



# ~~Special Environments~~

**There are 2 hard problems in computer science:  
cache invalidation, naming things, and off-by-1  
errors.**

**Thursday, July 28 2022**

**Mats Rynge**


# How do I ...

---

1. ... run GPU jobs?
2. ... find out if my home institution is sharing resources on the OSPool?
3. ... target operating systems?
4. ... my jobs run on the fastest nodes?
5. ... target CPU architecture / extensions?
6. ... find details about the execution points?

# What is a GPU?

- GPU = Graphical Processing Unit
- Has hundreds to thousands of “cores” that can be used to parallelize work.



Created by Ideologo Studio  
from Noun Project



# GPU Use Cases

---

- Programs that map well to GPUs include:
  - Deep learning
  - Molecular dynamics
  - Anything with lots of number crunching (like matrix operations) and low(er) data load.



# Submit File options

---

- Request GPUs with “request\_gpus”
- Can use custom requirements

```
request_gpus = 1
```

```
requirements = (GPUs_Capability >= 8.0)
```



```
requirements = (GPUs_Capability >= 8.0)

requirements = (OSGVO_OS_STRING == "RHEL 8")

requirements =
(HAS_CVMFS_oasis_opensciencegrid_org == TRUE) &&
(IsOsgVoContainer != True)
```

# Classads

Jobs and machines are described in classads

key = value

Matchmaking is two ways - jobs can have requirements on the machines, and the machines can have requirements on the jobs





# condor\_status

## condor\_status

```
. . .
slot1_1@glidein_47384_72490500@wn-ha-40.gina.surfsara.nl      LINUX      X86_64
Claimed      Busy
slot1@glidein_47378_404190132@wn-ha-41.gina.surfsara.nl    LINUX      X86_64
Unclaimed    Idle
slot1_1@glidein_47378_404190132@wn-ha-41.gina.surfsara.nl  LINUX      X86_64
Claimed      Busy
```

	Total	Owner	Claimed	Unclaimed	Matched	Preempting	Backfill	Drain
X86_64/LINUX	36630	0	33892	2691	0	47	0	0
Total	36630	0	33892	2691	0	47	0	0





# Pick one entry, use -l

```
condor_status -l slot1@glidein_39459_618000737@uct2-c566.mwt2.org
```

```
CPU = 24
GLIDEIN_ResourceName = "MWT2"
GLIDEIN_Site = "UChicago"
has_avx2 = true
JobStarts = 60
LoadAvg = 1.0
Memory = 64293
Microarch = "x86_64-v3"
Mips = 23578
OSGVO_OS_STRING = "RHEL 8"
OSGVO_CPU_MODEL = "AMD EPYC 7402 24-Core Processor"
START = (Owner == "rynge")
```



# Autoformat

```
condor_status -autoformat GLIDEIN_Site | sort | uniq -c
```

```
190 UConn-HPC
 39 University of Arkansas - Fayetteville
 33 University of Arkansas for Medical Sciences
105 University of Kansas
102 University of South Dakota
2497 UNL-PATH
 524 USD-Lawrence
   7 USF
  48 UTC-Epyc
  34 UW-IT
  25 Wichita State University
1514 Wisconsin
2490 WISC-PATH
  65 WTAMU-HPC
```



# Fastest machine

---

```
condor_status -autoformat Mips GLIDEIN_ResourceName | sort -n | head
```

```
38446 Colorado
```



# Only run jobs on faster than...

---

job.sub

```
requirements = Mips > 25000
```



# Operating Systems

```
condor_status -autoformat OSGVO_OS_STRING | sort | uniq -c
```

```
1078 DEBIAN 12
13809 RHEL 7
28388 RHEL 8
2550 RHEL 9
183 UBUNTU 20
1078 UBUNTU 22
```



# Pick an OS

---

job.sub

```
requirements = (OSGVO_OS_STRING == "RHEL 8")
```



# Or bring an OS

---

job.sub

```
container_image = "/cvmfs/singularity.opensciencegrid.org/htc/debian:12"
```



# Microarch == "x86\_64-v3"

The first of these microarchitecture levels, x86-64-v2, assumes the following on top of base level x86\_64 instructions:

```
CMPXCHG16B, LAHF-SAHF, POPCNT, SSE3, SSE4.1, SSE4.2, SSSE3.
```

This basically raises the processor feature level requirement to around Intel Nehalem, and supports any x86\_64 processor made in the last decade.

The x86-64-v3 microarchitecture requires the following instruction sets:

```
AVX, AVX2, BMI1, BMI2, F16C, FMA, LZCNT, MOVBE, XSAVE.
```

That is close to a Haswell processor, but does exclude some recent low end Intel CPU that removed AVX support.

Finally, x86-64-v4 requires:

```
AVX512F, AVX512BW, AVX512CD, AVX512DQ, AVX512VL
```





---

# Exercises