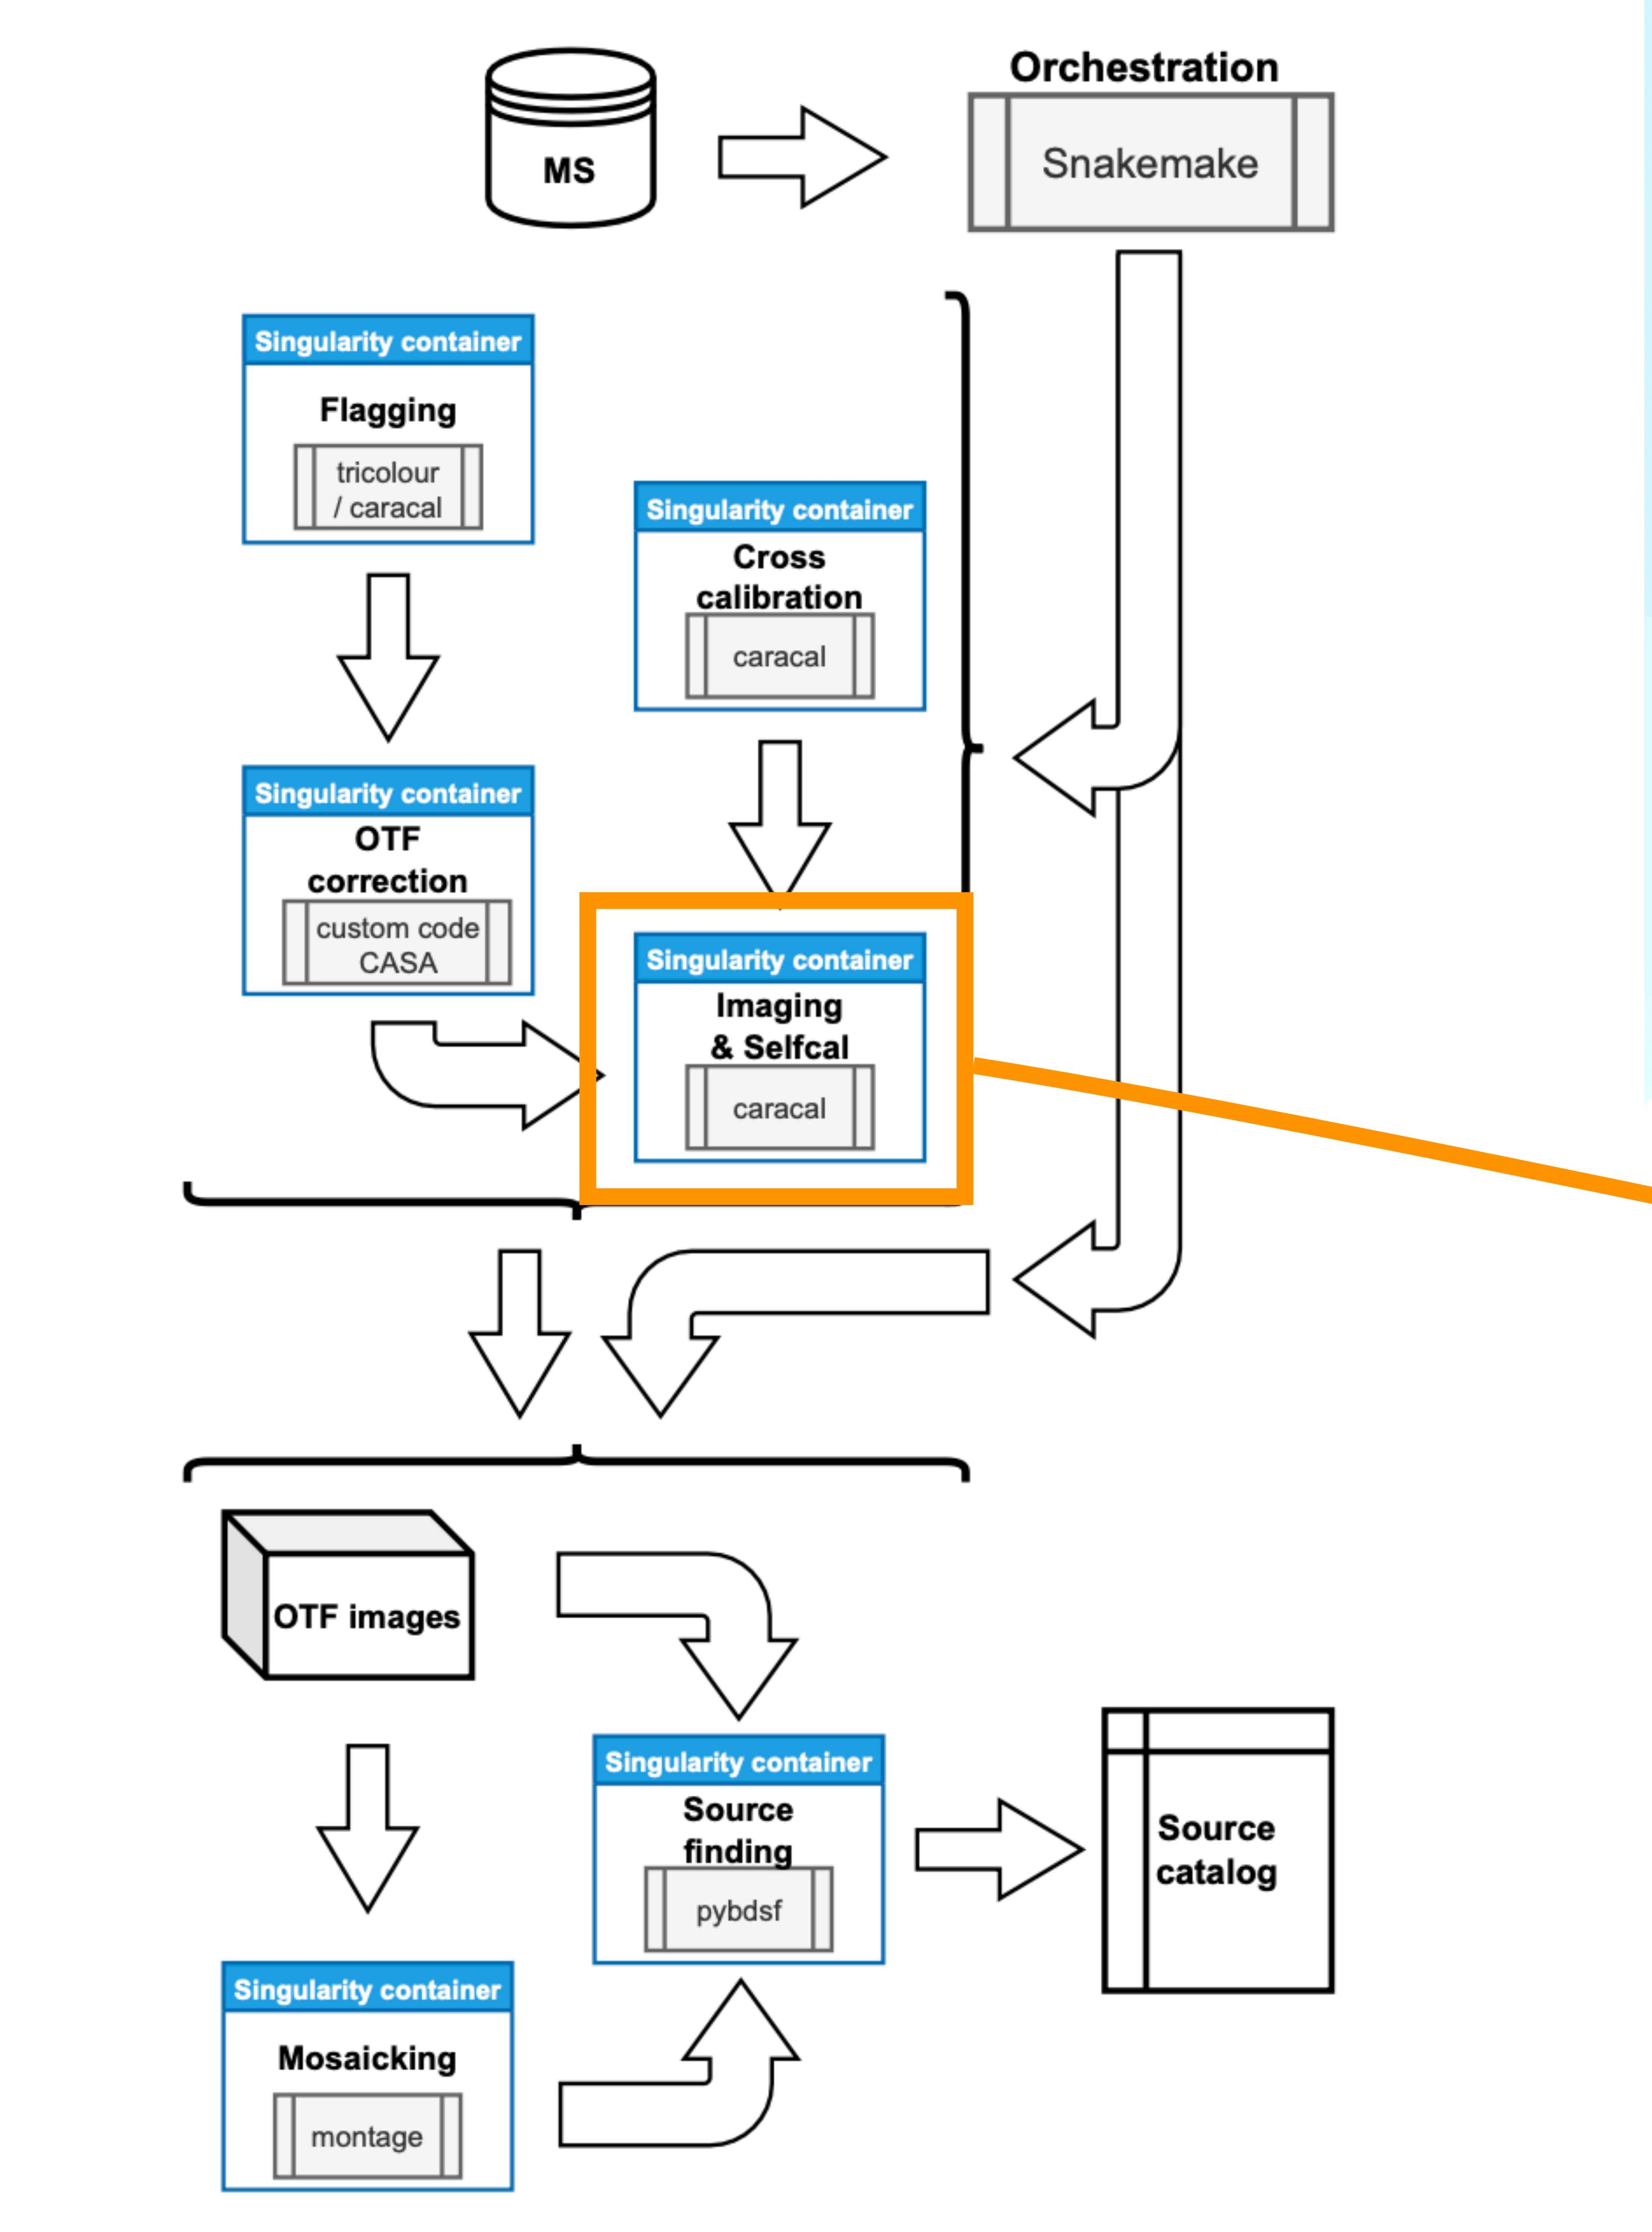

Biggest challenge faced so far: **how to use pre-containerized software**?

Imaging portion needs the software caracal.

- To be run on Compute4PUNCH caracal needs to be containerized.
- PROBLEM: caracal uses stimela which provides containerized software for specific tasks (called “cabs”).
- Need of a “container in container” solution.

Achievements - Imaging

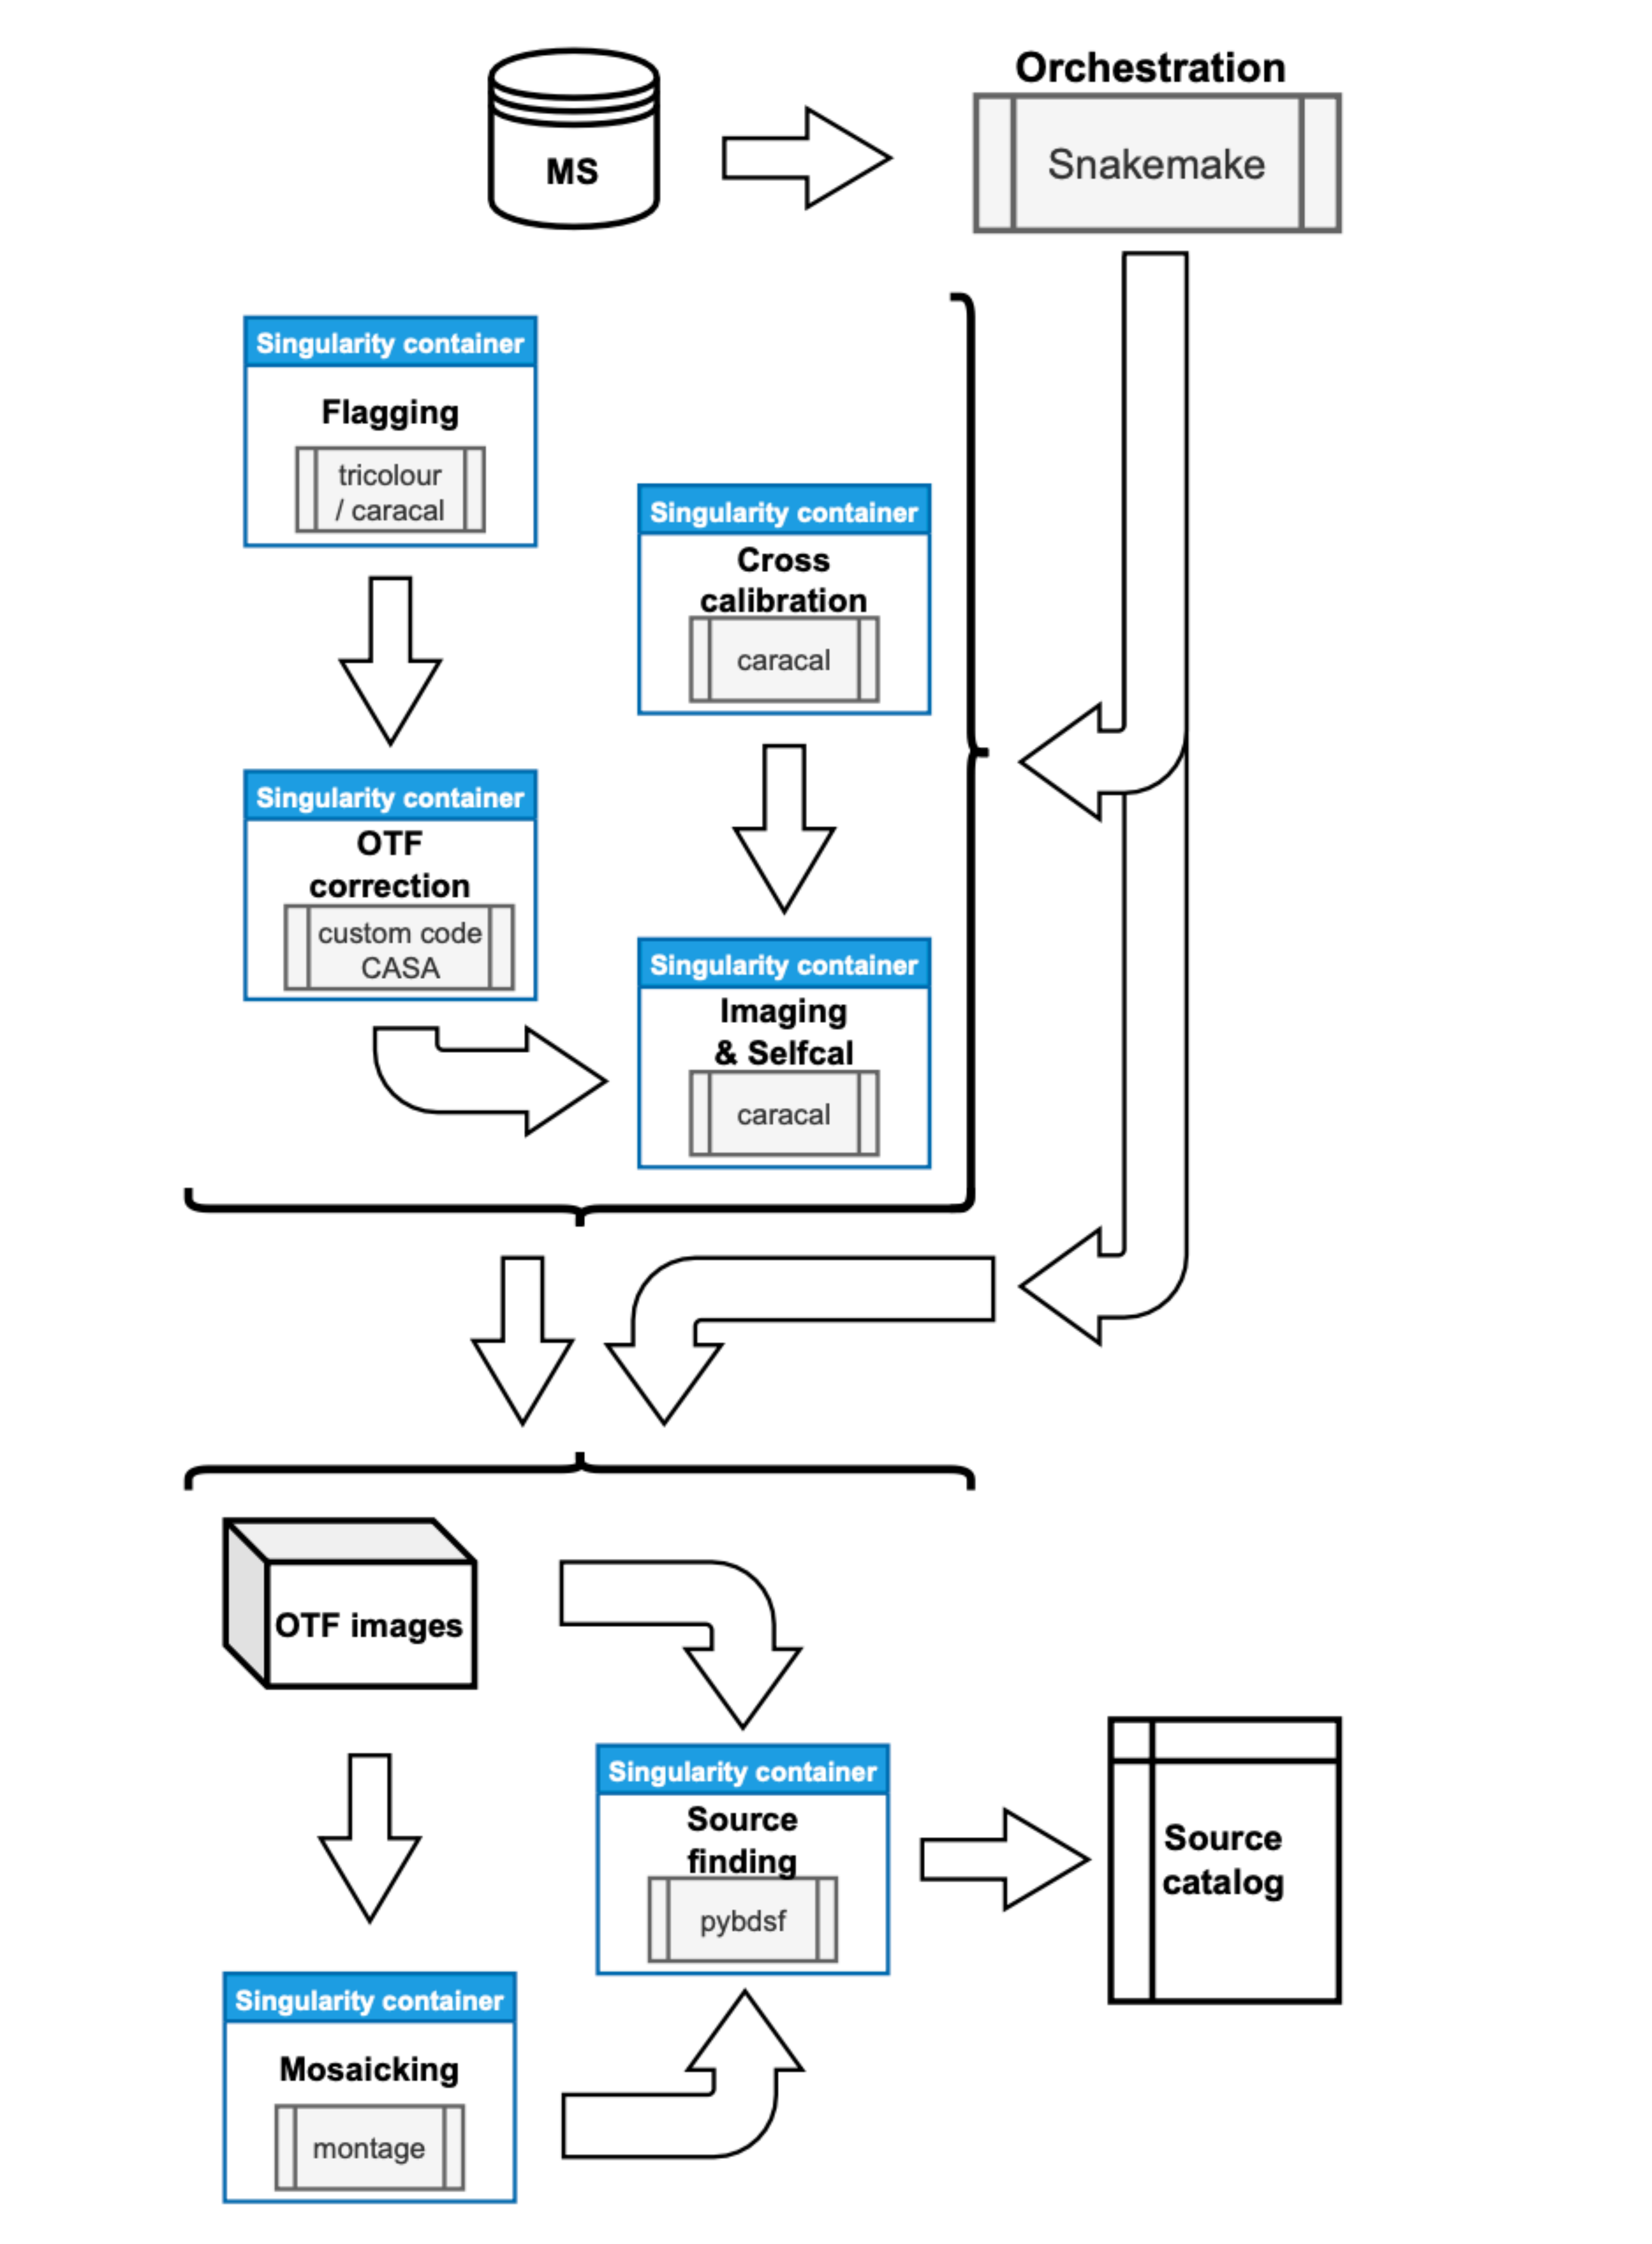

Our solution: workaround for container (cabs) access and modification of the caracal software.

Achievements - Imaging

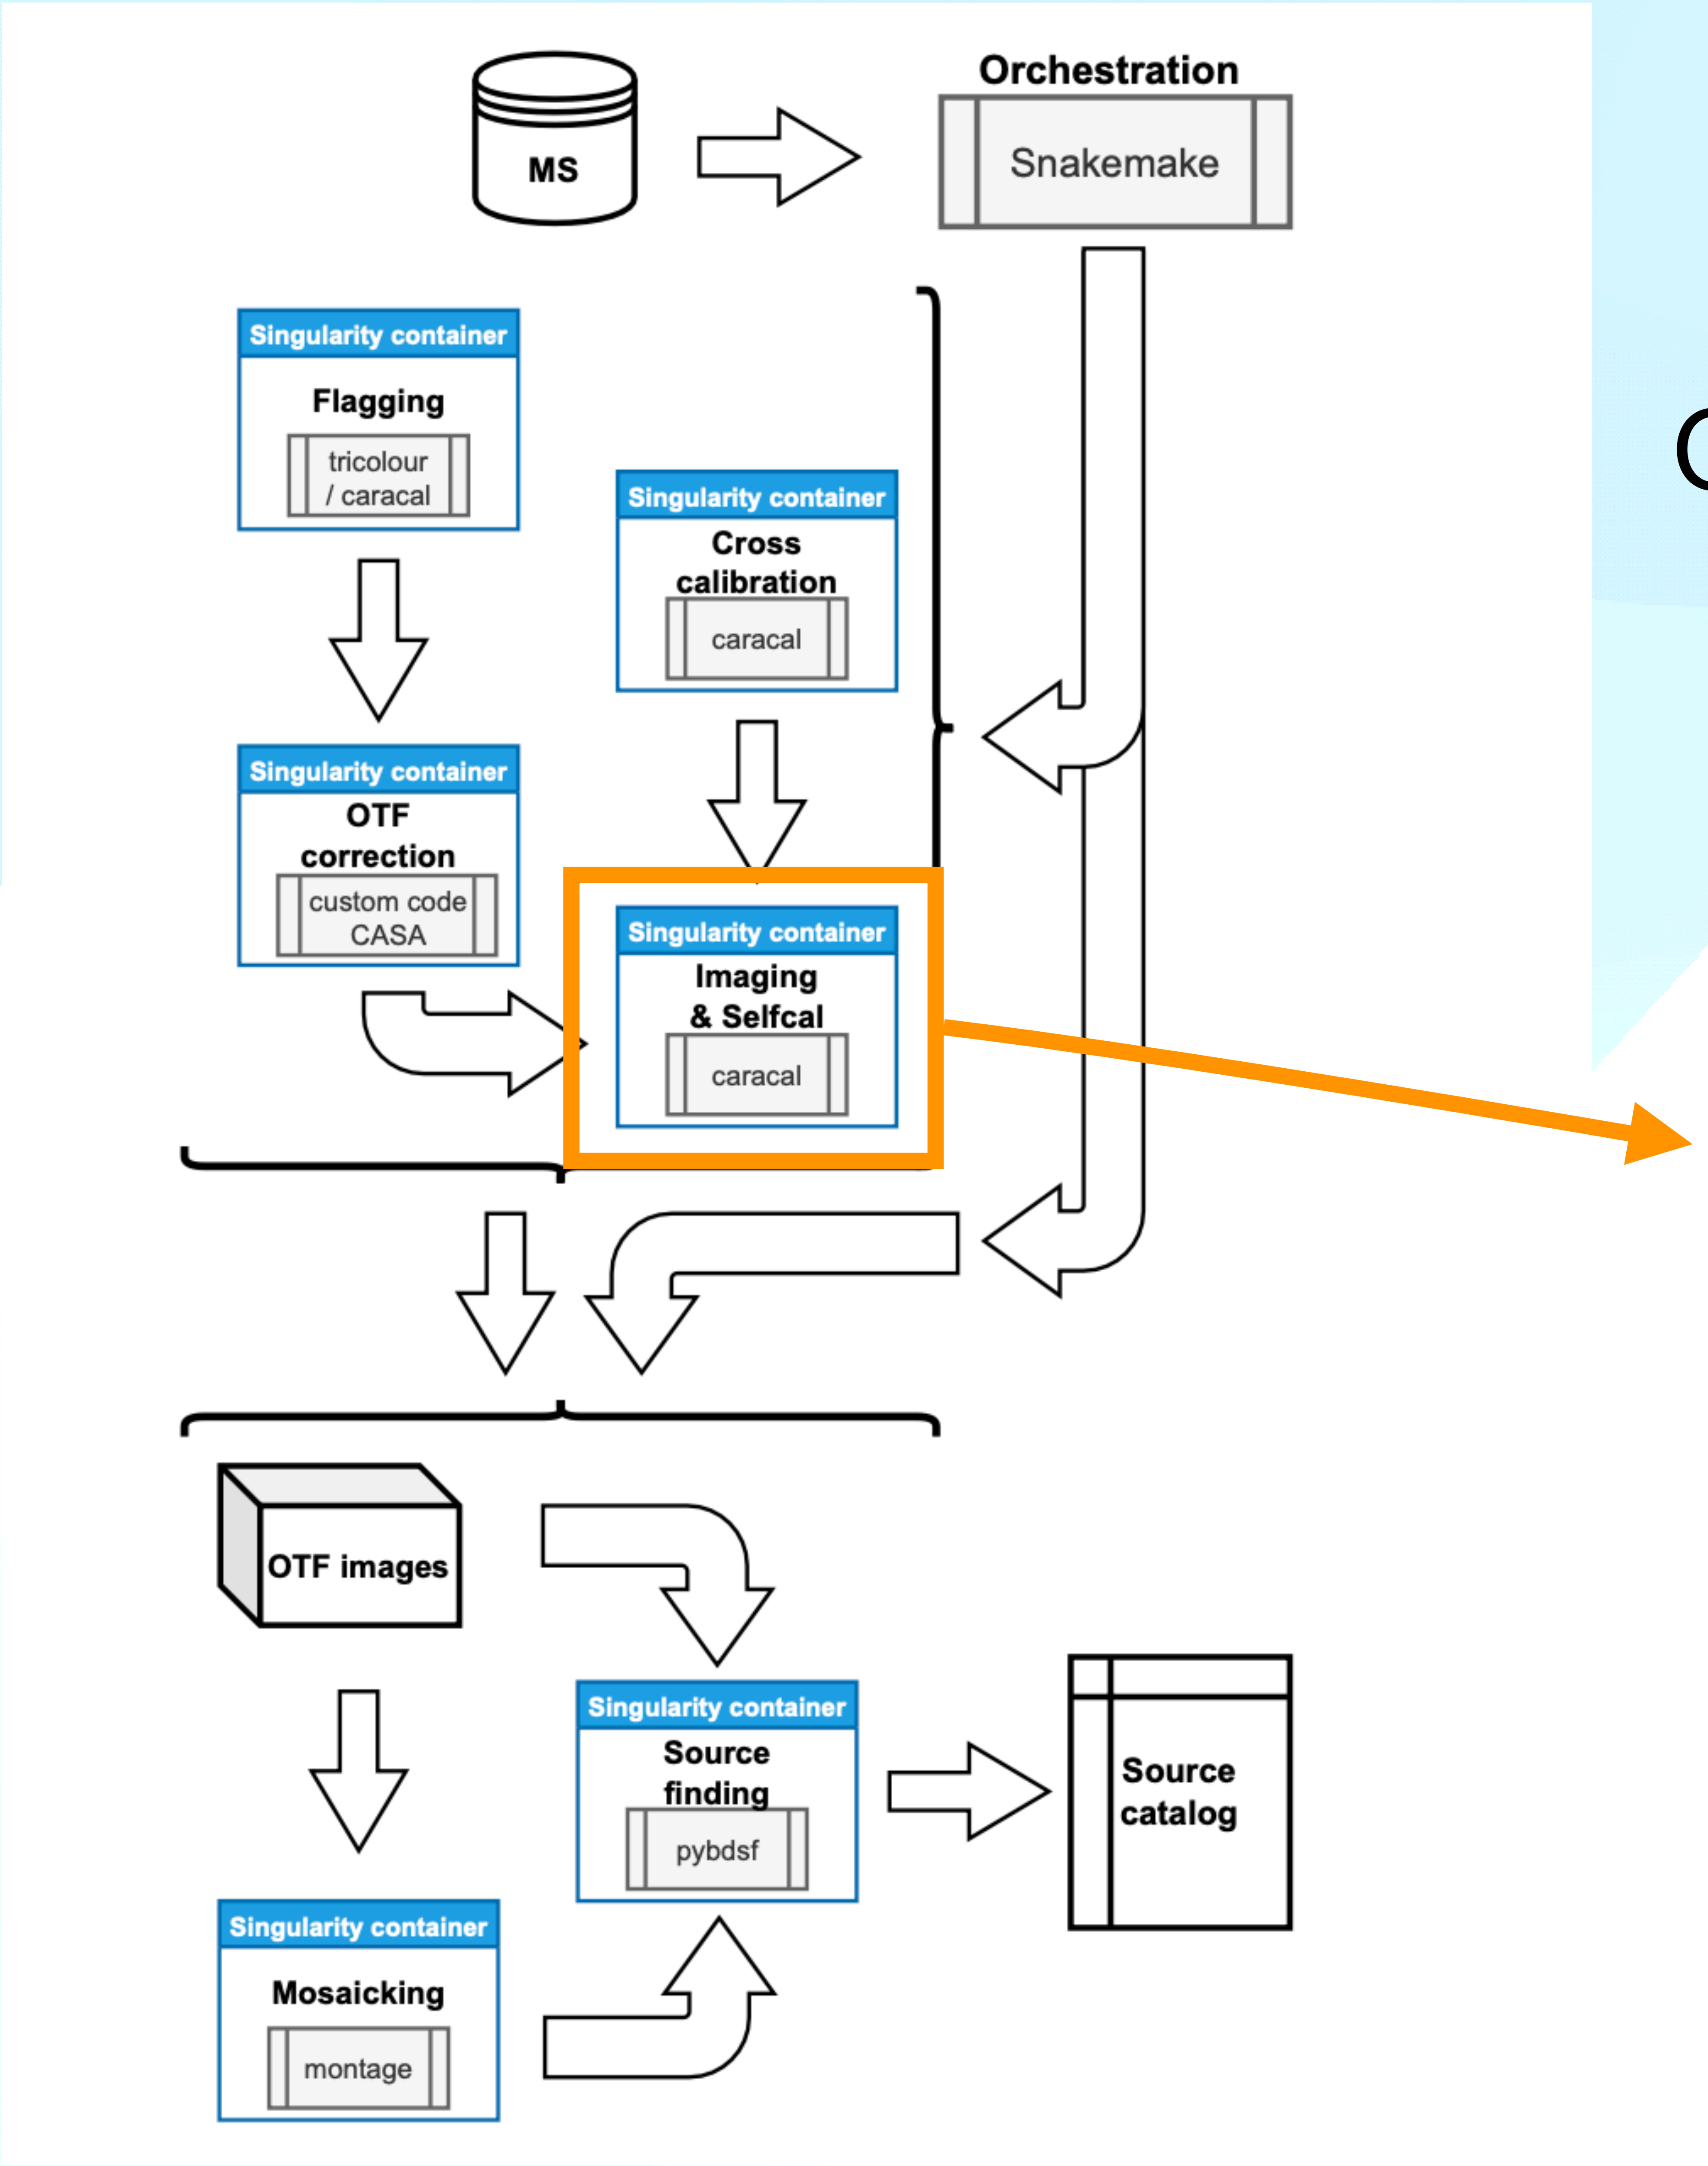

Our solution: workaround for container (cabs) access and modification of the caracal software.

Achievements - Imaging

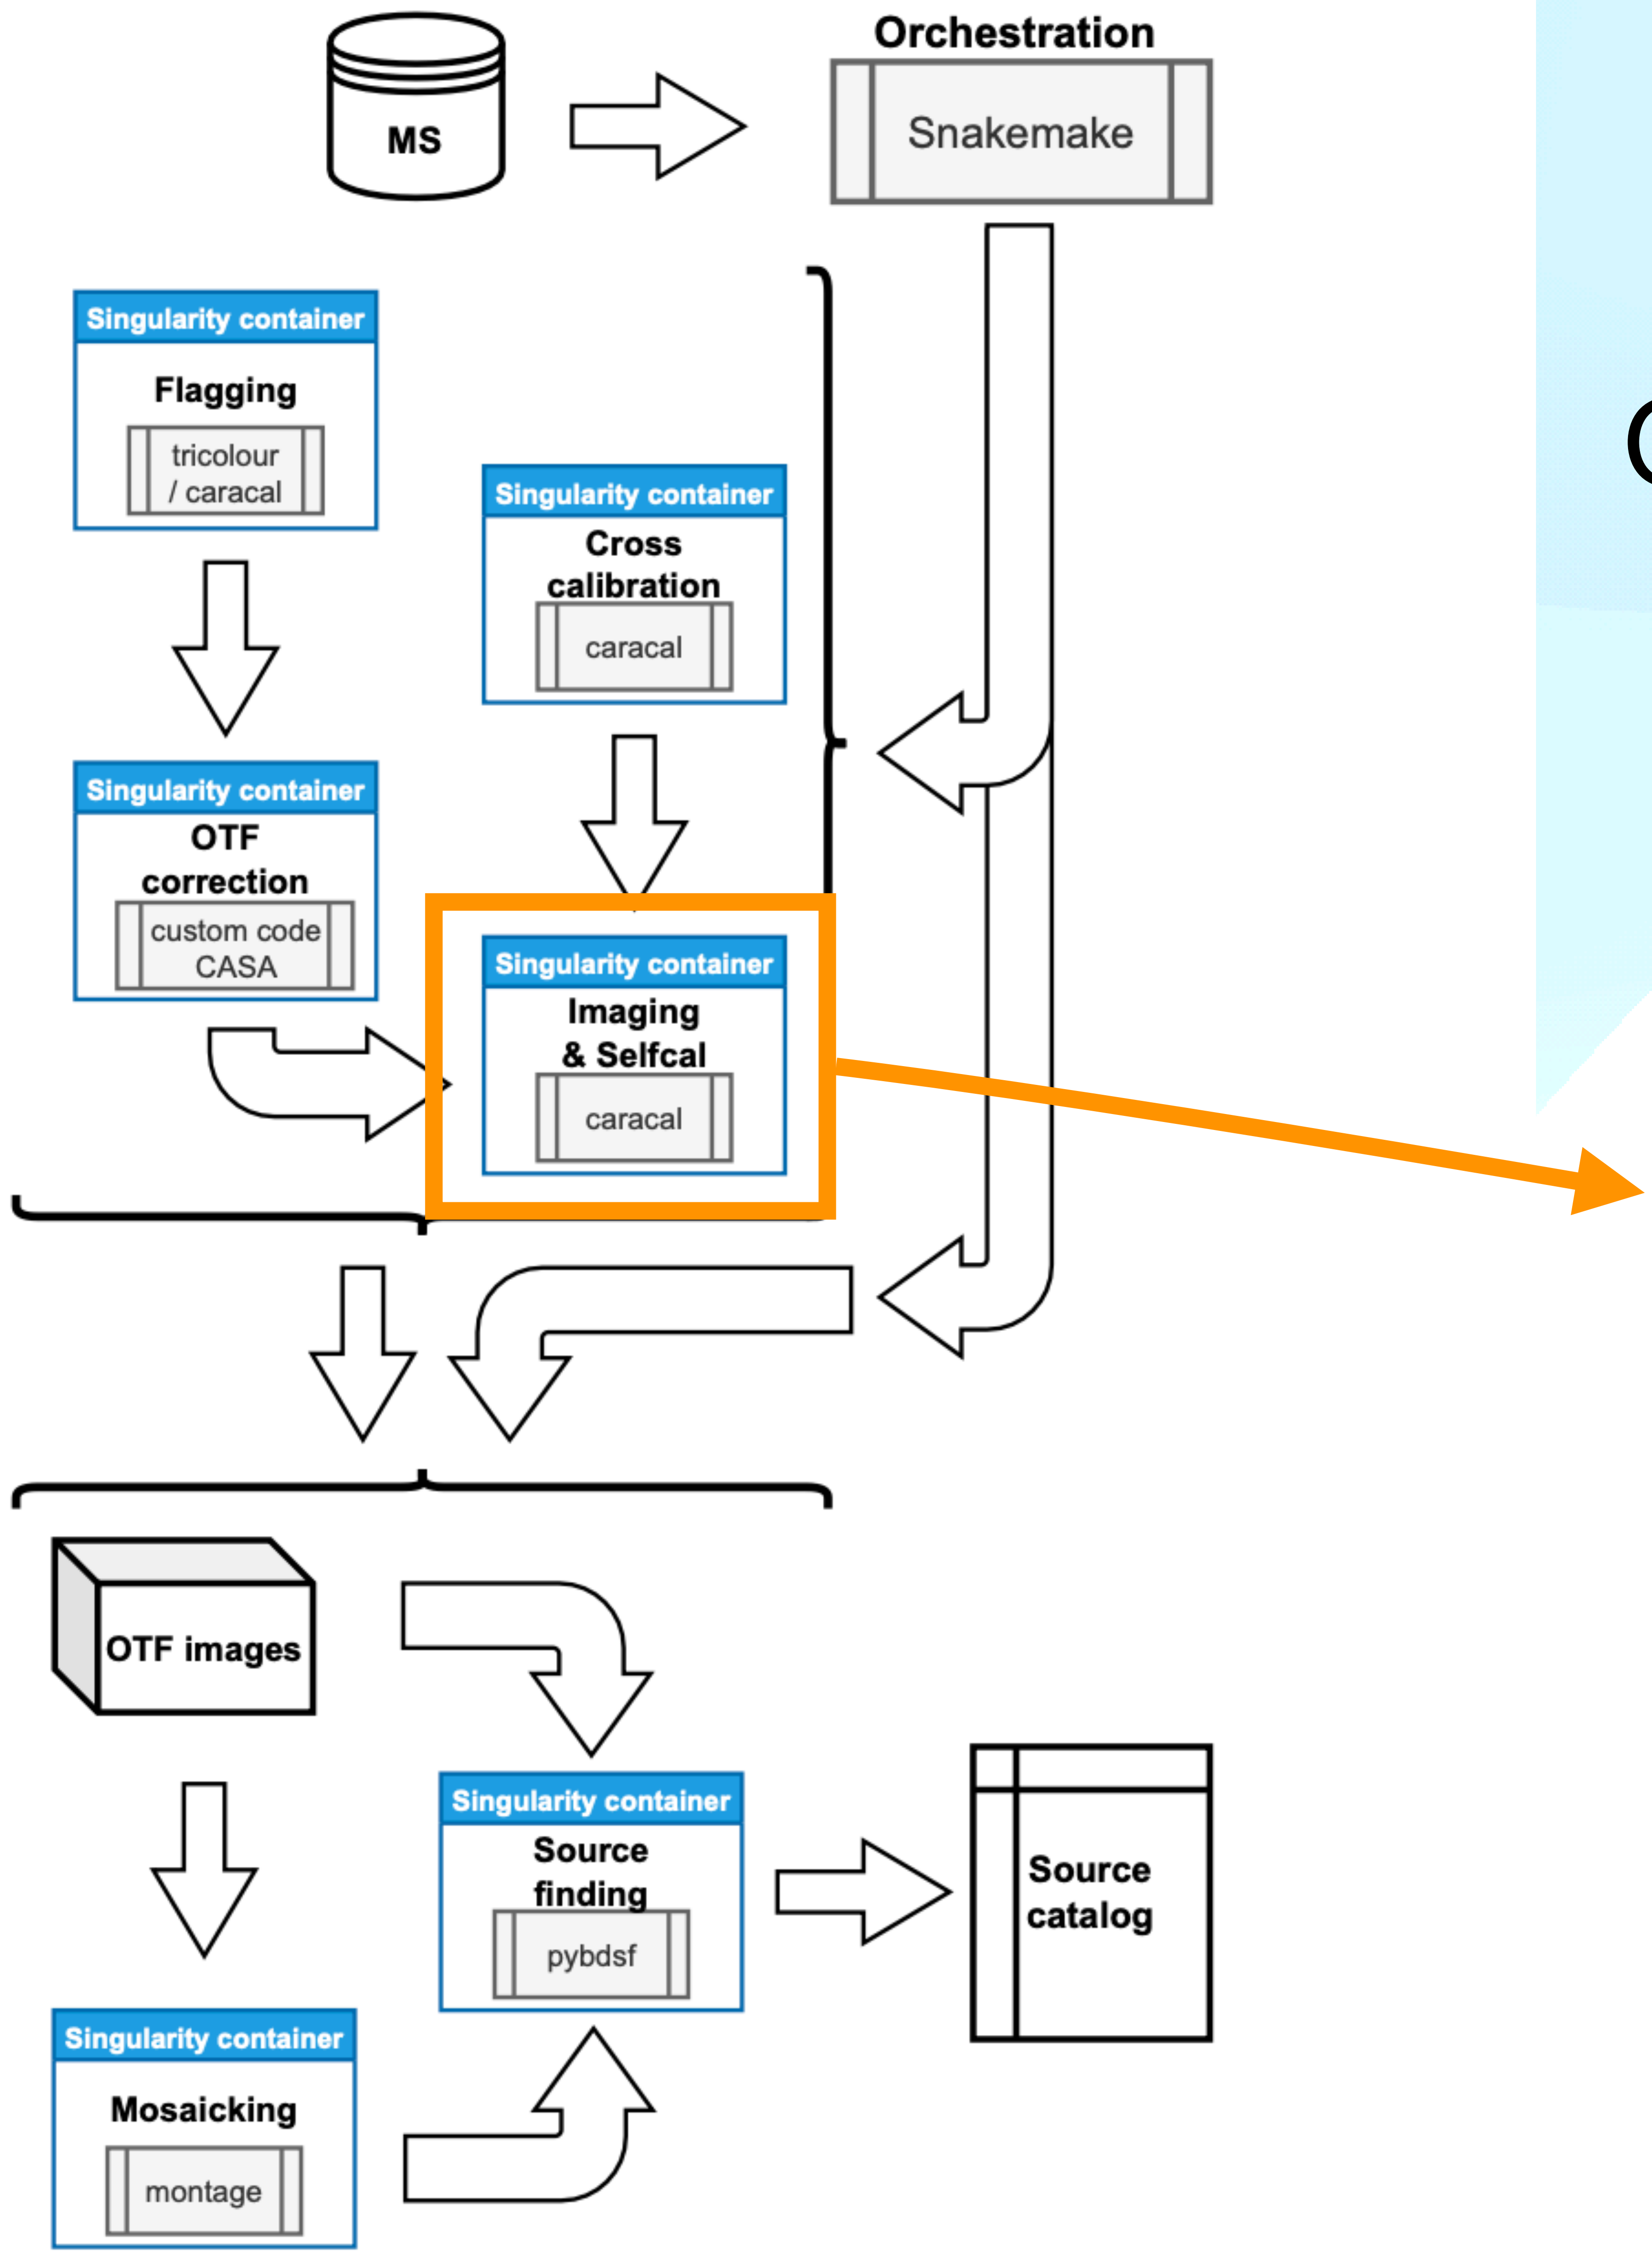

Our solution: workaround for container (cabs) access and modification of the caracal software.

- Caracal can use stimela to access containers in a local folder instead of docker registry.

Achievements - Imaging

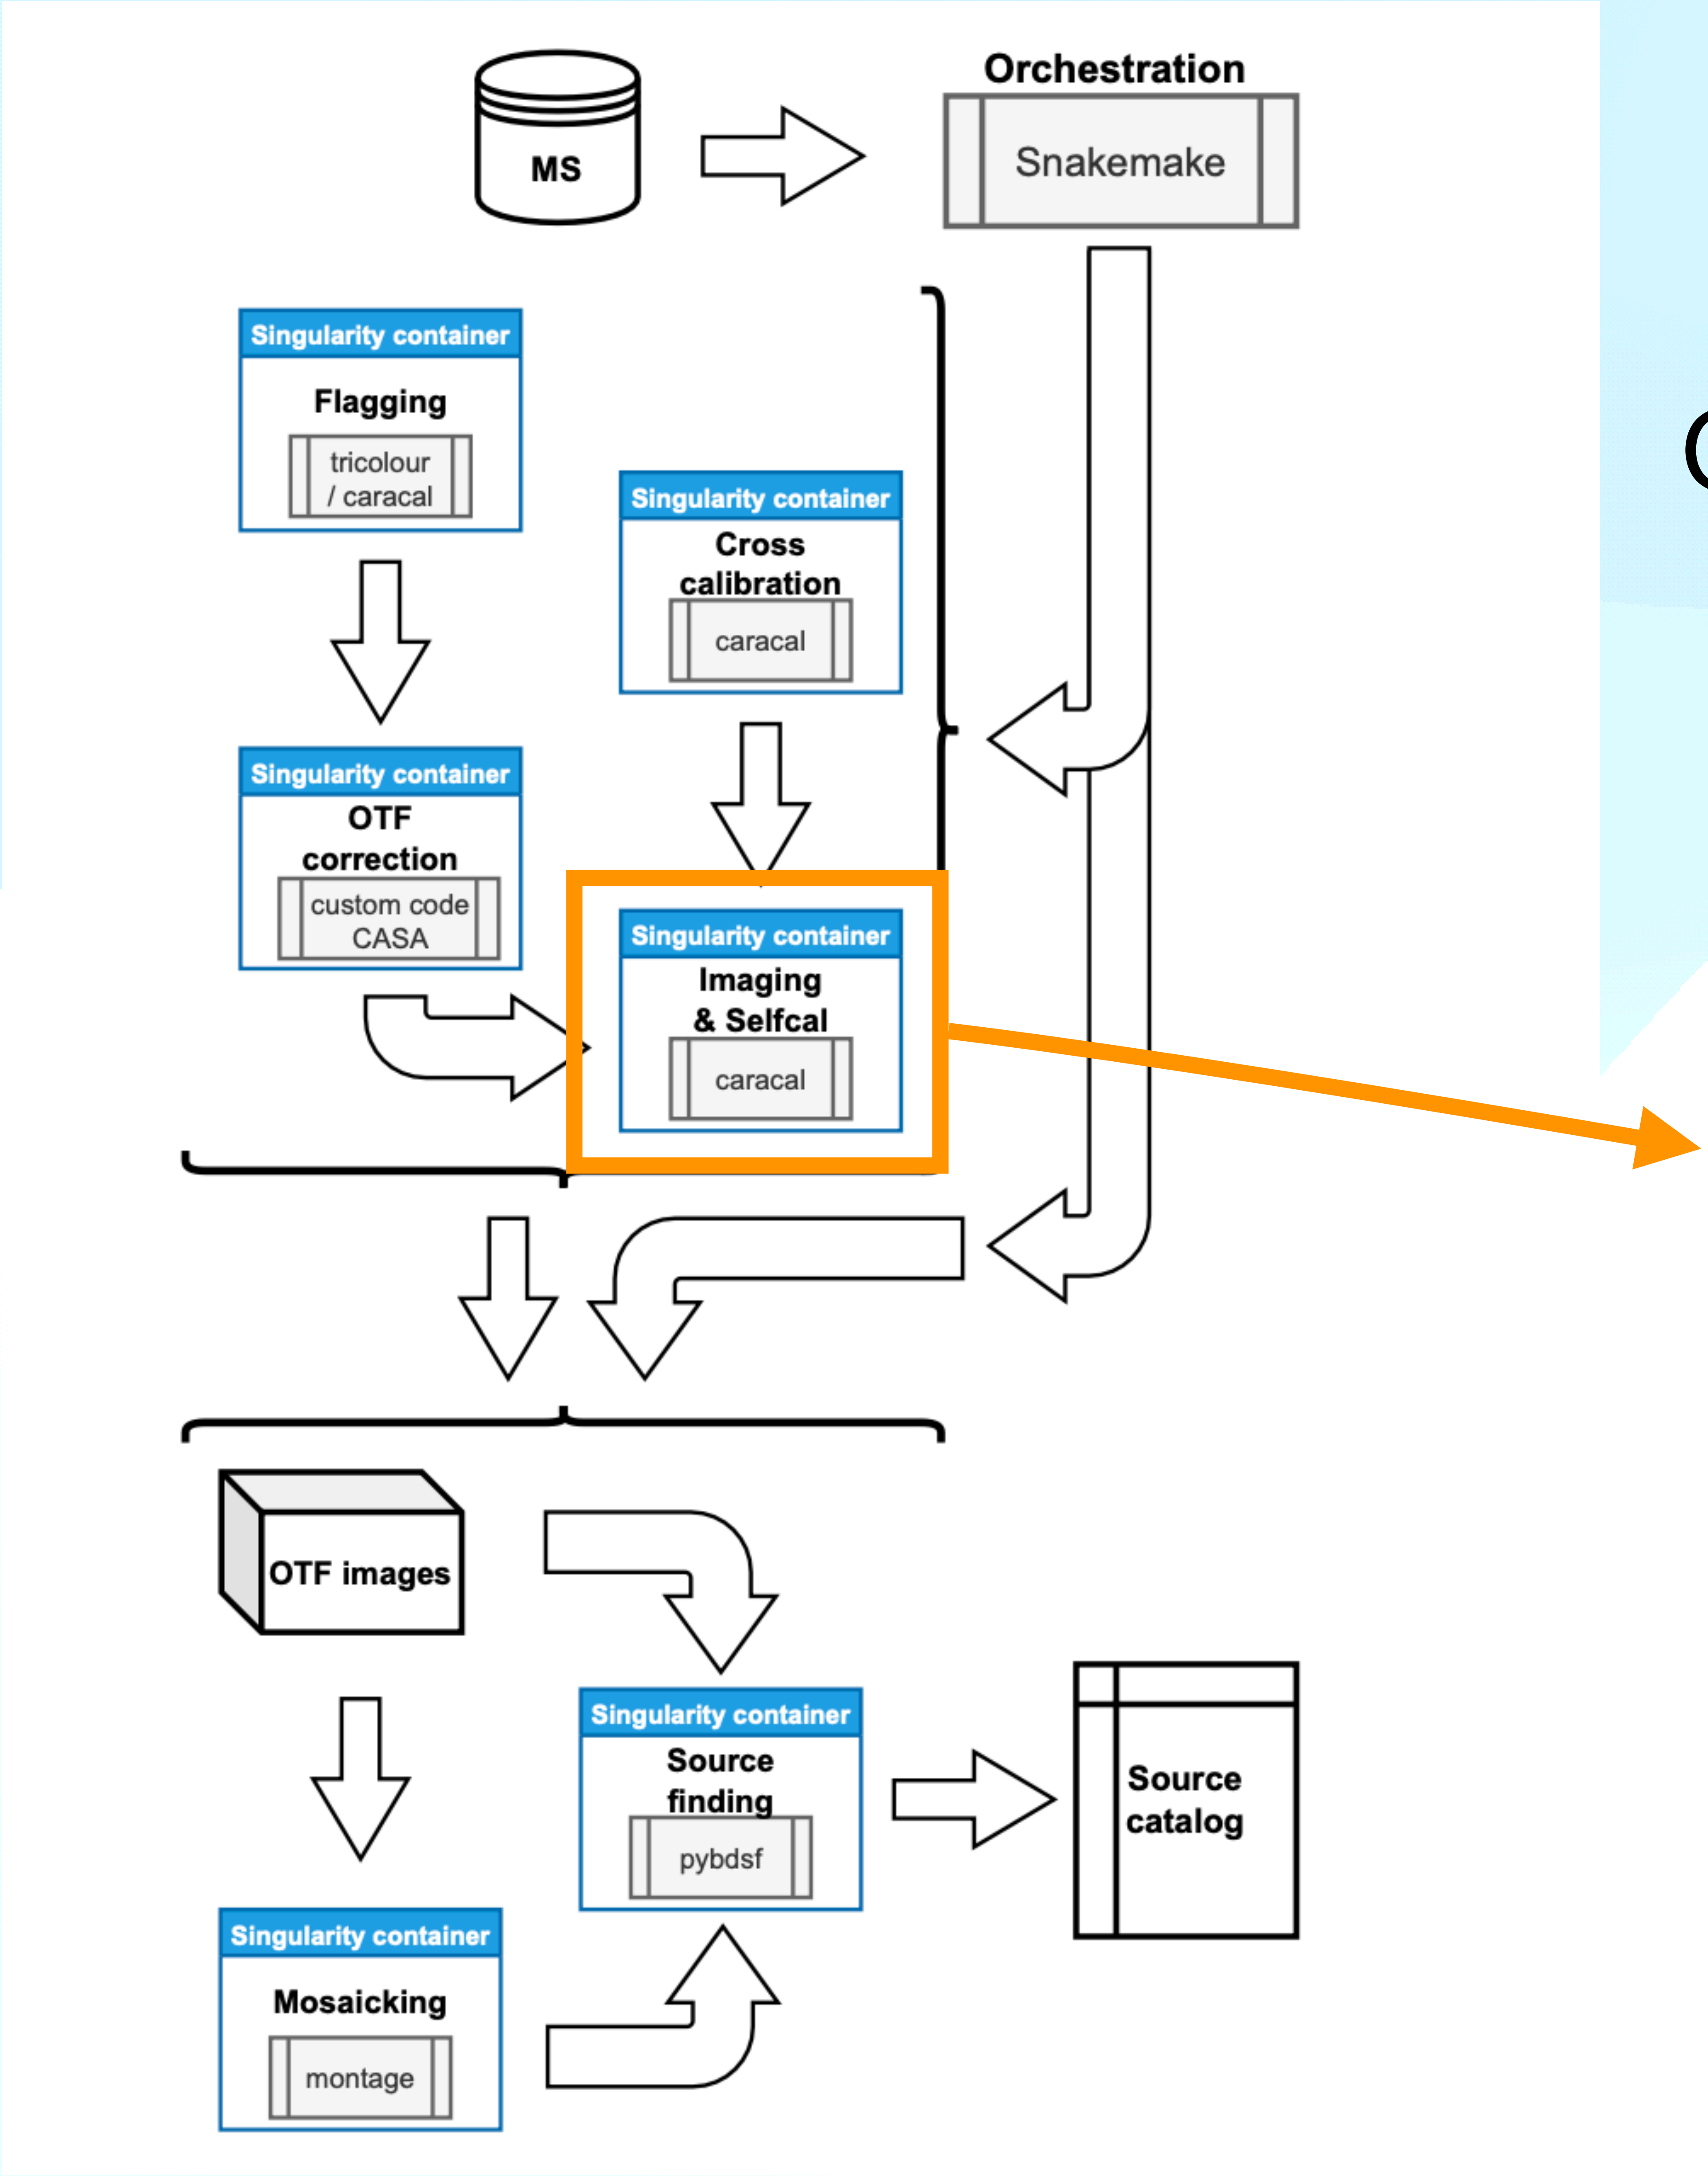

Our solution: workaround for container (cabs) access and modification of the caracal software.

- Caracal can use stimela to access containers in a local folder instead of docker registry.
- Caracal developers modified caracal version so that local folder can be CVMFS.

Achievements - Imaging

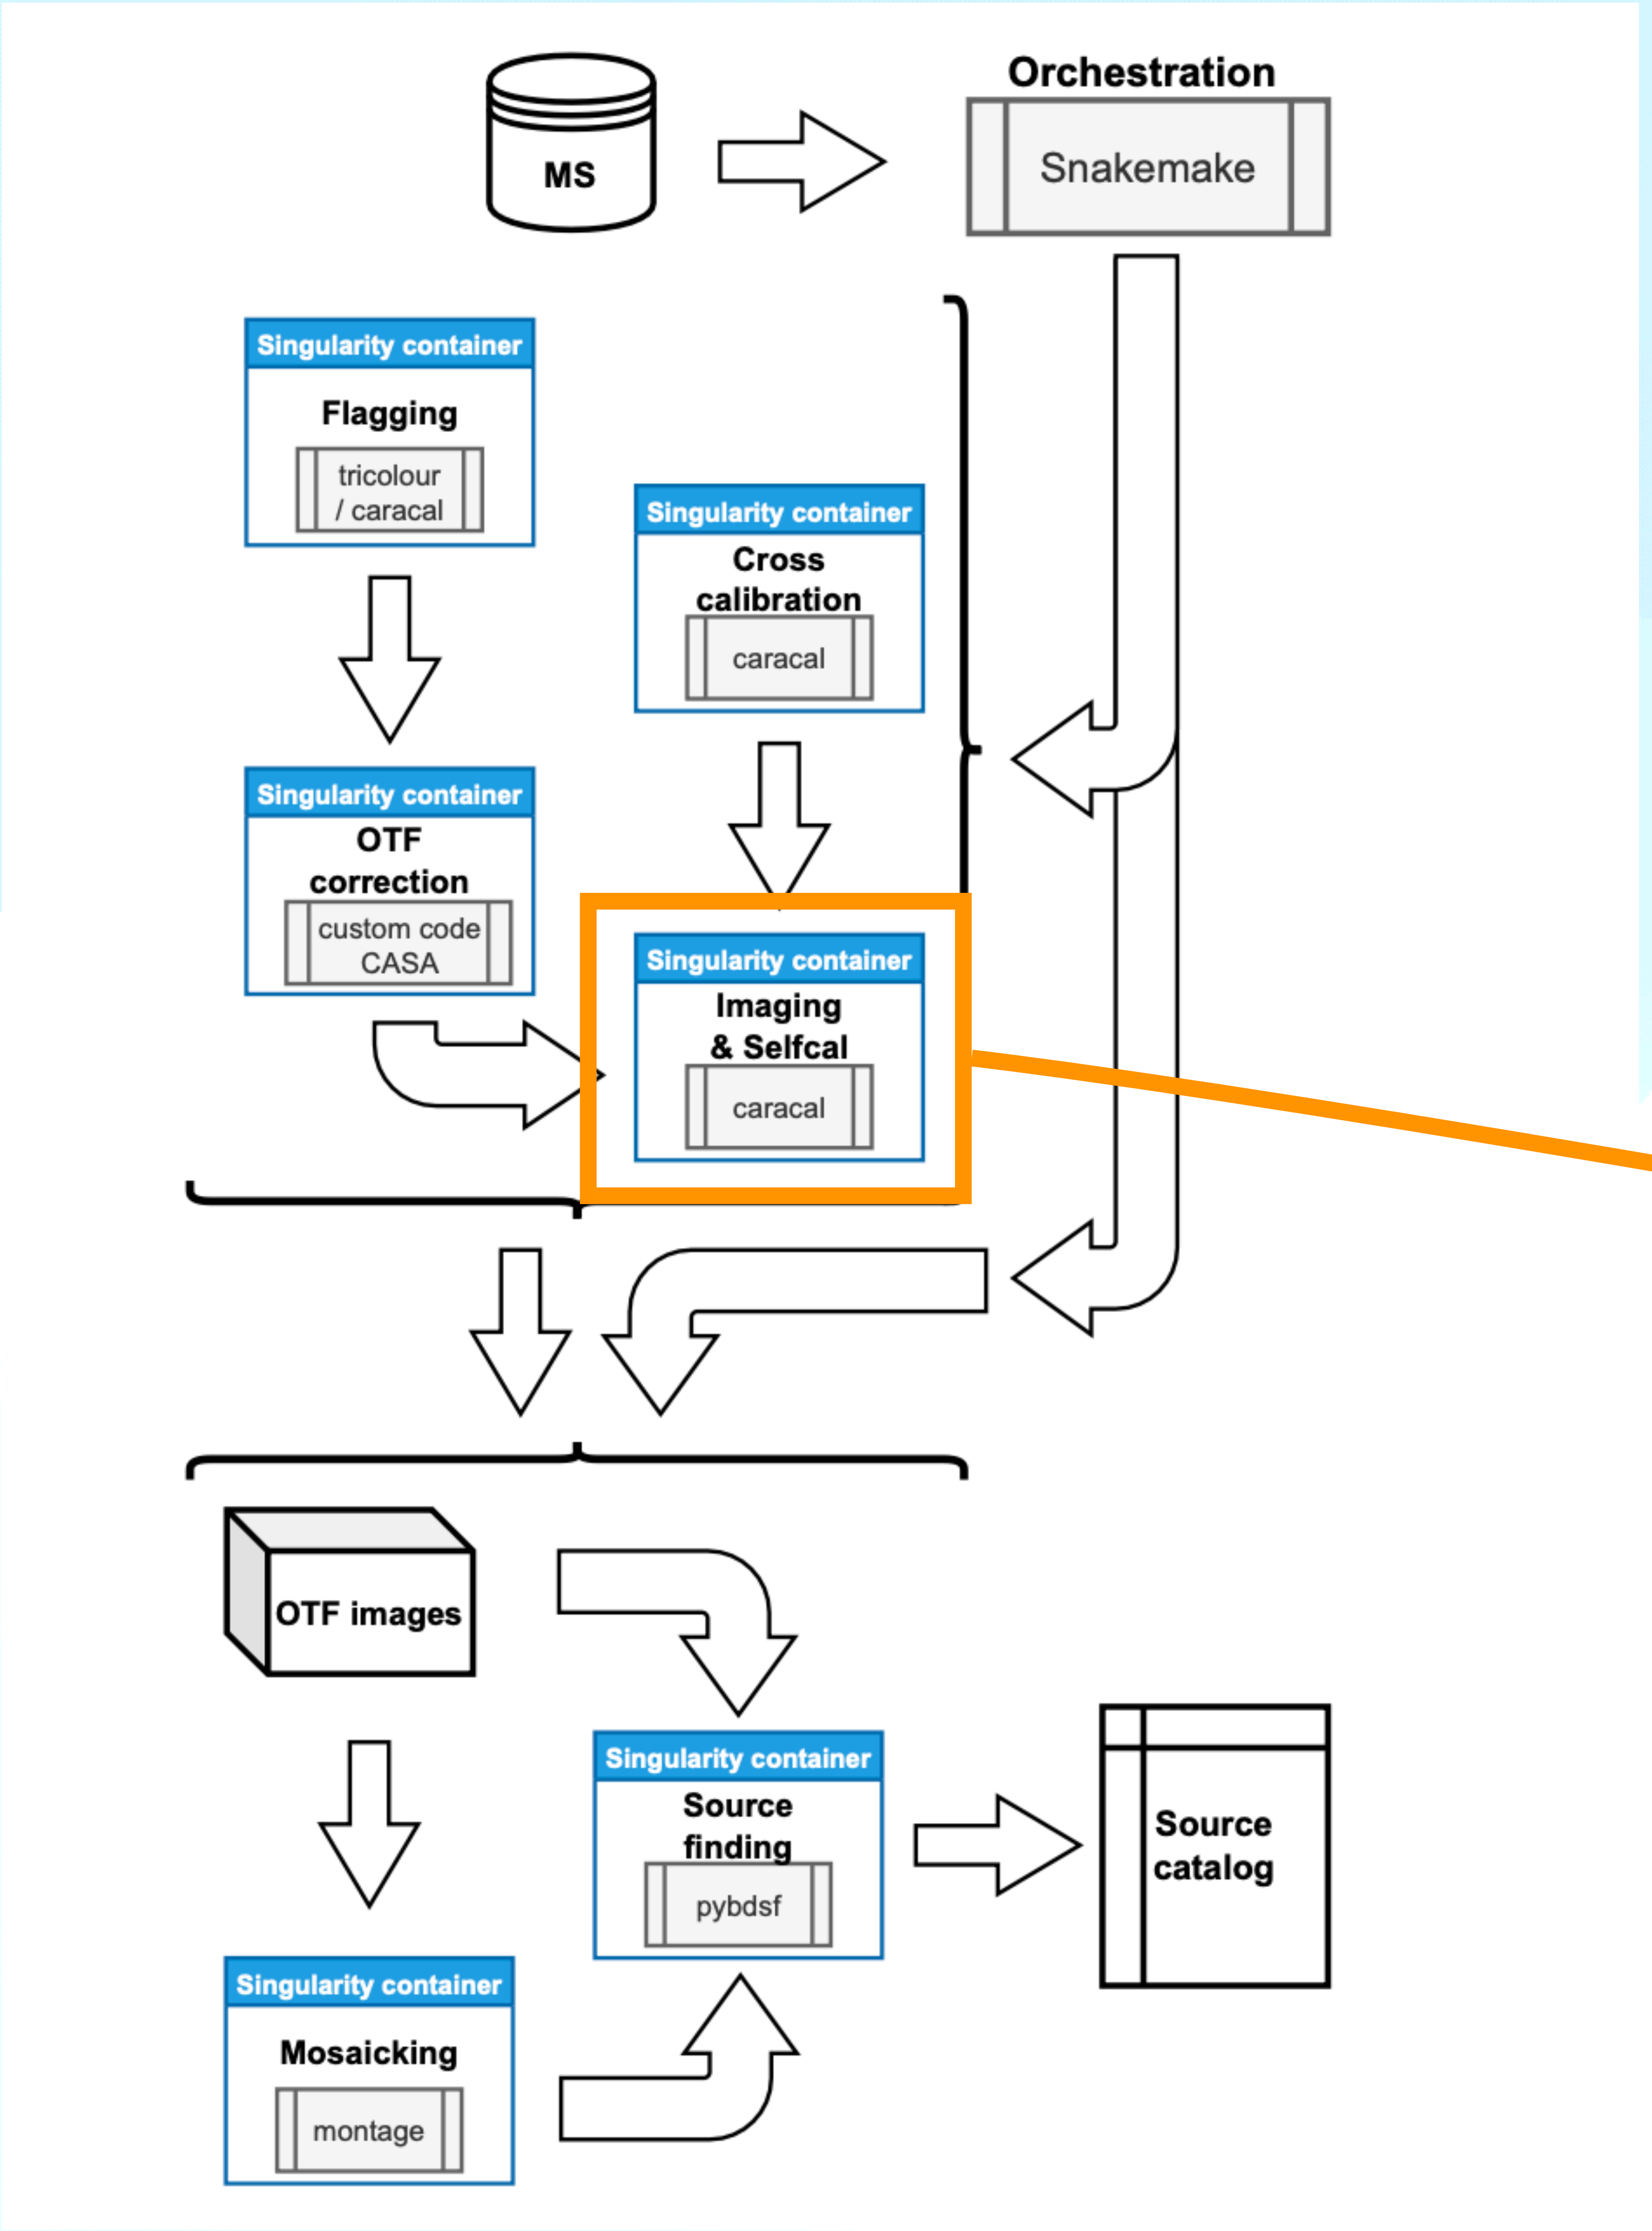

Our solution: workaround for container (cabs) access and modification of the caracal software.

- Caracal can use stimela to access containers in a local folder instead of docker registry.
- Caracal developers modified caracal version so that local folder can be CVMFS.
- Caracal now accesses cabs in CVMFS, available from C4P nodes.

Achievements - Imaging

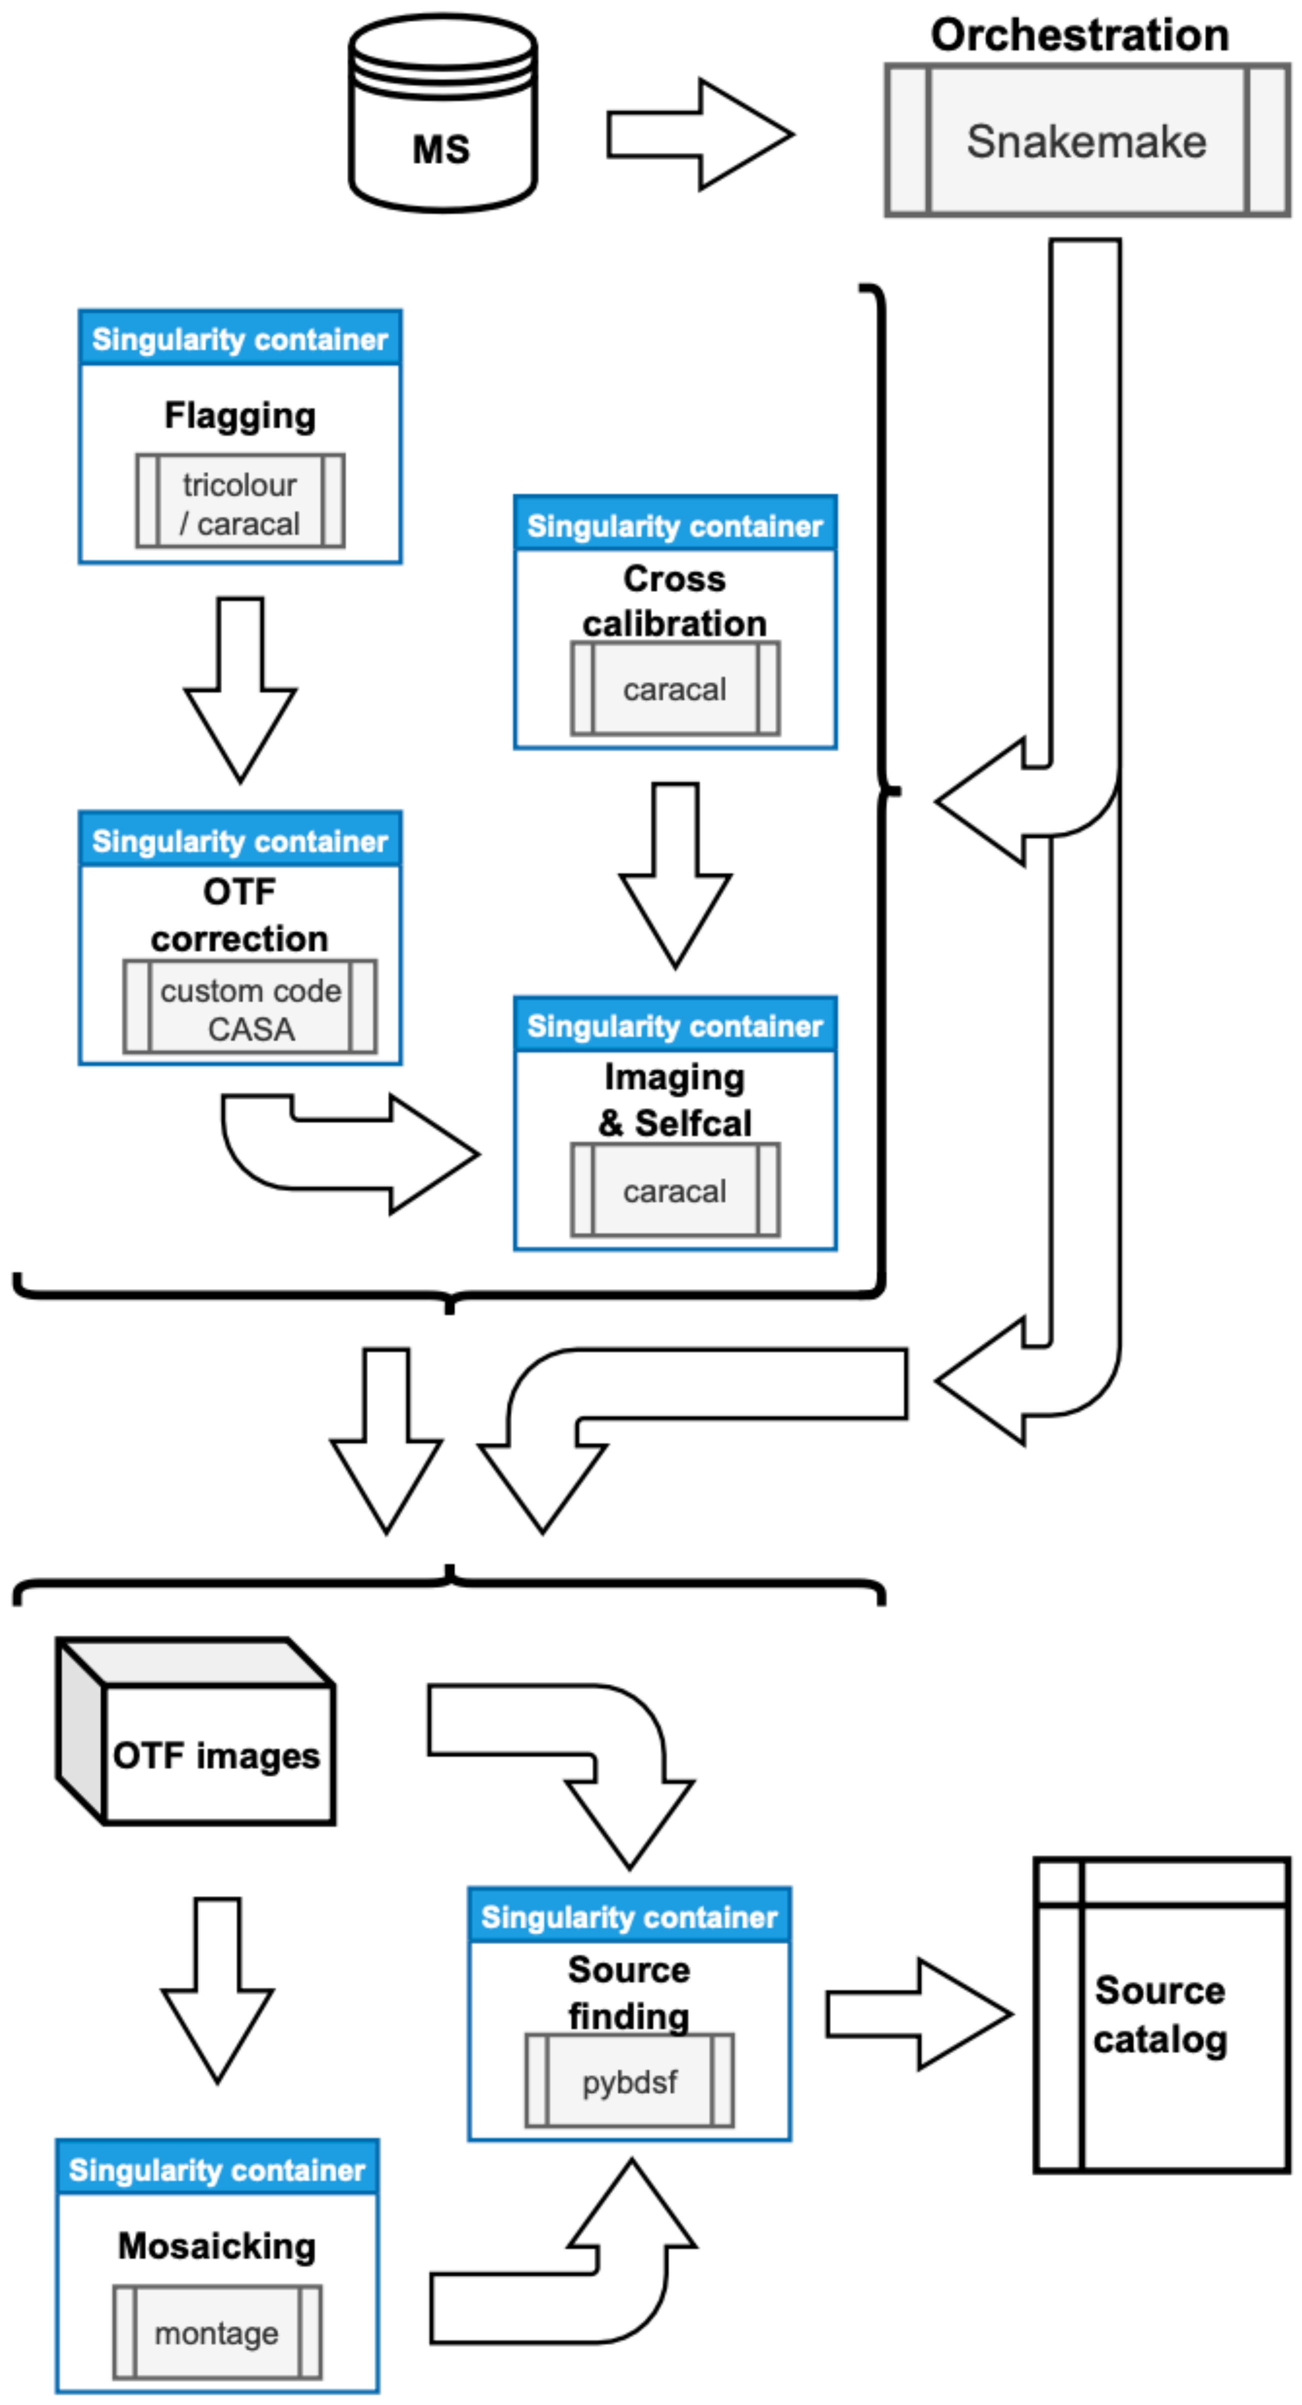

Achievements - Imaging

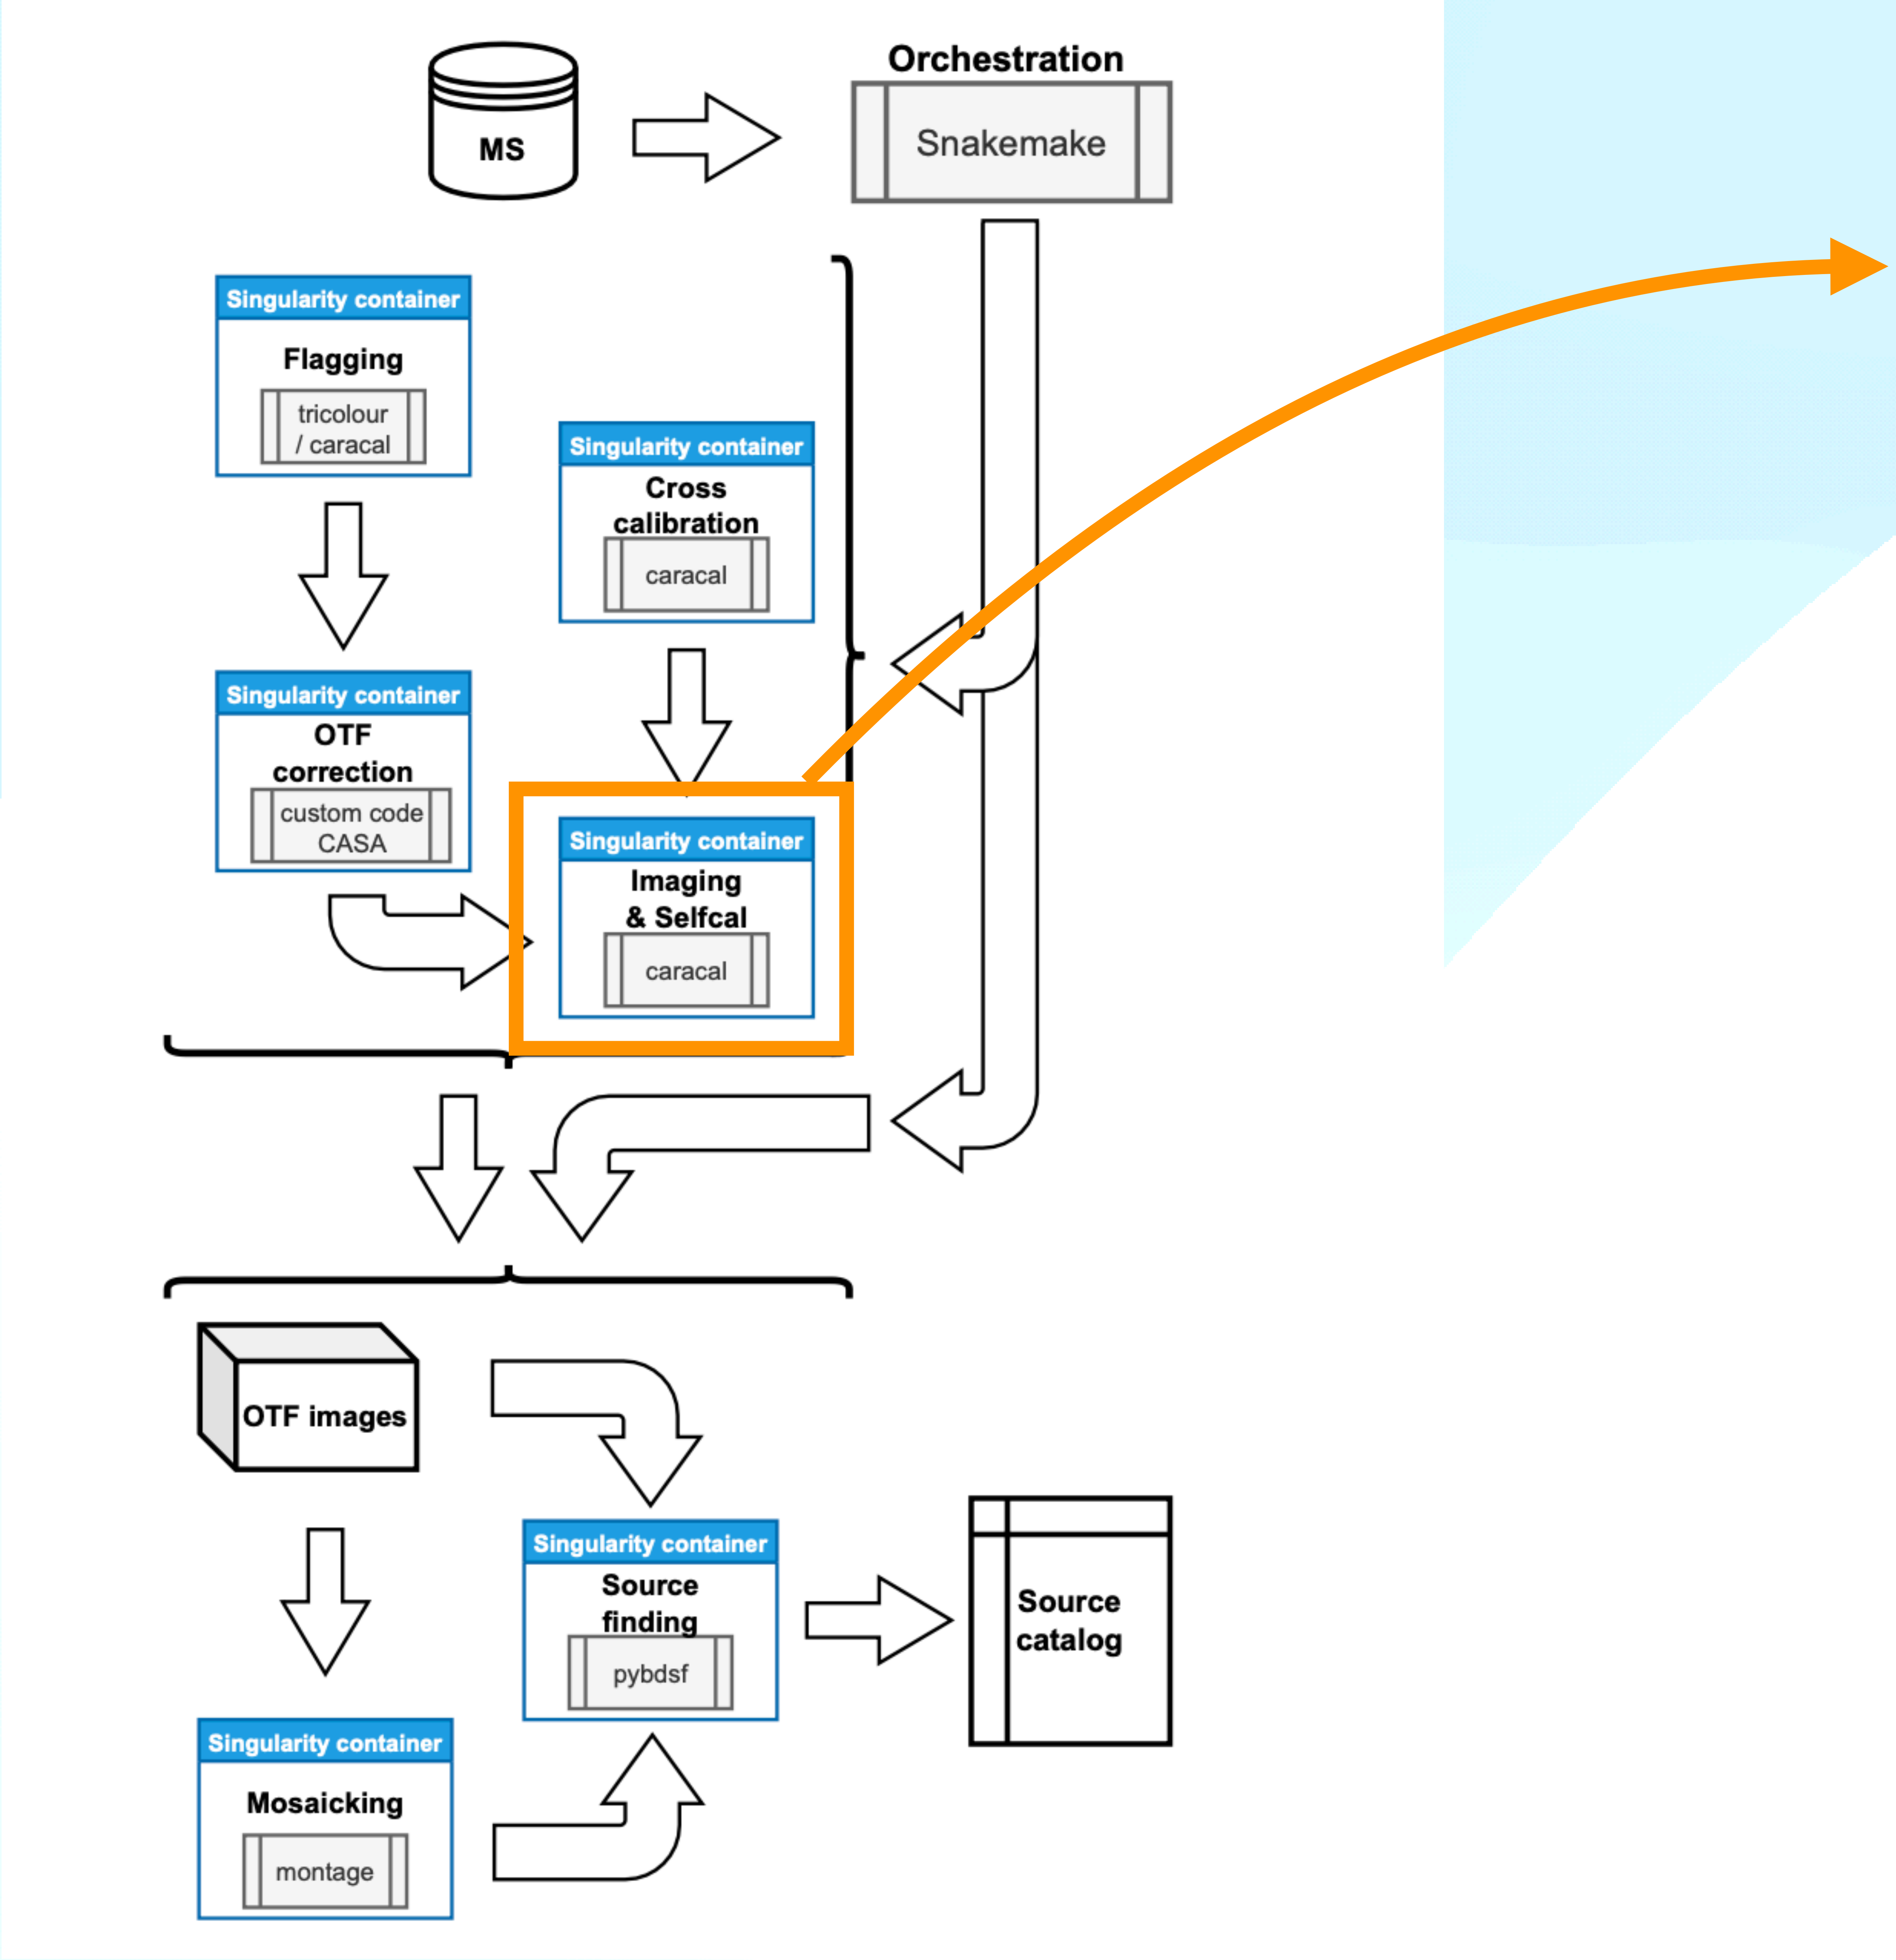

Achievements - Imaging

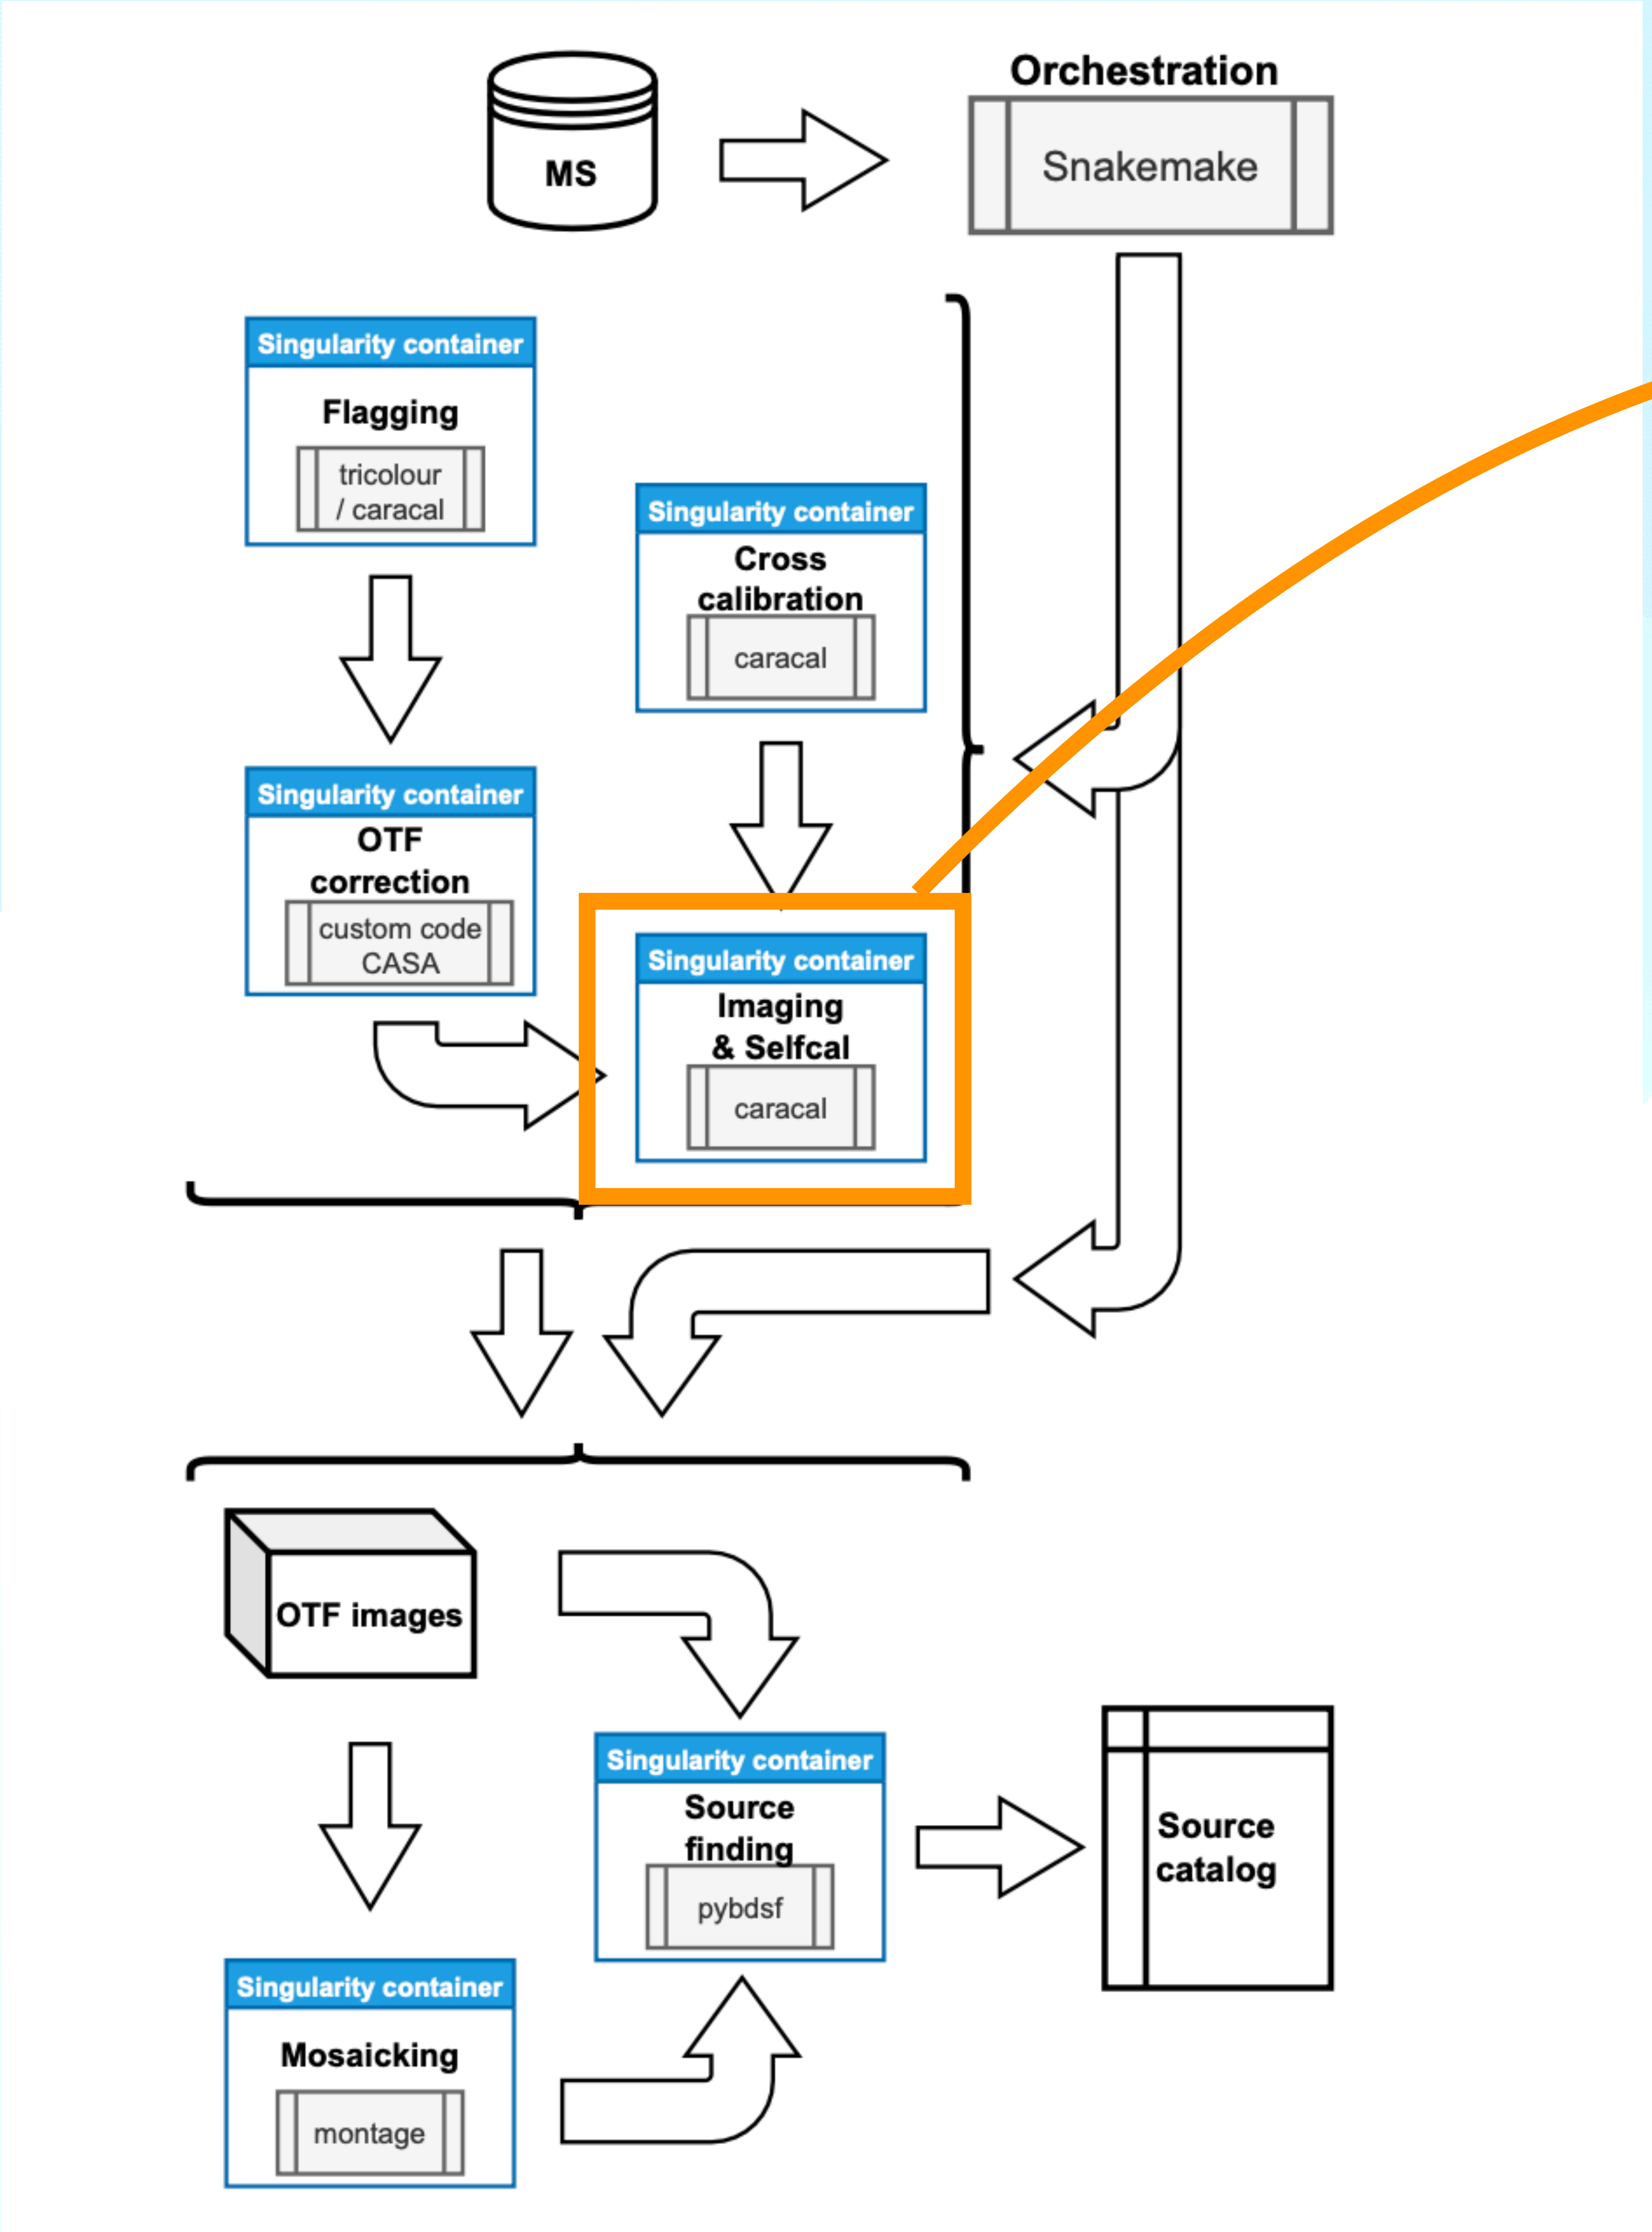

Host system (local laptop)

Achievements - Imaging

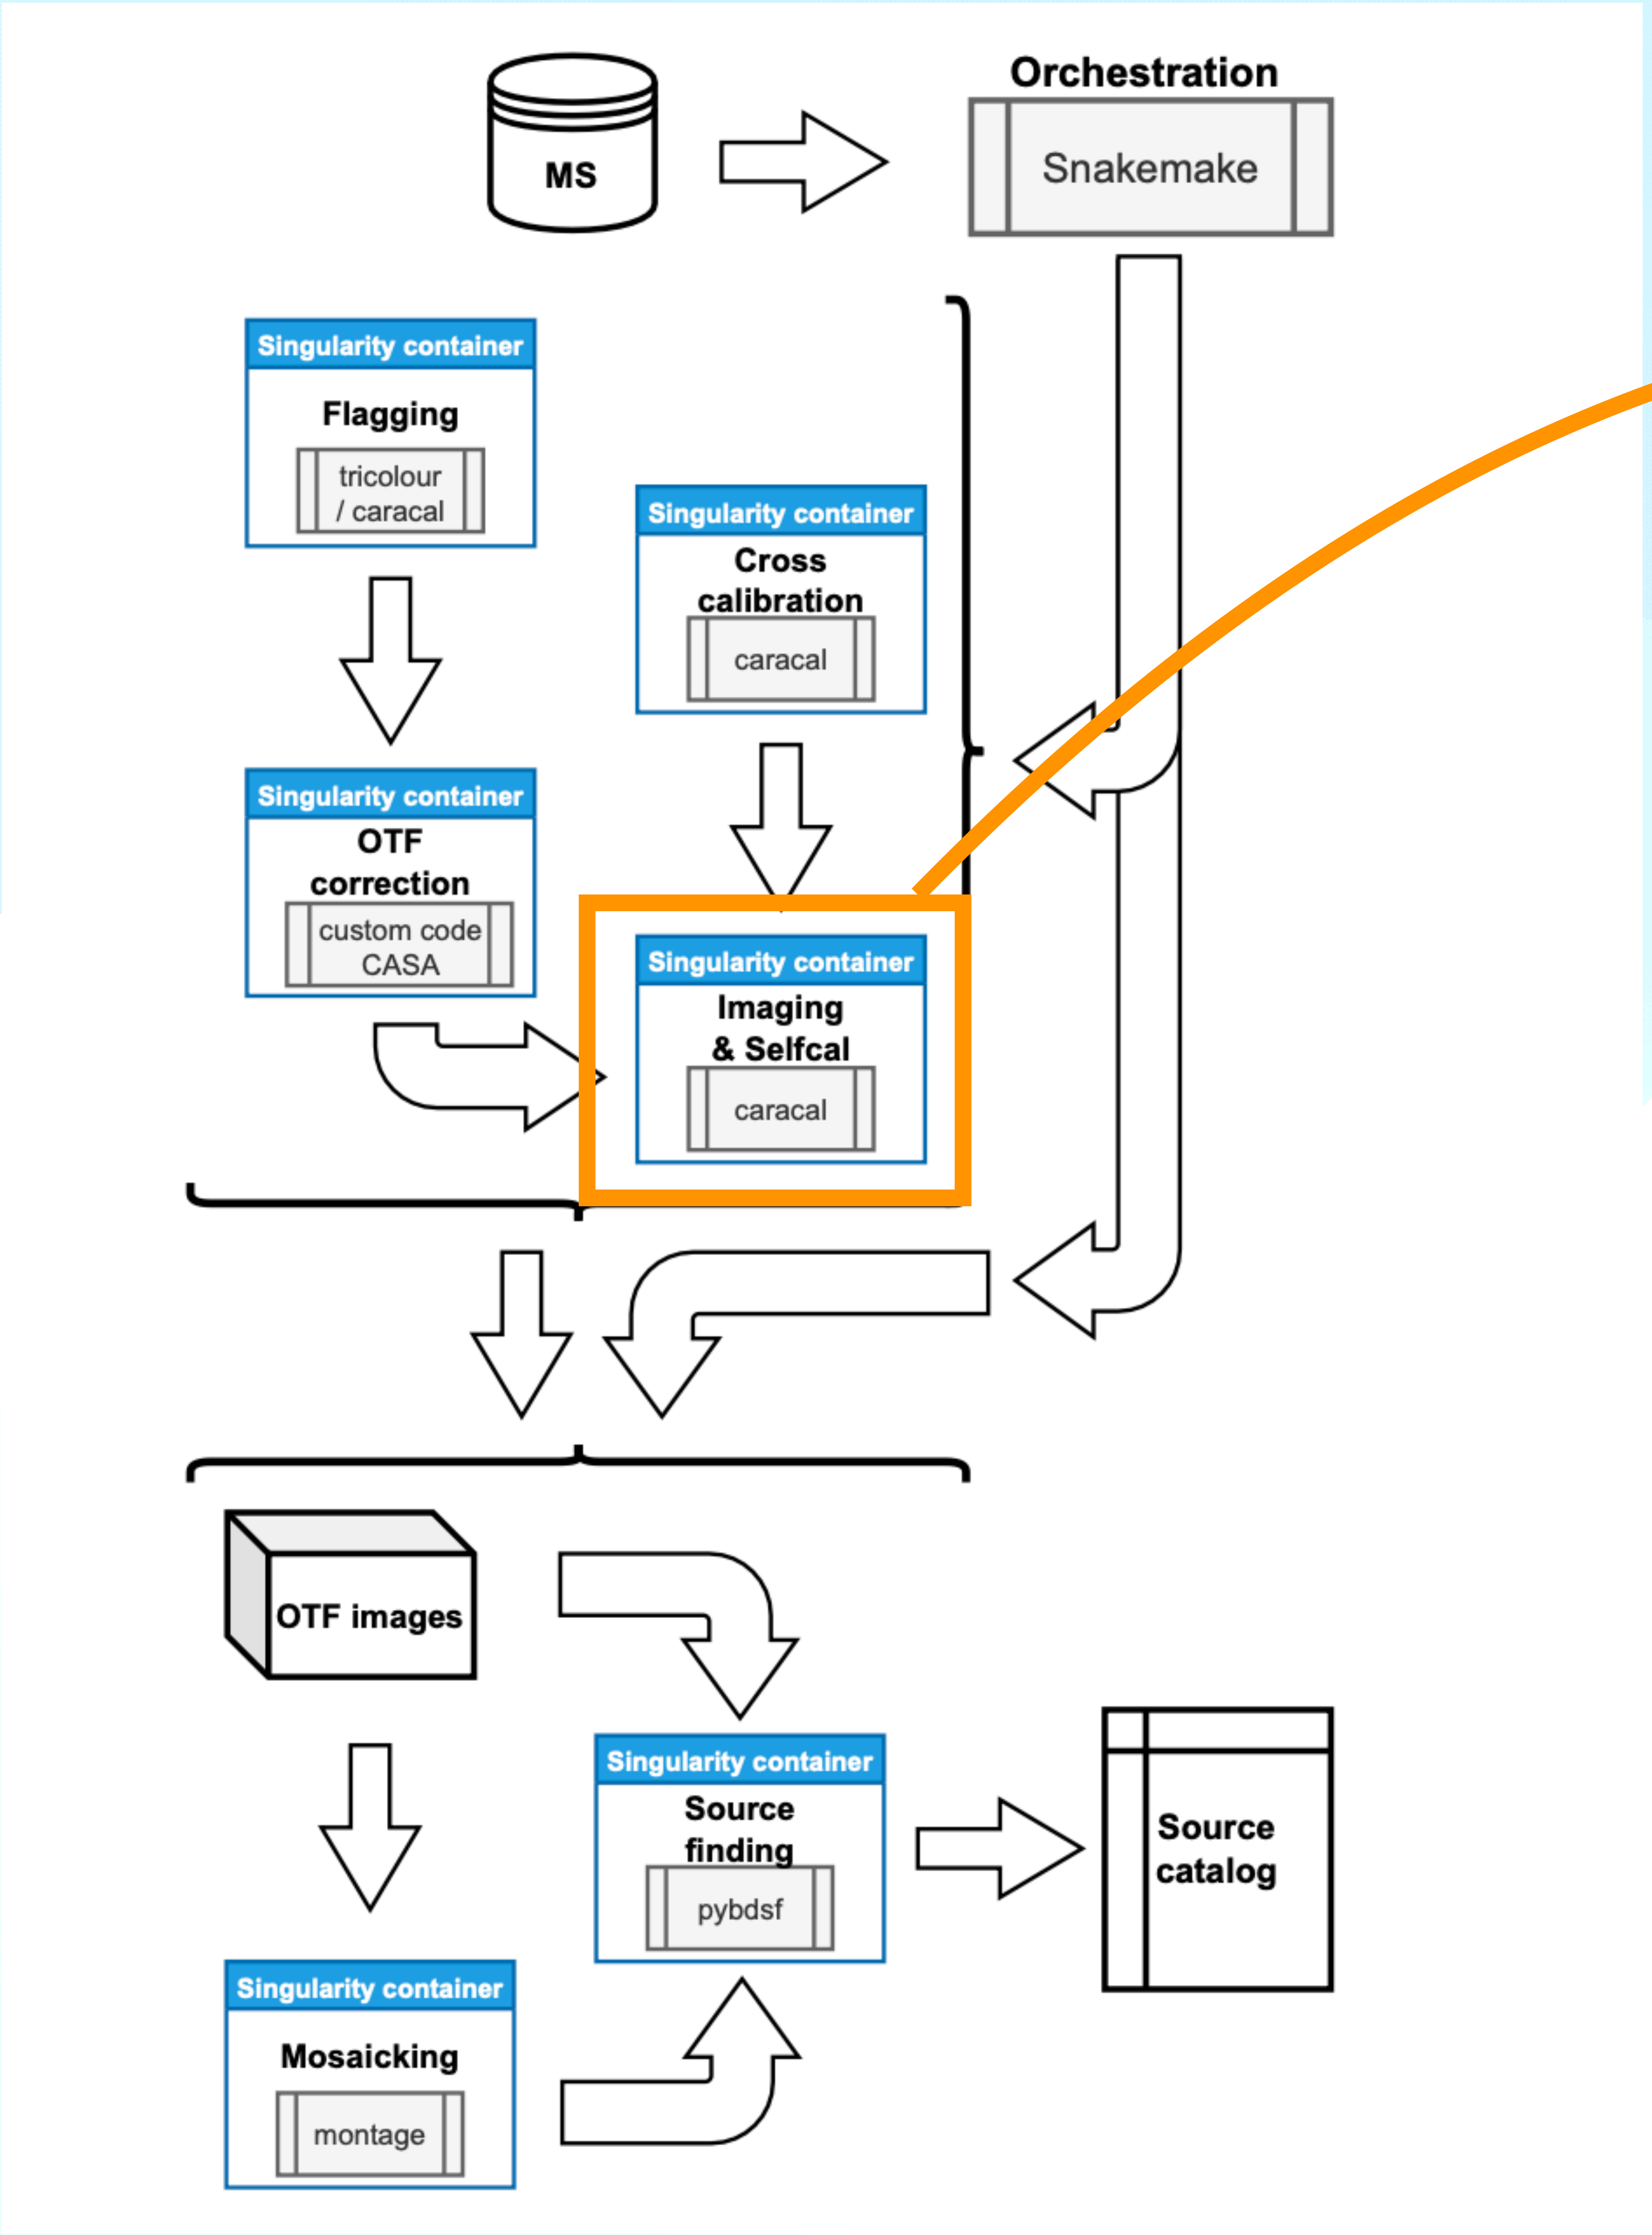

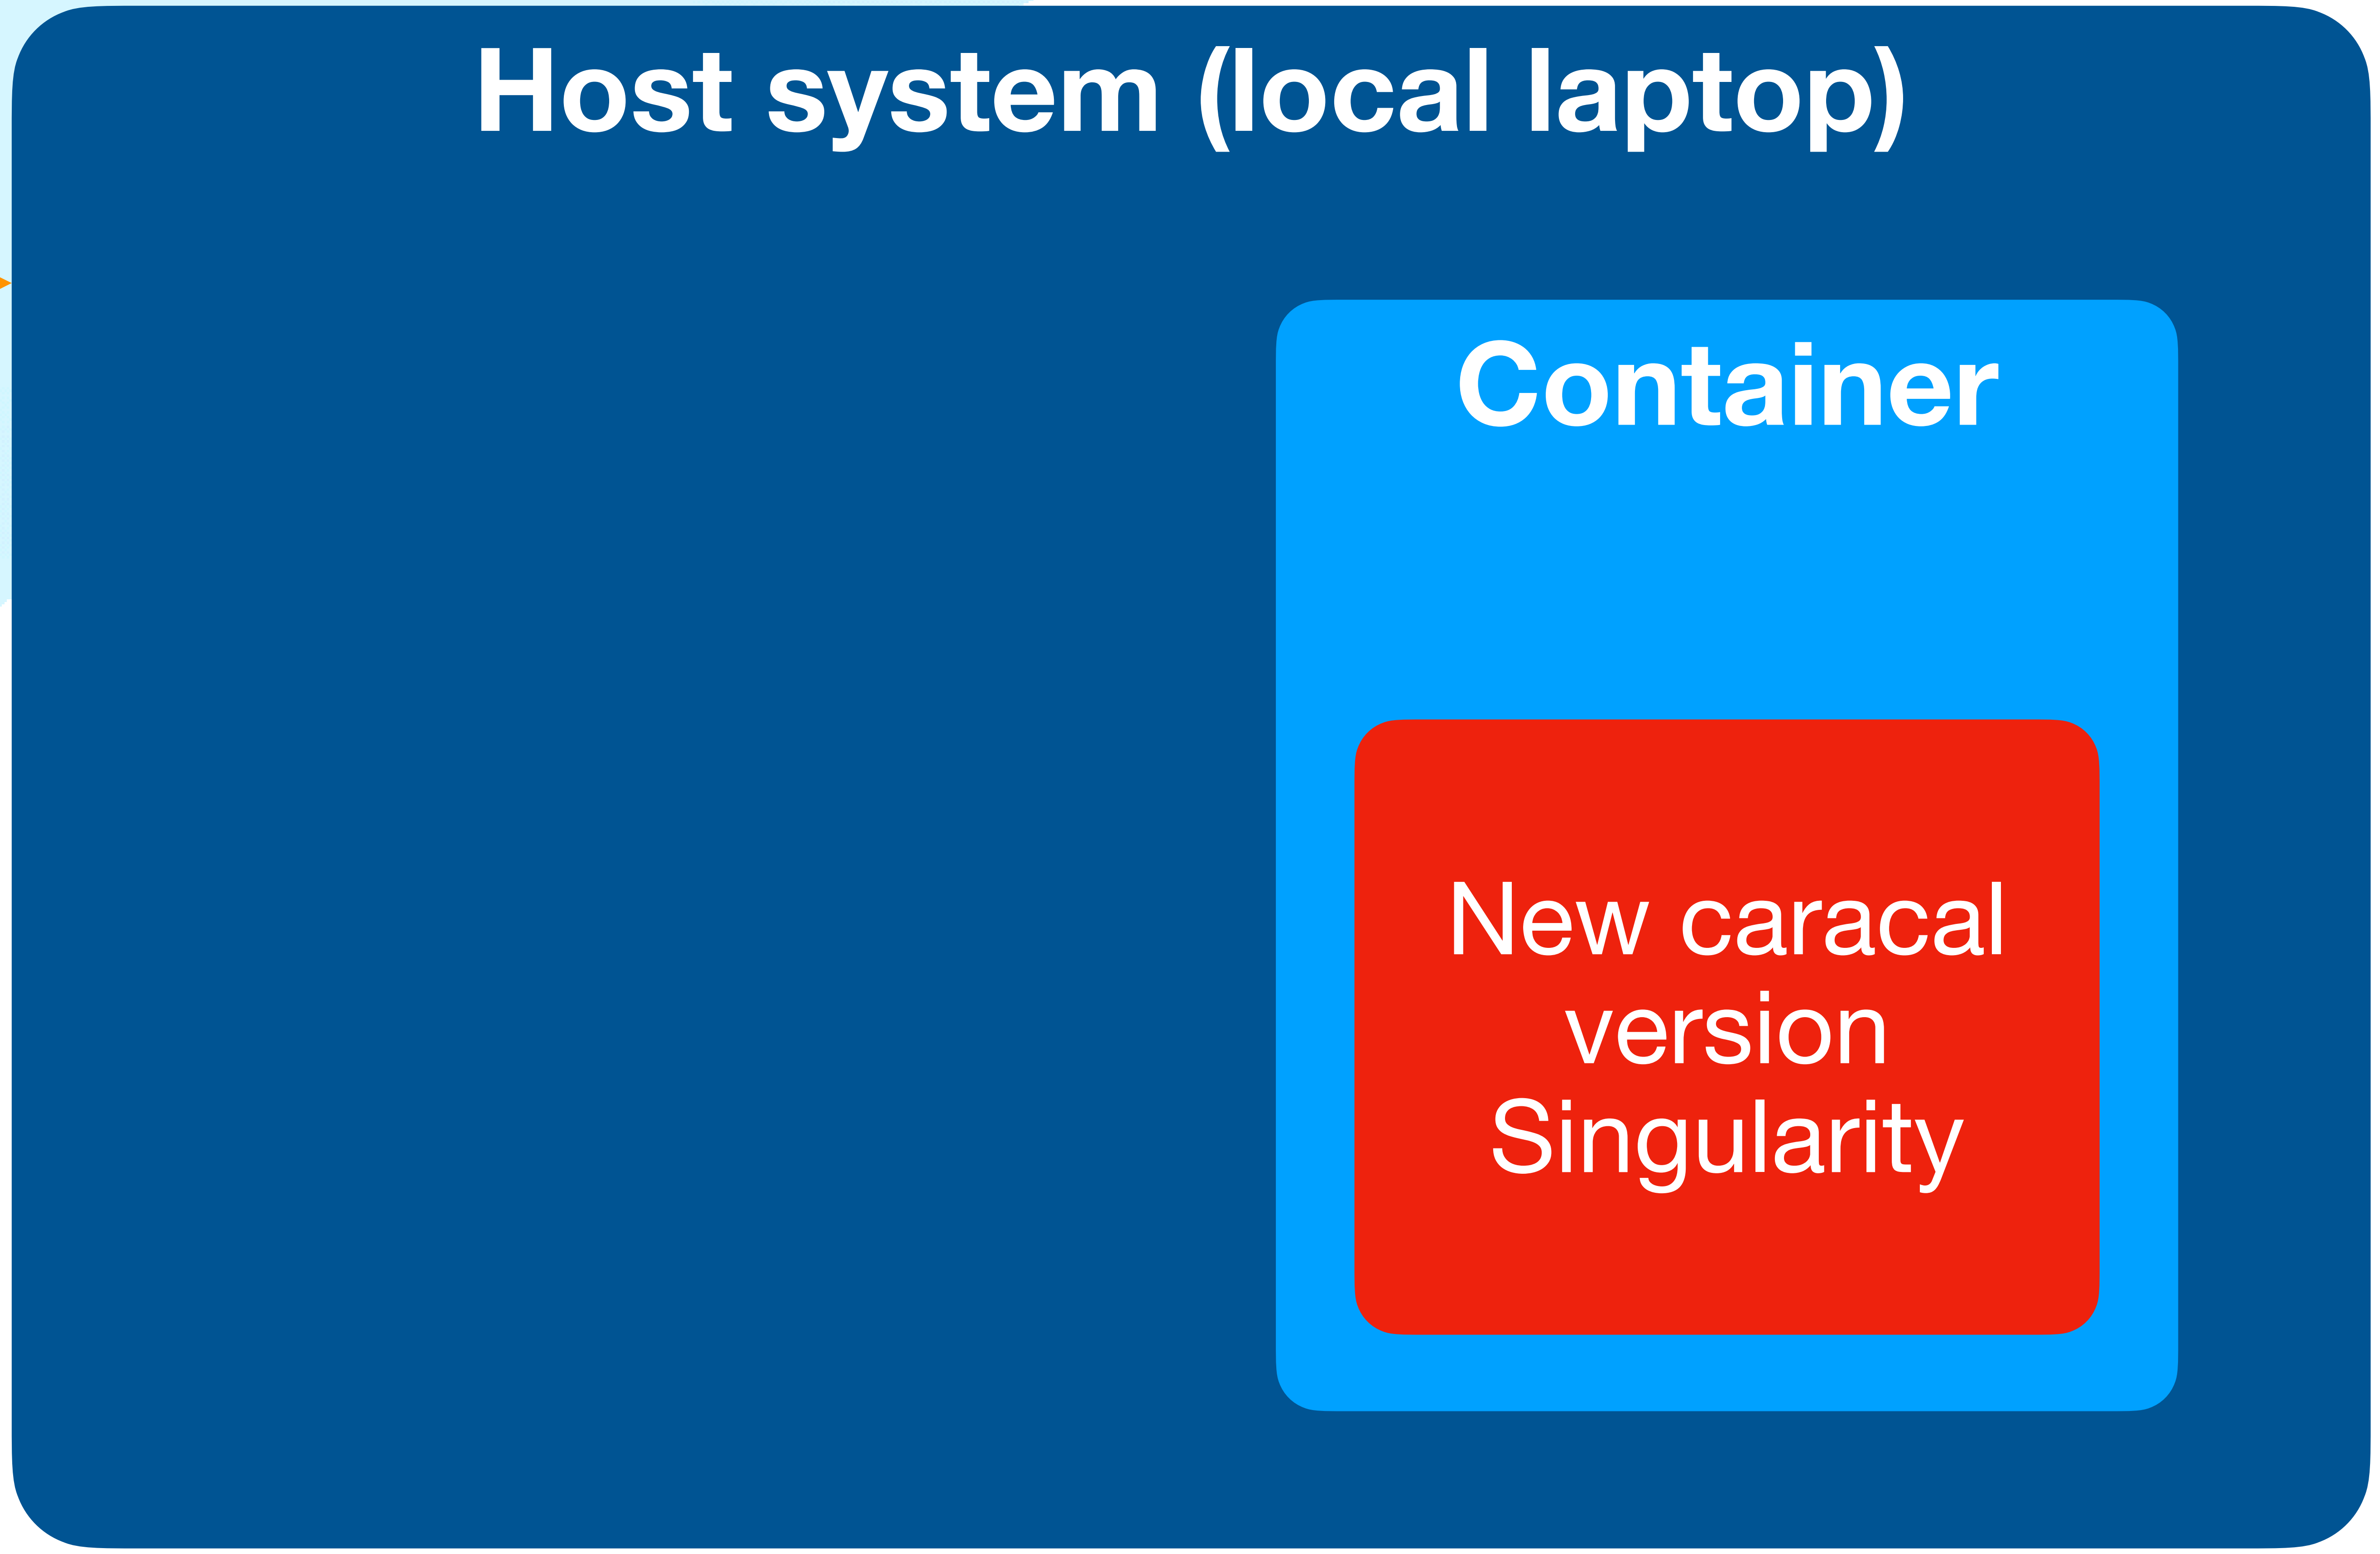

Achievements - Imaging

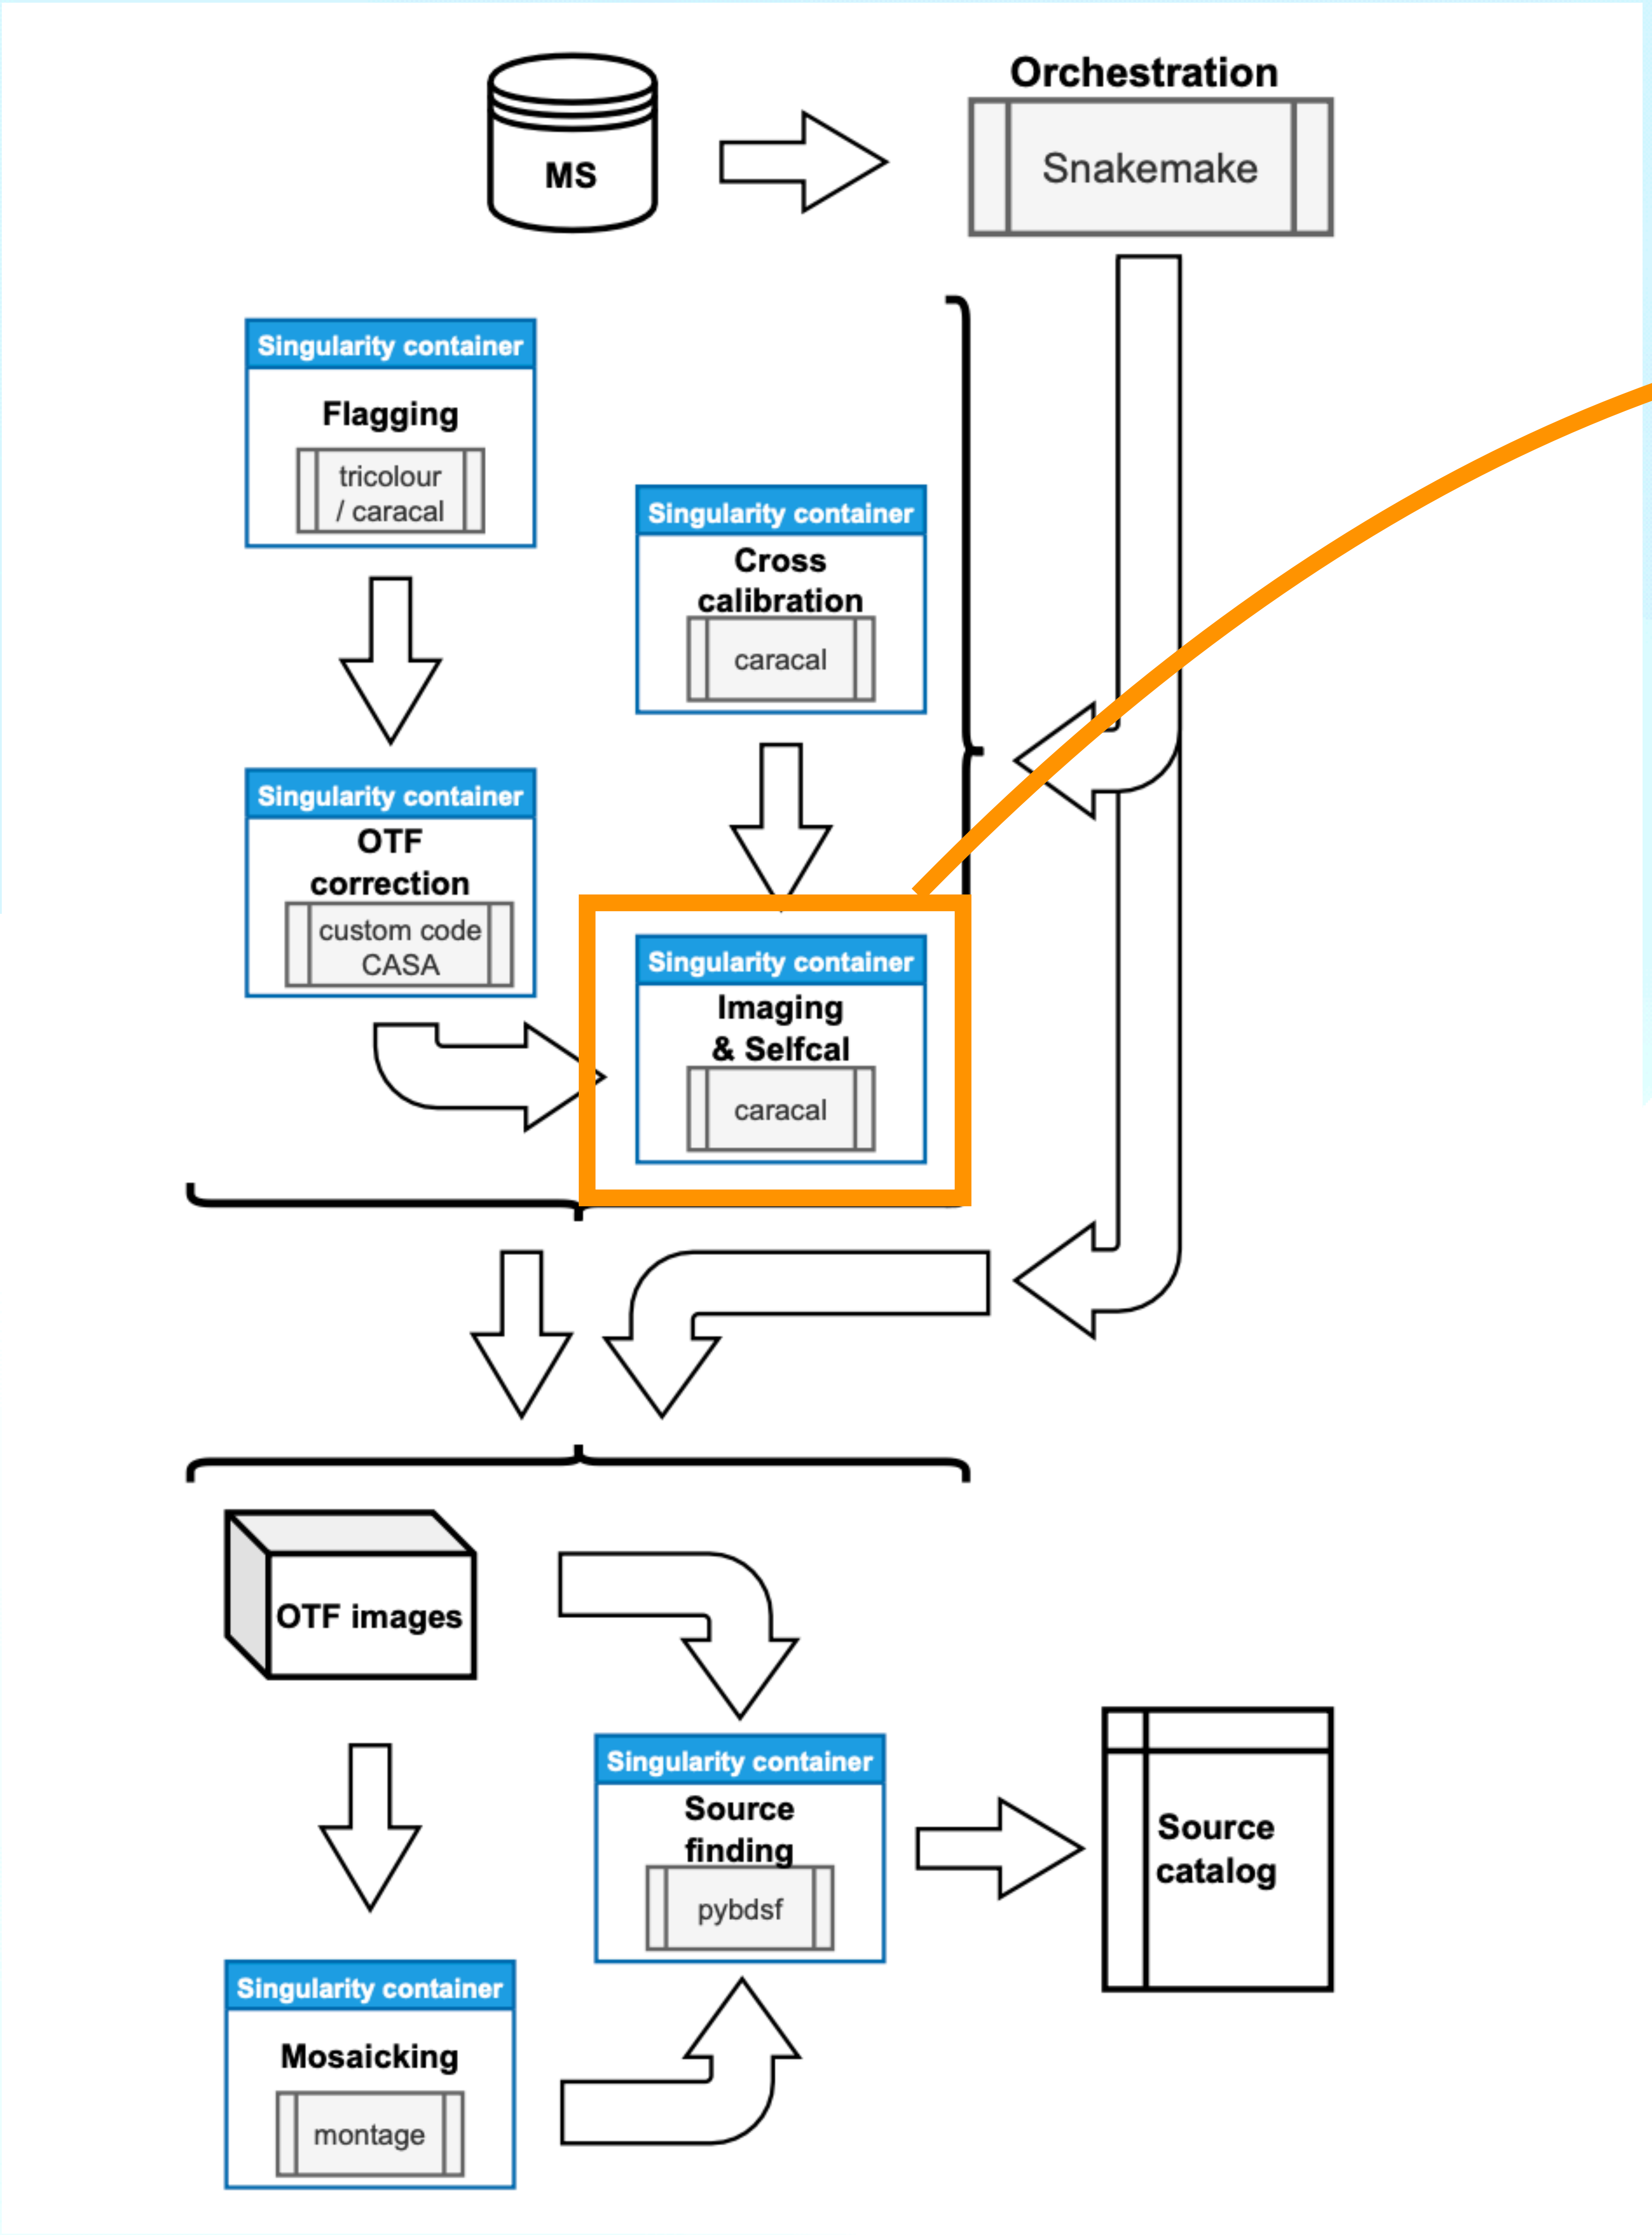

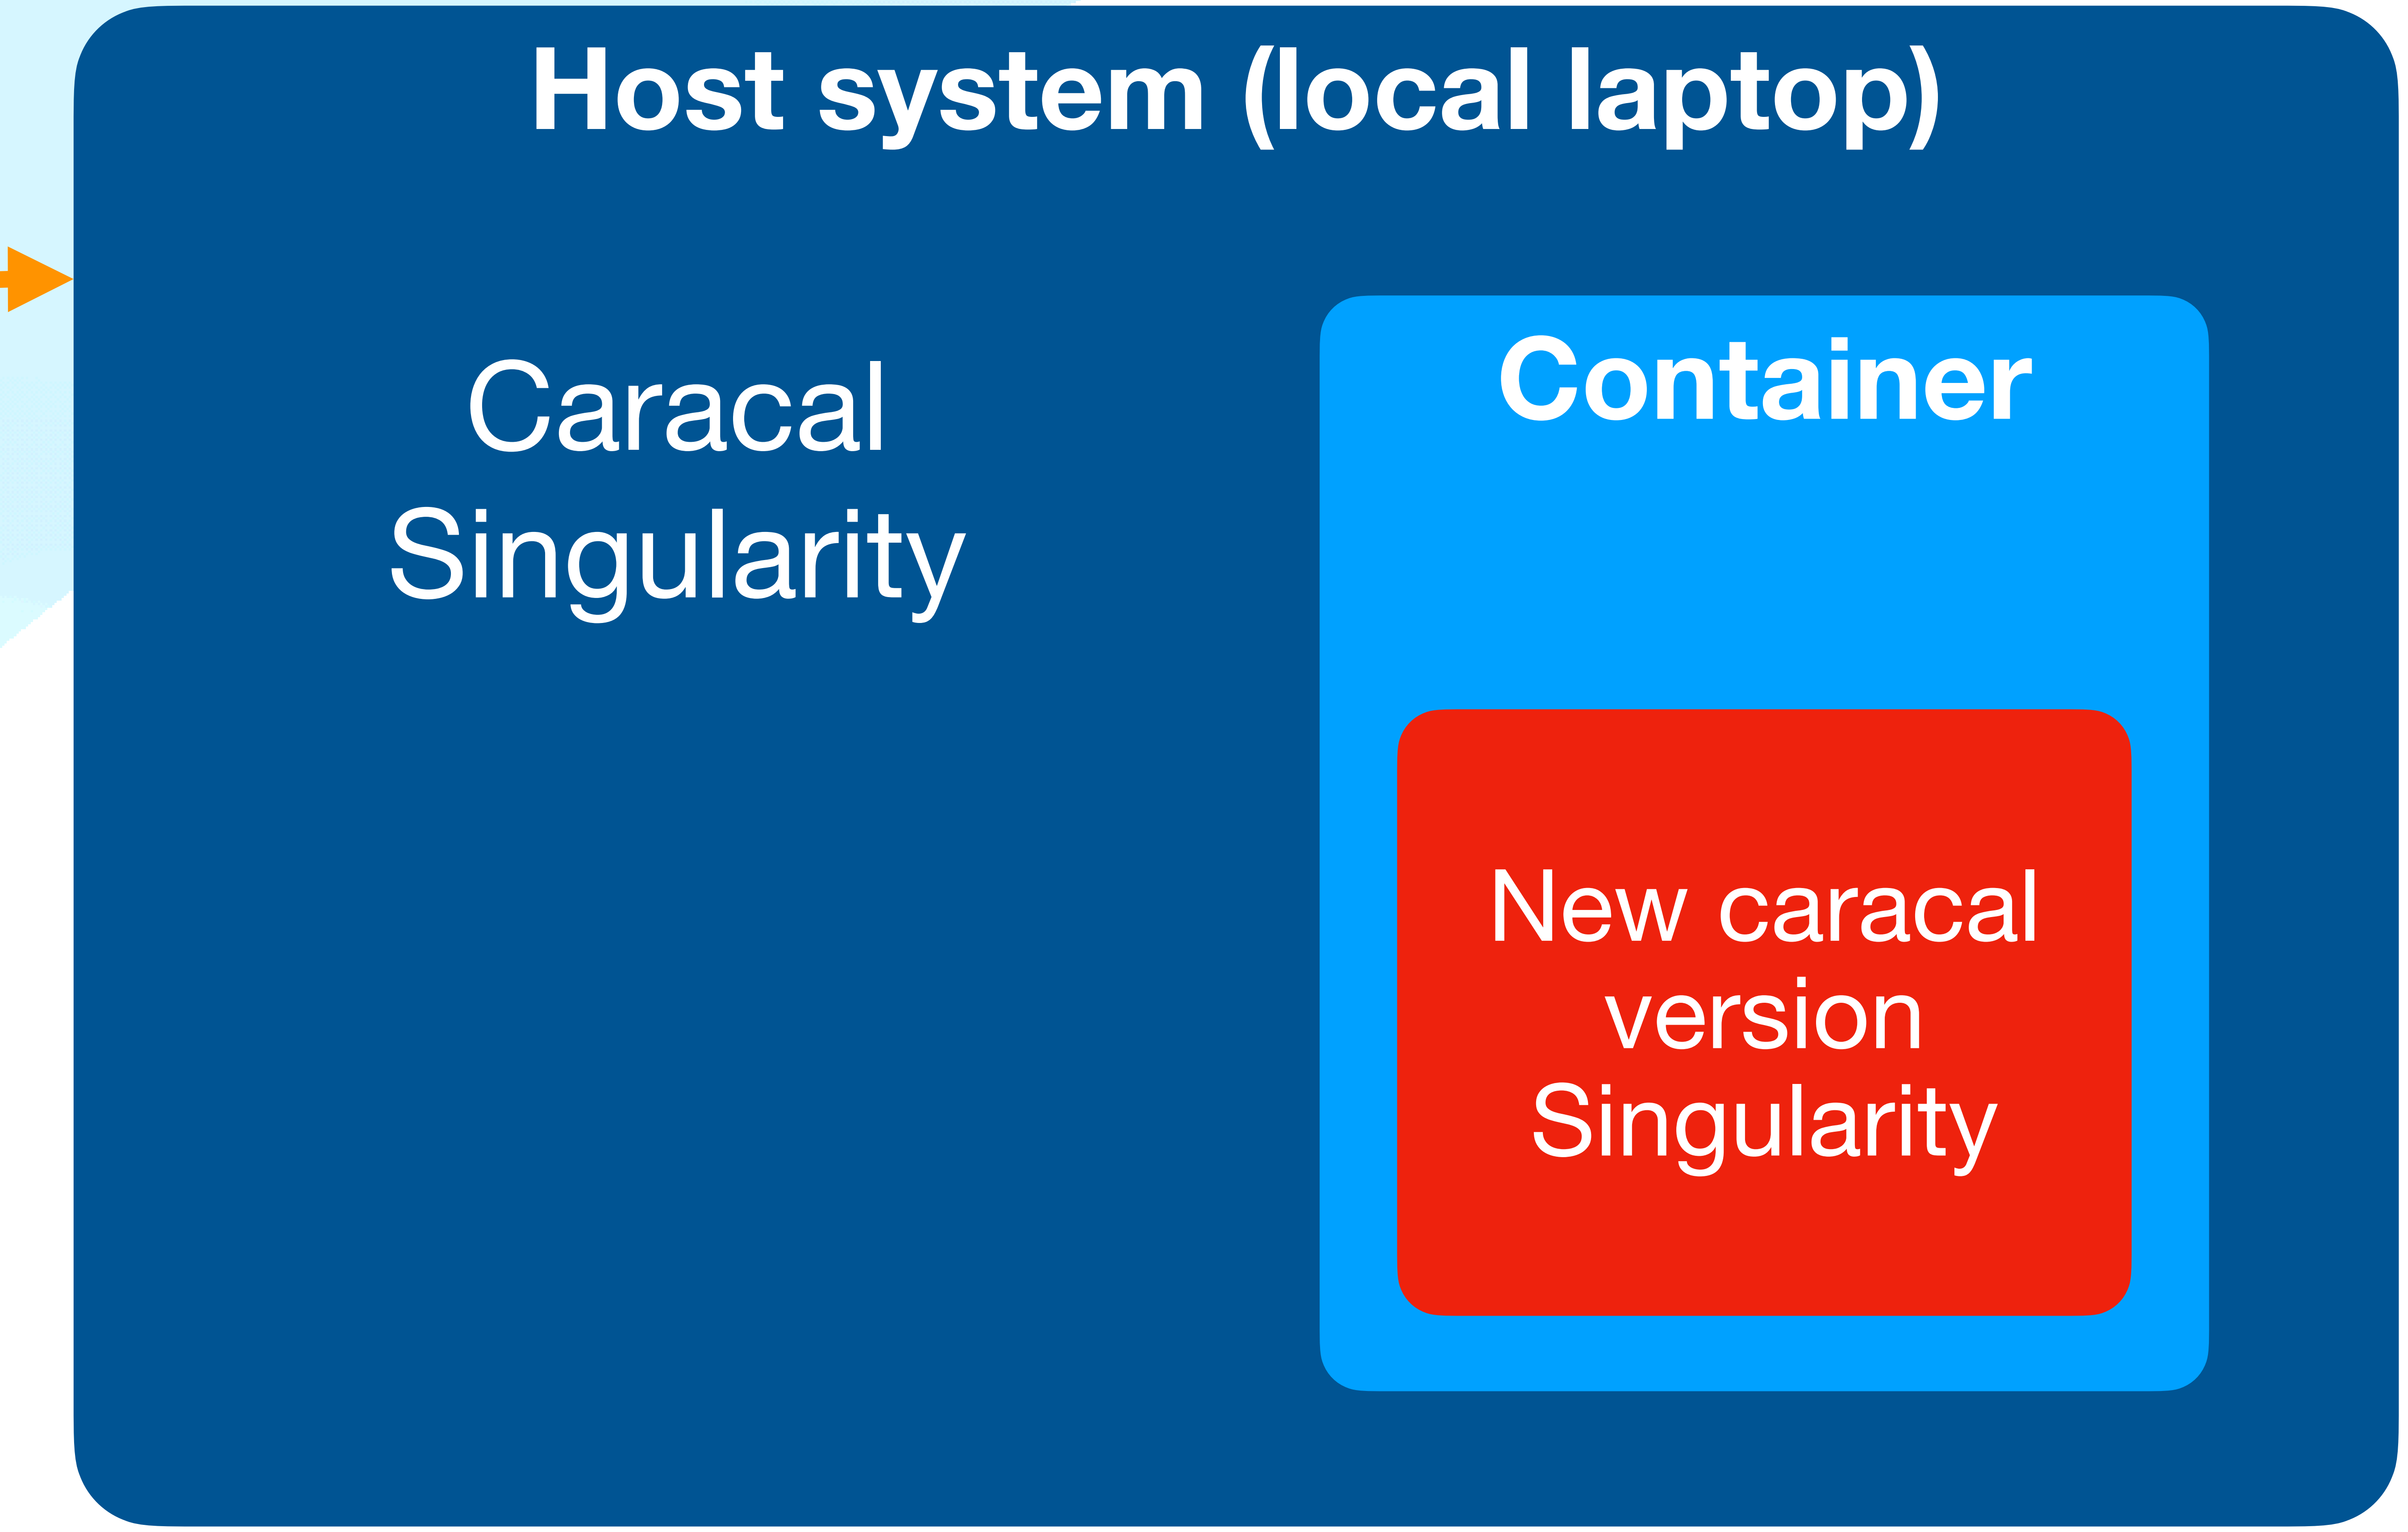

Achievements - Imaging

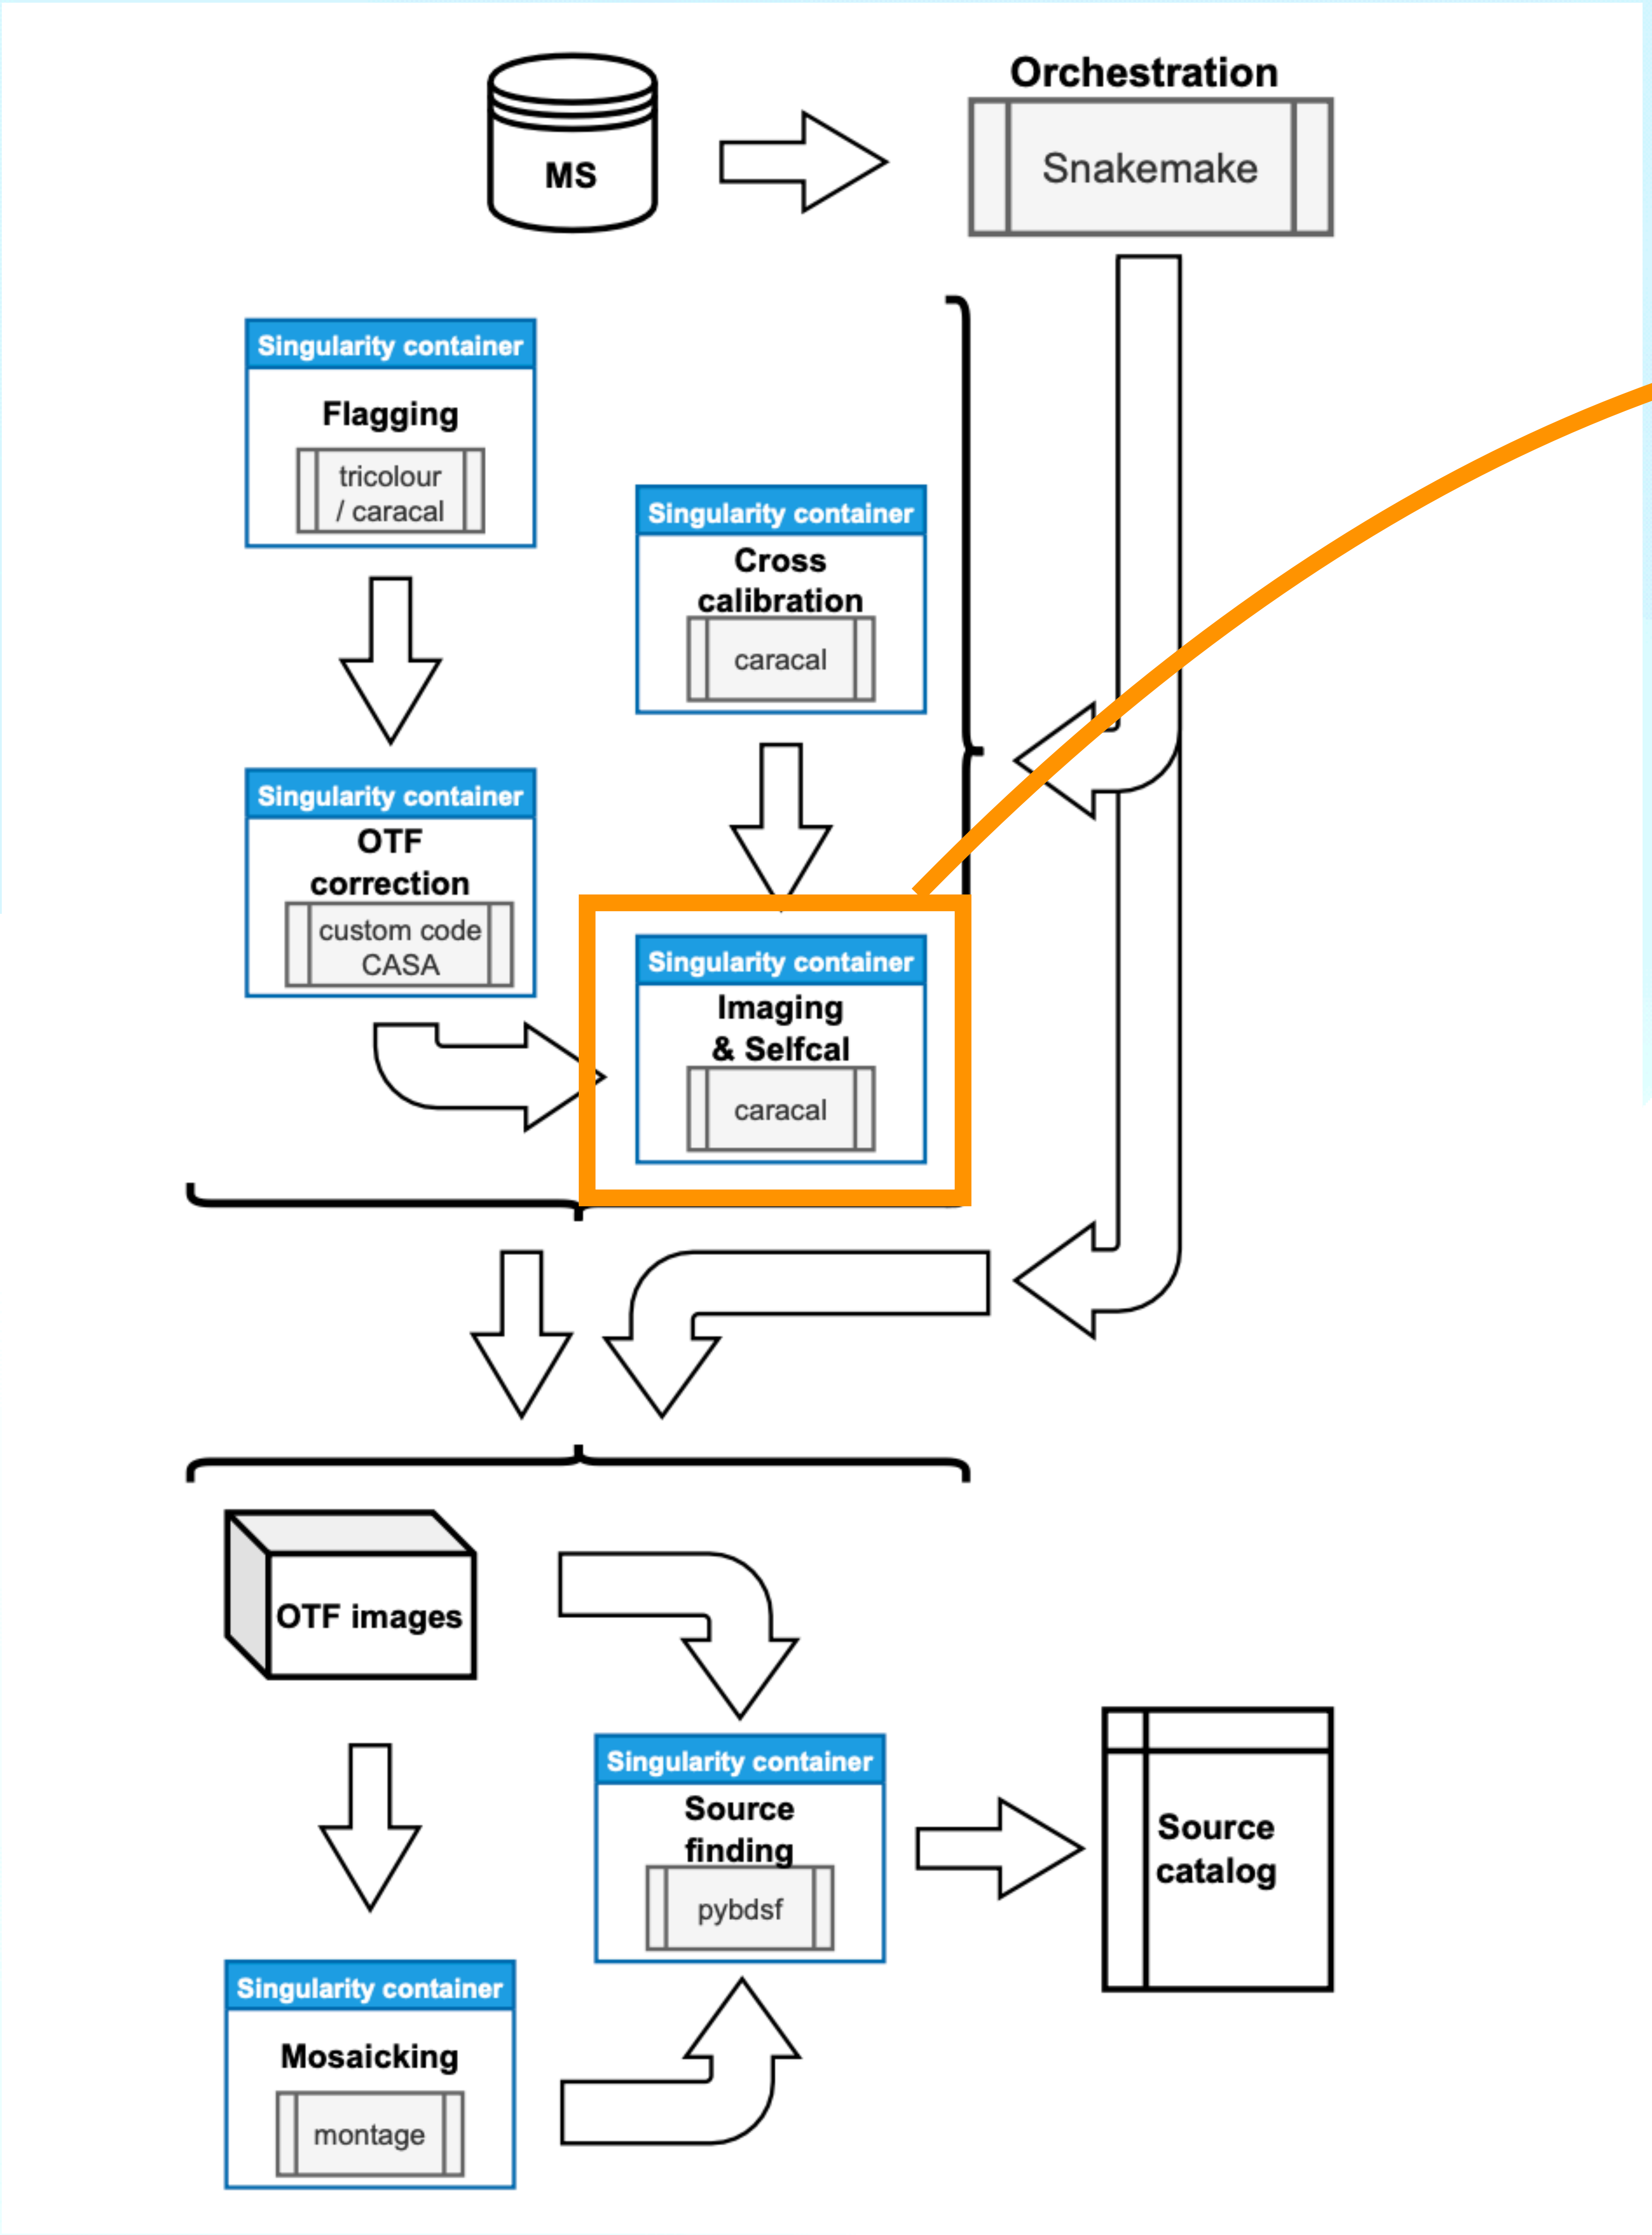

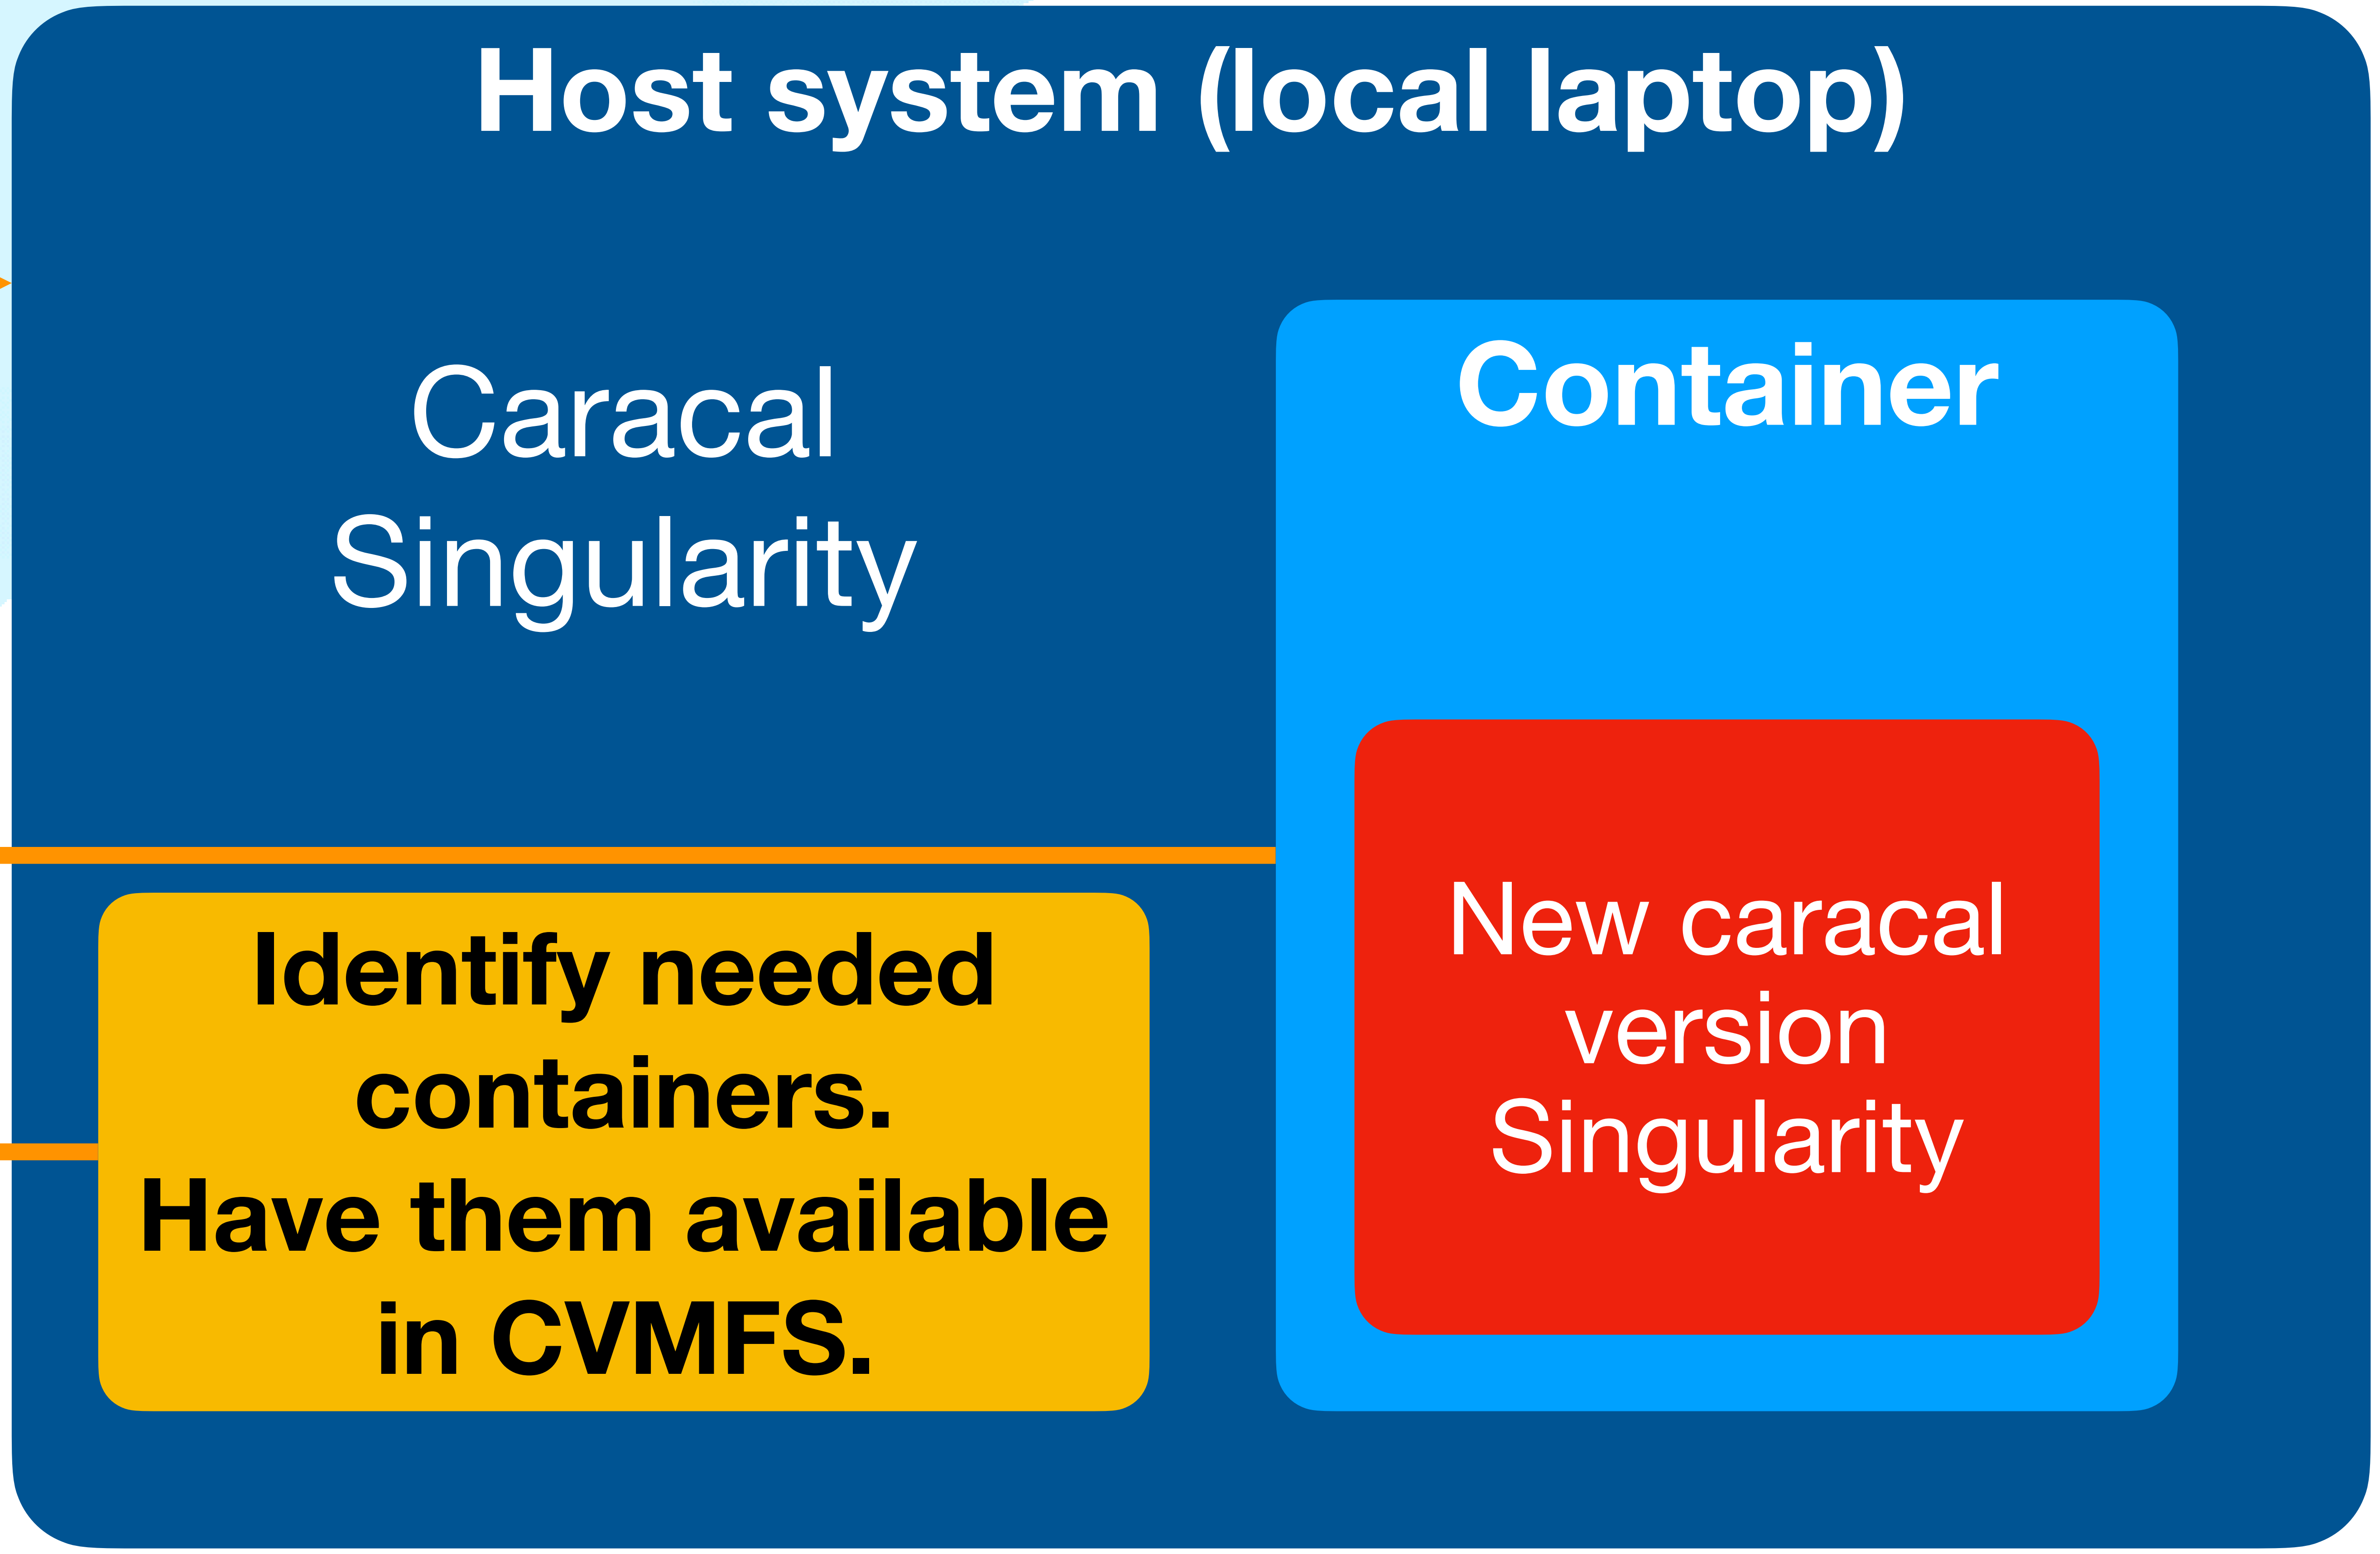

- Provide to Compute4PUNCH.
- Caracal in container uses containerized singularity to access pre-pulled cabs in CVMFS.

Achievements - Data management

In its current implementation, Storage4PUNCH and Compute4PUNCH are disconnected. Compute4PUNCH does not have a shared file system.

This presented a problem for running the pipeline. It is data intensive and data need to be copied to and from the compute nodes each time!

BUT… it spurred new work to include a new cluster (at LMU, hosting the data) into Compute4PUNCH.

Dynamic integration of a new node in C4P successful. Provided feedback to C4P developers and acted as a pathway for new clusters to be added in the future.

Conclusions

- PUNCH4NFDI is a new consortium aiming at creating infrastructure and services to improve scientific research in Germany
- Compute4PUNCH and Storage4PUNCH are federated computing and storage solutions which require new ways to implement workflows
- Implementation of complex pipelines on this new infrastructure is promising
- Stay tuned for future results