

**compute**  
canada

# **2017 Resource Allocations Competition Results**


# Table of Contents

Executive Summary.....	3
Computational Resources.....	5
<i>CPU Allocations</i> .....	5
<i>GPU Allocations</i> .....	6
<i>Cloud Allocations</i> .....	6
Storage Resources.....	6
Acceptance Rate .....	7
Allocation Process.....	8
<i>Scaling for Compute Requests</i> .....	8
Monetary Value of the 2017 Allocations .....	9
About Compute Canada.....	9
Definitions .....	10
<i>Computational Resource Categories:</i> .....	10
<i>Storage</i> .....	10
<i>Storage Categories</i> .....	10

# Executive Summary

With funding from the Canada Foundation for Innovation (CFI), Compute Canada is leading the renewal of the national advanced research computing platform. This \$75 million investment will increase capacity for computation, cloud computing and storage. We are pleased that we are able to fully satisfy 2017 requests for cloud computing and data storage in line with peer review recommendations. However, delivering sufficient resources to accommodate demand for CPU and GPU is an ongoing challenge. We anticipate a growth in capacity as new systems are in production. By the end of 2017, Compute Canada will have all four Stage One national systems operational. The Niagara system at the University of Toronto is expected to have at least 30,000 more CPU cores available than the system it is replacing. This will help mitigate the CPU shortfall. We anticipate further mitigation of compute shortfalls through deployment of CFI Cyberinfrastructure Stage-2 resources, some of which should be in place in time for the start of the 2018 RAC year.

Compute Canada is Canada's national provider of advanced research computing resources serving the needs of more than 11,000 federally funded researchers including more than 3,100 Canadian faculty. Through our Sustainable Planning for Advanced Research Computing initiative, over the next five years Compute Canada conservatively estimates 7x increase in data storage requirements and a 5x increase in computing requirements. Each year, the RAC process becomes more competitive as demand for Compute Canada's resources grows. The 2017 competition received the highest number of applications to date. In anticipation of this growth the CFI has initiated funding for a second round of capital investments ("Stage 2") to continue to grow capacity to meet demand.



**Compute Canada reserves 80% of its compute resources to the Resource Allocation Competitions (RAC),** leaving 20% for use via our [Rapid Access Service](#). Please note that, due to high demand for CPU resources in 2017, more than 80% of capacity has been allocated to RAC recipients on the largest systems (e.g., 88% on Cedar, 85% on Graham, 92% on GPC) in order to meet as much demand as possible. Unlike other countries around the world, in Canada, we do not have a specialized provider serving very high-end needs, separate from general-purpose need. We have done an analysis of our largest users compared to the bulk of our users (who do not have RAC awards) and have found the quality of their publications (by field-weighted citation impact) to be similar. In other words, both users with large CPU needs and those with more modest CPU needs achieve high levels of scientific impact with their Compute Canada supported research. View our bibliometrics report [here](#).

**Table 1: Applications submitted to the Resource Allocation Competitions between 2011 and 2017**

<b>Year</b>	<b>Resources for Research Groups (RRG)</b>	<b>Research Platforms and Portals (RPP)</b>	<b>CFI Cyberinfrastructure Challenge 1</b>	<b>Total</b>
<b>2017</b>	<b>345</b>	<b>64</b>	<b>5</b>	<b>414</b>
<b>2016</b>	<b>324</b>	<b>42</b>		<b>366</b>
<b>2015</b>	<b>335</b>	<b>15</b>		<b>350</b>
<b>2014</b>	<b>291</b>			<b>291</b>
<b>2013</b>	<b>211</b>			<b>211</b>
<b>2012</b>	<b>159</b>			<b>159</b>
<b>2011</b>	<b>135</b>			<b>135</b>




# Computational Resources

## CPU Allocations

Based on available computing resources for 2017, Compute Canada was only able to allocate 58% of the CPU (core years) requests. As **Table 2** shows, the allocation success rate for CPU improved slightly in 2017 compared to 2016. **Note that more than 50,000 of the cores available in 2017 are new and are higher performance than those they replace.** Note also that the total core-years allocated in RAC 2017 is the highest amount of core-years allocated in Compute Canada’s history.

**Table 2: Historical CPU demand vs. supply (core years)**

Year	Total CPU Capacity	Total CY Requested	Total CY allocated	Allocation success rate
2017	182,760	254,002	147,211	0.58
2016	155,952	237,862	128,463	0.54
2015	161,888	191,690	123,699	0.65
2014	190,466	172,989	133,508	0.77
2013	187,227	142,106	126,677	0.89
2012	189,024	103,845	87,312	0.84
2011	132,316	72,848	75,471	1.04



## GPU Allocations

Constraint in GPU resources was greater than in CPU. As **Table 3** shows, the demand for GPUs has increased 4.5x since 2015. In spite of the increase, the allocation success rate was 38%, compared to 20% in 2016. GPUs in the newest systems have much greater performance than legacy GPU devices.

**Table 3: Historical GPU demand vs. supply (GPU years)**

Year	Total GPU Capacity	Total Requested	Total allocated	Allocation success rate
2017	1,420	2,785	1,042	0.38
2016	373	1,357	269	0.2
2015	482	608	300	0.49

## Cloud Allocations

The installation of Arbutus at the University of Victoria campus increases our cloud computing capacity from 104 nodes to 290 plus 36 nodes at the Sherbrooke University Storage was quadrupled, to over 2.2 petabytes. We received requests for 9,152 VCPU's and our capacity was 23,040 VCPU's. As the need and awareness for cloud computing resources grows we anticipate more requirements in this area.

## Storage Resources


The incorporation of the new systems Cedar (SFU), Graham (Waterloo), and Arbutus (Victoria) made possible for Compute Canada to meet the storage demand in 2017, as **Table 4** shows.

**Table 4: 2017 Storage Supply vs. Demand by Storage Type (TB)**

Storage type	2017 Cluster Capacity	2017 Total Requested	2017 Total Allocated	Allocation success rate
Project	43,151	31,335	30,146	0.96
Nearline	83,333	16,640	16,892	1.02
Cloud	560	408.5	408.5	1.00
<b>Total</b>	<b>127,044</b>	<b>48,383.5</b>	<b>47,446.5</b>	<b>0.98</b>

## Requested vs. Allocated Storage: 2011-2017

Scale: Petabytes



## Acceptance Rate

Submissions are evaluated for both technical feasibility and scientific excellence. For the 2017 competitions, 414 applications were submitted and 390 allocations were awarded. Note that virtually all of applicants are requesting resources to support research programs and HQP **that are already funded through the three federal granting agencies.**

This year's resource allocations competition awarded 58% of the total compute requested and 98% of the total storage requested. Due to the competitiveness of the proposals and the limited amount of computing resources available, all projects, across all disciplines, received final allocations less than their original request.

**Table 5: Requests vs. Allocations (broken down by resource)**

RAC 2017	Number of Requests Received	Number of Requests Granted
Storage	282	271
CPU	351	314
GPU	42	34
Cloud (VCPU)	46	41


# Allocation Process

- Compute Canada Technical staff review each proposal;
- A peer review panel evaluates each proposal:
  - Each proposal receives multiple independent reviews;
  - Scientific committees meet to discuss the applications;
  - The peer review panel may or may not recommend specific cuts for an application;
  - The peer review panel gives a final science score on a 5-point scale;
- The committee of RAC chairs endorses a scaling function based on science score. That scaling function is applied to all compute requests.

## Scaling for Compute Requests

As in previous years, in 2017 the available compute resources were not enough to satisfy the demand. This is because a considerable number of legacy systems are being removed from service at the same time that the new systems are coming online.

The scaling function applied to the 2017 competition (see chart below) was set so that only applications with a science score of 2.25 or higher received an allocation, with a maximum of 87.5% for those with a score of 5. Note that those who did not receive a compute allocation can still make opportunistic use of system via our [Rapid Access Service](#).



# Monetary Value of the 2017 Allocations

These values represent an average across all Compute Canada facilities and include the total capital and operational costs incurred by Compute Canada to deliver the resources and associated services. These are not commercial or market values. For the 2017 competition, the value of the resources allocated was calculated on a per-year basis using the following rates:

- \$188.84 / core-year
- \$566.52 / GPU-year
- \$128.00 / TB-year
- \$40.50 / VCPU-year
- \$178.50 / cloud storage TB-year (Ceph)

Please note that the valuation of each of these resources goes down each year as older, more expensive, resources are retired and replaced with newer, more cost effective, resources.

## About Compute Canada

Compute Canada, in partnership with regional organizations [ACENET](#), [Calcul Québec](#), [Compute Ontario](#) and [WestGrid](#), leads the acceleration of research and innovation by deploying state-of-the-art advanced research computing (ARC) systems, storage and software solutions. Together we provide essential ARC services and infrastructure for Canadian researchers and their collaborators in all academic and industrial sectors. Our world-class team of more than 200 experts employed by 37 partner universities and research institutions across the country provide direct support to research teams. Compute Canada is a proud ambassador for Canadian excellence in advanced research computing nationally and internationally.

# Definitions

**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. Compute Canada compute allocations are based on core year allocations.

## Computational Resource Categories:

**CPU:** (pronounced as separate letters): is the abbreviation for central processing unit. Sometimes referred to simply as the central processor, but more commonly called processor, the CPU is the brains of the computer where most calculations take place.

**GPU:** computing is the use of a graphics processing unit (GPU) to accelerate [deep learning](#), [analytics](#), and [engineering](#) applications for example.. GPU accelerators now power energy-efficient data centers in government labs, universities, enterprises, and small-and-medium businesses around the world. They play a huge role in accelerating applications in platforms ranging from artificial intelligence to cars, drones, and robots.

**VCPU:** stands for virtual central processing unit. One or more VCPUs are assigned to every Virtual Machine (VM) within a cloud environment. Each VCPU is seen as a single physical CPU core by the VM's operating system.

## Storage

**Terabytes (TB):** Terabytes are most often used to measure the storage capacity of large storage devices. One terabyte (abbreviated "TB") is equal to 1,000 gigabytes and precedes the petabyte unit of measurement.

## Storage Categories

**Project:** The project filesystem is of medium performance disk and generally not available to compute nodes on a clustered system. This filesystem is larger in available storage than a home directory, and in most systems, backed up regularly. This filesystem is generally used to store frequently-used project data.

**Nearline:** The nearline filesystem is made up of medium to low performance storage in very high capacity. This filesystem should be used for storage of data that is infrequently accessed that needs to be kept for long periods of time. This is not true archival storage in that the datasets are still considered "active".

**Compute Canada Cloud:** is a pool of hardware supporting virtualization. This can be thought of as Infrastructure as a Service (IaaS). There are currently 2 geographically separate clouds: West and East, with more coming on-line with GP2 and GP3.

**Cloud storage:** Persistent cloud storage provides virtual disk functionality to virtual machines running in the cloud. Persistent cloud storage is very reliable and scalable, made possible by specialized software (Ceph) running on a highly-redundant physical disk array.



**compute** | **calcul**  
canada | canada

155 University Avenue, Suite 302, Toronto, Ontario, Canada M5H 3B7  
[www.computecanada.ca](http://www.computecanada.ca) | [www.calculcanada.ca](http://www.calculcanada.ca) | [@ComputeCanada](https://twitter.com/ComputeCanada)  
416-228-1234 | 1-800-716-9417