


# HTCondor-CE: What's new and upcoming?



Brian Lin

Center for High Throughput Computing


# Resource Allocation Requests (RARs)




Access Point



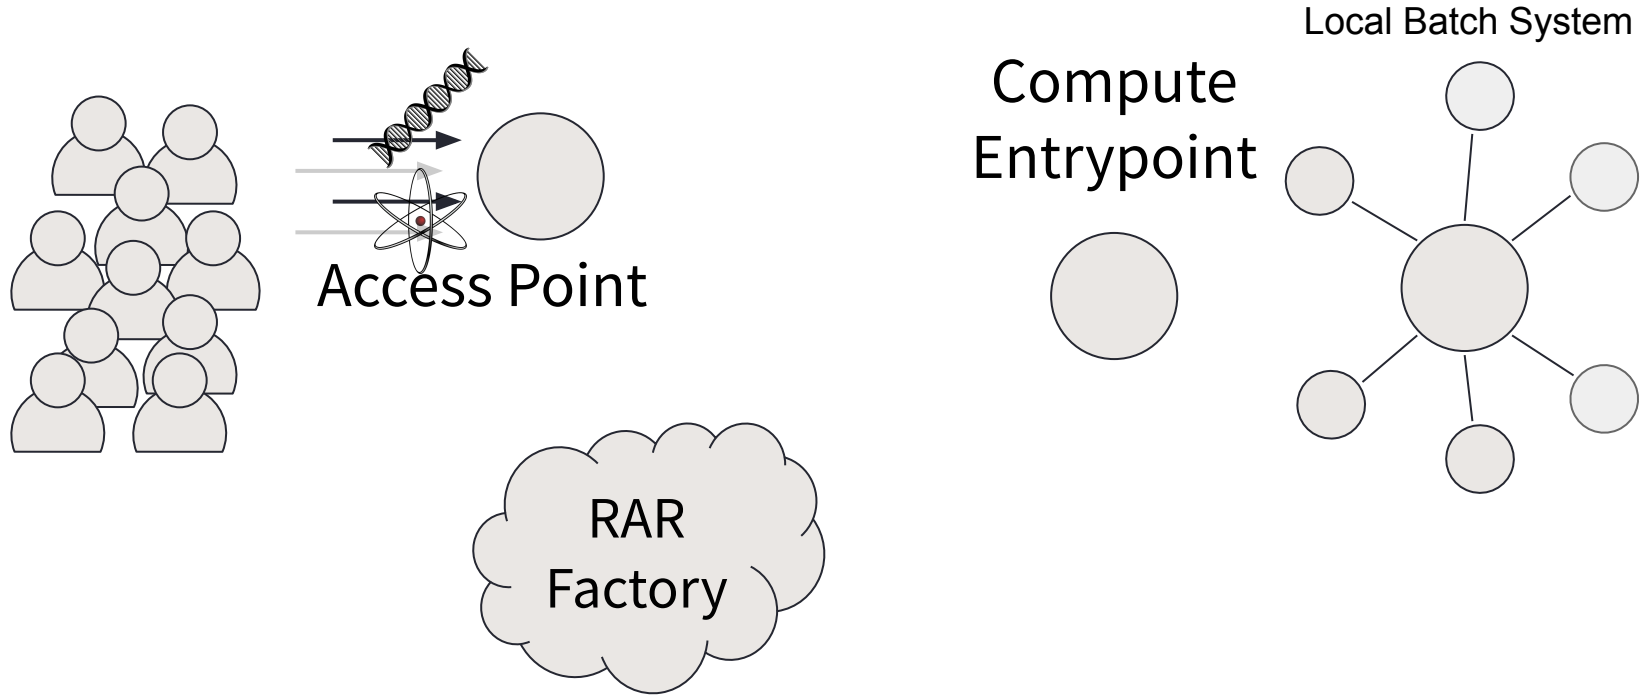
Compute  
Entrypoint



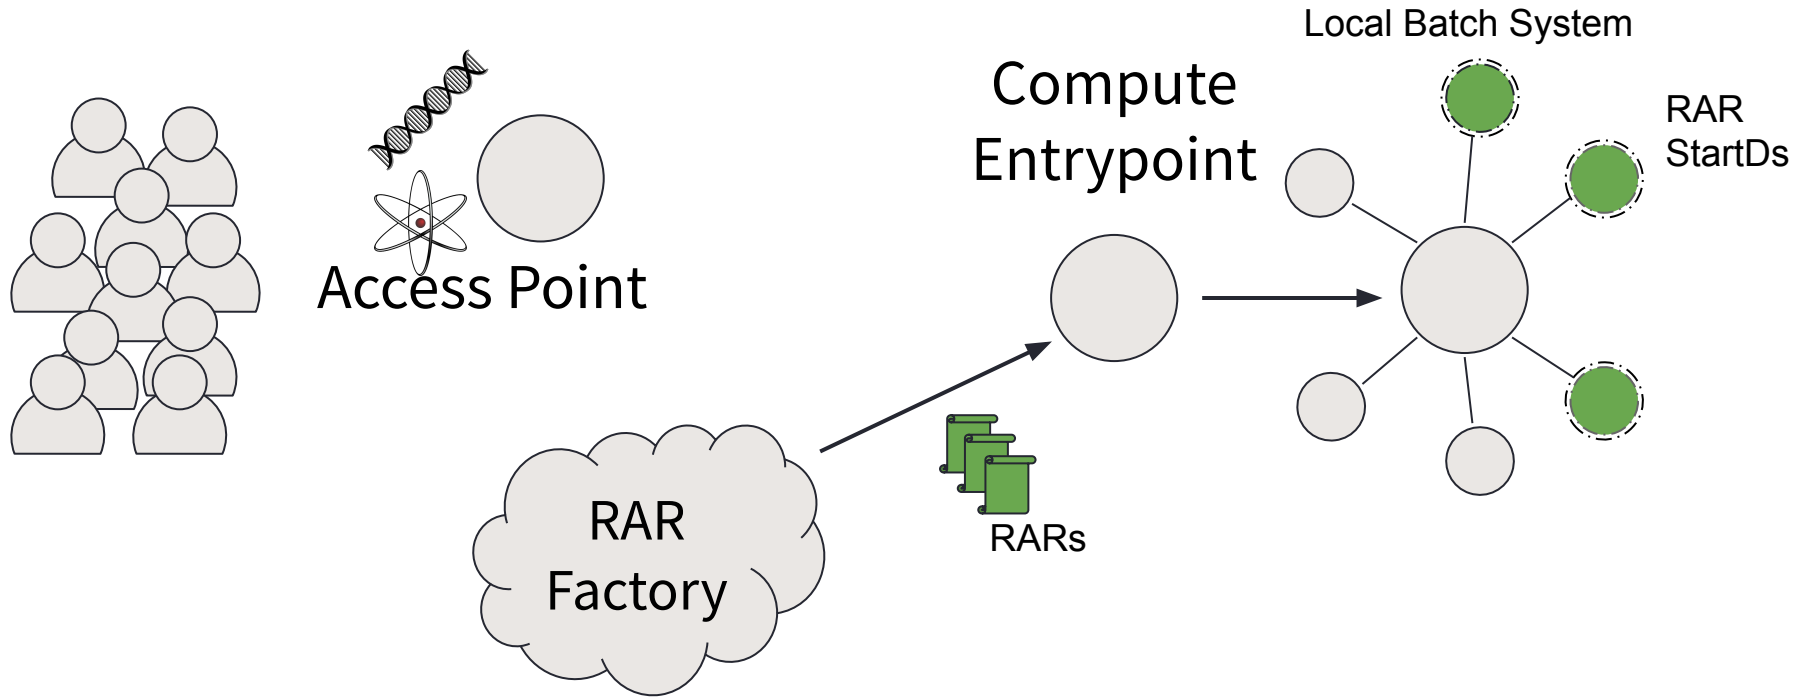
Local Batch System




# Resource Allocation Requests (RARs)



# Resource Allocation Requests (RARs)



# Resource Allocation Requests (RARs)



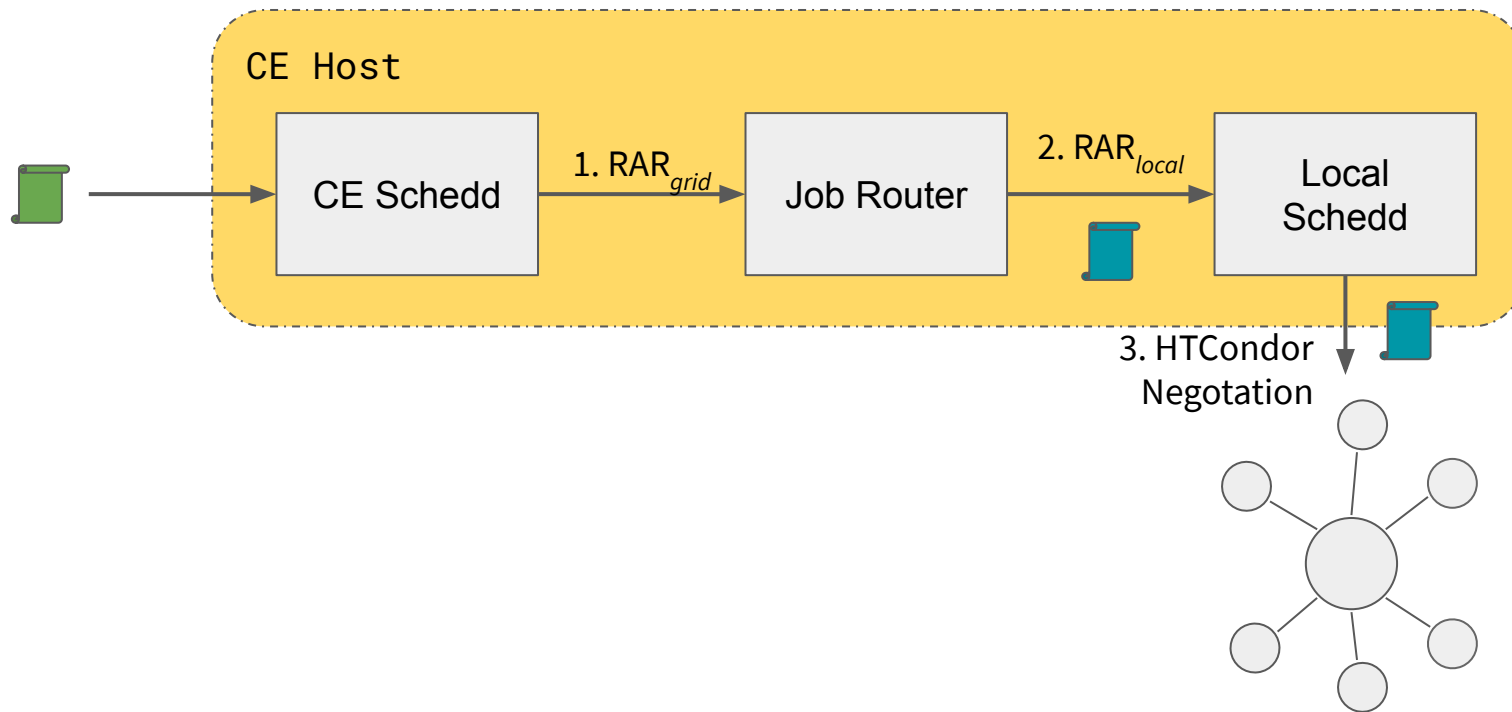
# HTCondor as a CE

- HTCondor-CE is HTCondor configured as a compute entrypoint
  - Same HTCondor binaries, description language (ClassAds), and configuration languages to provide the remote API
  - Relevant HTCondor tools are wrapped to use the HTCondor-CE configuration (e.g., `condor_ce_q`, `condor_ce_status`, etc.)
  - Separate `condor-ce` service
- Remote RAR submissions authorized via SciTokens/WLCG tokens or GSI/VOMS proxies
- Supports RAR submission to HTCondor, PBS, LSF, SGE, and Slurm. Also supports submission to these same batch systems using SSH
- Does not manage workers directly (no Negotiator or StartD)

# Job Router

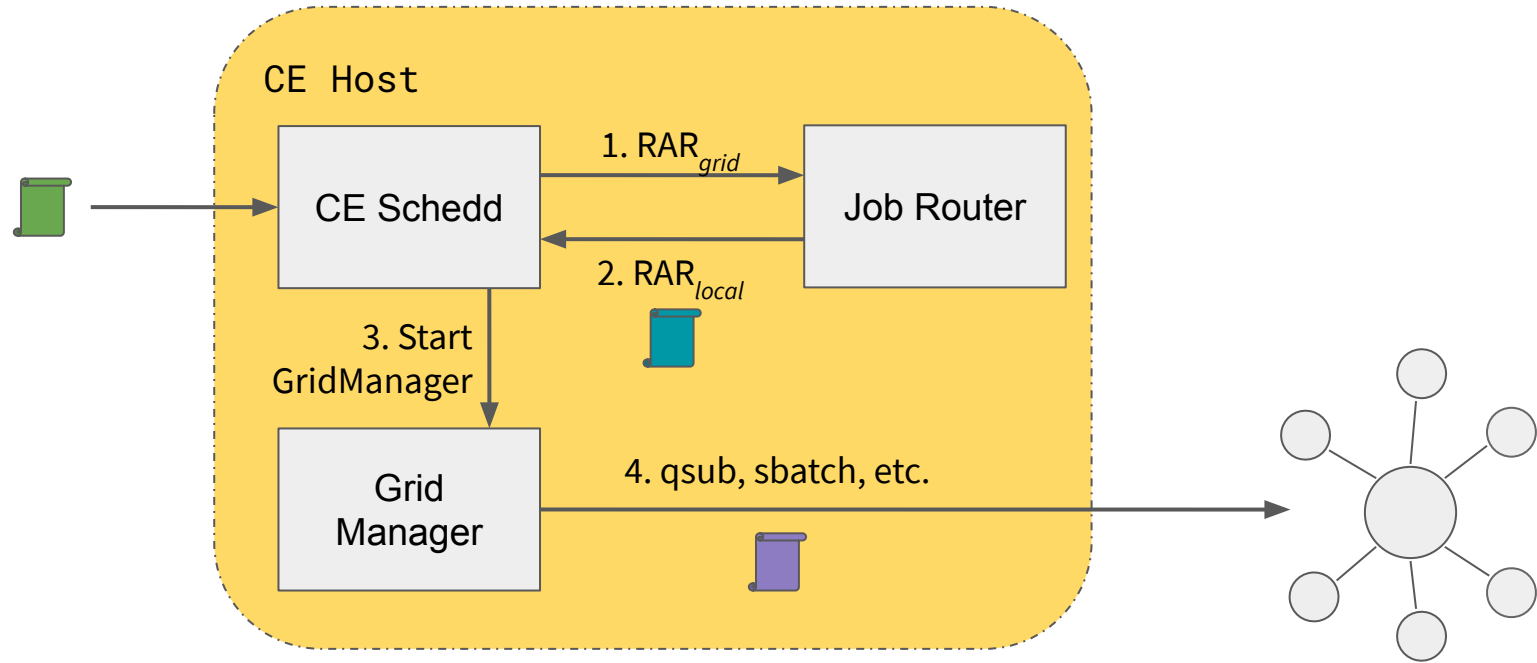
- The Job Router is responsible for taking a job, creating a copy, and changing the copy according to a set of rules
  - When running an HTCondor batch system, the copy is inserted directly into the batch SchedD. Otherwise, the copy is inserted back into the CE SchedD
  - Each chain of rules is called a “job route” and is defined by a transform or ClassAd
- Once the copy has been created, attribute changes and state changes are propagated between the source and destination jobs

# Local HTCondor Batch System





# Other Local Batch System



# HTCondor-CE 5!

- The new major version available alongside the HTCondor 9.0 and 9.1 series
- SCITOKENS authentication preferred over GSI (see Brian Bockelman's SciTokens talk from Tuesday)
- Support added for an improved Job Router configuration syntax: Job Router Transforms (see TJ's talk up next!)
- `condor_mapfile` formatting and location changes
- Most configuration moved to `/usr/share/condor-ce/config.d/`
- Moved to Python 3 and added support for Enterprise Linux 8 based operating systems
- `$HOME` is no longer set by default in the RAR environment
- Upgrade instructions and detailed release notes available at the redesigned docs (<http://htcondor.org>): <https://htcondor.github.io/htcondor-ce/v5/releases/>

# HTCondor-CE 5 Mapping Configuration

- CE administrators must now place their own RAR credential → local user mappings in a dedicated directory, `/etc/condor-ce/mapfiles.d/`
- Old customizations will be saved in `/etc/condor-ce/condor_mapfile.rpmsave`
- Downstream packagers can include default mappings in `/usr/share/condor-ce/mapfiles.d/`
  - OSG ships default SciToken → local user mappings in a separate package
  - Required HTCondor-CE mappings are stored in `/usr/share/condor-ce/mapfiles.d/50-common-default.conf`. Any lines in this file can be removed from your custom mappings!

# HTCondor-CE 5 Mapping Configuration

- Credential matching strings must be converted to regular expressions
  - In the old syntax, a mapping based on the SciToken issuer:  
SCITOKENS "https://scitokens.org/osg-connect" osgpilotuser
  - In the new syntax:  
SCITOKENS /^https:\\/\\/scitokens\\.org\\/osg-connect,/ osgpilotuser
  - Additionally, SciTokens can be mapped based on the SciToken's subject (e.g., "testing"):  
SCITOKENS /^https:\\/\\/scitokens\\.org\\/osg-connect,testing\$/ osgtestuser

# HTCondor-CE 5.1.2

- Job Router transform config fixes and improvements:
  - Fixed default CPU and memory requests
  - Fixed default per-route max idle and total job limits
  - CE startup verification checks for common misconfigurations
- Fix max runtime request expression for non-HTCondor batch systems
- Improved SciTokens support in `condor_ce_trace` and `condor_ce_run`, including WLCG token discovery (except for BEARER\_TOKEN)
- Added a standalone development container:  
<https://github.com/htcondor/htcondor-ce/blob/V5-branch/tests/containers/README.dev.md>

# What's Next?

- SCITOKENS troubleshooting guide
- Payload job auditing without GSI
  - Pilot jobs are currently allowed to advertise StartD ads back to the CE's Collector with GSI
  - Payload job information is then logged to `/var/log/condor-ce/AuditPayloadLog`
- Quality of life improvements
  - Job router transform debugging tool
  - Better SciTokens integration
- The ability to provision resources from a Kubernetes cluster
- Improved patch release cadence

# SciToken Transition

- **Oct 2021:** HTCondor 9.3.0 released without the Grid Community Toolkit. See detailed plan here:  
<https://htcondor-wiki.cs.wisc.edu/index.cgi/wiki?p=PlanToReplaceGridCommunityToolkit>
- **Feb 2022:** OSG retires Grid Community Toolkit from its software stack. See detailed plan here:  
<https://opensciencegrid.org/technology/policy/gridftp-gsi-migration/>
- **12 Oct 2021:** Pre-GDB on token/WebDAV transitions:  
<https://indico.cern.ch/event/876809/>
- **14-15 Oct 2021:** OSG Token Transition workshop:  
<https://opensciencegrid.org/events/Token-Transition-Workshop/>

# Questions?

This material is based upon work supported by the National Science Foundation under Grant Nos. 1148698, 1321762, 1836650, and 2030508. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.