NASA
Office of Inspector General

Audit of NASA's High-End Computing Capabilities

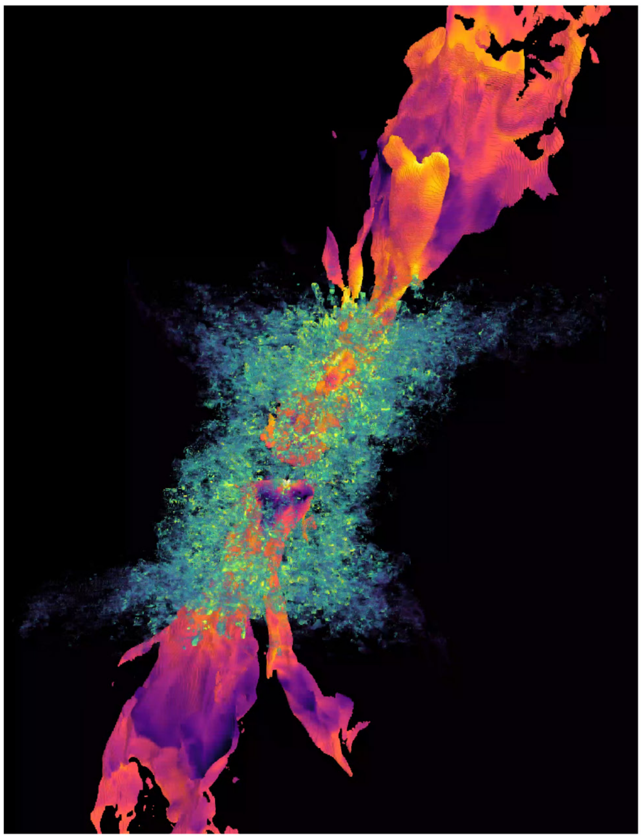

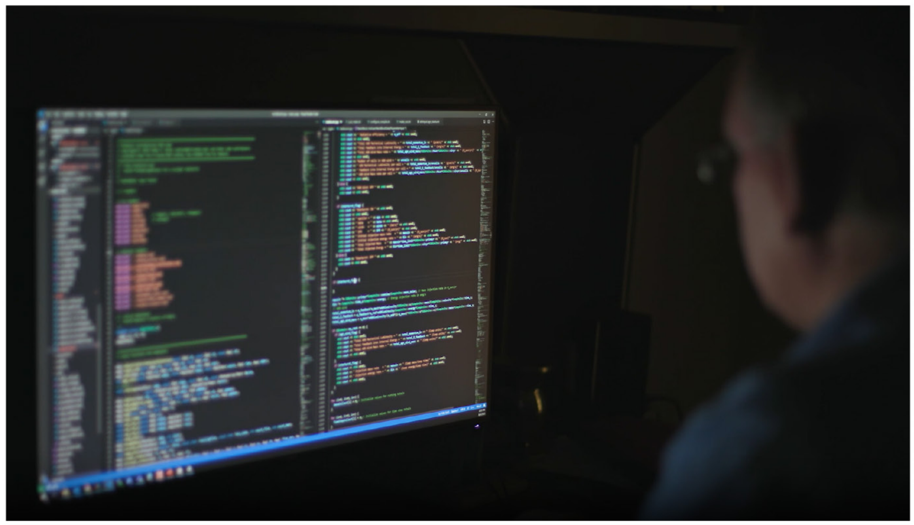

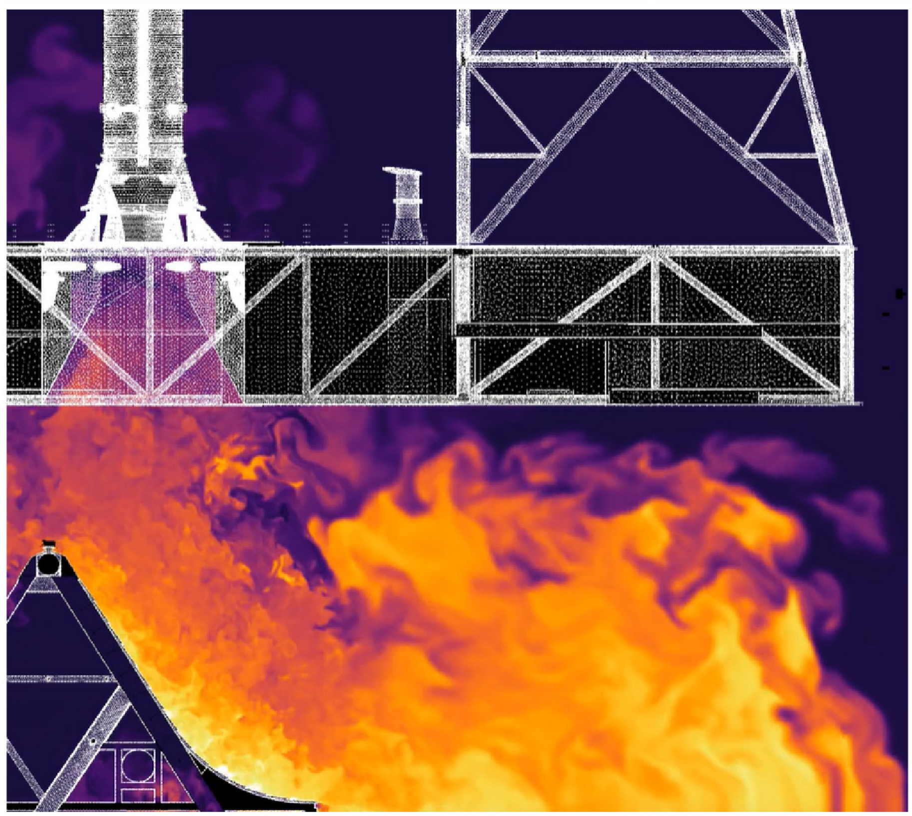

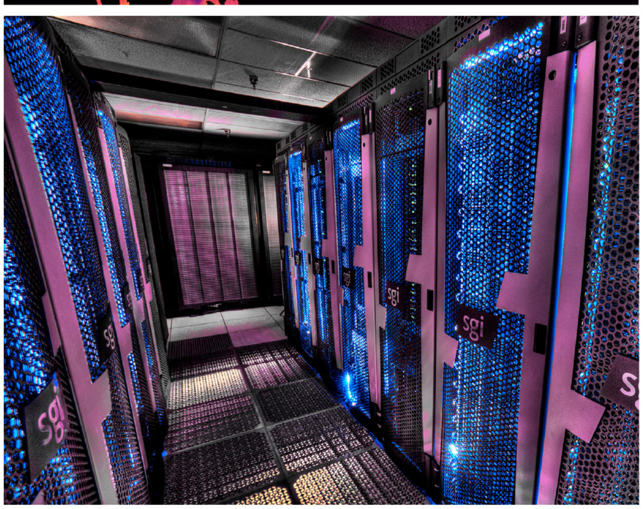

March 14, 2024

IG-24-009

Office of Inspector General

To report, fraud, waste, abuse, or mismanagement, contact the NASA OIG Hotline at 800-424-9183 or 800-535-8134 (TDD) or visit https://oig.nasa.gov/hotline.html. You can also write to NASA Inspector General, P.O. Box 23089, L'Enfant Plaza Station, Washington, D.C. 20026. The identity of each writer and caller can be kept confidential, upon request, to the extent permitted by law.

To suggest ideas or request future audits, contact the Assistant Inspector General for Audits at https://oig.nasa.gov/aboutAll.html.

COVER IMAGES

Top Left—A computer simulation of jets of particles and radiation, moving at near light speed, created by supermassive black holes.
Top Right— A NASA astrophysicist works on computer code to create the black hole simulations using NASA supercomputers.
Bottom Left—Two rows of Discover supercomputer racks at the NASA Center for Climate Simulation.
Bottom Right—A simulation of Space Launch System launch ignition NASA Kennedy Space Center's Launch Complex 39B.

RESULTS IN BRIEF

Audit of NASA's High-End Computing Capabilities

March 14, 2024 IG-24-009

WHY WE PERFORMED THIS AUDIT

Since its inception, NASA has pioneered many high-end computing (HEC) technologies and techniques that have become standard. HEC, or supercomputing, provides the critical processing power and time-saving capabilities that allow NASA to gain insight from large amounts of data that would take normal computers much longer to assess. A broad spectrum of employees, researchers, partners, external collaborators, and NASA Mission Directorates use the Agency's HEC capabilities. For example, NASA is currently using HEC capabilities to model the Agency's planned human landing on Mars, as well as to process and analyze the physics and environmental data critical to a successful landing. NASA's two main HEC facilities are the NASA Center for Climate Simulation at Goddard Space Flight Center and the NASA Advanced Supercomputing facility at Ames Research Center, but—thanks to remote access and cloud computing—NASA's HEC resources are used across NASA Centers and by authorized external partners around the world.

NASA manages HEC systems differently than its other computer systems. As one of five capability portfolios—a collection of functionally similar, site-specific capability components—NASA policy calls for HEC to be managed in an integrated manner and within budget constraints to meet certain requirements and strategic needs. Although HEC resource requirements are overseen by the Science Mission Directorate (SMD), each Mission Directorate has the autonomy to manage their own data and user access based on their individual requirements. Further, securing HEC systems is challenging due to their size; performance requirements; complex hardware, software, and applications; varying security requirements; and the nature of shared cyber resources.

In this audit, we assessed NASA's overall management of its HEC capabilities. Specifically, we focused on relevant policies, processes and controls, capacity planning, stakeholder engagement, and cybersecurity. We reviewed applicable policies and criteria; interviewed key officials, personnel, and stakeholders; evaluated capacity planning and success metrics; and reviewed techniques to identify and mitigate HEC cybersecurity risks. We met with officials and stakeholders from Headquarters and NASA Centers and benchmarked best practices with the Department of Energy and National Oceanic and Atmospheric Administration. We also participated in the National Institute of Standards and Technology (NIST) HEC security working group meetings to increase our comprehension of supercomputing technical architecture and cybersecurity techniques.

WHAT WE FOUND

Despite a history of innovation in HEC, NASA needs a renewed commitment and sustained leadership attention to reinvigorate its HEC efforts. Without key changes, the Agency's HEC is likely to constrain future mission priorities and goals. NASA's HEC is not managed as a program or centralized Agency strategic service; instead, resources are managed within the Earth Science Research Program within SMD, and this organizational placement hinders NASA's HEC efforts. One scientist within that Program is responsible for HEC capabilities at both HEC facilities, in addition to their Earth science responsibilities. This disjointed organization and management of HEC resources exacerbates several issues, including oversight, monitoring, and the foreign national accreditation access process. Although NASA has identified HEC as a capability portfolio, we found that key guiding documents and frameworks are absent, such as a management plan—an agreement detailing how the portfolio will be managed—and a commitment agreement designed to engage all relevant stakeholders and identify HEC as a strategic activity. Furthermore, while the Office of the Chief Information Officer (OCIO) has some oversight of HEC, it is not directly engaged in HEC activities or governance. Lastly, NASA is not keeping up with technological developments and advanced research computing requirements, in part due to these organizational and funding constraints.

NASA's HEC resources are oversubscribed and overburdened—in other words, Mission Directorates are requesting more computing time than existing capacity can provide. This scarcity drives schedule delays and often leads to NASA teams purchasing their own HEC resources to meet deadlines. For example, the Space Launch System team invests about $250,000 annually to purchase and locally manage their own HEC clusters rather than waiting for existing HEC resource availability. Agency officials told us that, except for Goddard Space Flight Center and Stennis Space Center, there are independent HEC assets installed at almost every NASA Center. NASA also lacks a comprehensive strategy for when to use HEC assets on the premises versus when to utilize cloud computing options—or a widespread understanding of the cost implications for each choice. Stakeholders told us that while they know NASA has HEC cloud computing options, they were hesitant to use them due to unknown scheduling practices or assumed higher costs.

NASA's decentralized HEC management also raises cybersecurity concerns. In addition to teams building their own HEC assets, OCIO-mandated cybersecurity controls are sometimes ignored or bypassed by Mission Directorates that view them as too stringent. OCIO's limited involvement with HEC system management can also result in duplicate spending (such as for software that OCIO already licenses) and difficulty in identifying and controlling access to HEC systems that are not included in OCIO's asset inventory tools. We also identified extensive use of NASA's HEC assets by external and foreign national parties without adequate user activity monitoring or a review process by security personal for gaining access to HEC systems. Finally, we found that individual Center HEC asset users are not steadily monitored, and there are no evaluations conducted to verify rights and accesses granted to international partners. Without an integrated HEC strategy and a more focused, security management approach, the Agency's trailblazing science and technology research will continue to be unnecessarily limited by NASA's disjointed HEC efforts.

What We Recommended

To establish executive leadership and strategically position NASA's HEC to meet the Agency's specialized needs, we recommended NASA's Associate Administrator (1) appoint executive leadership to determine the appropriate definition, scope, ownership, organizational placement, and structure for NASA's HEC. Additionally, we recommended that the Associate Administrator establish a tiger team to collaborate and strategize on HEC issues, including: (2) develop enterprise-wide stakeholder requirements to validate commitment agreements as required by policy; (3) identify technology gaps essential for meeting current and future needs and strategic technological and scientific requirements; (4) develop a strategy to improve prioritization and allocation of HEC assets, including on-premises versus cloud resources; (5) evaluate cyber risks to determine oversight and monitoring requirements; (6) implement an HEC classification designation for identifying HEC assets; (7) develop an inventory of enterprise-wide HEC assets; (8) document risk impact, classification, and categorization for all HEC jobs; and (9) identify and mitigate gaps in the foreign national accreditation access process.

We provided a draft of this report to NASA management who concurred with Recommendation 1 and partially concurred with Recommendations 2 through 9. NASA described planned actions to address Recommendation 1 and stated that a tiger team will be established to collaborate and strategize on HEC issues; subsequently, the tiger team will determine the implementation of Recommendations 2 through 9. While we consider management's comments responsive, OIG requests regular updates to understand the planned actions and timeline for implementation and to monitor progress towards implementation of Recommendations 2 through 9. The recommendations are resolved and will be closed upon completion and verification of the proposed corrective actions.

For more information on the NASA Office of Inspector General and to view this and other reports visit https://oig.nasa.gov/.

TABLE OF CONTENTS

Acronyms

AI	artificial intelligence
ARMD	Aeronautics Research Mission Directorate
AWS	Amazon Web Services
CFD	Computational Fluid Dynamics
CP	capability portfolio
CPU	central processing unit
DOE	Department of Energy
ESDMD	Exploration Systems Development Mission Directorate
FLOPS	floating point operations per second
GAO	Government Accountability Office
GPU	graphics processing unit
HEC	high-end computing
IT	information technology
NAS	NASA Advanced Supercomputing
NCCS	NASA Center for Climate Simulation
NIST	National Institute of Standards and Technology
NOAA	National Oceanic and Atmospheric Administration
NPR	NASA Procedural Requirements
NSCI	National Strategic Computing Initiative
OCIO	Office of the Chief Information Officer
OIG	Office of Inspector General
RISCS	Risk Information Security Compliance System
SLS	Space Launch System
SMCE	Science Managed Cloud Environment
SMD	Science Mission Directorate
TCP/IP	Transmission Control Protocol/Internet Protocol

INTRODUCTION

Serving a broad spectrum of employees, customers, researchers, partners, and external collaborators, NASA's high-end computing (HEC) capability—previously known as "supercomputers"—provides a comprehensive set of resources and services for the Agency's Mission Directorates as well as the NASA Engineering and Safety Center.[1] Because HEC environments differ from traditional information technology (IT) systems, it is more difficult to strike the right balance between scientific and engineering requirements and cybersecurity protections.

Defining High-end Computing

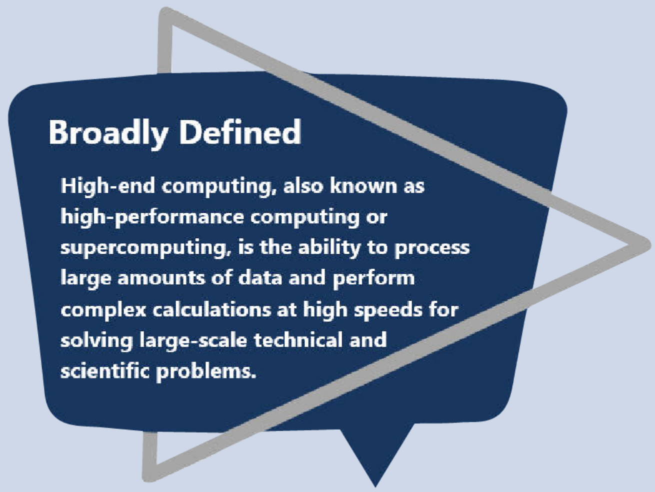

Source: NASA

Analyzing immense amounts of data would be impossible without HEC capabilities—a typical computer simply cannot process data rapidly enough to meet Agency scientific and engineering demands. For example, HEC capabilities enabled NASA to ingest and analyze months of Kepler observation data to more quickly identify evidence pointing to the existence of exoplanets—many located in the habitable zone.[2] Currently, NASA is using HEC capabilities to model the Agency's planned human landing on Mars, as well as to process and analyze the physics and environmental data critical to a successful atmospheric entry and landing on the Martian surface.

Because HEC systems transfer user data into the HEC environment, the system owner is responsible for ensuring the confidentiality, integrity, and availability of that data by employing effective cybersecurity controls.[3] However, traditional cybersecurity tools can be either incompatible or overly disruptive in an HEC environment. As a result, cybersecurity considerations may take a back seat to the processing speed and system capabilities needed for scientific research.

---

[1] For this report, we broadly define high-end computing, also known as high-performance computing or supercomputing, as the ability to process large amounts of data and perform complex calculations at high speeds for solving large-scale technical and scientific problems. In this document, we will use the term HEC throughout.

[2] The Kepler Mission, launched on March 6, 2009, was the first space mission dedicated to the search for Earth-sized and smaller planets in the habitable zone of other stars in our neighborhood of the galaxy. An exoplanet is any planet beyond our solar system. A habitable zone is the distance from a star at which liquid water could exist on orbiting planets' surfaces. Analyses typically take about 3 days to run on HEC versus more than a month on Kepler Science Operations Center computers.

[3] An HEC environment is a system that connects scientific computing users, data, applications software, and middleware and integrates them into a single research environment. Data transfer is the process of using computing techniques to move data from one computer to another. A cybersecurity control is a safeguard or countermeasure designed to protect the confidentiality, integrity, and availability of information and to meet defined security requirements.

In this audit, we assessed NASA's overall management of its HEC capabilities. Specifically, we focused on relevant policies, processes and controls, capacity planning, stakeholder engagement, and cybersecurity. See Appendix A for details on the audit's scope and methodology.

Background

In a complex, rapidly changing world, government agencies increasingly rely on HEC to manage and process massive volumes of data to solve mission-critical challenges. With the ability to process large amounts of data and perform calculations at high speeds, HEC helps solve large-scale technical and scientific problems, enabling researchers to study subjects that would otherwise be impractical, or impossible, to investigate due to their complexity or the danger they pose. In recent years, HEC has helped the federal government develop treatments for COVID-19, advance scientific research and discovery, and conduct high fidelity operational simulations.

Designed in 1964 by Seymour Cray—known as the "father of supercomputing"—the Control Data Corporation 6600 is considered the first high-end computer. Today, the 6600 is dwarfed when compared to the computing power of common smartphones.

Assessing computing power is highly complex. To simplify for illustration and ease of understanding: the computer's processor clock speed determines how quickly the central processing unit can retrieve and interpret instructions. Today, mobile devices have more computing power than HEC assets through the turn of the 21st century—including the computational ability of the Apollo 11 Guidance Computer that helped put astronauts on the Moon more than half a century ago. See Figure 1 for an illustration comparing processing power through the years and Appendix B for additional information.

Figure 1: Example Comparing Computer Processing Power

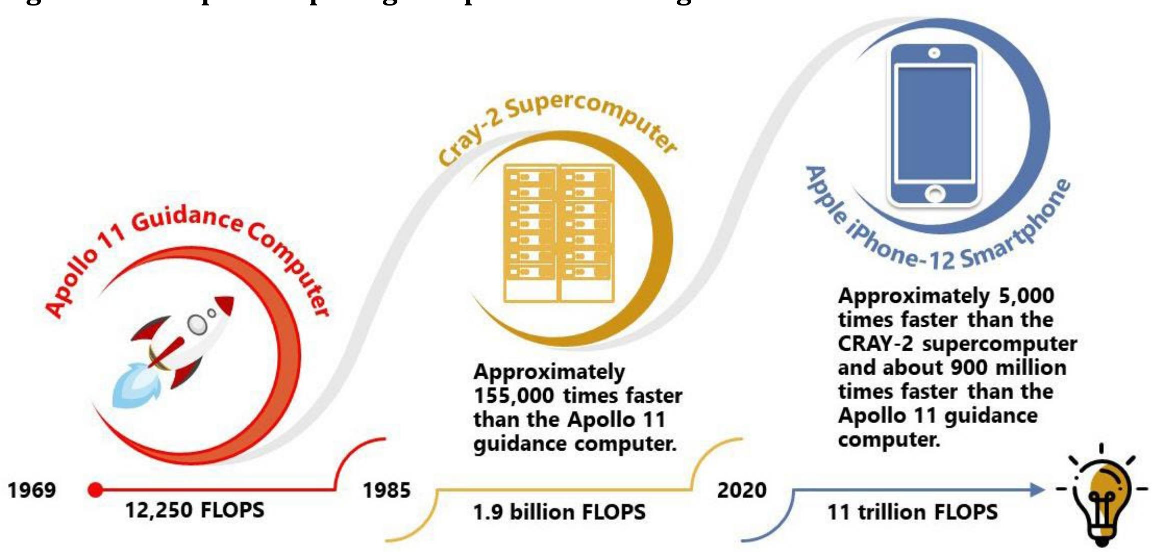

Source: NASA Office of Inspector General (OIG) presentation of Adobe information.

Note: "FLOPS" stands for floating point operations per second.

However, while interesting, this comparison is not precise. It's like comparing the first airplanes designed by the Wright Brothers and a fighter jet—both could fly, but the two are, technologically, worlds apart. Making a side-by-side comparison of computing power is difficult because there are many ways to measure computational performance, such as the speed at which a system calculates floating point operations per second, known as FLOPS, or the speed at which a system can run graphics-intensive applications. Performance also depends on several other factors, such as the amount and speed of memory, network performance, and how well the computer code utilizes the system hardware.

In 1991, when the Federal High-Performance Computing and Communications initiative began ramping up, NASA was tasked with conducting basic and applied research in networking and information technology, particularly in the field of computational science, with emphasis on aerospace sciences, Earth and space sciences, and remote exploration and experimentation.[4]

In 2015, an Executive Order established the National Strategic Computing Initiative (NSCI) and a whole-of-government HEC strategy.[5] The NSCI, in collaboration with industry and academia, led to the creation of a cohesive, multi-agency vision and federal investment to maximize the benefits of HEC for the United States. As such, NASA was established as a deployment agency—meaning that NASA, along with other deployment agencies, participates in the co-design process to integrate the special requirements of its missions and to influence the early design stages for new HEC systems, software, and applications.[6]

In 2021, the Government Accountability Office (GAO) conducted a review of agencies charged with HEC responsibilities in the NSCI Executive Order.[7] Although GAO noted that NASA and the other agencies made advances in HEC research, development, and deployment activities, stakeholders cited ongoing challenges related to the uncertainty over how to meet future funding needs, cloud utilization, and software maintenance.

Since the NSCI was established, the use of HEC for modeling, simulation, and artificial intelligence (AI) has grown substantially.[8] NASA, for example, uses a combination of high-end computers, AI, and satellite data to process extremely large volumes of imagery transmitted by the Transiting Exoplanet Survey Satellite.[9] Once the high-end computer processes the raw data, the data is used to investigate vast regions of space for valuable scientific data hidden within multiple star systems. As more agencies rely on HEC to process their data, these systems become high-profile targets for attackers. To protect information, the *Executive Order on Improving the Nation's Cybersecurity* mandates that data be encrypted in transit and at rest.[10]

---

[4] Public Law 102-194, *High-Performance Computing Act of 1991* (December 9, 1991).

[5] Executive Order 13702, *Creating a National Strategic Computing Initiative* (July 29, 2015).

[6] The Department of Defense, Department of Energy, and the National Science Foundation are NSCI lead agencies in developing the next generation of integrated HEC hardware and software capability, as well as workforce development. In addition to NASA, the other four NSCI deployment agencies are the Federal Bureau of Investigation, the National Institutes of Health, the Department of Homeland Security, and the National Oceanic and Atmospheric Administration.

[7] GAO, *High-Performance Computing Advances Made Towards Implementing the National Strategy, But Better Reporting and A More Detailed Plan Are Needed* (GAO 21-104500, September 2021).

[8] AI is generally thought of as the capability of a machine to imitate intelligent human behavior.

[9] The Transiting Exoplanet Survey Satellite is a NASA mission to discover Earth-size worlds around nearby stars.

[10] Executive Order 14028, *Improving the Nation's Cybersecurity (*May 12, 2021). Encryption, at its most basic level, is the process of protecting information or data by using mathematical models to scramble it in such a way that only the parties who have the key can unscramble it. As the name implies, data in transit is data moving from one location to another. Data at rest refers to inactive data, meaning it is not moving between devices or networks.

Cybersecurity Versus Performance is a Constant Struggle

As the global cyberthreat environment grows, protecting critical systems and sensitive information from digital attacks can overwhelm even the most sophisticated cybersecurity professionals. This is equally true in HEC environments, especially those providing a user front end (computer interface) that is exposed to the internet and its common threats of vulnerabilities and attacks. Securing HEC systems is challenging due to their size; performance requirements; diverse and complex hardware, software, and applications; varying security requirements; and the nature of shared cyber resources.

Mirroring the commercial industry, government HEC threats break down into four themes: exploiting humans, exploiting software, exploiting protocols, and insecure design. Exploiting humans includes issues like phishing attacks and insider threats.[11] Exploiting software includes taking advantage of bugs that, if exploited, can lead to memory being overwritten to redirect the flow of a program or information to an attacker. Exploiting protocols includes things such as brute-force password cracking attacks.[12] HEC users often consider cybersecurity valuable only as long as it does not significantly slow down system performance and impede research.

Underscoring the importance of HEC cybersecurity, in February 2023, the National Institute of Standards and Technology (NIST) released inaugural draft guidance to standardize and facilitate the sharing of HEC security information and knowledge through the development of an HEC system reference model.[13] The reference model divides an HEC system into four function zones: high-performance computing, data storage, access, and management. Because the NIST draft guidance is currently under review, the lack of standards contributes to the challenges of consistently applying cybersecurity in an HEC environment.

HEC Is Key to NASA's Mission Success

For over three decades, NASA's use of HEC resources has been a major and growing factor in the effectiveness and execution of its missions. For example, in October 2004 NASA's high-end computer Columbia, located at Ames Research Center (Ames), gained worldwide recognition as the most powerful computer in the world at the time. The combined speed and productivity increased the Agency's HEC capability tenfold for missions in aeronautics, space exploration, and Earth and space sciences. Notably, Columbia—named to honor the crew of the Space Shuttle Columbia STS-107, lost in 2003—was critical to NASA's post-Columbia return to flight effort and for near-real-time simulations that contributed to the safety of subsequent shuttle missions. The HEC Columbia remained active until early 2013 when it was phased out to make way for newer, upgraded HEC assets. See Appendix C for additional information on how various NASA missions use HEC to assist their work.

In a recent publication NASA described five uses of HEC related to mission success:

[11] Phishing is an attempt by an individual or group to solicit personal information from unsuspecting users by employing social engineering techniques. Phishing emails are crafted to appear as if they have been sent from a legitimate organization or known individual and entice users to click on a link that will take them to a fraudulent website that also appears legitimate. Cybersecurity threats posed by an organization's employees and contractors are commonly referred to as insider threats.

[12] A protocol is a set of rules or procedures for transmitting data between electronic devices, such as computers.

[13] NIST Special Publication 800-223, Draft, *High-Performance Computing (HPC) Security* (February 2023).

1. Designing safe, efficient air taxis, such as the Urban Air Mobility tiltwing single-main rotor concept vehicle, by simulating aerodynamic performance.[14]
2. Keeping planetary rovers safe during risky landings by simulating and analyzing various scenarios of supersonic parachute inflation.
3. Modeling spacecraft heat shield materials at the microscale to generate high-resolution 3D images.
4. Predicting weather and climate to keep humans safe by simulating weather and climate events.
5. Exploring the past, present, and future of planets inside and outside of our solar system by simulating the climates of planets.

NASA's HEC Helps Battle COVID-19

In March 2020, as the world came to a standstill due to the COVID-19 pandemic, HEC gained recognition for its role in helping accelerate science, therapeutics development, and COVID-19-related patient care. Recognizing the potential of HEC to accelerate progress in developing a vaccine for COVID-19, the HEC community—including NASA—created a public-private consortium between government, industry, and academia to aggregate compute time and resources on consortium member's HEC assets and to make them available to assist in researching the virus. Acting as a single point of access, the COVID-19 High-Performance Computing Consortium provided scientists state-of-the-art HEC resources to accelerate and enable pandemic research and response.[15]

As a consortium member, NASA assessed research proposals for technical feasibility and benefits, then matched approved proposals to resource providers based in part on HEC requirements. For example, using computing resources at the Ames facility, scientists from around the United States conducted research to understand the virus and develop treatments and vaccines. Running computationally intensive data derived from COVID-19 patient samples, NASA HEC assets identified potential biomarkers indicative of disease severity.[16] The results of this analysis allowed for identification of unique human sequence variation associated with high risk of morbidity for COVID-19 patients. Likewise, NASA HEC assets were used for COVID-19 environmental impact studies. As part of Earth science pandemic research efforts, one such study involved analyzing air pollution reduction due to a decrease in airline flights.

NASA Pioneered Many HEC Firsts

NASA builds its HEC systems differently than its counterparts at national labs and research universities by playing the long game and adding performance over time. Instead of commissioning a monolithic machine with a defined operational period, NASA uses a modular strategy, gradually adding additional computing capacity. For example, the HEC Aitken has been upgraded yearly since 2019, including

[14] Urban Air Mobility vehicles support vertical lift research aimed at capabilities and technologies that will eliminate barriers for quiet, safe, efficient, autonomous vehicles operating in both urban and rural environments. Using NASA's powerful supercomputers, researchers are simulating the aerodynamic performance to understand complex airflow around aircraft components.

[15] NASA, along with the Office of Science and Technology Policy, Department of Energy, National Science Foundation, IBM, Massachusetts Institute of Technology, Rensselaer Polytechnic Institute, Amazon Web Services, Hewlett Packard Enterprise, Google, and Microsoft joined the High-Performance Computing Consortium as founding members.

[16] A biomarker is a biological molecule found in blood, other body fluids, or tissues that is a sign of a normal or abnormal process, or of a condition or disease. It may be used to see how well the body responds to a treatment for a disease or condition.

recently in November 2023. By upgrading its HEC assets over time, NASA attempts to keep pace with the technology needed to tackle complex subjects such as Artemis missions and simulating weather on Mars.[17]

Since its inception, NASA has pioneered many technologies and techniques that have become standards within the HEC ecosystem.[18] For instance, in 1994 the first Beowulf Linux commodity cluster was constructed at NASA's Goddard Space Flight Center (Goddard) with its origins part of HEC folklore.[19] Beowulf linked personal computers, Linux, ethernet, and network drivers into a supercomputer that accelerated the global scientific community's move away from expensive, proprietary supercomputer systems and toward the adoption of computer clusters and open-source software. This method of building HEC clusters continues today. Currently, most of the world's top 500 computer systems use NASA's Beowulf approach. Significantly, in April 2022 Beowulf cluster computing was inducted into the Space Technology Hall of Fame.

NASA's architectural and operational models also led the world in a series of firsts, such as:

- First to put the UNIX operating system on HEC.
- First to implement TCP/IP networking in HEC environment.[20]
- First to link HEC hardware and workstations together to distribute computation and visualization (what is now known as client/server).
- Developed Aeronet, the first high-speed wide-area network connecting HEC resources to remote customer sites.
- Developed the first UNIX-based hierarchical mass storage system.

Figure 2 highlights select milestones in NASA's HEC evolution over more than 30 years.

Figure 2: Select NASA HEC Milestones (as of June 2023)

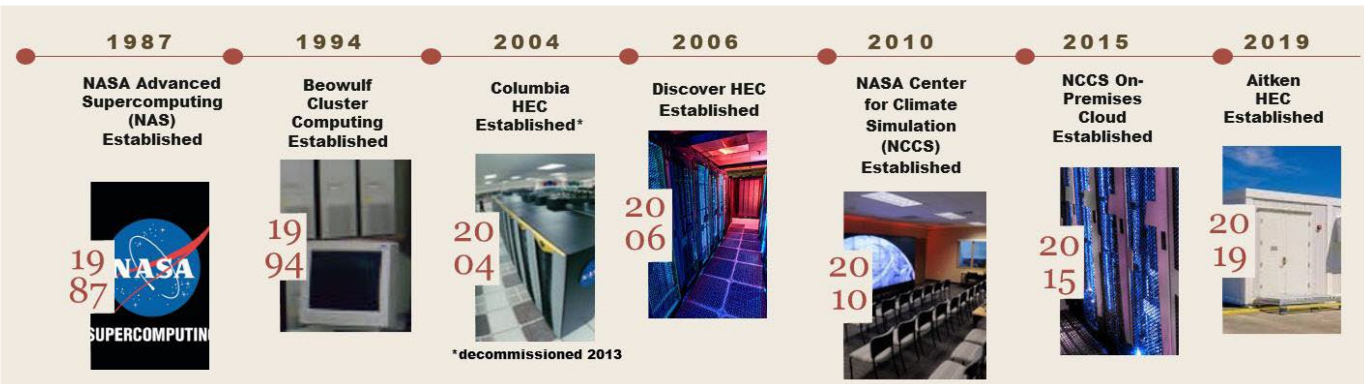

Source: NASA OIG presentation of Agency information.

[17] NASA's Artemis campaign seeks to return humans to the Moon's surface in 2025 before sending crewed missions to Mars in the 2030s.

[18] An HEC ecosystem includes servers, storage, software, and technical support.

[19] A computer cluster is a set of connected computers (nodes) that work together as if they are a single (much more powerful) machine.

[20] TCP/IP stands for Transmission Control Protocol/Internet Protocol—a suite of communication protocols used to interconnect network devices on the internet or intranet.

NASA's HEC Facilities

Leveraging on-premises HEC at two locations—the NASA Advanced Supercomputing (NAS) facility at Ames in California and the NASA Center for Climate Simulation (NCCS) at Goddard in Maryland—provide a range of assistance from designing safer aircraft and space vehicles to solving challenging weather and climate prediction problems.[21] At both the Ames and Goddard facilities, computing time is available to NASA-sponsored scientists or engineers for high-performance processing and analysis using the HEC assets. The Ames facility has four HEC assets while the Goddard facility houses one. Figure 3 provides additional information on NASA's HEC assets.

Figure 3: NASA HEC Assets (as of June 2023)

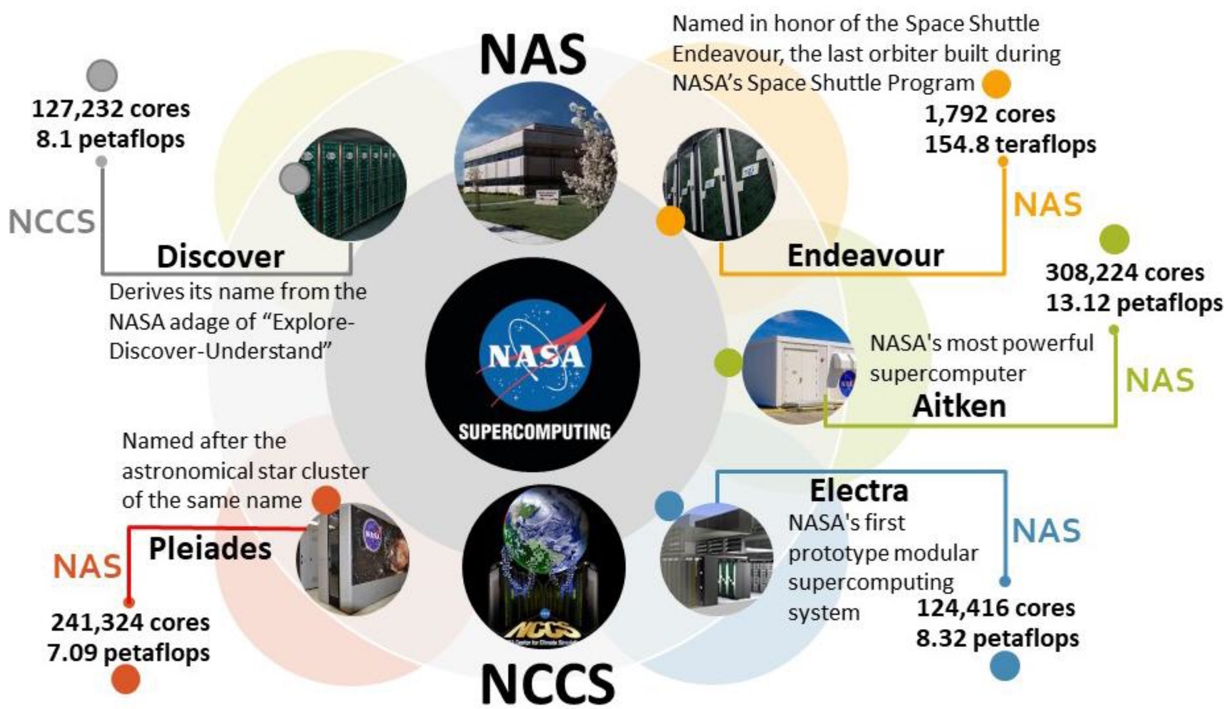

Source: OIG presentation of NASA information.

Note: A core is essentially a small central processing unit (CPU) built into a big CPU. It can independently perform or process all computational tasks. A teraflop is a rate of computing speed that achieves one trillion FLOPS. One petaflops equals one quadrillion operations per second and represents an extremely fast computing speed for a single machine.

In the early 1980s, Congress established the NAS facility at Ames with a charter to provide HEC capabilities for numerical simulations of proposed commercial and military aircraft designs. NASA's premier HEC facility began operation in 1987, and today NAS continues advancing key technologies in aeronautics and space explorations, such as rotorcraft (helicopter) stall dynamics and asteroid impact simulations for planetary defense.

[21] On-premises, also known as "on-prem," simply means that the IT infrastructure is hosted on-site.

As of June 2023, the four HEC systems at the NAS facility—Aitken, Electra, Pleiades, and Endeavour—are used by 1,885 researchers from NASA Centers, universities, and industry to support the Agency's mission to explore space and to understand our planet.[22]

Likewise, Goddard's HEC facility—with 1,410 current users—has provided resources to NASA scientists and engineers for more than 30 years. With an emphasis on climate simulation, the NCCS analyzes and models data on winds, clouds, precipitation, and atmospheric pollutants to enable scientific discoveries that will benefit humankind. As the centerpiece of Goddard's NCCS, the Discover HEC is capable of calculations at nearly 8.1 petaflops.[23] Discover is particularly suited for large, complex, communications-intensive problems employing large matrices and science applications.

As of November 2023, the on-premises HEC systems at Ames were ranked among the world's top 500 systems—Aiken #85, Pleiades #132, and Electra #143, respectively. NAS's Endeavour and Goddard's Discover on-premises HEC are not ranked in the *TOP500*.[24] See Appendix D for additional details.

Additionally, to augment NASA's on-premises HEC resources, both the NAS and NCCS offer commercial HEC cloud services provided by Amazon Web Services (AWS) on a pay-for-use basis.[25] Pricing depends on computing processor time, space used, the frequency of request, and the amount of data transferred out of AWS. The NAS HEC commercial cloud, launched in August 2019, allows NASA scientists, engineers, and collaborators to perform workloads that require additional resources to complement on-premises resources. Likewise, through the NCCS, the Science Managed Cloud Environment (SMCE) provides access to the commercial HEC cloud for NASA researchers and their collaborators. Managed by personnel at Goddard, SMCE utilizes commercial AWS by enabling access to HEC cloud resources for rapid prototyping and collaboration. The SMCE is specifically designed for data in a Federal Information Security Modernization Act low-security environment providing a "ready to go" cloud science analytics platform with access to high-performance computing.[26]

NASA's HEC Organizational Structure and Users

With a fiscal year 2023 budget of about $74 million, along with additional funding received from each of the Mission Directorates, the HEC capability is governed by the Mission Support Council and managed by the Agency's Science Mission Directorate (SMD).[27] The use of directorate funds enables HEC assets to operate as an Agency-wide resource. While SMD oversees and coordinates resource requirements (time and capacity) for all of NASA, each of the Mission Directorates has the autonomy to manage their allocated time, capacity, data, and user access based on project requirements.

---

[22] NAS deployed a fifth HEC system, Cabeus, after the completion of our review. Cabeus is not subject to our review.

[23] To match what a 1-petaflops computer system can do in just one second, you would have to perform one calculation every second for 31,688,765 years.

[24] Begun in 1993 and updated every 6 months, the *TOP500* provides a ranked list of the most powerful general purpose computer systems in common use for high-end applications.

[25] Cloud computing services is the practice of using a network of remote servers hosted on the internet for centralized data access and storage to computer services or resources such as Microsoft Office 365 or Google Apps.

[26] According to the Federal Information Security Modernization Act of 2014, Pub. L. No. 113-283, (2014), if the loss of confidentiality, integrity, or availability is projected to have a limited harm on the organization's activities, assets, or persons, the potential impact is considered low.

[27] The Mission Support Council serves as the Agency's senior decision-making body regarding the integrated mission support portfolios such as IT, capability portfolios, and associated investments.

More than a decade ago, the Agency identified HEC as one of five capability portfolios (CP)—a collection of functionally similar, site-specific capability components—in accordance with NASA policy.[28] NASA Procedural Requirements (NPR) 8600.1 provides guidance on how to strategically and centrally manage CPs that are determined to be of importance to the Agency and calls for HEC to be managed in an integrated manner and within budget constraints to meet current and future requirements and strategic needs. Below, Figure 4 displays the HEC organizational structure.

Figure 4: NASA's HEC Organizational Structure (as of December 2023)

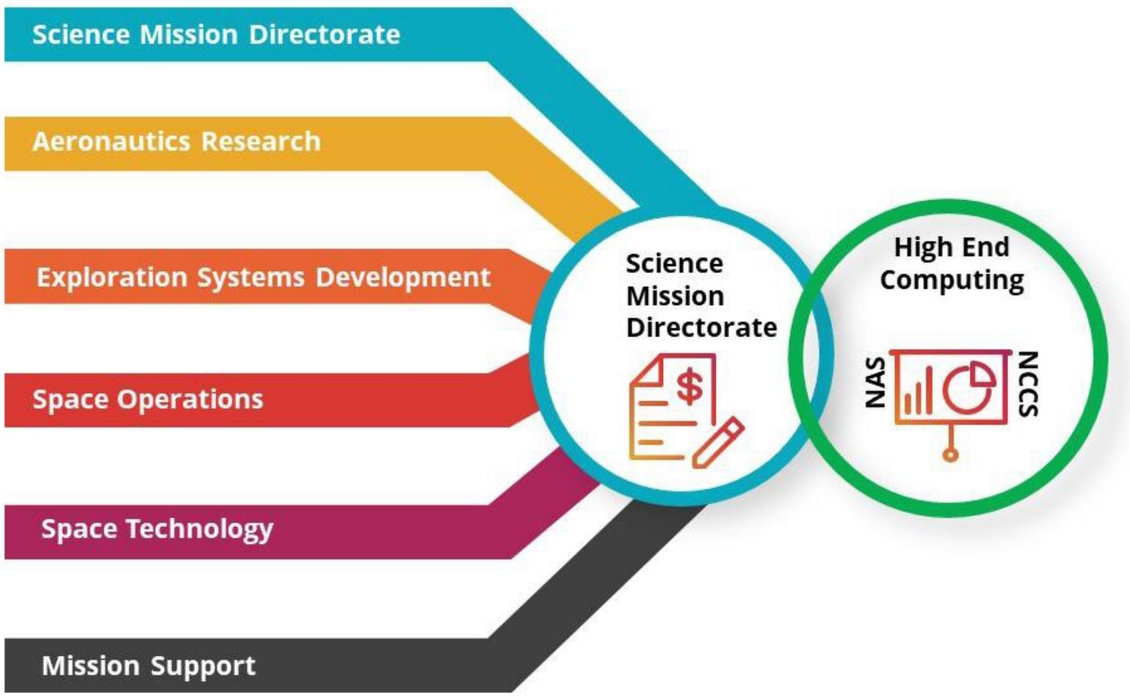

Source: NASA OIG presentation of the HEC organizational structure.

During fiscal year 2023, NASA's HEC resources were utilized by more than 3,000 users—of which more than half were attributed to external parties consisting of authorized foreign nationals and educational partnerships. Users typically participate in NASA-affiliated research, research in partnership with NASA or research that utilizes NASA data. Scientists and researchers who are considered foreign nationals—anyone who is neither a U.S. citizen, a U.S. lawful permanent resident, nor a U.S. protected individual—must adhere with U.S. export control policies and regulations. Generally, Agency HEC use by approved foreign nationals is permitted regardless of whether access to HEC resources is from within the U.S. or abroad, and according to Agency officials, NASA follows federal mandates and restrictions on foreign nationals.

[28] NPR 8600.1, *NASA Capability Portfolio Management Requirements* (April 22, 2019). NASA's CPs are Aerosciences Evaluation and Test Capabilities; High-End Computing Capability; Rocket Propulsion Test Program; Space Environments Testing Management Office; and Aircraft Management Advisory Board.

Sustained Leadership Attention and Renewed Commitment Are Needed to Reinvigorate NASA's HEC

HEC permeates NASA activities by providing computing services to engineers, scientists, and researchers alike, however, we found NASA's HEC stewardship and competitiveness is waning. For example, the Agency's HEC organizational structure contributes to ineffective and undisciplined management, resource oversubscription negatively impacts mission work, and the cloud approach and cybersecurity practices lack a comprehensive strategy. In our view, without elevating HEC leadership to an executive level position, better engaging relevant stakeholders, and employing an integrated Agency-wide HEC strategy, NASA's science and research objectives will continue to be constrained, likely impacting future mission priorities and goals.

Organizational Structure Hinders NASA's HEC Efforts

Although NASA's mission is unique, the challenges the Agency faces in managing a decentralized HEC environment are not. We examined HEC organizational placement and executive leadership at two agencies—the National Oceanic and Atmospheric Administration (NOAA) and the Department of Energy (DOE)—that, like NASA, are geographically diverse, independently managed, and support numerous users.[29] Unlike at NASA, NOAA and DOE HEC capabilities are managed at an executive level, which provides the capability with greater organization visibility. At NASA, the Associate Administrator for SMD is responsible for directing and overseeing five broad scientific pursuits: Earth science, planetary science, biological and physical sciences, heliophysics, and astrophysics; HEC is buried in Earth science under the Earth Science Research Program despite being used across much of the Agency. Within SMD, a scientist wearing dual hats—Earth science programmatic compute activities and HEC capability manager—manages HEC at both the NAS facility and NCCS.

NASA. HEC is not managed as a program or centralized Agency strategic service, and organizational placement resides several layers deep within SMD—hindering the Agency's ability to efficiently allocate resources, adapt to varying mission needs, and quickly respond to technology advancements. Furthermore, while the Office of the Chief Information Officer (OCIO) has some oversight of HEC, similar to its involvement with large procurements and governance through control boards, it is not directly engaged in HEC activities or governance.

According to Agency policy, CPs provide NASA with the ability to manage components strategically and centrally in a way that balances the needs and demands of programs and projects and external

[29] Executive leadership guides decision-making by influencing activities such as fulfilling organizational goals, strategic planning, funding, and integrating new technologies into agency processes. For the purpose of this report, we further define executive leadership as leadership at NASA's Associate Administrator level.

partners.[30] Although NASA identified HEC as a CP, we found that key guiding documents and frameworks are absent, such as a management plan—an agreement between the sponsoring Mission Directorate and the CP manager detailing how the portfolio will be managed—and a commitment agreement designed to engage all relevant stakeholders and identify HEC as a strategic activity. While the sponsoring Mission Directorate and CP manager are responsible for periodically evaluating the efficiency, effectiveness, and performance of the portfolio and its capability components, during the course of our audit, HEC internal stakeholders expressed frustration at the lack of a management plan detailing how HEC resources are allocated. In short, without a management plan and commitment agreement, important aspects of the administration of HEC resources remain obscured to stakeholders.

Additionally, while NASA has engaged in HEC activities for more than 30 years, the Agency is not keeping up with today's rapid technological developments and specialized scientific and advanced research computing requirements due to funding and organizational constraints. NASA's current HEC infrastructure is built almost exclusively on cores of central processing units (CPUs), which lack the computational capabilities of the more modern graphics processing units (GPUs)—currently in use within the field worldwide.[31] The NAS facility, for instance, has more than 18,000 CPUs and only 48 GPUs, with an even larger disparity observed at the NCCS. HEC officials raised multiple concerns regarding this observation, stating that the inability to modernize NASA's systems can be attributed to various factors such as supply chain concerns, modern computing language (coding) requirements, and the scarcity of qualified personnel needed to implement the new technologies. Ultimately, this inability to modernize its current HEC infrastructure will directly impact the Agency's ability to meet its exploration, scientific, and research goals. One such example is the Agency's use and development of AI, which we examined in a May 2023 report.[32] HEC officials explained that, based on increasing stakeholder demand, they are "starting to tackle AI workloads that require new and different hardware." While HEC can support some small AI projects, the Agency's current HEC ecosystem cannot support projects that require a massive data stream.

The lack of organizational structure and strategic focus also affects the evolution of AI technologies—meaningfully impacting NASA's mission success. As a primary user of HEC capabilities, NASA's AI users create processes, refine data, and develop capabilities to help solve HEC limitations and improve data-driven enhancements. NASA's AI community is expanding its own organizational structure and implementing the use of "tiger teams" to help assist with the alignment of AI leadership across the Agency. AI officials discussed their intent for their tiger teams to engage with HEC leadership in efforts to advocate for areas that HEC could assist with AI advancement. However, in our judgement, such efforts are likely to prove less fruitful given the limitations of the Agency's HEC organizational structure and capability deficiencies.

NOAA and DOE. At both NOAA and DOE, HEC is managed at an executive level. At NOAA, HEC is the cornerstone of weather forecasting, climate and weather research, and understanding coastal issues. HEC investments are managed as an integrated enterprise by the NOAA Office of the Chief Information Officer's High-Performance Computing and Communications Program. Additionally, NOAA's High

---

30 NPR 8600.1.

31 The term "processor" is often used to refer to a CPU. A GPU is an electronic circuit that can perform mathematical calculations at high speed. A GPU's design allows it to perform the same operation on multiple data values in parallel, increasing its processing efficiency for compute-intensive tasks.

32 NASA OIG, *NASA's Management of Its Artificial Intelligence Capabilities* (IG-23-012, May 3, 2023).

Performance Computing Board provides strategic guidance and oversight for the execution of the enterprise and the allocation of HEC resources.

At DOE, the Office of Science—a program office—oversees world-class HEC to advance energy and scientific technology. The Office of Science is led by a presidentially nominated, Senate-confirmed director and two senior-career federal deputy directors. DOE's Frontier supercomputer and its scientists are leading the world in science and technology innovation.[33] The Frontier supercomputer, which can model and simulate complex problems associated with the nation's energy sector, helps maintain U.S. leadership in AI and HEC systems.

HEC Oversubscription Impacts Missions

NASA's HEC is currently oversubscribed and overburdened by demand and competing priorities. HEC officials use an Agency website for capacity planning and management—balancing user requirements against resource availability. Quarterly updates to the 5-year Capacity Plan—in conjunction with a Data Management Plan—provide a catalyst encouraging scientists to "think up front," identify their needs, and understand HEC resource limitations. Generally, HEC capacity allocations are determined and prioritized at the Mission Directorate level. Each organization prioritizes requests against criteria such as mission priority, anticipated science workload, available capacity, and utilization time. Importantly, operational mission work is given HEC resource priority over research and development activities.

According to HEC officials, resources are currently operating at capacity and are oversubscribed—meaning more time is requested by missions than capacity can provide—with demand sometimes exceeding supply by as much as threefold. To make the most of the schedule ebb and flow, resources are typically allocated at 140 percent of HEC resource capacity to eliminate potential idle time. While this oversubscription permits smaller jobs—which typically take less time—to take advantage of availability gaps, larger, more intensive compute jobs are often delayed. HEC officials explained they are concerned that vital resources needed for mission innovation are strained, receiving three times more requests than capacity allows. Given that wait times vary in length, depending on resource availability, trickle-down schedule delays can cause impacts to a launch or critical manufacturing. For example, the Mars Ascent Vehicle, an element in the Mars Sample Return Program, has experienced schedule delays waiting on NAS facility resources. Specifically, a 3-week delay was encountered for critical Computational Fluid Dynamics (CFD) simulations needed for key milestone reviews.[34]

Similarly, Aeronautics Research Mission Directorate (ARMD) stakeholders identified NAS resource oversubscription as a formal risk due to allocated HEC compute resources falling short of the organization's needs. According to the stakeholders, ARMD programs are unable to meet their performance objectives due, in part, to capacity limitations driven by NAS demand exceeding availability. For example, the oversubscription risk and shortage of HEC resources affected NASA's

[33] Frontier, the most powerful supercomputer in the world, is located at the DOE's Oak Ridge National Laboratory in Tennessee.

[34] CFD is a powerful tool that creates a digital simulation visualizing the flow of fluids and the way they are affected by objects. It shows how temperature, pressure, and velocity are going to behave in a design.

Sustainable Flight Demonstrator—a partnership between NASA and The Boeing Company—causing ARMD to consider purchasing their own HEC resources at a projected cost of $3.1 million.[35]

During discussions with about a dozen HEC stakeholders, we identified additional examples of oversubscription and competing priorities. Officials within the Office of the Chief Engineer explained that job delays in resource availability and the subsequent impacts across the Agency have not been studied. Accordingly, NASA officials do not know whether paying more to run a job sooner, on an HEC commercial cloud resource for example, outweighs the cost of idle scientists and mission impacts while waiting for on-premises NASA HEC system availability.

Table 1 below depicts some HEC issues impacting NASA missions.

Table 1: Examples of HEC Issues Impacting Missions

Category	Issue and Impact
Hardware/System	Bad nodes (slow, crash, or reboot) cause jobs to fail; users have to work with HEC support to take those nodes offline.[36]
	Network issues degrade inter-node communication and cause jobs to fail.
	File system stability causes very slow response times (several minutes); file systems won't mount on compute nodes causing job to fail.
Resource Availability	High system usage has led to long queue times (several days to weeks for typical size jobs of ~2000 cores).
	Forced to get a reservation to meet project schedules after getting little throughput in the normal queues; impacted systems: Commercial Crew Program, Human Landing System, Space Launch System Exploration Upper Stage, and Mars Ascent Vehicle.
Large Resource Requests	Several high node count/long duration reservations have dominated system resources and impacted Exploration Systems Development Mission Directorate and Space Operations Mission Directorate.
	Directorates make allocation requests to HEC management; stakeholders are not engaged to justify resource allocation and prioritize mission share; it is not clear how the determination of mission priority is made.
Program Schedule	Human Landing System has experienced multiple delays in analyzing critical components and CFD simulations due to lack of available resources.

Source: NASA.

To satisfy unmet HEC needs, individual NASA organizations have circumvented the queue by purchasing additional HEC assets directly from vendors. While essential to overcoming the technological challenges inherent to many NASA projects, this scenario also leads to using alternative means to "get the job done" outside of the current HEC structure and capability. For example, the Space Launch System (SLS) Program purchases and operates their own HEC clusters due to lack of availability and reliability of NAS

[35] The purpose of the Sustainable Flight Demonstrator project is to engage with industry, academia, and other government organizations to identify, select, and mature key airframe technologies, such as new wing designs, that have a high probability of transition to the next generation single-aisle seat class airliner.

[36] A node is any computer or other device connected to a network that sends, receives, or redistributes data. For example, computers, file servers, network-connected printers, and routers are all nodes.

systems.[37] SLS officials explained that they invest about $250,000 annually to purchase and locally manage HEC clusters to ensure availability—necessitated, in part, to combat issues such as a 2-year delay in NAS resource availability required for aerosciences database analysis. Simply put, due to HEC resource oversubscription, the SLS Program does not reliably trust they can run jobs in time to meet their schedule needs.

On much smaller scales, HEC clusters at Johnson Space Center and Marshall Space Flight Center operate for the sole purposes of supporting site-specific engineering projects within their purview, including those for the Artemis campaign such as Orion (crew capsule), Gateway (orbiting lunar outpost), and SLS. Center officials explained that individual HEC assets are acquired for numerous reasons, such as the lack of access to existing HEC assets, lack of allocation availability, or lack of required software/technology for a specific project or function.

Frequently, these assets are unidentified and not managed as part of existing enterprise-wide HEC capabilities. Multiple NASA officials across organizational boundaries said that if an HEC asset was purchased directly by a mission or project, the OCIO does not know or manage the asset, and that project teams are doing what they want—essentially operating undisciplined, in a "free-for-all" HEC environment. We found that multiple independent HEC assets exist outside the HEC capability portfolio and are located at various NASA Centers. For example, Langley Research Center's (Langley) Computational Research Facility pays for and maintains K-Cluster as a mid-range HEC designed to be a proxy for the Pleiades asset at Ames.[38] As a local on-demand resource, K-Cluster offers Langley developers a way to supplement unavailable NAS resources. According to NASA officials, Center-owned HEC assets are a "third rail"—meaning a controversial and unresolved issue. Unanswered organizational questions persist, such as should Centers have their own HEC resources; does it make sense to incorporate and manage those clusters at the Agency level; and should Center assets be a shared or exclusive resource.

In fact, Agency officials told us that, except for Goddard and Stennis Space Center, there are independent HEC assets installed at almost every NASA Center. In the case of Center-managed HEC, the purchasing costs of the HEC assets are frequently paid for by program-specific funds while power consumption for these Center-specific assets is funded by the Centers. Moreover, the complex cooling mechanisms required to maintain proper HEC operating temperatures adds not only real estate required to house the HEC systems in the local data centers, but also additional water and power infrastructure and consumption. The cost of maintaining Center HEC assets represents a hidden cost that is not accounted for in Agency HEC budgets or reporting. As a result, competing priorities and an unpredictable HEC oversubscription strategy have created schedule delays, encouraged the development of shadow assets—information technology on the Agency's network that the Chief Information Officer did not purchase or authorize for use. Figure 5 displays the HEC infrastructure, including cooling pipes, at Langley.

---

[37] The SLS is a two-stage, heavy-lift rocket that launches the Orion Multi-Purpose Crew Vehicle into space.

[38] A mid-range supercomputer has, at most, one order of magnitude lower performance level than a high-end supercomputer.

Figure 5: Langley's High-End Compute Infrastructure (as of August 2023)

Source: OIG HEC photographs taken at Langley on August 24, 2023.

NASA's HEC Lacks a Comprehensive Commercial Cloud Strategy

NASA's HEC and, by extension, the Agency's HEC user base, does not have an overarching strategy on commercial cloud versus on-premises use. SMD, as the manager and largest user of NASA's on-premises HEC facilities, lacks a single, cohesive cloud approach needed to match evolving scientific and technological mission goals. Although NASA acquired commercial HEC cloud computing beginning in 2019 to complement and expand the Agency's current HEC on-premises capabilities, we found that cloud adoption among HEC users has been lackluster because of misconceptions among the various stakeholders about cloud computing related to cost, scheduling, and capacity. For instance, stakeholders explained that although they were aware of NASA's HEC cloud computing capabilities, they were hesitant to use the cloud due to unknown scheduling practices or assumed greater costs. However, OCIO officials stated that cloud pricing has improved dramatically from what was determined in previous cost studies and that a mechanism for capacity planning and scheduling was already in use by the HEC community. HEC and OCIO officials told us they were hopeful that with a more a strategic and integrated approach, cloud misconceptions can be resolved. Given that SMD currently operates five cloud computing environments with redundant capabilities, a centrally managed HEC cloud environment would benefit NASA's scientific goals by enabling a more optimized approach to managing HEC resources.

To address HEC needs, SMD officials explained they are focusing on combining the commercial cloud and on-premises resources to maximize the benefits of both—the former is more expensive and unrestricted, the latter is more affordable and constrained. Depending on user requirements, resource availability, and cost, the overall goal is to "marry" the best architecture required to support HEC applications. For example, a small weather simulation forecast can be quickly operational in the cloud, whereas it could take days to run a similar model with NASA HEC on-premises resources due to time and availability constraints.

Currently, NASA's ability to process large amounts of data from projects, such as heliophysics, is limited due to HEC resource capacity constraints.[39] According to SMD officials, resource challenges need to be met not only with investments in hardware, but also expertise in software and code modernization. Because about half of HEC resources are used for modeling, SMD management initiated a study to further evaluate augmenting HEC resources with cloud computing.[40] Chartered in 2022, the Data and Computing Architecture Study addressed the question of whether a coordinated cloud and on-premises computing infrastructure can meet SMD's data and computing needs, enable efficiencies, and support SMD's transition to open-source science. The study results were released in November 2023. The SMD Management Advisory Council agreed with the study recommendations to improve coordination and alignment of HEC, but a timeline for implementation has not yet been established.

Both NOAA and DOE use various commercial cloud service providers, such as AWS and Google Cloud, to supplement their on-premises HEC resources as a way to meet specific resource-intensive peak computing demands. Commercial cloud environments can provide additional computational resources to meet dynamic HEC user demands, including immediate access to new and novel technologies such as new chip architectures, graphics processing units, and other parallel file systems.[41] However, NASA's lack of charging policy contributes to cloud misconceptions and uncertainty about when it is appropriate—in terms of cost, time, and schedule—to use HEC cloud versus on-premises resources for the task at hand.

Diluted Cybersecurity Practices Expose HEC Risks

Cybersecurity controls in any HEC environment are atypical of those applied in a standard IT enterprise environment and are difficult to implement. NASA further compounds these issues through a myriad of shortcomings stemming from the decentralized deployment and management of its HEC assets. The difficulty of balancing performance demands for HEC projects against cybersecurity controls and applications is not unique to NASA. Compared to a typical IT environment, HEC systems are by definition designed and optimized to run at a high-performance level and operate exotic software to solve complex scientific problems by processing large amounts of data and performing calculations at high speeds. For example, HEC systems run complex AI software to model climate research, interpret geospatial data from satellite images, and use CFD to simulate environments impacting space vehicles.[42] Given that peak performance is the metric most HEC users are concerned with, OCIO officials explained that security controls are historically treated as an afterthought and implemented such that they do not interfere with usability or performance.

---

[39] Heliophysics studies the nature and dynamic interactions of the Sun, the heliosphere, the plasma environments of the planets, and interstellar space.

[40] Modeling allows high-fidelity simulations of systems in environments that are difficult or impossible to create on Earth, allows removal of humans from experiments in dangerous situations, and provides visualizations of datasets that are extremely large and complicated.

[41] A chip architecture is the overall design of a microchip encompassing both the hardware and software components—defining how they work together. A parallel file system is a type of storage system designed to store data across multiple networked servers and to facilitate high-performance access through simultaneous, coordinated input/output operations between clients and storage nodes.

[42] Geospatial data directly or indirectly references a specific geographical area or location related to the Earth.

Most federal agencies and enterprises standardize their cybersecurity practices by implementing the controls published by NIST.[43] While these controls can be generalized to a degree, they are descriptive in their design and purpose—asking the system owner to implement the most stringent measures available such as strict access controls, user monitoring, multi-factor authentication, and constant or frequent asset scanning. For complex reasons, overly stringent controls do not work effectively in an HEC environment. Specifically, monitoring software can slow processing speed, and scanning tools can negatively impact the delicate system environment utilized by HEC servers and nodes. However, instead of seeking to find a balance for compensatory controls that could meet some, if not most, cybersecurity requirements, the mission-oriented HEC user base's approach typically involves repeated waiver requests—known as risk-based decisions—to bypass the intended control, or built-in generic controls. Consequently, neither overly stringent security measures, nor waivers, offer practical solutions to protect the Agency's HEC environments from cyber threats.

During our review, HEC officials explained that HEC system administrators conduct manual audit log analysis on a sporadic basis—mainly to identify performance bottlenecks instead of detecting potential cybersecurity threats or vulnerabilities. We found this practice counterintuitive from a cybersecurity perspective for multiple reasons. For example, best practices, such as separation of duties, require independent analysis to prevent log manipulation by a potential insider threat. Also, manual log reviews are impractical, slow, and prone to human error such as not identifying anomalous patters or indicators of questionable or malicious activity.

HEC officials said that an automated monitoring solution, such as Splunk, is under consideration to convert data into operational intelligence using reports, charts, and alerts.[44] We found that NAS facility personnel are in the process of procuring a standalone instance of Splunk at an annual cost of $34,000. While an automated monitoring and analysis solution is ideal, we question the need to purchase a separate Splunk license when the OCIO holds an enterprise-wide license. NASA OIG previously reviewed the Agency's software asset management practices and found that duplicative spending and mismanagement of licenses is a problem throughout NASA.[45] In our view, this is a missed opportunity to consolidate analysis between the OCIO and Mission Directorates for improved cyber-situational awareness and software license cost avoidance.

NASA further exacerbates these cybersecurity issues as the result of poor management practices and limited oversight. In addition to the propensity for projects to build their own HEC clusters as previously discussed, we observed OCIO-mandated cybersecurity controls being ignored or bypassed by the Mission Directorates because they perceive the controls as overly stringent. Since these projects did not engage with OCIO in meaningful discussions about how best to implement various, essential tools—such as BigFix—needed for IT asset discovery and inventory are absent, resulting in the inability to identify and/or classify HEC systems.[46] Additionally, the high level of use of NASA HEC assets by external or

[43] NIST Special Publication 800-53, Revision 5, *Security and Privacy Controls for Information Systems and Organizations* (September 2020, includes updates as of December 10, 2020).

[44] Splunk is a software platform to search, analyze, and visualize machine-generated data gathered from the websites, applications, sensors, and devices that make up the IT infrastructure.

[45] NASA OIG, *NASA's Software Asset Management* (IG-23-008, January 12, 2023).

[46] BigFix provides computer management services, including asset inventory/discovery, security vulnerability detection and remediation, security policy enforcement, software distribution, IT compliance reporting, patch management, and software license management. This is NASA's primary system management tool.

foreign users further compounds cyber risk. Collectively, these issues raise serious questions about the overall security posture of NASA's HEC environment.

As previously discussed, while NASA's main HEC capabilities are found within the NAS and NCCS computing facilities, several other organizations operate individual HEC clusters for their own specific purposes. These individual HEC assets are owned and managed independent of NASA's primary systems, and each owner employs their own methodologies for system management, including the common tenets of cybersecurity such as monitoring and vulnerability mitigation/patching. These individual HEC clusters are not always placed on NASA's or public-facing networks, which limits the effectiveness of BigFix to discover, inventory, and monitor these assets for any known vulnerabilities that may need to be patched or mitigated. Also, the owners do not always annotate that these systems are high-performance or HEC in nature within the Agency's current asset system of record, known as Risk Information Security Compliance System (RISCS), and there is no automatic or unique classification mechanism within the system to do so.[47] This further complicates the Agency's ability to inventory and monitor its systems as there is no way to manually search RISCS for systems that are undetectable by BigFix. Similarly, these issues are not limited to the outlying, independent, Center HEC clusters. We found that NASA's HEC infrastructure likely poses additional risks due to insufficient asset monitoring and poor inventory, regardless of location or owner.

More concerning, we identified extensive use of NASA's HEC assets by external and foreign national parties. As of May 2023, a total of more than 700 foreign national users (including legal permanent residents living in the United States on a green card) were reported as currently having access to the NAS facility or NCCS assets for their respective projects and studies, which is likely only a snapshot of the overall picture. We found that individual Center HEC asset users are not steadily tracked, and interconnection security agreement evaluations are not conducted to verify rights and accesses granted to international partners.[48] Furthermore, multiple security personnel associated with the HEC capability discussed the lack of user activity monitoring and their lack of involvement in the approval or review process for external parties to gain access to HEC systems or NASA datasets. As a result, limited cybersecurity practices expose the Agency's HEC resources at a higher risk than necessary. Without a refocused effort to implement better cybersecurity safeguards, NASA's HEC resources will continue to be high-value targets to adversaries—potentially posing risks to not just a singular asset, but to NASA's network and IT inventory as a whole.

---

[47] RISCS is a data repository that contains an inventory of the Agency's hardware and software, including system security and contingency plans for each information system.

[48] Interconnection security agreements regulate security-relevant aspects of an intended connection between an agency and an external system.

CONCLUSION

Preparing for future space exploration is not just about developing large rockets, exploring Mars, and making scientific discoveries. It is also about novel concepts, innovation, and trailblazing—using HEC to solve the most pressing problems with data-driven solutions. Since NASA's inception in 1958 to present day, the Agency's history is written with each unique scientific and technological achievement. NASA has landed people on the Moon, visited every planet in the solar system, touched the Sun, and solved some of the mysteries of our home planet. Although NASA's HEC ecosystem is vital to sustaining mission priorities and advancing science, we found that executive leadership and stakeholder engagement is absent, lacking a cohesive approach to managing HEC across organizational boundaries. NASA's HEC is currently oversubscribed and overburdened by demand and competing priorities, which impact the wait time for missions to use the assets. The Agency also lacks a strategy on commercial cloud use, contributing to stakeholder misconceptions about cost, scheduling, and capacity. Finally, due to the unique nature of the HEC environment, security controls are often bypassed or not implemented, increasing the risk of cyberattacks. Without an integrated HEC strategy, a more focused management approach, and advocacy, the Agency's trailblazing science and technology research is severely limited while cyber risk is elevated.

Recommendations, Management's Response, and Our Evaluation

To establish executive leadership and to strategically position NASA's HEC capability to better meet the needs of specialized scientific and advanced research computing requirements, we recommended NASA's Associate Administrator:

1. Appoint executive leadership to determine appropriate definition/scope, ownership, organizational placement, and structure (e.g., portfolio, program, enterprise service) of HEC within NASA.

In addition, we recommended that the NASA Associate Administrator establish a tiger team to collaborate and strategize on wide-ranging HEC issues to:

2. Develop enterprise-wide HEC stakeholder requirements to validate commitment agreements as required in NPR 8600.1.
3. Identify technology gaps, such as GPU transition and code modernization, essential for meeting current and future needs and strategic technological and scientific requirements.
4. Develop a strategy to improve HEC asset allocations and prioritization for usage, including the appropriate use of on-premises versus cloud resources.
5. Evaluate cyber risks associated with HEC assets to determine oversight and monitoring requirements, establish risk appetite, and address control deficiencies. Consider using NASA's Splunk enterprise platform as a shared resource.
6. Implement an HEC classification/category designation within RISCS for identifying HEC assets.
7. Develop an inventory of enterprise-wide HEC assets and formalize procedures for hardware and software life-cycle management.
8. Document data risk impact levels, classification, and export control categorization for all HEC jobs.
9. Identify and mitigate gaps in the foreign national accreditation access process.

We provided a draft of this report to NASA management who concurred with Recommendation 1 and partially concurred with Recommendations 2 through 9. NASA described planned actions to address Recommendation 1 and stated that a tiger team will be established to collaborate and strategize on HEC issues; subsequently, the tiger team will determine the implementation of Recommendations 2 through 9. While we consider management's comments responsive, OIG requests regular updates to understand the planned actions and timeline for implementation and to monitor progress towards implementation of Recommendations 2 through 9. The recommendations are resolved and will be closed upon completion and verification of the proposed corrective actions.

Management's comments are reproduced in Appendix E. Technical comments provided by management and revisions to address them have been incorporated as appropriate.

Major contributors to this report include Tekla Colón, Mission Support Audits Director; Scott Riggenbach, Assistant Director; Joseph Cook; Linda Hargrove; Christopher Reeves; Vincent Whitfield, and Jaidan Williams. Courtney Daniels provided editorial support.

If you have questions about this report or wish to comment on the quality or usefulness of this report, contact Laurence Hawkins, Audit Operations and Quality Assurance Director, at (202) 358-1543 or laurence.b.hawkins@nasa.gov.

Digitally signed by George Scott
Date: 2024.03.13 12:32:38 -04'00'

George A. Scott

Acting Inspector General

Appendix A: Scope and Methodology

We performed this audit from March 2023 through February 2024 in accordance with generally accepted government auditing standards. Those standards require that we plan and perform the audit to obtain sufficient, appropriate evidence to provide a reasonable basis for our findings and conclusions based on our audit objectives. We believe that the evidence obtained provides a reasonable basis for our findings and conclusions based on our audit objectives.

Our audit scope encompassed NASA's HEC capabilities. Specifically, we focused on relevant processes and controls related to policy framework, capacity planning, stakeholder engagement, and cybersecurity. During our audit we identified and reviewed applicable policy framework and criteria; interviewed responsible officials, key personnel, and stakeholders; evaluated the adequacy of the Agency's capacity planning tools and success metrics; and reviewed the capabilities and techniques employed to identify and mitigate HEC cybersecurity risks.

Methodology

We divided the audit into four subject areas: (1) policy framework, (2) capacity planning, (3) stakeholder engagement, and (4) cybersecurity. Key work completed for each subject area is summarized below.

To gain a holistic view of NASA's HEC capabilities, we focused on relevant processes and cybersecurity controls related to the HEC ecosystem. For each subject area, we collected and reviewed numerous federal and Agency policies, regulations, guidance, and industry best practices. We researched the *TOP500* rankings of HEC. We interviewed responsible NASA officials and stakeholders from the OCIO, Cybersecurity and Privacy Division, Kennedy Space Center's Chief Technology Officer, Mission Directorates, and officials responsible for overseeing HEC activities. We met with officials from Ames, Goddard, Johnson Space Center, Langley, and the Jet Propulsion Laboratory regarding the facilities unique to their Centers. We met with the Networking and Information Technology Research and Development Program's NASA representative for an overview of the Agency's participation in the program. We benchmarked best practices with DOE's Oak Ridge National Laboratory and NOAA. We participated in NIST's high-performance computing security working group biweekly meetings to increase our comprehension of the HEC technical architecture and functional cybersecurity overlays. Additionally, we conducted site visits at the Ames, Goddard, and Langley HEC facilities. Collectively, this informed our understanding and helped us assess the overall management of NASA's HEC capabilities.

Assessment of Data Reliability

We used limited computer-processed data extracted from NASA's IT systems during the course of this audit. Although we did not independently verify the reliability for all of the information provided, we compared it with other available supporting documents to determine data consistency and reasonableness. From these efforts, we believe the information we obtained is sufficiently reliable for this report.

Review of Internal Controls

We assessed internal controls and compliance with laws and regulations to determine NASA's overall HEC program management and cybersecurity preparedness. Control weaknesses and misaligned management practices are identified and discussed in this report. Our recommendations, if implemented, will improve NASA's overall HEC program and those identified weaknesses.

Prior Coverage

During the last 5 years, the NASA Office of Inspector General (OIG) has not issued any reports of relevance to the subject of this audit. Additionally, GAO has issued an ancillary report of interest to this topic. Unrestricted reports can be accessed at https://oig.nasa.gov/ and https://www.gao.gov/ respectively.

Government Accountability Office

High-Performance Computing: Advances Made Towards Implementing the National Strategy, But Better Reporting and A More Detailed Plan Are Needed (GAO 21-104500, September 2021).

APPENDIX B: ANATOMY OF AN HEC

A high-end computer is a complicated machine. When scientists want to study something that is impossible to explore in a lab—like an exploding star or a fast-forming hurricane—the computer has to be super. As shown in Figure 6 below, various components enable HEC assets to produce massive amounts of computational power.

Figure 6: Anatomy of a High-End Computer

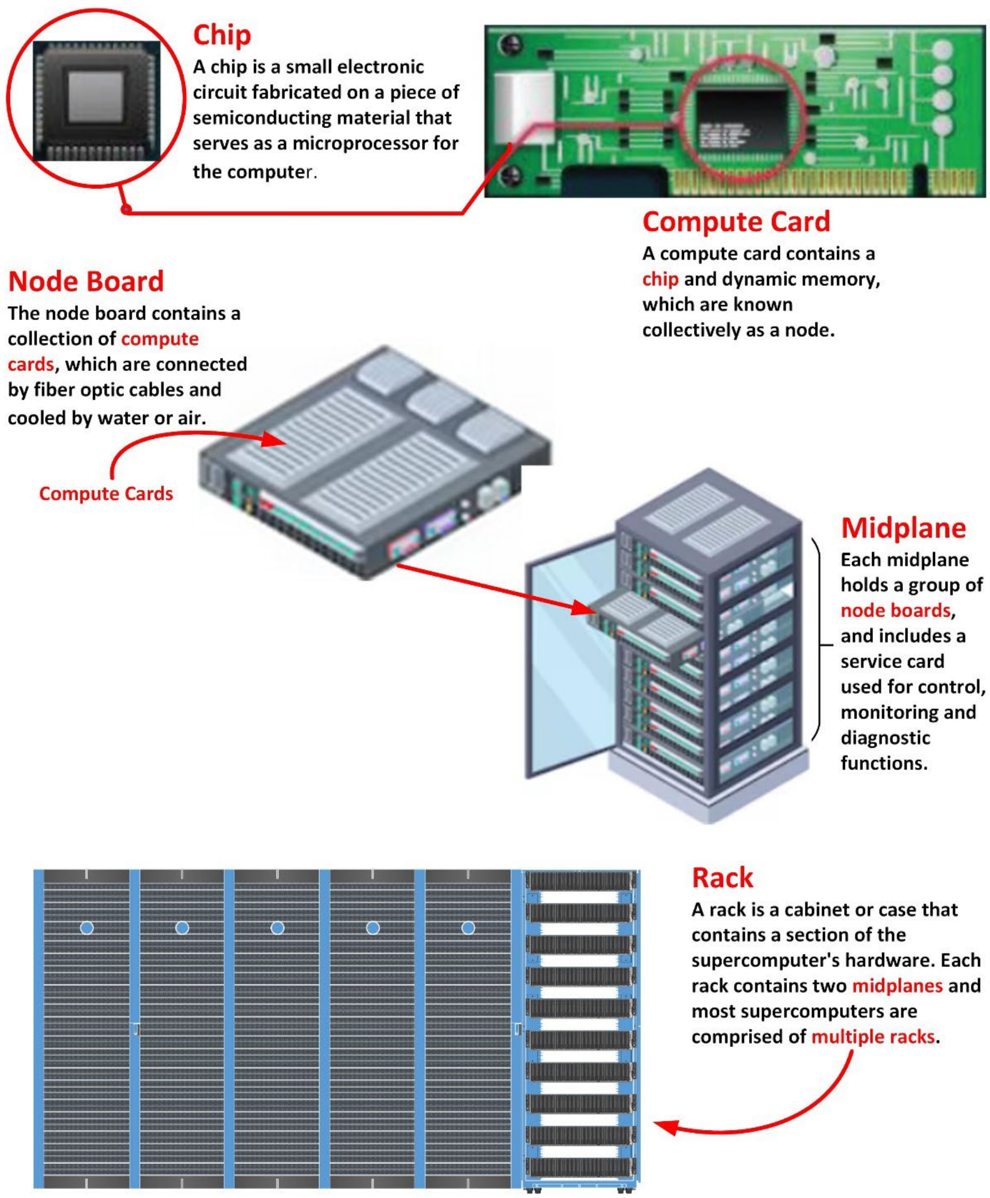

Source: OIG representation of DOE information.

APPENDIX C: EXAMPLES OF NASA'S HEC IN ACTION

Below is a representation of how HEC supports various NASA Mission Directorates in their work.

Science Mission Directorate (SMD)

Illustration of the Planetary System K2-138

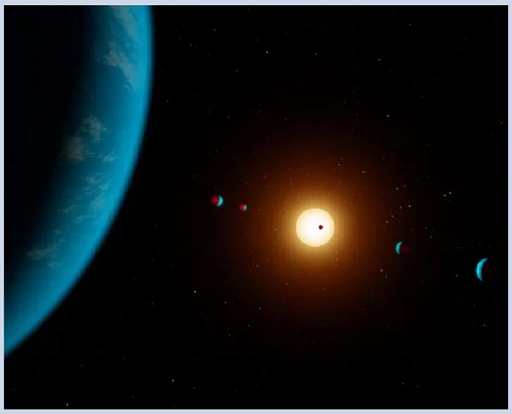

K2-138 was discovered in 2017 using data from NASA's Kepler space telescope.

Source: NASA

SMD leveraged NASA's Pleiades supercomputer and ExoMiner AI network to analyze data provided by the Kepler space telescope in search of new planets. These combined technologies were able to comb through huge amounts of data to distinguish between real imagery and false positives—leading to the discovery of 301 exoplanets seemingly all at once.

Aeronautics Research Mission Directorate (ARMD)

Using NASA's supercomputing facilities, ARMD is simulating the aerodynamic performance of several promising air taxi vehicle configurations that will someday carry passengers and cargo in urban and suburban areas. The highly complex simulations will be used to help design and develop these future air taxis—also called Advanced Air Mobility vehicles—that will be safe, quiet, and efficient.

Air Flow Vizualization for an Advanced Air Mobility Vehicle

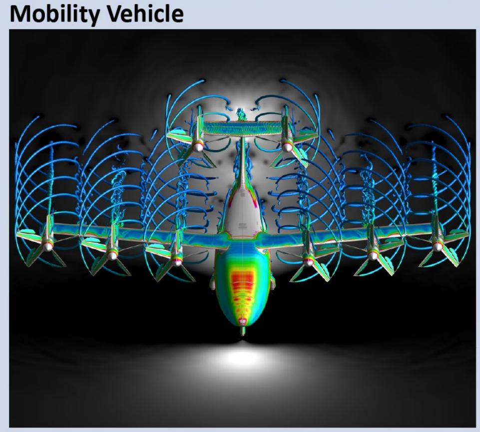

Source: NASA

Space Launch System on the Launchpad

Source: NASA

Exploration Systems Development Mission Directorate (ESDMD)

NASA's ESDMD has the responsibility of developing the SLS, which will send astronauts to the Moon in the next few years. They leverage NASA's HEC systems to conduct simulations generating terabytes of data that will be crucial for the final development and construction of the SLS to ensure crew safety and mission success.

Space Operations Mission Directorate

Working in tandem with the ESDMD, the Space Operations Mission Directorate is utilizing NASA's HEC systems to analyze wide ranges of data necessary to modernize Kennedy's Launch Pad 39B, bringing it up to date from the space shuttle era. Kennedy's launch pads will be used in NASA's future crewed Artemis missions to the Moon and Mars.

Simulation of SLS Launch Ignition

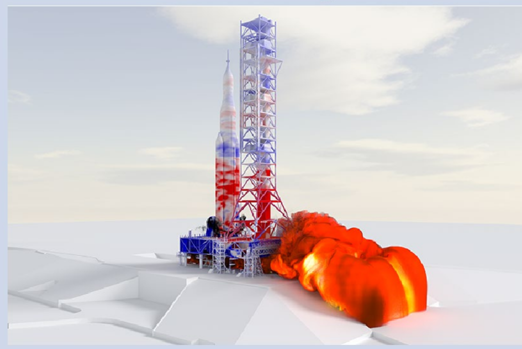

Simulation of a launch at Kennedy's Launch Pad 39B

Source: NASA

Entry Systems Modeling

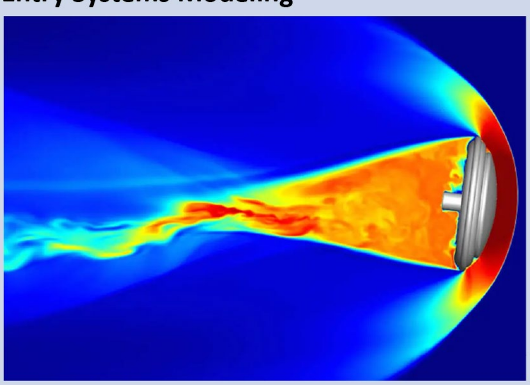

Modeling used to design heatshields for spacecraft

Source: NASA

Space Technology Mission Directorate

The Space Technology Mission Directorate uses HEC resources to simulate a wide variety of conditions related to the development of space technologies—such as those used for guidance, navigation, and control of spacecraft. The Entry Systems Modeling project is one such simulation, using supercomputers to predict spacecraft performance during the extreme conditions of atmospheric entry.

APPENDIX D: WORLD TOP500 HEC RANKINGS

Begun in 1993 and updated every six months, the *TOP500* provides a ranked list (using the LINPACK benchmark) of HEC systems throughout the world.[49] Statistics on HEC are of major interest to manufacturers, users, and potential users. Such statistics facilitate the establishment of collaborations, the exchange of data and software, and a better understanding of high-performance computing.

As of November 2023, the United States and China make up 53 percent of the machines on the *TOP500* supercomputer list—with a total of 161 and 104 HEC systems respectively. The world's top 10 HEC systems, as of November 2023, are listed below in Table 2. According to the November 2023 *TOP500* supercomputer world rankings, NASA HEC systems at Ames, Aiken, Pleiades, and Electra, are ranked numbers 85, 132, and 143, respectively. Ames' Endeavour and Goddard's Discover are not ranked in the *TOP500.*

Table 2: World Top HEC Rankings (as of November 2023)

Rank	Name	Owner	Location
1	Frontier 	Department of Energy	Oak Ridge National Laboratory, Oak Ridge, Tennessee, USA
2	Aurora 	Department of Energy	Argonne National Laboratory, Argonne, Illinois, USA
3	Eagle 	Microsoft Azure	Redmond, Washington, USA
4	Fugaku 	RIKEN Center for Computational Science	Kobe, Japan
5	LUMI 	EuroHPC	Kajaani, Finland
6	Leonardo 	EuroHPC	Cineca, Italy

[49] The LINPACK benchmark is a measure of a computer's floating-point rate of execution. It is determined by running a computer program that solves a dense system of linear equations.

7	Summit 	Department of Energy	Oak Ridge National Laboratory, Oak Ridge, Tennessee, USA
8	MareNostrun 5 ACC 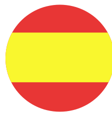	EuroHPC	Barcelona, Spain
9	EOS NVIDIA DGX Super POD 	NVIDIA Corporation	Santa Clara, California, USA
10	Sierra 	Department of Energy	Livermore, California, USA

Source: OIG representation of *TOP500* rankings.

Appendix E: Management's Comments

National Aeronautics and Space Administration

Office of the Administrator
Mary W. Jackson NASA Headquarters
Washington, DC 20546-0001

March 8, 2024

TO: Assistant Inspector General for Audits

FROM: Associate Administrator

SUBJECT: Agency Response to OIG Draft Report, "Audit of NASA's High-End Computing Program" (A-23-05-00-MSD)

The National Aeronautics and Space Administration (NASA) appreciates the opportunity to review and comment on the Office of Inspector General (OIG) draft report entitled, "Audit of NASA's High-End Computing Program" (A-23-05-00-MSD), dated February 8, 2024.

In this draft report, the OIG found that executive leadership and stakeholder engagement is absent, lacking a cohesive approach to managing high-end computing (HEC) across organizational boundaries. The OIG opined that without an integrated HEC strategy, a more focused management approach, and advocacy, the Agency's trailblazing science and technology research is severely limited. The OIG makes nine recommendations addressed to NASA's Associate Administrator.

The OIG recommends the following to the NASA Associate Administrator to establish executive leadership and strategically position NASA's HEC to meet the Agency's specialized needs:

Recommendation 1: Appoint executive leadership to determine appropriate definition/scope, ownership, organizational placement, and structure (e.g., portfolio, program, enterprise service) of HEC within NASA.

Management's Response: NASA concurs. The NASA Associate Administrator will convene a working group to perform a strategic assessment of the HEC Capability Portfolio (CP) to determine the appropriate definition/scope, ownership, organizational placement, and structure within NASA. The working group will be led by the Science Mission Directorate as the organization where the HEC CP is currently managed. The composition of the working group will include officials drawn from such places as Chief Program Management Officer, Aeronautics Research Mission Directorate, Office of the Chief Information Officer, Office of Strategic Infrastructure, Ames Research Center, Goddard Space Flight Center, and Langley Research Center. The working group will provide findings to the NASA Associate Administrator.

Estimated Completion Date: March 29, 2024.

In addition, the OIG recommends "…that the NASA Associate Administrator establish a tiger team to collaborate and strategize on wide-ranging HEC issues..."

NASA concurs with the spirit and intent of the OIG recommendation and will establish a tiger team to collaborate and strategize on the issues identified by the OIG in recommendations 2 through 9 listed below.

Recommendation 2: Develop enterprise-wide HEC stakeholder requirements to validate commitment agreements as required in NASA Procedural Requirements 8600.1, NASA Capability Portfolio Management Requirements.

Recommendation 3: Identify technology gaps, such as graphics processing unit transition and code modernization, essential for meeting current and future needs and strategic technological and scientific requirements.

Recommendation 4: Develop a strategy to improve HEC asset allocations and prioritization for usage, including the appropriate use of on-premises versus cloud resources.

Recommendation 5: Evaluate cyber risks associated with HEC assets to determine oversight and monitoring requirements, establish risk appetite, and address control deficiencies. Consider using NASA's Splunk enterprise platform as a shared resource.

Recommendation 6: Implement an HEC classification/category designation within Risk Information Security Compliance System for identifying HEC assets.

Recommendation 7: Develop an inventory of enterprise-wide HEC assets and formalize procedures for hardware and software life-cycle management.

Recommendation 8: Document data risk impact levels, classification, and export control categorization for all HEC jobs.

Recommendation 9: Identify and mitigate gaps in the foreign national accreditation access process.

Management's Response to Recommendations 2-9: NASA partially concurs with each recommendation. NASA concurs with the spirit and intent of the OIG recommendations and will establish a tiger team noting however, that the implementation of each recommendation should be determined by the tiger team.

Estimated Completion Date: The tiger team will establish estimated completion dates for each recommendation numbered 2 through 9 as soon as practical and will communicate those dates to the OIG by December 31, 2024.

We have reviewed the draft report for information that should not be publicly released. As a result of this review, we have not identified any information that should not be publicly released.

Once again, thank you for the opportunity to review and comment on the subject draft report. If you have any questions or require additional information regarding this response, please contact Kevin Gilligan at (202) 358-4544.

James M. Free

Appendix F: Report Distribution

National Aeronautics and Space Administration

Administrator
Deputy Administrator
Associate Administrator
Chief of Staff
Associate Administrator for Science Mission Directorate
Chief Information Officer
Chief Engineer

Non-NASA Organizations and Individuals

Office of Management and Budget
    Deputy Associate Director, Climate, Energy, Environment and Science Division

Government Accountability Office
    Director, Contracting and National Security Acquisitions

Office of Science and Technology Policy
    Deputy Director, Science and Society

Congressional Committees and Subcommittees, Chairman and Ranking Member

Senate Committee on Appropriations
    Subcommittee on Commerce, Justice, Science, and Related Agencies

Senate Committee on Commerce, Science, and Transportation
    Subcommittee on Space and Science

Senate Committee on Homeland Security and Governmental Affairs

House Committee on Appropriations
    Subcommittee on Commerce, Justice, Science, and Related Agencies

House Committee on Oversight and Accountability
    Subcommittee on Government Operations and the Federal Workforce

House Committee on Science, Space, and Technology
    Subcommittee on Investigations and Oversight
    Subcommittee on Space and Aeronautics

(Assignment No. A-23-05-00-MSD)