HST/JWST Data Processing Performance under HTCondor/OWL

Michael Swam, Mary Romelfanger

mswam@stsci.edu, mary@stsci.edu

SPACE TELESCOPE SCIENCE INSTITUTE
3700 San Martin Drive
Baltimore, MD 21218

Abstract:

The workflow management system now used by the science data processing pipelines for the Hubble Space Telescope (HST) and the James Webb Space Telescope (JWST) is called HTCondor/OWL, and consists of the widely-used HTCondor batch processing software, and the Open Workflow Layer (OWL), developed at the Space Telescope Science Institute (STScI). This poster describes early performance results obtained under the HTCondor/OWL system when reprocessing large subsets of the HST archive collection, and JWST ground test data. Description of the computer hardware configurations and data, the performance metrics gathered, and analysis of the overheads and efficiencies for the system are presented, along with some future plans.

The Processing Model

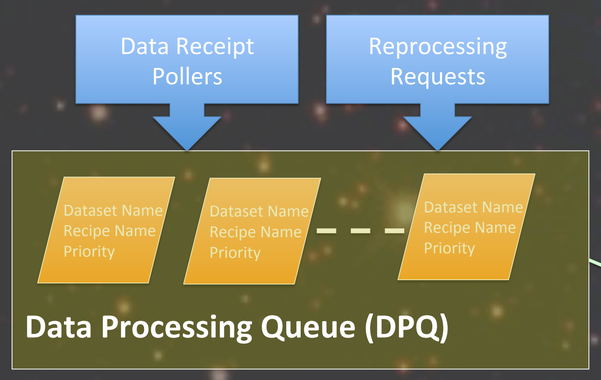

HTCondor workflow job step priorities are defined in OWL processing templates to ensure that higher priority workflows gain earlier access to CPUs

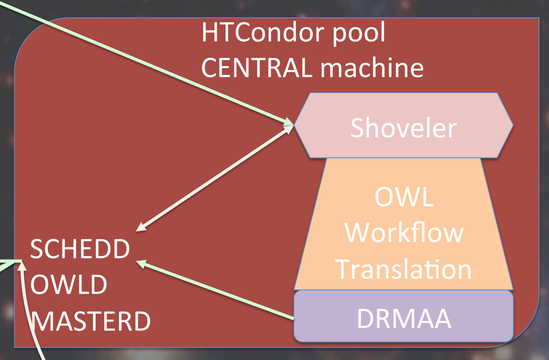

HST pool:
448 cpu cores
902 GB memory
6 Dell machines

JWST test pool:
64 cpu cores
376 GB memory
6 Vmware ESX Virtual machines

- **Data Receipt Pollers detect new downloads from the Observatory**
- **Reprocessing Requests are generated by operator tools that identify archived datasets that could be improved by recent calibration data and software improvements**
- **Initial download processing and reprocessing are run on the same hardware, using HTCondor priorities to control access to CPU cores**
- **OWL "Shoveler" task reads highest priority DPQ workflow requests for particular datasets to process using targeted processing algorithms (recipes); reprocessing is lower priority than new observatory downlinks**
- **OWL Shoveler queries HTCondor SCHEDD daemon to check pool capacity to accept new jobs (proceeds if queue space is available)**
- **OWL Shoveler calls OWL Workflow Translation layer to convert workflow request for a dataset into DAGMan workflow control files**
- **OWL submits workflow control files to DRMAA workflow plugin layer**
- **DRMAA translates workflow control files to HTCondor format and submits workflow to HTCondor**
- **HTCondor SCHEDD daemon, in consultation with the COLLECTOR and NEGOTIATOR daemons, chooses a STARTD worker machine to run the job and starts the job**
- **HTCondor MASTERD restarts any failed daemons**

Data Characteristics

HST (Active Instrument Data)

Instrument (primary mode)	Min/Max : Avg. Input Dataset size	Min/Max : Avg. Output Dataset size	Number of Archived Datasets (count all in a full reprocessing run)
ACS (imaging)	81 Kb / 68 Mb: 19 Mb	54 Kb/ 1054 Mb: 196 Mb	183505
COS (spectro)	36 Kb / 143 Mb: 4 Mb	28Kb/ 2435 Mb: 311 Mb	23505
STIS (spectro)	44 Kb / 237 Mb: 2 Mb	83Kb/ 3208 Mb: 27 Mb	124897
WFC3 (imaging)	23 Kb / 127 Mb: 18 Mb	64Kb/ 1768 Mb: 114 Mb	187545

JWST (Ground Test Data)

Instrument (primary mode)	Min/Max: Avg. Input Dataset size	Min/Max: Avg. Output Dataset size	Number of Test Datasets (SIC + DIL)
MIRI (both)	900 Kb/ 816 Mb: 65 Mb	1600 Kb/ 815 Mb: 51 Mb	452
NIRCAM (imaging)	410 Kb/ 907 Mb: 45 Mb	1054 Kb/ 176 Mb: 55 Mb	1089
NIRISS (both)	16 Kb/ 453 Mb : 71 Mb	43 Kb/ 453 Mb: 71 Mb	108
NIRSPEC (spectro)	7 Kb/ 739 Mb : 56 Mb	17 Mb/ 738 Mb: 91 Mb	702

Metrics and Conclusions

- **Since Oct.2015 all initial processing of HST datasets (51513 so far) on downlink from the observatory has occurred under HTCondor/OWL**
- **Major HST reprocessing runs for calibration data or software updates since HTCondor/OWL was installed in Dec.2014 are listed below**

Date	Instrument	Number of Datasets	Processing Time (hrs)
Jun.2015	COS	18135	50
Jul.2015	STIS	42047	100
Aug.2015	WFC3 (no CTE)	90959	133
Dec.2015	COS	16236	28
*** Processing pool expanded by 200 cores ***			
Jul.2016	COS	11232	44
Aug.2016	STIS	45666	41
Sep.2016	WFC3 (IR)	54040	94

- **JWST cryovac3 (CV3) Day In the Life (DIL) and Science Instrument Characterization (SIC) test suites were collected during observatory ground testing with all instruments live in Jan.2016**
  - **Data have been processed many times, using JWST test pool, as calibration data and software have improved**
  - **Aug.2016: all instruments, 1442 datasets processed in 5.0 hr**
  - **Oct.2016: all instruments, 1467 datasets processed in 6.7 hr**

- **Overheads computed from logs/database metrics collection**
  - **0.34 sec avg. spent in OWL Shoveler->HTCondor job submit**
  - **48 sec avg. spent in HTCondor job startup**
  - **Job start-up efficiency under HTCondor still warrants more tuning and study**
- **Continued HTCondor tuning is also being performed to optimize large HST background processing runs, but still allow the systems to respond well to higher-priority, on-demand processing**
- **The flexibility of easily adding machines to the HTCondor pool, as needed, for major reprocessing runs is a great benefit**
  - **No changes in processing software, just config files**
- **Many more JWST data processing test runs will be made over the next 2 years to complete the delivery of full pipeline capabilities , and tune the system for production runs during observatory operations starting in Oct. 2018**
  - **Additional production hardware purchases will be made soon to expand the JWST HTCondor pool for full observatory data processing**
  - **Sharing of data processing resources will be supported between HST and JWST, where one mission may have temporary, unused capacity that the other could borrow**