



# Power plus Open OnDemand: facilitating the transition from new to Power user

Robert Settlege  
Advanced Research  
Computing @Virginia Tech  
Oct 2020

**OPEN**  **nDemand**

 University at Buffalo  
Center for Computational Research



Ohio Supercomputer Center

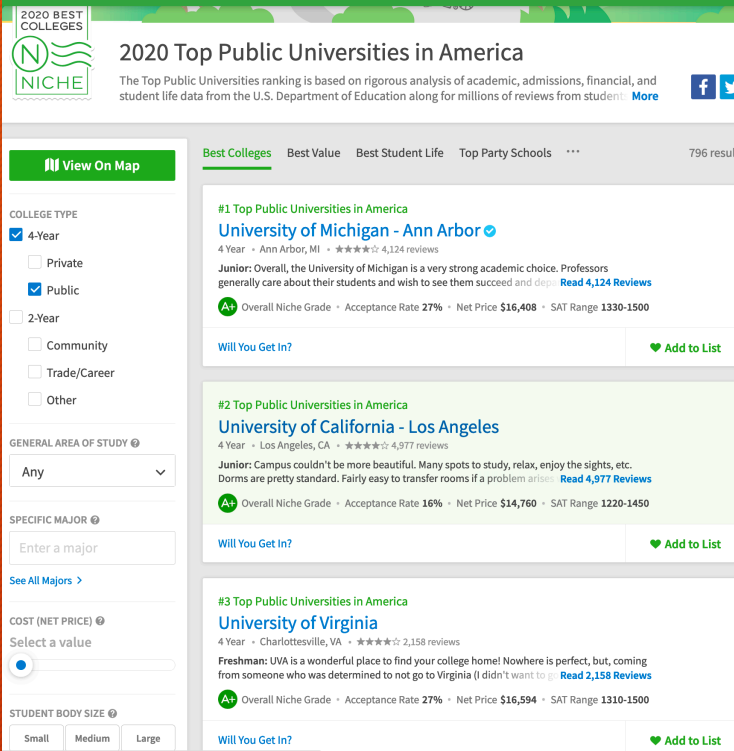
 **VIRGINIA  
TECH™**

It's (almost) that time of year...



# Who am I

- Researcher
- Collaborator
- Educator
- Nvidia DLI Ambassador (CUDA, CV)
- Chemist (PhD, University of Virginia)
- Statistician (MS, Virginia Tech)
- Computational Scientist
- UVA fan



**2020 BEST COLLEGES**  
NICHE

## 2020 Top Public Universities in America

The Top Public Universities ranking is based on rigorous analysis of academic, admissions, financial, and student life data from the U.S. Department of Education along for millions of reviews from students. [More](#)

796 results

**View On Map**

**COLLEGE TYPE**

- 4-Year
- Private
- Public
- 2-Year
- Community
- Trade/Career
- Other

**GENERAL AREA OF STUDY**

Any

**SPECIFIC MAJOR**

Enter a major

[See All Majors >](#)

**COST (NET PRICE)**

Select a value

**STUDENT BODY SIZE**

Small Medium Large

---

**#1 Top Public Universities in America**  
**University of Michigan - Ann Arbor**

4 Year · Ann Arbor, MI · ★★★★★ 4,124 reviews

**Junior:** Overall, the University of Michigan is a very strong academic choice. Professors generally care about their students and wish to see them succeed and dep... [Read 4,124 Reviews](#)

Overall Niche Grade · Acceptance Rate **27%** · Net Price **\$16,408** · SAT Range **1330-1500**

[Will You Get In?](#) [Add to List](#)

---

**#2 Top Public Universities in America**  
**University of California - Los Angeles**

4 Year · Los Angeles, CA · ★★★★★ 4,977 reviews

**Junior:** Campus couldn't be more beautiful. Many spots to study, relax, enjoy the sights, etc. Dorms are pretty standard. Fairly easy to transfer rooms if a problem arises. [Read 4,977 Reviews](#)

Overall Niche Grade · Acceptance Rate **16%** · Net Price **\$14,760** · SAT Range **1220-1450**

[Will You Get In?](#) [Add to List](#)

---

**#3 Top Public Universities in America**  
**University of Virginia**

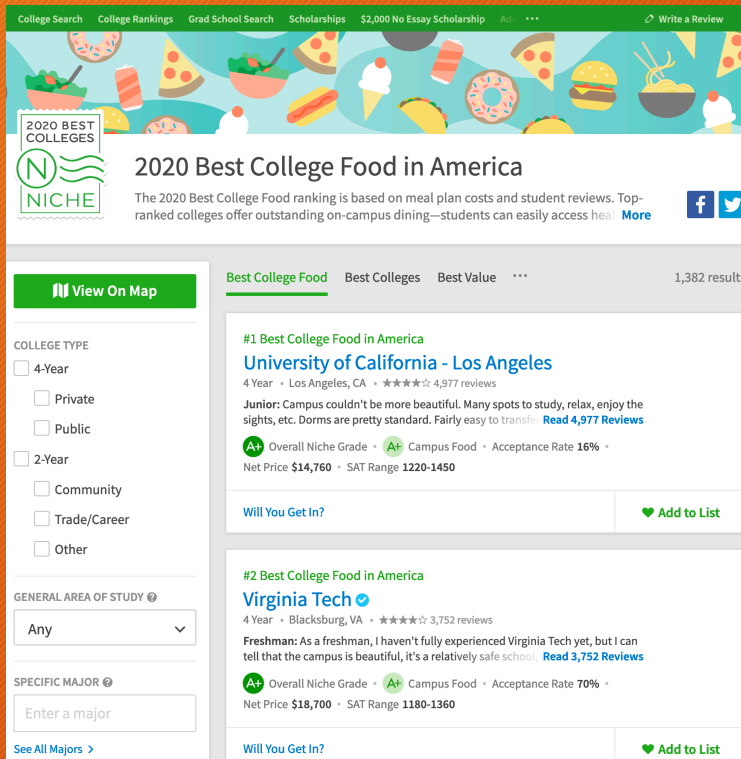
4 Year · Charlottesville, VA · ★★★★★ 2,158 reviews

**Freshman:** UVA is a wonderful place to find your college home! Nowhere is perfect, but, coming from someone who was determined to not go to Virginia (I didn't want to go... [Read 2,158 Reviews](#)

Overall Niche Grade · Acceptance Rate **27%** · Net Price **\$16,594** · SAT Range **1310-1500**

[Will You Get In?](#) [Add to List](#)

# We are Virginia Tech



2020 Best College Food in America

The 2020 Best College Food ranking is based on meal plan costs and student reviews. Top-ranked colleges offer outstanding on-campus dining—students can easily access health...

**#1 Best College Food in America**  
**University of California - Los Angeles**  
 4 Year • Los Angeles, CA • ★★★★★ 4,977 reviews  
 Junior: Campus couldn't be more beautiful. Many spots to study, relax, enjoy the sights, etc. Dorms are pretty standard. Fairly easy to transfer. [Read 4,977 Reviews](#)  
 Overall Niche Grade: A+ • Campus Food: A+ • Acceptance Rate: 16% • Net Price: \$14,760 • SAT Range: 1220-1450

**#2 Best College Food in America**  
**Virginia Tech**  
 4 Year • Blacksburg, VA • ★★★★★ 3,752 reviews  
 Freshman: As a freshman, I haven't fully experienced Virginia Tech yet, but I can tell that the campus is beautiful, it's a relatively safe school. [Read 3,752 Reviews](#)  
 Overall Niche Grade: A+ • Campus Food: A+ • Acceptance Rate: 70% • Net Price: \$18,700 • SAT Range: 1180-1360

## The only school to:

- ✓ Sequence and assemble the genome/transcriptome of our mascot
- ✓ Serve our mascot on Saturdays
- ✓ And save said mascot in November

Symposium  
 Applying Next-Generation Sequencing to Solve Poultry Problems

### Next-generation sequencing strategies for characterizing the turkey genome <sup>1</sup>

Rami A. Dalloul \*✉, Aleksey V. Zimin †, Robert E. Settlage †, Sungwon Kim \*, Kent M. Reed §

[Show more](#)

<https://doi.org/10.3382/ps.2013-03560> [Get rights and content](#)



### Presidential Turkeys return to roost at Virginia Tech

For the fourth year, the turkeys will travel from the White House to Gobblers Rest. Meet the newest Hokies at a Dec. 1 open house.  
 November 22, 2019




And well fed

# Agenda

- VT Advanced Research Computing
- The User Spectrum
- Enabling Users with Open OnDemand
- Extending OnDemand to Power
- Providing Feedback on Performance to Improve Resource Utilization (WIP)

# VT-Advanced Research Computing (ARC)

- Unit within the Office of the Vice President of Information Technology
- Provide centralized resources for:
  - Research computing
  - Visualization
- Staff to assist users
- Website: <http://www.arc.vt.edu>



# ARC - what we do

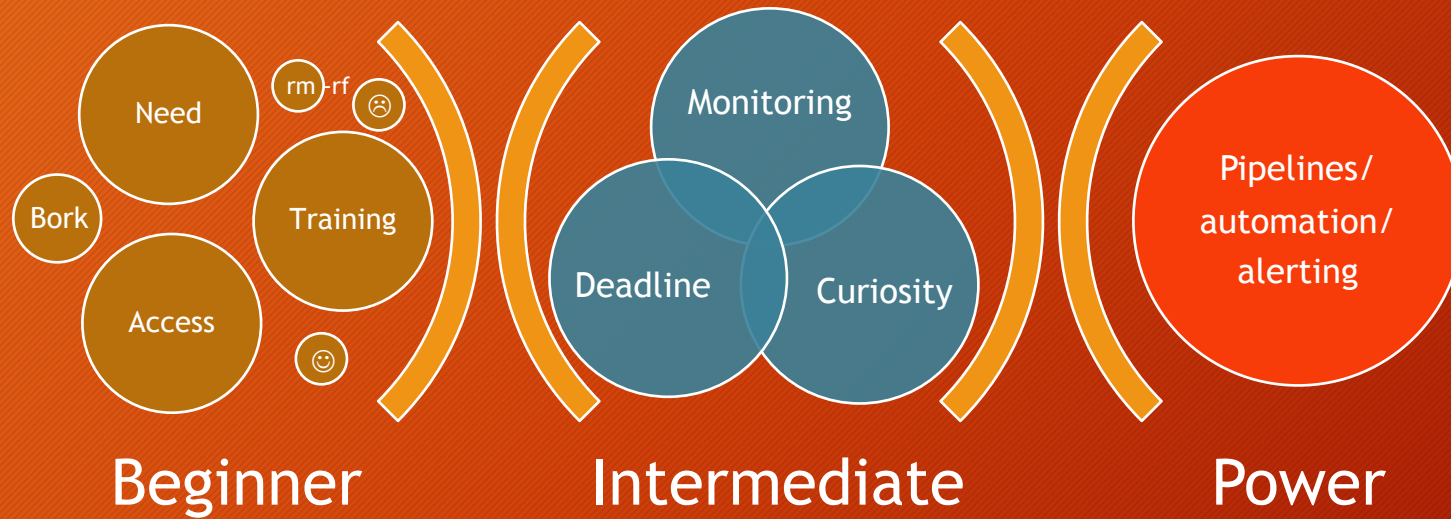
- Advance the use of computing and visualization in VT research
- Centralize resource acquisition, maintenance, and support for research community
- Provide support to facilitate usage of resources and minimize barriers to entry
- Enable and participate in research collaborations between departments
- Provide resources for instruction

# Student = user (research powerhouse)

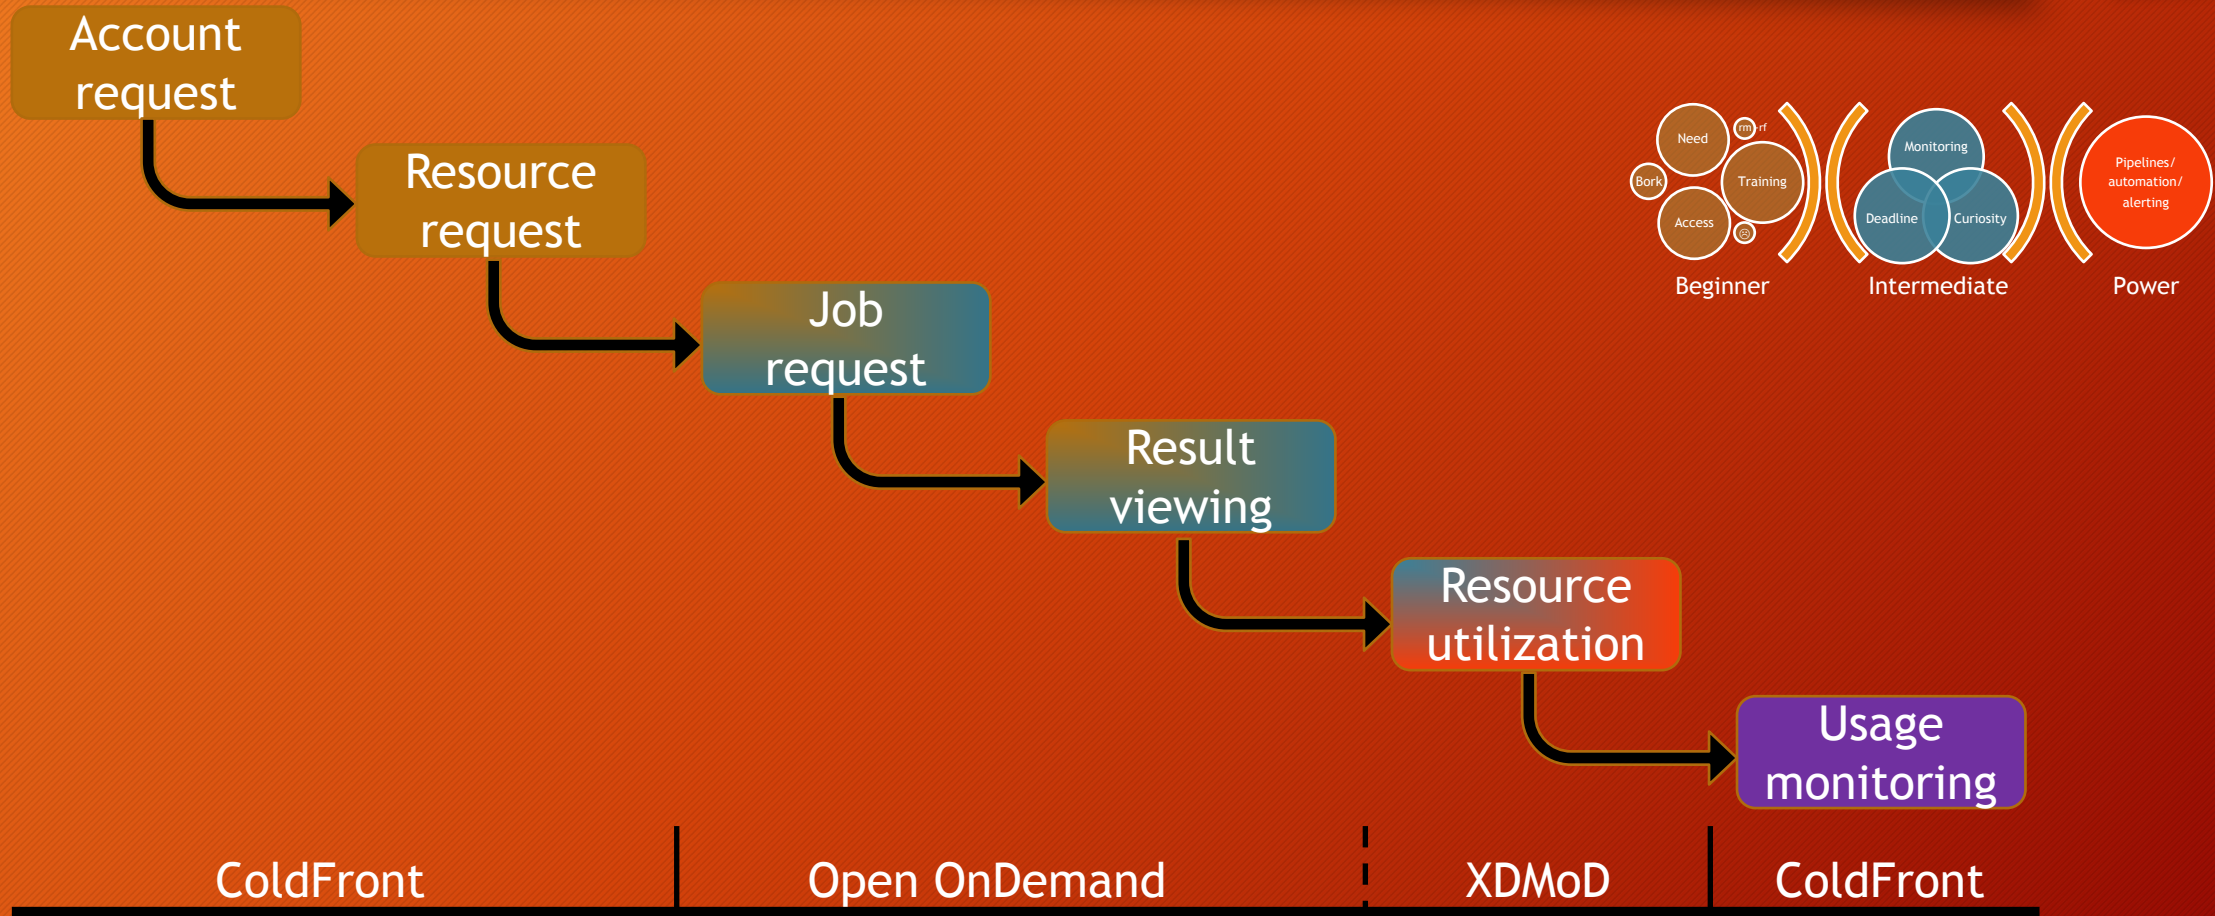
- Odd form of life that tries random stuff to meet random goals while accidentally learning



# User spectrum



# GOAL: Facilitate and speed the progression



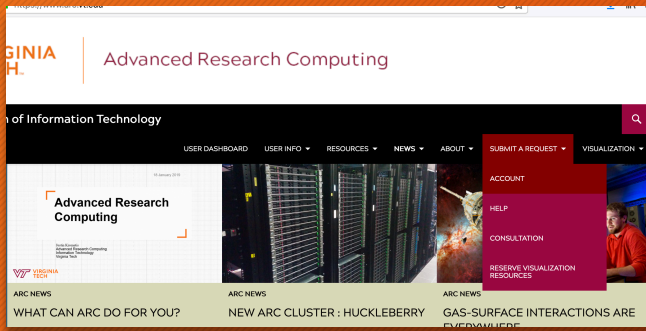
# Getting started on ARC

## What do you need?

- An account
- An allocation

## Coldfront.arc.vt.edu

1. Create a project
2. Create an allocation



Project	Name	Resource	Status
Systems Biology 3036 Teaching	SYSB3036	Compute (Cluster)	Active
STAT3615 Spring 2020	Irjohn-STAT3615-sp20	Compute (Cluster)	Active
Bovine DNA methylation	bDNA	Compute (Cluster)	Active
ARC-training	arc-train4	Compute (Cluster)	Active
root	root	Compute (Cluster)	Active

## Open OnDemand 2.0

**Principal Investigator:** Robert Settlage (rsettlag) ✉

**Description:** Open OnDemand setup and configuration

**Field of Science:** Advanced Research Computing

**Project Status:** Active

**Created:** Sep. 04, 2019

Username	Name	Email	Role	Status	Enable Notifications	Actions
bbooker	Brandon Booker	bbooker@vt.edu	User	Active	<input checked="" type="checkbox"/>	
brownm12	Matthew Brown	brownm12@vt.edu	User	Active	<input checked="" type="checkbox"/>	
bsawyers	Brandon Sawyers	bsawyers@vt.edu	User	Active	<input checked="" type="checkbox"/>	
efranz	Suppressed Suppressed	efranz@osc.edu	User	Active	<input checked="" type="checkbox"/>	
jdstrick	Justin Strickland	jdstrick@vt.edu	User	Active	<input checked="" type="checkbox"/>	
jkrometi	Justin Krometis	jkrometis@vt.edu	User	Active	<input checked="" type="checkbox"/>	
rsettlag	Robert Settlage	rsettlag@vt.edu	Manager	Active	<input checked="" type="checkbox"/>	

Resource Name	Resource Type	Information	Status	End Date	Actions
Compute	Cluster	slurm_account_name: ondemand2	Active	May. 22, 2020	

## ARC.vt.edu

- Requests
- account request

# Enabling Users with Open OnDemand



**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

### 2020-03-16 - 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

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

**Message of the Day**

This system is for authorized users only. Users accessing this system consent to the monitoring, recording and/or disclosure of all activity while using this system. Usage of this system is subject to the terms of the [Virginia Tech Acceptable Use Guidelines](#).

Diskspace core and node availability at 18 October, 2020 - 09:30:21PM

	dev_q	normal_q	preemptible_q	largemem_q	h30_q	v100_normal_q	v100_dev_q
total cores/node	32	32	32	72	32	24	24
total gpu/node	-	-	-	-	4	2	2
nodes in partition	190	188	190	2	4	38	40
reserved/down nodes	16	16	16	0	2	0	0
cores available	257	225	257	27	16	204	382
gpu available	-	-	-	-	8	16	20
1/2 nodes available	9	8	9	0	0	15	15
full nodes available	2	1	2	0	0	0	2
Notification covers/nodes	95,439.8	93,929.4	95,439.8	81,237.0	79,700	65,473.0	60,279

Diskspace core and node availability at 18 October, 2020 - 09:30:21PM

	dev_q	normal_q	preemptible_q
total cores/node	24	24	24
total gpu/node	-	-	-
nodes in partition	47	42	47
reserved/down nodes	0	0	0
cores available	209	90	209
1/2 nodes available	8	3	8
full nodes available	5	1	5
Notification covers/nodes	81,539.4	81,107.8	81,539.4

**NEW USERS:** Please run this script in a terminal window (under Clusters menu) before submitting jobs to the cluster: `./util/ccr/bin/ssh_no_password.sh`  
**MACHINE STATUS:** Academic cluster status and Industry cluster status or use 'sqsstat' command in terminal window.  
**NEXT DOWNTIME:** Tuesday, October 27, 2020 [More details](#)  
**SUMMER SYSTEM UPDATES:** We've been busy! Check out all the changes - many directly affect how you run jobs!  
**VIRTUAL WORKSHOPS:** Check out our library of virtual workshops [More info here](#)  
**FOLLOW US!** CCR is now on Twitter! Get system status updates, helpful hints, & highlights of interesting research done at CCR

**University at Buffalo**  
**Center for Computational Research**

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

## Message of the Day

Welcome to the Center for Computational Research!

You are accessing a University at Buffalo (UB) - Center for Computational Research (CCR) computer system that is provided for CCR-authorized users only. By using this system (which consists of any device attached to this machine, including compute nodes, remote visualization software and hardware, storage and database resources), you have implicitly agreed to abide by the highest standards of responsibility to your colleagues - the students, faculty, staff, and external users who share this environment. You are required to comply with ALL University at Buffalo policies (<http://www.buffalo.edu/ubit/policies/it-policies-a-to-z/computer-and-network-use.html>), as well as state and federal laws concerning appropriate use of information technology.

- CCR is not responsible for the loss or misuse of data on our systems.
- CCR systems are NOT HIPAA-compliant. Storage of any personally identifiable Protected Health Information (PHI) on our systems is a violation of the Health Insurance Portability and Accountability Act (HIPAA) of 1996 Privacy and Security Rules. If in doubt, contact CCR before transferring your data.
- All CCR systems are monitored for administrative and security reasons. Use of this system constitutes consent to this monitoring for these purposes.

By continuing to use this system you indicate your awareness of and consent to these terms and conditions of use. Non-compliance of these terms is considered a breach of University policy and may result in disciplinary and/or legal action.

**Jobs Efficiency Report - 2020-09-16 XDMoD**

92.3% efficient 7.7% inefficient

1635 inefficient jobs/21184 total jobs

**Core Hours Efficiency Report - 2020-XDMoD**

42.7% efficient 57.3% inefficient

937234.6 inefficient core hours/1636543.2 total core hours

**Recently Completed Jobs - 2020-09-XDMoD**

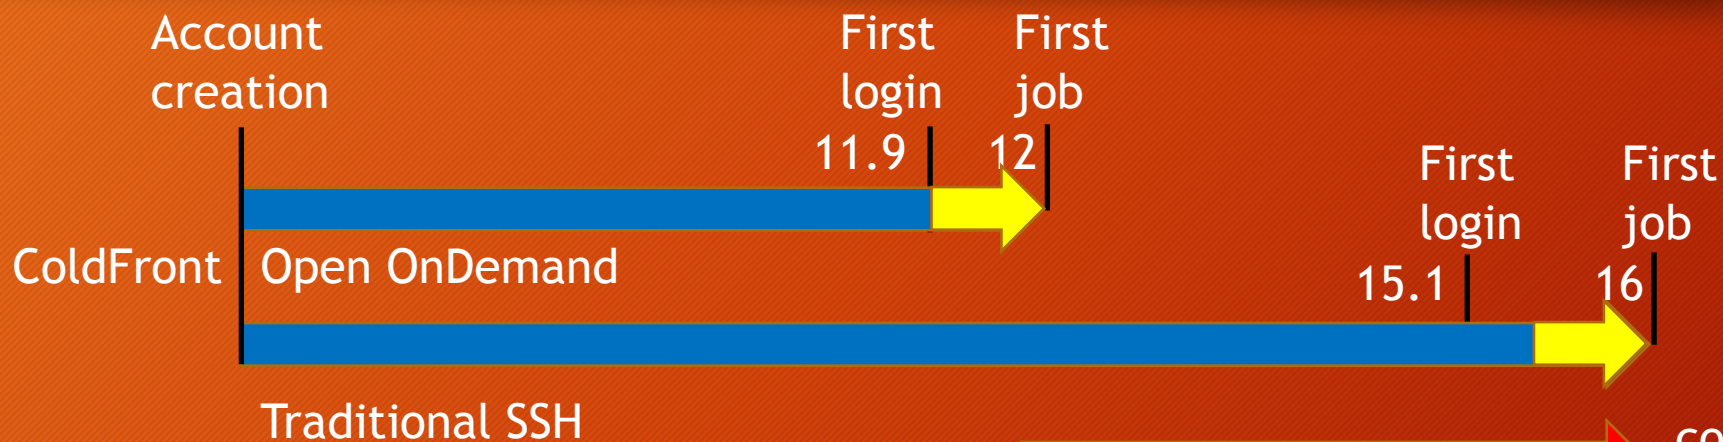
ID	Name	Date	CPU
4274244	SimRunArray_12	10/5	35.5
4270864	slurm_s64	10/5	98.3
4274168	SimRunArray_12	10/5	14.0
4271250	slurm_e64	10/5	99.9

- Browser based HPC access
- Zero user install
- Site customizable

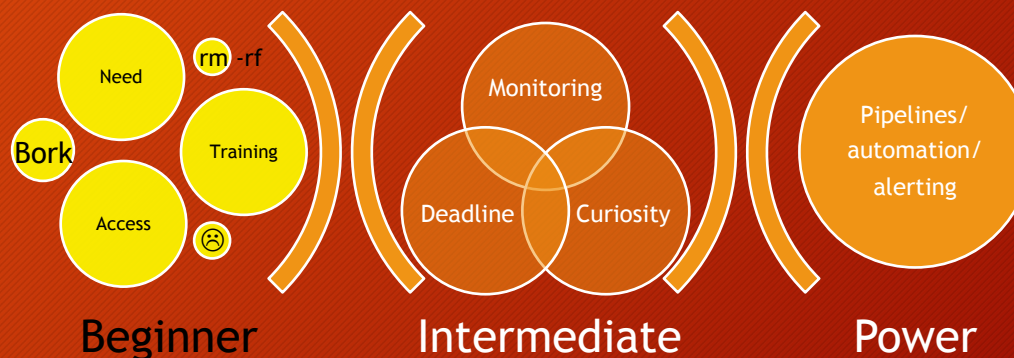
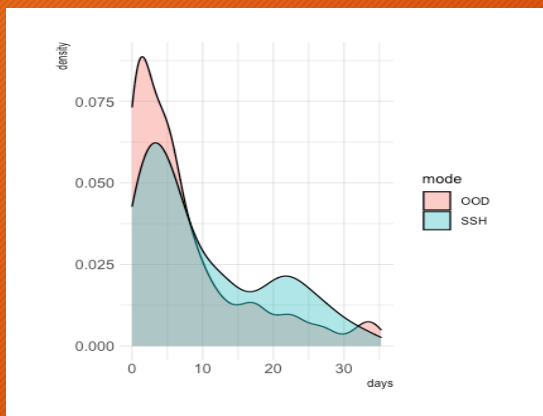
# Facilitate and speed the progression

76

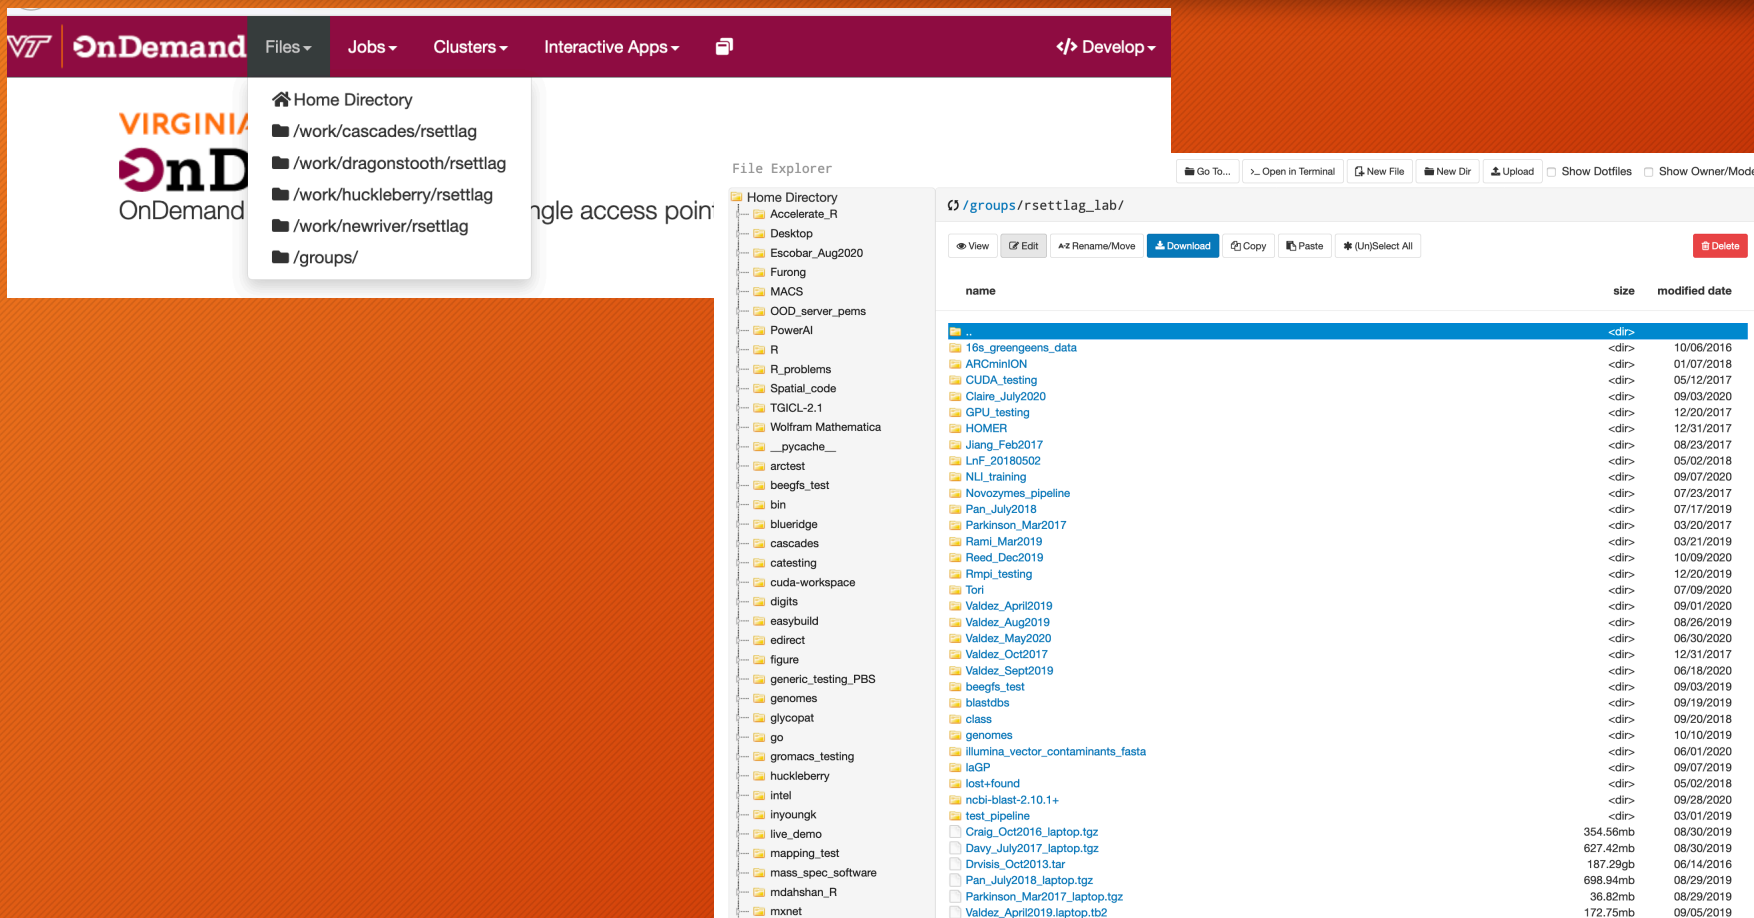
Less hours  
To first job



competitive advantage



# Open OnDemand: file management



- up/download
- Move/copy/delete
- Edit
- Soon Globus integration

# Open OnDemand: SSH access

OnDemand provides an integrated, single access point

Acceptable Use Guidelines (<http://www.policies.vt.edu/acceptableuse.php>)

NOTE: VT Enterprise Directory Password authentication requires a DUO second factor challenge. After your password is provided, you will receive a DUO challenge.

Last login: Fri Oct 16 18:00:38 2020 from ondemand-preprod.arc-systems.nis.dit.isb-1.opc.vt.edu

Information on and examples of how to use huckleberry and other ARC systems, plus forms for requesting accounts or submitting help tickets, are available at <http://www.arc.vt.edu>.

data usage:

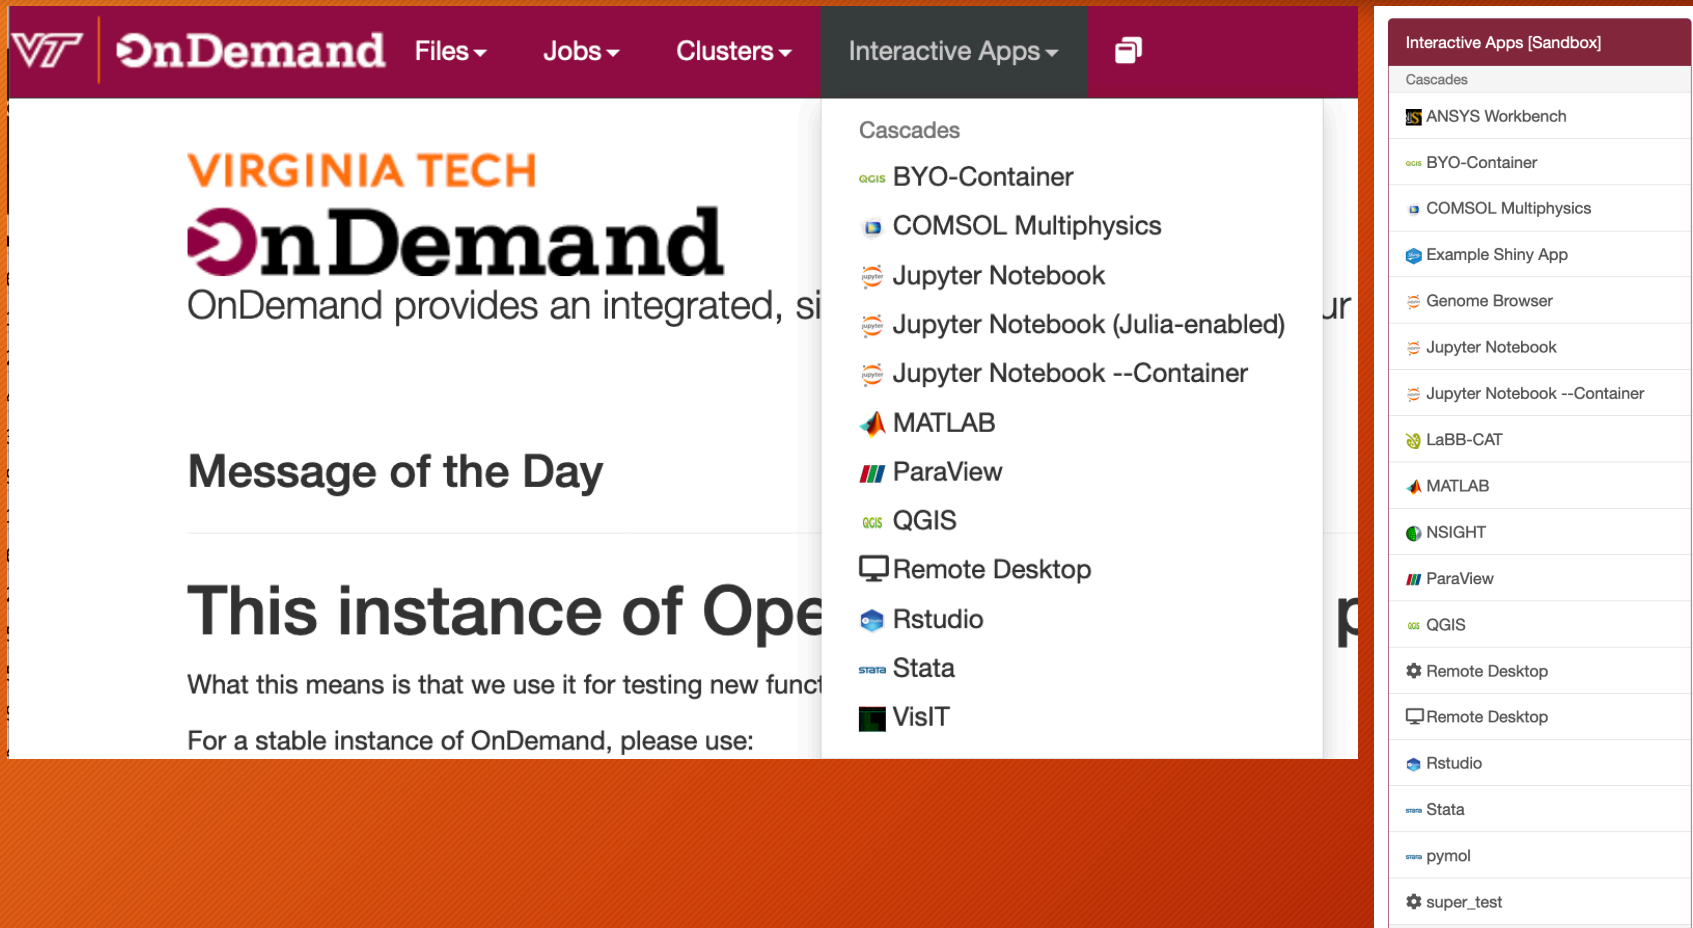
USER	FILESYS/SET	DATA (GiB)	QUOTA (GiB)	FILES	QUOTA	NOTE
rsettlag	/home	355.4	640	-	-	
rsettlag	/work	318.0	20480	495877	6291456	
rsettlag	/groups/arcadm/	101.0	20480	147912	10485760	
rsettlag	/groups/dea/	229236.0	256000	32430041	104857600	
rsettlag	/groups/rsettlag_lab/	6161.0	10240	168990	5242880	
rsettlag	/groups/dslade_lab/	4256.0	10240	24799	5242880	
rsettlag	/groups/songbird_transcriptome/	2024.0	10240	981446	5242880	
rsettlag	/groups/dalloul_lab/	20119.0	40960	4	10485760	
rsettlag	/groups/predictHPC/	0.0	10240	204	5242880	
rsettlag	/groups/DogsRule/	0.0	10240	9	5242880	
rsettlag	/groups/shorebirds/	55.0	10240	2228	5242880	
rsettlag	/groups/mohsan88_lab/	171.0	10240	1780	5242880	
rsettlag	/groups/michalak_lab/	7696.0	10240	2663467	5242880	
rsettlag	/groups/Lahmers_lab/	5924.0	10240	2043980	10485760	
rsettlag	/groups/DataLightHouse/	311.0	10240	6353	5242880	
rsettlag	/groups/jaketu_lab/	27598.0	30720	2981335	10485760	
rsettlag	/groups/inyoungk_lab/	0.0	10240	579	10485760	
rsettlag	/groups/hojiang_lab/	920.0	10240	83018	10485760	
rsettlag	/groups/arcsingularity/	101.0	20480	147912	10485760	
rsettlag	/groups/lrjohn-STAT3615-sp20/	0.0	10240	2	10485760	

USER	ALLOCATION	CLUSTER	QUOTA (hrs)	LEFT (hrs)	EXPIRES	NOTE
rsettlag	valdez_seq	all	Infinity	Infinity	2021-02-27	
rsettlag	ondemand	all	Infinity	Infinity	2021-10-16	
rsettlag	stat5014-fall20	all	Infinity	Infinity	2020-12-30	
rsettlag	arctest	all	Infinity	Infinity	2021-09-15	
rsettlag	mabrownlab	all	Infinity	Infinity	2021-10-16	
rsettlag	ondemand2	all	Infinity	Infinity	2021-10-16	
rsettlag	IGAL2018	all	Infinity	Infinity	2021-06-05	
rsettlag	ARC-train2	all	Infinity	Infinity	2021-10-15	
rsettlag	MinION_seq	all	Infinity	Infinity	2021-08-11	
rsettlag	valdez_seq_2	all	Infinity	Infinity	2021-02-27	
rsettlag	aylwardlab	all	Infinity	Infinity	2021-09-23	

- OOD knows who you are
- Authentication is done via Browser
- SSH window doesn't re-2FA

# Open OnDemand: interactive apps



The screenshot shows the Open OnDemand web interface. The top navigation bar includes 'Files', 'Jobs', 'Clusters', and 'Interactive Apps'. The main content area features the Virginia Tech OnDemand logo and a list of available applications. A sidebar on the right, titled 'Interactive Apps [Sandbox]', provides a detailed list of these applications.

**Interactive Apps [Sandbox]**

- Cascades
- ANSYS Workbench
- BYO-Container
- COMSOL Multiphysics
- Example Shiny App
- Genome Browser
- Jupyter Notebook
- Jupyter Notebook --Container
- LaBB-CAT
- MATLAB
- NSIGHT
- ParaView
- QGIS
- Remote Desktop
- Remote Desktop
- Rstudio
- Stata
- pymol
- super\_test

- System wide apps
- User developed apps



# Open OnDemand: interactive apps

RStudio Server (Owens and Pitzer) (11421197.owens-batch.ten.osc.edu)

1 node | 1 core | Running

Host: [>\\_o0456.ten.osc.edu](#)

Delete

Created at: 2020-10-16 18:09:41 EDT

Time Remaining: 1 hour

Session ID: [d842589d-cc50-41bb-8a87-32b2d62cf49b](#)

Connect to RStudio Server

MATLAB (11421211.owens-batch.ten.osc.edu)

1 node | 1 core | Running

Host: [>\\_o0467.ten.osc.edu](#)

Delete

Created at: 2020-10-16 18:11:57 EDT

Time Remaining: 1 hour

Session ID: [5d9be97c-6480-4e0d-b27d-248aebec1e97](#)

noVNC Connection

Native Instructions

Compression

0 (low) to 9 (high)

Image Quality

0 (low) to 9 (high)

Launch MATLAB

View Only (Share-able Link)

- Reverse proxy setup automated
- NoVNC
  - Image/connection settings
- Direct SSH to node
- View only sharable link

# Open OnDemand: form based jobs

## BYO-Container

This app will build a user [Singularity container](#) GUI on the [Cascades cluster](#). The app will exit once the build is complete.

### Container to build e.g. docker://alpine:latest

- This should be the full container name from the registry
- This will run: `singularity pull --force container`
- Creating: `/work/cascades/PID/singularity/container_tag.sif`

use cache if checked -- COMING SOON

Unchecked adds flag `--disable-cache` to pull command

### Slurm account

Slurm account, allocation, to use.

### Number of hours

Choose the partition to run in. This will be an 8 core job.

I would like to receive an email when the session starts

Launch

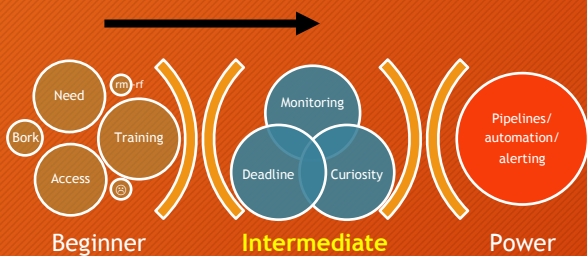
\* All BYO-Container session data is generated and stored under the user's home directory in the corresponding [data root directory](#).

- Same mechanism for GUI app launch can drive scriptable jobs
- Further simplify standard workflows
  - parameter sweeps
  - Production codes
  - Monitoring workflows

# Open OnDemand: job composer

Support transition from GUI prototyping to scripting  
(and back for result exploration)

- Job composer brings familiar tools to scripting
- Web based interaction with schedulers
- Web based script editing
- Scripts from templates
- Web based job monitoring



VT OnDemand / Job Composer Jobs Templates Help

Job was successfully updated.

### Jobs

+ New Job - Create Template

Edit Files Job Options Open Terminal Submit Stop Delete

Show 25 entries Search:

Created	Name	ID	Cluster	Status
September 17, 2020 10:10am	(default) Simple Sequential Python Job		Huckleberry	Not Submitted
August 10, 2020 1:22pm	(default) Simple Sequential Job		Cascades	Not Submitted
July 22, 2020 11:24am	(default) Simple Sequential Python Job	418809	Cascades	Completed
July 21, 2020 2:43pm	(default) Simple Sequential Python Job	418397	Cascades	Completed
September 13, 2019 9:30am	(default) Simple Sequential Job	64053	Huckleberry	Completed
July 3, 2019 6:24am	(default) Simple Sequential Job	64177	Huckleberry	Completed
May 20, 2019 2:40pm	(default) Simple Sequential Job	94442	Cascades	Completed
May 16, 2019 10:39am	(default) Simple Sequential Job	53507	Huckleberry	Completed

Showing 1 to 8 of 8 entries Previous 1 Next

---

VT OnDemand / Active Jobs

Your Jobs - All Clusters -

### Active Jobs

Show 50 entries Filter:

ID	Name	User	Account	Time Used	Queue	Status	Cluster
> 323223	qstat.30	rsettlag	arctest		normal_q.dev_q	Queued	Dragonstooth
> 323222	qstat.20	rsettlag	arctest		normal_q.dev_q	Queued	Dragonstooth
> 323220	qstat.10	rsettlag	arctest		normal_q.dev_q	Queued	Dragonstooth
> 323219	qstat.00	rsettlag	arctest		normal_q.dev_q	Queued	Dragonstooth
> 323218	qstat.50	rsettlag	arctest		normal_q.dev_q	Queued	Dragonstooth
> 323216	qstat.40	rsettlag	arctest		normal_q.dev_q	Queued	Dragonstooth

Showing 1 to 6 of 6 entries Previous 1 Next

---

### Job Details

Job Name: (default) Simple Sequential Python Job

Submit to: Huckleberry

Account: Not specified

Script location: /home/rsettlag/ondemand/data/sys/myjobs/projects/default/14

Script name: main\_job.sh

Folder Contents: /main\_job.sh

---

### Submit Script

main\_job.sh

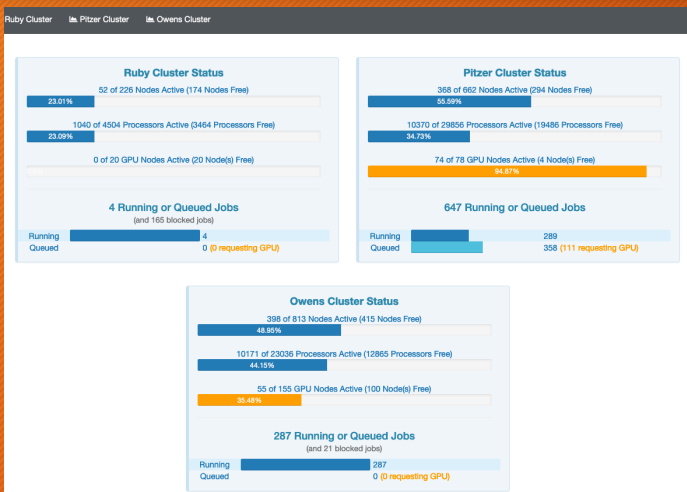
Script contents:

```
#!/bin/bash
#SBATCH --job-name=python_script
#SBATCH --time=01:00:00
#SBATCH -n 1

# A Basic Python Serial Job
#
# The following lines set up the Python environment
module load Anaconda/5.2.0
source activate pearc20
#
# Move to the directory where the job was submitted from
# You could also 'cd' directly to your working directory
cd $SLURM_SUBMIT_DIR
#
# Run Python
#
python hello.py
```

Open Editor Open Terminal Open Dir

# Open OnDemand - XDMoD integration



**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**  
 2020-03-16 - 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.

**Jobs Efficiency Report - 2020-09-19 to 2020-10-19**  
 75.0% efficient, 25.0% inefficient  
 2 inefficient jobs / 8 total jobs

**Core Hours Efficiency Report - 2020-09-19 to 2020-10-19**  
 98.1% efficient, 1.9% inefficient  
 0.0 inefficient core hours / 2.0 total core hours

**Recently Completed Jobs - 2020-09-19 to 2020-10-19**

ID	Name	Date
11421211	ondemand/roys	10/6
	/dashboard	
	/sys/bc_osc_matlab	

**Jobs Efficiency Report - 2020-09-16 to 2020-10-16**  
 92.3% efficient, 7.7% inefficient  
 1635 inefficient jobs / 21184 total jobs

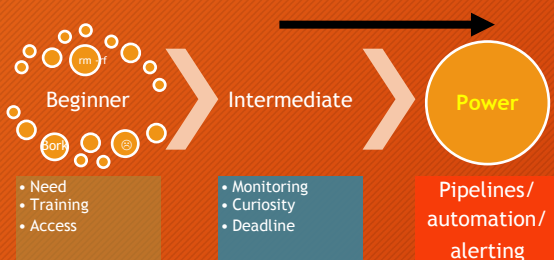
**Core Hours Efficiency Report - 2020-XDMoD 09-16 to 2020-10-16**  
 42.7% efficient, 57.3% inefficient  
 937234.6 inefficient core hours / 1636543.2 total core hours

**Recently Completed Jobs - 2020-09-XDMoD 16 to 2020-10-16**

ID	Name	Date	CPU
4274244	SimRunArray_12	10/5	25.5
4270864	slurm_s64	10/5	98.5
4274168	SimRunArray_12	10/5	14.0
4271259	slurm_s64	10/5	98.5

- XDMoD provide job level performance data analytics
- Power users search for available hardware features
  - Power users are cognizant of hardware capabilities

- mobile device friendly



# OnDemand - XDMoD job stats

**XDMoD** Hello, Robert Settlege (logout) Ohio Supercomputer Center

Dashboard Usage Metric Explorer App Kernels Report Generator Job Viewer About

Jobs - 2020-09-19 to 2020-10-19

Job Identifier	Start	End	CPU
owens-11421211	2020-10-16 18:12:26	2020-10-16 18:12:37	●
owens-11421208	2020-10-16 18:11:23	2020-10-16 18:11:36	●
owens-11421197	2020-10-16 18:10:20	2020-10-16 18:11:13	●
owens-11419592	2020-10-16 10:46:12	2020-10-16 10:50:23	●
owens-11419588	2020-10-16 10:41:57	2020-10-16 10:42:36	●
owens-11419469	2020-10-16 10:09:43	2020-10-16 10:09:46	N/A
owens-11419164	2020-10-16 09:59:42	2020-10-16 10:01:04	●
owens-11419162	2020-10-16 09:58:18	2020-10-16 09:58:20	N/A

Wait Hours - 2020-09-19 to 2020-10-19

Wait times by queue - 2020-09-19 to 2020-10-19

Job Efficiency Report - 2020-09-19 to 2020-10-19

<b>Job Efficiency</b>	<b>Total Job Count</b>	8
	<b>Inefficient Job Count</b>	2
<b>Core Hour Efficiency</b>	<b>Total Core Hours</b>	2
	<b>Inefficient Core Hours</b>	0.04

**XDMoD** Hello, Robert Settlege (logout) Ohio Supercomputer Center

Dashboard Usage Metric Explorer App Kernels Report Generator Job Viewer About

Search History

owens-11421211 | owens-11419592

CPU User: 0.69 | WaitTime Accuracy: 0.139 | CPU User Balance: 0.792 | Memory Headroom: 0.955

**XDMoD** Hello, Robert Settlege (logout) Ohio Supercomputer Center

Dashboard Usage Metric Explorer App Kernels Report Generator Job Viewer About

Search History

owens-11421211 | owens-11419592

CPU User: 0.69 | WaitTime Accuracy: 0.139 | CPU User Balance: 0.792 | Memory Headroom: 0.955

**XDMoD** Hello, Robert Settlege (logout) Ohio Supercomputer Center

Dashboard Usage Metric Explorer App Kernels Report Generator Job Viewer About

Search History

owens-11421211 | owens-11419592

CPU User: 0.69 | WaitTime Accuracy: 0.139 | CPU User Balance: 0.792 | Memory Headroom: 0.955


# Open OnDemand + Power AI

- Huckleberry
  - 14 node Power 8 cluster
- ECE recruitment tool
- Teaching tool
- AI workhorse
- DeepLearning Hackathons

reserved/down nodes	14	14	14	0	2	0	0
cores available	587	523	587	72	0	352	376
gpus available	-	-	-	-	8	8	11
1/2 nodes available	18	16	18	1	0	15	16
full nodes available	17	15	17	1	0	0	1
%utilization cores/nodes	89.6/90.3	90.6/91.4	89.6/90.3	50/50	100/100	61.4/100	60.8/97.5

Dragonstooth core and node availability at 19 October, 2020 -- 11:20:24AM

	dev_q	normal_q	preemptable_q
total cores/node	24	24	24
nodes in partition	47	42	47
reserved/down nodes	0	0	0
cores available	107	36	107
1/2 nodes available	5	2	5
full nodes available	3	1	3
%utilization cores/nodes	90.5/93.6	96.4/97.6	90.5/93.6



Huckleberry core and node availability at 19 October, 2020 -- 11:20:24AM

	normal_q
total cores/node	20
total gpus/node	4
nodes in partition	14
reserved/down nodes	3
cores available	0
gpus available	0
1/2 nodes available	0
full nodes available	0
%utilization cores/nodes	100/100

# Huckleberry Power8 AI

Power is just another choice in OnDemand.

To further enable, we have surfaced the IBM Power AI stack through Anaconda. To facilitate use, allow loading of the environment via Open OnDemand.

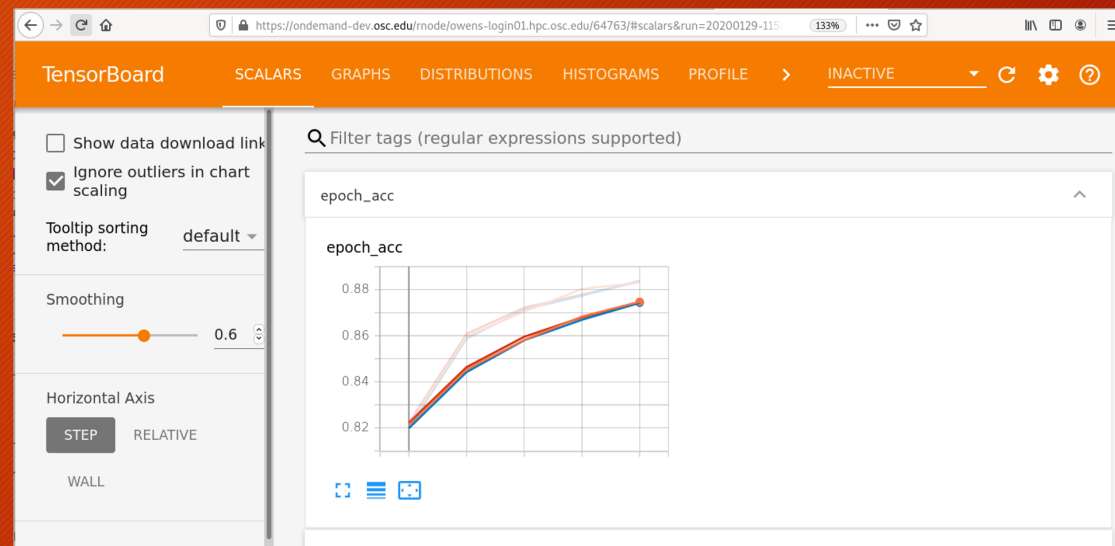
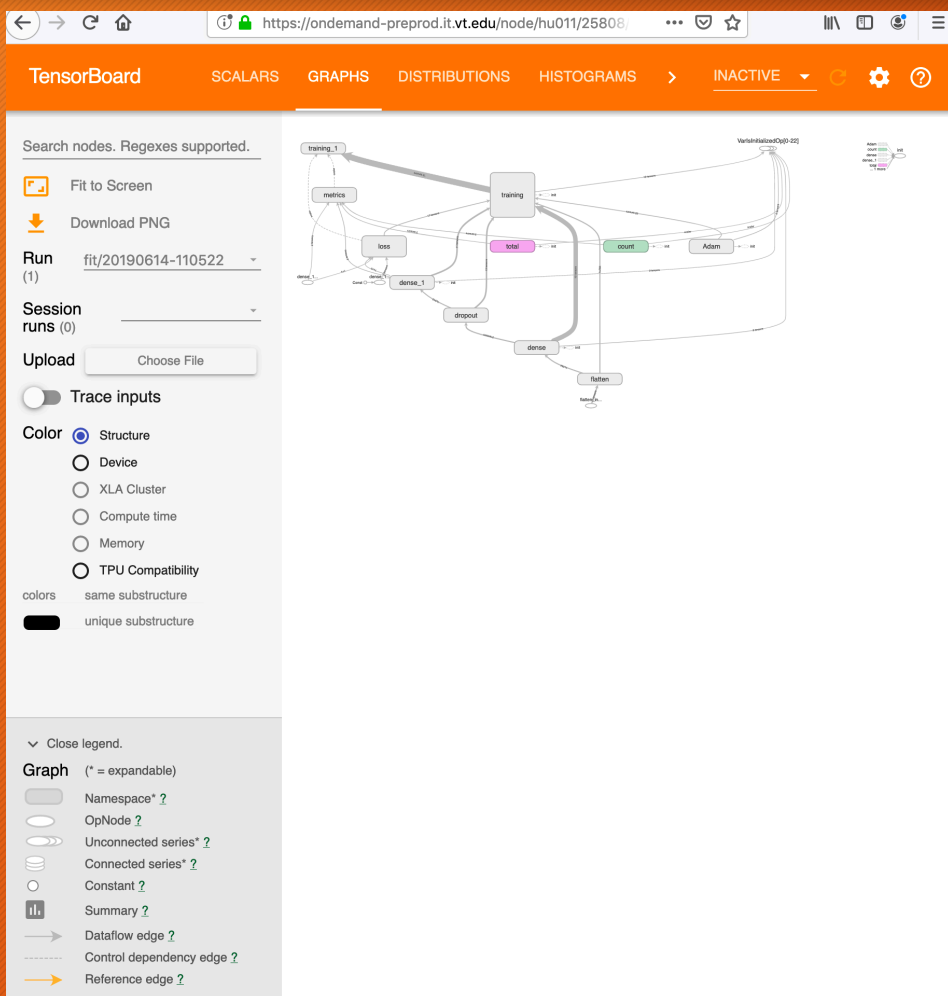
The screenshot displays the configuration page for a Jupyter Notebook application in the Open OnDemand environment. On the left, a list of 'Interactive Apps' includes various software like COMSOL, MATLAB, and ParaView, with 'Jupyter Notebook' selected. The main panel is titled 'Jupyter Notebook' and contains the following configuration sections:

- Account:** A dropdown menu set to 'ondemand2' with a note: 'The allocation you would like to use for SLURM.'
- Reservation:** An empty text input field.
- Partition:** A dropdown menu set to 'normal\_q' with a note: 'To request a GPU enabled queue, preface it with v100\_. Example: v100\_normal\_q'.
- Number of hours (min-1, max-48):** A spinner control set to '1'.
- Number of nodes (min-1, max-2):** A spinner control set to '1'.
- Number of cores per node (min-1, max-20):** A spinner control set to '10'.
- Number of GPUs per node (min-0, max-4):** A spinner control set to '1' with a note: 'If requesting GPU nodes, you must enter a GPU-enabled Partition above or the job will fail.'
- Required modules:** A text input field containing 'gcc/7.3.0 cuda/10.1.105 jdk/8.0.5.31 Anaconda3/2019.03' with a note: 'This should be spaced list of -compatible- modules.'
- Conda environment:** A dropdown menu set to 'powerai16\_ibm' with a note: 'This should be a Conda environment you wish to load. The conda'.

Note: traditional ssh; module load; conda activate works!

# AI/ML/DL benefit from Tensorboards

- Monitor training/learning progress
- Troubleshoot network design





# Hackathons, education, new users

- Power + Open OnDemand!!
- Hackathons are short duration learning challenges.
  - Move the focus from platform access to the hackathon topic - immediately!!
- Education
  - Don't create the day 1 bottleneck of “how do I ssh to ...”
- New users
  - Gentle introductions, all traditional modes of access are maintained

# What types of AI projects are running at VT

- Multi-robot coordination and navigation
- Smart city communications and security
- Crop identification and management
- Physics informed predictive turbulence
- Auto-tuning of CFD algorithms
- Prosthetic tuning for walking
- Improvements in communications systems
- Autonomous vehicles
- Search and rescue via drone swarms
- ...

# Power group: Jian-Bin Huang, ECE@VT

## Research Goal

Leverage large amounts of visual data that are readily available on the web to **learn, understand, and synthesize** our visual world.

*Teach machines to see!*

## Motivation

Machine learning (deep learning) models work well if we have lots of labeled data.

Supervised learning  
•Not scalable!

Unlabeled data  
•Readily available!

Our approach



•Use both labeled and unlabeled data for visual learning



## Learning Visual Reconstruction

How do we recover the missing data?



Image completion




Video completion



Blur removal

## Learning Dense Correspondence

How do we match and align a pair images?



Stereo matching




Motion estimation




Semantic correspondence

## Learning Generalizable Representation

How do train models that learn new task efficiently?



Temporal order



Lifelong learning

## Selected Publications

-  MaskRNN: Instance Level Video Object Segmentation  
Y.-T. Hu, [J.-B. Huang](#), and A. Schwing, NIPS 2017
-  Semi-Supervised Learning for Optical Flow with Generative Adversarial Networks  
W.-S. Lai, [J.-B. Huang](#), and M.-H. Yang, NIPS, 2017
-  Unsupervised Representation Learning by Sorting Sequences  
H.-Y. Lee, [J.-B. Huang](#), M. Singh, and M.-H. Yang, ICCV, 2017
-  Deep Laplacian Pyramid Networks for Fast and Accurate Super-Resolution  
W.-S. Lai, [J.-B. Huang](#), N. Ahuja, and M.-H. Yang, CVPR 2017
-  Temporally Coherent Completion of Dynamic Video  
[J.-B. Huang](#), S. B. Kang, N. Ahuja, and J. Kopf, SIGGRAPH Asia 2016 (Oral Presentation)
-  Deep Joint Guided Image Filtering  
Y. Li, [J.-B. Huang](#), N. Ahuja, and M.-H. Yang, ECCV 2016
-  Tracking Persons-of-Interest via Adaptive Discriminative Features  
S. Zhang, Y. Gong, [J.-B. Huang](#), J. Wang, J. Lim, N. Ahuja, and M.-H. Yang, ECCV 2016
-  Unsupervised Visual Representation Learning by Graph-based Consistent Constraints  
D. Li, W.-C. Hung, [J.-B. Huang](#), S. Wang, N. Ahuja, and M.-H. Yang, ECCV 2016
-  Detecting Migrating Birds at Nights  
[J.-B. Huang](#), R. Caruana, A. Farnsworth, S. Kelling, and N. Ahuja, CVPR 2016
-  A Comparative Study for Single Image Blind Deblurring  
W.-S. Lai, [J.-B. Huang](#), Z. Hu, N. Ahuja, and M.-H. Yang, CVPR 2016 (Spotlight Oral)
-  Weakly Supervised Object Localization with Progressive Domain Adaptation  
D. Li, [J.-B. Huang](#), Y. Li, S. Wang, and M.-H. Yang, CVPR 2016
-  Hierarchical Convolutional Features for Visual Tracking  
C. Ma, [J.-B. Huang](#), X. Yang, and M.-H. Yang, ICCV 2015
-  Single Image Super-Resolution with Transformed Self-Exemplars.  
[J.-B. Huang](#), A. Singh, and N. Ahuja, CVPR 2015 (Oral Presentation)
-  Image Completion using Planar Structure Guidance.  
[J.-B. Huang](#), S. B. Kang, N. Ahuja, and J. Kopf, SIGGRAPH 2014 (Oral Presentation)
-  Towards Accurate and Robust Cross-Ratio based Gaze Trackers Through Learning From Simulation.  
[J.-B. Huang](#), Q. Cai, Z. Liu, N. Ahuja, and Z. Zhang, ETRA 2014 (Best Paper Award)
-  Transformation Guided Image Completion.  
[J.-B. Huang](#), J. Kopf, N. Ahuja, and S. B. Kang, ICCP 2013 (Oral Presentation)
-  Saliency Detection via Divergence Analysis: A Unified Perspective.  
[J.-B. Huang](#) and N. Ahuja, ICPR 2012 (Best Paper Award)

# Open OnDemand: Adoption is growing

- Features:
  - ✓ Plugin-free web experience
  - ✓ Easy file management
  - ✓ Command-line shell access
  - ✓ Job management and monitoring
  - ✓ Graphical desktop environments
  - ✓ GUIs
- Install base is growing
- Industry and academia
- ~140 unique US locations
- ~70 unique international locations

## Production Deployments



## In Process of Installing



## Interested / Evaluating



# OPEN OnDemand

Questions?

Robert Settlage  
Advanced Research  
Computing @Virginia Tech  
Oct 2020