

# UManitoba Fall HPC and Cloud Workshop

November 6,7, 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 National resources

- **Beginner introduction to use**

- Linux shell; High Performance Computing with SLURM ; Software and Containers
  - OpenStack cloud

- **Further information, support, documentation**



Photo of Grex by Jin Michael Uy, ECE , 2022

# Programme of the Workshop

<b>Nov 6</b>	<b>Title</b>	<b>Presenter</b>	<b>Start time</b>	<b>End time</b>
1	Updates on Research computing resources	Grigory Shamov	10:30	11:10
2	Getting accounts, connecting, using Duo MFA	Ali Kerrache	11:10	11:50
	(break)		11:50	12:00
3	Basics of Linux Shell	Stefano Ansaloni	12:00	12:50
	(break)		12:50	13:00
4	Beginner running HPC jobs with SLURM	Ali Kerrache	13:00	14:00
5	Using HPC with OpenOnDemand Web portal	Grigory Shamov	14:00	14:30

# Programme of the Workshop

<b>Nov. 7</b>	<b>Title</b>		<b>Start time</b>	<b>End time</b>
1	Beginner HPC software overview	Ali Kerrache	10:30	11:30
	(break)		11:30	11:35
3	Using containers in HPC (Docker, Singularity, Apptainer)	Grigory Shamov	11:00	12:40
	(break)		12:40	12:50
4	Basics of using OpenStack cloud	Stefano Ansaloni	12:50	1:50
5	Getting most of HPC system	Ali Kerrache	1:50	2:30

# 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”




- 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
- Renewal of the ARC:
  - Includes infrastructure renewal (~220M) for HPC and Cloud
  - <https://www.alliancecan.ca/en/initiatives/dri-investments>
- Improving security of the National DRI
  - Develops and publishes Security policies, data handling standards etc.
  - Multi-factor authentication (MFA)

# Canada's Advanced Research Computing Platform

DRAC. <https://alliancecan.ca>

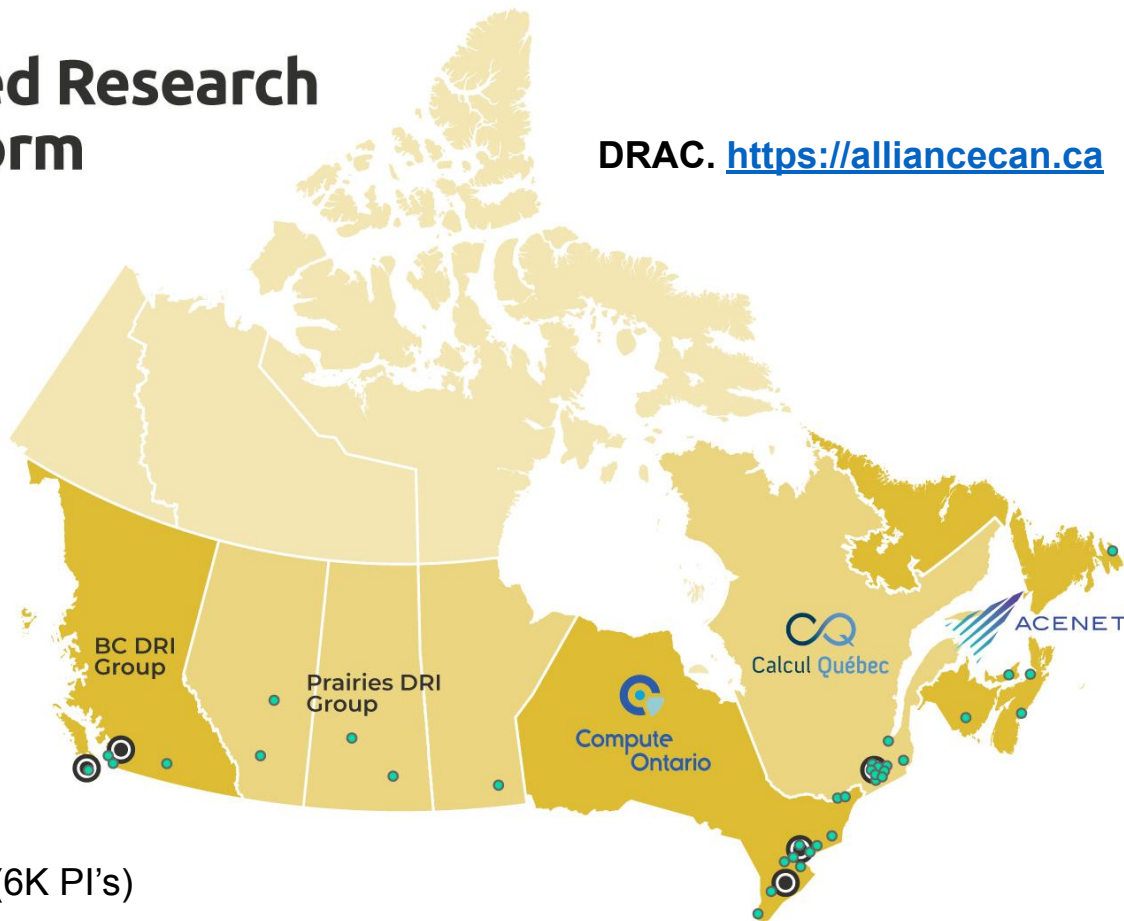


Digital Research Alliance of Canada



🕒 National Host Sites

● Support Sites



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

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




# 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
	<b>Total:</b>	<b>295,872</b>	<b>3,126</b>	<b>66.2</b>

# HPC vs Cloud computing



- HPC clusters are *shared* systems with *remote 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



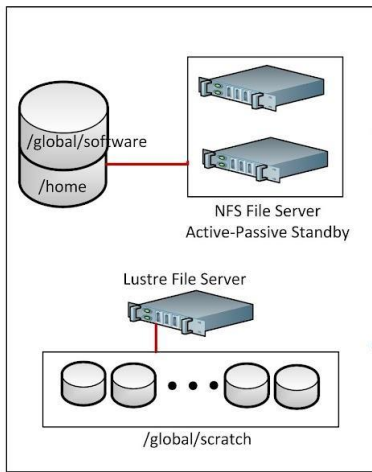


# Parts of an HPC system (like Grex)


Scheduler ; user software

D  
D  
N  
S  
2  
a  
-  
9  
0  
0  
0

Parallel Storage



Grex Storage System



Legend	
	Gigabit Link
	Infiniband Link
	Multipath Link

“Fast” Interconnect

Many high-end workstations (Compute nodes)

One or more Login nodes

# Services provided by the AF / DRAC

- HPC resources (SLURM, software)
- CVMFS CC software stack
  - Includes software, genomics datasets, AlphaFold etc.
- Globus file transfer service (still using CC endpoints and CC org)
- Platforms access to HPC and/or cloud
  - JupyterHub; <https://jupyterhub.cedar.computecanada.ca/>
  - GenAP project <https://genap.ca/>
- Grant consultation (CFI IF, JELF)
- Support; single National Helpdesk at [support@tech.alliancecan.ca](mailto:support@tech.alliancecan.ca)
  - Has ~100 HPC Analysts, may have some domain specific knowledge
  - HEP/Subatomic, MolModSim, Genomics Teams

# Updates for DRAC ecosystem : Renewal

- DRAC's proposal was accepted , ~200M including provincial matching funds
  - Only including existing HPC and Cloud hosting sites
    - UVic, SFU, Montreal, UofT, UWaterloo
    - Node-to-node replacements are expected
    - Procurement process is about to start, will take a year or so
    - Expected to deliver before the end of 2025 fiscal year.
  - New AI institutes (UofA, UofT and Laval) will have new GPU AI/ML systems
    - Procurement already about to start
- Enhancement of user support
  - Couple of HPC Analyst and RDM professionals will be added in 2024-25

# Updates for DRAC ecosystem : Security

- Security
  - DRAC is working on increasing the organizational policies
  - Extra Security Analysts staff to monitor vulnerabilities etc.
  - Improving security by adding SSH key auth to CCDB
  - Implementing Duo MFA for SSH for general users.
    - Everyone is invited, about 10% of users so far enrolled
    - Will become mandatory (eventually)
    - Accommodating automated workflows is tricky


[https://docs.google.com/presentation/d/1nm\\_EV1eJLBR-uNlbt9bJnWMhYQazVCgC23e6GmAEdTA/edit#slide=id.g285b4e901fa\\_0\\_21](https://docs.google.com/presentation/d/1nm_EV1eJLBR-uNlbt9bJnWMhYQazVCgC23e6GmAEdTA/edit#slide=id.g285b4e901fa_0_21)

- Technology requirements and compliance for Medical Research are more difficult

# What is multifactor authentication ?

Something you know


ex: password, passphrase for SSH key



```
ssh-rsa AAAA
QCp0IeTMpMS8
VT6UegrSs260
Wq+Ls8nhjCu1
4DA61h...0Zr
```

Something you have

ex: single use code (OTP), hardware key



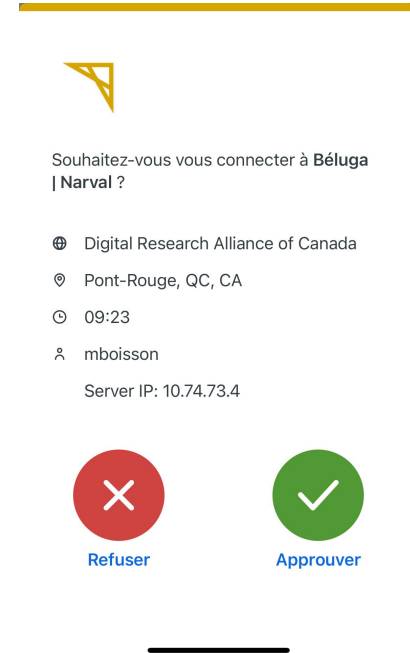
Something you are

ex: fingerprint, retinal scan, facial recognition



*Slides by: Maxime Boissonneault, RSNT Team Lead, AF*

# What are options of second factor ?



Slides by: Maxime Boissonneault, RSNT Team Lead, AF



# What are options of second factor ?



<https://www.yubico.com/products/yubikey-5-overview/>

*Slides by: Maxime Boissonneault, RSNT Team Lead, AF*



## 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 and helpdesk
  - 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 (OpenOnDemand)



# GreX status updates

- Security and compliance
  - Updated monitoring systems
  - Added local backup for mgmt systems, tape backup for /project
- Increased memory on 45 new compute nodes (96 GB to 192 GB)
- More contributed systems (840 core Compute, 1PB Storage added by Agriculture)
- SISF funding to replace 3000 old cores with new, in 2024.
- <https://ccdb.alliancecan.ca> is used on GreX as well
  - The system allows for user/password
  - Possible to deploy users' SSH public keys
  - The system allows registering devices for Duo MFA

# Access to National and Grex resources

- Access to Canadian user Database
  - Resources are allocated by PI
  - Eligible PIs can “sponsor” any group members or collaborators
  - All National / DRAC systems are using CCDB via LDAP
- <https://ccdb.alliancecan.ca>
  - The system allows for user/password
  - Possible to deploy users' SSH public keys
  - The system allows registering devices for Duo MFA

# Resource Allocation Calls (RAC, RRG, RPP)

- “Default allocation” of HPC CPU core-years and storage quota
  - Anyone who has a CCDB account
  - Part of the systems left “unallocated” gets divided between “default” users
- RAS, Rapid access service
  - Extra storage and Cloud resources on request throughout the year
- National RAC process, RRG (HPC) and RPP (Cloud, Portals projects)
  - Annual application through <https://ccdb.alliancecan.ca>
  - Evaluated on merit, scored by Scientific Review panels
  - Scaled (down) to fit into existing resources based on the score
  - RAC 2023 closes **November 7, 2023**.
- Local (GreX) RAC process
  - Annual , by email to [ARC@umanitoba.ca](mailto:ARC@umanitoba.ca)
  - Local RAC 2023 closes **November 21, 2023**

# Using HPC and cloud computing

- HPC is almost totally running Linux
  - User level skills of Command line , Scripts
- Remote access (SSH, etc.) or Portals/ Gateways
  - OpenOnDemand, Jupyter
- HPC software stacks
- HPC scheduler (SLURM)
- Cloud computing comes in many kinds
  - IaaS, PaaS, SaaS; can run many operating systems but Linux is a majority
  - Skill Sets differ widely (sysadmin, DevOp, user of a particular Software)
  - We will talk about IaaS here (Linux command line)



**University  
of Manitoba**