

Jetstream Overview

A national research and education cloud

Jeremy Fischer – jeremy@iu.edu

Senior Technical Advisor, Collaboration and Engagement Support
UITS Research Technologies



NSF Funding Areas in HPC

Traditionally concentrated on enabling petascale capability

- Blue Waters – 13.3 petaflops, 2012
- Stampede – 9.6 petaflops, 2013
- Comet – ~2.0 petaflops, 2014

Has funded research into building clouds and computer science

- CloudLab
- Chameleon

Now funding clouds to do research

- Bridges (Hybrid system)
- Jetstream



funded by the National Science Foundation
Award #ACI-1445604



Jetstream - Expanding NSF XD's reach and impact

Around 350,000 researchers, educators, & learners received NSF support in 2015

- Less than 2% completed a computation, data analysis, or visualization task on XD/XSEDE program resources
- Less than 4% had an XSEDE Portal account
- 70% of researchers surveyed* claimed to be resource constrained

Why are the people not using XD/XSEDE systems not using them?

- Perceived ease of access and use
- HPC resources – the traditional view of what XSEDE offers - are often not well-matched to their needs
- They just don't need *that much* capability



funded by the National Science Foundation
Award #ACI-1445604



What is Jetstream and why does it exist?

- NSF's first production cloud facility
- Part of the NSF eXtreme Digital (XD) program
- Provides on-demand *interactive* computing and analysis
- Enables *configurable* environments and architectures
- User-friendly, widely accessible cloud environment
- User-selectable library of preconfigured virtual machines



funded by the National Science Foundation
Award #ACI-1445604



What is Jetstream, continued...

- Focus on ease-of-use, broad accessibility
- Interactive computing and data analysis “on demand”
- Will support persistent gateways (iPlant, Galaxy, generic “SciGAP” build-a-gateway image)
- Reproducibility: Share VMs and then store, publish via IU Scholarworks (DOI)

Who uses Jetstream?

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



funded by the National Science Foundation
Award #ACI-1445604



21st Century Workforce Development

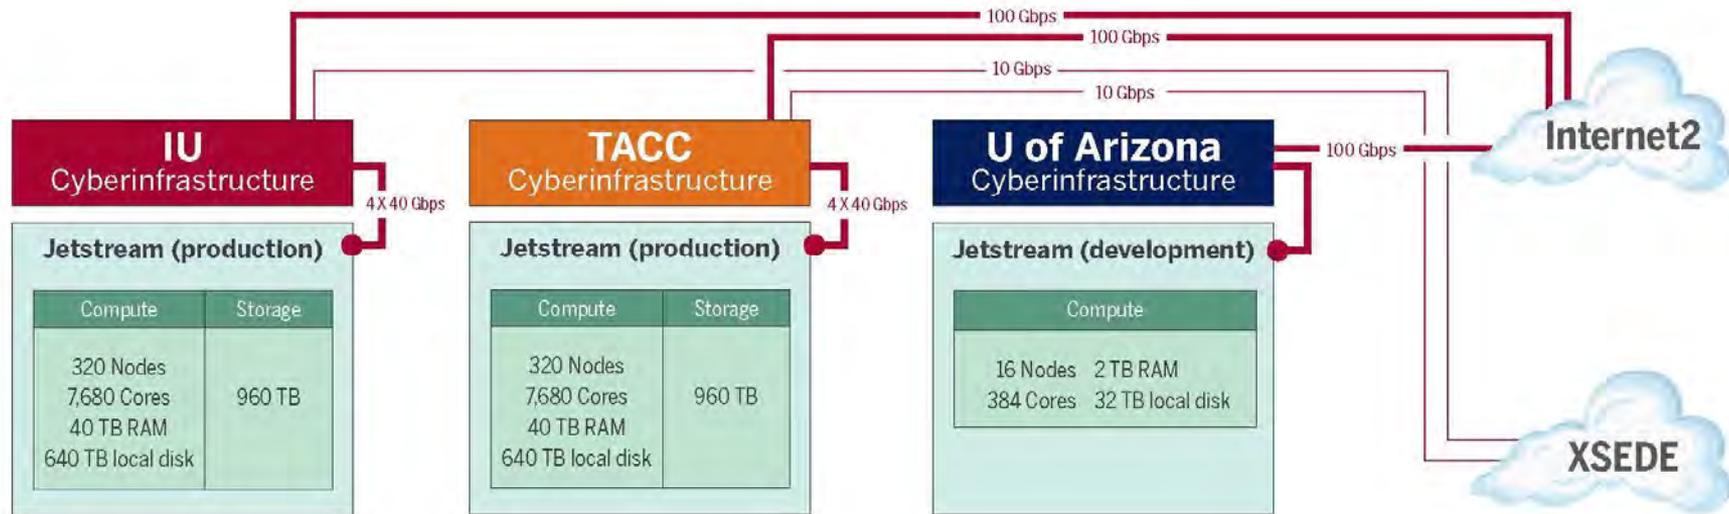
- Specialized virtual Linux desktops and applications to enable research and research education at small colleges and universities
- HBCUs (Historically Black Colleges and Universities)
- MSIs (Minority Serving Institutions)
- Tribal colleges
- Higher-education institutions in EPSCoR States



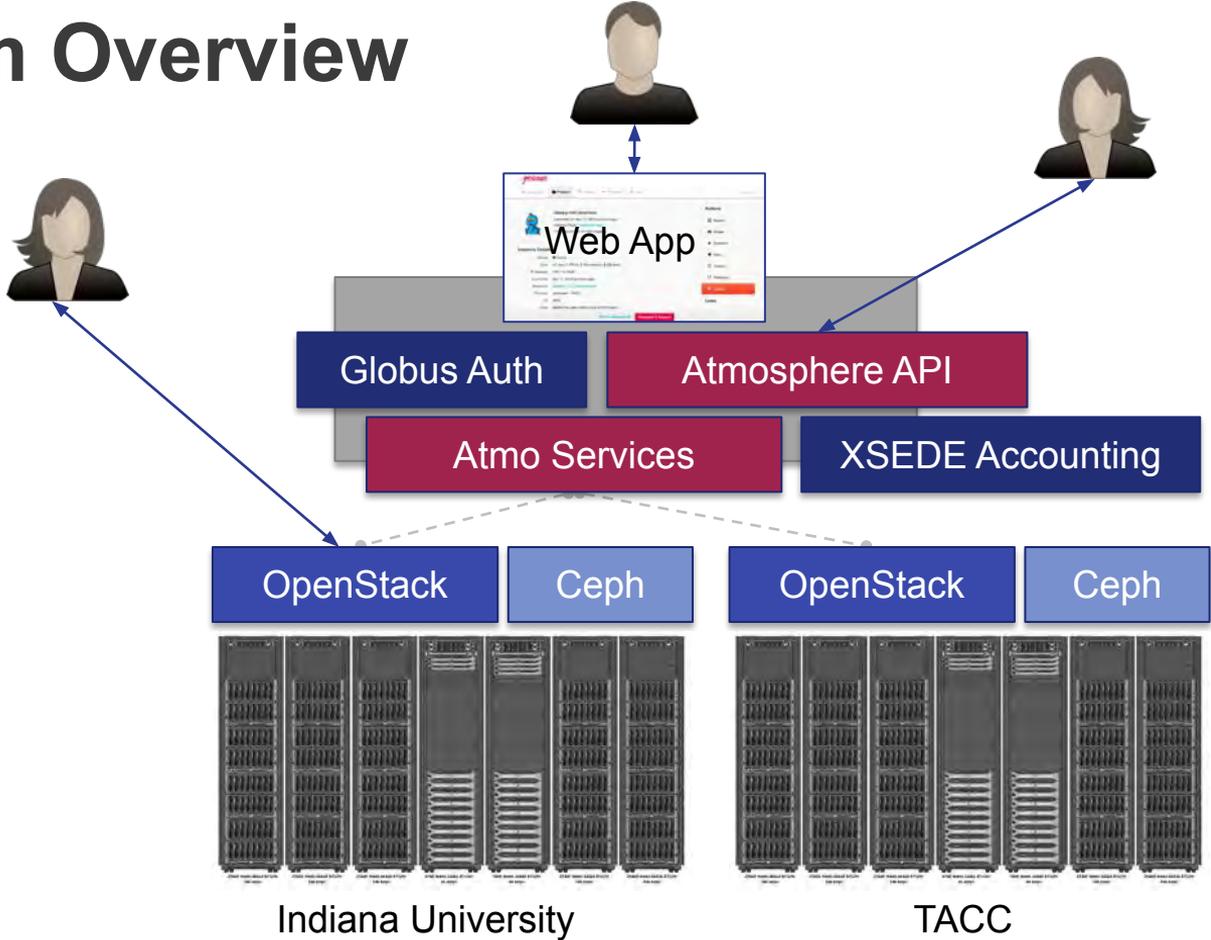
funded by the National Science Foundation
Award #ACI-1445604



Jetstream System Overview



Platform Overview



Hardware and Instance "Flavors"

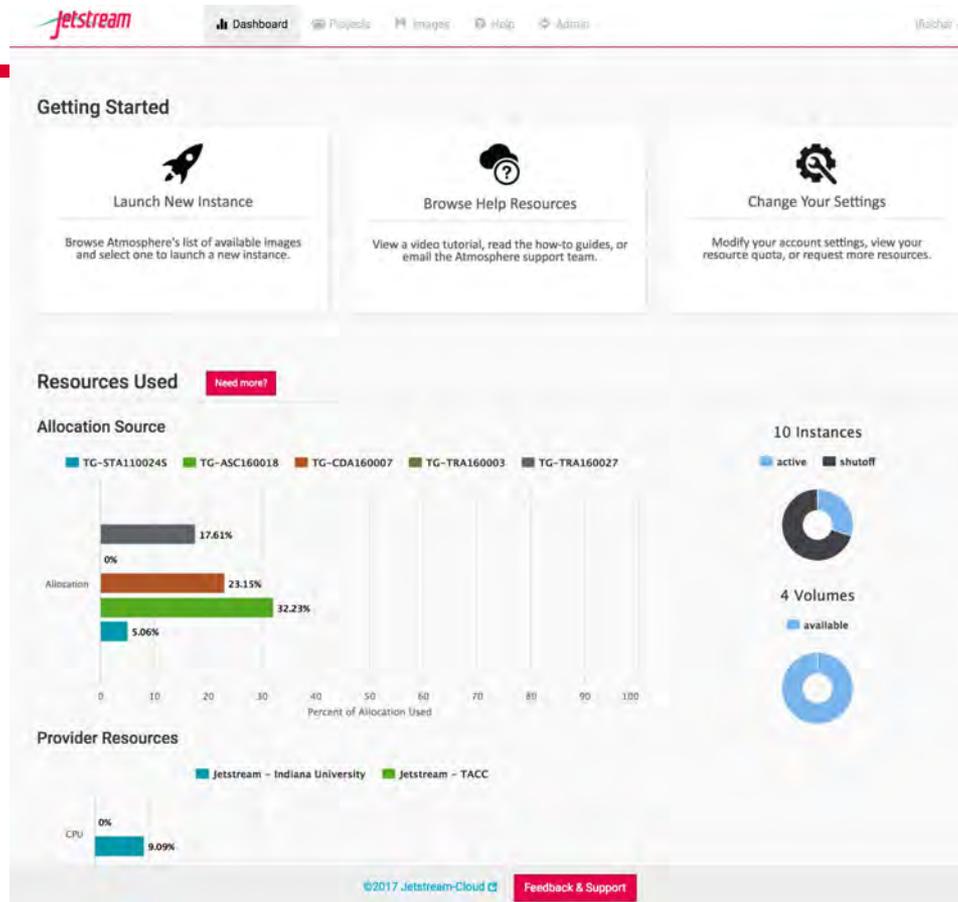
VM Host Configuration

- Dual Intel E-2680v3 "Haswell"
- 24 physical cores/node @ 2.5 GHz (Hyperthreading on)
- 128 GB RAM
- Dual 1 TB local disks
- 10GB dual uplink NIC
- Running KVM Hypervisor

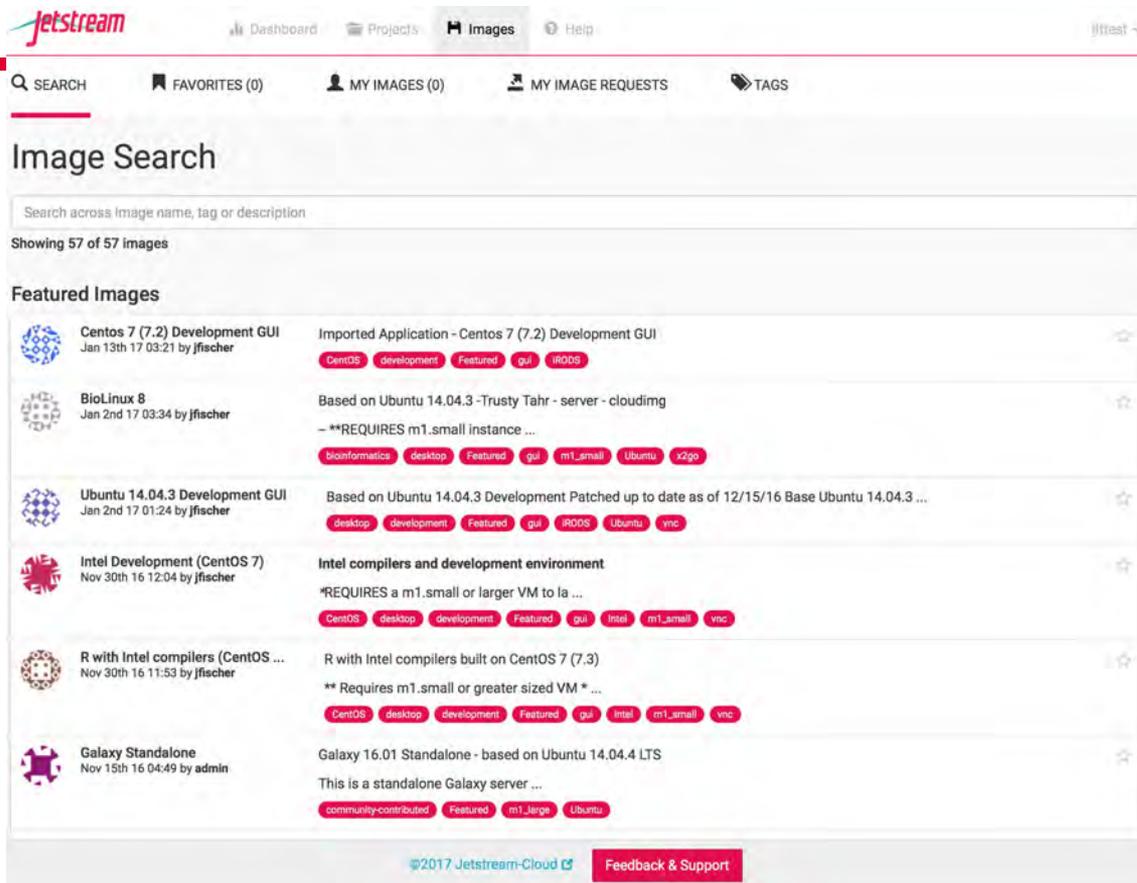
Flavor	vCPUs	RAM	Storage	Per Node
m.tiny	1	2	8	46
m.small	2	4	20	23
m.medium	6	16	60	7
m.large	10	30	120	4
m.xlarge	24	60	240	2
m.xxlarge	44	120	480	1

- Short-term storage comes as part of launched instance
- Long-term storage is XSEDE-allocated
- Implemented on backend as OpenStack Volumes
- Each user gets 10 volumes up to 500GB total storage
- Piloting object storage as well after recent update

The Jetstream Atmosphere web interface



The Jetstream Atmosphere web interface



The screenshot displays the Jetstream Atmosphere web interface. At the top, there is a navigation bar with the Jetstream logo and links for Dashboard, Projects, Images, and Help. Below this is a secondary navigation bar with a search bar and links for FAVORITES (0), MY IMAGES (0), MY IMAGE REQUESTS, and TAGS. The main content area is titled "Image Search" and contains a search input field with the placeholder text "Search across Image name, tag or description". Below the search field, it indicates "Showing 57 of 57 Images". The "Featured Images" section lists several image entries, each with a thumbnail icon, title, author, date, and a list of tags. The entries include:

- Centos 7 (7.2) Development GUI** by jfischer, Jan 13th 17 03:21. Tags: CentOS, development, Featured, gui, @RODS.
- BioLinux 8** by jfischer, Jan 2nd 17 03:34. Tags: bioinformatics, desktop, Featured, gui, m1_small, Ubuntu, x2go.
- Ubuntu 14.04.3 Development GUI** by jfischer, Jan 2nd 17 01:24. Tags: desktop, development, Featured, gui, @RODS, Ubuntu, vnc.
- Intel Development (CentOS 7)** by jfischer, Nov 30th 16 12:04. Tags: CentOS, desktop, development, Featured, gui, intel, m1_small, vnc.
- R with Intel compilers (CentOS ...)** by jfischer, Nov 30th 16 11:53. Tags: CentOS, desktop, development, Featured, gui, intel, m1_small, vnc.
- Galaxy Standalone** by admin, Nov 15th 16 04:49. Tags: community-contributed, Featured, m1_large, Ubuntu.

At the bottom of the interface, there is a footer with the text "©2017 Jetstream-Cloud" and a "Feedback & Support" button.

Using Jetstream VMs

Manipulating Jetstream VMs:

- Jetstream Atmosphere web interface
- Direct API access via OpenStack command line or Horizon access
 - API access enables Science Gateways and other always on services or on demand use cases; e.g. elastic compute techniques

Primary methods of logging into Jetstream VMs to work

- Interactive user access via web interface with VNC/SSH
- Direct VNC/SSH to individual instances



funded by the National Science Foundation
Award #ACI-1445604



HPC vs Cloud

Adapting to a different environment:

- No reservations, no queueing
- More interactive use and less/no batch queuing
- What? No parallel filesystem?!?
- Being your own admin – hey, we have root!
- You really can have almost any (linux) software you want**
- Constantly getting new features (<https://www.openstack.org/software/project-navigator/>)

** Here there be dragons...



funded by the National Science Foundation
Award #ACI-1445604



Jetstream for Education

Jetstream has been used in multiple IU Informatics Graduate Courses

- INFO 535 – Management, Access, and Use of Big and Complex Data
- INFO 590 – Topics in Informatics

BlueWaters Workflow Workshop

Multiple Research Data Alliance Workshops

Upcoming workshops/classes on Galaxy, data analysis in finance using R, security and intrusion detection, and principles in cloud computing



funded by the National Science Foundation
Award #ACI-1445604



Jetstream Timeline...what comes next?

- Transitioned to full operations on September 1, 2016
- As of January 1, 2017: 260 XSEDE projects covering 49 fields of science and 740+ active users representing 148 institutions
- Soliciting Research allocation requests plus Startup and Education allocations – including Science Gateways!
- Adding services as deemed useful/mature (heat, ceilometer, magnum, trove, manila, etc)
- Atmosphere enhancements
- Working on partnerships with groups like HubZero and others to extend the value of Jetstream

Requesting access to Jetstream

- You can request startup allocations anytime. (Startups are simple!)
- You can request allocations for educational use anytime.
- Next submission period for large allocations is 15 March 2017– 15 April 2017.
- We are happy to help you prepare a request and create a successful proposal.
- You do not have to have prior use of Jetstream to be successful.

Where can I get help?

Wiki / Documentation: <http://wiki.jetstream-cloud.org>

User guides: <https://portal.xsede.org/user-guides>

XSEDE KB: <https://portal.xsede.org/knowledge-base>

Email: help@xsede.org

Campus Champions: <https://www.xsede.org/campus-champions>

Training Videos / Virtual Workshops (TBD)



funded by the National Science Foundation
Award #ACI-1445604



Jetstream Fun: Happy cluster / Angry Cluster



Jetstream Partners



INDIANA UNIVERSITY

PERVASIVE TECHNOLOGY INSTITUTE



JOHNS HOPKINS
UNIVERSITY

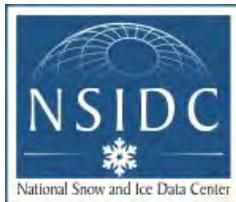


THE UNIVERSITY
OF ARIZONA

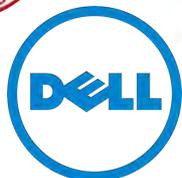
TACC



THE UNIVERSITY OF
CHICAGO



JSU | JACKSON
STATE
UNIVERSITY®



UNC
THE ODUM INSTITUTE



funded by the National Science Foundation
Award #ACI-1445604



Questions?

Project website: <http://jetstream-cloud.org/>

Project email: help@jetstream-cloud.org Direct email: jeremy@iu.edu

License Terms

- Fischer, Jeremy. January 18, 2017. Jetstream Overview – University of Texas San Antonio. Also available at: <http://jetstream-cloud.org/publications.php>
- Jetstream is supported by NSF award 1445604 (Craig Stewart, IU, PI)
- XSEDE is supported by NSF award 1053575 (John Towns, UIUC, PI)
- This research was supported in part by the Indiana University Pervasive Technology Institute, which was established with the assistance of a major award from the Lilly Endowment, Inc. Opinions presented here are those of the author(s) and do not necessarily represent the views of the NSF, IUPTI, IU, or the Lilly Endowment, Inc.
- Items indicated with a © are under copyright and used here with permission. Such items may not be reused without permission from the holder of copyright except where license terms noted on a slide permit reuse.
- Except where otherwise noted, contents of this presentation are copyright 2015 by the Trustees of Indiana University.
- This document is released under the Creative Commons Attribution 3.0 Unported license (<http://creativecommons.org/licenses/by/3.0/>). This license includes the following terms: You are free to share – to copy, distribute and transmit the work and to remix – to adapt the work under the following conditions: attribution – you must attribute the work in the manner specified by the author or licensor (but not in any way that suggests that they endorse you or your use of the work). For any reuse or distribution, you must make clear to others the license terms of this work.



funded by the National Science Foundation
Award #ACI-1445604

