From Per-Job Data to an Aggregated Workload Insight:

A Toolkit for Profiling HTCondor Workload

Kashika Mahajan

Problem

Researchers want to know how their workload is doing without having to dig for details

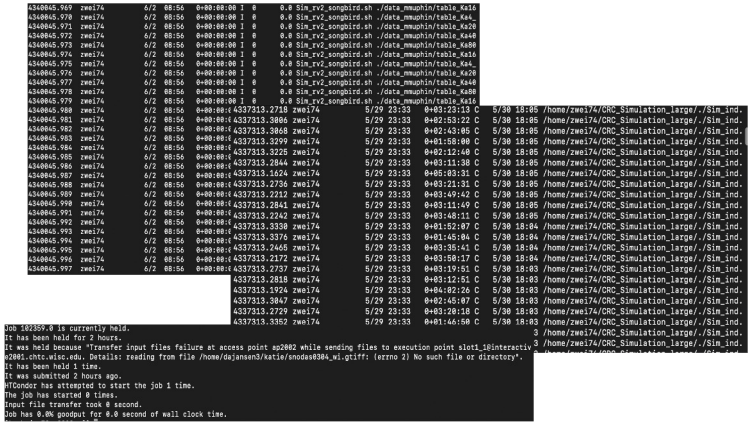

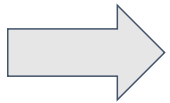

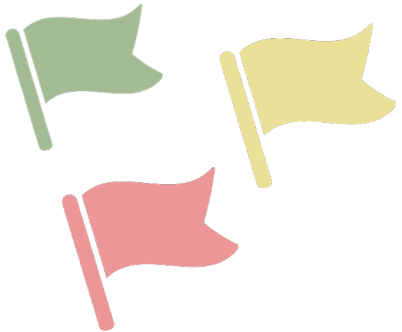

CHTC HTCondor Software Suite PATh

Objective of Project

Make a tool (or tools) that gives researchers simple feedback about their workload status

Visual

Evaluative

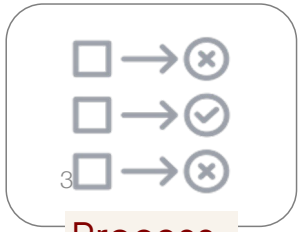

Process

CHTC HTCondor Software Suite PATh

What do Researchers want to know?

- Status
- Time
- Resources
- Jobs on Hold

How is my workload doing?

How we are providing this information:

- Status -> Dashboard
- Time -> Runtime Visualizations
- Resources -> Resource Utilization Report
- Jobs on Hold -> Hold Classifier

How is my workload doing? Cluster Health Report

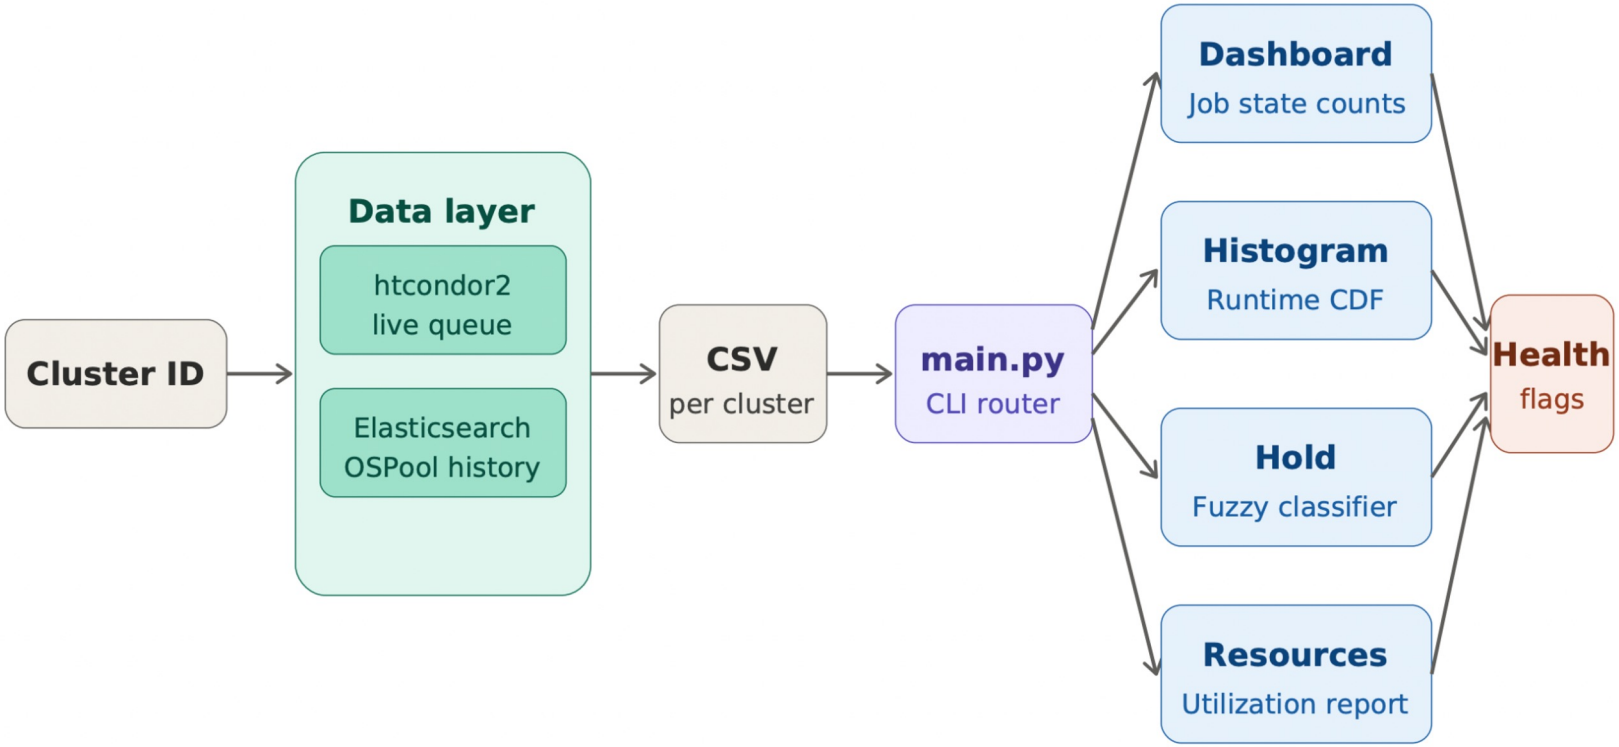
Cluster ID
Data layer
htcondor2
live queue
Elasticsearch
OSPool history
CSV
per cluster
main.py
CLI router
Dashboard
Job state counts
Histogram
Runtime CDF
Hold
Fuzzy classifier
Resources
Utilization report
Health
flags

Cluster Status Dashboard

Benefits:

- Visual representation of the status of their jobs
- Accumulates data about jobs from both queue and history

CHTC HTCondor Software Suite PATh

Runtime Visualization

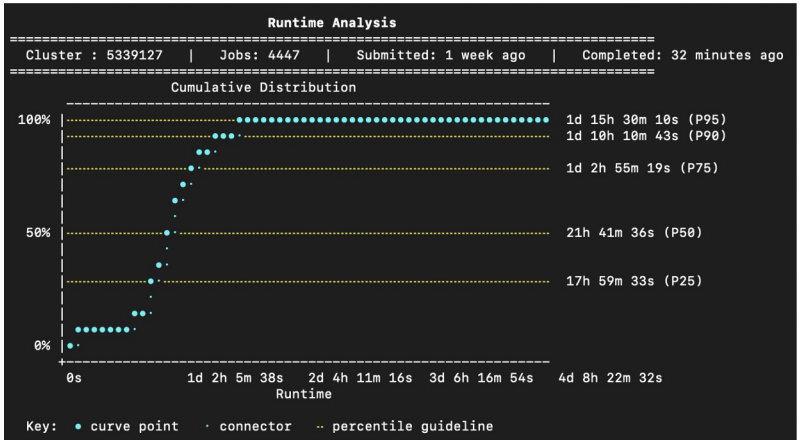

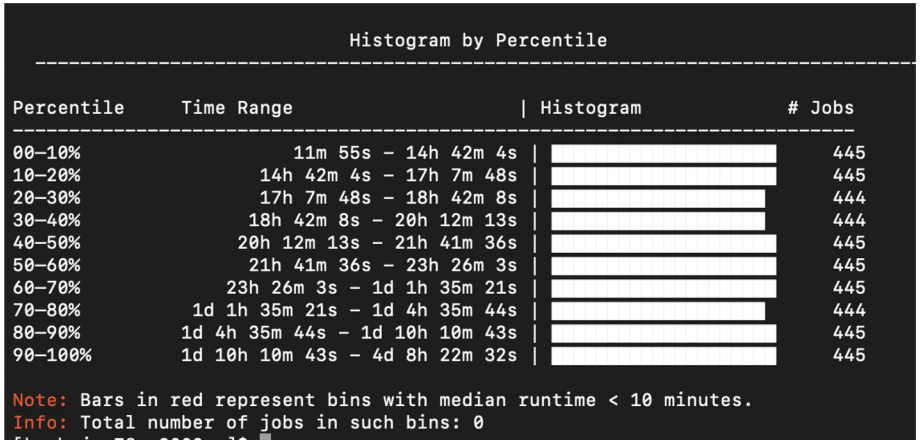

Benefits:

- Visual Feedback about runtimes of completed jobs
- Helps classify trends in workloads

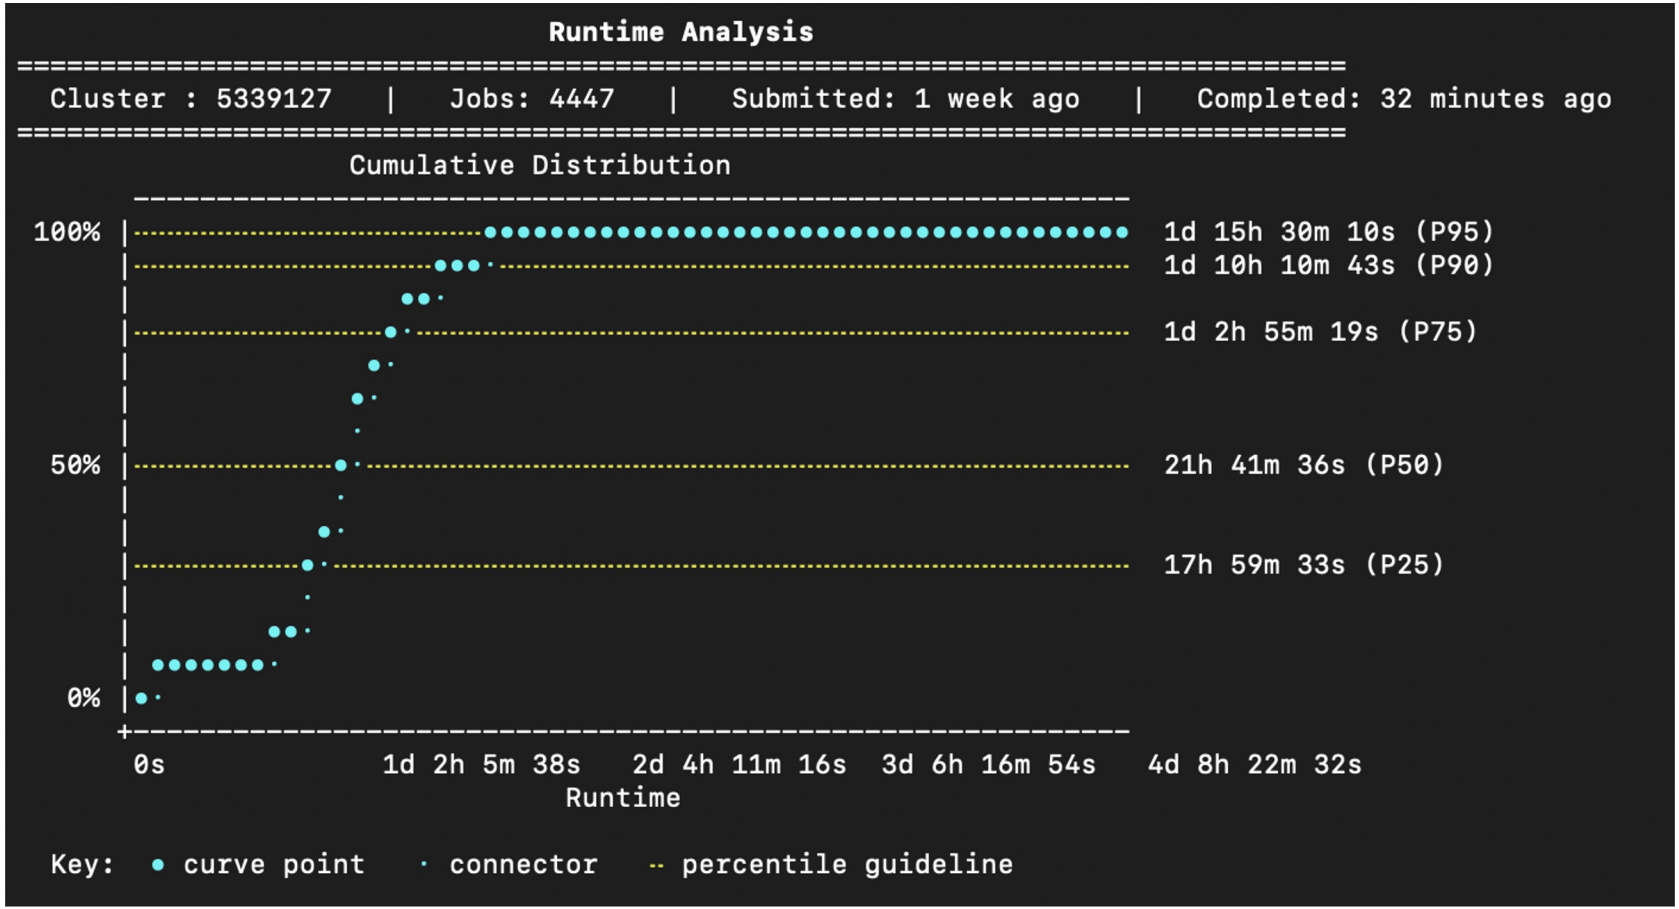
Runtime Analysis
Cluster : 5339127 | Jobs: 4447 | Submitted: 1 week ago | Completed: 32 minutes ago
Cumulative Distribution
100%
50%
0%
1d 15h 30m 10s (P95)
1d 10h 10m 43s (P90)
1d 2h 55m 19s (P75)
21h 41m 36s (P50)
17h 59m 33s (P25)
0s
1d 2h 5m 38s
2d 4h 11m 16s
3d 6h 16m 54s
4d 8h 22m 32s
Runtime
Key: curve point connector percentile guideline

Histogram by Percentile

Percentile	Time Range	Histogram	# Jobs
00–10%	11m 55s – 14h 42m 4s		445
10–20%	14h 42m 4s – 17h 7m 48s		445
20–30%	17h 7m 48s – 18h 42m 8s		444
30–40%	18h 42m 8s – 20h 12m 13s		444
40–50%	20h 12m 13s – 21h 41m 36s		445
50–60%	21h 41m 36s – 23h 26m 3s		445
60–70%	23h 26m 3s – 1d 1h 35m 21s		445
70–80%	1d 1h 35m 21s – 1d 4h 35m 44s		444
80–90%	1d 4h 35m 44s – 1d 10h 10m 43s		445
90–100%	1d 10h 10m 43s – 4d 8h 22m 32s		445

Note: Bars in red represent bins with median runtime < 10 minutes.
Info: Total number of jobs in such bins: 0

CHTC HTCondor Software Suite PATh

```
Histogram of Job Runtimes by Percentiles:

Percentile          Time Range                    | Histogram              # Jobs
-------------------------------------------------------------------------------
00-10%              6s          - 7s              |                        32
10-20%              7s          - 8s              |                        104
20-30%              8s          - 1m 45s          |                        88
30-40%              1m 45s      - 3m 33s          |                        75
40-50%              3m 33s      - 12m 14s         |                        75
50-60%              12m 14s     - 20m 57s         |                        74
60-70%              20m 57s     - 33m 8s          |                        75
70-80%              33m 8s      - 1h 4m 1s        |                        75
80-90%              1h 4m 1s    - 2h 32m 17s      |                        75
90-100%             2h 32m 17s  - 14h 10m 28s     |                        75

Note: Bars in red represent bins with median runtime < 10 minutes.
Info: Total number of jobs in such bins: 374
```

Benefits:

- Binned on the basis of percentiles
- Points out how many jobs have a runtime <10 mins, along with the option to print their JobIds
- Can be used to understand the range of runtimes for a cluster

Hold Classifier

```
Fetching held jobs from cluster...
Found 1925 held jobs

Cluster ID: 5339127
Held Jobs in Cluster: 1925

Time Analysis:
  First held: 2026-06-04 17:10:47
  Last held:  2026-06-10 13:33:24
  Duration:   140.4 hours
  Avg hold:   3d 6h

+----------------------+-----------+---------------------------+-----------------+----------------------------------------------------------------------------------------------+
| Hold Reason Label    |   SubCode | % of Held Jobs (Count)    | Avg Hold Time   | Example Reason                                                                               |
+======================+===========+===========================+=================+==============================================================================================+
| JobOutOfResources    |       102 | 82.0% (1578)              | 3d 6h           | Job has gone over cgroup memory limit of 4096 megabytes                                      |
+----------------------+-----------+---------------------------+-----------------+----------------------------------------------------------------------------------------------+
| TransferOutputError  |         2 | 17.1% (329)               | 3d 9h           | Transfer output files failure at execution point slot1_61@e2602.chtc.wisc.edu while sending files to |
|                      |           |                           |                 | access point ap2002                                                                          |
+----------------------+-----------+---------------------------+-----------------+----------------------------------------------------------------------------------------------+
| StartdHeldJob        |         0 | 0.9% (18)                 | 2d 0h           | Job failed to complete in 72 hrs                                                             |
+----------------------+-----------+---------------------------+-----------------+----------------------------------------------------------------------------------------------+

Legend:
+--------+----------------------+--------------------------------------------------------------------------+
|   Code | Label                | Reason                                                                   |
+========+======================+==========================================================================+
|     12 | TransferOutputError  | An error occurred while transferring job output files or self-checkpoint files. |
+--------+----------------------+--------------------------------------------------------------------------+
|     21 | StartdHeldJob        | The job was put on hold because WANT_HOLD in the machine policy was true.  |
+--------+----------------------+--------------------------------------------------------------------------+
|     34 | JobOutOfResources    | Memory usage exceeds a memory limit.                                     |
+--------+----------------------+--------------------------------------------------------------------------+
```

Gives a tabular summary for the held jobs, with the Hold Reason Label, Subcode, #Jobs, and an Example error message for each bucket

```
Cluster ID: 5339127
Held Jobs in Cluster: 1925

Time Analysis:
  First held: 2026-06-04 17:10:47
  Last held:  2026-06-10 13:33:24
  Duration:   140.4 hours
  Avg hold:   3d 6h
```

Hold Reason Label	SubCode	% of Held Jobs (Count)	Avg Hold Time	Example Reason
JobOutOfResources	102	82.0% (1578)	3d 6h	Job has gone over cgroup memory limit of 4096 megabytes
TransferOutputError	2	17.1% (329)	3d 9h	Transfer output files failure at execution point slot1_61@e2602.chtc.wisc.edu while sending files to access point ap2002
StartdHeldJob	0	0.9% (18)	2d 0h	Job failed to complete in 72 hrs

Legend:

Code	Label	Reason
12	TransferOutputError	An error occurred while transferring job output files or self-checkpoint files.
21	StartdHeldJob	The job was put on hold because WANT_HOLD in the machine policy was true.
34	JobOutOfResources	Memory usage exceeds a memory limit.

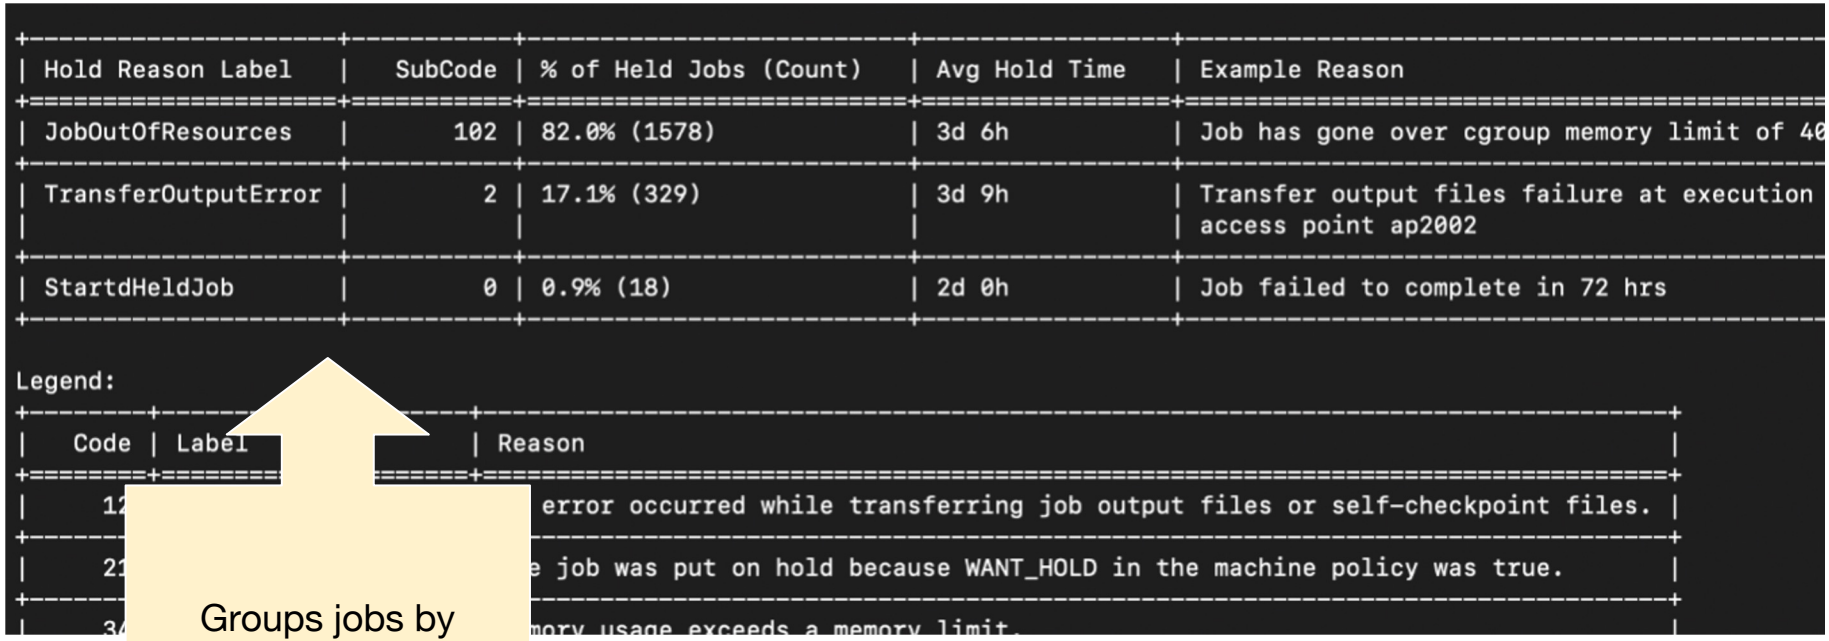

Hold Reason Label	SubCode	% of Held Jobs (Count)	Avg Hold Time	Example Reason
JobOutOfResources	102	82.0% (1578)	3d 6h	Job has gone over cgroup memory limit of 40
TransferOutputError	2	17.1% (329)	3d 9h	Transfer output files failure at execution access point ap2002
StartdHeldJob	0	0.9% (18)	2d 0h	Job failed to complete in 72 hrs

Legend:

Code	Label	Reason
12		error occurred while transferring job output files or self-checkpoint files.
21		e job was put on hold because WANT_HOLD in the machine policy was true.
34		mory usage exceeds a memory limit.

Groups jobs by HoldReasonCode and SubCode

..

d/TotalJobs	ExampleReason
420	Error from slot1_3@herringa2001.chtc.wisc.edu: Job failed to complete in 72 hrs
20	Error from slot1_29@herringa2001.chtc.wisc.edu: memory usage exceeded request_memory
420	Error from slot1_80@herringa2001.chtc.wisc.edu: Job has gone over cgroup memory limit of 5120 megaby
20	The job exceeded allowed execute duration of 3+00:00:00

and also HoldReason (by string comparison)

CHTC HTCondor Software Suite PATh

```
Legend:
+--------+---------------------+-----------------------------------------------------------------------------------+
|   Code | Label               | Reason                                                                            |
+========+=====================+===================================================================================+
|     12 | TransferOutputError | An error occurred while transferring job output files or self-checkpoint files.  |
+--------+---------------------+-----------------------------------------------------------------------------------+
|     21 | StartdHeldJob       | The job was put on hold because WANT_HOLD in the machine policy was true.         |
+--------+---------------------+-----------------------------------------------------------------------------------+
|     34 | JobOutOfResources   | Memory usage exceeds a memory limit.                                              |
+--------+---------------------+-----------------------------------------------------------------------------------+
```

Gives Legend of HoldReason, Label and the reason

```
Cluster ID: 5339127
Held Jobs in Cluster: 1925

Time Analysis:
  First held: 2026-06-04 17:10:47
  Last held:  2026-06-10 13:33:24
  Duration:   140.4 hours
  Avg hold:   3d 6h
```

Hold Reason Label	SubCode	% of Held Jobs (Count)	Avg Hold Time	Example Reason
JobOutOfResources	102	82.0% (1578)	3d 6h	Job has gone over cgroup memory limit of 4096 megabytes
TransferOutputError	2	17.1% (329)	3d 9h	Transfer output files failure at execution point slot1_61@e2602.chtc.wisc.edu while sending files to access point ap2002
StartdHeldJob	0	0.9% (18)	2d 0h	Job failed to complete in 72 hrs

Legend:

Code	Label	Reason
12	TransferOutputError	An error occurred while transferring job output files or self-checkpoint files.
21	StartdHeldJob	The job was put on hold because WANT_HOLD in the machine policy was true.
34	JobOutOfResources	Memory usage exceeds a memory limit.

Benefits:

- Easy to understand summary with example error reasons for better problem solving
- Adaptive to changes in the system with information being pulled from a table

Resource Utilization Report

Gives a comprehensive report about resource request and utilisation for a cluster

- Statistics
- Usage Distribution
- Optimization Recommendations
- Efficiency Summary

Benefits:

- Aggregated data from all jobs
- Better resource utilization

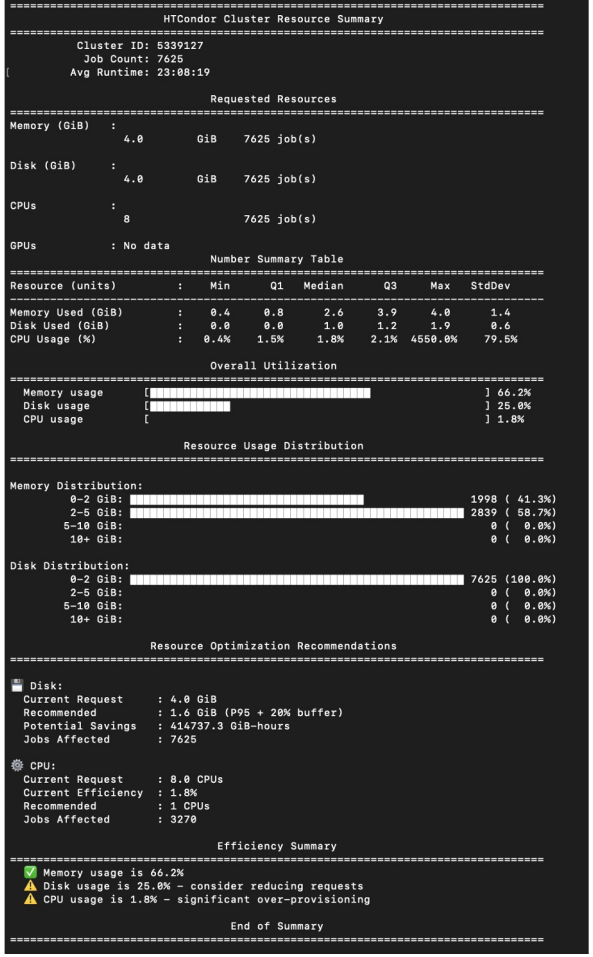

```
================================================================================
                        HTCondor Cluster Resource Summary
================================================================================
          Cluster ID: 5339127
           Job Count: 7625
         Avg Runtime: 23:08:19

                                Requested Resources
================================================================================
Memory (GiB)    :
                  4.0         GiB     7625 job(s)

Disk (GiB)      :
                  4.0         GiB     7625 job(s)

CPUs            :
                  8                   7625 job(s)
```

A table summarizing the resource request and the count of jobs

Number Summary Table

Resource (units)	Min	Q1	Median	Q3	Max	StdDev
Memory Used (GiB)	0.4	0.8	2.6	3.9	4.0	1.4
Disk Used (GiB)	0.0	0.0	1.0	1.2	1.9	0.6
CPU Usage (%)	0.4%	1.5%	1.8%	2.1%	4550.0%	79.5%

A table summarizing the resource utilization and statistics such as quantiles, and StdDev

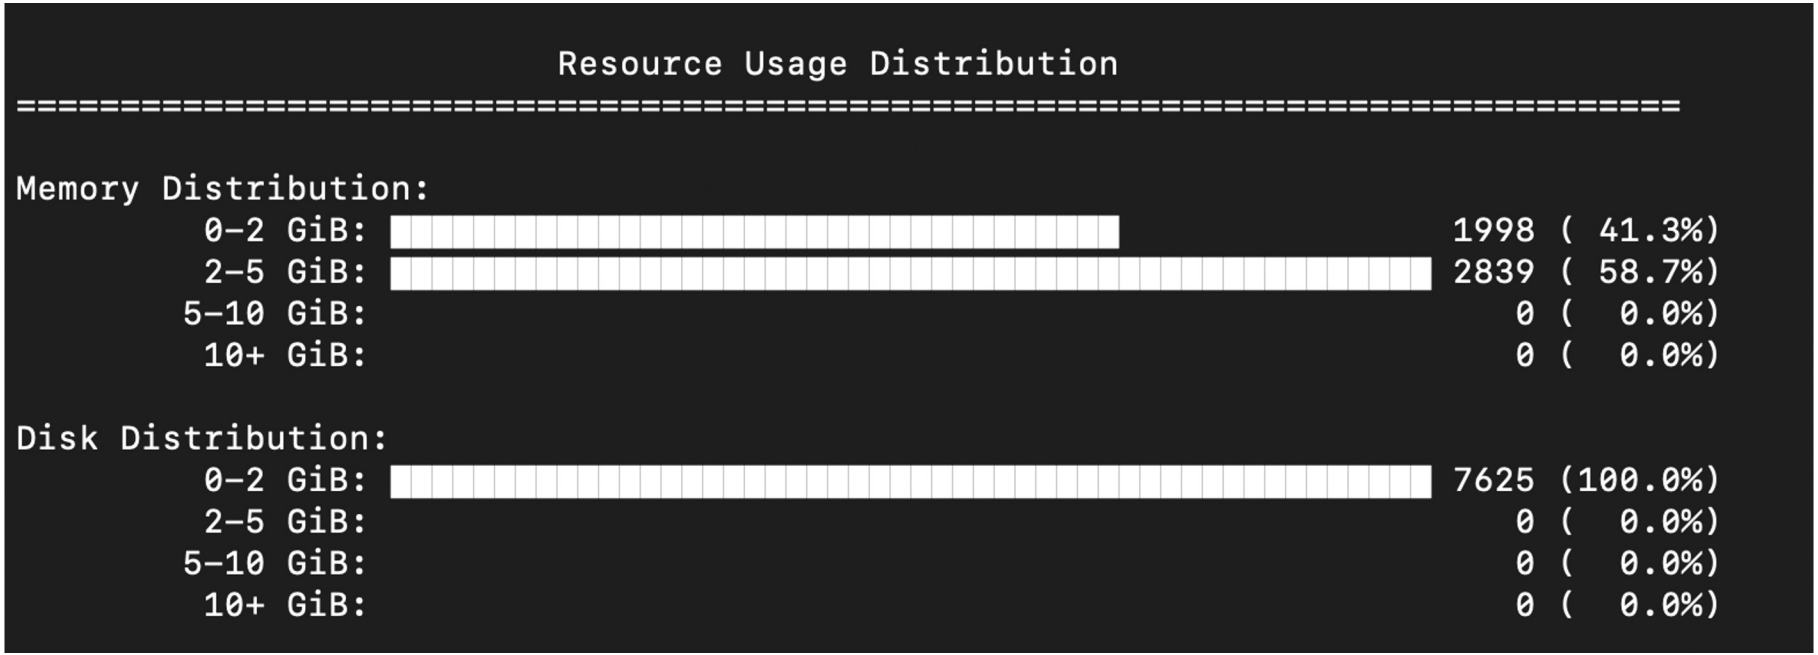

A visual representation of overall utilization

CHTC HTCondor Software Suite PATh

```
                    Resource Optimization Recommendations
================================================================================

💾 Disk:
  Current Request     : 4.0 GiB
  Recommended         : 1.6 GiB (P95 + 20% buffer)
  Potential Savings   : 414737.3 GiB-hours
  Jobs Affected       : 7625

⚙️ CPU:
  Current Request     : 8.0 CPUs
  Current Efficiency  : 1.8%
  Recommended         : 1 CPUs
  Jobs Affected       : 3270

                              Efficiency Summary
================================================================================
  ✅ Memory usage is 66.2%
  ⚠️ Disk usage is 25.0% - consider reducing requests
  ⚠️ CPU usage is 1.8% - significant over-provisioning
```

Summary and Recommendations

CHTC HTCondor Software Suite PATh

Cluster Health Report

```
================================================================================
                    HTCondor Cluster Health Report
================================================================================
Cluster ID: 5339127
Report Time: 2026-06-10 13:32:38
================================================================================

Overall Status: 🔴 CRITICAL
Critical Issues: 3 | Warnings: 1

+--------------------+------------+---------+-----------------------------------------+-------------------------------+
| Aspect             | Status     | Value   | Reason / Details                        | Tool for Details              |
+====================+============+=========+=========================================+===============================+
| Memory Efficiency  | 🟡 WARNING | 66.2%   | Low efficiency (66.2%)                  | python analytics.py 5339127   |
+--------------------+------------+---------+-----------------------------------------+-------------------------------+
| Disk Efficiency    | 🔴 CRITICAL| 25.0%   | Severe over-provisioning (25.0% efficiency) | python analytics.py 5339127 |
+--------------------+------------+---------+-----------------------------------------+-------------------------------+
| CPU Efficiency     | 🔴 CRITICAL| 1.8%    | Severe over-provisioning (1.8% efficiency)  | python analytics.py 5339127 |
+--------------------+------------+---------+-----------------------------------------+-------------------------------+
| Held Jobs          | 🔴 CRITICAL| 1924    | 1924 jobs held (25.2%), 3 unique reasons    | python hold_bucket.py 5339127 |
+--------------------+------------+---------+-----------------------------------------+-------------------------------+
| Fast Jobs          | 🟢 HEALTHY | 0.0%    | 1 jobs < 10min (0.0%)                   | python histogram.py 5339127   |
+--------------------+------------+---------+-----------------------------------------+-------------------------------+
| Runtime Consistency| 🟢 HEALTHY | CV=0.43 | Consistent runtimes (CV=0.43)           | python histogram.py 5339127   |
+--------------------+------------+---------+-----------------------------------------+-------------------------------+
| Job Status         | ℹ️ INFO    | 7624    | 1341 completed, 1571 running, 2788 idle | python dashboard.py 5339127   |
+--------------------+------------+---------+-----------------------------------------+-------------------------------+

================================================================================
Recommended Next Steps
================================================================================

🔴 HIGH PRIORITY:
  • Disk Efficiency: python analytics.py 5339127
  • CPU Efficiency: python analytics.py 5339127
  • Held Jobs: python hold_bucket.py 5339127

🟡 REVIEW:
  • Memory Efficiency: python analytics.py 5339127
```

CHTC HTCondor Software Suite PATh

```
================================================================================
                        HTCondor Cluster Health Report
================================================================================
Cluster ID: 5339127
Report Time: 2026-06-10 13:32:38
================================================================================

Overall Status: 🔴 CRITICAL
Critical Issues: 3 | Warnings: 1
```

Aspect	Status	Value	Reason / Details	Tool for Details
Memory Efficiency	🟡 WARNING	66.2%	Low efficiency (66.2%)	python analytics.py 5339127
Disk Efficiency	🔴 CRITICAL	25.0%	Severe over-provisioning (25.0% efficiency)	python analytics.py 5339127
CPU Efficiency	🔴 CRITICAL	1.8%	Severe over-provisioning (1.8% efficiency)	python analytics.py 5339127
Held Jobs	🔴 CRITICAL	1924	1924 jobs held (25.2%), 3 unique reasons	python hold_bucket.py 5339127
Fast Jobs	🟢 HEALTHY	0.0%	1 jobs < 10min (0.0%)	python histogram.py 5339127
Runtime Consistency	🟢 HEALTHY	CV=0.43	Consistent runtimes (CV=0.43)	python histogram.py 5339127
Job Status	ℹ️ INFO	7624	1341 completed, 1571 running, 2788 idle	python dashboard.py 5339127

Next Steps

- Make these tools ready for the researchers to use within HTCondor

- Improve the tool on any feedback or issues we receive

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant No. 2030508. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

CHTC HTCondor Software Suite PATh