

Introduction to HTC Facilitation

Christina Koch


OSG Virtual School 2021




Introductions

In the chat, share your institution,
your role, and who you support.


Facilitating HTC Workloads



Identify



Communicate



Support

Magnifying Glass by Adam Novantio from the Noun Project
Conversation by Maxim Kulikov from the Noun Project
sprout by lastspark from the Noun Project

Meeting a Researcher

Amina is fitting models to brain scan data using Matlab.

What questions would you ask to learn more?

Meeting a Researcher

Amina is fitting models to brain scan data using Matlab.

What questions would you ask to learn more?

Finding HTC Workloads

Tell me about
your research

...and how does
computing fit in?

What is your
near-term
bottleneck?

How are you running it now?

- compute/data requirements?
- your computing background?

How much
would you like
to run?

Finding HTC Workloads

Tell me about
your research

...and how does
computing fit in?

Takes two weeks
to analyze all
images (70) and
parameters (40)

- Running on a local cluster
- Each job using 16 cores,
runs for a day

Finish current
experiment
before defense.
Maybe 3-4 more?

Meeting a Researcher

Amina is fitting models to brain scan data using Matlab.

She's currently using a small local cluster. Each job fits many models (40 parameter combinations) to a single image. To speed things up, she's used Matlab's "parfor" option to parallelize this process, so each job uses 16 cores and takes about a day. There are 70 images to process this way, and she can only have about 5 jobs running at once on the cluster.

She'd like to run the current experiment in less time, especially since her defense is scheduled about 6 months from now. If possible, there's actually 3-4 experiments she'd like to run but with her current resources, she'd be happy just with one.

**Main question:
can it be broken up?**

Challenges and Opportunities

1. HTC is new
 - Requires thinking of work in a different way (splitting jobs; optimizing for throughput rather than single-job performance).
2. (OSG specific): OSG has unique job logistics
 - Data is transferred (no shared filesystem).
 - Software portability.
 - Jobs can be interrupted.
3. An HTC system can provide previously impossible scalability
 - On a small shared cluster, using 100s or 1000s of cores at once might be rare, or require long wait times. On the OSG, it's an everyday occurrence.


Communicating Outcomes

If you ...

- > execute your work ***this way*** (HTC)
- > requiring ***these learning steps*** and ***this much work***.

Then you ...

- > can reach your research goal after ***this much time and effort***, and
- > could achieve ***this much more research outcome***.



Outcomes for Amina

- **Jobs**

- 1 image, 40 combinations / job
- 24 hours / job (on 16 cores)
- 70 jobs total

- **Current Resources**

- 80 cores (5 nodes)
- 1920 core hours / day

- **Throughput**

- 5 jobs / day
- Overall time to completion: 2 weeks

- **Jobs**

- 1 image + 1 parameter
- 9.6 hours / job (on 1 core)
- 2800 jobs total

- **On an HTC system**

- 1000+ cores
- 24,000 core hours / day

- **Throughput**

- 2500 jobs / day
- Overall time to completion: 1 day + few hours

Support: Troubleshooting and Redirection

“I’m using a cron to submit a batch of jobs. When they complete, the cron checks and submits the next batch, that continue what the first jobs were doing. Jobs submitted by the cron aren’t running - can you help?”

Researcher has a script that generates 100 directories, each with their own “executable” with one different value in each. The script then loops through the directories, submitting one job at a time.

“My job is running for too long and getting put on hold. Is there a way to run for longer than the time limit?”

“Is it better to do one hundred 4 core jobs or four hundred 1 core jobs?”

Support: “Teach to Fish”

Provide an explanation of how you solved the problem.

R: *“My job isn’t starting, even though it’s been several hours.”*

F: *“The ‘-better-analyze’ flag for condor_q can give some information about idle jobs. I ran it with the job ID, and from that output, it looks like the job isn’t starting because there was a typo in the submit file.”*

Refer to existing documentation or training opportunities.

R: *“I’m getting errors from my jobs, saying that I’m over quota.”*

F: *“You can see this page for more information on checking your quota.”*

What are some of your strategies for making researchers self-sufficient?

The Big Picture

The goal is always a better research outcome for the researcher.

- **Choosing a “best-fit” computing approach that scales as much as possible**
 - HTC workloads compatible with the OSG can get a lot of computing.
- **Communicating the benefits and effort investment of that approach**
 - Is running on the OSG as “simple” as your desktop? Nope! Is it more powerful? Yes!
 - Quantify things! “You can achieve 10,000 runs in 2 days.” “You’re starting with 4,000, but you could easily run 40,000!”
- **Always keep this goal in mind when troubleshooting and providing guidance.**
 - Before solving a researcher’s question, stop and ask: is this the *real* problem?
 - Introduce skills that will help researchers help themselves.

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant No. 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.