

TACC: SITE UPDATE



John Cazes

Director, High Performance Computing

The University of Texas at Austin

September 21, 2023

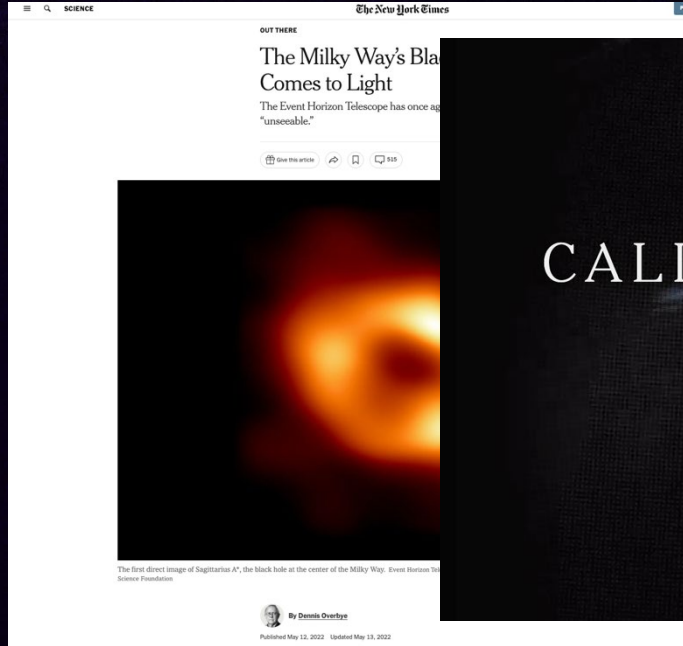
TACC IN A NUTSHELL

- ▶ We operate the Frontera, Stampede2, Jetstream, and Chameleon systems for the National Science Foundation
- ▶ Lonestar6 for our Texas academic and industry users.
- ▶ Altogether, ~20k servers, >1M CPU cores, 1k GPUs
- ▶ About seven billion core hours over several million jobs per year – for 3,000 projects and ~40,000 users per year.
- ▶ 190 Staff (~70 PhD)



ALL FOR UNCLASSIFIED, OPEN SCIENCE

WIDEST VIEW OF EARLY UNIVERSE HINTS AT GALAXY AMONG THE EARLIEST EVER DETECTED



The New York Times

THE COMING CALIFORNIA MEGASTORM

A different 'Big One' is approaching.
Climate change is hastening its arrival.

By **Raymond Zhong** | Graphics by **Mira Rojanasakul**
Photographs by **Erin Schaff**

Aug. 12, 2022

supercomputers enable scientists to combine myriad
/ resulting in single image

rt, CNS / Faith Singer, TACC

LinkedIn Email



Members of the CEERS collaboration explore the first wide, deep field image from the James Webb Space Telescope at the Texas Advanced Computing Center's Visualization Lab on the UT Austin campus on July 21, 2022. Credit: Nolan Zunk/UT Austin

Two new images from NASA's James Webb Space Telescope show what may be among the earliest galaxies ever observed. Both images include objects from more than 13 billion years ago, and one offers a much wider field of view than Webb's [First Deep Field image](#), which was released amid great fanfare July 12, 2022.

NEXT-GENERATION WEATHER MODELS

- ▶ NOAA Hydrometeorology Testbed
- ▶ This year's Winter Weather Experiment had three key science goals: to subjectively gauge the utility of convection-allowing model (CAM) forecasts to improve two-to-three day snowfall forecasts; objectively score the snowfall forecasts using community standard verification systems; and determine the optimal combination of physics to use in next-generation models. Recent testbed programs have also included the important task of evaluating NOAA's next-generation weather model, the FV3 model.
- ▶ Keith Brewster, Ming Xue, University of Oklahoma.



Figure 6 - A real-time forecast from the 2021 Winter Weather experiment, enabled by Frontera.

16 NVDIMM Nodes

Each node contains:

- 4 Intel Xeon Platinum 8280M chips
- 2x 28 core 2.2 Ghz Xeon cores
- 384 GB DRAM
- 2 TB NVMe RAM
- 4 TB NVMe disk



8368 Cascade Lake Nodes

Each node contains:

- 2 Intel Xeon Platinum 8280 chips
- 2x 28 core 2.2 Ghz Xeon cores
- 192 GB DRAM
- Mellanox HDR Infiniband



90 GPU Nodes

Each node contains:

- 4 NVIDIA QUADRO RTX 5000 GPUs
- 2 Intel Xeon E5-2620 v4
- 192 GB DRAM

Frontera

Dell 8000+ node cluster
40 Pflops
43 PB Lustre filesystem

FOURTH YEAR OF PRODUCTION ON FRONTERA

- ▶ In the last 12 months:
 - ▶ Uptime of 98%
 - ▶ Average Utilization of 95%
 - ▶ 1M jobs completed
 - ▶ 72M SUs delivered
 - ▶ Zero security incidents
- ▶ On the bright side, we are always full.
- ▶ On the downside, no way to squeeze anything else in.



USAGE

- ▶ **>2,200 jobs were > 512 nodes**
- ▶ >300 jobs at half or full system scale (Consider if all jobs were full scale, and averages 24 hours, we'd only run 365 jobs a year, as opposed to 1M jobs).
- ▶ Flex jobs, used for backfill, represent 15% of the jobs run, but represent less than 0.5% of SUs delivered (263K out of 70M).
- ▶ Small jobs represent ~30% of jobs, but less than 2% of cycles delivered.
 - ▶ So **97% of time goes to jobs >2 nodes.**
 - ▶ Average jobs size about **6x that of Stampede2** – this machine **is** used differently.
- ▶ We tune the scheduling policy multiple times a year... essentially adjusting to demand.

224 Icelake Nodes

Each node contains

- **2 Intel Xeon Platinum 8380 chips**
- **2x 40 core 2.3 Ghz Xeon cores**
- **256 GB DRAM**




3752 KNL Node

Each node contains:

- **1 Intel Xeon Phi 7250 chip**
- **68 1.4 Ghz cores**
- **96 GB DRAM + 16 GB MCDRAM**

RETIRED



Stampede 2

Dell 6000+ node cluster
18 Pflops
20 PB Lustre filesystem

1,000+ projects
5,000+ users

1736 Skylake Nodes

Each node contains

- **2 Intel Xeon Platinum 8160 chips**
- **2x 24 core 2.2 Ghz Xeon cores**
- **192 GB DRAM**

224 Icelake Nodes

Each node contains

- **2 Intel Xeon Platinum 8380 chips**
- **2x 40 core 2.3 Ghz Xeon cores**
- **256 GB DRAM**

1000 Skylake Nodes

Each node contains

- **2 Intel Xeon Platinum 8160 chips**
- **2x 24 core 2.2 Ghz Xeon cores**
- **192 GB DRAM**

Stampede3

- **10 PB VAST Filesystem**
- **400 Gb/sec Omni-Path Express**

560 Sapphire Rapids HBM Nodes



Each node contains:

- **2 Intel Max 9480 chips**
- **2x 56 1.9 Ghz cores**
- **2x 64 GB HBM2e**


20 Ponte Vecchio Nodes

Each node contains

- **4 Intel Data Center GPU Max GPUs**
- **128 X^e cores**
- **128 GB HBM2e**



Single-Node Application Performance Ratios





FRONTERA

TACC | NSF | TEXAS

cazes@tacc.utexas.edu