

**A** and Education   **L** Applications and Software   **B** BoF   **D** Diversity   **P** Panels

**S** Systems and System Software   **T** Training   **R** Tutorial   **W** Workforce Development   **K** Workshop

## JULY 19 • MONDAY

8:00am – 11:00am   **R** **TUTORIAL: An Introduction to Advanced Features in MPI**   Pathable Platform

*Speakers: Amit Ruhela, Victor Eijkhout*

The MPI library is now in version 3, with version 4 forthcoming, but most programmers use mechanisms from MPI-1 or 2 at best.

This half-day tutorial, aimed at current MPI programmers, will discuss a number of MPI-3 features that offer more flexibility, a more elegant expression of algorithms, and higher performance.

There will be lab sessions to exercise the material.

This tutorial was taught at PEARC20 where it was attended by 50 people, receiving a 4.3/5 rating.

New for this year will be a brief discussion of some features of the upcoming MPI-4 standard.

8:00am – 11:00am   **R** **TUTORIAL: Deploying Science Gateways with Apache Airavata**   Pathable Platform

*Speakers: Suresh Marru, Marlon Pierce, Eroma Abeysinghe, Marcus Christie, Dimuthu Wannipurage, Sudhakar Pamidighantam*

The authors present the Apache Airavata framework for deploying science gateways, illustrating how to request, administer, modify, and extend a basic science gateway tenant to hosted Apache Airavata middleware. We further show how to use this gateway to execute scientific software on XSEDE supercomputers and university-operated computing clusters. This tutorial builds on successful tutorials that the authors have presented at previous PEARC and Gateways conferences, including a completely online version at Gateways 2020; see <https://www.youtube.com/watch?v=FhAHkOoVGh4> for a recording.

8:00am – 11:00am   **R** **TUTORIAL: Empowering Research Computing at Your Organization Through the Open Science Grid**   Pathable Platform

*Speakers: Christina Koch, Lauren Michael, Mats Rynge, Emelie Fuchs*

This training will provide researcher-facing cyberinfrastructure professionals with the information and hands-on skills they need to engage with the Open Science Grid. Attendees will explore the capabilities of the OSG through hands-on activities, discuss strategies for engaging researchers, and hear from organizations who have already partnered with the OSG. Attendees should leave this training with a clear understanding of where the OSG can be transformative for research, the process for moving work onto the OSG, and what options exist for engaging further with the OSG.

8:00am – 11:00am   **R** **TUTORIAL: Security Log Analysis: Real world hands on methods and techniques to detect attacks**   Pathable Platform

*Speakers: Mark Krenz, Ishan Abhinit*

The goal of security log analysis is to more efficiently leverage log collection in order to identify threats and anomalies in your research organization. This half-day training will help you tie together various log and data sources to provide a more rounded, coherent picture of a potential security event. It will also help you understand log analysis as a life cycle (collection, event management, analysis, response) that continues to become more efficient over time. Interactive demonstrations will cover both automated and manual analysis using multiple log sources, with examples from real security incidents. 45% of the sessions will be devoted to hands-on exercises where students will analyze real log files in order to find security incidents. A knowledge of Unix commands such as grep, awk and wc are ideal for this class, but not required as the algorithmic methods can be applied to other systems. A brief primer on these commands will be provided. We have expanded our exercise this time to include both command line and Elastic Stack based analysis. This will be an interactive session allowing Q&A

and also will feature interactive polls to enhance the audience's learning experience.

- 
- 8:00am – 11:00am R **TUTORIAL: Visualize, Analyze, and Correlate Networking Activities for Parallel Programs on InfiniBand HPC Clusters using the OSU INAM Tool** Pathable Platform  
*Speakers: Dhableswar Panda, Hari Subramoni, Pouya Kousha*  
As the computing, networking, heterogeneous hardware, and storage technologies continue to evolve in HEC platforms, Understanding the full-stack performance tradeoffs and interplay between HPC applications, MPI libraries, the communication fabric, the file system, and the job scheduler becomes more challenging endeavor. Such understanding will enable all involved parties to understand the bottlenecks to maximize the efficiency and performance of the individual components that comprise a modern HPC system and solve different grand challenge problems. Through this tutorial, the participants will learn how to use the OSU InfiniBand Network Analysis and Monitoring (INAM) tool in conjunction with live jobs running on various remote clusters at OSC and OSU to visualize, analyze, and correlate how the MPI runtime, high-performance network, I/O filesystem, and job scheduler interact and identify potential bottlenecks online. Emphasis is placed on how tools are used in combination for identifying performance problems and investigating optimization alternatives. We will request remote access to the Pitzer system at OSC and the RI/RI2 clusters at OSU for hands-on exercises. This will help to prepare participants to locate and diagnose performance bottlenecks in their own clusters and parallel programs.
- 
- 8:00am – 11:00am K **EXHIBITOR WORKSHOP: Dell - Win More Grants** Pathable Platform  
*Speakers: Kenton McHenry*  
Find out how you can win more grants! Get tips and strategies from Dell Technologies, AMD, Grants Office LLC and your PEARC peers. There will be short, impactful presentations, panel discussions and audience Q&A throughout the workshop. Also hear from the Dell Technologies HPC & AI Innovation Lab and from the Chief Technology Office regarding their research and forward-looking revelations on new technologies.
- 
- 8:00am – 11:00am K **EXHIBITOR WORKSHOP: Intel Parallel Studio XE has become Intel oneAPI Toolkits How and Why!** Pathable Platform  
*Speakers: James Reinders, Chris Allison*  
In this workshop, PEARC21 conference attendees will learn how to take advantage of oneAPI for modern cross-architecture application development, analysis, and tuning. Attendees will learn Intel's oneAPI Toolkits are a useful upgrade to the popular and widely used Parallel Studio XE tools, to embrace the vision of oneAPI, making heterogeneous programming open and ubiquitous. The session will educate participants on current resources, oneAPI Toolkits, and development vehicles available to the PEARC21 community. Join us to learn more! NOTE: This session is ideal for research computing system administrators, application specialists, application developers, and other HPC and AI software stakeholders. Speakers: James Reinders, co-author of the recent book on Data Parallel C++ (and nine other HPC focused books), will be joined by other Intel engineers and experts to teach this workshop and answer questions from participants.
- 
- 8:00am – 11:00am K **EXHIBITOR WORKSHOP: NVIDIA - Best Practices for Operating a GPU system** Pathable Platform  
*Speakers: Gianluca Castellani, Oded Green, Jeff Layton, Scott McMillian, Bradley Palmer, Kaleb Smith, Craig Tierney*  
Hear about all the tools and techniques we use at NVIDIA to operate our multiple research computing clusters. We will demonstrate how you can use these same tools and techniques on your systems to help solve your users' research computing challenges. We will describe the toolsets that system administrators can use to improve utilization of GPUs and GPU nodes, including some live demonstrations. We will also review how research computing support staff can help their users make more efficient and effective use of GPU acceleration.
- 
- 8:00am – 11:00am K **WORKSHOP: ACM SIGHPC SYSPROS Symposium 2021** Pathable Platform  
*Speakers: Matthew Cawood, John Blaas*  
In order to meet the demands of high-performance computing (HPC) researchers, large-scale computational and storage machines require many staff members who design, install, and maintain these systems. These HPC systems professionals include system engineers, system administrators, network administrators, storage administrators and operations staff who face problems that are unique to high performance computing systems. While many conferences exist for the HPC field and the system administration field, none exist that focus on the needs of HPC systems professionals. Support resources can be difficult to find to help with the issues encountered in this specialized field. Often systems staff turn to the community as a support resource and opportunities to strengthen and grow those relationships are highly beneficial.

This Workshop is designed to share solutions to common problems, provide a platform to discuss upcoming technologies, and to present state of the practice techniques so that HPC centers will get a better return on their investment, increase performance and reliability of systems, and researchers will be more productive. Additionally, this Workshop is affiliated with the systems professionals' chapter of the ACM SIGHPC (SIGHPC SYSPROS Virtual ACM Chapter). This session would serve as an opportunity for chapter members to meet face-to-face, discuss the chapter's yearly workshop held at SC, and continue building our community's shared knowledge base.

8:00am –  
11:00am

**K WORKSHOP: Building a Strategic Plan for your Research Computing and Data Program**

*Speakers: Patrick Schmitz, Claire Mizumoto, Dana Brunson, Doug Jennewein, Scotty Strachan* Pathable Platform

This workshop will gather Research Computing and Data (RCD) professionals to learn leading practices for developing effective strategic plans for their Research Computing and Data programs. The workshop is open to RCD professionals who are familiar with issues around supporting Research Computing and Data, have experience contributing to strategic planning, and have some exposure to the RCD Capabilities Model. Attendees will hear the experiences of universities (Arizona State University, University of Nevada, Reno, University of Hawaii, plus one more) who are currently using the RCD Capabilities Model as part of their strategic planning work, including lessons learned. Attendees will discuss the range of RCD strategic planning models across the community, and approaches to building a strong strategic planning practice. Finally, participants will define requirements for a new effort to develop a shared community resource to support strategic planning for RCD, identifying potential elements of such a resource and a near-term roadmap for development. Workshop organizers will document the findings of the workshop in a report shared with the community.

8:00am –  
11:00am

**K WORKSHOP: Fifth Workshop on Trustworthy Scientific Cyberinfrastructure (TrustedCI@PEARC21)**

Pathable Platform

*Speakers: Jim Basney*

The Fifth Workshop on Trustworthy Scientific Cyberinfrastructure (TrustedCI@PEARC21) provides an opportunity for sharing experiences, recommendations, and solutions for addressing cybersecurity challenges in research computing. The half-day workshop provides a forum for information sharing and discussion among a broad range of attendees, including cyberinfrastructure operators, developers, and users.

Implementing cybersecurity for open science across the diversity of scientific research projects presents a significant challenge. There is no one-size-fits-all approach to cybersecurity for open science that the research community can adopt. Even NSF Major Facilities, the largest of the NSF projects, struggle to develop effective cybersecurity programs. To address this challenge, practical approaches are needed to manage risks while providing both flexibility for project-specific adaptations and access to the necessary knowledge and human resources for implementation. This workshop brings community members together to further develop a cybersecurity ecosystem, formed of people, practical knowledge, processes, and cyberinfrastructure, that enables research projects to both manage cybersecurity risks and produce trustworthy science.

8:00am –  
3:00pm

**R TUTORIAL: A Deep Dive into Constructing Containers for Scientific Computing and Gateways**

*Speakers: Eric Coulter, Peter Vaillancourt, Jeremy Fischer, Stephen Bird, Sanjana Sudarshan, Suresh Marru* Pathable Platform

In recent years, using containers has been rapidly gaining traction as a solution to lower the barriers to using more software on HPC and cloud resources. However, significant barriers still exist to actually doing this in practice, particularly for well-established community codes which expect to run on a particular operating system version or resource. Additional barriers exist for researchers unfamiliar with containerization technologies. While many beginner tutorials are available for building containers, they often stop short of covering the complexities that can arise when containerizing scientific computing software. The goal of this full-day tutorial is to demonstrate and work through building and running non-trivial containers with users. We will containerize community scientific software, exhibit how to share with a larger community via a container registry, and then run on a completely separate HPC resource, with and without the use of a Science Gateway. The subject matter will be approachable for intermediate to advanced users, and is expected to be of interest to a diverse audience including researchers, support staff, and teams building science gateways.

8:00am –  
3:00pm

**R TUTORIAL: Build a quick, effective coding tutorial**

Pathable Platform

*Speakers: Christina Maimone, Colby Witherup Wood*

Research Computing and Data professionals are often called upon to train others in technical skills, but creating

effective workshop and training materials is time-consuming, especially when instructors need to adapt to both virtual and in-person formats. This tutorial will introduce participants to a template for creating a successful, one-hour, interactive coding workshop that can easily transition between remote and in-person instruction. Through a combination of individual work sessions, peer feedback, and instructor advice, participants will refine their topic, outline their examples, and write effective exercises. They will also write a description of their workshop, think through the technology needed, and share tips and ideas with others. In addition to ending the day with materials for their new workshop, participants will leave with a quick template for building future effective workshops.

---

8:00am – 3:00pm R **TUTORIAL: Lucata Pathfinder-S Tutorial: Next-generation Computation with the Rogues Gallery** Pathable Platform  
*Speakers: Jeffrey Young, Semir Sarajlic, William Powell, Janice McMahon, Jason Riedy*

The Rogues Gallery is a new experimental testbed hosted at Georgia Tech that is focused on tackling "rogue" architectures for the post-Moore era of computing including those in areas like high-performance, near-memory, neuromorphic, and quantum computing. More recently, the Rogues Gallery has been awarded an NSF grant to serve as a novel architecture testbed as part of the CISE Community Research Infrastructure (CCRI) program. This tutorial will provide an introduction to this new community resource and will focus on hands on development with the signature architecture, Lucata's newly designed Pathfinder-S.

This virtually hosted tutorial will present a brief overview of the Rogues Gallery testbed and how NSF researchers can access and utilize unique hardware in the neuromorphic, smart networking, HPC, and near memory spaces to carry out related research goals for CISE-oriented research. Attendees will have an opportunity to learn about and program for the Lucata Pathfinder-S system, a near-memory computing architecture for sparse applications that has applications for database operations, graph analytics, and machine-learning related techniques. We will provide and work through a set of demonstration codes and will provide details on potential workflows that users can explore after the tutorial. Attendees will have an opportunity to continue their investigation into using the Pathfinder-S by requesting a free account to access the Rogues Gallery at the end of the tutorial.

---

8:00am – 3:00pm R **TUTORIAL: Managing HPC Software Complexity with Spack** Pathable Platform  
*Speakers: Gregory Becker, Robert Blake, Massimiliano Culpo, Tammy Dahlgren, Adam Stewart, Todd Gamblin*

The modern scientific software stack includes thousands of packages, from C, C++, and Fortran libraries, to packages written in interpreted languages like Python and R. HPC applications may depend on hundreds of packages spanning all of these ecosystems. To achieve high performance, they must also leverage low-level and difficult-to-build libraries such as MPI, BLAS, and LAPACK. Integrating this stack is extremely challenging. The complexity can be an obstacle to deployment at HPC sites and deters developers from building on each other's work.

Spack is an open source tool for HPC package management that simplifies building, installing, customizing, and sharing HPC software stacks. In the past few years, its adoption has grown rapidly: by end-users, by HPC developers, and by the world's largest HPC centers. Spack provides a powerful and flexible dependency model, a simple Python syntax for writing package build recipes, and a repository of over 5,000 community-maintained packages. This tutorial provides a thorough introduction to Spack's capabilities: installing and authoring packages, integrating Spack with development workflows, and using Spack for deployment at HPC facilities. Attendees will learn foundational skills for automating day-to-day tasks, as well as deeper knowledge of Spack for advanced use cases.

---

8:00am – 3:00pm R **TUTORIAL: Open OnDemand, Open XDMoD, and ColdFront: an HPC center management toolset** Pathable Platform  
*Speakers: Dori Sajdak, Andrew Bruno, Eric Franz, Jeff Ohrstrom, Ryan Rathsam*

The University at Buffalo Center for Computational Research (UB CCR) and Ohio Supercomputer Center (OSC) team up to offer HPC systems personnel a step-by-step tutorial for installing, configuring and using what many centers now consider vital software products for managing and enabling access to their resources. UB CCR offers two open source products - an allocations management system, ColdFront, and an HPC metrics & data analytics tool, Open XDMoD. OSC provides the open source OnDemand portal for easy, seamless web-based access for users to HPC resources. These three products have been designed to work together to provide a full package of HPC center management and access tools. In this tutorial the system administrators and software developers from OSC and UB CCR will walk attendees through the installation and configuration of each of these software packages. We'll show how to use these three products in conjunction with each other and the Slurm job scheduler.

We will begin the tutorial with a short overview of each software product and how they tie together to provide

seamless management of an HPC center. We'll spend the first part of the tutorial demoing the installation and configuration of ColdFront and Open XDMoD. The second half will be spent on the installation of Open OnDemand and examples of configuring interactive apps. We'll end with instructions on how to tie together Open XDMoD with Open OnDemand for access to job metrics within OnDemand.

---

8:00am – 3:00pm R **TUTORIAL: Programming and Profiling Modern Multicore Processors** Pathable Platform  
*Speakers: Amit Ruhela, Lars Koesterke, John Cazes, Matthew Cawood*

Modern processors, such as Intel's Xeon Scalable line, AMD's EPYC architecture, ARM's ThunderX2 design, and IBM's Power9 architecture are scaling out rather than up and increasing in complexity. Because the base frequencies for the large core count chips hover somewhere between 2-3 GHz, researchers can no longer rely on frequency scaling to increase the performance of their applications. Instead, developers must learn to take advantage of the increasing core count per processor and learn how to eke out more performance per core.

To achieve good performance on modern processors, developers must write code amenable to vectorization, be aware of memory access patterns to optimize cache usage, and understand how to balance multi-process programming (MPI) with multi-threaded programming (OpenMP). This tutorial will cover serial and thread-parallel optimization including introductory and intermediate concepts of vectorization and multi-threaded programming principles. We will address CPU as well as GPU profiling techniques and tools and give a brief overview of modern HPC architectures.

The tutorial will include hands-on exercises in parallel optimization, and profiling tools will be demonstrated on TACC systems. This tutorial is designed for intermediate programmers, familiar with OpenMP and MPI, who wish to learn how to program for performance on modern architectures.

---

8:00am – 3:00pm R **TUTORIAL: Python 201: Building Better Scientific Software in Python** Pathable Platform  
*Speakers: Geoffrey Lentner, Amiya Maji, Lev Gorenstein*

Scientists, engineers, researchers, and other CI professionals continue to be put in the position of being software developers. Nearly every pursuit includes the development of some final domain specific code, even if on top of robust core libraries and frameworks.

Writing code in the pursuit of science and data analysis brings with it the challenge of making that code deployable and accessible to collaborators. Novice programmers often find themselves writing difficult to maintain, difficult to manage code that their peers and collaborators have trouble using. This challenge is a hurdle to open, reproducible science.

The goal of this tutorial is to expose researchers to several best practices in scientific software engineering that may otherwise take several years to become acquainted with. Though the implementation of these lessons is Python-specific, the essential ideas can be exported to other languages or platforms.

The tutorial builds on several years of iteration and multiple instructors with a polished set of materials, hosted on GitHub ([glentner.github.io/python201](https://glentner.github.io/python201)). Delivered in previous years, including PEARC '18 and PEARC '20, we believe the content here has been well received and remains as relevant and in demand as ever.

The tutorial is hands-on with full examples outlined in a "readthedocs" style website of materials. Participants are expected to already be familiar with the Python language to the extent that they understand what Python is, how to write code files in a text editor that they are comfortable with, and run that code at the command-line on their platform of choice (Linux, macOS, Windows).

Topics covered: Python packaging, automated testing, documentation management, logging, command-line interfaces, performance profiling and optimization.

---

12:00pm – 3:00pm R **TUTORIAL: Boosting Performance of Machine Learning/Deep Learning and Dask Applications using the MVAPICH2-GDR Library** Pathable Platform  
*Speakers: Dhabaleswar Panda, Hari Subramoni, Arpan Jain, Aamir Shafi*

The recent advances in Machine Learning (ML) and Deep Learning (DL) have led to many exciting challenges and opportunities for CS and AI researchers alike. Modern ML/DL and Data Science frameworks including TensorFlow, PyTorch, Dask, and several others have emerged that offer ease of use and flexibility to train, and deploy various types of ML models and Deep Neural Networks (DNNs). In this tutorial, we will provide an overview of interesting trends in ML/DL and how cutting-edge hardware architectures and high-performance

interconnects are playing a key role in moving the field forward. We will also present an overview of different DNN architectures and ML/DL frameworks. Most ML/DL frameworks started with a single-node design. However, approaches to parallelize the process of model training are also being actively explored. The AI community has moved along different distributed training designs that exploit communication runtimes like gRPC, MPI, and NCCL. We highlight new challenges and opportunities for communication runtimes to exploit high-performance CPU and GPU architectures to efficiently support large-scale distributed training. We also highlight some of our co-design efforts to utilize MPI for large-scale DNN training on cutting-edge CPU and GPU architectures available on modern HPC clusters. The tutorial covers training traditional ML models including---K-Means, linear regression, nearest neighbors---using the cuML framework accelerated using MVAPICH2-GDR. Also, the tutorial presents accelerating GPU-based data science applications using the MPI4Dask package, which provides an MPI-based backend for Dask. Throughout the tutorial, we include hands-on exercises to enable attendees to gain first-hand experience of running distributed ML/DL training and Dask on a modern GPU cluster.

12:00pm –  
3:00pm

R **TUTORIAL: Deploying XSEDE Endpoints Using Globus Connect Server version 5** Pathable Platform  
*Speakers: Vas Vasiliadis, John-Paul Navarro, Lee Liming*

The XSEDE project operates Globus endpoints as a means for managing data on storage systems at service providers such as TACC and SDSC. The XSEDE community moves and shares tens of petabytes of data each year via these endpoints, using the Globus Connect Server (GCS) software to enable access to the Globus service. XSEDE endpoints currently run Globus Connect Server version 4 (GCSv4) and, in some cases, bespoke deployments of GridFTP. In 2021 XSEDE will introduce Globus Connect Server version 5 (GCSv5), with the intent of simplifying deployment and administration, en route to phasing out support for GCSv4 and other legacy GridFTP implementations in the future.

This tutorial will explore important new features of the GCSv5 software and help Globus endpoint operators prepare for the transition to GCSv5. We will compare and contrast the differences between GCSv4 and GCSv5, highlighting those changes that may require changes to user-facing and systems management processes. Participants will deploy a Globus endpoint using GCSv5 and will experiment with common configuration options. We will illustrate concepts using examples from the XSEDE ecosystem, but the material is equally relevant to system administrators at university research computing centers, national laboratories, and other advanced computing facilities that use Globus for data management.

12:00pm –  
3:00pm

R **TUTORIAL: Engineering your Application for Peak Performance with TAU and MVAPICH2** Pathable Platform  
*Speakers: Hari Subramoni, Sameer Shende*

This tutorial presents tools and techniques to optimize the runtime tunable parameters exposed by the MPI using the TAU Performance System® [<http://tau.uoregon.edu>]. MVAPICH2 [<http://mvapich.cse.ohio-state.edu>] exposes MPI performance and control variables using the MPI\_T interface that is now part of the MPI-3 standard. The tutorial will describe how to use TAU and MVAPICH2 for assessing the application and runtime system performance. We present the complete workflow of performance engineering, including instrumentation, measurement (profiling and tracing, timing, and PAPI hardware counters), data storage, analysis, and visualization. Emphasis is placed on how tools are used in combination for identifying performance problems and investigating optimization alternatives. We will request remote access to the Stampede system at TACC for hands-on exercises. We will also provide the ECP E4S OVA image [<https://e4s.io>] containing all of the necessary tools (running within a virtual machine) for the hands-on sessions. Participants will learn how to use the TAU Performance System with MPI, OpenMP (OMPT), CUDA, HIP, and OneAPI runtimes, and use the MPI-T interface from the MVAPICH2 library on the Frontera system at TACC and on the VM. This will help to prepare participants to locate and diagnose performance bottlenecks in their own parallel programs.

12:00pm –  
3:00pm

R **TUTORIAL: Interactive Scientific Computing on the Anvil Composable Platform** Pathable Platform  
*Speakers: Erik Gough, Eric Adams, Alan Chalker, Brian Werts, Sam Weekly, Stephen Kelley*

XSEDE capacity systems have traditionally provided batch access to large scale computing systems, meeting the high-performance computing (HPC) needs of domain scientists across numerous disciplines. New usage patterns have emerged in research computing that depend on the availability of custom services such as notebooks, databases, elastic software stacks, and science gateways alongside traditional batch HPC. Anvil, an XSEDE capacity system being deployed at Purdue University, integrates a high capacity, high performance computing cluster with a comprehensive ecosystem of software, access interfaces, programming environments, and composable services to form a seamless environment able to support a broad range of science and engineering applications. In this introductory-level tutorial, participants will get hands-on experience with interactive scientific computing using Anvil's Thinlinc remote desktop and Open OnDemand (OOD) services as well as the Anvil Composable Platform, a service providing web-based access to a Kubernetes-based private

- 
- 12:00pm – 3:00pm
- R TUTORIAL: Modern Tools for Supercomputers** Pathable Platform
- Speakers: Chun-Yaung Lu, Si Liu, Robert McLay, Kent Milfeld*
- Powerful supercomputers have played an important role in the computational research community. However, the increasing complexity of modern systems may defer or hinder their work. A large amount of precious time and effort has been spent unnecessarily managing the user environment, reproducing standard workflow, handling large scale I/O work, profiling and monitoring users' jobs, understanding and resolving unnecessary system issues, etc. To help supercomputer users focus on their scientific and technical work and to minimize the workload for the consulting team, we designed and developed a series of powerful tools for supercomputer users. These tools are portable and effective on almost all supercomputers and are now serving thousands of supercomputer users of TACC, XSEDE, and other institutions every day.
- In this tutorial, we will present and practice with supercomputer tools specifically designed for complex user environment (LMod, Sanity Tool), tools for workflow management (ibrun, launcher, launcher-GPU, pylauncher), tools for job monitoring and profiling (Remora, TACC-Stat, core\_usage, amask, etc.), and several other convenient tools. Attendees will learn how these tools are designed and used in their daily work. Detailed hands-on exercises are prepared beforehand and will be executed mainly on the Stampede2 and Frontera supercomputers at the Texas Advanced Computing Center (TACC).
- 
- 12:00pm – 3:00pm
- R TUTORIAL: Securing Science Gateways with Custos Services** Pathable Platform
- Speakers: Suresh Marru, Marlon Pierce, Eroma Abeysinghe, Jim Basney, Enis Afgan*
- The authors present a tutorial on Custos, a cybersecurity service based on open source software that helps science gateways manage user identities, integrate with federated authentication systems, manage secrets such as OAuth2 access tokens and SSH keys needed to connect to remote resources, and manage groups and access permissions to digital objects. This tutorial will provide an overview of Custos's capabilities, provide hands-on exercises on using its features, demonstrate to gateway providers how to integrate the services into their gateways with software development kits for the Custos API, introduce developers to the code and how to review and contribute to it, supply gateway providers with information on how Custos services are deployed for high availability and fault tolerance, and how Custos operations handle incidence response. This tutorial builds on a successful online tutorial presented by the authors at Gateways 2020; see <https://www.youtube.com/watch?v=CuBvFj194Kg> for a recording.
- 
- 12:00pm – 3:00pm
- K WORKSHOP: Fourth Workshop on Strategies for Enhancing HPC Education and Training (SEHET21)** Pathable Platform
- Speakers: Dana Brunson, Nitin Sukhija, Scott Lathrop, Julia Mullen, Kate Cahill, Susan Mehringer*
- High performance computing is becoming central for empowering scientific progress in the most fundamental research in various science and engineering, as well as society domains. It is remarkable to observe that the recent rapid advancement in the mainstream computing technology has facilitated the ability to solve complex large-scale scientific applications that perform advanced simulations of the implementation of various numerical models corresponding to numerous complex phenomena pertaining to diverse scientific fields. The inherent wide distribution, heterogeneity, and dynamism of the today's and future computing and software environments provide both challenges and opportunities for cyberinfrastructure facilitators, trainers and educators to develop, deliver, support, and prepare a diverse community of students and professionals for careers that utilize high performance computing to advance discovery.
- The SEHET21 workshop is an ACM SIGHPC Education Chapter coordinated effort aimed at fostering collaborations among the practitioners from traditional and emerging fields to explore strategies to enhance computational, data-enabled and HPC educational needs. Attendees will discuss approaches for developing and deploying HPC training, as well as identifying new challenges and opportunities for keeping pace with the rapid pace of technological advances - from collaborative and online learning tools to new HPC platforms. The workshop will provide opportunities for: learning about methods for conducting effective HPC education and training; promoting collaborations among HPC educators, trainers and users; and for disseminating resources, materials, lessons learned and good/best practices.
- 
- 12:00pm – 3:00pm
- K WORKSHOP: Refining Your Research Computing Pitch** Pathable Platform
- Speakers: Jane Combs, Richard Knepper*
- Speak above the noise. In the age of over communication via many channels, one common need among research service providers, particularly Research Computing Centers, is reaching their audience and getting the

word out that their services exist. We propose a PEARC workshop for professionals (center leaders, facilitators, faculty champions, etc) of existing or emerging research computing organizations (even those “one-person shops”) to get feedback on and develop their communication materials: new faculty handouts or introductory slides that are used to provide the first communication to faculty, students and administrators to familiarize them with research computing resources available at our institutions. Some existing materials will be reviewed and workshop members can provide feedback, then individual groups will convene to improve some sample materials provided by collaborators. The group will reconvene and identify lessons learned and feedback on the overall workshop. We hope through this workshop to provide a clearing-house of template materials that will make campus outreach easier for research computing professionals, which can be reviewed and improved by participants over time.

---

12:00pm –  
3:00pm

**K WORKSHOP: What Does it Mean to be a Campus Champion?**

Pathable Platform

*Speakers: Dana Brunson, Shelley Knuth, Jay Alameda, Torey Battelle, Thomas Cheatham, Douglas Jennewein, Bj Lougee, Julie Ma, Timothy Middelkoop, Henry Neeman*

The importance of research computing and data infrastructure for scientific discovery and scholarly achievement has grown as research questions become more complicated and datasets get larger. Many higher education institutions have built infrastructure utilized by researchers and supported by local research computing staff. Some of these research computing groups are large, with several infrastructure and support staff, while others may be supported by only one or two staff members. For both of these groups, the Campus Champions program has provided an opportunity for knowledge exchange, professional development, and growth (Brazil 2019). Over the past twelve years, the Campus Champions program has grown to nearly 720 Champions at over 300 research institutions. A significant number of Champions attend the PEARC conference; in 2020, over 23% of PEARC attendees were Champions.

With on average a net gain of 40 new Champions per year, a cohesive and all-encompassing onboarding program must be in place to ensure that the Champions get exposure to information and resources to assist members at their institutions as well as their peers in the community. For several years, an informal onboarding process has served as a basic introduction, with members of the community relying on the web or their colleagues to determine where appropriate resources and information might be located. A more extensive process that includes mentoring and exposure to resources is desired.

Given the number of Campus Champions who attend PEARC, and its vision to foster “the creation of a dynamic and connected community of advanced research computing professionals who promote leading practices and the frontiers of research, scholarship, teaching, and industry application” [PEARC 2019], hosting a workshop dedicated to the onboarding of Campus Champions would have maximum impact.

At PEARC 2020, this workshop was hosted for the first time (also entitled, “What Does It Mean to be a Campus Champion?”), with approximately 60 participants remaining engaged throughout the day. Anecdotally, we received a lot of positive feedback about this workshop from the participants. While the targeted audience was new or recent Campus Champions, more experienced Champions as well as those who were not Champions benefitted from the workshop. The more experienced Champions learned about changes in the program or about resources of which they were not aware, while people who were not Champions were either introduced to the community or learned practices to best support researchers at their institutions (if they support infrastructure) or to best utilize infrastructure (if they do not).

---

12:00pm –  
4:00pm

**K EXHIBITOR WORKSHOP: Optimize your Science & Simulations with AMD HPC Solutions**

*Speakers: Jason Hogan O’Neill, Brock Taylor, Mathieu Gontier, Derek Bouius, Michael Klemm, Katrina Gundal-Zaidi* Pathable Platform

Come interact with specialists from AMD’s HPC Center of Excellence who will share insights and techniques in optimizing for AMD platforms. Topics will cover broad elements of architecture and software development for 3rd generation AMD EPYC™ processors and AMD Instinct™ MI100 GPUs. AMD experts will provide their real-world experiences and discuss hardware and software elements that help build developer acumen. This session will dive into considerations for CPU based applications including characterization. It will expand on AMD’s open environment for developing accelerated applications and will include a hands-on view of analyzing applications using the AMD  $\mu$ Prof tool.

---

12:00pm –  
4:00pm

**K EXHIBITOR WORKSHOP: Google - Tropical Cyclone Intensity Estimation using a Deep Convolutional Neural Network**

Pathable Platform

*Speakers: Kenton McHenry*



**A** and Education    **L** Applications and Software    **B** BoF    **D** Diversity    **P** Panels

**S** Systems and System Software    **T** Training    **R** Tutorial    **W** Workforce Development    **K** Workshop

---

**JULY 20 • TUESDAY**

---

9:30am – 9:40am    **L**    **Accelerating key bioinformatics tasks 100-fold by improving memory access**    Pathable Platform  
*Speakers: Igor Sfiligoi, Rob Knight, Daniel McDonald*

Most experimental sciences now rely on computing, and biological sciences are no exception. As datasets get bigger, so do the computing costs, making proper optimization of the codes used by scientists increasingly important. Many of the codes developed in recent years are based on the Python-based NumPy, due to its ease of use and good performance characteristics. The composable nature of NumPy, however, does not generally play well with the multi-tier nature of modern CPUs, making any non-trivial multi-step algorithm limited by the external memory access speeds, which are hundreds of times slower than the CPU's compute capabilities. In order to fully utilize the CPU compute capabilities, one must keep the working memory footprint small enough to fit in the CPU caches, which requires splitting the problem into smaller portions and fusing together as many steps as possible. In this paper, we present changes based on these principles to two important functions in the scikit-bio library, principal coordinates analysis and the Mantel test, that resulted in over 100x speed improvement in these widely used, general-purpose tools.

9:30am – 9:40am    **S**    **Best practice of IO workload management in containerized environments on supercomputers**    Pathable Platform  
*Speakers: Lei Huang, Yinzhi Wang, Chun-Yaung Lu, Si Liu*

With the increasing adoption of containerization among HPC applications and workflows to achieve the ever-increasing demand in portability, flexibility, and customizability, the IO pattern and workload from such an isolated environment have evolved with a significant amount of complexity. The steadily increasing demand for powerful storage resources has become a problem when running on supercomputers that usually have a shared underlying filesystem. The IO-intensive workload from a single container can lead to performance degradation, or sometimes a complete breakdown, of the whole filesystem. While some IO workload management tools can well address such issues for conventional HPC applications, none of them were designed and tested to accommodate the cases of such isolated environments to the best knowledge of the authors. In this paper, we present a feasibility study using the Optimal Overloaded IO Protection System (OOOPS) tool to reduce the IO impact from containerized environments and discuss the best practice of IO workload management on supercomputers.

9:30am – 9:40am    **W**    **An Evaluation of Cyberinfrastructure Facilitators Skills Training in the Virtual Residency Program**    Pathable Platform  
*Speakers: Dana Brunson, Henry Neeman, Sandra Gesing, Lorna Rivera, Lizanne DeStefano, Hussein Al-Azzawi, Patrick Clemins, Dirk Colbry, Calvin Frye, Joshua Gyllinsky, Anna Klimaszewski-Patterson, Anchalee Phataralaoha, Todd Price, Mohammed Tanash, Daniel Voss*

Cyberinfrastructure (CI) Facilitation amplifies the productivity of researchers engaged in computing-intensive and data-intensive investigations. CI Facilitators help researchers to adapt their workflows to CI resources, and teaches these researchers how to use these systems, bridging between researchers and technology experts. The importance of CI Facilitators is well understood, and there is broad consensus about their need for ongoing training, especially because of not only their highly diverse career backgrounds and domains of expertise, but also the rapid advancement of CI technologies. The Virtual Residency Program (VRP) has been addressing this training gap since 2015, offering summer workshops that teach key CI Facilitation skills. These workshops are community driven and community developed. The 2020 VRP workshop, for the first time, featured an external evaluation of the efficacy and value of the workshop. This paper examines the results of that evaluation, to determine the value of the VRP to the CI Facilitator community, and thus to the many researchers these CI Facilitators serve.

9:30am – 11:00am    **P**    **Research Software Engineer (RSE) Careers - State of the Profession, Opportunities, and Challenges in Academia**    Pathable Platform  
*Speakers: Kenton McHenry, Sandra Gesing, Ian Cosden, Daniel S. Katz, Christina Maimone, Alan Sussman*

Research is becoming increasingly digital and reliant on software. As a result of this reliance on software, a new role within the research ecosystem has emerged: the Research Software Engineer (RSE). Research Software

Engineers, whose professional focus is centered around developing and contributing to research software, bring software engineering skills and practices to research to create more robust, manageable, and sustainable research software. These RSEs are contributing in an increasingly meaningful way to the computational science and engineering research ecosystem.

As institutions and organizations have begun to formally support and recognize the role, the number of individuals identifying and associating with the RSE role has steadily increased. Since its formation in 2018, the US Research Software Engineer Association, currently with almost 800 members, has seen a significant increase in both activity and membership. Despite the recent growth and community acceptance of the RSE concept, the RSE career path remains new within the academic research world and as a result, RSE career paths remain confusing and opaque. This panel is intended to provide attendees with an opportunity to hear from established RSEs and RSE leaders about the current state of the profession as well as the career opportunities and challenges currently facing the RSE field within academic institutions.

---

9:30am – 11:00am	<b>P Scientific Research Enabled by Cerebras CS-1 Systems</b>	Pathable Platform
	<i>Speakers: Paola A. Buitrago, Sergiu Sanielevici, Natalia Vassilieva</i>	
	The Cerebras CS-1 system is a promising, unique, innovative, and advanced compute server built around the world's largest processor, the Wafer Scale Engine (WSE) and is designed to deliver unprecedented performance on artificial intelligence (AI) as well as on suitable HPC workloads. CS-1 systems are now deployed at several advanced computing facilities serving academic, national laboratory, and commercial researchers worldwide, including the Pittsburgh Supercomputing Center (PSC), Argonne National Laboratory, (ANL) and Lawrence Livermore National Laboratory (LLNL). This PEARC21 panel will, for the first time, bring together researchers who have used these systems, to share their experiences with each other, with interested members of the PEARC community, and with Cerebras Systems. The organizers will set the stage by summarizing the state of the technology, its evolution, and its promise for research applications. Each panelist will then briefly present their research, followed by a discussion moderated by the organizers. Audience members will have the opportunity to explore how their own research could benefit from the WSE-powered systems and how to gain access to one of these limited servers. The organizers hope that this session will seed the development of a persistent user group and foster information exchange on this promising technology.	
9:40am – 9:50am	<b>L Adaptive Plasma Physics Simulations: Dealing with Load Imbalance using Charm++</b>	Pathable Platform
	<i>Speakers: Diego Jiménez, Esteban Meneses, V.I. Vargas</i>	
	High Performance Computing (HPC) is nearing the exascale era and several challenges have to be addressed in terms of application development. Future parallel programming models should not only help developers take full advantage of the underlying machine but they should also account for highly dynamic runtime conditions, including frequent hardware failures. In this paper, we analyze the porting process of a plasma confinement simulator from a traditional MPI+OpenMP approach to a parallel objects based model like Charm++. The main driver for this effort is the existence of load imbalanced input scenarios that through pure OpenMP scheduling cannot be solved. By using Charm++ adaptive runtime and integrated balancing strategies, we were able to increase total CPU usage from 45.2% to 80.2%, achieving a 1.64x acceleration, after load balancing, over the MPI+OpenMP implementation on a specific input scenario. Checkpointing was added to the simulator thanks to the pack-unpack interface implemented by Charm++, providing scientists with fault tolerance and split execution capabilities.	
9:40am – 9:50am	<b>S Secure Research Infrastructure Using tiCrypt</b>	Pathable Platform
	<i>Speakers: Erik Deumens, Thomas Samant, Eric Byrd, Samuel Porter, Axel Haenssen, Curtis W. Hillegas, Elizabeth Holtz, Irene Kopaliani</i>	
	Digital data is increasingly important and becoming more prevalent in research. As such, a significant and growing fraction of that data is subject to complex security and privacy standards for its storage and processing. With the introduction of contractual requirements for NIST Special Publication 800-171 and other compliance frameworks, universities are under expanded pressure to secure research infrastructure to ensure their researchers comply with laws, regulations, and contractual requirements. The University of Florida developed tiCrypt, a bundled security middleware product, and implemented their ResVault system to meet the growing need for secure research infrastructure. The University of Maryland and Princeton University, both seeking to meet compliance requirements for secure research, have also implemented systems using tiCrypt. As new institutions adopt tiCrypt, the Tera Insights development team continues to add functionality to meet the needs of a broad range of research workflows.	
9:40am –	<b>W Research Liaisons: the next layer of Facilitation</b>	Pathable Platform

---

9:50am

*Speakers: Deb McCaffrey, John Brussolo, Ryan Echlin, John Herlocher, Einor Jacobsen, Lovida Roach, Amy Yamasaki, Erin Dietrich, Dan St. Pierre*

Research productivity has been greatly enhanced by Research Computing Facilitation teams to help them maximize their use of advanced cyberinfrastructure. However, researchers have more technology needs than advanced cyberinfrastructure, such as data management and instrument device support. To address this, the Academic Engagement team in Michigan Medicine added Research Liaisons as another layer of human support on top of the Facilitation team.

The Liaisons are relationship builders. They are assigned to departments to build deep relationships with them and start proactively addressing labs' technology needs. They also build relationships with other teams, notably enterprise storage, enterprise networking, and research core facilities. These relationships allow Liaisons to provide a connective tissue between the researchers and IT teams.

---

9:50am –  
10:00am

**L Tools and Guidelines for Job Bundling on Modern Supercomputers**

Pathable Platform

*Speakers: Si Liu, Victor Eijkhout, John Cazes*

As many scientific computation tasks focus on solving large-scale and computationally intensive problems, a wide range of problems involving High-Throughput Computing (HTC) paradigms and data-oriented algorithms emerge. Solving these HTC problems efficiently on modern supercomputers usually requires efficient and convenient job bundling. In this research, we evaluate multiple handy tools and workflows that are used to realize efficient and convenient job bundling. We also provide some practice guidelines for users when job bundling is required.

---

9:50am –  
10:00am

**S Research Computing on Campus - Application of a Production Function to the Value of Academic High-Performance Computing**

Pathable Platform

*Speakers: Stephen Lien Harrell, Preston Smith*

High-performance computing is used by research institutions worldwide for managing research data, modeling, simulation, and big data analysis. This capability is critical for higher education institutions to attract top researchers and faculty, maximizing competitiveness for future funding, and training students in current data-intensive research methods.

As a corollary, procurement and operation of advanced cyberinfrastructure incurs substantial operating and capital expenses to an organization. In this paper we present an analysis of value metrics gathered at Purdue University to measure the return on investment (ROI) of institutional investment in cyberinfrastructure resources; and an application of an economic production function to measure the cyberinfrastructure's impact on the institution's academic, financial, and reputational output.

---

9:50am –  
10:00am

**W Assessing the Landscape of Research Computing and Data Support**

Pathable Platform

*Speakers: Patrick Schmitz*

We describe the first Research Computing and Data Capabilities Model Community Dataset, aggregating the assessments of 41 Higher Education Institutions. These assessments were completed using the 1.0 version of the Research Computing and Data Capabilities Model (RCD CM), over a period of several months in the Spring and Summer of 2020. The RCD CM allows organizations to self-evaluate across a range of RCD services and capabilities for supporting research, leveraging a shared vocabulary to describe RCD support. The Model supports a range of stakeholders and provides structured input to guide strategic planning and enable benchmarking relative to peer institutions. This Community Dataset provides insight into the current state of support for RCD across the community and in a number of key sub-communities. The dataset shows stark differences between Public and Private institutions, between institutions with a larger and smaller share of national funding, etc. In many cases, the patterns in the data confirm common perceptions about support across the community, but the dataset provides a quantitative baseline for understanding current RCD support, as well as granular insights to groups of institutions that are seeking to collaborate on shared solutions and strategies to advance RCD support. Over time, longitudinal data will provide additional insight into trends, and a means of evaluating the impact of programs designed to increase RCD support.

---

10:00am –  
10:10am

**L Comparing the behavior of OpenMP Implementations with various Applications on two different Fujitsu A64FX platforms**

Pathable Platform

*Speakers: Benjamin Michalowicz, Eric Raut, Yan Kang, Tony Curtis, Barbara Chapman, Dossay Oryspayev*

The development of the A64FX processor by Fujitsu has allowed the re-emergence of vectorized processors and the birth of the first supercomputer to achieve top speeds in all 5 of the major HPC benchmarks. We use a

variety of tools to analyze the behavior and performance of several OpenMP applications with different compilers, and how these applications scale on the A64FX processor on A64FX clusters at Stony Brook University and RIKEN in Japan.

- 
- 10:00am – 10:10am S **Jetstream2: Accelerating cloud computing via Jetstream** Pathable Platform  
*Speakers: Jeremy Fischer, Suresh Marru, Marlon Pierce, David Y. Hancock, John Michael Lowe, Winona Snapp-Childs, J. Eric Coulter, Matthew Vaughn, Brian Beck, Nirav Merchant, Gwen Jacobs, Edwin Skidmore*  
Jetstream2 will be a category I production cloud resource that is part of the National Science Foundation's Innovative HPC Program. The project's aim is to accelerate science and engineering by providing "on-demand" programmable infrastructure built around a core system at Indiana University and four regional sites. Jetstream2 is an evolution of the Jetstream platform, which functions primarily as an Infrastructure-as-a-Service cloud. The lessons learned in cloud architecture, distributed storage, and container orchestration have inspired changes in both hardware and software for Jetstream2. These lessons have wide implications as institutions converge HPC and cloud technology while building on prior work when deploying their own cloud environments. Jetstream2's next-generation hardware, robust open-source software, and enhanced virtualization will provide a significant platform to further cloud adoption within the US research and education communities.
- 
- 10:00am – 10:10am W **Broadening the Reach for Access to Advanced Cyberinfrastructure - Accelerating Research and Education** Pathable Platform  
*Speakers: Forough Ghahramani, John Hicks, Barr von Oehsen*  
Many smaller, mid-sized and under-resourced campuses, including MSIs, HSIs, HBCUs and EPSCoR institutions, have compelling science research and education activities along with an awareness of the benefits associated with better access to cyberinfrastructure (CI) resources. These schools can benefit greatly from resources and expertise to augment their in-house efforts. The Eastern Regional Network's (ERN) Broadening the Reach (BTR) working group is addressing this by focusing on learning directly from the under-resourced academic institutions in the region on how best to support them for research collaboration and advanced computing requirements. ERN BTR findings and recommendations will be shared based on engagement with the community, including workshop and survey results, as part of the NSF sponsored CC\*CRIA: The Eastern Regional Network Award OAC-2018927.
- 
- 10:10am – 10:20am L **kEDM: A Performance-portable Implementation of Empirical Dynamic Modeling using Kokkos** Pathable Platform  
*Speakers: Keichi Takahashi, Wassapon Watanakeesuntorn, Kohei Ichikawa, Joseph Park, Ryousei Takano, Jason Haga, Gerald M. Pao, George Sugihara*  
Empirical Dynamic Modeling (EDM) is a state-of-the-art non-linear time-series analysis framework. Despite its wide applicability, EDM was not scalable to large datasets due to its expensive computational cost. To overcome this obstacle, researchers have attempted and succeeded in accelerating EDM from both algorithmic and implementational aspects. In previous work, we developed a massively parallel implementation of EDM targeting HPC systems (mpEDM). However, mpEDM maintains different backends for different architectures. This design becomes a burden in the increasingly diversifying HPC systems, when porting to new hardware. In this paper, we design and develop a performance-portable implementation of EDM based on the Kokkos performance portability framework (kEDM), which runs on both CPUs and GPUs while based on a single codebase. Furthermore, we optimize individual kernels specifically for EDM computation, and use real-world datasets to demonstrate up to 5.5x speedup compared to mpEDM in convergent cross mapping computation.
- 
- 10:10am – 10:20am S **Building Tapis v3 Streams API Support for Real-Time Streaming Data Event-Driven Workflows** Pathable Platform  
*Speakers: Smruti Padhy, Anagha Jamthe, Sean B. Cleveland, Jack A. Smith, Joe Stubbs, Christian Garcia, Michael Packard, Steve Terry, Julia Looney, Richard Cardone, Maytal Dahan, Gwen A. Jacobs*  
The Tapis framework, an NSF-funded project, is an open-source, scalable API platform that enables researchers to perform distributed computational experiments securely and achieve faster scientific results with increased reproducibility. Tapis Streams API focuses on supporting scientific use cases that require working with real-time sensor data. The Streams Service, built on the top of the CHORDS time-series data service, allows storing, processing, annotating, querying, and archiving time-series data. This paper focuses on the new Tapis Streams API functionality that enables researchers to design and execute real-time data-driven event workflow for their research. We describe the architecture and design choices towards achieving this new capability with Streams API. Specifically, we demonstrate the integration of Streams API with Kapacitor, a native data processing engine for time-series database InfluxDB, and Abaco, an NSF Funded project, web service, and distributed computing platform providing function-as-a-Service (FaaS). The Streams API, which includes a wrapper interface for the

Kapacitor alerting system, can define and enable alerts. Finally, simulation results from the water-quality use case depict that Streams API's new capabilities can support real-time streaming data event-driven workflows.

- 
- 10:10am – 10:20am      **W Comprehensive Evaluation of XSEDE's Scientific Impact using Semantic Scholar Data** Pathable Platform  
*Speakers: Gregor von Laszewski, Fugang Wang*  
The United States science and engineering community face multiple challenges related to funding and funding policies for science and engineering. A framework is needed to evaluate the impact scientific facilities and instruments. In this paper, we demonstrate such an activity through our comprehensive work evaluating the scientific impact of XSEDE using the Semantic Scholar Data. In contrast to other studies, our study includes the bibliographic references of all recorded papers related to XSEDE over the entire performance period till April of 2021. This makes this study unique and distinguishes it from our earlier work while using (a) over 180 million papers as a comparison to our peer analysis, (b) include all publications reported, and (c) conduct the study repeatedly over several years to verify its validity.
- 
- 10:20am – 10:30am      **L Building Detection with Deep Learning** Pathable Platform  
*Speakers: Matthew Kusz, Justin Peters, Laura Huber, Jefferson Davis, Scott Michael*  
Deep learning frameworks have been widely used in image classification and segmentation tasks. In this paper we outline the methods used to adapt an image segmentation model, U-Net, to identify buildings in geospatial images. The model has been trained and tested on a set of orthophotographic and LiDAR data from the state of Indiana. Its results are compared to the results achieved by a ResNet101 and RefineNet model trained with the same data, excluding the LiDAR data. This tool has a wide range of potential uses in research involving geospatial imagery. We discuss these use cases and some of the challenges and pitfalls in tuning a model for use with geospatial data.
- 
- 10:20am – 10:30am      **S Reimagining Account and Allocation Management for Advanced Research Computing** Pathable Platform  
*Speakers: Tabitha Samuel, Adam Howard*  
Though they are the backbone of a center's infrastructure and store some of its most vital information, often times, the databases responsible for tracking project allocations, job accounting, storage commitments, and other core center information aren't designed with the proper attention and thought to scale and adapt with the center as it grows and changes. Off-the-shelf solutions for these data stores are limited, and come with their own assumptions. As the field of research computing and thus the demands on the center continue to evolve, so must the accounting solutions be able to accommodate the changing mission and goals. Many centers develop their own solutions in house, which requires additional effort to build and maintain. In this work, we present a data model which attempts to provide a more generalized solution to this problem, as well as a feature-rich set of tools provided by leveraging the powerful and widely-used Django web framework and ORM, which we call OpenAcct. We explore other solutions to this data management problem including both vendor-provided and in-house developed options. We then discuss the details of the data model, before demonstrating the capabilities of the application layer. Finally, future plans for the project are discussed.
- 
- 10:20am – 10:30am      **W The Connect.Cyberinfrastructure Portal - Creating Opportunities for Collaboration and Cohesion in the Research Computing Ecosystem** Pathable Platform  
*Speakers: Dana Brunson, Shelley Knuth, Torey Battelle, Thomas Cheatham, Douglas Jennewein, Julie Ma, Timothy Middelkoop, Eric Brown, John Goodhue, James Griffioen, Kaylea Nelson, Andrew Sherman, Scott Valcourt, Dhruva Chakravorty, Christopher Simmons, Sarah Akbar, Ermal Toto, Bj Lougee, Julia Sheats*  
The Connect.Cyberinfrastructure.org (Cnct.CI) Portal, originally known as the Cyberteam Portal, was developed to support the management of project workflows and to capture project results for the Northeast Cyberteam (NECT). The NECT is an NSF-sponsored program that aims to make regional and national cyberinfrastructure more readily accessible to researchers in small and mid-sized institutions in northern New England by providing research computing facilitation support and aggregating access to knowledge resources. More recently, the Cnct.CI Portal has expanded to provide support for a variety of programs in the Research Computing ecosystem, creating opportunities for collaboration and leveraging a consistent, cohesive approach to common challenges. As reported at SC20, a pilot was launched in July 2020 to enable six additional Cyberteam programs to explore the use of the Portal as a management tool for their related programs. In addition, in January of 2021 the leadership team of the Extreme Science and Engineering Discovery Environment (XSEDE) Campus Champions decided to use the Portal to modernize participant management and onboarding functions. Portal details, preliminary results of these expansion efforts and future plans are discussed.
-

11:15am –  
12:15pm

**B Assessing and communicating cyberinfrastructure readiness at EPSCoR and under-resourced institutions**

Pathable Platform

*Speakers: Patrick Schmitz, Dana Brunson, Scotty Strachan, Venice Bayrd, Gwen Jacobs, Pips Veazey*  
This Birds of a Feather (BoF) group session is aligned with an accepted PEARC21 workshop associated with the Research Computing and Data (RCD) Capabilities Model. This model allows institutions to assess their support for computationally- and data-intensive research, to identify potential areas for improvement, and to understand how the broader community views RCD support. Application of this model within National Science Foundation Established Program to Stimulate Competitive Research (NSF-EPSCoR) jurisdictions presents a set of potential challenges beyond typical expectations of model engagement and response quality. Institutions classified as NSF-EPSCoR-eligible are typically underfunded for research infrastructure and very often have minimal organized cyberinfrastructure support and institutional awareness. As a result of this, participation in the RCD Capabilities Model assessment survey requires additional preparation and often external assistance to ensure effective survey engagement by institutional research and administrative personnel. This BoF, which is part of a series of EPSCoR-focused engagement events for 2020-2022, brings together RCD Capabilities Model Working Group members and research technology support personnel from NSF-EPSCoR institutions to review ongoing challenges to assess and communicate institutional cyberinfrastructure (CI) readiness, discuss current model engagement efforts, provide feedback to the Model development process, and brainstorm potential cooperative CI efforts across the EPSCoR program. Non-EPSCoR institutions with interest in the Capabilities Model are welcome to participate and contribute.

11:15am –  
12:15pm

**B Open OnDemand User Group Meeting**

Pathable Platform

*Speakers: Alan Chalker, David Hudak, Robert DeLeon, Robert Settlage*

Open OnDemand is an NSF-funded open-source HPC platform currently in use at over 200 HPC centers around the world. It is an intuitive, innovative, and interactive interface to remote computing resources. Open OnDemand helps computational researchers and students efficiently utilize remote computing resources by making them easy to access from any device. It helps computer center staff support a wide range of clients by simplifying the user interface and experience.

This BoF is meant to be an open discussion to guide the future roadmap for Open OnDemand in the near term, by getting feedback from the community on the prioritization of the various tasks planned for the next few years.

As there are many people that attended previous iterations of this BOF at PEARC'18, PEARC'19 and PEARC'20 and spoke highly of them, the session leaders intend to replicate the same BOF format (with appropriate updates regarding what has been done in the past year and the roadmap for the current NSF award) and anticipate comparable numbers of attendees to the previous BOFs. A report summarizing the status of current installations and additional feature requests will be generated and distributed to anyone that has expressed interest in Open OnDemand.

11:15am –  
12:15pm

**B OpenHPC Community BoF**

Pathable Platform

*Speakers: Karl Schulz, Christopher Simmons, Nirmala Sundarara, David Brayford, Derek Simmel*

Over the last several years, OpenHPC has emerged as a community-driven stack providing a variety of common, pre-built ingredients to deploy and manage an HPC Linux cluster. Formed initially in November 2015 and formalized as a Linux Foundation project in June 2016, OpenHPC has been adding new software components and now supports multiple OSes/architectures. At this BoF, speakers from the OpenHPC Technical Steering Committee will provide technical highlights from the latest 2.x release first introduced in Q4 2020. We will then highlight potential community plans and solicit feedback regarding the recent distro announcement to discontinue CentOS8 at the end of 2021. Finally, we will invite open discussion giving attendees an opportunity to provide feedback on current conventions, packaging, request additional components and configurations, and discuss general future trends.

11:15am –  
12:15pm

**B Use Cases for ColdFront Resource Allocation Management System**

Pathable Platform

*Speakers: Dori Sajdak, Andrew Bruno, James Hong, Bd Kim, Robert Settlage, Raminder Singh*

ColdFront is an open source resource allocation management system designed by the University at Buffalo Center for Computational Research to provide a central portal for administration, reporting, and measuring scientific impact of cyberinfrastructure resources. ColdFront is designed to manage allocations for a diverse set of resources including HPC clusters, co-located departmental and lab servers, software licenses, digital storage, scientific instrumentation, cloud subscriptions, and data access requests. It is designed to complement existing tools, such as Slurm and OpenStack, which manage access to individual hardware and software components. Though every institution will have its own policies for granting access to a resource, how that access gets

provided is nearly universal. Through the flexibility of ColdFront's plug-ins, centers can easily automate their access policies with the back-end systems resulting in significant time management cost savings for campuses both small and large. Interest in ColdFront has exploded in the last 18 months and many HPC centers are either evaluating it or using it in production. In this BoF, attendees will hear from a panel of ColdFront adopters at three academic HPC centers who will discuss the various ways they're using ColdFront to manage access to their resources. CCR staff will provide a roadmap for future development and we'll engage attendees in a Q&A period to round out the session.

---



utilization of the HPC clusters.

Our model achieved over 85% accuracy in predicting the amount of time and the amount of memory for both SUMMIT and Beocat HPC resources. Our results show that our model helps dramatically reduce computational average waiting time (from 380 to 4 hours in RMACC Summit and from 662 hours to 28 hours in Beocat); reduced turnaround time (from 403 to 6 hours in RMACC Summit and from 673 hours to 35 hours in Beocat); and achieved up to 100% utilization for both HPC resources.

---

9:20am – 10:50am	<p><b>P Campus Research Computing Consortium (CaRCC) Town Hall</b> <span style="float: right;">Pathable Platform</span></p> <p><i>Speakers: Lauren Michael, Claire Mizumoto, Dana Brunson, Thomas Cheatham, Ruth Marinshaw, Patrick Schmitz, Scott Yockel</i></p> <p>CaRCC – the Campus Research Computing Consortium – is an organization of dedicated professionals developing, advocating for, and advancing campus research computing and data and associated professions. CaRCC advances the frontiers of research by improving the effectiveness of research computing and data (RCD) professionals, including their career development and visibility, and their ability to deliver services and resources for researchers. CaRCC connects RCD professionals and organizations around common objectives to increase knowledge sharing and enable continuous innovation in research computing and data capabilities. This panel will gather CaRCC leaders and community members to discuss recent products and significant activities that CaRCC has supported as well as new initiatives for 2021 and beyond. CaRCC is always interested to hear the concerns of our community and their ideas for what CaRCC (in partnership with other community organizations) can do to better support RCD Professionals. We will ensure there is plenty of time for audience questions and discussion, and we welcome all who are interested in (and/or involved with) CaRCC and our partner organizations in the community.</p>
9:20am – 10:50am	<p><b>P Cloud Facilitation in Research Computing and Data</b> <span style="float: right;">Pathable Platform</span></p> <p><i>Speakers: Dana Brunson, Timothy Middelkoop, Allissa Dillman, Alper Kinaci, Ananya Ravipati, Chris Lalande, Rob Fatland, John Mulligan</i></p> <p>Facilitating research in the cloud presents a unique set of challenges for Research Computing and Data facilitators, researchers, students, and institutions. In this panel we focus on the challenges and obstacles faced by facilitators and some of the tools, methods, and approaches used at various institutions to overcome them and lessons learned. Areas of discussion will include cloud training for facilitators and researchers, what workloads are appropriate (and not) for the cloud, the importance of adapting and fundamentally changing workloads for the cloud, managing the plethora of services and vendor lock in, cost management, and some of the valuable lessons learned and fundamental changes in thinking along the journey to the cloud. This panel will take the perspective of utilizing the most popular public cloud providers but will be relevant to most cloud-native environments (including on-prem cloud).</p>
9:20am – 10:50am	<p><b>P PEARC21 Panel on Student Competitions</b> <span style="float: right;">Pathable Platform</span></p> <p><i>Speakers: Elizabeth Leake</i></p> <p>Research computing and data (RCD) science conferences can be overwhelming for first-time attendees; especially for students. Even so, they're one of the best ways for students to get acquainted with advanced technologies, science drivers, and the global community of caring that leads to academic acceleration, internships and employment.</p> <p>Some students enjoy participating in co-located competitions where they can showcase their skills and meet like-minded colleagues. The environment serves as a conference on-ramp experience; a safe place where they can shine, and help others improve their skills while they're at it.</p> <p>There are co-located student cluster competitions (SCC's) at the Supercomputing Asia Conference (SCA, early March), International Supercomputing Conference (ISC, mid-June), Supercomputing Conference (SC, late November), Centre for High Performance Computing National Meeting (CHPC, early December), and others. ISC and SC contests were virtual in 2020 (ISC'21); the inaugural Winter Classic virtual contest is expected to return February 2022. The National Science Foundation's Science Gateways Community Institute (NSF-SGCI) Hackathon was virtual in 2020 and 21; it will conclude one week prior to PEARC21 so that students can take full advantage of the conference workshops and tutorials. The week's respite will help prevent Zoom fatigue.</p> <p>When interviewed, competitors who prevailed at SC and ISC-SCC's testified to having participated in multiple events each year; the added exposure appears to be an advantage. Most agree we should offer more such opportunities—real and virtual—throughout the year, so that it's geographically favorable and possible for more</p>

to participate.

Is ACM-PEARC a candidate for a co-located student competitive training forum? It would offer US students another affordable domestic option, and would augment what the February WC and July virtual Hackathon events foster among students from MSIs. SC21 is planning to offer a menu of competitive options for students; perhaps PEARCXX could, too?

This panel of international RCD experts has more than 60 years of collective experience with student competitions, and others have been invited. Each will present for 8-10 minutes, followed by a Q&A session for the balance of 90 minutes. Panelists will share lessons learned from managing student competitions, and potential silver linings they've observed from facilitating virtual events—a necessity during the global pandemic. Because they've been involved with competitions for a decade or more, they have witnessed the longitudinal benefits; some first became engaged as students. Each will share their own career “arcs” and highlight how the experience helped students they've mentored.

9:30am –  
9:40am

**L A Heterogeneous MPI+PPL Task Scheduling Approach for Asynchronous Many-Task Runtime Systems**

Pathable Platform

*Speakers: John Holmen, Martin Berzins, Damodar Sahasrabudhe*

Asynchronous many-task runtime systems and MPI+X hybrid parallelism approaches have shown promise for helping manage the increasing complexity of nodes in current and emerging high performance computing (HPC) systems, including those for exascale. The increasing architectural diversity, however, poses challenges for large legacy runtime systems emphasizing broad support for major HPC systems. Performance portability layers (PPL) have shown promise for helping manage this diversity. This paper describes a heterogeneous MPI+PPL task scheduling approach for combining these promising solutions with additional consideration for parallel third party libraries facing similar challenges to help prepare such a runtime for the diverse heterogeneous systems accompanying exascale computing. This approach is demonstrated using a heterogeneous MPI+Kokkos task scheduler and the accompanying portable abstractions~\cite{SCI:Hol2019b} implemented in the Uintah Computational Framework, an asynchronous many-task runtime system, with additional consideration for hypre, a parallel third party library. Results are shown for two challenging problems executing workloads representative of typical Uintah applications. These results show performance improvements up to 4.4x when using this scheduler and the accompanying portable abstractions~\cite{SCI:Hol2019b} to port a previously MPI-Only problem to Kokkos::OpenMP and Kokkos::CUDA to improve multi-socket, multi-device node use. Good strong-scaling to 1,024 NVIDIA V100 GPUs and 512 IBM POWER9 processor are also shown using MPI+Kokkos::OpenMP+Kokkos::CUDA at scale.

9:30am –  
9:40am

**S Real-World, Self-Hosted Kubernetes Experience**

Pathable Platform

*Speakers: Joe Stubbs, Michael Packard, Justin Drake, Christian Garcia*

Containerized applications have exploded in popularity in recent years, due to their ease of deployment, reproducible nature, and speed of startup. Accordingly, container orchestration tools such as Kubernetes have emerged as resource providers and users alike try to organize and scale their work across clusters of systems. This paper documents some real-world experiences of building, operating, and using self-hosted Kubernetes Linux clusters. It aims at comparisons between Kubernetes and single-node container solutions and traditional multi-user, batch queue Linux clusters. The authors of this paper have experience first running traditional HPC Linux clusters with batch queues, and later virtual machines using technologies such as Openstack. Much of the experience and perspective below is informed by this perspective. We will also provide a use-case from a researcher who deployed on Kubernetes without being as opinionated about other potential choices.

9:30am –  
9:40am

**S Using Single Sign-On Authentication with Multiple Open OnDemand Accounts: A Solution for HPC Hosted Courses**

Pathable Platform

*Speakers: Ping Luo, Benjamin Evans, Tyler Trafford, Jay Kubeck, Thomas Langford, Kaylea Nelson, Andrew Sherman*

Open OnDemand (OOD) greatly lowers the barrier to entry to high performance computing (HPC) resources and facilitates usage for new and experienced users. Moreover, using OOD for courses with computational components enables lecturers and teaching assistants to focus on the course instead of on-boarding students. To these ends the Yale Center for Research Computing (YCRC) adopted OOD to make its advanced cyberinfrastructure more accessible to researchers and students across the university. While the use of single sign-on authentication in OOD makes HPC clusters easier to access, as implemented this authentication method limits users to a single HPC account per university-provided identity. The YCRC provides separate and temporary course-specific accounts to isolate course usage, which presented a challenge for supporting users

with pre-existing research allocations or who are participating in multiple courses. In this paper, we present an easy to implement and maintain solution that separates traffic for each course and maps a single user identity to multiple HPC accounts. This solution can be used for other situations other than courses; wherever one single sign-on identity should access multiple HPC accounts.

---

9:40am – 9:50am      **L Integrity Protection for Research Artifacts using Open Science Chain's Command Line Utility**      Pathable Platform  
*Speakers: Manu Shantharam, Subhashini Sivagnanam, Scott Sakai, Kai Lin*

Scientific data, its analysis, accuracy, completeness and reproducibility play a vital role in advancing science and engineering. Open Science Chain (OSC) is a cyberinfrastructure platform built using the Hyperledger Fabric (HLF) blockchain technology to address issues related to data reproducibility and accountability in scientific research. OSC preserves integrity of research datasets and enables different research groups to share datasets with the integrity information. Additionally, it enables quick verification of the exact datasets that were used for a particular published research and tracks its provenance. In this paper, we describe OSC's command line utility that will preserve the integrity of research datasets from within the researchers' environment or from remote systems such as HPC resources or campus clusters used for research. The python-based command line utility can be seamlessly integrated within research workflows and provides an easy way to preserve the integrity of research data in OSC blockchain platform.

---

9:40am – 9:50am      **S Common Resource Descriptions for Interoperable Gateway Cyberinfrastructure**      Pathable Platform  
*Speakers: Suresh Marru, Eric Franz, Dimuthu Wannipurage, Sudhakar Pamidighantam, John-Paul Navarro, Joe Stubbs, Daniel Mejia, Steve Black, Claire Stirm, Maytal Dahan, Marlon Pierce, Michael Zentner*

Science gateway projects face challenges utilizing the vast and heterogeneous landscape of powerful cyberinfrastructure available today, and interoperability across technologies remains poor. This interoperability issue leads to myriad problems: inability to bring multiple heterogeneous specialized resources together to solve problems where different resources are optimized for different facets of the problem; inability to choose from multiple resources on-the-fly as needed based on characteristics and available capacity; and ultimately a less than optimal application of nationally-funded resources toward advancing science. This paper presents version 1.0 of the Science Gateways Community Institute (SGCI) Resource Description Specification – a schema providing a common language for describing storage and computing resources utilized by science gateway technologies – as well as an Inventory API and software development kits for incorporating resource definitions into gateway projects. We discuss multiple gateway integration design options, with trade offs regarding robustness and availability. We detail the adoption to date of the SGCI Resource Specification by several prominent projects, including Apache Airavata, HUBzero®, Open OnDemand, Tapis, and XSEDE. The XSEDE adoption is worth highlighting explicitly as it has led to a new API within the XSEDE Information Services architecture which provides SGCI resource descriptions of all active XSEDE resources. Additionally, we show how the use of the SGCI Resource Specification provides interoperability across resource providers and projects that adopt it. Finally, as a proof of concept, we present a multi-step analysis that runs Quantum ESPRESSO and visualizes the energy band structures of a Gallium Arsenide (GaAs) crystal across multiple resource providers including the Halstead cluster at Purdue University and the Stampede2 supercomputer at TACC.

---

9:40am – 9:50am      **S System Integration of Neocortex, a Unique, Scalable AI Platform**      Pathable Platform  
*Speakers: Paola A. Buitrago, Nicholas A. Nystrom, Julian A. Uran*

To advance knowledge by enabling unprecedented AI speed and scalability, the Pittsburgh Supercomputing Center (PSC), a joint research center of Carnegie Mellon University and the University of Pittsburgh, in partnership with Cerebras Systems and Hewlett Packard Enterprise (HPE), has deployed Neocortex, an innovative computing platform that accelerates scientific discovery by vastly shortening the time required for deep learning training and inference, fostering greater integration of deep AI models with scientific workflows, and providing revolutionary hardware for the development of more efficient algorithms for artificial intelligence and graph analytics. Neocortex advances knowledge by accelerating scientific research, enabling development of more accurate models and use of larger training data, scaling model parallelism to unprecedented levels, and focusing on human productivity by simplifying tuning and hyperparameter optimization to create a transformative hardware and software platform for the exploration of new frontiers. Neocortex has been integrated with PSC's complementary infrastructure. This paper shares experiences, decisions, and findings made in that process. The system is serving science and engineering users via an early user access program. Valuable artifacts developed during the integration phase have been made available via a public repository and have been consulted by other deployments that have seen Neocortex as an inspiration.

---

9:50am –      **L Experiences in building a user portal for Expanse supercomputer**      Pathable Platform

10:00am

*Speakers: Subhashini Sivagnanam, Michael Norman, Shawn Strande, Mahidhar Tatineni, Martin Kandes, Scott Sakai, Dmitry Mishin, Mary Thomas, Christopher Irving*

A User Portal is being developed for NSF-funded Expanse supercomputer. The Expanse portal is based on the NSF-funded Open OnDemand HPC portal platform which has gained widespread adoption at HPC centers. The portal will provide a gateway for launching interactive applications such as MATLAB, RStudio, and an integrated web-based environment for file management and job submission. This paper discusses the early experience in deploying the portal and the customizations that were made to accommodate the requirements of the Expanse user community.

9:50am –  
10:00am

**S Managing Cloud networking costs for data-intensive applications by provisioning dedicated network links**

Pathable Platform

*Speakers: Igor Sfiligoi, Michael Hare, David Schultz, Frank Würthwein, Benedikt Riedel, Tom Hutton, Steve Barnet, Vladimir Brik*

Many scientific high-throughput applications can benefit from the elastic nature of Cloud resources, especially when there is a need to reduce time to completion. Cost considerations are usually a major issue in such endeavors, with networking often a major component; for data-intensive applications, egress networking costs can exceed the compute costs. Dedicated network links provide a way to lower the networking costs, but they do add complexity. In this paper we provide a description of a 100 fp32 PFLOPS Cloud burst in support of IceCube production compute, that used Internet2 Cloud Connect service to provision several logically-dedicated network links from the three major Cloud providers, namely Amazon Web Services, Microsoft Azure and Google Cloud Platform, that in aggregate enabled approximately 100 Gbps egress capability to on-prem storage. It provides technical details about the provisioning process, the benefits and limitations of such a setup and an analysis of the costs incurred.

9:50am –  
10:00am

**S Skyway: A Seamless Solution for Bursting Workloads from On-Premises HPC Clusters to Commercial Clouds**

Pathable Platform

*Speakers: Yuxing Peng, Jonathan Skone, Callista Christ, Hakizumwami Runesha*

Interest in cloud computing within the HPC community continues to grow, but some barriers to wide-spread adoption remain due to aspects such as cost, support of low network latency applications, and data management. It has, however, tremendous potential to enhance HPC through its high scalability and elasticity. HPC workloads and hardware technologies are rapidly changing and becoming increasingly diversified. However, the configuration and capacity of on-premises resources, once deployed, are not easily changeable. In contrast, a hybrid infrastructure, with the ability of “burst-to-cloud”, can combine advantages from the on-premises and cloud resource spaces. The Research Computing Center (RCC) at The University of Chicago has developed such a solution, named Skyway, that incorporates multi-cloud computing resources as elastic extensions of its on-premises HPC infrastructure. An in-house system software is also developed to interface with the SLURM scheduler and cloud SDKs, that can burst HPC workloads to the cloud, providing users a seamless experience when interacting with both on-premises and cloud systems. Skyway currently can interface with Amazon AWS and Google GCP, and is being extended to other cloud providers (Microsoft Azure and Oracle), demonstrating the flexibility of using various cloud providers with the on-premises resources.

10:00am –  
10:10am

**L A Vision for Science Gateways: Bridging the Gap and Broadening the Outreach**

Pathable Platform

*Speakers: Sandra Gesing, Steven Brandt, Shannon Bradley, Mark Potkewitz, Kerk Kee, Noreen Whysel, Mark Perri, Sean Cleveland, Jack Smith, Annelie Rugg*

The future for science gateways warrants exploration as we consider the possibilities that extend well beyond 'science' and high-performance computing into new interfaces, applications and user communities. In this paper, we look retrospectively at the successes of representative gateways thus far. This serves to highlight existing gaps gateways need to overcome in areas such as accessibility, usability and interoperability, and in the need for broader outreach by drawing insights from technology adoption research. We explore two particularly promising opportunities for gateways - computational social sciences and virtual reality – and make the case for the gateway community to be more intentional in engaging with users to encourage adoption and implementation, especially in the area of educational usage. We conclude with a call for focused attention on legal hurdles in order to realize the full future potential of science gateways. This paper serves as a roadmap for a vision of science gateways in the next ten years.

10:00am –  
10:10am

**S Doing more with less: Growth, improvements, and management of NMSU's computing capabilities**

Pathable Platform

*Speakers: Strahinja Trecakov, Nicholas Von Wolff*

Deployed in 2015, Discovery is New Mexico State University's commonly-available High-Performance Computing (HPC) cluster. The deployment of Discovery was initiated by Information and Communication Technologies (ICT) employees from the Systems Administration group who wanted to help researchers run their computations on a more powerful system than one they had sitting in their offices. Over the years, the cluster has grown 6 times, and as of March 2021 has 52 compute nodes, 1480 CPU cores, 17 Terabytes of RAM, 30 GPUs, and 1.5 Petabytes of usable storage. Discovery's hardware is acquired using a combination of university funds, condo-model based funds, and grant funds, causing Discovery to be a heterogeneous system containing several CPU generations. This paper discusses our growth and administration experiences on this heterogeneous system, as well as our outreach and contribution to the HPC community.

- 
- 10:00am – 10:10am      **S Ookami: Deployment and Initial Experiences** Pathable Platform  
*Speakers: Benjamin Michalowicz, Eric Raut, Yan Kang, Tony Curtis, Andrew Burford, Alan Calder, David Carlson, Barbara Chapman, Firat Coskun, Catherine Feldman, Robert Harrison, Eva Siegmann, Daniel Wood, Robert DeLeon, Mathew Jones, Nikolay Simakov, Dossay Oryspayev, Joseph White*  
Ookami is a computer technology testbed supported by the United States National Science Foundation. It provides researchers with access to the A64FX processor developed by Fujitsu in collaboration with RIKEN for the Japanese path to exascale computing, as deployed in Fugaku, the fastest computer in the world. By focusing on crucial architectural details, the ARM-based, multi-core, 512-bit SIMD-vector processor with ultrahigh-bandwidth memory promises to retain familiar and successful programming models while achieving very high performance for a wide range of applications. We review relevant technology and system details, and the main body of the paper focuses on initial experiences with the hardware and software ecosystem for micro-benchmarks, mini-apps, and full applications, and starts to answer questions about where such technologies fit into the NSF ecosystem.
- 
- 10:10am – 10:20am      **L The LROSE Science Gateway: One-Stop Shop for Weather Data, Analysis, and Expert Advice** Pathable Platform  
*Speakers: Brenda Javornik, Jennifer C. DeHart, Hector Santiago III*  
Nexrad data along with software to convert between binary formats, perform quality control, analyze, and visualize the data are all public and open for access and download. The missing pieces are the knowledge of how to use the available components with reproducible results. A science gateway is a web-based platform to bring all the components together for a novice to learn from experts, and for expert researchers to customize the tools for science.
- 
- 10:10am – 10:20am      **S BenchTool: a framework for the automation and standardization of HPC performance benchmarking** Pathable Platform  
*Speakers: Matthew Cawood, Stephen Lien Harrell, Todd Evans*  
The procurement of large and complex HPC systems necessitates an extensive planning process. This process should ideally include the independent performance evaluation and validation of available solutions to maximize system throughput for expected workloads. This assessment procedure typically requires the arduous task of compiling, configuring and running multiple application benchmarks on every competing architecture, where any combination of hardware and software may produce unique idiosyncrasies and time consuming pitfalls. In addition, the possible lack of standardization and communication within a research group may produce undesirable differences in evaluation methodology, affecting consistency between results. The BenchTool utility was developed to provide a framework to automate the compilation of applications, the execution of benchmarks using said applications, and the collection of results including supplementary provenance data.
- 
- 10:10am – 10:20am      **S INAM: Cross-stack Profiling and Analysis of Communication in MPI-based Applications** Pathable Platform  
*Speakers: Dhableswar Panda, Hari Subramoni, Pouya Kousha, Arpan Jain, Aamir Shafi, Kamal Raj Sankarapandian Dayala Ganesh Ram, Mansa Kedia, Trey Dockendorf, Heechang Na, Karen Tomko*  
Understanding the full-stack performance trade-offs and interplay among HPC applications, MPI libraries, the communication fabric, and the job scheduler is a challenging endeavor. Unfortunately, existing profiling tools are disjoint and only focus on profiling one or a few levels of the HPC stack limiting the insights they can provide. In this paper, we propose a standardized approach to facilitate near real-time, low overhead performance characterization, profiling, and evaluation of communication of high-performance communication middleware as well as scientific applications using a cross-stack approach by INAM. The profiling capabilities are supported in two modes of with and without modifications to the application depending on the scope of the profiling session. We design and implement our designs using an MPI\_T-based standardized method to obtain near real-time insights for MPI applications at scales of up to 4,096 processes with less than 5% overhead. Through experimental evaluations of increasing batch size for DL training, we demonstrate novel benefits of INAM for

cross-stack communication analysis in real-time to detect bottlenecks and resolve them, achieving up to 3.6x improvements for the use-case study. The proposed solutions have been publicly released with the latest version of INAM and currently being used in production at various HPC supercomputers.

- 
- 10:20am – 10:30am      **L Defining Performance of Scientific Application Workloads on the AMD Milan Platform**      Pathable Platform  
*Speakers: Geoffrey Lentner, Amiya Maji, Sam Weekly, Tsai-Wei Wu, Stephen Lien Harrell, Alex Younts, Zoey Mertes, Preston Smith, Xiao Zhu*  
Understanding the capabilities of new architectures is key to informing system purchases and good long-term ROI for cluster installations. The newest AMD architecture, Milan, has become available first on Microsoft Azure and we use this opportunity to measure the performance of this 3rd generation AMD EPYC processor. In this paper single node performance is gathered for seven popular scientific applications and benchmark test-suites. Quantitative comparisons are carried out between two independent platforms, Milan and its architectural predecessor Rome, for performance evaluations. Our results have shown that the changes in the Milan architecture have improved performance and met our projections.
- 
- 10:20am – 10:30am      **S Simulation vs Actual Walltime Correction in a Real Production Resource-Constrained HPC**      Pathable Platform  
*Speakers: Jessi Christa Rubio, Aira Villapando*  
Today's increase in computational resource demand requires robust job scheduling performance in shared HPC systems, which have varying workload characteristics. We applied a previous theoretical work on a walltime corrective scheduling to full-scale production-level HPC. The resulting performance was analyzed in comparison with an old cluster from which the theoretical work was modeled and with a simulation modeled after the current implementation. The results rendered us substantial insights in adjusting the scheduling policies accordingly for effective resource management that similar HPC clusters could benefit from.
- 
- 11:00am – 12:00pm      **B Marketing and Promotion of Scientific Software for Sustainability Purposes**      Pathable Platform  
*Speakers: Alan Chalker, Enis Afgan, Claire Stirm, Michael Zentner, Sandra Gesing, Vas Vasiliadis*  
As many in the scientific software community struggle with developing, maintaining, and sustaining their projects, a key element required is generating awareness of the project, its benefits, and methods by which it can be adopted. All of these are key for sustainability, and yet have the elements of marketing and promotion which are often not in the skillsets of project personnel. The Science Gateways Community Institute is routinely approached by software projects seeing science gateways as one possible path toward sustainability. However, science gateways alone are only one element on that path and can easily fail without marketing and promotion. This Birds-of-a-Feather session will provide a forum for learning about existing approaches to marketing and promoting scientific software for purposes of increasing its chances of sustainability, and generating new ideas for services and techniques that may be adopted by the scientific software community.
- 
- 11:00am – 12:00pm      **B Quantifying the Research Computing and Data Professional Community for Attracting, Retaining, and Diversifying RCD Professionals**      Pathable Platform  
*Speakers: Christina Maimone, Scott Yockel, Amy Neeser, Ashley Stauffer, Timothy Middelkoop, Kimberly Grasch*  
As facilitating research continues to increase and broaden, it is important to understand the state of Research Computing and Data (RCD) staffing across the country through more than anecdotal reports. Most institutions are isolated from one another with regards to sharing basic information on staffing roles, position types, and pay grades. Participants in the National Science Foundation Virtual Workshop on the Research Innovation Workforce for Cyberinfrastructure identified recruiting and sustaining a diverse and inclusive workforce as a key challenge for the future of the RCD field, yet no systematic data is currently available on the current composition of the RCD workforce. A breakout group from this workshop created a joined the CaRCC RCD Professionalization Working Group, and set out to better understand quantitatively the state of RCD staffing across the United States by designing, testing, and implementing an RCD workforce survey tool and using that to conduct a national survey. The main objective of this group is to conduct a US national survey of RCD professionals to capture specific information on their personal background, current position, and outlook on RCD as a profession. During the Birds of Feather, the working group will provide an overview of the current state of this work and solicit feedback from the audience on a number of topics.
- 
- 11:00am – 12:00pm      **B Science Gateways and HPCs: Next Generation Access**      Pathable Platform  
*Speakers: Suresh Marru, Marlon Pierce, Sudhakar Pamidighantam, Jim Basney, Maytal Dahan, Joe Stubbs, Mark Krenz*  
This proposed Birds of a Feather (BOF) session will discuss policies and best practices for High Performance

Computing (HPC) centers to enable remote access for science gateways and related cyberinfrastructure. The BOF will also promulgate the efforts of the Science Gateways Next Generation HPC Access Working Group, co-organized by the Science Gateways Community Institute (SGCI, sciencegateways.org) and Trusted CI (trustedci.org). This working group is an open-community effort with a charter to draft a set of general recommendations for academic research computing centers and science gateway providers to securely enable science gateways. The working group will base its recommendations on a broad understanding of HPC center cybersecurity requirements and concerns, and on a classification and description of common science gateway access mechanisms.

---

11:00am –  
12:00pm

**B Towards Inclusive Terminology in Advanced Research Computing**

Pathable Platform

*Speakers: Susan Mehringer, Jay Alameda, Sergiu Sanielevici, Linda Akli, Katharine Cahill, Stephen Deems, John Holly, Boswell Hutson, Susan Lindsey, Deb Nigra, Robert Panoff, Robert Whitten*

In the past year, many projects and organizations active in advanced research computing have become acutely aware of the need to ensure that the language they use in formal and informal writing and speech, from technical standards, presentations and educational materials to the language used in the workplace, is free from terminology that is contrary to their commitment to foster an inclusive environment for all community members.

In particular, the Extreme Science Engineering and Discovery Environment (XSEDE) project is addressing these concerns by forming a Terminology Task Force (TTF) to review, address, and define processes to eliminate offensive terms in their materials. Members of the XSEDE TTF propose to organize this BoF session to briefly review their progress in order to stimulate discussion and exchange of best practices among interested PEARC21 participants. Several groups facing this complex problem are being invited to participate in the discussion by sharing their own approaches and experiences.

---

**A** and Education   **L** Applications and Software   **B** BoF   **D** Diversity   **P** Panels

**S** Systems and System Software   **T** Training   **R** Tutorial   **W** Workforce Development   **K** Workshop

---

**JULY 22 • THURSDAY**

---

- 8:00am – 8:10am   **L** **Anomaly Detection in Scientific Workflows using End-to-End Execution Gantt Charts and Convolutional Neural Networks** Pathable Platform  
*Speakers: Patrycja Krawczuk, George Papadimitriou, Shubham Nagarkar, Mariam Kiran, Ewa Deelman, Anirban Mandal*  
Fundamental progress towards reliable modern science depends on accurate anomaly detection during application execution. In this paper, we suggest a novel approach to tackle this problem by applying Convolutional Neural Network (CNN) classification methods to high-resolution visualizations that capture the end-to-end workflow execution timeline. Subtle differences in the timeline reveal information about the performance of the application and infrastructure's components. We collect 1000 traces of a scientific workflow's executions. We explore and evaluate the performance of CNNs trained from scratch and pre-trained on ImageNet [7]. Our initial results are promising with over 90% accuracy.
- 
- 8:00am – 8:10am   **S** **Securing CHEESEHub: A Cloud-based, Containerized Cybersecurity Education Platform** Pathable Platform  
*Speakers: Mike Lambert, Rajesh Kalyanam, Rob Kooper, Baijian Yang*  
The Cyber Human Ecosystem for Engaged Security Education (CHEESEHub) is an open web platform that hosts community contributed containerized demonstrations of cybersecurity concepts. In order to maximize flexibility, scalability, and utilization, CHEESEHub is currently hosted in a Kubernetes cluster on the Jetstream academic cloud. In this short paper, we describe the security model of CHEESEHub and specifically the various Kubernetes security features that have been leveraged to secure CHEESEHub. This ensures that the various cybersecurity exploits hosted in the containers cannot be misused, and that potential malicious users of the platform are cordoned off from impacting not just other legitimate users, but also the underlying hosting cloud. More generally, we hope that this article will provide useful information to the research computing community on a less discussed aspect of cloud deployment: the various security features of Kubernetes and their application in practice.
- 
- 8:00am – 9:30am   **P** **Diversity in the student pipeline and professional staff: challenges, success stories, and resources** Pathable Platform  
*Speakers: Tabitha Samuel, Craig Stewart, Gwen Jacobs, Paul Rad, Marisa Brazil, Jeaime H Powell, Gabriella Arellano*  
In 2021 most people understand that systemic racism and deeply engrained social biases are real and sometimes seemingly intractable problems in the US in general and in the advanced computing community in particular. However, many people lack knowledge about what they can do to change the situation we face today. The purpose of this panel will be to bring together recognized leaders in Advanced Research Computing (ARC) and hear from them about their own experiences cultivating diversity in their workforces, retaining that diversity, and offering advice on resources that all can use to create a more diverse, vibrant, inclusive, and effective HPC community for the future.
- 
- 8:00am – 9:30am   **P** **The NSF Computing Ecosystem: Category 1 and 2 Resources for Accelerating Research for the Community.** Pathable Platform  
*Speakers: Shawn Strande, Paola A. Buitrago, Amit Majumdar, Sergiu Sanielevici, Shawn Brown, David Hancock, X. Carol Song, Robert Harrison, William Gropp*  
The landscape of computational research has become quite complex, with a large, broad, and diverse research community requiring a variety of different computational capabilities. To address this growing need, the National Science Foundation has awarded eight new resources (Expanse, Bridges-2, Anvil, Delta, Jetstream2, Neocortex, Voyager, and Ookami) that are becoming available to the community throughout 2021, with three of the platforms being test-bed architectures for providing unique never-before available resources to the community. In this panel, we will introduce each of these architectures to the attendees, outlining the overall goals of the resource, description of the computational and data components of the resources, the use-cases for the resource, and any unique aspects of the resources that may be useful for the research community to know.
-

8:10am – 8:20am	L <b>Collecting and analyzing smartphone sensor data for health</b>	Pathable Platform
	<p><i>Speakers: Justin Drake, Karl Schulz, Radek Bukowski, Kelly Gaither</i></p> <p>Modern smartphones contain a collection of energy-efficient sensors capable of capturing the device's movement, orientation, and location as well characteristics of its external environment (e.g. ambient temperature, sound, pressure). When paired with peripheral wearable devices like smart watches, smartphones can also facilitate the collection/aggregation of important vital signs like heart rate and oxygen saturation. Evidence suggests that signatures of health and disease, or digital biomarkers, exist within the heterogeneous, temporally-dense data gathered from smartphone sensors and wearable devices that can be leveraged for medical applications. Here we discuss our recent experiences with deploying an open-source, cloud-native framework to monitor and collect smartphone sensor data from a cohort of pregnant women over a period of one year. We highlight two open-source integrations into the pipeline we found particularly useful: 1) a dashboard--built with Grafana and backed by Graphite--to monitor and manage production server loads and data collection metrics across the study cohort and 2) a back-end storage solution with InfluxDB, a multi-tenant time series database and data exploration ecosystem, to support biomarker discovery efforts of a multidisciplinary research team.</p>	
8:10am – 8:20am	S <b>ColdFront: Resource Allocation Management System</b>	Pathable Platform
	<p><i>Speakers: Andrew Bruno, Dori Sajdak</i></p> <p>ColdFront is an open source resource allocation management system, designed to provide a central portal for administration, reporting, and measuring scientific impact of cyberinfrastructure resources. It enables managing access to a diverse set of resource types across large groups of users and provides a rich set of extensible meta data for comprehensive reporting and integration with legacy systems. In this paper we introduce ColdFront, describe three main stakeholders, present the features and integrations currently implemented, and make the case for why ColdFront is vital to enabling research and reducing time-to-science. ColdFront is freely available at <a href="https://github.com/ubccr/coldfront">https://github.com/ubccr/coldfront</a>.</p>	
8:20am – 8:30am	L <b>RESIF 3.0: Toward a Flexible &amp; Automated Management of User Software Environment on HPC facility</b>	Pathable Platform
	<p><i>Speakers: Sebastien Varrette, Emmanuel Kieffer, Frederic Pinel, Ezhilmathi Krishnasamy, Sarah Peter, Xavier Besseron, Hyacinthe Cartiaux</i></p> <p>HPC is increasingly identified as a strategic asset and enabler to accelerate the research and the business performed in all areas requiring intensive computing and large-scale Big Data analytic capabilities. The efficient exploitation of heterogeneous computing resources featuring different processor architectures and generations, coupled with the eventual presence of GPU accelerators, remains a challenge. The University of Luxembourg operates since 2007 a large academic HPC facility which remains one of the reference implementation within the country and offers a cutting-edge research infrastructure to Luxembourg public research. The HPC support team invests a significant amount of time in providing a software environment optimised for hundreds of users, but the complexity of HPC software was quickly outpacing the capabilities of classical software management tools. Since 2014, our scientific software stack is generated and deployed in an automated and consistent way through the RESIF framework, a wrapper on top of EasyBuild and Lmod meant to efficiently handle user software generation. A large code refactoring was performed in 2017 to better handle different software sets and roles across multiple clusters, all piloted through a dedicated control repository. With the advent in 2020 of a new supercomputer featuring a different CPU architecture, and to mitigate the identified limitations of the existing framework, we report in this state-of-practice article RESIF 3.0, the latest iteration of our scientific software management suit now relying on streamline EasyBuild. It permitted to reduce by around 90% the number of custom configurations previously enforced by specific Slurm and MPI settings, while sustaining optimised builds coexisting for different dimensions of CPU and GPU architectures. The workflow for contributing back to the EasyBuild community was also automated and a current work in progress aims at drastically decrease the building time of a complete software set generation. Overall, most design choices for our wrapper have been motivated by several years of experience in addressing in a flexible and convenient way the heterogeneous needs inherent to an academic environment aiming for research excellence. As the code base is available publicly, and as we wish to transparently report also the pitfalls and difficulties met, this tool may thus help other HPC centres to consolidate their own software management stack.</p>	
8:20am – 8:30am	S <b>CloudBank: Managed Services to Simplify Cloud Access for Computer Science Research and Education</b>	Pathable Platform
	<p><i>Speakers: Michael Norman, Shawn Strande, Vince Kellen, Shava Smallen, Brian DeMeulle, Ed Lazowska, Naomi Alterman, Rob Fatland, Sarah Stone, Amanda Tan, Katherine Yelick, Eric Van Dusen, James Mitchell</i></p> <p>CloudBank is a cloud access entity founded to allow the computer science research and education communities</p>	

to harness the profound computational potential of public cloud platforms. It does so by delivering a set of managed services designed to alleviate common points of friction associated with cloud adoption, serving as an integrated service provider to the research and education community. These services include front line help desk support, cloud solution consulting, training, account management, automated billing, and cost monitoring support. It functions using a multi-cloud pay-by-use billing model, and aims to serve the spectrum of cloud users from novice to advanced.

---

8:30am – L **Research Cloud Bazaar** Pathable Platform

8:40am

*Speakers: Michael Lau, Stuti Trivedi, Zhenhua He, Tri Pham, Dhruva Chakravorty, Lisa Perez*  
Research workflows will benefit from a hybrid computing environment that offers seamless integration between on-campus and off-campus cloud resources. Commercial and Federal clouds offer researchers access to novel computing platforms that may not be available on their campus, offering opportunities to improve workflows and reduce the time to research. The large number of cloud offerings, however, makes cost management, and workflow transitions to appropriate platforms challenging. Successfully mapping workflows from on-campus resources to the cloud, and leveraging the available cost structures to find economical cost models are critical steps to enabling researcher access to this vast resource. To ameliorate these concerns, here we introduce the Research Computing Bazaar (RCB) software application for resource mapping and cost estimation. RCB is a software-as-a-service platform is an elastic, scalable, and fault tolerant system. It is developed using actual data from research computing workloads, and can be easily configured to be used by users or system administrators in Slurm-based on-premise computing environments. In this pilot, we inform researchers about opportunities to leverage flexible workload orchestration in managing cloud costs on a major cloud service provider. An extension into predictive capacities with machine learning mechanisms is being developed.

---

8:30am – S **Expanse: Computing without Boundaries** Pathable Platform

8:40am

*Speakers: Manu Shantharam, Shawn Strande, Mahidhar Tatineni, Dmitry Mishin, Haisong Cai, Wayne Pfeiffer, Christopher Irving, Amit Majumdar, Robert Sinkovits, Mike Norman, Nicole Wolter, Trevor Cooper, Ilkay Altintas, Marty Kandes, Ismael Perez, Mary Thomas, Subhashini Sivagnanam, Thomas Hutton*

We describe the design motivation, architecture, deployment, and early operations of Expanse, a 5 Petaflop, heterogenous HPC system that entered production as an NSF-funded resource in December 2020 and will be operated on behalf of the national community for five years. Expanse will serve a broad range of computational science and engineering through a combination of standard batch-oriented services, and by extending the system to the broader CI ecosystem through science gateways, public cloud integration, support for high throughput computing, and composable systems. Expanse was procured, deployed, and put into production entirely during the COVID-19 pandemic, adhering to stringent public health guidelines throughout. Nevertheless, the planned production date of October 1, 2020 slipped by only two months, thanks to thorough planning, a dedicated team of technical and administrative experts, collaborative vendor partnerships, and a commitment to getting an important national computing resource to the community at a time of great need.

---

8:40am – L **Investigating the Genomic Distribution of Phylogenetic Signal with CloudForest** Pathable Platform

8:50am

*Speakers: Reid Wagner, Benjamin S. Touns, Zhifeng Deng, Thomas McGowan, Kyle A. Gallivan, Jeremy M. Brown, James C. Wilgenbusch*

A central focus of evolutionary biology is inferring the historical relationships among species and using this context to learn about how evolution has shaped diverse organisms. These historical relationships are represented by phylogenetic trees, and the methods used to infer these trees have been an active area of research for several decades. Despite this attention, phylogenetic workflows have changed little, even though extraordinary advances have occurred in the scale and pace at which genomic data have been collected in the past 20 years. Modern phylogenomic datasets have also raised fascinating new questions. Why do different parts of a genome often support different relationships among species? How are these different signals distributed across chromosomes? We developed a new computational framework, CloudForest, to tackle such questions. CloudForest is flexible, efficient, and tightly integrates a diverse set of tools. Here, we briefly describe the architecture of CloudForest, including the advantages it provides, and use it to investigate the distribution of phylogenetic signal along the entire X chromosome of 24 cat species.

---

8:40am – S **Results from a second longitudinal survey of academic research computing and data center usage: expenditures, utilization patterns, and approaches to return on investment** Pathable Platform

8:50am

*Speakers: Alan Chalker, Sharon Broude Geva, Curt Hillegas, Donald Petravick, Craig Stewart, Alan Sill*  
Availability of cloud-based resource delivery modes is transforming many areas of computing. Academic research computing and data (RCD) support largely remains based on on-premises delivery and has adopted

commercial clouds more slowly than the private sector for a variety of stated reasons including factors related to cost efficiency, return on investment, institutional requirements, high costs for bulk commercial cloud computing usage, and funding patterns. Other factors involved in selection of computing resource delivery modes include capabilities and applications that are available only in or best adapted to specific computing environments. It is important for the higher education and research communities to be able to learn from each other as institutions and individuals to make optimum use of appropriate modes of delivery for RCD resources. This paper reports an overview of results from the second annual community-wide survey conducted by the Coalition for Advanced Scientific Computation on patterns of funding, usage, and return on investment for academic research computing and data resources. The results show that on-premises delivery continues to remain the preferred mode for RCD resources for most responding institutions as found in the first survey, but that commercial cloud usage is beginning to be reported for production use by a small number of respondents to the survey. Reasons for these preferences are further explored in the survey and initial high-level results are reported here.

---

8:50am –  
9:00am

**L Powering Plasma Confinement Simulations: from Classic to Photorealistic Visualizations**

*Speakers: Diego Jiménez, Luis Campos-Duarte, Esteban Meneses, Ricardo Solano-Piedra, Esteban Pérez, Ernesto Rivera-Alvarado, V.I. Vargas* Pathable Platform

As the world moves away from traditional energy sources based on fossil fuels, several alternatives have been explored. One promising clean energy source is nuclear fusion. The fusion of hydrogen atoms may provide generous consumable energy gains. However, nuclear fusion reactors are not ready to become a productive mechanism yet. To accelerate the required breakthroughs in that community, numerical simulations and scientific visualizations over high-performance computing systems are mandatory. The results from the simulations and a proper display of the data are key to design and tune up nuclear fusion reactors. We explore the scientific visualization of plasma confinement in this paper, presenting two different dimensions. First, we revisit how visualizations help scientists understand the phenomena behind the fundamental processes of nuclear fusion. Second, we explore how visualization may also work as scientific communication tools for the general public to grasp the impact of this endeavor. We introduce a computer-graphics model that uses the output of numerical simulations to create visually plausible images of plasma confinement. The model is based on a combination of computer graphics techniques implemented on a ray-tracing framework.

---

8:50am –  
9:00am

**S A Scalable Cloud-based Architecture to Deploy JupyterHub for Computational Social Science Research** Pathable Platform

*Speakers: Da Li, Runchao Jiang, Robert Pyke*

With the increasing popularity of computational approaches to conduct social science research, building a scalable and efficient computing platform has become a topic of interest for academia to empower research labs and institutions to analyze large-scale data. While social science researchers have been very excited about the advancement of emerging technologies in big data, deep learning, computer vision, network analysis, etc., they are also constrained by the available computing resources to analyze data. This paper describes a scalable solution to deploy JupyterHub for computational social science research on the cloud. We use a reference architecture on AWS to walk through the design principles and details. Our architecture has helped facilitate several collaborations between Facebook and academia. The case study (Facebook Open Research and Transparency platform) shows that our architecture, using technologies like containerization and serverless computing, can support thousands of users to analyze web-scale datasets.

---

9:00am –  
9:10am

**L DELTA-Topology: A Science Gateway for Experimental and Computational Chemical Data Analysis using Topological Models** Pathable Platform

*Speakers: Eric Coulter, Sudhakar Pamidighantam, Marcus Christie, Aurora Clark*

Chemical data are diverse and complex, are obtained from experimental and computational modeling, and may encode large degrees of freedom of movement of particles such as whole assemblies, clusters, molecules, atoms, and even nuclei and electrons. To derive knowledge from this data requires analyses using a variety of techniques including approximation, dimensionality reduction, principal component analysis, and topological analysis. In this manuscript we describe the DELTA Science Gateway that integrate several types of mathematical and topological analysis software for chemical data analysis. The focus is on energy landscape data derived from experimental and computational modeling techniques towards understanding the principals involved in structure and function of molecular moieties, particularly in delineating the mechanism of catalytic activity. The gateway design, creation and production deployment will be discussed. The DELTA gateway is hosted under the SciGaP project at Indiana University powered by Apache Airavata gateway middleware framework. The gateway provides an integrated infrastructure for simulations and analysis on XSEDE and IU HPC resources and interactive visualization through locally deployed VNC client and a JupyterHub deployed on the XSEDE Jetstream cloud using virtual clusters. The gateway provides intuitively simple user interfaces for

providing simulation input data, combines available model data, and enables users to set up and execute the simulation/analyses at the HPC systems.

---

9:00am – 9:10am S **Bridges-2: A Platform for Rapidly-Evolving and Data Intensive Research** Pathable Platform  
*Speakers: Shawn T. Brown, Paola Buitrago, Edward Hanna, Sergiu Sanielevici, Nicholas A. Nystrom, Robin Scibek*

Today's landscape of computational science is evolving rapidly, with a need for new, flexible, and responsive supercomputing platforms for addressing the growing areas of artificial intelligence (AI), data analytics (DA) and convergent collaborative research. To support this community, we designed and deployed the Bridges-2 platform. Building on our highly successful \Bridges\ supercomputer, which was a high-performance computing resource supporting new communities and complex workflows, Bridges-2 supports traditional and nontraditional research communities and applications; integrates new technologies for converged, scalable high-performance computing (HPC), AI, and data analytics; prioritizes researcher productivity and ease of use; and provides an extensible architecture for interoperability with complementary data intensive projects, campuses, and clouds. In this report, we describe Bridges-2's hardware and configuration, user environments, and systems support and present the results of the successful Early User Program.

---

9:10am – 9:20am S **Identifying Research Collaboration Challenges for the Development of a Federated Infrastructure Response** Pathable Platform

*Speakers: Maureen Dougherty, James Barr von Oehsen, Michael Zink*

In this paper we present the key collaboration challenges and recommendations identified by targeted research communities during the Eastern Regional Network (ERN) Architecture and Federation Virtual Workshop, for validation of the base design of the ERN Federated OpenCI Labs collaborative infrastructure model. The workshop was designed to stimulate open discussions surrounding key aspects of collaborative scientific research and workflows. A brief summary of the key data gathered is provided here. The findings from this workshop have led to a re-evaluation of the design of ERN Federated OpenCI Labs infrastructure.

---

9:20am – 9:30am S **Experiences Migrating the Agave Platform to A Kubernetes Native System on the Jetstream Cloud** Pathable Platform

*Speakers: Steven Brandt, Rion Dooley, Karen Liang, Eric Tanner*

The Agave Platform is a Science-as-a-Service (SaaS) platform for reproducible science. The current production deployments of the platform are deployed and managed using Ansible and Docker Compose. While capable, this has historically led to operational complexity for those adopting the platform. Over the last year, we have worked to migrate the platform to a Kubernetes native deployment. In this paper we discuss our experiences evolving the platform, its architecture, and getting the most out of the Jetstream cloud.

---

9:45am – 10:45am B **Science DMZ Experiences: Sharing Considerations and Approaches on Deploying and Maintaining a Science DMZ** Pathable Platform

*Speakers: David Wheeler, Richard Angeletti, Kathy Benninger, Thomas Hutton, Scott Sakai, Michael Lambert*

A Science DMZ is a part of the network specifically built to facilitate the high speed file transfers needed for many of today's high performance science applications and instruments. Built near the high performance network edge, these networks areas are unlike the typical enterprise or campus networks and often are new to the cyberinfrastructure engineers deploying them. The hardware, configurations, and security policies must be taken into consideration for their planning, deployment, and operations. This can result in a scalable and performant Science and Engineering (S&E) experience capable of addressing the security posture of the organization while still supporting improved time to discovery for the S&E users.

The goal of this session is to bring together these cyberinfrastructure engineers who are working on building or maintaining Science DMZs to discuss the various aspects of their deployments for the collective benefit of the group.

---

9:45am – 10:45am B **SOLAR Consortium: Accelerated Ray Tracing for Scientific Simulations** Pathable Platform

*Speakers: Paul Navratil, Estelle Dirand, Christiaan Gribble, Pascal Grosset*

Many physical simulations incorporate vector mathematics to model phenomena such as radiative transfer and to compute behavior such as particle advection. Hardware-optimized ray tracing engines, tuned by processor manufacturer engineers, can accelerate simulation critical sections that depend on ray-based traversals and intersections. This BOF will serve as a venue for developers and users of simulations, visual analysis codes, and

ray tracers to discuss interfaces, capabilities, and performance. This meeting will continue the conversation toward standardization of ray tracer interfaces to facilitate their expanding role throughout scientific workflows, including data evaluation, insight formulation, discovery communication, and presentation-quality artifact generation.

---

9:45am –  
10:45am

**B Tapis User Meeting**

Pathable Platform

*Speakers: Anagha Jamthe, Joe Stubbs, Julia Looney, Richard Cardone, Steve Black, Sean Cleveland, Joseph Meiring*

This session will provide an overview of the Tapis API Platform, an NSF-funded project for reproducible, distributed computational research, to current and potential users. Tapis offers a set of authentication, authorization, data transfer, job management, and execution services that can span multiple data centers and manage batch, interactive and streaming jobs. More than 15 active, funded projects across a wide range of domains of science and engineering rely on Tapis to run 100,000s of computational jobs for 1000s of users to accomplish their research objectives. The Tapis team is completing the second year of its five year NSF grant and will use this session to provide an update on the current status of the project and discuss the roadmap for the third year. The session will include a mix of short presentations and demos from the core Tapis team followed by Q&A sessions with participation from the audience.

---

9:45am –  
10:45am

**B The Connect.Cyberinfrastructure Portal (formerly known as the Cyberteam Portal): A Shared, yet Independent, Platform for Community Collaboration - Year 1 Update**

Pathable Platform

*Speakers: Dana Brunson, Shelley Knuth, Torey Battelle, Thomas Cheatham, Douglas Jennewein, Bj Lougee, Julie Ma, Timothy Middelkoop, John Goodhue, James Griffioen, Kaylea Nelson, Andrew Sherman, Dhruva Chakravorty, Christopher Simmons, Bhushan Chitre, Adrian Del Maestro, Anthony Elam, Vikram Gazula, Bruce Segee, Scott Valcourt, Ralph Zottola*

The Connect.Cyberinfrastructure.org (Cnct.CI) Portal, originally known as the Cyberteam Portal, was developed to support the management of project workflows and to capture project results for the Northeast Cyberteam (NECT). Launched in 2017, the NECT is an NSF-sponsored program that aims to make regional and national cyberinfrastructure more readily accessible to researchers in small and mid-sized institutions in northern New England by providing research computing facilitation support and aggregating access to knowledge resources. More recently, the Cnct.CI Portal has expanded to provide support for a variety of programs in the Research Computing and Data ecosystem, creating opportunities for collaboration and leveraging a consistent, cohesive approach to common challenges.

As reported at our BoF at PEARC20, a pilot was launched in July 2020 to enable six additional Cyberteam programs to explore the use of the Portal as a management tool for their related programs. Now, one year since launching the expansion effort, CAREERS (Cyberteam to Advance Research and Education in Eastern Regional States) and TRECIS (Texas Research and Education Cyberinfrastructure Services) programs are using the Portal in daily operation to manage their communities; and a new community, the XSEDE (Extreme Science and Engineering Discovery Environment) Campus Champions have funded a new view of the Portal which will be used to manage participants and streamline onboarding in this mainstay community of practice. The Great Plains CyberTeam, Rocky Mountain Advanced Computing Consortium (RMACC), SWEETER, and Kentucky Cyberteam programs are at varying stages of adoption of the platform to showcase projects and connect their communities. We will share current status of all of these efforts at this BoF and invite discussion and additional participation from the community.

---