2022–2023 Highlights

ACCESS

Innovating Access to Our Nation's Cyberinfrastructure

The national research cyberinfrastructure (CI) ecosystem is essential to computational and data-intensive research across all of 21st-century science and engineering (S&E). It's driven by rapid advances in a wide range of technologies, increasing volumes of highly heterogeneous data and escalating demand by the research community.

Research CI is a key catalyst for discovery and innovation. It plays a critical role in ensuring U.S. leadership in S&E, economic competitiveness and national security, consistent with the National Science Foundation's mission. The NSF, through the Office of Advanced Cyberinfrastructure (OAC), has published a vision that calls for the broad availability and innovative use of an agile, integrated, robust, trustworthy and sustainable CI ecosystem that can drive new thinking and transformative discoveries in all areas of S&E research and education.

The NSF's ACCESS (Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support) program builds upon the successes of the 11-year XSEDE project while also expanding the ecosystem with capabilities for new modes of research and further democratizing participation.

The program encompasses five service teams working in unison toward the goals of the ACCESS program:

- Allocations
- Operations
- Coordination
- Support
- Metrics

By tasking each team with a specific focus, ACCESS endeavors to bring greater innovation to the democratization of our nation's cyberinfrastructure. Opening the door wider to welcome more researchers solving the grand challenges that drive discovery and progress.

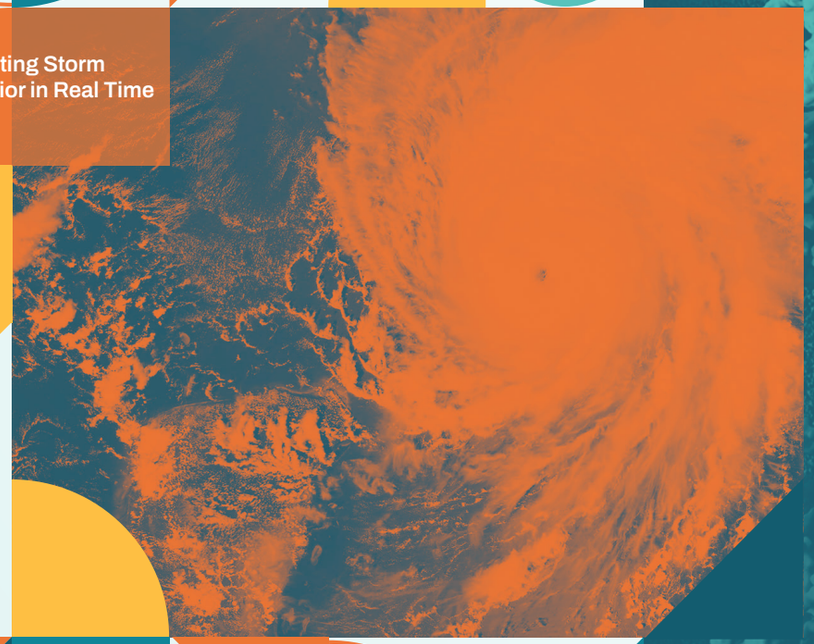

Predicting Storm Behavior in Real Time
— 11

The Ways of Water
— 14

The Best-dressed Dinner Guest
— 17

A Pore Solution
— 15

The Search for Better Battery Life
— 13

Contents

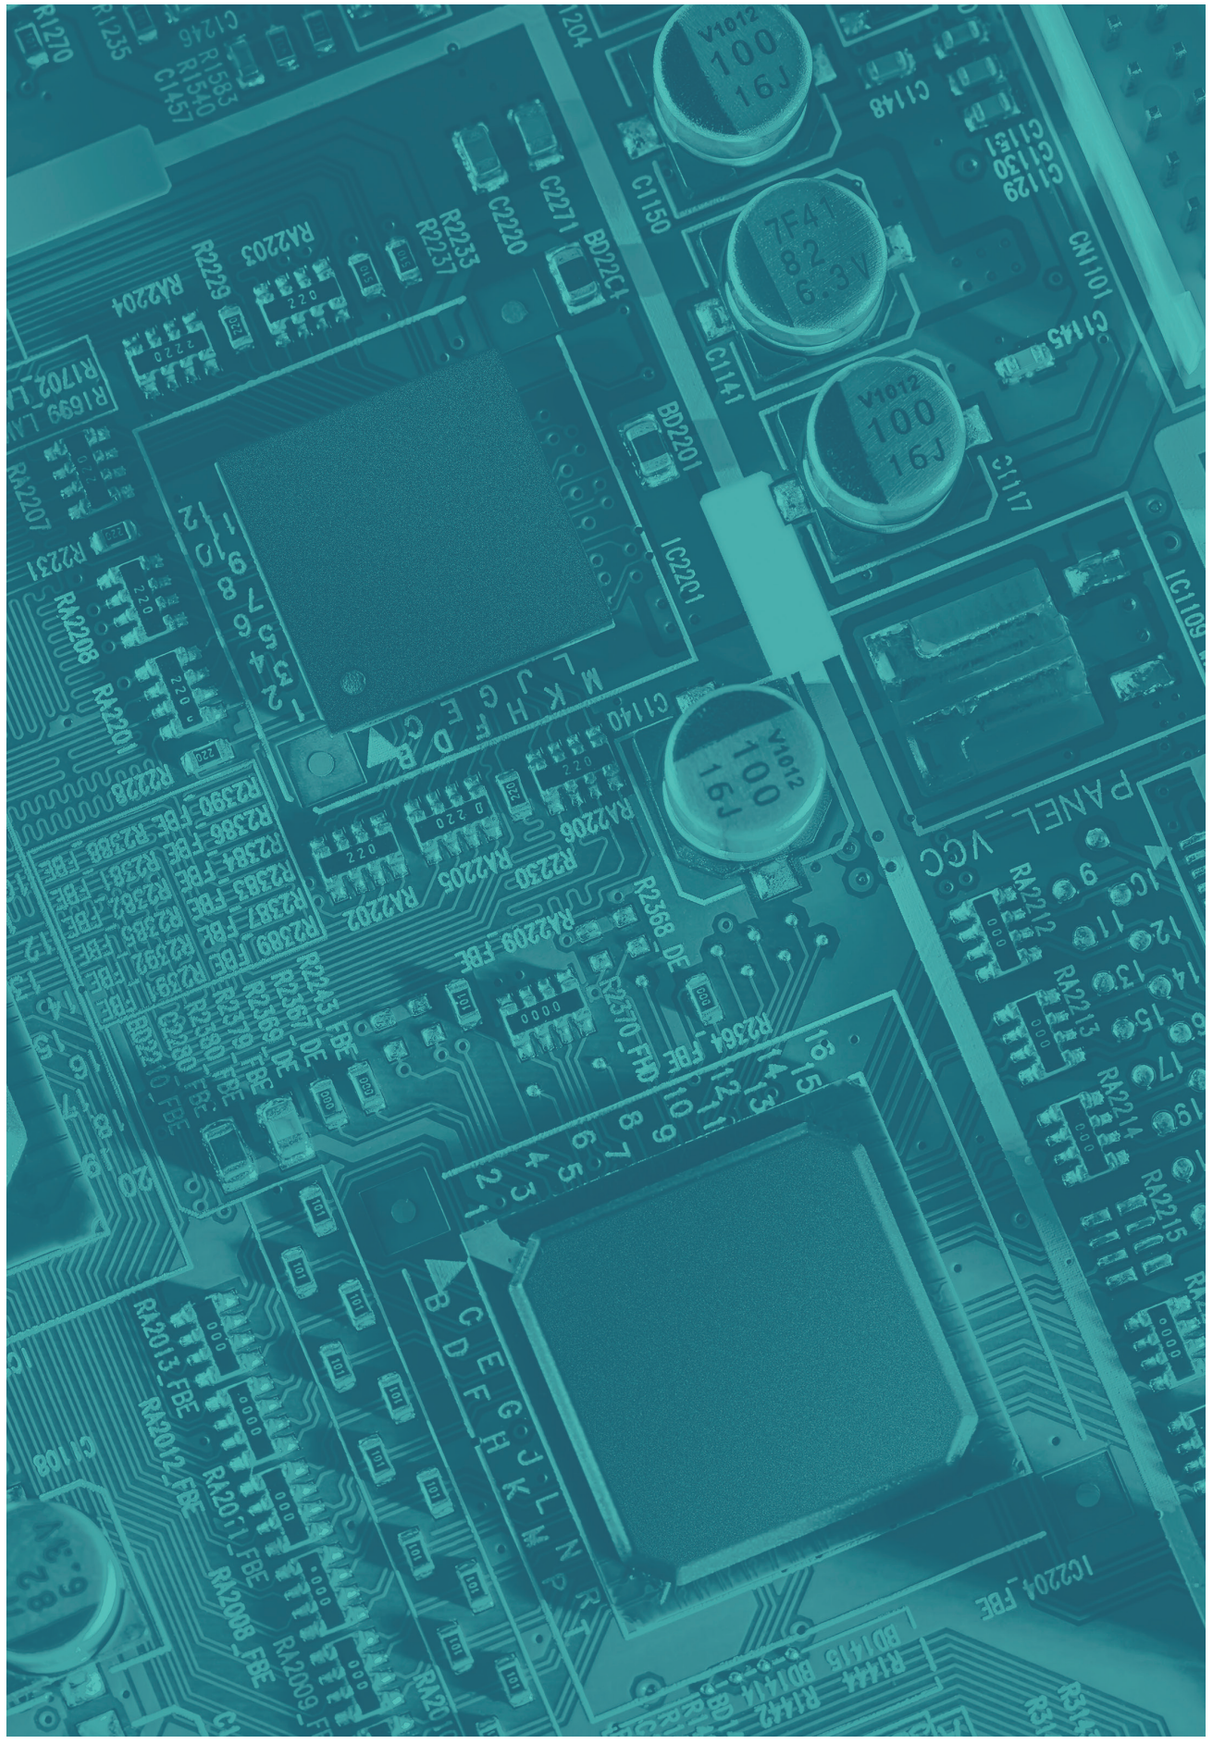

Welcome Letters from

ACCESS PIs

Stephen Deems
Principal Investigator | ACCESS Allocations

To all the community members who have played a pivotal role in the success of our inaugural year of operations, we send our sincere thanks and appreciation for your continued contributions.

We owe the research and educational communities a tremendous amount of thanks for their unwavering patience and understanding during our transitional period. Your feedback on how to improve our offerings and lower barriers to entry has been invaluable in propelling this program forward. The organic evangelization of our services that we've witnessed from our community members has been truly humbling and inspiring.

Without the work of the Resource Providers, our communities would lack the necessary infrastructure and expertise to carry out their respective initiatives on national CI resources. Your participation and input have been key in improving the experience for all our stakeholders and are integral to future innovations.

We also extend our appreciation to the reviewers who have generously volunteered their time and expertise to provide guidance to researchers in pursuit of their computing and storage needs.

To our fellow ACCESS team members, the dozens of partner institutions, and our program officers – who have all been instrumental to our initial successes – we are deeply thankful for your commitment and tireless efforts. Cheers to an even brighter future and a more robust cyberinfrastructure ecosystem for all.

Shelley Knuth

Principal Investigator | ACCESS Support

My favorite part of this project is knowing that we are able to exact change in scientific discovery by making computational resources more easily accessible by a broader population. The part of my work that drives and motivates me is expanding our community, and, therefore, research discoveries by engaging with wider populations and integrating new and innovative ideas. The role that ACCESS plays in this effort is tremendously important. Our team has worked hard in the last year to ensure that our resources are discoverable, accessible and that access to them is democratized. There is a great sense of pride in what we've been able to accomplish to this point. This absolutely could not have been done without our many community collaborators – the Resource Providers, the Campus Champions, members of CaRCC and CASC, research software engineers, cyberinfrastructure professionals, outreach experts and, of course, the researchers themselves. As we embark on the next phase of this project, our focus is on discovery of the needs of the communities that we serve and that support our mission, and developing solutions to better support computational research using ACCESS resources. I am excited to talk with these communities and learn more about their challenges and to understand better how to build our ecosystem. The incorporation of new and fresh ideas, combined with storied expertise, will lead us in exciting directions that we can't yet imagine. I am excited to see what the next phase of our project holds. Thank you for all of your support.

Tom Furlani

Principal Investigator | ACCESS Metrics

Over the past year, the ACCESS team has strived to integrate five separate awards into a seamless program to support science and engineering research in the U.S. While not perfect, and with more work ahead of us, we have nonetheless successfully transitioned from XSEDE to ACCESS. I take great pride in being part of this team and want to ensure all stakeholders in the community (researchers, educational institutions, Resource Providers, advisory committee members and NSF leadership) that the entire ACCESS team is fully committed to the continual improvement of the services we provide and doing so through direct involvement with the entire stakeholder community. I look forward to the many advancements and improvements planned for the upcoming year.

Tim Boerner

Principal Investigator | ACCESS Operations

The ACCESS Operations team put forward a tremendous amount of effort during our first year in ensuring a seamless transition into operations for the infrastructure and services that provide value to the NSF cyberinfrastructure community. Beyond the initial frenzy of transition, pivoting into early operations proved exciting. We have launched new conversations with the community of Resource Providers and worked to improve our resource integration processes, migrated to a new virtual network in partnership with Internet2, and maintained cyberprotection for the ecosystem. The first cohort for the Student Training and Engagement Program had a tremendously positive experience and proved to be a highlight as we closed out the first year of the program.

As we move into our second year and look farther into the future of the program, we see numerous possibilities to advance and support the national cyberinfrastructure ecosystem. We look forward to the many new engagements and partnerships needed to realize the vision for the ACCESS program!

John Towns

Principal Investigator | ACCESS Coordination Office

As others noted, the first year of the program was intense for the whole ACCESS team. It was not until we started to move into the second year of the program that we paused to look back over the first year and see the tremendous progress we have made and the fantastic new capabilities and access options that have been put in place. Concurrently, the broad team involved in the program has also very much come together. The support of the community as we navigated the transition between programs has been tremendous and I very much look forward to the new and exciting capabilities we are planning in the coming months and years. As always, we welcome your thoughts and input as to how we can best collaborate with the national research community and look forward to the exciting days ahead.

Impacts on Science

Highlights from our first year

Understanding the Evolution of Religion

Carnegie Mellon University and Santa Fe Institute researchers use Bridges-2 to reveal how some religions persist or change while others die out.

tinyurl.com/evolution-religion

PSC: Bridges-2
Pittsburgh Supercomputing Center

Carnegie Mellon University; Santa Fe Institute

Collecting data from hundreds of religions across the globe presents some challenges. For example, more information is known about some belief systems than others. All parts of the world are not equally represented by data. The historical record includes loads of documentation about the three major religions – Christianity, Hinduism and Islam – but only fragments of text or material remain from other religions that played a role in world history. When it comes to research outcomes, this unevenness poses a risk of bias toward the known.

Predoctoral fellow Victor Møller Poulsen, working with Professor Simon DeDeo's team at Carnegie Mellon University's (CMU) Department of Social and Decision Sciences, set out to examine how human culture changes over time through the lens of cultural landscapes. This approach was rooted in a methodology begun during the 20th century when scientists studying biological evolution used mathematical tools to outline a type of landscape that explained species' evolution.

The evolutionary landscape approach described paths of natural selection as crossing into low-lying valleys and climbing to higher peaks. For example, a particular species may not develop into a more stable form because of a valley-like barrier. The mathematics behind this way of explaining evolution proved effective in understanding species and how they change.

To test this approach against data concerning worldwide religious beliefs, Poulsen, now at the Santa Fe Institute in New Mexico, turned to the Pittsburgh Supercomputing Center's (PSC) National Science Foundation-funded Bridges-2 supercomputer. The study required the application of mathematical tools that corrected for incomplete and biased data – presenting a real computational challenge. Being rigorous about what they did and did not know meant exploring a vast number of different possibilities in the computer, with each unknown increasing the number of calculations. The powerful nodes of Bridges-2 enabled this computational exercise.

> **"There's a trade-off here. By being rigorous as to how we treat missing data, by being rigorous about how we handle finite data, we get an exponentially increasing computational demand…if we simulate histories of religions evolving, we can see if our algorithm can infer the known data correctly. Bridges-2 gave us this ability.**
>
> Simon DeDeo, Carnegie Mellon University

The results, reported in the journal *Entropy*, cast light on how religious beliefs emerge and develop. For example, state-endorsed religions experienced stability. Evangelical religions, non-state-sponsored religions and mystery religions each had unique stability – existing on a "floodplain" that offers stability and the ability to change. Extreme traditions, such as those involving human sacrifice, were unstable and did not persist over time.

Researchers involved in this project intend to continue studying human beliefs and practices with the cultural landscape approach. Poulsen and DeDeo believe their work could uncover new knowledge and answer questions about world religions that we know little about. Their team is also interested in testing their model on other cultural and political practices.

Predicting Storm Behavior in Real Time

One of the first ACCESS allocations helps researchers forecast Hurricane Ian's flooding impacts.

Image:
Flooding of a Florida neighborhood hit by Hurricane Ian.

tinyurl.com/storm-real-time

PSC: Bridges-2
Pittsburgh Supercomputing Center

University of North Carolina, Chapel Hill; The Water Institute of the Gulf

Hurricane Ian, a monstrous Category 4 storm that made landfall in the United States on Sept. 28, 2022, is credited as being one of the deadliest hurricanes to hit Florida since the "Labor Day" Hurricane of 1935.

While Hurricane Ian claimed over 140 lives, the vast network of emergency and recovery services that sprang into action likely saved countless more with their critical decisions about where people should evacuate, what support services would be affected by the storm and where to stage relief efforts.

To make these decisions, officials at dozens of locations across the state and country needed the best information possible about the storm, local capabilities and response resources. In particular, they needed to know where the storm would hit, what direction it would be coming from and how powerful a punch it would pack.

Image:
Satellite view of Hurricane Ian over Florida.

Above all, they needed this information in real time because a delay of even a couple of hours would be too late.

Fortunately, scientists at the University of North Carolina, Chapel Hill and The Water Institute of the Gulf were up to the challenge. Thanks, in part, to an early research allocation via ACCESS, the team was able to use the Pittsburgh Supercomputing Center's Bridges-2 system to make storm surge forecasts. These accurate, real-time predictions were shared with emergency management decision-makers, aiding them in evacuation and disaster-response decisions that may have saved lives.

> **Because PSC has set us up with a dedicated virtual machine, we were able to test the latest forecast system developments in the months leading up to an event like Ian. It's exciting that in its first real-time usage, we were able to deliver a product that was usable [by the stakeholders].**
>
> Zach Cobell, The Water Institute of the Gulf

The Search for Better Battery Life

With the help of Purdue's Anvil supercomputer, Yale researchers are trying to find ways to improve lithium batteries.

tinyurl.com/battery-search

RCAC: Anvil
Rosen Center for Advanced Computing

Yale University

When your only source of power is a lithium battery, the time you can go between charges is often on your mind. Will you make it through that presentation before your laptop dies? Or can you take that trip to Florida in your new electric vehicle? Researchers at Yale University have been working towards the goal of making more stable, long-lasting batteries. Creating batteries that last longer will not only improve the lives of consumers who use technology on the go but would help keep battery use sustainable and less of a demand for rare earth metals.

Aakash Kumar, a postdoctoral associate in the Department of Mechanical Engineering and Materials Science at Yale University and his team collaborated on efforts that first involved work in the lab to find an ideal combination of material layers to maximize the stability and longevity of lithium batteries through the testing of these materials.

Purdue's supercomputer, Anvil, was then used to aid in the calculations and analysis of the data the lab work produced. "In my mind, running calculations on Anvil is literally like striking a hammer," says Kumar. "There is so much power." He added that one of the features of Anvil, its large memory resources, was particularly useful for his research. "Because there is so much memory associated with each core, I can let them work on more memory-intensive problems."

> **Running calculations on Anvil is literally like striking a hammer. There is so much power.**
>
> Aakash Kumar, Yale University

An allocation with ACCESS provides more than the hardware researchers need; ACCESS also offers the technical support scientists require to use these complex computer systems. Kumar was grateful for all the support he received while working on this project with his team. "I have to say that the support staff was amazing," says Kumar.

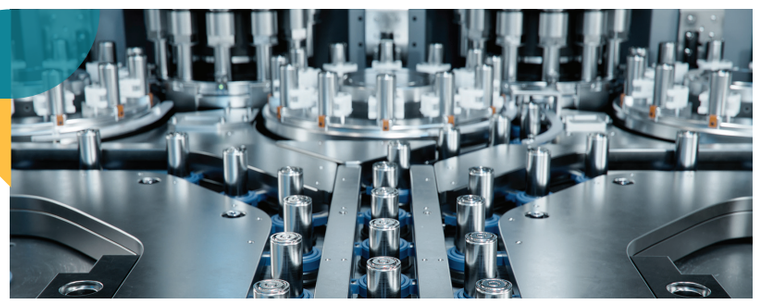

Image:
Industrial production line of lithium batteries.

The Ways of Water

Researchers from UC San Diego study water phases with the help of ACCESS resources at NCSA and SDSC.

tinyurl.com/water-ways

NCSA: Delta
National Center for Supercomputing Applications

SDSC: Expanse
San Diego Supercomputer Center

University of California, San Diego

Despite its ubiquitous presence in the everyday environment, our understanding of water has only slowly improved over the years. You may remember the scene in the 1993 movie Jurassic Park where chaos theory is explained using a dripping water faucet. This scene is based on real science. Certain properties of water remain notoriously difficult to predict but with recent advances in cyberinfrastructure and advanced resources like NCSA's Delta GPU-based supercomputer and SDSC's Expanse, two powerful ACCESS resources, simulations of even the most difficult models can be attempted.

Machines aren't the only things that have advanced far enough to make this breakthrough possible. The methodology has improved in the last three decades as well. Utilizing a method called first principles data-driven quantum simulations, the research team led by Francesco Paesani, chemistry and biochemistry professor at UC San Diego, was able to create a highly detailed many-body model of water called MB-pol. Their approach involved using the fundamental laws of quantum mechanics to predict the behavior of water molecules.

"Water's simple formula belies its complex behavior," Paesani said. "Using MB-pol, we've been able to model water across a wide range of temperatures and pressures, providing insights into how factors such as enthalpic, entropic and nuclear quantum effects shape its free-energy landscape. This work illustrates how recent advancements in first-principles, data-driven simulations have opened the door to realistic computational studies of complex molecular systems."

> **"Given the computational demands of our extensive simulations, it is important to underscore that our comprehensive explorations of water's phase diagram, spanning wide temperature and pressure ranges over long timescales, would not have been possible without the availability of high-performance computing resources.**
>
> Sigbjørn Bore, UC San Diego

Studies like these show the importance and power of using data-driven quantum simulations. By utilizing advanced computational methods, scientists can study complex systems like water with remarkable accuracy and detail in a highly customizable environment allowing for a greater understanding of complex systems. But there's also a more immediate impact to this research. Using the MB-pol model, scientists could study how water might behave in a deep sea station or even on Mars. If humans ever plan to inhabit other planets, it'll be simulations like these that ensure the fundamental necessities of life are taken care of for future explorers.

A Pore Solution

TACC supercomputers help researchers study some of the tiniest pores in the human body.

tinyurl.com/pore-solution

TACC: Frontera, Stampede 2
Texas Advanced Computing Center

University of Illinois Urbana-Champaign

Everything that comes and goes from a cell's nucleus is controlled by a membrane. Like your skin, this membrane has pores that allow materials to pass through from within and without. These materials are called macromolecules. Compared to other molecules, macromolecules are large, formed by thousands of covalently bonded atoms. Yet they are still so tiny that they can't be seen with visible light. Even using X-rays to see them is tricky, making observing them in action a challenge.

By programming simulations on supercomputers and using detailed computational models, scientists such as David Winogradoff, a post-doctorate from the University of Illinois Department of Physics, can see what's going on inside the nucleus of a cell – including how the dual-layered membrane of a nucleus works. Winogradoff and his research team have been studying how these membranes function using the supercomputers Frontera, Stampede2 and Blue Waters. Their models

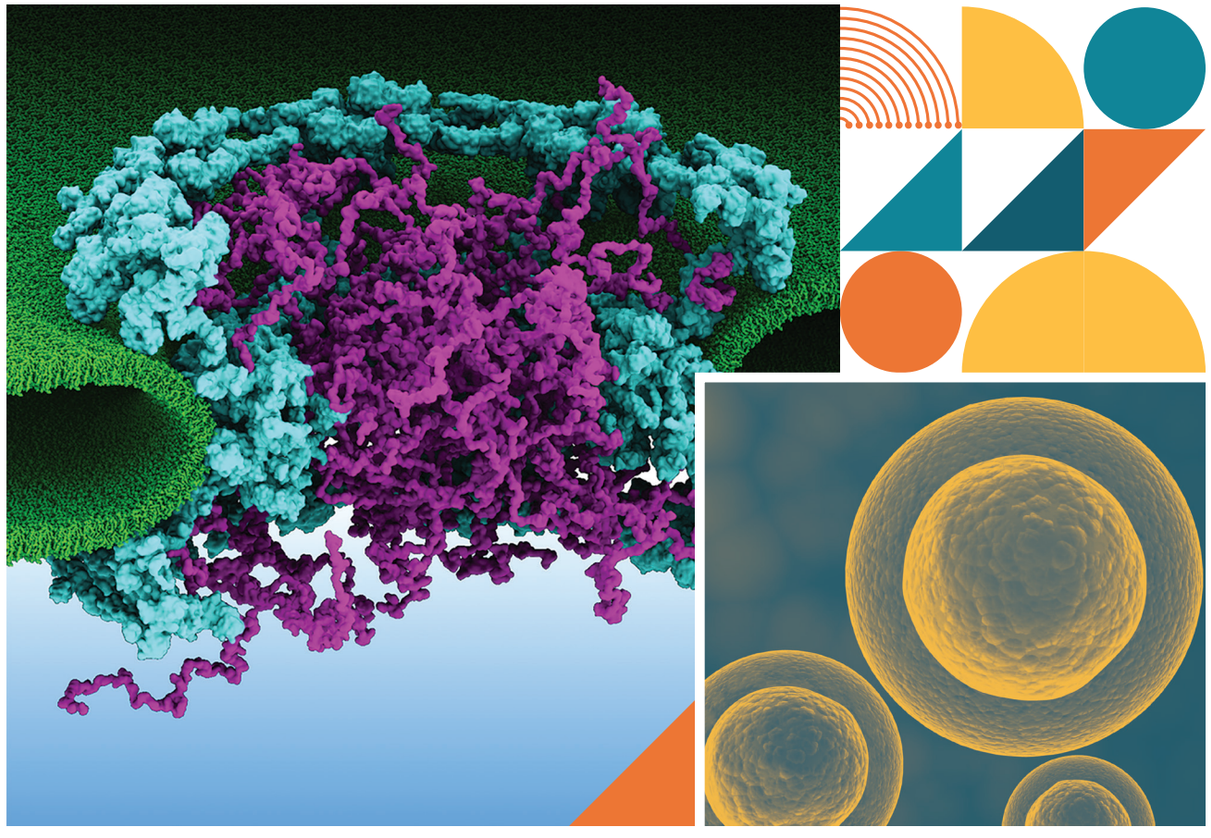

Image:
Side view shown of the simulation systems.
Winogradoff, D., Chou, HY., Maffeo, C. et al.

showed how the nuclear pore complexes (NPC) operate within the nucleus, yielding important discoveries, including the passive transport of biomolecules through the NPC.

"Our main finding is that a mesh-like interior of the nuclear pore exhibits a switch-over behavior based on protein size changing from a soft barrier for small proteins to a hard barrier beyond a certain threshold, essentially making it very difficult for proteins to get through," said study co-author David Winogradoff.

The hope is that this study, published in the journal *Nature Communications*, will lead to new advancements in drug therapy. Having access to supercomputing resources allowed for this research to be completed quickly and more accurately.

> **"What really helped in having Frontera and Stampede2 resources available was that we were able to simulate a large range of sizes of proteins. It sped up the process, which just took days versus months of/from local resources. Being able to run multiple replicas under different conditions gave a lot more weight to the robustness of our results.**
>
> David Winogradoff, University of Illinois, Physics Department

The Best-dressed Dinner Guest

Researchers use ACCESS resource Ookami to study the interconnectedness between krill connectivity and penguin populations.

tinyurl.com/dressed-guest

IACS: Ookami
Institute for Advanced Computational Science

Stony Brook University

Penguins always come to supper with their finery on. It hardly matters that their dinner rarely varies. *Pygoscelis penguin*, a genus that covers three distinct species of penguin, also called brush-tailed penguins, are much smaller than the famous emperor penguin. They live in the southernmost places on Earth, with some colonies living as far north as the Falkland Islands, a small collection of islands just off the southern tip of South America. These penguins primarily eat krill during the austral summer, which falls between November and March. This is why researchers at Stony Brook University have been using computational models run on the ACCESS resource Ookami to study the relationship between krill migration and *Pygoscelis penguins*.

Krill are tiny sea creatures with a dynamic population, their numbers and location changing dramatically between years and seasons. They have cycles of abundance that occur every four to six years in the West Antarctic Peninsula (WAP) region. Similar to how you might notice more cicadas during certain years, krill populations can explode due to a number of factors. If all a penguin ate was krill, you can imagine these types of fluctuations might affect where they travel or how successful they are in raising their young. Where the krill go, in theory, the penguins follow.

Image:
Adelie penguins in action.

Creating a computational model of krill populations and movement is tricky work. Krill exist within a delicate system balanced by cyclical changes in ice dynamics. In addition, krill fishing can impact this system as well.

Researchers have been studying penguins in the WAP region and have noted a trend in penguin populations around an area called the Adélie Gap. In particular, they have observed that penguin populations are growing or stable north of the Gap while they struggle elsewhere. Katherine Gallagher, a biological oceanographer, NSF Office of Polar Programs Postdoctoral Research Fellow and the PI of this NSF-funded project, is part of a team researching these penguin populations. Gallagher's team has hypothesized that krill connectivity plays an important role in *Pygoscelis penguin* population dynamics north and south of the Adélie Gap. They recently published a paper in *Scientific Reports* about their research.

"While the metrics that we calculated on large spatial scales weren't statistically related to penguin population dynamics, we did find some interesting trends that support hypotheses that prey, in this case, krill, may become limiting for penguins and other predators around the Adélie Gap," explained Gallagher. "This means that between the dynamic krill populations, predator demands and the growing krill fishery, that there might not be enough krill to go around. We're hoping that this new paper will help inform some future management decisions to help protect the krill and their predators."

Studying the migration of actual krill would be difficult work. You can't tag individual krill and because their life cycle takes them to hard-to-reach places, observing populations migrate would be like trying to track and follow individual specks of dust flying around in the air. However, computational modeling can fill in those gaps by simulating krill populations. To create the model, Gallagher's team used observational data on a krill behavior called diel vertical migration and used it to model krill in the Regional Ocean Modeling System (ROMS). This vertical migration behavior helps krill avoid visual predators like penguins during the day. ROMS was developed by researchers at Rutgers University and University of California Los Angeles. It's a highly complex model researchers use to test how things like how temperature or weather affect a specific area of the ocean.

The more complex the variables entered into a model like this, the longer simulations take to run. Supercomputers can reduce simulation time by factors of a hundred or more if the model is tuned correctly. For their results, Gallagher's team relied on the power of Ookami. With the help of data analysts assigned to the project, the code was made efficient for running on Ookami, speeding up the time it took to run simulations.

"We worked with Dr. Gallagher and her team on porting and tuning their application to the cutting-edge architecture of Ookami (Fujitsu A64FX)," said Eva Siegmann, lead research scientist for Ookami. "Porting software to a new system can be a challenging task. Our support team is happy to help and share their expertise. ROMS now runs seamlessly on Ookami, its performance is very good and its power consumption is exceptional (using two to three times less power than on a standard x86 system)."

> **"This was my first experience running physical ocean models such as ROMS by myself, instead of relying on others to do the model runs for me, so I am incredibly thankful to the Ookami team for their support during this learning process and for their help optimizing ROMS for the Ookami operating system. Thanks to their help, we were able to get to a point where we are able to simulate nearly 200,000 krill in the coastal ocean over a five-month period in less than two days.**
>
> Katherine Gallagher, NSF Office of Polar Programs Postdoctoral Research Program (PRFP) Fellow, Stony Brook University

Imagery:
Adelie penguins on the shoreline in Antarctica (top) and krill (below).

ACCESS
A Look Back and Ahead

Enabling Program Success

The ACCESS Coordination Office

With the rest of our ACCESS co-awardees, the ACCESS Coordination Office (ACO) spent the early part of our first program year standing up the structures to help facilitate governance, communications, community building and engagement, and evaluation across the program. These include many behind-the-scenes items, from establishing emails for all staff to setting up wiki spaces and Google drives for teams to work within. These early efforts helped ensure a smooth transition from NSF's XSEDE program for staff and the greater research community. Among the most important orders of business was to establish program governance.

By the time ACCESS launched to the public in September 2022, the ACO facilitated the seating of an Executive Council made up of service team principal investigators, each with voting rights and additional personnel to help facilitate the meetings along with our NSF program officer and an External Advisory Board member. The External Advisory Board was chartered shortly after launch and the ACO set about the work of calling for nominations and selecting a committee of 10 members representing a diverse range of backgrounds and perspectives to provide expert advice to the Executive Council on a range of topics key to the program's success. Finally, the ACO assisted with the establishment of various committees and working groups made of members across service teams to accomplish the important day-to-day work of the program. Those teams report regularly to the Executive Council and receive guidance as needed.

With a program of this complexity, having communications channels and strategies established early provides the framework within which internal audiences share information and report to the greater research community. Prior to launch, a brand was developed along with a website so people could connect with ACCESS on day one. Over the course of the first year, communications tools included the launch of internal and external newsletters and social media platforms, leveraging subscriptions and followers from the previous program. At PEARC, leading up to the launch of ACCESS, PIs presented on the program and once launched, ACCESS had a prominent presence at SC, the annual high-performance computing conference. ACO personnel oversee a committee and working group dedicated to communications and web presence across the program.

Community building and engagement is central to the program's mission of democratizing access to the national cyberinfrastructure. To this end, the ACO has created and oversees a team defining communities and opportunities for engagement, and determining how to expand participation to new and underserved communities as well as deepening participation from those already engaged with ACCESS. CB&E activities intersect with both communications and evaluations, and personnel from these teams meet regularly to discover how to best support the efforts of each. Among the engagement opportunities facilities by the ACO are participation at AIHEC23, CASC, Boise State Research Computing Days, CNI Spring Meeting, MS-CC Annual Meeting and RMACC23.

Finally, to set standards for achievement across the program, the ACO has established a team that meets regularly to discuss and develop measurement tools and processes. This cross-team group has worked together to create KPIs and metrics, identify commonalities, redundancies and gaps, and create strategies for data collection and sharing.

There is still much work to do at this point in the ACCESS program lifespan. From an ACO perspective, that translates into coordinating and facilitating activities supporting the program goals both internally and externally.

Internally, that means assisting the EC, working groups and standing committees with process recommendations to improve efficiency, cohesion and communication. The ACO will also work across service teams to understand the needs for evaluation and support those activities. Supporting DEI efforts across the programs is a high priority.

Externally, the ACO envisions working with service teams to help improve their communication and engagement with stakeholders. We anticipate continuing efforts to strengthen the experience for visitors across the ACCESS web presence. We'll highlight stories of ACCESS-enabled research at under-served institutions and support engagement with those communities.

ACO by the Numbers

September 2022–August 2023

84,856

Social media impressions

350

Total social media posts

38

Stories published on our site, access-ci.org

11,575

New subscribers to the ACCESS Advance newsletter, totaling 56,326

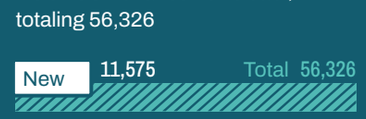

Access and Equity

Allocations Services

The first year of ACCESS Allocations was a busy one, with nearly 1,200 project requests received and processed as we began pursuing our vision of making access to the NSF-funded national cyberinfrastructure accessible and equitable for all researchers no matter the size of the institution, the scale of the planned work, the discipline of the research or the demographics of the requestor.

Following an initial scramble to prepare to launch the ACCESS program on September 1, 2022, the ACCESS Allocations team rolled out a new experience for U.S.-based researchers and educators that included not only a brand new portal experience but also new allocation policies designed to create an open, inviting and democratized allocations marketplace. The policies defined an efficient, scalable and simplified request-and-review framework that is supported by a robust, decentralized and flexible software platform – the eXtensible Resource Allocation Service (XRAS).

We defined these goals to maximize the impact of the NSF-funded cyberinfrastructure and keep our focus on getting researchers onto the resources they need to achieve their scholarly and instructional objectives. ACCESS, the Allocations team and the community continued to move quickly, and we were excited to report that the typical project in the first year of ACCESS experienced an ecosystem access time of just 12.8 days. That timeline included a median of less than a day to approve most project requests, about five days to complete an exchange of ACCESS Credits for a resource allocation and just 7.1 days on average to go from credit exchange to recording their first resource usage.

Another key to our Plan Year 1 progress has been an increased focus on diversity, equity and inclusion (DEI) among the ACCESS community. It's not uncommon for researchers to feel their project isn't right for high-end compute resources. Maybe your datasets aren't particularly large, or you only need a small amount of computing time for classroom activities. The Allocations team and our dedicated DEI facilitator are engaging with the community to help us explain how many more types of activities can benefit from what ACCESS offers. Our goal is to make it clear that ACCESS resources are ready to serve a diverse community of researchers with many different cyberinfrastructure needs. As part of our comprehensive DEI plan, we have attended outreach events, developed and delivered facilitated DEI discussions and supporting materials as part of the review process, conducted community focus groups, ACCESS DEI policy review, ACCESS-wide collaboration, and developed DEI metrics and goals.

We received many positive comments about the new allocation policies and approach throughout the year. Particularly popular changes included allowing graduate students to request and lead a small-scale project to support their dissertation or thesis work and permitting researchers to align their allocated projects more closely to their funded research awards, including awarding projects for the duration of a supporting grant. A few of our more visible changes had more mixed reviews: the four escalating project types and the introduction of ACCESS Credits. On the whole, however, results from a survey of individuals who interacted with the allocations system showed that we kept the community relatively satisfied – tallying a 4.1 rating on a 1 to 5 scale – with the allocations process as a whole and the user experience with the portal.

But we're not resting on these laurels. Within the Allocations team, we know we can do better. We're pressing forward in Plan Year 2 with three key areas of improvement – the Allocations web presence overall; helping researchers select resources that match their needs in collaboration with the Support and Operations teams; and system features for researchers, reviewers and administrators that allow us to more efficiently implement the ACCESS allocation policies. By the end of Plan Year 2, we plan to have a completely overhauled portal experience for researchers who come to ACCESS to advance their research, with dynamic pages that make it easier to create, view and manage projects and access to ecosystem resources.

And we want you to help us make the ACCESS experience even better. Visit the Allocations area of the ACCESS website and send us your feedback or volunteer to participate in the allocations process as a reviewer or focus group member. Together, we can advance innovation and take your research or classroom activities to the next level.

Allocations by the Numbers

September 2022–August 2023

10,962
Active individuals

741
Active institutions

99
Volunteer allocation reviewers

2,290
Active projects

Transformative Assistance

Support Services

Positioning ACCESS Support for evolution and growth has been a key focus during the first year of the project. Although our primary goal was to ensure the user experience was not negatively impacted during the XSEDE transition, our longer-term goals were focused on the future, ensuring that the new services being implemented would be transformative, easy to navigate, access and understand. To ensure success, our team worked closely with other service teams as we navigated the ticketing system, integrating the knowledge base, identity and access management, incorporating our services with the allocations process, and integrating metrics. We also worked with the Resource Providers (RPs) to develop sound processes, such as routing tickets and writing and posting documentation, to ensure the user needs would be served. We connected with community groups and provided information on our tiered support model to help researchers and cyberinfrastructure professionals discover the variety of resources available to them.

Deployment of our tools was a key component of our successful Year 1 plan. The ACCESS Support team leveraged a variety of existing community tools, such as Open OnDemand, the Pegasus data workflow tool and the Connect.CI portal as a part of ACCESS' new portfolio. Integration of these tools together, as well as with the processes of other service teams, allowed for the Support team to quickly formulate the user experience for ease of access to Resource Provider resources.

In addition to tool deployment, we also introduced new services to the larger community. Our MATCH Plus and MATCH Premier projects offered a new and novel way to help researchers remove roadblocks created by computational inefficiencies or other issues on their projects. Our community grants program (called CCEP) provided travel opportunities to community members who contributed to the ACCESS ecosystem. Each of these services, in addition to the democratization of access to resources provided by our tools, promoted diversity, equity, inclusion and accessibility, which is the backbone of a successful user experience.

As ACCESS Support moves into the next stage of the project, our goals have shifted to be focused less on the transition or standing up new services and more on the growth of those services and improving the user experience. With an eye toward our primary goals – providing scalable support, democratizing access, decreasing time to science and supporting research – we are pushing forward several initiatives to ensure success. A large component of this involves working with community members to solicit feedback and acting on that feedback to improve tools, language, accessibility or other components of our services. Some aspects of improvement include updates to the website to enable the discovery of information, leveraging Carnegie codes to enable better reporting and enhancing tag taxonomy to assist users in discovery. The continual evolution of Open OnDemand includes expanding deployments on ACCESS resources, and centralizing the deployment and working with the ACCESS Metrics team to collect metrics from ACCESS resources. Pegasus's evolution includes evolving accessibility, which includes developing example workflows, developing videos and training materials, and tighter integration of software components of Pegasus ACCESS to improve ease of use.

We are also expanding the resources available in the Knowledge Base, which is closely integrated with an evolution of our CCEP program as we make more specific requests to the community for material. An expansion of training topics and materials also supports an enhancement of our Knowledge Base.

Our concierge-level opportunities are undergoing expansion and evolution by working with community groups, NSF awardees and the Resource Providers. MATCH Plus and Premier, as well as CCEP, are learning directly from our current experiences as well as feedback from surveys and conversations with existing users to provide an even broader impact to our community and researchers.

ACCESS Support is also seeking out ways to broaden participation among potential community members, particularly focusing on those from under-resourced institutions, under-represented groups or scientific disciplines who are not traditionally users of advanced cyberinfrastructure. The importance of language and verbiage when describing these services, whether on the webpage or during talks, is incredibly critical to the successful inclusion and sense of belonging among these groups.

Support by the Numbers

September 2022–August 2023

7 of 12

Site installations of OnDemand services, with active in-process discussions at 2 additional sites

43 out of 89

CCEP awards made, with 89 total applications received

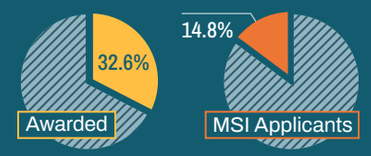

Applicant Institutional Stats

Distinct academic institutions:	46
Minority-serving institutions:	8

3,000

Unique user average per month through the Support Portal

Page views per month: 8,500–10,500

A Multifaceted Infrastructure

Operations Services

ACCESS Operations provides the core infrastructure for the ACCESS project, including operational support, networking and data support, and cybersecurity support.

During the first year of ACCESS, the Operations team celebrated many successes. The team lost no time ramping up for the September 1, 2022 go-live date. During the first four months of the project, the cybersecurity team built and deployed a comprehensive identity and access management infrastructure to support 100,000+ user accounts and established active cybersecurity vulnerability scanning. The networking team implemented a seamless network migration path for Resource Providers (RPs) to use CONECTnet. The operational support team seamlessly migrated key enterprise services, implemented an interim ACCESS ticketing system and co-led the effort to create integration roadmaps to allow for the seamless transition of all XSEDE-allocated resources to ACCESS with zero interruptions to the user community.

While the team took a moment to celebrate a successful transition from XSEDE to ACCESS operations on September 1, there was still much work to be done. The RP Coordinator began conversations leading to the XSEDE Service Provider Forum's rechartering for ACCESS and conducted an RP survey which led to improved Integration Roadmaps being developed. The operational support team implemented a new ticketing system in collaboration with other ACCESS program areas. The networking team developed improved IP address and DNS name management. The cybersecurity team updated security and access policies. Late in the first year, the team successfully recruited a diverse group of students (including over 66% female) for the inaugural Student Training & Engagement Program (STEP).

As the project matures during Plan Year 2 (PY2), Operations' plans are largely focused on working to expand the community ecosystem to new and novel resources and partners and improving on existing infrastructure. PY2 plans for the operational support team include expanding and enhancing infrastructure discovery services, conducting a pilot of an AI Chatbot and CI Pathfinder, improving Integration Roadmaps and operations software communications, implementing information-sharing, platform-secure public APIs and improving usage tracking capability. The network team plans to finalize the transition to CONECTnet L3VPN and add new sites, integrate new RP perfSONARs to the project's MaDDash dashboard, update to a new perfSONAR database format and investigate new perfSONAR service capabilities. The networking team also plans to engage the ACCESS Metrics team to evaluate data storage handoff as well as investigate new metrics reporting options for RPs with Internet2. During PY2, the cybersecurity team will ensure ingestion of security log data by OmniSOC and establish a functioning Security Event and Incident Management (SEIM) capability, implement non-user OAuth tokens for inter-service authentication, and develop a system security baseline compliance process with configuration management tool support. The project will turn more attention to broadening participation and DEI efforts by engaging a DEI consultant to help outline plans and accomplish DEI goals. It will also carry out a successful Student Training & Engagement Program that aims to build student interest in pursuing careers in cyberinfrastructure operations, data and networking, and cybersecurity. Early results of this program have been extremely positive.

Operations Components

Operations by the Numbers

September 2022–August 2023

0 Interruptions

Successful transition of all XSEDE-allocated resources to ACCESS with 0 interruptions to the user community

Implemented a seamless network migration path for RPs to use CONECTnet

100,000+

Built and deployed a comprehensive identity and access management infrastructure to support 100,000+ user accounts

66%+

Recruited a diverse group of students for the inaugural Student Training & Engagement Program (STEP)

Measuring Success

Metrics Services

The first year of ACCESS Monitoring and Measurement Service (MMS) kept the Metrics team busy both with the transition from a monitoring service that was independent of XSEDE, NSF's program prior to ACCESS, to one that is now fully integrated within the ACCESS program as well as the ongoing development of XDMoD and Open XDMoD.

A critical goal of the transition was to minimize disruption to the CI research community so that researchers were able to seamlessly utilize the CI resources within the ACCESS environment from day one of the go-live. Demonstrating the effectiveness of the ACCESS transition is where the Metrics team, with its extensive repository of usage data for XSEDE, comes in. To demonstrate that there was little, if any, change in the use of CI resources by researchers, ACCESS Metrics leveraged XDMoD to provide a comparison of system usage and job throughput between XSEDE and ACCESS. The analysis included metrics such as the number of jobs run, number of service units consumed, breakdowns of usage by resource, type of user and field of science. Below is a representative plot that shows usage by NSF Directorate pre-and post-ACCESS go-live (September 2022).

With the development and deployment of novel CI architectures, such as ARM, there comes a need to evaluate their overall performance with respect to the research applications that will run on them. To this end, a novel study of application performance and energy efficiency for ARM-based processors as compared to x86 processors was carried out and the results presented at the ARM HPC User Group SC22 symposium and in an accepted paper for the International Workshop on ARM-based HPC: Practice and Experience (IWAHPCE-2023). ARM performance, while generally slower, was nonetheless shown in many cases to be comparable to current x86 counterparts and often outperforms previous generations of x86 CPUs. In terms of energy efficiency, which considers both power consumption as well as execution time, ARM was shown in most cases to be more energy efficient than x86 processors. Given ongoing improvements in ARM compilers and libraries, these results hold promise for the application of future generations of ARM-based systems in high-performance computing.

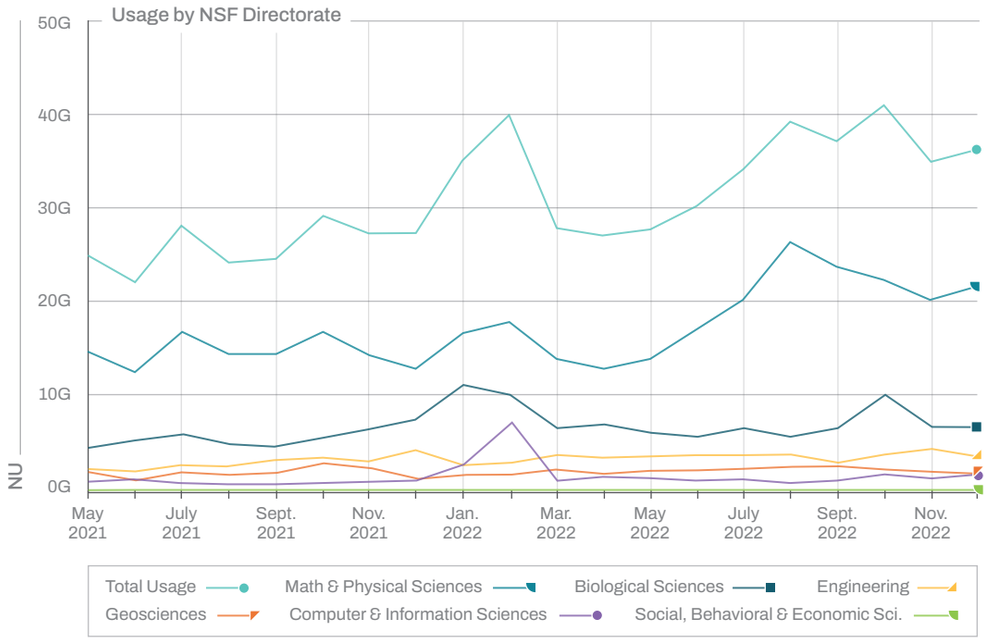

Another exciting development during the first year of the Metrics team's work, which has continued into the second year, is the development of our new data analytic framework that allows researchers to directly access the repository of usage and performance data in the ACCESS XDMoD data warehouse using Python through Jupyter notebooks. While ACCESS XDMoD itself provides a powerful interface for data analytics, the new analytic framework allows data scientists and other interested parties to carry out studies using the analytical tools they are most familiar with.

The second year will build on momentum from the first year. The new data analytic framework will be incorporated into the production release of ACCESS XDMoD and therefore widely available to the research community. In addition, ACCESS XDMoD and ACCESS OnDemand will be integrated so that researchers will be able to directly access job-level performance data through the ACCESS OnDemand portal and OnDemand usage metrics will be available in ACCESS XDMoD. Finally, in collaboration with ACCESS Allocations and Support, we will work to provide a much-desired tool to help guide users in the optimal CI resource to carry out their research.

Metrics by the Numbers

September 2022–August 2023

17,106,097
Total jobs run

9,653
Active users

3,054,923,444
CPU-hours used and counting

Funded Partnerships

ACCESS wouldn't be possible without our extensive network of collaborators. Made possible through these lead institutions – Carnegie Mellon University, University of Colorado Boulder, University of Illinois Urbana-Champaign and State University of New York at Buffalo – the ACCESS partnership also includes the following institutions:

Case Western Reserve University

Center for Advanced Computing (CAC) at Cornell University

Florida International University

Information Sciences Institute at the University of Southern California

Internet2

Massachusetts Green High Performance Computing Center

National Center for Atmospheric Research (NCAR) at the University Corporation for Atmospheric Research (UCAR)

National Center for Supercomputing Applications (NCSA) at the University of Illinois Urbana-Champaign

National Institute for Computational Sciences (NICS) at the University of Tennessee, Knoxville

OU Supercomputing Center for Education & Research at the University of Oklahoma

Ohio Supercomputer Center

Pervasive Technology Institute (PTI) at Indiana University

Pittsburgh Supercomputing Center (PSC), a joint effort of Carnegie Mellon University and the University of Pittsburgh

Rosen Center for Advanced Computing at Purdue University

San Diego Supercomputer Center (SDSC) at the University of California, San Diego

Shodor

Southeastern Universities Research Association (SURA)

Texas Advanced Computing Center (TACC) at the University of Texas at Austin

Tufts University

University of Chicago

University of Kentucky

Contributors and Credits

Managing Editor

Dina Meek, NCSA

Design

Megan Janeski, NCSA

Writing and Data Contributions

Megan Johnson, NCSA; Kimberly Mann Bruch, SDSC; Andrew Helregel, NCSA; Stephen Deems, PSC; Leslie Froeschl, NCSA; John Towns, NCSA; Tim Boerner, NCSA; Shelley Knuth, UC Boulder; Tim Boerner, NCSA; Tom Furlani, SUNY Buffalo.

Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) is supported by the National Science Foundation.

ACCESS-CI.org • @ACCESSforCI