



**PERVASIVE
TECHNOLOGY INSTITUTE**



RESEARCH TECHNOLOGIES
UNIVERSITY INFORMATION TECHNOLOGY SERVICES



**PERVASIVE
TECHNOLOGY INSTITUTE**



RESEARCH TECHNOLOGIES
UNIVERSITY INFORMATION TECHNOLOGY SERVICES

Jetstream2 as part of the NSF Ecosystem

David Y. Hancock – Indiana University

Director for Advanced Cyberinfrastructure

Jetstream & Jetstream2 Primary Investigator

PEARC21 Panel – July 22, 2021

The NSF Computing Ecosystem: Category 1 and 2 Resources
for Accelerating Research for the Community



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>



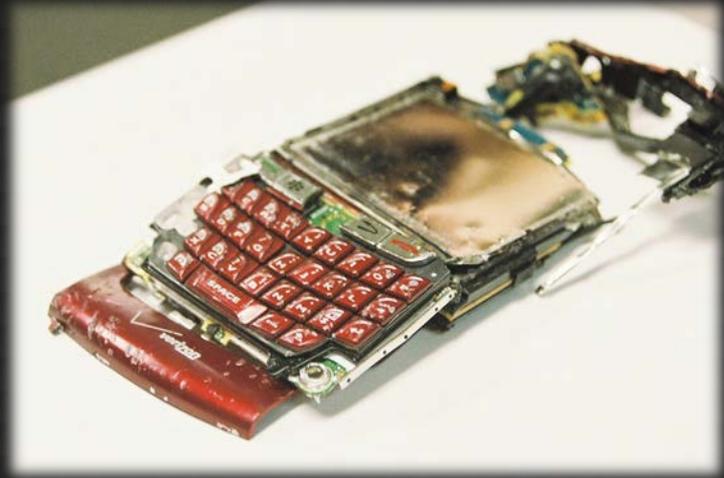
Are we there yet?



Jetstream2

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

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

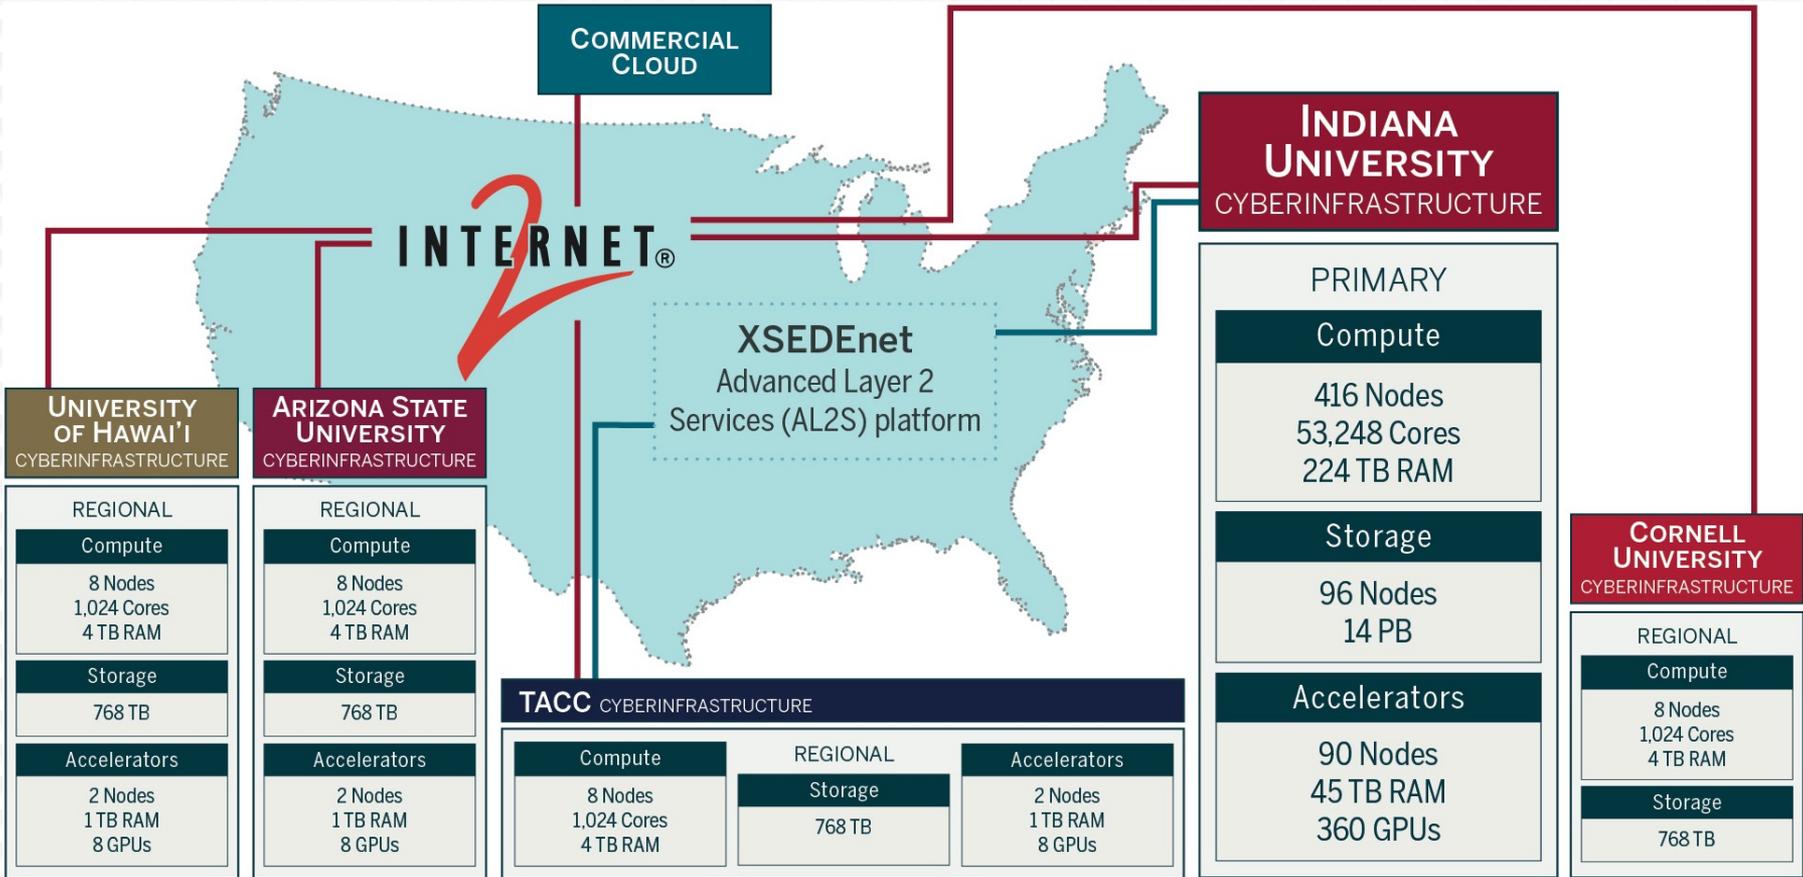
Revamped User Interface

- Unified instance management
- Multi-instance launch

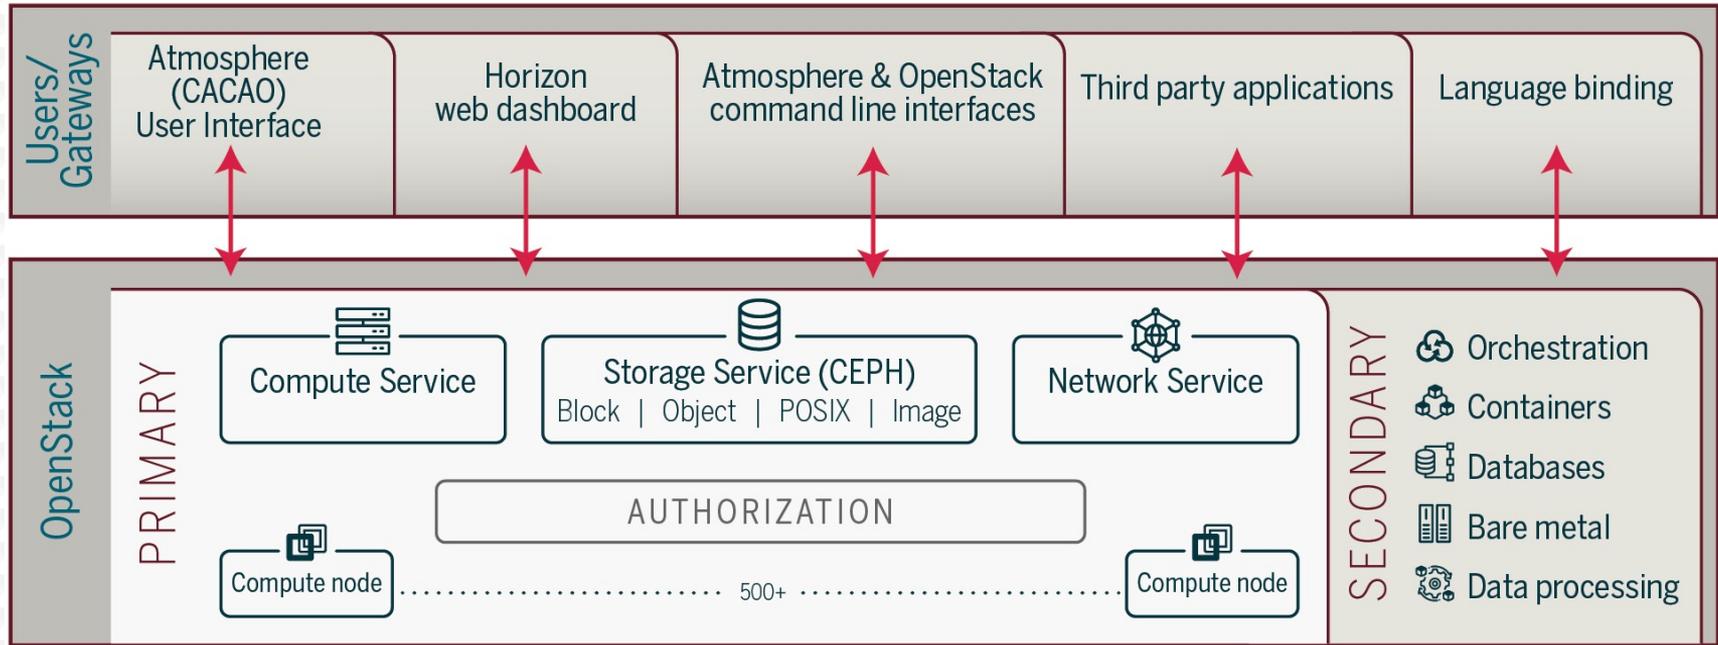


Feb 12, 2019 – Jet stream region called “Jet N6”
NASA/JPL-Caltech/SwRI/MSSS/Kevin M. Gill

- >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

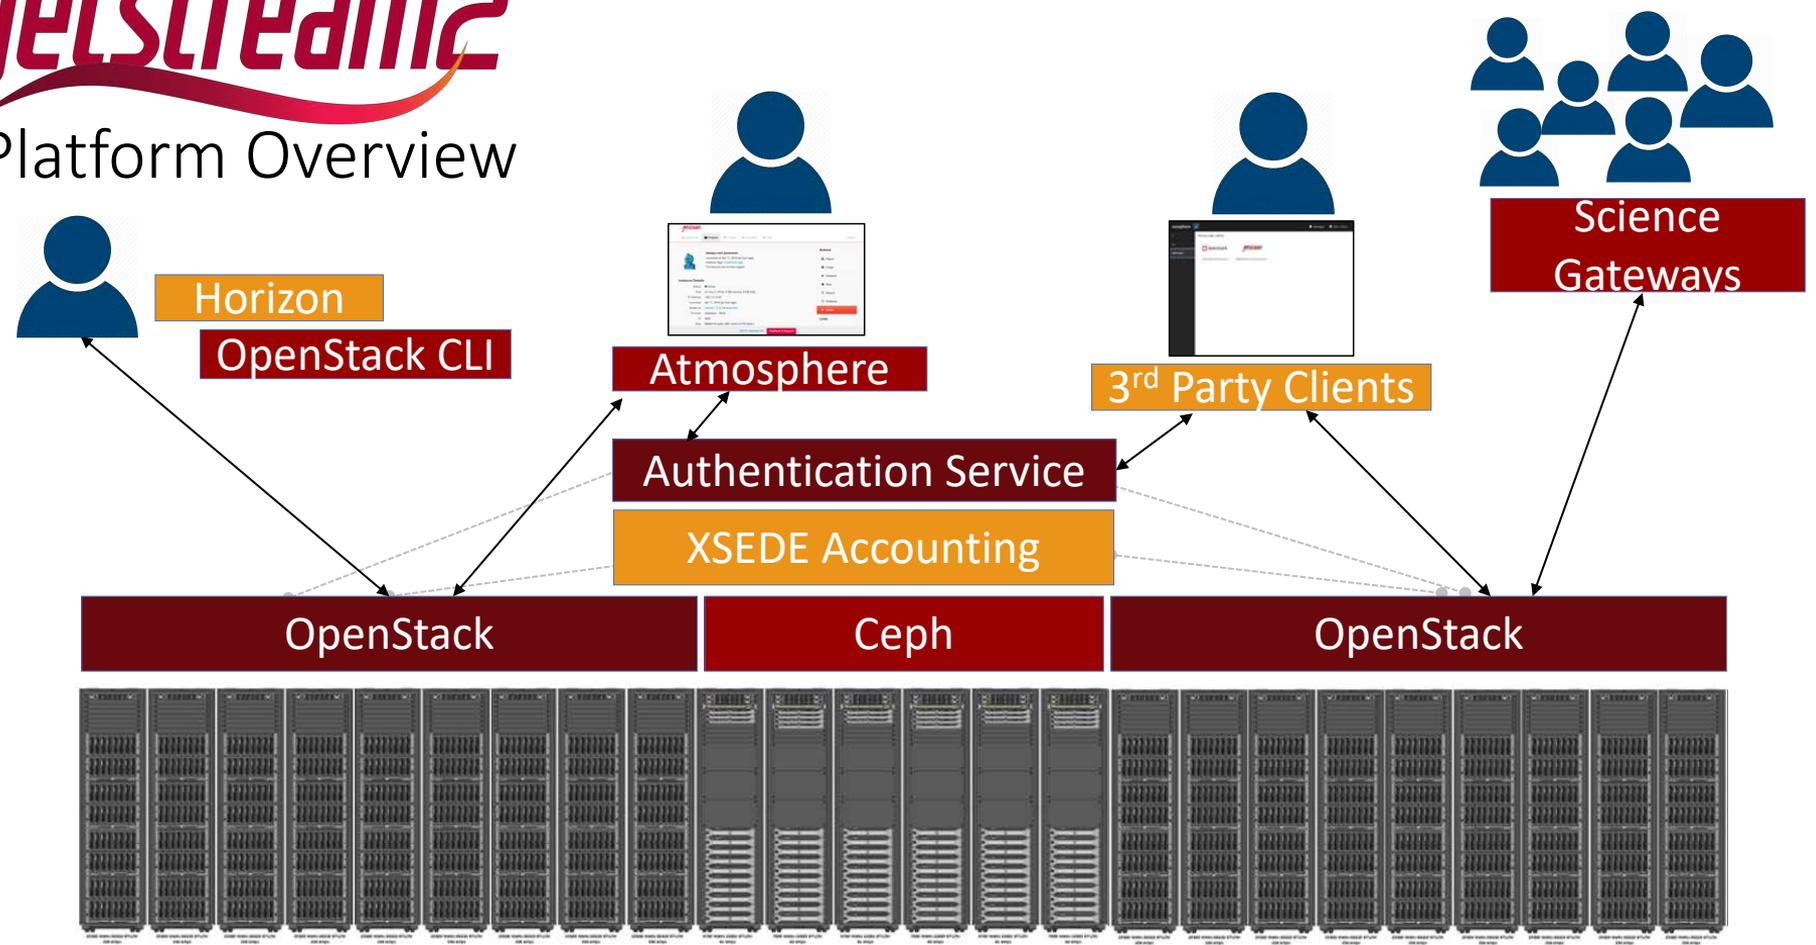


Conceptual Jetstream2 Architecture



Jetstream2

Platform Overview



Jetstream2 Core System

Timeline

- Jetstream now in 5th year of operations
- Jetstream extension granted by the NSF through November 2021
- Extension through March 2022 in process
- 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, and George Turner



**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



**PERVASIVE
TECHNOLOGY INSTITUTE**



RESEARCH TECHNOLOGIES
UNIVERSITY INFORMATION TECHNOLOGY SERVICES