

# **Jetstream Annual User Assessment - 2018 Summary Report**

*supported by  
National Science Foundation  
Office of Advanced Cyberinfrastructure  
Award #1445604*

*prepared for*

**David Y. Hancock  
Winona Snapp-Childs**

Research Technologies  
Indiana University Pervasive Technology Institute

**Jetstream Advisory Board**

*prepared by*

**Julie A. Wernert**

Indiana University Pervasive Technology Institute

April 2019  
Version 1.0

## Table of Contents

A.	Document History.....	3
B.	Document Scope and Findings.....	4
B.1.	Findings.....	5
C.	Survey Methodology and Results.....	6
C.1.	Training.....	10
D.	Appendix.....	11
D.1.	Appendix Part 1 - Final Questionnaire.....	11
D.2.	Appendix Part 2 — Open text responses for: Please tell us more about your dissatisfaction with Jetstream services and/or resources. ....	28
D.3.	Appendix Part 3 — Text Comments for: Please tell us how or where you are.....	30
D.4.	Appendix Part 4 — Text Comments for: Please offer any comments you may have about how or in what context you would utilize Jetstream GPUs. In particular, please identify why GPUs on Jetstream are desired as opposed to integration with a more traditional batch system such as Bridges..	32
D.5.	Appendix Part 5 — Text Comments for: Please share with us any general comments about your experiences with Jetstream. (For example, are there any features and/or services not presently available through Jetstream that would benefit your work? Are there are current features that you find particularly useful? Are there features that you do not find helpful? Etc.).....	34
D.7.	Appendix Part 7 — Text Comments for: What impact has your use of Jetstream had on your research, teaching and/or training efforts? .....	38
D.8.	Appendix Part 8 — Text Comments for: Please share any comments about.....	42
	the value derived from the National Science Foundation’s investment in Jetstream? .....	42

## A. Document History

---

Relevant Sections	Version	Date	Changes	Author
Entire Document	1.00	03/15/2019	Document creation	J. Wernert
Entire Document	1.00	4/30/2019	Draft submitted for review	J. Wernert

## B. Document Scope and Findings

---

This document summarizes the results of the Jetstream Annual User Assessment conducted among 1542 users from October 23, through December 10, 2018. This was the third survey of Jetstream users. It was designed as a tool to gauge broad, overall satisfaction with Jetstream activities and services — a basic “report card.” Other surveys and feedback mechanisms are deployed throughout the project to gather more granular, project-specific, and/or resource- and service-specific data.

As with any survey instrument or resulting report, one should exercise caution in inferring too much with regard to specific results, either positive or negative. The authors made every effort to summarize and convey the survey results accurately and as received so as to not introduce any bias. Nonetheless, readers should pay specific attention to the data collection methodology, especially sources of survey error, detailed in Section C.

Please direct any questions regarding the methods used in the administration of this survey and/or the summarization of responses provided in this report to Julie Wernert at the Indiana University Pervasive Technology Institute, [jwernert@iu.edu](mailto:jwernert@iu.edu).

## B.1. Findings

This report provides a basic analysis of the 2018 Jetstream Annual User Assessment. Section C describes the data collection methodology of the survey. The survey consisted of quantitative and qualitative questions designed to determine user satisfaction of Jetstream services and resources.

Findings from the Jetstream Annual User Assessment include:

- Overall, over ninety (90) percent of all respondents report they are either “satisfied” or “extremely satisfied” in their experience using the Jetstream system. Mean satisfaction with Jetstream is 4.32 (on a 5.0 scale).
- All service areas evaluated scored above 4.0 (on a 5.0 scale.)
- Users note the highest levels of satisfaction with the quality of response to questions via direct email to Jetstream staff (4.55/5.0), followed by the speed of response to questions via direct email to Jetstream staff (4.51/5.0), the speed of response to questions via help@xsede.org (4.47); and the availability of VM images to solve my problems (4.34).
- Some sixty-seven (67) percent of respondents report that Jetstream is “very important” or “essential” to their research activities with a mean importance rating of 3.88 (on a 5.0 scale).
- Nearly sixty (60) percent indicate the resource is “very important” or “essential” to their education activities, with a mean importance rating of 3.62 (on a 5.0 scale).
- Over ninety-one (91) percent of users responding indicate the ability to preserve VMs with Digital Object Identifiers (obtained through Indiana University’s IUScholarWorks) archive as at least “moderately useful,” with a mean usefulness score of 3.92 (on a 5.0 scale).
- Nearly eighty (80) percent of users report that if Jetstream had GPUs available they would find the functionality at least “moderately useful” to their research and/or educational activities.
- Users were largely neutral to positive about about the training methods they were asked to rate in relation to their use of Jetstream, but showed a clear preference for the ability to self-serve through the use of just-in-time, online resources. In particular, web documentation and self-paced online tutorials are preferred by Jetstream users.

## C. Survey Methodology and Results

---

The Jetstream Annual User Assessment was conducted from October 23 through December 10, 2018, and was available to all Jetstream users with an active account, without regard to allocation status. (The survey questionnaire is included in Section D.1.) Users were not required to participate and were able to opt out of the survey. Those who chose to participate were only eligible to do so once during the deployment period; the survey was accessed via a custom link specific to each member of the population, and after completing the survey (or opting out), the survey was no longer accessible. Users were sent via email an initial letter of invitation, followed by up to five reminder messages sent to non-respondents and partials (those who had started the survey but had not yet completed it).









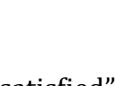
The total number of users in the population was 1542 (excluding those users whose invitation was returned as undeliverable due to outdated contact information), with 282 choosing to participate at some level. Respondents were not required to answer any questions and were free to skip questions they did not wish to answer and, therefore, the N varies from question to question. The effective rate of response was 18.2%.

The data collection instrument was submitted (under protocol #1703859488A002) to the Indiana University Institutional Review Board, and was granted “exempt” status. As such, the resulting data may be published externally in reports, presentations, and other documents, so long as data are presented in aggregate form and no identifying information is divulged. While all identifying information is redacted for public consumption, please note that responses were initially associated with the respondent’s email address and cannot be considered anonymous; in some cases, this may have a bearing on the responses provided and should be considered when reviewing and analyzing the data.

Surveys of this kind are sometimes subject to types of inaccuracies for which precise estimates cannot be calculated. For example, findings may be influenced by events that take place while the survey is in the field. Events occurring since the time the surveys were completed could have changed the opinions reported here. Sometimes questions are inadvertently biased or misleading. The views of people who responded to the survey may not necessarily replicate the views of those who refused to respond to the survey


In this survey, users were asked to rate their levels of satisfaction, ranging from “extremely dissatisfied” to “extremely satisfied,” with Jetstream services in 10 areas (Table 1), as well as overall satisfaction with Jetstream (Table 1). All areas evaluated scored above 4.0 (on a 5.0 scale.) Users note the highest levels of satisfaction with the quality of response to questions via direct email to Jetstream staff (4.55/5.0), followed by the speed of response to questions via direct email to Jetstream staff (4.51/5.0), the speed of response to questions via help@xsede.org (4.47); and the availability of VM images to solve my problems (4.34).

**Table 1. Satisfaction with Jetstream services**

Please rate your satisfaction with the following aspects of Jetstream on a scale of 1 to 5, with 1 being “extremely dissatisfied” and 5 being “extremely satisfied.” If you have no basis for rating your satisfaction, please select “Not applicable.”									
	Mean	Number of Applicable Responses	Distribution (1 = extremely dissatisfied, 5 = extremely satisfied)					Number providing no response	Histogram
			1	2	3	4	5		
Availability of VM images to solve my problems	4.34	206	1.0%	1.9%	5.3%	45.6%	46.12%	76	
Speed (responsiveness) of Jetstream	4.19	215	0.5%	7.0%	8.8%	40.9%	42.79%	67	
Documentation about Jetstream	4.06	205	1.5%	2.9%	18.0%	43.4%	34.15%	77	
Atmosphere, the Jetstream User Portal	4.05	199	2.5%	4.0%	15.6%	42.2%	35.68%	83	
Speed of response to my questions via help@xsede.org	4.47	182	1.1%	1.1%	5.5%	34.6%	57.69%	100	
Quality of response to my questions via help@xsede.org	4.47	182	1.6%	1.1%	4.4%	34.1%	58.79%	100	
Speed of response to my questions via direct email to Jetstream staff?	4.51	156	1.3%	0.6%	5.1%	31.4%	61.54%	126	
Quality of response to my questions via direct email to Jetstream staff?	4.55	156	1.3%	0.6%	3.2%	31.4%	63.46%	126	
Quality of in-person workshops and training	4.12	83	2.4%	1.2%	19.3%	36.1%	40.96%	199	
Quality of online workshops and training	4.03	89	2.2%	5.6%	15.7%	39.3%	37.08%	193	

Overall, over ninety (90) percent of all respondents report they are either “satisfied” or “extremely satisfied” in their experience using the Jetstream system. Applying a standard Likert scale to the responses offered, with “1” being “extremely dissatisfied” and “5” being “extremely satisfied,” the mean satisfaction is 4.32 on a 5.0 scale. The mean satisfaction and rating distribution are presented in Table 2.



**Table 2. Overall satisfaction with Jetstream**

Please rate your overall satisfaction with with Jetstream on a scale of 1 to 5, with 1 being “extremely dissatisfied” and 5 being “extremely satisfied.” If you have no basis for rating your satisfaction, please select “Not applicable.”								
Mean	Number of Applicable Responses	Distribution (1 = extremely dissatisfied, 5 = extremely satisfied)					Number providing no response	Histogram
		1	2	3	4	5		
4.32	216	1.4%	2.3%	6.0%	43.5%	46.76%	66	

Respondents who indicated they were either “extremely dissatisfied” or “dissatisfied” were asked to indicate why they were not satisfied. Respondents cited issues related to difficulties shelving and unshelving instances, to the lack of GPUs, to limited documentation and tutorials as reason for dissatisfaction; open text responses are included in Appendix D.2.

Respondents were also asked to indicate how important the Jetstream system is to their respective research and education activities (Table 3). Some sixty-seven (67) percent of respondents indicate that Jetstream is “very important” or “essential” to their research activities with a mean importance rating of 3.88 (on a 5.0 scale). Nearly sixty (60) percent indicate the resource is “very important” or “essential” to their education activities, with a mean importance rating of 3.62 (on a 5.0 scale). While one should take care not to make definitive conclusions without concrete data, it should be noted that twenty-five (25) percent fewer individuals responded to the importance items regarding education activities, further supporting the notion that Jetstream is generally less important to users’ educational activities.

**Table 3. Importance of Jetstream to users’ research and educational activities**

Importance of Jetstream to research and educational activities on a scale of 1-5, with 1 being “not important at all” and 5 being “essential.”									
	Mean	Number of Applicable Responses	Distribution (1 = not important at all, 5 = essential)					Number providing no response	Histogram
			1	2	3	4	5		
Research Activities	3.88	216	3.7%	5.1%	24.5%	33.3%	33.3%	66	
Education Activities	3.62	169	7.7%	10.1%	22.5%	32.5%	27.22%	113	

Jetstream allows its users to preserve VMs with Digital Object Identifiers (obtained through Indiana University’s IUScholarWorks) archive, which enables the sharing of results, reproducibility of analyses, and new analyses of published research data. With this as context, users were asked to rate how useful sharing and preserving VMs is in their use of Jetstream (Table 4). Users report that this feature is highly useful with over ninety-one (91) percent of those responding indicating that it was at least “moderately useful,” with a mean usefulness score of 3.92 (on a 5.0 scale).



**Table 4. Usefulness of the ability to share and preserve VMs with Digital Object Identifiers (via IU ScholarWorks)**

Usefulness of the ability to share and preserve VMs with Digital Object Identifiers (via IU ScholarWorks)								
Mean Usefulness	Number of Applicable Responses	Distribution (1 = not useful at all, 5 = extremely useful)					Number providing no response	Histogram
		1	2	3	4	5		
3.92	116	3.4%	5.2%	19.0%	40.5%	31.9%	166	

Respondents were also asked to share how (or where) they are presently sharing VMs and if they are using other DOI-sharing services for their products. Users who presently share VMs indicate they use a variety of services, with no clear preferred service. Full-text responses for this question are available in Appendix D.3.

Further, while the Jetstream system does not currently have GPUs, its project leadership was interested in knowing how useful having GPUs would be to its users’ research or educational activities (Table 5). Nearly eighty (80) percent of users report that this functionality would be at least “moderately useful” to their research and/or educational activities.

**Table 5. Usefulness of having GPUs available in Jetstream**







Usefulness of having GPUs available in Jetstream								
Mean Usefulness	Number of Applicable Responses	Distribution (1 = not useful at all, 5 = extremely useful)					Number providing no response	Histogram
		1	2	3	4	5		
3.55	168	8.3%	11.9%	23.2%	29.2%	27.4%	114	

Users were asked to share detailed comments about how or in what context they might utilize Jetstream GPUs. In particular, they were asked to identify why GPUs on Jetstream are desired as opposed to integration with a more traditional batch system such as Bridges. Full-text responses for this question are available in Appendix D.4.

## C.1. Training

Respondents were asked to consider their use of Jetstream and then rate their preferred training delivery formats. Users were largely neutral to positive about about the training methods they were asked to rate, but showed a clear preference for the ability to self-serve through the use of just-in-time, online resources. Table 6 presents Jetstream users' preferred training methods.

**Table 6. Respondents' preferred training methods**

Considering your use of Jetstream, please rate your preferences for training delivery on scale of 1-5, with 1 being "strongly do not prefer" and 5 being "strongly prefer."									
	Mean	Number of Applicable Responses	Distribution (1 = extremely dissatisfied, 5 = extremely satisfied)					Number providing no response	Histogram
			1	2	3	4	5		
Web documentation	4.64	135	0.0%	2.2%	14.8%	0.0%	82.96%	147	
Live, in-person tutorials or workshops	2.86	169	5.9%	26.0%	56.2%	0.0%	11.83%	113	
Live, online webinars (tutorials, workshops, etc.)	3.24	160	3.1%	15.0%	59.4%	0.0%	22.50%	122	
Recording of webinars (with minimal editing)	3.47	141	2.1%	14.2%	51.1%	0.0%	32.62%	141	
High-quality training videos	3.76	147	0.7%	9.5%	46.3%	0.0%	43.54%	135	
Self-paced, online tutorials	4.30	128	0.0%	1.6%	32.8%	0.0%	65.63%	154	

When examining training preferences by population type and/or role, preferences map closely to those of the overall population, with Web documentation and self-paced online tutorials being the most preferred methods. Data does not suggest that any particular field of study or professional role overwhelmingly affects one's preferred method of training delivery.

## D. Appendix

---

### D.1. Appendix Part 1 - Final Questionnaire

## 2018 Jetstream User Satisfaction Survey

---

### Start of Block: Block 1 - Informed Consent

Indiana University Institutional Review Board Protocol #1703859488A002/Exempt

**INFORMED CONSENT:** You are invited to participate in the Jetstream Annual User Satisfaction Survey conducted by principal investigators of the National Science Foundation-funded Jetstream project. We ask that you read this statement and ask any questions you may have before agreeing to take part in the survey. This survey is administered on behalf of the Jetstream project by the Indiana University Pervasive Technology Institute and is funded, in part, by the National Science Foundation.

**PURPOSE:** The purpose of this survey is aimed at assessing current levels of usage and satisfaction with the Jetstream environment and its associated resources and services (e.g., training, support, etc.). Survey information will be used to improve and expand the services provided by Jetstream and to aid in the decision-making processes related to resource allocation and service expansion and improvements. Survey results may also be used to inform scholarly publications, presentations, and/or funding proposals.

**PROCEDURES FOR THE STUDY:** If you agree to take part in the study, you will complete an online survey, for which there is no compensation. You will receive via email an initial letter of invitation, followed by up to four (4) reminder messages. After the initial letter of invitation, only those who have not responded will receive subsequent messages. You will have the opportunity to opt out of all future communications upon receipt of the initial letter of invitation. You will have the option of providing your name and contact information if future contact is desired. Future contact may be in the form(s) of telephone call, video-conference, in-person interview, and/or focus group. The survey should not take more than 10 minutes to complete.

**CONFIDENTIALITY:** Every effort will be made to keep any personal information that you disclose inadvertently or for the purpose of follow-up contact, as well as institutional data used to identify population members, confidential. All survey results will be reported in the aggregate and your identity will be held in confidence in reports in which the survey results may be published and/or databases in which results may be stored. Should the resulting data set be made public, it will be redacted of all identifying information. Archived data will be redacted of all identifying information and stored on secure university systems. However, we cannot guarantee absolute confidentiality. Your personal information may be disclosed if required by law. Organizations that may inspect and/or copy survey records for quality assurance and data analysis include groups such as the study investigator and his/her research associates, the Indiana University Institutional Review Board or its designees, and (as allowed by law) state or federal agencies, specifically the Office for Human Research Protections (OHRP).

**CONTACTS FOR QUESTIONS OR PROBLEMS:** For questions about this assessment, contact Indiana University Information Manager Julie Wernert at (812) 856-5517 or [jwernert@iu.edu](mailto:jwernert@iu.edu). For questions about your rights as a participant or to discuss problems, complaints, or concerns about the assessment; to obtain information, or to offer input, please contact the IU Human Subjects Office at (812) 856-4242 or by email at [irb@iu.edu](mailto:irb@iu.edu).

**VOLUNTARY NATURE OF STUDY:** Taking part in this assessment is voluntary. You may choose not to participate, to skip any questions you do not wish to answer, and/or to leave the assessment at any time. Leaving the assessment will not result in any penalty. Your decision whether or not to participate in this assessment will not affect your current or future relations with the Jetstream Project, the Pervasive Technology Institute, Indiana University, or the National Science Foundation. This study was approved by the Indiana University Institutional Review Board on October xx, 2018. Please reference protocol #1703859488A002/Exempt when inquiring.

**Do you agree to participate?**

Yes

No

*Skip To: End of Block If Jetstream Annual User Satisfaction Survey – Informed Consent Indiana University Institutional Re... = Yes*

*Skip To: End of Survey If Jetstream Annual User Satisfaction Survey – Informed Consent Indiana University Institutional Re... = No*

**End of Block: Block 1 - Informed Consent**

---

**Start of Block: Block 2 - Survey Questions**

In the past 12 months, approximately how many times have you used the Jetstream system?

I have not used Jetstream in the past 12 months

1-5

6-10

11-20

21-50

51-100

More than 100

*Skip To: Q6 If In the past 12 months, approximately how many times have you used the Jetstream system? = I have not used Jetstream in the past 12 months*

---

In which areas do you use the Jetstream system? *Please select all that apply.*

Research

Teaching

Training

Other:

---

Which of the listed Jetstream services do you primarily use?

Atmosphere, the Jetstream User Portal

Jetstream API services based on OpenStack

Jetstream OpenStack Horizon Portal

Not sure

---

Please rate your satisfaction with the following aspects of Jetstream on a scale of 1 to 5, with 1 being “extremely dissatisfied” and 5 being “extremely satisfied.” If you have no basis for rating your satisfaction, please select “Not applicable.”

	1- Extremely dissatisfied	2- Dissatisfied	3-Neither satisfied or dissatisfied	4-Satisfied	5- Extremely satisfied	Not applicable
Availability of VM images to solve my problems						
Speed (responsiveness) of Jetstream						
Documentation about Jetstream						
Atmosphere, the Jetstream User Portal						
Speed of response to my questions via help@xsede.org						
Quality of response to my questions via help@xsede.org						
Speed of response to my questions via direct email to Jetstream staff?						
Quality of response to my questions via direct email to Jetstream staff?						
Quality of in-person workshops and training						
Quality of online workshops and training						
Overall performance of Jetstream						

*Display This Question:*

*If Please rate your satisfaction with the following aspects of Jetstream on a scale of 1 to 5, with... = 1- Extremely dissatisfied*

*Or Please rate your satisfaction with the following aspects of Jetstream on a scale of 1 to 5, with... = 2- Dissatisfied*

Please tell us more about your dissatisfaction with Jetstream services and/or resources.

[open text]

---

Please rate the importance of Jetstream to your research activities on a scale of 1-5, with 1 being “not important at all” and 5 being “essential.” If you have no basis for rating Jetstream's importance to your research activities, please select “Not applicable.”

- 1-Not important at all (1)
  - 2-Somewhat unimportant (2)
  - 3-Moderately important (3)
  - 4-Very important (4)
  - 5-Essential (5)
  - Not applicable (6)
- 

Please rate the importance of Jetstream to your educational activities on a scale of 1-5, with 1 being “not important at all” and 5 being “essential.” If you have no basis for rating Jetstream's importance to your educational activities, please select “Not applicable.”

- 1-Not important at all (1)
  - 2-Somewhat important (2)
  - 3-Moderately important (3)
  - 4-Very important (4)
  - 5-Essential (5)
  - Not applicable (6)
- 

Page Break

---



*Display This Question:*

*If Please rate the importance of Jetstream to your research activities on a scale of 1-5, with 1 being = 3-Moderately important*

*Or Please rate the importance of Jetstream to your research activities on a scale of 1-5, with 1 being = 4-Very important*

*Or Please rate the importance of Jetstream to your research activities on a scale of 1-5, with 1 being = 5-Essential*

*Or Please rate the importance of Jetstream to your educational activities on a scale of 1-5, with 1... = 3-Moderately important*

*Or Please rate the importance of Jetstream to your educational activities on a scale of 1-5, with 1... = 4-Very important*

*Or Please rate the importance of Jetstream to your educational activities on a scale of 1-5, with 1... = 5-Essential*

Jetstream allow users to preserve VMs with Digital Object Identifiers (obtained through IU's archive called IUScholarWorks), which enable sharing of results, reproducibility of analyses, and new analyses of published research data. Please rate how useful sharing and preserving VMs is to you on a scale of 1-5, with 1 being "Not at all useful" and 5 being "Extremely useful." If you have no basis for evaluating the usefulness of such a community forum, please select "Not applicable."

- 1-Not at all useful
- 2-Slightly useful
- 3-Moderately useful
- 4-Very useful
- 5-Extremely useful
- Not applicable

---

Please tell us how or where you are sharing VMs or if you are using another DOI-sharing service for products.

[open text]

---

Page Break

Jetstream does not currently have GPUs. Please rate how useful having GPUs would be to your research or educational activities on a scale of 1-5, with 1 being "Not at all useful" and 5 being "Extremely useful." If you have no basis for evaluating the usefulness of GPUs to your research activities, please select "Not applicable."

- Not at all useful
- Slightly useful
- Moderately useful
- Very useful
- Extremely useful
- Not applicable

---

*Display This Question:*

*If Jetstream does not currently have GPUs. Please rate how useful having GPUs would be to your resea... = Slightly useful*

*Or Jetstream does not currently have GPUs. Please rate how useful having GPUs would be to your resea... = Moderately useful*

*Or Jetstream does not currently have GPUs. Please rate how useful having GPUs would be to your resea... = Very useful*

*Or Jetstream does not currently have GPUs. Please rate how useful having GPUs would be to your resea... = Extremely useful*

Please offer any comments you may have about how or in what context you would utilize Jetstream GPUs. In particular, please identify why GPUs on Jetstream are desired as opposed to integration with a more traditional batch system such as Bridges.

---

---

---

---

---

---

Page Break

Considering your use of Jetstream, please rate your preferences for training delivery formats on a scale of 1 to 5, with 1 being “strongly do not prefer” and 5 being “strongly prefer.”

	1-Strongly do not prefer	2-Do not prefer	3-Neutral	4-Prefer	5-Strongly prefer
Web documentation					
Live, in-person tutorials/workshops					
Live, online webinars (tutorials, workshops, etc.)					
Recording of webinars (with minimal editing)					
High-quality training videos					
Self-paced, online tutorials					

Page Break

End of Block: Block 2 - Survey Questions

---

Start of Block: Block 3 - Demographics

Please indicate the **primary** discipline in which you conduct research activities.

▼ Arts and Humanities (1) ... Not applicable (19)

---

Please indicate the **primary** discipline in which you conduct educational activities.

▼ Arts and Humanities (1) ... Not applicable (19)

---

Please indicate your status within the research community. *Please select all that apply.*

- University faculty or equivalent
  - University/Center research staff or equivalent (non-postdoctoral)
  - University/Center non-research support staff (or equivalent)
  - Postdoctoral fellow
  - Graduate student
  - Undergraduate student
  - XSEDE staff
  - Jetstream staff
  - XSEDE Campus Champion
  - Executive leadership (e.g., dean, director, vice president, CIO, etc.)
  - Other:
-

Please describe your institution. *Please select all that apply.*

- EPSCoR institution
  - Minority-serving institution
  - Associate's college (all degrees are at the associate's level)
  - Baccalaureate college/university
  - Master's college/university
  - Doctorate-granting university
  - Teaching-focused institution
  - Research-focused institution
  - Government lab or center
  - Non-profit organization (non-academic)
  - Corporate or industrial organization
  - Non-US institution
  - Other:
-

What is your gender?

Male

Female

Non-cisgender

Other:

Prefer not to disclose

---

What is your ethnicity?

Hispanic or Latino

Not Hispanic or Latino

Prefer not to disclose

---

What is your race? *Please select all that apply.*

Asian

Black or African-American

Caucasian

Native American (including Alaska Native)

Native Hawaiian or Other Pacific Islander

Other:

Prefer not to disclose

**End of Block: Block 3 - Demographics**

---

**Start of Block: Block 4 - Contact**

If necessary, may we contact you for additional feedback about your Jetstream user experience?

Yes

No

*Skip To: End of Block If necessary, may we contact you for additional feedback about your Jetstream user experience? = No*

*Display This Question:*

*If necessary, may we contact you for additional feedback about your Jetstream user experience? = Yes*

Please provide your name and preferred contact information.

Name:

Institution:

Phone number:

Preferred email address:

If additional follow-up is necessary, may we share your survey responses and comments with the Jetstream principal(s) who will contact you? (Survey responses will not be shared with those who may be contacting you for additional feedback unless you grant permission.)

Yes

No

**End of Block: Block 4 - Contact**

**Start of Block: Block 5-Open Text**

Please share with us any general comments about your experiences with Jetstream. (For example, are there any features and/or services not presently available through Jetstream that would benefit your work? Are there are current features that you find particularly useful? Are there features that you do not find helpful? Etc.)

[open text]

What impact has your use of Jetstream had on your research, teaching and/or training efforts?

[open text]



---

Please share any comments about the value derived from the National Science Foundation's investment in Jetstream?

[open text]

End of Block: Block 5-Open Text

---

Start of Block: Block 5 -Citations

Please tell us about any products (as defined by the National Science Foundation) you have produced that have benefited in some way from use of Jetstream. *Please select all that apply.*

- Not applicable
- Journal articles
- Books
- Book Chapters
- Thesis/Dissertations
- Conference Papers and Presentations
- Other Publications
- Technologies or Techniques
- Patents
- Inventions
- Licenses
- Websites
- Other:

*Skip To: End of Survey If Please tell us about any products (as defined by the National Science Foundation) you have produc... = Not applicable*

*Skip To: Q25 If Please tell us about any products (as defined by the National Science Foundation) you have produc... = Journal articles*

*Skip To: Q25 If Please tell us about any products (as defined by the National Science Foundation) you have produc... = Books*

*Skip To: Q25 If Please tell us about any products (as defined by the National Science Foundation) you have produc... = Book Chapters*

*Skip To: Q25 If Please tell us about any products (as defined by the National Science Foundation) you have produc... = Thesis/Dissertations*

*Skip To: Q25 If Please tell us about any products (as defined by the National Science Foundation) you have produc... = Conference Papers and Presentations*

*Skip To: Q25 If Please tell us about any products (as defined by the National Science Foundation) you have produc... = Other Publications*

*Skip To: Q25 If Please tell us about any products (as defined by the National Science Foundation) you have produc... = Technologies or Techniques*

*Skip To: Q25 If Please tell us about any products (as defined by the National Science Foundation) you have produc... = Patents*

*Skip To: Q25 If Please tell us about any products (as defined by the National Science Foundation) you have produc... = Inventions*

*Skip To: Q25 If Please tell us about any products (as defined by the National Science Foundation) you have produc... = Licenses*

*Skip To: Q25 If Please tell us about any products (as defined by the National Science Foundation) you have produc... = Websites*

*Skip To: Q25 If Please tell us about any products (as defined by the National Science Foundation) you have produc...(Other:) Is Not Empty*

---

Please provide citations for the products that you indicated have benefited from the use of Jetstream. Please include title, authors, publication and date, when and where presented, URL, and/or digital object identifier (DOI).

---

---

---

---

---

**End of Block: Block 5 -Citations**

---

## D.2. Appendix Part 2 — Open text responses for: Please tell us more about your dissatisfaction with Jetstream services and/or resources.

- Documentation is limited and hard to follow along with. The system works well for its intended use but expanding it to use it to support data services is a difficult task and one that can be served better elsewhere.
- Sometimes my VMs would get stuck when starting up. I'd contact support, but not hear back for more than a workday.
- Downtime of services
- We need more extensive documentation to cover all aspects of atmosphere and the basic stuff in API/Horizon
- Jetstream turns out to be a poor substitute for a true HPC system when doing distributed (MPI) computing. It is nice to have VMs but VMs do not work well for our MPI application, and a lot of adaptive work is necessary to even make it workable. Once it is working, it is slower than other resources like Stampede and Comet, however, availability is much better than Stampede.
- No GPUs, no enough tutorial for deep learning
- After shelving an instance when I get back to unshelve it, it takes forever.
- Atmosphere has big design problems. I can't easily find which of my VMs is running, shelved etc. I don't have my grants memorized so I have to look up which allocation to use for the Jetstream VM. Some of the images are broken in undocumented ways (such as one of the Ubuntu builds not supporting disk sizes > 10gb), but are still available to users to use. The way the "new" web shell handles copy and paste is atrocious. The web GUI is usually too slow to be usable.
- Instances are difficult to unshelve and often that volume is difficult to attach/unattach
- I cannot find a good image for my project to deal with sequence data. But that's not a big obstacle since I can start over building one for myself. Just take extra time.
- The I/O speed limits of data from attached volumes is a problem for some programs.
- I would like to see more user institutions added to the live webinars/training. It can be difficult to access material and/or upskill in this regard because there are a limited number of hosts for XSEDE training events.
- I allocated the largest machine with largest amount of cores and memory but somehow it cannot run jobs which can be handled by a 120GB memory desktop. I am still trying to understand the system fully and so I am a bit frustrated at the moment. However, I will keep persevering.
- Jetstream instances are often conky - especially during shelving/resumption. There are periods of such conkiness, interspersed by periods when everything works great
- No good was to manage use by an entire class; no accessible interface for our visually impaired student.
- When I start a session and suspend it, I am not able to resume it. I opened several tickets and the problem never solved. So, I was not able to use the system really.
- The only problem is that sometimes it takes up to 45 min to launch an instance.
- To be honest, I said I have used it 1-5 times and that is true but I have yet to successfully log in. I always work on getting it to work, either get frustrated or send a help email get an answer that I don't understand and then have to do something else, coming back to this project later. Then when I come back I have the same problems. I spend several hours, get a few steps further, then get stuck, send an email, don't understand the answer and have to stop and do something else. I did this ~5 times over the last year and although I have an allocation I have never had the time to actually get this to work so I can use it. Instead, I go back to things that work (a server I have access to) to do my work. I would like to use this for teaching but I don't think this is going to work with a class of 16 beginners if I can't get it to work myself. It's too bad. It's a great idea.

- Several areas of my help requests were categorically dismissed as being out of scope for user documentation and were thus condescending, inadequate, and unprofessional.
- Overall performance and availability of the system is less than I would expect from a national resource. Staff seems to be overworked and short-handed. The system is missing the auxiliary services and support needed to operate as a production resource.
- Limited resources for large scale simulations.
- Instance launching is too inconsistent
- Usability of the Atmosphere web portal can be improved. One major issue I was facing is there's no way to control (start/stop etc.) group of cluster nodes together. Need to go one by one.
- Not enough images, especially for the most recent distribution releases
- Some sessions wouldn't shut off, costing us valuable hours.
- The atmosphere interface is poor, it has lots of quirks, and misses many useful features

### D.3. Appendix Part 3 — Text Comments for: Please tell us how or where you are sharing VMs or if you are using another DOI-sharing service for products

- I primarily work with containers, and distribute through Docker Hub (or similar).
- Our institution uses VMs in classroom settings.
- I have used NSERC and Zenodo for materials in support of publications because I did not know this was available for Jetstream.
- Plan on using them in new project - IMLS funded LG-70-18-0202
- Tried to utilize this as an ftp service for archived data. It doesn't share because of firewall restrictions.
- Directly through Jetstream
- Research and testing of new applications, features, databases.
- our website has links to the images
- I'm not using this but I believe it should be useful
- I was not aware sharing VM by DOI was possible, but this is a feature we will definitely find extremely useful.
- As a service provider, our clients make use of Jetstream and we find the availability of such a platform in the NSF portfolio invaluable. The ability to transfer VMs amongst platforms makes Jetstream very useful.
- I share VMs for my class
- No sharing it
- I have put VM images in OpenStack Horizon & Atmosphere for others to use. e.g.  
[https://genapp.rocks/wiki/wiki/instance\\_jetstream\\_atmosphere](https://genapp.rocks/wiki/wiki/instance_jetstream_atmosphere) &  
[https://genapp.rocks/wiki/wiki/instance\\_jetstream\\_api](https://genapp.rocks/wiki/wiki/instance_jetstream_api)
- We are in the planning stages of using this feature (not ready to be archived yet), but we are looking forward to having a DOI for VMs!
- I am not using any
- We are planning to use DOIs to curate anatomical atlases derived from 3D scans of specimens (pending grant application)
- IU ScholarWorks, F1000
- It sounds very useful to me, but I never use it. Didn't know it until answering this question.
- Have not actually shared a VM. Didn't know it was possible!
- NA
- eventdata.utdallas.edu
- A Science Gateway allows users to run experiments in Jetstream.
- The Jetstream cluster is made available to our scientific gateway (ultrascan3) users.
- Zenodo
- We currently use zenodo, but I like the idea of VM preservation for this purpose, and would use it and support it.
- I am not using VMs.
- I usually create my images using the Horizon Dashboard at IU. This serves to share my personalized instances with dozens of students and researches which allows speeding up the learning time or sharing an environment preconfigured for our research activities.
- Docker hub
- We are doing this for our papers and book chapters from the Brendel Group- so that our analyses and workflows can be reproduced.
- Having the ability to archive a VM and then restart it is very useful.
- NA
- not applicable

- I share VMs with my students using the 'adduser' function
- I was not even aware of this feature.
- My students typically use AWS, Azure, or Google Cloud, because that is what they are familiar with from Industry. Last year I used Comet and Bridges to support a course, and the students struggled to get up and running. I think Jetstream might lower that barrier, so I would like to try using it, even though I haven't used it yet.
- We are using DataCite for DOIs
- I did not know that this was available. Our group is working to improve our reproducible research practices, currently use GitHub and uploaded data (with DOI into institutional repositories (for the UC system - DASH) to accompany papers. In future we're interested to made data available using mybinder.org. The way I've seen this done so far is to mint a DOI for a version of a release on a GitHub repository (e.g. with Zenodo) and have a link to open a Binder there. If in the future it was possible to set up a binder with associated code and data and host the compute on a service like Jetstream that would be good to know - I will stay tuned for what develop since am still learning about best practices for different types of projects.
- N/A
- Currently not using this service. Would start exploring the same.
- Would like to get a shared VM with the software needed for RNAseq or mapping etc. for research and teaching so I don't have to install it myself.
- Not particularly useful for \*me\*, but I always encourage researchers to use this functionality.
- I didn't know that this feature was available. Thank you!
- IU ScholarWorks
- We don't share VMs, but we do share Docker images
- I have not utilized the DOI-minting service yet but imagine it would be extremely useful once I have a VM that I would have others use in order to reproduce results generated on it.
- I am not currently sharing VMs
- I don't down
- N.A.
- My co-worker wrote a python script to transfer data to and from the VM. The script also performed a Finite-Element calculation on Jetstream.

#### **D.4. Appendix Part 4 — Text Comments for: Please offer any comments you may have about how or in what context you would utilize Jetstream GPUs. In particular, please identify why GPUs on Jetstream are desired as opposed to integration with a more traditional batch system such as Bridges.**

- Thanks for making this resource available! It has been immensely useful for our two NSF funded projects (Interactive Parallelization Tool and BOINC@TACC). Our students have benefitted immensely in developing their cloud computing skills through our allocation on Jetstream. We are also using Jetstream as the backend infrastructure for a mobile application from the health informatics domain. Expect when Jetstream goes down, perhaps due to unavoidable reasons, we are making constant use of it in our projects.
- Dynamic distributed deep learning system architectures are a must to research how advances statistical algorithms such as deep learning algorithms scale in a cloud environment. As we, as scientists and engineers, explore the boundaries of science and computation to build these algorithms at scale - to aid humans in a variety of tasks, an academic cloud such as Jetstream, in my humble opinion, should provide a pool of state-of-the-art GPUs connected through NVIDIA NVLink interconnects to provide unparalleled research compute hours for serious researchers.
- Thank you, Jetstream has been 4 to us.
- N/A. I need much more memories for my bioinformatics works.
- Thank you for providing this service
- I think Jetstream is a great service and wish I had more time to explore ways of using it effectively. I recommend it to researchers here at UC frequently. Please keep up the great work!
- Visualization with VisIt
- We could compare results across platforms. Some of our models are trained faster on GPUs.
- My research is shifting towards machine learning and assistant-based interfaces. Having GPU availability would enable faster processing and better throughput across the board.
- Would use to have faster Machine Learning and Image processing available to workflows
- Many of our services (gateways) are in support of researchers, often supporting ad/hoc and interactive analysis environments. For researchers that require GPUs, integration with batch systems is certainly an option, but requires bridging cloud-to-batch environments, managing interactive job submissions, and routing web traffic.
- GPU's look to be a very good fit for machine-learning problems that could be applied to the large datasets generated by (or accessed via) Jetstream nodes.
- Incorporation of more GPUs into Jetstream will be very helpful.
- Genome assembly using MEGAHIT- GPU VMs could potentially accelerate assembly time and reduce the need for longer-running VMs. Further, use of Jetstream is preferable because of its' "on demand" capability, and the level of interaction I can have with running jobs/VMs.
- Jeremy Fischer has been an essential resource in enabling my lab group's use of Jetstream. A heartfelt thanks to him and the rest of the Jetstream team.
- GPUs are essential for modern machine learning
- GPUs would be good to allow access to additional research areas.
- I appreciate very much the sincere effort of Jetstream staff to work with us to help us setup our gateway jobs on Jetstream.
- Bridges has GPUs which is why I rated 3 as the option for this question. There is no strong compelling reason to upgrade Jetstream with GPUs with Bridges available.
- Would use GPUs for traditional brute-force computations of wave-propagation and dynamics.



- Same advantages Jetstream CPUs have over Bridges: ability to do the work over the web, e.g. for teaching and tutorial, which would have a steep learning curve for attendees if using Bridges.

**D.5. Appendix Part 5 — Text Comments for: Please share with us any general comments about your experiences with Jetstream. (For example, are there any features and/or services not presently available through Jetstream that would benefit your work? Are there are current features that you find particularly useful? Are there features that you do not find helpful? Etc.)**

- I haven't used it.
- The web access makes using the resources easy, and I'm therefore more likely to use them.
- Its good. I WISH we had bigger memory instances
- Charging SUs for stopped VMs is mean
- Jetstream as it is fills an essential computational niche for my research group- wherein relatively short-running but highly distributable code in run on traditional HPC, and long-running, less distributable, and relatively memory intensive codes are run on Jetstream. As such, it is an essential component for us. One potential upgrade that could be of utility for us would be an option for a larger VM size with higher memory/ process per node capacity, as with some genomics applications (generally moderate benefit of core count, and high memory) I still rely on large-memory condo nodes attached to our local HPC.
- It was an assigned project for work. My interest was minimal. I did my job and moved on.
- Customized VM configuration (CPU, RAM, disk and etc.) could be very helpful.
- I would like to be able to start and stop large batches of vms
- I had to set up NFS to share volumes across multiple servers. This was especially painful because each time I restarted instances, they got different IPs, which meant that I had to update the IPs in my NFS configuration. It would be nicer to have a more easily configuration solution for shared storage, like Amazon Elastic File System.
- Extremely glad to have the service available for things I can't do on a traditional cluster
- Better documentation of Openstack APIs
- Thanks so much for making this resource available! It has allowed me to teach a research-based genomics class in which students are working on some pretty neat projects. Occasionally there are hiccups (e.g., instances fail to launch successfully or are slow to launch), but I know Jetstream is relatively new and I'm hopeful that these will be worked out as it matures and the user base grows.
- Overall Jetstream is great. There are some hiccups here and there with the pages taking longer to load. Support staff is extremely helpful and very quick to respond to any questions. It would be great if there would be bare bone VM images. Just the packages to start up a VM. It has nothing to do with space or the performance but more with the burden and the security. More stuff comes with a VM image, more stuff to maintain and worry about. I am not a system administrator, so looking though all the installed packages does not really tell me what stays and what goes. And I want to make sure my VM comes back up after all the cleaning.
- Overall, great service. VMs launched via Atmosphere interface can take "a while" (e.g., > 10 minutes) to start.
- Large performance deviation between the IU instances (very fast) and the TACC instances (very slow)
- Overall, we have been very satisfied with the services available on Jetstream.
- It would be really nice to have consistency between which VMs are available on Atmosphere and which are available through the OpenStack CLI. As of right now, they are not necessarily the same, which can be inconvenient at the least.
- Very good experience so far and they are very responsive
- VPN capabilities

- Would prefer it operated as production infrastructure with fewer outages.
- Most of my SUs were eaten up when my VM was idle
- Well, I already mentioned GPUs. Jetstream has been a great resource. I use it for Science Gateway production and development activities, training thru hosting demos, it was used to start up dedicated VMs for summer students and postdoc staff.
- More data storage, it is a huge limitation to only have a small amount of data storage available on the VMs.
- Transferring data is in general a hurdle especially from institutions that requires VPN in first.
- I use Jetstream to teach Singularity container. So far I have to start each VM manually. It would be useful to be able to fire them all up with a single command or button.
- The old web shell and GUI are pretty useful for fast access to the running instance. I think Atmosphere has a lot of potential, but it needs some dedicated user interaction work done to make it less painful to use. More documentation on getting started with the API would be good - the current documentation is good but it assumes a level of familiarity with the system already.
- I am very satisfied with, and grateful, what I can access now. I like the atmosphere most just because it's easy enough for my non-informatics background to do some NGS/genomic studies.
- We would benefit by having static IP address, since we provide a gateway access to some of our data. Also Guarantee that our data will be backup.
- It has been much easier to use Jetstream than AWS both for research and training.
- The feature that I love about Jetstream is that the team is very responsive and knowledgeable. So, they are able to help with our questions/requests in a timely fashion.
- Overall it works well. I would like to keep using it.
- It is awesome. Currently, I am working on the publication along with my collaborators to indicate XSEDE's Jetstream in our educational course.
- There are occasional problems with availability or responsiveness of compute nodes.
- Jetstream is extremely important for ongoing services (web services) and research computing. Mounting Wrangler disks are very important and hopefully such services are continued.
- Good
- Jetstream user support (hi Jeremy!) has been VERY helpful. Also I've encountered issues in customizing my Jetstream VMs (installing new packages), which we haven't fully resolved yet.
- I would be great to have more documentation/VM to run Jupyterhub servers launching custom single-user Docker containers
- It seems that Jetstream is not as good as Big Red 2 at IU
- I would love to see more RAM available to users, especially so that we can use a smaller 24 core instance rather than our current XxLarge config. We do this solely to access the RAM we need to run our analysis. I would love to be able to conduct hi-res model simulations in the near future, but the system is still somewhat limited in this capacity.
- I have not yet used Jetstream.
- My most limiting feature is storage space. As a student I have access to 100GB only, and with the explosive expansion of biological data this is hardly enough. Recently I have been using Jetstream for developing and testing my tools, but I can never do a complete research experiment using my data (~750GB).
- Jetstream has been extremely useful and reliable. There have been several issues that I have encountered that made using Jetstream difficult.
- I am still new to using the system and so it is a bit premature for me to comment. I would like to see more useful online training.
- Current platform offers a variety of services that allows me to speed up my research progress. I hope we could have Tap-as-a-Service and IoTronic or similar services for the development of IoT based research.

- The concept of SUS in Jetstream is somewhat confusing for users that are not in tune with HPC systems
- An interface for visually impaired users is sorely lacking.
- training: easy to access and use.
- I have encountered recurring problems unshelving shelved instances. This problem is quickly fixed when I report it, but this has been a source of difficulty, particularly for in class teaching.
- I want my students to access my VM so we can use the same resources (datasets, software) etc. but this does not seem possible.
- We are trying to replicate services we have on our existing servers. Because some features like NFS are not available we are having to spend time working around with other solutions.
- I was hoping that more basic Ubuntu or other Linux distro images could be included on Jetstream to start building from.
- Some suggestions: 1) A few times I was having trouble with instances failing to unshelve, it seems that this has been resolved (and someone sent a prompt and detailed reply to my service request). 2) It looks like the pop-up messages for Suspend, Shelve, and Stop on the Jetstream GUI don't match the new (since Oct 1) unit charge rates. 3) To set up RStudio (w/ tidyverse + sparklyr) in Jetstream, I didn't feel like I found good instructions in the Jetstream documentation, but was successful using these tutorials: <https://angus.readthedocs.io/en/2017/visualizing-blast-scores-with-RStudio.html> 5) If we use the credits more slowly than anticipated, I wish there can be a way to extend the length of time I can use the initial allocation over
- Would prefer if VM's shell defaulted to a more standard Linux prompt.
- Brilliant effort by IU/PTI to fill a unique niche that appears to be growing rapidly now...partly due to the success and adoption of JS!
- very reliable
- I thought the system worked well for what we were using it for.
- Overall is good. But I did face one fatal problem with Jet that one day my image cannot be recovered, and staff told me that it is permanently lost. I lost some results and some parts of codes. Other than this, Jetstream is good.
- I need more memories than Jetstream can provide.
- Even if certain resources for training are not made available via Jetstream documentation, pointers and links to pertinent resources would be greatly appreciated.
- Auto-scaling would be nice to have
- Please Can a window of support be given to researchers from Africa who may not easily get collaborators in the US?
- Very pleased with Jetstream. Please keep it going and expand.
- The reason why I did not answer many questions is because as PI, I do not personally use Jetstream. My students do. So I am not qualified to answer these questions.
- The staff has been amazing. Increase the support staff and invest in operational hardware to run the infrastructure. Reliability and performance have cost me significant opportunities and time lost over the past year.
- More storage (both for data storage and ephemeral disk) and the addition of GPUs
- Overall, I've been very satisfied with Jetstream. We do not use Atmosphere at all. We have encountered problems at both sites and our tendency is to move services from one to the other. Early in our use, we found IU to be more stable. Now we found TACC to be easier to provision on.
- GPU's; dynamic provisioning of VM's
- I find the Jetstream interface very user friendly and easy to navigate. From my use thus far I have no further suggestions for changes.
- The ability to launch GPU instances, in either single or multi GPU configurations would be greatly beneficial to advanced research questions and objectives
- Good features, easy tutorial would be nice

- It was an unexpected computational resource for our lab. The user interface and online documentation could be much more user friendly. Once I had a python workflow, it was a useful resource for my research.
- Missing feature: Ability to control multiple VMs together. Ex: start a cluster of 20 nodes using one click.
- Maybe this isn't the place to say this, but I went to watch a live webinar at Purdue University. They seemed a little shocked that someone outside of Purdue would attend, and at first had no way for me to connect to the internet. They ended up loaning me a laptop so I could use someone else's account. Nice that they worked things out, but also surprising as an advertised host site that they weren't ready.
- I'm too new to comment
- Haven't had the opportunity to use the resources yet as Bridges has been doing well for me.
- Better instance management

## D.7. Appendix Part 7 — Text Comments for: What impact has your use of Jetstream had on your research, teaching and/or training efforts?

- None.
- It has increased the breadth of my education arsenal. I'm able to teach many more concepts.
- Terrific - we published 3 papers this year built on work that we conducted on Jetstream
- It has helped provide VMs for running custom code that would otherwise require a paid service such as AWS or GCE.
- Greatly increased productivity by decreasing wait time (e.g. for public queues on traditional HPC) and making it possible to have multiple long-running codes at once, which previously would have been limited to 1-2 local workstations.
- Frustration
- Jetstream has been really useful in training biologists to start running programs using command line
- With Jetstream, I can easily distribute my work into many cores without worrying about dependency issues of my custom code.
- Please see my XSEDE summary and renewal for a detailed report!
- It allowed me to perform a large-scale analysis which would have been cost prohibitive otherwise.
- Has enabled new areas of research, especially in real time analytics, that would be very difficult on a traditional cluster or local resources
- Essential to provide students with a scalable Jupyter notebook environment for a data analysis class
- I previously taught this class using another HPCC, on which students didn't have root access, which made it difficult to install custom software. Jetstream has opened up many new possibilities for us and simplified the process for students who are new to the command line (e.g., no need to worry about job schedulers and submission scripts anymore). Thanks again for all you're doing.
- It gives me flexibility to test different tools for my projects.
- It has allowed Unidata to make a serious exploration and transition into cloud computing for the benefit of our community.
- Helped me convince my management to establish a similar service
- Very useful for educational purposes when needed.
- This has been a great asset for standing up training/workshop environments.
- Jetstream resources are crucial to delivering cloud-based data and software services to our community.
- It has increased opportunities to explore and try new things.
- Easily sharing virtual machines with other people.
- We learned a lot
- Jetstream is facilitating our high-performance computing education efforts significantly.
- Jetstream is great for overflow work when other clusters are busy and have long wait times, since Jetstream seems to be undersubscribed and therefore highly available. This greatly facilitates our workflows, I just wished Jetstream was an ordinary HPC cluster instead. I have also migrated our web-based environment to a permanent VM hosted on Jetstream, but a different hosting service could be used if Jetstream were not available. I am not convinced the expense of Jetstream is warranted by the usage pattern, at least for me. Resources may be better used for different, more urgently needed configurations.
- Potentially providing a scalable solution to teach more students
- Greatly accelerated the pace of my research

- Jetstream has been an extremely valuable asset to the success of the Science Gateways Community Institute
  - 5%
  - Jetstream has simplified my work. It has minimized my reliance on dedicated hardware and provided an easy way to support new developments and gateways. It would be a real effort to have to setup servers that are publically accessible. e.g. We have a current collaboration between NSF's GenApp project and their Seedme2 project. We simply setup a dedicated VM image and we are testing integration there. Another example was changes to the UTHSCSA infrastructure which required moving the NSF & NIH funded UltraScan project servers with little notice. I assisted PI Demeler with the transition he was up and running in relatively short order. It would require significant expense in \$ and effort to continue research without Jetstream or a similar could service. Please pardon any typos, as this is being written in a single line form ;)
  - Wonderful resource, which has had a tremendous on research, teaching, and training. I use Jetstream powered VMs in my course right now, use it for my research, and am working on a NSF CyberTraining grant using Jetstream resources.
  - It has saved our research progress as our physical server had issue that IT teams couldn't fix yet. Also the flexibility of increasing computing power and storage is very helpful so we don't need to limit our analysis to match our current system.
  - Jetstream allows me to give students root access which is not possible on traditional clusters and which is essential in teaching such topics as Singularity
  - Our team has been able to finish tasks more quickly by running test and devel. activities through Jetstream. We save lots of time not doing all of the sys-admin work required to spin up an OS.
  - I think the usefulness of on-demand VMs with a fresh operating system can't be understated. I find myself spinning up systems to test software and run small jobs all the time.
  - I am allowed to learn the skills to process seq data, apply the skills to my current research, and discover some exciting biology from genomic points of view. This is definitely crucial for my research career.
  - It helped us to share our calculations with contributor around the world.
  - Jetstream allowed me to teach a bog data course without the need for me to purchase a cluster or use AWS (which is, contrary to Jetstream, not easy to use and much more expensive)
  - It has enabled a paradigm shift in my research
  - It has helped graduate students publish findings in top venues and graduate with doctoral degrees. The training on Jetstream, Atmosphere, and OpenStack has helped my students get jobs in the hi-tech industry.
  - It sped up development at some points because I didn't have to wait in queue for HPC resources.
  - Has been extremely helpful in hosting an interactive website on a virtual machine used worldwide with very good response time.
  - It has directly led to research publications.
  - Our CUNY students are able to learn cloud based research computing. It would never happen without the Jetstream allocation and services provided to us.
- 
- Major instrumentation and API support for NSF funded project.
  - I will probably use Jetstream for education if I, as an instructor, have a tutorial session to teach students and other researchers.
  - Jetstream has been fundamental in the development of a new optimization framework that relies on evolutionary algorithms. Without access to a resource like Jetstream the development and deployment of such a tool would be extremely challenging.

- It has been a valuable additional asset.
- It has been vital for the research, training and educational aspects.
- I work all the time on it
- Huge help to my research efforts, the primary channel by which I get computational text analysis work done.
- We used Jetstream for the hands-on programming sessions of a summer school on inverse problems. We run a Jupyterhub via Docker containers so that students could interact with our Jupyter notebooks directly from their browser.
- Reduced time to science. Offloading tasks from HPC to Jetstream to enable more efficient use of limited resources. A great teaching tool because it enables student access in an efficient and streamlined manner.
- Under development.
- N/A
- We are able to check our web application in time.
- As a bioinformatics student, Jetstream is a very valuable platform to build, test, collaborate with other teammates, and distribute our workflows. It proved to be very helpful for biologists who aim to utilize our workflows and tools and do not have the required programming/Unix experience themselves. Publishing our work for others to make use of is very rewarding, yet not-so-trivial task. Thus, Jetstream provides a mean to this end.
- Jetstream let us process the final large set of data. The processing would've taken months instead of a few weeks. I am also about to save the stored data in one place because there is enough room. I did not have to split it and put it across the other small servers that I have.
- Still awaiting.
- Great impact. With Jetstream we are able to enable user with no Linux background with powerful and ease of use systems hosted on top of VMs such as Jupyter Lab and Elasticsearch.
- Moderate- I have not had as much time as I would have liked to use the system
- Extremely useful.
- My use of Jetstream has had positive impact on all three areas of effort.
- It has played a central and critical role in the bioinformatics class I am teaching this semester, and is a fantastic resource. Thank you!
- Helped me transition into the text-based interface and the flexibility for using different software
- Our faculty at Doane love the service that Jetstream provides. The ability to make vms easy to access to students is very helpful in our curriculum.
- It's let me access a powerful Linux-based development environment anywhere I can access the internet.
- As a grad student with a windows laptop and a limited budget in a group without a server, it is fabulously helpful to be able to work 1) on a Linux machine 2) with much more RAM than my laptop, 3) That can install and run software I haven't been able to get working on Windows (sparklyr for Spark in RStudio) and do it so much faster to boot! Version control makes it easy to push / pull to external repositories from Jetstream. I was fortunate to learn about Jetstream through a class with Carl Boettiger, and am now using Jetstream for my own research work. I recommend Jetstream to other students (my faculty didn't know it existed), though I don't know that any have tried it yet. I am definitely still towards the beginning of learning to \*fully\* use Jetstream's capabilities, and it took me a while to figure out how to set things up on my own, but now that it's up I have been using it every 1 or 2 days recently notice I am happy and excited to log into it. Seriously, thank you so much for making this available and the work you do to maintain it.
- Jetstream has allowed us the ability to produce results that we would be unable to afford in the commercial cloud space
- Allowed us to set up VM resources without cost and effort of using in-house research IT resources.
- N/A



- I could build an application that could not be easily built on the available XSEDE HPC systems.
- High impact. I rely on Jetstream for all computational needs.
- It was required for the labs in one of my courses.
- Most of my research results pops out from Jetstream.
- I could not have finished my PhD without it. Other platforms were not able to install the programs I needed.
- Our primary teaching related compute is handled using Jetstream cloud. Each student gets a VM, spun using a custom image we develop and maintain in our lab. These VMs contain a Jupyter Lab interface which provides students with all required packages for the coursework throughout the semester.
- N/A
- It really helped in getting my experiments done faster.
- Here in Nigeria, when I started with whole genome sequencing, I needed speed that Atmosphere brought. I was able to do my first whole genome sequencing. That has changed my career. I hope to use it for my long reads annotation next year.
- High impact; very worthwhile resource.
- Jetstream has saved me on many occasions and given me the flexibility try new things quickly and repeat-ably. It has overall been positive and allowed me to continue exploring areas that would have otherwise been closed to me.
- Extremely important to my EOT and Research efforts
- Jetstream has been an essential resource for many of the projects I work on.
- It's been an essential part of the NSF-funded research project I have been involved in, which has had as one its primary goals to "bring the computing, analysis, and visualization" all together on one cloud-based platform.
- Jetstream has been indispensable for my research. Our lab is predominantly wet lab with no computational infrastructure set up. Having access to computational resources via Jetstream has enabled me to pursue new avenues in research and has resulted in multiple collaborative projects.
- It has been central to my research enablement on XSEDE. It has helped me secure a faculty position in my career advancement, and has helped secure additional research grant funding within and outside NSF.
- Useful for uninvested students
- I used it for analysis. Specifically, performing FEM calculations.
- Currently very important for research; in future, likely to be important for teaching
- None as of yet.
- It will have a large impact on the upcoming eRPID grant. This will allow developing an Open Stack installation of important data fabric services.

## D.8. Appendix Part 8 — Text Comments for: Please share any comments about the value derived from the National Science Foundation's investment in Jetstream?

- It's a valuable resource that adds value to my instruction.
- Its high. We could not be doing the research we are doing without it.
- For researchers requiring cloud computing resources (which is increasingly the compute model being employed), Jetstream is an invaluable resource.
- It's definitely a worthwhile allocation of the NSF's funding for computing resources. Many researchers in my field are stretching their budgets to purchase time on private cloud systems (e.g. Amazon) because the previous XSEDE-funded resources were not applicable to their uses, so if anything I think Jetstream could benefit from better advocacy in certain scientific communities.
- It has no value to me
- It has been amazing and should be continued
- Very helpful in supporting a faculty who taught
- Without services like Jetstream, it would be impossible to teach cutting-edge, hands-on computational classes like this at undergraduate-focused institutions that can't afford to have HPCC's of our own. This is definitely a worthwhile investment because it opens up opportunities for students at non-R1 institutions to get practical skills that are in high demand these days.
- Jetstream gives a pathway for the Unidata community to start using the cloud in cost-realistic manner.
- High value as there are few other options.
- This is an extremely useful service, especially for a domain that is becoming largely computational, yet lacks computational training.
- This is very useful for short term research/development environment for our users as well as creating reproducible environments and training resources.
- As stated earlier, Jetstream resources are crucial to the provision of our data services and we greatly appreciate NSF's investment in science cloud infrastructure. We ask that a science-focused cloud infrastructure be sustained.
- It's extremely good tool
- This service is invaluable to both my research and educational activities.
- An essential part of modern research.
- Jetstream facilitates our workflows.
- Very pleased to see a national resource like this made "free" for educational purposes
- Jetstream has saved me months of lost time by providing the computational infrastructure I needed when it wasn't available anywhere else. I think it's a stellar system/program that is much needed right now by the research community, and will become much more needed in the future before it becomes less needed.
- Having a cloud resource where services can be hosted is very useful.
- Jetstream is a fantastic resource by NSF. There should be more such systems available publicly to complement a cloud market dominated by private vendors.
- Not aware of that
- Jetstream has enhanced my team's productivity on work for my recent NSF award and engagements activities with other NSF funded projects.
- A huge value. I have used Jetstream in my class (taught in the College of Liberal Arts) to grad/undergrad students who have never heard of cyberinfrastructure before. I am using it in my research and it makes collaborative science much easier to have these resources available.

- I think it's a great investment that will revolutionize how labs conducting research, especially reducing cost for spending on server and maintenance, and guarantee for a consistent viable environment.
- Very useful for cases when root access is needed. For example, teaching system administration, containers, etc.
- We like using the system and have grown to count on it. I hope NSF keeps this around instead of "dropping the ball" like they tend to do. How would the Interstate highway system work for us if it was built once and then never had another nickel spent on it. We'd be back in the stone age under that operating model. Some computing is "infrastructure" and if we don't maintain and support it, cool things won't happen "down the road". NSF needs to go visit the dept. of Transportation and learn what it takes to keep infrastructure going.
- Having a research cloud that is owned by research institutions makes a lot more sense than paying for time on systems optimized for large corporate markets.
- It's invaluable for me. I believe it would be the same value for many.
- It is nice to have a common place where we can create servers and service. In the past many institution provide the same thing, but it got research staff confuse where things stand. Once we have a unique place where everyone goes to get a VM, I believe it makes the whole process easy for the entire NSF community.
- A platform such as Jetstream was a missing in the set of platforms provided by XSEDE. I cannot imagine teaching my class on data analytics on a high-end cluster. I tried AWS in the past but the effort was time demanding and unsuccessful. The use of Jetstream was quite a different experience with great support, low learning barrier for my students, and great training support.
- This kind of asset is key to future atmospheric science research
- Jetstream infrastructure is vital to the research conducted by the PhD students in my group.
- Increasing access to HPC resources, especially for scientists who aren't affiliated with an institution that has its own HPC resources, is a tremendous benefit.
- Whatever you're investing, please double it. Jetstream's research value is high and Indiana University is doing a fantastic job of running it.
- Beyond compare.
- High value -- saves us from redundant resources.
- Substantial
- The investment from the National Science Foundation in HPC resources has unlocked our ability to answer complex questions about the sustainability and resilience of our natural resources under different management strategies, a critical issue to society. Moreover, resources like Jetstream present an opportunity for education and training of future researchers that would otherwise not be possible.
- Valuable
- It is a great resource and hope NSF continues and expands such virtual environments for research, education and training.
- Very helpful, thanks!
- It supported a lot in developing IrrigWise web application.
- Very useful!
- Greatly valuable. Jetstream is helping out to accelerate research and educational efforts at unprecedented speeds. I wish I had this power in my hands when I was doing my Master's Degree
- Investment in Jetstream is very valuable for researchers who do not usually need access to a full-blown cluster
- Jetstream is an extremely valuable resource that I'm grateful the NSF has invested in
- Very valuable as it allows people from resource-limited institutions and those with limited knowledge to have access to and use HTC facilities.

- It is a model system. Wish it could be funded further and provide a more stable resource for continuing research projects.
- This is an enabling and critical resource which drives not only research but training of scientists in quantitative sciences
- Jetstream is an extremely valuable resource for training and research; the ROI in terms of opportunity for budding computational researchers easily justifies NSF's investment.
- Jetstream is an excellent resource in my field (genomics) and has been invaluable for my research and teaching. I think the NSF's investment in Jetstream is extremely well justified - perhaps one of the best uses of NSF funds I can think of.
- Could not have done any Big Data type of work as my institution does not provide any such infrastructure
- I feel it is important for NSF to provide facilities like this so that better use can be made of computing resources (both hardware and human) and it also provide ability for researchers from having to spend money on resources that can be used jointly with others, reducing the cost of having to do research.
- Teaching computer science and computational science is much easier when our entire class can spin up VMs on atmosphere and learn in a controlled environment with their classmates.
- See previous response for the value of Jetstream generally
- Jetstream has allowed us the ability to produce results that we would be unable to afford in the commercial cloud space
- Jetstream resources saved us cost and time.
- This is one of those instances when the NSF investment on something "experimental" (at the time it was proposed) has paid off substantially! In fact, institutions are referencing the JS model to deploy similar but smaller systems on their campuses!
- Extremely valuable
- A great resource for educational communities, and classroom work. I wish I had more time to explore and do my own research using this tool.
- Thanks.
- providing great and essential infrastructure for data-intensive research
- Very valuable support for my research
- Very useful
- This is a VERY important project and necessary for research and teaching. However, as I described earlier, the way it currently works is not user friendly enough to make it widely useable for non-experts.
- Jetstream is still a young product that needs to spend time realizing its niche and focusing on supporting the full gamut of the "long-tail of science" that it claims to focus upon.
- NSF must continue these kinds of efforts. It is essential that we put compute resources in everyone's hands regardless of financial considerations.
- Great value and investment should be continued.
- Jetstream is an important computational option for researchers
- While it is easy to quantify the value of Jetstream with respect to time savings, the value of the community that has latched onto it and the opportunities it opens up to leverage new technologies and approaches to research has been invaluable. If anything I believe NSF should double down on the investment and make it a first class, persistent resource for the national community. This should be the bridge for academia to and from the commercial cloud.
- I encourage NSF to renew funding for this resource such that it can expand its capacity to meet the cloud needs of researchers.
- I work on multiple NSF-funded cyberinfrastructure projects that lack budget for VM infrastructure. Jetstream has provided a simple, reliable, and consistent mechanism to acquire essential resources for these projects.

- JETSTREAM IS ESSENTIAL!
- NSF has made a very wise investment in Jetstream. It provides an ideal venue for researchers to get acquainted with cloud computing, without having to jump through bureaucratic hoops in order to work on exclusively commercial cloud platforms.
- NSF's investment in Jetstream is extremely important. It has enabled availability of high performance computational resources to researchers like me who otherwise would not be able to afford access on HPC/cloud systems. In addition, by allowing for merit based projects, the allocations are awarded to well designed and thoughtful projects, further helping to refine a study which is not merely monetary based.
- Great job guys
- The use of Jetstream has been invaluable for my research. Money well spent by the NSF.
- Useful for managing and maintaining our software
- Platforms like Jetstream are critical to the advancement of data intensive scientific research.
- Overall, I appreciate the NSF investing in supercomputing. As a lab, we are still learning how to make the best use of Jetstream.
- It's great to have a resource my students can log into and practice on!
- Jetstream is a valuable cloud resource for the XSEDE community.
- Very valuable for my research



•