

Jetstream

Accessing Jetstream via the OpenStack Command Line Interface

George Turner, Chief Systems Architect
Pervasive Technologies Institute, UITS/RT, Indiana University

Open Cloud Institute
UTSA, San Antonio, TX

12-Oct-2017



funded by the National Science Foundation
Award #ACI-1445604

Terms & Definitions



funded by the National Science Foundation
Award #ACI-1445604



Cloud Computing Terms ...simplified

- **Image:** a file on a disk. It will be booted to create an...
- **Instance:** a running virtual server; i.e. something you can log into.
- **Running:** the *instance* is up & running
- **Suspended:** the *instance* is memory resident but not running
- **Stopped:** the *instance* is shutdown akin to powering down
- **Shelved:** the *instance* is shutdown, backedup, and stored

Cloud Computing Terms ...simplified

- **Flavor:** the size of a running instance; i.e. #core, RAM, disk
- **Hypervisor:** the thing the instance runs on; something akin to a software defined hardware compute server.
- **Snapshot:** the process of taking an instance and turning it to an image.
- **State:** something worth remembering; i.e. the state of the system

Cloud Computing Terms ...simplified (Cont.)

- **Object store:** a blob of bits; it has a starting address & a size. There may be metadata associated with the object. The data is consumed in a streaming manner.
- **Block store:** a software defined entity akin to an unformatted hardware disk drive.
- **Filesystem:** hierarchical in nature, directories & files, ability to open, seek, read, write.
- **Persistent storage:** If you pull the plug, it will still exist when power is restored. Safe to store data or state here.
- **Ephemeral storage:** If you pull the plug, it no longer exists. (Don't put your data here!!!)

What is Jetstream



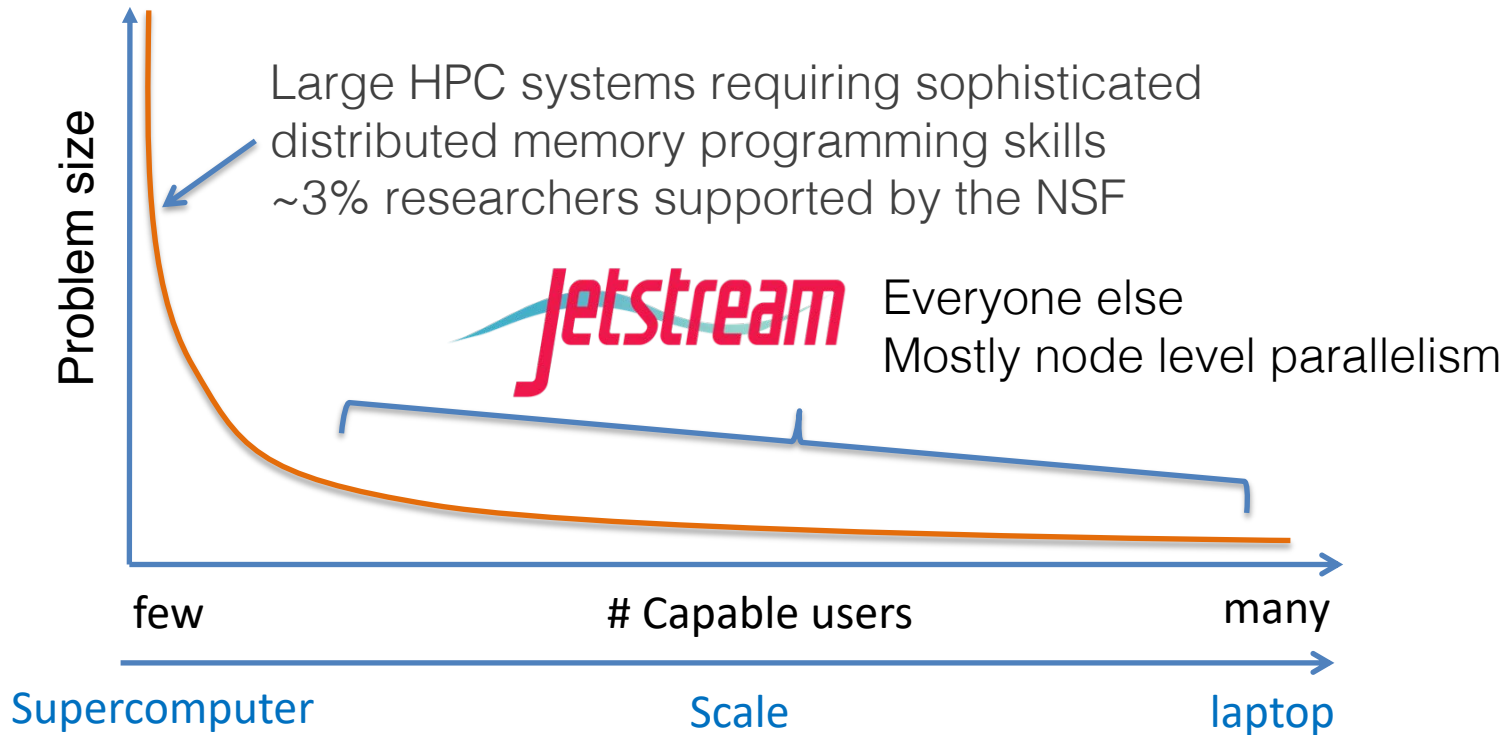
funded by the National Science Foundation
Award #ACI-1445604



What is Jetstream?

- **User-friendly**, widely accessible cloud environment
 - **User-selectable library** of preconfigured virtual machines
 - Interactive computing
 - Software maintained by domain specialist
 - No need for system administration skills
 - The “Atmosphere” side
 - **Programmable cyberinfrastructure**
 - Go beyond batch computing
 - Implement modern cloud computing techniques
 - Common modality for science gateways
 - The “API” side

“Long tail” of the Science



What is Jetstream?

- **Primary goal** is to **expand the user base** of NSF's e**Xtreme Digital (XD)** program resources beyond the current community of users.
- **Lowering the hurdle** to onboard to XSEDE resources
 - Working to **ease the allocation request** process
 - **Easy-Button**; quick access but limited ability. (Beta)
- **Making Science Easy** for domain researchers, engineers, & educators
 - Domain software installed & maintained by the professionals
 - No sys-admin skills necessary

What is Jetstream?

(Cont.)

- **Creating communities**
 - Domain developers **create, install, and maintain** the software
 - **Encourage collaboration** within the domains
 - **Operating system** level software is professionally **patched and maintained**
- **Repeatability**: store & publish images via IU Scholarworks & create a DOI
- **Science Gateways**:
- **Programmatic Cyberinfrastructure**
 - More on this in a minute
 - What this class will be covering today

What is Jetstream?

- **Cloudy Technologies:** clouds are more the just virtual machines (VM)
 - **Old way:** robust (expensive) infrastructure, weak (cheap) software
 - You expect the hardware to not fail
 - State in 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
 - **Cows, not pets:**
 - pets have **state**, you name them, you get attached to them, you put forth great amount of care and effort
 - cows **do not** have **state**, you expect to have high turnover, you do not get attached to them, you give them numbers instead of names

What is Jetstream?

- **Software layers**

- **Atmosphere** web interface (*covered this morning*)
 - 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.
- **Operating systems**: CentOS, Ubuntu, Windows?
- **Applications**; e.g. software developed by the domain specialist, gateways, etc.

API Access to Jetstream

- What was **unexpected**
 - **Demand** for **programmable cyberinfrastructure**
 - Great platform for learning system administration skills
 - Great platform for **teaching & learning cloudy technologies**
- **Command line clients**
- **Horizon dashboard** very popular; but, incomplete
- **Programmatic control**; python is popular
- **Slack channel** for collaboration API users of Jetstream

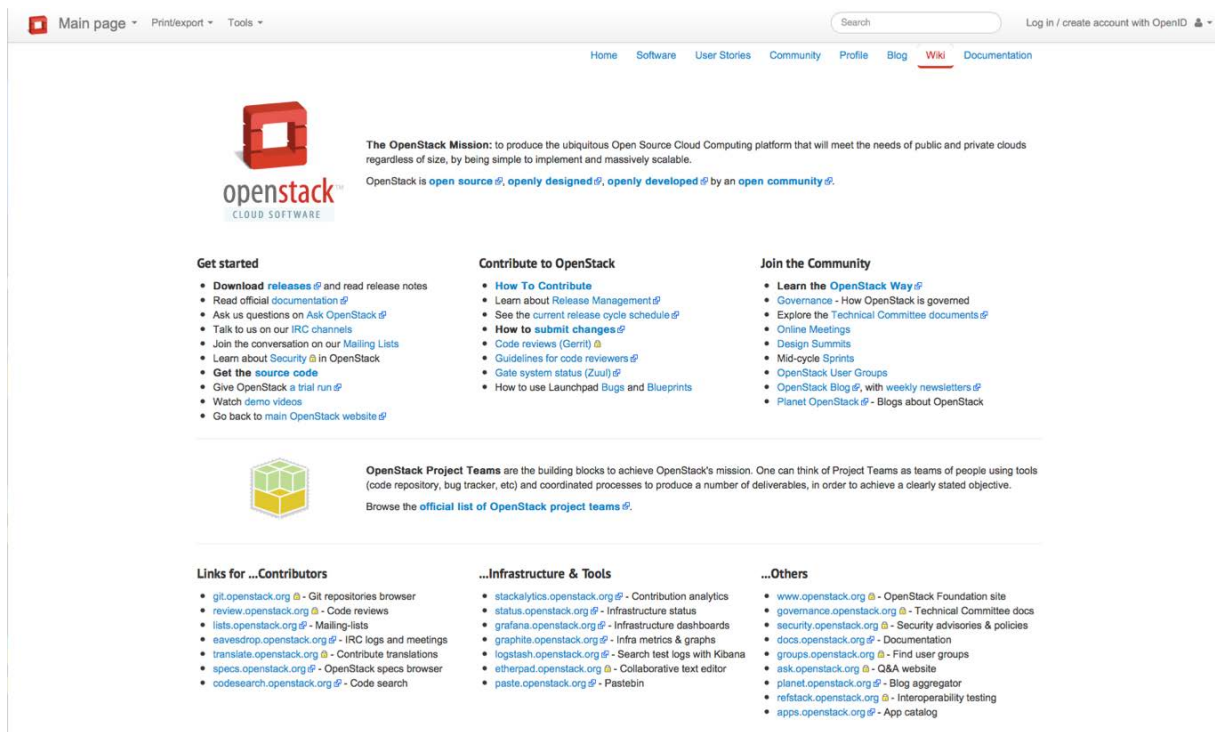
OpenStack



funded by the National Science Foundation
Award #ACI-1445604



OpenStack Organization



The screenshot shows the OpenStack website homepage. At the top, there is a navigation bar with links for 'Main page', 'Print/export', 'Tools', 'Home', 'Software', 'User Stories', 'Community', 'Profile', 'Blog', 'Wiki', and 'Documentation'. A search bar and a 'Log in / create account with OpenID' link are also present. The main content area features the OpenStack logo and a mission statement: 'The OpenStack Mission: to produce the ubiquitous Open Source Cloud Computing platform that will meet the needs of public and private clouds regardless of size, by being simple to implement and massively scalable.' Below this, it states 'OpenStack is open source, openly designed, openly developed by an open community.' The page is organized into several sections: 'Get started' with links to releases, documentation, IRC channels, mailing lists, and source code; 'Contribute to OpenStack' with links for how to contribute, release management, current release cycle, submitting changes, code reviews, guidelines for reviewers, gate system status, and using Launchpad; 'Join the Community' with links to learn the OpenStack way, governance, technical committee documents, online meetings, design summits, mid-cycle sprints, user groups, blog, and planet; 'OpenStack Project Teams' with a description of their role and a link to the official list; 'Links for ...Contributors' with links to git repositories, code reviews, mailing lists, IRC logs, translations, specs browser, and code search; '...Infrastructure & Tools' with links to analytics, infrastructure status, dashboards, metrics, search logs, collaborative text editor, and pastebin; and '...Others' with links to the foundation site, technical committee docs, security advisories, documentation, user groups, Q&A website, blog aggregator, interoperability testing, and app catalog.

OpenStack : the Project Navigator

The screenshot shows the OpenStack Project Navigator website. At the top, there is a blue navigation bar with the OpenStack logo, a search bar, and links for SOFTWARE, USERS, COMMUNITY, MARKETPLACE, EVENTS, LEARN, DOCS, JOIN, and LOGIN. Below the navigation bar, the main content area is titled "Software" and features a "PROJECT NAVIGATOR" tab. A search bar is present with the text "Enter a keyword". The main content area displays "Browse All OpenStack Projects" and provides a brief description of the Project Navigator's purpose. Below this, there is a section titled "Core Services (6 Results)" which contains six project cards. Each card displays the project name, a brief description, and three circular gauges representing Adoption, Maturity, and Age. The cards are: NOVA (Compute), NEUTRON (Networking), SWIFT (Object Storage), CINDER (Block Storage), KEYSTONE (Identity), and GLANCE (Image Service).

Project Name	Category	Adoption	Maturity	Age
NOVA	Compute	93%	8 of 8	6 yrs
NEUTRON	Networking	84%	8 of 8	5 yrs
SWIFT	Object Storage	52%	7 of 8	6 yrs
CINDER	Block Storage			
KEYSTONE	Identity			
GLANCE	Image Service			

<http://www.openstack.org/software/project-navigator/>

Openstack Projects ...the core services

Service	Name	Adoption	Maturity	Age
Identity	Keystone	96%	7/8	5 yrs
Images	Glance	95%	6/8	7 yrs
Block device	Cinder	88%	7/8	5 yrs
Networking	Neutron	93%	7/8	5 yrs
Compute	Nova	95%	8/8	7 yrs
Object device	Swift	52%	7/8	7 yrs

<https://www.openstack.org/software/project-navigator/>

Openstack Projects ...some other services

Service	Name	Adoption	Maturity	Age
Dashboard	Horizon	87%	6/8	5 yrs
Telemetry	Ceilometer	55%	1/8	4 yrs
Orchestration	Heat	67%	6/8	4 yrs
Containers	Magnum	11%	2/8	2 yrs
Map/Reduce	Sahara	10%	3/8	3 yrs

<https://www.openstack.org/software/project-navigator/>



funded by the National Science Foundation
Award #ACI-1445604



Openstack Projects ...some other services

Service	Name	Adoption	Maturity	Age
Shared Filesystems	Manila	14%	5/8	3 yrs
Workflow	Mistral	5%	1/7	1 yr
Load Balancing as a Service	Octavia	>0%	1/7	1 yr

<https://www.openstack.org/software/project-navigator/>



funded by the National Science Foundation
Award #ACI-1445604



Hardware & Infrastructure



funded by the National Science Foundation
Award #ACI-1445604



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 2 to 1 Fat Tree topology

Benchmarks – single node



funded by the National Science Foundation
Award #ACI-1445604



HPCC results : VM vs BareMetal Comparison

Benchmark	VM/ BareMetal	Units	What's tested
HPL	97%	FLOPS	floating point execution rate for solving a system of linear equations
DGEMM	98%	FLOPS	floating point execution rate for double precision real matrix-matrix multiplication
Bandwidth	88%	B/s	bytes/unit_time it takes to transmit a 2MB message from one node to another
Latency	97%	s	time required to send an 8-byte message from one node to another

HPCC results : VM vs BareMetal Comparison (Cont.)

Benchmark	VM/ BareMetal	Units	What's tested
Random	80%	up/s	rate of random updates of memory
Stream	77%	B/s	sustained memory bandwidth
MPI-FFT	67%	FLOPS	floating point rate of execution of double precision complex one-dimensional Discrete Fourier Transform
Ptrans	64%	B/s	rate of transfer for large arrays of data from multiprocessor's memory

Jetstream System Details



funded by the National Science Foundation
Award #ACI-1445604



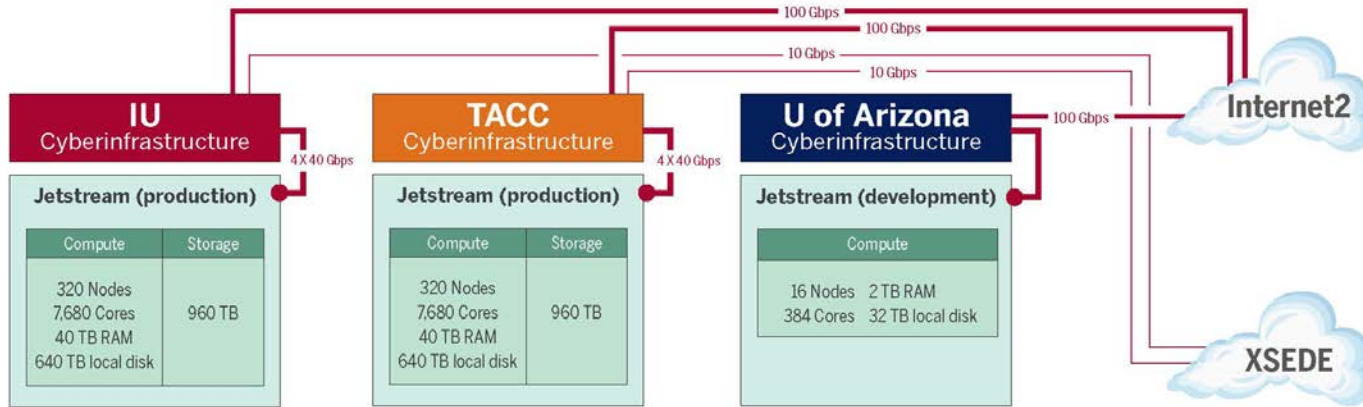
VM Instance Sizes (Flavors)

Instance Type	vCPUs	RAM(GB)	Storage(GB)	Instances/Node
Tiny	1	2	8	46
Small	2	4	20	23
Medium	6	16	60	7
Large	10	30	120/60*	4
X-Large	22	60	240/60*	2
XX-Large	44	120	480/60*	1


Node config: dual Intel E-2680v3 “Haswell”, 24 physical cores/node @ 2.5 GHz, 128 GB RAM, dual 1 TB local disks.

* Effective 29-Mar-2017

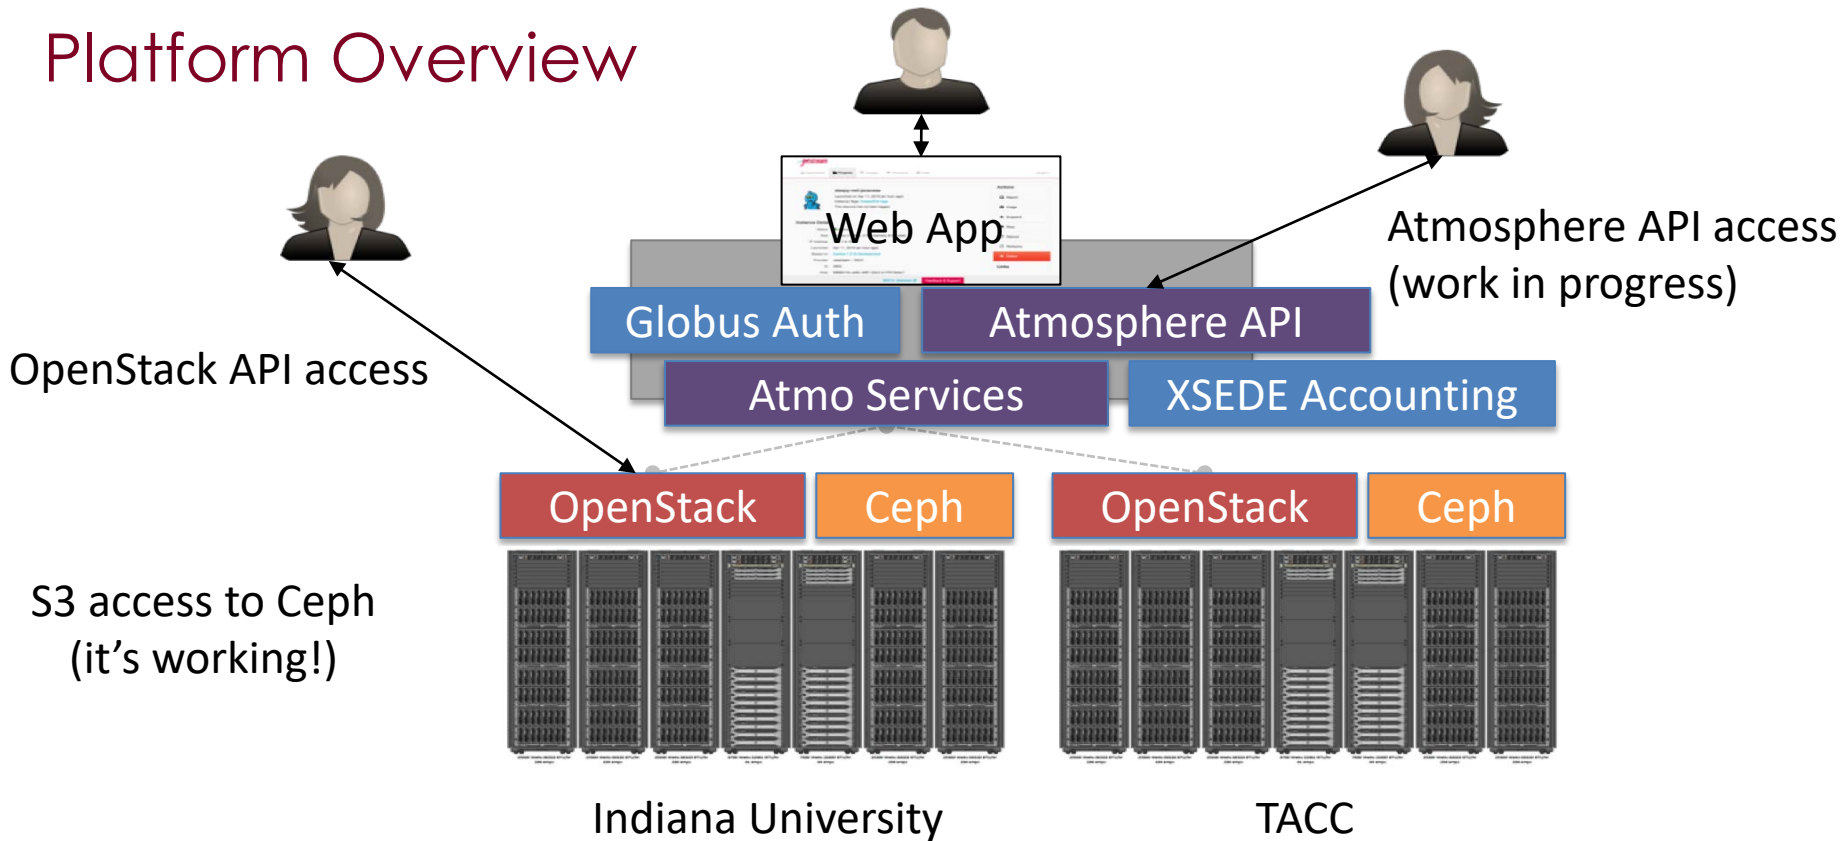
Jetstream Systems



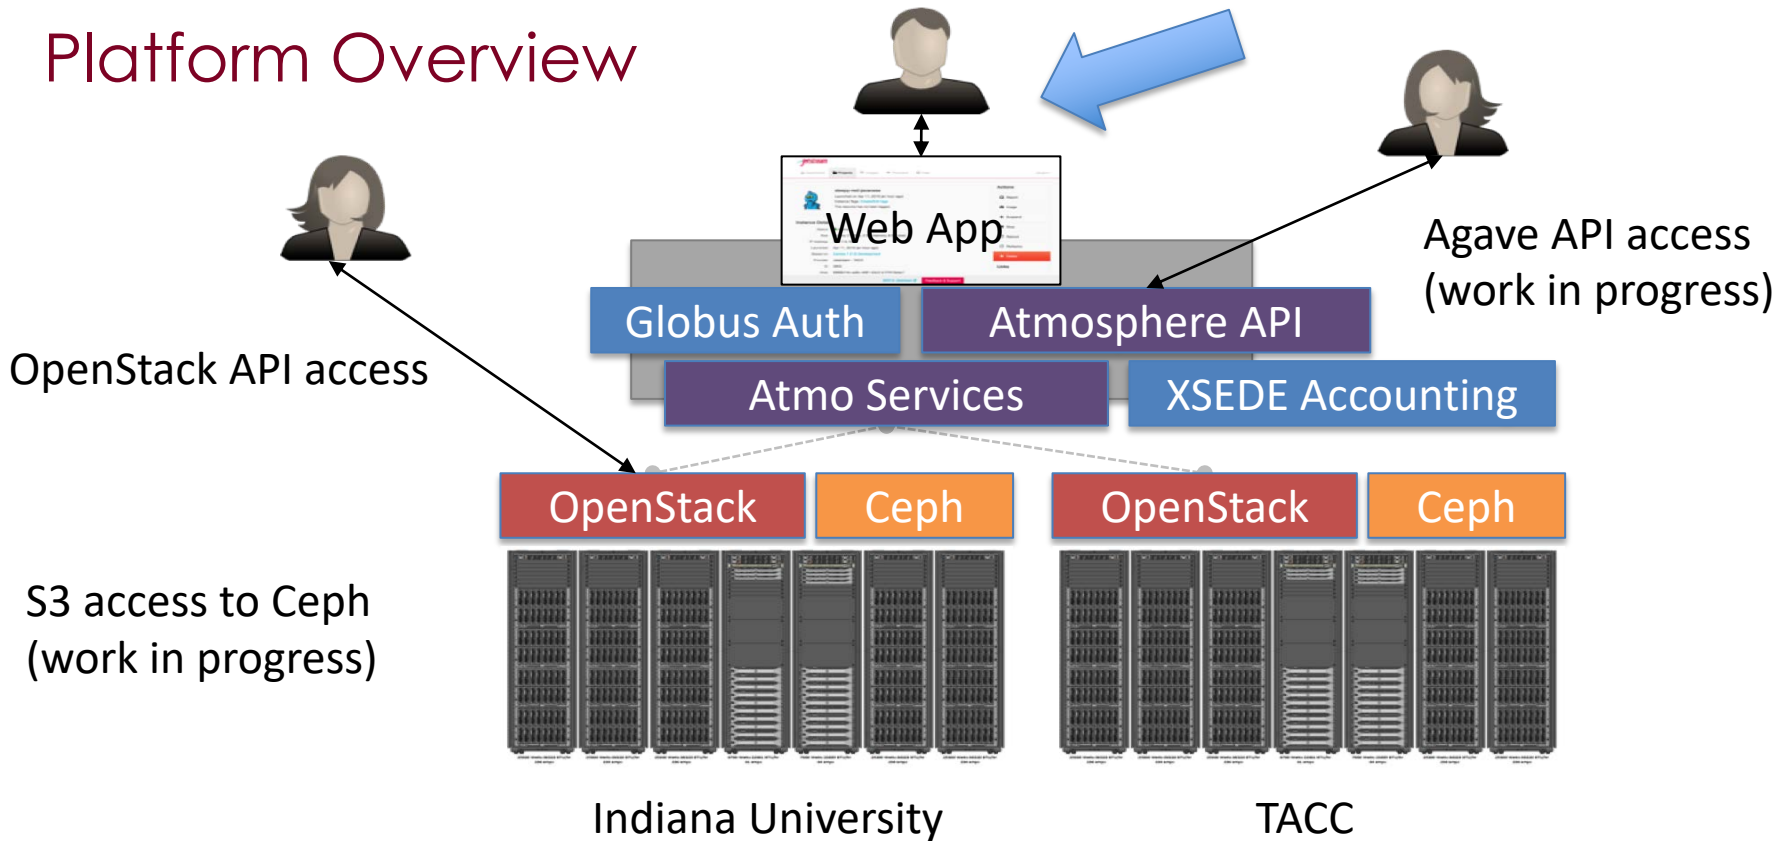
Jetstream Systems (cont.)



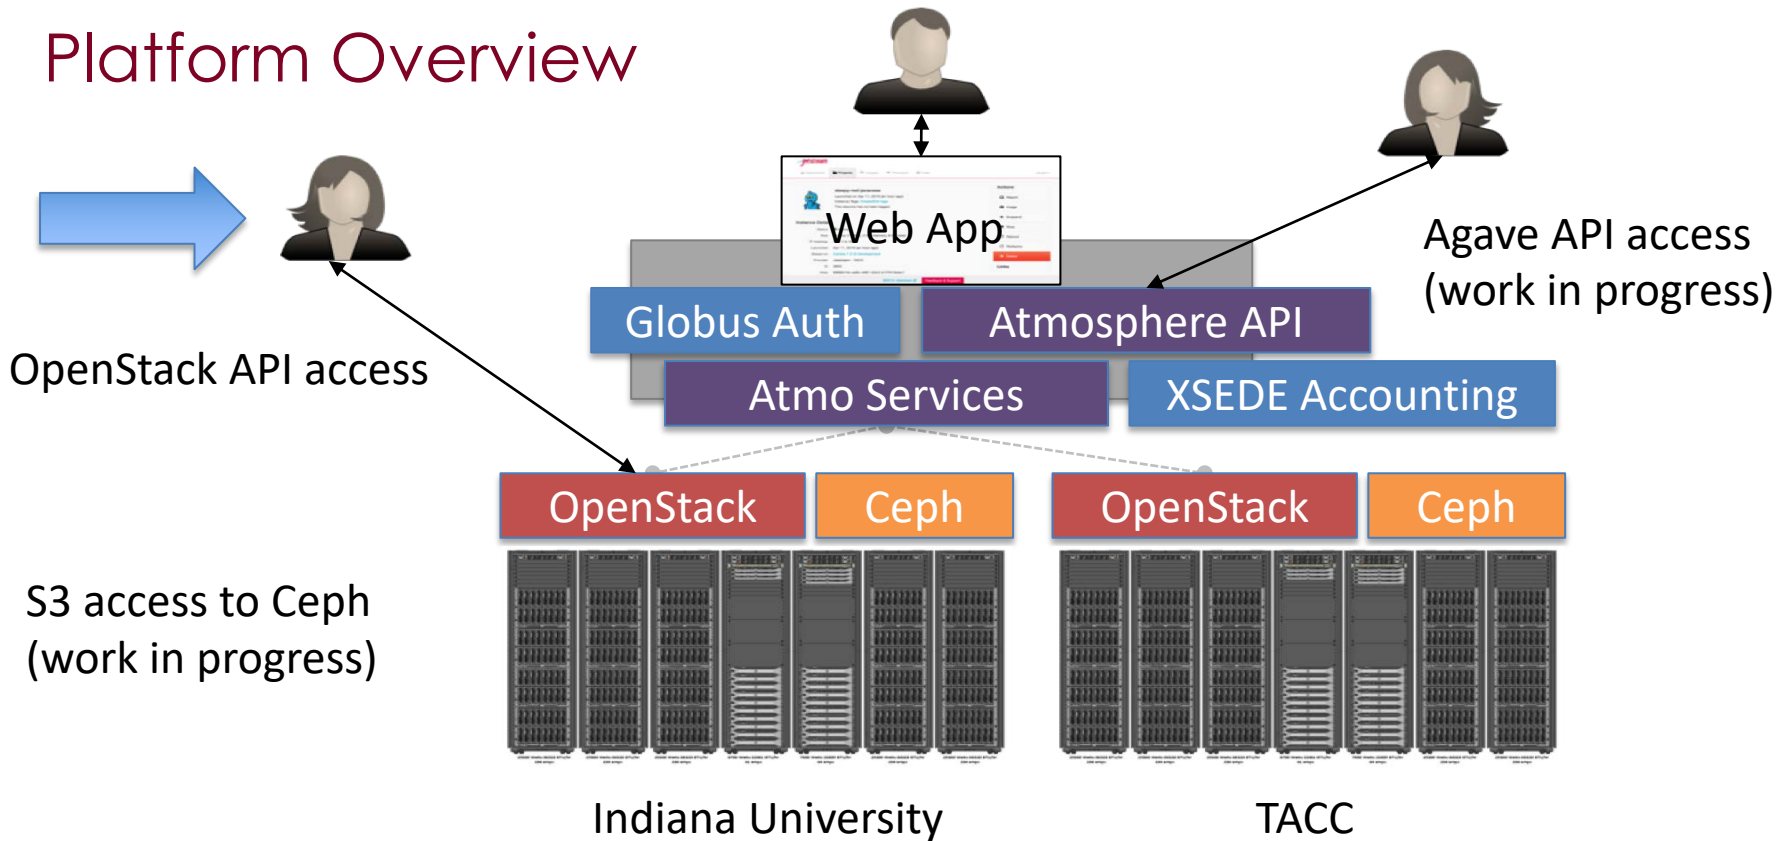
Platform Overview



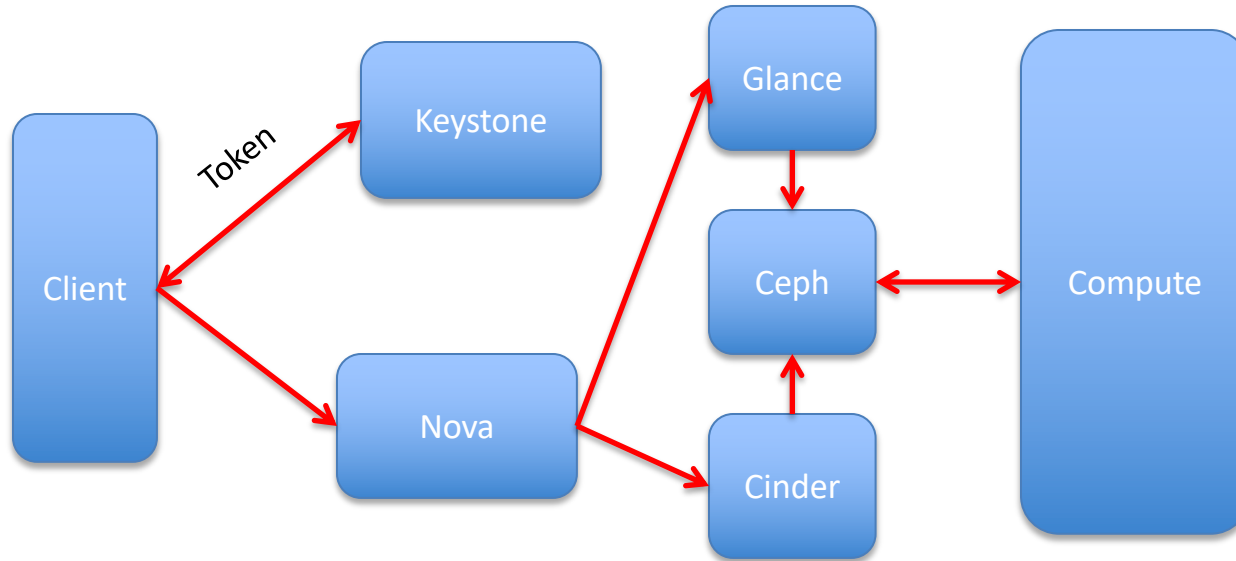
Platform Overview



Platform Overview



OpenStack Overview



How do we onboard users onto Jetstream?

- An XSEDE User Portal (XUP) account is required. They are free! Get one at <https://portal.xsede.org>
- Work with your XSEDE Campus Champion.
- Submit an **allocations request**
 - Read the Allocations Overview - <https://portal.xsede.org/allocations-overview>
 - Writeup an allocation request – ***start with a Startup or Education request*** - <https://portal.xsede.org/successful-requests>
- **Easy Button**: instant access to small, limited instances while the allocation request is processed and the user is vetted.

Jetstream Information Sources

- Twitter: @jetstream-cloud
- Jetstream's web interface: <https://use.jetstream-cloud.org/>
No login required to browse image library
- XSEDE User Portal account is required to actually login:
<https://portal.xsede.org> Create account in seconds.
- Jetstream Home page: <https://jetstream-cloud.org/>
- Jetstream's public documentation: <https://wiki.jetstream-cloud.org>

Jetstream Information Sources (Cont.)

- [Paper describing Jetstream Jetstream: A self-provisioned, scalable science and engineering cloud environment](#)
- Configuration management: <https://github.com/jetstream-cloud/Jetstream-Salt-States>


For questions, comments, etc. of any manner
help@jetstream-cloud.org



funded by the National Science Foundation
Award #ACI-1445604



Jetstream Partners



funded by the National Science Foundation
Award #ACI-1445604



Questions?

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

Project email: jethelp@iu.edu

Direct email: turnerg@iu.edu

License Terms

- Turner, G.. 2017. Jetstream: Accessing Jetstream via the OpenStack Command Line Interface; Open Cloud Institute UTSA, San Antonio, TX Also 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



Resources for this class

- The **Computational Science in the Cloud Institute** class
 - https://tacc.github.io/CSC2017Institute/docs/day4/intro_to_openstack.html
 - Utilizes Jupyter notebooks
 - Cut - & Paste examples
 - Use only the “new” commands
 - The “old” commands give insight into the various OpenStack Projects

Resources for this class

- General API access to Jetstream information
 - <https://iujetstream.atlassian.net/wiki/spaces/JWT/pages/39682057/Using+the+Jetstream+API>
- Setting up the openrc
 - <https://iujetstream.atlassian.net/wiki/spaces/JWT/pages/39682064/Setting+up+openrc.sh>
- Openstack Command Line examples for today's class
 - <https://iujetstream.atlassian.net/wiki/spaces/JWT/pages/35913730/OpenStack+command+line>
 - Cut-&-Paste examples

Resources for this class

- Openstack Command Line example
 - <https://iujetstream.atlassian.net/wiki/spaces/JWT/pages/35913730/OpenStack+command+lin>
- Jetstream-cloud.org
 - Click **Get Started**
 - Click **Jetstream wiki** in first paragraph
 - In the left column, scroll to near bottom
 - Click on **Using the Jetstream API**
 - Entry will expand
 - ~3 lines down, click on **OpenStack Command Line**

Getting started

- `ssh utsa-class-cmd.jetstream-cloud.org`
- username and password on paper
- `cat openrc`
- `source openrc`
- `printenv | grep OS | grep -v PASS`
- openstack image list
 - Is it working?