UManitoba Spring HPC and Cloud Workshop

May 17-19, 2023

University of Manitoba

TRADITIONAL TERRITORIES ACKNOWLEDGEMENT

The University of Manitoba campuses are located on original lands of Anishinaabeg, Cree, Oji-Cree, Dakota, and Dene peoples, and on the homeland of the Métis Nation.

We respect the Treaties that were made on these territories, we acknowledge the harms and mistakes of the past, and we dedicate ourselves to move forward in partnership with Indigenous communities in a spirit of reconciliation and collaboration.

Research Computing Resources for UM

- **ARC resources available locally and nationally**
  - Organizations providing DRI in Canada
  - High-performance computing resources, local UM and National
  - Community cloud resources
- **Beginner introduction to use**
  - Linux shell; High Performance Computing with SLURM ; OpenStack cloud
- **Intermediate notes and best practices**
- **Further information, support, documentation**

Photo of Grex by Jin Michael Uy, ECE , 2022

Programme of the Workshop

May 17	Title	Presenter	Start time	End time
1	Updates on Research computing resources	Grigory Shamov	10:30	11:00
2	Basics of Linux Shell	Stefano Ansaloni	11:00	11:50
3	(break)		11:50	12:00
4	Beginner introduction to using HPC machines	Ali Kerrache	12:00	1:30

Programme of the Workshop

May 18	Title		Start time	End time
1	Beginner HPC software overview	Ali Kerrache	10:30	11:50
2	(break)		11:50	12:00
3	Using OpenOnDemand Web portal on Grex, demo	Grigory Shamov	12:00	12:40
4	Using HPC clusters efficiently (SLURM, kinds of jobs)	Ali Kerrache	12:40	1:20

Programme of the Workshop

May 19	Title		Start time	End time
1	Advanced HPC software (Containers, CVMFS)	Ali Kerrache	10:30	11:00
2	Using GPUs on HPC systems, demo	Grigory Shamov	11:00	12:00
3	(break)		12:00	12:10
4	Beginner's introduction using Cloud computing (OpenStack)	Stefano Ansaloni	12:00	13:30
5	Overview of further training and materials available	Ali Kerrache	1:30	1:40

What is Advanced Research Computing (ARC)?

Advanced Research computing (ARC) refers to both “resources” and “modes of use”

- Focuses on enabling computational research.
- Provides capabilities that are not available with common (desktop, enterprise server) computing environment:
    - CPU time and memory,
    - storage capacity and performance
    - network resources optimized for research
- Has a specialised set of tools like parallel computing, data analysis, visualization, and resource management
- ARC resources tend to be large and thus expensive (but efficient)

The two most popular modes of delivery are “Traditional **HPC**” and “**Cloud** computing”

Digital Research Alliance of Canada

Digital Research Alliance of Canada

- Replaces an earlier National DRI organization, ComputeCanada 2020-2025
  - Mandate to integrate ARC, RDM, and RSE
- Took over CC's operations/security of National Hosting sites, National Teams etc.
  - https://www.alliancecan.ca/en/services/advanced-research-computing/account-management/policies
- Works with DRAC Federation to support the National cloud and HPC systems

- Submited a Funding Proposal in late 2022, approved in 2023
  - Includes infrastructure renewal (~220M) for HPC and Cloud
  - https://www.alliancecan.ca/en/initiatives/dri-investments

- Works on improving security of National DRI
  - Develops and publishes Security policies, data handling standards etc.
  - Multi-factor authentication (MFA) is coming!

Canada's Advanced Research Computing Platform

DRAC. https://alliancecan.ca

Digital Research Alliance of Canada

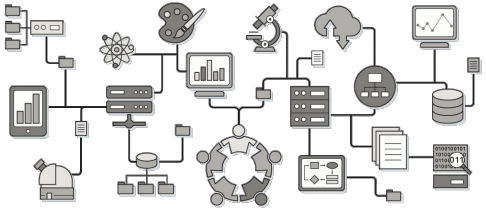

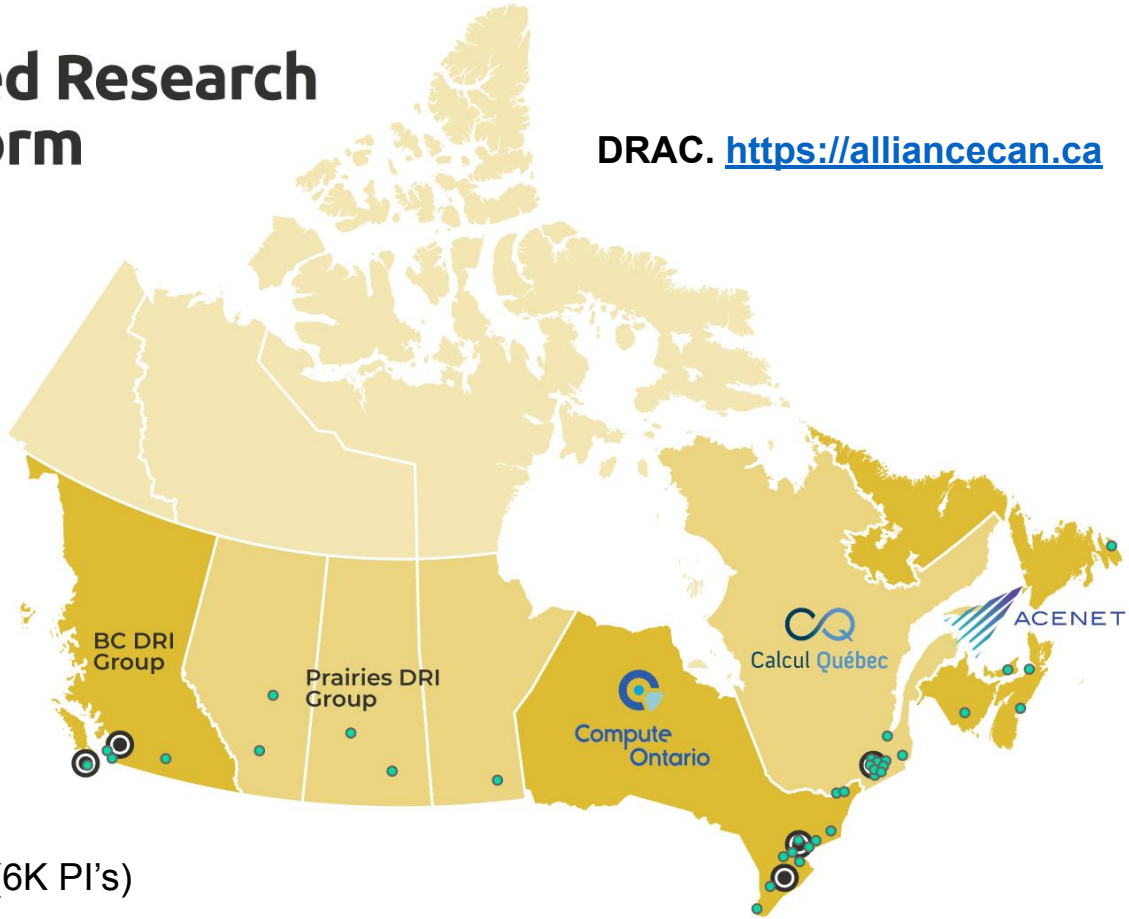

Number of "roles" in CCDB : **22K** (6K PI's)

Number of "roles" in Prairies: **2766** (800 PIs), of them in UM 347 (113 PIs)

What is Advanced Research Computing (ARC)?

Advanced Research computing (ARC) refers to both “resources” and “modes of use”

- Focuses on enabling computational research.
- Provides capabilities that are not available with common (desktop, enterprise server) computing environment:
    - CPU time and memory,
    - storage capacity and performance
    - network resources optimized for research
- Has a specialised set of tools like parallel computing, data analysis, visualization, and resource management
- ARC resources tend to be large and thus expensive (but efficient)

The two most popular modes of delivery are “Traditional **HPC**” and “**Cloud** computing”

Parts of an HPC system (like Grex)

Scheduler ; user software

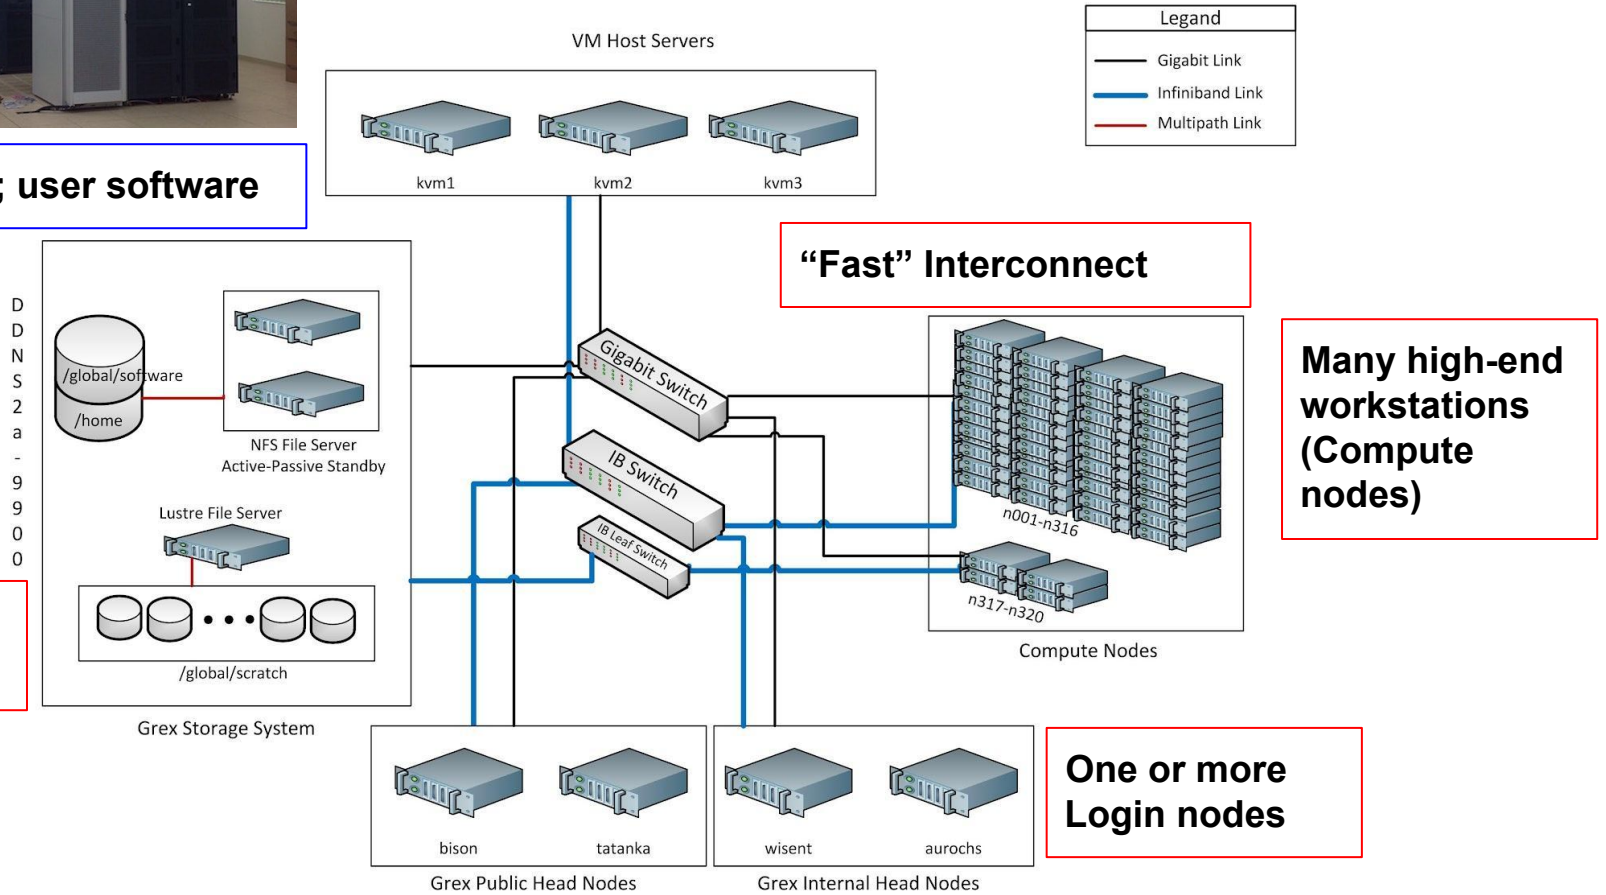

Parallel Storage

Many high-end workstations (Compute nodes)

One or more Login nodes

HPC vs Cloud computing

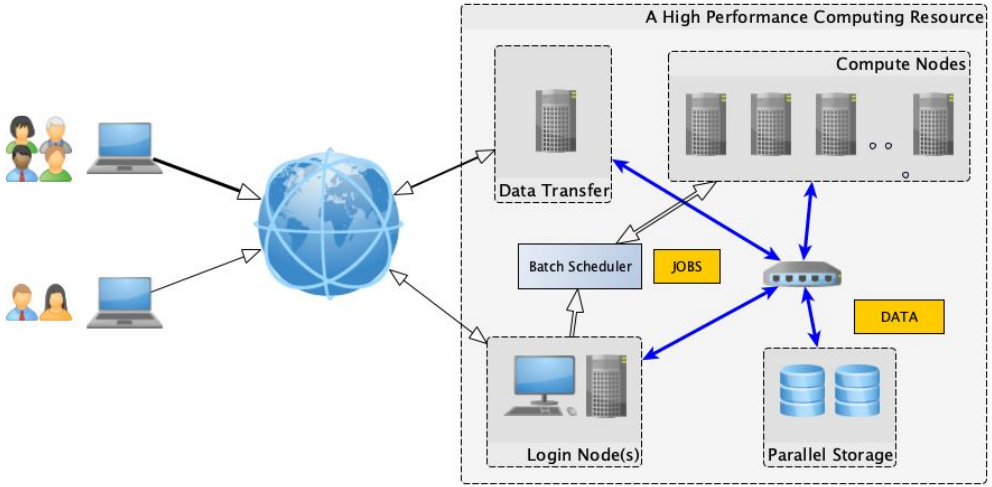

- HPC clusters are *shared* systems with r*emote access*
- Batch mode of usage
- A central Software delivery on HPC
- Dealing with Data (storage, transfer etc.)

- Public and Community Cloud computing
- Flexible, elastic, Provides isolation of tenants
- SDN, SDS, Virtualized compute
- Self-service

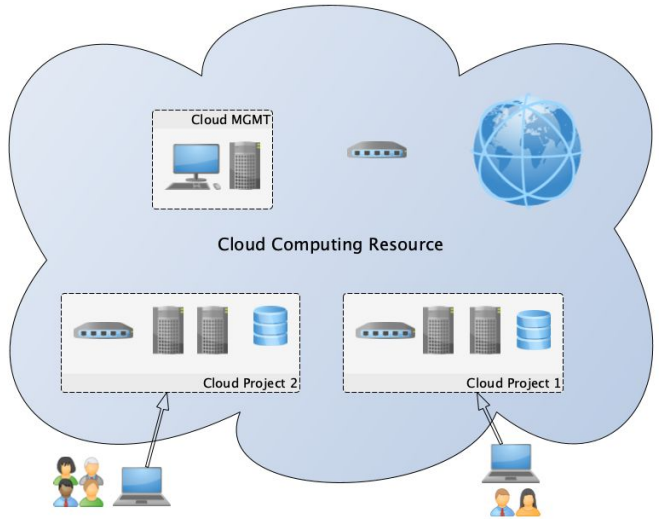

Size of National DRI systems

Digital Research Alliance of Canada

System, kind	Commission date	CPU cores	GPU devices	Project storage, PB
Arbutus, Cloud	09. 2016	16,008	108	12
Beluga , HPC	09. 2019	28,960	688	17
Cedar, HPC	03. 2017	94,528	1352	19.5
Graham, HPC	06. 2017	34,784	498	13
Narval, HPC	09. 2021	61,760	524	14
Niagara / Mist HPC	03. 2018	75,840	64	2.7
	Total:	**295,872**	**3,126**	**66.2**

Canada's Advanced Research Computing Platform

DRAC. https://alliancecan.ca

Digital Research Alliance of Canada

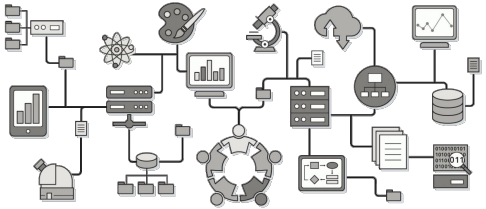

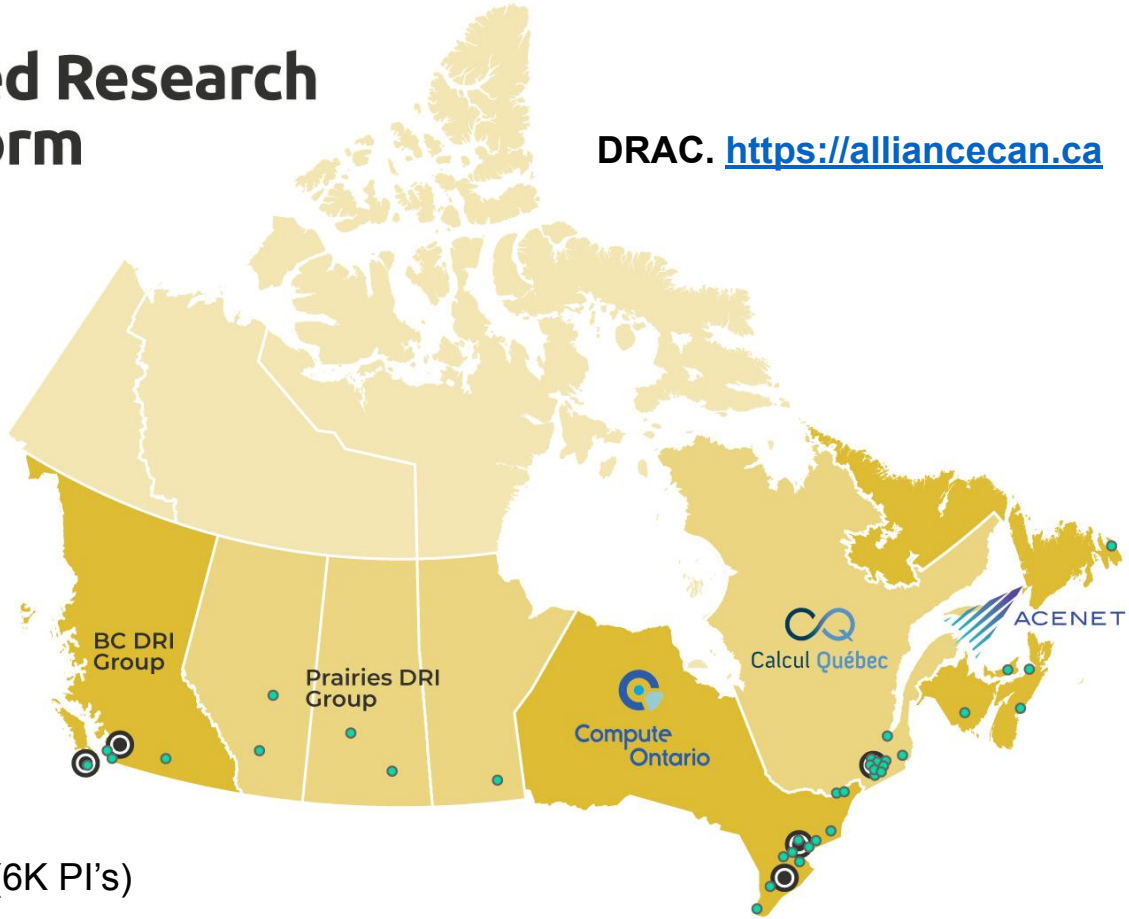

Number of "roles" in CCDB : **22K** (6K PI's)

Number of "roles" in Prairies: **2766** (800 PIs), of them in UM 347 (113 PIs)

Grex, our local HPC system

A step-up machine for UM users to DRAC resources

Hardware:

- The “legacy” CPU part from 2010, **~2400** cores left
- UM added **2820** cores and 2 GPU nodes in 2019-2021
- Contributed systems (GPU and CPU nodes), 32 GPUs
- Storage : **15 TB** NVMe (/home);
  1.1 PB Parallel Lustre FS (/project replaced old /global/scratch)

- Authentication and support mainly through CC/DRAC systems
- Managed by the same local DRAC Federation team.
- Provides both a local and the ComputeCanada software stacks
- Some serviced that DRAC does not yet provide (OnDemand)

Storage is not just data space: ARC vs RDM

1. ARC focuses on Active Research stage, when data is actively processed/analyzed/generated
   a. Research is some “modeling” typically involving intensive data access
   b. “*HPC is a way to convert computational bottlenecks into storage bottlenecks*”
   c. “*~90% of HPC outages are due to a storage problem*”
   d. The motto is “*moving compute close to storage*”
   e. Also, data transfer to/from HPC systems is optimized (Globus, ScienceDMZs)
2. In the RDM world, things like FAIR Principles etc. developed first in Humanities and Libraries, where Research can be, collecting and reading papers.
3. In the Enterprise IT world, typical is NAS storage, a.k.a. network disk, which is large, has Enterprise features like backup and high uptime
4. Sharing storage like DropBox, OneDrive and NextCloud is for collaboration on documents, and very popular.

(How do 1 - 4 above connect to each other?)

Security: are ARC systems secure?

A shared responsibility between researchers and system admin staff of ARC systems

- Security entails both organizational and technical controls.
  - DRAC is working on increasing the organizational infosec policies
  - Extra Security Analysts staff to monitor vulnerabilities etc.
  - Improving security by adding SSH key auth to CCDB (DRAC, Niagara, Grex)
  - Implementing Duo MFA for SSH everywhere (for DRAC staff in place).

- Technology requirements and compliance for Medical Research are more difficult
  - Data Leak protection
  - Logging every access to data
  - Encryption at rest
  - Hard to accommodate within a single, shared HPC system


University
of Manitoba