

# Genomics and HPC: Use cases and challenges

**hpc-ch forum on User-Centric view on HPC**

**October 6, 2022**  
**Campus Biotech**

- » DAIP activities
- » Why and how do we use HPC?
- » Are we happy with it?

# 1.

## Data Analysis and Interpretation

Activities and the need for HPC

## Data Analysis and Interpretation: Activities

4

- » Standardized/accredited sample processing e.g.
  - » Human genome/exome sequencing (germline DNA)
  - » RNA sequencing (gene expression)
  - » COVID-Seq (monitoring variants)
- » Non- or less- standardized processing:
  - » Downstream analysis for research projects
  - » Internal analytics and R&D
- » Tangential to HPC:
  - » Data delivery and storage
  - » Variant interpretation platform



## Data Analysis and Interpretation: Need for HPC

5

- » Many tasks are computationally expensive e.g.
  - » Read mapping
  - » Variant calling
- » Requirements
  - » Large projects (250 samples in 6 weeks)
  - » Processing time (COVID-Seq, Rapid WGS, internal QC)
  - » Clinical data must be on-premise, privacy concerns
  - » Ease of use

# 2.

## Infrastructure overview

What do we mean by «HPC»?

- » Multiple computation nodes
- » Shared storage
- » Scheduler
- » Common software stack

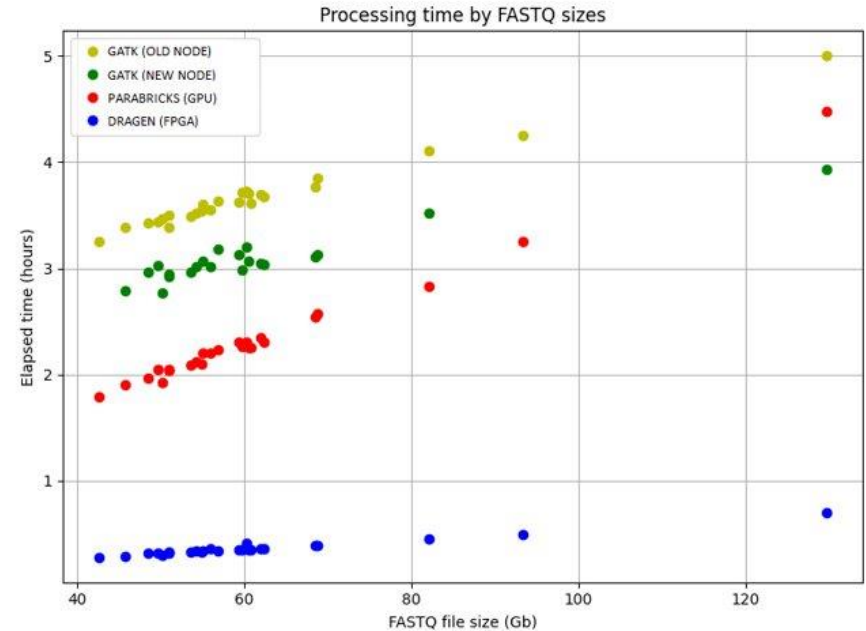
## Internal HPC cluster

8

- » 20 + 1 nodes
  - » 20 CPU nodes
    - 28 cores Xeon E5-2680
    - 250GB RAM
    - 1.5TB local scratch
  - » 1 FPGA node



- » 1 Isilon: NFS, 1.7PB
- » SLURM



Ilya Kolpakov, Arkadiy Shevrikuko @GenomeCenter



- » environment-modules
- » Installed with spack
- » Need a way for reproducible installations
- » Mess of modulefiles


```
[ahungler@frt1 ~]$ module av 2>&1 | wc -l  
1201  
[ahungler@frt1 ~]$ time module av samtools  
real 0m10.157s
```



## Network overview


10

- » Split internal/DMZ
- » Sensitive data and main workflow in internal network
- » ISO 15189 Medical laboratories
- » Data delivery in DMZ
- » Other non-accredited applications in DMZ




## Data flow

11




- » Size can vary
- » For Whole Genome Sequencing (WGS), order of 1TB/run





- » Customer brings the expertise (Tom McKee @HUG, Renku, Congenica)
- » We run them in the DMZ
  - » Virtual machine
  - » Or kubernetes cluster
  - » Or... depending on community needs



- » Need for more flexible usage (containers, VDI, ...)
- » Computation nodes: 5 hardware servers
- » Shared storage: (full-flash providing NFS, k8s volumes)
- » Scheduler: k8s or VM or manual
- » Software depending on usage:
  - » Container
  - » Stack from internal cluster
  - » User managed on VM



# 3.

## HPC Usage

- Shared FS, centrally deployed software and the scheduler makes leveraging nodes straightforward
- » Good fit for research/incremental development

Downsides for "production":

- » The cluster is a "snowflake":  
Developing/testing elsewhere is difficult
- » FS is a poor fit for atomic transactions:  
When something is "computed"?
- » Monitoring is rudimentary, debugging hard  
(unless jobs standardized and documented)



- » Data management system enabling standardized data capture, querying and removal
- » Tools for running production workflows at higher level of abstraction
- » Workflows implemented in terms of reasonably structured code amenable to isolated testing

- » Managing the complexity of workflows consisting of multiple processing steps
- » Typical features:
  - » Workflow definition language or library
  - » Caching system
  - » Abstraction over the computing resources
  - » Monitoring and debugging support

- » Based on a task library handling scheduling, caching (excellent API except for debugging)
- » Workflows are *strongly typed data* moving over a *static flow graph* (Akka streams)
- » *Rigid system difficult to change*

- » External workflow system (Camunda BPM)
- » DB-backed storage for data/tasks/results (caching, querying, garbage collection)
- » Workers pull tasks from Camunda and push results to the DB (via HTTP API)

# HPC Usage Camunda BPM (Workflow System)

21

The screenshot displays the Camunda Cockpit interface. At the top, the navigation bar includes 'Camunda Cockpit', 'Processes', 'Decisions', 'Human Tasks', and 'More'. The current view is 'Dashboard > Processes > ... > DragenSampleAnalysis : 61db9e43-4492-11ed-890c-02423d8ef1a2 : Runtime'. The main area shows a workflow diagram with three tasks: 'Prepare for dragen call', 'Run dragen', and 'Post computation results'. The 'Run dragen' task has a blue '1' and a red '1' indicator. Below the diagram is a 'Variables' section with a table of process variables.

Name	Type	Value	Scope	Acti...
computation_uid	Json	"14025d90-b93a-0af2-50a2-6e87f0...	DragenSampleAnalysis	✓ 🗑
dragen_argv	Json	["/opt/edico/bin/dragen", "--enable-...	DragenSampleAnalysis	✓ 🗑
output_dir	String	/tmp/store/61/db/9e/61db9e43-449...	DragenSampleAnalysis	✓ 🗑

Date and Time displayed in local timezone: Europe/Zurich

Powered by Camunda Platform / v7.16.0

- » Traditional "HPC" is a foundation
  - » Suitable for interactive usage or relatively simple workflows
- » For production one needs data and workflow management systems
  - » "HPC" tools alone insufficient at addressing operational and implementation complexity