IndySCC 2025 on Jetstream2

Zachary Graber
Research Cloud Infrastructure
UITS Research Technologies, Indiana University

PERVASIVE TECHNOLOGY INSTITUTE
RESEARCH TECHNOLOGIES
UNIVERSITY INFORMATION TECHNOLOGY SERVICES

What is Jetstream2?

- Jetstream2 is a **cloud computing environment** funded by the U.S. National Science Foundation (NSF)
  - Allocated at zero cost to US-based researchers and educators through ACCESS
  - ~67,000 CPU cores, ~300 TiB RAM, 488 GPUs, ~14 PB raw storage available through “instances” (virtual machines)
    - IndySCC teams will be using AMD Milan 7713, NVIDIA A100 40 GB
  - 100 Gbps Ethernet from compute nodes to leaf switches
- Jetstream2 is **not** a traditional HPC cluster
  - No resource placement guarantees
  - No ultra-high-performance network fabric (e.g. InfiniBand)

Jetstream2

PERVASIVE TECHNOLOGY INSTITUTE

RESEARCH TECHNOLOGIES
UNIVERSITY INFORMATION TECHNOLOGY SERVICES

Who operates Jetstream2?

PERVASIVE TECHNOLOGY INSTITUTE

RESEARCH TECHNOLOGIES
UNIVERSITY INFORMATION TECHNOLOGY SERVICES

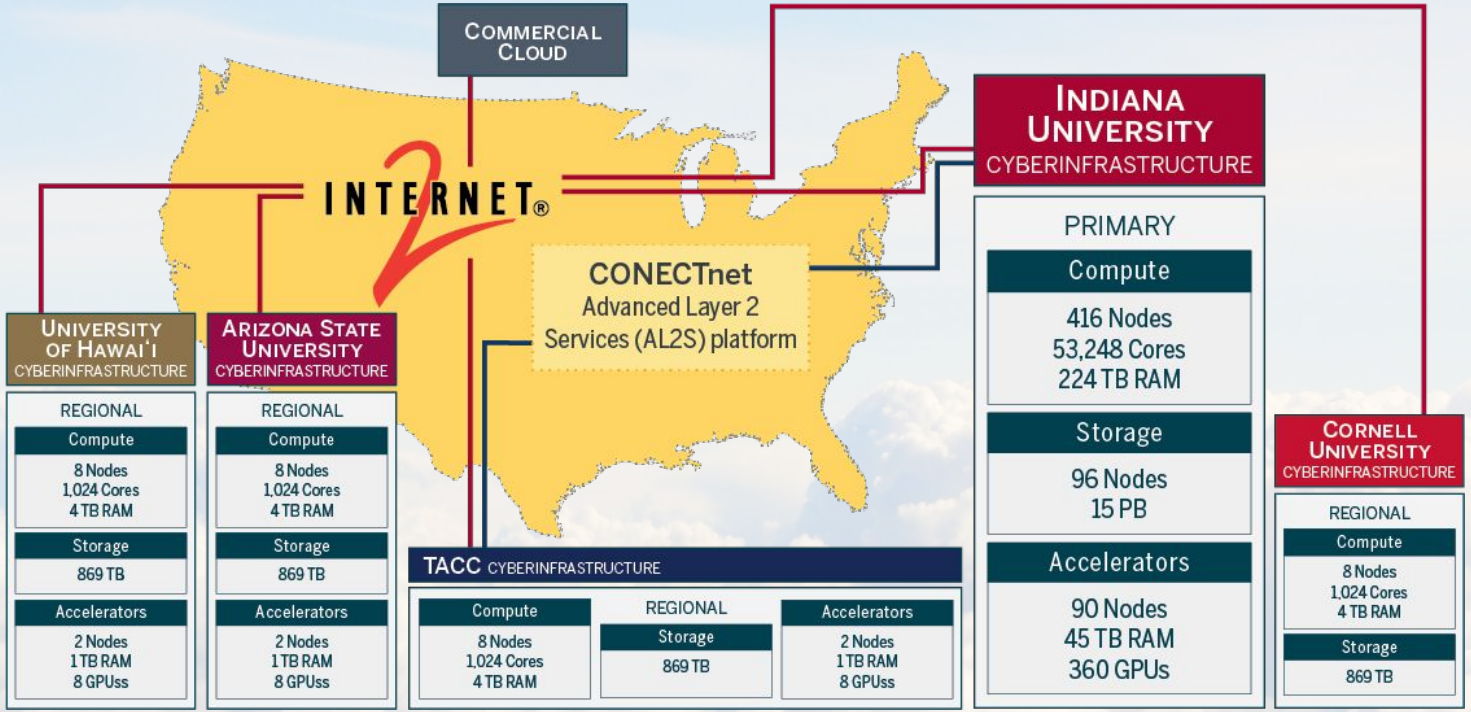

Arizona State University · The University of Arizona · Cornell University · University of Hawai'i · Indiana University · Johns Hopkins University · TACC · UCAR

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

What do IndySCC teams get on Jetstream2?

- **Three phases**
  - Exercise/prep
  - “Hero” benchmarking runs
  - Final 48-hour competition
- **Exercise/prep phase:**

Instance	Flavor	Count	# CPU	RAM (GB)	Default Storage	Other
Login/Head	m3.quad	1x	4	15 GB	20 GB	
CPU Compute	m3.2xl	3x	64	250 GB	60 GB	
GPU Compute	g3.large	1x	16	60 GB	60 GB	50% of an NVIDIA A100 40G

Jetstream2

PERVASIVE TECHNOLOGY INSTITUTE

RESEARCH TECHNOLOGIES
UNIVERSITY INFORMATION TECHNOLOGY SERVICES

What do IndySCC teams get on Jetstream2? (cont.)

- **“Hero” benchmark runs (24 hours):**
  - Plan is ~400-600 CPU cores per team
  - Exact instance flavors and counts TBD

- **Final competition:**
  - These counts are not final and may change in accordance with availability, etc.

Instance	Flavor	Count	# CPU	RAM (GB)	Storage	Other
CPU Compute	TBD	2-3x	128-192	TBD	60 GB	
GPU Compute	g3.xl	1x	32	120 GB	60 GB	1x **full** NVIDIA A100 40G

Jetstream2

PERVASIVE TECHNOLOGY INSTITUTE

RESEARCH TECHNOLOGIES
UNIVERSITY INFORMATION TECHNOLOGY SERVICES

Note: GPU placement on Jetstream2

Example: Prep phase (g3.large)

Example: Competition phase (g3.xl)

What storage options do teams have?

- **Default quota is 1 TB**
  - Teams can request more (up to 5 TB) via team liaison / Google group
- **All** storage on Jetstream2 is networked–even instance root disks!
  - Backed by NVMe on our distributed Ceph cluster
  - Unlikely to see throughput gains by striping across instances (e.g. BeeGFS)
- **Volumes** are single-attach, raw block storage devices
  - Choose whatever filesystem type you prefer (ext4, xfs, etc.)
  - Can only be mounted to a single instance at a time–can be shared externally (e.g. via NFS)
- **Manila Shares** are ready-to-use CephFS network filesystems
  - Can be mounted to many instances at a time–good for shared storage
- See https://docs.jetstream-cloud.org/general/storage/

SU consumption on Jetstream2

- Cloud compute billing typically based on **power state** of instance, not activity.
  - Forget about a running instance on AWS EC2? 💸
  - Active Jetstream2 instances consume SUs (service units) from your allocation **even if you're not logged in or running anything**
- On Jetstream2, users **shelve** instances to stop SU consumption when they're done working
  - Shelving removes instance from physical compute host (hypervisor)
- **Teams** are responsible for managing their allocations
  - Run out of SUs? You're out!
  - SU balance in Exosphere is only updated once every 24 hours!
  - Zach will demo how to monitor SU consumption later

How do teams install software and run jobs?

- Users have **full root (sudo) access** to their instances (nodes).
- No central scheduler to submit to
  - You can install/configure one (e.g. Slurm)!
  - OR just distribute manually with mpirun, GNU Parallel, etc.
- When installing/configuring, keep scalability in mind
- Partial GPU flavors (g3.large, etc.) require a special NVIDIA GRID driver (preinstalled on officially supported Ubuntu/Rocky OS images).
  - **Warning:** do not uninstall the `nvidia-linux-grid` deb/rpm package or install a different driver on a g3.large/g3.medium

What will Jetstream2 staff & the committee support?

- **Operating systems**
  - We can only guarantee support for official Jetstream2 OS images (Ubuntu 24/22 or Rocky Linux 9/8)
  - Advanced teams can provide any generic cloud compatible (e.g. qcow2) Linux OS image, but support isn't guaranteed
- **Filesystems**
  - Common POSIX filesystems such as ext4, xfs, or btrfs are recommended
  - The committee won't provide support for "exotic" installations (e.g. BeeGFS)

Jetstream2

PERVASIVE TECHNOLOGY INSTITUTE

RESEARCH TECHNOLOGIES
UNIVERSITY INFORMATION TECHNOLOGY SERVICES

Performance considerations

- **Networking**
  - Jetstream2 is not optimized for low-latency inter-node communication
- **CPU and memory**
  - Cores are not oversubscribed–all your vCPUs are yours alone
  - Cores may not be consecutively placed
  - Memory is statically partitioned, not ballooned
- **GPUs**
  - vGPU flavors (e.g. g3.large) opportunistically use empty cycles on the card
  - VRAM is statically partitioned

Demo

https://jetstream2.exosphere.app

Jetstream2

PERVASIVE TECHNOLOGY INSTITUTE

RESEARCH TECHNOLOGIES
UNIVERSITY INFORMATION TECHNOLOGY SERVICES

Where should teams get help with Jetstream2 resources?

- Ask in the **Google Group** !
  - Please do not contact IndySCC committee members directly
  - Please do not contact Jetstream2 support directly
- https://docs.jetstream-cloud.org/

Jetstream2

PERVASIVE TECHNOLOGY INSTITUTE

RESEARCH TECHNOLOGIES UNIVERSITY INFORMATION TECHNOLOGY SERVICES

Questions ?

JETSTREAM2 BENEFITS/IMPACT

- **Full sudo/admin access** to choose your operating system and install software you need
- **Broad range of resource "flavors"** including Large Memory and NVIDIA vGPU slicing
  - Easy and equitable access to GPUs helps spearhead AI/ML research nationally
- **Available on demand** without sharing or queues, and no runtime limit
- **Free to use** through our support from the National Science Foundation and ACCESS
  - "ACCESS allocations are available to any researcher or educator at a U.S. academic, non-profit research, or educational institution."
  - Available at any level from community college to R1 university

NSF ACCESS Advancing Innovation

Jetstream2

PERVASIVE TECHNOLOGY INSTITUTE

RESEARCH TECHNOLOGIES UNIVERSITY INFORMATION TECHNOLOGY SERVICES