



Research Computing on Cloudy Platforms

Jetstream: A national research and educational cloud

J. Michael Lowe (jomlowe@iu.edu)

Jetstream System Engineer
IU High Performance Systems

George Turner (turnerg@iu.edu)

Chief Systems Architect
IU Research Technologies

Operating Innovative Networks Workshop, Indiana University - Bloomington, 12-July-2016



funded by the National Science Foundation
Award #ACI-1445604 Jetstream
Award #ACI-1445606 Bridges

National Science Foundation

- Traditionally concentrated on enabling petascale capability
 - Blue Waters – 13.3 PetaFLOPS, 2012
 - Stampede – 9.6 PetaFLOPS, 2013
 - Comet – ~2.0 PetaFLOPS, 2014
- Has funded research into building clouds
 - Cloumlab
 - Chameleon
- Now funding clouds to do research
 - Bridges
 - Jetstream



funded by the National Science Foundation
Award #ACI-1445604 Jetstream
Award #ACI-1445606 Bridges



Expanding NSF XD's reach and impact

- Around 350,000 researchers, educators, & learners received NSF support in 2015
 - Only <2% completed a computation, data analysis, or visualization task on XD program resources
 - Less than 4% had an XSEDE Portal account
 - 70% of researchers surveyed* claimed to be resource constrained
- Why aren't they using XD systems?
 - Activation energy is pretty high
 - HPC resources are scarce and not well-matched to their needs
 - They just don't need *that much* capability

What is Jetstream?

- User-friendly, widely accessible cloud environment
- User-selectable library of preconfigured virtual machines; no need for system administration skills.
- NSF's first production cloud facility supporting all areas of science and engineering within NSF's scope
- Enable discoveries across disciplines such as biology, atmospheric science, economics, network science, observational astronomy, and social sciences.

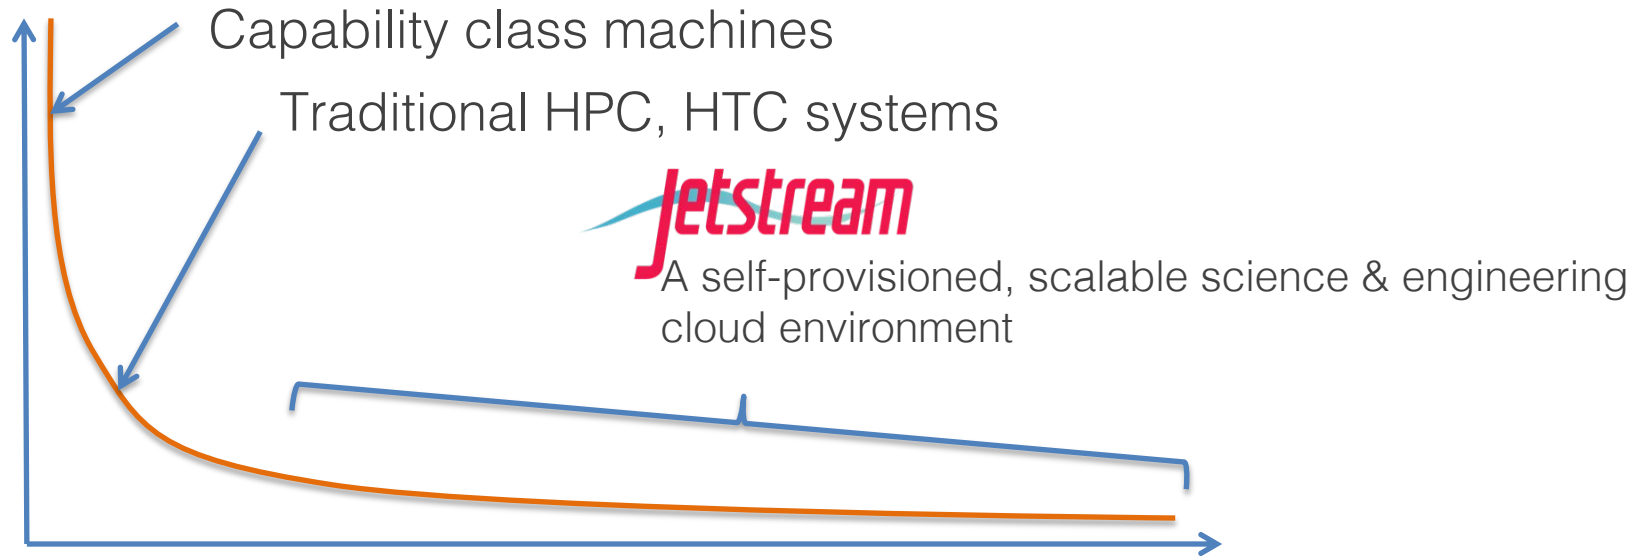


funded by the National Science Foundation
Award #ACI-1445604 Jetstream
Award #ACI-1445606 Bridges



What is Jetstream? (cont)

- Particularly focused on researchers working in the “long tail” of science with born digital data



Who will use Jetstream?


- Researchers & students needing access to interactive computing and data analysis resources on demand. “A few processors now instead of thousands next week.”
- Researchers & software developers creating & maintaining domain specific software packages
- As a backend supporting science gateways



funded by the National Science Foundation
Award #ACI-1445604 Jetstream
Award #ACI-1445606 Bridges



Jetstream hardware



funded by the National Science Foundation
Award #ACI-1445604 Jetstream
Award #ACI-1445606 Bridges



Hardware Specifics

VM Host Configuration

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

Flavor	vCPUs	RAM	Storage	Per Node
m.tiny	1	2	20	46
m.small	2	4	40	23
m.medium	6	16	130	7
m.large	10	30	230	4
m.xlarge	22	60	460	2
m.xxlarge	44	120	920	1

CEPH Storage

- 20x Dell 730xd per cloud
- 2x10Gbs bonded NIC per 730xd
- Running CEPH 0.94.5 Hammer
- Configured as OpenStack Storage
- Storage is XSEDE-allocated
- Implemented on backend as OpenStack Volumes
- Each user gets 10 volumes up to 500GB total storage
- Exploring object storage as well but that’s in the future



funded by the National Science Foundation
Award #ACI-1445604 Jetstream
Award #ACI-1445606 Bridges



Performance numbers:

Bare metal:

- LINPACK 86.5% efficient 0.6971 TFOPS
- STREAM 88.5GBs



Performance numbers (Cont):

HPCC results VM vs Bare Metal

- **HPL** **97%** FLOPS floating point execution rate for solving a system of linear equations
- **Ptrans** **64%** FLOPS floating point execution rate for double precision matrix multiplication
- **MPIFFT** **67%** FLOPS floating point execution rate for double precision complex one-dimensional Discrete Fourier Transform
- **Random** **80%** up/s rate of random updates of memory
- **Stream** **77%** B/s sustained memory bandwidth
- **DGEMM** **98%** B/s rate of transfer for large arrays of data from multiprocessor's memory
- **Bandwidth** **88%** B/s bytes/time it takes to transmit a 2,000,000 byte
- **Latency** **97%** s time required to send an 8-byte message

Performance numbers (Cont):

Units for previous slide

- FLOPS : Floating Point Operations per Second
- up/s : updates per second (random access; not a bandwidth measurement)
- B/s : Bytes per second (contiguous transfers; bandwidth measurement)
- s : microseconds

Details at <http://icl.cs.utk.edu/hpcc/faq/>



funded by the National Science Foundation
Award #ACI-1445604 Jetstream
Award #ACI-1445606 Bridges

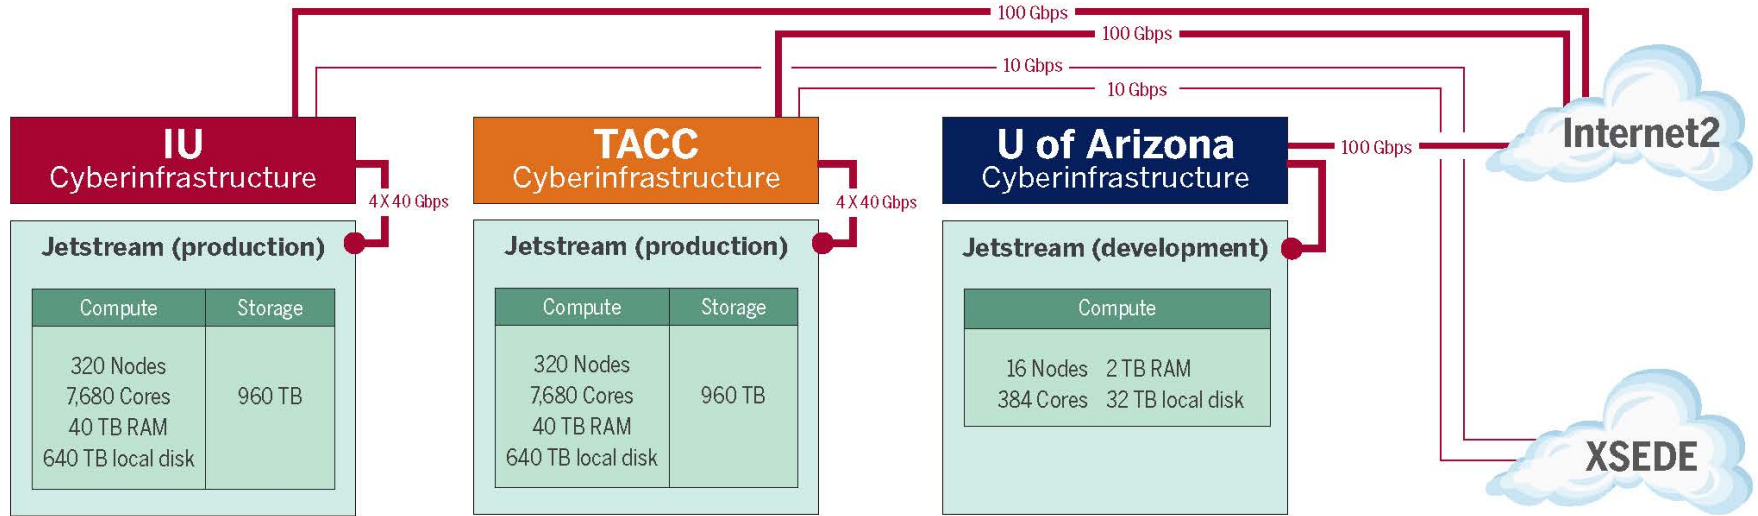


Performance Numbers


Rally

Response Times (sec)								
Action	Min (sec)	Median (sec)	90%ile (sec)	95%ile (sec)	Max (sec)	Avg (sec)	Success	Count
nova.boot_server	8.725	9.43	10.374	10.467	10.56	9.548	100.0%	10
nova.delete_server	2.608	2.679	2.788	2.83	2.872	2.696	100.0%	10
total	11.387	12.152	13.145	13.2	13.256	12.244	100.0%	10

Jetstream System Overview




Compute Chassis & Network Topology

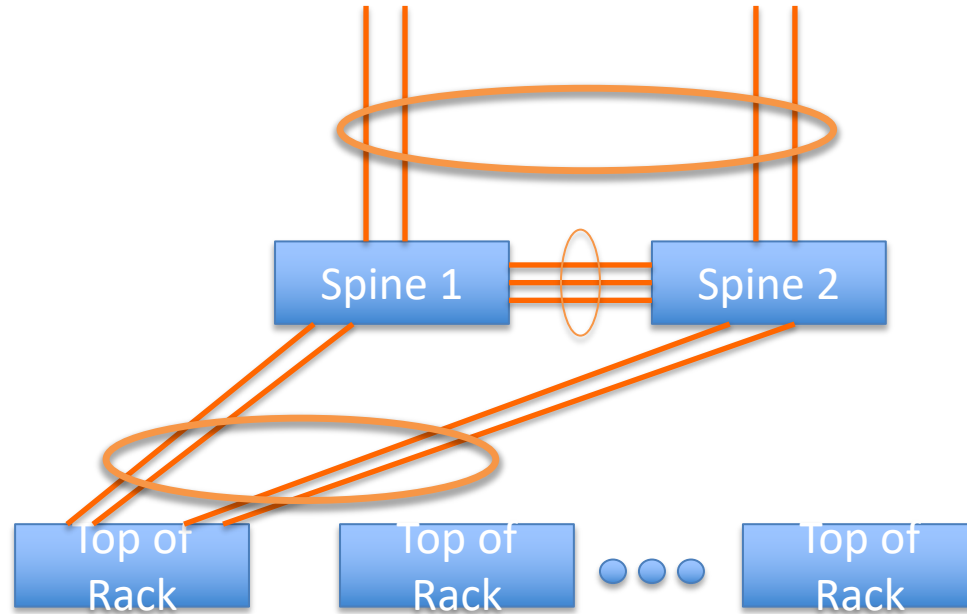


Network Topology

Chassis to Top of Rack



Network Topology ...cont.



Service nodes



Software

- Software layers
 - **Atmosphere** web interface
 - library of images, generic, domain specific
 - simplify VM administration
 - **OpenStack**: software tools for building and managing cloud computing platforms for public and private clouds.
 - **KVM** hypervisor: what the VMs run on
 - **Ceph**: storage platform that stores data on a single distributed computer cluster, and provides interfaces for **object-**, **block-** and *file-level* storage.

Things we did like everybody else

- Ceph backed Cinder and Glance
- Nova compute w/ KVM
- Deploy the most recent OpenStack release (Currently Mitaka)

Things some people do

- Linuxbridge and vxlan networking
- Intel x710 nics w/ vxlan offload
- Keystone Domains
- Use as many HA features as we could afford
- Use SaltStack



Things nobody else does

- Atmosphere web UI
- Water cooled doors
- 100Gb/s Provider network



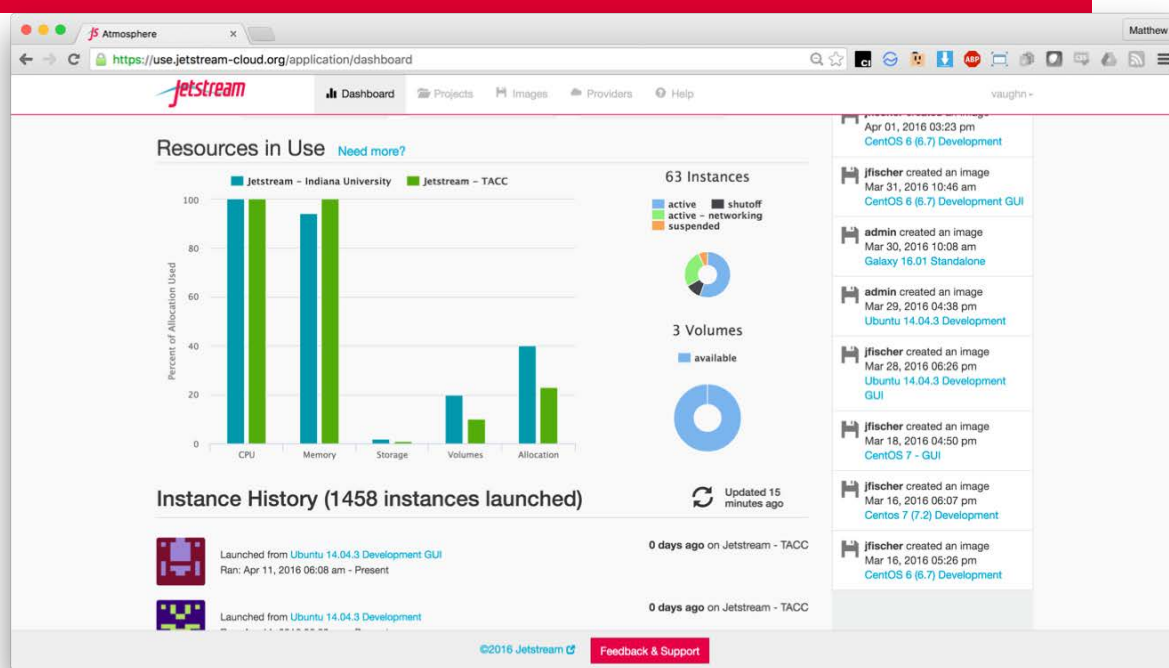
funded by the National Science Foundation
Award #ACI-1445604 Jetstream
Award #ACI-1445606 Bridges




Happy Cluster – Mad Cluster



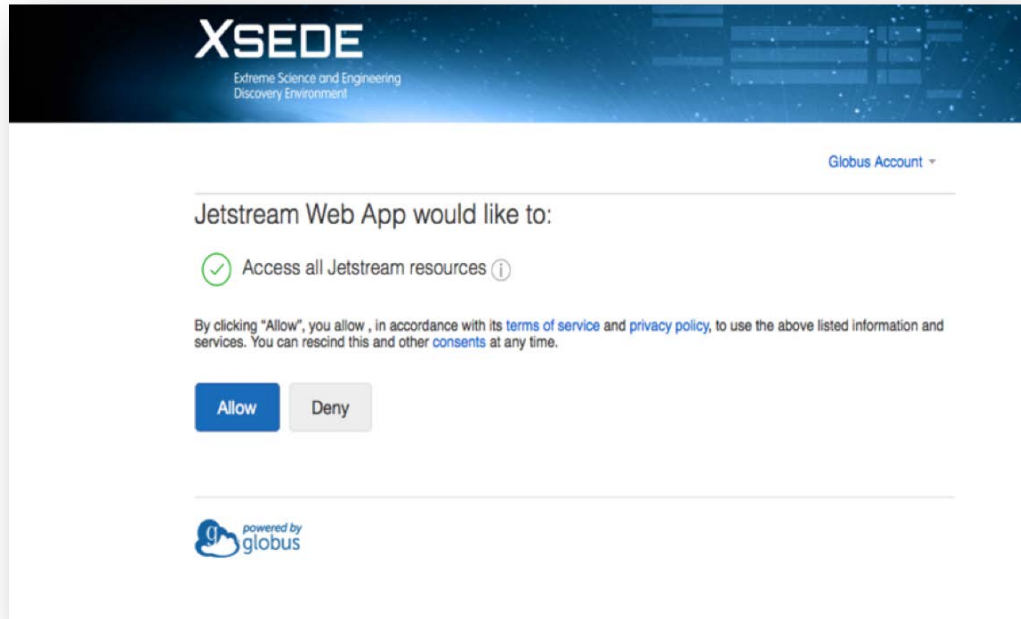
Jetstream's Atmosphere UI



Jetstream's Atmosphere UI



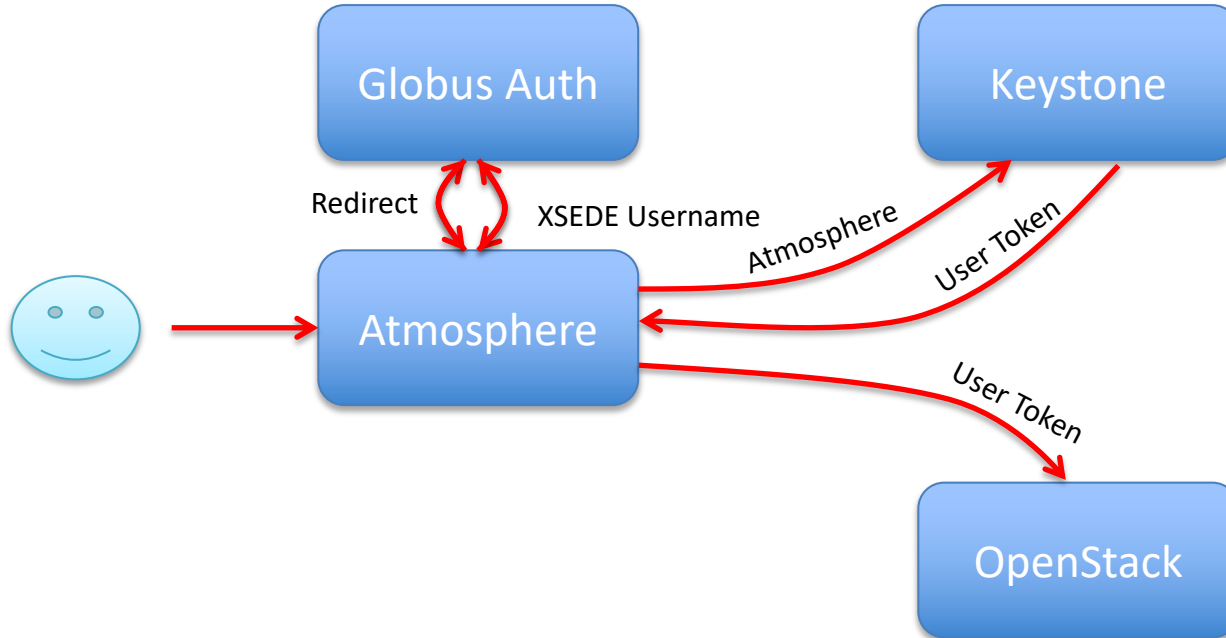
Integration with XSEDE via Globus Auth



Atmosphere Web App uses and Globus Auth implements industry-standard Oauth2

- Leaves us *flexibility* on identity and access
- Globus Auth implements (in beta) password grant Oauth flow, which means *Jetstream access can be entirely scripted*

Authentication



Jetstream Partner Organizations

Initial construction partners



THE UNIVERSITY
OF ARIZONA.



THE UNIVERSITY OF
CHICAGO



JOHNS HOPKINS
UNIVERSITY

Management & Operations partners



Application / community lead partners



UNIVERSITY
of ARKANSAS
AT PINE BLUFF
—1873—



UNC
THE ODUM INSTITUTE



funded by the National Science Foundation
Award #ACI-1445604 Jetstream
Award #ACI-1445606 Bridges



Jetstream Links

- Jetstream: <https://use.jetstream-cloud.org/>
- XSEDE User Portal is required to actually login: <https://portal.xsede.org>
- User guide: <http://jetstream-cloud.org/training.php>
- Paper describing Jetstream [Jetstream: A self-provisioned, scalable science and engineering cloud environment](http://jetstream-cloud.org/training.php)
- <http://github.com/jetstream-cloud/Jetstream-Salt-States>

Questions?

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

Project email: help@jetstream-cloud.org Direct email: jomlowe@iu.edu , turnerg@iu.edu

License Terms

- Lowe, Mike. Turner, George. July 12, 2016. Research Computing on Cloudy Platforms. Available at: <http://jetstream-cloud.org/publications.php>
- Jetstream is supported by NSF award 1445604 (Craig Stewart, IU, PI)
- XSEDE is supported by NSF award 1053575 (John Towns, UIUC, PI)
- This research was supported in part by the Indiana University Pervasive Technology Institute, which was established with the assistance of a major award from the Lilly Endowment, Inc. Opinions presented here are those of the author(s) and do not necessarily represent the views of the NSF, IUPTI, IU, or the Lilly Endowment, Inc.
- Items indicated with a © are under copyright and used here with permission. Such items may not be reused without permission from the holder of copyright except where license terms noted on a slide permit reuse.
- Except where otherwise noted, contents of this presentation are copyright 2015 by the Trustees of Indiana University.
- This document is released under the Creative Commons Attribution 3.0 Unported license (<http://creativecommons.org/licenses/by/3.0/>). This license includes the following terms: You are free to share – to copy, distribute and transmit the work and to remix – to adapt the work under the following conditions: attribution – you must attribute the work in the manner specified by the author or licensor (but not in any way that suggests that they endorse you or your use of the work). For any reuse or distribution, you must make clear to others the license terms of this work.



funded by the National Science Foundation
Award #ACI-1445604 Jetstream
Award #ACI-1445606 Bridges

