Submission to

The Canada Foundation for Innovation

for the

Mid-Term Review

of

Compute Canada

ADDITIONAL QUESTIONS

April 28, 2010

Mid-Term Review: Additional Questions

a) The Ontario and Québec provinces provide funds for the consortia in their respective province. Please provide a breakdown of the funds against the consortia. Please list the main requirements and constraints on the use of the funds.

CLUMEQ and RQCHP receive a total of $400,000 per year from FQRNT (Québec). This amount is split equally between the two consortia ($200,000 each). It covers operations and there is no strong constraint on usage: it can be used for infrastructure support, user support and/or management. Renewal for 2010/11 has not been confirmed yet, even though the new fiscal year began April 1, 2010. There is no guarantee that funding will be available for the 2010-2011 fiscal year.

Traditionally, the Government of Ontario has matched the CFI grants to the consortia.

SHARCNET has received operating funds from the provincial government twice. The first grant, for $8.5million covered SHARCNET during its initial start-up (2001-2005). Additional funding was secured for the five year period 2006-2010, in the amount of $10.9million in order to further expand SHARCNET, and supplemented a successful CFI grant at that time. A request to ORF-RE Round 4 for renewed funding in 2009, was recently turned down.

The operating funds provided by ORF allows SHARCNET to pay for a host of expenses not covered by the strict eligibly rules of CFI-IOF or MRS. ORF has much more flexibility in terms of what it will pay for and so long as the expense supports the project, it is generally considered eligible. This includes:

- management compensation and expenses (for positions like the Scientific Director, Associate Director(s), Director of Finance and Administration, Administrative Assistant, Communications Officer, Site Leaders, etc.) (20%)
- funding for technical salaries not covered under the current grants (i.e. for any new schools who joined the SHARCNET consortium after the granting period, or to cover short-falls in the IOF or MRS funds for technical positions) (10%)
- supplies/materials (10%)
- facilities/equipment/software (35%)
- promotions/outreach, events, conferences (20%)
- travel, training (5%)

Because of the industry matching component, a significant portion of the grant (~35%) is allocated to facilities/equipment/software, which has allowed SHARCNET to pay for new members added to the ORION network lease, the purchase of new equipment, such as "*Saw*", and the smaller IBM blade system for industry outreach, and renewal of software licenses not covered under IOF.

It should also be pointed out that most of the cash overhead contribution from the ORF grant (which could easily be taken by the institutions for back-end, indirect costs) is redirected to SHARCNET. In addition, while space rental and power costs would be eligible expenses against the ORF, SHARCNET does NOT pay for these costs directly, as they are currently covered by the institutions.

Having these additional operating funds has allowed SHARCNET to provide the most comprehensive research support services (number of technical staff), training, and outreach programs.

HPCVL received a 5-year $11.5million grant from the Ontario Research Fund – Research Excellence program with a start date of January 2006. A portion of this grant ($3.3million) is meant to cover the indirect costs of support the research infrastructure and is assigned to the institutions leaving $8.2million for operations and other costs such as, renovations and some equipment. Eligible expenses may be found at the program web site: http://www.mri.gov.on.ca/english/programs/orf/re/program.asp.

Funding from this program has allowed HPCVL to support User Support and Systems Engineers, Management and Administration, and various programs, such as, the workshop program. In addition, HPCVL provides licensed software and personnel have developed programming tools and training materials. As HPCVL is located off-campus, funding from this program pays for the rental of 9,000 sq. ft. for the offices, computer room and supporting infrastructure space.

Unlike other sites, CFI has only allowed HPCVL to purchase 3 years of extended warranty on purchased equipment using CFI Innovation Fund money. HPCVL has paid for years 4 and 5 of the extended maintenance costs from this ORF-RE grant as well.

SciNet received an $8million, 5-year grant from ORF-RE which is critical for ongoing operation. This grant is used to pay some of the technical staff (38% of the ORF total) as well as covering all management and administration costs (21%) and the rental of the off-campus data centre space (13%). An agreement with the University of Toronto ensures that SciNet receives the full indirect costs included in the $8million (28%) which are then applied primarily towards power costs.

b) Table 12 gives a list of research areas supported by the Compute/Calcul Canada infrastructure. Please indicate which of these areas (or subareas) fall within the scope of the CIHR and the SSHRC.

Some, but not necessarily all, of the research in the following areas could be supported although Compute Canada has no way of knowing whether or not researchers are receiving grants from any of the funding agencies.

NSERC: Statistics

CIHR: Biochemistry, Biology/Life Sciences, Biophysics/Medical, Medicine.

SSHRC: Arts, Social Science

c) The call for applications to NRAC indicates that up to 40% of the NPF-1 systems would be available for successful applicants. Please indicate how this figure was determined?

CFI provides only 40% of the funding for equipment. CFI rules require matching funding from other sources. Therefore, it is reasonable that Compute Canada can allocate to the maximum of CFI's contribution. It should be noted that for the first NRAC call, **all** CFI-funded equipment was included for 40% allocation. It was not limited to NPF-1 funded equipment.

d) The annual power costs seem to be very high in Tables 1 and 2 (pages 19-20). Please provide a break-down of these costs for each major system in 2009.

Table of estimated power costs per system (some minor systems not shown)

Name	Institution	type	est. power (kW)	PUE	net power (kW)	power rate ($/kWh)	est. power cost
GPC	UofT	hybrid cluster	800	1.16	835	0.0866	633 596
coming NPF equipment (2011)	McGill	capability	700	1.5	945	0.0580	480 136
TCS	UofT	capability	430	1.16	449	0.0866	340 558
coming NPF equipment (2010)	UofC	capability	262	1.5	354	0.1000	309 841
coming NPF equipment (2011)	UdeM	capability	350	1.45	457	0.0507	202 870
coming NPF Phase 2 equipment (2010)	UBC	capacity	220	1.8	356	0.0627	195 753
Angel	Waterloo	capability	200	1.5	270	0.0800	189 216
Requin	McMaster	capability	200	1.5	270	0.0800	189 216
Sun SPARC Enterprise M9000	Queen's	Large SMP cluster	176	1.5	238	0.0900	187 324
Colosse	Laval	capability	350	1.15	362	0.0580	184 052
coming NPF equipment (2010)	UdeS	capability	500	0.889	400	0.0507	177 686
coming NPF equipment (SMP & phase 2 capacity) 2010/2011	UofA	SMP	200	1.7	306	0.0629	168 481
Glacier	UBC	capacity	170	1.5	230	0.0627	126 053
Whale	Waterloo	capacity	130	1.5	176	0.0800	122 990
Bull	UWO	capability	120	1.5	162	0.0800	113 530
Sun Fire 25000	Queen's	SMP cluster	104	1.5	140	0.0900	110 691
coming NPF Phase 2 equipment (2011)	Uvic	storage & compute	120	1.5	162	0.0777	110 266
Hermes, Nestor & Pleiades	Uvic	capacity & capability	120	1.5	162	0.0777	110 266
Placentia	MUN	capability	134	*		0.0892	105 128
Orcinus	UBC	capacity	110	1.9	188	0.0627	103 314
Narwhal	Guelph	capability	100	1.5	135	0.0800	94 608
coming NPF Phase 2 equipment (2010)	SFU	cluster & storage	127	1.5	171	0.0627	94 169
coming NPF equipment (2010)	Uman	compute	150	1.5	203	0.0500	88 695
Mammouth-parallèle	UdeS	capability	212	1.015	194	0.0507	86 017
Brasdor	StFX	capacity	120	*		0.0942	64 712
Checkers	UofA	capacity	75	1.7	115	0.0629	63 180
Snowpatch	SFU	cluster	73	1.5	99	0.0627	54 129
–	UofT	storage	60	1.16	63	0.0866	47 520
Mammouth-série II	UdeS	hybrid cluster	115	1.015	105	0.0507	46 660
Nexus & Cortex	UofA	SMP	50	1.7	77	0.0629	42 120
Altix 4700	UdeM	SMP	66	1.45	86	0.0507	38 255
Bugaboo	SFU	capacity	50	1.5	68	0.0627	37 075
Sun SPARC Enterprise T5140	Queen's	SMP cluster	33	1.5	45	0.0900	35 123
Fundy	UNB (Fred.)	capacity	57	*		0.0697	34 803
Silky	Laurier	SMP	35	1.5	47	0.0800	33 113
Glooscap	Dalhousie	capacity	60	*		0.0616	32 356
–	Queen's	storage	30	1.5	41	0.0900	31 930
Mahone	St.Mary's	capacity	58	*		0.0616	31 493
Saw	Waterloo	capability	30	1.5	41	0.0800	28 382
Cottos	UdeM	capability	45	1.45	59	0.0507	26 083
Beowulf Cluster	Carleton	capacity	15	1.5	20	0.0900	15 965
Total			**6 957**		**8 027**		**5 187 357**

Note: *The PUE is the factor the power must be multiplied with to take into account the overhead consumption due to cooling. If heat is recycled, then this is also factored in the column called PUE. When the PUE is not known, the value 1.5 was used as a an approximation. Some power rates are also estimates. Systems are shown in decreasing order of estimated cost.*

**ACEnet has no way of easily measuring the actual PUE achieved, but all machines are cooled with typical CRACs with no heat recovery or outside air.*

The following table provides the power costs for data centres in the United States as a comparison for the numbers provided for the Compute Canada data centres.

INSTITUTION	POWER COSTS	REMARKS
DOE Laboratory (US)	3.5-4 cents/kwh	"cheap power" in this area of the US
PSC	9.3 cents/kwh	significant increase in last few years
SDSC	8 cents/kwh	an average of many sources
US total average	8.14 cents/kwh	range is from 5.01-13.95 cents/kwh excluding Hawaii

e) Table 18 gives a comparative breakdown of the number of user groups and analysts. Please breakdown the total figures for each of the consortia of Compute Canada.

The number of analysts below includes some individuals who are listed as "user support" although they do not work on people's code. As a result, in comparative terms, the number of groups per analyst may be even higher than is shown.

Consortium	# of Groups	# of Analysts	Groups per Analyst
ACEnet	199	5.5	36
CLUMEQ	68	4.3	16
HPCVL	129	4	32
RQCHP	110	6.5	17
SciNet	86	3.4	26
SHARCNET	322	9.8	33
WestGrid	257	8.2	31

f) Table 1 provides the annual costs of operation (2009) showing the funds for infrastructure support staff (CFI-IOF). Please provide the number of staff supported by these funds and the number of staff at each consortium.

It must be noted that while the total number of infrastructure support staff is 36.1, not all of these individuals are supported by IOF funds. Consortia fund their staff from a blending of the various sources of funds available to them. The break-out by consortium is as follows:

Consortium	Infrastructure Support Staff
ACEnet	2
CLUMEQ	3
HPCVL	4
RQCHP	5.1
SciNet	2.8
SHARCNET	7.9
WestGrid	11.3
TOTAL	**36.1**

g) NSERC-MRS provides the funds for user support staff which have been described in section 4. The funds (~$4m) appear to support around 40 effective full-time staff, indicating an average support of around $100,000 per full-time staff member. Please indicate the total average cost for a staff member including salaries, on-costs and overheads.

NSERC-MRS only provides $2million per year, not $4million per year. $4million per year is the cost, to which NSERC only contributes half on average. These funds are *not* enough to pay for the user support costs. As has been stated previously, we have a complex financial structure, with some funds (from provinces and institutions) paying for more than one category of staff. Please note that the total number of FTEs in this category, from the operations spreadsheet, is 47.98 (~ 48) and the total cost is $ 3,982,651 for an average cost per full-time staff member of $83,000.

Supplementary Questions

The submission outlines the need for Compute/Calcul Canada to provide additional user and infrastructure support staff and the allocation of staff resources (Section 4b). The main justification for additional staff is based on the comparative 'user-support' ratios. Compute/Calcul Canada might provide further justification for the additional staff (Table 19).

The complexity and sophistication of the projects requiring the support of Compute Canada staff is much greater than in the past. For example, Compute Canada is providing hardware and highly advanced grid middleware at centres across the country to support ATLAS. This requires new skills and must meet the expectations of our international partners. In addition, there are now researchers planning to run parallel codes across 20,000 (or more) cores simultaneously. Such simulations are an order-of-magnitude larger than anything that has ever been run on a Canadian system and the sheer scale necessarily requires significant expertise in parallel programming, visualization, networking, file systems, and systems administration making this a truly collaborative effort between the researchers and the Compute Canada staff. There are major shifts happening with new PGAS programming languages and GPUs which hold great promise. As a result, we need to develop expertise in UCP, OpenCL, CUDA, etc..

The experience within Compute Canada has been that increased computing resources have inspired researchers to attempt problems they couldn't even imagine tackling a year ago. Larger systems plus that new ambition creates new "challenges" in terms of systems administration, programming, visualization, and data management among others.

The current level of support has not kept up with the increase in the number of user groups. From 2008-2009 to 2009-2010 there has been an increase of 175 in the number of user groups and an increase of 700 in the number of users.

Canada is currently a part of the G-8 funding agencies call for exascale proposals. In order to play an effective role we need to assist research teams that are submitting proposals and to begin to build the expert support now to assist our researchers to port applications to this level of equipment in the future.

As the medical community becomes a larger component of the Compute Canada research base, we will need to be able to provide the specialized support they will require and which we do not have at the present time. We know we need to add expertise in the areas of database support and visualization.

Existing support staff serve users from across the country although there are not currently enough staff to provide significant in-depth support for individual large projects. They are managing and maintaining the resources but there is limited capacity to provide the kind of service to the research community that would increase their capabilities and productivity and to provide the level of assistance required by non-traditional users. In addition, there is, at times, significant demand for their expertise from the NRAC and Compute Canada.

The submission indicates that there is a need for archival storage support (p. 17). Compute/Calcul Canada might indicate the extent to which it is able to develop and market these capabilities to national research communities that are dependent on data management and access services.

It should be noted that the storage support referred to on page 17 is intended to mean “long term storage”. Compute Canada does not have the capability at the present time to offer a truly archival storage service with all that entails (curation, preservation, etc). In a recent JISC 10 discussion, it was noted that 70% of digital archive and storage costs are staff costs.

Compute Canada does not have sufficient funding in initiate a marketing function. We undertake outreach activities within the limit of our resources within the HPC community. In order to develop and market archival storage support, particularly beyond the existing community, Compute Canada would require a service-provision model with professional staff. There would be significant resource implications with respect to both people and hardware.