NM STATE | High Performance Computing

NEWSLETTER

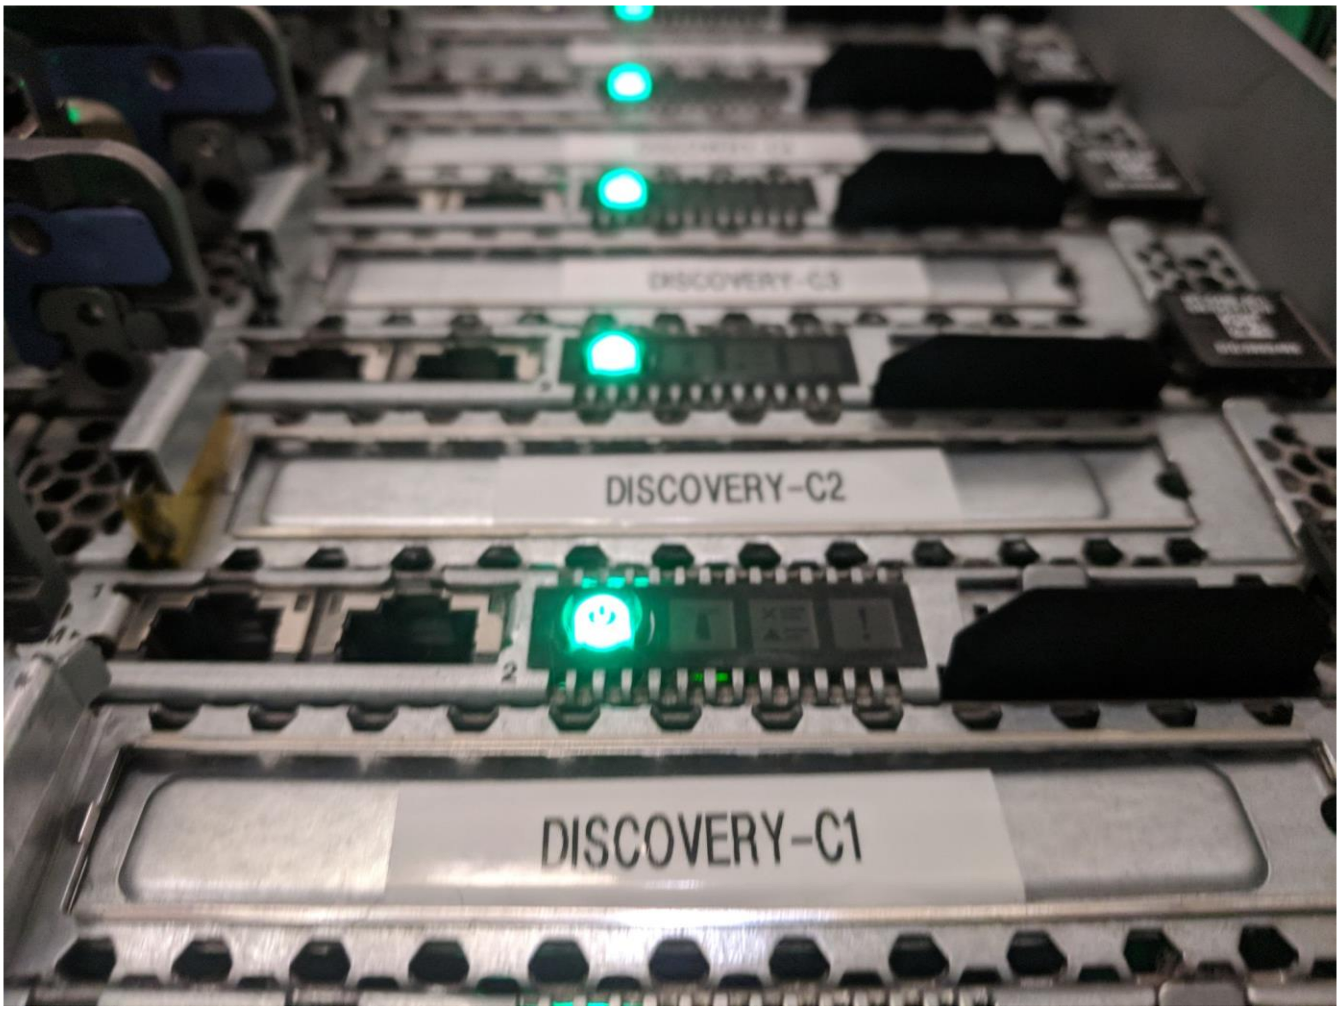

22 April | 2022 ISSUE: 0001 | VOL №: 01

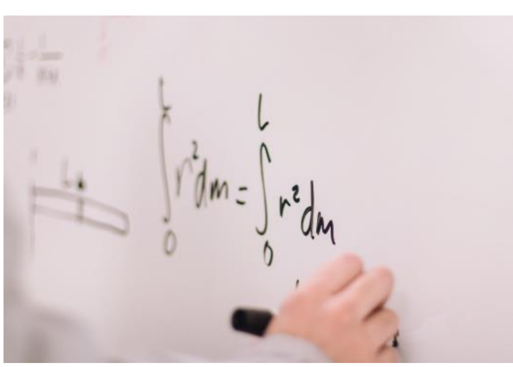

What's HPC?

A supercomputer is a high-level performance computer in comparison to a general-purpose computer. Supercomputers are used for computationally intensive tasks in various fields: quantum mechanics, weather forecasting, climate research, molecular modeling, physical simulations, and much more.

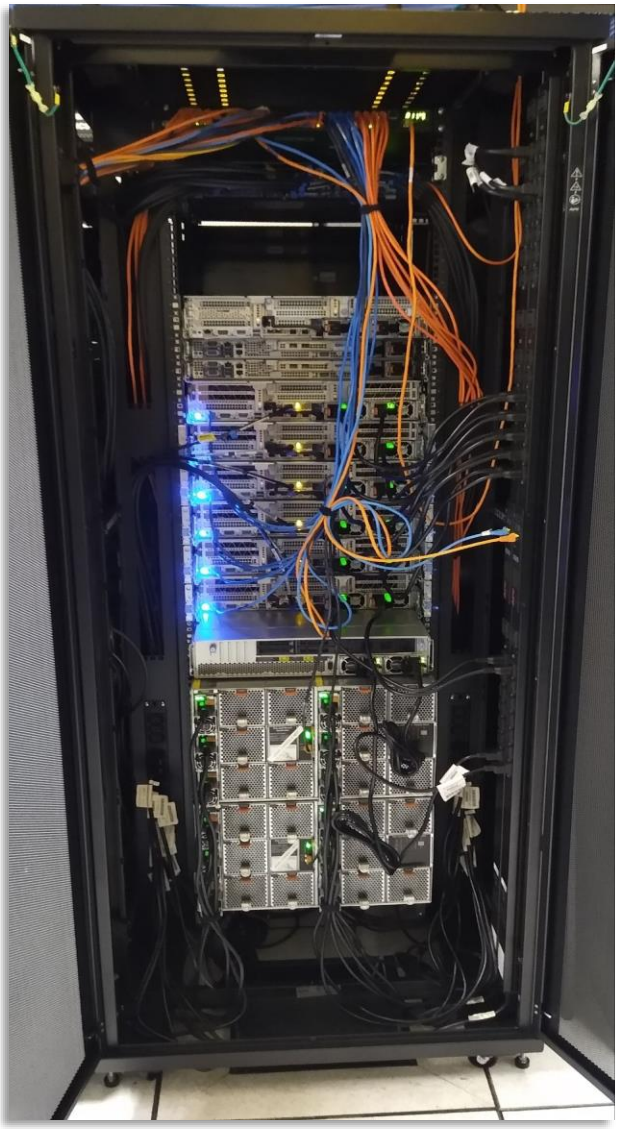

Supercomputers were first developed in the 1960s, and Seymour Cray's Control Data Corporation (CDC), Cray Research, and subsequent companies bearing his name or monogram were the fastest for decades. The earliest machines of this type were finely tuned conventional designs that outperformed their more general-purpose contemporaries. Increasing quantities of parallelism were introduced during the decade, with one to four processors being usual. Vector processors that operate on massive data arrays were popular in the 1970s. The Cray-1, which was released in 1976 and was a huge success, is a great example. Throughout the 1990s, vector computers were the most popular design.

The Discovery cluster is a collection of computers, or nodes, that communicate using InfiniBand, making it an ideal location to scale computational analysis from your personal computer.

HPC has been a critical part of academic research and industry innovation for decades. HPC helps engineers, data scientists, designers, and other researchers solve large, complex problems in far less time and at less cost than traditional computing.

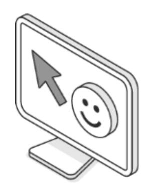

The high-performance computing (HPC) cluster is free to use for current students, faculty, and staff for both research, and teaching.

How did it begin?

The ICT Supercomputer came as a project of two ICT employers from the systems group. They came with the idea to help researchers run their computations on a more powerful machine that they had sitting in their offices. In the Summer of 2015 Student Technology Advisory Committee approved funding for an ICT Supercomputer called Joker. With initial funding, 7 Lenovo nx360 nodes were bought from which one of those nodes had 2 GPUs, the Nvidia K40s. As of May 17th, 2018, the ICT Supercomputer called Joker changed its name to Discovery. As of today, Discovery has 54 compute nodes (36 CPU nodes and 16 GPU nodes), with a total of 1,536 cores.

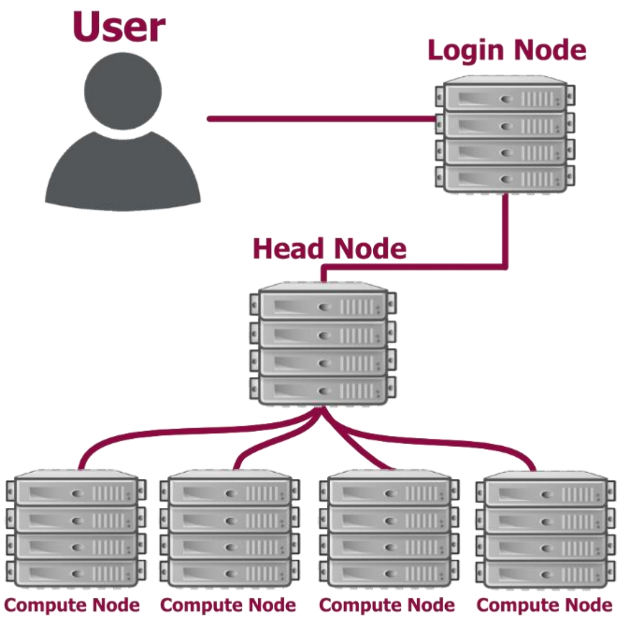

What are Partitions?

The cluster, *Discovery,* is divided into several partitions. These partitions are meant to divide up the resources in the cluster, to balance application workload across multiple machines, and application objects. Partitions are work queues that have a set of rules/policies and computational nodes included in it to run the jobs.

Partition names must be globally unique on all nodes in the cluster. The available partitions HPC has are normal, gpu, interactive, class, backfill, epscor, and other partitions.

Take the interactive partition for example, its ideal only to run interactive jobs to connect to a compute node and work on that node directly. This allows you to change how your jobs might run (i.e., test that commands run as anticipated before putting them in a script) and do heavy tasks that cannot be done on the login nodes, like the use of a lot of cores.

Open Science Grid

The Open Science Grid (OSG) promotes access to distributed large throughput computing for research. The resources provided through the OSG are shared by the community, organized by OSG, and administered by the OSG consortium.

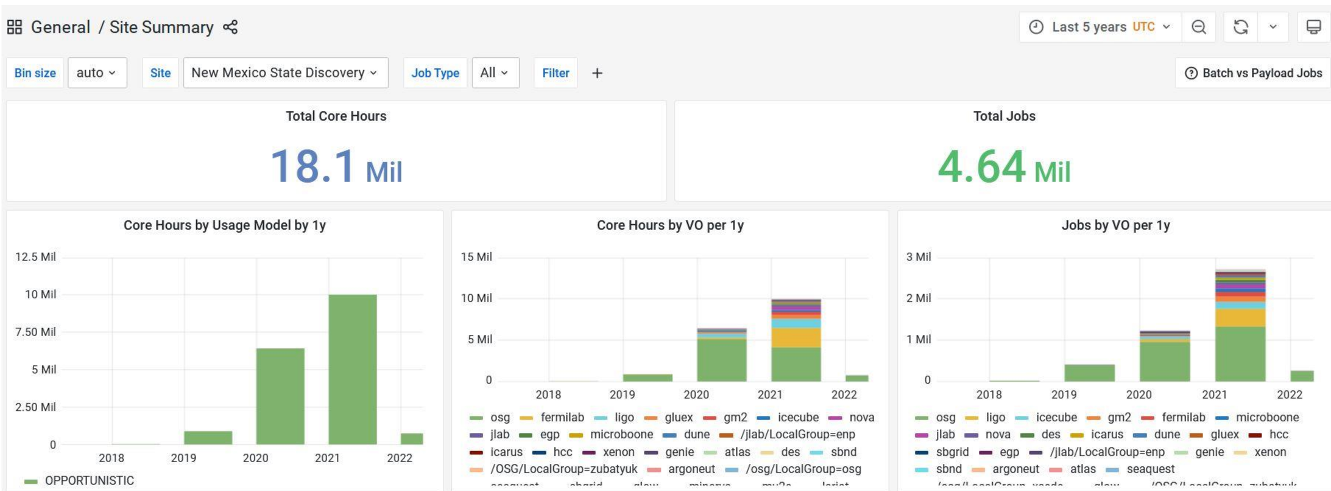

(*Total Core Hours and Jobs in the last 5 years*)

In the last year, the OSG has provided more than 1.2 billion CPU hours to researchers across a variety of projects. The OSG consists of computing and storage elements at over 100 individual sites across the United States. The OSG's Global Research Accounting (GRACC) system provides measures of capacity contributions and usage across reporting OSG pools, including specific data for other institutions and projects.

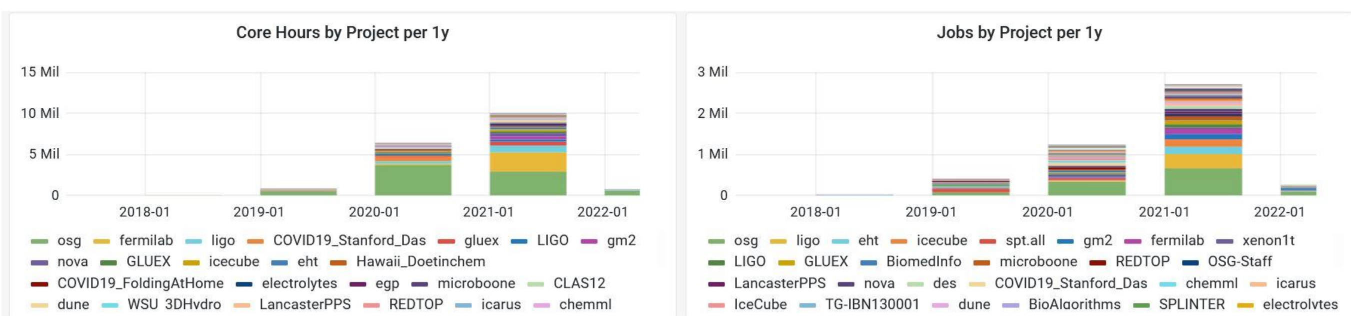

(*Total Core Hours and Jobs by project in the last 5 years*)

These sites, primarily at universities and national labs, range in size from a few hundred to tens of thousands of CPU cores. New Mexico State University has two sites through OSG's GRACC, Discovery and Aggie Grid. Through these sites researchers could see how the HPC cluster is performing, how many cores are available and how many jobs are submitted, along with other information. Some of that information includes Core hours by Project, Jobs by project, fields of science and more!

Wait, what about the Aggie Grid?

HPC systems tend to focus on parallel jobs, and as such they must execute within a particular site with low-latency interconnects. Conversely, HTC systems are independent, sequential jobs that can be individually scheduled on many different computing resources across multiple administrative boundaries. HTC systems achieve this using various grid computing technologies and techniques.

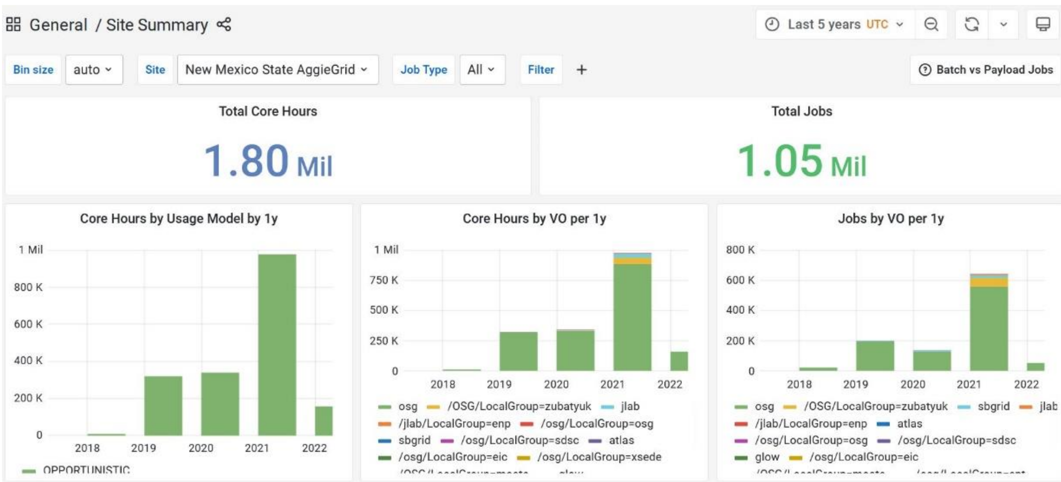

(*Total Core Hours and Jobs in the last 5 years*)

A few years ago, the HPC Team stood up a decentralized High-Throughput computing cluster called the Aggie-Grid to prove such a concept. And it worked! The cluster was a success. It consists of student lab machines that join the cluster when they are idle (when it's not being used). Around the same time NMSU started contributing its computational hours to the Open Science Grid from both the Aggie Grid and Discovery clusters from NMSU. We have contributed over 10.6 million core hours and more than 2.52 million jobs have run on our resources since 2018 including 807K core hours of COVID 19 related research.

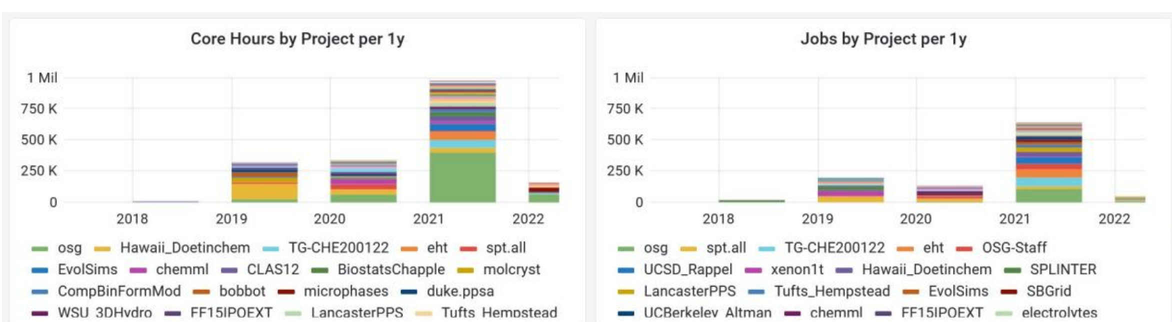

(*Total Core Hours and Jobs by project in the last 5 years*)

Traditionally, computing grids composed of resources (clusters, workstations, and volunteer desktop machines) have been used to support HTC. The general profile of HTC applications is that they are made up of many tasks of which the execution can last for a considerable amount of time. Examples of such applications are scientific simulations or statistical analyses. It is quite common to have independent tasks that can be scheduled in distributed resources because they do not need to communicate. HTC systems measure their performance in terms of jobs completed per month.

Open XDMoD

The University at Buffalo Center for Computational Research (CCR) created the open-source tool. The Open XDMoD (XD Metrics on Demand) tool provides HPC centers with the ability to easily obtain detailed operational metrics of HPC systems, such as Discovery, coupled with extensive analytical capability to optimize performance at the system level, job level, and provides accurate data to guide system upgrades and acquisitions.

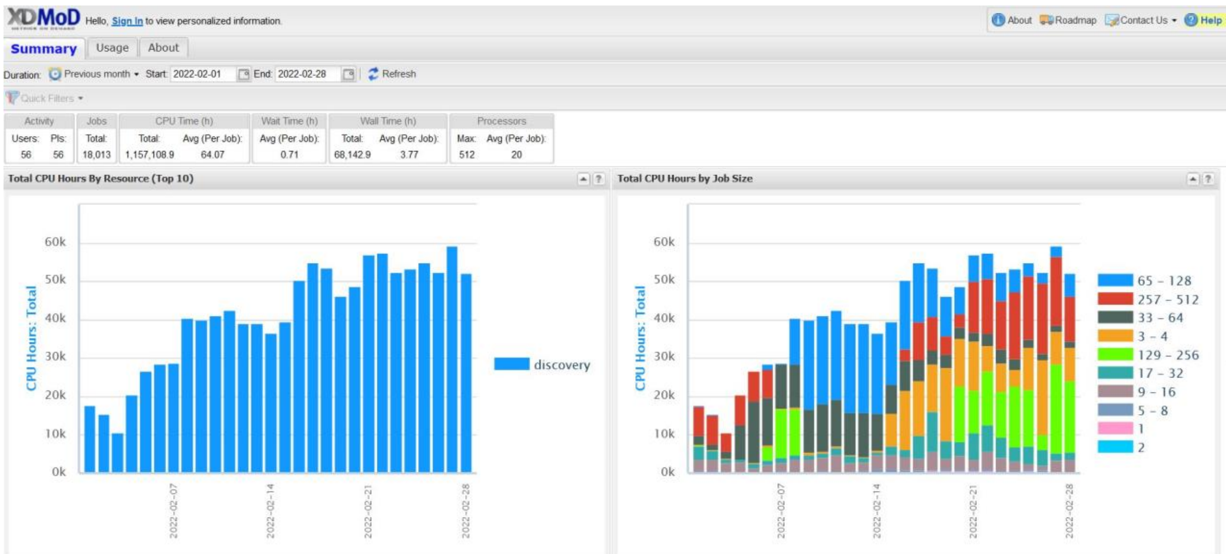

Open XDMoD (https://open.xdmod.org) is open source and funded by the National Science Foundation. It is designed to provide similar capability to academic and industrial HPC centers.

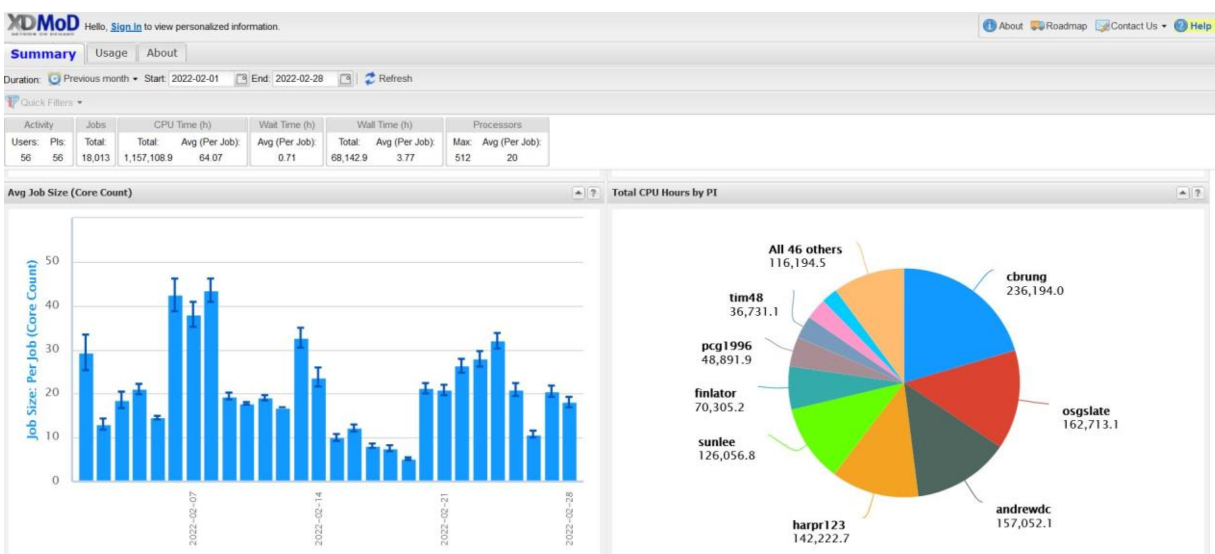

Below you will find our Open XDMoD for Discovery. Here you will know how much activity is present as well as jobs, CPU time, and Wall times. https://xdmod.nmsu.edu/

Open OnDemand

Open OnDemand is an interactive interface to remote computing resources. It helps researchers and students utilize remote computing resources by making them easy to access from any device.

Open OnDemand is an NSF-funded open-source HPC portal based on OSC's original OnDemand portal. The goal of Open OnDemand is to provide an easy way for system administrators to provide web access to their HPC resources, such as you! Luckily for you we have a workshop that covers how to use Open OnDemand on our YouTube Channel. So, if you want to know about the File manager, shell access, interactive desktops on compute nodes, VS Code editor, and more, we suggest you check our workshop out!

(*Click the image above for our YouTube Channel, or go to https://www.youtube.com/watch?v=5gboulTMN3Y*)

Discourse

The next-generation forum software for online discussions. As we move forward, we want to hear and reach out to our researchers. That is why we are implementing an online discussion forum. This will help us and you with any problems related to HPC or the community. In here you will find discussions ranging from introduction to python to MPI Parallelization. Discourse is 100% open-source discussion platform. Use it as a mailing list, discussion forum, long-form chat room, and more! If you have a question, you can post it here in our own Discourse community, where you can get answers from other members. Make sure to check it out!

(*Click the image above for our online discussion forum, or go to https://hpc-forum.nmsu.edu/*)

Future Upgrades and more!

As we move forward, we plan to upgrade our HPC software stack CentOS 7 to Rocky Linux 8 operating system. Hopefully this will be fully integrated by the end of the summer or possible winter. As more details arise, we will let the community know. For now, make sure to follow us on our social media and other platforms. Links are provided at the end.

Meet the Team

The HPC team in the NMSU Information and Communication Technologies (ICT) department oversees the Discovery cluster. The team is made up of a couple of system developers and a few graduate assistants.

Diana Dugas is our Cyber Infrastructure Architect whose mission is to bring more attention and resources to the need for computational support for researchers on campus, she answers directly to our CIO.

Nicholas Von Wolff and ***Strahinja Trecakov*** are the system developers who are behind the Discovery administration.

The rest of the team are all graduate assistants (GAs) working on outreach, on-boarding and other technical projects for the HPC Team.

Emanuel Medina works on outreach for the HPC team. His current graduate studies are in instructional design with an undergraduate degree in Animation & VFX.

Chidambaram Crushev works on on-boarding and other technical projects for the team. His graduate work revolves around Human Computer Interaction and is a huge soccer (Manchester United) fan.

Mohammad Al-Tahat, is working on a Computer Science master's degree. His specialization revolves around Java, SQL, Database, and object-orientated programming.

Essa Imhmed is currently working on an interdisciplinary Ph.D. majored in computer science and electrical engineering. His research interests lie primarily around memory system performance from both hardware and software perspectives.

Interesting fact is that we all come from different countries, so we have a diverse community. This helps us learn how we can provide a better service to a variety of students and researchers.

Together the HPC team provides a variety of ways to help the researchers and students in computational tasks. ICT provides the cyber-infrastructure needed for cluster computing, while the compute resources are owned and expanded by the members of the NMSU community.

Acknowledgement

Publications that feature work that relied on NMSU's HPC or HTC computing resources should cite the following publication. [Strahinja Trecakov and Nicholas Von Wolff. 2021. Doing more with less: Growth, improvements, and management of NMSU's computing capabilities. In Practice and Experience in Advanced Research Computing (PEARC '21), July 18–22, 2021, Boston, MA, USA. ACM, New York, NY, USA 4 Pages. https://doi.org/10.1145/3437359.3465610]

Please use the following language to acknowledge Discovery and the New Mexico State University High Performance Computing group in any published or presented work for which results were obtained using the Discovery cluster.

This work utilized resources from the New Mexico State University High Performance Computing Group, which is directly supported by the National Science Foundation (OAC-2019000), the Student Technology Advisory Committee, and New Mexico State University and benefits from inclusion in various grants (DoD ARO-W911NF1810454; NSF EPSCoR OIA-1757207; Partnership for the Advancement of Cancer Research, supported in part by NCI grants U54 CA132383 (NMSU)).