

Jetstream

Jetstream: Adding Cloud-based
Computing to the National
Cyberinfrastructure

Matthew Vaughn (@mattdotvaughn)

ORCID 0000-0002-1384-4283

Director, Life Science Computing
co-PI, Jetstream | Cyverse | Araport
Texas Advanced Computing Center



funded by the National Science Foundation
Award #ACI-1445604

What is Jetstream?

- A resource to *expand the community of users* who benefit from NSF investment in shared cyberinfrastructure
- *Production cloud system* supporting all domains of science and engineering research sponsored by the NSF
- Provide on-demand *interactive* computing and analysis
- Enable *configurable* environments and architectures
- Support computational *reproducibility and sharing*
- Democratizes access to *cloud-native* technology and software
- Focuses on ease of use, but maintaining flexibility



Funded by the National Science Foundation
Award #ACI-1445604



Expanding NSF XD's reach and impact

Around 299,000 researchers, educators, & learners received NSF support in 2012-2013


- Only 1.5% completed a computation, data analysis, or visualization task on XD program resources
- Less than 3% 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

* <https://www.xsede.org/xsede-nsf-release-cloud-survey-report>

Expanding NSF XD's reach and impact



Funded by the National Science Foundation
Award #ACI-1445604

Who do we expect will use Jetstream?

- Researchers, developers, and scientists who...
 - Require somewhere between 4 and a few hundred cores RIGHT NOW and for the foreseeable future, but not forever
 - Want the ability to fully customize the OS and configuration of their computing setup
 - Wish to move cloud-native workflows to/from an academic environment
 - Need interactive mode access to a computing and analysis system



Funded by the National Science Foundation
Award #ACI-1445604



Who do we expect will use Jetstream?

- Also...
 - Science gateway operators using Jetstream as either the frontend or processor for scientific jobs
 - Anyone who is evaluating or experimenting with software not traditionally supported on XD systems
 - STEM Educators teaching on a variety of subjects



Funded by the National Science Foundation
Award #ACI-1445604



21st century workforce development

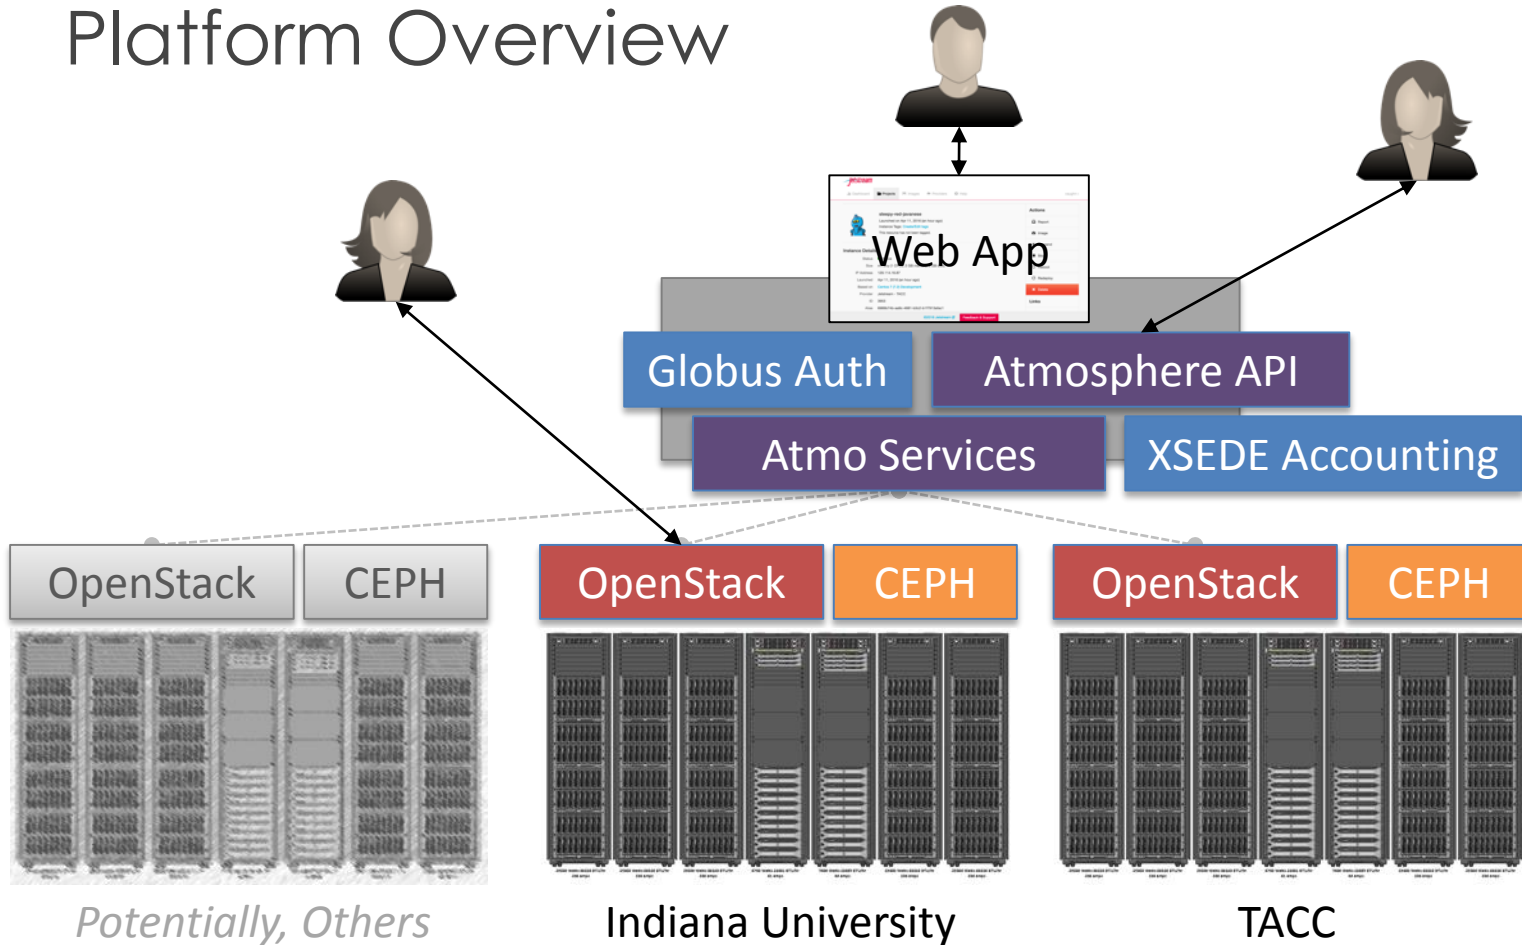
- Specialized virtual Linux desktops and applications to enable research and research education at small colleges and universities
 - HBCUs (Historically Black Colleges and Universities)
 - MSIs (Minority Serving Institutions)
 - Tribal colleges
 - Higher-education institutions in EPSCoR States
- Also, *complete democratization* of access to cloud-native technologies and approaches



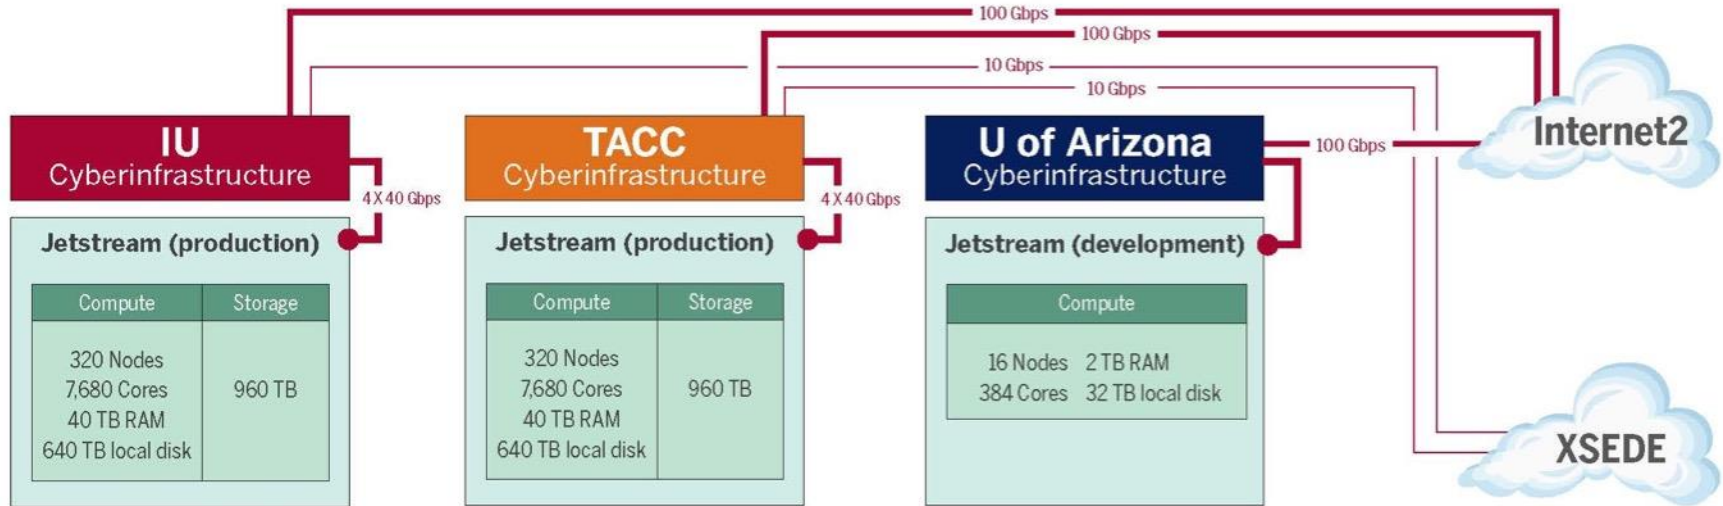
Funded by the National Science Foundation
Award #ACI-1445604



Platform Overview



Systems Overview



Funded by the National Science Foundation
Award #ACI-1445604

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

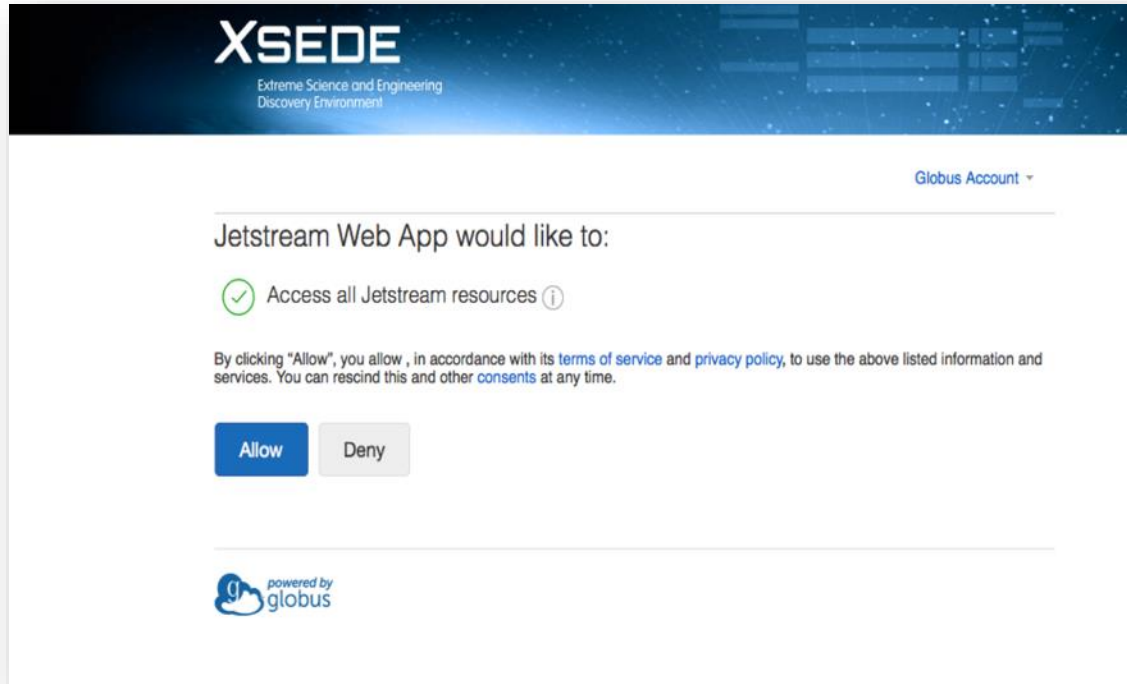
CEPH Storage

- 20x Dell 730xd per cloud
- 2x10Gbs bonded NIC per 730xd
- Running CEPH 0.94.5 Hammer
- Configured as OpenStack Storage

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

- 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

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*

Atmosphere x Matthew

https://use.jetstream-cloud.org/application/dashboard

Jetstream Dashboard Projects Images Providers Help vaughn -

Resources in Use [Need more?](#)

Resource	Jetstream - Indiana University (%)	Jetstream - TACC (%)
CPU	100	100
Memory	~95	100
Storage	~2	~1
Volumes	~20	~10
Allocation	40	~22

63 Instances

- active
- shutoff
- active - networking
- suspended

3 Volumes

- available

Updated 15 minutes ago

Instance History (1458 instances launched)

Launched from [Ubuntu 14.04.3 Development GUI](#)
Ran: Apr 11, 2016 06:08 am - Present

Launched from [Ubuntu 14.04.3 Development](#)

0 days ago on Jetstream - TACC

0 days ago on Jetstream - TACC

Apr 01, 2016 03:23 pm
[CentOS 6 \(6.7\) Development](#)

Mar 31, 2016 10:46 am
[CentOS 6 \(6.7\) Development GUI](#)

Mar 30, 2016 10:08 am
[Galaxy 16.01 Standalone](#)

Mar 29, 2016 04:38 pm
[Ubuntu 14.04.3 Development](#)

Mar 28, 2016 06:26 pm
[Ubuntu 14.04.3 Development GUI](#)

Mar 18, 2016 04:50 pm
[CentOS 7 - GUI](#)

Mar 16, 2016 06:07 pm
[Centos 7 \(7.2\) Development](#)

Mar 16, 2016 05:26 pm
[CentOS 6 \(6.7\) Development](#)

©2016 Jetstream [Feedback & Support](#)

Atmosphere x Matthew

https://use.jetstream-cloud.org/application/images





Jetstream Dashboard Projects Images Providers Help vaughn

SEARCH FAVORITES (5) MY IMAGES (0) MY IMAGE REQUESTS TAGS

centos

Showing 9 of 9 images for 'centos' List Grid

All Images

 Wrangler iRODS -- Cent... Apr 7th 2016 12:28 pm CDT by hannac	 CentOS 6 (6.7) Develo... Apr 1st 2016 03:23 pm CDT by jfischer	 CentOS 6 (6.7) Develo... Mar 31st 2016 10:46 am CDT by jfischer	 Ubuntu 14.04.3 Develo... Mar 29th 2016 04:38 pm CDT by admin
--	---	--	---

©2016 Jetstream Feedback & Support

Atmosphere x
https://use.jetstream-cloud.org/application/images/37

Jetstream Dashboard Projects Images Providers Help vaughn

SEARCH FAVORITES (4) MY IMAGES (0) MY IMAGE REQUESTS TAGS

← Centos 7 (7.2) Development

Name: Centos 7 (7.2) Development

Created: 3/16/2016 06:07 pm CDT

Created by: jfischer


Description: Based on Centos 7 Cloud 1601 Image, Development tools added, Patched up to date as of 3/16/16

Tags: CentOS development Featured

Launch

+ Add to Project

Versions:

	1.0 3/25/2016 01:08 pm CDT by jfischer
---	--

©2016 Jetstream Feedback & Support

Atmosphere x Matthew

https://use.jetstream-cloud.org/application/projects/48/resources

Jetstream Dashboard Projects Images Providers Help vaughn -

RESOURCES DETAILS OPTIONS

atmospherv2.apiary.io

NEW [Refresh] [Stop] [Folder]

Instances

Name	Status	IP Address	Size	Provider
<input type="checkbox"/> sleepy-cyan-badger	Active	129.114.104.247	m1.medium	Jetstream - TACC
<input type="checkbox"/> ugly-tomato-chamols	Active	129.114.104.246	m1.xlarge	Jetstream - TACC
<input type="checkbox"/> scummy-magenta-bobcat	Active	129.114.104.240	m1.medium	Jetstream - TACC
<input type="checkbox"/> sloppy-wheat-tang	Active	129.114.104.239	m1.tiny	Jetstream - TACC
<input type="checkbox"/> lousy-plum-salamander	Shutoff	N/A	m1.small	Jetstream - TACC
<input checked="" type="checkbox"/> ugly-magnolia-sponge	Active	129.114.104.244	m1.medium	Jetstream - TACC
<input type="checkbox"/> droopy-persimmon-bulldog	Active	129.114.104.242	m1.medium	Jetstream - TACC
<input type="checkbox"/> thirsty-xanthic-collie	Active	129.114.104.233	m1.medium	Jetstream - TACC
<input type="checkbox"/> zippy-cream-spaniel	Active	129.114.104.245	m1.tiny	Jetstream - TACC
<input type="checkbox"/> beady-jade-leopard	Active	129.114.104.230	m1.tiny	Jetstream - TACC

DETAILS

ID: 3846
Status: Active
Size: m1.medium (6 CPUs, 16 GB memory, 60 GB disk)
IP Address: 129.114.104.244
Launched: Apr 11, 2016 (5 hours ago)
Based on: CentOS 6 (6.7) Development
Provider: Jetstream - TACC

©2016 Jetstream Feedback & Support

Atmosphere x


Matthew

https://use.jetstream-cloud.org/application/projects/48/instances/3845

Jetstream Dashboard Projects Images Providers Help vaughn -

atmospherev2.apiary.io

NEW Resources > droopy-persimmon-bulldog



droopy-persimmon-bulldog
Launched on Apr 11, 2016 (5 hours ago)
Instance Tags: [Create/Edit tags](#)
This resource has not been tagged.

Instance Details

Status	● Active
Size	m1.medium (6 CPUs, 16 GB memory, 60 GB disk)
IP Address	129.114.104.242
Launched	Apr 11, 2016 (5 hours ago)
Based on	CentOS 6 (6.7) Development GUI
Provider	Jetstream - TACC
ID	3845
Alias	c368939d-26e9-4300-94bd-dc253f538e3f

Actions

- Report
- Image
- Suspend
- Stop
- Reboot
- Redeploy
- Delete**

Links

- Open Web Shell**

©2016 Jetstream [Feedback & Support](#)

```
Atmosphere
Gate One - 129-114-104
https://use.jetstream-cloud.org/web_shell?location=129-114-104-242&ssh=ssh://vaughn@129.114.104.242:22
[Press Shift-F1 for help]
Host/IP or ssh:// URL [localhost]: ssh://vaughn@129.114.104.242:22
Connecting to ssh://vaughn@129.114.104.242:22

The following SSH identities are being used for this connection:
  vaughn_default

Last login: Mon Apr 11 07:31:48 2016 from 129.114.104.5

Atmosphere

Jetstream
For assistance, contact help@xsede.org
[vaughn@js-104-242 ~]$ uname -a
Linux js-104-242.jetstream-cloud.org 2.6.32-573.22.1.el6.x86_64 #1 SMP Wed Mar 23 03:35:39 UTC 2016 x86_64 x86_64 x86_64 GNU/Linux
[vaughn@js-104-242 ~]$ whoami
vaughn
[vaughn@js-104-242 ~]$ sudo su -
-bash-4.1# whoami
root
-bash-4.1#
```

3: vaughn@js-104-242:~

Atmosphere Gate One - 129-114-104- x Matthew

https://use.jetstream-cloud.org/web_shell?location=129-114-104-242&ssh=ssh://vaughn@129.114.104.242:22

```
(5/7): nss-tools-3.19.1-8.el6_7.3.21.0-0.3.el6_7.x86_64. | 302 kB 00:00
(6/7): nss-util-3.19.1-5.el6_7.3.21.0-0.3.el6_7.x86_64.d | 22 kB 00:00
(7/7): nss-util-devel-3.19.1-5.el6_7.3.21.0-0.3.el6_7.x8 | 18 kB 00:00
Finishing rebuild of rpms, from deltarpms
<delta rebuild> | 1.8 MB 00:02
Presto reduced the update size by 65% (from 1.8 M to 661 K).
Package(s) data still to download: 46 k
nss-sysinit-3.21.0-0.3.el6_7.x86_64.rpm | 46 kB 00:00
Running rpm_check_debug
Running Transaction Test
Transaction Test Succeeded
Running Transaction
Updating : nspr-4.11.0-0.1.el6_7.x86_64 1/16
Updating : nss-util-3.21.0-0.3.el6_7.x86_64 2/16
Updating : nspr-devel-4.11.0-0.1.el6_7.x86_64 3/16
Updating : nss-util-devel-3.21.0-0.3.el6_7.x86_64 4/16
Updating : nss-sysinit-3.21.0-0.3.el6_7.x86_64 5/16
Updating : nss-3.21.0-0.3.el6_7.x86_64 6/16
Updating : nss-devel-3.21.0-0.3.el6_7.x86_64 7/16
Updating : nss-tools-3.21.0-0.3.el6_7.x86_64 8/16
Cleanup : nss-devel-3.19.1-8.el6_7.x86_64 9/16
Cleanup : nss-tools-3.19.1-8.el6_7.x86_64 10/16
Cleanup : nss-util-devel-3.19.1-5.el6_7.x86_64 11/16
Cleanup : nss-3.19.1-8.el6_7.x86_64 12/16
Cleanup : nss-sysinit-3.19.1-8.el6_7.x86_64 13/16
Cleanup : nspr-devel-4.10.8-2.el6_7.x86_64 14/16
Cleanup : nss-util-3.19.1-5.el6_7.x86_64 15/16
Cleanup : nspr-4.10.8-2.el6_7.x86_64 16/16
Verifying : nss-util-3.21.0-0.3.el6_7.x86_64 1/16
Verifying : nss-sysinit-3.21.0-0.3.el6_7.x86_64 2/16
Verifying : nss-devel-3.21.0-0.3.el6_7.x86_64 3/16
Verifying : nss-util-devel-3.21.0-0.3.el6_7.x86_64 4/16
Verifying : nss-3.21.0-0.3.el6_7.x86_64 5/16
Verifying : nspr-devel-4.11.0-0.1.el6_7.x86_64 6/16
Verifying : nss-tools-3.21.0-0.3.el6_7.x86_64 7/16
Verifying : nspr-4.11.0-0.1.el6_7.x86_64 8/16
Verifying : nss-devel-3.19.1-8.el6_7.x86_64 9/16
Verifying : nss-3.19.1-8.el6_7.x86_64 10/16
Verifying : nspr-4.10.8-2.el6_7.x86_64 11/16
Verifying : nspr-devel-4.10.8-2.el6_7.x86_64 12/16
Verifying : nss-util-3.19.1-5.el6_7.x86_64 13/16
Verifying : nss-tools-3.19.1-8.el6_7.x86_64 14/16
Verifying : nss-sysinit-3.19.1-8.el6_7.x86_64 15/16
Verifying : nss-util-devel-3.19.1-5.el6_7.x86_64 16/16

Updated:
nspr.x86_64 0:4.11.0-0.1.el6_7
nspr-devel.x86_64 0:4.11.0-0.1.el6_7
nss.x86_64 0:3.21.0-0.3.el6_7
nss-devel.x86_64 0:3.21.0-0.3.el6_7
nss-sysinit.x86_64 0:3.21.0-0.3.el6_7
nss-tools.x86_64 0:3.21.0-0.3.el6_7
nss-util.x86_64 0:3.21.0-0.3.el6_7
nss-util-devel.x86_64 0:3.21.0-0.3.el6_7

Complete!
-bash-4.1#
```

Terminal Sharing Info
Viewers: 5 Settings

1: vaughn@js-104-242:~

6:31:22 AM







Atmosphere x Matthew

https://use.jetstream-cloud.org/application/settings

Jetstream Dashboard Projects Images Providers Help vaughn

Appearance

Select the Image and Instance icon set you would like to use:


 Identicons  Retro  Robots  Unicorns  Monsters  Wavatars


Advanced

SSH Configuration

Enable ssh access into launched instances.

name	public key
Ulfur_New	ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQCMzP0ywiQNWh5Lvw0Ybu639vIHjxAGly+LelH6BgfOwt3dr3M6+VXt66OnZp6JKqIcaUiPFvU77Bg9bYHvW+K2 43qqO8YalAA+Ghcd mwvaughn@Fjall

+ 

©2016 Jetstream  Feedback & Support

Atmosphere x Matthew

https://use.jetstream-cloud.org/application/providers

Jetstream Dashboard Projects Images Providers Help vaughn

Jetstream - Indiana University

42% (6592/15840) AUs AUs currently used	34 instances Number of instances consuming allocation	101 hours Time remaining before allocation runs out	92 AUs/hour Rate at which AUs are being used
---	---	---	--

Jetstream - TACC

26% (4089/15840) AUs AUs currently used	21 instances Number of instances consuming allocation	93 hours Time remaining before allocation runs out	126 AUs/hour Rate at which AUs are being used
---	---	--	---

©2016 Jetstream [Feedback & Support](#)

Programmatic Access



Web Service APIs

- Openstack - Official and unofficial clients + libs (i.e. boto)
- Atmosphere - Very beta. Getting language libraries “soon”
 - Preview @ <http://docs.atmospherev2.apiary.io/>

Automation, Orchestration, and Workflow

- Marathon/MESOS
 - <https://www.youtube.com/v/VzZfwHLmcL0>
- Docker Machine* + Swarm
- Apache Airavata
- ClusterMan & Elasticcluster*

Configuration Management Tools

- Vagrant & Terraform (Hashicorp)
- Chef

**Still quite finicky*

Example: Marathon, MESOS, Unidata

The image displays three overlapping browser windows illustrating a distributed computing environment:

- Marathon (Left):** Shows the Marathon web interface for the application `/idv-cloud-withpsw`. It includes controls for `Suspend` and `Scale`, and a list of tasks with their IDs and status.
- Mesos (Middle):** Shows the Mesos web interface for the cluster `scigap-1r.jetstream-cloud.org:5050/#`. It displays cluster details (Server: 10.0.0.4:5050, Version: 0.25.0) and lists **Active Tasks** and **Completed Tasks**. A table of **Slaves** is also visible, showing 5 activated and 0 deactivated nodes.
- Unidata (Right):** Shows a web browser window displaying a visualization of a globe with atmospheric data overlays, including contour lines and color-coded regions, likely generated by a task in the Mesos cluster.

What comes next?

- Both cloud systems + all software components installed, configured, and operational
- “Early operations mode” is underway
- End of March 2016: **38 XSEDE projects and 250+ users**
- Acceptance review scheduled with NSF in early May
- Full, unrestricted operations after system is accepted
- Soliciting *Research* allocation requests NOW plus *Startup* and *Education* allocations



Funded by the National Science Foundation
Award #ACI-1445604



XSEDE USER PORTAL
Extreme Science and Engineering
Discovery Environment

MY XSEDE RESOURCES DOCUMENTATION ALLOCATIONS TRAINING USER FORUMS HELP

Summary Allocations/Usage Accounts Jobs Profile Publications Tickets Change Password Add User Community

Welcome, Maytal!
Last login: Sat 04/09/16
at 03:39:17 PM CST

Profile Allocations
Accounts Training

NEW! Share your XSEDE
Science Achievements

Trial Allocations

Trial Jetstream Access

Trial Comet Access

My XSEDE Resources

Resource	Status	Load	Username
----------	--------	------	----------

In The Past 7 Days

XD SUs Charged: Total: by Field of Science

Field of Science	Value
Biophysics	31,750,889.0
Materials Research	8,426,303.0
Biochemistry and Molecular Structure and Function	6,745,370.0
Molecular Biosciences	5,824,900.0
Astronomical Sciences	5,413,394.0
Chemistry	4,307,433.0
Fluid, Particulate, and Hydraulic Systems	3,219,320.0
Gravitational Physics	2,331,651.0
Physical Chemistry	2,157,332.0
Extragalactic Astronomy and Cosmology	1,190,115.0
All 77 others	27,424,244.0

The “Easy Button”

Allow any user with active XSEDE User Portal account to use a small (but functional) slice of Jetstream

- Sign up for XUP Account
- Sign into User Portal
- Click “Trial Jetstream Access” button
- Have (restricted) access to Jetstream in ~30 minutes or less

Analogous to the free tier that makes it so easy to get started on the public cloud

Partner Organizations

Construction



THE UNIVERSITY
OF ARIZONA.




THE UNIVERSITY OF
CHICAGO



JOHNS HOPKINS
UNIVERSITY

Management & Operations Partners



Application / Community Leads



UNIVERSITY
of ARKANSAS
AT PINE BLUFF
—1873—



UNC
THE ODUM INSTITUTE



Funded by the National Science Foundation
Award #ACI-1445604



QUESTIONS?

vaughn@tacc.utexas.edu
@mattdotvaughn

[http:// use.jetstream-cloud.org](http://use.jetstream-cloud.org)
@jetstream_cloud

<http://portal.xsede.org>

How can I use Jetstream?

- An XSEDE User Portal (XUP) account is required. They are free! Get one at <https://portal.xsede.org>
- Read the Allocations Overview - <https://portal.xsede.org/allocations-overview>
- Write a successful allocation request – start with a Startup or Education request - <https://portal.xsede.org/successful-requests>



Where can I get help or learn more?

- Production:
 - 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>
 - Training Videos / Virtual Workshops (TBD)
- Early use:
 - <http://jetstream-cloud.org/>
 - Early use: jethelp@iu.edu

