

Breakout Session 3: Track B

Exploration of Cloud-based High Performance Computing

Dr. Albert Lai (Moderator)
Professor, Washington University in St. Louis

Exploration of Cloud-based High Performance Computing

Albert M. Lai, PhD, FACMI, FAMIA

Chief Research Information Officer, *Washington University School of Medicine*

Deputy Director, *Institute for Informatics, Data Science, and Biostatistics*

Professor, *Department of Medicine*

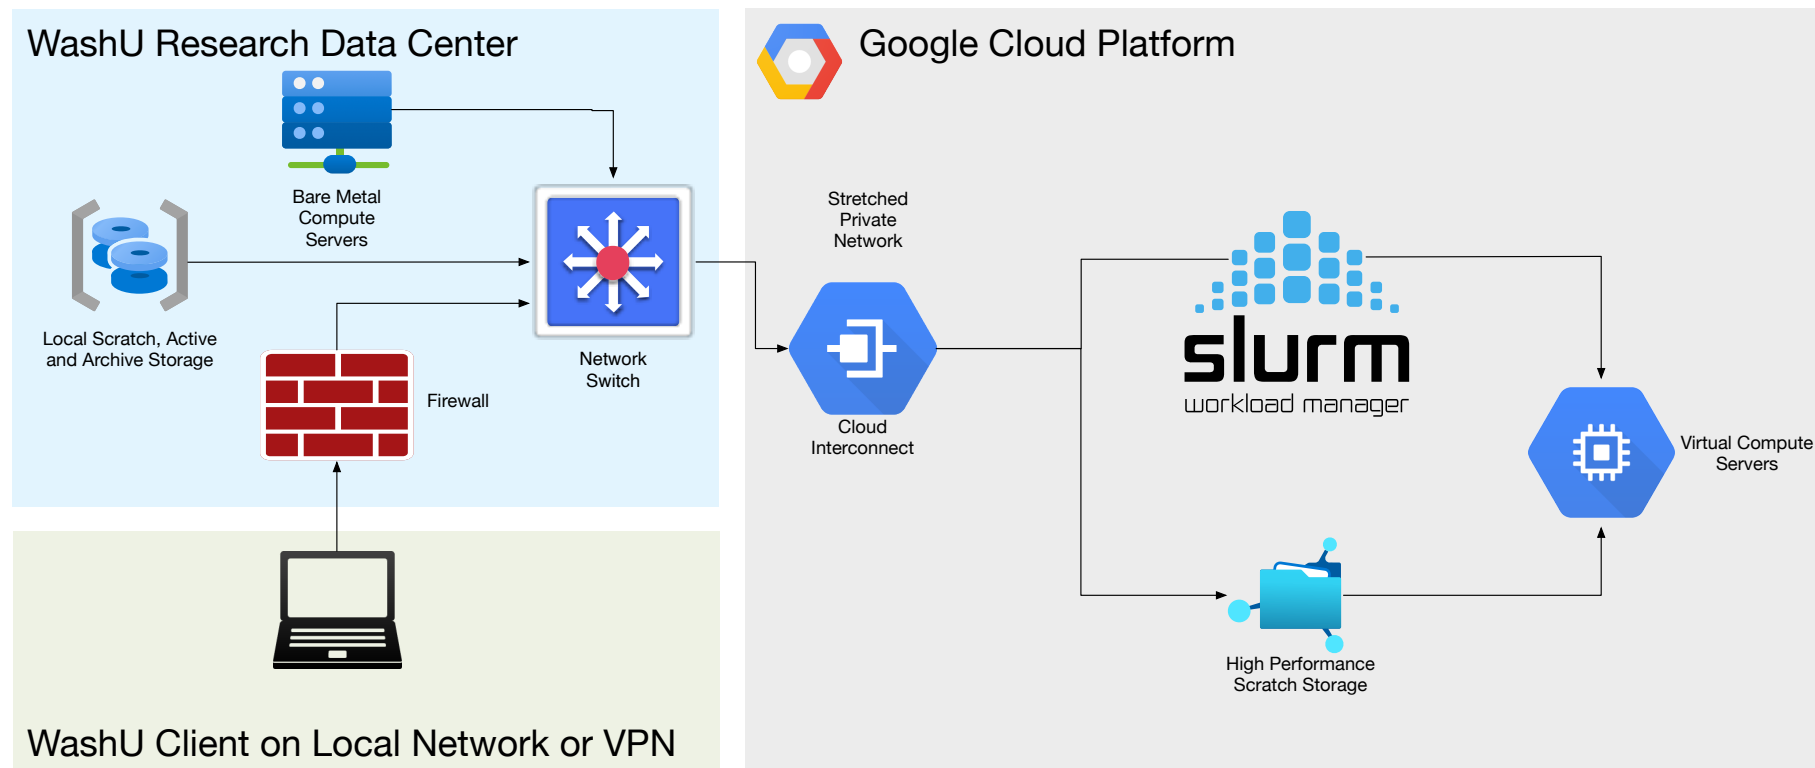
Professor, *Department of Computer Science & Engineering*

Specific Aims

- **Aim 1:** Deploy a hybrid-cloud HPC cluster to the public cloud that is seamlessly integrated with on-premises worker nodes and data storage.
- **Aim 2:** Run tests on a range of HPC workloads, including Data Simulations, Bioinformatics, Artificial Intelligence/Machine Learning & Imaging (GPU accelerator-focused), and general interactive workloads. We will compare the performance of workloads running on cloud and on-premises nodes, utilizing job run time and cost as the primary metrics.
- **Aim 3:** Evaluate the viability of deploying HPC clusters and nodes to the public cloud. We will complete this by examining the cost of running the workloads described in Aim 2 and comparing them to the projected costs of traditional on-premises approaches.

Hybrid HPC Cluster using GCP Toolkit

- Given the short timeline, needed a fast way to deploy an HPC Cluster in the Cloud.
- GCP Toolkit simplifies the deployment and management of (HPC) clusters on GCP
- Automates cluster setup, configuration, scaling, and monitoring.
- Reduces deployment complexity, accelerates time-to-results, and enhances resource efficiency using the Slurm workload manager



Workloads	Description / Workload type
ArCCH-WGS-VariantCalling	Variant calling pipeline for low Variant Allele Frequency (VAF) Clonal Hematopoiesis
GATK-SV	Structural variation discovery pipeline for WGS
Tensor Machine Learning Model	GPU-focused
Peptide Simulation using CAMPARI	CPU-focused simulation workload
Weather Research & Forecasting Model (WRF)	MPI-type workload
VNC and noVNC Desktops	Interactive workload

Experiments

Workloads	# Experiments Executed	# Tasks
Dr. Kelly Bolton Lab	13	2,486
Dr. Peter Jin Lab	78	275,572
Dr. Joshua Swamidass Lab	4	63
Dr. Alex Holehouse Lab	4	11,916
Dr. Jian Wang Lab	4	909
Interactive	2	12
Total	105	290,958

Experiments

- 110 Experimental clusters executed
- Permutations of
 - workload type
 - # data sets
 - dynamic/preconfigured server
 - preemptible/persistent configurations

Early Observations

- GCP Infrastructure
 - Limit of 7 Slurm clusters per GCP Project (subscription) metadata limitation
 - Regular service disruptions cause job failures every Sunday
 - 70% of Preemptible VMs run less than an hour
 - 14% of Preemptible VMs run for at least 22-24 hours

Early Experimental Results

- 56 Experimental clusters completed successfully
 - Preemptible VMs and infrastructure instability resulted in incomplete experiments
- CPU utilization is higher with VMs configured to match on-premises hosts
- Experiment runtime is less for dynamic CPU resources (VMs sized based on job resource parameters)
- ~50% savings on preemptible VMs
- Compute runtime impacts storage and total costs of experiments
- GCP Filestore (ephemeral scratch storage) cleanup impacts costs of experiments
 - Lack of policy or automation not integrated into workloads replicating on-premise use case

Questions?



Office of Health Information
and Data Science

sites.wustl.edu/ohids