

PROCESS

Use of the HPC Containers in the way towards exascale

Jan Meizner¹, Marian Bubak^{1,2}, Jan Kapała¹, Piotr Nowakowski¹,
Patrik Wójtowicz¹

¹ACC Cyfronet AGH, Kraków, Poland



²Department of Computer Science, AGH University of Science and Technology

<http://dice.cyfronet.pl/>



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



Lufthansa Systems



inmark
europa



ÚSTAV INFORMATIKY
SLOVENSKÁ AKADEMIA VIED



- Motivation and objectives
- Project Use Cases
- Road Towards Exascale
- Why Singularity?
- PROCESS Platform Architecture
- Singularity Step Implementation
- Conclusions and Further Plans

The PROCESS project aims to:

- Pave the way towards exascale by providing scalable platform
- Enable deployment of services on heterogenous infrastructures
- Support different domains of science and business

Our objective is to:

- Build the Container based platform based on Singularity
- Integrate the HPC resources across multiple countries
- Provide effortless user experience via the WebUI

Project Use Cases

Health

- **Content-based search and classification**
- HAUTE ECOLE SPECIALISEE DE SUISSE OCCIDENTALE (CH)

Astronomy

- **Square Kilometre Array SKA**
- STICHTING NETHERLANDS ESCIENCE CENTER (NL)

Risk Management

- **Supporting innovation based on global disaster risk data**
- LUDWIG-MAXIMILIANS-UNIVERSITAET MUENCHEN (DE)

Airlines Operations

- **Ancillary pricing for airline revenue management**
- LUFTHANSA SYSTEMS GMBH & CO KG (DE)


Agriculture

- **Agro-Copernicus**
- LUDWIG-MAXIMILIANS-UNIVERSITAET MUENCHEN (DE)

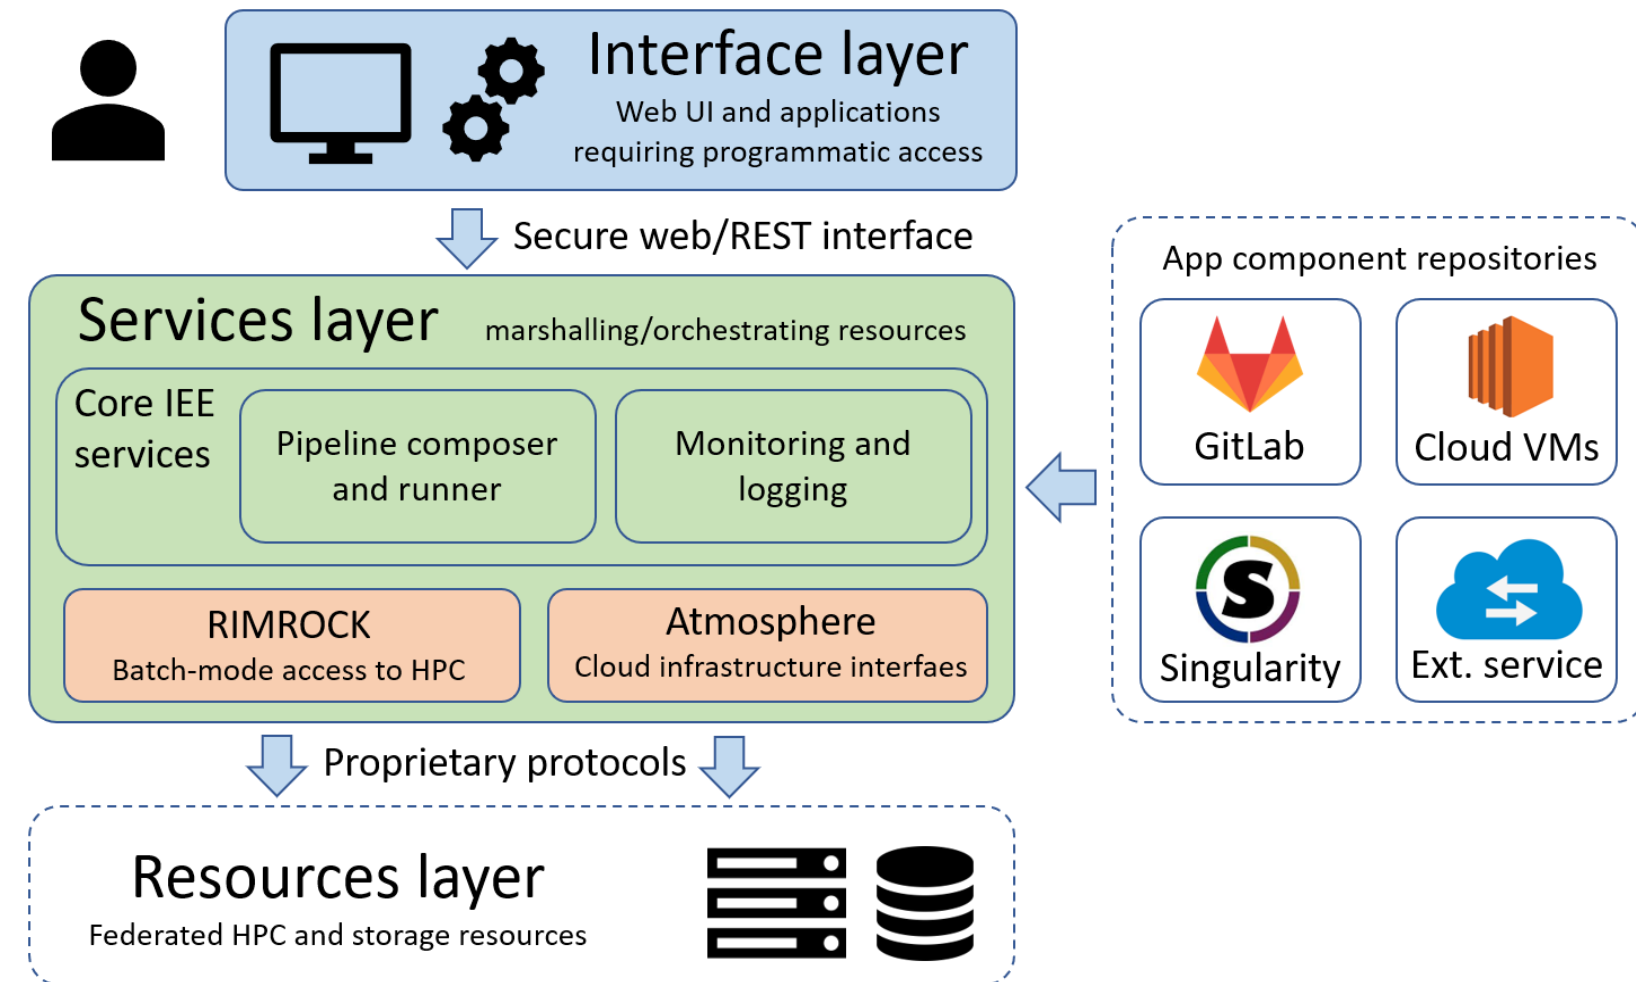
Rank	Site	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	DOE/SC/Oak Ridge National Laboratory United States	Summit - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband IBM	2,282,544	122,300.0	187,659.3	8,806
2	National Supercomputing Center in Wuxi China	Sunway TaihuLight - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway NRCP	10,649,600	93,014.6	125,435.9	15,371
3	DOE/NNSA/LLNL United States	Sierra - IBM Power System S922LC, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband IBM	1,572,480	71,610.0	119,193.6	
4	National Super Computer Center in Guangzhou China	Tianhe-2A - TH-IVB-FEP Cluster, Intel Xeon E5-2692v2 12C 2.2GHz, TH Express-2, Matrix-2000 NUDT	4,981,760	61,444.5	100,678.7	18,482
5	National Institute of Advanced Industrial Science and Technology (AIST) Japan	AI Bridging Cloud Infrastructure (ABCI) - PRIMERGY CX2550 M4, Xeon Gold 6148 20C 2.4GHz, NVIDIA Tesla V100 SXM2, InfiniBand EDR	391,680	19,880.0	32,576.6	1,649

- No single HPC system capable of exceeding exaflop **for generic computations**
- Fastest known: < 0.2 Eflop/s
- 1-5 still way bellow 1 Eflop/s
- We may need to pool resources of multiple Centers
- Has been done for years but we want to:
 - Enable running single computation on multiple sites
 - Provide support for non-IT scientists to prepare/run codes on such massive scale

Why Singularity?



- Why container ?
 - Small footprint
 - Less overhead
 - Quick launch
 - Manageable images
- Why Singularity?
 - Built for the HPC
 - Integrated with SLURM
 - Unprivileged / secure
 - Support for MPI, GPU, ...



- User accesses WebUI
- Service layer is used to:
 - select inputs
 - choose code version
 - prepare and run computations
- Computations are scheduled on the HPC via RIMROCK
- Computations may be in form of classical scripts or Singularity (new)

1

• SIGN-IN

2

• DASHBOARD

3

• CREATE PROJECT


4

• CREATE PIPELINE

5

• RUN COMPUTATION

← → ↻ https://process-dev.cyfronet.pl/users/sign_in

 You need to sign in or sign up before continuing.

Login form


Email

Password

☐ Remember me

Log in

Forgot your password?

Log in with 

New to site? [Create account](#)

PROCESS

©2016 All Rights Reserved.

1

• SIGN-IN

2

• DASHBOARD

3

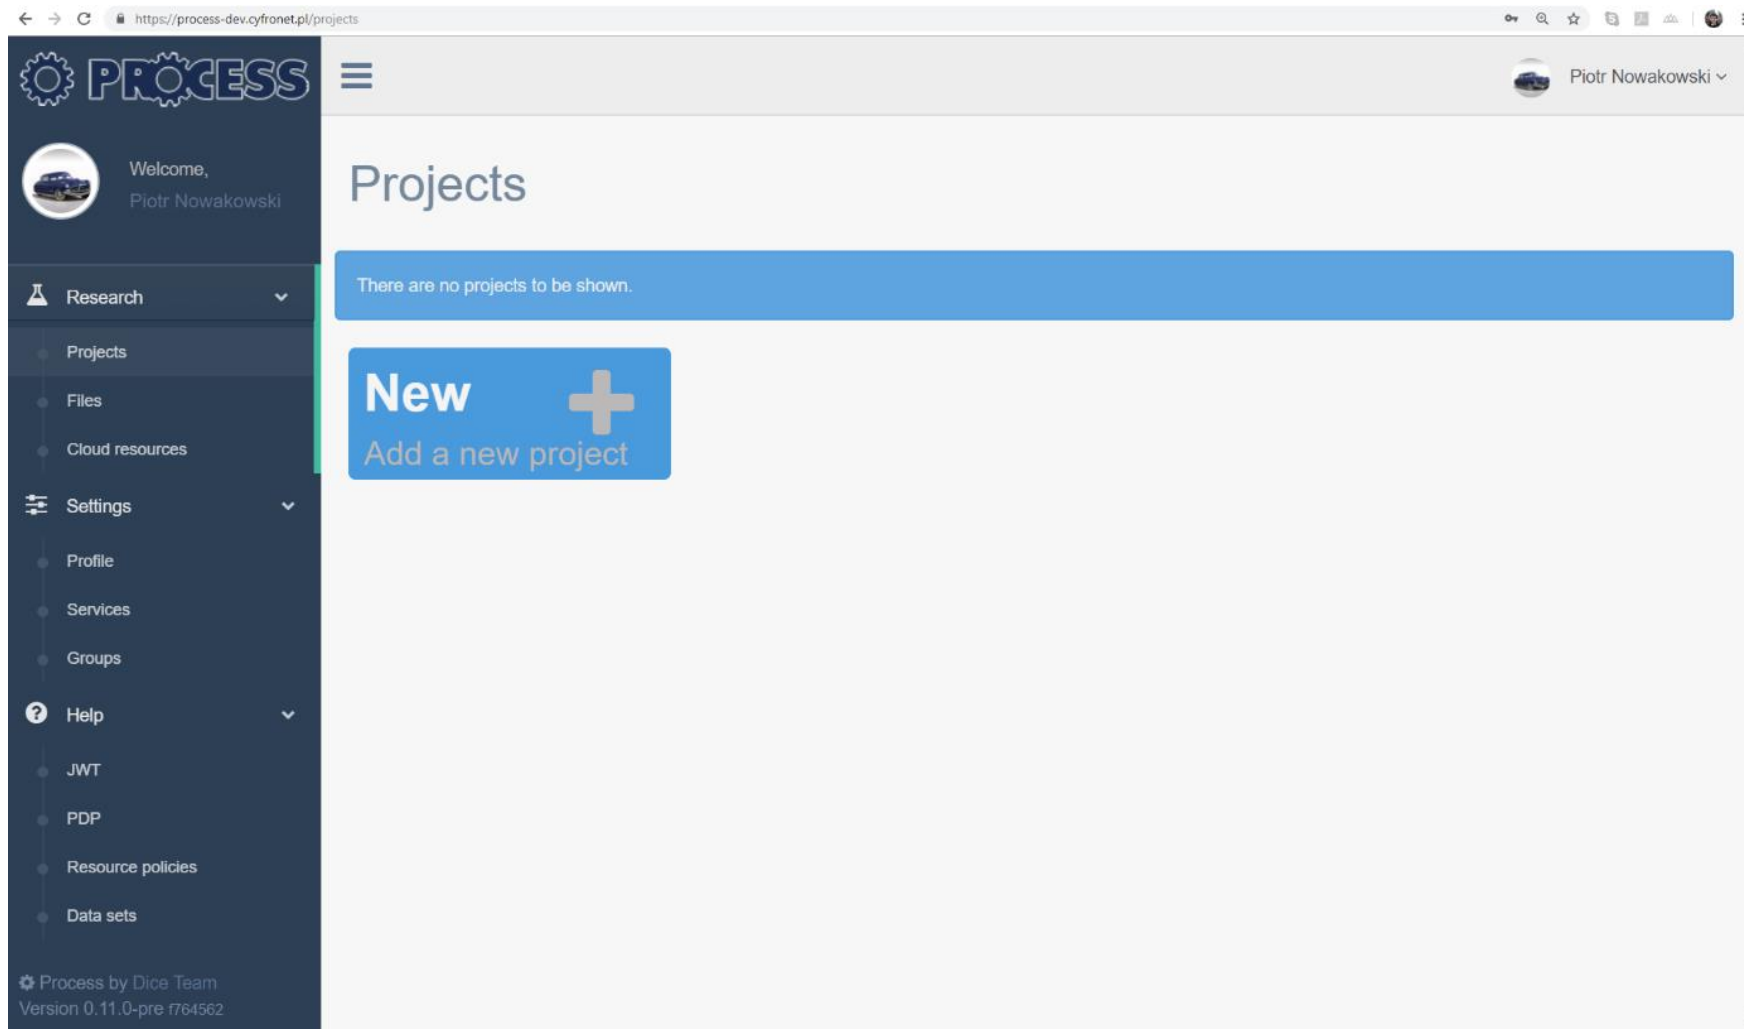
• CREATE PROJECT

4

• CREATE PIPELINE

5

• RUN COMPUTATION



1

• SIGN-IN

2

• DASHBOARD

3

• CREATE PROJECT

4

• CREATE PIPELINE

5

• RUN COMPUTATION

The screenshot displays the PROCESS web application interface. The browser address bar shows the URL: `https://process-dev.cyfronet.pl/projects/my_computation`. The interface includes a dark sidebar on the left with a navigation menu containing: Research (with a dropdown), Projects, Files, Cloud resources, Settings (with a dropdown), Profile, Services, Groups, Help (with a dropdown), JWT, PDP, Resource policies, and Data sets. At the bottom of the sidebar, it says "Process by Dice Team" and "Version 0.11.0-pre 68f2b36".

The main content area features a green notification banner at the top that reads "New project added." with a checkmark icon. Below this, the user's profile is shown as "Welcome, Piotr Nowakowski" with a dropdown arrow. The project name "my_computation" is displayed as "A Process project", accompanied by "Delete this project" and "See other projects" buttons. A section titled "Pipelines for this project" includes a "Set up new pipeline" button and a table with headers: Id, Name, Owner, and Status. The table is currently empty.

Below the pipelines section, there are two panels. The left panel, titled "my_computation project inputs", contains an "Upload" button and a dashed box labeled "Drop files or folders here" with the text "File Store Browser v0.21.3-SNAPSHOT" and "Empty" below it. The right panel, titled "All pipelines files", also shows "File Store Browser v0.21.3-SNAPSHOT" and "Empty".

1

• SIGN-IN

2

• DASHBOARD

3

• CREATE PROJECT

4

• CREATE PIPELINE

5

• RUN COMPUTATION

The screenshot displays the PROCESS web application interface. The top navigation bar includes the PROCESS logo, a hamburger menu, and the user's name 'Piotr Nowakowski'. The left sidebar contains a navigation menu with options: Research, Projects, Files, Cloud resources, Settings, Profile, Services, Groups, Help, JWT, PDP, Resource policies, and Data sets. The main content area shows the 'Run pipeline' form for a project named 'my_computation'. The form includes a 'Select pipeline' dropdown set to 'Singularity pipeline', a 'Name' field with 'pipeline_01', and a 'Mode' dropdown set to 'automatic'. Below the form, there is a 'Pipeline steps' section with a 'Singularity container computation' dropdown. The dropdown menu is open, showing options: 'Postpone selection until execution' (highlighted), 'branches', 'master', and 'tags'. Buttons for 'Delete this project' and 'See other projects' are visible in the top right. At the bottom of the sidebar, it says 'Process by Dice Team Version 0.11.0-pre 68f2b36'.

1

• SIGN-IN

2

• DASHBOARD

3

• CREATE PROJECT

4

• CREATE PIPELINE

5

• RUN COMPUTATION

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, Help, JWT, PDP, and Resource policies. The main content area shows the 'my_computation' project page. At the top, it says 'my_computation A Process project' with buttons for 'Delete this project' and 'See other projects'. Below this is the 'pipeline_01' section, identified as a 'Singularity pipeline (automatic pipeline)' owned by 'Piotr Nowakowski'. A blue message box states: 'All required input files are available. You can start Singularity container computation for this project.' Below this is a table with columns: Start time, Model version, Revision, Execution time, Outputs, and Status. The table contains one row with a status of 'Waiting for execution'. To the right of the table are buttons for 'Save Singularity container computation version' and 'Reload tags and branches'. At the bottom, there are two panels: 'pipeline_01 pipeline inputs' showing an 'input.csv' file uploaded on 10/10/18 at 14:30 (5 B), and 'pipeline_01 pipeline outputs' which is currently empty. The footer of the interface indicates 'Process by Dice Team Version 0.11.0-pre 682b36'.

- Conclusions:
 - New type of the pipeline step based on the Singularity containers was created for the Interactive Execution Environment
 - It enables execution of user provided applications on the HPC
 - Tested on the HPC system at Cyfronet (CPU and GPU)
- Future plans:
 - Extending solution for other Clusters in PROCESS (in Germany, Slovakia and the Netherlands)
 - Enable running same computation on multiple sites
 - Handle site-to-site communication
 - Add support for the private and public clouds

<http://dice.cyfronet.pl>



<http://www.process-project.eu>

PROCESS H2020 Project 777533