PATh Facility: Dedicated HTC Capacity

Brian Lin, OSG Software Area Coordinator
Center for High Throughput Computing
University of Wisconsin–Madison

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

HT CENTER FOR HIGH THROUGHPUT COMPUTING

Pictured: poorly photoshopped image of a farcical PATh Facility informercial

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

OSG

HT CENTER FOR HIGH THROUGHPUT COMPUTING

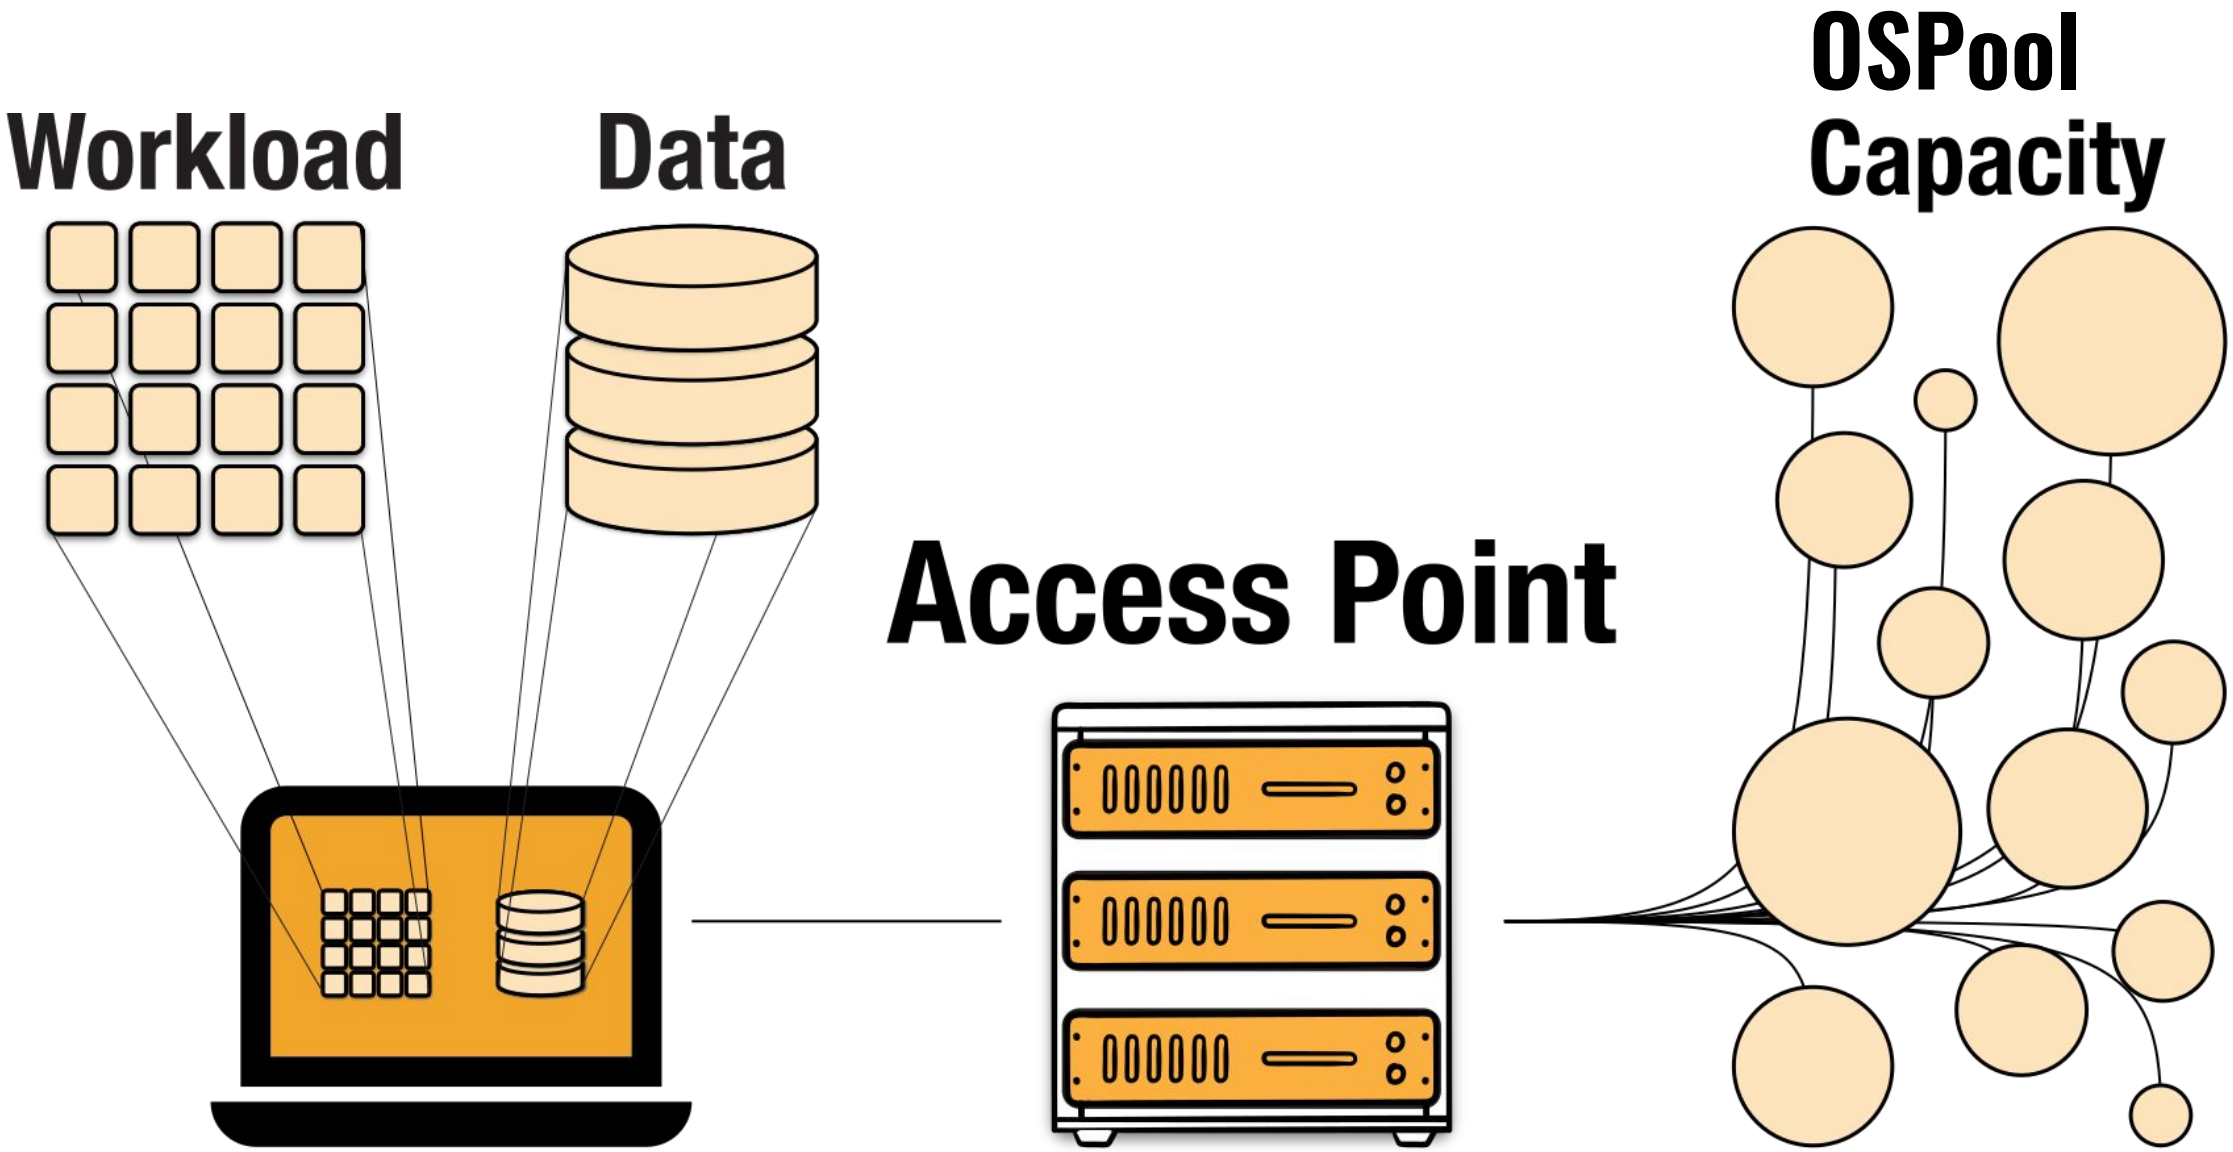
Workload
Data
Access Point
OSPool
Capacity

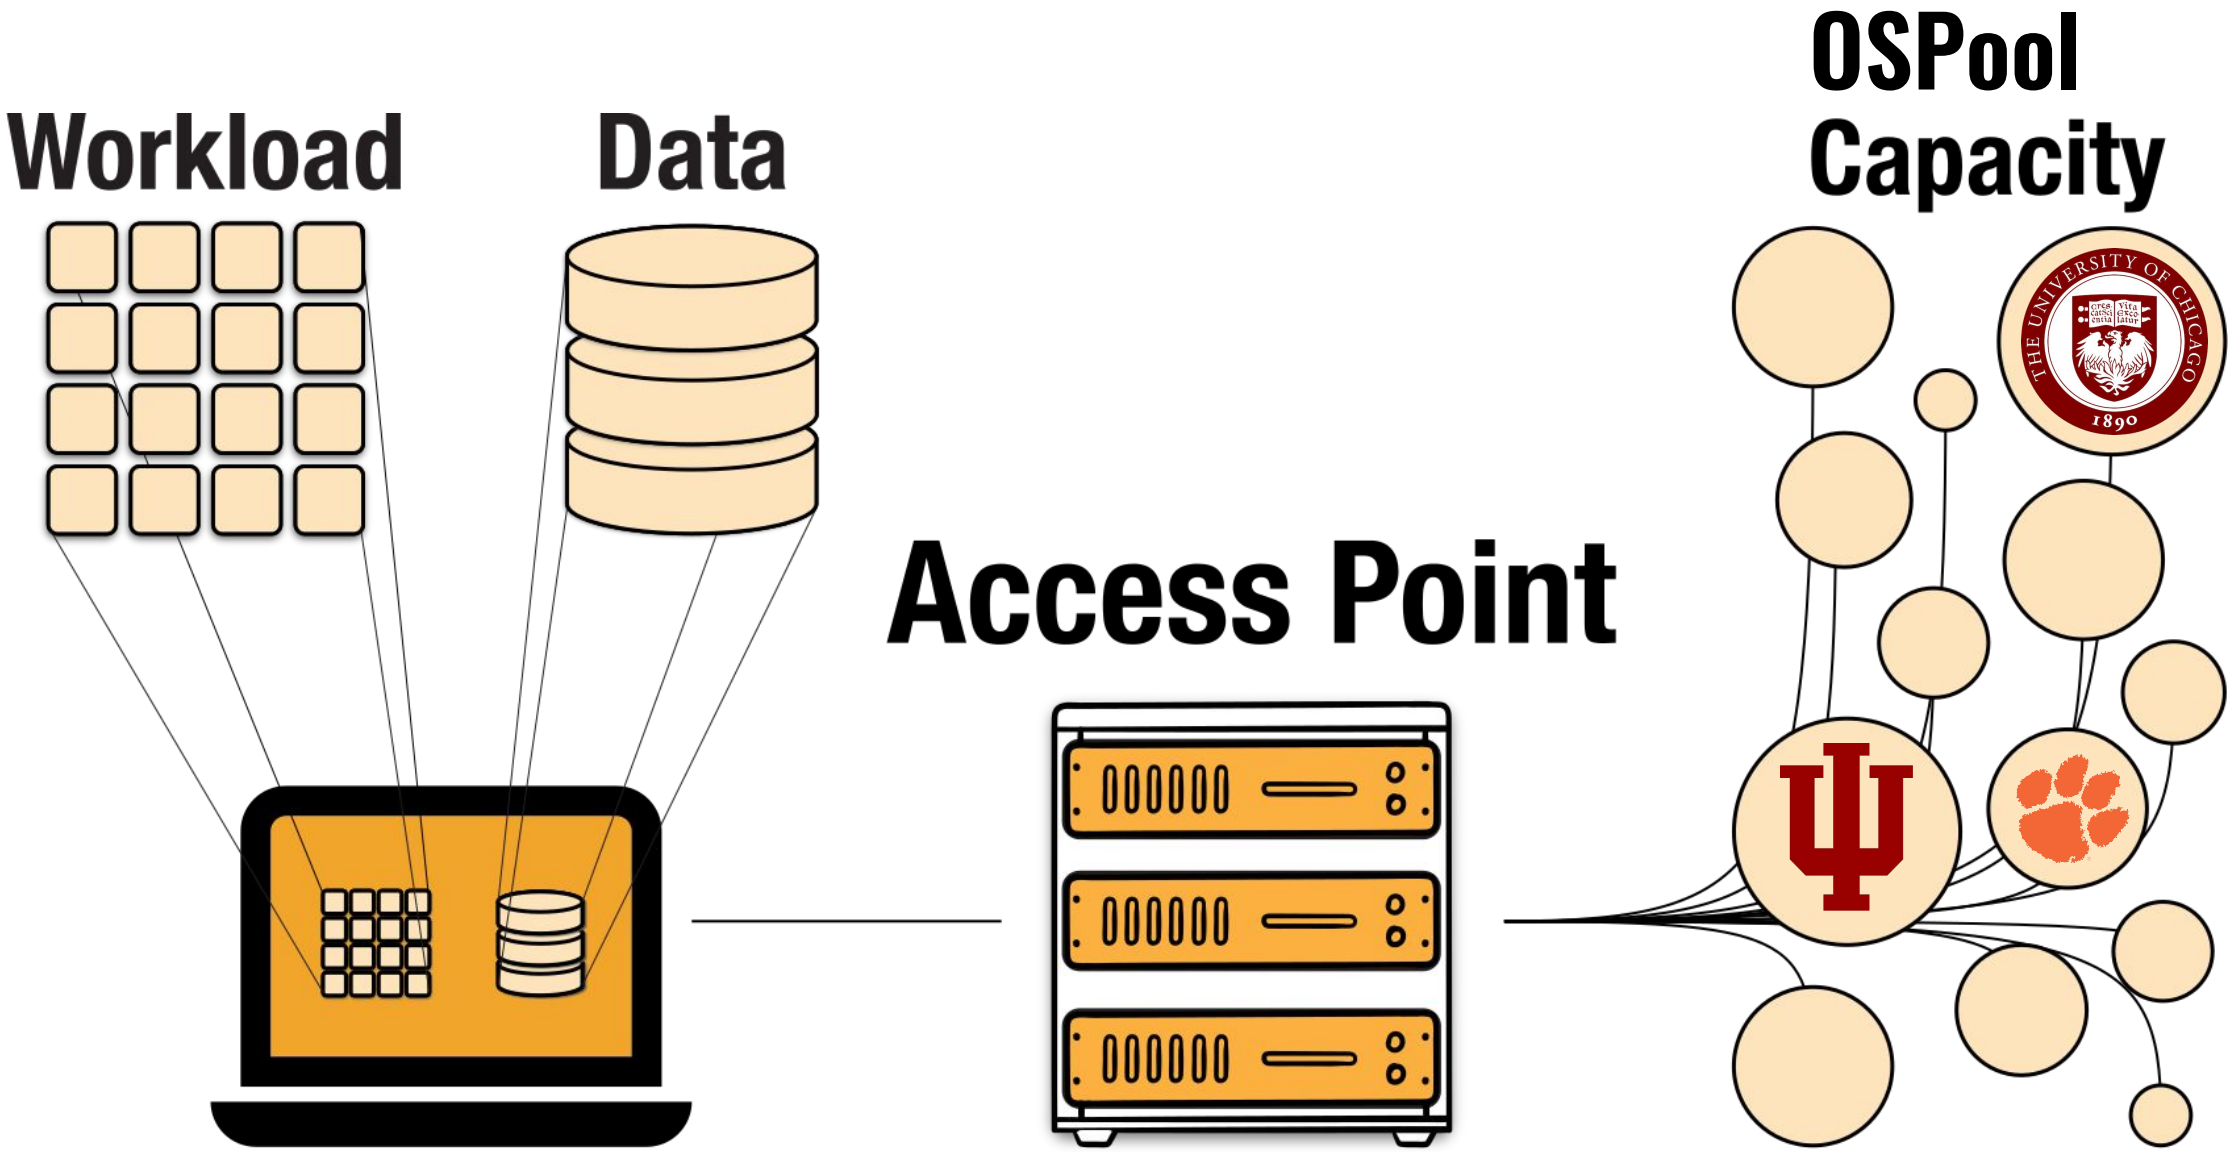

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

OSG

HT CENTER FOR HIGH THROUGHPUT COMPUTING

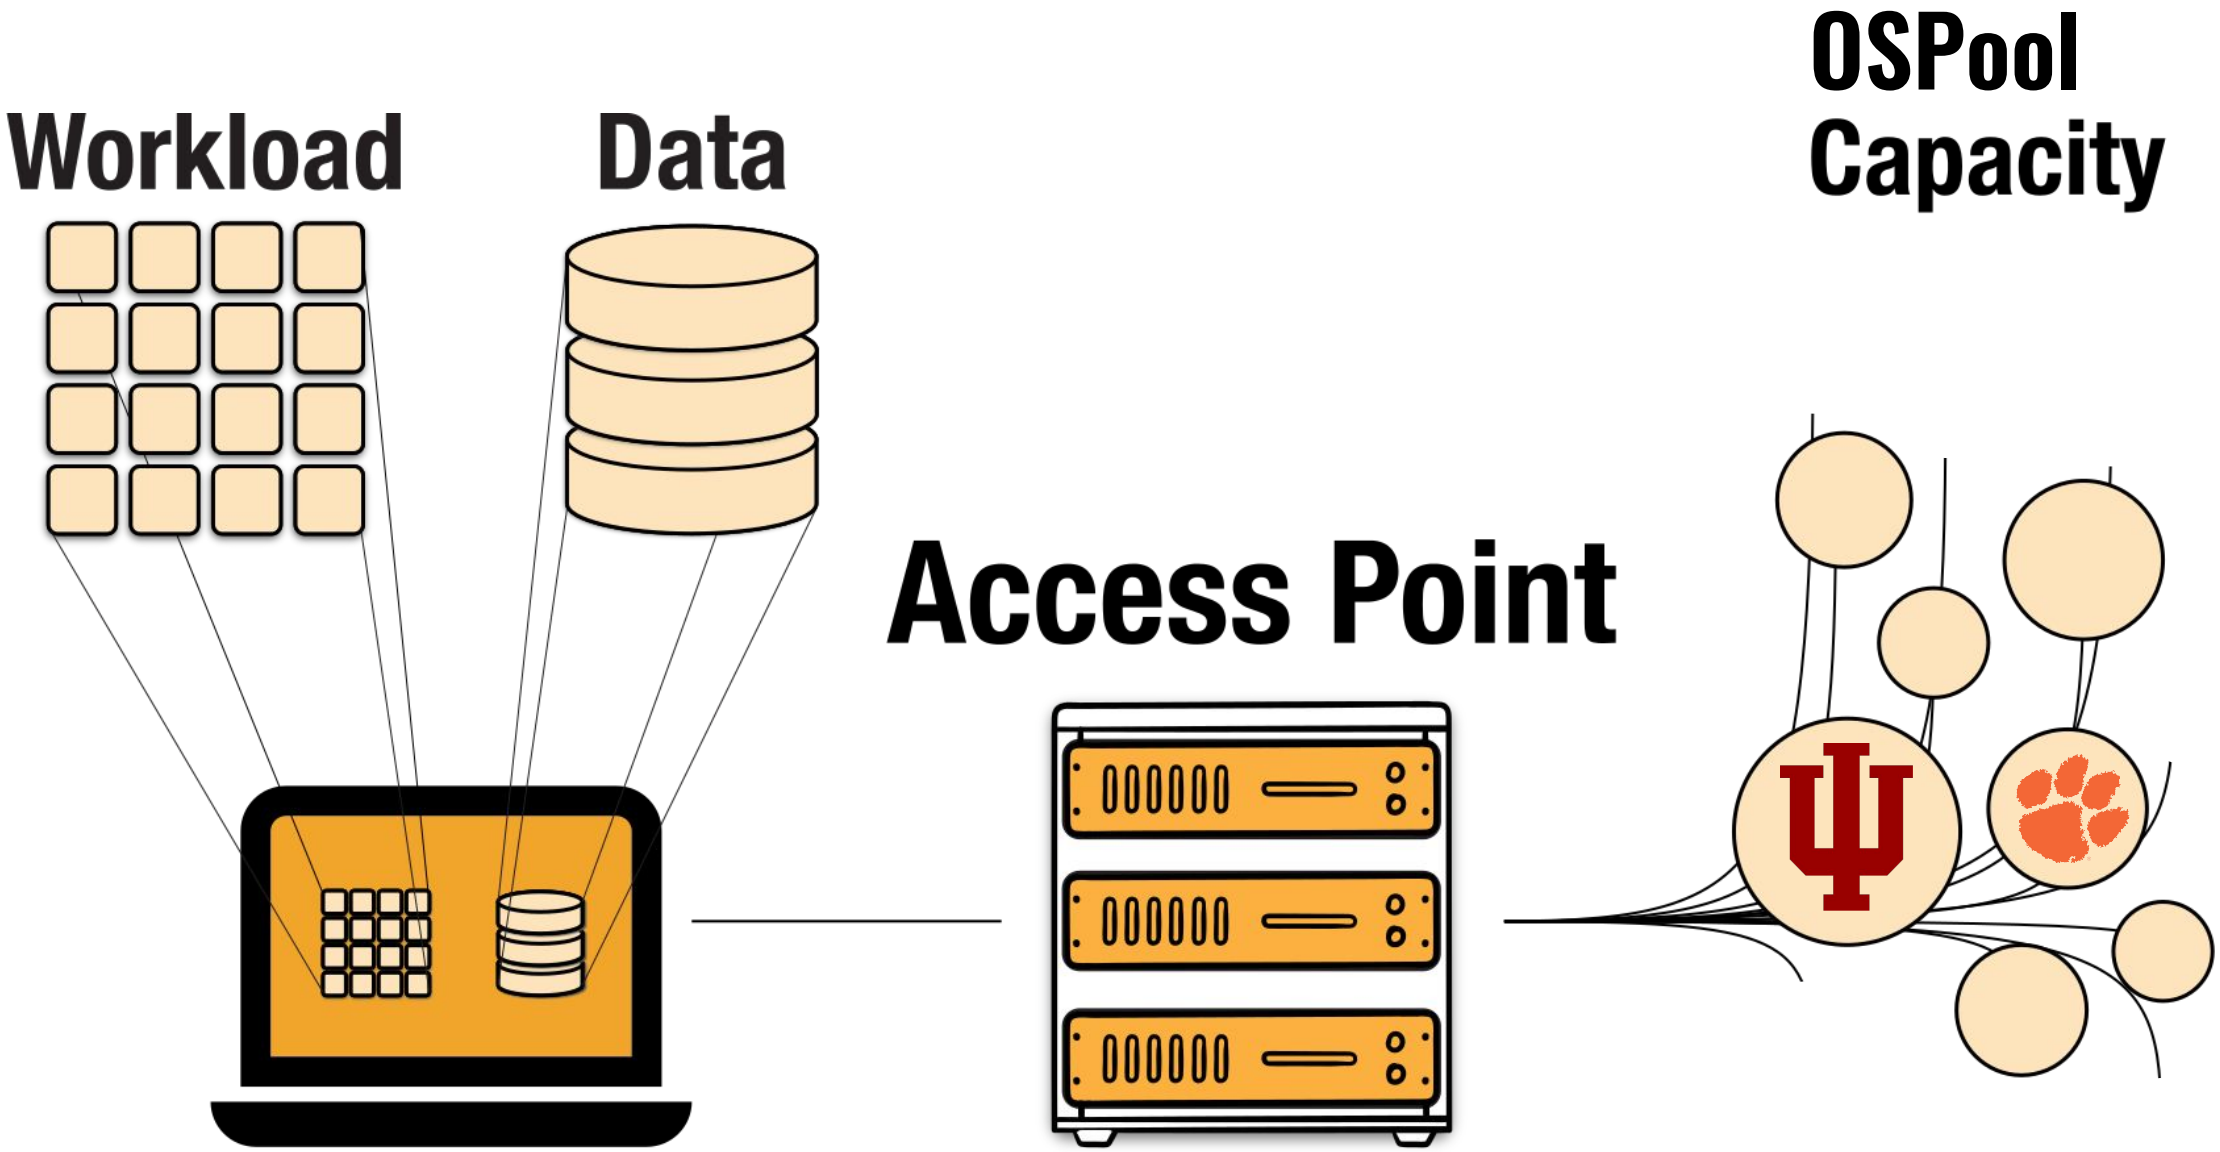
Workload
Data
OSPool
Capacity
Access Point

The PATh Facility

Powering distributed high throughput computing

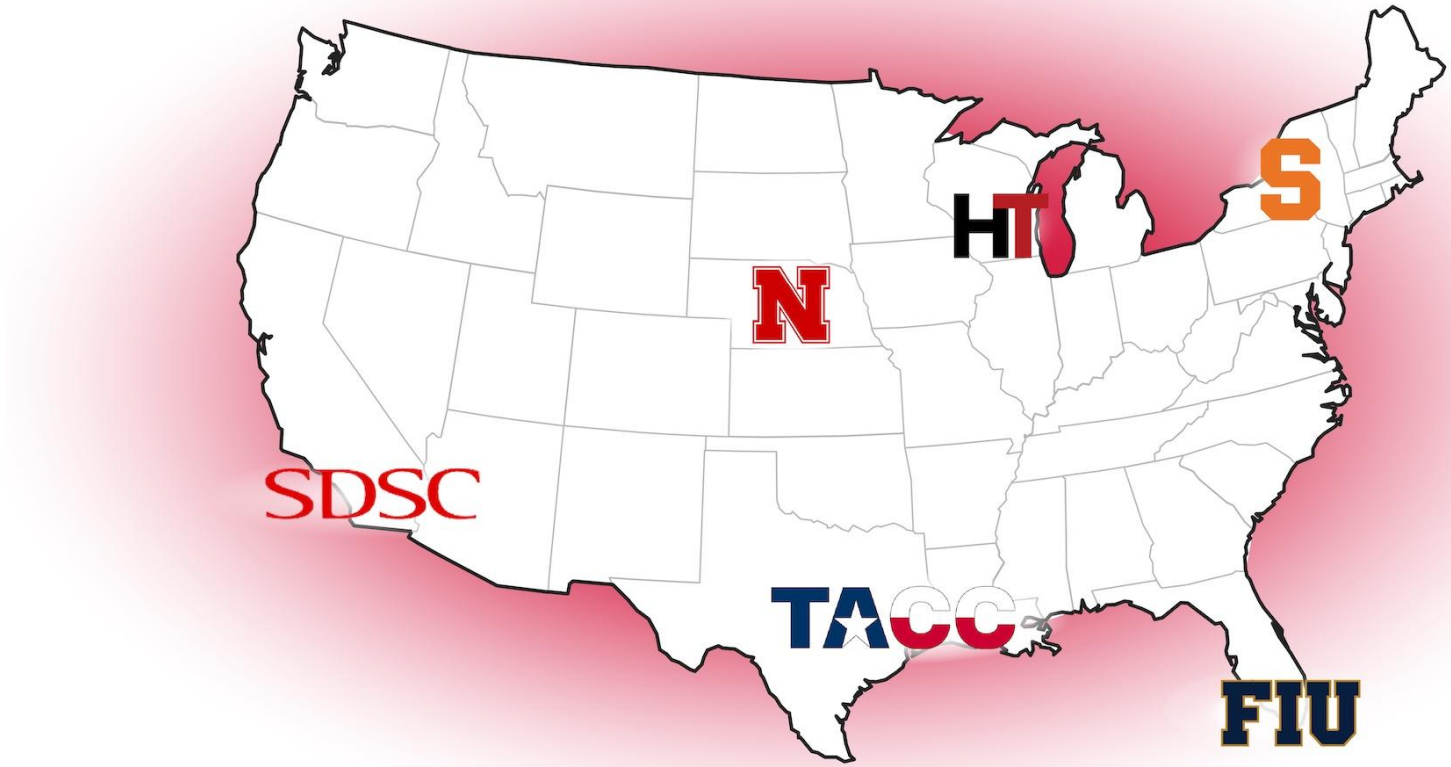

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

OSG

HT CENTER FOR HIGH THROUGHPUT COMPUTING

PATh Facility

- **Dedicated** distributed High-Throughput Computing
  → Longer job run times!
- Computational time charged to a credit accounting system
- Same HTCondor-based system as the OSPool
- Same friendly facilitators, operators, and developers that support the OSPool





Pictured: OSG/PATh Facilitation Team

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

OSG

HT CENTER FOR HIGH THROUGHPUT COMPUTING

Over 30,000 cores total!

128 and 80 cores per machine!

2-4 GB RAM per core!

More than 30 A100 GPUs!

Juggling too many requirements?

Real PATh Facility projects

- Simulations using 24-32 cores on a single node
- Modeling jobs needing only 1 core and 70 GB RAM

Pictured: man carrying a comical amount of equipment to wash a car, dropping it all and falling over

Tired of long jobs never completing?

```
[blin@submit4 ~]$ condor_watch_q
BATCH      IDLE  RUN  DONE  TOTAL  JOB_IDS
ID: 5556   500    -     -    500   5556.0 ... 5556.99 [----------------------------------------]

[----------------------------------------------------------------------------]

Total: 500 jobs; 500 idle

Updated at 2023-07-08 19:41:47
Input ^C to exit
```

Pictured above: condor_watch_q output indicating all of a user's jobs being idle
Pictured left: a skeleton wearing glasses, a lab coat, and a bowtie (Artur Tumasjan, Unsplash)

Get a PATh Facility credit account for 4 easy payments of $0

https://portal.path-cc.io/application

Credit System

- Computational time is charged to your credit account
- To be eligible, your group must have a relationship with an affiliated NSF program:
  - Have an active or approved award,
  - Preparing a proposal,
  - Or have a proposal under review
- Submit a credit request to your NSF program officer

Affiliated NSF Programs

- Computational and Data-Enabled Science and Engineering
- Cyberinfrastructure for Sustained Scientific Innovation
- Innovation: Bioinformatics
- Neural Systems
- Collaborative Research in Computational Neuroscience
- Astronomy and Astrophysics Research Grants
- Chemical Theory, Models, and Computational Methods
- Condensed Matter and Materials Theory
- Atomic, Molecular and Optical Physics – Theory
- Nuclear Physics – Theory
- Geoinformatics (GI)
- Geophysics (PH)
- Arctic Research Opportunities
- Antarctic Research

How many credits do you need?

CPU Credits

Cores per job	Credit charge per core hour
1	1.0
2-8	1.2
8-32	1.5
>32	2.0

GPU Credits

GPUs per job	Credit charge per GPU hour
1	1.0
2	1.2
3	1.5
4	2.0

Source: https://portal.path-cc.io/documentation/overview/references/credit-account-charges/

How many credits do you need?

Additional CPU Credit Charge

Memory Per Core

When more than 2GB per core of memory is requested by the job, there's an additional per-GB charge for memory.

Memory (GB) per job	Credit charge per hour, per GB
Up to 2GB per core ("nominal")	No charge
2-8 GB greater then nominal	0.125
8-32	0.25
32-128	0.375
128-512	0.50

Additional GPU Credit Charge

CPUS per GPU

When more than 16 cores per GPU is requested by the job, there's an additional per-core charge for the CPU cores beyond the nominal.

CPUs per GPU	Credit charge per hour, per core
Up to 16 cores per GPU ("nominal")	No charge
16-48 cores per GPU	0.125
48-64	0.20

Memory per GPU

When more than 2 GB per core of memory is requested by the job, there's an additional per-GB charge for memory for the beyond-nominal memory usage.

Memory (GB) per job	Credit charge per hour, per GB
Up to 128GB per GPU ("nominal")	No charge
128-384 GB per GPU	0.012
384-512	0.020

Source: https://portal.path-cc.io/documentation/overview/references/credit-account-charges/

How many credits do you need?

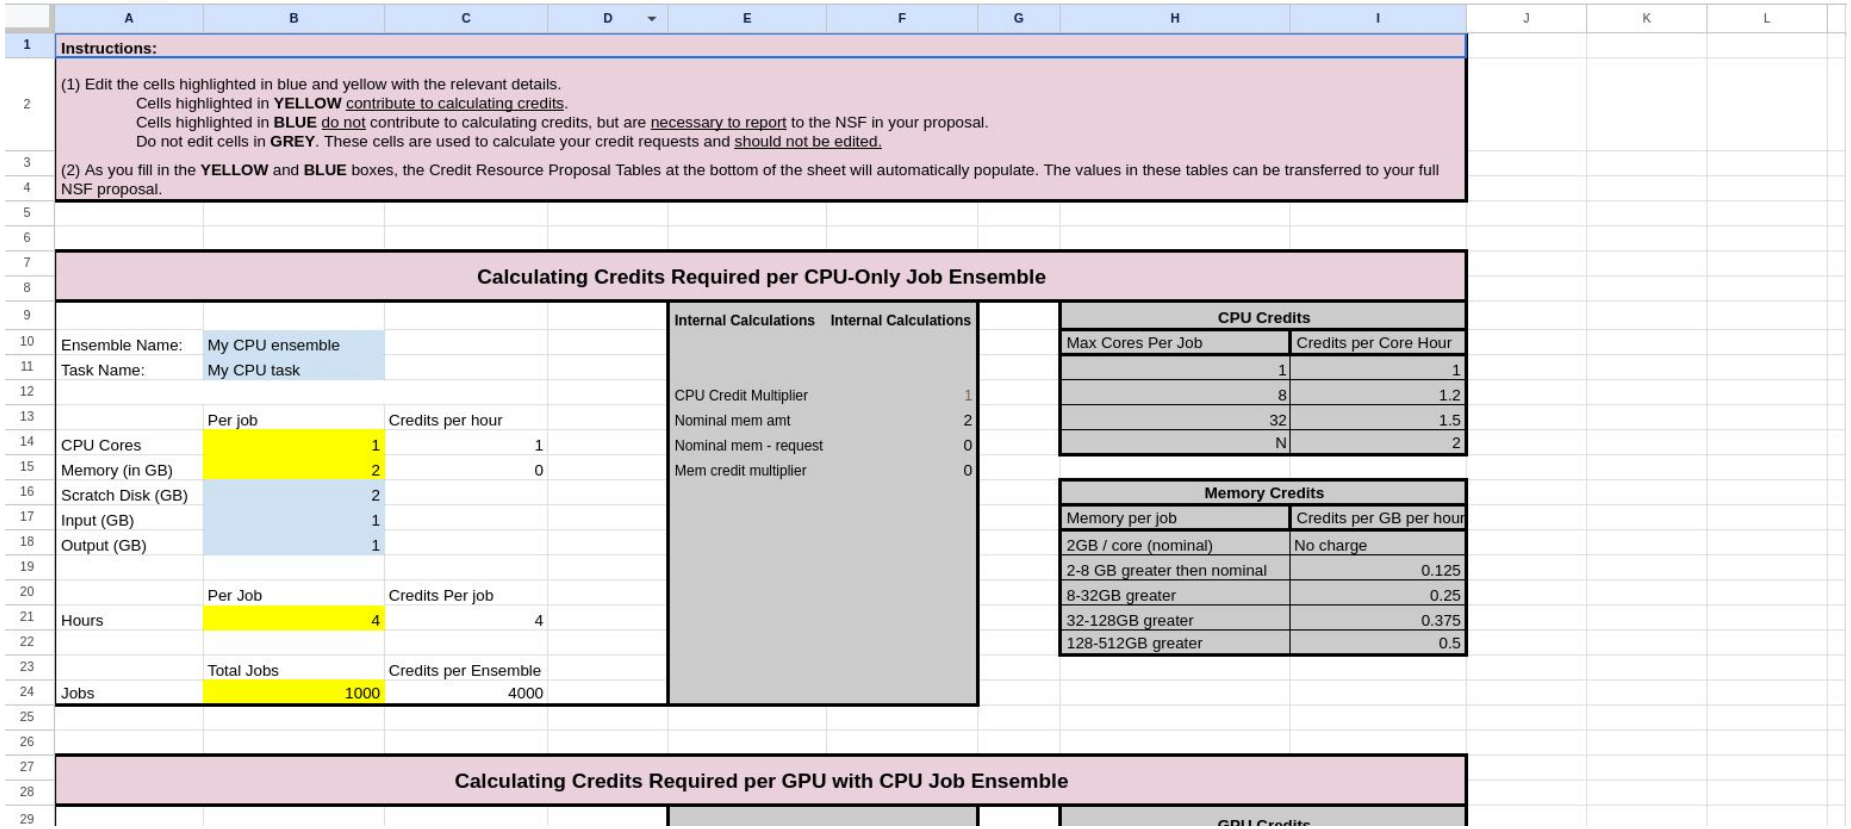

Instructions:

(1) Edit the cells highlighted in blue and yellow with the relevant details.
Cells highlighted in **YELLOW** contribute to calculating credits.
Cells highlighted in **BLUE** do not contribute to calculating credits, but are necessary to report to the NSF in your proposal.
Do not edit cells in **GREY**. These cells are used to calculate your credit requests and should not be edited.

(2) As you fill in the **YELLOW** and **BLUE** boxes, the Credit Resource Proposal Tables at the bottom of the sheet will automatically populate. The values in these tables can be transferred to your full NSF proposal.

Calculating Credits Required per CPU-Only Job Ensemble

Ensemble Name:	My CPU ensemble		
Task Name:	My CPU task		
	Per job	Credits per hour	
CPU Cores	1	1	
Memory (in GB)	2	0	
Scratch Disk (GB)	2		
Input (GB)	1		
Output (GB)	1		
	Per Job	Credits Per job	
Hours	4	4	
	Total Jobs	Credits per Ensemble	
Jobs	1000	4000	

Internal Calculations	Internal Calculations
CPU Credit Multiplier	1
Nominal mem amt	2
Nominal mem - request	0
Mem credit multiplier	0

CPU Credits	
Max Cores Per Job	Credits per Core Hour
1	1
8	1.2
32	1.5
N	2

Memory Credits	
Memory per job	Credits per GB per hour
2GB / core (nominal)	No charge
2-8 GB greater then nominal	0.125
8-32GB greater	0.25
32-128GB greater	0.375
128-512GB greater	0.5

Calculating Credits Required per GPU with CPU Job Ensemble

Source: https://portal.path-cc.io/documentation/overview/references/credit-account-charges/

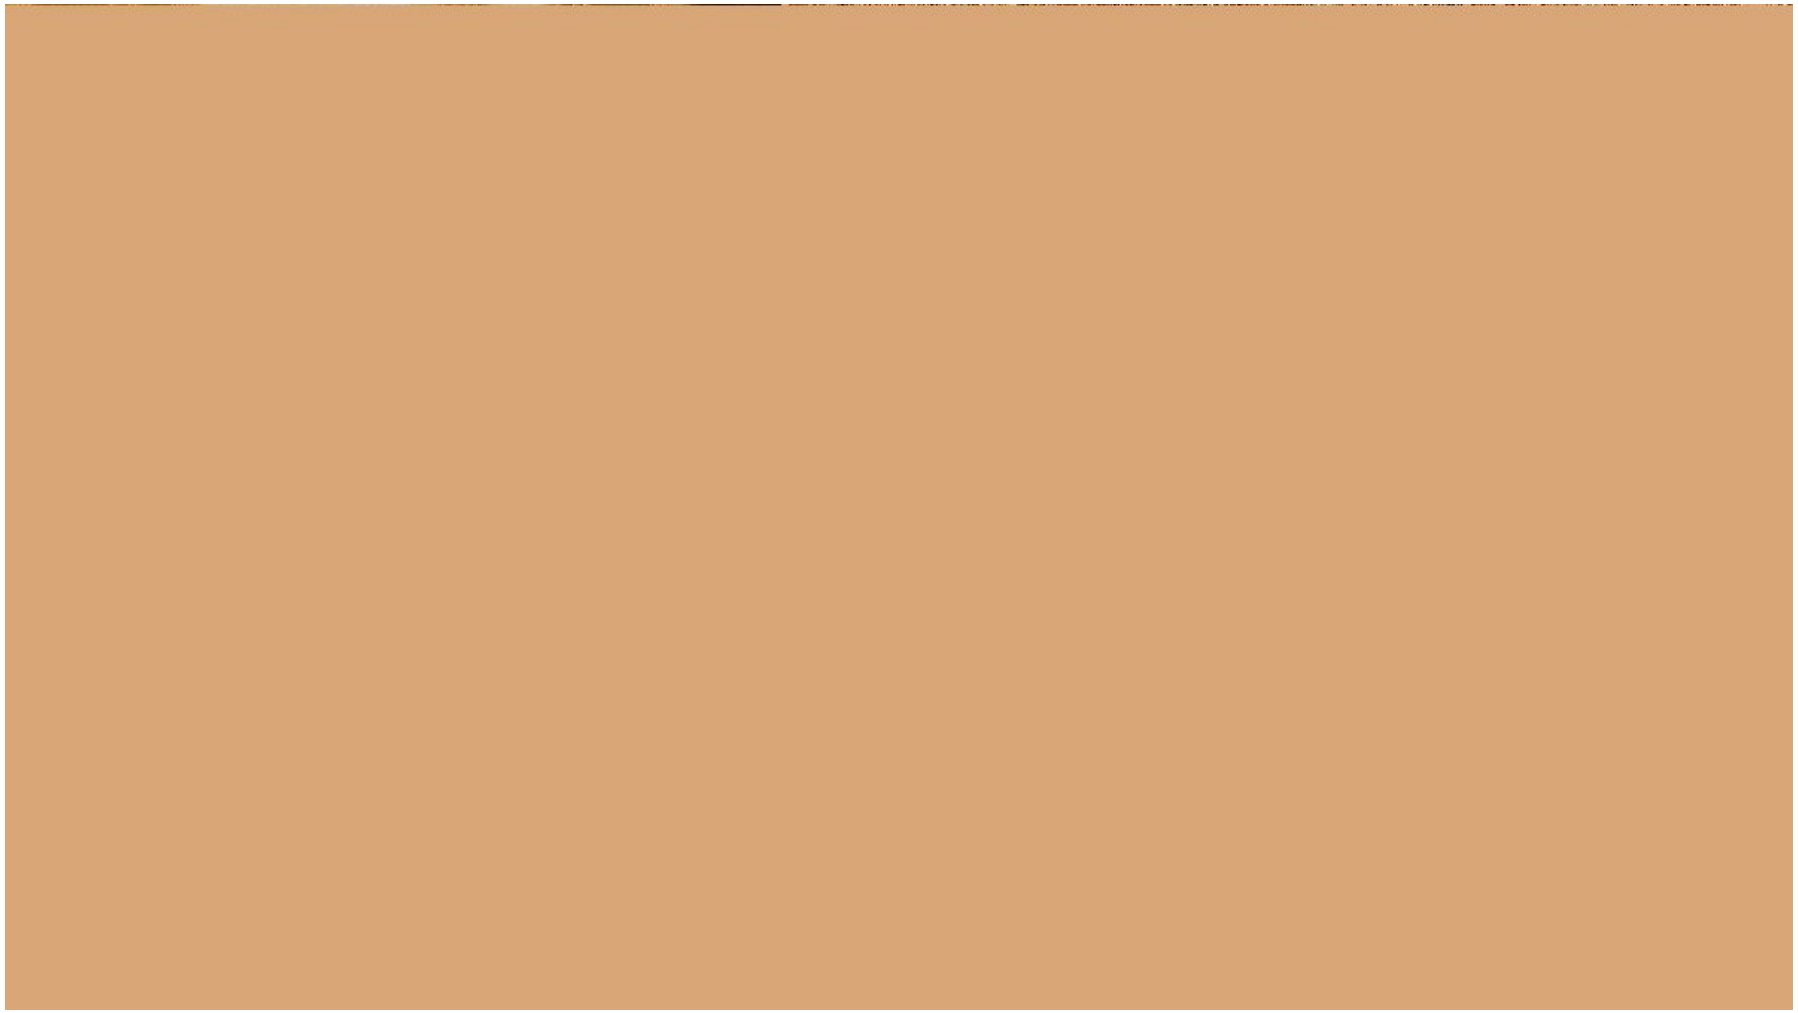

Pictured: Zach Galifianakis from the Hangover (2009)

Not sure if the PATh Facility is for you?

Request a free trial (aka starter credits)!

https://portal.path-cc.io/application

But wait, there's more!

Call* in the next 20 minutes† and receive

500 GB

of OSDF storage quota!!!

** Actually, just visit and apply at https://portal.path-cc.io/application*

† Offer available at all times

PATh Partnership to Advance Throughput Computing

OSG

HTCondor Center for High Throughput Computing

Schedule a consultation with our Facilitation Team today!

https://portal.path-cc.io/application

Pictured: poorly photoshopped image of a farcical PATh Facility informercial slogan

PATh PARTNERSHIP to ADVANCE THROUGHPUT COMPUTING

OSG

HT CENTER FOR HIGH THROUGHPUT COMPUTING

Questions

This material is based upon work supported by the National Science Foundation under Grant No. 2030508. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.