

HTCondor Workshop Autumn 2024 in Amsterdam



Report of Contributions

Contribution ID: 1

Type: **not specified**

Final project update

Thursday 26 September 2024 09:50 (20 minutes)

This year has been eventful for our research lab, New hardware that brought along a host of challenges, we will share network, architecture and recent challenges that we are facing. It's all about scale.

Desired slot length

20

Speaker release

No

Author: HANDELMAN, David

Presenter: HANDELMAN, David

Session Classification: Workshop Session

Track Classification: HTCondor user presentations

Contribution ID: 2

Type: **not specified**

Integrating an IDE with HTCondor

Thursday 26 September 2024 10:10 (20 minutes)

Graphical code editors such as Visual Studio Code (VS Code) have gained a lot of momentum in the last years among young researchers. To ease their workflows, we have developed a VS Code entry point to harness the resources of an HTC cluster within their IDE.

This entry point allows users to have a “desktop-like” experience within VS Code when editing and testing their code while working in batch job environments. Furthermore, VS Code extensions such as Jupyter notebooks and Julia packages can directly leverage cluster resources.

In this talk we will explain the use case of this entry point, how we implemented it and show some of the struggles we encountered along the way. The developed solution can also scale out to federated HTCondor pools.

Desired slot length

15

Speaker release

Yes

Authors: HUBNER, Michael (University of Bonn (DE)); FREYERMUTH, Oliver (University of Bonn (DE))

Presenter: HUBNER, Michael (University of Bonn (DE))

Session Classification: Workshop Session

Track Classification: HTCondor user presentations

Contribution ID: 3

Type: **not specified**

Practical experience with an interactive-first approach to leverage HTC resources

Friday 27 September 2024 09:25 (20 minutes)

Development and execution of scientific code requires increasingly complex software stacks and specialized resources such as machines with huge system memory or GPUs. Such resources are present in HTC/HPC clusters and used for batch processing since decades, but users struggle with adapting their software stacks and their development workflows to those dedicated resources. Hence, it is crucial to enable interactive use with a low-threshold user experience, i.e. offering an SSH-like experience to enter development environments or start JupyterLab sessions from a web browser.

Turning some knobs, HTCondor unlocks these interactive use cases of HTC and HPC resources, leveraging the resource control functionality of a workload manager, wrapping execution within unprivileged containers and even enabling the use of federated resources crossing network boundaries without loss of security.

This talk presents the positive experience with an interactive-first approach, hiding the complexities of containers and different operating systems from the users, enabling them to use HTC resources in an SSH-like fashion and with their JupyterLab environments. It also provides a short outlook on scaling this approach to a federated infrastructure.

Desired slot length

20

Speaker release

Yes

Authors: HUBNER, Michael (University of Bonn (DE)); FREYERMUTH, Oliver (University of Bonn (DE))

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

Session Classification: Workshop Session

Track Classification: HTCondor user presentations

Contribution ID: 4

Type: **not specified**

Lenovo's Cooler approach to HTC Computing

Thursday 26 September 2024 15:10 (20 minutes)

Breakthroughs in computing systems have made it possible to tackle immense obstacles in simulation environments. As a result, our understanding of the world and universe is advancing at an exponential rate. Supercomputers are now used everywhere—from car and airplane design, oil field exploration, and financial risk assessment, to genome mapping and weather forecasting.

Lenovo's High-Performance Computing (HPC) technology offers substantial benefits for High Transaction Computing (HTC) by providing the necessary computational power and efficiency to handle large volumes of transactions. Lenovo's HPC solutions, built on advanced hardware such as the ThinkSystem and ThinkAgile series, deliver exceptional processing speeds and reliability. These systems are designed to optimize data throughput and minimize latency, which are critical factors in transaction-heavy environments like financial services, e-commerce, and telecommunications. The integration of Lenovo's HPC technology into HTC environments enhances the ability to process transactions in real-time, ensuring rapid and accurate data handling. This capability is crucial for maintaining competitive advantage and operational efficiency in industries where transaction speed and accuracy are paramount. Additionally, Lenovo's focus on energy-efficient computing ensures that these high-performance systems are also sustainable, aligning with broader environmental goals.

By leveraging Lenovo's HPC technology, organizations can achieve significant improvements in transaction processing capabilities, leading to better performance, scalability, and overall system resilience. According to TOP500.org, Lenovo is the world's #1 supercomputer provider, including some of the most sophisticated supercomputers ever built. With over a decade of liquid-cooling expertise and more than 40 patents, Lenovo leverages experience in large-scale supercomputing and AI to help organizations deploy high-performance AI at any scale.

Desired slot length

30

Speaker release

Yes

Author: Mr KOOPMAN, Rick**Presenter:** Mr KOOPMAN, Rick**Session Classification:** Workshop Session**Track Classification:** Miscellaneous

Contribution ID: 5

Type: **not specified**

Transitioning the CMS pools to ALMA9

Friday 27 September 2024 11:40 (20 minutes)

The Submission Infrastructure team of the CMS experiment at the LHC operates several HTCondor pools, comprising more than 500k CPU cores on average, for the experiment's different user groups. The jobs running in those pools include crucial experiment data reconstruction, physics simulation and user analysis. The computing centres providing the resources are distributed around the world and dynamically added to the pools on demand.

Uninterrupted operation of those pools is critical to avoid losing valuable physics data and ensure the completion of computing tasks for physics analyses. With the announcement of the end-of-life of CentOS 7, the CMS collaboration decided to transition their infrastructure, running essential services for the successful operation of the experiment, to ALMA 9.

In this contribution, we outline CMS's federated HTCondor pools and share our experiences of transitioning the infrastructure from CentOS 7 to ALMA 9, while keeping the system operational.

Desired slot length

Speaker release

Yes

Authors: PEREZ-CALERO YZQUIERDO, Antonio (Centro de Investigaciones Energéticas Medioambientales y Tecnológicas); KHAN, Farrukh Aftab (Fermi National Accelerator Lab. (US)); VON CUBE, Florian (KIT - Karlsruhe Institute of Technology (DE)); KIM, Hyunwoo (Fermi National Accelerator Lab. (US)); MASCHERONI, Marco (Univ. of California San Diego (US)); ZOKAITE, Vaiva (Vilnius University (LT))

Presenter: VON CUBE, Florian (KIT - Karlsruhe Institute of Technology (DE))

Session Classification: Workshop Session

Track Classification: HTCondor user presentations

Contribution ID: 6

Type: **not specified**

HTCondor setup @ ORNL, an ALICE T2 site

Friday 27 September 2024 09:45 (20 minutes)

ALICE experiment at CERN runs a distributed computing model and it is part of the Worldwide LHC Computing Grid (WLCG). WLCG uses a tiered distributed grid model. As part of the ALICE experiment's computing grid we run two Tier2 (T2) sites in the US, at Oak Ridge National Laboratory and Lawrence Berkeley National Laboratory. Computing resource usage and delivery are being accounted through OSG via GRATIA probes. This information is then forwarded to the WLCG. With the OSG software update and deprecation of some GRATIA probes we had to update the setup for the OSG accounting. To do so we have recently started to move our existing setup to HTCondor based workflow and new GRATIA accounting. I will present the setup for our T2 sites and HTCondor configuration escapade.

Desired slot length

Speaker release

Yes

Author: CHAKABERIA, Irakli (Lawrence Berkeley National Lab. (US))**Presenter:** CHAKABERIA, Irakli (Lawrence Berkeley National Lab. (US))**Session Classification:** Workshop Session**Track Classification:** HTCondor user presentations

Contribution ID: 7

Type: **not specified**

HPC use case through PIC

Friday 27 September 2024 11:00 (20 minutes)

In this contribution, I will present an HPC use case facilitated through gateways deployed at PIC. The selected HPC resource is the Barcelona Supercomputing Center, where we encountered some challenges, particularly in the CMS case, which required meticulous and complex work. We had to implement new developments in HTCondor, specifically enabling communication through a shared file system. This contribution will detail the setup process and the scale we were able to achieve so far.

Desired slot length

20

Speaker release

Yes

Author: FLIX MOLINA, Jose (CIEMAT - Centro de Investigaciones Energéticas Medioambientales y Tec. (ES))

Presenter: FLIX MOLINA, Jose (CIEMAT - Centro de Investigaciones Energéticas Medioambientales y Tec. (ES))

Session Classification: Workshop Session

Track Classification: HTCondor user presentations

Contribution ID: 8

Type: **not specified**

Abstracting Accelerators Away

Tuesday 24 September 2024 14:00 (20 minutes)

Currently more and more frameworks appear to perform offloaded compute to accelerators, or accelerating ML/AI workloads using CPU accelerators or GPUs. However right now the user it self still needs to figure out or decide how and what is the best execution library or acceleration system to execute there workloads.

How can we model this abstraction the best for htcondor so for our users the overhead to use the acceleration?

Desired slot length

15

Speaker release

Yes

Author: KOOISTRA, Emily**Presenter:** KOOISTRA, Emily**Session Classification:** Workshop Session**Track Classification:** HTCondor user presentations

Contribution ID: 9

Type: **not specified**

Practical considerations for GPU Jobs

Tuesday 24 September 2024 11:35 (30 minutes)

Practical considerations for GPU Jobs

Desired slot length

30

Speaker release

Yes

Author: OWEN, Andrew

Presenter: OWEN, Andrew

Session Classification: Workshop Session

Track Classification: HTCondor presentations and tutorials

Contribution ID: **10**

Type: **not specified**

Troubleshooting: What to do when things go wrong

Tuesday 24 September 2024 11:00 (30 minutes)

Troubleshooting: What to do when things go wrong

Desired slot length

30

Speaker release

Yes

Author: OWEN, Andrew

Presenter: OWEN, Andrew

Session Classification: Workshop Session

Track Classification: HTCondor presentations and tutorials

Contribution ID: 11

Type: **not specified**

On-boarding new users to HTCSS: getting off to a good start

On-boarding new users to HTCSS: getting off to a good start

Desired slot length

20 - 30

Speaker release

Yes

Author: OWEN, Andrew

Presenter: OWEN, Andrew

Contribution ID: 12

Type: **not specified**

Philosophy and Architecture: What the Manual Won't tell You

Tuesday 24 September 2024 09:30 (40 minutes)

Philosophy and Architecture: What the Manual Won't tell You

Desired slot length

40

Speaker release

Yes

Author: LIVNY, MIRON

Presenter: LIVNY, MIRON

Session Classification: Workshop Session

Track Classification: HTCondor presentations and tutorials

Contribution ID: 13

Type: **not specified**

HTCSS New Edition: No more job submits. Only Lists.

HTCSS New Edition: No more job submits. Only Lists.

Desired slot length

25

Speaker release

Yes

Author: LIVNY, MIRON

Presenter: LIVNY, MIRON

Contribution ID: 14

Type: **not specified**

Placement Tokens and the Role of the AP

Placement Tokens and the Role of the AP

Desired slot length

25

Speaker release

Yes

Author: LIVNY, MIRON

Presenter: LIVNY, MIRON

Contribution ID: 15

Type: **not specified**

WLCG Token Transition Update (incl the illustrious return of x509)

Friday 27 September 2024 09:00 (20 minutes)

WLCG Token Transition Update (incl the illustrious return of x509)

Desired slot length

20

Speaker release

Yes

Author: BOCKELMAN, Brian Paul (University of Wisconsin Madison (US))

Presenter: BOCKELMAN, Brian Paul (University of Wisconsin Madison (US))

Session Classification: Workshop Session

Track Classification: HTCondor presentations and tutorials

Contribution ID: 16

Type: **not specified**

HTCSS Versioning and Upgrades

HTCSS Versioning and Upgrades

Desired slot length

30

Speaker release

Yes

Author: BOLLIG, Cole

Presenter: BOLLIG, Cole

Contribution ID: 17

Type: **not specified**

Managing Storage at the EP

Wednesday 25 September 2024 09:35 (30 minutes)

Managing Storage at the EP

Desired slot length

30

Speaker release

Yes

Author: BOLLIG, Cole

Presenter: BOLLIG, Cole

Session Classification: Workshop Session

Track Classification: HTCondor presentations and tutorials

Contribution ID: **18**

Type: **not specified**

The new HTCSS Python API: Python Bindings Version 2

Thursday 26 September 2024 11:00 (20 minutes)

The new HTCSS Python API: Python Bindings Version 2

Desired slot length

20

Speaker release

Yes

Author: BOLLIG, Cole

Presenter: BOLLIG, Cole

Session Classification: Workshop Session

Track Classification: HTCondor presentations and tutorials

Contribution ID: 19

Type: **not specified**

Opportunities and Challenges Courtesy Linux Cgroups Version 2

Thursday 26 September 2024 14:00 (25 minutes)

Opportunities and Challenges Courtesy Linux Cgroups Version 2

Desired slot length

25

Speaker release

Yes

Author: BOCKELMAN, Brian Paul (University of Wisconsin Madison (US))

Presenter: BOCKELMAN, Brian Paul (University of Wisconsin Madison (US))

Session Classification: Workshop Session

Track Classification: HTCondor presentations and tutorials

Contribution ID: **20**

Type: **not specified**

GPUs - Whats Changing

GPUs - Whats Changing

Desired slot length

20

Speaker release

Yes

Author: TANNENBAUM, Todd

Presenter: TANNENBAUM, Todd

Session Classification: Registration

Contribution ID: 21

Type: **not specified**

HTCondor: Whats New / Whats coming up

Thursday 26 September 2024 11:25 (45 minutes)

HTCondor: Whats New / Whats coming up

Desired slot length

45

Speaker release

Yes

Author: TANNENBAUM, Todd

Presenter: TANNENBAUM, Todd

Session Classification: Workshop Session

Track Classification: HTCondor presentations and tutorials

Contribution ID: 22

Type: **not specified**

DAGman: I didn't know it could do that!

Thursday 26 September 2024 09:00 (45 minutes)

DAGman: I didn't know it could do that!

Desired slot length

45

Speaker release

Yes

Author: BOLLIG, Cole

Presenter: BOLLIG, Cole

Session Classification: Workshop Session

Track Classification: HTCondor presentations and tutorials

Contribution ID: 23

Type: **not specified**

Exploring Job Histories with ElasticSearch and HTCondor AdStash

Wednesday 25 September 2024 16:30 (20 minutes)

Exploring Job Histories with ElasticSearch and HTCondor AdStash

Desired slot length

20

Speaker release

Yes

Author: TANNENBAUM, Todd

Presenter: TANNENBAUM, Todd

Session Classification: Workshop Session

Track Classification: HTCondor presentations and tutorials

Contribution ID: 24

Type: **not specified**

Dealing with sources of Data: Choices and the Pros/Cons

Wednesday 25 September 2024 09:00 (30 minutes)

Dealing with sources of Data: Choices and the Pros/Cons

Desired slot length

30

Speaker release

Yes

Author: BOCKELMAN, Brian Paul (University of Wisconsin Madison (US))

Presenter: BOCKELMAN, Brian Paul (University of Wisconsin Madison (US))

Session Classification: Workshop Session

Track Classification: HTCondor presentations and tutorials

Contribution ID: 25

Type: **not specified**

HTCondor CE Dashboard: My CE is giving capacity to who?!?!??

HTCondor CE Dashboard: My CE is giving capacity to who?!?!??

Desired slot length

20

Speaker release

Yes

Author: TANNENBAUM, Todd

Presenter: TANNENBAUM, Todd

Contribution ID: 26

Type: **not specified**

CHTC Vision: Compute and Data Together

Wednesday 25 September 2024 11:25 (15 minutes)

CHTC Vision: Compute and Data Together

Desired slot length

15

Speaker release

Yes

Author: LIVNY, MIRON

Presenter: LIVNY, MIRON

Session Classification: Workshop Session

Track Classification: HTCondor presentations and tutorials

Contribution ID: 27

Type: **not specified**

Pelican Intro

Wednesday 25 September 2024 11:45 (20 minutes)

Pelican Intro

Desired slot length

30

Speaker release

Yes

Author: BOCKELMAN, Brian Paul (University of Wisconsin Madison (US))

Presenter: BOCKELMAN, Brian Paul (University of Wisconsin Madison (US))

Session Classification: Workshop Session

Track Classification: HTCondor presentations and tutorials

Contribution ID: 28

Type: **not specified**

PANEL and Discussion - Pelican and Condor: Flying Together, Birds of a Feather, Don't drop your data!

Wednesday 25 September 2024 12:05 (25 minutes)

PANEL and Discussion - Pelican and Condor: Flying Together, Birds of a Feather, Don't drop your data!

Desired slot length

30

Speaker release

Yes

Authors: BOCKELMAN, Brian Paul (University of Wisconsin Madison (US)); LIVNY, MIRON; TANNENBAUM, Todd

Presenters: BOCKELMAN, Brian Paul (University of Wisconsin Madison (US)); LIVNY, MIRON; TANNENBAUM, Todd

Session Classification: Workshop Session

Contribution ID: 29

Type: **not specified**

Dynamic resource integration with COBalD/TARDIS

Wednesday 25 September 2024 14:00 (20 minutes)

With the continuing growth of data volumes and computational demands, compute-intensive sciences rely on large-scale, diverse computing resources for running data processing, analysis tasks, and simulation workflows.

These computing resources are often made available to research groups by different resource providers resulting in a heterogeneous infrastructure.

To make efficient use of those resources, we are developing COBalD/TARDIS, a resource management system for dynamic and transparent integration.

COBalD/TARDIS provides an abstraction layer of resource pools and sites and takes care of scheduling and requesting those resources, independent of their sites local resource management systems. Through the use of adapters, COBalD/TARDIS is able to interface with a range of resource providers, including OpenStack, Kubernetes, and others, as well as support different overlay batch systems, with current implementations for HTCondor and SLURM.

In this contribution we present the general concepts of COBalD/TARDIS, several setups, with a focus on those using HTCondor, in different university groups, as well as WLCG sites.

Desired slot length

Speaker release

Yes

Authors: KUHN, Eileen (KIT - Karlsruhe Institute of Technology (DE)); VON CUBE, Florian (KIT - Karlsruhe Institute of Technology (DE)); QUAST, Gunter (KIT - Karlsruhe Institute of Technology (DE)); GIFFELS, Manuel (KIT - Karlsruhe Institute of Technology (DE)); SCHNEPF, Matthias Jochen; FISCHER, Max (Karlsruhe Institute of Technology)

Presenter: VON CUBE, Florian (KIT - Karlsruhe Institute of Technology (DE))

Session Classification: Workshop Session

Track Classification: HTCondor user presentations

Contribution ID: 30

Type: **not specified**

HTC from the user perspective

Wednesday 25 September 2024 16:00 (30 minutes)

HTC from the user perspective - to be chosen from former material

Desired slot length

Speaker release

Author: BOLLIG, Cole

Presenter: BOLLIG, Cole

Session Classification: Workshop Session

Track Classification: HTCondor presentations and tutorials

Contribution ID: 31

Type: **not specified**

HTCondor System Administration Introduction

Wednesday 25 September 2024 16:50 (30 minutes)

Quick overview of HTCondor for system administrators

Desired slot length

Speaker release

Author: TANNENBAUM, Todd

Presenter: TANNENBAUM, Todd

Session Classification: Workshop Session

Track Classification: HTCondor presentations and tutorials

Contribution ID: 32

Type: **not specified**

The new CMS Tier 1 in Serbia

Ben will hopefully contribute something

Speaker release

Desired slot length

Track Classification: HTCondor user presentations

Contribution ID: **33**

Type: **not specified**

Bens CERN talk

hence the name ...

Speaker release

Desired slot length

Track Classification: HTCondor user presentations

Contribution ID: 34

Type: **not specified**

Christophs introduction to Jupyter notebooks

might include a live-demo

Speaker release

Desired slot length

Author: BEYER, Christoph

Presenter: BEYER, Christoph

Contribution ID: 35

Type: **not specified**

Adapting Hough Analysis workflow to run on IGWN resources

Wednesday 25 September 2024 14:25 (20 minutes)

The computing workflow of the Virgo Rome Group for the CW search based on Hough Analysis has been performed for several years using storage and computing resources mainly provisioned by INFN-CNAF and strictly tied with its specific infrastructure. Starting with O4a, the workflow has been adapted to be more general and to integrate with computing centers in the IGWN community. We discuss our work toward this integration, the encountered problems, our solutions and the further steps ahead.

Desired slot length

10 min

Speaker release

Yes

Author: DAL PRA, Stefano (Universita e INFN, Bologna (IT))

Co-authors: PIERINI, Lorenzo (INFN Rome); Dr SILVESTRI, Lorenzo (INFN-CNAF); ASTONE, Pia; Dr BAGNASCO, Stefano (Istituto Nazionale di Fisica Nucleare, Torino)

Presenter: DAL PRA, Stefano (Universita e INFN, Bologna (IT))

Session Classification: Workshop Session

Track Classification: HTCondor user presentations

Contribution ID: 36

Type: **not specified**

Moving from Torque to HTCondor on the local cluster

Thursday 26 September 2024 12:10 (20 minutes)

Nikhef operates a local compute facility of around 6k cores. For the last two decades, Torque has been the batch system of choice on this cluster.

This year the system has been replaced with HTCondor; in this talk we share some of the concerns, design choices and experiences of the transition from the operator's perspective.

Desired slot length

Speaker release

Yes

Author: Mr VAN DOK, Dennis (Nikhef)

Presenter: Mr VAN DOK, Dennis (Nikhef)

Session Classification: Workshop Session

Track Classification: HTCondor user presentations

Contribution ID: 37

Type: **not specified**

GPUs in the Grid

Thursday 26 September 2024 14:50 (20 minutes)

In this presentation we will go over GPU deployment at the NL SARA-MATRIX Grid site. An overview of the setup is shown, followed by some rudimentary performance numbers. Finally, the user adoption and how the GPU is used is discussed.

Desired slot length

15

Speaker release

Yes

Author: Dr NAUTA, Lodewijk (SURF)**Presenter:** Dr NAUTA, Lodewijk (SURF)**Session Classification:** Workshop Session**Track Classification:** HTCondor user presentations

Contribution ID: 38

Type: **not specified**

Welcome, Introduction and Housekeeping

Tuesday 24 September 2024 09:00 (10 minutes)

Speaker release

Desired slot length

Presenters: BEYER, Christoph; HESTER, Mary

Session Classification: Workshop Session

Track Classification: Miscellaneous

Contribution ID: 39

Type: **not specified**

Nikhef Welcome

Tuesday 24 September 2024 09:10 (20 minutes)

Speaker release

Desired slot length

Session Classification: Workshop Session

Track Classification: Miscellaneous

Contribution ID: 40

Type: **not specified**

Round the room introductions

Tuesday 24 September 2024 10:10 (20 minutes)

Who are you, where are you from and what do you hope to get out of the workshop?

Speaker release

Desired slot length

Session Classification: Workshop Session

Track Classification: Miscellaneous

Contribution ID: 41

Type: **not specified**

Workshop Wrap-Up and Goodbye

Friday 27 September 2024 12:20 (10 minutes)

Presenter: BREW, Chris (Science and Technology Facilities Council STFC (GB))

Session Classification: Workshop Session

Track Classification: Miscellaneous

Contribution ID: 42

Type: **not specified**

Fun with Condor Print Formats

Wednesday 25 September 2024 15:10 (20 minutes)

During the 20 years history of the Torque batch system at Nikhef, we constructed several command line tools providing various overviews of what was going on in the system. An example: a tool that could tell us “what are the 20 most recently started jobs?”

```
mrstarts | tail -20
```

With HTCondor we wanted the same kind of overviews. Much of this can be accomplished using the HTCondor “print formats” associated with the `condor_q`, `condor_history`, and `condor_status` commands. In this talk I’ll present and discuss some examples, advantages and disadvantages of the approach, and along the way present some HTCondor mysteries we haven’t solved.

Desired slot length

Speaker release

Yes

Author: TEMPLON, Jeff (Nikhef National institute for subatomic physics (NL))

Presenter: TEMPLON, Jeff (Nikhef National institute for subatomic physics (NL))

Session Classification: Workshop Session

Track Classification: HTCondor user presentations

Contribution ID: 43

Type: **not specified**

Implementing OSDF Cache in SURF - MS4 Service

Friday 27 September 2024 10:10 (20 minutes)

In this presentation there will be a brief mention of the environment that hosts the OSDF Cache, the setup and suitable software for MS4 service. The presentation will lay out in a bit more depth the process of installing the OSDF cache and the challenges that arose during the installation.

Desired slot length

Speaker release

No

Author: COLO, Jasmin**Presenter:** COLO, Jasmin**Session Classification:** Workshop Session**Track Classification:** HTCondor user presentations

Contribution ID: 44

Type: **not specified**

AMD INSTINCT GPU CAPABILITY AND CAPACITY AT SCALE

Thursday 26 September 2024 14:25 (20 minutes)

The adoption of AMD Instinct™ GPU accelerators in several of the major high-performance computing sites is a reality today and we'd like to share the pathway that lead us here. We'll focus on characteristics of the hardware and ROCm software ecosystem, and how they were tuned to match the required compute density and programmability to make this adoption successful, from the discrete GPU to the supercomputer that tightly integrate massive amounts of these devices.

Desired slot length

Speaker release

Yes

Author: ANTAO, Samuel (AMD)

Presenter: ANTAO, Samuel (AMD)

Session Classification: Workshop Session

Track Classification: Miscellaneous

Contribution ID: 45

Type: **not specified**

Monte Carlo simulations of extensive air showers at NIKHEF

Tuesday 24 September 2024 14:45 (20 minutes)

This presentation will show how the Comic Rays group at Nikhef is using HTCondor in their analysis workflows on the local pool.

Desired slot length

10-15 min

Speaker release

Yes

Author: CHEMINANT, Kevin (Radboud University / NIKHEF)

Presenter: CHEMINANT, Kevin (Radboud University / NIKHEF)

Session Classification: Workshop Session

Track Classification: HTCondor user presentations

Contribution ID: 46

Type: **not specified**

NetApp DataOps Toolkit for data management

Wednesday 25 September 2024 10:10 (20 minutes)

The NetApp DataOps Toolkit is a python library that makes it easy for developers, data scientists and data engineers to perform various data management tasks. These tasks include provisioning new data volumes or developing workspace almost instantaneously. It improves flexibility in development's environment management. In this presentation, we will go over some examples and showcase how these libraries can be leveraged for different data management use cases.

Desired slot length

Speaker release

Author: GAVA, Didier (NetApp)

Presenter: GAVA, Didier (NetApp)

Session Classification: Workshop Session

Track Classification: Miscellaneous

Contribution ID: 47

Type: **not specified**

Storage Solutions with AI workloads

Wednesday 25 September 2024 11:00 (20 minutes)

Various AI workloads, such as Deep Learning, Machine Learning, Generative AI or Retrieval Augmented Generation, require capacity, compute power or data transfer performance. This presentation will show how simple a hardware / Software stack solution deployment, can leverage and/or become part of an AI infrastructure based on Ansible scripts. In addition, I will discuss two use cases, one on video surveillance and the second on real-time language processing, powered by an AI infrastructure setup.

Desired slot length

Speaker release

Author: GAVA, Didier

Presenter: GAVA, Didier

Session Classification: Workshop Session

Track Classification: Miscellaneous

Contribution ID: 48

Type: **not specified**

Kubenettes ↔ HTC

Wednesday 25 September 2024 14:45 (20 minutes)

Operating HTCondor with kubenettes

Desired slot length

Speaker release

Yes

Author: BOCKELMAN, Brian Paul (University of Wisconsin Madison (US))

Presenter: BOCKELMAN, Brian Paul (University of Wisconsin Madison (US))

Session Classification: Workshop Session

Track Classification: HTCondor presentations and tutorials

Contribution ID: 49

Type: **not specified**

HTCondor in Einstein Telescope

Friday 27 September 2024 11:20 (20 minutes)

The Einstein Telescope (ET) is currently in the early development phase for its computing infrastructure. At present, the only officially provided service is the distribution of data for Mock Data Challenges (using the Open Science Data Federation + CVMFS-for-data), with GitLab used for code management. While the data distribution infrastructure is expected to be managed by a Data Lake using Rucio, the specifics of the data processing infrastructure and tools remain undefined. This exploratory phase allows for a detailed evaluation of different solutions. Drawing from the experiences of 2nd-generation gravitational wave experiments LIGO and Virgo, which began with modest computational needs and expanded into distributed computing models using HTCondor, ET aims to build upon these foundations. LIGO and Virgo adopted, for their offline data analyses, the LHC grid computing model through a common computing infrastructure called IGWN (International Gravitational-Wave Observatory Network), incorporating systems like glideinWMS, which works on top of HTCondor, to handle high-throughput computing (HTC) tasks. Despite this, challenges such as the reliance on shared file systems have limited the migration to grid-based workflows, with only 20% of jobs currently running on the IGWN grid. For ET, the plan is to adapt and evolve from the IGWN grid computing model, making sure workflows are grid-compatible. This includes exploring Snakemake, a framework for reproducible data analysis, to complement HTCondor. Snakemake offers the ability to run jobs on diverse computing resources, including grid, Slurm clusters, and cloud-based infrastructures. This approach aims to ensure flexibility, scalability, and reproducibility in ET's data processing workflows, while overcoming past limitations.

Desired slot length

Speaker release

Yes

Author: TABASSO, Luca**Presenter:** TABASSO, Luca**Session Classification:** Workshop Session

Track Classification: HTCondor user presentations

Contribution ID: 50

Type: **not specified**

An ATLAS researcher's experience with HTCondor.

Tuesday 24 September 2024 14:25 (20 minutes)

A new users experience of switching to HTCondor

Desired slot length

Speaker release

Yes

Author: WOLFFS, Zef (Nikhef National institute for subatomic physics (NL))

Presenter: WOLFFS, Zef (Nikhef National institute for subatomic physics (NL))

Session Classification: Workshop Session

Track Classification: HTCondor user presentations

Contribution ID: 51

Type: **not specified**

Heterogeneous Tier2 Cluster and Power Efficiency Studies at ScotGrid Glasgow

Friday 27 September 2024 12:00 (20 minutes)

With the latest addition of 4k ARM cores, the ScotGrid Glasgow facility is a pioneering example of a heterogeneous WLCG Tier2 site. The new hardware has enabled large-scale testing by experiments and detailed investigations into ARM performance in a production environment.

I will present an overview of our computing cluster, which uses HTCondor as the batch system combined with ARC-CE as the front-end for job submission, authentication, and user mapping, with particular emphasis on the dual queue management. I will also touch on our monitoring and central logging system, built on Prometheus, Loki, and Grafana, and describe the custom scripts we use to extract job information from HTCondor and pass it to the node_exporter collector.

Moreover, I will highlight our research on power efficiency in HEP computing, showing the benchmarks and tools we use to measure and analyze power data. In particular, I will present a new figure-of-merit designed to characterize power usage during the execution of the HEP-Score benchmark, along with an updated performance-per-watt comparison extended to the latest x86 and ARM CPUs (Ampere Altra Q80 and M80, NVidia Grace, and recent AMD EPYC chips). Within this context, we introduce a Frequency Scan methodology to better characterize performance/watt trade-offs.

Desired slot length

15-20 minutes

Speaker release

Yes

Author: SIMILI, Emanuele**Co-authors:** BORBELY, Albert Gyorgy (University of Glasgow (GB)); BRITTON, David (University of Glasgow (GB)); STEWART, Gordon; SKIPSEY, Samuel Cadellin**Presenter:** SIMILI, Emanuele**Session Classification:** Workshop Session**Track Classification:** HTCondor user presentations

Contribution ID: 52

Type: **not specified**

HTCondor + Nikhef - A History of Productive Collaboration

Tuesday 24 September 2024 15:10 (20 minutes)

Author: LIVNY, MIRON

Presenter: LIVNY, MIRON

Session Classification: Workshop Session

Track Classification: HTCondor presentations and tutorials