

Jetstream Overview

A national research and education cloud

9th workshop on Scientific Cloud Computing (ScienceCloud) – June 11, 2018 – Tempe, AZ

John Michael Lowe – iomlowe@iu.edu

Senior Cloud Engineer,
UITS Research Technologies

Lowe, J. (2018). Jetstream Overview: A national research and education cloud. Tempe, AZ. Retrieved from <https://jetstream-cloud.org/research/publications.php>



NSF Funding Areas in HPC

Traditionally concentrated on enabling petascale capability

- Blue Waters – 13.3 petaflops, 2012 (under re-compete)
- Stampede – 9.6 petaflops, 2013 (extended to Stampede2, going into service now)
- Comet – ~2.0 petaflops, 2014

Has funded research into building clouds and computer science

- CloudLab
- Chameleon (renewed for 2nd phase)

Now funding clouds to do research

- Bridges (Hybrid system)
- Jetstream



funded by the National Science Foundation
Award #ACI-1445604



Jetstream - Expanding NSF XD's reach and impact

Lots of stats below –

tl;dr summary: no one has enough computing resources. Ever. But they need easy access and use.

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

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

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

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

* XSEDE Cloud Survey Report - <http://dx.doi.org/10.2142/45766>



funded by the National Science Foundation
Award #ACI-1445604



What is Jetstream and why does it exist?

- NSF's first production cloud facility
- Part of the NSF eXtreme Digital (XD) program
- Focus on ease-of-use, broad accessibility
- User-selectable library of preconfigured virtual machines
- Provides on-demand *interactive* computing and analysis or persistent gateways (SEAGrid, Galaxy, GenApp NAMDRunner, CIPRES and others)
- Enables *configurable* environments and **programmable cyberinfrastructure**
- Reproducibility: Share VMs and then store, publish via IU Scholarworks (DOI)



funded by the National Science Foundation
Award #ACI-1445604



Who uses Jetstream?

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



funded by the National Science Foundation
Award #ACI-1445604



What Jetstream isn't...

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

HPC vs Cloud

Adapting to a different environment:

- No reservations, no queueing – more interactive usage
- Being your own admin – hey, we have root!**
- You really can have almost any (linux) software you want**
- Constantly getting new features (<https://www.openstack.org/software/project-navigator/>)

** Here there be dragons...



funded by the National Science Foundation
Award #ACI-1445604




Thinking about VMs...



Cows, not pets: pets take great amount of care, feeding, and you name them; cows you intend to have high turnover and you give them numbers.

Jetstream System Overview



Production cloud hardware (per site)

Hardware	Number	Specifications	Function (IU)
Dell PowerEdge M630 blades	320	2x Intel E5-2680v3 "Haswell" 24 cores @ 2.5 GHz 128 GB RAM 2 TB local disk	Compute hosts OpenStack services
Dell PowerEdge R630 1U server	7	2x Intel E5-2680v3 "Haswell" 24 cores @ 2.5 GHz 128 GB RAM 2 TB local disk	Cluster management High Availability Databases RabbitMQ
Dell PowerEdge R730xd 2U servers	20	2x Intel E5-2680v3 "Haswell" 24 cores @ 2.5 GHz 64 GB RAM 48 TB storage for Ceph pool	~1 PB Ceph storage
Dell S6000-ON network switches	9	32+2 40 Gb/s ports	Top of Rack Spine

Hardware and Instance "Flavors"

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

** s1.* storage-rich instances are not eligible to be saved into a customized image

- Short-term *ephemeral* storage comes as part of launched instance
- Long-term storage is XSEDE-allocated
- Implemented as OpenStack Volumes and object storage
- Default storage is modest, but more is available via allocation



funded by the National Science Foundation
Award #ACI-1445604



Using Jetstream VMs

Manipulating Jetstream VMs:

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

Primary methods of logging into Jetstream VMs to work

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



funded by the National Science Foundation
Award #ACI-1445604



Discipline or area of interest	#of Jetstream allocations	SUs allocated on Jetstream	% of SUs allocated on Jetstream	% of all SUs allocated on other XSEDE-supported systems
Astronomy	2	1,108,096	3.04%	8.61%
Atmospheric Sciences	4	2,752,400	7.55%	3.73%
Biological Sciences	57	5,199,000	14.27%	4.95%
Campus/Domain Champions	123	6,105,500	16.76%	0.09%
Computational Science	11	1,150,000	3.16%	0.92%
Computer Science	15	4,944,302	13.57%	1.8%
Education Allocations	24	2,847,600	7.82%	0.01%
Engineering	1	100,000	0.27%	3.81%
Geosciences	10	1,978,400	5.43%	2.87%
Humanities/Social Sciences	10	560,000	1.54%	0.45%
Molecular Biosciences	8	4,647,520	12.75%	17.65%
Network Science	3	200,000	0.55%	0.06%
Ocean Science	3	230,000	0.63%	1.30%
Physics	4	2,252,400	6.18%	16.43%
Training & Development	11	2,362,000	6.48%	0.16%

Jetstream usage highlights – 1 May 2018

- 313 active XSEDE projects covering 71 fields of science and **2411 active users** representing **205 institutions**
- **80%** of Jetstream users have **not used any other XSEDE system**
- >117M CPU hours allocated to XSEDE projects since June 2016
- 14 active science gateways
- 41 education/teaching allocations serving over 850 students
- 1151 (avg concurrent) active VMs in current qtr, 955 in PY2*
- **Highest** user satisfaction in most recent XSEDE survey
- **100%** system availability, **99.2%** cap availability in PY2*
- **98.8%** “job” completion in PY2*



funded by the National Science Foundation
Award #ACI-1445604



*M&O PY2 to date 11/12 months

Power

- 2MW supplied by generators tested weekly with minimum 7 days fuel storage
- ½ ton flywheels to condition and provide ~20 seconds of power
- Battery UPS with ~20 min run time


Physical Security

- 24/7/365 operations staff
- Reinforced concrete on six sides with earth berm rated to withstand F5 tornado
- Biometric access controls

Jetstream Fun: Happy cluster / Angry Cluster



Infrared image of Jetstream



ChilledDoor™
Rack Cooling System
by
motivair™

Jetstream
<http://jetstream-cloud.org/>



funded by the National Science Foundation
Award #ACI-1445604



Networking

- Every blade has fully redundant paths to top of rack switch via a pair of stacked chassis switches
 - Every top of rack switch has redundant paths to a pair of spine switches
 - Spine switches have redundant paths to core
- ** Top of rack is a single point of failure

Cloud Services

- All cloud services (vm, volume, image, network create/read/update/delete: cloud controller) are run in triplicate with rack level failure domains
- Two load balancers monitor health, balance requests across 3 cloud controllers, will take over for each other if one load balancer is down
- Backed by 3 way replicated galera database cluster and 3 way replicated rabbitmq message bus

Data Storage

Ceph self healing replicated object storage

- Host level failure domain
- Root and image block storage is 2 x replication, volumes 3x
- Cephfs (experimental) and S3/Swift object storage is erasure coded

Where can I get help?

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

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

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

Email: help@xsede.org

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

Introduction to Jetstream Virtual Workshop: <https://cyw.cac.cornell.edu/ietstream/>

Jetstream Allocations Virtual Workshop: <https://cyw.cac.cornell.edu/JetstreamReq/>



funded by the National Science Foundation
Award #ACI-1445604



Jetstream Partners



funded by the National Science Foundation
Award #ACI-1445604



Questions?

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

Project email: help@ietstream-cloud.org Direct email: iomlowe@iu.edu

License Terms

- Lowe, John Michael. June 11, 2018. Jetstream Overview – Science Cloud ‘18. Also available at: <http://ietstream-cloud.org/research/publications.php>
- Jetstream is supported by NSF award 1445604 (David Y. Hancock, IU, PI)
- XSEDE is supported by NSF award 1053575 (John Towns, UIUC, PI)
- This research was supported in part by the Indiana University Pervasive Technology Institute, which was established with the assistance of a major award from the Lilly Endowment, Inc. Opinions presented here are those of the author(s) and do not necessarily represent the views of the NSF, IUPTI, IU, or the Lilly Endowment, Inc.
- Items indicated with a © are under copyright and used here with permission. Such items may not be reused without permission from the holder of copyright except where license terms noted on a slide permit reuse.
- Except where otherwise noted, contents of this presentation are copyright 2015 by the Trustees of Indiana University.
- This document is released under the Creative Commons Attribution 3.0 Unported license (<http://creativecommons.org/licenses/by/3.0/>). This license includes the following terms: You are free to share – to copy, distribute and transmit the work and to remix – to adapt the work under the following conditions: attribution – you must attribute the work in the manner specified by the author or licensor (but not in any way that suggests that they endorse you or your use of the work). For any reuse or distribution, you must make clear to others the license terms of this work.