



HTCondor stats

John (TJ) Knoeller
Condor Week 2016

General statistics

- › Daemons track count and runtime stats
- › Stats are assigned verbosity levels
 - 0 = Chatty
 - 1 = Normal
 - 2 = Verbose
 - 3 = Never
- › Stats are published in daemon ad
STATISTICS_TO_PUBLISH = <cat>:<verb>

STATISTICS_TO_PUBLISH

› Basic control of stats to collector

- Categories are

- DC – Commands, Timers, Signals, DNS Resolver

- SCHEDD – Job counters

- TRANSFER – Transfer Queue counters

- COLLECTOR – Update counters

- ALL – All categories

- DEFAULT – Any category not otherwise specified

STATISTICS_TO_PUBLISH = DEFAULT:0, SCHEDD:1

STATISTICS_TO_PUBLISH_LIST

- › Fine control of stats to collector
- › Set to a list of attribute names
- › Changes verbosity of stats that match to equal the STATISTICS_TO_PUBLISH verbosity

Condor_status -statistics

`Condor_status -statistics <cat>:<verb>`

- › Override STATISTICS_TO_PUBLISH
- › Use with `condor_status -direct -schedd`
 - ‘live’ DC, SCHEDD and/or TRANSFER stats
- › Use with `condor_status -collector`
 - ‘live’ DC and/or COLLECTOR stats

Statistics

- › Time spent on DNS Lookups
 - Special counter for 'slow' lookups
- › Counter for ResourceRequestsSent
- › Per-user transfer stats in Submitter ads

Types of statistics probes

- › Absolute values
 - ShadowsRunning
- › Counters – count up as long as the daemon is alive
 - JobsCompleted
- › Runtime counters – count up and accumulate time
 - DCNameResolve / DCNameResolveRuntime
- › Moving averages – count up with windowed averages
 - FileTransferUploadBytesPerSecond_1h
- › “Miron” probes – count / min / max / avg / std
 - DCCCommand_NEGOTIATERuntimeAvg

Most stats are multiple attrs

› Runtime counters

- Foo = count, FooRuntime = runtime

› Recent window stats

- Prefix “Recent” is stat over RecentWindowMax

› Moving averages

- Foo = count, FooPerSecond_<time> = avg
- Can be multiple <time> windows

Most stats are multiple attrs(2)

- › “Miron” runtime probes
 - Foo = count
 - FooRuntime = accumulated time
 - FooRuntimeMin
 - FooRuntimeMax
 - FooRuntimeAvg
 - FooRuntimeStd
- › See DCCommand_RESCCHEDULE

Querying statistics

- › Use `condor_status` to get recent-ish stats
 - “Recent” prefix attributes are quick ‘n dirty
 - `RecentWindowMax` – span of “Recent” time
 - `RecentStatsTickTime` – last “Recent” update
- › Use `condor_status –direct –statistics`
 - Difference of 2 queries is most accurate view
 - Works with `schedd`, `submitter` ads
 - `StatsLastUpdateTime` – last stats update

Taxonomy of stats names

- › DC category stats start with DC
 - Mostly runtime stats
 - DCCommand, DCTimer, DCPumpCycle, etc
 - RecentDamonCoreDutyCycle
- › SCHEDD category stats start with
 - Jobs – job counters updated on shadow exit
 - SC – Schedd runtime stats
- › TRANSFER category stats start with
 - FileTransfer – byte counts, transfer rates
 - TransferQueue – use counts for transfer queue

What do we measure

- › Running, Idle and Held Job counts by universe
- › “Jobs” prefix
 - Counts by exit code
 - Runtime & Badput time
- › “FileTransfer”
 - Bytes and rates in and out
 - Bytes and rates for each submitter

What do we measure (2)

- › “Transfer” prefix
 - Transfer queue usage
- › “SC” prefix
 - Schedd task execution time
- › “Updates” prefix
 - Collector update rates
- › “DC” prefix
 - execution time for commands, DNS, fsync
 - Daemon health

Show only my jobs

- › For ordinary users, condor_q will default to showing only jobs from that user unless
 - User is a queue superuser (root, condor)
 - -allusers arg is used
 - 8.5.1 or earlier SCHEDD
 - CONDOR_Q_ONLY_MY_JOBS = false
 - For tool OR for SCHEDD config

condor_q -batch

- › One line of output for each batch of jobs where 'batch' is
 - A DAG and all of it's children
 - A Cluster and all of it's procs
 - A Clusters that have the same Owner and Cmd attributes. (Cmd is Executable)
- › -batch will be the new default in 8.6
 - -unbatch to get pre 8.6 output

Random stuff

- › Accountant ads in the collector
- › Revamped keyboard / mouse detection
 - Using X screen saver api
 - Changed defaults to protect all slots
- › Condor_status & condor_q pay attention to terminal width, truncate data less.