Overview of ACCESS allocated cyberinfrastructure usage

Joseph P. White
Robert L. DeLeon
Aaron Weeden
Matthew D. Jones
Thomas R. Furlani
jpwhite4@buffalo.edu
rldeleon@buffalo.edu
aaronwee@buffalo.edu
jonesm@buffalo.edu
furlani@buffalo.edu
Center for Computational Research, University at Buffalo
Buffalo, NY, USA

ABSTRACT

ACCESS is a program established and funded by the National Science Foundation to help researchers and educators use the NSF national advanced computing systems and services. Here we present an analysis of the usage of ACCESS allocated cyberinfrastructure over the first 16 months of the ACCESS program, September 2022 through December 2023. For historical context, we include analyses of ACCESS and XSEDE, its NSF funded predecessor, for the ten-year period from January 2014 through December 2023. The analyses include batch compute resource usage, cloud resource usage, science gateways, allocations, and users.

CCS CONCEPTS

• **General and reference** → **Measurement**; **Evaluation**; • **Information systems** → *Computing platforms*.

KEYWORDS

XDMoD, Data Analytics Framework, ACCESS, XSEDE

ACM Reference Format:
Joseph P. White, Robert L. DeLeon, Aaron Weeden, Matthew D. Jones, and Thomas R. Furlani. 2024. Overview of ACCESS allocated cyberinfrastructure usage. In *Practice and Experience in Advanced Research Computing (PEARC '24), July 21–25, 2024, Providence, RI, USA*. ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3626203.3670521

1 INTRODUCTION

ACCESS [1] is a program established and funded by the U.S. National Science Foundation (NSF) to help researchers and educators use the NSF national advanced computing systems and services. ACCESS consists of a variety of diverse cyberinfrastructure (CI) resources including compute resources, storage systems, and networking infrastructure; the main focus of this paper is compute resources. This paper provides a series of usage analyses intended to illuminate ACCESS usage and contextualize it with its predecessor program XSEDE. These types of study are important to show how well programs such as ACCESS are serving the scientific research and education community, pointing out both strengths and weakness, and to provide insight how to improve the program as it progresses.

2 BACKGROUND

The ACCESS program is composed of several tightly-coupled teams that provide services to the U.S. research community. As of March 2024, there are five teams: Allocations, Operations, Support, Metrics, and a Coordination Office. ACCESS follows the highly successful NSF eXtreme Science and Engineering Discovery Environment (XSEDE) [9] and XD Metrics Service (XMS) programs that provided access to and monitoring of national CI resources, respectively. The ACCESS program (and XSEDE previously) does not directly manage national CI resources. These are typically NSF-funded resources run by various resource providers (RPs) at different institutions across the U.S.

The ACCESS program awards began in April 2022, and there was a six-month transition period from XSEDE to ACCESS. The switchover from XSEDE to ACCESS occurred on September 1, 2022, when XSEDE-operated services shut down and ACCESS entered national production. The transition from XSEDE to ACCESS was structured to make it as seamless as possible for the researchers using CI. There were numerous changes to services; such as websites, login mechanisms, and allocation policies. However, many components remained; for example, the various resources that were integrated with XSEDE continued to be allocated under ACCESS. Researchers who had existing active allocations under XSEDE could continue to use them under ACCESS.

The ACCESS Monitoring and Measurement Services (Metrics) team serves in the important role of monitoring ACCESS-integrated CI to ensure optimal performance, robustness, and usage (including compute, cloud, storage, networking, software/data services, etc.). ACCESS Metrics also provides services to several other related NSF-funded programs such as Campus Cyberinfrastructure (CC*) [7] and

PEARC '24, July 21–25, 2024, Providence, RI, USA

ACM ISBN 979-8-4007-0419-2/24/07
https://doi.org/10.1145/3626203.3670521

Cyberinfrastructure for Sustained Scientific Innovation (CSSI) [8]. ACCESS Metrics follows the successful Technology Auditing Service (TAS) and XD Metrics Service (XMS) programs that monitored the resources allocated by XSEDE and tracked historical usage data dating back as far as the NSF TeraGrid [2] program. The main utility ACCESS Metrics uses for ingesting and reporting data is XDMoD [4], which provides a web-based portal with capabilities for data exploration, visualization, and export. The ACCESS XDMoD data warehouse contains information about the usage and performance of ACCESS allocated resources and the historical usage data collected under both the XMS and TAS programs. This facilitates analysis of usage trends over a large time period (the earliest data are from January 2004).

3 METHODS

The main tool used for the data analysis in this paper is ACCESS XDMoD, whose graphical interface allows users to make charts rapidly and easily, export data, and construct reports on various aspects of CI usage and functionality. A recent major addition to XDMoD is the Data Analytics Framework [10], which provides API access to the XDMoD data warehouse via Python and Jupyter Notebooks.

The data in ACCESS XDMoD are obtained from a variety of sources. These include information about project allocations from the ACCESS Allocations eXtensible Resource Allocation Service (XRAS) database, information about compute usage reported to the ACCESS Allocations service by RPs, job scheduler log files for batch (HPC and HTC) resources provided to ACCESS Metrics from RPs, OpenStack usage logs from cloud resources, node-level job performance data files provided to ACCESS Metrics from RPs, science gateway usage information, and resource specification data from ACCESS Operations' CyberInfrastructure Description Repository (CiDeR). The ACCESS XDMoD data warehouse also contains information curated by the Metrics team from publicly available datasets and from information exchanged directly between RPs and ACCESS.

In this paper, usage information is presented for multiple different compute resources with different hardware capabilities. In order to make meaningful comparisons of compute usage on disparate systems, we use the ACCESS Credit Equivalent service unit (ACE SU). One ACCESS Credit Equivalent is defined as one CPU Hour on the Expanse resource from San Diego Supercomputer Center, which is a compute resource based on AMD EPYC 7742 processors. Each RP determines the exchange rate between an ACCESS Credit and a service unit on their CI resources. RPs in ACCESS define a base (constant) exchange rate. ACCESS also supports variable exchange rates so that RPs can change the effective cost of their resource to help manage demand. The fixed base ACCESS Credit exchange rate is the value used in this paper to compute the ACCESS Credit Equivalent conversion for each resource. The ACCESS Credit Equivalent allows comparison between usage of node-allocated, core-allocated and GPU-allocated compute resources. It also allows a comparison between resources with different compute power per core.

This paper will generally cover the 16-month period of time from the start of ACCESS production in September 2022 through December 2023. To present ACCESS in the historical context of NSF sponsored computational resources, we also display a ten-year history of usage, metrics and the various analyses that covers the period of XSEDE and the starting 16 months of ACCESS where such data are available.

4 RESULTS

The ecosystem of ACCESS allocated resources consists of a diverse variety of different types of cyberinfrastructure (CI) including batch HPC/HTC compute, clouds, storage hardware, associated software, allocations procedures, and the user base. It is not possible in a single paper to present a comprehensive analysis of all of these systems and sub-systems. This paper will focus on a broad overview of batch compute, gateway, and cloud usage; briefly discuss allocations and innovative CI systems; and present some features of users and usage patterns.

4.1 Batch Compute Usage

In this section we present compute usage from batch processing compute resources, which includes large compute clusters designed to support large parallel compute jobs (high performance computing HPC resources such as TACC Stampede2) and batch computing clusters tuned to support large numbers of serial compute jobs (high throughput computing HTC tuned resources such as SDSC Comet). Figure 1 shows the ten-year history of HPC/HTC usage for both number of jobs and the ACCESS Credit Equivalent SUs consumed by those jobs; these data are from the *Jobs* realm in ACCESS XDMoD. There are two main features to note in this figure. Interestingly, the two plots are reasonably linear with R^2 values of 0.97 and 0.87 for the jobs and ACE data respectively. The second feature is that the numbers of jobs and ACEs track together over the ten-year period including XSEDE and ACCESS. Dividing the ACEs by the number of jobs to obtain a mean job size, ACE per job, produces a plot (not shown) with some scatter but no significant trend. It is worthy of note that over the ten-year span, with the many changes in resources and user base, that the mean job size should remain so stable.

A similar plot of the ten-year history of GPU usage is shown in Figure 2. Note that the GPU ACEs and jobs run over the last three years have increased rapidly as GPU resources have become increasingly available and have been increasingly adopted by users. In contrast to the batch compute CPU usage, the GPU mean job size has increased significantly over the ten-year period by approximately a factor of three as seen in Figure 3. The GPU jobs and GPU ACEs are still a relatively small fraction of the total batch compute usage as seen in Figure 4. However, partly due to the increase in GPU job size, the fraction of GPU ACE usage is growing rapidly; it has increased by a factor of three in the last three years.

Application usage is tracked by the *Job Performance (SUPREMM)* realm in XDMoD. Unfortunately, historically not all resources have reported performance data. However, comparisons of the *Jobs* realm, which essentially does have universal reporting, and the *Job Performance* realm show similar historical trends. Therefore, we will use the *Job Performance* realm application usage as representative of the full ACCESS ecosystem. As shown in Figure 5, the fractions of Python jobs and ACEs are growing rapidly. Figure 6 shows that the fraction of Python GPU usage compared to all usage is growing

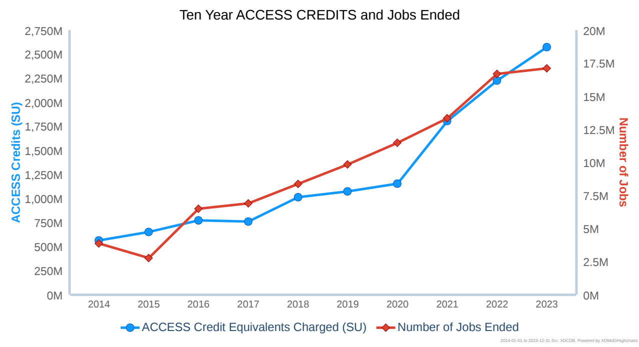

Figure 1: XDMoD plot of the ten-year history of HPC/HTC compute usage.

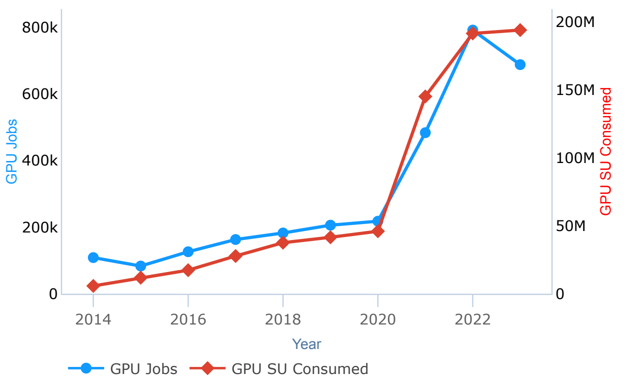

Figure 2: Plot of the ten-year history of batch compute usage on GPU resources.

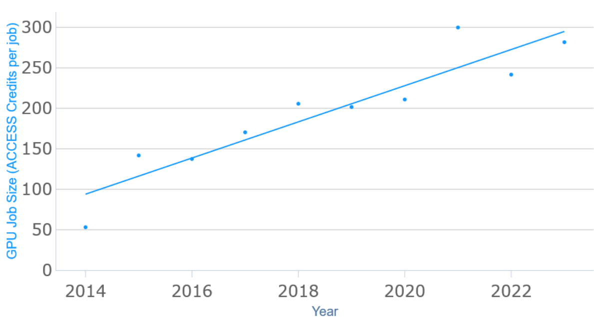

Figure 3: GPU mean job size in ACEs per job. The solid line is a linear regression with $p = 0.0001$ and $R^2 = 0.86$. Note the increasing trend in job size.

very rapidly. Finally, Figure 7 shows that the Python GPU job size is much larger than the Python CPU job size and is also increasing. Note that the *Job Performance* Python GPU usage data are only acquired for the Bridges-2 GPU resource from Pittsburgh Supercomputing Center. If one looks at the trend for other non-Python jobs for applications such as molecular dynamics, one does not see this

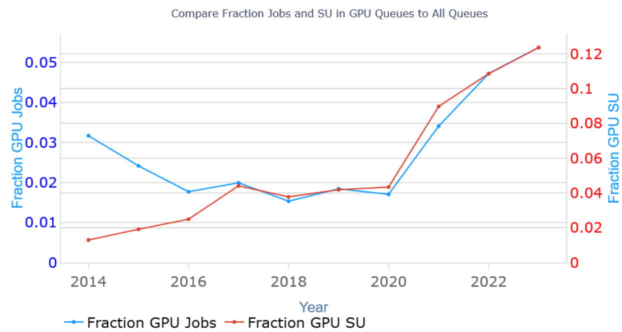

Figure 4: Fraction of GPU jobs and ACEs of total batch compute jobs. Although the GPU usage is relatively small, the fraction of ACEs it is growing rapidly. In the last three years the fraction of GPU jobs has gone up by a factor of three.

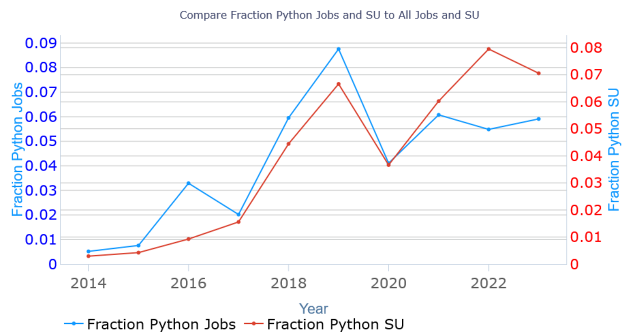

Figure 5: The fraction of Python usage compared to all usage.

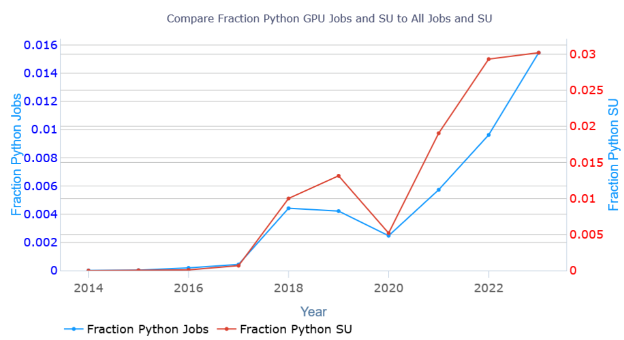

Figure 6: The fraction of Python GPU jobs and ACEs is increasing rapidly.

trend of increasing job size and usage. For example, for the NAMD and GROMACS molecular dynamics codes, which both have efficient GPU versions, there is very little growth in usage or change in job size. We think the trend in Python GPU usage is likely due to Artificial Intelligence (AI) code usage such as Keras (the Python-based API for TensorFlow), PyTorch, and other Python-based AI codes.

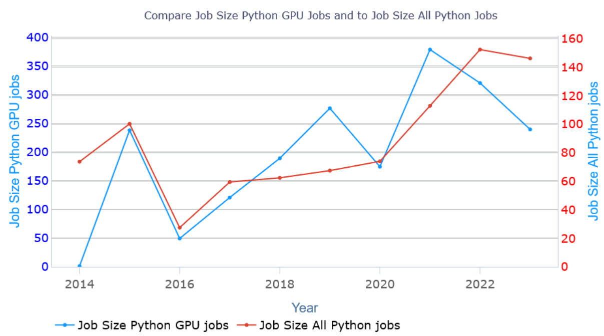

Figure 7: The job size of Python GPU jobs is increasing.

4.2 Cloud usage

Cloud resources were first made available in XSEDE with the Jetstream resource from Indiana University and the Texas Advanced Computing Center (TACC), which was in production from 2016 to 2022. The successor Jetstream2 resource entered production in 2022. Cloud resources only represent a small fraction of the total ACCESS usage (for example, in 2023 there were ~120 M ACCESS Credit Equivalents in cloud usage compared to ~2,600 M for batch compute jobs); however, cloud resources fulfill a need to support science gateways (see Section 4.3); provide an interactive, as opposed to batch, computational option; and support scaling for specialized projects. ACCESS cloud usage is shown in Figure 8.

The breakdown of allocation types using cloud resources is fundamentally different from the breakdown using batch compute resources (see Section 4.5 for details on the different XSEDE and ACCESS allocation types). Figure 9 shows the breakdown of batch compute usage by allocation type. It is dominated by the larger allocation types, *Maximize* and *Research*. In contrast, the breakdown of cloud usage, shown in Figure 10, is much more balanced among the various allocation types. This could be a reflection of the different users and project types on the cloud vs batch compute. HPC batch users tend to do larger, more heavily computational projects, whereas cloud compute may be doing smaller projects that support users' domain work.

4.3 Science gateway usage

Science gateways are services that provide access to CI tailored to the needs of a specific scientific community. Gateways typically have a web-based front end that is used for interacting with the service and software that manages scientific computations and data flows. There are two main mechanisms for gateway usage in ACCESS: gateways can run batch compute jobs on HPC and HTC resources, and/or gateways can use persistent cloud virtual machines (VMs) to run the web services and/or the compute jobs.

The numbers of gateway end users are provided directly from gateway providers (gateway compute usage as seen on a batch resource is all executed under a single "Community User" account). Not all gateways report user identities for the jobs that were run; this means some results in this section under-count users. There is also no federated identity provider between gateways; this means some results in this section count some users multiple times if they submitted jobs via multiple gateways.

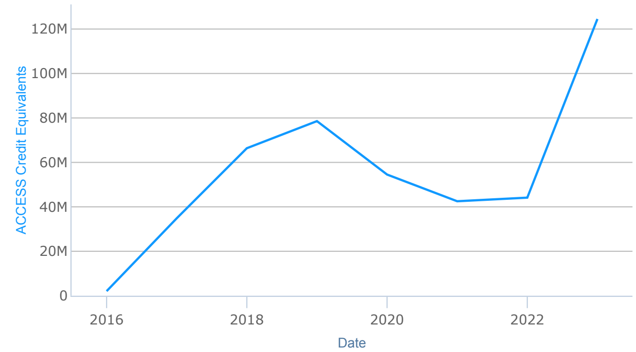

Figure 8: ACCESS cloud usage, in ACEs from 2016 to 2023. The Jetstream resource was in production from 2016 to 2022, and Jetstream2 entered production after Jetstream was retired.

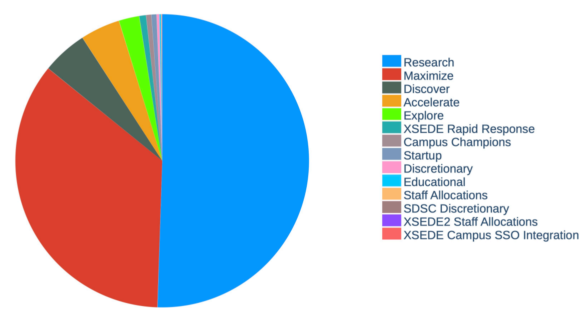

Figure 9: ACCESS batch compute usage broken down by allocation type. The batch compute usage is dominated by the larger allocation types.

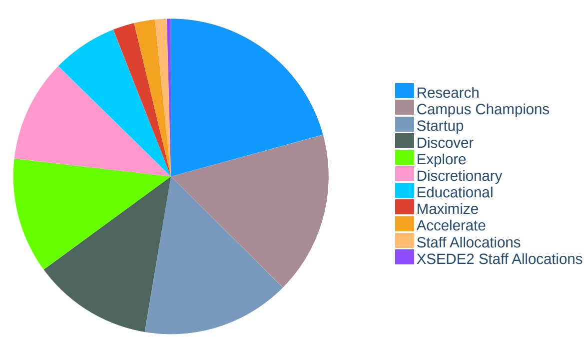

Figure 10: ACCESS cloud usage broken down by allocation type. In contrast to the batch compute usage, the cloud usage is much more balanced among the ACCESS allocation types.

4.3.1 *Batch compute jobs submitted via gateway.*

Gateways running batch compute jobs on ACCESS allocated resources from September 1, 2022, through December 31, 2023, supported a relatively small number of jobs but a large number of users. A total of 2,122,560 batch jobs were run on ACCESS allocated resources via gateways during this period. This is 9% of the total number of batch jobs run on ACCESS allocated resources during this period. The gateway jobs consumed 122,328,421 ACCESS Credit Equivalents (ACEs), which is 3% of the total ACEs consumed by batch jobs during this period. Figure 11 shows the monthly counts of ACEs consumed for gateway batch jobs and overall.

The jobs were run on 25 different gateways by 35,185 users, which is three times the number of users who ran non-gateway batch jobs during this period (11,524). Figure 12 shows the number of gateway users running jobs month by month eclipsing the number of users running non-gateway jobs. The University of Michigan hosted two of the gateways (I-TASSER and COSMIC2) that together constituted 77% of the jobs and 21% of the ACEs. UC San Diego hosted five of the gateways that together constituted 11% of the jobs and 17% of the ACEs. The remaining gateways were hosted by 16 other institutions across 10 U.S. states.

The gateway that ran the largest number of jobs and reported the largest number of users running jobs was I-TASSER (76% of gateway jobs; 17,577 users) followed by the Cipres gateway from UC San Diego (10% of gateway jobs; 12,180 users). The remaining 23 gateways each ran fewer than 3% of the jobs. Two of these reported a relatively large number of users: COSMIC2 (3,123) and ChemCompute from Sonoma State University (1,873). Of the remaining 21 gateways, eight reported fewer than 120 users running jobs, and 13 did not provide user information.

94% of the gateway jobs and 97% of the reported users ran on SDSC Expanse. There were five other ACCESS allocated resources on which jobs were run (PSC Bridges-2, TACC Stampede2, Anvil from Purdue University, Rockfish from Johns Hopkins University, and Delta from the National Center for Supercomputing Applications), each with fewer than 3% of the jobs. 54% of the ACEs charged for gateway jobs were on SDSC Expanse, 24% were on PSC Bridges-2, 17% were on Purdue Anvil, and the remaining three resources each had fewer than 4%.

93% of the jobs and 43% of the ACEs charged via gateway were for projects in the Biological Sciences. 48% of the ACEs charged were for projects in the Physical Sciences. The largest jobs in terms of ACEs charged per job were in the Physical Sciences (1,861 ACEs per job); other science areas each charged fewer than 300 ACEs per job.

4.3.2 Gateways running on cloud VMs. There were 22,759,887 ACEs consumed by ACCESS gateway allocations running on virtual machines on the Jetstream2 resource from Indiana University from September 1, 2022, through December 31, 2023. This is 14% of the total ACEs consumed by ACCESS allocations running VMs on Jetstream2 during this period. There were 53 gateway allocations running on cloud VMs, which is twice the number of gateways running batch compute jobs. The top four gateway allocations; together consuming 41% of the ACEs run via gateway on cloud VMs; were Galaxy (16%), EDGE COVID-19 (10%), brainlife.io (8%), and tRFtarget (7%).

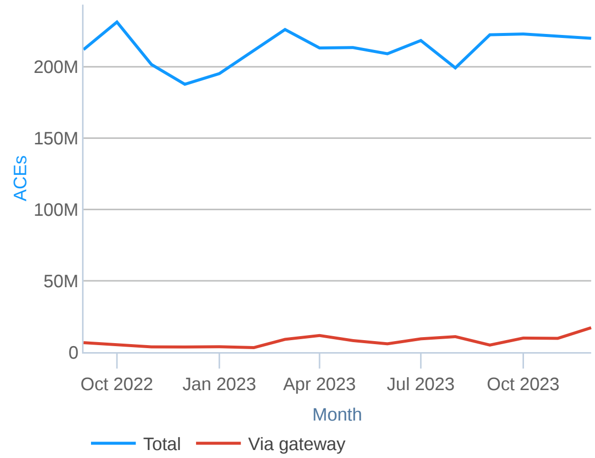

Figure 11: Batch compute jobs run via gateway only consume 3% of the ACCESS Credit Equivalents (ACEs) consumed by batch compute jobs overall.

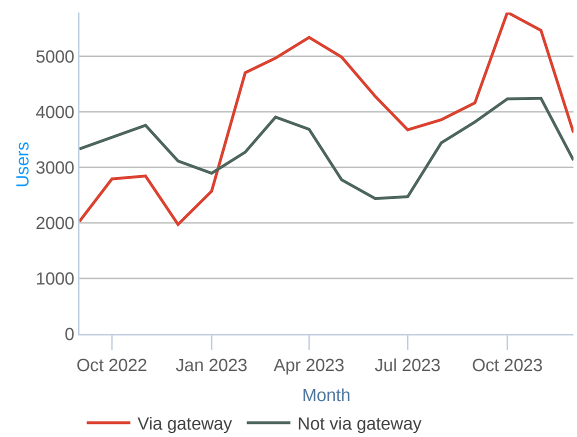

Figure 12: Gateways serve a disproportionately large user community compared to their moderate computational footprint. Starting in February 2023, the monthly number of ACCESS users running batch compute jobs via gateway eclipsed the number of users running batch jobs not via gateway.

4.4 Innovative Compute

NSF supports a number of smaller resources whose main function is to test how a variety of innovative computational modalities could provide complementary services to the user community. A number of these types of resources are included in the ACCESS program. Their role is not to compete with the larger mainstream resources in providing more compute hours but to serve specialized user needs and/or provide test beds for future CI. Examples include FASTER from Texas A&M University with its composable GPUs (~370,000 ACEs reported), and the ARM-based Ookami HPC resource (~5.4 M ACEs reported) in the Institute for Advanced Computational Science at Stony Brook University. Purdue Anvil and SDSC Expanse have Kubernetes composable subsystems that support container-based scientific workflows. These are not yet fully integrated into

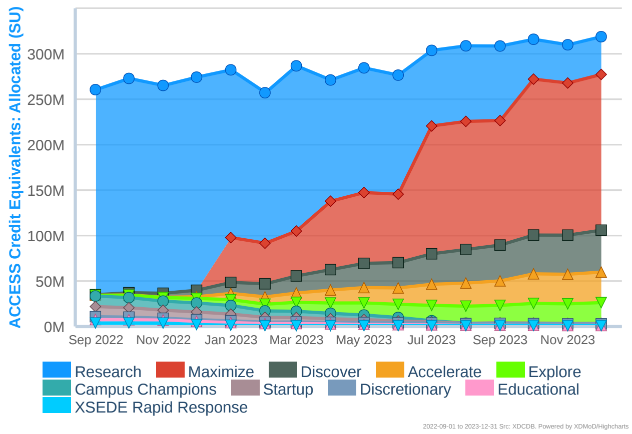

Figure 13: ACCESS Credit Equivalents allocated per month from the start of ACCESS September 2022 through December 2023.

ACCESS reporting. The SDSC Expanse Kubernetes cluster has only reported test jobs at the time of writing (~390 ACEs reported as of March 2024).

ACCESS supports allocations of distributed High Throughput Computing (dHTC) jobs on the OSG [5] Open Science Pool Resource. OSG also allows researchers to sign up directly for compute without interacting with ACCESS. Only a tiny fraction of the hundreds of millions of core hours on OSG Open Science Pool are via ACCESS allocations (OSG has reported ~53,000 ACEs consumed between September 2022 and December 2023).

4.5 Allocations

In XSEDE, the majority of service units were awarded via an allocations process that required PIs to write a detailed multi-page proposal that was reviewed by an allocations board (these allocations labelled *Research* in Figure 13). The board met quarterly to review proposals and award allocations. XSEDE also had *Startup* allocations to award small amounts of resources for researchers to get started and *Education* allocations for class instruction. In ACCESS, the allocation request procedure changed. Requests for large amounts of resource time (*Maximize* allocations) are now managed via bi-annual board reviews, and three other allocation types (*Discover*, *Accelerate*, and *Explore*) have a lightweight application process with smaller award amounts. Figure 13 shows the breakdown of active allocations each month by allocation type since ACCESS began. The plot shows the steady decrease in active allocations from XSEDE and a steady growth in the new ACCESS allocation types. There has been a steady, linear growth of *Explore* ~1.5 M ACE per month ($R^2 = 0.99$), *Accelerate* ~2.6 M ACE per month ($R^2 = 0.98$), and *Discover* ~3.3 M ACE per month ($R^2 = 0.99$) allocations. The majority of allocations are awarded via the *Maximize* review process.

4.6 Users and Usage Patterns

The number of active batch compute users, defined to be users who have submitted one or more batch compute jobs, has been

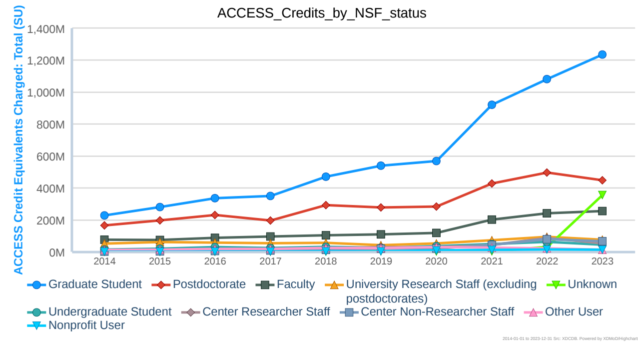

Figure 14: ACCESS Credit Equivalent usage grouped by academic status. Clearly graduate usage is increasing rapidly compared to all other categories.

increasing linearly over the past ten years, as seen in the blue trace in Figure 15. This plot is strictly batch compute users and does not include cloud usage. While it does include gateway usage, as described in Section 4.3, the gateway usage is a very small portion of the total usage. In contrast, the number of active PIs, defined as the principal investigator associated with submitted compute jobs, is flat, as seen in the red trace in Figure 15. The mean group size, defined to be the number of active users per active PI, is therefore also linearly increasing, as shown in Figure 16. We can dig down further into the details of this increase by plotting user academic status as a function of time over the same ten-year period, as shown in Figure 14. It seems the increase in active user group size is mainly due to graduate student compute usage, as there are only moderate increases in usage by postdoc and faculty numbers. Unfortunately, the information requested from new users during account registration changed between XSEDE and ACCESS; under XSEDE the academic status was a required field for all users, but in ACCESS this field has to date been optional. Data for user academic status has not been acquired, hence the rise in the "Unknown" status for the ACCESS period. The two outstanding features of this user analysis are the growth in mean group size and the growth of graduate usage. There are two possible explanations which both may contribute to these increases. First, computing research groups may simply be growing larger with more graduate students per group. Second, there may be a change in the way research computing is performed over the years with more graduate students and even undergraduate students submitting jobs relative to the professors doing the submissions themselves.

Figure 17 shows the number of PIs and active allocations over time from Jan 2014 through Dec 2023. In XSEDE a PI was only permitted one allocation of each type so the ratio of allocations to PIs was not larger than 1.1. ACCESS permits PIs to have multiple allocations of the same type. Figure 18 shows an uptick in the allocation PI ratio indicating that the number of PIs who are receiving multiple allocations is increasing.

The batch compute jobs submitted by each of the ~11.5K ACCESS users and the ACEs consumed by those jobs were acquired from the XDMoD data warehouse for the period of time from September 1, 2022, through December 31, 2023. A histogram of the number of

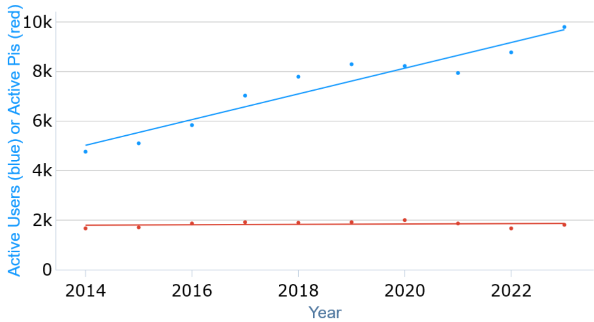

Figure 15: The number of active users (blue) and active PIs (red), that is, those users and PIs associated with submitted batch compute jobs. The lines are linear regressions. Note that while the number of users has increased over the ten year period, the number of PIs has remained constant.

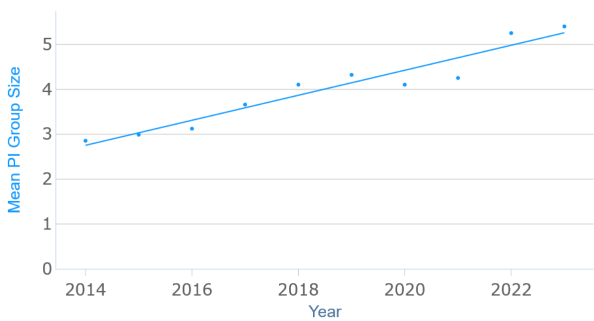

Figure 16: Since the number of active users has been increasing while the number of PIs has remained constant, the mean group size has increased substantially. The line is a linear regression.

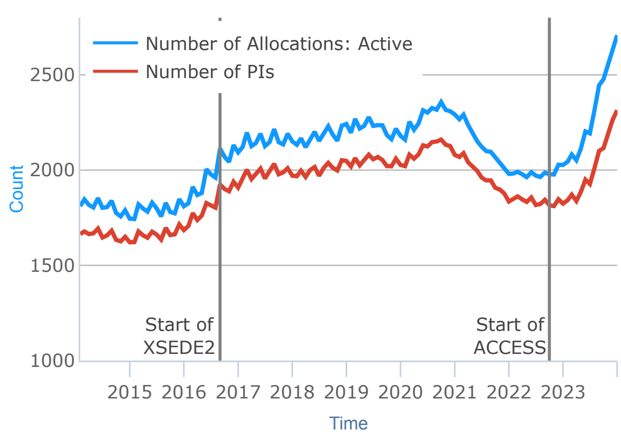

Figure 17: Number of active PIs and active allocations over time from Jan 2014 through Dec 2023. The two vertical lines indicate the time when the projects changed from XSEDE to XSEDE2 and from XSEDE2 to ACCESS.

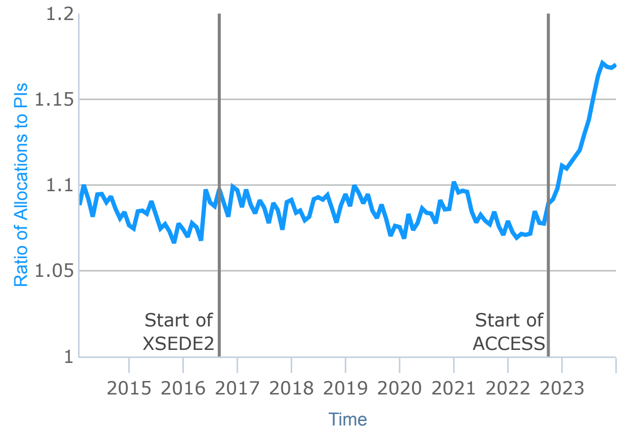

Figure 18: Ratio of number of allocations to PIs over time from Jan 2014 through Dec 2023. The two vertical lines indicate the time when the projects changed from XSEDE to XSEDE2 and from XSEDE2 to ACCESS.

jobs submitted by each user is shown in Figure 19. Most users have submitted fewer than 50 jobs with very few submitting more than a few hundred. A histogram of the number of jobs submitted under the auspices of each ACCESS PI, which gives a more project-centric view, is shown in Figure 20. The broader distribution reflects that multiple users are required to accomplish research project goals.

The job usage was ordered from largest to smallest, and a cumulative plot was made of user jobs submitted and ACEs consumed (Figure 21). Both traces show a similar trend in that a small majority of the users, approximately 10%, submit the large majority of the jobs, approximately 90%, and consume the large majority of the ACEs, approximately 90%. This is a familiar usage pattern to that of XSEDE users [6]. A more project-centric view of job submissions can be obtained by looking at the PI on each job rather than the user since, to a large extent, each PI is in charge of one project/allocation. Figure 22 shows the PI job submissions and the ACEs consumed by the jobs. The PI job submissions show a sharp rise with approximately 10% of the PIs accounting for 90% of the jobs submitted. The PI ACE curve rises a bit slower with a bit fewer than 80% of the ACEs consumed by 10% of the PIs. It should be emphasized that Figure 21 and Figure 22 are applicable only to the batch compute resource usage. They do not include cloud usage. More importantly, they do not include all of the gateway external users; they group each gateway's users into a single "Community User" for that gateway. We also did the analogous cumulative user analysis, filtering out the gateway usage. The curves change a little, becoming a bit less steep, but the main conclusion does not change. Looking at the cumulative curves for the number of jobs run, one would conclude that the usage is dominated by a small number of users running a large number of possibly small jobs. This is partly confirmed by the fact that the ACE curve does not rise as fast as the jobs curve since the larger number of small jobs is partly offset by the larger ACE consumption of the larger jobs. This conclusion is supported by a a cumulative usage analysis we did on individual resources for TACC Stampede2, which runs primarily larger jobs,

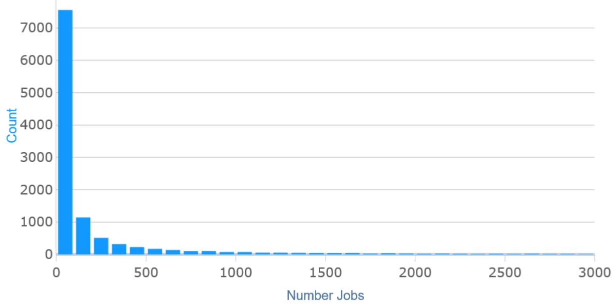

Figure 19: Histogram of the number of jobs submitted by ACCESS users. Most users have submitted fewer than 50 jobs with very few submitting more than a few hundred.

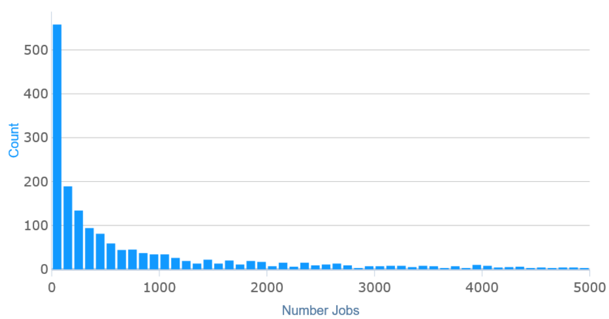

Figure 20: Histogram of the number of projects binned by the total number of jobs that were submitted for each project. The project histogram is substantially broader than the User histogram due to multiple users for each project.

and SDSC Expanse, which on average runs significantly smaller jobs. As expected, the SDSC Expanse slope, especially for the job analyses, was significantly steeper than the TACC Stampede2 slope.

ACCESS is a national-level NSF program, and the geographic distribution of the projects and PIs supported is therefore of interest. The ACE usage of PI institution locations was grouped by U.S. state and is displayed in Figure 23. Note that the usage by state is dominated by populous states such as California, Illinois, and New York. To remove the population factor, the per capita usage by state was also computed, see Figure 24. Although there are still disparities in the geographic distribution of usage, the per capita usage is more uniform and reflects real differences rather than simply higher state populations.

Figures 25 and 26 show the user count and ACEs for batch jobs by the Carnegie 2021 Basic Classification [3] of the institution of U.S. based users. The "N/A" category includes organizations such as government labs, private companies, and private institutions that do not have Carnegie Classifications. The majority of the usage and users are based in Doctoral Universities with high or very high research activity (R1 and R2).

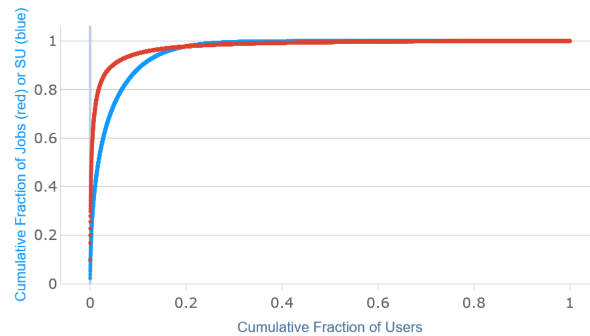

Figure 21: Cumulative plot of the fraction of jobs submitted (red) and ACEs consumed (blue) by ACCESS users. 0.9 of the ACEs are consumed by 0.1 of the users. For jobs even more than 0.9 are submitted by 0.1 of the users.

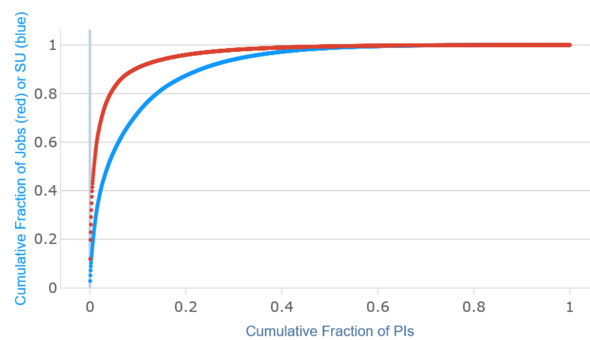

Figure 22: Cumulative plot of the fraction of jobs submitted (red) and ACEs consumed (blue) by ACCESS PIs. 90% of the jobs are submitted under the auspices of 10% of the PIs. For ACEs it is somewhat less steep with less than 80% of the ACEs consumed by the top 10% active PIs.

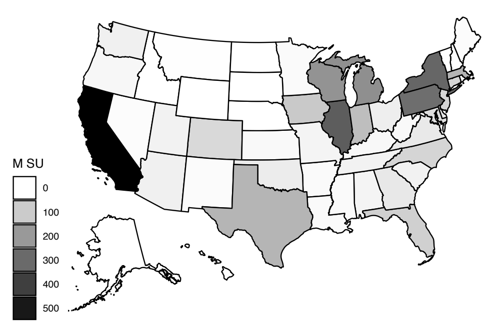

Figure 23: The usage of PIs is grouped by the location of their home institute and displayed on the U.S. map in millions of ACEs.

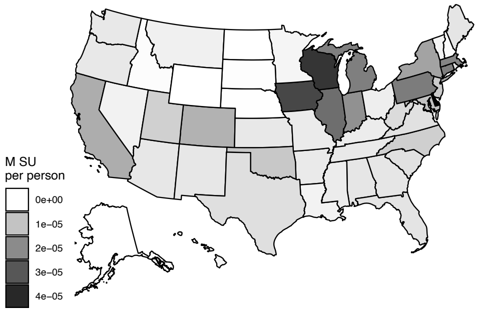

Figure 24: The usage of PIs is grouped by the location of their home institute. The usage is then divided by state population, that is, the per capita usage is displayed on the U.S. map in millions of ACEs.

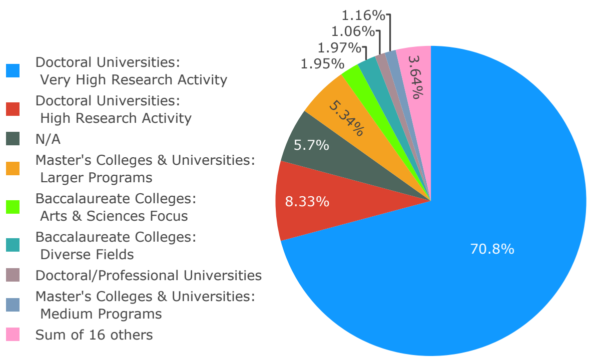

Figure 25: Pie chart of the number of active users grouped by the Carnegie 2021 Basic Classification of their Institution from January 2014 through December 2023.

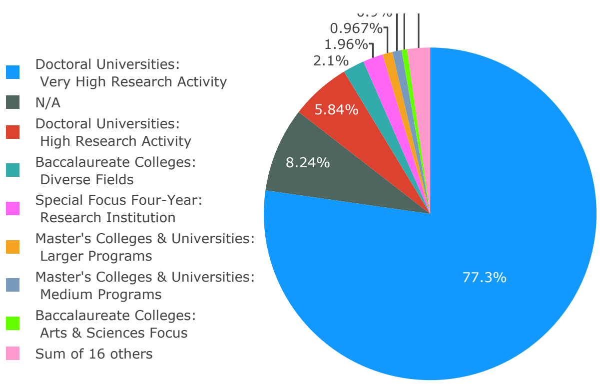

Figure 26: Pie chart of the amount of ACCESS Credit Equivalents grouped by the Carnegie 2021 Basic Classification of the Institution of the user who ran the jobs from January 2014 through December 2023.

5 CONCLUSIONS

ACCESS is still a relatively new program; the current analysis covers only the first 16 months of operation. In many ways, the usage shows a linear continuation of the XSEDE program. Many of the ten-year historical plots, such as the simple increase in ACEs, show linear increases as may be expected from improving computational hardware. Of particular interest is the increased usage of GPUs coupled with increased Python usage as an indicator of the growing importance of AI projects. Cloud usage is seen to be increasingly supportive of science gateways. Cloud usage is more balanced among the various allocation types as opposed to batch compute usage which is dominated by the two largest allocation types.

This study also focused on users and usage patterns. Although the number of PIs and projects has remained steady, the size of PI groups has increased mainly due to the increased usage of ACCESS computational resources by graduate and undergraduate students. Other usage patterns have remained similar to those of the XSEDE program. A relatively small fraction of the users submit most of the jobs and consume most of the ACEs. The geographical usage pattern of ACCESS usage is similar to that of XSEDE. The usage is primarily dominated by R1 institutions as classified by the Carnegie Classification.

6 FUTURE WORK

The current study provides an initial baseline to judge future progress on how well ACCESS serves the community. This study must be regarded as preliminary both from the point of view of the short time span of the program and the broad but shallow scope of the study. In addition to the desirability of doing a future study further into the program, almost every aspect of this study could easily be expanded into a full length study. In addition, other aspects of research workflows should be examined, such as data storage, the NSF Awards database, and data transfer/networking data now being acquired from NetSage.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. (2137603).

REFERENCES

[1] Timothy J. Boerner, Stephen Deems, Thomas R. Furlani, Shelley L. Knuth, and John Towns. 2023. ACCESS: Advancing Innovation: NSF's Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support. In *Practice and Experience in Advanced Research Computing* (Portland, OR, USA) *(PEARC '23)*. Association for Computing Machinery, New York, NY, USA, 173–176. https://doi.org/10.1145/3569951.3597559

[2] Charlie Catlett, William E Allcock, Phil Andrews, Ruth Aydt, Ray Bair, Natasha Balac, Bryan Banister, Trish Barker, Mark Bartelt, Pete Beckman, et al. 2008. *Teragrid: Analysis of organization, system architecture, and middleware enabling new types of applications.* Technical Report. IOS press.

[3] Indiana University Center for Postsecondary Research. [n. d.]. The Carnegie Classification of Institutions of Higher Education. https://carnegieclassifications.acenet.edu/carnegie-classification/ 2021 edition.

[4] Jeffrey T. Palmer, Steven M. Gallo, Thomas R. Furlani, Matthew D. Jones, Robert L. DeLeon, Joseph P. White, Nikolay Simakov, Abani K. Patra, Jeanette M. Sperhac, Thomas Yearke, Ryan Rathsam, Martins Innus, Cynthia D. Cornelius, James C. Browne, William L. Barth, and Richard T. Evans. 2015. Open XDMoD: A tool for the comprehensive management of high-performance computing resources. *Computing in Science and Engineering* 17, 4 (2015), 52–62. https://doi.org/10.1109/MCSE.2015.68

[5] Ruth Pordes, Don Petravick, Bill Kramer, Doug Olson, Miron Livny, Alain Roy, Paul Avery, Kent Blackburn, Torre Wenaus, Frank Würthwein, Ian Foster, Rob Gardner, Mike Wilde, Alan Blatecky, John McGee, and Rob Quick. 2007. The open science grid. *Journal of Physics: Conference Series* 78 (2007), 012057. https://doi.org/10.1088/1742-6596/78/1/012057
[6] Nikolay A. Simakov, Joseph P. White, Robert L. DeLeon, Steven M. Gallo, Matthew D. Jones, Jeffrey T. Palmer, Benjamin Plessinger, and Thomas R. Furlani. 2018. A Workload Analysis of NSF's Innovative HPC Resources Using XDMoD. arXiv:1801.04306 [cs.DC]
[7] The U.S. National Science Foundation. 2023. Campus Cyberinfrastructure (CC*). https://new.nsf.gov/funding/opportunities/campus-cyberinfrastructure-cc/. Accessed: 2023-09-20.
[8] The U.S. National Science Foundation. 2023. Cyberinfrastructure for Sustained Scientific Innovation (CSSI). https://new.nsf.gov/funding/opportunities/cyberinfrastructure-sustained-scientific. Accessed: 2023-09-20.
[9] John Towns, Timothy Cockerill, Maytal Dahan, Ian Foster, Kelly Gaither, Andrew Grimshaw, Victor Hazlewood, Scott Lathrop, Dave Lifka, Gregory D. Peterson, Ralph Roskies, J. Ray Scott, and Nancy Wilkins-Diehr. 2014. XSEDE: Accelerating Scientific Discovery. *Computing in Science & Engineering* 16, 5 (2014), 62–74. https://doi.org/10.1109/MCSE.2014.80
[10] Aaron Weeden, Joseph P. White, Robert L. DeLeon, Ryan Rathsam, Nikolay A. Simakov, Conner Saeli, and Thomas R. Furlani. 2024. The Data Analytics Framework for XDMoD. *SN Computer Science* 5, 462 (2024). https://doi.org/10.1007/s42979-024-02789-2