

Middleware Meeting Notes 2014-06-10

Date

June 10, 2014

Attendees

- LSST: [Kian-Tat Lim](#), [Greg Daues](#); HTCCondor: Greg Thain, Todd Tannenbaum

Discussion Items

GregD has been working on getting HTCCondor jobs assigned to slots according to data that they have cached locally. He has been working on two job assignment scenarios:

- 1) One best slot for each job, plus empty spares that are worse.
- 2) Best slot, empty spares, and already-used slots that are worse than both.

GregD:

- * Static ClassAds situation using job rank, scenario 1
 - + Submitting jobs slowly works well
 - + As jobs submitted faster, some go to "wrong" machine

GregT:

- * Job rank doesn't preempt
- * Negotiator pre-job rank may work better
 - + Use `condor_config_val NEGOTIATOR_PRE_JOB_RANK` to look at it
 - + Answer: `RemoteOwner =?= Undefined` (machines not currently used)

GregD:

- * Any way to look at ranks assigned?
 - + Tried debug in ClassAd expression, didn't see additional info

GregT:

- * Look in negotiator log

Todd:

- * Could also pass evaluated rank to job using `$$[]`
- * Machine rank would trump this if it is set

GregD:

- * Saw machine rank at 0 for all machines, which is correct

Todd:

- * What is negotiator interval?

GregD:

- * Not set, is default
- * Was also not seeing jobs going to second-best slot in scenario 2
 - + Saw it going to non-preferred slot instead

Todd:

- * Aside: shouldn't preempt a job by same (user) owner
- * When a job completes, slot goes to "claimed idle"
 - + Highest priority job gets assigned
- * Could force claim release
 - + Should be OK even with 2 minute jobs on 200 slots
- *** Will try to reproduce scenario 1 problem

GregD:

- * From last month: need to do slot-based dynamic ClassAds in HawkEye

Todd:

- * TJ has ticket to enhance HawkEye to enable specification of slots
-

GregD:

- * Will be trying fault tolerance scenarios
- * Possible for just one slot to disappear?

Todd:

- * All slots on a node are present or absent together
 - + Controlled by one process
- * Other scenarios to consider:
 - + Could have infinite loops in application code
 - + Could have processes block on filesystem or NIS or DNS
- * Killing processes gives cleaner notice to other side
- * Should notice within 20 min but possibly up to 2 hrs
 - + Collector ages machines out
 - + Collector option to keep track of absent machines

[blocked URL](#)

Action Items