Autumn 2025 HPC and Cloud Workshop

October 14-17, 2025

Advanced research computing team/

Alliance Federation support site at UManitoba

Digital Research Alliance of Canada

University of Manitoba

Traditional Territories Acknowledgement

The University of Manitoba campuses are located on original lands of Anishinaabeg, Ininiwak, Anisininewuk, Dakota Oyate and Dene, and on the National Homeland of the Red River Métis.

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.

Digital Research Alliance of Canada

University of Manitoba

- Research Computing Resources
  - HPC, cloud, and (or) AI / ML
  - National Digital Research Infrastructure
  - Local HPC system
- Introduction to using
  - 1. Linux shell; Data transfers;
  - 2. HPC software (Lmod and Containers)
  - 3. Batch HPC with SLURM
  - 4. OpenStack Cloud and Object Storage

Program of the Workshop

Oct 14	Session title	Presenter	Start time	End time
1	Updates and overview of the Digital Alliance of Canada and UManitoba computing resources	Grigory Shamov	12:30	13:00
2	Housekeeping: how to connect to training resources for the Workshop		13:00	13:10
3	Basics of Linux Shell (hands-on)	Stefano Ansaloni	13:10	14:15
		Break		
4	Data transfer with SSH and with Globus (hands-on)	Stefano Ansaloni	14:30	15:30
5	OpenOnDemand HPC Web portal: File Transfer , Remote Desktop and interactive GUI applications	Stefano Ansaloni	15:30	16:30

Program of the Workshop

Oct. 15	Session title	Presenter	Start time	End time
1	Housekeeping: how to connect to training resources for the Workshop		12:30	12:40
2	Intro to HPC software, Lmod modules tool (hands-on)	Dr. Ali Kerrache	12:40	13:30
3	Intro to HPC software, Lmod modules tool (hands-on)	Dr. Ali Kerrache	13:30	14:15
		break		
4	Using Jupyter notebooks on HPC Systems: OpenOnDemand and Jupyter Hub	Grigory Shamov	14:30	15:00
5	Containers in HPC: using Singularity/Apptainer and Podman containers	Stefano Ansaloni	15:00	16:30

Program of the Workshop

Oct. 16	Session title	Presenter	Start time	End time
1	Housekeeping: how to connect to training resources for the Workshop		12:30	12:40
2	Running HPC jobs with SLURM scheduler (hands-on)	Dr. Ali Kerrache	12:40	13:30
3	SLURM topics: optimizing your job throughput on HPC machines	Dr. Ali Kerrache	13:30	14:15
		break		
4	Running HPC Batch jobs with OpenOnDemand portal	Dr. Ali Kerrache	14:30	15:00
5	Running Python and Jupyter AI / Machine Learning workloads on HPC (hands-on)	Grigory Shamov	15:00	16:30

Program of the Workshop

Oct. 17	Session title	Presenter	Start time	End time
1	Housekeeping: how to connect to training resources for the Workshop		12:30	12:40
2	Cloud computing with OpenStack Cloud: IaaS VMs (hands-on)	Stefano Ansaloni	12:40	13:30
3	Cloud computing with OpenStack Cloud: provisioning a service on the cloud (hands-on)	Stefano Ansaloni	13:30	14:15
		break		
4	Using S3 Object Storage on OpenStack Cloud (hands-on)	Stefano Ansaloni	14:30	15:30
5	Closing remarks	Grigory Shamov	15:30	16:00

HPC and Cloud computing resources: an overview of Canadian DRI ecosystem

Grigory Shamov

Oct 14-17, 2025

Digital Research Alliance of Canada

University of Manitoba

DRI (Digital Research Infrastructure)

“Advanced Research Computing” (ARC), HPC, Cloud etc.

- Focuses on enabling computational research / users
- Provides capabilities that are not available with common (desktop, enterprise server) computing environment:
    - CPU time and memory, fast interconnect
    - Accelerators (GPU, TPU hardware)
    - Fast and scalable storage capacity
    - Network/data transfer resources optimized for research
- A specialised set of hardware and software
- Managed differently from Enterprise IT, more research-specific
- ARC systems tend to be large and thus expensive (but efficient)

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

HPC as a technology, architecture

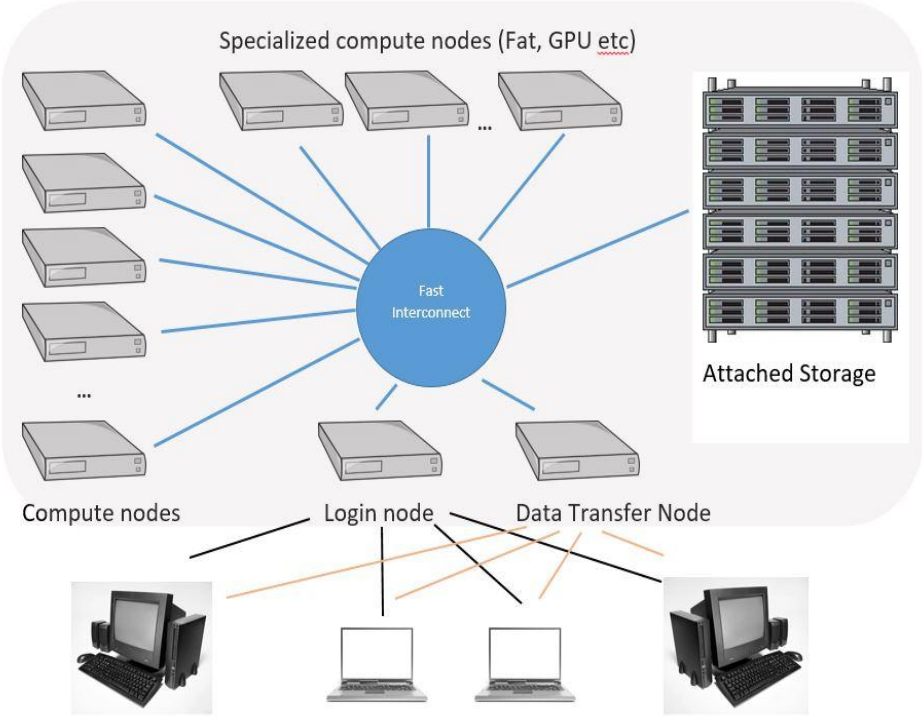

HPC vs Cloud computing

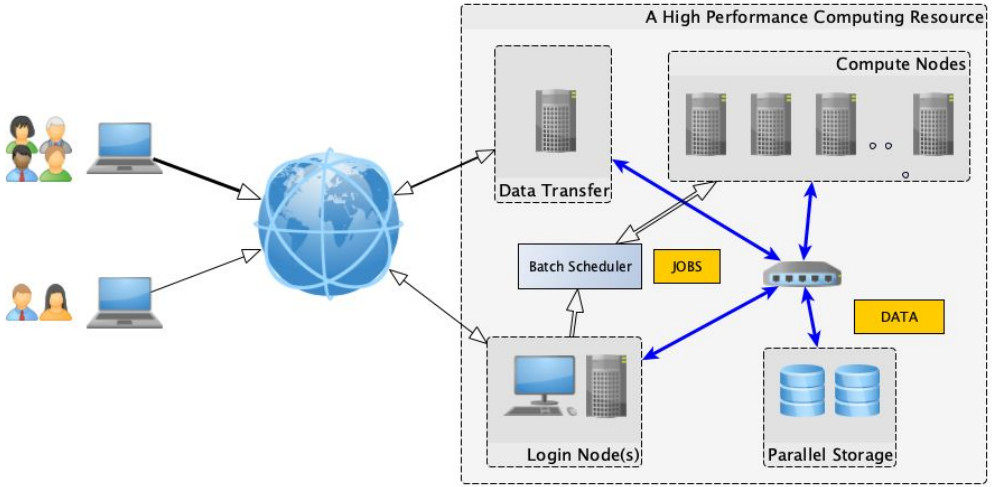

- HPC clusters are *shared* systems with r*emote access*
- Batch mode of usage
- A central Software delivery on HPC
- Dealing with Data (storage, transfer etc.)

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

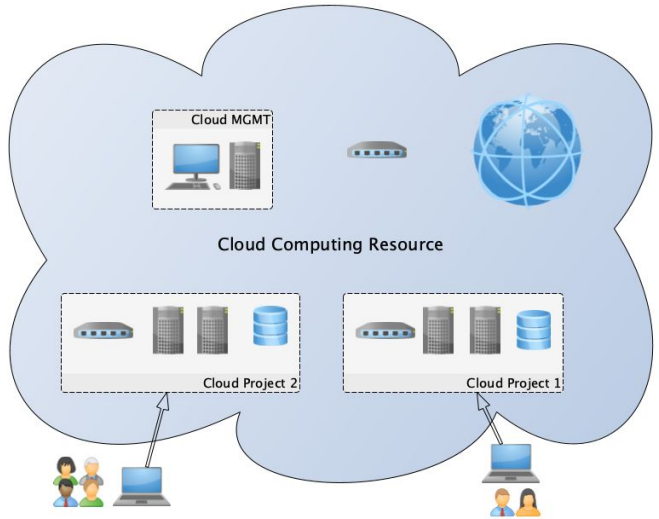

DRI (Digital Research Infrastructure)

- **HPC** is about efficiently organizing shared, bare-metal resources for top performance:
    - CPU, accelerators, scalable storage,
    - High-speed Interconnect.
    - High-density datacentres
    - Exascale HPC and High-Throughput computing / streaming

- **Cloud** computing originally focused on Enterprise computing
    - Business flexibility : virtualization, software-defined services
    - Standard datacentre,
    - Running persistent services.
    - Container orchestration platforms (Kubernetes)


Artificial Intelligence and Machine Learning revolution
https://static.googleusercontent.com/media/research.google.com/en//pubs/archive/35179.pdf "The Unreasonable Efficiency of Data"
http://www.incompleteideas.net/IncIdeas/BitterLesson.html "The Bitter Lesson"
https://dl.acm.org/doi/10.5555/3295222.3295349 "Attention is all you need"
https://blogs.nvidia.com/blog/what-are-foundation-models/ NVidia on foundational models.

DRI (Digital Research Infrastructure) for AI

IN MANY WAYS, AI VINDICATES THE "HPC WAY"

- ***AI needs fast interconnects***. We had them, the cloud and the enterprise did not.
  - Microsoft deployed 40,000 KM of *Infiniband*, in 2023, built for the HPC market ~1999,
- ***AI needs message passing***. MPI, the message passing interface, was built Open Source in the HPC community, ~1993
  - Now the standard library for transformer-based generative AI (e.g. ChatGPT, DeepSpeed, OpenAI etc.).
- ***AI needs heterogeneity*** – GPUs for general purpose computing – the hardware building block for AI - came out of the HPC world ("GPGPU" ~2004).
- ***AI needs fast, large scale filesystems –*** not object stores
- ***AI needs liquid cooling*** – even 5 years ago, many datacenter providers were convinced they could just use air, now none are. HPC systems switched to liquid cooling a long time ago.
- This means AI needs HPC hardware (probably good) and HPC programmers (good if you are one, bad if you need to hire one).

Slide by Dan Stazione, Director of TACC, US

***tamIA**, a real AI supercomputer of LavalU*

DRI requires a large capital investment

DRI requires an operation to run it:

- staff, expertise, real estate (datacentre), electricity

In Canada, there are two levels of DRI above “under your desk”:

- Institutional or “local” system.
  - Funded by Universities and / or projects (Faculties, Departments, PI grants)
- National DRI Organization / Consortium / Federation.
  - Funded by the Government of Canada with Provincial and/or Institutional matching

Canada's Advanced Research Computing Platform

DRAC. https://alliancecan.ca

CCDBhttps://ccdb.alliancecan.ca

Digital Research Alliance of Canada

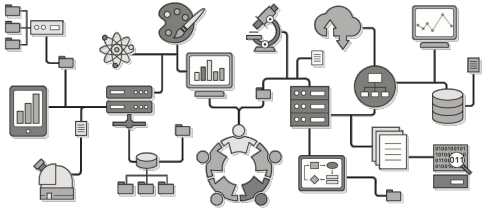

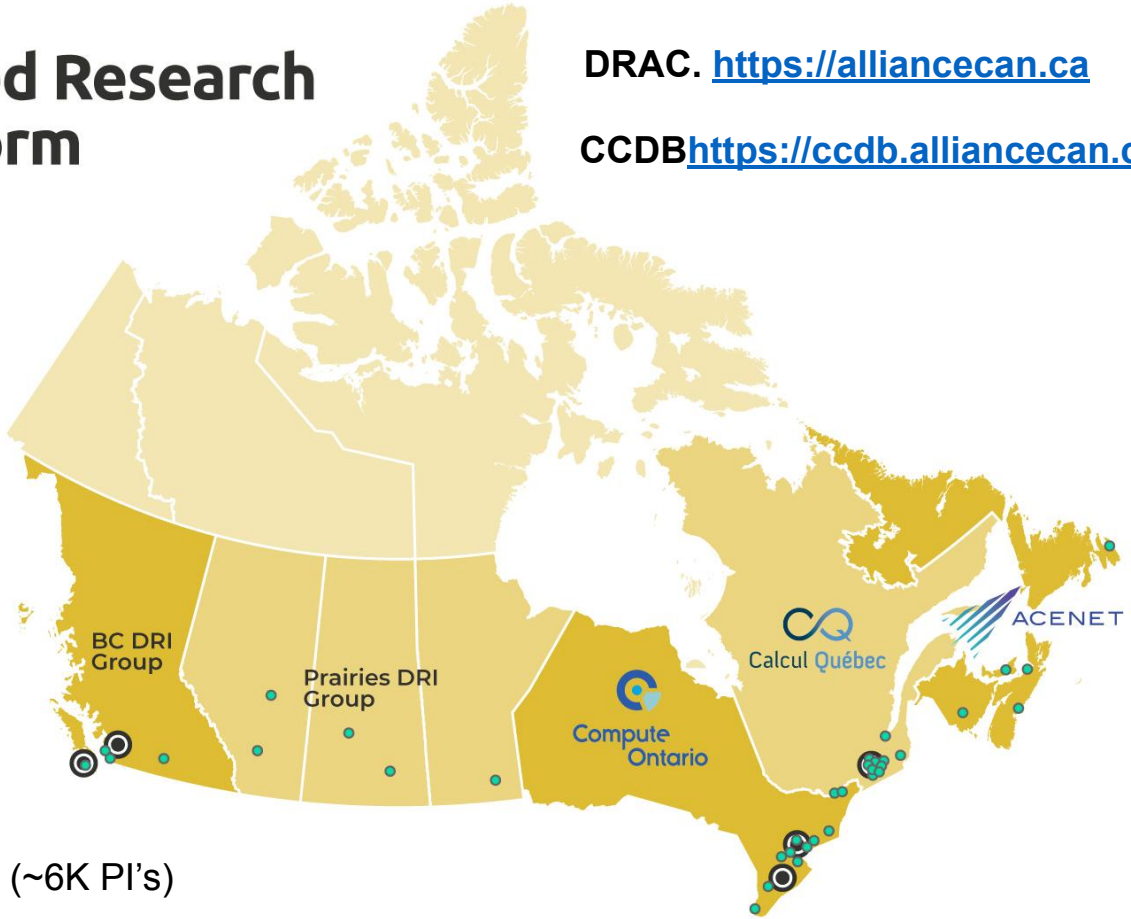

National Host Sites

Support Sites

Number of "roles" in CCDB : **~24K** (~6K PI's)

Number of "roles" in Prairies: **2873** (843 PIs), of them in Manitoba **490** (**153** PIs)

Digital Research Alliance of Canada

Digital Research Alliance of Canada

- Replaces an earlier National DRI organization, ComputeCanada (2015-2022)
  - Mandate to integrate ARC, RDM, and RSE

- https://alliancecan.ca

- Infrastructure renewal (**~220M**) for HPC and Cloud in 2022-2025.
  - Completed in September 2025!
  - Total: **5 175** nodes, **776, 912 CPU** cores, **2 152** GPUs (A100, H100)

- DRAC Works with Alliance Distributed workforce and Hosting Sites to support the National cloud and HPC systems
  - Coordinating, Funding and some operational roles (security).
  - Alliance consultations on the Sovereign AI Government's call

Size of National HPC systems 2025

Digital Research Alliance of Canada

Infrastructure renewal (**~220M**) for HPC and Cloud in 2022-2025.

- Fir (SFU)
- Nibi (UWaterloo)
- Rorqual (ETS Montreal)
  - Narval , same DC
- Trillium (UToronto)
- Arbutus Cloud (UVictoria)
  - Rolling update

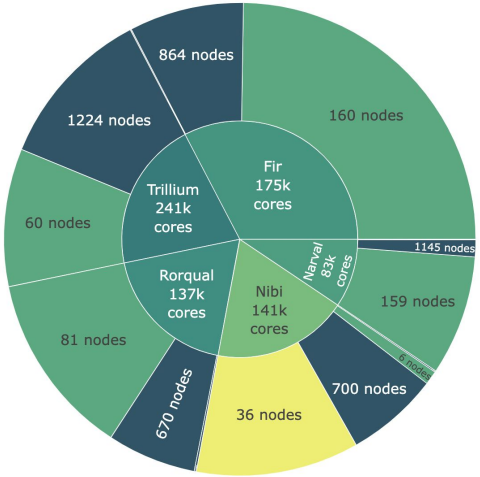

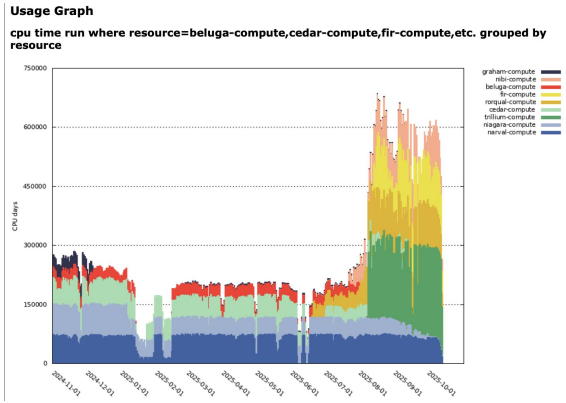

https://wgpages.netlify.app/files/clusters.html , *an infographics by Alex Razoumov @ SFU*

Example of a new HPC system Fir

Digital Research Alliance of Canada

https://docs.alliancecan.ca/wiki/Fir

- **175,104** CPU cores, **640** GPUs – **#78** in July's Top500 HPC for GPU-only
- 864 base CPU nodes:
  - 192 cores, 750 GB memory, 2x AMD EPYC 9655 Turin CPUs
- 8 large-memory CPU nodes:
  - 192 cores, 6000 GB memory, 2x AMD EPYC 9654 Turin CPUs
- 160 GPU nodes:
  - 48 cores, 1125 GB memory, 1 x AMD EPYC 9454 Genoa CPU
  - 4x NVidia H100 GPUs 80GB VRAM, SXM5
- **51 PB** high-performance Storage (DDN Lustre FS, all-SSD)
  - */home* , */scratch* , */project* partitions
- InfiniBand NDR interconnect **400Gb/s**

Grex and HPCC SISF Updates 2024

A step-up machine for UM users to DRAC resources

Grex used to be a Westgrid / National HPC machine

- 2011 - 2018, had 3840 CPU cores,QDR Infiniband.
- HPCC datacentre.

kept and maintained by the University of Manitoba as a local system

- Authentication and support through CC/DRAC systems
- Managed by the same local DRAC Federation team.
- Provides both a local and the ComputeCanada software stacks
- *Hosts CFI and other user-contributed hardware*

Renewed with help of the IST and VPRI funding and SISF 2024

- Serves ~140 research groups, ~450 users

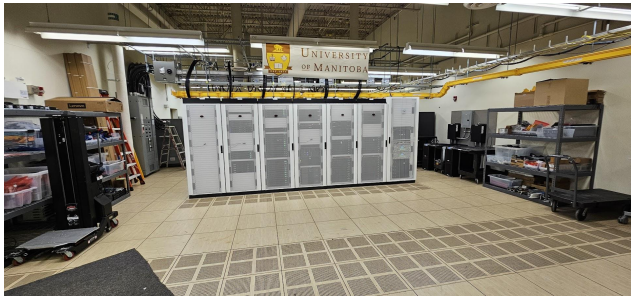

Renewal of Grex, local HPC system

System, kind (2011-2024)	CPU cores	GPU devices
Legacy SSE4.2 Intel	3840 to **~2000**	-
AVX512 Intel CascadeLake	2820	36 V100
AVX2 AMD	112	10 A30
	~ 5K	**46**

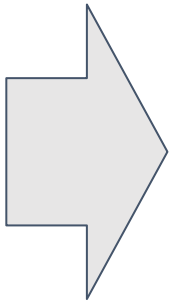

System, kind (2024-2025)	CPU cores	GPU devices
AVX512 Intel CascadeLake	2820	36 V100
AVX512 AMD Genoa	6964	2 L40s
AVX2 AMD	112	10 A30
	~10K	**48**

Using ARC resources

Access, Costs of Alliance Resources?

- Free for Eligible PI's : https://ccdb.alliancecan.ca
- Access through CCDB (ComputeCanada user DB)
  - The unit of resource allocation is "Research Group".
    - Not a university, not individual users
  - Thus, a Principal Investigator has to apply first, then "sponsor" anyone:
    - Undergraduate students research projects
    - External collaborators
    - Grad students, postdocs, support staff
- Eligible use: for academic research
- A user can belong to more than one research group.

Access, Costs of Alliance Resources?

- Free for Eligible PI's : https://ccdb.alliancecan.ca

- Default allocation (CPU, storage)
  - Immediately available for active accounts
- RAS (Rapid Access Service), on request any time of the year.
  - Mainly for Storage and Cloud resources.
- RAC (Resource allocation call) for increased requests, annually
  - Application process through CCDB
  - Proposal needed for CPU and GPU years, Storage TBs, cloud resources
  - AI systems follow their own Tiered allocation process
- **RAC 2026 is on, Deadline for PI's November 3, 2025 !**

HPC as a technology, architecture

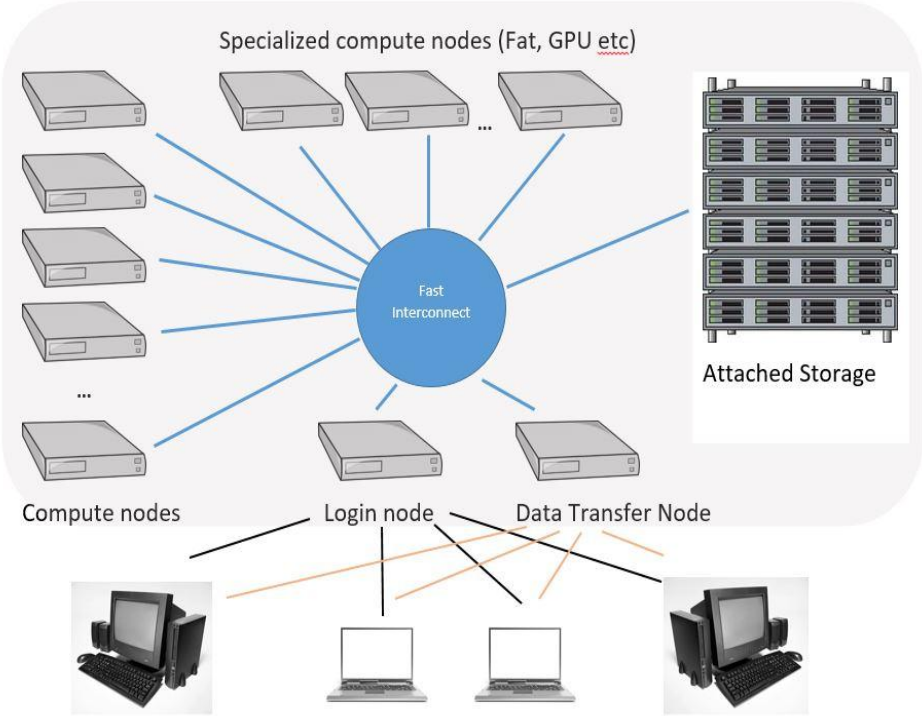

HPC workflow (from a user PoV)

- Working with a remote system
  - Remote Linux Shell
  - Data transfer to and between local and remote Systems
- Batch/text mode , with some provision for interactive and GUI
  - SLURM scheduler is most popular (PBS, LFS, etc.)
  - Make a “job script” → “Submit” → “”wait for completion” → “analyze results”
- Optimized software stacks for best utilization of baremetal hardware
  - Find a HPC software using Modules
  - Run Containers (that are HPC friendly : Podman and Singularity).

- Interactive GUI / Jupyter as jobs, OOD and JH portals
- Persistent/Server workloads can be served by OpenStack Cloud (UVic)
  - Infrastructure as a Service (Iaas); Object Storage


University
of Manitoba

Digital Research Alliance of Canada

Digital Research Alliance of Canada

- **Artificial Intelligence?**

- Pan-Canadian AI Compute Environment (**PAICE**) AI / ML Sites
  - AMII (UAlberta), Vector (UToronto) and Mila (LavalU)
  - ~ 40M initial investment, 3 new systems ready for production
  - Shared AF services: CCDB, software stack, Helpdesk

- DRAC Mandate renewal: Awarded further ~80M for 2026-2030
  - Infrastructure expansion for HPC and AI sites

- Participates in ISED Sovereign AI Strategy Call in 2025 for ~700M.
  - Submitted the Alliance's "Statement of Intent"
  - Consultations, coordinations of other SoI's from Academia and Industry

Size of new PAICE systems

Digital Research Alliance of Canada

System, kind (2025)	# GPU nodes	GPUs per node layout	Interconnect	Storage, PB
TamIA , HPC (Laval)	42 (H100)	**4 x** NVIDIA HGX H100 SXM	**4 x** HDR200 Infiniband, **non-blocking**	?
Vulcan, HPC (UofA)	205 (L40s)	4 x NVIDIA L40s	1x100Gbps Ethernet	5PB
Killarney, HPC (UofT)	168 (L40s) 10(H100)	4 x NVIDIA L40s, 8 x NVIDIA H100 SXM	1x HDR100, 2x HDR200	1.5 PB
HPC systems ?				
Fir, HPC	160	4 x NVidia H100 SXM	1x HDR200 Infiniband, blocking	51PB
Nibi, HPC	36	8 x Nvidia H100 SXM	1x Nokia 200/400G Ethernet	25PB
Trillium, HPC	60	4 x NVidia H100 SXM	1x NDR200/ NDR400 Infiniband	29PB

Closing Remarks!

Grigory Shamov

October 14-17, 2025

Digital Research Alliance of Canada

University of Manitoba

Canada's Advanced Research Computing Platform

DRAC. https://alliancecan.ca

CCDBhttps://ccdb.alliancecan.ca

Digital Research Alliance of Canada

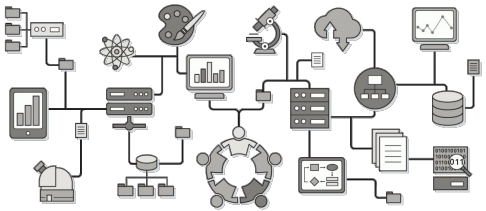

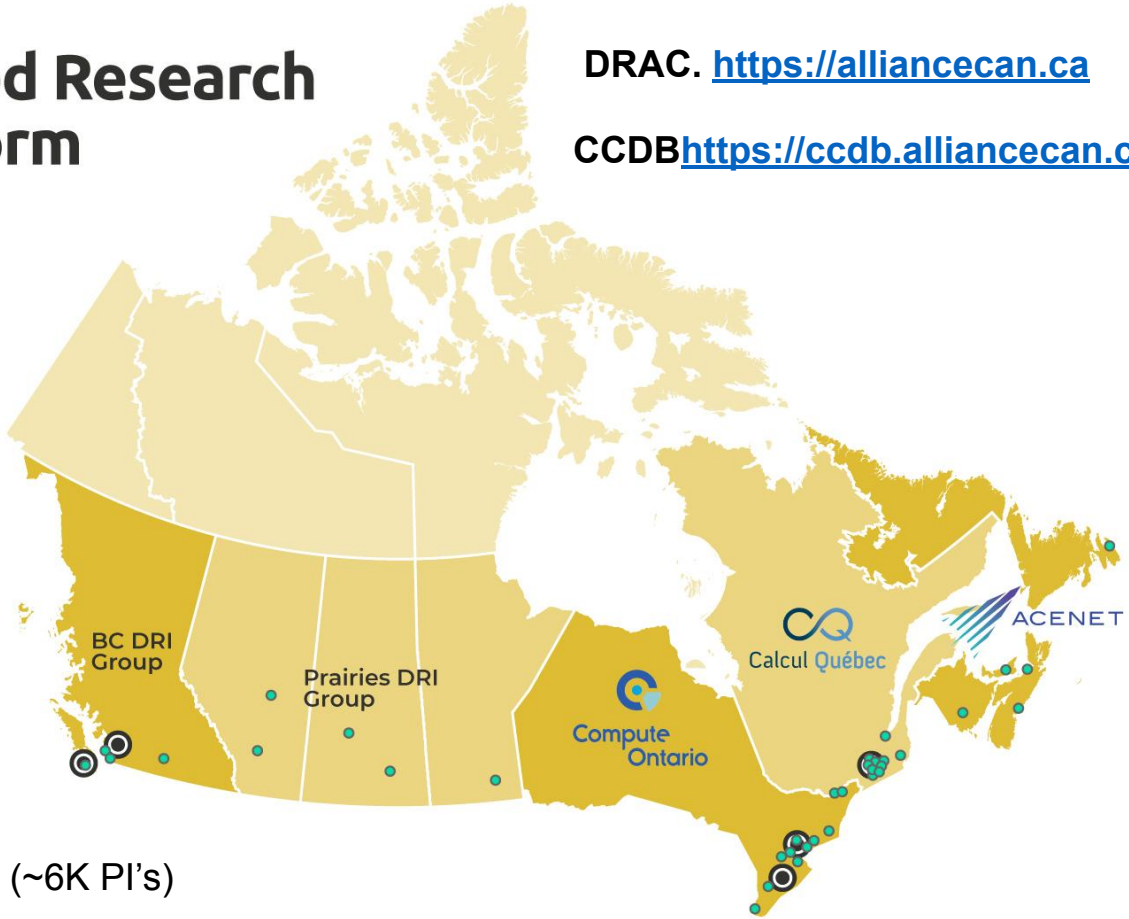

Number of "roles" in CCDB : **~24K** (~6K PI's)

Number of "roles" in Prairies: **2873** (843 PIs), of them in Manitoba **490** (**153** PIs)

Closing Remarks

Digital Research Alliance of Canada

- **Documentation**:
- User Wiki at the Alliance: https://docs.alliancecan.ca/wiki/Technical_documentation
  - https://docs.alliancecan.ca/wiki/National_systems
  - https://docs.alliancecan.ca/wiki/Running_jobs
  - Alliance systems' status page Status | Alliance

- UM Grex system, technical documentation site
  - https://um-grex.github.io/grex-docs/grex/
  - User-extendable: github.com/um-grex/grex-docs
  - Grex status page https://grex-status.netlify.app

- National Helpdesk contact
  - support@tech.alliancecan.ca : general support for HPC
  - cloud@tech.alliancecan.ca : OpenStack cloud questions

Closing Remarks

Digital Research Alliance of Canada

- **Support contacts**:

- National Helpdesk contact
  - support@tech.alliancecan.ca : general support for HPC
  - cloud@tech.alliancecan.ca : OpenStack cloud questions
  - globus@tech.allancecan.ca : Globus questions

- UManitoba IST support
  - support@umanitoba.ca

- IST Service catalogue for research computing
  - https://umanitoba.ca/information-services-technology/research-computing
  - Advanced Research Computing there

Closing Remarks

Digital Research Alliance of Canada

- **More training!**

- Western training courses (former WestGrid)
  - SFU: https://training.westdri.ca (/contact to subscribe)
  - UAlberta Research Computing Bootcamps | Information Services and Technology (IST)
  - Other regions (AceNet, SharcNet, SciNet, CalculQuebec)

- National Training Discovery portal at the Alliance
  - https://explora.alliancecan.ca