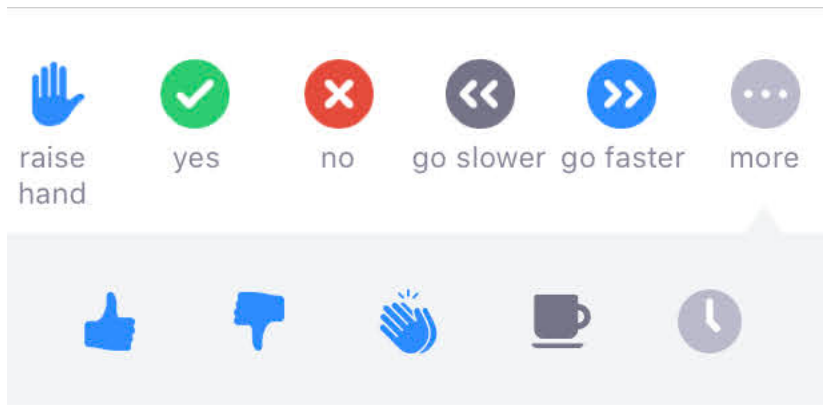


# Decorum for HPC Management Toolset Virtual Tutorial Gateways 2020

- Zoom will be live at least 15 minutes in advance for people to join and get settled
- Set your name to “First Last (Institution)”
  - Click on the “Participants” button, find your name (should be at the top of the list) and click on the “Rename” button. On iPad app, you’ll need to touch your name to see the list of options.
- Only presenters and moderators will broadcast video and screen share
- Participants’ microphones and cameras will be off upon entry to the meeting and we ask you to only unmute after you’ve raised your hand and been called on.
- Moderators will mute participants & shut off video, if necessary
- Zoom feedback buttons:
  - Click on the “Participants” button and you should see the feedback icons:, find your name (should be at the top of the list) and “Raise Hand” button.
  - Click on the “Raise hand” button to indicate you would like to speak, and a moderator will unmute you
  - As we work through steps of the tutorial, it would be helpful to get feedback from you in the form of “yes, I got it, move on” or “no, still working on it.” Also “go slower” and “go faster” reactions are available to you to click on so we can gauge how we’re doing.
  - If you would like, you may use the reactions to indicate when you are away or need a break, though this isn’t required.



- Please do NOT use the Zoom chat but rather use our Slack channel for communication during the tutorial and afterwards, unless you can’t access the Slack channel:  
<https://tinyurl.com/hpctoolset-slack>
- We will be adhering to the Gateways 2020 Code of Conduct  
(<https://sciencegateways.org/web/gateways2020/conduct>) Anyone in violation of this will be removed from the tutorial.



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

<https://openondemand.org/>

<https://buffalo.edu/ccr>

<https://open.xdmod.org/>

<https://coldfront.io>

<https://arc.vt.edu/>

## Important Info:

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

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

## Other Places You'll Find us at Gateways20:

**Monday, October 19, 3-3:15pm**

Improving Science Gateway Monitoring, Caveats and Goals - <https://sched.co/dcY7>  
*Jeanette Sperhac, UB CCR - XDMoD*

**Poster Session – Link to Posters in QiqoChat:**

Cloud HPC with Open OnDemand and Cloudy Cluster  
*Jeff Ohrstrom, OSC - OnDemand*





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



University at Buffalo

Center for Computational Research



**VIRGINIA  
TECH™**

# WELCOME!

This tutorial will use a cluster-in-a-container Docker infrastructure. Check out this repo for more information and setup instructions, if you'd like your own environment to test out these products:

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

Join the Slack channel for the tutorial  
<https://tinyurl.com/hpctoolset-slack>



# Open OnDemand, Open XDMoD, and ColdFront: An HPC center management toolset

Tutorial presented at Gateways20 by staff from:  
Ohio Supercomputing Center  
UB Center for Computational Research  
Virginia Tech Advanced Research Computing



**Ohio Supercomputer Center**

An OH·TECH Consortium Member




University at Buffalo

Center for Computational Research



**VIRGINIA  
TECH™**





## Tutorial Staff:

Andrew Bruno, UB

Alan Chalker, OSC

Andrew Collins, OSC

Robert DeLeon, UB

Trey Dockendorf, OSC

Eric Franz, OSC

David Hudak, OSC

Matt Jones, UB


Jeff Ohrstrom, OSC

Ben Plessinger, UB

Dori Sajdak, UB

Bob Settlage, VT

Joseph White, UB








# Agenda




- Getting Started
- Tutorial Goals
- Brief intro on all three products
- Tutorial technology
- Part 1: ColdFront
- Part 2: Open XDMoD
- Part 3: Open OnDemand & interactive app configuration
- Post Workshop – Zoom sessions & slack channel





## Getting Started

- View our “meeting decorum” document  
<https://tinyurl.com/gateways-hpctoolset>
  - Please mute yourself & leave your video off
  - Use the “raise hand” button if you have a question & our moderator will unmute you
- Join the Slack channel for the tutorial  
<https://tinyurl.com/hpctoolset-slack>
  - Use Zoom chat only if having trouble with Slack
- Clone the tutorial repo and follow instructions for starting containers  
<https://github.com/ubccr/hpc-toolset-tutorial>
- What to do if you’re having a technical problem – refer to Gateways support





# Tutorial Goals:

- Provide participants with an overview of each product & how they are installed
- Point out a few “gotcha!”s to look out for
- Show off the new 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








## ColdFront – Allocations Management Portal


- 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, central authentication, job statistics (XDMoD) that enable automation of access to or removal from resources
- Reports for center management to demonstrate the center's impact (publications, grants, research output)
- Used as the source of record in a HPC center to ensure security of the systems





## Open XDMoD

- 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
- New features such as the user & system report cards provide immediate feedback





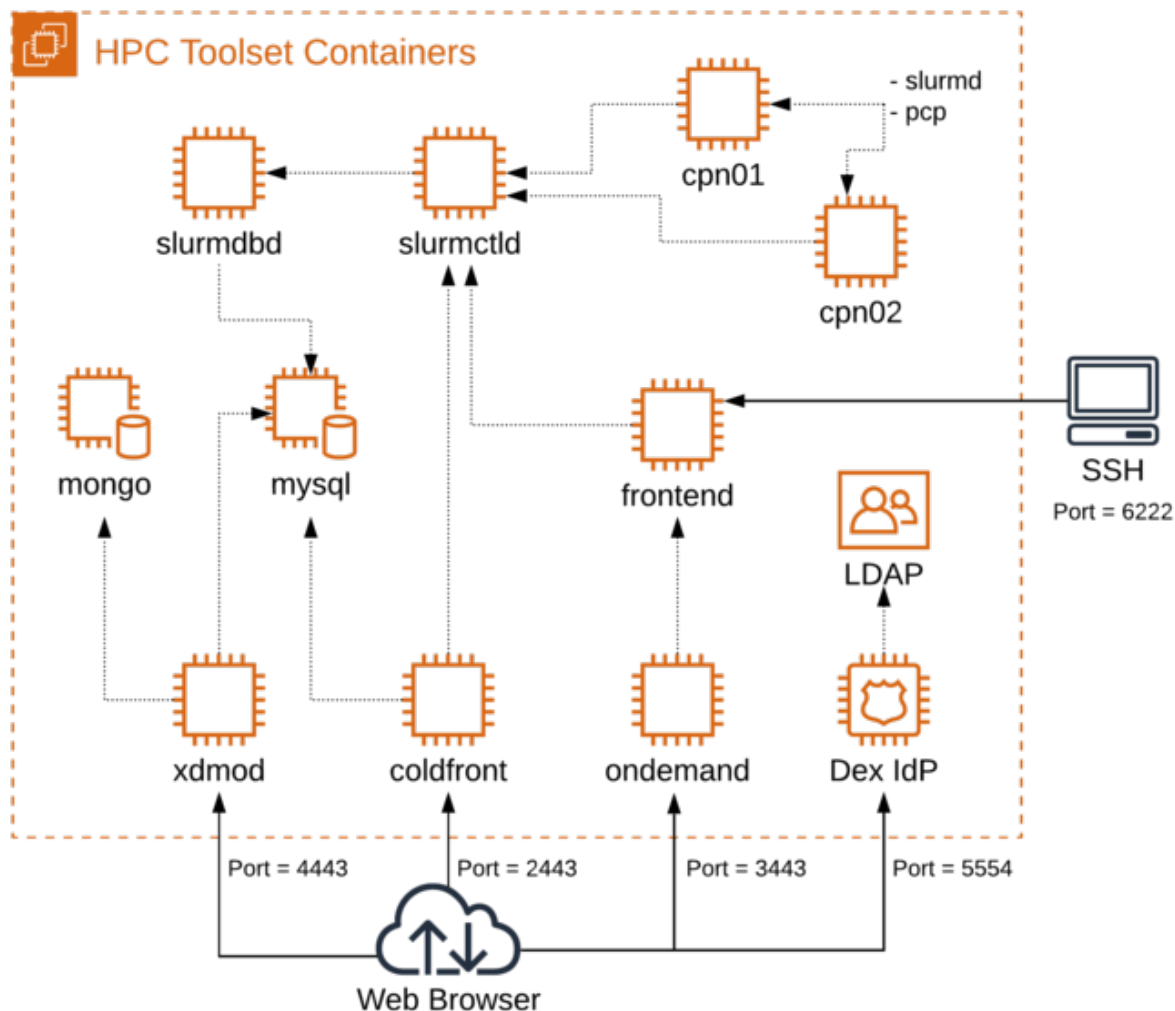
## Open OnDemand

- 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 to 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
- New! Display Open XDMoD job statistics in the OnDemand dashboard





# Tutorial Container Architecture




**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 ~12GB 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.





# Tutorial Walk Through

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

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

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





# **ColdFront:** OpenSource HPC resource **allocation portal** for **users, system admins, &** **center staff**

Tutorial presented at Gateways20 by:  
Andrew Bruno, UB  
Dori Sajdak, UB



**Ohio Supercomputer Center**

An **OH·TECH** Consortium Member

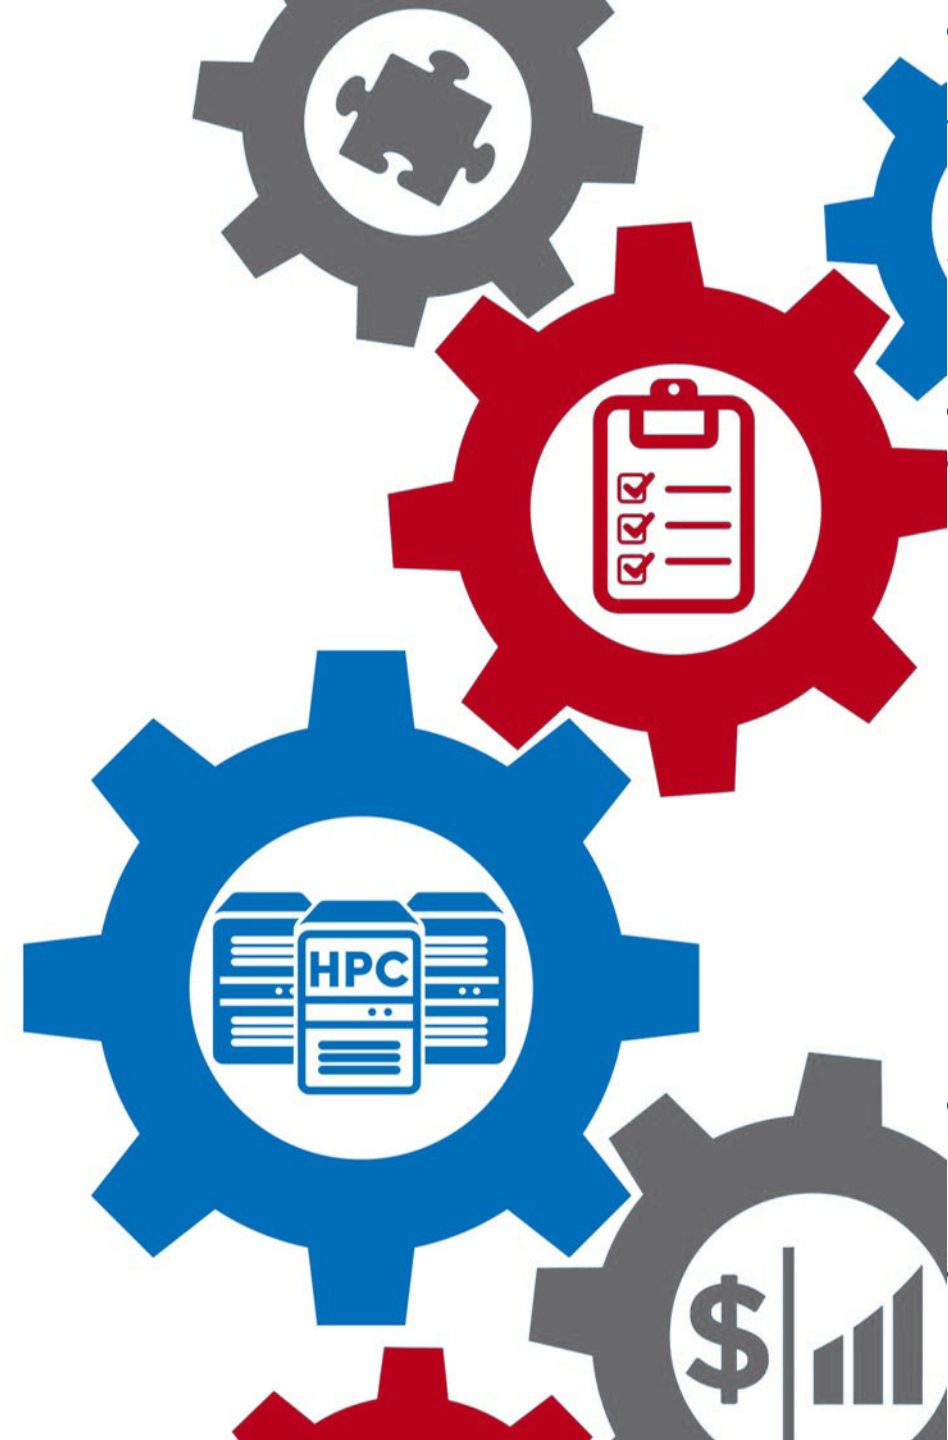


University at Buffalo

Center for Computational Research



**VIRGINIA  
TECH™**



# History of ColdFront


- System Administrators wanted:
  - More automation, less manual error
  - One location for access management of all resources
  - Allow PIs to self-service 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





# Automate access to your HPC resources

Manage access to all your resources in one place





## Resources

- CCR currently has nearly 100 distinct resources to track
- Our resources include:
  - clusters, storage, cloud, servers, and software licenses
- Some resources have limits:
  - storage (GB),
  - software (seats),
  - cloud (subscriptions)
  - these are all customizable
- We track attributes on the resources:
  - Private or public
  - Groups allowed access to it
  - Whether extra payment is required
  - Warranty expiration dates






## Allocations

- Determines what resource an account has access to
- Any limits/attributes associated with that access
  - Expiration date
  - CPU hours
  - Scheduler account name
  - UNIX group
- Users emailed when expiration dates approach
- Resource access can be removed/locked when an allocation expires








## Projects

- Project = users, allocations for resources, reportable data (publications, grants)
- Role based logins allow for:
  - read-only views for users,
  - additional capabilities for managers,
  - full project access for PIs
  - HPC center staff have access to tools for:
    - Allocation review & approval,
    - Usage reports
    - Other policy-driven tools
- PIs (group leads) can add/remove users on their project & allocations





■ National Institutes of Health (NIH): \$78,599,277 (33)  
■ National Science Foundation (NSF): \$52,283,068 (73)  
■ Other: \$12,161,778 (49)

Center Directors are able to better demonstrate the center's impact

**Report on resources & allocations**

**Collect publication information**

**Collect grant information**






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



Vendor APIs





# Sneak Peek! ColdFront – OnDemand Integration

## Allocations »

Project	Resource	Status
Demo Funding	University HPC (Cluster)	Active
Demo Funding	Chemistry-cgray (Cluster Partition)	New
Demo Funding	Budgetstorage (Storage)	Active

### Allocations 3

Show 10 entries

Search:

Resource Name	↑↓ Resource Type	↑↓ Information	↑↓ Status	↑↓ End Date	↑↓ Actions
Budgetstorage	Storage		Active	2021-10-15	
Chemistry-cgray	Cluster Partition		New		
University HPC	Cluster		Active	2021-10-09	

Showing 1 to 3 of 3 entries

Previous 1 Next



## Contact Info:

[Andrew Bruno - aebruno2@buffalo.edu](mailto:aebruno2@buffalo.edu)

[Dori Sajdak - djm29@buffalo.edu](mailto:djm29@buffalo.edu)

<https://coldfront.io>

Subscribe to the ColdFront mailing list:

Send an email to [listserv@listserv.buffalo.edu](mailto:listserv@listserv.buffalo.edu) with no subject, and the following command in the body of the message:

```
subscribe ccr-open-coldfront-list@listserv.buffalo.edu first_name last_name
```

More about UB CCR:

<https://buffalo.edu/ccr>

<https://twitter.com/ubccr>



**Ohio Supercomputer Center**

An **OH·TECH** Consortium Member




University at Buffalo

Center for Computational Research



**VIRGINIA  
TECH™**





A photograph of a server room with blue lighting. The room is filled with rows of server racks. In the background, three people are standing and talking. The floor is a perforated metal grating. The overall atmosphere is high-tech and professional.

# XDMoD

Gateway Conference 2020  
October 2020

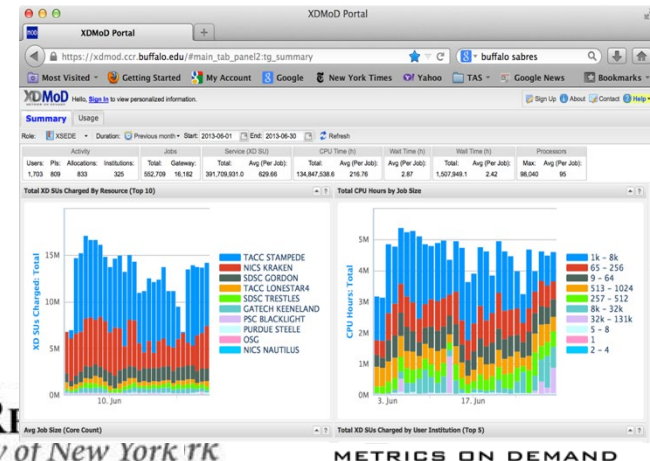
# Outline

- XDMoD Introduction
- XDMoD & Open XDMoD
  - Portal
  - Measuring Quality of Service with Application Kernels
  - Measuring job level performance with Job Viewer
- Updates & New Features
- XMS Team & Contact Information
- XDMoD Tutorial/Demo



# XDMoD: A Comprehensive Tool for HPC System Management

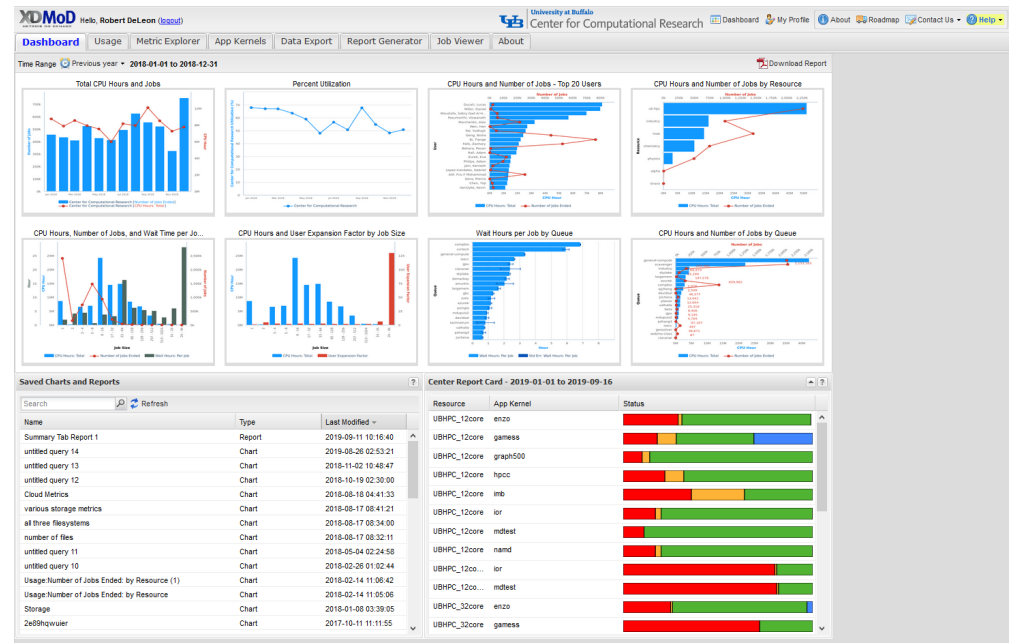
- **XD Net Metrics Service (XMS) NSF Award to CCR**
  - Following 5 year TAS award on-going 5-year XMS award
  - Develop **XDMoD** (**XD Metrics on Demand**) Tool
- **Open XDMoD: Open Source version for Data Centers**
  - Used to measure and optimize performance of HPC centers
  - 200+ academic & industrial installations worldwide
- **Goal: Optimize Resource Utilization and Performance**
  - Provide instantaneous and historical information on utilization
  - Measure Quality of Service
  - Enable data driven upgrades and procurements
  - Measure and improve job and system level performance





# Primary Components of XDMoD

- XDMoD web Portal
  - Metrics Explorer for reports on system usage and efficiency
  - Job Viewer to Measure and improve job performance
- Application Kernels
  - Measure Quality of Service



# XDMoD Portal


- Web-based
- Display metrics
  - Utilization, performance, scientific impact
  - Role based access
  - Custom Report Builder





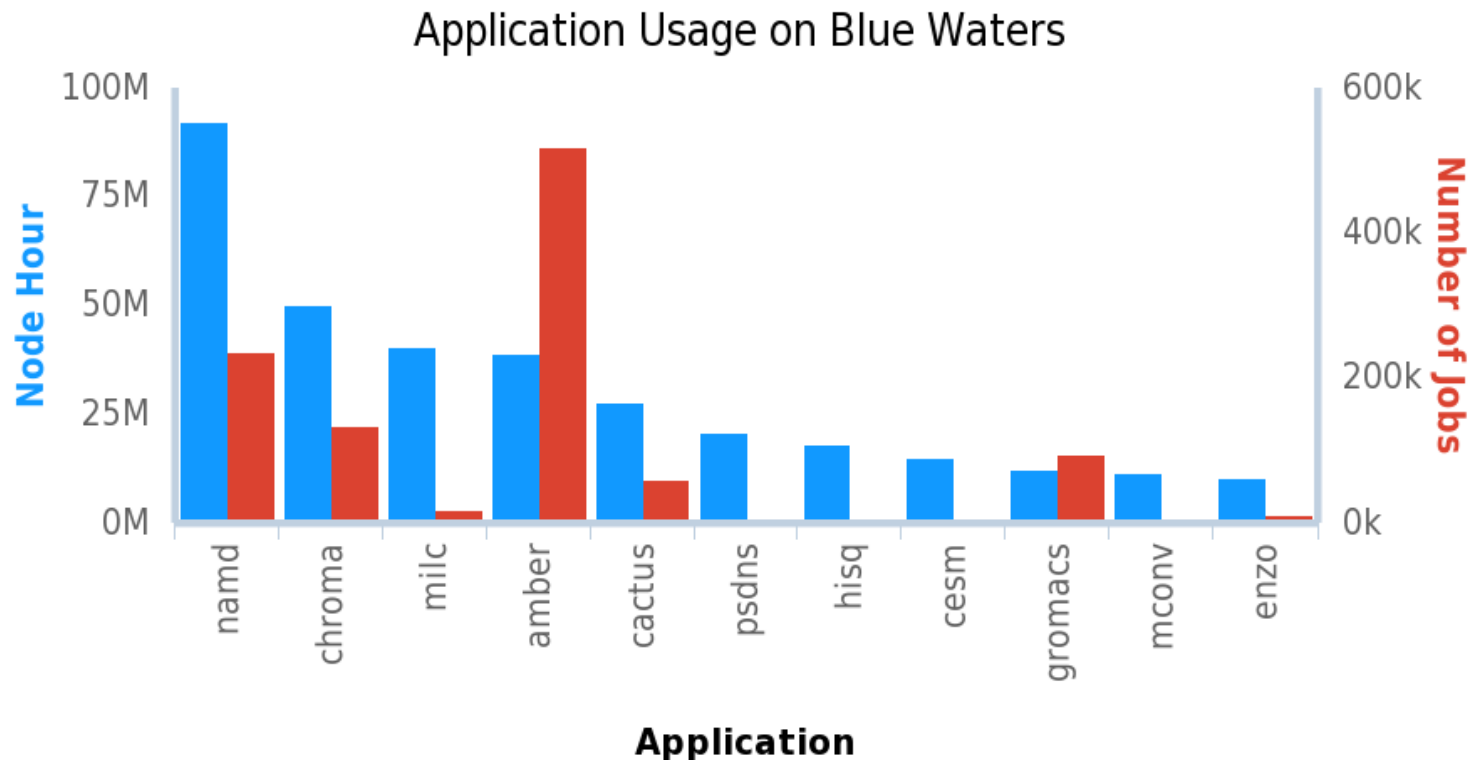
# Easily Obtain Utilization Metrics

- CPU hours consumed by campus units



# Uses: Application Analysis

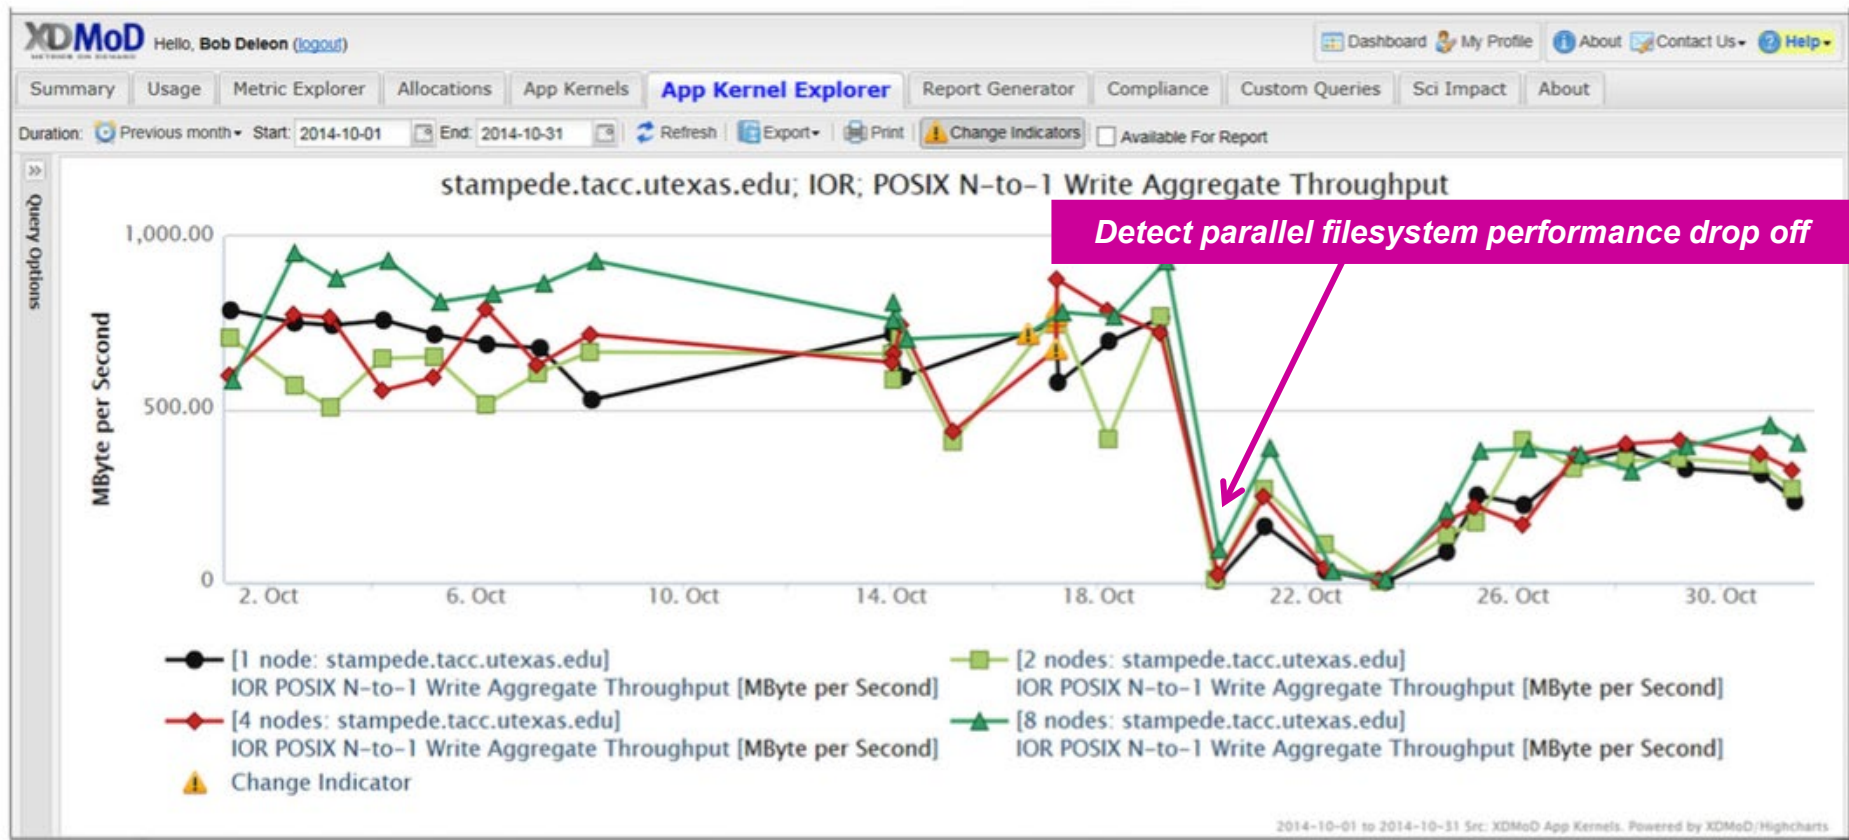
- Determine what are the mostly widely used applications running on Blue Waters (large NSCA HPC resource)




2013-03-01 to 2016-12-31 Src: SUPREMM. Powered by XDMoD/Highcharts

# QoS: Application Kernels

- Computationally lightweight benchmarks or applications
  - Run periodically or on demand to actively measure performance
- Measure system performance from User's perspective
  - Local scratch, global filesystem performance, local processor-memory bandwidth, allocatable shared memory, processing speed, network latency and bandwidth
- Proactively identify underperforming hardware and software

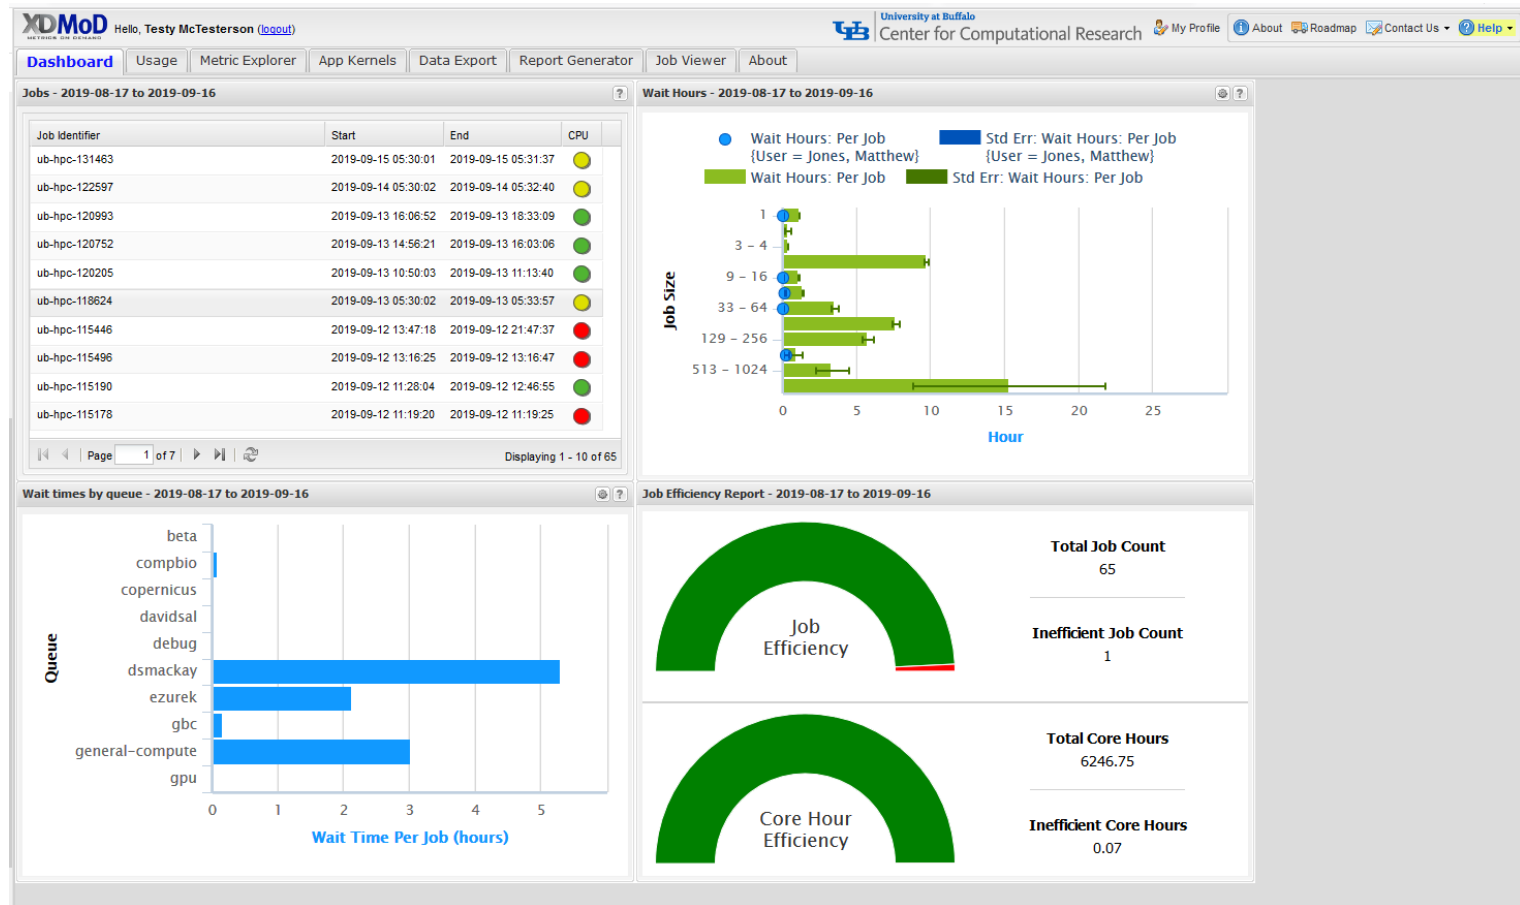


# Job Viewer: Measuring Job Performance

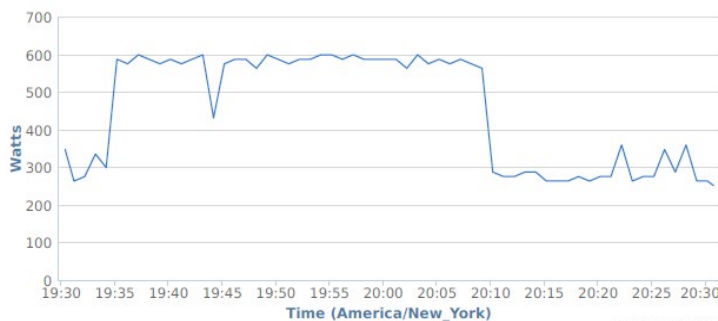


# XDMoD Updates


- Dashboard: start using XDMoD faster and easier
- Job efficiency report: How efficient are your jobs?



# XDMoD Job Level Energy/Power



(a) Node power usage over time.



(b) Overall CPU usage over time.

- Job level energy/power consumption monitoring has been added to XDMoD
- PEARC20 paper won the Phil Andrews award
  - <https://pearc20.sched.com/event/cnVp/monitoring-and-analysis-of-power-consumption-on-hpc-clusters-using-xdmod>





# XDMoD OnDemand Integration

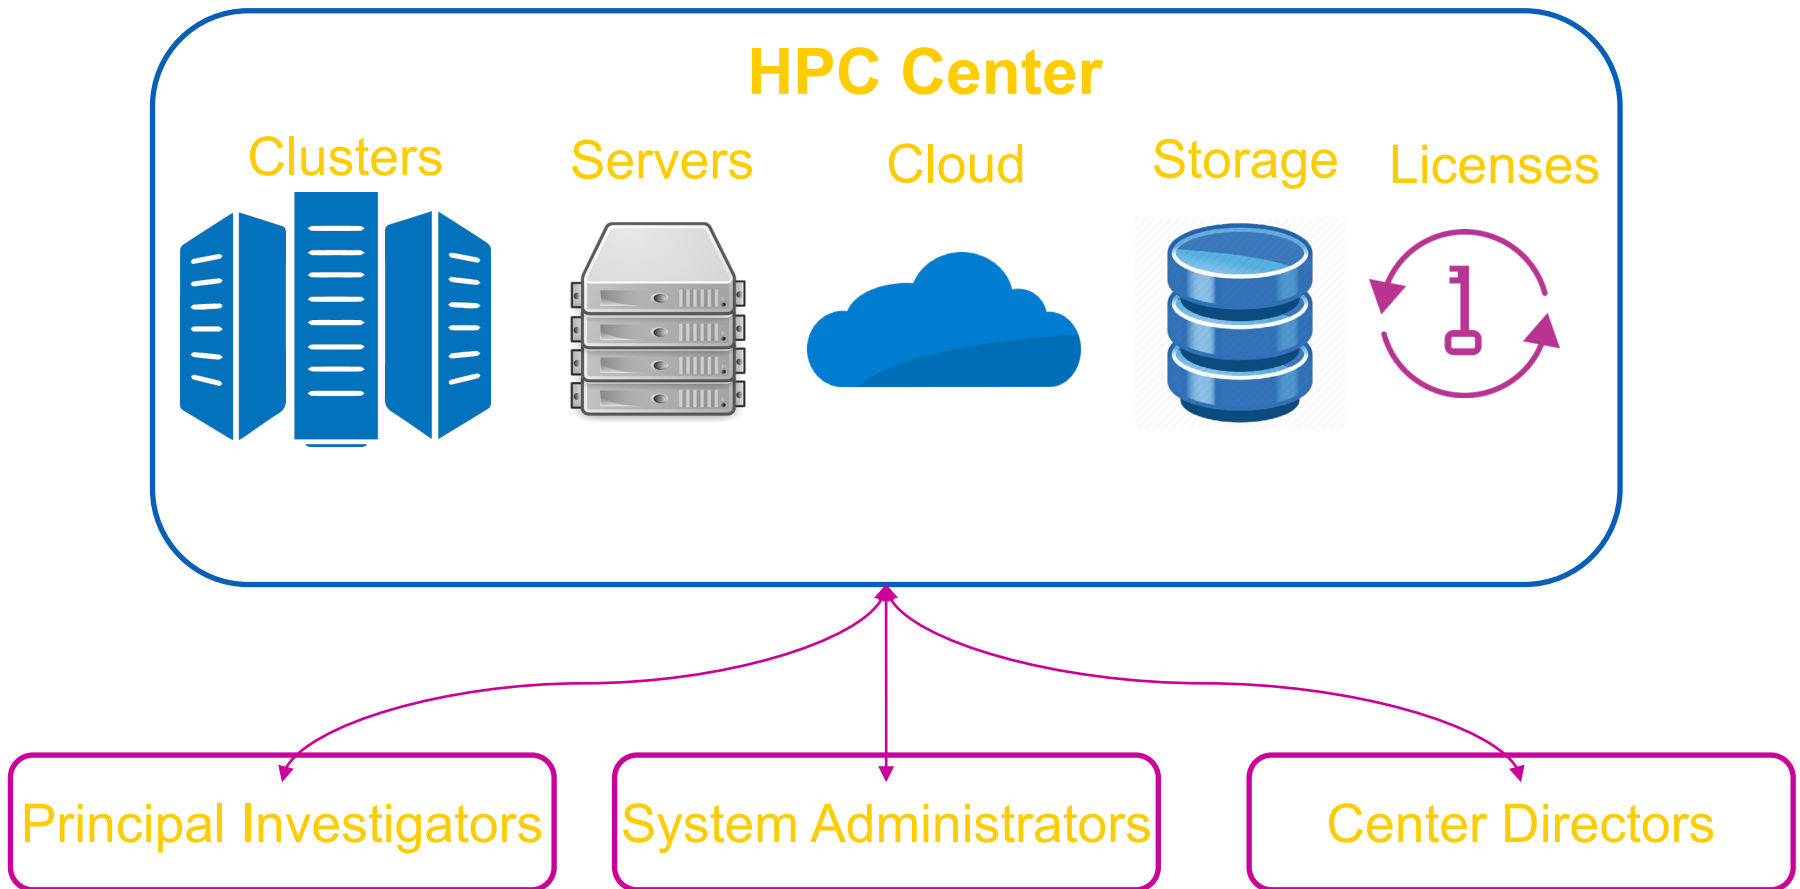
- Open OnDemand: “one stop shop for quick access to HPC”
  - OSU (Dave Hudak PI) & CCR (Tom Furlani coPI)
- Key program goal to integrate XDMoD into Open OnDemand.

The screenshot displays the Ohio Supercomputer Center (OSC) OnDemand dashboard. The top navigation bar includes links for Dashboard, Files, Jobs, Clusters, Interactive Apps, My Interactive Sessions, and All Apps. The main content area features a 'Message of the Day' section with news about COVID-19 support and classroom resources. To the right, there are three summary cards: 'XDMoD - Job Efficiency Report' showing 44.1% efficient and 55.9% inefficient jobs; 'XDMoD - Core Hours Efficiency Report' showing 71.2% efficient and 28.8% inefficient core hours; and 'XDMoD - Jobs - 2020-05-27 to 2020-06-26' which includes a table of job details.

ID	Cluster	Date	Time Used	CPU
10573480	owens	6/5	00:14:00	98.3
10573479	owens	6/5	00:02:08	29.7
10573468	owens	6/5	00:04:40	01.4
10572847	owens	6/5	00:30:55	00.8
10572013	owens	6/5	00:31:27	06.0
10572010	owens	6/5	00:00:03	N/A
10572007	owens	6/5	00:00:04	N/A
10539927	owens	6/1	02:00:41	08.1
10539673	owens	6/1	00:17:23	10.3
10539512	owens	6/1	00:44:25	09.3

# XDMoD Integration with ColdFront

## ColdFront Allocations utility




# Frontera Open XDMoD Instance

- XMS / TACC collaboration
  - XMS Frontera Open XDMoD instance established
- A data pipeline between TACC and local instance of Open XDMoD maintained by XMS
- Frontera data from the jobs realm and job performance realm are available in the XMS Frontera Open XDMoD instance
  - Frontera data validated
  - Full production instance has been established
- TACC personnel provided access to XMS Frontera Open XDMoD instance
- Accounts on XMS Frontera Open XDMoD instance created for NSF



# Frontera Open XDMoD Instance

- XMS / TACC collaboration: XMS Frontera Open XDMoD instance established
- Frontera data from the jobs realm and job performance realm are available in the XMS Frontera Open XDMoD instance
- From Feb-June 2020 Frontera managed ~90% daily utilization rate



2019-07-06 to 2020-07-01 Src: HPcDB. Powered by XDMoDHighcharts

# XMS Team

- **XD Metrics Service (XMS)**
  - **CCR:** Tom Furlani, Matt Jones, Steve Gallo, Bob DeLeon, Joe White, Jeff Palmer, Nikolay Simakov, Jeanette Sperhac, Ryan Rathsam, Ben Plessinger, Gregory Dean, Cynthia Cornelius, Abani Patra
  - **Indiana:** Gregor von Laszewski, Fugang Wang
  - **TACC:** Bill Barth, Todd Evans
- **NSF**
  - TAS: OCI I025159, SUPReMM: OCII203560
  - XMS: ACI-I445806,



# For More Information

- Monday Oct 19 at 3PM Jeanette Sperhac:  
“Improving XDMoD Science Gateway Monitoring”
- XSEDE XDMoD
  - <https://xdmod.ccr.buffalo.edu/>
- Open XDMoD
  - <https://open.xdmod.org/>
- XDMoD Help
  - [ccr-xdmod-help@buffalo.edu](mailto:ccr-xdmod-help@buffalo.edu)
- XDMoD/Open XDMoD Mailing List
  - <https://listserv.buffalo.edu/cgi-bin/wa?SUBED1=ccr-xdmod-list&A=I>





# **OPEN** **nDemand**

## Open, Interactive HPC via the Web

Alan Chalker, OSC

Eric Franz, OSC

Trey Dockendorf, OSC

Jeff Ohrstrom, OSC

Bob Settlage, VT

OSC has a job  
opening on the Open  
OnDemand team!

Please contact us if  
interested.



**Ohio Supercomputer Center**

An **OH-TECH** Consortium Member

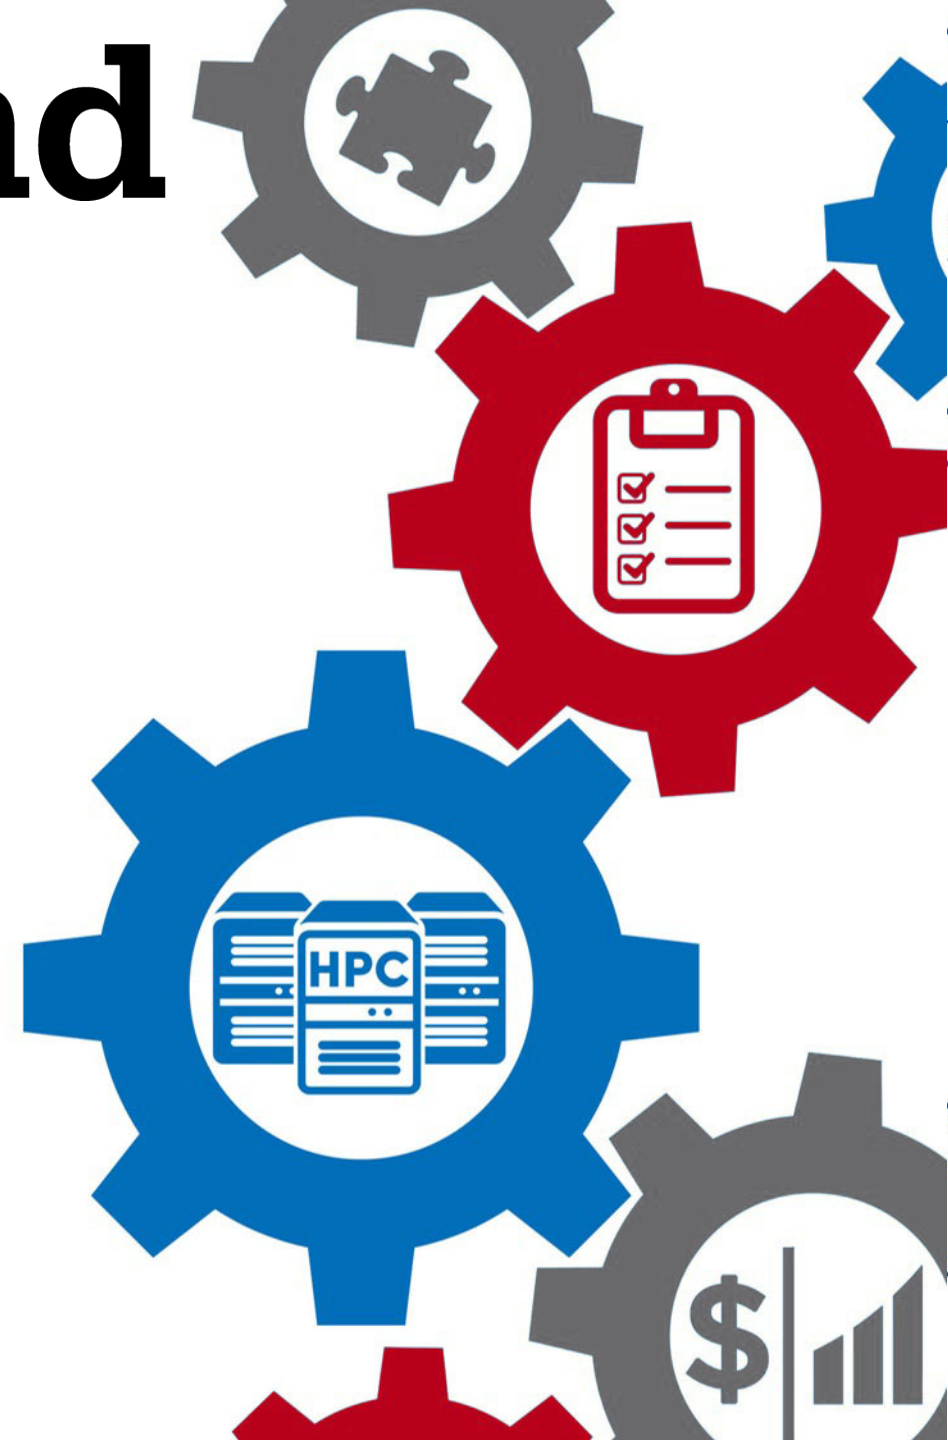


**VIRGINIA  
TECH™**



University at Buffalo

Center for Computational Research



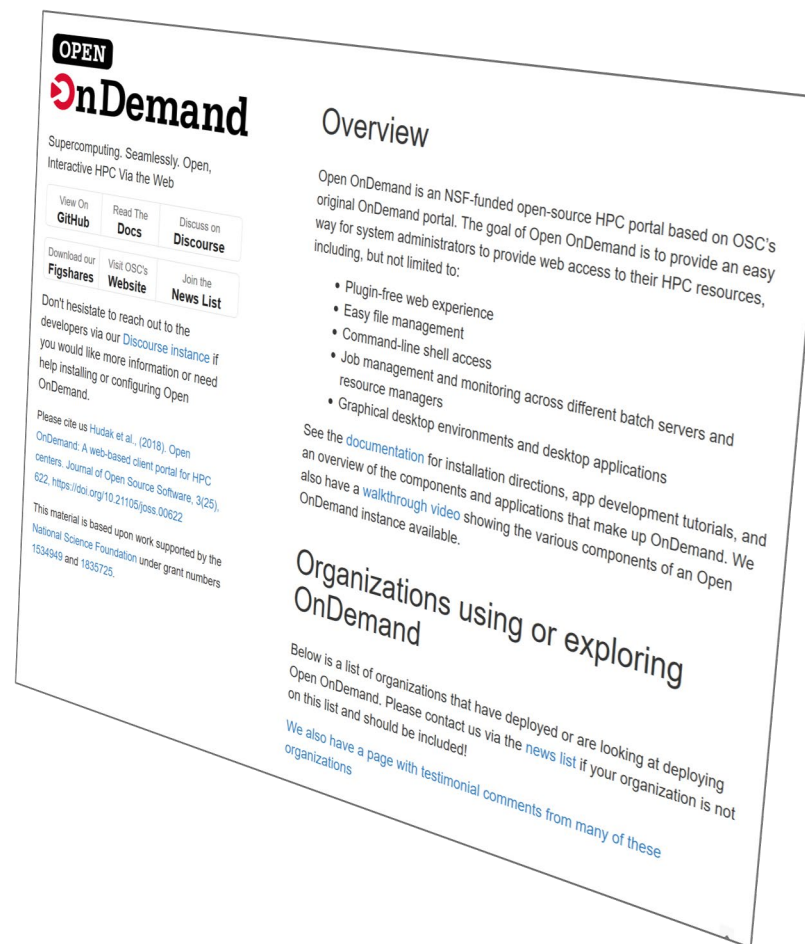


# OPENONDEMAND.ORG

Use our Discourse instance for help

Join our mailing list for updates

Our webinars are roughly quarterly

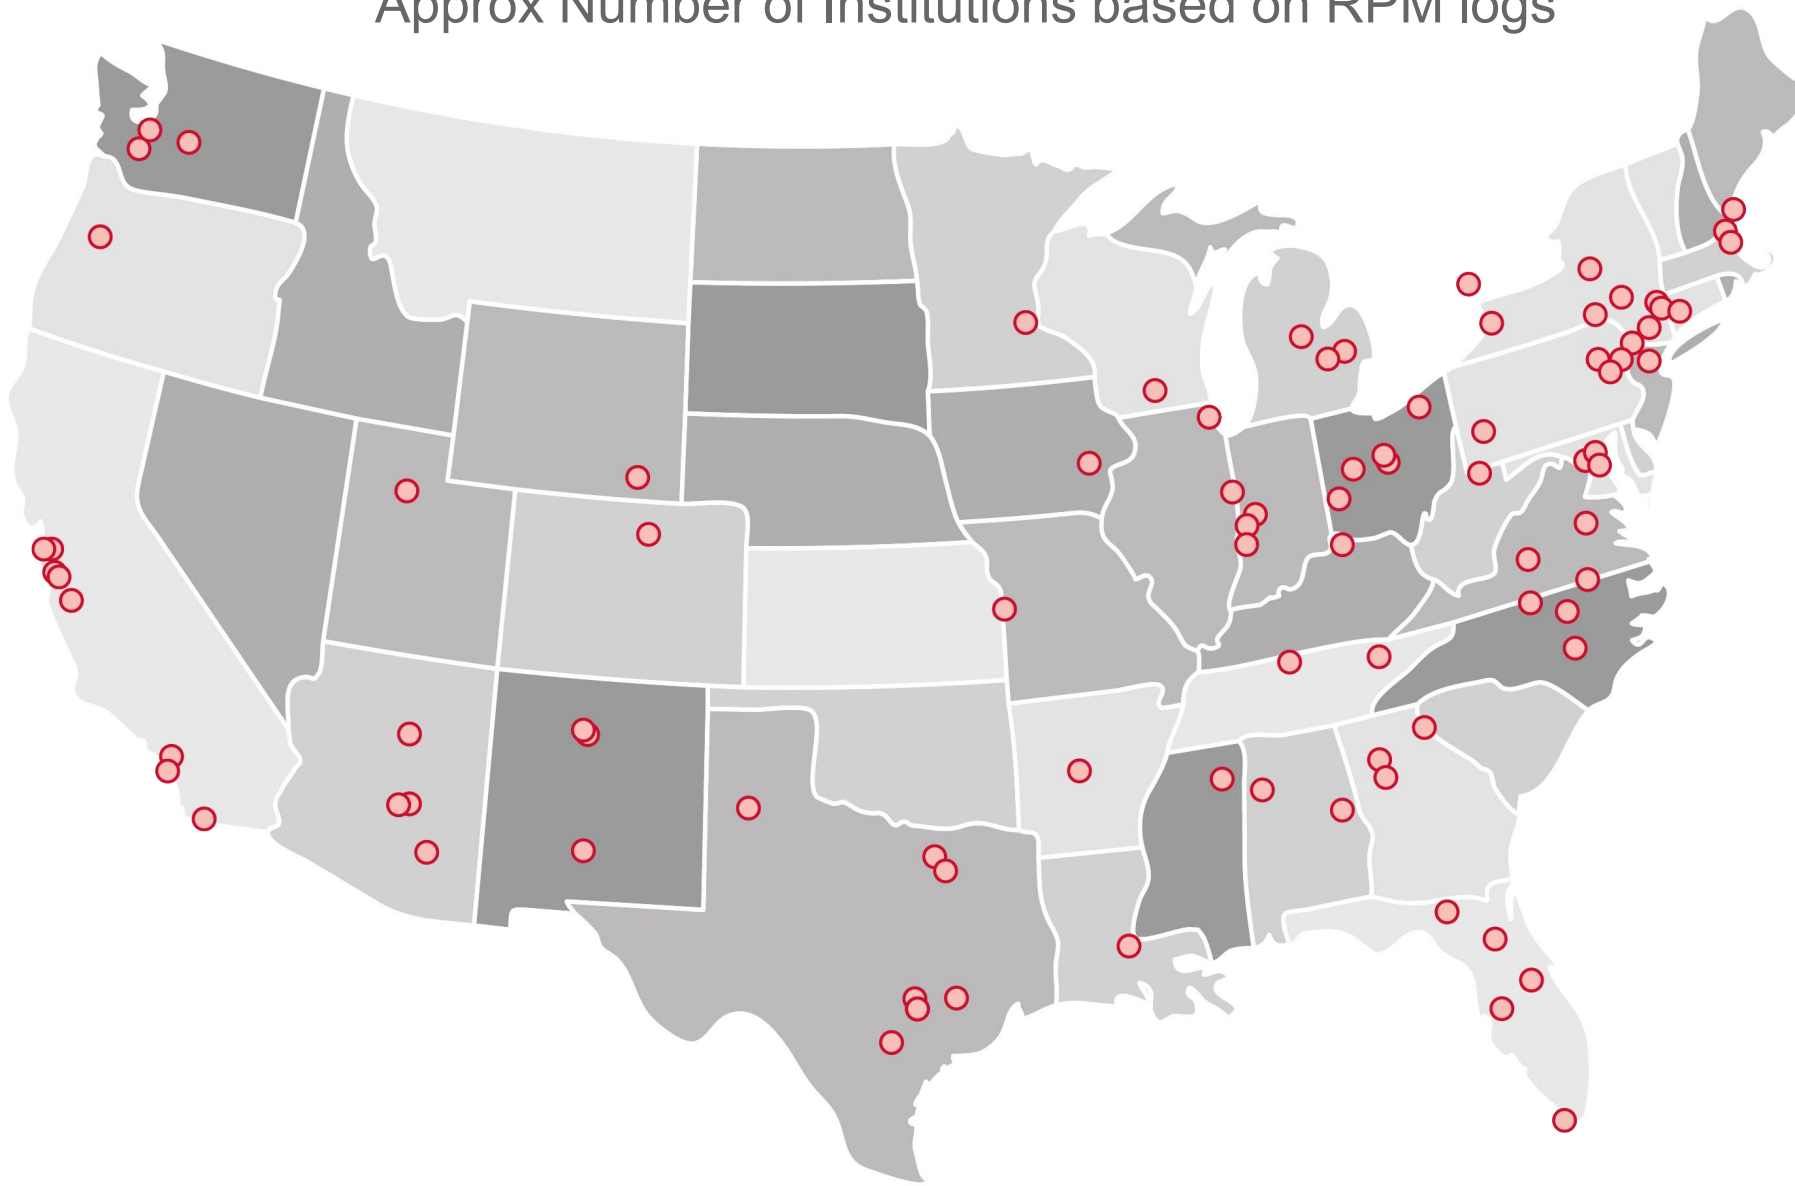


This work is supported by the National Science Foundation of the United States under the awards NSF SI2-SSE-1534949 and CSSI-Software-Frameworks-1835725.





Approx Number of Institutions based on RPM logs



- 136 unique US locations
- 70 unique international locations





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



University at Buffalo

Center for Computational Research



In Process of Installing

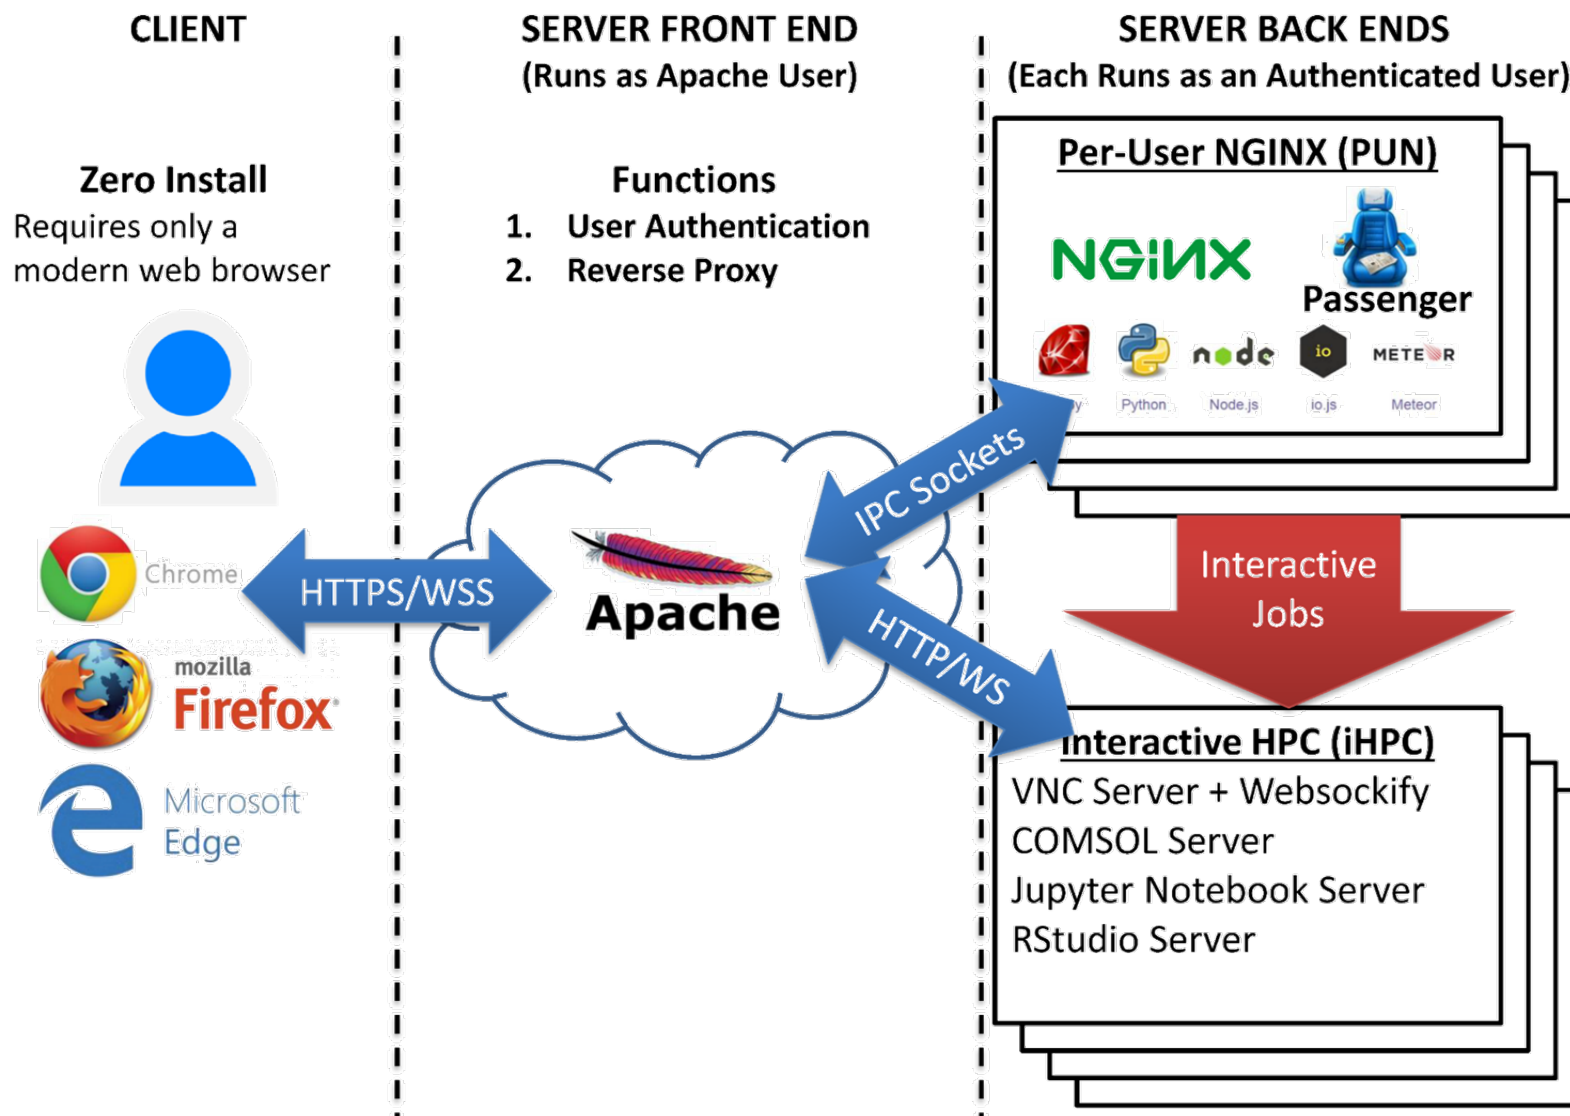
## Example Current Engagements and Deployments



## Production Deployments







# Open OnDemand 2.0 Project Overview

Previous three year NSF SI2 award (#1534949) to develop OnDemand 1.x

Awarded follow on NSF CSSI award (#1835725) to develop OnDemand 2.x

Project runs from Jan 2019 to Dec 2023

Collaborators include SUNY Buffalo and Virginia Tech

## Four areas

**Visibility:** Enhancing resource utilization visibility by integrating the existing Open XDMoD platform

**Scalability:** support more types of computing resources and software

**Accessibility:** appeal to more scientists in more fields of science

**Engagement:** establish community of departmental, campus and national HPC users and administrators







# Items ‘Coming Soon’ or Recently Added

Version	System Stuff	Apps
V1.7 (June 1)	Linux host adapter Keycloak identity brokering Ansible role OpenHPC integration	
V1.8 (August 17)	Dashboard with XDMoD DEX authentication Easier debugging interactive apps	App submission to cluster set Job composer with XDMoD Visual Studio Code app (beta)
V2.0 (December 1)	Kubernetes adapter SSL+auth abstraction for apps	New launch interface UX Files app replacement Dashboard widgets
Current “OSC only” features	System status with GPUs OpenStack Globus Integration	Stata, Tensorboard, QGIS, Render, Galaxy, Visual Studio Code Server, R Shiny



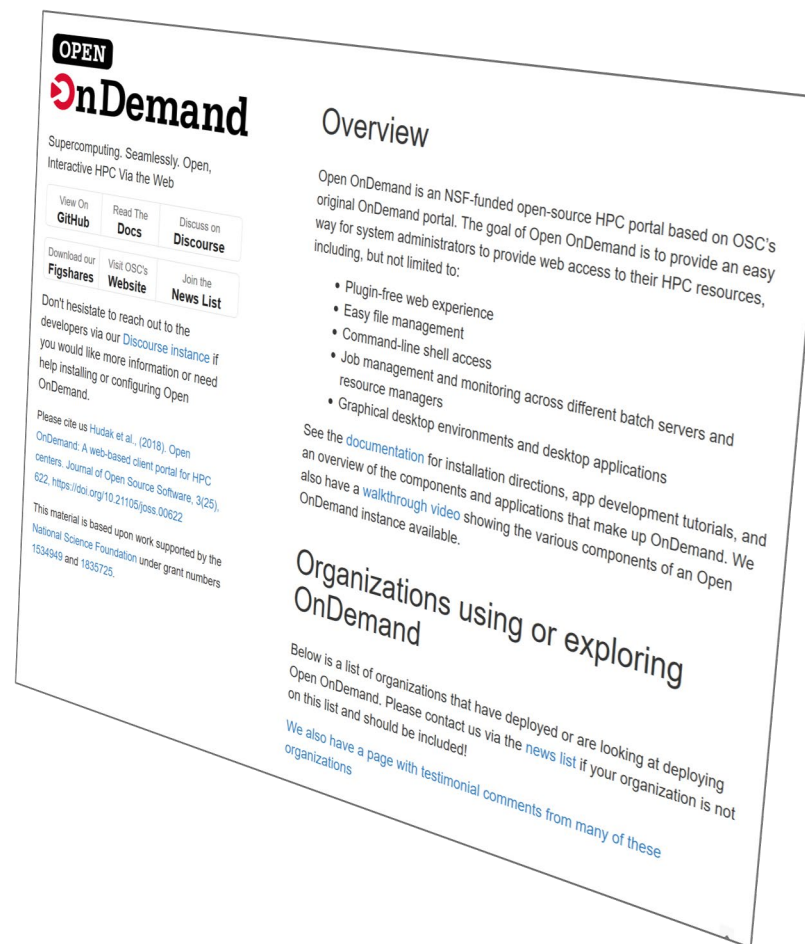


# OPENONDEMAND.ORG

Use our Discourse instance for help

Join our mailing list for updates

Our webinars are roughly quarterly



This work is supported by the National Science Foundation of the United States under the awards NSF SI2-SSE-1534949 and CSSI-Software-Frameworks-1835725.





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



University at Buffalo

Center for Computational Research



**VIRGINIA  
TECH™**

# Usage Overview Demo

Bob Settlage, VT





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



University at Buffalo


Center for Computational Research



**VIRGINIA  
TECH™**

# Installing OnDemand

Eric Franz, OSC





## Installation: requirements

- **Hardware**
  - OSC – VMware VM w/ 16 vCPU & 64GB RAM (~120 concurrent users)
- **Operating System**
  - Red Hat based OS – version 7 or 8
- **Batch client software**
  - Needs to be able to submit and query jobs (sbatch, squeue, qsub, qstat, etc)
- **Shared filesystems**
  - Home directories and other cluster filesystems
- OSC uses NFSv4 home and GPFS over NFS using CES
  - **An OnDemand server is closer to a login node than a standard web server**





## Installation: Batch schedulers supported

- Slurm
- Torque
- PBSPro
- SGE / UGE
- LSF
- Linux Host (processes started via ssh, tmux, and Singularity)
- Kubernetes (in development)
- CloudyCluster CCQ (in development)



# Installation: packages

- **YUM repos and RPMs for RHEL7/8**
  - Enable Software Collections (SCL) repos
  - Install OnDemand YUM repo then RPM package
    - We provide Passenger and NGINX as SCL versions were deprecated
  - yum install ondemand
- **Ansible role for other OSs**
- **Separate rpm packages for default authentication using Dex and enabling SELinux support**
- **Puppet module available**



# Installation: authentication

- **Three steps:**
  1. Configure Apache module
  2. Setup User Mapping
  3. Configure Logout
- Examples
  - OpenID Connect using Keycloak or Dex
  - CAS
  - Shibboleth
  - CILogon



# Installation: configuration & customization

- **Configuration files**
  - YAML file per cluster – used to define login and batch environment
  - YAML file to generate Apache configuration
  - YAML to configure Per User NGINX environment
- **All configuration files under /etc/ood**
- **App specific customization mainly managed using env variables**
  - /etc/ood/config/apps/dashboard/env
    - MOTD\_PATH="/etc/motd"
    - MOTD\_FORMAT="markdown"





## Installation: Branding


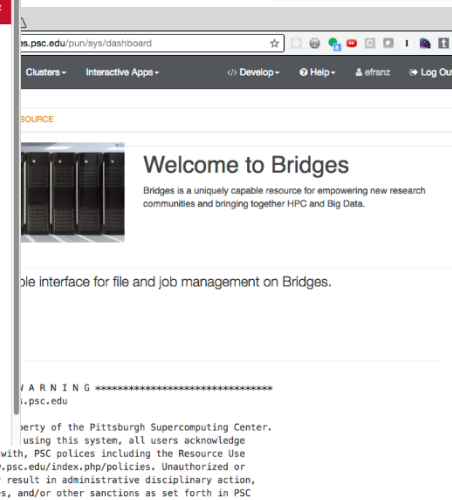
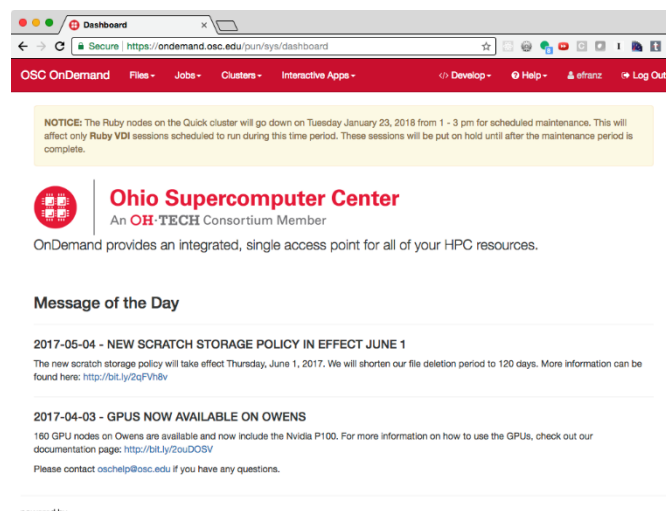
- Institution logo
- Navbar color
- Portal name
- Display MOTD on front page (/etc/motd)
- Display announcements on front page
- Display quota and balance warnings



OPEN

OnDemand

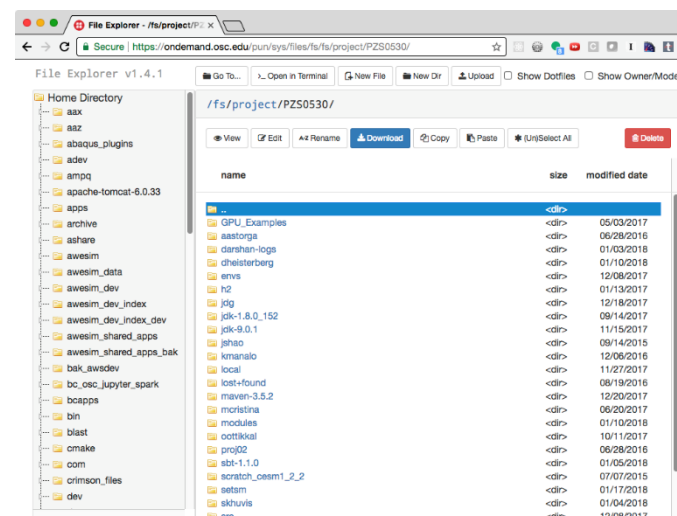
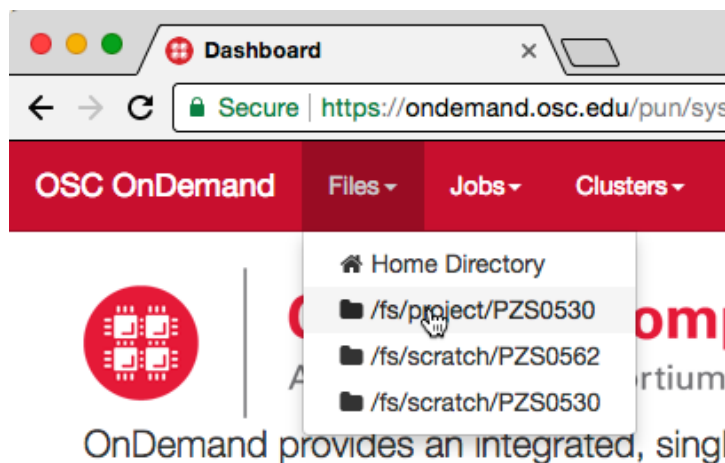
OnDemand provides an integrated, single access point for all of your HPC resources.





## Installation: Files shortcuts

- Add directory shortcuts to open Files app in home directory, scratch space, and project space





## Installation: considerations for production

- **Ensure short wait times for Interactive apps**
  - OSC uses “quick” batch environment with dedicated nodes
    - Moab tuned for 3s scheduling interval
  - SLURM sites could use dedicated partition or high priority QoS
  - Another possible solution is over subscription
  - Management of interactive access to HPC resources is not a solved problem
- **Separate test and production deployments**
  - Upgrade and verify test deployment first





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



University at Buffalo

Center for Computational Research



**VIRGINIA  
TECH™**

# Configuring software to be available in OnDemand

Eric Franz and Jeff Ohrstrom OSC



## Configuring software to be available in OnDemand

- New software is made available through OnDemand by adding new “apps”
- Users can develop and run apps in their home directory
- Admins can publish apps by copying them to the OnDemand web host’s local disk



# Configuring software: Types of apps

- **Interactive App Plugins**
  - Consists of a job template and configuration files
  - Submits a batch job which launches VNC GUI app or web server on compute node and provides user link to connect
- **Passenger web apps written in Python, Ruby, or Node.js**
  - run as the user - they are acting behalf of the user
  - do not need to manage authentication or authorization
  - write any app specific data to user dirs (\$HOME, \$SCRATCH)



# Demonstration: Create a Jupyter “Interactive App Plugin”





## BACKUP: Jupyter Tutorial: Get the App working

- Jupyter example application doesn't work out of the box
  - Configure it to use this cluster
  - Configure it to use the correct Jupyter installation
- The card is shown when a successful Jupyter job is launched

**HPC Tutorial Jupyter (2)** 1 node | 1 core | Running

**Host:** >\_cpn01 Delete

**Created at:** 2020-07-21 19:27:37 UTC

**Time Remaining:** 59 minutes

**Session ID:** b71ea2ba-83ec-40ea-9011-7dd5b834b31f

[Connect to Jupyter](#)





# BACKUP: Jupyter Tutorial: Modify the Partition

- Change the partition element to be a select dropdown instead of a text field

Partition

Compute

Compute

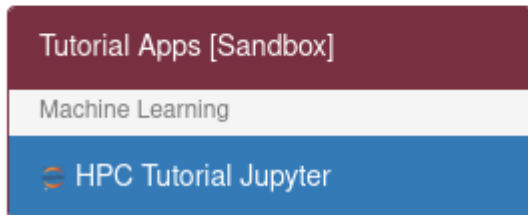
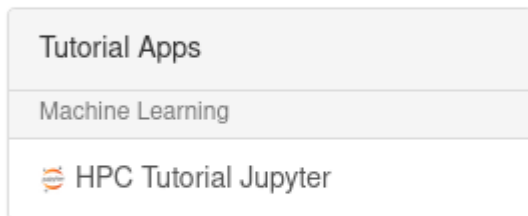
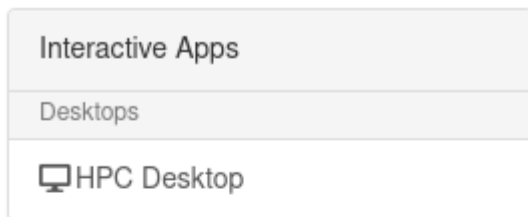
Debug





## BACKUP: Jupyter Tutorial: Deploy to production

- Deploy the app to production for other users





## BACKUP: Jupyter Tutorial: Set the memory request for the job

- Use the script.native attributes to set the --mem SLURM argument

Memory (MB)

600



RSS Memory

Launch

\* The HPC Tutorial Jupyter session data for this session can be accessed under the [data root directory](#).



# BACKUP: Jupyter Tutorial: Limit the number of nodes

- Put an upper limit on the number of nodes allowed

**Number of nodes**

17

⬆  
⬇  
⬆

Please select a value that is no more than 2.

the session starts

Launch

\* The HPC Tutorial Jupyter session data for this session can be accessed under the [data root directory](#).





## BACKUP: Jupyter Tutorial: Add a checkbox to start JupyterLab

- Add a checkbox so users can boot JupyterLab or Jupyter Notebook

☐ Use JupyterLab instead of Jupyter Notebook?

JupyterLab is the next generation of Jupyter, and is completely compatible with existing Jupyter Notebooks.

Launch

\* The HPC Tutorial Jupyter session data for this session can be accessed under the [data root directory](#).





## BACKUP: Jupyter Tutorial: Delete unused fields

- Delete unused fields to clean up the form

### Partition

Compute

### Number of hours

1

### Number of nodes

1

### Memory (MB)

600

RSS Memory

☐ Use JupyterLab instead of Jupyter Notebook?

JupyterLab is the next generation of Jupyter, and is completely compatible with existing Jupyter Notebooks.

Launch

\* The HPC Tutorial Jupyter session data for this session can be accessed under the [data root directory](#).





# Create a Passenger app

- “Status apps” that execute a command and display formatted output to the user can be very useful and easy to develop
- Many third party apps written in Python, Ruby, NodeJS can be configured to run via Passenger in OnDemand





## Example Status apps: Tufts many custom report apps



Galaxy

Globus File Transfer

Module List

Quota Increase

⚙ Quota Report

⚙ Scheduler Info

⚙ Show Groups

Reports

⚙ Inventory

⚙ Top Users

⚙ Utilization by User

Open OnDemand / Show Groups

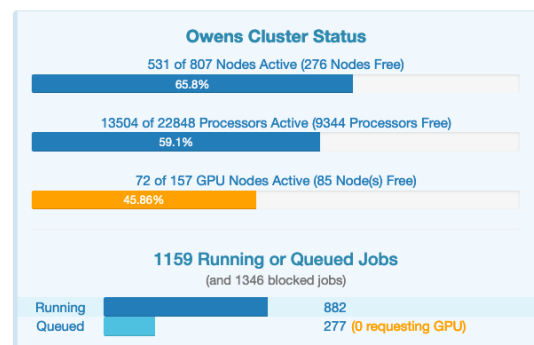
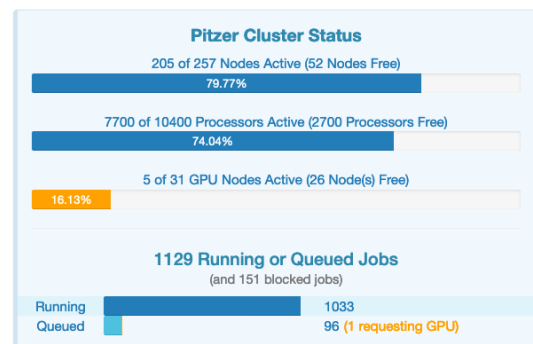
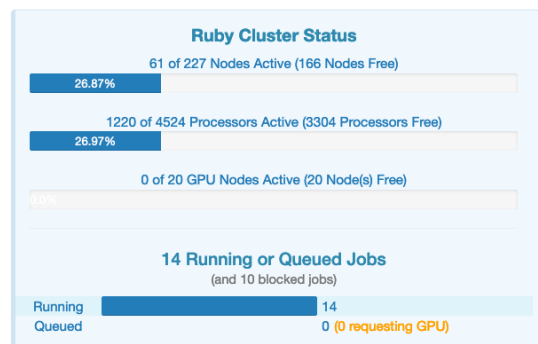
Using layout: Mon Jan 29 2018 13:59:28 GMT-05

```
chbe193
cohnlab
datalab
duchinlab
facstaff
fmri
galaxydev
gaussian
georgakoudilab
grj
gromacs
heldweinlab
hep
hrilab
isberg
levinlab
marchesini_lab
math150inst
perseus
rgts
schwob_lab
sokolevlab
student
train01
ttsworkshop
us@tlas1
vireos
```





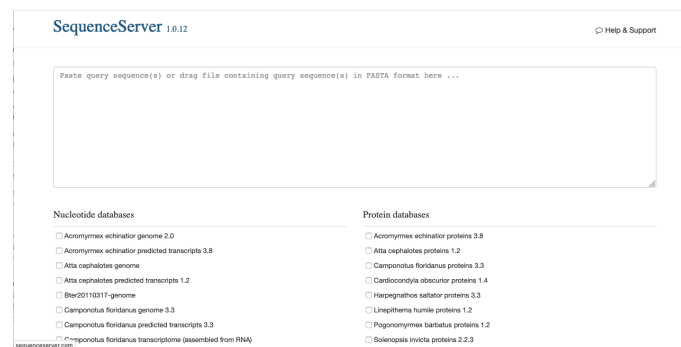
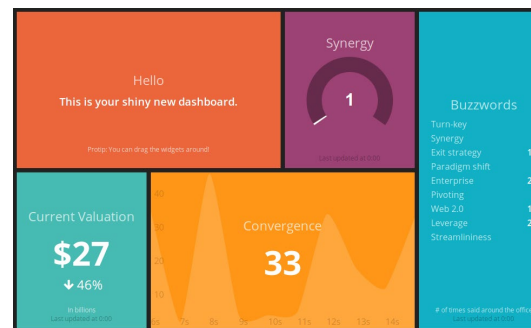
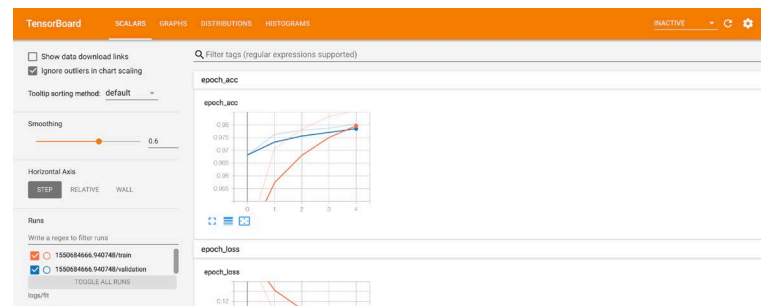
# Example Status apps: OSC System Status app





## Example third party apps:

- Tensorboard (in Python)
- Smashing dashboard  
(formerly Dashing from  
Shopify, in Ruby)
- SequenceServer (in Ruby)
- In OnDemand 2.0 we hope  
to expand support for  
launching Passenger  
apps in any language



## Demonstration: Create a Passenger app

1. Create a simple status app using Ruby
2. Review examples of apps in Python and NodeJS
3. Deploy the status app to production
4. Learn about authorization in OnDemand



# XDMoD and OnDemand Integration

Eric Franz, OSC



**Ohio Supercomputer Center**

An **OH·TECH** Consortium Member




University at Buffalo

Center for Computational Research



**VIRGINIA  
TECH™**





# Overview of integration

- Presents job efficiency reports from XDMoD directly on the OnDemand dashboard
- Integration enabled by user being logged into both XDMoD and OnDemand
- Only works if authenticated using same OpenID Connect or SAML Identity Provider
- This should be available in OnDemand 1.8 and XDMoD 9





## Enabling the XDMoD reports on OnDemand dashboard

1. Configure OnDemand with XDMoD host URL in PUN  
`/etc/ood/config/nginx_stage.yml`
2. Configure OnDemand with XDMoD resource id in each cluster config  
`/etc/ood/config/clusters.d/hpc.yml`
3. Add OnDemand host as domain to XDMoD portal settings for CORS  
`/etc/xdmod/portal_settings.ini`
4. Configure identity provider to include OnDemand host in HTTP Content-Security-Policy for frame-ancestors since OnDemand uses iFrames to trigger SSO with XDMoD when a user logs in





# XDMoD reports on the OnDemand dashboard

Dashboard

OSC OnDemand Files Jobs Clusters Interactive Apps Develop Help Logged in as efranz Log Out

Ohio Supercomputer Center  
An OH-TECH Consortium Member

OnDemand provides an integrated, single access point for all of your HPC resources.

Message of the Day

2020-07-31 - Backup failures for Project on August 1st and 2nd

OSC experienced backup failures on our GPFS file systems (both Project file systems, /fs/project and /fs/less) the mornings of August 1st and 2nd. The underlying cause was identified and backups were operating as expected the morning of August 3rd. As a result of these failed backups, OSC will not be able to complete some file restore requests for files changed between approximately 2020-07-31 02:30 through 2020-08-02 02:30.

2020-03-16 - OSC support during COVID-19 crisis

The Ohio Supercomputer Center serves as a critical resource for the public good and, as such, is striving to provide extraordinary support in light of the ongoing COVID-19 crisis. OSC staff are currently working from home but fully expect clients will see no disruption in our services to support this effort.

Examples of the types of special support OSC can provide include: - Priority, unbilled access to OSC computational and storage resources for COVID-19 research - Flexible billing terms and prices for clients anticipating negative economic impacts - Remote, virtual computing lab resources for classroom instructors and educators - Connections to domain experts in academia and industry

Please don't hesitate to contact OSC at [oschelp@osc.edu](mailto:oschelp@osc.edu) or (800) 686-6472 for more information on this initiative. Please also distribute this message via any communication channel you to which you might have access so that it can be distributed as widely as possible.

CLASSROOM RESOURCES FOR DISTANCE LEARNING

If your class has lost or limited access to computer labs, the Ohio Supercomputer Center might be able to help by providing no-cost access to cloud computing resources. Classes and workloads of any size can gain access. OSC's web-browser interface to its substantial Linux computer systems provides novice users with virtual desktops preloaded with applications, such as MATLAB, RStudio, or Jupyter Notebook.

As an example, an OSU undergrad statistics class recently used iPads to remotely access RStudio on OSC systems. We can provide online demonstrations or evaluations and potentially add additional software packages.

Please contact [OSChelp@osc.edu](mailto:OSChelp@osc.edu) to talk to OSC about distance-learning support options available to you.

2020-03-09 - Huge memory nodes partial scheduling

Beginning on Tuesday, March 10th users are able to run jobs using less than a full huge memory node on both the Owens and Pitzer clusters.

Please consider your request more carefully when you plan to use a huge memory node, and specify the resources based on what you will use. Please check our documentation for more detailed guidance: [https://www.osc.edu/resources/technical\\_support/supercomputers/owens/batch\\_limit\\_rules](https://www.osc.edu/resources/technical_support/supercomputers/owens/batch_limit_rules) [https://www.osc.edu/resources/technical\\_support/supercomputers/pitzer/batch\\_limit\\_rules](https://www.osc.edu/resources/technical_support/supercomputers/pitzer/batch_limit_rules)

2019-10-03 - Wall-time Accuracy Information

Due to current usage patterns on OSC's Owens Cluster, users may benefit from improving the wall-time accuracy of their submitted jobs. A job has high wall-time accuracy if the requested wall time comes near the actual wall time when the job completes.

The benefit of improving one's wall-time accuracy is that the requested wall time is taken into account when scheduling jobs to run. Usually, jobs that request shorter wall time, ceteris paribus, will wait for less time than jobs requesting longer wall times. One can investigate the accuracy of their previous jobs using the following HOW TO webpage:

Jobs Efficiency Report - 2020-07-25 to 2020-08-24

54.5% efficient 45.5% inefficient

15 inefficient jobs / 33 total jobs

Core Hours Efficiency Report - 2020-07-25 to 2020-08-24

18.8% efficient 81.2% inefficient

65.5 inefficient core hours / 80.7 total core hours

Recently Completed Jobs - 2020-07-25 to 2020-08-24

ID	Name	Date	CPU
11030363	ondemand/rys /dashboard	8/6	99.4
11025209	wlag/weld_predictor /sys/bc_psc_jupyter	8/5	99.4
14021	ondemand/rys /myjobs/default	8/5	N/A
14020	ondemand/rys /myjobs/default	8/5	N/A
10978802	ondemand/rys /dashboard	8/4	91.7
10974801	wlag/weld_predictor	8/3	91.8
10974800	wlag/weld_predictor	8/3	99.3
10936746	STDIN	8/3	99.3
10935957	ondemand/rys /myjobs/default	8/3	99.8
10935689	ondemand/rys /myjobs/default	8/3	99.8

Showing first 10 of 33 jobs. See your XDMoD dashboard for more information.

Open XDMoD

<https://xdmod-test.hpc.osc.edu/index.php>

XDMoD Hello, Eric Franz (logout)

Ohio Supercomputer Center My Profile About Roadmap Contact Us Help

Dashboard Usage Metric Explorer App Kernels Report Generator Job Viewer About

Job Efficiency Report - 2020-07-25 to 2020-08-24

Wait Hours - 2020-07-25 to 2020-08-24

Job Size

Wait Times by queue - 2020-07-25 to 2020-08-24

Queue

Wait Time Per Job (hours)

Jobs - 2020-07-25 to 2020-08-24

Job Identifier	Start	End	CPU
owens-11030363	2020-08-21 15:21:38	2020-08-21 16:15:15	●
owens-11025209	2020-08-20 17:03:52	2020-08-20 17:05:10	●
pitzer-exp-14021	2020-08-20 10:27:05	2020-08-20 10:27:07	N/A
pitzer-exp-14020	2020-08-20 10:26:50	2020-08-20 10:26:53	N/A
owens-10978802	2020-08-19 10:45:04	2020-08-19 10:52:35	●
owens-10974801	2020-08-18 14:51:48	2020-08-18 14:53:01	●
owens-10974800	2020-08-18 14:51:17	2020-08-18 14:52:31	●
owens-10936746	2020-08-11 15:53:32	2020-08-11 17:54:31	●
owens-10935957	2020-08-11 11:24:53	2020-08-11 12:36:10	●
owens-10935689	2020-08-11 10:10:49	2020-08-11 11:11:32	●

Page 1 of 4

Displaying 1 - 10 of 33



# XDMoD job links in the OnDemand Job Composer

Job Composer

OSC OnDemand / Job Composer Jobs Templates Help

## Jobs

+ New Job + Create Template

Edit Files Job Options Open Terminal Submit Stop Delete

Show 25 entries Search:

Created	Name	ID	Cluster	Status
May 28, 2020 10:37am	(default) Simple Sequential Job		Pitzer	Not Submitted
May 28, 2020 10:34am	(default) Simple Sequential Job		Owens	Not Submitted
May 19, 2020 3:28pm	auto - Simple Sequential Job	10276299.owens-batch.ten.osc.edu - XDMoD	Owens	Completed
May 19, 2020 3:28pm	auto - Simple Parallel Job	10276298.owens-batch.ten.osc.edu - XDMoD	Owens	Completed
May 19, 2020 3:28pm	auto - Parallel job with MPI and OpenMP	10276297.owens-batch.ten.osc.edu - XDMoD	Owens	Completed
May 19, 2020 3:28pm	auto - MPI Hello World	10276296.owens-batch.ten.osc.edu - XDMoD	Owens	Completed
May 19, 2020 3:28pm	auto - Hello OpenMP World in C	10276295.owens-batch.ten.osc.edu - XDMoD	Owens	Completed
May 19, 2020 3:28pm	auto - Basic Star-CCM+ Serial Job - Owens	10276294.owens-batch.ten.osc.edu - XDMoD	Owens	Completed
May 19, 2020 3:28pm	auto - Basic Star-CCM+ Parallel Job - Owens	10276293.owens-batch.ten.osc.edu - XDMoD	Owens	Completed

### Job Details

10276299.owens-batch.ten.osc.edu - XDMoD

Job Name: auto - Simple Sequential Job

Submit to: Owens

Account: Not specified

Script location: /users/PZ50562/efranz/ondemand/src/ondemand/apps/myjobs/data

Script name: sequential\_job.sh

Folder Contents:

- basic\_sequential.o10276299
- message.in
- message.out
- myscript.sh
- sequential\_job.sh

Open XDMoD

https://xdmod-test.hpc.osc.edu/index.php/job\_viewer?realm=SUPREMM&recordid=1&jobid=12936789&info

XDMoD Hello, Eric Franz (logout)

Dashboard Usage Metric Explorer App Kernels Report Generator Job Viewer About

Search History

- owens-4143443
- owens-4173938
- owens-4952616
- owens-1113203
- Linked Search
  - pitzer-931595
  - owens-8366777
  - owens-8500630
  - owens-8518552
  - pitzer-1136030
  - owens-9951541
  - owens-9951392
  - owens-10276299
  - Accounting data
  - Job script
  - Executable information
  - Peers
    - owens-10573479
    - owens-10573480
    - owens-4721524
    - owens-7153965
    - owens-4721524



owens-10276299

CPU User: N/A Walltime Accuracy: 0.133 CPU User Balance: N/A Memory Headroom: N/A

Metric Missing: Unknown Reason

Accounting data Job script Executable information Peers


Key	Value
Category: Administration (11 Items)	
Account	PZ50714
Local Job Id	10276299
Organization	Ohio Supercomputer Center
Resource	Owens
Hierarchy Bottom Level	Unknown
PI	PZ50714
PI Institution	Ohio Supercomputer Center
Timezone	US/Eastern
User	efranz
User Institution	Ohio Supercomputer Center
Username	efranz
Category: Allocated resource (6 Items)	
Shared	True
Cores	1
Nodes	1
Total Cores Available	0
Cpu Time	8 seconds
Node Time	8 seconds
Category: Executable (5 Items)	
Name	ondemand/sys/myjobs/basic_sequential
Application	NA
Executable	NA
Exit Status	0
Working directory	NA
Category: Requested resource (3 Items)	
Requested Nodes	1
Requested Wall Time	1 minute
Queue	serial

Description

- Accounting data: Shows information about the job that was obtained from the resource manager. This includes timing information such as the start and end time of the job as well as administrative information such as the user that submitted the job and the account that was charged.

# Benefits of integrating XDMoD and OnDemand

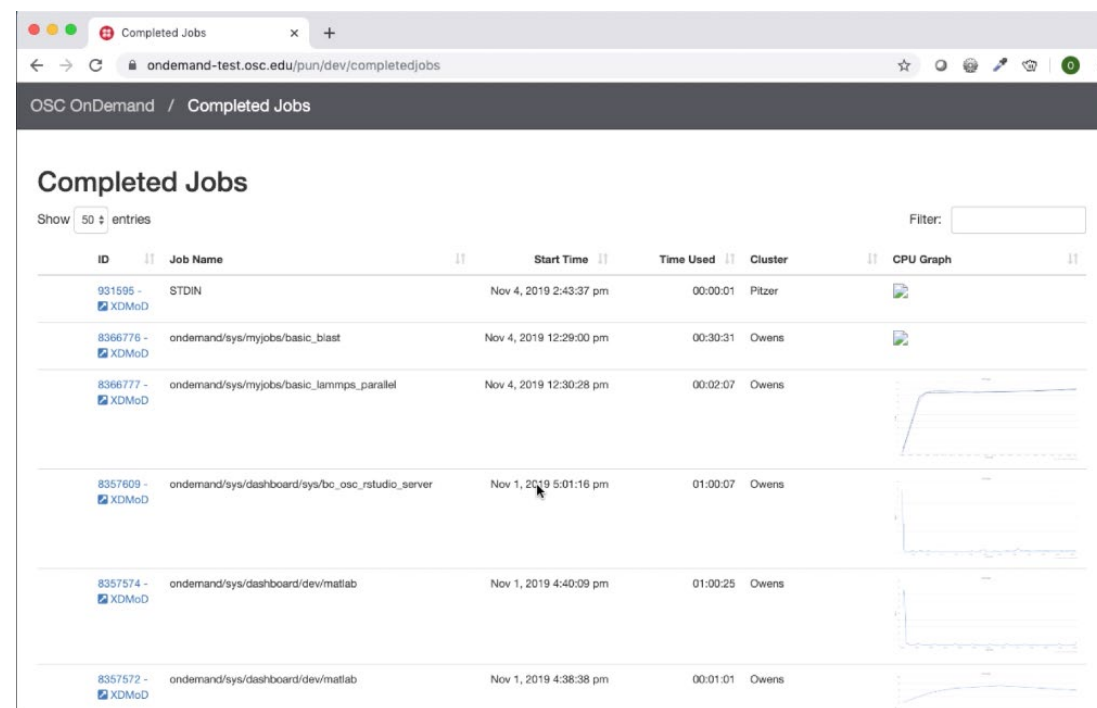
- 1. Encourage users, even those new to HPC to access to historical job information through XDMoD
- 2. Provide faster access to relevant job information XDMoD from OnDemand
- 3. Ensure that users with poorly performing jobs are encouraged to fix them by presenting reports with red graphs every time they log into OnDemand





## Future plans

1. Completed Jobs App
2. Server side integration to address Single Sign On problems by moving communication to the server (pending feedback)
3. Provide OnDemand usage metrics through XDMoD



The screenshot shows a web browser window with the URL `ondemand-test.osc.edu/pun/dev/completedjobs`. The page title is "OSC OnDemand / Completed Jobs". Below the title, there is a "Completed Jobs" section with a "Show 50 entries" dropdown and a "Filter:" input field. The main content is a table with columns: ID, Job Name, Start Time, Time Used, Cluster, and CPU Graph. The table lists several jobs, including "STDIN", "ondemand/sys/myjobs/basic\_blast", "ondemand/sys/myjobs/basic\_lammps\_parallel", "ondemand/sys/dashboard/sys/bo\_osc\_rstudio\_server", "ondemand/sys/dashboard/dev/matlab", and "ondemand/sys/dashboard/dev/matlab". Each job entry has a small "XDMoD" icon next to the ID. The "CPU Graph" column shows a small line graph for each job.

ID	Job Name	Start Time	Time Used	Cluster	CPU Graph
931595 - XDMoD	STDIN	Nov 4, 2019 2:43:37 pm	00:00:01	Pitzer	
8366776 - XDMoD	ondemand/sys/myjobs/basic_blast	Nov 4, 2019 12:29:00 pm	00:30:31	Owens	
8366777 - XDMoD	ondemand/sys/myjobs/basic_lammps_parallel	Nov 4, 2019 12:30:28 pm	00:02:07	Owens	
8357609 - XDMoD	ondemand/sys/dashboard/sys/bo_osc_rstudio_server	Nov 1, 2019 5:01:16 pm	01:00:07	Owens	
8357574 - XDMoD	ondemand/sys/dashboard/dev/matlab	Nov 1, 2019 4:40:09 pm	01:00:25	Owens	
8357572 - XDMoD	ondemand/sys/dashboard/dev/matlab	Nov 1, 2019 4:38:38 pm	00:01:01	Owens	

Find more ways to help users optimize their jobs!







## 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](#)
- We gratefully acknowledge the partnership with [Virginia Tech](#) on our current joint NSF project



# Thank you...

OSC, VT, and UB staff and students for helping with the tutorial today!



**Ohio Supercomputer Center**

An **OH·TECH** Consortium Member




University at Buffalo

Center for Computational Research



**VIRGINIA  
TECH™**



# Tutorial Staff:

Andrew Bruno, UB  
Alan Chalker, OSC  
Andrew Collins, OSC  
Robert DeLeon, UB  
Trey Dockendorf, OSC  
Eric Franz, OSC  
David Hudak, OSC

Matt Jones, UB  
Jeff Ohrstrom, OSC  
Ben Plessinger, UB  
Dori Sajdak, UB  
Bob Settlege, VT  
Joseph White, UB



**Ohio Supercomputer Center**

An OH·TECH Consortium Member




University at Buffalo

Center for Computational Research



**VIRGINIA  
TECH™**





## How to reach us:

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

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

ColdFront - <https://github.com/ubccr/coldfront>

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

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

Virginia Tech – Advanced Research Computing - <https://arc.vt.edu/>



# Other places you'll find us at Gateways20

**Monday, October 19, 3-3:15pm**

Improving Science Gateway Monitoring, Caveats and Goals - <https://sched.co/dcY7>  
*Jeanette Sperhac, UB CCR - XDMoD*

**Poster Session:**

Cloud HPC with Open OnDemand and Cloudy Cluster  
*Jeff Ohrstrom, OSC - OnDemand*



**Join the staff & developers of  
each product immediately  
following this tutorial**

**Zoom links are in QiqoChat**



**Ohio Supercomputer Center**

An **OH·TECH** Consortium Member




University at Buffalo

Center for Computational Research



**VIRGINIA  
TECH™**





# Thank you for attending!



**Ohio Supercomputer Center**

An **OH·TECH** Consortium Member



University at Buffalo

Center for Computational Research



**VIRGINIA  
TECH™**

