Special Report

Measuring and Assessing the Value and Scientific ROI of NERSC's Leadership Computing in Advancing Science

Mark Nossokoff, Jaclyn Ludema, and Earl Joseph
October 2024

HYPERION RESEARCH OPINION

The value and impact of the discoveries uncovered as a result of research done utilizing NERSC resources cannot be overstated. For 50 years, researchers and scientists from academic institutions and research organizations across the country have relied on NERSC to attain goals requiring ever-increasing computational capabilities to conduct their research. One measure of the extent of their accomplishments is the seven Nobel prizes awarded for achievements enabled by NERSC resources.

While Nobel prizes are one visible measure of success from scientific discovery, other measures are not so visible or tangible. Oftentimes it can take decades before the results of fundamental research being conducted come to fruition in the form of commercial or societal impact, if at all. Indeed, while science for science's sake is critically important for moving forward to address society's grand challenges, the tangible value of scientific research can be incredibly difficult to identify.

In support of its 50th anniversary celebration, NERSC commissioned Hyperion Research to measure and convey the impact, value, and scientific return on research (RoR) conducted at NERSC's leadership supercomputing facility. The study built upon Hyperion Research's existing framework for assessing scientific RoR conducted at global leadership computing sites and tailored it to NERSC's mission. The results of the study place the value and impact of the work accomplished at NERSC among the best in the world.

NERSC HPCs Throughout the Years

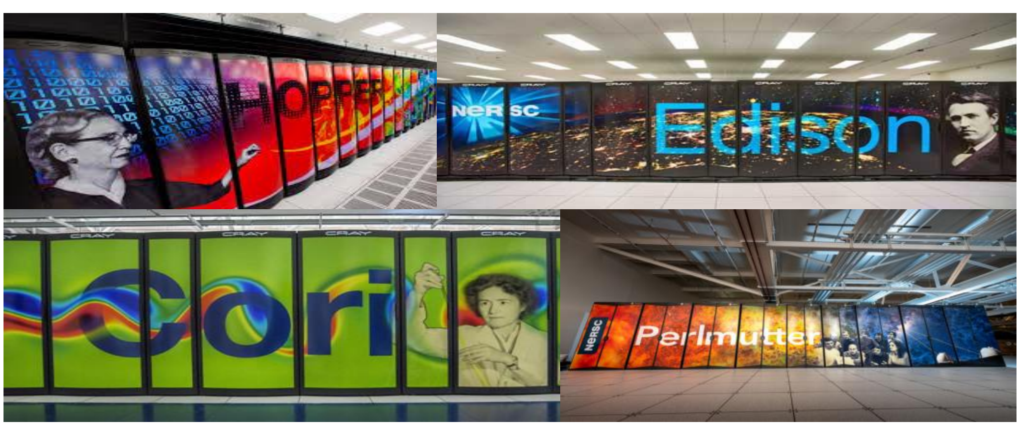

Source: Hyperion Research, 2024

TABLE OF CONTENTS

TABLE OF CONTENTS — Continued

LIST OF TABLES

LIST OF FIGURES

EXECUTIVE SUMMARY

Based on the results of this study, the value and impact of the research conducted at NERSC is among the best of global leadership computing sites. The study included reviews of 6 Nobel prize-winning research endeavors and interviews with 42 leading scientists and researchers who have utilized NERSC resources over its 50 years of operation.

Hyperion Research's framework to capture the importance and impact of research conducted using leadership class HPC computational capabilities was employed and focused on the return of predominantly scientific research, as opposed to heavily applied research for engineering or business process advancements. The study consistently demonstrated NERSC research provided the highest levels of importance and impact on the framework's innovation class index mapping than delivered at many other HPC leadership sites.

The breadth of scientific accomplishments delivered by research conducted at NERSC supported all ranges of scientific contributions defined in the value framework. NERSC science programs covered by the study participants included:

- Chemical sciences, geosciences
- Earth and environmental systems
- Fusion energy sciences
- Materials science and engineering
- Nuclear physics
- Scientific user facilities

While each of the science programs above primarily supported the framework contribution area of “advancing science” and to a lesser extent, “advancing engineering”, the other framework contribution areas of “economic growth”, “saving lives”, and “national security (non-defense-related)” were also represented.

Beyond the impact and value of the research conducted at NERSC, study participants also considered its operational characteristics to be the standard by which all other scientific computing facilities should aspire to.

Key Findings

Value and impact of research performed at NERSC is among the best of global leadership computing sites

Reflective of both the caliber of researchers and scientists utilizing NERSC resources and the capabilities of the resources over its 50 years of operation, the value and impact of the research and resulting scientific discoveries are among the best in the world.

FIGURE 1

Innovation Class Research Mapping

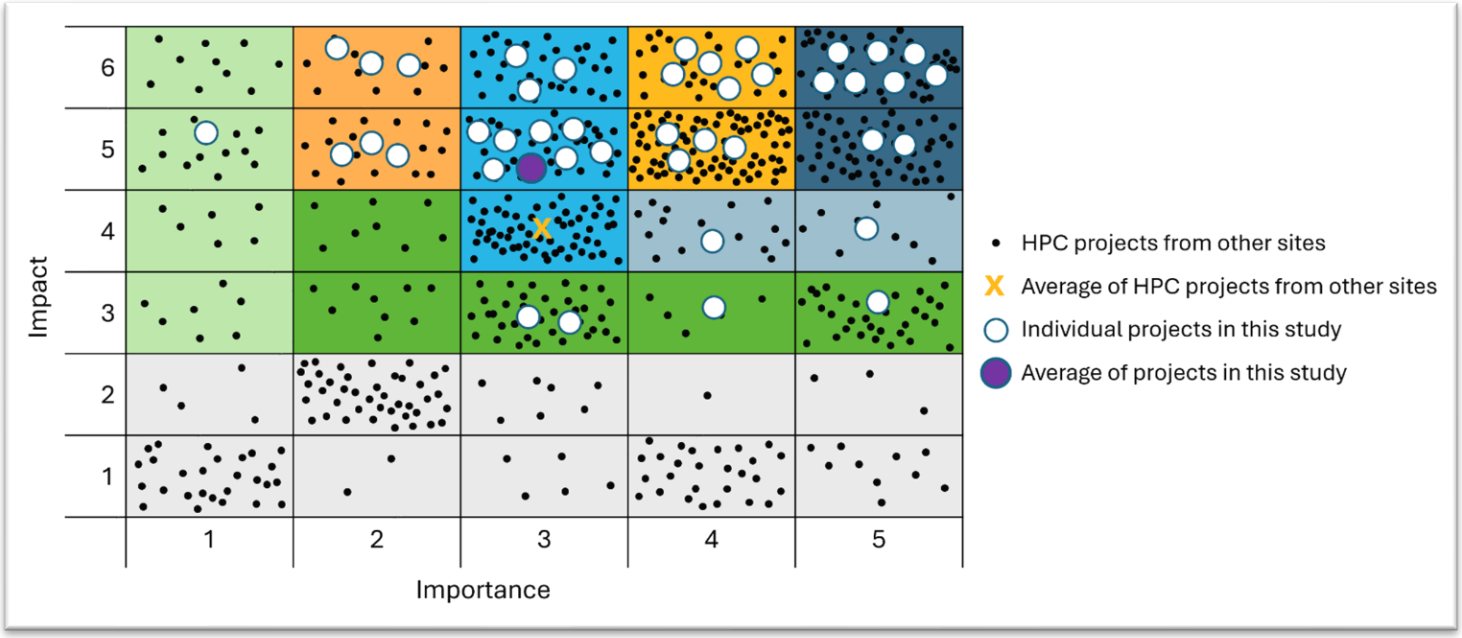

Source: Hyperion Research, 2024

Research performed at NERSC reflects multi-dimensional diversity

Scientific Coverage Areas

NERSC research represented a breadth of scientific discovery areas. Nearly all aspects of science were advanced by the scientific categories being explored by the researchers in the study, including:

- Chemistry (catalysis, chemical physics, physical chemistry)
- Earth systems (coupled systems)
- Energy (general, fusion)
- Geoscience (geochemistry)
- Materials Science (biomolecular materials, energy storage, general)
- Physics (accelerator science, astrophysics, condensed matter, cosmology, high energy [experimental, theoretical], nuclear physics [experimental, theoretical], plasma physics)

Scientific Organizations

Participating principal investigators (PIs) come from a wide range of organizations including leading academic research institutions, multiple US Department of Energy (DOE) national laboratories, and DOE Office of Science user facilities. Broadly collaborating with scientists at additional US-based and global universities, PIs also leveraged additional global computational capabilities at other DOE labs and European supercomputer centers.

Advancing science is the clear mission of all researchers

The scientific value framework utilized in this study defines five scientific contribution areas: advancing science, advancing engineering, economic growth, saving lives, and national security (non-defense-related). Advancing science was identified by all participants as the primary scientific contribution area, with advancing engineering identified by almost half. The other three, listed above, were also identified by several participants.

NERSC accomplishments go beyond scientific discovery

While scientific discovery is the primary objective of the research conducted at NERSC, there are other notable accomplishments as a result of 50 years of research.

Breaking the terascale performance barrier

Similar to the recent fanfare from the first exascale machine coming online (Frontier at ORNL) and related applications that could leverage exaflop performance at scale, the first at-scale application to leverage teraflop performance capabilities occurred at NERSC in 1998. As part of a DOE Grand Challenge project, researchers were able to model 1,024 atoms of a metallic magnet. Terascale computing allowed modelling of 1,024 atoms, significantly larger than what could've been modelled before.

International collaboration

Collaboration is a hallmark of the scientific community, and this was well-reflected in the study results. PIs worked closely with researchers at other academic and government research institutions, both within the U.S. and globally. And while the primary computational aspects of the research included in this study were performed at NERSC, PIs also leveraged the computing capabilities of their host organizations, as well as the institutions of the other researchers on their teams and of other U.S. and global HPC sites.

Bolstering global scientist and researcher bench strength

Multiple study participants identified several research outcomes that amplifies the capabilities of fellow scientists and researchers. *Increases the productivity, efficiency, and success of scientists and researchers* and *building a more diverse and stronger scientific workforce* were the most commonly recognized outcomes for a number of the research projects.

NERSC was resoundingly praised by scientists and researchers

Beyond the impact and value of the research conducted at NERSC, study participants also considered its operational characteristics to be the standard by which all other scientific computing facilities should aspire to. NERSC performs at a top leadership level across all dimensions valued by global leading scientists and researchers:

- Continued priority and excellence in computational capabilities
- Quality, reliability, and uptime of the site's resources
- Straightforward process by which to obtain time on its machines
- Superior support whenever questions or issues arise

HYPERION RESEARCH FRAMEWORK FOR ASSESSING VALUE AND RETURN OF SCIENTIFIC RESEARCH

With support from the U.S. Department of Energy (DOE), Hyperion Research created the HPC Return on Investment (ROI) framework to help provide insight into HPC investments by examining individual HPC projects and measuring the amounts spent on the HPC resources compared with the projects' financial and innovation returns. The ROI framework includes a macroeconomic model that depicts the way HPC investments result in economic advancements in the form of ROI in revenue and profits (and cost savings), as well as a Return on Research (ROR) innovation class index. The ROR index provides a means of measuring and comparing innovation levels based on two parameters: the importance of each innovation combined with how many organizations can use the innovation.

DOE HPC systems face unique challenges in demonstrating their return on investment due to the complex interplay of costs, utilization patterns, and diverse user needs. The massive scale of these systems requires substantial initial investment along with ongoing operational costs for power, cooling, and maintenance. Yet measuring direct value creation is complicated by factors like shared resource allocation across multiple research domains, varying project timelines, and the difficulty of quantifying scientific breakthroughs in monetary terms.

For this study on NERSC research, Hyperion Research emphasized the ROR metric because the ROI measurements fall short in capturing the full value of government-funded scientific computing. While private sector HPC investments can be evaluated through direct financial returns, the NERSC facility primarily advances fundamental science and engineering - areas where breakthroughs may take decades to translate into economic impact, if at all. The ROR framework focuses on each project's innovation importance and impact, providing a more suitable metric for DOE's mission-driven research portfolio. It also allows for meaningful comparison of diverse projects. This approach also acknowledges that many NERSC projects create foundational knowledge that enables subsequent innovations across multiple sectors, making their true value much broader than immediate financial returns would suggest.

Innovation Return on Research (ROR) Metrics

Hyperion Research employs a comprehensive rating system to evaluate innovations resulting from HPC projects, focusing on both the importance and impact of these innovations. This approach allows for a two-dimensional understanding of how significant each innovation is within its field over the past decade and the extent to which it benefits various organizations.

Importance

The importance of innovations is measured on a scale from 1 to 5, where higher ratings indicate greater significance. An innovation rated as a 5 is recognized as one of the top three discoveries in the field, while a rating of 1 places the innovation among the top fifty. This categorization reveals a concentration of noteworthy advancements, as many entries fall within the higher ranges:

The importance ratings indicate the following:

- Rating 5: Top 3 discoveries in the last decade
- Rating 4: Top 5 discoveries in the last decade
- Rating 3: Top 10 discoveries in the last decade
- Rating 2: Top 25 discoveries in the last decade

- Rating 1: Top 50 discoveries in the last decade

Impact

The impact ratings, ranging from 1 to 6, with higher scores indicating broader utility across multiple organizations. Innovations rated 6 are deemed beneficial to over fifty organizations, while a rating of 1 suggests recognition only among experts:

The impact ratings indicate the following:

- Rating 6 (Useful to over 50 organizations)
- Rating 5 (Useful to 10 to 49 organizations)
- Rating 4 (Useful to 6 to 9 organizations)
- Rating 3 (Useful to 2 to 5 organizations)
- Rating 2 (Only useful to 1 organization)
- Rating 1 (Recognized ONLY by experts)

Innovation Class Index

To synthesize these Importance and Impact evaluations, Hyperion Research assigns an overall Innovation Class rating, which combines importance and impact. Innovations are classified from Class 1, representing the highest tier of significance and utility, to Class 8, which encompasses all other innovations.

- Class 1 innovations – One of the top 1-3 innovations in a field over the last ten years PLUS useful to over 10 organizations
- Class 2 innovations -- One of the top 5 innovations in a field over the last ten years PLUS useful to over 10 organizations
- Class 3 innovations – One of the top 5 innovations in a field over the last ten years PLUS useful to over 5 organizations
- Class 4 innovations – One of the top 10 innovations in a field over the last ten years PLUS useful to over 5 organizations
- Class 5 innovations – One of the top 25 innovations in a field over the last ten years PLUS useful to over 10 organizations
- Class 6 innovations – One of the top 25 innovations in a field over the last ten years PLUS useful to at least 2 organizations
- Class 7 innovations – One of the top 50 innovations in a field over the last ten years PLUS useful to at least 2 organizations
- Class 8 innovations – All other innovations

This structured approach highlights the critical role of HPC in driving meaningful advancements and underscores the importance of evaluating both significance and utility in the assessment of technological innovations. The figure below shows the structured relationship of importance and impact ratings to the Innovation Class index.

Innovation Class Research Mapping

Mapping the Innovation Class areas within a scatter plot of importance and impact ratings brings the story together. Displaying the research projects data this way allows for a clear and immediate visual representation of the relationship between importance and impact. Incorporating innovation classes into the plot adds a layer of meaning that helps contextualize the projects and is an intuitive way for the audience to understand how the various projects compare to one another.

FIGURE 2

Innovation Class Research Map

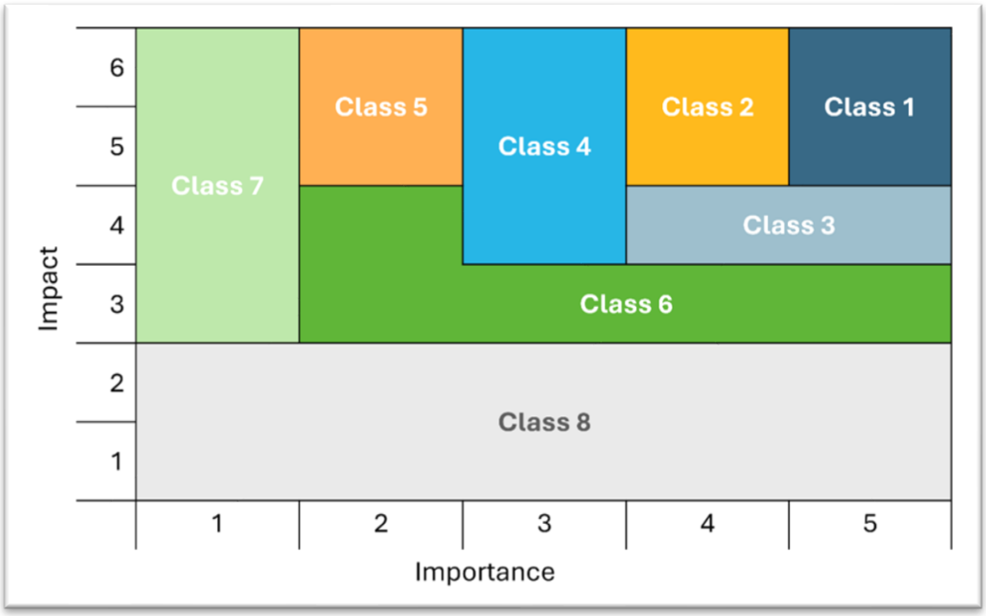

Source: Hyperion Research, 2024

This visualization creates an effective method of comparing dissimilar research projects, which works for the purposes of comparing the 42 NERSC projects in this study, as well as comparing the NERSC projects to a broader range of HPC projects surveyed overall by Hyperion Research.

STUDY RESULTS

The following are the results of the 42 NERSC projects studied. Results were collected via in-depth interviews with the individual project PIs.

Results of Importance Evaluations

The NERSC innovations are categorized based on their importance as follows:

- Rating 5 (Top 3 discoveries in the last decade): 11 research projects
- Rating 4 (Top 5 discoveries in the last decade): 12 research projects
- Rating 3 (Top 10 discoveries in the last decade): 12 research projects

- Rating 2 (Top 25 discoveries in the last decade): 6 research projects
- Rating 1 (Top 50 discoveries in the last decade):1 research project

23 research projects (approximately 55% of the total projects studied) are rated 4 or 5, suggesting that a significant portion of the research described in these interviews represents leading discoveries that have shaped their respective fields. Notably, 11 innovations fall into the highest category (5), indicating they are among the top 3 discoveries in the last decade.

Results of Impact Evaluation

The impacts of the NERSC innovations were assessed with the following distribution:

- Rating 6 (Useful to over 50 organizations): 19 research projects
- Rating 5 (Useful to 10 to 49 organizations): 17 research projects
- Rating 4 (Useful to 6 to 9 organizations): 2 research projects
- Rating 3 (Useful to 2 to 5 organizations): 4 research projects
- Rating 2 (Only useful to 1 organization): 0 research projects
- Rating 1 (Recognized ONLY by experts): 0 research projects

37 innovations (approximately 88%) are rated 5 or higher for impact, indicating broad utility. The high number of innovations rated 6 (19 entries) suggests a substantial capacity for widespread application.

Innovations with lower impact scores (e.g., ratings of 1, 2 or 3) suggest that while they may hold importance in their field, their practical utility may be limited to fewer organizations. This could indicate niche applications such as emerging technologies that have not yet gained widespread adoption, or "science for the sake of science".

Class Distributions

The results from the Innovation Class categorization of 42 NERSC research projects studied reveal a compelling narrative about the significance and utility of these innovations. The average importance score of 3 and an impact score of 5 indicate that these projects are not only noteworthy within their individual fields but also possess broad applicability across multiple organizations.

The distribution of projects across innovation classes is particularly interesting, with 9 projects in Class 1 and 10 in Class 2, highlighting that a substantial portion of the innovations are both critically important and impactful. Additionally, the presence of 11 projects in Class 4, and 6 in Class 5 illustrates a diverse range of innovations that, while not at the very top tier, still contribute significantly to their respective domains.

FIGURE 3

Innovation Class Research Mapping: NERSC Projects

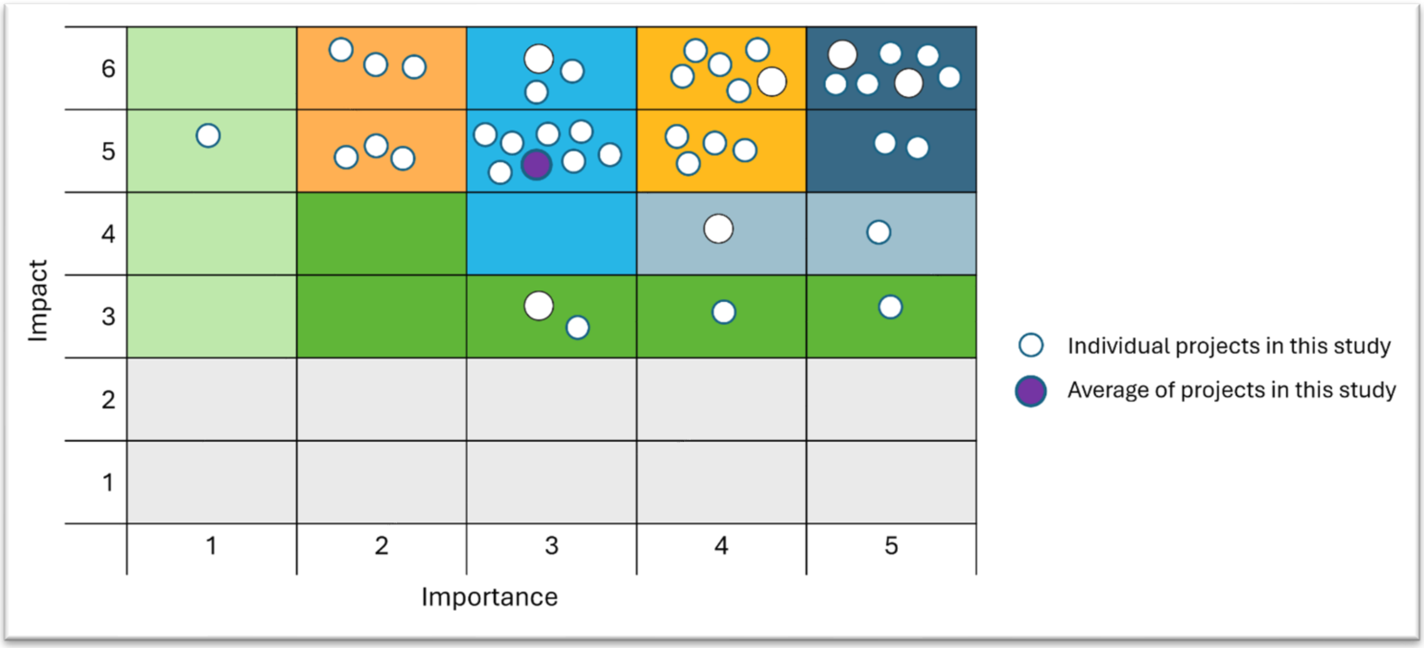

Source: Hyperion Research, 2024

Notably, a considerable portion of the research assessed in this study shows a high potential for substantial societal impact, indicating that many of these innovations not only advance their respective fields but also provide significant benefits to a broad range of organizations.

Comparison with Other HPC Data Centers

When comparing the NERSC study to the larger Hyperion Research dataset of 650 HPC projects conducted at HPC centers around the world, the unique class distribution becomes even more pronounced. Nearly a 20% of the larger HPC project dataset falls into Class 8, while none of the NERSC projects are in this lowest rating area. Class 8 projects can have important discoveries but are only useful to 1 organization or only recognized by experts in a specific field. This indicates that NERSC projects are intentionally built with the goal of sharing innovations with larger scientific communities.

This point is also illustrated when looking at the average importance and impact rating of NERSC projects compared to the larger Hyperion Research dataset of 650 HPC projects. NERSC projects average an importance score of 3, and an impact score of 5. The larger HPC dataset on the other hand has an average importance score of 3, and an impact score of 4. While the two project datasets share an average importance score, NERSC's impact score is higher. Again, this suggests that NERSC project innovations are intentionally made to be shared, growing the scientific community's understanding of basic research discoveries.

FIGURE 4

Innovation Class Research Mapping: Comparison with Other HPC Projects

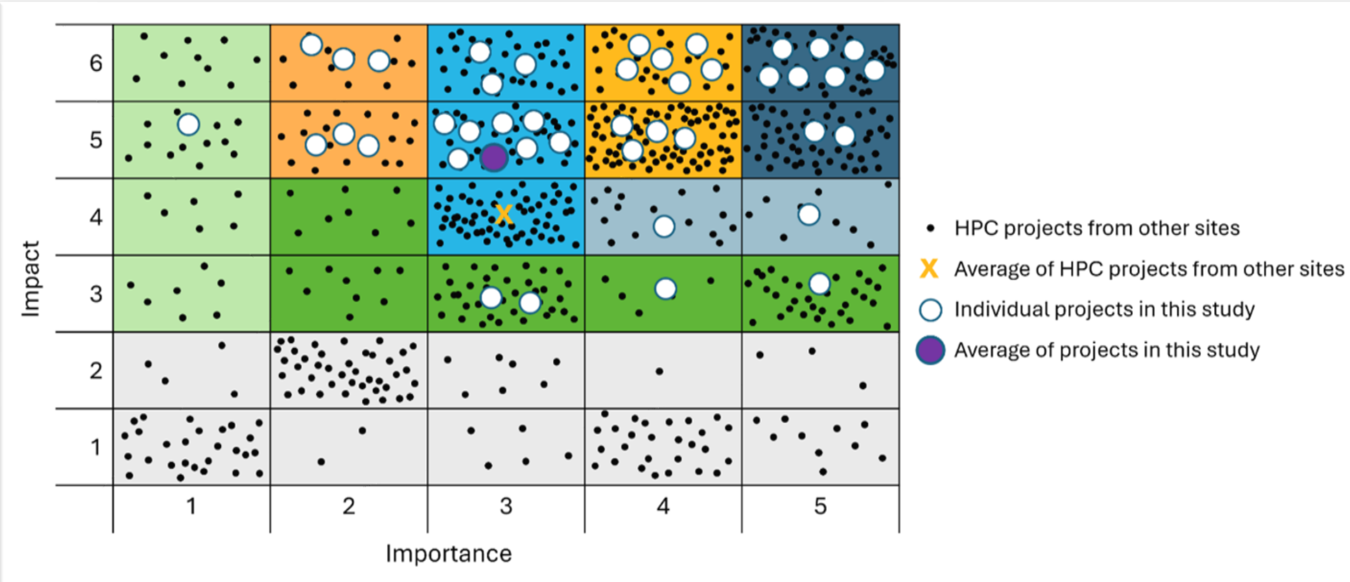

Source: Hyperion Research, 2024

Overall, the findings from this study, supported by the structured evaluation of importance and impact, highlight the significant role of NERSC in driving meaningful scientific advancements. By mapping these innovations within the Innovation Class framework, the scatter plot not only visualizes the data effectively but also provides a comprehensive understanding of how these projects compare within the broader landscape of research. This approach offers valuable insights into the potential of these innovations to drive progress, emphasizing the necessity of evaluating both significance and utility in assessing technological impact.

Select Research Project Profile and Summaries

The following projects are highlighted for their high ranking within the Innovation Class framework (Class 1's) and/or were featured projects in the NERSC 50th anniversary key findings report. These projects exemplify cutting-edge research that leverages NERSC computing to address critical challenges in various scientific fields.

Frontiers in Accelerator Design: Advanced Modeling for Next-Generation BES Accelerators

Robert Ryne's project at Lawrence Berkeley National Laboratory (LBNL) focuses on the design and development of particle accelerators, particularly those utilized as light sources by the Department of Energy's Office of Basic Energy Sciences (BES). These accelerators come in various configurations, including circular and linear designs, and support both the current and future light sources. The research emphasizes creating high-quality particle beams, improving beam quality, making accelerators more compact and cost-effective, and developing innovative methods for manipulating particle beams, commonly referred to as "beam gymnastics." Key facilities involved in this project include the Advanced Light Source (ALS) at Berkeley National Laboratory, the Advanced Photon Source (APS) at Argonne National Laboratory (ANL), and the Linac Coherent Light Source (LCLS) at Stanford University.

The implications of Ryne's work extend across multiple scientific disciplines and practical applications. Particle accelerators have evolved from tools primarily for particle physics research to essential instruments in materials science, chemistry, biology, environmental research, and national security. The light sources developed through this project enable researchers to investigate the structure and behavior of materials at molecular and atomic levels. In biosciences, these accelerators facilitate the determination of protein structures, aiding pharmaceutical drug design. Additionally, ultra-short X-ray pulses from free-electron lasers allow real-time imaging of biochemical processes, providing unprecedented insights into essential biological mechanisms. Overall, this project significantly advances science and engineering, fosters economic growth, and contributes to life-saving medical advancements.

Development of New Force Fields Using Machine Learning and First Principles Physics

Subramanian Sankaranarayanan's project at Argonne National Laboratory (ANL) focuses on creating new force fields for molecular dynamics simulations by integrating machine learning with the first principles of physics. This research aims to enhance the understanding of atomic interactions at the atomistic and nanoscale levels, facilitating the evolution of systems over time. By employing machine learning algorithms to parameterize physics-based models derived from computationally intensive ab initio calculations, the team has significantly reduced the model development time from years to mere weeks or days. This efficiency allows for larger system simulations over extended time scales, yielding insights into a variety of materials, including 2D materials and complex systems like water.

The implications of this research extend beyond theoretical advancements, enhancing the efficiency of experimental facilities and informing experimental design. By utilizing these models, scientists can gain a deeper understanding of material properties and behavior before conducting experiments, optimizing the use of resources at facilities like the Advanced Photon Source. This project has also produced significant real-world applications, such as a model that accurately describes water's thermodynamic behavior and anomalies, which has relevance in mechanical engineering and environmental science. Furthermore, the development of eco-friendly lubricants for automotive engines demonstrates the potential to replace hazardous additives, highlighting the project's impact across various industries, from microelectronics to energy efficiency. Through these advancements, the research not only propels scientific discovery but also fosters engineering innovations.

Three-Dimensional Simulations of Core-Collapse Supernovae

Adam Burrows at Princeton University is focused on simulating core-collapse supernovae, the explosive deaths of massive stars, to understand the mechanisms behind these cosmic events that lead to the formation of neutron stars or black holes. The project involves developing computational tools to solve intricate equations of radiation hydrodynamics, modeling the turbulent processes that occur during a star's core collapse and subsequent rebound, which generates a shockwave. These three-dimensional simulations require significant supercomputing resources to capture the complex interplay of fluid dynamics, nuclear physics, and neutrino interactions, aiming to accurately depict the chaotic evolution from core collapse through explosion and the ejection of stellar material into space.

The implications of this research extend far beyond the immediate study of supernovae, as these cosmic explosions are responsible for creating and disseminating many of the heavy elements essential for life on Earth. By enhancing our understanding of supernovae, the project provides insights into the chemical evolution of galaxies and the broader universe. The computational methods and tools developed for this work also have applications across various scientific and technological fields, bridging astrophysics, nuclear physics, and high-performance computing. Although the immediate practical applications may be limited, this fundamental research enriches human knowledge and may inspire future technological advancements, highlighting the critical role of supercomputing in addressing complex scientific challenges.

Cluster and Nanostructure for Energy and Bio Applications

Puru Jena's research at Virginia Commonwealth University focuses on theoretical physics and materials science, particularly in the context of energy production and storage. Utilizing high-performance computing resources at NERSC, the team conducts fundamental calculations to study materials relevant for hydrogen production, solar energy, and battery storage. A key aspect of their work is the development of novel catalysts, including innovative "super atom catalysts" designed to replace expensive noble metals with more efficient and cost-effective alternatives. Additionally, the research explores improvements in battery technology, investigating sodium-based and solid-state batteries while optimizing properties such as conductivity and interface stability.

The potential impact of this research is substantial, especially in advancing clean energy and sustainable technologies. By creating more efficient catalysts, the work could lead to significant breakthroughs in hydrogen production, CO2 conversion, and other vital chemical processes, aiding the transition to a hydrogen-based economy and mitigating climate change. In energy storage, the development of safer, more efficient battery technologies could revolutionize electric vehicles and grid-scale energy solutions, making renewable energy more feasible and reducing dependence on fossil fuels. The team's work on halogen-free electrolytes shows promise for creating safer, less toxic battery technologies, contributing to solutions for pressing energy and environmental challenges facing society today.

A Framework for Improving Analysis and Modeling of Earth System and Intersectoral Dynamics at Regional Scales

Project Hyperfacets, led by Paul Ulrich at the University of California, Davis (UC-Davis), aims to enhance the analysis and modeling of Earth system dynamics by collaborating with stakeholders to identify their climate data needs and align them with scientific inquiries. The project investigates various climate datasets to assess their applicability to stakeholder questions, focusing on significant historical events that have influenced decision-making. Utilizing multiple data sources, including NASA and simulations from the Energy Exascale Earth System Model (E3SM) on NERSC supercomputers, Project Hyperfacets works primarily with water managers and electrical system operators across four major U.S. watersheds. By generating high-resolution simulations that reduce computational costs while providing detailed data, the project employs a "Storyline framework" to recreate past events and project how they may evolve under future climate conditions.

By bridging the gap between climate science and practical decision-making, Project Hyperfacets empowers stakeholders to understand their vulnerability to climate change and assess their resilience. This innovative approach not only delivers historical and projected climate data but also explores the underlying science, directly impacting decisions in water management and electrical systems operation. By simulating potential future changes in high-impact events, the project aids stakeholders in evaluating their current strategies, ultimately contributing to disaster reduction, enhancing energy efficiency, and fostering climate change preparedness to save lives and prevent economic losses.

High Precision Calculations of the Nucleon Structure for Fundamental Symmetries

Rajan Gupta's project at Los Alamos focuses on understanding the fundamental properties of neutrons and protons, collectively known as nucleons, through complex computer simulations based on Quantum Chromodynamics (QCD). The research aims to elucidate how nucleon properties emerge from their constituent quarks and gluons, with particular emphasis on explaining the spin of these particles and the distribution of their electric charge. A critical objective of the project is the search for an electric dipole moment in neutrons, which, if detected, would indicate a significant violation of charge conjugation and parity (CP) symmetry, suggesting physics beyond the current Standard Model.

The implications of this research are profound, as it has the potential to advance our understanding of fundamental physics and reveal new insights into the nature of the universe. By probing the structure of matter at its most basic level, Gupta's work could help explain the matter-antimatter asymmetry that dominates our universe, offering a possible mechanism for this phenomenon. Such discoveries would not only validate and extend existing theories in particle physics but also bridge the gap between particle physics and cosmology, leading to a more comprehensive understanding of the universe's evolution and the fundamental laws governing it.

Structure Refinement of Spike-Protein of SARS-COV-2

Dr. Wai-Yim Ching's project at the University of Missouri - Kansas City focuses on the structural refinement of the spike protein of SARS-CoV-2, the virus responsible for COVID-19. Utilizing an in-house computational package, Dr. Ching and his team perform complex calculations to ensure the accuracy and optimization of the spike protein's structure, analyzing detailed atomic bonding and properties. This research, which has been ongoing for three years, emphasizes a holistic approach to understanding the entire complex structure of the virus, collaborating with biological modeling expert Dr. Puja Adhikari to enhance their insights.

While the immediate applications of this work may not yet be fully realized, it holds significant long-term potential for understanding the molecular functioning of SARS-CoV-2, which could inform strategies for combating COVID-19 and similar viruses in the future. This foundational research may lead to increased citations and recognition, ultimately contributing to improved healthcare outcomes and saving lives by providing critical knowledge for researchers and medical professionals. Dr. Ching's interdisciplinary approach bridges physics and biology, addressing complex biological systems and potentially inspiring similar methodologies in other scientific fields.

SUNCAT-FWP Catalysis NERSC Project

Michal Bajdich's research at Stanford University focuses on catalysis for key chemical reactions, particularly carbon dioxide reduction, nitrogen reduction and water splitting, utilizing electrocatalysis, thermal catalysis, and advanced machine learning techniques. A notable project involves the development of an iridium-based catalyst optimized through computational methods, demonstrating exceptional performance in water splitting and surpassing Department of Energy (DOE) activity requirements. This breakthrough has led to collaboration with Plug Power, the largest hydrogen electrolyzer manufacturer in the U.S., and garnered additional DOE funding for further research optimization.

Bajdich's work significantly contributes to the development of efficient and sustainable energy technologies, particularly in enhancing hydrogen production, which is vital for the transition to a clean energy economy. By improving the efficiency of water electrolysis, this research promotes the adoption of hydrogen as a clean fuel source, helping to reduce reliance on fossil fuels and mitigate climate change. Beyond hydrogen production, the project has implications for metal-air batteries and eco-friendly metal extraction, potentially leading to cleaner industrial processes. Overall, Bajdich's innovative research is poised to address critical global energy and environmental challenges while advancing scientific and engineering practices.

First Principles Simulations of Nanostructures

Giulia Galli's research at the University of Chicago employs quantum simulations to explore matter at the atomic level, focusing on solids, liquids, molecules, and nanostructures. Her work encompasses a range of applications, including solar energy harvesting, water splitting for fuel generation, and the development of materials for quantum technologies. Over the past two decades, Galli and her team have adapted their computational simulations, grounded in quantum mechanics, to interpret experimental results and design materials tailored for specific applications, emphasizing a close collaboration with experimentalists to ensure practical outcomes.

The societal impact of Galli's research is profound, particularly in advancing quantum technologies that could revolutionize computing, sensing, and communication. By identifying and designing materials with optimal properties for quantum applications, her work is crucial for creating more powerful computers, ultra-sensitive sensors, and secure communication systems. Additionally, her contributions to sustainable energy through improved materials for the oxygen evolution reaction enhance the performance of catalysts critical for water splitting, thereby facilitating the transition to clean energy technologies. Galli's interdisciplinary approach not only accelerates the path from fundamental science to practical applications but also advances computational methods in materials science, fostering breakthroughs across various scientific and engineering fields.

Energy Exascale Earth System Modeling (E3SM)

Ruby Leung's project at Pacific Northwest National Laboratory (PNNL) involves the development of the Energy Exascale Earth System Model (E3SM), a large-scale initiative that collaborates with eight national laboratories to enhance understanding of the interactions between energy use and climate change. Launched in FY15, E3SM aims to leverage exascale computing capabilities to create a sophisticated Earth system model that can run at high resolutions (down to 3 km globally), incorporating both natural Earth systems and human activities. This model features variable resolution capabilities, enabling targeted analysis of specific regions, such as North America and Antarctica.

The implications of the E3SM project are significant for addressing climate change and energy challenges, as it provides detailed simulations that inform predictions of future climate scenarios and their impacts on energy infrastructure and renewable sources. This information is vital for policymakers and industry leaders in shaping energy policies and climate mitigation strategies. Moreover, the project's success in utilizing exascale computing marks a breakthrough in climate science, recognized by a Gordon Bell award, and allows for more realistic simulations of the complex interplay between human society and the environment. By fostering effective strategies for balancing energy needs with environmental sustainability, E3SM contributes to the transition toward cleaner energy sources and more resilient infrastructure in the face of climate change.

High Performance Simulations for Regional Scale Earthquake Hazard and Risk Assessments

David McCallen's project, EQ SIM (Earthquake Simulation Framework), at Lawrence Berkeley National Laboratory (LBNL) focuses on modeling earthquake scenarios, particularly in the San Francisco Bay Area, with an emphasis on the Hayward Fault due to its significant seismic hazard. The framework simulates the entire process from fault rupture to ground motion propagation and its interaction with infrastructure, aiming to understand the distribution of ground motions across various earthquake scenarios and predict potential damage patterns. Utilizing high-performance computing resources, including GPU-accelerated systems such as Frontier and Perlmutter, the team runs multiple simulations to account for uncertainties in earthquake dynamics, with results compiled into an Open Access database hosted by the Pacific Earthquake Engineering Research Center (PEER).

This research is crucial for enhancing earthquake preparedness and improving infrastructure design in seismically active regions. By providing high-fidelity simulations, EQ SIM addresses critical gaps in observational data, particularly for faults like the Hayward Fault that have not experienced major earthquakes in over 150 years. The Open Access database will serve as a valuable resource for researchers and engineers, enabling better insights into earthquake behavior and facilitating the design of resilient structures. Ultimately, this work contributes to disaster reduction efforts by improving the understanding of earthquake hazards, informing effective mitigation strategies, and potentially saving lives while reducing economic impacts from future seismic events.

Synthetic Spectra of Astrophysical Objects

Eddie Baron's project at the University of Oklahoma focuses on modeling the spectra of Type Ia supernovae, which are thermonuclear explosions of white dwarf stars. The research aims to understand the origins of these explosions, the stellar systems they arise from, and the stellar evolution processes that produce the observed characteristics of supernovae. By utilizing spectroscopy to analyze the elements, velocities, temperatures, and densities of these explosive events, the team seeks to develop a canonical model that enhances understanding of these complex phenomena.

This research substantially advances knowledge of stellar evolution, element formation, and cosmic processes, providing insights into how the elements that constitute our world are created and distributed throughout the universe. Additionally, the project trains a diverse group of individuals in advanced computing techniques, fostering a skilled workforce adept at solving complex computational challenges. The computational methods developed are applicable beyond astrophysics, contributing to high-performance computing and offering transferable skills relevant to sectors such as national security and machine learning. Consequently, Baron's work not only enriches scientific understanding but also has broad implications for technology and economic growth.

High Performance Computing in Support of the DIII-D National Fusion Facility

Sterling Smith's research at General Atomics supports fusion energy studies conducted using the DIII-D tokamak device, a donut-shaped vacuum chamber that employs powerful magnetic fields to confine and heat plasma to temperatures exceeding those at the sun's core. The primary objective is to understand how to efficiently confine and heat plasma to achieve fusion, where atomic nuclei combine to release vast amounts of energy. This research involves complex simulations of plasma behavior, with a particular focus on turbulence and heat transport, utilizing significant computational resources provided by NERSC. The team employs NERSC's real-time queue and Superfacility API to conduct rapid analyses between plasma experiments, which are conducted every 15 minutes during operational days.

The implications of this research are profound, as it contributes to the overarching goal of making fusion a viable and practical energy source. Fusion holds the promise of providing clean, safe, and virtually limitless energy, which could revolutionize the global energy landscape and aid in combating climate change. By enhancing our understanding of plasma dynamics and validating theoretical models, this work instills confidence in predictions regarding future fusion devices. High-fidelity equilibrium reconstructions performed at NERSC are crucial for advancing statistical validations of various plasma models, ultimately paving the way for the design and construction of future fusion reactors capable of producing significant energy outputs. While still in the research phase, initiatives like this are steadily advancing the field toward realizing fusion as a transformative energy solution for the world.

Joint Genome Institute - Production Sequencing and Genomics

Kjiersten Fagnan's work at the DOE Joint Genome Institute (JGI) focuses on supporting scientific research in areas such as nutrient cycling, biofuels, microbial communities, and plant genomics through the generation of high-quality sequence data for various organisms. As a national user facility, JGI also produces metabolomics data and engages in synthetic biology initiatives. Since partnering with the National Energy Research Scientific Computing Center (NERSC) in 2010, JGI has scaled up its data production capabilities, enabling support for more users and the development of new analytical approaches, notably the Ex Biome project, which created a massively parallel genome assembler for complex samples like metagenomes.

The collaboration between JGI and NERSC has significantly advanced the field of genomics, enhancing our understanding of nutrient cycling, improving biofuel development, and providing insights into microbial community dynamics, all of which are crucial for addressing challenges in agriculture, environmental science, and sustainable energy. JGI's extensive data resources and user-friendly web interfaces support a vast community of around 2,000 scientists, facilitating research that explores complex metagenomic samples and opens new avenues for inquiry. This work has broad implications for various fields, including environmental science and human health, as it deepens our understanding of microbiomes and their effects on ecosystems and living organisms.

From Molecules to Continuum: Developing a Universal Approach for Accurate Description of X-ray Photo

Jin Qian's research at Lawrence Berkeley National Laboratory (LBNL) focuses on developing predictive tools for spectroscopic signatures, specifically X-ray Photoelectron Spectroscopy (XPS), to analyze nanoscale phenomena. By employing a specialized form of Density Functional Theory (DFT) known as real space Kohn-Sham DFT, the project can simulate systems with tens of thousands of atoms, greatly expanding beyond the limitations of conventional DFT, which typically handles around 100 atoms. This research addresses critical problems in chemistry and materials science, with potential applications in catalysis and battery research, including the ability to distinguish between various chemical environments, such as different oxidation states of gold or distinguishing between carbonate and bicarbonate species.

Notably, the research has produced surprising findings relevant to carbon capture, revealing an unexpected distribution of carbonate and bicarbonate species in liquid droplets that serve as proxies for ocean surfaces. This discovery challenges classical electrostatic theory, showing a higher concentration of carbonate on the surface of nano-droplets, which could significantly impact our understanding of carbon capture processes at ocean interfaces. The broader implications of Qian's work extend into multiple fields, offering tools to predict and understand nanoscale phenomena that could lead to advancements in materials science, more effective catalysts, improved battery technologies, and enhanced carbon capture methods. While still foundational, this research positions itself as a major contributor to critical scientific and technological advancements over the coming decade.

Scientific Contribution Areas

The value framework utilized in this study identifies five primary scientific contribution areas: advancing science, advancing engineering, economic growth, savings lives, and national security (non-defense). Participants were asked to select which contribution area best applied to their research; multiple choices were allowed. Additionally, for each contribution area selected, participants were asked to select from a series of contribution-area-specific outcomes best applied to the contribution area(s) they chose; again, multiple choices were permitted.

Advancing science was the clear mandate for the participants and was identified by each researcher. Advancing engineering was also selected by almost half of the participants. Economic growth, saving lives, and national security were each selected but with far less frequency.

The following figures show the distribution of selections for the contribution areas, as well as the distribution of the various combinations of contribution area selected.

FIGURE 5

Scientific Contribution Areas - Individual Mentions

Q. Which of the following areas best describes the primary scientific contribution of your research? (Select all that apply)

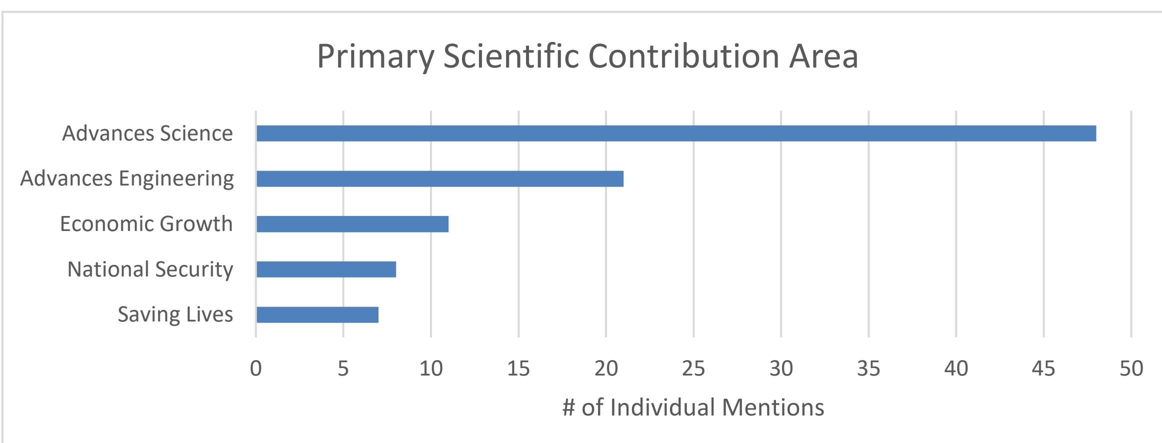

Source: Hyperion Research, 2024

FIGURE 6

Scientific Contribution Areas - Area Groupings

Q. Which of the following area best describes the primary scientific contribution of your research? (Select all that apply)

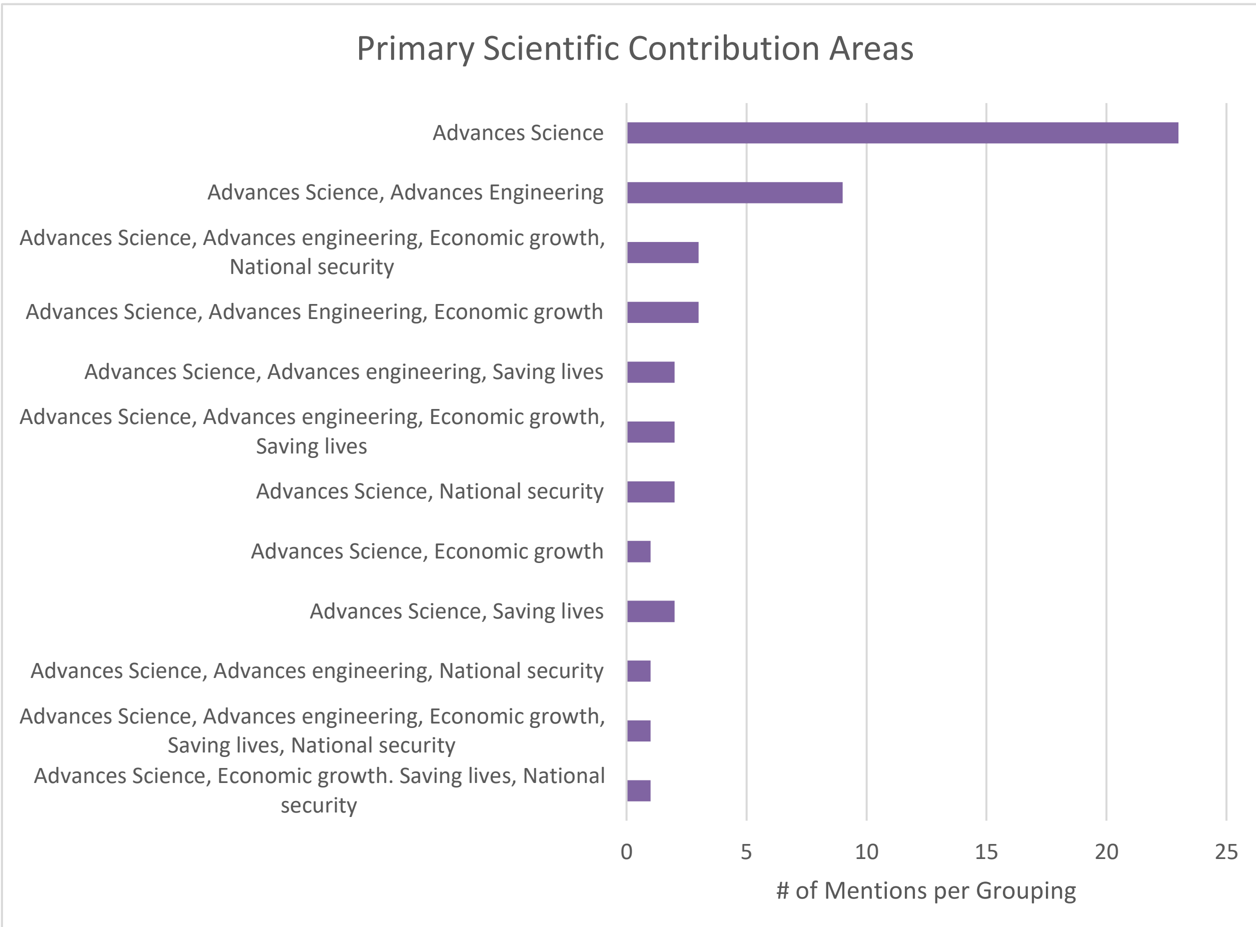

Source: Hyperion Research, 2024

Advancing Science

Advancing science covers a broad range of topics. More common examples include:

- Measurements of scientific discoveries, innovations, and research results
  - The level of importance of the research results and how broadly useful it is to other researchers/organizations
  - Enhancing fundamental science as a prerequisite for advancements in applied research and development

    - Number of scientific discoveries/results/publications and how often the results are used or cited by others
- Better understanding of how things work
- Speeding up the time-to-solution on critical large problems
- Testing otherwise untestable theories, either due to physical limitations or undesirable (potentially adverse) human impact
- Finding ways to reduce climate change impacts, improving the health of our planet, improving energy efficiency, and creating more sustainable products and approaches
- Increasing the productivity, efficiency and success of scientists and researchers -- this can allow each researcher to be more productive

Better understanding of how things work, speeding up the time-to-solution on critical large problems, and *testing otherwise untestable theories* were the top advancing science outcomes selected.

FIGURE 7

Advancing Science Outcomes

Q. Which of the following best reflects the outcome of your research? (Select all that apply)

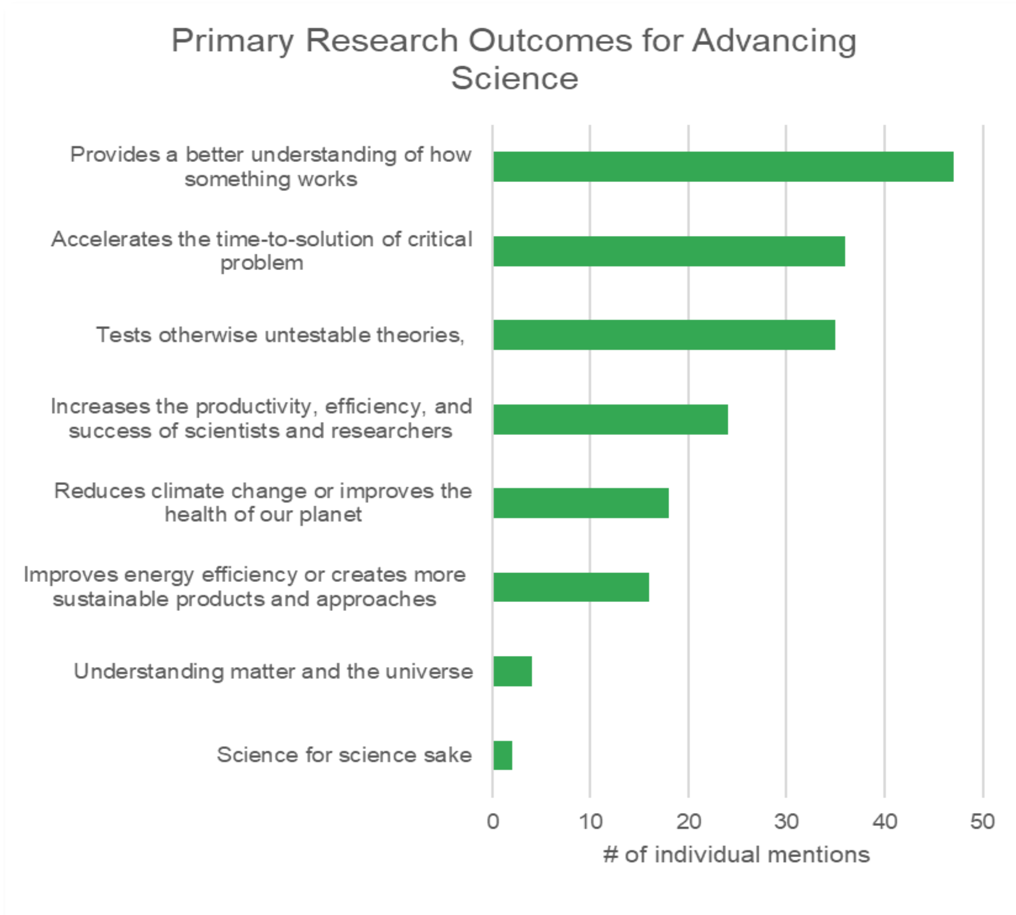

Source: Hyperion Research, 2024

Advancing Engineering

Advancing engineering also has a broad range of potential outcomes. Examples include:

- Creates better products and supports developing new types of products
- Designing better physical devices (e.g., designing better computers and making them more useable [e.g., GPUs])
- Accelerates the time from discovery to building products
- Testing otherwise untestable theories, either due to physical limitations or undesirable (potentially adverse) human impact
- Technology innovation transfer to industry (e.g., speeding up the time from discovery to products, improved understanding of how to build better products and services, developing new types of products, supporting industrial partnerships)

- Engineering better computing technologies (hardware, software, AI, etc.) for first-of-a-kind technologies (e.g. designing better physical devices, designing better computers and making them more useable [e.g., GPUs], creation of better software, data libraries/repositories and computing tools for broader use)
- Improving the design and safety of power plants and developing technologies to exploit alternative energy sources

Testing otherwise untestable theories and *engineering better computing technologies* were the top two advancing engineering outcomes.

FIGURE 8

Advancing Engineering Outcomes

Q. Which of the following best reflects the outcome of your research? (Select all that apply)

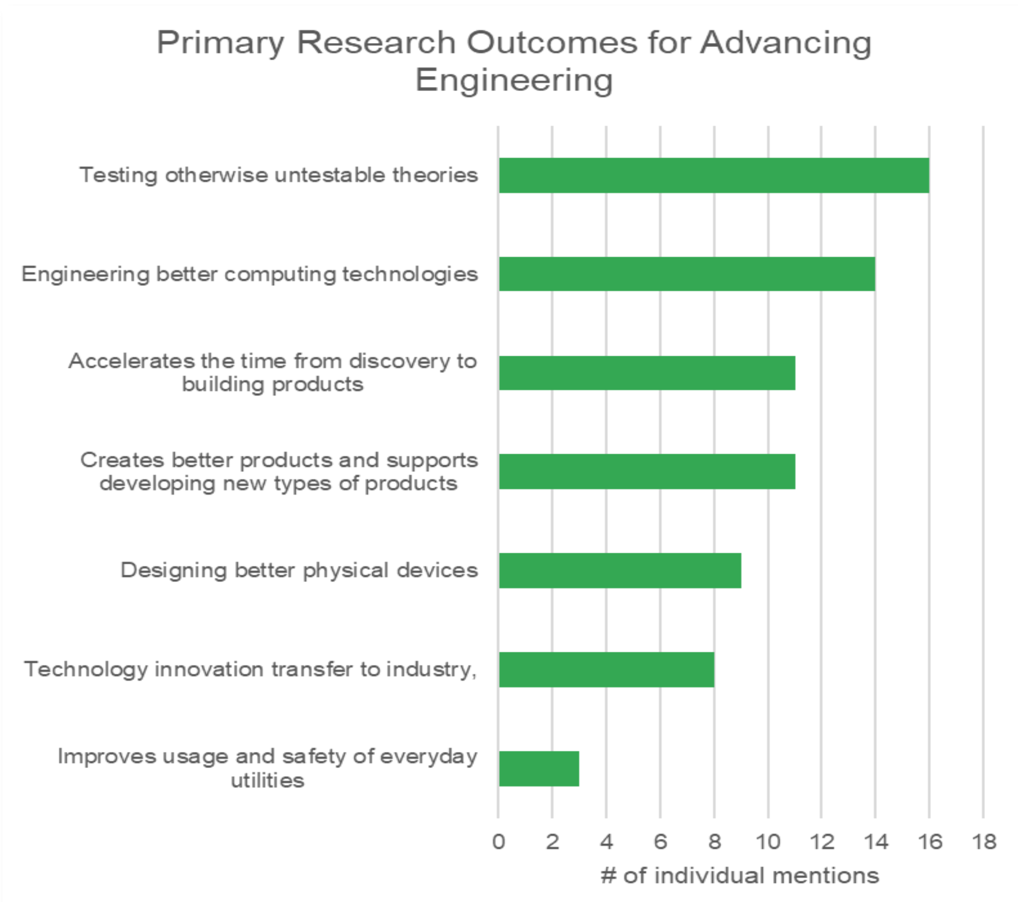

Source: Hyperion Research, 2024

Economic Growth

Economic growth had fewer outcomes in the context of this scientific research. Examples include:

- Helping to increase jobs and grow the nation's GDP
- Technology innovation transfer to industry (e.g., speeding up the time from discovery to products, improved understanding of how to build better products and services, developing new types of products, supporting industrial partnerships)
- Avoiding costs/saving money (e.g., simulation vs. high-cost testing can save substantial funds and have a faster time-to-solution)
- Building a more diverse and stronger scientific workforce

Building a more diverse and stronger scientific workforce and *technology innovation transfer to industry* were the top economic growth outcomes.

FIGURE 9

Economic Growth Outcomes

Q. Which of the following best reflects the outcome of your research? (Select all that apply)

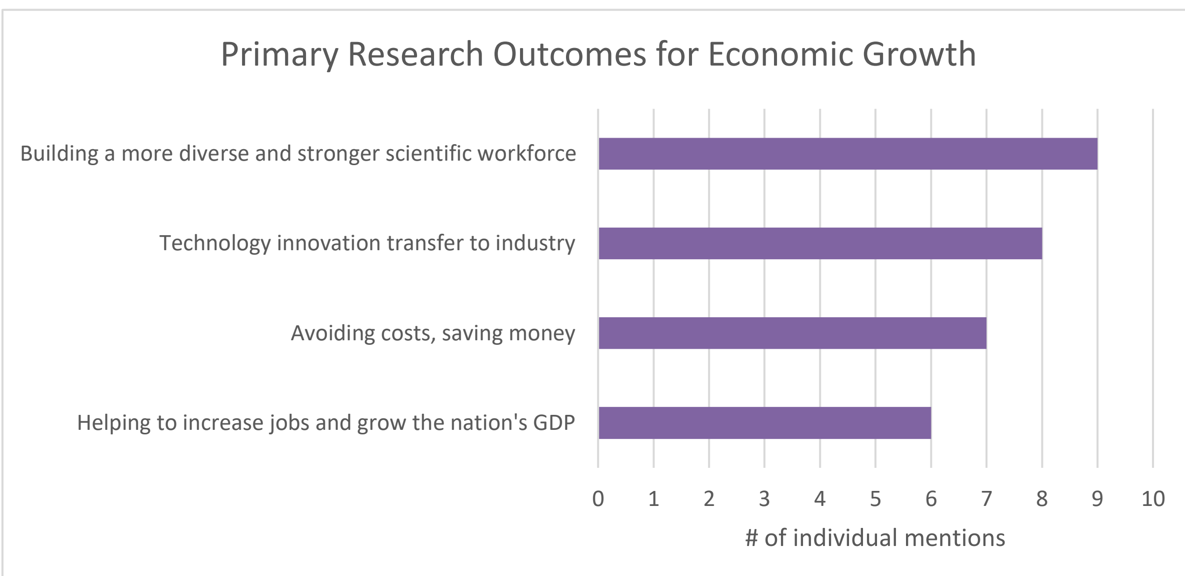

Source: Hyperion Research, 2024

Savings Lives

Several outcomes supported the saving lives contribution area, with *weather and disaster reduction projects* the most common selection.

FIGURE 10

Saving Lives Outcomes

Q. Which of the following best reflects the outcome of your research? (Select all that apply)

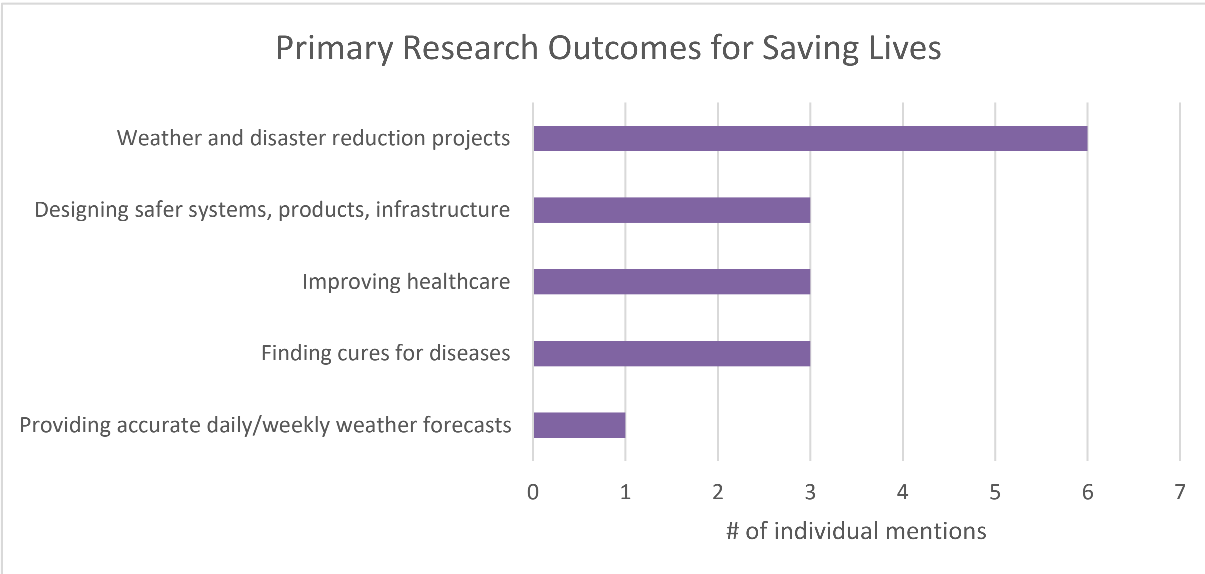

Source: Hyperion Research, 2024

National Security (non-defense)

Non-defense related outcomes for the national security contribution area include:

- Creating a framework that can be used without external help (e.g., in a classified mode)
- Additional metrics could focus on improvements in cyber security, making AI safer, tracking bad players, etc.
- Energy security and independence
- Food security

Energy security and independence was the top outcome for national security (non-defense).

FIGURE 11

National Security (non-defense) Outcomes

Q. Which of the following best reflects the outcome of your research? (Select all that apply)

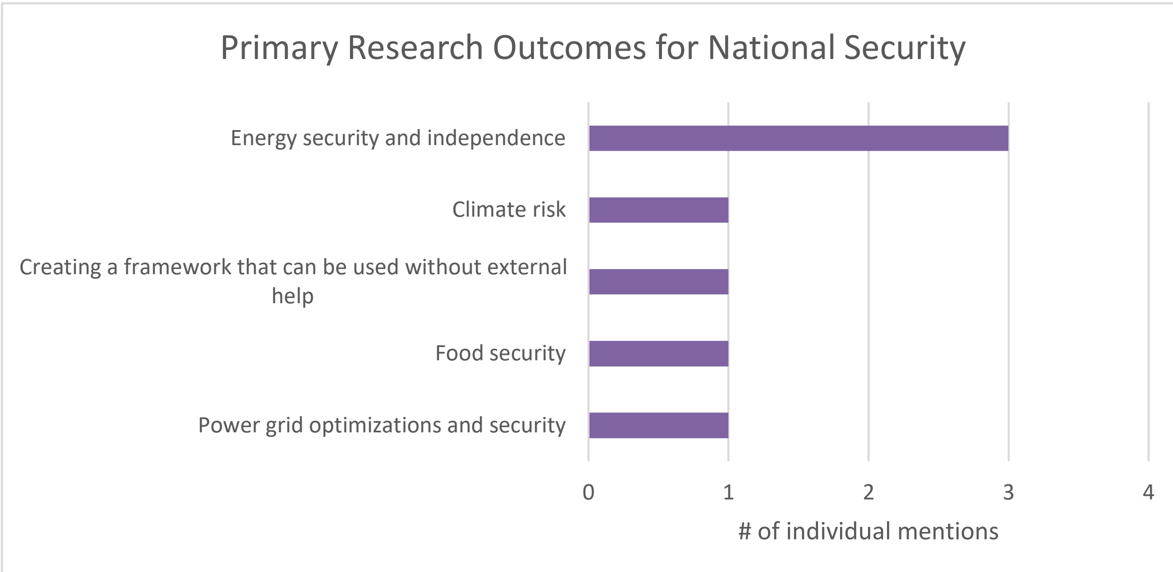

Source: Hyperion Research, 2024

NERSC Utilization

The utilization of NERSC resources is quantified through a metric known as a "NERSC hour," which serves as a normalized measure of computational usage. This metric facilitates comparisons across different time periods and architectural configurations, allowing for a consistent evaluation of resource consumption. For instance, if a project utilizes 1/365 of the available hours at NERSC, it earns credit for 24 NERSC Center Hours, equating to one day of usage. This standardized approach not only provides clarity in measuring usage but also enables stakeholders to assess the efficiency and impact of various projects within the NERSC ecosystem.

In a current snapshot of NERSC utilization, the center supports over 10,000 users and more than 1,000 distinct projects. The 42 projects evaluated in this study are projected to consume nearly 20% of the total NERSC hours, indicating that these are some of the most computationally demanding projects supported by the NERSC facility. The projects in this study span a range of durations, from 3 to 24 years, indicating that many are long-term endeavors contributing to ongoing research efforts.

Among the projects analyzed, the top three stand out in terms of NERSC hour consumption. The Advanced Simulation of Pore Scale Flow and Transport Processes in Nanoporous Materials, led by David Trebotich, accumulated an impressive 8,945 NERSC hours over 24 years. Similarly, the Computational Studies in Plasma Physics and Fusion Energy, headed by Nathan Howard, utilized 2,729 NERSC hours over 23 years, while the Quarkonia in Hot Medium project, led by Peter Petreczky, consumed 1,363 NERSC hours within a 4-year span.

The figure below shows the computational usage of all projects in this study that have been calculated using the NERSC hour metric. Due to the wide range of NERSC hour requirements, the chart below is represented on a logarithmic scale.

FIGURE 12

NERSC Hours for Research Projects

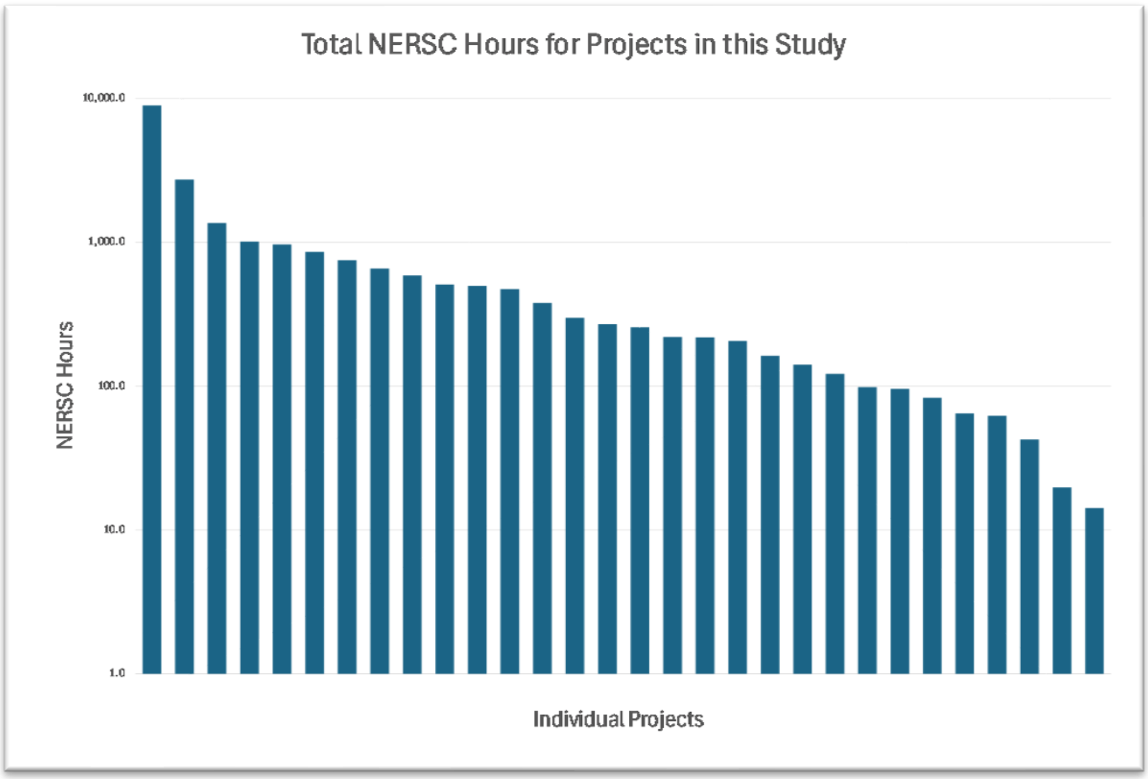

Source: Hyperion Research, 2024

Themes from Scientist and Researcher Feedback Regarding NERSC

Study participants were asked to share any specific thoughts they had about working with NERSC from any perspective. Based on unanimous, consistent feedback, the conclusion was that NERSC operations are the standard by which all other scientific computing facilities should aspire to. NERSC performs at a top leadership level across all dimensions valued by global leading scientists and researchers. Several key themes emerged from the feedback:

- Continued priority and excellence in computational capabilities
- Quality, reliability, and uptime of the site's resources
- Straightforward process by which to obtain time on its machines
- Superior support whenever questions or issues arise

Excellence in computational capabilities

Researchers commented that the cadence at which NERSC is able to deliver succeeding generations of performance continuously accelerates the progress of their research and delivery of discoveries.

Reliability of NERSC resources

Beyond the computational capabilities of the infrastructure, the researchers greatly appreciated the uptime of NERSC resources. They commented it was the rare occasion that resources went down and interrupted or delayed their research activities.

Straightforward process to obtain time on the machines

Researchers were beyond grateful at the ease and simplicity with which to apply for time on NERSC resources. They shared that the grant process was clear with minimal excess overhead.

Superior support

On the rare occasion issues did arise, researchers commented that NERSC support personnel were readily available, prompt with responses, and diligent in working until the issue was resolved.

CONCLUSION AND FUTURE OUTLOOK

The value and impact of the discoveries uncovered as a result of research done utilizing NERSC resources cannot be overstated. For 50 years, researchers and scientists from academic institutions and research organizations from across the country have relied upon NERSC to attain goals requiring ever increasing computational capabilities to conduct their research on almost every area of science. NERSC is clearly in the top echelon of leadership supercomputer sites in the world.

Many of the scientists who participated in the study have been conducting their research at NERSC over a long period of time, in some cases over 20 years. This is a testament to the computational capabilities of the steady cadence and evolution of machines over the years as well as to the quality, consistency and responsiveness of the personnel responsible for operations at NERSC.

Scientific discovery is never complete. New insights and discoveries are built upon what has been achieved by prior scientists, immortalized by Isaac Newton's statement, "If I have seen further, it is by standing on the shoulders of giants." As computational capabilities continue to advance, so too does the inquisitiveness of researchers imagining what they can do with it. Should NERSC continue to execute its mission, it is well-positioned to support current and future researchers in their quest to understand the nature of the universe and address society's grand challenges for another 50 years.

APPENDICES

Appendix 1 Methodology

A 3-phase approach was employed for the study.

Phase 1: Adjust the Hyperion Research framework that measures the value of scientific computing research to fit NERSC's unique situation. This involved selecting metrics that best fit the research conducted at NERSC.

Phase 2: Measure the historic value by surveying individual projects and accomplishments. This included conducting 42 interviews with scientists and researchers who performed their work on NERSC machines.

- NERSC provided researchers' contact information, research topics.
- NERSC introduced Hyperion Research and requested support for the project.
- Hyperion Research invited researchers to participate and schedule 30-minute phone conversations.

Phase 3: Analyze the results and compare with other HPC sites.

Appendix 2 Study Participants

Table 1

Study Participants

PI	Organization	Research Title
Hussein Aluie	University of Rochester	Computational Studies Across Time and Length Scales of Multifunctional Ionic Polymers Membranes
Alexander Austregesilo	Jefferson Lab	Analysis and Simulation for the GlueX Detector: studying the strong nuclear force and searching for exotic configurations of quarks and gluons
Michal Bajdich	SLAC National Accelerator Laboratory	SUNCAT-FWP CATALYSIS PROJECT: focusing on CO2 reduction, nitrogen reduction, and water splitting
Edward Baron	University of Oklahoma	Synthetic Spectra of Astrophysical Objects Radiative Transfer in Supernovae, r-process, and stellar evolution
Adam Burrows	Princeton University	Three-Dimensional Simulations of Core-Collapse Supernovae: to understand how supernova explosions occur, resulting in the formation of neutron stars or black holes.
Paolo Calafiura	Lawrence Berkeley National Laboratory	Detector Simulation of the CERN ATLAS Detector on NERSC HPCs
Roberto Car	Princeton University	Chemistry in Solution and at Interfaces: to fit the potential energy of interaction among atoms.
Wai-Yim Ching	University of Missouri - Kansas City	Structure refinement of Spike-protein of SARS-COV-2: This approach allows them to analyze the entire complex structure of the virus, rather than breaking it down into smaller separate pieces.
Norman Christ	Columbia Univsersity	High Energy Particle Physics - Symmetry of Nature. The equivalence of particles and antiparticles. Right handed vs. left handed
Carleton DeTar	University of Utah	Flavor physics: the focus is on understanding physics beyond the Standard Model
Kjiersten Fagnan	DOE Joint Genome Institute (JGI)	Joint Genome Institute - Production Sequencing and Genomics: researching areas such as nutrient cycling, biofuels, microbial communities, and plant genomics.
Giulia Galli	University of Chicago	First Principles Simulations of Nanostructures: to study matter at the atomic level
Vikram Gavini	University of Michigan	PRISMS: Integrated multiscale modeling of Mg structural alloys: The goal is to understand and improve the mechanical properties of magnesium alloys

Table 1

Study Participants

PI	Organization	Research Title
Vassiliki Glezakou	Oak Ridge National Laboratory	Carbon management and energy efficiency through multifunctional catalysis
Rajan Gupta	Los Alamos National Laboratory	Lattice QCD search for physics beyond the standard model High Precision Calculations of the Nucleon Structure for Fundamental Symmetries: to create a better understanding the properties of neutrons and protons.
Forrest Hoffman	Oak Ridge National Laboratory	Reducing Uncertainties in Biogeochemical Interactions through Synthesis and Computation (RUBISCO): to study the interactions between biogeochemistry and the climate system.
Dr. Nathan Howard	Massachusetts Institute of Technology	Computational Studies in Plasma Physics and Fusion Energy: Multi-scale Turbulence in Tokamak Reactors: to reduce heat and particle losses in fusion plasmas.
Stephen Jardin	Princeton Plasma Physics Laboratory	3D Extended MHD simulation of fusion plasmas: for tokamak fusion reactors.
Puru Jena	Virginia Commonwealth University	Cluster and Nanostructure for Energy and Bio Applications: for improved energy production and storage.
Lai-Yung Ruby Leung	Pacific Northwest National Laboratory	Energy Exascale Earth System Modeling (E3SM): improving the understanding of how energy use contributes to climate change and how climate change may affect future energy use and infrastructure. Water Cycle and Climate Extremes Modeling (WACCEM): to understand how various aspects of the water cycle, particularly precipitation, have changed in the past and how they might change in the future.
Manos Mavrikakis	University of Wisconsin - Madison	First-Principles Catalyst Design for Environmentally Benign Energy Production: to uncover why certain materials are effective catalysts while others are not.
David Mccallen	Lawrence Berkeley National Laboratory	High Performance Simulations for Regional Scale Earthquake Hazard and Risk Assessments
Maria Elena Monzani	SLAC National Accelerator Laboratory	LZ - LUX ZEPLIN experiment, which aims to detect dark matter.
Gottlieb Ore Alexander (Sasha) Tchekhovskoy	Northwestern University	Effects of remnant neutron star lifetime on compact object merger outflows, nucleosynthesis, and emission: The goal is to connect the merger event itself with the observable electromagnetic signatures, allowing scientists to extract valuable physics information from these observations.

Table 1

Study Participants

PI	Organization	Research Title
Dvora Perahia	Clemson University	Computational Studies Across Time and Length Scales of Multifunctional Ionic Polymers Membranes: The focus is on polymers with ionic groups that can conduct ions, generate electricity, and trap nanoparticles.
Peter Petreczky	Brookhaven National Laboratory	Quarkonia in Hot Medium: the properties of matter under extreme conditions.
Jin Qian	Lawrence Berkeley National Laboratory	From Molecules to Continuum: Developing a Universal Approach for Accurate Description of X-ray Photo: for nanoscale phenomena.
Diana Qiu	Yale University	Electronic and Optical Properties of Layered Materials for Energy Applications: to develop new theoretical and computational tools for understanding how materials interact with light.
David Radice	The Pennsylvania State University	Nuclear Astrophysics with Numerical Relativity: the collision between neutron stars
Robert Ryne	Lawrence Berkeley National Laboratory	Frontiers in Accelerator Design: Advanced Modeling for Next-Generation BES Accelerators: the design and development of particle accelerators.
Masao Sako	University of Pennsylvania	The Dark Energy Survey Supernova Search: The project aimed to measure the properties of around 300 million galaxies and discover thousands of new supernova explosions.
Subramanian Sankaranarayanan	Argonne National Laboratory	Development of New Force Fields using Machine Learning and First Principles Physics
Sahar Sharifzadeh	Boston University	Large-Scale Many-Body Perturbation Theory Simulations of Optoelectronic Materials: to understand the electronic properties of materials at the atomic scale.
Sterling Smith	General Atomics	Computing in support of the DIII-D National Fusion Facility: the goal is to study how to confine and heat plasma efficiently enough to achieve fusion
David Trebotich	Lawrence Berkeley National Laboratory	Advanced Simulation of Pore Scale Flow and Transport Processes in Nanoporous Materials: examining how fluids interact with materials at a microscopic level (around 1-100 microns).
Paul Ullrich	University of California - Davis	A Framework for Improving the Analysis and Modeling of Earth System and Intersectoral Dynamics at Regional Scales
Josh Vermaas	Michigan State University	Large Scale Molecular Simulations to Support Photosynthesis and Carbon Fixation

Table 1

Study Participants

PI	Organization	Research Title
Hailong Wang	Pacific Northwest National Laboratory	High-Latitude Application and Testing (HiLAT) of Earth System Models & Regional Arctic System Model
Xifan Wu	Temple University	Theoretical spectra calculations of liquid water and ion solutions

Source: Hyperion Research, 2024

Appendix 3 Research Synopses of NERSC Projects

The following are brief synopses of the 42 NERSC research projects investigated during this study:

1. Computer simulations of plasma turbulence in fusion reactors, revealing how large- and small-scale turbulence interact, leading to increased electron heat loss and providing insights into long-standing experimental observations.

2. High-resolution simulations of fluid interactions in nanoporous materials, utilizing advanced algorithms to model processes relevant to geological systems and engineered materials, with significant applications in CO2 sequestration.

3. Investigating properties of quark-gluon plasma under extreme conditions akin to those shortly after the Big Bang, using numerical simulations to explore heavy quark diffusion and the potential energy between quarks in this unique state of matter.

4. Enhancing the design of particle accelerators used as light sources, improving beam quality and compactness while developing innovative methods for particle manipulation to support advancements in accelerator technology.

5. Compact object mergers research seeking to understand their implications for nuclear physics, using advanced simulations to interpret gravitational wave data and predict future observations from upcoming detection instruments.

6. Studying catalysis targeting CO2 reduction and water splitting, developing an iridium-based catalyst that excels in performance and collaborates with industry to advance hydrogen production technologies.

7. Exploring the strong nuclear force and exotic quark configurations, utilizing high-performance computing resources to analyze vast datasets generated by accelerator experiments.

8. Conducting quantum simulations to study matter at the atomic level, collaborating closely with experimentalists to develop materials for sustainable energy and quantum technologies through interdisciplinary research.

9. Studying catalysis at the atomic level, employing quantum mechanics and simulations to improve catalytic materials for hydrogen production and fuel cells, bridging fundamental science with practical applications.

10. Understanding the water cycle and climate extremes using advanced modeling techniques, leveraging high-performance computing resources to simulate historical and future precipitation patterns and analyze extreme weather events.

11. A large-scale initiative involving multiple national laboratories aims to develop the Energy Exascale Earth System Model (E3SM) to enhance understanding of how energy use impacts climate change and vice versa, utilizing exascale computing capabilities for high-resolution modeling.

12. Studying fundamental symmetries in particle physics, successfully predicting a subtle asymmetry known as direct CP violation within the Standard Model, utilizing advanced computational resources to achieve unprecedented precision in theoretical predictions.

13. The Earthquake Simulation framework (EQ SIM), designed to model earthquake scenarios in the San Francisco Bay Area, using high-performance computing to simulate ground motion propagation and its interaction with infrastructure for various fault scenarios.

14. Developing new force fields for molecular dynamics simulations through machine learning, significantly reducing model development time and enabling the simulation of larger systems, including recently discovered 2D materials.

15. Research on core collapse supernovae aims to understand the explosive deaths of massive stars by creating complex simulations that model the turbulent processes involved, requiring substantial computational resources due to the multifaceted physics at play.

16. Computational chemistry integrating machine learning with traditional methods to investigate reactivity in complex systems, developing a global optimization code that improves the handling of large molecular systems and enhances prediction accuracy for various chemical reactions.

17. Using machine learning and AI to accelerate molecular dynamics simulations based on quantum mechanics, introducing a method that allows for larger system simulations while maintaining accuracy, and further expanding this work to mesoscale simulations.

18. The LUX-ZEPLIN experiment seeks to detect dark matter using a large liquid xenon detector, aiming to identify Weakly Interacting Massive Particles (WIMPs) through advanced detection techniques in one of the most sensitive searches for dark matter to date.

19. Investigating molecular structure of water, aiming to confirm the tetrahedral model against alternative theories by employing advanced quantum mechanics calculations and supercomputing resources to replicate experimental findings.

20. Theoretical physics focusing on materials for energy applications, exploring novel catalysts and alternative battery technologies through high-performance computing to improve energy production, storage, and hydrogen generation from water.

21. Simulating the response of high-energy physics detectors, specifically for the ATLAS experiment at CERN, aiming to "unfold" detector effects to accurately analyze particle behavior, requiring an estimated 10 billion CPU hours annually.

22. Developing and improving magnetohydrodynamic (MHD) simulation codes for tokamak fusion reactors, utilizing a 3D code to study plasma instabilities and predict operational limits, enhancing the design of safe and efficient fusion reactors.

23. Models the spectra of Type 1A supernovae to understand their origins and the stellar systems from which they emerge, using spectroscopy to analyze the characteristics of these thermonuclear explosions.

24. Investigating the production of r-process elements like gold and platinum in the universe, exploring whether these elements originate from core-collapse supernovae, kilonovae, or a combination of both.

25. The PRISMS Project develops multi-scale modeling methods for magnesium alloys to enhance their mechanical properties, utilizing density functional theory (DFT) calculations to study defects and their macroscopic effects.

26. Research on "flavor physics" using quantum chromodynamics (QCD) to explore phenomena beyond the Standard Model, conducting precise calculations to predict rare decay processes and the magnetic moment of the muon.

27. Fusion studies at General Atomics using the D3D tokamak, focusing on complex simulations of plasma turbulence and heat transport, leveraging high-performance computing resources for real-time analysis of experimental data.

28. The RUBISCO project conducts simulations using existing Earth system models to investigate the interactions between biogeochemistry and climate change, including the effects of human activities on the carbon cycle.

29. Generating high-quality DNA sequence data for various organisms and has partnered with a computing center to enhance data production and analysis capabilities, facilitating new research initiatives.

30. Project Hyperfacets collaborates with stakeholders to identify climate data needs, using high-resolution simulations to investigate historical events and their implications for future climate conditions, primarily serving water managers and electrical system operators.

31. Developing tools to predict X-ray Photo Electron Spectroscopy (XPS) signatures for nanoscale phenomena using a specialized form of Density Functional Theory (DFT) that can simulate systems with tens of thousands of atoms, addressing challenges in chemistry and materials science.

32. Creating theoretical and computational tools to understand material interactions with light, particularly in novel two-dimensional materials, utilizing large-scale calculations to study complex electron-hole interactions and high harmonic generation.

33. Simulating the merger of binary neutron star systems, exploring post-merger evolution and how these events contribute to the synthesis of heavy elements, aiming to connect merger events with observable electromagnetic signatures.

34. Investigating the properties of neutrons and protons using quantum chromodynamics (QCD), employing numerical simulations to explore the electric dipole moment of the neutron and potential violations of charge conjugation and parity (CP) symmetry.

35. Understanding nucleon properties through Quantum Chromodynamics (QCD) simulations, specifically investigating neutron and proton spin and electric charge distribution. A key goal is to search for an electric dipole moment in neutrons, which could indicate a violation of charge conjugation and parity (CP) symmetry, suggesting new physics beyond the Standard Model.

36. Studying photosynthesis by simulating synthetic carboxysomes, protein structures that enhance carbon fixation efficiency, and exploring the engineering of these shells to improve metabolic processes for applications like fertilizer production.

37. The HILAT project investigates Arctic amplification by utilizing high-resolution climate models to simulate historical climate change and future scenarios, with a focus on understanding tipping points in the Arctic climate system.

38. Structural refinement of the SARS-CoV-2 spike protein using computational modeling to analyze its complex structure, ensuring accuracy in atomic bonding and properties.

39. The Dark Energy Survey supernova research project aimed to measure properties of galaxies and discover new supernovae, particularly Type 1A supernovae, to better understand dark energy and the universe's expansion history.

40. Designing lightweight multifunctional materials, particularly polymers with ionic groups, for applications in clean energy and biotechnology, using neutron sources and computational modeling to bridge atomic-scale understanding with practical applications.

41. Studying turbulent flows and instabilities in high energy density systems focuses on modeling turbulence in contexts like supernovae and nuclear fusion, identifying biases in 2D models and exploring magnetized flows to improve understanding of astrophysical phenomena.

42. Investigating the electronic properties of materials at the atomic scale, focusing on how they interact with light through quantum mechanical equations. Utilizing supercomputing resources, the project explores phenomena like optical excitations and electron density redistribution, with applications in quantum communication and carbon nanotube technology.

Appendix 4 Participant Quotes

Table 2

Participant Quotes

Quote	Attribution
The insights we gained into multi-scale turbulence in fusion plasmas would not have been possible without the use of NERSC resources.	Nathan Howard, a Fusion Energy researcher from MIT Plasma Science and Fusion Center
NERSC is an extremely valuable resource for researchers supported by the DOE because it's accessible for any researcher. Continuous access to HPC resources is critical and they provide extremely good support. I really like to work with NERSC and without them, the research would not have been possible.	A researcher from Brookhaven National Laboratory
NERSC is best partner for energy science. I'm very happy with the NERSC allocation process and continued support over the years.	A researcher from SLAC National Accelerator Lab at Stanford University
The application process is nicely streamlined and pretty straightforward. We complement NERSC for making a system that is not a huge burden on users."	Alexander Austregesilo, a researcher from Thomas Jefferson National Accelerator Facility
Our experience with NERSC has been very positive and the impact of NERSC on our scientific research has been critical; without NERSC, our publications would not have been possible.	A researcher from University of Chicago
We appreciate what a super-efficient organization NERSC is for supporting world-class research. It's amazing to see how such a complex enterprise can facilitate world-class research. It's a real jewel for the computational sciences and engineering community.	Manos Mavrikakis, a Chemical Engineering researcher
NERSC will remain a really important resource for us.	Ruby Leung, an Earth Systems researcher from PNNL
NERSC provides very powerful and accessible capabilities for research, successful support for using the machines, and is a very effective place to work."	A researcher from Columbia University
Perlmutter and NERSC have been a tremendous asset to our project both from the standpoint of code development on GPU systems as well as running large regional cases of ground motion and infrastructure response.	David McCallen a researcher from the Energy Geosciences Division at Berkeley Laboratory
I cut my teeth in HPC on NERSC allocations that in retrospect helped to launch me as a computational astrophysicist.	Adam Burrows a researcher from Princeton University
I am beyond grateful - we honestly could not have done this research without NERSC.	A researcher from Oak Ridge National Laboratory (ORNL)
I am extremely grateful for the support provided by NERSC. Much of what we have done would not have been possible.	Puru Jena, a Chemical Physics researcher from Virginia Commonwealth University

Table 2

Participant Quotes

Quote	Attribution
What makes NERSC different from other global HPCs, it's not run just as a research facility but is the easiest to use. Documentation is excellent and support is good. It's HPC being run with the researchers in mind.	A researcher from Lawrence Berkeley National Laboratory (LBNL)
NERSC has been a stable computing platform for my entire career and has been phenomenal in terms of consistency and increasing capabilities.	Stephen Jardin, a Fusion researcher from the Princeton Plasma Physics Laboratory
I cannot overstate how important NERSC has been to my career and my research. From establishing the research, performing the work, and providing training, NERSC has been invaluable to my career. It's such a well-run organization and a pleasure to work with.	Eddie Baron, an Astrophysics researcher from the University of Oklahoma
NERSC by far, of all the HPC resources used in conjunction with this research, has the best infrastructure and the most responsive staff. The research couldn't have been done without NERSC.	Diana Qiu, a physics researcher from Yale University
By far NERSC is the best HPC resource I've worked with for my research in terms of on-line documentation, performance of the computing systems and the unbelievable staff who provide help whenever I need it - and they do it quickly. I've never had to wait more than half a day to get some problem resolved.at NERSC.	Masao Sako, a cosmology physics researcher from University of Pennsylvania
NERSC has constantly provided incredibly useful resources and support to my research in climate modeling and analysis for the past 15 years	Hailong Wang, an Earth Scientist from Pacific Northwest National Laboratory (PNNL)
I had a great experience with nurse help and nurse has enabled research that I couldn't do anywhere else. So I'm a big fan of NERSC.	Anonymous

Source: Hyperion Research, 2024

Appendix 5 Study Participant NERSC Demographic Profile

Table 3

NERSC Office

Office	#
Basic Energy Sciences	16
High Energy Physics	7
Nuclear Physics	7
Biological and Environmental Research	6
Fusion Energy Sciences	4
NERSC Directors Reserve	2

Source: Hyperion Research, 2024

Table 4

NERSC Program

Program	#
Chemical Sciences, Geosciences, & Biosciences (CSGB)	9
Nuclear Physics	7
High Energy Physics	6
Materials Sciences and Engineering	6
Earth and Environmental Systems Sciences Division	5
Fusion Energy Sciences	4
Biological Systems Science	1
DDR Campaign	1
DDR Scale	1
Scientific User Facilities	1

Source: Hyperion Research, 2024

Table 5

NERSC Science Category

Category	#
Physics: Astrophysics	4
Chemistry: Catalysis	3
Chemistry: Chemical Physics	3
Earth Systems: Coupled Systems	3
Energy: Fusion	3
Physics: High Energy Physics (Theory)	3

Table 5

NERSC Science Category

Category	#
Physics: Nuclear Physics (Theory)	3
Earth and Environmental Systems	2
Materials Science: General	2
Physics: Condensed Matter	2
Physics: High Energy Physics (Experimental)	2
Biosciences: Genomics	1
Biosciences: Molecular Science	1
Chemistry: Physical Chemistry	1
Energy: General	1
Geoscience: Geochemistry	1
Materials Science: Biomolecular Materials	1
Materials Science: Energy Storage	1
Materials Science: Soft Matter	1
Physics: Accelerator Science	1
Physics: Cosmology	1
Physics: Nuclear Physics (Experimental)	1
Physics: Plasma Physics	1

Source: Hyperion Research, 2024

Appendix 6 Survey Guide

1. Briefly provide an overview description of your research for a novice in the field. (open-ended)
2. Please describe in novice terms how the accomplishment impacted the world and why it's important to the world or your organization.
3. Which of the following areas best describes the primary scientific contribution of your research (please select all that apply and provide details):
   a) Advances Science (if selected, which of the following best reflects the outcome of your research; please select all that apply)
      i. Provides a better understanding of how something works
      ii. Accelerates the time-to-solution of a critical large problem
      iii. Tests otherwise untestable theories, either due to physical limitations or undesirable, potentially adverse human impact
      iv. Reduces climate change or impacts, or improves the health of our planet
      v. Improves energy efficiency or creates more sustainable products and approaches
      vi. Increases the productivity, efficiency, and success of scientists and researchers
      vii. Other (please specify)
   b) Advances Engineering (if selected, which of the following best reflects the outcome of your research; please select all that apply)
      i. Creates better products and supports developing new types of products
      ii. Designing better physical devices (e.g., Designing better computers and making them more useable (e.g., GPUs)
      iii. Accelerates the time from discovery to building products
      iv. Testing otherwise untestable theories, either due to physical limitations or undesirable (potentially adverse) human impact
      v. Technology innovation transfer to industry, (e.g., speeding up the time from discovery to products, improved understanding of how to build better products and services, developing new types of products, supporting industrial partnerships)
      vi. Engineering better computing technologies (hardware, software, AI, etc.) for first of the kind technologies, (e.g. designing better physical devices, designing better computers and making them more useable (e.g., GPUs), creation of better software, data libraries/repositories and computing tools for broader use)
      vii. Improving the design and safety of power plants and developing technologies to exploit alternative energy sources

viii. Other (please specify)

c) Economic Growth (if selected, which of the following best reflects the outcome of your research; please select all that apply)

   i. Helping to increase jobs and grow the nation's GDP

   ii. Technology innovation transfer to industry, (e.g., speeding up the time from discovery to products, improved understanding of how to build better products and services, developing new types of products, supporting industrial partnerships)

   iii. Avoiding costs/saving money (e.g., simulation vs. high-cost testing can save substantial funds and have a faster time-to-solution)

   iv. Building a more diverse and stronger scientific workforce

   v. Leadership computing sites can be a brain magnet for attracting talent and can help motivate people to go into science

d) Saving Lives (if selected, which of the following best reflects the outcome of your research; please select only one)

   i. Weather and disaster reduction projects, finding cures for diseases

   ii. Providing accurate daily/weekly weather forecasts needed by the transportation, agricultural, and tourism industries

   iii. Improving healthcare, e.g.,

      i. Finding ways to take advantage of the revolutionary advances in biotechnology

      ii. Addressing major diseases and widespread pandemics (e.g., cures, vaccines) and reducing overall impacts

   iv. Designing safer systems, products, infrastructure, etc.

e) National Security (non-defense) (if selected, which of the following best reflects the outcome of your research; please select all that apply)

   i. Creating a framework that can be used without external help (e.g., in a classified mode)

   ii. Additional metrics could focus on improvements in cyber security, making AI safer, tracking bad players, etc.

4. Over what period of time did you conduct your research?

5. What organization or institution were you with when you conducted your research? Please provide the organization's name and location.

6. Was the research performed solely with NERSC or were other organizations or partners involved? If "others", who else participated in and supported the research?

7. How many CPU hours were used to accomplish your research success? Please only account for the time on this specific project.

8. How many papers were published as a result of this research?

9. How many citations referenced these publications?

10. Would you allow any of the response to the prior question to be utilized as part of the NERSC 50th anniversary celebration in any kind of promotional material (e.g., press release, presentations, website descriptions, social media, marketing collateral, posters)?

    a) Yes (open ended)

    b) No

11. Did this accomplishment generate a financial ROI, an Innovation, or both?

    a) ___ Financial ROI (Please complete all questions)

    b) ___ Innovation only (Jump to Q18)

    c) ___ Both (Please complete all questions)

FOR PROJECTS THAT CREATED AN INNOVATION:

12. What was the general category of the innovation? (Please select just one)

    a) ___ Basic Research, including major discoveries and pioneering breakthroughs

    b) ___ Applied Research, including incremental innovations and process improvements

13. How would you rate the IMPORTANCE of you research in this field over the last ten years, using a scale of 1 to 5? (Please select just one, with 5 being the highest)

    a) ____5 One of the top 3 discoveries in your field in the last decade

    b) ____ 4 One of the top 5 discoveries in your field in the last decade

    c) ____ 3 One of the top 10 discoveries in your field in the last decade

    d) ____ 2 One of the top 25 discoveries in your field in the last decade

    e) ____ 1 One of the top 5 discoveries in your field in the last decade

14. How would you rate the IMPACT of your research to multiple organizations, using a scale of 1 to 6? (Please select just one, with 6 being the highest)

    a) ____ 6 It is useful to over 50 organizations

    b) ____ 5 It is useful to 10 to 49 organizations

c) ____ 4 It is useful to 6 to 9 organizations

d) ____ 3 It is useful to 2 to 5 organizations

e) ____ 2 It is only useful to 1 organization

f) ____ 1 It is recognized ONLY by experts in the field

15. Would you be willing to provide a direct quote to be utilized as part of the NERSC 50th anniversary celebration in any kind of promotional material (e.g., press release, presentations, website descriptions, social media, marketing collateral, posters)? The quote could relate to any one or more of the following:

    - Your experience working with NERSC
    - How using the NERSC facility helped your research to be successful
    - How your research could only have been accomplished using the NERSC facility

    a) Yes (open ended) Please provide up to 4 sentences for your quote.

    b) No

16. May we follow up with you if we have any further questions regarding your research as we progress with the study?

    a) Yes

    b) No

About Hyperion Research, LLC

Hyperion Research provides data-driven research, analysis and recommendations for technologies, applications, and markets in high performance computing and emerging technology areas to help organizations worldwide make effective decisions and seize growth opportunities. Research includes market sizing and forecasting, share tracking, segmentation, technology and related trend analysis, and both user & vendor analysis for multi-user technical server technology used for HPC and HPDA (high performance data analysis). Hyperion Research provides thought leadership and practical guidance for users, vendors, and other members of the HPC community by focusing on key market and technology trends across government, industry, commerce, and academia.

Headquarters

365 Summit Avenue
St. Paul, MN 55102
USA
612.812.5798
www.HyperionResearch.com and www.hpcuserforum.com
