RAC 2019 Changes & Best Practices

V1.1 Sep 15, 2018

Note - this document is a WestGrid-provided best practices guide for the Compute Canada Resource Allocation Competition (RAC). It does not replace the official RAC documents and PI's should carefully review the Compute Canada official documentation. In cases of conflict between this Best Practices Guide and the official RAC documents, the official documents take precedence.

1 Introduction

The central Compute Canada web site has the official documentation for the Resource Allocation Competitions (RAC's).This Best Practices Document is intended to amplify and extend the official documentation.

Summary presentations are also available:

- University of Manitoba Presentation
- University of Victoria presentation
- (Others will be added when available)

It is worth emphasizing that the official documentation should be carefully read!

- Resource Allocation Competitions Program Guide
- RPP 2019 Application Guide
- RRG 2019 Application Guide
- List of Allocatable Systems

1.1 Proposal Writing/Review Support

The regions provide extensive support for RAC proposals.

WG has support staff at most sites who are very experienced and knowledgeable. We strongly recommend that PI's consult with support staff at your institution.

To set up a consultation contact your local support staff or send a request to support@westgrid.ca.

1.2 Lack of Resources for 2019

Statistics for RAC 2018 are presented in the appendix. The key takeaway is that we have insufficient resources, and expect similar numbers for 2019.

Current 2019 predictions suggest that Compute Canada will be able to satisfy at most 50% of CPU asks and 30% of GPU asks. RAC 2019 will be very competitive!

1.3 Need for Competitions

The basic driver is insufficient resources needed to satisfy all the requests. A related driver is CFI's general focus on "excellence". Therefore CFI provided resources must be allocated through a national, **competitive** process emphasizing high quality scientific research.

CC's RAC process is primarily concerned with the scientific review which provides a quantitative "science score" for each proposal. This science score is then used to scale the computational asks.

2 Competitions: RRG and RPP

There are two competitions (CC Resource Allocation Competitions).

Resources for Research Groups (RRG)	• Annual requests for primarily job-based resources and storage on the national compute clusters. • Cloud resources can also be requested. • **Fast-track:** previous RRG holders will be invited to re-use their previous proposal (some limitations apply).

Research Platforms and Portals (RPP)	● Primarily aimed at scientific gateways (platforms and portals) which are based in the CC cloud. ● Generally RPP asks are for cloud virtual machines and storage ● Computational resources in the cloud or from the clusters can be included as components of the platform. ○ In some cases platforms have been hosted externally ● Longer-term: up to 3 years.

2.1 RAC 2019 Schedule

2.2 Schedule

Competition	Opens	Closes
Fast track (RRG)	Sep. 27, 2018	Oct. 25, 2018 11:59 PM EST
Fast-track Invitations sent	Sep. 12, 2018	
RPP	Sep. 27, 2018	Nov. 8, 2018 11:59 PM EST
RRG	Sep. 27, 2018	Nov. 8, 2018 11:59 PM EST

2.3 Q&A Events

Events	Date
Q&A (English)	Oct. 4, 12:00-13:30 EST
Q&A (French)	Oct. 5, 12:00-13:30 EST

2.4 Process

Process	Date
Science and Technical Reviews	Dec 2018 - Jan 2019
RAC face-to-face meetings	Feb, 2019
Award Letters	Mar, 2019

3 Fast Track

WestGrid strongly recommends that PI's consider fast-track if you receive an invitation. If you are reasonably happy with your 2018 allocation then fast-track is very straightforward and requires no further work on your part.

Eligible PI's will receive an invitation to submit a fast-track on Sept. 12. The submission procedure is very straightforward and the invitation includes instructions.

If you require more than 1,999 CY then you will have to submit a new proposal. Also note that you cannot increase your request.

You are not required to use the fast-track program. For instance you may not have been happy with your previous RAC allocation, and would like an increase. In this case you should submit a new proposal. Keep in mind the following:

- There is no guarantee that a revised or updated application will meet with success.
- You should of course review your previous proposal's comments and respond to them.
- It would be worthwhile asking a WG technical reviewer or a co-worker to critically review your old proposal.

Also note that accepting a fast-track does not mean that you receive exactly the same allocation as previously. Instead your score from the previous year is included in the process. In practice we do not expect big changes but this is still a competitive process.

4 Migration and Defunding

Most of the old legacy systems will be defunded by Mar.31/2019. The only remaining WestGrid system is Orcinus at UBC. Orcinus will be defunded on Mar.31, 2018 and will not be available for RAC 2019.

If you elect to fast-track and your allocation is on a to-be-defunded legacy system your allocation will be moved to a new system.

Migration details are on the Compute Canada Docs wiki.

5 Rapid Access Service and Early Stage Researchers

Early-stage researchers can be at a disadvantage in the competitions as they have not yet developed their own research reputations. To ameliorate this CC reserves approximately 20% of resources for non-competitive purposes, including the “Rapid Access Service” (RAS) which is aimed at new or early-stage researchers and smaller projects.

All users get a CC account, and this immediately gives access to the RAS on the national general purpose systems (currently cedar and graham), and on request to the CC cloud. RAS users share the 20% evenly with all other RAS users, so this is opportunistic. If there are a lot of RAS users then an individual user’s priority will be correspondingly lower.

The expected scenario is

1. Early-stage or new users get an account and learn about the systems.
2. Run prototypes or test jobs, and if possible production jobs
3. Acquire performance statistics
4. Predict future requirements
5. Use the above information to create a well-justified RAC application.

Additionally the scientific review committees identify particularly promising early-stage projects and will recommend special consideration to the main RAC team. This is a qualitative process so early-stage researchers should make an effort to provide a very clear idea of their projects and potential advantages. The ask should be modest, and aimed more at ramp-up rather than an unjustified huge ask.

5.1 RAS Limits

RAS has the following maximum limits. Requests for anything over these limits should submit a full RAC proposal.

Service	**Max RAS**
Compute	50 Core Years
Storage	10 TB
GPU	10 GPU years
Cloud Compute	80 vCPUs

Cloud Persistent	10 vCPUs
Cloud (block) storage	1 TB

RAS compute is **not an allocation**. Priorities are evenly distributed across the 20% default, so use is opportunistic. Usually there are a limited number of such RAS users and most of them in practice can get more than the above limits.

6 Out-of-Round Applications and New Users

New faculty may join an institution during the year. These prospective users can of course use the RAS, but can also submit an “out-of-round” RRG proposal.

Out-of-round proposals are also entertained for sudden successes where additional resources are immediately needed to take advantage of a breakthrough.

Out-of-rounds are full RRG proposals, and follow the same review process. The review can take a few weeks as proposals are sent to external reviewers. As usual support@westgrid.ca can help with the details.

Note that there is no RPP out-of-round process. New users with scientific gateway projects may submit an RRG out-of-round for an initial development phase with cloud resources, but will then need to submit a full RPP at the next competition.

7 Common Curriculum Vita (CCV)

Keep your CCV up to date!

1. It is worth emphasizing that the CCV is used by the review committees during their consideration of the quality of the PI and his/her research. PI’s must for instance keep their bibliography up-to-date to ensure that reviewers are aware of the latest (and greatest!) work.
2. CC relies on the CCV for bibliographic analyses used both for annual reporting, and for funding proposals. The Field Weighted Citation Index (FWCI) is a key component of the metrics presented by CC, and is completely dependent on having complete bibliographic information.

As PI's know the CCV was introduced a couple of years ago as a consistent CV that is expected to become universally accepted, and specifically could be used for RAC and CC reporting requirements (mandated by CFI). This was a major effort as all users had to create and complete their CCV, and was the cause of quite a few complaints due to the complexity of the process. However the results have been critical to CC's reporting requirements, and significantly improved the science review process.

So CC will continue to use the CCV!

Note that users are asked to update their CCV's during the winter/spring renewal process, so hopefully most PI's already have an updated CCV.

The CCV update process is unfortunately not trivial. Please see the CCV Submission Guide for details.

"For future allocation requests, please consider adding references to published work in the proposal, progress made in the previous year, and clear description of how HQP will be involved in achieving the milestones."

"Progress over the past year is missing. Based on the CCV, most of the group's publications were enabled through Compute Canada resources, so that information should have been accessible and would have been relevant to include."

8 Review Process

8.1 No Appeal

There is **no appeal process**.

So it is necessary to be clear, complete and focussed in your application to ensure that the reviewers have a complete understanding of your project.

8.2 Oversight and Final Allocations

The RAC committee oversees the RAC process, and undertakes the final scoring and makes the final recommendation to the board for allocations. A key task is the derivation of the scaling function for compute requests, and the final review of cloud and storage requests.

The RAC committee is chaired by the CC RAC project coordinator, and includes:

- CC CTO

- Chair of the Science Leadership Council (SLC).
- Expert Review subcommittee chairs (discipline-based)
- Technical Review leads
- Regional CTO's (or delegates)
- Regional CEO's (or delegates)
- Cloud National Team chair.
- Scheduling National Team chair.
- RAC Project Coordinator

8.3 Science Expert Review Committees

See the CC RAC document (about ½ way down the page) for details of the RAC science review committees.

8.4 Review Process

The overall process followed by the RAC committee is:

Preparation and setup	Review results and issues from previous year. Work with Science Leadership Council and Senior Management to revise and update.
Appoint expert review committees	Approximately 80 reviewers are appointed to various disciplinary review panels.
Assign proposals for science review.	Proposals assigned to expert reviewers. ● One reviewer for each proposal. ● Large proposals may have 2 reviews
Assign proposals for technical review	Internal process. Carried out by the Compute Canada tech review team.
Proposal rating	Expert review committees meet to discuss and rate their proposals.
Face-to-Face Allocation meetings	Expert review committee chairs and the RAC committee ● Decide on a scaling function for compute. ● Recommend storage and cloud allocations. ● Re-allocate as necessary to even out system totals. ● make final allocation recommendations (mid-February).

9 RAC 2019 Changes and Updates

9.1 RPP 2019 Focus on Gateways

The Research Platforms and Portals (RPP) Competition enables communities to develop research projects that improve access to shared datasets, enhance existing online research tools and facilities, or advance national or international research collaborations.

Note that the emphasis is on **creating scientific gateways which provide services to a community** of users. Compute Canada will no longer entertain RPP requests for pure compute or storage.

Pure compute and storage requests must now be in the RRG competition. Continuing multi-year asks should make use of the fast-track process.

9.2 RPP Annual Report Requirements

Existing RPP's are asked to submit a progress/status report usually in January. The requirements have been tightened.

IMPORTANT: **If the RPP annual report is <u>not received</u> by the deadline, *then the allocation for the current year will be left to expire and not renewed*.**

Once the progress report is completed and reviewed, the science score of the original RPP application is inserted into the normal allocation process, together with any requested updates to the original resource request. **All new and continuing platforms and portals are then allocated through the normal process.**

Note that this process does not guarantee an allocation identical to the previous year. CC strives to keep the allocations consistent from year-to-year, but compute asks may be scaled based on Compute Canada's ability to meet the demand.

The progress report is reviewed by Compute Canada staff. It is not supposed to be particularly onerous, but Compute Canada would really like to see evidence of uptake by the community of users. Particular items considered by the review team include:

9.2.1 For a new development:

1. Creation of the development team.
2. Development effort with prototypes in operation.

3. Identification of issues and challenges that have been encountered, and plans for mitigation including revisions to the project plan and schedule.
4. Evidence of uptake through preliminary usage statistics: number of hits, number of users, number of downloads/uploads, etc.

9.2.2 For a mature gateway

1. Evidence of continuing usage. Number of hits, number of users, total downloads/uploads, storage utilization, etc.
2. Evidence of marketing effort.
3. User breakdown between Canadian and international users.
4. Research outcomes associated with the platform (for example, number of papers acknowledging or citing the gateway).
5. Usage (compute, storage, and cloud) within 50% of the predicted estimates. Usage below this limit may result in a prorated cut to the next year's allocation.
6. Expected changes in support levels from CC. In particular requests for significant increases in CC support need to be justified.

9.2.3 Acceptable changes

An increase of up to 10% with a short justification to be included in the progress report. Increase requests greater than 10% may require further review. In this case, more detailed justification should be provided in the progress report.

Projects are also encouraged to inform Compute Canada if they require **fewer resources** than originally requested so that limited resources can be redistributed to ensure optimal utilization.

9.3 RPP Streams Removed

RPP 2018 had two streams (regular and large). In practice this was not successful as it was difficult for applicants to understand the difference and for reviewers to take into account shorter proposals. Therefore the two streams have been removed, and there is only one stream in RPP 2019, with one template.

9.4 Page Limits and PDF Only

The technical justification document can only be submitted as a pdf doc. Page limits will be applied.

9.5 Secondary Contact Person

PI's are encouraged to appoint a secondary contact person so that important communications do not go astray.

We had a few PI's who did not respond to important requests (like the RPP status report) and asked that notes be sent to their project manager or second-in-command. We have added a "secondary contact person" field for this purpose.

9.6 No Generic Resources

With the large number of generally similar legacy systems the RAC had "generic" categories for users who could run on any one of a class of systems. With the consolidation to a small number of national systems the "generic" classes have been removed. Please choose an actual system from the Compute Canada list of available resources.

We generally assume that new users have at least tried various sites under a RAS account, so should already have some idea of the right system. Please contact support@westgrid.ca for further consultation.

10 Evaluation Criteria

10.1 Introduction

The following is a summary, but remember to have a look at the details in the application guides.

- RPP 2019 Application Guide
- RRG 2019 Application Guide

(The evaluation criteria are at the end of the application guides.)

10.2 Scoring Matrix

A score of "Strong" to "Very Strong" was only just enough for an allocation last year!

The scoring matrix is described as:

Descriptor	Score	Our Comments
Exceptional	5	This is almost never given, and is only for the highest quality research and proposal.
Outstanding	4	Excellent science, excellent technical approaches with high quality proposals.

Very strong	3	Small impressions can make a difference between “Strong” and “Very Strong”.
Strong	2	**RAC 2018 cutoff was 2.0, so “Strong” proposals are insufficient for an allocation!**
Moderate	1	A passable proposal with no significant issues.
Insufficient	0	Significant issues either scientific or technical. Very few proposals receive scores <1.

Unfortunately this is fine-grained at the higher end, with scores from 2 to 5 being “strong” or better. Small differences can therefore have a significant effect on the overall score.

10.3 RRG Evaluation Criteria

10.3.1 Overview

Just minor changes from last year. As usual the emphasis is on the quality of the science as evidenced by the project and by the applicant.

Merit of the proposal	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 ● Unique opportunities for HQP Training
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

10.3.2 RRG Template Outline

This year there are no streams, and only 1 template.

- General presentation page limits are strictly enforced (12 pages)
- Must be in PDF format.
- See details in the Appendix.

	Section	Page	Description

		Limit	
1	Research Problem and Justification	1.5	Outline the research problem for which Compute Canada resources are being requested, its importance/relevance as well as your general objectives.
2	Research Justification	3	In-depth discussion of the problem(s), your specific projects, your methodology, timelines, and specific goals.
3	Training and Support of HQP	1.5	Describe how this allocation will support the training of Highly Qualified Personnel (HQP) that are reported on the online form.
4	Technical Justification (See below for details!)	4	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. Typically the entire section 4 will be 1-2 pages long but projects involving several key codes and/or with complicated storage requirements may need more.
5	Progress over last year	2	Highlight any notable RAC-enabled research that you have performed. This may be directly linked to a publication in your CCV publication record, or may be a work in progress. *The applicant has already included a CCV so this is an opportunity to highlight any particular contributions, and also to add anything new.*

10.3.3 RRG Technical Justification

4.1 Compute Requests (complete only if required)	**4.1.1 System Selection** *Justification for the preferred system requested in the online form. If possible mention other systems that might also meet your needs.* **4.1.2 Code Details** *Provide details for any special software requested in the online form. For example 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** *Describe code performance (e.g. # iterations/timesteps/Flops per hour of wallclock time and the*

	type of system used for this measure, RAM required per job/process, etc.) and particularly suitable systems or processor architectures. Provide 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 about 5 months (e.g. longer than for simple post-processing)? *In the case of parallel codes, please discuss scaling efficiency and justify the typical job sizes you will run. A detailed description of scaling efficiency is required if parallel jobs will use 256 cores or more and is preferred even for smaller jobs sizes (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** *Provide special memory requirements in addition to the basic per-core or per-node requests in the form.* **4.1.5 Size of Request** *Justify the computing needs that you asked for in the online form, and describe your level of confidence and experience with computing use. Explain how you estimated the total amount of compute time required for this project.*
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. 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** *In the online form, you have been asked to specify your storage requirements. Justify the stated storage needs and describe your level of confidence and experience using each of these storage types. Explain how you estimated the total amount of storage required for this project.*
5. Progress Over Past Year	*In addition to the CCV highlight any notable RAC-enabled research that you have performed. This may be directly linked to a publication in your CCV publication record, or may be a work in progress.*

10.4 RPP Evaluation Criteria

The overall approach is similar, but the format and scoring details have changed significantly. See the RPP 2019 Application Guide for details. The following is a summary highlighting major changes.

10.4.1 RPP Definitions and Eligibility

RPP's as the name suggests provide a platform or portal usually based around a dataset or toolset that is of more general interest to a particular scientific community. This gateway would generally be provided on Compute Canada cloud resources, but may also be provided on external infrastructure with, for instance, use of Compute Canada compute or storage. The key point however is that there must be a community of users for which the gateway provides data or services.

- Platforms like CANFAR (Astronomy) and Atlas (Particle Physics) that provide datasets and analysis tools through a comprehensive, world-wide network.
- Groups engaging in international agreements to provide multi-year computing or storage solutions based in Canada.
- Groups that are providing shared data sets accessible using a third party (non-Compute Canada) interface.

Allocations in this competition may be awarded for a maximum of three years. Multi-year allocations are subject to an annual review prior to the start of subsequent years and to availability of resources.

Generally platforms and portals are collaborative projects. It is useful to emphasize the collaborative nature, and give details of the contributions from each PI and team. PI's may be part of multiple proposals.

Compute Canada used to distinguish between portals (smaller data-only) and platforms (larger with compute toolsets) but in practice this was a confusing distinction. This year the split has been dropped, and we refer more generally to scientific gateways as including platforms and portals.

10.4.2 Criteria: High-level Aims

There are 2 top-level aims:

1. Provide resources to a larger research community via a set of cloud-based tools, applications, and/or data.
2. Be able to develop, operate and manage the proposed portal or platform with minimal support from Compute Canada.

And of course the usual HQP description.

In addition to being a useful gateway with a user community the project team must have the skills and ability to develop, operate and manage the gateway.

Note that the technical reviewers (expert CC staff) will explicitly consider the abilities of the team and make recommendations to the review committees.

Merit of the Strategic Plan	40%
Merit of the Management Plan	50%
Highly Qualified Personnel	10%

10.4.3 Criteria: Merit of the Strategic Plan (40%)

Research Problem and Justification	● The research problem or need that the platform/portal will address is clearly presented. ● The importance/relevance of the platform/portal for Canada is well justified. ● The general objectives of the platform/portal are clear.
Goal, Alignment and Impact	● The project goal is clearly stated and aligns with the goals of Compute Canada. ● The research area of focus is of importance and will generate benefits to Canada.

	● The expected impacts have been clearly explained.
Use of the Platform/Portal	● The applicant has clearly explained the added value from the creation of the proposed platform or portal for the identified communities. ● Creation of the research platform/portal is being driven by the research community targeted. ● If applicable – The application details the level of interaction between Canadian and international research groups.
Expected Outcomes	● The application presents a clear timeline for the delivery of the anticipated outcomes over the entire duration of the requested allocation and has indicated the means by which they will be measured. ● The outcomes presented are of relevance and importance, and will benefit the users of the platform/portal

10.4.4 Criteria: Merit of the Management Plan (50%)

Development and operations of the Platform/Portal	● The team assembled to develop and operate the platform have the right combination of skills (where positions are not yet filled, a description of the position has been included). ● The proposed methods and technologies are suitable and scientifically justified for the services to be provided by the platform. ● The approach to sharing data sets across the platform/portal is well detailed and the application addresses any potential accessibility issues.
Management of the Platform/Portal	● The team assembled to manage the platform have the right combination of skills (where positions are not yet filled, a description of the position has been included). ● The proposed management of the resources is well defined and will provide broad access to the research communities. ● The process for resource access is well defined and the application identifies a credible plan to maintain or increase the population accessing resources. ● If applicable, the scientific evaluation of selecting projects will result in supporting excellent research with the resources available. ● The reporting framework presented will ensure information collected on the users of the portal will provide the ability to track impacts (scientific, HQP, social, etc.) and the participants are maximizing the benefits of the resources to meet the expected outcomes of the application.

10.4.5 Criteria: Highly Qualified Personnel (10%)

	• The project will generate significant benefits to HQP and has outlined these benefits in the application. • The anticipated number of HQP participating in the platform/portal has been clearly delineated across academic levels. • The platform/portal will create unique training opportunities for the participating HQP. • The application identifies strong examples of cross-pollination between HQP disciplines resulting from the platform/portal.

11 Compute Canada Resources

Good RAC proposals are quite carefully aimed at specific CC resources. So it is important to that the proposal shows a good understanding of the CC resources being requested. A short overview is given in the appendix, and the details are in the CC documentation wiki.

We strongly recommend breaking up your project into subprojects or phases, and provide tables to organize your requests. The RRG and RPP templates have explicit examples.

Something like:

Project	**Team Members**	**Estimated Core-Years**	**/project Storage**	**Memory/ core**	**Comments**
Project 1	Student X	10,000	100TB	4 GB	...
Project 2	Students Y, Z	5,000	50TB	32 GB	...
Totals:		15,000	150TB		

where Project 1 and Project 2 are explicitly identified and justified in the justification text.

"*The technical justification did not show a calculation of the computing and storage needs. Providing a table with this information, as suggested in the guidelines, would have made this section stronger.*"

12 General Comments

12.1 Introduction

We generally expect PI's to be experienced research proposal writers who know what they need to present. There are many sources for best practices in writing research proposals, and most institutions have local experts who can help. Keep in mind that a RAC is for ARC resources, not for funding, and there may also be differences between disciplines and areas.

So the rest of this section may be familiar to most researchers, but it is worthwhile reviewing the basics.

12.2 Application Form vs Technical Justification

Last year there were quite a large number of discrepancies between the ask in the application form and the ask in the technical justification. We caught quite a few, but a number of incorrect allocations were made, with a lot of effort and some grief required to sort them out.

It was also clear that with the large number of legacy systems still in use the forms were complicated. This year (RAC 2019) there are only the new national systems and the forms have been considerably simplified.

Be very careful with resource asks in the application form. The application form is used to feed the master spreadsheet, so the **form takes precedence**. The free-form technical justification should be consistent!

12.3 Reviewer Expertise

Don't have lots of detailed jargon. That's very hard for someone outside the specific area. Write for a discipline expert, but not for the specific area.

Science reviewers are experts in their disciplines, but are not necessarily experts in the specific sub-discipline or area of any particular project. The very wide range of RAC proposals do not allow for area experts, so you cannot assume that your proposal is going to be read by someone who works directly in the area or field of the project.

Technical reviewers are Compute Canada staff. They are ARC experts, and will concentrate on the technical justification. But they are not discipline experts. They are generally looking for problems that can be flagged and discussed by the RAC review committee.

12.4 Use the Templates

We strongly recommend that PIs use the templates.

- RRG Application template: Word Version, LaTex version
- RPP Application Template

Summaries and details are in the following sections.

12.5 Address the Evaluation Criteria

Provide the information being asked for - address each item in the Evaluation Criteria.

Poor proposals generally do not provide sufficient information, or have mixed technical requirements into the research justification. This is very difficult to decipher for both science and technical reviewers, and results in poor scores.

13 Research Excellence And Justification Best Practices

13.1 General

Describe the science that will be done. Since reviewers are not area experts ensure that overall position and significance of the project within the research area is addressed. Highlight the significance of the research, and as usual emphasize any particular successes.

Citation rates of recent work help justify science.

Explain why it is challenging. Write a more technical subsection explaining the scientific methods for addressing the challenges. Use more general language for the technical aspects if possible. This will feed into the technical justification.

“Please improve motivation of why the proposed calculations are important, and what is to be learned and/or what other science depends on the results.”

13.2 Use a Thread or Narrative

There should be a thread or narrative used to present a well-connected and justified story.

Please do not give the proposal to a post-doc or grad student to write! In many cases it is clear that the supervisor has asked grad students to write sections without understanding or describing the big picture. This results in very piecemeal proposals without oversight or unified approaches.

Of course your students or post-docs maybe the experts in developing and running your applications, and therefore could for instance write individual sections. This can be an excellent approach but only if there is a guiding hand giving the overall picture.

14 Technical Justification Best Practices

14.1 Overview

Both the science and technical review committees will look at the technical justification. Internal experts in the tech review make comments which are summarized for the overall review. They inform the science review, are discussed by both the expert review panels and the final RAC committee deliberations, and can effect the overall rating.

Be clear and focussed in defining and justifying the resources required.

The technical justification connects the science to the actual ask. In practice it receives a lot of discussion by review committees.

Ensure that the online forms (requested resources) and the technical justification are consistent.

We have seen considerable differences between the forms, the ask and the technical justification! That is really just waving a red-flag for the reviewers.

The science reviewers are looking for good technical solution approaches and techniques that are indeed applicable to the challenges identified in the science.

The technical reviewers flag applications with issues or particular characteristics that need special consideration. They also comment both specifically and generally on the quality and

completeness of the technical justification. These comments and flags go to both the science review committees and the RAC committee.

Particular considerations of the technical review:

1. Flag justifications which are poor or not clear.
2. Review impact of scaling and flag applications which need special consideration.
3. Characterize storage requests (necessary, scalable, poorly justified, ..) and flag those that, for instance, are not eligible for scaling
4. Consider high-memory requests. This is a limited resource and needs to be well-justified.
5. Review specialized software requests. Most software is generally available but in particular some commercial software may only be licenses on specific systems.
6. Review GPU requests. This is a limited resource and needs to be well-justified.
7. Reviewers generally try to remain cognizant of inflated asks. Poorly justified asks come under intense scrutiny, and can be drastically cut if the reviewers are not happy.
8. Recommend specific systems. In particular the reviewer will consider and existing allocations or usage, and special requests, and recommend specific systems as necessary.

14.2 General Best Practices

Provide performance estimates.

- RAS is available to provide default resources to those who do not currently run on CC resources. So some resources are available for testing and prototyping.
- Performance tables really help the reviewers, and make the proposal look good!

Explain the projects being worked on and the proposed research plan.

- Include the number and size of the needed runs/jobs/virtual machines
- Again tables or details really add to the proposal
- If it is difficult to predict usage then emphasize the areas of uncertainty. This might include early research issues where details are still to be worked out, or administrative issues like onboarding graduate students or postdocs who have not yet worked out a detailed research plan.

Provide details of the numerical and computational techniques used.

- Describe the numerical and computational techniques. Again use reasonably accessible language. You can expect all the reviewers to be computationally literate, but not experts in your particular field.

- Provide details of any software packages required.
- It is advantageous to show that the users are experienced and familiar with the performance characteristics of such packages.
- If the packages are new to the research team then describe the approach for familiarization. If possible describe any tests that the team may undertake.

14.3 Memory Requirements for HPC jobs

We appreciate getting memory details with good justifications:

- Memory use estimates from actual runs and tests.
- Required memory footprints for shared-memory programs.
- Identify memory per core for highly distributed programs.

And it is fair to say that applicants generally give very inadequate details.

This is always a difficult issue for reviewers and final allocations.

Job schedulers will generally use **core equivalent** metrics to try to take into account memory use. The “cost” of a job is the core-equivalent use, not straight cores.

As an example at one extreme a job may require 1 core, but all the memory on a node. Since other jobs cannot run without memory this means that the remaining node cores are idle. So on a 2 cpu, 32 core machine this job would use 32 core-equivalents. **The scheduler will charge such jobs for 32 core equivalents!**

This kind of effect is in practice the single largest cause of unused resources, and dwarfs for instance the effects of power outages or system upgrades. So the technical reviewers, who are by nature very concerned with using systems efficiently, will carefully scrutinize memory requests. Also the RAC committee carefully considers total memory requests so that we do not over-allocate resources. Note that the RAC committee will juggle:

- Total cores.
- Total memory.
- Total storage.

During the RAC year be aware that running jobs which request a lot of memory will factor into the priority of following jobs. If you consistently run jobs which have greater memory requests than those you provided in your application then you will end up with much lower priority, and will not be able to use your whole allocation. Note in particular that we do receive complaints from users who towards the end of an allocation year realize that they are not getting their full allocation, and this can be due to memory issues in the priority calculation.

- The technical review team therefore works with Core Equivalents (CE) and must take into account the fact that there may be unavoidable idle resources.
- This is complicated as there are a range of memory nodes. Currently 7 different types of nodes with different memory footprints.
- Core Equivalents can be as much as 6x actual core requirements.
- Similarly for GPU nodes.

So it's very important to look carefully at your memory requirements. As usual a clear description of memory requirements with justifications (test results) considerably adds to the proposal.

15 RPP and Cloud Best Practices

15.1 Introduction and Definitions

Persistent cloud requests are usually concerned with the development of web-based portals and platforms (RPP competition) but RRG applications may also request cloud-based resources. In most cases these cannot be scaled as a specific server architecture is necessary.

It is important to understand that the current (RAC 2018) cloud resources are IaaS ("Infrastructure as a Service") resources. So it is completely up to the project team to design and implement a suitable architecture.

So as usual the application should provide the details of the scientific worth and quality, but focus more on the user community and the research that will be carried out by this community.

1. Identify the audience/community. Who would be interested in using the proposed platform or portal?
2. What is the size of the community? Please justify any such estimates.
3. What kinds of research would you expect the community to carry out? It would be good to point out explicit examples of exciting projects that would make use of the portal/platform.
4. Is this community international?
5. Are there any agreements in place that would put conditions on the request? For instance the Atlas High-Energy physics project is part of the Large Hadron Collider collaboration, and allocations must satisfy international agreements.

Applications should emphasize a number of different technical aspects, and show that the team is capable of developing and maintaining the portal or platform.

1. Carefully describe the server architecture.
2. If possible describe the security/privacy concerns, and the steps that will be taken to satisfy such requirements.
3. Is the architecture scalable? Refer to the size of the community to justify any scalability issues or plans.

4. Demonstrate the expertise and capability of your staff. A portal or platform requires system management so try to identify the staff resources necessary to both develop and operationally manage your proposed system.
5. Identify special requirements
    a. Cloud or tape based backups
    b. Particularly large memory requirements
    c. Performance or response requirements
    d. Redundancy and Reliability requirements. For instance multi-site architectures.
    e. ...

15.2 Cloud-based Platform and Portal Requests

Persistent cloud requests are usually concerned with the development of web-based portals and platforms (RPP competition) but RRG applications may also request cloud-based resources. In most cases these cannot be scaled as a specific server architecture is necessary. For larger asks based on a predicted large user community reviewers may recommend scaling if predicted usage is not adequately supported.

There have been quite enthusiastic portal/platform ideas with big asks, but in practice the user community is quite specialized and uptake is dependent on the user interfaces and the services offered. Such asks will be critically reviewed and the allocation decreased.

So we encourage phased plans, with for instance performance indicators showing a well-justified and strong understanding if the issues involved with not only technical development, but also marketing and communications.

If possible define performance and success metrics. It's always useful to have a nice table defining your Key Performance Indicators (KPIs)

16 HQP Training

Both RPP and RRG request details of the training of Highly Qualified Personnel.

- quality of HQP contributions
- impact of participation on HQP
- unique training opportunities for HQP
- potential cross pollination between disciplines of HQP

Tables are very convenient for reviewers. You should provide a table showing the expected HQP at each level (ugrad, masters, doctorate, ..).

If you have had any particularly excellent graduates then consider including a couple of

sentences highlighting their contribution.

You should highlight any training opportunities beyond that of normal academic teaching. For instance you may expect to hire a junior developer as part of a team and provide specific training during the project.

"For future allocation requests… describe in detail the involvement of HQP in past projects that utilized Compute Canada resources. The current proposal states that one PDF will be involved and reviewers were wondering if there were plans for this PDF to mentor junior graduate or undergraduate students. If so, mention this in the proposal."

"This proposal was very clearly articulated with an integrated HQP training plan. The number of HQPs is impressive."

17 Impact of Scaled Requests

17.1 Job-based HPC Requests

Compute requests are scaled based on the science score, and storage requests are individually considered due to the complexity of such requests. For instance an existing database or repository cannot be scaled, so decisions are accept/deny. But an output archive with historical data may be scalable, or at least moved to (cheap) tape.

The application form asks about the impact of scaling. Due to the above process this is a very important section. It is reviewed by technical staff during the technical review phase, and is used to flag or identify special compute and storage requests. These flags and comments are considered during the final allocation discussions.

So it is important to carefully address the impact of scaling. The technical reviewers receive a lot of requests for minimal scaling which are not well justified!

Similarly it is important to re-emphasize the significance of a well-justified and clearly described resource request. Technical reviewers are very aware of the issues involved with inflated requests. Both the technical justification and the impact of scaling sections contribute to a tech reviewer's understanding. Poorly justified requests are flagged and commented, and recommendations made to the RAC committee for consideration during the final RAC discussions.

We really appreciate an upfront, clear presentation of the impact of scaling on the project. What would you as a researcher do if you were given some percentage of your ask?

- Impact on hiring graduate students.
- Impact on hiring post-docs.
- Issues with collaborative projects for the other co-PI's or co-workers.
- Critical or short-term projects that are necessary for further activities.
- Identification of lower-priority projects or activities that could be postponed.
- Projects which are critical to, for instance, an early-stage researcher who is applying for tenure.

18 Appendix - Technical Review Process

18.1 Introduction

The Tech Review committee allocates proposals to regional technical staff for technical review. Generally this review is aimed at ensuring technical reasonableness and identifying special requests, and results in a set of comments and recommendations made to the RAC committee. *The technical review does not rank or score applications.*

In particular applications with special requests, or with insufficient justifications, are flagged for further consideration by the RAC committee.

In many cases the technical reviewer will attempt to contact the PI for further clarification. This has had issues in the past as PI's have not been responsive. If you are contacted please make an effort to respond.

18.2 Process

Applications are sent to the region of the PI's home institution. The regional leads appoint a local technical expert to undertake the review. Recommendations and comments are reviewed and summarized by the Tech Review Team lead, and integrated into the master spreadsheet for consideration by the RAC committee during the allocation discussions.

At any point in the process an additional technical review can be requested. This is usually for particularly complex proposals, for conflicts-of-interest, or when the original reviewer feels that their expertise or experience is not suitable.

18.3 Job-based Applications

These are mostly from the RRG competition, but there are also RPP proposals involving job-based computation.

The Technical review includes the following

- Compare requested resources with resources described in the technical justification.
- Evaluate the technical justification and identify any issues. Flag applications which are vague or unjustified.
- Identify inflated or poorly justified asks and flag them.
- Identify particularly complex requests with multiple resource requirements.
- Identify special requests.
- Review large memory requests with particular emphasis on the necessity for large memory nodes. These are in short supply.
- Identify requests for interactive use, usually involved with some sort of complicated pipeline or workflow.
- Identify particularly bursty or non-uniform requirements. Some users for instance have bursts of incoming experimental data, or expect to run a few extremely large jobs after a lengthy development cycle.
- Recommend specific systems if necessary.
- Identify storage requests that are archival in nature and could be moved to tape only.

18.4 Cloud-based Applications

These generally come from the RPP competition and are usually longer-term (up to 3 years).

- Compare requested resources with resources described in the technical justification.
- Identify persistent vs computational requests
- Carefully review the resource justification, and provide comments about the reasonableness of the request. Flag any issues or problems.
- Identify special requests, for instance of virtual systems with very large memory requirements, extensive public IPs, ..
- Carefully review storage requests. Block (file) storage in the cloud is limited

19 Appendix: RAC 2018 Statistics

See Compute Canada RAC 2018 results for further details.

19.1 Allocation Success Rate

Year	Total CPU Capacity	Total CY Requested	Total CY allocated	Allocation success rate
2018	211,020	287,957	158,632	55%
2017	182,760	254,251	147,384	58%
2016	155,952	237,862	128,463	54%
2015	161,888	191,690	123,699	65%
2014	190,466	172,989	133,508	77%
2013	187,227	142,106	126,677	89%
2012	189,024	103,845	87,312	84%
2011	132,316	72,848	75,471	104%

19.2 GPU Allocations

Year	Total GPU capacity	Total Requested	Total allocated	Allocation success rate
2018	976	4,092	1,840	21%
2017	1,420	2,785	1,042	38%
2016	373	1,357	269	20%
2015	482	608	300	49%

19.3 Cloud Allocations

Due to major upgrades in the arbutus cloud and new cloud partitions at graham and cedar it has not been possible to provide detailed statistics. RAC 2018 received a 36% increase in requests, and was able to allocate about 95% of the total virtual CPUs requested.22.6 Storage Allocations

20 Appendix - Resources

Good RAC proposals are quite carefully aimed at specific CC resources. So it is important to have an idea of the overall characteristics, and also of the details.

The resources that CC offers cover three broad approaches:

1. **Job-based HPC resources.** These are essentially large clusters to which users submit computational jobs.
    a. General purpose systems which are aimed more at smaller jobs of something like 1,000 cores or less. These systems include GPU nodes, and mixes of low and high memory nodes.
    b. Large parallel systems for big distributed memory jobs (MPI). These systems have very fast interconnects with as flat an architecture as possible, and generally have less memory per node as the focus is on providing the largest number of cores for the price.
2. **Long-term storage.** Very large, backed up storage systems.
    a. These are generally attached to the clusters, but are primarily long-term storage that can be used independently of the attached cluster.
3. **Cloud**
    a. Long-term persistent instances for Platforms and Portals. These are usually over-subscribed under the assumption that services like web-servers or databases are not under continuous load.
    b. Short-term instances for special-purpose computational jobs.

21 Appendix - RRG Template

Word version on Compute Canada web site

Purpose of this Document

This Technical Justification document will be reviewed for research quality and merit by a committee made up of your peers (ie. Canadian faculty members with expertise in the subject area). It will also be reviewed by Compute Canada technical staff in order to ensure that Compute Canada resources will be used appropriately and efficiently.

General Presentation

The document must have this format:

- ***Maximum of 12 pages** (extra pages will not be considered)*
- *Letter paper size*
- *Minimum of 2 cm margins on all sides*
- *Single line spacing*
- *Font in Times New Roman (12 pts or more) or Arial (11 pts or more)*

- *Header on all pages includes the title of the project and the name of the lead PI*
- *All pages numbered consecutively*
- *Text in italics may be removed*

Submission

- *This document must be submitted* ***in pdf format only.***
- ***Important note: remember to fill in the online application form with all your requests. There should be no discrepancies between requests in this document and the online form. Failing to include ALL your requests in the online application form might impact your final allocation, if your application is successful. In case of discrepancy, the online form will prevail.***

Please consult the Compute Canada Technical Glossary.

Evaluation Criteria

Merit of the Proposal (60%)

Considerations include:

- *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*
- *unique training opportunities for HQP*

Quality of the Applicant(s) (40%)

Considerations include (for the PI and co-PIs):

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

Resources for Research Groups Competition

Technical Justification Template 2019

1. Introduction to the Research Problem and Objectives

Maximum: 1.5 pages.

Outline the research problem for which Compute Canada resources are being requested, its importance/relevance as well as your general objectives.

2. Research Justification

Maximum: 3 pages.

In-depth discussion of the problem(s), your specific projects, your methodology, timelines, and specific goals.

3. Training and Support of HQP
Maximum: 1.5 pages.
Describe how this allocation will support the training of Highly Qualified Personnel (HQP) that are reported on the online form.

4. Technical Justification
Maximum: 4 pages
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 4 will be 1-2 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 (complete only if required)
Note: Please ensure ALL resources requested in this document are properly reflected in the Resources Request section of the online application form on CCDB. Failing to do so may negatively impact the evaluation and final allocation of successful applications.

4.1.1 System Selection
In the online form, you have been asked to indicate if there is a preferred system on which your project should run. Provide reasoning for your choice of system(s). If possible mention other systems that might also meet your needs. Applicants should use the resource identification tables on the Compute Canada website to indicate the resources requested.

4.1.2 Code Details
In the online form, you have been asked to specify any special software needed to support your project (including home-grown codes). 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 post-processing)?

In the case of parallel codes, please discuss scaling efficiency and justify (in terms of performance) the typical job sizes you will run. A detailed description 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

In the online form, you have been asked to specify memory requirements either per-core or per-node. If there is any additional information you would like to provide regarding memory requirements, please provide additional information here - otherwise leave blank.

4.1.5 Size of Request

In the online form, you have been asked to specify your compute requirements. In this document, please justify the stated computing needs and describe your level of confidence and experience with computing use. Explain how you estimated the total amount of compute time required for this project.

It is strongly recommended to include a list (or table) of projects and a justification for your resource request for each. For example:

Project	***Team Members***	***Estimated Number of core years***	***Associated Storage***
Project 1	*Student X*	*10,000*	*100TB*
Project 2	*Students Y, Z*	*5,000*	*50TB*
Totals:		*15,000*	*150TB*

(where Project 1 and Project 2 are explicitly identified and justified in the preceding text.) This allows reviewers to determine the consequences of any cuts they apply to your allocation and helps demonstrate the reasonableness of your request.

4.2 Storage Requests (complete only if required)

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

In the online form, you have been asked to specify your storage requirements, for a variety of different storage types (e.g., /PROJECT, /NEARLINE). In this document, please justify the stated storage needs and describe your level of confidence and experience using each of these storage types. Explain how you estimated the total amount of storage required for this project.

It is strongly recommended to include a list (or table) of projects and a justification for your resource request for each. For example:

Project	***Team Members***	***/PROJECT***	***/NEARLINE***
Project 1	*Student X*	*1,000 TB*	*0*
Project 2	*Students Y, Z*	*0*	*1,000 TB*
Totals:	*3 students*	*1,000TB*	*1,000 TB*

(where Project 1 and Project 2 are explicitly identified and justified in the preceding text.) This allows reviewers to determine the consequences of any cuts they apply to your allocation and helps demonstrate the reasonableness of your request.

5. Progress Over Past Year

Maximum: 2 pages

In the online form, you have been asked to provide an up-to-date CCV and to identify which of your publications were enabled by your use of Compute Canada resources. This is important for the future funding of Compute Canada and is useful for your review committee.

In this document, please highlight any notable RAC-enabled research that you have performed. This may be directly linked to a publication in your CCV publication record, or may be a work in progress.

22 Appendix - RPP Template

RPP Application Template (MS Word)

General presentation

The Strategic Plan and Resource Justification document should not exceed 15 pages. **Page limits will be enforced**, and any information after page 15 will be ignored.

Follow this formatting:

- *Page size: 8½ x 11 inches*
- *Margins: no less than ¾ of an inch on all sides*
- *Font: Times New Roman (no smaller than 12 pts) or Arial (no smaller than 11 pts)*
- *Single spaced*
- *Header: include the project title and the name of the PI at the top of every page*
- *All pages should be consecutively numbered*
- *You may remove all explanations (text in italics) if desired.*

Submission

- *This document must be submitted **in pdf format only.***

Please consult the Compute Canada Technical Glossary.

Evaluation Criteria

Merit of the Strategic Plan (40%)

- *Research Problem and Justification*
- *Goal, Alignment and Impact*
- *Use of the Platform/Portal*
- *Expected Outcomes*

Merit of the Management Plan (50%)

- *Development and Operations of the Platform/Portal*
- *Management of the Platform/Portal*

Highly Qualified Personnel (10%)

Sections 1-3: Maximum 10 pages in total

1. Strategic Plan

1.1. Research Problem and Justification

Discuss the research problem or need that each platform or portal listed in this proposal will address, its importance/relevance for Canada and your general objectives.

1.2. Goal, Alignment and Impact

- *List the specific goal(s) of the portal or platform*
- *Alignment with the strategic goals of Compute Canada as outlined in the Strategic Plan 2014-2019*
- *Describe the expected impacts of the research resulting from the platform/portal.*

1.3. Use of the Platform/Portal

- *Describe the need for a platform/portal and the added value that will be created for the user community*
- *Describe the user community for which the platform will provide services. Include if possible estimates of the size and activity of the user community*
- *If applicable - Describe the level of interaction between Canadian and international research groups and platforms*

1.4. Expected Outcomes

- *Present a clear timeline for the delivery of the anticipated outcomes over the entire duration of the requested allocation (see table below)*
- *Provide a detailed list of the outcomes resulting from the use of Compute Canada resources and the expected schedule for achieving these outcomes*
- *Explain the means by which the expected outcomes will be measured.*
- *Explain how other partners, researchers or scientific areas will benefit from the outcomes of the research*

The table below is provided as a suggestion. Please adjust according to the needs of your project. You can also include a Gantt chart:

Phase	**Year 1** 2019-2020	**Year 2** 2020-2021	**Year 3** 2021-2022	**Dependent on CC resources?**
Phase 1	*[goal, milestone or outcome]*			*Yes / No*
Phase 2				
Phase 3				

Note: Keep in mind that RAC resources are typically allocated from April to April each year.

2. Management Plan

2.1. Development and Operation of the Platform/Portal

- *Describe the team assembled to develop and operate the platform, and show that they have the right combination of skills (where positions are not yet filled, a description of the position should be included.)*

The table below is provided as a suggestion. Please adjust according to the needs of your project.

Full name	**Position**	**Allocation to project (%)**	**Brief description of the position**
Full name 1	*Position 1*	*100%*	*Brief description of position 1*
Full name 2	*Position 2*	*50%*	*Brief description of position 2*

- *Describe and justify the methods and technologies selected for the services to be provided by the platform.*
- *Outline the proposed approach to sharing data sets across or within the platform/portal*

2.2. Management of the Platform/Portal

- *The team assembled to manage the platform have the right combination of skills (where positions are not yet filled, a description of the position has been included).*

The table below is provided as a suggestion. Please adjust according to the needs of your project.

Name(s)	**Position**	**Allocation to project (%)**	**Brief description of the position**
Full name 1	*Position 1*	*100%*	*Brief description of position 1*

- *Provide a detailed description of the management team with particular focus on the team's expertise and ability to develop and operate the proposed platform.*
- *Explain how the research community will access the resources and how the platform/portal will maintain/increase the population accessing resources*
- *If applicable, describe the evaluation mechanisms used to provide scientific rigor in selecting projects*
- *Detail the type of information collected and the regularity of the project reporting to track the progress towards the anticipated expected outcomes*

3. Highly Qualified Personnel

- *Describe the potential impacts on HQP through their involvement in research activities using Compute Canada's resources*
- *Provide a table of the total number of HQP directly engaged in projects utilizing Compute Canada Resources and across academic levels (e.g. technicians, technologists, undergraduate students, graduate students, postdoctoral fellows and other trainees)*

HQP involved	# Canadian HQP	# International HQP
Undergraduate		
Master's		
PhD		
PostDoc		
Etc ...		

- *Please also provide the following information:*

Canadian collaborators

Name (or #) of user(s)	Role (in the project)	Position	Affiliation
	Co-PI, collaborator, HQP, user..*	*Professor, Assistant Professor, Masters, PhD, PostDoc, etc.*	*Institution*

**Only people with Compute Canada accounts can be listed as Co-PIs.*

International collaborators *(if applicable)*

Name (or #) of user(s)	Country	Role (in the project)	Position	Affiliation
		Collaborator, HQP, user..	*Professor, Assistant Professor, Masters, PhD, PostDoc, etc.*	*Institution*

- *Describe any unique training opportunities for HQP resulting from the platform/portal*
- *Provide any examples of potential for cross pollination between disciplines of HQP within the platform/portal community*

Section 4: Maximum 5 pages

4. Resource Justification

4.1. Resource request summary

Please summarize your ask here, and then provide a detailed justification in the next sections. Make sure that all resources requested in the table are also requested in the online application form. Feel free to modify the table to suit the needs of your project.

Resource Type	Year 1	Year 2	Year 3
CPU (core years)			
GPU (GPU years)			
/project storage (TB)			
/nearline storage (TB)			
Persistent Cloud (VCPU)			
Compute Cloud (VCPU)			
Cloud storage (GB)			

Important note: *remember to fill in the online application form with all your requests. There should be no discrepancies between requests in this document and the online form. Failing to include ALL your requests in the online application form might impact your final allocation, if your application is successful. In case of discrepancy, the online form will prevail.*

4.2. Expected resource usage

If you expect to request cloud resources in multiple locations, please provide details in this section.

4.2.1 Description of storage and compute needs.

4.2.2 Rationale for the choice of specific computing and storage resources (if applicable).

If you have any preferred location(s) for your cloud allocation, please mention it here and justify it. Check the list of Available Resources for RAC to know how these must be requested in the online form.

4.2.3 Description of high demand periods (if applicable).

4.2.4 Description of who is expected to use the resources (domestic proportion across country, international proportion and distribution).

23 Document History

Sept 11, 2018	v0.1	Copied from RAC 2018 version

Sept 14, 2018	v1.0	Updated to RAC 2019. First complete version.
Sept 15, 2018	v1.1	Fixed the templates.