



nDemand

Alan Chalker, Morgon Rodgers, Steve Gallo, Eric Franz, Doug Johnson, Bob Settlage

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

User Group BoF Agenda



Ohio Supercomputer Center



1. **About Open OnDemand**
2. Open OnDemand 2.0 Project Roadmap
3. Key Items of Note
4. Open Floor Discussion



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

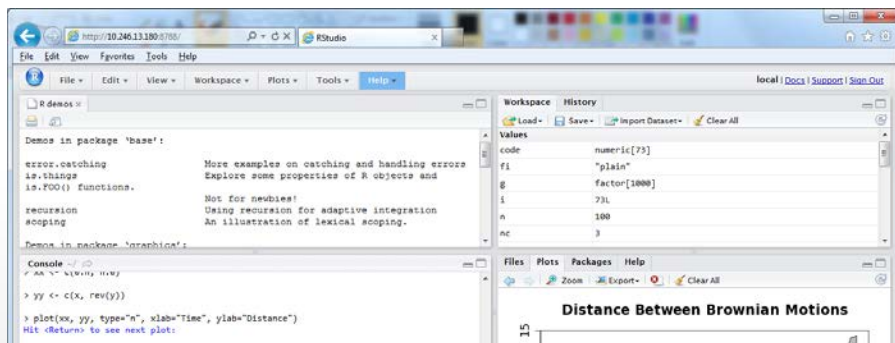


Interactive Apps

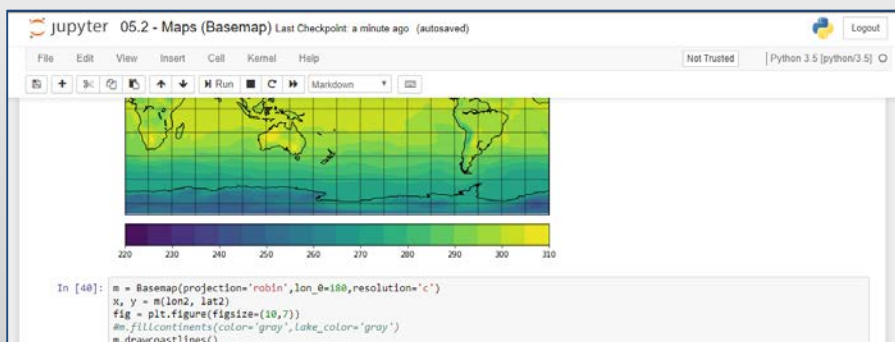
&

Cluster Access

RStudio Server – R IDE

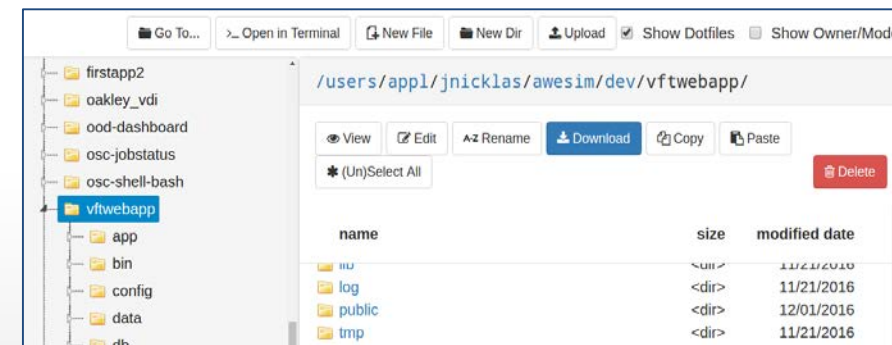


Jupyter Notebook – Python IDE



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

File Access (browse, edit, etc)



Manage Jobs (view, submit, etc)

Active Jobs									
Show 50 entries									
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	PAA0026	140:50:24	serial	Running	Owens		
> 3130446.owe...	RASPA_convert	osu1842	PAA0026	138:30:25	serial	Running	Owens		
> 3130447.owe...	RASPA_convert	osu1842	PAA0026	138:09:22	serial	Running	Owens		
> 3133647.owe...	high_jp_PIV_N_80_choke_wo_tm	osu9725	PAS1136	17:36:02	parallel	Running	Owens		
> 3137260.owe...	Case42	osu8290	PAA0008	96:36:34	longserial	Running	Owens		
> 3137265.owe...	Case195	osu8290	PAA0008	163:01:59	longserial	Running	Owens		
> 3137292.owe...	Case261	osu8290	PAA0008	165:44:57	longserial	Running	Owens		

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

Example Current Engagements and Deployments

Production Deployments



In Process of Installing



User Group BoF Agenda



Ohio Supercomputer Center



- ~~1. About Open OnDemand~~
- 2. Open OnDemand 2.0 Project Roadmap**
3. Key Items of Note
4. Open Floor Discussion



Open OnDemand 2.0 Project Overview

- OnDemand 1.x: 3 year NSF SI2 award (#1534949)
- OnDemand 2.x: 5 year NSF CSSI award (#1835725)
 - Collaborators: SUNY Buffalo, Virginia Tech
- Four areas
 - **Visibility:** Enhancing resource use visibility by integrating Open XDMoD
 - **Scalability:** support more types of computing resources and software
 - **Accessibility:** appeal to more scientists in more fields of science
 - **Engagement:** establish community of HPC users and administrators

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
 - DL/ML automated detection to optimize user efficiency


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: HPCoDB, File system storage logs. Powered by XDMoD/Highcharts



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

Scalability: 5 year goals (bolded year 1 focus)

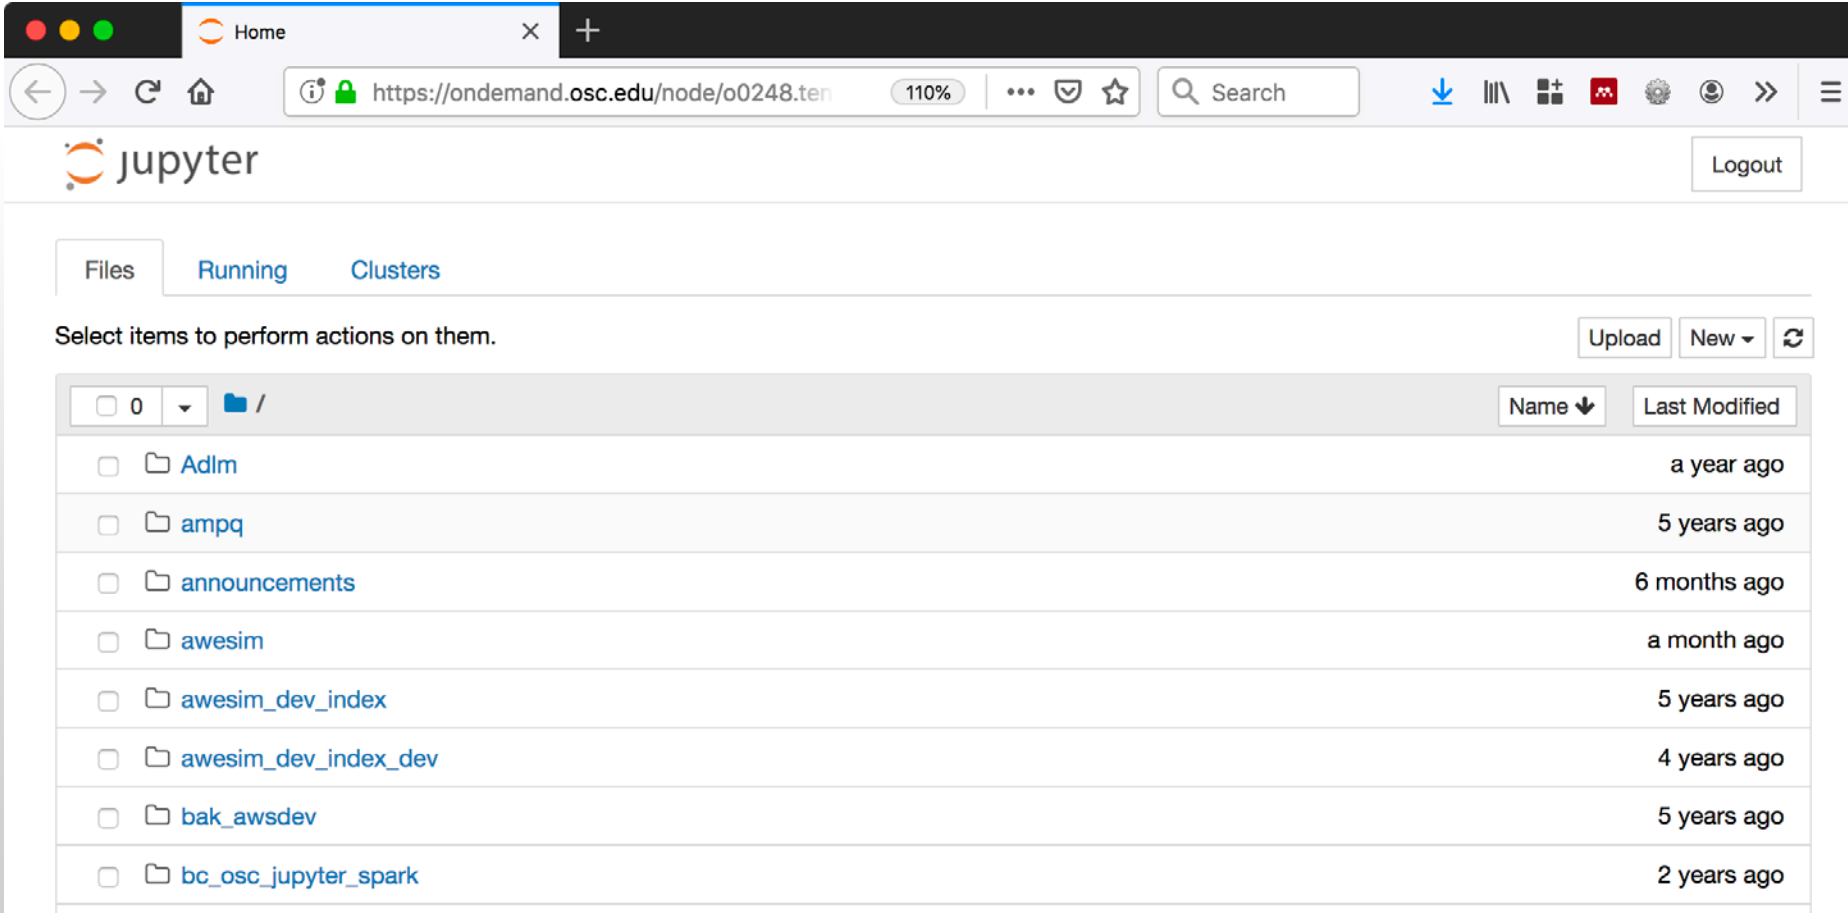
1. Extendibility

- **Interactive work without a batch scheduler**
- **Increase number of languages to build and deploy apps in**
- Support app development with REST API or ported App Kit

2. Performance

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

Scalability: Interactive without a batch scheduler - immediate access



Home

https://ondemand.osc.edu/node/o0248.ten

jupyter

Logout

Files Running Clusters

Select items to perform actions on them.

Upload New ↻

<input type="checkbox"/> 0	/	Name ↓	Last Modified
<input type="checkbox"/>	Adlm		a year ago
<input type="checkbox"/>	ampq		5 years ago
<input type="checkbox"/>	announcements		6 months ago
<input type="checkbox"/>	awesim		a month ago
<input type="checkbox"/>	awesim_dev_index		5 years ago
<input type="checkbox"/>	awesim_dev_index_dev		4 years ago
<input type="checkbox"/>	bak_awsdev		5 years ago
<input type="checkbox"/>	bc_osc_jupyter_spark		2 years ago



Accessibility: 5 year goals (bolded year 1 focus)

- 1. Improve Job Management**
- 2. Reduce Administrative Load (installation, config., etc)**
- 3. Streamlining interface (reduce steps to accomplish a task)**
4. Integrate Job/File Management, and App Launching interface
5. Integrated IDE for job and app development

Complete Software Catalog and Integrated App Plugin

Demand  My Interactive Sessions



Ohio Supercomputer Center

An **OH-TECH** Consortium Member

OnDemand provides an integrated, single access point for all of your I

Filters:

All

Servers


GUIs



ANSYS Workbench



Abaqus/CAE



COMSOL Multiphysics



Jupyter



Jupyter (Pitzer)



Jupyter + Spark



MATLAB

Launch MATLAB in interactive or batch mode.

Input File

Browse...

No file selected.

or

Choose system file

☐ Interactive - access interactive MATLAB GUI through VNC Session

Advanced Options: Environment

Advanced Options: Platform

Submit

OpenFOAM

Submit OpenFOAM serial or parallel job.

Case Directory Path or Tar

Browse...

No file selected.

or

Choose system file

Advanced Options: Environment

Advanced Options: Platform

Submit

Engagement: 5 year goals

1. Targeting non traditional HPC disciplines
2. Advocating for the beginner user
3. Outreach
4. Ensure the project is community guided (Advisory Group)
 - Members create local focus groups of users to gather feedback
 - Help identify apps to expand access to non traditional domains

Leveraging OnDemand in Gateway Opportunities (LOGO)

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

User Group BoF Agenda



Ohio Supercomputer Center

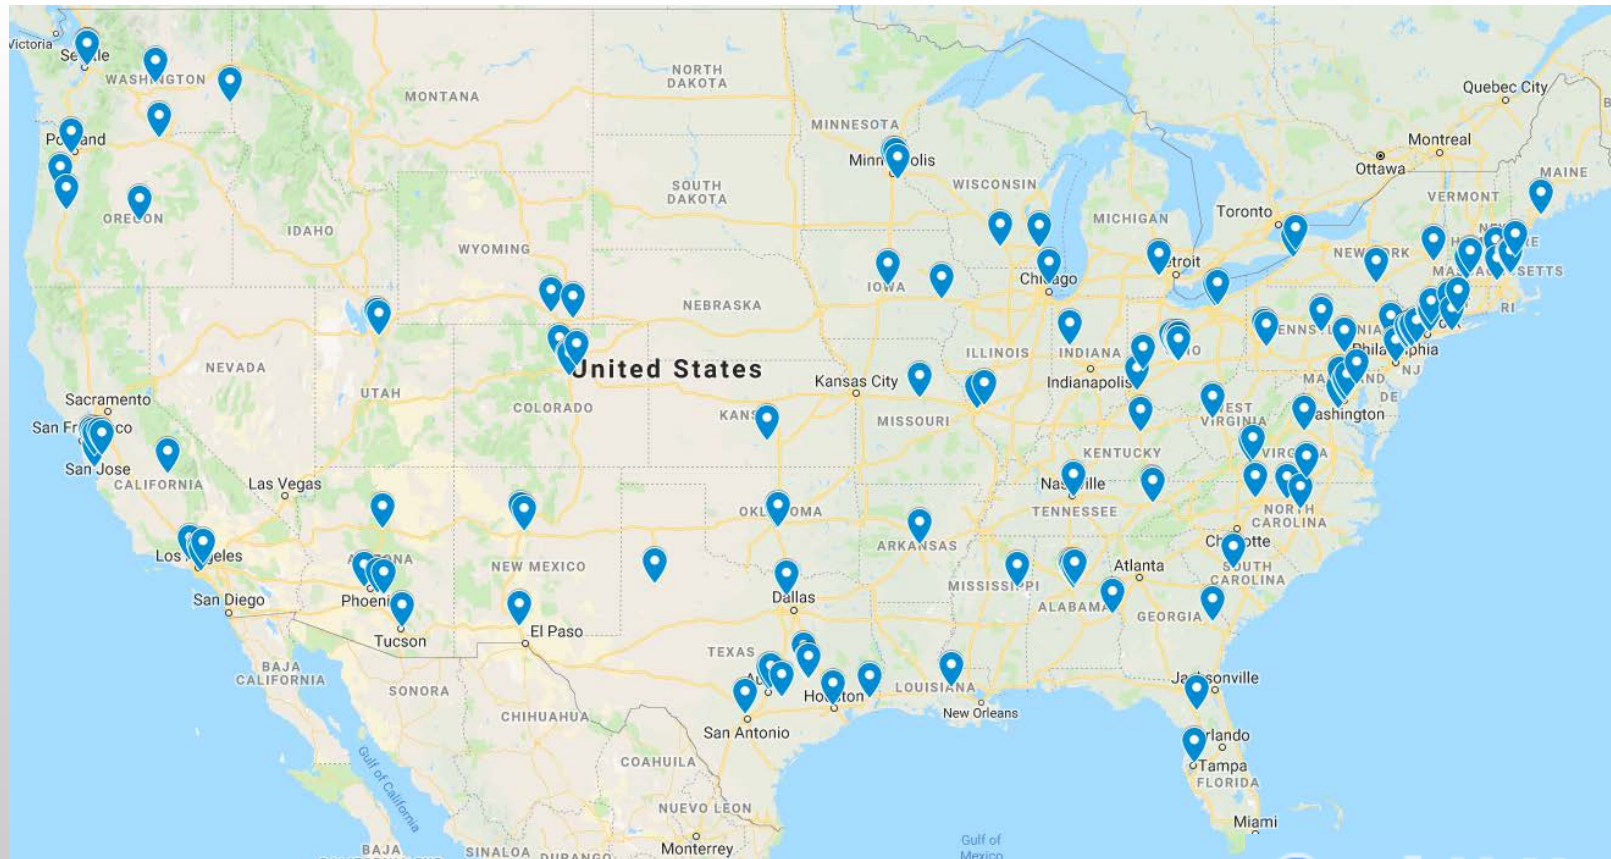


- ~~1. About Open OnDemand~~
- ~~2. Open OnDemand 2.0 Project Roadmap~~
- 3. Key Items of Note**
4. Open Floor Discussion



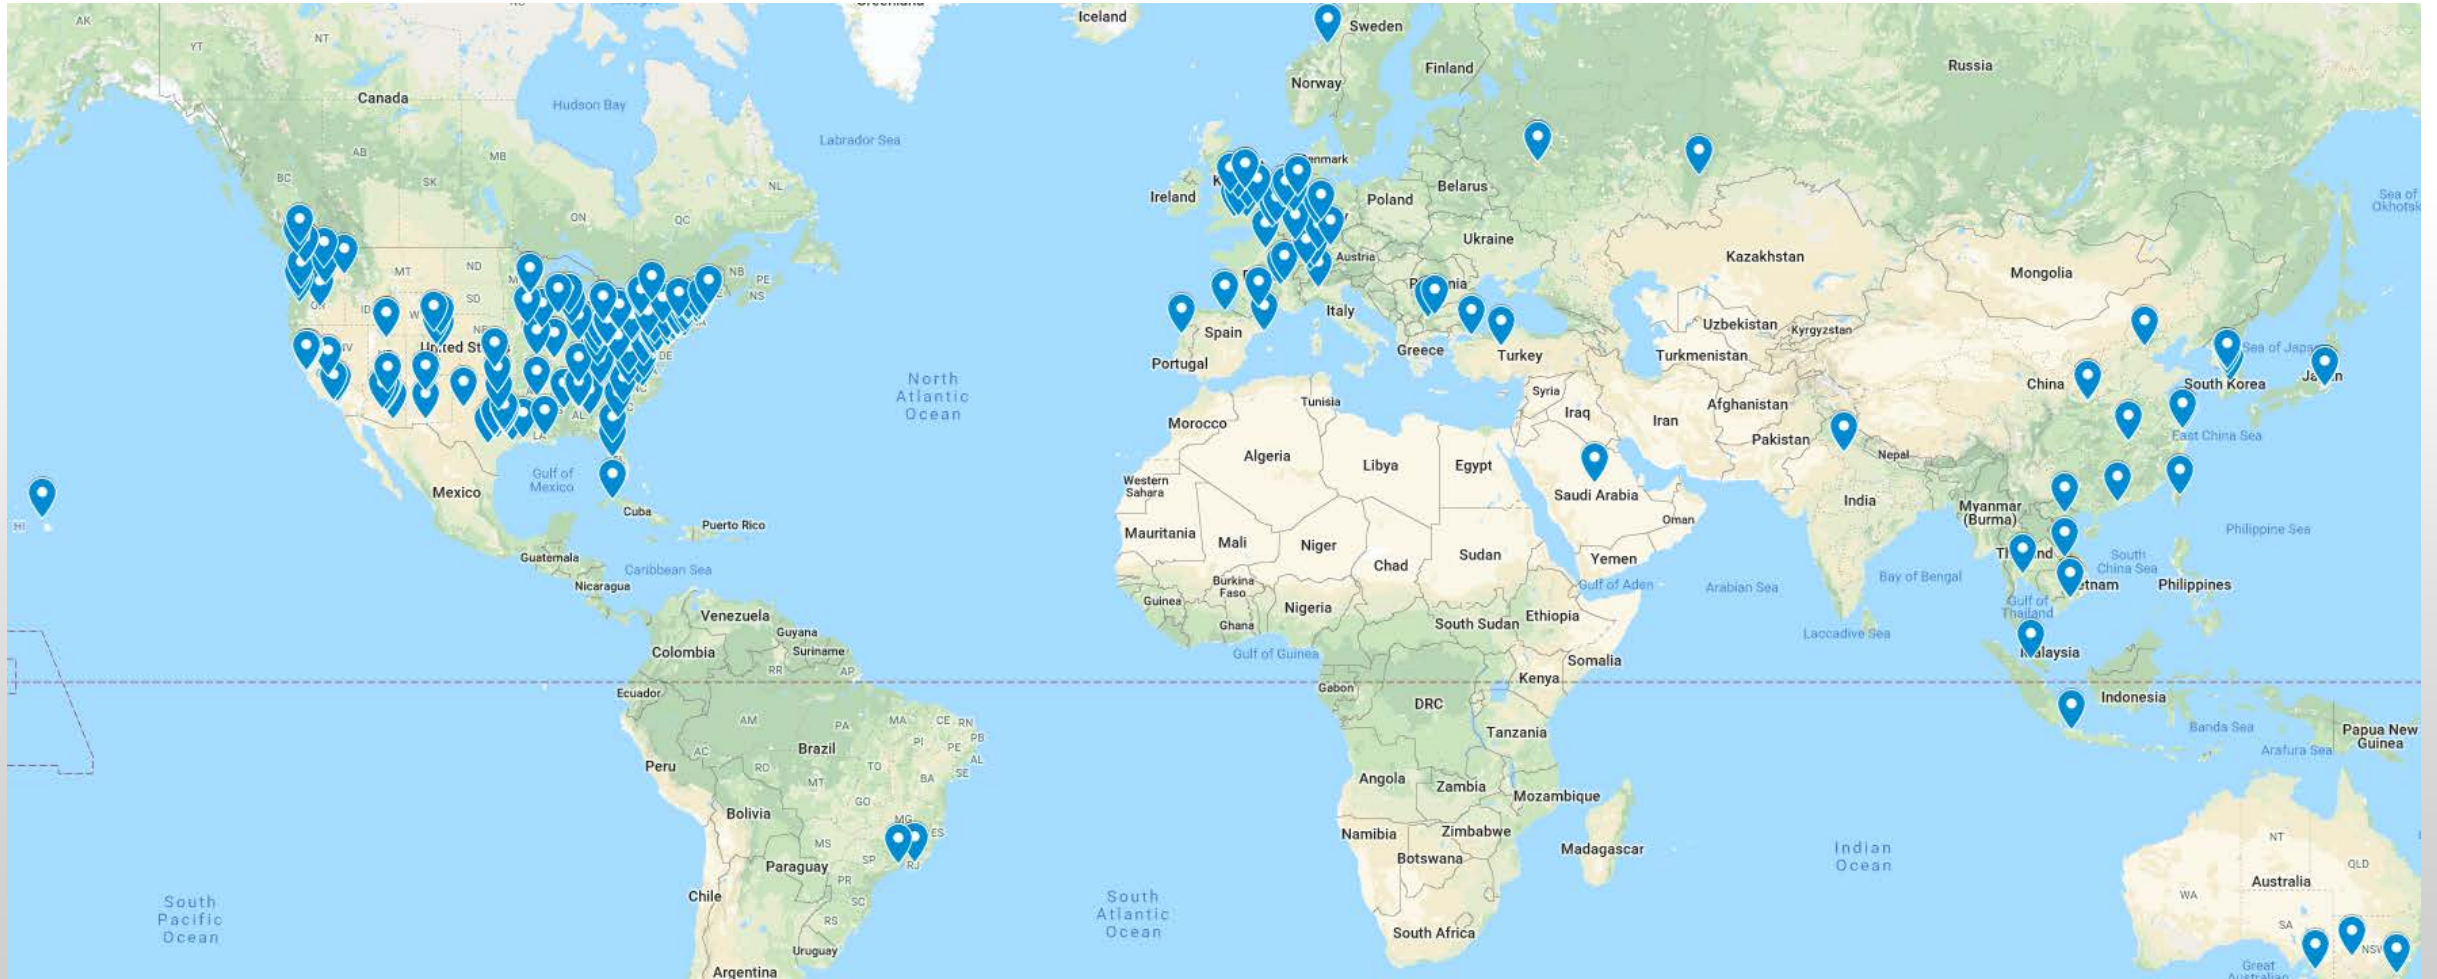
Approx Number of Institutions based on RPM logs

- 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

User Survey Results

- Questions included:
 - Usefulness, design, functionality ratings
 - How would you change the design, functionality
- 28 responses (1 best, 5 worst)
 - Overall OOD usefulness: 1.73 mean
 - OOD Interface design efficacy: 2.31 mean
 - OOD Interface functionality: 2.08 mean
- Conclusion: OOD is useful but there is room for improvement

Nor-Tech



AN OPEN DOOR TO OPEN SOURCE

NOW EVEN HPC NEOPHYTES CAN TAKE ADVANTAGE OF THE ENORMOUS COST-SAVINGS OF OPEN SOURCE.

Open OnDemand makes Open Source HPC as easy as using a desktop or website.

We are the only enterprise offering Open OnDemand and the only company able to integrate it into OpenHPC.

BENEFITS INCLUDE:

- Keeps initial and recurring licensing costs way down
- No ongoing licensing fees for cluster management software
- Much easier to use than command lines

SERVICES AVAILABLE

Nor-Tech has been a primary innovator in the HPC cluster space for more than a decade. Our elite staff of expert engineers, averaging more than a decade of hands-on experience, are consistently conquering industry challenges that our competitors long abandoned as insurmountable. The result is a suite of cost-effective NT-EZ branded services that remove the classic obstacles to high performance computing

We are continually evaluating additional areas where we can innovate in order to bring more value to our clients

The MIT License (MIT): Copyright © 2016 Ohio Supercomputer Center.
The software is provided "as is" without warranty of any kind, express or implied, including but not limited to the warranties of merchantability, fitness for a particular purpose and non-infringement, in no event shall the authors or copyright holders be liable for any claim, damages or other liability, whether in an action of contract, tort or otherwise, arising from, out of or in connection with the software or the use or other dealings in the software.

Nor-Tech's NT-EZ Suite

•**NT-EZ Remote Visualization:** Storage access too slow? With NT-EZ Remote Visualization, big data files don't have to be transferred to the user's location to be viewed for post processing. Only keyboard, mouse and screen bits travel to the user location.

•**NT-EZ Storage Guard:** This monitors cluster storage usage and alerts the user if available storage is reaching critically low levels. This keeps clusters stable and reliable. Without this feature, the client will run out of storage without warning causing the cluster job to crash.

•**NT-EZ SATM (System Ambient Temperature Monitor):** This prevents cluster failure due to excessive heat. SATM emails alerts to the client warning that the server room ambient temperature is too high. When the ambient temperature climbs above a preset threshold the cluster.

•**NT-EZ RCR Rapid Cluster Recovery/Bare Metal Backup:** Run manual cluster backups or schedule backups seamlessly. This not only backs up the cluster data, but also backs up an image of both the data and the OS. In the event of a disaster the whole cluster software and applications environment can easily be recovered. This adds a level of reliability to our clusters

•**NT-EZ Remote Monitoring & Management:** This provides an easy way for our clients to use and manage complex equipment such as a High Performance Compute Cluster without HPC expertise in-house. This service is provided on either a scheduled basis or an as needed basis. The scheduled service allows us to find and fix small problems before they become disasters. This service is sold in blocks of hours ahead of when the service may be needed.

•**NT-EZ Open OnDemand® Plus:** This delivers a GUI through a browser instead of only a CLI. Use NT-EZ Open OnDemand Plus instead of cryptic file editors and Linux utilities, which are not user friendly. Nor-Tech adapted Open OnDemand Plus to work with Open HPC our most prevalent Open Source cluster management suite.



User Group BoF Agenda



Ohio Supercomputer Center



- ~~1. About Open OnDemand~~
- ~~2. Open OnDemand 2.0 Project Roadmap~~
- ~~3. Key Items of Note~~
- 4. Open Floor Discussion**



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

