

PATh – A Partnership of CHTC and OSG

Miron Livny

John P. Morgridge Professor of Computer Science
OSG Technical Director
UW-Madison/CHTC



The UW-Madison Center for High Throughput Computing (**CHTC**) was established in **2006** to bring the power of High Throughput Computing (**HTC**) to all fields research, and to allow the future of HTC to be shaped by insight from all fields of research



- **CHTC** is home for the HTCondor Software Suite (**HTCSS**)
- **CHTC** is operating a HTCondor pool and a SLURM cluster that are open (fair share) to any campus researcher and their collaborators
- **CHTC** is supporting sharing across ~10 campus HTCondor pools (CMS T2, IceCube, BioState, SpaceScience, ...)
- **CHTC** is providing campus researchers with Research Computing Facilitation services

Some CHTC numbers:

- **HTCSS** is ~1.1M LoC released ~12 times a year
- **CHTC** HTCondor pool delivered ~110M core hours in the past 12 month
- **CHTC** coordinated resources delivered ~360M core hours to more than 17K users from more than 330 projects in the past 12 month
- **CHTC** delivered more than 1.7K Research Computing Facilitation hours in the past 12 month

OSG Statement of Purpose

OSG is a consortium dedicated to the advancement of open science via the practice of distributed High Throughput Computing (dHTC), and the advancement of its state of the art.

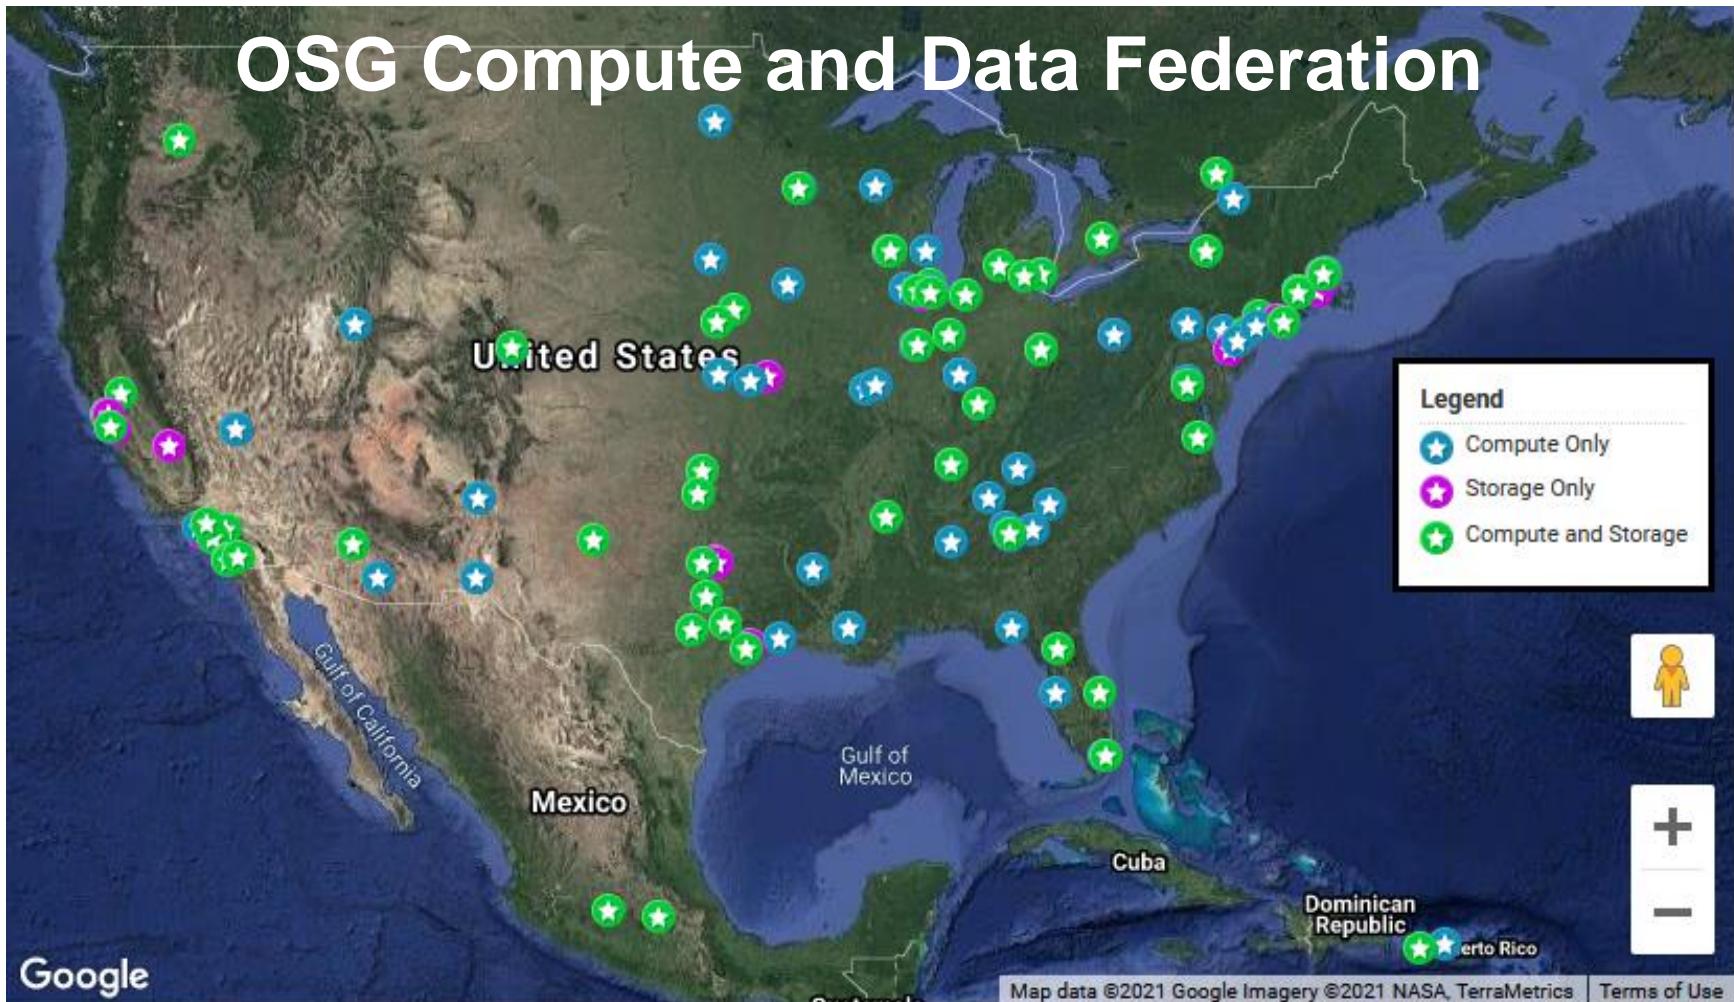
OSG Consortium

- Established in 2005, the OSG is a consortium governed by a **council**
- Consortium Members (Stakeholder) include campuses, research collaborations, software providers and compute, storage, networking providers
- The OSG provides a **fabric of dHTC Services** to the consortium members and to the broader US Science and Engineering (S&E) community
- While **members own** and operate resources, the consortium does not own or operate any resources
- Council elects the OSG **Executive Director** who appoints an **Executive team**. Together they steer and manage available effort

OSG Fabric of Services

- Organized under three main thrusts – *Community Building, Research Computing Facilitation, and Operation*
- Designed and operated to assure, **scalability, trustworthiness, reproducibility**.
- OSG claims its services enabled in the past **12** month more than **2B** core hours across more than **130** clusters located at more than **70** sites and more than **200TB** of data cached across **17** caches worldwide.

OSG Compute and Data Federation



A Partnership Launched!

On October 1, 2020 we started the 5 years, \$4.5M annual budget NSF “**Partnership for Advanced Throughput Computing (PATh)**” project

“The Partnership to Advance Throughput Computing (PAPath) project will expand Distributed High Throughput Computing (dHTC) technologies and methodologies through innovation, translational effort, and large-scale adoption to advance the Science & Engineering goals of the broader community.”



A strong partnership

- Partnership between the UW-Madison **CHTC** and the **OSG** Consortium
- Two main elements of the project are the **HTCondor Software Suite (HTCSS)** and the **Fabric of Capacity Services (FoCaS)** offered by the OSG
- Involves 40 individuals at seven institutions
- Builds on decades of collaboration, common vision and shared principals
- Committed to community building (HTCondor Week and the OSG school)
- Aligned with the NSF Cyberinfrastructure blueprint

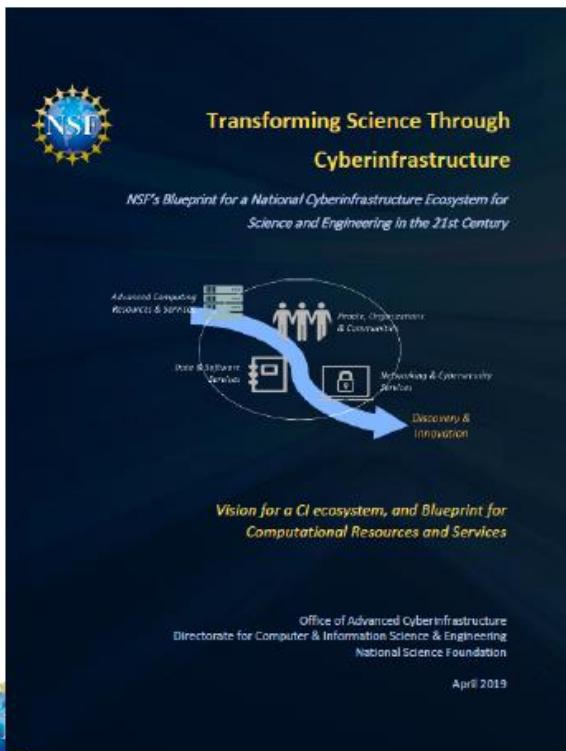


Manish Parashar

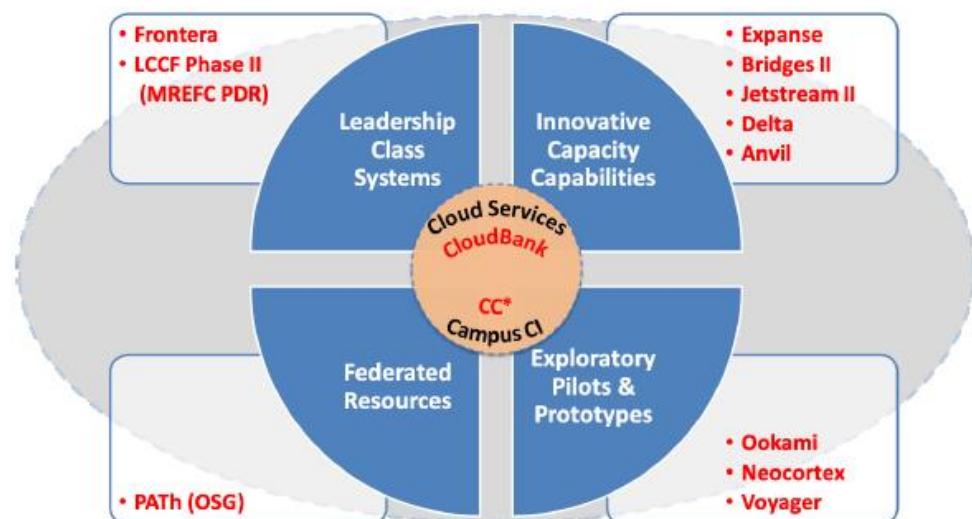
Director, Office of Advanced
Cyberinfrastructure,
Directorate for Computer & Information
Science & Engineering
National Science Foundation

November 30, 2020

NSF's Vision for a National CI Ecosystem



Computational Ecosystem: Elements



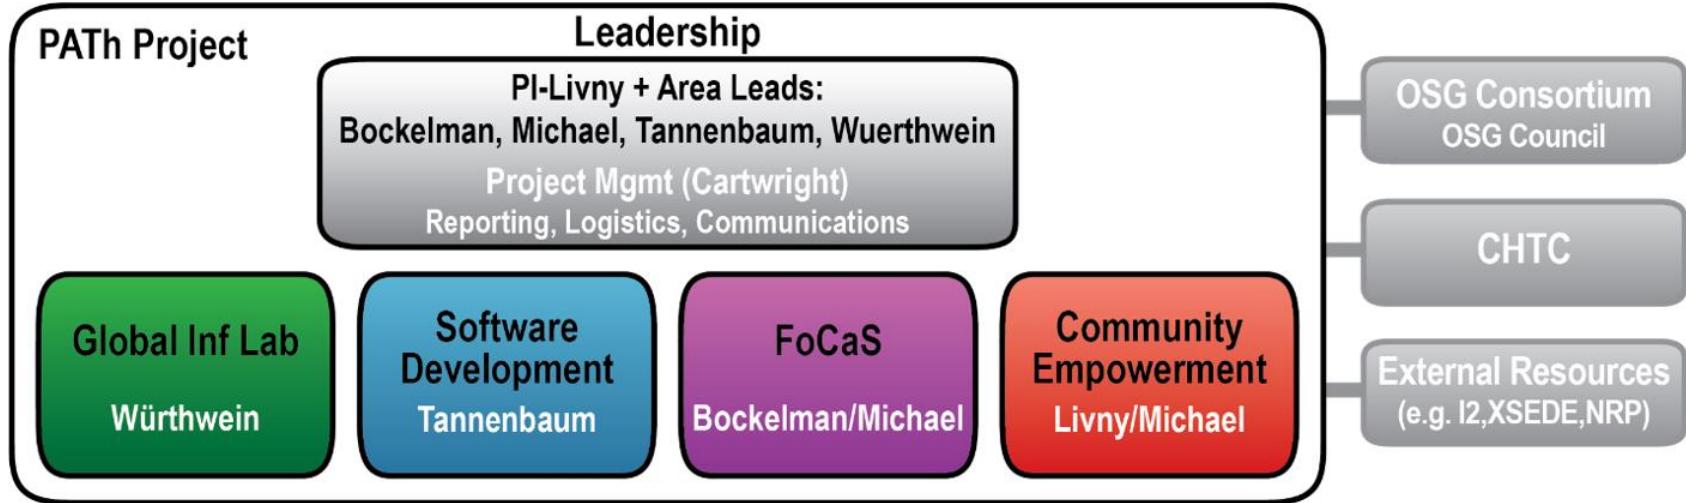


Figure 2: PATh Organization

- Strong ties to the OSG Executive Team
- OSG Council serves as Advisory Board
- Cooperative Agreement with the NSF

OSG Effort

OSG effort is managed by the OSG Executive Director (Frank) with the help of the Executive Team that includes individuals from BNL and Fermi

- 68% is contributed by PATh
- 22% is contributed by IRIS-HEP
- 10% is contributed by “other”

The Open Science Pool (OSPool)

One of the OSG services are **Access Points (APoint)** that are open to any US researcher and a distributed HTCondor pool that is managed under a fair-share scheduling policy

OSG Compute Federation sites contribute resources to the OSPool by (dynamically) running **Execution Points (XPoints)** according to autonomous site policies

- APoint provides workload automation, auditing, and workflow management (DAGMan, Pegasus) capabilities designed to accommodate High Throughput applications in a distributed environment
- Data for input sandboxes is staged at the APoint or placed in the OSG data federation
- Output sandbox data is staged at the APoint
- On 04/20/2021, the OSPool completed ~200K jobs from 29 projects submitted by 33 users that consumed ~640K core hours

How do sites contribute to the OSPool?

When a site in the OSG Compute Federation wants to contribute the capacity of a server to the OSPool, it runs an XPoint on the server

- Site can start and stop the XPoint at any time
- XPoint needs to establish trust with APoint
- OSG provides services to remotely activate XPoints through a Compute EntryPoint (CE) that submits activation requests to the batch system of the cluster
- OSG provides services to automate remote activation of XPoints
- XPoint prefers to have out going network connectivity



Open Science Grid

The new GRACC page



[OSG Connect Communities](#)

Individual PIs and small groups via OSG

Core Hours - 24 Hours

1.203 Mil

Core Hours - 30 Days

27.3 Mil

Core Hours - 12 Months

262 Mil

[Multi-Institutional Projects](#)

Including IceCube, LIGO, DES, glueX,

Core Hours - 24 Hours

689 K

Core Hours - 30 Days

16.6 Mil

Core Hours - 12 Months

225 Mil

[Campus Organizations](#)

Individual PIs and small groups via their

Core Hours - 24 Hours

95.5 K

Core Hours - 30 Days

3.974 Mil

Core Hours - 12 Months

28.7 Mil

[GPU Utilization](#)

This includes a heterogeneous mix of GPU models.

Core Hours - 24 Hours

38.5 K

Core Hours - 30 Days

741 K

Core Hours - 12 Months

7.685 Mil

LHC Usage

ATLAS and CMS

Core Hours - 24 Hours

4.704 Mil

Core Hours - 30 Days

118 Mil

Core Hours - 12 Months

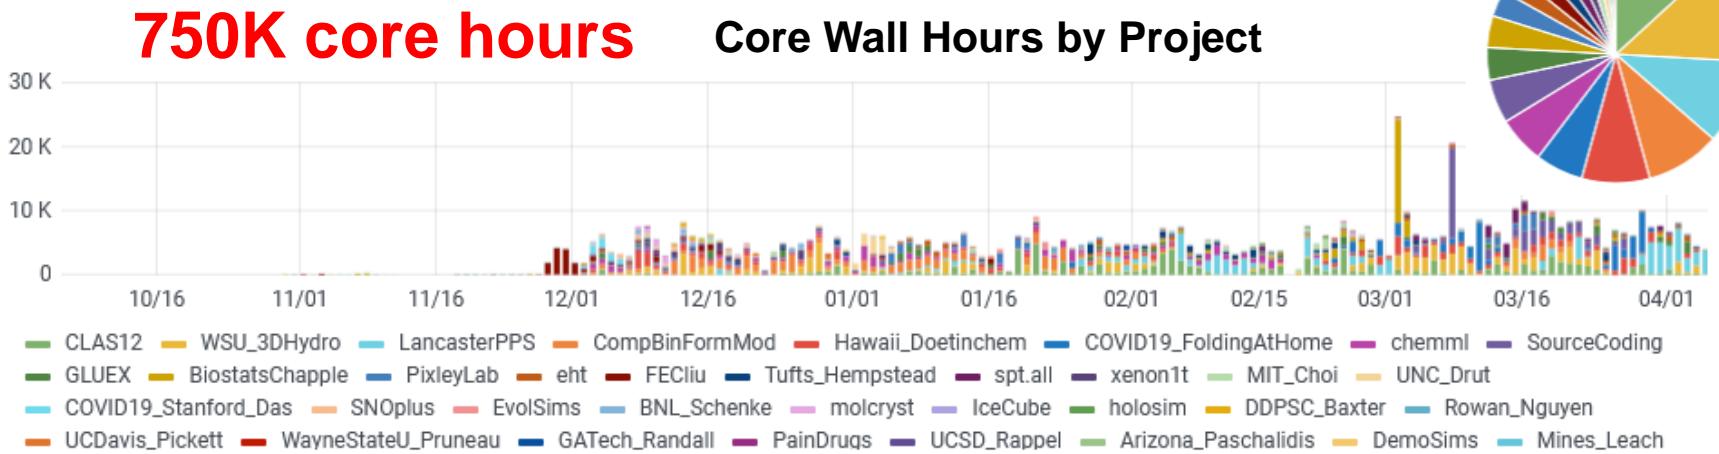
1.457 Bil

Highlighting the 3 distinct groups, our GPU offerings, and the LHC.

OpenStack Environment

The **Jetstream** cloud is an OpenStack, NSF-funded academic cloud operated by the Indiana University and TACC.

Operators decide when to start and terminate Virtual Machines (or containers) that run OSPool XPoints.



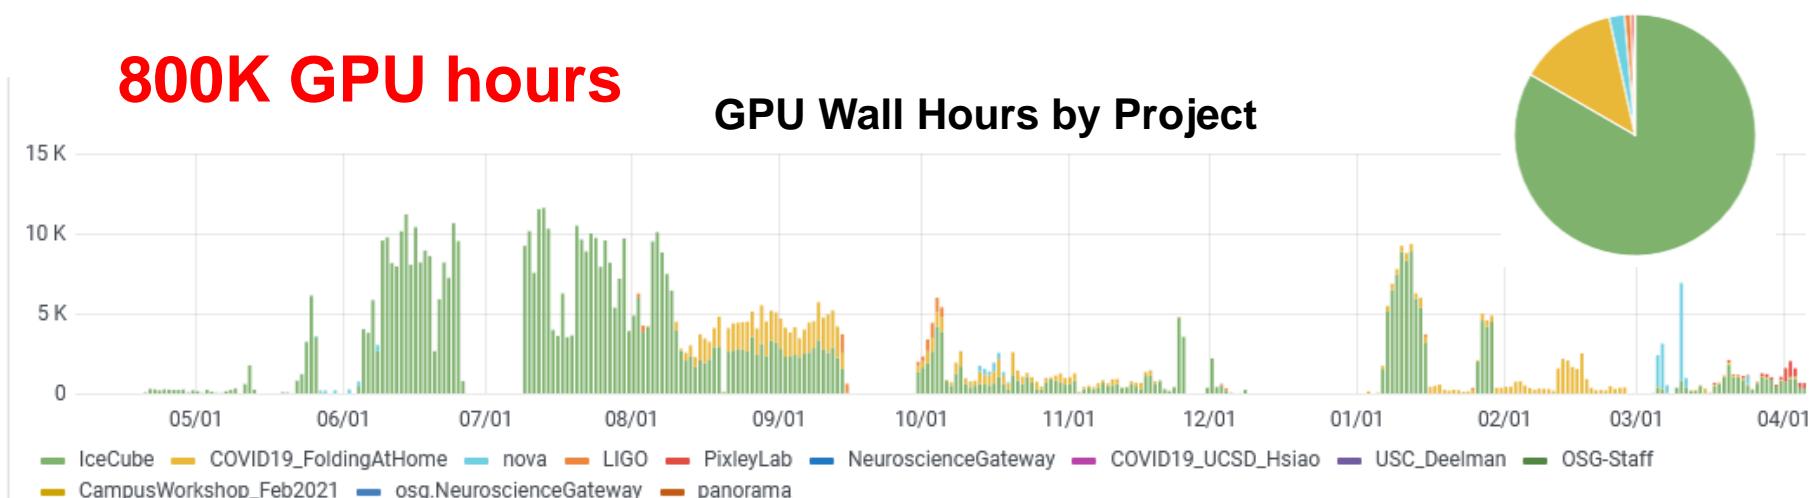
Work with NSF funded **Chameleon** is progressing nicely

k8s Cluster

The NSF funded Pacific Research Platform **PRP** operates a k8s cluster across more than 35 location worldwide.

A PRP HTCondor pool grows and shrinks under the control of the k8S cluster pod scheduler

A CE submits requests to HTCondor to deploy OSPool XPoints.



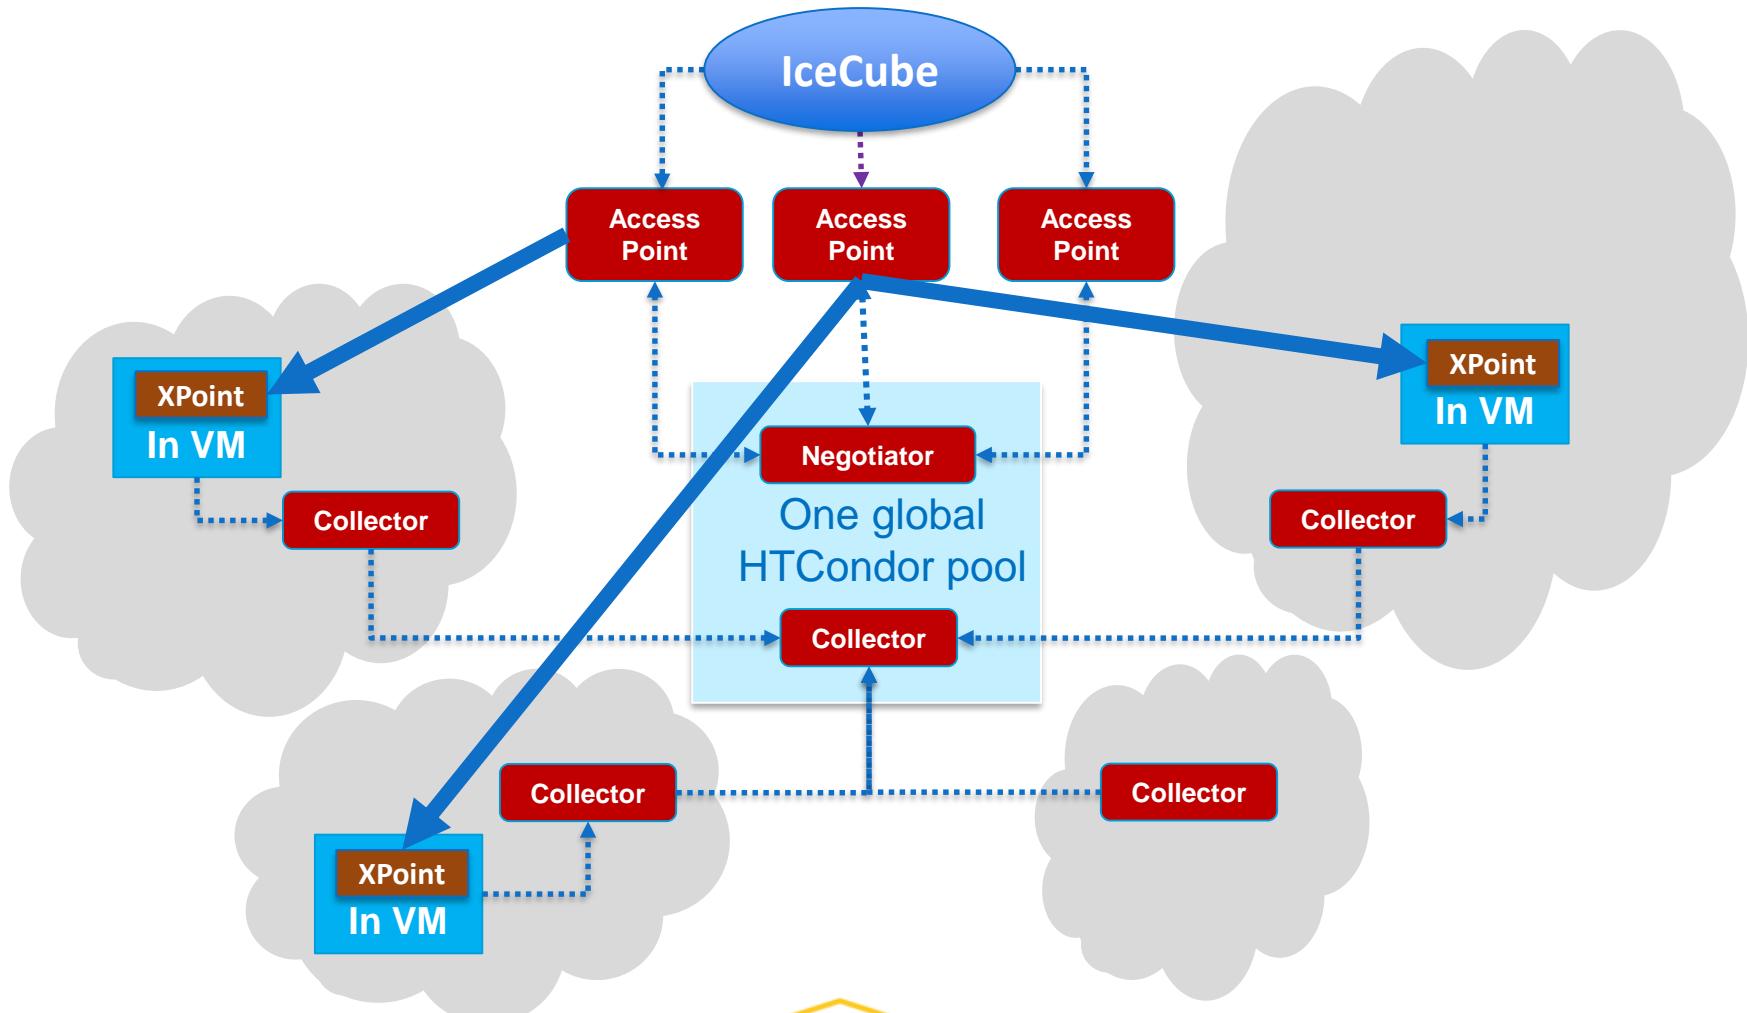
OSPool is not the only pool!

Organizations like science collaborations (CMS, LIGO, IceCube) and campuses (UCSD, UNL, UW-Madison, JLab) leverage OSG services to deploy and operate private distributed HTCondor pools

- APoints are private
- Deployment of XPoints negotiated by the organization with sites
- Resource acquisition and allocation policies defined by the organization

* A site can contribute to different HTCondor pools

HTCondor dHTC CI of >51K GPUs in the Cloud @ SC19



PATh Executive Summary

“Broader Impact – We firmly believe in dHTC as an accessible computing paradigm which supports the democratization of research computing to include researchers and organizations otherwise underrepresented in the national CI ecosystem. Our work is founded on universal principles like sharing, autonomy, unity of purpose, and mutual trust.”

Democratizing Access

In her presentation at the NSFNET 35th Anniversary NSF CISE AD Margert Martonosi articulated the challenge of **Democratizing Access to National Research Computing Resources.**



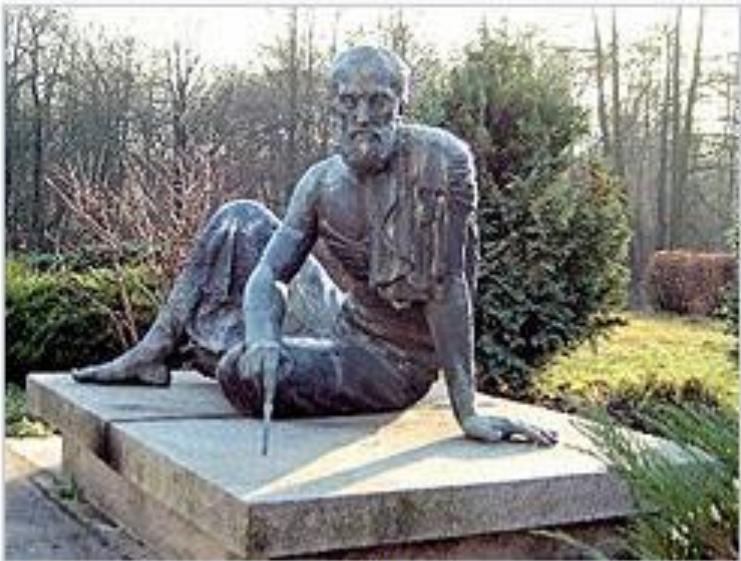
National Computing Research Resource



We view Access Points as holding the key to addressing this national challenge

- OSPool APoints can be deployed and operated by a single PI laboratory or by organizations like campuses and science collaborations
- APoints can be used to manage deployment of XPoints
- APoints can scale out to accommodate large HTC workloads

The Open Science Pool
(OSPool) is our
answer to
democratization of
access challenge!



Give me a place to stand, and I shall move the world.



Archimedes of Syracuse was a Greek mathematician, philosopher, scientist and engineer.



Give me a place to run an XPoint and I shall run your job.

Frank Würthwein is a Physics professor at UCSD and the Executive Director of the OSG