STScI | SPACE TELESCOPE SCIENCE INSTITUTE

EXPANDING THE FRONTIERS OF SPACE ASTRONOMY

Using HTCondor to Calibrate and Archive HST and JWST Data

Matthew Burger &

The STScI Data Processing & Archive Services (DPAS) Branch

05.23.2018

Our Vision

Our Mission

STScI | SPACE TELESCOPE SCIENCE INSTITUTE

expanding the frontiers of space astronomy

We help humanity explore the universe with advanced space telescopes and archives

Our Strategic Goals

Excel in the science operations of NASA's current and future astrophysics flagship missions

Advance state-of-the-art astronomical data, archives, and tools for scientific discovery

Make the world's astronomical information accessible to all

Science Operations for the Hubble Space Telescope (HST)

- ★ Proposal (observation) selection
- ★ Planning and scheduling
- ★ Data calibration
- ★ Data archives

HST on Twitter @spacetelelive

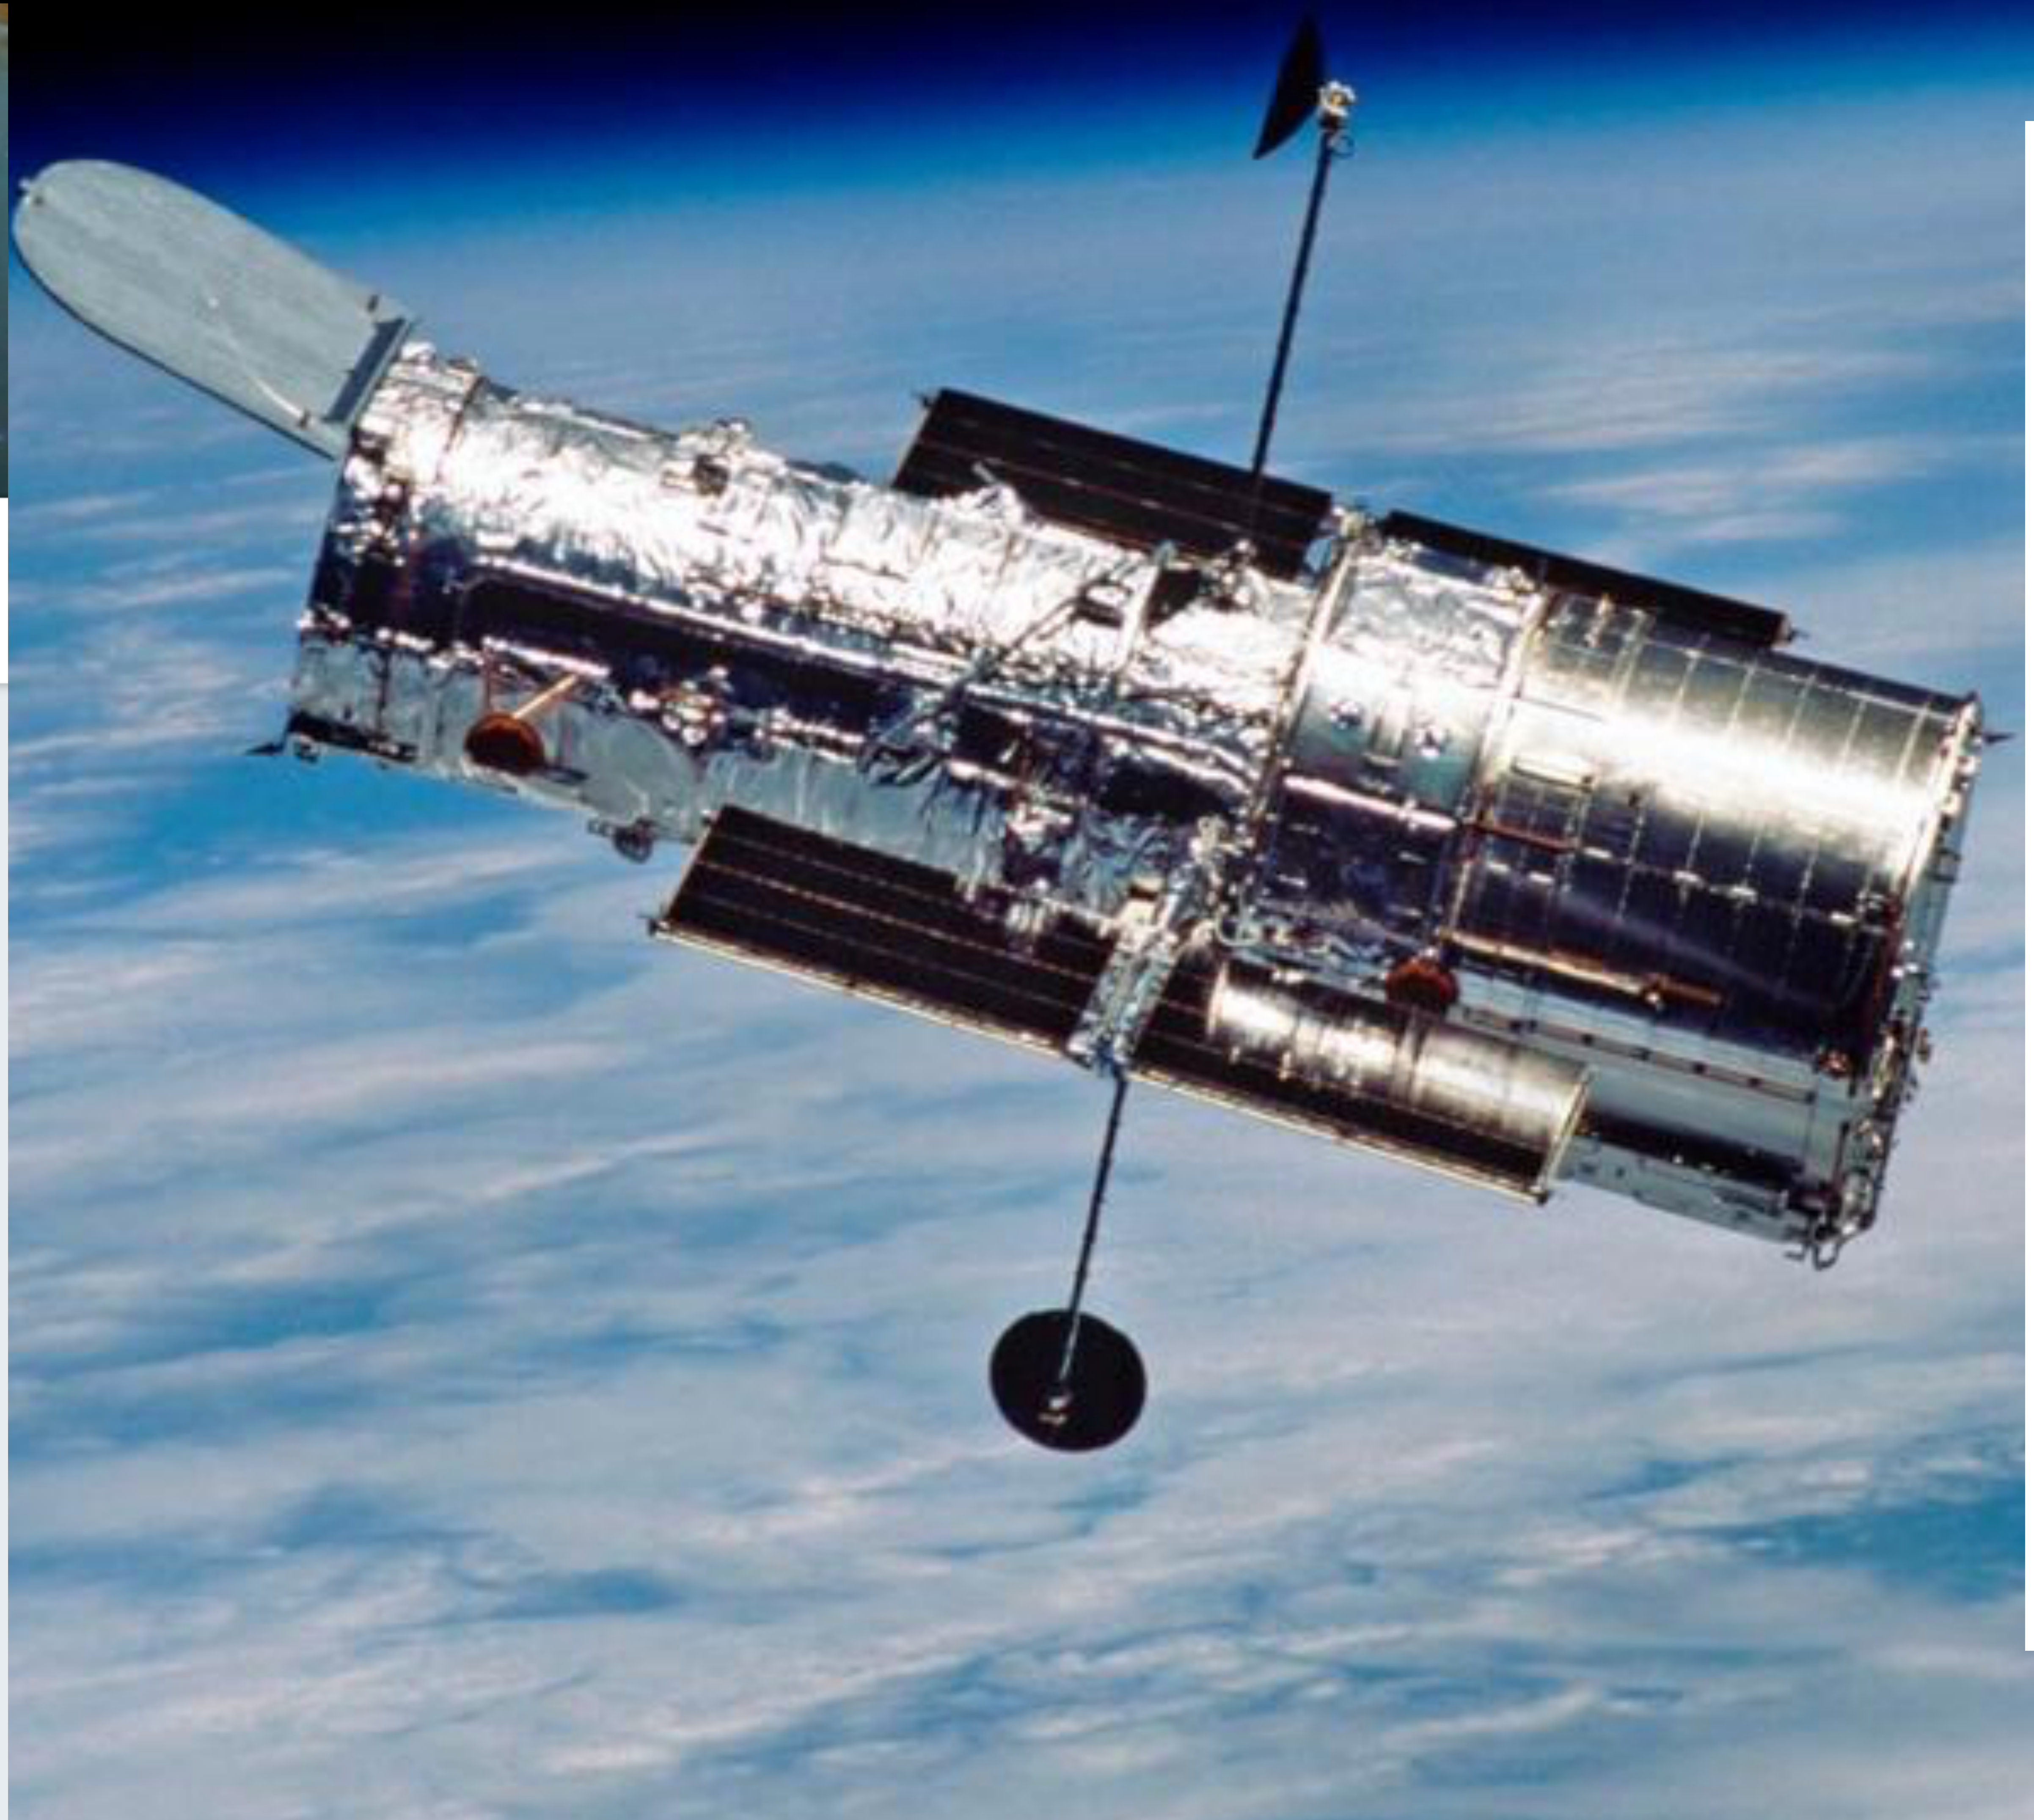

Space Telescope Live
@spacetelelive

What is @HubbleTelescope looking at? – Reference images from @STScI Digitized Sky Survey and Sloan Digital Sky Survey Data Release 13 (not actual Hubble images)

Low Earth Orbit
spacetelescopelive.org
Joined August 2016
Born on April 24, 1990

Tweet to Space Telescope Live

Space Telescope Live @spacetelelive Following

I am looking at the Galaxy NGC0613 with Wide Field Camera 3 for Dr. Peter Erwin. spacetelescopelive.org/2018-05-21T12:...

9:30 AM - 21 May 2018

6 Retweets 14 Likes

6 14

Spacecraft & Science Operations for the James Webb Space Telescope (JWST)
(Scheduled for launch in Spring 2020)

- ⭐Proposal (observation) selection
- ⭐Planning and scheduling
- ⭐Data calibration
- ⭐Data archives
- ⭐Mission Operations Center
- ⭐Wavefront Sensing and Control

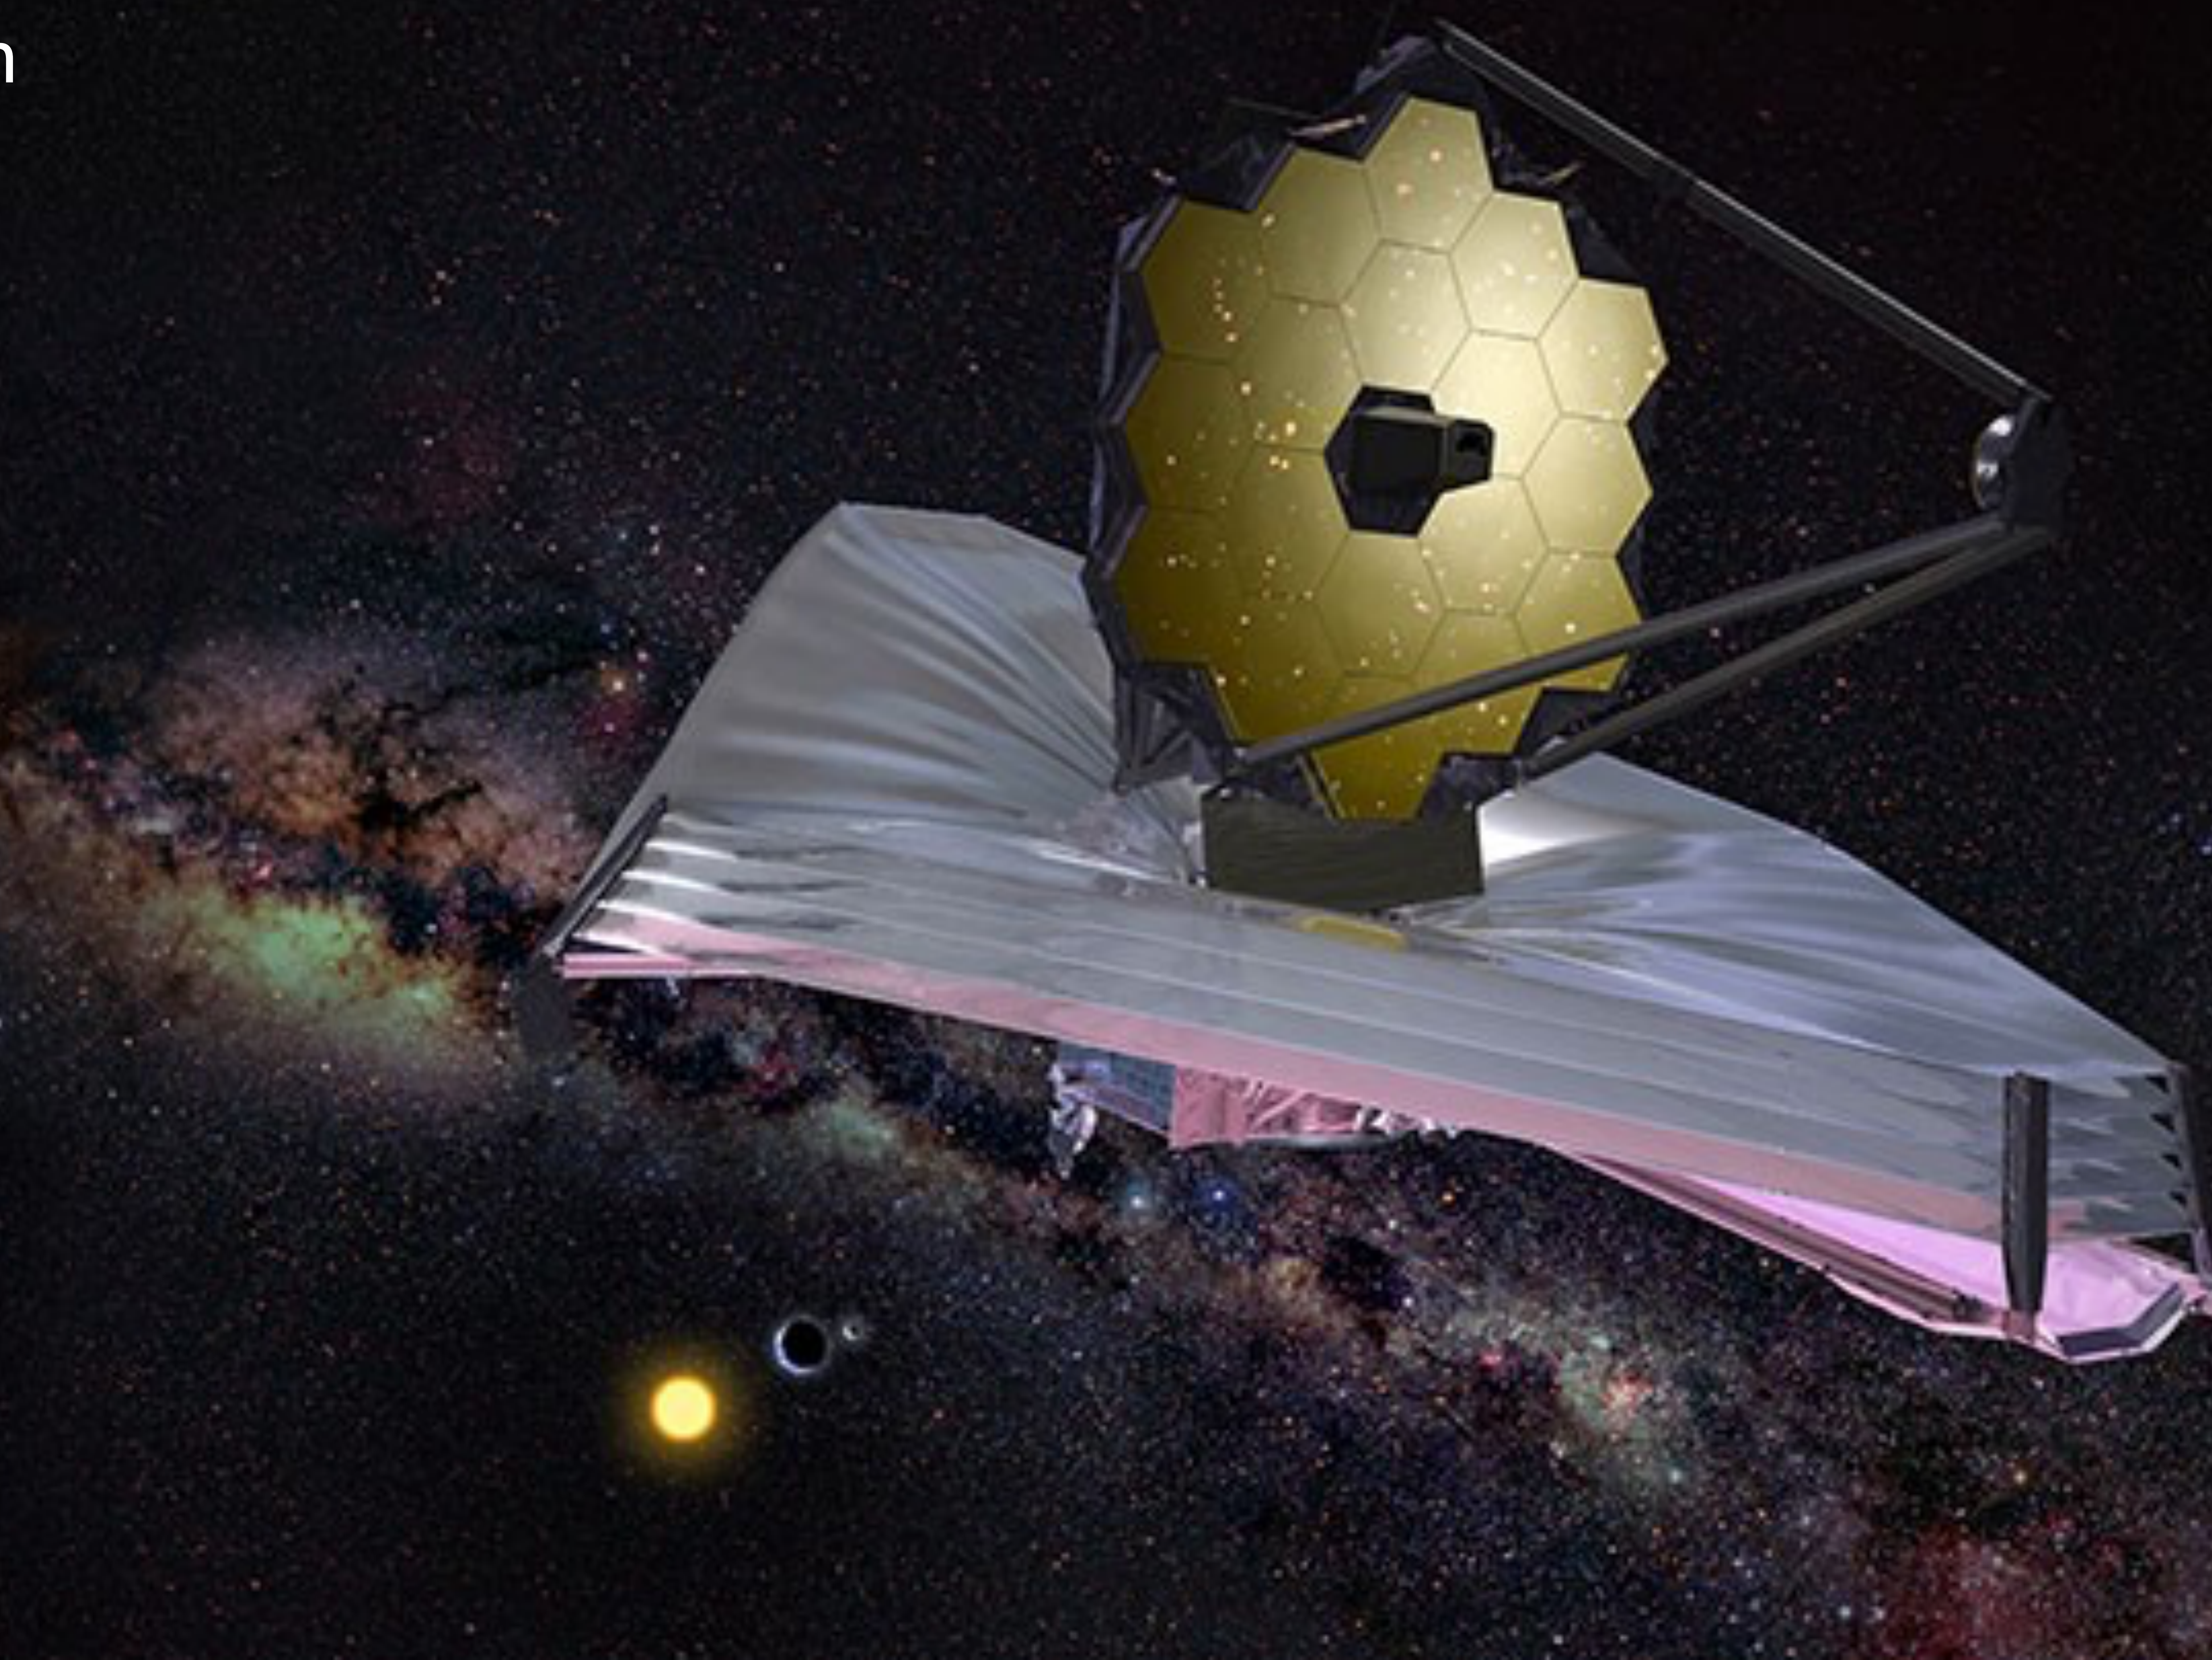

http://mast.stsci.edu/

- MAST Observations: Millions of observations from Hubble, Kepler, TESS, GALEX, IUE, FUSE, and more (to include JWST and WFIRST).
- Virtual Observatory: Search thousands of astronomical data archives from around the world for images, spectra, and catalogs.
- Hubble Source Catalog: A master catalog with a hundred million measurements of objects in Hubble images.
- MAST Catalogs: Access to catalog data such as Gaia and TESS Input Catalog, with more coming soon.
- All data is publicly available except for most HST data < 6 months old and JWST data < 12 months old

Data Management System (DMS)

- The data management system controls the flow of science and engineering data through the data calibration pipelines and into the archive (MAST).
- Multiple data pipelines (workflows):
  - New science from the telescope (four active instruments on HST)
  - Reprocessed science data due to changes in software or calibration data
  - Engineering data including spacecraft meta-data and jitter
  - Mission schedules (planned observations)
  - Spacecraft ephemerides

Data Calibration

GOAL: Convert an image from raw pixel counts on the detector to a source flux (number of photons emitted).

To solve this problem, one needs to account for:

- Detector bias: counts in a detector readout when no exposure is taken
- Dark current: spurious counts during an exposure due to noise in the detector
- Flat-field: pixel-to-pixel variations in detector efficiency
- CTE correction: pixel-to-pixel charge-transfer efficiency
- Combined/drizzled exposures: Long observations are broken into multiple exposures and combined in processing to increase the signal-to-noise ratio

Raw Data

Advance Camera for Surveys (ACS) image of M101 (Pinwheel Galaxy)

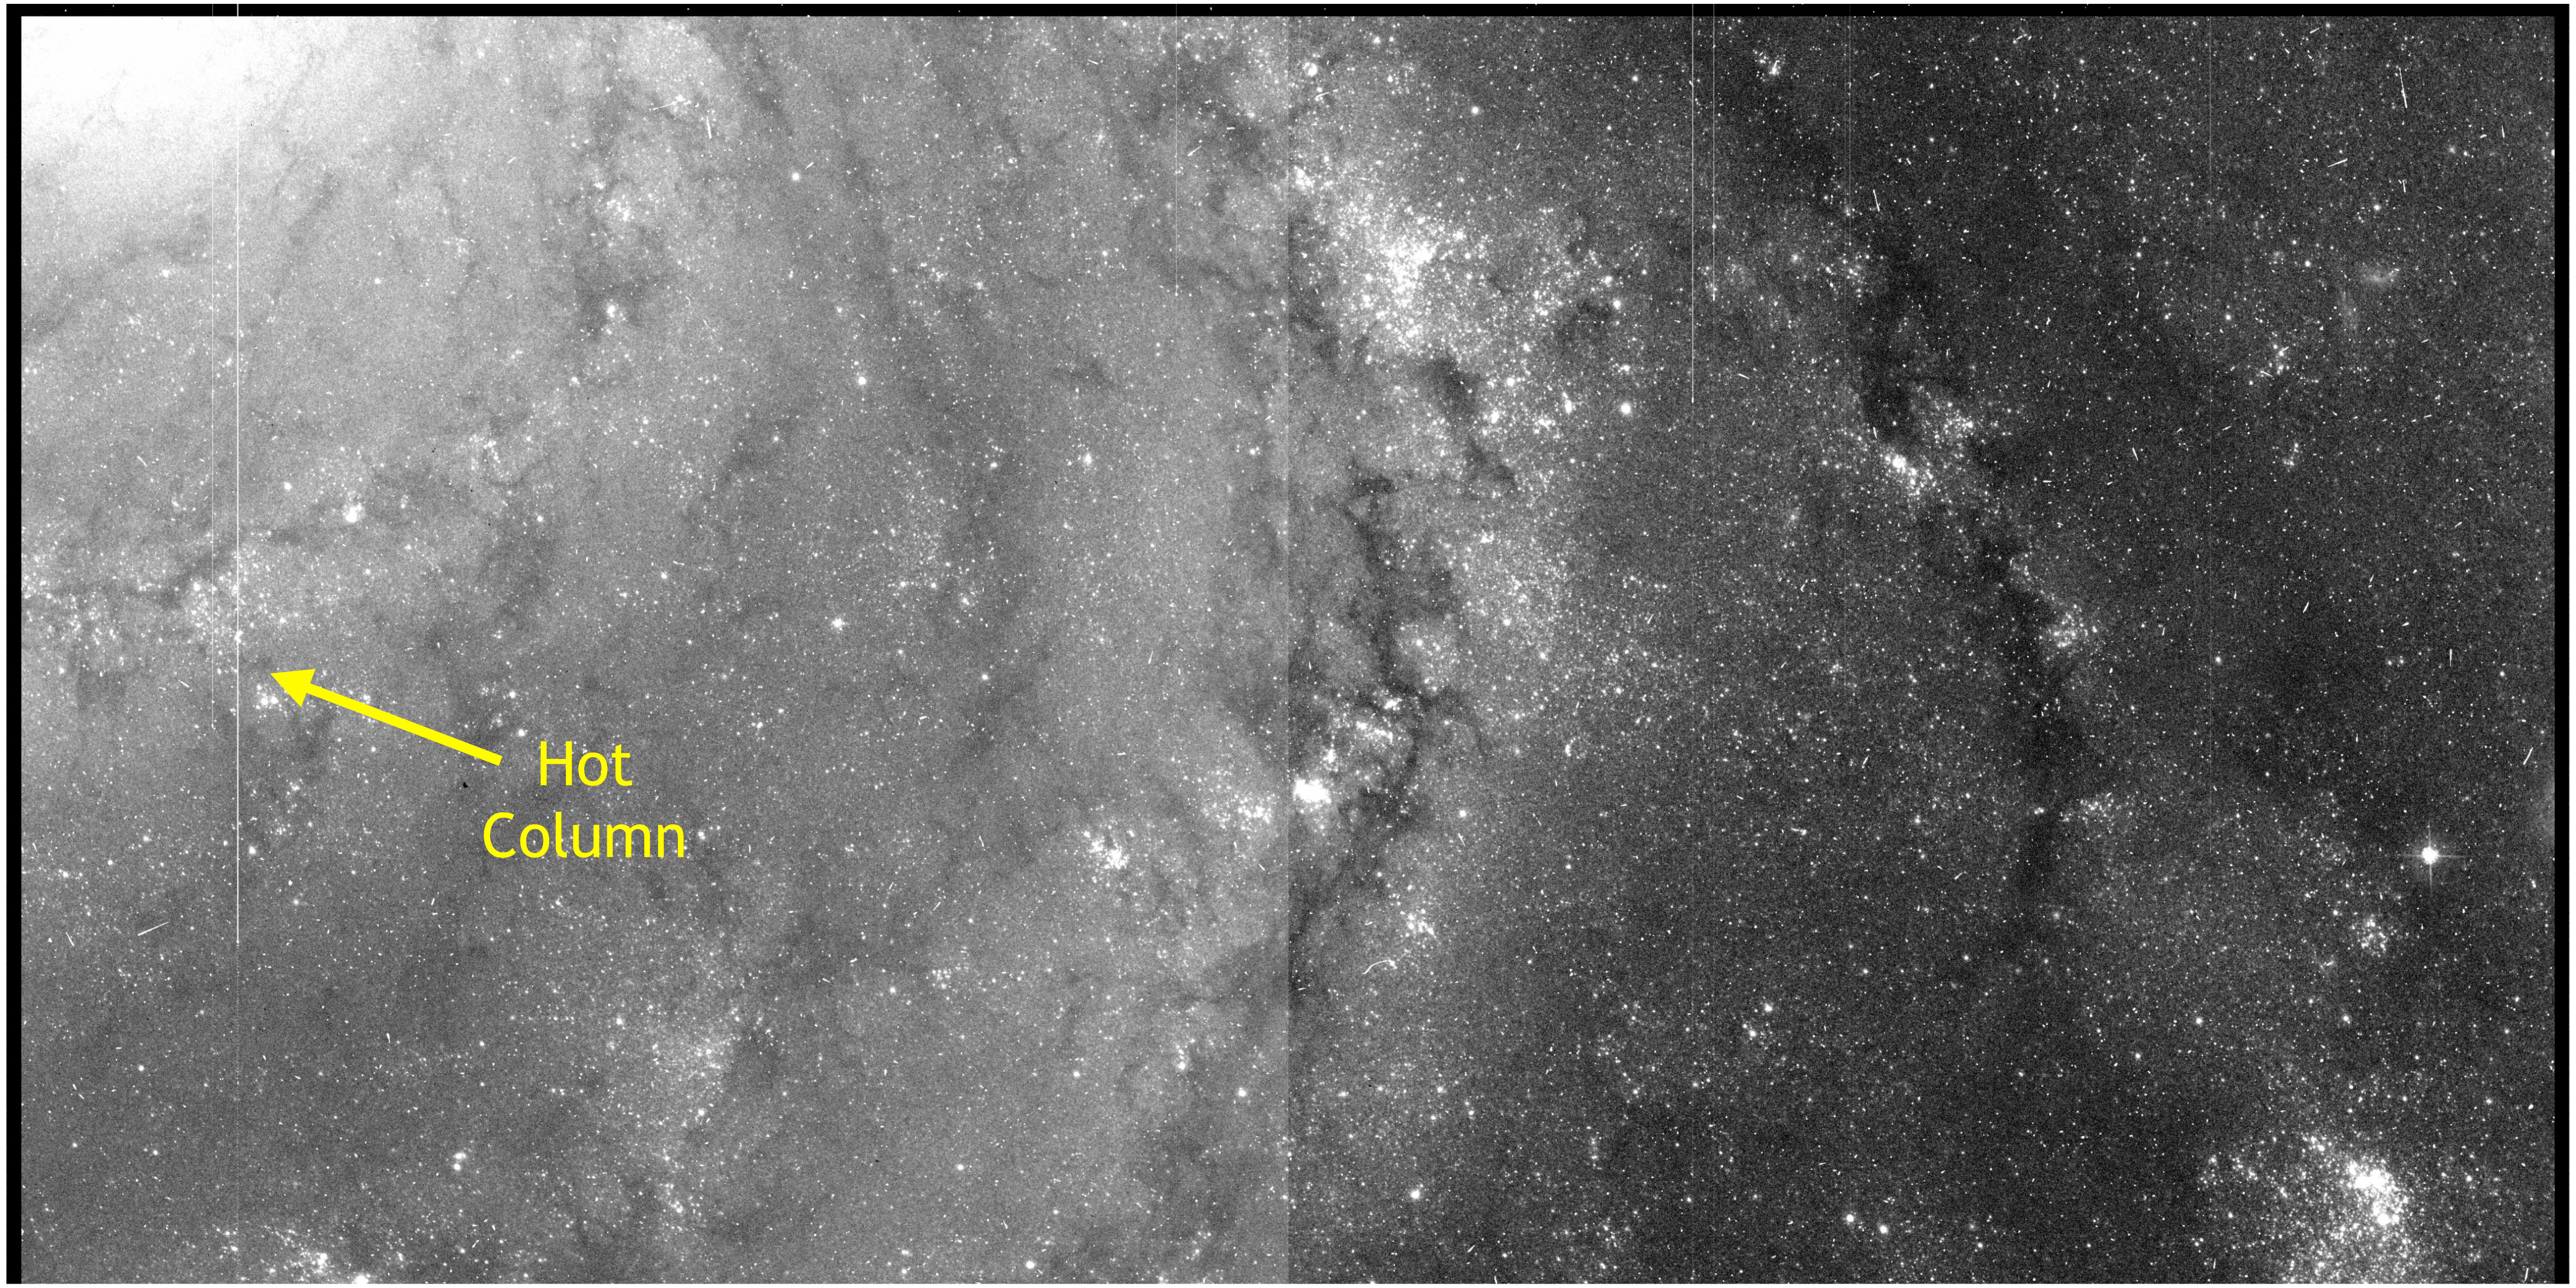

Two amplifiers with different bias levels

Bias Corrected & Flat-Fielded

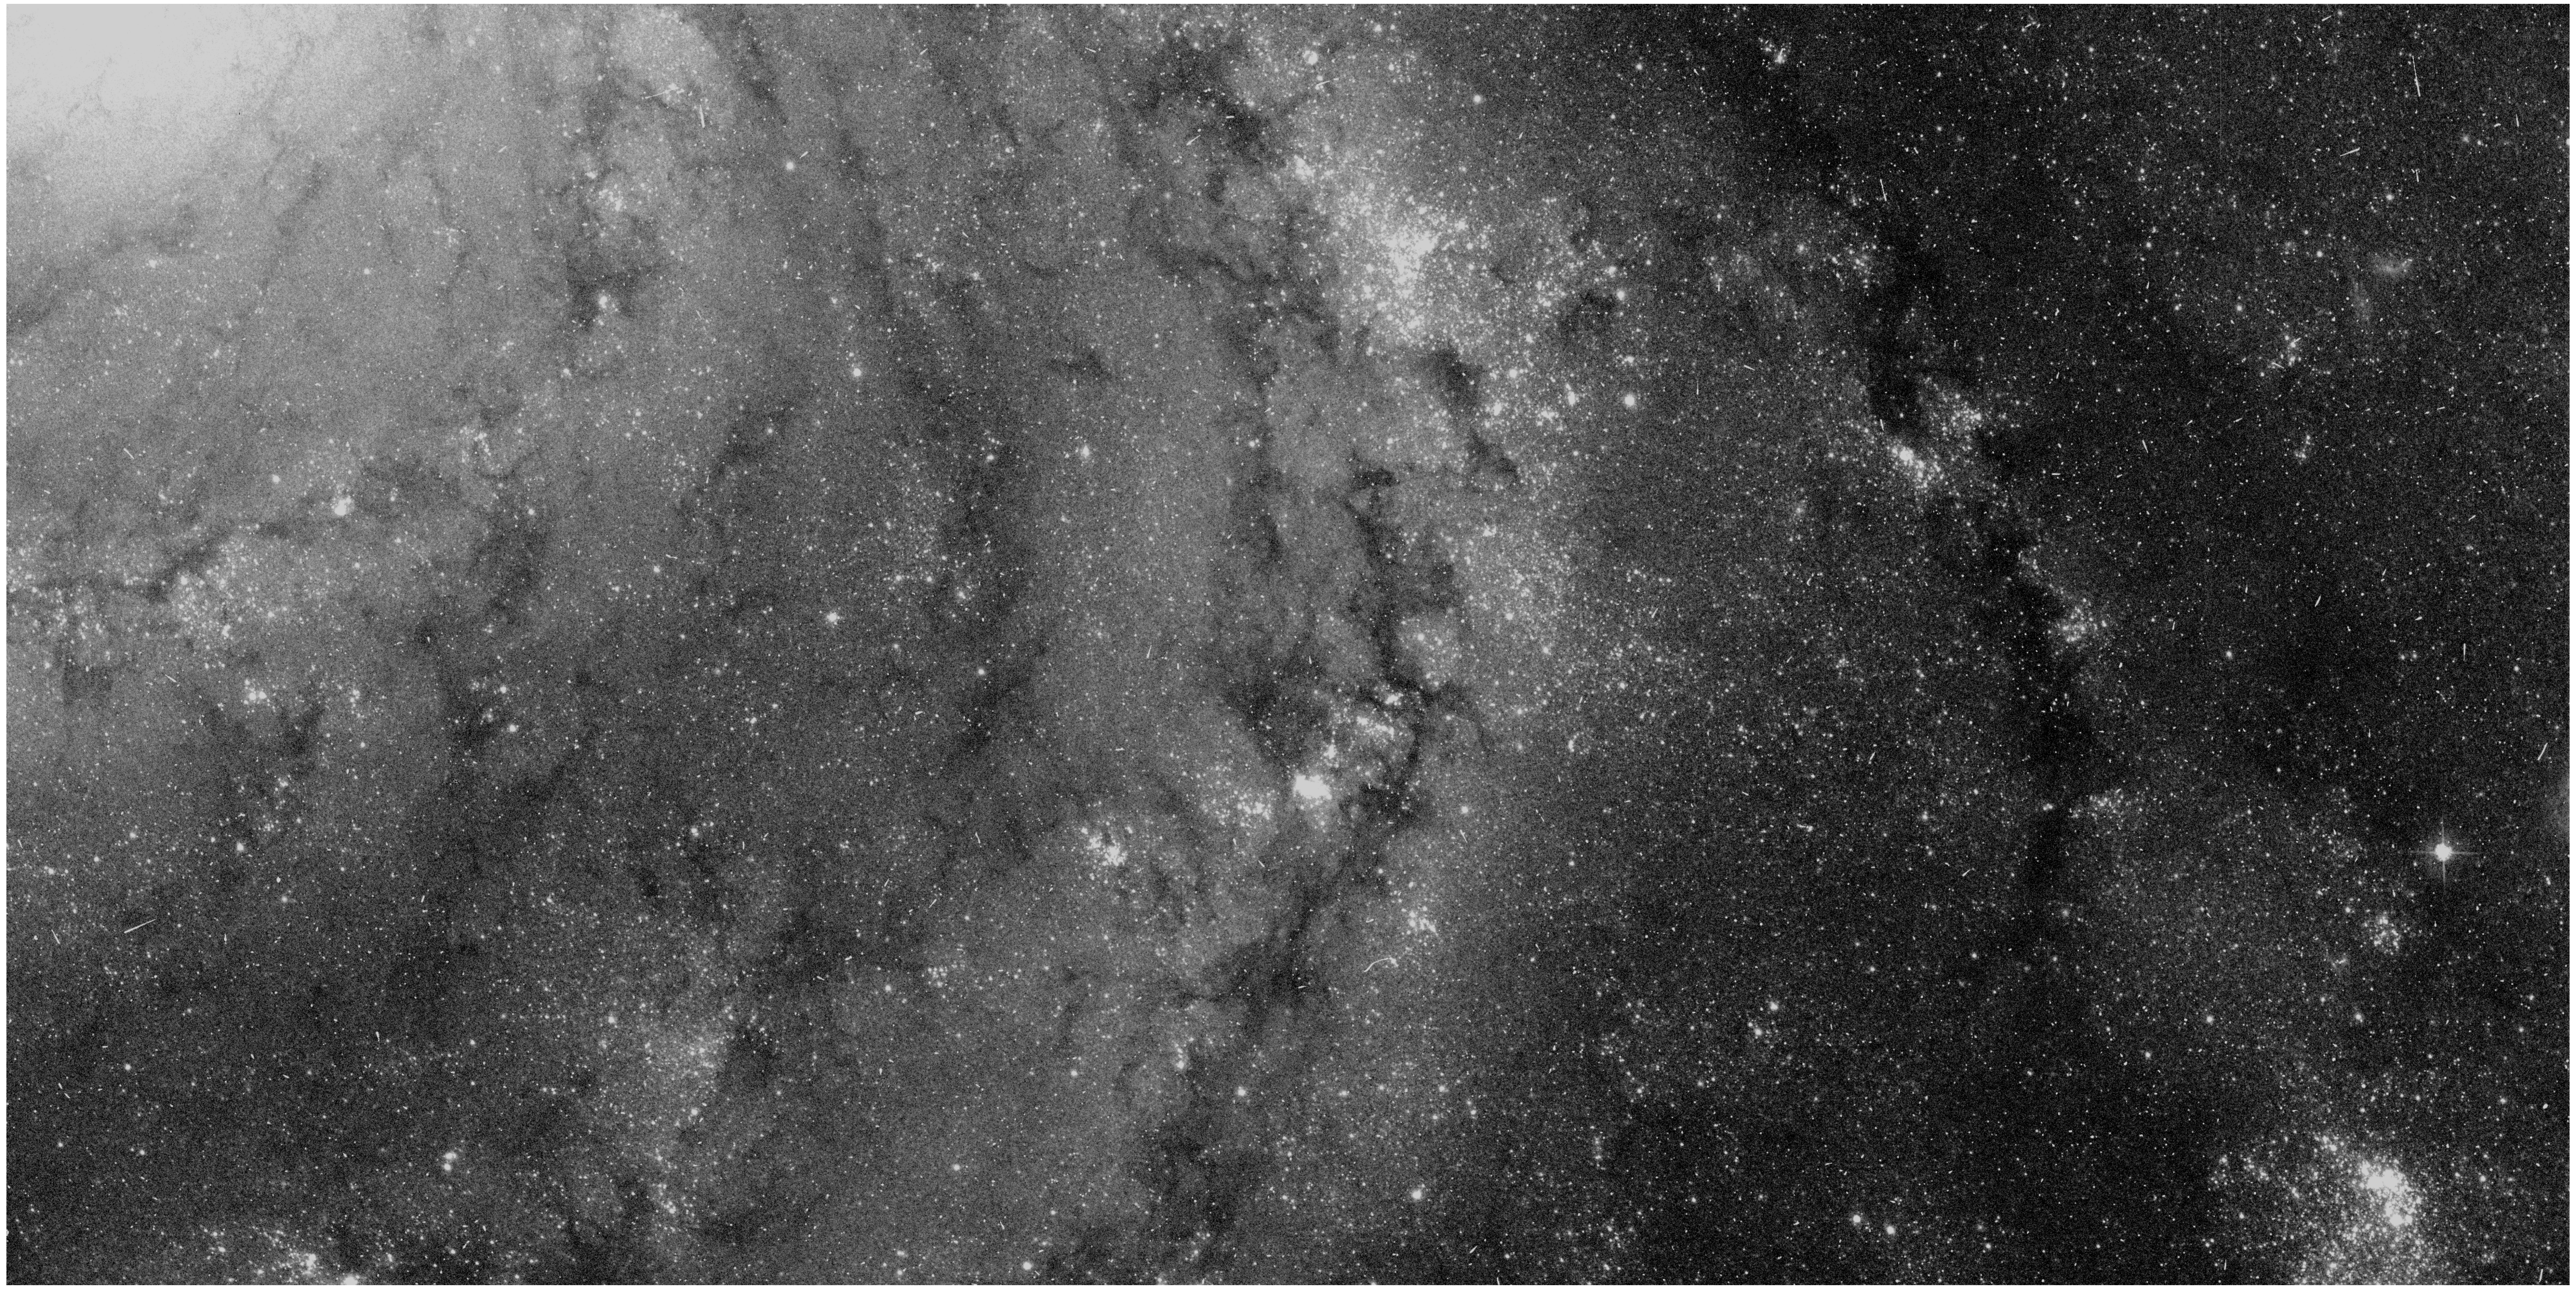

Cosmetic artifacts in the image have been removed

Final Product

Multiple images are combined to form the final product.

Cosmic ray hits have been removed.

This shows both CCDs in the instrument (previous slides only had the lower one).

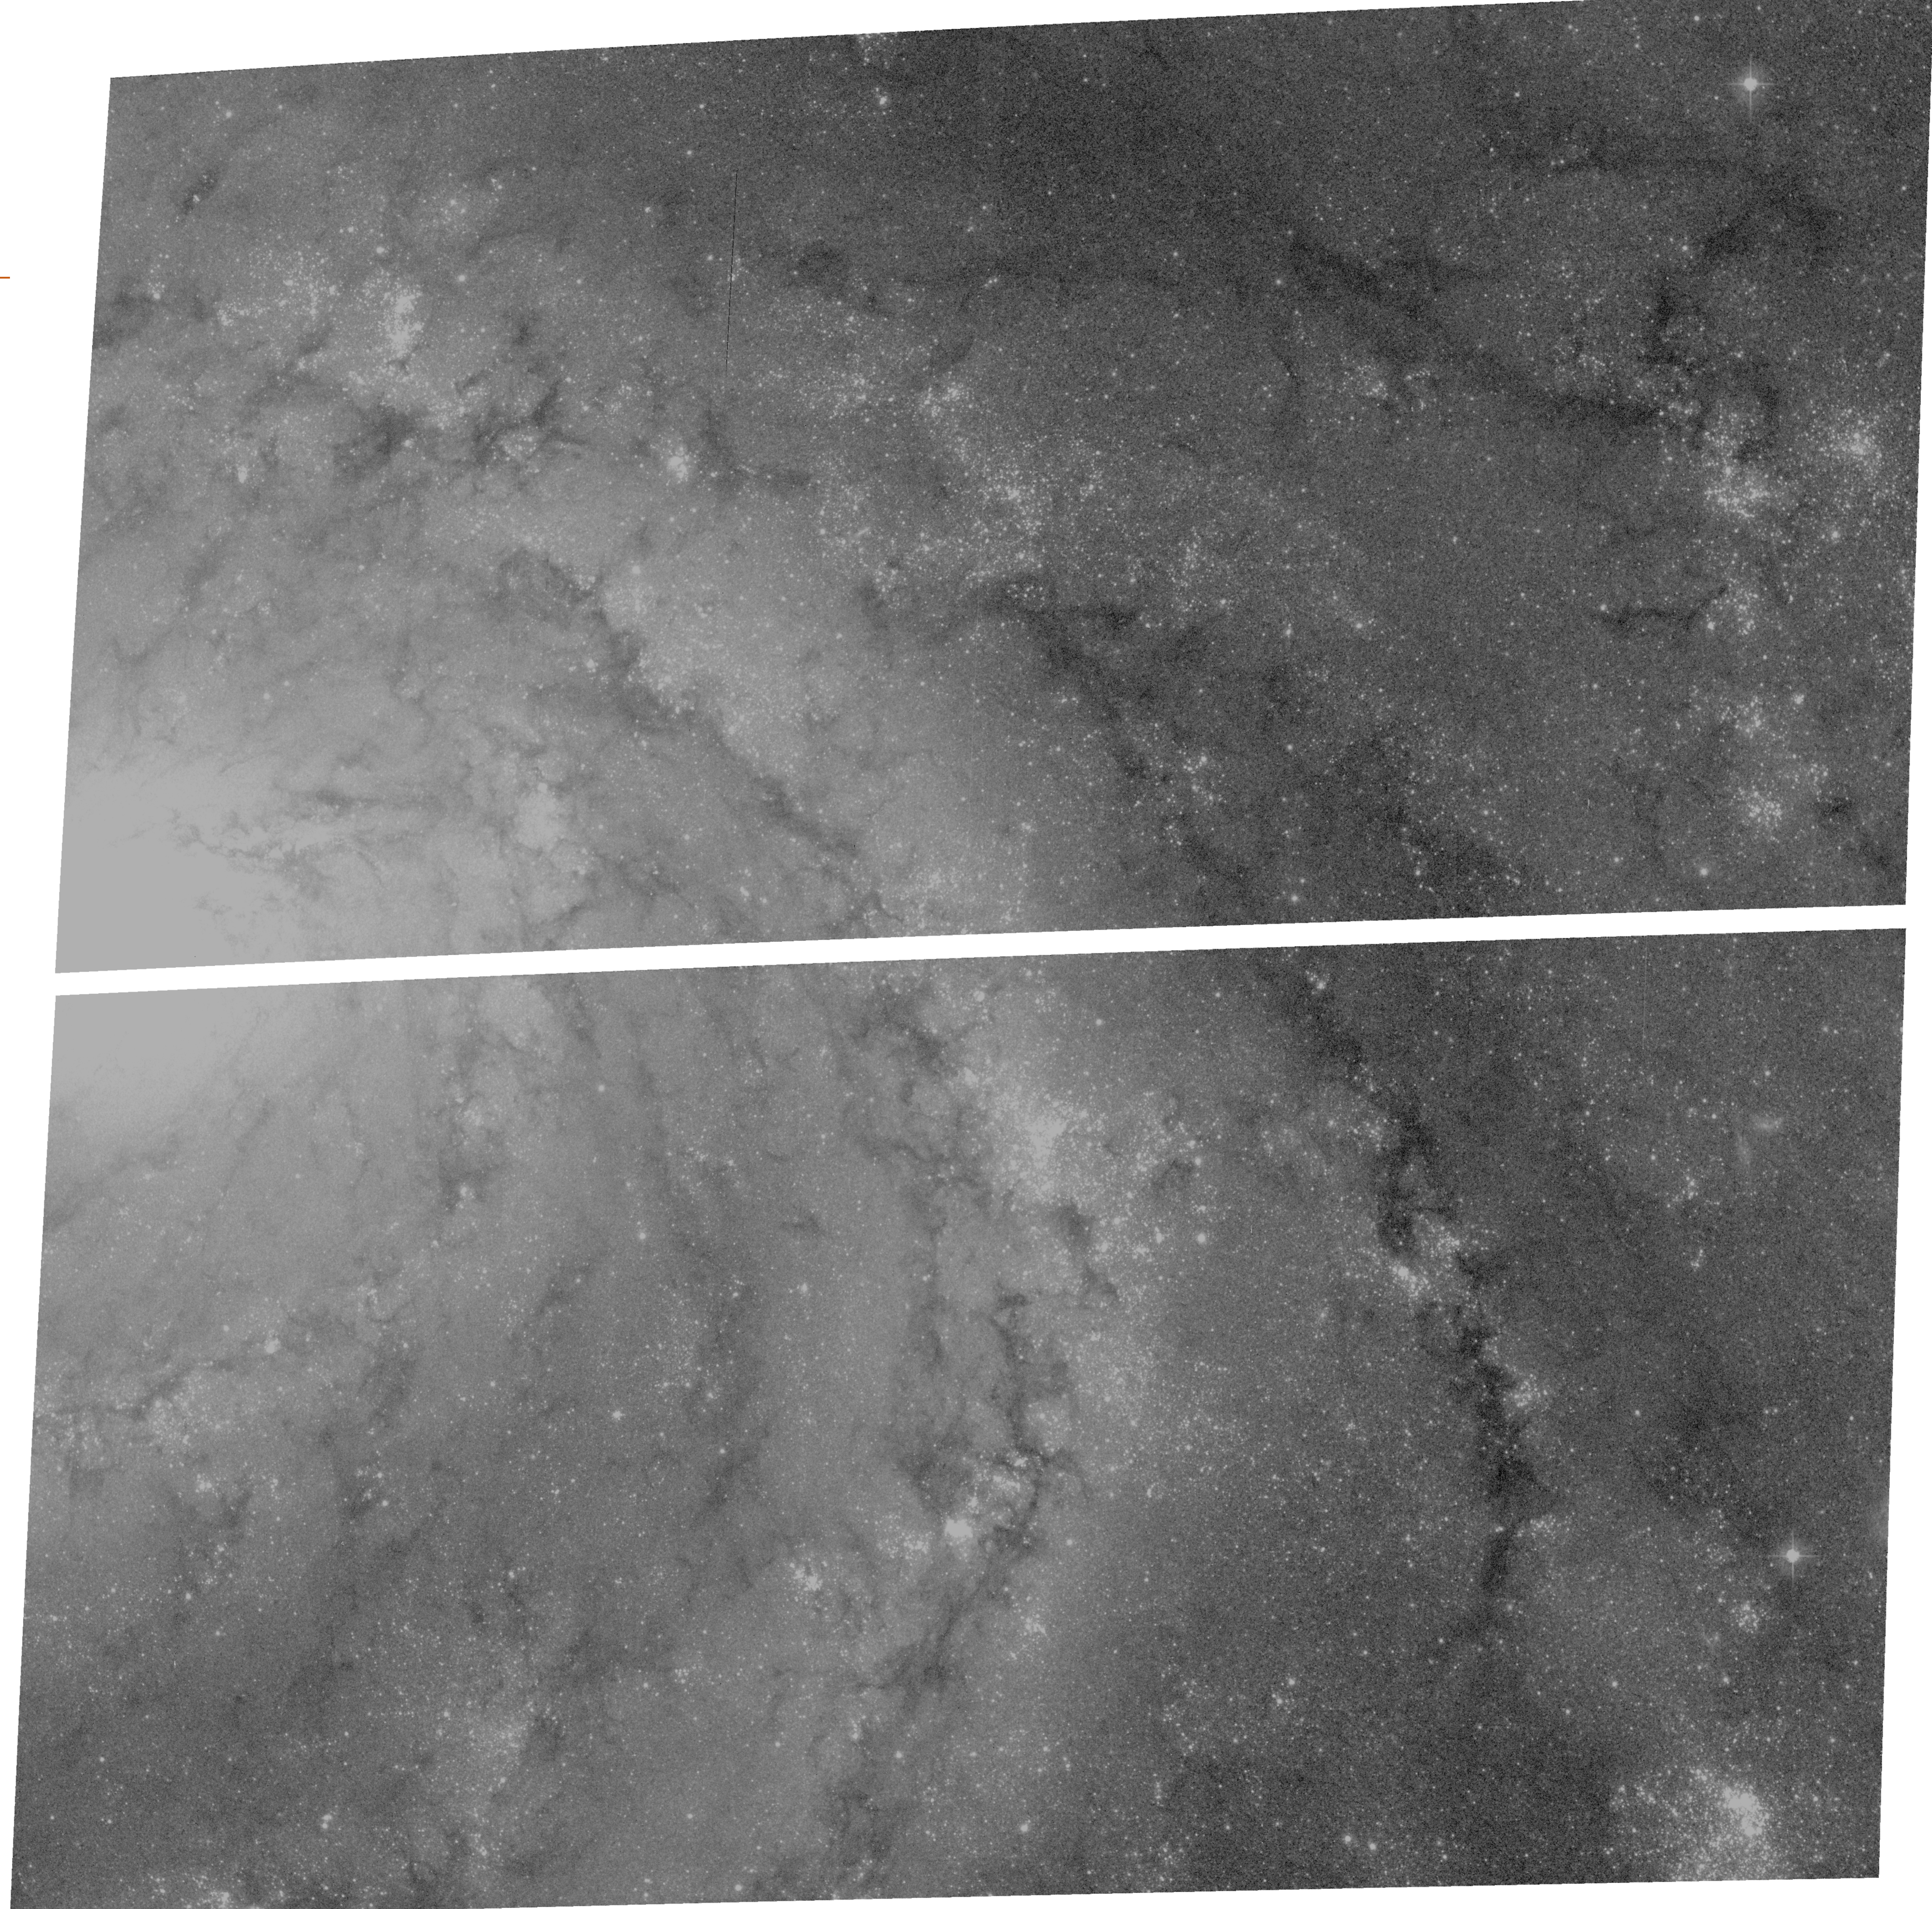

HTCondor Use at STScI

- HST previously used an in-house workflow manager called OPUS
- OPUS featured
  - Distributed processing on a pool of machines (shared disk access)
  - A blackboard paradigm for distributing work
- OPUS was reevaluated under a trade-study in 2011 to choose the workflow management system for JWST
  - HTCondor was chosen for use with custom-made OWL (Orchestrated Workflow Layer) add-on
  - The HST Mission Office decided to allow OPUS to be phased out for HST processing in favor of HTCondor/OWL
- OPUS was gradually phased out and was fully retired in Feb. 2018.

Why HTCondor?

- Better performance and flexibility for large processing runs
- Capability to add machines when needed, then release them for normal operational loads
- A more maintainable, reliable system in the future
  - OPUS was developed in-house, but the expertise had left
  - Huge HTCondor user base (with conferences and everything)
- OPUS could not handle the large data processing needs of JWST. Converting HST to HTCondor had the advantages of
  - Operators and developers only need to know one system
  - Use HST to gain HTCondor expertise before JWST launch and operations

What and Why is OWL?

- OWL = Orchestrated Workflow Layer
- HTCondor manages compute resources extremely well
- Lacks services for managing and tracking the data being processed
- What OWL provides
  - A job-tracking database table (the blackboard) that captures every step in the workflow populated and updated by HTCondor job hooks
  - Template-driven workflow generation (“DAGs on-the-fly”) using the Jinja2 template engine

DAGs-on-the-Fly

- DAG templates

```
# Job definitions.
JOB 2FITS sdp_edt2fits_{{ dataset }}.job
JOB RF sdp_bestref_{{ dataset }}.job
JOB BC sdp_before_calib_{{ dataset }}.job
JOB CA sdp_calibration_{{ dataset }}.job
JOB MD sdp_astrodrizzle_{{ dataset }}.job
JOB AC sdp_after_calib_{{ dataset }}.job
JOB INGEST_SCI archive_submit_{{ dataset }}.job
JOB PVW sdp_preview_{{ dataset }}.job
JOB CL sdp_clean_{{ dataset }}.job

# Relationships.
PARENT 2FITS CHILD RF
PARENT RF CHILD BC
PARENT BC CHILD CA
PARENT CA CHILD MD
PARENT MD CHILD AC
PARENT AC CHILD INGEST_SCI
PARENT INGEST_SCI CHILD PVW
PARENT PVW CHILD CL

# Condor priorities
PRIORITY 2FITS      50
PRIORITY RF         55
PRIORITY BC         60
PRIORITY CA         65
PRIORITY MD         70
PRIORITY AC         75
PRIORITY INGEST_SCI 85
PRIORITY PVW        80
PRIORITY CL         90
```

sdp_prearch_wf3.dag

DAGs-on-the-Fly

- DAGs are created for specific datasets
  (here WFC3 datasets idkv02010 and idkv02020)

```
[redacted]> condor_q


-- Schedd: [redacted] @ 05/21/18 15:57:15
OWNER   BATCH_NAME                           SUBMITTED    DONE    RUN     IDLE   TOTAL JOB_IDS
[redacted] sdp_prearch_wf3_idkv02010.dag+5640228   5/21 15:28      3      1       _       9 5640362.0
           sdp_prearch_wf3_idkv02020.dag+5640230   5/21 15:28      4      1       _       9 5640525.0

4 jobs; 0 completed, 0 removed, 0 idle, 4 running, 0 held, 0 suspended
[redacted]> condor_q --nobatch


-- Schedd: [redacted] @ 05/21/18 15:57:26
 ID         OWNER             SUBMITTED     RUN_TIME ST PRI SIZE  CMD
5640228.0  [redacted]        5/21 15:28   0+00:29:13 R  0     0.3 condor_dagman -f -l . -Debug 3 -Lockfile sdp_prearch_wf3_idkv02010.dag.lock -AutoRescue 1 -DoRescu
5640230.0                    5/21 15:28   0+00:29:13 R  0     0.3 condor_dagman -f -l . -Debug 3 -Lockfile sdp_prearch_wf3_idkv02020.dag.lock -AutoRescue 1 -DoRescu
5640362.0                    5/21 15:29   0+00:28:21 R  65  977.0 calwf3_OWL.sh
5640525.0                    5/21 15:56   0+00:00:44 R  70    0.0 md_xxx_OWL.sh

4 jobs; 0 completed, 0 removed, 0 idle, 4 running, 0 held, 0 suspended
```

What and Why is OWL?

- OWL = Orchestrated Workflow Layer
- HTCondor manages compute resources extremely well
- Lacks services for managing and tracking the data being processed
- What OWL provides
  - A job-tracking database table (the blackboard) that captures every step in the workflow populated and updated by HTCondor job hooks
  - Template-driven workflow generation (“DAGs on-the-fly”) using the Jinja2 template engine
  - A web app (OWL GUI) for monitoring dataset processing status and other system features

OWL GUI (Data Flow)

	Fileset	Process Name	Status	Rescue	Start Time	Completion Time	Controls	State	Exit Code
▸	iccz13q0q	INGEST		RESCUE	2018-05-11 18:33:50.0	1970-01-01 00:00:00.0		Starting	
▸	lcd903010	REPRO_CL		RESCUE	2018-05-11 18:33:50.0	1970-01-01 00:00:00.0		Starting	
▸	iccz13p1q	CL			2018-05-11 18:33:13.0	2018-05-11 18:33:15.293		Exited	0
▸	iccz13ptq	INGEST			2018-05-11 18:33:12.0	1970-01-01 00:00:00.0		Running	
▸	iccz13psq	INGEST			2018-05-11 18:33:01.0	1970-01-01 00:00:00.0		Running	
▸	iccz13p2q	CL			2018-05-11 18:32:59.0	2018-05-11 18:33:04.783		Exited	0
▸	iccz13oeq	CL			2018-05-11 18:32:53.0	2018-05-11 18:32:54.13		Exited	0
▸	iccz13pwq	INGEST			2018-05-11 18:32:50.0	1970-01-01 00:00:00.0		Running	
▸	iccz13ozq	CL			2018-05-11 18:32:24.0	2018-05-11 18:32:29.06		Exited	0
▸	idlx08030	CL			2018-05-11 18:32:22.0	2018-05-11 18:32:28.147		Exited	0
▸	iccz13pqq	INGEST			2018-05-11 18:32:22.0	1970-01-01 00:00:00.0		Running	

OWL GUI (Data Flow For a Single Dataset)

idlx08030	CL			2018-05-11 18:32:22.0	2018-05-11 18:32:28.147		Exited	0

Show 25 entries

Search:

End

Fileset	Process Name	Status	Start Time	Completion Time	Lasted Rescued	State	Exit Code	Owner
idlx08030	CL	Exited	2018-05-11 18:32:22.0	2018-05-11 18:32:28.147	None	Exited	0	
idlx08030	INGEST	Exited	2018-05-11 18:20:36.0	2018-05-11 18:32:08.713	None	Exited	0	
idlx08030	CL	Exited	2018-05-11 17:07:09.0	2018-05-11 17:07:10.33	None	Exited	0	
idlx08030	PVW	Exited	2018-05-11 17:05:01.0	2018-05-11 17:07:01.373	None	Exited	0	
idlx08030	INGEST_SCI	Exited	2018-05-11 17:03:01.0	2018-05-11 17:04:52.457	None	Exited	0	
idlx08030	AC	Exited	2018-05-11 17:02:48.0	2018-05-11 17:02:49.967	None	Exited	0	
idlx08030	MD	Exited	2018-05-11 17:00:06.0	2018-05-11 17:02:39.637	None	Exited	0	
idlx08030	CA	Exited	2018-05-11 14:49:43.0	2018-05-11 16:59:57.987	None	Exited	0	
idlx08030	BC	Exited	2018-05-11 14:49:30.0	2018-05-11 14:49:33.723	None	Exited	0	
idlx08030	RF	Exited	2018-05-11 14:49:09.0	2018-05-11 14:49:21.61	None	Exited	0	
idlx08030	2FITS	Exited	2018-05-11 14:48:45.0	2018-05-11 14:48:57.177	None	Exited	0	
Fileset	**Process Name**	**Status**	**Start Time**	**Completion Time**	**Lasted Rescued**	**State**	**Exit Code**	**Owner**

Start

Showing 1 to 11 of 11 entries

Previous 1 Next

Managing Datasets

- OWL allows us to specify and feed data processing runs into the system
  - Data Processing Queue (DPQ) database table holds workflow requests

Fri, 11 May 2018 18:11:21 UTC 131 DOY

Queue Processing

DESCRIPTION: Shows DPQ_DB.DpQueue table entries for workflows that have not yet been started by the Shoveler task, or for informational message workflowTypes that communicate status of processing where a workflow could not yet be started for a fileset. The Actions column provides a trash icon that, when selected, will result in deletion of the table row from the DPQ_DB.DpQueue database table. This can be used to remove informational message lines that have been looked into and are no longer needed. Once the Shoveler picks up an entry from the DPQ_DB.DpQueue table, the row on this display will disappear, and migrate to the DPQ_DB.DpQueueHist table, and be visible on the History tab.

Refresh | Clear Filters | Show 25 rows | Show/Hide columns

Show 25 entries Search:

Actions	File Set	Priority	Workflow Type	Insert Date	Shovel Date
	idp507czq	15	wf3_single_repro	2018-05-09 20:36:18.767	
	idp507d0q	15	wf3_single_repro	2018-05-09 20:36:18.913	
	idp507d1q	15	wf3_single_repro	2018-05-09 20:36:19.03	
	idp507d2q	15	wf3_single_repro	2018-05-09 20:36:19.18	
	idp507d3q	15	wf3_single_repro	2018-05-09 20:36:19.327	
	idp507d4q	15	wf3_single_repro	2018-05-09 20:36:19.48	

Managing Datasets

- OWL allows us to specify and feed data processing runs into the system
  - Data Processing Queue (DPQ) database table holds workflow requests

HTCondor — mburger@dmsops1:~ — ssh dmsops1 — 103×25

```
Last login: Fri May 11 16:48:41 UTC 2018 on pts/7
[redacted]> shoveler_status.py


------------------------------------------------------------
DPQ entries for [redacted] as of  Fri May 11 18:12:49 2018
priority        type    unshoveled entries

         77      ingest_pvw      92
         15      cos_asn_repro   3148
         15      cos_single_repro        4034
         15      reingest_pvw    8765
         15      wf3_asn_repro   8864
         15      wf3_single_repro        24435
         0       awaiting_gsa_jit        5
         0       awaiting_sci_jit        3
         0       collecting_asn_sci      1
         0       collecting_eng          28


------------------------------------------------------------


TOTAL DPQ entries:              49375
[redacted]> 
```

Managing Datasets

- OWL allows us to specify and feed data processing runs into the system
  - Data Processing Queue (DPQ) database table holds workflow requests
  - The DPQ is populated by Pollers which watch target directories for new files. When new files are found, workflows are added to the DPQ table.

OWL Pollers

Start/Stop	Status	Poller	Last Update	Interval
		repro_manager	2017-09-26 11:04:06.8070000	120
		rescue_me	2017-09-26 11:05:26.6630000	120
		SHOVELER	2017-09-13 18:19:42.8470000	120
		CRDS	2017-09-25 19:35:13.0600000	60
		CSUM_COS	2017-09-26 11:05:18.6400000	60
		CSUM_STI	2017-03-17 18:56:40.1270000	60
		dan	2017-09-26 11:05:19.0500000	60
		dlg	2000-01-01 00:00:00.0000000	60
		DP_ACS_OWL	2017-09-26 11:05:58.0500000	60
		DP_COS_OWL	2017-09-26 11:05:51.9200000	60
		DP_STI_OWL	2017-09-26 11:05:55.4870000	60
		DP_WF3_OWL	2017-09-26 11:05:58.0870000	60
		jit	2000-01-01 00:00:00.0000000	60
		pnm	2017-09-26 11:05:17.3830000	60
		tnm	2000-01-01 00:00:00.0000000	60

Managing Datasets

- OWL allows us to specify and feed data processing runs into the system
  - Data Processing Queue (DPQ) database table holds workflow requests
  - The DPQ is populated by Pollers which watch target directories for new files. When new files are found, workflows are added to the DPQ table.
  - The Shoveler governs the order and rate at which DPQ entries are sent through the OWL workflow template generator to be transformed into executing DAGMan jobs on the HTCondor pool
- The rescue server receives rescueDAG requests to re-try a failed workflow

A Failed Workflow

	Fileset	Process Name	Status	Rescue	Start Time	Completion Time	Controls	State	Exit Code	Remote Host	Owner
	DAN14448	dan_receipt		RESCUE	2017-09-20 04:26:41.0	2017-09-20 04:26:41.417		Exited	1		

Show 10 entries

Search:

Fileset	Process Name	Status	Start Time	Completion Time	Lasted Rescued	State	Exit Code	Owner
DAN14448	dan_receipt	Exited	2017-09-20 04:26:41.0	2017-09-20 04:26:41.417	None	Exited	1	
Fileset	**Process Name**	**Status**	**Start Time**	**Completion Time**	**Lasted Rescued**	**State**	**Exit Code**	**Owner**

Showing 1 to 1 of 1 entries

Previous 1 Next

Future STScI Operations

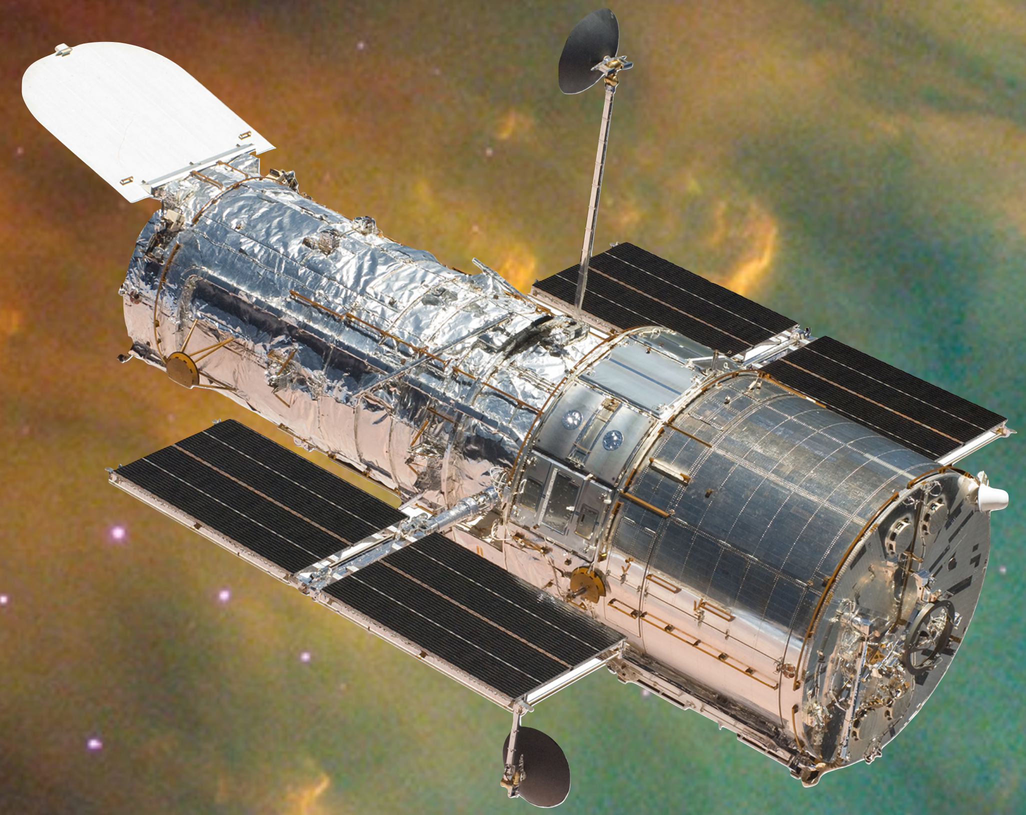

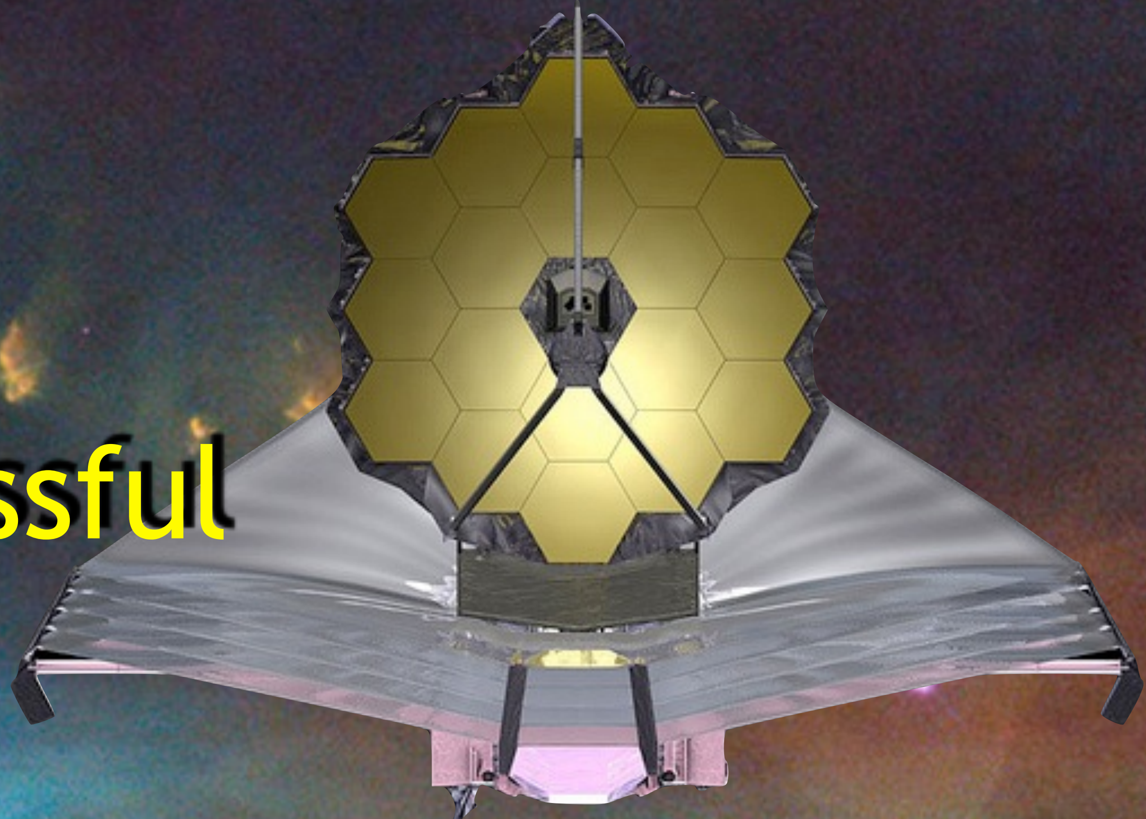

- ★ The Hubble Space Telescope is the most successful scientific observatory in history
  - ★ Operations are expected to continue into the mid 2020’s
- ★ The James Webb Space Telescope will be the premier observatory of the next decade
  - ★ Expected to launch spring 2020
  - ★ 6.5 meter segmented mirror optimized for the infrared
- ★ HTCondor + OWL provide a unified system for the data processing pipelines, allowing for efficient calibration and archiving operations.