

Open OnDemand, Open XDMoD, and ColdFront: An HPC center management toolset **Welcome!**

Tutorial presented at PEARC 2023 by staff from:

Ohio Supercomputer Center

University at Buffalo Center for Computational Research

IF YOU HAVE NOT ALREADY DONE SO,
PLEASE FOLLOW SETUP INSTRUCTIONS!

<https://github.com/ubccr/hpc-toolset-tutorial>




Ohio Supercomputer Center

An OH·TECH Consortium Member



University at Buffalo

Center for Computational Research





Ohio Supercomputer Center
An OH·TECH Consortium Member



University at Buffalo

Center for Computational Research


In-Person Tutorial Staff:

Alan Chalker, OSC
Jeff Ohrstrom, OSC
Ryan Rathsam, UB
Travis Ravert, OSC
Dori Sajdak, UB
Andrew Stoltman, UB

Contributing Tutorial & Product Dev Staff:

Andrew Bruno, UB
Alan Chalker, OSC
Andrew Collins, OSC
Robert DeLeon, UB

Trey Dockendorf, OSC
David Hudak, OSC
Matt Jones, UB
Joseph White, UB





Ohio Supercomputer Center
An OH-TECH Consortium Member




University at Buffalo

Center for Computational Research

Agenda

- Tutorial Goals
- Brief intro on all three products
- Tutorial technology & getting help
- Part 1: ColdFront
- **9:45-10:15am: Break**
- 10:15-10:45am: ColdFront (con't)
- 10:45-11:30am: Part 2: Open OnDemand
- **11:30am-1pm: Lunch**
- 1-2:15pm: OnDemand (con't)
- 2:15-3pm: Part 3: Open XDMoD
- **3-3:30pm: Break**
- 3:30-4:20pm: XDMoD (con't)
- 4:20-4:30pm: Wrap up & Q&A





Ohio Supercomputer Center
An OH·TECH Consortium Member




University at Buffalo

Center for Computational Research

Tutorial Goals:

- First presented at PEARC 2020 (virtually)
- Provide participants with an introduction to each product
- Demonstrate the features that allow the products to work together
- Give participants a cluster in a container to practice using these products
- Supply participants access to the developers of these products as a resource for questions & help





Ohio Supercomputer Center
An OH·TECH Consortium Member




University at Buffalo

Center for Computational Research

ColdFront – Managing Access

- Used as the source of record in an HPC center to ensure security & continuity of the systems
- Provides center staff ability to manage center resources & who has access to them
- Portal for users to manage their access to center resources & report on their research
- Plug-ins for job scheduler (Slurm), central authentication, job statistics (XDMoD), OnDemand, that enable automation of access to or removal from resources
- Reports for center management to demonstrate the center's impact (publications, grants, research output)





Ohio Supercomputer Center
An OH·TECH Consortium Member



University at Buffalo


Center for Computational Research

Open OnDemand – Easy Access

- Web-based portal for accessing HPC services that removes the complexities of HPC system environments from the end-user

Includes:

- Files app for upload/download & editing of files
- Terminal app (no need for separate SSH client)
- Job app to create/edit/submit/monitor jobs
- Interactive apps to run GUI applications. Users can create and share apps. Centers can publish apps for all users





Ohio Supercomputer Center
An OH·TECH Consortium Member




University at Buffalo

Center for Computational Research

Open XDMoD – Usage & Performance Metrics

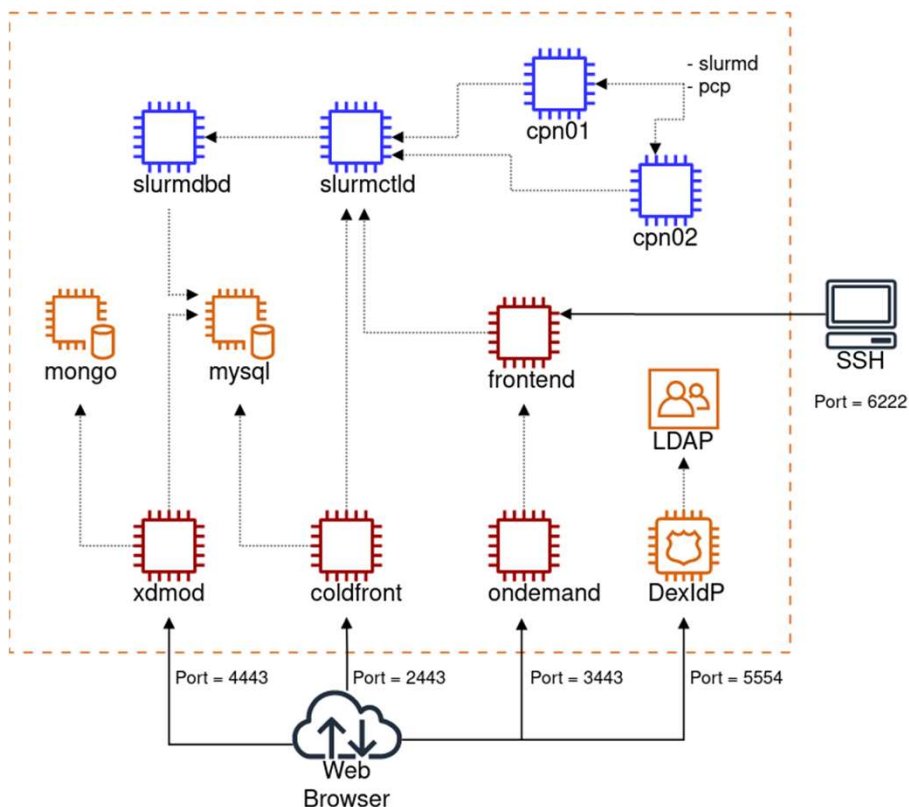
- Tool that aggregates job data & system performance metrics to inform system users, system staff & center decision makers
- Web portal providing job & system metrics, including utilization, quality of service metrics designed to proactively identify underperforming system hardware and software, and job level performance data for every job
- Role-based access to data with different levels of granularity, including job, user, or on a system-wide basis
- Ingest OnDemand logs into new OnDemand realm in XDMoD





Tutorial Container Architecture

HPC Toolset Containers



Requirements: <https://github.com/ubccr/hpc-toolset-tutorial/edit/master/docs/requirements.md>

Clone the Github Repo:

```
git clone https://github.com/ubccr/hpc-toolset-tutorial
cd hpc-toolset-tutorial
./hpcts start
```

* The first time you do this, you'll be download 20+GB worth of containers from Docker Hub. This can take a long time depending on your network speeds. After downloaded, the containers are started, and services launched.

WARNING!!! DO NOT run these containers on production systems. This project is for educational purposes only. The container images we publish for the tutorial are configured with hard coded, insecure passwords and should be run locally in development for testing and learning only.





Ohio Supercomputer Center
An OH·TECH Consortium Member



University at Buffalo

Center for Computational Research

Tutorial Walk Through

<https://github.com/ubccr/hpc-toolset-tutorial>

Keep the applications page open for easy access to account
usernames/passwords and portal URLs:

<https://github.com/ubccr/hpc-toolset-tutorial/blob/master/docs/applications.md>





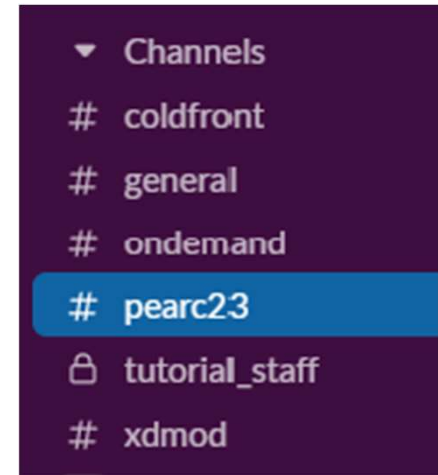
Getting Help

- Join the Slack organization for the tutorial

<https://tinyurl.com/hpctoolset>


- What to do if you're having a technical problem:

Slack us or raise your hand & we'll do our best to help out




Remember: all of this is available after the tutorial so you can just sit back and watch





Tutorial presented at PEARC 2023 by:

Dori Sajdak
Senior Systems Administrator, UB CCR



Ohio Supercomputer Center

An OH·TECH Consortium Member



University at Buffalo

Center for Computational Research

What is ColdFront?



HPC resource allocation management system

Integrates with 3rd party apps for automation & access control

Intuitive Portal for users, system admins, & center staff

Self-service project management for faculty, ROI data collection for you

Source of truth for your center

Centralized info, better security, customized for your workflow and policies

Open source, written in python, extensible with Django apps, **active community** that contributes to the project



Ohio Supercomputer Center

An OH·TECH Consortium Member

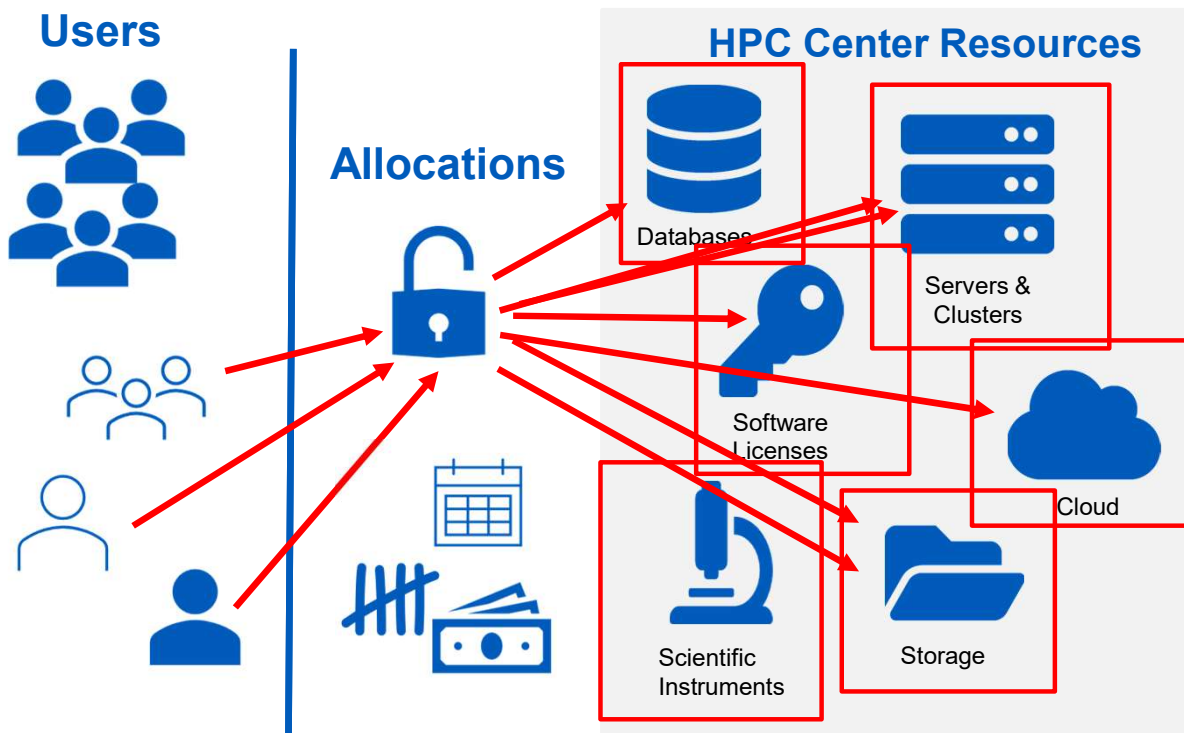


University at Buffalo

Center for Computational Research



Center Policies



Access Control

UNIX/AD Groups

Slurm Accounts

Database Accounts

Homegrown Scripts



Proprietary Software




ORACLE





Why We Developed ColdFront:

- System Administrators wanted:
 - More automation, less manual error
 - One location for access management of all resources
 - Allow PIs to self-service their access to resources
- Center Director wanted:
 - To require PIs to update project info annually
 - Consistent reporting of publication & grant info
 - Easy displays of usage for annual reporting





Ohio Supercomputer Center
An OH·TECH Consortium Member



University at Buffalo

Center for Computational Research

ColdFront – 3 Major Parts

Resources

Anything you want to control access to and/or monitor usage of

Allocations

Determine what resource an account has access to & for how long

Projects

Contain project info, users, allocations, & reportable data (publications, grants, etc)





Resources

- Resources might include:
 - clusters, cluster partitions, storage platforms, cloud, servers, scientific instruments, databases, software licenses
- Resources have attributes
 - Some might restrict access & tie into plugins for automation:
 - Cluster config options – Slurm plugin
 - System access or filesystem access restricted by UNIX group - FreeIPA plugin
 - Some might be informational:
 - End user license agreements
 - Warranty expiration dates
- Metadata allow for fine grain control
 - Is the resource private or public?
 - Available only to certain users/groups?
 - Is the resource a child resource of a larger (parent) resource?
- The attributes set on resources are inherited by allocations
- These attributes are customizable





Allocations

- Determines what resource a user account has access to & for how long
- Allocation attributes may set limits, restrict access, and/or tie into the ColdFront plugins. Examples include:
 - Slurm account name
 - Total number of CPU/core hours allowed
 - UNIX group
 - Storage quota
- Allocation Change requests allow requesting & updating attribute values
- Allocation attributes are customizable
- Allocation metadata includes start & end, creation & last modified dates, status, description, associated resource(s) & justification
- Users emailed when end date is imminent – configurable time spans
- If allocation isn't renewed & expires access to the resource is removed for everyone on the allocation





Projects

- Project = users, allocations for resources, reportable data (publications, grants), project attributes, description, field of science
- PIs (group manager) can request allocations for resources, add/remove users on their project & allocations, upload research info, complete annual project review, view group usage
- Role based logins allow for:
 - full project access for PIs
 - additional capabilities for managers assigned by PIs,
 - read-only views for users,
 - HPC center staff have access to tools for:
 - Allocation review, approval, & configuration
 - Annual project review approval
 - Other policy-driven tools





System administrator views of allocation requests

Allocation Requests

#	Date Requested/ Last Modified	Project Title	PI	Resource	Project Review Status	Status	Allocation Actions
1	Apr. 01, 2021	my project	cgray cgray (cgray)	hpc cluster (Cluster)		New	Activate Deny

Allocation Requests & Change Requests Can be Viewed by System Administrators

Allocation Change Requests

For each allocation change request below, there is the option to activate the allocation request and to view the allocation change's detail page. If a change request is only for an extension to the allocation, they can be approved on this page. However if the change request includes changes to the allocation's attributes, the request must be reviewed and acted upon in its detail page.

#	Requested	Project Title	PI	Resource	Extension	Actions
24	Jan. 13, 2022	Testing Storage 1	Dori Sajdak (djm29)	ProjectStorage (Storage)	60 days	Approve Details
27	Jan. 14, 2022	Testing Storage 1	Dori Sajdak (djm29)	BudgetStorage (Storage)	60 days	Approve Details





Annual Project Reviews

You cannot request a new allocation because you have to review your project first.

You need to review this project. [Review Project](#)

Test Project

Manage Project

Reviewing Project: Test Project

CCR requires faculty to review their project information annually in order to renew their group's accounts. The information provided by researchers is compiled and used to help make the case to the University for continued investment in CCR. Up-to-date and accurate information is crucial to our success. [Questions?](#) [Contact us](#)

Please update the following information:

1. [Verify](#) your project description is accurate
2. [Add](#) Publications
3. [Add](#) Grants
4. [Verify](#) the user accounts in your group and remove any that should no longer have access to CCR resources

Grants Last Updated:	Sep. 11, 2018
Publications Last Updated:	Sep. 11, 2018
Users in project:	Dori Sajdak

Reason for not updating project information*

If you have no new information to provide, you are required to provide a statement explaining this in this box. Thank you!

Completed Annual Project Reviews Can be Viewed by Center Director and System Admins


Pending Project Reviews

Project Title	Date Review Submitted	PI	Grants Last Updated	Publications Last Updated	Reason for not Updating Project	Project Review Actions
My Test Project	May. 13, 2021	Dori Sajdak (djm29)	May. 13, 2021	May. 13, 2021		Mark Complete Email





User Grants Summary



- National Institutes of Health (NIH): \$204,138,700 (93)
- National Science Foundation (NSF): \$86,283,804 (121)
- Department of Energy (DOE): \$49,513,248 (19)
- Other: \$16,838,729 (78)
- Department of Defense (DoD): \$8,455,898 (15)
- National Aeronautics and Space Administration (NASA): \$4,001,872 (7)
- New York State Department of Health (DOH): \$3,020,564 (1)
- Empire State Development's Division of Science, Technology and Innovation (NYSTAR): \$1,028,724 (2)
- Empire State Development (ESD): \$1,000,000 (1)
- New York State (NYS): \$746,245 (11)

Grants Total: \$375,027,785

Grants Total PI Only: \$179,766,586


Grants Total CoPI Only: \$152,122,982

Grants Total Senior Personnel Only: \$43,138,217

Center Directors can better demonstrate the center's impact

- Collect publication & grant info
- Research output – other work product
- Enforce annual updates of info – project description, field of science
- Other ROI information?

User Entered Publications



Total Publications: 1979



Ohio Supercomputer Center
An OH-TECH Consortium Member



University at Buffalo

Center for Computational Research

Extensible plug-in architecture
allows for **integration of
nearly anything!**



Vendor APIs





Ohio Supercomputer Center
An OH·TECH Consortium Member



University at Buffalo

Center for Computational Research

<https://www.osc.edu/> <https://openondemand.org/>

<https://buffalo.edu/ccr> <https://open.xdmod.org/> <https://coldfront.io>

NOW: Break: 30 minutes

Coming Up!

Part 1: ColdFront (con't)

Part 2: Open OnDemand

Important Info:

Tutorial Repo: <https://github.com/ubccr/hpc-toolset-tutorial>

Join us on Slack: <https://tinyurl.com/hpctoolset>

Other Places You'll Find us at PEARC23:

Open OnDemand User Group Meeting (BOF): Tues, 7/25 1:30-2:30pm – **F150**

Account & Access Management Needs & Approaches: A Community Conversation (BOF): Tues, 7/25 4:45-5:45pm – **E143**

XDMoD BOF: Wed, 7/26 1:30-2:30PM – **E146**

Open OnDemand Booth - During Exhibition Hall hours



OPEN **nDemand**

Overview & Walk-through

Alan Chalker

Travis Ravert

Jeff Ohrstrom


Ohio Supercomputer Center



Ohio Supercomputer Center
An OH·TECH Consortium Member



University at Buffalo
Center for Computational Research





Why Run Open OnDemand?

Zero installation

Run Open OnDemand entirely in your browser. No client software installation required.

Easy to use

Start computing immediately. A simple interface makes Open OnDemand easy to learn and use.

Compatible with any device

Launch on any device with a browser—even a mobile phone or tablet.

Access your organization's supercomputers through the web to compute from anywhere, on any device.

openondemand.org/run





Why Install Open OnDemand?

Low barrier to entry

Empower users of all skill levels by offering an alternative to command-line interface.

Free and open source

Install Open OnDemand for free, and gather knowledge from our large open-source community.

Configurable and flexible

Create and deploy your own applications to meet your users' unique needs.

Administer remote web access to your supercomputers to transform the way users work and learn.

openondemand.org/install





Ohio Supercomputer Center
An OH-TECH Consortium Member



University at Buffalo
Center for Computational Research

Example Deployments



Don't see an organization?
The more the merrier!
Let us know any that belong on the list



openondemand.org/orgs



Community Events



Tips and tricks calls

Hosted by the larger Open OnDemand community, tips and tricks webinars share best practices for setting up and using Open OnDemand. They take place on the first Thursday of every month at 1 p.m. ET.



Open office hours

Hosted by our development team, Zoom open office hours are the perfect opportunity to ask questions or make a suggestion. They are held on the second Tuesday of every month from 11:15 a.m. to 12:45 p.m. ET.



PEARC Booth / BoF

Visit the Open OnDemand table (#10)
Attend our User Group BoF,
Tues (July 25), 1:00 – 2:30 p.m.
Room F150

openondemand.org/events





Need Support?



Discuss on Discourse


The Get Help thread on our Discourse forum features user and admin questions and answers. Browse recent questions or leave one of your own.



Documentation

Our documentation outlines installation steps, app development guidelines, release notes, and more. Search for a specific topic or browse for general info.

openondemand.org/support





Ohio Supercomputer Center
An OH·TECH Consortium Member



University at Buffalo
Center for Computational Research

Demo Topics

Dashboard landing page

File management

Batch jobs / job composer

Viewing active jobs

Interactive applications

Dashboard profiles

openondemand.org/demo





Ohio Supercomputer Center
An **OH·TECH** Consortium Member



University at Buffalo

Center for Computational Research

<https://www.osc.edu/> <https://openondemand.org/>

<https://buffalo.edu/ccr> <https://open.xdmod.org/> <https://coldfront.io>

NOW: Lunch – 90 minutes

Coming Up – starting at 1:30pm

Part 2: Open OnDemand (con't)

Part 3: Open XDMoD

Important Info:

Tutorial Repo: <https://github.com/ubccr/hpc-toolset-tutorial>

Join us on Slack: <https://tinyurl.com/hpctoolset>

Other Places You'll Find us at PEARC23:

Open OnDemand User Group Meeting (BOF): Tues, 7/25 1:30-2:30pm – **F150**

Account & Access Management Needs & Approaches: A Community Conversation (BOF): Tues, 7/25 4:45-5:45pm – **E143**


XDMoD BOF: Wed, 7/26 1:30-2:30PM – **E146**

Open OnDemand Booth - During Exhibition Hall hours



XDMoD Overview

Ryan Rathsam, Andrew Stoltman
Center for Computational Research
University at Buffalo



Ohio Supercomputer Center

An OH·TECH Consortium Member

XDMoD
METRICS ON DEMAND




University at Buffalo

Center for Computational Research

 **ACCESS**
Metrics





XDMoD: Metrics on Demand

- **Comprehensive framework for CI system management**

- Compute Jobs, Storage, Cloud, and Networking

- **Understand and optimize resource utilization and performance**


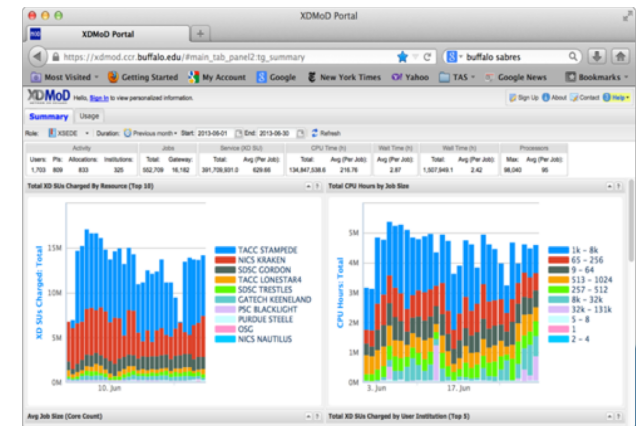
- Provide instantaneous and historical information on utilization
- Measure Quality of Service of CI systems and applications
- Measure and improve job and system level performance
- Inform computing system upgrades and procurements

- **ACCESS XDMoD tool**

- Analytics Framework for XSEDE/ACCESS

- **Open XDMoD*: Open Source version for CI centers**

- Used to measure and optimize performance of CI centers
- 400+ academic, governmental, & commercial installations worldwide
- <https://open.xdmod.org/>





Ohio Supercomputer Center
An OH-TECH Consortium Member



University at Buffalo

Center for Computational Research



Benefits for Stakeholders

- **PI and End User**
 - Account management, resource selection, application tuning, improved throughput
- **Systems Administrator**
 - System diagnostic and performance optimization, application tuning
- **Center Support Staff**
 - Diagnose Job Failure / Performance Issues
 - Identify users who may need additional help
- **CI Center Director**
 - Comprehensive resource management and planning tool
 - Return on Investment Metrics





Why Improving CI Performance Matters

- **CI systems are oversubscribed**

- Improving application or system performance improves overall job throughput
- Frees up otherwise wasted CPU cycles for useful work

- **Small improvements in performance can have high impact**

- Every 1% increase in system performance on the resources provided through XSEDE translates into the ability to allocate an additional 101 M CPU hours annually
- Corresponds to a savings of \$5M*





XDMoD Portal

- **Web-based**
 - Point and click drill down capability
- **Display metrics**
 - Utilization, performance, scientific impact
- **Role based access**
 - User
 - Principal Investigator
 - Center Staff
 - Center Director
- **Custom Report Builder**





Ohio Supercomputer Center
An OH·TECH Consortium Member



University at Buffalo

Center for Computational Research

<https://www.osc.edu/> <https://openondemand.org/>

<https://buffalo.edu/ccr> <https://open.xdmod.org/> <https://coldfront.io>

NOW: Break: 30 minutes

Coming Up!

Part 3: Open XDMoD (con't)

Wrap up and Q&A

Important Info:

Tutorial Repo: <https://github.com/ubccr/hpc-toolset-tutorial>

Join us on Slack: <https://tinyurl.com/hpctoolset>

Other Places You'll Find us at PEARC23:

Open OnDemand User Group Meeting (BOF): Tues, 7/25 1:30-2:30pm – **F150**

Account & Access Management Needs & Approaches: A Community Conversation (BOF): Tues, 7/25 4:45-5:45pm – **E143**

XDMoD BOF: Wed, 7/26 1:30-2:30PM – **E146**

Open OnDemand Booth - During Exhibition Hall hours



Thank you for attending!

Please fill out the post-tutorial survey

We value your opinions!




Ohio Supercomputer Center

An OH·TECH Consortium Member



University at Buffalo

Center for Computational Research





Funding and other acknowledgements:

- OnDemand is supported by the National Science Foundation – award numbers [NSF#1534949](#) and [NSF#1935725](#)
- Open XDMoD is supported by the National Science Foundation – award numbers [ACI 1025159](#) and [ACI 1445806](#) and [OAC 2137603](#)
- We gratefully acknowledge the partnership with [Virginia Tech](#) on our current NSF project





Ohio Supercomputer Center
An OH·TECH Consortium Member



University at Buffalo

Center for Computational Research

Other Places You'll Find us at PEARC23:

Open OnDemand User Group Meeting (BOF): Tues, 7/25 1:30-2:30pm – **F150**

Account & Access Management Needs & Approaches: A Community Conversation (BOF): Tues, 7/25 4:45-5:45pm – **E143**

XDMoD BOF: Wed, 7/26 1:30-2:30PM – **E146**

Open OnDemand Booth - During Exhibition Hall hours

Staff are available after the tutorial concludes for specific questions. Please continue to engage with us and the user communities on Slack!





Ohio Supercomputer Center
An OH·TECH Consortium Member



University at Buffalo

Center for Computational Research

QUESTIONS?

How to reach us:

Center for Computational Research – <https://buffalo.edu/ccr>

Open XDMoD - <https://open.xdmod.org/>

ColdFront – <https://coldfront.io>

Ohio Supercomputer Center - <https://www.osc.edu/>

OnDemand - <https://openondemand.org/>

