



This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 777533.

PROviding Computing solutions for ExaScale Challenges

D6.2	Second prototype with Demonstration using chosen pilot application		
Project :	PROCESS H2020 – 777533	Start / Duration:	01 November 2017 36 Months
Dissemination¹:	PU	Nature²:	R+O
Due Date:	Month 24	Work Package:	WP 6
Filename³	PROCESS_D6.2_ SecondPrototype_v1.0.docx		

ABSTRACT

The second prototype of the PROCESS platform presents the extensions made to the IEE and related components required to enable the execution of the actual applications provided by the Use Case owners. In this document we are demonstrating the prototype with two applications, one running on HPC (with multi-site support) and one targeting the Cloud environment.

¹ PU = Public; CO = Confidential, only for members of the Consortium (including the EC services).

² R = Report; R+O = Report plus Other. Note: all "O" deliverables must be accompanied by a deliverable report.

³ eg DX.Y_name to the deliverable_v0xx. v1 corresponds to the final release submitted to the EC.

Deliverable Contributors:	Name	Organization	Role / Title
Deliverable Leader⁴	Hluchý, Ladislav	UISAV	Coordinator
Contributing Authors⁵	Bubak, Marian	AGH / AGH-UST	Writers
	Meizner, Jan	AGH / AGH-UST	Writers
	Nowakowski, Piotr	AGH / AGH-UST	Writers
Reviewer(s)⁶	Guggemos, Tobias	LMU	Reviewer
	Habala, Ondrej	UISAV	Reviewer
Final review and approval	Höb, Maximilian	LMU	Coordinator

Document History

Release	Date	Reasons for Change	Status⁷	Distribution
0.1	2019-09-19	Initial version	Draft	2019-09-19
0.2	2019-09-20	Basic structure	Draft	2019-09-20
0.3	2019-10-04	Initial content	Draft	2019-10-04
0.4	2019-10-18	Added overview and theoretical descriptions	Draft	2019-10-18
0.5	2019-10-28	Added demonstration details	Draft	2019-10-28
0.6	2019-10-29	Internal review finished	In Review	2019-10-28
1.0	2019-10-31	Final version following internal review	Released	2018-10-31

⁴ Person from the lead beneficiary that is responsible for the deliverable.

⁵ Person(s) from contributing partners for the deliverable.

⁶ Typically person(s) with appropriate expertise to assess the deliverable quality.

⁷ Status = "Draft"; "In Review"; "Released".

Table of Contents

Executive Summary	4
List of Figures.....	5
1 Overview	6
2 Current status of the PROCESS platform	6
2.1 Implementation approach	6
2.2 PROCESS development environment	8
2.3 PROCESS from the user perspective	8
2.4 The PROCESS IEE workbench – structure and operations.....	9
3 Chosen pilot application.....	10
4 Additional applications tested with the prototype	11
5 Conclusion	11
6 Future plans	11
References	12
Appendix A: Demonstration report.....	13
Demonstration of the LOFAR pilot application.....	13
Demonstration of the ancillary pricing for airline revenue management application.....	16

Executive Summary

This deliverable follows up on the work described in the scope of the previous WP6 deliverable – namely, D6.1 [D6.1]. As before we are showing the software platform and integrated set of services building the PROCESS infrastructure, however in this document we are focusing on extensions that made running the real Use Case provided code possible.

The document is structured as follows:

- Section 1 provides a general overview of the current status of the Work Package 6, and serves as an introduction to subsequent sections.
- Section 2 highlight extensions of the platform components made to the platform to provide the functionality needed by the Second Prototype.
- Section 3 briefly presents the application provided by the Use Case 2 acting as the pilot application for the demonstration described in this document.
- Section 4 showcases the second application shown during the demonstration focusing on the Cloud capabilities of the platform – provided by the Use Case 4.
- Section 5 contains closing remarks.
- Section 6 discusses future plans related to development of the Interactive Execution Environment and integration with other components of the PROCESS infrastructure.
- Appendix A shows details of both demonstrations.

List of Figures

<i>Figure 1: Environment used for the second prototype deployment</i>	<i>7</i>
<i>Figure 2: Accessing the PROCESS Interactive Execution Environment - login screen</i>	<i>9</i>
<i>Figure 3: IEE user workbench.....</i>	<i>9</i>
<i>Figure 4: Overview of the LOFAR application pilot application</i>	<i>10</i>
<i>Figure 5: Ancillary pricing for airline revenue management in the Cloud demo.....</i>	<i>11</i>
<i>Figure 6: New LOFAR project created.....</i>	<i>13</i>
<i>Figure 7: Creating the new LOFAR pipeline step.....</i>	<i>14</i>
<i>Figure 8: Starting the LOFAR computation</i>	<i>14</i>
<i>Figure 9: Computation queued.....</i>	<i>15</i>
<i>Figure 10: Running the computation on the HPC.....</i>	<i>15</i>
<i>Figure 11: Finished LOFAR pipeline.....</i>	<i>16</i>
<i>Figure 12: Setting up a new project for the UC#4 pipeline.....</i>	<i>16</i>
<i>Figure 13: Configuring a new Cloudify pipeline.....</i>	<i>17</i>
<i>Figure 14: Starting the UC#4 pipeline.....</i>	<i>17</i>
<i>Figure 15: Running the UC#4 pipeline.</i>	<i>18</i>
<i>Figure 16: The UC#4 pipeline finished</i>	<i>18</i>

1 Overview

The goal of this document is to show the second prototype of the PROCESS platform including progress beyond the state presented in the D6.1 [D6.1] deliverable that shown the first prototype. The work described here still follows the basic architecture design by the Work Package 4 with necessary updates as described in WP4 deliverables. The main focus is obviously on the new elements of the Compute Platform implemented in the scope of Work Package 6, but also on necessary integration with Data Services developed in the scope of Work Package 5 and Orchestration platform developed in the scope of Work Package 7. As in the case of the first prototype the second one enables running applications supported by the usage scenario identified in the scope of Work Packages 2 and 3 (described in the D2.1 [D2.1]), with special focus on demonstration of the Square Kilometre Array SKA code (Use Case #2).

The initial prototype of the PROCESS infrastructure implemented key features and aspects of the platform's functionality and was intended for validation based on the basic use case requirements [D4.1]. Following up on that work we have significantly extended the Compute Platform, as well as integrated augmented versions of the Storage and Orchestration platforms to build a viable second prototype capable of running chosen pilot application on the fully-fledged PROCESS infrastructure. The first prototype features the following components:

- Rimrock component extended to support multi-site deployment capable of scheduling jobs on Prometheus (Cyfronet) and CoolMUC (LRZ);
- Data services upgraded to include capability to move data between original source location and multiple sites mentioned above;
- Set of orchestration templates for the Cloudify tool capable of scheduling codes on the Compute Cloud;
- Augmented version of the Interactive Execution Environment (IEE) integrated with the above tools and providing User Interface enabling control of the new features (such as site selection).
- Monitoring and benchmarking platform integrated with all platform components mentioned above.

The following sections describe the chosen pilot application provided by the Use Case developers as well as the current status of the PROCESS platform implementing the features described in this section.

2 Current status of the PROCESS platform

2.1 Implementation approach

The Second Prototype architecture follows the original one described in D4.2 [D4.2] and used for the creation of the first prototype; however minor tweaks have been applied to it as decided by WP4 to further reduce the overhead of the platform.

As a result, the environment has been integrated as depicted in Figure 1. This design still features IEE as the entry point for the platform. Rimrock has been augmented to support multiple HPC sites. To reduce platform overhead Atmosphere was replaced with a direct link between IEE and Cloudify which enables scheduling of the cloud VMs from it.

D6.2: Current status of the PROCESS platform

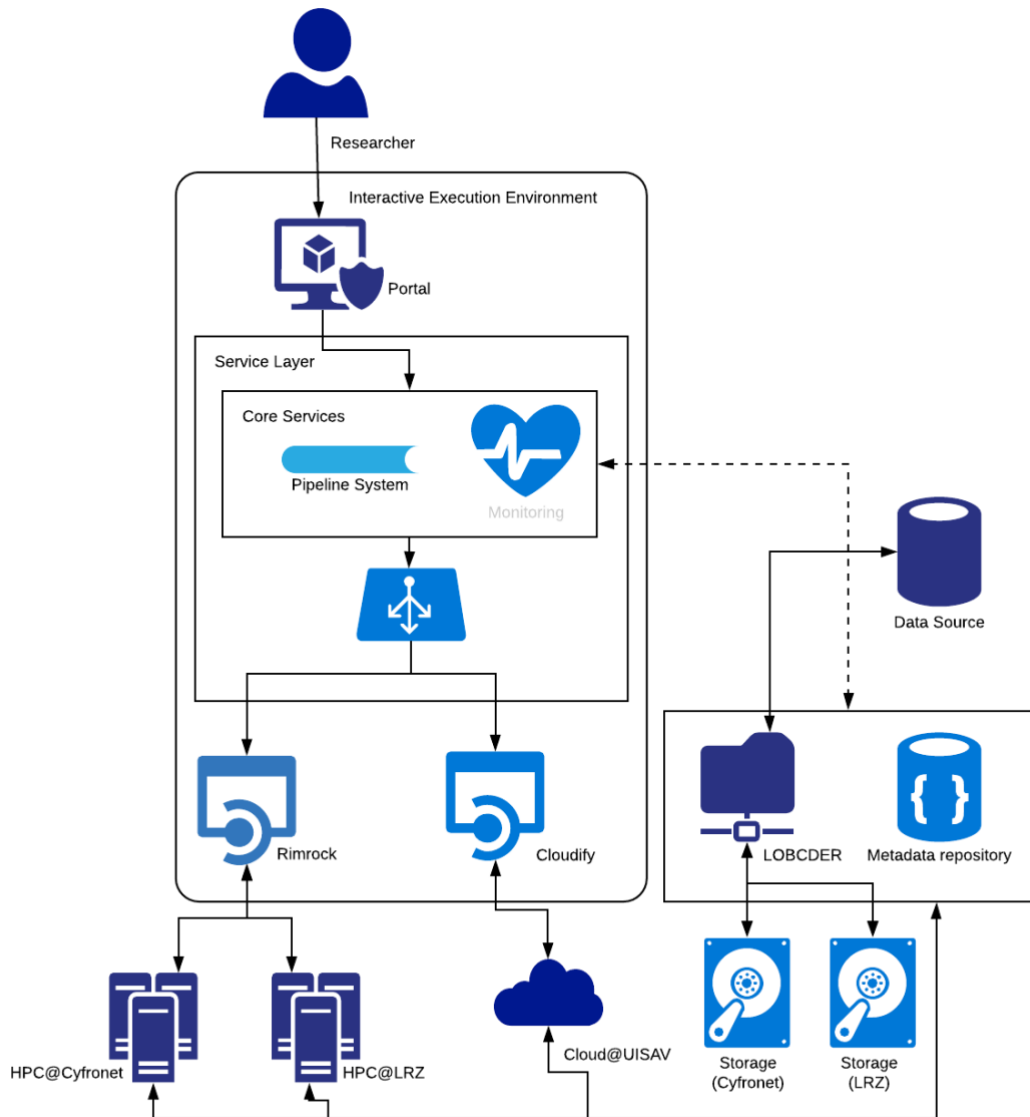


Figure 1: Environment used for the second prototype deployment

For the second demonstrator we have decided to implement the full pipeline for the Use Case #2 processing the LOFAR dataset [D2.1]. This enables the means to select the HPC site and then facilitate process of staging in data from the LOFAR Long Term Archives and moving it to the relevant site, running relevant computation using site-specific settings (Queuing and Container systems) and finally staging out the results.

For this case we've decided to utilize two separate HPC sites – Prometheus supercomputing cluster at ACC Cyfronet AGH in Kraków (same as for the first prototype), as well as a new site, namely the CoolMUC Cluster at LRZ in Garching/Munich. In addition to those HPC sites, to showcase the full capabilities of the platform, we also included a demo of the additional application from the UC #4 (Ancillary pricing for airline revenue management [D2.1]) deployed on the OpenStack cloud at UISAV in Bratislava via the Cloudify component.

D6.2: Current status of the PROCESS platform

To prepare for this demonstration we have:

- Extended the code of the Interactive Execution Environment to enable selection of the Compute Site – including providing Site-specific configuration.
- Support for Charliecloud container framework (used at LRZ) was added to enable running containers as HPC jobs on CoolMUC
- Adaptation of Slurm and other HPC runtime scripts
- Integration with the Cloud platform via Cloudify was performed and tested
- An augmented version of IEE was deployed and tested for the whole workflow in multiple scenarios composed of both data transfer and compute steps.

2.2 PROCESS development environment

The PROCESS development environment is the same as for the First Prototype including:

- The production environment available at <https://process.cyfronet.pl>
- The development environment at <https://process-dev.cyfronet.pl>

The augmented Second Prototype environment, of course, provides the same level of access to the PLGrid [PLGrid] services at Cyfronet, as the First Prototype (based on standard PLGrid accounts), however at the same time, it allows running code at three sites: two batch HPC sites and one cloud site. Access to those additional systems is enabled by integration of IEE with the security systems enforced at each site.

2.3 PROCESS from the user perspective

As in the case of the demo of the first prototype, in order to acquire access to the PROCESS IEE, each prospective member must first register with the PL-Grid project, either as a Polish scientist or a foreign collaborator thereof. This is done by creating an account at <https://portal.plgrid.pl> and requesting to join the PROCESS user team. Details of the registration process are out of scope of this deliverable; however, ACC Cyfronet AGH, in its capacity as a Project partner, provides ongoing assistance and help in case of registration problems, particularly for members of use case teams. We are also providing solution to delegate access to other HPC and Cloud Sites so access to IEE is sufficient for the whole PROCESS platform.

Once an account has been set up, it can be used to access the PROCESS development IEE at <https://process-dev.cyfronet.pl>, as illustrated in Figure 2.

D6.2: Current status of the PROCESS platform

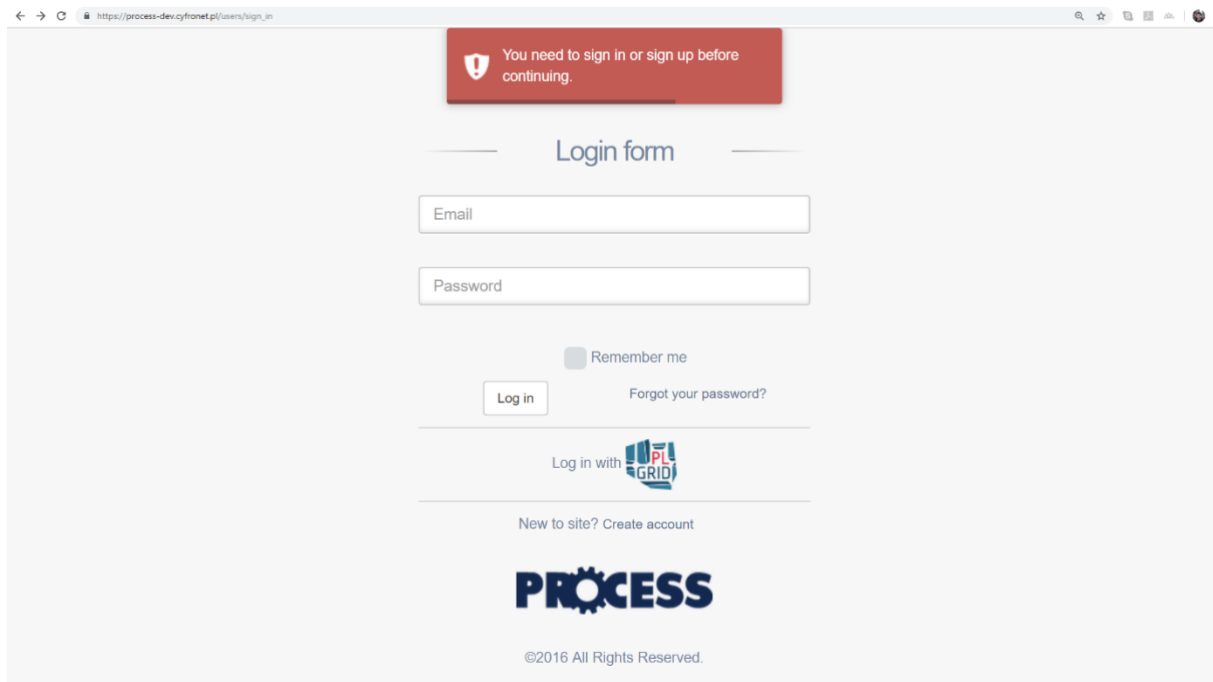


Figure 2: Accessing the PROCESS Interactive Execution Environment - login screen

Rather than directly typing account credentials in the provided form (which is a development-only feature meant to facilitate access for platform maintainers), the user should instead select the “Log in with PL-Grid” option, which will automatically redirect them to the PL-Grid login site as described in more detail previously in D6.1.

2.4 The PROCESS IEE workbench – structure and operations

The basic view of the IEE user workbench, which is presented to the user upon login, is shown in Figure 3.

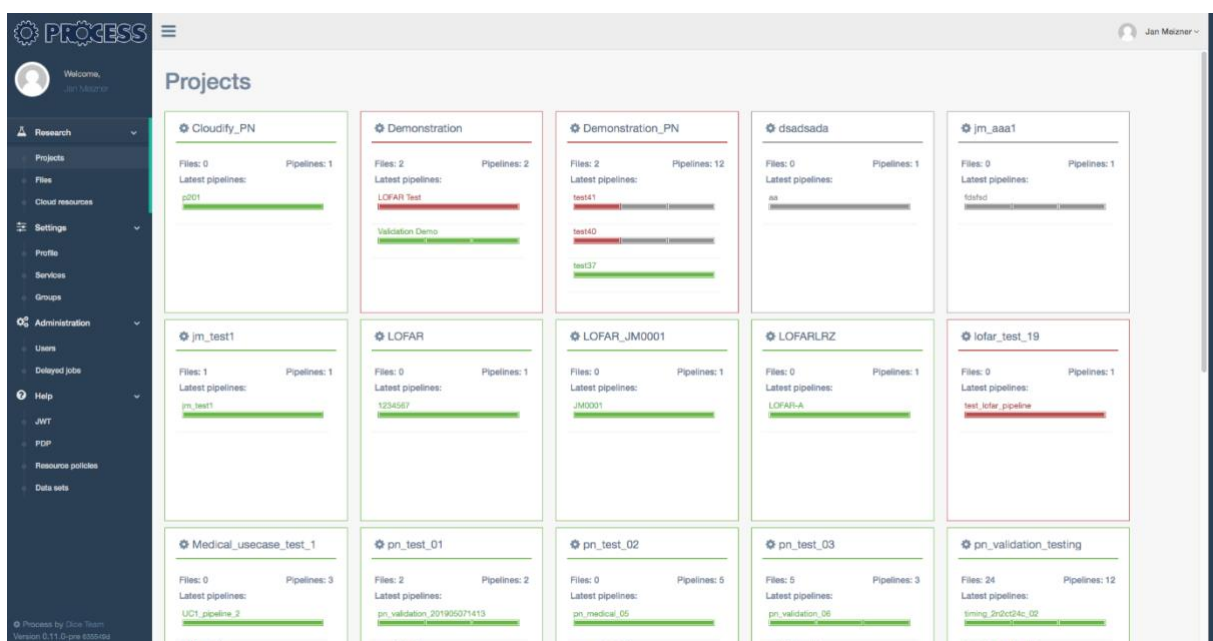


Figure 3: IEE user workbench

The basic concept of the PROCESS Interactive Execution Environment have not change since the First Prototype – and as such was described in detail in the D6.1. Use Case applications are still organized as projects composed of multiple steps – either computational or related to data transfer.

For the need of the second prototype the IEE has been augmented to:

- support Cloud infrastructures via integration with Cloudify,
- allow running HPC computations on additional sites in addition to Prometheus,
- support additional credentials other than PLGrid proxy certificates – such as the SSH key-pairs for access to non-PLGrid infrastructures,
- provide a GUI for the augmentations mentioned above

3 Chosen pilot application

We have chosen the LOFAR application developed in scope of Use Case 2 as the pilot application for the second prototype. This application requires significant resources to process large amounts of data representing observations stored on tapes in the Long Term Archive (LTA). During the demonstration of the prototype, we run a pipeline which utilizes an integrated PROCESS environment, including data stage-in, processing and stage-out. A schematic overview of the demonstration is shown in Figure 4.

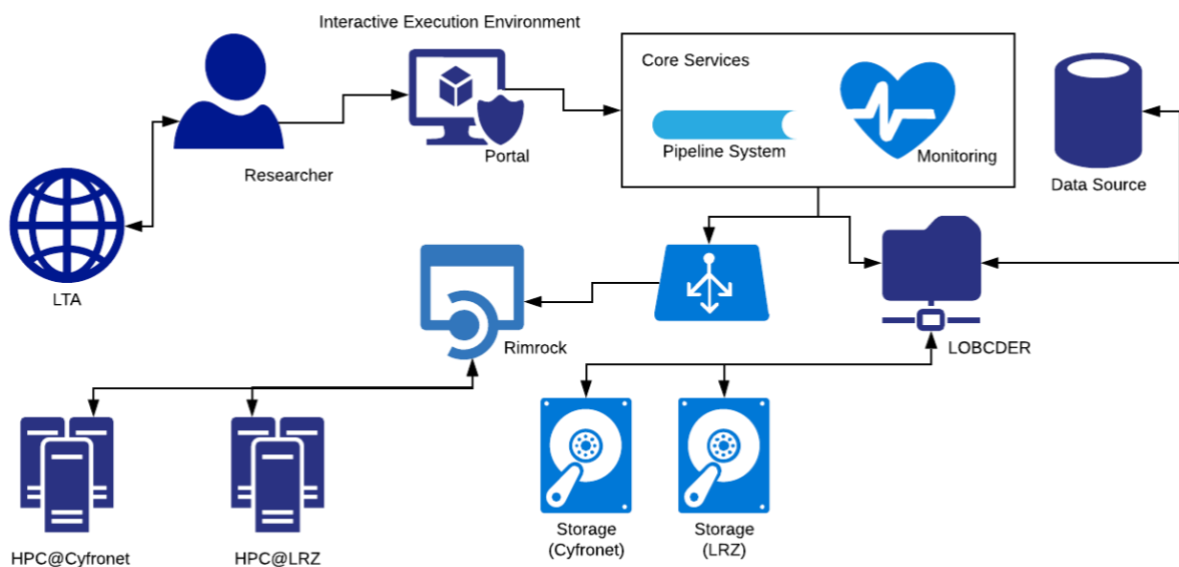


Figure 4: Overview of the LOFAR application pilot application

4 Additional applications tested with the prototype

In addition to the pilot application described in the Section 3 the demonstration includes additional pipeline featuring ancillary pricing for “airline revenue management” application. This application is built as a set of Docker containers running in the cloud. During the demonstration we utilize integration of the PROCESS IEE with the private cloud platform in UISAV, Bratislava, via the Cloudify component. Our goal is to showcase completeness of the platform featuring both HPC and Cloud resources. A schematic overview of the demo is shown in Figure 5.

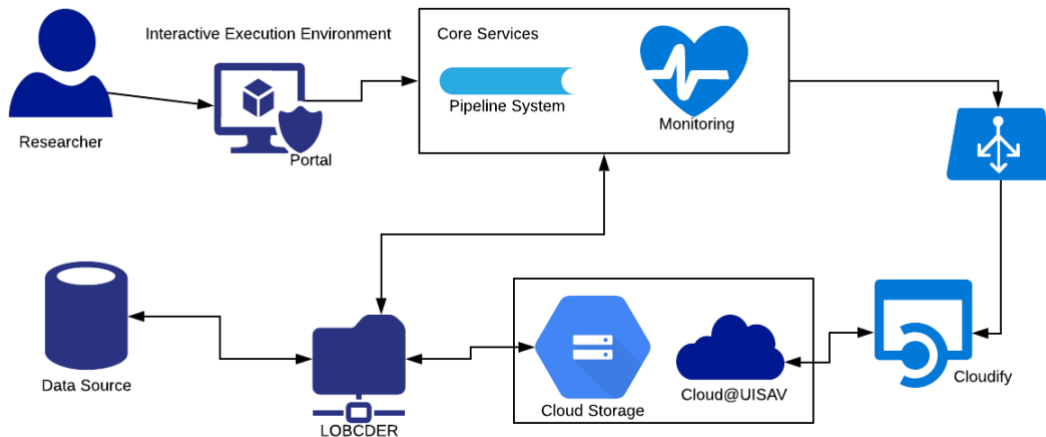


Figure 5: Ancillary pricing for airline revenue management in the Cloud demo

5 Conclusion

The platform presented in this document supports a much wider range of cases than the first prototype described in D6.1. In particular, the ability to run containerized applications as part of HPC pipelines has been extended to support multiple HPC Sites as well as other container technologies. Regardless of this heterogeneity which is inherent for independent IT systems operated by multiple organizations in different EU Member States, our platform still provides a unified experience with a single endpoint via the IEE. Moreover, in addition to HPC systems, we now also supporting cloud environments, with standardized deployment templates (TOSCA), also managed from the IEE.

6 Future plans

The second prototype featured integration of all components of the PROCESS platform both in the aspect of data transfer and computations performed on this data on various types of infrastructures. The platform was tested using the selected pilot application.

Following this prototype, we will focus on further solidifying the platform and testing it on various applications provided by all Use Cases present in the Project. We will also work on making the platform even more straightforward for all groups of users, even those less versed in the operations of the HPC and Cloud systems.

D6.2: References

The implementation and integration effort will now focus on safeguarding the common goal of the Project to produce the first packaged version of the software which is going to be described in D8.3 (due in 33rd month of the Project) as well as the final release of the platform which is going to be described in the scope of D8.4 (due in the 36th month of the Project).

References

[D2.1] PROCESS Deliverable D2.1, Progress Report (UC#1-5); internal document

[D4.1] PROCESS Deliverable D4.1, Initial state of the art and requirement analysis; internal document

[D4.2] PROCESS Deliverable D4.2, Report on Architecture Evaluation; internal document

[D6.1] PROCESS Deliverable D6.1, First prototype; internal document

[PLGrid] The PL-Grid National Computing Infrastructure, www.plgrid.pl.

Appendix A: Demonstration report

We have run the demonstration scenarios as described above – the results with relevant screencast is summarized below.

Demonstration of the LOFAR pilot application

The demonstration scenario involves running computation using the container provided by the Use Case 2. In the beginning, similarly to the First Prototype, we create a new project, as shown in Figure 6.

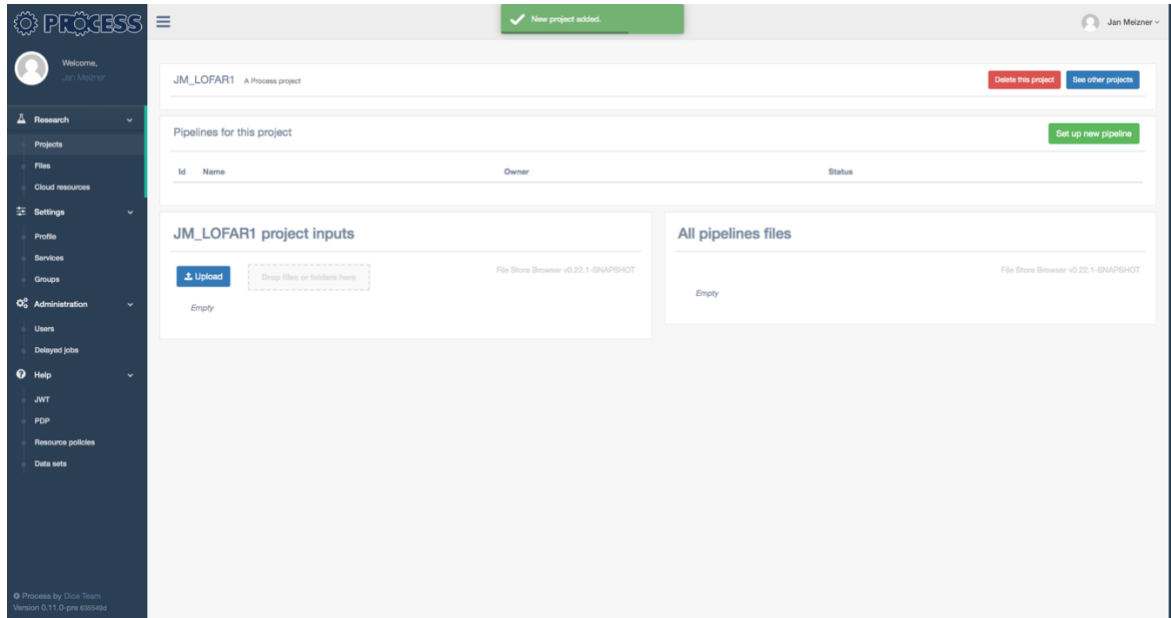


Figure 6: New LOFAR project created

Next, we need to launch a new PROCESS pipeline, providing the required information specific for this application, such as:

- Name of the container containing the most recent release of the Use Case code as defined by its builder,
- Selected HPC Site – e.g. Prometheus or CoolMUC
- Various parameters specific for the LOFAR computation, such as Visibility ID

D6.2: Appendix A: Demonstration report

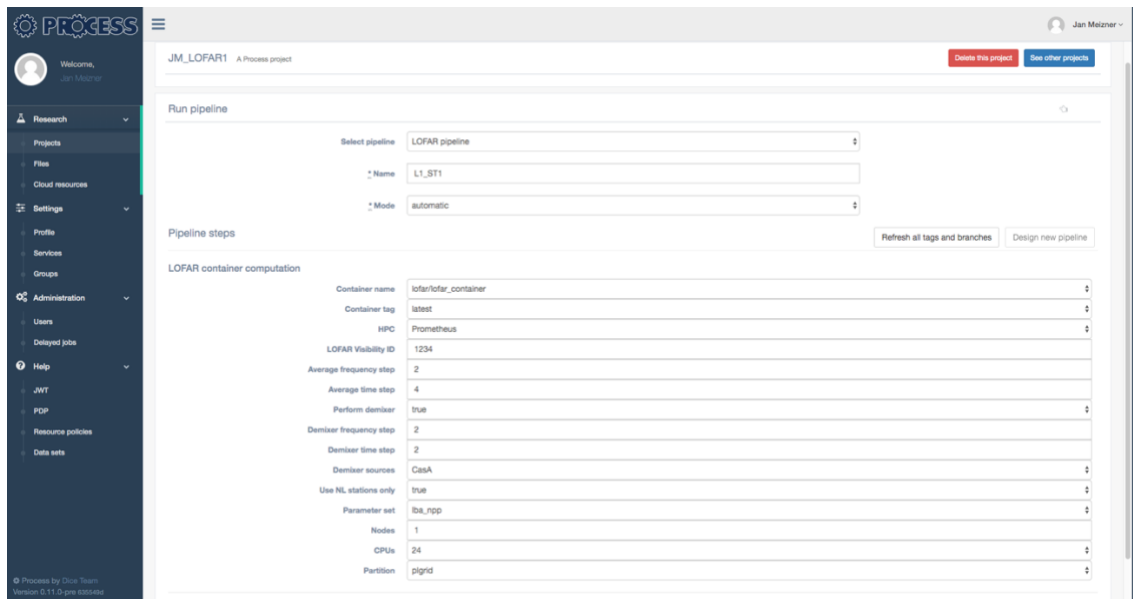


Figure 7: Creating the new LOFAR pipeline step

Once the computation is defined it will go through successive phases:

- Setup – computation created but not yet scheduled – see Figure 8
- Queued – waiting for execution on the HPC, as shown in Figure 9
- Running – computation has started on the HPC, as shown in Figure 10
- Finished – Computation has finished, as shown in Figure 11

As shown in Figure 11, once the computation has finished (which, in this instance, took 58 minutes and 9 seconds), outputs are staged out as a tar file.

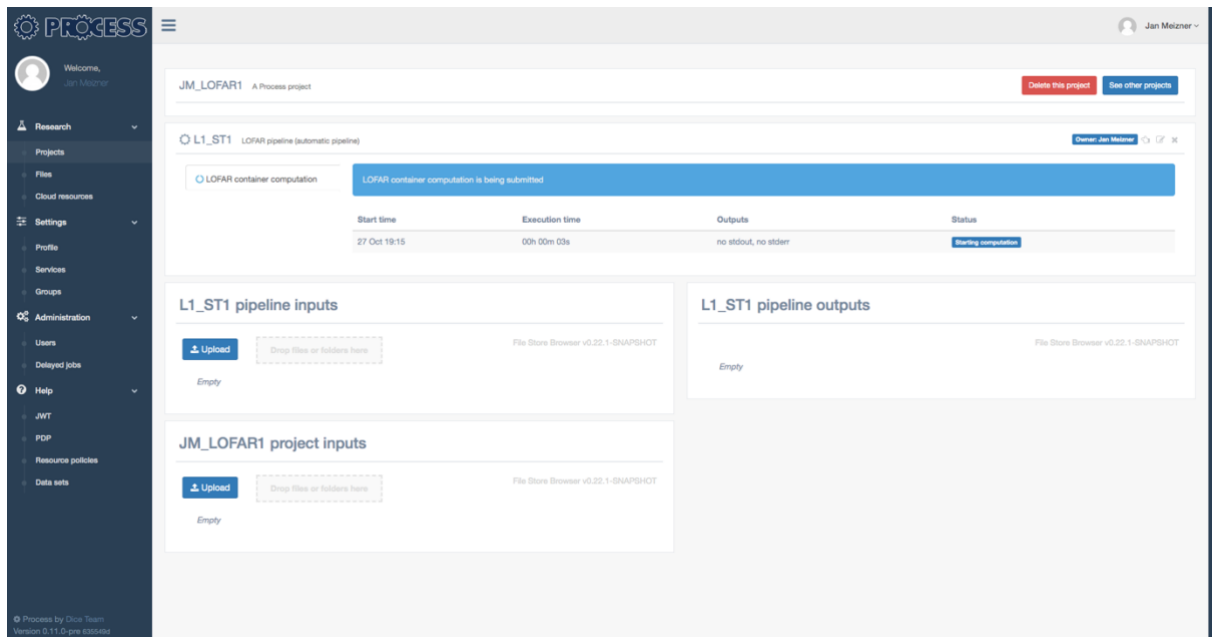


Figure 8: Starting the LOFAR computation

D6.2: Appendix A: Demonstration report

The screenshot displays the Process web interface. On the left is a dark sidebar with navigation links: Research, Projects, Files, Cloud resources, Settings, Profile, Services, Groups, Administration, Users, Delayed jobs, Help, JWT, PDP, Resource policies, and Data sets. The main content area shows the 'JM_LOFAR1' project. A blue progress bar at the top indicates 'LOFAR container computation is queued'. Below this, a table shows the start time as '27 Oct 19:15', execution time as '00h 00m 10s', outputs as 'stdout, stderr', and status as 'Queued'. Further down, there are sections for 'L1_ST1 pipeline inputs', 'JM_LOFAR1 project inputs', and 'L1_ST1 pipeline outputs', each with an 'Upload' button and a 'Drop files or folders here' area. The footer of the sidebar indicates '© Process by DICE Team Version 0.11.0-pre released'.

Figure 9: Computation queued

This screenshot shows the same Process web interface as Figure 9, but the computation status has changed. The blue progress bar now says 'LOFAR container computation is running'. The table below it shows the execution time as '00h 00m 22s' and the status as 'Running'. The input and output sections remain the same, with 'Upload' buttons and 'Drop files or folders here' areas. The sidebar and footer are identical to the previous figure.

Figure 10: Running the computation on the HPC

D6.2: Appendix A: Demonstration report

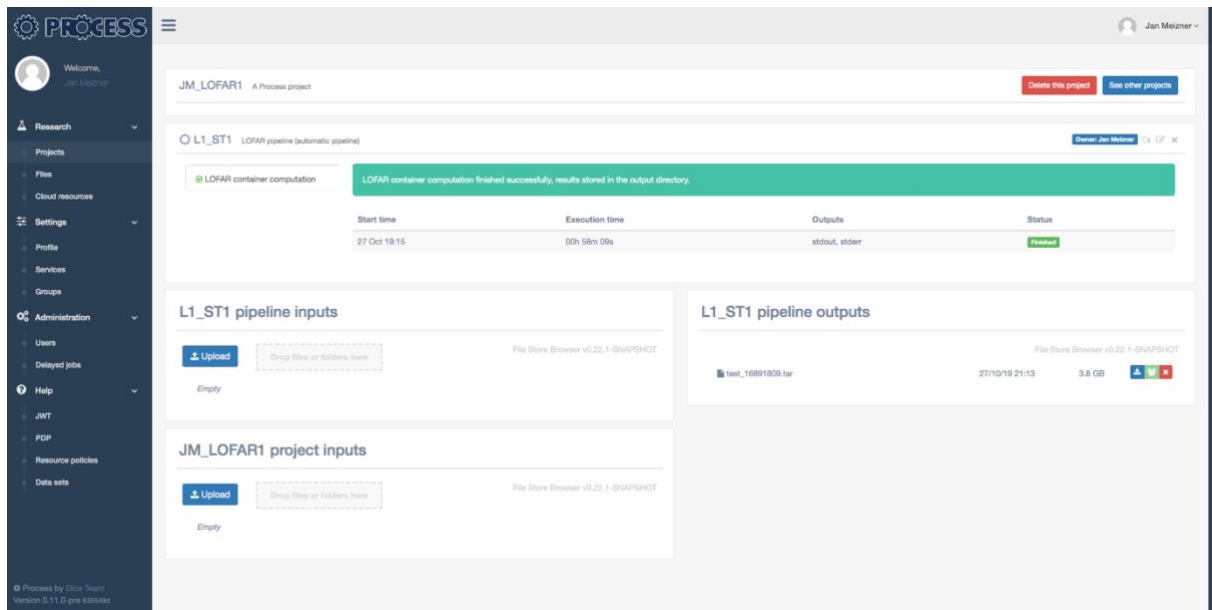


Figure 11: Finished LOFAR pipeline

Demonstration of the ancillary pricing for airline revenue management application

In addition to the demonstration of the pilot application shown above, we have prepared a demonstration of Use Case 4, featuring the usage of the cloud part of the PROCESS platform. As in the previous case, we begin by creating a dedicated project for this demonstration, as shown in Figure 12.

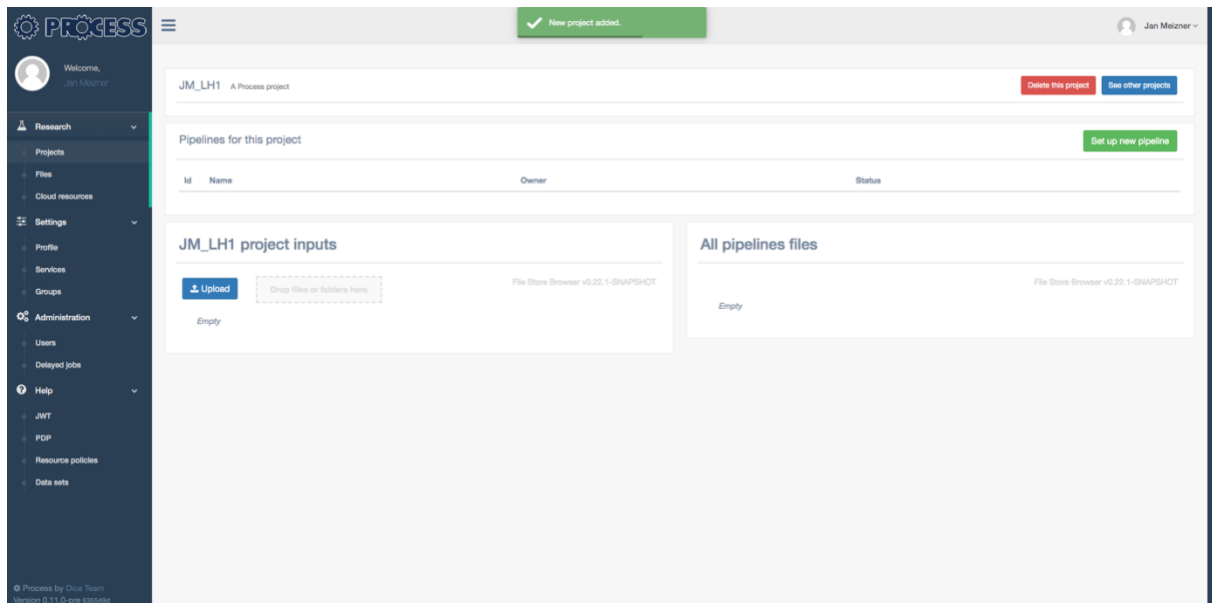


Figure 12: Setting up a new project for the UC#4 pipeline

The computation is already Defined in the form of TOSCA template for Cloudify, so no additional setup is needed. The IEE platform automatically handles the task of asking Cloudify to construct the necessary deployment and, thereupon, execute the code of the application within the environment set up for this purpose. Once processing is complete, IEE also instructs

D6.2: Appendix A: Demonstration report

Cloudify to tear down the deployment to avoid consuming resources which are no longer needed. Having provided standard settings common for all pipelines (name and mode), we may launch the pipeline, as shown in Figure 13.

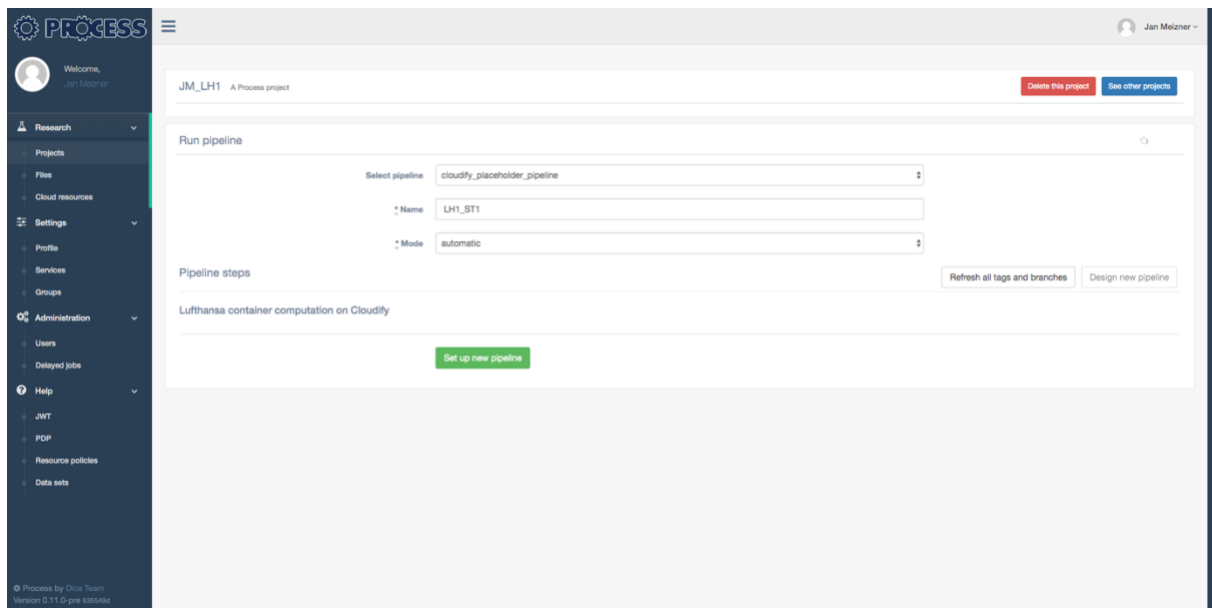


Figure 13: Configuring a new Cloudify pipeline.

Once launched, the computation proceeds through the following phases:

- Setup – before the computation is scheduled – see Figure 14
- Running – directly afterwards, as there is no queue in the cloud – see Figure 15
- Finished – all work is done – see Figure 16

As shown in Figure 16, the computation successfully completes in 22 minutes and 42 seconds.

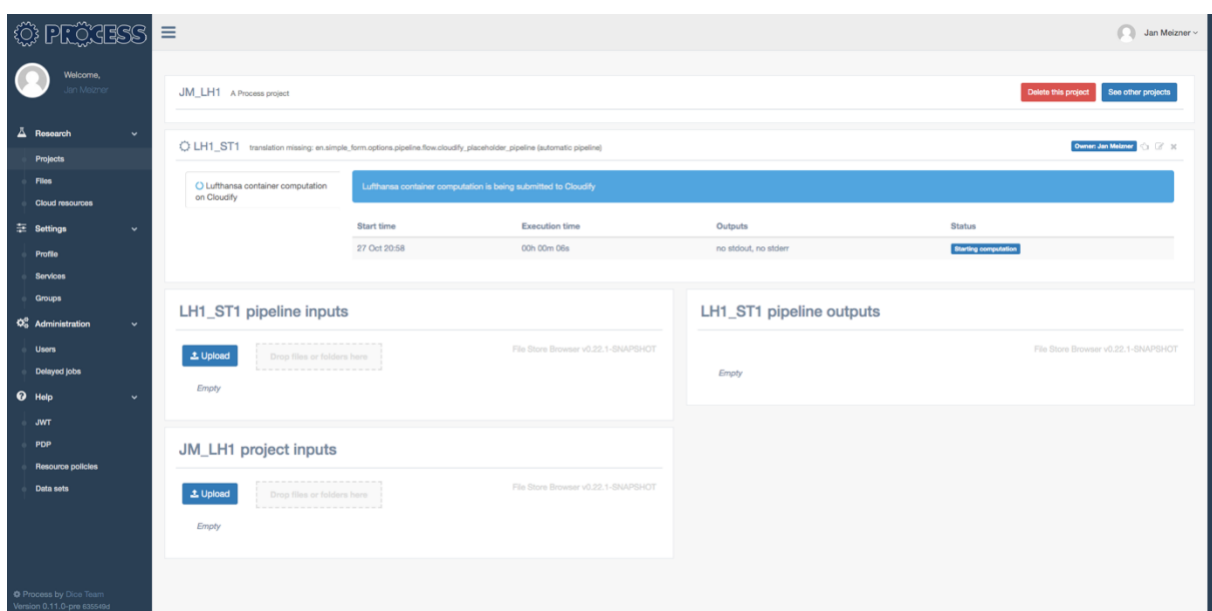


Figure 14: Starting the UC#4 pipeline

D6.2: Appendix A: Demonstration report

The screenshot shows the PROCESSTOOL web interface. On the left is a dark sidebar with navigation links: Research, Projects, Files, Cloud resources, Settings, Profile, Services, Groups, Administration, Users, Delayed jobs, Help, JWT, PDP, Resource policies, and Data sets. The main content area displays the 'JM_LH1' project. A blue progress bar indicates 'Lufthansa container computation is running in Cloudify'. Below this, a table shows the execution details:

Start time	Execution time	Outputs	Status
27 Oct 20:58	00h 00m 16s	no stdout, no stderr	Running

Below the table are sections for 'LH1_ST1 pipeline inputs', 'LH1_ST1 pipeline outputs', and 'JM_LH1 project inputs', each with an 'Upload' button and a 'Drop files or folders here' area. The status is 'Empty' for all.

Figure 15: Running the UC#4 pipeline.

The screenshot shows the same PROCESSTOOL interface, but the pipeline has completed. The progress bar is now green and says 'Lufthansa container computation finished successfully, results stored in the output directory.' The table below shows the final status:

Start time	Execution time	Outputs	Status
27 Oct 20:58	00h 22m 42s	no stdout, no stderr	Finished

The input and output sections remain empty.

Figure 16: The UC#4 pipeline finished