

National Aeronautics and Space Administration

High-End Computing Capability

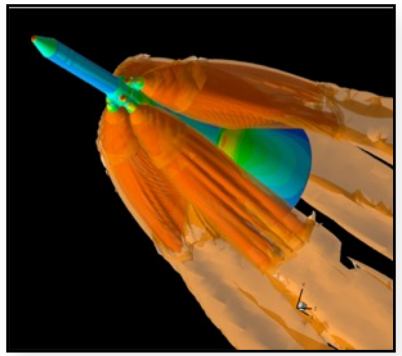
William Thigpen
Manager, HECC Component
william.w.thigpen@nasa.gov

May 22, 2024



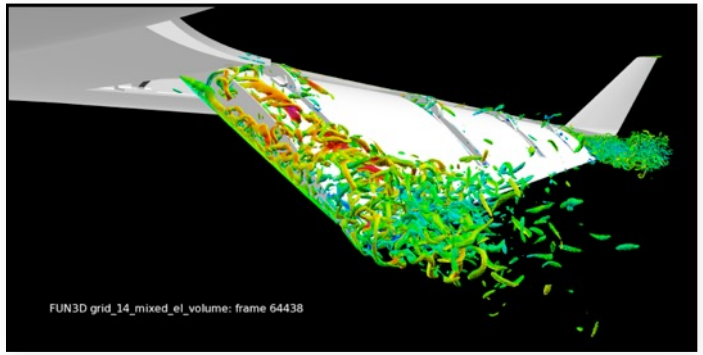
NASA's HEC Requirements: Capacity

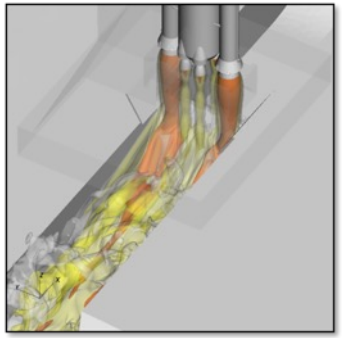
Engineering-related work requires HEC resources that can handle large numbers of relatively-low CPU-count jobs (1,000 – 4,000 cores) with quick turnaround times.



Over 1500 simulations utilized ~2 million processor hours to study launch abort systems on the next-generation crew transport vehicle.

The formation of vortex filaments and their roll-up into a single, prominent vortex at each tip on a Gulfstream aircraft.



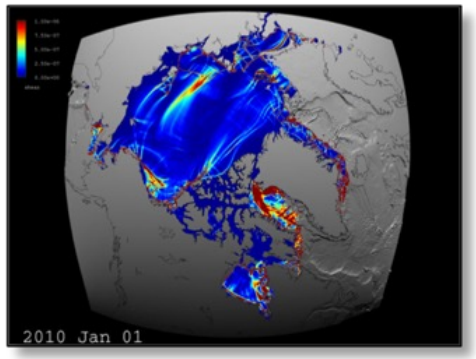


Over 4 million hours were used over a four-month project to evaluate future designed of the next generation launch complex at Kennedy Space Center.



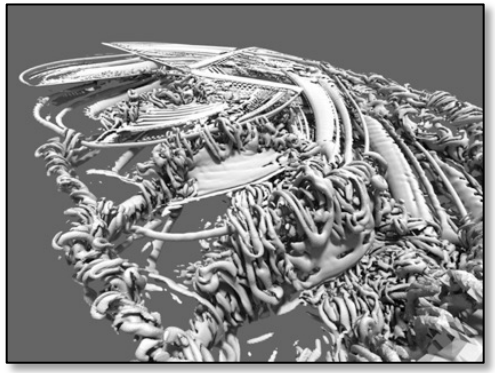
NASA's HEC Requirements: Capability

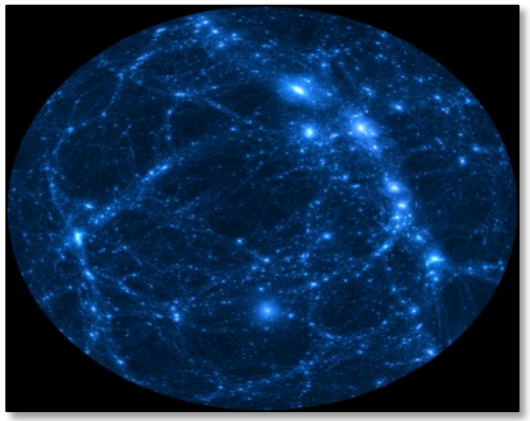
Aeronautics and science related work require HEC resources that can handle high fidelity relatively-large CPU-count jobs (50,000+ cores) with minimal time-to-solution. Capability enables work that wasn't possible on previous architectures.



NASA looks at the oceans, running 100s of jobs using up to 10,000 processors. Looking at the role of the oceans in the global carbon cycle is enabled by access to large processing and storage assets.

For the first time, the Figure-of-Merit has been predicted within experimental error for the V22 Osprey and Black Hawk helicopter rotors in hover, over a wide range of flow conditions.



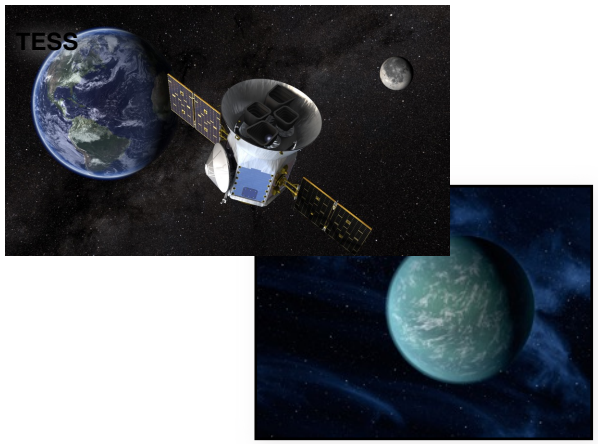


To complete the Bolshoi simulation, which traces how the largest galaxies and galaxy structures in the universe were formed billions of years ago, astrophysicists ran their code for 18 straight days, consuming millions of hours of computer time, and generating massive amounts of data.



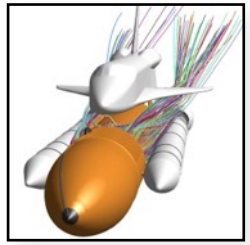
NASA's HEC Requirements: Time Critical

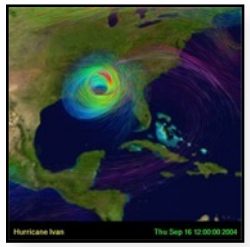
NASA needs HEC resources that can handle time-sensitive, mission-critical applications on demand (maintain readiness).



HECC worked with TESS to design a data pipeline that incorporated lessons learned from earlier missions enabling enormous planetary transit searches to be completed in less than a day with accuracy and effectiveness of the software pipeline.

REENTRY





STORM PREDICTION

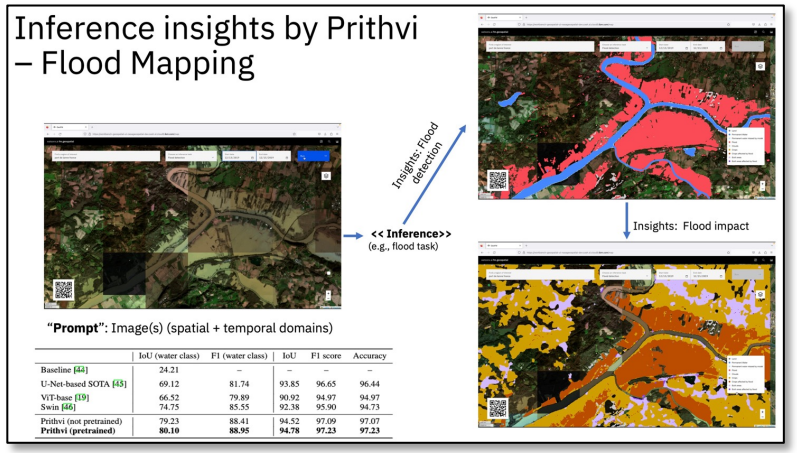


UAVSAR produces polarimetric (PolSAR) and interferometric (repeat-pass InSAR) data that highlight different features and show changes in the Earth over time.



NASA's HEC Requirements: Artificial Intelligence

NASA needs HEC resources that can enable discovery using computers to explore massive data sets.



HECC is working with the Earth Science Division in the Science Mission Directorate, the Department of Energy and industry to develop an Earth Science foundation model to do predictions and analysis.



Across NASA, teams are evaluating, testing and utilizing AI for design, scheduling, identifying areas of interest and predicting future states based on current conditions.



Meeting Today's Requirements

HECC Hardware Assets

Five Compute Clusters:

- **Aitken:** 20 Racks; 3,200 nodes; 13.1 PF; 10,204 SBU/hr
- **Pleiades:** 143 Racks; 10,310 nodes; 6.83 PF; 7,600 SBU/hr
- **Cabeus:** 25 Racks; 187 nodes; 756 GPUs; 7.76 PF; 6,441 SBU/hr
- **Electra:** 24 Racks; 3,456 nodes; 8.32 PF; 4,815 SBU/hr
- **Endeavour:** 2 Racks; 2 nodes; 155 TF; 81.2 SBU/hr

Other Systems:

- **1 Visualization Cluster:** 1 billion pixels; 128 nodes; 2.61 PF
- **13 Lustre Filesystems:** 118.1 PB
- **3 NFS Filesystems:** 3.3 PB
- **2 BeeGFS Filesystems:** 1.5 PB
- **1 VAST Filesystem:** 2.5 PB
- **Archive System:** 1,572 PB

NAS Facility Extension

- One-acre 30 MW site to house HPC systems in modules.

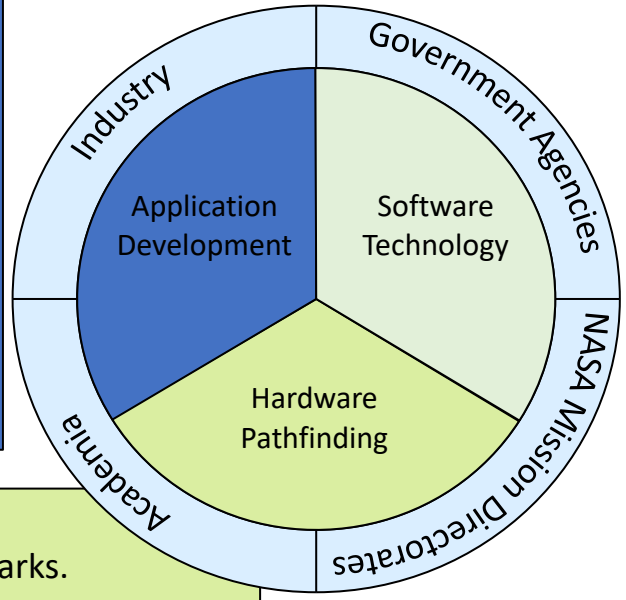
HECC provides a suite of complimentary services to the user community to enhance scientific and engineering results obtained from the hardware assets.

- **Systems:** Customized solutions including compute and storage solutions to meet specific project or mission requirements. Cloud access for immediate or non-standard computing.
- **Cloud Resources:** Customized cloud access for hybrid data pipelines, traditional simulations and exploratory research.
- **Application Performance and Productivity:** Software solutions provided to research/engineering teams to better exploit installed systems.
- **Visualization and Data Analysis:** Custom visualization during traditional post-processing or concurrent during simulation to understand complex interactions of data.
- **Networks:** End-to-end network performance enhancements for user communities throughout the world.
- **Data Analytics:** Exploitation of data sets through neural nets and emerging new techniques.
- **Machine Learning:** Custom environments to enable learning through advanced data techniques.
- **Custom Data Gateways:** Custom data portals to support diverse programs and projects.



Preparing for the Future

- New algorithms to match hardware realities.
- Mixed precision and machine learning/AI in computation.
- New methods in visualization and data analytics.
- Application enhancements to exploit I/O advancements.
- Fault resilience.
- Performance optimization.



- Examine requirements.
- Develop necessary benchmarks.
- Assess I/O approaches.
- Develop models to predict performance.
- Make recommendations for pathfinding systems.

- Research and develop programming approaches, for accelerated and non-accelerated architectures.
- Optimize and port techniques/tools.
- Optimize and port techniques/tools for visualization, data analytics, AI, and machine learning.
- Research and develop efficient math libraries.
- New file system software technologies.
- Research and develop machine learning techniques for problem solving and system error tracking/prediction.
- Evaluate software packaging technology, such as containers.



FY23 Highlights

HECC supported 794 projects across the majority of NASA:

- **ARMD** (350 users) – 153 Projects – 43,191,839 Standard Billing Units (SBUs)*
- **ESDMD** (126 users) – 38 Projects – 24,758,478 SBUs
- **SMD** (1,003 users) – 548 Projects – 67,797,882 SBUs
- **SOMD** (23 users) – 14 Projects – 1,054,840 SBUs
- **STMD** (80 users) – 34 Projects – 5,618,412 SBUs
- **NESC** (8 users) – 7 Projects – 10,706,595 SBUs



*The SBU is a way of standardizing work across dissimilar architectures. The current value of 1 represents the work done on a 2-processor Broadwell node in 1 hour. The FY23 value was $0.30/SBU.

FY23 Highlights (cont.)

- Increased General Purpose Graphics Processing Unit (GPGPU) Cluster by 324% providing a platform for machine learning, artificial intelligence and accelerated computational fluid dynamics.
- Increased system storage by 47.3 PB (quadrillion bytes).
- Improved the tools available to users for searching databases shared through the online portal and improved performance and reliability enabling the sharing of data externally without incurring the cost of duplicate storage.
- Identified an inefficiency on a JPL parallel application code through active monitoring and worked with the science team to achieve a 2.5x speedup.
- Provided cloud computing resources and mentors for the 2023 Winter Classic Student Cluster Competition providing introduction and access of high-end computing to students at minority serving institutions.
- Worked with NASA Langley and NVIDIA to jointly organize the 2023 NASA GPU hackathon bringing together application developers and computer experts to focus on migrating important NASA applications to accelerators improving time to solution and reducing cost and power necessary to obtain the solution.
- Worked with the Jet Propulsion Laboratory (JPL) on new approaches and capabilities that enable the next-generation, hybrid mission data processing systems to leverage the cloud and HECC facility's computational and storage assets in support of large-scale reprocessing campaigns for current and future Earth Observing Systems (SWOT, NISAR, SBG)



FY24 Plans

- Deploy second compute module on R&D-099, growing the NASA-leading energy-effective modular technology to facilitate further growth. Move the four Milan compute racks from N258 to R&D-099, saving 90% of the energy and water necessary to cool them.
- Release the RFP for the R&D-099 data module, which will provide failover capability for the HECC enclave and enable portions of the enclave to be in maintenance while maintaining computational capability overall during schedule or unscheduled outages.
- Deploy the next-generation hyperwall with one billion pixels to provide an environment for scientific discovery. Scientists have stated that they can accomplish more in a day using the hyperwall than they could do in 6 months on their workstation.
- Deploy Aitken and Cabeus expansions.
- Work with science teams, engineering teams, and mission directorates to develop a plan to procure and deploy an agency-scaled high-end computer environment to meet the future requirements for advanced computing, reducing programmatic risks and improving time to solution.



Proposed Changes

- The following eight slides discuss a plan that was presented as part of the budget process for FY26.
- The plans were presented to center and portfolio management, and then to the HEC Board of Advisors.
- We received concern from some of the mission directorates, and so have scheduled or will schedule meetings to discuss their concerns.
- These slides should be viewed as an unapproved plan at this time, and so you can provide input here or through your management chain if you have concerns.



HECC Evolution

System Evolution

- The growth pattern of HECC has evolved into a highly complex system of systems with very old equipment that makes maintenance and support unnecessarily complicated and difficult or in some cases impossible to perform essential patches increasing the vulnerability of the overall environment.
- In FY24 we have made expansions to both CPU and GPU systems improving the work accomplished per watt.
- In FY25 we plan to decommission Pleiades and the oldest part of Electra freeing up $2.5M in electrical cost. This loss of computing will be offset by growth in both the CPU and GPU clusters ending with an overall growth in compute capability.
- A major redesign of the InfiniBand networks will begin in FY24 and deployed early in FY25 improving stability and removing dependencies between our various compute clusters.
- FY25 – FY30 funds will be used to alternately expand the CPU and GPU clusters. Mission Directorates and programs will be given the opportunity to augment these purchases.
- Purchasing GPUs can have the largest impact on work done, but most users still use CPUs.



Storage Evolution

- As with systems, old storage hardware presents unique challenges in operation and maintenance.
- In late FY24 we will decommission our older file systems.
- FY25 – FY30 Storage augmentations are planned annually.
- If an organization elects to purchase computing, they will be charged an additional 17.5% for storage that will be added to the next storage purchase.

Removing Old Equipment

Pleiades and Electra Broadwell Nodes

- Pleiades is comprised of 13 Sandy Bridge racks (Installed 2013), 73 Ivy Bridge racks (Installed 2014), 28.5 Haswell racks (Installed 2015), and 28 Broadwell Racks (Installed 2016). Electra has 16 Broadwell Racks (Installed 2016 (operational 01/2017)).
- The InfiniBand network of Pleiades is made up of multiple generations of InfiniBand.
- Removing Pleiades removes half of our node types.
- Pleiades provide 7,600 SBU/hour. The Electra Broadwell nodes provide 1,152 SBU/hr. The Milan expansion adds 2,161 SBU/hour. The next gen CPU procurement in early FY25 adds 4,300 SBUs. On the GPU side the GH200 will add 4,800 SBUs.

Network

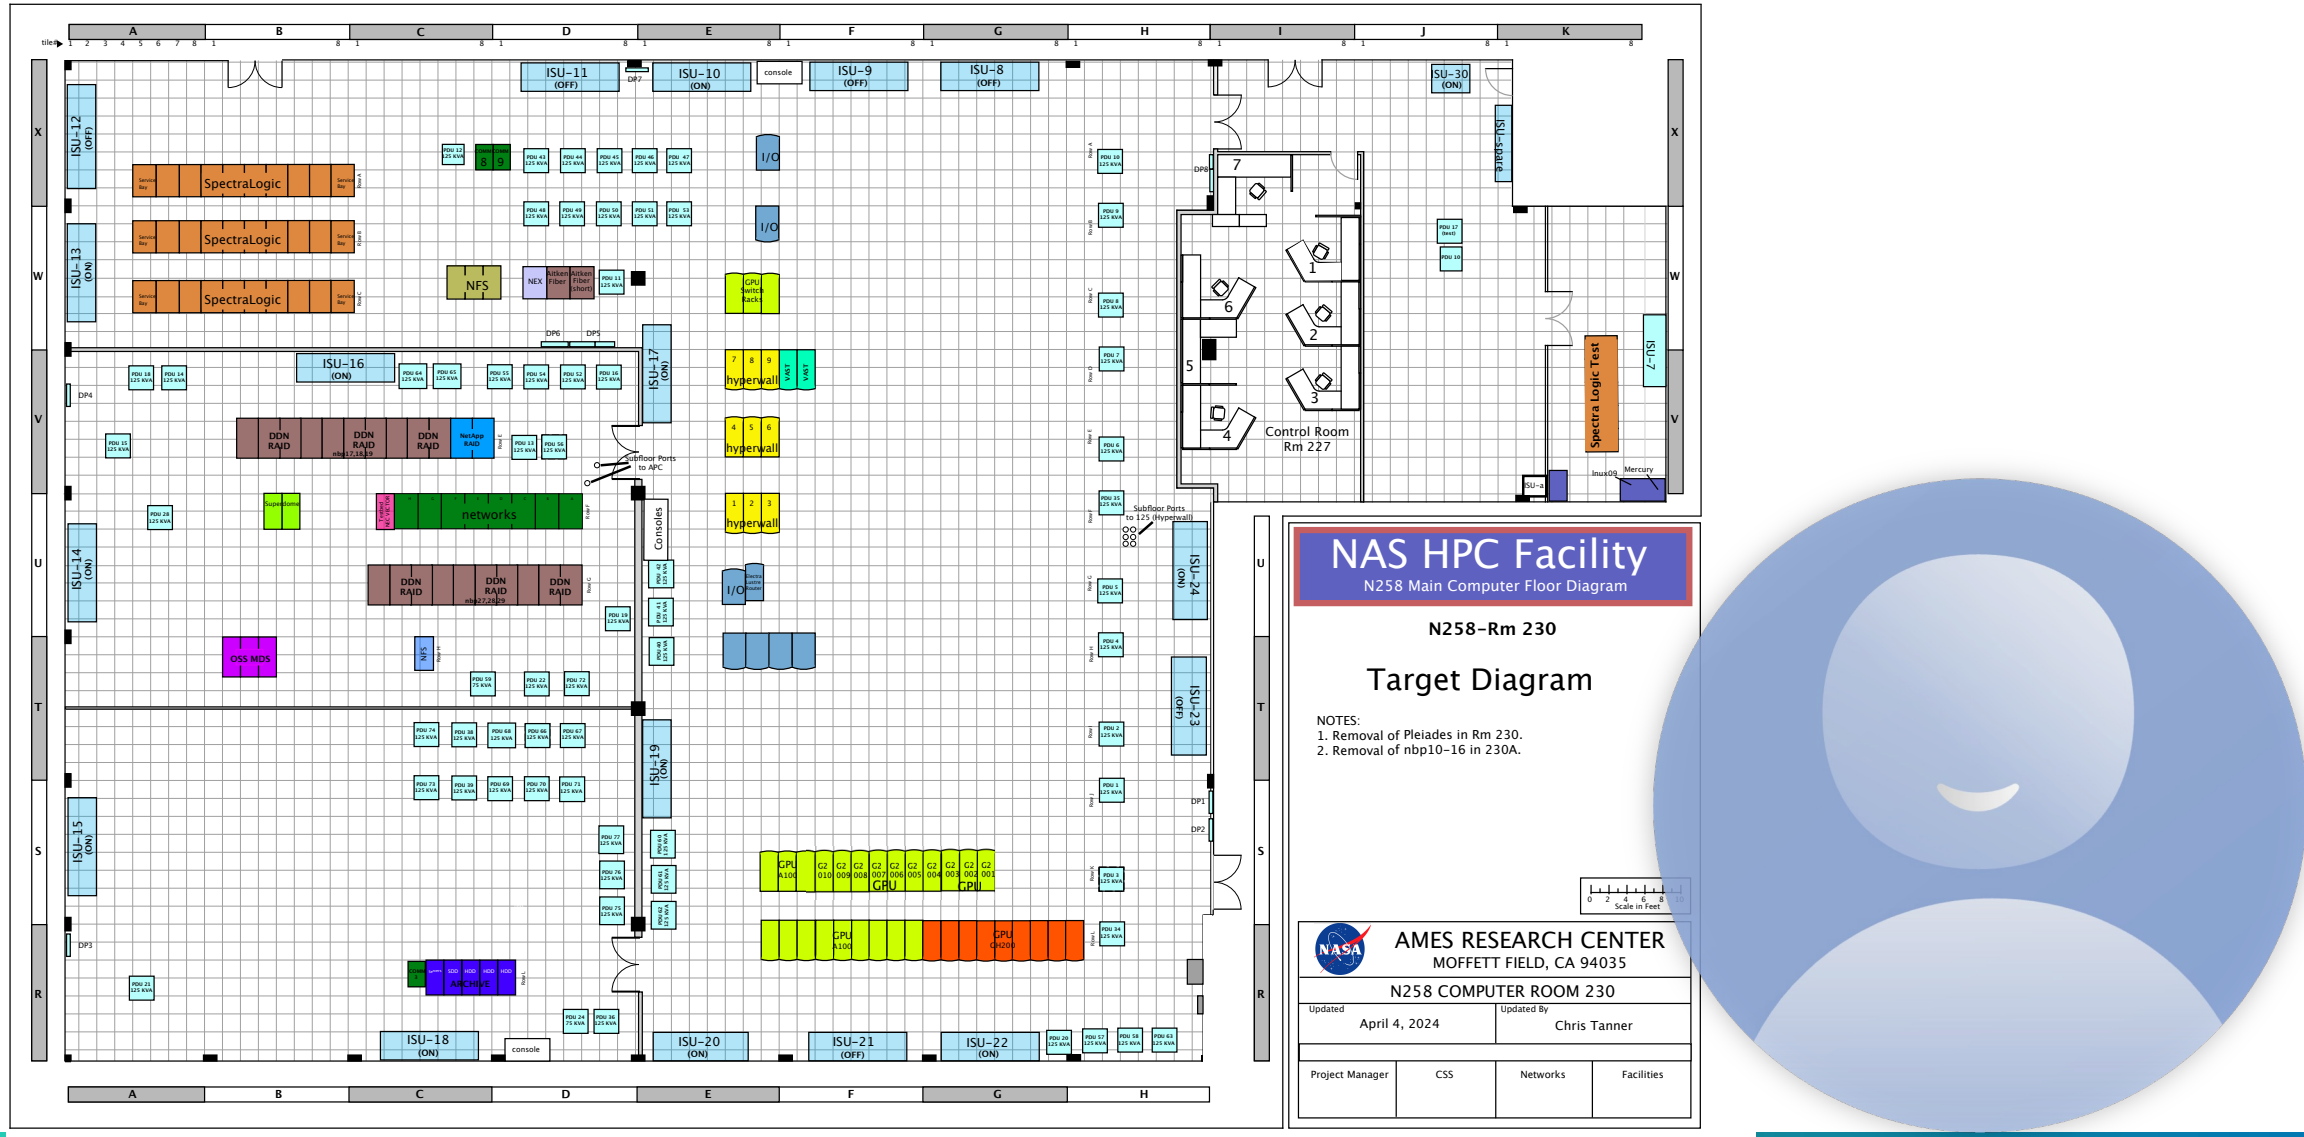
- The organic growth of the InfiniBand network has led to an unnecessarily complex network with tightly coupled unrelated networks.
- There is a network redesign that is ongoing that will separate the I/O network from all compute networks and remove dependencies in each compute network from other compute networks.
- This is highly simplified if Pleiades is removed from the equation.

Storage

- As in compute we have very old file systems that are problematic to operate and maintain. The 7 file systems provide 22.76 PB of storage.
- This storage will be made up with current storage and procurements in FY25 and FY26.



N258/230 Evolution



Installing New Equipment

Aitken Expansion

- Undergoing friendly user testing.
- 512 nodes
  - 2.6 Petaflops (PF) (15.7 PF total system)
  - 2,161 SBU/hour (Peak); 12,365 SBU/hr total system

Cabeus Expansion

- RFQ released – responses due 15 May – award 24 May.
- 300+ NVIDIA GH200 GPU nodes
  - 10.2 PF (12.2 PF total system)
  - 5,679 SBU/hour (Peak); 12,196 SBU/hr total system

Next Generation CPU Expansion

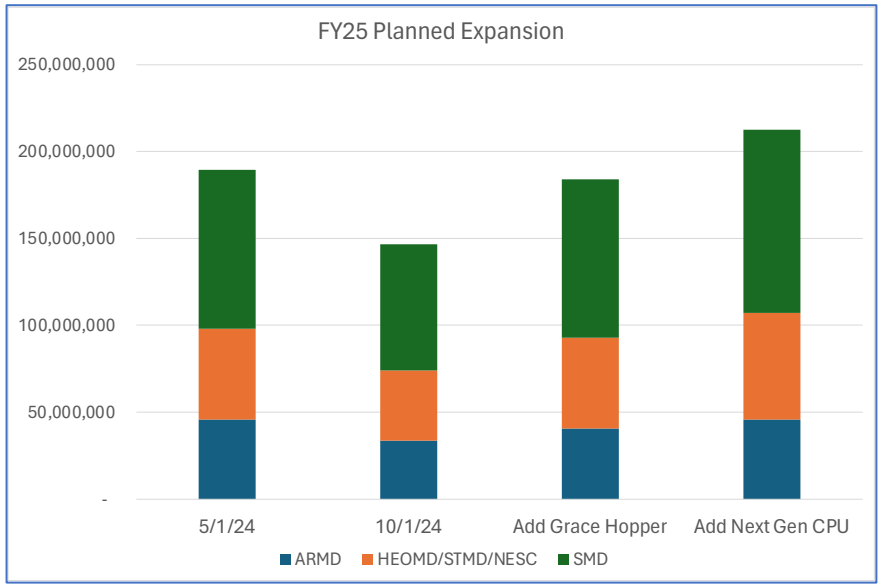
- RFP being prepared for late August/early September release. Award in early October.
- 512 nodes
  - More than 3.8 PF
  - More than 4,300 SBU/hour

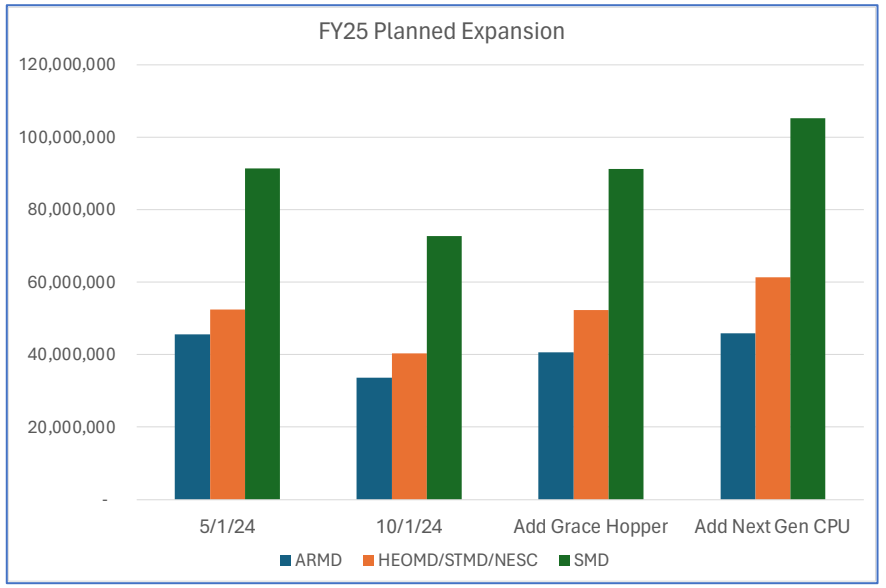
Storage

- RFP being prepared for late August/early September release. Award in early October.
- 20 Petabytes



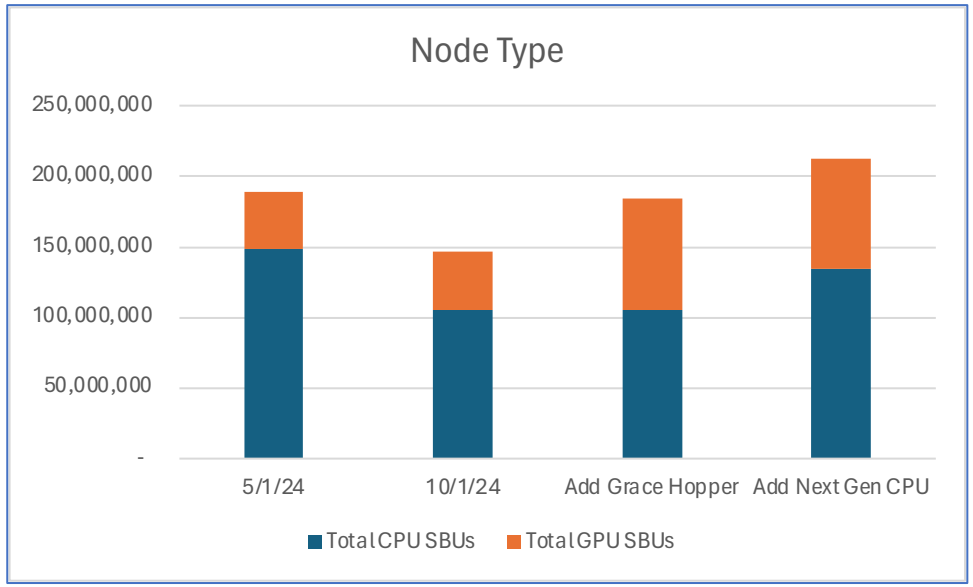
Impact on Total SBUs

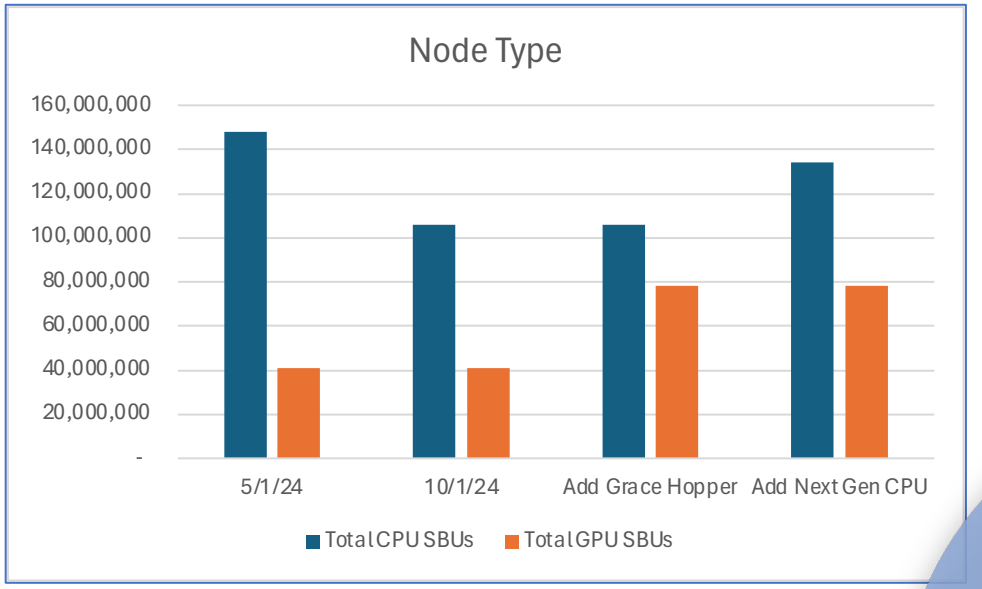


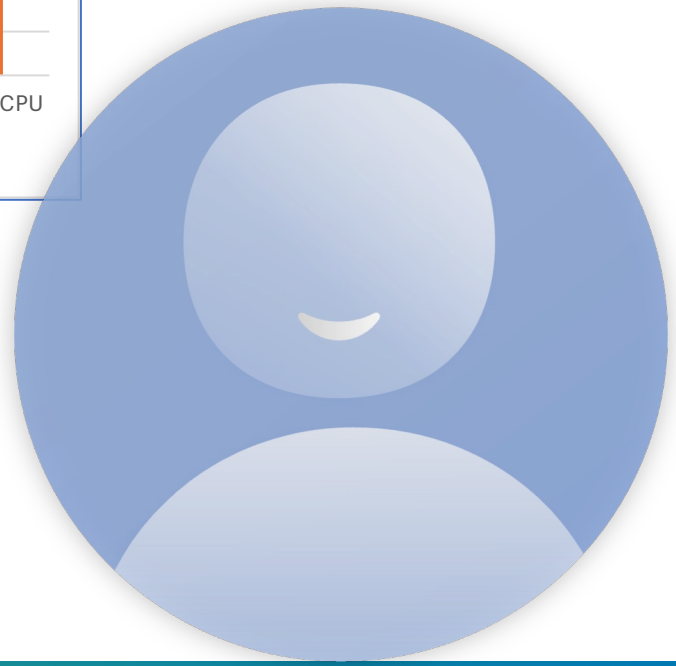




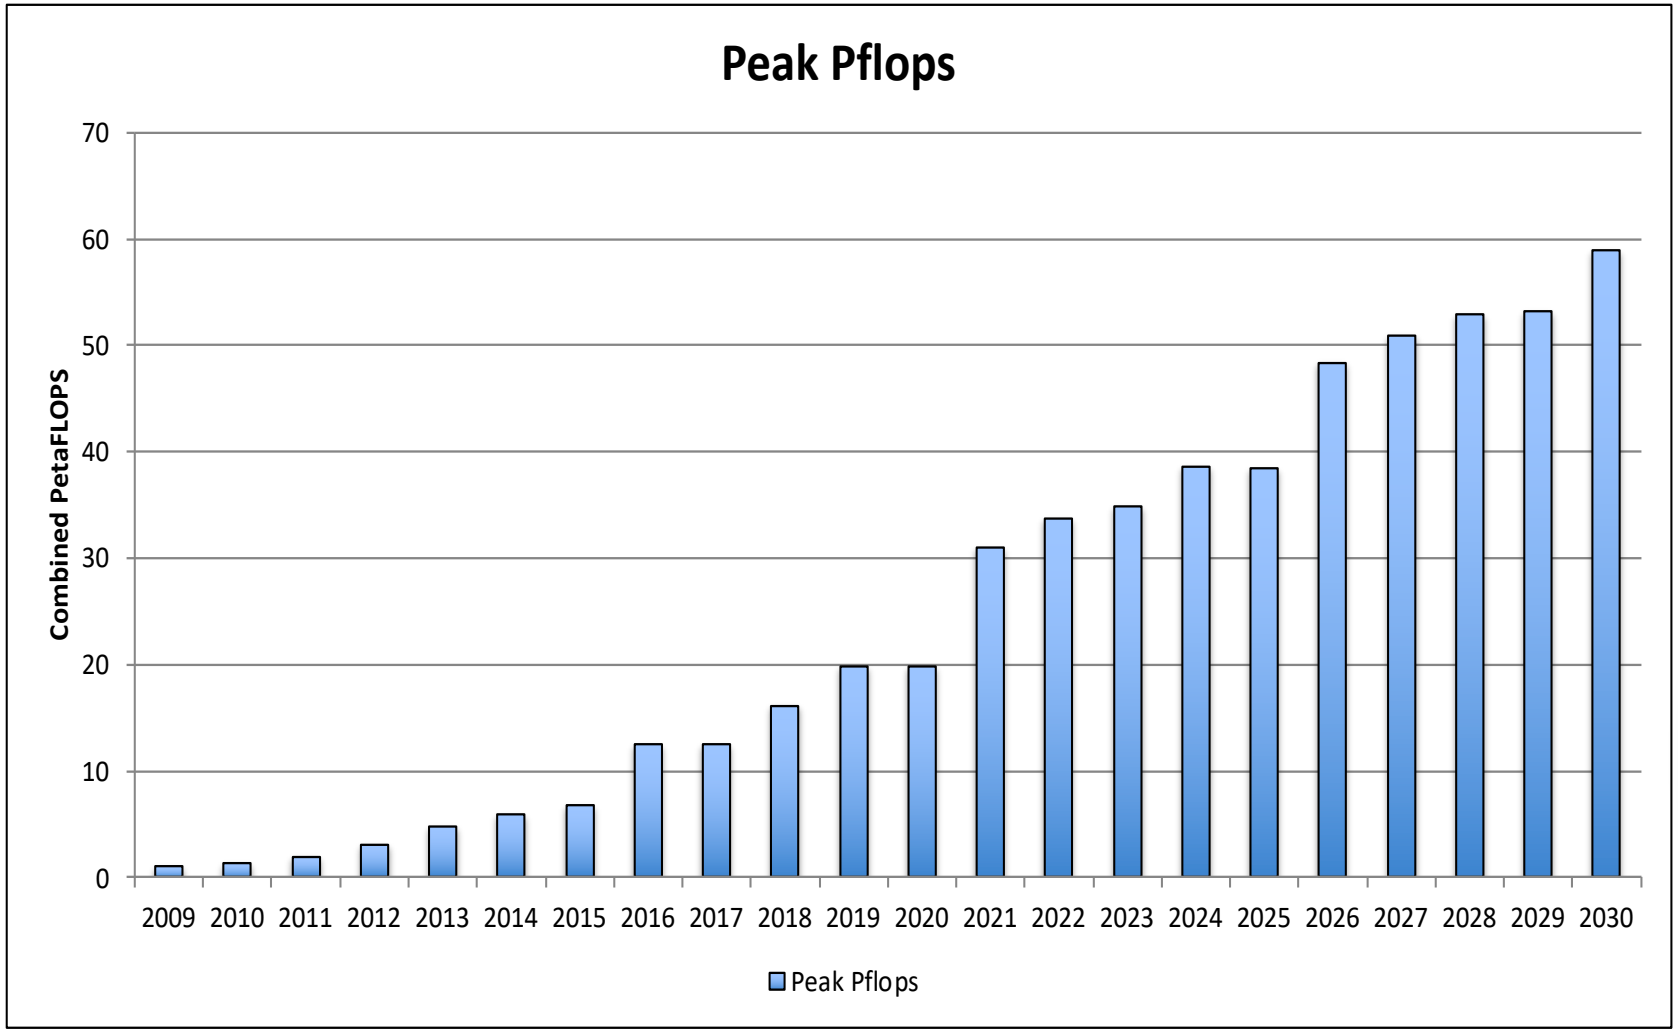
Impact on Total SBUs







Historic and Predicted Compute Growth



Totals through 2024 show systems that were or are installed. 2025 onward are based on baseline budget projections.



Migration to New Technology

GPU vs CPU SBUs

- The goal of the SBU rate is that an SBU represents a normalized unit of work. So, theoretically, it should take the same number of SBUs to get the desired output from a given input.
- This is in NO way a value measurement.

Cost / SBU

- Based on last 2 procurements with 5 years of operations.
  - A100: $0.0614 / SBU
  - Milan: $0.1350 / SBU

Watt / SBU

- Based on last 2 procurements with 5 years of operations.
  - A100: 52.83 W / SBU
  - Milan: 231.27 W / SBU

Summary

For users that can take advantage of the GPU systems, we are able to provide more than twice the capability while using less than a quarter of the power.



National Aeronautics and Space Administration

Ideas?
Questions?
Suggestions?

We want to hear from you!
support@nas.nasa.gov
1 (800)-331-8737

