

Jetstream: Building and Operating a First of Kind System (What could go wrong?)

RMACC HPC Symposium – May 21, 2019 – Boulder, CO

John Michael Lowe – jomlowe@iu.edu

Senior Cloud Engineer,
UITS Research Technologies

Lowe, J. (2019). Jetstream: Building and Operating a First of Kind System. Boulder, CO. 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 in 2017 – 18 petaflops)
- Comet – ~2.0 petaflops, 2014

Has funded research into building clouds and computer science

- CloudLab (renewed for 2nd phase)
- 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...but most aren't using XSEDE in any capacity at all.

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://hdl.handle.net/2142/45766>

Identifying the potential users

“But I really don’t have research needs...I don’t need the national research cyberinfrastructure.”

--- multiple researchers at a number of small colleges and universities

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
- Provides on-demand *interactive* computing and analysis or persistent gateways
- Enables *configurable* environments and **programmable cyberinfrastructure**

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



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...we'll get to this later)
- 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...



Jetstream and way of the cloud...

- **Cloudy Technologies:** clouds are more than just virtual machines (VM)
 - **Old way:** robust (expensive) infrastructure, weak (cheap) software
 - You expect the hardware to not fail
 - State is maintained in volatile data structures
 - **Cloudy way:** commodity infrastructure, robust software
 - Expect & plan for infrastructure to fail
 - Put intelligence into the software to handle infrastructure failure
- **And my favorite...**

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.

**some caveats for gateways...

Image source unknown

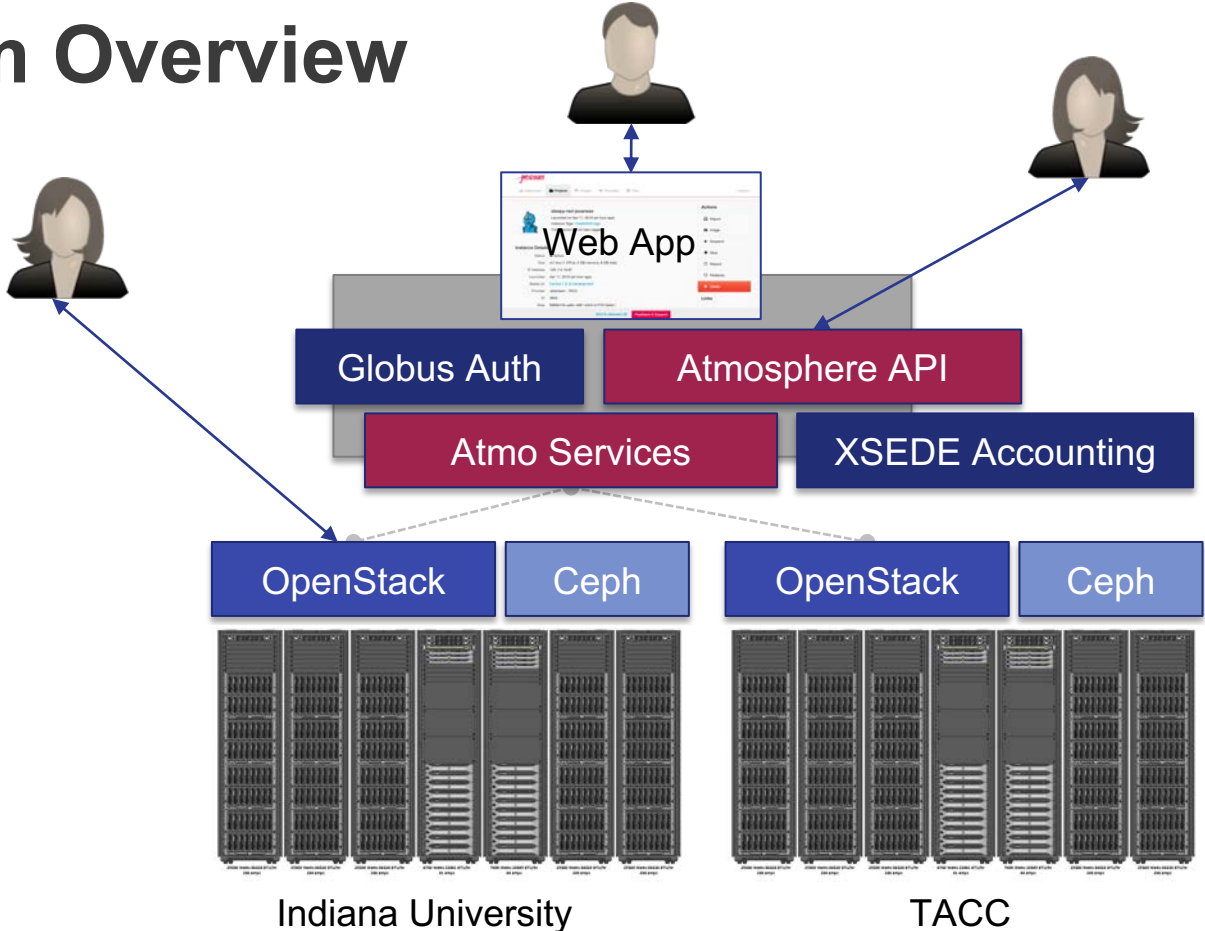
Jetstream System Overview



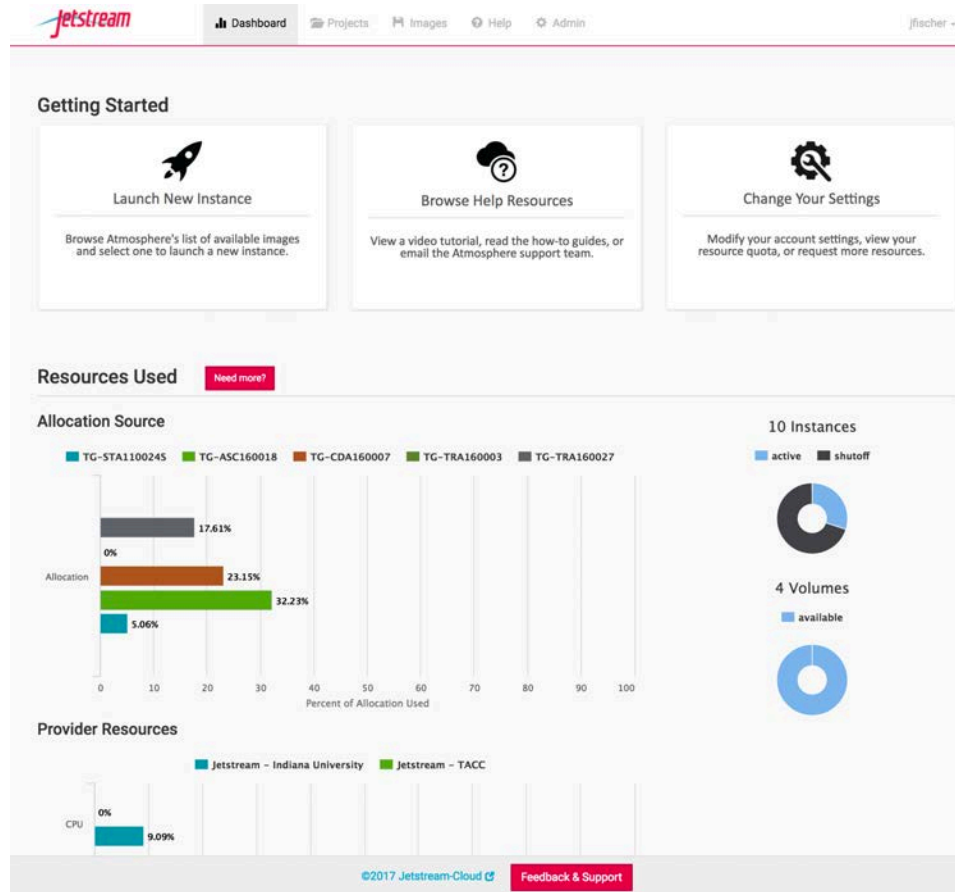
Production cloud hardware (per site)

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

Platform Overview



The Jetstream Atmosphere web interface



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

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

Discipline or area of interest	#of Jetstream allocations	SUs allocated on Jetstream	SU Increase/Decrease on Jetstream over previous year	% of SUs allocated on Jetstream	% of all SUs allocated on other XSEDE-supported systems
Behavioral Sciences	6	3,465,516	100% Increase	4.24%	0.61%
Biological Sciences	89	15,041,928	72.18% Increase	18.40%	3.59%
Biophysics	86	3,627,026	44.15% Decrease	4.44%	13.56%
Computer Science	72	6,883,269	33.28% Increase	8.42%	2.98%
Earth Sciences	37	5,476,250	37.06% Increase	6.70%	4.60%
Education and Training	128	16,599,512	2.62% Increase	20.31%	4.66%
Engineering	13	520,690	71.21% Increase	0.64%	1.75%
Materials Science	6	1,035,508	100% Increase	1.27%	13.89%
Mathematics	13	688,505	150.37% Increase	0.84%	0.90%
Molecular Science/Biochemistry	21	4,254,643	10.15% Decrease	5.20%	5.83%
Neuroscience	19	4,708,180	327.87% Increase	5.76%	1.98%
Physics	10	2,440,581	15.58% Decrease	2.99%	8.65%
Social Sciences and Humanities	28	2,409,633	192.27% Increase	2.95%	0.81%

Not just the usual suspects...

Physics, chemistry, and other “usual” HPC suspects are represented, but Jetstream also is home to projects on:

- Financial analysis / Economics
- Political science
- Humanities / Text analysis
- Network analysis
- Computer Science / Machine learning
- Satellite data analysis

Jetstream for Education

Jetstream has been used in multiple graduate and undergraduate courses

- Management, Access, and Use of Big and Complex Data
- Multiple informatics and general bioinformatics courses
- Business Intelligence (big data and analysis)
- Research Topics in Music
- Multiple genetics and sequencing courses
- Multiple information security and assurance courses
- ...and others...

Multiple Research Data Alliance Workshops, multiple workshops/classes on Galaxy, data analysis in finance using R, security and intrusion detection, and principles in cloud computing and more!



Another Use Case: Galaxy riding Jetstream

Galaxy is a platform for biomedical research, focused on accessibility, transparency and reproducibility

- The main project instance (usegalaxy.org) has more than 100,000 registered users executing 300,000+ jobs each month
- Many users need more capacity than the public quota, or other customizations (e.g., new tools)

Use Jetstream as a *bursting* platform

- From Galaxy Main, offload jobs onto a remote Slurm cluster running on Jetstream instances
- Run Galaxy Interactive Environments (i.e., Dockerized IPython/RStudio containers) in an isolated environment on a Swarm cluster running on Jetstream

Use Jetstream as a *self-service* platform

- Pre-built Galaxy image configured with hundreds of tools and access to TBs of genomic reference data, available via the self-launch model within minutes
- Allows users to acquire (free) resources, and gives them complete control

Jetstream Gateway Highlights

- IRIS
 - Serving large scale earthquake and geographical data for analysis
- Unidata
 - Providing distribution and analysis of meteorological data
- OpenMRS
 - Providing medical records systems for the resource-constrained
- SEAGrid
 - Computational chemistry, molecular and fluid dynamics, and structural mechanics gateway
- NAMDRunner
 - Based on the GenApp gateway – over 1 million computing hours used to date for MD
- ChemCompute Gateway
 - Providing a computational chemistry gateway for educational use
- Coming gateways: The Neuroscience Gateway, UltraScan III, and others

Jetstream usage highlights – 1 April 2019

- 414 active XSEDE projects covering 73 fields of science and **2300 active users** representing **207 institutions**
- **80%** of Jetstream users have **not used any other XSEDE system**
- >190M CPU hours allocated to XSEDE projects since June 2016
- 24 active science gateways
- 54 education/teaching allocations serving over 904 students
- 1458 (avg concurrent) active VMs in current qtr, 1590 peak active VM count
- **Highest** user satisfaction in most recent XSEDE survey

Requesting access to Jetstream

- Trial allocations available TODAY
<http://wiki.jetstream-cloud.org/Jetstream+Trial+Access+Allocation>
- You can request startup allocations anytime. (Startups are simple!)
<http://wiki.jetstream-cloud.org/Jetstream+Allocations>
- You can request allocations for educational use anytime.
- Next submission period for large allocations is 15 Sept - 15 Oct 2018.
- Research allocation: Main project description (up to 10 pages) and Scaling doc (up to 5 pages) – We can help!

Where can I get help?

Wiki / Documentation: <http://wiki.jetstream-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://cvw.cac.cornell.edu/jetstream/>

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

Jetstream Fun: Happy cluster / Angry Cluster

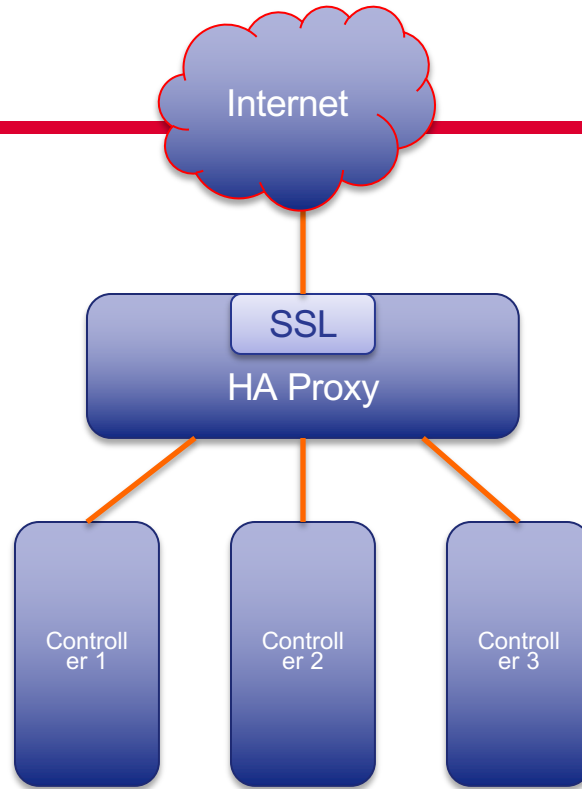


Infrared image of Jetstream



ChilledDoor™
Rack Cooling System
by
motivair™





Load
Balancer 1

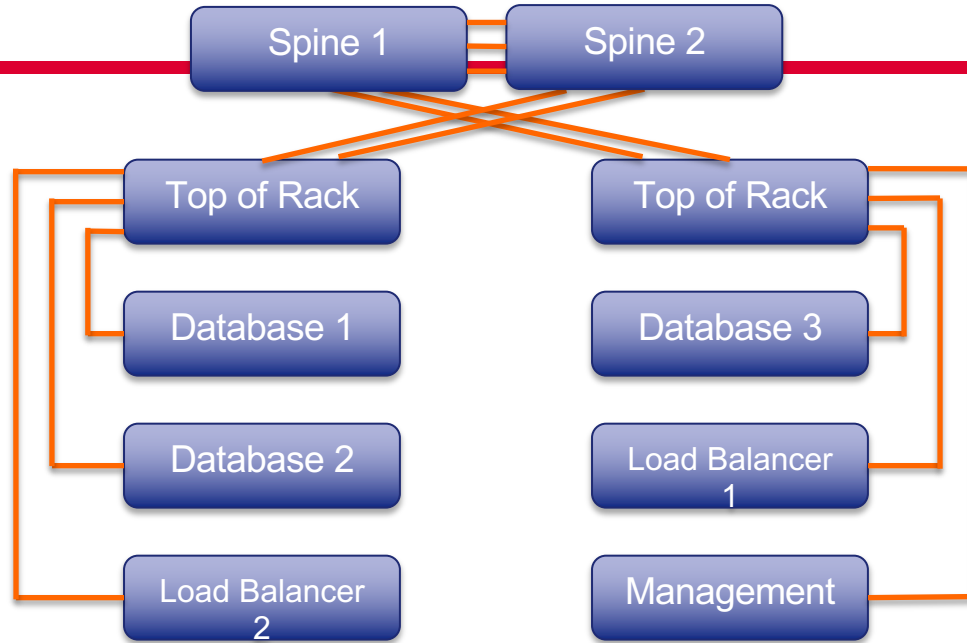


Load
Balancer 2

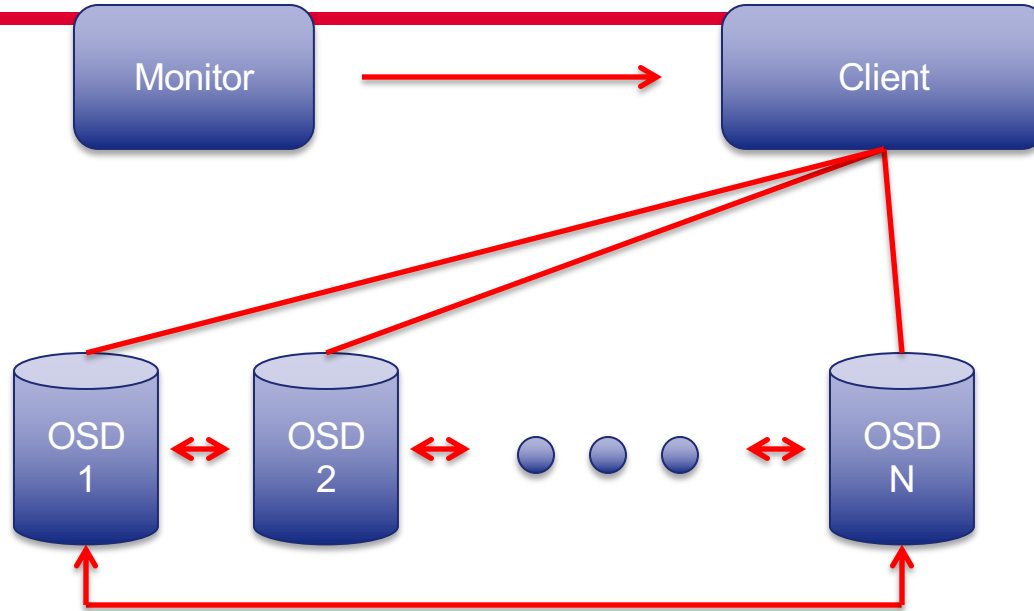
Keep Alive


DNS Round Robin
IP1 – IP2

High Availability layout for the databases

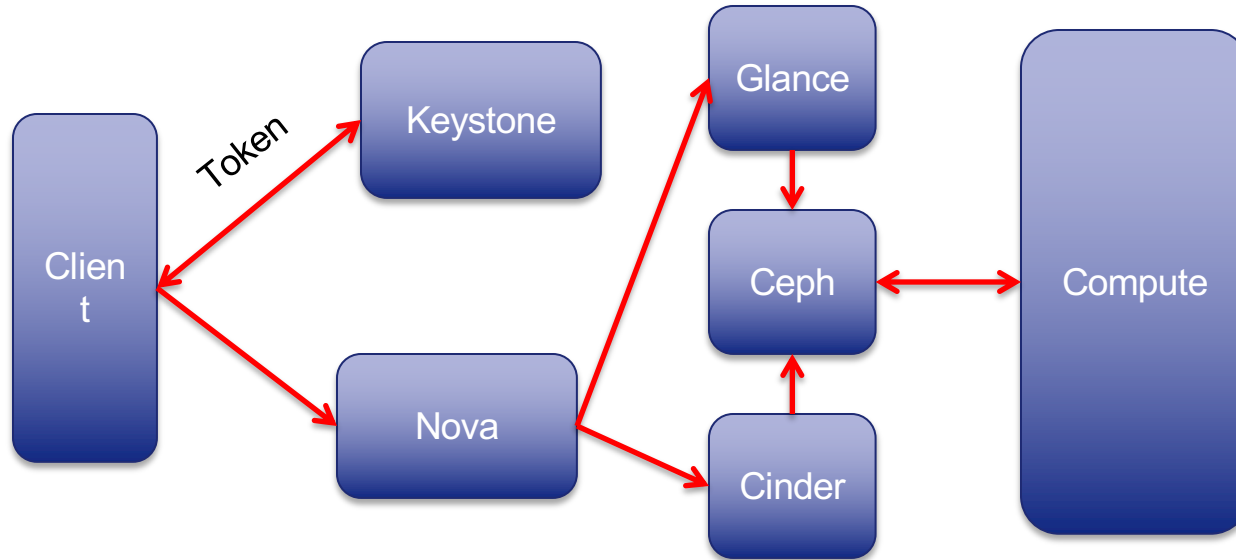


Glance - Cinder - Ceph







OpenStack Overview



VXLAN Packet

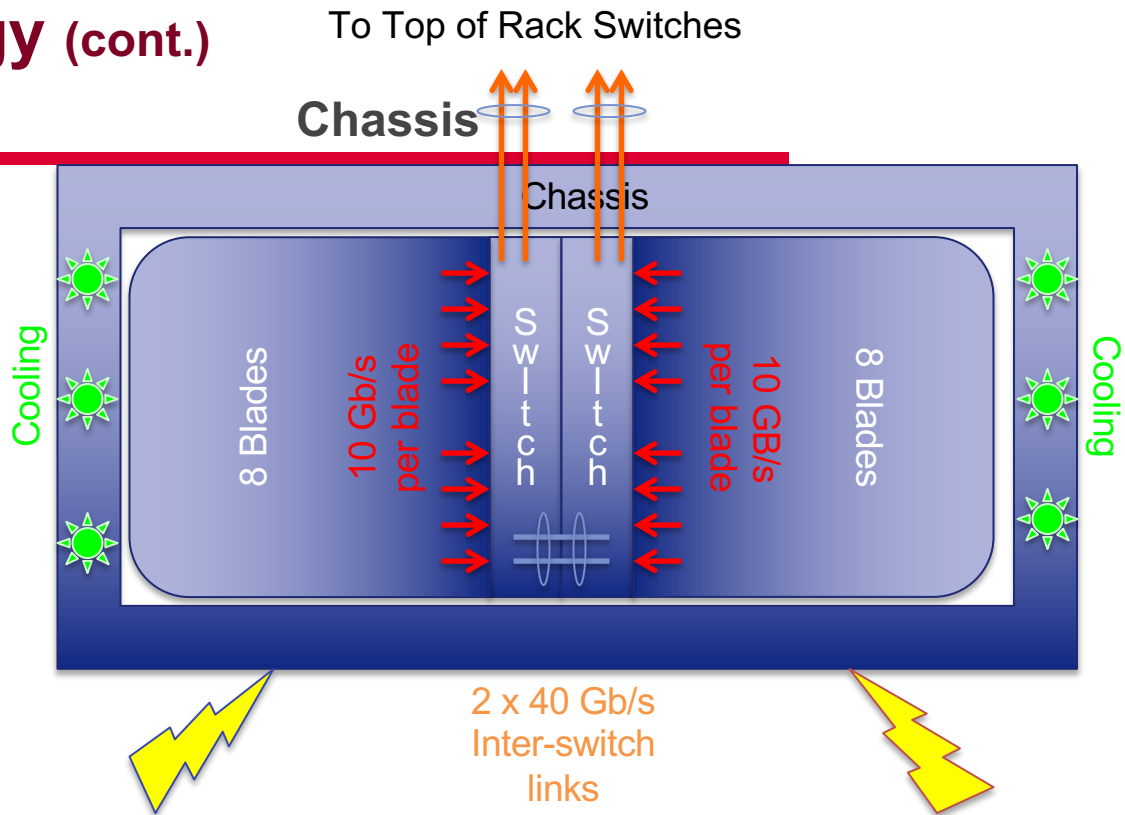



Neutron Networking




Network Topology (cont.)

Sixteen blades per chassis
Two switches per chassis



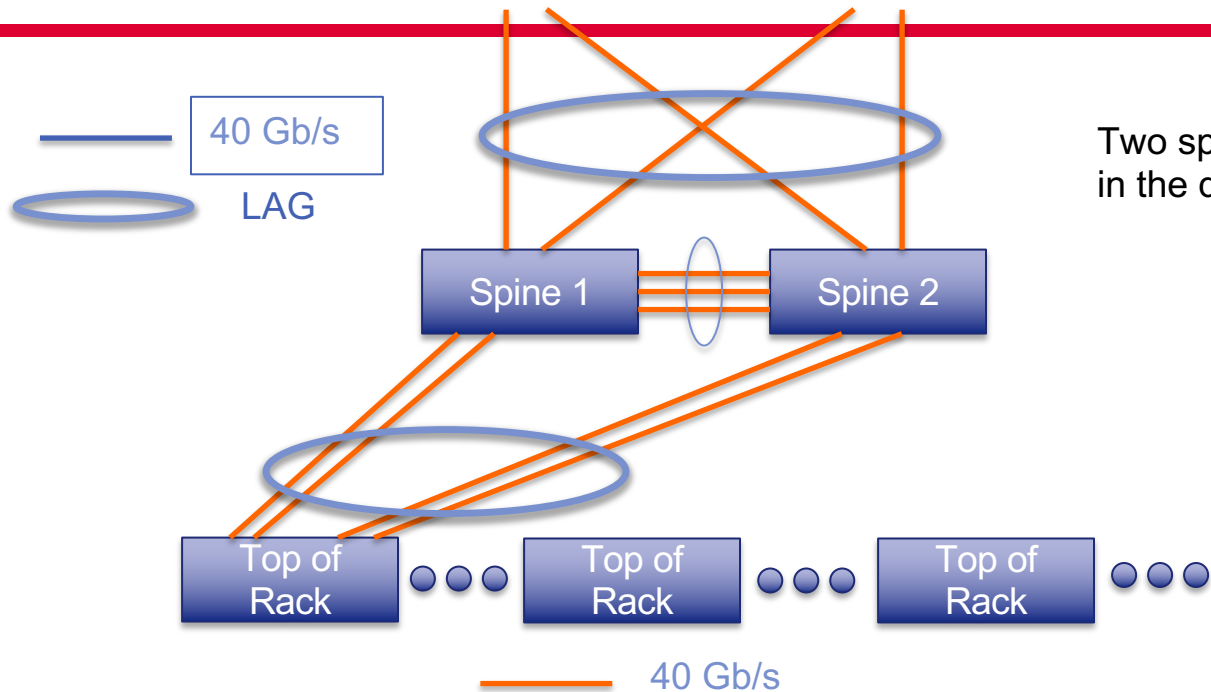
Network Topology (cont.)

Chassis to Top of Rack



Network Topology (cont.)

100 Gb/s uplink to Internet2



Two spines tie into two network blades in the datacenter switch

Seven racks tie into the two spine switches

Jetstream Partners



INDIANA UNIVERSITY

PERVASIVE TECHNOLOGY INSTITUTE



JOHNS HOPKINS
UNIVERSITY




THE UNIVERSITY
OF ARIZONA


TACC



THE UNIVERSITY OF
CHICAGO



JSU | JACKSON
STATE
UNIVERSITY



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



funded by the National Science Foundation
Award #ACI-1445604



Questions?

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

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

License Terms

- Lowe, Joh Michael. May 21, 2019. Jetstream Building and Operating a First of Kind System – RMACC HPC Symposium. Also available at: <http://Jetstream-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.

