University Office of the President

UC Shared Research Computing Services (ShaRCS) Program Proposal

8/11/2011

Table of Contents

Executive Summary 3
Background 5
Business Need 5
ShaRCS Business Model 7
  Condo Model 7
  ShaRCS Timeline 9
  Non-Participating Researcher Use of ShaRCS 9
Value Proposition 10
  Total Cost of Ownership 10
  Benefits to Principal Investigators 11
  Benefits to UC Administration 12
Recommendations 14
Appendix A: Pilot Projects and Oversight Board 16
  Pilot Projects 16
  Oversight Board 17
Appendix B: Feedback Regarding Pilot Project 19
  Campus and PI Surveys 19
    Campus CIOs 19
    Pilot PIs 20
  PI and Campus Benefits 20
Appendix C: Detailed Cost Analysis Assumptions and Comparisons 22

Executive Summary

The Shared Research Computing Services (ShaRCS) is a UC system-wide initiative to create a centralized high performance research computing capability to support computationally-intensive research, enhance the ability to compete for large, multi-disciplinary grants, and save on facilities and energy. ShaRCS was borne out of UC cyberinfrastructure planning initiatives to *maximize UC research competitiveness while reducing costs*. As a research initiative, ShaRCS is a focal point for (1) coalescing resources to competitively seek grant funding; (2) ensuring national leadership capability; and (3) sharing computational research capability, capacity and expertise throughout the UC system. It achieves economies of scale by aggregating demand for research computing clusters across UC into integrated high-performance computing (HPC) systems based in energy efficient facilities, utilizing professional system administration and user support, and sharing common infrastructure. It also promotes collaborations by researchers across campuses to pursue large, multi-disciplinary grants.

With $5.6M approved in 2008 by UCOP for a ShaRCS pilot project, two mid-sized (272 nodes each) clusters at a northern (LBNL/Berkeley) and a southern (UCSD) site have been deployed and operated for UC pilot principal investigators (PIs). "Mid-sized" refers to a cluster that is at or beyond typical campus level facilities but not as large as national leadership class facilities such as those funded by the National Science Foundation (e.g. UCSD/SDSC) or Department of Energy (e.g. LBL/NERSC). The mid-sized clusters provide an alternative/upgrade path for users of small clusters and can provide a steppingstone to national-class (DOE, NSF) systems, permitting development and testing of codes without lengthy queue waits or limited windows for running codes. The ShaRCS clusters are much larger than all but a few individual researcher clusters within UC.

The initial pilot project has been used to demonstrate the operation of an integrated system-wide facility and to develop a ShaRCS business model for sustained operations. ShaRCS has been in operation for more than 18 months, supporting about two dozen projects across nine UC campuses and research labs. Overall, the system has been running as designed and feedback from PIs has been positive. The initial UCOP funding will sustain operations through approximately March 2012. The purpose of this document is to present a proposal, developed in conjunction with the system-wide ShaRCS Oversight Board, for sustaining ShaRCS beyond the present pilot phase.

To achieve a sustainable ShaRCS program, the Oversight Board is recommending a "condo" cluster business model in which the computing system is purchased and owned by participating PIs, who each receive computing time commensurate with their funding contribution, and annual operating costs are paid by the UC Administration[1] and/or participating PIs. Significant evidence exists within several UC campuses and many other research universities across the country that a "condo" cluster model is the most successful model for shared research computing, particularly when the university administration

[1] This document uses the term "UC Administration" to refer to UCOP and UC campuses in the aggregate, since the specific sources and allocations of funds for ShaRCS are currently under discussion.

pays all or most of the operating costs to incentivize PI participation[2]. The condo model enables resources to be shared across the participating PIs, and to potentially be used by the broader university community when not in use by the owners. The key feature of the condo model is that it aggregates extramural equipment funds from individual PIs to cover the bulk of the capital investment in an HPC system of significant scale; these capital funds comprise approximately half the total life cycle cost of the system.

The current budget proposal has ShaRCS in a self-sustaining financial situation when 512 or more PI-purchased nodes are installed at each site (1024 nodes across two sites). Based on experiences at other universities, this level of adoption should be readily achievable across the UC campuses. However there will be a ramp-up phase for adoption by PIs, and a commitment is required for the UC Administration to guarantee the fixed operating costs and cover potential deficits in the initial years. If demand exceeds 512 nodes per site, the program remains self-sustaining and in fact, annual operating fees would likely decrease due to economies of scale.

A detailed cost model has been developed, both for the ShaRCS business model and for comparison with the prevalent “autonomous” PI-operated clusters. The assumptions used for comparing costs are open to discussion, with the primary drivers being the relative energy/facilities costs and the level of labor required for the two systems. However, reference scenarios demonstrate that a PI and his home campus could achieve roughly 20% savings in the total cost-of-ownership by participating in ShaRCS instead of operating an autonomous cluster in a typical “server room” on campus. Much of the calculated savings in cost-of-ownership result from energy efficiencies and avoidance of facility retrofits or new construction. However, the most significant benefits of ShaRCS over autonomous clusters result from the facts that PIs have access to a larger, more capable, and professionally-maintained resource, and that the levels of service, support, availability and utilization are much higher in ShaRCS.

In addition to direct financial and operational benefits, ShaRCS will provide UC researchers with a powerful research computing resource to improve competitiveness on grants and enable greater scientific discovery. The financial investments required of the UC Administration will be offset not only by utilities/facilities savings, but also by a higher success rate on computationally-enabled extramural grants and the attendant increase in IDC generation. In addition, the availability of a shared computing program is an important tool for faculty recruitment/retention. Many universities across the country have established shared research computing programs with significant investments by their university administration, and UC must remain competitive in this environment.

Moving beyond a services program, the board has been analyzing the criteria and acts on leveraging campus expertise and capabilities by including the potential for another partner to address demand for the condo model, unique opportunities and/or demand not met by the initial two sites. The trajectory of this model may extend beyond LBNL and SDSC to other UC campuses that have existing investments in growth, expertise and management capacity.

[2] NSF-Sponsored Workshop Report: Sustainable Funding and Business Models for Academic Cyberinfrastructure Facilities, November 2010, available at http://www.cac.cornell.edu/~lifka/Downloads/SRCC/NSF-sponsored-Sustainable-Cyberinfrastructure-Report.pdf

Background

The Shared Research Computing Services (ShaRCS), a high-performance computing cluster project, is a result of cyberinfrastructure planning focused on maximizing UC's competitive position in research by providing faculty extended research capability and capacity through a robust and accessible HPC service for scientific research computing, realizing efficiencies in facilities and energy use, increasing grant fulfillments, and increasing researcher and staff collaboration across UC.

The pilot project began with a call for proposals in April, 2008 for computationally-intensive projects that would 1) advance UC research in priority areas, such as global health and environmental science; 2) make the selected projects more competitive for obtaining extramural support; and 3) nucleate new communities of cyber-enabled researchers in areas like the social sciences, arts, and humanities. Over 30 proposals were received from almost all campuses and national laboratories, and about two dozen were projects deemed suitable for inclusion in the pilot.

The initial funding commitment from UCOP was made in 2008 for a pilot project, to include $3.4M for computing, networking and storage capital purchases, and $2.2M for two years of recurring operating costs. There are northern and southern clusters, one located at the UC Berkeley data center and operated by staff from Lawrence Berkeley National Laboratory (LBNL) and the other located in the San Diego Supercomputer Center (SDSC) regional colo facility at UC San Diego and operated by SDSC personnel. The clusters are connected by a dedicated CENIC network link and are both accessible from all UC campuses. The pilot project began operations in early 2010, and the $5.6M total funding is expected to sustain operation through approximately March, 2012.

Currently, ShaRCS is supporting 23 pilot projects with 155 users representing nine campuses. An Oversight Board involving CIOs, VCRs, faculty PIs, and HPC staff from multiple campuses is charged with overseeing the project and defining a sustainable operating and long-term funding model for ShaRCS. This document discusses the business need for ShaRCS and recommends a funding model to sustain this program beyond the pilot phase. *If UC is to benefit from the investment made to date, a plan must be defined very soon for the next phase in order to ensure continuity of staff, facilities, and operations should the program continue beyond the pilot phase.*

Business Need

The urgency of creating a world-class UC Research Cyberinfrastructure (RCI) is reflected in a rapidly escalating need for greater computing capability and capacity to do competitive research as well as the escalating energy costs and facility needs associated with these resources. Investments by peer institutions in high performance research computing are impacting the competitiveness of UC campuses in attracting the best faculty and in garnering extramural support. Many UC faculty, the UC Vice Chancellors for Research (VCRs) and Chief Information Officers (CIOs) have previously indicated strong support for a system-wide UC RCI.

A mid-range shared computing resource such as ShaRCS is envisioned to fill a gap in the spectrum of research computing capabilities. For researchers using single workstations or small clusters, ShaRCS represents a capability upgrade when needed to take on more complex computations or to increase throughput for faster time-to-results. ShaRCS is also an energy-efficient, cost-effective alternative to users or would-be users of small clusters, as explained below. For current and future users of national-class (DOE, NSF) supercomputing facilities, ShaRCS can serve as a steppingstone for developing & testing codes prior to seeking allocations at national facilities.

In the absence of shared UC research computing facilities, there has been a proliferation of distributed computing clusters throughout the campuses, including systems in non-data center space, which often requires expensive retrofit or is sub-optimal in terms of energy efficiency. Such arrangements are extremely costly in terms of facilities, power, cooling, and space, costing 3 to 4 times more than efficient data center space. More importantly, such project-specific clusters are often significantly underutilized and it is difficult to harness them to create a greater computing resource or serve a broader community. If UC continues with "business as usual," such costs will continue to spiral upward and imperil the research enterprise.

We currently estimate that clusters distributed throughout the campuses occupy a footprint of over 25,000 sq.ft., displacing often scarce instructional and office space across UC. Collectively, distributed clusters already exceed the computing power at major supercomputer centers. Industry research indicates that cluster computing is growing at approximately 30% per year. At UC, we project an increase in computing power by 2015 that will be the equivalent of adding at least a second major supercomputer. Additionally, UC can expect to add as much as 80,000 ft^2 in new data center space at a cost of $160 - $320M. Over half of these resources will be needed for research computing and storage. Power costs, even with efficient data center designs, are expected to grow by $10 - 15 million in annual costs across UC. Apart from the significant cost savings associated with centralizing distributed clusters in datacenters, continuing distributed compute clusters throughout the system puts UC at a significant disadvantage by not being able to leverage the extended research computing capability and capacity possible through aggregation of resources into a major HPC system. For example, a recent study[3] of HPC investments by U.S. universities found that institutions with a system on the Top500[4] list achieved a subsequent increase of $2.4M per year in NSF funding (result confidence level of 95% with a confidence interval of $770K-$4M). Availability of large shared HPC resources also has the potential to increase critical researcher collaboration across UC campuses, reducing duplication and better utilizing expertise across UC, raising the potential to more effectively compete for large, multi-disciplinary grants[5].

In contrast to UC's *de facto* investment in distributed resources, many UC competitors are now investing heavily in common computational resources for their researchers. Many institutions, like University of Illinois, University of Texas, USC, Purdue, Clemson and University of New Mexico have recently made

[3] Apon, A., Ahalt, S., et al, High Performance Computing Instrumentation and Researcher Productivity in U.S. Universities, Journal of Information Technology Impact, Vol. 10, No. 2, pp. 87-98, 2010. Also Apon, A. and Ahalt, S., Investment in High Performance Computing, a Predictor of Research Competitiveness in U.S. Academic Institutions, Sep 2010, available at www.casc.org/meetings/10sep/Apon-Ahalt.ppt, April 2010.

[4] The Top500 list is a worldwide ranking of HPC systems based on benchmark results. See www.top500.org.

[5] Unless otherwise cited, much of the information in this paragraph was compiled originally for the2008 document, "UC Cyber-infrastructure: The Business Case for a Pilot."

investments of $5-15M each to centralize their resources into a large capability and capitalize upon the resulting economies of scale. UC will simply not be competitive without investing in large-scale research computing, an endeavor best undertaken with a centralized approach.

ShaRCS Business Model

The Oversight Board is proposing a “condo” cluster business model in which the computing system is purchased and owned by participating PIs, who each receive computing time commensurate with their funding contribution, and annual operating costs are paid by the UC Administration and/or participating PIs. The condo model enables resources to be shared across the participating PIs, and to potentially be used by the broader university community when not in use by the owners. The key feature of the condo model is that it aggregates extramural equipment funds from individual PIs to cover the bulk of the capital investment in an HPC system of significant scale, shown earlier to be key in attracting new grant funding; these capital funds comprise approximately half the total life cycle cost of the system. As an Oversight Board member stated, “If a university believes it needs to invest in research computing infrastructure, this program is a way to get 50% off.” Other significant benefits accrue to the UC system and campuses through energy efficiencies, avoidance of facility retrofits or new construction, and better utilization of facilities. The primary benefit to PIs is access to a larger, more capable, and professionally-maintained resource than would be possible through a single project or grant. If the annual fees are subsidized by the UC Administration, there is also a strong financial incentive for a PI to essentially receive cost-sharing for ~50% of the total cost of ownership (TCO) of a cluster, enhancing the competitiveness of extramural proposals.

Condo Model

Figure 1 illustrates the proposed condo business model for ShaRCS.

The ShaRCS program would be operated as a recharge activity, requiring income (from annual fees, explained below) and expenses to be balanced annually. A deficit in budgeted income would need to be made up by the UC Administration; this is identified as the “guarantee” in the figure and is explained further below.

Referring to the figure, PIs “join” the condo by purchasing compute nodes (servers) for the cluster using equipment purchase funds from grants. A typical PI might purchase 10-50 nodes comprising 80-600 computing “cores” (processors). Two node configurations are offered - a “basic” node at ~$3,900/node (Ethernet interconnect and NFS storage only) and a “high-performance” node at $5,300/node (basic node plus an Infiniband interconnect and Lustre parallel file system). This capital cost includes the compute node itself (estimated ~$3,500) and its network interfaces, as well as a *pro rata* share of the common storage and switching infrastructure. The hardware purchase is primarily a pass-through cost for the recharge activity; only the portion associated with the common disk and switch infrastructure is accounted for in the recharge.[6]

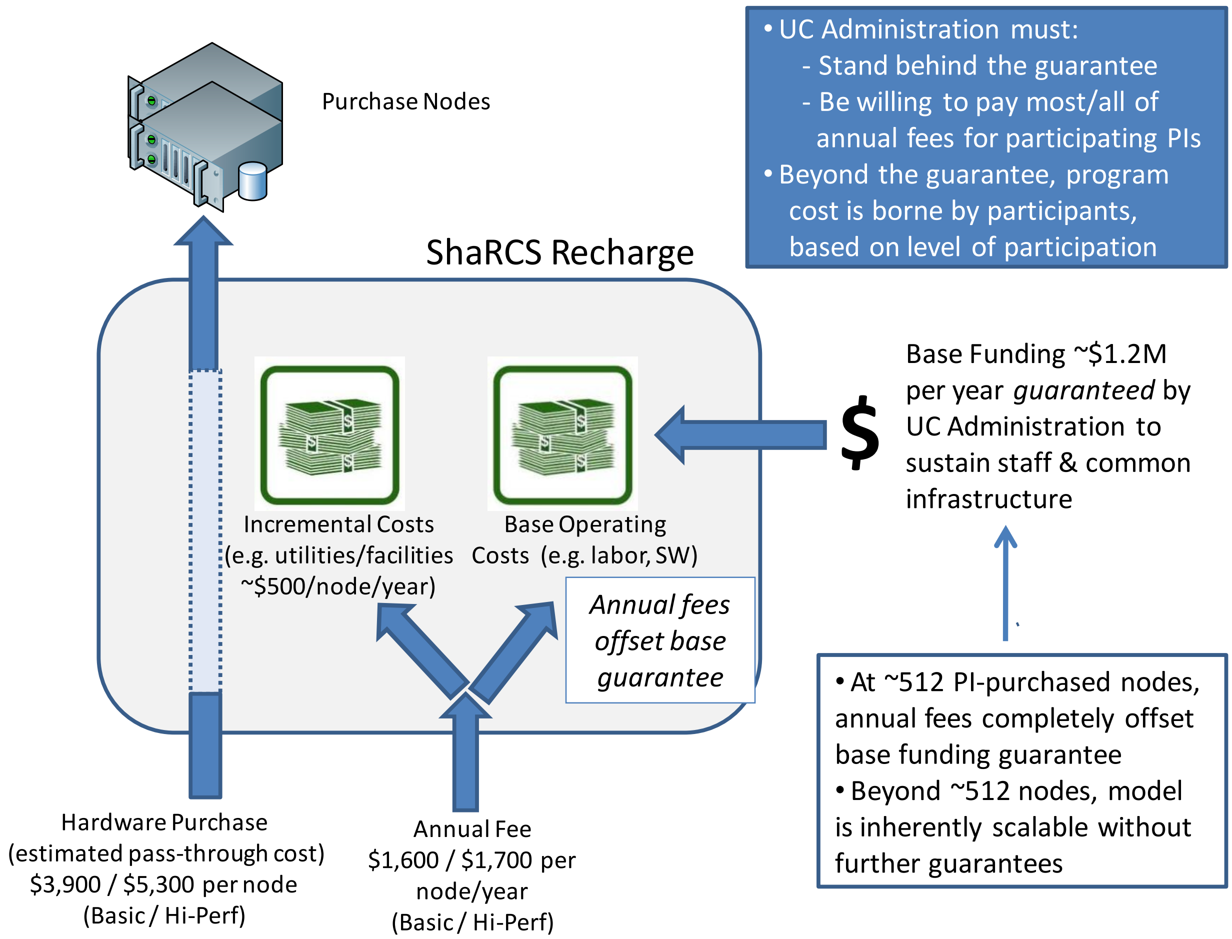

Figure 1. ShaRCS Business Model

Once nodes have been purchased and installed in the cluster, they are operated and maintained for the planned life cycle (nominally three years[7]). The operational costs are paid as an annual fee per node covering the *pro rata* share of labor (system administration, user support), common software licenses, facilities, and utilities. The annual fee amounts to $1,600 per year for basic nodes and $1,700 per year for high-performance nodes. This amount may be paid by the UC Administration and/or the PI. However based on experience at other universities, in order to incentivize participation and recognize the savings in facilities and utilities costs, it is strongly recommended that the UC Administration cover a significant portion of or even the entire annual fee.

[6] Participating PIs 'own' the nodes they purchased and, for the cost of shipping, may request their nodes be sent to them at the end of the ShaRCS operational period or if they wish to withdraw from ShaRCS for any reason. (There is no claim by the PI to the portion of the capital costs used for the common disk and switch infrastructure.)
[7] The operations period might be extended beyond three years, but due to Moore's law, there are diminishing efficiencies in utilities/facilities costs for older equipment. Especially if there are infrastructure limits on the number of nodes that can be hosted (e.g. switch size), there are incentives to limit operations to three years.

The proposed budget for the ShaRCS recharge activity is such that the activity is self-sustaining (income and expenses balance) at a cluster size of 512 nodes per site (1024 total across both sites). Should there be a shortfall in planned income, i.e., less than 512 PI-purchased nodes in the cluster, the UC Administration would need to make up the deficit. The maximum financial exposure, assuming no adoption/annual fees, paid would be $1.2M per year – this is the “guarantee” for base (fixed) operating costs necessary to ensure the cluster can be operated for participating PIs. Up to 512 nodes, the guarantee scales linearly. At or beyond 512 nodes per cluster, no funding from the guarantee is required and the recharge activity would be self-sustaining wherein annual fees would recover all operating costs (perhaps with some reduction in the annual fee due to further economies of scale). It is anticipated that the recharge activity would run a deficit in the first 2-3 years due to ramp-up time to solicit PIs, organize purchases, and begin receiving annual fees.

ShaRCS Timeline

Following is a high-level timeline of ShaRCS program evolution over five years:

	FY 11-12	FY 12-13	FY 13-14	FY 14-15	FY 15-16
Pilot (in progress)					
Pilot Extension Funding					
Transition to Condo					
Condo in Production					
# Purchased Nodes per site(year-end, cumulative) (illustrative scenario)		128	256	512	512
Recharge Deficit (both sites, covered by UC Administration) (illustrative, based on above adoption)		1,000,000	760,000	315,000	0

Table 1. ShaRCS Timeline

The last two rows in the timeline show an *assumed* ramp of PI-purchased nodes in the clusters over a three-year period (up to the self-sustaining level of 512 nodes per site or 1024 nodes total) and the deficit in the recharge balance that would need to be covered in this scenario until the program reaches the self-sustaining level. This adoption profile is an illustrative example only and not intended to be the plan of record.

Non-Participating Researcher Use of ShaRCS

It is likely the ShaRCS clusters will, on average, not be fully utilized by participating PIs. This opens the possibility for non-participating faculty and researchers to be allowed to access the clusters on a “space available" basis. This access could be structured as a paid service (e.g., cost per core-hour of computing time used) to generate marginal income for the recharge activity (reducing the guarantee or annual fee) or access could be provided for free in recognition of the campus contribution to annual fees. Permitting non-participating faculty and researchers to use the ShaRCS clusters is another way for the UC system and campuses to realize value from the investment in a centralized resource.

Value Proposition

Successful long-term adoption of ShaRCS requires buy-in from both PIs and the UC Administration based on the perceived merits and value propositions of the program. Surveys and interviews were conducted to help identify the value proposition for the stakeholders (see Appendix B). This section outlines the value proposition from the perspective of (a) total cost of ownership, (b) benefits to participating PIs, and (c) benefits to the UC Administration.

Total Cost of Ownership

Table 2 provides an illustrative comparison of the capital and operating costs of participating in ShaRCS as compared to an autonomous cluster. The table considers the case of a modest cluster (20 nodes) under different scenarios including installation in an inefficient "server room" with relatively high electricity costs and installation in an energy-efficient regional co-location facility such as the one currently operated by SDSC. Both the basic and high-performance compute node configurations are considered.

Key assumptions are provided in the table, with more detail found in Appendix C. While these estimates are obtained from a thoughtfully-constructed model, the results are sensitive to key assumptions - not only electricity rates and PUEs, but particularly labor levels which represent a large fraction of the operational costs. For example, the ShaRCS model assumes 3.0/3.5 (Basic/High-Performance) FTEs of professional staff for a robust service level of system administration and user support for each 512-node cluster, while the autonomous 20-node cluster is assumed to be operated by 25%/35% of a postdoc's time, likely at a much lower level of service and availability. While the ShaRCS team believes these illustrative examples are credible, we recognize others may make different assumptions with varying results.

As can be seen, *under the assumptions made in the examples*, 3-year cost savings on the order of 20% in total cost of ownership are achievable when participating in ShaRCS as compared to using an inefficient server room. In the case of the autonomous cluster being sited an efficient co-location facility, savings are calculated at 8-12%. While the total cost of ownership for a small cluster in an efficient co-location facility is close to the cost of participating in ShaRCS at the same level, it must be remembered that in the case of ShaRCS, the PI is gaining access to a more capable, professionally-administered cluster with the ability to run jobs at much larger core counts, enabling solution of more complex problems and/or higher throughput. In addition, the staffing levels and expertise of the ShaRCS staff will provide a higher level of availability and utilization of the system, and the ShaRCS user support staff will assist the PI in porting and optimizing application codes to better utilize the system. Finally, using ShaRCS permits the PI's research team to focus on their domain research rather than becoming system administrators.

Scenario	20-node basic cluster in inefficient server room	20-node high performance cluster in inefficient server room	20-node basic cluster in energy-efficient regional co-location facility	20-node high-perf cluster in energy-efficient regional co-location facility
Energy Assumptions	PUE 2.0 $0.14/kW-hr	PUE 2.0 $0.14/kW-hr	PUE 1.35 $0.08/kW-hr	PUE 1.35 $0.08/kW-hr
Autonomous Capital Cost	78,000	105,000	77,000	104,000
Autonomous Operating Cost (3 yrs)	137,000	162,000	107,000	132,000
ShaRCS Capital Cost for 20 nodes	75,000	106,000	75,000	106,000
ShaRCS Op. Cost for 20 nodes	95,000	103,000	95,000	103,000
Total Savings (added cost) over 3 years	45,000	58,000	14,000	28,000
% Savings (added cost) over Total Cost of Ownership (cap+ops)	21%	22%	8%	12%

Table 2. Cluster Cost Comparisons

Benefits to Principal Investigators

The main benefits of ShaRCS for PIs includes:

- Access to a larger and more capable HPC service to achieve faster time to results or to attack larger or more complex computational problems.
- By providing experienced on-call HPC staff to provide system administration, ShaRCS provides a higher level of availability and utilization for the cluster, and enables PIs to focus on their domain research as opposed to system administration of HPC systems.
- The ShaRCS user support staff assists the PI's team in using the system, including porting and optimizing codes onto an HPC system, improving time productivity and code efficiency.
- If the UC Administration supports all or most of the ShaRCS annual operating costs, the cost of obtaining required computing resources is reduced for PIs.
- The reduced costs, higher utilization, and increased productivity significantly improve the competitiveness of a PI's extramural proposals, resulting in a higher success rate/funding.
- Program is "out-sourced" – only capital investment by PI is required, with no hiring/facilities/wind-down for PI.

- ShaRCS facilitates increased collaboration among PIs across the entire system doing similar and related research leading to increased efficiencies and pursuit of larger, multi-disciplinary grants.

Benefits to UC Administration

The main benefits of ShaRCS for UCOP and campus administrations include:

- UC Administration can share the costs of creating a research computing infrastructure (~50%/50%) with funding agencies (~50%/50%) via the condo model and PI extramural grants.
- Program results in utilities and facilities savings, due to utilizing energy-efficient regional datacenters, avoiding costly facilities retrofits/construction, and higher utilization/fewer nodes in shared system than autonomous systems.
- Program increases research competitiveness of participants, resulting in increased indirect revenue.
- Computing resources and facilities are an important element of faculty recruitment/retention, and campuses can allocate startup package funds to ShaRCS.
- Program is scalable and self-sustaining by its participants once it hits target size.
- System potentially supports more researchers than just condo participants (through scavenging spare cycles for free or pay-as-you-go usage).

It is proposed that the annual operating fees be paid, substantially or in whole, by the UC Administration. While some of the key benefits of ShaRCS cannot be quantified financially, Table 3 details some estimates of the financial return on investment (ROI), assuming that the UC Administration pays 100% of the operating costs. Again, many elements of this table require assumptions and are subject to legitimate debate. However, it is important to note that ~$500 of the $1700/node/year annual fee represents simply a transfer of the minimal level of utilities/facilities costs on any campus, normally hidden costs, to the ShaRCS program. Furthermore, there can be significant savings in utilities/facilities above and beyond this minimal cost, as a system is moved from a “closet” to an energy-efficient facility, especially if retrofits/construction are required. We also note that $270 of the $1700 annual fee represent the utilities/facilities costs for continuing to operate the “seed” nodes already purchased by UCOP for the pilot hardware; the annual fee could be reduced by not operating these nodes, but there is still substantial useful life in these nodes and the program can provide substantial additional computing to participants. (Longer-term, these pilot nodes will be retired, and the annual fee will be reduced accordingly.)

An important element of the ROI is the projected increase in IDC revenue. The example in the table states that if a ShaRCS PI with 20 nodes increases their funding by $50K/year due to being more competitive by participating in the ShaRCS program, the increase in IDC equates to ~$900/node/year. Although actual experience will vary, there are many PIs within UC that have sizeable grants to support ShaRCS participation so this level of increase is certainly feasible and across all participating PIs, the increase in IDC revenue could be quite significant. Also, the Apon-Ahalt study referenced in footnote 3 found a strong correlation between university investments in HPC and increased NSF research funding.

Cost/savings elements	$/node/yr	Comments
ShaRCS business model annual fee	1,700	
Utility and facility costs in efficient regional colo facility (included in annual fee)	500	This is typically a hidden cost, but represents lower limit to actual costs to campus for hosting autonomous cluster.
Utility cost SAVINGS relative to energy-inefficient facility	0 – 800	Level of savings depends on relative cost of energy and PUE for hosting autonomous cluster. Is $500 for $0.14/kw-hr and PUE=2; is zero if autonomous cluster is in regional colo facility.
Utility savings due to improved system utilization	0 - 200	Experience at LBL that available cycles on autonomous systems are used ~50% while shared systems are ~70%. For same amount of computing, this better utilization provides an energy savings.
Facility SAVINGS relative to "closet"	0 - $$?	Highly variable depending on circumstances, but closet retrofits or new facilities are very expensive.
Increased IDC revenue due to enhanced competitiveness for PIs proposals by participating in ShaRCS	0 – 1000+ ?	Difficult to quantify but e.g., if participating PI increases grants by $2500/node/year due to ShaRCS, this generates ~$880/node/yr additional IDC (equates to $50K/yr more grant funding for 20-node cluster).
Net return (transfers/savings/revenue)	500 → >2500	
Utilities/facilities costs for operating pilot "seed" nodes from existing pilot hardware	260	This is optional cost, currently included in annual fee. Continuing to operate these nodes leverages pilot investment and provides 25-50% more cycles to participants for no additional capital costs.

Table 3: Financial Return on Investments by UC Administration

Some UC campuses already have various forms of local shared computing services, while others do not have shared facilities. For campuses with either no shared HPC service or modest shared HPC service, ShaRCS provides a professionally-administered HPC system without the need for a large capital investment from the campus. For campuses with an existing comparable HPC facility, ShaRCS may provide an augmentation capability allowing higher performance computing capability beyond what is currently offered by the campus or expansion/overflow capacity without facility retrofit or new construction.

In summary, UC as a whole benefits from ShaRCS through increased research competitiveness, collaboration, and productivity of PIs across all campuses. Conversely, if UC does not support centralized research computing models, PIs that rely on computing will be increasingly less competitive compared to their peers at the many institutions across the country that are making this strategic investment (see footnote 2).

Recommendations

Based on the success of the pilot project to date, the documented need for a robust research computing cyberinfrastructure, and the value propositions to both researchers and the administration, the recommendations are as follows:

1. Following the pilot phase, support continuation of the program under the condo business model and recharge activity as described in this document.
2. PIs will purchase compute nodes using equipment funds for approximately $3,900 per node for the basic configuration and $5,300 per node for the high performance configuration. (Note: The ShaRCS program will negotiate the best possible pricing for nodes of a given configuration and pass any savings on to PIs.)
3. A combination of the PI and UC Administration will pay annual operating fees for nodes in the condo at $1,600 per node per year for basic nodes and $1,700 per node per year for high performance nodes. It is recommended the UC Administration cover a significant share, preferably all, of the annual fees.
4. The UC Administration will provide financial support to cover any deficits in the ShaRCS recharge, up to $1.2M per year. A commitment to support the program will be *required* for 3 years following the date of the latest PI-purchased nodes, i.e., the initial commitment should last at least 5 years based on the review cycle suggested in (5) below.
5. At the end of two years, re-evaluate service based on the rate of adoption, user feedback and actual costs. At that point, a decision will be made whether the UC Administration will extend the guarantee beyond five years (thus allowing additional PIs to continue to join the program) or to operate the system only for the remaining life cycle of the last PI-purchased nodes that were installed.
6. The UC Administration may need to provide access to loans for capital purchases of common infrastructure equipment (e.g., switch chasses, storage, etc.) by the recharge account, costs that will be recouped from PIs as nodes are purchased.

The following are the next steps for the ShaRCS Program:

- Approve pilot extension funding for period of March-June 2012 as required, depending on timing of decision regarding program and implementation of necessary financial structure.

- Obtain approval of program structure and financial guarantee by UC Administration.
- Begin promoting program on campuses.
- Establish recharge activity and supporting financial arrangements (target start date of 1 Jul 2012).

The Oversight Board believes that a condo model provides an outstanding opportunity to acquire a significant scale HPC resource by leveraging extramural funds to purchase the computing hardware. This business model will transition the ShaRCS pilot to a reliable and stable research computing service. Solidifying support for this service creates the computing infrastructure that will yield substantial dividends for the research environment that lies ahead. By supporting this business model, UC can provide greater research capability, maintain UC's competitiveness, and significantly lower energy and facility costs system-wide.

Appendix A: Pilot Projects and Oversight Board

Pilot Projects

The following projects across the University of California were selected for the initial phase of ShaRCS.

Project	Campus	Principal Investigator(s)
Climate Modeling Capacity	Berkeley	John Chiang, Thomas Zach Powell, Inez Fung, Ron Cohen
Comparative Genomics Cyberinfrastructure Needs; Understanding Diversity in Microbial Community Sequencing	Berkeley	Steven Brenner
Phylogenomics Cyberinfrastructure for Biological Discovery	Berkeley	Kimmen Sjolander, Steven Brenner, Jasper Rine
Optimized Materials and Nanostructures from Predictive Computer Simulations	Davis	Giulia Galli, Francois Gygi
Hydrology Analysis Cyber-infrastructure Proposal	Irvine	Soroosh Sorooshian, Sue Bryant, Bisher Imam
Simulation and Modeling of biological molecules	Irvine	Doug Tobias
Speeding the Annotation and Analysis of Genomic Data for Biofuels and Biology Research	LBNL	Adam Arkin, Dylan Chivian, Paramvir Dehal, Paul Adams
CCSM to Study New Biofuels with Carbon Cycles	LBNL, Berkeley	Bill Collins
Research in the Physics of Real Materials at the Most Fundamental Level Using Atomistic First Principles (or ab initio) Quantum-Mechanical Calculations	LBNL, Berkeley	Steven Louie, Jeffrey Neaton
Universe-Scale Simulations for Dark Energy Experiments	LBNL, Berkeley	Martin White, David Schlegel
Nano-system Modeling and Design of Advanced Materials	Los Angeles	Nasr Ghoniem
Organic Reaction Mechanisms and Selectivities, Enzyme Design, and Material and Molecular Devices	Los Angeles	K.N. Houk
Particle-in-cell Simulations of Plasmas	Los Angeles	W.B.Mori, V.K.Decyk, F.S.Tsung, P.Pritchett, J.Tonge
Space Plasma Simulations	Los Angeles	Maha Ashour-Abdalla
Oceanic Simulation of Surface Waves and Currents	Los Angeles, Santa Barbara	J.C. McWilliams, A.F. Shchepetkin, and Yusuke Uchiyama

Dynamics and Allosteric Regulation of Enzyme Complex	Riverside	Chia-en Angelina Chang
Functional Theory for Multi-Scaling of Complex Molecular Systems and Processes	Riverside	Jianzhong Wu
Establishing CI Capable of Capture and Analysis of Next-Generation Sequencing Data	San Diego	Trey Ideker
Physics-Based Protein Structure Prediction	San Francisco	Ken Dill
Computational Chemistry and Chemical Engineering Projects	Santa Barbara	Joan Shea, Baron Peters
Development and Mathematical Analysis of Computational Methods	Santa Barbara	Paul Atzberger
California Current System	Santa Cruz	Christopher Edwards
Convection and Magnetic Field Generation	Santa Cruz	Gary Glatzmaier

Oversight Board

The Oversight Board, comprising UC faculty, technical staff, and administrators, has the responsibility for the on-going guidance, direction, and evaluation of this pilot.

Member	Campus
Chia-En Angelina Chang – PI, Asst Professor, Chemistry	Riverside
Jim Davis - Vice Provost - IT & Chief Academic Technology Officer	Los Angeles
David Ernst - Associate Vice President & Chief Information Officer	UCOP
Laine Farley - Executive Director of California Digital Library	UCOP
Gary Glatzmaier – PI, Professor, Earth and Planetary Sciences	Santa Cruz
Bernd Hamann – Associate Vice Chancellor for Research	Davis
John Huelsenbeck – PI, Professor, Biology	Berkeley
Richard Moore, Deputy Director, SDSC (co-chair)	San Diego
Warren Mori - PI and co-chair, Professor, Plasma Physics	Los Angeles
Sarah Nelson – PI, Professor, Radiology and Biomedical Imaging	San Francisco

Dana Roode – Chief Information Officer	Irvine
David Schlegel – PI, Senior Scientist, Astronomy	LBNL
Joan Shea – PI, Professor, Chemistry & Biochemistry, Physics	Santa Barbara
Larry Smarr – PI, Professor, Computer Science and Information Technologies	San Diego
Sam Traina - Vice Chancellor for Research and Graduate Dean	Merced
Michael Witherell - Vice Chancellor for Research	Santa Barbara
Peter Yellowlees – PI, Professor, Psychiatry and Behavioral Sciences	Davis Health Services

Appendix B: Feedback Regarding Pilot Project

Campus and PI Surveys

Campus CIOs

The CIOs were polled with a questionnaire about Shared Computing Service options at their campus in Q3 of 2009. The goal of the survey was to get more information about the options available to campus PIs for their scientific high performance computing needs. This information provided insight into the appropriate long-term funding requirements for ShaRCS and provided guidance on how best to market and position ShaRCS.

Campus with Shared Computing Resource(s)

Of the 12 campuses that responded to our questionnaire, 8 had some form of shared computing service (LBNL, Berkeley, Riverside, Santa Cruz, Irvine, San Diego, Santa Barbara [not at campus level], UCLA). HPC specifications varied from campus to campus but in many cases, ShaRCS provided a more powerful system than what was available to the campus. For the most part, many of the campus clusters were running at high capacities resulting in growing wait times for running jobs. There are some campus clusters, however, that could be scaled up to help address this issue. There is a mixture of "condo" and "hotel" models in these campus clusters and the price to the PI ranged from no charge to a low price point. For condo model clusters, PIs generally paid for the nodes. Cost seems to be the primary determinant of PI satisfaction. User support along with shorter wait times are also important to the PIs.

Campuses without a Shared Computing Resource

Of those campus that do not have a shared computing resource (ANR, Merced, Davis, San Francisco), all have considered in some way using a shared HPC resource. Price is the single most important factor in considering such a model followed by capacity and service.

How CIOs felt we could best position ShaRCS

Most campuses agreed that ShaRCS is best positioned between smaller scale clusters and larger clusters. Campuses with a shared computing service that is equal or greater in computing ability to ShaRCS felt that it would be best positioned as an augmentation capability. For those campuses that had an HPC inferior to ShaRCS or lacked such a shared computing resource, ShaRCS provided them the ability to run more complex and intensive jobs. This suggests that ShaRCS addresses a niche market in between small scale local clusters and the national HPC centers such as NSF and DOE. ShaRCS gives PIs the ability to run mid-level jobs without long turnaround times, a problem often associated with the national clusters. It was suggested that ShaRCS could be used as a cluster where researchers could develop and test their code prior to moving to these national centers.

Most campuses highlighted that funding ShaRCS and figuring out the "fair" allocation of cost to the PI and the institution (campus and UCOP) would be instrumental in its adoption.

Pilot PIs

In addition to continuously informally getting feedback about the pilot users, the PIs were formally surveyed in Q4 2010, half a year into the pilot. The survey questions were related to user satisfaction and how ShaRCS compared to different HPC options that users had access to.

Of the 23 pilot projects, 14 responded. Almost all users rated ShaRCS services to have at least met expectations. 90%+ of users noted that high performance computing is critical when they are evaluating a HPC service and most have considered dedicated HPC and national center as high performance computing options. 70%+ noted that the future availability of ShaRCS is critical for their research and 93% of the users who responded are very likely or somewhat likely to participate in ShaRCS long term. Of those users, 47% of them will be ready to participate in 2012.

PI and Campus Benefits

Benefit and Selling Point	PI	Campus
Access to professionally-administered resources	✓	
Access to common infrastructure, e.g., high performance interconnect, parallel file system	✓	
Opportunity to run larger-scale jobs than possible on project cluster	✓	
Guaranteed access to purchased nodes (within agreed-upon timeframe)	✓	
Short lead times for project startup	✓	
Avoid use of researchers/grad students for system administration	✓	
Higher system uptimes than with project-specific cluster	✓	
Improved grant competitiveness (lower cost to grant, cost-sharing opportunity, positive funding agency review for efficiency/"green" aspects	✓	

Bridge gap between technical workstations & national-scale systems	✓	
Reduced on-campus energy costs		✓
Avoid costly facilities modifications (HVAC, raised floor, power)		✓
Avoid duplication of secondary infrastructure (e.g., storage systems, switches)		✓
Possibility of increasing grant win rate & IDC recovery		✓
Computational resource for multi-campus grant proposals	✓	✓
Avoidance of building out additional campus capacity		✓
Possible excess capacity for specific research needs / projects	✓	✓

Appendix C: Detailed Cost Analysis Assumptions and Comparisons

Below is a list of input parameters used for the business model in the scenarios cited in this document. The business model is codified in a spreadsheet that uses these independent parameters as inputs and calculates a number of dependent parameters.

	AUTONOMOUS CLUSTERS					ShaRCS	
SCENARIO	20 node "basic" in inefficient server room	20-node "high-perf" in inefficient server room	20-node "basic" in energy-efficient regional colo facility	20-node "high-perf" in energy-efficient regional colo facility		512 node "basic" ShaRCS node (in regional colo facility)	512 node "high-perf" ShaRCS node (in regional colo facility)
Independent parameters to specify scenario							
# of compute nodes purchased by PIs	20	20	20	20		512	512
# years of operations	3.0	3.0	3.0	3.0		3.0	3.0
Cost per compute node, w/ 3-year warranty ($K)	3.50	3.50	3.50	3.50		3.50	3.50
Cost of storage/node ($K) - assume 5-year life in ShaRCS infrastructure, replace in autonomous	0.25	0.70	0.25	0.70		0.27	0.63
Cost of HCA+cable for IB connectivity ($K)	0.00	0.58	0.00	0.58		0.00	0.58
Cost of Infinband switch for cluster ($K), assume 4-year life in ShaRCS infrastructure, replace in autonomous - note: must scale appropriately with size of cluster	0.00	6.50	0.00	6.50		0	524
Cost of GigE switch for cluster ($K) - assume 4-year life in ShaRCS infrastructure, replace in autonomous - note: must scale appropriately with size of cluster	2.00	2.00	2.00	2.00		60	60
Cost of a 42U rack and PDU ($K) - five-year life in ShaRCS; no cost if at SDSC colo as it's included in the rate	2.00	2.00	0.00	0.00		0	0
Cost of software licenses for life of machine ($K)	15.00	15.00	15.00	15.00		135	135
Cost per FTE (salary and personnel-related costs; not benefits or IDC) ($K/yr)	60	60	60	60		90	90
Benefits rate for labor	40%	40%	40%	40%		40%	40%
# FTEs for sys admin and support	25%	35%	25%	35%		300%	350%
Power /node (incl pro-rated overhead factor for ancillary equipment) (KW)	0.33	0.33	0.33	0.33		0.33	0.33
Electricity cost ($/kw-hr)	0.14	0.14	0.08	0.08		0.08	0.08
Are utilities costs hidden (YES or NO)?	YES	YES	NO	NO		NO	NO
Facility PUE	2.00	2.00	1.35	1.35		1.35	1.35
Facility colocation cost/rack ($K/yr) without utilities - "true costs" - see notes above	6.50	6.50	6.50	6.50		6.50	6.50
Are facilities costs hidden (YES or NO)?	YES	YES	NO	NO		NO	NO
CENIC link ($K/yr)	0	0	0	0		7	7
Home directory backups ($K/yr)	0	0	0	0		50	50
Are backups optional service (yes) or included in annual fee	OPTIONAL	OPTIONAL	OPTIONAL	OPTIONAL		OPTIONAL	OPTIONAL
# seed nodes operated (not owned by PIs, but available to them) - incurs utility/facility costs	0	0	0	0		272	272
Applicable IDC rate (54.5% extramural, 0% state)	54.5%	54.5%	54.5%	54.5%		54.5%	54.5%
Cores/node (used ONLY for conversion to $/CPU-hour)	12	12	12	12		12	12
Average utilization (%)	50%	50%	50%	50%		70%	70%
% of ShaRCS annual fee paid by administration	N/A	N/A	N/A	N/A		100%	100%

Below is a summary of the cost elements for the scenarios cited in this document, using the assumptions in the previous table. Note that some assumptions and independent parameters are subjects of legitimate debate, so the results are dependent on the assumptions.

Summary of Cost Elements	**AUTONOMOUS CLUSTERS**					**ShaRCS**	
Scenarios	20 node "basic" in inefficient server room	20-node "high-perf" in inefficient server room	20-node "basic" in energy-efficient regional colo facility	20-node "high-perf" in energy-efficient regional colo facility		512 node "basic" ShaRCS node (in regional colo facility)	512 node "high-perf" ShaRCS node (in regional colo facility)
Total Cost-of-Ownership ($/node)	$10,739	$13,349	$9,214	$11,824		$8,652	$10,971
Capital ($/node)	$3,850	$5,200	$3,850	$5,200		$3,891	$5,841
Compute node and interfaces	$3,500	$4,075	$3,500	$4,075		$3,500	$4,075
Switch and storage infrastructure	$350	$1,125	$350	$1,125		$391	$1,766
Annual Operating Costs ($/node/year)	$2,296	$2,716	$1,788	$2,208		$1,587	$1,710
Utilities/facilities - PI nodes	$996	$996	$488	$488		$488	$488
Utilities/facilities - seed nodes	$0	$0	$0	$0		$259	$259
Labor	$1,050	$1,470	$1,050	$1,470		$738	$861
Software, CENIC connection	$250	$250	$250	$250		$102	$102