



**RESEARCH TECHNOLOGIES**  
UNIVERSITY INFORMATION TECHNOLOGY SERVICES



# Jetstream Overview: A national research and education cloud – Lightning Edition

**ACNN Lightning Talks**  
**September 4, 2020 – Bloomington, IN. (Webinar)**

*Jeremy Fischer – [Jeremy@iu.edu](mailto:Jeremy@iu.edu) - Indiana University  
Manager, Jetstream Cloud, UITS Research Technologies*

*Fischer, J. (2020). Jetstream Overview: A national research and  
education cloud. Bloomington, IN. Retrieved from  
<https://jetstream-cloud.org/research/publications.php>*

# What is Jetstream and why does it exist?

- *NSF's first production cloud facility*
- *Focus on ease-of-use, broad accessibility*
- *User-selectable library of preconfigured virtual machines*
- *Provides **on-demand** interactive computing and analysis or persistent services such as gateways*
- *Enables configurable environments; **programmable cyberinfrastructure***



# Who uses Jetstream?

- *The researcher needing a handful of cores (1 to 44/vCPU)*
- *Software creators and researchers needing to create their own VMs and workflows*
- *Science gateway creators using Jetstream as either the frontend or processor for scientific jobs*
- *STEM Educators teaching on a variety of subjects*






# What Jetstream isn't...

- *It's not traditional HPC*
- *There's no shared filesystem (think cloudy!)*
- *There's no high-end interconnect fabric (keep thinking cloudy!)*
- *There aren't GPUs widely available*
- *It isn't Amazon, Azure, or GCE (similar, but...)*




# Jetstream System Overview



<http://wiki.jetstream-cloud.org/Network+configuration+and+policies>

# Platform Overview



# What is Jetstream – a closer look

---

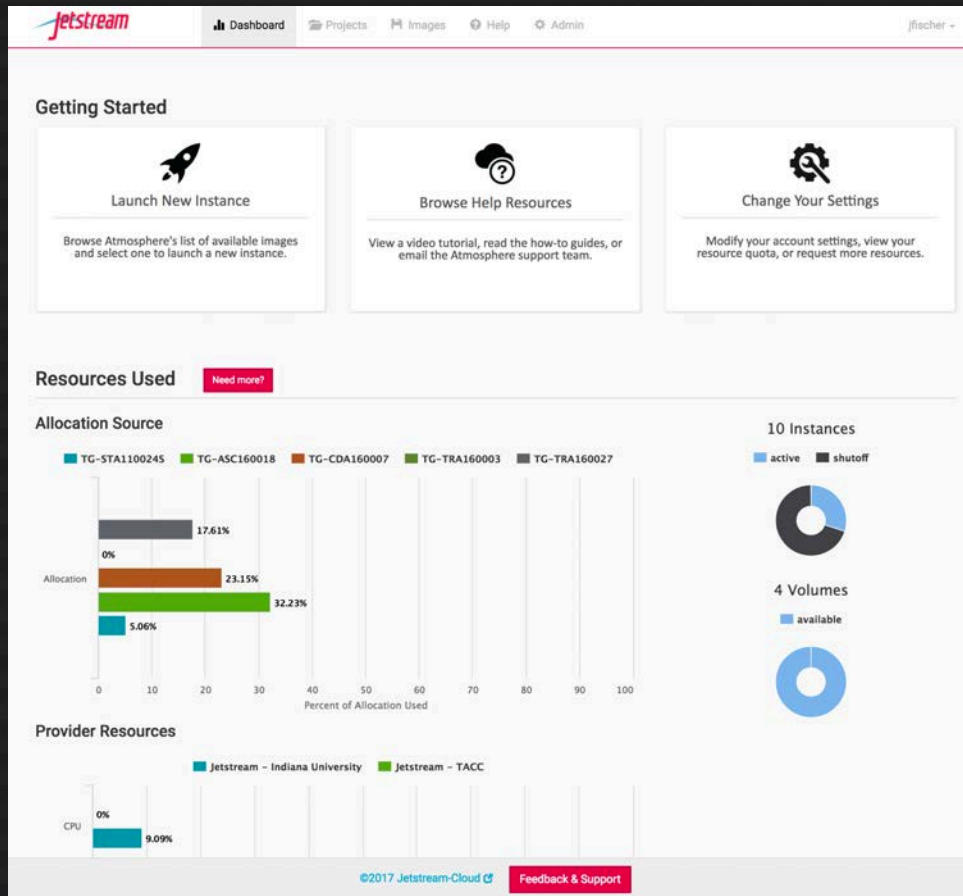
- **Software layers**

- **Atmosphere** web interface
  - library of images, generic, domain specific
  - simplify VM administration
- **OpenStack**: software tools for building and managing cloud computing platforms for public and private clouds.
- **KVM** hypervisor: what the VMs run on
- **Ceph**: storage platform that stores data on a single distributed computer cluster, and provides interfaces for **object**-, **block**- and *file-level* storage.
- **Operating systems**: CentOS, Ubuntu, Windows(?)
- **Applications**; e.g. software developed by the domain specialist, gateways, etc.





# The Jetstream Atmosphere web interface



# API Access to Jetstream

- *What was unexpected*
  - **Demand** for **programmable cyberinfrastructure**
  - Great platform for learning **system administration skills**
  - Great platform for **teaching & learning cloudy technologies**
- *Command line clients*
- *Horizon dashboard very popular; but, incomplete*
- *Programmatic control; python is popular*  
(<https://docs.openstack.org/openstacksdk/latest/>)
- *Slack channel for collaboration API users of Jetstream*



# Jetstream usage highlights – 1 July 2020

- 406 XSEDE projects covering 78 fields of science and over 2100 active users representing 206 institutions
- **80%** of Jetstream users have **not used any other XSEDE system**
- >337M CPU hours allocated to XSEDE projects since June 2016
- 38 active science gateways
- 43 education/teaching allocations serving over 700 students presently and over 4900 through July 2020
- 1189 mean active VMs in previous qtr, 1632 peak active VM count
- **Highest** user satisfaction in most recent XSEDE survey



# Jetstream2 Capabilities

- Enhancing IaaS model of Jetstream:
- Improved orchestration support
- Elastic virtual clusters
- Federated JupyterHubs
- Commitment to **>99%** uptime
- Critical for science gateway hosting
- Hybrid-cloud support
- Revamped User Interface
- Unified instance management
- Multi-instance launch



Feb 12, 2019 – Jet stream region called “Jet N6”  
NASA/JPL-Caltech/SwRI/MSSS/Kevin M. Gill

- **>57K** cores of next-gen AMD EPYC processors
- **>360** NVIDIA A100 GPUs will provide vGPUs via NVIDIA’s MIG feature
- **>18PB** of storage (NVMe and disk hybrid)
- 100GbE Mellanox network





COMMERCIAL  
CLOUD

INTERNET®

INDIANA  
UNIVERSITY  
CYBERINFRASTRUCTURE

XSEDEnet  
Advanced Layer 2  
Services (AL2S) platform

UNIVERSITY  
OF HAWAI'I  
CYBERINFRASTRUCTURE

ARIZONA STATE  
UNIVERSITY  
CYBERINFRASTRUCTURE

REGIONAL

REGIONAL

COMPUTE

8 Nodes  
1,024 Cores  
4 TB RAM

COMPUTE

8 Nodes  
1,024 Cores  
4 TB RAM

STORAGE

869 TB

STORAGE

869 TB

TACC CYBERINFRASTRUCTURE

COMPUTE

8 Nodes  
1,024 Cores  
4 TB RAM

REGIONAL

STORAGE

869 TB

ACCELERATORS

2 Nodes  
1 TB RAM  
8 GPUs

PRIMARY

COMPUTE

416 Nodes  
53,248 Cores  
224 TB RAM

STORAGE

96 Nodes  
15 PB

ACCELERATORS

90 Nodes  
45 TB RAM  
360 GPUs

CORNELL  
UNIVERSITY  
CYBERINFRASTRUCTURE

REGIONAL

COMPUTE

8 Nodes  
1,024 Cores  
4 TB RAM

STORAGE

869 TB

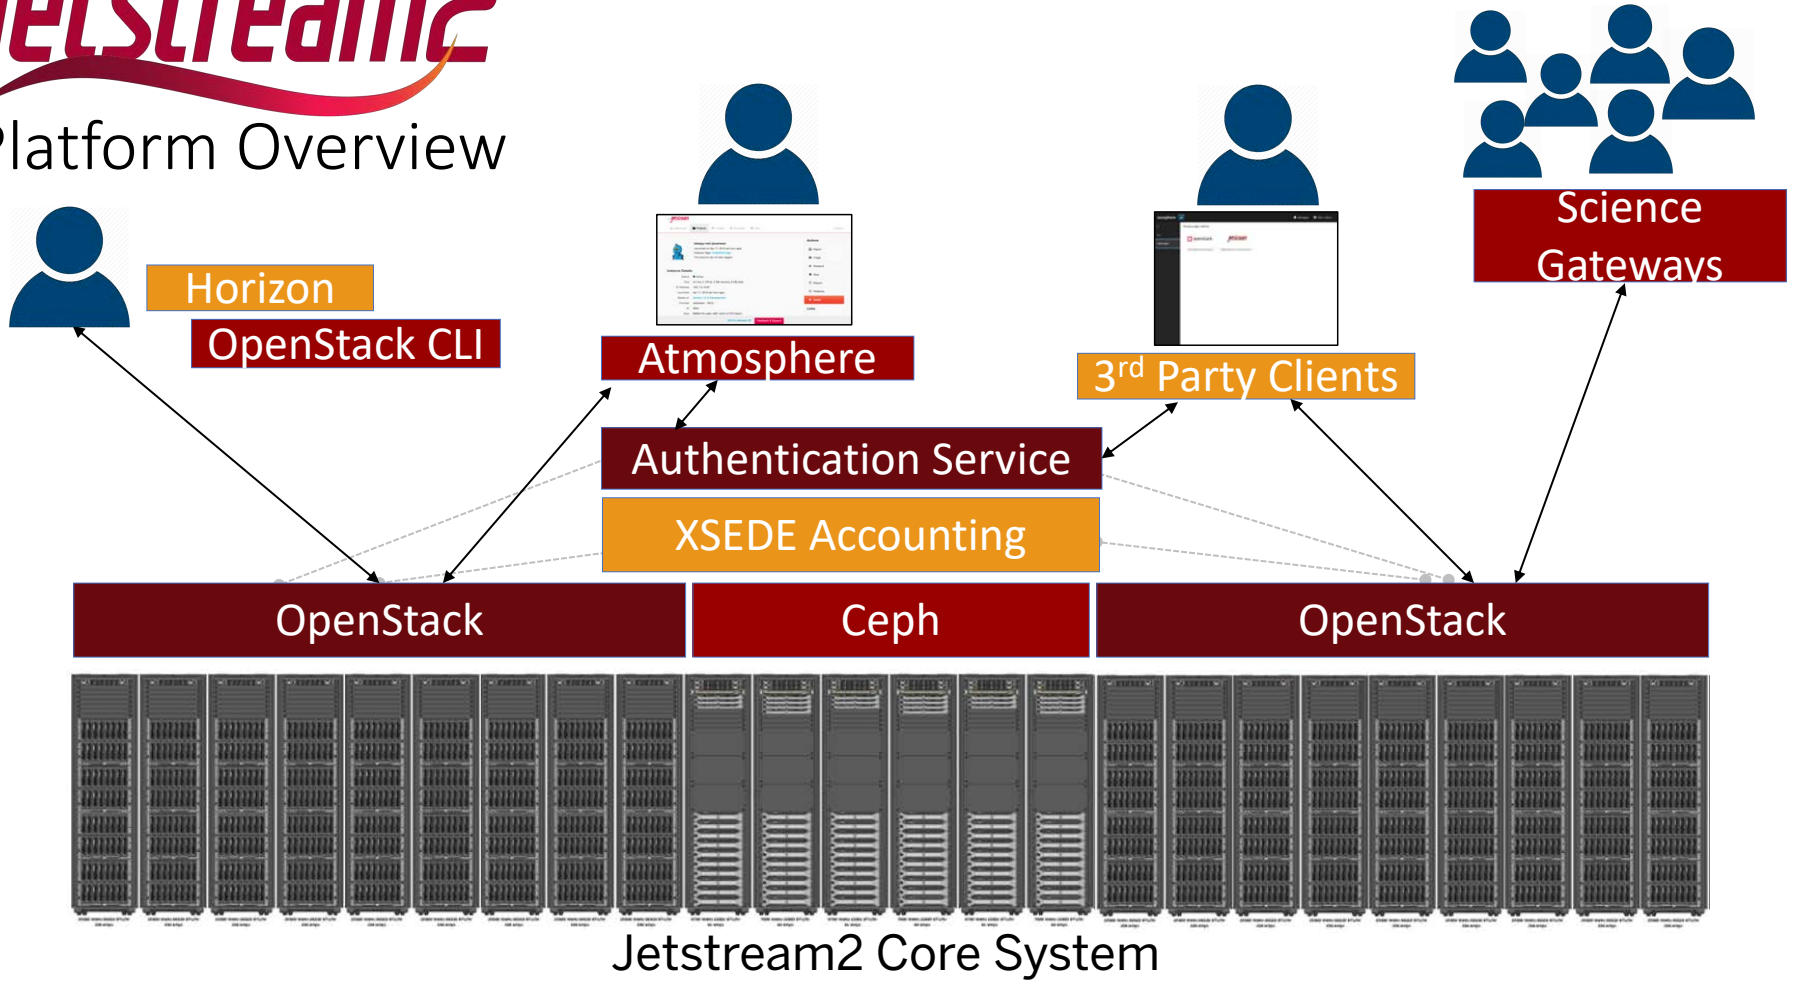


RESEARCH TECHNOLOGIES

UNIVERSITY INFORMATION TECHNOLOGY SERVICES

# Jetstream2

## Platform Overview



# Timeline

- *Jetstream now in 5th year of operations*
- *Jetstream extension requested through November 2021*
- *Jetstream2*
  - *Early operations planned for August 2021*
  - *Production operations by October 2021*



Flickr user Oiluj Samall Zeid - Lejos de Yulín



# Where can I get help?


- *Wiki / Documentation:* <http://wiki.jetstream-cloud.org>
- *API CLI Tutorial:* <https://github.com/jlf599/JetstreamAPITutorial>
- *User guides:* <https://portal.xsede.org/user-guides>
- *XSEDE KB:* <https://portal.xsede.org/knowledge-base>
- *Email:* [help@xsede.org](mailto:help@xsede.org)





# Jetstream Partners

---



funded by the National Science Foundation  
Award #ACI-1445604



# Jetstream2 partners

---



# 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*

- David Y. Hancock, J. Michael Lowe, Therese Miller, Maria Morris, Winona Snapp-Childs, and George Turner*

