THE GEORGE WASHINGTON UNIVERSITY
WASHINGTON, DC

Advancing data intensive science at GWU

Clark Gaylord
Director, Research Technology Services
George Washington University
cgaylord@gwu.edu

Thanks to: Glen MacLachlan, Adam Wong, Jarda Flidr, Terrence Lewis

THE GEORGE WASHINGTON UNIVERSITY
WASHINGTON, DC

Research Technology Services Portfolio

- Traditional HPC cluster: Pegasus 210 node, hybrid CPU/GPU, Slurm, IB, etc
- Teaching HPC cluster: Cerberus 12 node, focus on use in classroom
- HTCondor: "pilot" 8 node deployment
  - Entre to, working with Open Science Grid
- CApital Area Research and Education Network (CAAREN) regional Internet2
- Cloud services: mostly engineered, purpose-built solutions for special use cases
- Many boutique solutions for research group clusters, databases, visualization
- Data management, protected data environments
  - REDCap
  - Globus with Google Drive, Box, S3 connectors
- Solution consultancy

THE GEORGE WASHINGTON UNIVERSITY
WASHINGTON, DC

A good fit for HTC

- Pegasus is heavily utilized – often over 85% allocated.
- Like many clusters, most jobs have relatively modest resources
  - More than 91% of our jobs finish in under 1 hour.
  - More than 97% of our CPU-based jobs run on a single node.
  - Approx 1/2 of a recent sample of our jobs require less than 10 GB of RAM.
  - [everything is pareto distributed]
- This profile of jobs fit nicely into the HTC mold.

THE GEORGE WASHINGTON UNIVERSITY
WASHINGTON, DC

Current configuration

GWU HTCondor is an on-prem 9-node Dell cluster (Power Edge R730)

- 1 Physical manager node
- 8 Physical worker nodes
- Each node features:
  - 28 Intel Xeon E5-2680 Cores
  - 128 GB Memory
  - 150 GB scratch SSD Storage

THE GEORGE WASHINGTON UNIVERSITY
WASHINGTON, DC

Current growth

Addition of P100/V100 GPUs repurposed from Pegasus login nodes, etc.

Augmenting the original HTCondor deployment for higher storage deployment.

- Additional VMs – highly scalable deployment (begin with 8)
  - OpenStack hosts gives flexibility for OS "pseudo-bare metal"
- Storage
  - 1.2PB Gluster over ZFS
  - 3.2GB/s aggregate I/O
- IPv6 connectivity through CAAREN (Equinix facility)
- 4 additional GPUs

THE GEORGE WASHINGTON UNIVERSITY

Pegasus HPC

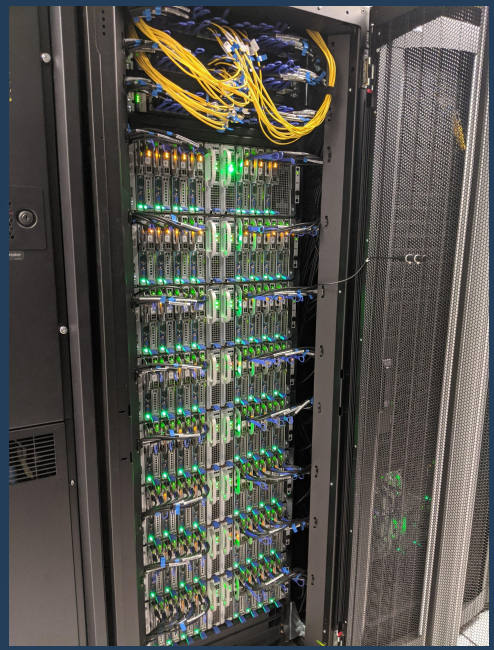

[actually this is Colonial One]

THE GEORGE WASHINGTON UNIVERSITY

WASHINGTON, DC

[insert Pegasus photo here]

THE GEORGE WASHINGTON UNIVERSITY
WASHINGTON, DC

A Science Driver from Genomics

- Prof. Anelia Horvath from Milken Institute School of Public Health at GW – Functional Analysis of SNVs (Single Nucleotide Variants) from single-cell RNA sequencing data.
- Single cell RNA sequencing (cell-level transcriptome studies) in cancer genomics.
  - Aim: develop methods to assess functionality of SNVs through their relationships with dynamic transcriptome traits from cancer scRNA-seq data.
  - Data source: generated by the 10x Genomics platform – a high-throughput single-cell sequencing platform.
  - Datasets: consisted of ~9K cells per individual, and ~150K sequencing reads per cell, and ~150nt/read; which produces a single file of ~1-1.5 billion sequencing reads.
  - Software: R/Python/cellRanger(from 10x Genomics).

THE GEORGE WASHINGTON UNIVERSITY
WASHINGTON, DC

A Science Driver from Genomics

- cellRanger is a set of analysis pipelines that process Chromium single-cell RNA-seq output to align reads, generate feature-barcode matrices and perform clustering and gene expression analysis.
- The pipelines of tools are:
  - cellRanger mkfastq – convert raw data fastq format data
  - cellRanger count – perform tasks such as alignments, filtering and data clustering
  - cellRanger aggr – aggregate results from multiple runs of cellRanger count
  - cellRanger reanalyze – perform secondary analysis by fine-tuning parameters

THE GEORGE WASHINGTON UNIVERSITY
WASHINGTON, DC

A Science Driver from Genomics

- The problem size (~1.5 billion sequencing reads) of this genomic data analysis demands for a great amount of computational resources.
- The data processing nature of such computational task makes it easily be broken into numerous independent prices (i.e. HTC-able) and thus it is a good candidate for a HTCondor cluster!
- A cellRanger workflow would fit well to HTCondor's Workflow Model where its DAGMan (Directed Acyclic Graph Manager) manages dependencies between jobs automatically.
- These cellRanager workflows of computation allow us to experience and learn about the advanced feature of HTCondor in workflow construction.

THE GEORGE WASHINGTON UNIVERSITY
WASHINGTON, DC

A Science Driver from Physics

- Prof. William J. Briscoe from the Physics Department at GW is working on an intermediate energy nuclear physics experiment at the Thomas Jefferson National Accelerator Facility (JLAB) using the CEBAF Large Acceptance Spectrometer (CLAS).
- JLAB can provide either continuous electron and photon beams with energies up to 12 Giga Electron Volts (12 GeV) to study properties of quarks in hadrons and mesons.
- Experiments are extremely expensive and time-consuming for hundreds of physicists. Careful simulation are paramount.

THE GEORGE WASHINGTON UNIVERSITY
WASHINGTON, DC

A Science Driver from Physics

- Researcher simulates part of an experiment using GEant4 Monte Carlo (GEMC)
- GEMC is a C++ application used by physicists to model the interaction of charged particles passing through matter such as incident, scattered, and recoil particles passing through targets and detectors.

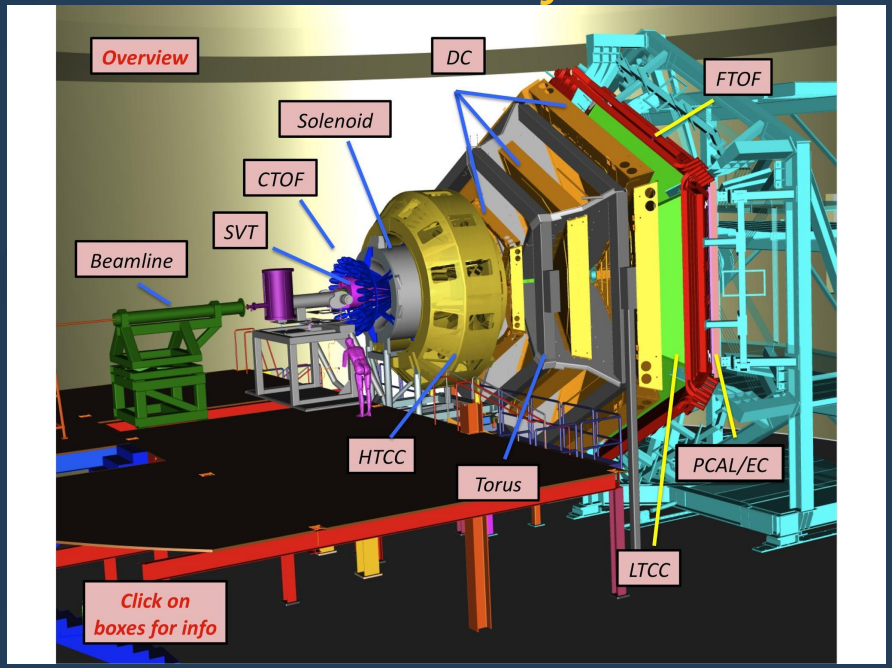

THE GEORGE WASHINGTON UNIVERSITY
WASHINGTON, DC

A Science Driver from Physics

- CLAS delivers GEMC and related data as both Docker and Singularity images via CVMFS.
- The simulation needed to be completed in 2-3 months and required approx 4.3 million core hours to complete which is a little more than GW could provide locally or what OSG could provide opportunistically.
- OSG allocated additional resources and simulation completed on schedule.

THE GEORGE WASHINGTON UNIVERSITY
WASHINGTON, DC

Omics data, Medicine and Public Health

Prof. Eric Vilain from the Children's National Hospital (DC) with the Institut National pour la Recherche Biomédical of the Democratic Republic of the Congo.

The objective is to ascertain microbial role in the Konzo disease and then to understand microbiota in settings such us different towns, cities and regions in Congo. Consequently, actionable genomics experiments could be performed to advance precision medicine in DRC

THE GEORGE WASHINGTON UNIVERSITY
WASHINGTON, DC

United States – DR Congo Collaboration

- GWU/ Children's National Hospital
  - Jonathan LoTempio, BS
  - D'Andre Spencer, BA, MPH
  - Neerja Vahist, BS
  - Matt Bramble, PhD
  - Eric Vilain, MD, PhD
- Institut National pour la Recherche Biomédical
  - Kizito Mosema, MD, MPH
  - Kevin Karume, BS
  - Désiré Tshala Katumbay, MD, PhD
  - Jean-Jacques Muyembe, PhD
  - Dieudonne Mumba, MD, PhD

THE GEORGE WASHINGTON UNIVERSITY
WASHINGTON, DC

Vast amount of omics data

To increase representation of Congolese individuals, about 100-1000 high quality references genomes will be generated via the advanced sequencing technologies including:

- High coverage Illumina sequencing
- High coverage PacBio HiFi sequencing
- Bionano optical mapping

THE GEORGE WASHINGTON UNIVERSITY
WASHINGTON, DC

Huge computation demand

- Mapping of sequences need to be done on 25 TB of data by end of 2021 and another 100 TB of data by the following 1-2 years.
- Common alignment tools such as Samtools and Minimap2 will be used.
- These can be carried out efficiently in a HTCondor cluster.

THE GEORGE WASHINGTON UNIVERSITY
WASHINGTON, DC

Serving the research community

- Working with HTCondor community (and OSG) leverages collaboration to bring resources to the table that are outside scope for HPC
- Continuing with HTCondor deployment to contribute to and partner with OSG

- Flagship HPC cluster continues to be strategic for many research applications
  - Not only MPI…
- Much of GWU research is naturally high throughput – which often *can* adapt to HPC but this isn't necessarily ideal

THE GEORGE WASHINGTON UNIVERSITY
WASHINGTON, DC

Thanks

We really appreciate the effort from the HTCondor Team in supporting our HTCondor project at GW. A special thanks goes to John M Knoeller (TJ) for meeting regularly with us and in helping us to solve any Condor related problems.