

European HTCondor Workshop 2019



Report of Contributions

Contribution ID: 1

Type: **not specified**

HTCondor and containers for Batch and Interactive use: (Mostly) a success story

Tuesday, September 24, 2019 5:00 PM (25 minutes)

An HTC cluster using HTCondor has been set up at Bonn University in 2017/2018. All infrastructure is fully puppetised, including the HTCondor configuration. Both interactive and batch jobs are run inside Singularity containers, and users only have to choose the desired OS via a job parameter from an offered collection of container images without setting up or building a container themselves.

The container images are rebuilt daily and are provided by a CernVM filesystem (CVMFS) along with various software packages. The data to be analysed is stored on a CephFS file system.

This talk presents the successful migration from a “classic” PBS cluster setup with separate login nodes providing an environment similar to the batch environment, causing many regular admin headaches, to a modern and flexible, HTCondor-based system offering the users several different environments both for interactive and batch usage. Both the successes and the several pitfalls and issues which were encountered will be discussed.

Desired slot length

20

Speaker release

Yes

Primary authors: WIENEMANN, Peter (University of Bonn (DE)); FREYERMUTH, Oliver (University of Bonn (DE))

Presenter: FREYERMUTH, Oliver (University of Bonn (DE))

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 2

Type: **not specified**

Automation & Monitoring in the CERN Condor pool

Wednesday, September 25, 2019 11:45 AM (25 minutes)

The CERN HTCondor pool is currently offering 200K cores of compute power to hundreds of users in the HEP community. Managing such cluster requires a significant effort in the daily operations, not only because of the scale, but also because of the diversity of the resources. In this scenario, the adoption of automation and monitoring tools becomes a strong requirement to optimize both the resource usage and the operators time.

This talk presents different projects and prototypes that have been developed and integrated in our infrastructure to make these daily operations more efficient.

Desired slot length

Speaker release

Yes

Primary author: FERNANDEZ ALVAREZ, Luis (CERN)

Presenter: FERNANDEZ ALVAREZ, Luis (CERN)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 3

Type: **not specified**

Edge cases at CERN

Wednesday, September 25, 2019 12:10 PM (25 minutes)

The overwhelming majority of batch workload at CERN is very similar. Experiment pilots, experiment tier-0 production, and local users running single core jobs. However one size doesn't fit all, and we now have a number of different edge cases that this talk will cover...

GPUs, from machine learning to software validation.

Users who are in a grey area between HPC & HTC

Lower priority preemptible workloads

Backfill of SLURM resources.

Running HTCondor on kubernetes

Desired slot length

20

Speaker release

Yes

Primary author: JONES, Ben (CERN)

Presenter: JONES, Ben (CERN)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 4

Type: **not specified**

Update on Implementation and Usage of HTCondor at DESY

Tuesday, September 24, 2019 5:25 PM (35 minutes)

In this talk we provide new details of the DESY configurations for HTCondor. We focus on features needed for user registry integration, node maintenance operations and fair share / quota handling. We are working on Docker, Jupyter and GPU integration into our smooth and transparent operating model setup.

Desired slot length

20 minutes

Speaker release

Yes

Primary authors: Mr FINNERN, Thomas (DESY); Mr BEYER, Christoph

Co-authors: Dr KEMP, Yves (DESY); Mr FLEMMING, Martin (DESY)

Presenters: Mr FINNERN, Thomas (DESY); Mr BEYER, Christoph

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 5

Type: **not specified**

'Startd flocking' and 'HT-HPC' c/o the Milan Physics Dept.

Wednesday, September 25, 2019 5:20 PM (25 minutes)

Group-owned and operated clusters at the Physics Department of the University of Milan (UNIMI) are pooled together via the HTCondor 'Startd flocking' feature. We describe the setup, its applications and possible use for parallel applications, with some preliminary performance results.

Desired slot length

Speaker release

Yes

Primary authors: PRELZ, Francesco (Università degli Studi e INFN Milano (IT)); REBATTO, David (Università degli Studi e INFN Milano (IT))

Presenter: PRELZ, Francesco (Università degli Studi e INFN Milano (IT))

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 6

Type: **not specified**

Large scale mapping of human settlements from earth observation data with JEO-batch of the JRC Earth Observation Data and Processing Platform

Friday, September 27, 2019 9:50 AM (25 minutes)

The spatial distribution of built-up areas and their expansion represent one of the most important forms of land use/ land cover changes confronting climate, environmental and socio-economic systems at a global scale. Characterizing the status and dynamics of built-up areas over large areas is technically feasible thanks to the availability of a panoply of earth observation data with different spatial, spectral and temporal characteristics. In the last 10 years, at the Joint Research Centre of the European Commission, the Global Human Settlement Layer project has been exploiting different sources of satellite imagery to monitor changes in the European and Global built-up landscapes to better inform policies and decision making (Corbane et al. 2017) (Florczyk et al. 2016).

To meet the demands of large-scale mapping of human settlements from space, not only mass storage infrastructures are needed but also novel data analytics combined with high-performance computing platforms have to be designed. The JRC EO Data and Processing Platform (JEODPP) developed in the framework of JRC Big Data Analytics (BDA) project provides petabyte scale storage coupled to high throughput computing capacities enabling and facilitating the extraction of built-up areas from large volumes of satellite data both at the European and Global scales (Soille et al. 2018). In this work, we present the JEO-batch feature of the JEODPP; a low-level batch processing orchestrated by a dedicated workload manager and its utility for the execution of two main automated workflows for extracting built-up areas over large zones. The first workflow is implemented on a pan-European coverage of Very High Resolution satellite data from the Copernicus contributing missions acquired in 2015; the second workflow exploits a global Sentinel-2 pixel-based composite from the Copernicus constellation of satellites acquired mainly in 2018.

Although the two workflows build on the same classifier, the number of images/tiles to be processed and their projections, the characteristics of remote sensing sensors, in particular their spatial resolutions and the derived outputs required different configurations of the workload automations. Taking advantage of the Docker universe of the JEO-batch that relies on the HTCondor architecture, the workflows, originally coded in Matlab, were compiled and successfully run on the JEODPP. The following table summarizes the main characteristics of the massive batch processing that allowed extracting built-up areas at the European and global scales:

Indico rendering error

Could not include image: [404] Error fetching image

References:

- Corbane, Christina, Martino Pesaresi, Panagiotis Politis, Vasileios Syrris, Aneta J. Florczyk, Pierre Soille, Luca Maffeni, et al. 2017. "Big Earth Data Analytics on Sentinel-1 and Landsat Imagery in Support to Global Human Settlements Mapping." *Big Earth Data* 1 (1–2): 118–44. <https://doi.org/10.1080/20964471.2017.1397899>.
- Florczyk, Aneta Jadwiga, Stefano Ferri, Vasileios Syrris, Thomas Kemper, Matina Halkia, Pierre Soille, and Martino Pesaresi. 2016. "A New European Settlement Map From Optical Remotely Sensed Data." *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 9 (5): 1978–92. <https://doi.org/10.1109/JSTARS.2015.2485662>.
- Soille, P., A. Burger, D. De Marchi, P. Kempeneers, D. Rodriguez, V. Syrris, and V. Vasilev. 2018. "A Versatile Data-Intensive Computing Platform for Information Retrieval from Big Geospatial Data." *Future Generation Computer Systems* 81: 30–40. <https://doi.org/10.1016/j.future.2017.11.007>.

Desired slot length

Speaker release

Yes

Primary author: Dr CORBANE, Christina (Joint Research Centre)

Co-authors: Dr RODRIGUEZ, Dario (European Commission - DG JRC); Mr SABO, Filip (Arhs Developments Italia S.r.l, Milano, Italy); Mr POLITIS, POLITIS (Arhs Developments S.A. Rue Nicolas Bové, 2b L-1253 Luxembourg); Dr SYRRIS, Vasileios (Joint Research Centre)

Presenters: Dr CORBANE, Christina (Joint Research Centre); Dr RODRIGUEZ, Dario (European Commission - DG JRC)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 7

Type: **not specified**

Planning for Future Scales and Complexity of HTCondor Pools in the CMS Experiment

Wednesday, September 25, 2019 5:45 PM (35 minutes)

The resource needs of high energy physics experiments such as CMS at the LHC are expected to continue to grow significantly over the next decade, and will be more and more satisfied by computing capacity with non-standard characteristics. This presents challenges not only of scale but of complexity in resource provisioning and allocation. In this contribution, we will present results of recent HTCondor scale tests we have conducted using the CMS Global Pool Integration Test Bed (ITB) employing the multi-threaded Negotiator, where we have pushed the size of the pool to the maximum limits with currently-available hardware and explored effective performance limitations of the submit nodes in our infrastructure with realistic payloads. We will also discuss recent integration of resource-specific job matching conditions to satisfy HPC and Cloud use cases, where resources may not be suitable for running all kinds of workflows. Finally, we will review some specific use cases that we have difficulty solving with the current implementation of HTCondor.

Desired slot length

20 minutes, unless you want to give us more.

Speaker release

Yes

Primary author: LETTS, James (Univ. of California San Diego (US))

Co-authors: PEREZ-CALERO YZQUIERDO, Antonio (Centro de Investigaciones Energéticas Medioambientales y Tecno); HALEEM, Saqib (National Centre for Physics (PK))

Presenter: LETTS, James (Univ. of California San Diego (US))

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 8

Type: **not specified**

HTCondor Philosophy and Architecture continued

Overview of HTCondor's architecture and sequence of events when scheduling jobs.

Desired slot length

30 minutes, but could go down to 20 minutes if needed

Speaker release

Yes

Presenter: THAIN, Gregory (University of Wisconsin-Madison)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 9

Type: **not specified**

Python Bindings Tutorial

Tuesday, September 24, 2019 11:45 AM (40 minutes)

An overview of HTCondor's official Python APIs, including job submission, monitoring, and what's new. Beyond just pointpoint files, this talk will attempt to allow people to learn live via Jupyter notebook sessions.

Desired slot length

40 minutes

Speaker release

Yes

Primary authors: TANNENBAUM, Todd (University of Wisconsin Madison (US)); BOCKELMAN, Brian Paul (University of Nebraska Lincoln (US))

Presenter: TANNENBAUM, Todd (University of Wisconsin Madison (US))

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 10

Type: **not specified**

What's new in HTCondor? What's coming up?

Wednesday, September 25, 2019 10:00 AM (35 minutes)

Overview of new features recently released as well as discussion about the HTCondor development roadmap.

Desired slot length

30 minutes

Speaker release

Yes

Primary author: TANNENBAUM, Todd (University of Wisconsin Madison (US))

Presenter: TANNENBAUM, Todd (University of Wisconsin Madison (US))

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 11

Type: **not specified**

HTMap and HTC Notebooks: Bringing HTC to Python

Wednesday, September 25, 2019 9:00 AM (35 minutes)

Introduction to some recent work by the HTCCondor team to enable Python code, including Python embedded in Jupyter Notebooks, to easily and naturally leverage high throughput computing via HTCCondor.

Desired slot length

30 minutes

Speaker release

Yes

Primary authors: TANNENBAUM, Todd (University of Wisconsin Madison (US)); KARPEL, Josh (University of Wisconsin-Madison)

Presenter: TANNENBAUM, Todd (University of Wisconsin Madison (US))

Session Classification: Workshop presentations

Track Classification: HTCCondor presentations and tutorials

Contribution ID: 12

Type: **not specified**

Moving Job Data

Wednesday, September 25, 2019 4:30 PM (25 minutes)

Discussion about how HTCCondor jobs can access their data, including some recent developments.

Desired slot length

20 minutes

Speaker release

Yes

Primary authors: TANNENBAUM, Todd (University of Wisconsin Madison (US)); BOCKELMAN, Brian Paul (University of Nebraska Lincoln (US))

Presenter: TANNENBAUM, Todd (University of Wisconsin Madison (US))

Session Classification: Workshop presentations

Track Classification: HTCCondor presentations and tutorials

Contribution ID: 13

Type: **not specified**

Philosophy of HTCondor

Tuesday, September 24, 2019 9:30 AM (1h 10m)

Provides an overview of HTCondor's design and the principles behind it

Desired slot length

30 minutes

Speaker release

Yes

Primary author: THAIN, Gregory (University of Wisconsin-Madison)

Presenter: THAIN, Gregory (University of Wisconsin-Madison)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 14

Type: **not specified**

HTCondor ClassAd Advanced Tutorial

Tuesday, September 24, 2019 11:10 AM (35 minutes)

An in-depth coverage of the HTCondor classad language with examples.

Desired slot length

Speaker release

Yes

Primary author: THAIN, Gregory (University of Wisconsin-Madison)

Presenter: THAIN, Gregory (University of Wisconsin-Madison)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 15

Type: **not specified**

HTCondor Negotiator Policy and Configuration

Tuesday, September 24, 2019 2:30 PM (1h 15m)

An explanation of the algorithms and policies of the HTCondor negotiator

Desired slot length

70 minutes

Speaker release

Yes

Primary author: THAIN, Gregory (University of Wisconsin-Madison)

Presenter: THAIN, Gregory (University of Wisconsin-Madison)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 16

Type: **not specified**

Federating HTCondor pools

Wednesday, September 25, 2019 12:35 PM (25 minutes)

Describes the several ways more than one condor pool can be joined.

Desired slot length

20 minutes

Speaker release

Yes

Primary author: THAIN, Gregory (University of Wisconsin-Madison)

Presenter: THAIN, Gregory (University of Wisconsin-Madison)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 17

Type: **not specified**

RESTful interfaces to HTCondor

Wednesday, September 25, 2019 4:55 PM (25 minutes)

Discusses work in program about possible RESTful interfaces to HTCondor

Desired slot length

20 minutes

Speaker release

Yes

Primary author: THAIN, Gregory (University of Wisconsin-Madison)

Presenter: THAIN, Gregory (University of Wisconsin-Madison)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 18

Type: **not specified**

Container support in HTCondor

Friday, September 27, 2019 11:35 AM (35 minutes)

Container support in HTCondor

Desired slot length

30 minutes

Speaker release

Yes

Primary author: THAIN, Gregory (University of Wisconsin-Madison)

Presenter: THAIN, Gregory (University of Wisconsin-Madison)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 19

Type: **not specified**

HTCondor-CE overview: from Clusters to Grids

Thursday, September 26, 2019 9:45 AM (20 minutes)

Overview of the design of the HTCondor CE

Desired slot length

15 minutes

Speaker release

Yes

Primary author: THAIN, Gregory (University of Wisconsin-Madison)

Presenter: THAIN, Gregory (University of Wisconsin-Madison)

Session Classification: Workshop presentations

Track Classification: HTCondor Compute Element (CE) presentations and tutorials

Contribution ID: 20

Type: **not specified**

HTCondor Advanced Job Submission

Tuesday, September 24, 2019 12:25 PM (35 minutes)

HTCondor Advanced Job Submission

Desired slot length

30 minutes

Speaker release

Yes

Primary author: THAIN, Gregory (University of Wisconsin-Madison)

Presenter: THAIN, Gregory (University of Wisconsin-Madison)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 21

Type: **not specified**

SciTokens and Credential Management

Friday, September 27, 2019 11:10 AM (25 minutes)

Presentation on work to integrate distributed authorization technologies such as SciTokens and OAuth 2.0 into HTCondor, and what this means for end-users and system administrators.

Desired slot length

20 minutes

Speaker release

Yes

Primary author: TANNENBAUM, Todd (University of Wisconsin Madison (US))

Presenter: TANNENBAUM, Todd (University of Wisconsin Madison (US))

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 22

Type: **not specified**

Security Mechanisms in HTCondor

Friday, September 27, 2019 10:15 AM (35 minutes)

Overview of the security mechanisms available in HTCondor, including discussion on security configuration and new security features being introduced in the HTCondor v8.9 series.

Desired slot length

Speaker release

Yes

Primary author: TANNENBAUM, Todd (University of Wisconsin Madison (US))

Presenter: TANNENBAUM, Todd (University of Wisconsin Madison (US))

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 23

Type: **not specified**

Reflections on HTC Needs and Directions

Thursday, September 26, 2019 5:40 PM (35 minutes)

Reflections on the common upcoming requirements and directions of the scientific high throughput computing community.

Desired slot length

30 minutes

Speaker release

Yes

Primary author: LIVNY, Miron (University of Wisconsin-Madison)

Presenter: LIVNY, Miron (University of Wisconsin-Madison)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 24

Type: **not specified**

From interactive to distributed computing of land parcel signatures using HTCondor

Wednesday, September 25, 2019 9:35 AM (25 minutes)

In the framework of the Common Agricultural Policy (CAP) of the European Union, a big technological shift is happening. For decades, the correct payment of subsidies to farmers was controlled by means of remotely sensed images, by doing visual interpretation and field visits, to assess that a randomly selected percentage of the land parcels respected all the rules. In recent years, we are witnessing a big evolution in Earth Observation, with the Copernicus Programme and the Sentinel satellites that provide coverage of every piece of land in high resolution every few days, together with the availability of cloud platforms capable of storing the big amount of data captured. Coupling this with the ease of use of cloud computing platforms and a wide number of tools for extracting valuable information from big geospatial datasets using machine learning techniques, puts the CAP controls sector in the edge of a revolution: from 2021, each single agricultural parcel in Europe will be constantly monitored for the full year. This involves the calculation and the constant update of temporal profiles that will monitor the vegetation status of land parcels in all their phases, from ploughing to sowing, from ripening to harvesting.

In this context, the JEODPP (Joint Research Centre Big Data Platform) group was involved in preliminary studies to assess the feasibility of the new CAP Monitoring. On a first stage, by accessing the full catalogue of Sentinel-2 images, we developed an interactive tool to calculate the Normalized Difference Vegetation Index profile for a single parcel at a time inside a JupyterLab notebook (s2explorer application).

When the algorithm was tested and verified, the need for scaling to regional or national level arose. This implied the need to process millions of vector polygons, each of them covered by more than 50 images per year, a perfect workspace for using the HTCondor workload manager services already available inside the JEODPP platform. The C++ routines developed for the interactive prototype were compiled in a standalone executable and the calculation was divided in three phases: 1) compilation of the list of satellite images involved in the selected spatial and temporal range, 2) creation of a job for each image, 3) collection of the result into a single binary file; a typical map-reduce schema. All these phases were performed by using HTCondor jobs, and, in particular, the second phase was heavy parallelized on the hundreds of cores available inside the JEODPP platform.

We executed a first test on a region in Hungary (10K parcels for a full year processed in less than half an hour) and then scaled to the full Catalonia (640K parcels processed in 4 hours).

The need to deeply evaluate the results of the batch processing generated the idea to “close the circle”, that is to provide an interactive tool to visually assess the calculations made by the HTCondor jobs. We developed a Python application running inside JupyterLab that could visualize all the land parcels involved in the calculation and, by clicking on each of them, immediately display the vegetation profile and the imageries extracted from each individual satellite acquisition date. The tool is widely used by the JRC D.5 unit and it is the base for the future characterization of the crops by means of machine learning algorithms, a key component of the CAP Monitoring.

This use case is an example of the involvement of HTCondor services in a complex environment where the need for interactive prototyping goes along with heavy distributed processing needs and contributes to create an integrated solution.

Desired slot length

Speaker release

No

Primary authors: Mr DE MARCHI, Davide (European Commission, Joint Research Centre (JRC) Directorate I. Competences. Unit I.3 Text and Data Mining); Mr WIRNHARDT, Csaba (European Commission, Joint Research Centre (JRC) Directorate D. Sustainable Resources. Unit D.5 Food Security); Mr SOILLE, Pierre (European Commission, Joint Research Centre (JRC) Directorate I. Competences. Unit I.3 Text and Data Mining)

Presenter: Mr WIRNHARDT, Csaba (European Commission, Joint Research Centre (JRC) Directorate D. Sustainable Resources. Unit D.5 Food Security)

Session Classification: Workshop presentations

Track Classification: HTCCondor presentations and tutorials

Contribution ID: 25

Type: **not specified**

HTCondor-CE Basics and Architecture

Thursday, September 26, 2019 10:05 AM (35 minutes)

An introduction to HTCondor-CE including an overview of its architecture and supported batch systems.

Desired slot length

30

Speaker release

Yes

Primary author: LIN, Brian Hua (University of Wisconsin - Madison)

Presenter: LIN, Brian Hua (University of Wisconsin - Madison)

Session Classification: Workshop presentations

Track Classification: HTCondor Compute Element (CE) presentations and tutorials

Contribution ID: 26

Type: **not specified**

HTCondor-CE Configuration

Thursday, September 26, 2019 11:40 AM (35 minutes)

An overview of HTCondor-CE configuration with an emphasis on Job Router semantics

Desired slot length

30

Speaker release

Yes

Primary author: LIN, Brian Hua (University of Wisconsin - Madison)

Presenter: LIN, Brian Hua (University of Wisconsin - Madison)

Session Classification: Workshop presentations

Track Classification: HTCondor Compute Element (CE) presentations and tutorials

Contribution ID: 27

Type: **not specified**

HTCondor-CE Integration with APEL and BDII

Thursday, September 26, 2019 12:40 PM (20 minutes)

An overview of the current support for APEL (accounting) and BDII (information services)

Desired slot length

20

Speaker release

Yes

Primary author: LIN, Brian Hua (University of Wisconsin - Madison)

Presenter: LIN, Brian Hua (University of Wisconsin - Madison)

Session Classification: Workshop presentations

Track Classification: HTCondor Compute Element (CE) presentations and tutorials

Contribution ID: 28

Type: **not specified**

HTCondor-CE Troubleshooting

Thursday, September 26, 2019 3:05 PM (45 minutes)

An overview of common issues encountered when supporting an HTCondor-CE, where to find additional information, and how to resolve them

Desired slot length

40

Speaker release

Yes

Primary author: LIN, Brian Hua (University of Wisconsin - Madison)

Presenter: LIN, Brian Hua (University of Wisconsin - Madison)

Session Classification: Workshop presentations

Track Classification: HTCondor Compute Element (CE) presentations and tutorials

Contribution ID: 29

Type: **not specified**

HTCondor-CE: What Next?

Thursday, September 26, 2019 3:50 PM (15 minutes)

Upcoming features and direction of HTCondor-CE

Desired slot length

15

Speaker release

Yes

Primary author: LIN, Brian Hua (University of Wisconsin - Madison)

Presenter: LIN, Brian Hua (University of Wisconsin - Madison)

Session Classification: Workshop presentations

Track Classification: HTCondor Compute Element (CE) presentations and tutorials

Contribution ID: 30

Type: **not specified**

HTCondor-CE status at CC-IN2P3

Thursday, September 26, 2019 12:15 PM (25 minutes)

At CC-IN2P3, CREAM-CEs are being used since grid was deployed. Now, that we are evaluating to move our grid production from UGE to HTCondor a logical step was to evaluate also HTCondor-CEs.

In this presentation we show our experience installing, configuring and monitoring the cluster, some results and also our future plans and requirements.

Desired slot length

Speaker release

Yes

Primary authors: HAMAR, Vanessa (CC-IN2P3); Dr VAMVAKOPOULOS, Emmanouil (Centre National de la Recherche Scientifique (FR)); ELOTO, Christelle; FOURNIALS, Nicolas (CC-IN2P3)

Presenter: HAMAR, Vanessa (CC-IN2P3)

Session Classification: Workshop presentations

Track Classification: HTCondor Compute Element (CE) presentations and tutorials

Contribution ID: 32

Type: **not specified**

Building a European wide, Bioinformatics jobs execution network

Friday, September 27, 2019 9:00 AM (25 minutes)

With more than 2,000 bioinformatics tools available, Usegalaxy.eu (<https://usegalaxy.eu>) is the biggest Galaxy instance in Europe covering most of the hottest bioinformatics topics and communities. After one year from its public launch into March 2018, Usegalaxy.eu has reached the important milestone of 5 million jobs executed and over 6 thousand registered users.

Several computer centers across Europe are currently sharing their remote computation power to support the Usegalaxy.eu load: IT, UK, CZ, DE, PT, ES,...

To create this network of shared computational resources, we leverage:

Pulsar (<https://pulsar.readthedocs.org>), a TES-like service written in Python that allows a Galaxy server to automatically interact with those remote systems,

VGCN (<https://github.com/usegalaxy-eu/vgcn>), a virtual image which has all of the required components to act as a galaxy compute node as part of an HTCCondor cluster.

Terraform (<https://github.com/usegalaxy-eu/terraform>), a set of scripts for safely and efficiently building the infrastructure into a modern cloud environment.

Galaxy's job destination framework allows job execution parameters to be determined dynamically at runtime, offering a flexible option for choosing the job endpoints, and the Pulsar layer ensures execution details are correctly exchanged to correctly perform jobs into the local and/or remote HTCCondor clusters.

Desired slot length

Speaker release

Yes

Primary authors: CUCCURU, Gianmauro (University of Freiburg); GRUENING, Bjoern (University of Freiburg); RASCHE, Helena (University of Freiburg)

Presenter: CUCCURU, Gianmauro (University of Freiburg)

Session Classification: Workshop presentations

Track Classification: HTCCondor presentations and tutorials

Contribution ID: 33

Type: **not specified**

Large-scale aerial photo processing for tree health monitoring with HTCondor

Friday, September 27, 2019 9:25 AM (25 minutes)

```

\documentclass{article}
\usepackage{filecontents}
\usepackage{authblk}
\usepackage{natbib}
\usepackage{natbib}
\bibliographystyle{abbrvnat}
\setcitestyle{numbers,open={},close={},citesep={,}}
\begin{filecontents}{\jobname.bib}
@article{phillips2008a,
author = {Phillips, S.J. and Dudík, M.},
title = {Modeling of species distributions with Maxent: new extensions and a comprehensive evaluation},
journal = {Ecography},
volume = {31},
number = {2},
year = {2008},
pages = {161–175},
url = {https://doi.org/10.1111/j.0906-7590.2008.5203.x},
language = {en}
}
@article{phillips2017a,
author = {Phillips, S.J. and Anderson, R.P. and Dudík, M. and Schapire, R.E. and Blair, M.E.},
title = {Opening the black box: an open-source release of Maxent},
journal = {Ecography},
volume = {40},
number = {7},
year = {2017},
pages = {887–893},
url = {http://doi.org/10.1111/ecog.03049},
language = {en}
}
@article{Soille,
author = {Soille, P. and Burger, A. and De Marchi, D. and Kempeneers, P. and Rodriguez, D. and Syrris, V. and Vasilev, V.},
title = {A versatile data-intensive computing platform for information retrieval from big geospatial data},
journal = {Future Gener. Comput. Syst},
volume = {81},
year = {2018},
pages = {30–40},
url = {https://doi.org/10.1016/j.future.2017.11.007},
language = {en}
}
@book{beck2019,
author = {Beck, P. S. A and Martínez-Sánchez, L. and Di Leo, M. and Chemin, Y. and Caudullo, G. and de la Fuente, B. and Zarco-Tejada, P. J.},
title = {The Canopy Health Monitoring (CanHeMon) project},

```

```

institution = {Joint Research Centre (European Commission)},
publisher = {Publications Office of the European Union, Luxembourg},
year = {2019},
isbn = {978-92-79-99639-9},
doi = {10.2760/38697},
pages = {79},
language = {en},
}
@book{marti,
author = {Beck, P. S. A and Martínez-Sanchez, L. and Di Leo, M. and Chemin, Y. and Caudullo, G.
and de la Fuente, B. and Zarco-Tejada, P. J.},
title = {Remote Sensing in support of Plant Health Measures–Findings from the Canopy Health
Monitoring},
institution = {Joint Research Centre (European Commission)},
publisher = {Publications Office of the European Union, Luxembourg},
year = {2019},
isbn = {978-92-76-02051-6},
doi = {10.2760/767468},
pages = {13},
language = {en},
}
@article{haralick1973a,
author = {Haralick, R.M. and Shanmugam, Kand Dinstein,I.},
title = {Textural Features for Image Classification},
journal = {IEEE Transactions on Systems, Man, and Cybernetics},
volume = {3},
year = {1973},
pages = {610–621},
url = {https://doi.org/10.1109/TSMC.1973.4309314},
number = {6},
language = {en}
}
@article{delafuente,
author = {de la Fuente, B. and Saura, S. and Beck, P. S. A.},
title = {Predicting the spread of an invasive tree pest: the pine wood nematode in Southern Europe},
journal = {Journal of Applied Ecology},
volume = {55},
year = {2018},
pages = {2374-2385},
doi = {10.1111/1365-2664.13177},
number = {5},
language = {en}
}
\end{filecontents}
\usepackage[utf8]{inputenc}

\begin{document}
\title{Large-scale aerial photo processing for tree health monitoring with HTCCondor}

\author1{Martinez-Sanchez, Laura}
\author1{Rodriguez-Aseretto, Dario}
\author1{Soille, Pierre}
\author1{Beck, Pieter S. A.}
\affil1{European Commission, Joint Research Centre (JRC)}
\date{September 2019}
\maketitle
\begin{abstract}

```

The Canopy Health Monitoring (CanHeMon) project ran at the Joint Research Centre of the European Commission from mid-2015 to mid-2018 and was funded by DG SANTE. DG SANTE is responsible, among other things, for the European Union's Plant Health legislation, which aims to put in place effective measures to protect the Union's territory and its plants, as well as ensuring trade is safe and the impacts of climate change on the health of EU crops and forests are mitigated. For specific harmful organisms that threaten its crops and forests, the EU takes emergency control measures. The Pine wood nematode (*Bursaphelenchus xylophilus*) is such a quarantine pest. It can kill European coniferous tree species and is spreading through Portugal since the end of the 1990s.

As part of the EU emergency measures against the pine wood nematode (PWN) (*Bursaphelenchus xylophilus*) Decision 2012/535/EU, Portugal should perform, outside and during the flight season of the PWN's vector, surveys of coniferous trees located in the 20 km wide buffer zone established along the Spanish border, with the aim to detect trees which are dead, in poor health or affected by fire or storm. These trees shall be felled and removed to avoid that they act as attractants for the longhorn beetle (*Monochamus* sp), the insect vector responsible for the spread of PWN [delafuente]. The CanHeMon project tasked the Joint Research Centre with analysing a portion of the buffer zone, using remote sensing data, to support detection on the ground of declining pine trees. During the project, a 400 km² area was imaged twice, in autumn 2015 and autumn 2016, at 15 cm resolution from aircraft, and individual declining tree crowns were detected using a MaxEnt-based [phillips2017a,phillips2008a], iterative image analysis algorithm, the performance of which was gauged through visual photointerpretation. The scalability of the automated methods was then tested using an image mosaic of the entire buffer zone at 30 cm resolution.

We sought an image analysis platform that could efficiently handle and parallelise the computations on the large (terabyte) volumes of image data in this project. The JRC Earth Observation Data and Processing Platform (JEODPP) [Soille], which was developed in parallel with the CanHeMon project [beck2019,marti], increasingly met these needs over the course of the project. Being an inhouse service of the EC, it facilitates processing of the data for which the licensing does not permit public distribution. It is a versatile platform that brings the users to the data through web access and allows for large-scale batch processing of scientific workflows, remote desktop access for fast prototyping in legacy environments, and interactive data visualisation/analysis with JupiterLab.

The storage and processing nodes underlying the JEODPP infrastructure consist of commodity hardware equipped with a stack of open source software. The storage service relies on the CERN EOS distributed file system which provides a disk-based, low latency storage service suitable for multipetabyte scale data. EOS is built on top of the XRootD protocol developed for high energy physics applications but also offers almost fully POSIX compliant access through a dedicated FUSE client called FUSEX that is suitable for other areas. As of summer 2019, the storage capacity of the EOS distributed file system of the JEODPP amounts to 14 PiB corresponding to a net capacity of 7 PiB given that all data are replicated once to ensure their availability and decrease the likelihood of data loss in case of disk failure. For all other services, the JEODPP relies on processing servers with a total of 2,200 cores distributed over 64 nodes. On average, 15 GB of RAM is available to each core. The batch processing service, called JEO-batch, is orchestrated with HTCondor. All applications running on the JEODPP are deployed within Docker containers to ease the management of applications having conflicting requirements in terms of library versions. Docker images are created by combining and modifying standard images downloaded from repositories. For the canopy health monitoring application, we created a Debian image with all the libraries needed to run the code (mainly R and GDAL libraries).

Covering the entire PWN buffer zone with 4-band images of 30 cm resolution, stored in 8-bit, generates 2.4 TB of data. The associated texture layers [haralick1973a] used in the analyses here added an additional 50 TB. The data were delivered and processed in 24,904 tiles measuring

1 km by 1 km. Processing a single tile in each iteration takes 40 to 55 minutes, with a memory usage of 5-7 GB on a regular CPU (with 2 to 8 cores). Processing the entire buffer zone on a single CPU would thus take more than a year. Assigning all of the 2,200 cores of the JEODPP to batch processing service and submitting the job with HTCCondor, the task would be completed in less than two hours. In practice, between 100 and 500 cores of the JEODPP were used in the processing at any one time. The results of this processing were used to support management of the area on the ground and make recommendations on the use of remote sensing for large-area surveys in the context of plant health.

```
\end{abstract}
```

```
\bibliography{\jobname}
```

```
\end{document}
```

Desired slot length

15

Speaker release

Yes

Primary authors: MARTINEZ SANCHEZ, Laura (JRC); RODRIGUEZ ASERETTO, Dario; Mr SOILLE, Pierre (European Commission, Joint Research Centre (JRC) Directorate I. Competences. Unit I.3 Text and Data Mining); Dr BECK, Pieter (European Commission, Joint Research Centre (JRC))

Presenter: MARTINEZ SANCHEZ, Laura (JRC)

Session Classification: Workshop presentations

Track Classification: HTCCondor presentations and tutorials

Contribution ID: 34

Type: **not specified**

Dynamic Batch service with HTCondor and Kubernetes

Thursday, September 26, 2019 9:00 AM (25 minutes)

One of the main challenges the JRC Big Data Platform-JEODPP 1 is to offer well consolidated computational services, such as the batch system or interactive data visualization on which users can process large scale geospatial data while ensuring a smooth user experience combined with easy administration of all resources from hardware to applications.

Due to the heterogeneity of the user requirements of the JEODPP platform, many services have dynamic demand over time. As an example users who want to visualise data interactively requires a considerable amount of resources just in short periods, mainly during the core time of the day. Considering that the batch service should use the full capacity whenever possible, the resources should be allocated dynamically. Moreover keeping a physical separation between htcondor, interactive nodes, and all our other services as in our previous setting is not satisfactory from service management, monitoring, and performance perspectives. In that fixed setting, a lot of idle time when some operation was ended was added up to a lack of resources in the peak of need. To address these issues, we are moving to implement all our services under the Kubernetes control.

Running batch workloads in Kubernetes is possible [2] and Kubernetes has a specific resource for that purpose (Job) but it is quite limited. To address this limitation, we present in this work the implementation of HTCondor within Kubernetes. HTCondor daemons are packaged inside pre-configured Docker images and deployed as a service in a container cluster handled by Kubernetes. As a proof of concept, we have successfully run an actual JEODPP use-case (atmospheric correction with Sen2Cor [3]) on Kubernetes and compared its performance with a standard HTCondor pool. Finally, we present some considerations to implement other services based on the namespaces of Kubernetes and some constraints to implement this solution in a production environment such as the JEODPP.

References

[1] P. Soille, A. Burger, D. Rodriguez, V. Syrris, and V. Vasilev.; *Towards a JRC earth observation data and processing platform* Proc. of the 2016 Conference on Big Data from Space (BiDS'16), pages 65-68, 2016. Available from doi: [10.2788/854791][4]

[2] D. Aiftimiei, T. Boccali, M. Panella, A. Italiano, G. Donvito, D. Michelotto, M. Caballer et al. *Geographically distributed Batch System as a Service: the INDIGO-DataCloud approach exploiting HTCondor*. In J. Phys. Conf. Ser., vol. 898, p. 052033. 2017.

[3] M. Main-Knorn, B. Pflug, J. Louis, V. Debaecker, U. Müller-Wilm and F. Gascon, *Sen2Cor for Sentinel-2*. In Image and Signal Processing for Remote Sensing XXIII, vol. 10427, p. 1042704. International Society for Optics and Photonics, 2017.

Desired slot length

Speaker release

Yes

Primary authors: RODRIGUEZ, Dario (European Commission - DG JRC); Mr MARLETTA, Luca (European Commission, Joint Research Centre (JRC)); Mr SOILLE, Pierre (European Commission, Joint Research Centre (JRC) Directorate I. Competences. Unit I.3 Text and Data Mining)

Presenter: Mr MARLETTA, Luca (European Commission, Joint Research Centre (JRC))

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 35

Type: **not specified**

The Natural Language Processing automatization at JRC for Knowledge Management

Thursday, September 26, 2019 9:25 AM (20 minutes)

In our Directorate F Health, Consumers and Reference Materials at JRC, the Knowledge for Health and Consumer Safety Unit F.7 deals with anticipating knowledge needs, mapping knowledge gaps and suggesting research topics to be carried out in the Directorate and possibly in the JRC. For example, thousands of publications are released every year on different topics where JRC has strong competence and a mandate for scientific advice to European Commission, keep the pace of this continuous growing knowledge is an increasing challenge. The velocity of literature production makes impossible to deliver state of the art answers for EC policy makers without automating the whole process. The only way to face this issue is to apply Machine Learning tools (AI) in the field of Natural Language. The same approach can be used on extracts from raw text, including those from speeches, to reveal sentiments and feelings that can be used to understand trends and (political) shifts that may improve JRC insight in policy developments. We will present examples of benchmarking use of this approach, made by combining tools like Natural Language Understanding IBM Watson and AllenNLP Machine Comprehension models, installed locally at the JRC.

Desired slot length

15 minutes

Speaker release

Yes

Primary authors: Mr PUERTAS GALLARDO, Antonio (European Commission Joint Research Centre); Mr PETRILLO, Mauro (European Commission Joint Research Centre)

Presenter: Mr PETRILLO, Mauro (European Commission Joint Research Centre)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 36

Type: **not specified**

Welcome, introduction to JRC

Tuesday, September 24, 2019 9:00 AM (20 minutes)

Desired slot length

Speaker release

Primary author: VAN DEN EEDE, Guy (EC-JRC)

Presenter: VAN DEN EEDE, Guy (EC-JRC)

Session Classification: Logistics JRC

Track Classification: Miscellaneous

Contribution ID: 37

Type: **not specified**

Workshop logistics

Tuesday, September 24, 2019 9:20 AM (10 minutes)

Desired slot length

Speaker release

Primary author: PUERTAS GALLARDO, Antonio (European Commission)

Presenter: PUERTAS GALLARDO, Antonio (European Commission)

Session Classification: Logistics JRC

Track Classification: Miscellaneous

Contribution ID: **38**

Type: **not specified**

Workshop wrap-up

Friday, September 27, 2019 12:10 PM (30 minutes)

Desired slot length

Speaker release

Primary author: MEINHARD, Helge (CERN)

Presenter: MEINHARD, Helge (CERN)

Session Classification: Miscellaneous

Track Classification: Miscellaneous

Contribution ID: 39

Type: **not specified**

Virgo: a computing evolution tale

Thursday, September 26, 2019 5:15 PM (25 minutes)

Desired slot length

Speaker release

Primary author: Mr FRONZE', Gabriele Gaetano (University e INFN Torino (IT), Subatech Nantes (FR))

Presenter: Mr FRONZE', Gabriele Gaetano (University e INFN Torino (IT), Subatech Nantes (FR))

Session Classification: Workshop presentations

Track Classification: HTCondor and HTCondor CE office hours

Contribution ID: 40

Type: **not specified**

HTCondor at DESY: GPUs, Jupyter, Docker, FaaS

Wednesday, September 25, 2019 11:25 AM (20 minutes)

Replay of second part of DESY presentation of Tue 24 Sep

Desired slot length

Speaker release

Presenter: BEYER, Christoph

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials