

(European) HTCondor workshop 2021: Questions, Comments, Answers

Monday 20 September 2021 - before coffee

Welcome and logistics (Helge Meinhard)

What's new in HTCondor? What is upcoming? (Todd Tannenbaum)

- (Matthew West, University of Exeter) Is there any intention to provide Puppet/Chef/Ansible config/setup files that would be akin to `get_htcondor`? Ala <https://github.com/HEP-Puppet/htcondor> but supported directly by the HTCSS team.
 - (Brian Lin, University of Wisconsin-Madison) the idea is that admins can run `get_htcondor` with the `--dry-run` option that prints out all the commands instead of running them and use the list of commands in their preferred configuration management tool.
- (Chris Brew - STFC RAL) I assume that there will be three sets of repos with the new version numbering scheme, stable, feature, and bleeding edge?
 - (Todd Tannenbaum, CHTC, Univ of Wisconsin-Madison) Yes, you assume correct! There is currently a separate repo for stable and new feature which will continue to hold software that has been beta tested in production for weeks before release. A new update repo will be added for the weekly patch releases.
 - (CB) Great, I was thinking of switching to the odd releases after 9, but not sure I'm ready for the update releases yet.
 - (Brian Lin, University of Wisconsin-Madison) Yes but we do retain the right to choose different names for the repos :)

What to measure and why (Miron Livny)

- (Matthew West, University of Exeter) Are there any equivalent projects to PATH/CC* in other nations to invest in *local* & *federated* dHTC resources? How could other agencies learn from and implement an equivalent scheme?
 - (Miron Livny, CHTC, Univ of Wisconsin-Madison) On this end, I will be happy to talk/support/guide any effort to "replicate" the PATH and CC* model in other places.

Monday 20 September 2021 - after coffee

New GPU architectures MIGs, multiple jobs per GPU, etc. (John Knoeller)

- (Gabriele Fronz , INFN) In Torino we are after an interface making available MIG, vGPU and MxGPU (AMD) partitioning under the exact same interface, to decouple vendor

specific code. Have you got any insights on why MIG doesn't allow to fully allocate the GPU cores and vRAM in most MIG configurations (either mixed profiles or not)?

- (Todd L Miller, University of Wisconsin) No, we don't know why MIG doesn't always fully allocate the card's resources.

Dealing with dynamic and mixed workloads (Stefano dal Pra)

- (Matthew West, University of Exeter) Is there anything to learn from how Slurm or any of the other schedulers do load balancing and resource allocation?

A new HTCondor monitoring for CNAF Tier-1 (Federico Versari)

Tuesday 21 September 2021 - before coffee

Auto-scaling in the cloud: Intelligent HTCondor resource management (Ross Thomson, Tom Downes)

- Request from speakers: reach out to drj@google.com or tpdowens@google.com if you are a Google Cloud customer; it helps us justify work to know about HTCondor usage

Introducing the HTCondor 9.0 series for users (Christina Koch)

- (Brian Bockelman, Morgridge) Do self-checkpointing jobs work well with draining? That is, will condor drain after checkpointing if the startd is draining?
 - (Todd L Miller, University of Wisconsin) That I recall, the self-checkpointing code is entirely unaware of draining. That should be easy to change and quite possibly worthwhile (to not restart the job after it finishes a checkpoint during draining).

Using SciTokens in HTCondor 9 (Brian Bockelman)

- (Thomas Hartmann, DESY) would it be possible for a user's job to bind-mount into a job a local path onto the issuer's namespace path? (according to some rules)
 - (Brian Lin, University of Wisconsin-Madison) Since the storage service would be the one authorizing the token, i think that would be fine. The path that appears to the user should be orthogonal to the path that the token is authorized to access
- (Brian Lin, University of Wisconsin-Madison) Note that WLCG token discovery in `condor_submit` isn't yet fully implemented: it respects the `scitokens_file` in the HTCondor submit file and the `BEARER_TOKEN_FILE` environment variable. I believe that the other discovery locations are expected in a future HTCondor release
 - (Dave Dykstra, FNAL) It does already know the default if `BEARER_TOKEN_FILE` is not set:
https://github.com/htcondor/htcondor/blob/master/src/condor_utils/condor_scitol
 - (Brian Lin, University of Wisconsin-Madison) Unfortunately, that piece of code is not used by `condor_submit`:
<https://opensciencegrid.atlassian.net/browse/HTCONDOR-634>
<https://opensciencegrid.atlassian.net/browse/HTCONDOR-634>

- (John Knoeller, University of Wisconsin-Madison) condor_submit needs a filename, and at the moment it will only look at the BEARER_TOKEN_FILE environment variable. There is a ticket to have it look in other WLCG standard file locations, but not BEARER_TOKEN. It is the responsibility of the submitter to insure that the token files remain valid as long as the job is in the queue; So it is probably not wise to use anything other than the submit keyword scitoken_file or the BEARER_TOKEN_FILE environment variable.
- (Dave Dykstra, FNAL) As I understand it, it's the responsibility of credd to keep the token files up to date for the length of the job, not the submitter. On the submit machine, I think that a token should only be needed at submit time and possibly later if the user wants to do something else such as condor_q when the schedd is on a different machine.

Tuesday 21 September 2021 - after coffee

Synthetic populations for personalized policy (Margherita Di Leo)

Operations in the HTCondor pool at CERN (Luis Fernandez Alvarez)

- (Thomas Hartmann, DESY) do you use for your containerization user namespaces? Do you map the UIDs:GIDs of the SE into the container namespaces or do you keep a flat namespace?

Running multiple experiment workflows on heterogeneous resources, the RAL experience (Alastair Dewhurst)

- (Matthew West, University of Exeter) So do the internal *clouds* run OpenStack enable HTCondor connect as backfill when they aren't used for specific tasks?

Wednesday 22 September 2021 - before coffee

Introducing HTCondor 9.0 for Admins (Greg Thain)

- (Matthew West, University of Exeter) I guess this is more a question about the future of the v9 series, will there be support additional architectures (ARM, Power10, etc) at some point?
 - If there is sufficient need, we will add support for new architectures. ARM seems to be in the lead at the moment. If people could mention their specific hardware and OS, that would help to make the decisions about what specific platforms to support -- e.g. not just ARM, but AArch64 running CentOS8, or whatever.
- (S Dal Pra) Is `classad_eval 'expression'` equivalent to `condor_status -compact -limit 1 -af 'expression'`?
 - (Todd L Miller, University of Wisconsin) Roughly. It's not clear from his example, but `classad_eval` allows you to construct a classed on the command-line (and of course, doesn't require a live collector).
https://htcondor.readthedocs.io/en/latest/man-pages/classad_eval.html
 - (S Dal Pra)

```
# classad_eval '3*2' 'True == False'
[ ]
6
false
```

- (Todd L Miller)
(The [] is an empty ClassAd)

```
$ classad_eval 'foo = 5' 'foo'
[ foo = 5 ]
5
```

- (Chris Brew - STFC RAL) Is it able to evaluate a ClassAd expression in the context of some Daemon? I often get an unevaluated expression back from condor_q or condor_status that it would be nice to be able to translate into it's final form. Ah, I see from the docs that you can extract a [machine|job] ClassAd and pass it to classad_eval
- (S Dal Pra) About condor_userprio --setceiling toddt 5: is there a way to read existing ceilings?
 - (TJ Knoeller, Univeristy of Wisconsin) The existing ceiling value is one of the columns of condor_userprio output.
- (David Cohen) Is there a removeceiling option?
 - (Greg Thain, University of Wisconsin) -1 means 'no ceiling'

Upgrading to HTCondor 9.0 (Todd Miller)

- (Name, affiliation) Question or comment

Wednesday 22 September 2021 - after coffee

HTCondor Integration with Hashicorp Vault for Oauth Credentials (Dave Dykstra)

- (Name, affiliation) Question or comment

Open Stage - Show us your toolbox

Note: this part will not be recorded

Oliver Freyermuth (U Bonn)

- Zabbix monitoring - usage, health checks, power, GPU/Disk/IO, file-system
- Foreman to config puppet to config HTCondor: marking nodes & users
- Extensive (interal) user documentation at Bonn with common useful commands

Luis Fernandez Alvarez, Ben Jones (CERN)

- website that allows one to manage (including ban) users by admins
- automation of some anomaly detection and notication methods (Prometheus)

K Scott Rowe (National Radio Astronomy Observatory)

- *summary*: wrote condor_qstat script to mirror Torque status messages
- (Brian Lin, University of Wisconsin-Madison) HTCondor-CE makes use of custom print formatting:
 - [Print format \(https://github.com/htcondor/htcondor-ce/blob/V5-branch/config/ce-status.cpf\)](https://github.com/htcondor/htcondor-ce/blob/V5-branch/config/ce-status.cpf)
 - [Relevant HTCondor configuration \(https://github.com/htcondor/htcondor-ce/blob/V5-branch/config/01-common-collector-defaults.conf#L12\)](https://github.com/htcondor/htcondor-ce/blob/V5-branch/config/01-common-collector-defaults.conf#L12)
- (Tim Theisen, University of Wisconsin-Madison) Print formatting is documented in [the HTCondor manual \(https://htcondor.readthedocs.io/en/latest/misc-concepts/print-formats.html\)](https://htcondor.readthedocs.io/en/latest/misc-concepts/print-formats.html)
- (Greg Thain, University of Wisconsin-Madison) FYI, to map a job status int to a string, you could use a classad list with an array reference, like {"I", "R", "H", "C", "X"} [JobStatus]
- (Jason Patton, University of Wisconsin-Madison) The Python bindings are also great for job and daemon query scripts, e.g. condor_watch_q

Christoph Beyer (DESY)

- icinga for alarming and checking service problems as well as node restarts
- batchnode script, reset & rebooting nodes

Chris Brew (UKRI-RAL)

- static webpage for condor_q & condor_status, long list of nodes
- startd cron jobs are wonderful (health checks for lots of things)
- schedule access to interactive GPU nodes & wrapper for condor_status for GPU status
- (Ben Jones, CERN) Maybe worth checking [CVMFS collectd plugin \(https://github.com/cvmfs/collectd-cvmfs\)](https://github.com/cvmfs/collectd-cvmfs) which can do per-repo cvmfs health reporting

Jason Patton (CHTC UW-Madison)

- Monitoring with Elasticsearch
- Setting up condor_adstash: <https://htcondor.readthedocs.io/en/latest/admin-manual/monitoring.html#elasticsearch>
- Email reports (highly site-specific!!!): <https://github.com/CHTC/JobAccounting/tree/python>

Thomas Hartmann (DESY)

Stefano dal Pra (CNAF)

- Configuration snippet to define holdreason in an alternative way
When a job is put on hold for one of several reason, the standard way to specify SYSTEM_PERIODIC_HOLD_REASON makes use of several nested IfThenElse(...) clauses, making the configuration difficult to read and maintain. We use an alternative way, as in the following snippet:

```

SecondStart = (NumJobStarts == 1 && JobStatus == 1)
TooMuchDisk  = (DiskUsage_raw > 35 * (CpusProvisioned ?: RequestCpus) *
1024000)
TooMuchRSS   = (ResidentSetSize_RAW > 40 * (CpusProvisioned ?: RequestCpus) * 1e6
)
TooMuchTime  = (jobstatus == 2 && (time() - JobStartDate > 86400 * 7))
#TooMuchMemory = (MemoryProvisioned > 6000)
#TooMuchImg  = (ImageSize_RAW/1e6 > 35 * CpusProvisioned)

SYSTEM_PERIODIC_HOLD = $(SYSTEM_PERIODIC_HOLD:False) || $(SecondStart) || \
$(TooMuchDisk) || \
$(TooMuchRSS) || \
$(TooMuchTime)
SYSTEM_PERIODIC_HOLD_REASON = strcat(Owner, {"", \
", Second start not allowed", \
", TooMuchDisk: 35GB/core", \
", TooMuchRSS: 20GB/core", \
", Job runtime > 20 days"}[max({\
int($(SecondStart)), \
int($(TooMuchDisk))*2, \
int($(TooMuchRSS))*3, \
int($(TooMuchTime))*4})])

```

The idea follows these steps:

- name each hold condition by its boolean expression: SecondStart, TooMuchDisk etc.
- define SYSTEM_PERIODIC_HOLD by or-ing together the conditions
- put your holdreasons in an array of strings (first one being empty string)
- use int(<condition>) to convert False/True to 0/1
- put these into an array, multiplied by their position (start from 1)
- take the max() from the array
- use that value as the index for the array of strings
- As a bonus, prepend the job Owner to the final holdreason. The result, as it appears in the SchedLog, looks as follows:

```

09/22/21 15:22:29 (pid:2029518) Job 3594460.0 put on hold: pillhcb031,
TooMuchRSS: 20GB/core
09/23/21 08:55:09 (pid:2029518) Job 3606258.0 put on hold: pillhcb031, Second
start not allowed
09/23/21 09:09:09 (pid:2029518) Job 3577762.0 put on hold: pilatlas030,
TooMuchDisk: 35GB/core

```

Thursday 23 September 2021 - before coffee

The CMS Submission Infrastructure deployment (Antonio Perez-Calero Yzquierdo)

- (Greg Thain, University of Wisconsin) The CMS Global pool is running HTCondor 9.1.2 today -- do you intend to stay with the "development" series of HTCondor, and can you talk about your experiences running the development version vs. the stable series?

Operations and Monitoring of the CMS HTCondor pools (Saqib Haleem)

- (Jason Patton, University of Wisconsin-Madison) When putting daemon ClassAd data into Elasticsearch, do you do anything to normalize the attribute names since Elasticsearch field names are case sensitive (and ClassAd attr names are insensitive)? This has been a problem when dealing with job ClassAd data, where different users/access points might use different casing (e.g. +WantGPULab vs. +WantGpuLab) to achieve the same result.

Self-Checkpointing Jobs in HTCondor (Christina Koch)

- (Dave Dykstra, Fermilab) FYI Checkpointing is planned to be [added to singularity](https://github.com/hpcng/singularity/issues/468) (<https://github.com/hpcng/singularity/issues/468>) in release 3.9 thanks to a contribution from [Lancium](https://lancium.com) (<https://lancium.com>) based on [DMTCP](https://github.com/dmtcp/dmtcp) (<https://github.com/dmtcp/dmtcp>). This should help a lot of people take advantage of this condor feature. The current implementation from Lancium does need some considerable updating for the latest singularity so this is still tentative.
 - (Matthew West) This is checkpointing the entire environment?
 - (Dave Dykstra) I believe it is checkpointing the entire container state.
 - (Thomas Hartmann) can I migrate frozen processes/cgroups between machines/kernels?
- (Matthew West, University of Exeter) When the checkpointed job restarts, does it resume on the same worker node?
 - (Todd Miller, University of Wisconsin-Madison) When a self-checkpointing job exits with the checkpoint code, it will be restarted on the same worker node (in the same slot). Once the job is interrupted, it will be restarted according HTCondor's usual scheduling.
- (Pepe Flix, PIC): What happens if you checkpoint a high mem job? The checkpointing file is big, one can run into troubles? i.e. filling the SPOOL? (it affects all of the users, your job, or all of your jobs?)
 - (Todd Miller, University of Wisconsin-Madison) The memory usage of a job isn't necessarily related to the size of its footprint, but if you fill up SPOOL, that could affect all users, depending on the configuration of the schedd. (You can configure per-user SPOOL directories, which could then be size-controlled using the usual quota mechanisms; see below.)
- (Matthew West, University of Exeter) The first result on [google](https://www.google.com/search?q=self-checkpointing+code&ei=9I1MYfTKCKGDhbIPmdOV0Ag&oq=self-checkpointing+code) (<https://www.google.com/search?q=self-checkpointing+code&ei=9I1MYfTKCKGDhbIPmdOV0Ag&oq=self-checkpointing+code>) for self-checkpointing code brings up [TheDocs](https://htcondor.readthedocs.io/en/latest/users-manual/self-checkpointing-applications.html) (<https://htcondor.readthedocs.io/en/latest/users-manual/self-checkpointing-applications.html>). Do you have any recommendations where researchers can find a good tutorial on introducing self-checkpoints in their own code?
 - (Christina Koch, CHTC/UW-Madison), I don't know of any. I tried to outline some of the key components on the slide about the Fibonacci code example: stopping condition, input logic and output logic. This might be in a CHTC guide sometime in the future.

- (From Zoom chat) can I as an admin limit the spool usage per user?
 - (Lauren Michael, CHTC/UW-Madison) yes, if you map each user's spool location to count against their quota (perhaps their /home quota), which is what we do now in CHTC.

Thursday 23 September 2021 - after coffee

Q / A / discussion: HTCondor philosophy and architecture (Greg Thain)

Q / A / Discussion: HTCondor Python Bindings (Jason Patton)

Q / A / Discussion: Negotiator policy and configuration (Greg Thain)

- (Will, BNL) Our slotweight is `ifThenElse(Cpus >= Memory/2048, Cpus, Memory/2048)`

Q / A / Discussion: Using IDTokens for authentication in HTCondor 9 (Todd Tannenbaum)

- (Matthew West, Uni of Exeter) *very poor summary of answer* - Buzzwords are useful when talking to high-level administrators. Zero-Trust making doing secure authentication and communication easier and by default.
- (Tom Downes, Google) asked about ways to bootstrap trust (e.g., in the cloud when there are highly transient machines) when using IDTOKENS. Can IDTOKENS requests be used here? He has an existing system that could determine if a request is reasonable/valid - is there a way to have condor make a "callout"?
 - (Brian Bockelman, Morgridge) One approach would be to have an entity that periodically queries for outstanding requests and programmatically approves them. We do something like this for a web application in OSG: <https://github.com/opensciencegrid/open-science-pool-registry>.
- (Stefano Dal Pra, INFN) How can I use IDTOKENS to enable remote submission (instead of REMOTE_FS)? Todd: Can, for example, generate an IDTOKEN on first login for each user based on PAM.
- (Stefano Dal Pra, INFN) How can I debug IDTOKENS? Todd: `condor_ping` is your friend to test out the security configuration! (Short demo ensued)
- (Chris Brew, RAL) Can IDTOKENS replace MUNGE for remote submit? (Not really)

Friday 24 September 2021 - before coffee

- (Chris Brew - RAL) - @ToddT, Following up from my question yesterday on the IDTokens, I think I'm looking for a service (or setuid binary) on the submit host, that I as an admin can authorise, that can then mediate getting a token for the user at submit time. I can stick with munge of course but a secure by default, nearly zero config for remote submission would be cool.

What's new / upcoming for the HTCondor CE? (Brian Lin)

- (Matthew West, University of Exeter) How does the purchased Globus software suite mesh with tokens? Some folks feel fine sticking with their payed services.

- (Brian Lin, University of Wisconsin-Madison) As I understand it, paid Globus endpoints are focused on user data transfer and is separate from the open source Grid Community Toolkit (née Globus Toolkit). So the move to bearer tokens (SciTokens/WLCG tokens) and IDTOKENS shouldn't have any effect on a site's paid Globus services.
- (Mary Hester, Nikhef)

There is a rather large gap in transitions to Scitokens vs WLCG tokens (I think they estimate a discontinuation of x509 certs is 2024 for WLCG?), are you worried that this could negatively impact the non-hep Vos that may not have planned for this change?

 - (Brian Lin, University of Wisconsin-Madison) one thing to tease out of the WLCG's May 2024 X.509 timeline is that this accounts for both compute and storage access and the latter has a longer timeline for X.509 support. For example, HTCondor 9.3.0 will still perform X.509 proxy delegation even without GCT libraries, meaning that pilot authentication and authorization should be performed with bearer tokens but the pilot job can still carry an X.509 proxy that is used for storage access. That being said, yes, I am concerned about the non-HEP VOs! We'll definitely be keeping an eye on how their transitions are going as we approach all of our milestones.
 - (Mary Hester, Nikhef) Thanks Brian! One reason that I ask is that we are a site that accepts jobs from VOs like ligo and dune, but we are an EGI site so the majority of the submissions are coming from Europe, which means we will need to support x509s until 2024. I think there may be a mismatch here and I'm not sure how we can make this a smooth transition at our site. (Also we're using ARC-CEs and I haven't seen anything about a migration to a version of ARC that will support WLCG tokens/Scitokens yet... maybe I've just missed this?) Basically I'm worried that with the timeline given, we will not have the infrastructure in place to support accepting Scitokens because of the timeline mismatch.
 - (Tim Cartwright, University of Wisconsin-Madison) Also, please note that token transition timelines are a major topic of the OSG Token Transition Workshop in mid-October: <https://opensciencegrid.org/events/Token-Transition-Workshop/>

Job router transforms (John Knoeller)

How to write a custom file xfer plugin (Mark Coatsworth)

- (Todd L Miller, University of Wisconsin-Madison) HTCondor *does* support S3, although it requires an https plug-in (available by default).
- (Matthew West, University of Exeter) Some sysadmin are likely to be supremely skeptical of the suggestion of letting users pull data remotely on their own. How would you allay their fears about this data transfer capability?
 - (Todd L Miller, University of Wisconsin-Madison) The plug-ins are run as the same user as the job, so there's no additional risk.
 - (Matthew West) The HPC admins I talk to regularly do not trust users in the slightest and given previous credential breaches are skeptical of any outside communication to their systems. I love these features and I expect other uses might as well, but it always feels like a hard sell to existing staff.
- (K. Scott Rowe, NRAO) Any good debugging techniques for plugins?
- (Ben Jones, CERN) I guess all the file xfer happens on the starter, so the Worker nodes

would need to be able to access anything directly (ie http pull etc)? Nothing happens via the schedd?

- (John Knoeller, HTCSS) When the job supplies the plugin, the plugin will be transferred from the schedd.

Friday 24 September 2021 - after coffee

Campus research and facilitation (Lauren Michael)

- (Matthew West, Exeter) What are the first objectives if one wishes to implement an OSG like system in one's own country?
 - Need local research champions. For tips on finding researchers with yet-to-be-realized HTC demand: <https://indico.fnal.gov/event/46925/>
 - Importance to connect with like-minded individuals at other institutions
 - Path & CHTC are willing to help formalize a strategy, if you're interested.

In silico detection of (CRISPR) spacers matching Betacoronaviridae genomes in gut metagenomics sequencing data (Gabriele Leoni)

Workshop wrap-up (Helge Meinhard)