



Open Science Grid

# Turning science problems into HTC jobs

Tuesday, Dec 7 2pm

Zach Miller

Condor Team

University of Wisconsin-Madison

# Overview

You now know how to run jobs using Condor, create basic workflows using DAGMan, and how to run a simple BLAST query. Let's put these pieces together to tackle larger problems.

This session will focus on how to break down and process large problems.



# Review of Blast Example

- First, run blast locally

Login to a treinamentoXX machine

```
source /opt/workshop/zkm_exercises/blast/blast.sh
```

```
blastp -db /opt/workshop/zkm_exercises/blastdb/yeast.aa -query query1
```

(You did this in the morning... This is intended as a sanity check to make sure your setup is still working)



# Think about running your application remotely

- What are the dependencies?
  - Architecture
  - OS // Linux Distro
  - Shared libraries
  - Input files
  - Environment variables
  - Scratch space
  - Available cpu ...

# Running BLAST under Condor

---

- Did we get all dependencies?
- Are we sized correctly?
- How long will the job run?
- How much data will it produce?
- What kind of load are we putting on various system resources?

# Sizing jobs for the grid

---

- If the job is too small, there's too much overhead
- If the job is too big, there's potentially for “badput”
- Badput is when a job runs for a while but is preempted and the output thrown away
- Rule of thumb: between 1 and 8 hours



# Other Considerations

---

- Besides how long a job will run, consider:
  - Memory requirements
  - Disk requirements
  - Network I/O
  - Consumable Licenses
- Try to identify all types of resources needed.

# Hands-on Exercise

---

- Processing multiple sequences
- You get all queries in one file
- Blast will accept input files with multiple queries
- Try running BLAST with the input file:
- `/opt/workshop/zkm_exercises/examples/five_inputs`
  - How long does it take?





# Hands-on Exercise

---

- Now, let's try running a much larger input  
`/opt/workshop/zkm_exercises/examples/larger_input`
- Note: it contains 6298 input files and will take 20 minutes or more! (Go ahead and submit it while I talk!)
  - `blastp -db yeast.aa -query large_input | grep '^Query=' | nl`
- Let's think about how we can do this more quickly...

# BLAST Input

---

- BLAST Input can be split at sequence boundaries
- Make a temporary directory in your home dir, and copy the file  
`/opt/workshop/zkm_exercises/examples/five_inputs`  
into the temporary directory
- Edit `five_inputs`, and look at the structure
- With only 5, it could be split manually, but...



# BLAST Input

---

- But with 6298 sequences, doing this manually is out of the question.
- We need some sort of automation!
- Write a script to split apart the input file, or copy  
`/opt/workshop/zkm_exercises/scripts/split_file.pl`
- Use your script first on `five_inputs` to create five input files:
  - `split_file.pl SMALL_TEST 1 < five_inputs`

# Submitting to Condor

---

- But how to submit?
- You could create a submit file for each new input file...
- Or you can use some fancy features in Condor to use a template submit file.
- Copy  
`/opt/workshop/zkm_exercises/examples/blast_template.sub`  
to a temporary directory and examine it

# Submitting to Condor

---

- Now run:
  - `condor_submit blast_template.sub`
- Watch it run with `condor_q`
- Examine the output (\*.out)
- Count the results:
  - `grep '^Query=' *.out | nl`

# Submitting to Condor

---

- Now do the same with a very large input:
  - Create a temporary directory and cd to it
  - `cp /opt/workshop/zkm_exercises/examples/large_input .`
  - `/opt/workshop/zkm_exercises/scripts/split_file.pl  
BIG_TEST 315 < large_input`
  - `cp /opt/workshop/zkm_exercises/examples/blast_template_2.sub .`
  - `condor_submit blast_template_2.sub`
- Again, watch it run:
  - `condor_q -dag`
  - `condor_q -run`



# Success?

---

- Problems with this approach:
  - Not completely automated
  - Requires editing template files
  - How do you know when the workflow is done?
  - How do you know it was all successful?

# Managing your Large Run

---

- Using DAGMan can help here!
- DAGMan brings the ability to implement several features:
  - Final notification when all pieces are complete
  - Verification that all results are present
  - Filtering or massaging the output into a final form
  - Throttling job submission for busy pools



# Managing your Large Run

---

- Once again, this can be automated
- Write a script that also generates a .dag file and also the submit files. The DAG will list all jobs but no relationships between them. (see `/opt/workshop/zkm_exercises/examples/example.dag`)
- Optionally, copy this script which does it for you:  
`/opt/workshop/zkm_exercises/scripts/gen_dag.pl`

# Managing your Large Run

---

- Try splitting the input into 20 pieces
- $6298 \text{ sequences} / 20 == 315 \text{ per file}$
- `gen_dag.pl LARGE_RUN 315 < large_input`
  - Look at what that script produced
- `condor_submit_dag LARGE_RUN.dag`
- Watch it run... How long is each piece taking?
- How would you change these numbers for an input of one million queries?

# Additional Work

---

- Modify your script to create a new node in the DAG that is a child of all other nodes
- Make a submit file for that node that runs a script. What should that script do?
  - Send email?
  - Verify results?
  - Concatenate results?
  - Compress results?
- How many of those you can implement?



# Notes on DAGMan

---

- Different nodes in the DAG can be different types of job: Vanilla, Grid, or Local
- Make the final node LOCAL universe (Check Condor Manual for details). You want it to run locally to verify that all results received are intact.
- Feel free to use and edit my script:  
`/opt/workshop/zkm_exercises/scripts/gen_dag_w_final.pl`



# Additional Work

---

- Run the individual pieces on OSG using Condor-G, instead of in the Vanilla universe
- Write another script that does everything using the pieces we just wrote:
  - Splits the input
  - Generates the DAG
  - Submits the workflow
  - Waits for completion (hint: `condor_wait`)
  - Combines results



# Additional Work

---

- Try running `five_inputs` and `large_input` against other databases. You'll need to:
  - Run a single input as a test. How long did it take?
  - Estimate how many compute hours this will take...
  - Decide how to split up the input appropriately
  - Get access to enough available resources to do this in a reasonable amount of time



# Questions?

---

- Questions about splitting, submitting, automating, etc?
- Break Time, resume after lunch