

Distributed Computing in IceCube

G. Merino, UW-Madison bootcamp
June 2016

Contact

Email:

- Problems/questions: help@icecube.wisc.edu
- discussion/information (mailing list): icecube-computing@icecube.wisc.edu

Slack channels:

- #icecube-it
- #software
- ...

A day in the life for an IceCube scientist




A day in the life for an IceCube scientist




IceCube computing at UW-Madison


The NPX cluster at UW-Madison is a great Collaboration resource

Currently ~7000 cpu cores... but it is normally very busy!





IceCube Distributed Resources



Funding Agencies

Fonds de la Recherche Scientifique (FRS-FNRS)
Fonds Wetenschappelijk Onderzoek-Vlaanderen (FWO-Vlaanderen)
Federal Ministry of Education & Research (BMBF)
German Research Foundation (DFG)


Deutsches Elektronen-Synchrotron (DESY)
Japan Society for the Promotion of Science (JSPS)
Knut and Alice Wallenberg Foundation
Swedish Polar Research Secretariat
The Swedish Research Council (VR)

University of Wisconsin Alumni Research Foundation (WARF)
US National Science Foundation (NSF)

Open Science Grid

A consortium of research communities who promote science via sharing of computing resources.


- Federation of more than 100 clusters
- Providing >100M cpuhours each month (~140k cpu cores)



Open Science Grid



OSG since Inception



Grid jobs vs local jobs

Remote execution environment will be different than that of the submit node

- No shared filesystems (no /data/exp, /data/ana/, /data/sim or /data/user)
- Jobs must transfer all input and output data: GridFTP
- Software accessed via CVMFS

Nodes are often access opportunistically. We are not owners.

- Some sites will evict your jobs if they need the nodes back.
- The less resources you need, the more cores you will get
 - Short jobs (<4 hours)
 - Small memory (<2GB)
 - Small disk (<20GB)

Accessing software: CVMFS

CernVM Filesystem (CVMFS)


Virtual software installation by means of an HTTP file system.

Advantages:

- Scalable
- Highly available
- Efficient
- Accessible from “anywhere”


Goal: facilitate the access of IceCube software anywhere (e.g. on Grid nodes).
No need for tarball making.

The same infrastructure can be used to access photonic tables for reconstruction.



(Image credit: Dan Bradley)

/cvmfs/icecube.opensciencegrid.org/



/cvmfs/icecube.opensciencegrid.org

Good news: great documentation at

http://software.icecube.wisc.edu/offline_trunk/metaproject/cvmfs.html

CVMFS

CVMFS (CernVM File System) is a shared filesystem available most anywhere around the world. It has become the standard for distributing software libraries in the physics community and beyond. For more details on CVMFS, see the [CERN documentation](#).

IceCube software is deployed at `/cvmfs/icecube.opensciencegrid.org`.

Quickstart ¶

From your shell (or in your login script) run:

```
eval $(/cvmfs/icecube.opensciencegrid.org/py2-v1/setup.sh)
```

This will detect the current OS and architecture and set your environment variables appropriately.

- If you are using bash, you can put the above line into your `.bash_profile` script to have the environment set on login.
- If you are using zsh, you can put the above line into your `.zlogin` script to have the environment set on login.
- If you are running tcsh, you should put the following line into your `.Login` script. (Note the use of backticks in the expression):



```
eval ` /cvmfs/icecube.opensciencegrid.org/py2-v1/setup.sh`
```

This should be near the end of any login script, so you don't accidentally put something else in front of CVMFS on the path.

Job Scheduling


HTCondor

HTCondor - NPX cluster



Glidein


Is just an execute node submitted as a job.



GlideinWMS

UW VO Frontend configuration

- Only IceCube jobs run at IceCube sites
- IceCube jobs can also run at any other OSG site



GridFTP


An extension of the FTP protocol for transferring large amounts of data in the WAN.

Adds needed features to fill the pipe:

- Buffer tuning
- Parallel streams


Authentication via Digital Certificate (X509)

Dedicated, high-performance Data Transfer Nodes: gridftp.icecube.wisc.edu



GlideinWMS: Authorization

Glideins carry X509 credentials (digital certificate) that identify them as “IceCube pilots” → IceCube group authorization



Let's give it a try...

```
ssh sub-1.icecube.wisc.edu
```

Simple job script

Advice: Job logfiles in network filesystems can generate instability in HTCondor.
Best to keep your job logfiles in local disk.

```
sub-1 ~ $ mkdir /scratch/gmerino/  
sub-1 ~ $ cd /scratch/gmerino/  
sub-1 ~ $ wget http://icecube.wisc.edu/~gmerino/bootcamp/job.sh
```

```
sub-1 ~ $ cat job.sh  
#!/bin/bash  
printf "Start time: "; /bin/date  
printf "Job is running on node: "; /bin/hostname  
printf "Job running as user: "; /usr/bin/id  
printf "Job is running in directory: "; /bin/pwd  
echo  
echo "Working hard..."  
sleep $1  
echo "Job complete!"
```

```
sub-1 ~ $ chmod +x job.sh  
  
sub-1 ~ $ ./job.sh 5
```

Submit File

```
sub-1 ~ $ cat job.sub
executable = job.sh
arguments = 10

log = job.log
output = job.out
error = job.err

request_cpus = 1
request_memory = 100MB
request_disk = 1GB

queue 1
```

List your executable and any arguments it takes.

Log: file created by HTCondor to track job progress

output/error: captures stdout/stderr

Request the appropriate resources for your job to run.

queue: keyword indicating “create a job”

Resource Request

If you do not specify resource requirements (cpus, memory, disk) the job will grab the default for the site where it will end up running.

Most sites will give you 1 cpu, ~2GB RAM, ~20GB disk... but this can vary a lot.

Some sites can have way smaller numbers as defaults! → your jobs will be killed if they overuse resources.

It is very important to request appropriate resources for a job

- Too little → your job might be killed for overusing resources
- Too much → your job will wait forever. There will be fewer “job slots” matching your requirements.

Run test jobs and use the log file to request the right amount of resources.

Submit File

```
sub-1 ~ $ cat job.sub
executable = job.sh
arguments = 10

log = job.log
output = job.out
error = job.err

request_cpus = 1
request_memory = 100MB
request_disk = 1GB

+WantGlidein = True
Requirements = IS_GLIDEIN &&
GLIDECLIENT_Group =?= "IceCube"

queue 1
```

Require to run in a “glidein node”

Require an “Icecube site” glidein

(normally friendlier for a tutorial...
properly configured CVMFS, etc)

Coming soon... a proper way to check for CVMFS

```
sub-1 ~ $ cat job.sub
executable = job.sh
arguments = 10

log = job.log
output = job.out
error = job.err

request_cpus = 1
request_memory = 100MB
request_disk = 1GB

+WantGlidein = True
Requirements = IS_GLIDEIN && HAS_CVMFS_icecube_opensciencegrid_org

queue 1
```

Submitting and monitoring

To submit jobs: `condor_submit`

```
sub-1 ~ $ condor_submit job.sh
Submitting job(s).
1 job(s) submitted to cluster 12898721.
```

To monitor submitted jobs: `condor_q`

```
sub-1 ~ $ condor_q gmerino
-- Schedd: sub-1.icecube.wisc.edu : <128.104.255.232:58276?...
  ID           OWNER           SUBMITTED      RUN_TIME ST PRI SIZE CMD
12898721.0    gmerino        6/13 19:32    0+01:45:39 R  0   1464.8 job.sh 100
```

http://research.cs.wisc.edu/htcondor/manual/v8.5/condor_submit.html

http://research.cs.wisc.edu/htcondor/manual/v8.5/condor_q.html

Submitting and monitoring

To submit jobs: `condor_submit`

```
sub-1 ~ $ condor_submit job.sh
Submitting job(s).
1 job(s) submitted to cluster 12898721.
```

To monitor submitted jobs: `condor_q`

```
sub-1 ~ $ condor_q gmerino
-- Schedd: sub-1.icecube.wisc.edu : <128.104.255.232:58276?...
  ID           OWNER           SUBMITTED      RUN_TIME ST PRI SIZE CMD
12898721.0     gmerino        6/13 19:32     0+01:45:39 R  0   1464.8 job.sh 100
```

A Job **Cluster** identifies a batch of jobs.

Individual jobs within a Cluster are identified by the number after the dot: the **Process Id.**

Submitting and monitoring

```
sub-1 ~ $ condor_submit job.sh
Submitting job(s).
1 job(s) submitted to cluster 12898721.
```

```
sub-1 ~ $ condor_q gmerino
-- Schedd: sub-1.icecube.wisc.edu : <128.104.255.232:58276?...
  ID          OWNER          SUBMITTED      RUN_TIME ST PRI SIZE CMD
12898721.0    gmerino          6/13 19:32    0+01:45:39 R  0   1464.8 job.sh 100
```

Most common job status cycle:

Idle (I): Job has not started yet... waiting in the queue.

→ **Running** (R): Job is running...

→ **Completed**: If the job has completed, it will not appear in condor_q

→ **Held** (H): Stalled jobs. Something YOU need to fix.

A job that goes on hold is interrupted (all progress is lost) and kept from running again. It remains in the queue in the “H” state.

Held jobs removed e-mail

```
From: root <root@sub-1.icecube.wisc.edu>
Date: 14 June 2016 at 04:17
Subject: [htcondor] sub-1: held jobs removed
To:

12925663.0    gmerino          6/13 23:54 Policy violation. Memory limit exceeded: 2004 MB
resident > 2000 MB requested. (by user condor)
12925674.0    gmerino          6/13 23:54 Policy violation. Memory limit exceeded: 2003 MB
resident > 2000 MB requested. (by user condor)
12925692.0    gmerino          6/13 23:54 Policy violation. Memory limit exceeded: 2005 MB
resident > 2000 MB requested. (by user condor)
12925704.0    gmerino          6/13 23:54 Policy violation. Memory limit exceeded: 2003 MB
resident > 2000 MB requested. (by user condor)
```

We periodically scan for held jobs in the queue, remove them and notify users via email.

Some other typical held reasons ...

```
12276425.0    gmerino          6/3  00:05 Error from slot1@e201.chtc.wisc.edu: Job failed to
complete in 72 hrs

12015071.5    gmerino          3/23 14:33 Error from glidein_7164_75405897@a0437: STARTER
at 10.80.2.181 failed to send file(s) to <128.104.255.232:59904>: error reading from
/home/icecu038/home_cream_966794713/CREAM966794713/glide_axZBCc/execute/dir_34476/_condo
r_stdout: (errno 2) No such file or directory; SHADOW failed to receive file(s) from
<134.93.174.12:35154>
```

Log File

```
sub-1 ~ $ cat job.log
000 (12898721.000.000) 06/13 21:50:14 Job submitted from host: <128.104.255.232:
58276?addrs=128.104.255.232-58276>
...
001 (12898721.000.000) 06/13 21:52:33 Job executing on host: <144.92.166.137:
27680?addrs=144.92.166.137-27680>
...
006 (12898721.000.000) 06/13 21:52:33 Image size of job updated: 1
    0 - MemoryUsage of job (MB)
    0 - ResidentSetSize of job (KB)
...
005 (12898721.000.000) 06/13 21:52:43 Job terminated.
    (1) Normal termination (return value 0)
        Usr 0 00:00:00, Sys 0 00:00:00 - Run Remote Usage
        Usr 0 00:00:00, Sys 0 00:00:00 - Run Local Usage
        Usr 0 00:00:00, Sys 0 00:00:00 - Total Remote Usage
        Usr 0 00:00:00, Sys 0 00:00:00 - Total Local Usage
314 - Run Bytes Sent By Job
281 - Run Bytes Received By Job
314 - Total Bytes Sent By Job
281 - Total Bytes Received By Job
Partitionable Resources :      Usage Request Allocated
    Cpus                  :              1          1
    Disk (KB)             :          12    102400    940361
    Memory (MB)           :              0         100         100
```

Finding Job Attributes

```
sub-1 ~ $ condor_q -long 12887688.0
JobStatus = 2
LastJobStatus = 1
User = "gmerino@icecube.wisc.edu"
Err = "/scratch/gmerino/Data/2013/logs/12887688.err"
Out = "/scratch/gmerino/Data/2013/logs/12887688.log"
NumJobStarts = 1
Args = "-g
/data/exp/IceCube/2013/filtered/level2/0505/Run00122300/Level2_IC86.
2013_data_Run00122300_0505_0_9_GCD.i3.gz -i
/data/exp/IceCube/2013/filtered/level2/0505/Run00122300/Level2_IC86.
2013_data_Run00122300_Subrun00000070.i3.bz2 -o
/data/ana/Cscd/StartingEvents/exp/IC86_2013/burnsample/13/00122300/Level2_IC
86.2013_data_Run00122300_Part00000070.i3.bz2"
RemoteHost = "slot1@e281.chtc.wisc.edu"
ResidentSetSize_RAW = 1200308
DiskUsage_RAW = 891690
RemoteUserCpu = 7669.0
```

Useful Job Attributes

JobStatus: number indicating Idle (1), Running (2), Held (5), etc

RemoteHost: where the job is running

ResidentSetSize_RAW: Maximum observed physical memory in use by the job in KiB while running.

DiskUsage_RAW: Maximum observed physical memory in use by the job in KiB while running.

RemoteUserCpu: The total number of seconds of user CPU time the job has used.

EnteredCurrentStatus: time of last status change

NumJobStarts: number of times the job started executing

http://research.cs.wisc.edu/htcondor/manual/v8.5/12_Appendix_A.html

Displaying Job Attributes

Use the “autoformat” option for condor_q

```
sub-1 $ condor_q gmerino -af JobStatus ClusterID ProcId RemoteHost  
ResidentSetSize_RAW  
  
2 12892531 0 slot1@glidein_235900_283164684@cabinet-0-0-7.t2.ucsd.edu 1272208  
2 12892986 0 glidein_10345_623188368@jux7c.zeuthen.desy.de 1290364  
2 12893002 0 slot1@e137.chtc.wisc.edu 1181296
```

The “-constraint” option can also be handy

```
sub-1 $ condor_q gmerino -c jobstatus==2 -af JobStatus ClusterID ProcId  
RemoteHost ResidentSetSize_RAW  
  
2 12892531 0 slot1@glidein_235900_283164684@cabinet-0-0-7.t2.ucsd.edu 1272208  
2 12892986 0 glidein_10345_623188368@jux7c.zeuthen.desy.de 1290364  
2 12893002 0 slot1@e137.chtc.wisc.edu 1181296
```

An IceTray Grid Job

IceTray remote file staging

IceTray can work with URLs as well as with local paths, and everything just works

```
gsiftp://gridftp.icecube.wisc.edu/data/exp/IceCube/2015/.../somefile.i3.bz2
```

Reading from remote files: File stager will issue the “gridftp download” command to stage the file into a temporary local directory, and provide a file handle to icetray.

Writing to remote files: IceTray will write the output into a temporary local directory, and File stager will issue a “gridftp upload” once the file handle is released.

IceTray-start wrapper script

http://software.icecube.wisc.edu/offline_trunk/metaproject/cvmfs.html

Self-contained IceTray Scripts

The CVMFS repository includes a wrapper script that you can use as the interpreter line in a Python script to automatically load an IceTray metaproject environment before execution. This makes the script easier to submit to a batch system while also documenting which metaproject it is intended to work with. To use it, put a line like the following at the top of your Python script:

```
#!/bin/sh /cvmfs/icecube.opensciencegrid.org/py2-v1/icetray-start
#METAPROJECT XXXXX
```

The metaproject specification XXXXX can either be

- a build directory:

```
#!/bin/sh /cvmfs/icecube.opensciencegrid.org/py2-v1/icetray-start
#METAPROJECT /data/user/you/metaprojects/icerec/build
```

- a tarball URL:

```
#!/bin/sh /cvmfs/icecube.opensciencegrid.org/py2-v1/icetray-start
#METAPROJECT http://convey.icecube.wisc.edu/data/user/you/tarballs/icerec-trunk
```

- one of the *[Pre-compiled metaprojects](#)* distributed through the CVMFS repository

Pre-compiled metaprojects

The CVMFS repository includes pre-built copies of commonly used versions of the offline-software, icerec, and simulation metaprojects. To use a pre-built metaproject, put a line like the following at the top of your Python script:

```
#!/bin/sh /cvmfs/icecube.opensciencegrid.org/py2-v1/icetray-start
#METAPROJECT: metaproject/VXX-YY-ZZ
```

IceTray Script

```
sub-1 ~ $ cat icetray-start_script.py
#!/bin/sh /cvmfs/icecube.opensciencegrid.org/py2-v1/icetray-start
#METAPROJECT offline-software/trunk

from argparse import ArgumentParser
from icecube import icetray, dataio
from I3Tray import *

parser = ArgumentParser()
parser.add_argument("-i", "--input", dest="INFILE")
parser.add_argument("-o", "--output", dest="OUTFILE")
args = parser.parse_args()

grid = 'gsiftp://gridftp.icecube.wisc.edu'
infile = grid + args.INFILE
outfile = grid + args.OUTFILE

tray = I3Tray()
tray.Add(dataio.I3Reader, Filename=infile)
tray.Add("Dump")
tray.Add("I3Writer", Filename=outfile)
tray.Execute(10)
```

```
sub-1 ~ $ chmod +x icetray-start_script.py
```

Submit many jobs with one submit file

HTCondor has built-in ways to submit multiple independent jobs with one submit file

Advantages

- Analyze multiple data files
- Test parameter combinations
- ...



Without having to create submit files for each job, and submit them individually.

IceTray script submit file

```
sub-1 ~ $ cat icetray.sub
executable = icetray-start_script.py
arguments = -i $(file) -o /data/user/gmerino/grid/$(Cluster)_$(Process)_$Fnx
(file)

log = /scratch/gmerino/logs/$(Cluster).log
output = /scratch/gmerino/logs/$(Cluster)_$(Process).out
error = /scratch/gmerino/logs/$(Cluster)_$(Process).out

request_cpus = 1
request_memory = 1GB
request_disk = 2GB

+WantGlidein = True
Requirements = IS_GLIDEIN && GLIDECLIENT_Group == "IceCube"

queue file from job_files.txt
```

```
cobalt ~ $ mkdir /data/user/gmerino/grid

cobalt ~ $ chmod a+wx /data/user/gmerino/grid
```

Icetray script submit file

```
sub-1 ~ $ cat icetray.sub
executable = icetray-start_script.py
arguments = -i $(file) -o
/data/user/gmerino/grid/$(Cluster)_$(Process)_$Fnx(file)
```

```
log = /scratch/gmerino/logs/$(Cluster).log
output = /scratch/gmerino/logs/$(Cluster)_$(Process).log
error = /scratch/gmerino/logs/$(Cluster)_$(Process).log
```

```
request_cpus = 1
request_memory = 1GB
request_disk = 2GB
```

```
+WantGlidein = True
Requirements = IS_GLIDEIN && GLIDECLIENT_Group
```

```
queue file from job_files.txt
```

\$F() - useful macro for manipulating file names

\$file = “./some/path/dir/filename.extension”

\$Fp(file) -> ./some/path/dir/

\$Fd(file) -> dir/

\$Fn(file) -> filename

\$Fx(file) -> .extension

\$Fnx(file) -> filename.extension

```
cobalt ~ $ mkdir /data/user/gmerino/grid
```

```
cobalt ~ $ chmod a+wx /data/user/gmerino/grid
```


Icetray script submit file

```
sub-1 ~ $ cat icetray.sub
executable = icetray-start_script.py
arguments = -i $(file) -o /data/user/gmerino/grid/$(Cluster)_$(Process)_$Fnx
(file)

log = /scratch/gmerino/logs/$(Cluster).log
output = /scratch/gmerino/logs/$(Cluster)_$(Process).out
error = /scratch/gmerino/logs/$(Cluster)_$(Process).out

request_cpus = 1
request_memory = 1GB
request_disk = 2GB

+WantGlidein = True
Requirements = IS_GLIDEIN && GLIDECLIENT_Group = 3

queue file from job_files.txt
```

Queue one job per entry in the job_files.txt file

Substitute \$file variable when expanding macro

```
cobalt ~ $ mkdir /data/user/gmerino/grid

cobalt ~ $ chmod a+wx /data/user/gmerino/grid
```

Files list for the submit script

```
sub-1 ~ $ cat job_files.txt
/data/sim/IceCube/2012/filtered/level2/CORSIKA-in-ice/11926/00000-00999/Level2_IC86.2012_corsika.011926.000000.i3.bz2
/data/sim/IceCube/2012/filtered/level2/CORSIKA-in-ice/11926/00000-00999/Level2_IC86.2012_corsika.011926.000001.i3.bz2
/data/sim/IceCube/2012/filtered/level2/CORSIKA-in-ice/11926/00000-00999/Level2_IC86.2012_corsika.011926.000002.i3.bz2
/data/sim/IceCube/2012/filtered/level2/CORSIKA-in-ice/11926/00000-00999/Level2_IC86.2012_corsika.011926.000003.i3.bz2
/data/sim/IceCube/2012/filtered/level2/CORSIKA-in-ice/11926/00000-00999/Level2_IC86.2012_corsika.011926.000004.i3.bz2
/data/sim/IceCube/2012/filtered/level2/CORSIKA-in-ice/11926/00000-00999/Level2_IC86.2012_corsika.011926.000005.i3.bz2
/data/sim/IceCube/2012/filtered/level2/CORSIKA-in-ice/11926/00000-00999/Level2_IC86.2012_corsika.011926.000006.i3.bz2
/data/sim/IceCube/2012/filtered/level2/CORSIKA-in-ice/11926/00000-00999/Level2_IC86.2012_corsika.011926.000007.i3.bz2
/data/sim/IceCube/2012/filtered/level2/CORSIKA-in-ice/11926/00000-00999/Level2_IC86.2012_corsika.011926.000008.i3.bz2
/data/sim/IceCube/2012/filtered/level2/CORSIKA-in-ice/11926/00000-00999/Level2_IC86.2012_corsika.011926.000009.i3.bz2
```

Submit 10 file processing jobs

```
sub-1 ~ $ cat job_files.txt
/data/sim/IceCube/2012/filtered/level2/CORSIKA-in-ice/11926/00000-00999/Level2_IC86.2012_corsika.011926.000000.i3.bz2
/data/sim/IceCube/2012/filtered/level2/CORSIKA-in-ice/11926/00000-00999/Level2_IC86.2012_corsika.011926.000001.i3.bz2
/data/sim/IceCube/2012/filtered/level2/CORSIKA-in-ice/11926/00000-00999/Level2_IC86.2012_corsika.011926.000002.i3.bz2
/data/sim/IceCube/2012/filtered/level2/CORSIKA-in-ice/11926/00000-00999/Level2_IC86.2012_corsika.011926.000003.i3.bz2
/data/sim/IceCube/2012/filtered/level2/CORSIKA-in-ice/11926/00000-00999/Level2_IC86.2012_corsika.011926.000004.i3.bz2
/data/sim/IceCube/2012/filtered/level2/CORSIKA-in-ice/11926/00000-00999/Level2_IC86.2012_corsika.011926.000005.i3.bz2
/data/sim/IceCube/2012/filtered/level2/CORSIKA-in-ice/11926/00000-00999/Level2_IC86.2012_corsika.011926.000006.i3.bz2
/data/sim/IceCube/2012/filtered/level2/CORSIKA-in-ice/11926/00000-00999/Level2_IC86.2012_corsika.011926.000007.i3.bz2
/data/sim/IceCube/2012/filtered/level2/CORSIKA-in-ice/11926/00000-00999/Level2_IC86.2012_corsika.011926.000008.i3.bz2
/data/sim/IceCube/2012/filtered/level2/CORSIKA-in-ice/11926/00000-00999/Level2_IC86.2012_corsika.011926.000009.i3.bz2
```

```
sub-1 ~ $ condor_submit icetray.sub
Submitting job(s).....
10 job(s) submitted to cluster 13075424.
```

```
sub-1 ~ $ tail -f /scratch/gmerino/logs/13075424.log
```

```
cobalt07 ~/bootcamp $ ls /data/user/gmerino/grid/13075424*
/data/user/gmerino/grid/13075424_0_Level2_IC86.2012_corsika.011926.000000.i3.bz2
/data/user/gmerino/grid/13075424_6_Level2_IC86.2012_corsika.011926.000006.i3.bz2
/data/user/gmerino/grid/13075424_1_Level2_IC86.2012_corsika.011926.000001.i3.bz2
/data/user/gmerino/grid/13075424_7_Level2_IC86.2012_corsika.011926.000007.i3.bz2
/data/user/gmerino/grid/13075424_2_Level2_IC86.2012_corsika.011926.000002.i3.bz2
/data/user/gmerino/grid/13075424_8_Level2_IC86.2012_corsika.011926.000008.i3.bz2
/data/user/gmerino/grid/13075424_4_Level2_IC86.2012_corsika.011926.000004.i3.bz2
/data/user/gmerino/grid/13075424_9_Level2_IC86.2012_corsika.011926.000009.i3.bz2
/data/user/gmerino/grid/13075424_5_Level2_IC86.2012_corsika.011926.000005.i3.bz2
```

Where did my jobs run?

condor_history command retrieves Class Ads for completed jobs. Accepts the same “autoformat” flag as condor_q

```
sub-1 ~ $ condor_history 13075424 -af LastRemoteHost
slot1@glidein_524087_132839132@cabinet-5-5-34.t2.ucsd.edu
slot1@glidein_2405516_480980500@cabinet-11-11-9.t2.ucsd.edu
slot1@glidein_2636731_630245952@red-c0821.unl.edu
slot1@glidein_1561967_43155000@cabinet-4-4-34.t2.ucsd.edu
slot1@glidein_752333_642282192@cabinet-8-8-6.t2.ucsd.edu
slot1@glidein_1563062_25053400@cabinet-4-4-34.t2.ucsd.edu
slot1@glidein_203500_32947863@cabinet-5-5-14.t2.ucsd.edu
slot1@glidein_2726240_44844460@cabinet-4-4-26.t2.ucsd.edu
slot1@glidein_466408_591189250@cabinet-4-4-16.t2.ucsd.edu
```

Contact

Email:

- Problems/questions: help@icecube.wisc.edu
- discussion/information (mailing list): icecube-computing@icecube.wisc.edu

Slack channels:

- #icecube-it
- #software
- ...

Class Ads for Computers

Our Grid submit machine “flocks to” several HTCondor pools

```
sub-1 ~ $ condor_config_val FLOCK_TO  
cm.chtc.wisc.edu, glidein2.chtc.wisc.edu, condor.hep.wisc.edu, condor.cs.wisc.  
edu, wid-cm.discovery.wisc.edu, condor.math.wisc.edu, condor1.doit.wisc.edu
```

glidein2.chtc.wisc.edu pool

```
gmerino@sub-1 ~/bootcamp $ condor_status -pool glidein2.chtc.wisc.edu
```

Name	OpSys	Arch	State	Activity	LoadAv	Mem	ActvtyTime
glidein_1941_12272	LINUX	X86_64	Claimed	Busy	1.000	2500	0+06:48:42
glidein_18606_6036	LINUX	X86_64	Claimed	Busy	15.930	4096	0+00:51:44
glidein_21369_3642	LINUX	X86_64	Claimed	Busy	16.020	4096	0+00:32:13
...							
glidein_62527_5628	LINUX	X86_64	Claimed	Busy	2.000	4096	0+01:16:30

Machines	Owner	Claimed	Unclaimed	Matched	Preempting
----------	-------	---------	-----------	---------	------------

X86_64/LINUX	5369	0	5109	260	0	0
--------------	------	---	------	-----	---	---

Total	5369	0	5109	260	0	0
-------	------	---	------	-----	---	---

glidein2.chtc.wisc.edu pool

```
gmerino@sub-1 ~/bootcamp $ condor_status -pool cm.chtc.wisc.edu
```

Name	OpSys	Arch	State	Activity	LoadAv	Mem	ActvtyTime
slot1@aci-001.chtc	LINUX	X86_64	Unclaimed	Idle	0.040	7021	0+00:29:40
slot1_10@aci-001.c	LINUX	X86_64	Claimed	Busy	1.000	3584	0+00:27:16
...							
slot1_8@wid-003.ch	LINUX	X86_64	Claimed	Busy	1.010	4096	1+14:50:21 32+19:05:10
slot2@vm-crane.cht	WINDOWS	X86_64	Owner	Idle	0.000	2047	32+19:05:11
Machines Owner Claimed Unclaimed Matched Preempting							
X86_64/LINUX	10668	0	10080	584	0	0	
X86_64/WINDOWS	2	2	0	0	0	0	
Total	10670	2	10080	584	0	0	

Machine Attributes

We can use the “-autoformat” with `condor_status`, like with `condor_q`

```
gmerino@sub-1 ~/bootcamp $ condor_status -pool glidein2.chtc.wisc.edu -af
RemoteOwner GLIDEIN_Site | grep gmerino | sort | uniq -c | sort -nr
  135 gmerino@icecube.wisc.edu DESY-ZN
   73 gmerino@icecube.wisc.edu SU-OG
   52 gmerino@icecube.wisc.edu UCSD
   52 gmerino@icecube.wisc.edu McGill
   45 gmerino@icecube.wisc.edu SPRACE
   19 gmerino@icecube.wisc.edu UCD
   19 gmerino@icecube.wisc.edu Nebraska
   17 gmerino@icecube.wisc.edu Caltech
   11 gmerino@icecube.wisc.edu Michigan
    9 gmerino@icecube.wisc.edu Omaha
    9 gmerino@icecube.wisc.edu NUMEP
    7 gmerino@icecube.wisc.edu NotreDame
    1 gmerino@icecube.wisc.edu Wuppertal
    1 gmerino@icecube.wisc.edu FIT
    1 gmerino@icecube.wisc.edu Clemson
```