Interactive Computing with Open OnDemand

Compute Ontario Colloquium

James Willis (SciNet)

December 11, 2024

Outline

- Motivation
- What is Open OnDemand?
- Key Features
- SciNet's Open OnDemand Portal
- Demo
- Summary
- Installation and Configuration

Motivation

- Terminal based interfaces can be very daunting for new users with little to no experience
- We need a way to make HPC more accessible to improve the learning curve
- Web-based interfaces are a good solution
- There are many available:
  - JupyterHub
  - RStudio
  - Galaxy
  - **Open OnDemand**

What is Open OnDemand?

- Open OnDemand is a web-based interface that provides access to HPC resources
- Open-source project developed by the Ohio Supercomputer Center (OSC)
- Funded by the National Science Foundation (NSF)
- Used worldwide across ∼400 HPC centres, including:
  - Grex - University of Manitoba
  - Bridges2 - University of Pittsburgh
  - Anvil - Purdue University
  - Expanse - San Diego
  - LUMI - IT Center for Science (Finland)

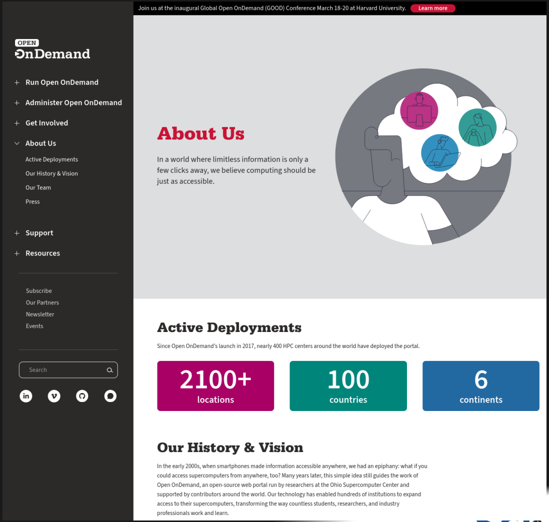

About Us

In a world where limitless information is only a few clicks away, we believe computing should be just as accessible.

Active Deployments

Since Open OnDemand's launch in 2017, nearly 400 HPC centers around the world have deployed the portal.

2100+ locations

100 countries

6 continents

Our History & Vision

In the early 2000s, when smartphones made information accessible anywhere, we had an epiphany: what if you could access supercomputers from anywhere, too? Many years later, this simple idea still guides the work of Open OnDemand, an open-source web portal run by researchers at the Ohio Supercomputer Center and supported by contributors around the world. Our technology has enabled hundreds of institutions to expand access to their supercomputers, transforming the way countless students, researchers, and industry professionals work and learn.

What is Open OnDemand?

- Designed to make HPC more accessible to users with no prior experience
- Built with Ruby and JavaScript
- Provides a graphical interface to computing resources as opposed to the command line
- Great tool for teaching and learning HPC skills incrementally
- Lowers the barrier to entry for new users who may be intimidated by the command line interface
- You can even access it from your mobile phone!

Open OnDemand Deployments

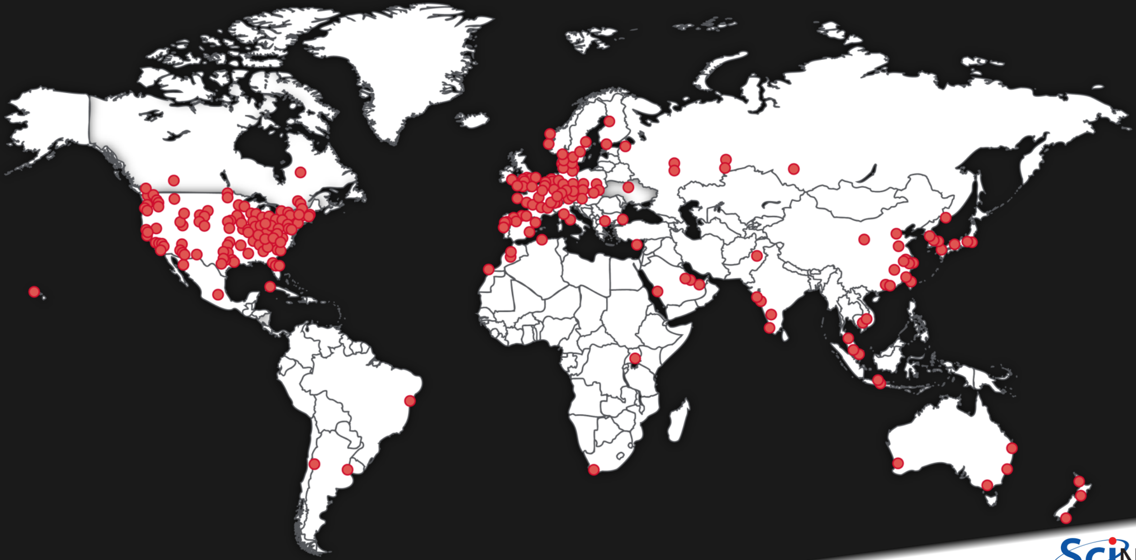

Key Features

Web-Based Access

- Users can access HPC resources through a web browser without needing to install specialised software or configure SSH connections manually

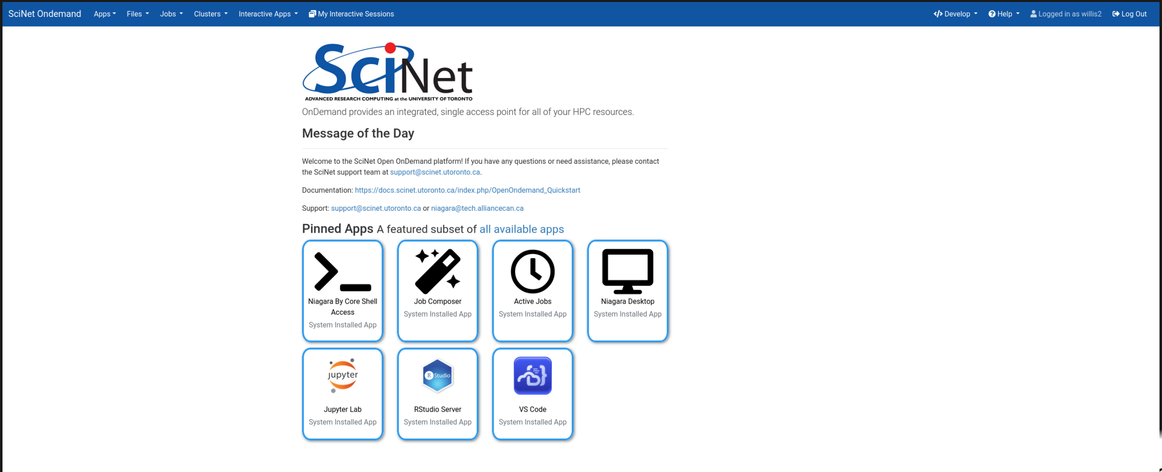

OnDemand provides an integrated, single access point for all of your HPC resources.

Message of the Day

Welcome to the SciNet Open OnDemand platform! If you have any questions or need assistance, please contact the SciNet support team at support@scinet.utoronto.ca.

Documentation: https://docs.scinet.utoronto.ca/index.php/OpenOndemand_Quickstart

Support: support@scinet.utoronto.ca or niagara@tech.alliancecan.ca

Pinned Apps A featured subset of all available apps

Job Management

- It provides an intuitive way to submit, monitor, and manage batch jobs using the SLURM scheduler. Users can easily create and customise job scripts using web forms

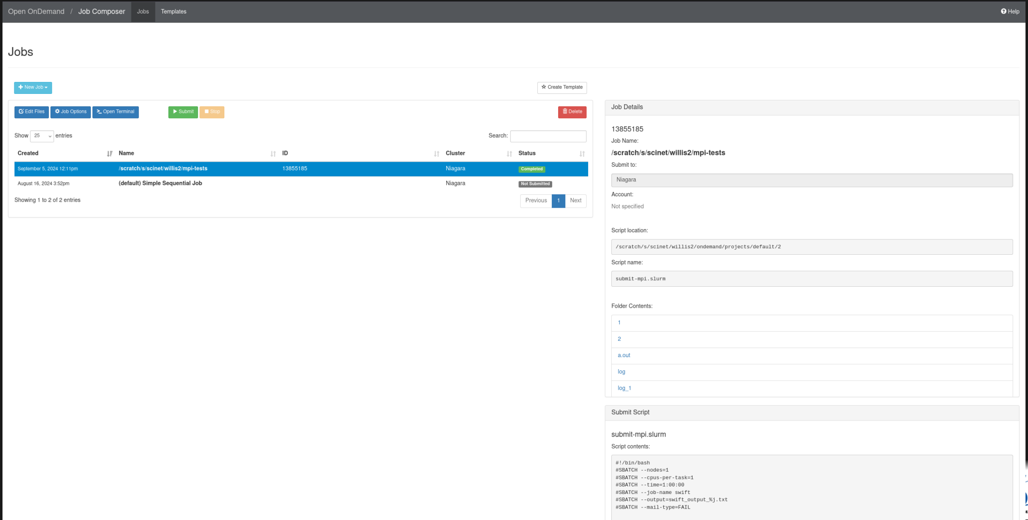

Jobs

Show 25 entries

Search:

Created	Name	ID	Cluster	Status
September 5, 2024 12:11pm	/scratch/s/scinet/willis2/mpi-tests	13855185	Niagara	Completed
August 16, 2024 3:52pm	(default) Simple Sequential Job		Niagara	Not Submitted

Showing 1 to 2 of 2 entries

Previous 1 Next

Job Details

13855185

Job Name:
/scratch/s/scinet/willis2/mpi-tests

Submit to:
Niagara

Account:
Not specified

Script location:
/scratch/s/scinet/willis2/ondemand/projects/default/2

Script name:
submit-mpi.slurm

Folder Contents:
- 1
- 2
- a.out
- log
- log_1

Submit Script

submit-mpi.slurm

Script contents:

```
#!/bin/bash
#SBATCH --nodes=1
#SBATCH --cpus-per-task=1
#SBATCH --time=1:00:00
#SBATCH --job-name swift
#SBATCH --output=swift_output_%j.txt
#SBATCH --mail-type=FAIL

#cd $SLURM_SUBMIT_DIR
```

Job Monitoring

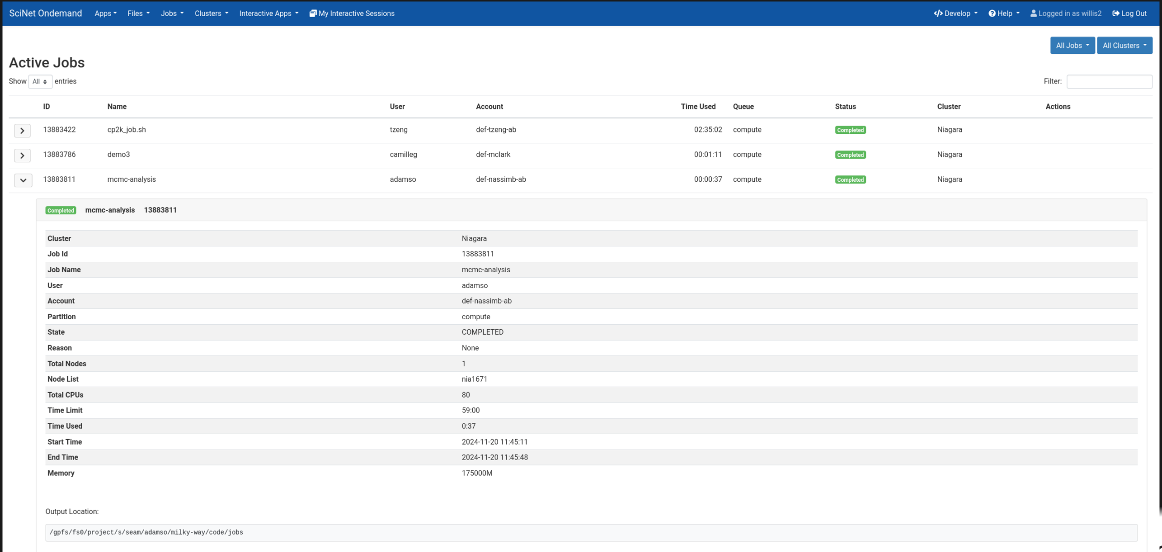

Active Jobs

	ID	Name	User	Account	Time Used	Queue	Status	Cluster	Actions
>	13883422	cp2k_job.sh	tzeng	def-tzeng-ab	02:35:02	compute	Completed	Niagara	
>	13883786	demo3	camilleg	def-mclark	00:01:11	compute	Completed	Niagara	
⌄	13883811	mcmc-analysis	adamso	def-nassimb-ab	00:00:37	compute	Completed	Niagara	

Completed **mcmc-analysis** **13883811**

Cluster	Niagara
Job Id	13883811
Job Name	mcmc-analysis
User	adamso
Account	def-nassimb-ab
Partition	compute
State	COMPLETED
Reason	None
Total Nodes	1
Node List	nia1671
Total CPUs	80
Time Limit	59:00
Time Used	0:37
Start Time	2024-11-20 11:45:11
End Time	2024-11-20 11:45:48
Memory	175000M

Output Location:

```
/gpfs/fs0/project/s/seam/adamso/milky-way/code/jobs
```

File Management

- Users can upload, download, edit, and manage files stored on the cluster through a built-in file browser

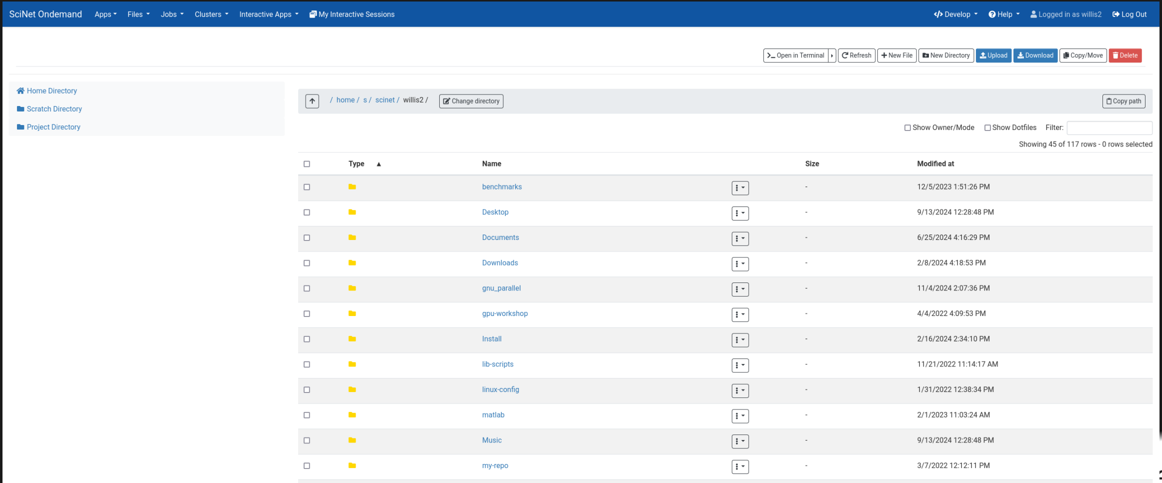

/ home / s / scinet / willis2 /

Showing 45 of 117 rows - 0 rows selected

☐	Type ▲	Name	Size	Modified at
☐	📁	benchmarks	-	12/5/2023 1:51:26 PM
☐	📁	Desktop	-	9/13/2024 12:28:48 PM
☐	📁	Documents	-	6/25/2024 4:16:29 PM
☐	📁	Downloads	-	2/8/2024 4:18:53 PM
☐	📁	gnu_parallel	-	11/4/2024 2:07:36 PM
☐	📁	gpu-workshop	-	4/4/2022 4:09:53 PM
☐	📁	Install	-	2/16/2024 2:34:10 PM
☐	📁	lib-scripts	-	11/21/2022 11:14:17 AM
☐	📁	linux-config	-	1/31/2022 12:38:34 PM
☐	📁	matlab	-	2/1/2023 11:03:24 AM
☐	📁	Music	-	9/13/2024 12:28:48 PM
☐	📁	my-repo	-	3/7/2022 12:12:11 PM

Interactive Applications

- It supports running interactive applications like Jupyter Notebooks, RStudio, VS Code or remote desktop sessions, making it useful for data analysis, visualisation, and interactive computational tasks

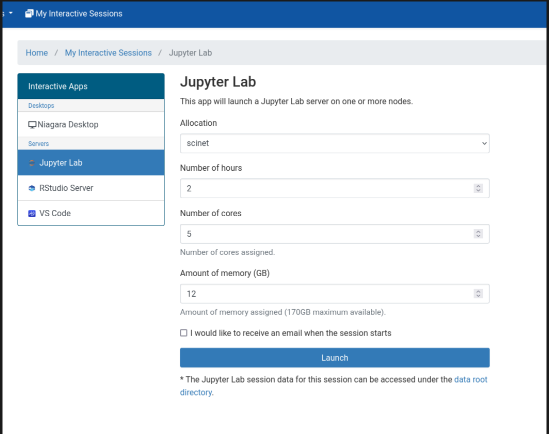

Jupyter Lab

This app will launch a Jupyter Lab server on one or more nodes.

Allocation: scinet

Number of hours: 2

Number of cores: 5

Number of cores assigned.

Amount of memory (GB): 12

Amount of memory assigned (170GB maximum available).

☐ I would like to receive an email when the session starts

* The Jupyter Lab session data for this session can be accessed under the data root directory.

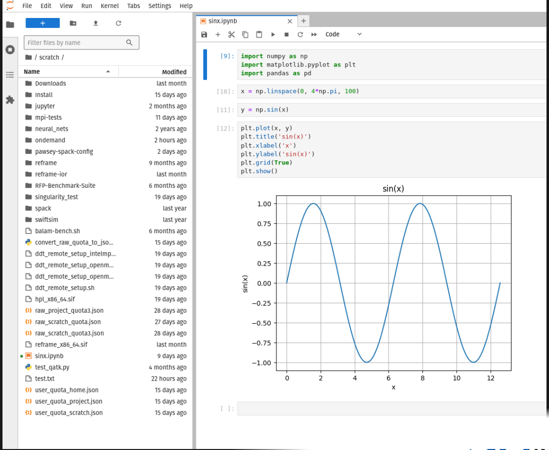

Name	Modified
Downloads	last month
install	15 days ago
jupyter	2 months ago
mpi-tests	11 days ago
neural_nets	2 years ago
ondemand	2 hours ago
pawsey-spack-config	2 days ago
reframe	9 months ago
reframe-ior	last month
RFP-Benchmark-Suite	6 months ago
singularity_test	19 days ago
spack	last year
swiftsim	last year
balam-bench.sh	6 months ago
convert_raw_quota_to_jso...	15 days ago
ddt_remote_setup_intelmp...	19 days ago
ddt_remote_setup_openm...	19 days ago
ddt_remote_setup_openm...	19 days ago
ddt_remote_setup.sh	19 days ago
hpl_x86_64.sif	19 days ago
raw_project_quota3.json	28 days ago
raw_scratch_quota.json	27 days ago
raw_scratch_quota3.json	28 days ago
reframe_x86_64.sif	last month
sinx.ipynb	9 days ago
test_qatk.py	4 months ago
test.txt	22 hours ago
user_quota_home.json	15 days ago
user_quota_project.json	15 days ago
user_quota_scratch.json	15 days ago

```
[9]:  import numpy as np
      import matplotlib.pyplot as plt
      import pandas as pd

[10]: x = np.linspace(0, 4*np.pi, 100)

[11]: y = np.sin(x)

[12]: plt.plot(x, y)
      plt.title('sin(x)')
      plt.xlabel('x')
      plt.ylabel('sin(x)')
      plt.grid(True)
      plt.show()

[ ]:
```

Interactive Applications

- Full list of supported apps:

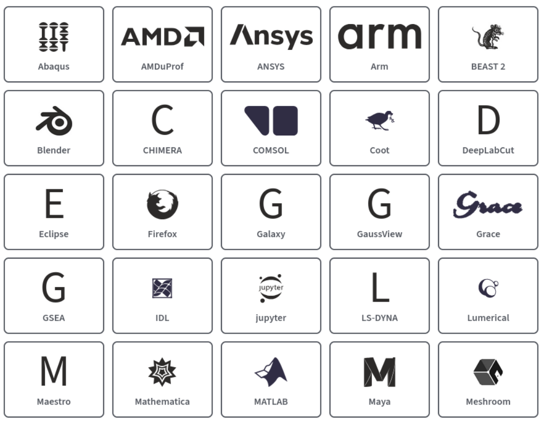

Abaqus	AMDuProf	ANSYS	Arm	BEAST 2
Blender	CHIMERA	COMSOL	Coot	DeepLabCut
Eclipse	Firefox	Galaxy	GaussView	Grace
GSEA	IDL	jupyter	LS-DYNA	Lumerical
Maestro	Mathematica	MATLAB	Maya	Meshroom

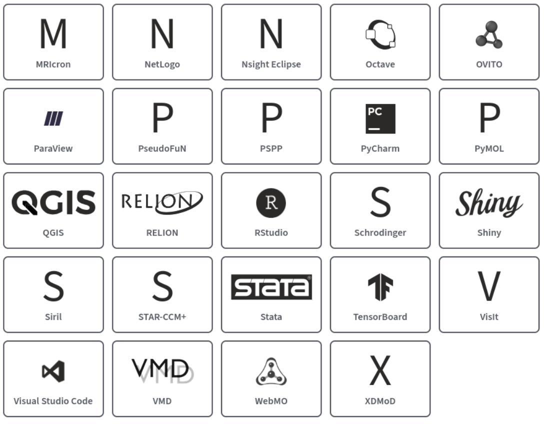

MRIcron	NetLogo	Nsight Eclipse	Octave	OVITO
ParaView	PseudoFuN	PSPP	PyCharm	PyMOL
QGIS	RELION	RStudio	Schrodinger	Shiny
Siril	STAR-CCM+	Stata	TensorBoard	VisIt
Visual Studio Code	VMD	WebMO	XDMoD	

- Unsupported applications can be developed and added to the growing community of Open OnDemand apps

Terminal Access

- It provides a web-based shell to run command-line tasks directly from the browser, which can ease new users into using a terminal

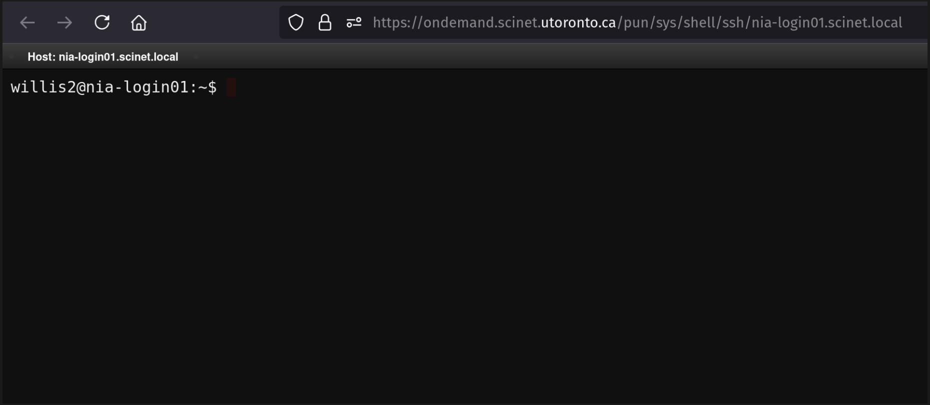

```
willis2@nia-login01:~$
```

Advantages over JupyterHub

- You can:
  - Browse the file system
  - Run terminal commands
  - Monitor active jobs

  without having to submit a job and waiting for a node to be allocated to you
- Open OnDemand compute nodes have internet access by default on Niagara
- There are large number of applications already supported within the community
- It's being actively developed and has a large user base
- High level of customisation and extensibility through configuration files (`.yaml`)
- Can customise the user experience based on their research needs

SciNet Deployment

- We have deployed Open OnDemand at SciNet and are currently in beta testing
- It provides access to Niagara currently, but will be expanded to other systems in the future, i.e. Trillium
- Users can access the SciNet Open OnDemand portal at: **https://ondemand.scinet.utoronto.ca**
- Login credentials are the same as your Alliance account
- MFA is enabled, so you will be prompted to authenticate with Duo
- There is a quickstart guide on our wiki: **https://docs.scinet.utoronto.ca/index.php/Open_OnDemand_Quickstart**
- We'd like to thank Grigory Shamov and Stefano Ansaloni from the University of Manitoba for their help with the initial setup

SciNet Deployment

- We have the following setup available on our portal so far:
  - Filebrowser
  - Job Submission
  - Job Monitoring
  - Terminal
  - Jupyter Lab
  - RStudio
  - VS Code
  - Remote Desktop
- The CC stack is used to load RStudio, VS Code, and Remote Desktop applications
- **Note: this is a work in progress and we are continuously adding new features/applications, fixing bugs and you may face outages**

SciNet Deployment - SLURM Scheduling

- We have setup a new scheduler for interactive Open OnDemand jobs
- Queue specs:
  - 2 nodes (`nia0073` and `nia0074`)
  - Scheduled by-core
  - Can request how much memory you need in GB
  - 3 day job time limit

Open OnDemand Demo

Feedback

- Please give Open OnDemand a try and let us know what you think!
- We welcome feedback and are open to suggestions for improvement
- Let us know if you'd like to see any specific applications or features added

Summary

- We introduced Open OnDemand as an alternative to terminal-based interfaces for accessing HPC resources
- Discussed its key features and advantages over other web-based interfaces
- Showed the SciNet Open OnDemand portal and the applications currently available
- Please email any questions to: **support@scinet.utoronto.ca**

Open OnDemand Installation

SciNet Installation and Configuration

- There are a couple of different components needed to get Open OnDemand up and running:
  1. Portal
  2. User authentication
  3. Compute nodes
  4. SLURM scheduler

Open OnDemand Portal

- The Open OnDemand portal runs on a dedicated server(s) separate from the compute nodes
- Orchestrated by Nomad (https://www.nomadproject.io)
- Runs as a Nomad service inside a container
- The portal itself is an Apache web server
- Filesystem is mounted via a CSI volume

User Authentication

- User authentication is handled by Authelia (https://www.authelia.com):
- Authelia is an open-source authentication and authorisation server
- Users log in with their Alliance credentials
- MFA is enabled, users will be prompted to authenticate with Duo/Yubikey

Compute Nodes

- In the *Job Composer* app, users can submit jobs to the Niagara compute nodes as normal
- However, Niagara compute nodes are scheduled by node meaning all 40 cores and 180GB memory are allocated to the user
- For Open OnDemand interactive jobs, we decided to have a separate set of nodes that are scheduled by-core
- So that we can make better use of the computing resources
- We set aside two nodes for Open OnDemand interactive jobs (likely to expand in the future)
- These nodes are scheduled by-core
- They are provisioned using *Puppet* (https://www.puppet.com) which is a configuration management tool
- We are transitioning away from *xCAT* to *Puppet*

SLURM Scheduler

- It is more efficient for SLURM to stick to one scheduling policy, i.e. either by-node or by-core
- So we have setup a separate scheduler for Open OnDemand jobs
- This is also provisioned using *Puppet*
- Jobs are scheduled via an `ssh` wrapper
- An ssh connection is established to the OOD *login* node and an `sbatch` command is executed

Network Diagram