

HTCondor EU WS 2023 - Highlights

Sep 18 – 22, 2023 IJCLab



Christoph Beyer, Thomas Hartmann, Helge Meinhardt

Workshop in numbers

Subheading, optional

- **41 Attendees** <https://indico.cern.ch/event/1274213>
 - 31 from this side of the Atlantic
 - Travel is still hard, many institutes only sending one person
 - Mix of new faces and old friends
- **20 Organizations from 9 Countries**
- **28 "Scheduled" Contributions by 16 Presenters**
- **Interesting wide range of talks, high quality presentations.**

The organizing committee

Chris Brew, Christoph Beyer, Helge Meinhard, Todd Tannenbaum, Mary Hester, Catalin Condurache, Antonio Puertas-Gallardo, Gabriele Fronz , Greg Thain, Josep Flix, Michel Jouvin, Matthew West

Local Organizer: Michel Jouvin

The Location

IJCLAB Orsay/France



The Group Picture :)

We picked the only rainy day for the pic



Interactive parts of the Program

“show us your toolbox” & “ Lightning talks”

- **Show us your toolbox by now a fixed program point & highlight**
 - Filled 1,5h easily without a single ‘prepared teaser’ by one of the usual suspects
 - 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 contribution
 - Could be splitted into several more specific sessions e.g. “show us your monitoring”, “show us your accounting” and a “general” 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
 - Still a good idea for the first day in order to fuel coffee break talks (not that those would need any fuel actually) ;)
 - Need to make purpose more clear
- **Both interactive parts very lively with plenty of ideas and tools shown, well received by the audience !**

Tuesday morning

Data & Security

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

Tuesday morning

HTCSS Vocabulary, Architecture, and User View

- Audience very admin focused for quite a while now
- Vocabulary talk replaces 'end-user-talk'
- provides overview and introduction of architecture

HTC = High Throughput Computing

HTCSS = HTCondor Software Suite


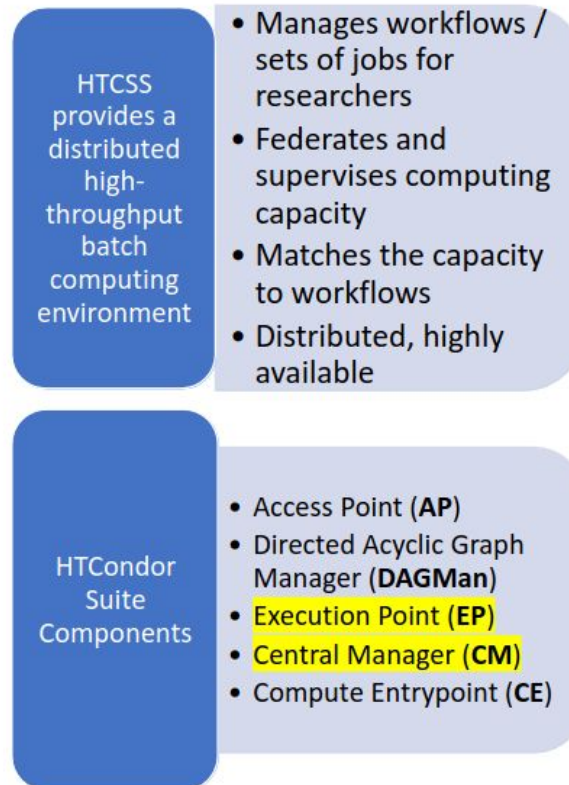
https://indico.cern.ch/event/1274213/contributions/5571128/attachments/2716862/4719087/intro_and_architecture.pdf

Tuesday morning

HTCSS Vocabulary, Architecture, and User View

- **SCHEDULER -> AccessPoint -> AP**
- **Condor-CE -> Compute Entrypoint -> CE**
- **STARTD -> Execute Point -> EP**
- **Condor Master -> Central Manager -> CM**

What is the HTCSS, what does it do?



9/19/2023

6

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.

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

Tuesday morning

Pelican: Advancing the Open Science Data Federation Platform



Over the last 12 months, the OSDF transferred

**64PB &
62 files/s**

Data used by

**15 science collaborations &
~120 OSPool users**



Tuesday morning

Pelican: Advancing the Open Science Data Federation Platform



Introducing: the Pelican Platform



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.

Tuesday afternoon

Subheading, optional

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		
17:00	<i>Auditorium Joliot Curie, IJCLab</i>		16:00 - 17:30

Tuesday afternoon

(Sci)Tokens authentication in HTCCondor-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>

Tuesday afternoon

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

- SciTokens & WLCG tokens accepted for job control
- Token discovery First location with a valid token is used
- Token can be in file or env (`$BEARER_TOKEN_FILE` & `use_scitokens = true`)
- Token claims in the job ad unalterable by the user
- Authorizing through mapfile & plugin if needed
- EGI CheckIn tokens supported out-of-the-box (plugin available)
- Simple testing scenario for SciTokens in the talk
- Quitting GSI & X.509 proxy support
 - X.509 proxy for user jobs supported in all versions

Wednesday morning

Jobs & 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's talk about how to gather informations about what the job actually does or did !
- Including a live demo analyzing 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 talks

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 talks / Site reports

- 3 interesting site report style talks (DESY, GRIF, CERN)
- Changes in the 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

- CERN
 - Energy savings over christmas
 - Collector and Negotiator got 'happier' over time despite growth of cores and jobs
 - Scaling easy due to HTC devs , CMS pool leading the way and "IRONIC" bare metal instead of VMs
 - Linux: skipping .*8 using .*9
 - 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...?

Wednesday afternoon

User talks CERN

Conditional containerisation of user jobs






JobImages:

- * e17 /cvmfs/unpacked.cern.ch/gitlab-registry.cern.ch/batch-team/containers/plusbatch/e17-full:latest
- * e18 /cvmfs/unpacked.cern.ch/gitlab-registry.cern.ch/batch-team/containers/plusbatch/e18-full:latest
- e19 /cvmfs/unpacked.cern.ch/gitlab-registry.cern.ch/batch-team/containers/plusbatch/e19-full:latest

JobOS:

- * e17 CentOS7
- * e18 CentOS8
- * e19 AlmaLinux9

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 ?

- “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 command line user interface
 - “*htcondor* <noun> <verb>” e.g. “*htcondor job submit*”
 - Legacy tools (condor_q etc.) not going anywhere

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

Thursday morning

What's new - what's coming up ?

New Job Status

\$ htcondor job status 123.45

Job 123.45 is currently running on host exec221.chtc.wisc.edu.

It started running again 2.1 hours ago.

It was submitted 3.6 hours ago.

Its current memory usage is 2.5 GB out of 4.0 GB requested.

Its current disk usage is 3.8 GB out of 5.5 GB requested.

It has restarted 2 times.

Goodput is 80% (0.5 hours badput, 2.1 hours goodput).

Thursday morning

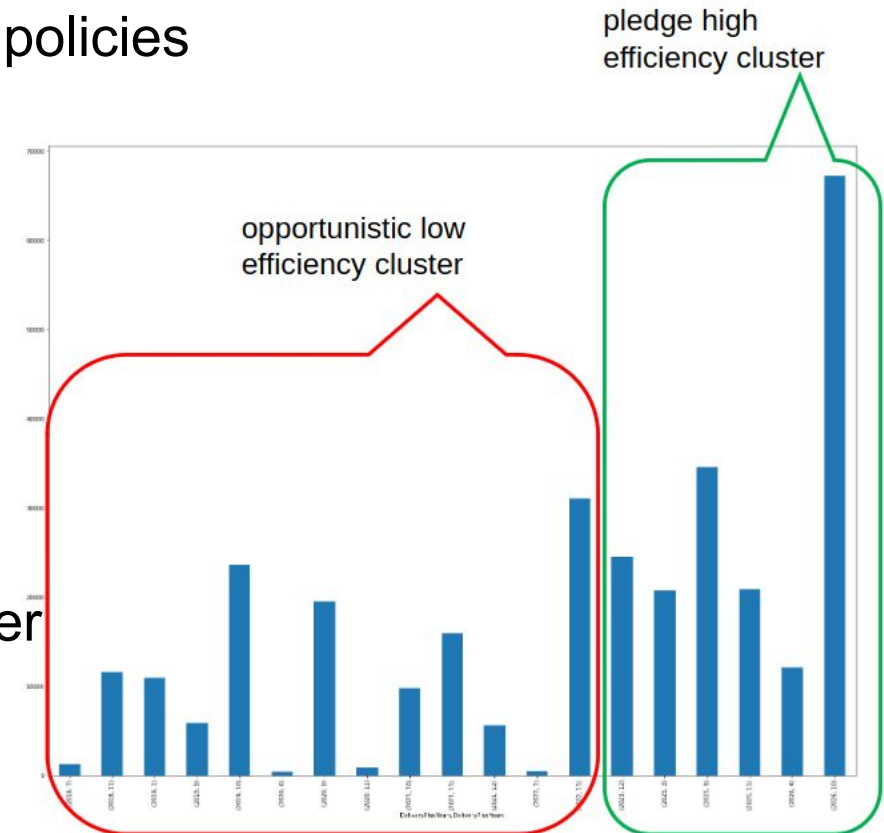
What's new - what's coming up ?

- Event Log holds more infos e.g. about glidein site
 - *"htcondor eventlog read my_log_file"*
- Can make "+CustomAttributes" first class with knobs
 - *+SiteJobType = "analysis" -> SiteJobType = "analysis"*
- Enterprise Linux 9 😊 😊 😊
 - Required work to support cgroups v2 and OpenSSL 3 Processors:
 - Intel/AMD ("x86_64")
 - ARM ("aarch64")

Thursday morning

Two talks about power-management

- Condor-rooster daemon meant for hibernating EPs can be used to shut off EPs
 - When/howto 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 high efficient EPS that fullfill pledges and 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 talks

- 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)





<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>

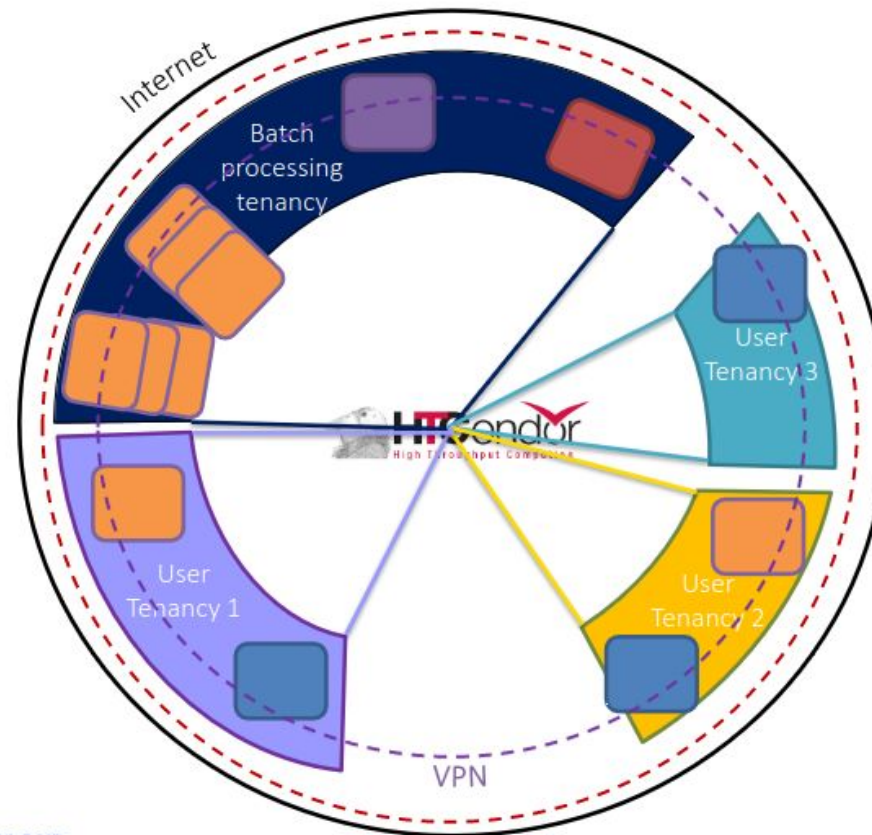
Thursday afternoon

user talks

Batch Processing on EWC: multitenancy



-  HTCondor manager
-  HTCondor submit node
-  HTCondor execute node
-  VPN



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




Thursday afternoon

user talks

- Just-in-time workflow matching for DUNE

What will DUNE be?

- Build huge tanks of liquid argon in a gold mine in South Dakota.
- Send neutrinos from Fermilab to South Dakota.
- Detect what neutrino flavours arrive in the liquid argon.
- Look out for supernova neutrinos and proton decays at the same time.



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

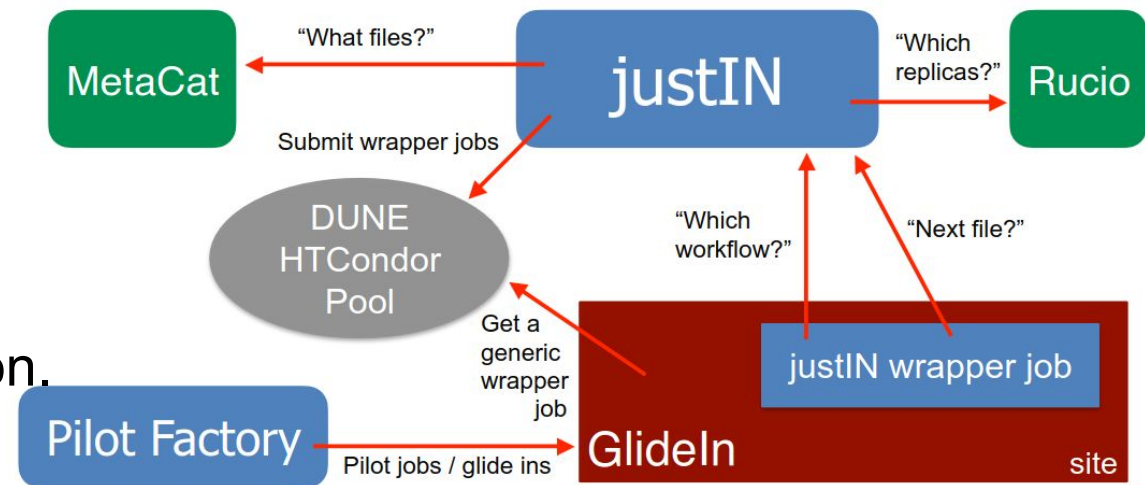
Thursday afternoon

user talks

- 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

justIN with generic wrapper jobs

- Central idea: once we get a job slot at a site, they choose and run a workflow that needs “nearby” unprocessed files, and we feed it the details (URLs, Rucio DIDs, ...) of nearby unprocessed files to work on.



Friday

Subheading, optional

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

Debugging Condor (and much more)

Introduction to perf, eBPF,... by Greg

<https://indico.cern.ch/event/1274213/contributions/5571148/>

- 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
- intro to **bpftool**
 - awk like syntax
- Check out Greg's user stories for inspiration



Getting ready for all Tokens

Debug and Test Token Access

- talks by Jamie and Brian on tokens in Condor
- debug (and impersonate) tokens
 - `condor_test_token`
 - *clone* (as admin) a token to trace the job flow
- enable X509 DN for direct submissions (no Globus)
 - how to submit w/o tokens for dedicated users/certs

<https://indico.cern.ch/event/1274213/contributions/5571156/>

<https://indico.cern.ch/event/1274213/contributions/5571155/>

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
- NIKHEF Amsterdam volunteered to host the next meeting
- Tentative next date: **calendar week 39 (23-09-2024 to 27-09-2024)**



Thank you