



**PERVASIVE  
TECHNOLOGY INSTITUTE**



**RESEARCH TECHNOLOGIES**  
UNIVERSITY INFORMATION TECHNOLOGY SERVICES



**PERVASIVE  
TECHNOLOGY INSTITUTE**



**RESEARCH TECHNOLOGIES**  
UNIVERSITY INFORMATION TECHNOLOGY SERVICES

# Jetstream2: Accelerating cloud computing via Jetstream

**Jeremy Fischer & Mike Lowe – Indiana University**

Manager, Jetstream2 / Senior Cloud Engineer

CANOPIE-HPC 2021 Workshop @SC21 – November 2021



# What is “the” Jetstream?

- Fast moving air currents
- Hot/Cold air boundaries
- An NSF-funded cloud environment
- A project re-defining state-of-the-ART



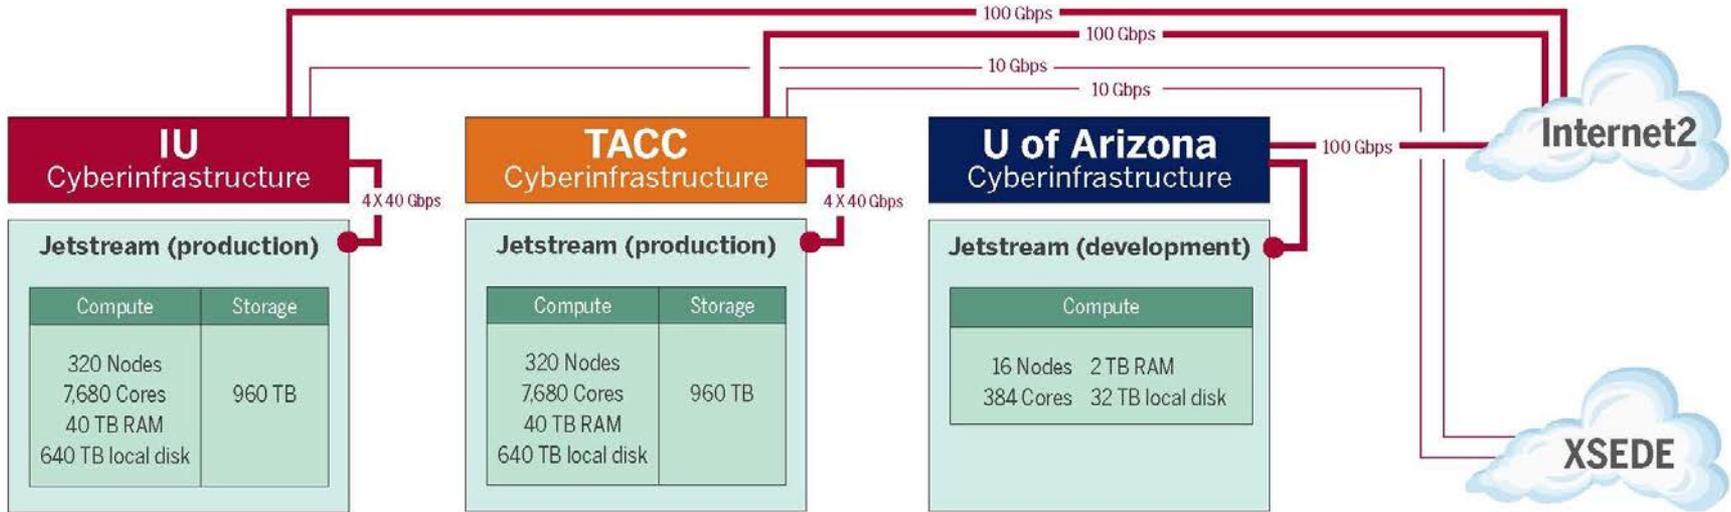
# What is Jetstream2 and why does it exist?

- Significant evolution of the Jetstream cloud resource
- Under 10% NSF investment → support for 24% of institutions, 23% of active PIs, and 32% of users\*
- Jetstream has provided **6x more** SUs than **any other** XSEDE resource for Education
- Emphasis on ease-of-use, broad accessibility, *AI for Everyone*
- Will provide **on-demand interactive** computing and persistent services for science gateways
- Enables *configurable* environments; *programmable cyberinfrastructure*



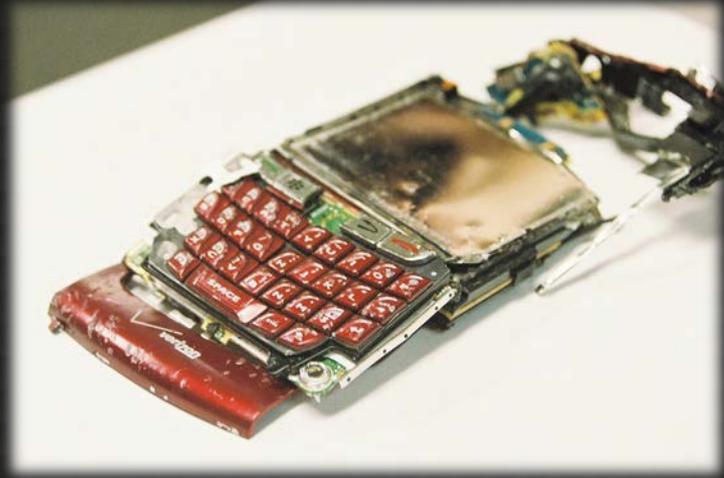
\*Based on XDMoD data at Workload Analysis Report: <http://arxiv.org/abs/1801.04306>

# Jetstream [1] System Overview



## What worked?

- Allowing API access and full control (root privileges)
- “Indefinite workflows” – allowing instances to run continuously – providing PIs renew their allocations
- Development of trial allocations



Flickr user MattHurst – Broken Blackberry

## What didn't work?

- Forcing small allocations into the research allocation process
- Lack of multi-year allocations
- Lack of shared data set storage

# Lessons learned

## Challenges -> Inspired changes

- Storage capacity -> Larger HDD pool and new flash storage
- Homogeneous hardware -> Inclusion of NVIDIA GPUs (w/MIG or vGPU) and memory diversity
- Separate OpenStack domains -> Unification of “Atmosphere” domain
- Virtual networking architecture/maintenance -> Increase offload capabilities via Cumulus Networks software and Mellanox hardware (NAT & simulation)
- Acceptance & integration into national CI ecosystem -> Changes to our metrics/KPIs & accounting processes
- Deployment diversity -> Leverage single technology for config management



D.Y. Hancock – Castello di Nipozzano 2017

# Jetstream2 Capabilities

Enhancing IaaS model of Jetstream:

- Improved orchestration support
- Elastic virtual clusters
- Federated JupyterHubs
- Ease storage sharing (CephFS w/Manilla)

Commitment to >99% uptime

- Critical for science gateway hosting
- Hybrid-cloud support

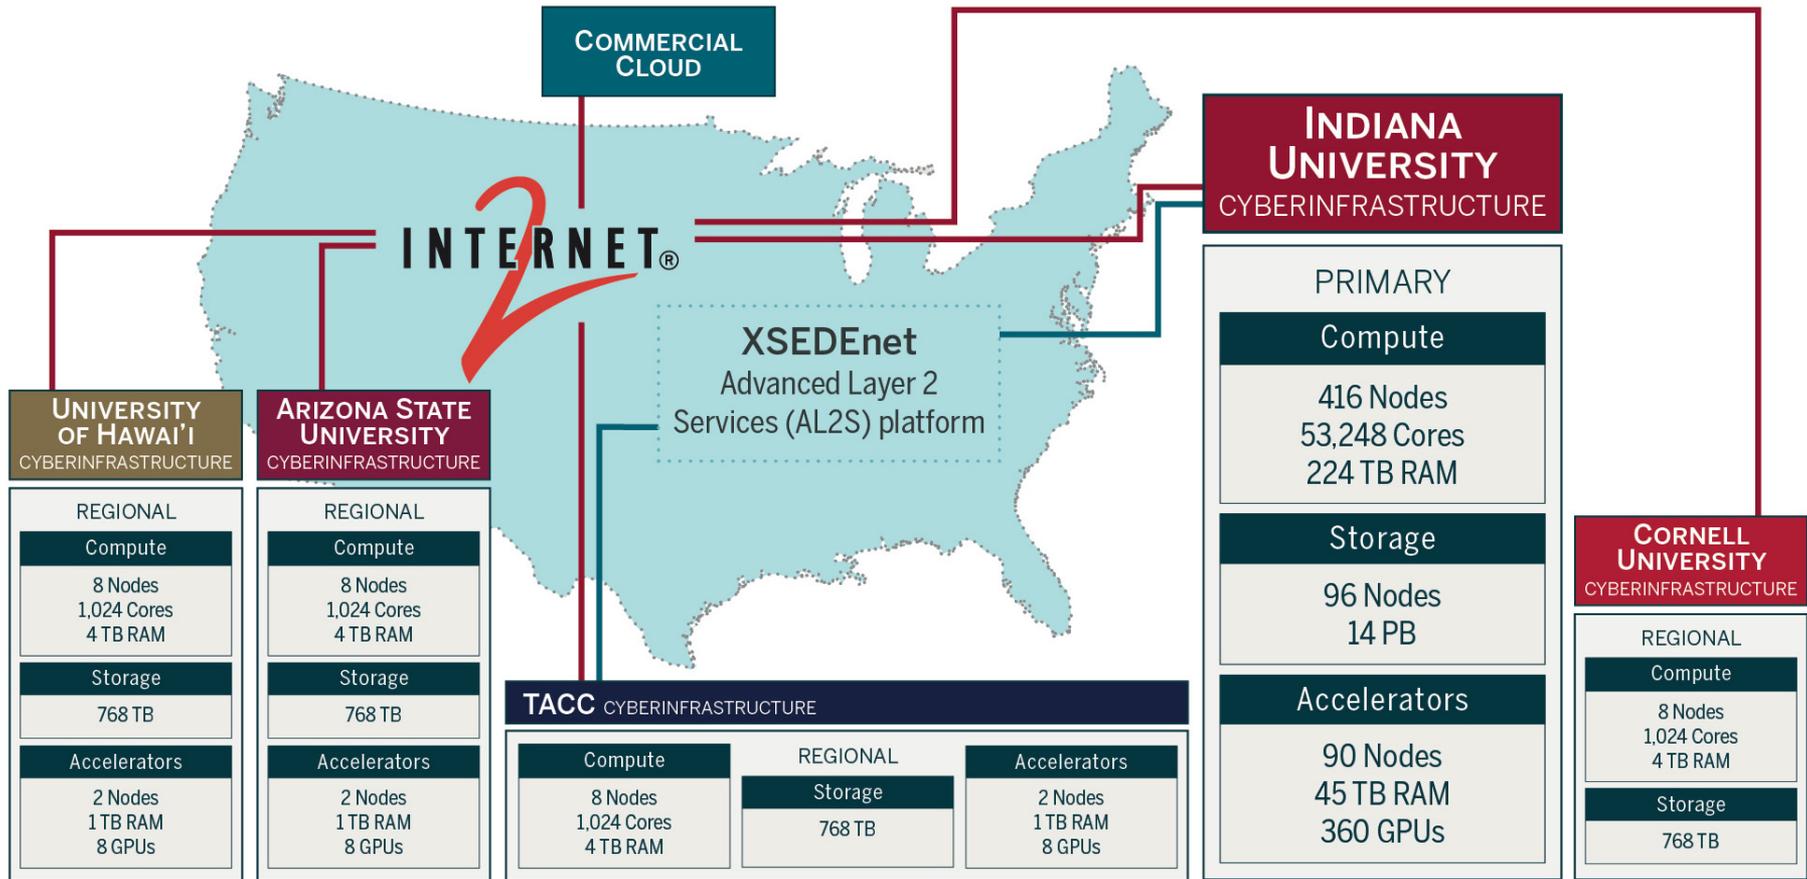
Revamped User Interface

- Unified instance management
- Multi-instance launch

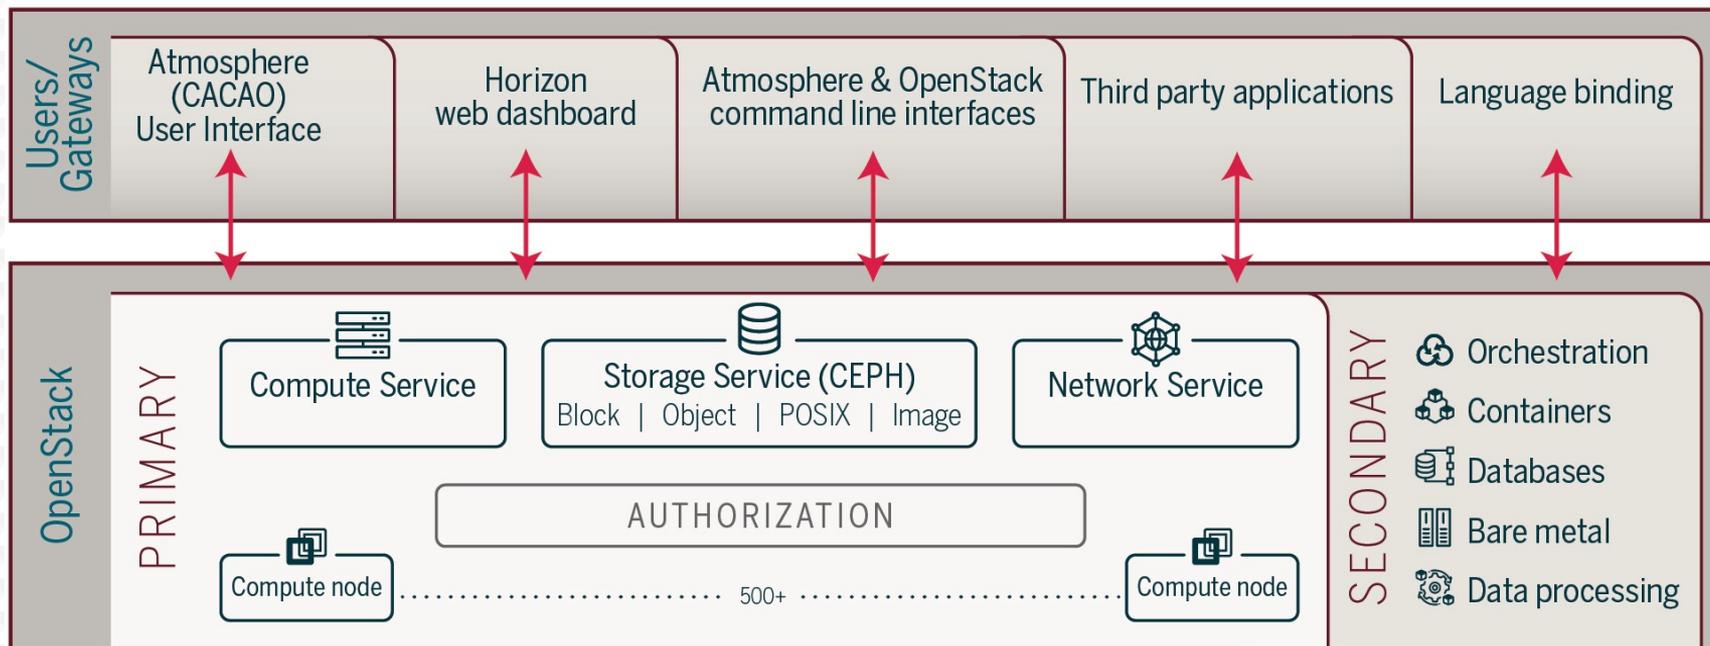


Feb 12, 2019 – Jet stream region called “Jet N6”  
NASA/JPL-Caltech/SwRI/MSSS/Kevin M. Gill

- >57K cores of next-gen AMD EPYC processors
- >360 NVIDIA A100 GPUs will provide vGPUs via NVIDIA's MIG feature
- >17PB of storage (NVMe and disk hybrid)
- 100GbE Mellanox network

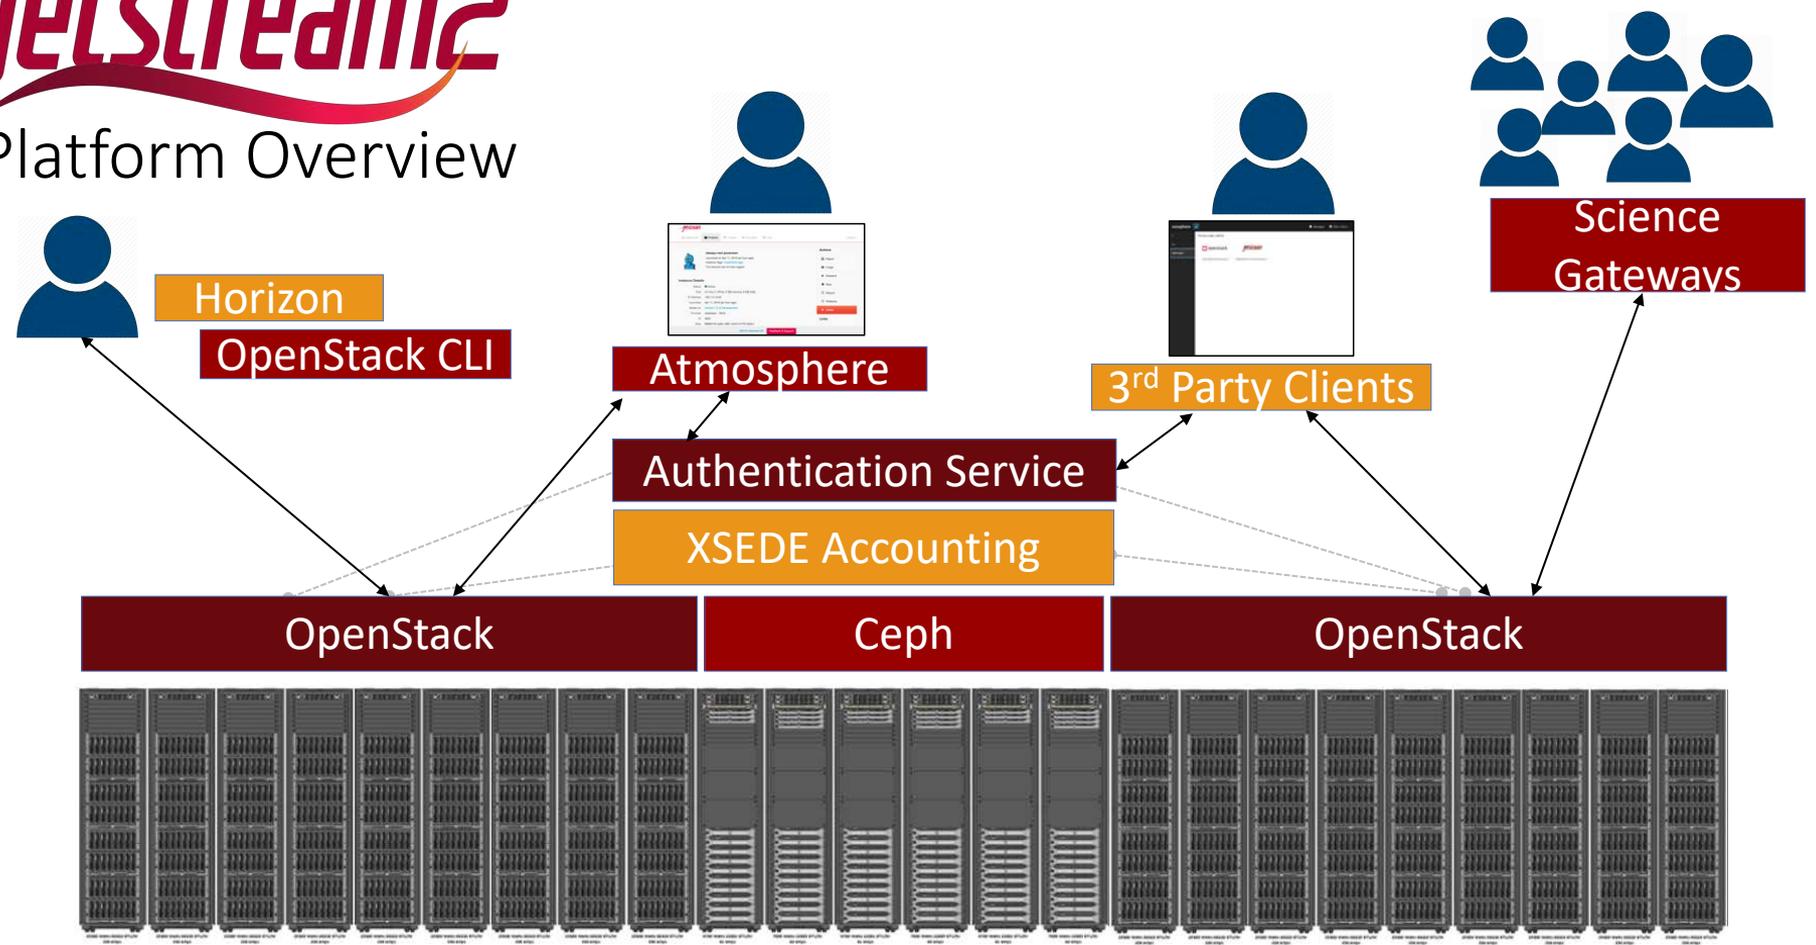


# Conceptual Jetstream2 Architecture



# Jetstream2

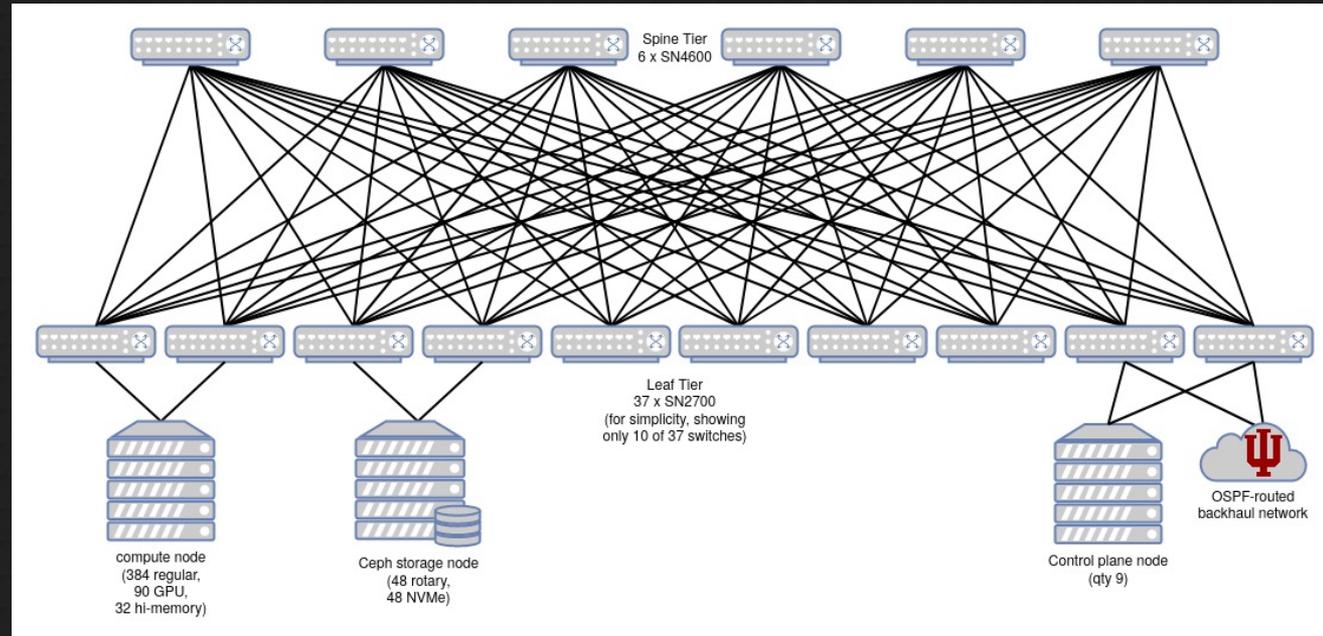
## Platform Overview



Jetstream2 Core System

# Networking in Jetstream 2

- Layer 3 fabric
- Routing on the Host
- Floating IP addresses happen on the compute nodes



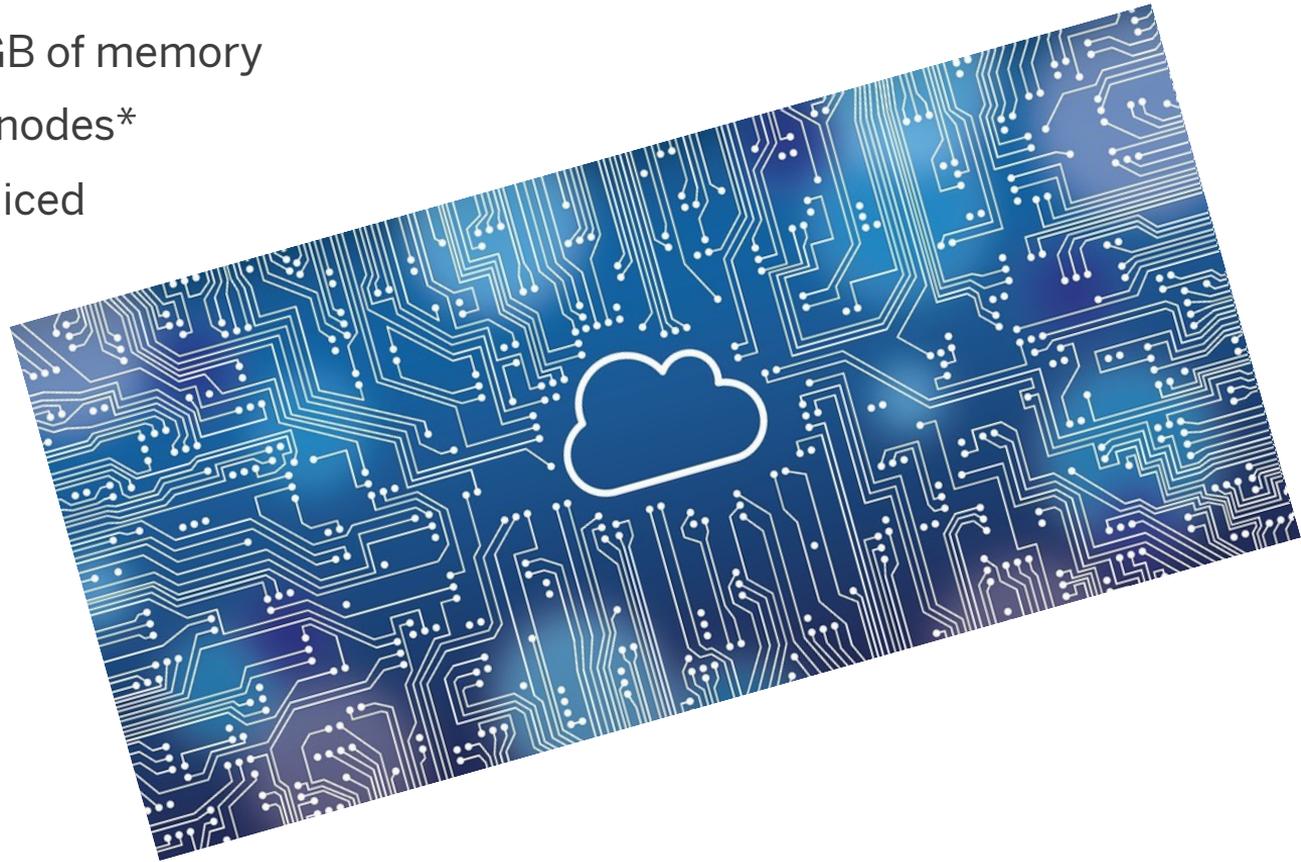


## Ceph and CephFS

- OpenStack Manila managed CephFS shares
- Mountable Anywhere\*
- Erasure Coding
- 3.6 PB of NVME
- 9 PB of spinning disk

# Big Memory, Larger Instances, GPUs

- Smallest node has 512GB of memory
- 32 Larger 1TB memory nodes\*
- A100 GPUs sliced and diced
- 128 Core nodes



# Exosphere

Jetstream

Messages Settings Get Support About Logou

Home > Project TG-CCR190024

iu.jetstream-cloud.org - TG-CCR190024

Remove Allocation

Create

## Instances

Instances used 11 of 25 total Cores used 26 of 132 total RAM used 100 of 388 GB

Select All

Ready formally\_trusty\_urchin

Shelved optionally\_certain\_longhorn with GUI

Ready wildly\_united\_mite

Hiding 8 instances created by other users

Show

## Volumes

Volumes used 2 of 10 total Storage used 279 of 1,100 GB

https://app.exosphere.localhost:8000/projects/ Exosphere 80%

Exosphere

Messages Settings Get Support About

iu

tacc

Add Project

Choose a root disk size

- 20 GB (default for selected size)
- Custom disk size (volume-backed)

How many Instances?

Your resource limits supports up to 15 of these.

1

Enable graphical desktop?

- No
- Yes

Any Binder™-compatible repository can be launched. See mybinder.org for more information

Launch a workflow in the instance

Git repository URL or DOI

https://github.com/binder-examples/minimal-dockerfile

Git ref (branch, tag, or commit) (optional)

HEAD

URL to open e.g. /rstudio (optional)

URL to open e.g. /rstudio (optional)

Remove workflow

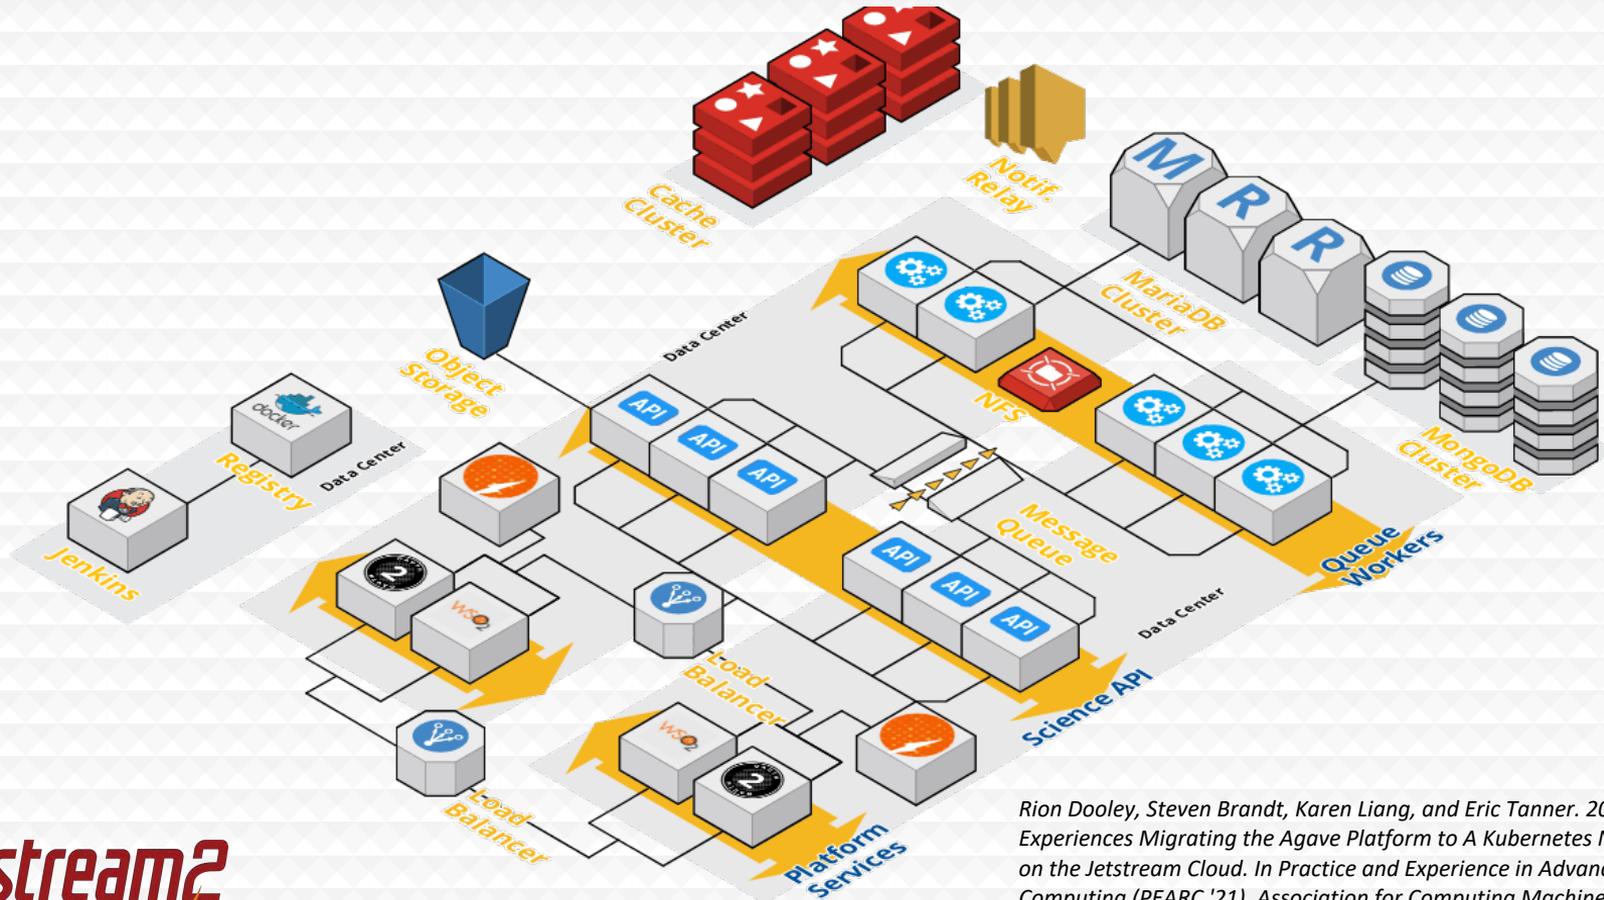
Advanced Options

- Hide
- Show

Jetstream2

<https://exosphere.Jetstream-cloud.org> or [try.exosphere.app](https://try.exosphere.app)

# Agave: Science-as-a-Service (SaaS) platform



Rion Dooley, Steven Brandt, Karen Liang, and Eric Tanner. 2021. Experiences Migrating the Agave Platform to A Kubernetes Native System on the Jetstream Cloud. In Practice and Experience in Advanced Research Computing (PEARC '21). Association for Computing Machinery, New York, NY, USA, Article 39, 1–4. DOI:<https://doi.org/10.1145/3437359.3465598>

# Jetstream for education – in action at AMS2020

- Unidata-led workshop at American Meteorological Society (AMS) 2020 conference
- 127 users actively participating
- Participants used a JupyterHub running on Jetstream (40 node Kubernetes cluster of 6 core m1.medium VMs) for a 90 minute Unidata PyAOS (Python for the Atmospheric and Oceanic Sciences) workshop
- The students were successfully able to run their interactive Python code notebooks as the instructors presented their material



# Scalable Galaxy Workloads Using Virtual Clusters

- Usegalaxy.org maintains elastic virtual clusters on Jetstream
  - Minimal footprint (headnode only) when not actively processing jobs
- **39,584** distinct users have executed jobs on Jetstream via usegalaxy.org (for period from 2016-01-01 through 2021-03-18)
  - **13,570** users in year 2020



# Timeline

- Jetstream now in 5th year of operations
- Jetstream extension granted by the NSF through November 2022
- Extending operations through March 2022
- Jetstream2
  - Early operations planned for December 2021
  - Production operations by January 2022



Flickr user Oiluj Samall Zeid - Lejos de Yulín



**PERVASIVE  
TECHNOLOGY INSTITUTE**



**RESEARCH TECHNOLOGIES**  
UNIVERSITY INFORMATION TECHNOLOGY SERVICES

# Acknowledgements

NSF Awards 1053575 & 1548562 (XSEDE), 1445604 (Jetstream) and 2005506 (Jetstream2)

This document was developed with support from the National Science Foundation. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF.

Special thanks to contributors & Jetstream2 partners

- Jeremy Fischer, J. Michael Lowe, Therese Miller, Maria Morris, Winona Snapp-Childs, George Turner, and Chris Martin.
- Vendors, particularly Dell and NVIDIA, also deserve recognition for their efforts



**PERVASIVE  
TECHNOLOGY INSTITUTE**



**RESEARCH TECHNOLOGIES**  
UNIVERSITY INFORMATION TECHNOLOGY SERVICES

# Jetstream2 partners



**JOHNS HOPKINS**  
UNIVERSITY



**UCAR**



<http://jetstream-cloud.org/>  
National Science Foundation  
Award #ACI-2005506