

Lies, !* lies, and Statistics



Job statistics in the SCHEDD

Motivation

- > My jobs keep failing!
- > I'm seeing IO spikes on network!
- > How much execution time is badput?
- > How big/long running are the jobs in the pool?

“Statistics” in the SCHEDD

- > Most are counters that always increase
- > Lost when the SCHEDD exits
- > Published to the Collector (or not)
- > `Condor_status -direct name -schedd`
- > more stats with `-statistics SCHEDD:2`

Instantaneous stats

- > Shadows/Jobs running
- > Using/Waiting for Transfer Queue
- > Histogram of memory use
64Kb, 256Kb, 1Mb, 4Mb, 16Mb, 64Mb, 256Mb, 1Gb, 4Gb, 16Gb, 64Gb, 256Gb
- > Histogram of accumulated run time
30Sec, 1Min, 3Min, 10Min, 30Min, 1Hr, 3Hr, 6Hr, 12Hr, 1Day, 2Day, 4Day, 8Day, 16Day

Aggregate stats

- > Job counts and accumulated time
- > Shadows recycled, reconnected
- > Histograms of memory and runtime
- > Over SCHEDD lifetime
- > Over last 1200* seconds
 - (yes there's a knob for that)

Attribute Name Pattern

- > Jobs* are over SCHEDD lifetime
- > Shadows* are over SCHEDD lifetime
- > RecentJobs* are over last 20 min
- > RecentShadows* are over last 20 min
- > TransferQueue* are current values

```
$ condor_status -direct submit -schedd -long | grep ^RecentJobs
```

```
RecentJobsStarted = 714
RecentJobsSubmitted = 1057
RecentJobsCompleted = 428
RecentJobsExited = 987
RecentJobsCheckpointed = 0
RecentJobsCoredumped = 0
RecentJobsDebugLogError = 0
RecentJobsExecFailed = 0
RecentJobsExitedAndClaimClosing = 18
RecentJobsExitedNormally = 428
RecentJobsExitException = 151
RecentJobsKilled = 1
RecentJobsMissedDeferralTime = 0
RecentJobsNotStarted = 46
RecentJobsShadowNoMemory = 0
RecentJobsShouldHold = 11
RecentJobsShouldRemove = 0
RecentJobsShouldRequeue = 350

RecentJobsAccumTimeToStart = 1019765
RecentJobsAccumRunningTime = 787867
RecentJobsAccumBadputTime = 7199
RecentJobsAccumPreExecuteTime = 6616
RecentJobsAccumExecuteTime = 791867
RecentJobsAccumPostExecuteTime = 6246

RecentJobsBadputRuntimes = "78, 0, 0, 4, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0"
RecentJobsBadputSizes = "74, 0, 0, 0, 0, 0, 4, 4, 0, 0, 0, 0, 0"
RecentJobsCompletedRuntimes = "246, 0, 0, 4, 1, 39, 14, 3, 0, 5, 0, 0, 0, 0"
RecentJobsCompletedSizes = "246, 0, 0, 0, 11, 0, 4, 43, 4, 4, 0, 0, 0"

JobsRuntimesHistogramBuckets = "30Sec, 1Min, 3Min, 10Min, 30Min, 1Hr, 3Hr, 6Hr,
12Hr, 1Day, 2Day, 4Day, 8Day, 16Day"
JobsSizesHistogramBuckets = "64Kb, 256Kb, 1Mb, 4Mb, 16Mb, 64Mb, 256Mb, 1Gb,
4Gb, 16Gb, 64Gb, 256Gb"
```



```
$ condor_status -direct submit -schedd -long -statistics SCHEDD:2 | grep Shadows
```

```
ShadowsRunning = 2347  
ShadowsRunningPeak = 4463  
ShadowsRecycled = 234773  
RecentShadowsRecycled = 378  
ShadowsStarted = 766014  
RecentShadowsStarted = 1015  
ShadowsReconnections = 0  
RecentShadowsReconnections = 0
```



SCHEDD_COLLECT_STATS_FOR_*

- > A ClassAd expression
 - Evaluated against the JobAd
 - Returns true or false
- > Produces a whole new set of stats
 - Attribute names begin with "*"
 - Counts only jobs where <expr> is true

SCHEDD_COLLECT_STATS_FOR_SOAR

```
SCHEDD_COLLECT_STATS_FOR_SOAR = Owner=="soaruser"
```

```
$ condor_status -l -direct name -schedd | sort | grep ^Soar
```

```
SoarJobsAccumBadputTime = 11434
```

```
SoarJobsAccumExecuteTime = 1982281
```

```
SoarJobsAccumPostExecuteTime = 1947
```

```
SoarJobsAccumPreExecuteTime = 9060
```

```
...and 50 more attribs...
```



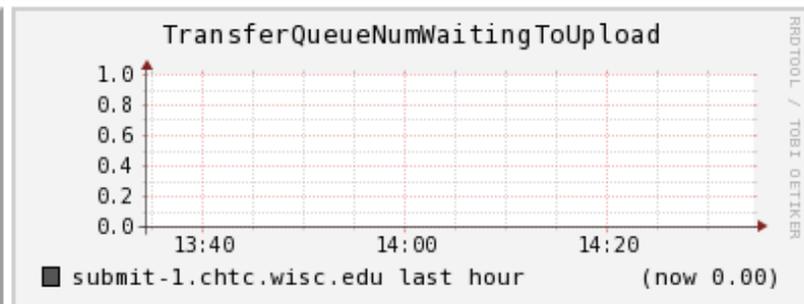
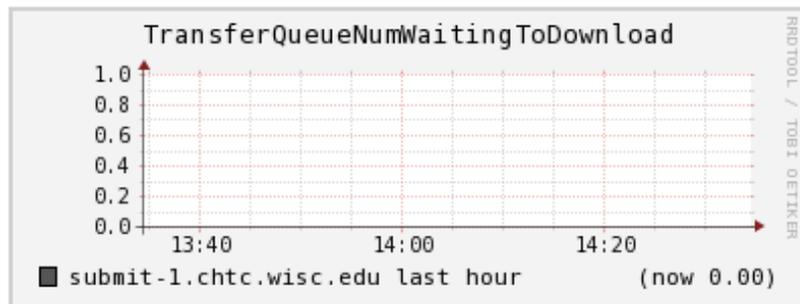
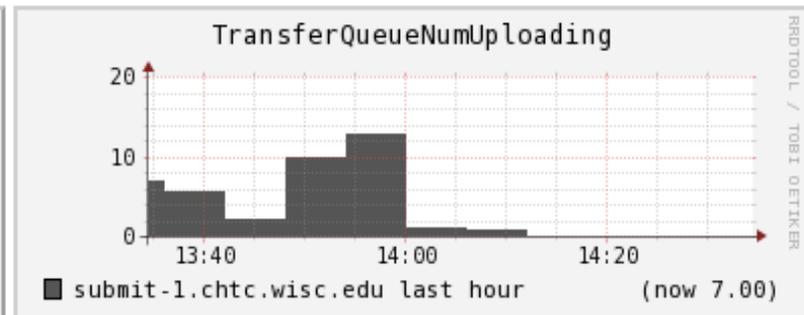
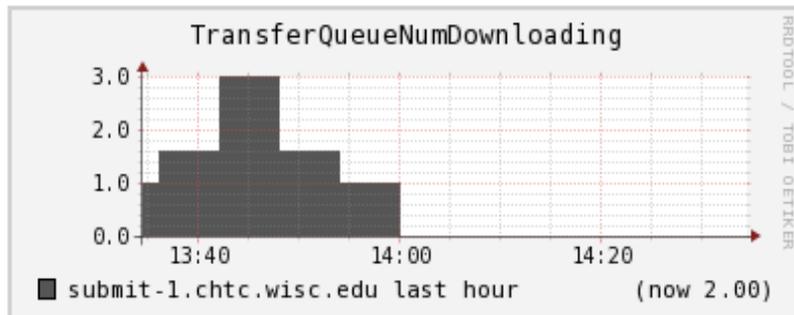
Graphing the stats

- > Some values are useful as-is
 - TransferQueueNumUploading
- > Some values should be shown as diffs
 - JobsStarted/JobsExited

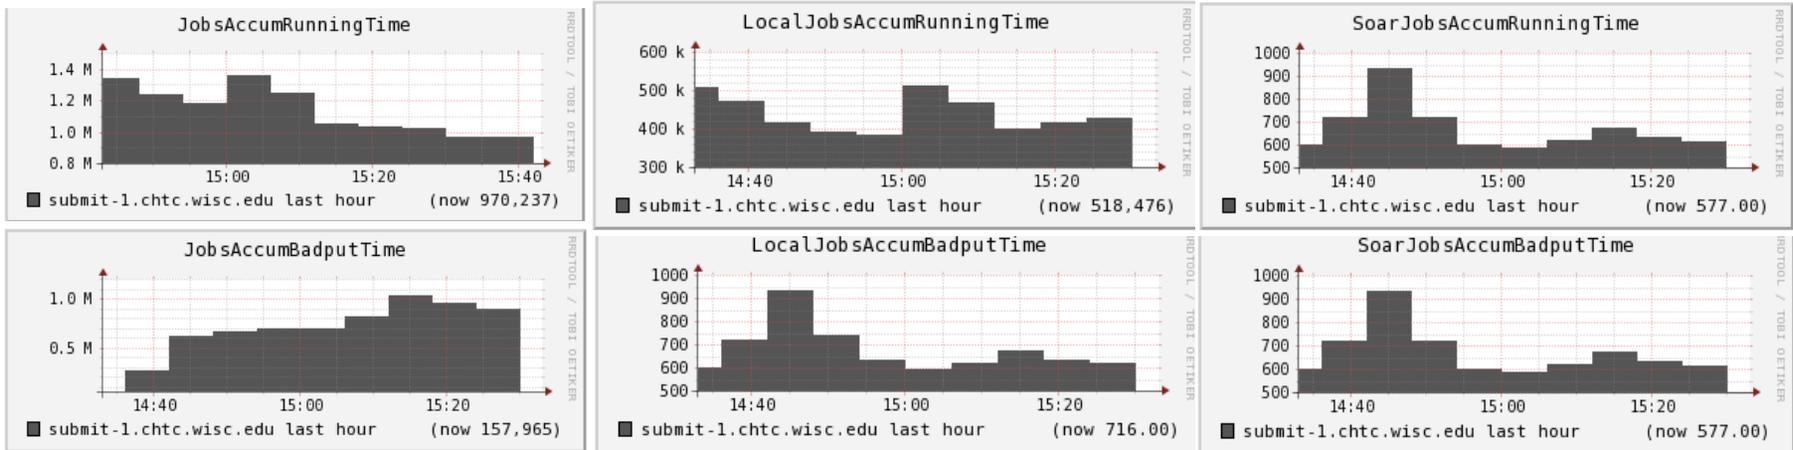
Making a Ganglia Plugin

- > Create a script to query Condor
- > Calculate deltas between current data and previous query
- > Publish the results to Ganglia

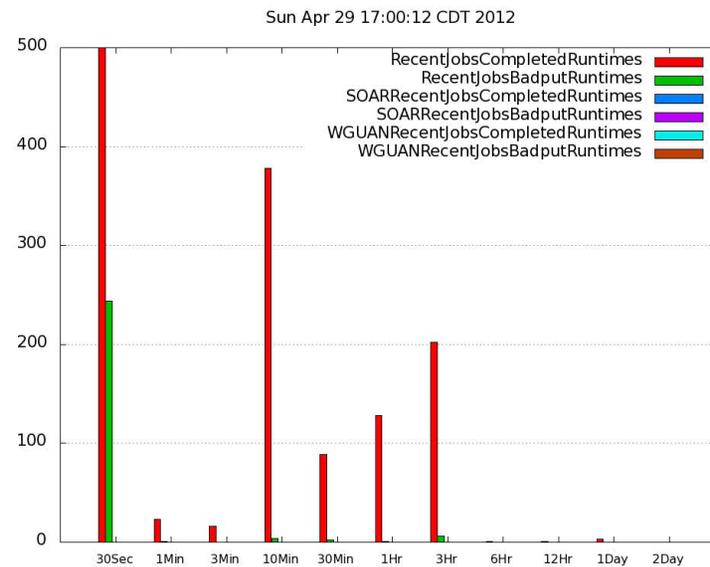
Ganglia Graphs - I/O issues



Ganglia Graphs - Badput



Histogram stats



Coming soon

> `SCHEDD_COLLECT_STATS_BY_* =`
`<expression>`

`SCHEDD_COLLECT_STATS_BY_OWNER = Owner`

```
$ condor_status -l -direct name -schedd | grep ^Owner
```

```
Owner_smith_RecentJobsCompleted = 0
```

```
Owner_jones_RecentJobsCompleted = 0
```

... and many many more...



www.cs.wisc.edu/Condor

