



**PERVASIVE  
TECHNOLOGY INSTITUTE**



**RESEARCH TECHNOLOGIES**  
UNIVERSITY INFORMATION TECHNOLOGY SERVICES

# Jetstream2: Accelerating science and engineering on-demand

**David Y. Hancock – Indiana University**

Director for Advanced Cyberinfrastructure

Jetstream2 Primary Investigator

CI Townhall for NSF AI Institutes

26 October 2022

*Jetstream2*

# What is “the” Jetstream(2)?

- NSF-funded production cloud environment
- Ease-of-use focus, rapid on-ramp to XSEDE/ACCESS
- **On-demand** interactive computing and persistent services for science gateways
- Enables configurable environments; **programmable cyberinfrastructure**



By Maria Morris: JS2 rear doors (lower) Banksy adaptation [non-commercial] (right)

Now with GPUs,  
large-memory,  
more faster PB!

**Jetstream2**





# NSF Vision and Blueprint

U.S. National Science Foundation (NSF) envisions an **agile, integrated, robust, trustworthy and sustainable CI ecosystem that drives new thinking and transformative discoveries in all areas of S&E research and education.**

- View CI more holistically...
- Recognize and support the translational research continuum...
- Develop a strategy that balances innovations with stability and continuity...
- Work closely with the diverse S&E communities to tightly couple discovery and innovation...
- Achieve new levels of usability by easing the pathways for discovering, accessing, understanding, and utilizing powerful CI capabilities...

From: OAC Vision & Blueprint: Overview and Computational Ecosystem (Apr 2019)

# What is Jetstream2 and why does it exist?

- Significant [**distinct**] evolution of the Jetstream cloud resource
- Under 10% NSF investment → support for 24% of institutions, 23% of active PIs, and 32% of users\*
- Jetstream has provided **6x more** SUs than **any other** XSEDE resource for Education
- Emphasis on ease-of-use, broad accessibility, *AI for Everyone*
- Provides **on-demand interactive** computing and persistent services for science gateways
- Enables *configurable* environments; *programmable cyberinfrastructure*

\*Based on XDMoD data at Workload Analysis Report: <http://arxiv.org/abs/1801.04306>



COMMERCIAL CLOUD

INTERNET®

INDIANA UNIVERSITY  
CYBERINFRASTRUCTURE

XSEDEnet  
Advanced Layer 2  
Services (AL2S) platform

UNIVERSITY OF HAWAI'I  
CYBERINFRASTRUCTURE

ARIZONA STATE UNIVERSITY  
CYBERINFRASTRUCTURE

REGIONAL
Compute
8 Nodes 1,024 Cores 4 TB RAM
Storage
768 TB
Accelerators
2 Nodes 1 TB RAM 8 GPUs

REGIONAL
Compute
8 Nodes 1,024 Cores 4 TB RAM
Storage
768 TB
Accelerators
2 Nodes 1 TB RAM 8 GPUs

TACC CYBERINFRASTRUCTURE		
Compute	REGIONAL	Accelerators
8 Nodes 1,024 Cores 4 TB RAM	Storage 768 TB	2 Nodes 1 TB RAM 8 GPUs

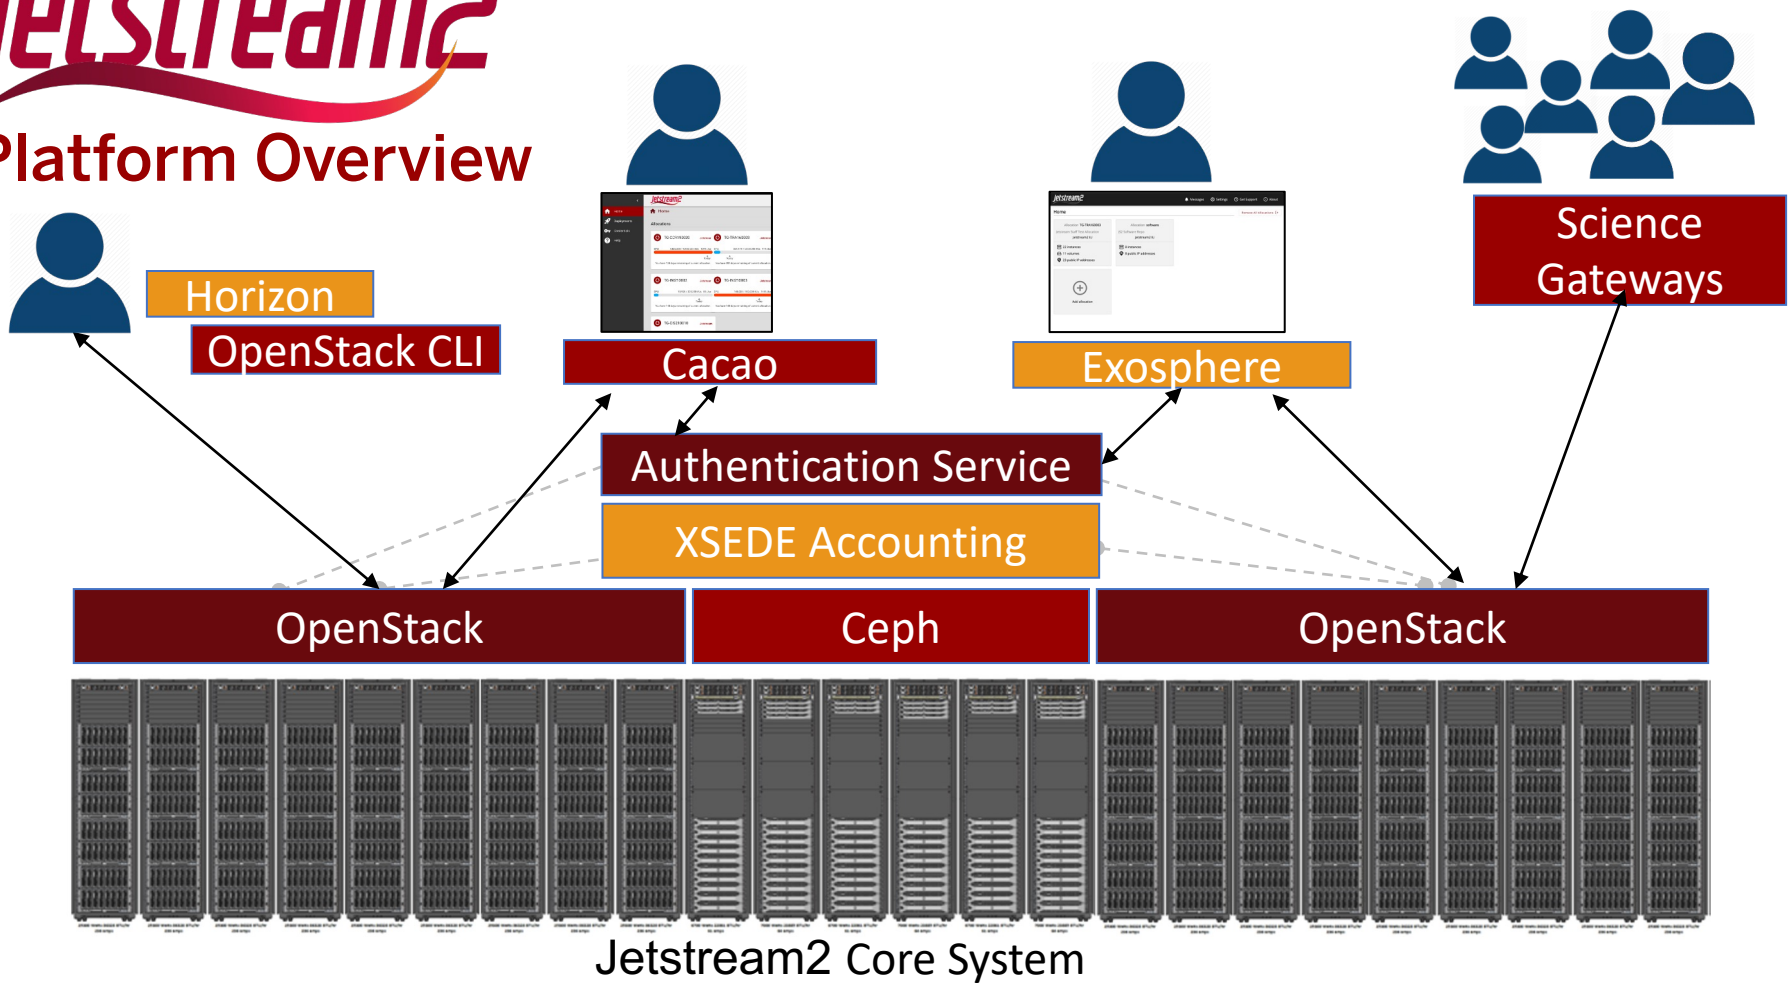
PRIMARY
Compute
416 Nodes 53,248 Cores 224 TB RAM
Storage
96 Nodes 14 PB
Accelerators
90 Nodes 45 TB RAM 360 GPUs

CORNELL UNIVERSITY CYBERINFRASTRUCTURE	
REGIONAL	
Compute	
8 Nodes 1,024 Cores 4 TB RAM	
Storage	
768 TB	



# Jetstream2

## Platform Overview





Remove Allocation Create

### Instances

Instances used 11 of 25 total Cores used 26 of 132 total RAM used 100 of 388 GB

Select All

- Ready `formally_trusty_urchin`
- Shelved `optionally_certain_longhorn with GUI`
- Ready `wildly_united_mite`

Hiding 8 Instances created by other users

Show

### Volumes

Volumes used 2 of 10 total Storage used 279 of 1,100 GB

Remove Allocation Create

### Instance formally\_trusty\_urchin

Created 19 minutes ago / by user tg836338 / from image JS-API-Featured-CentOS8-Latest

Status **Ready**  
UUID 2bc77f59-73bf-470f-95b6-51dc31d7577f  
Flavor m1.small  
SSH Public Key Name smart  
IP addresses  
Public IP Address 149.165.157.3  
Unassign

IP Details

### Volumes Attached

(none)

Attach volume

### Interactions

- Web Shell
- Web Desktop
- Native SSH: `exouser@149.165.157.3`
- Console

### Password

Try logging in with username "exouser" and the following password:

Show password

### Actions


- Lock Prevent further instance actions until it is unlocked
- Suspend Save execution state to disk
- Shelve Shut down instance and offload it from compute host
- Image Create snapshot image of instance
- Reboot Restart instance
- Delete Destroy instance

### Action History

Action	Time
create	19 minutes ago (2021-10-26 20:10:54 UTC)

### System Resource Usage

#### CPU Usage



#### Memory Usage



Home

Deployments

Credentials

Help

## Home

### Jetstream2 Alpha Release

Cloud Automation & Continuous Analysis Orchestration

Thank you for participating in the Jetstream2 Alpha release. Please use the following links to view known issues and submit any additional feedback as you use the application.

[VIEW KNOWN ISSUES](#) [SUBMIT FEEDBACK](#)

#### Allocations

Allocation ID	Project Name	CPU	GPU	Large Memory	Remaining
TRA220028	Jetstream2 Affiliated Development Projects	49,279 / 1,000,000 SUs (5% Used)	12,083 / 1,000,000 SUs (1% Used)	0 / 1,000,000 SUs (0% Used)	288 days
TRA160003	Jetstream Staff Test Allocation	146,965 / 2,000,000 SUs (7% Used)	95,323 / 2,000,000 SUs (5% Used)	0 / 2,000,000 SUs (0% Used)	87 days
CIS220046	Deep Learning Tutorial for Translational AI Center at Iowa State University	-	424,222 / 600,000 SUs (71% Used)	0 / 2,000,000 SUs (0% Used)	115 days

#### Featured Learning

- [Continuous Analysis 101](#)
- [Jetstream2 Basics](#)
- [Manage Resources](#)

### New Deployment

JETSTREAM 2 / BIO220047

Select a template that best describes what you want to do:

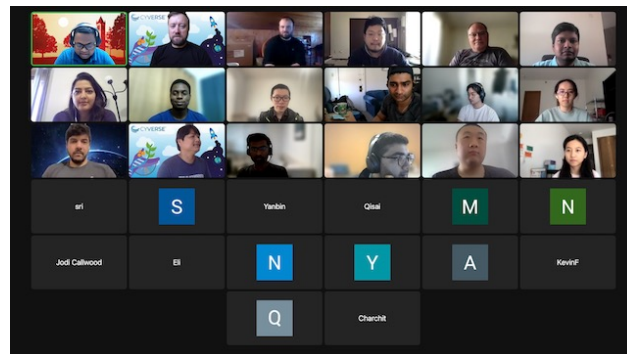
- [simple launch of one or more vms](#)  
openstack-single-image
- [launch a multi-vm zero-to-jupyterhub](#)  
jupyterhub
- [launch a multi-vm kubernetes cluster \(k3s\)](#)  
single-image-k3s
- [launch instances for a workshop](#)  
vm4workshop

[GO](#)



# Deep dive into DL

- First-of-a-kind workshop to apply DL techniques to agricultural data sets in April 2022
- AIIRA, AI Institute for Resilient Agriculture, intends to distribute the digital twin built on JS2 for community re-use
- Allows community training and inference
- Provided via Terraform templates and customized UI through CACAO



*...we were able to easily provide so many students with a GPU-enabled container so quickly. Normally, getting GPU resources on an HPC scheduler, like OnDemand, takes time, and the high demand for GPUs makes finding 40 or more unoccupied resources an impossibility.*

– Tyson Swetnam, CyVerse Co-PI and workshop instructor



[https://cyverse.org/deep\\_learning\\_workshop](https://cyverse.org/deep_learning_workshop)

# Early Operations Projects & Activity

- First PI invitations and projects added in early February 2022
- Followed push for all projects to migrate in May – July.
- Sept 2022: 322 projects and 1700 individuals (568 students)
- Includes multiple science gateways and education/training allocations
- Full production after recent NSF approval
- Retired Jetstream[1] in July/August 2022



"Bike Exchange - 2009 IU Women's Little 500" by  
Indiana Public Media  
Flickr CC BY-NC 2.0



**PERVASIVE  
TECHNOLOGY INSTITUTE**



**RESEARCH TECHNOLOGIES**  
UNIVERSITY INFORMATION TECHNOLOGY SERVICES

# Acknowledgements

NSF Awards 1053575 & 1548562 (XSEDE), 1445604 (Jetstream) and 2005506 (Jetstream2)

This document was developed with support from the National Science Foundation. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF.

Special thanks to contributors & Jetstream2 partners

- Jeremy Fischer, J. Michael Lowe, Therese Miller, Maria Morris, Winona Snapp-Childs, George Turner, and Chris Martin.
- Vendors, particularly Dell and NVIDIA, also deserve recognition for their efforts





**PERVASIVE  
TECHNOLOGY INSTITUTE**




**RESEARCH TECHNOLOGIES**  
UNIVERSITY INFORMATION TECHNOLOGY SERVICES

# Jetstream2 partners



THE UNIVERSITY  
OF ARIZONA®



JOHNS HOPKINS  
UNIVERSITY



UCAR



<http://jetstream-cloud.org/>  
National Science Foundation  
Award #ACI-2005506