

*The CI Landscape:  
Systems, Service Providers,  
Technologies*



*Jacob Fosso Tande,*

The University of North Carolina at Greensboro,

Campus Champions Leadership Team,

Virtual Residency

June 28, 2022



# Talk *OUTLINE*



*NSF/XSEDE → ACCESS  
Systems: Anvil, Bridges2,  
Expanse, Jetstream2,  
Frontera  
Open Science Grid (OSG)*

*Science Gateways*

*DOE National Labs:  
OLCF's Summit &  
Frontier and NERSC's  
Perlmutter & Cori*


*ESnet and EPOC*

*Internet2*

*International Orgs:  
GEANT, Compute  
Canada, PRACE*

*Tools: Globus, XMod,  
Spack, Open OnDemand,  
ask.ci, etc*

# Striking the Right Balance



Services and Systems are omitted because of:

1. Expanded landscape
2. Time
3. Where to find?

## The following are omitted:

- NSF/ACCESS: Darwin, Rockfish, Open Storage Network(OSN), Stampede2, FASTER
- DOE National Lab resources such as the ALCF's Aurora & Polaris test bed.
- NIH resources such as BioWulf, Helix and Helixweb
- Cloud resources such as CloudBank, CloudLab, GENI, emulab and Chameleon.
- Commercial Cloud: Google Cloud, Azure, AWS, IBM cloud etc..
- Regional Networks and THE QUILT

## *Why Bother?*

- You are from a resource limited institution.
- Your users have needs that are not present at your institution.
  - Extreme memory compute systems
  - Testbeds
  - Distributed data sharing (Open Storage Network)
- Demand surpasses offer
- Highly collaborative research projects:
  - IceCube project
  - Event Horizon Telescope Collaboration

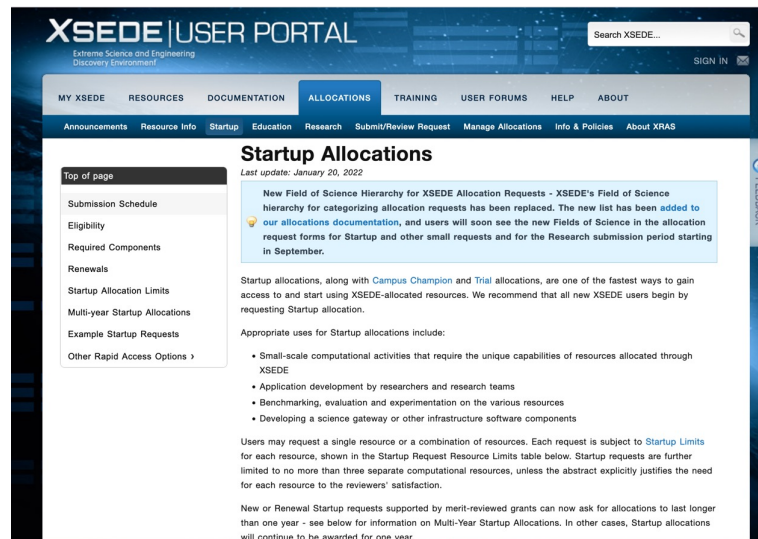
What is more important? Knowing all about the systems and services or knowing where to find these systems and services?

# *NSF/XSEDE → ACCESS Systems*

## Startup Allocations

### Why get a Startup Allocation?

- Small-scale computational activities that require the unique capabilities of resources allocated through XSEDE
- Application development by researchers and research teams
- Benchmarking, evaluation and experimentation.
- Developing a science gateway or other infrastructure software components



## *NSF/XSEDE → ACCESS Systems*

### Startup Allocations

***Share your thoughts in the chat: If you have/had a startup allocation, how do/would you use it?***

## *Anvil HPC*



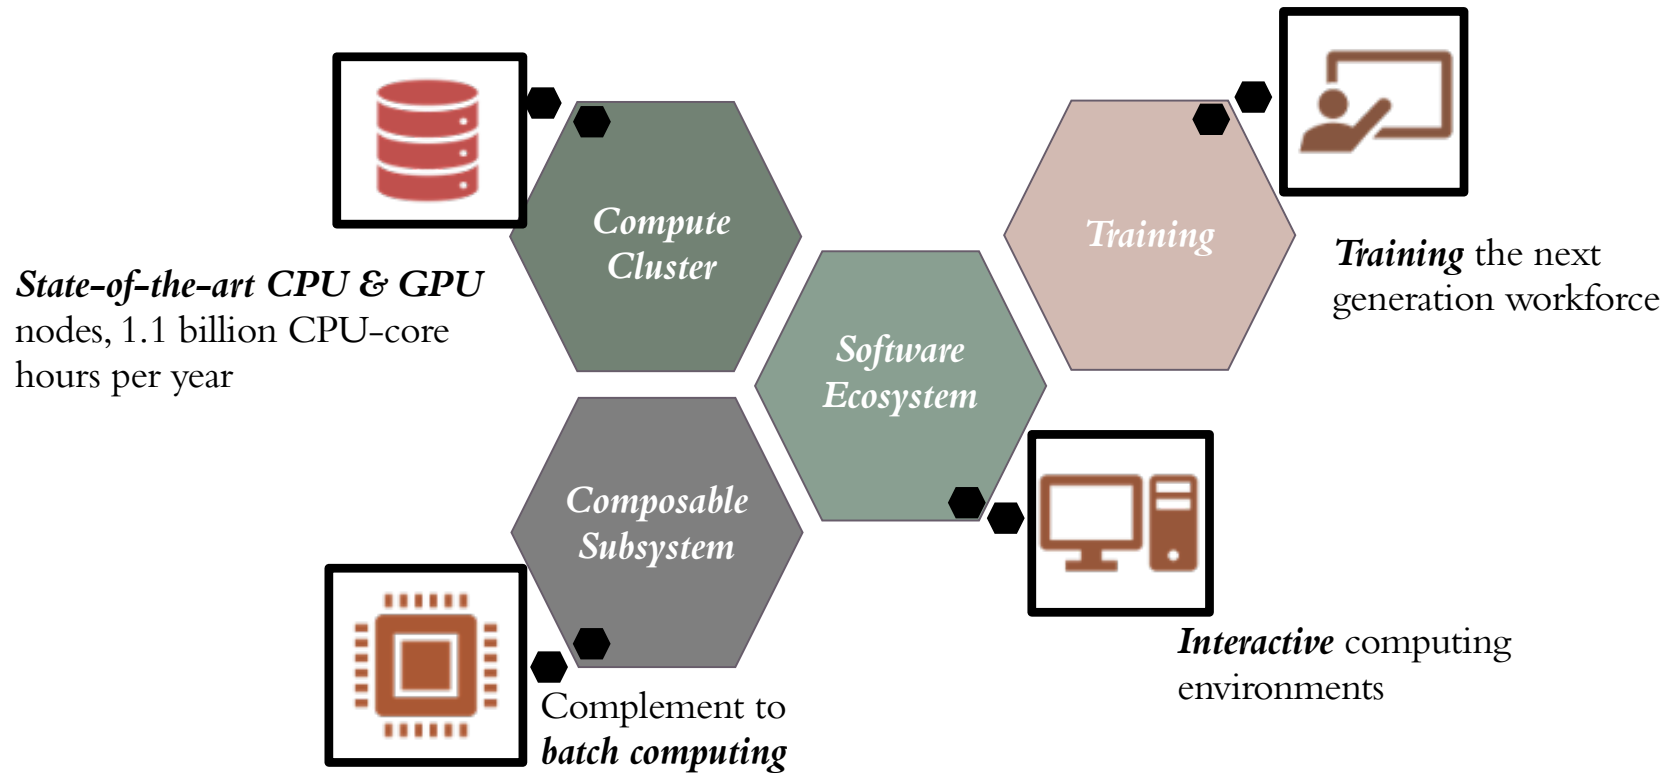
“ ... reflects the Purdue Boilermakers' strength and workmanlike focus on producing results ...”

*Category I*: A National Composable  
Advanced Computational Resource for the Future  
of Science and Engineering

5-year Production Operations

Started Feb 2022

# Anvil's Ecosystem





# The Anvil System

- **1000 Compute Nodes**
- 128 core 3rd Gen AMD EPYC™ processors
- 5.1 PF peak performance

*High -  
performance*

- 16 nodes with 4 NVIDIA **A100 GPUs** each
- 32 large memory nodes with 1 TB of RAM

*GPU/Large-  
memory*

- **8 large memory & storage nodes**
- Kubernetes – Rancher for SciOps
- Cloud bursting to Azure

*Composable  
subsystem*

*Storage*

- **Multi-tier storage** (including archival & object storage)
- 10 PB of parallel file system, and 3 PB of all-flash storage
- Globus data transfer

# *The Anvil Software Stack*

## *Programming Libraries and Compilers*

- Various popular programming languages, GNU, Intel and AOCC compilers, message passing libraries
- Workflow and data management and analysis tools
- Debugging and profiling tools

## *Scientific Applications*

- General purpose mathematics and statistics modeling tools, numerical libraries, visualization tools
- Broad application base with installs and modules from various science and engineering domains

## *Containers and Datasets*

- Support for Singularity containerization and execution
- Efficient access to various databases (e.g., NCBI)

# Bridges-2 HPC



Provides transformative capability for rapidly evolving, computation-intensive and data-intensive research, creating opportunities for collaboration and convergence research.

*Connecting new communities to advanced research computing.*

***More Science:***

Approximately 3x larger than Bridges

***Faster Computing:***

Latest AMD EPYC processors

***Faster Storage:***

Fast flash array and tiered data management

***Smarter Science:***

Designed for Full System AI and data-centric computing

***Scalable:***

Interoperability with cloud and campus resources

“PSC’s newest supercomputer”

**Carnegie Mellon University**



**University of  
Pittsburgh**

# *Bridge-2's Concepts*

## *Bridges-2 core concepts:*

- Converged HPC + AI + Data
- Custom fat tree Clos topology optimized for data-centric HPC, AI, and HPDA
- Heterogeneous node types for different aspects of workflows
- CPUs and AI-targeted GPUs
- 3 tiers of per-node RAM (256GB, 512GB, 4TB)
- Extremely flexible software environment
- Community data collections & Big Data as a Service

## *Innovations beyond Bridges:*

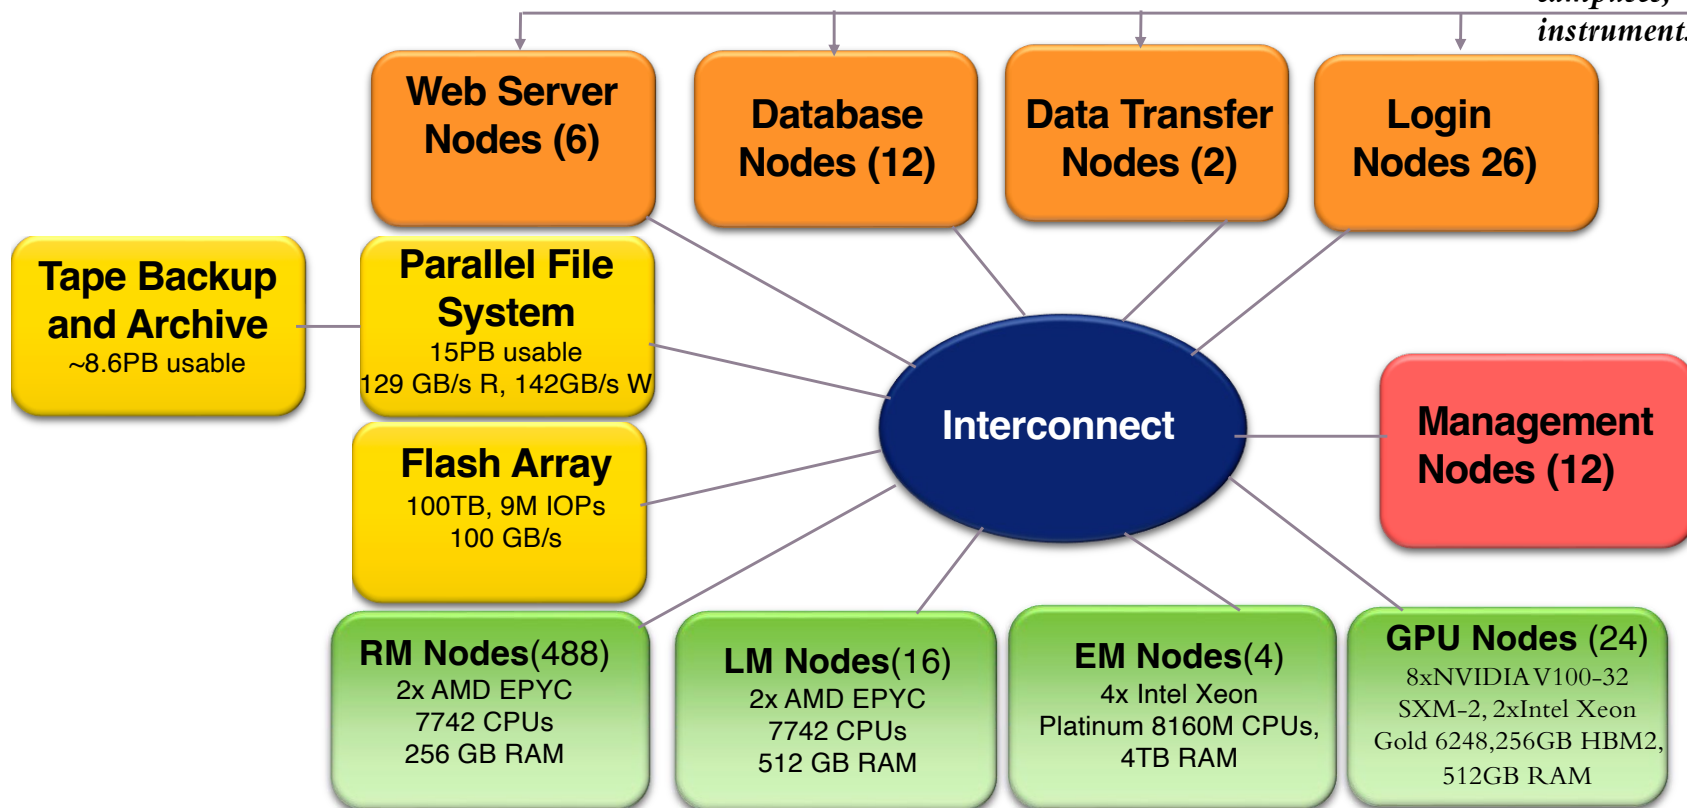
- AMD EPYC 7742 CPUs: 64-core, 2.25–3.4 GHz
- AI scaling to 192 V100-32GB SXM2 GPUs
- 100TB, 9M IOPs flash array accelerates deep learning training, genomics, and other applications
- Mellanox HDR-200 InfiniBand doubles bandwidth & supports in-network MPI-Direct, RDMA, GPU Direct, SR-IOV, and data encryption
- Cray ClusterStor E1000 Storage System
- HPE DMF single namespace across disk and tape for data security and expandable archiving

Carnegie Mellon University



# Bridge-2's Architecture

*Users, XSEDE,  
campuses,  
instruments*



## Using Bridge-2

- Interactivity
- Popular languages and frameworks: Python, Anaconda, R, MATLAB, Java, Spark, Hadoop
- AI frameworks: TensorFlow, Caffe2, PyTorch, etc.
- Containers and virtual machines (VMs)
- Databases
- Gateways and distributed (web) services
- Large collection of applications and libraries



Carnegie Mellon University

PSC

University of Pittsburgh

# EXPANSE

COMPUTING WITHOUT BOUNDARIES  
5 PETAFL0P/S HPC and DATA RESOURCE

## HPC RESOURCE

13 Scalable Compute Units  
728 Standard Compute Nodes  
52 GPU Nodes: 208 GPUs  
4 Large Memory Nodes

## LONG-TAIL SCIENCE

Multi-Messenger Astronomy  
Genomics  
Earth Science  
Social Science


## DATA CENTRIC ARCHITECTURE

12PB Perf. Storage: 140GB/s, 200k IOPS  
Fast I/O Node-Local NVMe Storage  
7PB Ceph Object Storage  
High-Performance R&E Networking

## REMOTE CI INTEGRATION

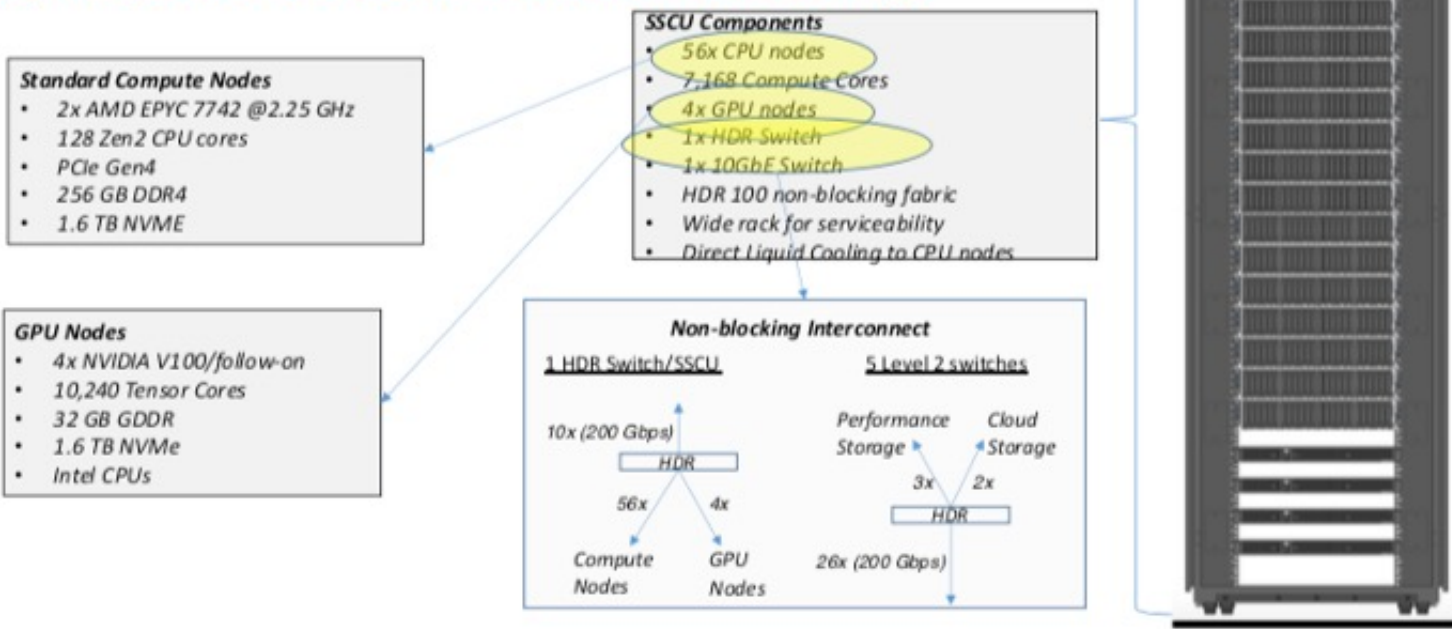
## INNOVATIVE OPERATIONS

Composable Systems  
High-Throughput Computing  
Science Gateways  
Interactive Computing  
Containerized Computing  
Cloud Bursting



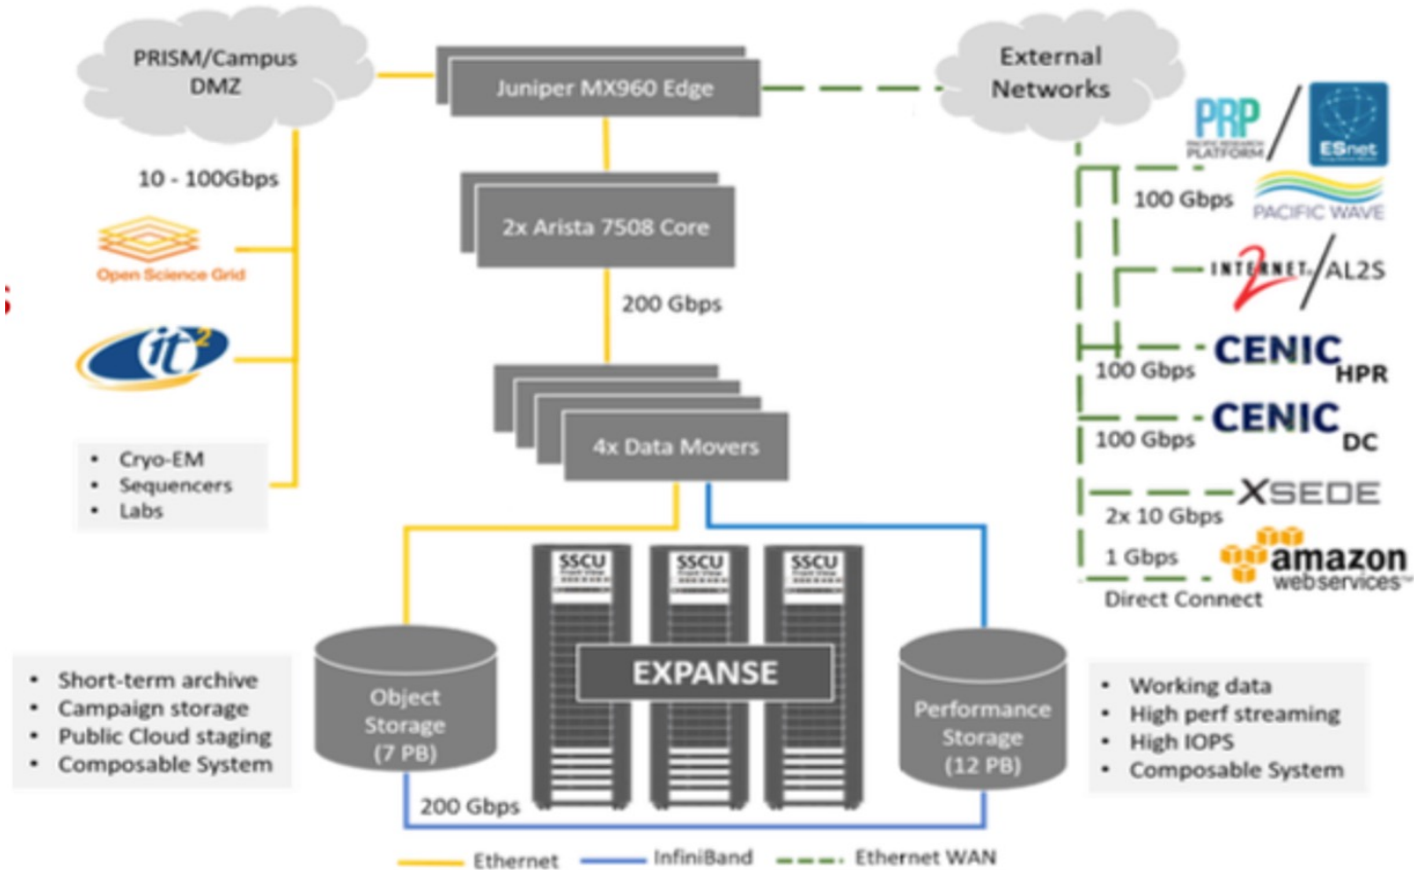
# Expanse Scalable Compute Unit

The SSCU is Designed for the Long Tail Job Mix, Maximum Performance, Efficient Systems Support, and Efficient Power and Cooling

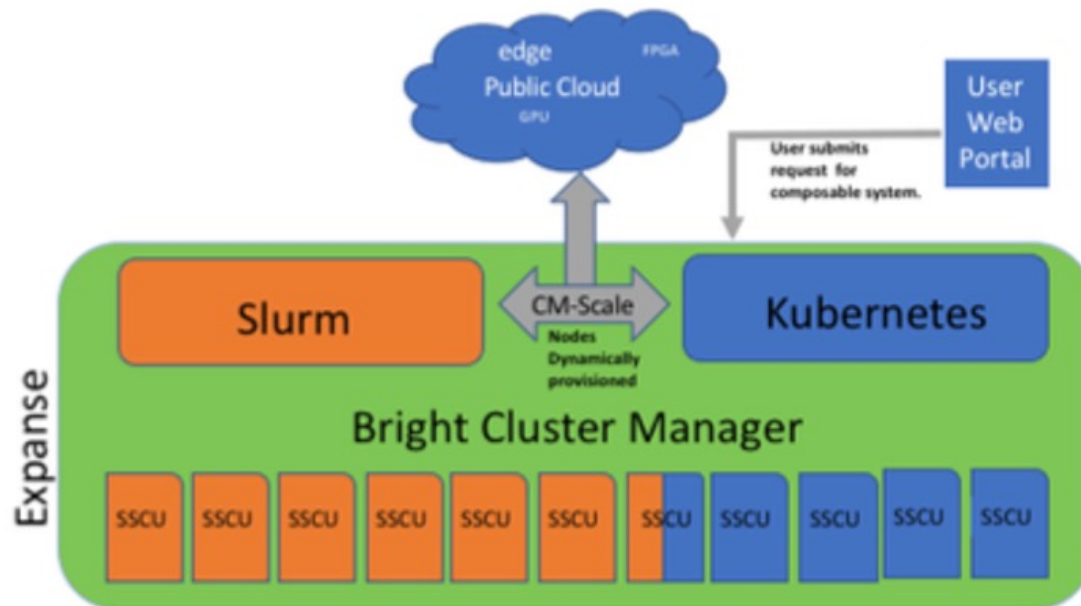




# Expanse Connectivity Fabric



# Expanse Composable Systems

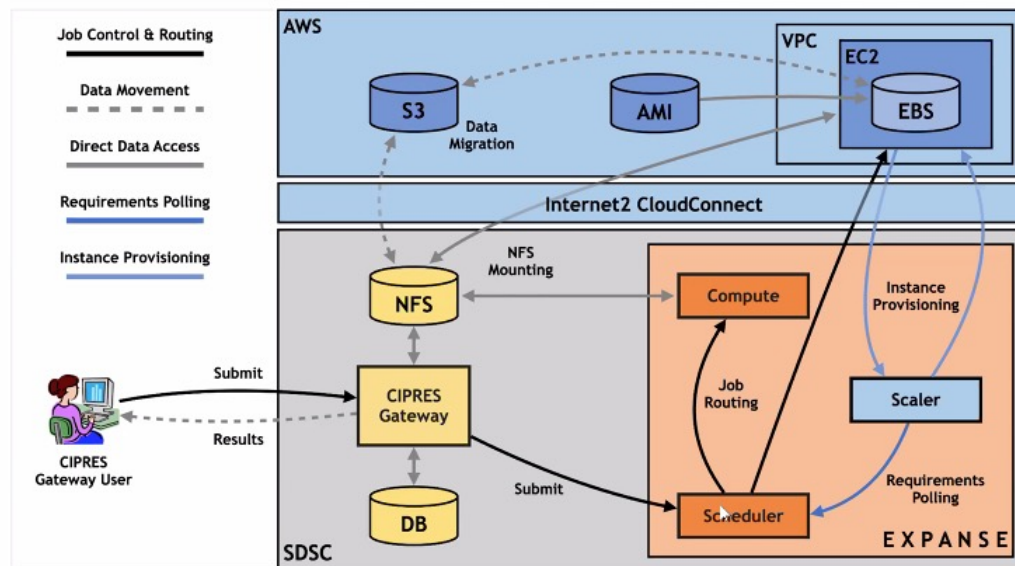


Bright Cluster Manager + Kubernetes  
Requests for a composable system will  
be part of an XRAC request

Composable Systems support complex, distributed, workflows – making Expanse part of a larger CI ecosystem. Advanced User Support resources is available to assist with projects

# Expanse Cloud Bursting

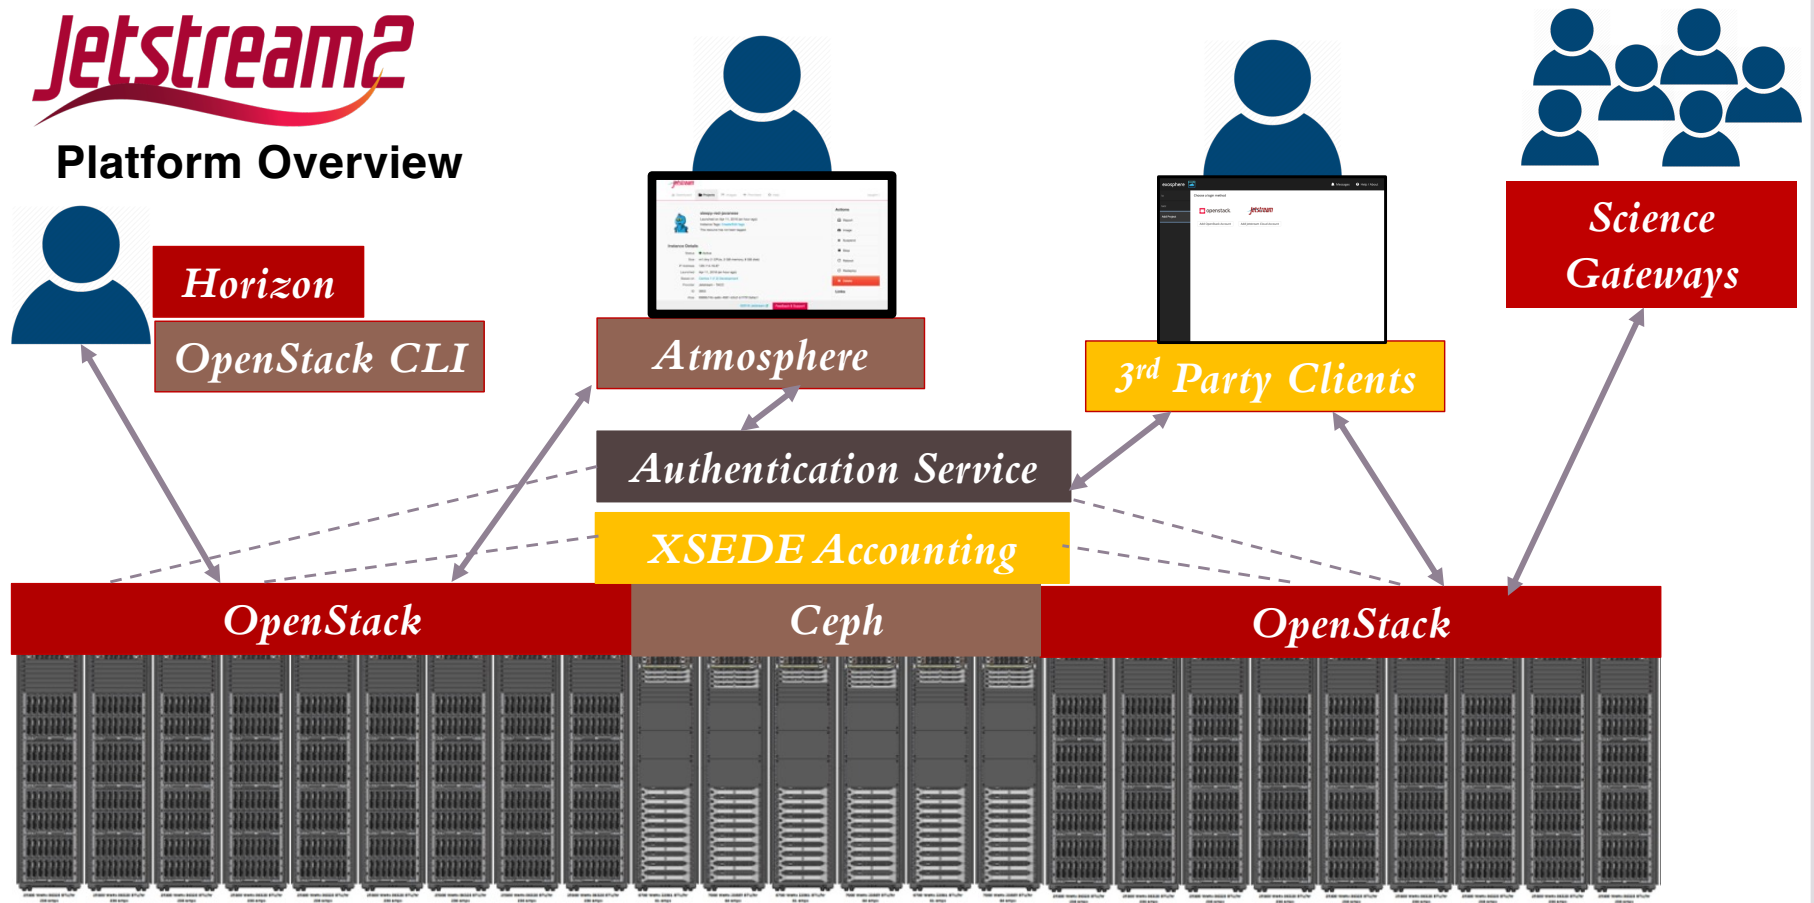
## Cloud Integration Technology



- Expanse supports integration with public clouds.
- Support for projects that share data, need access to novel technologies, and integrate cloud resources into workflows
- Users submit directly via Slurm, or as part of a composed system.
- Options for data movement: data in the cloud; remote mounting of file systems; cached filesystems (e.g., StashCache), and data transfer during the job.

# Jetstream2

## Platform Overview



Jetstream2 Core System

# Jetstream2 Architecture

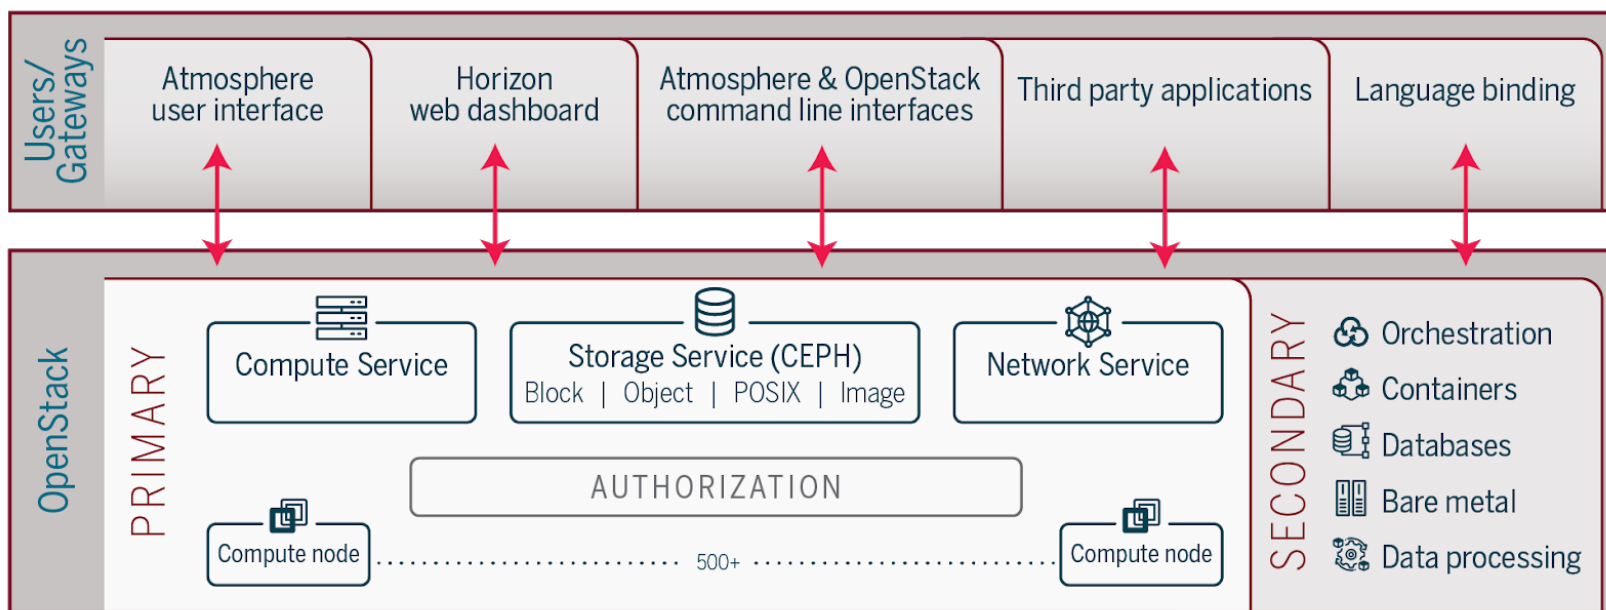
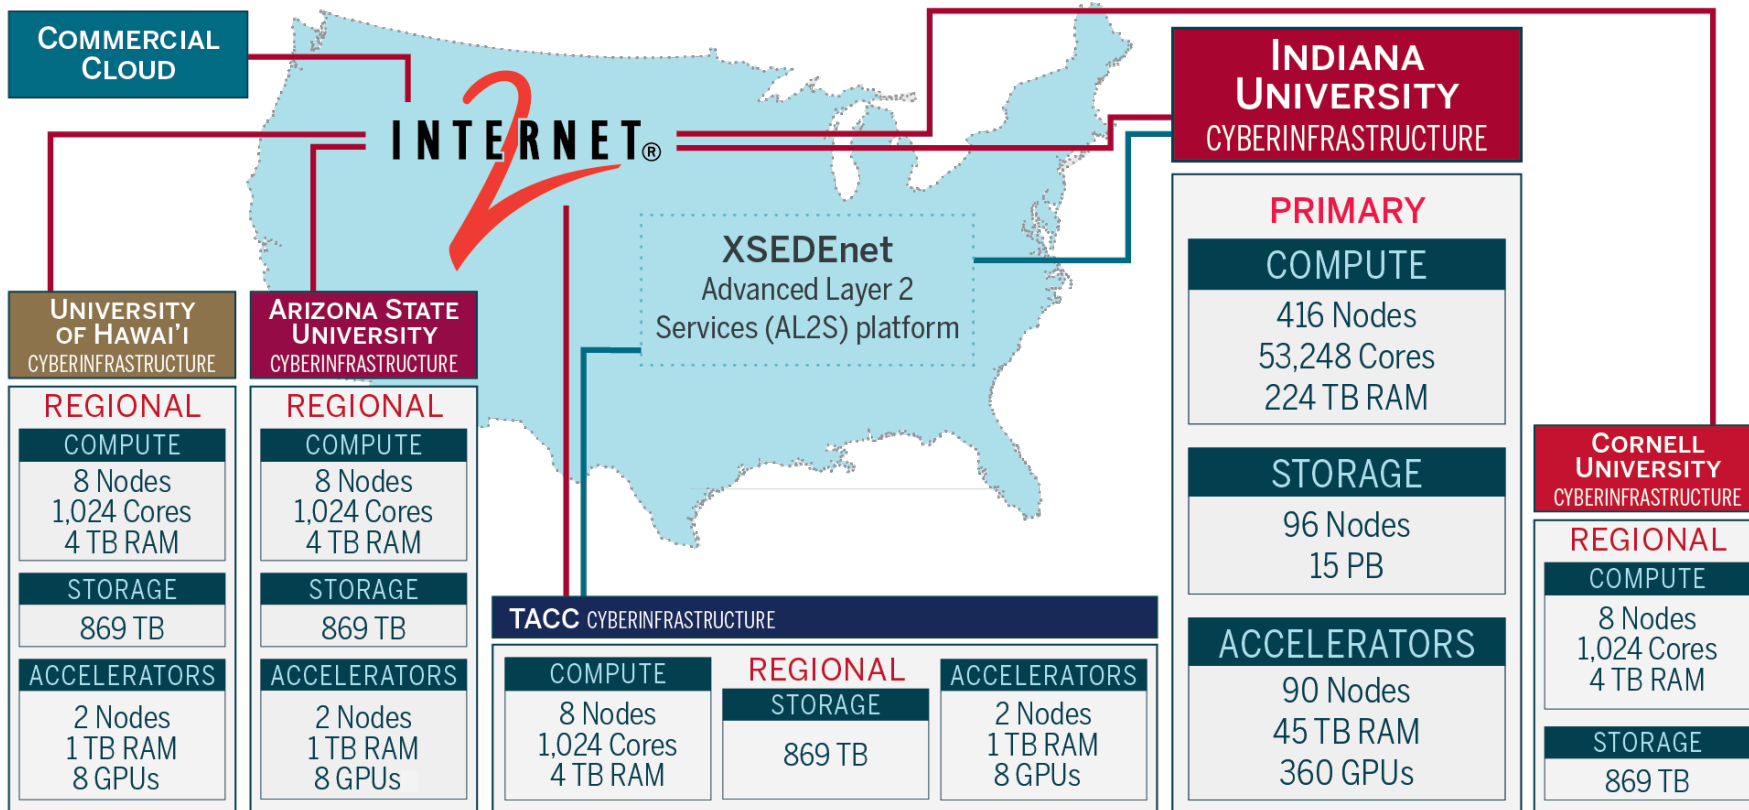


Figure 2: Conceptual and architectural diagram of the system environment

# Jetstream2 Cloud



# Jetstream2 Capabilities

Jetstream2

Home > Project TG-CHE200061 > Create Instance

Jetstream2 IU - TG-CHE200061 (logged in as jfossota@xsede.org)

## Create Instance

Name \*

Image: Featured-Ubuntu20

### Flavor

#### General-purpose

	Name	CPUs	RAM	Root Disk	Ephemeral Disk
<input checked="" type="radio"/>	m3.tiny	1	3 GB	20 GB	none
<input type="radio"/>	m3.small	2	6 GB	20 GB	none
<input type="radio"/>	m3.quad	4	15 GB	20 GB	none
<input type="radio"/>	m3.medium	8	30 GB	60 GB	none
<input type="radio"/>	m3.large	16	60 GB	60 GB	none
X	m3.xl	32	125 GB	60 GB	none
X	m3.2xl	64	250 GB	60 GB	none

#### Large-memory ⓘ

	Name	CPUs	RAM	Root Disk	Ephemeral Disk
X	r3.large	64	500 GB	60 GB	none
X	r3.xl	128	1,000 GB	60 GB	none

#### GPU ⓘ

	Name	CPUs	RAM	Root Disk	Ephemeral Disk
<input type="radio"/>	g3.small	4	15 GB	60 GB	none
<input type="radio"/>	g3.medium	8	30 GB	60 GB	none
<input type="radio"/>	g3.large	16	60 GB	60 GB	none
X	g3.xl	32	125 GB	60 GB	none

### Enhancing IaaS model of Jetstream:

- Improved orchestration support
- Elastic virtual clusters
- Federated JupyterHubs Commitment to >99% uptime
- Critical for science gateway hosting
- Hybrid-cloud support Revamped User Interface
- Unified instance management
- Multi-instance launch

- >57K cores of next-gen AMD EPYC processors
- >360 NVIDIA A100 GPUs will provide vGPUs via NVIDIA's MIG feature
- >18PB of storage (NVMe and disk hybrid)
- 100GbE Mellanox network

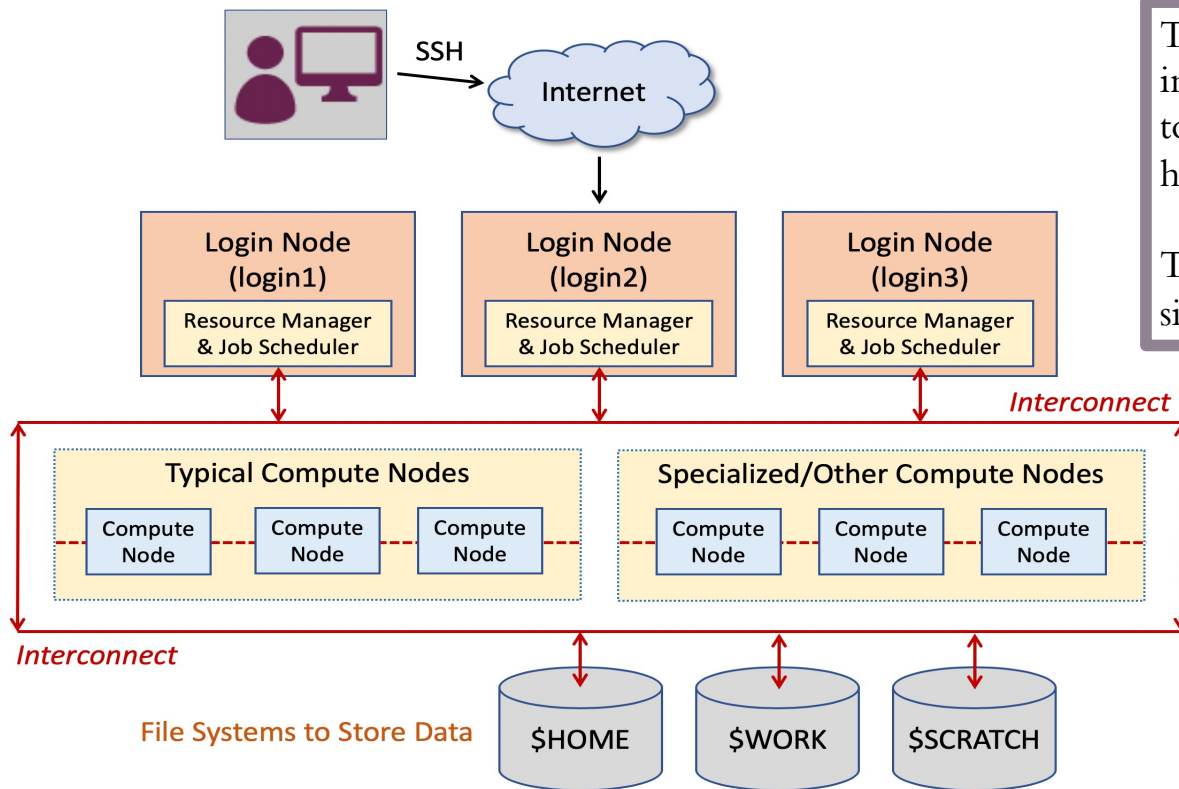
# FRONTERA



<i>Object</i>	<i>Frontera</i>
<i>Compute Node Type</i>	<i>Intel Xeon Platinum 8280 ("Cascade Lake")</i>
<i>Total Nodes</i>	<i>8,368</i>
<i>Cores/Node</i>	<i>56 (28 cores/socket with 2 sockets)</i>
<i>Hardware Threads/core</i>	<i>1</i>
<i>Clock Rate</i>	<i>2.7 GHz Max Turbo: 4.0 GHz</i>
<i>Memory</i>	<i>192 GB DDR4-2933</i>
<i>Peak Memory Bandwidth<sup>6</sup></i>	<i>282 GB/s</i>
<i>L1 Cache</i>	<i>32 KB per core</i>
<i>L2 Cache</i>	<i>1 MB per core</i>
<i>L3 Cache</i>	<i>38.5 MB per socket</i>
<i>Local Storage (/tmp)</i>	<i>144 GB partition on 240 GB SSD</i>



# FRONTERA System Architecture



The Mellanox InfiniBand interconnect ties all the components to each other and to the various high-performance storage systems.

The login nodes are configured similarly to the compute nodes

# FRONTERA *Software Stack*


 **ABAQUS**

Open  **FOAM**

**GROMACS**  
*fast, flexible & free*



**NAMD**  
Scalable Molecular Dynamics



**UASP**  
b-initio  
ackage  
ienna imulation

**NWChem**  
high performance computational chemistry software



**MATLAB**



**Caffe**

**PETSc**

**SLEPc**

The above graphic displays just a small subset of supported applications available on Frontera. *Singularity* is available through the *tacc-singularity* module on Frontera. It cannot be built on the system but can be developed locally, then upload to Frontera.

## *NSF/XSEDE → ACCESS Systems*

Efficient and Effective Use of Resources

*Share your thoughts in the chat:* state observed habits adopted by your users that have resulted in the wasteful usage of resources. What steps have you taken to reduce its occurrence?

# The OSG Consortium

## **Open Science Pool (OSPool)**

provides researchers with *fair-share* access to computing and data capacity powered by *distributed high-throughput* computing (dHTC) technologies.



## **OSG All-Hands Meetings (AHM)**

provides the consortium stakeholders and the broader dHTC community with a venue to *share ideas and exchange information*.


## **Open Science Data Federation (OSDF)**

enables users and institutions to *share data files and storage capacity*, making them both accessible in dHTC environments such as the OSPool.



# The OSPool

Core Hours by Facility



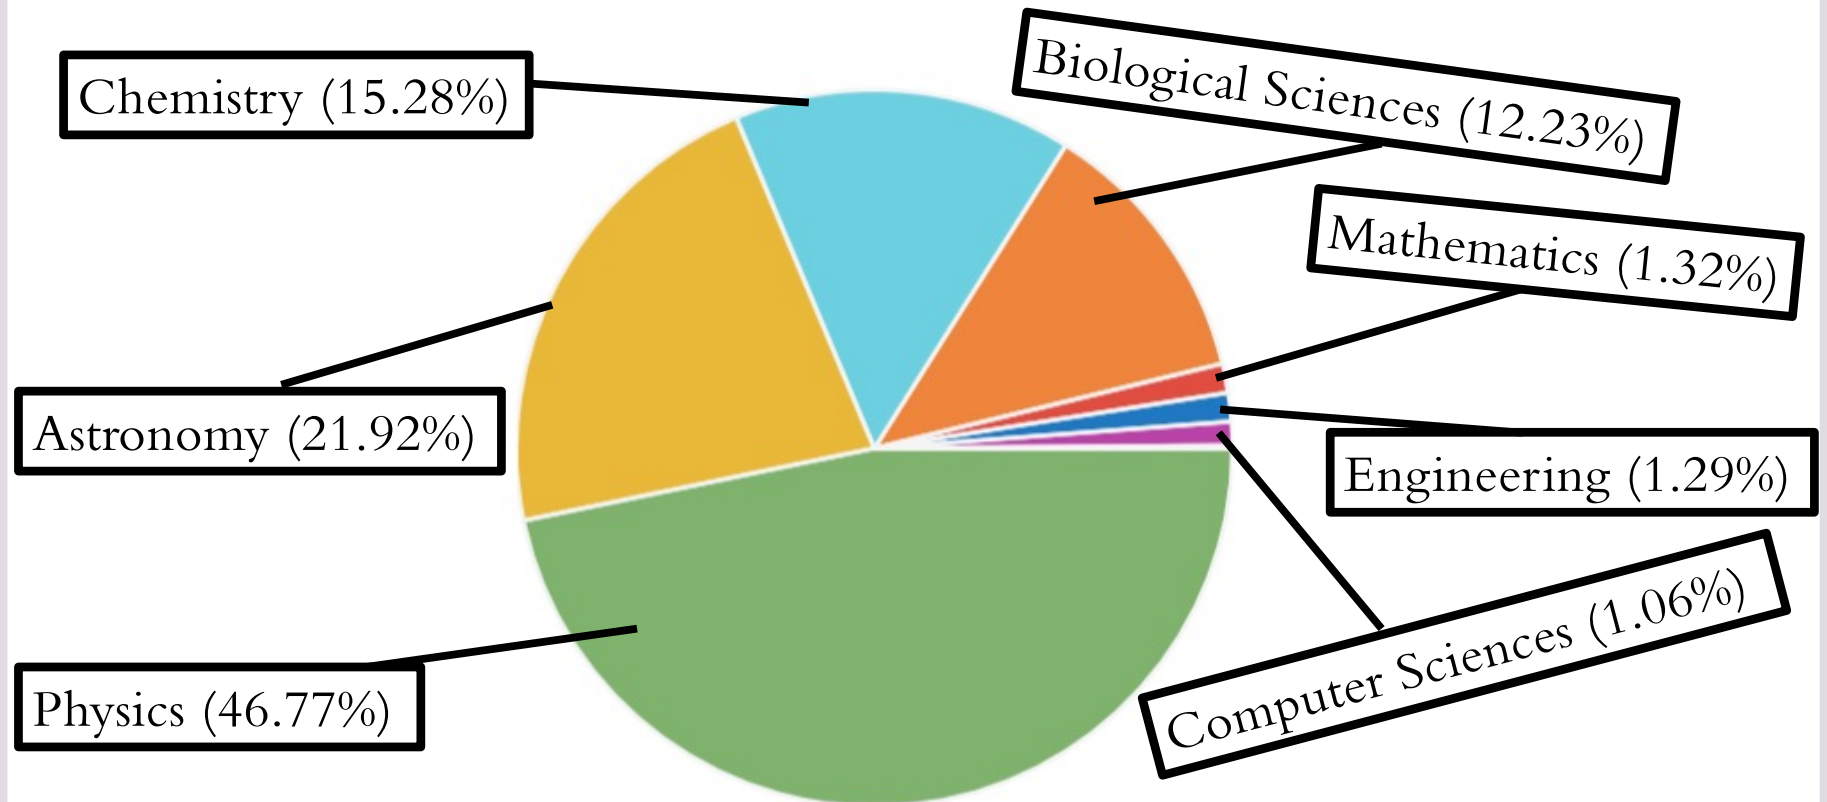
	total
Lancium	4 Mil
MWT2 ATLAS UC	1 Mil
Purdue Geddes	1 Mil
GP-ARGO	606 K
SU ITS	604 K
GLOW	582 K
FermiGrid	409 K
UConn-OSG	399 K
Texas Advanced Computing Center	367 K
UColorado_HEP	358 K
Pervasive Technology Institute	332 K
Rice-CRC	246 K
AGLT2	245 K
WTAMU_HPC	232 K
UTC - Epyc	212 K
Lehigh - Hawk	203 K
SIUE - CC	195 K
UIUC-ICC-SPT	187 K
BNL ATLAS Tier1	165 K
AMNH	158 K



## *Open to All*

- Open to providers at all scales
  - from small colleges to large national labs
- Open to user communities at all scales
  - from individual students to large research communities
    - domain science specific and across many campuses
    - campus specific and across many domain sciences
- Open to any business model
  - sharing, allocations, purchasing
  - preemption is an essential part of operations

## *Core Hours by NSF Field of Science*



## *The OSG Consortium*

### Compute Resource Consolidation

*Share your thoughts in the chat:* have you consolidated compute resources on your campus? What are the benefits of such an initiative?




## *NSF/XSEDE → ACCESS Systems*

### *Systems Information*

- ❑ *Allocated Resources*: general description of the resource and its recommended use.
- ❑ *Non-Allocated Resources*: information about other resources in the XSEDE ecosystem
- ❑ *Allocation Requests*: general information about resources and allocation requests.
- ❑ *Link to information*: <https://portal.xsede.org/allocations/resource-info>

# DOE National Labs:



## OLCF Summit Overview

### *The system includes:*

- 4,608 nodes
- Dual-port Mellanox EDR InfiniBand network
- 250 PB IBM file system transferring data at 2.5 TB/s

### *Each node has:*

- 2 IBM POWER9 processors
- 6 NVIDIA Tesla V100 GPUs
- 608 GB of fast memory (96 GB HBM2 + 512 GB DDR4)
- 1.6 TB of NV memory



### *System Performance:*

- Peak of 200 Petaflops (FP64) for modeling & simulation
- Peak of 3.3 ExaOps (FP16) for data analytics and artificial intelligence

# DOE National Labs:



<i>Feature</i>	<i>Summit</i>
Peak FLOPS	200 PF
Max possible Power	13 MW
Number of Nodes	4,608
Node performance	42 TF
Memory per Node	512 GB DDR4 + 96 GB HBM2
NV Memory per Node	1.6 TB
Total System Memory	2.8 PB + 7.4 PB NVM
System Interconnect	Dual Port EDR-IB (25 GB/s)
Interconnect Topology	Non-blocking Fat Tree
Bi-Section Bandwidth	115.2 TB/s
Processors om node	2 IBM POWER9™ 6 NVIDIA Volta™
File System	250 PB, 2.5 TB/s, GPFS




## OLCF Summit Specs

DOE National Lab



Image courtesy of [wccftech](https://www.wccftech.com)

### OAK RIDGE NATIONAL LABORATORY'S FRONTIER SUPERCOMPUTER



- 74 HPE Cray EX cabinets
- 9,472 AMD EPYC CPUs, 37,888 AMD GPUs
- 700 petabytes of storage capacity, peak write speeds of 5 terabytes per second using Cray Clusterstor Storage System
- 90 miles of HPE Slingshot networking cables

Source: May 10, 2022 Top500 release

Ranking	Performance
<b>#1</b> TOP500	1.1 exaflops of performance on the May 2022 Top500 list.
<b>#1</b> GREEN500	52.23 gigaflops/watt power efficiency.
<b>#1</b> HPL-AI	6.88 exaflops on the HPL-AI benchmark.

**Press Release:** The 59th edition of the TOP500 revealed the Frontier system to be the first true exascale machine with an HPL score of 1.102 Exaflop/s.

# DOE National Lab



<i>Feature</i>	<i>Frontier</i>
Peak FLOPS	> 1.5 EF
Cabinets	> 100
Node	1 HPC and AI Optimized 3rd Gen AMD EPYC CPU 4 Purpose Built AMD Instinct 250X GPUs
CPU-GPU Interconnect	AMD Infinity Fabric
System Interconnect	Multiple Slingshot NICs providing 100 GB/s network bandwidth. Slingshot network which provides adaptive routing, congestion management and quality of service.
Storage	2-4x performance and capacity of Summit's I/O subsystem. Frontier will have near node storage like Summit.

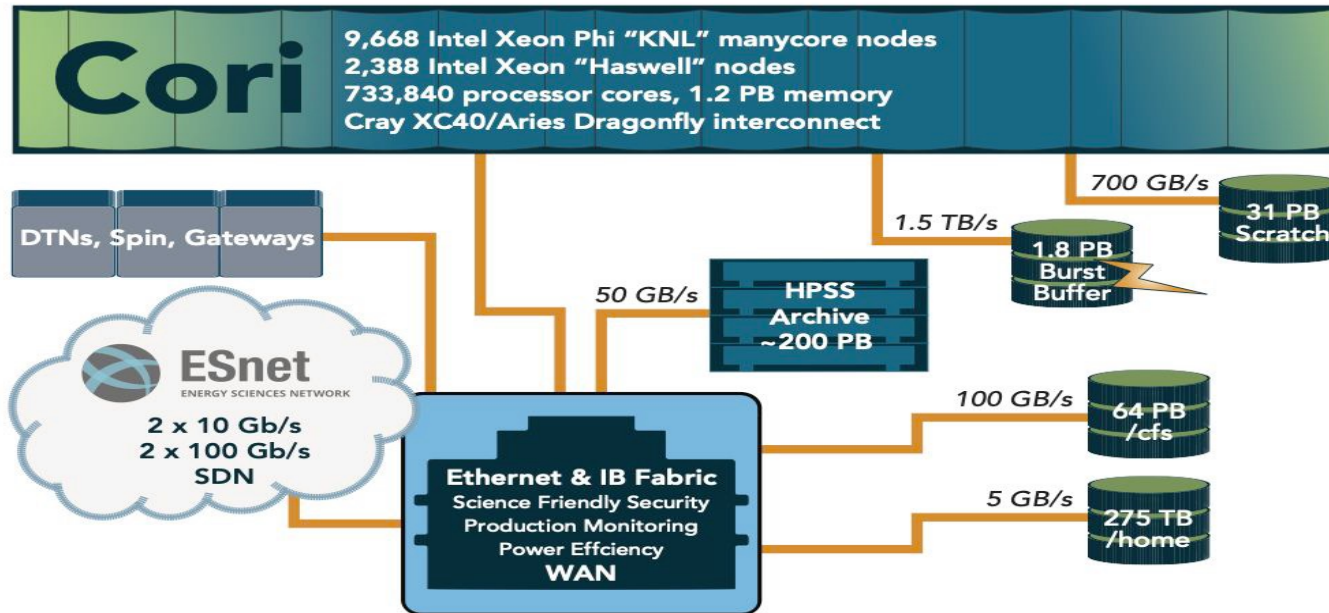


## OLCF Frontier Specs

NB: Summit has 250 PB, 2.5 TB/s, GPFS


# DOE National Lab: NERSC Cori

## NERSC Systems



# DOE National Lab: NERSC Cori

## Chemistry & Materials Science Applications



BerkeleyGW

**NAMD**  
Scalable Molecular Dynamics

*abinit.org*

**GÅMESS**



**WANNIER90**

CPMD consortium page **CPMD**



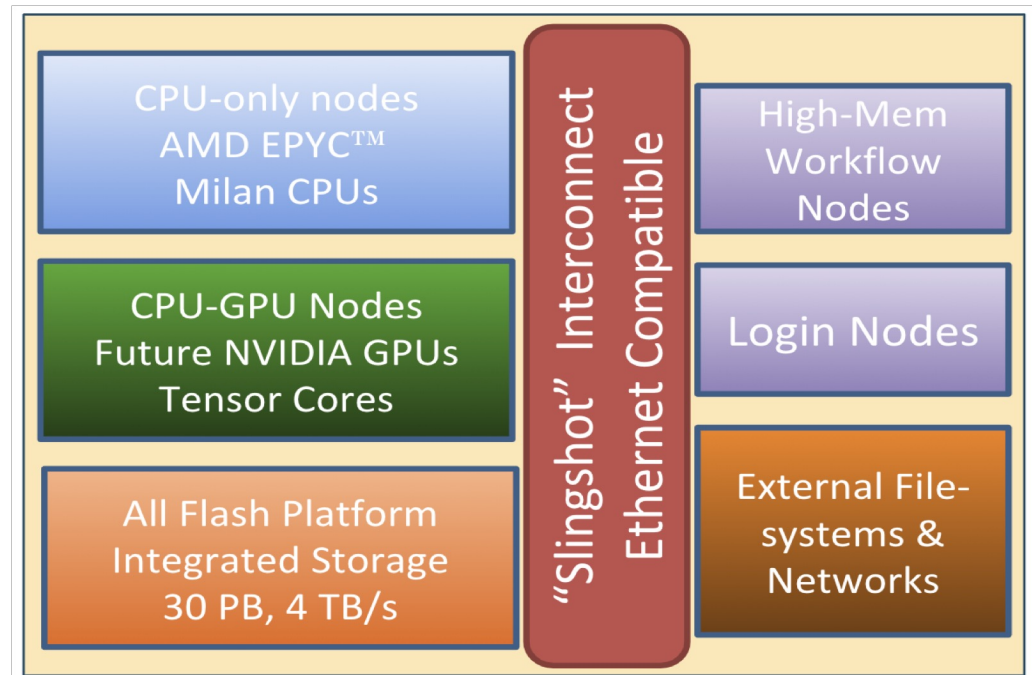
**More than 13.5 million lines of source code Compiled, Optimized, and Tested**

**LAMMPS**

# DOE National Lab: NERSC Perlmutter

## A System Optimized for Science

- GPU-accelerated and CPU-only nodes meet the needs of large-scale simulation and data analysis from experimental facilities
- Cray “Slingshot” - High-performance, scalable, low-latency Ethernet compatible network
- Single-tier All-Flash Lustre based HPC file system, 6x Cori’s bandwidth
- Dedicated login and high memory nodes to support complex workflows





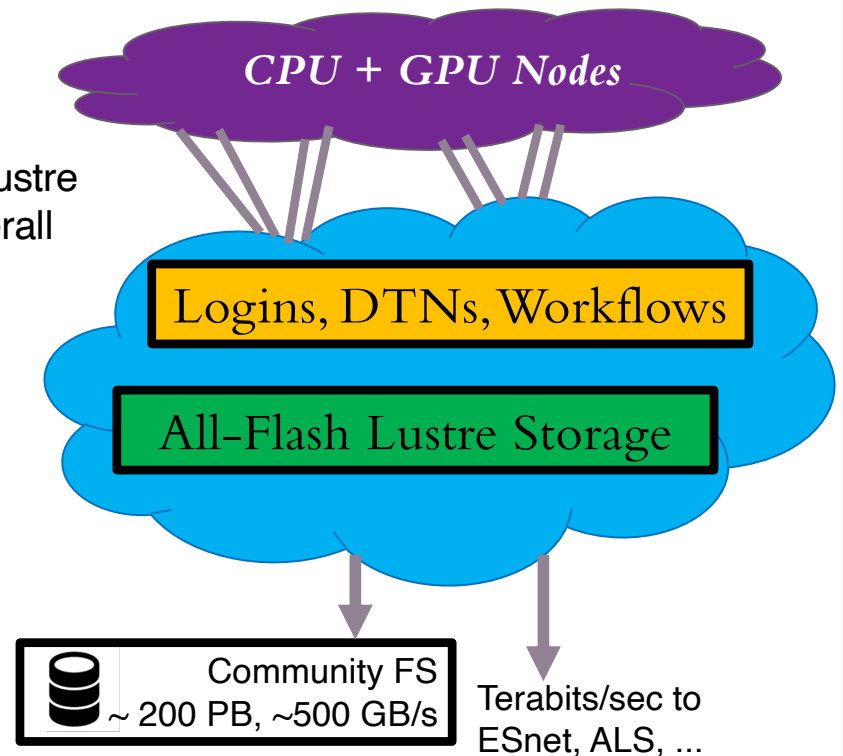
# DOE National Lab: NERSC Perlmutter

## A System Optimized for Science

- ***Fast across many dimensions***
  - 4 TB/s sustained bandwidth
  - 7,000,000 IOPS
  - 3,200,000 file creates/sec
- ***Usable for NERSC users***
  - 30 PB usable capacity
  - Familiar Lustre interfaces
  - New data movement capabilities
- ***Optimized for NERSC data workloads***
  - NEW small-file I/O improvements
  - NEW features for high IOPS, nonsequential I/O



4.0 TB/s to Lustre  
>10 TB/s overall



## *DOE National Lab*

*Share your thoughts in the chat: What are the similarities/differences between the DOE and XSEDE → ACCESS Systems?*

## ***DOE National Lab HPC Centers and Systems***

- ***The National Renewable Energy Laboratory (NREL)***  
<https://www.nrel.gov/hpc/>
- ***Livermore Computing: HPC at LLNL*** <https://hpc.llnl.gov/>
- ***Los Alamos NL High Performance Computing***  
<https://www.lanl.gov/org/ddste/aldsc/hpc/index.php>
- ***Idaho NL HPC*** <https://hpc.inl.gov/SitePages/Home.aspx>
- ***Oak Ridge NL's Leadership Computing Facility***  
<https://www.olcf.ornl.gov/>
- ***Argonne NL's Leadership Computing Facility*** <https://www.alcf.anl.gov/>
- ***Sandia NL HPC*** <https://hpc.sandia.gov/access/>

# Science Gateways

The screenshot shows the 'Science Gateways Catalog' website. At the top, there is a teal header with the title 'Science Gateways Catalog' and the tagline 'Connect people and resources to accelerate discovery by empowering the science gateway community'. Below the header are three navigation buttons: 'FIND A GATEWAY', 'CATALOG', and 'HELP/ABOUT'. The main content area features a search bar with a 'Search' button and a 'Sort by' dropdown menu set to 'A-Z'. On the left, there is a 'FILTER BY:' section with a 'Category: Gateway' filter applied. Below this, a list of categories with their respective counts is shown: GATEWAY (539), SGC CLIENT (34), USED IN CLASSROOM (45), PHYSICAL (264), LIFE (257), SOCIAL (83), APPLIED (157), INTERDISCIPLINARY (157), FORMAL (24), and PHILOSOPHY (33). The main content area displays two gateway entries. The first is 'Arctos: Collaborative Collection Management Solution' with the URL 'http://arctos.database.museum/home.cfm'. It includes a description: 'Arctos is an ongoing effort to integrate access to specimen data, collection-management tools, and external resources on the internet.' and categories: 'Social, Anthropology, Archaeology, History'. The contact info is 'arctos.database@gmail.com' and tags are 'Arctos' and 'museum collection'. The second entry is '20th Century Reanalysis Project' with the URL 'http://portal.nersc.gov/project/20C\_Reanalysis/'.


- Science Gateways simplify access to computing resources by hiding infrastructure complexities.
- Science Gateways provide higher level user interface for XSEDE resources that are tailored to specific scientific communities. (<https://www.xsede.org/ecosystem/science-gateways/gateways-listing>)
- A Science Gateway is a community developed set of tools, applications, and data that are integrated via a portal or a suite of applications, usually in a graphical user interface, that is further customized to meet the needs of a specific community

# Science Gateways

Are you building websites that serve your discipline?

Do you wish you could connect with and learn from others, doing the same thing?

SCIENCE  
GATEWAYS  
.ORG



**SGCI** Science Gateways  
Community Institute


# Science Gateways Listing

XSEDE: <https://www.xsede.org/web/site/ecosystem/science-gateways/gateways-listing>

SGCI Catalog: <https://catalog.sciencegateways.org/#/home>

## SGCI Catalog Summary

- GATEWAY (264)
- SGCI CLIENT (17)
- USED IN CLASSROOM (28)
- PHYSICAL (264) ▾
- LIFE (89) ▾
- SOCIAL (45) ▾
- APPLIED (72) ▾
- INTERDISCIPLINARY (73) ▾
- FORMAL (18) ▾
- PHILOSOPHY (26) ▾



## Pinned Apps A featured subset of [all available apps](#)



Jupyter Notebook

System Installed App

Gromacs

fast  
flexible  
free




Gromacs Desktop

System Installed App



RStudio Server (R-4.1.0)

System Installed App



MATLAB

System Installed App

## Welcome to OnDemand, a Data Science platform and portal to Longleaf

OnDemand provides a web-based interface to the [Longleaf](#) compute cluster with interactive apps such as Jupyter Notebooks, R Studio, Matlab, Stata, and more. These interactive apps allow you to work directly with your files on ITS-RC systems such as your home directory and `/proj`.

For any comments or questions, please contact us at [research@unc.edu](mailto:research@unc.edu).

powered by




OnDemand version: v2.0.20

# Science Gateways Example


UNC INFORMATION TECHNOLOGY SERVICES

Apps Files Jobs Clusters Interactive Apps My Interactive Sessions Help Logged in as jfossota Log Out


## Pinned Apps A featured subset of all available apps




Jupyter Notebook  
System Installed App



Gromacs Desktop  
System Installed App



RStudio  
System Installed App



MATLAB  
System Installed App

## Welcome to OnDemand, a Data Science Gateway


OnDemand provides a web-based interface to the Longleaf Desktop and more. These interactive apps allow you to work directly on the Longleaf Desktop.

For any comments or questions, please contact us at [research@unc.edu](#)

## Portal to Longleaf

Access various apps such as Jupyter Notebooks, R Studio, Matlab, Stata, and more. You can also access files such as your home directory and `/proj`.

- Desktops
  - Longleaf Desktop
  - Gromacs Desktop
- GUIs
  - 3D Slicer
  - COMSOL
  - FSL
  - Firefox
  - Freeview
  - GView
  - MATLAB
  - Mathematica
  - Pymol
  - RStudio Desktop
  - SAS
  - Spyder
  - Stata
- Servers
  - Jupyter Notebook
  - RStudio Server (R-4.1.0)

powered by 

OnDemand version: v2.0.20



# Science Gateways Example



INFORMATION  
TECHNOLOGY SERVICES

Apps ▾

Files ▾

Jobs ▾

Clusters ▾

Interactive Apps ▾

My Interactive Sessions

Help ▾

Logged in as jfossota

Log Out

Open in Terminal ▾

+ New File

New Directory

Upload

Download

Copy/Move

Delete

Home Directory

/proj

/datacommons

/pine/scr/jf/jfossota



/ nas / longleaf / home / jfossota /

Change directory

Copy path

Show Owner/Mode

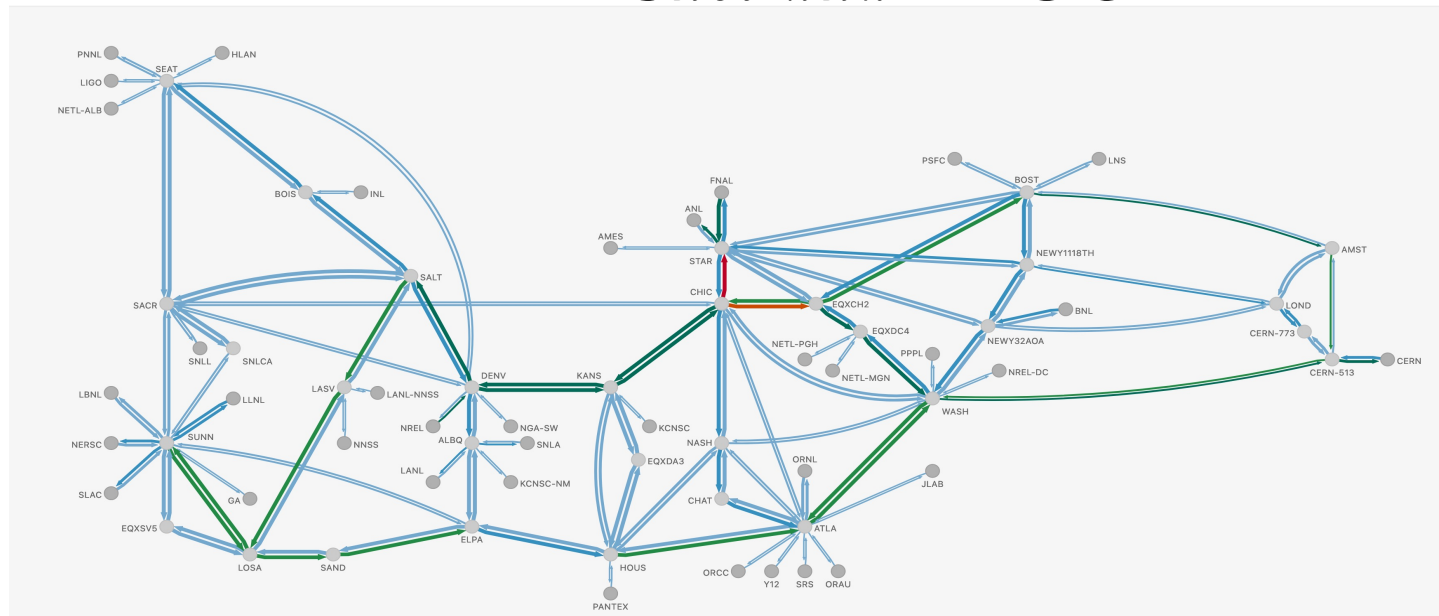
Show Dotfiles

Filter:


Showing 12 of 41 rows - 0 rows selected

	Type	↑↓	Name	↑↓	Size	↑↓	Modified at	↑↓
<input type="checkbox"/>	Folder		Desktop	⋮	-		4/13/2022 11:10:22 AM	
<input type="checkbox"/>	Folder		DMTCP-Tutorial	⋮	-		10/13/2021 3:34:02 PM	
<input type="checkbox"/>	Folder		dogwood	⋮	-		3/31/2020 7:14:47 AM	
<input type="checkbox"/>	Folder		LleafClusterJobs	⋮	-		7/24/2020 1:53:28 PM	
<input type="checkbox"/>	Folder		longleaf	⋮	-		3/31/2020 7:29:52 AM	
<input type="checkbox"/>	Folder		Matlab_2019a_Tests	⋮	-		7/22/2020 11:09:49 AM	

# ESnet and EPOC



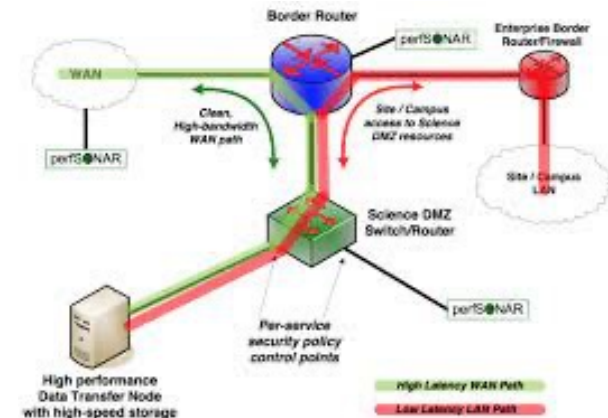
- *ESnet is a high-performance, unclassified network*
- Built to support scientific research.
- *Funded by the U.S. Department of Energy's Office of Science (SC)*
- Managed by Lawrence *Berkeley National Laboratory*.
- Provides services to more than *50 DOE research sites*



# ESnet and EPOC

## Esnet: Technical and Consulting Services

- *The Science DMZ Architecture and Security*
- *Data Transfer Nodes*
- *perfSONAR*
- *Data transfer: Tools and Platforms*
- *Modern Research Data Portal.*
- *Network Requirements Gathering*
- *Collaboration*
- *Training Programs and Workshops*
- *Knowledge Base*
- *CI Engineering Lunch & Learn Series*



<https://www.es.net/science-engagement/technical-and-consulting-services/>

# *ESnet and EPOC*

## *EPOC: Six Main Activities*

- *Roadside Assistance and consultation* via a coordinated Operations Center to resolve network performance problems with end-to-end data transfers reactively;
- *Application Deep Dive* to work more closely with application communities to understand full workflows for diverse research teams in order to evaluate bottlenecks and potential capacity issues
- *Network Analysis* enabled by the *NetSage* monitoring suite to proactively discover and resolve performance issues
- The *Data Mobility Exhibition* and associated work with our simplified portal to check transfer times against known “good” end points
- *Provision of managed data services* via support through the *IU GlobalNOC* and our *Regional Network Partners*
- Coordinated Training to ensure effective use of network tools and science support



### **NSF CC\* Specific Resources**


### **Including EPOC in a Proposal**

<https://epoc.global/proposal-collaborations/>

# *Internet2*

## *About Internet2*

- Internet2 is a community:
  - Higher education and Research institutions
  - Government entities, corporations and cultural organizations
- Cyberinfrastructure provider:
  - secure high-speed network, cloud solutions,
  - research support, services tailored for research and education
- Through InCommon provides:
  - Security and privacy,
  - IAM tools for research and education,
  - single sign-on (SSO) for access to cloud local services and roaming wi-fi



# *Internet2*

## **BY THE NUMBERS**

ADVANCING SCHOLARSHIP • ACCELERATING DISCOVERY

**320+**

HIGHER EDUCATION MEMBERS

**950+**

EDUROAM SUBSCRIBERS

**350+**

NET+ SUBSCRIBERS

**1000+**

INCOMMON PARTICIPANTS

**80,000+**

COMMUNITY ANCHOR INSTITUTIONS

**50+**

INDUSTRY MEMBERS

**60**

AFFILIATE & GOVERNMENT MEMBERS

**100+**

COUNTRIES & RESEARCH NETWORKS CONNECTIONS

**750+**

NET+ CLOUD CONTRACTS

**46**

REGIONAL & STATE NETWORKS

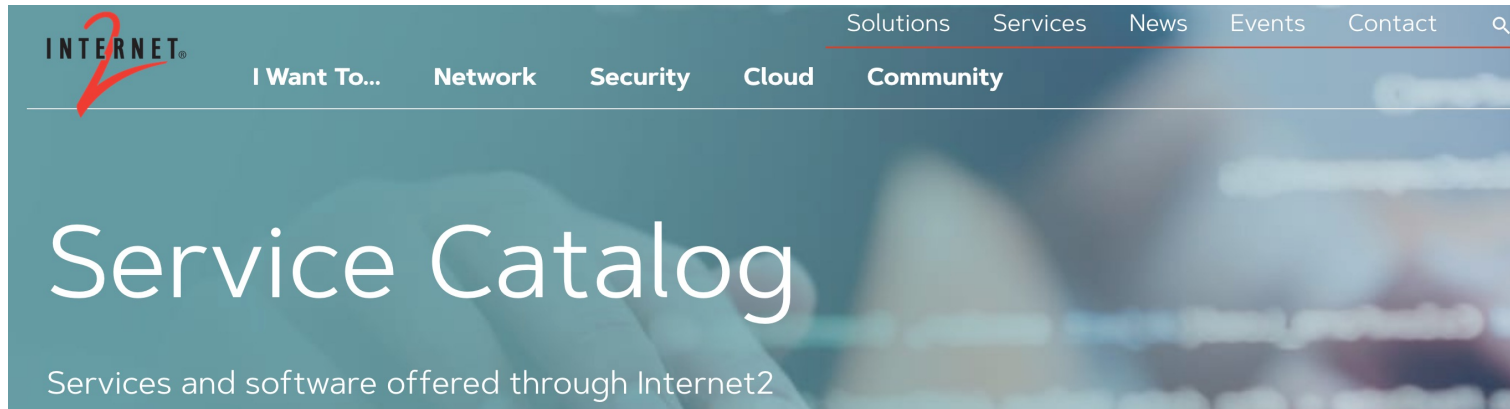
**800G+**

WAVELENGTHS OF NETWORK CAPABILITY

**32Tbps**

CAPACITY PER LINK

# Internet2



The Internet2 community works together to develop, vet, and provide services and software to meet the needs and requirements of research and education.

- **Cloud:** AWS, DocuSign, Duo Security, GCP, ServiceNow, Zoom, Splunk
- **InCommon:** COmanage, eduroam, InCommon Federation, Shibboleth
- **Network:**  
Cloud Connect: Extend your data center to the cloud using the combination of your regional network and the Internet2 high-speed national


# *International Orgs: Compute Canada*

FRANÇAIS ENGLISH

Bienvenue chez  
**calcul**  
canada

SERVICES ET EXPERTS

DÉCOUVREZ CALCUL CANADA





# *International Orgs: Compute Canada*

## *Research Portal*

### Account Management

Apply for an account

Account Renewals

Login to CCDB

Compute Canada CCV Submission Guide

User Roles to Access Resources and Services of the Compute Canada Federation

### National Services

Clusters

Compute Canada Cloud

Storage

Data Movement (Globus)

Humanities and Social Sciences

Visualization

### Accessing Resources

Resource Allocation Competitions

Rapid Access Service

Available Resources for RAC 2022

Technical Glossary

Acknowledging Compute Canada

### Technical Support

Getting Help

Training

Documentation

# *International Orgs:* GÉANT

## Services

GÉANT develops and operates a range of connectivity, cloud, and identity services that ensure a safe and secure environment for researchers, educators, and students.

## Community

The GÉANT Community Programme facilitates collaboration and knowledge sharing between NRENs, user organisations, R&E institutions, and the commercial sector.

## Digital Inclusion



**+50%**

Bringing fibre connectivity to more of European research and education



## Network

Europe's essential terabit network for research and education interconnects Europe's NRENs and links them to over 100 countries in every region of the world.

## Data Growth




**+30%**

Average annual increase in network traffic over last five years

# *International Orgs: GÉANT*

The GN4 Phase 3  
Network project  
(GN4-3N)



# *International Orgs:*



**PARTNERSHIP FOR ADVANCED  
COMPUTING IN EUROPE**

## **INFRASTRUCTURE SUPPORT**



## **HPC IN EUROPE PORTAL**



The mission of PRACE is to enable high-impact scientific discovery and engineering research and development across all disciplines to enhance European competitiveness for the benefit of society.

## **INDUSTRY & SMEs**



## **TRAINING & SUPPORT**



## *Cyberinfrastructure Tools*

*Share your thoughts in the chat:*

What is a cyberinfrastructure tool?

# Cyberinfrastructure Tools

ASK.CYBERINFRASTRUCTURE

Transfer your data.



As a rule, tools can be characterized as devices that help scientists do what they know they must do.

Ref: <https://www.ncbi.nlm.nih.gov/books/NBK25460/>

OPEN

 **nDemand**

**XDMoD**  
METRICS ON DEMAND



I Am Grateful For  
Your Attention.  
Thank You.