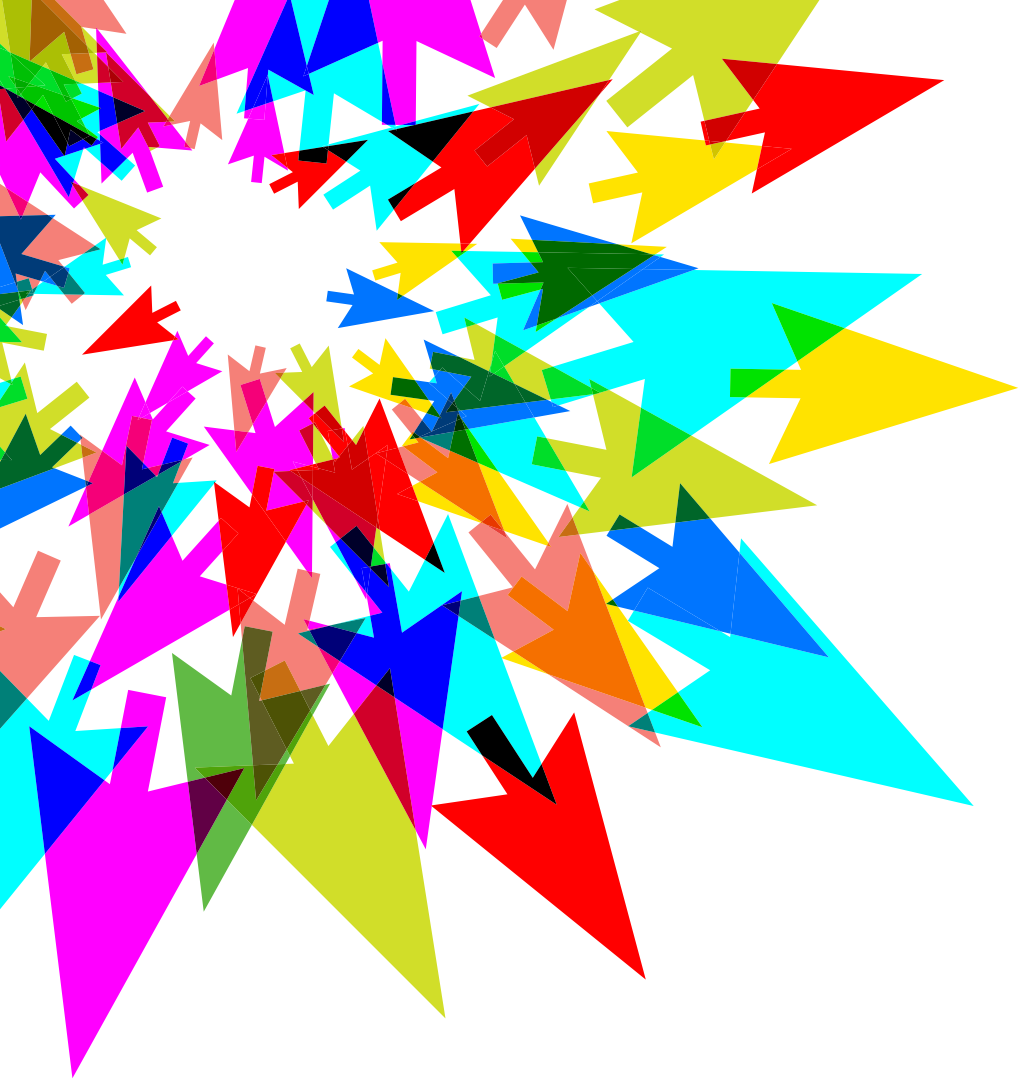

compute
canada

Compute Canada Resource Allocation Competition 2016

Call for Proposals | September 2015

Contents

36 York Mills Road, Suite 505, Toronto, Ontario, Canada M2P 2E9
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada

Context

Compute Canada, in partnership with regional organizations ACENET, Calcul Québec, Compute Ontario and WestGrid, provides essential advanced research computing (ARC) services and infrastructure for Canadian researchers and their collaborators in all academic and industrial sectors.

Each fall, Compute Canada hosts a national Resource Allocation Competition (RAC) for researchers who require greater than the default usage level on any of its systems. The competition is open to projects from all disciplines — from humanities to engineering — and allocations are awarded based on scientific merit, quality of the research team, and development of highly-qualified personnel (HQP). This competition is peer-reviewed and **the allocations are available for one (1) year** starting in early January 2016. Although all users have access to our team of Compute Canada technical experts, we do not allocate dedicated support (ie. programmers, developers) as a part of this resource competition.

Compute Canada has seen a significant growth in demand since its first competitive national call in 2010. Over the last five years, the number of applications has increased by 159%, from 135 applications in 2011 to 350 applications in 2015.

In 2016, between the Resource Allocation Competition and Research Platforms and Portals (RPP) competitions, we expect to be able to allocate 124,000 core-years in computing power and 15.5 PB of storage.

To learn more about the RAC and the resources available at Compute Canada, please refer to the Frequently Asked Questions page on our website.

For other opportunities for resource allocations, please see Appendix A and consult the Compute Canada website.

Eligibility

Applicants

Researchers at Canadian academic institutions who are eligible to apply to national granting councils for funding can apply to the RAC for access to ARC resources on Compute Canada systems. Typically, this means that regular faculty members are eligible to apply while postdoctoral fellows and graduate students are not. Individuals who have adjunct status may apply for an allocation if the project for which the allocation is requested is eligible for funding from the Canada Foundation for Innovation (CFI) or other national granting council. A lead Principal Investigator (PI) cannot submit more than one individual application but can be a participant in other submissions.

Request Size

Users who are applying for a resource allocation should verify that their application exceeds the default allocation available through the various Compute Canada sites. A complete list of sites and default allocations can be found here. Any application for less than the default level at a particular site will be directed to access a default allocation, which can be done by simply registering with the Compute Canada DataBase (CCDB). Please note: default allocations are not a guaranteed level of resource availability, but rather our best estimate based on usage in the previous year.

36 York Mills Road, Suite 505, Toronto, Ontario, Canada M2P 2E9
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada

Guidelines for Completing a Resource Allocation Competition Application

Consultation

It is strongly recommended that all applicants consult with the Compute Canada technical staff responsible for the systems they intend to use, well before completing their proposal. It is within the best interest of the applicants and ensures the technical aspects of the proposals are appropriately matched with the requested systems.

Applicants who have not previously used Compute Canada systems MUST discuss their proposal with Compute Canada technical staff before submission. Please refer to the Questions & Additional Information section in this guide for contact details.

RAC consultations must happen prior to September 30, 2015, to allow adequate time for support by Compute Canada technical staff.

Compute Canada Database (CCDB) Application

All applications to the RAC are submitted using the CCDB. Users must register for an account or log in using an existing account to create a new application.

Under the heading "RAC Application" select "Resource Application". Navigate down the page to the title "Current resource applications" and select "Apply for an allocation". This will open the Resource Application page where applicants can complete a RAC application. **Note that no other formats of applications are acceptable.**

Applicants must complete all sections of the application and include the following information:

- Title of the project
- Research summary
- Research area
- Research image & caption (optional)
- Co-Principal Investigators (all Co-PIs must have created a CCDB login to be added to an application)
- Principal Investigators will be asked to upload their Canadian Common CV
- Research and Technical Justification (download and use the appropriate template – see Appendix B)
- Details on any previous allocation requests and usage
- Resource requests (specification of the resources needed across the various systems available)

Note that scaling tests are required for any parallel job using more than 256 cores. Please review the form and the template well in advance to make sure you are able to provide all the information required. If you have any questions or need technical assistance, please send an email to rac@computecanada.ca.

36 York Mills Road, Suite 505, Toronto, Ontario, Canada M2P 2E9
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada

Evaluation Process

Applications submitted to the RAC are evaluated for both technical feasibility and scientific excellence. The technical review is conducted by Compute Canada staff who will:

- ensure the appropriate system is requested by the PI and the required software is available
- ensure the code is efficient (and scales well when applicable)
- identify groups that may need help with code and workflow optimization
- identify discrepancies between the online request and the complete description of the project
- provide a technical opinion on the reasonability of the request

During the Technical Review, staff may require additional information from applicants and will engage them directly. In order to ensure an application can progress beyond the Technical Review, applicants should respond to requests from the Compute Canada technical team within 48 hours of the request.

Following the Technical Review, the application will undergo a Scientific Review by a Compute Canada Expert Review Committee. The Expert Review Committees are composed of researchers within a respective field and are populated annually based on the number of submissions in a given area.

Each Expert Review Committee will evaluate proposals based on the quality of the science and the quality of the team. The research evaluation criteria are weighted as follows: 60% for the quality of the science and 40% for the quality of the team. The specific review considerations can be found below in the Evaluation Criteria section.

Applications will be scored by an Expert Review Committee based on the merit of the proposal and may be further discussed by a committee comprised of the chairs for each Expert Review Committee and representatives from Compute Canada. Final decisions on allocations will be based on science merit of the proposals and the availability of resources.

We understand that some applicants may have already undergone an evaluation through the various granting councils previously; however, the RAC evaluation ensures the most effective use of the computing resources to achieve scientific goals and ensures fairness, as not all applicants receive funding through only the tri-councils (eg. National Institutes of Health).

36 York Mills Road, Suite 505, Toronto, Ontario, Canada M2P 2E9
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada

Evaluation Criteria

The research evaluation criteria are divided into two categories: Quality of the Science and Quality of the Applicant(s). Considerations in each area are as follows:

Quality of the Science (60%)

- originality and innovation
- significance and expected contributions to research
- clarity and scope of objectives
- clarity and appropriateness of methodology
- feasibility
- discussion of relevant issues
- impact of the research
- HQP Training
    - number of HQP directly engaged in the project across academic levels (e.g. Undergraduate, Masters, PhD, PostDoc, etc.)
    - quality of HQP contributions
    - impact of participation on HQP
    - unique training opportunities for HQP
    - potential cross pollination between disciplines of HQP

Quality of the Applicant(s) (40%)

- knowledge, expertise, and experience
- quality of contributions to, and impact on, the proposed and other areas of research
- importance of contributions

Submission Procedures & Deadlines

Proposals must be submitted electronically through the CCDB no later than **October 8, 2015 at 4 p.m. (EDT)**. It is the responsibility of the applicant to ensure the application is complete with all additional documentation uploaded. Incomplete applications or applications submitted after the deadline will not be accepted.

Following the review process, applicants will be informed regarding the status of their applications via e-mail in December 2015.

36 York Mills Road, Suite 505, Toronto, Ontario, Canada M2P 2E9
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada

Questions & Additional Information

For any questions or for more information on the RAC, please contact us at rac@computecanada.ca.

For server-specific questions, please contact your regional support team:

ACENET:	support@acenet.ca
Quebec:	support@calculquebec.ca
Ontario:	
HPCVL:	help@hpcvl.org
SciNet:	support@scinet.utoronto.ca
SHARCNET:	help@sharcnet.ca
WestGrid:	rac@westgrid.ca

Privacy Policy

Confidentiality of Information

Compute Canada safeguards the information it receives from applicants. All reviewers are required to sign a Nondisclosure Agreement and Compute Canada's Conflict of Interest policy, and they are instructed to keep all proposal information confidential and to use it only for review purposes. All proposals are available for review by all Compute Canada reviewers and the RAC Administrative Committee.

Use of Personal Information

Any personal information collected by Compute Canada is used only to review applications. Such information may be shared with relevant officials in the relevant consortium and/or with their research institution.

Public Information

Compute Canada will post the following information on its website for allocations provided:

- *Applicant(s) Name(s)*
- *Department/Institution/Organization*
- *Project Title*
- *Summary*
- *Allocation*

36 York Mills Road, Suite 505, Toronto, Ontario, Canada M2P 2E9
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada

Appendix A - Other Resource Opportunities

Other Compute Canada Resource Opportunities

Default Allocations

Default allocations on Compute Canada compute and storage systems provide researchers with a gateway into the use of ARC for their research. Researchers looking to access these available resources are encouraged to register for a CCDB account. There is no formal application process for accessing default allocations. See the complete list of sites and default allocations.

Fast Track (By Invitation Only)

In certain situations, there are researchers and small groups who wish to maintain their existing allocation with no change to the scientific approach or minimal changes to the amount of resources needed. To reduce the need to re-evaluate these proposals, Compute Canada has created a “Fast Track” process.

This process will be available for existing users who fall within the following criteria:

- Have moderate-sized allocations.
- Request the same amount or less of the resources allocated (increase requests will not be considered).
- Received better than average science score from the RAC.
- Have used their allocation consistently over the previous year.
- Have never applied for resources via the Fast Track process *or* have received an allocation via Fast Track and used it for only one year.

If this applies to your group, you will receive a letter explaining how to proceed. For more information on the Fast Track process, please contact rac@computecanada.ca.

Please note that applicants who have received resource allocations via Fast Track for two consecutive years MUST apply for resources by submitting a full application to the RAC or RPP competitions.

Research Platforms and Portals Competition

The Research Platforms and Portals (RPP) competition focuses on providing ARC resources to a larger community of researchers driven towards creating a platform for collaboration or a portal for broader use of Compute Canada resources. Allocations may be granted for multiple years and are evaluated using specialized review criteria. Groups are encouraged to use this competition if their application falls within any of the following categories:

36 York Mills Road, Suite 505, Toronto, Ontario, Canada M2P 2E9
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada

- Resource requests on behalf of a large community of users that will be reallocated to individuals and small groups following the award.
- Applications that provide a public platform that will make use of Compute Canada computing or storage.
- Groups engaging in international agreements to provide multi-year computing or storage solutions based in Canada.
- Recipients of funding from the Major Science Initiative from the Canada Foundation for Innovation.
- Groups that are providing shared data sets accessible using a third party (non-Compute Canada) interface.

For more information, see the RPP Application Guide in the Research Portal section of the Compute Canada website.

36 York Mills Road, Suite 505, Toronto, Ontario, Canada M2P 2E9
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada

Appendix B

Compute Canada Resource Allocation Request, Template 2016

This proposal for Compute Canada resources will be reviewed for research quality and merit by a committee made up of faculty members with expertise in the subject area (or one closely related). It will also be reviewed by Compute Canada technical staff on the basis of technical justification and expertise in order to ensure that Compute Canada resources will be used appropriately and efficiently.

You must use the section headings below, leaving 2 cm (0.8 in) in all the margins and using Arial or Times New Roman fonts (12 points) . You may remove all explanations (text in italics) if desired.

Please consult the Glossary in Appendix C to complete your application.

1. Introduction to the Research Problem

Outline the research problem and its importance/relevance as well as your goals. Typically 0.5 - 1 page long.

2. Research Justification

In-depth discussion of the problem, your methodology, timelines and specific goals. Typically 1-2 pages long.

3. Training and Support of HQP

Describe how this allocation will support the training of Highly Qualified Personnel (HQP). This should include a description of the research group and estimate of the number of graduate students, undergraduate students, and postdoctoral fellows who will be supported by the allocation.

4. Technical Justification

This section addresses the technical details of your computational and/or storage needs. Compute Canada needs enough information to ensure that compute cycles and storage are used as efficiently as possible, that resource requirements are estimated reasonably, and that the appropriate systems are being used. If you are not requesting an allocation of compute time then don't complete section 4.1. If you are not requesting a storage allocation then don't complete section 4.2. Typically the entire section will be 10-15 pages long but projects involving several key codes and/or with complicated storage requirements may need more. If you do not know the meaning of the technical questions or need help benchmarking your code, do not hesitate to contact us at rac@computecanada.ca (or contact your local Compute Canada support person).

4.1 Compute Requests

4.1.1 System Selection

Provide reasoning for your choice of system(s). Mention other systems that might also meet your needs. If a "close" or local system was requested, provide justification of why that may be a better choice than a remote system. Applicants should use the resource identification tables on the Compute Canada website for computing and storage to indicate the resources requested.

36 York Mills Road, Suite 505, Toronto, Ontario, Canada M2P 2E9
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada

4.1.2 Code Details

If requesting compute cycles, provide details about the codes (e.g. name, key reference publication, essential numerical methods used, whether it is serial/parallel, the type of parallelism [if any], etc). Indicate whether the code is private (written by you or a collaborator), community, Open Source, or commercial, and whether there are any licensing requirements.

4.1.3 Code Performance & Utilization

If requesting compute cycles, discuss code performance (e.g. how many iterations/timesteps/Flops per hour of wallclock time and the type of system used for this measure, how much RAM is required per job/process, etc.) and whether there are particular system or processor architectures for which the code is best suited. Discuss numbers and sizes of files expected to be produced per job. How much of the resulting data needs to be kept on the system longer than 23 weeks (e.g. longer than for simple postprocessing)?

In the case of parallel codes, please discuss scaling efficiency and justify (in terms of performance) the typical job sizes you will run. A plot of scaling efficiency is ***required*** *if parallel jobs will use 256 cores or more and is preferred even for smaller jobs sizes (note that Compute Canada can make cycles available for this type of code performance testing). Also indicate how much (temporary) disk storage will be required for actually running your jobs (not for storing/archiving results).*

Is checkpoint restart implemented?

4.1.4 Memory Requirements

Describe memory needed and the level of confidence and experience with memory use. Describe whether this is per-core, per-node, per-job, etc., including situations in which memory requirements might change. Please consult the Glossary in Appendix C to complete this section.

4.1.5 Size of Request

Explain how you estimated the total amount of compute time required for this project. Discuss what other computing systems you have access to (list both non-Compute Canada systems and any systems operated for you by Compute Canada) and why the allocation requested is needed in addition to these resources. It is most effective if this section contains a list (or table) of projects and a justification for each. For example:

Project	***Team Members***	***Estimated Number of Core Years***	***Associated Storage (Temporary)***
Project 1	*Student X*	*100,000*	*100TB*
Project 2	*Students Y, Z*	*500,000*	*20TB*
Totals		***600,000***	***120TB***

36 York Mills Road, Suite 505, Toronto, Ontario, Canada M2P 2E9
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada

4.1.6 Impact of a Cut

Due to the oversubscription of our systems, our process is very competitive and we may not be able to allocate the full amount of resources requested. What would be the effect on the project's research goals if the compute cycles requested were to be cut? Please address two concrete scenarios: a) 25% and b) 50% cuts.

4.2 Storage Requests

4.2.1 Storage Details

Explain why a storage allocation is required as opposed to making use of scratch or other space made available by Compute Canada for running jobs. Is the allocation being requested in order to store codes and data files (this would be typical of many requests) or are there additional special requirements (e.g. for databases, web access, availability from multiple sites/systems etc)? Roughly how many individual files will be stored and what is their size distribution?

State whether the data being stored must be directly accessible by running jobs, or if it can be on remote network accessible server. Indicate if the data being stored is the only copy of the data that exists and state what would be required in order to regenerate the data if it was lost. Does this data need to be backed up by the site?

4.2.2 Storage Performance & Utilization

Will storage requirements vary during the year (e.g. will all requested storage be needed immediately in January or can the allocation grow/vary during the year)? Will the storage allocation be required to persist into the following year? Is storage performance (e.g. bandwidth and IOPS) critical to the project and, if so, what estimated I/O and IOPS rates are required and why?

Are you using data compression? If not, is compression possible?

Compute Canada provides many types of storage. Many long-term storage needs can be met by tape at much lower cost than disk. Tape is still accessible on a short-timescale and is suitable for data that is not being very frequently accessed. If you know that you cannot use tape storage for your longer-term needs, please explain.

4.2.3 Size of Request

Explain how you estimated the total amount of storage required for this project. Explain what other storage systems you have access to and why the allocation requested is needed in addition to these. It is most effective if this section contains a list (or table) of projects and a justification for each. For example:

Project	*Team Members*	*Associated Storage (long term)*
Project 1	*Student X*	*100TB*
Project 2	*Students Y, Z*	*20TB*
Totals		***120TB***

36 York Mills Road, Suite 505, Toronto, Ontario, Canada M2P 2E9
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada

4.2.4 Impact of a Cut

Due to the oversubscription of our systems, our process is very competitive and we may not be able to allocate the full amount of resources requested. What would be the effect on the project's research goals if the storage requested were to be cut? Please address two concrete scenarios: a) 25% and b) 50% cuts.

4.2.5 Data Privacy Considerations

Please indicate if your datasets are subject to special privacy considerations (e.g., personal health information). If such considerations are needed, please provide additional details as to the regulations surrounding your data.

5. Progress Over Past Year

If you had a RAC allocation last year, please summarize the progress made by your group as a result of your allocation. Please note the research results that were made possible as a result of your allocation (i.e. specific discoveries, publications, awards, etc). This section is very important for Compute Canada's reporting to funding agencies on the use of the infrastructure.

36 York Mills Road, Suite 505, Toronto, Ontario, Canada M2P 2E9
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada

Appendix C - Glossary

General

Job – A job is the basic execution object managed by the batch system. It is a collection of one or more related computing processes that is managed as a whole. Users provide a resource description on the job creation, by queuing it into the batch system. A job description includes a resource request, such as the amount of required memory, the duration of the job, and how many compute cores this job will require. Based on the resources utilized, jobs can be either serial (running on one compute core) or parallel (running on multiple compute cores).

Parallel job – A job that runs on multiple CPU cores. Parallel jobs can be roughly classified as threaded/SMP jobs running on a single compute node and sharing the same memory space and distributed memory jobs that can run across several compute nodes.

Serial job – A job that requires one compute CPU core to run.

Cluster – A group of interconnected compute nodes managed by a resource manager acting like a single system.

Compute node – A computational unit of the Cluster, one or more of which can be allocated to a job. A node has its own operating system image, one or more CPU cores and some memory (RAM). Nodes can be used by the jobs in either exclusive or shared manner depending on a system.

Core year – The equivalent of using 1 CPU core continuously for a full year. Using 12 cores for a month, or 365 cores for a single day are both equivalent to 1 core year.

Memory

Memory per core – The amount of memory (RAM) per CPU core. If a compute node has 2 CPUs, each having 6 cores and 24GB (Giga Bytes) of installed RAM, then this compute node will have 2GB of memory per core.

Memory per job – The total amount of memory that a user job will require to finish its computation. If the job is serial, or SMP-parallel, the memory per job cannot be greater than memory per node. If the job is a distributed memory parallel, the memory usage is determined by memory user per process. The compute nodes will be best utilized if the memory per process for the job matches the memory per core on the nodes the job runs.

Memory per node – The total amount of installed RAM in a compute node.

36 York Mills Road, Suite 505, Toronto, Ontario, Canada M2P 2E9
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada

Storage

Frequently Accessed Storage/Infrequently Accessed Storage – Scratch and home directory are considered frequently accessed storage. Off-site storage (including tape archives) are considered infrequently accessed storage.

Long-Term Storage – Storage for active research projects which must persist for more than a few months. This might be frequently or infrequently accessed. This is not true archival storage in that the datasets are still considered “active”.

High-Performance Storage – Accessible on compute nodes, the high-performance storage provides a scalable and high throughput file system where thousands of programs can read from or write to that storage.

Home Directory – Each user has a folder on the non-local storage called “home directory”. The home directory is persistent, smaller than scratch and, in most systems, backed up regularly.

Local Storage – This refers to the hard drive in a compute node that can be used to temporarily store programs, input files, or their results. The local storage is not persistent, so the files created on the local storage should be moved to non-local storage to avoid data loss.

Non-Local Storage – A permanent storage that is physically located outside the compute nodes but can be usually accessed from a program running on a compute node. The non-local storage usually includes home directory and scratch and may include other file systems at particular sites.

Off-Site Storage – A storage facility located outside of the site where a program is run. The off-site storage may be used for long-term or infrequently accessed storage.

Scratch – The scratch storage, accessible on compute nodes, can be either local or non-local and is used to run programs, to keep input files and to store temporary or permanent output files. A non-local scratch storage is usually larger than a home directory but is not backed up. Scratch storage is usually subject to periodic “cleaning” according to local system policies.

Site – A member of one of Compute Canada’s regional consortia providing advanced research computing (ARC) resources (such as high-performance computing clusters, Clouds, storage, and/or technical support).

Disk – A disk or hard drive is a permanent storage (compared to a computer’s main memory or RAM) that holds programs, input files, output results, etc.

Tape – Tape is a storage technology used to store long-term data that are infrequently accessed. It is considerably lower in cost than disk and is a viable option for many use cases.

36 York Mills Road, Suite 505, Toronto, Ontario, Canada M2P 2E9
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada