

Philosophy and Architecture

Miron Livny

John P. Morgridge Professor of Computer Science


Vials Research Professor

Director of CHTC

Technical Director of OSG



I founded The **UW-Madison Center for High Throughput Computing (CHTC)** in 2006 on the principles of Translational Computer Science where **innovation** and **services** that advance scientific discovery feed each other in a reenforcing partnership



Aga Khan Palace



"We should Magnify our own Errors"

"Kasturba trust would have even a bigger place in national service when freedom was there, than it had today. For all would go to the winds, if the women were not properly trained. Every worker in the trust would have a great deal to give to the new government. But we have got to generate that strength within

us. We may not be frightened of making mistakes. Man is born to make mistakes, but the great thing is to see out mistakes and learn from them. We should magnify our own errors, so as to be deterred from falling into them again. Those who make imagine that they never make mistakes are to be faked..."

- M K Gandhi

*Mohandas
Karamcha
nd Gandhi
(2 October
1869 - 30*




CHTC is hiring, view the new position on the jobs page and apply now! [View Job Posting](#)

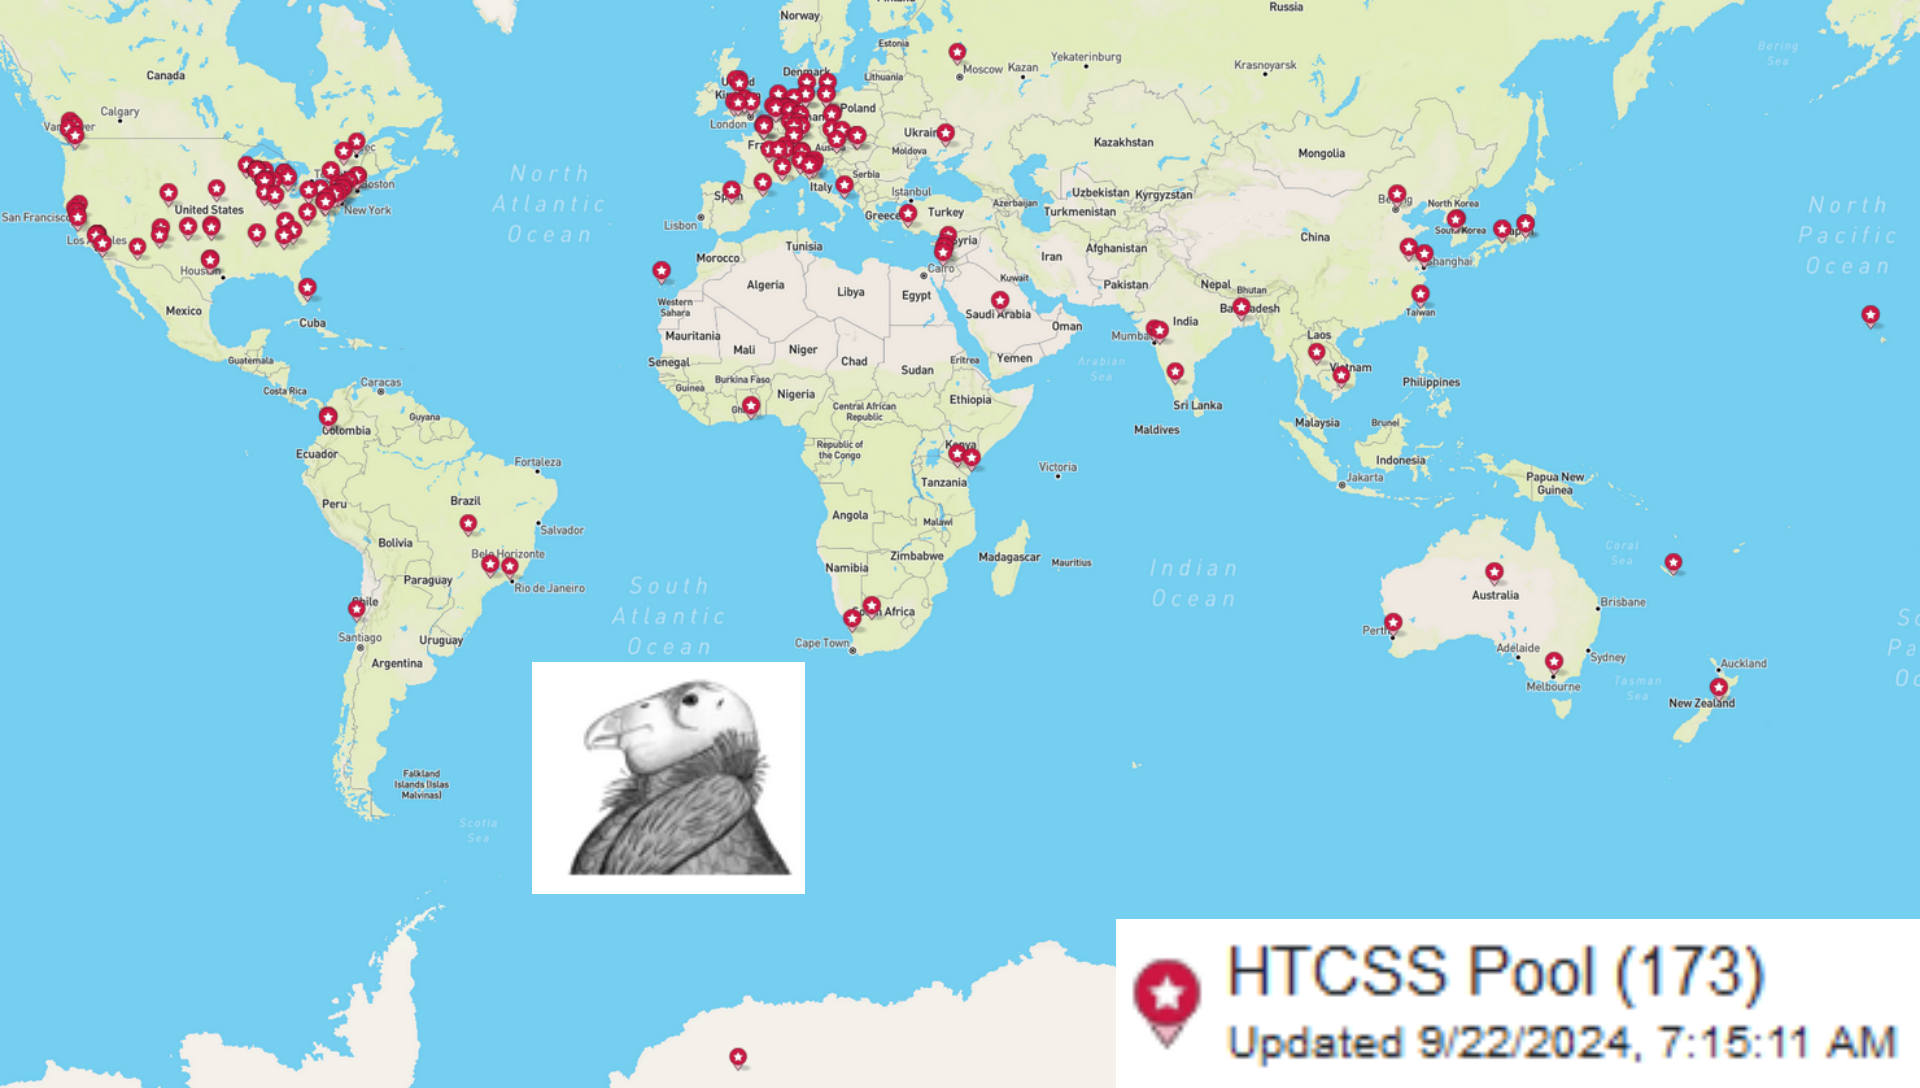
Last Year Serving UW-Madison



College	Projects Supported	HTC Core Years	HPC Core Years	Facilitator Interactions
Agricultural and Life Sciences	62	8,038	54	159
Education	5	216	-	10
Engineering	66	3,342	4,221	316
Law	1	181	-	1
Letters and Sciences	112	11,645	1,244	440
Medicine and Public Health	42	7,172	54	154
Off-Campus Collaborations	14	2,852	138	-
Pharmacy	5	75	258	7
Veterinary Medicine	5	21	-	-
Total	312	33,542	5,969	1,087

CHTC is home for the HTCondor Software Suite (HTCSS) and the Pelican Platform





The Open Science Pool (OSPool) is an HTC OSG service powered by the HTCondor Software Suite (HTCSS) that federates **Open** computing capacity contributions in support of **Throughput Computing** workloads



 OSPool Contributor
86 Sites, 54 Institutions

OSPool: Serving Open Science throughput computing

On August 30

830K jobs completed

Placed by **43 researchers**

Triggering **10M file transfers**

Consuming **673K core hours**

On August 30

830K jobs completed


Harnessing capacity from **58 institutions.**

OSP Pool

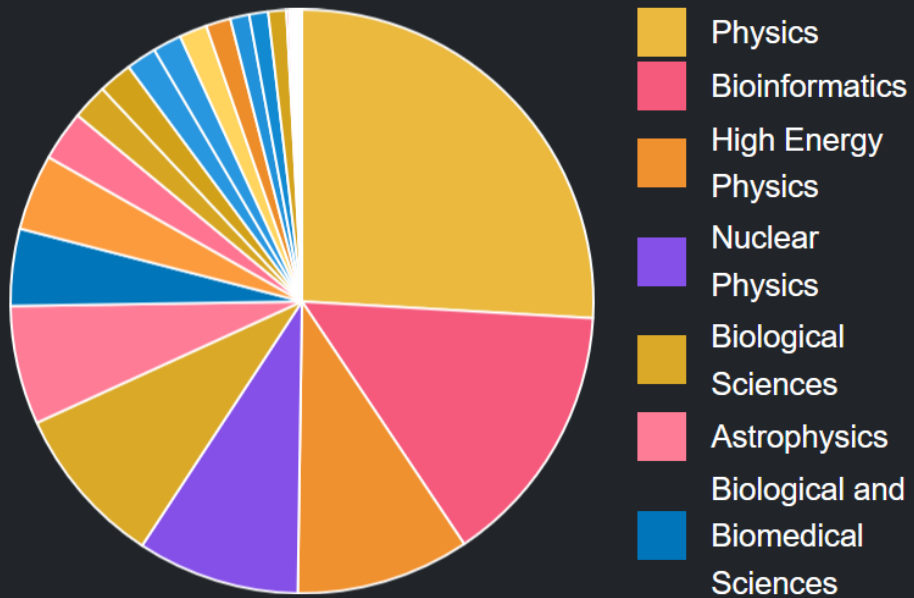
Throughput over the last 12 month

Month Starting	Jobs Completed	Core Hours	Files Transferred	Unique Users	Unique Projects	Unique Institutions Benefiting	Unique Institutions Contributing
TOTAL	192,392,377	293,101,942	4,886,141,481	509	225	101	68

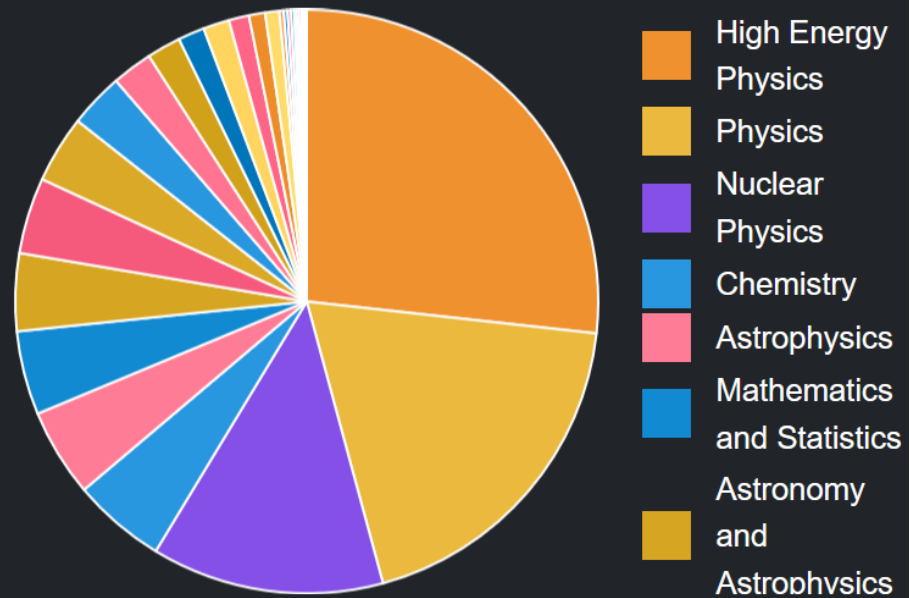
$192,392,377 / (365 * 24 * 60 * 60) = 6.10$ jobs per second
 $4,886,141,481 / (365 * 24 * 60 * 60) = 154.94$ file transfers per second




Fields of Science by Jobs ?



Fields of Science by CPU Hours ?





Bring data to processing capacity Store results at Object Stores



US Map featuring the locations of current OSDF architectural components.

Open Science Data Federation

The Open Science Data Federation (OSDF) is an OSG service designed to support the sharing of files staged in autonomous “origins”, for efficient access to those files from anywhere in the world via a global namespace and network of caches. The OSDF may be used either standalone - allowing

- There are 9 namespaces like this, and all 9 belong to international collaborations
=> See Panel Discussion Tuesday Afternoon

name	Read	Unique data
LIGO IGWN	40 PB	203 TB
IceCube	10 PB	66 TB
LIGO users	4 PB	28 TB
IGWN shared	1.7 PB	11 TB
KOTO	8 PB	3.5 TB

name	Read	Unique data
Einstein Telescope	1.5 PB	3.2 TB
Nova	5 PB	3 TB
MicroBoone	12 PB	1.7 TB
IGWN CIT	17 PB	1.2 TB

Gravitational Wave Observatories Community dominates unique data

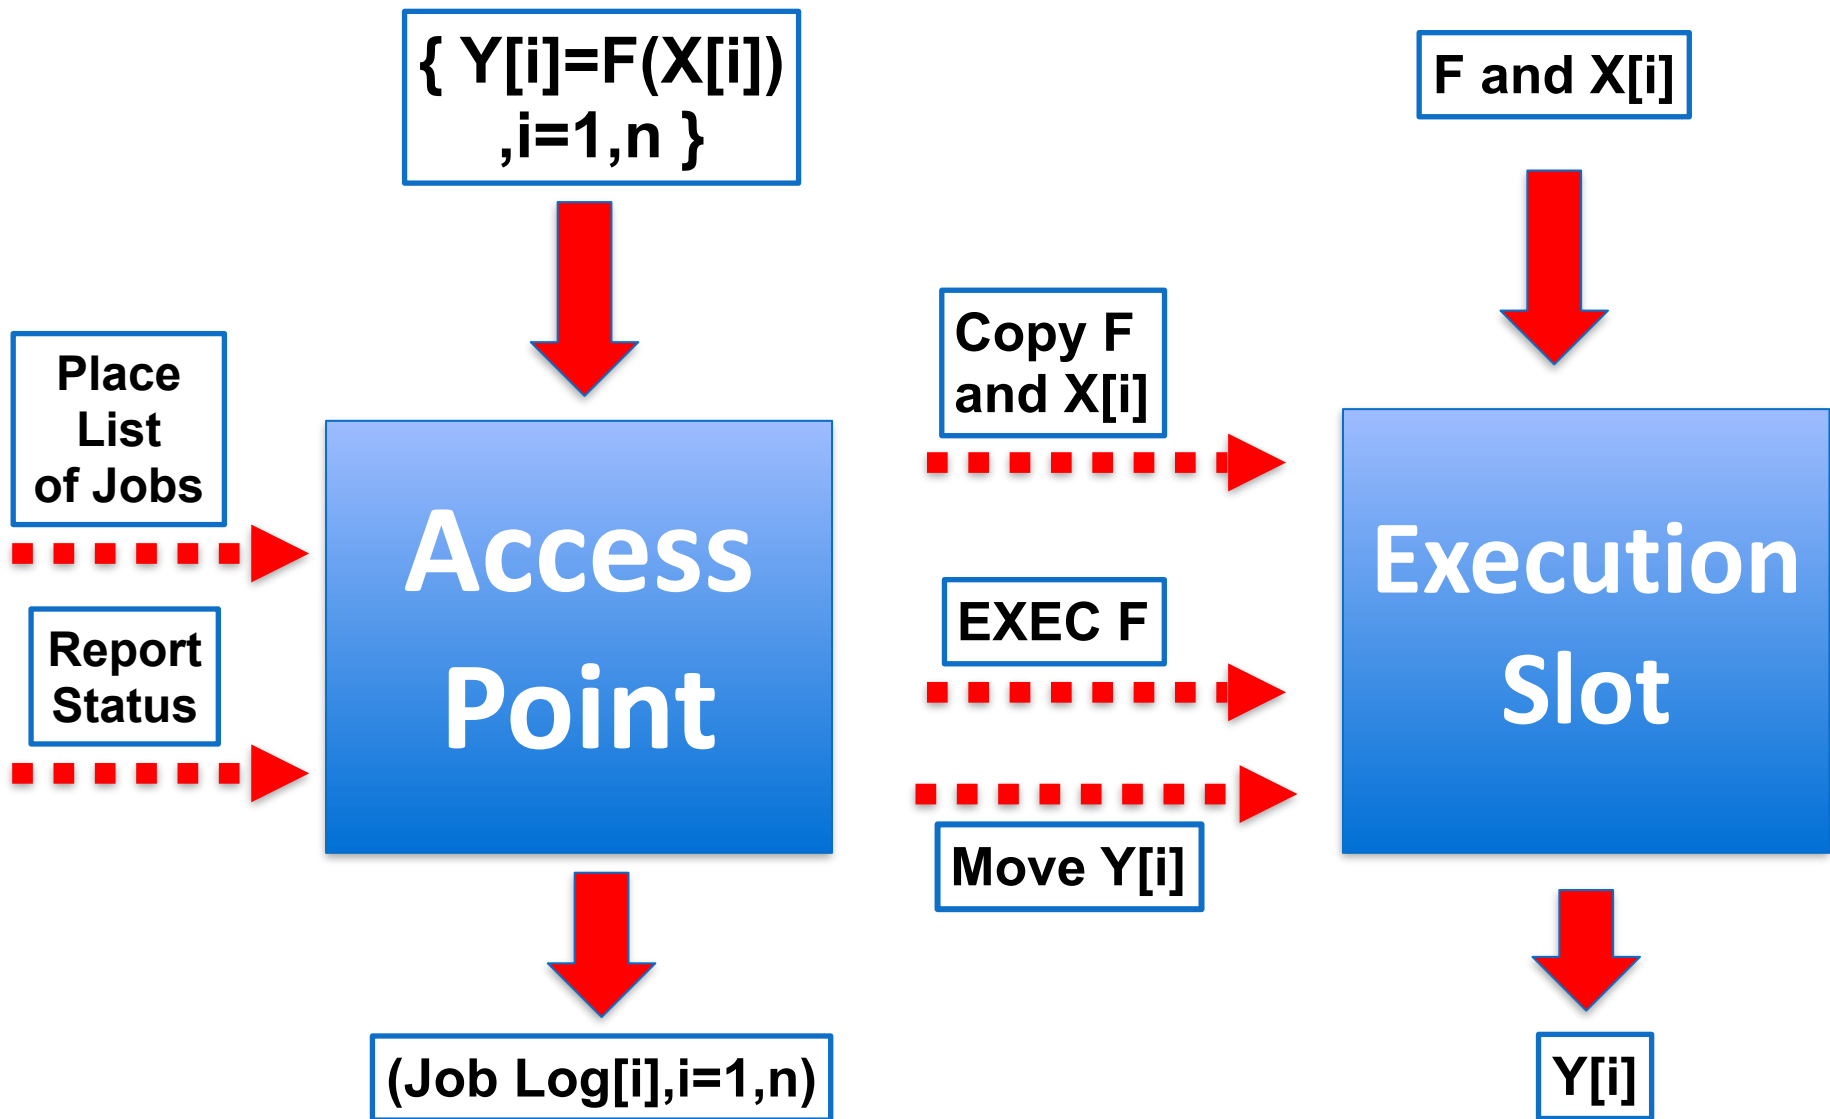
Next come neutrino physics experiments (IceCube, Nova, MicroBoone)


The words of Leonardo di ser Piero da Vinci

(15 April 1452 - 2 May 1519)

*Simplicity is
the ultimate
sophistication*





The five years  project – is a Partnership between the UW-Madison Center for High Throughput Computing (**CHTC**) – home of the HTCondor Software Suite (HTCSS) and the **OSG** Consortium – provider of the **OSG fabric of distributed High Throughput Computing (dHTC)** services with a national scale mission

“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.”

PATh Proposal 04/21/2020

Aligned with NSF Cyberinfrastructure blueprint



“... many fields today rely on **high-throughput computing for discovery.**”

Committee on Future Directions for NSF Advanced Computing Infrastructure to Support U.S. Science in 2017-2020; Computer Science and Telecommunications Board; Division on Engineering and Physical Sciences; National Academies of Sciences, Engineering, and Medicine



*The words of Koheleth son of David, king
in Jerusalem ~ 200 A.D.*

Only that shall happen

Which has happened,

Only that occur

Which has occurred;

There is nothing new

Beneath the sun!



Ecclesiastes, (קֹהֵלֶת, *Kohelet*, "son of David, and king in Jerusalem" alias Solomon, Wood engraving Gustave Doré (1832–1883)

Claims for “benefits” provided by Distributed Processing Systems

P.H. Enslow “What is a Distributed Data Processing System?” IEEE Computer, January 1978

- High Availability and Reliability
- High System Performance
- Ease of Modular and Incremental Growth
- Automatic Load and Resource Sharing
- Good Response to Temporary Overloads
- Easy Expansion in Capacity and/or Function

Confucius

(traditionally 28 September 551 B.C.
– 479 B.C.)

- 名不正，則言不順
- If names be not correct, language is not in accordance with the truth of things.
- Paraphrased as a Chinese proverb "**The beginning of wisdom is to call things by their proper name.**"



HTCondor Access Points (APs) are capable and dependable handlers of Throughput Computing workloads

- Large job ensembles
- Bring Your Own Capacity (BYOC)
- Support Job-Lists presented as tables
- Job dependencies defined by Directed Acyclic Graphs (DAG)
- Workflows managed by the Pegasus WFMS
- Supports access to remote datasets via HTTP and OSDF

An **HTCondor Execution Point (EP)** is hosted by a server and represents the capacity of the server to HTCondor Access Points through a collection of **HTCondor Execution Slots (ES)**.

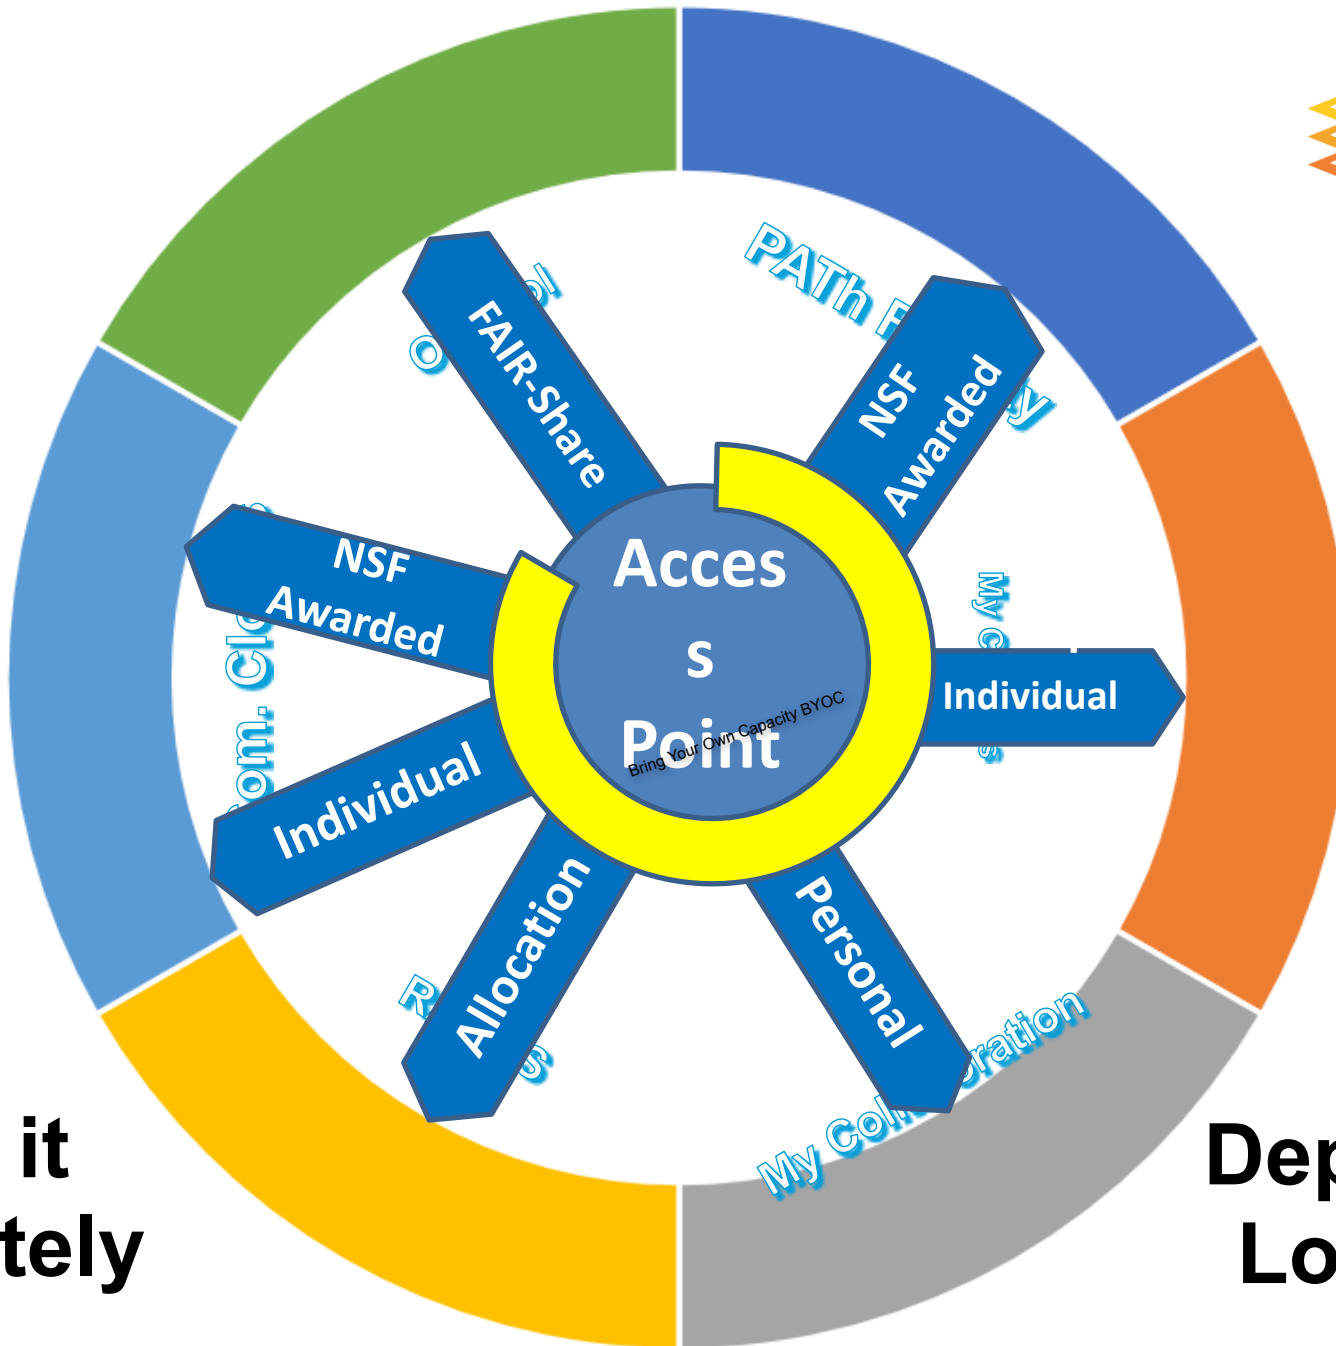
- Number and capacity of Slots is managed by by EP.
- EP is responsible of isolation of ES

A collection of Execution Points forms a **HTCondor Pool**

**HTCondor Pool +
HTCondor AP
= HTCondor System**


IGWN HTCondor Pool is deployed and maintained by the PATH operation team using the OSG Fabric of dHTC services

- **IGWN APs** are deployed and operated by CalTech.
- The AP of the IUACC HTCondor System can be configured to “**flock**” to the IGWN HTCondor Pool



**Use it
Remotely**

**Deploy it
Locally**





HTCondor



The IGWN Computing Grid is based on top of [HTCondor](#), a specialized workload management system for compute-intensive jobs. HTCondor is used to specify discrete work units (jobs) you want completed that are then distributed across the available resources with sophisticated scheduling, prioritisation, monitoring, and reporting capabilities.

Basic HTCondor Usage

For an excellent introduction to HTCondor usage, please refer to the official HTCondor Users' Manual:

Copyright © 2021 the KAGRA Collaboration, the LIGO Scientific Collaboration, and the Virgo Collaboration



Made with IGWN Material for MkDocs .

Last updated September 06, 2024.

IGWN AP Weekly Usage Report 2024-09-09

	User	Num Uniq Job Ids	All CPU Hours	% Good CPU Hours	Job Unit Hours	% Ckpt Able	% Rm'd Jobs
37	TOTAL	955,907	1,499,007	80.9	3,000,868	1.0	4.0

IGWN AP Weekly Usage Report 2024-08-26

AP Weekly Usage Report 2024-08-19

	User	Num Uniq Job Ids	All CPU Hours	% Good CPU Hours
24	TOTAL	486,450	942,926	73.3

	User	Num Uniq Job Ids	All CPU Hours	% Good CPU Hours
22	TOTAL	305,268	439,270	86.2

Collaborations

Scientific collaborations from across the globe rely upon HTCSS to advance their research.

Statistics from these pools generated in the last 365 days.

<p>CMS ☺</p> <p>19,308,785 Jobs</p> <p>1,257,389,376 CPU Core Hours</p> <p>62,087 GPU Hours</p>	<p>REDTOP ☺</p> <p>7,538,162 Jobs</p> <p>42,965,331 CPU Core Hours</p>	<p>IceCube ☺</p> <p>41,916,525 Jobs</p> <p>69,286,625 CPU Core Hours</p> <p>1,478,056 GPU Hours</p>	<p>IGWN ☺</p> <p>9,517,636 Jobs</p> <p>74,648,105 CPU Core Hours</p> <p>565,549 GPU Hours</p>
<p>South Pole Telescope ☺</p> <p>1,522,901 Jobs</p> <p>1,267,722 CPU Core Hours</p>	<p>Xenon ☺</p> <p>11,071,788 Jobs</p> <p>5,380,927 CPU Core Hours</p>	<p>Future Colliders ☺</p> <p>7,160 Jobs</p> <p>42,588 CPU Core Hours</p>	<p>KOTO ☺</p> <p>9,680,807 Jobs</p> <p>25,812,590 CPU Core Hours</p>
<p>CLAS12 ☺</p> <p>6,355,967 Jobs</p> <p>50,507,272 CPU Core Hours</p>	<p>GLUEX ☺</p> <p>1,218 Jobs</p> <p>715 CPU Core Hours</p>	<p>EIC ☺</p> <p>547 Jobs</p> <p>37 CPU Core Hours</p>	<p>MOLLER ☺</p> <p>42,515 Jobs</p> <p>57,441 CPU Core Hours</p>
<p>DUNE ☺</p> <p>20,356,384 Jobs</p> <p>45,653,147 CPU Core Hours</p> <p>11 GPU Hours</p>	<p>EHT ☺</p> <p>2,481,012 Jobs</p> <p>2,801,633 CPU Core Hours</p>	<p>ePIC ☺</p> <p>4,269,301 Jobs</p> <p>13,050,920 CPU Core Hours</p>	



<https://path-cc.io/contact/>



Contact

PATH is a unique partnership between the Center for High Throughput Computing (CHTC) and the OSG Consortium.

- For enquiries about the *PATH project*, please contact the [PATH leadership](#).
- For help with *CHTC technologies* such as the HTCondor Software Suite (HTCSS), contact chtc@cs.wisc.edu.
- *Campuses* interested in providing resources to the [Open Science Pool \(OSPool\)](#) can contact support@osg-htc.org
- *Users* interested in using an Access Point to leverage resource like the OSPool can contact support@osgconnect.net.
- *PIs* interested in getting credit accounts on PATH-managed hardware should visit the [dedicated page](#).

This work is supported by the National Science Foundation under Cooperative Agreements OAC-2030508, OAC-2331480. Any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.