Good Morning!

Good Afternoon!

Good Evening!

Good Night!

Good Late Night!

Good early Morning!

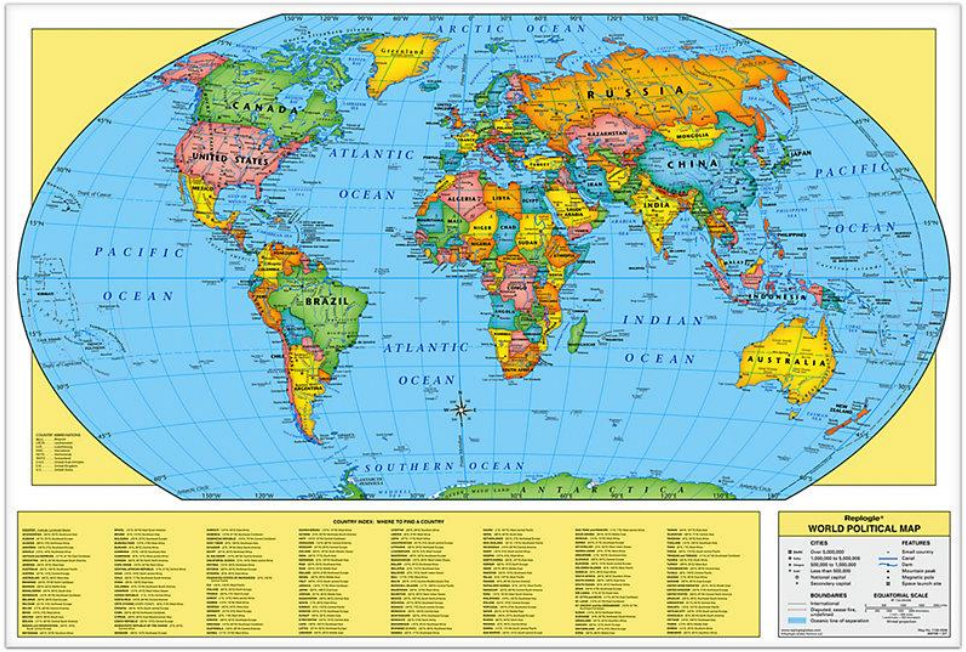


# (-2 to +15)

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

Open Science Grid

HTC CENTER FOR HIGH THROUGHPUT COMPUTING

# Thank you for joining us for the 22th


# Week

# (year 37 of the (HT)Condor project)

# (year 15 of the CHTC)

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

Open Science Grid

HT CENTER FOR HIGH THROUGHPUT COMPUTING

ICDCS2020: Opening and awards

11/2020


IEEE COMPUTER SOCIETY

Technical Committee on Distributed Processing

**2020 Outstanding Technical Achievement Award**

# Professor Miron Livny

For Influential Contributions of the Condor System to Distributed and High Throughput Computing

22:21 / 43:14

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

Open Science Grid

HTC CENTER FOR HIGH THROUGHPUT COMPUTING

**“The Contributions of the Condor system to Distributed and High Throughput Computing is the result of the ingenuity, imagination, commitment and hard work of many individuals. I am standing here to receive this award as a representative of this special group of researchers, developers and students. These contributions are also due to the researchers who trusted us and adopted Condor to serve their high throughput computational needs. Their guidance and feedback have been invaluable in getting us to where we are today in terms of understanding distributed systems and High throughput computing technologies.”**

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING


HT CENTER FOR HIGH THROUGHPUT COMPUTING

THE OVERFLOW

Find something

Latest Newsletter Podcast Company

code-for-a-living MAY 11, 2021

# Testing software so it's reliable enough for space

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

Open Science Grid

HT CENTER FOR HIGH THROUGHPUT COMPUTING

The CI environment is based on [HT Condor], a workload management system for compute-intensive jobs that originated with the High Throughput Computing Group at the University of Wisconsin Madison. It's prized at SpaceX for its powerful queueing, job prioritization, and resource management capabilities – particularly for the HITL testbeds (more on those later).

Condor manages workloads similar to a traditional batch system, but can make better use of idle computer resources. Ishimoticha says, "we run about a million CI builds a month."

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING
Open Science Grid
HT CENTER FOR HIGH THROUGHPUT COMPUTING

# The NMI Build & Test Laboratory: Continuous Integration Framework for Distributed Computing Software

Andrew Pavlo, P. Couvares, Rebekah Gietzel, Anatoly Karp, Ian D. Alderman, M. Livny, Charles Bacon less •

Published in LISA 2006 • Computer Science

We present a framework for building and testing software in a heterogeneous, multi-user, distributed computing environment. Unlike other systems for automated builds and tests, our framework is not tied to a specific developer tool, revision control system, or testing framework, and allows access to computing resources across administrative boundaries. Users define complex software building procedures for multiple platforms with simple semantics. The system balances the need to continually... Expand

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

Open Science Grid

HT CENTER FOR HIGH THROUGHPUT COMPUTING

# Oracle Case Study – Implementation Timeline

- 1998-2002
  - ☑ Validation of existing products
  - ☑ Solution Validation
    - Working Prototype

- 2002-Present
  - ☑ Founded Optena
    - Mission is to provide Scalable, Manageable High Throughput Computing Solutions
    - Partnered with UW to get a source license of Condor in order to build a complete solution
    - Plan was to build additional capability into Condor and around it
  - ☑ Deployed to Oracle
    - Initial 400 node production Grid Dec 2003
    - Grew to 3000+ node Grid by Apr 2004
    - Running on Windows, Linux and Solaris

optena DO IT SMART

# 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 (PATh) 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."*

**Aligned with NSF Cyberinfrastructure blueprint**

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

Open Science Grid

HTC CENTER FOR HIGH THROUGHPUT COMPUTING


**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

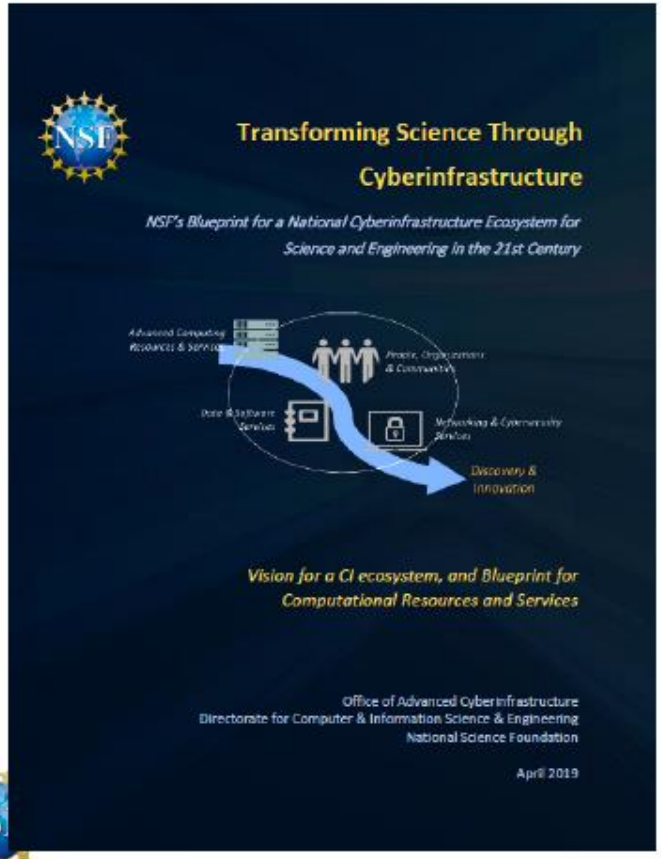


## Computational Ecosystem: Elements

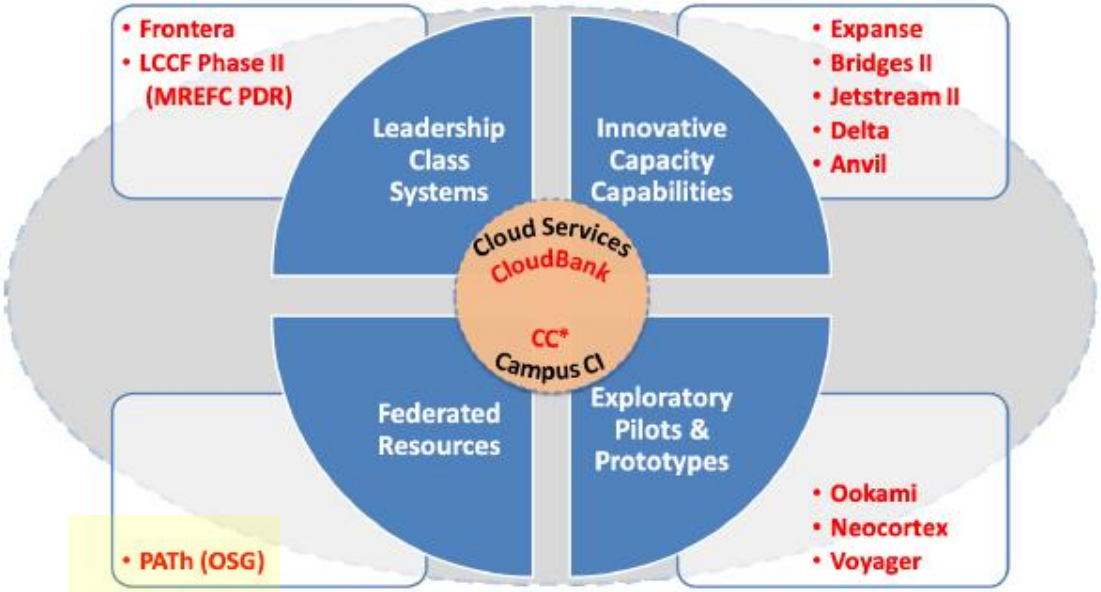

# An organic partnership

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

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

Open Science Grid

HT CENTER FOR HIGH THROUGHPUT COMPUTING

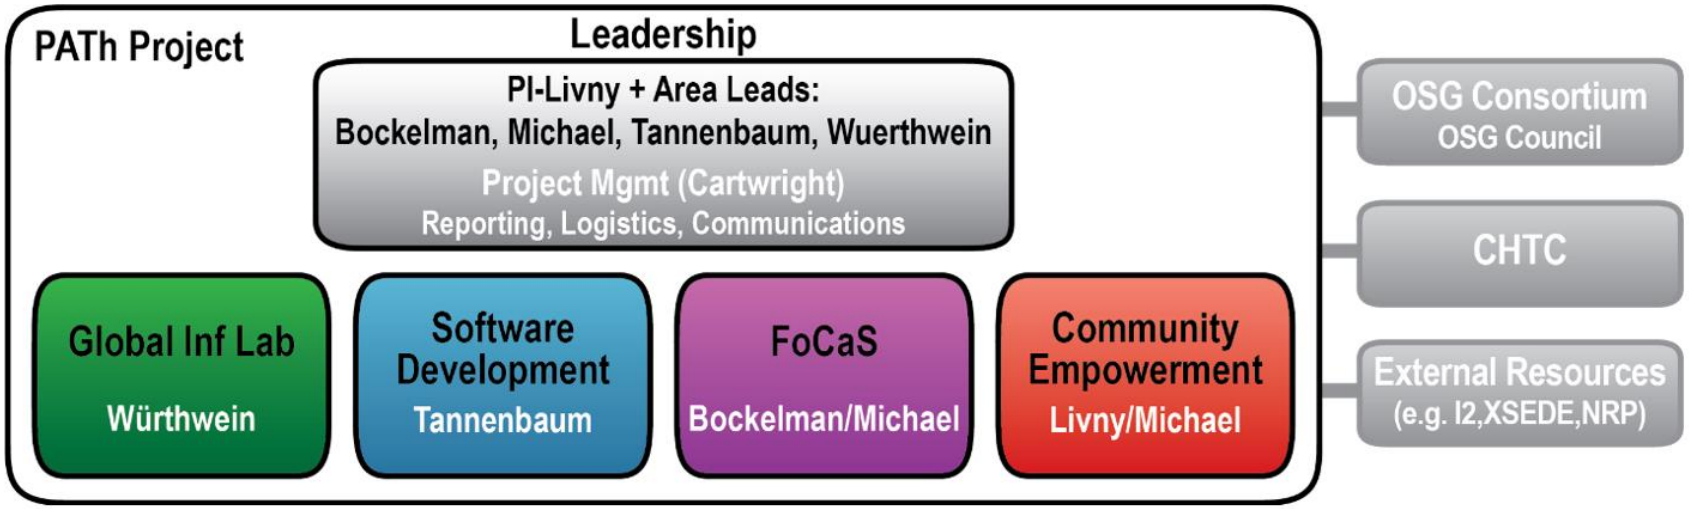


*Figure 2: PATh Organization*

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

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

Open Science Grid

HTC CENTER FOR HIGH THROUGHPUT COMPUTING

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**

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

Open Science Grid

HT CENTER FOR HIGH THROUGHPUT COMPUTING

- **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

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

Open Science Grid

HT CENTER FOR HIGH THROUGHPUT COMPUTING

# 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

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

Open Science Grid

HT CENTER FOR HIGH THROUGHPUT COMPUTING

# 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.**

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

Open Science Grid

HT CENTER FOR HIGH THROUGHPUT COMPUTING

# 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.

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

Open Science Grid

HTC CENTER FOR HIGH THROUGHPUT COMPUTING

# OSG Compute and Data Federation

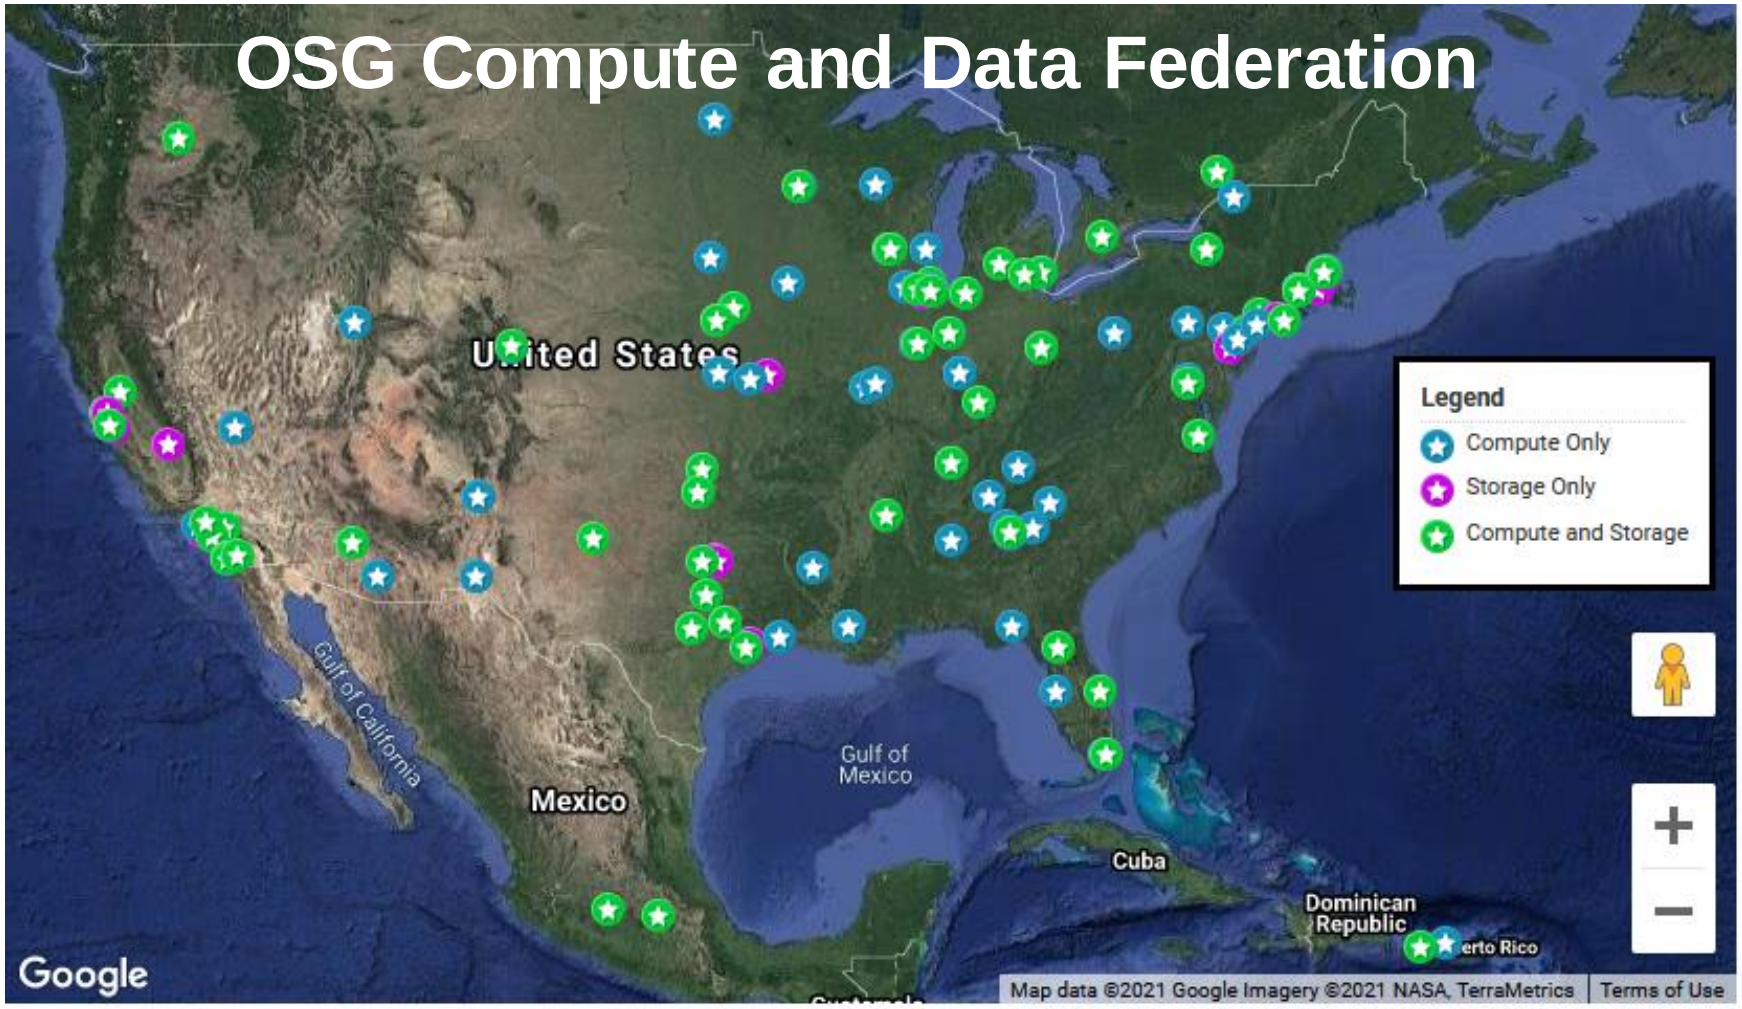


PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

Open Science Grid

HTC CENTER FOR HIGH THROUGHPUT COMPUTING

# How do Cloud Resources fit into the Distributed High Throughput Computing (dHTC) model of the fabric of Services provided by the Open Science Grid (OSG)?

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING


HT CENTER FOR HIGH THROUGHPUT COMPUTING

# Naturally!

**It is all about offering**
**Access Points**
**and services to deploy**
**Execution Points**

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING


HT CENTER FOR HIGH THROUGHPUT COMPUTING

# The Open Science Pool (OSPool)

A service provided by the OSG 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

- 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

# 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."

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

Open Science Grid

HT CENTER FOR HIGH THROUGHPUT COMPUTING

All this leads to the democratization of computing-

You do not have to be a "super" person in order to benefit from "super" computing power

# 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**.


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

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

Open Science Grid

HT CENTER FOR HIGH THROUGHPUT COMPUTING

# Bring Your Own Resources (BYOR)

The members of **PATh** are working on enhancing BYOR services for institutional clusters, HPC systems and commercial cloud resources

- Enable users to define, create, manage, and control usage of collections of resources obtained via batch systems, HPC (XRAC) allocations, NSF CloudBank accounts or purchased directly from commercial cloud providers. We refer to such a collection as an personal **Annex**
- Interface with workload and workflow services to manage acquisition of resources (deployment of XPoints) from different providers of computing capacity

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

Open Science Grid

HTC CENTER FOR HIGH THROUGHPUT COMPUTING

# 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 eXecution Point (**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

# 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.

**750K core hours**

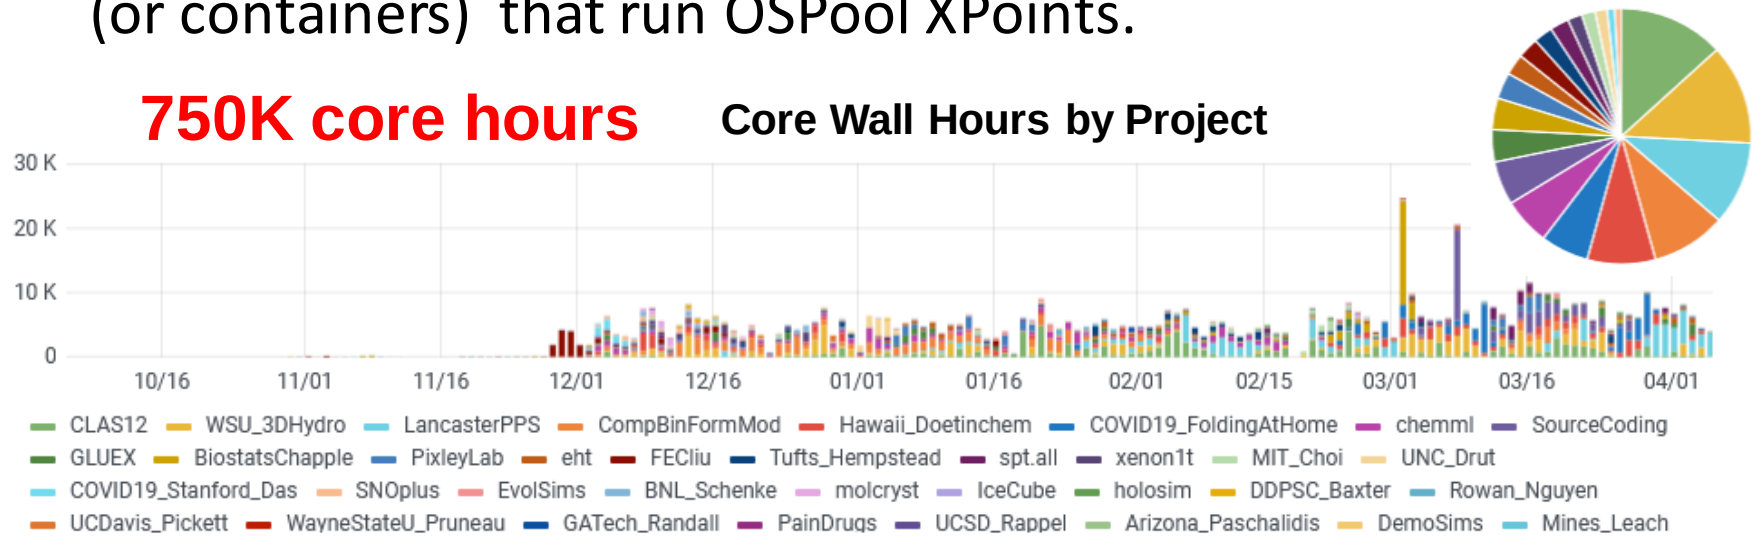


Work with NSF funded **Chameleon** is progressing nicely

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

Open Science Grid

HT CENTER FOR HIGH THROUGHPUT COMPUTING


**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**

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

Open Science Grid

HT CENTER FOR HIGH THROUGHPUT COMPUTING

# HTCondor dHTC CI

## of >51K GPUs in the Cloud @ SC19

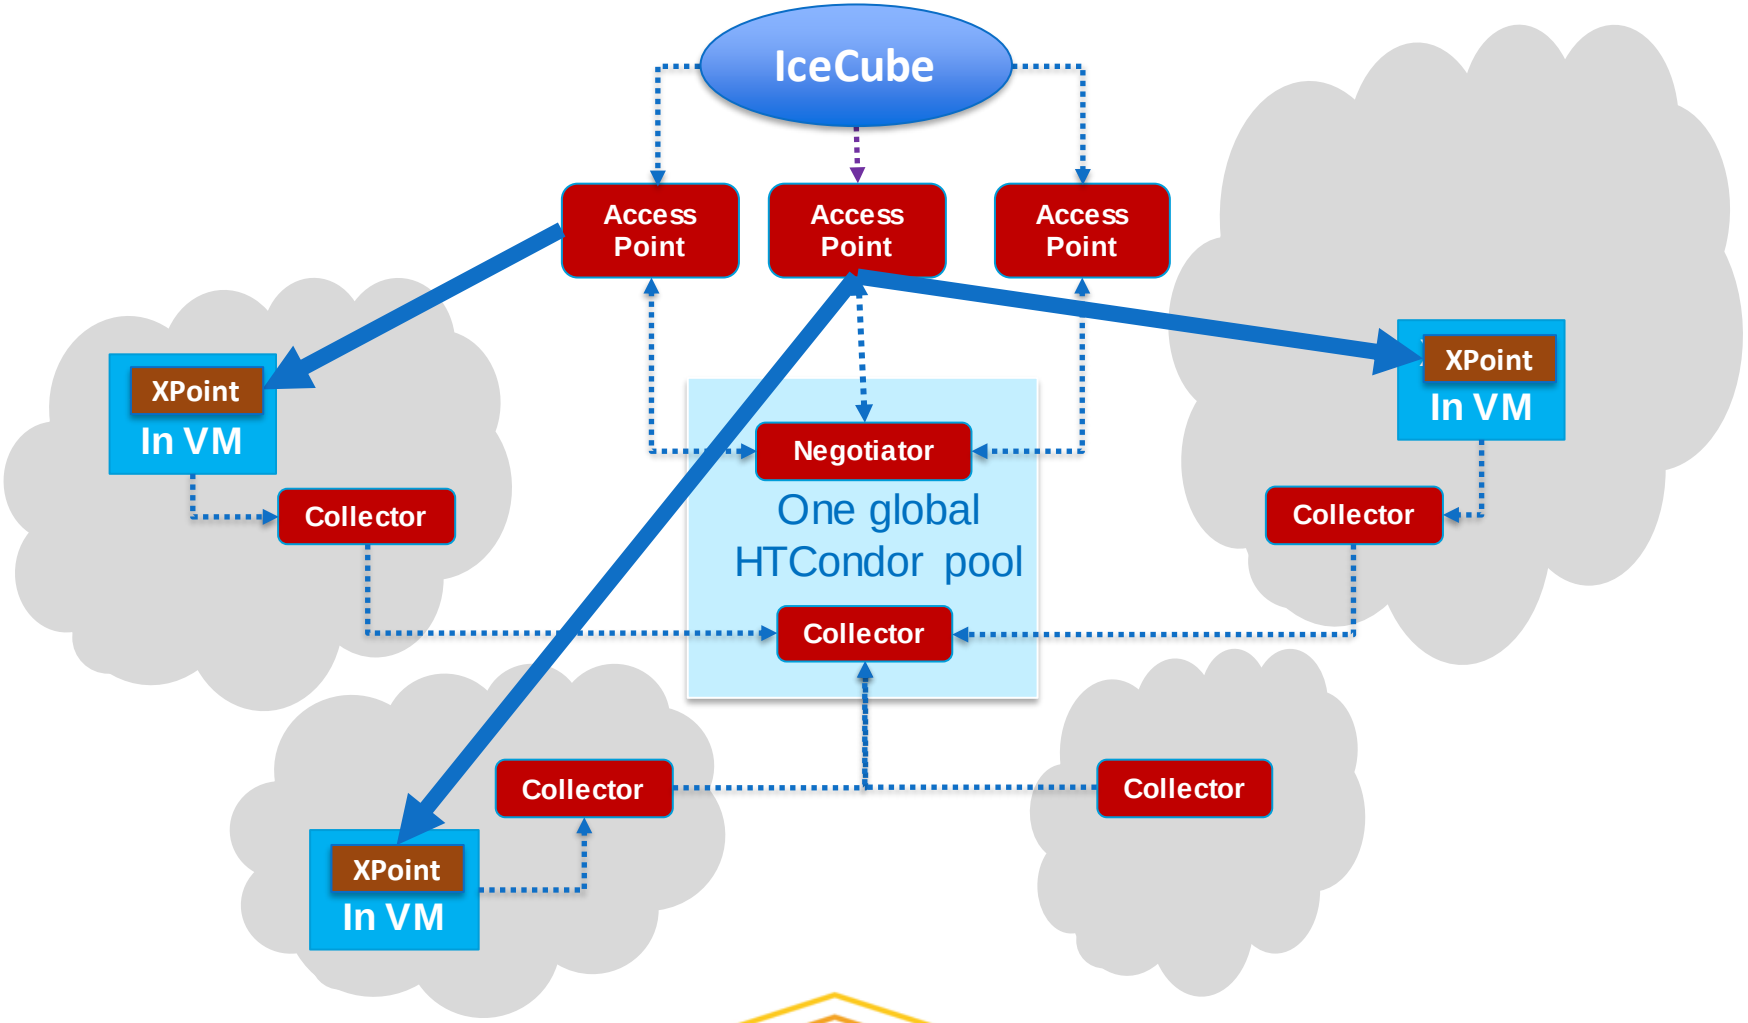


PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

Open Science Grid

HT CENTER FOR HIGH THROUGHPUT COMPUTING

# Thank you for building such

a wonderful and thriving
HTC community