The Open Science Pool – An OSG dHTC Service for the S&E Community

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

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

Open Science Grid

HTC 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 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 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

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
- Yesterday (04/20/2021), the OSPool completed ~200K jobs from 29 projects submitted by 33 users that consumed ~640K core hours

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 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

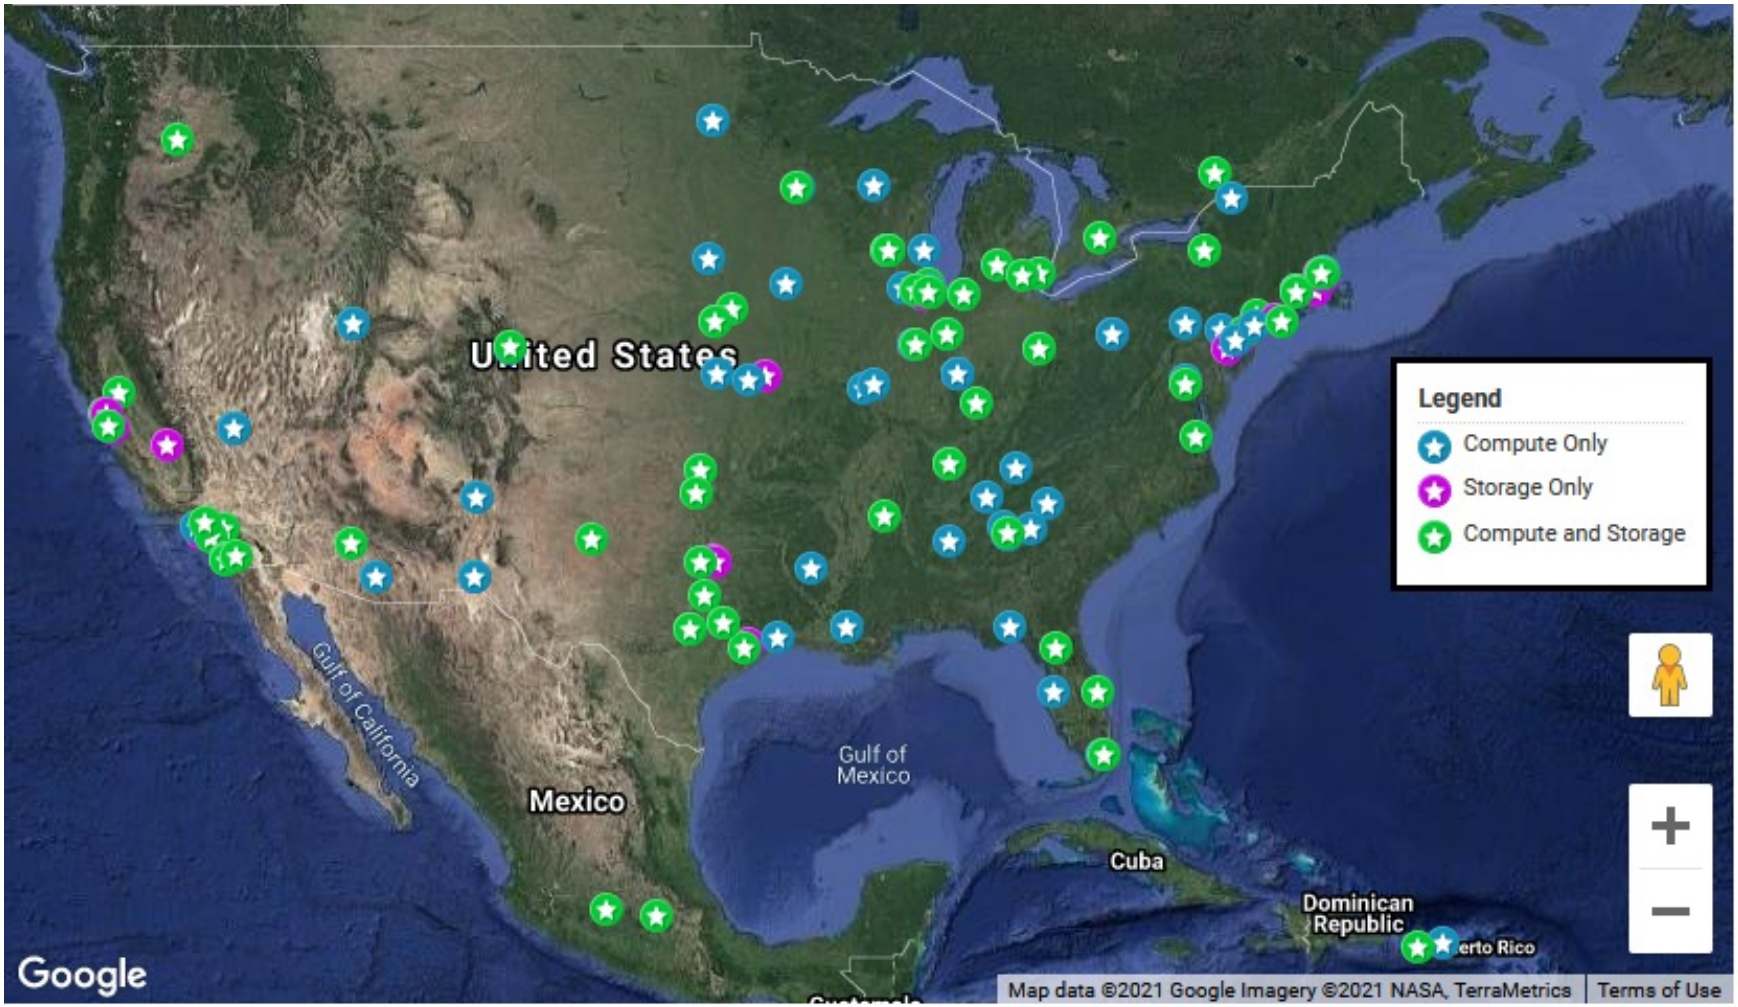

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

Open Science Grid

HTC CENTER FOR HIGH THROUGHPUT COMPUTING

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

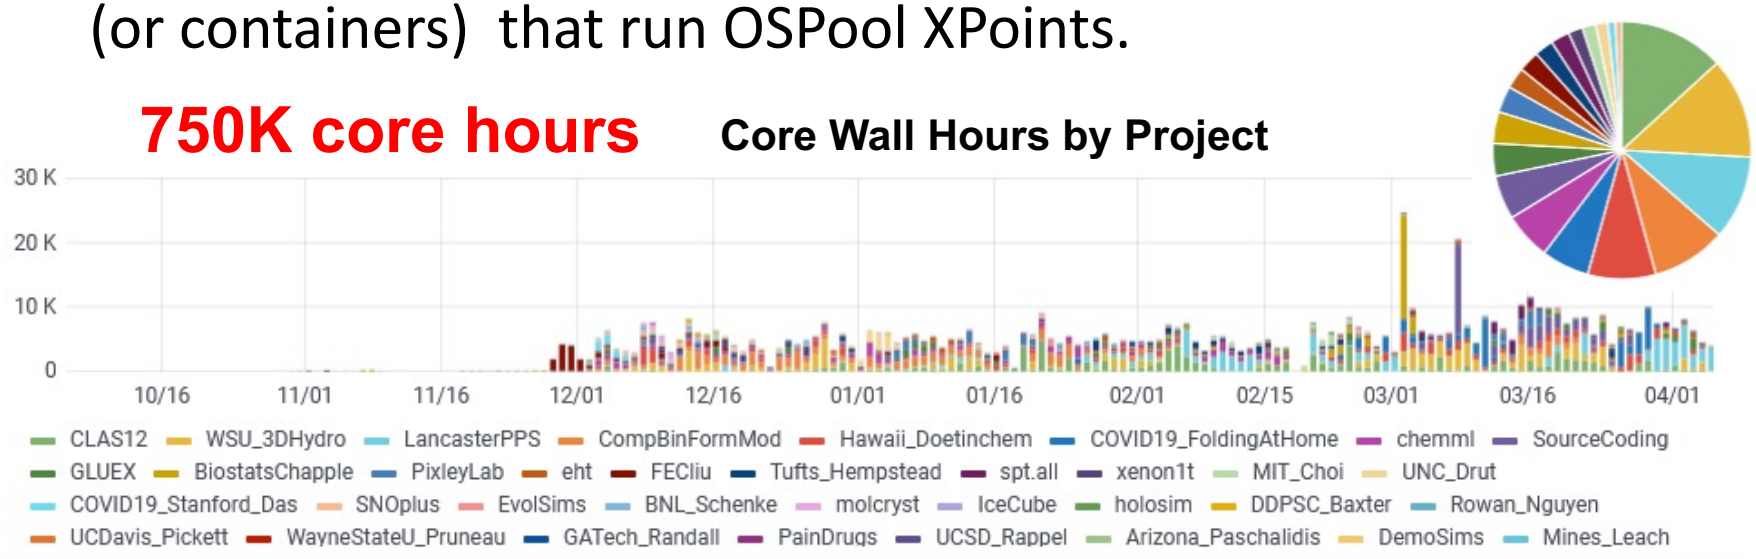

Work with NSF funded **Chameleon** is progressing nicely

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

Open Science Grid

HT CENTER FOR HIGH THROUGHPUT COMPUTING

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.

800K GPU hours

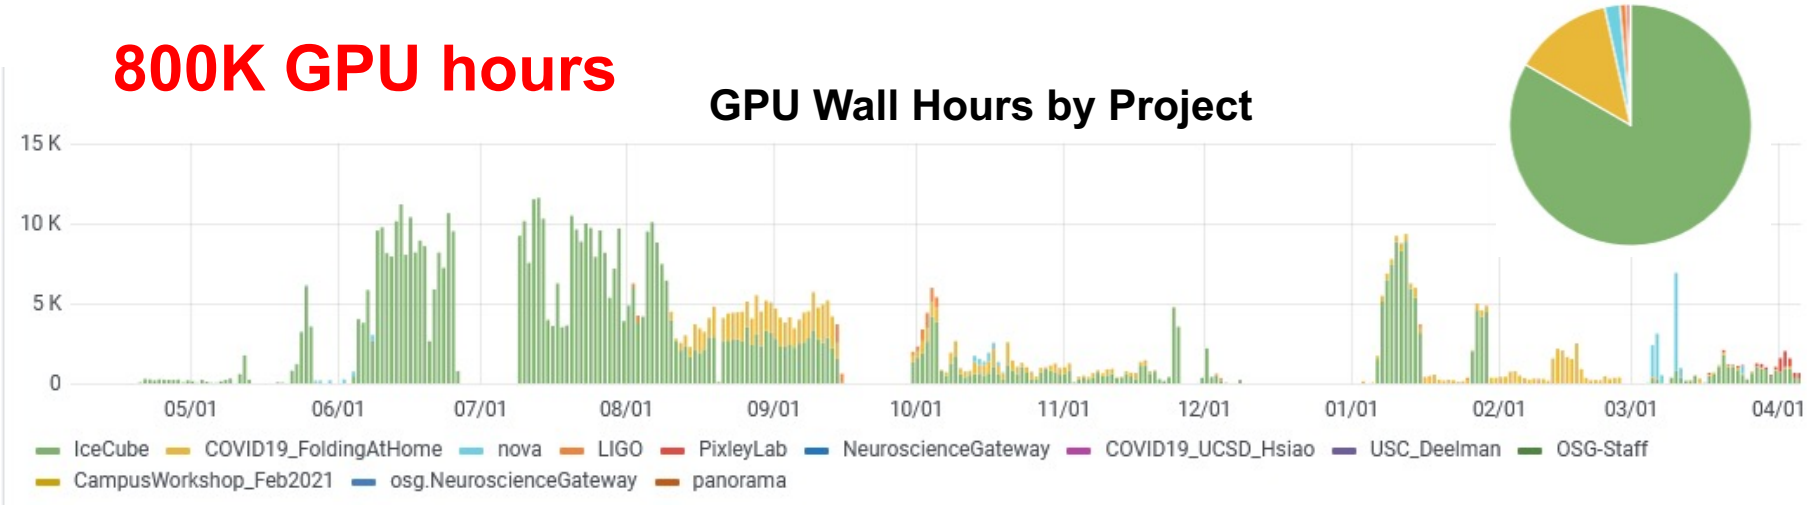

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

Open Science Grid

HT CENTER FOR HIGH THROUGHPUT COMPUTING

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

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

HTCondor dHTC CI

of >51K GPUs in the Cloud @ SC19

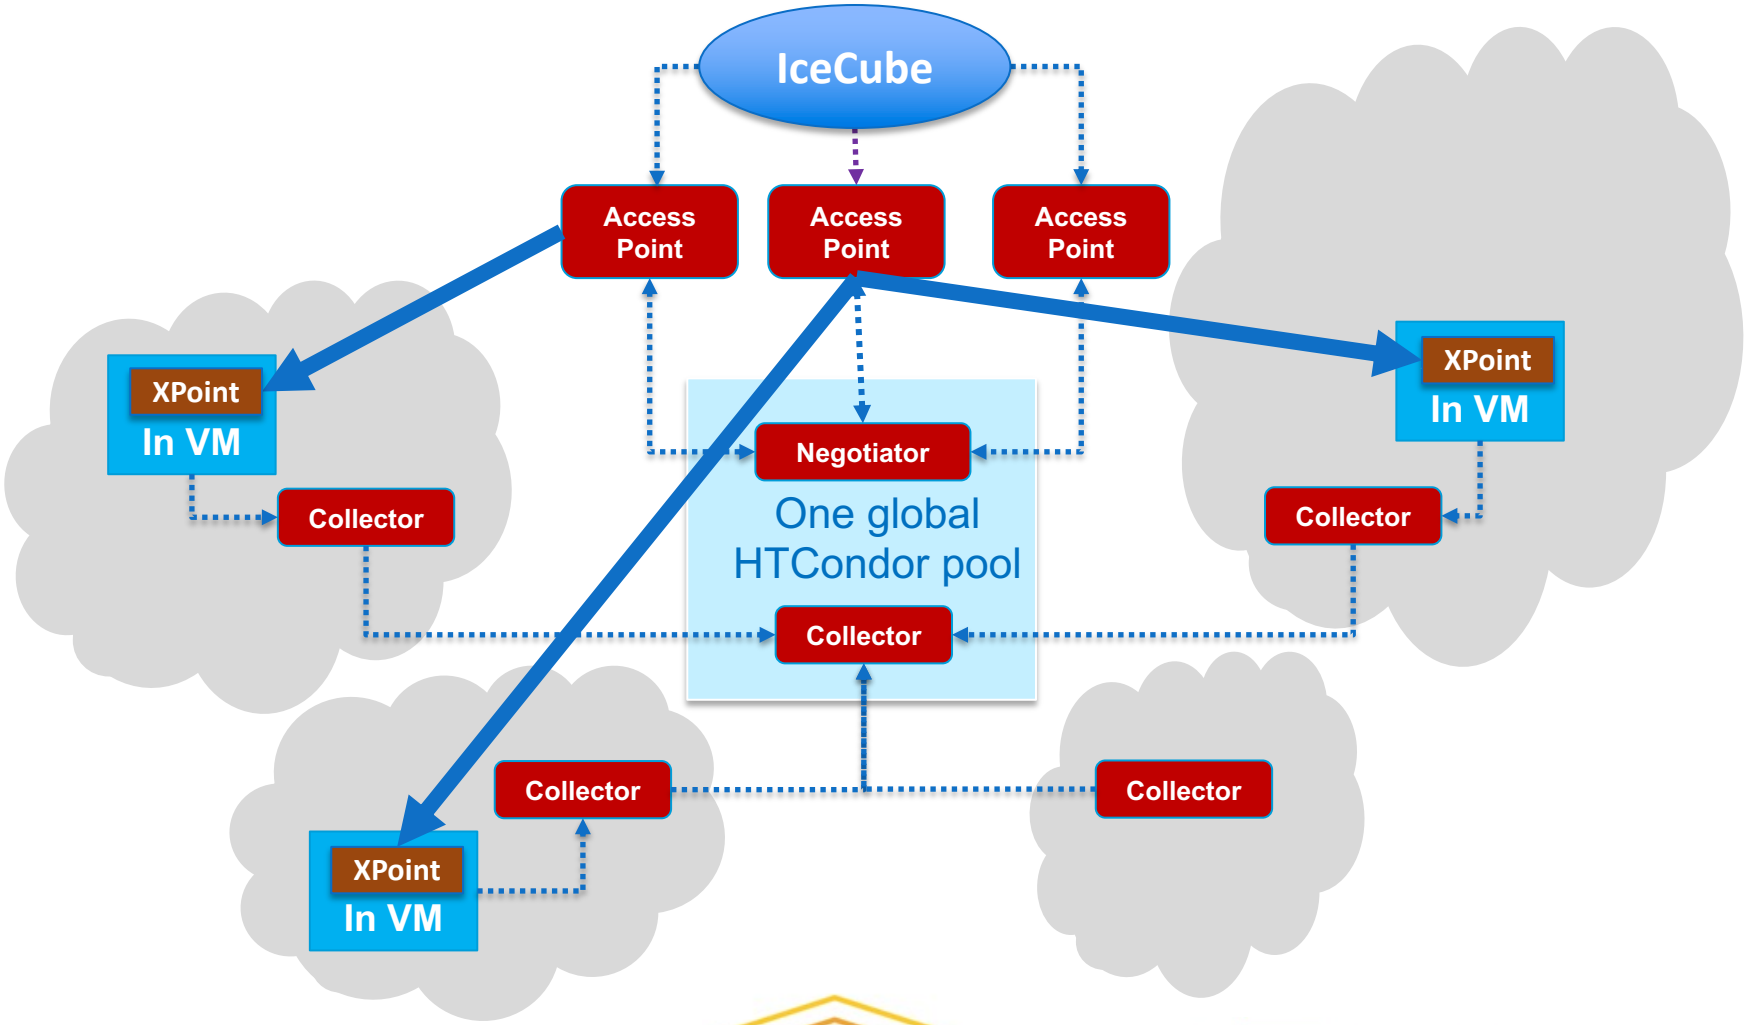

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

Open Science Grid

HT CENTER FOR HIGH THROUGHPUT COMPUTING

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

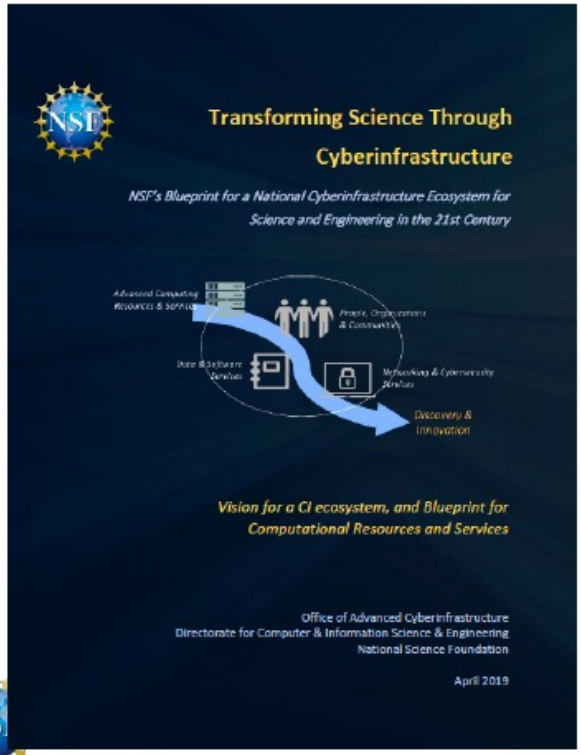

Computational Ecosystem: Elements

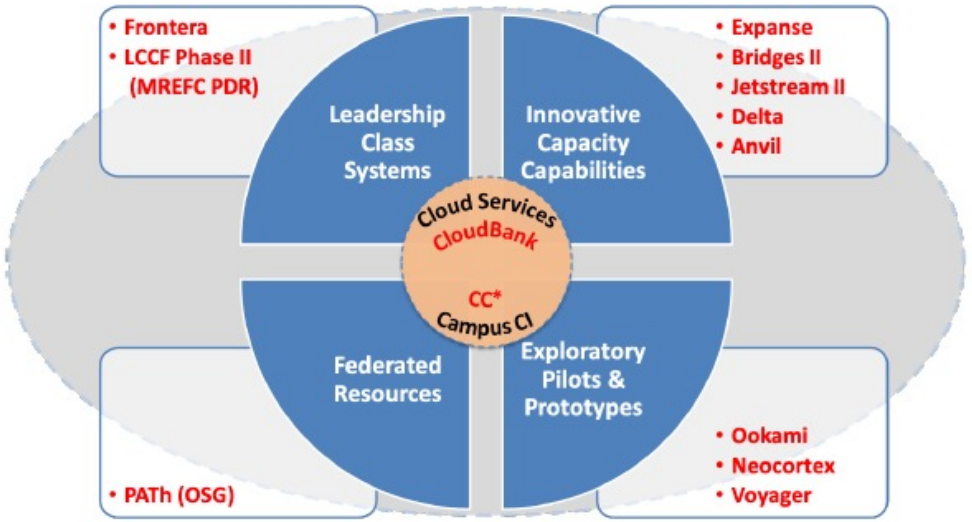

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

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