rt.iu.edu

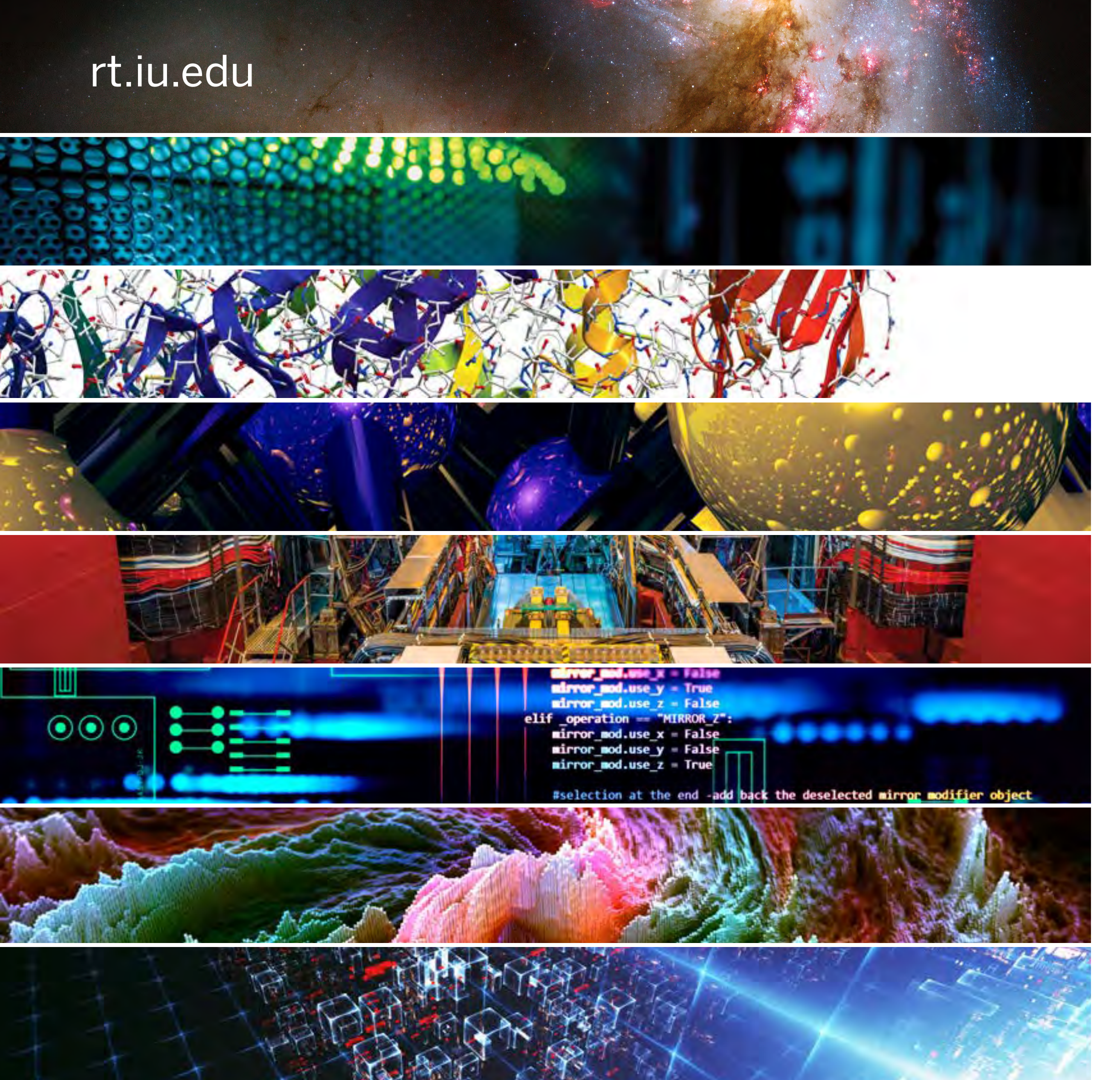

```
    mirror_mod.use_x = False
    mirror_mod.use_y = True
    mirror_mod.use_z = False
elif _operation == "MIRROR_Z":
    mirror_mod.use_x = False
    mirror_mod.use_y = False
    mirror_mod.use_z = True

#selection at the end -add back the deselected mirror modifier object
```

RESEARCH TECHNOLOGIES

UNIVERSITY INFORMATION TECHNOLOGY SERVICES

Jetstream Security Quick Look

Jetstream REU Program – Indiana University
June 14, 2021 – Bloomington, IN.

Jeremy Fischer – Jeremy@iu.edu - Indiana University

Manager, Jetstream Cloud, UITS Research Technologies

Fischer, J. (2021). Jetstream Security Quick Look. Bloomington, IN. Retrieved from https://jetstream-cloud.org/research/publications.php

OpenStack Overview

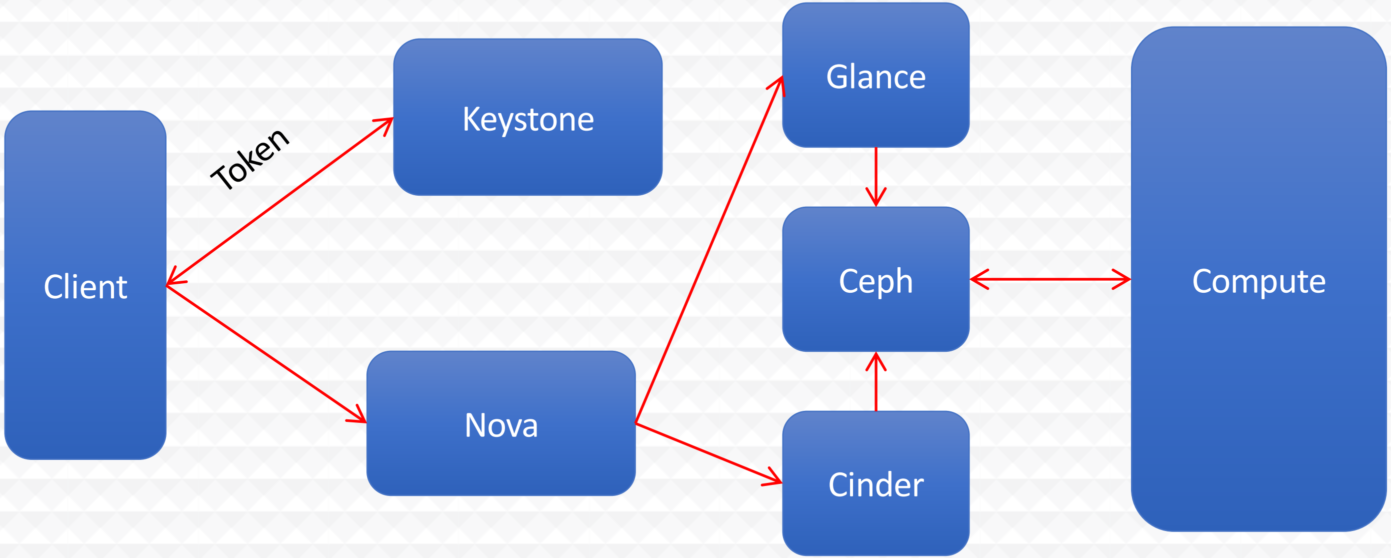

HPC vs Cloud

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**

** 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 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
  - **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.

-- Mike Lowe (Jetstream architect)

**some caveats for gateways...

What is Jetstream – a closer look

- **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.
  - **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 (https://docs.openstack.org/openstacksdk/latest/)
- **Slack channel** for collaboration API users of Jetstream
- Paved the way for 3rd party interfaces like Exosphere

Using the OpenStack CLI on Jetstream

What an openrc file looks like:

```
export OS_AUTH_URL=https://iu.jetstream-cloud.org:35357/v3
export OS_PROJECT_NAME="TG-ABC190028"
export OS_USER_DOMAIN_NAME="tacc"
export OS_USERNAME=”taccusername”
export OS_IDENTITY_API_VERSION=3
# export OS_PASSWORD=’string’
read -sr OS_PASSWORD_INPUT
export OS_PASSWORD=$OS_PASSWORD_INPUT
```

- Please do not publish the AUTH URLs anywhere
- CLI is python based – reads this information from the environment.
- Horizon can generate an openrc file for you (see the Wiki docs)
- Common pitfall – make sure you specify the correct Project (allocation) if you have more than one!

Installing the client

- Simple on most Mac OS X and Linux hosts (a single pip command)
- Less simple, but still do-able on Windows
  - Once you have a python installed, becomes a simple pip install
- Latest python-openstackclient (> 4.0.0) works with Python 3
- Best practice – use a virtual environment like virtenv for your install
- Docs on the wiki for this!
- Other CLI clients are available – e.g. python-swiftclient (Swift and S3), python-heatclient (Heat templates), etc
  - These are optional and not necessary for basic operations!

CLI / API Interface

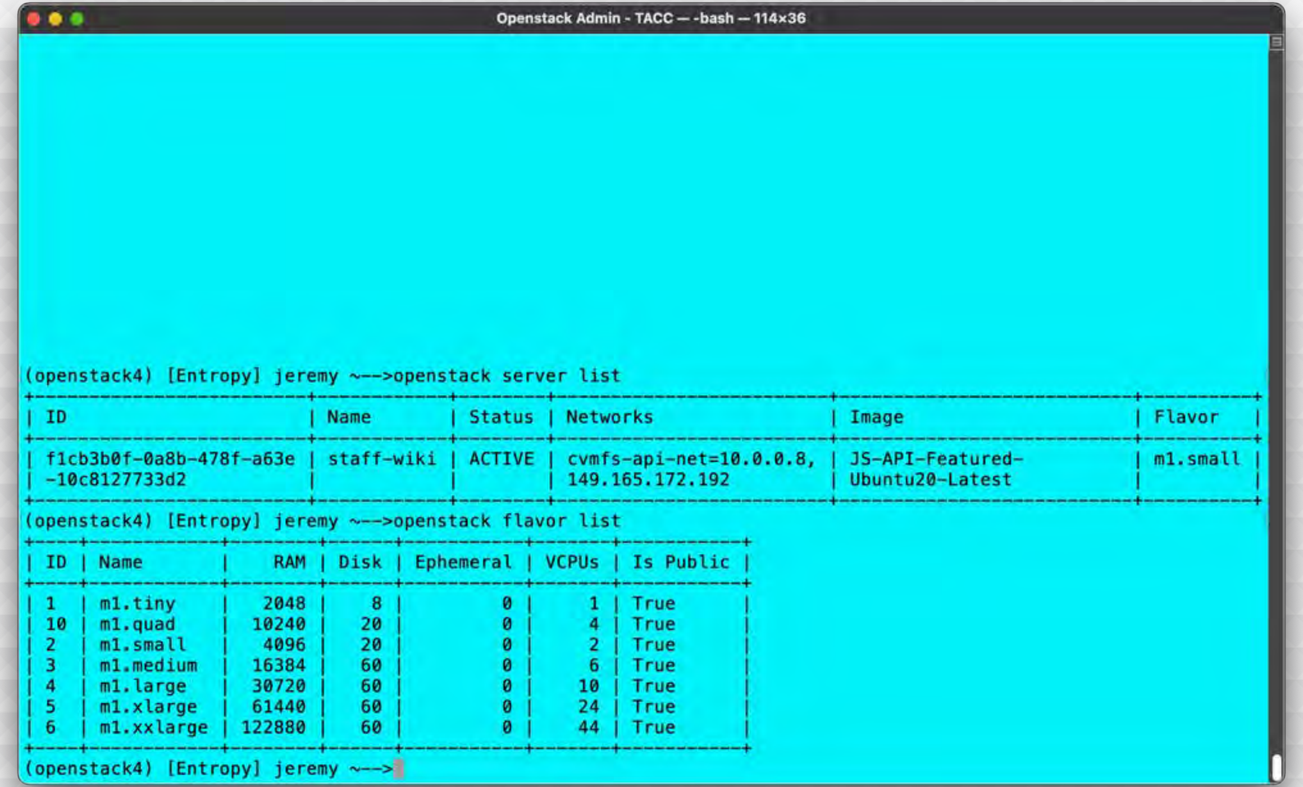

```
(openstack4) [Entropy] jeremy ~-->openstack server list
+-------------------------+-----------+--------+-------------------------+------------------------+----------+
| ID                      | Name      | Status | Networks                | Image                  | Flavor   |
+-------------------------+-----------+--------+-------------------------+------------------------+----------+
| f1cb3b0f-0a8b-478f-a63e | staff-wiki | ACTIVE | cvmfs-api-net=10.0.0.8, | JS-API-Featured-       | m1.small |
| -10c8127733d2           |           |        | 149.165.172.192         | Ubuntu20-Latest        |          |
+-------------------------+-----------+--------+-------------------------+------------------------+----------+
(openstack4) [Entropy] jeremy ~-->openstack flavor list
+----+-----------+--------+------+-----------+-------+-----------+
| ID | Name      |    RAM | Disk | Ephemeral | VCPUs | Is Public |
+----+-----------+--------+------+-----------+-------+-----------+
| 1  | m1.tiny   |   2048 |    8 |         0 |     1 | True      |
| 10 | m1.quad   |  10240 |   20 |         0 |     4 | True      |
| 2  | m1.small  |   4096 |   20 |         0 |     2 | True      |
| 3  | m1.medium |  16384 |   60 |         0 |     6 | True      |
| 4  | m1.large  |  30720 |   60 |         0 |    10 | True      |
| 5  | m1.xlarge |  61440 |   60 |         0 |    24 | True      |
| 6  | m1.xxlarge | 122880 |   60 |         0 |    44 | True      |
+----+-----------+--------+------+-----------+-------+-----------+
(openstack4) [Entropy] jeremy ~-->
```

Horizon GUI interface

- Allows most things you can do from the CLI
- Nice for some tasks
    - Network visualizer is something we tend to use as a troubleshooting tool
    - Easier to look at security groups on Horizon (IMHO)
- Downsides:
    - considerably slower than using CLI
    - not all features are present that are in CLI
    - can't do things programmatically

Log In

Domain: TACC

User Name:

Password:

Exosphere

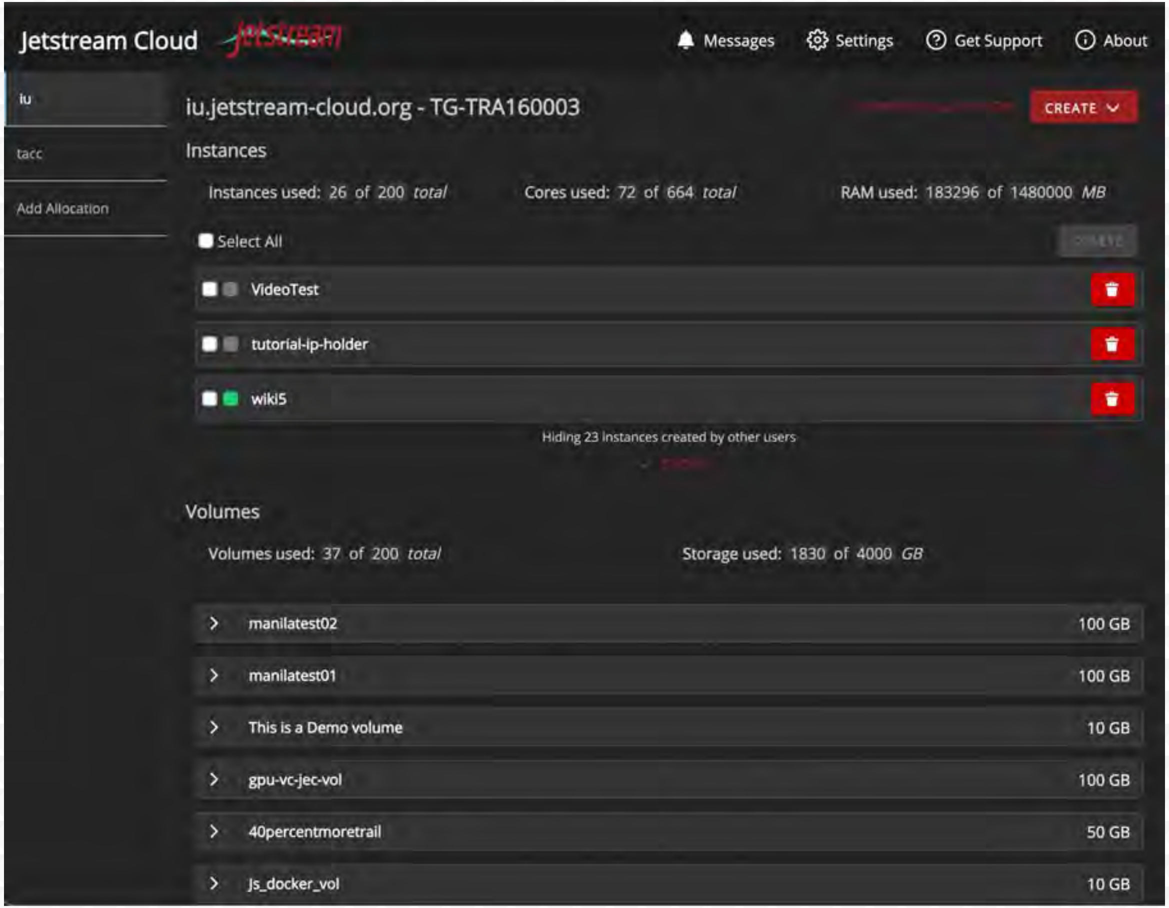

iu.jetstream-cloud.org - TG-TRA160003

Instances

Instances used: 26 of 200 *total*

Cores used: 72 of 664 *total*

RAM used: 183296 of 1480000 *MB*

□ Select All

- □ VideoTest
- □ tutorial-ip-holder
- □ 🟩 wiki5

Hiding 23 instances created by other users

Volumes

Volumes used: 37 of 200 *total*

Storage used: 1830 of 4000 *GB*

manilatest02	100 GB
manilatest01	100 GB
This is a Demo volume	10 GB
gpu-vc-jec-vol	100 GB
40percentmoretrail	50 GB
Js_docker_vol	10 GB

Exosphere GUI interface

- 3rd party GUI interface for OpenStack clouds
- Developers have a past connection to Jetstream but are working with multiple cloud providers
- Attempting to fill the gap between interfaces built for system administrators like OpenStack Horizon, and intuitive-but-proprietary services like DigitalOcean
- More about Exosphere:
  - https://gitlab.com/exosphere/exosphere

Getting started with the API

Things you'll set up once (hopefully):

- SSH keys
- Security groups (though you'll build on the basics as you do more advanced things)
- Create a network
- Create a subnet
- Create a router

Things you'll likely do many times:

- Create and launch instances
- Screw up and delete instances
- Launch more instances
- Expand security groups

API CLI Tutorial walkthrough: https://github.com/jlf599/JetstreamAPITutorial
API Horizon walkthrough: http://wiki.jetstream-cloud.org/Using+the+OpenStack+Horizon+GUI+Interface

API General Best Practices

- Jetstream-specific – don’t use Atmosphere images on the API side (start with JS-API-Featured-* images)
- Think about your security groups and only open what you REALLY need to open.
- Give objects unique and descriptive names
- When in doubt, use the universally unique identifier (UUID)
- When deleting items, use the universally unique identifier (UUID)
- Before deleting anything, though, “measure twice, cut once”
- Understand that an allocation/tenant lets you see everyone else’s things. Be aware and be ware of deleting things – do unto others…
- Put your toys away if you’re done with them

Security Best Practices

- Think about your security groups and only open what you REALLY need to open. (yes, it’s in the slides twice…on purpose…)
- In a production system, you’d likely want to also run a host-based firewall in addition to security groups (defense in depth!)
- Update often! Unattended security upgrades should be turned on in JS-API-Featured-* images…but still…
- Turn off any services/listeners you do not need
- For any service you run on a host, limit the access as much as possible – if it’s world accessible, make sure permissions and privileges are as limited as possible
- Limit the number of people that interactively login – and create accounts for them instead of using shared accounts (e.g. centos or ubuntu account)
- Monitor the logs – lots of tools out there to help with this!

Security groups…some thoughts

- Security groups layer – best to do in small, logical chunks for readability and management
- Security group updates happen in REAL TIME!
- Security group rules are OPPOSITE of traditional unix firewalls
- Make changes in small bites
- Conflicting rules can happen (and will)
- When restricting by network (slash) notation, that last number is crucial!
- It’s tempting to just completely open access – think carefully
- Security groups from the command line can be daunting at first

Troubleshooting and verifying your rules

- Starting simple usually works
  - Ping, ssh, telnet
- Tools like nmap (Network Mapper) are your friends
  - https://nmap.org/

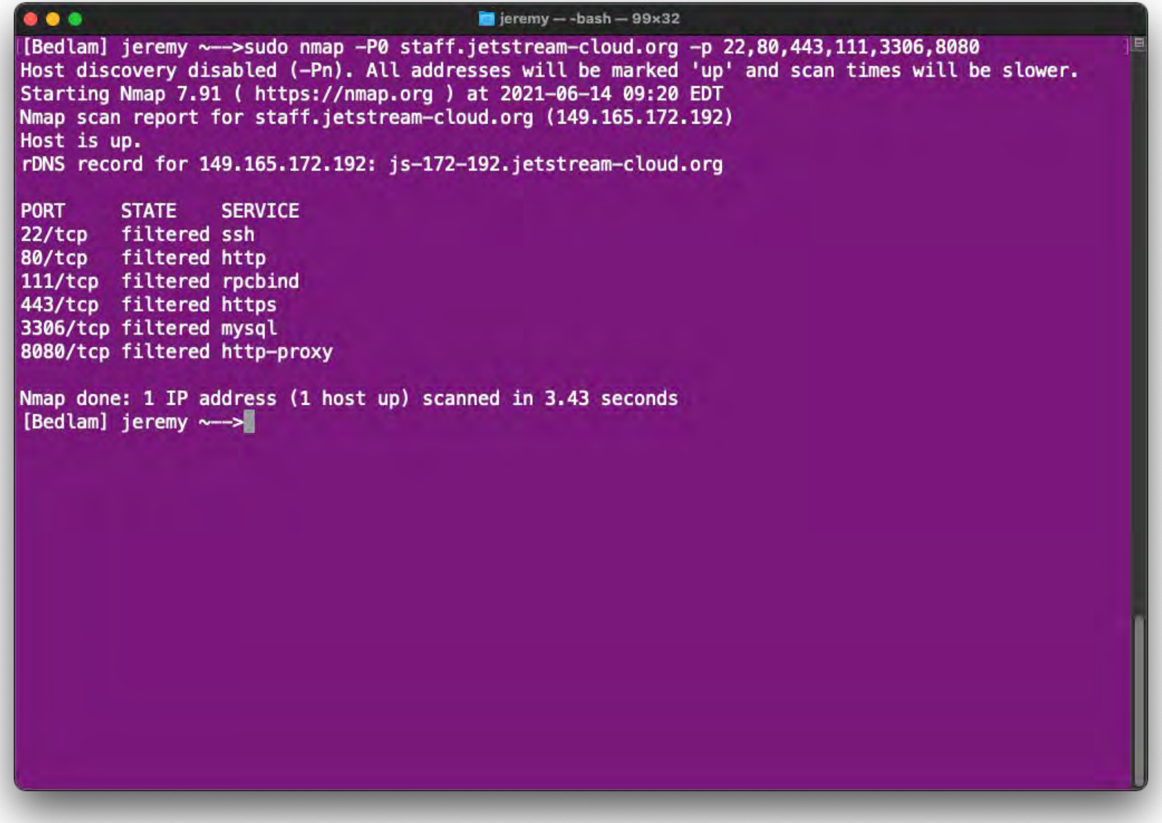

```
[Bedlam] jeremy ~-->sudo nmap -P0 staff.jetstream-cloud.org -p 22,80,443,111,3306,8080
Host discovery disabled (-Pn). All addresses will be marked 'up' and scan times will be slower.
Starting Nmap 7.91 ( https://nmap.org ) at 2021-06-14 09:20 EDT
Nmap scan report for staff.jetstream-cloud.org (149.165.172.192)
Host is up.
rDNS record for 149.165.172.192: js-172-192.jetstream-cloud.org

PORT     STATE    SERVICE
22/tcp   filtered ssh
80/tcp   filtered http
111/tcp  filtered rpcbind
443/tcp  filtered https
3306/tcp filtered mysql
8080/tcp filtered http-proxy

Nmap done: 1 IP address (1 host up) scanned in 3.43 seconds
[Bedlam] jeremy ~-->
```

Where can I get help?

- **Wiki / Documentation: http://wiki.jetstream-cloud.org**
- API CLI Tutorial: https://github.com/jlf599/JetstreamAPITutorial
- User guides: https://portal.xsede.org/user-guides
- XSEDE KB: https://portal.xsede.org/knowledge-base
- Email: help@xsede.org

Jetstream Partners

INDIANA UNIVERSITY
PERVASIVE TECHNOLOGY INSTITUTE

JOHNS HOPKINS
UNIVERSITY

THE UNIVERSITY OF ARIZONA®

TACC

THE UNIVERSITY OF CHICAGO

CORNELL UNIVERSITY FOUNDED A.D. 1865

THE UNIVERSITY OF TEXAS UTSA AT SAN ANTONIO

DELL

UNIVERSITY OF HAWAI'I MĀLAMALAMA 1907 UA MAU KE EA O KA 'ĀINA I KA PONO

NSIDC
National Snow and Ice Data Center

JSU 1877 JACKSON STATE UNIVERSITY®

UNC THE ODUM INSTITUTE

Jetstream
http://jetstream-cloud.org/

NSF funded by the National Science Foundation
Award #ACI-1445604

Questions?

- Project website: http://jetstream-cloud.org/
- Project email: help@jetstream-cloud.org Direct email: jeremy@iu.edu

License Terms

- Fischer, Jeremy. June 14, 2020. Jetstream Security Quick Look – Jetstream REU Program – Indiana University. 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.