

HTCondor Workshop in Europe 2023 - Highlights

19 – 22 September 2023 at IJCLab

<https://indico.cern.ch/e/htcondor2023>

Christoph Beyer, Thomas Hartmann, Helge Meinhard

- Ninth workshop in this series, which started with a pre-GDB in December 2014 (two workshops were held on-line only)
- Traditional time format: Tuesday morning to Friday lunchtime
- 41 attendees
 - 31 from this side of the Atlantic
 - Travel is still hard, many institutes under budget pressure, sending one person only
 - Considerably higher attendance than at 2022 edition in Cuneo, on a good way back to pre-Covid levels
 - Mix of new faces and old friends (good!)
- 20 organisations from nine countries
- 28 "scheduled" contributions by 16 presenters
- Interesting wide range of talks, high quality presentations, intense discussions

The organizing committee

***Chris Brew (co-chair), Christoph Beyer (co-chair), Helge Meinhard, Todd Tannenbaum, Mary Hester, Greg Thain, Michel Jouvin, Matthew West
Catalin Condurache, Antonio Puertas-Gallardo, Gabriele Fronzé, Josep Flix***
Local organiser: Michel Jouvin

The Location

IJCLab Orsay/France



The Group Photo :)

We picked the only rainy day for the pic



Spontaneous Parts of the Programme

“Show us your toolbox” & “ Lightning talks”

- **“Show us your toolbox” by now a fixed programme point and highlight**
 - Filled 1.5 hours easily with spontaneous wishes to give a short presentation, had even to turn down 2 more candidates
 - Good mix of entries reflecting tools and ideas of a lively community
 - Monitoring and energy saving amongst the contributions
 - Could be split into several more specific sessions, e.g. “show us your monitoring”, “show us your accounting” and a “general/other/miscellaneous” one
- **Lightning talks**
 - Meant as a chance for a short introduction of what people are using HTC for, like a micro-site report with one slide
 - Worked out well but was more like an additional toolbox session, need to explain purpose more clearly
 - Still a good idea for the first day in order to fuel coffee break discussions even further
- **Both spontaneous parts very lively with plenty of ideas and tools shown, well received by the audience**

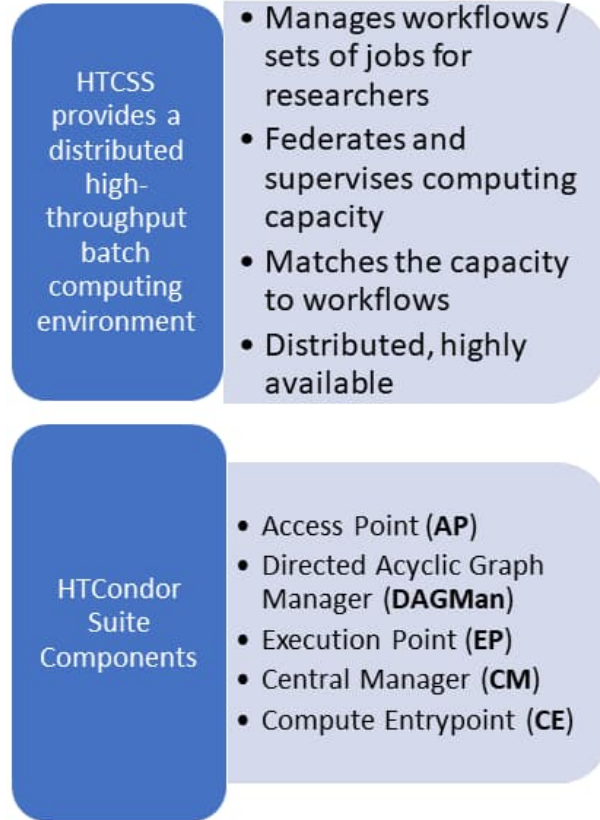
Tuesday Morning: Intro / Data and Security

09:00	Welcome Auditorium Joliot Curie, IJCLab	Christoph Beyer 09:00 - 09:05
	Welcome and housekeeping Auditorium Joliot Curie, IJCLab	Michel Jouvin  09:05 - 09:15
	IJCLab in a nutshell Auditorium Joliot Curie, IJCLab	Valerie CHAMBERT et al.  09:15 - 09:30
	HTCSS Vocabulary, Architecture, and User View Auditorium Joliot Curie, IJCLab	Todd Tannenbaum  09:30 - 10:20
10:00	Coffee Auditorium Joliot Curie, IJCLab	10:30 - 11:00
11:00	Data and Security intro Auditorium Joliot Curie, IJCLab	MIRON LIVNY 11:00 - 11:10
	Pelican: Data Federations and HTCSS Auditorium Joliot Curie, IJCLab	Brian Paul Bockelman  11:10 - 11:30
	IGWN Data Story Auditorium Joliot Curie, IJCLab	Peter Couvares et al.  11:35 - 12:05
12:00	Token Taxonomy Auditorium Joliot Curie, IJCLab	Brian Paul Bockelman  12:10 - 12:30

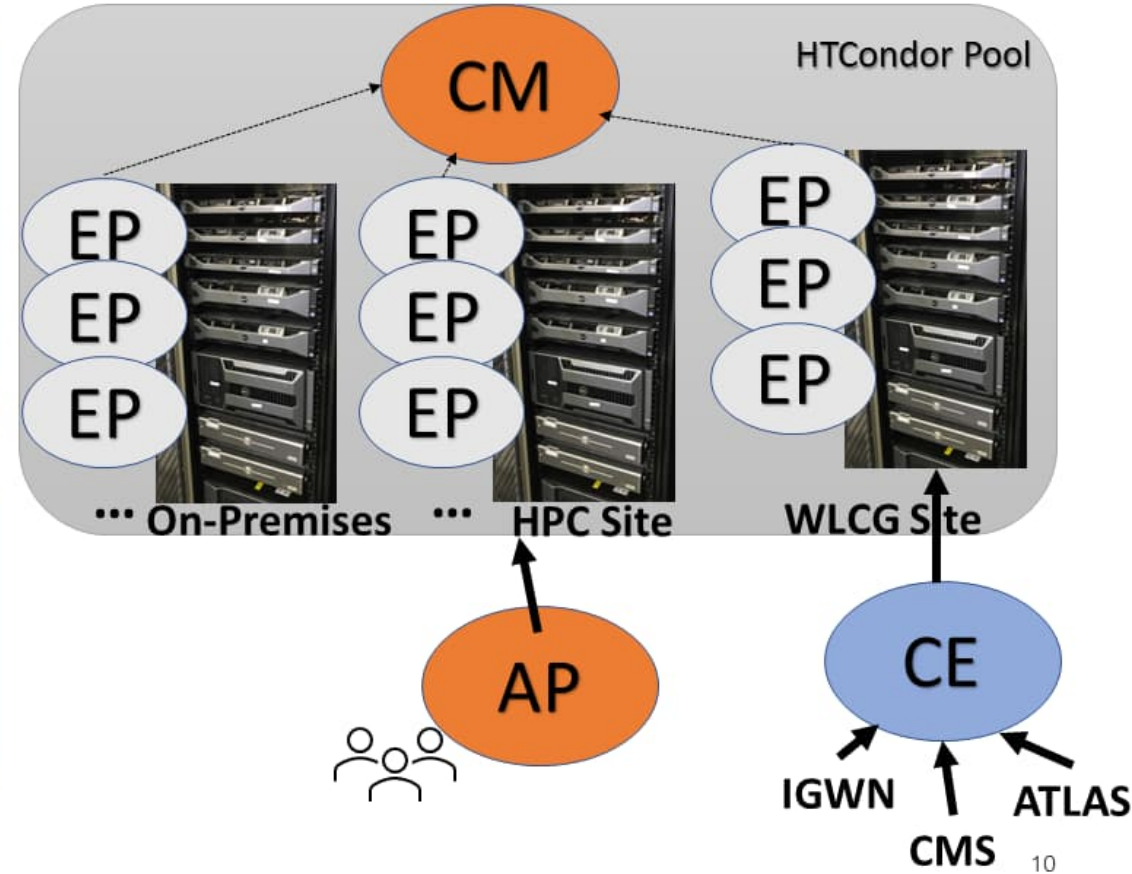
Tuesday Morning

HTCSS vocabulary, architecture and user view

- Audience traditionally very admin-focused at European workshops – this one was no exception in this sense
 - Vocabulary talk replaced ‘end-user-talk’
- Overview, introduction of architecture, vocabulary



10/15/2023



HTC = High Throughput Computing

HTCSS = HTCondor Software Suite

https://indico.cern.ch/event/1274213/contributions/5571128/attachments/2716862/4719087/intro_and_architecture.pdf

Tuesday Morning

Pelican: Advancing the Open Science Data Federation Platform



Dataset Sharing Powers Open Science

- Datasets are the lifeblood of science. Without them, the cleverest software or algorithms are essentially useless.
- To *utilize* these datasets, they need to be effectively connected to a broad set of computing types and capacity.
- Engineering this access is an immensely difficult problem:
 - Challenges in both software & infrastructure scale.
 - Inherently distributed problem: the archive and the computing may be separate resources.

The **Open Science Data Federation (OSDF)** federates dataset repositories with a reach of all of open science.

The Pelican Platform is a project recently funded by the NSF to help advance the OSDF. Scale: \$7M over 4 years starting 1 September.

Three main goals:

1. Strengthen and Advance the OSDF
2. Expand the types of computing where OSDF is impactful.
3. Expand the science user communities.
 - Particularly, Pelican is funded as part of the “National Discovery Cloud for Climate” (NDC-C) so there will be an early focus on climate datasets.



Over the last 12 months, the OSDF transferred

64PB &
62 files/s



Data used by

15 science collaborations &
~120 OSPool users



<https://indico.cern.ch/event/1274213/contributions/5571141/attachments/2716597/4719454/European-HTCondor-Workshop-2023-Pelican.pdf>

Tuesday Afternoon

14:00	(Sci)Tokens authentication in HTCondor-CE (incl new features) <i>Auditorium Joliot Curie, IJCLab</i>	James Frey 	14:00 - 14:30
	Getting Your Data with HTCondor <i>Auditorium Joliot Curie, IJCLab</i>	Brian Paul Bockelman 	14:35 - 15:05
15:00	Data and Security Wrap Up <i>Auditorium Joliot Curie, IJCLab</i>	MIRON LIVNY	15:10 - 15:25
	Coffee <i>Auditorium Joliot Curie, IJCLab</i>		15:30 - 16:00
16:00	Lightning Talks <i>Auditorium Joliot Curie, IJCLab</i>		16:00 - 17:30
17:00			

Tuesday Afternoon





(Sci)Tokens authentication in HTCondor-CE (incl. new features)

Why Move to Tokens?

- › HTC has used X.509 credentials for user authentication
 - Proposed by grid community with extensions
- › Tokens are more widely supported
 - Proxies never embraced by industry
- › Tokens allow a better security model
 - Authorization via capability instead of identity

<https://indico.cern.ch/event/1274213/contributions/5571156/attachments/2716926/4719205/SciTokensOnTheCE.pdf>

Wednesday Morning: Jobs and Access Points

09:00	Jobs and Access Points Introduction <i>Auditorium Joliot Curie, IJCLab</i>	MIRON LIVNY 09:00 - 09:10
	Talk of all the Tools <i>Auditorium Joliot Curie, IJCLab</i>	Gregory Thain  09:10 - 09:55
10:00	Visit of ThomX and SUPRATECH facilities <i>Auditorium Joliot Curie, IJCLab</i>	10:00 - 10:30
	Coffee <i>Auditorium Joliot Curie, IJCLab</i>	10:30 - 11:00
11:00	Linux tuning for large scale APs <i>Auditorium Joliot Curie, IJCLab</i>	Gregory Thain  11:00 - 11:30
	What Are My Jobs Doing? <i>Auditorium Joliot Curie, IJCLab</i>	James Frey  11:35 - 12:00
12:00	Whirlwind Tour: Job Placement Options <i>Auditorium Joliot Curie, IJCLab</i>	Gregory Thain  12:05 - 12:30

Wednesday Morning

Whirlwind tour of HTCondor tools

- Excellent entry point for administrators and power users alike

Every* HTCondor tool in 40
minutes

Greg Thain
Center for High Throughput Computing
University of Wisconsin - Madison



<https://indico.cern.ch/event/1274213/contributions/5573363/attachments/2715115/4721039/whirlwind-tour.pdf>

Wednesday Morning

What Are My Jobs Doing?

- Any user would benefit of Christina Koch's talk (given by Jaime Frey) about how to gather informations about what the job actually does or did
- Included a live demo analysing running and finished jobs

What Are My Jobs Doing?





HTCondor Week 2022

Christina Koch



<https://indico.cern.ch/event/1274213/contributions/5571145/attachments/2717809/4721048/HTCondorJobInfo-CKoch.pdf>

Wednesday Afternoon: User Presentations

14:00	HTCondor Services for Non-Humans <i>Auditorium Joliot Curie, IJCLab</i>	James Frey 	14:00 - 14:20
	Jobs and Access Points Wrap Up <i>Auditorium Joliot Curie, IJCLab</i>	MIRON LIVNY	14:25 - 14:45
15:00	News from the DESY Clusters <i>Auditorium Joliot Curie, IJCLab</i>	Thomas Hartmann 	14:50 - 15:10
	Deployment of HTCondor at GRIF <i>Auditorium Joliot Curie, IJCLab</i>	Dr Emmanouil Vamvakopoulos 	15:15 - 15:40
16:00	Coffee <i>Auditorium Joliot Curie, IJCLab</i>		15:40 - 16:10
	CERN Update <i>Auditorium Joliot Curie, IJCLab</i>	Ben Jones 	16:10 - 16:35
17:00	Office Hours <i>Auditorium Joliot Curie, IJCLab</i>		16:40 - 17:30

Wednesday Afternoon

User presentations

- Three interesting site report style talks (DESY, GRIF, CERN)
- Changes in LINUX ecosystem are an issue for everyone
- DESY
 - Analysing job-events in KAFKA
 - Scale-out of Jupyter notebooks with DASK
 - Adapting to energy challenges by load shedding
- GRIF
 - Early adaptor of HTC (2014)
 - Currently upgrading to HTC 10.X with pure token auth
 - ARC CE still in use at IRFU/CEA, HTC CE in the distributed common pool

<https://indico.cern.ch/event/1274213/contributions/5570407/attachments/2717928/4721764/DESYHH.pdf>

https://indico.cern.ch/event/1274213/contributions/5573382/attachments/2718282/4722008/htcondor@GRIF_2023-3.pdf

<https://indico.cern.ch/event/1274213/contributions/5570468/attachments/2718124/4722076/condoreu23lxbatch.pdf>

Wednesday Afternoon

- CERN
 - Energy savings over Christmas
 - Collector and Negotiator got 'happier' over time despite growth of cores and jobs
 - Scaling successful due to HTCSS devs, CMS pool leading the way
 - Now using bare metal (via Openstack Ironic) instead of VMs; Linux: skipping EL/* 8, using EL/* 9
 - Helped by move of many user communities to containers
 - Long tail of "local" user-submitted jobs is 80% of the effort
 - Possible transition scenario between interactive to batch use case:
 - Login to SWAN, use notebook on small dataset, when happy and ready to submit to batch...?

Thursday Morning

09:00	What's New, What's Coming Up <i>Auditorium Joliot Curie, IJCLab</i>	<i>Todd Tannenbaum</i>	09:00 - 09:40
	Energy Savings: Roosters and Condors, oh my! <i>Auditorium Joliot Curie, IJCLab</i>	<i>Todd Tannenbaum</i>	09:45 - 10:00
10:00	Visit of ThomX and SUPRATECH facilities <i>Auditorium Joliot Curie, IJCLab</i>		10:05 - 10:35
	Coffee <i>Auditorium Joliot Curie, IJCLab</i>		10:35 - 11:00
11:00			
	Reaching new scales in the CMS Global pool <i>Auditorium Joliot Curie, IJCLab</i>	<i>Marco Mascheroni</i>	11:25 - 11:45
12:00	The Virgo Data Quality Reports: an HTCondor automated framework to vet gravitational-wave candidates <i>Dr Nicolas Arnaud</i>		
	Energy savings by power modulation of a HTC pool <i>Auditorium Joliot Curie, IJCLab</i>	<i>Christoph Beyer</i>	12:10 - 12:30

Thursday Morning

What's new - what's coming up ?

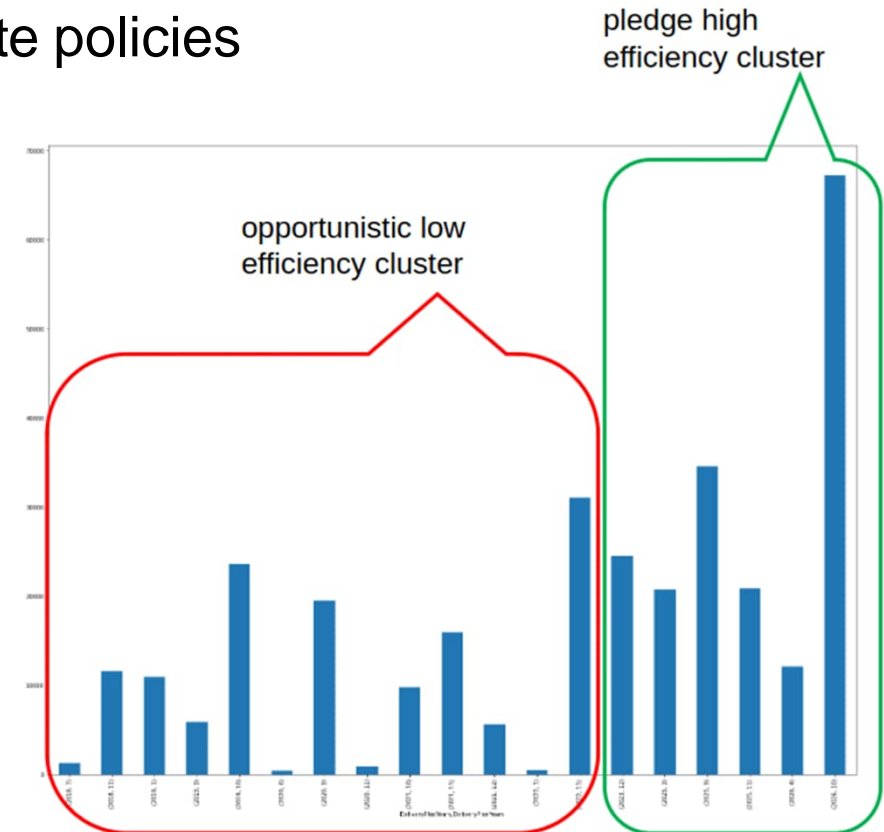
- New (as of 2022) version number schema
 - “Stable series” -> LTS release
 - Bug fixes, 10.0.0, 10.0.1, ...
 - “Developer series” -> Feature Releases
 - Bug fixes & new features, 10.1.0, 10.2.0, ...
 - New (as of 2023): Major version number = year (23.0.0, 23.1.2 etc.)
- New “standard” command line user interface being developed
 - “*htcondor <noun> <verb>*” e.g. “*htcondor job submit*”
 - Legacy tools (condor_q etc.) not going anywhere
- Event Log holds more infos e.g. about glidein site
 - “*htcondor eventlog read my_log_file*”
- Enterprise Linux 9 required work to support cgroups v2 and OpenSSL 3, now good for Intel/AMD (“x86_64”) and ARM (“aarch64”)

https://indico.cern.ch/event/1274213/contributions/5571127/attachments/2718736/4722885/Whats_New_european_workshop_2023%20-%20Copy.pdf

Thursday Morning

Two talks about power-management

- Condor-rooster daemon meant for hibernating EPs can be used to shut off EPs
 - When and how to shut off easy to implement in Hibernate policies
- DESY is preparing to modulate the power consumption following a given profile off cheap/sustainable energy
- Split pool in highly-efficient EPs that fulfill pledges and an opportunistic low-efficiency part
- Currently load shedding empty EPs using condor-rooster



https://indico.cern.ch/event/1274213/contributions/5571151/attachments/2718737/4722887/Power_Management_european_workshop_2023.pdf

https://indico.cern.ch/event/1274213/contributions/5570403/attachments/2718752/4722916/power_modulation_condor_pool.pdf

Thursday Afternoon

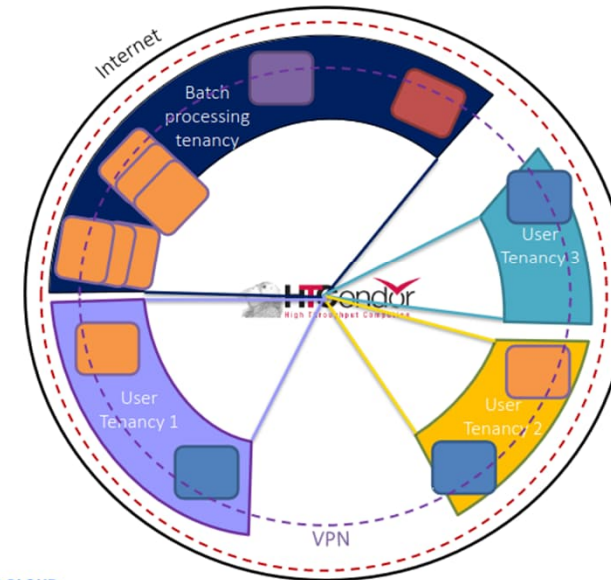

14:00	Multi-tenancy HTCondor for the European Weather Cloud <i>Auditorium Joliot Curie, IJCLab</i>	<i>Francesco Murdaca</i> 	13:45 - 14:05
	Just-in-time matching of workflows for the DUNE experiment <i>Auditorium Joliot Curie, IJCLab</i>	<i>Andrew McNab</i> 	14:10 - 14:30
	The journey to modern storage. <i>Auditorium Joliot Curie, IJCLab</i>	<i>David Handelman</i>	14:35 - 14:55
15:00	Thoughts on EDIA in dHTC Cyber-Infrastructure <i>Auditorium Joliot Curie, IJCLab</i>	<i>Matthew West</i> 	15:00 - 15:10
	Discussion on EDIA <i>Auditorium Joliot Curie, IJCLab</i>		15:10 - 15:30
	Coffee <i>Auditorium Joliot Curie, IJCLab</i>		15:30 - 16:00
16:00	Show your Toolbox		
17:00	<i>Auditorium Joliot Curie, IJCLab</i>		16:00 - 17:30

Thursday Afternoon

User presentations

- European Weather Cloud: Batch Processing (EWC)
- Aims to be the cloud-based collaboration platform for meteorological application development and operations in Europe and enables the digital transformation of the European Meteorological Infrastructure (EMI)
- Challenging use cases, large datasets, millions of relatively small jobs
- HTCondor already used internally at EUMETSAT for full mission dataset- reprocessing activities (Climate Team)
- Physically separated tenancies will be joined through an overall HTC pool (currently beta)

Batch Processing on EWC: multitenancy



Advantages:

- Use Idle machines
- Small tenancies can benefit from external large amount of resources
- User can submit job from their tenancy with their submit node
- Easy to deploy/maintain/scale in each tenancy
- Retain individual privacy



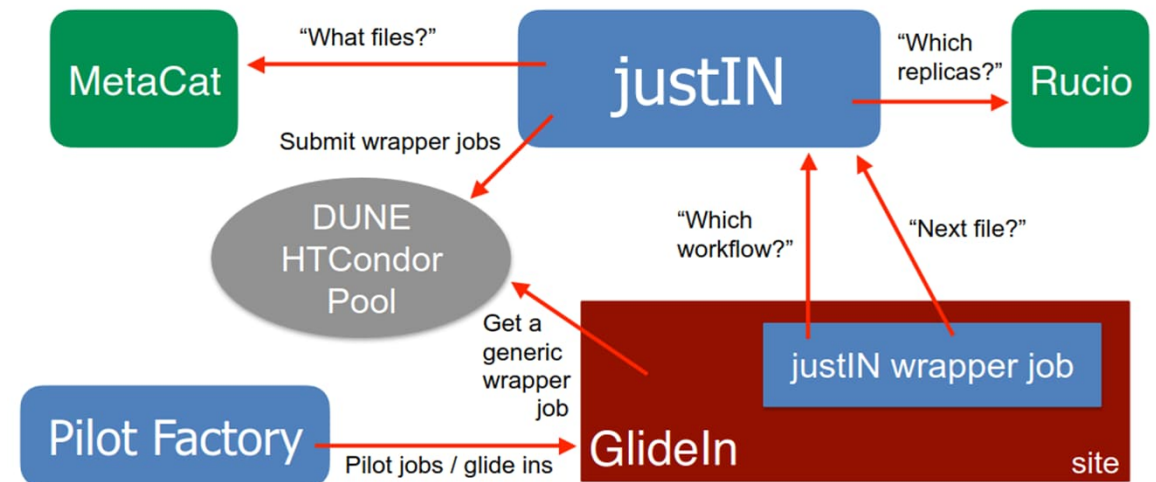
<https://indico.cern.ch/event/1274213/contributions/5570408/attachments/2719001/4723803/Batch%20Processing%20on%20EWC%20-%20HTCondor%20Europe%20Workshop%202023%20-%20Francesco%20Murdaca,%20Mike%20Grant.pdf>

Thursday Afternoon

User presentation: Just-in-time workflow matching for DUNE

- Due to DUNE's scale and specific requirements, in 2022 Fermilab set up a dedicated HTCondor Global Pool for DUNE to allow DUNE-specific configurations
- This uses HTCondor schedds at Fermilab and offsite: BNL and RAL-PPD so far; CC-IN2P3 and CERN are planned.
- To satisfy these Rucio/MetaCat/LArSoft/GlideInWMS constraints they designed a just-in-time workflow management system: justIN
- Central idea: once we get a job slot at a site, they choose and run a workflow that needs “nearby” unprocessed files, and feed it the details (URLs, Rucio DIDs, ...) of nearby unprocessed files to work on.

justIN with generic wrapper jobs



<https://indico.cern.ch/event/1274213/contributions/5570409/attachments/2719062/4723510/20230921-mcnab-dune-justin.pdf>

Friday Morning

09:00	Pools and Execution Points Intro <i>Auditorium Joliot Curie, IJCLab</i>	<i>MIRON LIVNY</i> 09:00 - 09:05
	Collector Overviews <i>Auditorium Joliot Curie, IJCLab</i>	<i>James Frey</i>  09:05 - 09:35
10:00	Microarchitecture <i>Auditorium Joliot Curie, IJCLab</i>	<i>Gregory Thain</i>  09:40 - 10:10
	What are glideins using my cluster for? Monitoring via the CE <i>Auditorium Joliot Curie, IJCLab</i>	<i>Todd Tannenbaum</i>  10:15 - 10:30
	Coffee <i>Auditorium Joliot Curie, IJCLab</i>	10:30 - 11:00
11:00	Advanced debugging with eBPF and Linux perf tools <i>Auditorium Joliot Curie, IJCLab</i>	<i>Gregory Thain</i>  11:00 - 11:50
12:00	Pools and Execution Points Wrap Up <i>Auditorium Joliot Curie, IJCLab</i>	<i>MIRON LIVNY</i> 11:50 - 12:10
	Wrap Up and Goodbye <i>Auditorium Joliot Curie, IJCLab</i>	<i>Chris Brew</i>  12:10 - 12:30

Friday Morning

Debugging HTCondor (and much more): introduction to perf, eBPF, ...

- perf: Powerful debugging tool
 - attach to kernel performance counters
 - get debugging info directly from the kernel
- eBPF: Trace the kernel yourself
 - build your own
 - monitoring, diagnostics, telemetry,...
 - performance, security,... analyses
 - replace your own kernel modules
- Introduction to **bpftool**
 - Awk-like syntax



After the Workshop is Before the Workshop

Tentative planning for 2024

- Feedback by attendees very good, workshop is considered useful
- Strong commitment by the HTC developers
- About 50% regular participants and 50% new faces
- Very lively community profiting from the rather small-scale workshop by interacting a lot
 - Proof of importance of in-person workshops
- NIKHEF Amsterdam volunteered to host the next meeting
- Tentative date: **calendar week 39 (23-Sep-2024 to 27-Sep-2024)**
 - Details to be fixed; presumably start Tuesday 24-Sep-2024



Thank You