National Aeronautics and
Space Administration

NASA EXPLORES
Advanced Computing

The NASA Advanced Supercomputing (NAS) Division has been the agency's primary resource for high performance computing, data storage, and advanced modeling and simulation tools for over 30 years. From the 1.9 gigaflop Cray-2 system installed in 1985 to the current petascale Pleiades, Aitken, and Electra superclusters, the NAS facility at NASA's Ames Research Center in Silicon Valley has housed over 40 production and testbed supercomputers supporting NASA missions and projects in aeronautics, human space exploration, Earth science, and space science.

www.nasa.gov

NASA ADVANCED SUPERCOMPUTING (NAS) DIVISION

The NASA Advanced Supercomputing (NAS) facility's computing environment includes the three most powerful supercomputers in the agency: the petascale Electra, Aitken, and Pleiades systems. Developed with a focus on flexible scalability, the systems at NAS have the ability to expand and upgrade hardware with minimal impact to users, facilitating the ability to continually provide the most advanced computing technologies to support NASA's many inspiring missions and projects without significant downtime.

Part of this hardware diversity includes the integration of graphics processing units (GPUs), which can speed up some codes and algorithms run on NAS systems. Recently, the latest Intel Skylake nodes, augmented with NVIDIA Tesla V100 GPUs, were added to the Pleiades supercomputer, providing dozens of teraflops of computational boost to each of the enhanced nodes.

The NAS facility also houses several smaller systems to support various computational needs, including the Endeavour shared-memory system for large-memory jobs, and the hyperwall visualization system, which provides a unique environment for researchers to explore their very large, high-dimensional datasets. The division supports both short-term RAID and long-term tape mass storage systems, providing more than 1,500 users running jobs on the facility's supercomputers with over an exabyte of data storage (with standard compression).

In addition to hardware resources, NAS also provides support services for users to take full advantage of the unique advanced computing environment. These services include application performance optimization, data analytics and machine learning tools, end-to-end networking, scienctific visualization, and data publication.

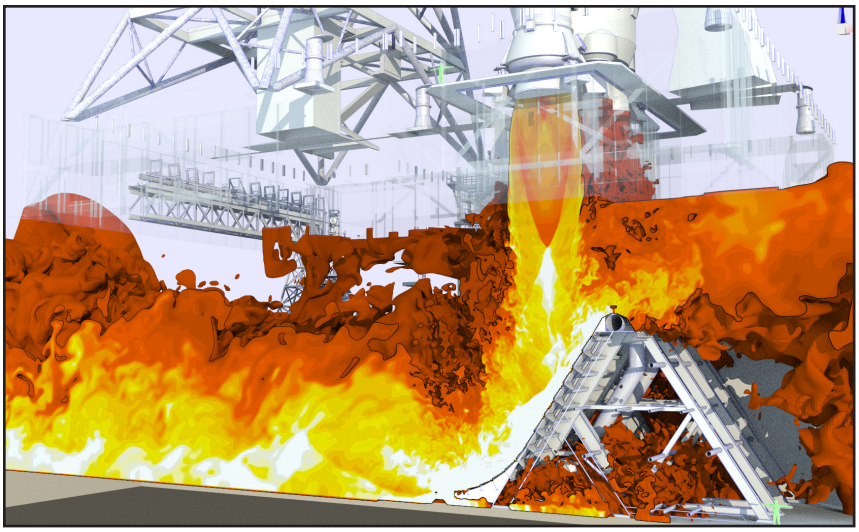

Image from a simulation of launch ignition for NASA's next-generation Space Launch System. Colors indicate temperature, where white is hotter and brown is cooler. The plume is contoured based on the air-mass fraction (that is, the fraction by mass of air vs. gas plume species). This simulation, run on the Pleiades and Electra supercomputers at the NASA Advanced Supercomputing (NAS) facility, generated hundreds of terabytes of data for each calculation. Data analysis and visualization were conducted by NAS visualization experts using custom in-house software and high-performance systems. Michael Barad, Tim Sandstrom, NASA/Ames

Pleiades Architecture Overview

- 158 SGI/HPE racks with Intel Xeon Broadwell, Haswell, Ivy Bridge, and Sandy Bridge processors
- 11,207 nodes, 241,324 cores, and 927 terabytes of memory
- 7.09 petaflops theoretical peak performance
- 5.95 petaflops sustained performance (November 2019)
- 83 GPU-enhanced nodes with 614,400 NVIDIA CUDA cores
  - 64 nodes of Intel Xeon Sandy Bridge processors enhanced with NVIDIA Tesla K40 GPUs
  - 19 nodes of Intel Xeon Skylake processors enhanced with NVIDIA Tesla V100 GPUs

Endeavour Architecture Overview

- 3 SGI UV2000 racks with Intel Xeon Sandy Bridge processors
- 2 nodes, 1,504 cores, and 5.9 terabytes of shared system memory
- 32 teraflops theoretical peak performance

Data Storage Systems

- 40 petabytes of Lustre and NFS "scratch" space
- 1,040-petabyte (1 exabyte) archival tape storage capacity
- 7.6 petabytes of archive disc cache; over 100 petabytes of unique data on tape

Aitken Architecture Overview

- 4 HPE E-Cells with Intel Xeon Cascade Lake processors
- 1,150 nodes, 46,080 cores, and 221 TB of memory
- 2.38 petaflops sustained performance (November 2019)
- 3.69 petaflops theoretical peak performance

Electra Architecture Overview

- 16 SGI D-Racks with Intel Xeon Broadwell processors; 8 HPE E-Cells with Intel Xeon Skylake processors
- 3,456 nodes, 124,416 cores, and 589 TB of memory
- 5.44 petaflops sustained performance (November 2019)
- 8.32 petaflops theoretical peak performance

Merope Architecture Overview

- 56 half-populated SGI Altix racks with Intel Xeon Westmere processors
- 1,792 nodes, 21,504 cores, and 86 TB of memory
- 252 teraflops theoretical peak performance

Visualization Systems: the hyperwall

- 128-screen tiled LCD wall arranged in an 8x16 grid
- 2,560 Intel Xeon Ivy Bridge cores
- 128 NVIDIA GeForce GTX 780 Ti GPUs