Cloud HPC Integration

SGCI Science Gateways Community Institute

Omnibond - Introduction

Leadership Team

- Former COO at NCSA
- CIO at Purdue and Clemson Universities
- Sw CTO - Clemson University
- Director Computing Engineering School Miami Ohio
- Consultant to Tribal Colleges and Minority Serving Institutions on cyberinfrastructure
- Founded ACI-REF http://www.aciref.org and CaRCC http://carcc.org
  (Emphasize the human element)

Software Products

- Cloud Orchestration - **CloudyCluster**
- Computer Vision & AI based **TrafficVision** analytics system.
- Parallel File System - **OrangeFS** included upstream in the Linux Kernel
- Identity Management - **Identity Manager Connectors**
  - products installed in thousands of customers worldwide

Why did so much Traditional IT go to the Cloud?

Uses

- eCommerce
- Seasonal Scalability
- Unsteady workloads
- Massive Scalability
- Cloud Specific Technology
- Dev / Test / QA

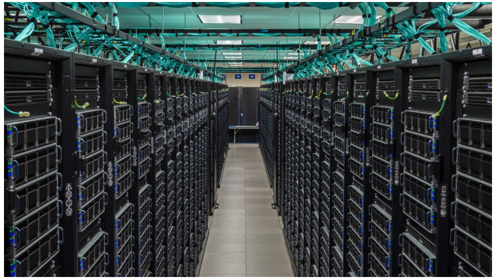

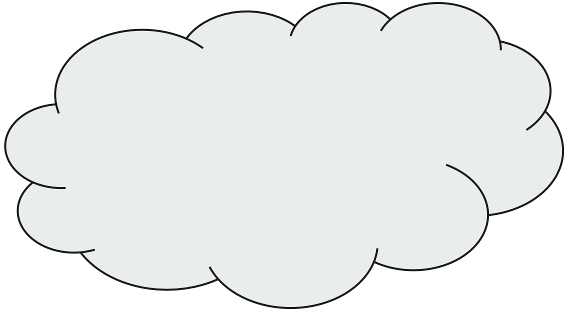

But HPC has been a more efficient use of hardware over time, though Campuses generally share data centers with HPC and Traditional IT, what happens if Traditional IT moves all to the cloud?

The Cloud HPC Opportunity?

What can it provide

- Instant Compute Availability
- Instant Storage Availability
- Diverse Resources (GPU, TPU, FPGA, etc..)
- Massive Scalability
- Cloud Specific Technologies

What about

- Latency?
- Scalability?
- Cost?

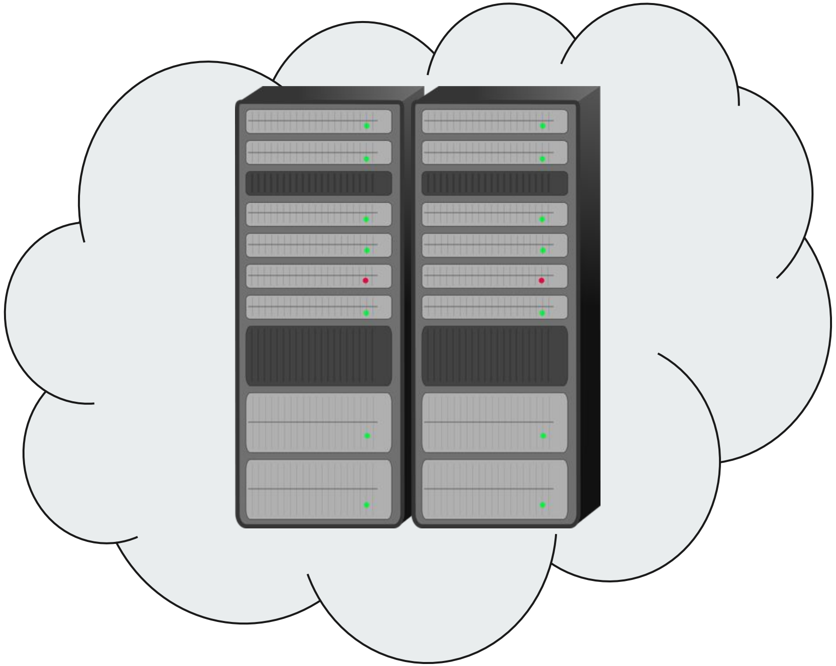

cloudyCluster

What would could Cloud HPC look like?

What should it provide?

- HPC Job Portability
- Same software as on-prem HPC clusters
- Support diverse resources
- Job expense tied to billing
- Jobs dynamically scale up and down
- Tiered long term data storage
- High Performance Scratch Storage
- Same components as on prem to make pivots easy
    - Login, Scheduler, Storage, etc…
- Optimized for Cloud specific features

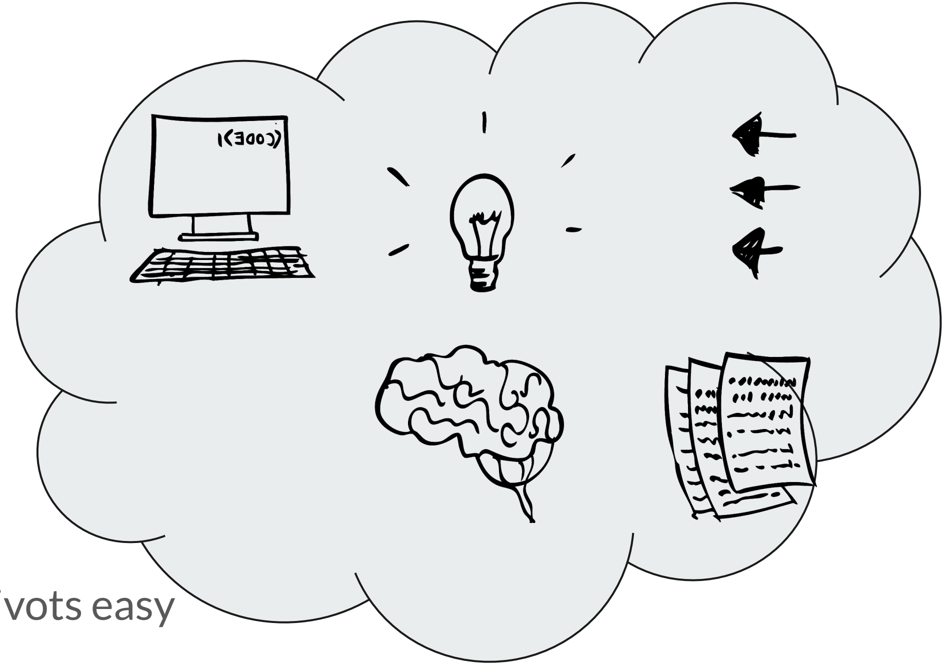

cloudycluster

Build it yourself?

Clouds provide powerful api and console driven infrastructure.

- Images
- Instances
- VPCs
- Firewall configs
- Security constructs
- Auto-scaling
- Keep up with the latest capabilities
- Object Storage
- Parallel Storage

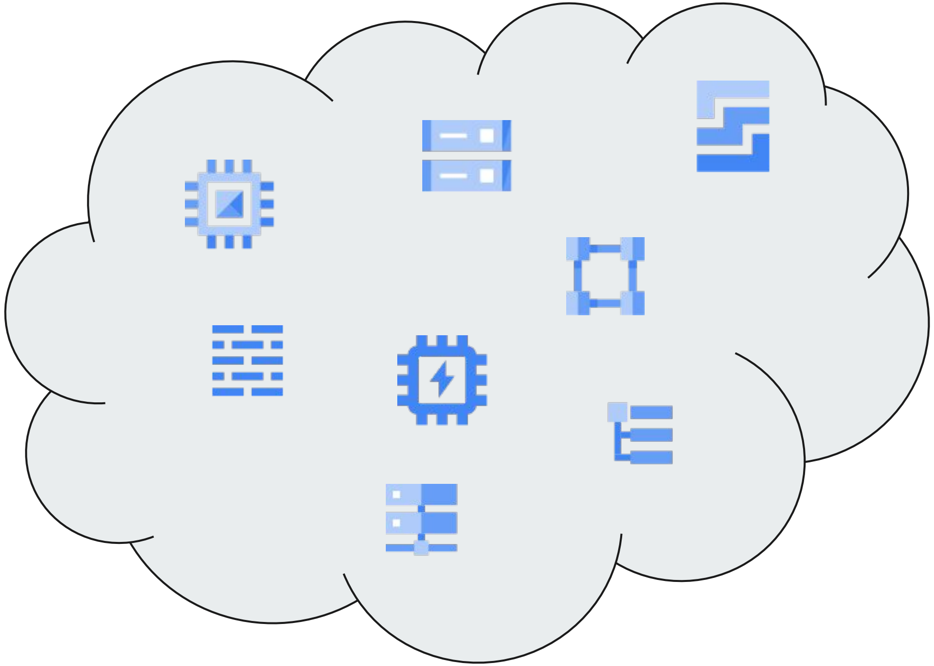

cloudyCluster

What we built into CloudyCluster

CloudyCluster = Automated Cloud CyberInfrastructure

- Automated Deployment of fully functioning HPC/HTC environment
  (Complete with VPC, firewall config, scheduler, login, parallel & object storage)
- Provides users a familiar HPC & HTC experience to on prem
- Over 300 HPC, HTC, and AI packages and libraries pre-configured
- Job scripts determine instances used (Standard, Preemptible, GPU)
- Slurm or Torque fronted by the powerful CCQ meta-scheduler
- Parallel and Object storage Options
- Operates in the customer’s own Cloud account (eliminating 3rd party BAAs)
- HPC / HTC job integration with GCP Billing Labels
- Automatically leverages Placement Policies for jobs using Supported (C2) Instances

Self-Service Elastic HPC & HTC

You Create a fully operational & secure computation cluster in minutes, complete with:

Encrypted Storage: GCS, OrangeFS on PD

Compute: Job Driven Elastic Compute through CCQ

Scheduler: Torque & SLURM with the CCQ Meta-Scheduler

Includes over 300 packages and libraries used in HPC including:

HPC Libraries:

Boost, Cuda Toolkit, Docker, FFTW, FLTK, GCC, Gengetopt, GRIB2, GSL, Hadoop, HDF5, ImageMagick, JasPer, mpich, NetCDF, NumPy, Octave, OpenCV, OpenMPI, PROJ, R, Rmpi, SciPy, SWIG, WGRIB, UDUNITS, .NET Core, Singularity, Queue, Picard, and xrootd

HPC Software:

Ambertools, ANN, ATLAS, BLAS, Blast, Blender, Burrows-Wheeler Aligner, CESM, GROMACS, JupyterLab, LAMMPS, NCAR, NCL, NCO, nwchem, OpenFoam, papi, paraview, Quantum Espresso, SAMtools, WRF, Galaxy, Vtk, Su2, Dakota, and Gatk

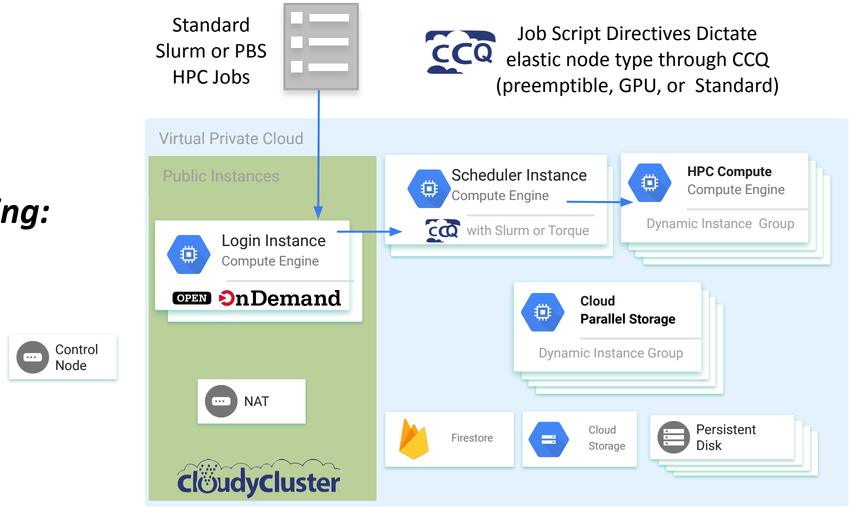

ML Software:

Mlpack, NuPIC, Octave, OpenCV, PICARD, Queue, Scikit-learn, Tensorflow and Theano

Customization: *You can easily add your own software to a custom Image.*

Scalability (HTC)

Google HPC Blog Post

Cloud against the storm: Clemson's 2.1 million VCPU experiment

https://cloud.google.com/blog/topics/hpc/clemson-experiment-uses-2-1-million-vcpus-on-google-cloud

Urgent HPC can Burst Affordably to the Cloud

https://www.nextplatform.com/2020/01/08/urgent-hpc-can-burst-affordably-to-the-cloud/

- 133,573 GCP Instances at peak
- 2,138,000 vCPUs at peak
- 6,022,964 vCPU hours

Processed 2,479,396 hours (~256TB) of video data

- ~4 hours of runtime
- ~1M vCPU within an hour
- ~1.5M vCPU within 1.5 hours
- 2.13M vCPU within 3 hours

Total Cost: $52,598.64 USD
Average cost of $0.008 USD per vCPU hour

CLEMSON UNIVERSITY · Google Cloud · Cloudy Cluster by Omnibond · intel · TrafficVision · burwood group, inc.

HPC Workloads for Google Cloud

If you look at the Campus resources a majority (>90%) of the jobs will run on <9 instances (given 30 cores per instance).

The MPI latency between <9 instances is only impactful on a very small number of job types.

A majority of workloads run on a campus cluster will be supported in GCP

The larger NSF resources require more investigation

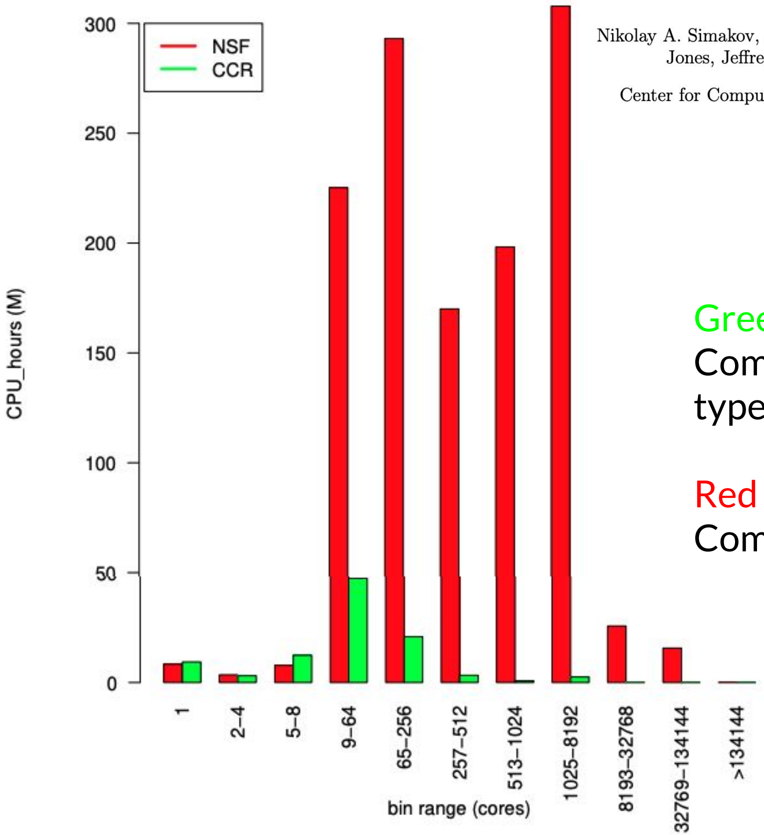

A Workload Analysis of NSF's Innovative HPC Resources Using XDMoD

Nikolay A. Simakov, Joseph P. White, Robert L. DeLeon, Steven M. Gallo, Matthew D. Jones, Jeffrey T. Palmer, Benjamin Plessinger, and Thomas R. Furlani

Center for Computational Research, University at Buffalo, Buffalo NY USA 14203

January 16, 2018

Green represents Campus Computing Resources and the job types (# cores)

Red represents National NSF Computing Resources

Google Working Toward Optimal HPC Performance

Google Cloud

Blog Menu

HPC

Introducing HPC VM images—pre-tuned for optimal performance

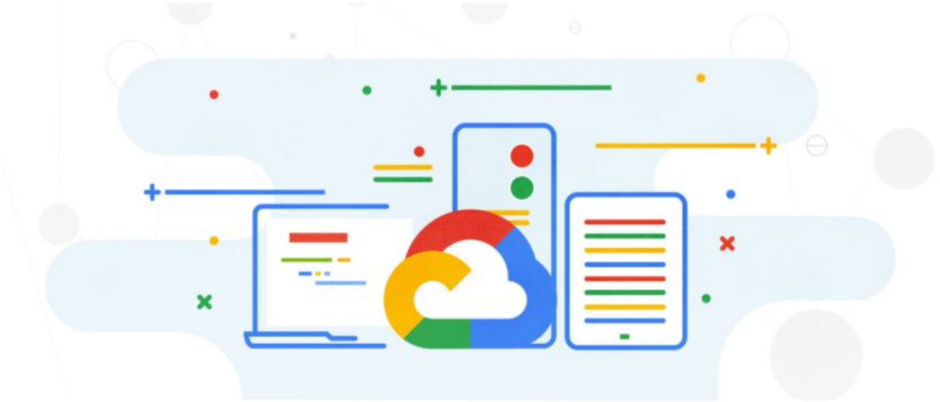

Pavan Kumar
Product Manager

Jason Zhang
Software Engineering Manager

February 2, 2021

Customer story: Scaling SDPB solver using CloudyCluster and HPC VM image

Walter Landry is a research software engineer in the Caltech Particle Theory Group working with the international Bootstrap Collaboration. The collaboration uses SDPB, a semidefinite program solver, to study Quantum Field Theories, with application to a wide variety of problems in theoretical physics, such as early universe inflation, superconductors, quantum Hall fluids, and phase transitions.

To expand the collaboration's computation capabilities, Landry wanted to see how SDPB would scale on Google Cloud. Working with Omnibond CloudyCluster and leveraging the HPC VM image, Landry achieved comparable performance and scaling to an on-premises cluster at Yale, based on Intel Xeon Gold 6240 processors and Infiniband FDR.

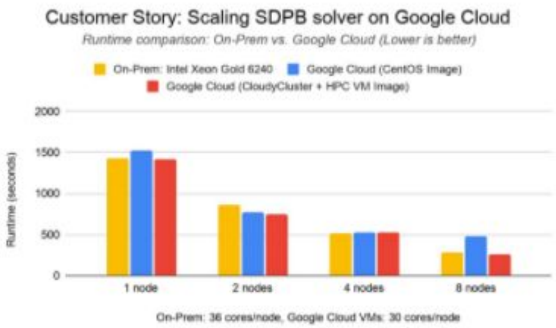

https://www.googlenewsapp.com/introducing-hpc-vm-images-pre-tuned-for-optimal-performance/

Community Partners

NET+ GCP program

- BAA vehicle
- Easiest access to NIH STRIDES program for GCP

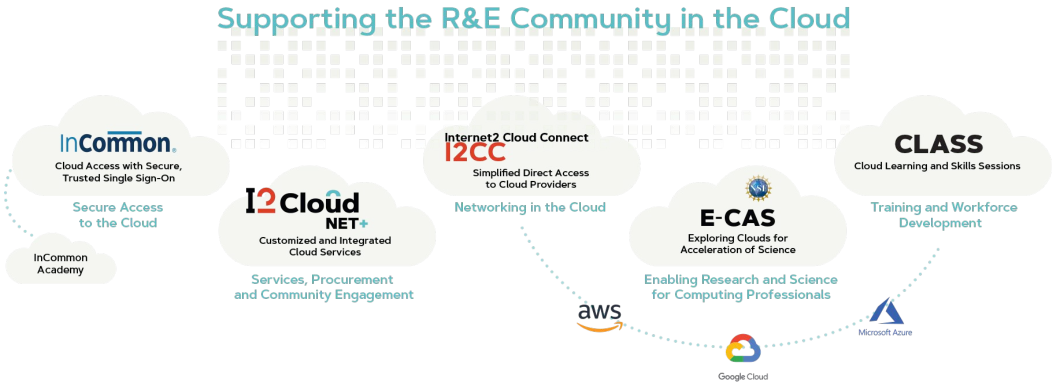

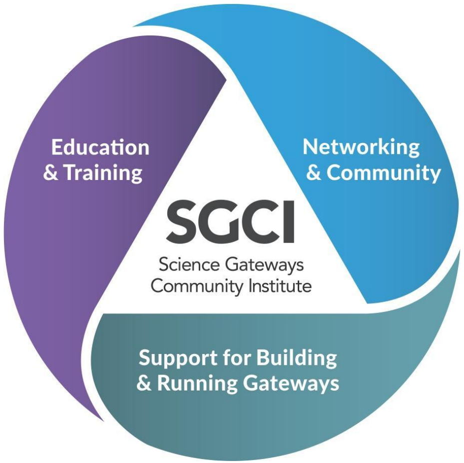

Cloudify Gateways Solicitation

sciencegateways.org/cloudify2020

Hackathons at SC20 & Pearc21

hackhpc.org

Consumption, Deployment & Engagement

CloudyCluster can be launched easily from the GCP Marketplace:

- Launch easily from the Cloud Marketplace
- Pricing is about 5% of on-demand costs
- Can be sold as a subscription for select customers.

The CloudyCluster Team Can Help with:

- HPC & HTC Discussions (Consulting and Pre-Sales)
    - Workloads that work well in the Cloud
    - Migrating HPC & HTC Jobs that can run in the cloud and on premises
    - Cloud HPC Cost Estimates
    - Proof of Concept (POC) workflow
- Additional Services available for tighter campus integration

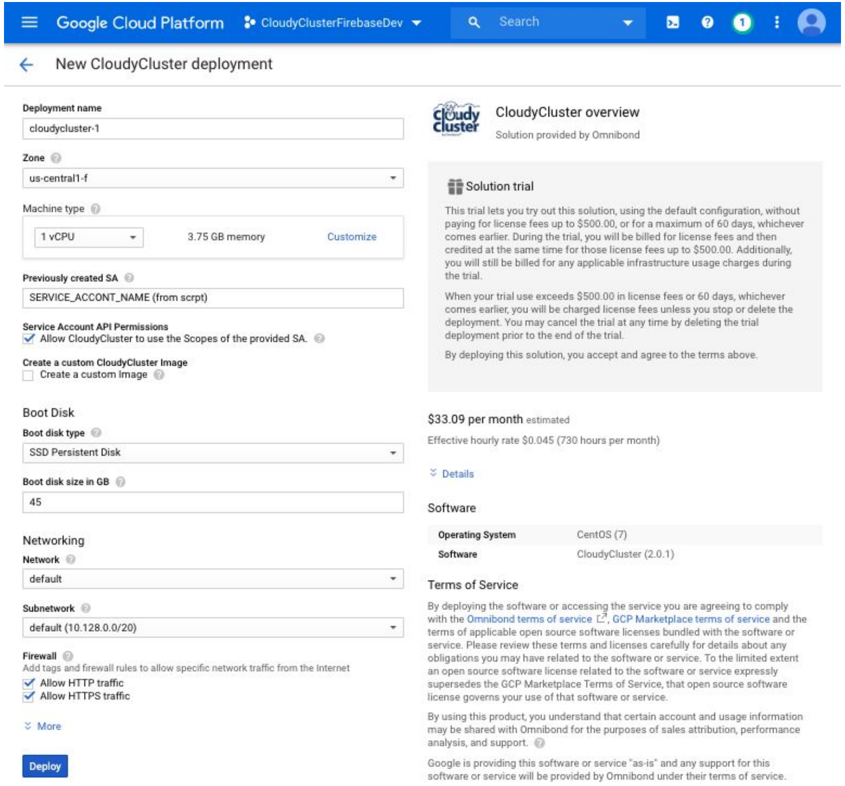

cloudyCluster

Gateway Configurations for Cloud

On-prem with Cloud option

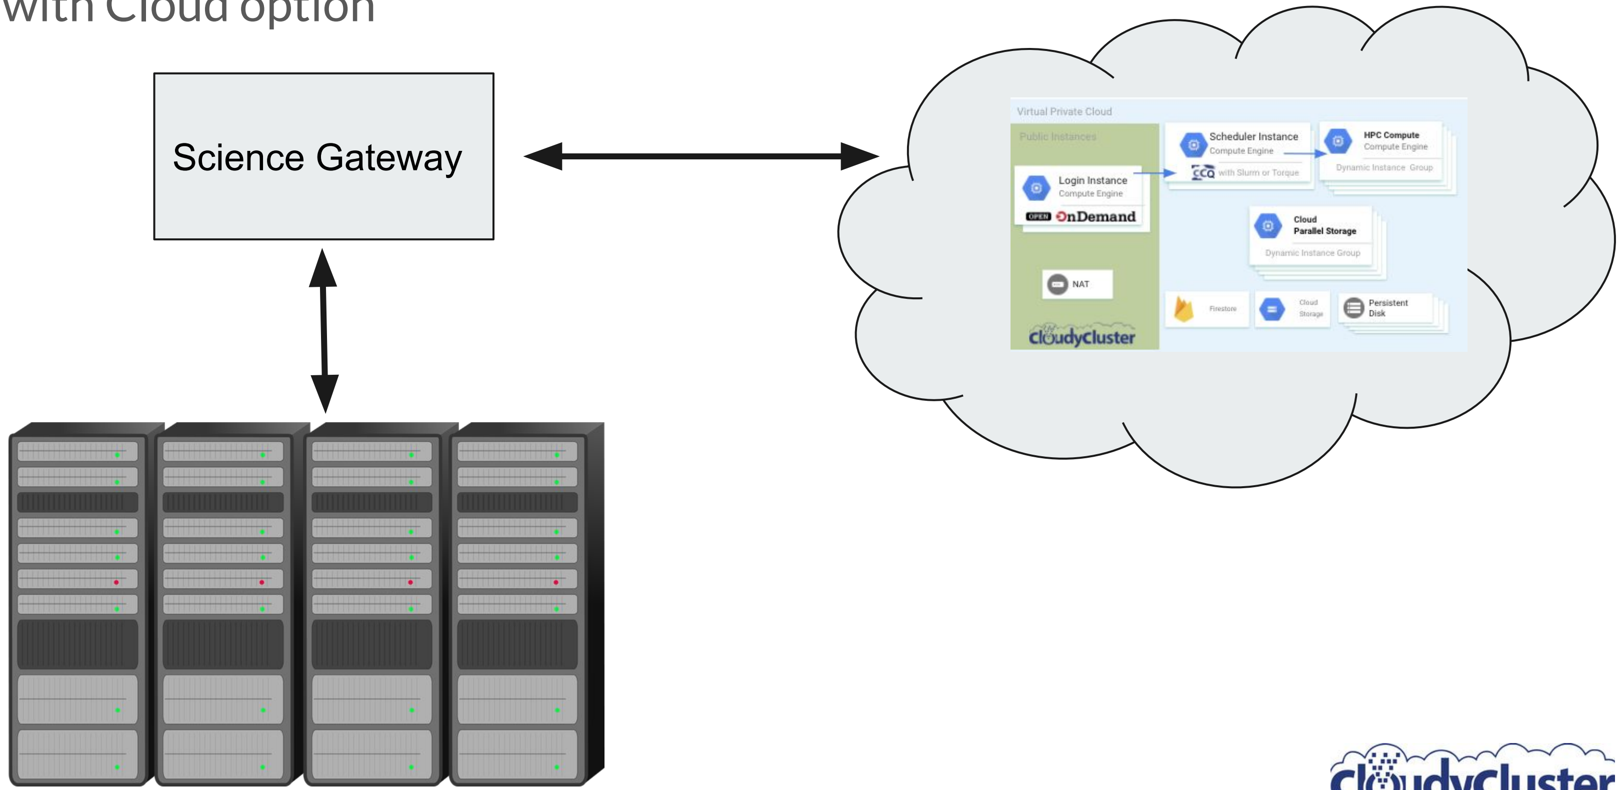

CloudyCluster

Gateway Configurations for Cloud

Cloud native option

Provides options for commercial partners

Broader Researcher Reach

Can build a turn-key gateway option

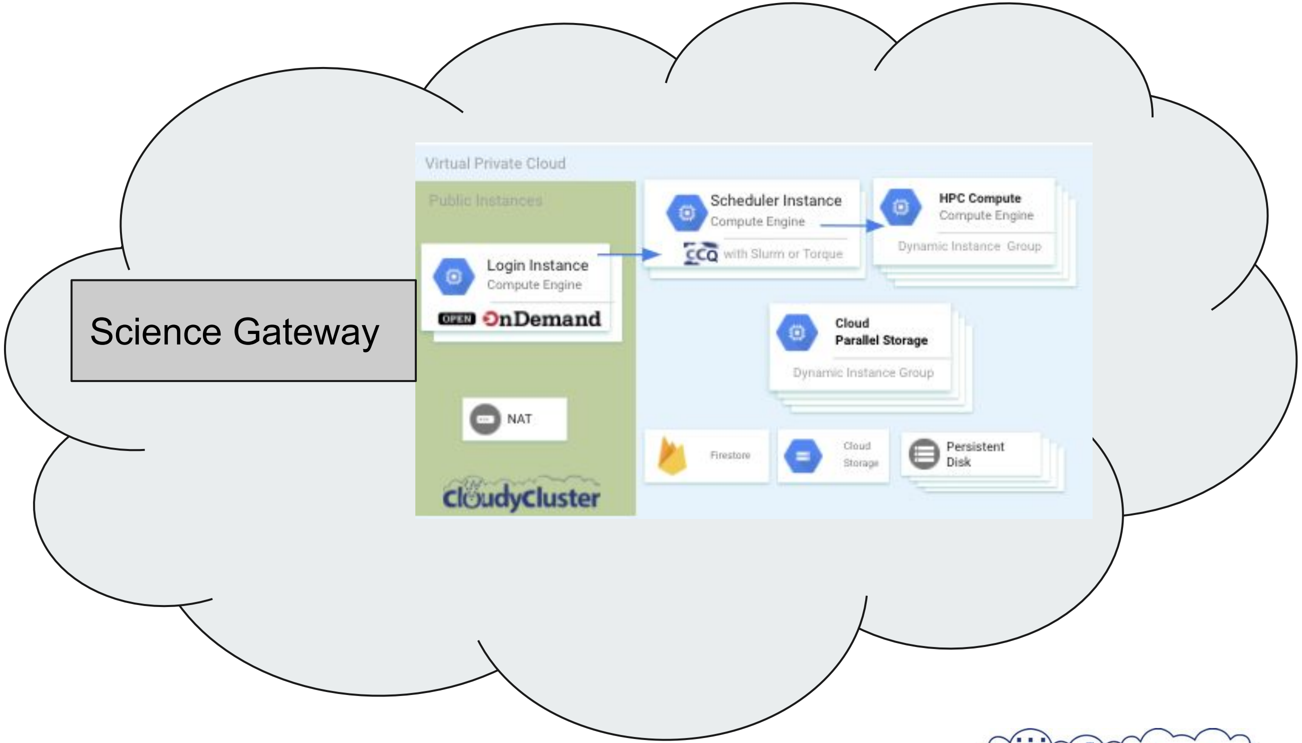

Cloudify Project Open OnDemand on CloudyCluster

Supercomputing. Seamlessly. Open, Interactive HPC Via the Web.

Project

- **Need**: Many HPC centers are evaluating the usage of cloud-based HPC resources

- **Problem**: OOD architecture based on traditional HPC resource managers and intra-node communication doesn't work in the cloud

- **Solution**: Combine the expertise of the OOD and CloudyCluster teams, with CCQ and Slurm running on the Google Cloud Platform

- **Outcome**: A demo of OOD running on GCP and supporting an example HPC client workflow

Ohio Supercomputer Center

University at Buffalo Center for Computational Research

VIRGINIA TECH.

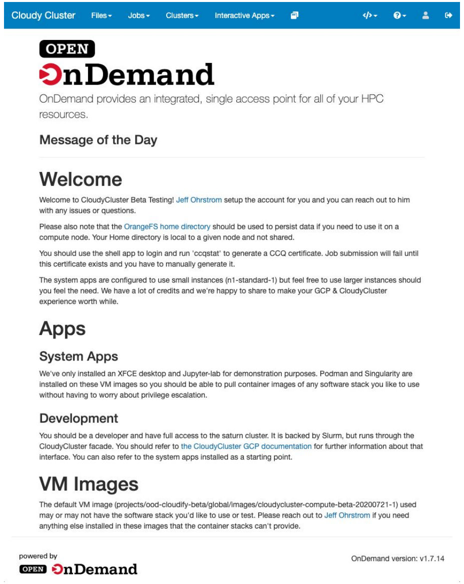

Education & Training | Networking & Community | SGCI Science Gateways Community Institute | Support for Building & Running Gateways

Google Cloud | OPEN OnDemand | cloudyCluster

Cloud Cost Comparison

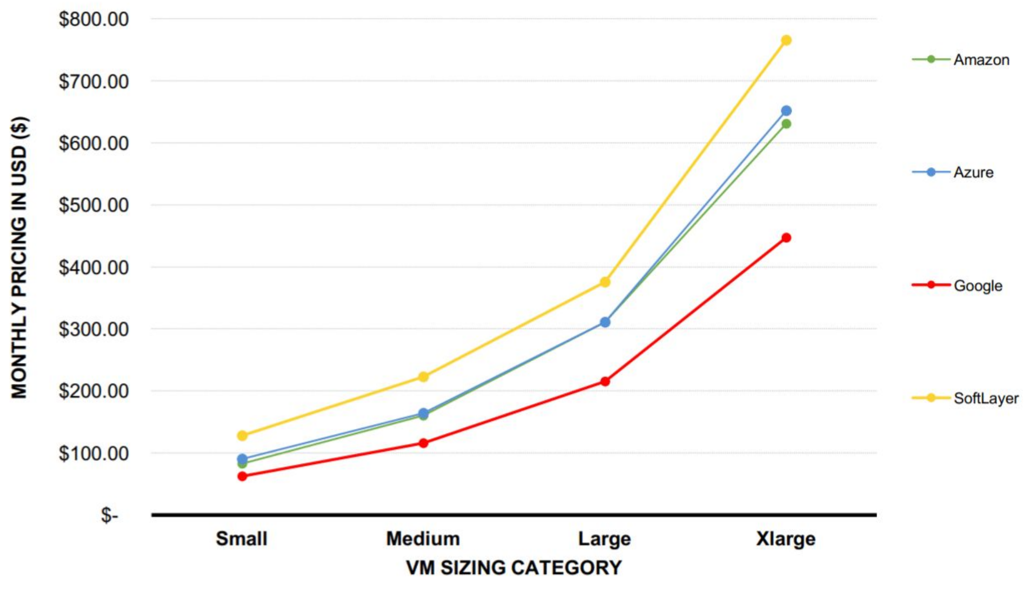

Table 7A: Monthly Cost of VMs Across CSPs

	Small	Medium	Large	Extra Large
Amazon	$83.00	$160.27	$310.54	$631.08
Azure	$90.19	$163.92	$310.65	$652.02
Google	$62.24	$115.98	$214.96	$446.92
SoftLayer	$128.00	$223.00	$376.00	$766.00

Source: Kinsta

Next Evolution in Cloud Pricing

Google Cloud Subscription Agreement for Public Sector

Provides the “right to use” any GCP product for a period of time for a defined, specified use case for a fixed price

Provides you with predictability of costs - no overage fees

Reduce your implementation risk - use whatever service you need, as much as you need, for your project

We’re in this with you! **Google Cloud has “skin in the game”** with you - we want your project to be successful

Google Cloud

Custom Integration Options

for Research Organizations

1 **Cloud Image customization**
for seamless operation
 - For a majority of workloads the cloud image can look like on premise image

2 **Extending on premise Login Node to a Cloud Queue**
 - Providing the ability to make the cloud simplified to the end researcher.

3 **Job Cloud Billing Integration**
 - From an auditing or bill back perspective

5 **IDM Integration**
of on Premises and Cloud HPC including HPC System Consolidation

3 **Data & Results Staging integration**

6 **Training & Tools**
to train the Research Computing Support Staff

University of Michigan
Cianfrocco Lab (Work in Progress)

RELION & CRYO-EM

Work in Progress

Working with the Cianfrocco Lab at University of Michigan to test data transfer and scalability to the Cloud for CRYO-EM projects using RELION

Also working with related Startups in the space that don’t have access to large on prem clusters.

Initially Introduced as part of the Cloudify Program through the SGCI.

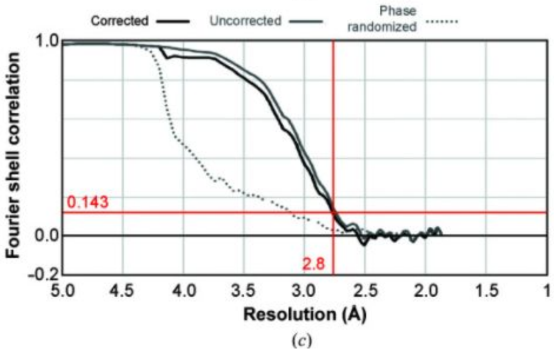

High-resolution cryo-EM using beam-image shift at 200 keV

Jennifer N Cash [1], Sarah Kearns [1], Yilai Li [1], Michael A Cianfrocco [1]

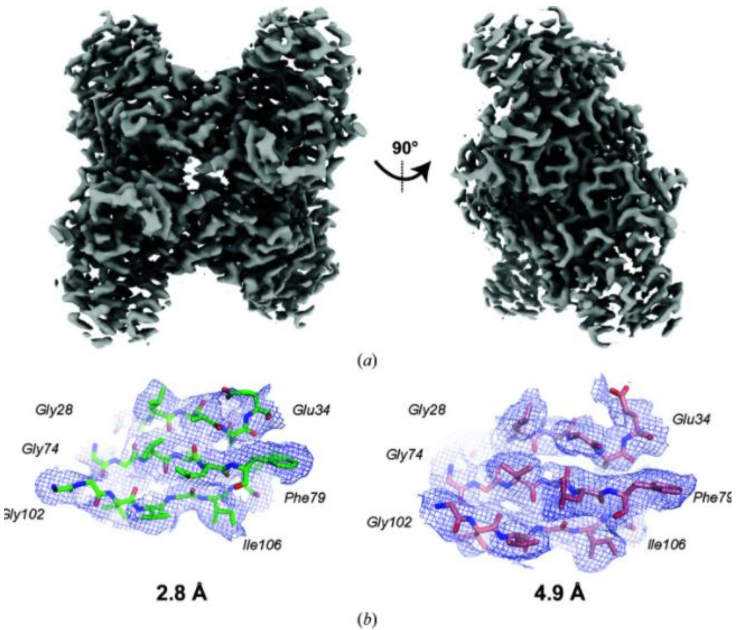

Figure 5 Final aldolase reconstruction at 2.8 Å resolution. (*a*) Sharpened aldolase reconstruction at 2.8 Å resolution. (*b*) Example densities and models for aldolase at 2.8 and 4.9 Å resolution. (*c*) FSC curve for the final reconstruction.

Outreach among Rutgers & ERN Researchers: Using Cloud Solutions to Extend and Diversify Compute

Outreach at Rutgers & ERN

- Distributed research support operations based at Rutgers and interconnected with **28 institutions**
- Sharing best practices, training, resources, computational research projects
- Faculty, staff, students onboarded with research computing + the range of options: on-prem, regional, national, cloud
    - Intro to Research Computing
    - Software (on-prem, regional, cloud)
    - Storage options (on-prem, cloud)
    - Compute options (on-prem, regional, cloud)
    - Use cases & demonstrations
    - Support resources

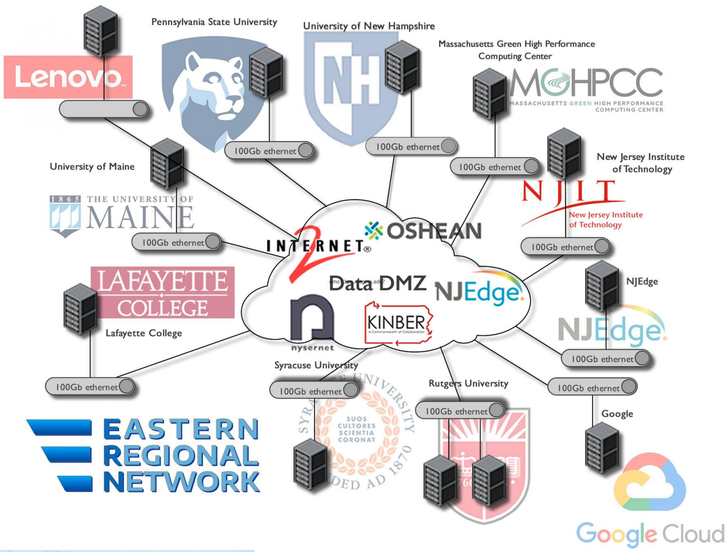

Rutgers & Cloud-Enabled Research

RUTGERS

- 8,000+ faculty, 300 research centers
- Part of my team’s mission is pairing researchers with the **right tools for their projects**
- All faculty, staff, students have free access to our central HPC resources, but there are **constraints**
- Central IT manages accounting for granted and discounted cloud credits, including STRIDES-provided resources
- Cloud marketplaces are a **key differentiator** when considering cloud vs. on-prem capabilities
- CloudyCluster is the primary recommendation for compute elasticity and scale

CloudyCluster in a Student's Toolbox: Preparing Emerging RCFs/REFs with Key Cloud Capabilities

Student Workforce Development

- Support and education tends to stem from the role of Research Computing Facilitator (RCF) or Research and Education Facilitator (REF)
- No formal, focused training program, so experiential student programs have emerged
- Critical skills and exposure to technologies not encountered elsewhere
    - Project management fundamentals
    - Research computing facilitation
    - Cloud storage selection & management
    - AWS SageMaker / GCP ML Engine
    - Software-defined networking & P4
    - HPC: traditional & cloudycluster

CAREERS CYBERTEAM

https://careers-ct.cyberinfrastructure.org

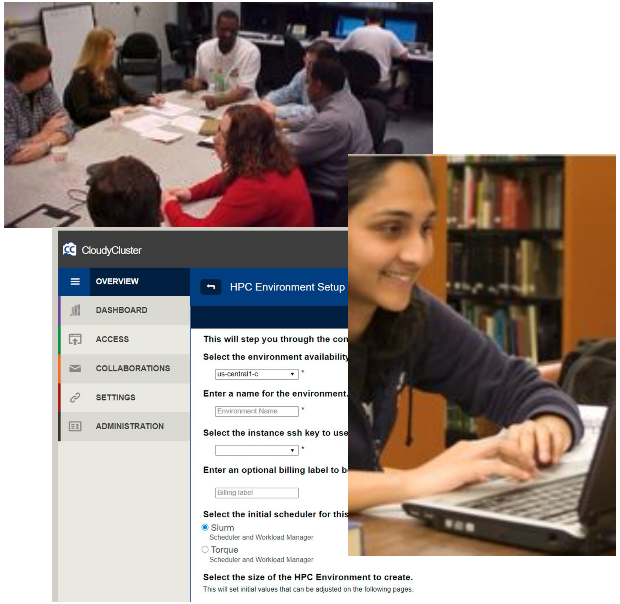

Student Projects Exploring CloudyCluster

- Electrostatics of the GABA(A) receptor pore: exploring a novel asymmetric hydrophobic mismatch mechanism that may be responsible for experimentally observed membrane deformation.
  PI: **Grace Brannigan** RUTGERS

- Parallel analysis of variants in multiple bear genomes: extending the GROM (Genome Rearrangement Omni-Mapper) core algorithm and software suite to identify all types of genome variants in a single run.
  PI: **Andrey Grigoriev** RUTGERS

- Developing software for analysis of electron microscopy images of fractal soot aggregates: creating numerical representations of aggregates for use in morphological and optical calculations.
  PI: **Alexei Khalizov** NJIT New Jersey Institute of Technology

- Training deep learning models for protein-protein-interaction prediction in absence of structural information about the protein complex.
  PI: **Guillaume Lamoureux** RUTGERS

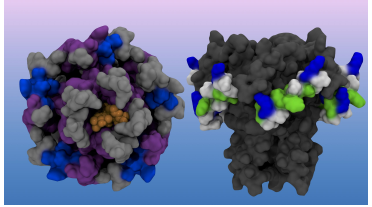

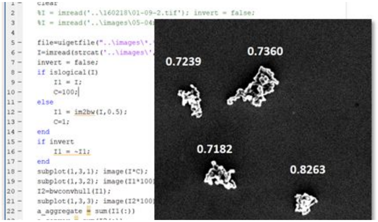

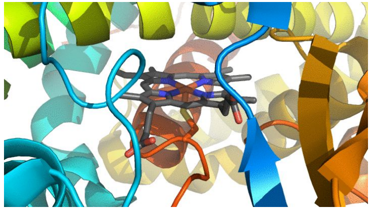

Research Credits

Transform research data into valuable insights and conduct large-scale analyses with the power of Google Cloud Platform and CloudyCluster

It is the intent of the program to provide seed funding in the form of credits to be used for research projects

- Not intended to fully fund research
- 12 month credit limit
- One award per calendar year per researcher
- One award per project
- Can apply for up to $5,000 in credits
- PhD candidates eligible **NEW**

- Omnibond can help determine project scope

Apply Now

Not sure where to start?

Access trainings

Explore our resource library

Join our faculty community

Questions? Email gcpresearchcredits@google.com

Google Cloud

Yates Lab Cystic Fibrosis Project

Background

The structure of DNA was proposed in 1953

After the translation of the RNA message into proteins they must fold into the proper conformation and proteins are modified with many different types of Post Translational Modifications (PTM) that are necessary to achieve and modulate biological activity.

These PTMS are a protein code that controls cellular functions. A component of protein function is determining how modifications are used to regulate or alter function.

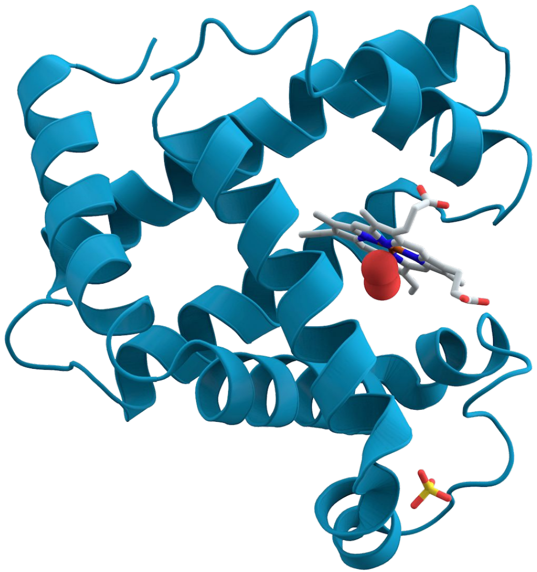

Through funding from the NIH (Grant ID: 1 R01 HL 131697-01A1), the Yates Laboratory at The Scripps Research Institute in La Jolla, led by professor John Yates. Researchers include Robin Park, Sandra Pankow, Titus Jung

Investigation

- The **majority of Cystic Fibrosis** cases are **caused by a deletion of phenylalanine 508** (ΔF508) in the CFTR protein.
- This deletion leads to an altered 3D structure of CFTR, resulting in loss of anion channel function, because the CFTR protein does not reach its proper location in the cell membrane.
- Furthermore, initial studies by the Yates lab indicated that post-translational modifications of the misfolded ΔF508 CFTR protein are different from wt-CFTR.
- **Identifying CFTR PTMs**, especially differentially regulated ones, is of **enormous importance to better understand the ΔF508 CFTR channel and cellular processing defects and could provide new strategies for therapeutic intervention.**
  - **Identifying CFTR PTMs comprehensively is difficult** because the concentration of CFTR in cells is very low and many different PTMs and combinations of PTMs are possible per protein. And while some PTMs may be quite abundant, others are just above the detection level.

Investigation

- An unbiased search for PTMs present on the protein **("Blind-PTM search") has the potential to identify all PTMs** (including SNPs and amino acid deletions) because peptide sequences are identified along with peptides that deviate from predicted molecular weights, signaling the presence of modifications.
- However, **identifying PTMs by searching extracted tandem mass spectra against the entire human genome (Blind-PTM) is computationally demanding**. Such queries can take up to two weeks on 600 node cluster for a single experiment.
- **CloudyCluster HPC enables higher parallelization and can significantly reduce run-times.**

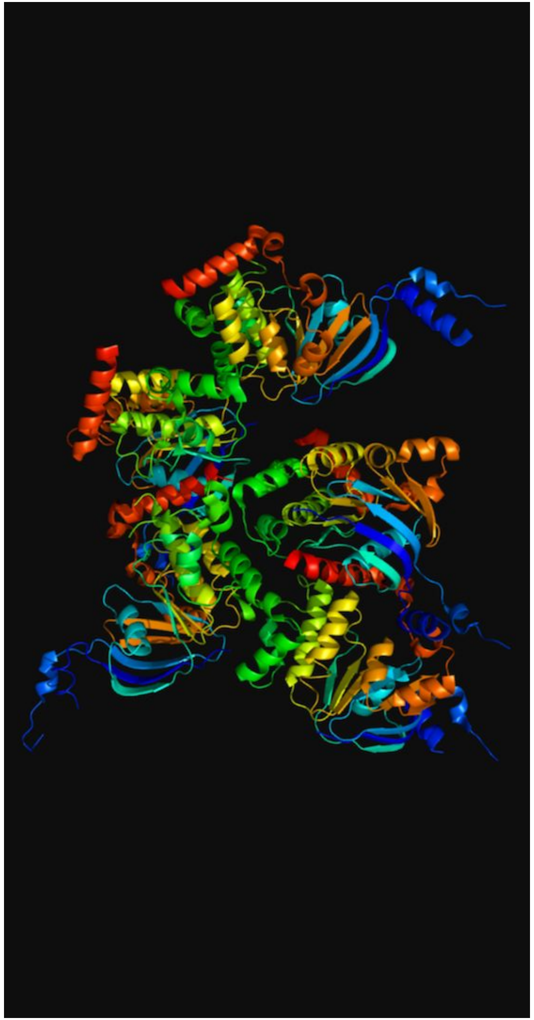

Scripps Workflow

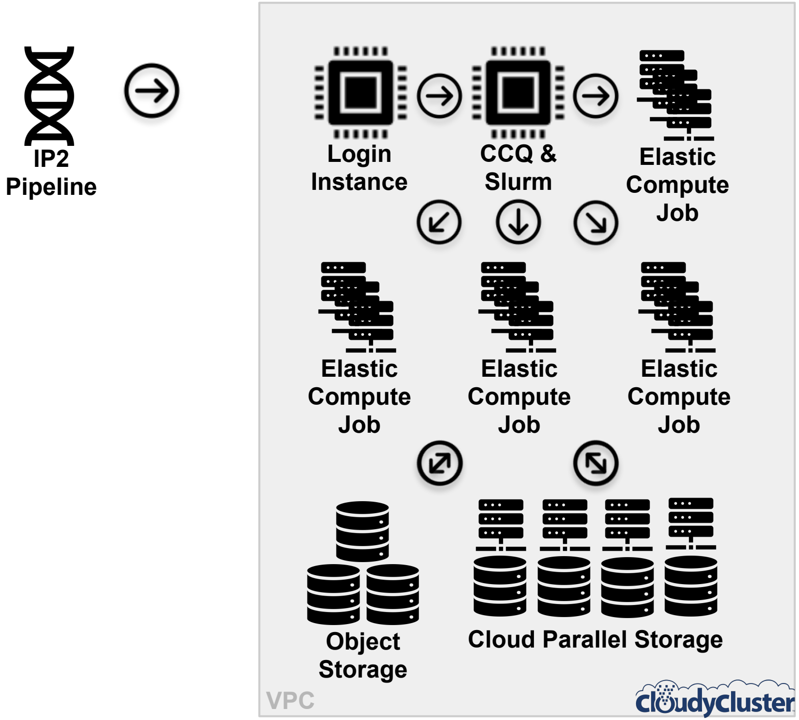

- The Integrated Proteomic Pipeline (IP2) was used.
- The IP2 is also currently being used at a variety of research labs, universities, pharmaceutical companies and research hospitals.

Data & Results

- Over 90 independent samples were analyzed on high mass accuracy mass spectrometers such as the LTQ-Orbitrap Elite (Thermo Fisher). Each experiment generated > 1 GB of data.
- About **15,000 - 20,000 spectra that were identified by the traditional ProLuCID** search approach. While the traditional search approach identified around 80,000 CFTR spectra, and revealed important PTMs that were quantitatively and qualitatively altered in Cystic Fibrosis,
- The Blind-PTM search revealed an additional 30,000 CFTR spectra that contained various PTMs. **Ultimately, it was the Blind-PTM search approach that revealed additional critical information on PTM** changes that would have been missed by traditional searches.
- The Yates lab is currently following up on several of these PTM changes to see if influencing some of these PTMs could aid in CF therapy and **a paper describing the results will be published soon**.

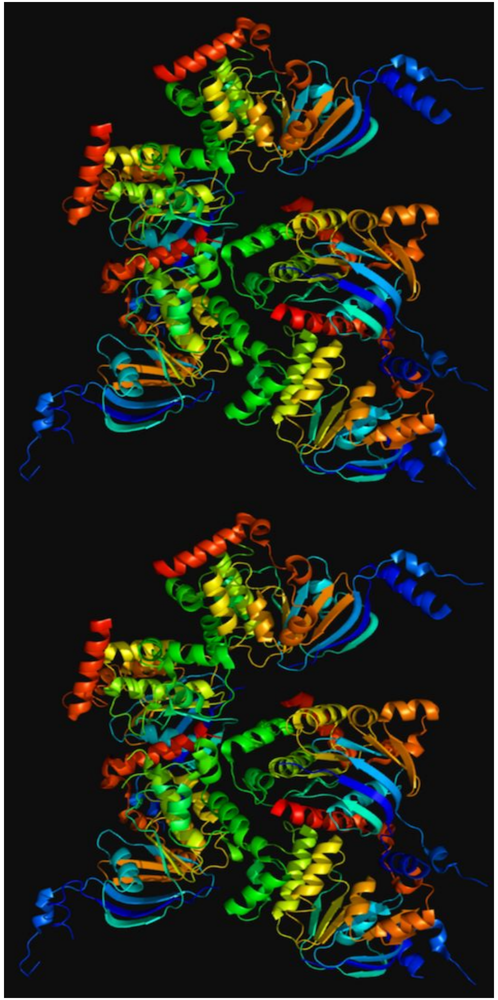

University of Colorado Denver Biomedical Ontologies Project

Background - OBO

The modern ecosystem of **biomedical ontologies** coupled with **Semantic Web** computational technology has the potential to **transform the knowledge-based interpretation** of biomedical research results.

At the forefront of this effort are the **Open Biomedical Ontologies (OBOs)**, a diverse, mostly orthogonal collection of ontologies that **hierarchically organize** and logically define concepts in a wide range of domains in biomedicine; these **domains include biological processes, cell types, phenotypes, diseases, chemicals**, among many others. Although the stated common goal of the OBOs is "to develop a family of **interoperable ontologies**"

Open Biomedical Ontologies (OBOs)

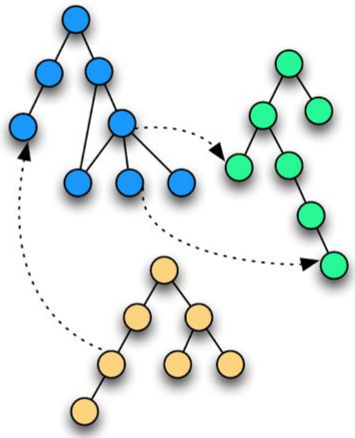

Bill Baumgartner, Center for Computational Pharmacology

School of Medicine
UNIVERSITY OF COLORADO ANSCHUTZ MEDICAL CAMPUS

Recent Work

Recent work shows that the **distributed** and **loosely coupled** nature of OBO development, where each ontology is typically authored by a distinct research group, **has resulted in unintentional conflicts** and **logical inconsistencies** between ontologies that **prevent them from being used in the unified manner** for which they were designed.

Solution

Using **CloudyCluster** we have **implemented a continuous integration system** for the **OBO development community** to monitor the OBOs **for both intra- and inter-ontology logical consistency**, as well as other quality assurance criteria. Development patterns differ among the OBOs with some ontologies changing on **a nightly basis**, while others are updated less frequently.

CloudyCluster provides the **flexibility** to **instantiate an appropriate amount of compute power** depending on the volatility of the **currently 187 ontology files being tracked**.

The continuous integration **system monitors** the ontology files for changes and typically submits **200-30,000 jobs on a periodic basis (nightly to weekly)** to keep up-to-date with OBO development. **This system is being used to unify the OBOs of today** and will serve as a **global quality control tool flagging changes with unintended consequences to maintain interoperability** among the OBOs as their development continues.

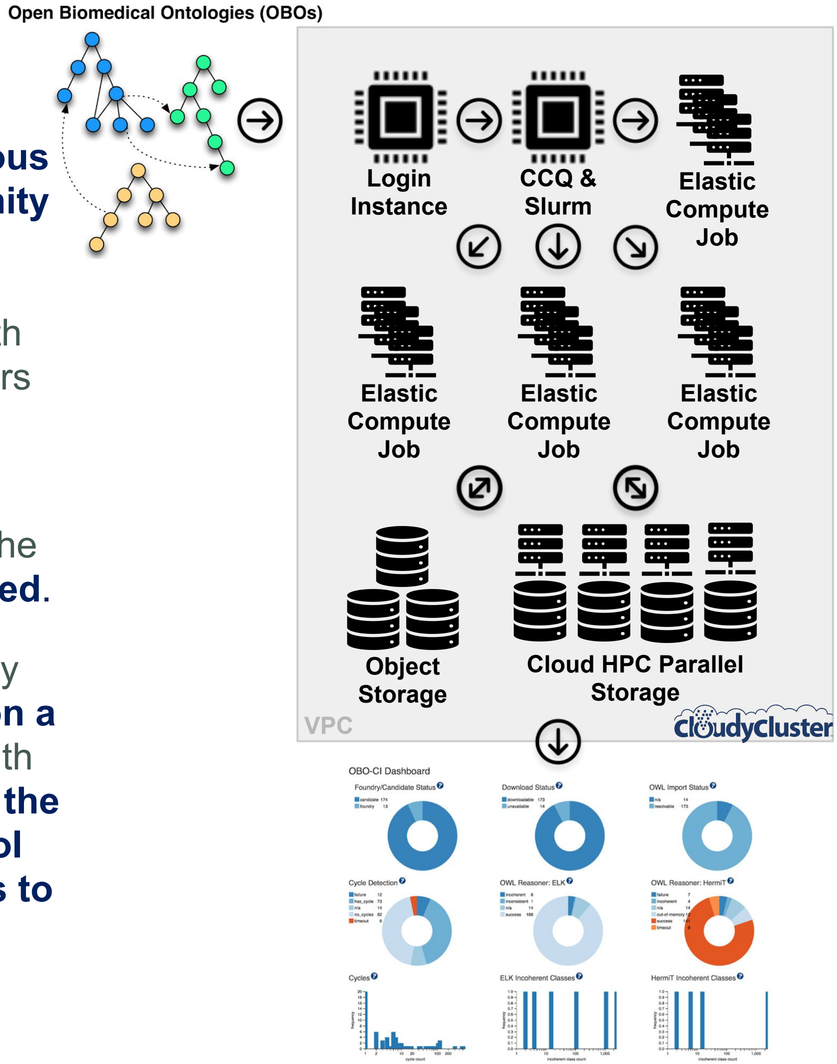

OBO CI Summary

Monitoring interoperability of the Open Biomedical Ontologies

Motivation

The modern ecosystem of biomedical ontologies has potential to transform knowledge-based interpretation of biomedical research results

Problem

Distributed development has resulted in interoperability issues that could be automatically detected

Solution

A continuous integration system, built on the CloudyCluster platform, to monitor OBO development globally and detect interoperability issues when they first arise

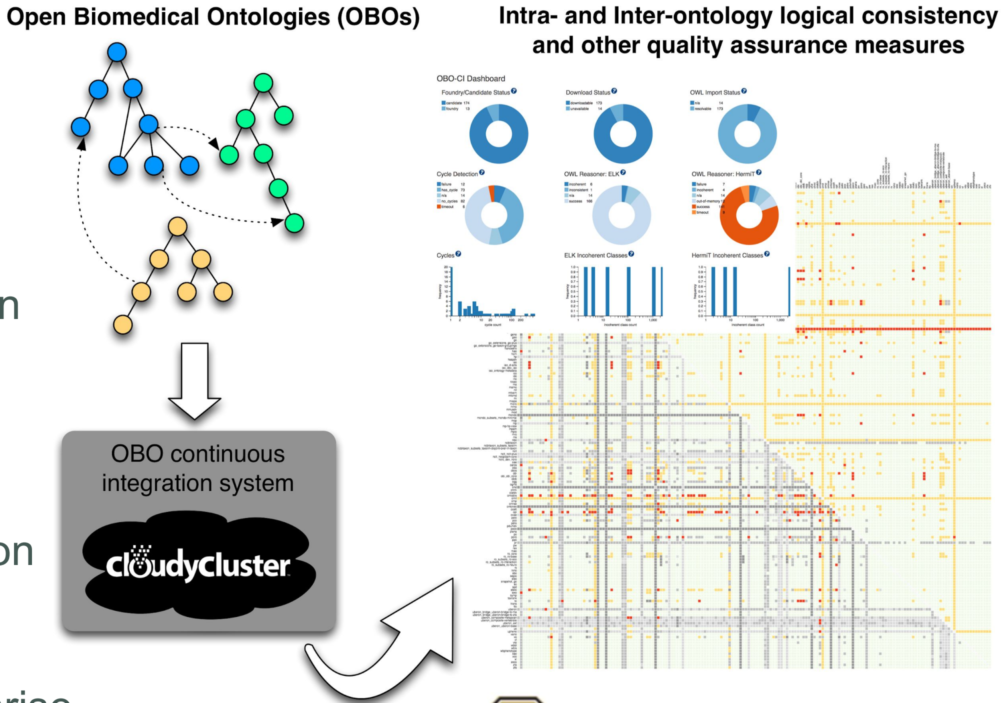

School of Medicine
UNIVERSITY OF COLORADO ANSCHUTZ MEDICAL CAMPUS

How to get started

1. Cloud HPC orchestration

What you need for a fully functioning HPC environment in GCP

- GCP Agreement and Account
- Launch a CloudyCluster instance from the GCP Marketplace
- setup Admin user and credentials
- Choose the size if your environment (Test, SM, M, L, XL)
- Customize options for the environment -> Select Create
- CloudyCluster automatically creates a Private Network, Login, Scheduler (Slurm or PBS), and Parallel Storage all ready to run HPC jobs.
- Upload your data / Tweak your HPC job script
- Run your job and Download Results.

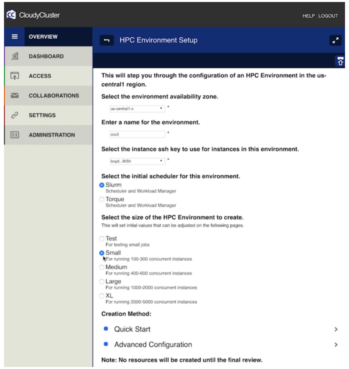

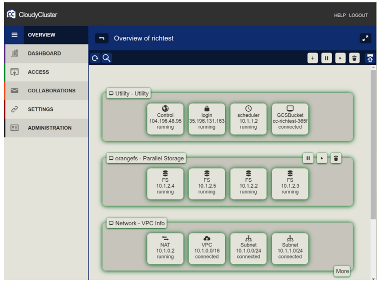

Google Cloud

CloudyCluster

DEMO - Video

1. Control Instance Creation
2. Environment Creation

DEMO - Video

3. Job submission through the CCQ Meta-Scheduler to SLURM
4. CCQ instance clean up once the Jobs have completed

Key Contacts

Jim Bottum - JB@omnibond.com

Boyd Wilson - boyd@omnibond.com

Amy Cannon (Sales) - AmyCannon@omnibond.com

Support - support@cloudycluster.com

CloudyCluster

Thank You

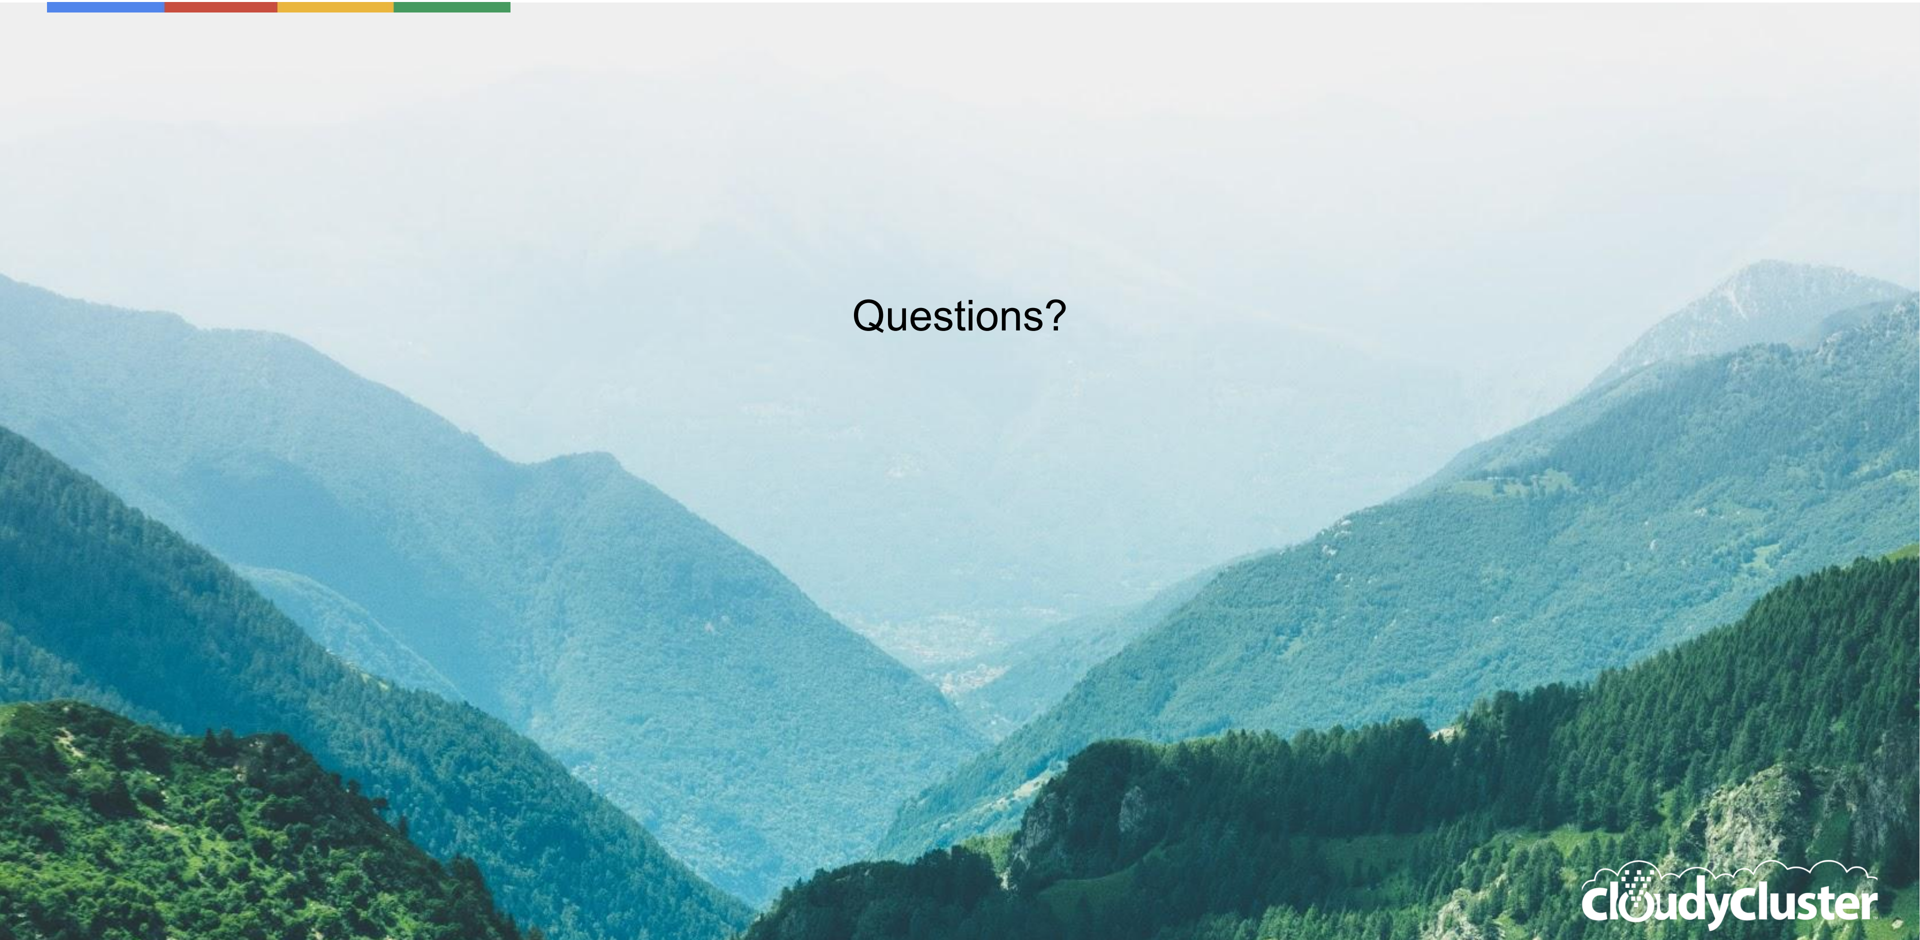

Questions?