



Fusion Data Space Provider

Andrew Lahiff

UK Atomic Energy Authority

EGI Conference 2021

Dissemination level: Public

Disclosing Party: UKAEA

Recipient Party: EGI-ACE Project

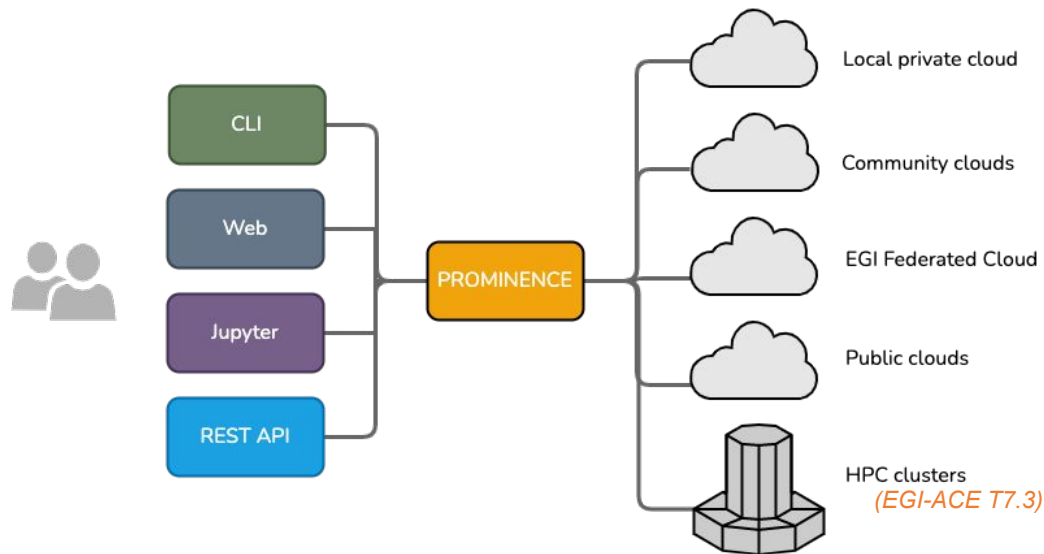


EGI-ACE receives funding from the European Union's Horizon 2020 research and innovation programme under grant agreement no. 101017567.

Introduction to PROMINENCE

Submit globally, run globally

- Platform allowing users to run batch jobs/workflows across multiple clouds
 - Appears like a traditional batch system to users
 - Users don't need to worry about provisioning clusters or dealing with infrastructure
 - Jobs are directed automatically to the appropriate resources



Introduction to PROMINENCE

Submit globally, run globally



Installing the CLI

```
pip install prominence-cli
```

Running a job

```
prominence create --cpus 8 --memory 16 alahiff/lammps-intel-avx512-2018 \  
"lmp_intel_cpu_intelmpi -in in.lj"
```

Listing jobs

```
prominence list jobs
```

View stdout/err in real time

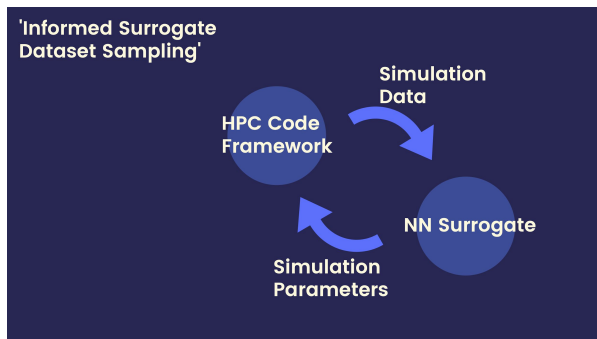
```
prominence stdout <job id>
```

GUI under development

The screenshot shows the Prominence web interface. On the left is a navigation sidebar with links for Home, Jobs, Workflows, Compute, Storage, and Documentation. The main content area is titled 'Jobs' and includes a '+ CREATE JOB' button. Below this is a table of jobs with columns for ID, Name, Status, Created, Elapsed, Image, Command, and Actions. The table contains 9 rows of job data, all with a 'Completed' status. At the bottom of the table, it says 'Showing 1 to 9 of 9 jobs' and has 'Previous' and 'Next' navigation buttons.




ID	Name	Status	Created	Elapsed	Image	Command	Actions
2178168	comprehens-ve-tesla/accelash-space-content-singularity	Completed	2021-02-24 20:26:11	1m	oemros.7	/bin/bash -c 'echo working ; du -s .'	
2178169	comprehens-ve-tesla/accelash-space-content-udocker	Completed	2021-02-24 20:26:11	45s	oemros.7	/bin/bash -c 'echo working ; du -s .'	
2178170	comprehens-ve-tesla/accelash-auth-singularity	Completed	2021-02-24 20:26:11	1m	reg.stry-github.com/alahiff/accelash-auth-test		
2178171	comprehens-ve-tesla/accelash-executable-input-file	Completed	2021-02-24 20:26:11	1m	oemros.7	du -s .	
2178172	comprehens-ve-tesla/accelash-executable	Completed	2021-02-24 20:26:11	42s	oemros.7	ls -l	
2178173	comprehens-ve-tesla/accelash-fact-sa3a11-mount-singularity	Completed	2021-02-24 20:26:12	1m	oemros.7	/bin/bash -c 'echo Singularity ; ls -l /data'	
2178174	comprehens-ve-tesla/accelash-fact-sa3a11-mount-udocker	Completed	2021-02-24 20:26:12	46s	oemros.7	/bin/bash -c 'echo udocker ; ls -l /data'	
2178175	comprehens-ve-tesla/output-file-object-store	Completed	2021-02-24 20:26:12	23s	oemros.7	/bin/bash -c 'echo Hello = file1.txt'	
2178176	comprehens-ve-tesla/output-directory-object-store	Completed	2021-02-24 20:26:12	1s	oemros.7	/bin/bash -c 'mkdir /outdir1 ; echo Hello = output1/fact1.txt'	

- **Building NN-based surrogate models in an efficient & effective manner**
 - Traditional approach involves arbitrary scans across a range of parameters
 - May be unaware of more complex behaviour & nuances
- **An alternative approach to overcome these difficulties**
 - Physics-informed neural network
 - Cyclic workflow:
 - Gaps in knowledge identified
 - Trigger the simulation code to generate additional data points



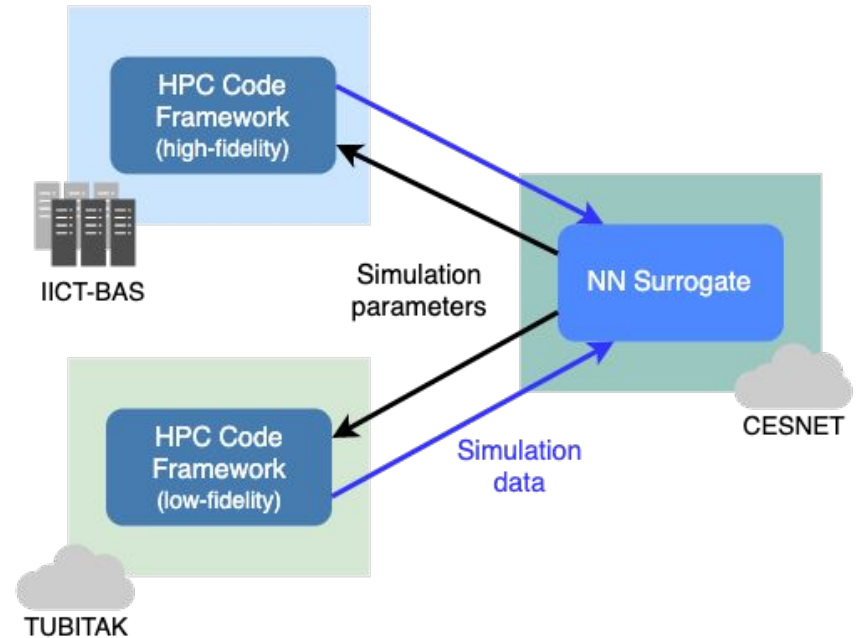
Building NN-based surrogate models

- Fusion example: JOREK
 - Simulation of MHD instabilities at the edge of Tokamak plasmas
 - Filamentary structures ejected from the edge
 - Dynamics highly dependent on edge plasma pressure

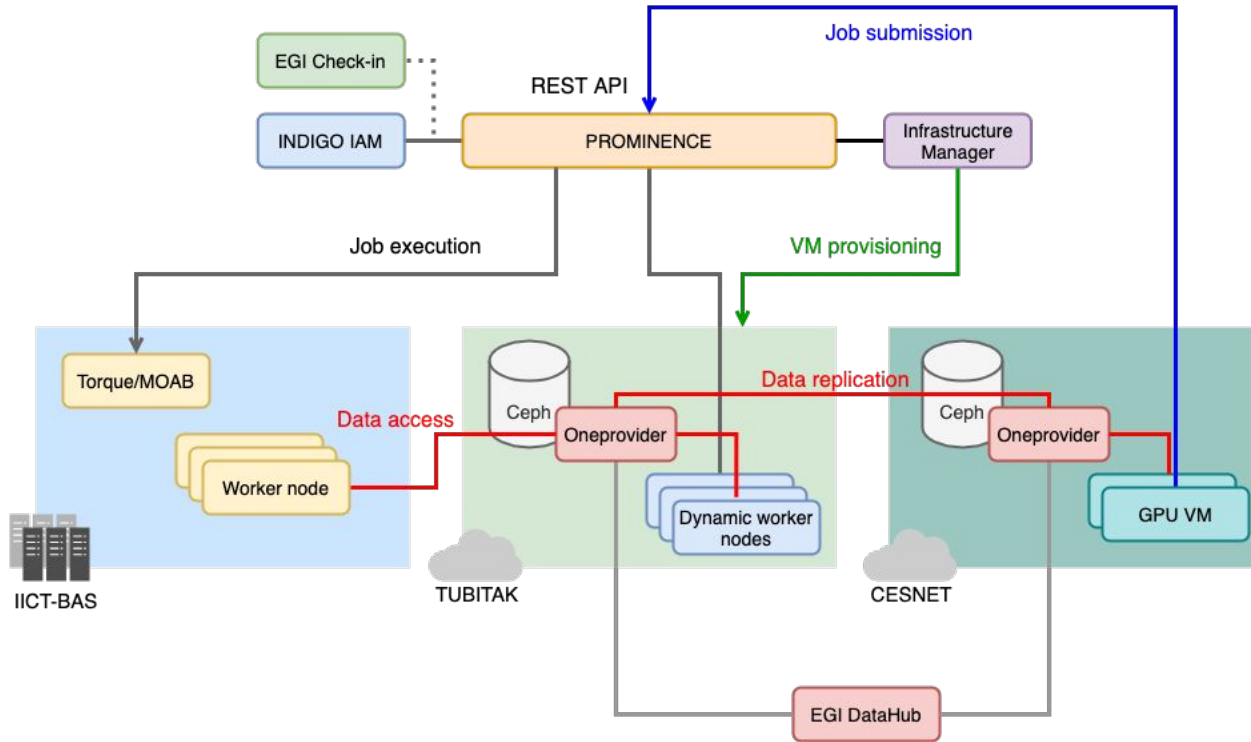


Use case mapped onto resources

- **JOREK is a HPC code**
 - For high-fidelity, requires large numbers of CPUs with low-latency interconnects
 - On typical clouds jobs limited to a single node
 - Limits fidelity
 - Limits physics
- **Compute**
 - HPC at IICT-BAS
 - Cloud (CPUs) at TUBITAK
 - Cloud (GPUs) at CESNET
- **Storage**
 - 50 TB TUBITAK
 - 10 TB CESNET
 - Access via OneData



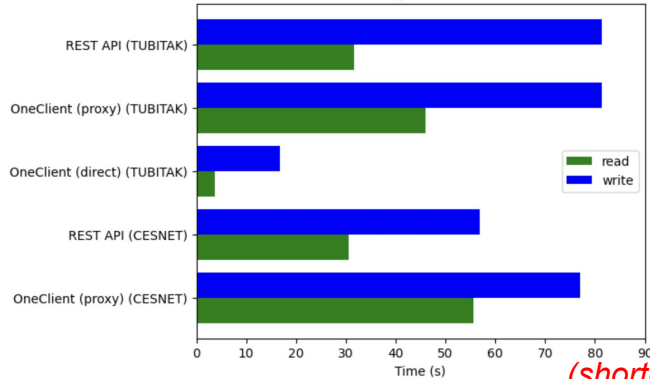
Architecture




- In PROMINENCE all jobs run in containers
 - Container runtime
 - PROMINENCE supports Singularity & udocker
 - udocker very useful
 - Singularity not installed on some HPC systems (e.g. AVITOHOL at IICT-BAS)
 - udocker can be installed by unprivileged users
- Portability vs performance
 - Containers do not guarantee portability when codes are optimised for specific architectures
 - Typically building images on Intel Skylake (or newer)
 - Our EGI-ACE resources:
 - TUBITAK: AMD Opteron 23xx
 - IICT-BAS: Intel Xeon E5-2650 v2
 - UNIV-LILLE: *Intel Westmere*

Data access & performance


VM @ TUBITAK



VM @ CESNET



WN @ ICT-BAS




(shorter is better)

Proxy I/O mode



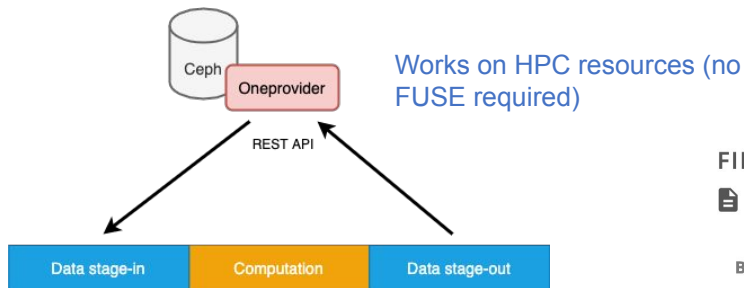
Direct I/O mode



- Comparing:
 - FUSE: Oneclient (proxy IO mode)
 - FUSE: Oneclient (direct IO mode) where possible
 - CDMI REST API
 - On HPC resources use of FUSE by normal users typically not possible
- Tests: copying a 1 GB file

Accessing & finding data

- **Data access in PROMINENCE**
 - Jobs can mount OneData/WebDAV storage at specified location
 - For OneData can use GeoIP to find nearest provider if necessary
 - Stage-in & stage-out specified input & output files to object storage
 - Modified to support OneData REST API as well (previously only S3)



- **Custom meta data in OneData**
 - JSON documents can be attached to files

FILE METADATA

T_rho_profiles.dat


Basic JSON RDF

```
1 {
2   "parameters": {
3     "rho_coef5": 0.96,
4     "rho_coef4": 0.08,
5     "T_coef5": 0.96,
6     "T_coef4": 0.057743,
7     "FF_coef8": 0.07,
8     "FF_coef7": 0.92,
9     "FF_coef6": -0.35
10  }
11 }
```

Distributed data

- Three options available for dealing with the two storage systems
 - OneData can automatically transfer files as necessary on the fly
 - Use QoS to specify replication rules
 - Trigger bulk replication of data when needed
- Example: replication of 16368 low-fidelity JOREK runs from TUBITAK to CESNET


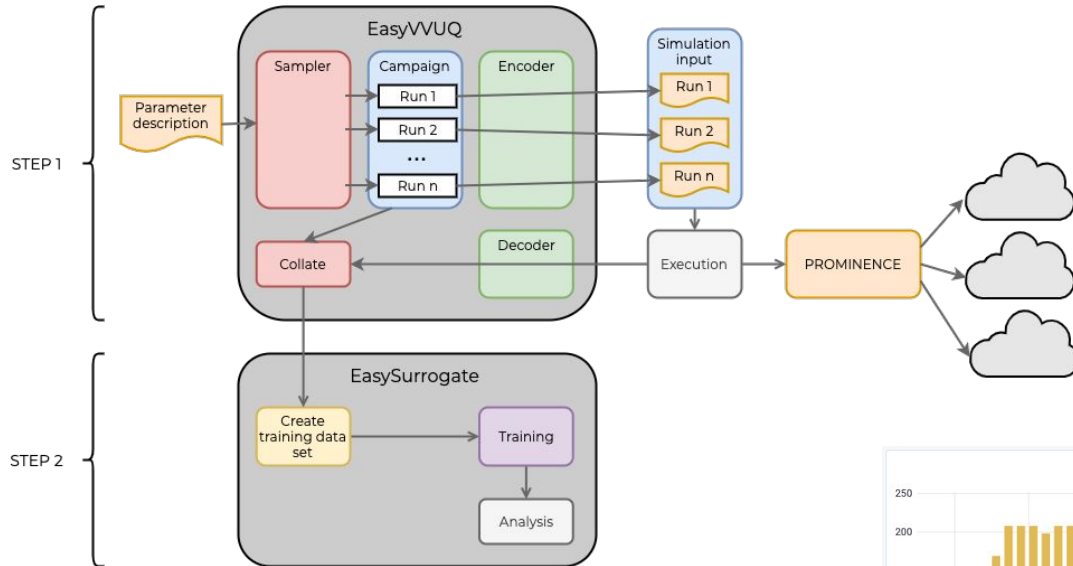
ONGOING TRANSFERS MAP




File/directory	Username	Destination	Started at	Finished at	Replicated	Type	Status
transfers	Andrew Lahiff	cesnet-oneprovider-01	19 Oct 2021 6:39:57	19 Oct 2021 21:20:03	699.1 GiB		✓

Basic proof of concept

(Very) simple surrogate model of JOREK using VECMA



Comparison between surrogate model & JOREK (test data)



400 JOREK runs executed by PROMINENCE at TUBITAK

Future plans



- JOREK portability & performance
 - Compile JOREK for best performance & portability
 - Get JOREK to run successfully on IICT-BAS
- Production & curation of initial JOREK dataset for NN training
- Development of initial surrogate model
- Metrics & methods for identifying knowledge gaps in models



Thank you!

Contact: egi-ace-po@mailman.egi.eu
Website: www.egi.eu/projects/egi-ace



[EGI Foundation](#)



[@EGI_eInfra](#)



EGI-ACE receives funding from the European Union's Horizon 2020 research and innovation programme under grant agreement no. 101017567.