

European HTCondor Workshop 2018



Report of Contributions

Contribution ID: 1

Type: **not specified**

HTCondor configuration with puppet

Wednesday, 5 September 2018 16:30 (25 minutes)

Configuring a condor cluster and keeping the configuration synchronised can be quite the chore. For this purpose, under the umbrella of HEP-Puppet, sysadmins have gathered to create a simple-to-use Puppet module. With just a few lines of YAML (hiera) you can configure your own HTCondor cluster within minutes (Puppet infrastructure provided). This talk will showcase the module with snippets from a real WLCG site configuration.

Primary author: Dr KRECZKO, Lukasz (University of Bristol (GB))

Co-authors: MOHAMMAD, Kashif (University of Oxford); WIENEMANN, Peter (University of Bonn (DE)); FREYERMUTH, Oliver; SCHAER, Frederic (CEA)

Presenter: Dr KRECZKO, Lukasz (University of Bristol (GB))

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 2

Type: **not specified**

HTCondor command line monitoring tool

Wednesday, 5 September 2018 09:00 (25 minutes)

The University of Oxford Tier-2 Grid cluster converted to using HTCondor in 2014. At that time, there was no suitable monitoring tool available. The Oxford team developed a command line tool, written in Python, that displays snapshot information about the running jobs. The tool provides the capability of reporting on the number of jobs running on a given node and the efficiency of each job. Further development resulted in a web-based display, which continuously updates the status of jobs running on the cluster. Details of the development of the tool and its features will be presented.

Primary author: Mr VIPUL, Davda (University of Oxford)

Presenter: Mr VIPUL, Davda (University of Oxford)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 3

Type: **not specified**

ClassAd Language Tutorial

Tuesday, 4 September 2018 14:35 (35 minutes)

HTCondor uses the ClassAd language in three different ways. This tutorial will cover the full syntax of the ClassAd language, the uses in HTCondor, and advanced topics in ClassAd usages for system administration and monitoring.

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: 4

Type: **not specified**

HTCondor Administration Tutorial

Tuesday, 4 September 2018 16:05 (1h 5m)

This tutorial covers the basic installation and configuration of the HTCondor system. Theory of operation, and system architecture is also covered.

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: 5

Type: **not specified**

Configuring Group Quotas, Policies, and Fair Share across Users with the HTCondor Negotiator

Wednesday, 5 September 2018 11:30 (1h 5m)

This tutorial covers HTCondor's "Fair Share" mechanisms for assigning resources to users, configuring groups of users with quotas, and other aspects of global policy via the HTCondor negotiator.

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: 6

Type: **not specified**

Networking Concepts in HTCondor

Wednesday, 5 September 2018 14:50 (35 minutes)

How HTCondor deals with network architecture difficulties.

Primary author: FREY, James (University of Wisconsin Madison (US))

Presenter: FREY, James (University of Wisconsin Madison (US))

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 7

Type: **not specified**

HTCondor-CE Overview and Architecture

Wednesday, 5 September 2018 10:00 (35 minutes)

The HTCondor-CE provides a remote API on top of a local site batch system.

Primary author: FREY, James (University of Wisconsin Madison (US))

Presenter: FREY, James (University of Wisconsin Madison (US))

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 8

Type: **not specified**

HTCondor Annex: Elasticity into the Public Cloud

Thursday, 6 September 2018 09:00 (35 minutes)

Learn how the Annex allows you to seamlessly expand your HTCondor pool using machines from Amazon EC2.

Primary author: FREY, James (University of Wisconsin Madison (US))

Presenter: FREY, James (University of Wisconsin Madison (US))

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 9

Type: **not specified**

Workflows with HTCondor's DAGMan

Thursday, 6 September 2018 14:00 (35 minutes)

DAGMan lets you manage large, complex workflows in HTCondor.

Primary author: FREY, James (University of Wisconsin Madison (US))

Presenter: FREY, James (University of Wisconsin Madison (US))

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 10

Type: **not specified**

Scaling HTCondor at CERN

Wednesday, 5 September 2018 11:05 (25 minutes)

HTCondor has been the primary production batch service at CERN for the last couple of years, passing the 100k core mark last year. The challenge has been to scale the service, in terms of course of the number of resources, but also in terms of the number of heterogenous use cases. The use cases involve dedicated LHC Tier-0 pools, dedicated resources within standard pools, special CE routes to dedicated Clouds and Storage pools, and managing a diverse user community. This talk will go through some of the different use cases, the technical decisions that have been taken, and the challenges that have been encountered.

Primary author: JONES, Ben (CERN)

Presenter: JONES, Ben (CERN)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 11

Type: **not specified**

Haggis: Accounting Group Management at CERN

Wednesday, 5 September 2018 14:25 (25 minutes)

Haggis is an information system used to map CERN users to HTCondor accounting groups as well as hold information about quota and priority allocation per accounting group as well as information relevant to resource usage accounting. It enforces a tree-like domain model that supports resource mapping under different compute pools. All the data stored in Haggis is completely manageable by the appropriate parties via a RESTful CRUD API, as well as a CLI client.

The data needed for HTCondor to operate can be injected into the system by using Haggis' delivery mechanism to generate the appropriate configuration files.

Haggis is based on a modular, layered and pluggable architecture that allows implementations with different delivery and management mechanisms, backend storage systems as well as different authorization policies. Thus, it can be easily tailored to accommodate different use cases and needs of different HTCondor setups.

In this presentation we will talk about how CERN uses Haggis to fit its accounting group management needs. Moreover we will demonstrate its software architecture and discuss the ways that Haggis can be modified, extended and deployed in order to be used in different HTCondor environments.

Primary authors: Mr TRIANTAFYLLIDIS, Nikolaos Petros (CERN); JONES, Ben (CERN)

Presenter: Mr TRIANTAFYLLIDIS, Nikolaos Petros (CERN)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 12

Type: **not specified**

Key Challenge Areas for Distributed High Throughput Computing

Thursday, 6 September 2018 09:35 (35 minutes)

Based on current trends and past experience, this talk will identify and discuss six key challenge areas that will continue to drive High Throughput Computing technologies innovation in the years to come.

Primary author: TANNENBAUM, Todd (Univ of Wisconsin-Madison, Wisconsin, USA)

Presenter: TANNENBAUM, Todd (Univ of Wisconsin-Madison, Wisconsin, USA)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 13

Type: **not specified**

What's New in HTCondor?

Wednesday, 5 September 2018 09:25 (35 minutes)

An overview of recent developments and future plans in HTCondor.

Primary author: TANNENBAUM, Todd (Univ of Wisconsin-Madison, Wisconsin, USA)

Presenter: TANNENBAUM, Todd (Univ of Wisconsin-Madison, Wisconsin, USA)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 14

Type: **not specified**

SciTokens: Moving away from identity credentials to capability tokens in HTCondor

Thursday, 6 September 2018 14:35 (35 minutes)

We believe that distributed, scientific computing community has unique authorization needs that can be met by utilizing common web technologies, such as OAuth 2.0 and JSON Web Tokens (JWT). The SciTokens team, a collaboration between technology providers including the HTCondor Project and domain scientists, is working to build and demonstrate a new authorization approach at scale.

Primary authors: TANNENBAUM, Todd (University of Wisconsin Madison (US)); BASNEY, Jim; BOCK-ELMAN, 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: 15

Type: **not specified**

Monitoring HTCondor

Friday, 7 September 2018 09:00 (25 minutes)

An overview of monitoring an HTCondor pool

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: 16

Type: **not specified**

Put your jobs into a box: Using Cgroups, Docker, Singularity

Friday, 7 September 2018 09:25 (35 minutes)

Overview of HTCondor's mechanisms in support of job isolation, including Docker, Singularity, cgroups, and namespace mounts.

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**

What defines a workload as High Throughput Computing

Tuesday, 4 September 2018 17:10 (35 minutes)

Distinguishing characteristics of High Throughput Computing (HTC), including how it contrasts with High Performance Computing (HPC). When is HTC appropriate, when is HPC appropriate? Also lessons and best practices learned from experiences running the Open Science Grid, a 100+ institution distributed HTC environment.

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: **18**

Type: **not specified**

Job and Machine Policy

Thursday, 6 September 2018 16:55 (45 minutes)

Discussion of policy expressions available to users when they submit their HTCondor jobs, and expressions available to Administrators when they configure HTCondor execute nodes. Time permitting, there will be a demonstration of special purpose execution slots.

Request 60 Minute slot.

Primary author: KNOELLER, John (University of Wisconsin-Madison)

Presenter: KNOELLER, John (University of Wisconsin-Madison)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 19

Type: **not specified**

Schedd Job Transforms

Friday, 7 September 2018 11:20 (35 minutes)

Discussion of the Job Transform language in the HTCondor Schedd.

Request 30 Minute time slot.

Primary author: KNOELLER, John (University of Wisconsin-Madison)

Presenter: KNOELLER, John (University of Wisconsin-Madison)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 20

Type: **not specified**

Config and Submit language

Thursday, 6 September 2018 11:55 (35 minutes)

Discussion of the language used by HTCondor for configuration and job submit files.

Request 30 Minute time slot.

Primary author: KNOELLER, John (University of Wisconsin-Madison)

Presenter: KNOELLER, John (University of Wisconsin-Madison)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 21

Type: **not specified**

Using Python to monitor and control HTCondor

Wednesday, 5 September 2018 15:25 (35 minutes)

Introduction to the HTCondor python bindings and their use to query HTCondor.

Primary author: KNOELLER, John (University of Wisconsin-Madison)

Presenter: KNOELLER, John (University of Wisconsin-Madison)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 22

Type: **not specified**

Using Python to submit jobs

Wednesday, 5 September 2018 16:55 (35 minutes)

Tutorial on using python to submit jobs to HTCondor, concentrating on the 8.7 series improvements in the HTCondor python bindings.

Primary author: KNOELLER, John (University of Wisconsin-Madison)

Presenter: KNOELLER, John (University of Wisconsin-Madison)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 23

Type: **not specified**

DESY Features on Top of HTCondor

Wednesday, 5 September 2018 14:00 (25 minutes)

The talk provides some details of special DESY configurations. It focuses on features we need for user registry integration, node maintenance operations and fair share / quota handling. With the help of job transformations defining job classes and proper job duration and memory setting, we setup a smooth and transparent operating model.

Primary author: FINNERN, Thomas (DESY)

Co-authors: BEYER, Christoph; KEMP, Yves (Deutsches Elektronen-Synchrotron (DE))

Presenter: FINNERN, Thomas (DESY)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 24

Type: **not specified**

Managing Cluster Fragmentation using ConcurrencyLimits

Tuesday, 4 September 2018 15:10 (25 minutes)

Clusters running differently sized jobs can easily suffer from fragmentation: Large chunks of free resources are required to run larger jobs, but smaller jobs can block parts of these chunks, making the remainder too small. For example, clusters in the WLCG must provide space for 8-core jobs, while there is a constant pressure of 1-core jobs. Common approaches to this issue are the DEFRAG daemon, custom scheduling ordering, and delays that protect free chunks.

At the GridKa Tier 1 cluster, providing roughly 30.000 cores and growing, we have developed a new approach to stay responsive and efficient at large scales. By tagging new jobs during submission, we can manage job groups using HTCondor's inbuilt ConcurrencyLimit feature. So far, we have successfully used this to enforce fragmentation limits for small jobs in our production environment.

This contribution highlights the challenges of fragmentation in large scale clusters. Our focus is on scalability and responsiveness on the one hand, as well as maintainability and configuration overhead on the other hand. We show how our approach integrates with regular scheduling policies, and how we achieve proper utilisation without micromanaging individual resources.

Primary author: FISCHER, Max (GSI - Helmholtzzentrum für Schwerionenforschung GmbH (DE))

Co-authors: ALEF, Manfred (Karlsruhe Institute of Technology (KIT)); PETZOLD, Andreas (KIT - Karlsruhe Institute of Technology (DE))

Presenter: FISCHER, Max (GSI - Helmholtzzentrum für Schwerionenforschung GmbH (DE))

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 25

Type: **not specified**

Bringing together HTCondor, Python, and Jupyter

Wednesday, 5 September 2018 17:30 (30 minutes)

Miron Livny would like to lead a discussion on how to best interface with HTCondor when working inside a Python environment, especially an interactive science-based environment such as Jupyter Notebook / Lab. We have been experimenting with some approaches at UW-Madison that we can share, but what we are looking for an open discussion of ideas, feedback, and suggestions.

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: 26

Type: **not specified**

Day-to-day HTCondor Operations at RAL

Thursday, 6 September 2018 10:10 (25 minutes)

RAL Tier-1 originally used the PBS batch system for its Grid related activities. Increased LHC operation requirements exposed scalability problems, therefore other batch systems were taken into consideration.

In this presentation we review the history of HTCondor at RAL and detail on how it evolved from an initial conventional setup with cgroups for resource control to current use of Docker containers that presents its own set of challenges.

We describe the integration of the batch farm with the Ceph storage system by means of dedicated Docker containers, and we discuss our experience with jobs bursting into the RAL cloud.

The presentation also comprises our consolidation plans, including future needs, especially ensuring a sustained number of multicore jobs on the batch farm.

Primary author: KELLY, John (S)

Co-author: CONDURACHE, Catalin (Science and Technology Facilities Council STFC (GB))

Presenter: KELLY, John (S)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 27

Type: **not specified**

General Integration Issues

Thursday, 6 September 2018 11:05 (25 minutes)

HTCondor is a product, but it is not an application. Like operating systems, networks, database management systems, and security infrastructures, HTCondor is a general system, upon which other applications may be built.

Extra work is needed to create something useful from HTCondor. The extra work depends on the goals of the designer. This talk identifies a few general areas that need to be addressed and gives specific ways that they were actually solved when adapting HTCondor to work in the grid environment.

Primary author: Mr JONES, Stephen (GridPP/Liverpool)

Presenter: Mr JONES, Stephen (GridPP/Liverpool)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 28

Type: **not specified**

Cloud scavenging with HTCondor in the EOSCpilot Fusion Science Demonstrator

Thursday, 6 September 2018 11:30 (25 minutes)

Access to both HTC and HPC facilities is vitally important to the fusion community, not only for plasma modelling but also for advanced engineering and design, materials research, rendering, uncertainty quantification and advanced data analytics for engineering operations. The computing requirements are expected to increase as the community prepares for the next generation facility, ITER. Moving to a decentralised computing model is vital for future ITER analysis where no single site will have sufficient resource to run all necessary workflows.

PROMINENCE is one of the Science Demonstrators in the European Open Science Cloud for Research Pilot Project (EOSCpilot) and aims to demonstrate that the fusion community can make use of distributed cloud resources. Here we will describe our proof-of-concept system, leveraging HTCondor, which enables users to submit both HTC and HPC jobs using a simple command line interface or RESTful API and run them in containers across a variety of cloud sites, ranging from local cloud resources, EGI FedCloud sites through to public clouds.

Primary author: Dr LAHIFF, Andrew (UKAEA)

Presenter: Dr LAHIFF, Andrew (UKAEA)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 29

Type: **not specified**

A versatile environment for large-scale geospatial data processing with HTCondor

Thursday, 6 September 2018 16:30 (25 minutes)

Geospatial data are one of the core data sources for scientific and technical support to the European Commission (EC) policies. For instance, the Copernicus programme of the European Union provides a vast amount of Earth Observation (EO) data for monitoring the environment through the Sentinel satellites operated by the European Space Agency. In terms of data management and processing, big geospatial data streams and other data sources have motivated the development of a petabyte-scale computational platform at the EC Joint Research Centre (JRC). This platform is called the JRC Earth Observation Data and Processing Platform (JEODPP) [1]. Thematic applications at the JRC rely on a variety of data sources each with their own data formats and protocols. In addition, experts from different domains build on different software, tools and libraries, making difficult knowledge sharing and the reproducibility of the experimental results. Taking into consideration all these challenges, the JEODPP has been designed by following the principles of modularity, parallelization and virtualization/containerization. In this way, it provides a flexible working environment where the users are able to deploy and optimize software and algorithmic workflows specialized for their tasks while fostering knowledge and data sharing.

Although there is no constraint on the type of data that can be processed, the main focus of the platform is currently on geospatial analysis and on the processing of satellite images. The Sentinel satellites are following a series of fixed orbits with image data delivered on a continuous basis and with a revisit time depending on the Sentinel mission type. The image data are stored in the form of flat files with each file mapping a given portion of the Earth surface. This drove both the architectural decisions and the physical/logical implementations regarding the JEODPP set up. In particular, the platform supports batch processing via mainly high-throughput computing where large collections of files are processed in parallel. Besides the batch farm, JEODPP offers other services such as interactive data analysis and visualization, data sharing, data storage, remote desktop access and experimental results dissemination. The operation of all these services is based on Docker containerisation.

HTCondor was chosen as workload manager, a versatile and robust job scheduler. Taking advantage of the Docker universe that HTCondor inherently supports, massive batch processing runs successfully on JEODPP since 2016. Besides, HTCondor functionalities allow a flexible combination of both types of nodes, workers, and managers. For example, it is possible for the user to submit jobs from different nodes, containers, or IPython notebooks using varying methods for authentication. Since it requires no external services for storage, HTCondor can use both the local and the network file system such as the EOS open source storage solution developed by CERN and deployed on the JEODPP. In practice, HTCondor shares features of a resource manager combined with those of a job scheduler. By integrating these features into a single system, it allows complex policy configurations and sophisticated optimizations. In this presentation, we show two applications that fully rely on HTCondor as workload manager and provide suggestions and lessons learnt based on our experience.

- Mosaicking Copernicus Sentinel-1 Data at Global level [2,3]: An algorithmic workflow for producing mosaics based on the dual polarisation capability of Sentinel-1 SAR imagery;
- Optimizing Sentinel-2 image selection in a Big Data Context [4]: An optimization scheme that selects a subset of the Sentinel-2 archive in order to reduce the amount of processing, while retaining the quality of the resulting output. As a case study, the focus is on the creation of a cloud-free composite, covering the global land mass and based on all the images acquired from January 2016 until September 2017.
- Marine ecosystem modelling in the SEACOAST project comprises types of modelling codes that are relevant to Marine Framework Strategy Directive [5], implemented on different spatial and

temporal scales, complemented by essential data (bathymetry, initial, boundary forcing, in and output) that are inherently coupled to each other. These models are implemented as an MPI application based on FORTRAN and it is running by using the parallel universe of HTCondor. We add a network file system NetApp beside EOS, which improves the performance of the MPI jobs over 80%.

In the near future, the possibility to combine HTCondor with Apache Mesos will be investigated. The aim is to provide a flexible, reconfigurable and extendable infrastructure to cover a wide range of different scientific computing use cases like HTC, HPC, Big Data analytics, GPU acceleration and Cloud technologies.

References

- [1] P. Soille, A. Burger, D. De Marchi, D. Rodriguez, V. Syrris, and V. Vasilev.; *A versatile data-intensive computing platform for information retrieval from big geospatial data*; Future Generation of Computer System, pages 30-40, 2018. Available from: <https://doi.org/10.1016/j.future.2017.11.007>
- [2] V. Syrris, C. Corbane, and P. Soille; *A global mosaic from Copernicus Sentinel-1 data* in Proc. Big Data Space, 2017, pp. 267–270. Available from: <http://dx.doi.org/10.2760/383579>
- [3] V. Syrris, C. Corbane, M. Pesaresi, and P. Soille; *A global mosaic from Copernicus Sentinel-1 data* IEEE Tr. on Big Data. Available from: <http://dx.doi.org/10.1109/TBDATA.2018.2846265>
- [4] P. Kempeneers and P. Soille.; *Optimizing Sentinel-2 image selection in a Big Data context*; Big Earth Data, pages 145-148, 2017. Available from: <https://doi.org/10.1080/20964471.2017.1407489>
- [5] D. Macias and E. Garcia-Gorriz and A. Stips.; *Productivity changes in the Mediterranean Sea for the twenty-first century in response to changes in the regional atmospheric forcing* Frontiers in Marine Science, pages 70, 2015. Available from: <https://doi.org/10.3389/fmars.2015.00079>

Primary authors: Dr RODRIGUEZ ASERETTO, Dario (European Commission); Dr SOILLE, Pierre (European Commission)

Presenter: Dr RODRIGUEZ ASERETTO, Dario (European Commission)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 30

Type: **not specified**

Sharing group-owned clusters at the via a 'super'-collector/negotiator at the UNIMI Physics Department.

Friday, 7 September 2018 10:00 (25 minutes)

A setup to share clusters that used to be owned and operated by experimental and theory sub-groups in the Physics Department of the University of Milan is described. Each sub-cluster is configured as a separate Condor Pool, reporting to one additional 'super'-collector. With a few assumptions on the available execution environment, plus mutually agreed priorities for 'local' jobs, this allows to match pending jobs to available slots in other clusters.

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: 31

Type: **not specified**

Pushing HTCondor boundaries: the CMS Global Pool experience

Thursday, 6 September 2018 15:10 (25 minutes)

In recent times, the CMS HTCondor Global Pool, which unifies access and management to all CPU resources available to the experiment, has been growing in size and evolving in its complexity, as new resources and job submit nodes are being added to the design originally conceived to serve the collaboration during the LHC Run 2. Having achieved most of our milestones for this period, the pool performs efficiently according to our present needs. However, looking into the coming years, and particularly into the HL-LHC era, a number of challenges are being identified and preliminarily explored. In this contribution we will present our current Global Pool setup and operational experience and how it is expected to extrapolate to meet the near and long-term future challenges.

Primary authors: PEREZ-CALERO YZQUIERDO, Antonio (Centro de Investigaciones Energéticas Medioambientales y Tecno); LETTS, James (Univ. of California San Diego (US)); DAVILA FOYO, Diego (Autonomous University of Puebla (MX))

Presenter: PEREZ-CALERO YZQUIERDO, Antonio (Centro de Investigaciones Energéticas Medioambientales y Tecno)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 32

Type: **not specified**

The CMS global pool, from a pilot-based to an heterogeneous mix of resources

Thursday, 6 September 2018 15:35 (25 minutes)

Nowadays computational resources come in a wide variety of forms from pilots running on sites, cloud resources and spare cycles on desktops, laptops and even phones through volunteer computing and our duty, as the Submission Infrastructure team at CMS, is to be able to use them all. When it comes to Integrate these different models into a single pool of resources, different challenges arise. In this talk we will talk about some of these cases, and how we have faced them using the flexibility provided by HTCondor.

Primary authors: DAVILA FOYO, Diego (Autonomous University of Puebla (MX)); LETTS, James (Univ. of California San Diego (US)); PEREZ-CALERO YZQUIERDO, Antonio (Centro de Investigaciones Energéticas Medioambientales y Tecnológicas)

Presenter: DAVILA FOYO, Diego (Autonomous University of Puebla (MX))

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 33

Type: **not specified**

Building a LIGO HTCondor site on top of a shared HPC cluster

Friday, 7 September 2018 10:55 (25 minutes)

All members of the LIGO Scientific Collaboration have access to a handful of dedicated LIGO Data Grid clusters which feature HTCondor, system-installed software, the LIGO and Virgo data, and other standard components. Cardiff University also host a LIGO Data Grid Site, but this is built on top of the shared institutional HPC cluster. In this talk I describe how I used HTCondor, Spack, Singularity, and CVMFS to create a LIGO site that aims to provide users with a “no-surprises” experience.

Primary author: Dr HOPKINS, Paul (Cardiff University)

Presenter: Dr HOPKINS, Paul (Cardiff University)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 34

Type: **not specified**

Welcome

Tuesday, 4 September 2018 14:00 (15 minutes)

Short introduction to UKRI, STFC, RAL

Primary author: SANSUM, Andrew (STFC)

Presenter: SANSUM, Andrew (STFC)

Session Classification: Workshop presentations

Contribution ID: 35

Type: **not specified**

Logistics

Tuesday, 4 September 2018 14:25 (10 minutes)

Workshop logistics

Primary author: CONDURACHE, Catalin (Science and Technology Facilities Council STFC (GB))

Presenter: CONDURACHE, Catalin (Science and Technology Facilities Council STFC (GB))

Session Classification: Workshop presentations

Contribution ID: **36**

Type: **not specified**

Workshop wrap-up

Friday, 7 September 2018 11:55 (20 minutes)

Primary author: MEINHARD, Helge (CERN)

Presenter: MEINHARD, Helge (CERN)

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials

Contribution ID: 37

Type: **not specified**

Welcome

Tuesday, 4 September 2018 14:15 (10 minutes)

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

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

Session Classification: Workshop presentations

Contribution ID: 38

Type: **not specified**

RAL Tier-1 strategy - Growing the UK community

Thursday, 6 September 2018 17:40 (20 minutes)

In 2013 the RAL Tier-1 switched its batch farm to using HTCondor. In the years following, several more UK sites have made the switch. The RAL Tier-1 batch farm is now well over 20000 job slots and HTCondor is a key service delivering our pledged resources to the WLCG, now and for the foreseeable future.

New funding opportunities are available to provide computing in the UK to the “long tail” of science. These are science experiments with only a handful of users but ever growing computing requirements. This talk will discuss how the RAL Tier-1 and other UK sites needs to evolve to meet these changing requirements.

Primary author: DEWHURST, Alastair (Science and Technology Facilities Council STFC (GB))

Presenter: DEWHURST, Alastair (Science and Technology Facilities Council STFC (GB))

Session Classification: Workshop presentations

Track Classification: HTCondor presentations and tutorials