



Analyzing Job/Machine Matches using `condor_q -analyze`

John (TJ) Knoeller
Condor Week 2013

Scenarios

- › Why is my job not running?
 - On any machine?
 - On a specific machine?
- › Why is this glide-in not running jobs?
- › How well do queued jobs match machines?
- › How well do machines match jobs?
- › Why does analysis sometimes “hang”?

Changes for 7.9 series

- › -better-analyze
- › -better-analyze:priority
- › -verbose
- › -machine <machine-name>
- › -mconstraint <constraint-expr>
- › -better-analyze:reverse
- › -analyze:summary

What's better about –better-analyze ?

- › `condor_q –analyze <jobid>`
 - Stops when it gets an answer (running, hold) or when job has not yet been considered
- › `condor_q –better <jobid>`
 - Same output as `–analyze` + more details
 - Always does match analysis

Analysis output sections

Request is running, held, etc.

Run analysis summary, of NNN machines

Requirements expression for your job is:

Your job defines the following attributes:

... your job reduces to these conditions:

Suggestions:

Example:

condor_q -analyze 42.0

42.000: Request is held.

Hold reason: Failed to initialize user log to
/home/work/gauss_rifle/42/process.log

Example:

condor_q -better 43.0

```
43.000:  Run analysis summary.  Of 3979 machines,  
        3979 are rejected by your job's requirements  
          0 reject your job because of their own requirements  
          0 match and are already running your jobs  
          0 match but are serving other users  
          0 are available to run your job  
        No successful match recorded.  
        Last failed match: Tue Nov 20 10:58:59 2012  
  
        Reason for last match failure: no match found
```

Ex: condor_q –better 43.0 (2)

WARNING: Be advised:

No resources matched request's constraints

The Requirements expression for your job is:

```
( Machine == "node99.chtc.wisc.edu" ) &&  
( MY.JobUniverse == 12 || MY.JobUniverse == 7 ||  
  MY.WantFlocking || MY.WantGlidein ||  
  TARGET.PoolName == "CHTC" ||  
  TARGET.COLLECTOR_HOST_STRING == "kiddiepool.cs.wisc.edu" ) &&  
( TARGET.Arch == "X86_64" ) && ( TARGET.OpSys == "LINUX" ) &&  
( TARGET.Disk >= RequestDisk ) &&  
( TARGET.Memory >= RequestMemory ) &&  
( TARGET.HasFileTransfer )
```


Ex: condor_q –better 43.0 (3)

Your job defines the following attributes:

JobUniverse = 5

WantFlocking = undefined

WantGlidein = undefined

DiskUsage = 22

ImageSize = 22

RequestDisk = 22

RequestMemory = 1

Ex: condor_q –better 43.0 (4)

The Requirements expression for your job reduces to these conditions:

	Slots	
Step	Matched	Condition
-----	-----	-----
[0]	4	Machine == "node99.chtc.wisc.edu"
[8]	3672	TARGET.PoolName == "CHTC"
[12]	0	[0] && [8]

Why is my job NOT running on this idle machine

› `condor_q -bet <jobid> -machine <name>`

› But what do you do when you get this?

...

0 are rejected by your job's requirements

1 reject your job because of their own requirements

When machines reject your job “because of their own requirements”

- › `condor_q -bet:reverse -mach <machine>`
 - Counts jobs whose Requirements match this machine
 - Counts jobs that match this machine's Requirements and START expressions
 - Analyzes machine's Requirements and START expressions

Example: condor_q -bet:reverse

```
> condor_q -pool glidein.chtc.wisc.edu -bet:rev -mach glidein_15533@cmswn1332.fnal.gov
```

...

The Requirements expression for this slot is

```
( START ) && ( IsValidCheckpointPlatform )
```

START is ((true) &&

```
( ( x509userproxysubject isnt undefined ) &&
```

```
( ( GLIDEIN_REQUIRE_VOMS is undefined ) ||
```

```
( GLIDEIN_REQUIRE_VOMS is false ) ||
```

```
( TARGET.x509userproxyfirstfqan isnt undefined ) ) && ( true ) ) &&
```

```
( TARGET.WantGlidein && ( ( isUndefined(osg_site_whitelist) ||
```

```
stringListMember(GLIDEIN_Site,osg_site_whitelist) is true ) &&
```

```
( isUndefined(osg_site_blacklist) ||
```

```
stringListMember(GLIDEIN_Site,osg_site_blacklist) isnt true ) ) ) &&
```

```
( x509UserProxySubject isnt undefined &&
```

```
( ( isUndefined(TARGET.estimated_run_hours) ||
```

```
TARGET.estimated_run_hours * 3600 <= GLIDEIN_ToDie - CurrentTime ) ) ) &&
```

```
( ( ( GLIDEIN_ToRetire is undefined ) ||
```

```
( CurrentTime < GLIDEIN_ToRetire ) ) )
```

Ex: condor_q -bet:reverse (2)

This slot defines the following attributes:

```
CheckpointPlatform = "LINUX X86_64 2.6.x normal N/A"  
CurrentTime = 1366660640  
GLIDEIN_REQUIRE_VOMS = true  
GLIDEIN_Site = "FNAL"  
GLIDEIN_ToDie = 1366822155  
GLIDEIN_ToRetire = 1366787355  
IsValidCheckpointPlatform = true
```

Ex: condor_q -bet:reverse (3)

The Requirements expression for this slot reduces to these conditions:

Step	Clusters Matched	Condition
[1]	0	x509userproxysubject isnt undefined
[5]	0	TARGET.x509userproxyfirstfqan isnt undefined
[11]	141	TARGET.WantGlidein
[12]	210	isUndefined(osg_site_whitelist)
[15]	210	isUndefined(osg_site_blacklist)
[21]	0	x509UserProxySubject isnt undefined
[22]	210	isUndefined(TARGET.estimated_run_hours)

Ex: condor_q -bet:reverse (4)

```
glidein_15533@cmswn12.fnal.gov: Run analysis summary of 1877 jobs.  
  0 (0.00 %) match both slot and job requirements.  
  0 match the requirements of this slot.  
782 have job requirements that match this slot.
```


How well do jobs match?

> condor_q -ana:sum [<user>]

```
> condor_q -ana:sum
```

```
Analyzing matches for 4099 slots
```

JobId	Autocluster Members/Idle	Matches Requirement	Machine Rejects	Job	Running Users	Job	Serving Other	User	Avail	Owner
23724675.0	1/0	3371	706		1/4		2330	0	sierra	
23724668.0	1/0	3455	779		1/9		2335	0	sierra	
24161055.0	7/7	16	0		16		0	0	tango	
24175421.0	100/0	3178	470		100/300		2118	0	zulu	
24175550.0	92/0	3178	470		92/300		2118	0	zulu	
23989910.0	47/0	3418	717		47/17		2362	0	november	
23861860.0	17/0	1380	430		0		646	0	charlie	
23559579.0	1/1	0	0		0		0	0	delta	
23724669.0	1/0	3455	779		1/9		2335	0	sierra	
23859855.0	8/0	1380	430		0		646	0	charlie	
24212524.0	7/0	40	0		7/17		0	0	thomas	
18151825.0	5/0	3181	509		0		2382	0	denver	

How well do machines match?

> `condor_q -ana:sum,rev [-mconst <const>]`

Analyzing matches for 1484 jobs

Name	Slot Type	Slot's Req Matches	Job's Req Job	Matches Slot	Both Match %
slot1@lab-008.wisc.edu	Stat	1469		0	0.00
slot2@lab-008.wisc.edu	Stat	1469		0	0.00
slot1@lab-037.wisc.edu	Stat	1469		0	0.00
slot2@lab-037.wisc.edu	Stat	1469		0	0.00
slot1@node002.chtc.wisc.edu	Part	1411		1378	92.79
slot1_5@node002.chtc.wisc.edu	Dyn	1271		1178	79.31
slot1_6@node002.chtc.wisc.edu	Dyn	84		72	4.85
slot1_8@node002.chtc.wisc.edu	Dyn	1271		1178	79.31
slot1@node006.chtc.wisc.edu	Part	0		1377	0.00
slot1_1@node006.chtc.wisc.edu	Dyn	1295		1263	85.04

Working offline

> -jobads <file>

- <file> is the output of `condor_q -long` or `condor_submit -dump`

> -slotads <file>

- <file> is the output of `condor_status -long`

> -userprios <file>

- <file> is the output of `condor_userprio -long`
- Use with `-better:priority`



Any Questions?