Ohio Supercomputer Center
An OH·TECH Consortium Member

University at Buffalo
Center for Computational Research

VIRGINIA TECH

https://www.osc.edu/ https://openondemand.org/ https://buffalo.edu/ccr https://arc.vt.edu/

OnDemand team is hiring: https://go.osu.edu/ood-job https://open.xdmod.org/ https://coldfront.io

Important Info:

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

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

Decorum Doc: https://tinyurl.com/pearc-hpctoolset

Join us on Slack: https://tinyurl.com/pearc-slack

Other Places You'll Find us at PEARC20:

- **Tues, 7/28 2:35-3:50pm PST**

Open OnDemand User Group Meeting https://sched.co/cnUi

- **Wed, 7/29 1:35-3:35pm PST**

Grendel: Bare Metal Provisioning System for HPC https://sched.co/cnVj
Monitoring & Analysis of Power Consumption on HPC clusters using XDMoD https://sched.co/cnVp
Informing the on/o-prem cloud discussion in higher education https://sched.co/cnVm

- **Thurs, 7/30 8-9:35am PST**

Cloud & Data Center usage, expenditures, & approaches to ROI: a survey of academic research computing centers https://sched.co/cnWS

Ohio Supercomputer Center
An OH·TECH Consortium Member

University at Buffalo
Center for Computational Research

VIRGINIA TECH

WELCOME!

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

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

- View our “meeting decorum” document
  https://tinyurl.com/pearc-hpctoolset
- Join the Slack channel for the tutorial
  https://tinyurl.com/pearc-slack

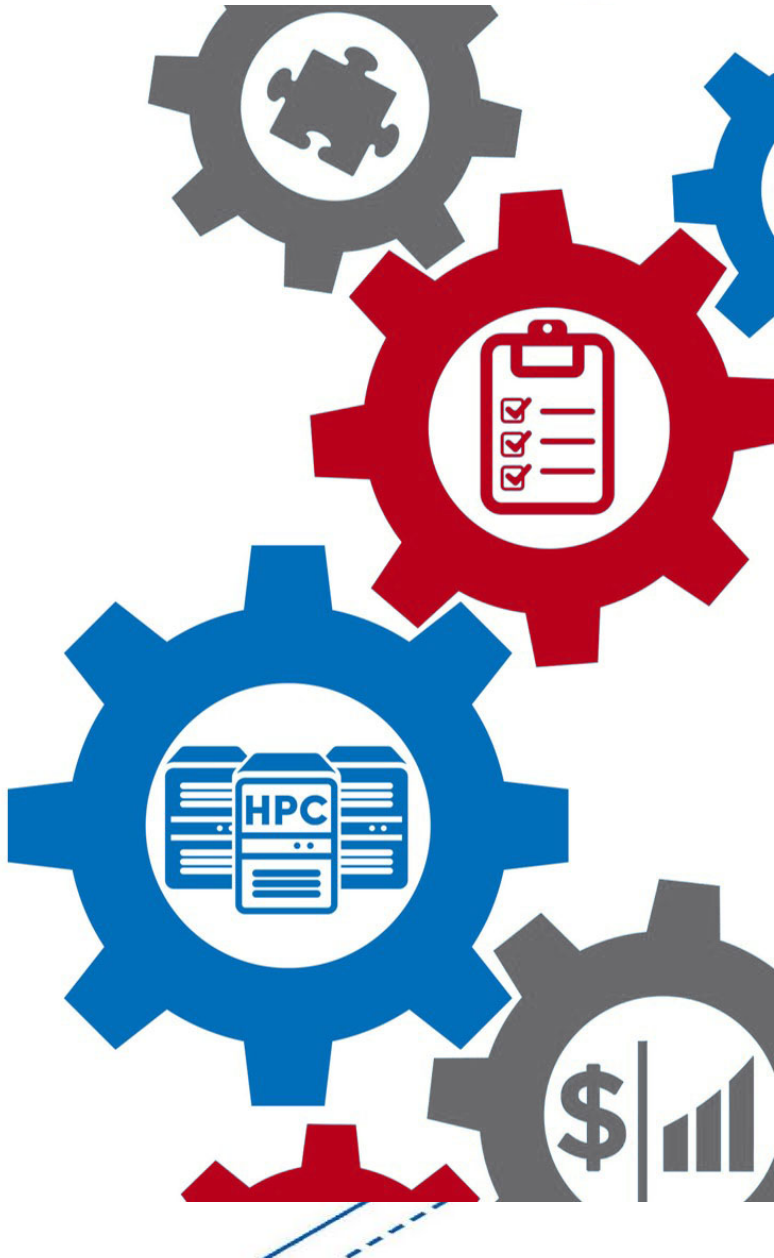

Open OnDemand, Open XDMoD, and ColdFront: An HPC center management toolset

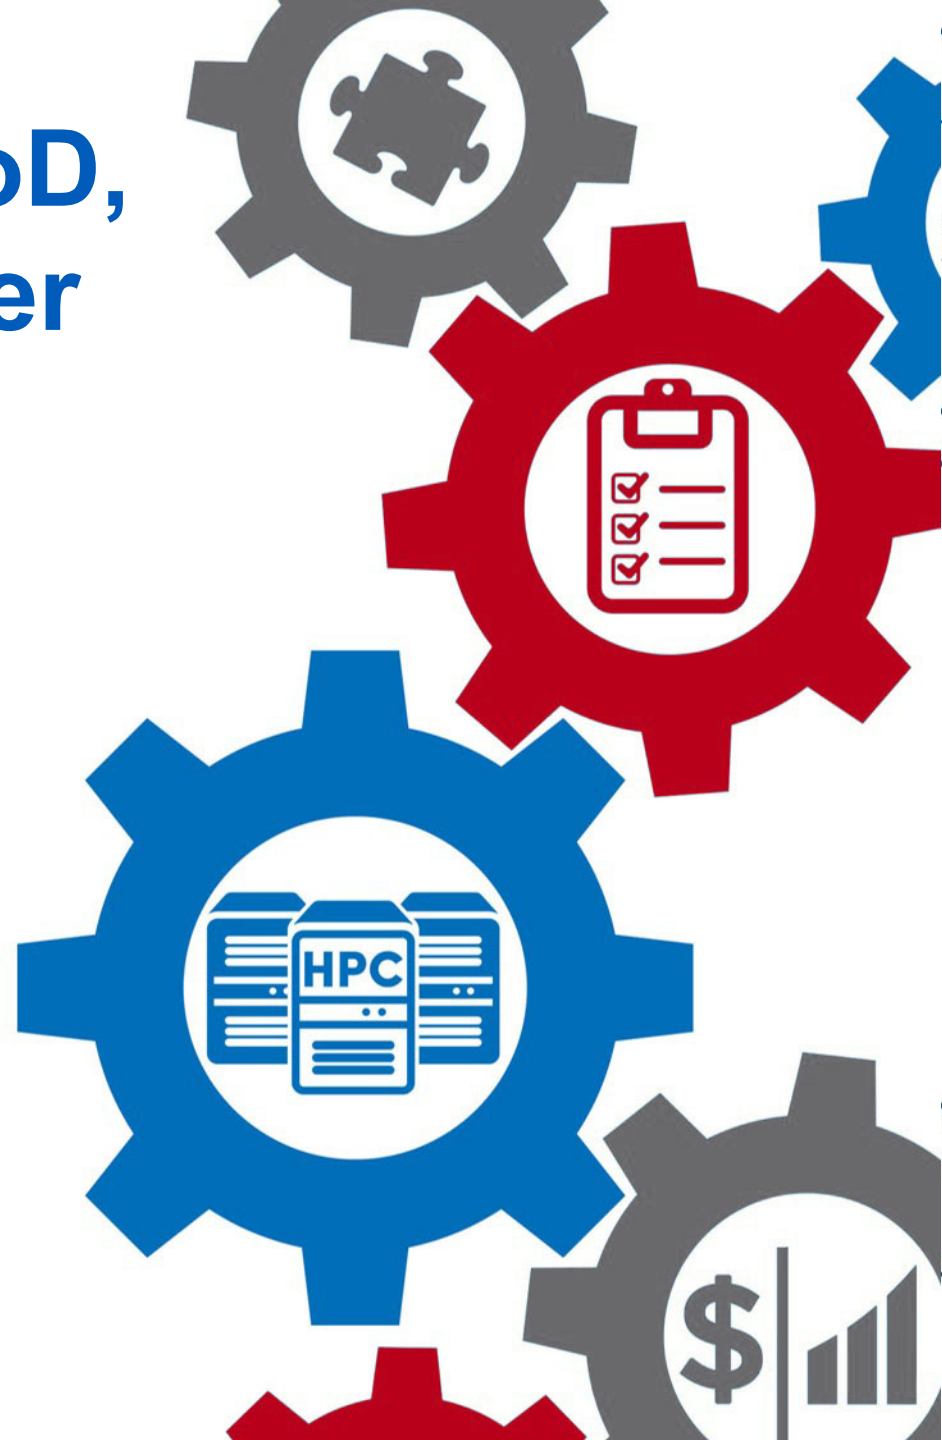

Tutorial presented at PEARC20 by staff from:

Ohio Supercomputing Center

UB Center for Computational Research

Virginia Tech Advanced Research Computing

Ohio Supercomputer Center
An OH·TECH Consortium Member

VIRGINIA TECH

University at Buffalo
Center for Computational Research

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
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

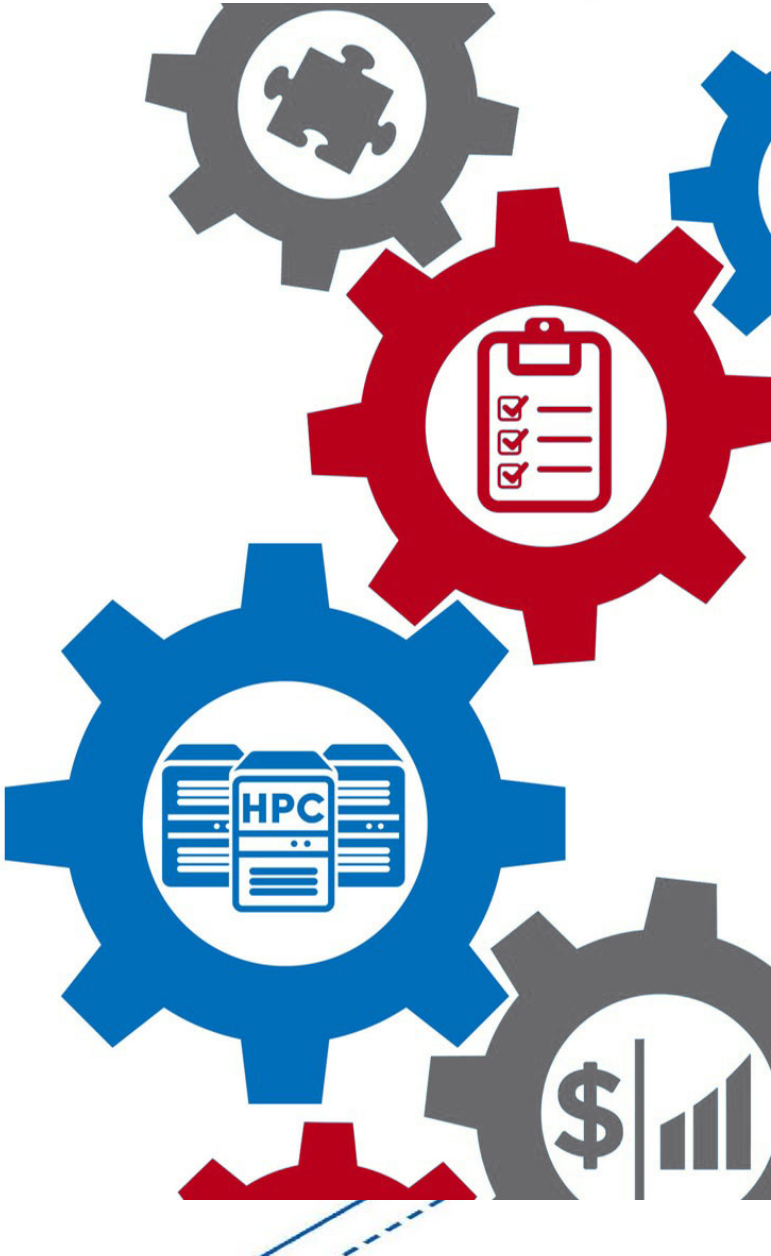

Agenda

- Getting Started
- Tutorial Goals
- Brief intro on all three products
- Tutorial technology
- Part 1: ColdFront
- Part 2: Open XDMoD
- “Lunch” Break – 30 minutes
- Part 3: Open OnDemand
- Part 4: Open OnDemand interactive app configuration
- Part 5: Configuring Open OnDemand to display Open XDMoD chart
- Post Workshop – breakout sessions & slack channel

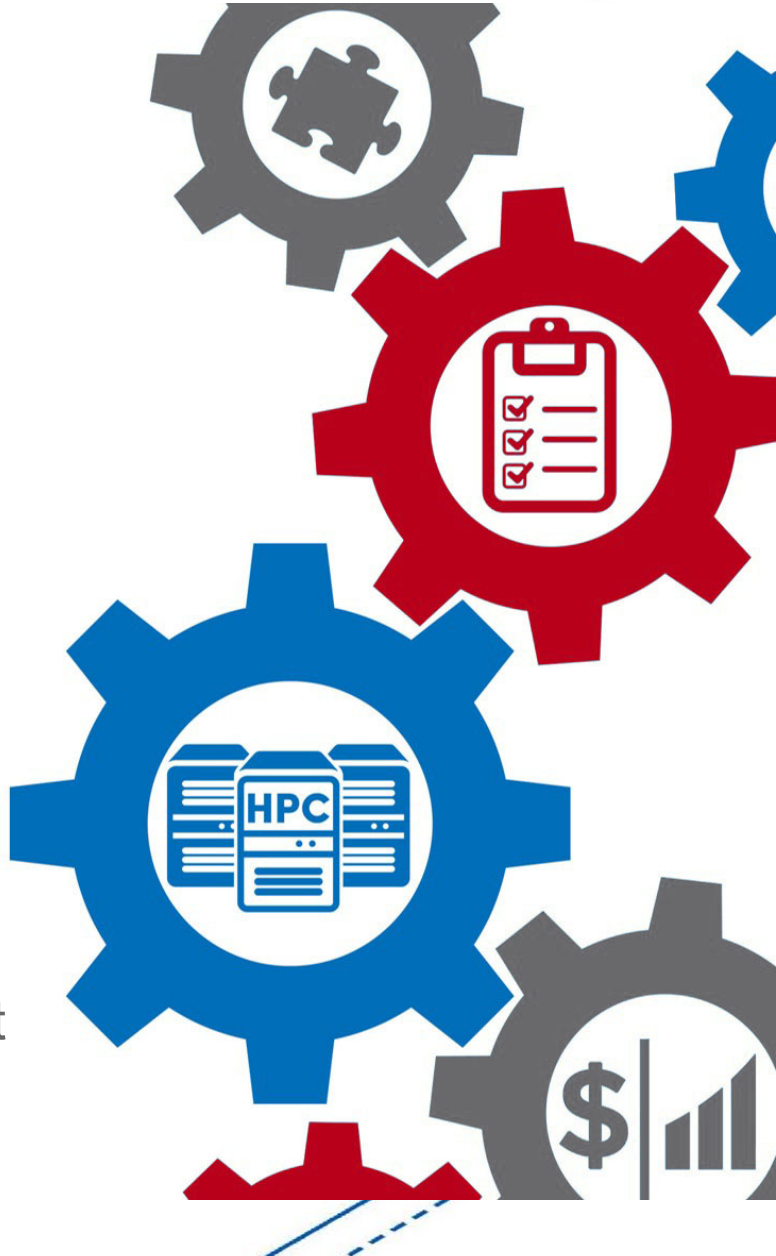

Getting Started

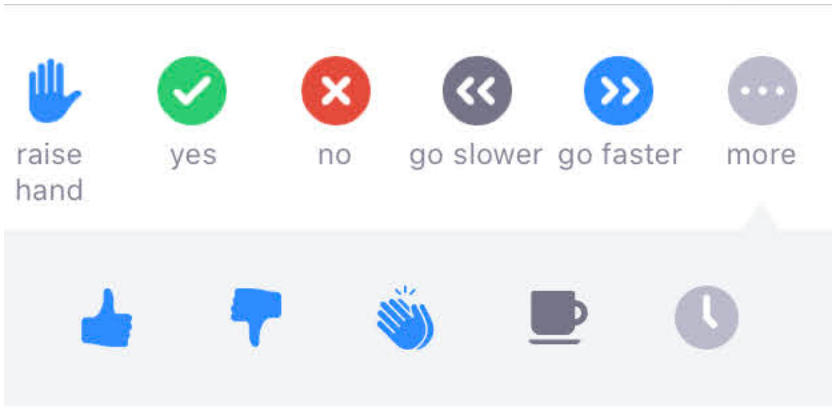

- View our “meeting decorum” document

  https://tinyurl.com/pearc-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/pearc-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 PEARC
- Breaks – sorry, no free lunch!

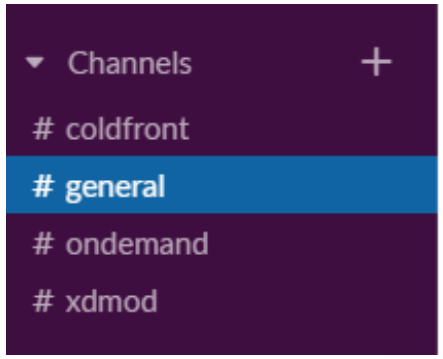

Tutorial Goals:

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

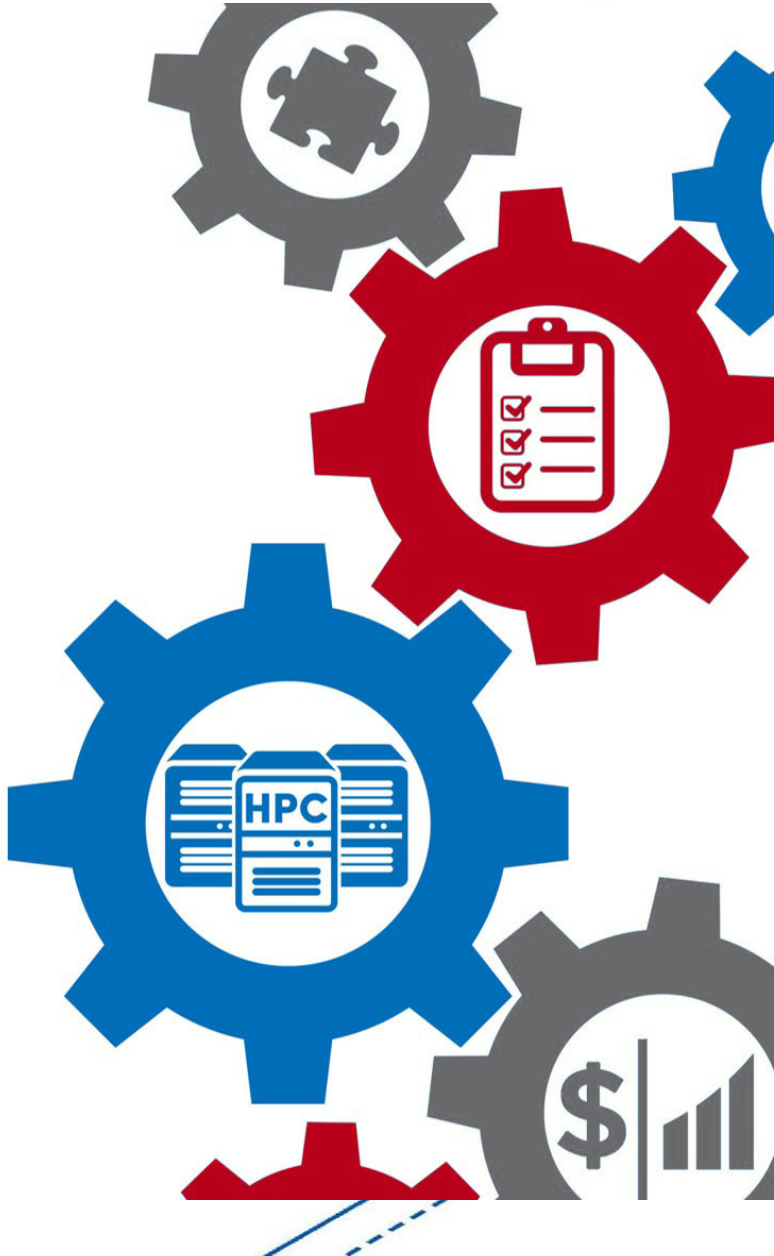

ColdFront – Allocations Management Portal

- 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

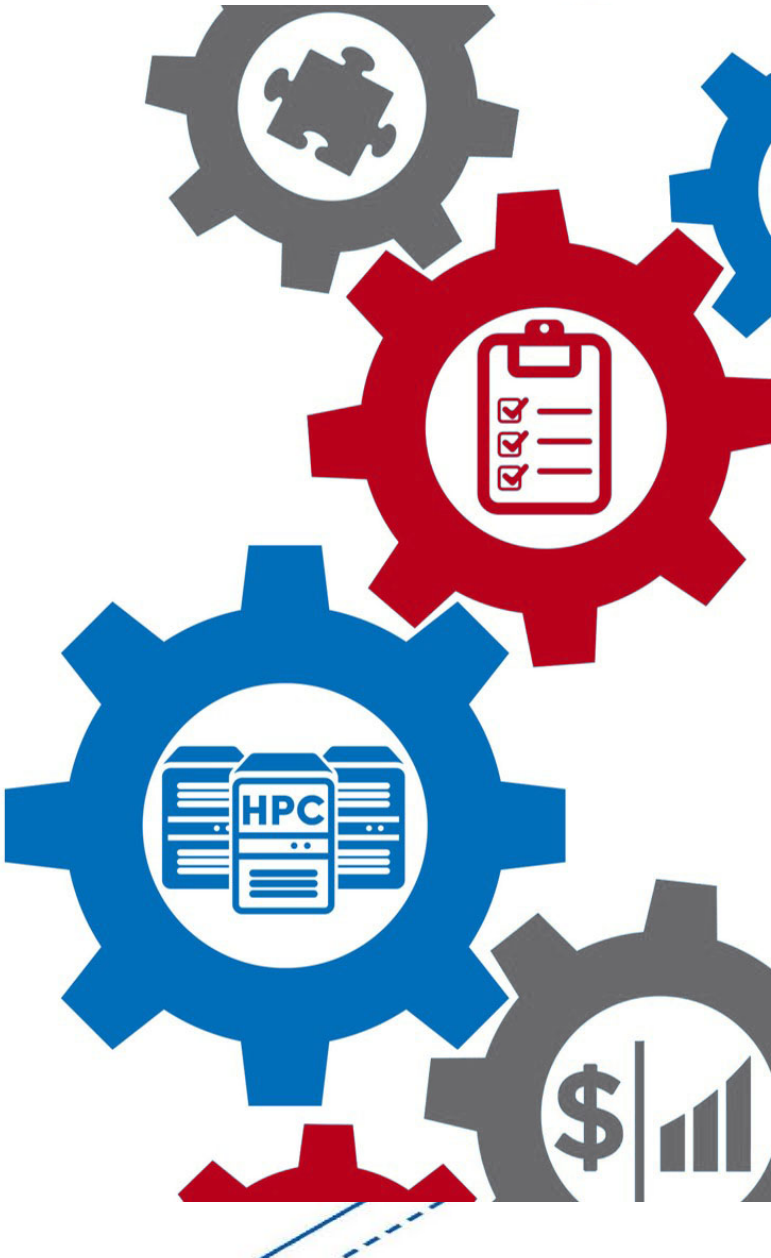

Open XDMoD

- 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

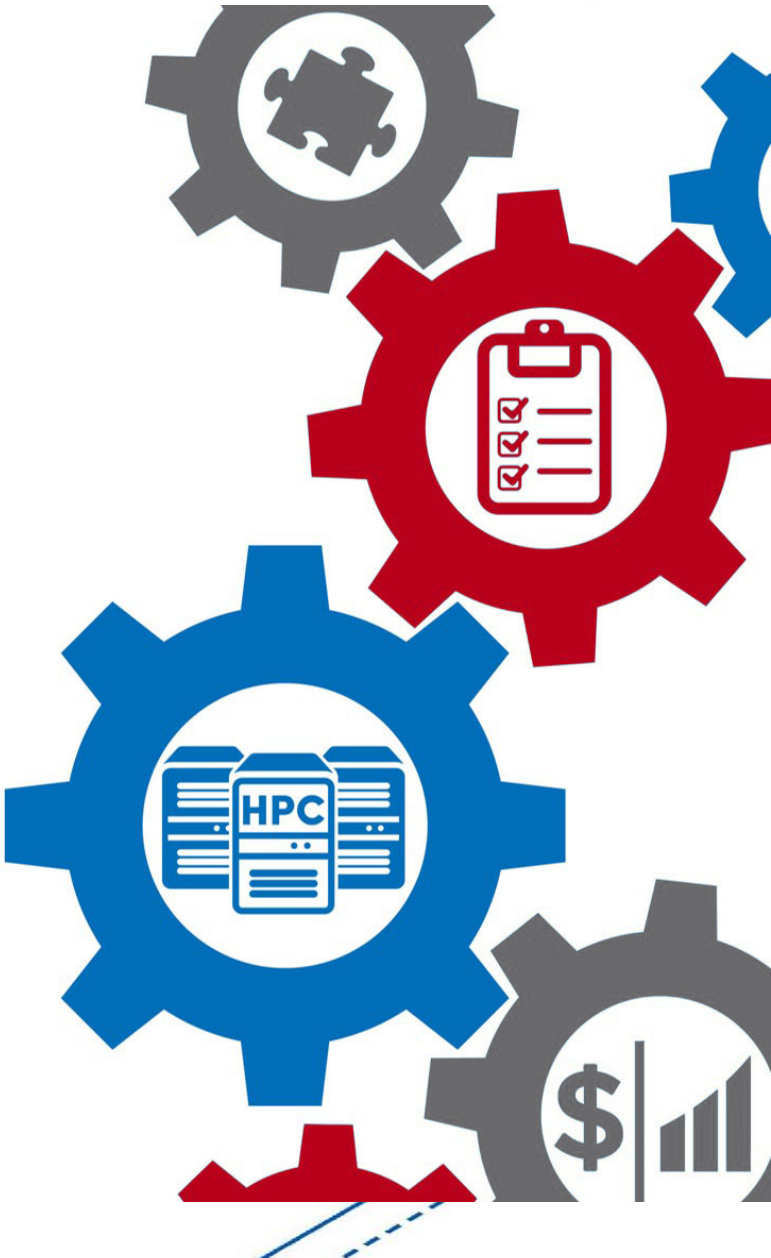

Open OnDemand

- 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

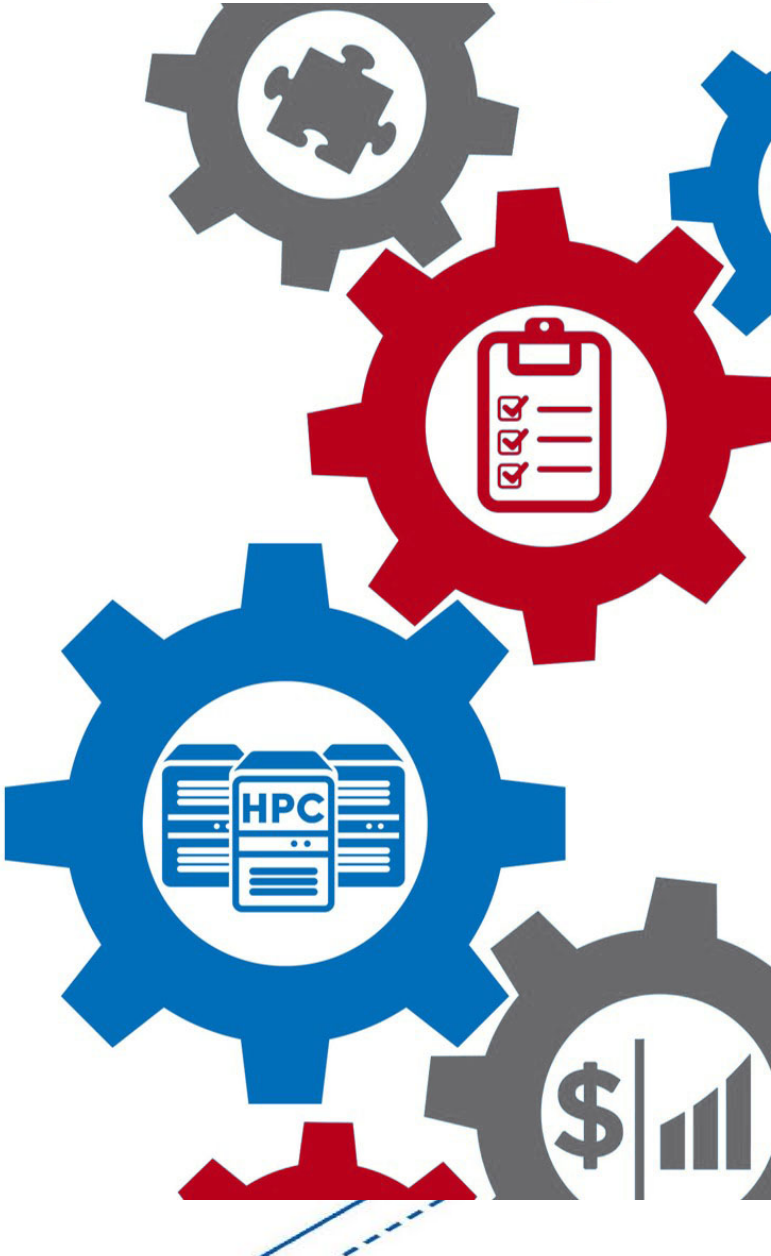

Tutorial Container Architecture

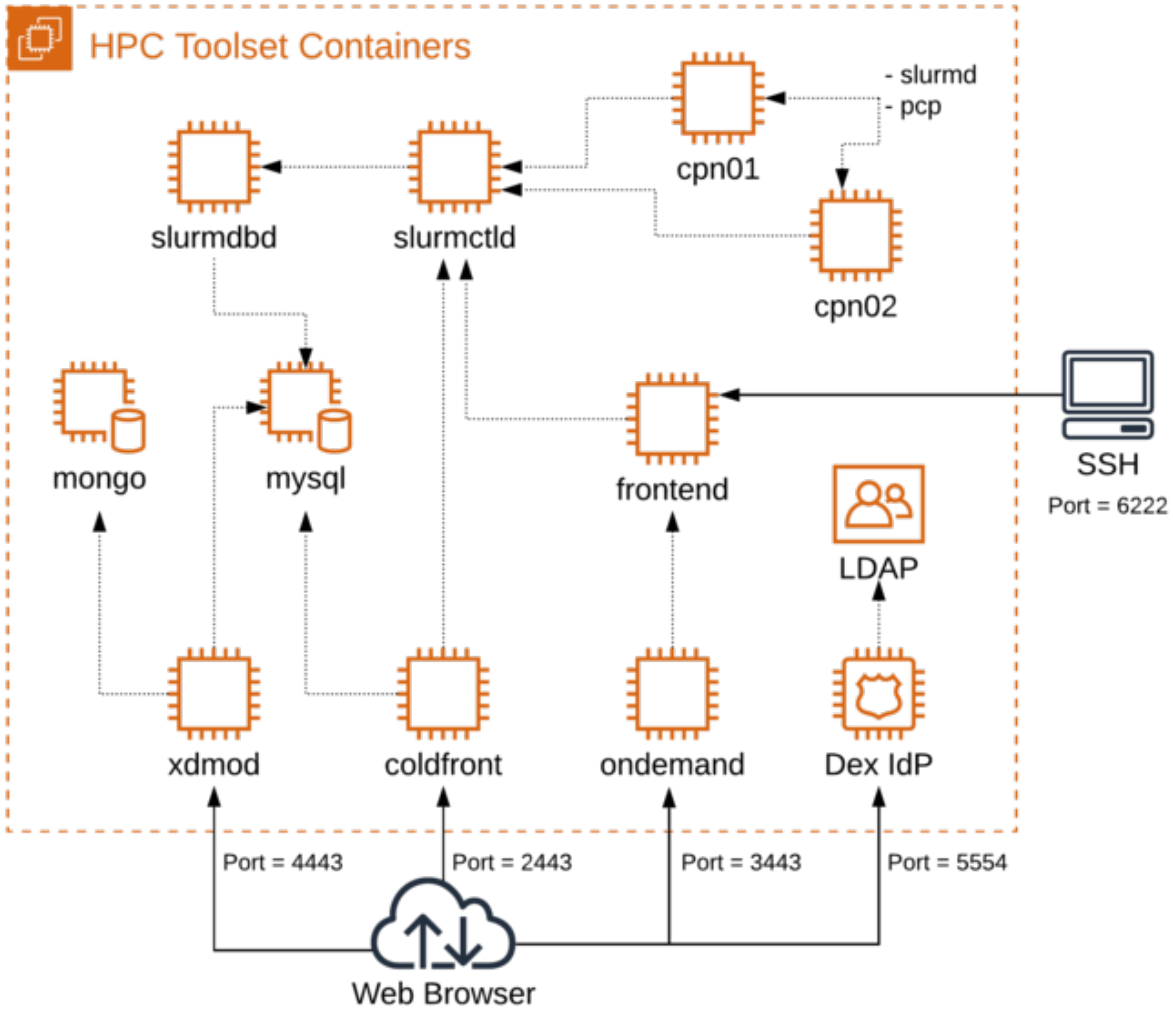

Requirements: https://github.com/ubccr/hpc-toolset-tutorial/edit/master/docs/requirements.md

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**

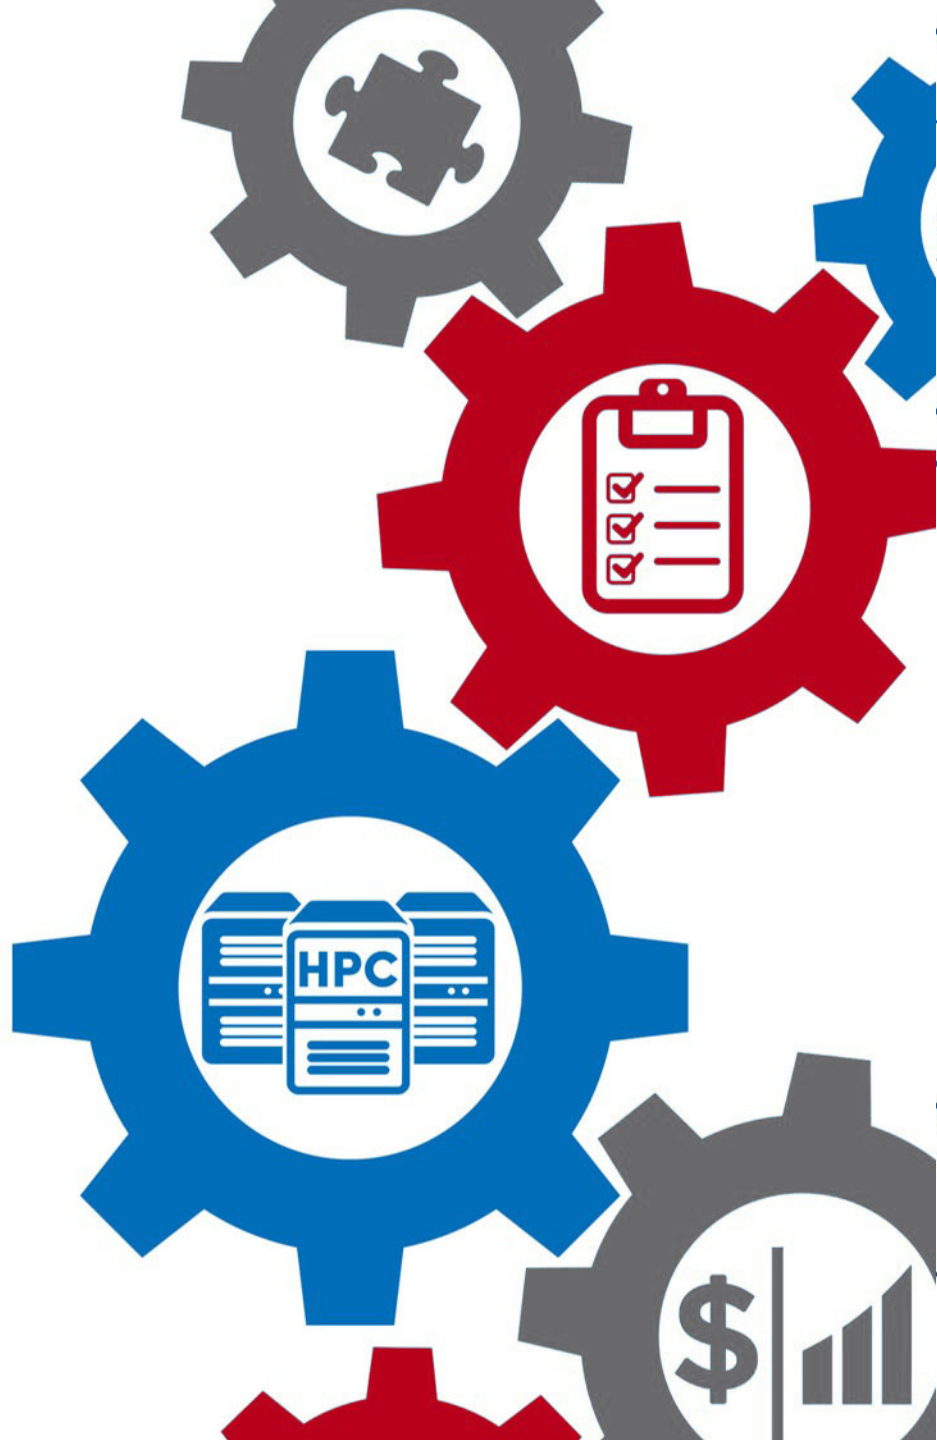

Tutorial presented at PEARC20 by:
Andrew Bruno, UB
Dori Sajdak, UB

Ohio Supercomputer Center
An OH·TECH Consortium Member

VIRGINIA TECH

University at Buffalo
Center for Computational Research

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

Users

Policies

Databases

ACLs

Cloud

System Administrators

Servers & Clusters

Software Licenses

Storage

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

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

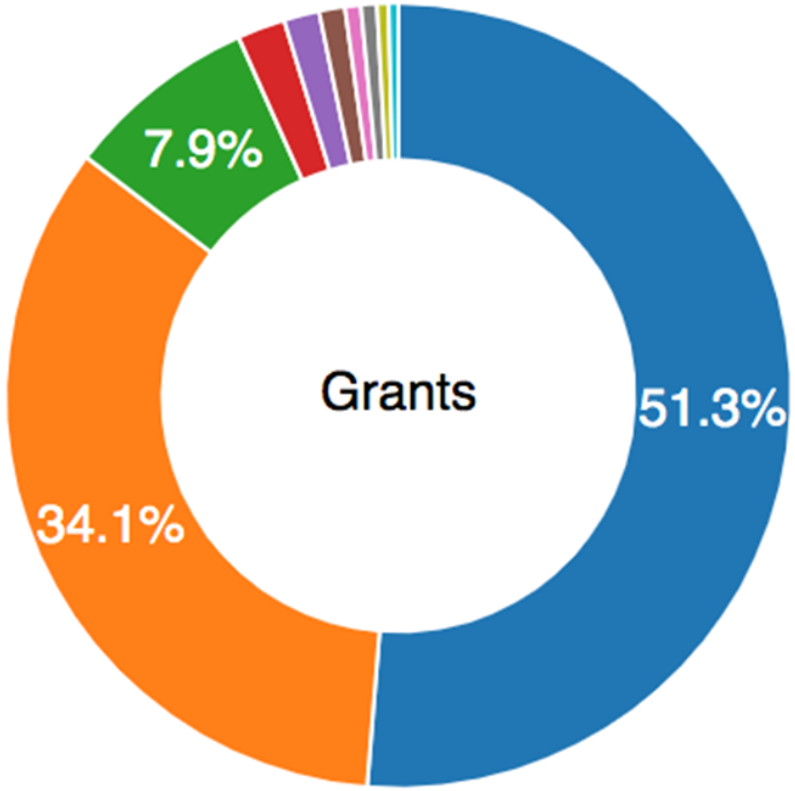

Report on resources & allocations

Collect publication information

Collect grant information

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

slurm workload manager

OpenID

FreeIPA
Open Source Identity Management Solution

XDMoD
METRICS ON DEMAND

Vendor APIs

Contact Info:

Andrew Bruno - aebruno2@buffalo.edu

Dori Sajdak - djm29@buffalo.edu

https://coldfront.io

Subscribe to the ColdFront mailing list:

Send an email to 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

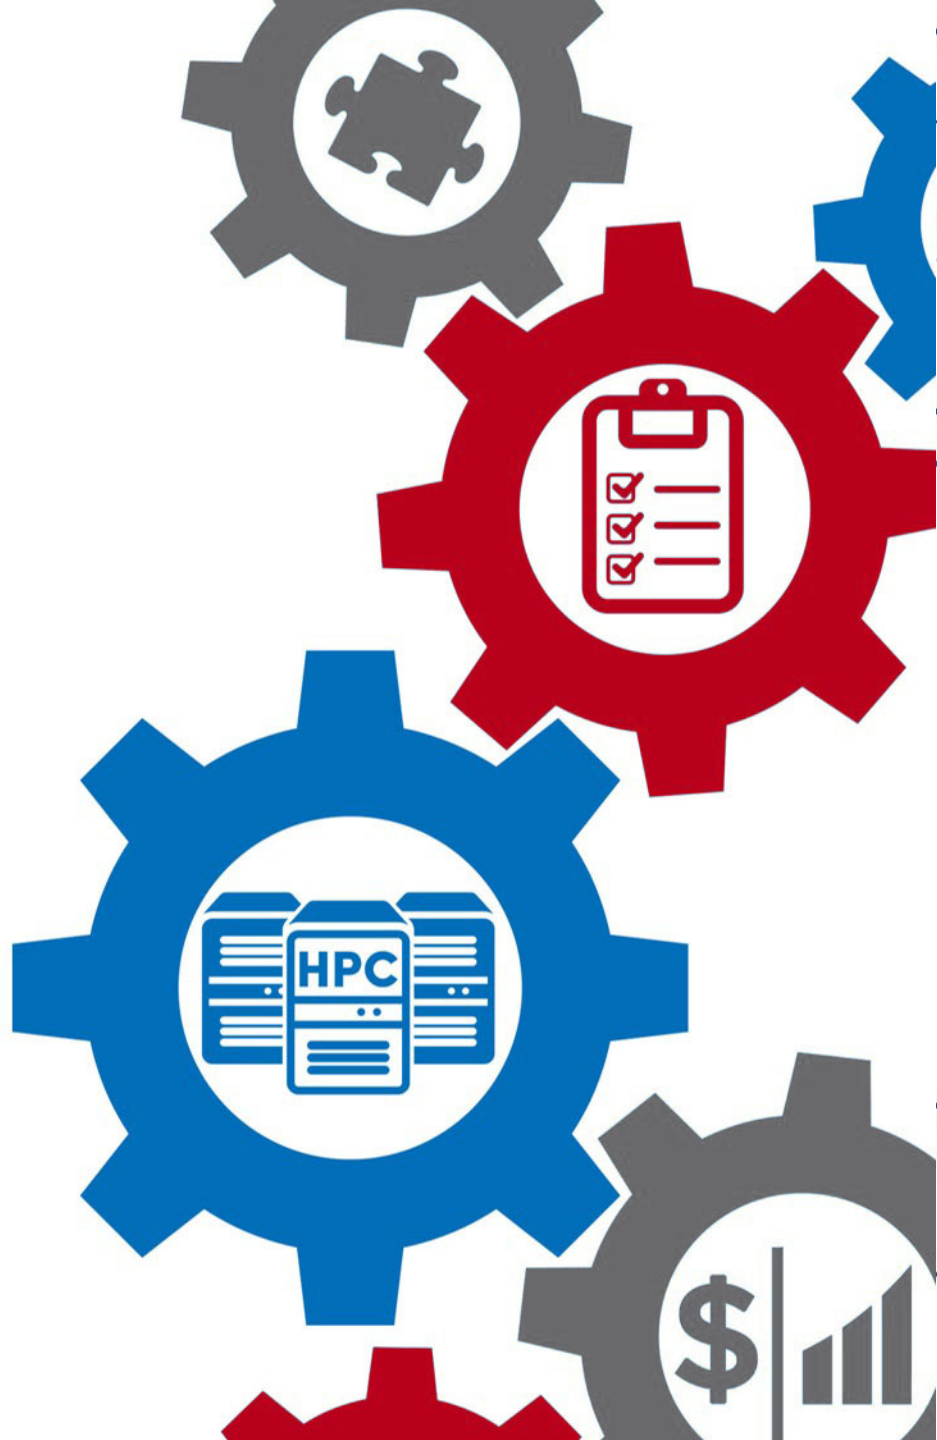

Ohio Supercomputer Center
An OH·TECH Consortium Member

VIRGINIA TECH

University at Buffalo
Center for Computational Research

Ohio Supercomputer Center
An OH·TECH Consortium Member

University at Buffalo
Center for Computational Research

VIRGINIA TECH

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

https://buffalo.edu/ccr
https://open.xdmod.org/

https://arc.vt.edu/
https://coldfront.io

Important Info:

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

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

Decorum Doc: https://tinyurl.com/pearc-hpctoolset

Join us on Slack: https://tinyurl.com/pearc-slack

Other Places You'll Find us at PEARC20:

- **Tues, 7/28 2:35-3:50pm PST**

Open OnDemand User Group Meeting https://sched.co/cnUi

- **Wed, 7/29 1:35-3:35pm PST**

Grendel: Bare Metal Provisioning System for HPC https://sched.co/cnVj
Monitoring & Analysis of Power Consumption on HPC clusters using XDMoD https://sched.co/cnVp
Informing the on/o-prem cloud discussion in higher education https://sched.co/cnVm

- **Thurs, 7/30 8-9:35am PST**

Cloud & Data Center usage, expenditures, & approaches to ROI: a survey of academic research computing centers https://sched.co/cnWS

Open XDMoD

Tutorial presented at PEARC20 by:
Benjamin D. Plessinger, UB

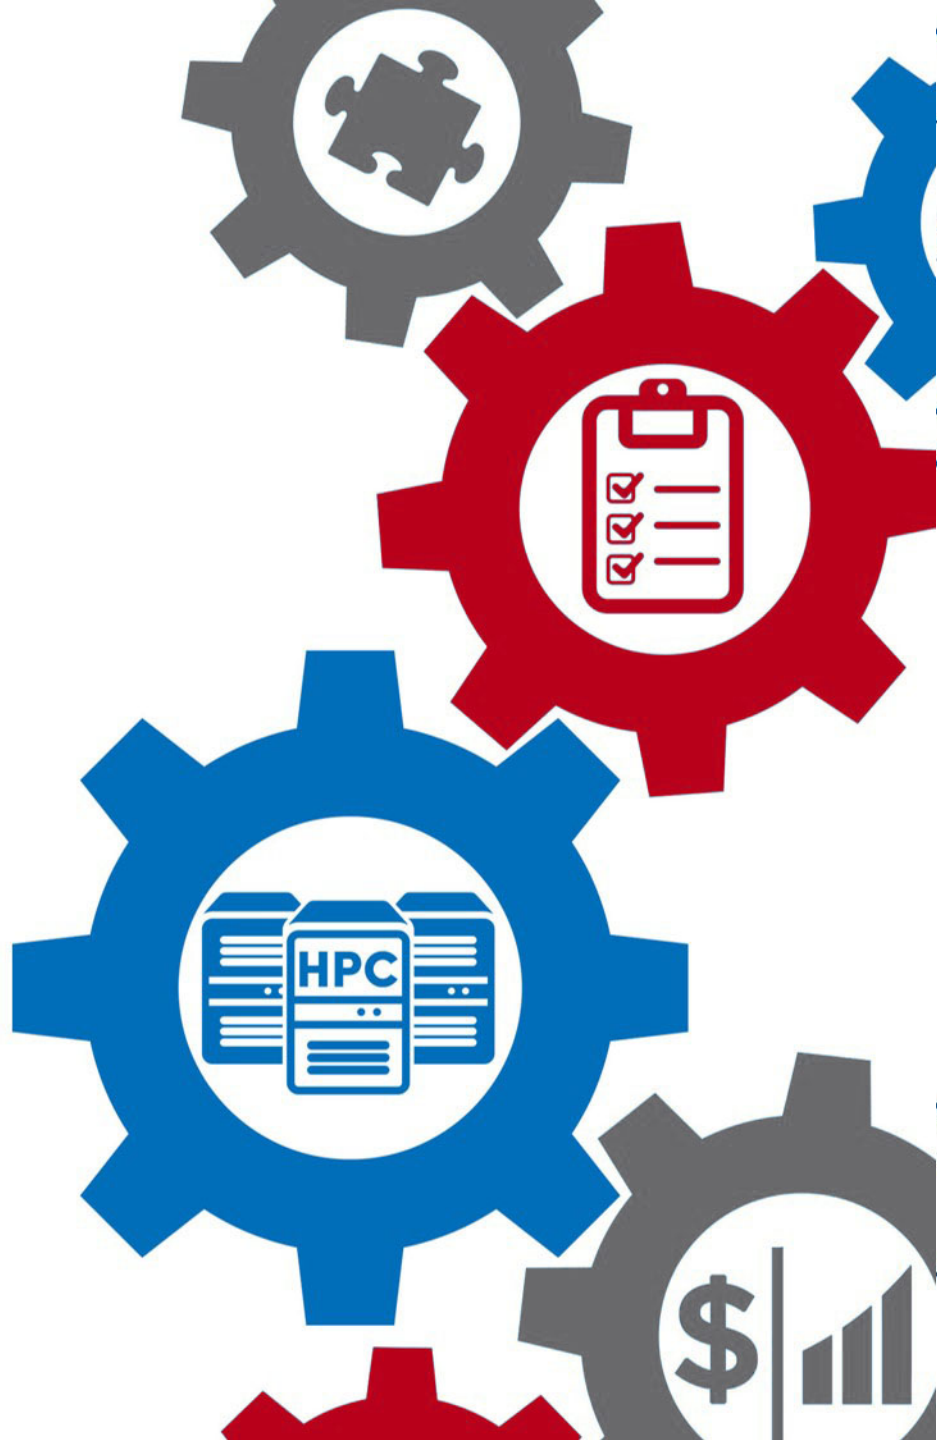

Ohio Supercomputer Center
An OH·TECH Consortium Member

VIRGINIA TECH

University at Buffalo
Center for Computational Research

Open XDMoD Tutorial Prerequisites

- Use the running containers that you setup in the ColdFront tutorial
- If just joining us:

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

Open XDMoD

- Comprehensive Framework for HPC Management
  - Utilization metrics
  - Measure QoS of HPC Infrastructure (App Kernels)
  - Job-level performance data

- 200 known installations worldwide
  - NIH, EPA, CDC, NASA, NCAR, and NYSDOH
  - Dow Chemical, Rolls Royce, Biogen, Radix Trading, Fujitsu Systems Europe, Regeneron Pharmaceuticals, ….
  - Germany, Spain, Belgium, England, Canada, Brazil, Italy, France, Netherlands, Australia, Saudi Arabia, China, Israel, Pakistan, Turkey, India, …..

XMS Team

- UB: Matt Jones, Bob DeLeon, Ryan Rathsam, Jeff Palmer, Joseph White, Jeanette Sperhac, Abani Patra, Nikolay Simakov, Cynthia Cornelius, Ben Plessinger, Greg Dean
- Roswell Park: Tom Furlani, Steve Gallo
- Indiana: Gregor von Laszewski, Fugang Wang
- TACC: Bill Barth, Todd Evans

Tutorial Container Architecture

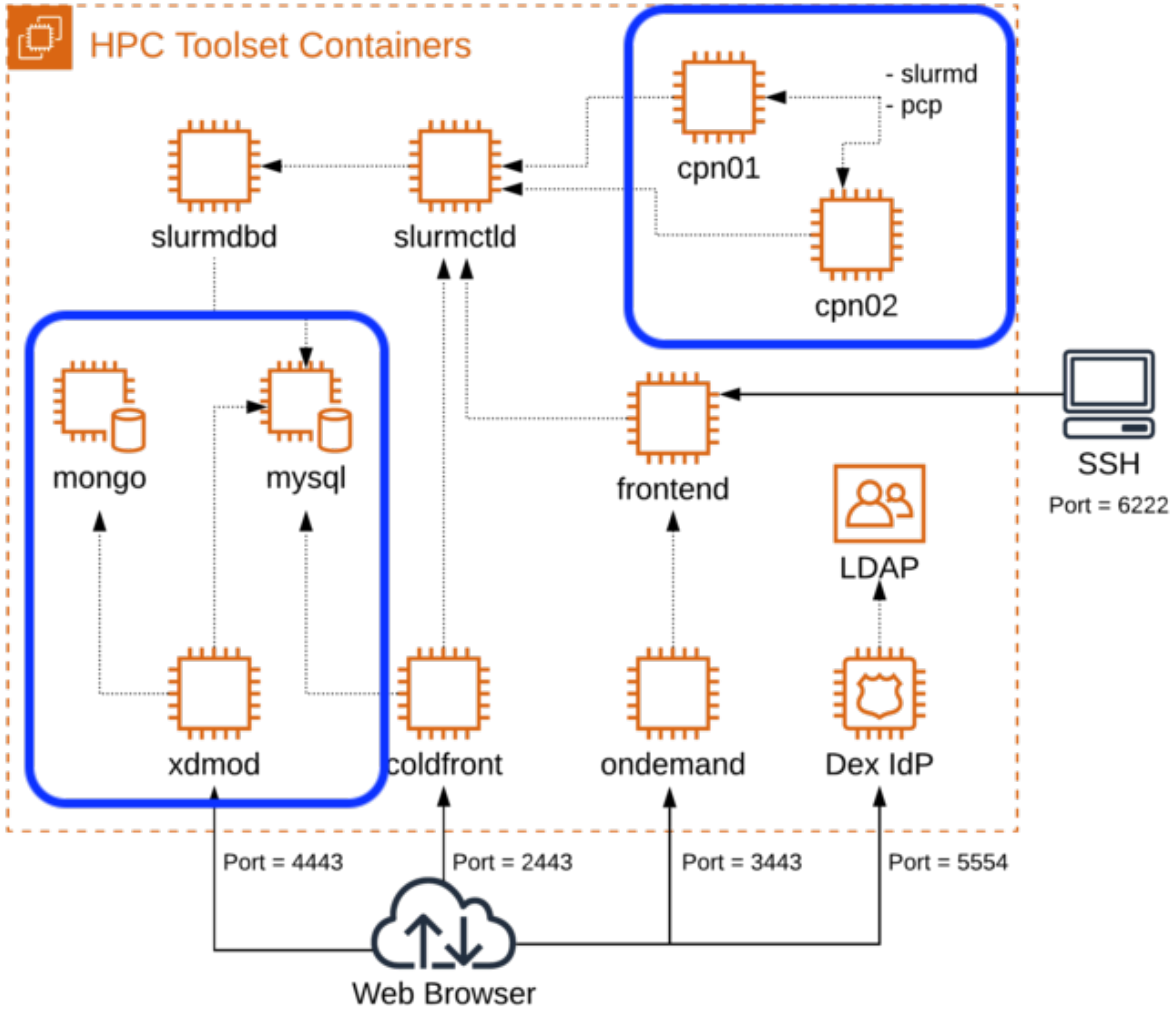

Requirements: https://github.com/plessbd/hpc-toolset-tutorial/tree/master/xdmod

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 Overview

- Using release candidate of Open XDMoD 9.0
  - Full release should be available to download next week
- Includes
  - Open XDMoD (core)
  - Open XDMoD Job Performance module
- Excludes
  - Open XDMoD Application Kernels module
  - Other data realms (Storage, Cloud)

OPEN OnDemand

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!

Full details are available here:
go.osu.edu/ood-job

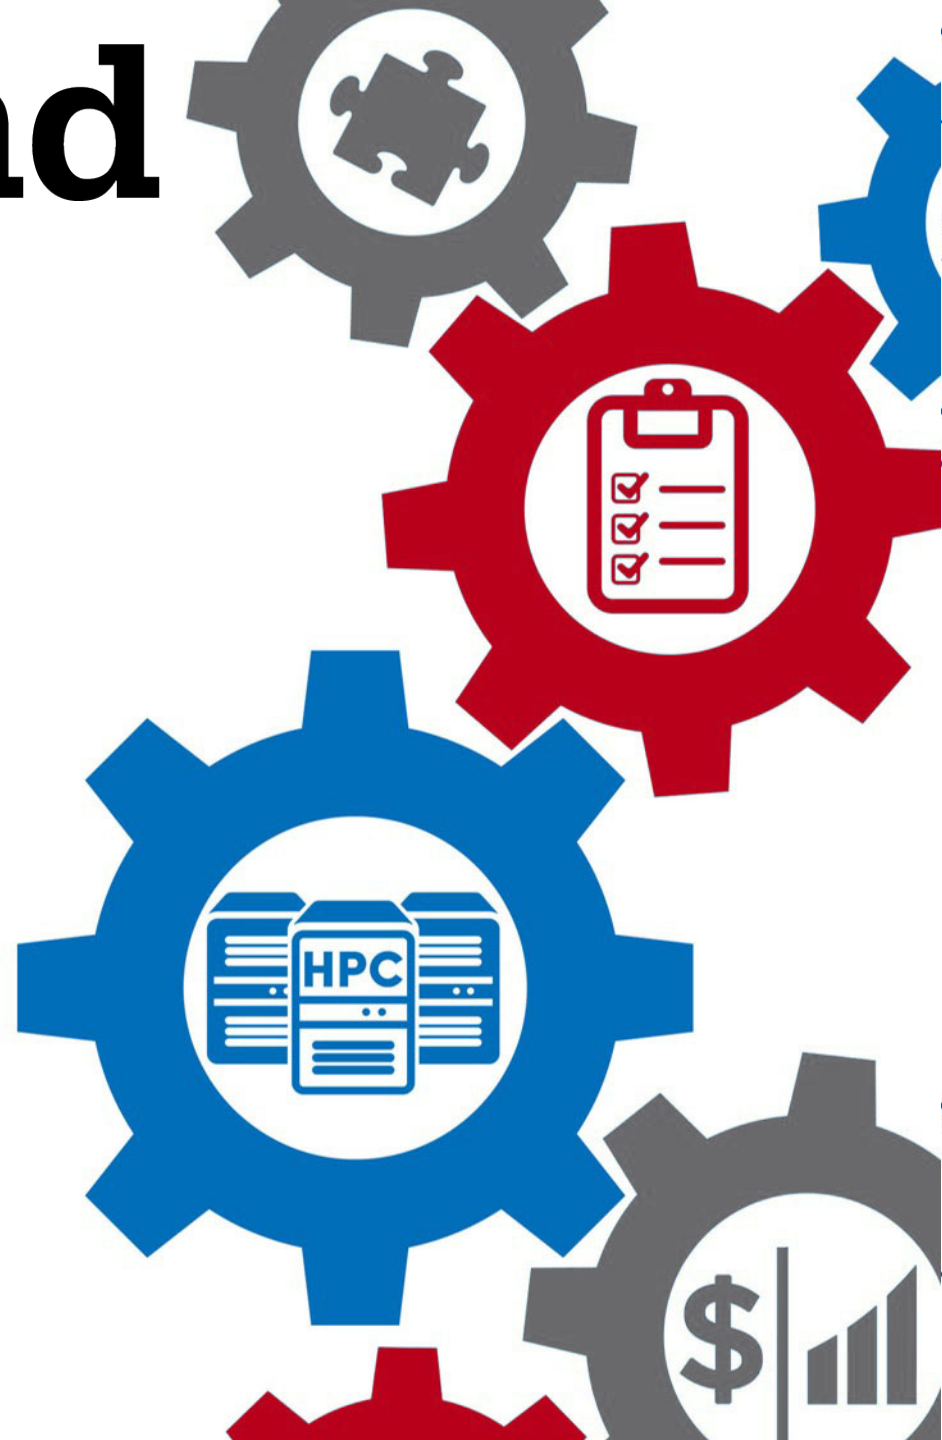

Ohio Supercomputer Center
An OH·TECH Consortium Member

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

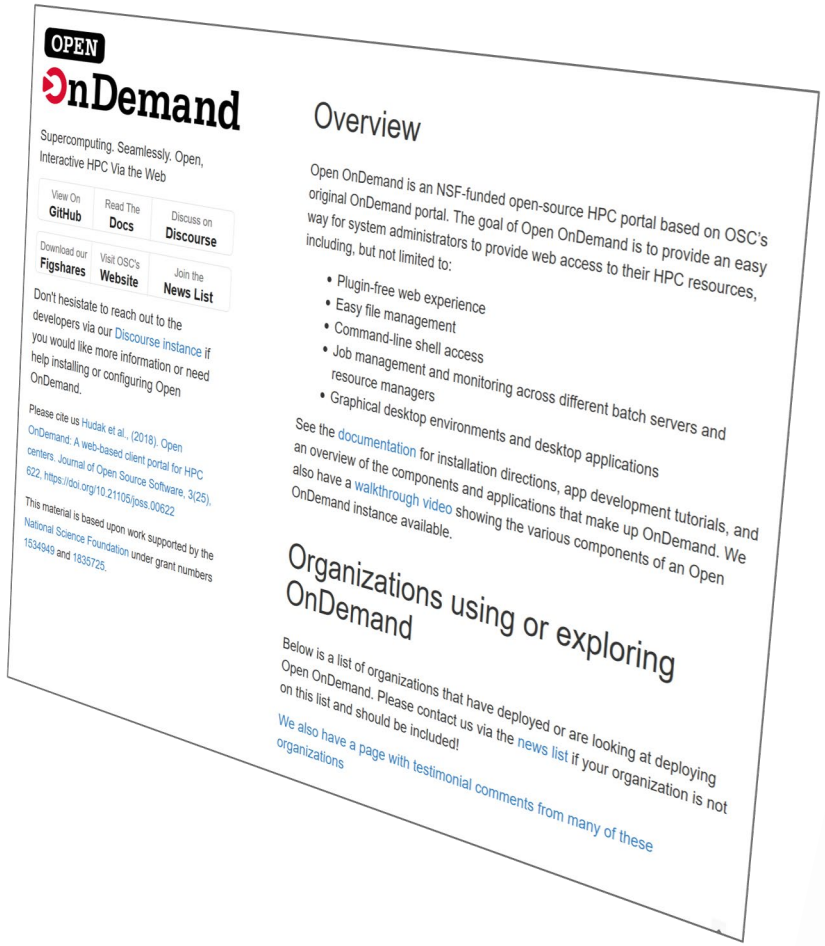

This work is supported by the National Science Foundation of the United States under the awards NSF SI2-SSE-1534949 and CSSI-Software-Frameworks-1835725.

Approx Number of Institutions based on RPM logs

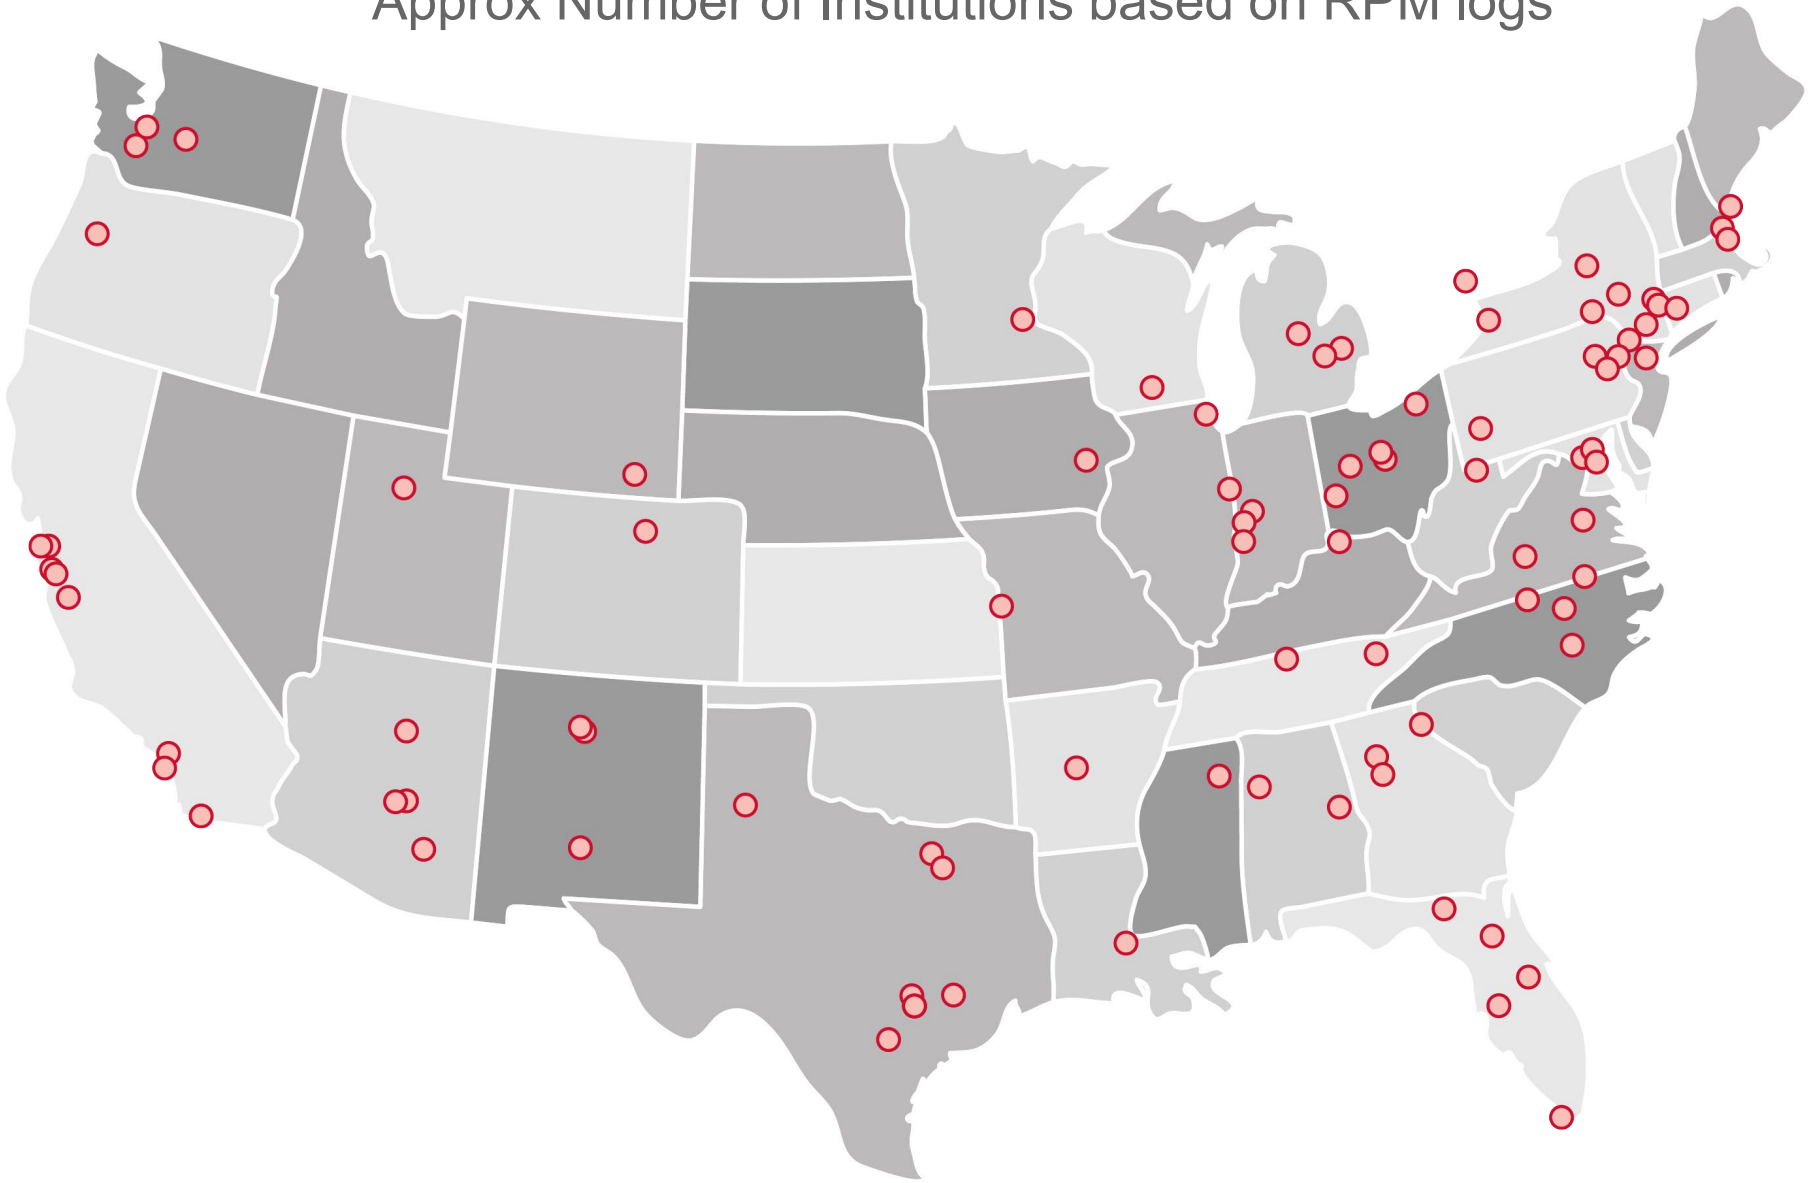

- 136 unique US locations
- 70 unique international locations

Ohio Supercomputer Center
An OH·TECH Consortium Member

University at Buffalo
Center for Computational Research

VIRGINIA TECH

Example Current Engagements and Deployments

Production Deployments

ROSWELL PARK COMPREHENSIVE CANCER CENTER
THE UNIVERSITY OF ARIZONA
VANDERBILT UNIVERSITY
TEXAS A&M UNIVERSITY
Stanford
Ohio Supercomputer Center
SHANGHAI JIAO TONG UNIVERSITY
CASE WESTERN RESERVE UNIVERSITY EST. 1826
PITTSBURGH SUPERCOMPUTING CENTER
ILLINOIS NCSA | National Center for Supercomputing Applications
JOHNS HOPKINS UNIVERSITY
university of groningen
RUTGERS THE STATE UNIVERSITY OF NEW JERSEY
PURDUE UNIVERSITY
THE UNIVERSITY OF ALABAMA AT BIRMINGHAM
THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL
TCNJ THE COLLEGE OF NEW JERSEY
HARVARD UNIVERSITY
University at Buffalo
NOR-TECH People Friendly Technology
UNIVERSITY OF MICHIGAN
THE UNIVERSITY OF UTAH
Idaho National Laboratory
WAGENINGEN UNIVERSITY & RESEARCH
PennState
TANDY SUPERCOMPUTING CENTER
The University of Vermont
University of Houston Clear Lake
GHENT UNIVERSITY
PRINCETON UNIVERSITY
SMU
Caltech
University of Pittsburgh
VIRGINIA TECH
WRIGHT STATE UNIVERSITY
BOSTON UNIVERSITY
SwRI SOUTHWEST RESEARCH INSTITUTE
Yale
Tufts UNIVERSITY
INSTITUUT LORENTZ
NVIDIA

In Process of Installing

Cincinnati Children's
UNIVERSITY OF MARYLAND
UNIVERSITY of VIRGINIA
University of Kentucky
NAVAL NUCLEAR LABORATORY
MGHPCC MASSACHUSETTS GREEN HIGH PERFORMANCE COMPUTING CENTER
Oregon State University
UNIVERSITY of WASHINGTON
TOTALSIM US
UNIVERSITY OF OREGON
THE OHIO STATE UNIVERSITY
University of Missouri
UMassAmherst
UNIVERSITY OF HULL
UTSA The University of Texas at San Antonio
جامعة الملك عبدالله للعلوم والتقنية King Abdullah University of Science and Technology
University of CINCINNATI
NORTHERN ARIZONA UNIVERSITY
LSU
OAK RIDGE National Laboratory
FEDERAL RESERVE BANK OF SAN FRANCISCO
X-ISS

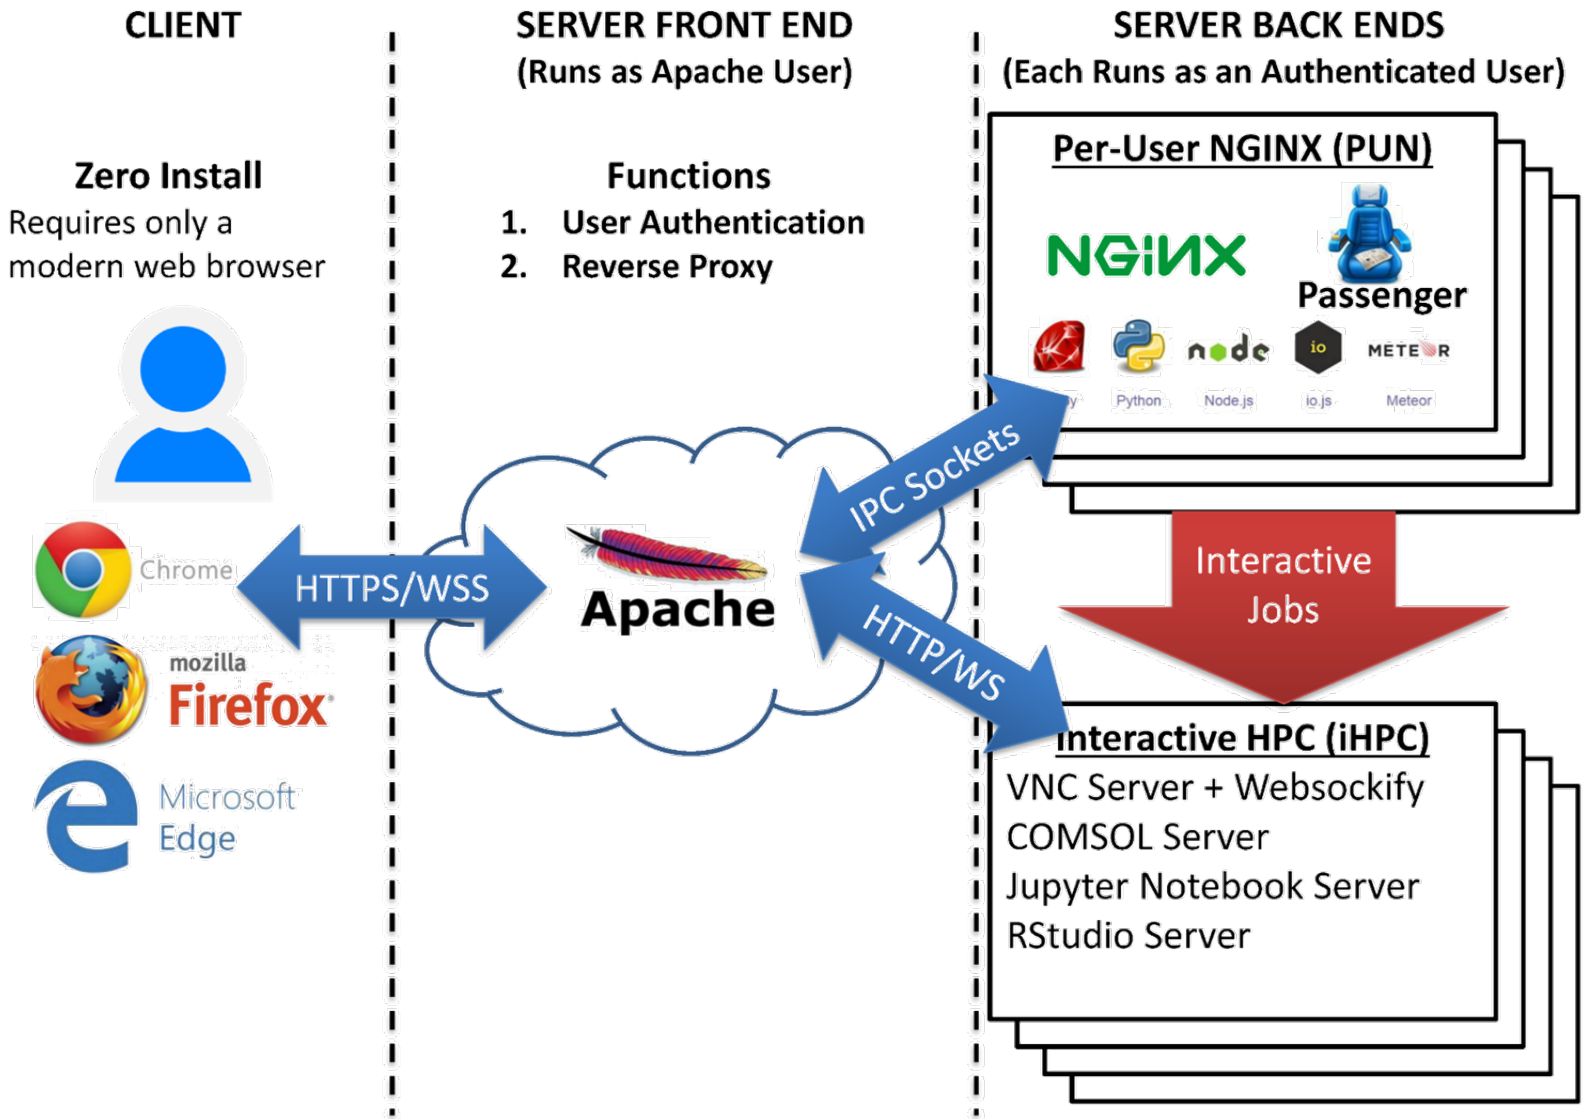
CLIENT
SERVER FRONT END
(Runs as Apache User)
SERVER BACK ENDS
(Each Runs as an Authenticated User)
Zero Install
Requires only a
modern web browser
Functions
1. User Authentication
2. Reverse Proxy
Per-User NGINX (PUN)
NGINX
Passenger
Python
Node.js
io.js
Meteor
IPC Sockets
Chrome
HTTPS/WSS
Apache
HTTP/WS
Interactive
Jobs
mozilla
Firefox
Microsoft
Edge
Interactive HPC (iHPC)
VNC Server + Websockify
COMSOL Server
Jupyter Notebook Server
RStudio Server

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 (November 9)	Kubernetes adapter SSL+auth abstraction for apps	New launch interface UX Files app replacement Completed Jobs App
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

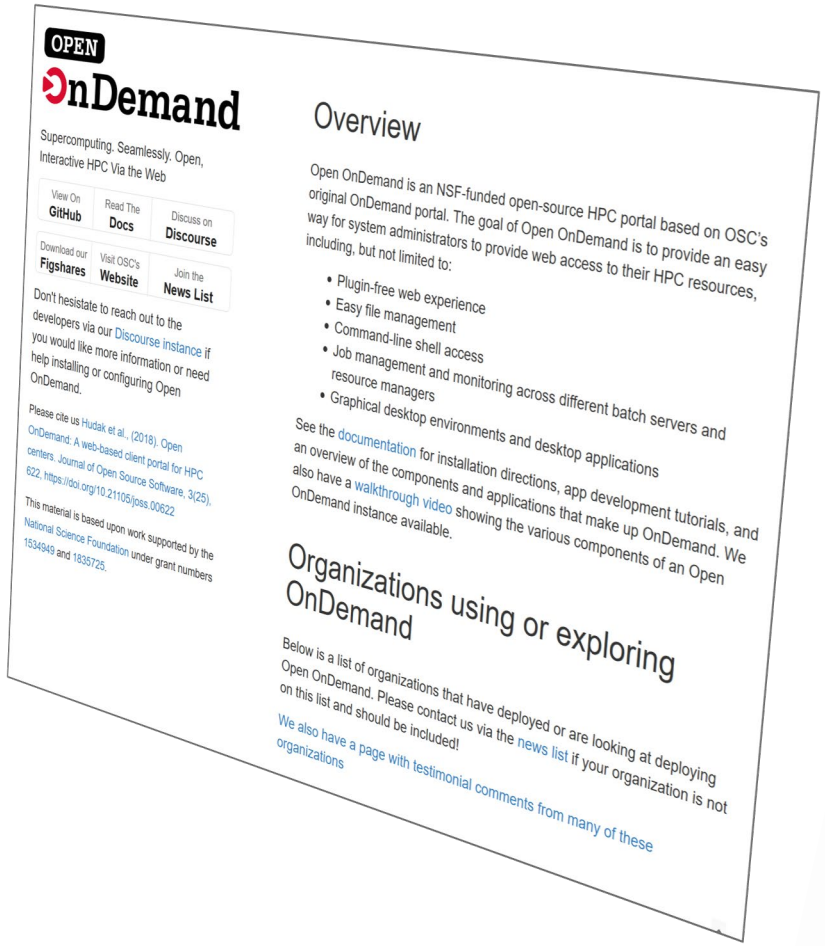

This work is supported by the National Science Foundation of the United States under the awards NSF SI2-SSE-1534949 and CSSI-Software-Frameworks-1835725.

Ohio Supercomputer Center
An OH·TECH Consortium Member

University at Buffalo
Center for Computational Research

VIRGINIA TECH

Usage Overview Demo

Bob Settlage, VT

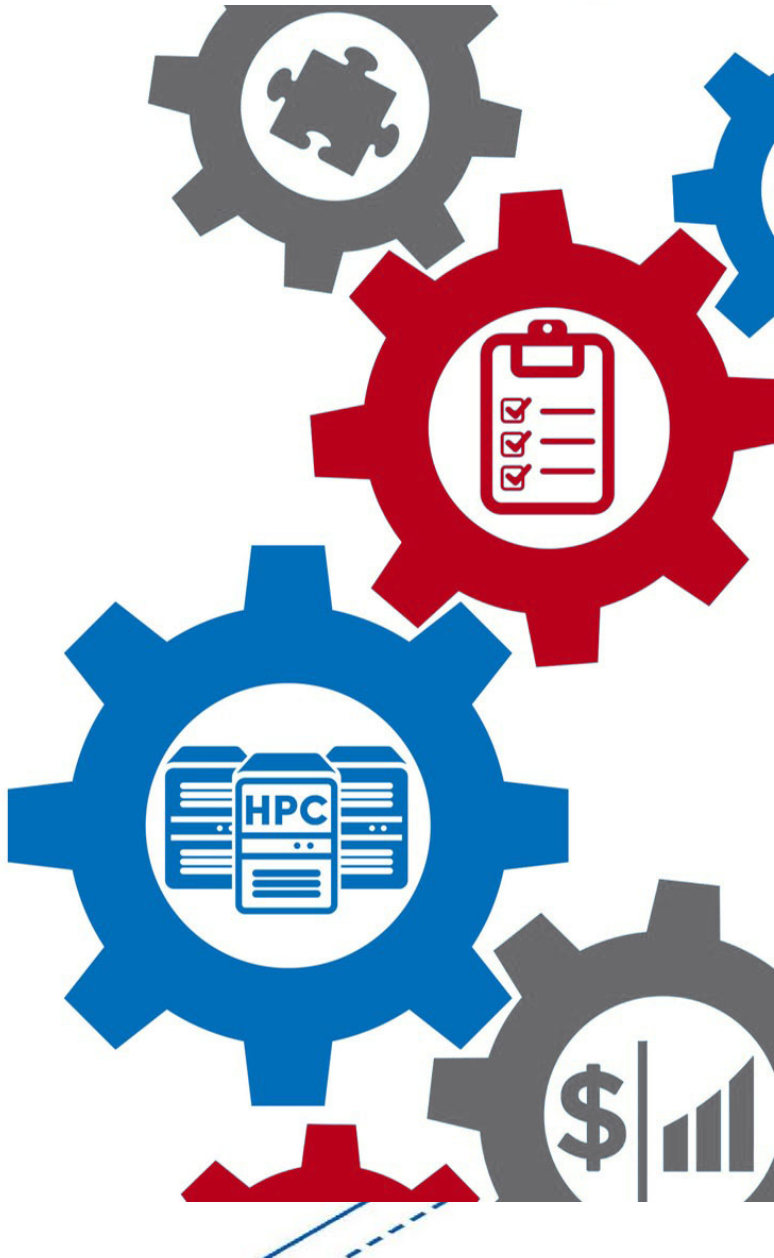

Ohio Supercomputer Center
An OH·TECH Consortium Member

University at Buffalo
Center for Computational Research

VIRGINIA TECH

Installing OnDemand

Eric Franz, OSC

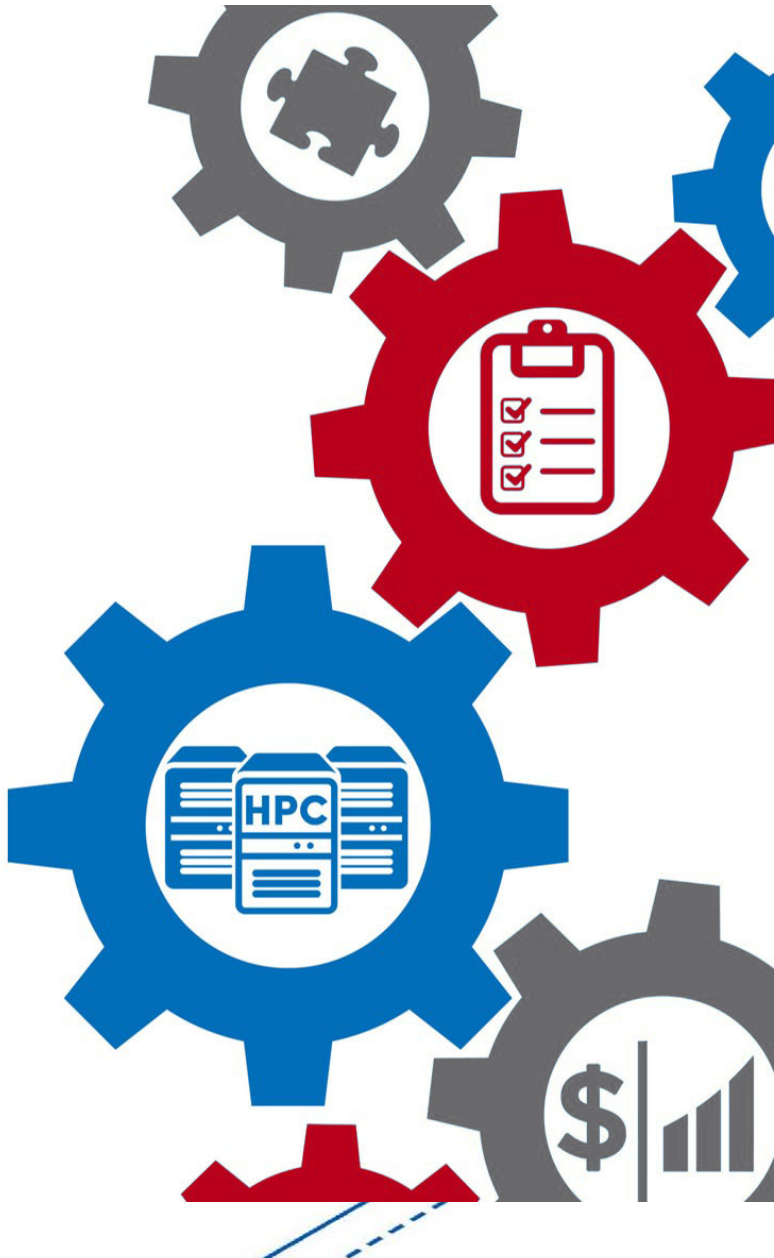

Installation: requirements

- **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

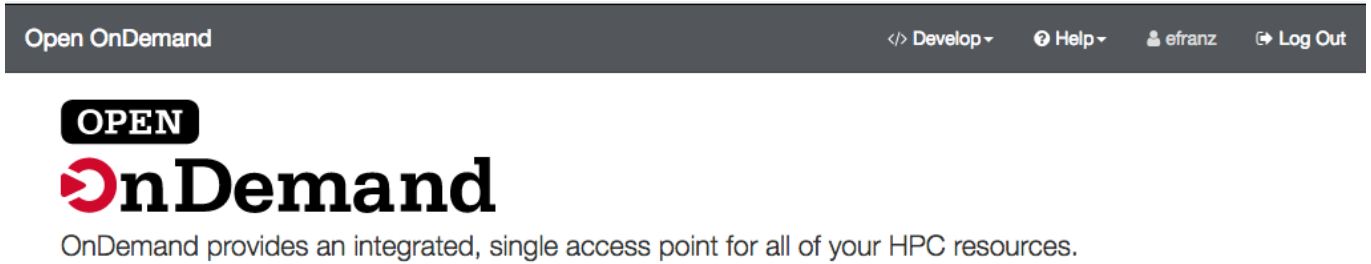

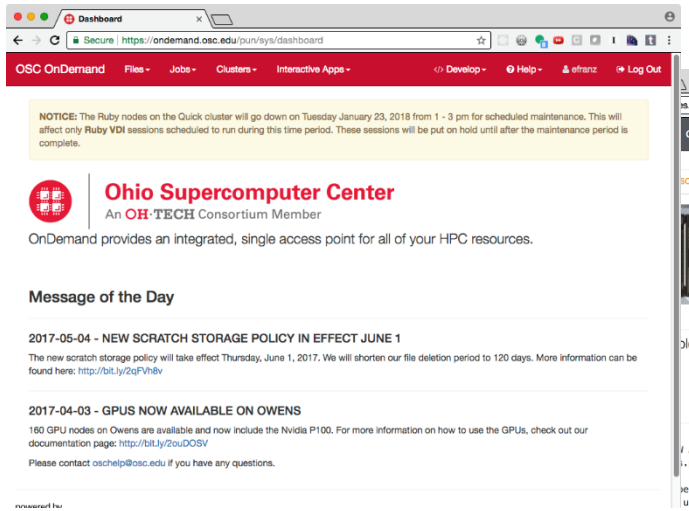

Installation: Files shortcuts

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

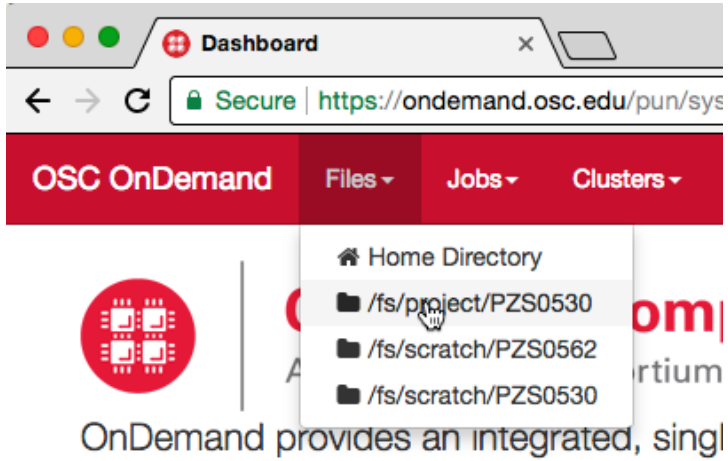

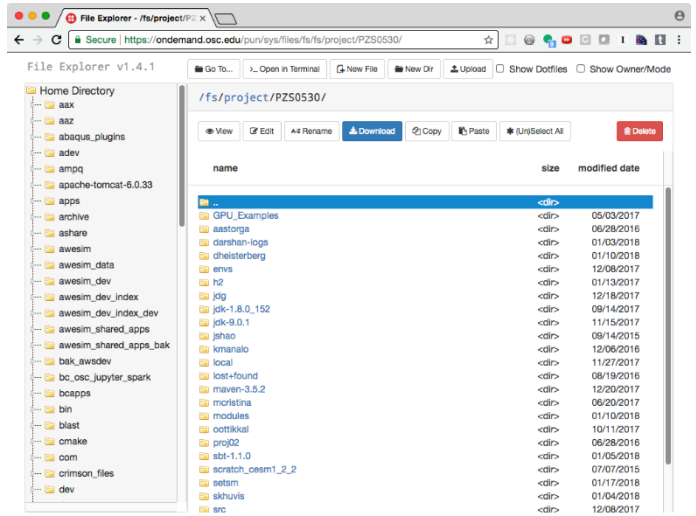

Installation: considerations for production

- **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

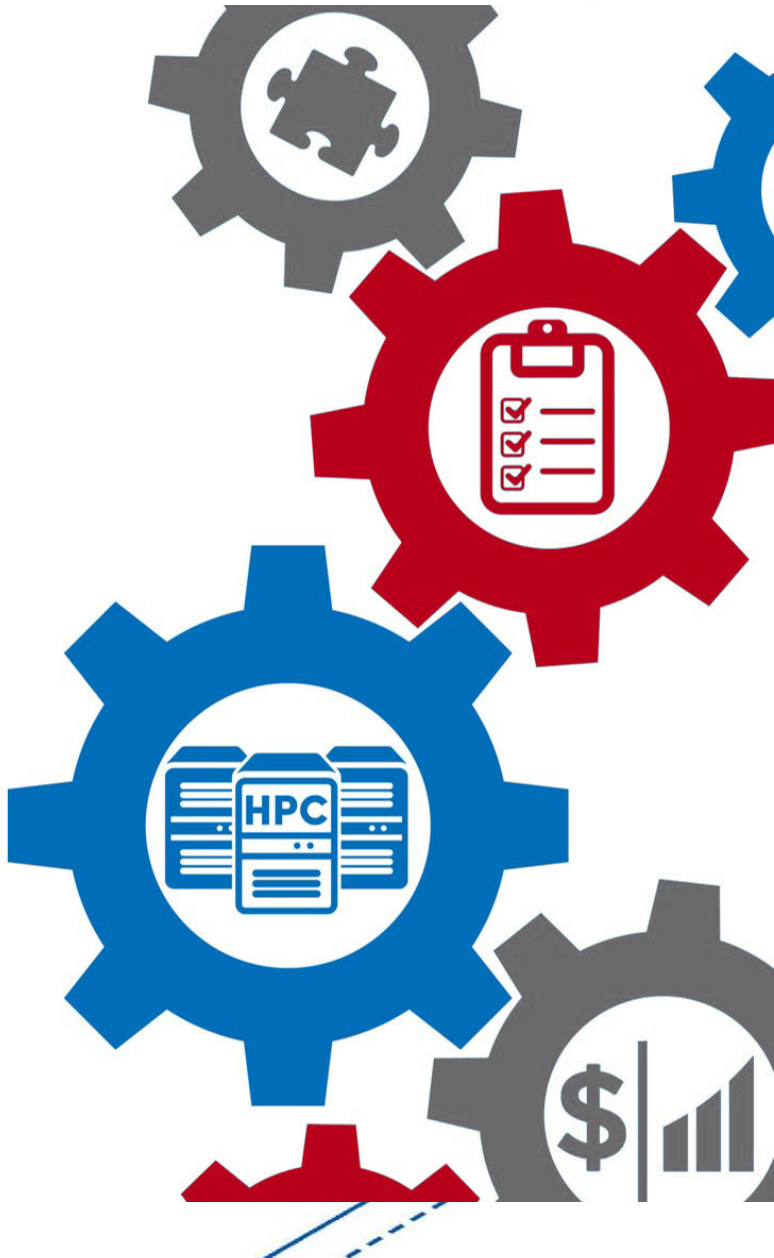

Configuring software to be available in OnDemand

- 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)

Configuring software: Tools

Developer Menu	One place for all developer features Located in top menu in Dashboard
Developer Documentation	Tutorials Example apps to copy and modify
My Sandbox Apps	List of apps I'm developing Create new app and delete app
App Editor	Edit app metadata Open file and shell apps to edit code Displays git status of app's source code
Published Apps	Apps installed here will be visible to every user

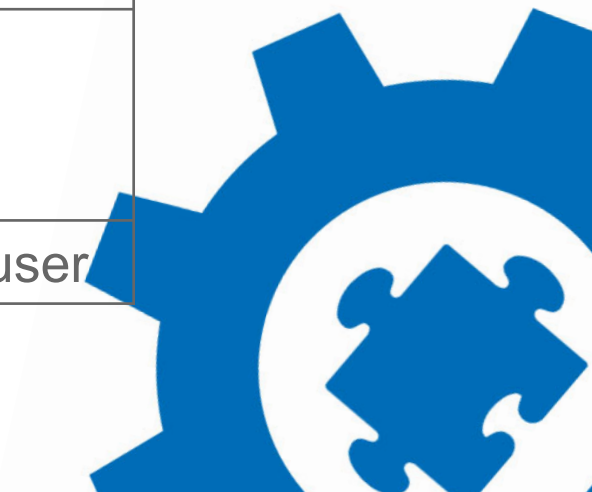

Hands on Tutorial: Create a Jupyter “Interactive App Plugin”

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**

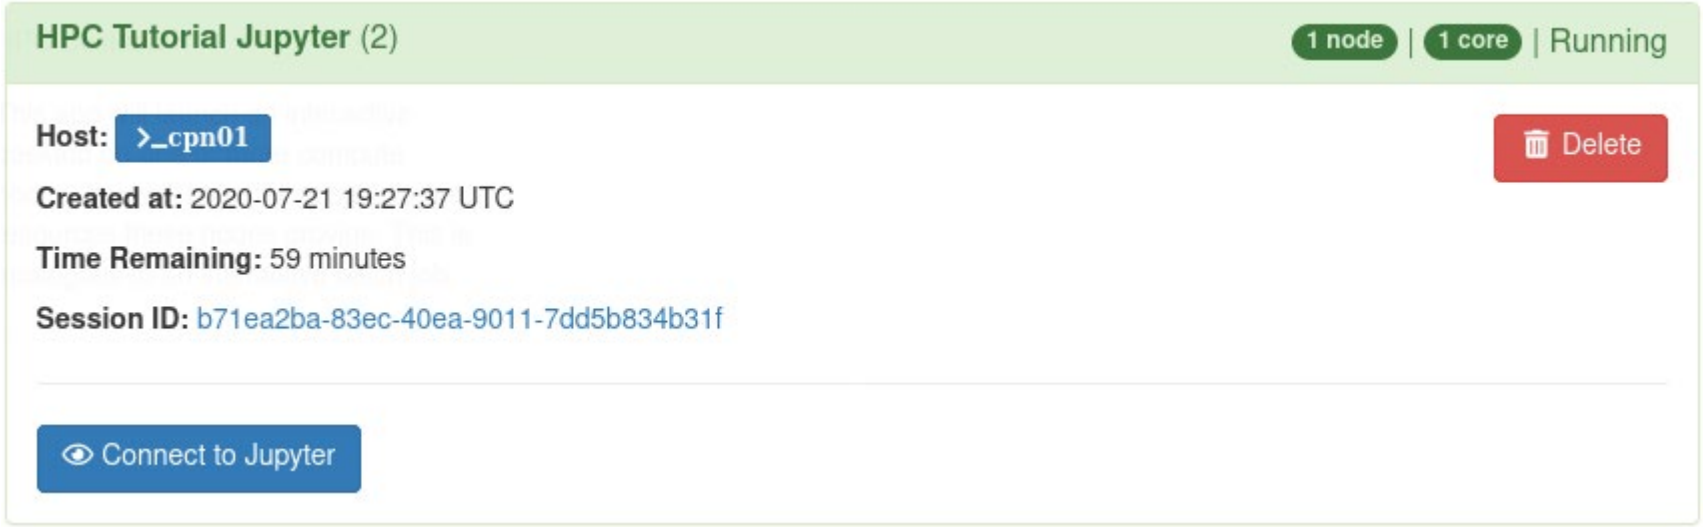

Jupyter Tutorial: Modify the Partition

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

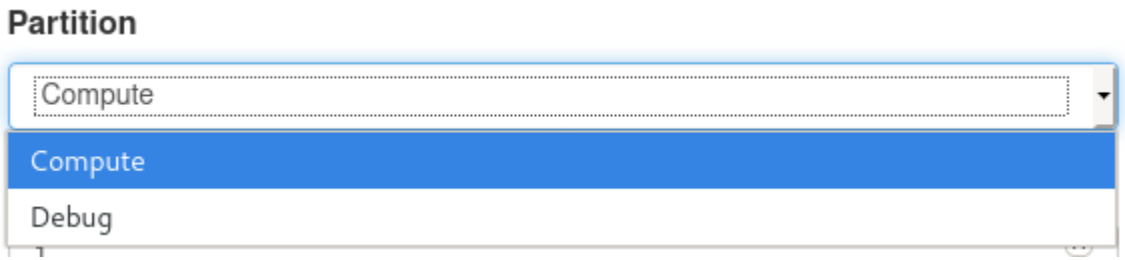

Jupyter Tutorial: Deploy to production

- **Deploy the app to production for other users**

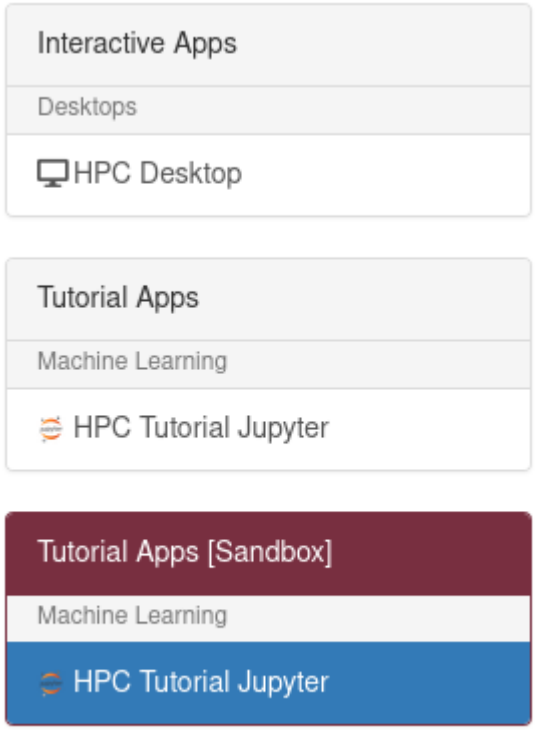

Jupyter Tutorial: Set the memory request for the job

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

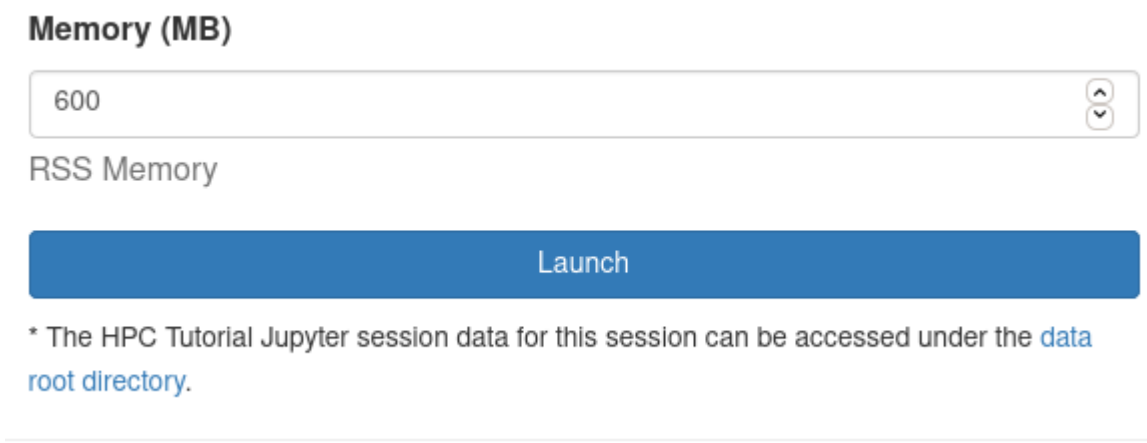

Jupyter Tutorial: Limit the number of nodes

- **Put an upper limit on the number of nodes allowed**

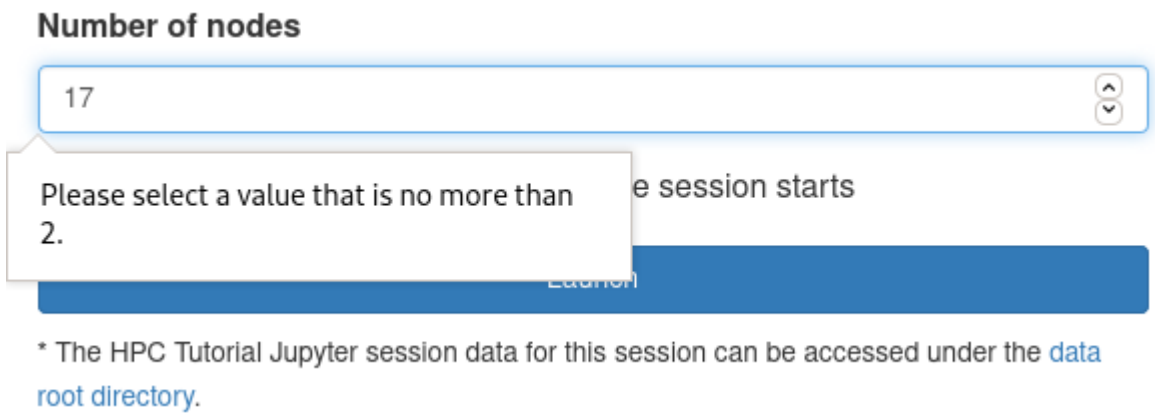

Jupyter Tutorial: Add a checkbox to start JupyterLab

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

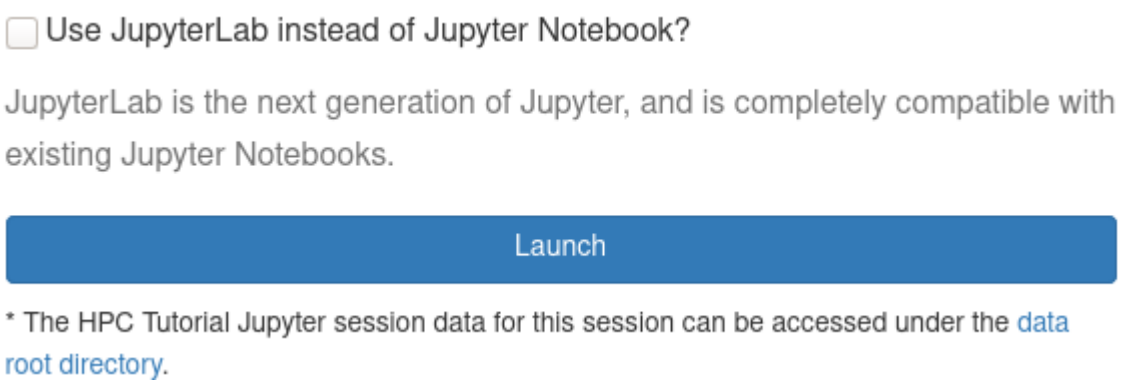

Jupyter Tutorial: Delete unused fields

- **Delete unused fields to clean up the form**

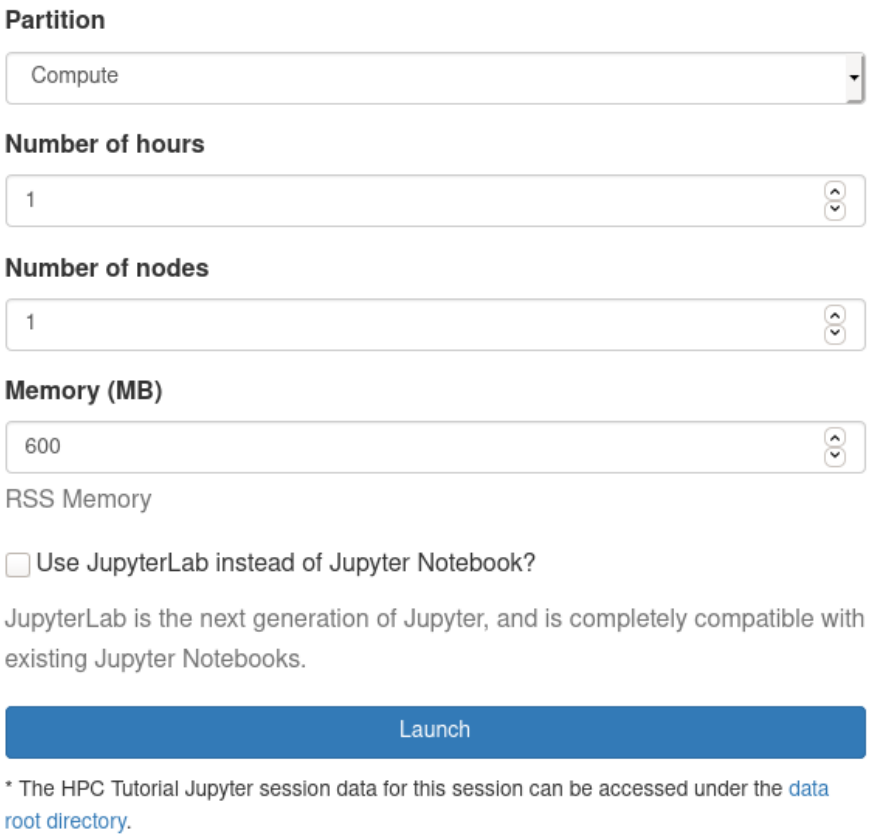

Break

OSC has a job opening on the Open OnDemand team! Full details here: **go.osu.edu/ood-job**

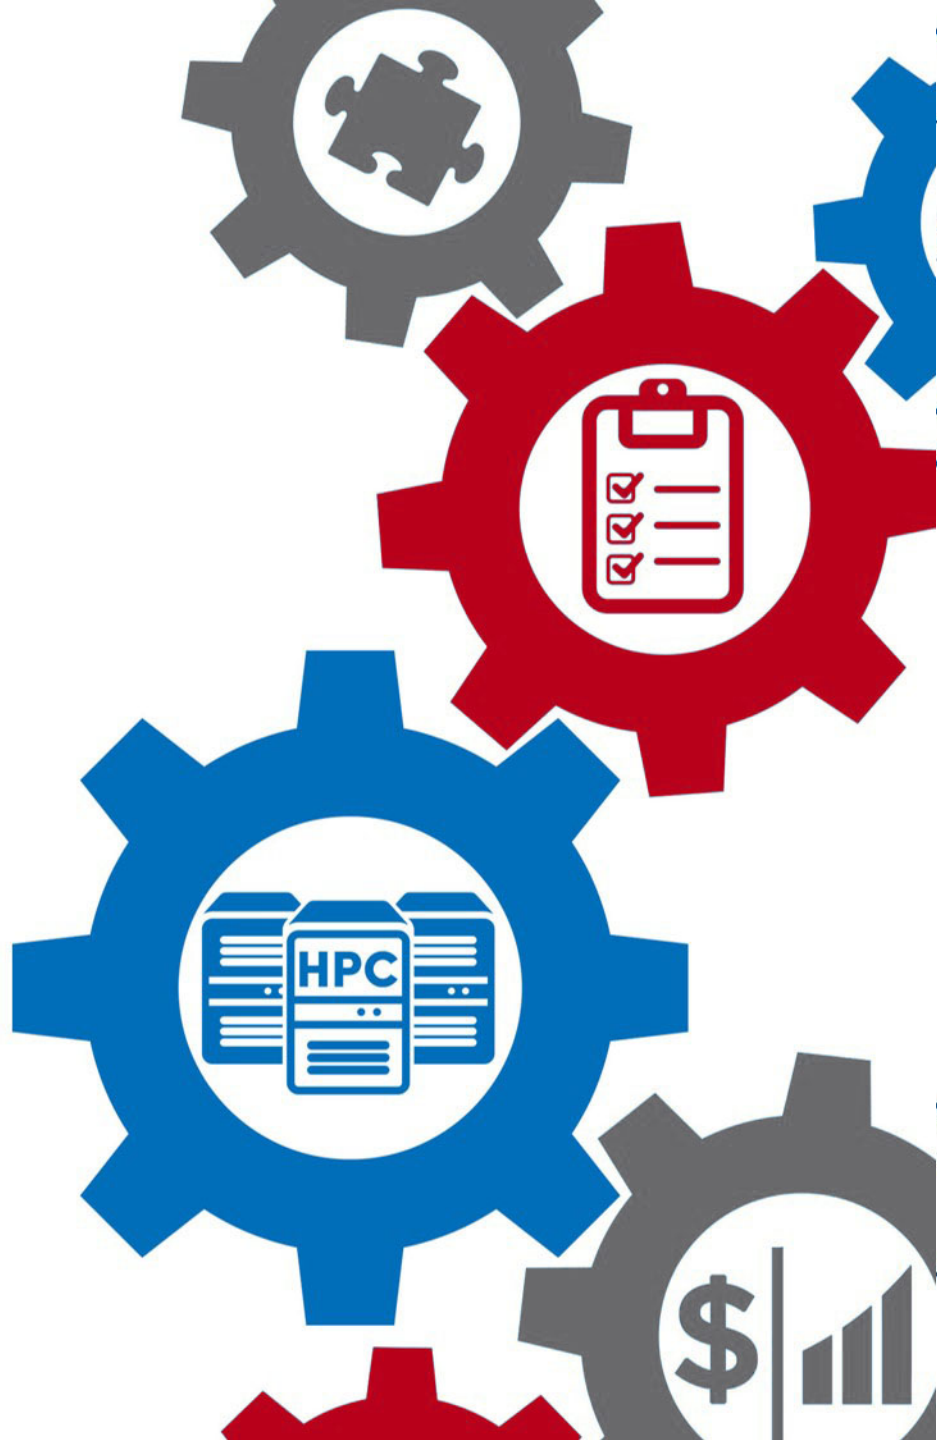

Ohio Supercomputer Center
An OH·TECH Consortium Member

VIRGINIA TECH

University at Buffalo
Center for Computational Research

Hands on Tutorial: 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

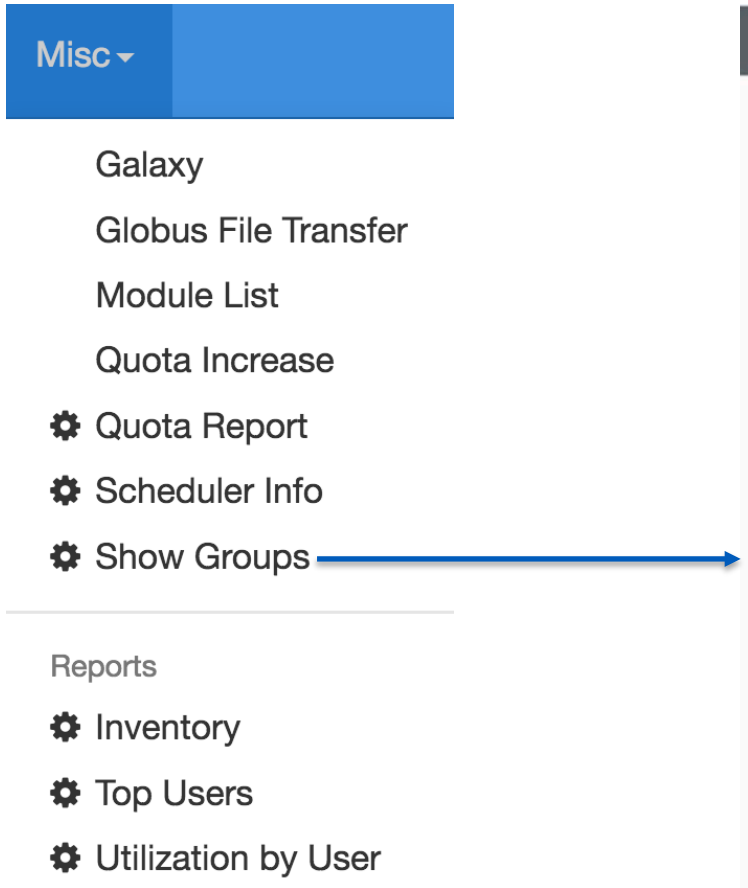

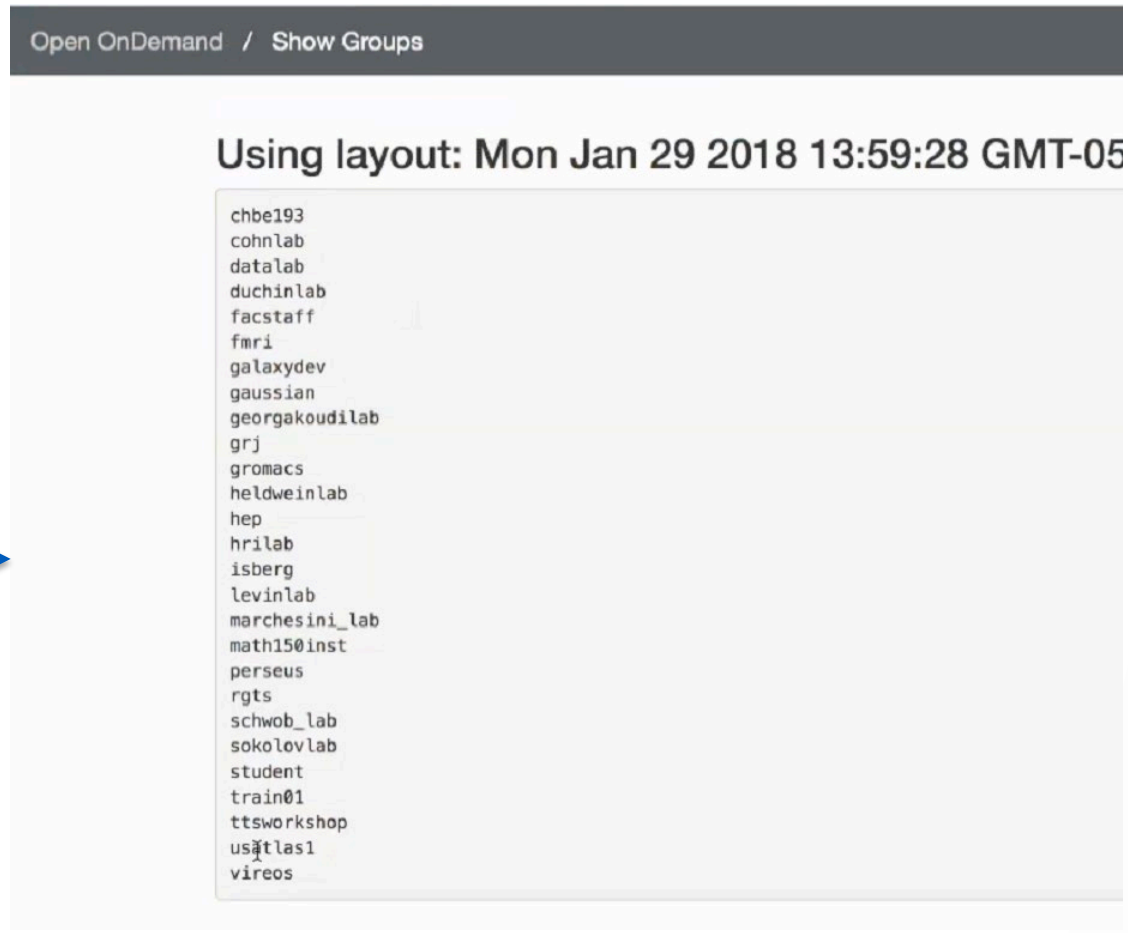

Example Status apps: OSC System Status app

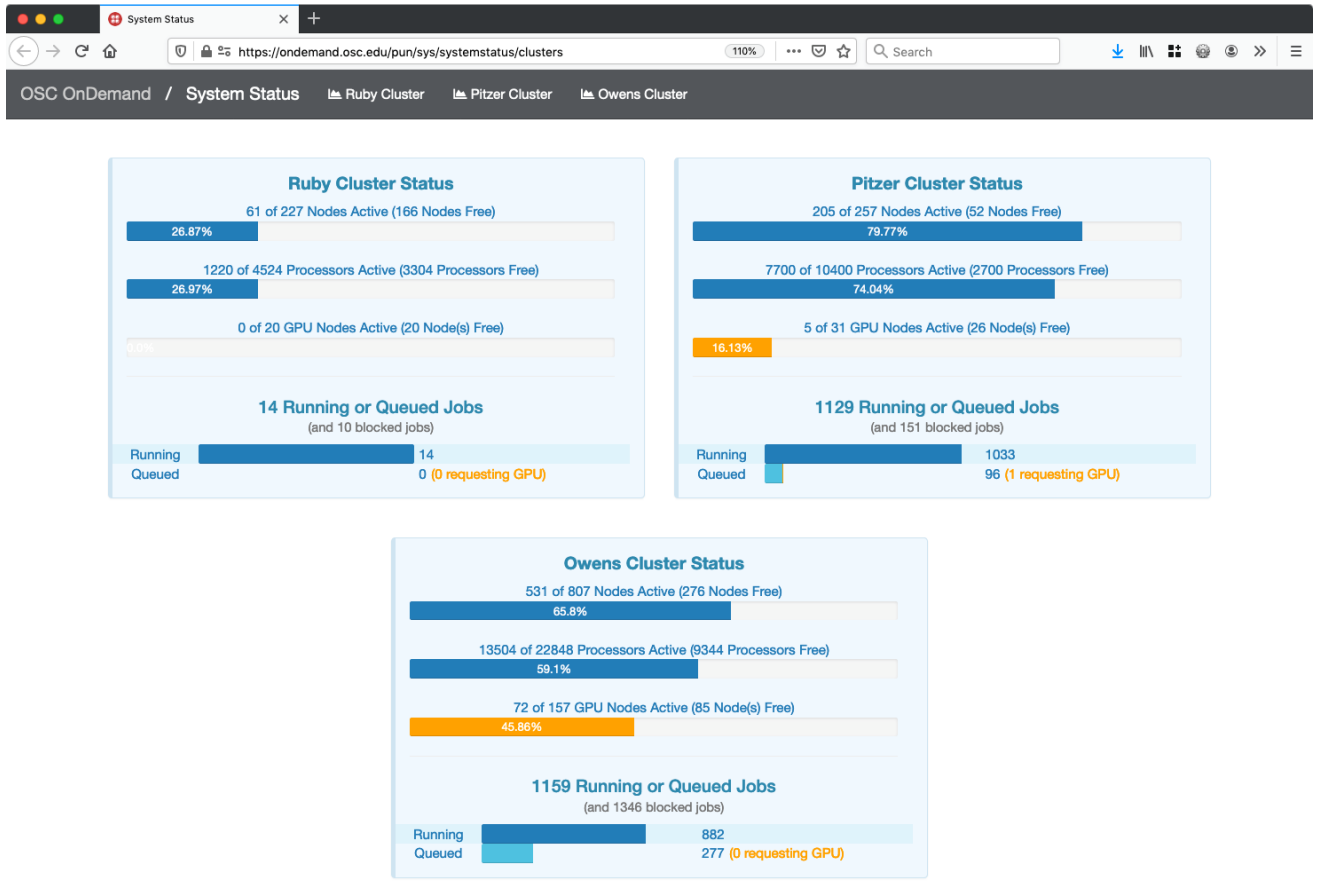

Example third party apps:

- 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

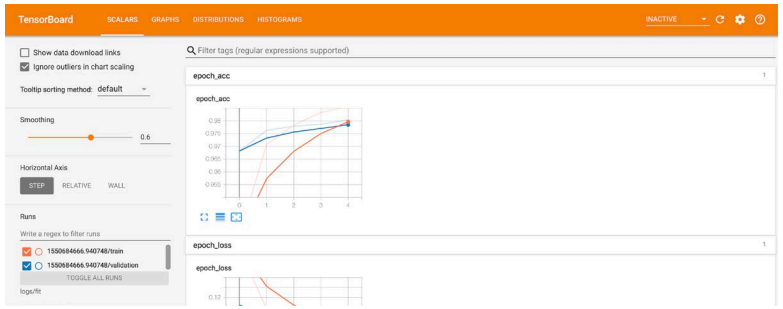

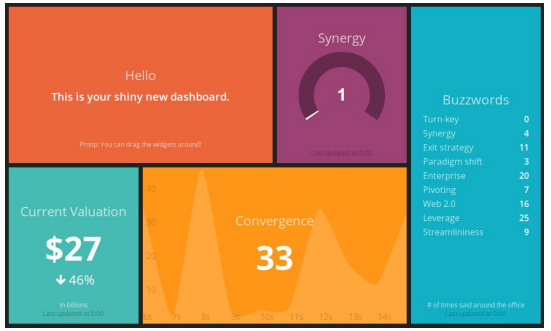

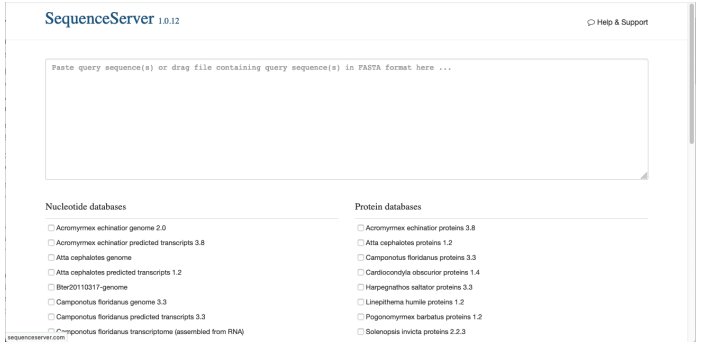

Hands on Tutorial: Create a Passenger app

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

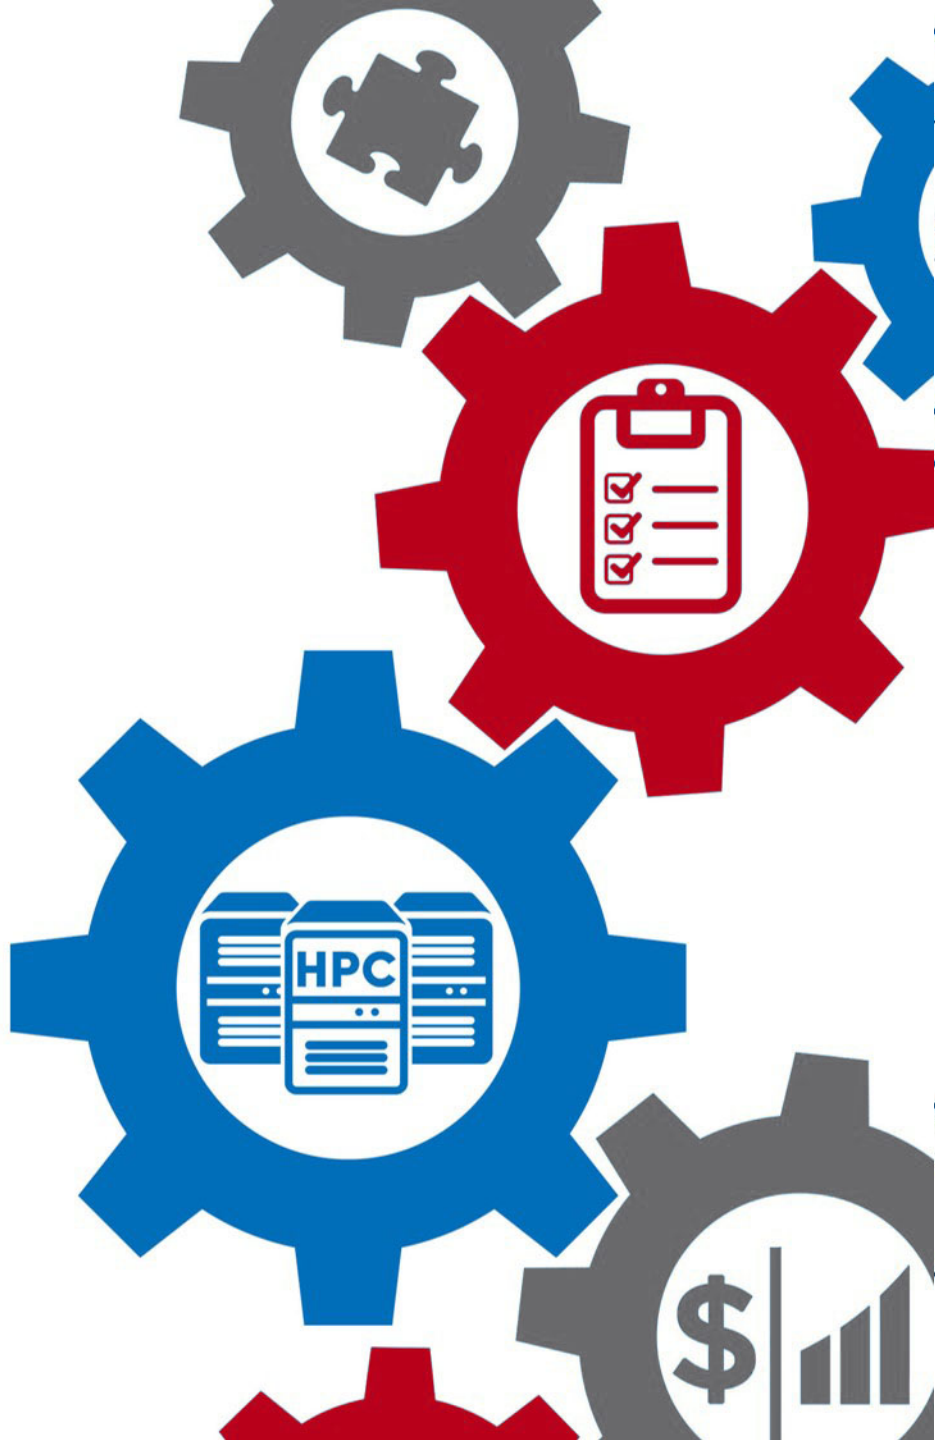

Ohio Supercomputer Center
An OH·TECH Consortium Member

VIRGINIA TECH

University at Buffalo
Center for Computational Research

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

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

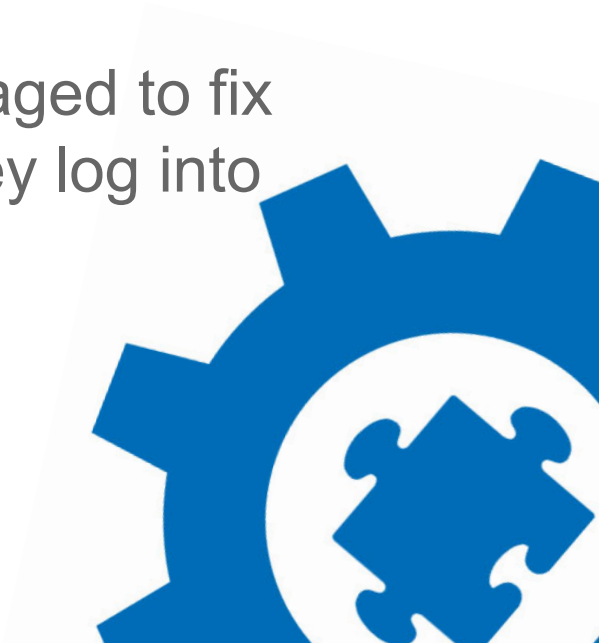

Future plans

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

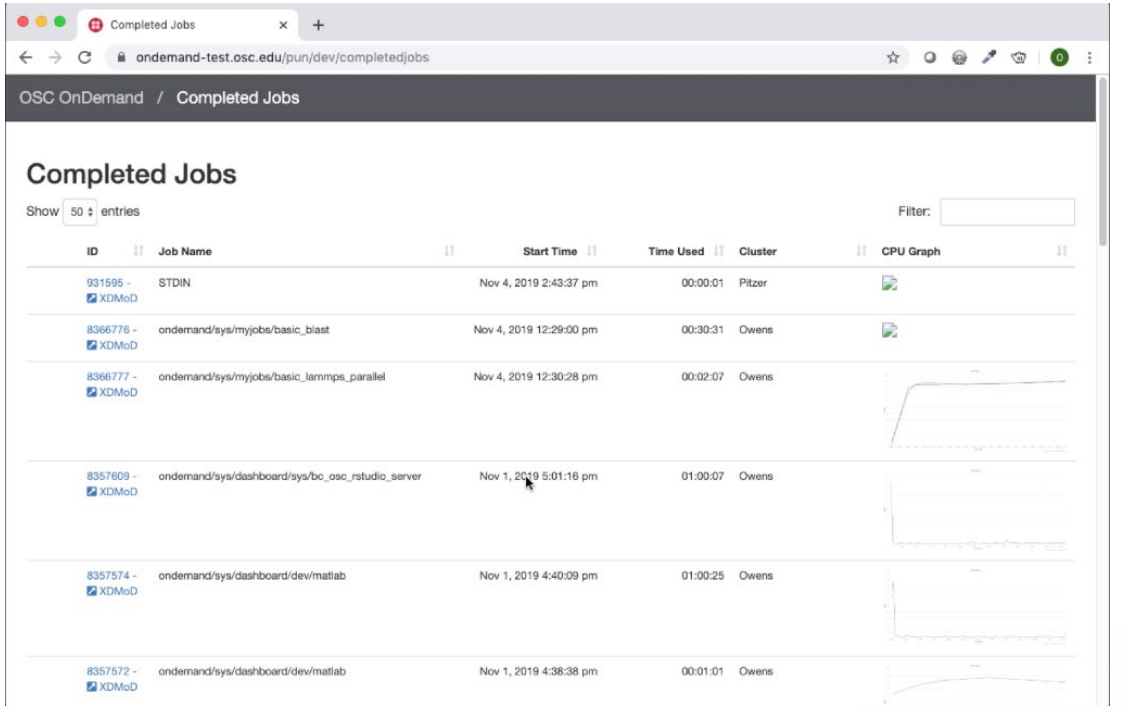

Find more ways to help users optimize their jobs!

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!

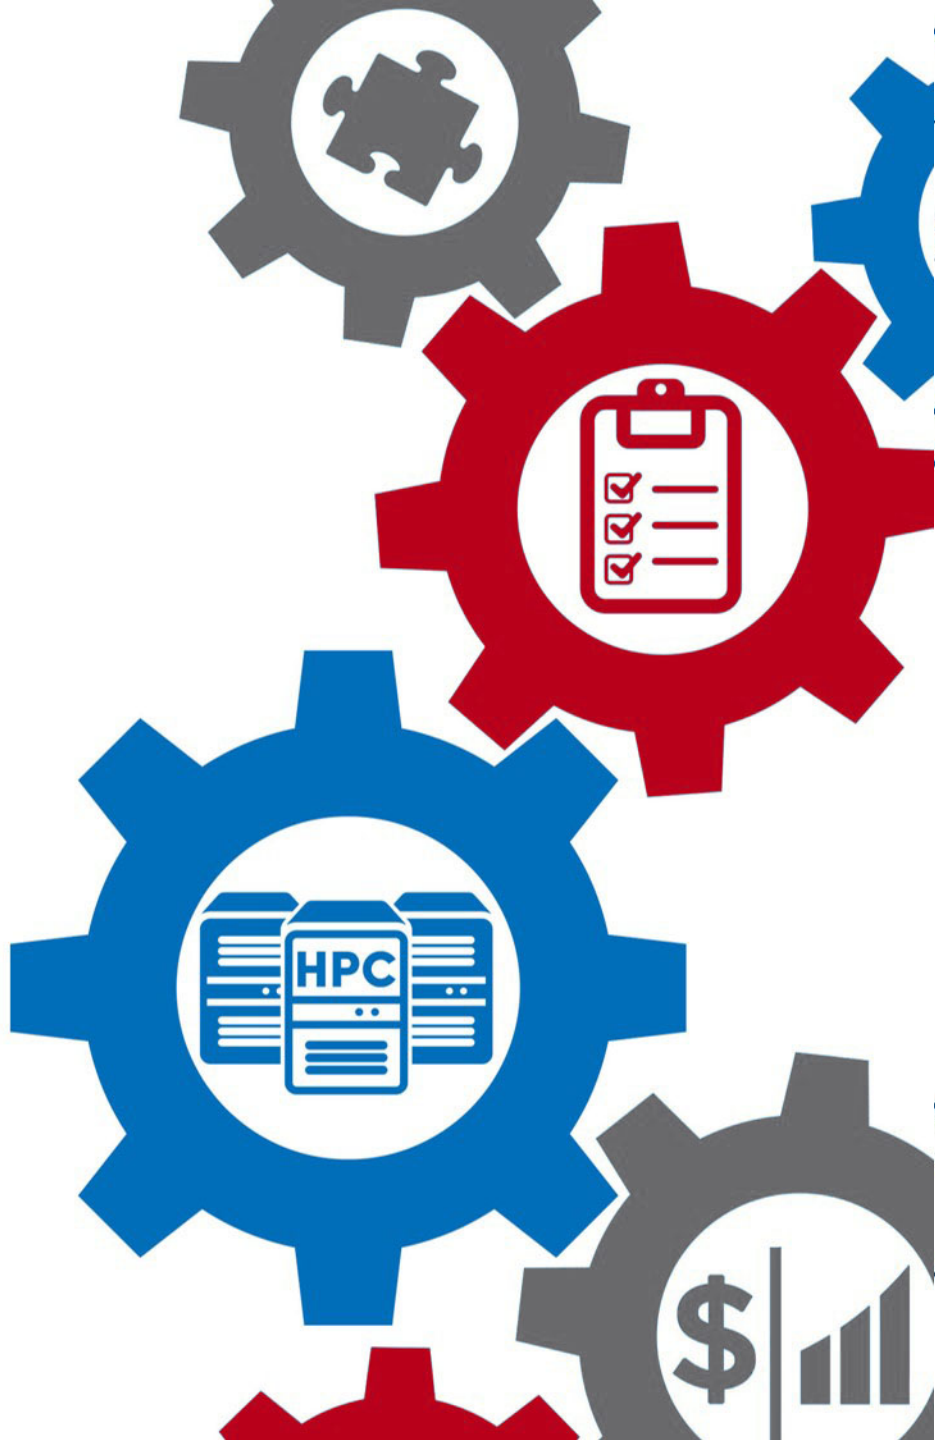

Ohio Supercomputer Center
An OH·TECH Consortium Member

VIRGINIA TECH

University at Buffalo
Center for Computational Research

Tutorial Staff:

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

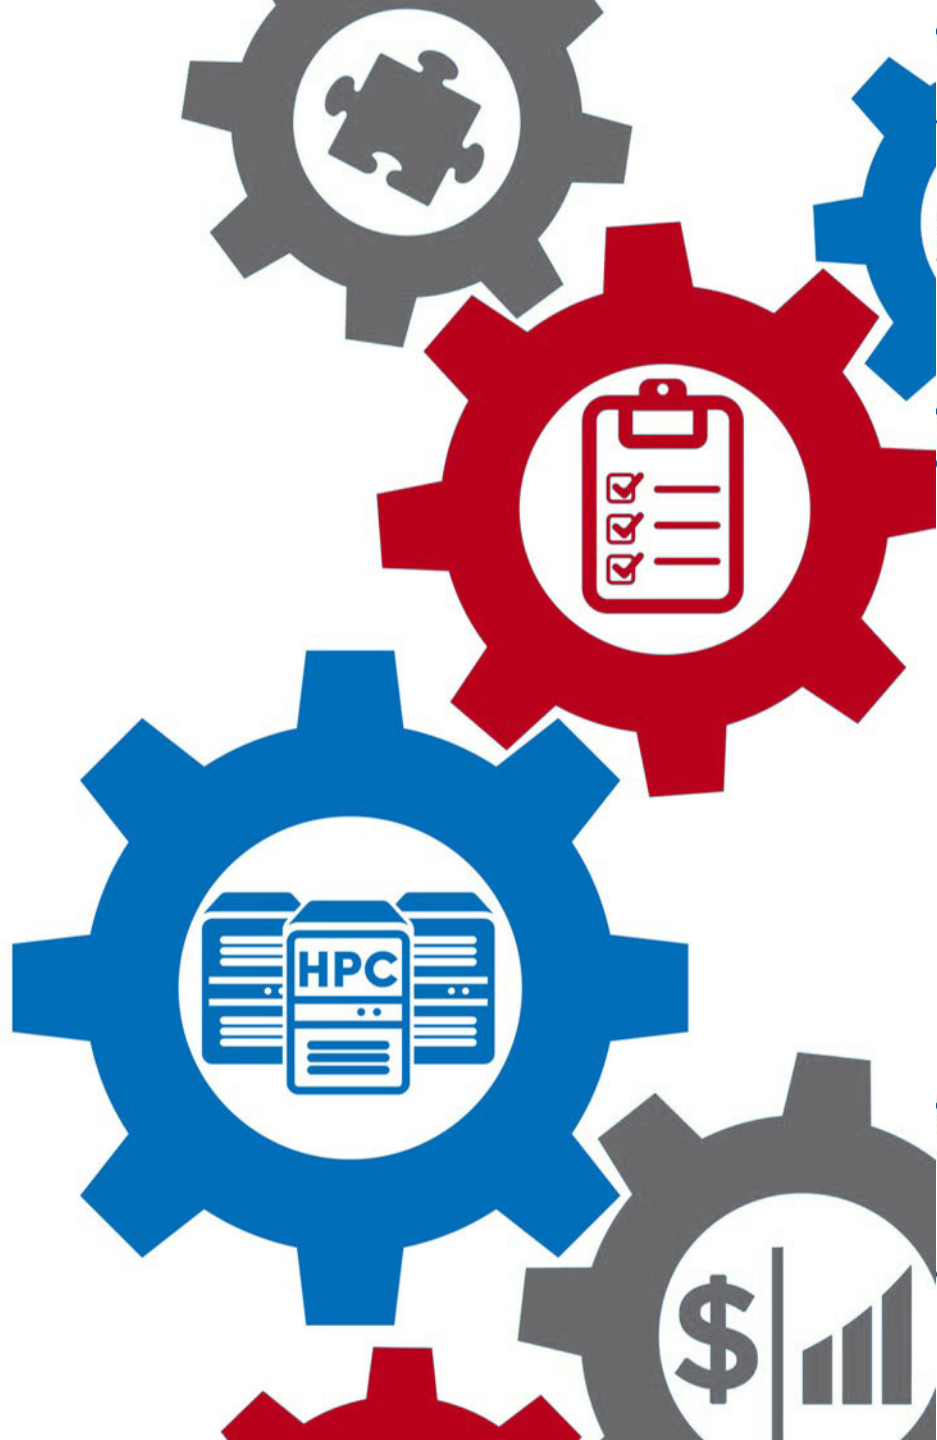

Ohio Supercomputer Center
An OH·TECH Consortium Member

VIRGINIA TECH

University at Buffalo
Center for Computational Research

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 PEARC20

- Tues, 7/28 2:35-3:50pm PST

Open OnDemand User Group Meeting https://sched.co/cnUi

- Wed, 7/29 1:35-3:35pm PST

Grendel: Bare Metal Provisioning System for HPC https://sched.co/cnVj

Monitoring & Analysis of Power Consumption on HPC clusters using XDMoD https://sched.co/cnVp

Phil Andrews Most Transformative Contribution Award and Best Paper in "Advanced research computing environments - systems & system software" Track Award

Informing the on/o-prem cloud discussion in higher education https://sched.co/cnVm

- Thurs, 7/30 8-9:35am PST

Cloud & Data Center usage, expenditures, & approaches to ROI: a survey of academic research computing centers https://sched.co/cnWS

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

Zoom Coordinates will be posted in Slack

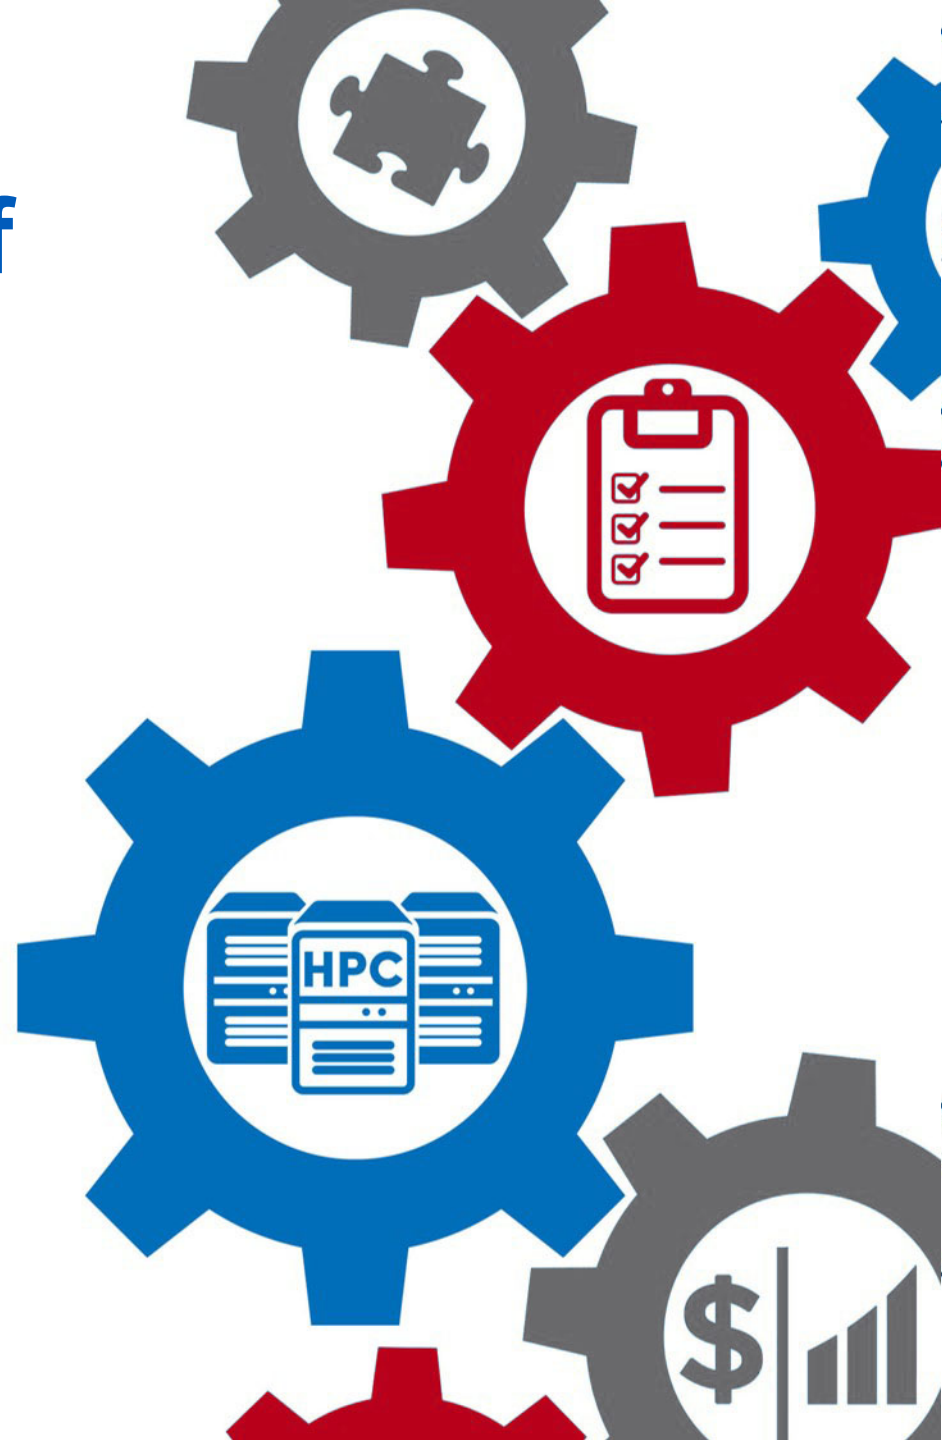

Ohio Supercomputer Center
An OH·TECH Consortium Member

VIRGINIA TECH

University at Buffalo
Center for Computational Research

Thank you for attending!

Please fill out the post-tutorial survey

We value your opinions!

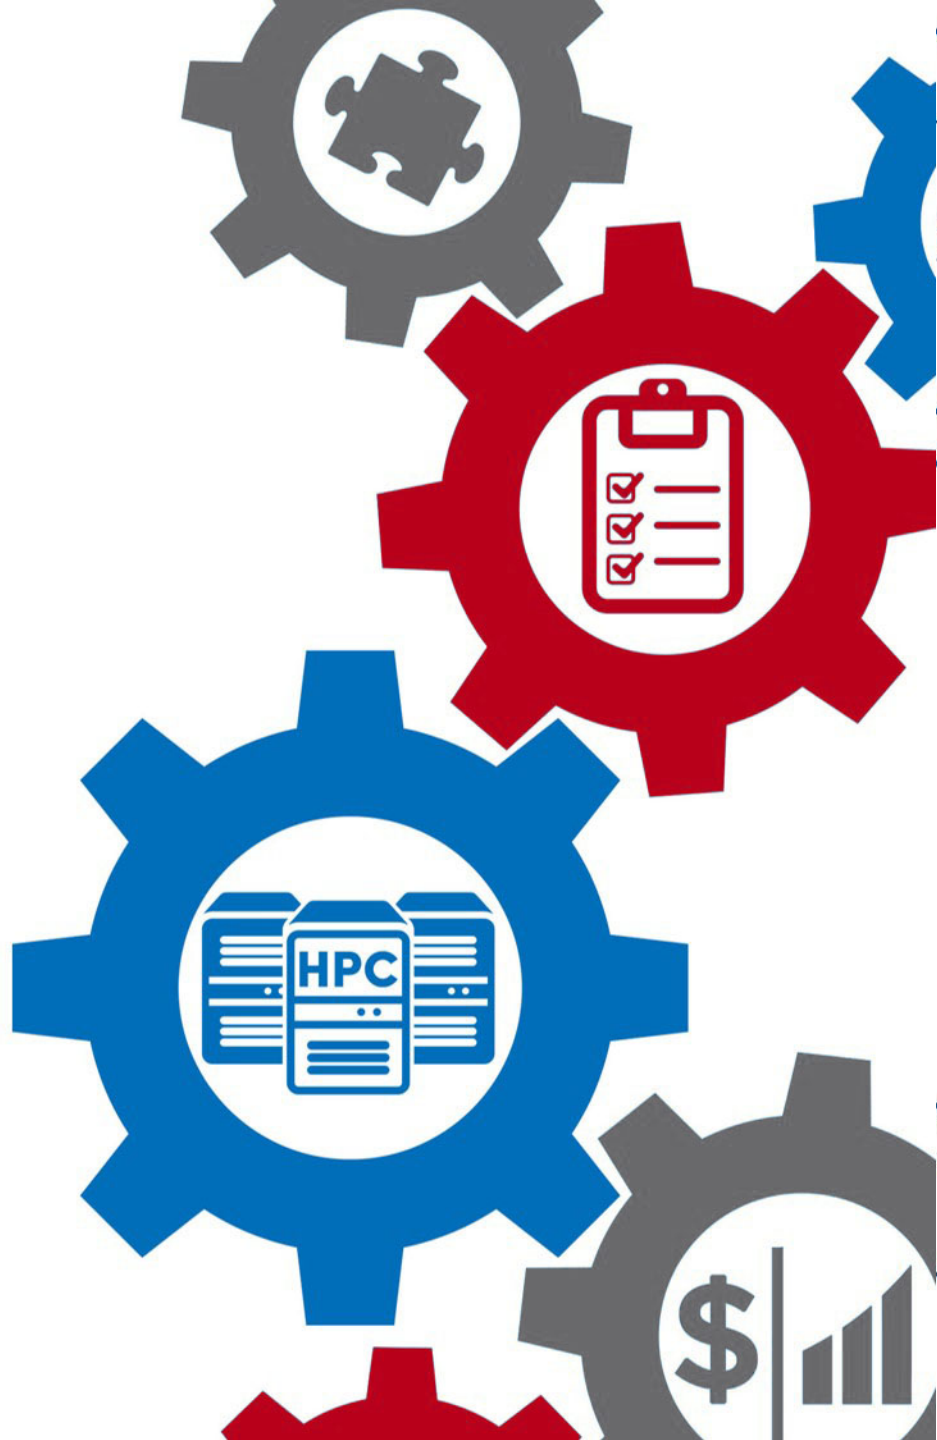

Ohio Supercomputer Center
An OH·TECH Consortium Member

VIRGINIA TECH

University at Buffalo
Center for Computational Research