



**PERVASIVE
TECHNOLOGY INSTITUTE**



RESEARCH TECHNOLOGIES
UNIVERSITY INFORMATION TECHNOLOGY SERVICES



**PERVASIVE
TECHNOLOGY INSTITUTE**



RESEARCH TECHNOLOGIES
UNIVERSITY INFORMATION TECHNOLOGY SERVICES

Panel: Ups and downs of cloud computing in open science

David Y. Hancock – Indiana University

Director for Advanced Cyberinfrastructure

Jetstream & Jetstream2 Primary Investigator

2020 CI/CS Workshop – August 20, 2020



What is Jetstream(2) and why does it exist?

- NSF's first production cloud facility (and successor)
- Provides *on-demand interactive* computing and analysis or persistent services (e.g. gateways)
- Under 10% NSF investment → support for 24% of institutions, 23% of active PIs, and 32% of users*
- Jetstream has provided **6x more** SUs than **any other** XSEDE resource for Education
- Emphasis on ease-of-use, broad accessibility, *AI for Everyone*
- Enables *configurable* environments; *programmable cyberinfrastructure*



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

M87 black hole: how cloud computing supports astronomy



Event Horizon Telescope



*M87 black hole image
generated by EHT*

- Event Horizon Telescope, a large number of scientists, NASA spacecraft, and a variety of computing resources enabled the first image of a black hole.
- For the M87 black hole image, two critical steps were done in the cloud and piloted on Jetstream
 - Correcting for anomalies, so that further image processing could occur, and
 - Large survey study of how image reconstruction algorithms affect the final images.
- Chi-kwan Chan, leader of the EHT Computations and Software Working Group, "The production run was actually carried out on Google Cloud, but much of the early development was on Jetstream. Without Jetstream, it is unclear that we would have a cloud-based pipeline at all."

<https://go.iu.edu/35yN>

Project TERRIER: modernizing political science research



Journalists, through their words, provide an invaluable service, sharing information about global events to which many of us would not otherwise have access.

They send missives directly from event sites, recording during protests, summits, and speeches. For political scientists, these articles offer a rich mine of data about these events.

Project Temporally Extended, Regular, Reproducible International Event Records ([TERRIER](#)) extracts event data from roughly 300M news articles and puts it into a form usable by researchers.

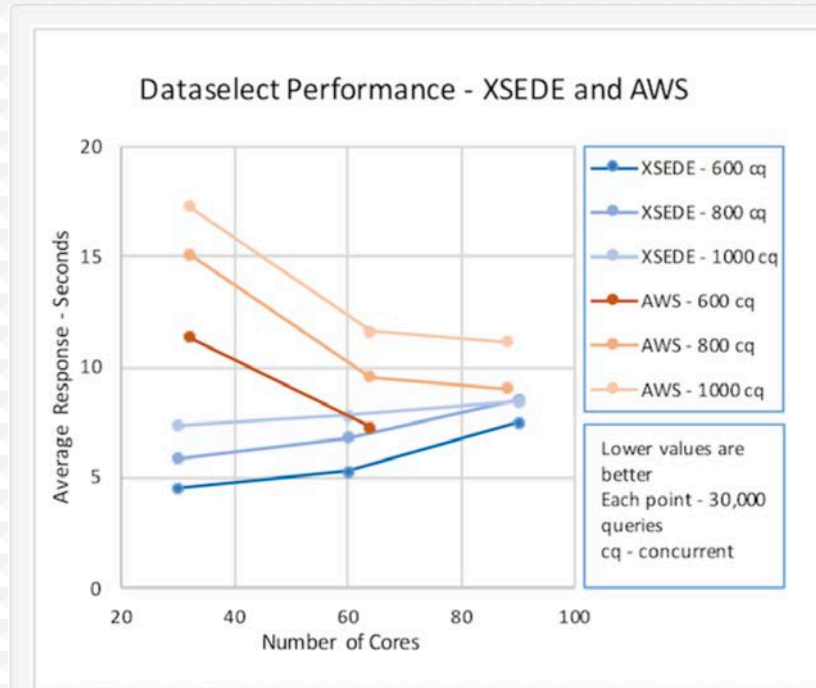
Jetstream provides the storage and structure to launch the pipeline and process news articles for TERRIER.



<https://go.iu.edu/35yM>

GeoSciCloud: Exploring the potential for hosting a geoscience data center in the cloud

- Project with EarthCube support to evaluate a data center in the cloud (IRIS and UNAVCO)
- Compared performance, cost, operations using a test data set and core data access services.
- IRIS is providing ~30 years of seismic data with a primary archive of ~500TB.
- Over 100K users annually.
- Details courtesy Chad Trabant of IRIS: <https://go.iu.edu/35Ao>



Comparison of raw data access response times for both cloud environments with a variety of supporting cores. This result shows that XSEDE provides faster access with lower core counts but roughly matching AWS at higher core counts. The increase of response time with an increase in cores at XSEDE is due to saturation of the connection between the virtual machines and the storage system. Note that the levels of concurrent requests (cq) are not something we can currently offer from the DMC.

Operational comparison, key points

XSEDE/Jetstream

- Somewhat known Openstack
- No-cost, high performance connectivity
- Possible to get highly valuable direct support
- *No large storage, needed Wrangler*
- *No guaranteed uptime, 95% w/NSF*
- *Saturated storage system path*
- *Short-term allocation model*

AWS

- The Amazon way
- Anything significant costs
- Free consulting for startup, EDU?
- Base level of support \$50/day, next-day
- No insight into internals
- Lots of options including global replication of data, cost sharing, etc.
- Limits (e.g. CPU) and complexity
- Potential for vendor lock in

Jetstream for education – in action at AMS2020

- Unidata-led workshop at American Meteorological Society (AMS) 2020 conference
- 127 users actively participating
- Participants used a JupyterHub running on Jetstream (40 node Kubernetes cluster of 6 core m1.medium VMs) for a 90-minute Unidata PyAOS (Python for the Atmospheric and Oceanic Sciences) workshop
- The students were successfully able to run their interactive Python code notebooks as the instructors presented their material



Return on Investment

- XSEDE ROI analysis led by Craig Stewart; co-authored subsequent papers
 - Added Jetstream and Big Red II analysis
 - Using ROI as a proxy using cost avoidance
- Big Red II (Cray XE6/XK7 over 6.5 years): 3.80 – 5.75*
- Big Red 3 (Cray XC30/XC40 over 3 years): 5.16 – 7.81*

IEEE UCC/BDCAT 2018 Paper: <https://doi.ieeecomputersociety.org/10.1109/UCC.2018.00031>

*using 1-year or 3-year up front reserved instance cost



Timeline

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



Flickr user Oiluj Samall Zeid - Lejos de Yulín



**PERVASIVE
TECHNOLOGY INSTITUTE**



RESEARCH TECHNOLOGIES
UNIVERSITY INFORMATION TECHNOLOGY SERVICES

Acknowledgements

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

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

Special thanks to contributors & Jetstream2 partners

- Jeremy Fischer, J. Michael Lowe, Therese Miller, Maria Morris, Winona Snapp-Childs, and George Turner, Chad Trabant



**PERVASIVE
TECHNOLOGY INSTITUTE**



RESEARCH TECHNOLOGIES
UNIVERSITY INFORMATION TECHNOLOGY SERVICES

Jetstream2 partners



JOHNS HOPKINS
UNIVERSITY



UCAR



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



**PERVASIVE
TECHNOLOGY INSTITUTE**



RESEARCH TECHNOLOGIES
UNIVERSITY INFORMATION TECHNOLOGY SERVICES

Jetstream2 Capabilities

Enhancing IaaS model of Jetstream:

- Improved orchestration support
- Elastic virtual clusters
- Federated JupyterHubs

Commitment to **>99%** uptime

- Critical for science gateway hosting
- Hybrid-cloud support

Revamped User Interface

- Unified instance management
- Multi-instance launch



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

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