



Cloud scavenging with HTCondor in the EOSCpilot Fusion Science Demonstrator

Andrew Lahiff

Culham Centre for Fusion Energy, UK Atomic Energy Authority

EOSCpilot
The European Open Science
Cloud for Research Pilot Project
www.eosc-pilot.eu



- Introduction
 - What is EOSC & EOSCpilot?
 - Computing in Fusion
- EOSCpilot Fusion Science Demonstrator
- Conclusions & future work




What is EOSCpilot?

- European Open Science Cloud (EOSC)
 - EOSCpilot
 - First phase in the development of the EOSC
 - EOSC-Hub
 - Providing a way for researchers to discover, access & use and re-use a broad spectrum of resources for data-driven research
- Science Demonstrators in EOSCpilot
 - Aim to show the relevance & usefulness of EOSCpilot services
 - 15 demonstrators from different disciplines
 - Environment, earth sciences, high energy physics, life sciences, social sciences, astronomy, ...
 - (Very) small-scale proof-of-concepts



- PROMINENCE is the Fusion Science Demonstrator



NASA Goddard Space Flight Centre
<https://creativecommons.org/licenses/by/2.0/legalcode>



- Fusion
 - Recreating the power of the sun in a small volume
 - Eventual aim to create electricity
- EUROfusion research community today
 - Isolated “islands” of data & computing resources
 - Difficult to access data & resources at other sites
- Challenges of ITER
 - Largest fusion experiment ever built
 - 2 PB of raw data will be generated per day
 - Data will need to be distributed globally
 - Need for access to more computing resources



© ITER Organization, <http://www.iter.org/>



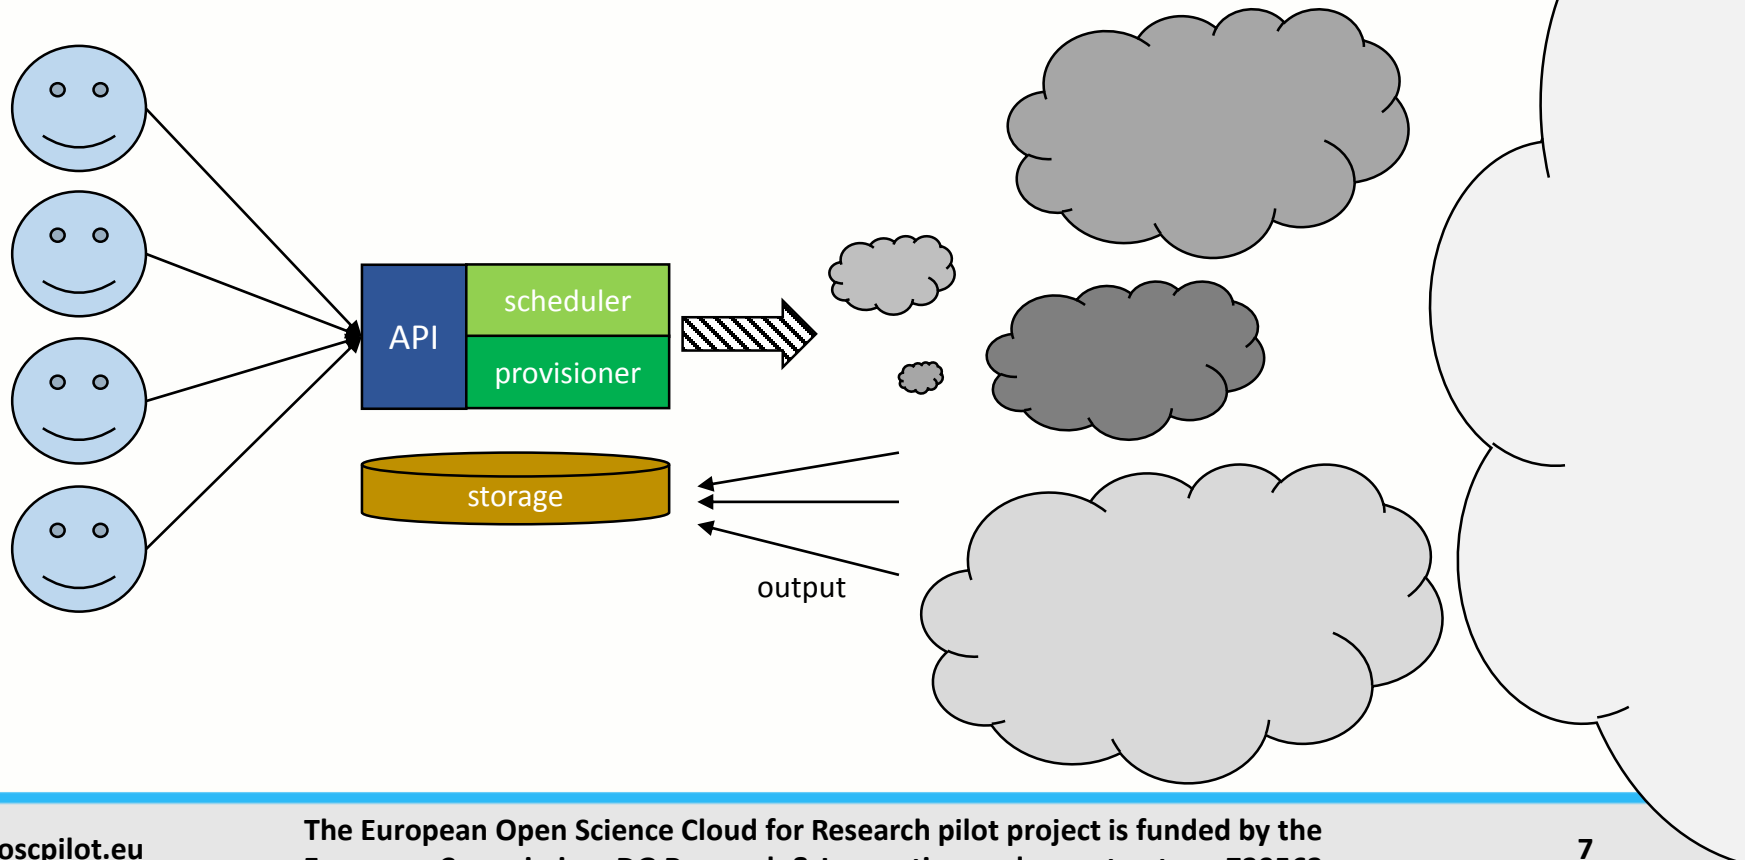
The Fusion Science Demonstrator

- Aims to address part of the fusion computing problem
 - How to make use of academic and public clouds for Fusion HPC applications
 - How to extend small scale on-premises facilities to cope with bursty workloads
 - How to get the same result wherever you run
- Data management is not being considered yet



The Fusion Science Demonstrator

- Allow users to run containerized HTC & HPC jobs on clouds
 - Bursting from local cloud resources onto external clouds
 - Easy access to output files for further analysis






Batch jobs in clouds

- Why develop yet another way to run batch jobs on clouds?
 - Most people assume you have a *Single Great Big Reliable Cloud*TM
- What if you only have access to lots of little clouds?
 - In particular only or mostly opportunistic access?
- Cloud bursting
 - You have a small local cloud, but want to burst out to external clouds automatically
 - Not just bursting to AWS (or Azure, or Google)
 - Burst first to national research cloud(s), then to public clouds
 - We need “hierarchical cloud bursting”

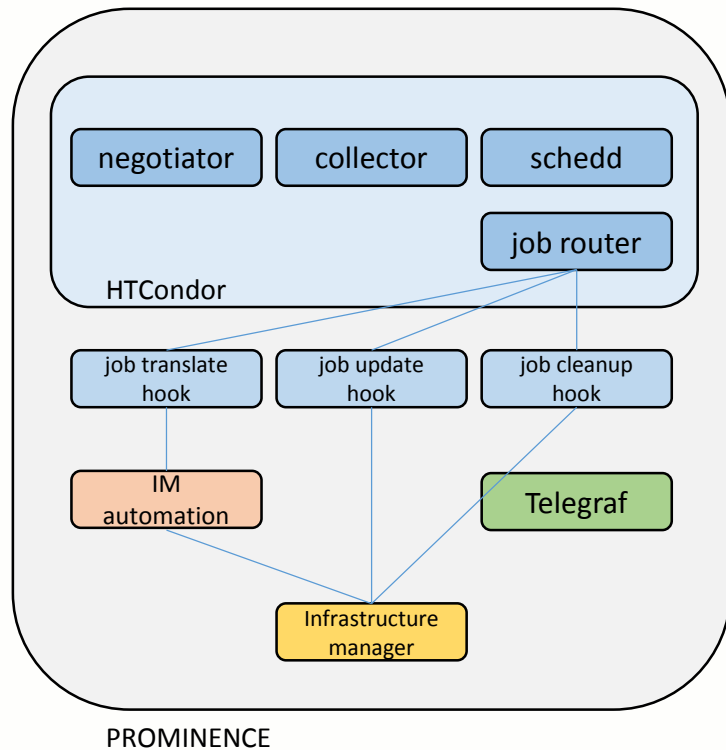


Hierarchical cloud bursting





Architecture



IAM


Swift
RAL

DockerHub

InfluxDB
Google Cloud

Grafana
grafana.net

Other services





HTCondor functionality used

- Job router + plugins
 - Provisions & destroys infrastructure for each job
- File transfer plugin
 - Downloads any artifacts & extracts archives if necessary
- Job prepare hook
 - Modifies job ClassAds to use Singularity
 - We're not using HTCondor's built-in support for Singularity
 - HTCondor doesn't appear to support jobs with no executable
- Job exit hook
 - Uploads output files (if necessary) to Swift
- Many custom ClassAd attributes
 - Using HTCondor as a database



Cloud deployment

- Deploy one VM (or group of VMs) per job
 - Expecting all jobs to be at least multi-core if not multi-node
 - Expecting medium job run-times (hours-days)
 - Expecting a wide variety of cores & memory requirements per job
- Infrastructure Manager used to deploy infrastructure on clouds
 - One of the few tools which supports OCCI, needed for EGI FedCloud
 - Can use either Ansible or cloud-init for contextualization
 - Successfully used EGI FedCloud, OpenStack, Google, Azure



- Automation layer around Infrastructure Manager
 - Finds the best cloud to provision infrastructure, given requirements & preferences
 - Requirements
 - Number of cores, memory, image (name or distribution/version/...)
 - Site or region
 - Preferences
 - List of regions in order of preferences
 - Prefer to run an MPI job on a cloud with low-latency interconnects
 - Handling failures
 - If multiple clouds meet requirements, if deployment fails on one clouds, will automatically try another
 - Avoid clouds temporarily which have failed recently



- Submitting jobs to HTCondor
 - ssh access to a submit node running schedd
 - Remote submission
- This is quite limiting
 - Giving users ssh access to a machine is never easy
 - Especially external users
- What about a RESTful API + JSON?
 - Submit easily from anywhere
 - Mac laptop, Windows desktop, Go app, ...
 - Provide a simple CLI which uses the RESTful API

~~*Submit locally, run globally*~~

Submit globally, run globally



- Almost like a Grid CE, but provides a “normal” API & doesn’t use X.509
- Python Flask + gunicorn + nginx
 - Uses HTCondor Python API
 - nginx provides SSL termination (Let’s Encrypt)
- Sandbox directory created per job
 - Input files, stdout/err
- Authentication
 - For proof-of-concept using hardwired list of usernames & passwords
 - Have successfully tested integration with the INDIGO-DataCloud IAM service
 - OpenID Connect tokens for authentication



Command line interface examples

- Submitting the simplest possible job:

```
[cloudadm@vnode-0 ~]$ prominence create alahiff/testpi:latest
{
  "id": 117
}
```

- Listing jobs:

```
[cloudadm@vnode-0 ~]$ prominence list
[
  {
    "id": 95,
    "status": "running",
    "image": "alahiff/cherab-jet:latest",
    "cmd": "python",
    "args": "batch_make_sensitivity_matrix.py 0 59"
  },
]
```



- Details about a job

```
[cloudadm@vnode-0 ~]$ prominence get 127
```

```
[  
  {  
    "id": 127,  
    "status": "running",  
    "image": "alahiff/raysect-demo-bunny:latest",  
    "cpus": 1,  
    "memory": 1,  
    "nodes": 1,  
    "disk": 10,  
    "runtime": 720,  
    "outputFiles": [  
      {  
        "url": "https://s3.echo.stfc.ac.uk/swift/v1/prominence-jobs/317c4f75-e109-459d-bded-e",  
        "name": "stanford_bunny.png"  
      }  
    ],  
    "events": {  
      "creation": "2018-08-29T15:01:53Z",  
      "containerCreationStart": "2018-08-29T15:12:42Z"  
    }  
  }  
]
```



Examples

- Have successfully run a variety of applications, including
 - Geant4 (MPI)
 - simulation of the passage of particles through matter
 - SERPENT2
 - continuous-energy Monte Carlo particle transport code
 - ASCOT (MPI)
 - Accelerated Simulation of Charged Particle Orbits in a Tokamak
 - Raysect
 - Scientific ray-tracing

Example: Breeder blanket design

- Breeder blankets are designed to ensure tritium self-sufficiency in fusion power plants
- A blanket design tool was created which automates and parameterizes the production of 3D CAD based neutronics models
- Machine learning was used to drive the parametric model creation & optimize the breeder blanket performance in terms of tritium production

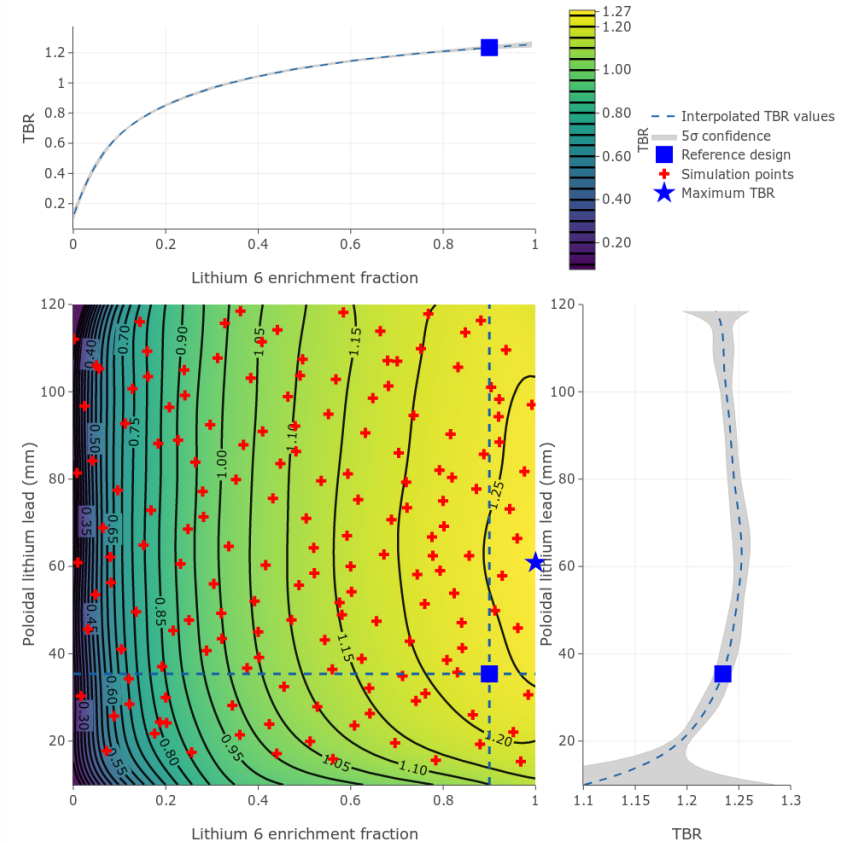
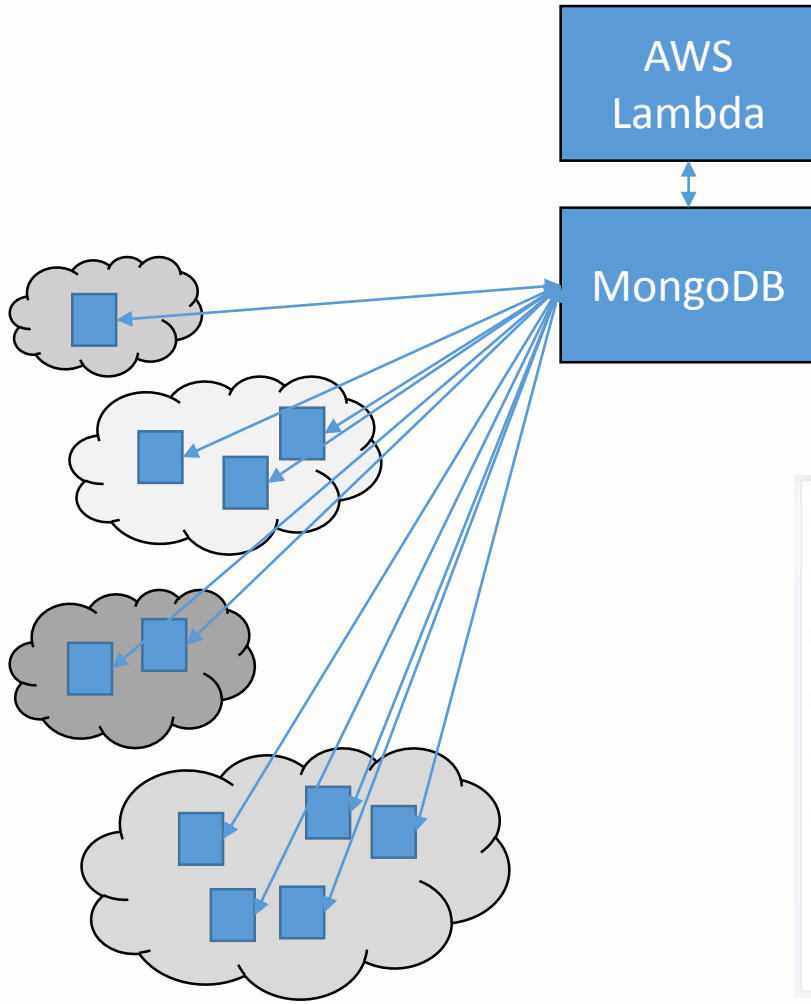


Image courtesy of Jonathan Shimwell




Example: Breeder blanket design




- Lots of identical jobs
- Each job obtains input parameters from MongoDB
- Output written to MongoDB
- An AWS Lambda function visualizes the progress of the parameter sweep

Cloud resources used






Example: Breeder blanket design



Full reactor design



The different blanket designs simulated

Images courtesy of Jonathan Shimwell



Conclusions & future work

- EOSCpilot has allowed us to successfully demonstrate the usage of cloud computing for fusion
 - Including MPI jobs
 - Makes extensive use of HTCondor
- Future work
 - EOSC-Hub Fusion Competence Centre
 - Workflows
 - Data management
 - Integration with a private Docker registry
 - Restrict access to container images to specific people/groups
 - Improvements & scalability
 - Moving towards making PROMINENCE a production service



EOSCpilot
The European Open Science
Cloud for Research Pilot Project
www.eoscpilot.eu

Questions?



© 2015 Kisha Gianni and Lyn Gianni