

# Open OnDemand 2.0 Roadmap



Ohio Supercomputer Center



Alan Chalker, Dave Hudak, Steve Gallo, Eric Franz, Doug Johnson, Bob Settlege

Ohio Supercomputer Center,  
University at Buffalo Center for Computational Research,  
Advanced Research Computing at Virginia Tech

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.

# Webinar Agenda



Ohio Supercomputer Center



1. **About Open OnDemand**
2. The 1.5 release & upcoming 1.6 release
3. Open OnDemand 2.0 Project Overview
4. Engagement
5. Visibility
6. Scalability
7. Accessibility



# Supercomputing. Seamlessly.

## Open OnDemand: Open, Interactive HPC Via the Web

Provides an easy to install and use, web-based access to supercomputers, resulting in intuitive, innovative support for interactive supercomputing.

Features include:

- Plugin-free web experience
- Easy file management
- Command-line shell access
- Job management and monitoring
- Graphical desktop environments and applications

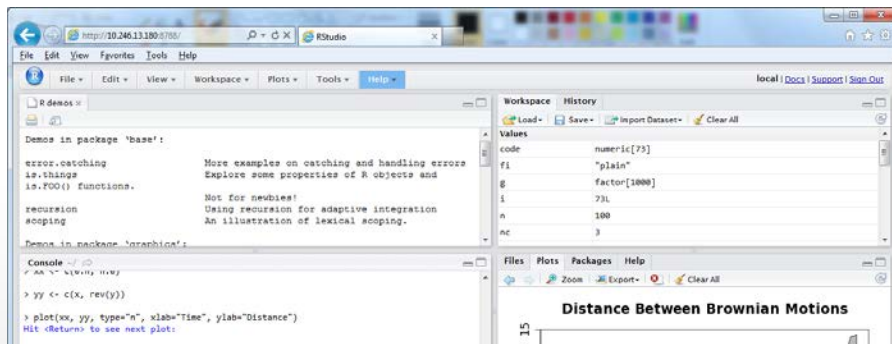


# Interactive Apps

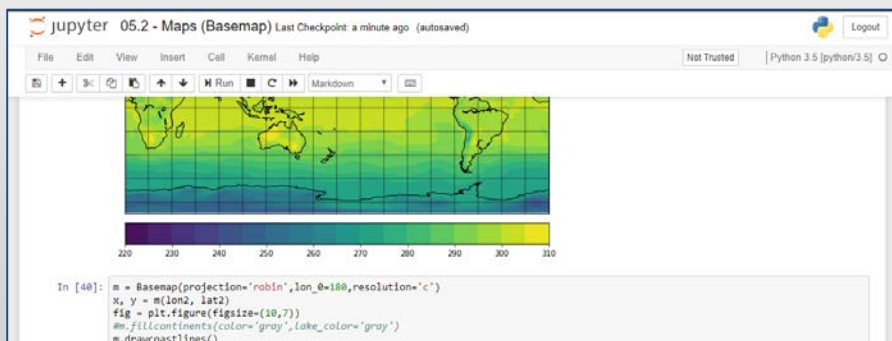
&

# Cluster Access

## RStudio Server – R IDE

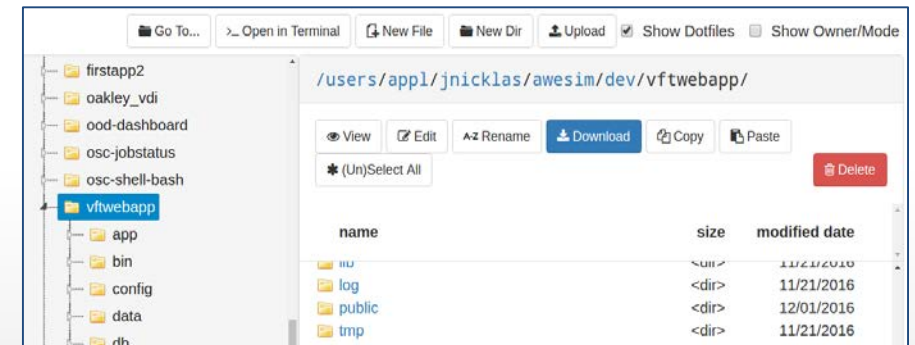


## Jupyter Notebook – Python IDE



And many more, such as ANSYS Workbench, Abaqus/CAE, MATLAB, Paraview, COMSOL Multiphysics

## File Access (browse, edit, etc)



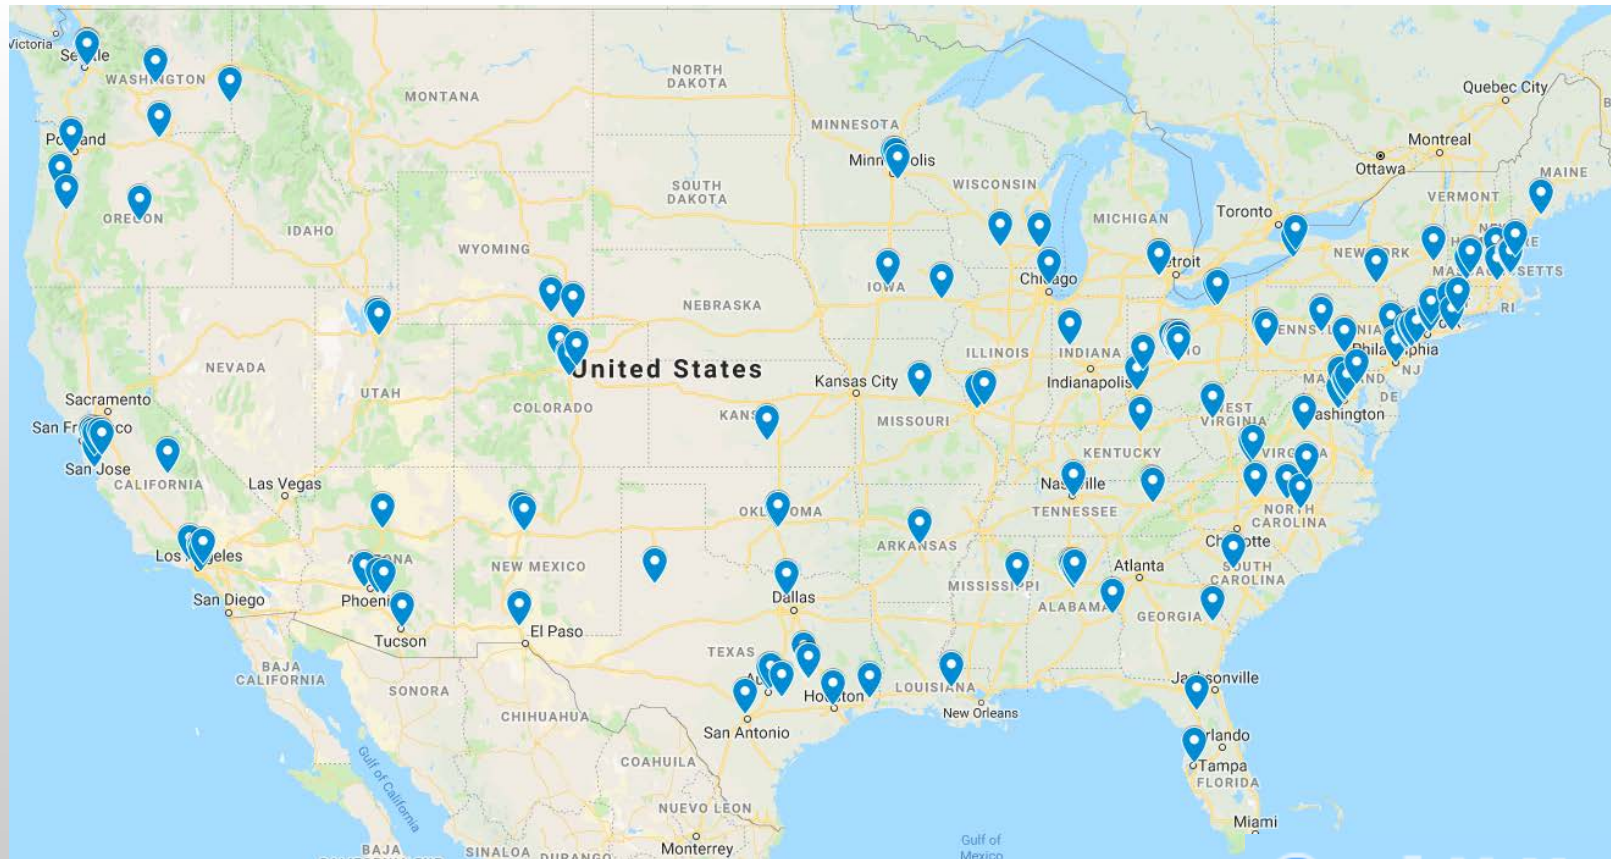
## Manage Jobs (view, submit, etc)

ID	Name	User	Account	Time Used	Queue	Status	Cluster
> 2057900.owe...	high_jp_PIV_N_80_PR_1_2_w_tm	osu9725	PAS1136		parallel	Running	Owens
> 3130444.owe...	RASPA_convert	osu1842	PA00026	140:50:24	serial	Running	Owens
> 3130446.owe...	RASPA_convert	osu1842	PA00026	138:30:25	serial	Running	Owens
> 3130447.owe...	RASPA_convert	osu1842	PA00026	138:09:22	serial	Running	Owens
> 3133547.owe...	high_jp_PIV_N_80_choke_wa_tm	osu9725	PAS1136	17:36:02	parallel	Running	Owens
> 3137260.owe...	Case42	osu8290	PA00008	96:36:34	longserial	Running	Owens
> 3137265.owe...	Case195	osu8290	PA00008	163:01:59	longserial	Running	Owens
> 3137262.owe...	Case261	osu8290	PA00008	165:44:57	longserial	Running	Owens

And many more, such as in-browser SSH terminal, job constructors, VNC desktops

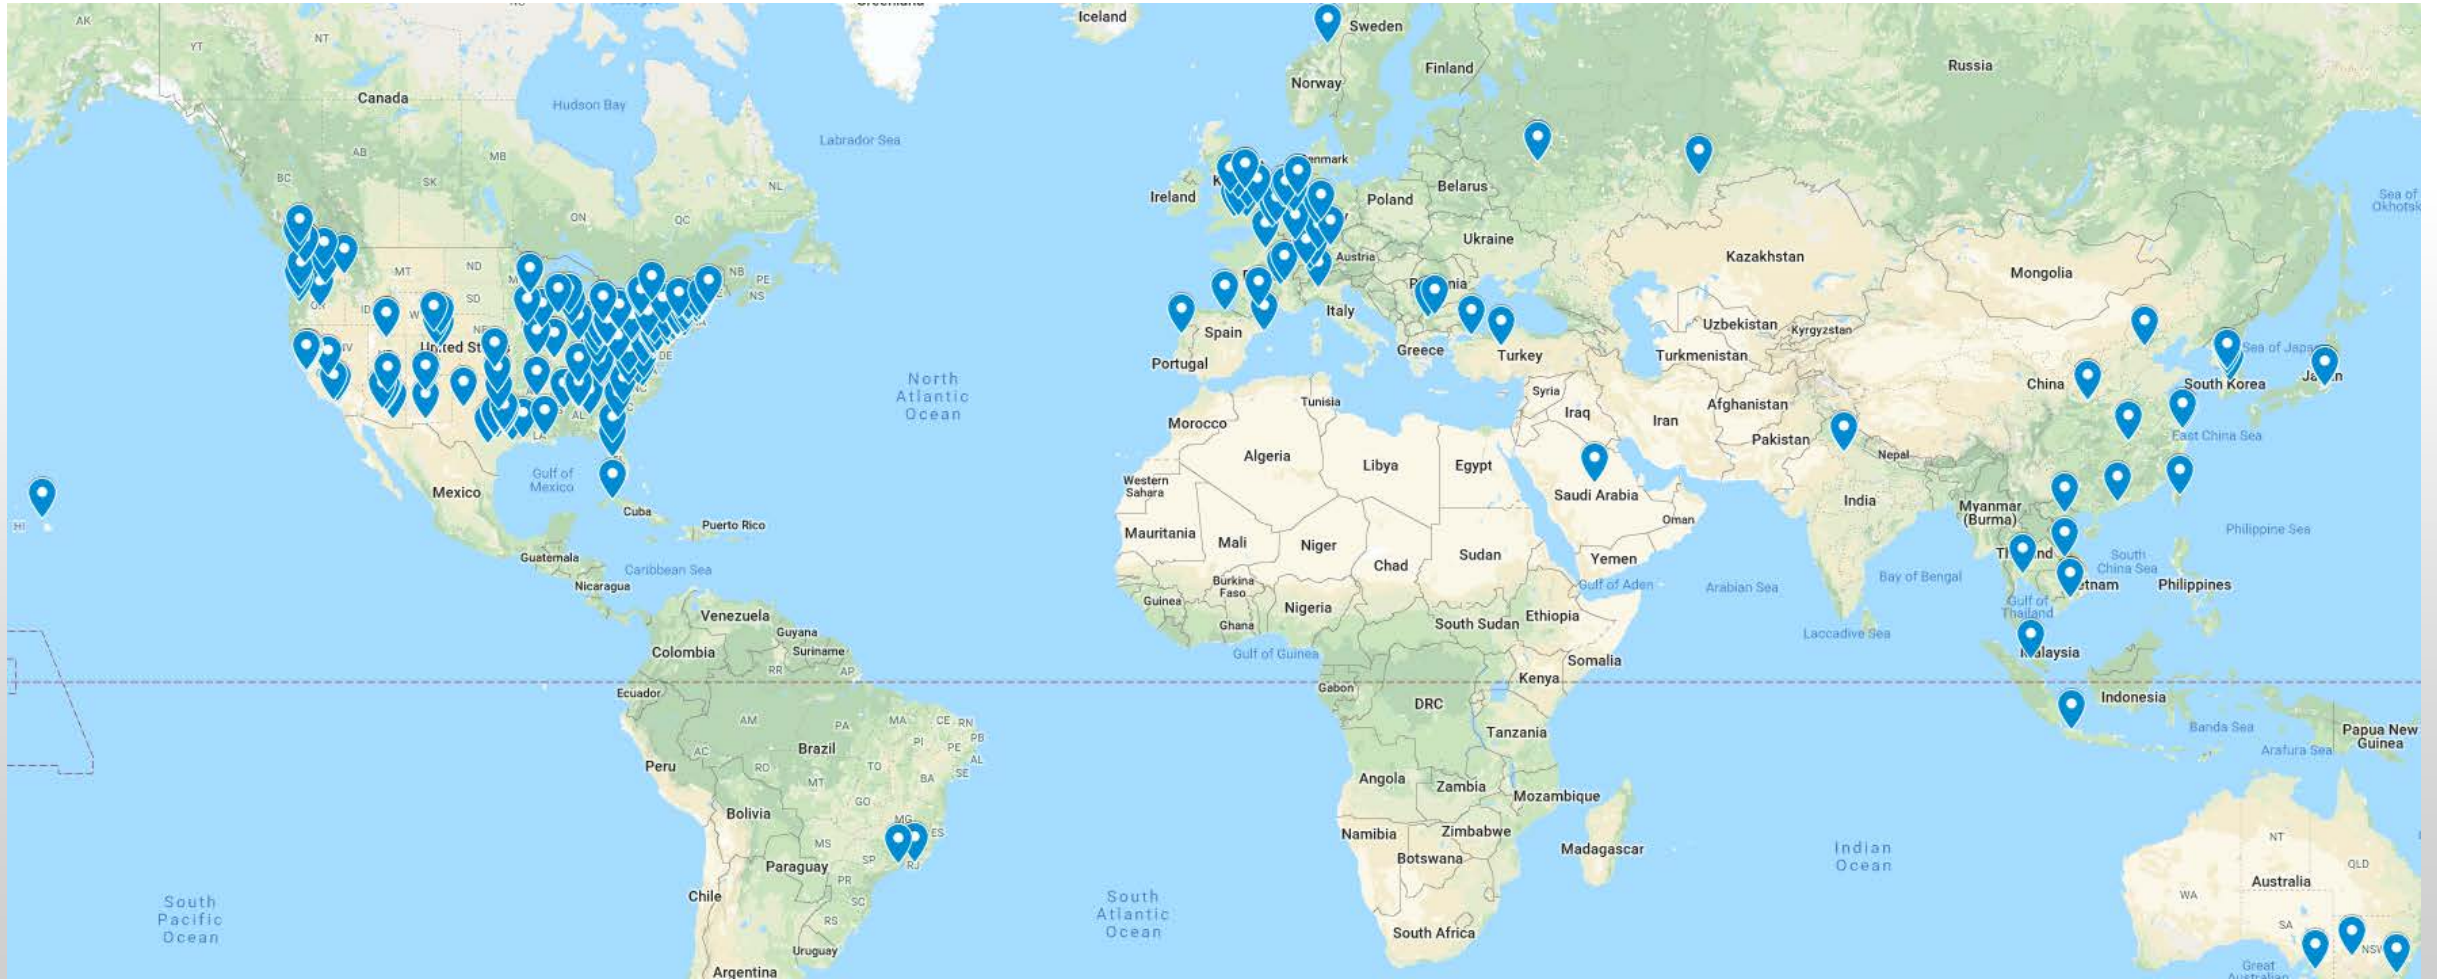
# Approx Number of Institutions based on RPM logs

- 136 unique US locations
- 70 unique international locations



- Map data @2019 Google, INEGI, ORION-ME

# Approx Number of Institutions based on RPM logs



- Map data @2019 Google, INEGI, ORION-ME

# Webinar Agenda



Ohio Supercomputer Center



- ~~1. About Open OnDemand~~
- 2. The 1.5 release & upcoming 1.6 release**
3. Open OnDemand 2.0 Project Overview
4. Engagement
5. Visibility
6. Scalability
7. Accessibility



## Open OnDemand Current Version (1.5 Release, Feb 21 2019)

- Basic Job Array support for Slurm, Torque, and SGE
- proot-less RStudio with varying versions of R
- Easier customization of home page text using localization
- Documentation for two factor authentication with Keycloak and Duo
- Documentation for peer to peer app sharing

## Open OnDemand Upcoming 1.6 Release

- Basic Job Array support for PBSPro and LSF
- More localization options
- Some performance improvements with the job adapters
- Bug fixes and some small requested features, including several contributions from the community
- Targeted for release at end of April

# Webinar Agenda



Ohio Supercomputer Center



1. ~~About Open OnDemand~~
2. ~~The 1.5 release & upcoming 1.6 release~~
3. **Open OnDemand 2.0 Project Overview**
4. Engagement
5. Visibility
6. Scalability
7. Accessibility

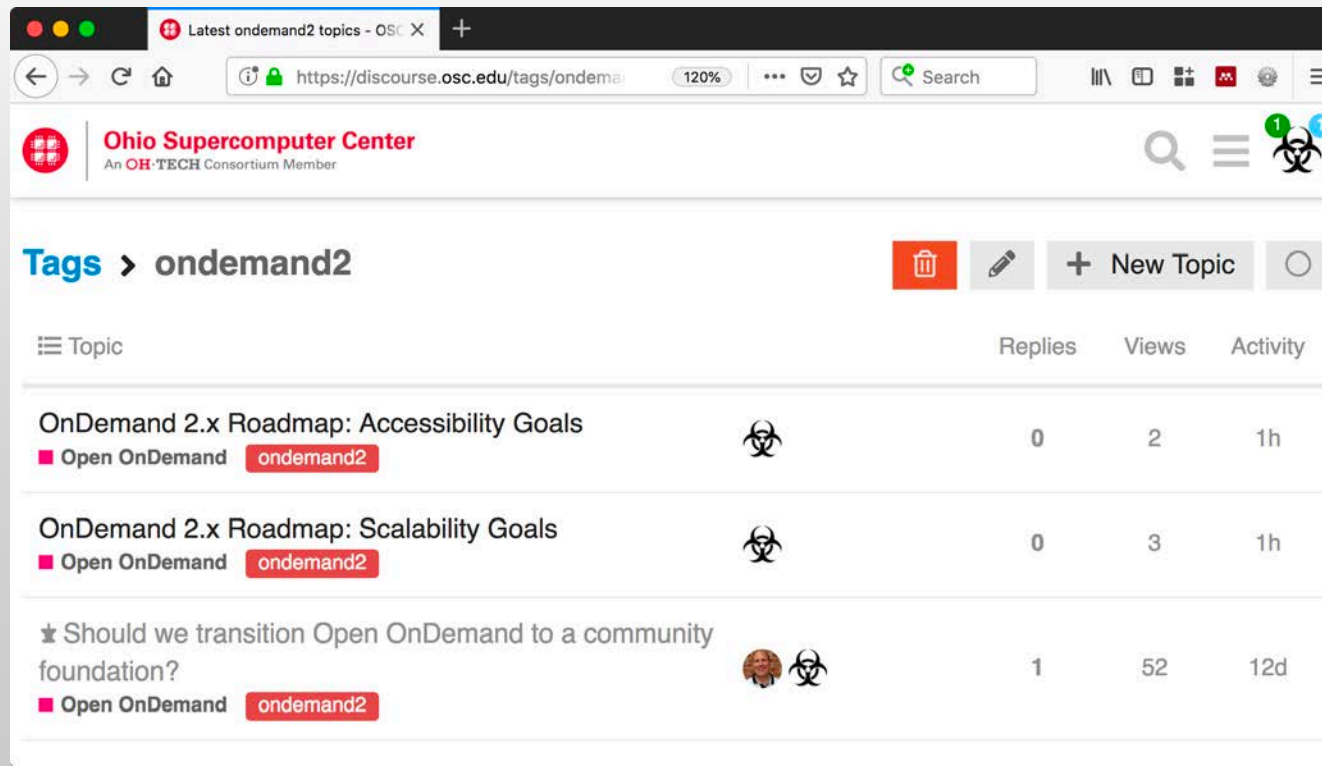


# Open OnDemand 2.0 Project Overview

- Just finished three year NSF SI2 award (#1534949) to develop OnDemand 1.x
- Awarded a 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

# Discourse for follow up discussions

- <https://discourse.osc.edu/tags/ondemand2>
- Pinned topics on Visibility, Accessibility, Scalability, and Engagement and other OnDemand 2.x roadmap discussion topics



The screenshot shows a web browser window displaying the Discourse forum page for the 'ondemand2' tag. The browser's address bar shows the URL 'https://discourse.osc.edu/tags/ondemand2'. The page header includes the 'Ohio Supercomputer Center' logo and the text 'An OH:TECH Consortium Member'. The main content area is titled 'Tags > ondemand2' and features a 'New Topic' button. Below this, a table lists three topics, each with a biohazard icon. The first two topics are 'OnDemand 2.x Roadmap: Accessibility Goals' and 'OnDemand 2.x Roadmap: Scalability Goals', both marked as 'Open OnDemand' and having 0 replies, 2 and 3 views respectively, and being 1 hour old. The third topic is 'Should we transition Open OnDemand to a community foundation?', also marked as 'Open OnDemand', with 1 reply, 52 views, and 12 days old. It features a profile picture and a biohazard icon.

Topic	Replies	Views	Activity
OnDemand 2.x Roadmap: Accessibility Goals ■ Open OnDemand ondemand2	0	2	1h
OnDemand 2.x Roadmap: Scalability Goals ■ Open OnDemand ondemand2	0	3	1h
★ Should we transition Open OnDemand to a community foundation? ■ Open OnDemand ondemand2	1	52	12d

# Webinar Agenda



Ohio Supercomputer Center



- ~~1. About Open OnDemand~~
- ~~2. The 1.5 release & upcoming 1.6 release~~
- ~~3. Open OnDemand 2.0 Project Overview~~
- 4. Engagement**
5. Visibility
6. Scalability
7. Accessibility



**nDemand**

# Engagement: Goals

- Targeting non traditional HPC disciplines
- Advocating for the beginner user
- Outreach
  - i.e. presented OnDemand to Mid-Atlantic Research Infrastructure Alliance (MARIA) HPC Users Group
- Ensure the project is community guided

# Engagement: Science & Client Advisory Group

- Lead by Dr. Robert Settlage of Virginia Tech
- Help guide the OnDemand roadmap
- Members create local focus groups of users to gather feedback on current OnDemand designs, planned changes, and development direction
- Help identify OnDemand apps that can be built to expand HPC access to non traditional domains (humanities, art, economics, linguistics, archeology, medicine)
- Members include representatives from Arizona, Miami University, University of Utah, Oak Ridge National Laboratory, Rutgers, Tufts University, Texas A&M University, NorTech, and San Diego Supercomputer Center

# Engagement: Leveraging OnDemand in Gateway Opportunities

- National-scale Science Gateway community emerging
  - Want to avoid duplication of effort
- NSF is interested in the “science of cyberinfrastructure”
  - OnDemand’s unique per-user web-server architecture is an opportunity for study
- LOGO: How should OnDemand integrate/extend existing Gateway solutions?
  - Sandstone HPC, Galaxy, Apache Airavata

# Engagement: User Survey Results

- Questions included:
  - Usefulness, design, functionality ratings
  - How would you change the design, functionality; other feature requests or comments
- 28 responses
  - Overall OOD usefulness (1 best, 5 worst): 1.73 mean
  - OOD Interface design efficacy (1 best, 5 worst): 2.31 mean
  - OOD Interface functionality (1 best, 5 worst): 2.08 mean
- Conclusion: OOD is useful but there is room for improvement on the design and functionality
- We are working through all the detailed comments and will include it in discussion on Discourse for this year's roadmap

# Engagement: Move to a community foundation?

- OnDemand adoption has led to two good outcomes:
  - A growing community of practice of adopters answering each others' questions on discourse
  - A notable uptick in developers creating useful software extensions to the platform
- Currently: standard open source GitHub-based workflow of “fork/extend/pull request”
- Should OnDemand move to a formal community foundation (e.g., Apache)?
- Are community ownership and engagement paths for developers (outside of project team) important for OnDemand?

# Questions about Engagement?

# Webinar Agenda



Ohio Supercomputer Center



1. About Open OnDemand
2. The 1.5 release & upcoming 1.6 release
3. Open OnDemand 2.0 Project Overview
4. Engagement
- 5. Visibility**
6. Scalability
7. Accessibility

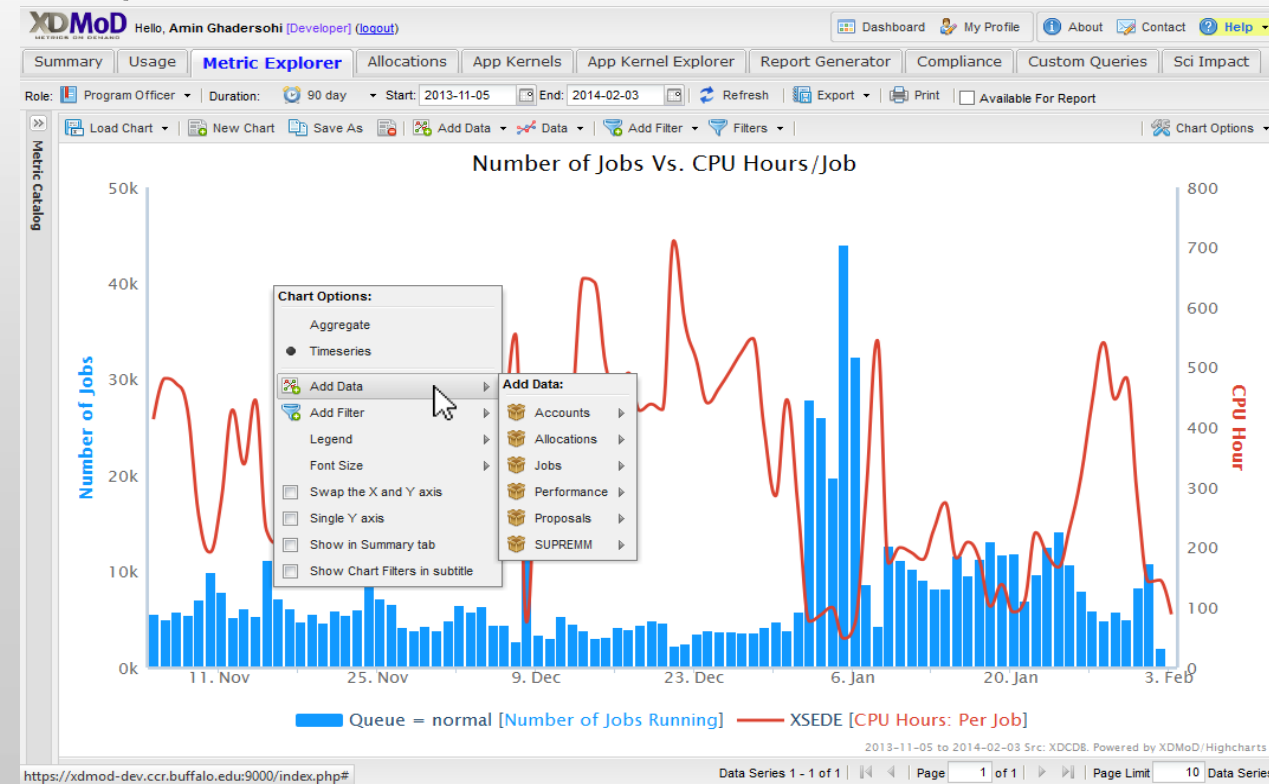


# Open XDMoD

- XDMoD: XD Metrics on Demand
- On demand access to job accounting & performance data
- Optimize resource utilization & performance
  - Utilization metrics
  - Measure infrastructure QoS
  - Job and Cloud level performance data
- 200+ academic & industrial installations worldwide
- <http://open.xdmod.org/>

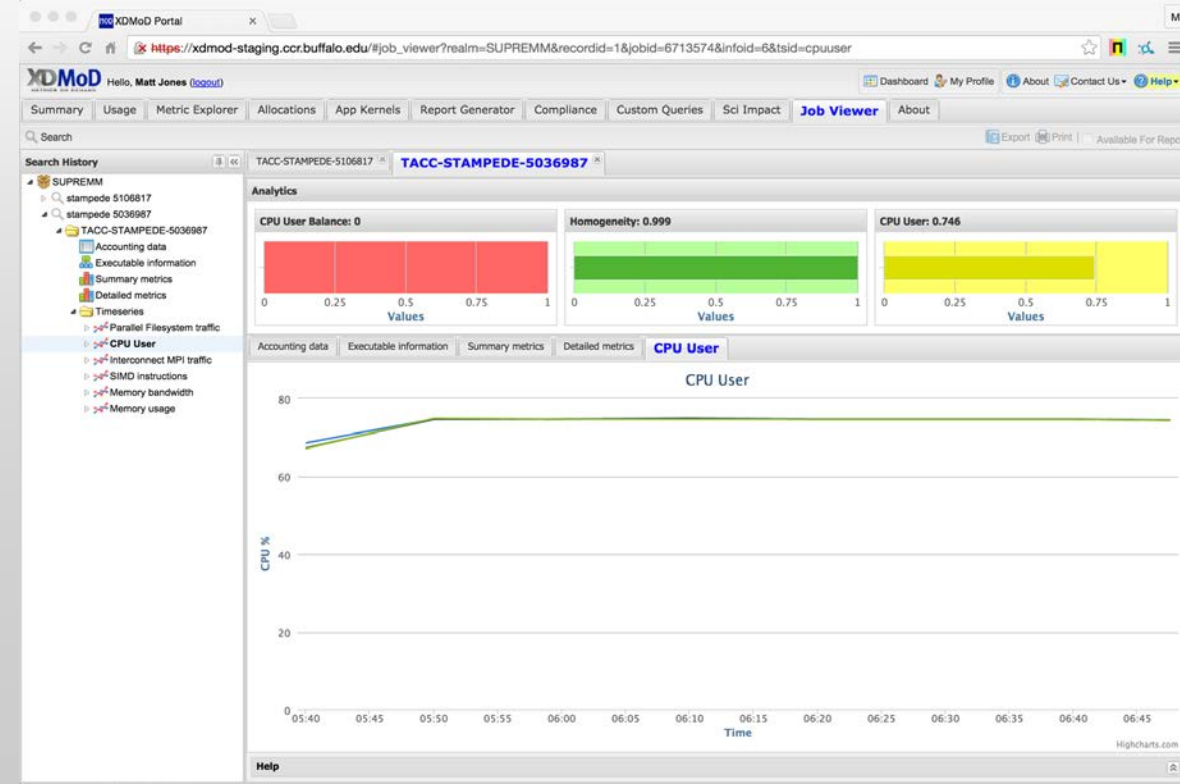
# Open XDMoD Portal

- Web-Based portal for displaying metrics
  - Accounting, utilization, performance
  - Custom report builder (PDF, Word)
  - Quality of Service
  - Role-based



# Open XDMoD Portal

- Integrated with monitoring frameworks (PCP, TACC\_Stats)
  - CPU usage, Memory usage, I/O usage
- Identify poorly performing jobs
  - HPC consultants & users
  - Diagnose problems



# Open XDMoD Demo

## Use Cases

- User
  - Job sizes with shortest queue wait times
- Center Director
  - CPU Hours by Department
- Center Staff
  - CPU Hours by Application
  - Identify users with (potentially) wasted CPU cycles

## Visibility: 5 year goals (bolded year 1 focus)

1. Integrate XDMoD Metrics into OnDemand
  - **Incorporate Job Accounting and Performance Summary**
  - **Develop direct link capability to XDMoD from OnDemand**
  - **Provide OnDemand usage metrics through XDMoD**
  - Link to XDMoD Job Viewer from Active Jobs app
2. Efficiency to Users
  - Develop User Report Card
  - Support user customization of their OOD dashboard
  - Improve summary data and report card based on user feedback
  - Deep learning/machine learning automated detection of application specific performance to optimize user efficiency and time to science

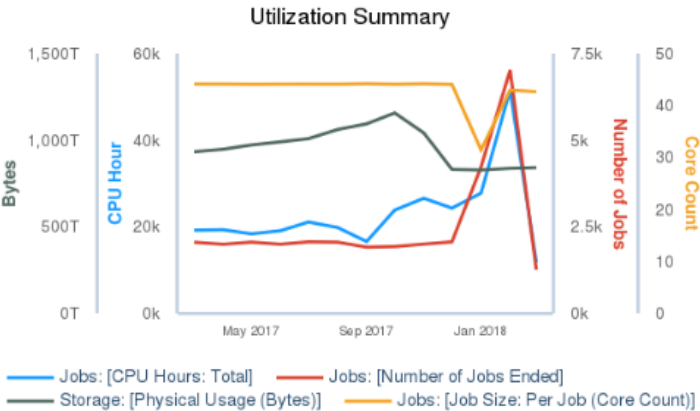
# Open XDMoD and OnDemand Integration

OnDemand provides an integrated, single access point for CCR's HPC resources

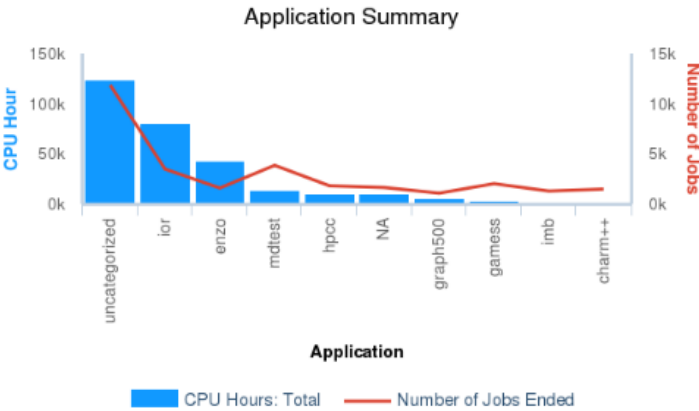
Users can transfer files, access a shell environment on the cluster front-end login server, launch interactive and remote visualization jobs, and monitor jobs all without installing any client software or web plug-ins. Access these features using the menus at the top of this page. Note that many of the apps will launch in a new tab or new browser window but the dashboard will remain open in the original window.

## Utilization Summary

	Previous Month	Previous Quarter	Year To Date
Total CPU Hours	51,541	74,617	298,725
Number of Jobs	7,017	5,973	32,551
Average Job Size (Cores)	42.1	43.9	44.1
Storage (GB)	834	1,008	964,150



2017-03-01 to 2018-03-23 Src: HPCDB, File system storage logs. Powered by XDMoD/Highcharts



2017-03-01 to 2018-03-23 Src: SUPREMM. Powered by XDMoD/Highcharts

# Open XDMoD and OnDemand Integration

- Active Jobs App links directly to Open XDMoD Job Viewer
- Detailed performance data
  - Per core CPU utilization
  - Network performance
  - Storage reads/writes

CCR OnDemand / Active Jobs

All Jobs ▾

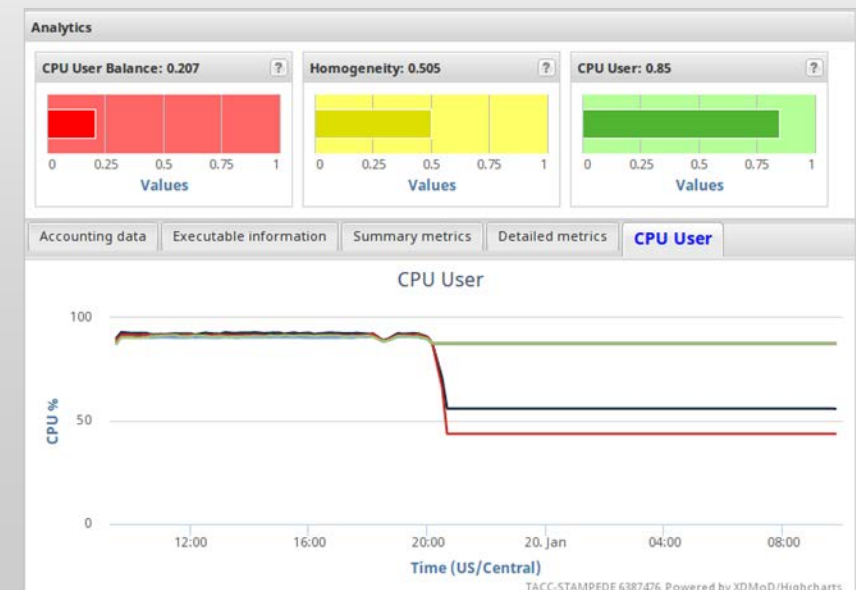
All Clusters ▾

## Active Jobs

Show 50 entries

Filter:

ID	Name	User	Account	Time Used	Queue	Status	Cluster	Performance
8524227	GaVCF_ALL	jw24	big		general-com...	Queued	UB HPC	
8518896	JCG_chr1:36746293	jw24	big	44:12:05	general-com...	Running	UB HPC	
8518944	JCG_chr1:26849169	jw24	big	44:09:54	general-com...	Completed	UB HPC	XDMoD
8518965	JCG_chr1:37553840	jw24	big	44:09:54	general-com...	Completed	UB HPC	XDMoD
8518969	JCG_chr1:45553844	jw24	big	44:09:54	general-com...	Completed	UB HPC	XDMoD
8519044	JCG_chr1:69553856	jw24	big	44:07:31	general-com...	Completed	UB HPC	XDMoD



# Questions about Visibility?

# Webinar Agenda



Ohio Supercomputer Center



- ~~1. About Open OnDemand~~
- ~~2. The 1.5 release & upcoming 1.6 release~~
- ~~3. Open OnDemand 2.0 Project Overview~~
- ~~4. Engagement~~
- ~~5. Visibility~~
- 6. Scalability**
7. Accessibility



# Scalability: 5 year goals (bolded year 1 focus)

## 1. Extendibility

- **Interactive work without a batch scheduler**
- **Increase number of languages to build and deploy apps in (currently limited to Ruby, Python, and Node; extend to PHP, Java, etc.)**
- Support app development with REST API or ported App Kit (i.e. only Ruby)

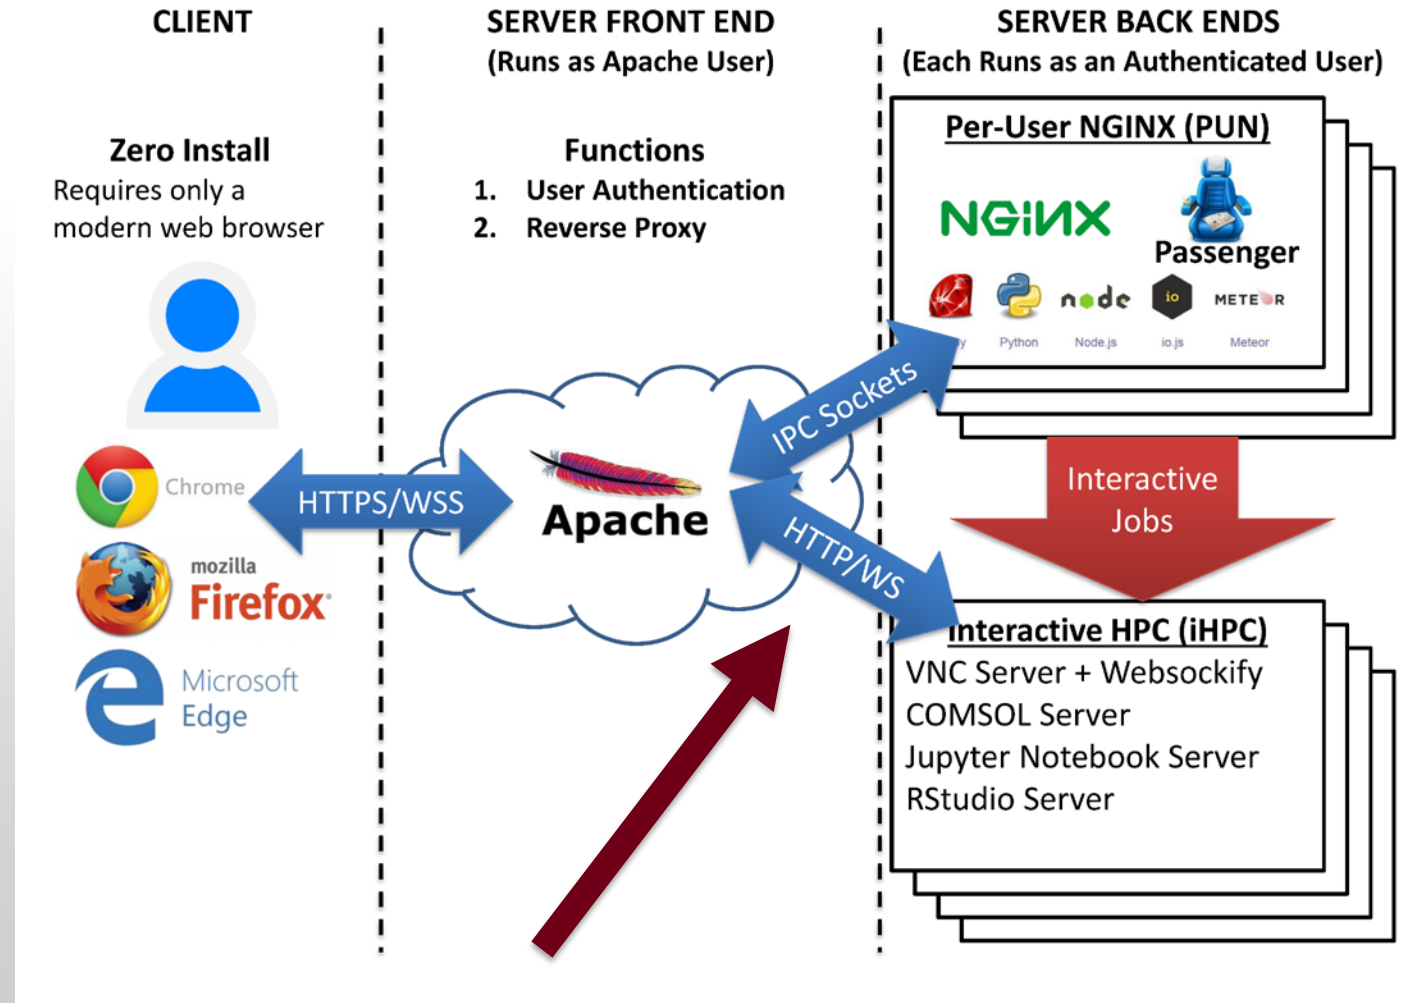
## 2. Performance

- **Reduce response times, start up times, memory usage**
- Support horizontal scaling of Per User NGINX processes
- Explore and implement persistence options to support things like caching

# Scalability: 5 year goals (continued)

## 3. Security

- Secure backend communication between web node and compute node (SSL)
- Common abstraction for authenticating interactive apps



# Scalability: Extendibility - interactive work without a batch scheduler

1. Kubernetes
2. SSH/Fork
3. OpenStack
4. Others?

# Webinar Agenda



Ohio Supercomputer Center



1. ~~About Open OnDemand~~
2. ~~The 1.5 release & upcoming 1.6 release~~
3. ~~Open OnDemand 2.0 Project Overview~~
4. ~~Engagement~~
5. ~~Visibility~~
6. ~~Scalability~~
7. **Accessibility**



## Accessibility: 5 year goals (bolded year 1 focus)

- 1. Improve Job Management**
- 2. Reduce Administrative Load (installation, configuration, debugging)**
- 3. Streamlining interface (reduce steps to accomplish a task)**
4. Integrate Job Management, File Management, and App Launching interface to support new workflows and simplify use
5. Integrated IDE for job and app development

## Accessibility: Job Composer's original goal

1. 99% of all jobs are copies of another job
2. User presented with set of template jobs, each containing
  1. Job script
  2. Input file
3. Workflow
  1. Use web controls to create copy of template as directory in user's home directory
  2. Customize files in directory through shell and files editor and files browser
  3. Use web controls to submit the job script or to stop the job after submission

# Accessibility: Job Composer's limitations

1. Job template customization is a manual step for the user. This is too difficult for beginner users.
2. It creates a new directory for every job – it treats a job as a directory. But a job is not a directory or a file:
  1. A job script is a file/program
  2. A job is an execution of a program
3. It imposes on users a view of their work as a collection of executions that are unsubmitted, queued, running, or completed

Too difficult for the beginner user.

Too limiting for the experienced user.

## Accessibility: goals to address Job Composer limitations

1. decouple job submission from the filesystem (allow submission with just a job script)
2. enable users to easily submit certain kind of jobs
3. provide support for job development “as a first-class citizen”
4. provide support for turning jobs into apps (taking a running job and turning it into a reusable tool)
5. provide a unified job view

# Accessibility: potential solutions

- Extend the Files app.
  - Job Submission: Allow a user to select one or more files and submit them to a scheduler.
  - App “shortcuts”: For example, allow a user to select a .m file and click a “go run MATLAB on this cluster” option. Automatically submits a 1-core, 59-minute job with that MATLAB file as input.
- Support “Batch Apps” in Dashboard (analogous to “Interactive Apps”)
  - Batch App (e.g., MATLAB batch app) would allow user to specify hardware, wall time, input file(s) and output location.
- Unified Jobs app “My Jobs”
  - Completed jobs (with stats from XDMoD)
  - Running jobs from scheduler (with stats from Ganglia)
  - Queued jobs from scheduler
- Job Script IDE – writing a job script is software development!
  - Incorporate version control
  - Syntax highlighting and checking
  - Wizards for hardware configuration and job environment configuration (Eclipse PTP may have examples)
  - Lint to check script for errors
  - Slurm Assist from Sandstone HPC
- Incorporate Galaxy or HubZero RAPPTURE for turning jobs into apps

# Accessibility: interim solutions – improve Job Composer app

Feature improvements could help some beginner and intermediate users make use of the Job Composer:

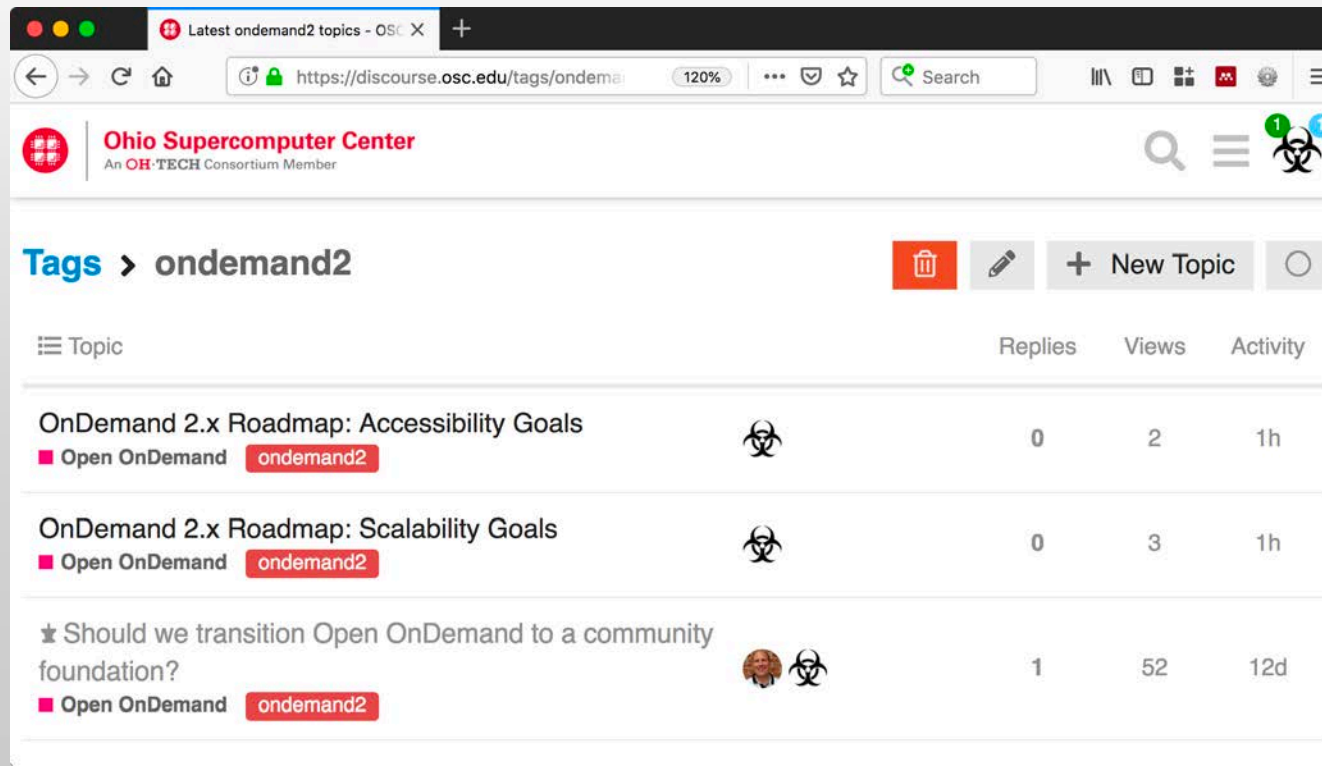
1. Group job directories together by “project”
2. Control where job directories are created
3. Control the name of the job directories when you make a copy of a template
4. Extend Job Options web form with all the resource manager options available and allow for site specific customizations
5. Separate between job directories and job executions in the interface

There are more feature requests from community on Discourse and GitHub

Questions about Scalability, Accessibility or anything else?

# Discourse for follow up discussions

- <https://discourse.osc.edu/tags/ondemand2>
- Pinned topics on Visibility, Accessibility, Scalability, and Engagement and other OnDemand 2.x roadmap discussion topics

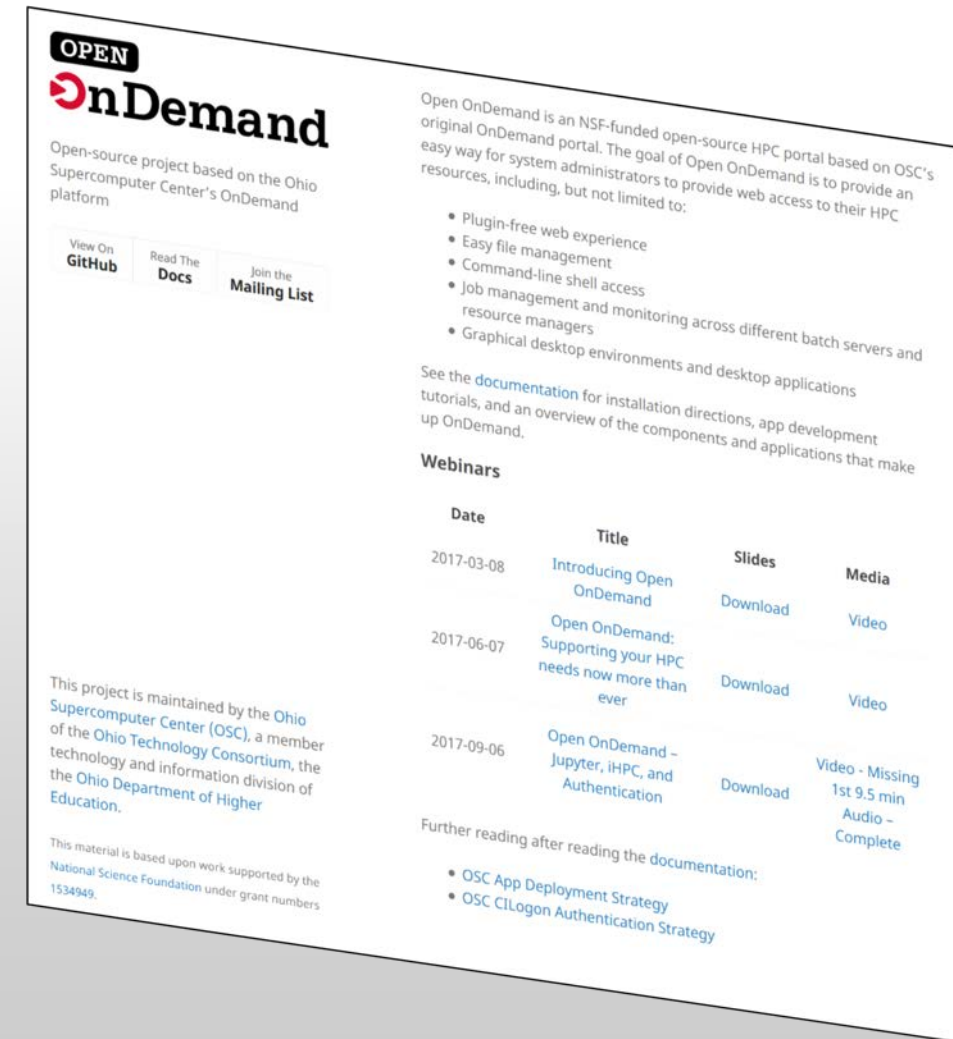


The screenshot shows a web browser window displaying the Discourse forum page for the 'ondemand2' tag. The browser's address bar shows the URL 'https://discourse.osc.edu/tags/ondemand2'. The page header includes the 'Ohio Supercomputer Center' logo and the text 'An OH:TECH Consortium Member'. The main content area is titled 'Tags > ondemand2' and features a 'New Topic' button. Below this, a table lists three topics, each with a biohazard icon. The first two topics are 'OnDemand 2.x Roadmap: Accessibility Goals' and 'OnDemand 2.x Roadmap: Scalability Goals', both marked as 'Open OnDemand' and having 0 replies, 2 and 3 views respectively, and being 1 hour old. The third topic is 'Should we transition Open OnDemand to a community foundation?', marked as 'Open OnDemand', with 1 reply, 52 views, and 12 days old. It also features a biohazard icon and a small profile picture icon.

Topic	Replies	Views	Activity
OnDemand 2.x Roadmap: Accessibility Goals ■ Open OnDemand ondemand2	0	2	1h
OnDemand 2.x Roadmap: Scalability Goals ■ Open OnDemand ondemand2	0	3	1h
★ Should we transition Open OnDemand to a community foundation? ■ Open OnDemand ondemand2	1	52	12d

# Staying in Touch

- Visit our website
  - <http://openondemand.org>
- Use our Discourse
  - <https://discourse.osc.edu/c/open-ondemand>
- Join our mailing list
  - <https://lists.osu.edu/mailman/listinfo/ood-users>
- Our webinars are planned roughly quarterly
  - Let us know what you'd like to learn about next



# Thank you! Any questions?



Alan Chalker, Ph.D.  
Director of Strategic Programs at OSC  
Ohio Supercomputer Center  
[alanc@osc.edu](mailto:alanc@osc.edu)

[openondemand.org](http://openondemand.org)