

Jetstream (NSF Award 1445604) Program Year 3 Annual Report (December 1, 2016 – November 28, 2017)

Craig A. Stewart¹
David Y. Hancock¹
Matthew Vaughn²
Nirav Merchant³
Jeremy Fischer¹
J. Michael Lowe¹
Lee Liming⁴
James Taylor⁵
Enis Afgan⁵
George Turner¹
Edwin Skidmore³
Michael Packard²
Brian W. Beck²
Ian Foster⁴

¹Indiana University Pervasive Technology Institute

²University of Texas at Austin, Texas Advanced Computing Center

³University of Arizona

⁴University of Chicago Computation Institute

⁵Johns Hopkins University

Stewart, C.A., Hancock, D.Y., Vaughn, M., Merchant, N., Fischer, J., Lowe, J.M., Liming, L., Taylor, J., Afgan, E., Turner, G., Skidmore, E., Beck, B.W. & Foster, I. (2017). Jetstream (NSF Award 1445604) Year Program Year 3 Annual Report (Dec 1, 2016 – Nov 30, 2017) (PTI Technical Report PTI-TR17-008). Bloomington, IN: Indiana University. Retrieved from <http://hdl.handle.net/2022/21854>.



INDIANA UNIVERSITY

PERVASIVE TECHNOLOGY INSTITUTE

1. Accomplishments	3
1.1. What are the major goals of the project?	3
1.2. What was accomplished under these goals?	3
1.2.1. Major Activities	3
1.2.2. Specific Objectives:	5
1.2.3. Significant results	8
1.2.4. Key outcomes or Other achievements	12
1.4. How have the results been disseminated to communities of interest?	14
1.5. What do you plan to do during the next reporting period to accomplish the goals?	24
2. Products (resulting from this project during the specified reporting period)	25
3. Participants	31
3.2. Partner organizations	32
3.3. Have other collaborators or contacts been involved?	35
4. Impact	35
4.1. What is the impact on the development of the principal discipline(s) of the project?	35
4.2. What is the impact on other disciplines?	36
4.3. What is the impact on the development of human resources?	37
4.4. What is the impact on physical resources that form infrastructure?	37
4.5. What is the impact on institutional resources that form infrastructure?	37
4.6. What is the impact on information resources that form infrastructure?	37
4.7. What is the impact on technology transfer?	38
4.8. What is the impact on society beyond science and technology?	38
5. Changes/ Problems	38
5.1. Changes in approach and reasons for change	38
5.2. Actual or Anticipated problems or delays and actions or plans to resolve them	39
5.3. Changes that have significant impact on expenditures	44
5.4. Significant changes in use or care of human subjects	44
5.5. Significant changes in the use or care of vertebrate animals	44
5.6. Significant changes in the use or care of biohazards	44

1. Accomplishments

1.1. What are the major goals of the project?

Jetstream is a configurable large-scale computing resource that leverages both on-demand and persistent virtual machine technology to support a much wider array of software environments and services than current NSF resources can accommodate. As a fully configurable "cloud" resource, Jetstream bridges a major gap in the current ecosystem, which has machines targeted at large-scale High-Performance Computing, high memory, large data, high throughput, and visualization resources. As the open cloud for science, Jetstream provides:

- "Self-serve" academic cloud services, enabling researchers or students to select a VM image from a published library, or alternatively to create or customize their own virtual environment for discipline- or task-specific personalized research computing.
- Hosting of persistent VMs to provide services beyond the command line interface for science gateways and other science services.
- New modes of sharing computations, data, and reproducibility.
- Expanded access to the NSF XSEDE (the eXtreme Science and Engineering Discovery Environment) ecosystem by making virtual desktop services accessible from institutions with limited resources.

While Jetstream is a hardware acquisition, the system is both a pilot and production system that is made novel by the software components and interfaces. The overarching goal of the Jetstream team is to continue to improve upon existing functionality throughout the duration of the award by increasing stability and adding new features. The ability at which the latter can be done will ultimately be determined by the management and operations budget along with supplemental awards.

1.2. What was accomplished under these goals?

1.2.1. Major Activities

Jetstream was recommended for PY2 funding after successfully completing a review of PY1 operations. The review panel's executive summary accurately reflects the state of the project.

The review panel met with the Jetstream team at the Cyberinfrastructure Building at Indiana University, July 19-20, 2017. The purpose of the review was to assess the Jetstream 1st year of production operations. The panel acknowledges that the Jetstream team has done a commendable job in bringing online and operating an innovative national cloud computing resource that is successfully engaging non-traditional research

computing communities. The panel appreciated the user-focused nature of the Jetstream team and their focus on ease of use and empowering non-traditional computational users. The panel believes that all Jetstream award expectations have been met and highly recommends that the funding be continued for a second year of production operations.

- The Jetstream Stakeholder Advisory Board (SAB) held their first meeting in Chicago on during February 2017. The SAB report was used as input to shape future Jetstream activities and prepare the team for the PY1 operations review.
- The Jetstream team conducted a webinar for outreach purposes to provide an overview of Jetstream and assist with the XSEDE application process for research allocations. The recording can be viewed at <https://youtu.be/q-o8paUT4KY>.
 - The webinar was used as the basis for a Cornell Virtual Workshop on allocations. The module is now available from the XSEDE portal training page, Cornell, and Jetstream websites: <https://cvw.cac.cornell.edu/JetstreamReq/>
- The Jetstream team hosted four REU students in the summer of 2017 through a supplemental award. One was from UT San Antonio, one was from Bethune-Cookman University, one was from University of Washington, one will be matriculating at IU in the fall. The main topics covered were: interoperability; efficient utilization of Jetstream for data analytics tasks; best practices with big data analytics using Jetstream.
- The Jetstream team received a supplemental award for 2-years to provide for an additional education, outreach, and training position. The supplement also provides funding for in person SAB meetings. This supplement has filled a need that enhances existing Jetstream partner effort. Sanjana Sudarshan was hired in August 2017 and is already having a positive impact on user support and outreach.
- The Jetstream team conducted multiple surveys during this period to help inform current and future CI needs. The targeted groups were Biological Field Stations, Educators that had used Jetstream for courses, and NSF-funded Engineering PIs.
- Jetstream has kept pace with the OpenStack release schedule by performing updates at both production sites on a regular basis. Currently running Ocata with Pike in testing Jetstream has performed upgrades at twice the rate originally planned to provide the best user experience and ensure the environment is well supported.
- The Jetstream web portal based on Atmosphere has received regular monthly updates and during October 2017 maintenance period important new beta features were released. These included an updated webshell and browser-based VNC client. These updates address multiple issues and features users have encountered and requested. After more users move to the new tools the existing webshell and VNC clients will be deprecated.
- The Jetstream Trial Allocation system went online mid June 2017. It allows users to get a modest allocation of 2000 SUs using m1.small or m1.tiny instances on each cloud by just creating an XSEDE user portal account. We saw 71 trial allocation users in the first quarter that it was available (PY2 Q2). At the close of Q4, there are now 183 trial allocation users.

- The project has enabled or supported significant new modalities of use of NSF resources. The publication on Portable Learning Environments demonstrates that use within the classroom and for workshops. Multiple science gateway publications also highlight the importance of Jetstream by providing virtual clusters to GenApp integration on OpenStack to Galaxy for Proteogenomics research.
- The Jetstream team through our partner at the University of Chicago team made regular, minor fixes and improvements to the Globus authentication services that Jetstream uses for Web authentication. They've provided weekly service updates without interrupting Jetstream availability and provided routine technical and user support for these authentication services.
- The Jetstream team through our partner at Johns Hopkins University has continued to support and expand Galaxy use on Jetstream:
 - Continued to burst significant portion of usegalaxy.org jobs to instances running on Jetstream, having run more than 92,500 jobs on behalf of more than 7,200 users while requiring no input from those users or altering their typical workflow.
 - Enabled use of Galaxy Interactive Environments from usegalaxy.org. This is a new Galaxy application feature that allows users to start instances of programmatic runtime environments within their Galaxy workspace, opening doors to exploratory analysis previously not possible.
 - Continued to provide updates to a standalone image of Galaxy for the self-serve mode that allows researchers to install custom tools and avoid quotas imposed on usegalaxy.org.
- The Jetstream team has made outreach to conferences, universities, and small workshops a significant priority. The outreach efforts for the current year include 47 events, covering 16 states plus Washington D.C. The lessons learned from this year of operations is helping the team re-focus efforts on locations where the team can have the most impact on new and underserved communities.

1.2.2. Specific Objectives:

Accomplishments relevant to achievement of goals for this project are described below:

Indiana University

- The technical team worked with vendors to secure NVMe devices for ingestion and processing of metrics from ceilometer. The original setup caused periodic performance issues on the Ceph storage environment and limiting our ability to collect fine-grained CPU utilization data. This was resolved at both IU and TACC during this project year.
- Repaired multiple water-cooled doors with recurring issues with sensors and hoses.
- Configured a new Intrusion Detection System for inspecting packets from Jetstream.
- Performed multiple OpenStack updates, currently at Ocata.
- Performed general hardware maintenance, updated firmware, and security software updates.

- Cloud-init scripts have been implemented on the IU side of Jetstream to automatically apply updates to VMs after they're started (in addition to the practice of regularly updating featured VMs). If the time/impact of this change remains reasonable the team will replicate the functionality at TACC. Currently the major complaint is that this increases launch times for new instances.
- The Jetstream team at IU successfully facilitated use of Wrangler and Jetstream together to support a half-dozen research projects with joint allocations. Users and science gateways can access Wrangler via dedicated NFS gateway services for larger storage volumes.
- The Jetstream team submitted a project improvement fund request to XSEDE to create a cloud library that can be used to deploy images at multiple resources providers using cloud-init. The proposal was approved but only for 25% of requested funding. The Jetstream team will work with partners from PSC, Cornell, and UC/ANL on implementing a searchable library and testing orchestration templates between multiple sites.

TACC

- Performed multiple OpenStack updates, currently at Ocata
- Performed general hardware maintenance, updated firmware, and security software updates.
- Improved reliability of integration with XSEDE AMIE system to support the Trial Application mechanism
- Developed protocols for power failure response that decreased downtime upon subsequent power events.
- Developed protocols for providing special connectivity for other TACC resources including Corral, Wrangler, and Stockyard.

University of Arizona

- Installed and incorporated the test and development openstack system at the University of Arizona as part of the Atmosphere development process for new OpenStack and Globus features.
- Implemented Atmosphere improvements to the existing integration with XSEDE allocations, including additional reporting, support for special allocations, and improved synchronization and validation on utilization with XSEDE
- Implemented "Trial Allocation" access within Atmosphere, providing a low-barrier entry to Jetstream
- Exposed OpenStack instance "shelve" within Atmosphere, a more resource-friendly way for storing inactive instances
- Enhanced imaging and the image catalog, including instance size selection based on image metadata, providing metrics about instance launches, restricting image access by email address or account (via Atmosphere API), and a new plugin framework to allow site operations to define custom logic in the imaging and validation workflow.
- Enhanced volumes through Atmosphere, including support for ext4 volumes.

- Implemented an Atmosphere API to provide reporting on Atmosphere data on instances, images, and users
- Enhanced Atmosphere's web desktop and web shell
- Enhanced Atmosphere user-defined boot scripts for instances, including user-defined execution (i.e. upon first launch or upon each deployment) and synchronous/asynchronous execution.
- Implemented numerous bug fixes and enhancements to stability, performance, security, and user experience

Johns Hopkins University (JHU)

- Integrated use of Jetstream resources with usegalaxy.org to expand the job processing capacity of that science gateway and enable new modes of computation via the web
- Facilitate self-serve model of utilizing Jetstream resources with Galaxy via a standalone image
- Delivered a number of training events throughout the year that utilized or relied on resources provisioned from Jetstream (e.g., a week-long Galaxy Admin training course)
- Developed three Python libraries for managing resources on Jetstream
 - SwarmScale: a library for managing Docker Swarm cluster based on Jetstream resources. It enables robust utilization of Jetstream resources from usegalaxy.org to power Galaxy Interactive Environments (GIEs);
 - SlurmScale: a library for managing resources composing a Slurm cluster (e.g., provision, configure, integrate, decommission), paving a path for more efficient utilization of Jetstream resources from usegalaxy.org science gateway;
 - CloudBridge: a multi-cloud library with a consistent API for multiple cloud providers allowing users to uniformly utilize resources from an OpenStack (Jetstream), Amazon, and Azure clouds.
- Developing an API and a portal (CloudLaunch) for launching Galaxy instances and, in the future, compute resources to be attached to a running Galaxy service.
- Facilitated use of Jetstream via Galaxy for a course taught at JHU that focuses on bioinformatics data analysis: Bioinformatics: Tools for Genome Analysis (course 410.635) <http://advanced.jhu.edu/academics/graduate-degree-programs/bioinformatics/course-descriptions/>). Among nine students and the instructor they consumed more than 280,000 SUs from March through May 2017. They each launched and managed their own instances.
- Submitted a major NIH proposal that will, if funded, expand the capabilities of Galaxy to more easily consume resources from Jetstream on behalf of individual researchers.

University of Chicago

In order to better support the Jetstream system and allow for more comprehensive integration testing, University of Chicago added a Preview environment that Jetstream has been using to test future releases of Globus authentication services. This advanced testing has ensured that

Jetstream’s web login interface, provided by Globus, works reliably and continuously for Jetstream users despite regular, incremental changes.

Cornell University

- Developed an online training module, “Jetstream Allocations”. This topic covers how to apply for an XSEDE allocation on the Jetstream computing resource. It includes details about Jetstream itself, a description of the necessary documents which comprise an application, and advice about how to prepare a successful allocation request. March 2017.
- Developed an online training module, “Introduction to Jetstream”. This topic is an Introduction to the Jetstream system and the Atmosphere cloud computing environment, covering how to create, use, and archive virtual machines on Jetstream. December 2017.

1.2.3. Significant results

We have the following significant technical results so far:

Metric	Goal per program year	Q1 Dec 2016 - Feb 2017	Q2 Mar 2017 - May 2017	Q3 Jun 2017 - Aug 2017	Q4 Sep 2017 - Nov 2017	Achieved Dec 2016 - Nov 2017	Notes
System availability (uptime of an element of the production hardware, as % of wall clock time)	95%	100%	100%	100%	100%	100%	
Capacity availability (% of the total capacity of Jetstream available for NSF use over time)	95%	99.33%	99.36%	99.64%	99.31%	99.41%	

Job completion success - % of jobs submitted should complete without having to be resubmitted as a result of a failure in the hardware or system software.	96%	98.00%	97.43%	96.77%	99.08%	97.82%	
Average number of active VMs	320	560	633	630	864	672	
CPU % utilization	6%	2.27%	3.11%	4.95%	5.43%	3.94%	IU cloud only. Acquisition of statistics at TACC has been problematic
Core cloud environment software will be upgraded to match current versions components such as operations systems and cloud software environments	Updated prudently, generally within ≤ 12 months of major releases	Newton	Newton	Newton	Ocata	Ocata	

Additional statistics of note:

Metric	Goal per program year	Q1 PY2	Q2 PY2	Q3 PY2	Q4 PY2	Notes
		Dec-Feb	Mar-May	Jun-Aug	Sep-Nov	
Capacity of	90%	>100%	>100%	>100%	>100%	

system allocated via NSF-specified allocation process						
Total number of distinct users	1,000	1,576	1,752	2,014	2,261	annual goal already exceeded
Total number of students having used Jetstream in an educational or training setting	100	397	534	604	820	annual goal already exceeded
Total number of science gateways using Jetstream (running total)	2	11	13	9	9	annual goal already exceeded
SUs available to user community per month	>1.47M	>1.47M	>1.47M	>1.47M	>1.47M	
% of SUs available to user community that were used per month	Goal not established	>100%	80.29%***	>100%	>100%	***There is a reporting problem with TAS not uploading complete data to XSEDE DB for April & May 2017

Total number of VM images and/or data sets published with a DOI via IUScholarWorks	10	1	0	1	2
--	----	---	---	---	---

Allocation statistics for September through November 2017:

	Q1 PY2 Dec-Feb 2017 Total	Q2 PY2 Mar-May 2017 Total	Q3 PY2 Jun-Aug 2017 Total	Q4 PY2 Sep-Nov 2017 Total	PY2 Totals
<i>Startup</i>					
Total requests	23	34	40	36	133
SUs requested	2,100,000	2,590,000	2,885,120	3,116,240	10,691,360
SUs awarded	2,100,000	2,590,000	2,885,120	3,116,240	10,691,360
<i>Educational</i>					
Total requests	6	9	11	10	36
SUs requested	877,000	798,840	1,250,000	1,648,140	4,573,980
SUs awarded	877,000	798,840	1,250,000	1,648,140	4,573,980
<i>Campus Champion/Staff</i>					
Total requests	29	32	46	33	140
SUs requested	1,450,000	1,600,000	2,728,000	1,650,000	7,428,000
SUs awarded	1,450,000	1,600,000	2,728,000	1,650,000	7,428,000
<i>Supplemental/Discretionary</i>					
Total requests	3	5	2	0	10

SUs requested	1,000,000	1,351,000	2,050,000	0	4,401,000
SUs awarded	1,000,000	1,351,000	2,050,000	0	4,401,000
<i>Research allocations</i>					
Total requests	N/A	6	5	3	14
SUs requested	N/A	3,884,528	2,408,720	6,500,000	12,793,248
SUs awarded	N/A	12,884,528	2,408,720	2,000,000	17,293,248
Total request and allocations					
Total requests	61	86	104	82	333
Total SUs requested	5,427,000	10,224,368	11,321,840	12,914,380	39,887,588
Total SUs awarded	5,427,000	19,224,368	11,321,840	8,414,380	44,387,588

1.2.4. Key outcomes or Other achievements

New (major) features:

- Implementation and integration of a new Intrusion Detection System at IU.
- Support for “shelving” instances was added to both IU and TACC with push-button functionality in the Jetstream portal.
- Wrangler mounts via NFS were enabled at IU for jointly allocated projects.
- Trial allocations were deployed for new users with ~4 hour turnaround time.
- New web terminal and remote desktop functionality was deployed within the Jetstream portal.
- Interactive Environments for Galaxy users were added.
- Users may use the “ezj” command to launch their own private Jupyter notebook after launching an instances within the Jetstream portal.
- The object storage gateway was made available at both IU and TACC.
- The Heat orchestration services was made available at both IU and TACC.
 - Users are actively leveraging this service to deploy virtual clusters or integrate with container orchestration engines such as Kubernetes.

New Research Allocations:

Award #	PI Name	Institution	SUs Awarded	Storage Awarded
MCB140147	James Taylor	Johns Hopkins University	2,000,000	N/A
BIO150062	Nirav Merchant	University of Arizona	1,000,000	1TB
MCB170057	Emre Brookes	University of Texas Health Science Center at San Antonio	628,000	26.5TB
SES170012	Patrick Brandt	University of Texas Dallas	630,720	N/A
CDA170003	Mark Perri	Sonoma State University	150,000	N/A
ASC170002	Burke Mamlin	Indiana University	1,086,240	1.26TB
MCA05S028	Aleksei Aksimentiev	University of Illinois at Urbana-Champaign	5,000,000	2TB
MCB070039N	Borries Demeler	University of Texas Health Science Center at San Antonio	250,000	N/A
MCB160013	Thomas Miller	California Institute of Technology	1,000,000	1TB
MCB170016	Niles Pierce	California Institute of Technology	4,000,000	1TB
MSS170003	Liming Xiong	Iowa State University	1,548,288	N/A
Total Awards			17,293,248	32.76TB

1.3. What opportunities for training and professional development has the project provided?

The Jetstream team has enhanced the training and development of many CI professionals, domain scientists, and students through outreach efforts. The outreach events under major activities are a key driver for the Jetstream team and will continue in full force during the next project year.

1.4. How have the results been disseminated to communities of interest?

Results have been communicated to communities of interest in the following ways (which are listed in the Products section): published papers in peer-reviewed technical conferences and journals, and gave conference presentations. In addition, Jetstream team members engaged in the following activities:

Outreach Efforts for 2017

Date	Event Name	Location (City, State)	Presenter(s)/ Attendees	Notes, Link to slides/materials, etc.
Jan 4	Jetstream for RStudio and GIS Modeling	USDA Agricultural Research Service, Southwest Watershed Research Center	Tyson Swetnam (CyVerse)	https://www.ars.usda.gov/pacific-west-area/tucson-az/southwest-watershed-research/
Jan 8	Jetstream for RStudio and CyVerse DataStore	Botanical Information and Ecological Network (BIEN)	Tyson Swetnam (CyVerse)	http://bien.nceas.ucsb.edu/bien/

Jan 9-12	MiniCourse in Statistics	Brandeis University, Waltham, MA	Francesco Pontiggia	https://kb.brandeis.edu/display/SCI/MiniCourse+in+Statistics https://kb.brandeis.edu/display/SCI/Resources+and+Computational+Tools https://kb.brandeis.edu/display/SCI/Jetstream+Cloud
Jan 14-18	PAG 2017	San Diego, CA	Matt Vaughn	
Jan 17-20	UTSA Meeting	San Antonio, TX	Jeremy Fischer, George Turner, Mike Lowe, Dave Hancock, Brian Beck, Mike Packard, Ashley Bucholz, Andy Lenards	Agenda - UTSA Meeting - Jan. 17-20, 2017
Jan 26	Data-driven Neuroscience with Python	UCSF, San Francisco, CA	Chris Holdgraf, Fatma Deniz, Ariel Rokem, Maryana Alegro, and Aaron Culich	https://bids.berkeley.edu/news/berkeley-and-uw-data-scientists-team-ucsf-researchers-deliver-data-driven-analysis-and-machine http://grinberglab.ucsf.edu/python
Feb 14-16	CyVerse / TERRA-REF / Phenome Hackathon	Tucson, AZ	Tyson Swetnam (CyVerse)	https://github.com/bjoyce3/Image_Data_rapid_QAQC
Feb. 20	AAAS Conference	Boston, MA	Brian Beck	https://aaas.confex.com/aaas/2017/webprogram/meeting2017-02-20.html

Feb 24	WEBINAR: WQ-MAKER: A flexible, scalable genome annotation pipeline on Jetstream Cloud	Online	Upendra Kumar Devisetty	http://www.cyverse.org/blog/events/webinar-wq-maker-flexible-scalable-genome-annotation-pipeline-jetstream-cloud YouTube Video of Webinar
Feb 27-Mar 3	SIAM CSE17 Broader Engagement	Atlanta, GA	Jeremy Fischer	http://shinstitute.org/broaden-engagement-at-siam-cse17/
Mar 18	Invited talk to Discovery Saturday	Santa Rita Experimental Range, Green Valley AZ	Tyson Swetnam (CyVerse)	http://prezi.com/ba-ij00save/?utm_campaign=share&utm_medium=copy&rc=ex0share https://cals.arizona.edu/sref/
Mar 25-28	GalaxyP Workshop at ABRF 2017	San Diego, CA	Tim Griffin, Jim Johnson, Pratik Jagtap	Workshop information Presentation Slides
Mar 27-30	DellHPC	Austin, TX	George Turner, Mike Lowe	Presentation Slides
Mar 31	Texas A&M Research and IT talk and tutorial	College Station, TX	Jeremy Fischer	https://jetstream-cloud.org/files/Jetstream-TexasAM-Presentation.pptx
April 11	GlobusWorld 2017	Chicago, IL	Lee Liming	https://www.globusworld.org/program Note especially the Lightning Talk "Login with XSEDE and Jetstream: Our Experiences with Globus Auth" https://www.globusworld.org/files/2017/06_XSEDE_Jetstream_Liming.pdf

April 25	LSST Data Science Fellowship		Nirav Merchant	Use of NSF Resources – link TBD https://www.lsst.org/news/raining-tomorrow UA presented a brief intro for Jetstream and some students utilized Jetstream during the impromptu hackathon opting for their own dockerized tools on Jetstream
April 26	Open Science Forum	Bloomington, IN	George Turner	Describe collaboration between Jetstream and IVMOOC Presentation Slides
April 28	Society of American Foresters Sectional Meeting	Flagstaff, AZ	Tyson Swetnam	https://prezi.com/view/DKmuPtvGJHfTolrerRuW/
May 2-3	Northern Lights talk	Minnesota, MN	Jeremy Fischer	https://sites.google.com/a/umn.edu/mcm-2017/meeting-agenda
May 3-4	Jackson State Outreach Event	Jackson, MS	David Walling	R/R Studio Using Jetstream – assisted with images. David did instruction.
May 11	Boston Open Research Cloud Workshop	Boston, MA	George Turner, Andrew Lenards, Jeremy Fischer, Mike Lowe, Steve Bird (attendees)	Panel discussion Presentation Slides

May 18-19	South Carolina State University Outreach Event	Orangeburg, SC	George Turner	http://span.scsu.edu Presentation Slides
May 31-Jun 2	University Corporation for Atmospheric Research	Boulder, CO	George Turner	http://www.unidata.ucar.edu/events/2017CloudModelingWorkshop/#home Presentation Slides
June 1-2	Great Plains Network Annual Meeting	Kansas City, MO	Jeremy Fischer	http://conferences.k-state.edu/gpn-gwla/ https://jetstream-cloud.org/files/Jetstream-GPN-Annual-Conference-Presentation.pptx
June 4-8	American Society for Mass Spectrometry (ASMS) annual conference	Indianapolis, IN	Jeremy Fischer	https://www.asms.org/conferences/annual-conference/annual-conference-homepage https://jetstream-cloud.org/files/Jetstream-ASMS-Intro-Presentation.pptx
June 5-9	PoreCamp USA	College Station, TX	Matt Vaughn, Nirav Merchant, and PoreCamp USA Team	https://porecamp.github.io/
June 16	Science Gateways Weekly Call	Online / Bloomington, IN	Jeremy Fischer	https://jetstream-cloud.org/files/Jetstream-Gateways-Call-June2017-Presentation.pptx
June 18 - July 21	Data Intensive Biology Summer Institute for Sequence Analysis at UC Davis	Davis, CA	C. Titus Brown, Upendra Kumar Devisetty	http://ivory.idyll.org/blog/2017-dibsi-xsede-request.html

June 21	NetSci 17	Indianapolis, IN	George Turner	
June 26-30	2017 Metagenomics Workshop	San Diego, CA	Rob Edwards	https://edwards.sdsu.edu/research/2017-metagenomics-workshop/
July 9-13	PEARC17	New Orleans, LA	Jeremy Fischer, Tassie Gnaidy, George Turner, Mike Lowe, Eric Coulter	https://jetstream-cloud.org/files/Jetstream-PEARC17-Presentation.pptx https://github.com/ECoulter/Tutorial_Practice
July 24-28	Computational Sciences in the Clouds	Austin, TX	George Turner, Mike Packard, (partial Brian Beck)	https://www.tacc.utexas.edu/education/institutes/computational-science-in-the-cloud
July 24-27	ESIP Summer Meeting	Bloomington, IN	Jeremy Fischer	http://www.esipfed.org/meetings/upcoming-meetings/esip-summer-meeting-2017 https://jetstream-cloud.org/files/Jetstream-ESIP-Intro-Presentation.pptx
Sept 8	Big Data Neuroscience Workshop (ACNN)	Bloomington, IN	Jeremy Fischer, Sanjana Sudarshan	http://www.neurosciencenetwork.org/ACNN_Workshop_2017.html
Sept 11-12	University of Arkansas Fayetteville - Jetstream + NCGAS Workshop	Fayetteville, AR	Jeremy Fischer, Sanjana Sudarshan, Tom Doak, Carrie Gannote, Bhavya Papudeshi	https://jetstream-cloud.org/files/Jetstream-UArk-Presentation.pptx

Sept 20	Purdue Jetstream Tutorial	West Lafayette, IN	Jeremy Fischer / Sanjana Sudarshan / Mike Lowe	https://jetstream-cloud.org/files/Jetstream-Purdue-Presentation.pptx
Sept 20-24	2017 Annual Organization of Biological Field Stations	Iltasca, Minnesota	Craig Stewart/Tom Doak	http://www.obfs.org/annual-meeting
Sept 28	School of Informatics, Computing & Engineering	Bloomington, IN	Craig A. Stewart & David Y. Hancock	http://hdl.handle.net/2022/21845
Oct 1-4	SIGUCCS17	Seattle, WA	Jeremy Fischer	Jetstream in Education paper presentation https://jetstream-cloud.org/files/Jetstream-SIGUCCS17-Presentation.pptx
Oct 10-11	UTSA Jetstream Workshop	San Antonio, TX	Jeremy Fischer, George Turner, Sanjana Sudarshan, Mike Lowe (all staff will be present for technical meeting following workshop)	https://jetstream-cloud.org/files/Jetstream-UTSA-Oct2017-Presentation.pptx https://jetstream-cloud.org/files/Jetstream-UTSA-171012.pptx
Oct 11	LBNL Labtech 2017	Berkeley, CA	Brian Beck	https://iu.box.com/s/t2tjtg1z9v6z4xk8udewcl4539t4n62
Oct 17	LANL Cloud and HPC day	Los Alamos, NM	Mike Lowe	
Oct 20	Statewide IT	Bloomington, IN	Steve Bird	https://iu.box.com/s/ftgjb12rltkj3c53odf7gjx8mscf8fl

Oct 23-25	Science Gateways Conference 2017	Ann Arbor, MI	Jeremy Fischer, Sanjana Sudarshan, Lee Liming	https://sciencegateways.org/web/gateways2017 http://bit.ly/auth-tutorial https://github.com/jlf599/GatewaysTutorial/blob/master/README.md NOTE: Enis Afgan (JHU) also gave a presentation that prominently mentioned Jetstream.
Oct 23	Science Gateway Architectures	Bloomington, IN	Mike Lowe, George Turner	https://iu.box.com/s/ruula9t9heqdro6ebnz5uecco3n5x3tq
Nov-8	GIS day	Bloomington, IN	George Turner, Sanjana Sudarshan	
Nov 10-11	CIS-IEEE EnCON Conference	Bloomington, IN	Steve Bird	Chaired by Katy Börner, http://www.cis-ieee.org/EnCON/
Nov 12-17	SC17	Denver, CO	IU Booth 601	http://sc17.supercomputing.org
Nov 11-15	Society for Neuroscience	Washington, DC	Brian Beck, Sanjana Sudarshan, Booth 106	https://www.sfn.org/annual-meeting/neuroscience-2017

Meetings with research groups interested in being future users of Jetstream:

- Afgan, E. - Began collaborative effort to establish a cross-national cloud infrastructure for use by the Galaxy project with Dr. Andrew Lonie and the Genomics Virtual Laboratory project from Melbourne, Australia. The intent is to standardize and integrate usage of the Jetstream and NeCTAR research clouds.
- Beck, BW “Expanding the use of TACC Computing Resources Available to the University of Texas at Dallas via the UTRC” with Associate Vice President for Research Rafael Martin at University of Texas at Dallas. Feb 27, 2017.

- Beck, BW; Turner, G; Fischer, J; Lowe, M; Packard, M; Lenards, A; Hancock, D “Jetstream UTSA/UTHSCSA Outreach Meeting” at University of Texas at San Antonio with University of Texas at San Antonio Open Cloud Institute and University of Texas Health Sciences Center San Antonio. Jan 16-19, 2017. Presented Jetstream and assisted the students, staff, and faculty of UTSA Faculty Rad, P and UTHSCSA Faculty Brookes, E. adapting their cloud offerings for use on Jetstream
- Beck, BW - Discussion with University of Texas at Dallas Faculty Brandt, P. about moving his journalism/political science gateway to a full research project on Jetstream. Feb 27, 2017.
- Beck, BW presented a tutorial, including Jetstream use. “TACC Computing Resources Available to UT System members via the UTRC” at University of Texas at Arlington Office of Information Technology Research Advisory Committee Meeting. Dec 1, 2016
- Beck, BW. Invited talk. "XSEDE - Jetstream: Accessible cloud computing for the national science and engineering communities” at Genome Biology Group, Department of Biology, University of Texas at Arlington. Feb 3, 2017
- Beck, BW Discussion with Vice President for Information Technology and Chief Information Officer Frank Feagans at University of Texas at Dallas about migrating/starting projects on Jetstream, particular focused on *Internet-of-Things* projects that UTD is developing as a center for excellence. May 3, 2017
- Beck, BW Discussion with University of Texas Health Science Center Houston Associate Dean for Academic Affairs Susan Fenton, Assistant Professor of the School of Biomedical Informatics Luca Giancardo, and Sr. Administrator of the School of Biomedical Informatics John Rizkallah about developing two courses using Jetstream: One focused on Biomedical Computing in the cloud and the second focused on teaching biomedical informatics methods based on Jetstream instances. May 8, 2017
- Beck, BW Discussion with Associate Vice President Dr. Victor Fishman and Principal Investigator Michael Karich of the Applied Research Center (ARC) at the University of Texas at Dallas. The ARC has a small cloud they use to develop industry-faculty partnerships and were interested in leveraging Jetstream to expand the reach and scope of these projects. May 9, 2017
- Beck, BW Discussion with Dr. Mina Ameri of the Jindahl School of Management at the University of Texas at Dallas about moving her marketing and economics model simulator to Jetstream. May 22, 2017
- Presentation to Dr. Thomas Spencer, Director of Operations: Academic and Administrative Information Resources, University of Texas Southwestern Medical Center. August 31, 2017. (Brian Beck)
- Presentation to Terry Montgomery, Senior Project Manager, Health Systems Information Resources, University of Texas Southwestern Medical Center. July 25, 2017. (Brian Beck)
- Presentation to R. Gabe Cavazos, Director, Research Information Systems, University of Texas at Dallas, Aug 18, 2017. (Brian Beck)
- Presentation to Dr. Andrew Wiedlea, Tammy Campbell, Kuldeep Chawla, Ian Vaino, Tareq Saif. Information Technology. Lawrence Berkeley National Laboratory, Aug 16

2017. Introduced Jetstream and discussed research opportunities for LBL staff. Invited to present at October 11, 2017 LBL meeting. (Brian Beck)

- Jeremy Fischer presented a webinar for the Science Gateways Community Group on June 16, 2017. Available at <https://jetstream-cloud.org/files/Jetstream-Gateways-Call-June2017-Presentation.pptx>
- Swetnam, T. - Presentation to the Botanical Information and Ecological Network (BIEN) on JS for RStudio and CyVerse DataStore.
- Swetnam, T. - Began collaborative effort with Dr. Jonathan Coop (Western Colorado University) and MS student Ryan Walker to use JS processing of 3D point clouds derived from lidar for post-fire ecosystem analysis. Jan 21, 2017 . url: <http://www.western.edu/people/jonathancoop>
- Swetnam, T. - Began collaborative effort with Dr. Greg Barron-Gafford (University of Arizona) with MS student Sean Hendryx to use JS processing of 3D point clouds derived from structure from motion for ecosystem characterization. Feb 6, 2017 urls:
- Swetnam, T. Meeting with USDA Agricultural Research Service, Southwest Watershed Research Center. Short presentation on using JS for RStudio and GIS modelling. Jan 4, 2017. url: <https://www.ars.usda.gov/pacific-west-area/tucson-az/southwest-watershed-research/>
- Swetnam, T. - CyVerse / TERRA-REF / Phenome Hackathon. Used JS for collaborating and developing code w/ Jupyter Notebook. Feb 14-16, 2017. url: https://github.com/bjoyce3/Image_Data_rapid_QAQC
- Discovery Saturday talk titled “Cyber-Cowboys on the Range” by Tyson Swetnam at the Santa Rita Experimental Range in Green Valley, AZ on March 18, 2017. http://prezi.com/ba-ij00savre/?utm_campaign=share&utm_medium=copy&rc=ex0share

Courses and workshops utilizing Jetstream resources:

- Jetstream partners at UTSA has taught the following courses during Spring 2017 by leveraging Jetstream Cloud resources for student projects and labs
 - Dr. Mo Jamshidi, Control Engineering , Cloud-Based Intelligent robotics
 - Dr. Paul Rad, Computer Engineering , Applied machine learning with Big Data
- Jetstream partners at JHU facilitated use of Jetstream via Galaxy for a course taught there that focuses on bioinformatics data analysis.
 - *Bioinformatics: Tools for Genome Analysis* (course 410.635, <http://advanced.jhu.edu/academics/graduate-degree-programs/bioinformatics/course-descriptions/>). Among nine students and the instructor they consumed more than 280,000 SUs from March through May 2017. They each launched and managed their own instances.
- Jetstream was utilized to teach the 2017 Computational Science in the Cloud Institute at TACC. July 24-28, 2017. <https://tacc.github.io/CSC2017Institute/>

- Matt Vaughn, Nirav Merchant, and PoreCamp USA Team assembled the PoreCamp USA workshop at Texas A&M University in College Station, TX June 5-9, 2017. <https://porecamp.github.io/>
- C. Titus Brown and Upendra Kumar Devisetty taught the Data Intensive Biology Summer Institute for Sequence Analysis at UC Davis in Davis, CA, June 18 - July 21, 2017. <http://ivory.idyll.org/blog/2017-dibsi-xsede-request.html>
- Rob Edwards gave the 2017 Metagenomics Workshop in San Diego, CA on June 26-30. <https://edwards.sdsu.edu/research/2017-metagenomics-workshop/>
- Jawon Song. Jetstream workshop presented to Rural Development Administration (RDA)-South Korea in June 26-28, 2017

1.5. What do you plan to do during the next reporting period to accomplish the goals?

Particular technical highlights of this work are as follows:

- Now that S3-style object storage services are beginning to be used on Jetstream, the University of Arizona and University of Chicago teams are working toward providing web-browser user interfaces that will enable easier, broader access for novice users. We hope that this will provide the scaffolding necessary to introduce non-expert researchers and educators to this new style of data storage.
- The University of Arizona will be extending Atmosphere to support the following functionality:
 - Persistent instances: The current focus on Atmosphere is the management of on-demand cloud resources. Atmosphere will be extended to support persistent instances by allowing users to manually assign static ip addresses and manage their OpenStack security groups. Allowing users to manage persistent instance functionality will especially benefit science gateway users who want create or scale their Jetstream instances using Atmosphere.
 - On-Demand virtual clusters: Atmosphere will allow users to launch multiple instances simultaneously and provide the tooling to organize instances into virtual clusters. Users will also be able to logically organize their instances into groups when the virtual clusters are initially defined as well as add instances to existing logical groups as their scale their virtual clusters.
 - Atmosphere's web services will be engineered to support "high availability" and to reduce the time needed for maintenances.
 - Atmosphere users will be able to share cloud resources within a project
 - Atmosphere users will be able to associate DOIs to images they create and search for other images that have DOIs.

- Johns Hopkins University will focus on enhancing utilization of Jetstream resources from Galaxy Main by incorporating dynamic scaling of the number of resources based on the current system workload.

Administrative aspects of this work include:

- Organizing the next Stakeholder Advisory Board Meeting
- Administering the next Jetstream user survey
- Continuing reporting and metric collection
- Considering additional supplements for undergraduate and graduate students

Upcoming training and outreach events include:

- Develop an online training module, “Using the Jetstream API (GUI and CLI) for advanced cloud functionality”. Cornell University
- Develop an online training module, “Scientific reproducibility of computations and analyses: publishing, archiving, and curating your VM images for posterity”. Cornell University
- A tutorial and presentation is planned at the Plant and Animal Genome Conference
- PEARC18 – Multiple submissions for papers (both staff and student-led)
- Workshops at UNC, UF, NC State are in plan.
- Additional outreach to the Neuroscience community is planned via multiple events.

2. Products (resulting from this project during the specified reporting period)

Journal Articles or Juried Conference Papers:

1. Chambers, M. C., Jagtap, P. D., Johnson, J. E., McGowan, T., Kumar, P., Onsongo, G., ... & Grüning, B. (2017). An Accessible Proteogenomics Informatics Resource for Cancer Researchers. *Cancer Research*, 77(21), e43-e46.
<http://cancerres.aacrjournals.org/content/77/21/e43.short>
2. Fischer, J., Hancock, D. Y., Lowe, J. M., Turner, G., Snapp-Childs, W., & Stewart, C. A. (2017, October). Jetstream: A Cloud System Enabling Learning in Higher Education Communities. In *Proceedings of the 2017 ACM Annual Conference on SIGUCCS* (pp. 67-72). ACM. DOI: <https://doi.org/10.1145/3123458.3123466>
3. Swetnam, T. L., P. D. Brooks, H. R. Barnard, A. A. Harpold, and E. L. Gallo. 2017. Topographically driven differences in energy and water constrain climatic control on forest carbon sequestration. *Ecosphere* 8(4):e01797.10.1002/ecs2.1797
<http://onlinelibrary.wiley.com/doi/10.1002/ecs2.1797/full>
4. Natasha E. Weiser, Danny X. Yang, Suhua Feng, Natallia Kalinava, Kristen C. Brown, Jayshree Khanikar, Mallory A. Freeberg, Martha J. Snyder, Györgyi Csankovszki,

- Raymond C. Chan, Sam G. Gu, Taiowa A. Montgomery, Steven E. Jacobsen, John K. Kim (2017). MORC-1 Integrates Nuclear RNAi and Transgenerational Chromatin Architecture to Promote Germline Immortality. *Developmental Cell*, Volume 41, Issue 4, 408 - 423.e7 DOI: <http://dx.doi.org/10.1016/j.devcel.2017.04.023>
5. Zhou, C., & Paffenroth, R. C. (2017, August). Anomaly Detection with Robust Deep Autoencoders. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 665-674). ACM. <https://doi.org/10.1145/3097983.3098052>
 6. **Submitted:** Gajamannage, K., Paffenroth, R., & Bollt, E. M. (2017). A Nonlinear Dimensionality Reduction Framework Using Smooth Geodesics. *arXiv preprint arXiv:1707.06757*. <https://arxiv.org/abs/1707.06757>
 7. Holdgraf, C., Culich, A., Rokem, A., Deniz, F., Alegro, M., & Ushizima, D. (2017). Portable Learning Environments for Hands-On Computational Instruction: Using Container- and Cloud-Based Technology to Teach Data Science. In *Proceedings of the Practice and Experience in Advanced Research Computing 2017 on Sustainability, Success and Impact* (PEARC17). ACM, New York, NY, USA, Article 32, 9 pages. DOI: <https://doi.org/10.1145/3093338.3093370>
 8. Emre Brookes and Alexey Savelyev. (2017). GenApp Integrated with OpenStack Supports Elastic Computing on Jetstream. In *Proceedings of the Practice and Experience in Advanced Research Computing 2017 on Sustainability, Success and Impact* (PEARC17). ACM, New York, NY, USA, Article 11, 8 pages. DOI: <https://doi.org/10.1145/3093338.3093356>
 9. Knepper, R., Coulter, E., Pierce, M., Marru, S., & Pamidighantam, S. (2017, May). Using the Jetstream Research Cloud to provide Science Gateway resources. In *Proceedings of the 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing* (pp. 753-757). IEEE Press. <https://doi.org/10.1109/CCGRID.2017.121>
 10. Afgan, E., Christie, M., Goonasekera, N., (2017 Oct). CloudLaunch as a Gateway for Discovering and Launching Cloud Applications. In *Proceedings of Gateways 2017*. Ann Arbor, MI, USA. <https://doi.org/10.6084/m9.figshare.5471800.v1>
 11. **Submitted:** Afgan, E., Lonie, A., Taylor, J., Goonasekera, N. CloudLaunch: Discover and Deploy Cloud Applications. *Future Generation Computer Systems* (FGCS).

Other Conference Presentations/Papers:

1. Afgan, E. (Feb 5, 2017). Talk titled "Resource planning on the Cloud: exploring the scalability spectrum" at the Galaxy Australasia Meeting (GAME) 2017 conference. <https://www.slideshare.net/afgane/resource-planning-on-the-amazon-cloud>
2. Afgan, E. (June 29, 2017). Training workshop at the Galaxy Community Conference (GCC) titled "Advanced accelerated Galaxy admin" by Enis Afgan, et al. in Montpellier, France. <https://galaxyproject.github.io/dagobah-training/2017-montpellier/00-intro/intro.html#1>
3. Fischer, J. (2017). "Jetstream overview - A national research and education cloud". Jeremy Fischer presented at the Northern Lights Network Annual Conference to a group

- of technical staff. Minneapolis, MN.
<http://www.jetstream-cloud.org/files/Jetstream-NorthernLights-Presentation.pptx>
4. Fischer, J. (2017). "Jetstream overview and tutorial" Jeremy Fischer made a short presentation and then conducted a hands-on tutorial at the Society of Industrial and Applied Mathematics (SIAM) 2017 Computational Science and Engineering (CSE17) conference. Atlanta, GA Available at
<http://www.jetstream-cloud.org/files/Jetstream-SIAM-CSE-Presentation.pptx>
 5. Fischer, J. (2017, June 1). Introduction to Jetstream. Presentation given at the Great Plains Network Annual Meeting, Kansas City, MO.
<https://jetstream-cloud.org/files/Jetstream-GPN-Annual-Conference-Presentation.pptx>
 6. Fischer, J. (2017, June 7). A brief introduction to Jetstream. Presentation given at the American Society for Mass Spectrometry (ASMS) annual conference, Indianapolis, IN.
 7. Fischer, J. & Coulter, E. (2017, July 9). Programmable Cyberinfrastructure: Introduction to Building Clusters in the Cloud. Presentation given at the annual meeting of the Practice & Experience in Advanced Research Computing Conference, New Orleans, LA. Retrieved from https://github.com/ECoulter/Tutorial_Practice
 8. Fischer, J. (2017, July 26). Introduction to Jetstream. Presentation given at the annual meeting of the Earth Science information Partners (ESIP), Bloomington, IN. Retrieved from <http://www.esipfed.org/meetings/upcoming-meetings/esip-summer-meeting-2017>
 9. Gniady, T. & Fischer, J. (2017, July 9). Using R with Jetstream. Presentation given at the annual meeting of the Practice & Experience in Advanced Research Computing Conference, New Orleans, LA. Retrieved from
<https://jetstream-cloud.org/files/Jetstream-PEARC17-Presentation.pptx>
 10. Hancock, D.Y., Lowe, J.M. (April 12, 2017). "Industry Insights: OpenStack from the Basics to Production Experiences on Jetstream". Talk given at the annual meeting of the HPC Advisory Council. Lugano, Switzerland.
<http://www.hpcadvisorycouncil.com/events/2017/swiss-workshop/pdf/Weds12April/DHancock.JMLowe.OpenStackBasicstoProductionExprncsonJetstream.Wed041217.pdf>
 11. Jones, C.J., Sudarshan, S., Mehta, I.D., Beck, B.W. (Feb 11-16 2017). "Different Applications for CAS in Functional Classification of Protein Interfaces". 61st Annual Biophysical Society Meeting, New Orleans, LA DOI:
<http://dx.doi.org/10.1016/j.bpj.2016.11.2645>
[http://www.cell.com/biophysj/fulltext/S0006-3495\(16\)33675-X](http://www.cell.com/biophysj/fulltext/S0006-3495(16)33675-X)
 12. Liming, R.L. (Apr 11, 2017). "Login with XSEDE and Jetstream: Our Experiences with Globus Auth." Presentation given at GlobusWorld 2017, Chicago, IL. Retrieved from
https://www.globusworld.org/files/2017/06_XSEDE_Jetstream_Liming.pdf
 13. Mehta, I.D., Beck, B.W. (Feb 11-16 2017) "Protein Energy Network Models to Classify and Predict Functionally Linked Interfaces of Proteins from Functionally Uncorrelated Interfaces". 61st Annual Biophysical Society Meeting, New Orleans, LA. DOI:
<http://dx.doi.org/10.1016/j.bpj.2016.11.1871>
[http://www.cell.com/biophysj/fulltext/S0006-3495\(16\)32901-0](http://www.cell.com/biophysj/fulltext/S0006-3495(16)32901-0)
 14. Nekrutenko, A. & Taylor, J. (June 29, 2017). "The Evolution of Galaxy: A Rough Timeline". Presentation given at the Galaxy Community Conference (GCC), Montpellier,

- France. Retrieved from
http://schd.ws/hosted_files/gcc2017/b3/S1_T1_Evolution_of_Galaxy.pdf
15. Pope, K. (2017, July 31). Green Infrastructure Increases Vegetation Growth. Presentation given at the Urban Water Innovation Network Annual Meeting, Fort Collins, CO. Retrieved from
https://erams.com/UWIN/wp-content/uploads/2017/08/Pope-Green-Infrastructure-Increases-Vegetation-Growth_compressed.pdf
 16. Swetnam, T. (April 28, 2017). "A Gentle Introduction to Forestry Science Workflows in the Era of Cloud Computing". Presentation to the Society of American Foresters Sectional Meeting, Albuquerque, NM. Retrieved from
<https://prezi.com/view/DKmuPtvGJHfTolrerRuW/>
 17. Taylor, J. (Feb 5, 2017). "A performance evaluation for Galaxy". Presentation given at the Galaxy Australasia Meeting (GAME) 2017 conference. Melbourne, Australia. Retrieved from <https://speakerdeck.com/jtx/game-2017-a-performance-evaluation-for-galaxy>
 18. Taylor J. (Feb 22, 2017). "Accessible, transparent, reproducible genomic analysis with Galaxy". Presentation given at the Bio Genomics 2017 conference. Washington D.C. Retrieved from <https://speakerdeck.com/jtx/bio-genomics-2017>
 19. Taylor, J. (July 22, 2017). "Supporting highly scalable scientific data analysis with Galaxy". Presentation given at the International Society for Computational Biology (ISMB) conference, Prague, Czech Republic. Retrieved from
<https://speakerdeck.com/jtx/ismb-2017-supporting-highly-scalable-scientific-data-analysis-with-galaxy>
 20. Vaughn, M. (Jan 13, 2017). "Jetstream: On-Demand Cloud Computing for Life Sciences Research and Education". Presentation given at the XXV Plant and Animal Genome Conference. San Diego, CA.

Other Publications:

1. Afgan, E. (Feb 17, 2017). Talk titled "Cloud Computing and Bioinformatics" at the University of Colombo, in Colombo, Sri Lanka.
<https://www.slideshare.net/afgane/cloud-computing-and-bioinformatics>
2. Afgan, E. (Feb 27, 2017). Talk titled "The Pulse of Cloud Computing" at the University of Colombo, in Colombo, Sri Lanka.
<https://www.slideshare.net/afgane/the-pulse-of-cloud-computing-with-bioinformatics-as-a-n-example>
3. Fischer, J. "Jetstream overview - A national research and education cloud" Jeremy Fischer presented to Texas A&M Research staff, IT staff, and researchers with a talk and tutorial in College Station, TX. Available at:
<http://www.jetstream-cloud.org/files/Jetstream-TexasAM-Presentation.pptx>
4. George Turner. "New Ventures in Research, Engineering, and Educational Computing" George Turner presented on the collaboration between the Jetstream and the Indiana University Network Science Institute at the [Open Science Forum](#) at Indiana University, Bloomington, IN. Apr 26, 2017. Retrieved from
<http://hdl.handle.net/2022/21537>

5. George Turner. "New Ventures in Research, Engineering, and Educational Computing" George Turner participated in a panel discussion at the [Boston Open Research Congress](#) an effort by organizations with interest in cloud computing to establish open standards for clouds to encourage interoperability. May 11-12, 2017. Retrieved from <http://hdl.handle.net/2022/21536>
6. George Turner. "New Ventures in Research, Engineering, and Educational Computing" George Turner presented at the [Southern Partnership in Advanced Networking \(SPAN\) Workshop](#) on Jetstream as an advanced cyberinfrastructure facility that could benefit SPAN's membership which is made up predominantly historically black colleges and universities. South Carolina State University, Orangeburg, SC. May 18-19, 2017. Retrieved from <http://hdl.handle.net/2022/21535>
7. George Turner. "Jetstream Cloud Facility for Advancing Scientific Communities" George Turner presented at the [UCAR Cloud Modeling Workshop](#) about how Jetstream could be utilized by the atmospheric science community for modeling, analysis, and data distribution. May 31 - June 2, 2017. Retrieved from <http://hdl.handle.net/2022/21534>
8. Merchant, N. (2017). Computing at scale: From laptop to cloud and HPC. Presentation given at the annual meeting of the Large Synoptic Survey Telescope Corporation Data Science Fellowship Program. Tucson, AZ. Retrieved from <http://hdl.handle.net/2022/21545>
9. Craig A. Stewart, Brian Beck, (Feb 15, 2017). "Jetstream - A Science Cloud Supporting US Research in All Science Disciplines". Poster given at the annual meeting of American Association for the Advancement of Science. <http://hdl.handle.net/2022/21248>
10. Craig A. Stewart et al. (Feb 28, 2017). "Jetstream Stakeholder Advisory Board Meeting February 2017: Presenters' Report". <http://hdl.handle.net/2022/21247>
11. Stewart, C.A. (2017). New trends in cyberinfrastructure for ecology and evolutionary biology. Presented at Indiana University, Bloomington, IN. Retrieved from <http://hdl.handle.net/2022/21264>
12. Stewart, C. A., Knepper, R., Link, M. R., Pierce, M., Wernert, E., & Wilkins-Diehr, N. (2017). Cyberinfrastructure, Cloud Computing, Science Gateways, Visualization, and Cyberinfrastructure Ease of Use. IGI Global. <http://hdl.handle.net/2022/21589>
13. Discovery Saturday talk titled "Cyber-Cowboys on the Range" by Tyson Swetnam at the Santa Rita Experimental Range in Green Valley, AZ on March 18, 2017. http://prezi.com/ba-ij00savre/?utm_campaign=share&utm_medium=copy&rc=ex0share
14. O'Riley, S. & Arya, A. (2017). Jetstream and Brain-Life: Creating a reproducible research platform on the Jetstream cloud. <https://doi.org/10.5967/K8MW2F9P>
15. Lwowski, B. (2017). Inter-operable, reproducible and efficient deployment of software and software changes in the Jetstream OpenStack environment. <https://doi.org/10.5967/K8RN361D>
16. Blake, S. (2017). Simulations to Test Data Assimilation through Feedback Nudging. <https://doi.org/10.5967/K8H41PK8>
17. Jones, C. (2017, August). Jetstream joyride. *IT Connections*. Retrieved from <https://itconnections.iu.edu/2017-august/jetstream.php>

Other Products:

Survey Instruments

1. Jetstream 2017 User Survey available at <http://hdl.handle.net/2022/21840>
2. Jetstream 2016 User Survey available at <http://hdl.handle.net/2022/21839>
3. Jetstream Field and Marine Station Computational Needs Survey available at <http://hdl.handle.net/2022/21838>
4. Jetstream Engineering Computational Needs Assessment Survey available at <http://hdl.handle.net/2022/21836>

Audio or video products:

1. Fischer, J. (2017, June 16). Jetstream for Science Gateways Community Group [Webinar]. Retrieved from <https://jetstream-cloud.org/files/Jetstream-Gateways-Call-June2017-Presentation.pptx>
2. Stewart, C.A. & Hancock, D.Y. (2017). Jetstream Overview - what it is, how to apply for use [Webinar]. Retrieved from <https://youtu.be/q-o8paUT4KY>.
3. Virtual workshop on allocations available at <https://cvw.cac.cornell.edu/JetstreamReq/>

Software or netware

1. Application: <https://github.com/brain-life/app-dipy-afq>
2. Application: <https://github.com/brain-life/app-dipy-tracking>
3. Application: <https://github.com/brain-life/app-dipy-csamodel>
4. Application: <https://github.com/brain-life/app-dipy-shore>
5. Application: <https://github.com/aarya22/app-recobundles>
6. Application: <https://github.com/aarya22/app-tract-profile>
7. Application: <https://github.com/brain-life/ui-surfaceview>
8. Application: <https://github.com/brain-life/ui-tractview>
9. Application: <https://github.com/brain-life/core-warehouse>
10. Application: <https://brain-life.org/warehouse>
11. Application: <https://github.com/brain-life/ui-lifeview>
12. Application: <https://github.com/brain-life/amaretti>
13. Application: <https://github.com/galaxyproject/cloudlaunch/>
14. Application: <https://github.com/galaxyproject/cloudlaunch-ui/>
15. Application: <https://github.com/gvlproject/cloudbridge/>

Other

1. Freeberg, Mallory, & Heydarian, Mohammad. (2017). Training material for de novo transcriptome reconstruction from RNA-seq data [Data set]. Zenodo. <http://doi.org/10.5281/zenodo.254485>
2. Freeberg, Mallory, & Heydarian, Mohammad. (2016). Training material for ChIP-seq analysis [Data set]. Zenodo. <http://doi.org/10.5281/zenodo.197100>

3. Participants

3.1. Individuals

First Name	Last Name	Most Senior Project Role	Nearest Person Month Worked	Email (if new to project)	Affiliation	Contribution
Craig	Stewart	PD/PI	3		IU	PI
Ian	Foster	Co-PD/PI	1		UC	Co-PI
Matthew	Vaughn	Co-PD/PI	1		TACC	Co-PI
Nirav	Merchant	Co-PD/PI	4		UA	Co-PI
James	Taylor	Co-PD/PI	1		JHU	Co-PI
Bret	Hammond	Other Professional	3		IU	System Administration
John	Lowe	Other Professional	8		IU	System Administration
Therese	Miller	Other Professional	4		IU	Project management
Enis	Afgan	Other	6		JHU	Developer
William (Joe)	Allen	Other	0		TACC	User Services
Brian	Beck	Other	8		TACC	User Services
Stephan	Bird	Other	4		IU	System administration
Margues	Bland	Other	1	mbland@tacc.utexas.edu	TACC	Project management
Adam	Brazier	Other	1	ab447@cornell.edu	Cornell University	Outreach
Ashley	Bucholz	Other	3		TACC	Project management
Tim	Cockerill	Other	4		TACC	Project management
Jeremy	Fischer	Other	12		IU	Outreach
Mallory	Freeberg	Other	1		JHU	Outreach
Steve	Gregory	Other	8		UA	Software engineer
David	Hancock	Other	4		IU	Acting PI / Senior investigator
Manu (Mary)	John	Other	2	mjohn@tacc.utexas.edu	TACC	Project management
Amit	Juneja	Other	5		UA	Developer
Andy	Lenards	Other	6		UA	Developer
Lee	Liming	Other	1		UC	Developer
Chris	Martin	Other	5		UA	Developer
Susan	Mehring	Other	1	shm7@cornell.edu	Cornell University	Outreach

					ty	
Nathaniel	Mendoza	Other	3		TACC	Information Security Officer
Mike	Packard	Other	6		TACC	System administration
Julian	Pistorius	Other	9		UA	Developer
Akhil	Seth	Other	4		TACC	Developer
Edwin	Skidmore	Other	4		UA	Project lead for iPlant/Atmo coordination
Sanjana	Sudarshan	Other	4		IU	Outreach
Patrick	Storm	Other	3		TACC	Security Engineer
Ben	Trumbore	Other	1	wbt3@cornell.edu	Cornell University	Outreach
George	Turner	Other	8		IU	Chief Architect
Aman	Arya	REU participant	2		UW	
Steven	O'Riley	REU participant	2		IU	
Sanchez	Blake	REU participant	2		Bethune-Cookman	
Brandon	Lwowski	REU participant	2		UTSA	

3.2. Partner organizations

Type of Partner organization: Academic Institution

Name: University of Colorado/CIRES

Location: Boulder, CO

Partner's contribution to the project: Collaborative research

Brian Johnson (University of Colorado-National Snow and Ice Data Center) has agreed to participate as a research collaborator using the Jetstream system for research in the polar science area.

Type of Partner organization: Academic Institution

Name: University of Texas at San Antonio

Location: San Antonio, TX

Partner's contribution to the project: Collaborative research

More detail on partner and contribution: Paul Rad, UTSA and Rackspace, will contribute software expertise for system integration.

Type of partner organization: Academic Institution

Name: University of Arizona

Location: Tucson, AZ

Partner's contribution to the project: Facilities; In-kind support

The University of Arizona develops the software product called Atmosphere, which is the primary, initial user interface to Jetstream and seamlessly orchestrates cloud resources across the two separate OpenStack sites, located at the Indiana University and TACC. Atmosphere was initially conceived and successfully used by CyVerse (formerly, named iPlant Collaborative) for their users, and has been adapted to Jetstream. The Atmosphere development team primarily focuses on integrating with the TACC Accounting System (TAS), modifying Atmosphere to the XSEDE allocations model, improving the user experience of the Atmosphere user interface, and numerous enhancements and bug fixes related to Jetstream's specific use cases and requirements. Another significant activity related to end-user and operational support tasks, when issues arise within the Jetstream system.

Type of Partner organization: Academic Institution

Name: Johns Hopkins University

Location: Baltimore, MD

Partner's contribution to the project: Collaborative research; In-kind support

More detail on partner and contribution: JHU has developed a system for setting up and integrating virtual machines running on Jetstream into a Docker Swarm cluster as a way to facilitate robust scheduling of containerized applications run on usegalaxy.org server. The developed system automatically provisions and configures the machines based on current workload. We have also provided an updated version of the Galaxy Standalone virtual machine with the latest version of Galaxy and a toolset to reflect recent advances in genomic tool development.

Type of Partner organization: Academic Institution

Name: University of Chicago

Location: Chicago, IL

Partner's contribution to the project: In-kind support; Collaborative research

More detail on partner and contribution: UChicago made certain that Jetstream's authentication/login function--which relies on Globus-based XSEDE user authentication--worked continuously and satisfied user needs. The UChicago-based Globus user support team actively monitored XSEDE tickets related to Jetstream logins and was involved in debugging and resolving a handful of specific user requests related to this function.

Type of Partner organization: Academic Institution

Name: University of Texas at Austin, Texas Advanced Computing Center (TACC)

Location: Austin, TX

Partner's contribution to the project: In kind support; Facilities; Collaborative research

More detail on partner and contribution: TACC supports the project by hosting a portion of the production hardware as well as offering personnel to operate and function as a tier-2 user support.

Type of Partner organization: Academic Institution

Name: Cornell University

Location: Ithaca, New York

Partner's contribution to the project: In-kind support

More detail on partner and contribution: Cornell has delivered a web-based training module (virtual workshop) with additional modules planned. These modules are comprised of short video clips, examples, exercises, and quizzes, with full text discussion incorporating an HPC glossary. The first module focused on the allocations process. The second module will be an introduction to the Jetstream system and the Atmosphere cloud computing environment, focused on using, creating, and archiving services and VMs.

Type of Partner organization: Academic Institution

Name: Penn State University

Location: University Station, Pennsylvania

Partner's contribution to the project: In kind support

More detail on partner and contribution: Anton Nekrutenko, Penn State University, development of the Galaxy software suite.

Type of Partner organization: Academic Institution

Name: University of Hawaii

Location: Honolulu, Hawaii

Partner's contribution to the project: Collaborative research

More detail on partner and contribution: Gwen Jacobs, University of Hawaii, serves as an exemplar for colleges and universities across the United States in utilizing Jetstream's unique features for faculty and students as well as developing VM images supporting the research needs in ocean science.

Type of Partner organization: Academic Institution

Name: University of North Carolina at Chapel Hill

Location: Chapel Hill, North Carolina

Partner's contribution to the project: Collaborative research

More detail on partner and contribution: Thomas M. Carsey, University of North Carolina at Chapel Hill, The Odum Institute, advises on software tools of interest to the social science research community; work to increase awareness of Jetstream to our user community, and serve as a pilot site for the implementation of Jetstream's distributed cloud services.

Type of Partner organization: Academic Institution

Name: Jackson State University

Location: Jackson, Mississippi

Partner's contribution to the project: Collaborative research

More detail on partner and contribution: Jessie Walker, Jackson State University, is an unfunded partner, leveraging the Jetstream project in support of academic and research endeavors.

Type of Partner organization: Commercial vendor

Name: Dell, Inc.

Location: Round Rock, Texas

Partner's contribution to the project: In-kind support

More detail on partner and contribution: Configure, manufacture, and support the Jetstream hardware utilized in the test and development as well as the production systems.

Type of Partner organization: Commercial vendor

Name: The MathWorks, Inc

Location: Natick, Massachusetts

Partner's contribution to the project: In-kind support

More detail on partner and contribution: Enable end users of the Jetstream system to run MATLAB™ (and any other MathWorks products the end user is licensed to use). This is accomplished through the “bring your own license” arrangement which MathWorks has piloted with a small number of facilities worldwide. In addition, MathWorks, Inc. is allowing Indiana University to purchase a set of academic licenses of MathWork products which can be made available to any academic user of the jetstream system.

3.3. Have other collaborators or contacts been involved?

No.

4. Impact

4.1. What is the impact on the development of the principal discipline(s) of the project?

Jetstream has had significant impact on the areas of computational science. In the area of computational science, as a first-of-a-kind production cloud funded by the NSF, Jetstream has been a pilot, a pathbreaker, and a tremendous learning experience. The Jetstream project has published lessons learned in peer-reviewed conference proceedings, which are important ways to have impact on the computational science community. In the first operational year of this project Jetstream may have had more impact on the community through extensive participation in many working groups, workshops, BOFs, and small meetings and discussions that are shaping the community approach to implementation, provision, and support of publicly-funded cloud services for open science.

A few particular examples of significant impact on the field of computational science are as follows:

- The Jetstream team has initiated and is leading a multi-organization discussion about mechanisms for management of VM libraries and repositories. The one thing we have learned so far is that at least two mechanisms are likely needed in the future: one is the repository of static, unchanging VM images that we have already established through use of IU's persistent digital repository (Scholarworks.iu.edu). However, for operational work, a scripted, “create on the fly” mechanism is essential so that you can create a

current version of a particular VM functionality but incorporate whatever is the current version of relevant components at any given time. The team is looking to leverage GitHub, Cloud-init scripts, and Heat templates to demonstrate this functionality.

- During PY2 hundreds of students took classes in computational sciences and did independent research in computational sciences based on use of Jetstream. The Jetstream team impacted over 800 students through classes and research uses of the system.
- Stewart et. al [2015] – the main paper describing plans for the Jetstream system - has already been cited by 20 other published technical papers, and read statistics from Researchgate.net indicate that it has been read more than 200 times just through that venue.
- Stewart et. al [2016] – the paper describing the implementation of, acceptance of, and early experiences with Jetstream – was given the best technical paper award at the XSEDE'16 conference, this indicating both its quality and impact as a contribution to the field of computational science. A follow on to this paper was invited to submit to a special issue journal that includes the notable papers from the PEARC '17 conference.
- The significant number of products from the Jetstream team and highlights from researchers that have used Jetstream during its first year of operations are included which both demonstrate the breadth of effort the team has put into outreach and the impact already being made within specific disciplines.

4.2. What is the impact on other disciplines?

The impact on other disciplines is as follows:

- Jetstream has been used by over 800 students through educational or startup allocations.
- Several papers have been published in peer-reviewed journals.
- Several students have completed/made considerable progress on their dissertations.
- Sanchez Blake (2017) REU participant presented some of his work (“Simulations to Test Data Assimilation through Feedback Nudging”) at the annual Indiana Mathematics REU program (West Lafayette, IN). The presentation can be found here: <https://doi.org/10.5967/K8H41PK8>
- At 2017 Urban Water Innovation Network Annual Meeting in Fort Collins, CO on July 31-August 1, 2017, undergraduate student Kayla Pope presented on her research, “Green Infrastructure Increases Vegetation Growth”, which resulted from using Jetstream. Kayla Pope received third place in a poster competition. Her mentor is Tyson Swetnam from the University of Arizona.
- 12 neuroscience applications were developed on Jetstream and are available on Brain-Life.org. Jetstream has also been selected as the cloud provider for NSF DBI#1707356 "NeuroNex Technology Hub: Enhanced resolution for 3DEM analysis of synapses across brain regions and taxa"

4.3. What is the impact on the development of human resources?

Jetstream is having significant impact as a resource in the development of human resources, and Jetstream dissemination efforts are playing an important role in providing opportunities for self-directed learning. Key points include:

- Jetstream has been used directly and indirectly by more than 10,000 people to date through specific allocations and jobs launched via science gateways.
- More than 2,000 people have attended a seminar, tutorial, or talk about Jetstream.
- Hundreds of people have now read peer-reviewed publications thus adding significantly to the development of these individuals as computationally-oriented scientists.
- Every XSEDE staff member was given a short term startup account on Jetstream, providing specific professional development activities for cyberinfrastructure professionals already working in this field with at least some part of their salary funded by the NSF award to operate the eXtreme Science and Engineering Discovery Environment (XSEDE).
- Jetstream now provides trial allocations to allow any user to test drive the system while considering their longer-term CI needs. This has significantly lowered the barrier to entry for users new to XSEDE.

4.4. What is the impact on physical resources that form infrastructure?

Nothing to report.

4.5. What is the impact on institutional resources that form infrastructure?

Nothing to report.

4.6. What is the impact on information resources that form infrastructure?

The pilot nature of Jetstream has been proved as a viable mechanism for large research projects to host and leverage as a service provider. Projects such as IRIS and UNAVCO currently leverage multiple physical data centers in different regions for service availability. Jetstream provides the opportunity for those type of projects to attain external out-of-region services. The services Jetstream provides indicate that through expansion the system could

provide widespread cloud service hosting that could replace multiple discrete data centers for domain specific projects within the NSF.

4.7. What is the impact on technology transfer?

There has been important transfer of information from the Jetstream project to major open source technology projects, particularly OpenStack. The configuration files and methods used to deploy Jetstream are also publically available and have been publically shared for community adoption. There have been no formal invention disclosures, but Jetstream staff and research faculty are playing an important role in the development of OpenStack. Members of the Jetstream team are also leading collaboration efforts in the scientific community through the OpenStack Scientific Working Group. OpenStack itself is continually transferred from its open source repository to anyone in the US that wishes to download, install, and use it.

4.8. What is the impact on society beyond science and technology?

Through our impact on the OpenStack project, we are playing an important role in the ongoing improvement of the open source cloud technology that is the most widely used open source cloud environment for small business and industry. This aids the development of the cloud-based economy in the US generally.

5. Changes/ Problems

5.1. Changes in approach and reasons for change

During this annual period the instances sizes available were changed to speed the imaging process. Users may only image instances with disk sizes of 60GB or smaller. Users may still elect to launch on larger sizes for computational work but those instances cannot be imaged for re-use.

During this annual period the placement strategy within OpenStack was changed to pack instances on compute nodes. This was done to ensure launches of full-node instance types were not blocked by tiny instances spread across multiple hardware resources.

During this annual period shelving was enabled in OpenStack and within the Atmosphere portal. Shelving frees local hardware resources unlike a shutdown or suspend operation and was not available at the time of production. It's anticipated the Jetstream team may change the charging strategy in the future to further incentivize shelving operations over a suspend or shutdown, neither of which currently incur any service units to be charged.

During this annual period SI David Hancock took over as acting PI for Craig Stewart due to health related issues. This was a documented risk that was acted upon.

5.2. Actual or Anticipated problems or delays and actions or plans to resolve them

All planned and unplanned outages for the annual reporting period are documented below. The anticipated issue during the next yearly period is increased usage of Jetstream storage. The overall storage pool is greater than 70% full. Old and unused images will be cleaned and users will be forced to migrate to newer image versions in some cases. In addition to policy changes it's likely that the Jetstream storage pool will need to be expanded to accommodate the growing user community.

Entity	Start	End	Duration	Planned/ Unplanned	Service, System, or Components	User Impact
TACC	2017.11. 02 10:40 EDT	2017.11. 02 11:05 EDT	0.42 hr (0h25m)	Unplanned	Network	TACC network DDoS. Intermittent connectivity to instances running at TACC. Atmosphere & IU cloud unaffected.
TACC	2017.10. 03 11:30 EDT	2017.10. 04 17:45 EDT	30.25 hr (30h15 m)	Planned	OpenStack upgrade	Upgrade is taking longer than expected
Atmo	2017.08. 29 12:00 EDT	2017.08. 29 20:00 EDT	8.00 hr 8h00m	Planned	Software maintenanc e & upgrades	Users were unable to control instances via Atmosphere. API users and running instances

						were unaffected.
Atmo	2017.07.25 12:00 EDT	2017.07.25 20:00 EDT	8.00 hr 8h00m	Planned	Software maintenance & upgrades	Users were unable to control instances via Atmosphere. API users and running instances were unaffected.
Atmo	2017.07.16 18:23 EDT	2017.07.16 21:08 EDT	2.75 hr 1h45m	Unplanned	TAS Outage - could not resolve logins to Atmosphere	New logins to Atmosphere failed.
IU	2017.07.04 12:27 EDT	2017.07.04 13:45 EDT	1.30 hr 1h18m	Unplanned	One controller had full disk and could not create lock files to process requests	Some networking was affected, new instance launches failed, some openstack calls failed
Atmo	2017.06.13 12:00 EDT	2017.06.24 20:00 EDT	8.00 hr 8h00m	Planned	Software maintenance & upgrades	Users were unable to control instances via Atmosphere. API users and running instances were unaffected.
IU	2017.05.23 19:00 EDT	2017.05.24 10:15 EDT	15.25 hr	Unplanned	Error in database sync script	OpenStack commands were failing.

			(15h15m)			Users would have had problems starting or stopping instances. Running instances were not affected.
TACC	2017.05.23 17:45 EDT	2017.05.23 19:55 EDT	2.17 hr (2h10m)	unplanned	Firewall rules	Inability of users to interact with the OpenStack APIs. Running instances were not affected
Atmo	2017.05.01 12:00 EDT	2017.05.01 20:00 EDT	8.00 hr (8h00m)	Planned	Software maintenance & upgrades	Users were unable to control instances via Atmosphere. API users and running instances were unaffected.
TACC	2017.04.12 11:20 EDT	2017.04.12 11:20 EDT	0.01 hr (0h01m)	Unplanned	Utility power, power distribution units	The utility supplying power to Jetstream-TACC had a momentary blip on one phase of its three phase circuit. Two-thirds (2/3) of the

						Jetstream-T ACC's compute hosts immediately ceased operating resulting in two-thirds (2/3) of user's running instances ceasing to exist. The event damaged one power distribution unit (PDU) which was replaced.
Atmo	2017.04.11 14:32 EDT	2017.04.11 15:43 EDT	1.06 hr (1h11m)	Unplanned	resource exhaustion on node serving Atmosphere	Users were unable to control instances via Atmosphere. API users and running instances were unaffected.
Atmo	2017.03.23 12:00 EDT	2017.03.28 20:00 EDT	8.00 hr (8h00m)	Planned	Software maintenance & upgrades	Users were unable to control instances via Atmosphere. API users and running instances were unaffected.

IU	2017.02.21 11:50 EST	2017.02.21 18:45 EST	6.92 hr (6h55m)	Unplanned	Loadbalancer	Intermittent connectivity issues relating to OpenStack commands. Users are having difficulties controlling instances; however, running instance are uneffected.
TACC	2017.01.24 13:30 EST	2017.01.24 17:05 EST	3.58 hr (3h35m)	Unplanned	run time cpu & memory resource issues with RabbitMQ & gnocchi	OpenStack commands failing. Running instances unaffected.
IU	2016.12.19 17:35 EST	2016.12.20 11:40 EST	18.08 hr (18h05m)	Unplanned	Complications from upgrading OpenStack required restarting stopped instances	Running instances stopped and needed restarting.
TACC	2016.12.10 9:34 EST	2016.12.10 10:06 EST	0.53 hr (32 min)	Unplanned	UT networking commodity uplinks went down. Links to Internet2 were up, so some connectivity (such as TACC-IU) stayed up.	Users were unable to start new instances or connect to existing instances on the TACC cloud. Running instances remained up.

TACC	2016.12.0 6 23:50 EST	2016.12.0 7 12:05 EST	12.25 hr (12h15m)	Unplanned	Zookeeper logs filled root partition	Openstack commands failing. Running instances unaffected.
------	-----------------------------	-----------------------------	--------------------------	-----------	--	--

5.3. Changes that have significant impact on expenditures

Nothing to report.

5.4. Significant changes in use or care of human subjects

Nothing to report.

5.5. Significant changes in the use or care of vertebrate animals

Nothing to report.

5.6. Significant changes in the use or care of biohazards

Nothing to report.