

# The Multi-Tier Assistance, Training, and Computational Help (MATCH) project, a Track 2 NSF ACCESS Initiative

Shelley L. Knuth

University of Colorado Boulder  
Boulder, CO  
shelley.knuth@colorado.edu

Julie Ma

Massachusetts Green High  
Performance Computing Center  
(MGHPCC)  
Holyoke, MA  
jma@mghpcc.org

Joel Adams

University of Kentucky  
Address, Address  
joel.adams@uky.edu

Alan Chalker

Ohio Supercomputer Center  
Columbus, OH  
alanc@osc.edu

Ewa Deelman

University of Southern California  
Marina del Rey, CA  
deelman@isi.edu

Layla Freeborn

University of Colorado Boulder  
Boulder, CO  
Layla.Freeborn@colorado.edu

Vikram Gazula

University of Kentucky  
Lexington, KY  
gazula@uky.edu

John Goodhue

Massachusetts Green High  
Performance Computing Center  
(MGHPCC)  
Holyoke, MA  
jtgoodhue@mghpcc.org

James Griffioen

University of Kentucky  
Lexington, KY  
griff@netlab.uky.edu

David Hudak

Ohio Supercomputer Center  
Columbus, OH  
dhudak@osc.edu

Andrew Pasquale

Massachusetts Green High  
Performance Computing Center  
(MGHPCC)  
Holyoke, MA  
andrew@elytra.net

Dylan Perkins

University of Colorado Boulder  
Boulder, CO  
dylan.perkins@colorado.edu

Alana Romanella

University of Colorado Boulder  
Boulder, CO  
alana.romanella@colorado.edu

Mats Rynge

University of Southern California  
Marina del Rey, CA  
rynge@isi.edu

## ABSTRACT

NSF-supported cyberinfrastructure (CI) has been highly successful in advancing science and engineering over the last few decades. During that time, there have been significant changes in the size and composition of the participating community, the architecture and capacity of compute, storage, and networking platforms, and the methods by which researchers and CI professionals communicate. These changes require rethinking the role of research support services and how they are delivered. To address these changes and support an expanding community, MATCH is implementing a model for research support services in ACCESS that comprises

three major themes: 1) leverage modern information delivery systems and simplify user interfaces to provide cost-effective, scalable, support to a broader community of researchers, 2) engage experts from the community to develop training materials and instructions that can dramatically reduce the learning curve, and 3) employ a matchmaking service that will maintain a database of specialist mentors and student mentees that can be matched with projects to provide the domain-specific expertise needed to leverage ACCESS resources. A new ACCESS MATCH Portal (AMP) will serve as the single front door for researchers to obtain guided support and assistance. AMP will leverage emerging curated and tagged taxonomies to identify and match inquiries with knowledge base content and expertise. Expert-monitored question and answer platforms will be created to ensure researcher questions are accurately answered and addressed in a timely fashion, and easy-to-use interfaces such as Open OnDemand and Pegasus will be enhanced to simplify CI use and provide context-aware directed help. The result will be a multi-level support infrastructure capable of scaling to serve a growing research community with increasingly specialized support

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Copyright ©JOCSE, a supported publication of the Shodor Education Foundation Inc.

needs, resulting in research discoveries previously hindered by researchers' inability to effectively utilize NSF CI resources. This talk will cover the components of the MATCH project and discuss how MATCH will engage and work with the ACCESS community.

## KEYWORDS

ACCESS, research support, knowledge base

## 1 INTRODUCTION

As science and engineering researchers continue to push boundaries in their respective fields, consumption of cyberinfrastructure (CI) in the form of compute, storage, networking, and software becomes increasingly important. Ensuring the availability of state-of-the-art CI to continue these successful research endeavors, the National Science Foundation (NSF) announced a new vision for access to NSF funded CI through the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) project. The goal of this project is to provide democratized utilization of CI, allowing for equitable use of these powerful systems. This vision builds from the success of the eXtreme Science and Engineering Discovery Environment (XSEDE) [3], which provided a robust ecosystem of resources and support to enable research discovery. The ACCESS project as outlined consists of five tracks focused on allocations, end user support services, operations and integration services, monitoring and measurement services, and technology translation services, tied together through an ACCESS Coordination Office.

Available hardware to enable research, while important, is only one component of a successful CI ecosystem. Without adequate support to understand how to use CI efficiently and appropriately, researchers can face challenges in successful research discovery. The Multi-Tier Assistance, Training, and Computational Help (MATCH) project is designed to support the end users of these systems, providing equitable, scalable support to best enable research on NSF funded CI. The goal of the MATCH project is to decrease the time to successful solution for computational research problems, enabling researchers to focus less on complex workflows and more on discovery, while also providing confidence in their results.

## 2 TIERED SUPPORT MODEL

To accomplish this goal, the MATCH project focuses on a tiered support model (Figure 1). The object of this model is to successfully decrease the time to solution while still maintain high quality service. The base of the pyramid represents maximum usage of CI by users while providing minimal human interaction, while the apex of the pyramid represents a small number of users who likely have special cases requiring more consultation. By providing a mixture of support solutions tailored to different researchers' needs, the MATCH project can help researchers resolve workloads quickly, efficiently, and successfully. It is important to recognize that, in practice, researchers will not step actually through the tiers; rather, the MATCH project has designed a variety of resources, outlined below, to help them accomplish their goals.

### 2.1 Tier 1 - Existing Tools

The Tier 1 support model, shown at the bottom of the pyramid in Figure 1, leverages existing tools in the high performance computing and data communities that better enable research through ease of use. The highly successful Open On-Demand (OOD) tool [2], developed at The Ohio Supercomputer Center, supports computational researchers by providing an interface that allows for the efficient use of remote computing resources. The OOD interface removes the underlying complexity of CI by providing a web interface to enable efficient file management, command line shell access, and job management and monitoring across a variety of servers and resource managers. Templates can be utilized to submit jobs to a scheduler, transfer data, and monitor services.

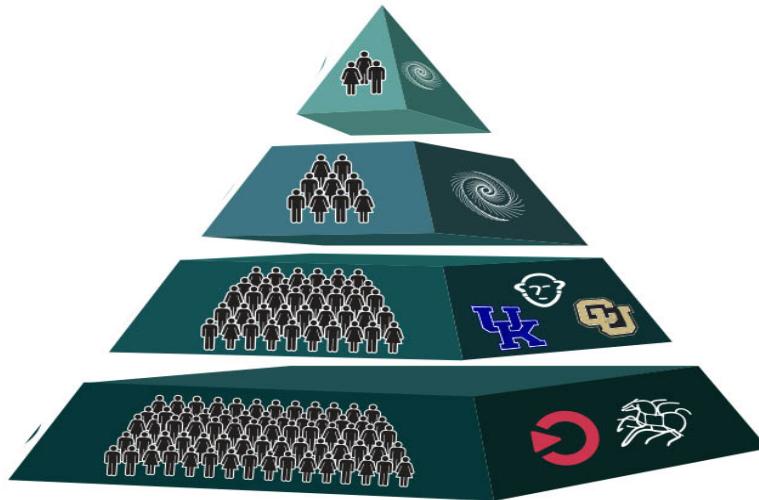
The MATCH project will integrate the OOD tool with the Pegasus workflow tool [1]. The Pegasus project simplifies complex data workflows by helping applications execute in a variety of environments. An application can be easily mapped onto available compute resources and execute a multi-step process in an appropriate order, managing workflows with millions of computational tasks. The Pegasus tool also eases complex workflows for end users through reproducibility, ensuring data integrity, and enabling provenance.

The OOD and Pegasus tools will be enhanced, improved, and integrated to further remove the complexity of using CI systems. Beginning in September 2022, three pilot projects involving three different institutions will inform the future of this effort.

The Pegasus and OOD tools will also be integrated with the ACCESS Support Portal (ASP), so that a user can login to one place to use all of these tools. The ASP has underpinnings in the Connect.CI portal. This tool provides a window into the ACCESS service, allowing users to monitor and apply for allocations, conduct account management activities, submit tickets, and join affinity groups. By providing a seamless, singular tool to utilize a variety of support and account information, researchers are able to work toward their goals quicker.

### 2.2 Tier 2 - Knowledge Base

The goal for Tier 2 of the MATCH project is to assist researchers with quickly finding answers to their computational problems by providing a singular location to source information. In the MATCH knowledge base, a researcher may search for a question and be provided with a number of answers from a variety of MATCH approved sources. These sources may include documentation or training written by ACCESS staff members, information resourced from existing HPC centers across the globe, documentation, code templates, trainings, or other materials provided by our community grants program (described below), or answers to questions that have previously been posed in a community forum (<https://ask.cyberinfrastructure.org/>). The MATCH program will work with members of the Computational Science Support Network (CSSN) to identify gaps in computational resources intended to support researchers, and create or resource new or existing materials to fulfill that gap. The MATCH effort will be able to provide financial assistance through a community grants program to members of the CSSN to assist with supporting filling these gaps in resources. All information contained in the ACCESS knowledge base can be found here: <https://support.access-ci.org/knowledge-base>. Membership in



**Figure 1: Tiered Support Model in MATCH**

the CSSN is as simple as filling out a short form indicating certain attributes about the individual.

The incorporation and development of new and existing affinity groups will play a strong role in this effort. Affinity groups are formed with community members who have a common, shared interest in a computational issue, scientific or engineering research endeavor, diversity and inclusion effort, or any other connection point that will appropriately serve the ACCESS community. The MATCH team will provide each affinity group with access to specific Slack channels, question and answer forums, news and outage alerts, events and trainings, and other relevant knowledge base resources. The affinity groups structure will play a significant role for all of the ACCESS teams to connect with a variety of community members, including researchers as well as those needing allocations support or ACCESS Resource Providers.

### 2.3 Tier 3 - MATCH Plus

The Tier 3 MATCH Plus engagement supports researchers in need of short term assistance to resolve a computational barrier that is slowing the time to resolution. Often times, these computational barriers, such as transitioning to a new resource or optimizing code, may resolve with short term assistance from a student. Researchers will be able to request a MATCH Plus engagement where a short form will ask for a description of the problem and any tags to help identify students with relevant skill sets who can assist. Students will also be matched with a mentor who will provide guidance to the student for resolution of the issue. The MATCH Plus engagements are intended to be short term commitments (on the order of six months or less). Following the successful similar Northeast Cyberteam student engagement project, organized by several members of the MATCH team, MATCH will facilitate regular touch points between the students and researchers, students and mentors, and among the students, ensuring completion of the project. This

model is intended to provide extra support to quickly facilitate completion of computational research projects, as well as provide professional development experience to students across a variety of disciplines. Students, researchers, and mentors who are part of the CSSN will be able to participate.

### 2.4 Tier 4 - MATCH Premier

Similarly to MATCH Plus, MATCH Premier provides support to researchers needing short term engagement to resolve a computational barrier being faced when conducting a research project. The primary difference between a Tier 3 and Tier 4 engagement is that the Tier 4 engagement is suitable for computational issues that may be more complex and require the assistance of a staff consultant rather than a student. Similarly to the XSEDE Extended Collaborative Support Services (ECSS) program, consultants are available to support removing these barriers. However, at this time, the MATCH program will only be able to provide the connection points via consultants who are members of the CSSN rather than the full suite of services the ECSS program provided. MATCH is working with researchers to identify consultants that can be funded through proposals being written outside of the ACCESS project. The MATCH team will provide suitable collaborators for the research projects via the same tag taxonomy that is part of the MATCH Plus engagement, and similarly to Tier 3, will facilitate regular connection points between the team members.

## 3 CONCLUSIONS

The goal of the MATCH effort is to assist the science and engineering research community with the appropriate computational resources, tools, and support to reduce the time to solution for research and allow researchers to focus on their primary interests - solving research questions, securing future grants, attracting quality faculty and students to their institution, publishing papers, and

graduating students. The MATCH approach relies on simplifying what can be a complex process for researchers who use large-scale computational resources as only one component of their research ecosystem. The approach described relies on best practices from existing support models but also moves the community into the next phase of supporting cyberinfrastructure. With a primary goal of providing a variety of groups with equitable and scalable support, which includes, but is not limited to, "traditional" HPC workflows, high throughput computing workflows, under-represented groups, or those in the social sciences, this tiered model approach can create connection points across the community that will best support and enable cutting edge research nationally. The MATCH team looks forward to evolving and growing with the needs of the community through the length of this project.

## ACKNOWLEDGMENTS

This project is supported by the National Science Foundation under OAC-2138286. The MATCH Team would like to acknowledge the work of the other ACCESS projects, with whom this project is in close collaboration.

## REFERENCES

- [1] Ewa Deelman, Karan Vahi, Gideon Juve, Mats Rynge, Scott Callaghan, Philip J. Maechling, Rajiv Mayani, Weiwei Chen, Rafael Ferreira Da Silva, Miron Livny, et al. 2015. Pegasus, a workflow management system for science automation. *Future Generation Computer Systems* 46 (2015), 17–35.
- [2] Dave Hudak, Doug Johnson, Alan Chalker, Jeremy Nicklas, Eric Franz, Trey Dockendorf, and Brian L. McMichael. 2018. Open OnDemand: A web-based client portal for HPC centers. *The Journal of Open Source Software* 3, 27 (2018), 622.
- [3] 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 and Engineering* 16, 5 (Sept. - Oct 2014), 62–74.